



Cisco Crosswork Network Controller 6.0 Solution Workflow Guide

First Published: 2023-08-16

Last Modified: 2023-09-21

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2023 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Solution Overview 1

Description 1

What's New in This Release 1

Supported Use Cases 9

Solution Components Overview and Integrated Architecture 11

Multi-Vendor Capabilities 17

Extensibility 18

CHAPTER 2

UI Overview 19

Log In 19

Dashboard 20

Navigation 20

Commit – Advanced Options 23

CHAPTER 3

Orchestrated Service Provisioning 27

Overview 27

Scenario: Implement and Maintain SLA for an L3VPN Service for SR-MPLS (using ODN) 29

Step 1 Create an ODN template to map color to an SLA objective and constraints 31

Step 2 Create an L3VPN Route Policy 33

Step 3 Create and provision the L3VPN service 36

Step 4 Enable Service Health monitoring 38

Step 5 Visualize the New VPN Service on the Map to See the Traffic Path 41

Step 6 Observe automatic network optimization 43

Step 7 Inspect a degraded service using Service Health to determine active symptoms 44

Summary and Conclusion 49

Scenario: Implement and Maintain SLA for an L3VPN Service for SRv6 (using ODN) 50

Step 1 Create an ODN template to map color to an SLA objective and constraints	51
Step 2 Create an L3VPN Route Policy	54
Step 3 Create and provision the L3VPN service	57
Step 4 Visualize the New VPN Service on the Map to See the Traffic Path	59
Step 5 Observe automatic network optimization	62
Summary and Conclusion	63
Scenario: Mandate a Static Path for an EVPN-VPWS Service using an Explicit MPLS SR-TE Policy	63
Step 1 Prepare for Creating a SID List	65
Step 2 Create the SID List in the Provisioning UI	66
Step 3 Create an explicit SR-TE policy for each VPN endpoint by importing a file	67
Step 4 Create and provision the L2VPN service	68
Step 5 Attach the SR-TE policies to the L2VPN Service	70
Step 6 Enable Service Health monitoring	70
Step 7 Visualize the L2VPN on the Map	73
Step 8 Inspect a degraded service using Service Health and Last 24Hr Metrics to identify issues	74
Summary and Conclusion	78
Scenario: Provision an L2VPN Service over an RSVP-TE Tunnel with Reserved Bandwidth	78
Step 1 Create an RSVP-TE tunnel for both directions of the L2VPN	79
Step 2 Create the L2VPN service and attach the RSVP tunnel to the service	81
Step 3 Visualize the L2VPN service on the map	83
Summary and Conclusion	83
Scenario: Provision a Soft Bandwidth Guarantee with Optimization Constraints	83
Step 1 Create a BWoD SR-TE Policy with the Requested Bandwidth and Optimization Intent	84
Step 2 Enable and Configure BWoD	87
Step 3 Verify that the policy's operational state is now Up and view the path on the map	87
Summary and Conclusion	88

CHAPTER 4**Bandwidth and Network Optimization 89**

Overview	89
Scenario: Use LCM to Reroute Traffic on an Overused Link	93
LCM Scenario: Assumptions and Prerequisites	94
LCM Scenario: Workflow	96
Step 1: Enable LCM and Configure the Utilization Thresholds	97
Step 2: View Link Congestion on the Map	99

Step 3: View TTE SR Policy Recommendations in the LCM Operational Dashboard	99
Step 4: Validate TTE SR Policy Deployment	101
Step 5: Remove TTE SR policies on LCM Recommendation	102
LCM Scenario: Summary and Conclusion	103
Scenario: Use CS-SR Policies to Reserve Bandwidth	103
CS-SR Scenario: Assumptions and Prerequisites	104
CS-SR Scenario: Workflow	112
Step 1: Enable Circuit Style Manager	113
Step 2: Configure Circuit Style SR-TE Policies Using Device CLI	118
Step 3: Configure Circuit Style SR-TE Policies Using Add	121
Step 4: Configure Circuit Style SR-TE Policies Using Import	123
Step 5: View Circuit Style SR-TE Policies on the Topology Map	126
Step 6: Verify Circuit Style SR-TE Policy Bandwidth Utilization	130
Step 7: Trigger Circuit Style SR-TE Path Recomputation	131
Summary and Conclusion	131

CHAPTER 5**Network Maintenance Window 133**

Overview	133
Scenario: Install an SMU During a Scheduled Maintenance Window	134
Step 1 Download Topology Plan Files for Impact Analysis	135
Step 2: Schedule the SMU Installation Playbook Run	136
Step 3 Verify the SMU Job Status	140
Summary and Conclusion	141

CHAPTER 6**Programmable Closed-Loop Remediation 143**

Overview	143
Scenario: Achieve Predictive Traffic Load Balancing Using Segment Routing Affinity	144
Workflow	145

CHAPTER 7**Automation of Onboarding and Provisioning of IOS-XR Devices Using ZTP 147**

Overview	147
Scenario: Use ZTP to Onboard and Provision New Devices Automatically	148
ZTP Scenario: Workflow	148

CHAPTER 8

Visualization of Native SR Path 151

Overview 151

Scenario: Troubleshoot Native SR IGP Paths Over Inter-AS Option C 152

Workflow: Native SR IGP Paths Troubleshooting 153

CHAPTER 9

Provision, Visualize, and Analyze Tree Segment Identifier Policies in Multipath Networks 157

Overview 157

Scenario: Provisioning, Visualizing, and Analyzing Tree-SID Policies in a Point-to-Multipoint L3VPN Service 158

Step 1 Create a Static Tree-SID Policy 159

Step 2 Visualize and Validate the new Static Tree-SID policy 163

Step 3 Associate the Static Tree-SID Policy with the newly created L3VPN service model 167

Step 4 Add the VPN nodes 171

Step 5 Visualize and Edit the Static mVPN Tree-SID Policy's L3VPN service model 173

Summary and Conclusion 176

CHAPTER 10

Transport Slice Provisioning 177

Overview 177

Scenario: Implement an Any-To-Any L3 eMBB Slice 183

Step 1 Create a Slice Template Catalog Entry 186

Step 2 Add Service Assurance into the Slice Template Catalog using the NSO CLI (optional) 188

Step 3 Create the Transport Slice Instance 189

Step 4 Deploy a Slice using NSO CLI (optional method) 198

Step 5 Visualize and Validate the New Slice Deployment 200

Summary and Conclusion 203



CHAPTER 1

Solution Overview

This section explains the following topics:

- [Description, on page 1](#)
- [What's New in This Release, on page 1](#)
- [Supported Use Cases , on page 9](#)
- [Solution Components Overview and Integrated Architecture, on page 11](#)
- [Multi-Vendor Capabilities , on page 17](#)
- [Extensibility, on page 18](#)

Description

The exponential growth of network traffic and the pressures of efficiently running network operations pose huge challenges for network operators. Providing quick, intent-based service delivery and optimal network utilization with the ability to react to bandwidth and latency demand fluctuations, in real time, is vital to success. Migration to Software-Defined Networks (SDNs) and automation of operational tasks is the optimal way for operators to accomplish these goals.

Cisco Crosswork Network Controller is an integrated network automation solution for deploying and operating IP transport networks that delivers increased service agility, cost efficiency, and optimization for faster time-to-customer value and lower operating costs. The solution combines intent-based network automation to deliver critical capabilities for service orchestration and fulfillment, network optimization, service path computation, device deployment and management, and anomaly detection with operator selected manual, or automated, remediation. Cisco Crosswork Network Controller delivers network optimization capabilities that are nearly impossible to replicate even with a highly skilled and dedicated staff operating the network.

The fully integrated solution combines functionality from multiple Crosswork components installed upon a common Crosswork infrastructure, as well as industry-leading capabilities from Cisco® Network Services Orchestrator (NSO), Cisco Segment Routing Path Computation Element (SR-PCE), and Cisco WAN Automation Engine (WAE). Its unified user interface provides a single pane of glass for real-time visualization of the network topology and services, provisioning, monitoring, and optimization.

What's New in This Release

The information below lists the primary new features and functionality introduced in Cisco Crosswork Network Controller 6.0.x.

Traffic Engineering

• Local Congestion Mitigation (LCM) feature pack:

- Automated Mode—This option allows LCM to automatically deploy TE tunnel recommendations based on thresholds that you configure.



Note Automated mode is accessible through Limited Availability. Engage your account team for further details.

- Manual Mode (default)—This option, which was available in previous releases, requires a user to view the LCM Operational Dashboard and decide whether to commit TE tunnel recommendations.
- Pause Mode—This option can pause LCM operations on a particular interface when LCM is in either Automated or Manual mode. Pausing operations in Automated mode are necessary in cases where deployed solutions do not result in the intended resolution, there is uneven ECMP traffic, there are policies that are not carrying traffic, or when an interface is continuously throttling between different solutions.



Note Pausing LCM operations removes all existing TE policies that were deployed for that interface.

• SR Circuit Style Manager (CSM) feature pack:

- Hop count is now available as a metric type when computing SR-TE Circuit Style policies.
- In response to feedback from customers, we have changed some events to alarms. For example, an alarm is triggered when policy traffic exceeds the reserved bandwidth pool size or threshold.
- APIs:
 - RESTCONF APIs—Manually re-optimize (single or multiple) SR-TE Circuit Style policies. These APIs can be initiated after network topology changes.
 - CSPolicyPathsOnLinks—Lists Circuit Style SR-TE policies on a specified link and filtered by its operational state (up,down,active, and unknown) of the specified policies.
 - AllCSPolicyPaths—Lists Circuit Style SR-TE policies filtered by its operational state and if it has hops (segment lists).
 - CSPolicyPathsonNode—Lists all Circuit Style SR-TE policies on specified nodes filtered by its operational state (up,down,active, and unknown).

To view API documentation, see [Cisco Devnet](#).

• Bandwidth on Demand feature pack:

- In previous releases, BWoD required protected adjacency SID constraints. Now user can elect BWoD to prefer to use protected (default option) or unprotected adjacency SIDs.
- The Policy Violation now has two options: Strict or Loose.

- The process of changing delegation from one PCE to another has been improved to guarantee a clean transfer of PCE roles.
- Enhanced batch processing of queued BWoD policy computations. The queue is initially cleared prior to running a list of new pending delegations/undelegations instead of running each delegation one at a time.

• **Flexible Algorithm:**

- You can now view Application-Specific Link Attribute ASLA Flexible Algorithm metrics (TE and Delay) in the link details:

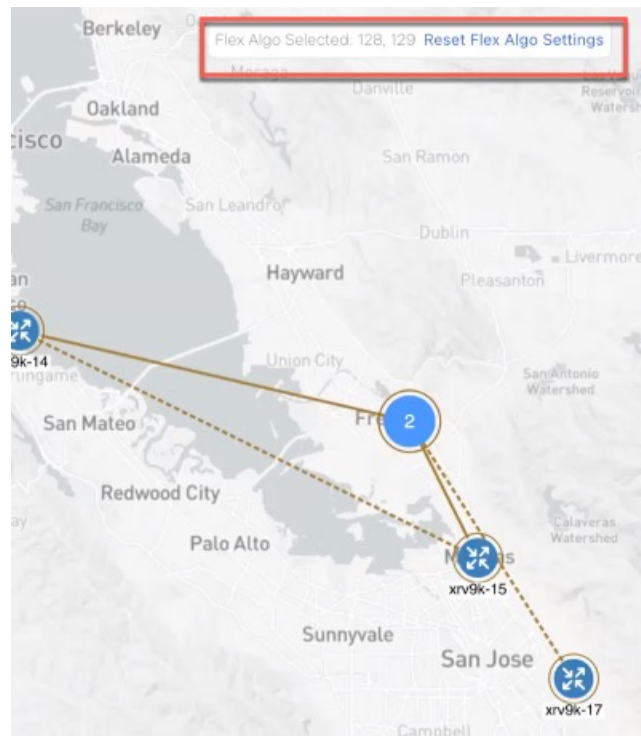


Note ASLA is supported on PCC and core routers that are Cisco IOS XR 7.4.1 or later versions.

1. From the Traffic Engineering topology map, click on a participating Flex Algorithm link.
2. From the Links page, click **Link_Type_entry** > **Traffic Engineering** tab > **General**. For example:

	A Side	Z Side
Node	xrv9k-15	xrv9k-13
IF Name	GigabitEthernet0/0/...	GigabitEthernet0/0/...
FA Affinities		
FA TE Metric	531	351
FA Delay Metric		
FA Topologies	128, 129, 130, 131, 1...	128, 129, 130, 131, 1...

- An overlay on the topology map has been added when Flexible Algorithms are selected. This is to help identify which Flexible Algorithms are selected more easily. For example:



- **Tree-SID:** PCE warnings and path compute elements are displayed in Tree-SID policy details.

Tree-SID Policy Details

Current
History

Summary

Admin State ↑ Up

Oper Status ↓ Down source node 192.168.0.2 not connected via PCEP

Label 9999

Type Static ?

Programming State None

Metric Type TE

Constraints Exclude-Any: -
Include-Any: -
Include-All: -

FRR Protected Disable

Node Count Leaf: 3 | Bud: 0 | Transit: 0

Path Compute Elements (SR-PCEs) 172.27.226.118 (Compute)

Last Update 01-Aug-2023 07:23:41 PM CDT

[See less](#) ^

- **Performance Metrics of TE policies:** When Service Health is installed and SR-PM collection is enabled, you can view KPI metrics (Delay, Jitter, and Liveness) from the Traffic Engineering table or from the TE tunnel details.

You can view the following KPI metrics for the policies:

SR-MPLS and SRv6 policies : Delay, Delay Variance (Jitter) or Liveness (Boolean value) along with traffic utilization.

RSVP-TE policy: Delay and Delay Variance (Jitter) along with Utilization.

- **Asymmetric delay for links**: In previous releases, only one side of the link delay value for an interface was considered during computation. When you configure delays on both remote and local nodes, the calculation of each delay on each interface is now taken into consideration when computing a path.



Note To configure link delay over an interface, refer to the device platform configuration guide. For example, [Segment Routing Configuration Guide for Cisco NCS 540 Series Routers](#).

- **Unique TE tunnel and device detail URLs**: TE tunnel or device details are now assigned unique URLs that can be shared. The URL sends the user to the Policy or Device Details page after logging in.
 - SR-MPLS, SRv6, Tree-SID, and RSVP-TE tunnels —From the Traffic Engineering table, click **Actions > View Details** for a particular row.
 - Devices—From the Traffic Engineering topology map, click on a device to view its details.
- **Increased performance and memory footprint**: Improvements made in topology discovery time, network model building, and processing cache, bandwidth, metric, and TE tunnel type information.
- **Transport Slicing**: Cisco Crosswork Network Controller offers direct support for network slicing at the transport level. This slice “instance” is a unique slice provisioned in the network but with a set of Service Level Requirements chosen from a set of pre-created Network Slice Templates (NST). The SMF in turn communicates with each sub-domain controller, called a Network Slice Subnet Management Function (NSSMF) which in turn provisions the corresponding domain specific slice instance across its own sub-domain boundaries (called a Network Slice Subnet Instance (NSSI) using a similar set of domain specific Slice Subnet Templates (NSST).

Cisco Crosswork Network Controller also offers:

- Slice design and deployment from the perspective of two user persona: Slice Designer and Slice Instance Requester.
- Deploying the Slice Catalog using the new slice template UI.
 - Adding Service Assurance into the catalog using the NSO CLI.
- Requesting a new Slice Instance by picking intent from the catalog using IETG Slice YANG model, select endpoints and submit.
 - You can deploy a slice instance after providing information, in the UI, after following four easy steps.
- Automated slice instance deployment:
 - QoS: The Slicing CFP will apply input and output QoS policy maps on all slice endpoint interfaces (policy-maps pre-deployed). Both L2 & L3 QoS supported.

- **Path Forwarding:** The Slicing CFP will deploy SR-TE ODN templates on all head-ends (metrics= latency, igp, TE, BWoD, FA, etc). Additionally, it will set BGP color community accordingly on all slice advertised VPN prefixes.
- **Service Assurance:** The Slicing CFP will setup:
 1. Cisco Crosswork Network Controller Heuristic packages for Cisco Crosswork Network Controller Automated Assurance/Service Health.
 2. Configure Y1731 probing for P2P L2 slices.
 3. Configure SR-PM probing for delay and liveness on all slice SR-TE tunnels.
- **Connectivity:** The Slicing CFP will use the L2/L3VPN IETF NM to setup L3 or L2 connectivity automatically across defined slice endpoints. All VPN parameters inferred and abstracted.
 - Setup eVPN VPWS for P2P L2 slices.
 - Setup eVPN any-to-any or hub-spoke for L2 multipoint or L3 multipoint slices.
 - Setup up “extranet” connectivity between dedicated and shared slice types. (more on this later).
 - Setup PE-CE eBGP for L3 based slices.



Note If you are creating an L2 point-to-point transport slice, prerequisites include the following: (a) The route-policy needs to be configured on the PE nodes (for example: L2-ATTACH). (b) In global-settings on NSO, configure this sample command: `set network-slice-services global-settings parent-rr-route-policy L2-ATTACH`.

- Using the UI, visualize the slice components: VPN, Transport, Health:
 - Display a slice on the map.
 - View Slice and VPN view along with Shared Slices and CE (Neighbor) connected in Logical View.
 - Visualize Shared Slices associated to dedicated slice.
 - From the VPN list, display VPN details including Assurance data if monitoring is enabled.
 - From the Transport list, display SR TE details including SR-PM data if SR-PM is enabled.
 - Using Health details, view symptom details and any failed subexpressions and metrics (which will provide information on any active symptoms and root causes).

Service Health

- **Introduced a new monitoring status - Monitoring Error:** Errors due to a component failures, operational errors or device errors are now displayed as **Monitoring Errors** on the UI. You can filter these errors using the mini-dashboard or the filters.
- **Ability to rate-limit monitoring requests:**

To efficiently manage service monitoring requests, Service Health has implemented a rate-limiting process. This means that there may be a delay in publishing service monitoring requests if the number of requests raised per minute exceeds a specific threshold. The thresholds are defined as follows:

- Basic monitoring requests – 50 services per minute
- Advanced monitoring requests – 5 services per minute
- Delete monitoring requests – 30 services per minute

The rate-limiting process also extends to the monitoring data, that is metrics and Events of Significance (EOS), sent by Crosswork Data Gateways to the Crosswork Tracker component. For example, during a restore process, when all Crosswork Data Gateways send metrics again to the Crosswork Tracker component, the rate at which the Crosswork Tracker processes this data and forwards it to Assurance Graph Manager is regulated. This may lead to a delayed reporting of Events of Significance (EOS) following the restore.

In the event of delays, an event is triggered with a severity level of 'Warning' and a corresponding description to notify you of the delay. The event is cleared once Service Health resumes normal publishing of monitoring requests.

- **Ability to monitor performance metrics of TE policies using SR-PM:** To measure the performance metrics of VPN services using the SR-MPLS or RSVP-TE Traffic Engineering policies, Service Health leverages Segment Routing Performance Measurement (SR-PM). This feature enables measuring metrics on the underlay SR-TE policy to enforce Service Level Agreements in VPN services.
- **Monitor service health with external probes from Accedian Skylight:** Crosswork Network Controller can leverage external probing, provided by Accedian Skylight, to measure metrics of the network services. The metrics are compared with the contracted SLA (defined in the Heuristic package), and the results are made available on the Crosswork Network Controller UI.

After an L3VPN service is provisioned and service monitoring is enabled, the probe intent and probe topology are learned (from provisioned service) and a probe session to monitor the service starts automatically by invoking relevant RESTConf APIs. Service Health processes the metrics and raises symptoms as needed to be displayed on the UI. You can view historical data for upto 24 hours from the Probe Sessions.

The maximum number of probe sessions per service are capped at 200 (for all connection types).



Note Accedian Skylight integration is available as a limited-availability feature in this release. Engage with your account team for more information.

Topology

- **Simplified Topology Rebuild Tool:** If the topology is not displaying status as expected, you can now place the system into maintenance mode and then choose to rebuild the topology. This will force the system to create a new topology model and avoid the complicated steps from previous versions.



Note Only users with write permission can Rebuild Topology.

Crosswork Data Gateway

- **Ability to reattempt the import of Controller Certificate file:** When Crosswork Infrastructure and Crosswork Data Gateway are deployed simultaneously, on the first reboot Data Gateway attempts to download the Controller Certificate file from Crosswork Infrastructure. If the Infrastructure deployment is in-progress, Crosswork Data Gateway may not find the certificate. In the past, you had to wait for the Data Gateway VM to restart before downloading the certificate through the Interactive Console menu.

With Crosswork Data Gateway's latest release, you can let Data Gateway retry the certificate download multiple times. If the file download fails, the Crosswork Data Gateway will now retry automatically.

For information on importing the certificate, see the *Import Controller Signing Certificate File* section in [Cisco Crosswork Network Controller 6.0 Installation Guide](#).

- **Dynamic reallocation of the Crosswork Data Gateway vCPU resources:** The Crosswork Data Gateway vCPU resources are now dynamically configured to meet the scaling requirements in response to the number of CPUs assigned to the VM.
- **Parameter to configure the CLI session timeouts for devices:** The **SSH Session Timeout** parameter is implemented to indicate the duration of the CLI connection on a device.

For information on how to configure the **SSH Session Timeout** parameter, see the *Configure Crosswork Data Gateway Global Parameters* section in [Cisco Crosswork Network Controller 6.0 Administration Guide](#).

- **Changes to the Crosswork Data Gateway APIs:**

The Crosswork Data Gateway APIs have been altered in the following ways:

- The new dg-manager APIs are compatible with the OpenAPI v2/v3 specification.
- The change logs include the deprecated APIs. In the subsequent release, the deprecated APIs are removed.
- A change log is created for each modified API. The change log includes the APIs that have been deprecated, removed, or updated.

For information on change logs, see [Cisco Devnet](#).

- **Netconf Collector support is decommissioned:** The NETCONF collector enabled data collection over the NETCONF protocol.

Support for the NETCONF collector has been discontinued in configurations, such as the base VM, application layer, Docker, and dg-manager.

Infrastructure

- **Device Level RBAC:** This release introduces role-based access control (RBAC) at a device granularity for provisioning and device configuration workflows. Each user must be assigned a role that determines what functions they can access along with a Device Group that determines on which devices they can manage or deploy services. For more information, see the *Manage Device Access Groups* section in the [Cisco Crosswork Network Controller 6.0 Administration Guide](#).
- **Geo Redundancy:** This release introduces the first phase of the geo redundancy solution for Crosswork Network Controller and its components in case of a region or data center failure. For more information, see the *Enable Geo Redundancy* section in the [Cisco Crosswork Network Controller 6.0 Installation Guide](#).



Note Geo Redundancy is accessible through Limited Availability. Engage your account team for further details.

Documentation

- An [Information Portal](#) is now available for Crosswork Network Controller 6.0. Information is categorized per functional area, making it easy to find and easy to access.
- [Cisco Crosswork Network Controller 6.0 Service Health Monitoring](#) is a new Crosswork Network Controller specific guide that provides information on monitoring the health of L2VPN and L3VPN services. It provides insights into analyzing and troubleshooting degraded services, as well as visualizing service health status and logical dependency trees.
- [Cisco Crosswork Network Controller 6.0 Traffic Engineering and Optimization](#) is a new Crosswork Network Controller specific guide that provides information on how to visualize and configure traffic engineering in Crosswork Network Controller.
- [Cisco Crosswork Network Controller 6.0 Network Bandwidth Management](#) is a new Crosswork Network Controller specific guide that provides information on how to use Crosswork Network Controller feature packs. Feature packs are tools that tackle congestion mitigation and the management of SR-TE policies to find and maintain intent based bandwidth requirements.

Supported Use Cases

Crosswork Network Controller supports a wide range of use cases allowing operators to manage many aspects of the network. The following describes specific use cases, with details about the Crosswork applications needed to deliver each capability.

- **Orchestrated service provisioning:** Provisioning of layer 2 VPN (L2VPN) and layer 3 VPN (L3VPN) services with underlay transport policies to define, meet, and maintain service-level agreements (SLA), using the UI or APIs. Using [Segment Routing Flexible Algorithm](#) (Flex-Algo) provisioning and visualizing to customize and compute IGP shortest paths over a network according to specified constraints.



Note An SLA defines the expectations set between a service provider and a customer. The SLA details the products or services that are to be delivered, the point of contact for end-user issues, and metrics by which the effectiveness of the process is both monitored and approved.

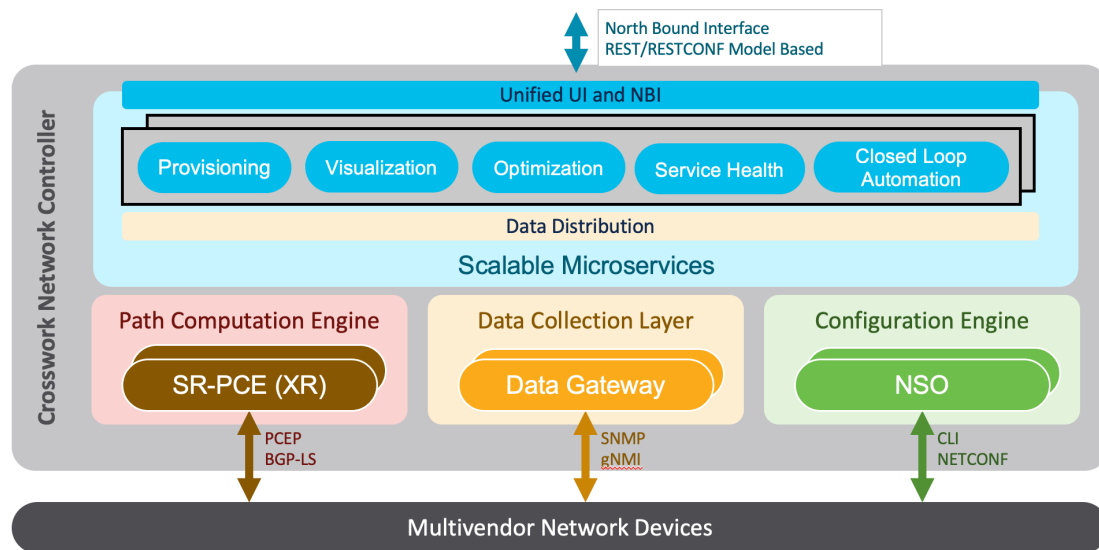
- **Real-time network and bandwidth optimization:** Intent-based closed-loop automation, congestion mitigation, and dynamic bandwidth management based on Segment Routing and RSVP-TE. Optimization of bandwidth resource utilization by setting utilization thresholds on links and calculating tactical alternate paths when thresholds are exceeded. The ability to provision SR-Circuit Style policies and visualize them in your network topology provides:
 - Straightforward verification of SR-Circuit Style policy configurations
 - Visualization of SR-Circuit Style details, bi-directional active and candidate paths

- Operational status details
 - Failover behavior monitoring for individual SR-Circuit Style policies
 - A percentage of bandwidth reservation for each link in the network
 - Manually triggered recalculations of existing SR-Circuit Style policy paths that may no longer be optimized due to network topology changes
- **Local Congestion Management:** Local Congestion Mitigation (LCM) provides localized mitigation recommendations within surrounding interfaces, with the use of standard protocols. Data is gathered in real-time and when congestion is detected, solutions are suggested. LCM supports deployment as either "human in the loop" or fully automated implementations allowing the operator to decide how they want to use the feature. See the Local Congestion Mitigation chapter in the Crosswork Network Controller 6.0 Network Bandwidth Management guide for more information.
 - **Visualization of network and service topology and inventory:** Visibility into device and service inventory and visualization of devices, links, and transport or VPN services and their health status on maps with logical or geographical contexts.
 - **Performance-based closed-loop automation:** Automated discovery and remediation of problems in the network by allowing Key Performance Indicator (KPI) customization and monitoring of pre-defined remediation tasks when a KPI threshold is breached. For this use case, Cisco Crosswork Health Insights and Cisco Crosswork Change Automation must be installed.
 - **Planning, scheduling, and automating network maintenance tasks:** Scheduling an appropriate maintenance window for a maintenance task after evaluating the potential impact of the task (using WAE Design). Automating the execution of maintenance tasks (such as throughput checks, software upgrades, SMU installs) using playbooks. For this use case, Cisco Crosswork Health Insights and Change Automation must be installed.
 - **Secured zero-touch onboarding and provisioning of devices:** Onboarding new IOS-XR devices and automatically provisioning Day0 configuration resulting in faster deployment of new hardware at lower operating costs. For this use case, Cisco Crosswork Zero Touch Provisioning must be installed.
 - **Visualization of native SR paths:** Visualizing the native path using the traceroute SR-MPLS multipath command to get the actual paths between the source and the destination can be achieved using Path Query. With Cisco Crosswork Network Controller, a traceroute command runs on the source device for the destination TE-Router ID and assists in retrieving the paths.
 - **Provision, Visualize, and Analyze Tree Segment Identifier Policies in Multipath Networks:** Creating and visualizing static Tree-SID policies using the UI. Static mVPN Tree-SID policies associated with existing, or newly created, L3VPN service models (SR MPLS point-to-multi-point) using the Crosswork Network Controller that can also be visualized and analyzed to assist in efficient management and troubleshooting of your multicast network. Configuring link affinities used to specify the link attributes that determine which links are suitable to form a path for the Tree-SID policy and maps each bit position or attribute with a color (making it easier to refer to specific link attributes). Modifying existing static Tree-SID policies and mVPN Tree-SID policies associated with an L3VPN service model – both edit and delete – using the UI.
 - **Transport Slice Provisioning:** Cisco Crosswork Network Controller offers direct support for network slicing at the OSI transport layer. Using this solution, network engineering experts can design slice profiles around customer intents and then add them to a catalog. Network line operators can then simply pick the slice that best meets the customer's needs, specify the slice endpoints, and (where needed) set any custom constraints or options built into the chosen slice. Using the UI, you can inspect the slice

details for active symptoms, failures, and root causes. In addition, the slice can be visualized on a geographical map.

Solution Components Overview and Integrated Architecture

The following diagram provides a high-level illustration of how the solution's components work together within a single pane of glass to execute the primary supported use cases.



The following components make up the Cisco Crosswork Network Controller 5.0 solution:

Cisco Crosswork Active Topology

Cisco Crosswork Active Topology's logical and geographical maps provide real-time visibility into the physical and logical network topology, service inventory, and SR-TE policies and RSVP-TE tunnels, all within a single pane of glass. They enable operators to see, at-a-glance, the status and health of the devices, services, and policies. Services and transport policies can be visualized end-to-end as an overlay within the context of the topology map. Cisco Crosswork Active Topology provides device grouping functionality so that operators can set up their maps to monitor exactly the set of devices, services, and locations for which they are responsible. In addition, operators can save custom views for quick and easy access to the views and functionality they use on an ongoing basis.

Cisco Crosswork Optimization Engine

Cisco Crosswork Optimization Engine provides real-time network optimization allowing operators to effectively maximize network capacity utilization, as well as increase service velocity. Leveraging real-time protocols, such as BGP-LS and Path Computation Element Communication Protocol (PCEP), SR-PCE and Crosswork Optimization Engine enables closed-loop tracking of the network state, reacting quickly to changes in network conditions to support a self-healing network.

Cisco Service Health

- Service Health substantially reduces the time required to detect and troubleshoot service quality issues. It monitors the health status of provisioned L2/L3 VPN services and enables operators to pinpoint why and where a service is degraded. It can also provide service-specific monitoring, troubleshooting, assurance, and proactive causality through a heuristic model that visualizes the:
 - Health status of sub-services (device, tunnel) to a map when a single service is selected
 - Service logical dependency tree and help the operator in troubleshooting in case of degradation by locating where the problem resides, an indication of possible symptoms, and impacting metrics in case of degradation
 - Historical view of service health status up to 60 days

Service Health also provides the following:

- Service Health monitoring is available for both Basic Monitoring and Advanced Monitoring options. For help selecting the appropriate monitoring option for your needs, see the section **Basic and Advanced Monitoring Rules**.
- Service Health provides **Internal Storage** of monitoring data up to a maximum limit of 50 GB. This data is stored on your system. If you exceed the limit of the internal storage, historical data will be lost. If you choose to extend Service Health storage capacity, you can optionally configure **External Storage** in the cloud using an Amazon Web Services (AWS) cloud account. By leveraging External Storage, all existing internal storage data will be automatically moved to the external cloud storage (see **Configuring Service Health External Storage Settings** appendix for more details) and your internal storage will act locally as cache storage. Configuring External Storage for Service Health ensures you will not lose historical data for services that continue to monitor a service's health, and will retain service health data for any service you choose to stop monitoring when you select the option to retain historical monitoring service for the data. For more information on Internal and External Storage, and how to retain historical monitoring service data when stopped, see the Appendix sections **Configuring Service Health External Storage Settings** and **Stopping Service Health monitoring**.



Note If you anticipate monitoring a large amount of Service Health services, Cisco recommends you configure External Storage after you install Service Health and before you begin monitoring services so to avoid exceeding the Internal Storage and losing historical data.

- To view subservices supported by Service Health L2VPN/L3VPN, see the **Service Health Supported Subservices** appendix section. Details are provided that define which subservices are supported by each VPN service flavor.
- Service Health supports point-to-point L2VPN.



Note Currently, Service Health does not support multipoint L2VPN.

- Service Health supports integration with standalone Network Services Orchestrator (NSO) or NSO Layered Service Architecture (LSA).

- NSO LSA support is limited to one CFS node and two RFS nodes. These additional NSO types serve as a high availability feature. By distributing your devices across the different types, the LSA feature in Service Health allows for dynamic configurations for assurance.

To manage the Service Health provider Access, select **Administration > Manage Provider Access**. The Providers screen appears. See the Crosswork Administration guide and NSO documentation for additional, detailed information.

- The Service Health Collection Jobs administrative option provides the capability to view Parameterized Jobs (template-based collection jobs) that supports a greater number of jobs, adding the ability to view CLI collection jobs. This is useful when troubleshooting collection job issues by examining details of individual devices using Parameterized Jobs. Devices are identified by their Context ID (protocol) to determine if they are GMNI, SNMP, or CLI-based jobs. Additionally, you may export the collection job information to review. The information is collected at the time the export is initiated and stored in a .csv file.



Note When exporting the collection status, you must fill in the information each time an export is executed. In addition, make sure to review the Steps to Decrypt Exported File content available on the Export Collection Status pop up to ensure you can access and view the exported information.

- Service Health provides expanded redundancy/High Availability (HA) for Assurance Graph Manager, Expression Orchestrator, and Crosswork Expression Tracker microservices (two instances are now available). To view, select **Administration > Crosswork Manager**. In the Crosswork Summary tab, select Crosswork Service Health to view the Application Details screen and Microservices.
 - For example, if you click the Assurance Graph Manager, two redundant/high availability instances appear. In certain situations, one of the instances will be in the active-active mode while the other is in the active-standby mode. This ensures that if one instance goes down, the second acts as a redundant, HA, backup.
- Heuristic Packages: Three additional Rules have been added to assist in Basic monitoring level rules, where a rule to generate Assurance Graph information, for example Basic L2VPN NM P2P services, can be used along with two sub services:
 - Rule-L2VPN-NM- Basic
 - Rule-L2VPN-NM-P2P-Basic
 - Rule-L3VPN-NM-Basic
- Heuristic Package Metrics now has the capability for CLI based metrics and GMNI filtering customizations of packages.

Cisco Crosswork Data Gateway

Cisco Crosswork Data Gateway is a secure, common collection platform for gathering network data from multi-vendor devices. It is an on-premise application deployed close to network devices. Crosswork Data Gateway supports multiple data collection protocols including MDT, SNMP, CLI, standards-based gNMI (dial-in), and syslog. Any type of data can be collected by Crosswork Data Gateway as long as it can be

delivered over one of the supported protocols. In this way, it can provide support for a growing set of use cases and customizations.

To address scale challenges, Cisco Crosswork Data Gateway is implemented as a number of VMs and designed with a distributed architecture in mind. Each lightweight VM manages a subset of the overall network and as the network grows, additional VMs can be added horizontally to address the new demands on the compute resources. It also supports a flexible redundancy configuration based on the operator's needs. After the initial setup, Cisco Crosswork Network Controller automatically orchestrates the collection across the multiple Cisco Crosswork Data Gateway VMs.

APIs and configuration examples are available to illustrate how to add new collection jobs (outside of those built for you by Cisco Crosswork Network Controller) to gather additional information from your network. The collected data can be published to approved destinations. Supported destinations are Kafka and gRPC messaging bus.

Crosswork Common UI and API

All Cisco Crosswork Network Controller's functionality are provided within a single, common graphical user interface. This common UI brings together the features of all Crosswork Network Controller's components, including common inventory, network topology and service visualization, service and transport provisioning, and system administration and management functions. When optional add-on Crosswork components are installed, their functionalities are also fully integrated into the common UI. Having all functionality within a common UI, instead of having to separately navigate individual application UIs, enhances the operational experience and increases productivity.

A common API enables Crosswork Network Controller's programmability. The common APIs provides a single access point for all APIs exposed by various built-in components. The API provides a REST-based Northbound Interface to external systems (e.g., OSS systems) to integrate with Cisco Crosswork Network Controller. RESTCONF and YANG data models are made available for optimization use cases. For details about the APIs and examples of their usage, see the [Cisco Crosswork Network Automation API Documentation on Cisco DevNet](#).

Crosswork Infrastructure and Shared Services

The Cisco Crosswork Infrastructure provides a resilient and scalable platform on which all Cisco Crosswork components can be deployed. This infrastructure and shared services provide:

- A single API endpoint for accessing all APIs of Crosswork applications deployed
- A shared Kafka bus to pass data between applications
- Shared database(s) (such as relational and graph) for applications to store data
- A single shared database to store all gathered time-series data from the network
- A robust Kubernetes-based orchestration layer to provide for process-level resiliency
- Tools for monitoring the health of the infrastructure and the cluster of virtual machines (VMs) on which it resides

Cisco Crosswork Health Insights and Cisco Crosswork Change Automation

Cisco Crosswork Health Insights and Cisco Crosswork Change Automation are components that can optionally be installed with Cisco Crosswork Network Controller.

Cisco Crosswork Health Insights performs real-time Key Performance Indicator (KPI) monitoring, alerting, and troubleshooting. Cisco Crosswork Health Insights enables programmable monitoring and analytics. It provides a platform dynamically for addressing changes to the network infrastructure. Cisco Crosswork Health Insights builds dynamic detection and analytics modules that allow operators to monitor and alert about network events based on user-defined logic.

Cisco Crosswork Change Automation automates the process of deploying changes to the network. Orchestration is defined via an embedded Ansible Playbook and then configuration changes are pushed to Cisco Network Services Orchestrator (NSO) to be deployed to the network.

These components within Cisco Crosswork Network Controller enable closed-loop discovery and remediation of problems in the network. Operators can match alarms to pre-defined remediation tasks, which are performed when a defined Key Performance Indicator (KPI) threshold is breached. This reduces the time it takes to discover and repair a problem while minimizing the risk of human error resulting from manual network operator intervention.

Cisco Crosswork Zero-Touch Provisioning (ZTP)

Cisco Crosswork ZTP can optionally be installed with Cisco Crosswork Network Controller.

Cisco Crosswork ZTP is an integrated turnkey solution for automatically onboarding and provisioning new IOS-XR devices, resulting in faster deployment of new hardware at lower operating costs. Operators can quickly and easily bring up devices using a Cisco-certified software image and a day-zero software configuration. After it is provisioned in this way, the new device is onboarded to the Crosswork device inventory where it can be monitored and managed along with other devices.

Cisco Crosswork ZTP offers Secure ZTP functionality in addition to the Classic ZTP functionality. Secure ZTP is based on RFC 8572 standards and uses secure transport protocols and certificates to verify devices and perform downloads. Secure ZTP is useful when public Internet resources must be traversed to reach remote network devices, or when the devices are from third-party manufacturers. With Secure ZTP, the device and the Cisco Crosswork ZTP bootstrap server authenticate each other using the device's Secure Unique Device Identifier (SUDI) and Crosswork server certificates over TLS/HTTPS. After a secure HTTPS channel is established, the Crosswork bootstrap server allows the device to request to download and apply a set of signed image and configuration artifacts adhering to the RFC 8572 YANG schema. After the image (if any) is downloaded and installed, and the device reloads with the new image, the device downloads configuration scripts and executes them.

Cisco Network Services Orchestrator

Cisco Network Services Orchestrator (NSO) is an orchestration platform that makes use of pluggable function packs to translate network-wide service intent into device-specific configuration. Cisco NSO provides flexible service orchestration and lifecycle management across physical network elements and cloud-based virtual network functions (VNFs), fulfilling the role of the Network Orchestrator (NFVO) within the ETSI (European Telecommunications Standards Institute) architecture. It provides complete support for physical and virtual network elements, with a consistent operational model across both. It can orchestrate across multi-vendor environments and support multiple technology stacks, enabling the extension of end-to-end automation to virtually any use case or device.

Cisco NSO has a rich set of APIs designed to allow developers to implement service applications. It provides the infrastructure for defining and executing the YANG data models that are needed to realize customer services. It is also responsible for providing the overall lifecycle management at the network service level.

Service and device models, written using YANG modelling language, enable Cisco NSO to efficiently 'map' service intent to device capabilities and automatically generate the minimum required configuration to be deployed in the network. This feature, facilitated by Cisco NSO's FASTMAP algorithm, is capable of comparing

current configuration states with a service's intent and then generating the minimum set of changes required to instantiate the service in the network.

All Crosswork components that are included in Cisco Crosswork Network Controller or are optional add-ons, with the exception of Cisco Crosswork ZTP, require integration with Cisco NSO.

Cisco Crosswork Network Controller requires the following Cisco NSO function packs:

- SR-TE core function pack (CFP) enables provisioning of explicit and dynamic segment routing policies, including SRv6, and on-demand SR-TE policy instantiation for prefixes with a specific color.
- Sample function packs for IETF-compliant L2VPN and L3VPN provisioning. These function packs provide baseline L2VPN and L3VPN provisioning capabilities, based on IETF NM models. Prior to customization, these sample function packs enable provisioning of the following VPN services:
 - L2VPN:
 - Point-to-point VPWS using Targeted LDP
 - Point-to-point VPWS using EVPN
 - Multipoint VPLS using EVPN (with service topologies ELAN, ETREE, and Custom)
 - L3VPN
- Sample IETF-compliant RSVP-TE function pack intended as a reference implementation for RSVP-TE tunnel provisioning, to be customized as required.



Note

- By default, the IETF-compliant NM models are used. If your organization wishes to continue to use the Flat models that were provided with the previous version, a manual setup process is required.
 - The Cisco NSO sample function packs are provided as a starting point for VPN service provisioning functionality in Cisco Crosswork Network Controller. While the samples can be used “as is” in some limited network configurations, they are intended to demonstrate the extensible design of Cisco Crosswork Network Controller. Answers to common questions can be found on Cisco Devnet and Cisco Customer Experience representatives can provide answers to general questions about the samples. Support for customization of the samples for your specific use cases can be arranged through your Cisco account team.
 - Cisco NSO currently does not support bundle ethernet (BE), route distinguisher (RD), or BGP route-target (RT) functions with L2VPN EVPN. Although it does support multihoming and L2VPN route policy, there is no option to specify an RD value in L2VPN for an EVPN ELAN/ETREE, nor is there an option to specify load balancing type. To perform these functions, contact your Cisco account team for a set of custom configuration templates and advice on configuring bundles manually.
-

Cisco Segment Routing Path Computation Element (SR-PCE)

Cisco SR-PCE is an IOS-XR multi-domain stateful PCE supporting both segment routing (SR) and Resource Reservation Protocol (RSVP). Cisco SR-PCE builds on the native Path Computation Engine (PCE) abilities within IOS-XR devices, and provides the ability to collect topology and segment routing IDs through BGP-LS, calculate paths that adhere to service SLAs, and program them into the source router as an ordered list of segments. A Path Computation Client (PCC) reports and delegates 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 and re-optimize paths where necessary.

Cisco SR-PCE can either reside on server resources using virtualized XRv9000, or as a converged application running within IOS-XR Routers.



Note Adding static routes for auto-discovering the scale nodes from SR-PCE after 2,000 nodes is not supported.

Multi-Vendor Capabilities

Today's networks have typically been built up over time and incorporate multiple vendors and multiple generations of hardware and software. Furthermore, there is a lack of industry standardization, making support for these networks using a single tool challenging.

Service providers require an integrated solution to manage third-party devices that will reduce operational expenses and maintenance overhead, as well as eliminate the need to build custom applications to deploy and maintain different vendor products for a single network.

Because it uses standards-based protocols, Cisco Crosswork Network Controller has multi-vendor capabilities for:

- Network service orchestration via Cisco Network Services Orchestrator using CLI and Netconf/YANG. Cisco Network Services Orchestrator is a YANG model-driven platform for automating provisioning, monitoring, and managing applications and services across multi-vendor networks.
- Telemetry data collection using SNMP with standards-based MIBs, syslog, and gNMI with standard OpenConfig models. Cisco Crosswork Data Gateway also supports Native YANG data models for external destinations and proprietary SNMP MIBs with custom packages.
- Topology and transport discovery via SR-PCE, using IGP and BGP-LS, with link utilization and throughput collected via SNMP using standard MIBs.
- Transport path computation using PCEP.



Note For third-party network device support, use cases must be validated by Cisco Customer Experience representatives in the customer's multi-vendor environment, especially if legacy platforms and non-standard devices or services are involved.

Extensibility

The Cisco Crosswork Network Controller provisioning functionality can be extended using the application programming interfaces (APIs). For more information about the APIs, see the [Cisco Crosswork Network Automation API Documentation on Cisco DevNet](#).

The provisioning UI is extensible as it is rendered based on the YANG model. When new services are introduced, they can be easily incorporated.



CHAPTER 2

UI Overview

This section explains the following topics:

- [Log In](#), on page 19
- [Dashboard](#), on page 20
- [Navigation](#), on page 20
- [Commit – Advanced Options](#), on page 23

Log In

Log into the web UI by entering the following URL in the browser's address bar:

```
https://<Crosswork Management Network Virtual IP (IPv4)>:30603/  
https://[<Crosswork Management Network Virtual IP (IPv6)>]:30603/
```



Note The IPv6 address in the URL must be enclosed with brackets.

In the Log In window, enter the username and password configured during installation and click **Log In**.

Self-signed certificate: At first-time access, 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 download the certificate, the browser accepts the server as a trusted site in all future login attempts.


CA signed certificate: For production use, a CA signed certificate may be installed and is recommended to avoid a warning that the site is untrusted.



Note For information on installing CA signed certificates, see the [Manage Certificates](#) topic in the *Cisco Crosswork Network Controller Administration Guide*.

Dashboard

After successful login, the Home page opens. The Home page displays the dashboard which provides an at-a-glance operational summary of the network being managed. The dashboard is made up of a series of dashlets. The specific dashlets included in your dashboard depend on which Cisco Crosswork applications you have installed. Links in each dashlet allow you to drill down for more details.



Note Your Dashboard may differ from this screen capture, which displays optional components you may not have installed.

Navigation

The main menu along the left side of the window provides access to all features and functionality in Cisco Crosswork Network Controller, as well as to device management and administrative tasks. The Dashboard, Topology, Services & Traffic Engineering, Device Management and Administration menu options are available when all native components of Cisco Crosswork Network Controller are installed. Additional menu options are available in the main menu depending on which Cisco Crosswork add-on applications are installed.

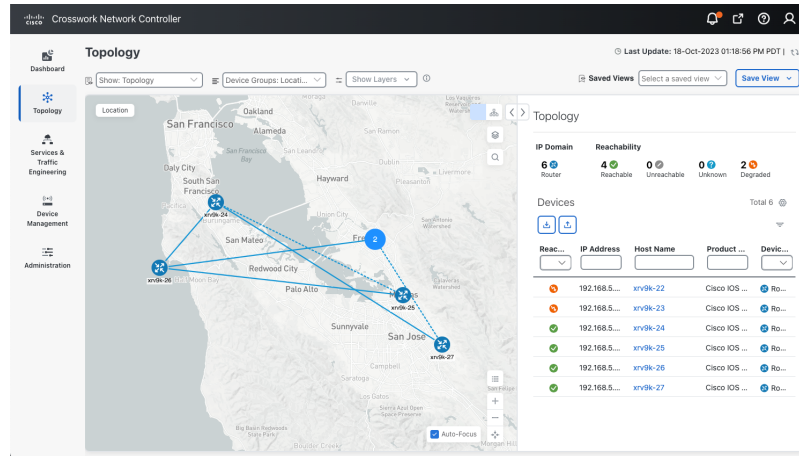
Dashboard

The home page contains the dashboard, as described in the [Dashboard](#) topic.

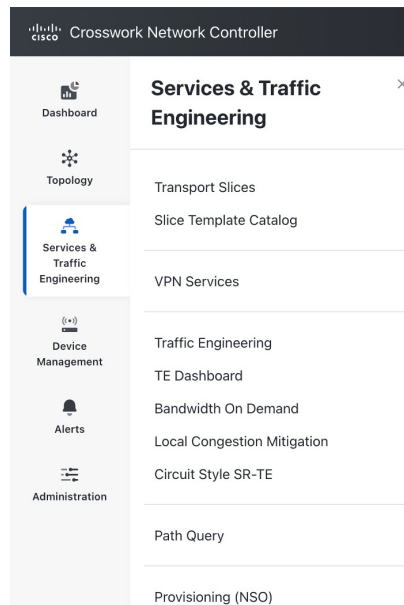
Topology

Users can display the network device and link topology on a logical map or a geographical (geo) map. The logical map shows devices and their links, positioned according to an automatic layout algorithm. The geo map shows single devices, device groups, 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). Operators supply this location information in CSV or KML files uploaded using the Crosswork UI.

The Topology page consists of a map showing managed devices and the links between them, along with a device table listing managed devices. In the map you can see the status and health of the devices at a glance. Clicking on a device in the table highlights the device on the map and shows details of the device and its associated links. Use the toggle buttons to switch between the geographical map (shown below) and the logical map. Clicking on the question mark in the map provides a detailed legend of the various symbols and their meaning.



Services & Traffic Engineering



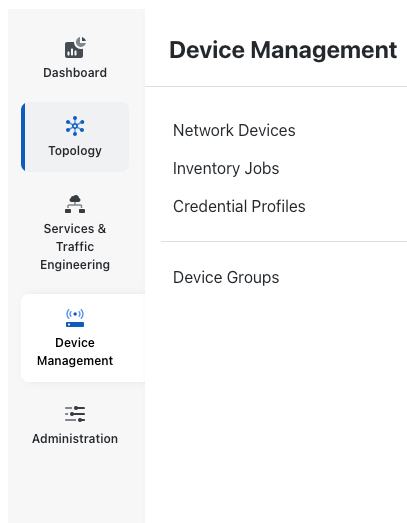
The Services & Traffic Engineering menu provides access to VPN and transport provisioning and visualization functionality, bandwidth management functionality, as well as access to the configuration pages used to enable Feature Packs. For more information, click [here](#) to see the Crosswork Optimization Engine 6.0 User Guide.

Choose **VPN services** or **Traffic Engineering** to see managed VPN services, SRv6 policies, or SR-TE policies/RSVP-TE tunnels within the context of a logical or geographical map.

Choose **Provisioning (NSO)** to access the provisioning UI rendered from the Cisco Network Services Orchestrator models. Here you can create L2VPN and L3VPN services, SR-TE policies, SRv6 policies, SR


ODN templates, and RSVP-TE tunnels. You can also create the resources required for these services and policies, such as resource pools, route policies for L2VPN and L3VPN services, and SID lists for SR-TE policies. SR-TE policies and RSVP-TE tunnels can be attached to VPN services to define and maintain SLAs by tracking network changes and automatically reacting to optimize the network.

Device Management



The Device Management menu provides access to device-related functionality, including adding, managing, and grouping devices, creating and managing credential profiles, and viewing a history of device-related jobs.

Administration



The Administration menu provides access to all system management functions, data gateway management, Crosswork cluster and application health, backup and restore, smart licensing and other setup and maintenance functions that are typically performed by an administrator.

Click [here](#) to see the Crosswork Network Controller 6.0 Administration Guide for information about these functions.

Commit – Advanced Options

Before committing a configuration, commit advanced options may be available. These options are for advanced users only. Use of the commit options is recommended after pre-validation in a lab environment.


In the Advanced Options window, the following options are available.

- **Commit Queue:** While the configuration change is committed to configuration database (CDB) immediately it is not committed to the actual device but rather queued for eventual commit in order to increase transaction throughput. This enables use of the commit queue feature for individual commit commands without enabling it by default. There are two operation modes: async, sync.
 - **async mode:** The async mode operation returns successfully if the transaction data has been successfully placed in the queue.
 - **sync mode:** The sync mode will cause the operation to not return until the transaction data has been sent to all devices, or a timeout occurs. If the timeout occurs the transaction data stays in the queue and the operation returns successfully. The timeout value can be specified with the timeout or infinity

option. By default the timeout value is determined by what is configured in `/devices/global-settings/commit-queue/sync`.

Commit

Advanced Options ^

 The below options are for advanced users only. Use of commit options is recommended after pre-validation in a lab environment.

Commit Queue

No Out of Sync Check

No Overwrite

[Cancel](#) [Commit](#)

- **Atomic:** The atomic option sets the atomic behavior of the resulting queue item. If this is set to false, the devices contained in the resulting queue item can start executing if the same devices in other non-atomic queue items ahead of it in the queue are completed. If set to true, the atomic integrity of the queue item is preserved.
 - **Block Others:** This option will cause the resulting queue item to block subsequent queue items which use any of the devices in this queue item, from being queued.
 - **Lock:** This option will place a lock on the resulting queue item. The queue item will not be processed until it has been unlocked, see the actions unlock and lock in `/devices/commit-queue/queue-item`. No following queue items, using the same devices, will be allowed to execute as long as the lock is in place.

Commit

Advanced Options ^

⚠ The below options are for advanced users only. Use of commit options is recommended after pre-validation in a lab environment.

Commit Queue

Atomic

Block Others

Lock

Tag

Timeout

Error Option

No Out of Sync Check

No Overwrite

In addition, the commit-queue flag has a number of other useful options that affects the resulting queue item.

- **Tag:** This option sets a user defined, opaque tag that is present in all notifications and events sent referencing the queue item..
- **Timeout:** The timeout value can be specified with the timeout or infinity option. By default the timeout value is determined by what is configured in `/devices/global-settings/commit-queue/sync`.
- **Error Option:** Depending on the Error Option selected, NSO will store the reverse of the original transaction to be able to undo the transaction changes and get back to the previous state. This data is stored in the `/devices/commit-queue/completed` tree from where it can be viewed and invoked with the rollback action. When invoked the data will be removed. There are two values available: `continue-on-error`, `stop-onerror`.
 - **continue-on-error:** The continue-on-error value means that the commit queue will continue on errors. No rollback data will be created.
 - **stop-onerror:** The stop-on-error means that the commit queue will place a lock on the failed queue item, thus blocking other queue items with overlapping devices to be executed. The lock must then either manually be released when the error is fixed or the rollback action under `/devices/commit-queue/completed` be invoked.
- **No Out of Sync Check:** Commit even if out of sync.
- **No Overwrite:** Do not overwrite modified data on the device.



CHAPTER 3

Orchestrated Service Provisioning

This section explains the following topics:

- [Overview, on page 27](#)
- [Scenario: Implement and Maintain SLA for an L3VPN Service for SR-MPLS \(using ODN\), on page 29](#)
- [Scenario: Implement and Maintain SLA for an L3VPN Service for SRv6 \(using ODN\), on page 50](#)
- [Scenario: Mandate a Static Path for an EVPN-VPWS Service using an Explicit MPLS SR-TE Policy, on page 63](#)
- [Scenario: Provision an L2VPN Service over an RSVP-TE Tunnel with Reserved Bandwidth , on page 78](#)
- [Scenario: Provision a Soft Bandwidth Guarantee with Optimization Constraints, on page 83](#)

Overview

By using the scenario workflows described in this section, we are providing examples of how to configure the system to deliver the operator's intended configuration. These scenarios do not fully demonstrate all of the capabilities of Crosswork Network Controller. They are intended to demonstrate the flexibility of the platform. Additional customization is possible either by leveraging the resources available on Cisco DevNet or through engagement with Cisco Customer Experience.

Objective

Provision a set of VPN services with underlay transport policies that will meet and maintain service-level agreements (SLAs) between the service provider and the customer. An SLA defines the service-delivery expectations agreed upon between the service provider and the customer. The SLA details the products or services that the provider is to deliver to the customer, the provider's point of contact to which the customer will bring service issues, and the metrics the provider and customer both use to monitor compliance with the SLA.

Challenge

The service-provider network state changes continuously and so quickly that it is difficult to track and react to network problems fast enough to avoid congestion and maintain SLA compliance. In a typical lifecycle, there is a feedback loop that traditionally requires manual monitoring and intervention, which is time- and resource-intensive.

Solution

With network automation, the objective is to automate the feedback loop to enable quicker reaction to and remediation of network events. With Crosswork Network Controller, network operators can orchestrate L2VPN

and L3VPN services across the transport network, via a programmable interface, in a very quick and efficient manner. Segment routing traffic engineering (SR-TE) policies can be configured to continuously track network changes and automatically react to optimize the network. These SR-TE policies can serve as the underlay configuration for the VPN services to automatically maintain the SLAs.

The services required for this solution can be created and managed using the Crosswork Network Controller UI. L2/L3 VPN Yang model-based service intents are implemented using the Cisco Network Services Orchestrator sample function packs, which provide sample service models that can be extended and fine-tuned to meet customer needs. Optionally, Service Health monitoring can be enabled to see which services are working as provisioned, if issues have been flagged, and what symptoms are detailed so to quickly address and fix.



Note The Network Services Orchestrator sample function packs are provided as a starting point for VPN service provisioning functionality in Crosswork Network Controller. While the samples can be used “as is” in some limited network configurations, they are intended to demonstrate the extensible design of Crosswork Network Controller. Answers to common questions can be found on Cisco Devnet and Cisco Customer Experience can provide answers to general questions about the samples. Support for customization of the samples for your specific use cases can be arranged through your Cisco account team.



Note Some scenario features and functions belonging to multiple components (such as Crosswork Optimization Engine, Crosswork Service Health, Crosswork Active Topology) will not be available as described unless all of the applications are successfully deployed.

How Does it Work?

1. User creates an SR-TE policy/On-Demand Next Hop (ODN) template with intent (e.g., bandwidth, latency) using the Cisco Crosswork Network Controller UI or APIs.
2. User creates a VPN service using the UI or APIs and specifies the following:
 - The endpoints participating in the VPN
 - Other required VPN parameters
 - The SR-TE policy/ODN template that is to be associated with the VPN service
3. During the provisioning process for the above steps, Cisco Network Services Orchestrator configures the SR-TE policy and the VPN service on the specified endpoints.
4. When the service is active, the network interacts with the SR-PCE to dynamically program the path that meets the intent in the configured SR-TE policy/ODN template. The headend device requests a path from the SR-PCE via PCEP (for dynamic SR-TE policies). If the request specifies bandwidth, the SR-PCE gets the path from Cisco Crosswork Optimization Engine.
5. The SR-PCE sends the path to the headend device via PCEP and updates the headend if path changes are required.

Usage Scenarios

We will walk you through the following usage scenarios that illustrate the execution of the orchestrated service provisioning use case using the Cisco Crosswork Network Controller UI:

- [Scenario: Implement and Maintain SLA for an L3VPN Service for SR-MPLS \(using ODN\)](#)
- [Scenario: Implement and Maintain SLA for an L3VPN Service for SRv6 \(using ODN\)](#)
- [Scenario: Mandate a Static Path for an EVPN-VPWS Service using an Explicit MPLS SR-TE Policy](#)
- [Scenario: Provision an L2VPN Service over an RSVP-TE Tunnel with Reserved Bandwidth](#)
- [Scenario: Provision a Soft Bandwidth Guarantee with Optimization Constraints](#)

Additional Resources

- For information about segment routing and segment routing policies, click [here](#) to see the Crosswork Optimization Engine User Guide.
- Cisco Network Services Orchestrator documentation is included in the latest Network Services Orchestrator image [here](#).

Scenario: Implement and Maintain SLA for an L3VPN Service for SR-MPLS (using ODN)

This scenario walks you through the procedure for provisioning an L3VPN service with a specific SLA objective: all traffic for this service must take the lowest-latency path. The customer requires this low-latency path for this service, as all of this service's traffic is high priority. The customer also wants to use disjoint paths; that is, two unique paths that steer traffic from the same source but to two unique destinations, avoiding common links so that there is no single point of failure.


We'll achieve this using Segment Routing (SR) On-Demand Next Hop (ODN). SR ODN allows a service headend router to automatically instantiate an SR-TE policy to a BGP next-hop when required (on-demand). We configure the headend with an ODN template with a specific color that identifies the SLA. Crosswork will optimize the traffic path when it receives a prefix with that SLA-specific color. We define prefixes in a route policy that is associated with the L3VPN.

Crosswork Network Controller continues to monitor the network and will automatically optimize the network based on the defined SLA, in a closed loop.

Within this workflow, we also have the option to enable Crosswork's Service Health monitoring, and to use Flex-Algo as a constraint on how paths are computed and visualized. With Service Health monitoring, operators can gather quick insights into degraded and down services and then use these insights to visualize, inspect, and troubleshoot for improved network optimization.

With Flex-Algo, we can customize IGP shortest-path computations using algorithms we define. IGP will compute paths based on a user-defined combination of metric types and constraints, and present a filtered topology view based on our specific Flex-Algo definitions.

The following topology provides the base for this scenario:



In this scenario, we will:

- Create a segment routing ODN template with a specific color on the endpoints to ensure that traffic is transported within an LSP (underlay) and that a best-path tunnel is created dynamically when a prefix with the specified color is received. The ODN template defines the SLA on which you want to optimize the path. In this case, we will optimize on latency.
- Specify that the computed paths be disjoint: they will not share the same link.
- Create a route policy on each endpoint to be used to bind the L3VPN to the ODN template. This route policy adds a color attribute to the customer prefixes and advertises via BGP to other endpoints. This color attribute is used to indicate the SLA required for these prefixes.
- Create an L3VPN service with 3 endpoints and enable Service Health monitoring.
- Visualize how this overlay/underlay configuration optimizes the traffic path and automatically maintains the SLA while monitoring your service's health.

Assumptions and Prerequisites

- To use ODN, BGP peering for the prefixes must be configured between the endpoints or PEs. Usually for L3VPN, this is the VPNv4 and VPNv6 address family peering.
- For Service Health enablement, Service Health must be installed. See the Crosswork Network Controller Installation Guide chapter, Install Crosswork Applications.
- Before using Service Health's Assurance Graph, ensure that topology map nodes have been fully configured and created with a profile associated to the service. If not, Subservice Details metrics will show that no value has yet to be reported. See the Crosswork Network Controller Service Health Guide for further details.
- L3VPN service monitoring supports XR devices and does not support XE devices. Thus, after an L3VPN service is created and Service Health monitoring is enabled, if a provider and devices are removed, and then added back, service monitoring remains in a degraded state with a METRIC_SCHEDULER error. To recover, service monitoring must be stopped and restarted.
- (Optional) Flexible Algorithms, and the IDs that are used, must be configured in your network.



Note Screen captures, showing services and data, are for example purposes only and may not always reflect the devices or data described in the workflow content.

Step 1 Create an ODN template to map color to an SLA objective and constraints

Disjointness constraints work by associating a disjoint group ID with the ODN template, and all tunnels with the same disjoint group ID will be disjoint, i.e., they will use different links, nodes and shared risk link groups depending on how the disjoint groups are configured.

We will create the following ODN templates:


- Headend PE-A, color 72, latency, disjoint path (link), group ID 16 - L3VPN_NM-SRTE-ODN_72-a
- Headend PE-A, color 71, latency, disjoint path (link), group ID 16 - L3VPN_NM-SRTE-ODN_71-a
- Headend PE-B and PE-C, color 70, latency - L3VPN_NM-SRTE-ODN_70
- Headend PE-B, color 72, latency - L3VPN_NM-SRTE-ODN_72-b
- Headend PE-C, color 71, latency - L3VPN_NM-SRTE-ODN_71-c

For example purposes, we will show how to create the first ODN template - L3VPN_NM-SRTE-ODN_72-a. The other ODN templates can be created using the same procedure.

Before you begin

In this step, we will create an ODN template on each endpoint. The ODN template specifies the color and the intent; in this case, latency and disjointness. This ODN template will be used to dynamically create tunnels (on-demand) when prefixes with matching colors are received via BGP. Traffic to these prefixes will be automatically steered into the newly created tunnels, thereby meeting the SLA objective and constraints intended for these prefixes and signaled using colors in the BGP routes.

Step 1 Go to **Services & Traffic Engineering > Provisioning (NSO) > SR-TE > ODN-Template**.

Step 2 Click  to create a new template and give it a unique name. In this case, the name is **L3VPN_NM-SRTE-ODN_72-a**. Click **Continue**.

You may also browse for an existing template on your system so to import the file. The information from the imported file is populated into the form.

SR-TE > ODN-Template



Name

Step 3 Choose the head-end device, **PE-A**, and specify the color **72**.

Step 4 Under dynamic, select **latency** as the metric-type. This is the SLA objective on which we are optimizing.

Step 1 Create an ODN template to map color to an SLA objective and constraints




Step 5 Select the **pce** check box to specify that the path should be computed by the SR-PCE, not by the Path Computation Client (PCC).

Step 6 Define the required constraints. In this case, we want the computed paths to be disjoint in that they must not share a link. Under disjoint-path, choose **link** as the type, and specify a numeric group ID, in this case, **16**, as the group-id.

Note You may choose the group ID. All paths requested with the same group-id will be disjoint from each other.

Note Optionally, you may configure Flex-Algo as a constraint.

head-end ⓘ ⓘ

name

PE-A

color * ⓘ

72 ⓘ

dynamic ⓘ

Enable dynamic

pce ⓘ

flex-alg ⓘ

metric-type ⓘ

latency

metric-margin ⓘ

affinity ⓘ

segments ⓘ

disjoint-path ⓘ

Enable disjoint-path


type*


link

group-id * ⓘ


16 ⓘ

Step 7 Commit your changes or click **Dry Run** to check what will be configured on the devices before you commit.

Step 8 Check that the new ODN template appears in the table and its provisioning state is **Success**. Click  in the Actions column and choose **Config View** to see the Yang model-based service intent data that details the ODN template you created.



Step 9 Create the other ODN templates listed above in the same manner.


Note You can save some time by using the Clone function to build the other policies needed to complete this scenario. Simply select **Clone** from the  Actions column, provide a new name for the clone, edit the values, and then select **Commit**.

Step 2 Create an L3VPN Route Policy

In this step, we will create a route policy for each endpoint, and we will specify the same color as defined in the ODN template for that endpoint. The route policy defines the prefixes to which the SLA applies. When traffic from the specified network with a matching color is received, paths are computed based on the SLA defined in the ODN template. We will create the following route policies by first setting the routing policy tag and routing policy destination prefix. The routing policy prefixes should match with the subnet prefix configured on the PE devices in the service:

- Color 70, IPv4 prefix 70.70.70.0/30 - L3VPN_NM-SRTE-RP-PE-A-7
- Color 71, IPv4 prefix 70.70.71.0/30 - L3VPN_NM-SRTE-RP-PE-B-7
- Color 72, IPv4 prefix 70.70.72.0/30 - L3VPN_NM-SRTE-RP-PE-C-7

Step 1 Go to **Services & Traffic Engineering > Provisioning (NSO) > L3VPN > Routing Policy Tag**.

Step 2 Click  to create a new routing policy tag and type the name of the tag set: **COLOR_70**. Click **Continue**.

This is used as a label to reference the set in actions and conditions.

Step 3 Under tag-value, click + and type the tag-value: **70**.

The screenshot shows two parts of the configuration interface. On the left, the 'Routing Policy Tag {COLOR_70}' form has 'name' set to 'COLOR_70' and 'tag-value' set to '70'. Below the 'tag-value' field, there are '+', '-', and 'Total 0' icons. A table below shows 'tag-value' with 'No Rows To Show'. On the right, a smaller form shows 'tag-value' set to '70' with a 'Continue' button.

The tag value may be a number between **1 – 4294967295** and should match to a color value.

Step 4 Click **Continue**. The new routing policy tag name with the new tag value is visible. Click **Commit changes**.

Create the other two routing policy tags (**COLOR_71** and **COLOR_72**) and tag values (**71** and **72**) by following the same steps above. Click **Continue**.

Now create the routing policy destination prefixes.

Step 5 Go to **Services & Traffic Engineering > Provisioning (NSO) > L3VPN > Routing Policy Destination Prefix**.

Step 6 Click **+** to create a new routing policy destination prefix and type the name: **DEST_PREFIX_SET_70**.

The name of the prefix set will reference the set in match conditions.

Step 7 For Mode, select **ipv4**.



Step 8 Expand prefixes and click **+** to add the ip-prefix to the prefix-list. Type **70.70.70.0/30** and click **Continue**.



The screenshot shows the 'Create L3VPN > Routing Policy Destination Prefix' configuration page. The title is 'Routing Policy Destination Prefix {DEST_PREFIX_SET_70}'. The 'name' field is 'DEST_PREFIX_SET_70'. The 'mode' dropdown is set to 'ipv4'. Under the 'prefixes' section, there is a 'prefix-list' field with a '+' icon. Below it, the 'ip-prefix' field is set to '70.70.70.0/30'.


Step 9 Create the other two routing policy destination prefixes (**DEST_PREFIX_SET_71** and **DEST_PREFIX_SET_72**) by following the same steps.



Now we are ready to create the first route policy - L3VPN_NM-SRTE-RP-PE-A-7. The other route policies can be created using the same procedure.



Step 10 Go to **Services & Traffic Engineering > Provisioning (NSO) > L3vpn > Routing Policy**.



- Step 11** Click  to create a new route policy and type a unique name for the top-level policy definition: **L3VPN_NM-SRTE-RP-PE-A-7**. Click **Continue**. The statements section appears.
- Note** The Route Policy statement defines the condition and action taken by the system.
- Step 12** Expand statements and click  to add the name of the policy statement (such as **stmt1**) and click **Continue**. The statement {stmt1} panel appears showing **conditions** and **actions** sections.
- Step 13** Expand conditions and then expand match-dest-prefix-set. In the prefix-set list, select or type the following: **DEST_PREFIX_SET_70**. This is what references a defined prefix set.
- Note** Once selected, the **Enable match-dest-prefix-set** toggle, which will match a referenced prefix-set according to the logic defined in the match-set-options list, switches on.
- Step 14** Expand actions and then expand bgp-actions.
- Step 15** For bgp-actions, slide the Enable bgp-actions toggle to the on position. By toggling bgp-actions on, it defines the top-level container for BGP-specific actions.
- Step 16** Now expand set-ext-community. Slide the Enable-set-ext-community toggle to the on position. By toggling set-ext-community on, it sets the extended community attributes.
- Step 17** For Method and reference, select the Ext-community-set-ref list and select **COLOR_70**. The Ext-community-set-ref references a defined extended community set by name.

statement{stmt1}  


name 



conditions  


match-source-prefix-set  



match-dest-prefix-set  

Enable match-dest-prefix-set



prefix-set 

actions  

policy-result 

bgp-actions  


Enable bgp-actions

set-ext-community  

Enable set-ext-community

Method

reference

ext-community-set-ref 

- Step 18** Click **X** in the top-right corner to close the statement {stmt1} panel and click **Commit changes**.
- Step 19** Create the other route policies (**L3VPN_NM-SRTE-RP-PE-B-7** and **L3VPN_NM-SRTE-RP-PE-C-7**) in the same manner prior to creating the L3VPN service.

After creating the L3VPN route policies, create the VPN profile for each route policy and then create and provision the L3VPN service. The VPN profile will be referenced from the L3VPN service. This will bind the route policy to the L3VPN service.

Step 3 Create and provision the L3VPN service

In this step, we will create the L3VPN service with three endpoints: PE-A, PE-B, and PE-C. Each endpoint will be associated with a vpn-instance-profile, which in turn points to a VPN profile that contains the route policy with the same color as specified in the ODN template. In this way, traffic that matches the specified prefixes and color will be treated according to the SLA specifications.

First, we will create the VPN profiles. The newly created VPN profiles will have the same names as the L3VPN routing policy names.

-
- Step 1** Go to **Services & Traffic Engineering > Provisioning (NSO) > L3VPN > VPN Profiles**.
- Step 2** Click **+** to create a VPN profile to be referenced in the VPN service.
- Step 3** Select the Id list and select **L3VPN_NM-SRTE-RP-PE-A-7**.
Now create and provision the L3VPN service.
- Step 4** Go to **Services & Traffic Engineering > Provisioning (NSO) > L3vpn > L3vpn-Service**.
- Step 5** Click **+** to create a new service and type a new vpn-id: **L3VPN_NM-SRTE-ODN-70**. Click **Continue**.
A VPN identifier uniquely identifies a VPN and has a local meaning (for example, within a service provider network).
- Step 6** Create vpn-instance-profiles, which is a container that defines the route distinguisher (RD), route targets, and the export/import route policy. We will create vpn-instance-profiles for each endpoint, as follows:
- L3VPN_NM_SR_ODN-IE-PE-A-7 with route distinguisher 0:70:70
 - L3VPN_NM_SR_ODN-IE-PE-B-7 with route distinguisher 0:70:71
 - L3VPN_NM_SR_ODN-IE-PE-C-7 with route distinguisher 0:70:72
- a. Expand vpn-instance-profiles and click **+** to create a new vpn-instance-profile profile-id: **L3VPN_NM_SR_ODN-I-PE-A-7**. Click **Continue**.
 - b. Enter the route distinguisher (rd) that will differentiate the IP prefixes and make them unique. For this scenario, we are using **0:70:70**.
 - c. For address-family, click **+** and select **ipv4** from the list. Click **Continue**.
 - d. Define the required VPN targets, including id, route-targets, and route-target-type (import/export/both).
 - e. Under vpn-policies, in the export-policy list, choose the relevant VPN profile (which contains the route policy: **L3VPN_NM-SRTE-RP-PE-A-7**). This forms the association between the VPN and the ODN template that defines the SLA.

- f. Click **X** in the top-right corner when you are done.
- g. Similarly, create the other vpn-instance-profiles.

Step 7 Define each VPN endpoint individually: PE-A, PE-B, and PE-C.

- a) Expand vpn-nodes and click **+** to select the relevant device from the list: **PE-A**. Click **Continue**.
- b) Enter the local-as number for network identification: **200**.
- c) Expand active-vpn-instance-profiles and click **+** to select the profile-id you created in the previously: **L3VPN_NM-SRTE-RP-PE-A-7**. Click **Continue**.
- d) Define the network access parameters for communication from the PE towards the CE:
 - Under vpn-network-accesses, click **+** to create a new set of VPN access parameters and provide a unique ID. Click **Continue**.
 - In the Interface-id field, type **Loopback70**. This is the identifier for the physical or logical interface. The identification of the sub-interface is provided at the connection level and/or the IP connection level.
 - Expand ip-connection > ipv4 and enter a local-address (**70.70.70.1**) and then prefix-length (**30**).
 - For routing-protocols, click **+** to create a unique id, set the type to bgp-routing, and then expand bgp to set the peer-as number (**70**), and the address-family (**ipv4**). In addition, set the bgp neighbor (**70.70.70.2**) and the multihop number (for example, **11**) that indicates the number of hops allowed between the bgp neighbor and the PE device.

type* ⓘ

bgp-routing

bgp ⓘ

peer-as * ⓘ


70

address-family ⓘ

ipv4

neighbor ⓘ

Total 1 ⚙️

+ 

neighbor

70.70.70.2

multihop ⓘ

11

- Click **X** in the top-right corner until you are back on the Create L3VPN screen.
- Similarly, create the other VPN nodes: **PE-B** and **PE-C**.

Step 8 Commit your changes or click **Dry Run** to check what will be configured on the devices before you commit.

Step 9 Check that the new L3VPN service appears in the table and its provisioning state is **Success**.

Step 4 Enable Service Health monitoring


After creating and provisioning the required L3VPN services, you can begin monitoring their health.

Before you begin

- (*Optional*) Ensure that Crosswork Service Health is installed. For details, see the "Install Crosswork Applications" chapter in the [Cisco Crosswork Network Controller Installation Guide](#). For more information on Service Health, see the Cisco Crosswork Network Controller Service Health Guide.
- The Service Health related steps assume you have excess capacity available. Requirements (such as available resources, storage capacity, etc.) may be beyond the scope explained in this guide. See the Crosswork Network Controller Service Health Guide for further details.

Select the newly created, unmonitored, service which will have a gray health indicator:

Step 1 From the main menu, choose **Services & Traffic Engineering > VPN Services**. The map opens on the left side of the page and the table opens on the right side.

Step 2 In the Actions column, click  for the service you want to start monitoring the health.

Step 3 Click **Start Monitoring**.

VPN Services Refined By: All Endpoints ▾

Provisioning: 5 Success, 0 Failed, 0 In-Progress

Health (Monitoring: 3 Services): 2 Good, 1 Degraded, 0 Down

Total 5

[Create ▾](#) ☰

Health	Service ...	Type	Provisioni...	Las... [ⓘ]	Actions
	L2VPN_N...	L2vpn-Ser...	Success	26-Jul-...	...
	L3NM-PR...	L3vpn-Ser...	Success	26-Jul-...	...
	L3NM-PR...	L3vpn-Ser...	Success	26-Jul-...	
	L3NM-PR...	L3vpn-Ser...	Success	26-J	View Details
	L3NM-PR...	L3vpn-Ser...	Success	26-J	Edit / Delete
					Start Monitoring

Note The Health column color coding indicates the health of the service:

- Blue = Initiated
- Green = Good
- Orange = Degraded
- Red = Down
- Gray = Not Monitoring

Step 4 In the Monitor Service dialog box, select the Monitoring Level. For help with selecting the appropriate monitoring level for your needs, see the Cisco Crosswork Network Controller 6.0 Service Health Guide.

Monitor Service

Name L3NM-PROBES-45-2-3-endpoint

Monitoring Level ?

Silver_L3VPN_ConfigProfile custom

Gold_L3VPN_ConfigProfile custom

Basic Monitoring

Advanced Monitoring


Thresholds to use for Silver L3VPN services

Cpu Threshold Max 80.5 %




Memfree Threshold Min 1000000000 bytes




Cancel **Start Monitoring**


Step 5 Click **Start Monitoring**.


Note Once you have started monitoring the health of the service, in the Actions column, if you click  to view additional Service Health options, you will see: Stop Monitoring, Pause Monitoring, Edit Monitoring Settings, and Assurance Graph.













VPN Services Refined By: All Endpoints

Provisioning: 5  Success, 0  Failed, 0  In-Progress

Health (Monitoring: 3 Services): 2  Good, 1  Degraded, 0  Down


Total 5 

Create 

Health	Service ...	Type	Provisioni...	Las... 	Actions
	L2VPN_N...	L2vpn-Ser...	 Success	26-Jul-...	...
	L3NM-PR...	L3vpn-Ser...	 Success	26-Jul-...	
	L3NM-PR...	L3vpn-Ser...	 Success		View Details
	L3NM-PR...	L3vpn-Ser...	 Success		Edit / Delete
	L3NM-PR...	L3vpn-Ser...	 Success		Stop Monitoring
					Pause Monitoring
					Edit Monitoring Settings

Step 6 Repeat these steps for each service that you wish to start health monitoring.

Step 5 Visualize the New VPN Service on the Map to See the Traffic Path

Step 1 In the L3VPN Service table, click on the service name or click  in the Actions column and choose **View Details** from the menu.

The map opens and the service details are shown to the right of the map.

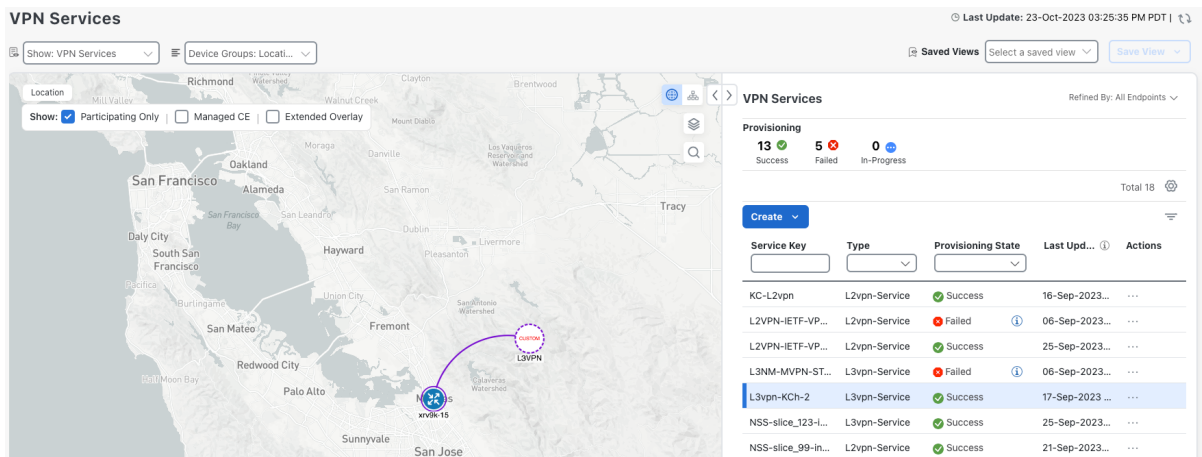
or

Go to **Services & Traffic Engineering > VPN Services**.

The map opens and a table of VPN services is displayed to the right of the map.

Click on the VPN in the Services table. If there are many services in the table, you can filter by name, type, or provisioning state to help locate the VPN.

In the map, you will see the VPN as an overlay on the topology. It shows a representation of the three endpoints and a dashed line that indicates that it is a virtual path.



VPN Services Last Update: 23-Oct-2023 03:25:35 PM PDT | ↕

Show: VPN Services Device Groups: Locati...

Location: Saved Views Select a saved view Save View

Show: Participating Only Managed CE Extended Overlay

Service Key	Type	Provisioning State	Last Upd...	Actions
KC-L2vpn	L2vpn-Service	Success	16-Sep-2023...	...
L2VPN-IETF-VP...	L2vpn-Service	Failed	06-Sep-2023...	...
L2VPN-IETF-VP...	L2vpn-Service	Success	25-Sep-2023...	...
L3NM-MVPN-ST...	L3vpn-Service	Failed	06-Sep-2023...	...
L3vpn-KCh-2	L3vpn-Service	Success	17-Sep-2023...	...
NSS-slice_123-l...	L3vpn-Service	Success	25-Sep-2023...	...
NSS-slice_99-in...	L2vpn-Service	Success	21-Sep-2023...	...

Select the **Show Participating Only** check box if you do not want to see the devices that are not involved in the selected VPN.

Step 2 In the Actions column, click



to drill down to a detailed view of the VPN service, including the device configurations and the computed transport paths.

Step 3 To see the computed paths for this VPN, click on the Transport tab in the Service Details pane. All the dynamically created SR-TE policies are listed in the Transport tab. Select one or more SR-TE policies to see the path from endpoint to endpoint on the map.

Step 5 Visualize the New VPN Service on the Map to See the Traffic Path

In this example, we are looking at the disjoint paths computed from PE-A to PE-B and from PE-A to PE-C.

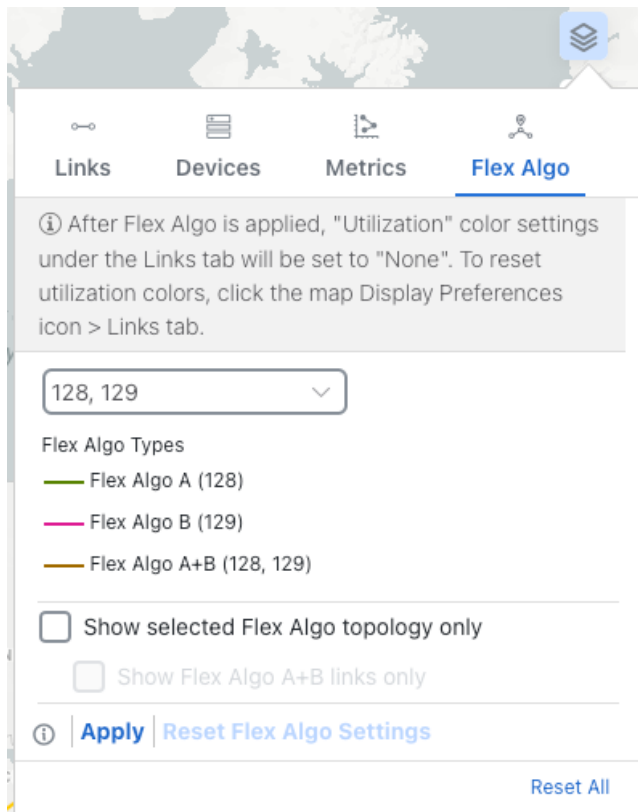


Step 4 To see the physical path between the endpoints, select the **Show IGP Path** check box in the top-left corner of the map. Hover with your mouse over a selected policy in the table to highlight the path in the map and show prefix SID and



routing information.

Step 5 To filter the topology to a specific Flex-Algo constraint and visualize nodes and links you have configured manually in your network, click the button at the top right of the map and do the following:



- Click the **Flex Algo** tab.
- From the drop-down list, choose up to 2 Flex-Algo IDs.
- View the Flex-Algo Types and confirm that the selection is correct. Also, note the color assignments for each Flex-Algo ID.
- (Optional) Check the **Show selected Flex Algo topology only** check box to isolate the Flex-Algo IDs on the topology map. When this option is enabled, SR policy selection is disabled.
- Check the **Show Flex Algo A+B links only** to show only those links and nodes that participate in both Flex-Algos.
If a selected Flexible Algorithm is defined with criteria but there are no links and node combinations that match it (for example, a defined affinity to include all nodes or links with the color blue), then the topology map will be blank. If a selected Flexible Algorithm is not configured on a node or link, then the default blue link or node color appears.
- Click **Apply**. You must click **Apply** for any additional changes to your Flex-Algo selections to see the update on the topology map.
- (Optional) Click **Save View** to save the topology view and Flexible Algorithm selections.

Step 6 Observe automatic network optimization

Observe automatic network optimization

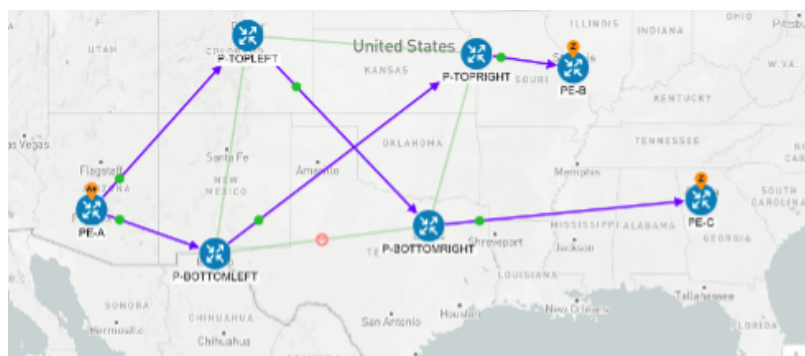
The SR-PCE constantly monitors the network and automatically optimizes the traffic path based on the defined SLA. For illustration purposes, let's look at what happens when one of the links goes down, in this case, the link between

Step 7 Inspect a degraded service using Service Health to determine active symptoms

P-BOTTOMLEFT and P-BOTTOMRIGHT. This means that the previous path from PE-A to PE-C is no longer viable. Therefore, the SR-PCE computes an alternative path, both from PE-A to PE-C and from PE-A to PE-B, to compensate for the link that is down and to maintain the disjoint paths.

Recomputed paths:

Source and Destination	Old path	New path
PE-A > PE-C	PE-A > P-BOTTOMLEFT > P-BOTTOMRIGHT > PE-C	PE-A > P-TOPLEFT > P-BOTTOMRIGHT > PE-C
PE-A > PE-B	PE-A > P-TOPLEFT > P-TOPRIGHT > PE-B	PE-A > P-BOTTOMLEFT > P-TOPRIGHT > PE-B



Step 7 Inspect a degraded service using Service Health to determine active symptoms

By analysing the root cause of reported active symptoms and impacted services, you can determine what issues must be addressed first to maintain a healthy setup and what requires further inspection and troubleshooting.

To view the active symptoms and root causes for a service degradation:

Before you begin

Ensure that service health monitoring is enabled for the service you want to inspect.



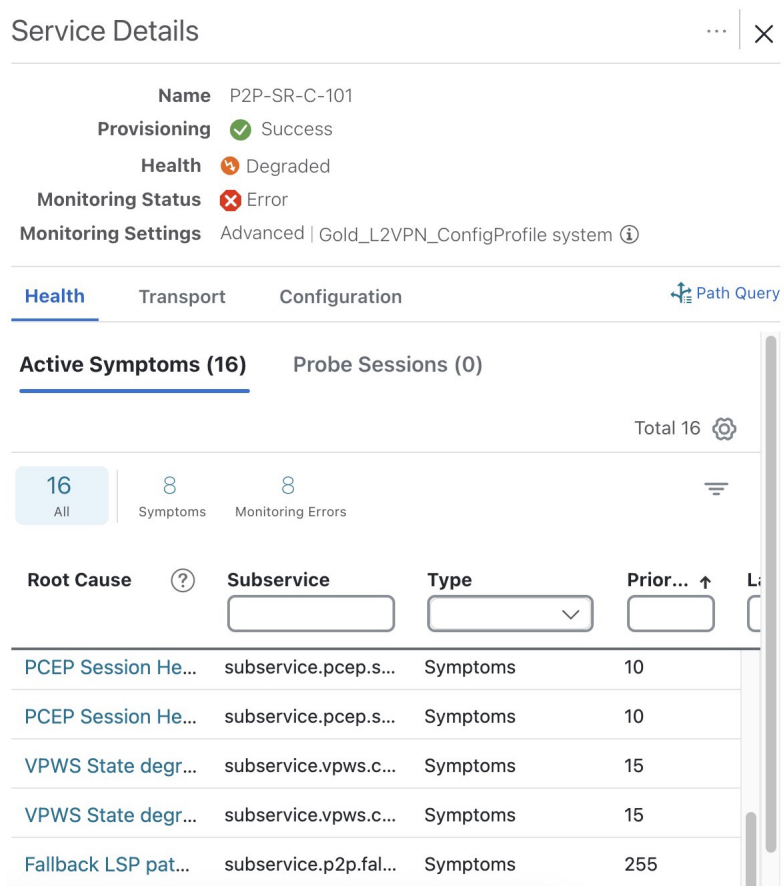
Note L3VPN service monitoring is supported on XR devices and not on XD devices. For an L3VPN service being monitored, if a provider and devices are deleted, and then added again, the monitoring status remain in the degraded state with a METRIC_SCHEDULER error. To recover from this error, stop and restart the service monitoring.

Step 1 From the main menu, choose **Services & Traffic Engineering > VPN Services**. The map opens on the left side of the page and the table opens on the right side.

Step 2 In the Actions column, click  and click **View Details**. The Service Details panel appears on the right side.

Step 3 Select the Health tab and click the **Active Symptoms** tab. The Active Symptoms table displays **Active Symptoms** and **Monitoring Errors** by default. To filter the table to show only the Active Symptoms, either click the **Symptoms** tab in the mini dashboard above the table or select **Symptoms** from the filter box under the **Type**. The table now shows a filtered list containing only the Active Symptoms.

Review the Active Symptoms for the degraded service (including the Root Cause, Subservice, Type, Priority, and Last Updated details).



The screenshot shows the Service Details panel for service P2P-SR-C-101. The Health tab is selected, and the Active Symptoms sub-tab is active. The table displays 16 symptoms, with 8 symptoms and 8 monitoring errors. The table columns are Root Cause, Subservice, Type, Priority, and Last Updated. The following table represents the data shown in the screenshot:

Root Cause	Subservice	Type	Prior...	Last Updated
PCEP Session He...	subservice.pcep.s...	Symptoms	10	
PCEP Session He...	subservice.pcep.s...	Symptoms	10	
VPWS State degr...	subservice.vpws.c...	Symptoms	15	
VPWS State degr...	subservice.vpws.c...	Symptoms	15	
Fallback LSP pat...	subservice.p2p.fal...	Symptoms	255	

Step 4 Click on a Root Cause and view both the **Symptom Details** and the **Failed Subexpressions & Metrics** information. You can expand or collapse all of the symptoms listed in the tree, as required. In addition, use the **Show Only Failed** toggle to focus only on the failed expression values.

Step 7 Inspect a degraded service using Service Health to determine active symptoms

Service Details ... ✕

Name P2P-SR-C-101

Provisioning ✔ Success

Health ⚠ Degraded

Monitoring Status ✖ Error

Monitoring Settings Advanced | Gold_L2VPN_ConfigProfile system ⓘ

Health | Transport | Configuration ↕ Path Query

Symptom Details ^

Name PCEP Session Health degraded. Device: CL2-PE-C, PCC-
Peer: 192.168.15.42

Sub Service subservice.pcep.session.health system

Last Updated 28-Jul-2023 11:29:20 PM IST

Failed Subexpressions & Metrics ^

Show Only Failed Expand All | Collapse All

Name	Expre
explabel	pcc_p
⚠ pcc_peer_state == 'up'	false

Step 5 Click the **Transport** and **Configuration** tabs and review the details provided.

Step 6 Click **✕** in the top-right corner to return to the VPN Services list.

Step 7 In the Actions column, click ⋮ for the required degraded service and click **Assurance Graph**. The topology map of services and subservices appear with the Service Details panel showing Service Key, Status, Monitoring Status, Monitoring Settings, Sub Services, and Active Symptoms details.

This may take up to 5-10 minutes to update after a service has been enabled for monitoring.

Metrics such as Jitter-RT (Jitter Round Trip), Latency-RT (Latency Round Trip), PktLoss-DS (Packet Loss from Destination to Source), and PktLoss-SD (Packet Loss from Source to Destination) also appear (information collected using Y.1731 probes). Additionally, a table of Active Symptoms listing Root Cause, Subservice, Priority, and Last Updated details is populated.

At the top-right of the map, select the stack icon to select the appearance option for the Subservices: **State + Icon + Label** or **State + Icon**.

Step 8 By default, the Assurance Graph displays a concise view with only the service and the top level subservices (aggregator nodes). Click the **+** icon in the nodes to expand the graph and to view the dependent details. To expand all the nodes at once, click the **Subservices: Expand All** check box at the top.

Step 9 Select a degraded subservice in the Assurance Graph. The Subservice Details panel appears with subservice metrics, as well as subservice specific Active Symptoms and Impacted Services details.


- **Active Symptoms:** Provides symptom details for nodes actively being monitored.
- **Impacted Services:** Provides information for services that are impacted by issues based on historical monitoring of health status.

Step 7 Inspect a degraded service using Service Health to determine active symptoms

Note At the top left of the map, check the **Down & Degraded only** or **Soft Dependencies** check boxes to further isolate the subservices. Soft Dependencies implies that a child subservice's health has a weak correlation to its parent's health. As a result, the degraded health of the child will not result in the parent's health degradation.

Note In some cases, the Summary node feature is available and summarizes the aggregated health status of child subservices and reports a consolidated health status to a service node. The Summary node feature is available in both L2VPN multipoint Basic and Advanced monitoring models.

- Basic monitoring subservices:
 - Device—Summarizes the health status of all underlying Devices participating in the given L2VPN service.
 - Bridge Domain—Summarizes the L2VPN service's Bridge Domain health status across all participating devices.
- Advanced monitoring subservices (in addition to what is also available with Basic monitoring):
 - EVPN—Summarizes the health status of all underlying subservices—BGP Neighbor Health and MacLearning Health across all participating PE endpoints and provides a consolidated overall EVPN health summary status.
 - Transport—Summarizes the health status of all underlying subservices—SR-ODN (dynamic), SR Policy (statically configured), and RSVP TE Tunnel, across all participating PE endpoints and provides a consolidated overall Transport health summary status.
 - SR-PCEP—Summarizes the health status of all the underlying subservices that are monitoring the PCEP sessions. Each underlying subservice monitors the PCEP session health on a particular device participating in the given VPN service.



Step 10 Inspect the Active Symptoms and Impacted Services information, and the root causes associated with the degraded service to determine what issues may need to be addressed to maintain a healthy setup.


To further troubleshoot a service health issue (such as a device that is degraded due to not properly fetching data), continue with the following steps to examine if the issue is associated with a collection job.

Step 11 From the main menu, choose **Administration > Collection Jobs**.

The Collection Jobs page appears.

Step 12 Click the **Parameterized Jobs** tab.

Step 13 Review the Parameterized Jobs list to identify the devices that may have service health degradation issues. By reviewing Parameterized Jobs, you can identify and focus on gNMI, SNMP, and CLI-based jobs by their Context ID (protocol) for further troubleshooting purposes.

Step 14 In the Job Details panel, select the collection job you want to export and click  to download the status of collection jobs for further examination. The information provided is collected in a .csv file when the export is initiated.

Note When exporting the collection status, you must fill in the information each time an export is executed. In addition, make sure to review the **Steps to Decrypt Exported File** content available on the Export Collection Status dialog box to ensure you can access and view the exported information.

Step 15 Click **Export**.

Step 16 To check the status of the exported collection job data, click **View Export Status** at the top right of the Job Details panel. The Export Status Jobs panel appears providing the status of the export request.

Step 17 Review the exported .csv file for collection job details and the possible cause of the degraded device.

Summary and Conclusion

As we observed in this example, operators can use Cisco Crosswork Network Controller to orchestrate L3VPNs with SLAs and to maintain these SLAs using SR-TE policies that continuously track network conditions and automatically react to optimize the network. This automation increases efficiency and reduces human error that is generally unavoidable with manual tasks. Enabling Service Health to monitor provisioned services allows for more detailed symptoms, metrics, and analysis of each service.

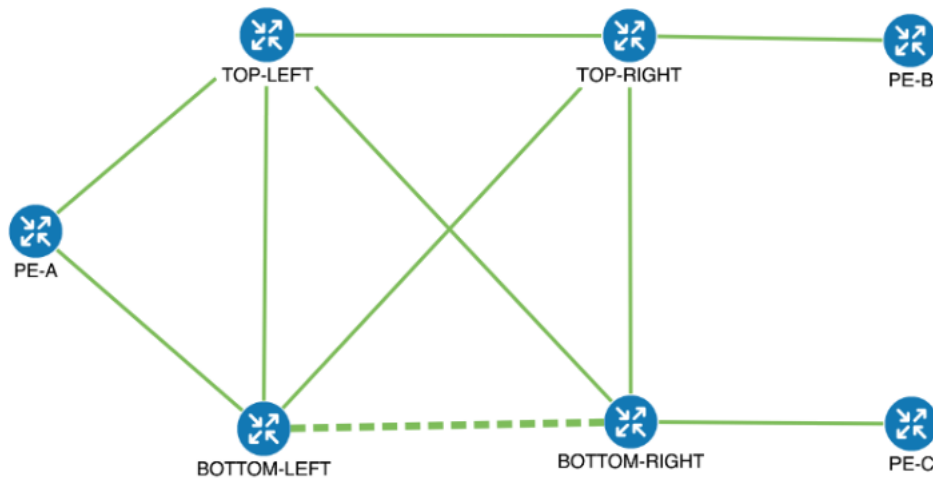
Scenario: Implement and Maintain SLA for an L3VPN Service for SRv6 (using ODN)

This scenario walks you through the procedure for provisioning an L3VPN service that requires a specific SLA objective. In this example, the lowest latency path is the SLA objective. The customer requires a low latency path for high priority traffic. The customer wants to use disjoint paths, i.e., two unique paths that steer traffic from the same source and to the same destination, avoiding common links so that there is no single point of failure. The customer also wants to enable SRv6, which utilizes the IPv6 protocol to handle packets with more efficiency, increase security and performance, allowing for a significantly larger number of possible addresses.

This is achieved using Segment Routing (SR) On-Demand Next Hop (ODN). ODN allows a service head-end router to automatically instantiate an SR-TE policy to a BGP next-hop when required (on-demand). The headend is configured with an ODN template with a specific color that defines the SLA upon which the traffic path will be optimized when a prefix with the specified color is received. Prefixes are defined in a route policy that is associated with the L3VPN.

Cisco Crosswork Network Controller continues to monitor the network and will automatically optimize the network based on the defined SLA, in a closed loop.

The following topology provides the base for this scenario:



In this scenario, we will:

- Create a segment routing ODN template with a specific color on the endpoints to ensure that traffic is transported within an LSP (underlay) and that a best-path tunnel is created dynamically when a prefix with the specified color is received. Enable SRv6 (IPv6) for service and link details. The ODN template defines the SLA on which you want to optimize the path. In this case, we will optimize on latency.
- Specify that the computed paths be disjoint: they will not share the same link.
- Create a route policy on each endpoint to be used to bind the L3VPN to the ODN template. This route policy adds a color attribute to the customer prefixes and advertises via BGP to other endpoints. This color attribute is used to indicate the SLA required for these prefixes.
- Create an L3VPN service with 3 endpoints: PE-A, PE-B, and PE-C. This is the overlay configuration.

- Visualize how this overlay/underlay configuration optimizes the traffic path and automatically maintains the SLA.

Assumptions and Prerequisites

- To use ODN with SRv6, BGP peering for the prefixes must be configured between the endpoints/PEs. Usually for L3VPN, this is the VPNv4 and VPNv6 address family peering, and this BGP peering is required to be over IPv6.

Procedure to Implement and Maintain SLA for an L3VPN Service for SRv6 Using ODN is detailed in this section.

Step 1 Create an ODN template to map color to an SLA objective and constraints

We will create the following ODN templates:

- Headend PE-A, color 72, latency, disjoint path (link), group ID 16 - L3VPN_NM-SRTE-ODN_72-a
- Headend PE-A, color 71, latency, disjoint path (link), group ID 16 - L3VPN_NM-SRTE-ODN_71-a
- Headend PE-B and PE-C, color 70, latency - L3VPN_NM-SRTE-ODN_70
 - With multiple headends in the SRv6 enabled ODN template, the same locator name should be configured on the headend routers. Otherwise, different ODN templates should be created for each headend.
- Headend PE-B, color 72, latency - L3VPN_NM-SRTE-ODN_72-b
- Headend PE-C, color 71, latency - L3VPN_NM-SRTE-ODN_71-c

For example purposes, we will show how to create the first ODN template - L3VPN_NM-SRTE-ODN_72-a. The other ODN templates can be created using the same procedure.

Before you begin

In this step, we will create an ODN template on each endpoint. The ODN template specifies the color and the intent; in this case, latency and disjointness. This ODN template will be used to dynamically create tunnels (on-demand) when prefixes with matching colors are received via BGP. Traffic to these prefixes will be automatically steered into the newly created tunnels, thereby meeting the SLA objective and constraints intended for these prefixes and signaled using colors in the BGP routes.

Disjointness constraints work by associating a disjoint group ID with the ODN template, and all tunnels with the same disjoint group ID will be disjoint, i.e., they will use different links, nodes and shared risk link groups depending on how the disjoint groups are configured.

Step 1 Go to **Services & Traffic Engineering > Provisioning (NSO) > SR-TE > ODN-Template**.

Step 2 Click  to create a new template and give it a unique name.

In this case, the name is **L3VPN_NM-SRTE-ODN_72-a**. Click **Continue**.

Step 1 Create an ODN template to map color to an SLA objective and constraints

SR-TE > ODN-Template



Name

Step 3 Choose the headend device, **PE-A**, and specify the color **72**.

head-end ⓘ ?



name

PE-A

color * ⓘ

 ⓘ

Step 4 Under srv6, select the **Enable srv6** toggle.

Step 5 Under locator, enter the required SRv6 **locator-name**.

The locator name should match what is configured on the router.

srv6 ⓘ

Enable srv6



locator ⓘ

Enable locator



locator-name * ⓘ

ALG0r5

behavior ⓘ

ub6-insert-reduced

binding-sid-type ⓘ

srv6-dynamic

Step 6 Under dynamic, select **latency** as the metric type. This is the SLA objective on which we are optimizing.

Step 7 Select the **pce** check box to specify that the path should be computed by the SR-PCE, not by the Path Computation Client (PCC).

Step 8 Define the required constraints. In this case, we want the computed paths to be disjoint in that they must not share a link.

Under disjoint-path, choose **link** as the type, and specify a numeric group ID, in this case, 16.

dynamic ⓘ

Enable dynamic

pce ⓘ

flex-alg ⓘ

metric-type ⓘ

latency

metric-margin ⓘ

affinity ⓘ

segments ⓘ

disjoint-path ⓘ

Enable disjoint-path

type*


link

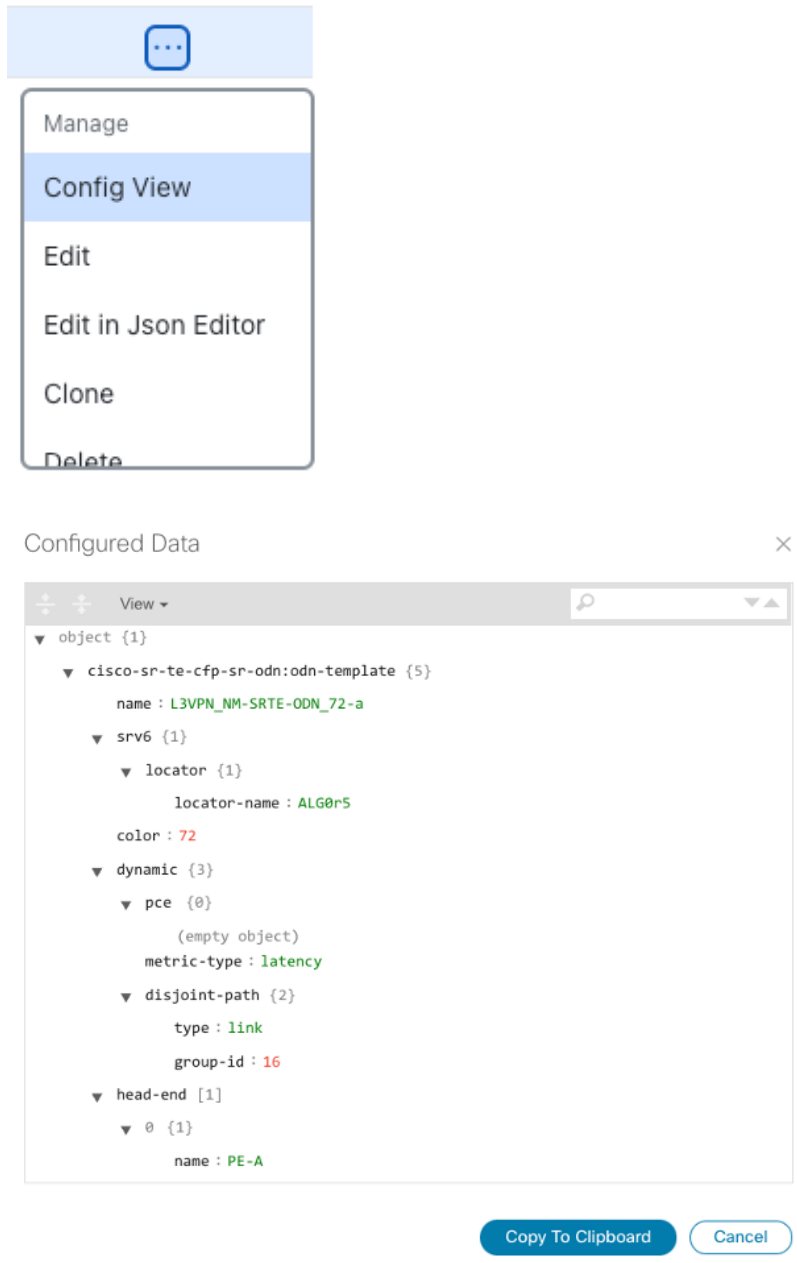
group-id * ⓘ

16

Commit your changes or click **Dry Run** to check what will be configured on the devices before you commit.

Step 9

Check that the new ODN template appears in the table and its provisioning state is **Success**. Click  in the Actions column and choose **Config View** to see the Yang model-based service intent data that details the ODN template you created.



The image shows a configuration menu and a 'Configured Data' window. The menu includes options: Manage, Config View (highlighted), Edit, Edit in Json Editor, Clone, and Delete. The 'Configured Data' window displays a hierarchical JSON structure:

```

object {1}
  cisco-sr-te-cfp-sr-odn:odn-template {5}
    name : L3VPN_NM-SRTE-ODN_72-a
    srv6 {1}
      locator {1}
        locator-name : ALG0r5
        color : 72
      dynamic {3}
        pce {0}
          (empty object)
          metric-type : latency
        disjoint-path {2}
          type : link
          group-id : 16
        head-end [1]
          0 {1}
            name : PE-A
  
```

At the bottom of the 'Configured Data' window, there are two buttons: 'Copy To Clipboard' and 'Cancel'.

Step 10 Create the other ODN templates listed above in the same manner.

Step 2 Create an L3VPN Route Policy

In this step, we will create a route policy for each endpoint, and we will specify the same color as defined in the ODN template for that endpoint. The route policy defines the prefixes to which the SLA applies. When traffic from the specified network with a matching color is received, paths are computed based on the SLA defined in the ODN template. We will create the following route policies:

- Color 70, IPv6 prefix 70:70:70::0/64 - L3VPN_NM-SRTE-RP-PE-A-7
- Color 71, IPv6 prefix 70:70:71::0/64 - L3VPN_NM-SRTE-RP-PE-B-7
- Color 72, IPv6 prefix 70:70:72::0/64 - L3VPN_NM-SRTE-RP-PE-C-7

For example purposes, we will show how to create the first route policy - L3VPN_NM-SRTE-RP-PE-A-7. The other route policies can be created using the same procedure.

First, we will create the routing policy tag and routing policy destination prefix. The routing policy prefixes should match with the subnet prefix configured on the PE devices in the service.

Step 1 Go to **Services & Traffic Engineering > Provisioning (NSO) > L3VPN > Routing Policy Tag**.

Step 2 Click **+** to create a new routing policy tag and type the name of the tag set: **COLOR_70**. Click **Continue**.

This is used as a label to reference the set in actions and conditions.

Step 3 Under tag-value, click **+** and type the Tag-value: **70**.

The tag value may be a number between **1 – 4294967295** and should match to a color value.

Step 4 Click **Continue**. The new routing policy tag name with the new tag value is visible. Click **Commit changes**.

Create the other two routing policy tags (**COLOR_71** and **COLOR_72**) and tag values (**71** and **72**) by following the same steps above.

Now create the routing policy destination prefixes.

Step 5 Go to **Services & Traffic Engineering > Provisioning (NSO) > L3VPN > Routing Policy Destination Prefix**.

Step 6 Click **+** to create a new routing policy destination prefix and type the name: **DEST_PREFIX_SET_70**.

The name of the prefix set will reference the set in match conditions.

Step 7 For Mode, select **ipv6**.

Step 8 Expand prefixes and click **+** to add the ip-prefix to the prefix-list.

Step 9 For Ip-prefix, type **70:70:70::0/64** and click **Continue**.

Create the other two routing policy destination prefixes (**DEST_PREFIX_SET_71** and **DEST_PREFIX_SET_72**) by following the same steps. Click **Commit changes**.

Now we are ready to create the first route policy L3VPN_NM-SRTE-RP-PE-A-7. The other route policies can be created using the same procedure.

Routing Policy Destination Prefix {DEST_PREFIX_SET_70}

name * ⓘ



DEST_PREFIX_SET_70

mode ⓘ

ipv4

prefixes ⓘ


prefix-list ⓘ ⓘ


ip-prefix

70:70:70::/64

Step 10 Go to **Services & Traffic Engineering > Provisioning (NSO) > L3VPN > Routing Policy**.

Step 11 Click  to create a new route policy and type a unique name for the top-level policy definition: **L3VPN_NM-SRTE-RP-PE-A-7**. Click **Continue**. The statements section appears.

Note The Route Policy statement defines the condition and action taken by the system.

Step 12 Expand statements and click  to add the name of the policy statement (such as **stmt1**) and click **Continue**. The statement {stmt1} panel appears showing **conditions** and **actions** sections.

Step 13 Expand conditions and then expand match-dest-prefix-set before selecting the prefix-set list and select **DEST_PREFIX_SET_70**. This is what references a defined prefix set.

Note Once selected, the **Enable match-dest-prefix-set** toggle, which will match a referenced prefix-set according to the logic defined in the match-set-options list, switches on.

Step 14 Expand actions and then expand bgp-actions.

Step 15 For bgp-actions, slide the Enable bgp-actions toggle to the on position. By toggling bgp-actions on, it defines the top-level container for BGP-specific actions.

Step 16 Now expand set-ext-community. Slide the enable-set-ext-community toggle to the on position. By toggling set-ext-community on, it sets the extended community attributes.

Step 17 For Method and reference, select the ext-community-set-ref list and select **COLOR_70**. The Ext-community-set-ref references a defined extended community set by name.

Note Creating routing-policy tag-set is mandatory and needs to be mapped here.

Step 18 Click **X** in the top-right corner to close the statement{stmt1} panel and click **Commit**

statement{stmt1}
↻ ×

name * ⓘ

stmt1

conditions ⓘ ^

match-source-prefix-set ⓘ v

match-dest-prefix-set ⓘ ^

Enable match-dest-prefix-set

prefix-set ⓘ

DEST_PREFIX_SET_70

actions ⓘ ^

policy-result ⓘ

v

bgp-actions ⓘ ^

Enable bgp-actions

set-ext-community ⓘ ^

Enable set-ext-community

Method

reference

ext-community-set-ref ⓘ

COLOR_70 v

changes.







Step 19 Create the other route policies (L3VPN_NM-SRTE-RP-PE-B-7 and L3VPN_NM-SRTE-RP-PE-C-7) in the same manner.






After creating the L3VPN route policies, create the VPN profile for each route policy and then create and provision the L3VPN service. The VPN profile will be referenced from the L3VPN service. This will bind the route policy to the L3VPN service.

Step 3 Create and provision the L3VPN service

In this step, we will create the L3VPN service with three endpoints: PE-A, PE-B, and PE-C. Each endpoint will be associated with a vpn-instance-profile, which in turn points to a VPN profile that contains the route policy with the same color as specified in the ODN template. In this way, traffic that matches the specified prefixes and color will be treated according to the SLA specifications.

First, we will create the VPN profiles. The newly created VPN profiles will have the same names as the L3VPN routing policy names.


-
- Step 1** Go to **Services & Traffic Engineering > Provisioning (NSO) > L3VPN > VPN Profiles**.
- Step 2** Click  to create a valid VPN profile to be referenced in the VPN service.
- Step 3** Select the Id list and select **L3VPN_NM-SRTE-RP-PE-A-7**.
Now create and provision the L3VPN service.
- Step 4** Go to **Services & Traffic Engineering > Provisioning (NSO) > L3vpn > L3vpn-Service..**
- Step 5** Click  to create a new service and type a new vpn-id: **L3VPN_NM-SRTE-ODN-70**.
A VPN identifier uniquely identifies a VPN and has a local meaning (for example, within a service provider network).
- Step 6** Click **Continue**.
- Step 7** Create vpn-instance-profiles, which is a container that defines the route distinguisher (RD), route targets, and the export/import route policy. We will create vpn-instance-profiles for each endpoint, as follows:
- L3VPN_NM_SR_ODN-IE-PE-A-7 with route distinguisher 0:70:70
 - L3VPN_NM_SR_ODN-IE-PE-B-7 with route distinguisher 0:70:71
 - L3VPN_NM_SR_ODN-IE-PE-C-7 with route distinguisher 0:70:72
- a. Expand vpn-instance-profiles and click  to create a new vpn-instance-profile profile-id: **L3VPN_NM_SR_ODN-I-PE-A-7**. Click **Continue**.
 - b. Enter the route distinguisher (Rd) that will differentiate the IP prefixes and make them unique: **0:70:70**.
 - c. For address-family, click  and select **ipv6** from the list. Click **Continue**.
 - d. Define the required VPN targets, including route targets and route target types (import/export/both).
 - e. Under vpn-policies, in the Export-policy list, choose the relevant VPN profile (which contains the route policy: **L3VPN_NM-SRTE-RP-PE-A-7**). This forms the association between the VPN and the ODN template that defines the SLA.
 - f. Click **X** in the top-right corner when you are done.
 - g. Expand srv6 and slide the Enable srv6 toggle to the on position and then click  under address-family.
 - h. Select **ipv6** from address family list and click **Continue**.
 - i. For Locator-name, type **ALG0r5**. The SRv6 locator name should match locators configured at a node-global level on each router. Click **X** in the top-right corner until you are back on the Create L3VPN screen.
 - j. Similarly, create the other vpn-instance-profiles.
- Step 8** Define each VPN endpoint individually: PE-A, PE-B, and PE-C.
- a) Expand vpn-nodes and click  to select the relevant device from the list: **PE-A**. Click **Continue**.
 - b) Enter the local autonomous system number for network identification: **200**.

- c) Expand active-vpn-instance-profiles and click  to select the Profile-id you created in the previously: **L3VPN_NM-SRTE-RP-PE-A-7**. Click **Continue**.
- d) Define the network access parameters for communication from the PE towards the CE:
- Under vpn-network-accesses, click  to create a new set of VPN access parameters and provide a unique ID. Click **Continue**.
 - In the Interface-id field, type **Loopback70**. This is the identifier for the physical or logical interface. The identification of the sub-interface is provided at the connection level and/or the IP connection level.
 - Expand ip-connection > ipv6 and enter a Local-address (**70:70:70::1**) and the Prefix-length (**64**).
 - Expand routing-protocols and click  before typing a unique identifier for the routing protocol: **EBGP**. Click **Continue**.
 - From the routing protocol Type list, select **bgp-routing**.
 - Expand bgp and for Peer-as, type **70**. This information indicates the customer's ASN when the customer requests BGP routing.
 - From the Address-family list, select **ipv6**.
 - Under neighbor, click  to create a neighbor IP address and type **70:70:70::2**. Click **Continue**.
 - Type the Multihop number: **11**. This describes the number of IP hops allowed between a given BGP neighbor and the PE.
 - For redistribute-connected, click  and select **ipv6** from the Address-family list. Click **Continue**.
 - Click **X** in the top-right corner until you are back on the Create L3VPN screen.
 - Similarly, create the other VPN nodes: **PE-B** and **PE-C**.

Step 9 Commit your changes or click **Dry Run** to check what will be configured on the devices before you commit.

Step 10 Check that the new L3VPN service appears in the table and its provisioning state is **Success**.

Step 4 Visualize the New VPN Service on the Map to See the Traffic Path

Step 1 In the L3VPN Service table, click on the service name or click  in the Actions column and choose **View Details** from the menu.

The map opens and the service details are shown to the right of the map.

or

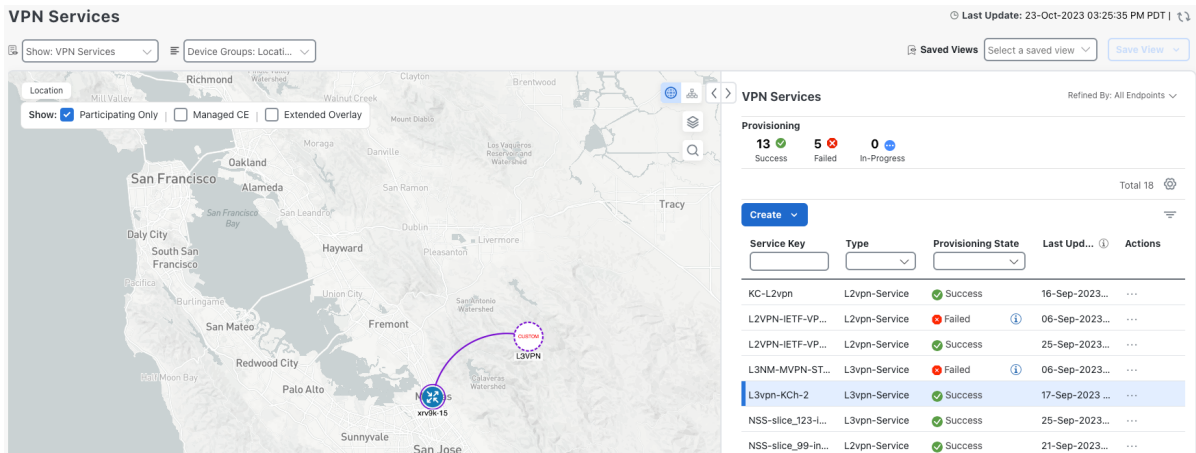
a) Go to **Services & Traffic Engineering > VPN Services**.

The map opens and a table of VPN services is displayed to the right of the map.

b) Click on the VPN in the Services table. If there are many services in the table, you can filter by name, type, or provisioning state to help locate the VPN.

Step 4 Visualize the New VPN Service on the Map to See the Traffic Path

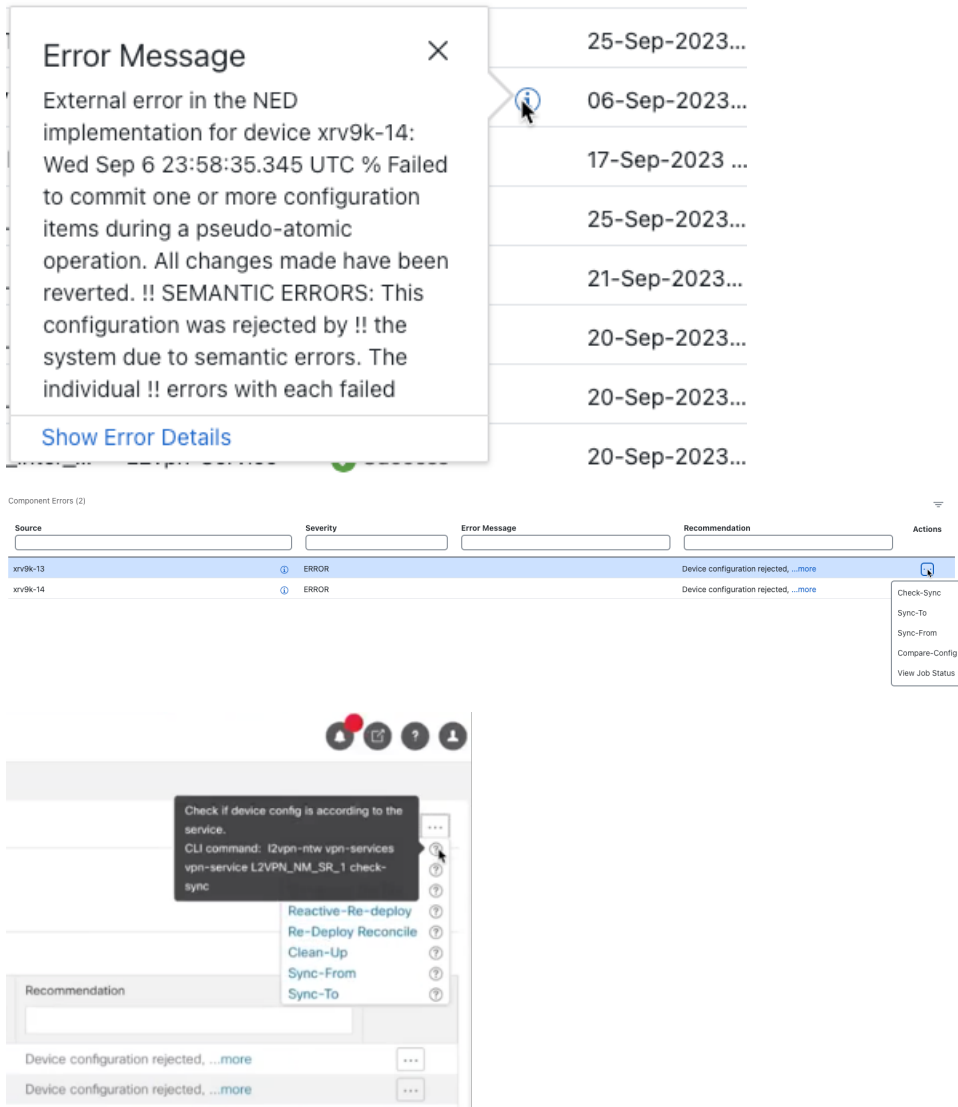
In the map, you will see the VPN as an overlay on the topology. It shows a representation of the three endpoints and a dashed line that indicates that it is a virtual path.




Select the **Show Participating Only** check box if you do not want to see the devices that are not involved in the selected VPN.

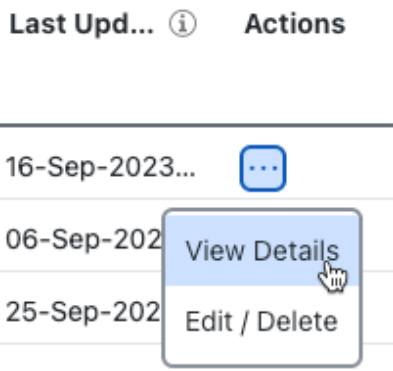
Note When a Provision State shows a Failed state, an information icon appears. This is true whether you are on the VPN Services, Service Details, and many of the Provisioning screens that show a table of services and their Provisioning status. If you select the icon, Error Message details appear describing the failure. You can also click the **Show Error Details** link to view the Component Errors screen and take action to fix the error. Each failed source provides further error message details and recommendations. For example, in the Action column for the failed source on the component Errors screen, you may click ⋮ for different options (such as, **Check-Sync**, **Sync-To**, **Sync-From**, **Compare-Config**, **View Job Status**) that will assist in fixing the error. Service level actions are also available for additional options (such as, **Re-Deploy**, **Reactive-Re-deploy**, **Re-Deploy Reconcile**, **Clean-up**, etc.) that will assist in fixing the service level error. Use the information icons that appear next to these options, as well, for further fix details.

vpn-Service	Success		16-Sep-2023...	...
vpn-Service	Failed	i	06-Sep-2023...	...
vpn-Service	Success		25-Sep-2023...	...




Step 2

In the Actions column, click  to drill down to a detailed view of the VPN service, including the device configurations and the computed transport paths.

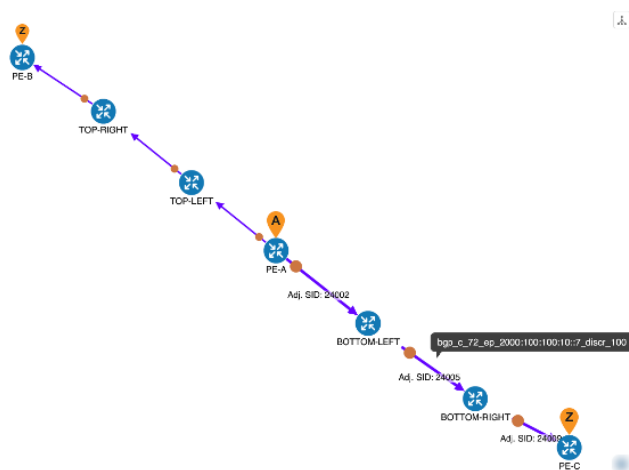


Step 3 To see the computed paths for this VPN, click on the Transport tab in the Service Details pane. All the dynamically created SR-TE policies are listed in the Transport tab. Select one or more SR-TE policies to see the path from endpoint to endpoint on the map.

In this example, we are looking at the disjoint paths computed from PE-A to PE-B and from PE-A to PE-C.



Step 4 To see the physical path between the endpoints, select the **Show IGP Path** check box in the top-left corner of the map. Hover with your mouse over a selected policy in the table to highlight the path in the map and show prefix SID and routing information.




Step 5 Observe automatic network optimization

The SR-PCE constantly monitors the network and automatically optimizes the traffic path based on the defined SLA. For illustration purposes, let's take a look at what happens when one of the links goes down, in this case, the link between P-BOTTOMLEFT and P-BOTTOMRIGHT. This means that the previous path from PE-A to PE-C is no longer viable. Therefore, the SR-PCE computes an alternative path, both from PE-A to PE-C and from PE-A to PE-B, in order to compensate for the link that is down and to maintain the disjoint paths.

Recomputed paths:

Source and Destination	Old path	New path
PE-A > PE-C	PE-A > BOTTOM-LEFT > BOTTOM-RIGHT > PE-C	PE-A > TOP-LEFT > BOTTOM-RIGHT > PE-C
PE-A > PE-B	PE-A > TOP-LEFT > TOP-RIGHT > PE-B	PE-A > BOTTOM-LEFT > TOP-RIGHT > PE-B



Summary and Conclusion

As we observed in this example, operators can use Cisco Crosswork Network Controller to orchestrate L3VPNs for SRv6 with SLAs and to maintain these SLAs using SR-TE policies that continuously track network conditions and automatically react to optimize the network. This automation increases efficiency and reduces human error that is generally unavoidable with manual tasks.

Scenario: Mandate a Static Path for an EVPN-VPWS Service using an Explicit MPLS SR-TE Policy

To ensure that mission-critical traffic within a VPN traverses the higher capacity interfaces, rather than the lower capacity interfaces, we will create a point-to-point EVPN-VPWS service and associate a preferred path (explicit) MPLS SR-TE policy on both endpoints for service instantiation. In this way, we will mandate a static path for the mission-critical traffic.

In this scenario, we will see how quick and easy it is to create SR-TE policies and VPN services by uploading a file containing all the required configurations. We will download sample files (templates) from the provisioning UI, fill in the required data, and then import the file via the UI. Lastly, we will use the Service

Health functionality to review the health of the services and view the Assurance Graph and Last 24Hr Metrics to better analyze our service's health details.



Note In this scenario, reference to SR-TE specifically means SR-TE over MPLS.

In this scenario, we will:

- Create a SID list - a list of prefix or adjacency Segment IDs, each representing a device or link along the path.
- Provision an explicit SR-TE policy, which will reference the SID list, thus creating a predefined path into which the EVPN prefix will be routed.
- Provision a point-to-point EVPN-VPWS service from PE-A to PE-C and attach the explicit SR-TE policy.
- Visualize the path of the service and review the health of the services.

Assumptions and Prerequisites

- For transport mapping to L2VPN service, devices must be configured with the **l2vpn all** command.
- For Service Health enablement, Service Health must be installed. See the Crosswork Network Controller Installation Guide chapter, Install Crosswork Applications.
- (Optional) Service Health provides **Internal Storage** of monitoring data up to a maximum limit of 50 GB. This data is stored on your system. If you exceed the limit of the internal storage, historical data will be lost. If you choose to extend Service Health storage capacity, you can configure **External Storage** in the cloud using an Amazon Web Services (AWS) cloud account. By leveraging External Storage, all existing internal storage data will be automatically moved to the external cloud storage (see **Configuring Service Health External Storage Settings** for more details) and your internal storage will act locally as cache storage. Configuring External Storage for Service Health ensures you will not lose historical data for services that continue to monitor a service's health, and will retain service health data for any service you choose to stop monitoring when you select the option to retain historical monitoring service for the data. For more information on Internal and External Storage, and how to retain historical monitoring service data when stopped, see **Configuring Service Health External Storage Settings** and **Stopping Service Health monitoring**.
- Before using Service Health's Assurance Graph, ensure that topology map nodes have been fully configured and created with a profile associated to the service. If not, Subservice Details metrics will show that no value has yet to be reported.
- For Service Health, you must configure 2 buckets on the Y1731 profile associated with the device. If you have fewer than 2 buckets configured, Service Health cannot report the Y1731 probes/KPIs on the service details page.

Step 2 Create the SID List in the Provisioning UI

```

{
  "cisco-crosswork-optimization-engine-sr-policy-operations:output": {
    "segment-list-hops": [
      {
        "step": 0,
        "sid": 23002,
        "ip-address": "100.100.100.7",
        "type": "node-ipv4"
      }
    ],
    "igp-route": [
      {
        "node": "PE-A",
        "interface": "GigabitEthernet0/0/0/0"
      },
      {
        "node": "P-TOPLEFT",
        "interface": "GigabitEthernet0/0/0/2"
      },
      {
        "node": "P-BOTTOMRIGHT",
        "interface": "GigabitEthernet0/0/0/3"
      }
    ],
    "state": "success",
    "message": ""
  }
}

```

Step 2 Create the SID List in the Provisioning UI

In this scenario, we will create a SID list for traffic from PE-C to PE-A and another SID list for traffic in the opposite direction.

Step 1 Go to **Services & Traffic Engineering > Provisioning (NSO) > SR-TE > SID-List**.

Step 2 Click **+** to create a new SID list and give it a unique name. For this example, the SID list name is **L2VPN_NM-P2P-SRTE-PE-C-240**. Click **Continue**.

Step 3 Under sid, click **+** to create a new SID index and give it a numeric value. Click **Continue**.

Step 4 Under mpls, enter the SID ID that was received in the API response in Step 1.

Create SR-TE > SID-List

The screenshot displays the configuration interface for a new SID list. The main configuration area shows the name 'L2VPN_NM-P2P-SRTE-PE-C-240' and a table with one SID index '1'. A right-hand pane shows the configuration for the selected SID index, including the type 'mpls' and the label '23002'.

Step 5 Click **X** in the top-right corner to return to the SID list. Your new SID appears in the index table.

Step 6 Repeat these steps to create additional SID indexes, as required.

- Step 7** Commit your changes.
- Step 8** Check that the new SID list appears in the table.
- Step 9** Create another SID list for the traffic from PE-A to PE-C. For this example, the SID list name is **L2VPN_NM-P2P-SRTE-PE-A-240**.

Step 3 Create an explicit SR-TE policy for each VPN endpoint by importing a file

In this step, we will provision two explicit SR-TE policies which will reference the SID lists created in Step 1.

The first SR-TE policy specifies PE-C as the headend and provides the IP address of PE-A as the tail end. The second SR-TE policy specifies PE-A as the headend and provides the IP address of PE-C as the tail end.


Instead of manually filling in each field in the provisioning UI, we will import an xml file containing all the configurations required to create the SR-TE policy.

Step 1 Go to **Services & Traffic Engineering > Provisioning (NSO) > SR-TE > Policy**.


Step 2 Click  to import.

Step 3 Download the sample .json or .xml file which will serve as a template for the required configuration. In the Import Service dialog, click the **Download sample .json and .xml files (.zip)** link

Import Service

 Sample xml or json files contains basic service parameter that can be modified in your local machine, and then imported back into crosswork to create a new service.

Search to identify service type of imported file

Policy 

File Name

Browse

[Download sample .json and .xml files \(.zip\)](#)

Cancel

Import

Step 4 Unzip the downloaded file and open sr-Policy.xml in an XML editor.



Step 5 Edit the xml file as required. Provide a name for the SR-TE policy, and specify the SID list to be associated with this policy. Save the xml file.

Step 4 Create and provision the L2VPN service

```

<config xmlns="http://tail-f.com/ns/config/1.0">
  <sr-te xmlns="http://cisco.com/ns/nso/cfp/cisco-tdsn-sr-te">
    <policies xmlns="http://cisco.com/ns/nso/cfp/cisco-tdsn-sr-te-sr-policies">
      <policy>
        <name>SR-Policy-1</name>
        <head-end>
          <name>iosxrv-5</name>
        </head-end>
        <tail-end>7.7.7</tail-end>
        <color>100</color>
        <binding-sid>100</binding-sid>
        <path>
          <preference>100</preference>
          <dynamic>
            <metric-type>te</metric-type>
            <metric-margin>
              <relative>40</relative>
            </metric-margin>
            <constraints>
              <sid-limit>10</sid-limit>
            </constraints>
          </dynamic>
        </path>
        <path>
          <preference>200</preference>
          <explicit>
            <sid-list>
              <name>mysidlist</name>
              <weight>10</weight>
            </sid-list>
            <constraints>
              <affinity>
                <rule>
                  <action>include-all</action>
                  <color>GREEN</color>
                  <color>RED</color>
                </rule>
              </affinity>
            </constraints>
          </explicit>
        </path>
      </policy>
    <sid-list>
      <name>mysidlist</name>
      <sid>
        <index>1</index>
        <mpls>
          <label>17001</label>
        </mpls>
      </sid>
    </sid-list>
  </policies>
</sr-te>
</config>

```

- Step 6** In the Import Service dialog, select **Policy** as the type of file to import, browse to the edited xml file, and click . If there are any errors in the file, you will be notified. If there are no errors, the file will be imported. The policy will be created and the devices will be configured accordingly.
- Step 7** Check whether the new SR-TE policy appears in the Policy table and its Provisioning State is **Success**.
- Step 8** Click  in the Actions column and choose **Config View** to see to see the Yang model-based service intent data that details the SR-TE policy you created. You can also check the devices themselves to make sure that they were provisioned correctly.

Step 4 Create and provision the L2VPN service

In this step, we will create and provision a P2P VPN service with PE-A and PE-C as the endpoints. The VPN service will reference the SR-TE policies we created in the previous step to ensure that the traffic traversing the VPN will follow the path defined in the SID lists.

As we did with the SR-TE policy, we will create the VPN service by importing an xml file containing all the required configurations. Once we have provisioned the VPN service, we will edit it in the provisioning UI in order to associate the SR-TE policies.

Step 1 Go to **Services & Traffic Engineering > Provisioning (NSO) > L2vpn > L2vpn-Service**.

Step 2 Click  to import.

Step 3 If you did not download the sample .json or .xml files in Step 3, do so now.


Step 4 Open l2nm.xml in an XML editor.

Step 5 Edit the xml file as required. Provide a name for the L2VPN, configure each endpoint, and define the VPN parameters.


This is the configuration for PE-A in our example:

```
<vpn-node-id>xrv9k-22</vpn-node-id>
<signaling-option>
  <ldp-or-l2tp>
    <pw-peer-list>
      <peer-addr>192.168.0.22</peer-addr>
      <vc-id>100</vc-id>
      <mpls-label xmlns="http://cisco.com/ns/nso/fp/examples/cisco-l2vpn-ntw">100</mpls-label>
    </pw-peer-list>
  </ldp-or-l2tp>
</signaling-option>
<vpn-network-accesses>
  <vpn-network-access>
    <id>300</id>
    <interface-id>GigabitEthernet0/0/0/1</interface-id>
    <connection>
      <encapsulation>
        <encap-type xmlns:vpn-common="urn:ietf:params:xml:ns:yang:ietf-vpn-common">vpn-common:dot1q</encap-type>
        <dot1q>
          <cvlan-id>100</cvlan-id>
        </dot1q>
      </encapsulation>
    </connection>
  </vpn-network-access>
</vpn-network-accesses>
<te-service-mapping xmlns="http://cisco.com/ns/nso/fp/examples/cisco-l2vpn-ntw">
  <te-mapping>
    <sr-policy>
      <policy-type>policy</policy-type>
      <policy>SR-300</policy>
    </sr-policy>
  </te-mapping>
</te-service-mapping>
</vpn-node>
<vpn-node>
  <vpn-node-id>xrv9k-23</vpn-node-id>
```

Step 6 Save the xml file.


Step 7 In the Import Service dialog, select **l2vpn service** as the type of file to import, browse to the edited xml file, and click . If there are any errors in the file, you will be notified. If there are no errors, the file will be imported. The service will be created and the devices will be configured accordingly.

Step 8 Check that the new L2VPN service appears in the L2VPN Service table and its Provisioning State is **Success**.

Step 9 Click  in the Actions column and choose **Config View** to see the Yang model-based service intent data that details the VPN service you created. You can also check the devices themselves to make sure that they were provisioned correctly.

Step 5 Attach the SR-TE policies to the L2VPN Service

At this stage, the provisioned L2VPN service you created does not have associated SR-TE policies that define the transport path. In this step, we will edit the L2VPN service in the provisioning GUI, attach the relevant SR-TE policies to each endpoint, and re-provision it.

-
- Step 1** Locate the L2VPN in the VPN Service table.
- Step 2** Click  in the Actions column and choose **Edit**.
- Step 3** Under vpn-nodes, select **PE-A** and click the **Edit** button above the table.
- Step 4** In the pane that opens on the right, open the **te-service-mapping > te-mapping** section.
- Step 5** In the sr-policy tab, in the policy field, enter the name of the SR-TE policy created for PE-A: **L2VPN_NM-P2P-SRTE-PE-A-240**.
- Step 6** Click **X** in the top-right corner to close the PE-A pane.
- Step 7** Repeat the above steps for PE-C and attach the SR-TE policy: **L2VPN_NM-P2P-SRTE-PE-C-240**.
- Step 8** Click **Commit Changes**.
-


Step 6 Enable Service Health monitoring

After creating and provisioning the required L2VPN services, you can begin monitoring their health.

Before you begin

- Ensure that Crosswork Service Health is installed. For details, see the "Install Crosswork Applications" chapter in the [Cisco Crosswork Network Controller Installation Guide](#).
- Ensure that the required L2VPN services are created and provisioned.

To enable service health monitoring, do the following:

-
- Step 1** From the main menu, choose **Services & Traffic Engineering > VPN Services**. The map opens on the left side of the page and the table opens on the right side.
- Step 2** In the Actions column, click  for the service you want to start monitoring the health.
- Step 3** Click **Start Monitoring**.

VPN Services Refined By: All Endpoints ▾

Provisioning Health (Monitoring: 3 Services)

5 Success
0 Failed
0 In-Progress
2 Good
1 Degraded
0 Down

Total 5

Create ▾ ☰

Health	Service ...	Type	Provisioni...	Las... <small>ⓘ</small>	Actions
	L2VPN_N...	L2vpn-Ser...	Success	26-Jul-...	...
	L3NM-PR...	L3vpn-Ser...	Success	26-Jul-...	...
	L3NM-PR...	L3vpn-Ser...	Success	26-Jul-...	
	L3NM-PR...	L3vpn-Ser...	Success	26-J	View Details
	L3NM-PR...	L3vpn-Ser...	Success	26-J	Edit / Delete
					Start Monitoring

Note The Health column color coding indicates the health of the service:

- Blue = Initiated
- Green = Good
- Orange = Degraded
- Red = Down
- Gray = Not Monitoring

Step 4

In the Monitor Service dialog box, select the Monitoring Level. For help with selecting the appropriate monitoring level for your needs, see the **Cisco Crosswork Network Controller 6.0 Service Health Guide > Basic and Advanced Monitoring Rules** section.

Monitor Service

Name L3NM-PROBES-45-2-3-endpoint

Monitoring Level ?

Silver_L3VPN_ConfigProfile custom

Gold_L3VPN_ConfigProfile custom

Basic Monitoring

Advanced Monitoring

Thresholds to use for Silver L3VPN services

Cpu Threshold Max	80.5 %
Memfree Threshold Min	1000000000 bytes

Cancel


Once you have started monitoring the health of this service, if you select the Actions column and click to view additional Service Health options, you will see: **Stop Monitoring**, **Pause Monitoring**, **Edit Monitoring Settings**, **Assurance Graph**.

Note If you select **Edit Monitoring Settings**, you may update the Monitoring Level setting from Basic Monitoring to Advanced Monitoring, or from Advanced Monitoring to Basic Monitoring, at any time.

Note If you later decide to **Stop Monitoring** a service that has already been started, you have the option to retain the historical service data for that stopped service. See the **Cisco Crosswork Network Controller 6.0 Service Health Guide > Stop Service Health Monitoring** section for more detailed information.







Step 5 Click **Start Monitoring**.


Note

Once you have started monitoring the health of the service, in the Actions column, if you click  to view additional Service Health options, you will see: Stop Monitoring, Pause Monitoring, Edit Monitoring Settings, and Assurance Graph.













VPN Services Refined By: All Endpoints ▾

Provisioning Health (Monitoring: 3 Services)

5  Success 0  Failed 0  In-Progress 2  Good 1  Degraded 0  Down

Total 5 

Create ▾ ☰


Health	Service ...	Type	Provisioni...	Las... 	Actions
	L2VPN_N...	L2vpn-Ser...	 Success	26-Jul-...	...
	L3NM-PR...	L3vpn-Ser...	 Success	26-Jul-...	
	L3NM-PR...	L3vpn-Ser...	 Success		<div style="border: 1px solid #ccc; padding: 5px; width: 150px;"> <p>View Details</p> <p>Edit / Delete</p> <p>Stop Monitoring</p> <p>Pause Monitoring</p> <p>Edit Monitoring Settings</p> </div>
	L3NM-PR...	L3vpn-Ser...	 Success		
	L3NM-PR...	L3vpn-Ser...	 Success		

Step 6 Repeat these steps for each service that you wish to start health monitoring.

Step 7 Click **X** in the top-right corner when you are done.

Step 7 Visualize the L2VPN on the Map

In this step we will take a look at the representation of the L2VPN on the map, and we'll see the paths the traffic will take from PE-A to PE-C and vice versa, based on the explicit SR-TE policies we created.

Step 1 In the L2VPN Service table, in the Actions column for the new VPN, click  and choose **ViewDetails** from the menu. The map opens and the service details are shown to the right of the map.

or


Go to  **Services & Traffic Engineering > VPN Services**.


The map opens and a table of VPN services is displayed to the right of the map.

- a) Click on the VPN in the Services table. If there are many services in the table, you can filter by name, type, or provisioning state to help locate the VPN.


Step 8 Inspect a degraded service using Service Health and Last 24Hr Metrics to identify issues

- b) In the map, you will see the VPN as an overlay on the topology. It shows a representation of the endpoints and a solid line that indicates that it is a virtual path.
- c) Select the **Show Participating Only** check box if you do not want to see the devices that are not involved in the selected VPN.



Step 2 Under the Actions column, click  and choose **View Details** to drill down to a detailed view of the VPN service, including the device configurations, the computed transport paths, and the health status for transport paths.

Step 3 In the Transport tab, select one or more SR-TE policies to see the path from endpoint to endpoint on the map. The image below shows the path for PE-C to PE-A. The **Show IGP Path** check box in the top left corner of the map is selected so the physical path is shown. The dashed line indicates that this link is being used to transport multiple services.




Step 8 Inspect a degraded service using Service Health and Last 24Hr Metrics to identify issues

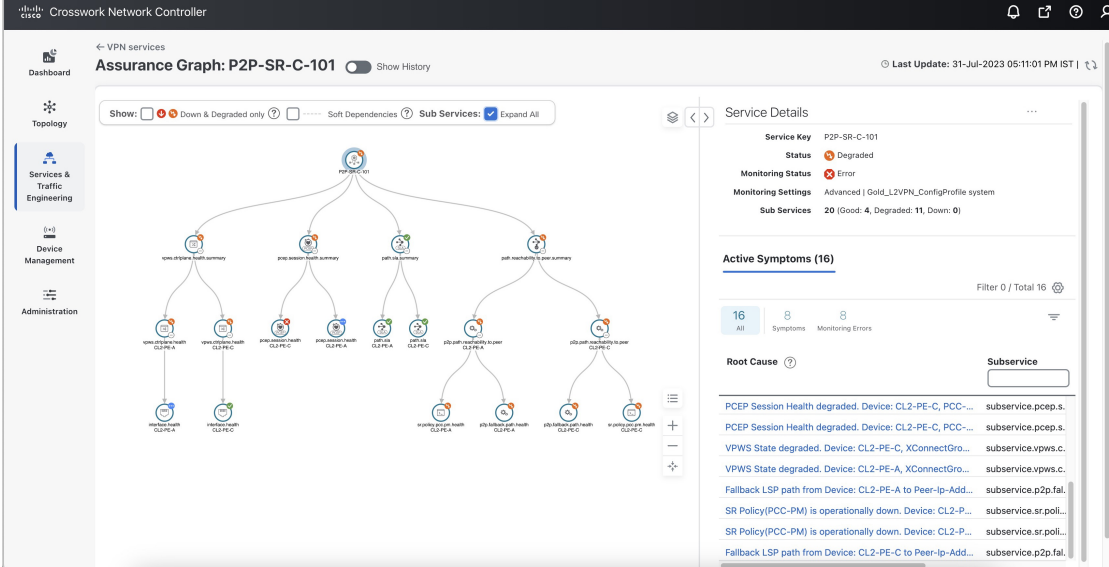
In this step, you can utilize the Last 24Hr Metrics to identify the issues with the degraded services within a specific time range. By isolating the issues within a specific time range, you can drill down on the details that may have caused the degraded (or down) service that can lead to troubleshooting the service or the node to address detailed symptoms.

Step 1 Return to the VPN Services list.


Step 2

In the Actions column, click  for the degraded service and click **Assurance Graph**. The topology map of services and subservices appear with the Service Details panel showing Service Key, Status, Monitoring Status, Monitoring Settings, Sub Services, and Active Symptoms details.

Note This may take up to 5-10 minutes to update after a service has been enabled for monitoring.


Step 3

At the top of the page, click the **Show History** mode toggle. The historical Date Range graph appears. This graph shows different ranges of historical health service monitoring details from one day (1d) up to sixty days (60d). When you hover over an event on the Date Range graph, a tool tip with information about that event appears (such as date and time of the event, and number of symptoms).

**Step 4**

Review the Root Cause information by clicking a particular event in the graph. The Service Details panel reloads, showing the active symptoms and the root causes associated with the event. Columns can be resized using your mouse or you can select the gear icon to deselect or select columns you want to appear.

Note Once you enable **Show History** mode, Root Cause information in the Active Symptoms table will start to show the blue Last 24Hr Metrics icon. Data from the device will be initially delayed, however, and may take some time before **Last 24Hr Metrics** begins to populate with data. Until then, the value of zero is reported.

Step 8 Inspect a degraded service using Service Health and Last 24Hr Metrics to identify issues

Service Details

Service Key L3VPN_U_NM-SRTE-ODN-1061

Status ⚠ Degraded

Monitoring Status ✖ Error

Monitoring Settings Advanced | Gold_L3VPN_ConfigProfile custom

Sub Services 27 (Good: 19, Degraded: 5, Down: 0)

Symptoms (4)

4 All | 2 Symptoms | 2 Monitoring Errors Total 4 ⚙️ ☰

Root Cause ⓘ **Subservice**

Root Cause	Subservice
Unable to get feed from device for metric(s): metric.inter...	subservice.interfa...
Unable to get feed from device for metric(s): metric.inter...	subservice.interfa...
eBGP Session to neighbor 10.10.10.238 is not up for D...	Last 24Hr Metrics ...
eBGP Session to neighbor 10.10.10.238 is not up for Device: CL2-PE-A, Vrf: L3VPN_U_NM-SRTE-ODN-1061	

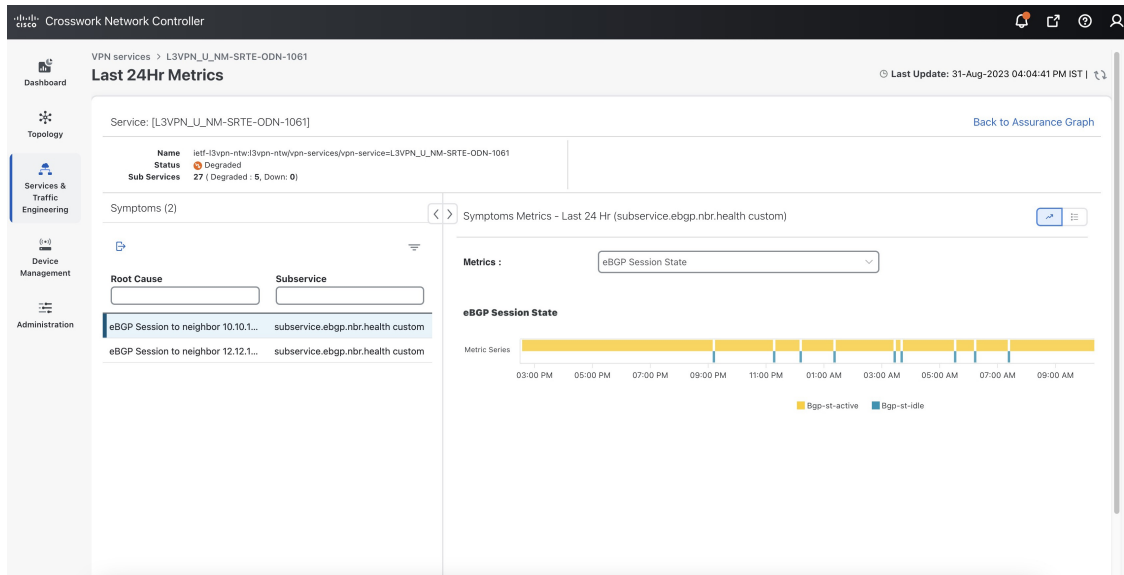
Step 5

To further isolate the possible issues and to utilize the **Last 24Hr Metrics**, perform the following steps:

- In the Date Range graph, use your mouse to select the range of historical health service monitoring details from one day (1d) up to sixty days (60d).

Note At the top-right of the Date Range graph, select the appropriate icons to either zoom in or out, horizontally scroll through the date ranges, or refresh the graph to go back to the most recent event, for example. You can also use your mouse to draw a rectangle over events to further zoom in on the degraded devices. Events that are consecutive may appear as a line of white space.

- Click on a degraded event in the graph. The Service Details panel reloads, showing any active symptoms and the root causes to be inspected. Expand the table and information as necessary for further details.




Step 6 Check the **Down & Degraded Only** check box at the top-left corner of the map to show only the Subservices which are degraded, along with other dependent but healthy subservices. Inspect the Service Details panel showing the active symptoms and their root cause. Uncheck the **Down & Degraded Only** check box and check the **Soft Dependencies** check box in the top-left corner of the map. Soft Dependencies implies that a child subservice's health has a weak correlation to its parent's health. As a result, the degraded health of the child will not result in the parent's health degradation.

Use the + or – symbols in the bottom-right corner of the map to zoom in or out on services mapped. Select the ? to view the Link Color Legend that explains all of the icons, symbols, badges, and colors and their definitions.

Step 7 Select the degraded subservice in the map to show the subservice details.

Step 8 Click the **Symptoms** tab to show any root causes for the service health details that are displayed and then click the **Impacted Services** tab to view the impacted services.

Step 9 Click **X** in the top-right corner to return to the VPN Services list and in the Actions column, click  for the degraded service in the list and click **Assurance Graph** to show the Service Details panel.

Step 10 Again, select the **Show History** toggle in the top-right corner of the Service Details panel before selecting the blue metrics icon in one of the Root Cause rows. The Symptoms Metrics – Last 24 Hr bar chart appears. This chart provides details of the metric patterns for different sessions states (such as active, idle, failed) for individual root cause symptoms with Status, Session, Start Time, and Duration information to assist in troubleshooting prevailing issues. Use your mouse to hover over the chart to view the different details.

Continue to troubleshoot a service health issue using Parameterized Jobs

To further troubleshoot a service health issue (such as a device that is degraded due to not properly fetching data), continue with the following steps to examine if the issue is associated with a collection job.

Step 11 From the main menu, choose **Administration > Collection Jobs**.

The Collection Jobs page appears.

Step 12 Click the **Parameterized Jobs** tab.

- Step 13** Review the Parameterized Jobs list to identify the devices that may have service health degradation issues. By reviewing Parameterized Jobs, you can identify and focus on gNMI, SNMP, and CLI-based jobs by their Context ID (protocol) for further troubleshooting purposes.
- Step 14** In the Job Details panel, select the collection job you want to export and download the status of collection jobs for further examination. The information provided is collected in a .csv file when the export is initiated.
- Note** When exporting the collection status, you must fill in the information each time an export is executed. In addition, make sure to review the **Steps to Decrypt Exported File** content available on the Export Collection Status dialog box to ensure you can access and view the exported information.
- Step 15** Click **Export**.
- Step 16** To check the status of the exported collection job data, click **View Export Status** at the top right of the Job Details panel. The Export Status Jobs panel appears providing the status of the export request.
- Step 17** Review the exported .csv file for collection job details and the possible cause of the degraded device.
-

Summary and Conclusion

In this scenario, we observed how simple it is to create explicit SR-TE policies and attach them to a L2VPN service in order mandate a static path for the mission-critical traffic. We saw how editing a pre-defined template and then importing it into the system enables quick and easy provisioning of services and SR-TE policies. We were then able to visualize the actual traffic paths on the map. Lastly, we used Service Health to monitor the health of the new service using the Assurance Graph, Last 24hr Metrics, and SubExpressions metrics to view when service may have been up, degraded, or down, and what the root causes were identified.

Scenario: Provision an L2VPN Service over an RSVP-TE Tunnel with Reserved Bandwidth

For the continuous stream transmission required for rich data media types, such as video and audio, bandwidth reservation is often required to provide higher quality of service. Cisco Crosswork Network Controller supports the creation and management of RSVP-TE tunnels to reserve guaranteed bandwidth for an individual flow. RSVP is a per-flow protocol that requests a bandwidth reservation from every node in the path of the flow. The endpoints, or other network devices on behalf of the endpoints, send unicast signaling messages to establish the reservation before the flow is allowed. If the total bandwidth reservation exceeds the available bandwidth for a particular LSP segment, the LSP is rerouted through another LSR. If no segments can support the bandwidth reservation, LSP setup fails and the RSVP session is not established.

In this scenario we will:

- Create RSVP-TE tunnels with reserved bandwidth.
- Enable Bandwidth on Demand functionality.
- Provision a VPN service from PE-A to PE-B and attach the RSVP-TE tunnels as underlay configuration.
- Visualize the path of the traffic when link utilization is below the bandwidth threshold. This path would change if the bandwidth utilization on the link crossed the specified threshold.


Assumptions and Prerequisites

Scenario 4 to provision an L2VPN service over an RSVP TE Tunnel with reserved bandwidth the following are the assumptions and prerequisites.

- For transport mapping to L2VPN service, devices must be configured with the **l2vpn all** command.
- For Service Health enablement and usage to monitor a services health, Service Health must be installed.
- For steps to enable Service Health during this scenario, see Scenario 3, [Step 6 Enable Service Health monitoring](#). For additional Service Health related details, see [Scenario: Implement and Maintain SLA for an L3VPN Service for SR-MPLS \(using ODN\)](#), [Scenario: Mandate a Static Path for an EVPN-VPWS Service using an Explicit MPLS SR-TE Policy](#).
- (Optional) Service Health provides **Internal Storage** of monitoring data up to a maximum limit of 50 GB. This data is stored on your system. If you exceed the limit of the internal storage, the least recently used historical data will be lost. If you choose to extend Service Health storage capacity, you can configure **External Storage** in the cloud using an Amazon Web Services (AWS) cloud account. By leveraging External Storage, all existing internal storage data will be automatically moved to the external cloud storage (see [Configuring Service Health External Storage Settings](#) for more details) and your internal storage will act locally as cache storage. Configuring External Storage for Service Health ensures you will not lose historical data for services that continue to monitor a service's health, and will retain service health data for any service you choose to stop monitoring when you select the option to retain historical monitoring service for the data. For more information on Internal and External Storage, and how to retain historical monitoring service data when stopped, see [Configuring Service Health External Storage Settings](#) and [Stopping Service Health monitoring](#).
- (Optional) For initializing a Heuristic Package to monitor health of a services, see [Initializing Heuristic Packages to monitor the health of a service](#) for detailed steps to be performed prior to starting monitoring.

Step 1 Create an RSVP-TE tunnel for both directions of the L2VPN

In this step, we will create an RSVP-TE tunnel from PE-A to PE-B and from PE-B to PE-A, and we'll reserve bandwidth of 1200 on the link.

-
- Step 1** Go to Services & Traffic Engineering > Provisioning(NSO) > **RSVP-TE** > **Tunnel**.
 - Step 2** Click  to create a new RSVP-TE tunnel and give it a unique name. Click **Continue**.
 - Step 3** In the Identifier field, enter a numeric identifier for the tunnel. You will use this identifier later when you associate this RSVP-TE tunnel with the L2VPN service. For this example, the identifier is **2220**.
 - Step 4** In the source and destination fields, enter the loopback0 IP address of the source (PE-A) and the destination (PE-B) devices. This is the TE router ID. To find the TE router ID, go to Topology and click on a device in the map or in the list of devices. The Device Details pane opens and the TE router ID is shown under the Routing section.

Step 1 Create an RSVP-TE tunnel for both directions of the L2VPN

Device Details
✕

Details
Links

Summary
^

Host Name	PE-A
Reachability	✔ Reachable
IP Address	192.168.11.79
Geo Location	Latitude 33.436665, Longitude -112.048822
Device Type	🔌 Router
Device Group	Location > All Locations > Unassigned Devices
Product Type	Cisco IOS XRv 9000 Router
Connect To Device	➤ Telnet IPv4 🔒 SSH IPv4
Last Update	22-Oct-2023 03:58:16 PM PDT

Routing
^

TE Router ID	100.100.10.5
IPv6 Router ID	2000:100:100:10::5
ISIS System ID	0000.0000.0005 Level-1/2
ASN	200

Step 5 Define the endpoints:

- a) Under head-end, select the headend device from the dropdown list.
- b) Under tail-end, select the tailend device from the dropdown list.

Step 6 Reserve bandwidth on the link. Under te-bandwidth > generic, enter the bandwidth threshold for the link.

Step 7 Define the path of the RSVP-TE tunnel.

You have the option to define an explicit path or to have the path locally computed by the participating devices. Alternatively, you can have the SR-PCE compute a path dynamically. For this scenario we will have the path locally computed.

- a) Under p2p-primary-paths, click + to create a new path.
- b) In the pane that opens on the right, give the path a name.
- c) Select the path computation method – **path-locally-computed**.
- d) Specify a numeric preference for the path. The lower the number, the higher the preference.

e) Define the optimization metric, in this case,

The screenshot shows the configuration interface for a primary path. The 'primary-path' section is expanded, showing a table with one entry: 'L2VPN_NM-P2P-RSVPTE-PE-A-2220' with 'path-computation-method' set to 'te-types:path-locally-computed' and 'preference' set to '1'. The 'optimizations' section is also expanded, showing 'explicit-route-objects-always'.

igp.

[Commit changes](#) [Dry Run](#) [Cancel](#)

Step 8

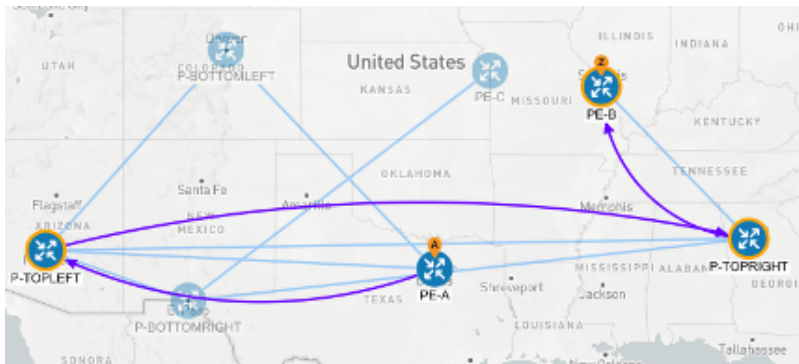
Click **Commit Changes**.

Step 9

Verify that the RSVP-TE tunnel appears in the list of tunnels and its Provisioning State is **Success**.

Step 10

Click on the tunnel name to visualize the tunnel on the map and to see the tunnel details.



Step 2 Create the L2VPN service and attach the RSVP tunnel to the service

In this step, we will create a P2P L2VPN service using the provisioning GUI. If you want to create the service by importing a template, refer to Scenario 3—Mandate a static path for an EVPN-VPWS service using an explicit SR-TE policy

- Step 1** Go to **Services & Traffic Engineering > Provisioning (NSO) > L2VPN > L2vpn Service**.
- Step 2** Click **+** to create a new service and give it a unique name. Click **Continue**.
- Step 3** Choose the vpn-type field.
- Step 4** Define each VPN endpoint individually – PE-A and PE-B.
- Under vpn-nodes, click **+**.
 - Select the relevant device from the vpn-node-id and ned-id dropdown lists and click **Continue**.
- Step 5** Define the LDP signaling options by creating one or more pseudowires. In this case, specify the TE router ID of the peer device (PE-B), and provide a unique numeric label to identify the pseudowire.
- Step 6** Attach the RSVP tunnel to the service:
- Under te-service-mapping > te-mapping, click the te-tunnel-list tab.
 - Click the **ietf-te-service** tab.
 - Enter the name of the RSVP-TE tunnel you want to attach to this L2VPN service. The tunnel ID will be extracted from the tunnel configuration.

te-service-mapping ⓘ ^

te-mapping ⓘ ^

Te

sr-policy **te-tunnel-list** odn

te-tunnel-list ^

Enable te-tunnel-list

Tunnel-te-id-source *

te-tunnel-id **ietf-te-service**

ietf-te-service ⓘ

ⓘ

fallback ⓘ

▾

Note If you have an RSVP-TE tunnel on the device that was configured externally to Crosswork Network Controller, you can provide the tunnel ID under the te-tunnel-id tab.

- Step 7** Define the VPN network access. In this case, we are using dot1q encapsulation and we have specified the physical interface (GigabitEthernet0/0/0/2) and the VLAN ID (2220).
- Step 8** Follow the above steps for PE-B as well.
- Step 9** Click **Commit Changes**. Verify that the L2VPN appears in the list of VPN services and that its Provisioning state is **Success**.

Step 3 Visualize the L2VPN service on the map

In this step we'll take a look at the representation of the L2VPN on the map and we'll see the paths the traffic will take from PE-A to PE-B and vice versa, based on the RSVP-TE tunnels we created.

Step 1 In the L2VPN Service table, click on the service name. The map opens and the service details are shown to the right of the map.

or

a) Go to **Services & Traffic Engineering > VPN Services**.

The map opens and a table of VPN services is displayed to the right of the map.

b) Click on the VPN in the Services table. When there are many services in the table, you can filter by name, type, or provisioning state to help locate the VPN.

In the map, you will see the VPN as an overlay on the topology. It shows a representation of the three endpoints and a dashed line that indicates that it is a virtual path.

Note The image below shows the VPN overlay in the geographical map. Use the buttons at the top right of the map to toggle between the logical and geographical maps.

Step 2 To see the hops in the route between PCC7_56 and PCC5_81, click the Transport tab and select one or more of the underlying TE tunnels to see the path from endpoint to endpoint on the map. The image below shows both RSVP-TE tunnels selected in the Transport tab and the route from PCC7_56 to PCC5_81 as shown on the logical map.

Step 3 As the RSVP-TE tunnels are configured with a reserved bandwidth, if the bandwidth utilization across the link exceeds the specified bandwidth, the path would be rerouted.

Summary and Conclusion

This scenario illustrated how to create RSVP-TE tunnels with reserved bandwidth and attach them to an L2VPN service to meet the high quality of service requirements for continuous streaming of rich data media. We observed the path on the map. This path would be recomputed if the bandwidth utilization on the link crossed the bandwidth reservation threshold.

Scenario: Provision a Soft Bandwidth Guarantee with Optimization Constraints

Service providers must be able to provide fast connections with the lowest latency possible to meet the needs of customers' peak traffic utilization times and to dynamically optimize services based on the customers' changing priorities throughout the day. For this purpose, the operator might need to reserve bandwidth on specific links to ensure a dedicated path that can handle a set amount of traffic with a specific optimization intent. The Bandwidth on Demand (BWoD) feature within Crosswork Network Controller enables this functionality. Paths with the requested bandwidth are computed when available. If a path that guarantees the requested bandwidth cannot be found, an attempt will be made to find a *best effort* path.

In this scenario, we will use BWoD to calculate the lowest TE metric path with a specified amount of available bandwidth between two endpoints.

This scenario uses the following topology as a base:



The goal is to create a path from F2.cisco.com to F7.cisco.com that can accommodate 250 Mbps of traffic while keeping the utilization at 80%. BWoD will initially try to find a single path to accommodate the requested bandwidth without exceeding the utilization threshold. If a single path cannot be found, BWoD may recommend splitting the path.

In this scenario we will:

- Orchestrate a new SR-TE policy with bandwidth and TE constraints.
- Configure and enable BWoD.
- Verify the state of the SR-TE policy and view the path on the map.

Step 1 Create a BWoD SR-TE Policy with the Requested Bandwidth and Optimization Intent

To create a BWoD SR-TE Policy with the Requested Bandwidth and Optimization Intent

Step 1 Go to **Services & Traffic Engineering > Provisioning (NSO) > SR-TE > Policy**.

Step 2 Click **+** to create a new SR-TE policy and give it a unique name. Click **Continue**.

Step 3 Define the endpoints:

- Under head-end, click **+** and select the headend device from the dropdown list and click **Continue**. Click **X** to close the Headend pane.
- Enter the IP address of the tail-end device.
- Enter a color to identify the traffic.

Step 4 Define the parameters on which the path will be computed:

- Under path, click **+**.
- Enter a path preference and click **Continue**.

- c) In the dynamic-path tab, select **te** in the metric-type dropdown list as the optimization objective.
- d) Select the **pce** check box to have the SR-PCE compute the paths for this policy.

path{123} ↻ ✕

preference • ⓘ
123 ⓘ

Sr-te-path-choice

explicit-path **dynamic-path**

dynamic ⓘ ^

Enable dynamic

pce ⓘ

metric-type ⓘ
te ▼

constraints ▼

metric-margin ⓘ ▼

- e) Click **X** to close the path pane.

Step 5

In the **Bandwidth** field enter the requested bandwidth in Kbps. In this case, we are requesting **250** Mbps or 250,000 Kbps.

Step 1 Create a BWoD SR-TE Policy with the Requested Bandwidth and Optimization Intent

head-end* ⓘ ⓘ

+ 🗑️

name

xrv9k-23

tail-end* ⓘ

192.168.0.25

color* ⓘ

15130 ⓘ

binding-sid ⓘ

path* ⓘ ⓘ

+ ✎ 🗑️

preference

123

bandwidth ⓘ

250000 ⓘ

Commit changes Dry Run Cancel


Step 6 Click **Commit Changes**. The new policy is created and appears in the list of SR-TE policies. The provisioning state should be **Success**.

SR-TE > Policy Total 5 | Last Refresh: 24-Oct-2023 03:03:05 PM PDT | 🔍 | 🔄

+ 📄

Name	Provisioning State	Date Created	Actions
Policy-HE-12	Success	06-Sep-2023 01:13:25 PM PDT	...
Policy-HE-13	Success	06-Sep-2023 01:22:54 PM PDT	...

Step 7 Verify the new policy by viewing its details and its representation on the map:

- Click  in the Actions column and choose **View**.
- The map opens with the SR-TE policy details displayed to the right of the map.

Note The operational state of the policy is down because the SR-PCE alone is not able to address bandwidth computations before the BWoD functionality within Crosswork Network Controller is enabled.

Step 2 Enable and Configure BWoD

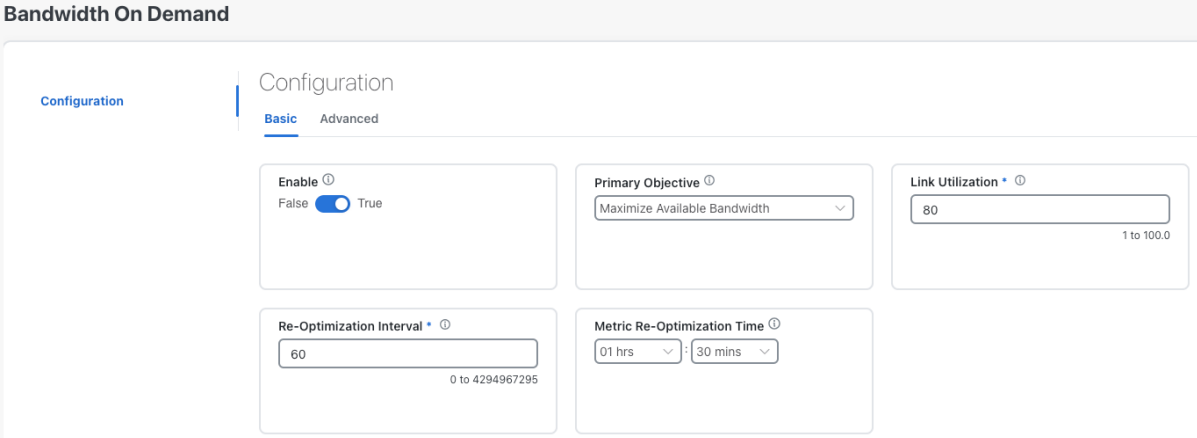
Procedure to enable and configure BWoD

Step 1 Go to **Services & Traffic Engineering > Bandwidth on Demand**.

Step 2 Toggle the Enable switch to True, and enter 80 to set the utilization threshold percentage. To find descriptions of other options, hover the mouse over.

Step 3 Click **Commit Changes**.

Bandwidth On Demand



Configuration

Basic Advanced

Enable [ⓘ]
False True

Primary Objective [ⓘ]
Maximize Available Bandwidth

Link Utilization [ⓘ]
80
1 to 100.0

Re-Optimization Interval [ⓘ]
60
0 to 4294967295

Metric Re-Optimization Time [ⓘ]
01 hrs : 30 mins

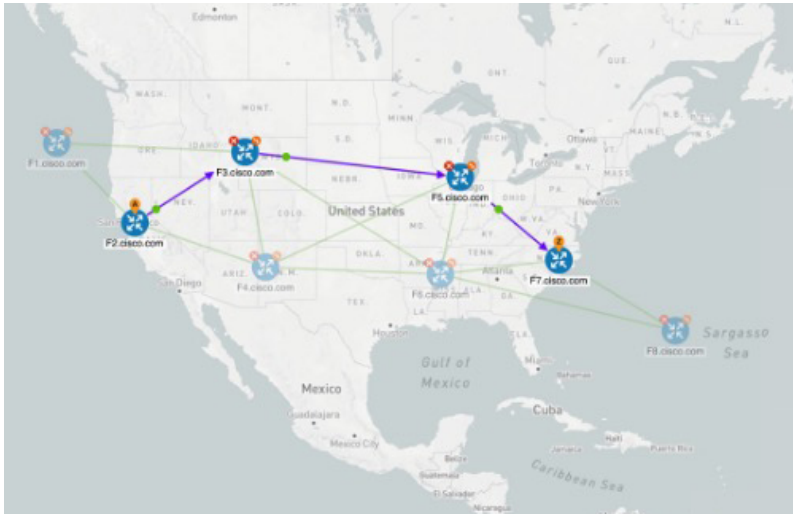
Step 3 Verify that the policy's operational state is now Up and view the path on the map

Procedure to verify that the policy's operational state is now Up and view the path on the map

Step 1 Go to **Services & Traffic Engineering > Provisioning (NSO)**.

Step 2 In the Policy table, locate and select the path computed for the endpoints.

Step 3 The path is shown as an overlay on the map. Select the **IGP Path** check box to see the physical path between the endpoints.



Summary and Conclusion

Operators can set and maintain bandwidth requirements based on optimization intent using the BWoD functionality provided in Cisco Crosswork Network Controller. This scenario illustrated how to provision an SR-TE policy with a specific bandwidth request. We saw how to enable BWoD functionality so that traffic is rerouted automatically to maintain bandwidth requirements. This automation alleviates the task of manually tracking and configuring paths to accommodate bandwidth requirements set by SLAs.



CHAPTER 4

Bandwidth and Network Optimization

This section explains the following topics:

- [Overview, on page 89](#)
- [Scenario: Use CS-SR Policies to Reserve Bandwidth, on page 103](#)

Overview

Objective

Network operators need a toolset to help automate bandwidth optimization, steer traffic with little operator intervention, and ensure that critical links always have sufficient bandwidth to avoid congestion.

Challenge

For service providers, managing bandwidth problems used to be a reactive and manual process. Pressure to solve it is huge. Network congestion leads to poor end-customer experiences. Congested links, high latency, and other network impairments lead to a poor perception of the services carried across your network or result in an inability to meet the service level agreements (SLAs) you have with your customers. In the worst-case scenario, your network issues lead to SLA or contract violations and the loss of your brand equity.

Solution

Using LCM and Circuit Style policies, SPs can now specify business-critical links with the intention to reserve bandwidth for these links. Identifying critical links and the operator's intention enables automatic optimization of the network in real time.

Cisco Crosswork Network Controller offers both:

- Local Congestion Mitigation (LCM) is a tactical solution for bandwidth management and congestion mitigation. It is best applied when you are attempting to solve congestion issues directly, on the devices themselves, without a full-scale traffic matrix or advanced planning.
- Circuit-Style Segment Routing (CS-SR) is a strategic traffic engineering solution that permits you to reserve bandwidth in advance for critical links, avoiding congestion issues entirely for these high-priority links.

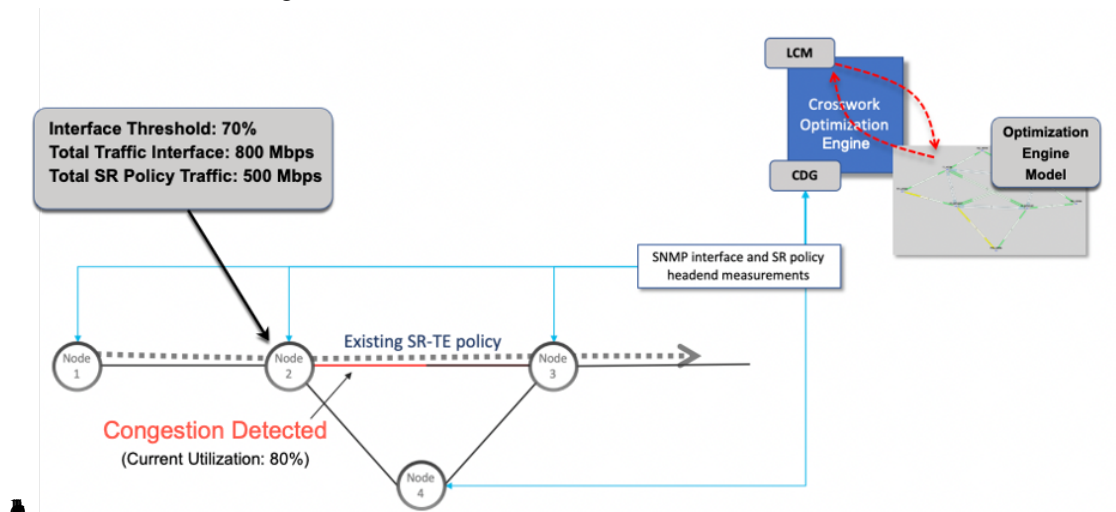
Local Congestion Mitigation (LCM)

Instead of optimizing for bandwidth resource in the network by rerouting traffic in the entire network (end-to-end path optimization), LCM checks the capacity locally, in and around the congested area, at an interface level and reroutes traffic between the endpoints of the congested interface (local interface-level optimization). Focusing on an issue locally eliminates the need for simulating edge-to-edge traffic flows in the network through a full traffic matrix, which is both cumbersome to create and is less scalable as node counts continue to increase.

When congestion is detected in the network, LCM provides recommendations to divert the minimum amount of traffic away from the congested interface. LCM performs the collection of SR-TE policy and interface counters through SNMP. It estimates the amount of traffic that may be diverted and, if the user approves, performs the mitigation through the deployment of Tactical Traffic Engineering (TTE) SR-TE policies. Mitigating congestion locally does not require the use of the full Segment Routing Traffic Matrix (SR-TM). TTE SR-TE policies are created at the device on only either side of the congested link, with the shortest paths possible that do not congest interfaces elsewhere.

How Does LCM Work?

1. First, network operators create domains that define "local" portions of the network. A domain can be the entire network, but more commonly a domain will match one or geographical areas or groups of device interfaces. In this example, we have defined a domain with four devices and all their interfaces. We also assume that all the links in this domain are 1Gbps.
2. Operator specifies a threshold defining what "congestion" means for a particular domain. In this example, the operator has set the domain's congestion threshold to 70%. The congestion threshold you decide on may vary. For guidance on how to determine what's congestion threshold is best for your network and its domain architecture, see [Cisco's Local Congestion Mitigation \(LCM\) White Paper](#).
3. LCM first analyzes the Optimization Engine Model (a realtime representation of the physical network, its topology and its traffic) on a regular cadence. After a congestion check interval, LCM detects congestion when Node 2 utilization goes above the 70% utilization



- 4.
5. LCM calculates how much traffic is eligible to divert. LCM will follow these rules and restrictions in its recommendations:

LCM only diverts traffic that is not already routed by an existing SR policy (for example: unlabeled, IGP-routed, or carried via FlexAlgo-0 SIDs). The traffic within an SR policy will not be included in LCM calculation and will continue to travel over the original programmed path.

LCM computes diversion-eligible traffic by taking the interface traffic statistics that account for all traffic on the interface and subtracting the sum of traffic statistics for all SR-TE policies that flow over the interface.

Total interface traffic – SR policy traffic = Eligible traffic that can be optimized

This process must account for any ECMP splitting of SR policies to ensure the proper accounting of SR policy traffic. In this example, the total traffic on congested Node 2 is 800 Mbps. The total traffic of all SR policies routed over Node 2 is 500 Mbps.

The total traffic that LCM can divert in this example is 240 Mbps. That is: 800 Mbps – 560 Mbps = 240 Mbps

6. LCM calculates the amount of traffic that must be sent over alternate paths by subtracting the threshold-equivalent traffic from the total traffic on the interface. In this example, the amount to be diverted is 100 Mbps:

$800 \text{ Mbps} - 640 \text{ Mbps (70\% threshold)} = 100 \text{ Mbps}$


LCM must route 100 Mbps of 300 Mbps (eligible traffic) to another path.

7. LCM determines how many TTE SR policies are needed and their paths. The ratio of how much LCM eligible traffic can stay on the shortest path to the amount that must be rerouted, will determine the number of TTE SR policies that are needed on the shortest versus alternate paths, respectively.

In this example, LCM needs to divert one-third of the total eligible traffic (100 Mbps out of 300 Mbps) away from the congested link. Assuming a perfect ECMP, LCM estimates that three tactical SR-TE policies are required to create this traffic split: one tactical SR-TE policy will take the diversion path and two tactical SR-TE policies will take the original path. There is sufficient capacity in the path between Node 2 and Node 4. Therefore, LCM recommends three TTE SR policies (each expected to route approximately 100 Mbps) to be deployed from Node 2 to Node 3 via SR-PCE:

- 2 TTE SR policies to take a direct path to Node 3 (200 Mbps)
- 1 TTE SR policy takes a path via Node 4 (100 Mbps)

These recommendations will be listed in the **LCM Operational Dashboard**.



Assuming you deploy these TTE SR policies, LCM continues to monitor the deployed TTE policies and will recommend modifications or deletions as needed in the LCM **Operational Dashboard**. TTE SR policy removal recommendations will occur if the mitigated interface would not be congested if these policies were removed (minus a hold margin). This helps to avoid unnecessary TTE SR policy churn throughout the LCM operation.

Circuit-Style Policies

Circuit-Style Segment Routing Policies (CS-SR, or CS policies) are connection-oriented transport services that you can use to implement what are sometimes referred to as "circuit emulations" or "private lines". Combining segment-routing architecture's adjacency SIDs with stateful PCEP path computation, CS policies provide:

- Persistent, dedicated, bi-directional, co-routed transport paths with predictable latencies and other performance metrics in both directions.
- Guaranteed bandwidth commitments for traffic-engineered services using these paths.
- End-to-end path protection to ensure there is no impact on Service Level Agreements.
- Automatic monitoring, maintenance and restoration of path integrity.
- Flexible operations, administration and management of Circuit-Style paths.
- A software-defined replacement for older CEM infrastructure, such as SONET/SDH.

How Do Circuit-Style Policies Work?

Initial configuration of CS policies follows these steps:

1. Crosswork Network Controller and its applications discover and map the network topology.
2. Crosswork users enable CS policy support, specifying the base bandwidth to be allocated to CS policies as a whole, and a threshold percentage of bandwidth usage which, when exceeded on any CS-calculated path, will generate an alarm. So, for example, on a 1 GB link with 20 percent of bandwidth reserved for Circuit Style use, CS policies can use up to 200 Mbps of that link. Note, however, that if the bandwidth

minimum threshold is set to the default of 80 percent, alarms will be generated as soon as 160 Mbps of the link is used.

3. Network operators create a CS policy for each set of nodes for which they want to establish a guaranteed path. The policy specifies the two nodes to be linked by the main path, the bandwidth to be reserved, and the backup path. To ensure bandwidth and path failures can be accommodated, the configuration must include bi-directionality, path protection, and performance-management liveness-detection settings.
4. When the operator commits the CS policy, the device-resident Path Computation Client (PCC) will request the Crosswork-resident PCE server to compute candidate Working and Protected paths that conform to the CS policy's bandwidth and other constraints (using a single PCEP request message).
5. The PCC computes both paths and deducts the CS policy-guaranteed bandwidth for them from the total available bandwidth allocated when CS policy support was enabled.
6. Crosswork replies to the PCC with the primary Working and Protected path lists and commits to, or "delegates", them. The topology map displays the current Active and Protected paths between the two nodes, using the colors configured when the CS policy was configured, and labels the two endpoint nodes so they can be identified as CS policy endpoints.


After the initial configuration:

1. Crosswork monitors the delegated path and the active CS policies. It updates the available and reservable bandwidth in the network in near real time.
2. Crosswork generates threshold-crossing alarms when bandwidth usage or additional CS policy requirements exceed the configured reserved bandwidth or bandwidth usage threshold.
3. If delegated paths fail for any reason, Crosswork recomputes paths as needed.

Scenario: Use LCM to Reroute Traffic on an Overused Link

In this scenario, we will enable Local Congestion Mitigation (LCM) and observe its congestion mitigation recommendations. LCM will recommend that we deploy Tactical Traffic Engineering Segment Routing (TTE SR) policies on a device's interfaces when usage exceeds a defined threshold. We will preview the recommended TTE SR policies before committing them.

This example uses the following topology:



Note If you are viewing the HTML version of this guide, click on the images to view them in full-size.

We will enable LCM with a configuration that results in the link between **cw-xrv60** and **cw-xrv60** becoming over-used. We will then review the mitigation solutions Crosswork calculates. In this example, it is left to the operator whether to apply the solution or not.

LCM Scenario: Assumptions and Prerequisites

The following sections list high-level requirements that must be met to ensure proper LCM operation.

Congestion Evaluation Requirements

LCM requires traffic statistics from the following:

- Interface traffic measurements
- Headend SR-TE policy traffic measurements

To ensure LCM is receiving these traffic statistics:

- Enable SNMP on the devices whose traffic you want to monitor, including the headend device. For more on this task, see [Configuring SNMP Support](#). Note that gNMI is also an option for collecting traffic measurements.
- Ensure that the SNMP-enabled devices are all reachable from the Crosswork Data Gateway. For more on this task, see [Check Connectivity to the Destination](#).
- Configure the headend device to use strict SID labels for SR policies. To perform this task:
 1. Enable segment routing on the headend device and configure the segment routing global block (SRGB) and the segment routing local block (SRLB) ranges. For example:

```
segment-routing
mpls
  global-block 16000 23999
  node-msd 16
```

```
!
srlb 15000 15999
```

2. Configure the SR policy candidate paths to use strict SID labels. You can use either explicit paths or dynamic paths with constraints. For example:

```
segment-routing
traffic-eng
policy COLOR-100-TO-10.0.0.1
color 100 end-point ipv4 10.0.0.1
candidate-paths
preference 100
explicit segment-list SL1
!
preference 200
dynamic
constraints
affinity include-any RED BLUE
sid-algorithm strict-spf
!
!
!
!
!
!
segment-list SL1
index 10 mpls label 16001 node 10.0.0.2 strict
index 20 mpls label 16002 node 10.0.0.3 strict
index 30 mpls label 16003 node 10.0.0.4 strict
!
```

3. Configure the SR policy headend behavior using the binding SID and the autoroute announce option. For example:

```
!segment-routing
traffic-eng
pcc
profile 1
autoroute
include ipv4 all
force-sr-include
!
!
!
!
```

Congestion Mitigation Requirements

The headend device must support PCE-initiated SR-TE policies with autoroute steering. However, LCM will not work if the headend is a Cisco NCS device and there is L2VPN traffic in the network.

Devices should be configured with `force-sr-include` to enable traffic steering into SR-TE policies with autoroute. For example:

```
segment-routing traffic-eng pcc profile ID autoroute force-sr-include
```

The `ID` parameter in this command identifies the PCC profile associated with the SR-TE policy that PCE has provisioned. The ID value can be any integer from 1 to 65535, but it must match the profile ID that PCE uses to instantiate the policy. If not, the policy will not be activated. For example, if PCE provisions a policy with profile ID 10, you must configure `segment-routing traffic-eng pcc profile 10 autoroute`

`force-sr-include` on the headend router to enable autoroute announcement for that policy. For more information, see the [Segment Routing Configuration Guide, Cisco IOS XE 17 \(Cisco ASR 920 Series\), COE-PCE Initiated SR Policy with OSPF and IS-IS SR-TE Autoroute Announce](#).



Note The ID that is configured under the PCC profile, must match the Profile ID option set in the LCM Configuration page.

The headend device must support Equal Cost Multi-Path (ECMP) across multiple parallel SR-TE policies. To verify that a device can support SR-TE policies using ECMP, check that the device has the following:

- Segment Routing is enabled and configured, with a Segment Routing Global Block (SRGB) that matches the SRGB of the SR-TE policy headend and tailend routers. Use the `show segment-routing mpls state` command to verify the SRGB configuration on the device.
- BGP-LS is enabled and configured to advertise and receive link-state information from the SR-TE policy headend and tailend routers. Use the `show bgp link-state link-state` command to verify the BGP-LS status and the `show bgp link-state link-state database` command to verify the link-state information on the device.
- ECMP is enabled and configured to load-balance traffic across multiple equal-cost paths based on flows. Use the `show ip route` command to verify the ECMP routes and the `show ip cef` command to verify the ECMP load-balancing algorithm on the device.

If all these conditions are met, then the device can support an SR-TE policy using ECMP.

Related Topics

For more information and examples on how to configure and verify SR-TE policies, see:

- [Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers](#)
- [Segment Routing Configuration Guide, Cisco IOS XE 17 | Access and Edge Routers](#)

LCM Scenario: Workflow

Workflow steps	Detailed procedure links
Step 1. Enable LCM and configure the global utilization thresholds	Step 1: Enable LCM and Configure the Utilization Thresholds, on page 97
Step 2. View link congestion on the map	Step 2: View Link Congestion on the Map, on page 99
Step 3. View TTE SR policy recommendations in the LCM Operational Dashboard	Step 3: View TTE SR Policy Recommendations in the LCM Operational Dashboard, on page 99
Step 4. Validate the TTE SR policy deployment	Step 4: Validate TTE SR Policy Deployment, on page 101
Step 5. Remove the TTE SR policies upon LCM recommendation	Step 5: Remove TTE SR policies on LCM Recommendation, on page 102

Step 1: Enable LCM and Configure the Utilization Thresholds

To enable LCM and configure the global utilization threshold:

Step 1 Go to **Services & Traffic Engineering > Local Congestion Mitigation > Domain-ID** and click **Configure**.

Step 2 Toggle the **Enable** switch to **True**, and enter the global utilization threshold you want to set. In this case, we set the threshold at 80%, and select the **Interfaces to Monitor > All Interfaces** option. In the **Advanced** tab, Operation mode is set to **Manual**. Manual mode allows you to view recommended TTE policies prior and decide whether or not to deploy them. To see information about other options for each configuration setting, hover the mouse over **i** (help icon).

Figure 1: Basic LCM Configuration

Configuration

Basic Advanced

<p>Enable ⓘ</p> <p>False <input checked="" type="checkbox"/> True</p>	<p>Color * ⓘ</p> <p>2000</p> <p>Range: 1 to 4294967294</p>	<p>Utilization Threshold * ⓘ</p> <p>80</p> <p>Range: 0 to 100</p>
<p>Utilization Hold Margin * ⓘ</p> <p>5</p> <p>Range: 0 to Utilization Threshold</p>	<p>Delete Tactical SR Policies when Disabled ⓘ</p> <p>False <input type="checkbox"/> True</p>	<p>Profile ID * ⓘ</p> <p>0</p> <p>Range: 0 to 65534</p>
<p>Congestion Check Interval * ⓘ</p> <p>900 seconds</p> <p>Range: 60 to 86400 seconds</p>	<p>Max LCM Policies per Set * ⓘ</p> <p>8</p> <p>Range: 1 to 8</p>	<p>Interfaces to Monitor ⓘ</p> <p><input type="radio"/> Selected Interfaces <input checked="" type="radio"/> All Interfaces</p>
<p>Description ⓘ</p> <p>LCM Startup Config</p>		

Step 1: Enable LCM and Configure the Utilization Thresholds

Figure 2: Advanced LCM Configuration (Manual Mode)

Configuration

Basic **Advanced**

Auto Repair Solution ⓘ

False True

Stay in Area ⓘ

False True

Adjacency Hop Type ⓘ

Unprotected

Operation Mode ⓘ

Manual Automated

Optimization Objective ⓘ

Minimize the IGP metric

Deployment Timeout * ⓘ

180

Range: 10 to 300

Congestion Check Suspension Interval * ⓘ

600

Range: 600 to 3600

Over-Provisioning Factor * ⓘ

Range: 0.0 to 100.0

Uneven ECMP Traffic Threshold * ⓘ

Range: 0.0 to 100.0

Throttle Mode Threshold * ⓘ

5

Range: 0 to 10

Step 3 Click **Commit Changes**.

Note After committing the configuration changes, LCM will display *recommendations* on the **LCM Operational Dashboard** if congestion occurs on any monitored interfaces. LCM will *not* commit or deploy new TTE policies automatically. Later, you will be able to preview the recommended TTE policies and decide whether or not to commit and deploy them onto your network.

Step 4 You can also define individual interface thresholds. Go to the **Customized Interface Thresholds** page (**Traffic Engineering > Local Congestion Mitigation > Domain-ID > ... > Interface Thresholds**).

See the following example and note the defined threshold for cw-xrv60 with interface GigabitEthernet0/0/0/1 is 16%.

Note The utilization thresholds used in this example are extremely low and are best used for lab environments.

Figure 3: Customized Interface Thresholds

Customized Interface Thresholds

Interfaces to Monitor: Selected Interfaces - LCM monitors only the interfaces with custom thresholds.

+ Create | Edit Mode: OFF

Node	Interface	Threshold (%)	Select for
cw-xrv60	GigabitEthernet0/0/0/1	16	<input type="checkbox"/>

Step 2: View Link Congestion on the Map

The link between **cw-xrv60** and **cw-xrv62** is now congested. Let's see that on the map.

Step 1 Go to **Services & Traffic Engineering > Traffic Engineering**.

Step 2 Click on the link to view link details, including utilization information. Usage has surpassed the custom LCM threshold defined at 16% for node **cw-xrv60** with interface GigabitEthernet0/0/0/1.

The screenshot displays a network map on the left and a 'Link Details' panel on the right. The map shows several nodes (cw-xrv57, cw-xrv58, cw-xrv59, cw-xrv60, cw-xrv61, cw-xrv62, cw-ncs8, cw-ssr903) connected by lines. A red dashed line highlights the link between cw-xrv60 and cw-xrv62. The 'Link Details' panel is divided into 'Summary' and 'Traffic Engineering' tabs. The 'Summary' tab shows the link name, state (Up), link type (L3 ISIS IPv4), ISIS level (2), and last update time. The 'Traffic Engineering' tab shows a table of link details for both sides (A Side and Z Side). The A Side details for cw-xrv60 show a utilization of 38.35% (383.5Mbps/1Gbps), which is highlighted in red. The Z Side details for cw-xrv62 show a utilization of 0% (1.44Kbps/1Gbps). A second table below the main details shows the A Side details for cw-xrv60, also highlighting the 38.35% utilization.

	A Side	Z Side
Node	cw-xrv60	cw-xrv62
TE Router ID	3.3.3.60	3.3.3.62
IPv6 Router ID	bb:bb:bb:3:3::60	bb:bb:bb:3:3::62
IF Name	GigabitEthernet0/0/0/1	GigabitEthernet0/0/0/1
IF Description		GigabitEthernet0/0/0/1
IF Alias		*** connect to xrv60 ***
Type		ETHERNETCSMACD
Local IF ID	3.3.3.60 (0)	3.3.3.62 (7)
Utilization	38.35% (383.5Mbps/1Gbps)	0% (1.44Kbps/1Gbps)
Packet Drops	0%	0%
Delay	0 μs	
Jitter	0 μs	
IGP Metric	1	
Delay Metric	1	
TE Metric	1	
Admin Groups		

	A Side
Node	cw-xrv60
TE Router ID	3.3.3.60
IPv6 Router ID	bb:bb:bb:3:3::60
IF Name	GigabitEthernet0/0/0/1
IF Description	
IF Alias	
Type	
Local IF ID	3.3.3.60 (0)
Utilization	38.35% (383.5Mbps/1Gbps)

Step 3: View TTE SR Policy Recommendations in the LCM Operational Dashboard

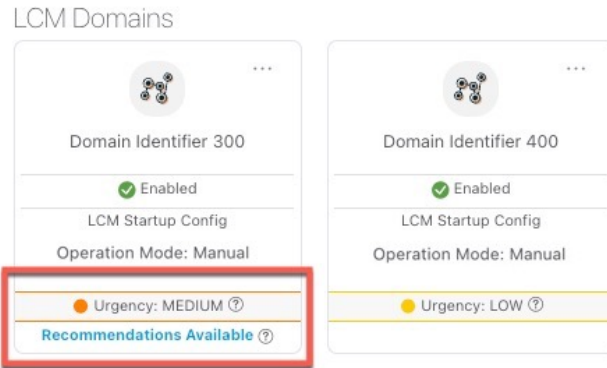
LCM has detected the congestion and computed tactical policies to mitigate the congestion, which we can preview and then decide whether or not to commit them.

Note that, in this scenario, the congested device is healthy, reachable and in sync with Crosswork. The actions we take and policies we implement will be different if, in addition to congestion, the device is down, unreachable or out of sync.

Step 1 Go to **Services & Traffic Engineering > Local Congestion Mitigation**.

Step 3: View TTE SR Policy Recommendations in the LCM Operational Dashboard

When congestion is detected, the domain displays the urgency type and recommendations that are available. Click the question mark icons to display more information about the urgency type and when the most recent recommendation was given.



Step 2 Open the Operational Dashboard (**Services & Traffic Engineering > Local Congestion Mitigation > Domain-ID > ... > Operational Dashboard**).


The dashboard shows that cw-xrv60 utilization has surpassed 16% and is now at 38.5%. In the **Recommended Action** column, LCM recommends the deployment of TTE policy solution sets (**Recommended Action - Update Set**) to address the congestion on the interface.

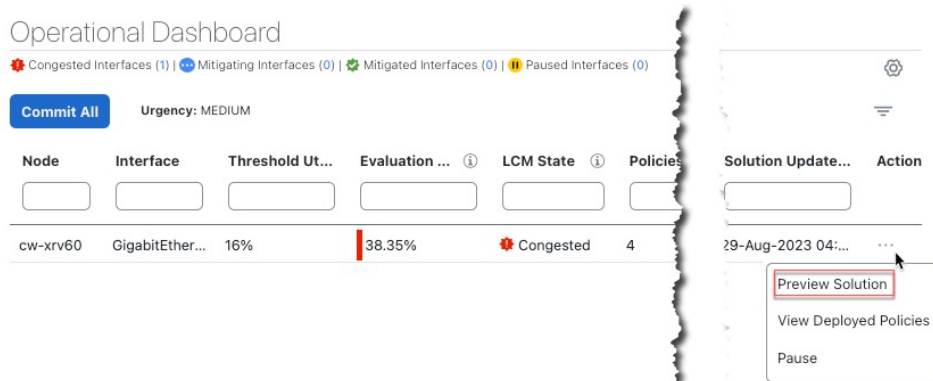
Operational Dashboard

Congested Interfaces (1) | Mitigating Interfaces (0) | Mitigated Interfaces (0) | Paused Interfaces (0)

Commit All Urgency: MEDIUM

Node	Interface	Threshold Ut...	Evaluation ...	LCM State	Policies De...	Policy Set ...	Recommende...	Commit ...	Expected Util...	Solution Update...
cw-xrv60	GigabitEther...	16%	38.35%	Congested	4	DEGRADED	Update Set	None	14%	29-Aug-2023 04:...

Step 3 Before committing TTE policies, you can preview the deployment of each TTE policy solution set. Click  in the **Actions** column and choose **Preview Solution**.



The resulting window displays the node, interface, and the recommended action for each TTE policy. From the Preview window, you can select the individual TTE policies, and view different aspects and information as you would normally do in the topology map. You can expand each policy to view individual segments. After reviewing the potential implications on your network, you can decide whether or not to deploy the bypass policies that LCM recommends.

The following figure shows the recommended TTE policies for node cw-xrv60.

The screenshot displays a map of North America with several network nodes marked. A panel titled 'Preview Recommended TTE Policies' is open, showing details for Node 'cw-xrv60' and Interface 'GigabitEthernet0/0/0/1'. It lists two candidate paths:

Headend	Endpoint	Color	Recommended Action
<input checked="" type="checkbox"/> cw-xrv60	cw-xrv62	2000	UPDATE

Seg...	Segme Type	Label	Algo	IP	No...	Interf...
0	Node SID	16562	1	3.3.3.62	cw...	

Se...	Segme...	La...	Algo	IP	N...	Interf...	Sl...
0	IGP ...	24...	0	12.1.14...	cw...	GigabitEthe	
1	Nod...	165...	1	3.3.3.57	cw...		Stri...
2	Nod...	165...	1	3.3.3.62	cw...		Stri...

A 'Back To LCM Dashboard' button is visible at the bottom of the panel.

Step 4 After you are done viewing the recommended TTE policies on the map, go back to the **Operational Dashboard** and click **Commit All**. The LCM **State** column changes to **Mitigating**.

All LCM recommendations per domain must be committed in order to mitigate congestion and produce the expected utilization as shown in the **Operational Dashboard**. The mitigating solution is based on *all* LCM recommendations being committed because of dependencies between solution sets.

The screenshot shows the 'Operational Dashboard' with the following details:

- LCM Domain Identifier: 300
- Operational State: Enabled
- Description: LCM Startup Config
- Mode: Manual


Summary: Congested Interfaces (0) | Mitigating Interfaces (1) | Mitigated Interfaces (0) | Paused Interfaces (0)

Urgency: LOW

Node	Interface	Threshold Util...	Evaluation UL...	LCM State	Policies Depl...	Policy Set S...	Recommended...	Commit St...
cw-xrv60	GigabitEther...	16%	31.01%	Mitigating	2	-	No Change	CONFIRMED

Step 4: Validate TTE SR Policy Deployment

To validate the TTE SR policy deployment, follow the steps given below:

Step 1 With the **Operational Dashboard** displayed, click the  at the top right of the user interface to open the **Alarms** window, then select the **Events** tab. You can use these two tabs to monitor LCM alarms and events. The **Events** shows you events for the LCM recommendations, the commit actions, as well as any exceptions.


Crosswork will report network events that are detected based on the policies and features you have enabled. For example, if a link drop causes an SR-TE policy to go down, or if LCM detects congestion, an event is displayed in the UI.

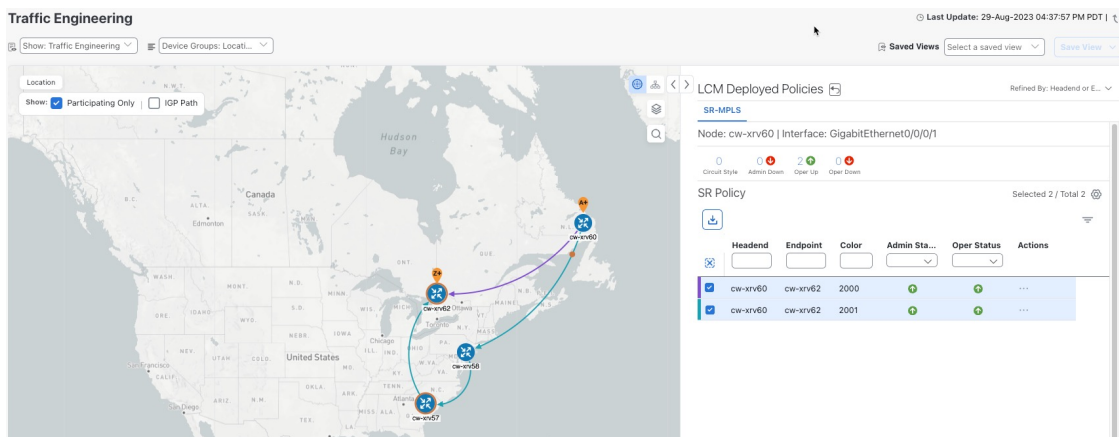
Step 2 Return to the **Operational Dashboard** to see that the LCM state changes to **Mitigated** for all TTE policy solution sets.

Note The LCM state change can take up to twice as much time as the SNMP cadence.

Step 5: Remove TTE SR policies on LCM Recommendation

Step 3


Confirm the TTE policy deployment by viewing the topology map. Click  in the **Actions** column and choose **View Deployed Policies**. The deployed policies are displayed in focus within the topology map. All other policies are dimmed.

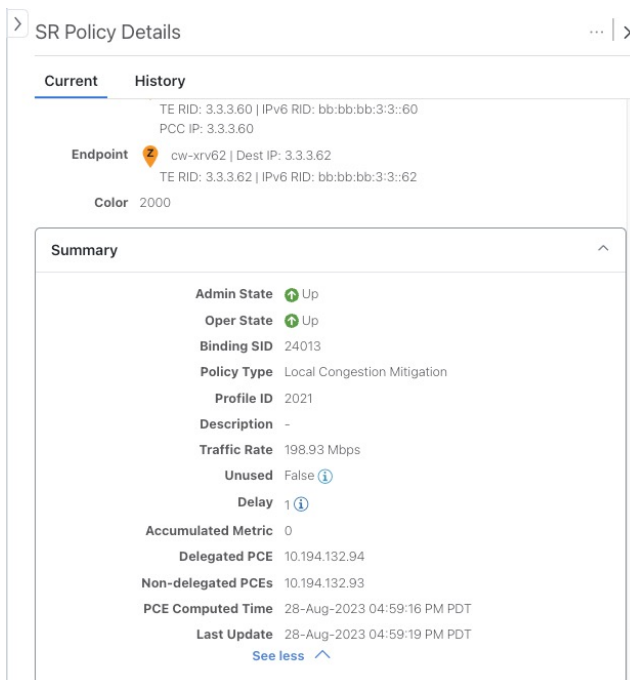


The screenshot shows the Traffic Engineering interface. On the left, a map of the United States displays network nodes (cw-xrv60, cw-xrv62, cw-xrv58, cw-xrv57) and their connections. On the right, the 'LCM Deployed Policies' panel is open, showing a table of SR Policies:

Headend	Endpoint	Color	Admin Sta...	Oper Status	Actions
cw-xrv60	cw-xrv62	2000		Up	...
cw-xrv60	cw-xrv62	2001		Up	...

Step 4

View the SR policy details. From the **Actions** column of one of the deployed policies, click  and choose **View Details**. Note that the **Policy Type** is **Local Congestion Mitigation**.



The screenshot shows the 'SR Policy Details' view. The 'Current' tab is active, displaying the following information:

- TE RID: 3.3.3.60 | IPv6 RID: bb:bb:bb:3:3:60
- PCC IP: 3.3.3.60
- Endpoint: cw-xrv62 | Dest IP: 3.3.3.62
- TE RID: 3.3.3.62 | IPv6 RID: bb:bb:bb:3:3:62
- Color: 2000


The 'Summary' section provides additional details:

- Admin State: Up
- Oper State: Up
- Binding SID: 24013
- Policy Type: Local Congestion Mitigation
- Profile ID: 2021
- Description: -
- Traffic Rate: 198.93 Mbps
- Unused: False
- Delay: 1
- Accumulated Metric: 0
- Delegated PCE: 10.194.132.94
- Non-delegated PCEs: 10.194.132.93
- PCE Computed Time: 28-Aug-2023 04:59:16 PM PDT
- Last Update: 28-Aug-2023 04:59:19 PM PDT

Step 5: Remove TTE SR policies on LCM Recommendation

After some time, the deployed TTE SR policies may no longer be needed. This occurs if utilization continues to stay under threshold without the LCM-initiated TTE policies. If this is the case, LCM generates new recommended actions to delete the TTE SR policy sets.

To remove the TTE SR policies upon LCM recommendation, follow the steps given below:

- Step 1** If needed: Display the topology map and click  in the **Actions** column. Choose **View Deployed Policies**.
- Step 2** Click **Commit All** to remove the previously deployed TTE SR policies.
- Step 3** Confirm the removal by viewing the topology map and SR Policy table.

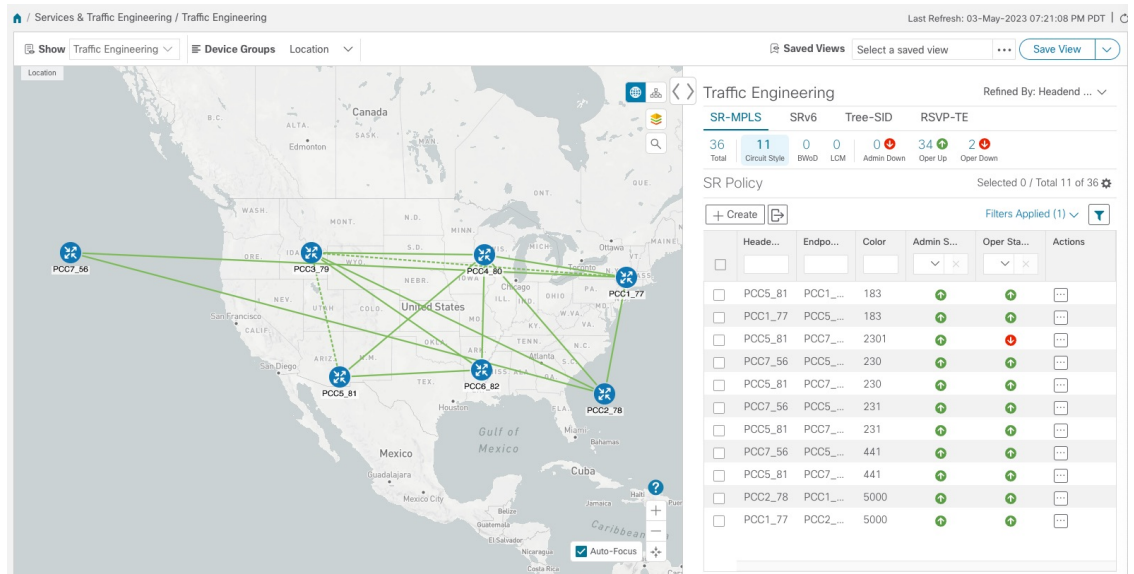
LCM Scenario: Summary and Conclusion

In this scenario, we observed how to leverage LCM to alleviate traffic congestion in the network. LCM takes the manual tracking and calculation out of your hands but at the same time gives you control as to whether to implement the congestion mitigation recommendations, or not. You can preview the recommendations and see how the potential deployment will take effect in your network before you deploy them. As traffic changes, LCM tracks the deployed TTE SR policies and decides whether or not they are still needed. If not, LCM recommends deleting them.

Scenario: Use CS-SR Policies to Reserve Bandwidth

In this scenario, we enable Circuit-Style Segment Routing Traffic Engineering (CS-SR, or CS SR-TE) policies and set bandwidth-reservation parameters, then configure a CS-SR policy and visualize it on the topology map. We will inspect the policy's details, including its computed Active (working) and Protected (protect) paths.

The examples in this scenario use the following topology:



The screenshot displays the Cisco Network Controller interface for Traffic Engineering. On the left, a map shows a network topology with nodes labeled PCC7_56, PCC3_79, PCC4_80, PCC1_77, PCC5_81, PCC6_82, and PCC2_78. On the right, the 'Traffic Engineering' section shows a summary of 36 total policies, with 11 in Circuit Style, 0 BWD, 0 LCM, 0 Admin Down, 34 Oper Up, and 2 Oper Down. Below this is an 'SR Policy' table with 11 rows of data.

Head...	Endpo...	Color	Admin S...	Oper Sta...	Actions
<input type="checkbox"/>	PCC5_81	PCC1_...	183		
<input type="checkbox"/>	PCC1_77	PCC5_...	183		
<input type="checkbox"/>	PCC5_81	PCC7_...	2301		
<input type="checkbox"/>	PCC7_56	PCC5_...	230		
<input type="checkbox"/>	PCC5_81	PCC7_...	230		
<input type="checkbox"/>	PCC7_56	PCC5_...	231		
<input type="checkbox"/>	PCC5_81	PCC7_...	231		
<input type="checkbox"/>	PCC7_56	PCC5_...	441		
<input type="checkbox"/>	PCC5_81	PCC7_...	441		
<input type="checkbox"/>	PCC2_78	PCC1_...	5000		
<input type="checkbox"/>	PCC1_77	PCC2_...	5000		

We will observe what happens when the Active bandwidth-reserved path between the NCS1 and NCS3 nodes fails. We will then re-optimize the failed path.

CS-SR Scenario: Assumptions and Prerequisites

The following sections provide a list of high-level requirements for proper CS-SR operation, including requirements and constraints on the policy attribute values set in each CS-SR policy, and the processing logic followed during path reversions.

In addition to the constraints discussed in the following sections:

- The Crosswork Circuit Style Manager (CSM) feature pack is a feature of the Crosswork Network Automation Essential Suite. All licensed features are available during the 90-day trial period. After the trial period, you must have a license for Crosswork Optimization Engine to continue using CSM.
- Circuit-Style policy configuration was introduced with Crosswork Network Controller (CNC) 5.0. To use it, you must have version 7.9.1 (or later) of the Cisco IOS-XR Path Computation Client (PCC) installed on your devices. If you have been using a previous version of CNC with IOS-XR version 7.7.1 or earlier, please upgrade to version 7.9.1 or later before attempting to configure CS-SR policies.
- When using CSM with Crosswork Network Controller, the UI navigation starts from **Traffic Engineering & Services**. When using CSM with Crosswork Optimization Engine, the navigation starts from **Traffic Engineering**.

CS Policy Attribute Constraints

In this scenario, we will build a CS policy between node NCS1 and node NCS 3. The policy will use the following settings and constraints:

- **PolicyName:** NCS1-NCS3
- **Headend Device:** NCS1
- **Headend IP Address:** 192.168.20.4
- **Tailend Device:** NCS3
- **Tailend IP Address:** 192.168.20.14
- **Color-choice:** 1000
- **Bandwidth:** 10000
- **path-protection:** Enabled
- **disjoint-path:** Enabled
- **disjoint-path forward-path type:** Link
- **disjoint-path forward-path group-id:** 531
- **disjoint-path reverse-path type:** Link
- **disjoint-path reverse-path group-id:** 5311
- **performance-measurement :** Enabled.
- **performance-measurement profile-type:** Liveness
- **performance-measurement liveness-detection:** Enabled
- **performance-measurement profile:** CS-active

- **working-path**: Enabled
- **working-path preference**: 100
- **working-path dynamic-path**: Enabled
- **working-path dynamic-path pce**: Enabled
- **working-path dynamic-path metric type**: **igp**
- **working-path dynamic-path bidirectional-association-choice**: Enabled
- **working-path dynamic-path bidirectional-association-id**: 230
- **working path dynamic constraints segments**: Enabled
- **working-path constraints segments protection**: **unprotected-only**
- **protect-path**: Enabled
- **protect-path preference**: 100
- **protect-path dynamic-path**: Enabled
- **protect-path dynamic-path pce**: Enabled
- **protect-path dynamic-path metric type**: **igp**
- **protect-path dynamic-path bidirectional-association-choice**: Enabled
- **protect-path dynamic-path bidirectional-association-id**: 231
- **protect-path dynamic constraints segments**: Enabled
- **protect-path constraints segments protection**: **unprotected-only**
- **restore-path**: Enabled
- **restore-path preference**: 100
- **restore-path dynamic-path**: Enabled
- **restore-path dynamic-path pce**: Enabled
- **restore-path dynamic-path metric type**: **igp**
- **restore-path dynamic-path bidirectional-association-choice**: Enabled
- **restore-path dynamic-path bidirectional-association-id**: 232
- **restore-path dynamic constraints segments**: Enabled
- **restore-path constraints segments protection**: **unprotected-only**

The following table shows all of the options you can choose from when building a policy. It is important to understand that the attributes described in the table act as constraints. Each of them corresponds to elements of the configuration that Cisco Crosswork uses to govern how Circuit-Style path hops are computed. Each value is effectively a path computation or optimization constraint, since they either specify a required property of a path or exclude possible choices for that path.

There are dependencies that must be met as well as combinations that are not allowed. The system will warn you when these sorts of issues arise. We encourage you to experiment to learn how to provision services in your network that match the types of services you want to deliver.

Table 1: Supported Circuit Style SR-TE Policy Attribute Values and Constraints

Attribute	Description
Policy Path Protection	The path protection constraint is required for both sides of a Circuit Style SR-TE policy.
Bandwidth Constraint	<p>The bandwidth constraint is required and must be the same on both sides of a Circuit Style SR-TE policy. Bandwidth changes can be made to existing policies, with these effects:</p> <ul style="list-style-type: none"> • Once you configure the new bandwidth on both sides, Crosswork will evaluate the path. This will not result in a recomputed path. • If the new bandwidth is higher, Crosswork checks the existing path to ensure sufficient resources. If all currently delegated paths can accommodate the new bandwidth, Crosswork returns the same path with the new bandwidth value, indicating to the path computation client (PCC) that it was successful. If any of the current paths cannot accommodate the new bandwidth, it returns the old bandwidth value indicating that it was unsuccessful. This evaluation will not be retried unless the bandwidth is changed again. • If the bandwidth is lower, Crosswork returns the same path with the new bandwidth value to indicate to the PCC that it was successful. <p>The user interface shows both the requested and reserved bandwidth under each candidate path when you view the policy details. These values can differ if the requested bandwidth is increased but there is insufficient available CS pool bandwidth along one or more of the paths.</p>
Candidate Paths and Roles	<p>The <code>Working</code> path is defined as the highest preference Candidate Path (CP).</p> <p>The <code>Protect</code> path is defined as the CP with the second highest preference.</p> <p>The <code>Restore</code> path is defined with the lowest preference CP. The headend must have <code>backup-ineligible</code> configured.</p> <p>CPs of the same role in each direction must have the same CP preference.</p>
Bi-Directional	<p>All paths must be configured as co-routed.</p> <p>Paths with the same role on both sides must have the same globally unique bi-directional association ID.</p>

Attribute	Description
Disjointness	<p>Working and Protect paths on the same PCC must be configured with a disjointness constraint using the same disjoint association ID and disjointness type.</p> <p>The disjointness association ID for a Working and Protect path pair in one direction must be unique when compared with the corresponding pair in the opposite direction.</p> <p>Only the <code>Node</code> and <code>Link</code> disjoint types are supported. The disjoint type used must be the same in both directions of the same policy.</p> <p>The Restore path must not have a disjointness constraint set.</p> <p>Crosswork follows strict fallback behavior for all Working and Protect path disjointness computations. This means that, if node type disjointness is configured but no path is available, Crosswork makes no automatic attempt to compute a less restrictive link type disjoint path.</p>
Metric Type	<p>Only the <code>TE</code>, <code>IGP</code> and <code>Latency</code> metric types are supported. The metric type used must match across Working, Protect and Restore paths in both directions.</p>
Segment Constraints	<p>All Working, Protect and Restore paths must have the following segment constraints:</p> <ul style="list-style-type: none"> • <code>protection unprotected-only</code> • <code>adjacency-sid-only</code> <p>To ensure persistence through link failures, configure static adjacency SIDs on all interfaces that might be used by Circuit Style policies.</p>
Supported Policy Changes	<p>The following constraints may be changed for an operationally "up" Circuit Style SR-TE policy that has been previously delegated:</p> <ul style="list-style-type: none"> • Metric type • Disjoint type • MSD • Affinities <p>Once configuration changes are made in a consistent manner across all CPs and both PCCs (for example: the new metric type is the same for all CPs and both sides), Crosswork will initiate a recompute, which can result in new Working, Protect and Restore paths.</p> <p>During any transitory period in which configurations are not in sync between paths on the same PCC or between PCCs, no path updates are sent to the PCCs.</p>

Attribute	Description
Path Computation	<p>Crosswork computes paths for circuit style policies only after a complete bi-directional, path-protected set of candidate paths has been delegated, including Working and Protect paths on both sides.</p> <p>Crosswork computes the Restore path only after the Working and Protect paths are down. The SR Circuit Style Manager feature pack configuration interface provides a configurable delay timer to control how long after Restore paths are delegated from both sides to wait before the path is computed. This delay allows topology and SR policy state changes to fully propagate to Crosswork, in cases where these changes triggered the Restore path delegation.</p> <p>Path computation is supported for Intra/Inter area/level and Intra/Inter IGP Domain (same AS).</p>
Reversion Behavior	<p>Reversion behavior is controlled by the configuration of the WTR lock timer option under the Protect and Revert paths (it is not relevant for the Working path):</p> <ul style="list-style-type: none"> • No lock configuration: Revert after a default 5-minute lock • Lock with no duration specified: No reversion • Lock duration <value>: Revert after the specified number of seconds

Unsupported CS Policy Options

The following table lists the CS policy options, attributes and constraints that are not supported in this version of CSM.

Table 2: Unsupported Circuit Style SR-TE Policy Options

Attribute	Description
Unsupported Configurations	<p>The following configurations are not supported:</p> <ul style="list-style-type: none"> • Metric-bounds • SID-Algo constraints • Partial recovery is not supported with 7.8.x. • State-sync configuration between PCEs of a high-availability pair. These are not required with Circuit Style SR-TE policies. Use of this feature may result in degraded performance. • Multiple Circuit Style SR-TE policies between the same nodes with the same color but different endpoint IPs.

Attribute	Description
Unsupported Policy Changes	<p>The following configuration changes to a previously delegated and operationally "up" Circuit Style SR-TE policy are not supported:</p> <ul style="list-style-type: none"> • CP preference • Disjoint Association ID • Bi-directional Association ID <p>To change these configurations for an existing policy, you must first shut down the policy on both sides, make the change (complying with restrictions as detailed above in terms of consistency) and then "no shut" the policy.</p>
Unsupported Path Computation	Automatic re-optimization is not supported for any paths based on changes in topology, LSP state, or any periodic event. Path computation is not supported for Inter-AS.

Path Reversion Logic

Path reversion depends on the initial state of the Working, Protect and Revert paths and the events affecting each path. The scenarios in the following table provide examples of typical reversion behavior.

Table 3: Path Reversion Scenarios

Initial State	Events	Behavior
Working path is down, Protect path is up/active	Working path comes back up	<ol style="list-style-type: none"> 1. Working path recovers to up/standby state. 2. Each PCC moves the Working path to active after the WTR timer expires. 3. Protect path moves to up/standby.
Working path is down, Protect path is down, Revert path is up/active	Working path comes back up, then Protect path comes back up	<ol style="list-style-type: none"> 1. Working path recovers and goes to up/active state 2. Revert path is removed 3. Protect path recovers and goes to up/standby

Initial State	Events	Behavior
Working path is down, Protect path is down, Revert path is up/active	Protect path comes back up, then Working path comes back up	<p>On side A: The Working path failure is local (the first Adj SID in the SegList is invalid):</p> <ol style="list-style-type: none"> 1. Protect path recovers and goes to up/active. 2. Recover path is removed. 3. Working path recovers and goes to up/standby. 4. Each PCC moves the Working path to active after the WTR timer expires, Protect path goes to up/standby. <p>On side Z: Working path failure is remote (first Adj SID in SegList is valid):</p> <ol style="list-style-type: none"> 1. Protect path recovers but is not brought up, Revert path remains up/active. 2. Working path recovers and goes up/active. 3. Revert path is removed. 4. Protect path goes to up/standby.

What Happens When Path Failures Occur?

Cisco Crosswork computes paths for CS policies only after a complete bidirectional, path-protected set of candidate paths has been delegated. A path can be considered to have "failed" due to a variety of reasons, including transient changes in workloads caused by congestion elsewhere in the network, or any condition that causes the path not to meet bandwidth expectations. Irrespective of the cause, there are three types of paths used during these kinds of failures. Crosswork activates them as needed, according to their preference settings:

- **Working**—This is the path with the highest preference value. Crosswork always tries to keep the operational (Oper Up) path with the *highest* preference as the *Active* path.
- **Protected**—This is the path with the second highest preference. If the Working path goes down, the Protected path (with the lower preference value) is activated. After the Working path recovers, the Protected path remains active until the default lock duration expires, then the Working path is activated.
- **Restore**—This is the path with the lowest preference path. Crosswork computes the Restore path only when the Working and Protected paths are both down. You can control how long after Restore paths are delegated to wait before the path is computed. This delay allows topology and policy state changes to fully propagate through the network and gives Crosswork a chance to gather and analyze telemetry to determine network health.

To address failures effectively and switch from the Working to the Protected path, be sure to configure Performance Measurement (PM) as part of your CS policy. For more information, see [Step 4: Configure Circuit Style SR-TE Policies Using Import, on page 123](#).

The following image shows that the Working and Protected paths of an example CS policy are operational. The *active* path is indicated by the "A" icon shown next to that path in the **State** column in the **Candidate Path** list.


Path Name	Pref	RoleState
> cfg_r1-r2_discr_100	100	Active
> cfg_r1-r2_discr_50	50	Protected

If the Working path performance falls below expectations, the Protected path becomes Active immediately (usually, under 50 milliseconds).

Path Name	Pref	RoleState
> cfg_r1-r2_discr_50	50	Active
> cfg_r1-r2_discr_100	100	Protected

When the Working path comes back up, the Protected path resumes the Protected role again and the Working path (with preference 100) becomes Active again.

If both the Working and Protected paths go down, Crosswork calculates a Restore path and makes it the active path. Note that the Restore path has the lowest preference value of 10. The Restore path only appears in this particular case. If either the Working or Protected paths become operational again, Crosswork will activate them, and the Restore path will disappear from the topology map and from the Candidate Path list.



CS-SR Scenario: Workflow

Workflow steps	Detailed procedure links
1. Enable the SR Circuit Style Manager (CSM) feature pack.	Step 1: Enable Circuit Style Manager, on page 113.
2-4. Configure Circuit Style SR-TE policies on the devices. Note If you haven't enabled the feature pack, the Circuit Style SR-TE policies you configure will appear operationally down.	You can configure Circuit Style SR-TE policies using any of the following methods: <ul style="list-style-type: none"> • On the device, using the CLI. See Step 2: Configure Circuit Style SR-TE Policies Using Device CLI, on page 118 • Using the user interface. See Step 3: Configure Circuit Style SR-TE Policies Using Add, on page 121 • Import a JSON or XML file. See Step 4: Configure Circuit Style SR-TE Policies Using Import, on page 123
5. Verify that the Circuit Style SR-TE policy appears in the SR Policy table and on the topology map.	See Step 5: View Circuit Style SR-TE Policies on the Topology Map, on page 126
6. Verify that the reserved bandwidth pool settings you defined in Step 1 are configured properly.	See Step 6: Verify Circuit Style SR-TE Policy Bandwidth Utilization, on page 130.
7. Trigger path re-computation after path failures.	See Step 7: Trigger Circuit Style SR-TE Path Recomputation, on page 131.

Step 1: Enable Circuit Style Manager

In order to manage and visualize Circuit Style SR-TE policies on the topology map, we must first enable SR Circuit Style Manager (CSM) and set bandwidth reservation settings. As soon as you define these settings, CSM computes the best bidirectional failover paths between the two nodes, while observing the requested CSM bandwidth and threshold settings, and the constraints defined in the Circuit Style SR-TE policy. The following steps show how to do this.

CSM tries to ensure that the total reserved bandwidth on all interfaces remains at or below the network-wide resource pool. When the total usage on all interfaces exceeds the threshold value you set, CSM generates a threshold-crossing alarm.

To help you estimate Circuit Style SR-TE bandwidth pool and threshold settings that are reasonable for your organization's implementation, this topic provides two examples showing how CSM handles policies that exceed either the bandwidth pool size or both the pool size and alarm threshold. For the purposes of this scenario, you can enter either one of these examples, or choose settings less likely to be exceeded in a practical implementation.

After enabling CSM, you will need to create Circuit Style SR-TE policy configurations. You can use any of the following methods to create Circuit Style SR-TE policies. In this scenario, we will create the same policy each time, but we will go through each method in order, so that you can decide which methods will best meet your needs:

- [Step 2: Configure Circuit Style SR-TE Policies Using Device CLI, on page 118](#)
- [Step 3: Configure Circuit Style SR-TE Policies Using Add, on page 121](#)
- [Step 4: Configure Circuit Style SR-TE Policies Using Import, on page 123](#)

- Step 1** From the main menu, choose **Services & Traffic Engineering > Circuit Style SR-TE > Configuration > Basic**.
- Step 2** Toggle the **Enable** switch to **True**.

The screenshot shows the 'Circuit Style SR-TE' configuration page. It has two tabs: 'Basic' and 'Advanced'. Under the 'Basic' tab, there are three main configuration areas:

- Enable:** A toggle switch is currently set to 'True' (indicated by a blue circle).
- Link CS BW Pool Size:** A text input field contains the value '10', followed by a '%' sign. Below the field, it says '0 to 100%'.
- Link CS BW Min Threshold:** A text input field contains the value '80', followed by a '%' sign. Below the field, it says '0 to 100%'.

At the bottom of the configuration area, there are three buttons: 'Commit Changes' (highlighted in blue), 'Get Default Values', and 'Discard Changes'.

- Step 3** Enter the required bandwidth pool size and threshold information, as explained in the table below. See also the examples below, and choose one of them to enter.

Field	Description
Link CS BW Pool Size	The percentage of each link's bandwidth reservable for Circuit Style SR-TE policies.
Link CS BW Min Threshold	The Link CS BW Pool utilization percentage beyond which Crosswork will generate a threshold-crossing event notification.

Step 4 Click **Commit Changes** to save the Basic configuration.

Step 5 (Optional): Click the **Advanced** tab to display additional CS-SR configuration values.

Circuit Style SR-TE

Configuration

Basic **Advanced**

Validation Interval * ⓘ Sec

10

5 to 3600 seconds

Timeout * ⓘ Sec

300

30 to 600 seconds

Restore Delegation Delay * ⓘ Sec

5

1 to 60 seconds

Debug Optimizer

Debug Optimizer ⓘ

False True

Debug Optimization Max Files * ⓘ

30

0 to 1024

Commit Changes
Get Default Values
Discard Changes

a) Change the values on the **Advanced** tab as explained in the table below.

Field	Description
Validation Interval	This is the interval that CSM will wait before the bandwidth that is reserved for an un-delegated policy is returned to the Circuit Style SR-TE policy bandwidth Pool.
Timeout	The duration CSM will wait for the delegation request, before generating a threshold-crossing alarm.
Restore Delegation Delay	The duration CSM will pause before processing a restore path delegation.
Debug Optimizer	Toggle the switch to True to turn on the Debug Optimizer for all CS-SR policies. The Debug Optimizer will write log files to the Crosswork file system whenever it calculates routes, up to the maximum number of files you specify.
Debug Optimization Max Files	Enter the maximum number of log files the Debug Optimizer will write. Once the maximum is reached, the Optimizer will overwrite existing files.

b) When you are finished entering Advanced configuration values, click **Commit Changes** to save the configuration.

Example

Example: Bandwidth Utilization Surpasses Defined Threshold

In this example, we assume the reserved bandwidth settings are as follows:

- Bandwidth Pool Size: 10%
- Bandwidth Pool Threshold: 1%

Our two nodes have 10 Gbps Ethernet interfaces, so the bandwidth pool size with these settings is 1Gbps and the alarm threshold is 100 Mbps.

1. We create a CS policy connecting node 5501-02 to node 5501-01 (r02 - r01), with a bandwidth of 100 Mbps.

Link Details 🗑️ ✕

Summary **Traffic Engineering**

General SR-MPLS SRv6 Tree-SID RSVP-TE

	A Side	Z Side
Node	5501-02	5501-01
IF Name	TenGigE0/0/0/0	TenGigE0/0/0/0
FA Affi...		
FA Top...		
∨ Circuit...		
Pool ...	1000 Mbps	1000 Mbps
Used	10 Mbps	10 Mbps
Avail...	990 Mbps	990 Mbps

2. Later, the requested bandwidth for the policy increases to 500 Mbps. The updated CS policy is created and operational (Oper State Up).

Link Details 🗑️ | ✕

Summary Traffic Engineering

General SR-MPLS SRv6 Tree-SID RSVP-TE

	A Side	Z Side
Node	5501-02	5501-01
IF Name	TenGigE0/0/0/0	TenGigE0/0/0/0
FA Affi...		
FA Top...		
∨ Circuit...		
Pool ...	1000 Mbps	1000 Mbps
Used	500 Mbps	500 Mbps
Avail...	500 Mbps	500 Mbps

- Since the bandwidth utilization of 500 Mbps with the updated policy is greater than the configured bandwidth threshold (100 Mbps), Crosswork triggers alerts.

Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for frankenrouter-02 TenGigE0/0/0/21
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for frankenrouter-02 TenGigE0/0/0/20
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-02 TenGigE0/0/0/2
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-02 TenGigE0/0/0/0
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-01 TenGigE0/0/1/0/1
Optima CSM App	Warning	Bandwidth pool allocation (500.000) exceeds pool threshold (100.00) for 5501-01 TenGigE0/0/0/0

Example: Bandwidth Pool Size and Usage Exceeded

In this example, we assume the reserved bandwidth settings are as follows:

- Bandwidth Pool Size: 10%
- Bandwidth Pool Threshold: 10%

The bandwidth pool size for the 10 Gbps Ethernet interfaces is 1Gbps and the alarm threshold is 100 Mbps.

- An existing CS-SR policy from node 5501-02 to node 5501-01 (*r02- r01*) uses a bandwidth of 500 Mbps.
- Later, a new policy requiring a bandwidth of 750 Mbps with a path from node 5501-02 to node 5501-01 to 5501-2 (*r02- r01- r2*) is requested. Since the existing policy and this new policy together exceed the bandwidth pool size and alarm threshold of 1 Gbps (750 Mbps + 500 Mbps = 1250 Mbps), the following behaviors occur:
 - The new CS-SR policy *r02- r01- r2* is created but not operational (Oper State Down).

Endpoint 5501-01 | TE RID: 10.255.255.1

Color 2000

Summary

- Admin State ↑ Up
- Oper State ↓ Down
- Binding SID 0
- Policy Type Circuit-Style
- Profile ID -
- Description -
- Traffic Rate 0 Mbps
- Unused True i

[See more](#) v

Candidate Path

Path Name	Pref	RoleState
<input checked="" type="checkbox"/> > cfg_r1-r2-2_discr_100	100	↓ A
<input type="checkbox"/> > cfg_r1-r2-2_discr_50	50	↓

- Alerts are triggered.

Source	Severity	Description
Optima CSM App	⚠ Warning	Unable to compute path for 10.255.255.1 <-> 10.255.255.2 color 2000 due to CsmUpdateStatus.NO_PATH
SR Policy [10.255.255.2#10.255...	⚠ Warning	Policy 'srte_c_2000_ep_10.255.255.1' has operational status as DOWN.
SR Policy [10.255.255.1#10.255...	⚠ Warning	Policy 'srte_c_2000_ep_10.255.255.2' has operational status as DOWN.

- Later, the CS-SR policy *r02- r01- r2* is updated and only requires 10 Mbps. The following behaviors occur:
 - Since the total bandwidth required for the two policies (10 Mbps + 500 Mbps = 510 Mbps) now requires less than the bandwidth pool size (1Gbps), CS-SR policy *r02- r01- r2* becomes operational (Oper State Up).

Step 2: Configure Circuit Style SR-TE Policies Using Device CLI

Endpoint 5501-01 | TE RID: 10.255.255.1

Color 2000

Summary

Admin State ↑ Up
 Oper State ↑ Up
 Binding SID 24005
 Policy Type Circuit-Style
 Profile ID -
 Description -
 Traffic Rate 0 Mbps
 Unused True i
[See more](#) v

Candidate Path

Path Name	Pref	RoleState
<input type="checkbox"/> > cfg_r1-r2-2_discr_50	50	↑
<input checked="" type="checkbox"/> > cfg_r1-r2-2_discr_100	100	↑ A

- Since the bandwidth utilization (10 Mbps) with the updated policy is below the configured bandwidth threshold (1 Gbps), alerts are cleared.

Source	Severity	Description
SR Policy [10.255.255.1#10.255.255.1]	✔ Clear	Policy 'srte_c_2000_ep_10.255.255.2' has operational status back to UP.
SR Policy [10.255.255.2#10.255.255.1]	✔ Clear	Policy 'srte_c_2000_ep_10.255.255.1' has operational status back to UP.

Step 2: Configure Circuit Style SR-TE Policies Using Device CLI

Prior to Cisco Crosswork, most network engineers created Circuit Style SR-TE policies directly on the devices themselves, using the appropriate network operating system CLI commands. This step of the scenario covers direct CLI policy configuration for a Cisco device. We present it only because this is a legitimate way to create these policies, and so you can see how the configuration implemented using this method matches the configuration for the other, Crosswork-native methods presented in this scenario.

Crosswork Network Controller's topology discovery will automatically recognize CS policy configurations implemented directly on devices, and will help you visualize them on the topology map. However, this method has some important drawbacks. To start with, you will need to be familiar with the CLI commands required to configure Circuit Style SR-TE policies properly. More importantly, Crosswork can *discover* Circuit Style SR-TE policies configured directly on a device, but cannot change or delete them. We encourage you to use instead the **Add** or **Import** methods, which allow you to manage and change your configuration using Crosswork. For help using these methods, skip this step and go on to [Step 3: Configure Circuit Style SR-TE Policies Using Add, on page 121](#) or to [Step 4: Configure Circuit Style SR-TE Policies Using Import, on page 123](#).

A Circuit Style SR-TE policy configuration must include the destination endpoint, the amount of requested bandwidth, and the bidirectional attribute. See [CS-SR Scenario: Assumptions and Prerequisites, on page 104](#) for additional requirements and notable constraints.

When configuring Circuit Style SR-TE policies directly on Cisco devices, make sure the configuration includes a Performance Measurement (PM) Liveness profile. A PM Liveness profile enables proper detection of candidate path liveness and effective path protection. Path Computation Clients (PCCs) do not validate past the first SID, so without PM Liveness, the path protection will not occur if the failure in the Circuit Style SR-TE policy candidate path occurs after the first hop in the segment list. Crosswork supports software-based and hardware-offload PM Liveness configuration methods. For more background on PM Liveness profiles and methods, see [Configuring SR Policy Liveness Monitoring](#).

Step 1 Use your preferred method to access the head-end device console and log in.

Step 2 If applicable, enable the hardware module on the device for PM configuration.

Example:

```
hw-module profile offload 4
reload location all
```

Step 3 Configure the Performance Measurement (PM) Liveness profile on the device. The following example uses a hardware-offload configuration.

Example:

```
performance-measurement
  liveness-profile sr-policy name CS-active-path
    probe
      tx-interval 3300
  !
  npu-offload enable    !! Required for hardware Offload only
  !
  !
  liveness-profile sr-policy name CS-protected-path
    probe
      tx-interval 3300
  !

npu-offload enable    !! Required for hardware Offload only
!
!
!
```

Step 4 Configure the Circuit Style SR-TE policy. All configuration entries shown are required in order for the Crosswork CSM feature pack to manage the policy. Entry values that you must change appropriately for your network (or for your PM Liveness profile) are shown in *italics*. See [CS-SR Scenario: Assumptions and Prerequisites, on page 104](#) for additional requirements and notable constraints.

Example:

```
segment-routing
  traffic-eng
    policy NCS1-NCS3

    performance-measurement
      liveness-detection
        liveness-profile backup name CS-protected    !! Name must match liveness profile defined for
        Protect path
```

```

    liveness-profile name CS-active                !! Name must match liveness profile defined for
Active path
!
!
bandwidth 10000
color 1000 end-point ipv4 192.168.20.4
path-protection
! Path protection is required on both ends of the candidate-paths
! Defining the Working path. Must have the highest CP preference
preference 100
  dynamic
    pcep
    !
    metric
      type igp
    !
  !
  constraints
    segments
      protection unprotected-only
      adjacency-sid-only
    !
    disjoint-path group-id 3 type node
  !
  bidirectional
    co-routed
    association-id 230
!
! Defining the Protect path. Must have second highest CP preference.
preference 50
  dynamic
    pcep
    !
    metric
      type igp
    !
  !
  constraints
    segments
      protection unprotected-only
      adjacency-sid-only
    !
    disjoint-path group-id 3 type node
  !
  bidirectional
    co-routed
    association-id 231
! Defining the restore path. It must have both the lowest CP preference and backup-ineligible
setting
preference 10
  dynamic
    pcep
    !
    metric
      type igp
    !
  !
  backup-ineligible
  !

  constraints
    segments
      protection unprotected-only
      adjacency-sid-only

```


You will need to change the device names and IP addresses you enter to match actual devices on your network.

Table 4: Example: Circuit Style SR-TE Policy Using Add

Expand this:	To specify this:
head-end	<ul style="list-style-type: none"> • Device: Enter NCS1. • Ip-address: Enter 192.168.20.4.
tail-end	<ul style="list-style-type: none"> • Device: Enter NCS3. • Ip-address: Enter 192.168.20.14.
disjoint-path	Click Enable disjoint-path .
disjoint-path > forward-path	<ul style="list-style-type: none"> • Type: Select Link. • group-id: Enter 531.
disjoint-path > reverse-path	<ul style="list-style-type: none"> • Type: Select Link. • group-id: Enter 5311.
performance-measurement	Click Enable performance-measurement .
performance-measurement > Profile-type	Click liveness .
performance-measurement > Profile-type > liveness-detection	Click Enable liveness-detection . Then: <ul style="list-style-type: none"> • Profile: Enter CS-active. • Backup: Enter CS-protected.
working-path	Click Enable working-path . Then select dynamic-path .
working path > dynamic	Click Enable dynamic-path . Then: <ul style="list-style-type: none"> • pce: Check the checkbox. • Metric-type: Select igmp • Bidirectional-association-choice: Select bidirectional-association-id and enter 230 in the field.
working path > dynamic > constraints > segments	Click Enable segments . Then in the Protection field, select unprotected-only .
protect-path	Click Enable protect-path . Then select dynamic-path .

Expand this:	To specify this:
protect-path > dynamic	Click Enable dynamic . Then: <ul style="list-style-type: none"> • pce: Check the checkbox. • Metric-type: Select igp • Bidirectional-association-choice: Select bidirectional-association-id and enter 231 in the field.
protect-path > dynamic > constraints > segments	Click Enable segments . Then in the Protection , field, select unprotected-only .
restore-path	Click Enable restore-path . Then select dynamic-path .
restore-path > dynamic	Click Enable dynamic-path . Then: <ul style="list-style-type: none"> • pce: Check the checkbox. • Metric-type: Select igp • Bidirectional-association-choice: Select bidirectional-association-id and enter 232 in the field.
restore-path > dynamic > constraints > segments	Click Enable segments . Then in the Protection field, select unprotected-only .

Step 7 When you are finished, click **Dry Run** to validate your changes and save them. Crosswork will display your changes in a popup window.

If you want to configure a service that has requirements that do not match those we describe in this example, contact Cisco Customer Experience.

Step 8 When you are ready to activate the policy, click **Commit Changes**.


Step 4: Configure Circuit Style SR-TE Policies Using Import

If your organization has already implemented Circuit Style SR-TE policies and wants to roll them out on more network devices, the easiest way to do so is using Crosswork Network Controller's **Import** function. You can use **Import** to download a policy template file from Crosswork. The template file will be in JSON or XML format. You can save the template under a new name, insert the policy values of your choice, and then import the modified file.

As well as being fast, using the **Import** function is a good way to standardize Circuit Style SR-TE policies across your network. You can set the same template files to establish CS-SR policies between multiple pairs of devices, varying only the endpoint names and addresses, and any other values as appropriate for each circuit.

Step 1 From the main menu, choose **Services & Traffic Engineering > Provisioning (NSO)**.

Step 2 In the **Services/Policies** column on the left, select **SR-TE > Circuit-Style Policy**.

- Step 3** Click . Then click the **Download sample JSON and XML files (.zip)** link. The downloaded ZIP file contains templates for all the Crosswork service types, including Circuit-Style, in JSON and XML formats.
- Step 4** Unzip `samplePayload.zip` and locate the `CS-Policy.json` and `CS-Policy.xml` template files.
- Step 5** Using the [JSON](#) or [XML](#) file editor of your choice, open the `CS-Policy` template file and save it under the name **cs1-cs4**.
- Step 6** If you are using the JSON template file, edit it so that it looks like the example below. If you are using the XML template, go on to the next step.

Example:**CS-SR Policy in JSON**

```
{
  "name": "NCS1-NCS3",
  "head-end": {
    "device": "NCS1",
    "ip-address": "192.168.20.4"
  },
  "tail-end": {
    "device": "NCS3",
    "ip-address": "192.168.20.14"
  },
  "color": 1000,
  "bandwidth": 10000,
  "disjoint-path": {
    "forward-path": {
      "type": "Link",
      "group-id": 531
    },
    "reverse-path": {
      "type": "Link",
      "group-id": 5311
    }
  },
  "performance-measurement": {
    "profile-type": "liveness", {
      "profile": "CS-active",
      "backup": "CS-protected"
    }
  },
  "path-protection": {},
  "working-path": {
    "dynamic": {
      "constraints": {
        "segments": {
          "protection": "unprotected-only"
        }
      },
      "pce": {},
      "metric-type": "igp",
      "bidirectional-association-id": 230
    }
  },
  "protect-path": {
    "dynamic": {
      "constraints": {
        "segments": {
          "protection": "unprotected-only"
        }
      },
      "pce": {},
      "metric-type": "igp",

```

```

        "bidirectional-association-id": 231
    },
    "revertive": true
},
"restore-path": {
  "dynamic": {
    "constraints": {
      "segments": {
        "protection": "unprotected-only"
      }
    },
    "pce": {},
    "metric-type": "igp",
    "bidirectional-association-id": 232
  }
}
}
}

```

Step 7 If you are using the XML template file, edit it so that it looks like the example below.

Example:

CS-SR Policy in XML

```

<config xmlns="http://tail-f.com/ns/config/1.0">
  <cs-sr-te-policy xmlns="http://cisco.com/ns/nso/cfp/cisco-cs-sr-te-cfp">
    <name>NCS1-NCS3</name>
    <head-end>
      <device>cs1</device>
      <ip-address>192.168.20.4</ip-address>
    </head-end>
    <tail-end>
      <device>cs4</device>
      <ip-address>192.168.20.14</ip-address>
    </tail-end>
    <color>1000</color>
    <bandwidth>10000</bandwidth>
    <disjoint-path>
      <forward-path>
        <type>Link</type>
        <group-id>531</group-id>
      </forward-path>
      <reverse-path>
        <type>Link</type>
        <group-id>5311</group-id>
      </reverse-path>
    </disjoint-path>
    <performance-measurement>
      <profile-type>liveness
        <profile>CS-active</profile>
        <backup>CS-protected</backup>
      </profile-type>
    </performance-measurement>
    <path-protection></path-protection>
    <working-path>
      <dynamic>
        <constraints>
          <segments>{
            <protection>unprotected-only</protection>
          </segments>{
        </constraints>{
          <pce></pce>
          <metric-type>igp</metric-type>
          <bidirectional-association-id>230</bidirectional-association-id>
        </dynamic>
      </working-path>
    </path-protection>
  </cs-sr-te-policy>
</config>

```


Step 5: View Circuit Style SR-TE Policies on the Topology Map

```

</working-path>
<protect-path>
  <dynamic>
    <constraints>
      <segments>
        <protection>unprotected-only</protection>
      </segments>
    </constraints>
  <pce></pce>
  <metric-type>igp</metric-type>
  <bidirectional-association-id>231</bidirectional-association-id>
</dynamic>
</protect-path>
<restore-path>
  <dynamic>
    <constraints>
      <segments>
        <protection>unprotected-only</protection>
      </segments>
    </constraints>
  <pce></pce>
  <metric-type>igp</metric-type>
  <bidirectional-association-id>232</bidirectional-association-id>
</dynamic>
</restore-path>
</cs-sr-te-policy>
</config>

```

Step 8 When you have finished editing the file and saved your changes, navigate to **Services & Traffic Engineering > Provisioning > SR-TE > Circuit-Style Policy** again.

Step 9 Click  again. In the **File Name** field, enter the path to and file name of your modified template file, or click **Browse** to locate and select it. Then click **Import**.

Step 5: View Circuit Style SR-TE Policies on the Topology Map

Next, we'll use Crosswork to visualize the NCS1-NCS3 Circuit Style SR-TE policy and isolate it on the map. To make this step more realistic and demonstrate how to focus on just one policy, the scenario assumes that we have multiple active Circuit Style SR-TE policies, not just the policy we created. We'll also view the Circuit Style SR-TE policy details, including endpoints, bandwidth constraints, IGP metrics, and candidate (Active/Working and Protect) paths.

Step 1 From the main menu, choose **Services & Traffic Engineering > Traffic Engineering > SR-MPLS**. Then click **Circuit Style**.

Traffic Engineering Refined By: Headend or E... ▾

SR-MPLS	SRv6	Tree-SID	RSVP-TE
90 Total	6 Circuit Style	2 BWoD	0 LCM
		0 ↓ Admin Down	31 ↑ Oper Up
			59 ↓ Oper Down

The **SR Policy** table lists all of the Circuit Style SR-TE policies.

Step 2 From the **Actions** column, click **Circuit Style SR-TE > View Details** for one of the Circuit Style SR-TE policies.

Note You cannot edit or remove Circuit Style SR-TE policy configurations that have been created directly on the device.

The screenshot shows the Traffic Engineering interface. On the left is a topology map of the San Francisco Bay Area with nodes like xrv9k-14, xrv9k-16, xrv9k-15, and xrv9k-17. A path is highlighted between xrv9k-16 and xrv9k-15. On the right is the 'Traffic Engineering' panel with a summary table and a list of SR Policies.

SR-MPLS	SRv6	Tree-SID	RSVP-TE
90	6	2	0
Total	Circuit Style	BWd	LCM
0	0	0	0
Admin Down	Oper Up	Oper Down	
31	59		

Head...	Endp...	Color	Admin...	Oper St...	Actions
<input checked="" type="checkbox"/>	xrv9k-16	xrv9k-15	11056	+	+
<input type="checkbox"/>	xrv9k-15	xrv9k-16	11056	+	+
<input type="checkbox"/>	xrv9k-16	xrv9k-15	4294...	+	+
<input checked="" type="checkbox"/>	xrv9k-15	xrv9k-16	4294...	+	+
<input checked="" type="checkbox"/>	xrv9k-...	xrv9k-12	5600	+	+
<input checked="" type="checkbox"/>	xrv9k-12	xrv9k-...	5600	+	+

The **Circuit Style Policy Details** window is displayed in the side panel. By default, the Active path is displayed in the topology map and shows the bidirectional paths (Bi-Dir Path checkbox is checked) on the topology map. The Candidate Path list displays the Active (path that currently takes traffic) and Protected paths.

The screenshot shows the 'Circuit Style Policy Details' window. On the left is a topology map of the Fremont area with nodes xrv9k-23 and xrv9k-25. A path is highlighted between them. On the right is the 'Circuit Style Policy Details' panel.

Current

- Headend** xrv9k-25 | TE RID: 192.168.0.15 PCC IP: 192.168.0.15 Source IP: 192.168.0.15
- Endpoint** xrv9k-23 | TE RID: 192.168.0.15 Dest IP: 192.168.0.15
- Color** 6905

Performance Metrics

- Traffic Rate** 0 Mbps avg.

Summary

Candidate Path

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_srte_c_6905_ep_192.16...	100	+	+
<input type="checkbox"/> cfg_srte_c_6905_ep_192.16...	50	+	+

Note The Bandwidth Constraint value can differ from the bandwidth you requested if the value was increased and insufficient resources existed to satisfy demand on all Active and Protected candidate paths.

Step 3 View Candidate path configuration details.

a) The **Circuit Style Policy Details** window allows you to drill down to view more information about the candidate paths. The operational (Oper State Up) candidate path with the highest preference will always be the Active path. In this example, the Protected path (with preference 50) is currently the Active path and is displayed on the topology map. Notice that it is designated with a green "A" icon under State to clearly indicate it is currently the operational Active path. Click **Expand All** to view more information about both paths.

Step 5: View Circuit Style SR-TE Policies on the Topology Map

The screenshot displays the Cisco Crosswork Network Controller 6.0 interface. On the left, a topology map shows a network path between two nodes labeled 'avw-25' in Fremont and San Jose. The path is color-coded: a purple link (highest preference), a blue link (second preference), and a pink link (third preference). The map includes various geographical markers and district names like Fremont, Mountain View, and San Jose.

On the right, the 'Circuit Style Policy Details' panel is open, showing a table of candidate paths. The table has columns for Path Name, Pref, Role, and State. The first path is 'cfg_srte_c_6905_ep_192.168.0.23...' with a preference of 100 and a state of 'Down'. The second path is 'cfg_srte_c_6905_ep_192.168.0.23...' with a preference of 90 and a state of 'Active'. The third path is 'cfg_srte_c_6905_ep_192.168.0.23...' with a preference of 50 and a state of 'Active'. The details for the active paths include Path Name, Oper State, Metric Type, Bandwidth, Bi-Dir Association ID, Disjoint Group, PCE Initiated, Affinity, Segment Type, and SID Algorithm.

- Note**
- First preference paths are shown as purple links.
 - Second preference paths are shown as blue links.
 - Third preference paths are shown as pink links.

If the Circuit Style SR-TE policy configuration was done through the UI, you have the option to view the Circuit Style SR-TE policy configuration. To see the configuration, click the link next to **Config ID**. For example:

Circuit Style Policy Details

Current History

Path Name	Pref	Role	State
<input checked="" type="checkbox"/> cfg_srte_c_6905_ep_192.168....	100		↑ A

Path Name cfg_srte_c_6905_ep_192.168.0.25_disc

Oper State ↑ Up | A Active

Metric Type IGP

Bandwidth Requested: 9.006 Mbps
Reserved: 0 Mbps

Bi-Dir Association ID 5906

Config ID [CS-CS-SR-WP-601-head-end-internal](#)


Disjoint Group ID: 567
Association Source: 0.0.0.0
Type: Node-disjoint

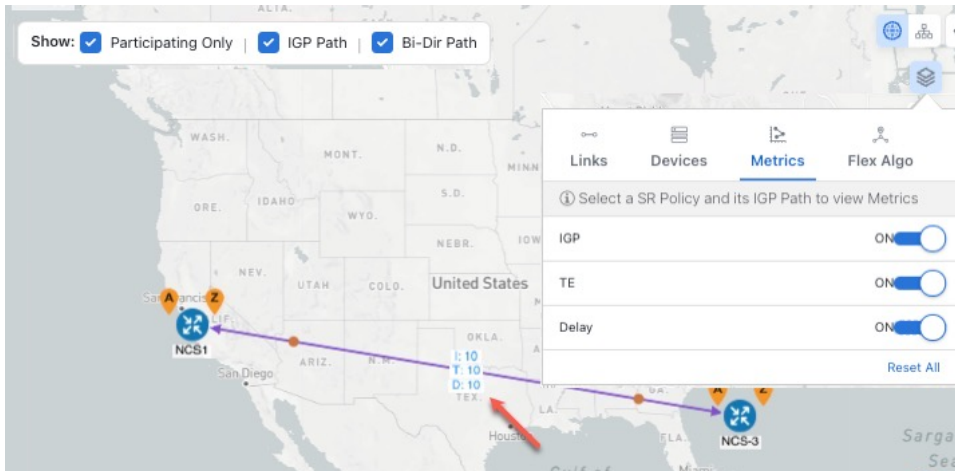
PCE Initiated false

Affinity Exclude-Any: -
Include-Any: -
Include-All: -

Segment Type Unprotected

SID Algorithm -

Step 4 To view the physical path and metrics between endpoints of the selected policies, click  to turn applicable metrics on and check the **IGP Path** checkbox.



The screenshot shows a map of the United States with a path highlighted between two nodes, NCS1 (San Francisco) and NCS-3 (Houston). A 'Metrics' panel is open on the right, showing checkboxes for IGP, TE, and Delay, all of which are turned on. The panel also includes a 'Reset All' button and a note: 'Select a SR Policy and its IGP Path to view Metrics'.

Step 6: Verify Circuit Style SR-TE Policy Bandwidth Utilization

Let's verify that the reserved bandwidth pool settings we defined when enabling Circuit Style SR-TE (see [Step 1: Enable Circuit Style Manager, on page 113](#)) are configured properly. We can also check how much bandwidth is either in use or still available.

- Step 1** From the main menu, choose **Services & Traffic Engineering > Traffic Engineering > SR-MPLS**. Then, under the **SR-MPLS** column, click **Circuit Style**. The **SR Policy** table lists all CS SR policies.
- Step 2** In the **SR Policy** table, check the check box next to the participating device whose details you want to see.
- Step 3** On the topology map, click on a participating Circuit Style SR-TE policy node to display the **Device Details** for that node.
- Step 4** On the **Device Details** page, click the **Links** tab to display the list of CS-SR and other links on the participating node. Then click on the link whose details you want to see. The **Link Details** list displays a **Summary** of the link information.
- Step 5** On the **Link Details** page, click on the **Traffic Engineering** tab, then the **General** tab. The **Link Details** list displays detailed information for the link.

Under **Circuit Style Bandwidth Pool**, you can see the reserved bandwidth pool size, the amount of bandwidth currently being used, and the amount of bandwidth (of the total allocated to Circuit Style SR-TE policies) is still available.

In this example, the reserved bandwidth pool size is displayed as 800 Mbps for NCS-3 and NCS1. The configured settings were earlier defined as 80% for the bandwidth pool size. Since the interfaces on this circuit are both 1 Gbps, we can confirm that Circuit Style SR-TE has correctly allocated 80 percent of bandwidth for these two interfaces.

Link Details



Link Details		
Summary		
Traffic Engineering		
General		
SR-MPLS		
SRv6		
Tree-SID		
RSVP-TE		
	A Side	Z Side
Node	NCS-3	NCS1
IF Name	GigabitEthernet0/0/0/2	GigabitEthernet0/0/0/0
FA Affinities		
FA Topologies		
<input type="checkbox"/> Circuit Style Bandwidth Pool		
Pool Size	800 Mbps	800 Mbps
Used	4 Mbps	4 Mbps
Available	796 Mbps	796 Mbps

Step 7: Trigger Circuit Style SR-TE Path Recomputation

Circuit-Style policies are static in nature, meaning once the paths are computed, Crosswork will not re-compute them automatically. Changes in your network topology or operational status may affect the previously computed Working and Protected paths to the extent that you want Crosswork to re-compute and optimize them for the new situation. In this step, we see a demonstration of how to re-optimize for paths to accommodate these types of changes.

For more details on the logic CSM employs in these cases, see [What Happens When Path Failures Occur?](#), on page 110.

Step 1 From the main menu, choose **Services & Traffic Engineering > Traffic Engineering > SR-MPLS** and click **Circuit Style**.



Step 2 The SR Policy table displays the status of each of the Active CS-SR policies. One of them is Operationally down.

Step 3 From the **Actions** column next to the CS-SR TE policies whose Operational State is **Down**, click ***** > View Details**.

Crosswork displays the **Circuit Style Policy Details** window in the side panel. By default, the topology map shows the Active path and the bidirectional paths on the topology map (for these to appear, the **Bi-Dir Path** checkbox in the topology map's **Show** panel must be checked). The **Candidate Path** list at the bottom of the side panel displays the Active (Working) and Protected paths.

In the Summary panel, click the **See more** link to get a closer look at the type of Summary details available. The Candidate Path list displays the Active and Protected paths.

Step 4 To have Crosswork re-optimize these paths: Click ******* at the top of the **Circuit Style Policy Details** panel and select **Re-optimize**. Click **Yes** when prompted to confirm your selection.

Summary and Conclusion

In this scenario, we observed how to use Circuit Style Segment Routing policies to reserve bandwidth for high-priority services and traffic in the network. CS-SR removes the need to manually track and calculate high-priority traffic paths, but still gives you control over how those paths are calculated and optimize bandwidth usage on each path. You can use these policies to ensure that available bandwidth is dedicated for these services. As traffic changes, Circuit Style policies warn you when your dedicated "circuit" paths fail, and allows you to re-optimize them as needed.



CHAPTER 5

Network Maintenance Window

This section explains the following topics:

- [Overview, on page 133](#)
- [Scenario: Install an SMU During a Scheduled Maintenance Window, on page 134](#)

Overview

Objective

Schedule and automate maintenance workflows with minimal network interruption and most efficient results.


Challenge

Maintenance activities typically require system downtime and temporary disruption of services. Keeping downtime and disruption to a minimum is critical but challenging. Therefore, maintenance activities must take place during a carefully calculated optimal time slot, usually when activity is at its lowest.

Solution

Cisco Crosswork Change Automation and Cisco Crosswork Health Insights are optional add-on applications that provide the functionality needed to automate the scheduling and execution of maintenance tasks. Planning the optimal time for maintenance activities can be done successfully using Cisco WAE Design to simulate “what-if” scenarios based on timed topology snapshots exported from Cisco Crosswork Network Controller using APIs.

How Does it Work?



- Using the Crosswork Network Controller APIs, you can create topology snapshots (plan files) which capture and represent topology state at a given point in time, including the IGP topology as well as interface level statistics (traffic load). For impact analysis purposes, these snapshots should be representative of a time period to be evaluated for an upcoming maintenance activity. For example, if you are planning a router upgrade at midnight on a Monday, you would take snapshots from several Mondays at midnight to evaluate typical traffic loads at this time. You can export these plan files to a central storage repository, where a library of topology plan files can be stored for a specified period of time.
- Cisco WAE Design allows you to explore “what-if” scenarios relevant to the planning of the maintenance window. For example, in the case of upgrading a router, Cisco WAE Design can simulate the resulting traffic load on the remaining devices after traffic is diverted from the device being upgraded. You can also explore the impact of deploying tactical traffic engineering policies to further optimize the topology during the maintenance window. For more information, contact your Cisco Customer Experience representative.

Additional Resources

[Cisco Crosswork Change Automation and Health Insights User Guide](#)

[Cisco WAE Design documentation](#)

Cisco Crosswork Network Automation API Documentation on [Cisco Devnet](#)

Scenario: Install an SMU During a Scheduled Maintenance Window

Scenario Context

In this scenario, we will first use Cisco WAE Design to evaluate the impact of removing a Provider node from the network during a specific time-frame in order to install a Cisco SMU (Software Maintenance Upgrade) on the device. We will choose a predefined Crosswork Playbook to automate the SMU installation on the device, and we will schedule it to run during the predetermined maintenance window.

Assumptions and Prerequisites

The high-level requirements that must be met to enable this scenario include:

- Cisco Crosswork Change Automation is installed and running.
- You have access to Cisco WAE Design.
- The Device Override Credentials are set for Crosswork Network Change Automation, so that running the Playbook is possible. Go to **Administration > Settings > System Settings > Network Automation** to set these credentials.

Step 1 Download Topology Plan Files for Impact Analysis

When considering when to take down a device for maintenance so that there will be the least impact to the network, you need information about the traffic trends around that device at the targeted time. Using the Cisco Crosswork Optimization API, you can download plan files that capture a snapshot of the network topology at that time. If you download plan files at the same time over a period of time, you can use Cisco WAE Design to analyze the traffic trends. Based on this analysis, you can decide whether the impact to the network would be acceptable or not.

Refer to [Cisco Crosswork Network Automation API Documentation on Cisco Devnet](#) for more information about the API.

The input for this scenario is as follows:

Step 1 Prepare the input required to download the plan file. You need to specify the version of Cisco WAE design that you will be using for analysis and the format in which you want the plan file, either `txt` or `pln`.

Note If you download the plan file as a `txt` file, you can view it in any text editor. If you download it as a `pln` file, you can open it only in Cisco WAE design.

The input for this scenario is as follows:

```
{
  "input": {
    "version": "7.3.1",
    "format": "txt",
  }
}
```

Step 2 Invoke the API on the Cisco Crosswork Network Controller server using the input prepared in the previous step. For example:

```
curl --location --request POST
'https://10.194.63.198:30603/crosswork/nbi/optima/v1/restconf/operations/cisco-crosswork-
optimization-engine-operations:get-plan \
--header 'Content-Type: application/yang-data+json' \
--header 'Accept: application/yang-data+json' \
--header 'Authorization: Bearer
```


Step 1 Go to **Network Automation > Run Playbook**.

Step 2 Browse the Available Playbooks list, and click the Install a SMU playbook. You can also filter using keywords to identify the playbook. Note that the playbook execution stages, supported software platform, software version, and individual play details are displayed on the right side.

Step 3 Click **Next** to go to the next task: **Select Devices**. All devices tagged with City: NY will be selected for SMU installation.

Step 4 Under the City tag on the left, click **NY**. The devices tagged with NY are listed on the right and are automatically selected.

Reachability St...	Operational State	Host name	Software Pla...	Provider	Unique Identifier
Reachable	OK	P-BOTTOMRIGHT	IOS XR		bcc1bc0c-d1cc-4932-90a7-30...
Reachable	OK	P-TOPRIGHT	IOS XR		ce944bd2-c476-4391-9c47-b...

Step 5 Click **Next** to go to the next task: **Parameters**.

Step 6 Edit the runtime parameters to execute the SMU playbook. Alternatively, you can upload a JSON file that contains the parameter values. The following values are used specifically for this scenario. You can change them as required:

- Under the **Install a SMU or an optional package on a router** play, set **collection_type** as **mdt**.
- Under the **Perform DLM node lock on device(s)** play, set **retry_count** and **retry_interval** as **3** and **5s** respectively.

Step 2: Schedule the SMU Installation Playbook Run

- c. Under the **Install add package(s)** play, enter the SMU package name in **item 1**.

- d. Under the **Install activate package(s)** play, click the piece of paper symbol, and set **action** to **Activate**.
- e. Under the **Install commit package(s)** play, set the **action** to **Commit**.
- f. Under the **Verify package in committed list on router** play, set **collection_type** to **mdt**.

Step 7 Click **Next** to go to the next task: **Execution Policy**.

Step 8 Set the **Execution Mode** to **Continuous**. This will set the playbook to run uninterrupted, with no pauses. Under **Failure policy**, select the action you want taken if the execution fails: **Abort** or **Complete Roll Back**.

Step 9 Under **Schedule**, set the playbook to execute for the optimal time calculated during the impact analysis stage. Uncheck the **Run Now** option. Note the calendar at the right, and the timers that let you **Schedule Pre-check** and **Schedule Perform** play execution dates and times.

Execution Mode

- Continuous**: Run the playbook without interruption.
- Single Stepping**: Run the Playbook one play at a time, and specify when to pause.
- Dry Run**: View the configuration changes without performing a commit.

Collect Syslog
 Yes No

Failure policy
 On failure:

Timeout

Schedule

Run Now

Schedule Pre-check (America/Los_Angeles)
 2023-10-31
 Increment hours: Increment minutes:
 Decrement hours: Decrement minutes:

Schedule Perform (America/Los_Angeles)
 2023-10-31
 Increment hours: Increment minutes:
 Decrement hours: Decrement minutes:

All Scheduled Jobs

October 2023

Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31	1	2	3	4

Buttons: Cancel, Previous, Next

Step 10 Click **Next** to go to the next task: Confirm Job.

Step 11 Review your job details. Label your job with a unique name.

Step 12 When you are finished, click **Run Playbook**. The SMU installation is now scheduled to run in the planned maintenance window.

Review your Job

Playbook: Install a SMU or an optional package on a router
 Continuous (0)
 Pre Maintenance (1)
 Maintenance (5)
 Post Maintenance (1)

Tag:

Mop Params

```
{
  "1": {
    "collection_type": "mdt"
  },
  "2": {
    "retry_count": "3",
    "retry_interval": "5s"
  },
  "3": {
    "packages": [
      "xrv-9k-base-2.0.0.144-r721.CSCcv93809x86.64.rpm"
    ]
  },
  "4": {
    "packages": [
      "xrv-9k-base-2.0.0.144-r721.CSCcv93809x86.64.rpm"
    ]
  },
  "5": {
    "action": "Activate"
  },
  "6": {
    "action": "Commit"
  }
}
```

Label your Job

Name:

Labels:

Buttons: Cancel, Previous, Run Playbook

Step 3 Verify the SMU Job Status

- Step 1** After the scheduled maintenance window time, go to **Network Automation > Automation Job History**. Under **Job Sets**, check that the job status icon on the SMU install job is green, indicating that the scheduled job has run successfully.
- Step 2** Select the SMU install job. Note the Job Set details on the right side. Click the link under the job **Execution ID** for job details.
- Step 3** Double-check that the correct SMU has been installed by executing the `show install active summary` and `show install committed summary` commands on the device. If the install was successful, the SMU you installed will appear in the list of packages. The following figure shows example outputs from these commands:

```

1 RP/0/RP0/CPU0:CX-AA-PE4#show install active summary
2 Mon Apr 12 11:09:20.198 EDT
3   Active Packages: 12
4     ncs5500-xr-6.6.3 version=6.6.3 [Boot image]
5     ncs5500-ospf-2.0.0.0-r663
6     ncs5500-mpls-2.1.0.0-r663
7     ncs5500-eigrp-1.0.0.0-r663
8     ncs5500-isis-2.2.0.0-r663
9     ncs5500-li-1.0.0.0-r663
10    ncs5500-mpls-te-rsvp-4.1.0.0-r663
11    ncs5500-mcast-3.1.0.0-r663
12    ncs5500-mgbl-3.0.0.0-r663
13    ncs5500-k9sec-3.1.0.0-r663
14    ncs5500-routing-4.0.0.17-r663.CSCvr43225
15    ncs5500-mpls-te-rsvp-4.1.0.17-r663.CSCvr43225
16
17 RP/0/RP0/CPU0:CX-AA-PE4#show install committed summary
18 Mon Apr 12 11:09:27.092 EDT
19   Committed Packages: 12
20     ncs5500-xr-6.6.3 version=6.6.3 [Boot image]
21     ncs5500-ospf-2.0.0.0-r663
22     ncs5500-mpls-2.1.0.0-r663
23     ncs5500-eigrp-1.0.0.0-r663
24     ncs5500-isis-2.2.0.0-r663
25     ncs5500-li-1.0.0.0-r663
26     ncs5500-mpls-te-rsvp-4.1.0.0-r663
27     ncs5500-mcast-3.1.0.0-r663
28     ncs5500-mgbl-3.0.0.0-r663
29     ncs5500-k9sec-3.1.0.0-r663
30     ncs5500-routing-4.0.0.17-r663.CSCvr43225
31     ncs5500-mpls-te-rsvp-4.1.0.17-r663.CSCvr43225
32
33 RP/0/RP0/CPU0:CX-AA-PE4#

```

Summary and Conclusion

In this scenario we saw how to plan for a maintenance window in which to bring down a device in order to install an SMU. The goal is to cause as little impact to the traffic in the network as possible. To analyze the impact on the network, we showed how to download snapshots of the network topology (plan files) at the target time for the maintenance window. The plan files can then be analyzed using Cisco WAE design.

Assuming that the impact was acceptable, we chose a predefined playbook to install the SMU on specific devices and we scheduled it for the planned maintenance window time when there would be the least impact to the network.



CHAPTER 6

Programmable Closed-Loop Remediation

This section explains the following topics:

- [Overview, on page 143](#)
- [Scenario: Achieve Predictive Traffic Load Balancing Using Segment Routing Affinity, on page 144](#)
- [Workflow, on page 145](#)

Overview

Objective

Detect anomalies and generate alerts that can be used for notifying an operator or triggering automation workflows.

Challenge

Discovering and repairing problems in the network usually involves manual network operator intervention and is time-consuming and error prone.

Solution

Incorporating Cisco Crosswork Change Automation and Cisco Crosswork Health Insights into Cisco Crosswork Network Controller gives service providers the ability to automate the process of discovering and remediating problems in the network by allowing an operator to match an alarm to pre-defined remediation tasks. These tasks will be performed after a defined Key Performance Indicator (KPI) threshold has been breached. Remediation can be implemented with or without the network operator's approval, depending on the setting and preferences of the operator.

Using such closed-loop remediation reduces the time taken to discover and repair a problem while minimizing the risk of making a mistake and creating an additional error through high-stakes manual network operator intervention.

How Does it Work?

Smart Monitoring

- The Smart Monitoring feature helps operators collect, filter, and present the data in useable formats, such as graphs and tables. Operators can remain focused on their business goals while the configuration required for the data collection is done by the Cisco Crosswork Network Controller and Cisco Crosswork Change Automation and Cisco Crosswork Health Insights using the feature Zero-touch telemetry.

- By using a common collector to collect network device data over SNMP, CLI, and model-driven telemetry, and making it available as modelled data described in YANG, duplicate data collection is avoided, optimizing the load on both the devices and the network.
- Recommendation Engine analyzes network device hardware and software, configuration, and employs a pre-trained model built from data mining, producing KPI relevant recommendations facilitating per use-case monitoring.
- KPIs cover a wide range of statistics from CPU, memory, disk, layer 1/2/3 network counters, to per protocol, LPTS and ASIC statistics.

Smart Filtering

- Cisco Crosswork Health Insights builds dynamic detection and analytics modules that allow operators to monitor and see alerts on network events based on user-defined logic (KPI).
- Key Performance Indicators (KPIs) Alerting Logic can be:
 - Simple static thresholds (TCA); e.g., CPU load above 90 percent.
 - Moving average, standard-deviation, and percentile based, etc., e.g., CPU load above mean and staying there for five minutes.
 - Streaming jobs which provide real-time alerts or batch jobs which run periodically.
 - Customized for threshold values and visualization dashboards.
 - Customized operator-created KPIs based on business logic.
 - TCAs can be exported or integrated with other systems via HTTP, Slack, and socket interfaces.
- KPIs are associated with dashboards, which provide real-time and historical views of the data and corresponding TCAs.
- KPIs also provide purpose-built dashboards that go beyond raw data and provide valuable information in various infographic style charts and graphs useful for triaging and root-causing complex issues.

Smart Remediation

- Health Insights KPIs can be associated with Cisco Crosswork Change Automation playbooks, which can be either executed manually or via auto-remediation. Remediation workflow could be used to fix the issue or collect more data from the network devices. By proactively remediating the situation, instead of resorting to ad hoc debugging and unscheduled downtime, operators can save time and money, providing better QOE to their customers.
- Health Insights does the correlation of alerts or anomalies on the topology of the network, allowing easy visualization of the impact of events.

Scenario: Achieve Predictive Traffic Load Balancing Using Segment Routing Affinity

Scenario Context

To maintain smooth and optimal traffic flow, operators need to be able to monitor traffic on the interfaces, identify errors such as CRC, watchdog, overrun, and then reroute the traffic so that the SLA is maintained. This process can be automated using Cisco Crosswork Network Controller with Cisco Crosswork Health Insights and Cisco Crosswork Change Automation applications installed.

Assumptions and Prerequisites

Cisco Crosswork Health Insights and Cisco Crosswork Change Automation must be installed and running.

Workflow

The following is a high-level workflow for executing this scenario:

-
- Step 1** Deploy Day0 ODN templates on edge nodes with dynamic path calculation delegated to SR-PCE and the ODN template configured to exclude links that are tagged with a specific affinity; for example, RED affinity. ODN allows a service head-end router to automatically instantiate an SR-TE policy to a BGP next-hop when required (on-demand). The ODN template defines the required SLA using a specific color.
- For an example procedure for creating an ODN template, refer to [Step 1 Create an ODN template to map color to an SLA objective and constraints, on page 31](#) in [Scenario: Implement and Maintain SLA for an L3VPN Service for SR-MPLS \(using ODN\), on page 29](#).
- Step 2** Create an L3VPN route policy to specify the prefixes to which the SLA applies and mark them with the same color used in the ODN template. When traffic from the specified network with a matching color is received, paths are computed based on the SLA defined in the ODN template.
- For an example procedure for creating a route policy, refer to [Step 1 Create an ODN template to map color to an SLA objective and constraints, on page 31](#).
- Step 3** Provision an L3VPN across the required endpoints and create an association between the VPN and the route policy. This makes the connection between the VPN and the ODN template that defines the SLA.
- For an example procedure for provisioning an L3VPN, refer to [Step 3 Create and provision the L3VPN service , on page 36](#).
- Step 4** Define and enable the KPIs on the devices. This will continuously monitor the uplink interfaces on the L3VPN endpoints.
- For information about defining KPIs, see the [Cisco Crosswork Change Automation and Health Insights User Guide](#).
- Step 5** When there is an error on monitored interfaces, mark the dirty link with RED affinity so that it will be excluded, based on the specifications of the ODN template. This is achieved by creating a custom playbook. Cisco Crosswork Network Controller learns the name of the interface generating the alert regarding the error and this is fed into the custom playbook so that the affinity configuration can be pushed to the relevant router, forming a closed-loop automation scenario. In this way, the customer should not experience outages.
- For information about defining playbooks, see the [Cisco Crosswork Change Automation and Health Insights User Guide](#).
- Step 6** Cisco Crosswork Network Controller continues to monitor the link and when there are no longer alerts, the RED affinity tag can be removed. Define another playbook for this purpose.
-



CHAPTER 7

Automation of Onboarding and Provisioning of IOS-XR Devices Using ZTP

This section explains the following topics:

- [Overview, on page 147](#)
- [Scenario: Use ZTP to Onboard and Provision New Devices Automatically, on page 148](#)
- [ZTP Scenario: Workflow, on page 148](#)

Overview

Objective

Allow users to quickly, easily, and automatically onboard new devices and provision them using a Cisco-certified software image and a day-zero software configuration.

Challenge

Deploying and configuring network devices is a tedious task. It requires extensive hands-on provisioning and configuration by knowledgeable personnel, which is time-consuming, expensive, and error-prone.

Solution

Automate onboarding of new devices using Crosswork Zero Touch Provisioning (Cisco Crosswork ZTP). Cisco Crosswork ZTP allows users to provision networking devices remotely, without a trained specialist on site. After establishing an entry for the device in the DHCP server and the ZTP application, all the operator needs to do is connect the device to the network, power on and press reset to activate the devices. A certified image and configuration are downloaded and automatically applied to the device. After it is provisioned in this way, the new device is onboarded to the Crosswork device inventory where it can be monitored and managed like other devices.

How Does it Work?

- **Classic ZTP:** The DHCP server verifies the device's identity based on the device's serial number, then offers downloads of the boot file and image. After the device is imaged, it downloads the configuration file and executes it.
- **Secure ZTP:** The device and the Cisco Crosswork ZTP bootstrap server authenticate each other using the device's Secure Unique Device Identifier (SUDI) and Crosswork server certificates over TLS/HTTPS. After a secure HTTPS channel is established, the Crosswork bootstrap server allows the device to request to download and apply a set of signed image and configuration artifacts adhering to the RFC 8572 YANG

schema. After the image (if any) is downloaded and installed, and the device reloads with the new image, the device downloads configuration scripts and executes them.

- Plug and Play (PnP) ZTP: The Cisco PnP agent on the IOS-XE device and the Cisco Crosswork PnP Server authenticate each other over HTTP using a PnP profile supplied on a TFTP server. They then establish a secure connection over HTTPS and the PnP agent downloads and installs image (optional) and configuration artifacts.

Additional Resources

Detailed information is available in the ZTP chapter in the [Cisco Crosswork Network Controller Administration Guide](#).

Scenario: Use ZTP to Onboard and Provision New Devices Automatically

Scenario Context

With the exponential growth of service provider networks and their rapid expansion into new customer sites and new locations, there is a need to connect an ever-increasing number of edge devices. At the same time, functional sophistication is increasing, requiring more time to configure those devices and activate new services. Manual processes limit a service provider's ability to rapidly scale networks and roll out new services in a cost-efficient way.

In this scenario, we will onboard the new IOS-XR devices required to set up a new customer site in a remote location and go live, without the need to send skilled technicians on time-consuming and costly on-site visits to complete the provisioning.

We will leverage the configuration of devices at existing customer sites that are already set up and operating to ensure that the Day0 configuration of the new devices includes whatever is necessary to get the devices up and running quickly and efficiently.

Assumptions and Prerequisites

- Crosswork ZTP must be installed in your Cisco Crosswork Network Controller setup.
- For Classic ZTP, Crosswork and the devices must be deployed in a secure network domain. Secure ZTP does not have this requirement; it is secure across public networks.
- The Crosswork server must be reachable from the devices, via an out-of-band management network or an in-band data network.
- If you want to onboard devices to Cisco NSO also, Cisco NSO must be configured as a Crosswork provider. When configuring the NSO provider, be sure to set the provider property key to *forward* and the property value to *true*.

ZTP Scenario: Workflow

This is a high-level workflow for onboarding IOS-XR devices using Cisco Crosswork Classic or Secure ZTP.

To onboard IOS-XE devices, or for more detailed information on these options, see the ZTP chapter in the [Cisco Crosswork Network Controller 5.0 Administration Guide](#).

Step 1 Assemble and upload ZTP assets

a) Assemble the following assets before you begin:

- (Optional) Software images. For Classic ZTP, you can use Cisco IOS-XR versions 6.6.3, 7.0.1, 7.0.2, 7.0.12, and 7.3.1 or later. For Secure ZTP, use Cisco IOS-XR 7.3.1 or later (except 7.3.2 and 7.4.1).
- Configuration Files: SH, PY, or TXT files. You can specify up to three different configuration files for Secure ZTP.
- Credentials of the devices to be onboarded
- Serial numbers of the devices to be onboarded

For Secure ZTP only, also assemble:

- Owner certificates - your organization's CA-signed end-entity certificates, installed on your devices and binding a public key to your organization.
- Pinned domain certificate - your organization's CA- or self-signed domain certificate, with its public key pinned to your organization's DNS network domain. The PDC helps your devices verify that images and configurations downloaded and applied during ZTP processing come from within your organization.
- Ownership vouchers - Nonceless audit vouchers that verify that devices being onboarded with ZTP are bootstrapping into a domain owned by your organization. Cisco supplies OV's when a request is submitted with your organization's PDC and device serial numbers.

- b) If applying software images: Upload the software images. Go to **Device Management > Software Images**.
- c) Upload the configuration files. Go to **Device Management > ZTP Configuration Files**.
- d) Upload device serial numbers. Go to **Device Management > Serial Number and Voucher** and click **Add Serial Number**.
- e) For Secure ZTP, upload your pinned domain certificate and owner certificates. Go to **Administration > Certificate Management** and add your certificates.
- f) For Secure ZTP, upload ownership vouchers. Go to **Device Manager > Serial Number and Voucher**.

Step 2 Create a ZTP profile combining an image file and configuration file

Crosswork uses ZTP profiles to automate imaging and configuration processes. While optional, creating ZTP profiles is recommended as the best way to combine a single image file and configuration file based on a product or device family, such as the Cisco ASR 9000 or Cisco NCS5500. We recommend that you create only one day-zero ZTP profile for each device family, use case or role the devices serve in the network.

To create ZTP profiles, go to **Device Management > ZTP Profiles**.

Step 3 Prepare ZTP device entries for the devices to be onboarded

Depending on how many devices you are onboarding, you can either prepare and import a CSV file or you can create device entries individually.

- a. Go to **Device Management > Devices**.
- b. Click the **Zero Touch Devices** tab. Then:

- To create a device entry file for many devices, click the **Import** icon and download the CSV template. Edit the template and add entries for each device you want to onboard. See the ZTP chapter for details on the file entries. Then click the **Import** icon again to import your device entry file.
- To create device entries one at a time, click the **Add** icon.

Step 4 Set up DHCP for Crosswork ZTP

Before triggering ZTP processing, you must update your organization's DHCP server configuration file with the IDs for your ZTP device entries and the paths to the image and configuration files stored in the ZTP repository. This allows Crosswork and DHCP to identify these ZTP devices and to respond correctly to each device's requests for connection to the network, and to download image and configuration files. For sample DHCP entries, see the ZTP chapter.

Step 5 Initiate ZTP processing to onboard the devices

Initiate ZTP processing by rebooting each of the devices to be provisioned: Power-cycle it, or press the chassis reset button.

Step 6 Monitor the ZTP processing status

You can monitor the progress of ZTP processing in the dashboard.

- Click **Home** in the main menu and take a look at the Zero Touch Provisioning dashlet.



- Click on the **View ZTP devices** link to view the status of individual devices.

Step 7 Verify your onboarded devices

Go to **Device Management > Devices**. Click the **Zero Touch Devices** tab. All of your onboard devices should be listed.

You may need to edit the information for some devices. Some of the information needed for a complete device record either is not needed in order to onboard the device, or not directly available through automation. For example, geographical location data defined using a set of GPS coordinates.

ZTP devices, after being onboarded, are automatically part of the shared Crosswork device inventory. You can edit them like any other device.



CHAPTER 8

Visualization of Native SR Path

This section explains the following topics:

- [Overview, on page 151](#)
- [Scenario: Troubleshoot Native SR IGP Paths Over Inter-AS Option C, on page 152](#)
- [Workflow: Native SR IGP Paths Troubleshooting, on page 153](#)

Overview

Objective

Visualize the actual path traffic flows physically through the topology map, even if traffic is on a native SR IGP path (not SR-policy) over inter-AS option C.

Challenge

Visualizing the native SR IGP path is often an operational challenge. Without access to a streamlined and simple-to-use interface, diagnosing and troubleshooting the native path requires you to repeatedly login to network devices without a solution to improve efficiency.

Solution

With the Path Query option, the objective is to visualize the native path using the traceroute SR-MPLS multipath command to get the actual paths between the source and the destination. With Cisco Crosswork Network Controller, a traceroute command runs on the source device for the destination TE-router ID and helps in retrieving the paths. By using native gRPC calls from the Crosswork server, you are able to get the paths from the device which assist in visualizing the native path through which the traffic flows. Since the traceroute command results in an operation that might take time to converge, Cisco Crosswork Network Controller provides an asynchronous user experience where you can send a request for such an operation and then be notified when the output is ready for inspection.

How Does it Work?

- Create a new path query, defining the headend and endpoint devices to find the available Native SR IGP paths.
- Visualize the available Native SR IGP paths as defined by the query on the topology map.
- Inspect the available paths and review the Output, Nexthop, Source, Destination, and Hop Index information.

- Create additional path queries as needed based on service type and instance and visualize the paths on the topology map.
- Troubleshoot any failed path queries.

Scenario: Troubleshoot Native SR IGP Paths Over Inter-AS Option C

Scenario Context

Visualization of the paths that traffic flows in is not readily available without manual tasks from different sources. Once attaining traffic-flow paths, the data is often out of date. Cisco Crosswork Network Controller supports the creation of Path Queries, which you define within the Crosswork GUI. This allows visualization of actual SR IGP paths between the source and destination on a topology map. Cisco Crosswork Network Controller provides an asynchronous user experience where the user is notified when results are ready for inspection. This facilitates rapid troubleshooting for issues with native traffic flows.

Assumptions and Prerequisites

- The device should have IOS XR version 7.3.2.
- The device should have gRPC (Remote Procedure Call) enabled. To check, run “show grpc” the in device and follow these steps:
 - For gRPC without a secure connection: If gRPC is showing as not enabled, enable gRPC using the following commands: `configure terminal; grpc; no-tls`.
 - For gRPC with a secure connection: Upload security certificates to Cisco Crosswork Network Controller in order to connect to the device using the following commands: `configure terminal; grpc`.
- Your Cisco Crosswork Optimization Engine server should have the devices imported with gNMI (Network Management Interface) capability and gNMI connectivity for the devices.
 - Make sure the credential profiles include connectivity information for gNMI. Go to **Device Management > Credential Profiles**. The Credential Profiles screen appears. Select a profile to edit. On the Edit Profile Devices screen, click + **Add Another**. For Connectivity Type, select **GNMI**. Add the User Name, Password, and Confirm Password information. Click **Save**.
 - Devices should have gNMI capability enabled in Cisco Crosswork Network Controller while attaching the device. Go to **Device Management > Network Devices**. Select the device to edit. The Edit Device Details screen appears. From the required Capability list, select **GNMI**. Click **Save**.
 - Devices should have the gNMI connectivity information enabled. Go to **Device Management > Network Devices**. Select the device to edit. On the Edit Device Details screen, under Connectivity Details, click + **Add Another**. For Protocol, select **GNMI** and add the IP Address / Subnet Mask information. Type the Port information and for Encoding Type, select **JSON**. Click **Save**.

Workflow: Native SR IGP Paths Troubleshooting

- Step 1** Select **Services & Traffic Engineering > Path Query**. The **Path Query** dashboard appears.
- Step 2** Click **+ New Query**. The **New Path Query** window appears, with query parameters displayed on the right and the topology map, with the mapped Device Groups, on the left.
- Step 3** Enter the device information in the required fields to find available native SR IGP service paths:
- Select the **Headend** device from the list. For this example, select **P-Edge-A1**
 - Select the **Endpoint** device from the list. For this example, select **P-Edge-B2**
- Step 4** Click **Get Paths**. The **Running Query ID** pop-up appears, indicating that Crosswork is processing your query. Path queries can take time to complete. While your query is running: if you want to review past queries select **View Past Queries**. Crosswork redisplay the **Path Query** dashboard. If you already had path queries in the list, you can view them while your new query continues to run in the background (indicated by the blue Running icon in the **Query State** column). When the new query's state turns green, click on the entry for it in the dashboard to view your results
- Step 5** When the pop-up changes to show that your query was successful, click **View Results**. The **Path Details** panel appears, with corresponding **Available Paths** details displayed below your **Headend** and **Endpoint** selections. Each found path will appear in its own section in the **Available Paths** portion of the panel. Meanwhile, the topology map on the left is filtered to display only the nodes and paths participating in the **Available Paths**.
- Step 6** Click on the **Available Paths** option whose details you want to review (for example, click **Path 0** or **Path 1**). Crosswork Status details for Output, Nexthop, Source, Destination, and Hop Index information. When you select one of the available paths, the map will update with the corresponding Device Groups topology mapping of Path 0 and Path 1.

Note Ensure that the **Show Participating Only** check box is selected in the top-right corner of the map.

Note There are three likely status outcomes to a path query. The screen captures below are independent examples not directly associated with the scenario's workflow:


- a. Non-Broken Path (path is complete):** Path Status shows as **Found** with path hop details and overlay shown.

The screenshot shows the Path Query interface with the following details:

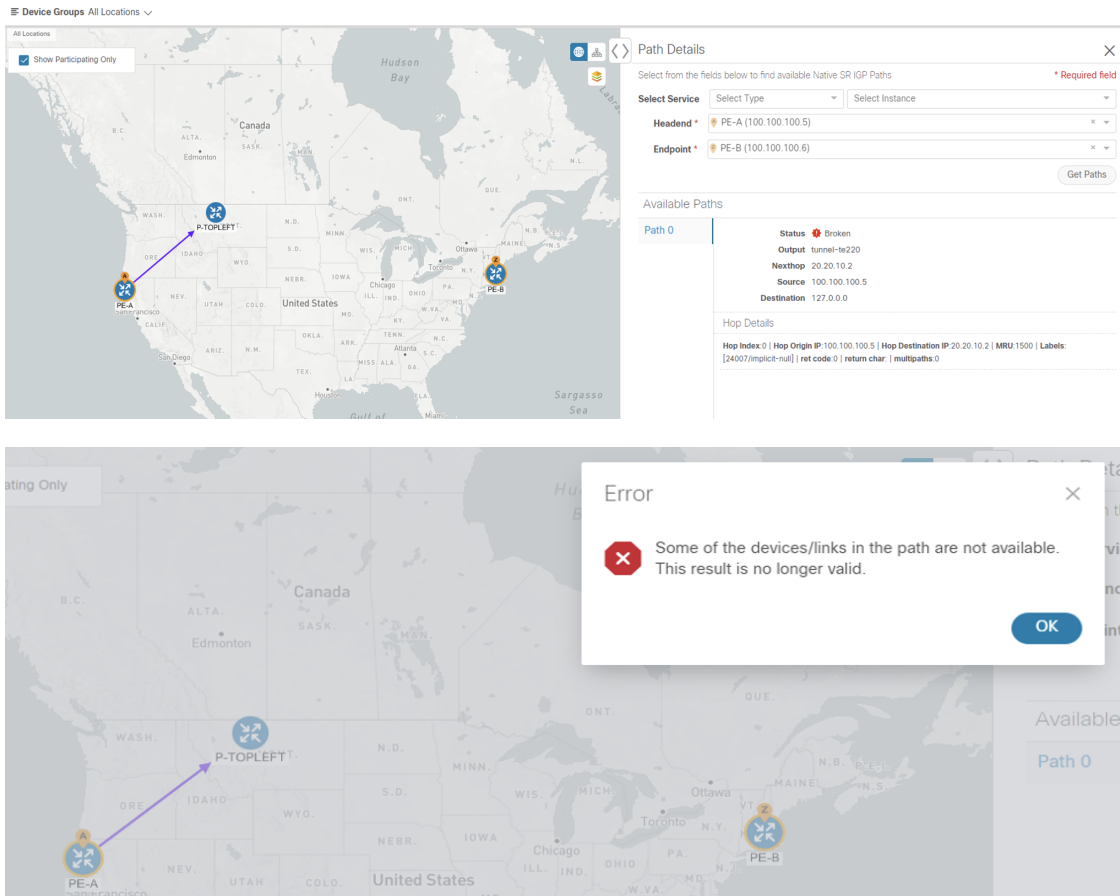
- Path Query Parameters:**
 - Select Service: L3vpn-Service
 - debug-issue-19-r2r9
 - Headend: R9 (9.9.9.9)
 - Endpoint: R2 (2.2.2.2)
- Available Paths:**
 - Path 0: Status Found
 - Path 1: Status Found
- Path 0 Details:**
 - Status: Found
 - Output: Bundle-Ether79
 - Nexthop: 10.7.9.1
 - Source: 9.9.9.9
 - Destination: 127.0.0.2
 - Hop Details:
 - Hop Index 0 | Hop Origin IP: 9.9.9.9 | Hop Destination IP: 10.7.9.1 | MRU: 1500 | Labels: [implicit-nut/36002] | ret code: 0 | multipaths: 0
 - Hop Index 1 | Hop Origin IP: 10.7.9.1 | Hop Destination IP: 1.5.7.5 | MRU: 9198 | Labels: [implicit-nut/36002] | ret code: 8 | return char: L | multipaths: 1
 - Hop Index 2 | Hop Origin IP: 1.5.7.5 | Hop Destination IP: 10.3.5.1 | MRU: 1500 | Labels: [17002] | ret code: 8 | return char: L | multipaths: 2
 - Hop Index 3 | Hop Origin IP: 10.3.5.1 | Hop Destination IP: 10.2.3.1 | MRU: 1500 | Labels: [implicit-nut] | ret code: 8 | return char: L | multipaths: 1
 - Hop Index 4 | Hop Origin IP: 10.2.3.1 | MRU: 0 | ret code: 3 | return char: | multipaths: 0

- b. Broken Path (Path is complete):** Path Status shows as Broken with path hop details and overlay shown.

Workflow: Native SR IGP Paths Troubleshooting




- c. **Broken Path (Path is not complete):** Path Status shows as Broken with path hop details partially shown (depending on gNMI output for traceroute – see Step 17 for troubleshooting details) and overlay details partially shown. An Error message will appear indicating that the devices and links are not available.



Step 7 Select **Services & Traffic Engineering > Path Query** to return to the **Path Query** dashboard.

Step 8 Ensure that the new path Query State column shows as completed with a green icon. The new path in the table will also show a Query ID link, both the corresponding Headend and destination Endpoint, and the Available Paths column will show 2 for both paths.

If a query state is broken, see the last step in the workflow for troubleshooting details.

Step 9 As needed, click on the **Query ID** link or click  and select **View Details** to again review the Path Details panel and map.

Step 10 Create additional path queries by selecting **Services & Traffic Engineering > Path Query**. The Path Query Dashboard appears where the previous path queries are listed by Query ID.

Note Make sure to set the **Automatically delete query older than every < X >** option within the number of hours needed from the Path Query Dashboard. The maximum number of hours provided is **24**.

Step 11 Click **New Query**. The New Path Query panel appears on the right with the mapped Device Groups panel on the left.

Step 12 For Select Service, select the Type from the list. In this example, select **L2VPN-SERVICE**.

By utilizing Select Service, when you later select the Headend and Endpoint, the options are conveniently identified according to the relevant VPN service type.

Step 13 For Select Service, select the Instance from the list. In this example, select **L2VPN_NM_P2P-NATIVE-210**.

The topology map will update to show the path between both servers. In this example, **P-Edge-B2** and **P-Edge-C3** are isolated on the map showing the logical path.

Step 14 Select the following from the list:

- a. Headend: **P-Edge-B2**.
- b. Endpoint: **P-Edge-C3**.

Step 15 Click **Get Paths**.

The Running Query ID pop-up appears.

Step 16 Click **View Results** when it becomes available. The Path Details panel appears with the corresponding Available Paths details, while the defined topology map appears with the available Native SR IG Paths on the left. This view shows the actual, physical hops between B2 and C3 that is carrying the traffic.

Step 17 To troubleshoot any Failed path queries appearing in the Path Query Dashboard's Query State column, select the "I" icon for error details.

In this example, the gNMI protocol is missing from the Connectivity Details for a previous path query with the Headend P-BOTTOMLEFT device and the Endpoint P-BOTTOMRIGHT devices. To troubleshoot the failed path query, do the following:

- a. Select Device Management > Network Devices.
- b. Find the device by Host Name and select the check box.
- c. Click the Edit icon at the top of the table. The Edit Device window appears.
- d. In this example, the Connectivity Details for Protocol is missing gNMI. Click + **Add Another** and type GNMI until it appears in the list. Select it.
- e. Enter the IP Address / Subnet Mask information and Port field information.
- f. Enter the Timeout field as **30**.

- g. In the Endcoding Type list, type **JSON** until it appears in the list. Select it and click **Save**.
 - h. Select Services & Traffic Engineering > Path Query. The Path Query Dashboard appears.
 - i. Click **New Query**. The New Path Query panel appears.
 - j. Select the following from the list:
 - 1. Headend device: **P-BOTTOMLEFT**.
 - 2. Endpoint device: **P-BOTTOMRIGHT**.
 - k. Click **Get Paths**. The Running Query ID pop-up appears.
 - l. Click **View Results** when it becomes available. The Path Details panel appears with corresponding Available Paths details, while the defined topology map appears with the available Native SR IG Paths on the left and is now in a Completed state.
-



CHAPTER 9

Provision, Visualize, and Analyze Tree Segment Identifier Policies in Multipath Networks

This section explains the following topics:

- [Overview, on page 157](#)
- [Scenario: Provisioning, Visualizing, and Analyzing Tree-SID Policies in a Point-to-Multipoint L3VPN Service, on page 158](#)

Overview

Allow users to provision and visualize Tree Segment Identifier (Tree-SID) Segment Routing policies easily and quickly before associating the policies with an L3VPN service model.

Objective

To provision, visualize, and update static Tree-SID policies within your network using Crosswork Network Controller and associate the (mVPN) policies with an L3VPN service model. By provisioning the Tree-SID policies using the Crosswork Network Controller UI and both visualizing and analyzing the multicast paths, root and leaf nodes, transit nodes, and view information about each link among the nodes, provides a holistic view of creating, visualizing, updating, and maintaining point-to-multipoint (P2MP) network configurations. These static Tree-SID policies can now be associated with an L3VPN service model and visualized and edited, as needed, using the Crosswork Network Controller UI.

Challenge

Keeping track of SR PCE and PE paths within networks is a challenge for video broadcasting and streaming service providers, who must use multipath protocols to replicate traffic and send it to different points in the network. Ensuring a high level of service quality forces providers to use difficult manual approaches to visualize, update and maintain their point-to-multipoint (P2MP) network configurations. This slows response to network problems and increases costs.

Solution

Tree-SID is a method of implementing tree-like multicast flows over a segmented routing network. Using Tree-SID, an SDN controller (a device running SR-PCE using PCEP), calculates the tree. Each node (device) in the tree has a specific role in routing the multicast data through the tree. These roles include Ingress for the root or headend node, Transit or Bud for midpoint nodes that are not leaf nodes, and Egress for destination leaf nodes. The tree itself is assigned a single SID label, which represents all of the tree segments and devices

in it. The SDN controller is highly flexible, as it understands the segmentation and can construct routing paths using any kind of constraints that network architects can specify.

The most interesting use case for constraint-based Tree-SID use is where routers are configured to deliver two P2MP streams with the same content over different paths. Here, the multicast flow is forwarded twice through the network, each copy following a unique path. The two copies never use the same node or link to reach the destination, reducing packet loss due to network failures on any one of the paths.

By using Crosswork Network Controller, you can now create Static Tree-SID policies using the UI, associate Static mVPN Tree-SID policies with a provisioned L3VPN service, visualize, analyze, and edit or delete your Tree-SID policies to actively manage your multicast network.



Note Static and Dynamic mVPN Tree-SID policies can be associated with a L3VPN service model. In this workflow tutorial, only a Static mVPN Tree-SID policy will be associated with a L3VPN service model, visualized, and analyzed.

How Does it Work?

- Create a Static Tree-SID policy using the Crosswork Network Controller UI
- Visualize and validate the new Static Tree-SID policy
- Associate your Static mVPN Tree-SID policy with an L3VPN service model (or import an existing static or dynamic Tree-SID policy)
- Visualize and analyze the performance details of your Static mVPN Tree-SID paths and nodes within the L3VPN service model
- Edit your existing Static mVPN Tree-SID policy to enhance performance or correct issues with your Tree-SID L3VPN service model

Scenario: Provisioning, Visualizing, and Analyzing Tree-SID Policies in a Point-to-Multipoint L3VPN Service

Scenario Context

Without Crosswork Network Controller, provisioning and visualizing Tree Segment Identifier (Tree-SID) point-to-multipoint traffic flows is possible only by using manual tasks from different sources. Restriction to manual tasks means that the creation of Tree-SID policies, associating a policy with an L3VPN service model, visualization, and editing of the policy and service is significantly hampered. By using Crosswork Network Controller, you can sidestep the time loss between manual setup and the visualization of the traffic flow paths and avoid data that is often out of date with manual configurations. Crosswork Network Controller supports both creation and discovery of the Tree-SID segmentation directly from network configurations and displays the data flow map immediately. This facilitates rapid troubleshooting for issues with Tree-SID traffic flows.

Crosswork Network Controller's topology services uses PCE topology and LSP data to discover and visualize pre-configured Tree-SID policies in your network. The PCE controller manages the layer 3 topology, LSP and Tree-SID data using BGP link state, and supports initial discovery and notifications for the Tree-SID trees. Static Tree-SID policies can also be configured and later associated with newly created, or previously configured, L3VPN services directly in the Crosswork Network Controller's UI. Likewise, based on the health

of the service and policies, editing capabilities are also performed using the UI to troubleshoot and optimize models operations.

Assumptions and Prerequisites

If your network has PCE and Tree-SID policies already configured on your devices, this workflow assumes, at a minimum, the following basic configuration options:

1. On all nodes involved in the Tree-SID path, irrespective of role:
 - a. Enable Path Computation Element Protocol (PCEP)
 - b. Enable Computation Client (PCC)
2. Under SR-PCE, on end points: Configure a P2MP SR static or dynamic Policy.
3. On all root and leaf nodes:
 - Enable multicast routing
 - Configure `interface vrf <vrf-number>`
 - Configure `router bgp` on topo nodes and PCE. On corresponding neighbors between PCE and PCC nodes, mention the configured `interface vrf <vrf-number>`.
 - Configure `route-policy <vrf-number>` and `PASS_ALL`
 - Under segment routing traffic engineering: Configure `ODN color <same as vrf-number>`
4. On all leaf nodes only: Configure `router PIM`, `route-policy TREESID_CORE`.

Step 1 Create a Static Tree-SID Policy

If you are using preconfigured Static or Dynamic Tree-SID policies already configured on your devices, skip to Step 2 in the workflow. If you are configuring Tree-SID policies using the Crosswork Network Controller's UI, this task first creates a Static Tree-SID policy, each representing a leaf or root node, before you have the option to associate the policies with a L3VPN service model that can be visualized and edited as necessary:

Step 1 Go to **Services & Traffic Engineering > Traffic Engineering**.

The logical map opens and the Traffic Engineering panel is displayed to the right of the map.

Step 2 In the Traffic Engineering panel, select the **Tree-SID** tab.

The Traffic Engineering Tree-SID Policy screen appears.

Step 1 Create a Static Tree-SID Policy

Traffic Engineering Refined By

SR-MPLS SRv6 **Tree-SID** RSVP-TE

6 0 6 0 6 0

Total Dynamic Static Admin Down Oper Up Oper Down

Tree-SID Policy Selected 0 / Total 6

[+ Create](#) [📄](#)

<input type="checkbox"/>	Ro...	Ro...	Name	Tre...	Label	Admi...	Oper ...	Actions
<input type="checkbox"/>	xrv9k...	192.1...	Disney	-	152001			...
<input type="checkbox"/>	xrv9k...	192.1...	MY_F...	-	15200			...
<input type="checkbox"/>	xrv9k...	192.1...	Tree-...	-	800			...
<input type="checkbox"/>	xrv9k...	192.1...	Tree-...	-	900			...
<input type="checkbox"/>	xrv9k...	192.1...	tesd	-	23			...
<input type="checkbox"/>	xrv9k...	192.1...	tesdt	-	34			...

Step 3 Click + Create.

The New Tree-SID Policy (Static) screen appears.

New Tree-SID Policy (Static) * Required Field

Name *

Tree-SID Label *

Root *

Selected - None

Leaf (s) *

Selected - None

[+ Add Another](#)

Optimization Objective*

LFA FRR

Enable Disable

Constraints

Affinity

[+ Add Another](#)

Step 4 To enter or select the required Static Tree-SID policy values, do the following:

- a) After providing a name for your new Static Tree-SID policy, in the Tree-SID Label field, assign the MPLS label associated with the Tree-SID policy (for example: **152001**).
The Tree-SID Label must be in the range from 16 to 1048575.
- b) In the Root field, enter the host name (for example: **xrv9k-26**) or select a node on the map or an existing device in the list. As you type or select the Root information, a Root label for the selected node appears on the map. Only PCC nodes with PCEP session to PCE can be added as a Root node.
- c) In the Leaf field, enter the host name (for example: **xrv9k-24**) or select a node on the map. As you type, or select, the Leaf information, Leaf label(s) for the selected nodes appear on the map.

Click + **Add another** to add additional constraints (for example: **xrv9k-27**).

- d) For Optimization Objective, select one of the following constraints: Interior Gateway Protocol (IGP) Metric; Traffic Engineering (TE) Metric; Latency (for example: **IGP**).
- e) For LFA FRR, either select **Enable** or **Disable** (for example: **Enable**).

By selecting Enable, the Loop Free Alternate Fast Reroute (LFA FRR) is enable on all of the nodes in the distribution tree.

- f) For additional Constraints, select one of the following Affinity options: **Exclude-Any**, **Include-Any**, **Include-All**.

In addition, from the Select or Create Mapping drop-down list, click **Manage Mapping**. The Affinity Mapping dialog box opens. For more information on Affinities, see the Configure Link Affinities section in the Crosswork Optimization Engine guide.

Affinity Mapping

+ Create

Name	Bit Position (0-31)	Actions

There are no TE link Affinities here, may be you should create new one.

Done

- g) For Affinity Mapping, type a Name (color) of the mapping and enter the Bit Position (**0 – 31**). Enter the same bit position that is used on the device interface. Click **Done**.

To create additional constraints, click + **Create**.

Step 1 Create a Static Tree-SID Policy

New Tree-SID Policy (Static) * Required Field

Name *

Tree-SID Label * ⓘ

Root * ⓘ

Selected - xrv9k-26 [192.168.0.26] [2001:192:168::26] [Edit](#)

Leaf (s) *

Selected - xrv9k-24 [192.168.0.24] [2001:192:168::24] ⓘ [Edit](#)

Selected - xrv9k-27 [192.168.0.27] [2001:192:168::27] ⓘ [Edit](#)

[+ Add Another](#)

Optimization Objective*

LFA FRR ⓘ

Enable Disable

Constraints

Affinity

<input type="text" value="Exclude-Any"/>	<input type="text" value="Select or Create Mapping"/>	<input type="button" value="Delete"/>
<input type="text" value="Include-Any"/>	<input type="text" value="Select or Create Mapping"/>	<input type="button" value="Delete"/>
<input type="text" value="Include-All"/>	<input type="text" value="Select or Create Mapping"/>	<input type="button" value="Delete"/>

[+Add Another](#)

- h) To commit the policy, click **Provision** to activate the policy on the network.

The newly provisioned Tree-SID policy may take some time to appear in the Tree-SID table depending on the network size and performance. The Tree-SID table is auto refreshed every 30 seconds. Once the request is successful, either select **View Tree-SID Policy List** or **Create New** to add additional policies. If you select **View Tree-SID Policy List**, the Tree-SID Policy screen appears showing the newly created policy in the table.

Step 2 Visualize and Validate the new Static Tree-SID policy

Step 1

Select the root Tree-SID policy check box from the list. In this example, select **xrv9k-26**.


Tree-SID Policy Selected 1 / Total 6

+ Create 📄

	Root ...	Root IP	Name	Tree ID	Label	Admin ...	Oper S...	Actions
<input checked="" type="checkbox"/>	xrv9k-...	192.16...	Disney	-	152001	⬆️	⬆️	...
<input type="checkbox"/>	xrv9k-...	192.16...	MY_Fl...	-	15200	⬆️	⬆️	...

If there is a large number of policies in the table, filter by Root IP, Name, Label, or other parameter, to help locate the policy you want to visualize.

In the map, you will see the selected Tree-SID policy as an overlay on the topology. It shows a representation of the Tree-SID policy routes, with icon flags indicating the root **R** node (**xrv9k-26**, also known as the ingress device) and the two leaf **L** nodes (**xrv9k-24** and **xrv9k-27**, also known as egress devices), with intermediary transit nodes between them. Administrative and operational status for each node is shown in the table.

Note Use the buttons at the top right of the logical map to toggle between the Logical Map and the Geo Map  views.

Step 2

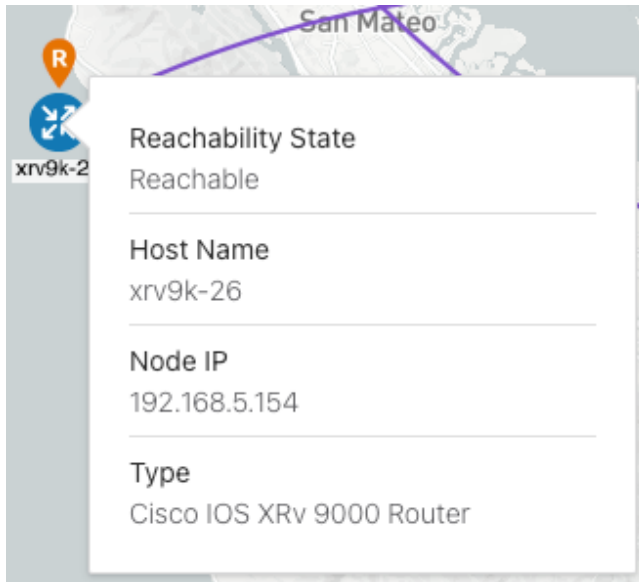
Select the **Geo Map** button to view the selected Tree-SID service topology overlaid on a world map.

Step 3

In the map, select the **Show: Participating Only** check box to hide underlay devices that are not participating in the selected Tree-SID policy. Then use your mouse to hover over the **xrv9k-26** root device to view its corresponding Reachability State, Host Name, Node IP, and device Type.

Check any participating Tree-SID device in the same fashion to view their corresponding details.

Step 2 Visualize and Validate the new Static Tree-SID policy

Step 4 In the map, click **xrv9k-24**.

The Device Details screen opens showing **xrv9k-24** information organized by Summary and Routing in the Details tab, and PCEP Sessions in the Traffic Engineering tab.

Device Details ✕

Details | Links | Traffic Engineering

Summary ^

Host Name	xrv9k-24
Reachability	✔ Reachable
IP Address	192.168.5.152
Geo Location	Latitude 37.621300, Longitude -122.379000
Device Type	🌐 Router
Device Group	Location > All Locations > Unassigned Devices
Product Type	Cisco IOS XRv 9000 Router
Connect To Device	🔑 SSH IPv4
Last Update	09-Oct-2023 12:32:35 PM PDT

Routing ^

TE Router ID	192.168.0.24
IPv6 Router ID	2001:192:168::24
ISIS System ID	0000.0000.0004 Level-2
ASN	65000

Device Details

Details Links Traffic Engineering

General SR-MPLS SRv6 Tree-SID RSVP-TE Flex Algo

Collapse All

IGP: Domain ID: 1000, ISIS System ID: 0000.0000.0004, Level: 2

SR-MPLS


SRv6

PCEP Sessions

PCE : 172.27.226.126, PCC/Source - 192.168.0.24

Stateful	true
Source Address	192.168.0.24
Capability Instantiate	true
Capability SR	true
PCE Address	172.27.226.126
Capability Update	true
MSD	10

Step 5 Click **X** in the top-right corner to return to the Tree-SID Policy table to close the Device Details screen and then again select the **Tree-SID** tab.

Step 6 In the Tree-SID Policy list for the selected **xrv9k-26** device, click  in the **Actions** column and select **View Details** to drill down to a current and detailed view of the Tree-SID policy.

The Tree-SID Policy Details screen appears.

Note To view all of the Tree-SID Policy Details, click **See more**.

Note When viewing Tree-SID Policy Details, if a Source Node is not available, a warning icon and message appear next to the Oper Status field (hover your mouse over the warning icon), detailing where the connection issue resides. For example:

The screenshot displays the 'Tree-SID Policy Details' window. At the top, there are tabs for 'Current' and 'History'. Below the tabs, the policy details are listed:

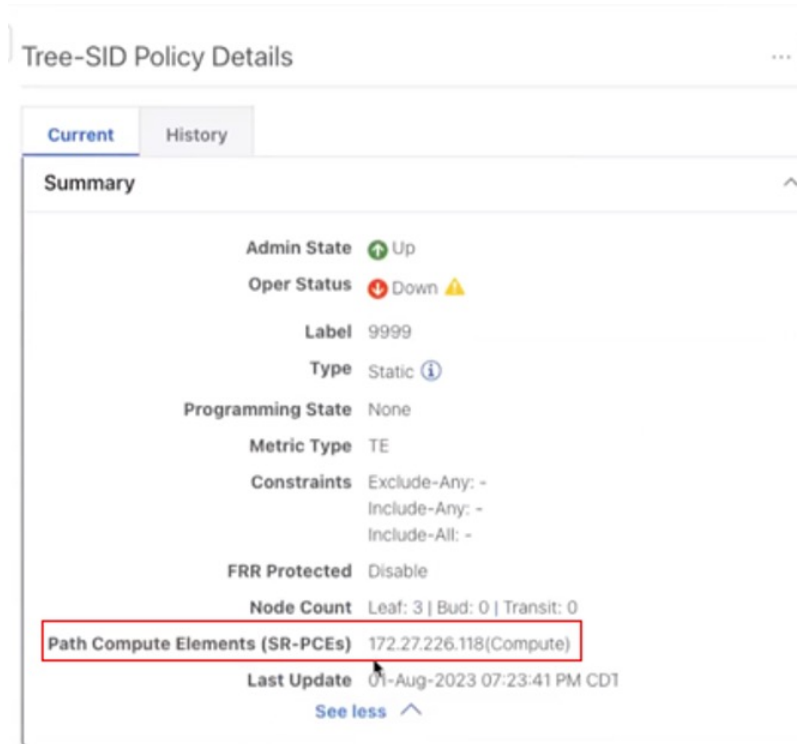
- Root:** xrv9k-12 | Root IP: 192.168.0.2
TE RID: 192.168.0.2 | IPv6 RID: 2001:192:168::2
- Name:** Test-TS2
- Tree ID:** - ⓘ

A 'Summary' section is expanded, showing the following details:

- Admin State:** Up (green arrow icon)
- Oper Status:** Down (red arrow icon) ⓘ. A tooltip is visible over the warning icon, stating: 'source node 192.168.0.2 not connected via PCEP'.
- Label:** 9999
- Type:** Static ⓘ
- Programming State:** None
- Metric Type:** TE
- Constraints:** Exclude-Any: -
Include-Any: -
Include-All: -
- FRR Protected:** Disable

At the bottom of the summary section, there is a 'See more' link with a downward arrow.

Note A (Compute) label, next to the SR-PCE field, details the SR-PCE used to create the policies. For example:



- Step 7** In the Tree-SID path section, click **Expand All** to view Tree-SID path names and IPs for the **xrv9k-24** and **xrv9k-27** leaf nodes. The list also shows details for the corresponding Root node, all Transit nodes, the two Leaf nodes, and their Egress Link's Local IP and Remote IP information.
- Step 8** Deselect the **xrv9k-22** check box to see Tree-SID path details for **xrv9k-24** and **xrv9k-27** devices only. The topology updates to show only the selected **xrv9k-24** and **xrv9k-27** Tree-SID routes.
- Step 9** Click **X** in the top-right corner to return to the Tree-SID Policy table.
- Step 10** Select the Root IP Tree-SID policy **xrv9k-26** check box from the list. Make sure the geographical map option is selected. The geographical map updates to show the policy and its disjunct routes. You can click on individual links and get details on the Tree-SID policies in which each link participates.

Step 3 Associate the Static Tree-SID Policy with the newly created L3VPN service model

- Step 1** Go to **Services & Traffic Engineering > Provisioning (NSO)**. The Provisioning screen appears showing available Services/Policies.
- Step 2** Select **L3VPN > L3vpn-Service**. The L3VPN > L3vpn-Service table appears.

Step 3 Associate the Static Tree-SID Policy with the newly created L3VPN service model

- Step 3** To create a new L3vpn-Service, click the **+** symbol.
The Create L3VPN > L3vpn-Service screen appears.
- Note** You may also click the **↓** symbol to import an existing L3vpn-Service.
- Step 4** In the vpn-id field, type the unique ID for the service (for example: **MVPN-TREE-SID-119**) and click **Continue**.
- Note** This identifier has a local meaning (such as within a service provider network).
- Step 5** In the vpn-service-topology drop-down list, select **custom** to define the service topology.
- Note** Point-to-point VPN service topology is not supported.
- Step 6** Expand the vpn-instance-profiles section and click the **+** symbol to add the profile ID.
The vpn-instance-profiles panel appears.
- Step 7** In the profile-id field, type the VPN instance profile identifier (for example: **MVPN-TREE-SID-119**) and click **Continue**.
The vpn-instance-profiles panel refreshes with additional fields to fill.
- Step 8** In the Rd-choice section, enter the directly-assigned rd that indicates an RD value that is explicitly assigned (for example, **0:70:70**).
- Step 9** For address-family, click the **+** symbol. The address-family panel appears and select **ipv4** from the address-family list and click **Continue**.
The address-family {ipv4} panel updates with vpn-targets section included.
- Step 10** For vpn-target, click the **+** symbol so to signify the VPN target id and route-target-type.
The vpn-target panel appears.
- Step 11** In the id field, enter the id (for example: **91**) and click

Continue..

- Step 12** In the vpn-target{91} panel, select the route-target-type drop down list and select **both**.
The address-family{ipv4} panel updates showing the vpn-target id (as **91**) and route-target-type (as **both**).
- Step 13** In the vpn-target{91} panel for route-targets, click the **+** symbol and type the route-target (for example, **0:70:70**) and click **Continue**. Click **X** to close the panel.
The route-target table updates with the new information. Click **X** in the top right to close all of the remaining panels.
Adding the vpn-instance-profiles is now complete.
- Step 14** Select **multicast** and then ipv4 to expand both sections.
- Step 15** Expand the mvpn-ipmsi-tunnel-ipv4 section and select **static-sr-mpls-p2mp** from the tunnel-type list.
The Enable ipv4 toggle is now switched on.

Note The sr-mpls-p2mp selection in the list is for a Dynamic Tree-SID policy.

vpn-instance-profile ⓘ ⓘ

+ **✎** **✖**

profile-id

MVPN-TREE-SID-119

vpn-nodes ⓘ

service-assurance ⓘ

probes ⓘ

multicast ⓘ

ipv4

Enable ipv4

static-sr-mpls-p2mp ⓘ

+ **✎** **✖**

policy-name

mvpn-ipmsi-tunnel-ipv4 ⓘ




tunnel-type ⓘ

static-sr-mpls-p2mp **⌵** ⓘ

mvpn-spmsi-tunnels-ipv4 ⓘ

Commit changes **Dry Run** **Cancel**

- Step 16** For static-sr-mpls-p2mp, click the **+** symbol.
The static-sr-mpls-p2mp panel appears.




- Step 17** In the policy-name field, type the previously created Static Tree-SID policy name (for example: **xrv9k-26**) and click **Continue**.
The static-sr-mpls-p2mp{Static-xrv9k-26} panel updates.
- Step 18** In the sr-p2mp-policy area for the group-address, click the  symbol to add the address.
The group-address panel appears.
- Step 19** In the Address field, type the IPv4 static multicast group address (for example: **1.1.1.1**) and click **Continue**.
The group-address{1.1.1.1} panel refreshes. Click **X** at the top right to close any remaining panels.
- Step 20** Click the  symbol in the multicast > ipv4 subsection to add the other policy name.
The static-sr-mpls-p2mp panel appears.
- Step 21** In the policy-name field, type the other previously created Static Tree-SID policy name (for example: **xrv9k-24**) and click **Continue**.
The static-sr-mpls-p2mp{xrv9k-24} panel updates.
- Step 22** In the sr-p2mp-policy area for the group-address, click the  symbol to add the address.
The group-address panel appears.
- Step 23** In the address field, type the IPv4 static multicast group address (for example: **2.2.2.2**) and click **Continue**.
The group-address{2.2.2.2} panel refreshes. Click **X** at the top right to close any remaining panels.
You have now successfully mapped the static Tree-SID policy to the L3VPN multicast service model. Next, you must add the VPN node details.
- Note** For advanced configurations, you may select mvpn-spmsi-tunnels-ipv4 subsection under the multicast section to define the tunnel-type, switch-wildcard-mode, switch-threshold, per-item-tunnel-limit, group-acl-ipv4 details.

multicast ⓘ

ipv4

Enable ipv4

static-sr-mpls-p2mp ⓘ

policy-name

xrv9k-24

xrv9k-26




mvpn-ipmsi-tunnel-ipv4 ⓘ

tunnel-type ⓘ

static-sr-mpls-p2mp

mvpn-spmsi-tunnels-ipv4 ⓘ

Step 4 Add the VPN nodes

- Step 1** In the vpn-nodes section, click the  symbol to add your VPN nodes set up in the Static Tree-SID policy (**xr9k-26**, **xr9k-24**, and **xr9k-27**).
- The vpn-node panel appears so to add the vpn-node-id.
- Step 2** From the vpn-node-id drop down, select the first of the VPN node (for example: **xr9k-26**) and click **Continue**.
- The vpn-node{xr9k-26} panel updates with additional fields.
- Step 3** In the Local-as field, type **65000**.
- Step 4** In the active-vpn-instance-profiles section, click the  symbol to add the VPN instance profile ID.
- Step 5** In the profile-id drop down list, the previously added profile ID appears. Select it (for example: **MVPN-TREE-SID-119**), click **Continue** and click **X** to close the panel.
- Step 6** In the vpn-node{xr9k-26} panel, select the vpn-network-accesses section and click the  symbol to add the vpn-network-access ID. In the Id field, add a number (for example: **1**) and click **Continue**.
- The vpn-network-access{1} panel updates with additional fields.
- Step 7** In the interface-id field, type the identifier for the physical or logical interface (for example: **Loopback70**).
- The identification of the sub-interface is provided at the connection level and/or the IP connection level.
- Step 8** In the ip-connection section, select the ipv4 subsection and in the local-address field, type the IP address used at the provider's interface (for example: **70.70.10.1**).
- Step 9** In the prefix-length field, type **30**.

Step 4 Add the VPN nodes

The subnet prefix length is expressed in bits. It is applied to both local and customer addresses.

The screenshot shows the configuration for a VPN network access profile. The left pane displays the 'vpn-network-accesses' section with a table containing one entry: 'vpn-network-access' with ID '1'. The right pane shows the configuration for 'vpn-network-access(1)'. Key fields include: 'local-address' set to '70.70.10.1', 'prefix-length' set to '30', and 'multihop' set to '11'. The 'routing-protocols' section is currently empty.

Step 10 In the routing-protocols section, click the **+** symbol to add the unique identifier for the routing protocol. In the Id field, type **bgp** and click **Continue**.

The routing-protocol{bgp} panel appears.

Step 11 In the type drop down list, select **bgp-routing**.

The routing-protocol{bgp} panel refreshes with additional sections.

Step 12 In the bgp section, for the peer-as field, type **70** to indicate the customer's ASN when the customer requests BGP routing, and in the address-family drop down list, select **ipv4**. This node contains the address families to be activated.

Note If you select dual-stack, it means that both ipv4 and ipv6 will be activated.



Step 13 In the multihop field, type **11** to describe the number of IP hops allowed between a given BGP neighbor and the PE.

The screenshot shows the configuration for the 'routing-protocol{bgp}' section. The left pane shows the 'routing-protocols' section with a table containing one entry: 'bgp'. The right pane shows the configuration for 'routing-protocol{bgp}'. Key fields include: 'peer-as' set to '70', 'address-family' set to 'ipv4', 'multihop' set to '11', and 'redistribute-connected' set to 'enable'. The 'neighbor' section is currently empty.

Step 14 For neighbor section, click the **+** symbol and in the Neighbor field, type the device address (for example: **70.70.10.2**) and click **Continue**.

Step 15 For redistribute-connected section, click the **+** symbol and from the address-family drop down list, select **ipv4** and click **Continue**.

The redistribute-connected{ipv4} panel appears.

- Step 16** In the enable field, select **true** to enable the redistribution of connected routes.
Close all panels (click X in the top right corner) until the Create L3VPN > L3vpn-Service screen appears.
- Step 17** In the vpn-nodes section, you will see xrv9k-26 listed in the vpn-node table. Select **xrv9k-26** and select the  symbol.
The vpn-node{xrv9k-26} panel appears.
- Step 18** Select the multicast section and click the  symbol to add the mapping of the policy for each node.
The static-sr-mpls-p2mp panel appears.
- Step 19** For the policy-name drop down list, select the policy you want to add to this node (either the source or the receiver).
Select **xrv9k-24** as a receiver and click **Continue**.
The static-sr-mpls-p2mp{xrv9k-24} panel updates with additional fields.
- Step 20** For the role drop down list, select **receiver**.
Close all additional panels (click X in the top right corner) until the Create L3VPN > L3vpn-Service screen appears.

static-sr-mpls-p2mp{xrv9k-24}

policy-name * ⓘ


xrv9k-24

role*

receiver

- Step 21** Repeat steps 1 – 20 to add the other two VPN nodes set up in the Static Tree-SID policy: **xr9k-24** and **xr9k-26**.
- Step 22** After all of the VPN nodes have been added, click **Commit changes**.

Step 5 Visualize and Edit the Static mVPN Tree-SID Policy's L3VPN service model

- Step 1** Go to **Services & Traffic Engineering > Provisioning (NSO)**.
The Provisioning screen appears showing available Services/Policies.
- Step 2** Select **L3VPN > L3vpn-Service**.
The L3VPN > L3vpn-Service table appears.
- Step 3** Locate the newly created L3VPN service ID in the table (**MVPN-TREE-SID-119**) and in the Actions column, click  and select **Config View**.
The Configured Data pop-up screen appears.

Step 5 Visualize and Edit the Static mVPN Tree-SID Policy's L3VPN service model



Step 4 In the Configured Data pop-up screen, review the data configuration and click **Copy To Clipboard** if you want to save a copy, or click **Cancel** to exit.

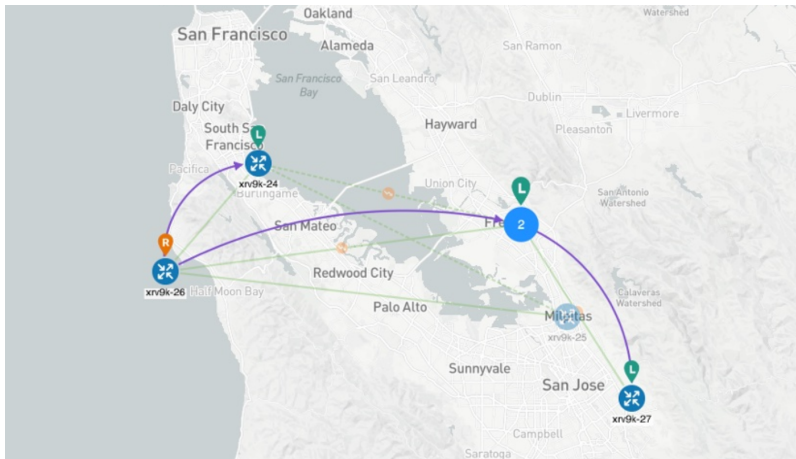
Step 5 To view the new Static mVPN Tree-SID policy associated with the L3VPN service model, click the name of the VPN Id in the table or in the Actions column, click and select **View**.

The Service Details screen appears with the geographical map showing the newly created L3VPN service and the associated nodes: xrv9k-26, xrv9k-24, xrv9k-27. On the right, the Service Details panel shows the details of the MVPN-TREE-SID-119 service model.

Step 6 In the Service Details panel, select the **Transport** tab to view the Tree-SID Policy information.

Step 7 In the table, select the check box next to xrv9k-26.

In the geographical map, the policy will appear showing the one Root, or source, node (xrv9k-26) and the two Leaf, or receiver, nodes (xrv9k-24 and xrv9k-27).



Step 8 Select the second check box next to xr9k-24.

The geographical map updates.

Step 9 Use your mouse to hover over the Tree-SID policy names in the table. Depending which policy your mouse hovers over, the geographical map will show the designated path(s) between the nodes so to differentiate them from each other.

Step 10 For the first policy in the table, in the Actions column, click and select **View Details**.

The Tree-SID Policy Details panel appears showing the policy's details such as the Name, a Summary section, and the Tree-SID path information that can be expanded to show additional detail. You may also select the History tab to view historical information for the policy.

Step 11 To edit, or add additional policies, go to **Service & Traffic Engineering > Provisioning (NSO)**, and select **L3VPN > L3vpn-Service**.

Step 12 For your L3VPN service, in the Actions column, click and select **Edit**.

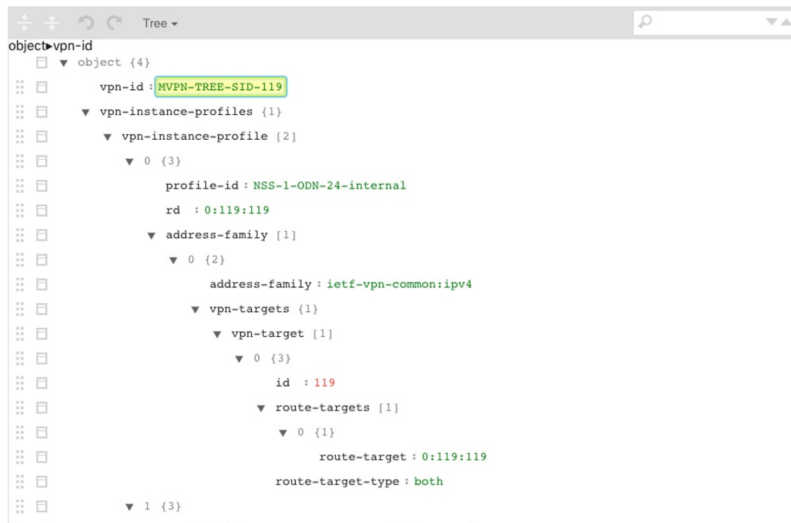
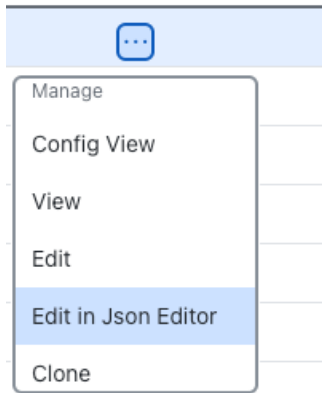
The Edit L3VPN > L3vpn-Service screen appears where you can make additional updates (such as adding VPN nodes so to replace a degraded path so to give it a different route) and modifications to existing details that make up the service.

While editing, to show all or hide the multiple fields that make up the service configuration, select the **Show all fields** toggle at the top right. Click on the toggle for Show all fields to be on. Click the toggle again for the Show all fields to be off, showing just a subset of the fields.

Step 13 In addition, from the **L3VPN > L3vpn-Service** screen, click in the Actions column and select **Edit in Json Editor** for your L3VPN service.

The json Configuration editor appears. Using the json Configuration editor, you can highlight different details that make up the service configuration and edit them directly in the json editor.

Actions



Step 14 Once completed, either click **Commit** to initiate the changes and update the service's configuration or click **Cancel**.

Summary and Conclusion

As we observed, you can provision new Static Tree-SID policies within the Crosswork Network Controller UI. Once provisioned, you can use the Tree-SID tab and its associated map to visualize Tree-SID defined routes, identify disjunct policy routes, and identify problems with transit nodes, interfaces and links that may affect traffic from the Root to the Leaf nodes. In addition, once the Tree-SID policies are associated with an L3VPN service model, similar capabilities are at hand to visualize and analyze Static Tree-SID policies associated with an L3VPN service model and edit in dynamic ways that improve efficiency, accuracy, and ease of use.



CHAPTER 10

Transport Slice Provisioning

This section explains the following topics:

- [Overview, on page 177](#)
- [Scenario: Implement an Any-To-Any L3 eMBB Slice, on page 183](#)
- [Step 1 Create a Slice Template Catalog Entry, on page 186](#)
- [Step 2 Add Service Assurance into the Slice Template Catalog using the NSO CLI \(optional\), on page 188](#)
- [Step 3 Create the Transport Slice Instance, on page 189](#)
- [Step 4 Deploy a Slice using NSO CLI \(optional method\), on page 198](#)
- [Step 5 Visualize and Validate the New Slice Deployment , on page 200](#)
- [Summary and Conclusion, on page 203](#)

Overview

Objective

Simplify Transport Service provisioning by focusing on the service's SLA intents (the "what" instead of the "how"). This implies a service-oriented view, leveraging the concepts of software-defined networking (SDN).

Challenge

Service providers face ever-growing demands from end users for highly customized, flexible network services with very different, sometimes contradictory, service level requirements: support for highly mobile smart cars, ultralow-latency AR and mobile gaming applications, secure machine-to-machine communication in logistics and manufacturing, and so on. Modern software-defined network (SDN) traffic engineering technologists have responded with a host of innovative protocols and features that offer many ways to engineer network traffic to meet these special needs. Crosswork support for these approaches, such as SR-TE services, Tree-SID and Local Congestion Mitigation, are featured elsewhere in this Guide.

The advent of 5G mobile networking has accelerated this trend, resulting in a new approach to network architecture: network slicing. This still-emerging standard enables network engineers to slice the 5G network's bandwidth into tranches that prioritize some services over others, instead of treating it as a single, monolithic network. The network engineer can design each network slice around the needs and intents of its users, allocating speed, latency, throughput and other resources to each slice as required. CNC delivers a rich and customizable tool set to make deploying these slices easier. When combined with Service Health, it provides the added ability to easily monitor the health of these services. The provider organization can then offer the slice itself as a service, helping to broaden the range of service offerings.

But how to make these services easy to provision? The design and coding of the sophisticated traffic engineering services that underlie network slicing require the skills of experienced network architects and deep knowledge of each provider's existing network infrastructure. Without some form of automation that allows line operators to instantiate the designed slices quickly and easily, network slices might remain a type of custom configuration, achievable only for a small set of important users, instead of a scalable commodity providers can monetize.

This is an evolving standard. At present, the Crosswork solution addresses the Transport-level Network Slice Management Functions (NSMF) only.

Solution

Cisco Crosswork Network Controller offers direct support for network slicing at the OSI transport layer. Using this solution, network engineering experts can design slices around customer intents and then add them to a catalog. Network line operators can then simply pick the slice intent that best meets the customer's needs, specify the slice endpoints, and (where needed) set any custom constraints or options built into the chosen slice.

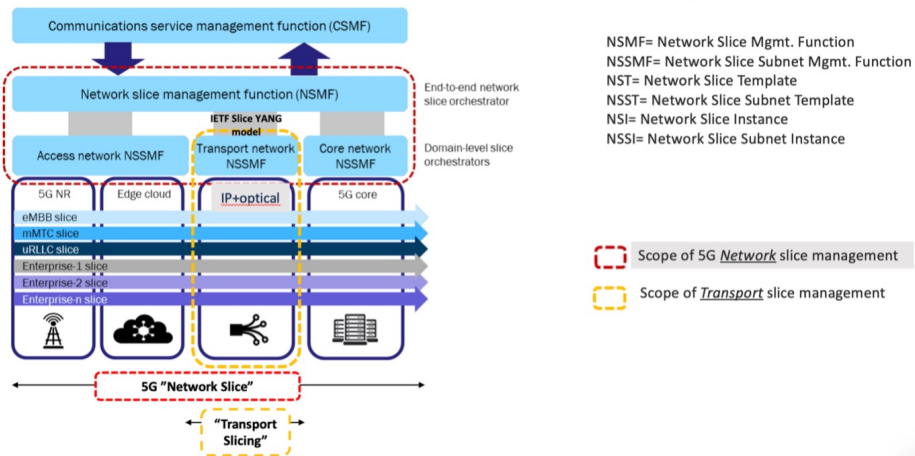
This is Cisco's initial offering in the network slicing arena, chosen because of our company's strengths at the transport layer. At present, the Crosswork solution provides a large catalog of slice template examples and an extensive customization for each template. This document offers a scenario that you can follow to create and (optionally) monitor a network slice.

How Does Transport Slicing Work?

It's important at the outset to understand the difference between 5G network slicing and generalized transport slicing. When operational, a 5G network slice is an end-to-end solution crossing multiple sub-domains. The 5G network authority 3rd Generation Partnership Project (3GPP) refers to each end-to-end network slice operating on the network as the Network Slice Instance (NSI). Each NSI is a unique entity, provisioned in the network with a set of Service Level Requirements chosen from a set of pre-created Network Slice Templates (NST).

All NSIs must be orchestrated by a controller called the Network Slice Management Function (NSFM). The NSMF in turn communicates with sub-domain controllers, referred to as Network Slice Subnet Management Functions (NSSMF). Each NSSMF in turn provisions the corresponding domain-specific slice instance across its own sub-domain boundaries (called a Network Slice Subnet Instance or NSSI). For the Transport domain, the IETF has defined the NSSI as an "IETF Network Slice" in order to differentiate slices in the transport domain from slices bridging other domains. The space occupied by transport slicing in this hierarchy is shown in the illustration below, where the CNC solution will provide the NSSMF functionality for the Transport domain. It is important to highlight that Cisco's Transport Slicing solution can be used independently from 5G use cases, as it's a generic solution for implementing any transport service based on intents.

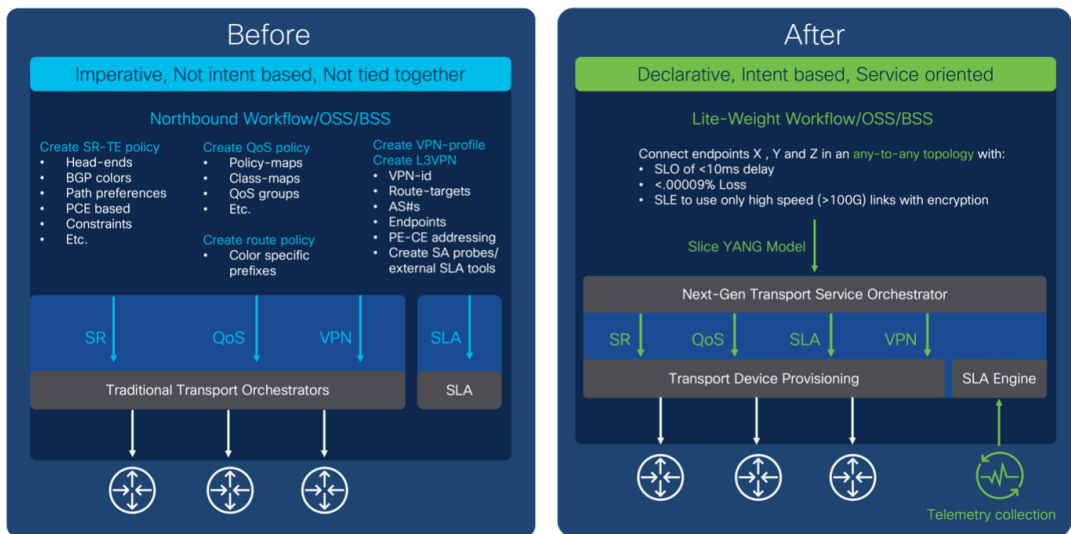
Defining Transport Slicing Scope: 3GPP reference architecture for 5G network slicing



Simplification and ease of use are key principles in transport slicing. The operator wants to start very simply, by requesting from a controller a service based on a desired service intent or outcome (such as supplying low latency to an AR application). He then wants the controller to build the service.

The controller must also monitor the built service to ensure it honors the operator's intent. Above all, the operator wants to avoid exposure to the many configuration parameters required to actually deploy the service at the device layer. Realizing that vision requires the creation of intent-based modularity for value-added transport services supporting the slice, using well-abstracted and declarative service-layer APIs. These service APIs must be maintained and exposed by a controller that can act in a declarative and outcome-based way, as shown in the following figure.

Abstracting the Service Intent



Monitoring the slice's fidelity to the intent involves a Service Level Agreement (SLA) between the customer and the slice provider. To ensure that this SLA both captures the slice intent and has concrete, actionable terms, it can be further defined as either an SLO or SLE:

- **Service Level Objective (SLO):** A desired, achievable target value or range of values for the measurements returned by observation of a Service Level Indicator (SLI). For example, an SLO may be expressed as "SLI <= target", or "lower bound <= SLI <= upper bound".
- **Service Level Expectation (SLE):** The expression of an unmeasurable service-related request that a customer makes of the provider. An SLE is distinct from an SLO because the customer may have little or no way of determining whether the SLE is being met, but will still contract with the provider for a service that meets the expectation (see the following table of sample SLEs).

Table 5: Sample Service Level Expectations

SLE	Description
Encrypted Link Services	Traffic must transit encrypted links only.
Disjoint Path Services	The network has multiple forwarding planes with no common nodes or links.
High speed links only	Traffic must transit high-speed links only. Links offering speeds greater than or equal to 100Gbps are typical for "elephant flows".
Lowest Latency	Always take the lowest latency path. No SLO would be specified in this case.
Regional Avoidance	Do not use nodes or links in specific regions or countries.
Trusted Nodes	Only use trusted nodes ("trusted" meaning verified and not in the common carrier space).
L4-L7 Services	Redirect to "in-line" L4-L7 service on traffic (typically used for security services).
Reliable Links	Use only transit links that have optical protection and L1 diversity.
"Circuit-Style" Services	Provide L1 circuit-like connectivity.
Gaming Services	Use network segments optimized for network gamers (low latency, high bandwidth)
Connected Car	Use network segments optimized for network-connected cars (low latency, close proximity)
Cloud Provider-Specific	Connect me to the secure "walled garden" for a cloud provider (such as AWS or Azure).

The SLA therefore sets key goals and measures to be applied for a given connectivity construct between a sending endpoint and the set of receiving endpoints. It may also describe the extent to which divergence from individual SLOs and SLEs can be tolerated, and specific consequences for violating these SLOs and SLEs.

What Makes Up a Cisco Transport Slice?

To build and deploy these highly abstract intents, Crosswork Network Controller must translate them into actual device configurations. Governing bodies like the IETF and 3GPP leave these decisions to vendors. Cisco can leverage a complete toolkit, built over many years of innovation, as shown in the following figure.

Review: Cisco Toolset for transport level slicing

- QoS and H-QoS: Core and edge
- Forwarding Planes: Shortest Path / SR policies (SRv6 / SR-TE / Flex-algo / Circuit-Style (future))
- SR underlay performance management tools (SR-PM)

Creating and managing the forwarding plane (underlay)

Combining these offer different levels of transport slice separation


- Virtual Private networks : L2 / L3 VPNs
- ODN and Automated traffic Steering (AS)
- VPN performance management tools (Y1731)

Endpoint selection, Slice isolation and mapping to slice forwarding planes. (overlay)

For a Cisco Transport Slice Instance, we categorize the features in the preceding illustration as follows: Service Assurance, Path Forwarding, QoS (PHB), and BGP-based EVPN. The configurations in these categories are what support the slice instance, as shown in the illustration below.

What is a Cisco Transport Slice Instance?

The four legs of the table that make up a Transport Slice Service Instance



RED= Defined in Slice Catalog (intents)
Black= Defined in Slice Instance (endpoints)

Important: A Transport Slice Instance (or Service) is the combination of all these components.

Scale goals:

Slice "types" defined in catalog = ~10-20?

Slice "instances" (differentiated by VPNs/endpoints) = ~1000s

The first three of these features (shown in red) are defined in the slice template catalog (this catalog is equivalent to what 3GPP calls the NSST). The slice catalog contains slice templates, each of which is defined once by a slice designer. Slice *templates* are just blueprints and are not instantiated in the network in any way. Slice *instances* are the instantiated services after they are deployed in the network. The end-user really doesn't need to know the details of what is inside the templates, just what the overall intent (or SLA) is for each slice template. The slice template catalog is thus a catalog of slice intents.

The fourth category – BGP-based VPNs – that makes up a Cisco Transport Slice Instance is the selection of endpoints and service types (L2 or L3 forwarding). Operators define these when deploying the Transport Slice Instance.

The benefit of this approach is to fully abstract the underlying configuration details of the various machinery components required to realize a Transport Slice Instance (aka the IETF Network Slice, or, in 3GPP parlance, the Network Subnet Slice Instance or NSSI).

To deploy a new slice instance, the operator executes the following steps:

1. Select a Slice Intent from the available Templates in the Slice Catalog.
2. Select slice endpoints and connectivity details, which drive the VPN configuration. Once committed, Crosswork Network Controller will then provision:
 - The forwarding plane policy details which drive the segment routing traffic engineering (SR-TE) configurations and BGP prefix coloring for ODN/AS.
 - The QoS profile details, which drive the ingress marking (for PHB treatment) and the egress scheduling.
 - The SLA details, which will drive the needed Service Assurance configurations.
 - The BGP based VPN connectivity requirements to provide endpoint connectivity.

The following illustration provides more detail on the parts of the slice template that automate slice instantiation.

So what is automated when deploying a Slice Instance?

1. **QoS:** The Slicing CFP can apply input and output QoS policy maps on all slice endpoint interfaces (policy-maps pre-deployed). Both L2 & L3 QoS supported.
2. **Path Forwarding:** The Slicing CFP can deploy SR-TE ODN templates on all headends (metrics= latency, igmp, TE, BWoD, FA, etc). Additionally, it will set BGP color community accordingly on all slice advertised prefixes.
3. **Service Assurance:** The Slicing CFP can setup:
 - CNC Heuristic packages for CNC Automated Assurance/Service Health
 - Configure Y1731 probing for P2P L2 slices
 - Configure SR-PM probing for delay and liveness on all slice SR-TE tunnels
4. **Connectivity:** The Slicing CFP will use the L2/L3VPN IETF NM to setup L3 or L2 connectivity automatically across defined slice endpoints. All VPN parameters inferred and abstracted.
 - Setup eVPN VPWS for P2P L2 slices
 - Setup eVPN any-to-any or hub-spoke for L2 multi-point or L3 multipoint slices.
 - Setup up “extranet” connectivity between dedicated and shared slice types. (more on this later).
 - Setup PE-CE eBGP for L3 based slices

© 2022 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

Transport Slice High Level Workflow

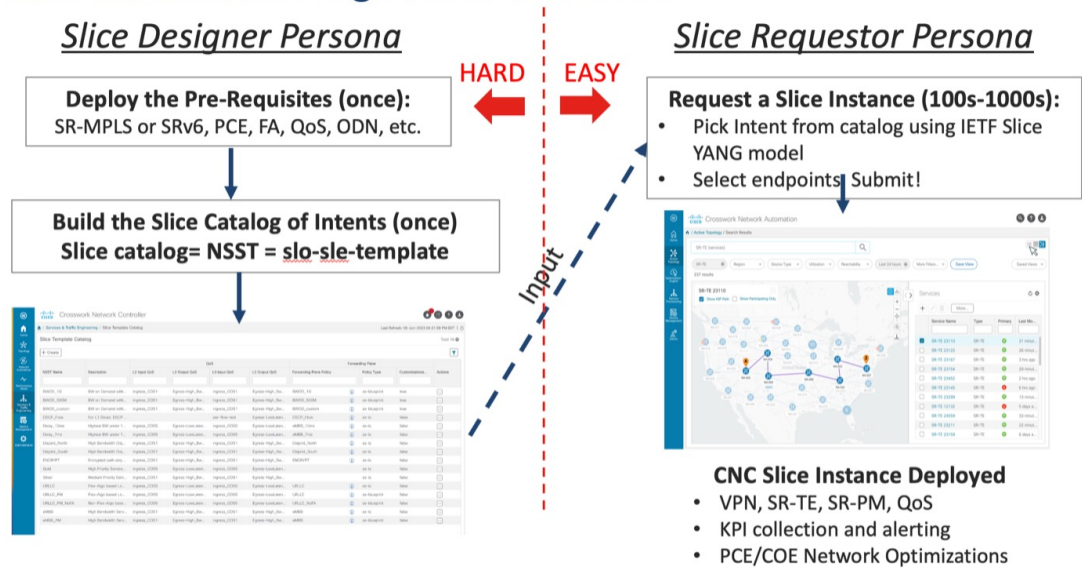
Transport slicing in Crosswork Network Controller is designed around two main user personas:

Slice Designer: The Designer understands the service requirements the provider organization want to offer to customer and is very familiar with the provider network’s underlying capabilities. This person has authorization to do one-time setup operations within the network and has a networking engineering background. They will set up the needed network pre-requisites and then build the slice template catalog, which offers a listing of available slice service offerings for network operators.

Slice Requestor: The Requestor requests new slice instances using the intent-based and simplified CNC user interface. They pick their desired slice type from the pre-built slice catalog, select their endpoints and transport options, and then click submit.

Cisco's objective in the Cisco Crosswork Network Controller Transport Slice solution is to make the user experience as simple as possible for the Requestor. This is the only slice deployment operation driving network service provisioning, and as it must be done constantly for a large SP network, it is a major contributor to provider OPEX. The Slice template catalog creation is done once by highly skilled designer personnel. While the design step is not automated, this approach leverages those skilled resources in a way that maximizes their value to the provider organization at a scale that cannot be realized if the designer must instantiate every slice by hand. The catalog creation requires a good understanding of the network and its capabilities, and requires pre-requisite configurations as shown in the figure below. Slice Designers must be familiar with all the pre-requisite configuration types listed in the illustration for this approach to work.

Slice Automation High Level Workflow



Scenario: Implement an Any-To-Any L3 eMBB Slice

In this scenario, you require a transport slice which has Layer3 any-to-any connectivity across three endpoints, using the intent defined in the catalog for Mobile Broadband (eMBB). The eMBB intent will provide the highest bandwidth available path (including proper QoS marking/scheduling treatment), along with some basic service assurance capabilities such as endpoint interface status and PE-CE route health. The eMBB intent will also enable you to specify:

- The highest bandwidth available path.
- Some basic service assurance capabilities.
- eBGP peer routing connectivity details to CE devices.

Assumptions and Prerequisites

This scenario assumes the network has already built out the required network capabilities for this intent. However, they will be briefly reviewed here for this scenario. For more detailed explanations, see the Cisco Transport Slice Automation Design Guide.

Slice Service Package Prerequisites

There are a few optional prerequisites used by the NSO slice services package that need to be bootstrapped into NSO. The need for these prerequisites are dependent on the types of slices and intents required.

First, if you plan on using any “as-blueprint” forwarding-plane-policy-types in your template catalog, then you need to create an NSO sr-color resource pool so that the slicing service can assign colors for dynamically create ODN policies. This pool should then be referenced by the slicing service.

Second, if creating point-to-point L2 slice service types, the route-policy map assigned to the BGP session for the route reflector is required to be identified to the slicing package. This policy map will be modified by the slicing package with new policies as needed for L2 services, which is a standard approach for VPWS.

In this scenario, these items are not required since you are using neither of these functions:

```
resource-pools id-pool sr-color-pool
range start 1000
range end 2000
!
network-slice-services cfp-configurations color-pool-name sr-color-pool
network-slice-services global-settings parent-rr-route-policy SET_COLOR_EVPN_VPWS
!
```

Path Forwarding Prerequisites

The following settings have been preconfigured with the NSO T-SDN SR-TE CFP for the eMBB ODN path-forwarding intent with these properties:

- Use Color 100 to identify the intent.
- PCE is responsible for dynamic path computation
- The dynamic path computation will be based on the IGP metric.

On NSO, this set of properties will look like the below example. At this stage, the ODN policy has not been pushed to the devices.

```
admin@ncs# show running-config cisco-sr-te-cfp:sr-te odn odn-template eMBB
cisco-sr-te-cfp:sr-te odn odn-templatee eMBB
color 100
dynamic pce
dynamic metric-type igp
!
```

QoS Prerequisites

As described, you (the Slice Designer) should have a good understanding of the network’s settings and device capabilities. You should have a well designed and implemented QoS design throughout your network. In the case of QoS treatment for the high BW business services and for example illustration purposes only, you have chosen to deploy these services with the network’s existing “Class of Service 1” traffic policy (called “ingress_COS1”).

The details of this policy are again provider specific, but in this example the policy will not examine or modify the ingress traffic’s IP DSCP setting, but simply mark all the traffic with an MPLS experimental bit (EXP) of 1 so that downstream core scheduling can provide the proper BW treatment. On egress from the provider

network, you have chosen a policy called “Egress-High_BW-Apps” which will assign 50% of the bandwidth to Class of Service 1 (COS1) marked traffic.

It is assumed these QoS policies are already deployed on all edge PE devices (but not yet on the customer facing interfaces, but have been built out and ready for use). Yet, you still need to identify that these QoS policies are available to the Slice Template catalog for use for Slice services. You will need to provide that mapping and will need to identify which policies are available for either Layer 2 or Layer 3 slice services, or both. Since QoS policies can be tailored specifically to Layer 3 or Layer 2 traffic (for example matching on L3 DSCP vs L2 ToS bits) the system allows you to specify the usage. In the case of the example above, since all traffic is being marked with EXP=1 regardless of DSCP or ToS, these policies are applicable to both L2 or L3 services.

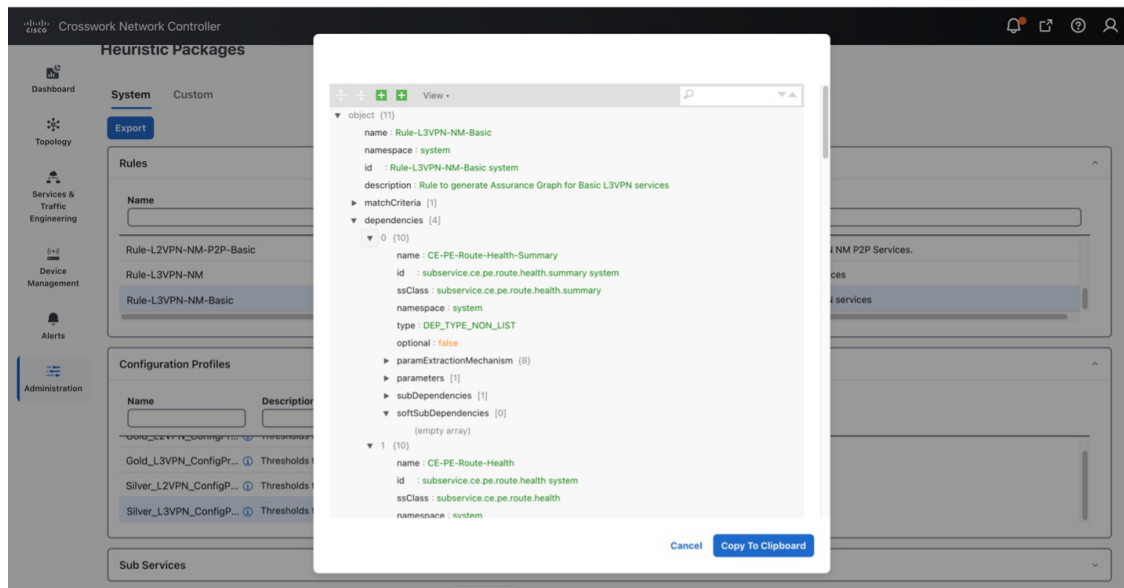
```
admin@ncs# show running-config network-slice-services slq-sle-templates qos-catalog
network-slice-services slq-sle-templates qos-catalog L2 output-qos-policy Egress-High_Bw_Apps
description "High BW egress"
!
network-slice-services slq-sle-templates qos-catalog L2 input-qos-policy ingress_COS1
description "Treat all as Business Data"
!
network-slice-services slq-sle-templates qos-catalog L3 output-qos-policy Egress-High_Bw_Apps
description "High BW egress"
!
network-slice-services slq-sle-templates qos-catalog L3 input-qos-policy ingress_COS1
description "Treat all as Business Data"
!
```

Slice Service Assurance Settings

In this scenario you only have basic service assurance requirements which are based on passive state monitoring (no active probing). You will be using Crosswork Network Controller's Service Health capability and the Crosswork Network Controller system’s pre-built heuristic packages (ConfigProfiles and Rules) which define the objects to be monitored. In the scenario you want to monitor basic device health and PE-CE route health which are included in the basic system package. When you build the slice catalog, you can define which packages to use for L2 point-to-point, L2 multipoint and/or L3 services.

The scenario above is requiring L3 services, but when you create your catalog for the eMBB intent (next step), you also need to consider future slices instances for eMBB services that are L2 service types, thus you can include all these pre-built system heuristic packages in the catalog entry for eMBB. Since these are all pre-built system packages, no prerequisite configurations are required. The system heuristic packages can be viewed via the CNC UI by selecting **Administration > Heuristic Packages** (see below figure).

System Heuristic Profile and Rule names used for eMBB	Usage
Silver_L3PN_ConfigProfile system	L3 profile-name
Rule-L3VPN-NM-Basic system	L3 rule-name
Silver_L2PN_ConfigProfile system	L2 multipoint profile-name
Rule-L2VPN-MP-Basic system	L2 multipoint rule-name
Silver_L2PN_ConfigProfile system	L2 point-to-point profile-name
Rule-L2VPN-NM-Basic system	L2 point-to-point rule-name



Step 1 Create a Slice Template Catalog Entry

You will now build out the slice catalog entry for eMBB intent-based slice services. This operation is done once and will use the above components as input and can cover both L2 and L3 slice instance requests. Once complete, you can move to deploying the slice service instance for this scenario and this entry will now be available for future slice instances requiring eMBB intent services (for example, all of the above prerequisite steps will not be required).

This step, performed by the Slice Designer, builds a slice catalog of intents, or slice types, that will be referred to when creating the actual slice instance. This catalog (along with the slice instances themselves) can be built in multiple ways:

- Using the Crosswork Network Controller UI
- Using the Crosswork Network Controller or NSO Slicing API
- Using the NSO CLI, including load merge from a text file

This scenario has a Service Assurance option that can only be created using the NSO CLI or the Crosswork Network Controller API. It will be shown in the section "Add Service Assurance into the Slice Template Catalog using the NSO CLI".


```
admin@ncs# show running-config network-slice-services slo-sle-templates slo-sle-template eMBB
network-slice-services slo-sle-templates slo-sle-template eMBB
template-description "High Bandwidth Service with basic SLA monitoring"
qos-policy L2 input-policy ingress_COS1
qos-policy L2 output-policy Egress-High_Bw_Apps
qos-policy L3 input-policy ingress_COS1
qos-policy L3 output-policy Egress-High_Bw_Apps
odn forwarding-plane-policy eMBB
odn forwarding-plane-policy-type as-is
service-assurance heuristics monitoring-state enable
service-assurance heuristics L2 point-to-point profile-name "Silver_L2VPN_ConfigProfile system"
service-assurance heuristics L2 point-to-point rule-name "Rule-L2VPN-NM-Basic system"
service-assurance heuristics L2 multipoint profile-name "Silver_L2VPN_ConfigProfile system"
service-assurance heuristics L2 multipoint rule-name "Rule-L2VPN-MP-Basic system"
service-assurance heuristics L3 profile-name "Silver_L3VPN_ConfigProfile system"
service-assurance heuristics L3 rule-name "Rule-L3VPN-NM-Basic system"
!
```

To create a slice template catalog entry using the Crosswork Network Controller UI, do the following:

Step 1 Go to **Services & Traffic Engineering > Slice Template Catalog**.

The Slice Template Catalog screen appears.

Note For the purpose of this scenario, the templates that appear in the image below have already been created on NSO.

NSST Name	Description	QoS				Forwarding Plane			Actions
		L2 Input QoS	L2 Output QoS	L3 Input QoS	L3 Output QoS	Forwarding Plane Policy	Policy Type	Customizati...	
BWOD_1G	BW on Demand wit...	ingress_COS1	Egress-High_Bw...	ingress_COS1	Egress-High_Bw...	BWOD_1G	as-blueprint	true	...
BWOD_500M	BW on Demand wit...	ingress_COS1	Egress-High_Bw...	ingress_COS1	Egress-High_Bw...	BWOD_500M	as-blueprint	true	...
BWOD_custom	BW on Demand wit...	ingress_COS1	Egress-High_Bw...	ingress_COS1	Egress-High_Bw...	BWOD_custom	as-blueprint	true	...
DSCP_Flow	For L3 Slices: DSC...			per-flow-test	Egress-LowLate...	DSCP_Flow	as-is	false	...
Delay_10ms	Highest BW under ...	ingress_COS5	Egress-LowLate...	ingress_COS5	Egress-LowLate...	eMBB_10ms	as-is	false	...
Delay_7ms	Highest BW under ...	ingress_COS5	Egress-LowLate...	ingress_COS5	Egress-LowLate...	eMBB_7ms	as-is	false	...
Disjoint_North	High Bandwidth Di...	ingress_COS1	Egress-High_Bw...	ingress_COS1	Egress-High_Bw...	Disjoint_North	as-is	false	...
Disjoint_South	High Bandwidth Di...	ingress_COS1	Egress-High_Bw...	ingress_COS1	Egress-High_Bw...	Disjoint_South	as-is	false	...
ENCRYPT	Encrypted path onl...	ingress_COS1	Egress-High_Bw...	ingress_COS1	Egress-High_Bw...	ENCRYPT	as-is	false	...
Gold	High Priority Servi...	ingress_COS5	Egress-LowLate...	ingress_COS5	Egress-LowLate...			false	...
Silver	Medium Priority Se...	ingress_COS1	Egress-High_Bw...	ingress_COS1	Egress-High_Bw...			false	...
URLLC	Flex-Algo based Lo...	ingress_COS5	Egress-LowLate...	ingress_COS5	Egress-LowLate...	URLLC	as-is	false	...
URLLC_PM	Flex-Algo based Lo...	ingress_COS5	Egress-LowLate...	ingress_COS5	Egress-LowLate...	URLLC	as-blueprint	false	...

Step 2 Click **+ Create** to create a new slice catalog entry. The New Slice Template screen appears.

Step 3 For Network Subnet Slice Template (NSST), type the new slice template name: **eMBB**. In addition, in the Description field, type a short description of the slice template's intent: **High Bandwidth Service**.

Step 4 Assign the QoS ingress and egress policies. This depends on the slice instance Service Type (L2 or L3 policy) you will define later when creating a new slice instance.

Note For the purpose of this scenario, the five fields at the bottom of the screen (L2 Input QoS, L2 Output QoS, L3 Input QoS, L3 Output QoS, Forwarding Plane Policy Template) have already been provisioned in NSO and automatically appear as an option in each list.

Note You may also refer to the table built earlier and found under QoS in the prerequisites.

Step 2 Add Service Assurance into the Slice Template Catalog using the NSO CLI (optional)

- For L2 Input QoS, select **ingress_COS1**
- For L2 Output QoS, select **Egress-High_Bw_Apps**
- For L3 Input QoS, select **ingress_COS1**
- For L3 Output QoS, select **Egress-High_Bw_Apps**

Step 5 For Forwarding Plane Policy Template, select the ODN template policy created earlier (see prerequisites and assumptions section) that complements the forwarding plane intent. For the purpose of this scenario, select **eMBB**.

Step 6 For Policy Type, determine if this template is to be used **as-is** or **as-blueprint**. For the purpose of this scenario, select **as-is**.

The **as-is** forwarding templates increase overall scalability as the SR-TE tunnels can be shared across multiple slice templates and instances. However, these ODN templates will not be dynamically modified by the slice package with additional functionality, including dynamic support for Performance Measurement or BWoD reservations.

Note When **as-blueprint** is selected, determine if you also want to allow further per-slice instance customizations. These settings determine the SR-TE infrastructure re-use and the scale. Once **as-blueprint** is selected for Policy Type, the **Allow Customizations** check box becomes available.

The screenshot shows the 'New Slice Template' configuration page in the Crosswork Network Controller. The page is titled 'New Slice Template' and has a breadcrumb '← Slice Template Catalog'. The configuration fields are as follows:

- NSST**: eMBB
- Description**: High Bandwidth Service
- L2 Input QoS**: ingress_COS1
- L2 Output QoS**: Egress-High_Bw_Apps
- L3 Input QoS**: ingress_COS1
- L3 Output QoS**: Egress-High_Bw_Apps
- Forwarding Plane Policy Template**: eMBB
- Policy Type**: as-is
- Allow Customizations**:

Annotations on the right side of the screenshot provide additional context:

- Enter string data: A unique transport Network Slice Template Name (NSST) and Description
- Enter Desired QoS polices, depending on Slice Instance Service Type the slicing package will select the proper L2 or L3 policy.
- Forwarding Plane Intent, select the desired Forwarding Plane Policy Template (from pre-created ODN templates). Determine if this forwarding template should be used 'as-is' or 'as-blueprint'. If 'as-blueprint' determine if you would like to allow further per-slice instance customizations (i.e. BWoD). These settings will determine the SR-TE infra re-use and ultimately the scale.

Step 7 Click Save.

Step 2 Add Service Assurance into the Slice Template Catalog using the NSO CLI (optional)

If Service Assurance is required for the Slice instance, then the Slice Designer can add the necessary Service Assurance functionality into the Slice Template using NSO CLI or API. Below is a template in NSO CLI with Service Assurance parameters to be used for this scenario. It can be added directly into NSO CLI or with the API.

There are three Service Assurance sections to the template settings:

- Step 1** Reference pointers to Crosswork Network Controller Service Health Heuristic packages to be used and monitoring state. This monitoring state cannot be changed at the Slice instance level at this time, it is set universally for all slice instances referencing this Slice template. Since different connectivity-types (pt-2-pt or multi-point) can be selected when provisioning a slice instance and different service types (L2 or L3), multiple heuristic package options are available, and the system will select the proper package depending on the slice instance requirements.
- Step 2** If the Slice Instance is a L2 service type with pt-2-pt connectivity, then Y1731 probe monitoring can be enabled. The settings required are shown in the below example. Slice SLA alarming and alerting can be configured if the proper settings are selected in the L2 Heuristic package for Service Health.
- Step 3** This scenario does not require Performance Measurement, so it is not included in the below template. But if desired, it can be enabled in the template and SR-PM will be dynamically configured on the SR-TE tunnel if the Slice Forwarding policy-type is set for **as-blueprint**. Slice SLA alarming and alerting can be configured if the proper settings are selected in the Heuristic package for Service Health.

```

admin@ncs# show running-config network-slice-services s1o-s1e-templates s1o-s1e-template eMBB
network-slice-services s1o-s1e-templates s1o-s1e-template eMBB
template-description "High Bandwidth Service with basic SLA monitoring"
qos-policy L2 input-policy Ingress_COS1
qos-policy L2 output-policy Egress-High_Bw_Apps
qos-policy L3 input-policy Ingress_COS1
qos-policy L3 output-policy Egress-High_Bw_Apps
odn forwarding-plane-policy eMBB
odn forwarding-plane-policy-type as-is
service-assurance heuristics monitoring-state enable
service-assurance heuristics L2 point-to-point profile-name "Silver_L2VPN_ConfigProfile system"
service-assurance heuristics L2 point-to-point rule-name "Rule-L2VPN-NM-Basic system"
service-assurance heuristics L2 multipoint profile-name "Silver_L2VPN_ConfigProfile system"
service-assurance heuristics L2 multipoint rule-name "Rule-L2VPN-MP-Basic system"
service-assurance heuristics L3 profile-name "Silver_L3VPN_ConfigProfile system"
service-assurance heuristics L3 rule-name "Rule-L3VPN-NM-Basic system"
service-assurance ethernet-service-oam md-name foo
service-assurance ethernet-service-oam md-level 4
service-assurance ethernet-service-oam y-1731 profile-delay Profile-Delay-1
!
    
```

See previous explanations from UI figure

CNC Heuristic packages to be used for Service health. Slicing package will pick proper package depending on Slice instance service type (L2 or L3) and L2 connectivity model (point-to-point or multi-point). User can also associate custom packages.

Y1731 probing specifications: For L2 point-to-point slice service types only. Ignored for other slice types

As previously highlighted, Slice Templates with Service Assurance parameters can only created using the NSO CLI or Crosswork Network Controller/NSO API at this time. This also means that these additional parameters will not be visible when viewing the Slice Template in the Crosswork Network Controller UI.

Step 3 Create the Transport Slice Instance

Once the slice type catalog has been created, we can now deploy the transport slice instance. The below table outlines the user data required to deploy this slice. The mandatory data consists of a series of string-data names (user defined), selection of the service type (L2 or L3), catalog intent selection and then defining the Service Demarcation Points (SDPs) which are the PE endpoints facing the customer. These PE endpoints will require IP information since this is a L3 slice and optionally since eBGP was desired for the PE-CE peering protocol, the CE eBGP information is required.

For the purpose of this scenario, use the sample data below:

Table 6: Required Parameter Values:

Parameter	User Value	Mandatory	Notes
slice-service-name	a_L3_A2A_ded	Y	String. Maximum 17 characters. Must be unique.

Parameter	User Value	Mandatory	Notes
description	“any string data”	N	Any string
customer	ACME	N	String meta data- user defined
service-tag	L3	Y	L2 or L3 forwarding
nssai	123459876	N	String meta data- could match 5G nssai assignment if provider desires
slo-sle-template	eMBB	Y	Selection from pre-built slice catalog
isolation	dedicated	N	The default is dedicated- the other option is shared
First SDP endpoint name	1	Y	String- unique within slice instance- At least one SDP must be created, the rest optional
Node-Name	Node-4	Y	PE Node-Name as defined in CNC topology
Attachment-circuit name	ac1	Y	String- unique within slice instance
Interface-ID	TenGigE0/0/0/10	Y	Customer facing PE Interface
VLAN ID	401	N	VLAN ID if using vlan sub-interfaces
Interface IP	172.16.2.1	Y	PE Interface IP address (since L3 service)
Interface IP Mask	29	Y	Interface prefix length (i.e. /29)
Peering protocol	BGP	N	PE-CE peering protocol (BGP or none)
BGP Neighbor ASN	65102	Y	Since bgp selected, peer ASN
BGP Neighbor Address	172.16.2.2	Y	Since bgp was selected, peer IP address

Parameter	User Value	Mandatory	Notes
Second SDP endpoint name	2	N	Additional SDPs are optional, but in this scenario we have three endpoints
Node-Name	Node-5		PE Node-Name as defined in CNC topology
Attachment-circuit name	Ac2		String- unique within slice instance
Interface-ID	TenGigE0/0/0/2		Customer facing PE Interface
VLAN ID	301		VLAN ID if using vlan sub-interfaces
Interface IP	172.16.1.1		PE Interface IP address (since L3 service)
Interface prefix length	29		Interface prefix length (i.e. /29)
Peering protocol	bgp		PE-CE peering protocol (bgp or none)
BGP Neighbor Address	172.16.1.2		Since bgp was selected, peer IP address
BGP Neighbor ASN	65101		Since bgp selected, peer ASN
Third SDP endpoint name	3	N	Additional SDPs are optional, but in this scenario, we have three endpoints
Node-Name	Node-2		PE Node-Name as defined in CNC topology
Attachment-circuit name	Ac3		String- unique within slice instance
Interface-ID	TenGigE0/0/0/2		Customer facing PE Interface
VLAN ID	601		VLAN ID if using vlan sub-interfaces
Interface IP	172.16.3.1		PE Interface IP address (since L3 service)

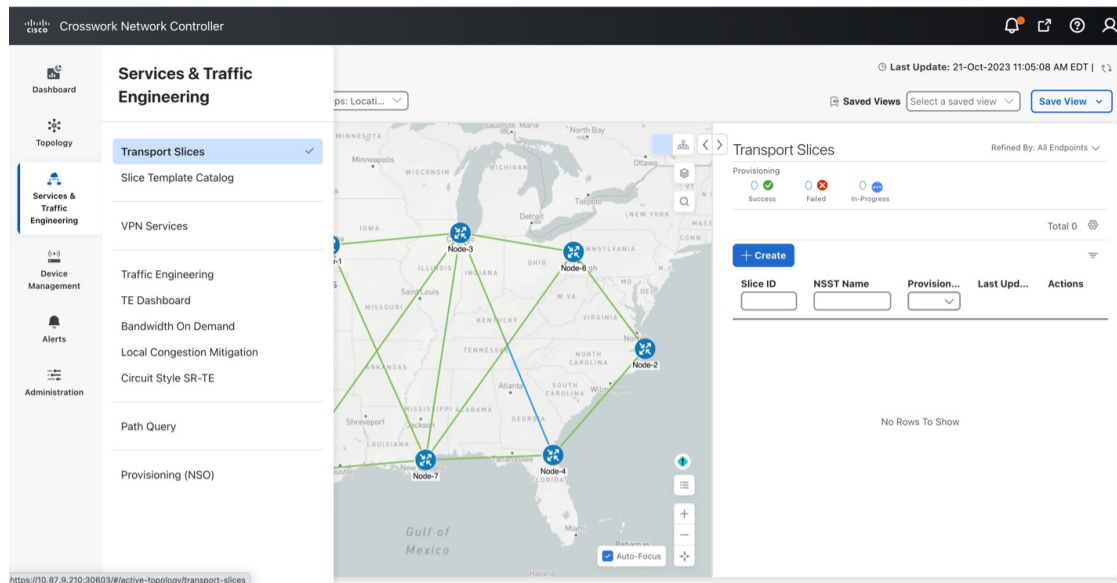
Step 3 Create the Transport Slice Instance

Parameter	User Value	Mandatory	Notes
Interface prefix length	29		Interface prefix length (i.e. /29)
Peering protocol	bgp		PE-CE peering protocol (bgp or none)
BGP Neighbor Address	172.16.3.2		Since bgp was selected, peer IP address
BGP Neighbor ASN	65103		Since bgp selected, peer ASN

The Slice Instance may be created using the Crosswork Network Controller UI, Crosswork Network Controller /NSO API, or NSO CLI. The below example demonstrates the UI steps and explains various fields (both required and optional).

Step 1 Go to **Services & Traffic Engineering > Transport Slices**.

The Transport Slices panel appears.



Step 2 Click **+ Create** to create a new slice.

The New Slice panel appears. At the top, four steps are displayed that tracks the creation of a new slice. The first step requires Basic Details of the new slice.



Step 3 Type the string data into the Slice ID, Customer, and Description fields. For example"

- Slice ID: **a_L3_A2A_ded**
- Customer: **ACME**
- Description: **L3 any-2-any dedicated slice**

Step 4 Select the Service Type: either Layer 2 (**L2**) or Layer 3 (**L3**) connectivity services. In this instance, select **L3**.

Step 5 Optionally, add a string-based Single-Network Slice Selection Assistance Information (S-NSSAI) for 5G mobility customers. This mobility slice-ID information is only used as meta-data by the orchestration system. For example, type **123459876**.

Step 6 Click **Next**.

The screenshot shows the 'New Slice' configuration interface. At the top, there's a progress indicator with four steps: 'Basic Details', 'NSST', 'Connectivity', and 'SDP'. The 'Basic Details' step is active. Below this, there are input fields for 'Slice ID' (value: a_L3_A2A_ded) and 'Customer' (value: ACME). A 'Description' text area contains 'L3 any-2-any dedicated slice'. Under 'Service Type', the 'L3' radio button is selected. An 'S-NSSAI' field contains the value '123459876'. At the bottom, there are 'Cancel' and 'Next' buttons.

Step 7 The second step requires Network Subnet Slice Template (NSST) details.

This information specifies which Slice Template to use from the pre-created Template catalog. So to match the 3GPP 5G naming convention, it is named the Network Slice Subnet Template (NSST). The description on these templates describe the intent specified by the Slice Designer. Depending on the Slice Service Type selected in the previous step (which was **L3**), the system pulls the appropriate L3 (or L2, if specified) based functionality referenced in the template (for example, QoS settings).

Step 8 Select the desired intent from the pre-created Slice Catalog: **eMBB**.

Step 3 Create the Transport Slice Instance

New Slice * Required Field

Basic Details **NSST** Connectivity SDP

Network Subnet Slice Template (NSST) * ⓘ

- BWOD_500M
- BWOD_custom
- DSCP_Flow
- Delay_10ms
- Delay_7ms
- Disjoint_North
- Disjoint_South
- ENCRYPT
- Gold
- Silver
- URLLC
- URLLC_PM
- URLLC_PM_NoFA
- eMBB

L3 Input QoS ingress_COS1

L3 Output QoS Egress-High_Bw_Apps

Forwarding Plane Policy Template eMBB

Policy Type as-is

Customization false

Cancel Previous Next

Note The Slice Catalog names, descriptions, and parameters are set by Slice Designer during the catalog creation phase

Step 9 Click **Next**.

Step 10 The third step requires Connectivity details.

This information builds the connectivity details for the slice by defining if the slice is dedicated or shared. If this is a dedicated slice, it can optionally connect to pre-created shared slices (if it is not a L2 P2P slice). Single Sided Control will allow for uniform bi-directional policies when connecting to the shared slice (i.e., the dedicated slice policies are used when connecting to shared slice endpoints).

Step 11 For Connectivity Group, the field will automatically show Default and cannot be changed.

Note The IETF Slice YANG model has the concept of Connectivity Groups with the idea that multiple Connectivity Groups can be built under a single Slice ID. Currently, only one Connectivity Group is supported.

Step 12 Determine the slice Isolation behavior by either selecting **Dedicated** or **Shared** for Isolation. In this instance, select **Dedicated**.

Note Unique to Crosswork Network Controller, **Dedicated** slices can connect to shared slices (for example, providing a VPN extranet connectivity model).

Note If **Shared** is selected, the remaining selections in the Connectivity step default to system details (for example, the Connectivity Type is set to **Any To Any**).

Step 13 For Connectivity Type, select **Any To Any**.

When selecting the connectivity requirements, choosing L2 or L3 services will determine the available Connectivity Type options.

- L3 Services: **Any To Any, Hub and Spoke.**
- L2 Services: **Any To Any, Hub and Spoke, Point To Point.**

Note If you select **Hub and Spoke**, the endpoint role is selected in a later step.

Step 14 In this instance, skip both Connectivity Shared Slices and Bandwidth Reservation.

Note Connectivity Shared Slices – If **Dedicated** was selected, the option to connect to an existing shared slice instance becomes available.

Note Bandwidth Reservation – If **Allow Customizations** is selected during the catalog entry (and the Policy Type selected is **as-blueprint**), thus having a customizable NSST, you can select Bandwidth Reservation per slice instance or enter a different value.

Step 15 For Single Sided Control, leave as the default, **True**.

Note If **True** is selected, it will force a dedicated slice path forwarding behavior towards shared slice endpoints (overriding shared slice path forwarding intent and will ensure the same bi-directional path forwarding).

Step 16 Select **Show advanced settings** to edit optional parameters.

Advanced settings are only available for non-L2 P2P based slices. This will allow for custom Route Target (RT) and Route Distinguisher (RD) settings. By default, these are set to auto and thus not required to configure. If connectivity type is any-to-any and if manual RT is selected then a box allowing for manual entry of the RT is presented, with this value being used uniformly across all sites to import/export. If the connectivity type is hub-spoke, then there will be two RTs required, one for hub and one for spoke.

Step 17 For Route Target Type, select **Auto**.

Step 18 For Route Distinguisher Type, select **Auto**.

Step 19 Click **Next**.

Step 3 Create the Transport Slice Instance

New Slice * Required Field

Progress: Basic Details | NSST | **Connectivity** | SDP

Connectivity Group [ⓘ]
Default

Isolation [ⓘ]
 Dedicated Shared

Connectivity Type [ⓘ]
Any To Any

Connectivity Shared Slices [ⓘ]
Select One or More

Single Sided Control [ⓘ]
 True False

Bandwidth Reservation [ⓘ] : None Selected

1 G 5 G 10 G 50 G 100 G

OR

Enter a value Gbps

[Hide advanced settings](#)

Route Target Type [ⓘ]
 Auto Manual

Route Distinguisher Type [ⓘ]
 Auto Manual

Step 20 The fourth step requires Slice Demarcation Point (SDP) details.

Here, the Slice Requester provides PE endpoint interface details. These endpoints are called SDPs per IETF Slicing standards.

Step 21 Enter the SDP ID and Attachment Circuit ID to configure the string data. Both entries must be unique within the slice service instance.

- For SDP ID, type **1**.
- For Attachment Circuit ID, type **1**.

Step 22 In Node ID, select a node from the list: **Node-4 [192.168.255.20]**.

This Node ID uniquely identifies an edge node of the SDP.

Step 23 Select the Interface Type, **TenGigE**, and type the Interface ID, **0/0/0/10**.

Step 24 For VLAN ID (optional), type **401**.

Step 25 For Interface IP, type **172.16..2.1** and type **29** for the mask. This defines the IP address of the attachment circuit.

Step 26 Since this is an L3 slice service, the following field is required. For Peering Protocol, select **BGP** as the SDP peering protocol to CE.

Step 27 Since the Peering Protocol was defined as BGP, the following fields are required.

- a. For Remote-AS, type **65102**.
- b. For Neighbor Address, type **172.16.2.2**.

Step 28

Click **+ Add Another** to add a second (Node-5) and third (Node-2) SDP endpoints. See the Slice Instance Required Data table for parameter values.

Slice Demarcation Point • ⓘ

Node ID	SDP ID	AC ID
Node-4	1	1

SDP ID • ⓘ
1

Attachment Circuit ID • ⓘ
1

Node ID • ⓘ
Node-4 [192.168.255.20] x ▾

Interface Type • ⓘ
TenGigE ▾

Interface ID • ⓘ
0/0/0/10

VLAN ID ⓘ
401

Interface IP • ⓘ
172.16.2.1 / 29
Range: 1 to 128

Peering Protocol • ⓘ
BGP ▾

Remote-AS •
65102

Neighbour Address •
172.16.2.2

[Show advanced settings](#)

[+Add Another](#)

[Cancel](#) [Previous](#) [Commit Changes](#)

Step 29

Click **Commit Changes**.

The new slice service is deployed.

Step 4 Deploy a Slice using NSO CLI (optional method)

The option to deploy a slice using the NSO CLI is also available. The below payload shows the details of deploying a slice using load merge when using the NSO CLI. The defaults are not displayed.

```

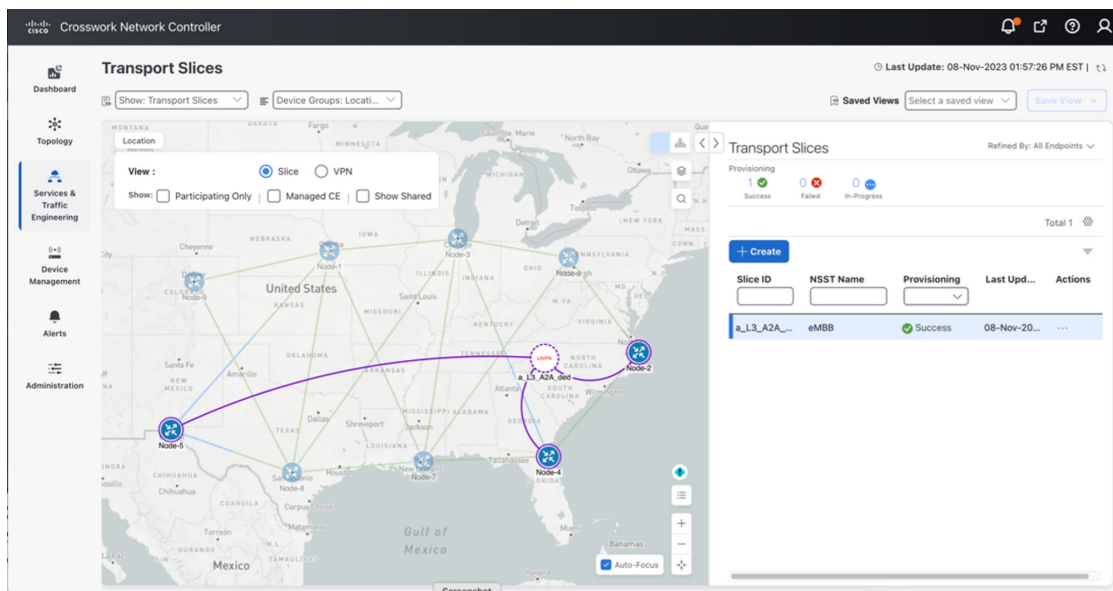
network-slice-services slice-service a_L3_A2A_ded
service-description "L3 any-2-any dedicated slice- a_L3_A2A_dedicated_eMBB.cli"
service-tags tag-type service-tag-customer
  value [ ACME_1 ]
!
service-tags tag-type service-tag-service
  value [ L3_1 ]
!
service-tags tag-opaque DSSAI
  value [ 123459876_1 ]
!
slo-sle-template eMBB
sdfs sdp 1
node-id Node-4
service-match-criteria match-criterion 1
  target-connection-group-id group1
!
attachment-circuits attachment-circuit ac1
ac-tp-id          TenGigE0/0/0/10
ac-ip-address     172.16.2.1
ac-ip-prefix-length 29
ac-tags ac-tags attachment-circuit-tag-vlan-id
  value [ 401_1 ]
!
sdp-peering protocol peering-protocol-bgp
bgp-attributes neighbor [ 172.16.2.2_1 ]
bgp-attributes remote-as 65102
!
!
!
sdfs sdp 2
node-id Node-5
service-match-criteria match-criterion 1
  target-connection-group-id group1
!
attachment-circuits attachment-circuit ac2
ac-tp-id          TenGigE0/0/0/2
ac-ip-address     172.16.1.1
ac-ip-prefix-length 29
ac-tags ac-tags attachment-circuit-tag-vlan-id
  value [ 301_1 ]
!
sdp-peering protocol peering-protocol-bgp
bgp-attributes neighbor [ 172.16.1.2_1 ]
bgp-attributes remote-as 65101
!
!
!
sdfs sdp 3
node-id Node-2
service-match-criteria match-criterion 1
  target-connection-group-id group1
!
attachment-circuits attachment-circuit ac3
ac-tp-id          TenGigE0/0/0/2
ac-ip-address     172.16.3.1
ac-ip-prefix-length 29
ac-tags ac-tags attachment-circuit-tag-vlan-id
  value [ 601_1 ]
!
sdp-peering protocol peering-protocol-bgp
bgp-attributes neighbor [ 172.16.3.2_1 ]
bgp-attributes remote-as 65103
!
!
!
connection-groups connection-group group1
connectivity-type any-to-any
!
!

```

Step 5 Visualize and Validate the New Slice Deployment

Step 1 Go to **Services & Traffic Engineering > Transport Slices**.

The Transport Slices panel appears with the new slice displayed. The Provisioning state should show as **Success**.



Step 2 Optionally, the slice service state can be verified, from the NSO CLI, that all stages were successfully provisioned with all plan states **reached**.


```
admin@ncs# show network-slice-services slice-service-plan_a_L3_A2A_ded
```

TYPE	NAME	BACK TRACK	GOAL	STATUS CODE	NODE	STATE	STATUS	WHEN	POST ACTION ref	STATUS
self	self	false	-	-	-	init ready	reached	2023-10-21T16:06:21	-	-
sdr	1	false	-	-	Node-4	init ready	reached	2023-10-21T20:51:31	-	-
sdr	2	false	-	-	Node-5	init ready	reached	2023-10-21T16:06:21	-	-
sdr	3	false	-	-	Node-2	init ready	reached	2023-10-21T20:51:31	-	-
sdr	3	false	-	-	Node-2	init ready	reached	2023-10-21T16:06:21	-	-
sdr	3	false	-	-	Node-2	init ready	reached	2023-10-21T16:06:21	-	-
sdr	3	false	-	-	Node-2	init ready	reached	2023-10-21T20:51:31	-	-

```

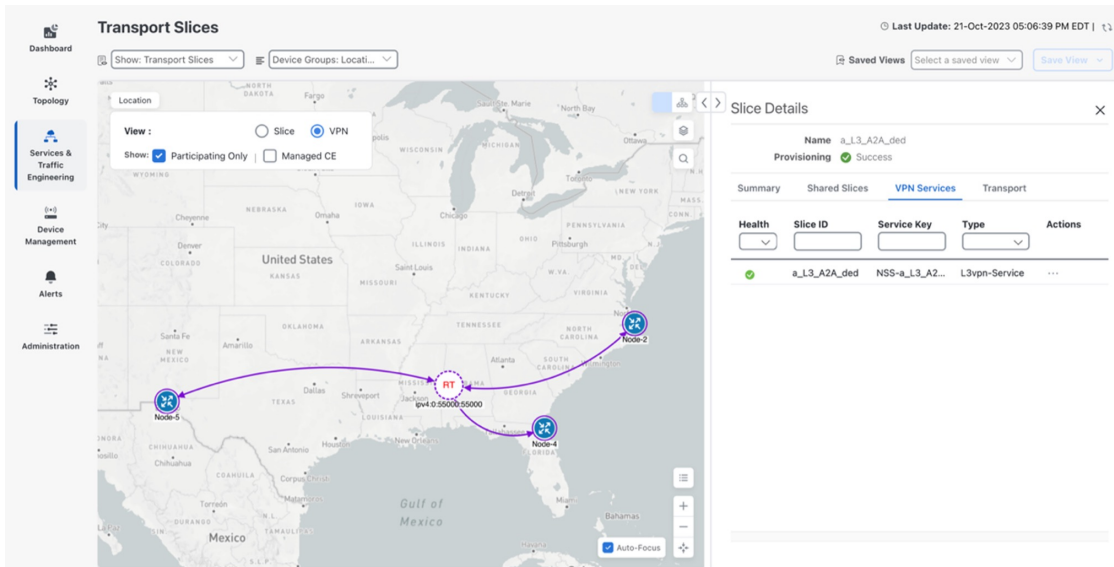
plan status color-allocation-data color 100
plan status service-tag-service [ L3 ]
plan status forwarding-plane-policy AMBB
plan status rt-allocation-data hub-rt 0:55000:55000

```

Step 3 From the Transport Slices screen in the UI, click  in the Actions column for the newly created slice, **a_L3_A2A_ded**, and select **View Details**.

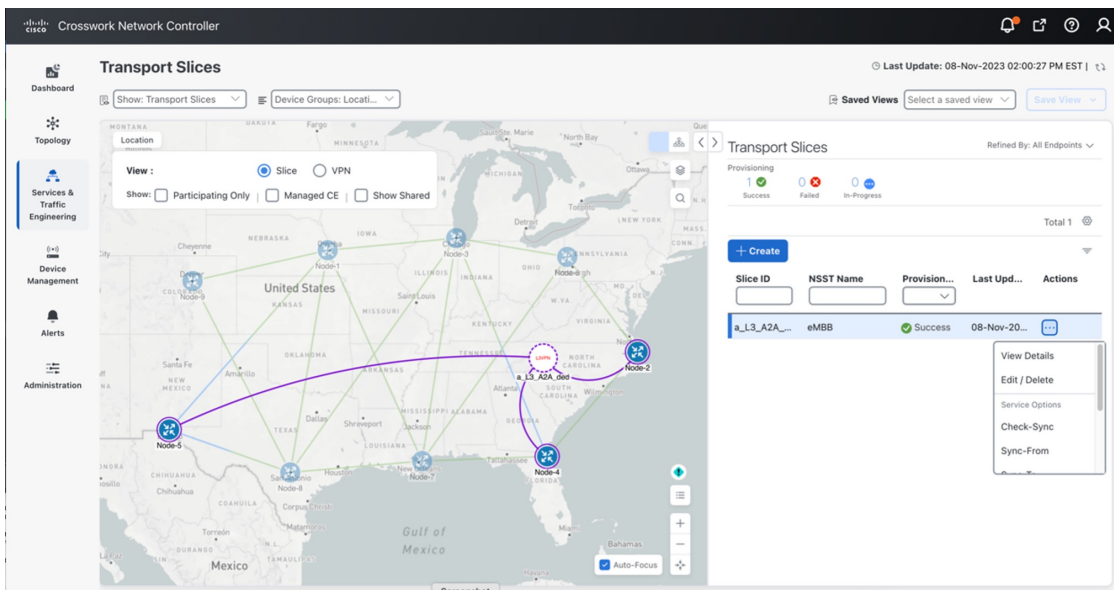
The Slice Details panel appears while the topology map updates to show the new slice.

Step 4 In the Slice Details panel, select the **VPN Services** tab, and in the topology map select **VPN** as the View. The information updates so you can see that the slice provisioned with auto-RTs of 55000:55000 and the service is healthy.



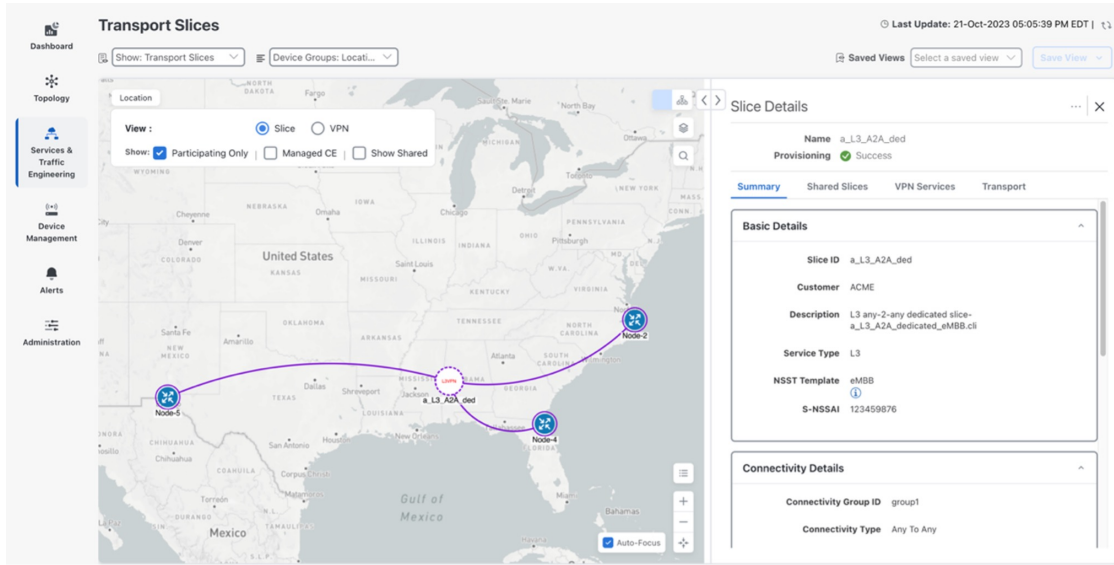
Step 5 In the Slice Details panel, click **X** to close the VPN Services tab and Slice Details panel.

Step 6 In the Transport Slices panel, click **⋮** in the Actions column for the newly created slice, **a_L3_A2A_ded**, and select **View Details**.



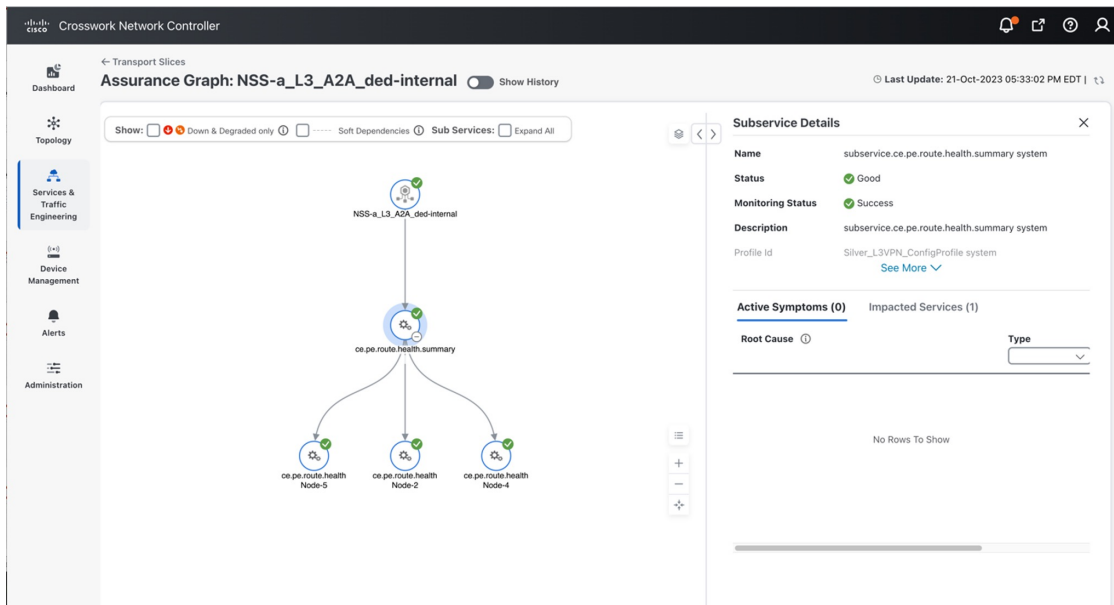
Step 7 In the Slice Details panel, select the **Summary** tab to view the slice details: Basic Details, Connectivity Details, Service Demarcation point (SDP).

Step 5 Visualize and Validate the New Slice Deployment



Step 8


Again, select the VPN Services tab. Click ... in the Actions column for the newly created slice, **a_L3_A2A_ded**, and select **Assurance Graph** so to view the Monitoring Status and the status of the Service Health components defined in the Service Assurance Heuristics Package that was included in the Slice eMBB intent.



Step 9

As additional slices are added, you can visualize and validate further details, such as forwarding path. For example, here is a sample forwarding path for slice traffic from Node-2 to Node-5. A few observations.

- The high BW link between Node-3 and Node-8 (IGP=5k) is used (which was the desired intent), but this link also has a delay of high delay of 10ms (in each direction). Since latency was not an intent objective, this is fine.
- Also notice that Equal Cost Multi-Pathing is used when available (multiple ECMP paths between Node-2 and Node-3).



Step 10

Using the example above, once additional slices are added, if external Accedian probes were also installed at the CPE sites connected to the Slice endpoints at Node-2 and Node-5, an ~11ms latency (each way) between the two sites can be seen (example below). This is accurate because the link between Node-3 and Node-8 has a latency of ~10ms in each direction.



Summary and Conclusion

As we observed in this example, users can utilize Cisco Crosswork Network Controller to create a transport slice which has Layer3 any-to-any connectivity across three endpoints, using the enhanced Mobile Broadband (eMBB) catalog intent. The eMBB intent provides the highest bandwidth available path (including proper QoS marking/scheduling treatment), along with some basic service assurance capabilities such as endpoint interface status and PE-CE route health.

