

Cisco Crosswork Network Controller 2.0 Solution Workflow Guide

Contents

Contents	2
1 Solution Overview	1
Description	1
Supported Use Cases	1
Solution Components and Integrated Architecture	2
Supported Device Software	7
Multi-Vendor Capabilities	7
Extensibility	8
UI Overview	8
2 Orchestrated Service Provisioning	12
Overview	12
Scenario 1 - Implement and Maintain SLA for an L3VPN Service (using ODN)	13
Scenario 2 - Mandate a static path for an EVPN-VPWS service using an explicit SR-TE policy	24
Scenario 3 - Provision an L2VPN service over an RSVP-TE tunnel with reserved bandwidth	32
Scenario 4 - Provision a Soft Bandwidth Guarantee with Optimization Constraints	39
3 Bandwidth and Network Optimization	45
Overview	45
Scenario 5 - Use Local Congestion Mitigation (LCM) to reroute traffic on an over-utilized link	47
4 Network Maintenance Window	55
Overview	55
Scenario 6: Perform a software upgrade on a provider device during a scheduled maintenance window ..	56
5 Programmable Closed-Loop Auto-Remediation	67
Overview	67
Scenario 7 - Achieve Predictive Traffic Load Balancing Using Segment Routing Affinity	68
6 Automation of Onboarding and Provisioning of IOS-XR Devices Using ZTP	70
Overview	70
Scenario 8 - Automatically onboard and provision new devices in the network	71

1 Solution Overview

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 with optimal network utilization and 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 daily operational tasks is the optimal way to become more efficient and competitive.

Cisco Crosswork Network Controller is a turnkey 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 cost. The solution combines intent-based network automation to deliver critical capabilities for service orchestration and fulfilment, network optimization, service path computation, device deployment and management, and anomaly detection and automatic remediation. Using telemetry gathering and automated responses, Cisco Crosswork Network Controller delivers network optimization capabilities that would be nearly impossible to replicate even with a highly skilled and dedicated staff operating the network.

The fully integrated solution combines core capabilities from multiple innovative, industry-leading products including Cisco® Network Services Orchestrator (NSO), Cisco Segment Routing Path Computation Element (SR-PCE), and the Cisco Crosswork suite of applications. Its unified user interface allows real-time visualization of the network topology and services, as well as service and transport provisioning, via a single pane of glass.

Supported Use Cases

- **Orchestrated service provisioning:** Provisioning of layer 2 (L2VPN) and layer 3 VPN (L3VPN) services with underlay transport policies in order to define, meet, and maintain SLAs, using the UI or APIs. The ability to replace manual CLI-based configuration activities to API- or UI-guided activities saves time and improves the accuracy of configuration changes.
- **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. Real-time telemetry is used to detect changes in network traffic and then changes in the network are automatically implemented to deliver on the operator's intent.
- **Local Congestion Management:** Cisco Crosswork Network Controller 2.0 introduces support for Local Congestion Mitigation (LCM) which 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 has a "human in the loop" aspect where the control of making changes in the network is in the hands of the operator.
- **Visualization of network and service topology and inventory:** Visibility into device and service inventory and visualization of devices, links, and transport/VPN services and their health status on logical maps or within their geographical context.
- **Performance-based closed-loop automation:** Automated discovery and remediation of problems in the network by allowing Key Performance Indicator (KPI) customization and monitoring and triggering of pre-

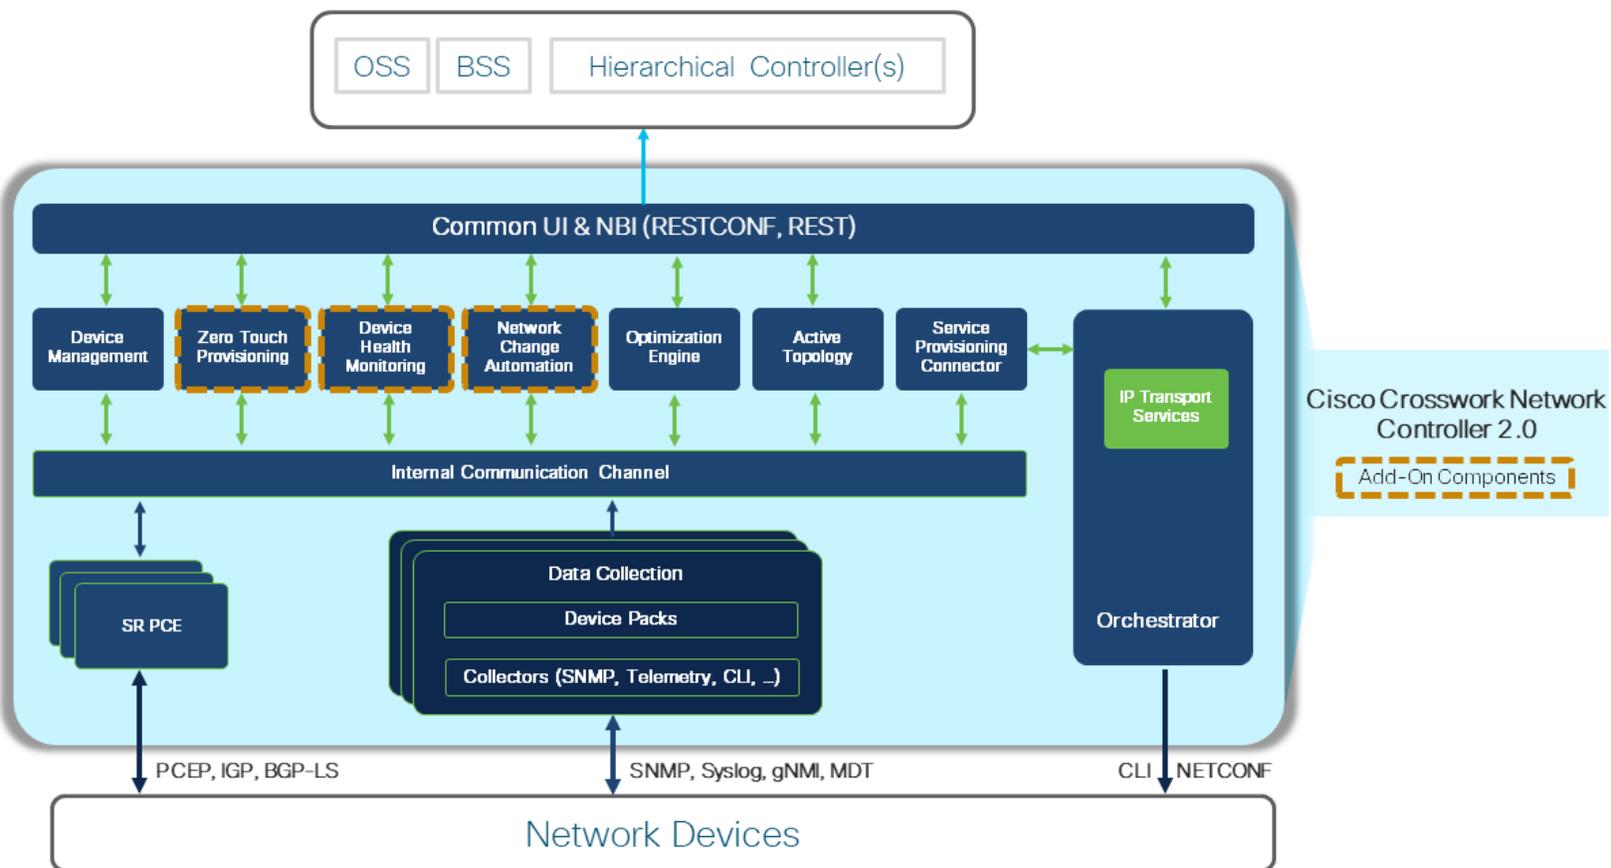
defined remediation tasks when a KPI threshold is breached. Cisco Crosswork Health Insights and Cisco Crosswork Change Automation must be installed for this use case.

- 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. Cisco Crosswork Health Insights and Change Automation must be installed for this use case.
- Secured zero-touch onboarding and provisioning of devices:** Onboarding new IOS-XR devices and provisioning Day0 configuration automatically, resulting in faster deployment of new hardware at a lower operating cost. Cisco Crosswork Zero Touch Provisioning must be installed for this use case.

Solution Components 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.

Note: Crosswork Zero Touch Provisioning, Device Health Monitoring (Crosswork Health Insights), and Crosswork Network Change Automation are optional add-on components.



The following components make up the Cisco Crosswork Network Controller 2.0 solution:

- [Cisco Crosswork Active Topology](#)

- [Cisco Crosswork Optimization Engine](#)
- [Cisco Crosswork Data Gateway \(CDG\)](#)
- [Crosswork Common UI and API](#)
- [Crosswork Infrastructure and Shared Services](#)
- [Cisco Crosswork Health Insights and Cisco Crosswork Change Automation](#)
- [Cisco Crosswork Zero-Touch Provisioning \(ZTP\)](#)
- [Cisco Network Services Orchestrator \(NSO\)](#)
- [Cisco Segment Routing Path Computation Element \(SR-PCE\)](#)

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 enable closed-loop tracking of the network state, reacting quickly to changes in network conditions to support a self-healing network.

Cisco Crosswork Data Gateway (CDG)

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 that 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 an undefined 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. After 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 of Cisco Crosswork Network Controller's functionality is provided within a single, common graphical user interface. This common UI brings together the features of all of 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 functionality is also fully integrated into the common UI. Having all functionality within a common UI instead of having to navigate individual application UIs enhances the operational experience and increases productivity.

Common APIs enable Crosswork Network Controller's programmability. The common API provides single access point for all APIs exposed by the 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 of the Cisco Crosswork applications can be deployed. It is a microservices-based platform that brings together streaming telemetry and model-driven application programming interfaces (APIs) to redefine service provider network operations. It retrieves real-time information from the network, analyzes the data, and uses APIs to apply network changes. It employs a cluster architecture to be extensible, scalable, and highly available.

Cisco Crosswork Health Insights and Cisco Crosswork Change Automation

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

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

The Cisco Crosswork Change Automation application 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 applications 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 automatically performed when a defined Key Performance Indicator (KPI) threshold has been breached. This reduces the time it takes to discover and repair a problem while minimizing the risk of human error through manual network operator intervention.

Cisco Crosswork Zero-Touch Provisioning (ZTP)

Cisco Crosswork ZTP is an application that can optionally be installed with Cisco Crosswork Network Controller.

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

Cisco Crosswork ZTP now 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 the 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. Once 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. Once 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 (NSO)

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 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 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 needed to realize customer services and 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 generating the minimum set of changes required to instantiate the service in the network.

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

Cisco Crosswork Network Controller is packaged with the following Cisco NSO function packs:

- SR-TE core function pack (CFP) enables provisioning of explicit and dynamic segment routing policies, including 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
- L3VPN

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. Consult your Cisco Customer Experience representative for more information.

- Sample IETF-compliant RSVP-TE function pack intended as a reference implementation for RSVP-TE tunnel provisioning, to be customized as required.

Note: 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 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.

Supported Device Software

OS	Version	SR-PCE	PCE-Init	PCC-Init	NSO+CFP		Crosswork Infrastructure 4.0	Optimization Engine	ZTP
					CLI	NETCONF			
IOS-XR	6.5.3				yes		yes	yes	
	6.6.3	yes	yes	yes	yes		yes	yes	
	7.0.2		yes	yes	yes		yes	yes	
	7.1.2	yes	yes	yes	yes		yes	yes	
	7.2.1		yes	yes	yes	yes		yes	yes
	7.3.1	yes	yes (Cisco ASR9000 Series only)	yes	yes	yes	yes	yes	yes
IOS-XE	17.4.1			yes	yes		yes	yes	

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 operational applications to deploy and maintain different vendor products for a single network.

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

- Network service orchestration via Cisco NSO using CLI and Netconf/YANG. Cisco NSO 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 CDG 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/services are involved.

Extensibility

The Cisco Crosswork Network Controller provisioning functionality can be extended using the product application programming interfaces (APIs). For more information about the product 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, so when new services are introduced, they can be easily incorporated.

UI Overview

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/
```

or

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

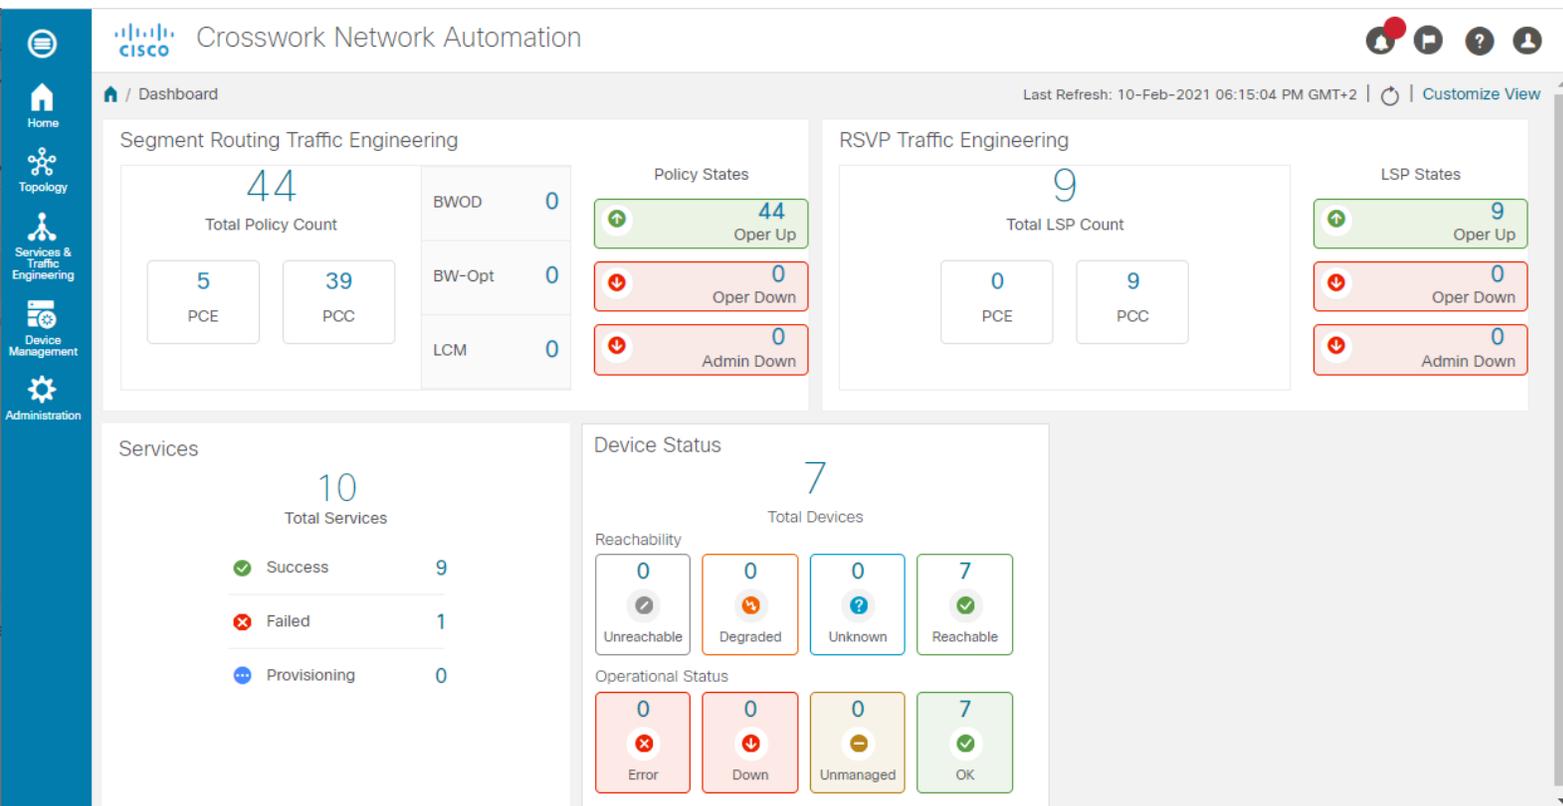
Note: The IPv6 address in the URL must be enclosed with brackets.

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

Upon 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 do this, the browser accepts the server as a trusted site in all future login attempts.

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, including reachability and operational status of devices, as well as transport policies and VPN services. The dashboard is made up of a series of dashlets. The specific dashlets included in the dashboard depend on which Cisco Crosswork applications are installed. Links in each dashlet allow you to drill down for more details.



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 Home, 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.

Home

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

Topology

The network topology can be displayed on a logical map or a geographical map (geo map), where the devices and links are shown in their geographic context. The logical map shows devices and their links, positioned according to an automatic layout algorithm, ignoring their geographical location. 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).

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 logical map (shown below) and the geographical map.

The screenshot displays the Cisco Crosswork Network Automation interface. On the left is a navigation sidebar with icons for Home, Topology, Services & Traffic Engineering, Device Management, and Administration. The main area is titled 'Topology' and shows a network diagram with nodes labeled PE-A, PE-B, PE-C, P-BOTTOMLEFT, P-BOTTOMRIGHT, P-TOPLEFT, and P-TOPRIGHT connected by dashed lines. On the right, a 'Devices' panel shows a table with 7 devices, all in an 'OK' operational state and 'Reachable'.

Host Name	Node IP	Operational State	Reachability State
P-BOTTOMLEFT	172.16.1.43	OK	Reachable
P-BOTTOMRIG...	172.16.1.44	OK	Reachable
P-TOPLEFT	172.16.1.41	OK	Reachable
P-TOPRIGHT	172.16.1.42	OK	Reachable
PE-A	172.16.1.45	OK	Reachable
PE-B	172.16.1.46	OK	Reachable
PE-C	172.16.1.47	OK	Reachable

Services & Traffic Engineering

The screenshot shows the 'Services & Traffic Engineering' menu. It is divided into three sections: 'Services & Traffic Engineering' (containing VPN Services and Traffic Engineering), 'Device Management' (containing TE Link Affinities, Bandwidth On Demand, and Bandwidth Optimization), and 'Admin' (containing Local Congestion Mitigation and Provisioning).

The Services & Traffic Engineering menu provides access to VPN and transport provisioning and visualization functionality, as well as bandwidth management functionality.

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

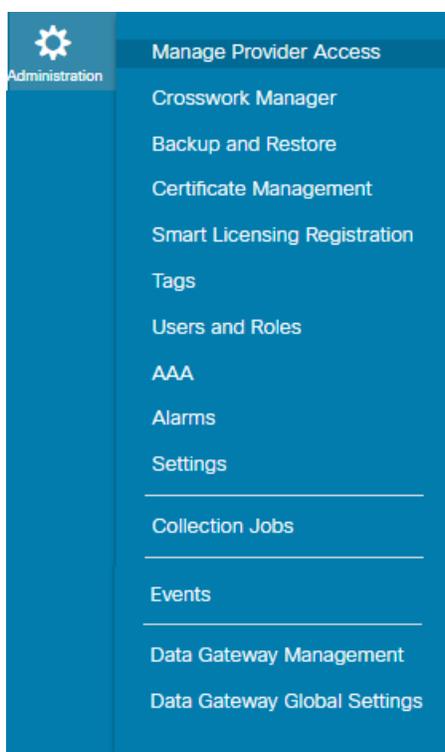
Choose **Provisioning** to access the provisioning UI rendered from the Cisco NSO models. Here you can create L2VPN and L3VPN services, SR-TE policies, SR ODN templates, and RSVP-TE tunnels. You can also create the resources required for these services and policies, such as resource pools, route policies 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 would typically be performed by an administrator.

Refer to the [Cisco Crosswork Infrastructure 4.0 and Applications Administration Guide](#) for information about these functions.

2 Orchestrated Service Provisioning

Overview

Objective

Provision VPN services with underlay transport policies in order to meet and maintain service-level agreements (SLA).

Challenge

The 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 SLAs. In a typical lifecycle, there is a feedback loop that traditionally requires manual 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 Cisco 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 in order to automatically maintain the SLAs.

The services required for this solution can be created and managed using the Cisco Crosswork Network Controller UI. L2/L3 VPN Yang model-based service intents are implemented using the Cisco NSO sample function packs which provide sample service models that can be extended and fine-tuned to meet customer needs.

Note: 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 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.

How Does it Work?

1. User creates an SR-TE policy/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 to be associated with the VPN service
3. During the provisioning process for the above steps, Cisco NSO 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 involves 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 1 - Implement and Maintain SLA for an L3VPN Service \(using ODN\)](#)
- [Scenario 2 - Mandate a static path for an EVPN-VPWS service using an explicit SR-TE policy](#)
- [Scenario 3 - Provision an L2VPN service over an RSVP-TE tunnel with reserved bandwidth](#)
- [Scenario 4 - Provision a Soft Bandwidth Guarantee with Optimization Constraints](#)

Additional Resources

- For information about segment routing and segment routing policies, refer to the [Cisco Crosswork Optimization Engine User Guide](#)
- Cisco NSO documentation is available here: <https://software.cisco.com/download/home/286323467/type/286283941/release/5.4.2>

Scenario 1 - Implement and Maintain SLA for an L3VPN Service (using ODN)

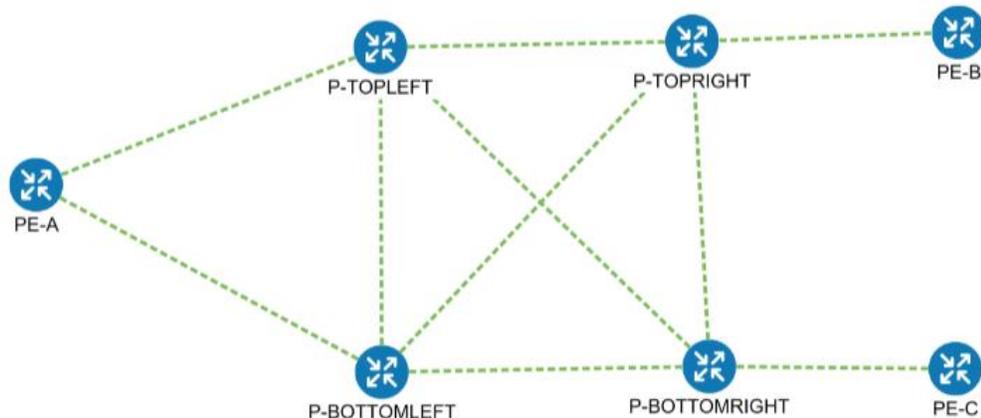
Scenario Context

This scenario walks you through the procedure for provisioning an L3VPN service that requires a specific SLA objective. In this example, lowest latency path is the SLA objective. The customer requires a low latency path for high priority traffic. The customer also wants to use disjoint paths, i.e., two unique paths that steer traffic from the same source and destination avoiding common links so that there is no single point of failure.

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. The ODN template defines the SLA upon which you want to optimize the path. In this case, we will optimize on latency.
- Request that the computed paths be disjoint – they will not share the same link.
- Create an L3VPN service with 3 endpoints – PE-A, PE-B, and PE-C (this is the overlay configuration).
- 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.
- Visualize how this overlay/underlay configuration optimizes the traffic path and automatically maintains the SLA.

Assumptions and Prerequisites

- To use ODN, BGP peering for the prefixes should be configured between the endpoints/PEs. Usually for L3VPN, this is the VPNv4 and VPNv6 address family peering.

Workflow

- [Step 1: Create an ODN template to map color to an SLA objective and constraints](#)
- [Step 2: Create a Route Policy](#)
- [Step 3: Create a VPN profile](#)
- [Step 4: Create and provision the L3VPN service](#)
- [Step 5: Visualize the New VPN Service on the Map to See the Traffic Path](#)
- [Step 6: Observe automatic network optimization](#)

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

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 create tunnels dynamically (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.

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.

Procedure

1. Go to **Services & Traffic Engineering > Provisioning > SR-TE > ODN Template**.
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**.
3. Choose the headend device, PE-A, and specify the color **72**.
4. Under dynamic, select **“latency”** as the metric type. This is the SLA objective on which we are optimizing.
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).

- 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.

L3VPN_NM-SRTE-ODN_72-a

head-end

name

PE-A

maximum-sid-depth

color *

72

bandwidth

dynamic

Enable dynamic

metric-type

latency

pce

flex-alg

> metric-margin

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.

- 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.

ODN-Template Selected 1 / Total 5 ⚙

+ ↻
⏮

Name	Provisioning State	Date Created	Actions
L3VPN_NM-SRTE-ODN_70	✔ Success	30-Jan-2021 05:56:12 PM GMT+2	...
L3VPN_NM-SRTE-ODN_71-a	✔ Success	30-Jan-2021 05:56:13 PM GMT+2	...
L3VPN_NM-SRTE-ODN_71-c	✔ Success	30-Jan-2021 05:56:13 PM GMT+2	...
L3VPN_NM-SRTE-ODN_72-a	✔ Success	30-Jan-2021 05:56:13 PM GMT+2	<ul style="list-style-type: none"> Config View Edit Delete Re-deploy
L3VPN_NM-SRTE-ODN_72-b	✔ Success	30-Jan-2021 05:56:13 PM GMT+2	...

Configured Data ×

View -
🔍

```

object {1}
  cisco-sr-te-cfp-sr-odn:odn-template {4}
    head-end [1]
      0 {1}
        name : PE-A
        color : 72
        name : L3VPN_NM-SRTE-ODN_72-a
    dynamic {3}
      disjoint-path {2}
        group-id : 16
        type : link
      pce {0}
        (empty object)
        metric-type : latency
    
```

Copy To Clipboard
Cancel

- Create the other ODN templates listed above in the same manner.

Step 2: Create a 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, 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

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

Procedure

1. Go to **Services & Traffic Engineering > Provisioning > L3vpn > L3vpn Route Policy**.
2. Click + to create a new route policy and give it a unique name.
3. Under Color, click +. Specify the same color that is specified in the ODN template for PE-A, **70**, and click **Continue**.
4. Enter the required IPv4 or IPv6 prefixes to identify the network traffic.
In this case, IPv4 prefix **70.70.70.0/30**.

The screenshot displays the configuration interface for a route policy. The main window is titled "L3VPN_NM-SRTE-RP-PE-A-7". It shows a "color" configuration pane with a table containing one entry: "70". A secondary pane titled "color{70}" is open, showing the "id" set to "70" and the "ipv4" section with "Enable ipv4" checked and a "prefix" table containing one entry: "70.70.70.0/30".

5. Click **X** in the top right corner to close the Color pane.
6. Commit your changes.
7. Check that the new route policy appears in the table.

8. Create the other route policies listed above in the same manner.

Step 3: Create a VPN profile

In this step, we will create a VPN profile for each route policy. The VPN profile will be referenced from the L3VPN service. This will bind the route policy to the L3VPN service.

We will create the following VPN profiles for each of the route policies we created in the previous step:

- L3VPN_NM-SRTE-RP-PE-A-7
- L3VPN_NM-SRTE-RP-PE-B-7
- L3VPN_NM-SRTE-RP-PE-C-7

Procedure

1. Go to **Services & Traffic Engineering > Provisioning > L3vpn > VPN Profiles**
2. Click + to create a new VPN profile.
3. From the dropdown list, choose the required route policy.
4. Commit your changes.
5. Check that the new VPN profile appears in the table.
6. Create the other VPN profiles listed above in the same manner.

Step 4: 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 an ie-profile which in turn points to a VPN profile which 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.

1. Go to **Services & Traffic Engineering > Provisioning > L3vpn > L3vpn Service**.
2. Click + to create a new service and give it a unique name. Click **Continue**.
3. Create an ie-profile which is a container that defines the route distinguisher (RD), route targets, and the export/import route policy. We will create an ie-profile 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. Under ie-profile, click + to create a new ie-profile and give it a unique name.
 - b. Enter the route distinguisher that will differentiate the IP prefixes and make them unique.
 - c. Define the required VPN targets, including route targets and route target types (import/export/both).

- d. Under vpn-policies, in the export policy dropdown list, choose the relevant VPN profile (which contains the route policy). This forms the association between the VPN and the ODN template that defines the SLA.

The screenshot displays two configuration panels. The left panel, titled 'L3VPN_NM-SRTE-ODN-70', shows the 'ie-profiles' section with a table of three profiles:

ie-profile-id	rd
L3VPN_NM_SR_ODN-IE-PE-A-7	0:70:70
L3VPN_NM_SR_ODN-IE-PE-B-7	0:70:71
L3VPN_NM_SR_ODN-IE-PE-C-7	0:70:72

The right panel, titled 'ie-profile{L3VPN_NM_SR_ODN-IE-PE-A-7}', shows the configuration for the selected profile. The 'vpn-policies' section has the 'export-policy' dropdown menu highlighted with a red box, showing the selection of 'L3VPN_NM-SRTE-RP-PE-A-7'.

- e. Click X in the top right corner when you are done.
 - f. Similarly, create the other IE profiles.
4. Define each VPN endpoint individually – PE-A, PE-B, and PE-C.
 - a. Under vpn-nodes, click +, select the relevant device from the dropdown list and click **Continue**.
 - b. Enter the local autonomous system number for network identification.
 - c. Select the ie-profile you created in the previous step.
 - 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 port-id field, provide the name of a loopback interface that will be dedicated for this VRF.
 - Under ip-connection > IPv4 > address-allocation-type, choose **static-address**.
 - Under static-addresses, there is a table where you can create a list of IP addresses for network access for this endpoint. After creating at least one address, you can select a primary address. In the address table, click + to create a new address and provide a unique ID. Click **Continue**. Specify the IP address and prefix length of the loopback interface.
 - Click **X** in the top right corner to return to the VPN network access parameters.
 - Select the address you just created from the dropdown list in the primary-address field.

- Define BGP routing protocol parameters, including the peer AS number and local AS number, IP address type (IPv4), the IP address of the BGP neighbor, and the number of hops allowed between the BGP neighbor and the PE device.
 - e. Click **X** in the top right corner when you are done.
 - f. Repeat these steps for each endpoint.
5. Commit your changes or click **DryRun** to check what will be configured on the devices before you commit.
 6. Check that the new L3VPN service appears in the table and its provisioning state is **Success**.

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

1. In the L3VPN Service table, click on the service name or click ... in the Actions column and choose **View** from the menu. The map opens and the service details are shown to the right of the map.

or

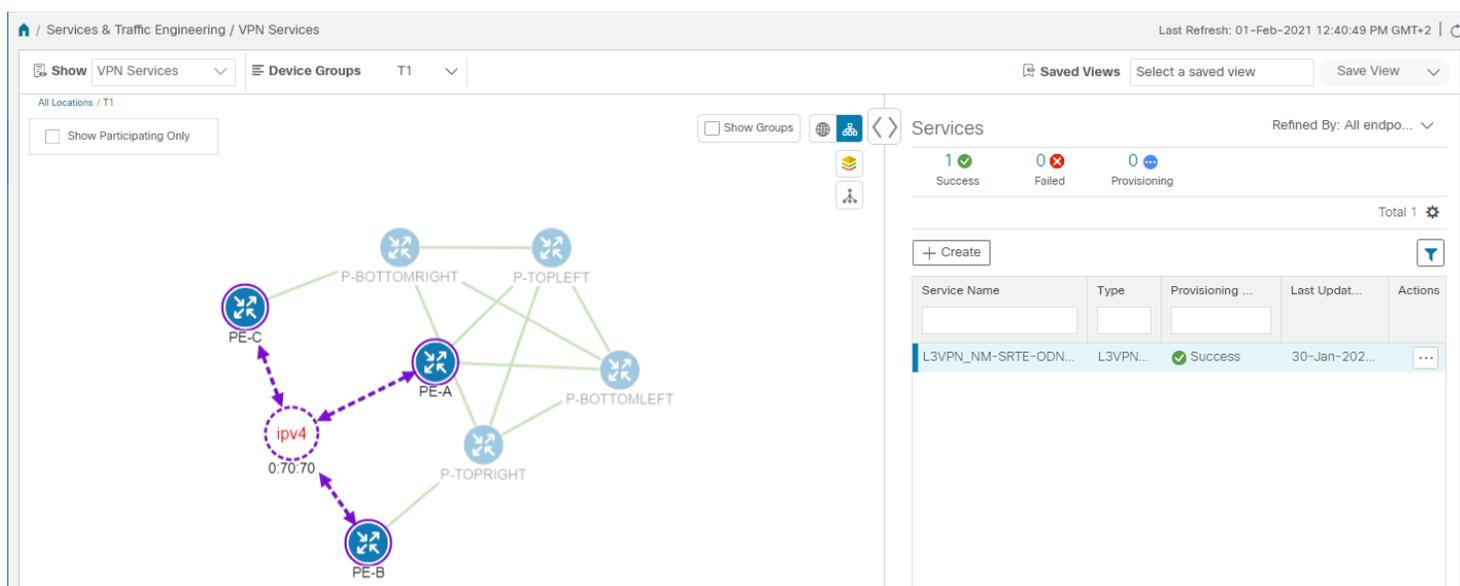
Go to **Services & Traffic Engineering > VPN Services**.

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.

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



The screenshot displays the 'Services & Traffic Engineering / VPN Services' interface. On the left, a logical map shows a network topology with several PE devices (PE-A, PE-B, PE-C) and intermediate nodes (P-BOTTOMRIGHT, P-TOPLEFT, P-BOTTOMLEFT, P-TOPRIGHT). A dashed purple line represents the VPN overlay, connecting PE-C, PE-A, and PE-B. A central node is labeled 'ipv4 0.70.70'. On the right, a 'Services' table is shown, filtered by 'All endpoints'. The table contains one entry: 'L3VPN_NM-SRTE-ODN...' with a provisioning state of 'Success' and a last update date of '30-Jan-202...'. Above the table, there are statistics for Success (1), Failed (0), and Provisioning (0).

Service Name	Type	Provisioning ...	Last Updat...	Actions
L3VPN_NM-SRTE-ODN...	L3VPN...	Success	30-Jan-202...	...

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

- 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.

Service Name	Type	Provisioning ...	Last Updat...	Actions
L3VPN_NM-SRTE-ODN...	L3VPN...	✓ Success		View Details Edit / Delete

- 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.

Services & Traffic Engineering / VPN Services Last Refresh: 01-Feb-2021 01:08:05 PM GMT+2

Show VPN Services | Device Groups T1 | Saved Views Select a saved view | Save View

All Locations / T1

Show IGP Path | Show Participating Only

Show Groups

The diagram illustrates a network topology with three edge routers: PE-A, PE-B, and PE-C. PE-A is the source, and PE-B and PE-C are destinations. There are four intermediate nodes: P-BOTTOMLEFT, P-BOTTOMRIGHT, P-TOPLEFT, and P-TOPRIGHT. Purple arrows indicate the computed paths. From PE-A, one path goes through P-TOPLEFT to PE-B, and another goes through P-TOPLEFT to P-BOTTOMLEFT, then P-BOTTOMRIGHT, and finally to PE-C. There are also direct paths from PE-A to P-TOPRIGHT and P-BOTTOMRIGHT.

Service Details

Name: L3VPN_NM-SRTE-ODN-70
 Provisioning State: ✓ Success

Summary Transport

SR POLICY Selected 2 / Total 6

	He...	En...	C...	A...	O...	A...
<input type="checkbox"/>	PE-B	PE-A	70	↑	↑	...
<input type="checkbox"/>	PE-C	PE-A	70	↑	↑	...
<input checked="" type="checkbox"/>	PE-A	PE-B	71	↑	↑	...
<input type="checkbox"/>	PE-C	PE-B	71	↑	↑	...
<input type="checkbox"/>	PE-B	PE-C	72	↑	↑	...
<input checked="" type="checkbox"/>	PE-A	PE-C	72	↑	↑	...

- 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.

The screenshot displays the network configuration interface. On the left, a network diagram shows a hierarchical structure of nodes: PE-A at the top, connected to P-BOTTOMLEFT and P-TOPLEFT. P-BOTTOMLEFT is connected to P-BOTTOMRIGHT, and P-TOPLEFT is connected to P-TOPRIGHT. P-BOTTOMRIGHT is connected to PE-C, and P-TOPRIGHT is connected to PE-B. The diagram includes labels for Adj. SID (24010, 24007, 24011) and a tooltip for the P-TOPRIGHT node showing 'bgp_c_71_ep_100.100.100.6_discr_100'. On the right, the 'Service Details' panel shows the service name 'L3VPN_NM-SRTE-ODN-70' with a 'Success' provisioning state. Below this is a table of SR policies.

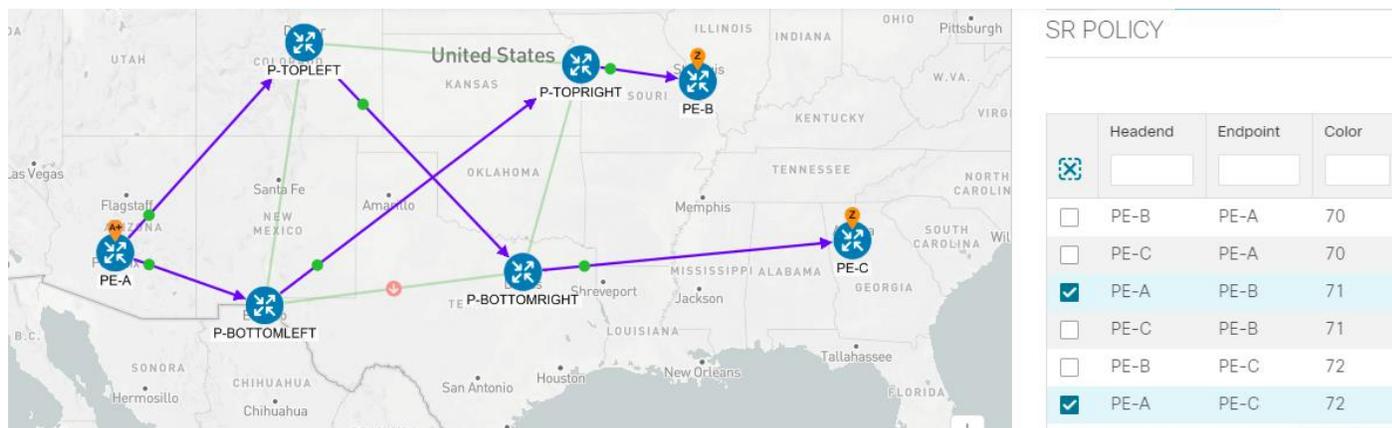
	Head...	End...	Co...	Ad...	Op...	Act...
<input type="checkbox"/>	PE-B	PE-A	70	↑	↑	...
<input type="checkbox"/>	PE-C	PE-A	70	↑	↑	...
<input checked="" type="checkbox"/>	PE-A	PE-B	71	↑	↑	...
<input type="checkbox"/>	PE-C	PE-B	71	↑	↑	...
<input type="checkbox"/>	PE-B	PE-C	72	↑	↑	...
<input checked="" type="checkbox"/>	PE-A	PE-C	72	↑	↑	...

Step 6: 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 > 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-C



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.

Scenario 2 – Mandate a static path for an EVPN-VPWS service using an explicit SR-TE policy

Scenario Context

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) 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.

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.

Assumptions and Prerequisites

For transport mapping to L2VPN service, devices must be configured with the “l2vpn all” command.


```
wiY2hhbmdlX3B3ZCI6ImZhbHNlIiwizXhwIjoxNjE2NDU5OTIwLzJpYXQiOjE2MTY0MzExMjAsImZpcnN0X25hbWU
i0iJqb2huIiwianRpIjoiU1QtODQtOFVlWXM5bEt3R2d1Z3RIYj14MzVmTF1NTGVVVRlp6OURyNGpoeFcxakhsV01V
YXdXSXgxbUdTd01aRC1tOEK1S2Z0amI2ZmlWTUhlYnBYYjBMMFZqRFc2wVppUFVUbHRpNFVpZnNUeG9aQ284WlpPW
Ec2VlFjS0Mwb291WjJhc3BwanMzYnA3bHo5VkhYSlBCTz15TDNGcFRIWXRPeWJtVi1jYXMtMSJ9.Vi4k0w8KsOv5M
_08zBqWochT3k9V9Pn2NjSn5ES9c5Pf-4ds0o4kk6xuZx5_cggauIEICuUMnzmRzneST-oAuA ' \
--data-raw '{
  "input": {
    "head-end": "100.100.100.5",
    "end-point": "100.100.100.7",
    "sr-policy-path": {
      "path-optimization-objective": "igp-metric"
    }
  }
}'
```

- Note the SID list ID in the API response. You will use this when creating the SID list in the next step. For example:

```
{
  "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.

Procedure

- Go to **Services & Traffic Engineering > Provisioning > SR-TE > SID-List**.

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**.
3. Under sid, click **+** to create a new SID index and give it a numeric value. Click **Continue**.
4. Under mpls, enter the SID ID that was received in the API response in [Step 1: Prepare for Creating a SID List](#).

The screenshot displays two side-by-side configuration panels. The left panel, titled 'Sid240', shows the 'name' field set to 'Sid240' and the 'sid' section with a table containing one index with the value '1'. The right panel, titled 'sid{1}', shows the 'index' field set to '1', the 'type' set to 'mpls', and the 'label' field set to '23002'. A red box highlights the 'mpls' section and the 'label' field.

5. Click **X** in the top right corner to return to the SID list. Your new SID appears in the index table.
6. Repeat these steps to create additional SID indexes, as required.
7. Commit your changes.
8. Check that the new SID list appears in the table.
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.

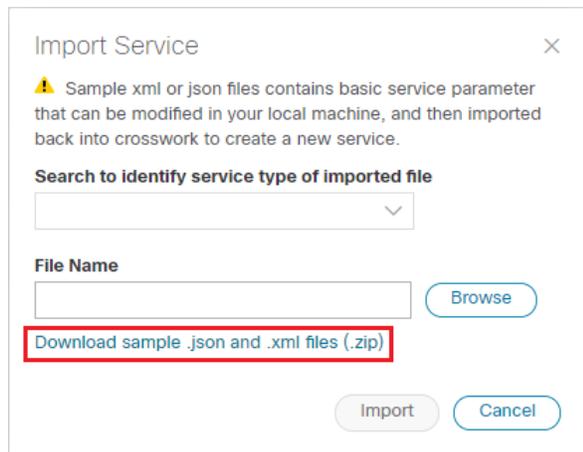
The first SR-TE policy specifies PE-C as the headend and provides the IP address of PE-A as the tail end. The second SR-TE policy specifies PE-A as the headend and provides the IP address of PE-C as the tail end.

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.

Procedure

1. Go to **Services & Traffic Engineering > Provisioning > SR-TE > Policy**.

- Click Import  above the table.
- 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.



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

```
<config xmlns="http://tail-f.com/ns/config/1.0">
  <sr-te xmlns="http://cisco.com/ns/nso/cfp/cisco-tdn-sr-te">
    <policies xmlns="http://cisco.com/ns/nso/cfp/cisco-tdn-sr-te-sr-
policies">
      <policy>
        <name>L2VPN_NM-P2P-SRTE-PE-C-240</name>
        <head-end>
          <name>PE-C</name>
        </head-end>
        <tail-end>100.100.100.5</tail-end>
        <color>240</color>
        <binding-sid>240</binding-sid>
        <path>
          <preference>1</preference>
          <explicit>
            <sid-list>
              <name>L2VPN_NM-P2P-SRTE-PE-C-240</name>
              <weight>1</weight>
            </sid-list>
          </explicit>
        </path>
      </policy>
    </policies>
  </sr-te>
</config>
```

- In the Import Service dialog, select **Policy** as the type of file to import, browse to the edited xml file, and click **Import**. 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.
- Check that the new SR-TE policy appears in the Policy table and its Provisioning State is **Success**.
- 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.

Procedure

1. Go to **Services & Traffic Engineering > Provisioning > L2vpn > L2vpn Service**.
2. Click Import  above the table.
3. If you did not download the sample .json or .xml files in [Step 3: Create an explicit SR-TE policy for each VPN endpoint by importing a file](#), do so now.
4. Open **l2nm.xml** in an XML editor.
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 : PE-A
  ▼ signaling-options [1]
    ▼ 0 {2}
      type : vpn-common:t-ldp
      ▼ t-ldp-pwe {1}
        ▼ ac-pw-list [1]
          ▼ 0 {2}
            peer-addr : 100.100.100.5
            vc-id : 240
        ▼ vpn-network-accesses [1]
          ▼ vpn-network-access [1]
            ▼ 0 {2}
              ▼ connection {2}
                encapsulation-type : vpn-common:dot1q
                ▼ dot1q-interface {2}
                  l2-access-type : vpn-common:dot1q
                  ▼ dot1q {2}
                    c-vlan-id : 240
                    physical-inf : GigabitEthernet0/0/0/2
                    id : 240
              ne-id : PE-A

```

6. Save the xml file.

7. In the Import Service dialog, select **l2vpn service** as the type of file to import, browse to the edited xml file, and click **Import**. 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.
8. Check that the new L2VPN service appears in the L2VPN Service table and its Provisioning State is **Success**.
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.

Procedure

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

Step 6: Visualize the L2VPN 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-C and vice versa, based on the explicit SR-TE policies we created.

Procedure

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

or

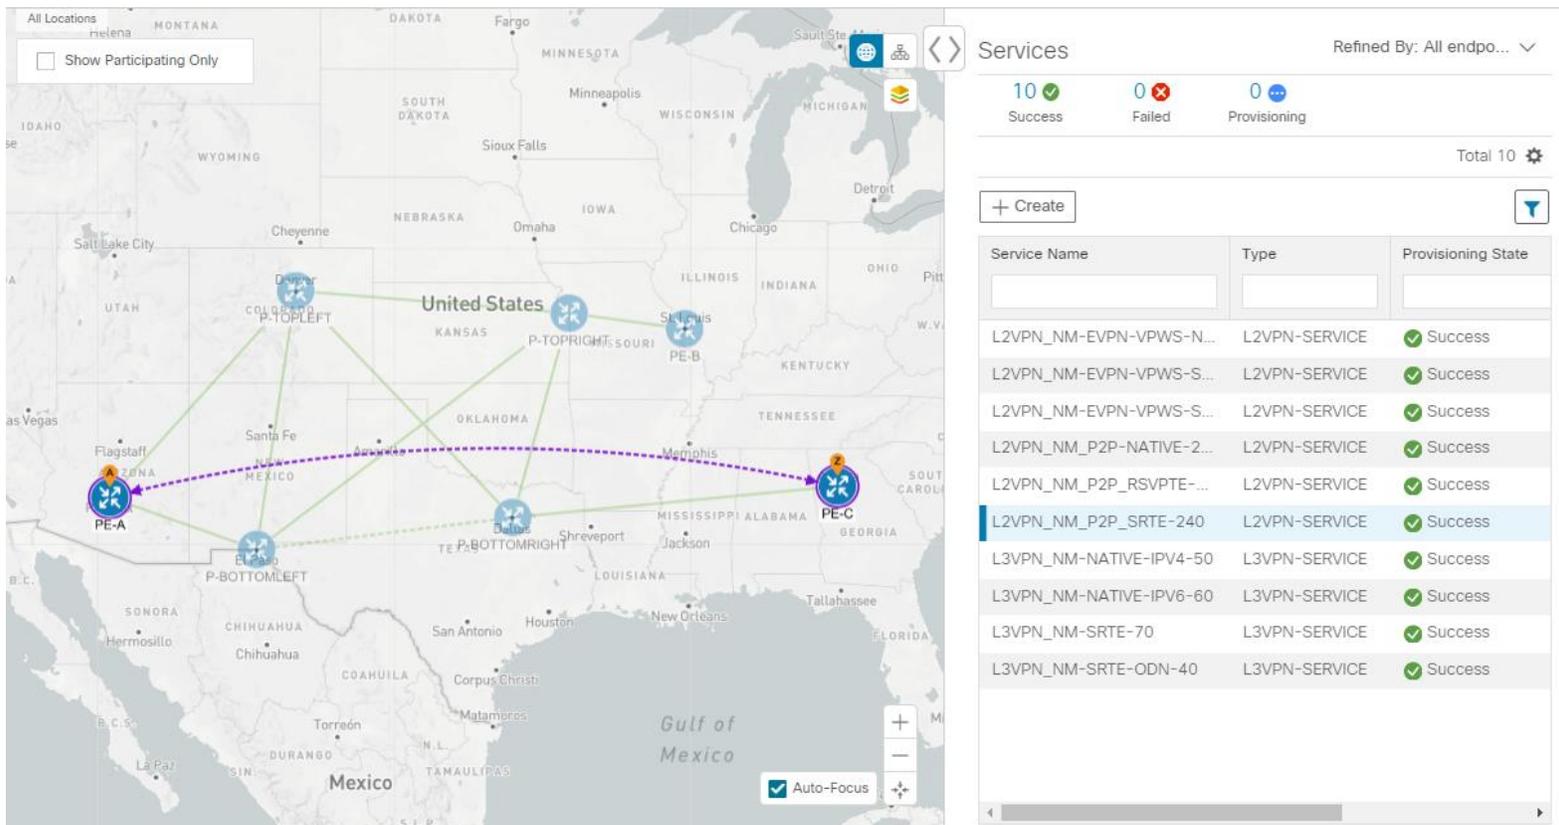
Go to **Services & Traffic Engineering > VPN Services**.

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 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.

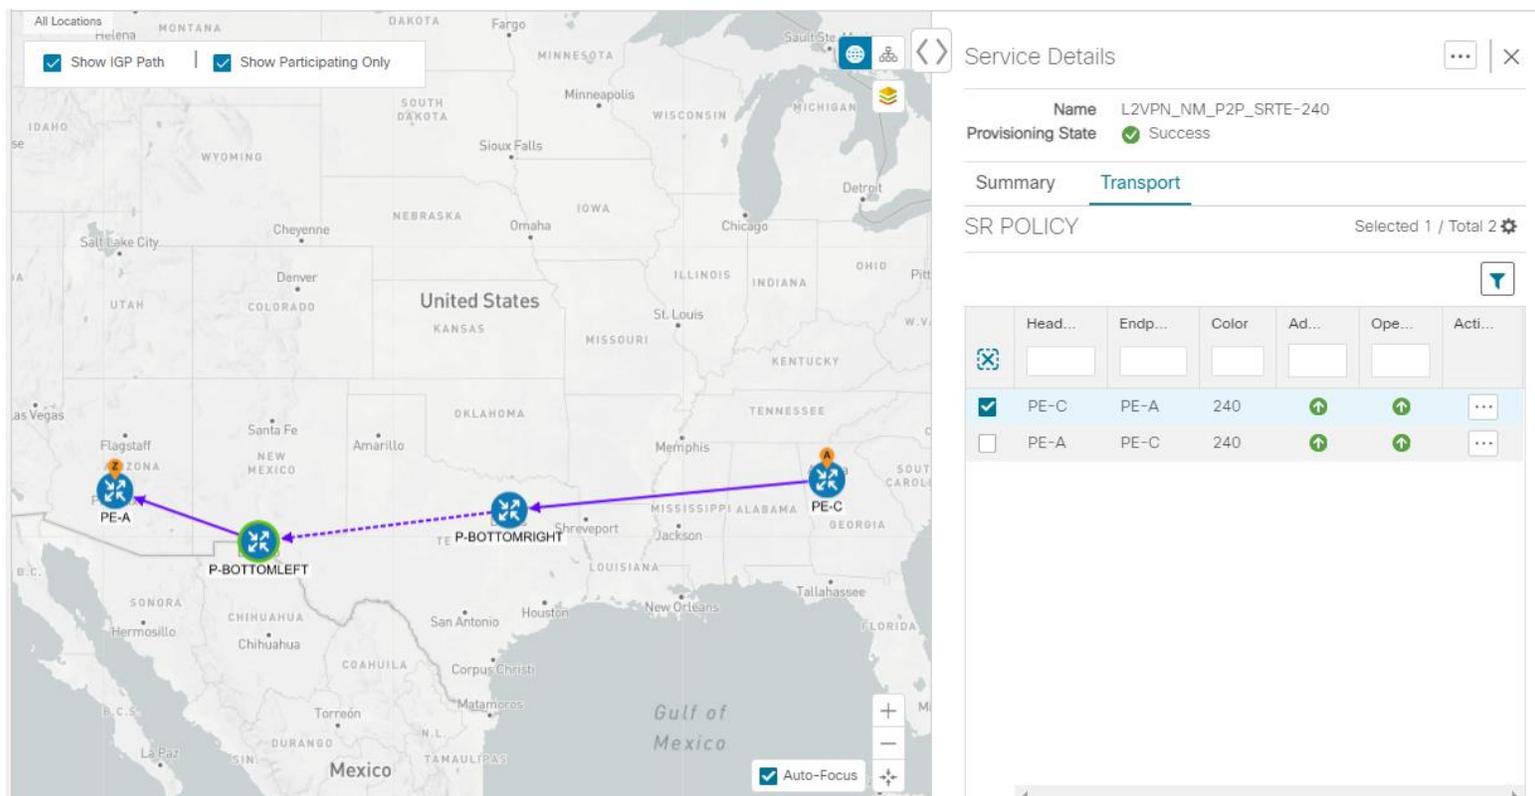


The screenshot displays a geographical map of the United States and Mexico with a VPN overlay. The map shows several endpoints (PE-A, PE-B, PE-C) and a network topology. A dashed purple line indicates the path from PE-C to PE-A. The interface includes a 'Services' panel on the right with a table of service provisioning states.

Service Name	Type	Provisioning State
L2VPN_NM-EVPN-VPWS-N...	L2VPN-SERVICE	Success
L2VPN_NM-EVPN-VPWS-S...	L2VPN-SERVICE	Success
L2VPN_NM-EVPN-VPWS-S...	L2VPN-SERVICE	Success
L2VPN_NM_P2P-NATIVE-2...	L2VPN-SERVICE	Success
L2VPN_NM_P2P_RSVPTE-...	L2VPN-SERVICE	Success
L2VPN_NM_P2P_SRTE-240	L2VPN-SERVICE	Success
L3VPN_NM-NATIVE-IPV4-50	L3VPN-SERVICE	Success
L3VPN_NM-NATIVE-IPV6-60	L3VPN-SERVICE	Success
L3VPN_NM-SRTE-70	L3VPN-SERVICE	Success
L3VPN_NM-SRTE-ODN-40	L3VPN-SERVICE	Success

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

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



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 to 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.

Scenario 3 – Provision an L2VPN service over an RSVP-TE tunnel with reserved bandwidth

Scenario Context

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

For transport mapping to L2VPN service, devices must be configured with the “l2vpn all” command.

Workflow

- [Step 1: Create an RSVP-TE tunnel for both directions of the L2VPN](#)
- [Step 2: Create the L2VPN service and attach the RSVP tunnel to the service](#)
- [Step 3: Visualize the L2VPN service on the map](#)

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.

1. Go to **Services & Traffic Engineering > Provisioning > RSVP-TE > Tunnel**.
2. Click **+** to create a new RSVP-TE tunnel and give it a unique name. Click **Continue**.
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.

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.

Device Details

Details Links

▼ Summary

Host Name	PE-A
Reachability State	✔ Reachable
Operational State	↑ OK
Node IP	172.16.1.45
Civic Address	Chennai, Tamilnadu, India, Asia, 600002
Geo Location	Latitude 30.000000, Longitude 80.000000
Device Group	All Locations > Unassigned Devices
Product Type	ciscoCRS16S
Connect To Device	🔒 SSH IPv4
Last Update	02-Mar-2021 10:55:13 PM GMT+2

▼ Routing

TE Router ID	100.100.100.5
ISIS System ID	0000.0000.0005 Level-1/2
ASN	1

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

RSVP-TE Tunnel {L2VPN_NM-P2P-RSVPTE-PE-A-2220}

signaling-type
te-types:path-setup-rsvp

head-end *
PE-A

tail-end
PE-B

te-bandwidth
technology
generic
generic
1200

p2p-primary-paths

traffic-steering

p2p-primary-path{L2VPN_NM-P2P-RSVPTE-PE-A-2220}

name *
L2VPN_NM-P2P-RSVPTE-I

path-computation-method
path-locally-computed

preference
1

optimizations

explicit-route-objects-always

Commit changes
Dry Run
Delete
Cancel

8. Click **Commit Changes**.

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

Services & Traffic Engineering / Provisioning

Services/Policies

Resource Pool

- ▼ L2VPN
 - ID-Pools
 - L2vpn Route Policy
 - L2vpn-Service
- ▼ L3VPN
 - L3vpn Route Policy
 - L3vpn-Service
 - VPN Profiles
- ▼ RSVP-TE
 - Tunnel

Tunnel

Total 5 | Last Refresh: 01-Apr-2021 11:30:58 AM GMT+3

Name	Provisioning State	Date Created	Acti...
IETF-RSVP-TE-1	Success	28-Mar-2021 09:55:47 AM G...	...
IETF-RSVP-TE-2	Failed	31-Mar-2021 12:32:28 AM G...	...
L2VPN_NM-P2P-RSVPTE-PE-A-2220	Success	17-Mar-2021 11:28:30 AM G...	...
L2VPN_NM-P2P-RSVPTE-PE-B-2220	Success	17-Mar-2021 11:28:32 AM G...	...
rsvp-TE-demeke	Success	17-Mar-2021 07:49:42 PM G...	...

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

The screenshot shows the 'RSVP-TE Tunnel Details' page. The left pane displays a map of the United States with several network devices (PE-A, PE-B, P-TOPLEFT, P-TOPRIGHT, P-BOTTOMLEFT, P-BOTTOMRIGHT) and a purple path connecting them. The right pane shows the 'RSVP-TE Tunnel Details' for tunnel ID 2220, including a summary and an explicit route object (ERO) table.

Summary

- Headend: PE-A (100.100.100.5)
- Endpoint: PE-B (100.100.100.6)
- Tunnel ID: 2220
- Description: -
- Path Name: L2VPN_NM-P2P-RSVPTE-PE-A-2220
- LSP ID: 2
- Path Type: Unknown
- Admin State: Up
- Oper State: Up

Explicit Route Object (ERO)

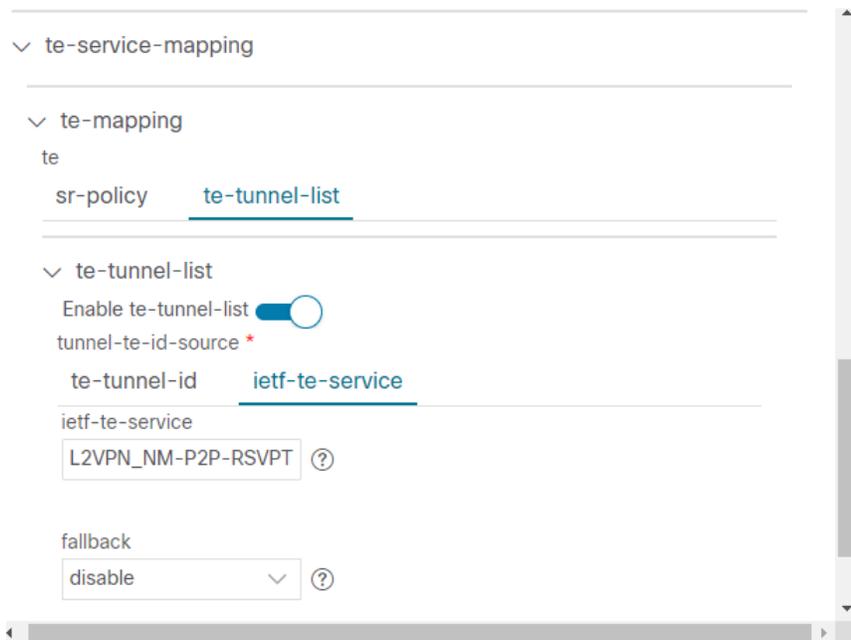
Hop	Node	IP	Interface Name
0	P-TOPLEFT	20.20.10.2	GigabitEthernet0/0/0
1	P-TOPRIGHT	20.20.10.14	GigabitEthernet0/0/0
2	PE-B	20.20.10.26	GigabitEthernet0/0/0
3	PE-B	100.100.100.6	

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 2](#) - Mandate a static path for an EVPN-VPWS service using an explicit SR-TE policy

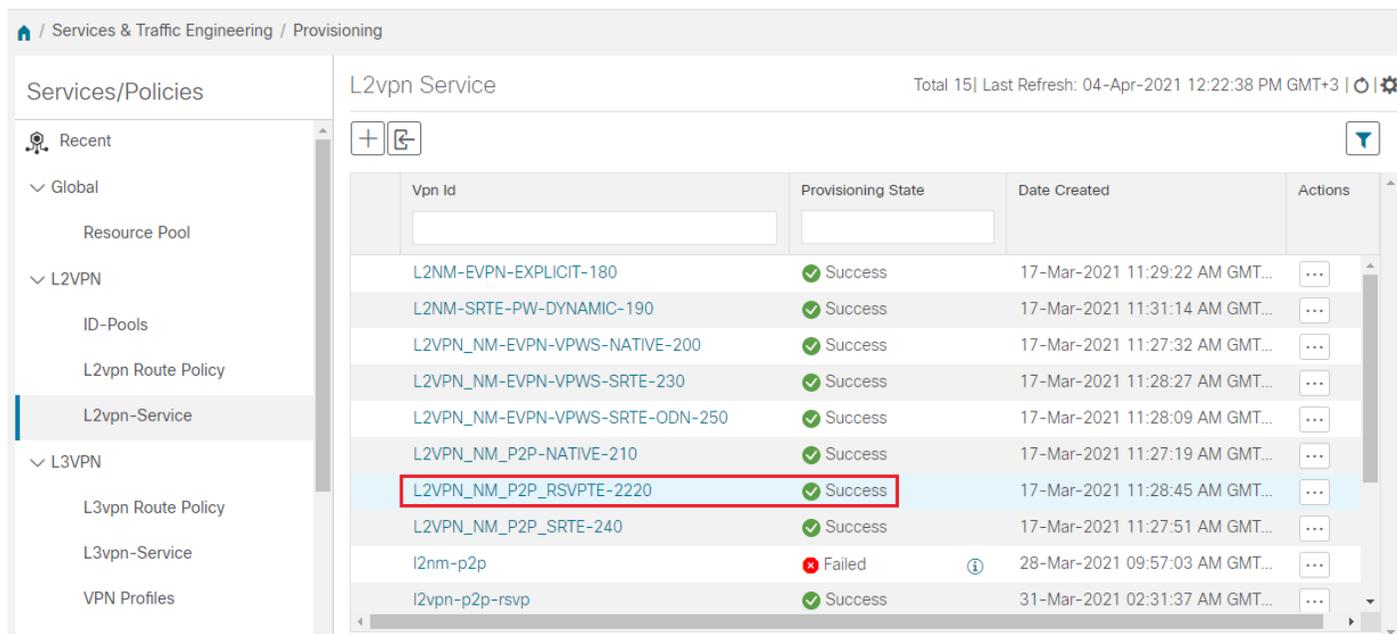
- Go to **Services & Traffic Engineering > Provisioning > L2VPN > L2vpn Service**.
- Click + to create a new service and give it a unique name. Click **Continue**.
- Choose **vpn-common:t-ldp** in the vpn-svc-type field.
- Define each VPN endpoint individually – PE-A and PE-B.
 - Under vpn-nodes, click +.
 - Select the relevant device from the vpn-node-id and ned-id dropdown lists and click **Continue**.
 - Enter the local autonomous system number for network identification.
- 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.
- 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.

- c. 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.



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

- 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).
- Follow the above steps for PE-B as well.
- 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.

Procedure

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

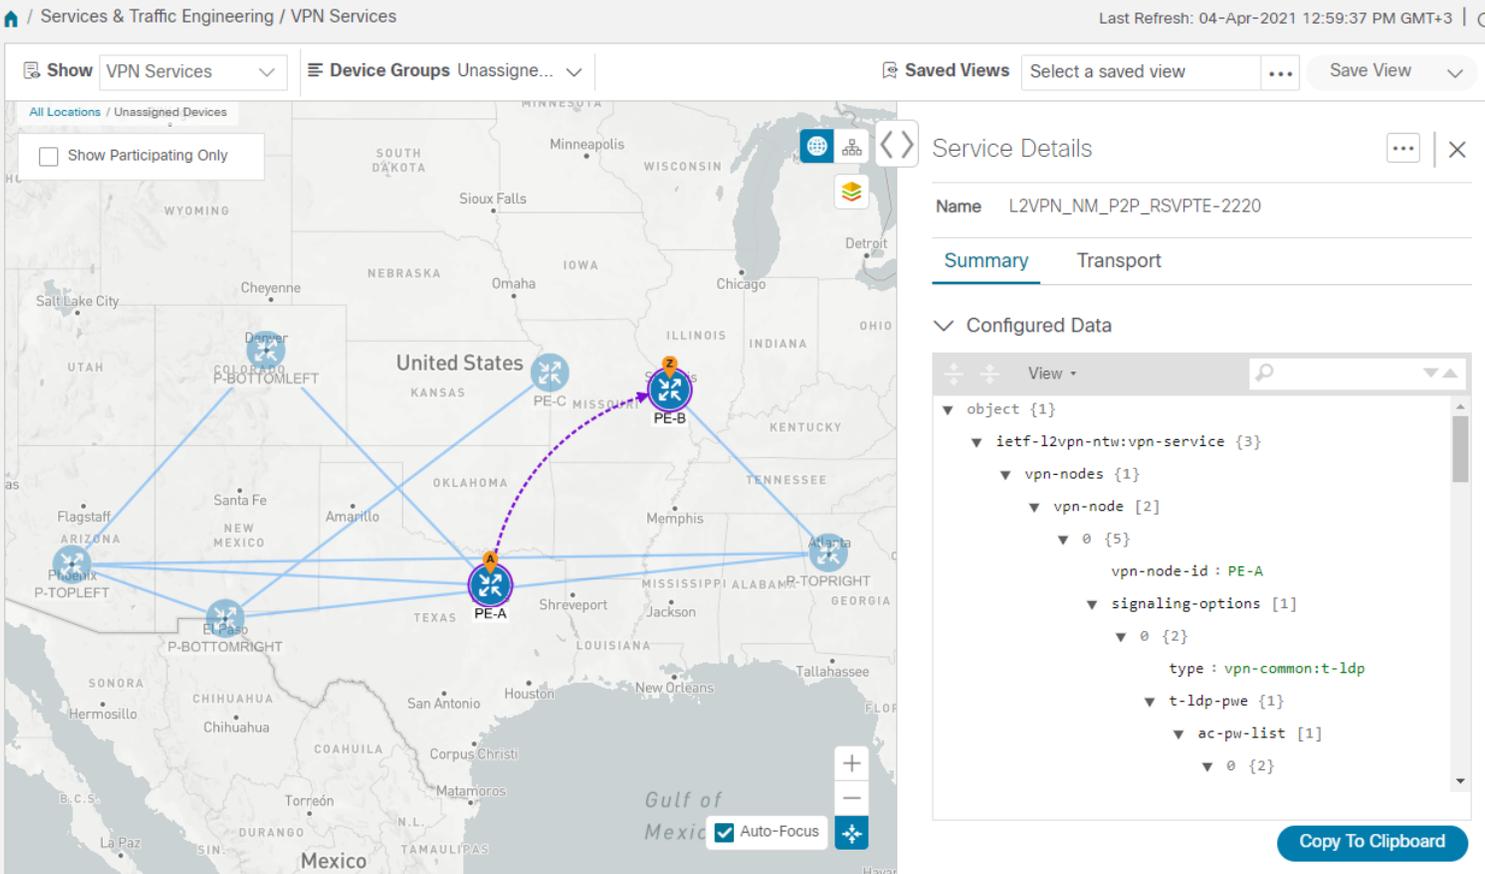
Go to **Services & Traffic Engineering > VPN Services**.

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.

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.



The screenshot shows the Cisco Crosswork Network Controller interface. The main area is a geographical map of the United States and Mexico. A VPN overlay is visible, showing a dashed purple line connecting PE-A and PE-B. The interface includes a 'Service Details' panel on the right with tabs for 'Summary' and 'Transport'. The 'Configured Data' section shows the configuration for the L2VPN service, including the VPN nodes and signaling options.

```

object {1}
  ietf-l2vpn-ntw:vpn-service {3}
    vpn-nodes {1}
      vpn-node {2}
        0 {5}
          vpn-node-id : PE-A
          signaling-options {1}
            0 {2}
              type : vpn-common:t-ldp
              t-ldp-pwe {1}
                ac-pw-list {1}
                  0 {2}

```

2. To see the hops in the route between PE-A and PE-B, 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 PE-A to PE-B and from PE-B to PE-A is shown on the logical map.

The screenshot displays the Cisco Crosswork Network Controller interface. On the left, a logical map shows a network topology with nodes PE-A, PE-B, P-TOPLEFT, P-TOPRIGHT, P-BOTTOMLEFT, and P-BOTTOMRIGHT. Purple arrows indicate the path of RSVP-TE tunnels between PE-A and PE-B. On the right, the 'Service Details' panel shows the service name 'L2VPN_NM_P2P_RSVPTE-2220' and the 'Transport' tab. Below the tab, two RSVP-TE tunnels are listed in a table:

Tunnel...	He...	En...	A...	O...	A...	
<input checked="" type="checkbox"/>	2220	PE-A	PE-B	↑	↑	...
<input checked="" type="checkbox"/>	2220	PE-B	PE-A	↑	↑	...

- 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 4 – Provision a Soft Bandwidth Guarantee with Optimization Constraints

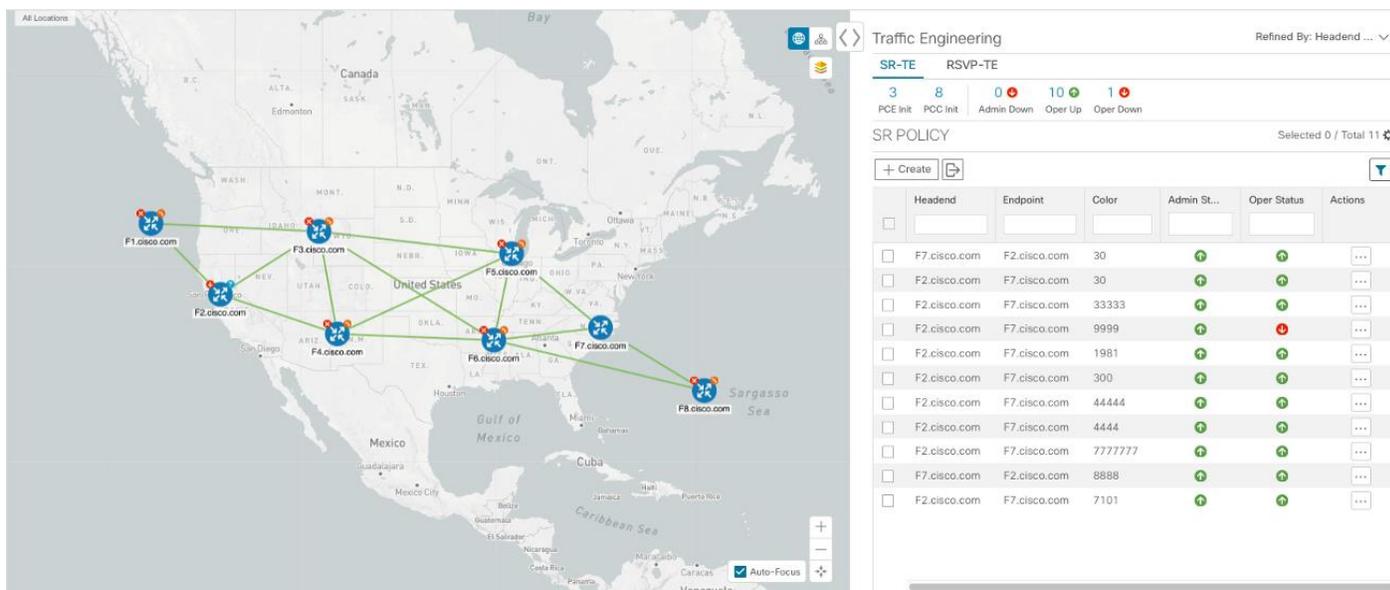
Scenario Context

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 Cisco 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

Workflow

- [Step 1: Create a BWoD SR-TE Policy with the Requested Bandwidth and Optimization Intent.](#)
- [Step 2: Enable and Configure BWoD](#)
- [Step 3: Verify that the policy's operational state is now Up and view the path on the map](#)

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

1. Go to **Services & Traffic Engineering > Provisioning > SR-TE > Policy**.
2. Click **+** to create a new SR-TE policy and give it a unique name. Click **Continue**.

3. Define the endpoints:

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

4. Define the parameters upon which the path will be computed:

- a. Under path, click +.
- b. 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
metric-type
te

pce ?

> metric-margin

> constraints *

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

- In the **Bandwidth** field enter the requested bandwidth in Kbps. In this case, we are requesting **250 Mbps** or 250000 Kbps.

head-end * Selected 0 / Total 1 ⚙

+ / ✎ / 🗑 ⏴

name

F2.cisco.com

tail-end * ?

color * ?

binding-sid ?

path * Selected 0 / Total 1 ⚙

+ / ✎ / 🗑 ⏴

preference

123

bandwidth ?

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

Policy

+ / 🔄

Name	Provisioning State	Date Created
bwOD-pcc	✔ Success	11-Feb-2021 03:27:17 AM PST
bwOD-pcc_F2_F7	✔ Success	11-Feb-2021 03:35:03 AM PST
srte_c_300_ep_100.100.100.3222222	✔ Success	10-Feb-2021 06:52:38 PM PST

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

The screenshot displays the Cisco Crosswork Network Controller interface. On the left, a map of the United States shows two locations marked with orange dots and labeled 'F2.cisco.com' and 'F7.cisco.com'. A purple line connects these two locations across the country. On the right, the 'SR Policy Details' panel is visible, showing a 'Summary' section with various parameters. The 'Oper State' is highlighted with a red box and set to 'Down'. Below the summary, there is a 'Path' table with columns for Segment, Segment Type, Label, IP, Node, Interface, and Skd Type, which currently shows 'No Rows To Show'.

Step 2: Enable and Configure BWoD

1. Go to **Services & Traffic Engineering > Bandwidth on Demand**.
2. Toggle the Enable switch to **True** and enter **80** to set the utilization threshold percentage. To find descriptions of other options, simply hover the mouse over .
3. Click **Commit Changes**.

The screenshot shows the 'Bandwidth On Demand' configuration page. The 'Configuration' tab is selected, and the 'Basic' sub-tab is active. The 'Enable' switch is set to 'True'. The 'Primary Objective' is set to 'Maximize Available Bandwidth'. The 'Link Utilization' is set to '80'. The 'Re-optimization Interval' is set to '60'. The 'Metric Re-Optimization Time' is set to '01 hrs : 30 mins'. At the bottom, there are three buttons: 'Commit Changes', 'Get Default Values', and 'Discard Changes'.

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

1. Go to **Services & Traffic Engineering > Provisioning**.
2. In the Policy table, locate and select the path computed for the endpoints.
3. The path is shown as an overlay on the map. Check the **Show IGPPath** check box to see the physical path between the endpoints.

The screenshot displays the Cisco Crosswork Network Controller interface. On the left, a map of the United States shows a path overlay connecting several nodes labeled F1.cisco.com through F8.cisco.com. The path starts at F1.cisco.com (San Diego) and ends at F7.cisco.com (New York). The path is highlighted in purple and green. On the right, the 'SR Policy Details' panel is visible. The 'Oper State' is highlighted in red and shows 'Up'. Below the details, a 'Path' table is shown:

Segment	Segment Type	Label	IP	Node	Interface
0	Node SID	16007	192.168.100.7	F7.cisco.com	

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.

3 Bandwidth and Network Optimization

Overview

Objective

Tactically optimize the network during times of congestion in real-time.

Challenge

Network congestion leads to poor end-customer experiences. If you have poor connections, slow streaming video, and packet loss, your users will be dissatisfied, which leads to a poor perception of your service in the marketplace. In the worst-case scenario, your network problems lead to service level agreement (SLA) or contract violations and the loss of your brand equity. Network operators need a toolset to help automate bandwidth optimization and efficiently steer traffic with little operator intervention.

Solution

Cisco Crosswork Network Controller provides local congestion mitigation (LCM) as a solution for bandwidth management and congestion mitigation. This is enhanced functionality introduced in Cisco Crosswork Network Controller 2.0.

Local Congestion Mitigation (LCM)

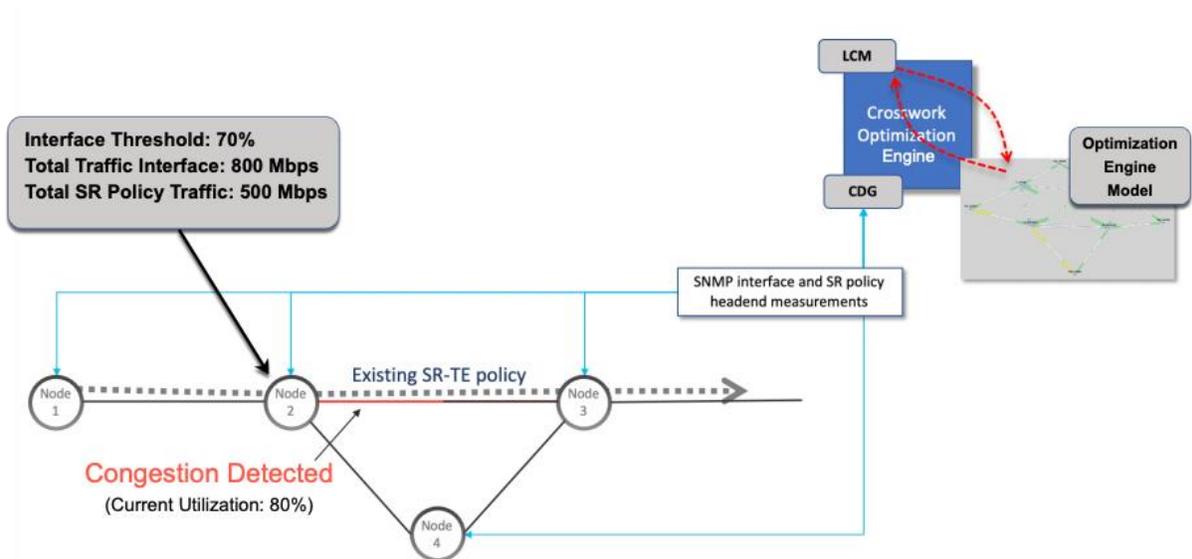
Instead of optimizing for bandwidth resource in the network by rerouting traffic in the entire network (end-to-end path optimization), LCM checks the capacity locally, in and around the congested area, at an interface level and reroutes traffic between the endpoints of the congested interface (local interface-level optimization). Focusing on the problem 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.

LCM provides recommendations to divert the minimum amount of traffic away from the congested interface to bring it out of congestion. LCM performs the collection of SR-TE policy and interface counters via 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) and TTE SR-TE policies are created only at the device on either side of the congested link, with the shortest paths possible that do not congest interfaces elsewhere.

How Does LCM Work?

1. LCM analyzes the Cisco Crosswork Optimization Engine Model (a real-time topology and traffic representation of the physical network) on a regular cadence. This cadence can be configured but should be greater than or equal to SNMP polling (5 mins).

- In the following example, during one of its polling checks, LCM detects congestion when Node 2 utilization goes above the 70% utilization threshold.



- LCM calculates how much traffic is eligible to divert.

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

Eligible traffic is computed 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-TE policy traffic = Eligible traffic that can be optimized

This process must account for any ECMP splitting of SR-TE policies to ensure the proper accounting of SR-TE policy traffic. In this example, the total traffic on congested Node 2 is 800 Mbps. The total traffic of all SR-TE policies routed over Node 2 is 500 Mbps. The total traffic that LCM can divert in this example is 300 Mbps:

$$800 \text{ Mbps} - 500 \text{ Mbps} = 300 \text{ Mbps}$$

- 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:

$$800 \text{ Mbps} - 700 \text{ Mbps (70% threshold)} = 100 \text{ Mbps}$$

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

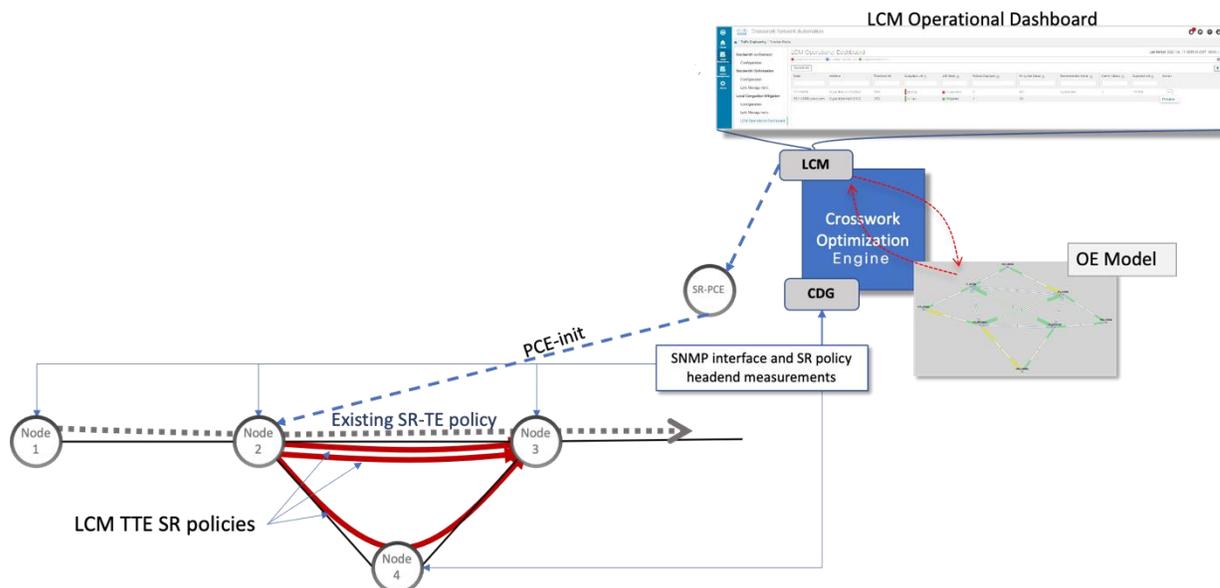
- LCM determines how many TTE SR-TE policies are needed and their paths. The ratio of how much LCM eligible traffic can stay on the shortest path to the amount that must be rerouted, will determine the number of TTE SR-TE policies that are needed on the shortest versus alternate paths, respectively.

In this case, LCM needs to divert 1/3 of the total eligible traffic (100 Mbps out of 300 Mbps) away from the congested link. Assuming a perfect ECMP, LCM estimates that 3 tactical SR-TE policies are required to create this traffic split: 1 tactical SR-TE policy will take the diversion path and 2 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 3 TTE SR-TE policies (each expected to route approximately 100 Mbps) to be deployed from Node 2 to Node 3 via SR-PCE:

- 2 TTE SR-TE policies to take a direct path to Node 3 (200 Mbps)
- 1 TTE SR-TE 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-TE policies, LCM continues to monitor the deployed TTE policies and will recommend modifications or deletions as needed in the LCM Operational Dashboard. TTE SR-TE 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-TE policy churn throughout the LCM operation.

Usage Scenarios

We will walk you through the following usage scenario that illustrates the execution of bandwidth-constrained optimization and LCM:

[Scenario 5 - Use Local Congestion Mitigation \(LCM\) to reroute traffic on an over-utilized link](#)

Additional Resources

[Cisco Crosswork Optimization Engine Documentation](#)

Scenario 5 - Use Local Congestion Mitigation (LCM) to reroute traffic on an over-utilized link

Scenario Context

In this scenario, we will enable LCM and observe the congestion mitigation recommendations to deploy TTE policies when utilization on a device's interfaces surpasses the defined utilization threshold. We will preview the recommended TTE policies before committing them to mitigate the congestion.

This example uses the following topology:



We will observe the actions taken when the link between PE1-ASR9k and P1-ASR9k becomes over-utilized. Note that there is currently no indication of congestion on that link.

Assumptions and Prerequisites

The following is a non-exhaustive list of high-level requirements for proper LCM operation:

Congestion Evaluation

LCM requires traffic statistics from the following:

- SNMP interface traffic measurements
- SNMP headend SR-TE policy traffic measurements

Congestion Mitigation

- The headend device must support PCE-initiated SR-TE policies with autoroute steering.
- The headend device must support Equal Cost Multi-Path (ECMP) across multiple parallel SR-TE policies.

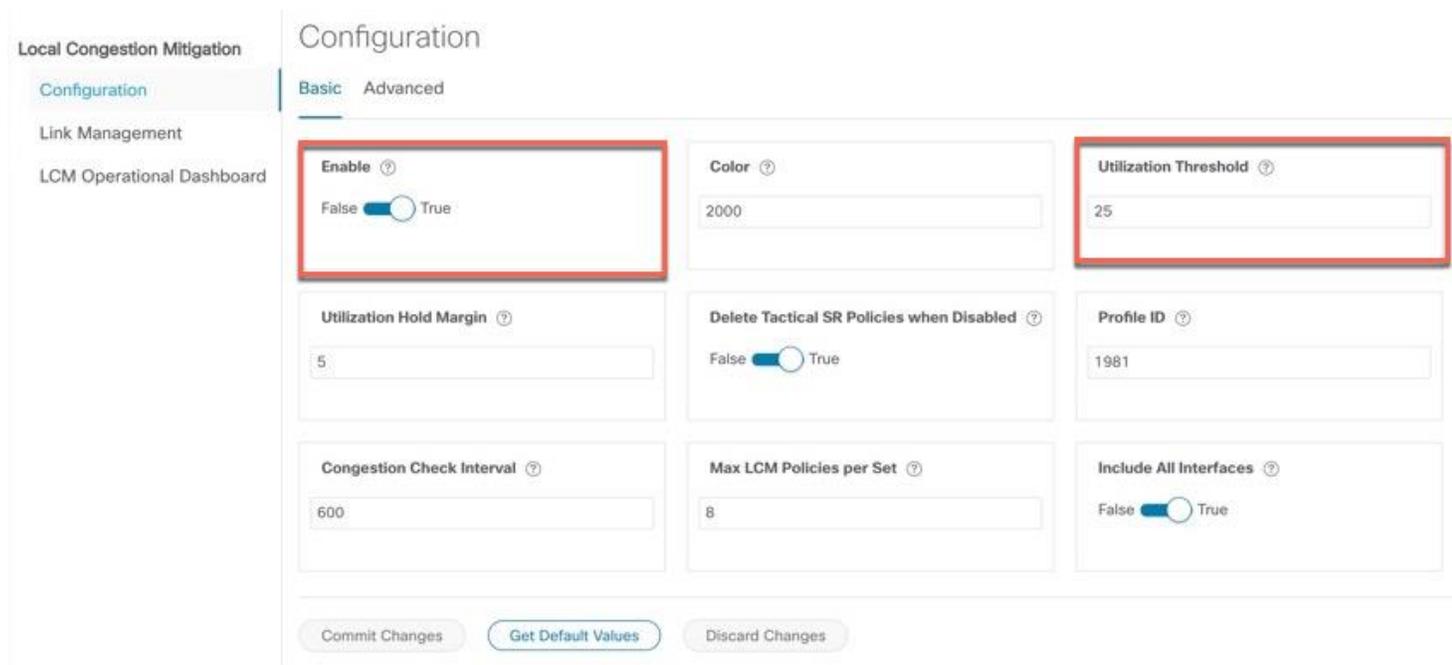
As an example, the Cisco ASR 9000 is a platform that fulfills all requirements and is supported in Crosswork Network Controller 2.0. For more information, please contact your Cisco Account representative.

Workflow

- [Step 1: Enable LCM and configure the utilization threshold](#)
- [Step 2: View link congestion on the map](#)
- [Step 3: View TTE SR-TE policy recommendations in the LCM dashboard](#)
- [Step 4: Validate the TTE SR-TE policy deployment](#)
- [Step 5: Remove the TTE SR-TE policies upon LCM recommendation](#)

Step 1: Enable LCM and configure the utilization threshold

1. Go to **Services & Traffic Engineering > Local Congestion Mitigation > Configuration**.
2. Toggle the Enable switch to **True** and enter the global utilization threshold you want to set. In this case, the threshold is set at 25%. To see information other configuration options, simply hover the mouse over .



The screenshot shows the 'Configuration' page for Local Congestion Mitigation. The 'Basic' tab is selected. The 'Enable' switch is set to 'True'. The 'Utilization Threshold' is set to '25'. Other fields include 'Color' (2000), 'Utilization Hold Margin' (5), 'Delete Tactical SR Policies when Disabled' (False), 'Profile ID' (1981), 'Congestion Check Interval' (600), 'Max LCM Policies per Set' (8), and 'Include All Interfaces' (False). Buttons for 'Commit Changes', 'Get Default Values', and 'Discard Changes' are at the bottom.

3. Click **Commit Changes**.

Step 2: View link congestion on the map

The link between PE 1-ASR9k and P1-ASR9k is now congested. Let's see that on the map.

1. Go to **Services & Traffic Engineering > Traffic Engineering**. Note that the link is Orange, indicating higher utilization.

- Click on the link to view link details, including utilization information. Note that utilization on the PE1-ASRk interfaces has surpassed 25%, the threshold that was defined in the LCM configuration.

Link Details

Summary Segment Routing RSVP-TE

Name GigabitEthernet0/0/0/4-GigabitEthernet0/0/0/1
 State Up
 Link Type L3 ISIS IPV4
 ISIS Level 2
 Last Update 2020-Dec-19, 10:34:46 (GMT +08:00)

	A Side	Z Side
Node	PE1-ASR9k	P1-ASR9k
TE Router ID	192.168.60.11	192.168.60.21
IF Name	GigabitEthernet0/0/0/4	GigabitEthernet0/0/0/1
IF Description	GigabitEthernet0/0/0/4	GigabitEthernet0/0/0/1
Type	ETHERNETCSMACD	ETHERNETCSMACD
IP Address	10.10.1.1	10.10.1.2
Utilization	25.4% (254.3Mbps/1Gbps)	0% (0Bps/1Gbps)
IGP Metric	1	1
Delay Metric	1	1

Step 3: View TTE SR-TE policy recommendations in the LCM 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.

- Go to **Services & Traffic Engineering > Local Congestion Mitigation > LCM Operational Dashboard**.

In this case, the dashboard also shows that the utilization on the PE1-ASRk interfaces has surpassed 25%. In the Recommended Action column, there is a recommendation to deploy 2 TTE policy solution sets to address the congestion. The Expected Util column shows the expected utilization of the interface if the recommended action is committed.

LCM Operational Dashboard

Local Congestion Mitigation

Configuration

Link Management

LCM Operational Dashboard

Last Refresh: 2020-Dec-19, 10:42:01 (GMT +08:00)

Committed Interfaces (2) | Mitigating Interfaces (0) | Mitigated Interfaces (0)

Commit All

Node	Interface	Thresho...	Eval... ↑	LCM State ↑	Pol... ↑	Pol... ↑	Reco... ↑	C... ↑	Expected ... ↑	Actions
PE1-ASR9k	GigabitEthernet0/0/0/5	25%	34.85%	Congested	0	-	Create Set	-	23.78%	...
PE1-ASR9k	GigabitEthernet0/0/0/4	25%	25.43%	Congested	0	-	Create Set	-	19.07%	...

- To preview the recommended TTE policies, click ... in the Actions column and choose **Preview**. You will see a list of TTE policies for the node and a specific interface, in this case, GigabitEthernet0/0/0/4. Select a TTE policy to see its representation on the map.

Recommended TTE Policies (Preview)

Node PE1-ASR9k
Interface GigabitEthernet0/0/0/4

Headend	Endpoint	Color	Recommended Action
PE1-ASR9k	P1-ASR9k	2000	CREATE
PE1-ASR9k	P1-ASR9k	2001	CREATE
PE1-ASR9k	P1-ASR9k	2002	CREATE
PE1-ASR9k	P1-ASR9k	2003	CREATE

Back To LCM Dashboard

- After you are done viewing the recommended TTE policies on the map, click **Back to LCM Dashboard**.
- If you are satisfied with the LCM recommendations, click **Commit All**. The LCM State column changes to **Mitigating**. Note that all LCM recommendations must be committed in order to mitigate congestion and produce the expected utilization as shown in the LCM Dashboard. The mitigation solution is based on all LCM recommendations being committed because of dependencies between solution sets.

Local Congestion Mitigation

Configuration

Link Management

LCM Operational Dashboard

LCM Operational Dashboard

Last Refresh: 2020-Dec-19, 10:43:37 (GMT +08:00) | ⚙

⚠ Congested Interfaces (0) | 🟡 Mitigating Interfaces (3) | 🟢 Mitigated Interfaces (0)

Commit All

Node	Interface	Thresho...	Eval... ?	LCM State ?	Policies D... ?	Policy Set... ?	Reco... ?	Com... ?	Expected ... ?	Actions
PE1-AS...	GigabitEt...	25%	34.85%	🟡 Mitigating	2	-	-	-	23.78%	⋮
PE1-AS...	GigabitEt...	25%	25.43%	🟡 Mitigating	4	-	-	-	19.07%	⋮

Step 4: Validate the TTE SR-TE policy deployment

1. Click  in the top right corner to open the **Events** window in which you can monitor LCM events. The Source column indicates the type of event, in this case, Optima LCM. You should see events for the LCM recommendations, the commit actions, as well as any exceptions.

Events Total 53

Description	Time	Severity	Source
Recommendation committed	2020-Dec-19, 10:43:52 (GMT +08:00)	INFO	Optima LCM
A new recommendation has been created: 6 creates, 0 updates, 0 delet...	2020-Dec-19, 10:40:17 (GMT +08:00)	INFO	Optima LCM
A new recommendation has been created: 2 creates, 0 updates, 0 delet...	2020-Dec-19, 10:18:48 (GMT +08:00)	INFO	Optima LCM
A new recommendation has been created: 7 creates, 0 updates, 0 delet...	2020-Dec-19, 10:08:48 (GMT +08:00)	INFO	Optima LCM
A new recommendation has been created: 5 creates, 0 updates, 0 delet...	2020-Dec-19, 09:58:47 (GMT +08:00)	INFO	Optima LCM
Recommendation committed	2020-Dec-17, 01:51:16 (GMT +08:00)	INFO	Optima LCM
A new recommendation has been created: 0 creates, 0 updates, 4 delet...	2020-Dec-17, 01:50:22 (GMT +08:00)	INFO	Optima LCM
Recommendation committed	2020-Dec-17, 01:44:49 (GMT +08:00)	INFO	Optima LCM
A new recommendation has been created: 4 creates, 0 updates, 0 delet...	2020-Dec-17, 01:36:29 (GMT +08:00)	INFO	Optima LCM
A new recommendation has been created: 4 creates, 0 updates, 0 delet...	2020-Dec-16, 22:03:10 (GMT +08:00)	INFO	Optima LCM
Recommendation committed	2020-Dec-16, 21:57:03 (GMT +08:00)	INFO	Optima LCM
Unable to fully deploy solution (delete isps) to mitigate interface: PE1-A...	2020-Dec-16, 21:57:03 (GMT +08:00)	MAJOR	Optima LCM
Exception deleting LCM tactical policy: PE1-AS9k : P1-AS9k : 2002 ...	2020-Dec-16, 21:57:03 (GMT +08:00)	MAJOR	Optima LCM
A new recommendation has been created: 0 creates, 0 updates, 8 delet...	2020-Dec-16, 21:55:40 (GMT +08:00)	INFO	Optima LCM

1 to 15 of 53 First Previous Page 1 of 4 Next Last

2. You can also view the LCM Dashboard to check that the LCM state changes to **Mitigated** for all TTE policy solution sets.

Local Congestion Mitigation

- Configuration
- Link Management
- LCM Operational Dashboard**

LCM Operational Dashboard Last Refresh: 2020-Dec-19, 11:24:46 (GMT +08:00) | 

● Congested Interfaces (0) | ● Mitigating Interfaces (0) | ● Mitigated Interfaces (3) 

Commit All 

Node	Interface	Thresho...	Eval... 	LCM State 	Policies D... 	Policy Set... 	Reco... 	Com... 	Expected ... 	Actions
PE1-AS...	GigabitEt...	25%	17.43%	✔ Mitigated	2	OK	-	-	-	
PE1-AS...	GigabitEt...	25%	15.81%	✔ Mitigated	6	OK	-	-	-	

3. Confirm the TTE policy deployment visually on the topology map and in the SR-TE policy table.
 - a. Go to **Services & Traffic Engineering > Traffic Engineering > SR-TE** tab.

- b. Select one of the new SR-TE policies and view the SR-TE policy details (click ... in the Actions column and choose **View**). Note that the Policy Type is Local Congestion Mitigation.

The screenshot displays the SR Policy Details pane for a Local Congestion Mitigation policy. The summary information is as follows:

- Headend:** PE1-ASR9k (192.168.60.11)
- Endpoint:** P1-ASR9k (192.168.60.21) / 192.168.60.21
- Color:** 2000
- Description:** -
- Path Name:** lcm_to_P1-ASR9k_c_2000
- Policy Type:** Local Congestion Mitigation
- Admin State:** Up
- Oper State:** Up
- Binding SID:** 1005023
- Profile ID:** 1981

The Path table below the summary shows the following details:

Segment	Segment Type	Label	IP	Node	In
0	Node SID	650501	192.168.6...	P1...	

- c. Close the SR-TE policy details pane to return to the topology map. Click on the link to view utilization details which should now be below the threshold.

Step 5: Remove the TTE SR-TE policies upon LCM recommendation

After some time, the deployed TTE policies might no longer be needed. This occurs if the utilization will continue to be under the threshold without the LCM-initiated TTE policies. In this case, LCM generates new recommended actions to delete the TTE policy sets, as shown below. Click **Commit All** to remove the deployed TTE policies.

Local Congestion Mitigation

- Configuration
- Link Management
- LCM Operational Dashboard

Last Refresh: 2020-Dec-19, 11:24:46 (GMT +08:00) |

LCM Operational Dashboard

Congested Interfaces (0) | Mitigating Interfaces (0) | Mitigated Interfaces (3)

Commit All

Node	Interface	Thresho...	Eval...	LCM State	Policies D...	Policy Set...	Reco...	Com...	Expected ...	Actions
PE1-AS...	GigabitEt...	25%	17.43%	✔ Mitigated	2	OK	Delete Set	-	17.43%	...
PE1-AS...	GigabitEt...	25%	15.81%	✔ Mitigated	6	OK	Delete Set	-	15.81%	...

Crosswork Network Automation

- Home
- Change Automation
- Traffic Engineering
- Device Management
- Admin

Last Refresh: 2020-Dec-19, 11:24:46 (GMT +08:00) |

LCM Operational Dashboard

Congested Interfaces (0) | Mitigating Interfaces (0) | Mitigated Interfaces (3)

Commit All

Node	Interface	Thresho...	Eval...	LCM State	Policies D...	Policy Set...	Reco...	Com...	Expected ...	Actions
PE1-AS...	GigabitEt...	25%	17.43%	✔ Mitigated	2	OK	Delete Set	-	17.43%	...
PE1-AS...	GigabitEt...	25%	15.81%	✔ Mitigated	6	OK	Delete Set	-	15.81%	...

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-TE policies and decides whether or not they are still needed. If not, LCM recommends deleting them.

4 Network Maintenance Window

Overview

Objective

Schedule and automate maintenance workflows with minimal network interruption and most efficient results.

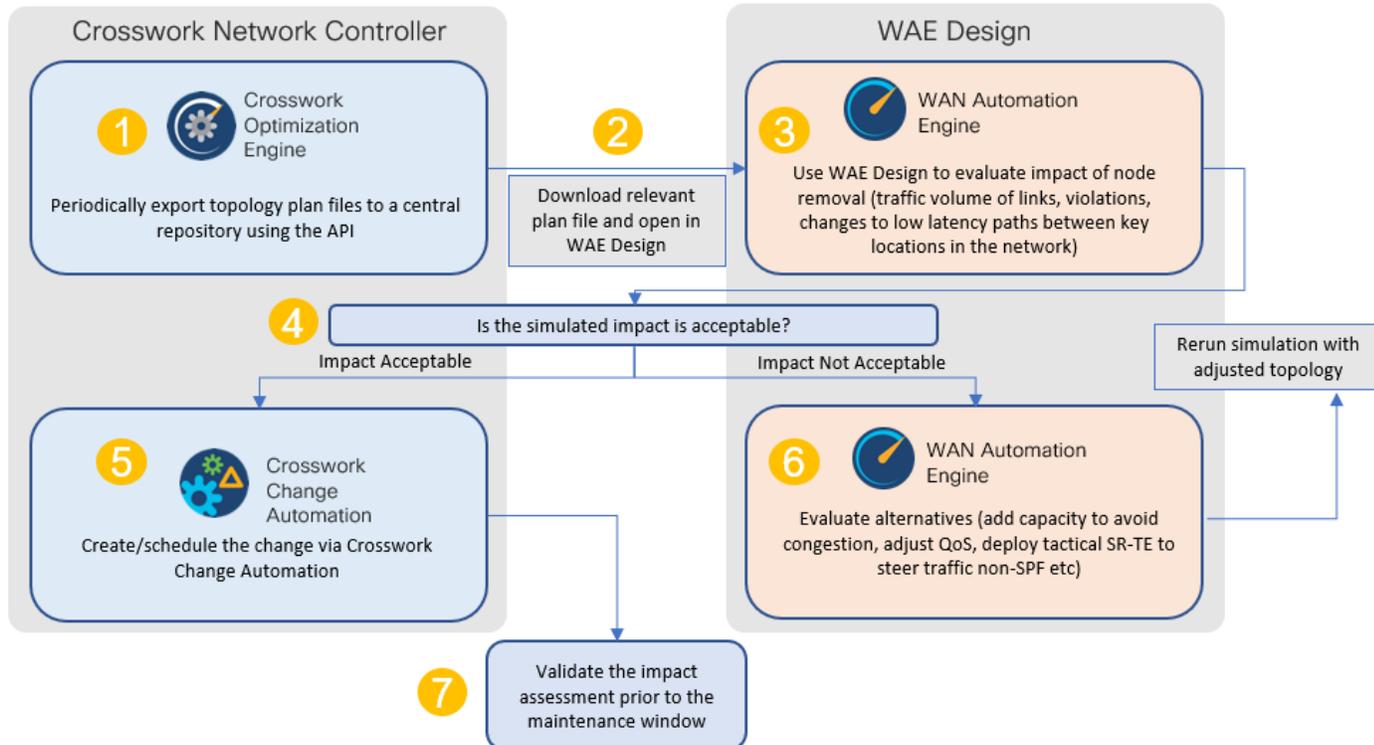
Challenge

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 the 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 once 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.

Usage Scenarios

[Scenario 6: Perform a software upgrade on a provider device during a scheduled maintenance window](#)

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 6: Perform a software upgrade on a provider device during a scheduled maintenance window

Scenario Context

This scenario assumes that Cisco WAE Design has been used to evaluate the impact of removing a P node from the network to perform a software upgrade during a specific timeframe. In this scenario, we will choose a predefined playbook to automate the execution of the SMU on the device, and we will schedule it to run during the predetermined maintenance window.

Assumptions and Prerequisites

- Cisco Crosswork Change Automation must be installed and running.
- You must have access to Cisco WAE Design.
- The Device Override Credentials must be set for Crosswork Network Change Automation to be functional. Go to **Administration > Settings > System Settings > Network Automation**.

Workflow

- [Step 1: Download Topology Plan Files for Impact Analysis](#)
- [Step 2: Schedule and execute the SMU by running a playbook](#)


```

    "version": "7.3.1",
    "format": "txt",
    "
  }
}'

```

- Note the plan file content in the API response. It is encoded for security purposes and must be decoded before you can view the content.

```

{
  "cisco-crosswork-optimization-engine-operations:output": {
    "status": "accepted",
    "plan-file content": "
PE5ldHdvcms+C1Byb3BlcnR5CVZhbHVlC1RpdGx1CQpwZXJzaW9uCTcuMy4xCgo8TmV0d29ya09wdGlvbnM+Ck9wd
GlvbglWYX1ZQpFbmZvcmlQWRqU01ETG9jYXpF0aW9uCVRSVUUKCjx0aXJjdWl0cz4KTmFtZQ1<<<>>Ob2Rl
QQlJbnRlcmZyY2VBCU5vZGVCCU1udGVyZmFjZUIJQ2FwYWNPdHkJRGVsYXkJRGlzdMjTmV0SW50U05NUF9FcnJvcg
l0ZXRJbnRtY29yY2UJUmV0SW50UkUwQ1BVmW0JmV0SW50UkUwQ1BVNWZpZXIjQWxnb3JpdGhtCVJmbGFncU5mbGF
nCVBmbGFncU5mbGFncVZmbGFncU5mbGFncG=="
  }
}

```

- Use a script to decode the plan file or copy the plan file content into a decoder. After decoding the plan file, you can see the content of the model to be used in Cisco WAE Design. It includes a full snapshot of the topology, including the devices, interfaces, links, LSPs, traffic levels, and other information.
- Open the plan file in Cisco WAE Design, simulate the device going down, and observe the impact on the network. Refer to the [Cisco WAE Design documentation](#) for more information.
- Based on the analysis, decide on a good time to execute the SMU.

Step 2: Schedule and execute the SMU by running a playbook

If the simulated impact is acceptable, you can create and schedule the change by running a playbook via Cisco Crosswork Change Automation. For this scenario, we will run a predefined playbook to install a Software Maintenance Update (SMU) on devices tagged under a certain geographic location (NY).

Note: If the predefined (stock) plays and playbooks do not meet your specific needs, you can create custom plays and playbooks. To create a custom play, go to **Network Automation > PlayList** and to create a custom playbook, go to **Network Automation > Playbook List**.

- Go to **Network Automation > Run Playbook**.
- 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.

3. Click **Next** to go to the next task: Select Devices. All devices tagged with City: NY will be selected for SMU installation.
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...

5. Click **Next** to go to the next task: Define Parameters.
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:

- a. Under “verify package consistency on the device” play, set **collection_type** as **mdt**.
- b. Under “perform DLM node lock on device” play, set **retry_count** and **retry_interval** as **3** and **5s** respectively.



✓ Install a SMU or an optional package on a router



✓ Verify package consistency on router ?



collection_type

Data collection type

✓ Perform DLM node lock on device(s) ?



retry_count

Number of time node lock will be retried

retry_interval

Time interval between subsequent retries for node lock. e.g. 10s, 1m, etc. Valid time units are 'ns', 'us' (or 'µs'), 'ms', 's', 'm', 'h'.

- c. Under “Install add package(s)” play, set **action** as **add**, and **optimize** as **false**. Enter the <SMU package name> in **item 1** and set **region** as **NODES**.



Install add package(s) ?



optimize

false

Whether or not to optimize the package list installation. If check mode is set the packages list will be available as facts.

packages ?



item 1

xrv-9k-base-2.0.0.144-r721.CSCuv93809x86_64.rpm

JSON List of SMU package names to be installed on the router, or a tar containing SMU packages

region

NODES

The region in which the host belongs.

- d. Set type as SCP, and enter values for the source, address, destination, and dlm_credential_profile.
- e. Under **Install activate package(s)**, click , select action, and set **action** to **Activate**.
- f. Under **Install commit package(s)**, set **action** to **Commit**.
- g. Under **Verify package in committed list on router**, set **collection_type** to **mdt**, and enter the <SMU package name> in **item 1**.



✓ Install activate package(s) ?



action

Activate

The install action to perform on the router

✓ Install commit package(s) ?



action

Commit

The install action to perform on the router

✓ Verify package in committed list on router ?



collection_type

mdt

Data collection type

✓ packages ?



7. Click **Next** to go to the next task: Define Execution Policy.
8. Select **Continuous** as the Execution mode so that the playbook will run uninterrupted with no pauses. Under Failure policy, select the action you want taken if the execution fails – abort or rollback.
9. Schedule the execution for the optimal time calculated during the impact analysis stage. Uncheck the **Run Now** option. Note the calendar and timer that are displayed to schedule pre-check and perform plays. Select the date and time for the scheduled maintenance.



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 ▼

Schedule

Run Now

Schedule Pre-check (Asia/Jerusalem) ?

Add date

▲ Increment hours ▲ Increment minutes

:

▼ Decrement hours ▼ Decrement minutes

Schedule Perform (Asia/Jerusalem) ?

Add date

▲ Increment hours ▲ Increment minutes

:

▼ Decrement hours ▼ Decrement minutes

All Scheduled Jobs

Show jobs for selected devices only

Previous	Today		April 2021				Month	Week
Next								Day
Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday		
28	29	30	31	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	1		

10. Click **Next** to go to the next task: Confirm Job.

- Review your job details. Label your job with a unique name. 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 [Change](#)
 Continuous (0)
 Pre Maintenance (1)
 Maintenance (4)
 Post Maintenance (1)

Tag NY [Change](#)

Mop Params

```
{
  "1": {
    "collection_type": "mdt"
  },
  "2": {
    "retry_count": "3",
    "retry_interval": "5s"
  },
  "3": {
    "optimize": false,
    "packages": [
      "xrv-9k-base-2.0.0.144-r721.CSCuv93809x86_64.rpm"
    ],
    "region": "NODES",
    "repository": {
      "type": "SCP",
      "source": "/root/smus",
      "address": "192.168.6.1",
      "destination": "harddisk:",
      "dml_credential_profile": "abc"
    }
  }
}
```

Label your Job

Name *

Labels

Step 3: Verify the SMU install job completion status

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.

Job Sets | 1 / 43

Job Set: smu_xrv-77993990ce

Status: Success

Job Set Tags

PlayBook Title: router_op_smu_upgrade

Created By: admin

All Jobs in the Set (1) Selected 1 / Total 1

Status	Device	Execution ID	Start Time	End Time
Succeeded	xrv9k-1	1613667141147-5b7e0cec-7c19-4368-b540-177d470add02	Thu, Feb 18, 2021, 08:55:5...	Thu, Feb 18, 2021, 09:20:0...

2. Select the SMU install job. Note the Job Set details on the right side. Click the **Execution ID** for job details.

Change Automation / Job History / Job Set: smu_xrv-77993990ce / 1613667141147-5b7e0cec-7c19-4368-b540-177d470add02

Playbook: Install a SMU or an optional package on a router

Device: xrv9k-1

SUCCEEDED 2021-Feb-18, 09:20:04 (GMT -08:00)

Parameters View

Execution Mode

Step	Description	Status
Pre Maintenance 1/1	1 Verify package consistency on router	Success
Maintenance 4/4	2 Perform DLM node lock on device(s)	Success
	3 Install add package(s)	Success
	4 Install activate package(s)	Success
	5 Install commit package(s)	Success
Post Maintenance 1/1	6 Verify package in committed list on router	Success

Events Syslog Console

GENERIC EVENT
2021-Feb-18, 09:20:04 (GMT -08:00) - Node Name : ["xrv9k-1"] - Event : {"description":"MoP job completed","status":"COMPLETED"}

MOP STATUS
2021-Feb-18, 09:20:04 (GMT -08:00) Status: SUCCEEDED - Description: maintenance phase succeeded

MOP TASK EVENT
2021-Feb-18, 09:20:04 (GMT -08:00) - Node Name : ["xrv9k-1"] - Task : Verify package in committed list on router - Result: SUCCESS - Description: Input package(s) given are present in committed package(s)

GENERIC EVENT
2021-Feb-18, 09:20:04 (GMT -08:00) - Node Name : ["xrv9k-1"] - Event : Input package(s) given are present in committed package(s)

NODE STATUS UPDATE
2021-Feb-18, 09:20:04 (GMT -08:00) - Node Name : ["xrv9k-1"] - Status : READY

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 and checking that the SMU you installed appears in the list. Some example outputs from these commands are shown below:

```
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.

5 Programmable Closed-Loop Auto-Remediation

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 automatically be performed once 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

- Zero-touch telemetry streamlines the operational and network management overhead of collecting and cleansing data, thereby allowing operators to focus on their business goals. As part of zero-touch telemetry, devices are automatically provisioned with telemetry configuration and tables/schema are created in a Time Series Database (TSDB).
- 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 going 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 (easily scripted with a domain specific language).
- TCAs can be exported/integrated with other systems via HTTP, Slack and socket interfaces.
- KPIs can be associated with dashboards, which provide real-time and historical views of the raw 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 (CCCA) playbooks (or webhooks), 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/anomalies on the topology of the network, allowing easy visualization of the impact of events.

Scenario 7 – 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, watch dog, 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

Following is a high-level workflow for executing this scenario:

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](#) in [Scenario 1](#) - Implement and Maintain SLA for an L3VPN Service (using ODN).

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](#) in [Scenario 1](#) - Implement and Maintain SLA for an L3VPN Service (using ODN).

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 4: Create and provision the L3VPN service](#) in [Scenario 1](#) - Implement and Maintain SLA for an L3VPN Service (using ODN).

4. Define KPIs and 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](#).

5. When there is a KPI failure, 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 KPI failure 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](#).

6. Cisco Crosswork Network Controller continues to monitor the link and when there are no longer KPI failures, the RED affinity tag can be removed. Define another playbook for this purpose.

6 Automation of Onboarding and Provisioning of IOS-XR Devices Using ZTP

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 applied to the device automatically. Once 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. Once 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. Once 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. Once the image (if any) is downloaded and installed, and the device reloads with the new image, the device downloads configuration scripts and executes them.

Usage Scenarios

[Scenario 8 - Automatically onboard and provision new devices in the network](#)

Additional Resources

Detailed information is available in the [ZTP chapters](#) in the [Cisco Crosswork Infrastructure 4.0 and Applications Administration Guide](#).

Scenario 8 – Automatically onboard and provision new devices in the network

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 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.
- The Crosswork server must be reachable from the devices, via an out-of-band management network or an in-band data network.
- 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*.

Workflow

This is a high-level workflow for onboarding devices using Cisco Crosswork ZTP. Detailed information is available in the [ZTP chapters](#) in the [Cisco Crosswork Infrastructure 4.0 and Applications Administration Guide](#).

- [Step 1: Assemble and upload ZTP assets](#)
- [Step 2: Create a ZTP profile combining an image file and configuration file](#)
- [Step 3: Prepare ZTP device entries for the devices to be onboarded](#)
- [Step 4: Set up DHCP for Crosswork ZTP](#)
- [Step 5: Initiate ZTP processing to onboard the devices](#)
- [Step 6: Monitor the ZTP processing status](#)
- [Step 7: Verify that the onboarded devices appear in the inventory of managed devices](#)
- [Step 8: Add/edit device information if necessary](#)

Step 1: Assemble and upload ZTP assets

1. Assemble the following assets before you begin:
 - Software images
 - Configuration Files
 - Credentials of the devices to be onboarded
 - Serial numbers of the devices to be onboarded
 - Owner certificates (for secure ZTP) - your organization's CA-signed end-entity certificates, installed on your devices and binding a public key to your organization.
 - Pinned domain certificate (for secure ZTP) - 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 (for secure ZTP) - Nonceless audit vouchers that verify that devices being onboarded with ZTP are bootstrapping into a domain owned by your organization. Cisco supplies OVs when a request is submitted with your organization's PDC and device serial numbers.
2. Upload the software images. Go to **Device Management > Software Images**.
3. Upload the configuration files. Go to **Device Management > Configuration Files**.
4. Upload device serial numbers. Go to **Device Management > Voucher Management** and click **Add Serial Number**.
5. For secure ZTP, upload your pinned domain certificate and owner certificates. Go to **Administration > Certificate Management** and add your certificates.
6. For secure ZTP, upload ownership vouchers. Go to **Device Manager > Voucher Management**.

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.

1. Go to **Device Management > Devices**.
2. Click the **Zero Touch Devices** tab.

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.

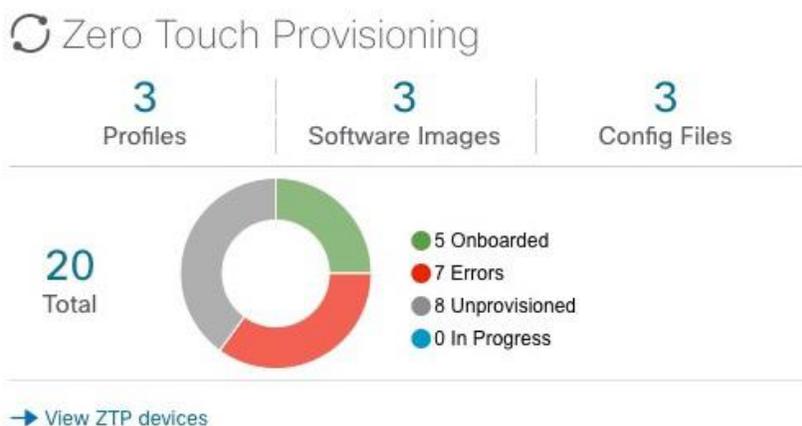
Step 5: Initiate ZTP processing to onboard the devices

Initiate ZTP processing by rebooting each of the devices to be provisioned.

Step 6: Monitor the ZTP processing status

You can monitor the progress of the ZTP processing in the dashboard.

1. Click **Home** in the main menu and take a look at the Zero Touch Provisioning dashlet.



2. Click on the **View ZTP devices** link to view the status of individual devices.

Step 7: Verify that the onboarded devices appear in the inventory of managed devices

Step 8: Add/edit device information if necessary

Some of the information needed for a complete device inventory record is either not needed in order to onboard the device, or not directly available via automation. For example, geographical location data defined using a set of GPS coordinates.

ZTP devices, once onboarded, are automatically part of the shared Crosswork device inventory, and can be edited like any other device.