



Get Started

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

- [Segment Routing, on page 1](#)
- [Inventory Management Concepts, on page 3](#)
- [Before You Begin, on page 4](#)
- [High-Level Workflows, on page 5](#)

Segment Routing

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

Segments

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

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

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

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

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

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

Segment Routing for Traffic Engineering

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

Segment Routing Policies

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



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

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

Dynamic SR Policy

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

Explicit SR Policy

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

Disjointness

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

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

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

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

**Note**

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

Inventory Management Concepts


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

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

Before You Begin

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

- **User Accounts** : Cisco recommends as a best practice that you create separate accounts for all of your users, so that there is an audit record of user activity on the system. Prepare a list of the people who will use Cisco Crosswork Optimization Engine. Decide on their user names and preliminary passwords, and create user profiles for them (see [Manage Users](#)).
- **User Roles**: Cisco recommends that you use role-based access control to confine users to just the software functions needed to perform their job duties. By default, every new user you create has full administrative privileges. Unless you want to extend the same privileges to every user, you will need to plan a system of user roles, create them, and assign them to the user profiles you create (see [Create User Roles](#)).
- **Credentials**: Gather access credentials and supported protocols that you will use to monitor and manage your devices. For providers, this always includes user IDs, passwords, and connection protocols. For devices, it includes user IDs, passwords, and additional data such as the SNMP v2 read and write community strings, and SNMPv3 auth and privilege types. You will use these to create credential profiles (see [Inventory Management Concepts, on page 3](#) and [Manage Credential Profiles](#)).
- **Tags**: Plan a preliminary list of custom tags to create when setting up the system, so that you can use them to group your devices when you first onboard them. As explained in [Inventory Management Concepts, on page 3](#), you will want to consider grouping devices by functionality. You need not have a complete list of tags at first, as you can always add more later, but please note that all the tags you do plan to use must be in place before you need them; you cannot create them "on the fly" (see [Manage Tags](#) and [Create Tags](#)).
- **Providers**: As explained in [Inventory Management Concepts, on page 3](#), providers do the basic work of direct interaction with network devices, so that Cisco Crosswork Optimization Engine can automate monitoring and responses to network events. At a minimum, Cisco Crosswork Optimization Engine must have an SR-PCE provider defined in order to discover devices and to distribute policy configuration to devices. You should determine the auto-onboarding mode and device profile you will use (if you auto-onboard devices). See [Add Cisco SR-PCE Providers](#).
- **Devices**: Decide how you are going to onboard your devices: manually, via the user interface, or automatically via synchronization or CSV import. This determines the amount of additional information you will need to onboard your devices, which is covered in [About Adding Devices](#).
- **External Data Destination(s)**: Decide which external data destination (Kafka or gRPC) you are going to use and ensure it is set up to receive input from Cisco Crosswork Data Gateway.

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

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

High-Level Workflows

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

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







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

Workflow: Auto-Onboard Devices

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

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





Step	For more information, see...
1. Ensure that your devices are configured properly for communication and telemetry.	Refer to the guidelines and sample configurations in: <ul style="list-style-type: none"> • Prerequisites for Onboarding Devices • Sample Configuration for Devices in Cisco NSO <p>Note Only if NSO is being used for device management.</p> <ul style="list-style-type: none"> • Prerequisites for Device Model Driven Telemetry
2. Create a device credential profile.	Create Credential Profiles

Step	For more information, see...
<p>3. Configure SR-PCE as a provider.</p> <p>Note The auto-onboard provider property value must be set to managed or unmanaged to enable automatic onboarding of devices. For more information see About Adding Devices and Auto-Onboard Property Descriptions.</p>	Add Cisco SR-PCE Providers
4. Validate communications with provider.	Get Provider Details
5. (Required if using NSO for device management) Configure NSO credential profile and provider.	<ul style="list-style-type: none"> • Create Credential Profiles • Add Cisco NSO Providers
<p>6. (Required if using NSO for device management, otherwise optional) To update device attributes (such as mapping a device to NSO, adding connectivity IP and geographical coordinates, and so on) export the CSV device list, make and save modifications, and import it back to the device inventory.</p> <p>Note If you wish to use the geographical topology map, you must add geographical location details.</p>	<ul style="list-style-type: none"> • Export Devices • Import Devices
7. (Optional) Create tags for use in grouping new devices.	Manage Tags
8. Add a Cisco Crosswork Data Gateway instance. At least one instance should have been added and enrolled during installation.	See the Cisco Crosswork Optimization Engine Installation Guide .
9. Confirm that the Cisco Crosswork Data Gateway instance is running. The operational state of the Cisco Crosswork Data Gateway instance must be Up before continuing to the next step.	Manage Cisco Crosswork Data Gateway Instances
10. Attach devices to Cisco Crosswork Data Gateway so that data can be collected.	Attach a Device to a Cisco Crosswork Data Gateway Instance
<p>11. View device list (Inventory Management > Devices) to check that devices have been added properly.</p> <p>If devices are unreachable, select and edit the device with connectivity details.</p>	Manage Network Devices Click  to investigate any device whose Reachability State is marked as  (unreachable),  (degraded), or  (unknown).
12. Confirm visualization of IGP topology (logical view).	Network Topology Map
13. Visualize discovered SR policies and RSVP-TE tunnels and create new tunnels.	Visualize SR Policies and RSVP-TE Tunnels

Workflow: Manually Import Devices

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

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

Step	For more information, see...
1. Ensure that your devices are configured properly for communication and telemetry.	Refer to the guidelines and sample configurations in: <ul style="list-style-type: none"> • Prerequisites for Onboarding Devices • Sample Configuration for Devices in Cisco NSO <p>Note Only if NSO is being used for device management.</p> <ul style="list-style-type: none"> • Prerequisites for Device Model Driven Telemetry
2. Create a device credential profile.	Create Credential Profiles
3. Configure the SR-PCE provider. Note Set auto-onboard property value to off for manual device onboarding. For more information see Auto-Onboard Property Descriptions .	Add Cisco SR-PCE Providers
4. (Required if using NSO for device management) Configure NSO credential profile and provider.	<ul style="list-style-type: none"> • Create Credential Profiles • Add Cisco NSO Providers
5. (Optional) Create tags for use in grouping new devices.	Manage Tags
6. Create a CSV file and import devices.	Import Devices
7. (Optional) Modify device details.	Edit Devices
8. Add a Cisco Crosswork Data Gateway instance. At least one instance should have been added and enrolled during installation.	See the Cisco Crosswork Optimization Engine Installation Guide .
9. Confirm that the Cisco Crosswork Data Gateway instance is running. The operational state of the Cisco Crosswork Data Gateway instance must be Up before continuing to the next step.	Manage Cisco Crosswork Data Gateway Instances
10. Attach devices to Cisco Crosswork Data Gateway so that data can be collected.	Attach a Device to a Cisco Crosswork Data Gateway Instance
11. View device list (Inventory Management > Devices) to check that devices have been added properly. If devices are unreachable, select and edit the device with connectivity details.	Manage Network Devices Click  to investigate any device whose Reachability State is marked as  (unreachable),  (degraded), or  (unknown).

Step	For more information, see...
12. Confirm visualization of IGP topology (logical view).	Network Topology Map
13. Visualize discovered SR policies and RSVP-TE tunnels and create new tunnels.	Visualize SR Policies and RSVP-TE Tunnels