



Cisco Crosswork Network Controller 3.0 Solution Workflow Guide

Contents

Solution Overview	4
Description	4
Supported Use Cases	4
Solution Components and Integrated Architecture	5
Supported Device Software	9
Multi-Vendor Capabilities	10
Extensibility	10
UI Overview	10
Orchestrated Service Provisioning	14
Overview	14
Scenario 1 – Implement and Maintain SLA for an L3VPN Service for SR-MPLS (using ODN)	16
Scenario 2 – Implement and Maintain SLA for an L3VPN Service for SRv6 (using ODN)	32
Scenario 3 – Mandate a Static Path for an EVPN-VPWS Service using an Explicit SR-TE Policy	43
Scenario 4 – Provision an L2VPN service over an RSVP-TE tunnel with reserved bandwidth.....	58
Scenario 5 – Provision a Soft Bandwidth Guarantee with Optimization Constraints.....	65
Bandwidth and Network Optimization.....	70
Overview	70
Scenario 6 – Use Local Congestion Mitigation (LCM) to reroute traffic on an over-utilized link.....	72
Network Maintenance Window	78
Overview	78
Scenario 7 – Perform a software upgrade on a provider device during a scheduled maintenance window.....	80
Programmable Closed-Loop Remediation	89
Overview	89
Scenario 8 – Achieve Predictive Traffic Load Balancing Using Segment Routing Affinity	90
Automation of Onboarding and Provisioning of IOS-XR Devices Using ZTP	91
Overview	91
Scenario 9 – Automatically onboard and provision new devices in the network	92
Visualization of native SR path	95
Overview	95
Scenario 10 – Troubleshooting paths between native SR paths over inter-AS Option C.....	95
Appendix	100
Initializing Heuristic Packages to monitor the health of a service	100

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 the ability to react to bandwidth and latency demand fluctuations in real time is vital to success. Migration to Software-Defined Networks (SDNs) and automation of operational tasks is the optimal way 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 costs. 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 are nearly impossible to replicate even with a highly skilled and dedicated staff operating the network.

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


Supported Use Cases

- **Orchestrated service provisioning:** Provisioning of layer 2 VPN (L2VPN) and layer 3 VPN (L3VPN) services with underlay transport policies to define, meet, and maintain SLAs, using the UI or APIs. Utilizing Flexible Algorithm (FlexAlgo) provisioning and visualizing to customize and compute IGP shortest path over a network according to specified constraints.
- **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.
- **Local Congestion Management:** Local Congestion Mitigation (LCM) provides localized mitigation recommendations within surrounding interfaces, with the use of standard protocols. Data is gathered in real-time and when congestion is detected, solutions are suggested. LCM has a “human-in-the-loop” aspect which ensures that 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 or 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 of pre-defined remediation tasks when a KPI threshold is breached. For this use case, Cisco Crosswork Health Insights and Cisco Crosswork Change Automation must be installed.
- **Planning, scheduling, and automating network maintenance tasks:** Scheduling an appropriate maintenance window for a maintenance task after evaluating the potential impact of the task (using WAE Design). Automating the execution of maintenance tasks (such as throughput checks, software upgrades, SMU installs) using playbooks. For this use case, Cisco Crosswork Health Insights and Change Automation must be installed.

- **Secured zero-touch onboarding and provisioning of devices:** Onboarding new IOS-XR devices and automatically provisioning Day0 configuration resulting in faster deployment of new hardware at lower operating costs. For this use case, Cisco Crosswork Zero Touch Provisioning must be installed.
- **Visualization of native SR paths:** Visualizing the native path using the traceroute SR-MPLS multipath command to get the actual paths between the source and the destination can be achieved using Path Query. With Cisco Crosswork Network Controller, a traceroute command runs on the source device for the destination TE-Router ID and assists in retrieving the paths.

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.



The following components make up the Cisco Crosswork Network Controller 3.0 solution:

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

- [Cisco Service Health](#)
- [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 Service Health](#)

Cisco Crosswork Active Topology

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

Cisco Crosswork Optimization Engine

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

Cisco 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 support multiple data collection protocols including MDT, SNMP, CLI, standards-based gNMI (dial-in), and syslog. Any type of data can be collected by Crosswork Data Gateway as long as it can be delivered over one of the supported protocols. In this way, it can provide support for a growing set of use cases and customizations.

To address scale challenges, Cisco Crosswork Data Gateway is implemented as a number of VMs and designed with a distributed architecture in mind. Each lightweight VM manages a subset of the overall network and as the network grows, additional VMs can be added horizontally to address the new demands on the compute resources. It also supports a flexible redundancy configuration based on the operator's needs. After the initial setup, Cisco Crosswork Network Controller automatically orchestrates the collection across the multiple Cisco Crosswork Data Gateway VMs. APIs and configuration examples are available to illustrate how to add new collection jobs (outside of those built for you by Cisco Crosswork Network Controller) to gather additional information from your network. The collected data can be published to approved destinations. Supported destinations are Kafka and gRPC messaging bus.

Crosswork Common UI and API

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

A common API enables Crosswork Network Controller's programmability. The common APIs provides a single access point for all APIs exposed by various built-in components. The API provides a REST-based Northbound Interface to external

systems (e.g., OSS systems) to integrate with Cisco Crosswork Network Controller. RESTCONF and YANG data models are made available for optimization use cases. For details about the APIs and examples of their usage, see the [Cisco Crosswork Network Automation API Documentation on Cisco DevNet](#).

Crosswork Infrastructure and Shared Services

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

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

Cisco Crosswork Health Insights and Cisco Crosswork Change Automation

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

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

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

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

Cisco Crosswork Zero-Touch Provisioning (ZTP)

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

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

Cisco Crosswork ZTP offers Secure ZTP functionality in addition to the Classic ZTP functionality. Secure ZTP is based on RFC 8572 standards and uses secure transport protocols and certificates to verify devices and perform downloads. Secure ZTP is useful when public Internet resources must be traversed to reach remote network devices, or when the devices are from third-party manufacturers. With Secure ZTP, the device and the Cisco Crosswork ZTP bootstrap server authenticate each other using the device's Secure Unique Device Identifier (SUDI) and Crosswork server certificates over TLS/HTTPS. After a secure HTTPS channel is established, the Crosswork bootstrap server allows the device to request to download and apply a

set of signed image and configuration artifacts adhering to the RFC 8572 YANG schema. After the image (if any) is downloaded and installed, and the device reloads with the new image, the device downloads configuration scripts and executes them.

Cisco Network Services Orchestrator (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 (European Telecommunications Standards Institute) architecture. It provides complete support for physical and virtual network elements, with a consistent operational model across both. It can orchestrate across multi-vendor environments and support multiple technology stacks, enabling the extension of end-to-end automation to virtually any use case or device.

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

Service and device models, written using YANG modelling language, enable Cisco NSO to efficiently ‘map’ service intent to device capabilities and automatically generate the minimum required configuration to be deployed in the network. This feature, facilitated by Cisco NSO’s FASTMAP algorithm, is capable of comparing current configuration states with a service’s intent and then generating the minimum set of changes required to instantiate the service in the network.

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

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

- SR-TE core function pack (CFP) enables provisioning of explicit and dynamic segment routing policies, including SRv6, and on-demand SR-TE policy instantiation for prefixes with a specific color.
- Sample function packs for IETF-compliant L2VPN and L3VPN provisioning. These function packs provide baseline L2VPN and L3VPN provisioning capabilities, based on IETF NM models. Prior to customization, these sample function packs enable provisioning of the following VPN services:

Note: The Service Health function pack should be independently installed apart from Cisco Crosswork Network Controller function packs.

- L2VPN:
 - Point-to-point VPWS using Targeted LDP
 - Point-to-point VPWS using EVPN
- L3VPN
- Sample IETF-compliant RSVP-TE function pack intended as a reference implementation for RSVP-TE tunnel provisioning, to be customized as required.

Note: By default, the IETF-compliant NM models are used. If your organization wishes to continue to use the Flat models that were provided with the previous version, a manual setup process is required. Consult your Cisco Customer Experience representative for more information.

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.

Cisco Service Health

Note: Service Health is not generally available yet. At this stage, it is available for pre-launch laboratory evaluation only. Engage your account team if you are interested in participating in the evaluation.

Service Health substantially reduces the time required to detect and troubleshoot service quality issues. It monitors the health status of provisioned L2/L3 VPN services and enables operators to pinpoint why and where a service is degraded. It can also provide service-specific monitoring, troubleshooting, assurance, and proactive causality through a heuristic model that visualizes the:

- Health status of sub-services (device, tunnel) to a map when a single service is selected
- Service logical dependency tree and help the operator in troubleshooting in case of degradation by locating where the problem resides, an indication of possible symptoms, and impacting metrics in case of degradation
- Historical view of service health status up to 60 days.

Supported Device Software

OS	Version	SR-PCE	PCE-Init	PCC-Init	NSO+CFP		Crosswork Infrastructure 4.1	Crosswork Optimization Engine	ZTP	Service Health
					CLI	NETCONF				
IOS-XR	6.5.3				yes		yes	yes		
	6.6.3		yes	yes	yes		yes	yes		
	7.1.2		yes	yes	yes		yes	yes		
	7.2.1		yes	yes	yes		yes	yes		
	7.3.1		yes (Cisco ASR9000 Series only)	yes	yes	yes	yes		yes	yes
	7.3.2	yes	yes	yes		yes	yes	yes	yes	yes

	7.4.1		yes	yes		yes	yes	yes	yes	yes
IOS-XE	17.4.1			yes	yes		yes	yes	yes	
	17.5.1			yes	yes		yes	yes	yes	
	17.6.1			yes	yes		yes	yes	yes	

Note: IOS XR versions below 7.x do not have gNMI support.

Note: PCE-Init is not supported for SRv6.

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 has multi-vendor capabilities 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 or 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; 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**.

Self-signed certificate: At first-time access, some browsers display a warning that the site is untrusted. When this happens, follow the prompts to add a security exception and download the self-signed certificate from the server. After you download the certificate, the browser accepts the server as a trusted site in all future login attempts.

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

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 your dashboard depend on which Cisco Crosswork applications you have installed. Links in each dashlet allow you to drill down for more details.



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

Navigation

The main menu along the left side of the window provides access to all features and functionality in Cisco Crosswork Network Controller, as well as to device management and administrative tasks. The 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](#).

Topology

The network topology can be displayed on a logical map or a geographical (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. 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. Clicking on the question mark in the map provides a detailed legend of the various symbols and their meaning.

Host Name	IP Address	Reachability	Device Type	Product
P-BOTTOMLEFT	192.168....	✓ Re...	Router	ciscoNCS...
P-BOTTOMRIGHT	192.168....	✓ Re...	Router	ciscoNCS...
P-TOPLEFT	192.168....	✓ Re...	Router	ciscoNCS...
P-TOPRIGHT	192.168....	✓ Re...	Router	ciscoNCS...
PCE-RR		?	Router	ciscoNCS...
PE-A	192.168....	✓ Re...	Router	CISCO-X...
PE-B	192.168....	✓ Re...	Router	CISCO-X...
PE-C	192.168....	✓ Re...	Router	CISCO-X...

Services & Traffic Engineering



The Services & Traffic Engineering menu provides access to VPN and transport provisioning and visualization functionality, bandwidth management functionality, as well as access to the configuration pages used to enable Feature Packs. For more information, see the [Cisco Crosswork Optimization Engine 3.0 User Guide](#).

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


Choose **Provisioning (NSO)** to access the provisioning UI rendered from the Cisco 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 are typically performed by an administrator.

Refer to the [Cisco Crosswork Infrastructure 4.1 and Applications Administration Guide](#) for information about these functions.

Orchestrated Service Provisioning

Overview

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

Objective

Provision VPN services with underlay transport policies 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 monitoring and intervention, which is time- and resource-intensive.

Solution

With network automation, the objective is to automate the feedback loop to enable quicker reaction to and remediation of network events. With 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) polices can be configured to continuously track network changes and automatically react to optimize the network. These SR-TE polices can serve as the underlay configuration for the VPN services to automatically maintain the SLAs.

The services required for this solution can be created and managed using the 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. Optionally, Service Health monitoring can be enabled to see which services are working as provisioned, if issues have been flagged, and what symptoms are detailed so to quickly address and fix.

Note: The 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 that is 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 specifies bandwidth, the SR-PCE gets the path from Cisco Crosswork Optimization Engine.
5. The SR-PCE sends the path to the headend device via PCEP and updates the headend if path changes are required.

Usage Scenarios

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

- [Scenario 1 – Implement and Maintain SLA for an L3VPN Service for SR-MPLS \(using ODN\)](#)
- [Scenario 2 – Implement and Maintain SLA for an L3VPN Service for SRv6 \(using ODN\)](#)
- [Scenario 3 – Mandate a Static Path for an EVPN-VPWS Service using an Explicit SR-TE policy](#)

- [Scenario 4 – Provision an L2VPN service over an RSVP-TE tunnel with reserved bandwidth](#)
- [Scenario 5 – 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 included in the Cisco NSO image here:
<https://software.cisco.com/download/home/286323467/type/286283941/release/5.5.2>

Scenario 1 – Implement and Maintain SLA for an L3VPN Service for SR-MPLS (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, the 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 to the same 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, at 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.

You also have the option within the workflow to enable Service Health monitoring and visualize Flexible Algorithm (Flex Algo) as a constraint. With Service Health, you may monitor a service's health status and use the insights provided into degraded and down services to visualize, inspect, and troubleshoot for improved network optimization.

Note: Service Health is not generally available yet. At this stage, it is available for pre-launch laboratory evaluation only. Engage your account team if you are interested in participating in the evaluation.

With Flex Algo, you may customize an IGP shortest path computation according to user-defined algorithms where the IGP computes paths on a user-defined combination of metric type and constraint, as well as view a filtered topology based on your specific Flex Algo definitions.

The following topology provides the base for this scenario:



In this scenario, we will:

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

Assumptions and Prerequisites

- To use ODN, BGP peering for the prefixes must be configured between the endpoints or PEs. Usually for L3VPN, this is the VPNV4 and VPNV6 address family peering.
- For Service Health enablement, Service Health must be installed.

Note: Service Health is not generally available yet. At this stage, it is available for pre-launch laboratory evaluation only. Engage your account team if you are interested in participating in the evaluation.

- Before using Service Health's Assurance Graph, ensure that topology map nodes have been fully configured and created with a profile associated to the service. If not, Subservice Details metrics will show that no value has yet to be reported.
- (Optional) Flexible Algorithms, and the IDs that are used, must be configured in your network.

Workflow

- [Step 1. Create an ODN template to map color to an SLA objective and constraints](#)
- [Step 2. Create an L3VPN Route Policy](#)
- [Step 3. Create and provision the L3VPN service](#)
- [Step 4. Enable Service Health monitoring](#)

- [Step 5. Visualize the New VPN Service on the Map to See the Traffic Path](#)
- [Step 6. Observe automatic network optimization](#)
- [Step 7. Inspect a degraded service using Service Health to determine active symptoms](#)

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

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

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

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

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 (NSO) > 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**. Click **Continue**.
3. Choose the head-end 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).
6. Define the required constraints. In this case, we want the computed paths to be disjoint in that they must not share a link.
Under disjoint-path, choose **link** as the type, and specify a numeric group ID, in this case, **16**, as the group-id.

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

Note: Optionally, you may configure Flexible Algorithm (flex-algo) as a constraint.

L3VPN_NM-SRTE-ODN_72-a

head-end

[+](#) [Edit](#) [Delete](#)

name	<input type="text"/>
PE-A	

maximum-sid-depth [?](#)

color * [?](#)

bandwidth [?](#)

dynamic

Enable dynamic

metric-type [?](#)

pce [?](#)

flex-alg [?](#)

> metric-margin

disjoint-path

Enable disjoint-path

type * [?](#)

group-id * [?](#)

7. Commit your changes or click **Dry Run** to check what will be configured on the devices before you commit.
8. Check that the new ODN template appears in the table and its provisioning state is **Success**. Click ... in the Actions column and choose **Config View** to see the Yang model-based service intent data that details the ODN template you created.

ODN Template

Total 5 | Last Refresh: 12-Oct-2021 04:10:25 PM PDT | ⚙️

Name	Provisioning State	Date Created	Action
L3VPN_NM-SRTE-ODN_70	✓ Success	12-Oct-2021 03:59:31 PM PDT	...
L3VPN_NM-SRTE-ODN_71-a	✓ Success	12-Oct-2021 03:57:33 PM PDT	...
L3VPN_NM-SRTE-ODN_71-c	✓ Success	12-Oct-2021 04:06:27 PM PDT	...
L3VPN_NM-SRTE-ODN_72-a	✓ Success	12-Oct-2021 03:53:41 PM PDT	...
L3VPN_NM-SRTE-ODN_72-b	✓ Success	12-Oct-2021 04:04:20 PM PDT	...

Manage
Config View Config View
Edit
Delete

Cross Launch
View In NSO
View Plan Data

Service Options
Check-Sync
Sync-From
Sync-To
Re-Deploy Dry Run
Re-Deploy
Re-Deploy Reconcile
Reactive-Re-deploy
Clean-Up

Configured Data

```

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

```

[Copy To Clipboard](#)[Cancel](#)

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

Step 2. Create an L3VPN Route Policy

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

We will create the following route policies:

- Color 70, 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, we will show how to create the first route policy - L3VPN_NM-SRTE-RP-PE-A-7. The other route policies can be created using the same procedure.

Procedure

1. Go to **Services & Traffic Engineering > Provisioning (NSO) > 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. For instance, expand ipv4 and click **+** to create a new prefix. Enter IPv4 prefix **70.70.70.0/30** and click **Continue**.

The screenshot shows the 'L3vpn Route Policy' configuration page. On the left, there's a table for 'color' with one row selected (id: 70). On the right, a detailed view pane shows the configuration for color 70. It includes fields for 'id' (70), 'exclusive' (false), 'description' (false), and an 'ipv4' section where 'prefix' is set to 70.70.70.0/30. A large 'X' button is visible in the top-right corner of the pane.

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.

Note: After creating an L3VPN Route Policy, the VPN profile for each route policy will be automatically created. The VPN profile will be referenced from the L3VPN service. This will bind the route policy to the L3VPN service. Thus, VPN profiles for each of the route policies we created in the previous step will be created:

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

Step 3. Create and provision the L3VPN service

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

1. Go to **Services & Traffic Engineering > Provisioning (NSO) > 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 shows two windows side-by-side. The left window is titled 'L3VPN_NM-SRTE-ODN-70' and contains a table for 'custom-template' with one row labeled 'name'. Below it, under 'ie-profiles *', there is a table for 'ie-profile' with three rows: L3VPN_NM_SR_ODN-IE-PE-A-7 (selected), L3VPN_NM_SR_ODN-IE-PE-B-7, and L3VPN_NM_SR_ODN-IE-PE-C-7. The right window is titled 'ie-profile[L3VPN_NM_SR_ODN-IE-PE-A-7]' and shows an 'ie-profile-id' field with 'L3VPN_NM_SR_ODN-IE-PE-A-7' and an 'rd' field with '0:70:70'. Under 'vpn-targets *', there is a table for 'vpn-target' with one row: 'id' 100 and 'route-target-type' both. Under 'vpn-policies *', there is a table for 'import-policy' with a dropdown menu and an 'export-policy' field containing '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 **Dry Run** 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 4. Enable Service Health monitoring

Note: Service Health is not generally available yet. At this stage, it is available for pre-launch laboratory evaluation only. Engage your account team if you are interested in participating in the evaluation.

1. Go to **Services & Traffic Engineering > VPN Services**. The map opens and a table of VPN Services is displayed to the right of the map.
2. In the Actions column, click ... for the new service you want to start monitoring health.
3. Select **Start Monitoring**.

Health	Service Key	Type	Provisioning State	Last Updated Ti...	Actions
✓	L2NM-EVPN-SRTE-105	L2vpn-Se...	✓ Success	27-Jul-2021 05:...	...
✗	L2VPN_NM_P2P_SRTE-...	L2vpn-Se...	✓ Success	07-Sep-2021 1...	...
✗	L3VPN_NM-SRTE-70	L3vpn-Se...	✓ Success	02-Aug-2021 0...	...
⌚	I2nm-evpn	L2vpn-Se...	✓ Success	02-Aug-2021 0...	...
⌚	I2nm-evpn-01_sr	L2vpn-Se...	✓ Success	02-Aug-2021 0...	...


Note: The Health column color coding indicates the health of the service: Green = Good; Orange = Degraded; Red = Down, Gray = Not Monitoring.

4. The Monitor Service pop-up appears. Select which service to monitor and again click **Start Monitoring**.

The screenshot shows the 'VPN Services' dashboard with a modal window titled 'Monitor Service'. The modal contains configuration for the 'Gold_L3VPN_ConfigProfile system', including CPU threshold (70.5%), memory threshold (10), tunnel throughput (50 mbps), and other parameters. Below the modal is a table of 7 VPN services, each with a status column indicating success or failure.

Service Key	Type	Provisioning State	Last Upda...	A...
i2vpn-115sr-119igp...	L2vpn-Service	Success	24-Oct-2021...	...
i2vpn-evpn-srte-255	L2vpn-Service	Success	25-Oct-2021...	...
i2vpn-evpn-srte_no...	L2vpn-Service	Success	24-Oct-2021...	...
i2vpn-p2p-250	L2vpn-Service	Success	24-Oct-2021...	...
i3vpn-260	L3vpn-Service	Success	25-Oct-2021...	...
i3vpn-270	L3vpn-Service	Success	24-Oct-2021...	...
i3vpn-srte-odn-70	L3vpn-Service	Success	24-Oct-2021...	...

Note: Now that you have started monitoring the health of this service, in the Actions column, click ... to view additional Service Health options: **Stop Monitoring**, **Pause Monitoring**, **Assurance Graph**.



5. Repeat this step for each service you wish to start health monitoring.
6. Click X in the top-right corner when you are done.

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 VPN overlay in the logical map. Use the buttons at the top right of the map to toggle between the logical and geographical maps.



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

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

Health	Service Key	Type	Provisioning State	Last Upda...	A...
	odn				
✓	I3vpn-srte-odn-70	L3vpn-Service	✓ Success		

- View Details
- Edit/Delete
- Stop Monitoring
- Pause Monitoring
- Assurance Graph

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

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

The screenshot shows the Cisco Network Services & Traffic Engineering interface. On the left, a map displays network nodes (A-PE5, A-PE6, A-PE3, PA1, PA4, AG3, AG4, AG2, AG1) and links. A purple path is highlighted between A-PE5 and A-PE6. On the right, the 'Service Details' panel is open for a service named 'i3vpn-site-odn-70'. It shows a table of SR policies with columns for Health, Headend, Endpoint, Color, Admin Status, Oper St..., and Actions. Two policies are selected: one from A-PE6 to A-PE5 (Color 71) and another from A-PE6 to A-PE3 (Color 72). The 'Transport' tab is selected.

- 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.

This screenshot shows the same interface after selecting the 'Show IGP Path' option. The highlighted path between A-PE5 and A-PE6 is now shown in red. A tooltip on the link between A-PE6 and PA4 displays prefix SID 17112 and route information: 'bgp_c_71_ep_192.168.0.17_discr_100'. The 'Service Details' panel shows the same service configuration and selected policies, with the 'Transport' tab still active.

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

- Click the **Flex Algo** tab.
- From the drop-down list, choose up to 2 Flexible Algorithm IDs.
- View the Flexible Algorithm Types and confirm that the selection is correct. Also, note the color assignments for each Flexible Algorithm.
- (Optional) Check the **Show selected Flex Algo topology only** check box to isolate the Flexible Algorithms on the topology map. When this option is enabled, SR policy selection is disabled.
- Check the **Show Flex Algo A+B links only** to show only those links and nodes that participate in both Flexible Algorithms.
- Click **Apply**. You must click **Apply** for any additional changes to Flexible Algorithm selections to see the update on the topology map.

Note: If a selected Flexible Algorithm is defined with criteria but there are no links and node combinations that match it (for example, a defined affinity to include all nodes or links with the color blue), then the topology map will be blank. If a selected Flexible Algorithm is not configured on a node or link, then the default blue link or node color appears.


- g. (Optional) Click **Save View** to save the topology view and Flexible Algorithm selections.

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 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, to compensate for the link that is down and to maintain the disjoint paths.

Recomputed paths:

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



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

Note: Service Health is not generally available yet. At this stage, it is available for pre-launch laboratory evaluation only. Engage your account team if you are interested in participating in the evaluation.

In this step, we will monitor the VPN services using Assurance Graph capabilities and inspect any services or related nodes that are degraded. By inspecting the root causes that lead to reported active symptoms and impacted services, you can determine what issues must be addressed first to maintain a healthy setup and what requires further inspection and troubleshooting.

1. Click **X** in the top-right corner to return to the VPN Services list.
 2. Click on the name of a service that shows as being degraded. The map will update to highlight the service you selected.
- Degraded services show an orange icon in the Health column. You can filter by health state (Down, Degraded,

Good) by clicking in the space at the top of the column and selecting the appropriate filter. To clear the filter, click the X next to the designated filter appearing in the space at the top of the column and it will remove all filtering and default to showing all VPN Services in the list.

Note: If a service is not yet being monitored, the icon in the Health column will show as the color grey. To enable monitoring for such a service, click ... and select **Start Monitoring**.

3. In the Actions column, click ... and click **View Details**. The Service Details screen appears on the right side.
4. With the Health tab selected, review Active Symptoms for the degraded service (including the Root Cause, Subservice, Priority, and Last Updated details) present in the Health tab if the service is being currently monitored.


The screenshot shows the 'Service Details' page for a service named 'I3vpn-srte-odn-70'. The service is listed as 'Provisioning Success' and 'Health Degraded' with a 'Monitoring Profile' of 'Gold_L3VPN_ConfigProfile sys...'. Below this, there are tabs for 'Health' (which is selected), 'Transport', and 'Configuration', along with a 'Path Query' button. The main content area is titled 'Active Symptoms (4)' and displays a table with four rows of data. The columns are 'Root Cause', 'Subservice', 'Prior...', and 'Last Updated'. The data is as follows:

Root Cause	Subservice	Prior...	Last Updated
VRF Route to peer-vpn-ip-address: 70.70.7...	subservice.vrf.plai...	255	26-Oct-2021...
VPN Interface GigabitEthernet0/0/0/6.70 Op...	subservice.interfa...	255	26-Oct-2021...
No connected routes reported for VRF: I3vp...	subservice.ce.pe....	255	26-Oct-2021...
VRF Route to peer-vpn-ip-address: 70.70.7...	subservice.vrf.plai...	255	26-Oct-2021...

5. Click on a Root Cause and view both the Symptom Details and the Failed Subexpressions & Metrics information.

Name	Expression Value
interface_oper == 'up'	false
explabel	oper_up
symptomMetrics	
metric.interface.oper system(device=A-PE6, gigEthIfId=GigabitEther...	down
subExps	
observedValue	down
symptomMetrics	
metric.interface.oper system(device=A-PE6, gigEthIfId=GigabitEthe...	down

6. Select the Transport and Configuration tabs and review the details provided.
7. To further isolate the degraded service details, click X in the top-right corner to return to the VPN Services list.
8. Again, click on the name of the degraded service in the list. The Service Details panel appears and the map updates, isolating the corresponding devices participating with that service.
9. Within the map, view further service health details doing the following:
 - At the top-left of the map, select the Show Participating Only check box so the map only shows the participating services.
 - In the map, hover your mouse over one of the devices and smaller badges that indicate health status and review the pop-up information relating to its Reachability State, Host Name, Node IP, and Type.
10. In the Actions column, click ... for the degraded service in the list and click **Assurance Graph**. The topology map of services and subservices appear with the Service Details panel showing Service Key, Status, Sub Services, and Active Symptoms details.




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

- In the topology map, select a degraded subservice. The Subservice Details panel appears with subservice metrics, as well as subservice specific Active Symptoms and Impacted Services details.

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

Note: Use your mouse to on subservices in the map for details on the degraded health. At the top left of the map, select Down & Degraded Only or Soft Dependencies to further isolate subservices.



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

Summary and Conclusion

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

Scenario 2 – Implement and Maintain SLA for an L3VPN Service for SRv6 (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, the lowest latency path is the SLA objective. The customer requires a low latency path for high priority traffic. The customer wants to use disjoint paths, i.e., two unique paths that steer traffic from the same source and to the same destination, avoiding common links so that there is no single point of failure. The customer also wants to enable SRv6, which utilizes the IPv6 protocol to handle packets with more efficiency, increase security and performance, allowing for a significantly larger number of possible addresses.

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

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

The following topology provides the base for this scenario:



In this scenario, we will:

- Create a segment routing ODN template with a specific color on the endpoints to ensure that traffic is transported within an LSP (underlay) and that a best-path tunnel is created dynamically when a prefix with the specified color is received. Enable SRv6 (IPv6) for service and link details. The ODN template defines the SLA on which you want to optimize the path. In this case, we will optimize on latency.

- Specify that the computed paths be disjoint: they will not share the same link.
- Create a route policy on each endpoint to be used to bind the L3VPN to the ODN template. This route policy adds a color attribute to the customer prefixes and advertises via BGP to other endpoints. This color attribute is used to indicate the SLA required for these prefixes.
- Create an L3VPN service with 3 endpoints: PE-A, PE-B, and PE-C. This is the overlay configuration.
- Visualize how this overlay/underlay configuration optimizes the traffic path and automatically maintains the SLA.

Assumptions and Prerequisites

- To use ODN with SRv6, BGP peering for the prefixes must be configured between the endpoints/PEs. Usually for L3VPN, this is the VPNv4 and VPv6 address family peering, and this BGP peering is required to be over IPv6.

Workflow

- [Step 1. Create an ODN template to map color to an SLA objective and constraints](#)
- [Step 2. Create an L3VPN Route Policy](#)
- [Step 3. Create and provision the L3VPN service](#)
- [Step 4. Visualize the New VPN Service on the Map to See the Traffic Path](#)
- [Step 5. 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 dynamically create tunnels (on-demand) when prefixes with matching colors are received via BGP. Traffic to these prefixes will be automatically steered into the newly created tunnels, thereby meeting the SLA objective and constraints intended for these prefixes and signaled using colors in the BGP routes.

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

We will create the following ODN templates:

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

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

Procedure

1. Go to Services & Traffic Engineering > Provisioning (NSO) > 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 **srv6**, select the **Enable srv6** toggle.
5. Under locator, enter the required SRv6 **locator-name**.
Note: The locator name should match what is configured on the router.
6. Under dynamic, select “**latency**” as the metric type. This is the SLA objective on which we are optimizing.
7. Select the **pce** check box to specify that the path should be computed by the SR-PCE, not by the Path Computation Client (PCC).
8. Define the required constraints. In this case, we want the computed paths to be disjoint in that they must not share a link.
Under disjoint-path, choose **link** as the type, and specify a numeric group ID, in this case, 16.

ODN-Template {L3VPN_NM-SRTE-ODN_72-a}

name *

custom-template

[+]

name
PE-A

head-end

[+]

name
PE-A

maximum-sid-depth

color *

bandwidth

source-address

> srv6

dynamic

Enable dynamic

metric-type

latency

pce

flex-alg

> metric-margin

disjoint-path

Enable disjoint-path

type *

link

group-id *

sub-id

> affinity

✓ **srv6**

Enable srv6 [?](#)

✓ **locator**

Enable locator [?](#)

locator-name * [?](#)

behavior [?](#)

binding-sid-type [?](#)

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

Total 5 | Last Refresh: 12-Oct-2021 04:10:25 PM PDT | [O](#) [I](#) [⚙](#)

Name	Provisioning State	Date Created	Actions
L3VPN_NM-SRTE-ODN_70	✓ Success	12-Oct-2021 03:59:31 PM PDT	...
L3VPN_NM-SRTE-ODN_71-a	✓ Success	12-Oct-2021 03:57:33 PM PDT	...
L3VPN_NM-SRTE-ODN_71-c	✓ Success	12-Oct-2021 04:06:27 PM PDT	...
L3VPN_NM-SRTE-ODN_72-a	✓ Success	12-Oct-2021 03:53:41 PM PDT	...
L3VPN_NM-SRTE-ODN_72-b	✓ Success	12-Oct-2021 04:04:20 PM PDT	...

+ ↶ ☰

Manage Config View Edit Delete

Cross Launch View In NSO View Plan Data

Service Options

Check-Sync ?
Sync-From ?
Sync-To ?
Re-Deploy Dry Run ?

Re-Deploy ?
Re-Deploy Reconcile ?
Reactive-Re-deploy ?
Clean-Up ?

Configured Data

```

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

Copy To Clipboard **Cancel**

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

Step 2. Create an L3VPN Route Policy

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

We will create the following route policies:

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

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

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 IPv6 prefixes to identify the network traffic.
In this case, IPv6 prefix **70:70:70::0/64**.

L3vpn Route Policy {L3VPN_NM-SRTE-RP-PE-A-7}

color

id	exclusive	description
70	false	

extra-policy

name	operation	address
No Rows To Show		

color{70}

id *
70

exclusive
false

description

> **ipv4**

< **ipv6**
Enable ipv6
ipv6-prefix

ipv6-prefix
70:70:70::/64

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.

Note: After creating an L3VPN Route Policy, the VPN profile for each route policy will be automatically created. The VPN profile will be referenced from the L3VPN service. This will bind the route policy to the L3VPN service. Thus, VPN profiles for each of the route policies we created in the previous step will be created:

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

Step 3. Create and provision the L3VPN service

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

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

- Under ie-profile, click + to create a new ie-profile and give it a unique name.
- Enter the route distinguisher that will differentiate the IP prefixes and make them unique.
- Define the required VPN targets, including route targets and route target types (import/export/both).
- 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.

L3VPN_NM-SRTE-ODN-70

vpn-id * L3VPN_NM-SRTE-ODN-70

custom-template Selected 0 / Total 0

name
No Rows To Show

ie-profile{L3VPN_NM_SR_ODN-IE-PE-A-7 }

ie-profile-id * L3VPN_NM_SR_ODN-IE-PE-A-7

rd 0:70:70

vpn-targets * vpn-target Selected 0 / Total 1

id	route-target-type
100	both

vpn-policies * import-policy

export-policy L3VPN_NM-SRTE-RP-PE-A-7

g. Click X in the top-right corner when you are done.

h. Similarly, create the other IE profiles.

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

- Under vpn-nodes, click +, select the relevant device from the dropdown list, and click **Continue**.
- Enter the local autonomous system number for network identification.
- Select the ie-profile you created in the previous step.
- 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 > IPv6 > 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 (IPv6), the IP address of the BGP neighbor, and the number of hops allowed between the BGP neighbor and the PE device. Enable SRv6 and provide the address-family name and locator-name. The locator-name must match what is configured on the L3VPN PEs.
 - 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 **Dry Run** 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 4. 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 VPN overlay in the logical map. Use the buttons at the top right of the map to toggle between the logical and geographical maps.




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

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


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

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

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



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




Step 5. Observe automatic network optimization

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

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

Recomputed paths:

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



Summary and Conclusion

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

Scenario 3 – 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. Lastly, we will use the Service Health functionality to review the health of the services and view the Assurance Graph and Last 24Hr Metrics to better analyze our service's health details.

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

In this scenario, we will:

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

Assumptions and Prerequisites

- For transport mapping to L2VPN service, devices must be configured with the **l2vpn all** command.
- For Service Health enablement, Service Health must be installed.

Note: Service Health is not generally available yet. At this stage, it is available for pre-launch laboratory evaluation only. Engage your account team if you are interested in participating in the evaluation.

- Before using Service Health's Assurance Graph, ensure that topology map nodes have been fully configured and created with a profile associated to the service. If not, Subservice Details metrics will show that no value has yet to be reported.
- For Service Health, you must configure 2 buckets on the Y1731 profile associated with the device. If you have fewer than 2 buckets configured, Service Health cannot report the Y1731 probes/KPIs on the service details page.

Workflow

- [Step 1. Prepare for Creating a SID List](#)
- [Step 2. Create the SID List in the Provisioning UI](#)
- [Step 3. Create an explicit SR-TE policy for each VPN endpoint by importing a file](#)
- [Step 4. Create and provision the L2VPN service](#)
- [Step 5. Attach the SR-TE policies to the L2VPN Service](#)
- [Step 6. Enable Service Health monitoring](#)
- [Step 7. Visualize the L2VPN on the Map](#)
- [Step 8. Inspect a degraded service using Service Health and Last 24Hr Metrics to identify issues](#)

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

Step 1. Prepare for Creating a SID List

The SID list consists of a series of prefix or adjacency SIDs, each representing a node or link along on the path. Each segment is an end-to-end path from the source to the destination, and it instructs the routers in the network to follow the specified path instead of the shortest path calculated by the IGP.

To build the SID list, you will need the MPLS labels of the desired traversing path. You can get these labels from the devices themselves or you can invoke the northbound Cisco Crosswork Optimization Engine API to retrieve this information.

Refer to [Cisco Crosswork Network Automation API Documentation on Cisco Devnet](#) for more information about the API.

Procedure

1. Prepare the input required to produce the SID list for the path from endpoint to endpoint. You will need the router ID of each endpoint, as follows:

```
'{  
  "input": {  
    "head-end": "100.100.100.7",  
    "end-point": "100.100.100.5",  
    "sr-policy-path": {  
      "path-optimization-objective": "igp-metric"  
    }  
  }  
}'
```

2. Invoke the API on the Cisco Crosswork Network Controller server by using the input prepared in the previous step. For example:

```
curl --location --request POST  
'https://10.194.63.198:30603/crosswork/nbi/optima/v1/restconf/operations/cisco-crosswork-  
optimization-engine-sr-policy-operations:sr-policy-dryrun' \  
--header 'Content-Type: application/yang-data+json' \  
--header 'Accept: application/yang-data+json' \  
--header 'Authorization: Bearer  
eyJhbGciOiJIUzUxMiJ9.eyJzdWIiOiJhZG1pbiiSImlzRnJvbU5ld0xvZ2luIjoidHJ1ZSIiInBvbGljeV9pZCI6I  
mFkbWluiwiYXV0aGVudGljYXRpb25EYXR1IjoiMjAyMS0wMy0yM1QxNjozODozNy43NDY2MTZaW0dNVF0iLCJzdWN  
jZXNzZnVsQXV0aGVudGljYXRpb25IYw5kbGVycyI6I1F1ZXJ5RGF0YWJhc2VBdXRoZW50aWNhdGlvbkhbmRsZXIiL  
CJpc3Mi0iJodHRwOlwvXC9sb2NhbGhvc3Q6NTQ40VwvU1NPIiwibGFzdF9uYW1lIjoiic21pdGgiLCJjcmVkZW50aWF  
sVHlwZSI6I1VzZXJuYW1lUGFzc3dvcmRDcmVkJw50aWFsIiwiYXVkJjoiHR0cHM6XC9cLzEwLjE5NC42My4xOTg6M  
zA2MDNcL2FwcC1kYXNoYm9hcmQiLCJhdXRoZW50aWNhdGlvbk1ldGhvZCI6I1F1ZXJ5RGF0YWJhc2VBdXRoZW50aWN  
hdGlvbkhbmRsZXIiLCJsb25nVGVybUF1dGh1bnRpY2F0aW9uUmVxdWVzdFRva2VuVXNjZCI6ImZhbHNlIiwiY2hhb  
md1X3B3ZCI6ImZhbHNlIiwiZXhwIjoxNjE2NDU5OTIwLCJpYXQiOjE2MTY0MzExMjAsImZpcnN0X25hbWUiOijqb2h  
uIiwanRpIjoiU1QtODQtOFV1wXMybEt3R2d1Z3RIYj14MzMvTF1NTGVVRlp6OURyNGpoeFcxakhsV01VYXdXSWgxb  
UdTd01aRC1tOEK1s2Z0amI2ZmlWTUh1YnBYYjBMMFZqRFc2WVppUFVUbHRpNFVpZnNUeG9aQ284WWpPWEc2V1FjS0M  
wb29lWjJhc3BWanMzYnA3bHo5Vkhys1BCTz15TDNGcFRIWXRPeWJtVi1jYXMtMSJ9.Vi4k0w8Ks0v5M_08zBqWochT  
3k9V9Pn2Njsn5ES9c5PF-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"  
    }  
  }  
}'
```

3. 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

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

The screenshot shows the Cisco NM-OS provisioning interface. On the left, a table lists SID entries with index 1 and value 1. On the right, a detailed configuration pane for 'sid{1}' is shown. The configuration includes:

- index ***: 1
- type ***: mpls (selected)
- mpls *** (highlighted with a red box):
 - label ***: 23002

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 (NSO) > SR-TE > Policy**.
2. Click Import  above the table.
3. Download the sample .json or .xml file which will serve as a template for the required configuration.
In the Import Service dialog, click the **Download sample .json and .xml files (.zip)** link.



4. Unzip the downloaded file and open sr-Policy.xml in an XML editor.
5. Edit the xml file as required. Provide a name for the SR-TE policy, and specify the SID list to be associated with this policy. Save the xml file.

```

<config xmlns="http://tail-f.com/ns/config/1.0">
  <sr-te xmlns="http://cisco.com/ns/nso/cfp/cisco-tsdn-sr-te">
    <policies xmlns="http://cisco.com/ns/nso/cfp/cisco-tsdn-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>
        </path>
      </policy>
    </sr-te>
  </config>

```

6. 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.
7. Check whether the new SR-TE policy appears in the Policy table and its Provisioning State is **Success**.
8. Click ... in the Actions column and choose **Config View** to see to see the Yang model-based service intent data that details the SR-TE policy you created. You can also check the devices themselves to make sure that they were provisioned correctly.

Step 4. Create and provision the L2VPN service

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

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

Procedure

1. Go to **Services & Traffic Engineering > Provisioning (NSO) > L2vpn > L2vpn-Service**.
2. Click Import  above the table.
3. If you did not download the sample .json or .xml files in Step 3, 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}
            12-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. Enable Service Health monitoring

Note: Service Health is not generally available yet. At this stage, it is available for pre-launch laboratory evaluation only. Engage your account team if you are interested in participating in the evaluation.

1. Go to **Services & Traffic Engineering > VPN Services**. The map opens and a table of VPN Services is displayed to the right of the map.
2. In the Actions column, click ... for the new service you want to start monitoring health.
3. Select **Start Monitoring**.




The screenshot shows the 'VPN Services' table in the Cisco SD-WAN Provisioning interface. The table has columns for Health, Service ID, Type, Provisioning Status, Last Update, and Action. A context menu is open over the first row, showing options: 'View Details', 'Edit/Delete', and 'Start Monitoring'. The 'Start Monitoring' option is highlighted with a blue background.

Health	Service ID	Type	Provisioning	Last Update	Action
Success	I2	L2vpn-Service	Success	26-Oct-2021	...
Failed	I2nm-p2p...	L2vpn-Service	Success	24-Oct-2021	...
In Progress	I2vpn-115...	L2vpn-Service	Success	24-Oct-2021	...
Good	I2vpn-evp...	L2vpn-Service	Success	26-Oct-2021	...
Degraded	I2vpn-evp...	L2vpn-Service	Success	24-Oct-2021	...
Down	I2vpn-p2p...	L2vpn-Service	Success	26-Oct-2021	...

Note: The Health column color coding indicates the health of the service: Green = Good; Orange = Degraded; Red = Down; Gray = Not Monitoring.

4. The Monitor Service pop-up appears. Select which service to monitor and gain click **Start Monitoring**.



Note: Now that you have started monitoring the health of this service, in the Actions column, click ... to view the additional Service Health options: **Stop Monitoring**, **Pause Monitoring**, **Assurance Graph**.



5. Repeat this step for each new service you wish to start health monitoring.
6. Click **X** in the top-right corner when you are done.

Step 7. Visualize the L2VPN on the Map

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

Procedure

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

or


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

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


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

- 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.



- 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.



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

Note: Service Health is not generally available yet. At this stage, it is available for pre-launch laboratory evaluation only. Engage your account team if you are interested in participating in the evaluation.

In this step, we will review the Service Health assurance graph and utilize the Last 24Hr Metrics to identify issues within a specific time range. By isolating the issues within a specific time range, you can drill down on the details that may have caused the degraded (or down) service that can lead to troubleshooting the service or the node to address detailed symptoms. For this example, we will inspect a degraded service.

1. Click **X** in the top right corner to return to the VPN Services list.
2. Click on the name of a service that shows as being degraded. The map will update to highlight the service you selected.
Degraded services show an orange icon in the Health column. You can filter by health state by clicking in the space at the top of the column and selecting the appropriate filter. To clear the filter, click the **X** next to the designated filter appearing in the space at the top of the column and it will remove all filtering and default to showing all VPN Services in the list.
Note: If a service is not yet being monitored, the icon in the Health column will show as grey. To enable monitoring for such a service, click ... and select **Start Monitoring**.
3. In the Actions column, click ... and click **View Details**. The Service Details panel appears on the right side where you can review Active Symptoms for the service (including the Root Cause, Subservice, Priority, and Last Updated details) present in the Health tab if the service is being currently monitored. Review the details provided.

Service Details

Name I2nm-p2p-240
Provisioning Success
Health Degraded | Monitoring Profile Gold_L2VPN_ConfigProfile sys... ⓘ

Health **Transport** **Configuration** **Path Query**

Active Symptoms (2)

Root Cause	Subservice	Prior...	Last Updated
VPWS State degraded. Device: 172.16.1.11...	subservice.vpws.c...	15	27-Oct-2021...
VPWS State degraded. Device: 172.16.1.11...	subservice.vpws.c...	15	27-Oct-2021...

- Click on a Root Cause and view both the Symptom Details and the Failed Subexpressions & Metrics information.

Service Details

Name I2nm-p2p-240
Provisioning Success
Health Degraded | Monitoring Profile Gold_L2VPN_ConfigProfile sys... ⓘ

Health **Transport** **Configuration** **Path Query**

Symptom Details

Name	VPWS State degraded. Device: 172.16.1.118, XConnectGroup: I2nm-p2p-240, XconnectName: I2nm-p2p-240
Sub Service	subservice.vpws.ctrlplane.health system
Last Updated	27-Oct-2021 03:56:16 PM GMT+3

Failed Subexpressions & Metrics

Show Only Failed [Expand All](#) | [Collapse All](#)

Name	Expression Value
xconnect_state == 'up' && ac_state == 'up' && evpn_state == 'up'	false
subExps	
xconnect_state == 'up' && ac_state == 'up' && evpn_state == 'up'	false
xconnect_state == 'up' && ac_state == 'up' && evpn_state == 'up'	false
symptomMetrics	
subExps	
xconnect_state == 'up' && ac_state == 'up' && evpn_state == 'up'	false
xconnect_state == 'up' && ac_state == 'up' && evpn_state == 'up'	false
symptomMetrics	
subExps	


- To further isolate the degraded service details, click **X** in the top-right corner to return to the VPN Services list.
- Again, click on the name of the degraded service in the list. The map will update and isolate the corresponding devices participating with that service.
Within the map, to view further service health details, do the following:

- a. In the map, hover your mouse over the device showing as degraded and its smaller badges that indicate health status and review the pop-up information relating to its Reachability State, Host Name, Node IP, and Type.
 - b. At the top-left of the map panel, you may also select the **Show Participating Only** check box so the map only shows the participating services.
7. In the Actions column, click ... for the degraded service in the list and click **Assurance Graph**. The topology map of services and subservices appear with the Service Details panel showing Service Key, Status, and Sub Services details. Metrics also appear, such as Jitter-RT (Jitter Round Trip), Latency-RT (Latency Round Trip), PktLoss-DS (Packet Loss from Destination to Source), and PktLoss-SD (Packet Loss from Source to Destination). Additionally, a table of Active Symptoms listing Root Cause, Subservice, Priority, and Last Updated details is populated.

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

8. At the top-right of the screen, select the **Show History** mode toggle. The historical Date Range graph appears. This graph shows different ranges of historical health service monitoring details from one day (1d) up to sixty days (60d). You can select the (+) icon at the top-right to zoom in on the event or use your mouse to draw a rectangle over events to further zoom. Events that are consecutive may appear as a line of white space.


Note: After selecting **Show History**, the Date Range graph opens by default to the range of one day (**1d**). If no events of significance have occurred over the last day, or even in the last several days (such as **1d to 3d**), it is possible no data is available to display. Select a different time range, starting with **5d** to broaden the data collection range that is available to display.



Note: When you select an event on the Date Range graph, a tool tip with information about that event appears (such as date and time of the event, and severity level and number of symptoms). Click anywhere within the chart to hide the tool tip.

9. Review the Root Cause information by either hovering your mouse over a particular row or click the arrow to expand the Service Details panel to full screen mode. Columns can be resized using your mouse or you can select the gear icon to deselect or select columns you want to appear.

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



10. You can also use the map and click on the degraded node to bring up Service Details information on both Active Symptoms and Impacted Services.

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



Note: If you view the Subservice Details panel, each subservice metric (Jitter-RT, Latency-RT, PktLoss-DS, PktLoss-SD) will initially report a value of zero. Based on a device's configuration, it may take up to 10 minutes for the metric values to begin reporting.

11. Use the active and impacted information to inspect the degraded service details to determine the issues that led to the degraded service.
12. To further isolate the possible issues and to utilize the **Last 24Hr Metrics**, perform the following steps:

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

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

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



13. Next, select the **Show: Down & Degraded Only** check box in the top-left corner of the map so only Subservices which are degraded, along with other dependent but healthy subservices, appear. Inspect the Service Details panel showing the active symptoms and their root cause.
14. Deselect the **Show: Down & Degraded Only** check box and select the **Soft Dependencies** check box in the top-left corner of the map. Soft Dependencies implies that a child subservice's health has a weak correlation to its parent's health. As a result, the degraded health of the child will not result in the parent's health degradation. Use the + or - symbols in the bottom-right corner of the map to zoom in or out on services mapped. Select the ? to view the Link Color Legend that explains all of the icons, symbols, badges, and colors and their definitions.
Note: You can also select the **Subservices** icon in the top-right corner of the map to show service appearance options.
15. Select the degraded service in the map to show the subservice details.
16. Select the **Active Symptoms** tab to show any root causes for the service health details that are displayed and then select the **Impacted Services** tab to show services where their health is degraded.

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

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. Lastly, we used Service Health to monitor the health of the new service using the Assurance Graph, Last 24hr Metrics, and SubExpressions metrics to view when service may have been up, degraded, or down, and what the root causes were identified.

Scenario 4 – 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.
- For Service Health enablement and usage to monitor a services health, Service Health must be installed.
- For initializing a Heuristic Package to monitor health of a services (optional), see the Appendix section, [Initializing Heuristic Packages to monitor the health of a service](#), for detailed steps to be performed prior to starting monitoring.

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 (NSO) > 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
<p>Summary</p> <p>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</p> <p>Routing</p> <p>TE Router ID 100.100.100.5 ISIS System ID 0000.0000.0005 Level-1/2 ASN 1</p>	

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.

- Under p2p-primary-paths, click + to create a new path.
- In the pane that opens on the right, give the path a name.
- Select the path computation method – **path-locally-computed**.
- Specify a numeric preference for the path. The lower the number, the higher the preference.
- Define the optimization metric, in this case, **igp**.

The screenshot shows the configuration of an RSVP-TE Tunnel and its associated p2p-primary-path. The tunnel configuration includes:

- RSVP-TE Tunnel {L2VPN_NM-P2P-RSPVTE-PE-A-2220}**
- signaling-type**: te-types:path-setup-rsvp
- head-end ***: PE-A
- tail-end**: PE-B
- te-bandwidth** (generic):
 - technology**: generic
 - generic**: 1200
- p2p-primary-paths**
- traffic-steering**

The p2p-primary-path configuration is as follows:

- p2p-primary-path{L2VPN_NM-P2P-RSPVTE-PE-A-2220 }**
- name ***: L2VPN_NM-P2P-RSPVTE-1
- path-computation-method**: path-locally-computed
- preference**: 1
- optimizations**
- explicit-route-objects-always**

At the bottom, there are buttons: **Commit changes**, **Dry Run**, **Delete**, and **Cancel**.

- Click **Commit Changes**.
- 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 | [O](#) [G](#)

	Name	Provisioning State	Date Created	Action
+	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-RSVPT-E-PE-A-2220	Success	17-Mar-2021 11:28:30 AM G...	...
-	L2VPN_NM-P2P-RSVPT-E-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.

RSVP-TE Tunnel Details

Summary

Headend	▲ PE-A (100.100.100.5)
Endpoint	● PE-B (100.100.100.6)
Tunnel ID	2220
Description	-
Path Name	L2VPN_NM-P2P-RSVPT-E-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
1	P-TOPRIGHT	20.20.10.14	GigabitEthernet0/0
2	PE-B	20.20.10.26	GigabitEthernet0/0
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 3—Mandate a static path for an EVPN-VPWS service using an explicit SR-TE policy.

1. Go to **Services & Traffic Engineering > Provisioning (NSO) > L2VPN > L2vpn Service**.
2. Click **+** to create a new service and give it a unique name. Click **Continue**.

3. Choose vpn-common:t-ldp in the vpn-svc-type field.
4. Define each VPN endpoint individually – PE-A and PE-B.
 - a. Under vpn-nodes, click +.
 - b. Select the relevant device from the vpn-node-id and ned-id dropdown lists and click **Continue**.
 - c. Enter the local autonomous system number for network identification.
5. Define the LDP signaling options by creating one or more pseudowires. In this case, specify the TE router ID of the peer device (PE-B), and provide a unique numeric label to identify the pseudowire.
6. Attach the RSVP tunnel to the service:
 - a. Under **te-service-mapping > te-mapping**, click the **te-tunnel-list** tab.
 - b. 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 provide the tunnel ID under the te-tunnel-id tab.

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

Services & Traffic Engineering / Provisioning

L2vpn Service

Total 15 | Last Refresh: 04-Apr-2021 12:22:38 PM GMT+3 | ⚙️ | ⚙️

Services/Policies	Vpn Id	Provisioning State	Date Created	Actions
Recent				
Global				
Resource Pool				
L2VPN				
ID-Pools	L2NM-EVPN-EXPLICIT-180	Success	17-Mar-2021 11:29:22 AM GMT...	...
L2vpn Route Policy	L2NM-SRTE-PW-DYNAMIC-190	Success	17-Mar-2021 11:31:14 AM GMT...	...
L2vpn-Service	L2VPN_NM-EVPN-VPWS-NATIVE-200	Success	17-Mar-2021 11:27:32 AM GMT...	...
L3VPN	L2VPN_NM-EVPN-VPWS-SRTE-230	Success	17-Mar-2021 11:28:27 AM GMT...	...
L3vpn Route Policy	L2VPN_NM-EVPN-VPWS-SRTE-ODN-250	Success	17-Mar-2021 11:28:09 AM GMT...	...
L3vpn-Service	L2VPN_NM_P2P-NATIVE-210	Success	17-Mar-2021 11:27:19 AM GMT...	...
VPN Profiles	L2VPN_NM_P2P_RSVPTE-2220	Success	17-Mar-2021 11:28:45 AM GMT...	...
	L2VPN_NM_P2P_SRTE-240	Success	17-Mar-2021 11:27:51 AM GMT...	...
	I2nm-p2p	Failed	28-Mar-2021 09:57:03 AM GMT...	...
	I2vpn-p2p-rsvp	Success	31-Mar-2021 02:31:37 AM GMT...	...

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

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

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



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.



- 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 5 – 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 (NSO) > 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 on 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.

5. 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 *
192.168.100.7

color *
787878

binding-sid

path * Selected 0 / Total 1

	preference
	123

bandwidth
250000

6. 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
	srtc_c_300_ep_100.100.100.322222	Success	10-Feb-2021 06:52:38 PM PST

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

- a. Click ... in the Actions column and choose **View**.

- b. 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 the BWoD functionality within Cisco Crosswork Network Controller is enabled.




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, hover over the mouse over **?**.
3. Click **Commit 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 IGP Path** check box to see the physical path between the endpoints.



Summary and Conclusion

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

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 issues 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 an enhanced functionality introduced in Cisco Crosswork Network Controller 3.0.


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 an issue locally eliminates the need for simulating edge-to-edge traffic flows in the network through a full traffic matrix, which is both cumbersome to create and is less scalable as node counts continue to increase.

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 through SNMP. It estimates the amount of traffic that may be diverted and, if the user approves, performs the mitigation through the deployment of Tactical Traffic Engineering (TTE) SR-TE policies. Mitigating congestion locally does not require the use of the full Segment Routing Traffic Matrix (SR-TM). TTE SR-TE policies are created at the device on only either side of the congested link, with the shortest paths possible that do not congest interfaces elsewhere.

How Does LCM Work?

1. 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).
2. 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).

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:

$$\text{Total interface traffic} - \text{SR-TE policy traffic} = \text{Eligible traffic that can be optimized}$$

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:

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


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

- two TTE SR-TE policies to take a direct path to Node 3 (200 Mbps)
- one 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 6 – Use Local Congestion Mitigation \(LCM\) to reroute traffic on an over-utilized link](#)

Additional Resources


[Cisco Crosswork Optimization Engine Documentation](#)

[**Scenario 6 – 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 P4-NCS5501 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.

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

```
segment-routing traffic-eng pcc profile <id> autoroute force-sr-include
```

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

For more information, contact your Cisco Account representative.

Workflow

- [Enable LCM and configure the utilization threshold](#)
- [View link congestion on the map](#)
- [View TTE SR-TE policy recommendations in the LCM dashboard](#)
- [Validate the TTE SR-TE policy deployment](#)
- [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 > Domain-ID > ... > 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 12%, and the Interfaces to Monitor > Selected Interfaces option is selected.

To see information about other configuration options, hover the mouse over

Configuration

Basic Advanced

The screenshot shows the 'Advanced' configuration page for Local Congestion Mitigation. It includes fields for Enable (True), Color (2000), Utilization Threshold (12), Utilization Hold Margin (5), Delete Tactical SR Policies when Disabled (False), Profile ID (0), Congestion Check Interval (900), Max LCM Policies per Set (8), Description (LCM Startup Config), and Interfaces to Monitor (Selected Interfaces). The 'Utilization Threshold' and 'Interfaces to Monitor' sections are highlighted with red boxes.

Utilization Threshold: 12

Interfaces to Monitor: Selected Interfaces


Commit Changes Get Default Values Discard Changes

3. Click **Commit Changes**.

Step 2. View link congestion on the map

The link between PE1-ASR9k and P1-ASR9k is now congested. Let's see that on the map.

1. Go to **Services & Traffic Engineering > Traffic Engineering**.
2. Click on the link to view link details, including utilization information. Note that utilization on the P4-NCS5501 interfaces has surpassed 12%, the threshold that was defined in the LCM configuration.



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.

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

When congestion is detected, a domain will have a timestamp of recommended actions.

LCM Domains

Domain Identifier 100 Enabled LCM Startup Config Last recommendation timestamp: 13-Oct-2021 01:43:10 PM PDT Recommendation Available	Domain Identifier 0 Disabled LCM Startup Config
--	---


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

The dashboard shows that the utilization has surpassed 12% and is at 13.01%. In the Recommended Action column, there is a recommendation to deploy TTE policy solution sets (Create Set) to address the congestion on the interface. The Expected Util column shows the expected utilization of each of the interface if the recommended action is committed.

Operational Dashboard

Last Refresh: 12-Oct-2021 09:14:03 AM PDT													
Congested Interfaces (0) Mitigating Interfaces (0) Mitigated Interfaces (0)													
Commit All		Last Recommendation: 12-Oct-2021 09:02:23 AM PDT		Urgency: MEDIUM									
Node	Interface	Threshold Util	Evaluatio...	LCM State	Pol...	Pol...	Recomm...	Com...	Expected Util	Solution Up...	Actions		
P4-NC55	GigabitEthernet0/0/0/1	12%	13.01%	Congested	0	-	Create Set	None	11.97%	12-Oct-202...	...		

3. 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. Note the Color IDs to later verify TTE policy deployment.



Recommended TTE Policies (Preview)

Node P4-NCS5501
Interface GigabitEthernet0/0/0/4

	Headend	Endpoint	Color	Recommended Action
<input type="checkbox"/>	P4-NCS5501	P1-ASR9k	1910	CREATE
<input checked="" type="checkbox"/>	P4-NCS5501	P1-ASR9k	1911	CREATE

- After you are done viewing the recommended TTE policies on the map, navigate back to the **Operational 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. Because of dependencies between solution sets, the mitigation solution is based on all LCM recommendations being committed.

Operational Dashboard

Last Refresh: 12-Oct-2021 09:14:03 AM PDT |

Commit All											
Last Recommendation: 12-Oct-2021 09:02:23 AM PDT Urgency: MEDIUM											
Node	Interface	Threshold Util	Evaluatio...	LCM State	Pol...	Pol...	Recomm...	Com...	Expected Util	Solution Up...	Actions
P4-NCS55	GigabitEthernet0/0/0/1	12%	13.01%	Mitigating	2	-	Create Set	None	11.97%	12-Oct-202...	

Step 4. Validate the TTE SR-TE policy deployment

- Click > **Events** tab to open the Events window in which you can monitor LCM events. You see events for the LCM recommendations, the commit actions, as well as any exceptions.

Alarms & Events

All System Network

Alarms **Events**

Selected / Total 929

Filters Applied (1)

Source	Severity	Description	Creation Time	Category	Correlated Alarm
LCM					
LCM for domain 100	Info	A new recommendation has been created: 2 creates, 0 updates, 0 delete...	30-AUG-2021 04:56:33 P...	System	NO
LCM for domain 100	Info	Recommendation committed.	30-AUG-2021 04:45:31 P...	System	NO
LCM for domain 100	Info	A new recommendation has been created: 0 creates, 0 updates, 6 delete...	30-AUG-2021 04:44:51 P...	System	NO
LCM for domain 100	Major	Mitigated interface F2.cisco.com GigabitEthernet0/0/0/5 is down.	30-AUG-2021 04:44:50 P...	System	NO
LCM for domain 100	Info	A new recommendation has been created: 0 creates, 2 updates, 4 delete...	30-AUG-2021 04:25:46 P...	System	NO
LCM for domain 100	Info	Recommendation committed.	30-AUG-2021 04:00:46 P...	System	NO
LCM for domain 100	Info	A new recommendation has been created: 1 creates, 5 updates, 0 delete...	30-AUG-2021 03:52:29 P...	System	NO
LCM for domain 100	Info	LCM is enabled	30-AUG-2021 03:52:11 P...	System	NO
LCM for domain 101	Info	LCM Worker with domain_id: '101' has started.	30-AUG-2021 03:52:04 P...	System	NO
LCM for domain 100	Info	LCM Worker with domain_id: '100' has started.	30-AUG-2021 03:52:04 P...	System	NO
LCM for domain 101	Info	LCM is disabled	30-AUG-2021 03:52:03 P...	System	NO

- 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) |

Commit All

Node	Interface	Thresho...	Evalu...	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-MPLS 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.

SR Policy Details

Headerend: P4-NCS5501 (TE RID: 192.168.60.24) PCC IP: 192.168.60.24
Source IP: 192.168.60.24

Endpoint: P1-ASR9k (TE RID: 192.168.60.21)
Dest IP: 192.168.60.21

Color: 1911

Summary

- Admin State:
- Oper State:
- Binding SID: 1004019
- Segment Type: Unprotected
- Policy Type: Local Congestion Mitigation** (highlighted with a red box)
- Profile ID: 1981
- Description: -
- Utilization: 0 Mbps
- Delay: 2
- BWOD Policy Bandwidth: 0 Mbps
- Accumulated Metric: 0
- Delegated PCE: 172.29.10.121
- Non-delegated PCEs: 172.29.10.122
- PCE Computed Time: 14-Oct-2021 12:39:21 PM PDT
- Last Update: 14-Oct-2021 12:39:21 PM PDT

Candidate Path

Path Name: Icm_to_P1-ASR9k_c_1911
Preference: 100
Path Type: Explicit

- 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 a period, 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.

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.

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 APIs.

How Does it Work?



- Using the Crosswork Network Controller APIs, you can create topology snapshots (plan files) which capture and represent topology state at a given point in time, including the IGP topology as well as interface level statistics (traffic load). For impact analysis purposes, these snapshots should be representative of a time period to be evaluated for an upcoming maintenance activity. For example, if you are planning a router upgrade at midnight on a Monday, you would take snapshots from several Mondays at midnight to evaluate typical traffic loads at this time. You can export these plan files to a central storage repository, where a library of topology plan files can be stored for a specified period of time.
- Cisco WAE Design allows you to explore “what-if” scenarios relevant to the planning of the maintenance window. For example, in the case of upgrading a router, Cisco WAE Design can simulate the resulting traffic load on the remaining devices after traffic is diverted from the device being upgraded. You can also explore the impact of deploying tactical traffic engineering policies to further optimize the topology during the maintenance window. For more information, contact your Cisco Customer Experience representative.

Usage Scenarios

Scenario 7 – 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 7 – 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

- [Download Topology Plan Files for Impact Analysis](#)
- [Schedule and execute the SMU by running a playbook](#)
- [Verify the SMU install job completion status](#)

Step 1. Download Topology Plan Files for Impact Analysis

When considering when to take down a device for maintenance so that there will be the least impact to the network, you need information about the traffic trends around that device at the targeted time. Using the Cisco Crosswork Optimization API, you can download plan files that capture a snapshot of the network topology at that time. If you download plan files at the same time over a period of time, you can use Cisco WAE Design to analyze the traffic trends. Based on this analysis, you can decide whether the impact to the network would be acceptable or not.

Refer to [Cisco Crosswork Network Automation API Documentation on Cisco Devnet](#) for more information about the API.

Procedure

1. Prepare the input required to download the plan file. You need to specify the version of Cisco WAE design that you will be using for analysis and the format in which you want the plan file, either txt or pln.

Note: If you download the plan file as a txt file, you can view it in any text editor. If you download it as a pln file, you can open it only in Cisco WAE design.

The input for this scenario is as follows:

```
'{  
  "input": {  
    "version": "7.3.1",  
    "format": "txt",  
  }  
}'
```

2. Invoke the API on the Cisco Crosswork Network Controller server using the input prepared in the previous step. For example:

```
curl --location --request POST
'https://10.194.63.198:30603/crosswork/nbi/optima/v1/restconf/operations/cisco-crosswork-
optimization-engine-operations:get-plan \
--header 'Content-Type: application/yang-data+json' \
--header 'Accept: application/yang-data+json' \
--header 'Authorization: Bearer
eyJhbGciOiJIUzUxMiJ9.eyJzdWIiOiJhZG1pbisImlzRnJvbU5ld0xvZ2luIjoidHJ1ZSiSInBvbGljeV9pZCI6I
mFkbWluIiwiYXV0aGVudGljYXRpb25EYXR1IjoiMjAyMS0wMy0yMlQxNjozODozNy43NDY2MTZaW0dNVF0iLCJzdWN
jZXNzNvVsQXV0aGVudGljYXRpb25IYw5kbGVycyI6I1F1ZXJ5RGF0YWJhc2VbdXRoZW50aWNhdGlvbkhbmRsZXiL
CJpc3Mi0iJodHRwOlwvXC9sb2NhbGhvc3Q6NTQ40VwvU1NPIiwibGFzdF9uYW1lIjoiic21pdGgiLCJjcmVkZW50aWF
sVHlwZSi6IlVzZXJuYW1lUGFzc3dvcmRDcmVkZW50aWFsIiwiYXVkJioiaHR0cHM6XC9cLzEwLjE5NC42My4xOTg6M
zA2MDNcL2FwcC1kYXNoYm9hcmQiLCJhdXRoZW50aWNhdGlvbk1ldGhvZCi6I1F1ZXJ5RGF0YWJhc2VBdXRoZW50aWN
hdGlvbkhbmRsZXiLCJs25nVGVybUF1dGh1bnRpY2F0aW9uUmVxdWVzdFRva2VuVXN1ZCi6ImZhbHN1IiwiY2hhb
mdlX3B3ZCi6ImZhbHN1IiwiZXhwIjoxNjE2NDU5OTIwLCJpYXQiOjE2MTY0MzExMjAsImZpcnN0X25hbWUiOjqb2h
uiIwianRpIjoiU1QtODQtOFV1WXMybEt3R2d1Z3RIYj14MzVmTF1NTGVVR1p6OURyNGpoefcxakhsV01VYXdXSWgxb
UdTd01aRC1t0Ek1S2Z0amI2Zm1WTUh1YnBYYjBMMFZqRFC2WVppUFVUbHRpNFVpZnNUeG9aQ284WwpPWEc2V1FjS0M
wb29lWjJhc3BWanMzYnA3bHo5VkhjS1BCTz15TDNGcFRIWXRPelWJtVi1jYXMtMSJ9.Vi4k0w8Ks0v5M_08zBqWochT
3k9V9Pn2Njsn5ES9c5PF-4ds0o4kk6xuZx5_cggauiEICuUMnzmRzneST-oAuA' \
--data-binary '{
  "input": {
    "version": "7.3.1",
    "format": "txt",
    ""
  }
}'
```

3. 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+C1Byb3B1cnR5CVZhHV1C1RpdGx1CQpWZXJzaW9uCTcuMy4xCgo8TmV0d29ya09wdGlvbnM+Ck9wdG
1vbglWYwx1ZQpFbmZvcmN1QWRqU01ETG9jYWxpeM0aw9uCVRSVUUKCjxDaXj jdW10cz4KTmFtZQ1<<>>Ob2R1Q0
1JbnR1cmZhY2VBCU5vZGVCCU1udGVyZmFjZUIJQ2FwYwNpdHkJRGVsYXkJRG1zdMJTmV0SW50U05NUF9Fcnjvcg10Z
XRJbnRTb3VyY2UJTmV0SW50UkUwQ1BVMW0JTmV0SW50UkUwQ1BVNWZpZXIJQwxb3JpdGhtCVJmbGFnCU5mbGFnCVB
mbGFnCUVmbGFnCVZmbGFnCuXmbGFnCg=="
  }
}
```

4. 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 an optimal 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 through 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 > Play List**, 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.

The screenshot shows the 'Run Playbook' interface. At the top, there are five tabs: 'Select Playbook' (highlighted with a blue dot), 'Select Devices', 'Parameters', 'Execution Policy', and 'Confirm'. Below these tabs, the interface is divided into two main sections. On the left, under 'Available Playbooks', there is a search bar and a list of playbooks. The 'All Playbooks' tab is selected. The list includes: Delete ACL based traffic steering configs, Delete DSCP based traffic steering configs, Enable/Disable traffic collector, **Install a SMU or an optional package on a router** (which is highlighted with a blue border), Interface State change for XE device, Interface State change on XR, Node State Snapshot, Prefix Set ADD/DELETE, Push config using NAMs, and Remove interface from Bundle Ether and delete bundle-ether. At the bottom left is a 'Cancel' button. On the right, under 'Install a SMU or an optional package on a router', it shows the last modified date (14-Oct-2020, 1:45 AM by Cisco), the software platform (IOS XR, Version 1.0.0), and a description: 'Install SMU or an optional package on a router'. It also displays the playbook's structure with three sections: 'Pre Maintenance (1)' containing '1 Verify package consistency on router'; 'Maintenance (4)' containing '2 Perform DLM node lock on device(s)', '3 Install add package(s)', '4 Install activate package(s)', and '5 Install commit package(s)'; and 'Post Maintenance (1)' containing '6 Verify package in committed list on router'.

- Click **Next** to go to the next task: Select Devices. All devices tagged with City: NY will be selected for SMU installation.
- Under the City tag on the left, click **NY**. The devices tagged with NY are listed on the right and are automatically selected.

○ Select Playbook **Select Devices** Parameters Execution Policy ○ Confirm

LIST Select Device Tag Select Device Manually Allow Bulk Jo

Select Tags* [Clear All](#) NY X Tags will be resolved dynamically at runtime to determine constituent de

City Devices with selected tag

Reachability St...	Operational State	Host name	Software Pla...	Provider	Unique Identifier
Reachable	OK	P-BOTTOMRIGHT	IOS XR	bcc1bc0c-d1cc-4932-90	
Reachable	OK	P-TOPRIGHT	IOS XR	ce944bd2-c476-4391-9	

Default

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.

○ Select Playbook Select Devices **Parameters** Execution Policy ○ Confirm

✓ Install a SMU or an optional package on a router



✓ Verify package consistency on router ?




collection_type

mdt

Data collection type

✓ Perform DLM node lock on device(s) ?



retry_count

3

Number of time node lock will be retried

retry_interval


5s

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.

Execution Policy

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) [?](#)

2021-04-09	Add date	
<input type="button" value="Increment hours"/>	<input type="button" value="Increment minutes"/>	
00	:	42
<input type="button" value="Decrement hours"/>	<input type="button" value="Decrement minutes"/>	

Schedule

Perform (Asia/Jerusalem) [?](#)

2021-04-09	Add date	
<input type="button" value="Increment hours"/>	<input type="button" value="Increment minutes"/>	
00	:	42
<input type="button" value="Decrement hours"/>	<input type="button" value="Decrement minutes"/>	

All Scheduled Jobs

Show jobs for selected devices only

Month Week Day

April 2021

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.
11. 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/smuss",
      "address": "192.168.6.1",
      "destination": "harddisk:",
      "dlm_credential_profile": "abc"
    }
  }
}
```

Label your Job

Name* smu_upgrd

Labels update

update X

[Cancel](#)

[Previous](#)

[Run Playbook](#)

Step 3. Verify the SMU install job completion status

- 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.

The screenshot shows the 'Job Sets' page with a list of jobs. On the left, there's a table with columns for Status, Name, and Id. A checkbox is checked for the first job, 'smu_xrv-77993990ce'. To the right, detailed information for this job is shown: 'Status Success', 'Job Set Tags' (empty), 'PlayBook Title router_op_smu_upgrade', and 'Created By admin'. Below this, a section titled 'All Jobs in the Set (1)' shows a table with columns for Status, Device, Execution ID, Start Time, and End Time. The single job listed is 'Succeeded' on device 'xrv9k-1' with execution ID '1613667141147-5b7e0cec-7c19-4368-bf'.

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

The screenshot shows the 'Job History' page for the job set 'smu_xrv-77993990ce'. It includes sections for Playbook (Install a SMU or an optional package on a router), Device (xrv9k-1), Status (SUCCEEDED), and Parameters (View). The main area displays the 'Execution Mode' with steps: Pre Maintenance 1/1 (Verify package consistency on router), Maintenance 4/4 (Perform DLM node lock on device(s), Install add package(s), Install activate package(s), Install commit package(s)), and Post Maintenance 1/1 (Verify package in committed list on router). The status for all steps is 'SUCCEEDED'. On the right, the 'Events' tab is active, showing logs for generic events, MOP status, MOP task events, and node status updates.

- 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:

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.

Programmable Closed-Loop 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 be performed after a defined Key Performance Indicator (KPI) threshold has been breached. Remediation can be implemented with or without the network operator's approval, depending on the setting and preferences of the operator.

Using such closed-loop remediation reduces the time taken to discover and repair a problem while minimizing the risk of making a mistake and creating an additional error through high-stakes manual network operator intervention.

How Does it Work?

Smart Monitoring

- The Smart Monitoring feature helps operators collect, filter, and present the data in useable formats, such as graphs and tables. Operators can remain focused on their business goals while the configuration required for the data collection is done by the Cisco Crosswork Network Controller and Cisco Crosswork Change Automation and Cisco Crosswork Health Insights using the feature Zero-touch telemetry.
- By using a common collector to collect network device data over SNMP, CLI, and model-driven telemetry, and making it available as modelled data described in YANG, duplicate data collection is avoided, optimizing the load on both the devices and the network.
- Recommendation Engine analyzes network device hardware and software, configuration, and employs a pre-trained model built from data mining, producing KPI relevant recommendations facilitating per use-case monitoring.
- KPIs cover a wide range of statistics from CPU, memory, disk, layer 1/2/3 network counters, to per protocol, LPTS and ASIC statistics.

Smart Filtering

- Cisco Crosswork Health Insights builds dynamic detection and analytics modules that allow operators to monitor and see alerts on network events based on user-defined logic (KPI).
- Key Performance Indicators (KPIs) Alerting Logic can be :
 - Simple static thresholds (TCA); e.g., CPU load above 90 percent.

- Moving average, standard-deviation, and percentile based, etc., e.g., CPU load above mean and staying there for five minutes.
 - Streaming jobs which provide real-time alerts or batch jobs which run periodically.
 - Customized for threshold values and visualization dashboards.
 - Customized operator-created KPIs based on business logic.
 - TCAs can be exported or integrated with other systems via HTTP, Slack, and socket interfaces.
- KPIs are associated with dashboards, which provide real-time and historical views of the data and corresponding TCAs.
 - KPIs also provide purpose-built dashboards that go beyond raw data and provide valuable information in various infographic style charts and graphs useful for triaging and root-causing complex issues.

Smart Remediation

- Health Insights KPIs can be associated with Cisco Crosswork Change Automation (CCCA) playbooks, which can be either executed manually or via auto-remediation. Remediation workflow could be used to fix the issue or collect more data from the network devices. By proactively remediating the situation, instead of resorting to ad hoc debugging and unscheduled downtime, operators can save time and money, providing better QOE to their customers.
- Health Insights does the correlation of alerts or anomalies on the topology of the network, allowing easy visualization of the impact of events.

Scenario 8 – Achieve Predictive Traffic Load Balancing Using Segment Routing Affinity

Scenario Context

To maintain smooth and optimal traffic flow, operators need to be able to monitor traffic on the interfaces, identify errors such as CRC, watchdog, overrun, and then reroute the traffic so that the SLA is maintained. This process can be automated using Cisco Crosswork Network Controller with Cisco Crosswork Health Insights and Cisco Crosswork Change Automation applications installed.

Assumptions and Prerequisites

Cisco Crosswork Health Insights and Cisco Crosswork Change Automation must be installed and running.

Workflow

The following is a high-level workflow for executing this scenario:

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 [Create an ODN template to map color to an SLA objective and constraints](#) in [Scenario 1 – Implement and Maintain SLA for an L3VPN Service for SR-MPLS \(using ODN\)](#).

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 [Create an ODN template to map color to an SLA objective and constraints](#) in [Scenario 1 – Implement and Maintain SLA for an L3VPN Service for SR-MPLS \(using ODN\)](#).

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

4. Define and enable the KPIs on the devices. This will continuously monitor the uplink interfaces on the L3VPN endpoints.

For information about defining KPIs, see the [Cisco Crosswork Change Automation and Health Insights User Guide](#).

5. When there is an error on monitored interfaces, mark the dirty link with RED affinity so that it will be excluded, based on the specifications of the ODN template. This is achieved by creating a custom playbook. Cisco Crosswork Network Controller learns the name of the interface generating the alert regarding the error and this is fed into the custom playbook so that the affinity configuration can be pushed to the relevant router, forming a closed-loop automation scenario. In this way, the customer should not experience outages.

For information about defining playbooks, see the [Cisco Crosswork Change Automation and Health Insights User Guide](#).

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

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 automatically applied to the device. After it is provisioned in this way, the new device is onboarded to the Crosswork device inventory where it can be monitored and managed like other devices.

How Does it Work?

- Classic ZTP: The DHCP server verifies the device's identity based on the device's serial number, then offers downloads of the boot file and image. After the device is imaged, it downloads the configuration file and executes it.
- Secure ZTP: The device and the Cisco Crosswork ZTP bootstrap server authenticate each other using the device's Secure Unique Device Identifier (SUDI) and Crosswork server certificates over TLS/HTTPS. After a secure HTTPS channel is established, the Crosswork bootstrap server allows the device to request to download and apply a set of signed image and configuration artifacts adhering to the RFC 8572 YANG schema. After the image (if any) is downloaded and installed, and the device reloads with the new image, the device downloads configuration scripts and executes them.
- Plug and Play (PnP) ZTP: The Cisco PnP agent on the IOS-XE device and the Cisco Crosswork PnP Server authenticate each other over HTTP using a PnP profile supplied on a TFTP server. They then establish a secure connection over HTTPS and the PnP agent downloads and installs image (optional) and configuration artifacts.

Additional Resources

Detailed information is available in the ZTP chapter in the [Cisco Crosswork Infrastructure 4.1 and Applications Administration Guide](#).

Scenario 9 – 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 IOS-XR devices required to set up a new customer site in a remote location and go live, without the need to send skilled technicians on time-consuming and costly on-site visits to complete the provisioning.

We will leverage the configuration of devices at existing customer sites that are already set up and operating to ensure that the Day0 configuration of the new devices includes whatever is necessary to get the devices up and running quickly and efficiently.

Assumptions and Prerequisites

- Crosswork ZTP must be installed in your Cisco Crosswork Network Controller setup.
- For Classic ZTP, Crosswork and the devices must be deployed in a secure network domain. Secure ZTP does not have this requirement; it is secure across networks.
- The Crosswork server must be reachable from the devices, via an out-of-band management network or an in-band data network.
- If you want to onboard devices to Cisco NSO also, Cisco NSO must be configured as a Crosswork provider. When configuring the NSO provider, be sure to set the provider property key to *forward* and the property value to *true*.

Workflow

This is a high-level workflow for onboarding IOS-XR devices using Cisco Crosswork Classic or Secure ZTP.

To onboard IOS-XE devices, or for more detailed information on these options, see the ZTP chapter in the [Cisco Crosswork Infrastructure 4.1 and Applications Administration Guide](#).

- [Step 1. Assemble and upload ZTP assets](#)
- [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 your onboarded devices](#)

Step 1. Assemble and upload ZTP assets

1. Assemble the following assets before you begin:

- (Optional) Software images. For Classic ZTP, you can use Cisco IOS-XR versions 6.6.3, 7.0.1, 7.0.2, 7.0.12, and 7.3.1 or later. For Secure ZTP, use Cisco IOS-XR 7.3.1 or later (except 7.3.2 and 7.4.1).
- Configuration Files: SH, PY, or TXT files. You can specify up to three different configuration files for Secure ZTP.
- Credentials of the devices to be onboarded
- Serial numbers of the devices to be onboarded

For Secure ZTP only, also assemble:

- Owner certificates - your organization's CA-signed end-entity certificates, installed on your devices and binding a public key to your organization.
- Pinned domain certificate - your organization's CA- or self-signed domain certificate, with its public key pinned to your organization's DNS network domain. The PDC helps your devices verify that images and configurations downloaded and applied during ZTP processing come from within your organization.
- Ownership vouchers - Nonceless audit vouchers that verify that devices being onboarded with ZTP are bootstrapping into a domain owned by your organization. Cisco supplies OVs when a request is submitted with your organization's PDC and device serial numbers.

2. If applying software images: Upload the software images. Go to **Device Management > Software Images**.
3. Upload the configuration files. Go to **Device Management > ZTP Configuration Files**.
4. Upload device serial numbers. Go to **Device Management > Serial Number and Voucher** 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 > Serial Number and Voucher**.

Step 2. Create a ZTP profile combining an image file and configuration file

Crosswork uses ZTP profiles to automate imaging and configuration processes. While optional, creating ZTP profiles is recommended as the best way to combine a single image file and configuration file based on a product or device family,

such as the Cisco ASR 9000 or Cisco NCS5500. We recommend that you create only one day-zero ZTP profile for each device family, use case or role the devices serve in the network.

To create ZTP profiles, go to **Device Management > ZTP Profiles**.

Step 3. Prepare ZTP device entries for the devices to be onboarded

Depending on how many devices you are onboarding, you can either prepare and import a CSV file or you can create device entries individually.

1. Go to **Device Management > Devices**.

2. Click the **Zero Touch Devices** tab. Then:

- To create a device entry file for many devices, click the **Import** icon and download the CSV template. Edit the template and add entries for each device you want to onboard. See the ZTP chapter for details on the file entries. Then click the **Import** icon again to import your device entry file.
- To create device entries one at a time, click the **Add** icon.

Step 4. Set up DHCP for Crosswork ZTP

Before triggering ZTP processing, you must update your organization's DHCP server configuration file with the IDs for your ZTP device entries and the paths to the image and configuration files stored in the ZTP repository. This allows Crosswork and DHCP to identify these ZTP devices and to respond correctly to each device's requests for connection to the network, and to download image and configuration files. For sample DHCP entries, see the ZTP chapter.

Step 5. Initiate ZTP processing to onboard the devices


Initiate ZTP processing by rebooting each of the devices to be provisioned: Power-cycle it, or press the chassis reset button.

Step 6. Monitor the ZTP processing status

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

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

Zero Touch Provisioning



[View ZTP devices](#)

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

Step 7. Verify your onboarded devices

Go to **Device Management > Devices**. Click the **Zero Touch Devices** tab. All of your onboard devices should be listed.

You may need to edit the information for some devices. Some of the information needed for a complete device record either is not needed in order to onboard the device, or not directly available through automation. For example, geographical location data defined using a set of GPS coordinates.

ZTP devices, after being onboarded, are automatically part of the shared Crosswork device inventory. You can edit them like any other device.

Visualization of native SR path

Overview

Objective

Visualize the actual path traffic flows physically through the topology map, even if traffic is on a native SR IGP path (not SR-policy) over inter-AS option C.

Challenge

Visualizing the native SR IGP path is often an operational challenge. Without access to a streamlined and simple to use interface, diagnosing and troubleshooting the native path requires you to repeatedly login to network devices without a solution to improve efficiency.

Solution

With the Path Query option, the objective is to visualize the native path using the traceroute SR-MPLS multipath command to get the actual paths between the source and the destination. With Cisco Crosswork Network Controller, a traceroute command runs on the source device for the destination TE-router ID and helps in retrieving the paths. By using native gRPC calls from the Crosswork server, you are able to get the paths from the device which assist in visualizing the native path through which the traffic flows. Since the traceroute command results in an operation that might take time to converge, Cisco Crosswork Network Controller provides an asynchronous user experience where you can send a request for such an operation and then be notified when the output is ready for inspection.

How Does it Work?

- Create a new path query, defining the headend and endpoint devices to find the available Native SR IGP paths.
- Visualize the available Native SR IGP paths as defined by the query on the topology map.
- Inspect the available paths and review the Output, Nexthop, Source, Destination, and Hop Index information.
- Create additional path queries as needed based on service type and instance and visualize the paths on the topology map.
- Troubleshoot any failed path queries.

Usage Scenarios

Scenario 10 –Troubleshooting paths between native SR paths over inter-AS Option C

Scenario Context

Visualization of the path traffic flows is not readily available without manual tasks from different sources. Once attaining traffic flow paths, the data is often out of date. Cisco Crosswork Network Controller supports the creation of Path Queries, which you define within the GUI. This allows visualization of actual SR IGP paths between the source destination on a

topology map. Cisco Crosswork Network Controller provides an asynchronous user experience where the user is notified when results are ready for inspection. This facilitates rapid troubleshooting for issues with native traffic flows.

Assumptions and Prerequisites

- The device should have IOS XR version 7.3.2.
- The device should have gRPC (Remote Procedure Call) enabled. To check, run “show grpc” the in device and follow these steps:
 - For gRPC without a secure connection: If gRPC is showing as not enabled, enable gRPC using the following commands: configure terminal; grpc; no-tls.
 - For gRPC with a secure connection: Upload security certificates to Cisco Crosswork Network Controller in order to connect to the device using the following commands: configure terminal; grpc.
- Cisco Crosswork Optimization Engine server should have the devices imported with gNMI (Network Management Interface) capability and gNMI connectivity for the devices.
 - Devices should have gNMI capability enabled in Cisco Crosswork Network Controller while attaching the device. Go to **Device Management > Network Devices**. Select the device to edit. The Edit Device Details screen appears. From the required Capability list, select **GNMI**. Click **Save**.
 - Devices should have the gNMI connectivity information enabled. Go to **Device Management > Network Devices**. Select the device to edit. On the Edit Device Details screen, under Connectivity Details, click **+ Add Another**. For Protocol, select **GNMI** and add the IP Address / Subnet Mask information. Type the Port information and for Encoding Type, select **JSON**. Click **Save**.
 - Make sure the credential profiles include connectivity information for gNMI. Go to **Device Management > Credential Profiles**. The Credential Profiles screen appears. Select a profile to edit. On the Edit Profile Devices screen, click **+ Add Another**. For Connectivity Type, select **GNMI**. Add the User Name, Password, and Confirm Password information. Click **Save**.

Workflow

1. Select **Services & Traffic Engineering > Path Query**. The Path Query dashboard appears.
2. Click **New Query**. The New Path Query panel appears on the right with the mapped Device Groups panel on the left.
3. Enter the device information in the required fields to find available Native SR IG Paths.
 - a. Select the Headend device from the list. For this example, select **P-Edge-A1**.
 - b. Select the Endpoint device from the list. For this example, select **P-Edge-B2**.
4. Click **Get Paths**. The Running Query ID pop-up appears.
Note: Path queries may take a moment to complete. When the Running Query ID pop-up appears, you can also select **View Past Queries** to return to the Path Query Dashboard. If you already had path queries in the list, you can view existing details as the new query continues to run in the background, which is indicated by the blue Running icon in the Query State column. When the new query state turn green, completed, it can be viewed.
5. Click **View Results** when it becomes available on the Running Query ID pop-up. The Path Details panel appears with corresponding Available Paths details while the defined topology map appears with the available Native SR IG Paths on the left.

6. Click on the Available Paths options (for example, **Path 0** and **Path 1**) to review Status details for Output, Nexthop, Source, Destination, and Hop Index information. When you select one of the available paths, the map will update with the corresponding Device Groups topology mapping of Path 0 and Path 1.

Note: Ensure that the **Show Participating Only** check box is selected in the top-right corner of the map.

Note: There are three likely status outcomes to a path query. The screen captures below are independent examples not directly associated with the scenario's workflow:

- a. **Non-Broken Path (path is complete):** Path Status shows as **Found** with path hop details and overlay shown.

The screenshot shows a map of Japan with several router nodes (R1 through R9) highlighted. A blue line connects R2 to R9, representing a found path. The 'Path Details' panel on the right shows the following configuration:

- Select Service: L3vpn-Service
- Headend: R9 (9.9.9.9)
- Endpoint: R2 (2.2.2.2)

The 'Available Paths' section displays two paths:

- Path 0:** Status: Found, Output: Bundle-Ether79, Nexthop: 10.7.9.1, Source: 9.9.9.9, Destination: 127.0.0.2. Hop details show four hops from R2 to R9.
- Path 1:** Status: Found, Output: Bundle-Ether79, Nexthop: 10.7.9.1, Source: 9.9.9.9, Destination: 127.0.0.2. Hop details show four hops from R2 to R9.

- b. **Broken Path (Path is complete):** Path Status shows as **Broken** with path hop details and overlay shown.


The screenshot shows a map of Japan with router nodes R1, R3, R5, and R7 highlighted. A blue line connects R1 to R7, representing a broken path. The 'Path Details' panel on the right shows the following configuration:

- Select Service: L3vpn-Service
- Headend: R1 (1.1.1.1)
- Endpoint: R7 (7.7.7.7)

The 'Available Paths' section displays two paths:

- Path 0:** Status: Broken, Output: Bundle-Ether131, Nexthop: 10.1.3.2, Source: 1.1.1.1, Destination: 127.0.0.0. Hop details show three hops from R1 to R7.
- Path 1:** Status: Broken, Output: Bundle-Ether131, Nexthop: 10.1.3.2, Source: 1.1.1.1, Destination: 127.0.0.0. Hop details show three hops from R1 to R7.

- c. **Broken Path (Path is not complete):** Path Status shows as **Broken** with path hop details partially shown (depending on gNMI output for traceroute – see Step 18 for troubleshooting details) and overlay details partially shown. An Error message will appear indicating that the devices and links are not available.



7. Select **Services & Traffic Engineering > Path Query** to return to the Path Query Dashboard.
8. Ensure that the new path Query State column shows as completed with a green icon. The new path in the table will also show a Query ID link, both the corresponding Headend and destination Endpoint, and the Available Paths column will show 2 for both paths.
- If a query state is broken, see the last step in the workflow for troubleshooting details.
9. As needed, click on the **Query ID** link or click ... and select **View Details** to again review the Path Details panel and map.
10. Create additional path queries by selecting **Services & Traffic Engineering > Path Query**. The Path Query Dashboard appears where the previous path queries are listed by Query ID.
Note: Make sure to set the **Automatically delete query older than every < X >** option within the number of hours needed from the Path Query Dashboard. The maximum number of hours provided is **24**.
11. Click **New Query**. The New Path Query panel appears on the right with the mapped Device Groups panel on the left.
12. For Select Service, select the Type from the list. In this example, select **L2VPN-SERVICE**.
By utilizing Select Service, when you later select the Headend and Endpoint, the options are conveniently identified according to the relevant VPN service type.
13. For Select Service, select the Instance from the list. In this example, select **L2VPN_NM_P2P-NATIVE-210**.
The topology map will update to show the path between both servers. In this example, **P-Edge-B2** and **P-Edge-C3** are isolated on the map showing the logical path.

14. Select the following from the list:
- Headend: **P-Edge-B2**.
 - Endpoint: **P-Edge-C3**.
15. Click **Get Paths**.
The Running Query ID pop-up appears.
16. Click **View Results** when it becomes available. The Path Details panel appears with the corresponding Available Paths details, while the defined topology map appears with the available Native SR IG Paths on the left. This view shows the actual, physical hops between B2 and C3 that is carrying the traffic.
17. To troubleshoot any Failed path queries appearing in the Path Query Dashboard's Query State column, select the "I" icon for error details.
In this example, the gNMI protocol is missing from the Connectivity Details for a previous path query with the Headend P-BOTTOMLEFT device and the Endpoint P-BOTTOMRIGHT devices. To troubleshoot the failed path query, do the following:
- Select Device Management > Network Devices.
 - Find the device by Host Name and select the check box.
 - Click the Edit icon at the top of the table. The Edit Device Details pop-up appears.
 - In this example, the Connectivity Details for Protocol is missing gNMI. Click + Add Another and type GNMI until it appears in the list. Select it.
 - Enter the IP Address / Subnet Mask information and Port field information.
 - Enter the Timeout field as 30.
 - In the Encoding Type list, type JSON until it appears in the list. Select it and click Save.
 - Select Services & Traffic Engineering > Path Query. The Path Query Dashboard appears.
 - Click New Query. The New Path Query panel appears.
 - Select the following the list:
 - Headend device: **P-BOTTOMLEFT**.
 - Endpoint device: **P-BOTTOMRIGHT**.
 - Click **Get Paths**.
The Running Query ID pop-up appears.
 - Click **View Results** when it becomes available. The Path Details panel appears with corresponding Available Paths details, while the defined topology map appears with the available Native SR IG Paths on the left and is now in a Completed state.

Appendix

Initializing Heuristic Packages to monitor the health of a service

Objective

Enabling Service Health and using system designed Heuristic Packages to monitor the newly created service or exporting them to your system to make adjustments before importing them back in Cisco Crosswork Network Controller, allows for customization on ongoing, detailed monitoring of your service's health.

Note: Service Health is not generally available yet. At this stage, it is available for pre-launch laboratory evaluation only. Engage your account team if you are interested in participating in the evaluation.

Workflow

Select either a system or custom Heuristic Package for ongoing, specialized Service Health monitoring of your new VPN service.

Initialize a Heuristic Package to monitor health of the new service.

1. Go to Administration > Heuristic Packages. The Heuristic Packages screen opens with System and Custom tabs. By default, a system defined Heuristic Package is used.
2. From the System tab, you can preview the package detail Rules, Configuration Profiles, Sub-Services, and Metrics by expanding each section for more information and hover your mouse over the information "I" icon for finer details.
3. You can click Export to download a System defined package to your system to make changes to the .json (?) files before importing them to CNC as a customized package.
4. If you exported a system file for customization, or you have custom packages on your system you want to import, click Import.
5. The Import Heuristic Packages screen opens and click Browse to find the name of your custom package on your system.
6. Select your custom package and click Import. Note: Your system performance might be impacted during heuristic package import due to high server resource consumption.
7. From the Import Heuristic Packages screen, click Preview to review the details of the package to be imported. Further information on the package's Rules, Configuration Profiles, Sub-Services, and Metrics appears.
8. Select each option to preview the details of the custom package. CNC will provide information on the details and if any details need to be updated before CNC will accept the new custom package and allowing it to be imported.
9. After importing the custom package, select it so the new rules and configuration details begin to monitor the ongoing health of your designated services.