

Crosswork Network Controller Circuit Style SR-TE

Contents

Introduction	3
Augmentation to Segment Routing (SR) Architecture	4
Cisco Solution Components and Platform Dependencies	5
Workflow - Day 0, Deployment, DayN	6
Summary	14
Appendix: CSM Configuration Panel	14

Introduction

Some of the challenges facing service providers include SONET/SDH equipment going end-of-life, with spares becoming very expensive to procure. This has increased overall operational costs to provide such services. IP networks are traditionally more cost-effective at higher data rates but are often perceived to be less reliable and more complex to operate. Both greenfield and brownfield operators have been looking for a technology with the same reliability as SONET/SDH but at a lower operating cost.

Segment Routing allows an operator to build a single network to accommodate different service types. IP-centric services are optimized for equal-cost multipath (ECMP), shared bandwidth consumption, and link protection, while circuit-centric services are optimized for a single path with bandwidth reservation and path protection. Circuit-style Segment Routing Traffic-Engineering (CS SR-TE) based on Segment Routing provides a natural evolution as it is simpler and scales better. Circuit-style SR-TE does not require the maintenance of network states at intermediate routers, nor does it require complex multiprotocols to be implemented.

Traditional Time Division Multiplexed (TDM) transport or circuit-switched services may be delivered using circuit-style SR-TE services. Dedicated circuit-type services require a reserved bandwidth per path and path protection to ensure reliability and no impact on service-level agreements (SLAs) due to changing network loads.

The following are some properties of circuit-style services:

- **Co-routed bidirectional** - Both forward and return paths are routed over the same path between a pair of ingress and egress nodes.
- **Persistence** - The working and protect paths, once calculated, are persistent throughout the lifecycle of the circuit-style SR-TE service. Path re-optimization is disabled and is only performed when explicitly requested by the user.
- **Guaranteed Bandwidth and Latency** - Strict bandwidth and latency requirements are essential to delivering circuit-type services traditionally provided over SONET/SDH networks. It is also essential for the delivery of critical services where the reserved bandwidth is always available for use.
- **End-to-end path protection and restoration** - End-to-end path protection and restoration ensures that bandwidth constraints can always be met, even in the event of link failure.
- **TDM, Circuit Emulation (CE), Private Line Emulation (PLE) Point-to-Point Services** - These point-to-point, guaranteed bandwidth services are traditionally delivered over a SONET/SDH network and can be delivered over circuit-style SR-TE services.

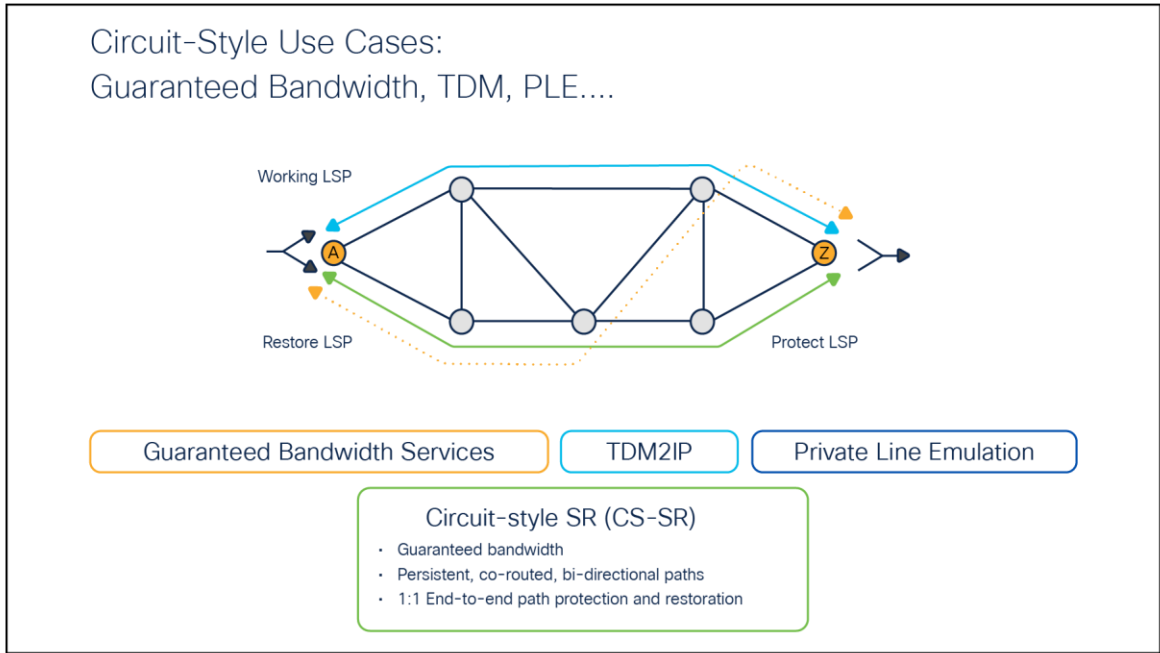


Figure 1.
Circuit Style use cases

Augmentation to Segment Routing (SR) Architecture

Segment Routing provides complete control over the forwarding paths by combining simple network instructions. It does not require any additional protocol. Instead, it removes unnecessary protocols simplifying your network.

- Segment routing networks embrace ECMP for efficient network utilization.
- Local TI-LFA protection (<50 msec) is used to recover from failures.
- (Temporary) congestion is dealt by
 - Local congestion mitigation approaches (Cisco® COE LCM)
 - Tactical (global) traffic optimization

Table 1. Comparison between IP vs Optical-like solution

IP solution	Optical-like solution
<ul style="list-style-type: none"> • ECMP and local protection (LFA FRR) • Bandwidth optimization • Soft bandwidth guarantee 	<ul style="list-style-type: none"> • No ECMP and path protection (Working, Protect, and Restore) • Hard bandwidth guarantee • Low bandwidth optimization

Cisco Solution Components and Platform Dependencies

- **Circuit-Style Manager (CSM)** – The **Circuit-Style SR-TE Feature Pack**, also known as **SR Circuit-Style Manager (CSM)**, provides a bandwidth-aware Path Computation Element (PCE) for computing circuit-style SR-TE policy paths. Circuit-style SR-TE policies comprise bidirectional, co-routed, and path-protected SR policies for carrying high-priority or critical services traffic over packet-based networks requiring committed bandwidth with protected paths.

CSM performs centralized bookkeeping of bandwidth resources across the network, and path computation of the circuit-style SR-TE policy paths. CSM ensures that the computed paths meet the requested strict bandwidth requirements while concurrently adhering to user-specified constraints such as disjointness and latency minimization. CSM also allows the user to monitor the bandwidth resource levels and to identify where bandwidth resources are running low.

- **Segment Routing Path Computation Element (SR-PCE)** – SR-PCEs running on IOS® XR provide stateful PCE functionality for the network. SR-PCEs perform topology and LSP collection, allowing the user to gain visibility and insights to the network in real time. Path Computation Clients (PCCs) can report and delegate control of head-end Label Switched Paths (LSPs) sourced from the PCC to a PCE peer. The PCE can request the PCC to update and modify the parameters of the LSPs it controls. In the context of circuit-style SR-TE, the PCCs delegate the path computation of CS SR-TE policies to the SR-PCE. These paths are subdelegated to the Circuit-Style Manager (CSM) for path computation.
- **XR platform/headend device** – Comprises routing platforms such as the ASR9000 and NCS5500 devices that support the **circuit-style SR-TE and SR Policy Liveness monitoring** features. Circuit-style SR-TE policies are PCC configured/initiated, and the bidirectional active, protect, and restore paths are delegated to the SR-PCE/CSM for path computation. Once the required paths are returned from the controller, the headend will encode the Adjacency-SIDs for the circuit-style SR-TE policies and monitor the bandwidth guarantee for SLA. In addition, the headend devices will perform liveness monitoring of the CS SR-TE policies and perform a switchover to a protect or restore path where required.
- **NSO:** Provisioning of various services using service packages for VPN – L2VPN, L3VPN and Traffic Engineering Tunnels – SR-TE, SRv6-TE, RSVP-TE, and CS-SR. NSO service packages use Network Element Drivers to provision device-level config for various vendors. A rich ecosystem of NEDs for Cisco and other vendors can be augmented to perform multivendor service provisioning.

Workflow – Day 0, Deployment, DayN

Day 0 - Requisites and Dependencies

To support the use of CS SR-TE policies, the following requisites need to be met.

- Segment Routing network infrastructure
 - SR-PCE running IOS-XR 7.9.1 or later
 - Headend devices/PCCs running IOS-XR 7.8.1 or later, with SR Policy Liveness monitoring
 - Manual/statically allocated Adjacency-SIDs
 - Network partitioning via MPLS EXP allocation, QoS policies
- Crosswork Optimization Engine/Crosswork Network Controller 5.0 or later

To support circuit-style services and guarantee end-to-end bandwidth, there are requisites to be met. Besides version dependencies on both SR-PCE, headend devices, and Cisco Crosswork®, manually allocated Adjacency-SIDs are desired. This allows for Adjacency-SID persistence over reloads and restarts. The network needs to be partitioned for IP and circuit-style services. There are a few methods to partition the network, and one of the simpler ones is to allocate one class of service for the circuit-style partition. For example, allocate MPLS EXP7 and keep MPLS EXP0 to EXP6 for the IP partition, and provision a reservable bandwidth for all links. This bandwidth will be the maximum reservable bandwidth that will be available for circuit-style services. Together with the per-link Quality of Service (QoS) configuration, this ensures bandwidth availability on the data plane and will protect circuit-style traffic from congestion.

Day 1 - Deployment

a. Provisioning circuit-style SR-TE policies

At this point in time, Crosswork has collected the topology and onboarded devices and has maximum reservable bandwidth defined on each link.

NSO CS-SR TE Provisioning: Circuit-style SR policies can be provisioned from CNC using CNC-NSO Function Packs. When submitting the configuration, NSO will perform provisioning of the circuit-style SR-TE policy at both headend devices. Intent is abstracted to deploy relevant configuration on devices using CLI- or NETCONF-based south-bound connectivity to the devices.

The following screenshot shows the options available under the **Crosswork - Services and Traffic Engineering / Provisioning (NSO) /SR-TE/ Circuit-Style Policy Panel**.

Circuit-style SR-TE policy may be provisioned with the intent below:

- **Auto-color** – Automatically picks color from the resource pool to be allocated per headend or assigned manually.
- **Bandwidth** – Requested bandwidth for circuit-style SR policy in below snapshot shows policy bandwidth requested for 1 Mbps.
- **Disjoint Path** – Compute a path that is disjoint from another path in the same disjoint group. Applicable for working and protect paths.
- **Working, Protect, and Restore Path** – Provide dynamic path metric and constraints for “W”, “P,” and “R” path. For “P” and “R” path, enable revertive and wait to revert timer.

Provisioning state will show success or failure based on response from NED and service.

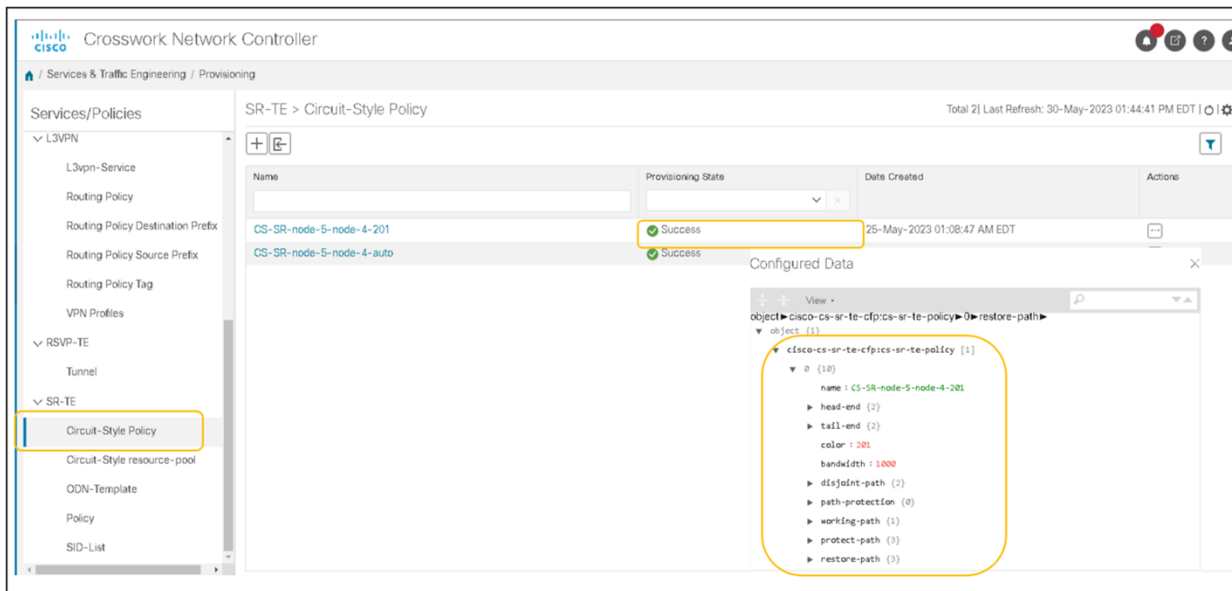


Figure 2.
Circuit-style policy panel

CNC/COE Circuit-Style SR-TE Visualization

Once the circuit-style SR-TE policy is provisioned on the headend devices, it will trigger a PCEP request for circuit-style SR-TE path computation to the SR-PCE. This request will be subdelegated to the Circuit-Style Manager (CSM) on Crosswork Optimization Engine. CSM will compute the bidirectional co-routed working and protect paths and reply both headend devices with SID lists comprised of Adjacency-SIDs. The two headend devices will install the circuit-style SR-TE policy for the working and protect paths with the SID lists.

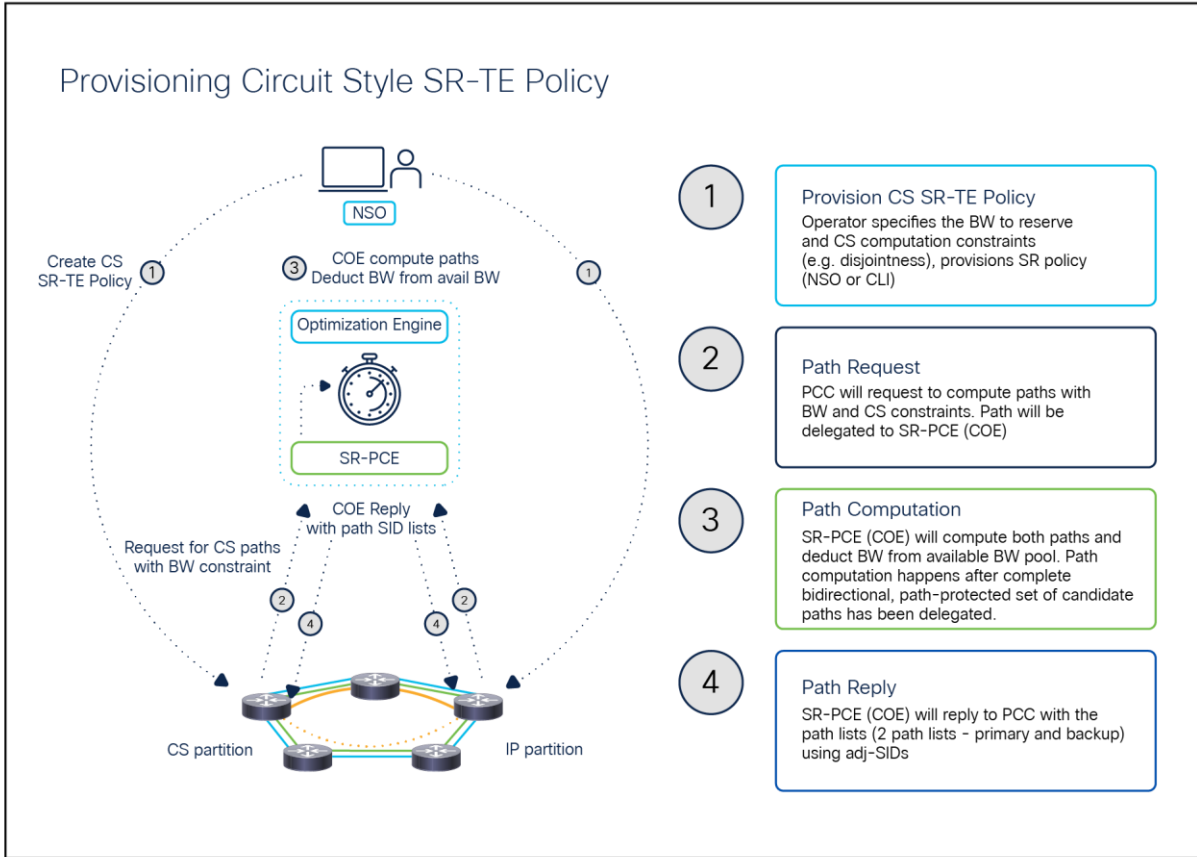


Figure 3. Provisioning circuit-style SR-TE policy workflow

The following screenshot shows visualization of circuit-style SR-TE policies available under **Crosswork - Services and Traffic Engineering / Traffic Engineering**.

Here you can see two paths computed as **Working Path** and **Protect Path** from each headend that depicts bidirectional and co-routed.

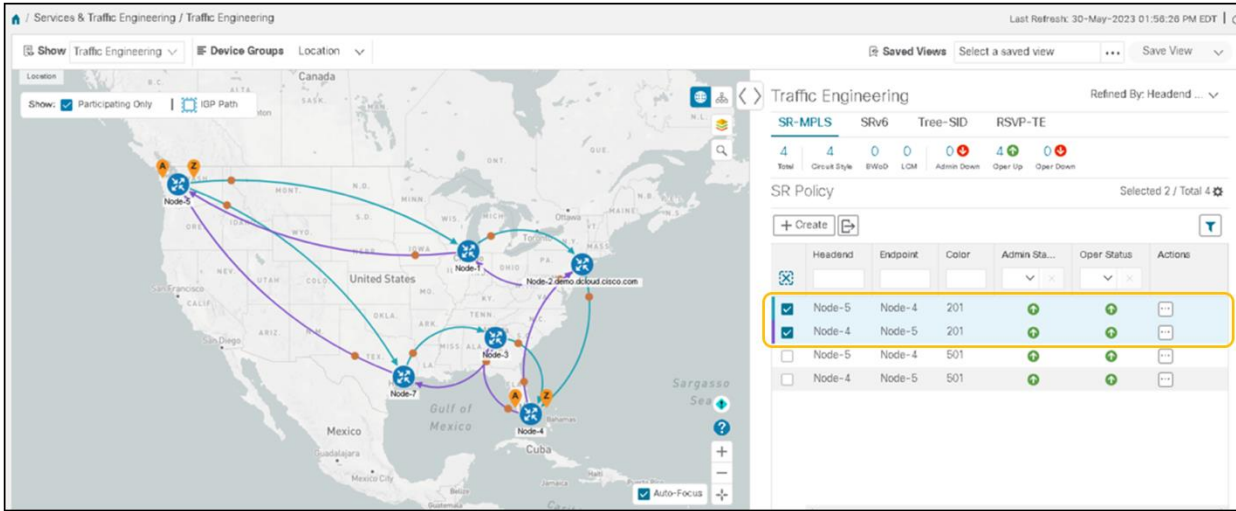


Figure 4. Visualization of circuit-style SR-TE polices on traffic engineering pane

On the same screen, we can visualize the circuit-style SR-TE policy with IGP paths from both headend by checking the IGP Path checkbox and enabling the metric display for TE, IGP, and Delay under the layers icon button.

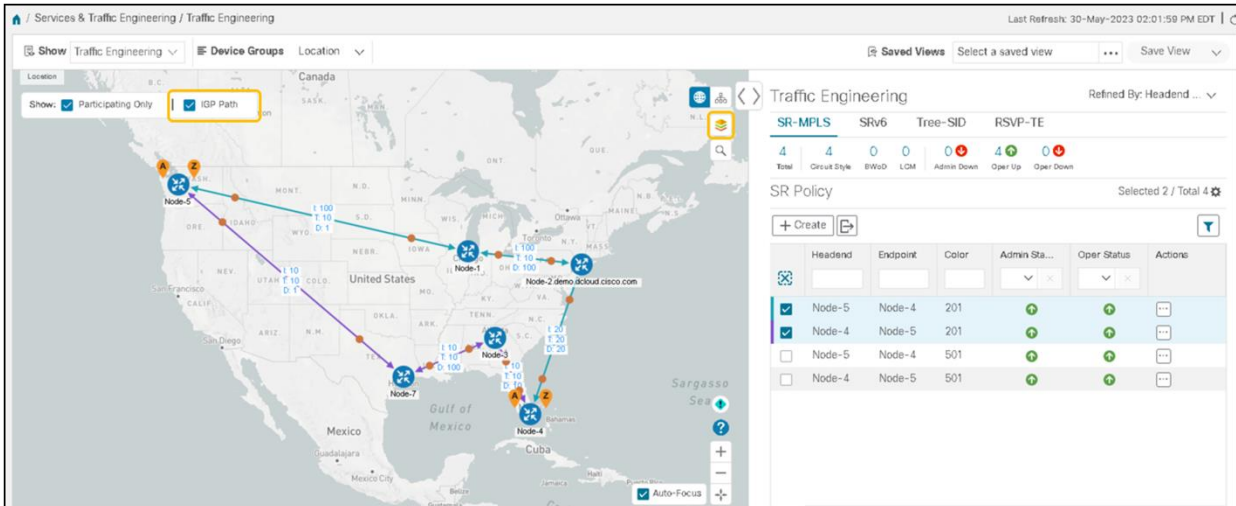


Figure 5. Visualization of circuit-style SR-TE policies with link metrics for working and protect paths

We can view candidate paths computed by CSM that are operational UP, and only the working path is active. In the same view, selecting the Actions tab shows Candidate Path 100 (W), which is Up and (A)ctive, and 50(P), which is Up.



Figure 6. Visualization of circuit-style SR-TE policy details with candidate path

By selecting the chevron (>), we provide an expanded view for candidate path preference 100. In the view shown below, per candidate path-based Adjacency-SID, Node, and respective interfaces are listed, showcasing metric type with requested and reserved bandwidth.

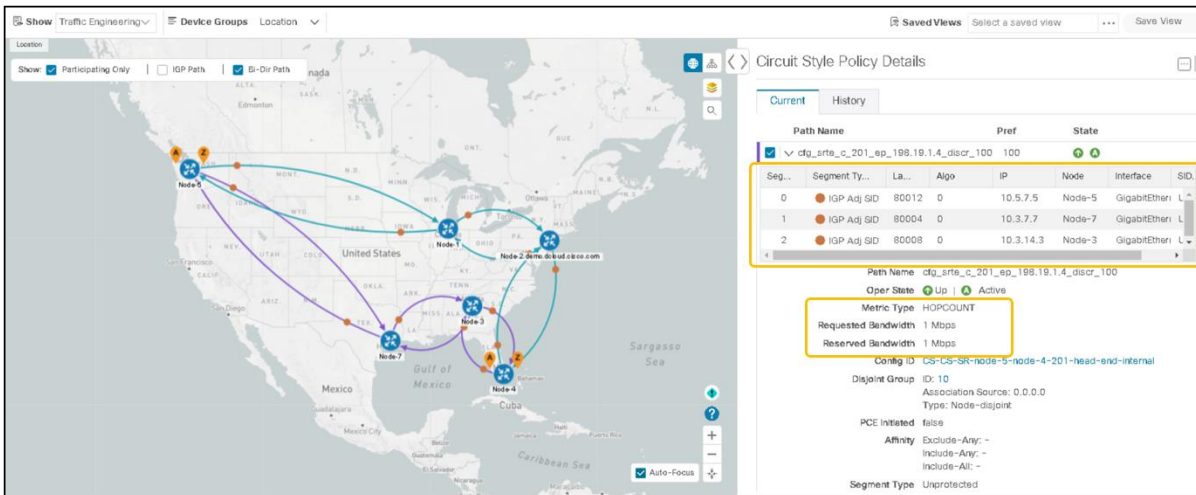


Figure 7. Visualization of circuit-style SR-TE policy candidate path SID list details

Bandwidth is reserved on links for the paths computed by CSM for this policy. This can be seen under the **Topology Tab-> dropdown - Traffic Engineering**. Selecting each link and navigating under Traffic Engineering, you can see the used bandwidth is marked as 1 Mbps and available is shown as 4 Mbps. Similarly this can be seen for all the links consumed by the W and P paths.

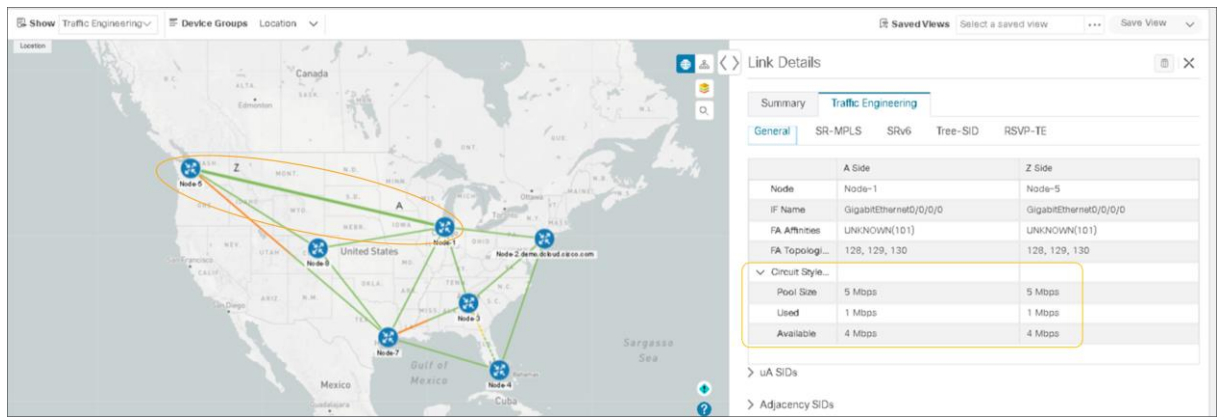


Figure 8. Visualization of circuit-style SR-TE per link bandwidth pool, utilization, and availability

Restore Path

At this state, R-restore path is not computed and no bandwidth is reserved. CSM will compute the “R” path only when both “W” and “P” paths are down (i.e., 1+1:R scenario).

In this case, by shutting down links on both the headends for “W” and “P” paths, CSM computes the “R” path, this can be seen in the same view as the list of candidate paths:

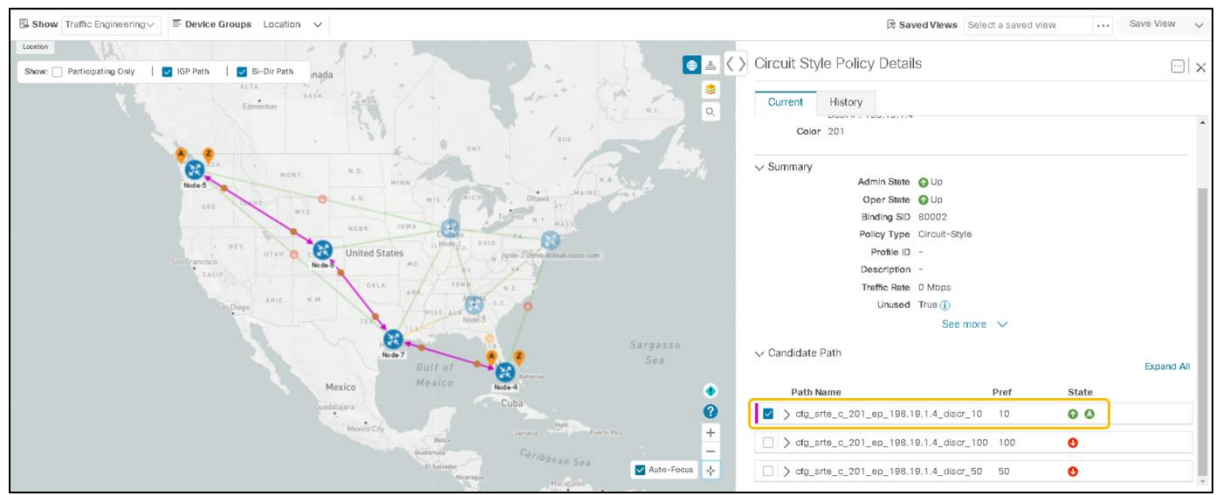


Figure 9. Visualization of circuit-style SR-TE policy restore path

Expanding the candidate path for “R” preference:10 shows nodes, the Adjacency-SID list, and respective interfaces with metric, requested, and reserved bandwidth.

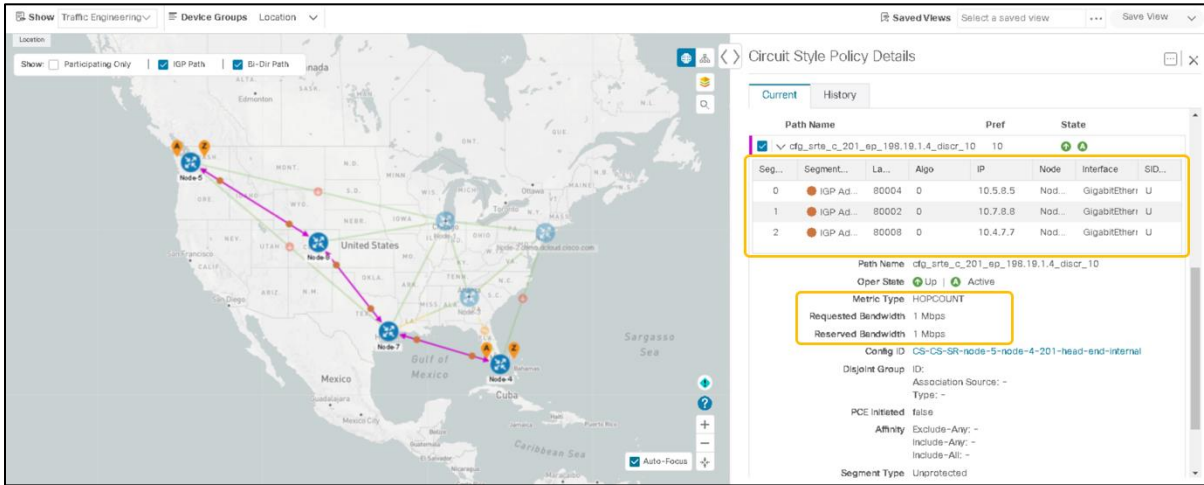


Figure 10. Visualization of circuit-style SR-TE restore path SID list details

CNC alarms view will notify the user that for this policy W and P paths are down. The R path is computed and active. This can be seen under the CNC/Administration/Alarms view by filtering on the source circuit-style SR-TE policy.

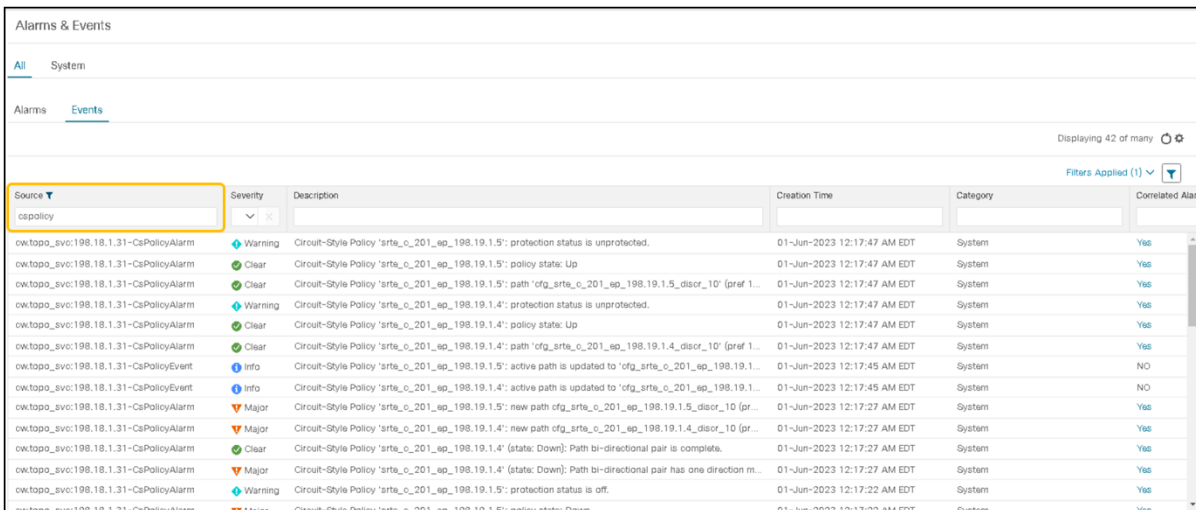


Figure 11. Visualization of circuit-style SR-TE alarms and events

Day N - Operations

a. Monitoring CS SR-TE per-device circuit-style bandwidth pool

The following screenshot shows per-link utilization of the circuit-style bandwidth pool under **Topology Tab-> dropdown - Traffic Engineering**. Selecting each link and navigating under Traffic Engineering, we can see as below:

- Pool Size: 50% of available link
- Used BW: Various circuit-style SR policies using this link
- Available BW: Future CS-SR policies that will be used by CSM to compute bandwidth

Similarly, we can see circuit-style bandwidth pool utilization for every link.



Figure 12.

Visualization of circuit-style SR-TE per-link bandwidth resource pool size, utilization, and availability

b. Maintenance

- Operating circuit-style SR-TE policy. During the lifecycle of the circuit-style SR-TE policy, network changes may have occurred, and this could mean that there could be newer and better paths available. Circuit-style SR-TE policies will be manually optimized by the operator using the Crosswork Optimization Engine (COE) user interface, or via APIs. When reoptimization is triggered, COE will recompute the path, perform bandwidth bookkeeping, and send the new path SID list to both headend nodes.
- Link and node maintenance. The operator may need to perform link and node maintenance. To do so, the operator will specify the link or node to exclude in the next computation and trigger a new path computation, which will send the updates to both headend devices. This could be performed using link affinities, or other mechanisms as the solution evolves. After the maintenance is complete, the operator may request that the previous path be used.
- Circuit-style SR-TE policy resizing. Due to changing bandwidth requirements, the operator may elect to resize the existing circuit-style SR-TE policy by submitting the changes. This will trigger a provisioning change on both headends and trigger new PCEP requests. CSM will recompute the policy resize, reserve bandwidth, and install newer path lists if required.

Summary

The circuit-style SR-TE solution is designed to meet the requirements of a connection-oriented transport network service, which was traditionally delivered over circuit-switched SONET/SDH networks. Circuit-style SR-TE policies allow a common network infrastructure to be used for both connection-oriented and connectionless IP-based services, which eliminates the need to run multiple parallel transport networks. This reduces both the operator's capital and operating expenditures.

Appendix: CSM Configuration Panel

The following screenshot shows the options available under **Crosswork - Services and Traffic Engineering / Circuit-Style SR-TE / Basic Configuration Panel**.

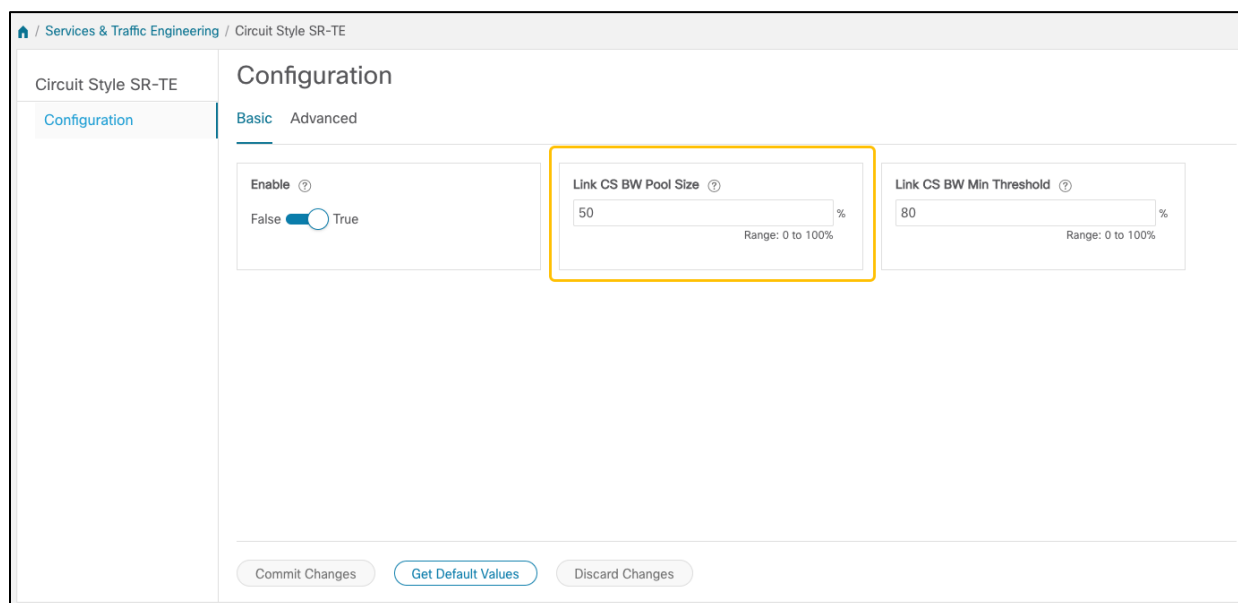


Figure 13.
Circuit Style SR-TE Basic Configuration Panel

- a. **Enable** - This toggle enables/disables the Circuit-Style SR-TE Function Pack (Circuit-Style Manager) on Crosswork.
- b. **Link CW BW Pool Size** - This global parameter specifies the percentage of each link's bandwidth that may be allocated to the Link CS BW Pool.
- c. **Link CS BW Min Threshold** - This global parameter specifies the Link CS BW Pool utilization threshold percentage beyond which a threshold crossing event notification will be generated.

The following screenshot shows the options available under **Crosswork - Services and Traffic Engineering / Circuit-Style SR-TE / Advanced Configuration Panel**.

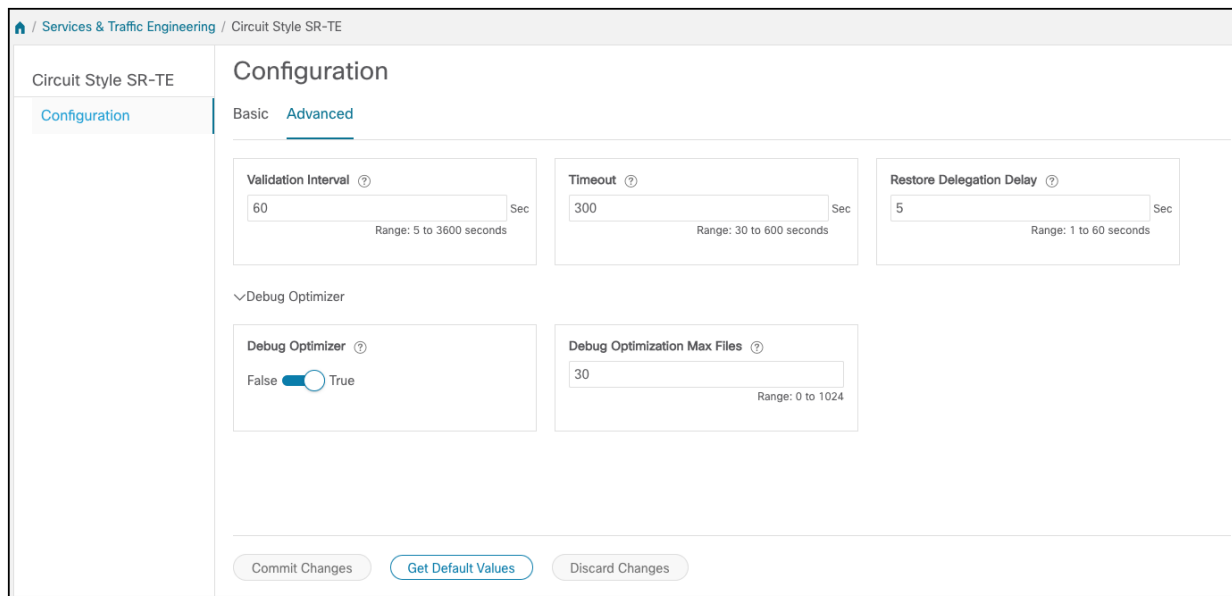


Figure 14.
Circuit Style SR-TE Advanced Configuration Panel

- a. **Validation Interval** – This specifies the duration that CSM will wait before the bandwidth that is reserved for an undelegated policy is returned to the circuit-style SR-TE policy bandwidth pool.
- b. **Timeout** – This specifies the duration that CSM will wait for the delegation request to generate a notification.
- c. **Restore Delegation Delay** – This specifies the duration that CSM will pause before processing a restore path delegation.
- d. **Debug Optimizer** – This toggle enables/disables the CSM Debug Optimizer. This should only be enabled for debugging purposes.
- e. **Debug Optimization Max Files** – This specifies the maximum number of debug files that will be generated by the debug optimizer.

Contributing authors

Fung Lim, Technical Marketing Engineer

Sujay Murthy, Technical Marketing Engineer

Eric Ortheau, Engineering Product Manager

Max Williams, Principal Engineer

Jose Liste, Distinguished Technical Marketing Engineer

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)