# Cisco Crosswork Planning Design 7.2.x User Guide

**Cisco Crosswork Planning 7.2**
Release: 7.2.x | Updated August 14, 2026

© 2026 Cisco Systems, Inc. All rights reserved.

# Full Cisco Trademarks with Software License

Full Cisco Trademarks with Software License

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

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

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

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

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

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

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

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

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

# Topics included

**1 Get Started.....15**
Core capabilities of Cisco Crosswork Planning .....16
System overview.....16
Types of Cisco Crosswork Planning Design use cases.....17
Core concepts in Cisco Crosswork Planning.....18
Log in to Cisco Crosswork Planning .....19
Dashboards.....20
Customize dashlets.....21
Custom dashboards.....22
Reset the General dashboard to the default view.....24
Copy dashboards.....24
Design engines.....25
Design engine instances resource management.....25
Create design engine instances.....27
Select design engine profiles.....28
Reset engine sessions.....28
Delete design engine instances.....29
Recreate design engine instances.....30

**Part I: Visualize Network Models.....31**

**2 Visualize Network Models.....33**
Create or import network models.....34
Import plan files from the local machine.....34
Import plan files from archive.....34
Create new network models manually.....35
Open plan files.....36
Delete plan files.....37
Download plan files.....38
Key features of the Network Models page.....38
Key features of the Network Design page.....41
Plan tables.....44
Plan table columns.....44
Table schema and plan tables.....45
Assign sites to nodes.....46
Site assignment rules.....46
Run Assign sites to nodes initializer.....48
Assign geographic locations to sites.....49

Set site location....50
Network topology....51
Network plot views....51
Site plots....52
Graphical network topology views....53
Grouping of nodes....53
Plan object representations....54
Set objects to the foreground in network plot....54
Traffic utilization colors....55
Example: Determining traffic utilization....56
Identify the most highly utilized interface....57
Utilization color profiles....57
Network summary tables....60
Work with tables and object selections....61
Show or hide tables or columns....62
Sort columns in tables....62
Filter options in network summary tables....63
Manage table views....64
User-defined columns....64
Enable dark mode....66

**3 Understand Plan Objects....67**
Plan objects....68
Nodes and sites....68
Parent sites and contained objects....69
PSN nodes....69
Node configuration fields....70
Site configuration fields....70
Merge nodes....71
Circuits and interfaces....72
Circuit and interface configuration fields....72
Merge circuits....73
SRLGs....73
Create SRLGs....73
Create SRLGs for circuits only....74
Ports, port circuits, and LAGs....75
Port configuration fields....76
Port circuit configuration requirements....76
Create LAGs....77
Set LAG simulation properties....77
Key operations you can perform on objects....77
Create objects....77
Edit objects....78
Delete objects....79

**Part II: Simulate Your Network....81**

**4 Simulation Overview....83**

Network simulations....84
Auto-resimulation....84
States of plan objects....85
Failed state....86
Active state....89
Operational state....90
Simulated capacity....91
Delay....92

**5 Simulate Traffic Flow from Source to Destination Using Demands....95**

Demands....96
Recommendations for demand sources and destinations....98
Demand meshes....98
Demand latency bounds....98
Use cases of demands....99
Create demands and demand meshes....99
Create a demand....99
Create a demand mesh....100
Create demands for LSPs....101
Set demand latency bounds....102
Visualize demands in the network plot....103
Traffic rerouting....104
View shortest IGP and latency paths, and overlapped paths....105
Demand traffic....106
Create and select traffic levels....107
Simulate traffic using demands....107
Modify demand traffic....109
Demand deduction....112
Estimate demand traffic using demand deduction....115

**6 Perform What-If Analysis....121**

Failure scenarios....122
Perform failure analysis....122
Example: Performing failure analysis....123
Perform impact analysis of topology changes....124
Example: Performing impact analysis of topology changes....124
Perform impact analysis of metric changes....126
Example: Impact analysis of metric changes....126

**7 Evaluate Impact of Worst-Case Failures....129**
Simulation analysis tool....130
Worst-case traffic utilization....130
Worst-case QoS violation....132
Fail circuits to worst-case utilization....132
Worst-case demand latency....133
Fail demands to worst-case latency....134
Run simulation analysis....135
Perform pairwise failure analysis....137
Protect objects....142
Simulation analysis report....143
Failure impact view....144
Configure parallelization for simulation analysis....145

**8 Evaluate Impact of Traffic Growth....147**
Demand groupings....148
Growth plans....148
Create growth plans....148
Growth plans from demand grouping....150
Create growth plans from demand growth....150
Create growth plans from interface growth....151

**9 Perform Capacity Planning....153**
Capacity planning optimization....154
Perform capacity planning....154
Optimization options....156
Advanced optimization options....157
Capacity planning optimization example....158
Parallel circuits design examples....159
Optimization report analysis....160

**10 Simulate IGP Routing Protocol....161**
IGP process ID....162
Configure an IGP process ID....162
Associate interfaces with an IGP process ID....162
OSPF and IS-IS differences....163
OSPF area simulation....163
IS-IS multi-level simulation....164
EIGRP....165
Simulate IGP multipath....166
Exclude ABR nodes....166
Configure IGP simulation....167

Multi-topology routing....167
Configure topologies....168
Add topologies to objects....168

**11 Simulate BGP Routing....171**
Multi AS network modeling for BGP routing....172
Internal and external AS types....172
Configure ASes....172
Create ASes....172
Associate nodes with an AS....173
Edit AS routing policies....174
Route demands between ASes....175
Route selection between internal ASes....175
Routes between external and internal ASes....175
External meshes....178
Know about BGP routing details....179
BGP multihop....179
Configure BGP load balancing....179
BGP next hop....180
How BGP routing works....180

**12 Simulate Quality of Service (QoS)....183**
Quality of Service (QoS)....184
QoS requirements....184
QoS bound and QoS violations....185
Policies and QoS bound calculations....186
Interface queue properties and QoS bound calculations....186
Interface QoS bound calculations using multiple QoS parameters....189
Service class QoS bound calculations using multiple QoS parameters....190
Queue and service class information....190
QoS bound and QoS violation columns....190
Create queues....191
Edit interface queue properties....192
Create service classes....192
Map service classes to queues....193
Create policy groups for interfaces....193
Create or edit service class policies....194
Apply QoS settings globally....195
How QoS violations are indicated in the topology....195

**13 Simulate VPN....197**
VPN model....198
VPN objects....198

VPN topology and connectivity....198
VPNs....199
Create VPNs....200
VPNs table....201
Identify interfaces used by VPNs....202
VPN node properties....202
Create VPN nodes....203
Add VPN nodes to VPNs....203
VPN nodes table....204
Layer 3 VPN example....205
VPN simulation analysis....207

**14 Simulate Advanced Routing with External Endpoints....209**
External endpoints....210
Use cases for routing with external endpoints....210
Create external endpoints and add members....210
External endpoint member attributes....211
Simulate routing with external endpoints....212
Traffic distribution....213

**15 Simulate Multicast....215**
Source-Specific Multicast (SSM)....216
Discovered versus simulated multicast flows....216
Configure multicast flow hops for simulations....217
Configure Cisco next hop for multicast flows....217
Multicast demands....217
Simulated multicast demands....218
View multicast flows....219
Create demands for multicast flows....220

**Part III: Traffic Engineering and Optimization....223**

**16 Optimize Metrics in the Network Core....225**
Metric optimization tools....226
Core and edge distinctions....226
Perform metric optimization....226
Metric optimization options....228
Tactical metric optimization....230
Metric optimization reports....231

**17 Configure MPLS Routing....233**
MPLS routing....234

Supported LSP types....234
Create and visualize LSPs and LSP paths....234
LSP paths....235
Path options and active path....235
Create LSP paths....236
Path latency calculations....237
Route demands through LSPs....237
Routing through intra-area LSPs....237
Routing demands through inter-area LSPs....238
Routing demands through private LSPs....239
Load sharing between LSPs....242
Global simulation parameters....243
Set LSP establishment order....244
LSP simulation diagnostics....244
Run LSP simulation diagnostics....244
Simulation diagnostics for LSP routing failures....245

**18 Optimize LSPs....247**
LSP disjoint path optimization....248
Disjoint path optimization inputs....248
LSP loadshare optimization....254
Run LSP loadshare optimization....254
LSP setup bandwidth optimization....258
Run LSP setup bandwidth optimization....259

**19 Configure RSVP-TE Routing....263**
Dynamic LSP routing and CSPF....264
Configure interface MPLS properties....264
View LSP reservations....266
Calculate LSP setup bandwidth....267
RSVP LSP path attributes and behavior....268
Standby paths....269
How LSP paths respond to failures....269
Named paths and explicit LSP routing....270
Named path hops example....270
Create named paths and their hops....271
Edit named paths and their hops....272
Actual paths....273
Deactivate actual paths for simulations....274
Affinities....274
Create and edit affinities....275
Assign affinities to interfaces....275
LSP affinity associations....276
Assign LSPs to affinities....277

Assign LSP paths to affinities....278
Assign affinities when creating LSP meshes....279
Global simulation parameters....279
Fast Reroute (FRR) simulations....280
FRR fundamentals....281
FRR LSP routing....283
Types of FRR LSP protection....284
Configure routing around failures....284
Set up fast reroute simulations....285
FRR LSPs initializer....286
Create FRR LSPs manually....288
Run fast reroute simulation....288
Autobandwidth-enabled LSPs....289
Autobandwidth convergence example with and without failures....290
Circuit Style RSVP LSPs....292
CS-RSVP LSP key features....293
CS-RSVP LSP paths....293
How CS-RSVP LSP path simulation works....294
Create CS-RSVP LSPs....295
CS-RSVP LSP optimizer....296
CS-RSVP LSP simulation diagnostics....300
CS-RSVP LSP visualization....302
Advanced RSVP-TE LSP simulations....306
Ignore reservable bandwidths for capacity planning....306
P2MP LSPs....310
P2MP LSP bandwidth....311
P2MP LSP demands....313
Requirements for P2MP LSPs and sub-LSPs....313
Create P2MP LSPs....313
Deletion behavior for P2MP LSPs and sub-LSPs....315
Unresolved LSP destinations and hops....316

**20 Optimize RSVP-TE Routing....317**
RSVP-TE optimization....318
RSVP-TE optimization input properties....318
LSP group options....320
Optimization parameters....321
Disjoint groups....322
Avoidance constraints....323
Post optimization options....324
Run RSVP-TE optimization....324
Optimization output....325

**21 Perform Explicit and Tactical Explicit RSVP LSP Optimization....327**

Explicit RSVP LSP optimization tool....328
Run Explicit RSVP LSP optimization....328
Explicit RSVP LSP optimization options....330
Tactical explicit RSVP-TE LSP optimization....334

22 Configure Segment Routing....337
Segment Routing LSPs....338
SR LSP segment types....338
Node segment list hops....338
Interface segment list hops....339
LSP segment list hops....340
Anycast group segment list hops....341
SR LSP paths....342
SR LSP routing rules....343
Inter-AS LSP routing....344
SR LSPs and their LSP paths....344
Flexible Algorithms....345
Application-Specific Link Attribute (ASLA)....345
Create segment lists....345
Create anycast groups....346
Create SIDs....347
Create a node SID....347
Create an SRv6 node SID....347
Create an interface SID....348
Create a SRv6 interface SID....348
Create flex algorithm....349
SR-TE protection....349
SR-TE tunnel simulation before convergence....350
Simulate SR-TE tunnels after convergence....350
Constraints for SR FRR and TI-LFA....350

23 Optimize Segment Routing....353
SR LSP optimization tools....354
SR-TE optimization....354
SR-TE optimization input properties....354
Run SR-TE optimization....356
SR-TE bandwidth optimizations....358
Operating modes for SR-TE bandwidth optimization....358
SR-TE bandwidth optimization input parameters....359
Optimize SR-TE bandwidth....362
Analyze the congestion under different failure sets....363
Bandwidth optimization report....364

**Part IV: Access Reports, Jobs, Add-ons, Patch Files, and Changeover Tool.....367**

**24 Access Reports.....369**

Reports in Cisco Crosswork Planning.....370
Plan comparison reports.....370
Create plan file comparison reports.....371
Report columns.....371
Report summary.....376
Traffic comparison reports.....377
Create traffic comparison reports.....378
View a generated report.....379

**25 View Jobs.....381**

Job manager.....382
User role permissions for job manager.....382
View job details.....382
Job manager column details.....383
Access output plan files from job manager.....384
Run tools or initializers using CLI.....385
Run external scripts.....387
Example: Run external scripts.....388

**26 Manage Add-ons.....389**

Add-ons.....390
Requirements and limitations for add-on configuration.....390
User role permissions for add-ons.....390
Create add-ons.....391
Add-on controls.....392
Run an add-on on a plan file.....393
Import an add-on.....394
Export an add-on.....395
Edit an add-on.....395
Sample add-on.....396

**27 Update Configuration from One File to Another.....397**

Changeover tool.....398
Run the changeover tool.....398
Report analysis.....400

**28 Create and Use Patch Files.....401**

Patch files.....402

Create patch files....402
Apply patch files....403
View or edit patch files....404

**29 Export Routes....405**
Export routes in Cisco Crosswork Planning....406
Route tables....406
Output table column details....406
Export route table example....407
Export routes for selected objects....407

**Appendix A: Cross Table Filters....409**
Supported cross table filters....410

**Appendix B: Legend for Plan Objects....415**
Interfaces and Circuits....416
Nodes....416
Sites....418
Demands....419
LSPs....419
LSP paths....420
LSP named paths....421
LSP actual paths....421
SR LSP paths....421

# 1 Get Started

**Topics:**


- Core capabilities of Cisco Crosswork Planning
- Core concepts in Cisco Crosswork Planning
- Log in to Cisco Crosswork Planning
- Dashboards
- Design engines


Outlines the core features and workflow of Cisco Crosswork Planning, dashboard customization, and design engine management to help users optimize, simulate, and monitor network environments effectively

**Note**

This guide applies to Cisco Crosswork Planning Releases 7.2.0 and 7.2.1 Security Patch (SP). If you are running Release 7.2.0, migrate to Release 7.2.1 SP, which supersedes the earlier releases.

# Core capabilities of Cisco Crosswork Planning

Explains the core capabilities of Cisco Crosswork Planning, highlighting how it enables continuous modeling, real-time monitoring, and effective optimization of your network environments for enhanced performance and reliability.

Cisco Crosswork Planning provides tools to create a model of the existing network by continuously monitoring the network and its traffic demands. At any given time, this network model contains all relevant information about a network, including topology, configuration, and traffic information. You can use this information as a basis for analyzing the impact on the network due to changes in traffic demands, paths, node and link failures, network optimizations, or other changes.

### Key features

Some important features of Cisco Crosswork Planning include:

- Traffic engineering and network optimization: Compute TE LSP configuration to meet service level requirements, perform capacity management, and perform local or global optimization in order to maximize efficiency of deployed network resources.
- Demand engineering: Examine the impact on network traffic flow of adding, removing, or modifying traffic demands on the network.
- Topology and predictive analysis: Observe the impact on network performance caused by changes in the network topology, which are driven either by design or by network failures.
- TE tunnel programming: Examine the impact of modifying tunnel parameters, such as the tunnel path and reserved bandwidth.
- Class of service (CoS)-aware bandwidth on demand: Examine existing network traffic and demands, and admit a set of service-class-specific demands between routers.

### Components

Cisco Crosswork Planning comprises two primary components:

- Cisco Crosswork Planning Collector: This component consists of a set of services that create, maintain, and archive a model of the current network. It achieves this through continual monitoring and analysis of the network and the traffic demands placed on it.
- Cisco Crosswork Planning Design: This component helps network engineers and operators predict growth in their network, simulate failures, and optimize the network design to meet performance objectives while minimizing cost.

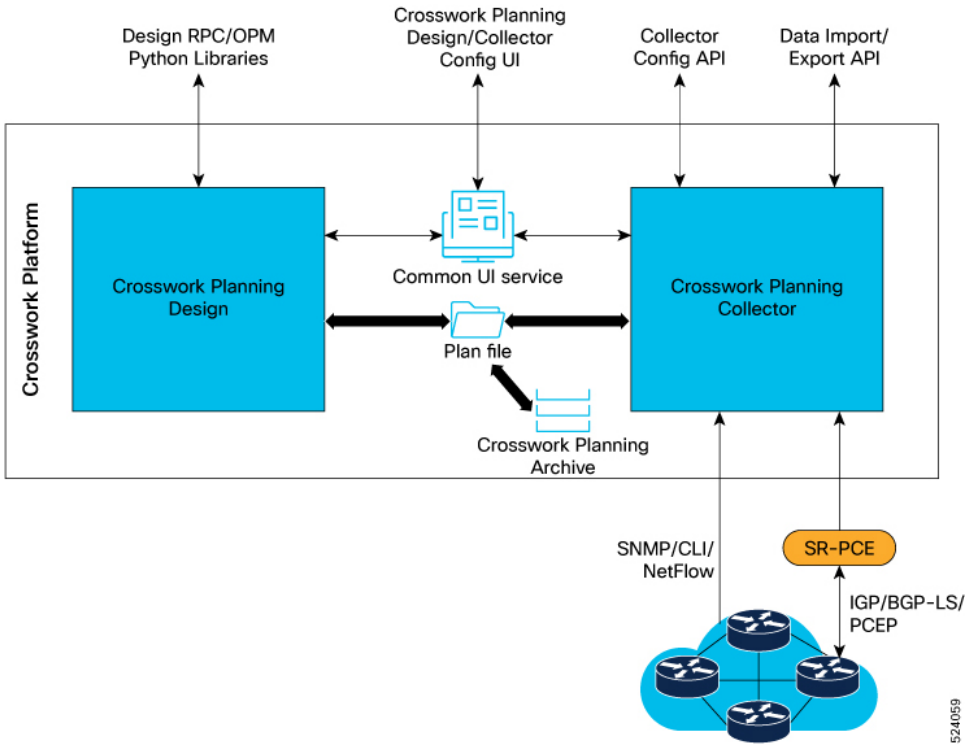
## System overview

This topic explains the main components of the Crosswork Planning system, describes their roles, and explains how they interact.

Cisco Crosswork Planning runs on the Cisco Crosswork infrastructure and is part of the Cisco Crosswork Network Automation suite of products.

Cisco Crosswork Planning Design and Cisco Crosswork Planning Collector are packaged as separate components and can be enabled/disabled as per your requirements. These two applications run independently of each other. The communication between Cisco Crosswork Planning Design and the archive on the Cisco Crosswork Planning Collector to import the network models happens over well-defined APIs.

**Figure 1: Cisco Crosswork Planning system overview**




## Types of Cisco Crosswork Planning Design use cases

This section lists the types of design use cases supported by Cisco Crosswork Planning, covering planning, engineering, and operations scenarios for modeling and optimizing network designs.

Cisco Crosswork Planning lets you model, simulate, and analyze failures, design changes, and traffic growth impacts, while optimizing your network for maximum efficiency.

With Cisco Crosswork Planning network models and tools, you can answer a wide variety of questions about traffic management.

**Planning**:

- Where is traffic going in the network?
- When and where will my network run out of capacity?
- How do I convert my growth forecasts into upgrade plans?

**Engineering**:

- How can I better distribute traffic throughout the network?
- How can I design for differentiated quality of service?
- Where is my network most vulnerable to failure, and how can I mitigate it?

**Operations**:

- How can I visualize the current health of the network?
- What caused the outage experienced earlier today?
- How would a current failure affect the network at peak time tomorrow?
- Which VPN customers will be affected by a planned outage or by an unplanned failure?

# Core concepts in Cisco Crosswork Planning

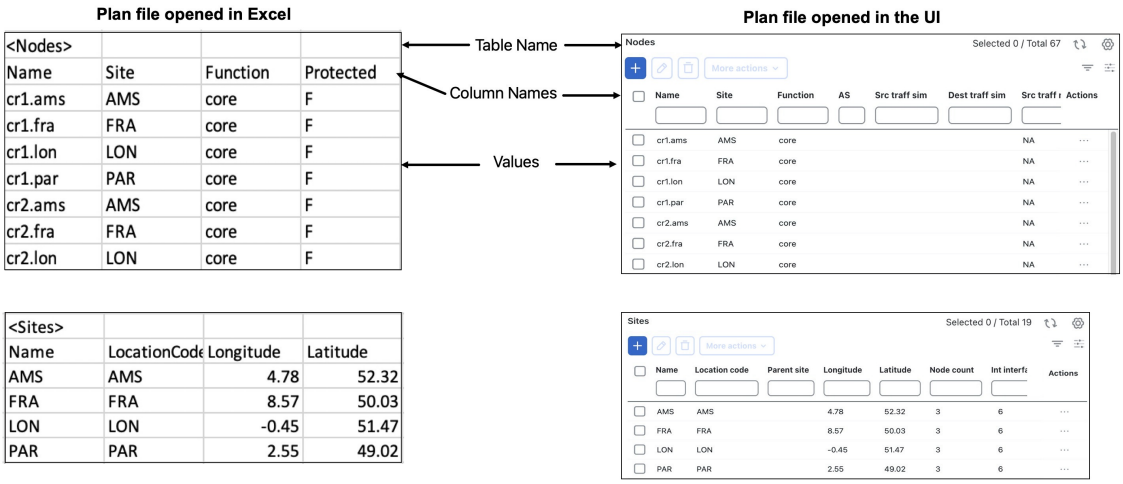
This topic explains the key concepts of Cisco Crosswork Planning to help you understand the essential components and capabilities of the product.

Before you begin using Cisco Crosswork Planning, be familiar with these key concepts.

### Plan files

A plan file consists of tables that describe network characteristics, including network topology, traffic, service classes, and routing protocols. In Cisco Crosswork Planning, each plan file represents one network.

**Figure 2: Opened plan file**




Cisco Crosswork Planning discovery tools initially capture and store the network topology and routing information in plan files. These plan files are the basis of information displayed in and used by Cisco Crosswork Planning.

Plan files include these types of information:

- Network configurations
- Visual layouts
- IP/MPLS routes, including multicast and LSPs
- Measured traffic
- Estimated end-to-end traffic matrix
- Operational state of network objects
- Results of analyses, such as worst-case failure analysis results

There are two formats for plan files:

- The .pln format is compact and can be quickly loaded to and saved in the Cisco Crosswork Planning UI.
- The .txt format contains tab-delimited columns. You can create and edit these directly in a text editor or spreadsheet like Excel to apply bulk changes efficiently. In these plan files, each table is labeled with angle brackets, such as <Nodes> and <Sites>.

Cisco Crosswork Planning Design application uses the information from plan files to visualize underlying network models, which are then used to perform simulations.

### Plan objects

An object is a representation of elements found in networks, such as nodes, sites, circuits, interfaces, LSPs, and so on. For example, a node represents a router in a network. A site is an object, and is a Cisco Crosswork Planning construct for simplifying the visualization of a network by grouping nodes within a site, or even by grouping sites within a site.

For more information, refer to Understand Plan Objects on page 67.

### Design engines

The design engine is the core component that manages simulations and optimizations. It allows performing tasks interactively, either as synchronous or as asynchronous jobs. When you log in and open a plan file, an engine is assigned to your session, subject to engine availability. Once assigned, this engine handles all user real-time activities. Each plan file is maintained as an individual network model, with only one network model active at any given time. You can switch the active network model among all opened network models; all activities apply only to the active network model.

- Engine spaces: Engines can be run in two spaces: user space and job space. Each space operates independently and uses only the resources assigned to it. Each engine runs on the same version of the image. However, depending on whether the engine operates in user space or job space, it runs in synchronous or asynchronous mode, respectively.
- Synchronous mode: When the Cisco Crosswork Planning engine operates in user space, it is in synchronous mode. It can handle user requests interactively and provide results in real-time.
- Asynchronous mode: When the Cisco Crosswork Planning engine operates in job space, it is in asynchronous mode. Depending on computation complexity, some operations may take longer time to complete. These operations are submitted as jobs and the engines running in job space process these job requests. These jobs run in the background without affecting other user activities.

  In asynchronous mode, the engine processes each assigned task independently. After a job is complete, it saves the results as .tar files in the Job Manager. You can download and extract this .tar file, and then import the updated file into the user space.

For more information, refer to Design engines on page 25.

### Patch files

A patch file is a compact way to represent the differences between plan files. These differences, or “patches”, can be applied to other plan files or deployed to the network. Patch files have a .plp format.

For information on creating and applying patches, refer to Create and Use Patch Files on page 401.

## Log in to Cisco Crosswork Planning

Explains how to log in to Cisco Crosswork Planning and access the UI for the first time after installation. This procedure guides you through entering your credentials, setting a new password, and reviewing the application's health status.

Cisco Crosswork Planning is a browser-based application. For more information on supported browser versions, refer to the *"Supported web browsers" section in the Cisco Crosswork Planning 7.2.x Installation Guide*.

Complete these steps to access the UI after installing Cisco Crosswork Planning.

### Procedure

1. Open a web browser and enter `https://<Crosswork Management Network Virtual IP (IPv4)>:30603/`.

   When you access Cisco Crosswork Planning from your browser for the first time, you may see a warning that the site is untrusted. If this occurs, follow the prompts to add a security exception and download the self-signed certificate

from the server. After you do this, the browser accepts the Cisco Crosswork Planning server as a trusted site in all subsequent logins.

2. Log in to Cisco Crosswork Planning.
   a) Enter the administrator username **admin** and the default password **admin**.
   b) Click **Login**.
   c) When prompted to change the administrator's default password, enter the new password in the fields provided and then click **OK**.

   > **Note**
   >
   > Use a strong VM password (a minimum of eight characters, including uppercase and lowercase letters, numbers, and at least one special character). Avoid using passwords similar to dictionary words (for example, "Pa55w0rd!") or relatable words.

   The **Crosswork Manager** page appears.
3. Click the **Crosswork Health** tab, and click the **Crosswork Platform Infrastructure** tab to view the health status of the microservices running on Cisco Crosswork Planning.
4. (Optional) Change the name assigned to the admin account to something more relevant.

---

You gain access to Cisco Crosswork Planning and you can begin planning or managing tasks as needed.

### What to do next

To log out, in the top-right corner of the main page, click > **Logout**.

> **Note**
>
> Logging out does not close the plan file you are working on; the file remains open.

## Dashboards

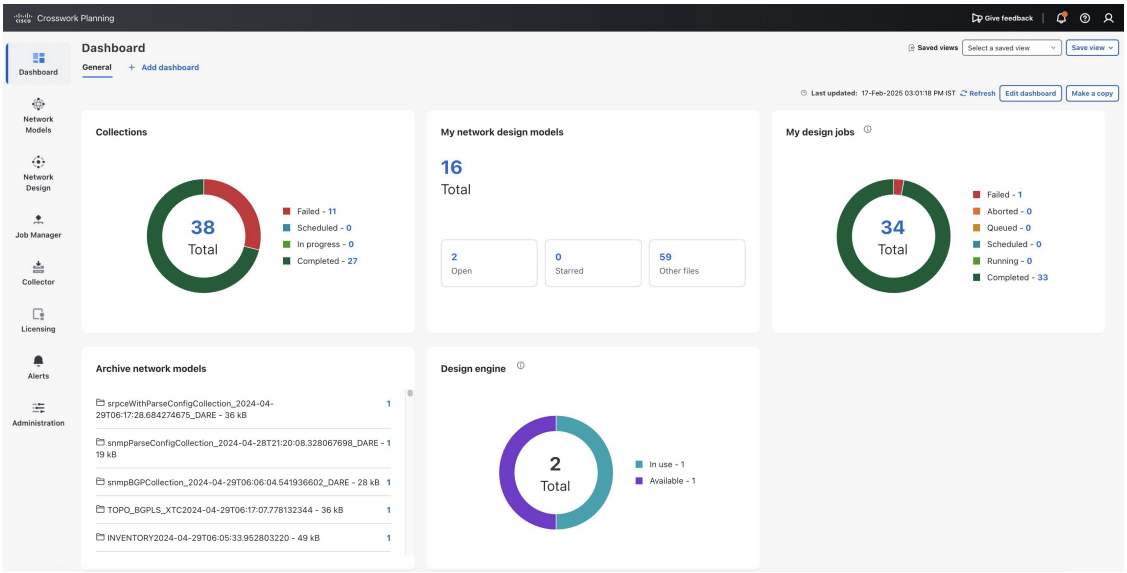
Outlines the dashboard page's role in Cisco Crosswork Planning, detailing its overview and customization options.

A **Dashboard** is a UI page that

- provides a quick operational summary of Cisco Crosswork Planning
- allows you to navigate directly to relevant pages and functions, and
- consists of multiple dashlets that vary according to the installed application.

The **Collections** and **Archive network models** dashlets appear only if you have installed the Cisco Crosswork Planning Collector application. Similarly, the **My network design models**, **My design jobs**, and **Design engine** dashlets appear only if you have installed the Cisco Crosswork Planning Design application.

**Figure 3: Dashboard view**




### Dashlet navigation

Links in each dashlet allow you to navigate to the desired pages easily. For example, in Figure 3: Dashboard view on page 20, link **2** in the **Open** tab in the **My network design models** dashlet indicates that there are two network models open in the UI. If you click this number **2**, the two opened network models are displayed in the **Network Design** page.

### Dashlet customization

Use the **Edit dashboard** button at the top right corner to customize how the dashlets appear.

## Customize dashlets

Guides you through customizing dashboard dashlets, including rearranging, adding, removing, renaming, and copying dashlets using the Edit dashboard feature.

Using **Edit dashboard** at the top right corner, you can

- rearrange the dashlets as per your requirement
- add dashlets to the dashboard
- remove dashlets from the dashboard
- edit the names of the dashlets, and
- make copies of the dashlets.

### Procedure

1. From the main menu, navigate to **Dashboard**.

   The Dashboard page opens with the default **General** tab. The configured custom dashboards appear as individual tabs.

2. Go to the required dashboard and click **Edit dashboard** at the top right. The page becomes editable.
3. To rearrange dashlets:
   a) Click ⋮⋮ in the dashlet you want to move, and drag and drop to the position you want.
   b) Save the changes.

4. To add dashlets to the dashboard:
   a) Click **+ Add dashlet** at the top right. The Add Dashlets window appears.
   b) From the list, select the dashlet you want.
   c) Click **+ Add**.
   d) Save the changes.
5. To edit dashlet names:
   a) In the text field displaying the current dashlet name, enter the new name for the dashlet. Ensure that the title is between 5 and 50 characters long.
   b) Save the changes.
6. To copy dashlets:
   a) Click ··· > **Copy** in the dashlet you want to copy.

      A copy of the dashlet appears in the dashboard.

   b) Save the changes.
7. To remove dashlets:
   a) Click ··· > **Delete** in the dashlet you want to delete.

      The selected dashlet is removed from the dashboard.

   b) Save the changes.

## Custom dashboards

Introduces custom dashboards, explaining how to create dashboard pages with selected dashlets, manage tabs by renaming or deleting them, and identifies the default, non-editable General dashboard tab.

A custom dashboard is a dashboard page that

- you can create to include dashlets of your choice
- is displayed as a separate tab within the Dashboard page, and
- can be renamed or deleted.

Cisco Crosswork Planning includes a default dashboard tab named General, which you cannot rename or delete.

### Create custom dashboards

Learn how to create custom dashboards and add relevant dashlets to personalize your monitoring and visualization experience.

Use this task to create a custom dashboard and to add relevant dashlets.

#### Procedure

1. From the main menu, navigate to **Dashboard**.

   The Dashboard page opens with the default **General** tab.
2. Click the **+ Add dashboard** at the top right to create a new dashboard tab.
3. Click the edit icon (✎) next to the default title, and enter a meaningful name for the dashboard. Ensure that the title is between 5 and 50 characters long.
4. Add the required dashlets:

a) Click **+ Add dashlet**.
b) In the Add Dashlets drawer, select the dashlet you want.
c) Click **+ Add**.
The selected dashlets are added to your dashboard.
d) Enter a unique, meaningful title for each dashlet.

5. Click **Save** to save the dashboard.

### Edit custom dashboards

Guides you through editing custom dashboards by renaming dashboard titles and adding or removing dashlets, including navigation steps and restrictions on modifying the default General dashboard.

You can modify a custom dashboard by renaming it, or by adding or removing dashlets.

> **Note**
>
> You cannot rename the default **General** dashboard. You can rename only the custom dashboards created by users.

#### Procedure

1. From the main menu, navigate to **Dashboard**.

   The Dashboard page opens with the default **General** tab. The configured custom dashboards appear as individual tabs.

2. Go to the required custom dashboard and click **Edit dashboard** at the top right.
3. Edit the dashboard settings as required:

   - Rename the dashboard: Click next to the dashboard title at the top left and enter a meaningful name. Ensure that the title is 5 to 50 characters long.
   - Modify dashlets: To modify dashlets, refer to Customize dashlets on page 21.

4. Click **Save** to save the dashboard.

### Delete custom dashboards

Learn how to remove custom dashboards that you no longer need.

Use this task to delete the custom dashboards that you no longer need.

> **Note**
>
> You cannot delete the default **General** dashboard. You can delete only the custom dashboards created by users.

**Procedure**

1. From the main menu, navigate to **Dashboard**.

   The Dashboard page opens with the default **General** tab. The configured custom dashboards appear as individual tabs.

2. Go to the required custom dashboard and click **Edit dashboard** at the top right.
3. Click **Delete dashboard** at the top right.
4. Click **Delete** in the confirmation dialog box.

The selected custom dashboard is permanently removed from the Dashboard page.

## Reset the General dashboard to the default view

Provides step-by-step instructions to reset the General dashboard to its default view.

> **Note**
>
> You cannot perform this procedure on custom dashboards.

Use this task to reset the General dashboard to the default view.

**Procedure**

1. From the main menu, navigate to **Dashboard**.
   The Dashboard page opens with the default **General** tab. The configured custom dashboards appear as individual tabs.
2. On the **General** dashboard, click **Edit dashboard** at the top right.
3. Click **Reset to default** at the bottom of the page.

   The Reset Dashboard confirmation dialog box appears.

4. Click **Reset** in the confirmation dialog box.

## Copy dashboards

Provides instructions for duplicating dashboards through the Dashboard interface, including navigating dashboard tabs and initiating the copy process.

Use this task to make a copy of the dashboard.

**Procedure**

1. From the main menu, navigate to **Dashboard**.

   The Dashboard page opens with the default **General** tab. The configured custom dashboards appear as individual tabs.

2. Go to the required dashboard and click **Make a copy** at the top.
   The system creates a copy of the dashboard named *dashboard*(COPY).

# Design engines

Outlines how design engine instances in Cisco Crosswork Planning manage simulations and optimizations, describe resource configurations, and explain options to support multiple concurrent users and efficient resource utilization.

A design engine is a core component that

- manages simulations and optimizations
- is assigned to user sessions depending on engine availability, and
- allows you to perform tasks interactively as synchronous or asynchronous jobs.

For an overview of design engines, refer to Core concepts in Cisco Crosswork Planning on page 18.

### Design engine instances

Post installation, two design engine instances are created in the system.

- cp-sync-engine-default: This engine is allocated for handling synchronous tasks.
- cp-async-engine-default: This engine is allocated for running asynchronous jobs.

By default, these engine instances are configured with 4 vCPU (4000m vCPU) and 16 Gi of memory for large profiles. For extra large (XLarge) profiles, they are configured with 8 vCPU (8000m vCPU) and 24 Gi of memory.

Cisco Crosswork Planning allows you to set up additional design engine instances, and allocate vCPU and memory based on your requirements. You can create up to 10 design engine instances in total, including both synchronous and asynchronous instances. By having multiple instances, you can run several tasks at the same time. For example, if there are eight asynchronous engine instances, it means that eight scheduled jobs can be run in parallel. Similarly, if the number of synchronous engine instances is five, it means that five users can run the synchronous jobs simultaneously.

If there are distinct concurrent users logged in to Cisco Crosswork Planning Design to perform simulations, then an equivalent number of synchronous design engine instances is required.

**Note**

Idle design engine instances do not consume compute resources until they are allocated.

## Design engine instances resource management

Outlines how to manage CPU and memory resource limits for design engine instances, describes allocation strategies, shares best practices for high concurrency and large network models, and illustrates resource management scenarios in multi-engine deployments.

When creating a design engine instance, you specify the maximum limit for CPU (vCPU) and memory (RAM) that it can use. The system enforces these limits, so the instance cannot go over the specified resource limit during runtime. These limits do not reserve or block resources for the instance; resources are shared, and their usage depends on what is available at runtime.

- CPU: The design engine instance cannot use more CPU than its specified maximum limit. If it reaches the limit, the system slows down the instance's CPU usage.
- Memory: If the design engine instance exceeds its memory limit, the system may terminate and restart it, especially if the node is low on memory.

**Resource allocation for design engine instances**

In large scale deployments, you may need to complete simulation analysis faster, conduct capacity planning more quickly, or perform tasks on huge plan files. In such cases, allocate higher vCPU and memory. The maximum vCPU and memory you can allocate to an instance depend on the system's total resources.

For example, on a system with 4 vCPUs and 64 Gi memory, a single instance can be allocated up to 4000m vCPU and 64 Gi memory.

To find the system's resource information, go to the **Administration** > **Crosswork Manager** > **Crosswork summary** > **System summary** page. The **Used** section reflects the current resource consumption.

The resource limits for each instance appear in the **Engine profile** column of the Design Engine Instances page (**Administration** > **Settings** > **Design settings** > **Design engine instances**).

**Best practices and key considerations**

- Resource limits are upper bounds, not reservations.
- If the total amount of resources set for all engine instances is more than what the system has, the engines may have to compete for resources, which can make them run slower or work less efficiently.
- The maximum number of concurrent users is determined by the number of synchronous engine instances.
- Each engine operates independently. Performance of one is not directly affected by others. However, the overall system performance depends on how much total resources are being used by all engines together. If you run all engine instances at their highest CPU and memory settings at the same time, the total resource usage might exceed the system's capacity. This condition can cause the engines to slow down or fail.
- For deployments with high concurrency or large network models, ensure sufficient CPU, memory, and disk resources for all configured engines.
- When planning deployments, consider both the network model size and the expected number of concurrent users to select the appropriate hardware profile and VM count.

**Example scenarios**

**Example 1:**

If your system has 16 vCPUs and you create 4 engines with a 4 vCPU limit each, then if all engines run at their maximum at the same time, they will use all 16 vCPUs, which is the total capacity of your system.

**Example 2:**

If you set up 8 engines with a 4 vCPU limit each on a system that only has 16 vCPUs, and all engines try to use their full limit at the same time, the system will limit the total usage. As a result, some engines may be slowed down.

# Create design engine instances

Provides instructions for creating design engine instances, including configuring CPU and memory allocation, choosing synchronous or asynchronous services, and verifying engine availability and integration within the dashboard and configuration interfaces.

Complete these steps to create design engine services as per your requirement.

## Before you begin

Ensure you have the "CW Planning Design Setting APIs" Role permission (**Administration** > **Users and Roles** > **Users**). For more information on managing user roles and permissions, refer to the *"Manage Users" section of the Cisco Crosswork Planning 7.2.x Collection Setup and Administration* document.

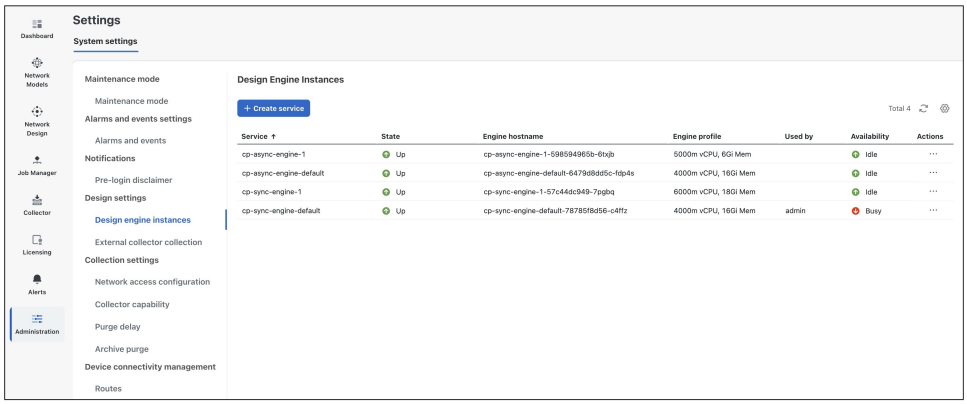
> **Note**
>
> You can allocate a maximum of 10 design engines for both synchronous and asynchronous jobs combined.

## Procedure

1. From the main menu, choose **Administration** > **Settings** > **System settings** > **Design settings** > **Design engine instances**.

   The Design Engine Instances page opens.

   **Figure 4: Design engine instances page**

2. Click **Create Service**.

   The Create Service page opens.

3. Enter the **Virtual CPU** in milli vCPU as per your requirement. For example, to allocate 5 vCPUs, enter 5000. Ensure the value is greater than 100 and does not exceed the system's virtual CPU capacity. For example, if the system has 6 vCPUs, do not enter a value more than 6000.
4. Enter the **Memory** in Gi (gibibyte) as per your requirement. Ensure that the entered value does not exceed the system's memory.
5. Choose whether to create synchronous or asynchronous instances.
6. Save the changes.

   The creation process takes approximately 1 to 2 minutes. On successful creation, the **Availability** column of the newly created service displays the status as "Idle".

If this is the first service you are creating, a service named either cp-sync-engine-1 or cp-async-engine-1 is created based on your selection in Step 5. Subsequently, cp-sync-engine-2 or cp-async-engine-2 is created, and the pattern will continue.

---

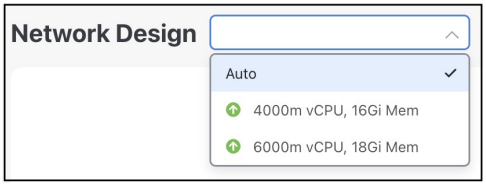
Notice that:

- The number of design engines increases in the **Design Engines** dashlet on the Dashboard page.
- The newly created synchronous engine profile appears under the **Auto** drop-down list on the **Network Design** page. For details on using the dynamically created engine profiles, refer to Select design engine profiles on page 28.
- The newly created asynchronous engine profile appears under the **Engine Profile** section on the **Run Settings** page while running the tools as scheduled jobs.

## Select design engine profiles

Describes how to change the design engine profile as per your requirement.

At the top of the **Network Design** page, all available synchronous engine profiles are displayed. By default, a synchronous engine profile with the highest memory is selected (Auto). If you need to work on a large plan file that requires more vCPU and memory, you can select the engine profile with higher vCPU and memory allocations. For information on design engine profiles, refer to Design engines on page 25.




### Before you begin

Ensure that the required engine profiles are available. For more information on creating design engine instances, refer to Create design engine instances on page 27.

Complete these steps to change the engine profile as per your requirement.

### Procedure

---

1. From the main menu, choose **Network Design**.
2. Close any open plan files.
3. Select the required synchronous engine profile from the drop-down list located next to the page title.

   Ensure that the engine profile's "State" is Up and "Availability" is Idle. If the does not meet this requirement, a red icon is displayed next to it in the drop-down. You can also check the status of the engine profile on the **Design Engine Instances** page.
4. Open the required plan files and start working on them. They will now use the selected engine profile for all synchronous tasks.

---

## Reset engine sessions

Provides instructions for resetting design engine sessions.

The **Design Engine Instances** page shows details about the design engine instance's availability and user status.

- The **Availability** column indicates if the engine is "Busy" or "Idle" based on user activity. If a user is working on a plan file, the corresponding sync instance displays "Busy" in the Availability column. Similarly, when a scheduled job runs, the async instance shows "Busy". When the engine instance is not in use, the Availability column is marked as "Idle."
- The **Used by** column specifies the user who is currently using the engine instance.

You can free up an engine instance from a user, allowing it to be reassigned to a different user. When a user session is ended, all open plans for the selected user are closed. Any operations triggered by the user are aborted, thereby freeing up the engine instance from the selected user. If the Cisco Crosswork Planning Design application becomes unresponsive or stuck due to certain operations, this feature helps to reset it.

**Before you begin**

Ensure you have the "CW Planning Design Setting APIs" Role permission (**Administration** > **Users and Roles** > **Users**). For more information on managing user roles and permissions, refer to the *"Manage Users" section of the Cisco Crosswork Planning 7.2.x Collection Setup and Administration* document.

Complete these steps to reset engine sessions.

**Procedure**

1. From the main menu, choose **Administration** > **Settings** > **System settings** > **Design settings** > **Design engine instances**.

   The Design Engine Instances page opens, displaying details of the available engine instances.
2. Under the **Actions** column, click ··· > **Reset session** for the engine instance whose session you want to reset.
3. Click **Reset** in the confirmation dialog box.

After resetting the session, notice that:

- The **Availability** column displays the status as "Idle".
- For sync engine instances, the opened plan files are closed.
- For async engine instances, the scheduled jobs are aborted.

## Delete design engine instances

Provides step-by-step instructions for deleting user-created design engine instances, including navigation to system settings and important restrictions on removing default instances.

> **Note**
>
> You cannot delete the default design engine instances.

Complete these steps to delete the user-created design engine instances.

**Before you begin**

Ensure you have the "CW Planning Design Setting APIs" Role permission (**Administration** > **Users and Roles** > **Users**). For more information on managing user roles and permissions, refer to the *"Manage Users" section of the Cisco Crosswork Planning 7.2.x Collection Setup and Administration* document.

**Procedure**

1. From the main menu, choose **Administration** > **Settings** > **System settings** > **Design settings** > **Design engine instances**.

   The Design Engine Instances page opens, displaying the details of the available engine instances.

2. Under the **Actions** column, click ··· > **Delete** for the service you want to delete.

   Ensure that the design engine service you selected is not the default one.

3. Click **Delete** in the confirmation dialog box.

   The deletion process takes approximately 1 to 2 minutes to complete.

## Recreate design engine instances

Provides instructions for recreating design engine instances after upgrades, configuration restores, or data migrations using the application's administration settings.

The user-created design engine instances are deleted when

- upgrading the Cisco Crosswork Planning Design application
- restoring the Cisco Crosswork Planning configuration data after a natural or human-caused disaster, or
- migrating data from an older version of Cisco Crosswork Planning to a newer version.

Complete these steps to recreate the design engine instances.

**Procedure**

1. From the main menu, choose **Administration** > **Settings** > **System settings** > **Design settings** > **Design engine instances**.

   The Design Engine Instances page opens, displaying the details of the available engine instances.

2. Under the **Actions** column, click ··· > **Recreate** for the service you want to recreate.
3. Click **Recreate** in the confirmation dialog box.

   The selected design engine instance is recreated.

# Part I

# Visualize Network Models

**Topics:**


- Visualize Network Models
- Understand Plan Objects

# 2 Visualize Network Models

**Topics:**


- Create or import network models
- Open plan files
- Delete plan files
- Download plan files
- Key features of the Network Models page
- Key features of the Network Design page
- Plan tables
- Assign sites to nodes
- Assign geographic locations to sites
- Network topology
- Traffic utilization colors
- Network summary tables
- Enable dark mode


Explains key concepts and procedures for creating, importing, viewing, managing, and customizing network models, including plan file operations, UI navigation, site assignments, topology visualization, and network summary tables.

# Create or import network models

Guides you through creating new network models in the UI or importing existing plan files into user space for use within the application.

The network models are typically generated using the Cisco Crosswork Planning Collector application and are saved as plan files. You can create new network models in the Cisco Crosswork Planning Design UI or import existing one from your local system. To use the plan files in the Cisco Crosswork Planning Design application, they must be available in the user space.

Use these procedures to import plan files into the user space or to manually create new network models.

## Import plan files from the local machine

Guides you through importing plan files with .txt, .pln, or .db extensions from your local machine into your user space, including optional tagging and tag creation for organization within network models.

This topic describes how to import the plan files into the Cisco Crosswork Planning user space from your local machine.

### Before you begin

Ensure the plan file you want to import has a .txt, .pln, or .db extension.

#### Procedure

1. From the main menu, choose **Network Models**.

   By default, **User space** > **My network models** page opens.

2. Click **Import plan file**.

   The Import Plan File page appears.

3. Click **Browse** and choose the plan file that you want to import.
4. (Optional) Select the required tags from the list (if available) or create new tags.

   Complete these steps to create new tags.

   a) Click **Add new tag**.
   b) Enter the tag name.
   c) Click the **+** icon next to the text field.

   > **Note**
   >
   > If a tag is removed from all available plan files, it is deleted from this list as well.

5. Click **Import**.

The plan file is imported into the **My user space** > **Network models** page.

## Import plan files from archive

Explains how to import plan files from local or remote archives, including handling different deployment scenarios.

This topic describes how to import plan files into the Cisco Crosswork Planning user space from Local or Remote archive.

### Before you begin

The archived network model from the Cisco Crosswork Planning Collector application is saved in a plan file format (.pln). The location of the archive differs based on whether the Cisco Crosswork Planning Design and Collector applications are installed on the same machine or different machines.

- If the Cisco Crosswork Planning Design and Collector applications are installed on the same machine, the archived network models appear under **Network Models** > **Local archive**.
- If the two applications are installed on different machines, connect to the machine where the Cisco Crosswork Planning Collector application is installed. After connecting, the archived network models appear under **Network Models** > **Remote archive** of the Cisco Crosswork Planning Design application. For more information, refer to the *"Scenario 2: When the Cisco Crosswork Planning Design and Collector Applications are Installed on Different Machines" section in the Cisco Crosswork Planning 7.2.x Collection Setup and Administration document.*

### Procedure

1. From the main menu, choose **Network Models**.

   By default, **My user space** > **Network models** page opens.

2. In the left pane, under **Local archive** or **Remote archive**, the archived collections are displayed. Select the desired collection name from the list. The right panel shows the plan files that have been created under this collection at different scheduled times. Use the **Last updated** column to know the time at which the plan file was created.
3. Select the required plan file from the right panel and click ··· > **Export to user space** under the **Actions** column.

   The Export Plan to User Space page appears.

4. (Optional) In the **Save as** field, enter a new name for the plan file.
5. (Optional) Select the required tags from the list (if available) or create new tags.

   Complete these steps to create new tags.

   a) Click **Add new tag**.
   b) Enter the tag name.
   c) Click the **+** icon next to the text field.

   > **Note**
   >
   > If a tag is removed from all available plan files, it is deleted from this list as well.

6. Click **Save**.

The plan file is imported into the **My user space** > **Network models** page.

## Create new network models manually

Details the manual creation process for new network models in user space, including naming conventions, mandatory file extensions, and organizing network models with custom tags for streamlined network design and administration.

This topic describes how to create new network models in the Cisco Crosswork Planning user space manually.

**Procedure**

1. From the main menu, choose **Network Models**.

   By default, **My user space** > **Network models** page opens.

2. Click **Create new**.

   The New Network Model page opens.

3. In the **Name** field, enter the name of the plan file.

   > **Note**
   >
   > Ensure that you include the .txt or .pln extension while entering the plan file name.

4. (Optional) Select the required tags from the list (if available) or create new tags.

   Complete these steps to create new tags.

   a) Click **Add new tag**.
   b) Enter the tag name.
   c) Click the **+** icon next to the text field.

   > **Note**
   >
   > If a tag is removed from all available plan files, it is deleted from this list as well.

5. Click **Save**.

The newly created network model opens in the **Network Design** page. By default, two nodes and interfaces are added to the network model.

**What to do next**

- Add nodes and site details to the network model you created. For more information, refer to Create objects on page 77.
- Open the plan file in the **Network Design** page to perform any actions, as per your requirement.

# Open plan files

Describes how to open plan files from the Network Models page and the Network Design page.

You can open the plan file either from the **Network Models** page or from the **Network Design** page. The plan file opens in the Network Design page. You can then perform various actions on the plan files as needed.

**Before you begin**

Ensure that the plan file is available in the User space. To create or import the plan file into the user space, use the procedures described in Create or import network models on page 34.

**Note**

- You can open only three plan files simultaneously, and only one plan file can be active at a time.
- If you log out while working on a plan file, the file does not close; it remains open.

**Procedure**

1. To open the plan file from the Network Models page:
   a) From the main menu, choose **Network Models**.

      The Network Models page opens listing all the available plan files.

   b) Open the required plan file using one of these methods:
      - Click the name of the required plan file.
      - Under the **Actions** column, click ••• > **Open** for the plan file you want to open.

2. To open the plan file from the Network Design page:
   a) From the main menu, choose **Network Design**.
   b) Click the **+** icon at the end of the plan files tab bar displayed at the top of the page.

      The Network Models page opens listing all the available plan files.

   c) Open the required plan file using one of these methods:
      - Click the name of the required plan file.
      - Under the **Actions** column, click ••• > **Open** for the plan file you want to open.

The plan file opens in the **Network Design** page.

- The network plot shows nodes connected by circuits. A circuit is two directly connected interfaces.
- You can click and drag nodes to change their position.
- The measured data in the model represents the period during which the data was collected from the network. This is the model used in the Cisco Crosswork Planning application.
- In the **My network models** page, the selection check box next to the opened plan file appears grayed out, preventing you from deleting it.

## Delete plan files

Guides you through deleting plan files from user spaces, including prerequisites such as file selection limits, admin user permissions, open file restrictions, and the removal of associated reports.

Use this task to delete plan files from the user spaces.

**Before you begin**

Consider these points when deleting plan files:

- You can delete up to 10 plan files at a time.
- Only admin users can delete the files that are available in other user spaces.
- You cannot delete the plan files that are currently open.
- When you delete plan files, the associated reports are also deleted.

**Procedure**

1. From the main menu, choose **Network Models**.
2. Select the plan files that you want to delete and click the icon at the top.

   If you are deleting a single file, you can use the ··· > **Delete** option under the **Actions** column.
3. Click **Delete** in the confirmation dialog box.

## Download plan files

Provides instructions for downloading one or more plan files to your local machine, including available options for single or multiple files, and limits on batch downloads.

Use this task to download plan files to your local machine.

**Procedure**

1. From the main menu, choose **Network Models**.
2. From the user space, select the plan files that you want to download and click the icon at the top.

   If you are downloading a single file, you can use the ··· > **Download** option under the **Actions** column.

   **Note**

   You can download up to 10 plan files at a time. When you download more than one plan file, the files are bundled together in a single .tar file.

The plan files are downloaded to your local machine.

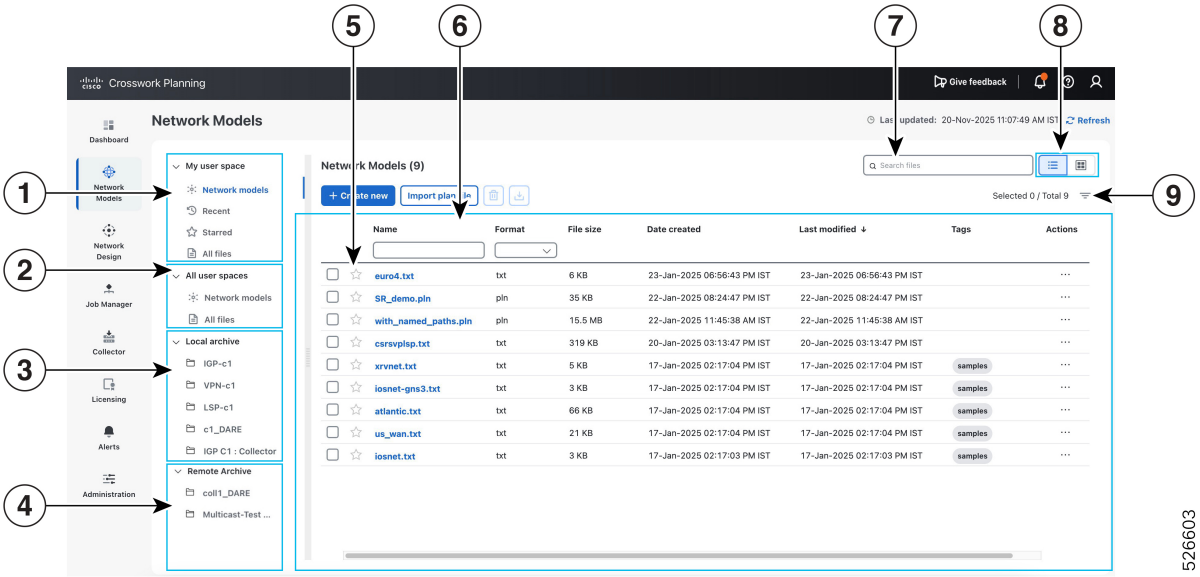
## Key features of the Network Models page

This section lists the various key features of the **Network Models** page.

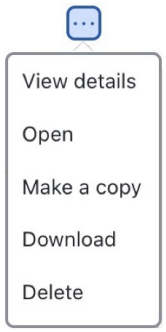
This topic describes the primary features and functions of the **Network Models** page.

All plan files, archived network models, and other files are listed on this page. To access this page, from the main menu, choose **Network Models**.

**Figure 5: Network models user interface**




| Callout No. | Description |
|---|---|
| 1 | **My user space**: My user space serves as an exclusive storage for the logged-in user, containing plan files and other files. Each user has their own space where they can view, edit, create, or delete plan files. Except for admin users, users cannot view the plan files used by other users.<br>By default, My user space contains a few sample plan files.<br>There are four subsections in this section:<br>• **Network models**: Lists all the plan files accessible to the logged-in user.<br>**Note**<br>In the right panel, the grayed out selection check box next to a plan file indicates that the plan file is currently open, and you cannot delete it.<br>• **Recent**: Displays the logged-in user's recently accessed plan files.<br>• **Starred**: Displays the plan files that are marked as favorites by the logged-in user.<br>• **All files**: In addition to plan files, Cisco Crosswork Planning supports a few other type of files. For example, patch files, output result files, and so on. The **All files** section displays these types of files (along with plan files) accessible to the logged-in user. |
| 2 | **All user spaces**: If you are an admin user, you can access the plan files from the user spaces of other users. The **All user spaces** section displays the plan files and all other files that are available in the user spaces of other users. |

| Callout No. | Description |
|---|---|
| 3 | **Local archive**: If the Cisco Crosswork Planning Collector and Cisco Crosswork Planning Design applications are installed on the same machine, the **Local archive** section displays the archived network models that are generated using this application. |
| 4 | **Remote archive**: Displays the archived network models that are generated by the external collectors. Access this archive when the Cisco Crosswork Planning Collector and Cisco Crosswork Planning Design applications are installed on different machines. |
| 5 | **Starred**: Click the ☆ icon to mark the plan files as favorites. The files marked as favorites appear under **My user space** > **Starred**. |
| 6 | **Plan files table**: The table lists all the plan files that are available in the user space.<br>Use the ··· icon under the **Actions** column to<br>• view the details of the plan file.<br>**Note**<br>Use this option to add or delete tags from the plan files.<br>• open the plan file in the Network Design page<br>• make a copy of the plan file in the user space<br>• download the plan file to your local machine, and<br>• delete the plan file from the user space.<br>**Figure 6: Actions menu**<br> |
| 7 | **Search plan files**: Use the Search bar at the top of the page to search for the required plan file. |
| 8 | **List and card layouts**: The plan files in the user space can be visualized in a list layout (default) or a card layout. Use these icons to toggle between list and card views. |
| 9 | **Filter**: Use ⚌ to filter the plan files based on file types or by tag names.<br>It also allows you to filter the data in the table based on the date range you want to view (specific date, 3 months, 1 month, 1 week, and 1 day). |

# Key features of the Network Design page

Lists the various key features of the **Network Design** page.

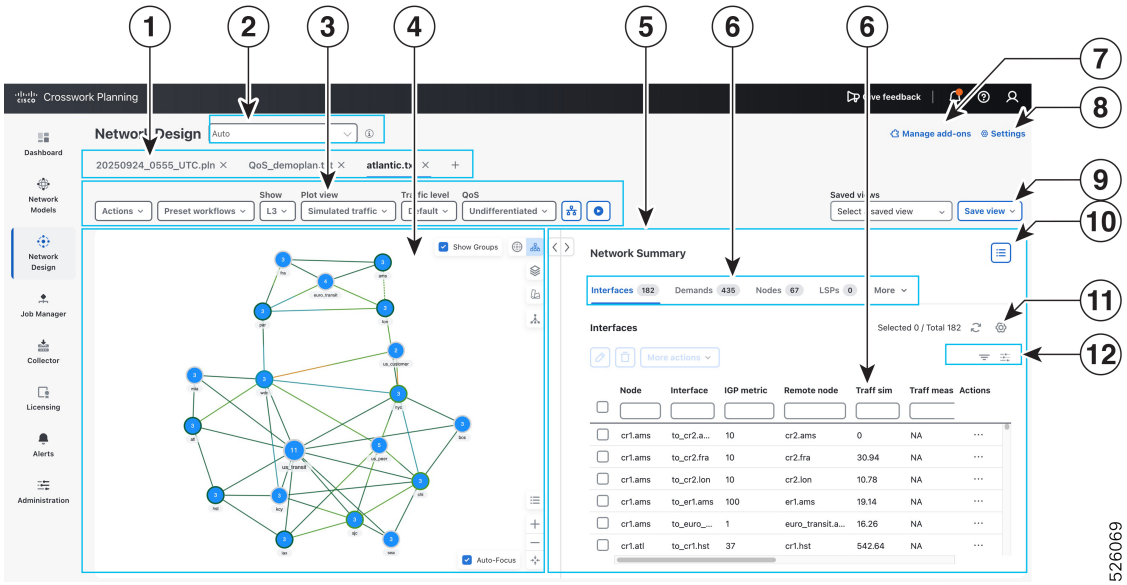
This topic describes the primary features and functions of the **Network Design** page.

When you click the plan file name in the Network Models page, it opens in the Network Design page. This page displays the network plot, which shows sites that appear to be connected by circuits. In reality, these circuits connect to nodes (routers) within each site. The nodes within each site are connected with intrasite circuits. Each circuit consists of two interfaces. The page also displays information related to each object in the network model in a tabular format.

There are four primary sections in the Cisco Crosswork Planning Network Design page.

- Plan file tabs
- Toolbar
- Network plot
- Network Summary panel

**Figure 7: Network Design user interface**




| Callout No. | Description |
|---|---|
| 1 | **Plan file tabs**: Identifies which plan files are currently open. You can open up to three plan files at a time. To bring a plan file to the forefront and make it active, click its associated tab. To open a new plan file, click the **+** icon at the end of this tab bar. |
| 2 | **Design engine profile selection**: Allows you to choose the design engine profile for synchronous tasks. For more information, refer to Select design engine profiles on page 28. |

| Callout No. | Description |
|---|---|
| 3 | **Toolbar**: Provides options for network modeling, simulation, and optimization.<br>• **Actions**: Contains options related to tools, initializers, and reports. Use the **Actions** > **Insert** option to create various plan objects.<br>• **Preset workflows**: Includes options that allow you to navigate to the simulation and optimization tools quickly.<br>  • **Evaluate impact of failures**: Helps you navigate to the Simulation analysis tool.<br>  • **Evaluate impact of traffic growth**: Helps you navigate to the Create growth plans tool.<br>  • **Perform capacity planning**: Helps you navigate to the Capacity planning optimization tool.<br>  • **Perform optimization**: Helps you navigate to the various LSP, SR LSP, and RSVP LSP optimization tools.<br>• **Show L3**: Indicates that the UI is currently displaying Layer 3 (L3) networks.<br>• **Plot view**: Lets you choose the type of traffic displayed in the plot. For more information, refer to Network plot views on page 51.<br>• **Traffic level**: Lets you create and apply traffic levels globally for each opened plan file.<br>• **QoS**: Lets you create and apply QoS settings globally for each opened plan file. For more information, refer to Apply QoS settings globally on page 195.<br>• **Network options** icon (): Lets you set global simulation modes, such as Full Convergence or Fast Reroute for MPLS, and whether to include multicast hops. You can also set IGP, BGP, and Multicast protocol options.<br>• **Simulate** icon (): Triggers resimulation. By default, any change that voids the current simulation does not automatically trigger resimulation. For example, a change in topology requires a resimulation. For more information, refer to Auto-resimulation on page 84. |

| Callout No. | Description |
|---|---|
| 4 | **Network plot**: Displays a visual representation of the network topology. The color on each interface indicates its percentage of traffic utilization (see Traffic utilization colors on page 55). When you select the objects in the Network Summary tables, such as demands or interfaces, they are highlighted in the plot.<br>• **and** icons: There are two views available to visualize the network topology graphically: Schematic view and Geographical view. For more information, refer to Graphical network topology views on page 53.<br>• **Show Groups** check box: In a schematic view, use this check box to group or ungroup the nodes. For more information, refer to Grouping of nodes on page 53.<br>• **Map options** icon ( ): Use this icon to plot different types of routes for Demands and LSPs. For example, Simulated paths, Shortest IGP paths, Shortest latency paths, and so on.<br>• **Utilization colors** icon ( ): Use this icon to modify the traffic utilization colors as needed. For more information, refer to Traffic utilization colors on page 55.<br>• **Auto-Focus** check box: When this option is selected, selecting an object in the Network Summary table automatically focuses on it in the network plot. This option is selected by default.<br>• **Map legend** icon ( ): Use this icon to display a legend of icons and line types in the plot, and their meanings.<br>• Zoom icons and icons: Use these icons to gradually zoom in and zoom out the network topology.<br>• **Zoom fit** icon ( ): Use this icon to automatically adjust the topology map to fit within the network plot. |
| 5 | **Network Summary panel**: Provides a tabular summary of various objects in the network. For more information, refer to Network summary tables on page 60. |
| 6 | **Object tabs and tables**: There are numerous tables available for specialized purposes, such as Multicast, P2MP LSPs, and Ports. The tables listed in Common tables are the most commonly used and are the defaults.<br>To bring a table to the forefront and make it active, click its associated tab. For example, to open the **LSPs** table, click the LSPs tab. If the tab you want to use is not visible, then click the Show/hide tables icon ( ) and check the check box corresponding to the required object.<br>The tables display a series of rows and columns, where the rows are objects and the columns are properties. |
| 7 | **Manage add-ons**: Provides options to configure add-ons. For more information, refer to Manage Add-ons on page 389. |
| 8 | **Settings** button ( Settings): Use this button to enable/disable the Auto-resimulate option. For more information, refer to Auto-resimulation on page 84. |
| 9 | **Save views**: Use the **Save view** button to save views of the commonly used tables. For more information, refer to Manage table views on page 64. |
| 10 | **Show/hide tables** icon ( ): Use this icon to show or hide one or more object tables from the Network Summary panel. For more information, refer to Show or hide tables or columns on page 62. |

| Callout No. | Description |
|---|---|
| 11 | **Show/hide table columns** icon ( ): Use this icon to show or hide one or more columns from the object tables. For more information, refer to Show or hide tables or columns on page 62. |
| 12 | **Filter icons**: There are three filter options available in Cisco Crosswork Planning: Floating filter, Advanced filter, and Cross table filter. For more information, refer to Filter options in network summary tables on page 63. |

# Plan tables

Outlines how plan files are structured using tables and guides you to access a complete list of tables and columns through the Schema Definition in the UI.

All aspects of a plan file are defined using a collection of tables. For a complete list of tables and their columns, click > **Schema Definition** in the Cisco Crosswork Planning UI.

## Plan table columns

Details three distinct categories of plan table columns, explaining their identification, purpose, and examples within the schema and UI.

Plan tables contain three distinct categories of columns:

- Key columns: Columns that uniquely identify the rows of the table. Each table has one or more key columns. For example, the <Nodes> table has one key column, Name, which is the unique name of the node. The <Interfaces> table has two key columns: Node and Interface. This pair must (jointly) be unique for all entries in the table. Another example is the <Demands> table, which has key columns: Name, Source, Destination, and ServiceClass. If there are two demands from the same source to the same destination with the same service class, they must have different names.

  On the Schema Definition page, key columns are indicated by an orange line next to their names (refer to Figure 8: Nodes table on page 44).

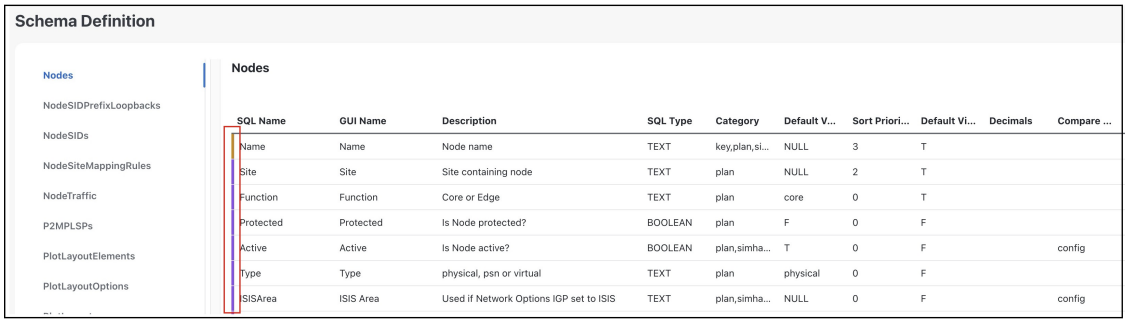
- Plan columns or Collected data: Columns that define or configure properties of entries in a table. For example, in the <Nodes> table, the Site column specifies the site that contains the node, and is therefore a plan column. Key columns are always plan columns.

  On the Schema Definition page, plan columns are indicated by a purple line next to their names (refer to Figure 8: Nodes table on page 44).

- Derived columns or Derived data: Columns that provide information that is derived from the plan columns in the same table or a different table. These are not stored in plan files, but are generated by the UI when the tables are displayed. For example, Remote Node in the <Interfaces> table is derived by looking up the remote node for the interface as defined in the <Circuits> table. Some derived information can be more complex to obtain. For example, the Traff sim column is a derived column that is the result of a simulation performed on the network. The entries in tables generated in the UI can depend on some pre-specified parameters. For example, in the <Interfaces> table, Traff meas column is the measured traffic on that interface for a specified traffic level. For a particular QoS selection the column can be overall (undifferentiated) traffic, traffic on a particular queue, or traffic for a particular service class.

  On the Schema Definition page, derived columns are indicated by a blue line next to their names (refer to Figure 8: Nodes table on page 44).

**Figure 8: Nodes table**




## Table schema and plan tables

Explains the schema and structure of the plan tables, demonstrating column properties and their roles in plan files, with examples for table identification and entry definition.

In a text plan file, each table starts with <TableName> to identify the name of the table. For example, <Nodes>.

The first row of the table body contains the column headings, followed by rows that describe the table entries. The order of the columns is irrelevant, and only the key columns must be present.

Each plan table is defined using an excerpt from the database schema, the part that defines that table. For example, Table 1: NamedPaths table schema on page 45 lists a table schema excerpt for the <NamedPaths> table. Table 2: <NamedPaths> example on page 45 shows an example of a <NamedPaths> table that has been populated within a plan file.

In Table 2: <NamedPaths> example on page 45, the first column is Name, which is described in the first row of Table 1: NamedPaths table schema on page 45. In this case, the Name of the named path (PathA or PathB) is the same in the plan file and UI, it's the name of the path, has a data type of text, and is a table key. Being a key means the Name column is among the columns that uniquely define a row. In this case, the Name and Source together define a unique row.

**Table 1: NamedPaths table schema**

| SQL Name | GUI Name | Description | SQL Type | Category |
|---|---|---|---|---|
| Name | Name | Name of the path | TEXT | key |
| Source | Source | Name of the source of the NamedPath | TEXT | key |
| Active | Active | Is the path active? | BOOLEAN | plan |
| Resolved | Resolved | True if all hops in path are resolved in plan, False if not, NA if no hops. | BOOLEAN | derived |

**Table 2: <NamedPaths> example**

| Name | Source | Active | Resolved |
|---|---|---|---|
| PathA | Router1.AcmeG | T | T |
| PathB | Router1.AcmeO | T | T |

# Assign sites to nodes

Describes the "Assign sites to nodes" initializer, which enables temporary or persistent mapping of network nodes to sites, improving network organization and visualization in Cisco Crosswork Planning.

**Assign sites to nodes** is an initializer in Cisco Crosswork Planning that enables you to:

- Reorganize your network by changing the node-to-site mappings
- Create a node-to-site association based on a simple mapping rule that applies to all nodes. This creates a temporary mapping that is not stored in the plan file.
- Create or customize a Node-to-site mapping table that uses regular expression substitutions to map nodes to sites. Because the table is stored in the plan file (as a <NodeSiteMappingRules> table), you can maintain and reuse it.

Although nodes do not have to be contained in sites, adding them to sites can simplify and improve the visualization of the network. For example, you typically place all nodes in the same geographic location or point-of-presence (PoP) into a single site.

## Site assignment rules

Outlines how the Assign sites to nodes initializer allocates sites to nodes.

This topic describes how the **Assign sites to nodes** initializer assigns sites to nodes.

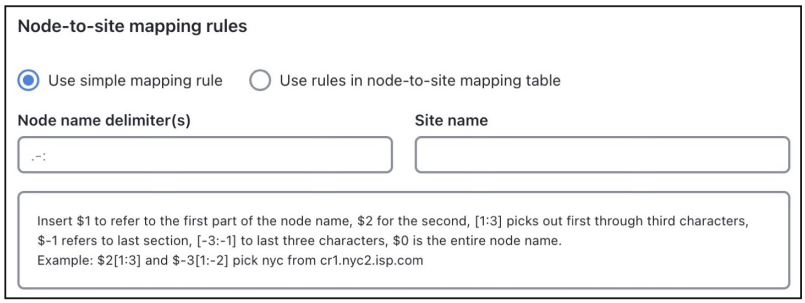
| If a node is ... | Then ... |
|---|---|
| **in an external AS** | Cisco Crosswork Planning assigns it to a site named after the AS name. If the AS name does not exist, Cisco Crosswork Planning assigns the node to a site named after its ASN. |
| **not in an external AS and is not a pseudo-node (PSN)** | Cisco Crosswork Planning assigns nodes to sites based on the Simple mapping rules or on the Node-to-Site mapping rules, depending on what you choose in the **Assign sites to nodes** window. |
| **not in an external AS and it is a PSN** | Cisco Crosswork Planning assigns it to the site that contains the most nodes connected to the PSN. In case of a tie, Cisco Crosswork Planning assigns the node to the site with the lowest lexicographic name. |

### Simple mapping rules

To create temporary node-to-site mappings that are not stored in the network model, click the **Use simple mapping rule** radio button on the **Assign sites to nodes** page.

The mapping is created using two fields:

- **Node name delimiter(s)**: Identifies which sections of the node names to use in the assignments. By default, these characters are a period, a hyphen, and a colon (.-:). New site names are derived from the sections between these characters. For example, by default Cisco Crosswork Planning parses the node name of acme.router into two sections: *acme* and *router*.




- **Site name**: Determines how to create the site names based on the node names. In this list, # equals any integer.
  - $#–Specifies the section reading left to right. For example, $1 matches *chicago* in *chicago.isp*. Note that $0 specifies the entire node name.
  - [#:#]–Specifies the character range reading left to right. For example, $1[1:3] matches *chi* in *chicago.isp*.
  - $-#–Specifies the section name reading right to left. For example, $-1 matches *jose* in *san.jose*.
  - [-#:-#]–Specifies the character range reading right to left. For example, both $2[-4:-1] and $-2[-4,-1] match *jose* in *san.jose.cr1*.

### Node-to-Site mapping rules

To create a Node-to-Site Mapping <NodeSiteMappingRules> table in the network model, particularly for use by the template:

1. Click the **Use rules in node-to-site mapping table** radio button on the **Assign Sites to Nodes** page.
2. Click ⊞ and enter the details.

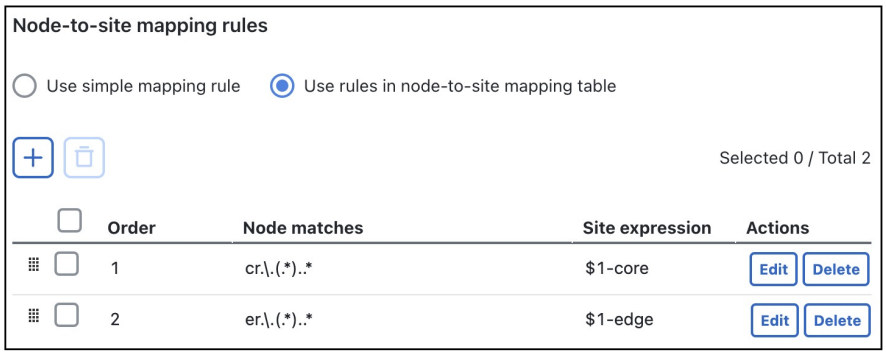
| Column | Description |
|---|---|
| Order | Identifies the order in which rules are applied. |
| Node matches | Regular expression matching the node names. |
| Site expression | Site name expression, which can use references in the Node matches rule. |

An attempt is made to match the node name to expressions in the **Node matches** column. If a match is found, the node is assigned to the site as defined by the corresponding **Site expression**. The order in which these matches are attempted is defined by the **Order** column.

Examples:

- Figure 9: Node-to-site mapping example on page 47 matches cr1.lax into lax-core and er1.lax into lax-edge.

**Figure 9: Node-to-site mapping example**




- This table describes the examples of various combinations of Node matches and Site expression values, and the corresponding results.

| Node matches | Site expression | Result |
|---|---|---|
| cr1.chi.isp.net | chi | Map node cr1.chi.isp.net to site chi. |
| .*\.(.*)\..* | $1 | Map node cr1.chi.isp.net to site chi as above. Additionally, map node cr1.okc.isp.net to site okc. |
| ..(.)\.(.*)..* | $2-$1 | Map node cr1.par.isp.net to site par-1. |

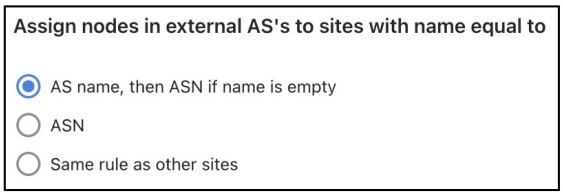
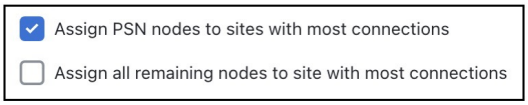
## Run Assign sites to nodes initializer

Guides you through running the Assign sites to nodes initializer, including selecting nodes, configuring mapping rules, and managing unmatched nodes.

Complete these steps to run the **Assign sites to nodes** initializer.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Choose **Actions** > **Initializers** > **Assign sites to nodes**.
3. Choose one or more nodes for which you want to assign sites and click **Next**.
4. Choose which method of node-to-site mappings to use.
   - To use a temporary node-to-site mapping that is not stored in the plan file, click **Use simple mapping rule**.
   - To use a Node-to-Site Mapping table that is saved in the file, click **Use rules in node-to-site mapping table**.
     - To add a new rule, click  and enter the details.
     - To edit an existing rule, click **Edit**.
     - To delete an existing rule, click **Delete**. To delete multiple rules, select the rules and click .
5. For those nodes that the node-to-site definition does not find or cannot create a matching site, you have the option to keep the nodes in their current sites (if applicable) or to remove them from sites. To keep them in their current sites, check the **Keep unmatched nodes in current sites** check box.
6. Choose the options as follows.

| Option | Description |
| --- | --- |
|  | • For nodes that are in external ASes, assign them to sites according to one of these options:<br>  • AS name, then ASN if the name is empty<br>  • ASN<br>  • Using the same rules as other sites |
|  | • Choose whether to assign PSN nodes and all remaining nodes to sites with most connections. |

7. Click **Next** and verify the assignments.
8. When you are satisfied with the mappings, submit the changes.

## Assign geographic locations to sites

Learn how to assign multiple sites to locations using the Cisco Crosswork Planning UI.

Cisco Crosswork Planning includes a database of worldwide city names and airport codes that identify the longitude and latitude of major cities. The **Actions** > **Initializers** > **Assign locations to sites** initializer accesses this database. It lets you quickly lay out sites within a network with geographic precision.

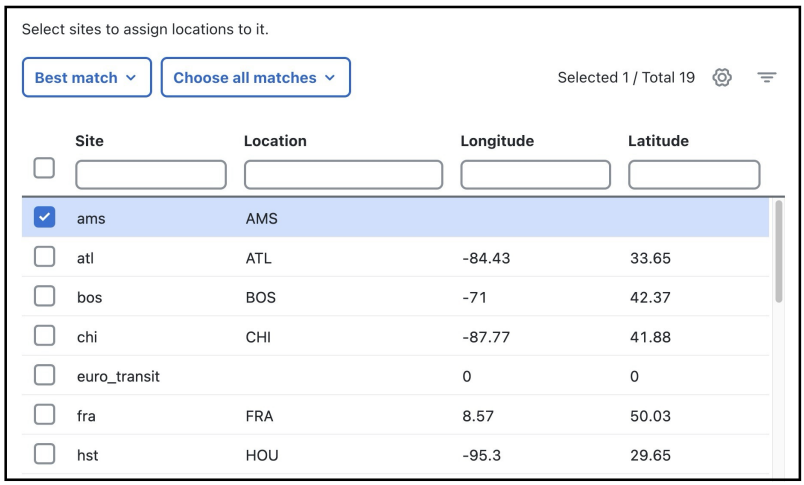
> **Note**
>
> Changing the geographic location of a site does not change the location of its children sites.

Complete these steps to assign multiple sites to locations.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Initializers** > **Assign locations to sites**.

**Figure 10: Assign locations to sites**




3. Set the longitude and latitude values using one of these methods. For each method, Cisco Crosswork Planning fills in the fields with the most closely associated airport or city code in its database. This process assigns longitude and latitude.
   - Select one or more rows from the table. Click either **Best match** > **Best match by site** or **Best match** > **Best match by location**. Cisco Crosswork Planning finds the airport or city codes that most suitably match the site or location.
   - Select a row from the table. Click either **Choose all matches** > **Choose all matches by site** or **Choose all matches** > **Choose all matches by location**. Cisco Crosswork Planning finds all airport or city codes that might be applicable to that site or location. Choose the one you want.

   If you match by site, locations might change based on the match. If you match by location, the site names do not change.
4. Click **Save** to accept the newly assigned locations.

## Set site location

Details the procedure for setting a site's location through the Edit page.

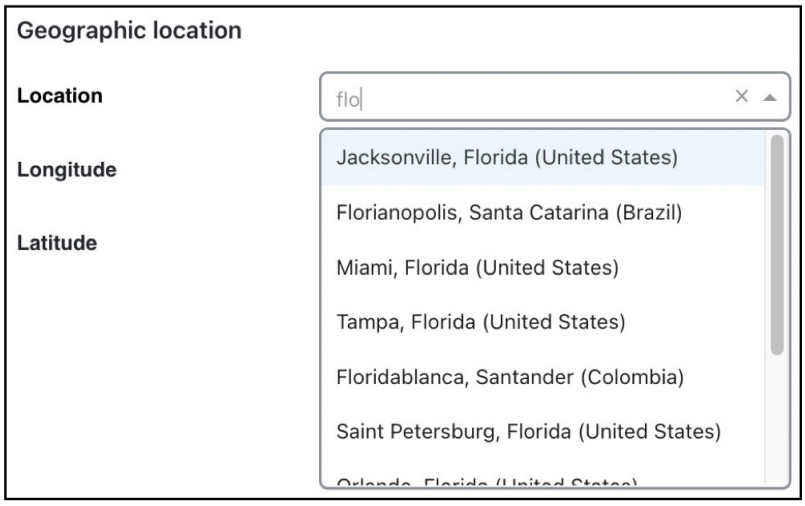
As an alternative to using the **Assign locations to sites** initializer, you can set a location from the Edit window. Use airport codes to improve accuracy.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select one or more sites from the **Sites** table.
3. Click .

   To edit a single site, you can also use the ••• > **Edit** option under the **Actions** column.
4. In the **Location** field, enter a geographic name, such as a city name. The system fills in the **Location** fields with the most closely associated airport or city code in its database. Select the location that best matches your entry.

   **Figure 11: Choosing geographic location**




5. Click **Save** to accept the newly assigned location.

# Network topology

Explains how network plots graphically represent the arrangement and connections of sites and nodes within Cisco Crosswork Planning to support visualization and analysis.

A network topology is a graphical representation that

- displays both sites and nodes within the Cisco Crosswork Planning platform, and
- allows you to visualize the arrangement and interconnections of network elements.

#### Network plot

The network plot, also referred to as the plot, is the main area showing the network topology. To view the network plots, choose **Network Design** from the main menu. The Cisco Crosswork Planning platform displays only the sites in the network plot and related tables, but you can examine these sites in more detail as needed.

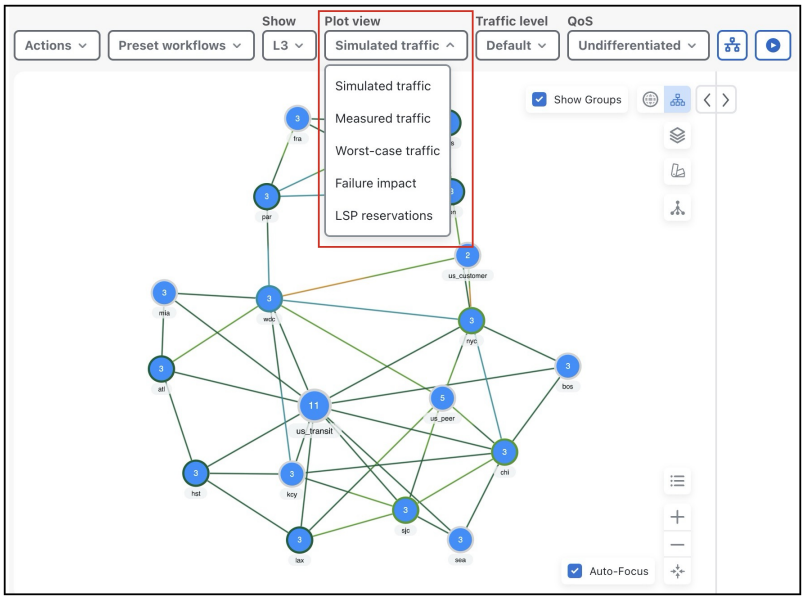
## Network plot views

Lists and describes the available network plot views in Cisco Crosswork Planning, including their features and intended use cases.

This topic describes the different network plot views available in Cisco Crosswork Planning.

The default view is **Simulated traffic**. To change the view, select it from the **Plot view** drop-down list.

**Figure 12: Network plot views**




- **Simulated traffic**: View of simulated traffic that can be used for capacity planning, what-if analysis, and failure planning. Circuits are sized in proportion to their capacity. Interfaces are colored according to the percentage of simulated traffic (in comparison to available capacity, which is defined by the size of the interface or the QoS bound).

  **Note**

  You can observe the colored links only if the auto-resimulation is enabled. To trigger the resimulation manually, click in the toolbar. For more information on updating the auto-resimulation setting, refer to Auto-resimulation on page 84.

- **Measured traffic**: View of traffic as measured from live views of the network. Circuits are sized in proportion to their capacity. Interfaces are colored according to the percentage of measured traffic (in comparison to available capacity, which is defined by the size of the interface or the QoS bound).
- **Worst-case traffic**: Interfaces are colored to show the maximum utilization of the interface over all failure scenarios, based on the most recent simulation analysis. For more information, refer to Evaluate Impact of Worst-Case Failures on page 129.
- **Failure impact**: Interfaces and nodes are colored to indicate how a failure of that interface or node affects other interfaces and nodes. For more information, refer to Evaluate Impact of Worst-Case Failures on page 129.
- **LSP reservations**: Circuits are sized in proportion to their LSP reservable bandwidth. Interfaces are colored according to the percentage of bandwidth (in comparison to what is reserved for the LSP). For information, refer to Configure RSVP-TE Routing on page 263.

## Site plots

Introduces site plots used for topology visualization, outlining their ability to display nodes and links within a site, and differentiate remote connections.

A site plot is a topology visualization that:

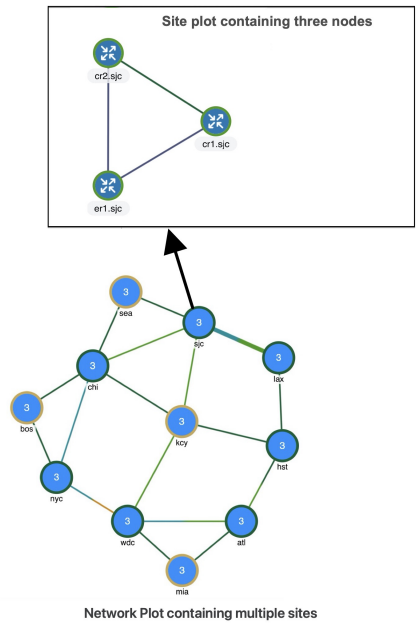
- displays nodes and links within a specific site
- identifies links that connect to remote nodes outside the selected site, and

- differentiates remote nodes from intra-site nodes using unique icons.

Site plots allow you to examine the topology of a specific site. To enhance visibility of network connectivity, site plots can be configured to show pointers to remote nodes, providing a clearer understanding of how a site connects to the broader network infrastructure.

To open a site plot, click a single site from the network plot.

**Figure 13: Site with three nodes**




## Graphical network topology views

Outlines available graphical views for network topology visualization, highlighting logical and geographic views, methods for repositioning nodes or sites, and procedures for persisting layout changes across sessions.

This topic lists the available graphical views for visualizing network topology.

- **Schematic View**: This view shows the network topology, positioned according to an automatic layout algorithm, ignoring their geographical location.

  Click in the network plot to visualize the topology in a schematic view. You can change the layout using . You can drag and move the nodes/sites in schematic view. However, the positions of the nodes/sites are not saved across user sessions.

- **Geographic View**: This view shows network topology, superimposed on a map of the world. Each device's location on the map reflects its GPS coordinates—longitude and latitude—as defined in the device inventory. If GPS coordinates are not available, this view becomes unavailable, and nodes will be placed in pre-determined locations.

  Click in the network plot to visualize the topology in a geographic view.

## Grouping of nodes

Outlines node grouping options in logical and geographic views, describing how nodes are aggregated based on site or location to aid hierarchical analysis of network topology.

Nodes can be grouped in different ways depending on the chosen network topology view:

- In the logical view (), the nodes are grouped based on the sites that they belong to. Use the **Show Groups** check box to group or ungroup the nodes based on sites. Clicking a site replaces the topology map with a new map that shows only the nodes within that site and the links that interconnect them.
- In the geographical view (), the nodes are grouped and ungrouped automatically based on the geographical location.

## Plan object representations

Explains how network plan objects are depicted with distinct visual symbols or icons in the network plot.

Plan objects in the network plot are displayed using distinct visual symbols or icons to help you identify their type quickly. To view all plan object representations in the network plot, click the  icon in the network plot.

For detailed descriptions of the visual representations, refer to Legend for Plan Objects on page 415.

## Set objects to the foreground in network plot

Describes how to emphasize selected objects in the network plot by bringing them to the foreground.

Use this task to highlight specific objects in the topology view for improved visibility and focus.

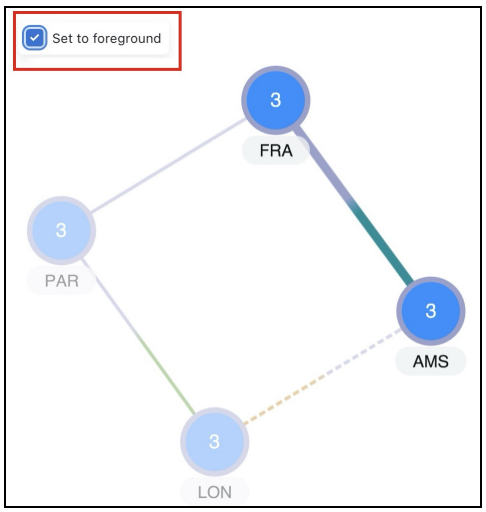
In Cisco Crosswork Planning, you can bring selected objects to the foreground for improved viewing. All other objects move to the background and appear transparent. You can also choose to hide unselected objects completely.

When you set objects to the foreground, they remain highlighted until you deselect the option. Other objects stay de-emphasized or hidden unless you include them as foreground objects. This feature helps you concentrate on relevant elements in complex topologies.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Select one or more objects (nodes, sites, interfaces, circuits, demands, or LSPs) from the network plot or the table.
3. Check the **Set to foreground** check box in the network plot.

   **Figure 14: Set to foreground option**

   Only the selected objects are highlighted in the network plot.
4. (Optional) Choose how to display the selected objects in the foreground. You can highlight the selected objects or hide the unselected items.

a) Click Settings on the Network Design page.
b) In the **Set to foreground** section, select the required option.
   - **Emphasize selected items**: This option highlights the selected items.
   - **Show only selected items**: This option shows only the selected items and hides the unselected items.

   The system saves your preference and next time you use the Set to foreground option on any plan file, this preference will be applied.

---

You can now view the selected objects more easily in the foreground, while all other objects are de-emphasized or hidden.

### What to do next

To remove objects from the foreground, uncheck the **Set to foreground** check box.

## Traffic utilization colors

Explains how traffic utilization colors visually represent bandwidth usage on network interfaces by mapping specific colors to traffic load thresholds, making it easy to monitor, distinguish, and analyze network performance at a glance.

Traffic utilization colors are network visualization indicators that

- visually represent outbound traffic load on interfaces
- use color coding to display bandwidth utilization for traffic leaving the interface proportional to capacity, and
- distinguish traffic thresholds at a glance for monitoring links.

Interfaces can have both measured and simulated traffic (in Mbps). Each link may display two different colors, with each color representing the utilization of a specific unidirectional interface.

> **Note**
>
> You can observe the colored links only when you trigger the resimulation. To trigger the resimulation, click in the toolbar. By default, the auto resimulation is off in Cisco Crosswork Planning. To change this default setting, refer to Auto-resimulation on page 84.

### Utilization color mapping

The colors for traffic utilization appear in the **Utilization colors** option ( ) in the network plot. There are six default colors representing thresholds from 0% to 100%.

This table shows the system's default mapping between link color and traffic utilization.

| Traffic utilization | Color | Example |
|---|---|---|
| Less than 30% | Dark green | |
| 30-50% | Light green | |
| 50-80% | Blue | |

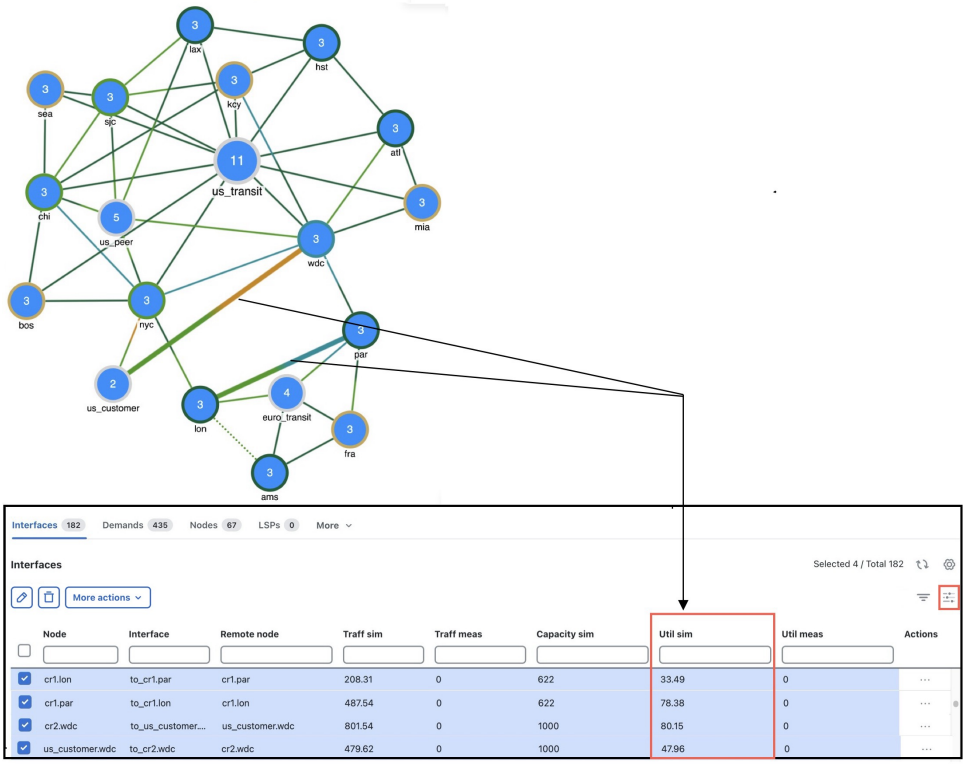
| Traffic utilization | Color | Example |
|---|---|---|
| 80-90% | Light orange | |
| 90-100% | Orange red | |
| More than 100% | Red | |
| 0% | Light slate blue | |
| No traffic | Slate blue | |

Use the **Edit colors** option to modify the traffic utilization values. You can also add or remove colors using this option. When you save the file, your changes are stored in the plan file. For more information, refer to Edit the utilization colors and threshold ranges on page 59.

## Example: Determining traffic utilization

Describes how to configure circuit filters and review utilization values to assess interface traffic between network sites.

In this example, you determine the traffic utilization of the interfaces between the sites, "lon" and "par", and "us_customer" and "wdc".




### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Click the link between the sites. In this example, click the links between "lon" and "par", and "us_customer" and "wdc". The circuits are selected in the Network Summary table.

3. In the Network Summary panel on the right side, choose > **Filter to circuits**. The Circuits table opens and shows only the selected circuits.
4. Select the circuits and choose > **Filter to interfaces**. The Interfaces table opens and shows only the interfaces included in the selected circuits.
5. Notice the utilization values in the **Util sim** column.
   - For the "to_cr1.par" interface, the Util sim value is 33.49%. This value falls within the 30-50% range and is denoted by light green. The interface displays green filling almost half of it. The other half of the circuit ("to_cr1.lon") shows a traffic utilization of 78.38% for the reverse direction. This value is within the 50-80% range and is denoted by Blue.
   - For the "to_us_customer" interface, the Util sim value is 80.15%. This value is within the 80-90% level and is denoted by light orange. The other half of the circuit ("to_cr2.wdc") displays the traffic utilization for the reverse direction. The value for this interface is 47.96%, which is within the 30-50% level and is denoted by light green.
6. To show all the interfaces, clear the applied filter.

## Identify the most highly utilized interface

Outlines procedures for analyzing interface utilization, sorting traffic values, and pinpointing the most congested link in a network plan for performance troubleshooting.

This topic describes how to quickly identify the most utilized interface in a network plan for performance analysis and troubleshooting.

Identifying the most utilized interface helps you analyze large, complex networks for potential congestion and optimize resource allocation.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Interfaces table, click the **Util sim** column heading to sort the interfaces by descending order. Click it two more times. The column sort toggles between ascending and descending order.
3. Notice that the most utilized interface is at the top row. Select this row. The interface highlighted in the plot is the most highly utilized interface.
4. Deselect the row to clear the selection if needed.

You now have identified and visualized the most highly utilized interface in the network plan.

## Utilization color profiles

This topic explains how utilization color profiles allow users to edit link colors for traffic ranges, define threshold profiles per plan, and manage color settings.

Utilization color profile is a setting that

- allows you to edit link colors for different traffic utilization ranges
- defines multiple link utilization coloring threshold profiles per plan file, and
- makes color settings available for each user.

In Cisco Crosswork Planning, there are three coloring profiles.

- System default: This is the default link utilization color profile defined by Cisco Crosswork Planning. When creating a new plan file, this color profile is applied by default.
- Default (plan file default): This is the utilization color profile stored in the default layout of the plan file. If you open an existing plan file, the links are displayed using the default colors stored within the plan file.
- Custom profiles: These are link utilization color profiles created by users. When you apply a new profile and then save the plan file, any user who accesses that plan file can use this saved profile.

### Create a utilization color profile

Guides through creating a custom color scheme for network link utilization, including specifying threshold ranges, assigning colors, and saving new profiles within plan files.

Use this task to create a custom color scheme for traffic utilization ranges in a plan file.

Cisco Crosswork Planning allows you to create new color profiles to visualize network link utilization with custom color settings. You can create up to 10 custom profiles per user.

#### Procedure

1. On the **Network Design** page, click **Settings**.
2. In the **Utilization color profiles** section, click **Create new profile**.
   The Create New Color Profile page appears.
3. Enter the profile name.
4. Specify threshold ranges and assign a color to each range.
   a) To select colors, click the color swatch and select the color from the palette. Alternatively, enter the hex code for your chosen color.
   b) Use **From (inclusive)** and **To** to specify the threshold range for each color.
   c) (Optional) To add a new color, click the plus sign and specify the color and the threshold range.

   > **Note**
   >
   > You must define at least two threshold ranges for a color profile. You can add as many utilization threshold ranges as needed within a profile.

5. Add the new color profile.

   You can now use the new custom utilization color profile for your plan.
6. Apply the color profile to the plan file. For more information, refer to Apply utilization color profiles on page 58.
7. Save the plan file to store the new color profile: From the **Network Design** toolbar, select **Actions** > **File** > **Save**.

The color profile is saved in the plan file and is available for all users who access this plan file.

### Apply utilization color profiles

Provides instructions for applying system default or custom utilization color profiles to the topology map, enabling visualization of link usage through selected color schemes.

After you create a utilization color profile, you can apply it to the topology map. This lets you visualize link utilization using the defined color scheme.

When you open the plan file, the default (plan file default) color profile is applied to the plan file. Complete these steps to apply the System default profile or a custom color profile to your plan file.

### Before you begin

To use the custom color profile, create the required custom profile. For instructions, refer to Create a utilization color profile on page 58.

#### Procedure

1. In the network plot, click the **Utilization color** icon ( ).
2. To apply a custom color profile, select the required custom color profile from the drop-down list.
3. To apply the system default color profile, select **System default**.

The colors of the links in the topology map update to reflect the selected color profile.

#### Edit the utilization colors and threshold ranges

Details steps for modifying colors and utilization threshold ranges in both default and custom profiles, allowing you to tailor visual network indicators according to their preferences.

This topic describes how to modify the colors and threshold ranges for utilization in the non-system default profiles. Adjusting these settings helps tailor visual indicators to your preferences.

> **Note**
>
> You can modify only the default and custom color profiles. You cannot modify the system default color profile.

#### Procedure

1. In the network plot, click the **Utilization color** icon ( ).
2. Select the required profile from the drop-down list.
3. Click **Edit colors**.
4. Update the color and the utilization threshold ranges as required.
5. Save the changes.

The selected color profile reflects your updated colors and threshold ranges.

#### Edit a custom color profile

Explains how to update colors and threshold values in existing custom color profiles, ensuring profiles remain relevant by reflecting current utilization requirements.

This topic describes how to modify the colors and threshold ranges in an existing custom color profile.

#### Procedure

1. On the **Network Design** page, click **Settings**.
2. In the **Utilization color profiles** section, click ••• > **Edit** for the utilization color profile you want to edit.

The Edit Color Profile page appears.

3. Modify the color or the utilization threshold ranges as needed.
4. Save the changes.

---

The selected custom color profile is updated with the new values.

**Support for color profiles during migration**

Describes retention and management of custom utilization color profiles during system backup, restore, and upgrade, and explains default profile settings for plan files migrated to Cisco Crosswork Planning.

This section describes how utilization color profiles are retained or set as default during migration from previous versions.

- Custom color profiles are retained during system backup, restore, or upgrade processes.
- If a plan file created with Cisco WAE Design or earlier versions of Cisco Crosswork Planning versions includes a saved color scheme, the system sets that color scheme as the default profile when the plan file is used in Cisco Crosswork Planning 7.2.x.

## Network summary tables

Introduces network summary tables that display plan objects and their properties on the Network Design page, offering a concise overview of the network's structure and configuration.

A network summary table is a table that

- displays a series of rows and columns, where the rows represent objects and the columns represent their properties
- appears on the right side of the Network Design page, and
- provides a summarized view of network information.

Network models also contain other tables that are not viewable from the UI. These tables contain more complex information, such as showing complex relationships between objects. If you open a network model using a .txt editor, each table is labeled with angle brackets, such as <Nodes> and <Sites>.

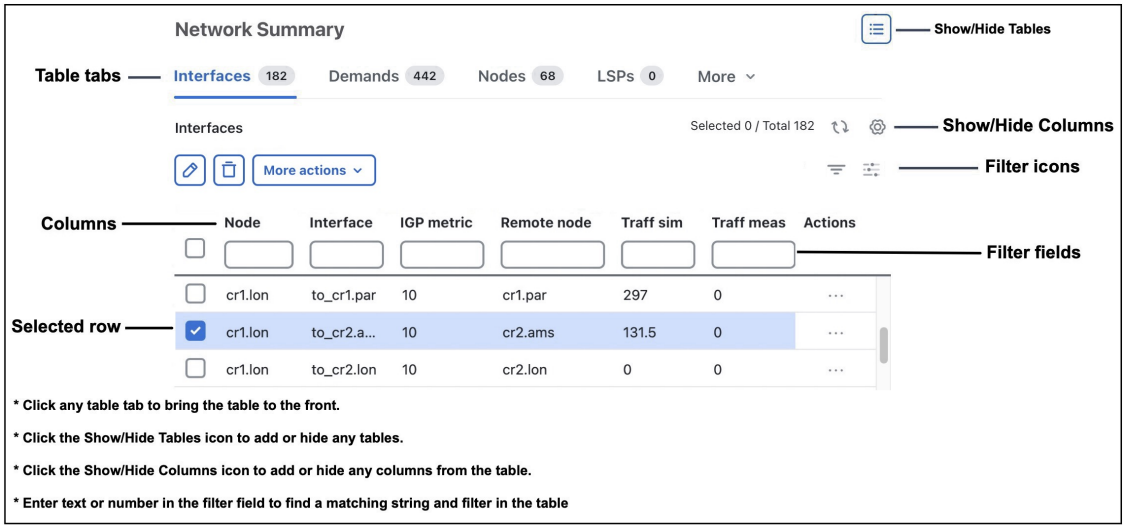
**Common tables**

There are numerous tables available for specialized purposes, such as Multicast, P2MP LSPs, and Ports. These tables are the most commonly used and are the defaults.

- Interfaces: A list of the interfaces in the network.
- Demands: A list of the demands in the network. Each demand specifies how much traffic is routed from a source (node, external AS, or external endpoint) to a destination (node, external AS, external endpoint, or multicast flow destination).
- Nodes: A list of network routers, which typically have names that suggest their location and function. For example, node cr1.atl is a core router 1 in the Atlanta site.
- LSPs: A list of MPLS LSPs in the network; each LSP contains both source and destination. Note that sub-LSPs within point-to-multipoint (P2MP) LSPs are listed in the LSP table.
- Sites: A list of network sites.
- SRLGs: A list of the shared risk link groups (SRLGs). An SRLG is a group of objects that might all fail due to a common cause.

- AS: A list of both internal and external autonomous systems (ASes). An AS is a collection of connected IP routing prefixes that are controlled by one operator.

**Figure 15: Common table functions**




## Work with tables and object selections

Explains the interaction between network summary tables and plan objects, demonstrating how selecting objects highlights them in the network plot and vice versa, and providing steps for selecting and verifying objects across table and topology views.

This topic familiarizes you with the basic dynamics between network summary tables and plan objects, such as nodes, sites, circuits, interfaces, and ports. When you select an object in a table, it is also highlighted in the network plot, and vice versa.

In this section, you show different tables and select objects. Each time you make a selection, observe what happens in the tables and in the network plot to see their relationships.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.

   The Interfaces tab is selected by default in the Network Summary panel on the right side.

2. Click the **Sites** tab (click **More** > **Sites**) to show the Sites table.
3. Select various sites. Notice that the color of the circle around the selected site changes to gray.
4. Show the Circuits table (which is not in the default table tabs) using the **Show/hide tables** icon ( ). For more information, refer to Show or hide tables or columns on page 62.
5. Select the circuits you want. The selected circuits are highlighted in the network plot. If you select circuits in the network plot, the corresponding rows are highlighted in the table.

   When you select two circuits, it selects four interfaces (two interfaces per circuit). To verify these selections, in the Circuits table, ensure that the circuits are selected and click > **Filter to interfaces**. The Interfaces table shows the corresponding four interfaces.

6. To select all objects, check the selection check box in the header row. The network plot highlights the selected objects. Uncheck this check box to deselect all objects.




## Show or hide tables or columns

Describes procedures for displaying or hiding specific tables and columns within the Network Summary panel.

This topic describes how to display or hide tables in the Network Summary panel to tailor the view to your needs. It also describes how to display or hide columns within a table to focus on specific data.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. To show or hide one or more network summary tables:
   a) In the Network Summary panel on the right side, click the **Show/hide tables** icon ( ).
   b) Check the check boxes for the objects that you want to show and uncheck the ones you want to hide. For example, if you want to display the Circuits table, then select the check box next to **Circuits**.

      **Note**

      Use the Search bar at the top to search for the table names quickly.

   c) Click **Apply**.
3. To show or hide one or more columns:
   a) In the Network Summary panel on the right side, select the required Object tab.
   b) Click the **Show/hide table columns** icon ( ).
   c) Check the check boxes for the columns that you want to show and uncheck the ones you want to hide. For example, if you want to display the Capacity column in the table, then select the **Capacity** check box.

      **Note**

      Use the Search bar at the top to search for the column names quickly.

## Sort columns in tables

Guides you to reorder network summary table data by sorting column headings in ascending or descending order, facilitating easier comparison and review of network attributes.

This topic describes how to reorder table data for easier review or comparison.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the required Network Summary table, click the heading of the column to toggle between ascending and descending order.

The table data is reordered based on the selected column and the sort direction.

## Filter options in network summary tables

Outlines available filter options, including floating, advanced, cross table, and selection-based filters, to efficiently refine and display network summary table data according to user criteria.

Network summary tables include multiple filter options to help you select specific values efficiently.

**Table 3: Filter options**

| Filter option | Description |
|---|---|
| Floating filter | Sets the filter criteria on one or more columns in the table. By default, this filter is enabled on all columns.<br>Click to toggle the floating filters at the top of each column.<br>To clear all applied filters, click the X icon in the Filters field that appears above the table. |
| Advanced filter | Filters the table to a list of specific search values.<br>To use this filter:<br>1. Select > **Advanced filter**.<br>2. Select the conditions from the drop-down lists and enter the values.<br>3. If you are filtering by more than one criterion, you can optionally use these controls:<br>• All (AND): (Default) Filters to only those rows that match all the specified criteria (default).<br>• Any (OR): Filters to rows that match any one of the criteria.<br>4. Click **Apply**. |
| Cross table filter | Filters the selection to show its associated objects. For example, you might want to filter a circuit to view its associated interfaces.<br>To use this filter:<br>1. Select the objects from the network summary table.<br>2. Click > **Filter to** *X*, where *X* is the name of the viable object. |
| Filter to selection | Displays only the selected objects in the network summary tables.<br>To use this filter:<br>1. Select the objects from the network summary table.<br>2. Click > **Filter to selection**. |

## Manage table views

Explains how to save, apply, rename, and delete custom table views in network models.

This topic describes how to save, apply, rename, and delete table views in network models to simplify access to frequently used tables.

The network models often contain multiple plan object tables. Cisco Crosswork Planning allows you to display only the tables you need by saving your preferred configurations as views. When you apply a saved view, the network model displays only the selected tables, making it easier to focus on the relevant data.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Use  to show or hide the tables as needed. For more information, refer to Show or hide tables or columns on page 62.
3. Click 

> **Save view** in the top right corner.

   The Save View As page appears.
4. Enter a name and click **Save**.

   The newly created view appears in the **Saved views** list.
5. To apply a saved view:
   a) Click the **Saved views** drop-down list.
   All saved views are displayed.
   b) To apply the selected view to the network model, click the view's thumbnail in the card layout or click the view name in the list layout.
6. After applying a view to a network model, you can rename it. To rename a view:
   a) Ensure that the view is applied (see Step 5).
   b) Click 

> **Rename view**. The Rename View page appears.
   c) Enter a new name for the view and click **Save**.
7. To delete a saved view:
   a) Click 

> **Manage views**.
   All saved views are displayed.
   b) Click  in the view that you want to delete.
   c) Click **Delete** in the confirmation page.

## User-defined columns

This topic explains user-defined columns, customizable fields users can add to network summary tables to organize and manage data without conflicts.

User-defined columns are customizable columns that

- users can create and add to the standard network summary tables
- are prefixed with a namespace grouping columns under a single name, and
- help to prevent conflicts with plan table names or existing plan columns.

#### Properties of user-defined columns

- User-defined column headings have the *Prefix*::*Name* format.
- User-defined column names must be unique.
- Once added, only the value of the user-defined column is editable.
- User-defined columns are hidden on the network summary page. To view these columns, manually select the desired columns each time you open the plan file.

#### Network summary tables supporting user-defined columns

This table lists the network summary tables (plan objects) which allow you to create user-defined columns.

| | | |
|---|---|---|
| Interfaces | P2MP LSPs | Ports |
| Circuits | Sites | Port Circuits |
| Demands | SRLGs | VPN |
| Nodes | AS | VPN Nodes |
| LSPs | External Endpoints | |
| LSP Paths | Interface Queues | |

**Note**

You can add a maximum of 50 user-defined columns to each supported plan object.

#### Add user-defined columns

Provides instructions for adding user-defined columns to network summary tables in plan files

This topic describes how to add user-defined columns to plan file objects, enabling you to capture additional information specific to your planning needs.

### Before you begin

You can create user-defined columns only in the network summary tables listed in Network summary tables supporting user-defined columns on page 65.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Show the network summary table in which you want to add the column.
3. Select the row you want to customize and click .
4. On the Edit page, click the **User** tab.
5. Click .
6. Enter the prefix, name, and select the type (text, real, integer, or boolean).

### Example

Table 4: User-defined column example on page 66 shows a user-defined column that was added to the **Sites** table. The first three columns are standard columns. The last is a user-defined column where the Prefix is Customer and the Name is Service. Combined, they form the user-defined column named Customer::Service. The values Voice and Video were entered on their respective rows.

**Table 4: User-defined column example**

| Name | Longitude | Latitude | Customer::Service |
|---|---|---|---|
| EMEA_SW | 2.55 | 49.02 | Voice |
| EMEA_Central | 4.78 | 52.32 | Video |

7. Enter a value matching the column type, and click **Save** under the **Actions** column.
8. (Optional) Repeat steps 4 through 6 for each user-defined column you want to add.
9. Click **Save** to save the changes.

---

The newly created column is added to the plan file. To view these columns, click ⚙ and check the *Prefix*::*Name* check box.

### What to do next

After you add a user-defined column to a plan file, you can manipulate it in the Cisco Crosswork Planning Design UI. You can

- filter them by prefix, name, type, or value in the **User** tab
- sort them by prefix, name, type, or value in the **User** tab, and
- delete them when they are no longer needed.

## Enable dark mode

Learn how to enable dark mode in the Cisco Crosswork Planning UI.

The dark mode changes the appearance of the UI from a light theme to a darker color scheme. This feature enhances readability in low-light environments. It also reduces eye strain and provides an alternative look and feel.

Complete these steps to enable dark mode in the Cisco Crosswork Planning UI.

### Procedure

---

1. Log in to the Cisco Crosswork Planning UI.
2. Click 👤 in the top-right corner of the main page.
3. Select **Classic dark**.

---

The UI color theme switches to dark mode.

### What to do next

To return to light mode, click 👤 > **Classic light**.

# 3 Understand Plan Objects

**Topics:**


- Plan objects
- Nodes and sites
- Circuits and interfaces
- SRLGs
- Ports, port circuits, and LAGs
- Key operations you can perform on objects


Describes the management of plan objects, including nodes, sites, circuits, ports, SRLGs, and LAGs, and guides users through creation, configuration, and deletion procedures critical to network planning and operations.

# Plan objects

Explains the types of objects in Cisco Crosswork Planning networks and how their properties are managed and visualized.

Cisco Crosswork Planning networks consist of objects, such as nodes (which represent routers), interfaces, circuits, SRLGs, LSPs, ports, and port circuits. A site is also an object, and is a Cisco Crosswork Planning construct for simplifying the visualization of a network by grouping nodes within a site, or even by grouping sites within a site.

Most objects are represented in the network plot, and all of them are represented in Network Summary tables on the right side. They have *properties* that identify and define them, many of which are discovered. They can also be manually added and changed. For example, all circuits have a discovered **Capacity** property that can be edited. Other properties are derived. For example, **Capacity sim** is derived from the **Capacity** property. Another example is that interfaces have a **Util sim** property that identifies the percentage of **Capacity sim** the simulated traffic is using. Properties are viewable and editable through the Edit window. These are represented by columns in the object's table, or by entries in tables of related objects.

Cisco Crosswork Planning has a Layer 3 (L3) view containing objects. Throughout this guide, the terms *node* and *circuit* refer to objects in the L3 view.

# Nodes and sites

Explains how nodes and sites organize network components for simulation and traffic analysis within the product. It clarifies the distinctions between nodes and sites, their roles, and visual representations.

A node or site is a network modeling construct that

- represents devices (nodes) and hierarchical groupings (sites), and
- provides visual cues and organizational structure for network planning and analysis.

**Node**: A node is a device in the network, which can be one of three types: physical, pseudonode (PSN), or virtual. The **Type** property distinguishes whether the node represents a real device or an abstraction, such as a single node representing a number of edge nodes connected in the same way to the network. A physical node is a Layer 3 device, or router. Physical and virtual nodes behave in the same way within Cisco Crosswork Planning. A pseudonode is typically represents a Layer 2 device or a LAN.

Nodes can reside both inside and outside of sites. The external arrangement could be useful for small networks where routers are not geographically dispersed.

**Site**: A site is collection of nodes and/or other sites, potentially forming a hierarchy of sites. Any site that contains other sites is called the *parent* site.

Both nodes and sites have simulated traffic, while nodes also have measured traffic.

### Network plot visualization

- Nodes appear as blue router icons (  ) in the network plot. The border color indicates the traffic sourced from and destined to the node. A light blue outline indicates that the node is selected. For more information, click the icon in the network plot.
- Sites appear as blue circles (  ) in the network plot. The border color indicates the traffic utilization of all the nodes and circuits inside the site, including all nested nodes and circuits. A light blue outline indicates that the site is selected. The number inside the circle indicates the number of nodes in the site. For more information, click the icon in the network plot.

  Note that sites can contain L3 nodes. Empty sites and sites containing L3 nodes appear in the L3 view.

## Parent sites and contained objects

Explains how parent sites organize network hierarchies and the nature of contained (nested) objects within them.

A parent site is a site that

- contains other sites, nodes, or circuits (known as contained objects)
- defines hierarchical relationships between sites through nesting, and
- enables filtered views and management of all objects within its hierarchy.

Contained or nested objects are nodes, sites, or circuits that reside within a parent site. These are also referred to as children and are often geographically co-located. For example, a site representing a Point of Presence (PoP) can be a parent site if it contains routers.

The site's **Parent site** property defines whether it is nested within another site. If it is empty, the site is not nested.

### Network plot visualization

- In the network plot, the parent site shows all egress inter-site interfaces of all nodes contained within it, no matter how deeply the nodes are nested. This visualization logic applies to each child site plot contained within the parent.
- Selecting a site from the network plot does not select the sites or nodes under it.

### Filtering to contained objects

To filter to the sites or nodes contained within a site, select the site and click  > **Filter to contained**.

- **Direct**: filters to only those objects contained immediately within it.
- **All**: filters to all objects contained within it, no matter how deeply they are nested.

### Site deletion options and behavior

Provides details on the behavior and options available when deleting a site, including how contained sites and nodes are handled.

When deleting a site from the network plot, all objects directly within the site are selected for deletion by default. However, you can choose to retain the contained sites and nodes in the confirmation dialog. If you opt to retain them, these actions occur:

- Objects directly contained within the deleted site are moved to the same level as the site that was removed.
- Sites and nodes that are more deeply nested maintain their original parent relationships.

## PSN nodes

Details the characteristics and modeling considerations of PSN nodes, including their role in IGP and BGP peer modeling, creation criteria, naming, and best practices for using PSN nodes in planning and simulation.

A PSN is a node that

- represents pseudonodes, which are used to model LANs or switches that connect more than two routers
- is used for IGP modeling and BGP peer modeling, and
- has a **Type** property of **PSN**.

In IGP networks, a LAN interconnecting multiple routers is represented by a PSN node, with circuits connected to each of the nodes representing the interconnected routers. Both OSPF and IS-IS have a built-in system. In this system, one of

the routers on the LAN is the designated router (DR) for OSPF or the designated intermediate system (DIS) for IS-IS. The PSN node uses the name of this designated router. Cisco Crosswork Planning creates nodes with a property Type of PSN automatically during IGP discovery.

When BGP peers are discovered, Cisco Crosswork Planning might find that a router is connected to multiple peers using a single interface. This is typical at switched Internet Exchange Points (IXPs). Cisco Crosswork Planning then creates a node with a property Type of PSN, and connects all the peers to it, each on a different interface.

#### Points to consider when working with PSN nodes

Consider these points when working with nodes that have PSN as a Type property:

- Two PSNs cannot be connected by a circuit.
- If a PSN node is created by Cisco Crosswork Planning, “psn” is added as a prefix to the designated router’s node name.
- When creating demand meshes, Cisco Crosswork Planning does not create demands with nodes of Type PSN as sources or destinations. Creating demands with nodes of Type PSN as sources or destinations is possible in manual demand creation, but this is not recommended. Cisco Crosswork Planning sets the IGP metric for all egress interfaces from a node of Type PSN to zero. This ensures that the presence of a PSN in a route does not add to the IGP length of the path.

## Node configuration fields

Describes node configuration parameters such as naming, addressing, site and AS assignment, BGP settings, and so on.

These fields are commonly used when configuring a node.

- Name: Specifies the name for the node.
- IP address: This value is often the loopback address used for the router ID.
- Site: Specifies the name of the site in which the node exists. If left empty, the node resides in the network plot. This method provides a convenient way to create a site when creating the node, move nodes from one site to another, or remove nodes from a site so it stands alone in the network plot.
- AS: Specifies the name of the AS in which this node resides, identifying its routing policy. This field can be left empty if no BGP is being simulated.
- BGP ID: Specifies the IP address that is used for BGP.
- Function: Indicates whether this is a Core or Edge node.
- Type: The node type, which is physical, PSN, or virtual. Because a PSN node represents a Layer 2 device or a LAN, interfaces on a PSN must all have their IGP metrics set to zero, and two PSN nodes cannot be directly connected to one another. If you change a node type to PSN, Cisco Crosswork Planning automatically changes the IGP metrics on its associated interfaces to zero.
- Longitude and Latitude: Indicates the geographic location of the node within the network plot. These values are relevant when using geographic backgrounds.

To create nodes, refer to Create objects on page 77, where *Object* is **Node**.

## Site configuration fields

Presents the primary configuration fields for sites, including naming conventions, parent site assignment, and display options.

These fields are commonly used when configuring a site.

- Name: Specifies the unique name for the site.
- Display name: Specifies the site name that appears in the plot. If this field is empty, the Name entry is used.
- Parent site: Indicates the site that immediately contains this site. If empty, the site is not contained within another one.
- Location: Select the location from the list of cities. To automatically place a site in its correct geographic location and update the Longitude and Latitude fields, enter the airport code and press Enter.
- Longitude and Latitude: Specifies the geographic location of the site within the network plot. These values are relevant when using geographic backgrounds.

To create sites, refer to Create objects on page 77, where *Object* is **Site**.

## Merge nodes

Explains the process for merging multiple nodes into a single virtual node, detailing methods, effects on circuits and demands, naming strategies, and simulation considerations to simplify core network design.

This topic describes how to merge multiple network nodes into a single virtual node to simplify network planning and accelerate calculations.

Real network topologies often include several nodes—typically edge nodes—that are connected to the network in the same way. For example, all nodes might connect to the same core node or to a pair of core nodes. For planning and design of the network core, it is often desirable to merge these physical nodes into a single virtual node, which simplifies the plan and accelerates the calculations and simulations performed.

> **Note**
>
> Merging nodes changes the actual plan, not just the visual representation.

The name of a newly merged node can be based on the site name, selected node name (base node), or have a new user-specified name. The effects of merging nodes are:

- Reattach circuits from other nodes to the base node.
- Move demands to or from other nodes to the base node.
- Move LSPs to or from other nodes to the base node.
- Set base node traffic measurements to the sum of measurements of the selected nodes.
- Delete other nodes.

Complete these steps to merge nodes.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions > Initializers > Merge nodes**.
3. Select the nodes that you want to merge in the Merge Nodes page. If you do not select any nodes, Cisco Crosswork Planning merges all nodes.
4. Click **Next**.
5. Select whether to merge the nodes per site or merge them into one node.

- Separate merge per site: Merges nodes on a per-site basis. For example, if you selected all nodes in the plan, the result would be one merged node per site. If you do not specify a new suffix, the default name is the same as the site.
- Merge all nodes together into base node: Merges all nodes selected into one node. For example, if you selected two nodes in one site and three nodes in another, the result would be a single node in the site and node combination selected as the base node.

  If you do not specify a new name, the default is to use the name of the base node.
- Merge nodes table: Merges nodes based on the file containing <MergeNodes> table. You can choose this file either from the user space or local machine.

6. Click **Next**.
7. Review the effects of node merge. If they are acceptable, then click **Merge**.

---

The selected nodes are merged into a single virtual node.

# Circuits and interfaces

Outlines the relationship between circuits and interfaces, including logical and physical mappings, circuit connections, interface properties, and traffic measurement.

In Cisco Crosswork Planning, an interface is either an individual logical interface or a LAG logical interface. If there is a one-to-one mapping between a logical and physical interface, then the interface contains both Layer 3 properties (for example, Metric) and physical properties (for example, Capacity). If there is a one-to-many mapping between logical and physical interfaces, then the interface is the logical LAG and the ports are included in the plan file as the physical ports in the LAG. For more information on ports and port circuits, refer to Ports, port circuits, and LAGs on page 75.

Each circuit connects a pair of interfaces on two different nodes. Therefore, an interface always has an associated circuit. Both the Edit Interface and Edit Circuit windows let you simultaneously edit properties for the pair of interfaces and the circuit.

Interfaces have both measured and simulated traffic. The traffic that appears in the Circuits table is the higher of the traffic in the two interfaces.

## Circuit and interface configuration fields

Outlines key configuration fields for circuits and interfaces, including capacity, SRLGs, and parallel group naming.

These fields are commonly used when configuring a circuit or an interface.

- Capacity: Specifies the total traffic this circuit can carry. The drop-down displays the most widely used capacities.
- SRLGs: Assigns the circuit to an SRLG by selecting from the list or creating a new one. For more information, refer to Create SRLGs for circuits only on page 74.
- Parallel group name: Enter the name to include this circuit in a new or existing parallel group.
- Interface A and B: Identifies the two interfaces connected by the circuit.

To create circuits, refer to Create objects on page 77, where *Object* is **Circuit**. As a result of creating a circuit, two interfaces are also created.

### Merge circuits

Explains how to merge circuits with identical endpoints to streamline network design, detailing procedures for selecting, reviewing, and consolidating parallel circuits during long-term capacity planning.

This topic describes how to merge circuits between the same endpoints to provide a simplified network design.

You can simplify the plan by merging circuits that have the same source and destination endpoints (nodes). This capability is useful in long-term capacity planning, where multiple parallel circuits can be ignored and only the site-to-site connections are of interest.

Complete these steps to merge circuits.

#### Before you begin

- Identify the plan file containing the circuits to be merged.
- Assess the implications of merging circuits. This operation changes the plan itself, not just the visual representation.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Initializers** > **Merge circuits**.
3. Select the circuits that have the same source and destination endpoints.
4. Review the listed effects of circuit merge.
   - Sets base circuit capacity to the sum of capacities.
   - Sets base circuit metric to minimum of metrics.
   - Sets base traffic measurements to the sum of measurements.
   - Deletes other circuits.
5. If these effects are acceptable, click **Submit**.

The selected circuits are merged and the plan file is updated.

## SRLGs

Outlines what SRLGs are and their purpose in network design.

A Shared Risk Link Group (SRLG) is a network grouping mechanism that

- contains network elements, such as links or circuits, that may all fail simultaneously due to a single cause, and
- is commonly used to improve network resiliency by identifying points of shared risk.

#### Example

An SRLG may contain all the circuits whose interfaces belong to a common line card. If the line card fails, all circuits in that SRLG will be affected at once.

### Create SRLGs

Guides you through the process of creating SRLGs in your network plan.

Complete these steps to create SRLGs.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Create SRLGs using one of these methods:
   - From the toolbar, choose **Actions** > **Insert** > **SRLG**.
   - In the Network Summary panel on the right side, navigate to the **SRLGs** tab and then click +.

     For more information on showing or hiding network summary tables, refer to Show or hide tables or columns on page 62.
3. Enter the SRLG name.
4. From the **Object** type drop-down list, choose the type of object that you want to include in the SRLG.
5. For each object to include in the SRLG, check the check box under the **Included** column.
6. Click **Add**.

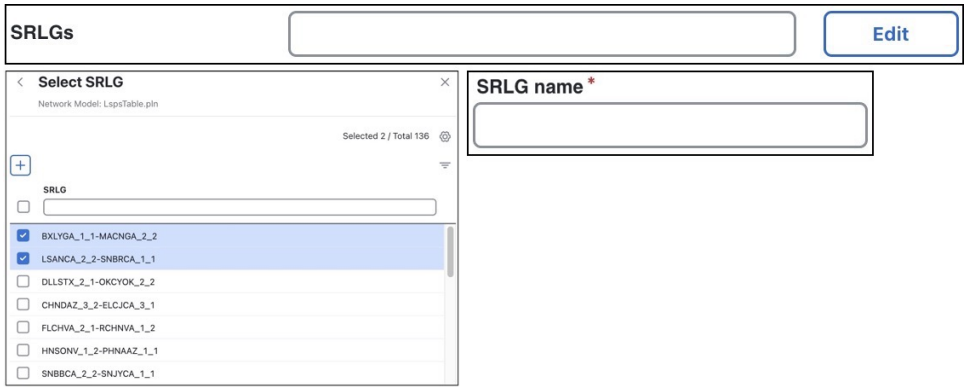
## Create SRLGs for circuits only

Outlines the relationships and configuration procedures for circuits and interfaces, covering logical and physical mappings, and essential configuration fields.

Complete these steps to create SRLGs for circuits.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Open the Add/Edit Circuits page using any of these methods:
   - To create a new circuit, use the steps in Create objects on page 77, where *Object* is **Circuit**.
   - To create SRLGs for selected circuits, select one or more circuits from the **Circuits** tab. Then, click ✎.
3. Click **Edit** button associated with the **SRLGs** field.




4. Associate the selected circuits with one or more existing SRLGs. Alternatively, create a new SRLG.

| To associate circuits with an existing SRLG ... | To create a new SRLG ... |
|---|---|
| a. For each SRLG in which you want to include the selected circuits, check the check box.<br>b. Save your changes. | a. Click +.<br>b. Enter the new SRLG name, and click **Save**. |

5. Click **Add** or **Save** in the Add or Edit Circuit page, as appropriate.

## Ports, port circuits, and LAGs

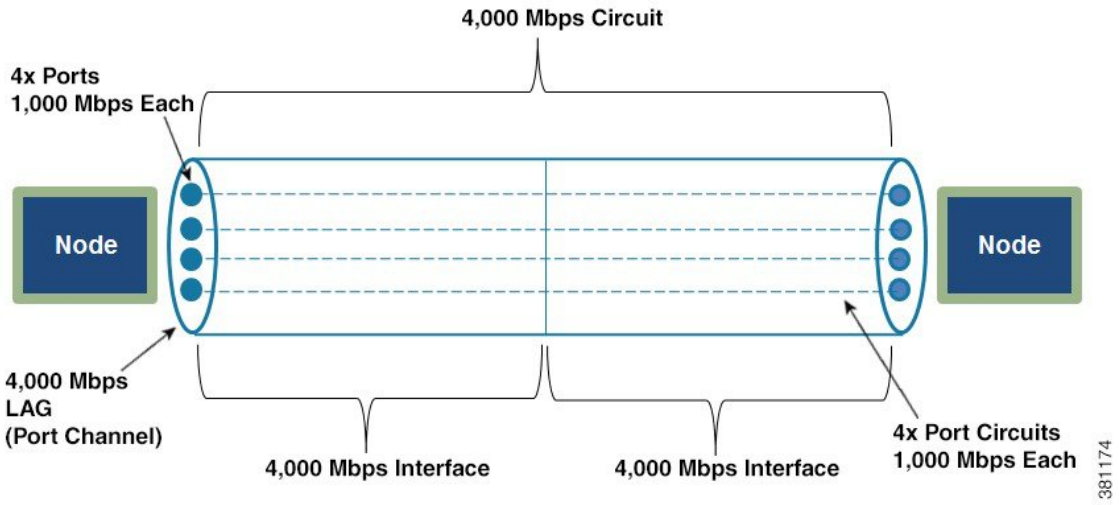
Describes the relationships between physical ports, port circuits, and LAGs in the system, including relevant configuration details.

Ports serve as physical interfaces. Port circuits define the logical paths provisioned over those connections. Link aggregation groups (LAGs) bundle physical interfaces to provide increased bandwidth and redundancy for efficient network management.

- **Port**: A physical interface on a device. In Cisco Crosswork Planning, ports and port circuits are used to model LAGs and port channels.
- **LAG**: A group of physical ports that are bundled into a single logical interface. It is also known as "bundling" or "trunking".

  By default, each logical interface listed in the Interfaces table corresponds to a single physical port. These ports need not be explicitly modeled. The exception is when the logical interface is a LAG, which bundles more than one physical port. In this case, the physical ports are listed in the Ports table.
- **Port circuit**: A connection between two ports. Ports are not required to be connected to other ports by port circuits.

**Figure 16: Ports, port circuits, and LAGs**




## Port configuration fields

Outlines key fields used in port configuration.

These fields are commonly used when configuring a port.

- Name: Name of the port.
- Site and Node: Site and node on which this port exists.
- Interface: Logical interface to which this port is mapped. This must be defined to create port circuits using this port.
- Capacity: The total traffic amount this port can carry. The drop-down list displays the most widely used capacities.

To create ports, refer to Create objects on page 77, where *Object* is **Port**.

## Port circuit configuration requirements

Outlines port circuit configuration requirements.

This topic provides requirements for configuring a port circuit, including required fields.

### Configuration requirements

To ensure successful configuration of a port circuit, these requirements apply:

- Both ports must exist and be mapped to interfaces that are connected by a circuit.
- When selecting two ports for the port circuit, if one port is assigned to an interface, the other must be assigned to the remote interface on the same circuit.

### Configuration fields

These fields are commonly used when configuring a port circuit.

- Site and Node: Site and node on which the port exists.
- Port: Name of the port.
- Capacity: The total traffic amount this port circuit can carry. The drop-down list displays the most widely used capacities.

To create port circuits, refer to Create objects on page 77, where *Object* is **Port circuit**.

### Create LAGs

Explains how to create a LAG by assigning ports to an existing interface in your network plan.

Complete these steps to make one of the existing interfaces into a LAG by assigning ports to it. If the interface does not contain ports, you must first create them (refer to Port configuration fields on page 76).

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the **Ports** network summary table, select all ports that belong to the LAG. You can quickly do this by filtering to their common interface.
3. Select one of the ports and click  to open the Edit page.
4. Select the interface that contains these ports.
5. Save the changes.

### Set LAG simulation properties

Guides you through configuring LAG simulation properties by setting minimum port and capacity parameters.

Each interface is treated as a LAG (port channel). If a LAG loses too much capacity because of non-operational ports, you can configure its properties to automatically take the entire LAG down.

Complete these steps to configure LAG properties.

**Procedure**

1. Open the Edit page of an interface or circuit (refer to Edit objects on page 78).
2. Click the **Advanced** tab.
3. In the **Port channel** section, set one or both of these parameters for one or both interfaces on the circuit:
   a) In **Min number of ports**, enter the minimum number of ports that must be active and operating for the LAG circuit to be up.
   b) In **Min capacity**, enter the minimum capacity that must be available for the LAG circuit to be up.
4. Save your changes.

## Key operations you can perform on objects

Outlines the available operations for managing plan objects, providing guidance on how to perform actions such as creation, modification, and deletion.

These sections describe the operations that you can perform on plan objects.

### Create objects

Guides the process of adding objects to a plan file, detailing methods for object creation, required property entry, and integration into the network design environment.

This topic describes how to add objects to a plan file.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Create objects using one of these methods:
   - From the toolbar, choose **Actions** > **Insert** > *Object*.
   - In the Network Summary panel, navigate to the *Objects* tab and then click .

     The required *Objects* tab may be available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon () and check the relevant *Object* check box.
3. Enter the required details. The properties differ for each object.
4. Add the object.

The object is added to the plan file.

## Edit objects

Explains how to modify plan object properties, covering single and bulk editing methods, workflow navigation, property management, and saving updated configurations within the network design page.

This topic describes how to edit objects in a plan file.

Most objects in the plot and in the associated tables have a set of properties that you can manage using the Edit page. The system uses these properties to define and simulate an object.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel, navigate to the desired *Object* tab, and edit the object properties.
3. To edit a single object, use one of these methods:
   - Select the required object and click .
   - In the **Actions** column, click ••• > **Edit** for the object whose properties you want to edit.
4. To edit multiple objects (bulk edit), select the required objects and click .

   **Note**

   In case of a bulk edit operation, you can
   - select all objects and edit them
   - filter objects, then select and edit them, or
   - select objects from the network plot and edit them.
5. Edit the properties as required.

6. Save the changes.

The selected objects are updated with the new properties.

## Delete objects

Provides instructions for removing objects from a plan file, describing procedures for individual and bulk deletion, navigation steps, and confirmation messages to update the network design.

This topic describes how to delete objects from a plan file.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel, navigate to the desired *Object* tab, and delete the objects.
3. To delete a single object, use one of these methods:
   - Select the required object and click .
   - In the **Actions** column, click ··· > **Delete** for the object whose properties you want to delete.
4. To delete multiple objects (bulk delete), select the required objects and click .

The object deletion successful message appears.

# Part II

## Simulate Your Network

**Topics:**


- Simulation Overview
- Simulate Traffic Flow from Source to Destination Using Demands
- Perform What-If Analysis
- Evaluate Impact of Worst-Case Failures
- Evaluate Impact of Traffic Growth
- Perform Capacity Planning
- Simulate IGP Routing Protocol
- Simulate BGP Routing
- Simulate Quality of Service (QoS)
- Simulate VPN
- Simulate Advanced Routing with External Endpoints
- Simulate Multicast

# 4 Simulation Overview

**Topics:**


- Network simulations
- Auto-resimulation
- States of plan objects
- Simulated capacity
- Delay


Outlines network simulation use cases, automated resimulation processes, and the management of plan object states.

## Network simulations

Explains how network simulations enable predictive analysis, capacity planning, and forecasting for network infrastructure.

Network simulation is a predictive analysis technique that

- enables what-if analysis to assess the outcomes of changes in the network model
- supports capacity planning through simulations
- helps with forecasting by projecting growth percentages onto demands, and
- calculates demand routings and traffic distributions throughout the network based on the given traffic demand, network topology, configuration, and state.

Simulation is the fundamental capability of Cisco Crosswork Planning on which most of the other tools are built, including those for planning, traffic engineering, and worst-case failure analysis. The system supports a range of protocols and models, such as IGP, MPLS RSVP-TE, BGP, QoS, VPNs, and Multicast.

**Examples**

- What-if analysis: You can examine what happens if you change any aspect of the network model, such as:
  - What happens if a link or a node fails?
  - What happens if you change a metric?
  - What happens if you change the topology?

  For more information, refer to Perform What-If Analysis on page 121.
- Capacity planning with resiliency analysis: You can simulate what happens if a node, SRLG, LAG, or a site fails. Cisco Crosswork Planning has the **Simulation analysis** tool to automate this process and provide the analysis. Running the tool displays "worst-case" scenarios that highlight areas most at risk of congestion. The "Failure impact" view displays failures causing the worst case. For more information, refer to Evaluate Impact of Worst-Case Failures on page 129.
- Capacity planning and forecasting: Using the **Create growth plans** tool, you can apply a growth percentage to a demand or set of demands and project that growth into the future. For more information, refer to Evaluate Impact of Traffic Growth on page 147.

## Auto-resimulation

Describes how auto-resimulation automatically triggers simulation when network routing changes occur.

Auto-resimulation is a feature that

- automatically triggers resimulation when any changes affect routing in a network, and
- keeps simulation results updated in response to ongoing modifications.

Auto-resimulation is disabled by default in newly opened plan files.

**Examples of events that trigger auto-resimulation**

- Changes in topology, such as adding and deleting objects or changing explicit paths
- Changes to an object state, such as failing an object or making it inactive

- Changes in numerous properties, such as metrics, capacities, and delay

#### Enable auto-resimulation by default

Complete these steps to enable auto-resimulation by default.

1. Click Settings at the top right corner.
2. Enable the **Auto-resimulate** slider under the **Simulator** section.

   After you update this setting, it applies only to the current plan file. Each time you open this particular file, the same setting is used.

#### Run simulation manually

When any change is made in the plan file that affects or invalidates the current simulation, you can trigger a re-simulation manually. To do this, click the  icon in the **Network Design** page.

## States of plan objects

Lists and describes the states of plan objects, including their definitions and how these states appear in the network summary tables.

The state of a plan object affects the simulation and determines whether that object is operational.

There are three states available in Cisco Crosswork Planning plan objects.

- **Failed**: Identifies whether the object is failed.
- **Active**: Identifies whether the object is available for use in the simulated network. For example, an object might be unavailable because it has been set administratively down.
- **Operational**: Identifies whether an object is operational. For example, an object might be non-operational because it is failed, is inactive, or because other objects on which it depends are not operational.

#### Columns in network summary tables

The **Failed** and **Active** columns in the network summary tables show a visual representation of their status. Likewise, the **Operational** column shows the calculated operational state. In each column, "true" indicates that the object is in that state. "false" indicates that it is not. The value of "true" or "false" for the Active state in the Interfaces table reflects the status of the associated circuit. The plot shows graphical representations of these states with either a white cross or a down arrow displayed inside a red circle.

These plan objects have the Active, Failed, and Operational columns:

- Circuits
- Nodes
- Sites
- Ports
- Port circuits
- SRLGs
- External endpoint members

## Failed state

Outlines the Failed state in network simulations, including the effects of interface and circuit failures on traffic, rerouting of demands, visual indicators, and options for disabling rerouting to manage L2 and L3 traffic behavior.

A Failed state is a condition in network simulations where objects such as interfaces or circuits experience failure.

The quickest way to see the effects of failures in a simulated traffic view is to have demands in place, and then fail an object. When a failure occurs, the simulation reflects the changes immediately.

- The plot immediately displays where the traffic increases as a result (Figure 17: Failed circuit on page 86)
- The **Util sim** column in the Interfaces table reflects the traffic changes, and
- Demands are then rerouted around the failure (Figure 18: Reroute of demand around a failed circuit on page 86).

If you select an interface to fail, you are actually failing its associated circuit. To view the complete list of objects that can be failed, refer to the list in States of plan objects.

**Figure 17: Failed circuit**

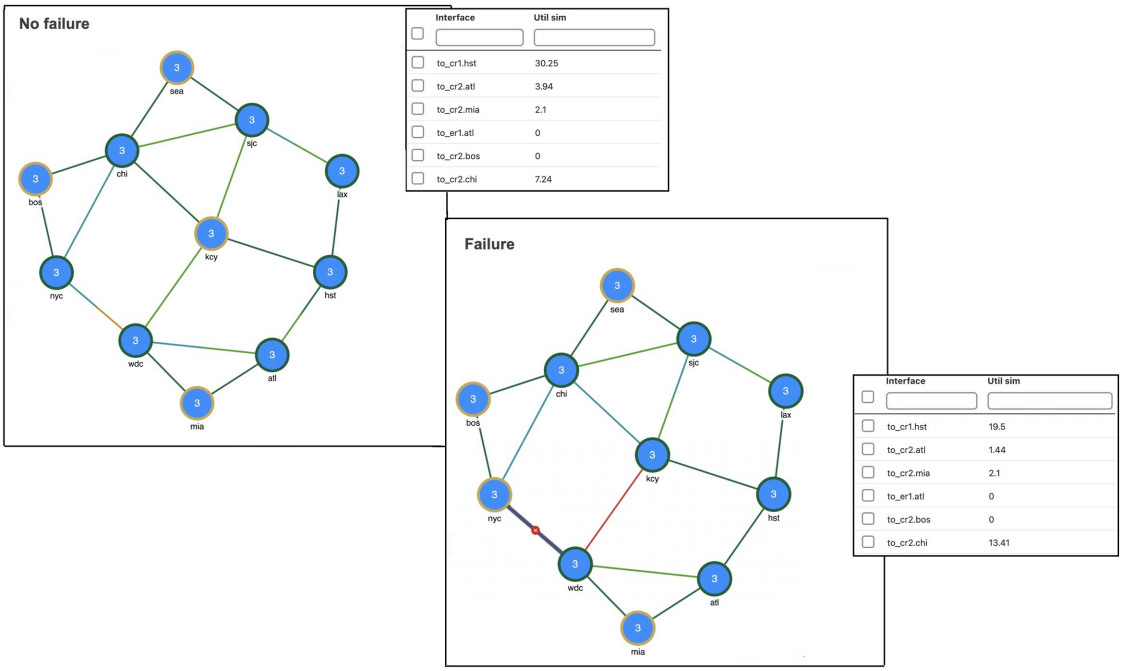
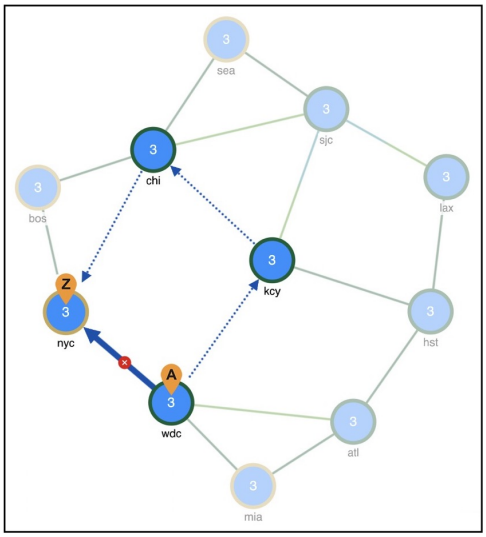



**Figure 18: Reroute of demand around a failed circuit**




### Disabling demand rerouting

To specify that a demand should not reroute around failures, uncheck the **Reroutable** check box in the demand's Edit page. This can be used as a way of including L2 traffic on an interface. For example, a one-hop, non-reroutable demand can be constructed over the interface to represent the L2 traffic. Other reroutable demands can be constructed through the interface as usual. If the interface fails, the L2 traffic is removed and the L3 traffic reroutes.

### Fail or recover plan objects

Guides you through updating the operational status of network objects by simulating failures or performing recovery.

Use this task to update the operational status of one or more network objects to simulate failures or perform recovery.

#### Before you begin

This table lists the objects that you can fail or recover.

| | |
|---|---|
| Interfaces | Ports |
| Circuits | Port circuits |
| Nodes | SRLGs |
| Sites | External endpoint members |

### Procedure

1. To fail a single object, click ••• > **Fail** under the **Actions** column.
2. After failing an object, the menu option changes to **Recover**. Use this option to recover the failed objects.
3. Follow these steps to fail or recover multiple objects:
   a) Select the required objects.
   b) Click **More actions** > **Fail** or **More actions** > **Recover** to fail or recover objects.

> **Note**
>
> You can
>
> - select all objects, and fail or recover them
> - filter objects, and fail or recover them, or
> - select objects from the network plot, and fail or recover them.

The failed object is marked with a red circle and white X in the network plot (for example, refer to Figure 17: Failed circuit on page 86). The icon disappears when you recover the object.

### Protect circuits from SRLG failures

Guides you through protecting circuits from SRLG failures and SRLG worst-case analysis by setting circuit properties and adjusting network options.

You can protect circuits from being included in SRLG failures and SRLG worst-case analysis. However, there are differences in behavior, such as:

- These circuits do not fail when an SRLG fails individually. However, they will fail if the circuit itself fails.
- These circuits are protected from being included in Simulation analysis regardless of whether they are in an SRLG.

> **Note**
>
> This setting does not affect the routing of FRR SRLGs. For information on FRR SRLGs, refer to Optimize RSVP-TE Routing on page 317.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Set the **Protected** property for circuits.
   a) In the Network Summary panel, select one or more circuits from the **Circuits** table.
   b) Click .

   > **Note**
   >
   > If editing a single circuit, you can also use the ••• > **Edit** option under the **Actions** column.

   c) In the **State** field, check the **Protected** check box.
   d) Save the changes.

3. Set the **Network options** property for protecting circuits included in SRLGs.
   a) Click  in the toolbar or choose **Actions** > **Edit** > **Network options**.
   b) In the **Redistribute routes across IGP process** section under the **Simulation** tab, check the **Exclude protected circuits from SRLG failure** check box.
   c) Save the changes.

## Active state

Introduces the concept of active state for objects, explains conditions for inactivity, details methods to change active state, and outlines the impact on demand routing and utilization simulation for various network objects.

An active state indicates whether an object is available for measured or simulated traffic calculations.

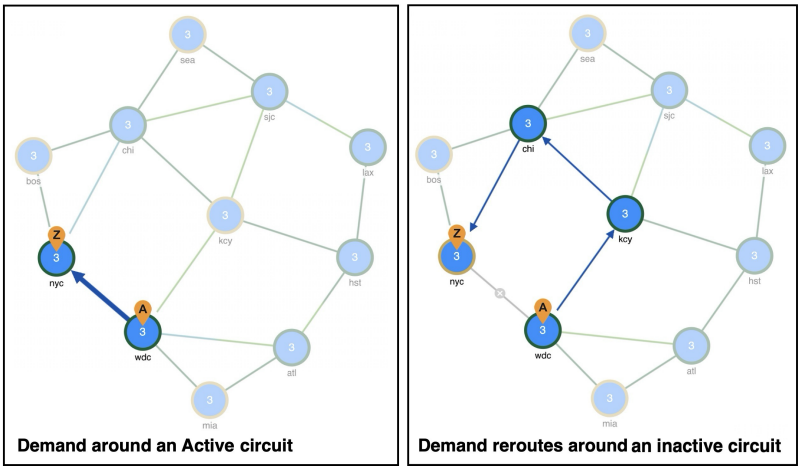
An object can be inactive for several reasons:

- It is administratively down.
- It is a placeholder. For example, you might be planning to install an object and want its representation in the network plot.
- It exists in a copied plan but was not present in the original discovered plan.

You can simultaneously change the active state of one or more objects. If you change the active state of an interface, you change its associated circuit.

Like failures, changing an object from active to inactive immediately affects demand routing and the **Util sim** column in the Interfaces tables (Figure 19: Inactive circuit on page 89).

**Figure 19: Inactive circuit**




### Objects that can be set as Active

This is the list of objects that can be set as Active:

- Circuits
- Nodes
- Sites
- Ports
- Port circuits

- SRLGs
- External endpoint members
- Demands
- LSPs
- LSP paths

#### Set objects to Active or Inactive state

Guides you through setting objects to Active or Inactive states within your network design by selecting, editing, and saving object status in the Network Summary panel.

Complete these steps to set the state of the objects to Active.

### Before you begin

Make a note of the list of objects whose State can be set to Active or Inactive by referring to the Active State section.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel, select one or more objects from their respective tables.
3. Click .

   **Note**

   If editing a single object, you can also use the ••• > **Edit** option under the **Actions** column.

4. In the **State** field, check the **Active** check box to toggle it on or off.

   A check mark indicates the object is active. To make it inactive, clear the check box.

5. Save the changes.

The system updates the selected objects to active or inactive, reflecting their current state in your network design. If an object is active, the **Active** column in the Network summary table displays "true" for that object.

## Operational state

Explains how operational state is determined for network objects based on their failed and active status, and outlines the impact of object failures or inactivity on the operational status of related components.

The operational state identifies whether the object is functioning. You cannot set an operational state. It is automatically calculated based on the failed and active states.

- Any object that is failed or inactive is operationally down.
- If the object relies on other objects to function, its operational state matches the state of those objects.

This table lists how failure or inactivity of objects impacts operational status.

| If this object fails or is inactive | These objects are operationally down |
|---|---|
| Node | circuits connected to the failed node |
| Site | sites, nodes, and circuits within the failed site |
| SRLG | objects within the failed SRLG |
| Port | port circuits that contain the failed port |

## Simulated capacity

Explains how simulated capacity represents the effective capacity of network objects, reflecting both configured and operational conditions to guide utilization calculations and planning.

A simulated capacity is a network capacity metric that

- reflects the effective capacity of an object
- takes into account network state, including any failures that may reduce capacity, and
- is represented by the **Capacity sim** column in the Network Summary tables.

All utilization figures in the Interfaces, Circuits, and Interface Queues tables are calculated based on this value.

The **Capacity** column shows the configured physical capacity of interfaces, circuits, ports, and port circuits. Each circuit, port, and port circuit has a physical capacity that you can set in the Capacity field of the Edit window. The interfaces have a configurable capacity that you can set in the **Configured capacity** field. From these properties, a simulated capacity (Capacity sim) is derived for each object.

#### Simulated capacity calculation rules

There are a few rules to note regarding Capacity sim calculation.

- If a circuit's Capacity is specified, this becomes the Capacity sim of the circuit, and all other capacities (interface and constituent port capacities) are ignored. Specify circuit capacities instead of interface capacities to easily modify existing capacities. This is useful for build-out planning.
- If the circuit has no Capacity, then its Capacity sim is the minimum of its constituent interface Capacity values. The Capacity of an interface is the sum of the Capacity values of the associated ports. If the interface has no ports or if the ports have no Capacity, it is the same as the interface's Configured capacity property.

> **Note**
>
> The field in the interface's Edit window is **Configured capacity**, while the column name in the Interfaces table is **Capacity**.

- If two ports are connected explicitly by a port circuit, the Capacity sim of the port circuit is set to the minimum capacity of the three, which effectively negotiates down the capacity of each side of the connection.
- In a LAG interface, if any of the constituent LAG members are operationally down, the interface Capacity sim column shows a value that is reduced by the aggregate capacity of all the LAG members that are down. For example, if a 1000-Mbps port of a four-port 4000-Mbps LAG is operationally down, the simulated capacity for that LAG interface becomes 3000 Mbps.

> **Note**
>
> If a pair of ports is considered in Capacity sim calculations, both must be operational to be considered.

## Delay

Explains the concept of delays in network interfaces and circuits, including how delay is represented, configured, and why asymmetric delay matters.

Delay is a property that

- indicates the time it takes for a data packet to travel from the source to the destination across a network path
- can be set in the Add or Edit page of the Interfaces and Circuits, and
- is indicated by the **Delay** column in the Interfaces and Circuit tables.

### Asymmetric delay

Starting with the 7.2.0 release, Cisco Crosswork Planning allows you to set different delay values on interfaces associated with circuits. In other words, each interface in a circuit can have its own delay metric, rather than a single delay value for the circuit as a whole.

### Columns related to delay

There are two columns related to Delay in the Interfaces and Circuits tables.

- Delay: One-way transmission latency in milliseconds (ms).

- Delay sim: (Derived) Simulated one-way transmission latency. All Cisco Crosswork Planning delay calculations that use the L3 circuit delay, such as Metric optimization, use the Delay sim value.

# 5 Simulate Traffic Flow from Source to Destination Using Demands

**Topics:**


- Demands
- Use cases of demands
- Create demands and demand meshes
- Visualize demands in the network plot
- Demand traffic


Explains demand modeling in network simulation, covering creation, configuration, traffic assignment, latency bounds, demand meshes, visualization, traffic modification, and deduction techniques to optimize network design and traffic engineering with Cisco Crosswork Planning.

# Demands

Explains how demands represent the source and destination of potential traffic flows across a network, and the properties that define them.

A demand is a network modeling object that

- describes the source and destination of a potential traffic flow across a network
- determines how traffic is routed through the simulated model, and
- contains properties and classifications that influence simulation and analysis accuracy.

Since demands determine how traffic is routed through the simulated model, creating realistic demands and demand meshes (refer to Demand meshes on page 98) is imperative to the accuracy of other information that can be derived from Cisco Crosswork Planning. As such, all defaults are set to create demands and demand meshes that best suit most network models.

A route simulation determines the routes that this traffic takes from the source to the destination, which are determined by the topology, the routing protocols, and the failure state of the network. To model IGP routing, these sources and destinations are nodes or interfaces within the topology. To model basic inter-AS routing, the sources and destinations are neighboring external ASes, peering nodes in these ASes, or interfaces to these peering nodes.

Each demand has a specified amount of traffic. The demand traffic is the basis for many of the Cisco Crosswork Planning simulation and traffic engineering tools. An accurate set of demands and demand traffic is essential for effective planning, designing, engineering, and operating of a network. Accurate knowledge of demands is essential for accurate traffic trending and traffic growth predictions. There are several methods for putting this traffic into the demands, including the **Demand deduction** tool, which calculates a realistic amount of per-demand traffic based on measured traffic.

### Demand properties and visualization

Each demand consists of unique properties (keys) that define it, additional properties, and traffic. This table summarizes these. For a complete list of properties, refer to the available columns in the Demands table.

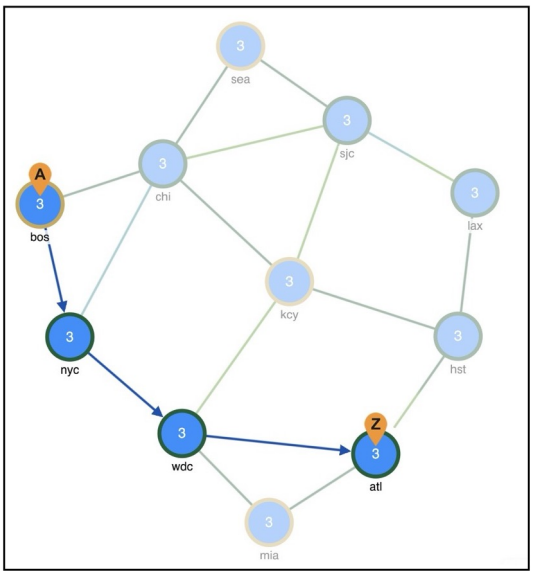
**Table 5: Demand properties**

| Property | Description |
|---|---|
| Unique properties (Keys) | Each demand is defined by a unique combination of these four properties:<br>• Name: By default, this is blank.<br>• Source: Nodes, interfaces, external ASes, or external endpoints.<br>• Destination: Nodes, interfaces, external ASes, external endpoints, or multicast destinations. |

| Property | Description |
|---|---|
| Commonly used properties | • Service class: User-defined classification of traffic, such as for voice or video.<br>• Latency bound: The policy that sets the maximum permissible latency on a demand during normal operation. This property is used by Cisco Crosswork Planning traffic engineering tools.<br>• Topology: Demands can be assigned to a specific IGP, and only route through interfaces that belong to that IGP.<br>• Private LSP: If a demand is associated with a private LSP, the demand can only route through that LSP, and the only demand that is permitted to cross that LSP is this demand.<br>You can associate an existing demand to an existing private LSP. The Private LSP drop-down list shows the private LSP that is currently associated with the selected demand. You can choose a different private LSP, or you can choose None to remove an associated LSP.<br>• Active: Only active demands are routed during simulations.<br>• Reroutable: You can enable or disable the routing of demands around failures. Disabling reroutes might be useful in some situations.<br>• Require LSP: If this option is selected, Cisco Crosswork Planning simulation only uses LSPs in routing that demand. If this is not possible, the demand will not be routed. By default, this option is disabled. |
| Traffic | By default, demands have zero traffic. You must add simulated traffic to them.<br>Demand traffic belongs to the service class of the demand.<br>Demand traffic can be set per traffic level. |

Selected demand paths are shown in blue color. An “A” labels the source and a “Z” labels the destination.

**Figure 20: Demand route**



## Recommendations for demand sources and destinations

Outlines key recommendations for specifying sources and destinations during demand routing.

Use these recommendations when creating sources and destinations for demand routing:

- For internal routing, use nodes as sources and destinations.
- For external ASes, use a combination of ASes, nodes, and interfaces. Specify the exact interface on which the demand traffic enters or exits a node by using interfaces.
- For complex routing requiring multiple sources, destinations, or failover scenarios, use external endpoints.
- For multicast routing, use multicast destinations.

If multiple interfaces are attached to a node and a demand is sourced to or destined for that node, traffic splits across one or more of those interfaces, depending on other properties, such as IGP metrics or BGP policies on a peering circuit. You can, however, specify just one of those interfaces.

> **Note**
>
> If using an interface as a source of a demand, it refers to the inbound interface. If using an interface as the destination of a demand, it refers to the outbound interface.

## Demand meshes

Explains demand meshes, the default behaviors for their creation, and the available options for network demand automation.

Demand meshes enable you to create numerous demands efficiently for the entire network. By default, Cisco Crosswork Planning creates a source-destination mesh among nodes, interfaces, external ASes, and external endpoints. There are also advanced options, such as the ability to use a different set of destinations to create the demand meshes.

## Demand latency bounds

Explains the concept of demand latency bounds and their role in traffic engineering and simulation analysis.

A demand latency bound is a traffic engineering policy that

- sets the maximum permissible latency on a demand under normal operation
- guides route selection of the traffic engineering tools, and
- helps the Simulation analysis tool determine if latency requirements are violated in worst-case failure scenarios.

### Demands table latency measurements

Each demand can have a latency bound. The Demands table provides several Latency columns. Key ones are:

- **Average latency**: Average latency over all ECMP subroutes.
- **Minimum latency**: Minimum latency over all ECMP subroutes.
- **Maximum latency**: Maximum latency over all ECMP subroutes.
- **Min possible latency**: Total latency of the shortest path that the demand could take.
- **Diff min possible latency**: Maximum latency minus the Minimum possible latency.
- **Latency bound**: Maximum permissible latency on a demand.

- **Diff latency bound**: Latency bound minus the Maximum latency.

## Use cases of demands

Lists the main purposes and practical applications of demands in network modeling, design, and analysis.

Demands are used in network planning to support these use cases.

**Use cases of demands**

| Purpose | Suggested Steps to Take |
|---|---|
| Model discovered networks | 1. Create a demand mesh based on where the traffic originates. For example, if all traffic is between edge routers, create a demand mesh between those edge routers. For more information, refer to Create a demand mesh on page 100.<br>2. Set the demand traffic manually or using the Demand Deduction tool. For more information, refer to Modify demand traffic on page 109 and Estimate demand traffic using demand deduction on page 115. |
| Model future usage in the network | 1. Create a demand mesh. For more information, refer to Create a demand mesh on page 100.<br>2. After setting the traffic, use the Cisco Crosswork Planning tools for growing the traffic and then analyze the effects on the network. You can import demand growth, you can modify selected demand traffic to emulate growth, or you can use demand groupings and other forecasting tools. For more information, refer to Evaluate Impact of Traffic Growth on page 147. |
| Design networks | 1. Create a demand mesh. For more information, refer to Create a demand mesh on page 100.<br>2. Set the demand traffic using methods described in Simulate traffic using demands on page 107. |
| Analyze existing plans | Use a variety of Cisco Crosswork Planning tools that rely on demand traffic. |

## Create demands and demand meshes

Outlines the process for creating demands and configuring demand meshes.

### Create a demand

Configure demand properties from the Network Design page to model traffic and forecasting scenarios.

Use this task to add a new demand to your network plan via the **Network Design** page.

#### Before you begin

All selections and entries are optional except for identifying the source and destination.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.

   From the toolbar, choose **Actions** > **Insert** > **Demands** > **Demand**.

   Alternatively, in the Network Summary panel on the right side, click + > **Demands** in the **Demands** tab.

2. Enter the demand name.
3. Specify the type of source and destination (node, interface, external AS, or external endpoint). Select additional Source and Destination details as needed.

   **Figure 21: Source and destination options**




4. Select a service class. If there are no service classes, create a new service class and select it. For instructions on creating a new service class, refer to Create service classes on page 192.
5. Enter a value for the latency bound.
6. Choose a topology to restrict the demand routes only to interfaces or LSPs belonging to that topology. The default is unrestricted routing.
7. Retain the **Active** default to include the demand in Cisco Crosswork Planning simulations, or uncheck **Active** to exclude this demand from simulations.
8. For each traffic level, click **Edit** to specify the amount of traffic. You can also leave it empty for the Demand deduction tool to complete. Specify the growth rate you want to use for forecasting purposes. For more information, refer to Evaluate Impact of Traffic Growth on page 147.
9. Click **Add** to create the demand.

The new demand is created and appears in your plan, ready for simulation and analysis.

## Create a demand mesh

Describes how to configure a demand mesh in your network design plan to streamline demand creation and enable targeted traffic modeling.

Complete these steps to add a demand mesh to your network plan via the **Network Design** page.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **Demands** > **Demand mesh**.

   Alternatively, in the Network Summary panel on the right side, click the **Demands** tab, and then click > **Demand mesh**.
3. In the **Demand mesh details** panel:
   a) Enter a demand name. By default, no name is set to prevent the creation of many demands using the same name. If you need to identify a specific area of the network, such as a VPN, assigning names can be helpful. However, leaving demands unnamed reduces the chance of creating a large number of demands that have the same name.
   b) Select a service class.
   c) Select a topology.
4. In the **Source** panel:

   Select one or more sources from the Source check boxes. Options include Nodes, External ASes, and External Endpoints. By default, all options are selected. Also, all available nodes, external ASes, and external endpoints are selected. To customize your selection, use the **Edit** button next to each option.
5. In the **Destination** panel:
   a) To create demands for destinations that differ from the sources, check the **Specify separate set of destinations** check box and provide the required details.
   b) To create demands in only one direction, uncheck the **Also create demands from destination to source** check box. This applies only when you select a different set of destinations.
6. Expand the **Other options** panel and make the required changes.
   - **Delete existing demands with same name**: Deletes all existing demands before new ones are created. The default (unselected) is to keep the existing demands and add only new ones.
   - **Use interface endpoints to/from external AS nodes**: When creating demands for external ASes, use a source/destination type of interface, and create a demand for all interfaces connected to each node in the external AS. For information on AS relationships and routing policies, refer to Simulate BGP Routing on page 171.
   - **Respect AS relationships**: If checked, keep the existing AS relationships defined by the Routing Policy (default). If unchecked, recreate the AS relationships. The Routing policy property is defined in the Edit AS Relationships window. For more information about AS relationships and routing policies, refer to Simulate BGP Routing on page 171.
   - **Respect external mesh settings**: If checked, keep the existing External mesh settings defined for external AS meshes (default). If unchecked, recreate the external AS mesh. The External mesh property is set in the Edit AS window.
   - **Include demands to self**: Creates demands that have the same source and destination node (default).
7. Save the changes.

## Create demands for LSPs

Explains how to create demands for LSPs in your network plan.

Use this task to add demands for LSPs in a plan file.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, click the **LSPs** tab.
3. Click  and choose **Demands for LSPs**.
4. Select the LSPs over which you want to run the demands.
5. Select the service class for the resulting demands.
6. Select the traffic for the newly created demands.
   - Traffic equal to the LSP setup bandwidth
   - Traffic equal to the LSP measurements
   - Zero, which is appropriate if you need to insert demand traffic in other ways, such as using Demand Deduction, importing it, or manually modifying it.
7. To remove the restriction of setting these demands to only these LSPs, uncheck **Mark LSPs as private**. If you do not uncheck this check box, the LSPs will use only the resulting demands.
8. Submit your changes.

Demands are created for selected LSPs and applied according to traffic settings and service class.

## Set demand latency bounds

Describes how to configure latency bounds for network demands to optimize traffic flow and ensure compliance with service requirements.

Use this task to set fixed latency bounds for demands within a network plan to control traffic latency and support performance objectives.

Use the Edit Demands page to set demand latency bounds to fixed values. All values are in milliseconds (ms).

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select one or more demands from the **Demands** table.
3. Click  and select the required **Demands**.

   **Note**

   If you are editing a single demand, you can also use the ••• > **Edit** option under the **Actions** column.

4. To set a fixed value for the latency bound, enter a value in the **Latency bound** field.

   To delete a latency bound, delete the text in this field.
5. Save your changes.

# Visualize demands in the network plot

Describes how to highlight demand paths in the network plot for better visualization and analysis.

Use this task to visualize demand paths in the network plot to understand traffic flows.

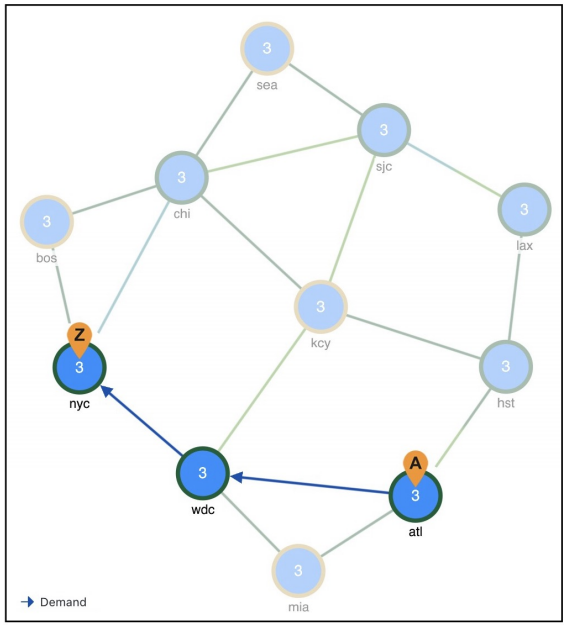
Demand paths represent specific data flows across the network. By visualizing these paths, you can quickly identify source and destination points and inspect how demands traverse through nested sites.

**Procedure**

1. In the Network Summary panel on the right side, select the demands you want to visualize from the **Demands** table.

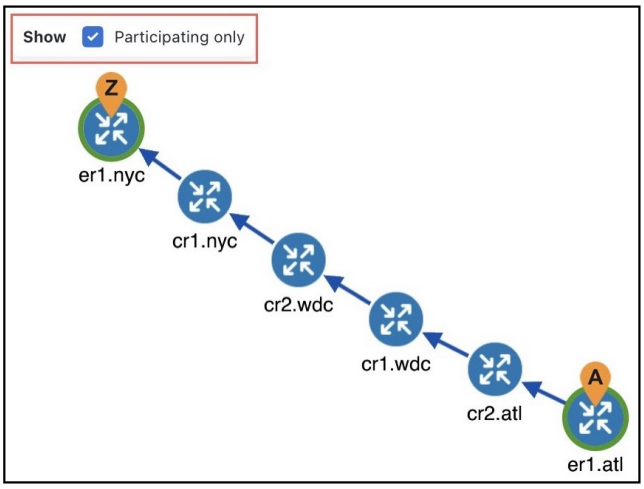
   Selected demands appear as blue arrows in the network plot. The source is labeled as “A” and the destination is labeled as “Z”. If sites are nested, these “A” and “Z” labels appear in all relevant child sites.

   **Figure 22: Demand route**

2. Hover over any selected demand in the Network Summary table to highlight its path in the network plot.
3. To view only the nodes participating in the selected demand, perform these steps.
   a) Check the **Participating only** check box in the network plot.

      **Figure 23: Participating only nodes**




b) To return to the previous view, uncheck the **Participating only** check box or check the **Show Groups** check box.

---

The network plot displays selected demand paths with clear source and destination labels, allowing for rapid analysis of traffic flows.

## Traffic rerouting

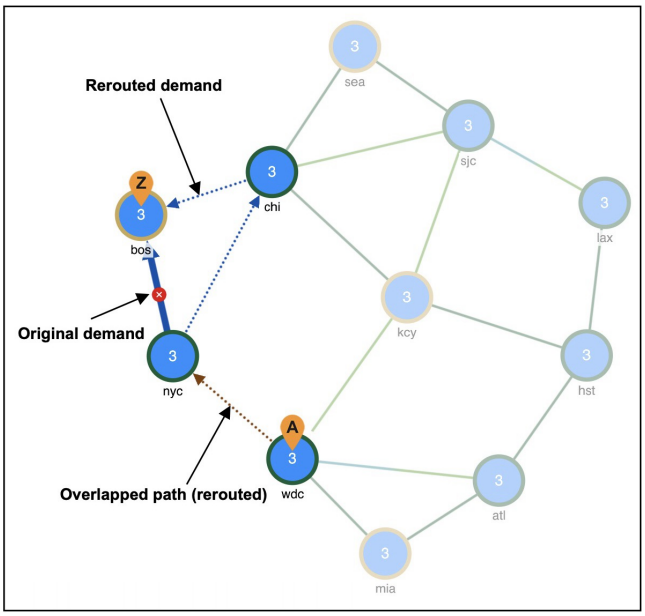
Explains how traffic reroutes operate to redirect demand around network failures and visualizes these behaviors using a diagram.

Traffic rerouting is a network behavior that

- automatically redirects demands when a failure occurs, and
- uses visual indicators to differentiate between rerouted and overlapping routes.

Demands are typically used to show how traffic reroutes around failures. In a network plot, a blue dotted line represents the rerouted demand, while a brown dotted line represents a route that overlaps with another route.

**Figure 24: Rerouted demands**



## View shortest IGP and latency paths, and overlapped paths

Describes how to configure map options to visualize shortest IGP paths, shortest latency paths, and identify overlapped paths for demands in the network plot.

Use this task to distinguish demand paths in the network plot based on IGP, latency, and path overlap.

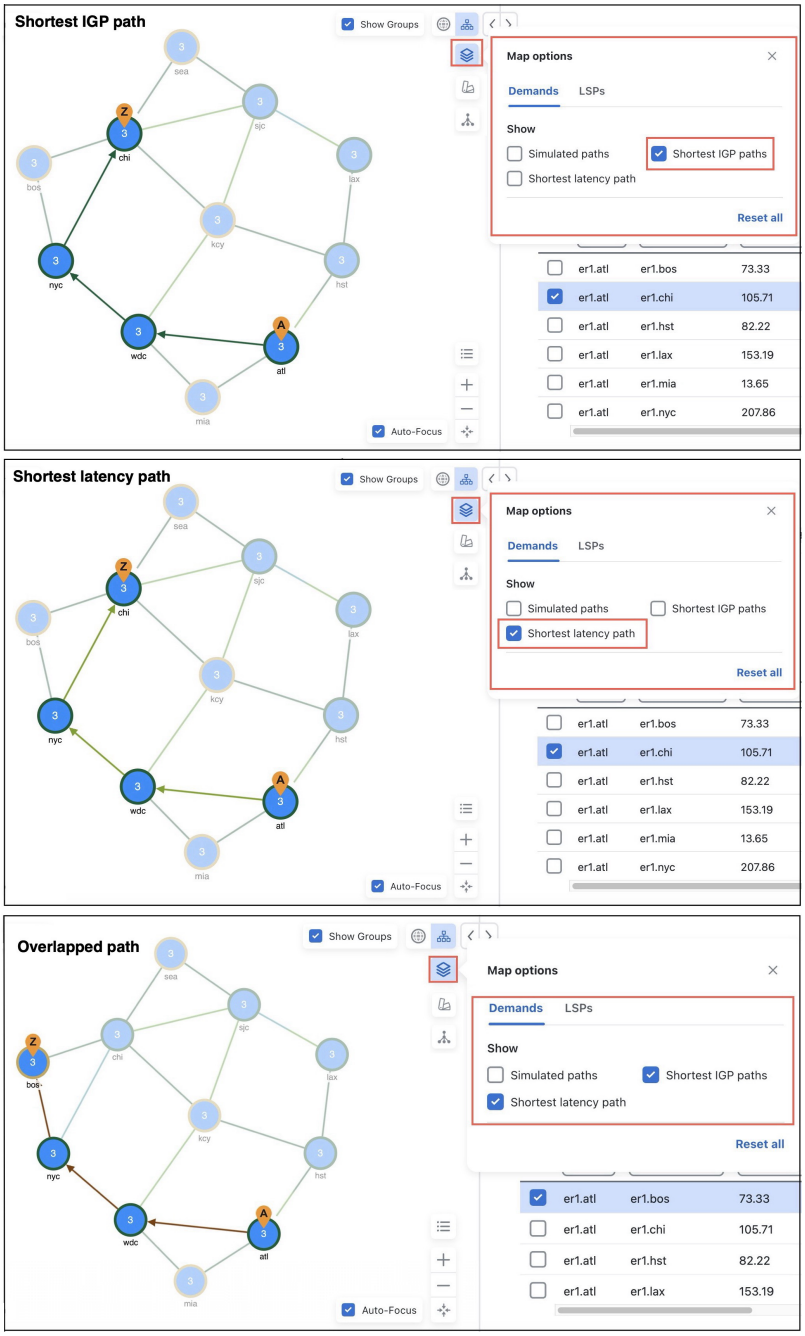
Viewing different path types in the network plot enables quick understanding of how demands traverse the network and aids in troubleshooting.

For a description of shortest IGP paths and shortest latency paths, refer to Simulate IGP Routing Protocol on page 161.

**Procedure**

1. In the network plot, click the **Map options** icon ( ).
2. Under the **Demands** tab, check the relevant options:
   - **Simulated paths**: Displays the simulated demand paths. By default, this option is selected.
   - **Shortest IGP paths**: Displays the shortest IGP path for demands.
   - **Shortest latency path**: Displays the shortest latency path for demands.
3. In the Network Summary panel, select the demands from the **Demands** table.
4. Observe the network plot:
   - If **Shortest IGP paths** is selected, the path is shown as a dark green arrow. Refer to Figure 25: Shortest IGP and latency paths, and overlapped path on page 105.
   - If Shortest latency path is selected, the path is shown as a light green arrow. Refer to Figure 25: Shortest IGP and latency paths, and overlapped path on page 105.
   - If multiple options are selected and if the routes overlap, the overlapped segment appears as a brown arrow. Refer to Figure 25: Shortest IGP and latency paths, and overlapped path on page 105.

**Figure 25: Shortest IGP and latency paths, and overlapped path**




The network plot displays the selected demand's path with distinct color-coded arrows: dark green for shortest IGP, light green for shortest latency, and brown for overlapped paths.

## Demand traffic

Explains the characteristics and purpose of demand traffic as it relates to network planning and simulation.

Demand traffic is the amount of traffic a demand is attempting to propagate through the network. For example, demand traffic is used to calculate interface utilizations during simulations. By default, demands have no traffic, and thus there is no simulated traffic.

The Demand deduction tool is the most complex and powerful way to add demand traffic. It estimates demand traffic from measured traffic values and enables accurate modeling in network simulations.

#### Traffic levels

Traffic levels are a way to represent and manage different traffic demand scenarios within the plan file. They enable flexible analysis and planning for different traffic scenarios.

You can create multiple traffic levels and associate demands with specific traffic levels. The traffic level displayed determines what is displayed in the plot and table calculations. For example, consider a totally static network that has no changes in topology or routing protocols. Throughout a day, the only thing that changes from one plan to the next is the traffic level. Thus, you can store traffic levels for different times of day all in a single plan file. Then, you can view the different states of the network from a single plan file by switching the view between different traffic levels.

## Create and select traffic levels

Describes how to configure and select traffic levels to manage planning scenarios in your network topology.

Use this task to create traffic levels that can be globally available for an opened plan file.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, select **Manage traffic levels** from the **Traffic level** drop-down list, or click **Actions** > **Edit** > **Traffic** > **Traffic levels**.
3. Click  to add a new traffic level.

   An empty row appears.
4. Enter the name and click **Save**.

   The newly created traffic level appears in the **Traffic level** drop-down list in the toolbar.
5. Select the traffic level from the **Traffic level** drop-down list.

> **Note**
>
> You can view a maximum of 200 traffic levels in the plan file.

The selected traffic level is used when visualizing the network topology and running computations in the network summary tables.

#### What to do next

(Optional) Assign the newly created traffic level to specific demands or use it to review network scenarios for different times of a day.

## Simulate traffic using demands

Guides you through simulating traffic using demands, including identifying demand routes and service classes.

In this task, we identify demand route in the network plot, determine which service classes are associated with demands, and read demand traffic and latency, and identify demand route in the network plot.

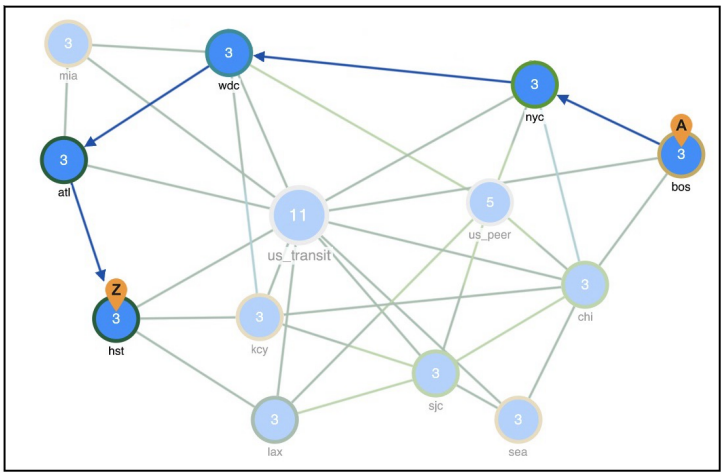
Several Cisco Crosswork Planning tables have columns that identify how much traffic is being carried or what percentage of the capacity is being utilized. For example, the Interfaces table has a **Util sim** column that reflects simulated traffic utilization. The two basic inputs to the simulation are the network configuration itself and a set of demands. A demand is a request for a specified amount of traffic to be sent from one node, the source, to another node, the destination. The routes taken are based on traffic, topology, network health, as well as the protocols used.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, click the **Demands** tab to show the Demands table.
3. Click the demand from er1.bos to er1.hst.

   The network plot shows this demand from "bos" to "hst" sites using a blue arrow to show the route, an A to show the source, and a Z to show the destination.

   **Figure 26: Demand route**




4. Show the **Service class** column:
   a) Click the **Show/hide table columns** icon ( ).
   b) In the Search field, enter the word “service”. The Service class column name appears.
   c) Select the **Service class** check box
5. Click the **Service class** column heading to sort demands by service class.
6. Read the **Traffic** column values to determine how much traffic each demand is attempting to route.
7. Determine the sum of the delays for all the interfaces on the longest path taken by each demand:
   a) In the Demands table, click the **Maximum latency** column heading and notice the values. For example, select a demand with maximum latency value of 23.
   b) Choose > **Filter to interfaces**. The Interfaces table opens and shows only the interfaces included in this demand.
   c) Click the **Show/hide table columns** icon ( ) and select the check box for the **Delay sim** column.

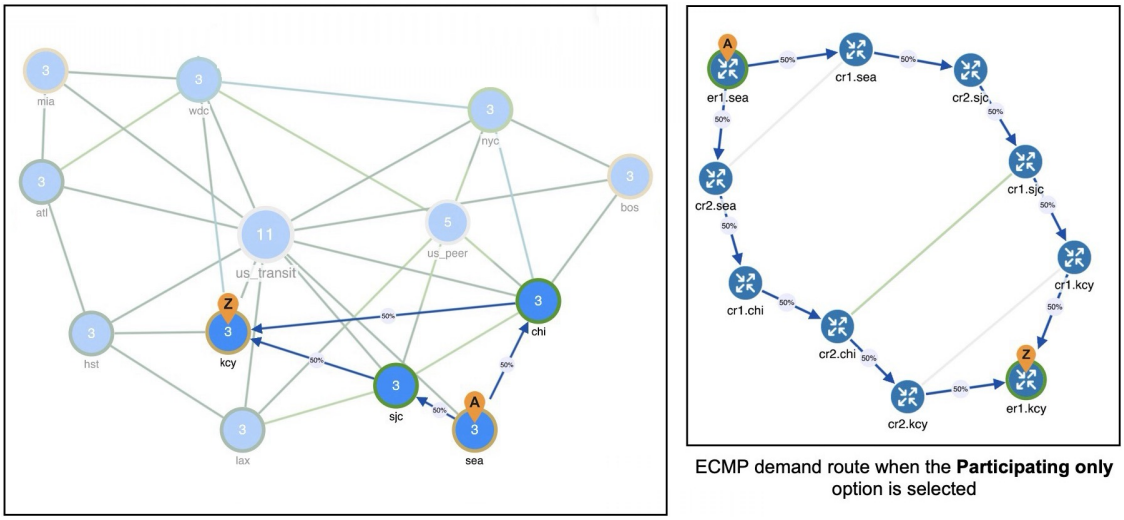
      The **Delay sim** column appears on the Interfaces table.

   d) Notice that the sum of Delay sim values of all interfaces is equal to the maximum latency of the corresponding Demand. In this case, 23.

   Clear the filter once you have completed the above steps.
8. In the Demands table, notice the er1.sea to er1.key demand takes four equal-cost multipath (ECMP) routes. The number 50% indicates that 50% of the split demand is flowing through each of these circuits.

Select the **Participating only** check box in the network plot to view the details of the demand route. When you select this option, the network plot displays only the nodes associated with the demand route.

**Figure 27: ECMP demand route**




ECMP demand route when the **Participating only** option is selected

## Modify demand traffic

Explains how demand traffic modifications can be applied to regions, sites, or the entire network to simulate network changes and plan for future growth.

You can modify demand traffic to determine the effects such changes in traffic have on the network. These modifications can be applied to regions or sites, or they can apply uniformly over the network. For example, you might increase the demand traffic to plan for future traffic growth by simulating an overall traffic growth trend. Another example might be to determine the network impact of an anticipated increase in sales of a particular service, such as video on demand.

You have numerous options for modifying demand traffic, all from the same window. The changes that you make apply to the selected demands for the current traffic level. You can modify demand traffic to either fixed values or to their relative values.

- To set a fixed value, use the Edit Demands page (Modify fixed demand traffic on page 109).
- To set a fixed or relative value, use the **Modify demand traffic** initializer (Modify fixed or relative demand traffic on page 110).

### Modify fixed demand traffic

Describes how to configure changes to fixed demand traffic in your network plan to simulate different traffic levels and optimize demand scenarios.

Use this task to modify fixed demand traffic.

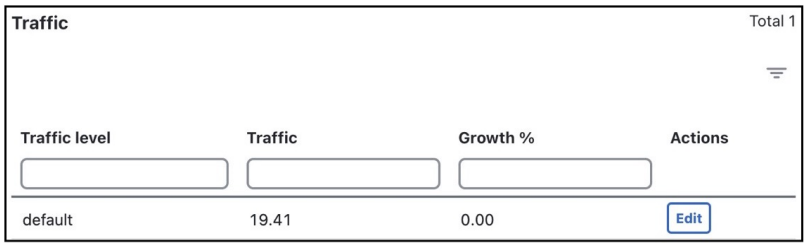
#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select one or more demands from the Demands table.
3. Click .

> **Note**
>
> If you are editing a single demand, you can also use the ••• > **Edit** option under the **Actions** column.

4. Under the **Traffic** section, click **Edit** under the **Actions** column.

**Figure 28: Modify demand traffic**




5. For each applicable traffic level, enter the desired amount of simulated traffic in the **Traffic** field and click **Save**.
6. Click **Save** on the Edit Demand page.

### Modify fixed or relative demand traffic

Explains how to configure demand traffic to a fixed or relative value using the Modify demand traffic initializer.

Use this task to modify the demand traffic to a fixed or relative value using the Modify demand traffic initializer.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the **Traffic level** drop-down list in the toolbar, select the traffic level you want the changes to apply to.
3. From the toolbar, click **Actions** > **Initializers** > **Modify demand traffic** to open the Modify Demand Traffic page.
4. Select the demands for which you want to modify the traffic. By default, all demands are selected. Deselect all and select the required demands.
5. Click **Next**.
6. Select one of the options listed in Demand traffic modification options on page 110 and choose a relevant value.
7. Submit your changes.

### Demand traffic modification options

Provides details on the available options for modifying demand traffic and how each option works.

This table lists the available options for modifying demand traffic. Except for the percentage option, all values are in Mbps.

**Table 6: Modify demand traffic options**

| Option | Description |
|---|---|
| Change traffic by __ % | Change traffic by a specified percentage. Positive percentages add to the traffic, and negative percentages subtract. For example, if the traffic were 1000 Mbps and you entered -10, the traffic would be reduced to 900 Mbps. |

| Option | Description |
|---|---|
| Add | • **Add __ Mbps in total, proportionally**: Add a set amount of traffic spread over all demands in proportion to their current traffic. For example, if one demand had 1000 Mbps of traffic and the other had 2000 Mbps, and if you added 50 Mbps proportionally, one would have 1016.67 Mbps and the other would have 2033.33 Mbps.<br>• **Add __ Mbps in total, uniformly**: Add a set amount of traffic uniformly to all the demands. For example, if one demand had 1000 Mbps of traffic and the other had 2000 Mbps, and if you added 50 Mbps uniformly, one would have 1025 Mbps and the other would have 2025 Mbps. |
| Set traffic to | • **Set traffic to __ Mbps each**: Set traffic to a fixed value.<br>• **Set traffic to __ Mbps in total, proportionally**: Set traffic to a specific value that is spread proportionally over all demands. For example, if one demand had 1000 Mbps of traffic and the other had 2000 Mbps, and if you set them to 4000 Mbps proportionally, one would have 1333.33 Mbps and the other would have 2666.67 Mbps.<br>• **Set traffic to __ Mbps in total, uniformly**: Set a specified amount of traffic, in Mbps, uniformly to all the demands. For example, if one demand had 1000 Mbps of traffic and the other had 2000 Mbps, and if you set them to 4000 Mbps uniformly, they would both be 2000 Mbps. |

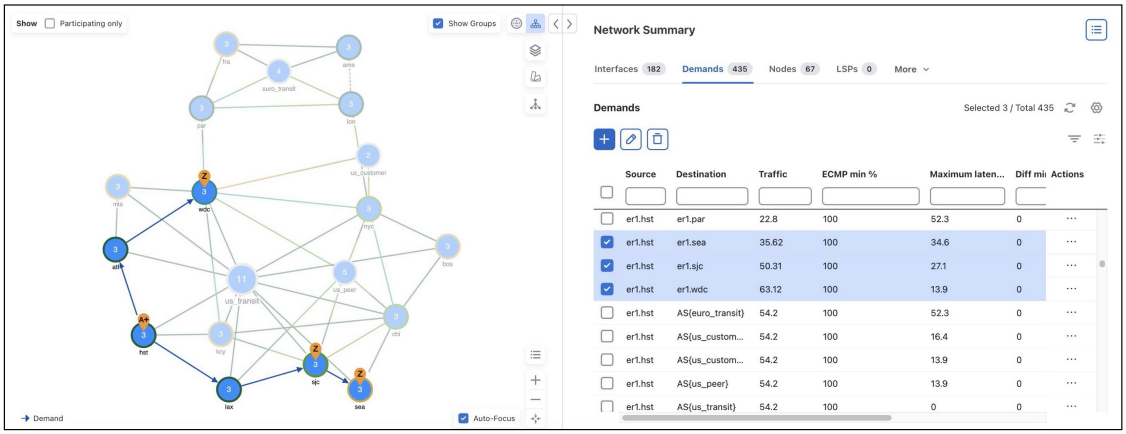
### Demand traffic modification example

Provides reference values and implications when increasing demand traffic by a specified percentage.

This example details the effects of modifying the demand traffic of the selected demands by 50%.

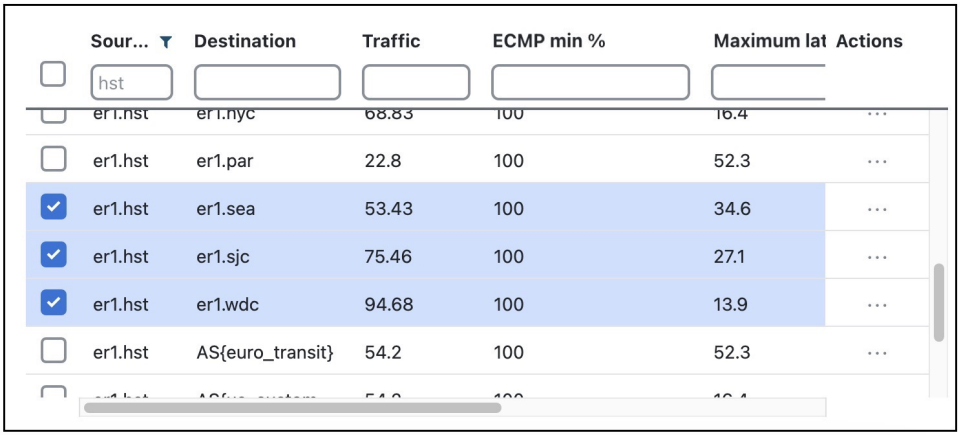
Note the values in the **Traffic** column for the demands: 35.62, 50.31, and 63.12 Mbps.

**Figure 29: Demand traffic**




To increase the demand traffic by 50%, enter **50** in the **Change traffic by ___ %** field. Notice that the values in the **Traffic** column have increased by 50% to 53.43, 75.46, and 94.68 Mbps.

**Figure 30: Demand traffic increased by 50%**




## Demand deduction

Explains how demand deductions estimate network demand traffic based on measured values from interfaces, queues, and LSPs, improving accuracy and planning.

A demand deduction is a Cisco Crosswork Planning tool that

- estimates demand traffic based on available measurements from network interfaces, queues, and LSPs, and
- adjusts results for accuracy depending on the type and availability of measurements.

Network models contain traffic measurements on the discovered network. Traffic can be measured on interfaces, interface queues, and RSVP LSPs, as well as on general traffic flows, such as from LDP LSPs. You can use the Demand deduction tool to estimate demand traffic based on any of these measurements.

The accuracy and usefulness of the results depend on factors such as the amount and type of available measured traffic. For example, interface measurements are most often available, but LSP measurements might provide more information. The results also depend on the accuracy of the demand mesh and the routing model.

Typically, you only have interface traffic measurements. In this case, the individual demands estimated by Demand deduction are not necessarily accurate. However, aggregates of demands can be highly accurate. For example, you can predict the overall utilization after a failure, a topology change, or a metric change, even if the underlying demands individually are not reliable.

For more accuracy of individual demands, include point-to-point measurements, such as for RSVP LSPs or LDP flow measurements. Also, it is useful to combine different types of measurements together for use in Demand Deduction. Interface measurements are generally the most accurate measurements available, and if included in a Demand deduction, can correct for missing or inaccurate LSP or flow measurements.

Note that you can also use Demand deduction to set **Traffic balance (%)** values for external endpoint members that are set to a **Deduce Traffic** type.

### Differences in measured and simulated traffic

Provides an overview of how demand deduction in Cisco Crosswork Planning uses measured and simulated traffic, explains key interface columns, and lists causes for discrepancies and related warnings.

Demand deduction relies on accurate topologies, demand meshes, and traffic measurements. These factors influence whether simulated traffic matches measured traffic and affect demand deduction accuracy.

You can see how close these values are by showing the **Abs meas diff** and **Meas diff/cap (%)** columns in the Interfaces table.

- **Abs meas diff**: The difference between measured traffic (Traff meas) and simulated traffic (Traff sim).
- **Meas diff/cap (%)**: The absolute measured difference expressed as a percentage of capacity.

### Causes for discrepancies between measured and simulated traffic

If these columns show large values, one of these situations likely exists:

- Inaccurate measurements: Different measurements, for example of traffic through different interfaces, can be made at slightly different points in time. Fluctuations in traffic levels might take place between the times that measurements are being taken. This means that measurements could be inconsistent with one another. Usually, these inconsistencies are small and do not seriously affect the Demand deduction results.
- Insufficient measurements: A network typically contains more demands than measurements. As a result, many solutions will fit the observed data well. Demand deduction chooses between possible solutions using knowledge of typical behavior of point-to-point traffic.
- Incorrect network configurations: If the network topology is incorrect in the plan file, the simulated routes would naturally be incorrect and measurements would not be adequately interpreted.
- Unbalanced ECMP: ECMP hashing can result in imperfect load balancing. Demand deduction, however, distributes traffic evenly across ECMPs.
- Static routes: Cisco Crosswork Planning does not model static routes. If these are present, demands routes might be simulated incorrectly, leading to deduction errors.
- Incomplete demand meshes: Demand meshes do not contain nodes even though traffic is routed between those nodes.
- Inappropriate priorities: In the Demand Deduction window, you have the option to set the priority for calculations as 1 or 2. Cisco Crosswork Planning first uses the measurements identified as Priority 1 to calculate the demands. Therefore, if the priority settings do not match the consistency of the traffic measurements in the network, the simulated traffic measurements will be less than optimal.

### Demand deduction warnings

Demand deduction displays warnings for misleading or undesirable results.

- AS “(AS name)” contains both dynamic LSPs and interface traffic. Interface traffic in AS has been ignored.

  Routing of dynamic LSPs is non-deterministic. So it is not possible to make use of both measured interface traffic and measured dynamic LSP traffic for LSPs that may (or may not) traverse these interfaces. If the network contains an AS with both dynamic LSPs and interface traffic, this warning is issued. The interface traffic is not used.

- Some interface measurements exceed capacities by as much as (percent).

  This warning is issued if a specified measurement exceeds the corresponding circuit capacity.

### Strategies for minimizing differences between measured and simulated traffic

Lists effective methods to reduce discrepancies between measured and simulated network traffic, with a focus on demand deduction accuracy.

Demand deduction estimates demands that predict interface utilizations under incremental changes to the topology, including failures, metric changes, or design changes, such as adding a new express route. If only interface measurements are available, you might fine-tune the Demand deduction calculations to get better results, especially for site-to-site traffic.

To enhance the accuracy of Demand deduction results, consider these suggestions:

- Include RSVP LSP or LDP measurements in the network discovery process.
- Restrict demand meshes to exclude demands that are known to be zero.

  For example, if you know that core nodes do not source traffic, then exclude core nodes when creating the demand mesh.

- Check the Nodes table to see if there is a node where the measured traffic going into it (**Dest traff meas**) and out of it (**Source traff meas**) are very different. Ensure these nodes are included in the demand mesh because they are either sources or destinations for traffic.
- On the Demand Deduction page, set the most consistent measurements to a Priority 1. The most reliable measurements are usually interface measurements. Likewise, LSP measurements are end-to-end, and thus also generally highly reliable. You can set multiple measurements to priority 1.

  For example, if flow measurements are inconsistent and interface measurements are very consistent, set interface measurements to Priority 1 and the flow measurements to Priority 2.
- If only a few measurements are available or if there are many inaccurate measurements, the tool sometimes estimates more traffic in a circuit than its capacity. To prevent this, on the Demand Deduction page, select the option to keep the interface utilization below 100%. This forces the resulting simulated calculations to be below the given percentage of circuit capacity.

### Flow measurements for demand deduction

Provides a summary of columns in the Flows table used for flow measurements relevant to demand deduction, including attributes and use cases.

Besides node, interface, and LSP traffic measurements, Demand deduction allows more general flow measurements. These flow measurements can be flows from (or through) a specified node, to (or through) another node. You can also use combinations of these node-to-node flows. This measurement format can be used to enter, for example, peer-to-peer flow measurements, or traffic measurements obtained from LDP routing or from NetFlow.

> **Note**
>
> You can only view the Flows table if is already present in your plan file. You cannot create, edit, or delete the Flows in the UI.

Flow measurements are entered in the plan file in the <Flows> table. They appear in the UI in the Flows table. Table 7: Flows table columns on page 114 lists some of the useful columns in the Flows table. The included traffic is defined by the **From type** and **To type** columns.

**Table 7: Flows table columns**

| Column | Description |
|---|---|
| From | Specifies the source node. |
| From type | • Source: Traffic originating at the From node is included in the flow.<br>• Interior: Traffic is included that passes through the From node, entering that node from another node in the same AS<br>• Border: Traffic is included that passes through the From node, entering that node from another node in a different AS. |
| To | Specifies the destination node. |
| To type | • Dest: Traffic that is destined for the To node.<br>• Interior: Traffic that passes through the To node to another node in the same AS.<br>• Border: Traffic that passes through the To node to another node in a different AS. |

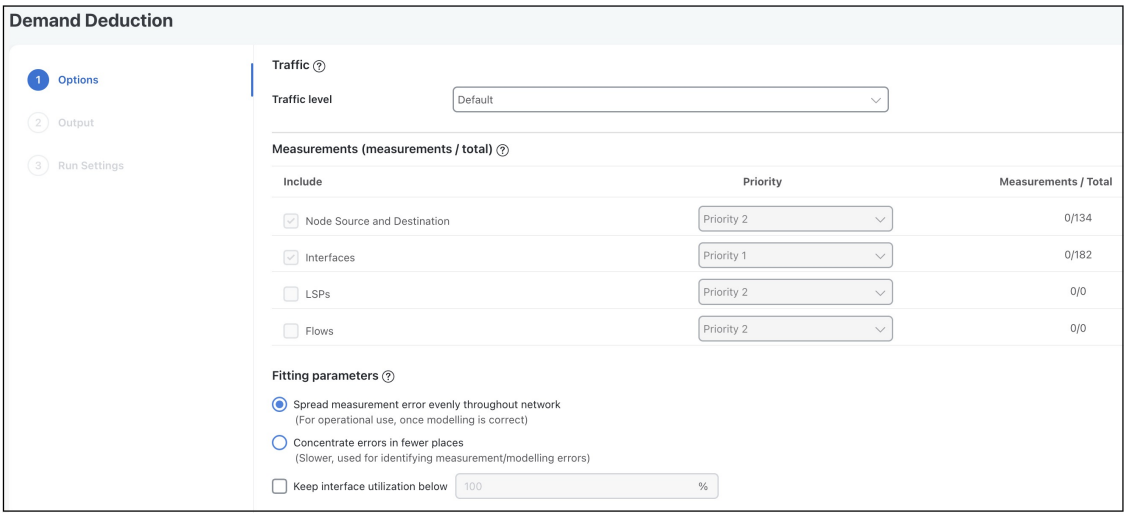
| Column | Description |
|---|---|
| Traff meas | Measured traffic used by Demand deduction in its calculations. If more than one node is included in either the From or To columns, this measurement is the sum of the traffic over all flows between individual pairs of From and To nodes. |

## Estimate demand traffic using demand deduction

Guides you through estimating demand traffic using the Demand deduction tool.

The Demand deduction tool calculates demand traffic when traffic measurements are available.

The options available can significantly affect the calculations. For information on improving accuracy of results, refer to Strategies for minimizing differences between measured and simulated traffic on page 113. For information on setting up external endpoint members to be included in Demand deduction calculations, refer to Simulate Advanced Routing with External Endpoints on page 209.




Complete these steps to estimate the demand traffic using the Demand deduction tool.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the **Traffic level** drop-down list in the toolbar, select the traffic level on which you are creating simulated traffic. The measurements used come from this traffic level, and the simulated traffic for each demand is placed in this traffic level.
3. From the toolbar, choose **Actions** > **Tools** > **Demand deduction**.
   a) (Optional) Modify the traffic measurements for one or more demands by clicking the **Edit** button in the **Measurements (measurements / total)** section. In the window that appears, you can modify measured traffic of interfaces.You can further change these measurements for a traffic level and interface queue.

      Another option in this window is to enter growth percents for use with the **Create growth plans** tool. For more information on creating growth plans, refer to Evaluate Impact of Traffic Growth on page 147.

   b) Identify one or more types of measurements used in the calculations: Nodes (source and destination), Interfaces, LSPs, and Flows.
   c) For each type, set its priority. Select Priority 1 for high priority and Priority 2 for lower priority. You can have multiple measurements of the same priority. Like priorities are calculated simultaneously with equal consideration for the measurements.

By default, all available measurements in the selected traffic set are used, and the interface measurements have priority over node, LSP, and flow measurements.

d) Choose the required **Fitting parameters**.
e) If you need to keep the traffic utilization below a different percentage than 100% (default), check the **Keep interface utilization below __ %** check box and enter a value.
f) Click **Next**.

4. Select the demands for use in constructing the demand calculations.
   - Use existing: Calculates demands using the existing demands only. This option is useful when simulating a pattern of demands that cannot be represented as a simple mesh between nodes. If you did not select one or more demands before opening this window, use this option.
   - Use selected: Calculates demands for the selected rows in the Demands table. This option is helpful when you want to recalculate some of the demands, for example, such as a VPN submesh.
5. Determine whether to fix multicast demands. If selected, the multicast demands are fixed at their current traffic value.
6. Determine whether to remove demands with zero traffic. The default is to remove them because Demand deduction typically estimates a significant percentage of the simulated traffic to be zero when a large number of point-to-point utilizations in a mesh are extremely small. Using this default can substantially improve simulation and optimization performance in large plans. Do not remove demands with zero traffic if all demand routes are of interest, irrespective of traffic. Click **Next**.
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

   The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

   > **Note**
   >
   > Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. Submit your changes.

---

The Demand deduction tool calculates the simulated traffic and lists the results in a Demand Deduction report.

### Demand deduction example

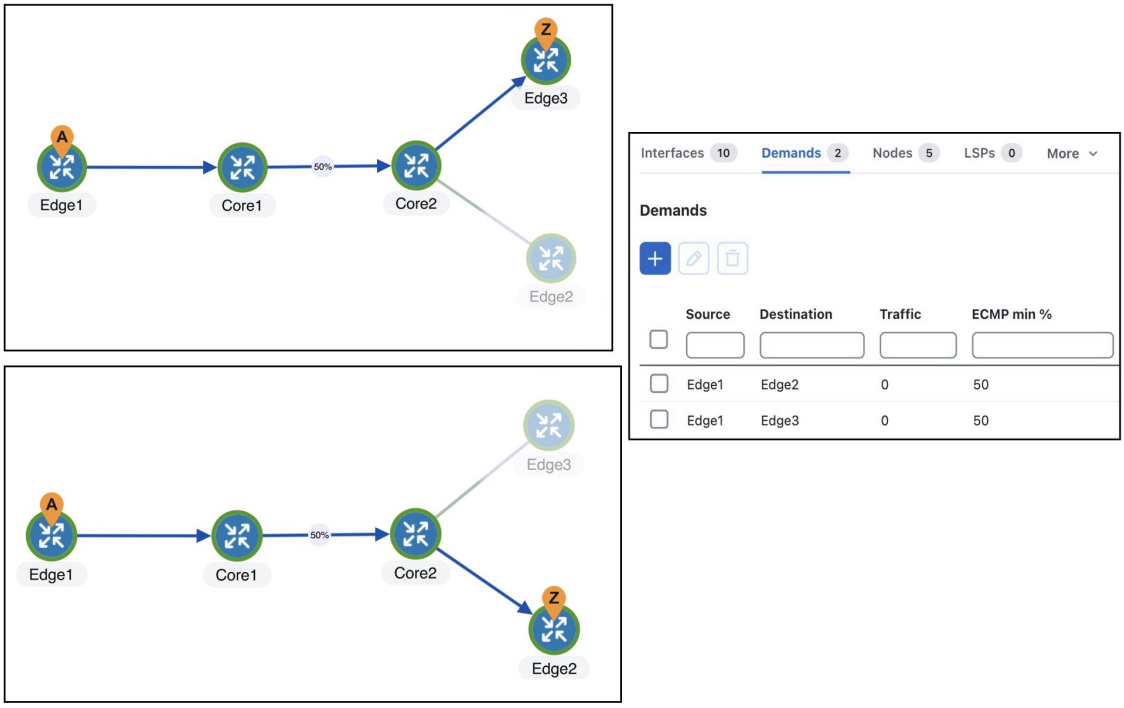
Provides a detailed example of how the Demand Deduction tool calculates estimated demand traffic in a simple network using measured interface values.

This example demonstrates results when using the Demand Deduction tool on a simple network.

### Scenario overview

Figure 31: Network containing two demands and no demand traffic on page 117 shows the routes of two demands in a network. These demands split between the two parallel core circuits due to an ECMP, and they have a common routing until the last hop. The Traffic column in the Demands table shows 0, as these demands do not contain any traffic yet.

**Figure 31: Network containing two demands and no demand traffic**




### Measured interface traffic

Figure 32: Measured traffic view and interfaces associated with the demands on page 117 shows the Measured traffic view and the five interfaces associated with the two demands, three of which have measured traffic.

- Edge1 to Core1 has 470 Mbps of measured traffic.
- One Core1 to Core2 interface has 210 Mbps, and the other has 240 Mbps, making a total of 450 Mbps. This unequal split results from imperfect load balancing of the ECMP.
- There is no traffic from Core2 to Edge2 or from Core2 to Edge3.

**Figure 32: Measured traffic view and interfaces associated with the demands**



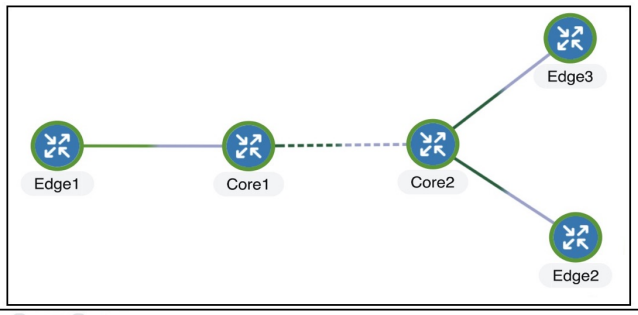

### Demand deduction simulation

Upon running Demand deduction with its default options, the Simulated traffic view appears. Other than the measured interface traffic, there is no other information about the demand traffic. So, Demand deduction first splits the difference between the measured 470 Mbps of traffic (Edge1 to Core1) and the measured traffic of 450 Mbps (Core1 to Core2) to get an estimated total demand traffic of 460 Mbps. In the absence of any other information, it divides this 460 equally to give 230 Mbps of traffic to each demand (Figure 33: Simulated view showing demand traffic on page 118). In the Interfaces table, the Traff sim column now has values and the network plot shows simulated traffic percentages on all five interfaces associated with the demands.

- Edge1 to Core1 has 460 Mbps of simulated traffic.
- Both Core1 to Core2 interfaces have 230 Mbps.
- Core2 to Edge2 and Core2 to Edge3 both have 230 Mbps.

  The Abs meas diff and Meas diff/cap (%) columns in the Interfaces table show mismatches between measured and simulated values.
- Edge1 to Core1 has a difference of 10 Mbps, or 1%.
- One Core1 to Core2 has a difference of 20 Mbps, or 2%, while the other has a difference of 10 Mbps, or 1%.
- Neither the Core2 to Edge2, nor the Core2 to Edge3 interfaces have values because they had no measured traffic.

**Figure 33: Simulated view showing demand traffic**




Interfaces 10 Demands 2 Nodes 5 LSPs 0 More

Interfaces

Selected 0 / Total 10

More actions

| Node | Interface | IGP metric | Remote node | Traff sim | Traff meas | Capacity sim | Util sim | Util meas | Abs meas diff | Meas diff/cap (%) | Actions |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Core1 | to_Core2 | 1 | Core2 | 230 | 210 | 1000 | 23 | 21 | 20 | 2 | ··· |
| Core1 | to_Core2[1] | 1 | Core2 | 230 | 240 | 1000 | 23 | 24 | 10 | 1 | ··· |
| Edge1 | to_Core1 | 1 | Core1 | 460 | 470 | 1000 | 46 | 47 | 10 | 1 | ··· |
| Core2 | to_Core1[1] | 1 | Core1 | 0 | NA | 1000 | 0 | NA | NA | NA | ··· |
| Core2 | to_Edge2 | 1 | Edge2 | 230 | NA | 1000 | 23 | NA | NA | NA | ··· |
| Edge2 | to_Core2 | 1 | Core2 | 0 | NA | 1000 | 0 | NA | NA | NA | ··· |
| Core2 | to_Edge3 | 1 | Edge3 | 230 | NA | 1000 | 23 | NA | NA | NA | ··· |
| Edge3 | to_Core2 | 1 | Core2 | 0 | NA | 1000 | 0 | NA | NA | NA | ··· |
| Core1 | to_Edge1 | 1 | Edge1 | 0 | NA | 1000 | 0 | NA | NA | NA | ··· |

### Alternate scenario

In this same example, if the Core2 to Edge2 interface had 50 Mbps traffic, the results would have been different. Because this interface is used only by the one demand, the measured 50 Mbps of traffic would be used as an estimate only for that one demand. Using the same logic as before, the demands should total to 460 Mbps, so the other demand is set to the difference, which is 410 Mbps.

# 6 Perform What-If Analysis

**Topics:**


- Failure scenarios
- Perform failure analysis
- Perform impact analysis of topology changes
- Perform impact analysis of metric changes


Describes methods for conducting what-if analysis by examining failure scenarios, performing failure and impact analyses of topology and metric changes, and illustrating these processes with practical examples.

## Failure scenarios

Explains how failure scenarios are used in network planning to evaluate the impact of changes or unexpected events on network performance.

Failure scenarios are simulated events that

- represent the loss or degradation of specific network elements such as links, nodes, or metrics
- help planners assess how minor or major changes affect the network's overall service and congestion, and
- support decision-making by predicting outcomes when network conditions or topology are intentionally altered or disrupted.

A frequent question in the planning process is how to test the value of a minor change that may prevent the need for a larger one. For instance, instead of a huge upgrade, a planner may consider whether adding a single link between two sites can mitigate congestion in part of the network.

In Cisco Crosswork Planning, you can examine what happens if you change any aspect of the network model. For example:

- What happens if a link or a node fails?
- What happens if you change a metric?
- What happens if you change the topology?
- What happens if a new customer or service is added?

## Perform failure analysis

Describes how to configure and test the failure and recovery of network objects to analyze changes in network traffic utilization and routing.

One of the most useful features of Cisco Crosswork Planning is its ability to show the network behavior when objects fail. For example, your network might show that utilization remains well under 100% under normal operation. However, what happens when a failure occurs?

Cisco Crosswork Planning lets you perform failure analysis, which can be useful in long-term planning and short-term preparation. It simulates rerouting and changes in traffic utilization under individual failure events. For an example of how failure analysis is performed in Cisco Crosswork Planning, refer to Example: Performing failure analysis on page 123.

Complete these steps to fail or recover objects.

**Procedure**

1. To fail a single object, click the ··· > **Fail** option under the **Actions** column.

   After failing an object, the menu option changes to **Recover**. Use this option to recover the failed objects.

2. To fail or recover multiple objects:
   a) Select the required objects.
   b) Click **More actions** > **Fail** or **More actions** > **Recover** to fail or recover objects.
3. Observe how traffic utilization and routing in the network plot change after objects are failed or recovered.

## Example: Performing failure analysis

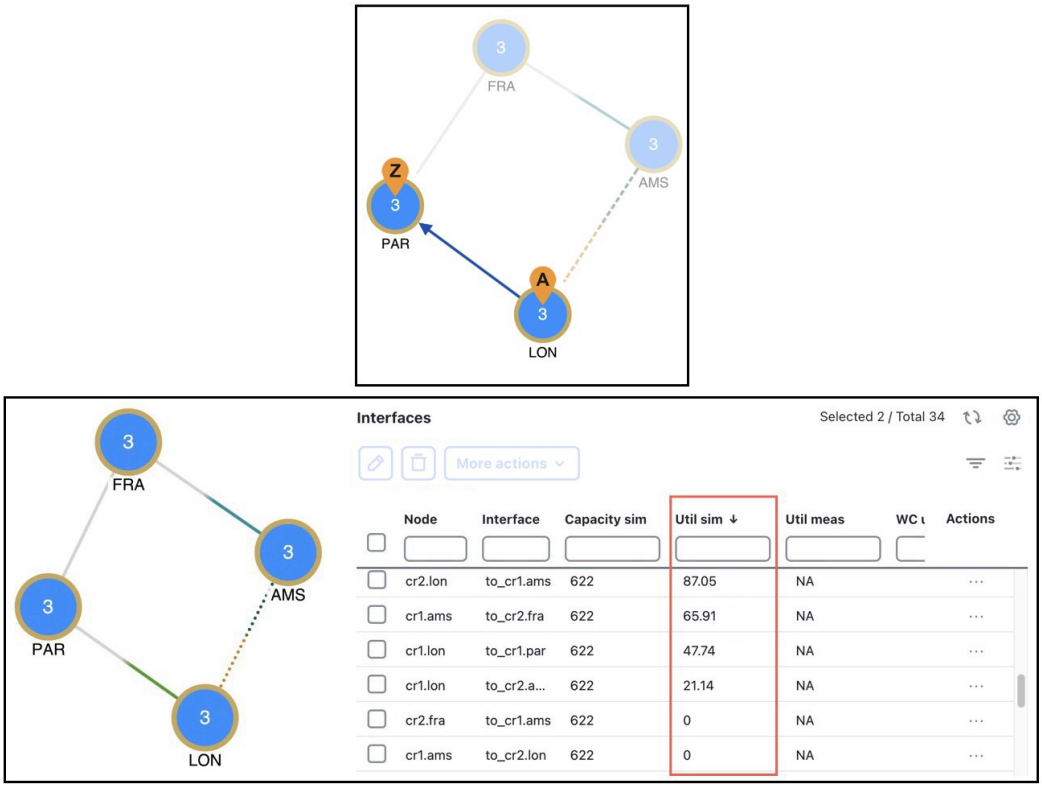
Guides you through performing failure analysis by simulating interface failures.

In this example, we determine the effects on the network if a single interface fails.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the network plot (refer to Figure 34: Network plot before failing the interface on page 123), observe that all traffic utilization is under 100%. You can tell this because none of the interfaces are red, which by default indicates over 100% utilization. Also, the **Util sim** column in the Interfaces table shows no values over 100.
3. In the Interfaces table, select the cr1.lon-cr1.par interface (LON-PAR circuit) and choose > **Filter to demands** > **Through all interfaces**.
4. Click the demand to see its route. The solid violet arrow indicates the demand path.

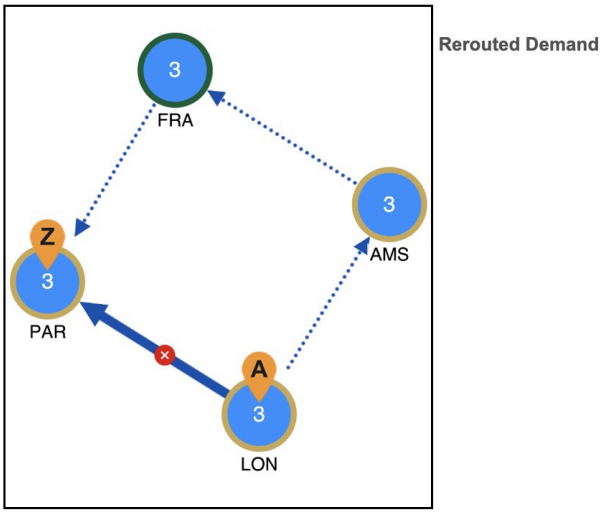
**Figure 34: Network plot before failing the interface**



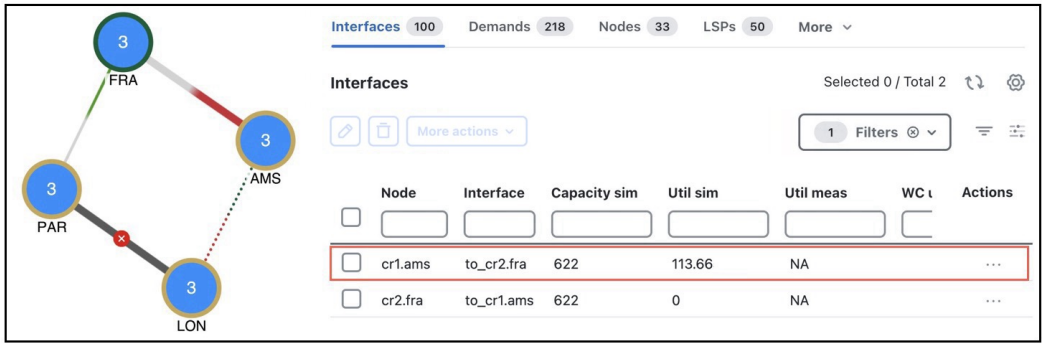

5. In the interfaces table, select the cr1.lon-cr1.par interface and choose ••• > **Fail** option under the **Actions** column. You can also use **More actions** > **Fail**.
6. Notice how the demand has been rerouted (refer to Figure 35: Network plot after failing the interface on page 123). The dotted line shows the current demand path under failure.

   The red cross on the link indicates a failure.
7. Clear the filter in the Demands table and notice the difference in traffic utilization colors in the plot. For example, the AMS-FRA circuit (cr1.ams-cr2.fra interface) is now congested, as indicated by red color.

**Figure 35: Network plot after failing the interface**




Interface Showing Congestion




8. To restore the failed circuit/interface, select it and choose ••• > **Recover** option under the **Actions** column. You can also use **More actions** > **Recover**.

---

# Perform impact analysis of topology changes

Explains how topology changes impact network performance and how you can evaluate their effects using modeling and analysis tools.

Being able to emulate and predict the impact of topology changes promotes Service Level Agreement (SLA) adherence and staff efficiency. Topology changes can affect traffic flows, congestion, and latency. Understanding these effects is valuable for operators, planners, and designers, and is critical for those with penalty clauses in customer SLAs.

Cisco Crosswork Planning lets you edit the network topology so that you can add, edit, and delete objects. Once topology changes are made, demands are rerouted, showing the resulting utilization changes in the new network. For an example of how adding a circuit between two sites can help in relieving the congestion, refer to Example: Performing impact analysis of topology changes on page 124.

## Example: Performing impact analysis of topology changes

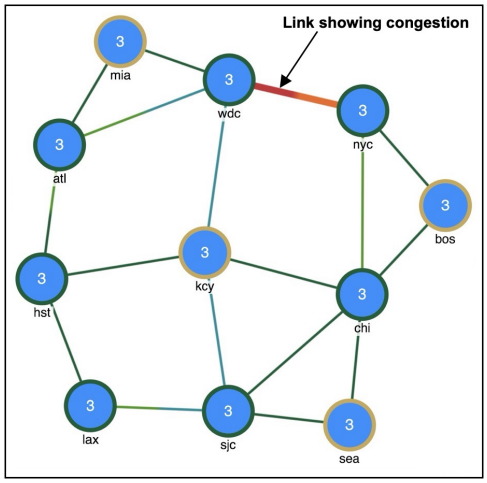
Guides you through performing impact analysis of topology changes by adding a circuit between two sites.

In this example, we see how adding a circuit between two sites ("wdc" and "nyc") helps in relieving the congestion.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Observe that the link between the sites "wdc" and "nyc" is congested, as indicated by the red interface.

   **Figure 36: Network plot before adding a circuit**




3. Add a circuit between "wdc" and "nyc".
   a) From the toolbar, choose **Actions** > **Insert** > **Circuit**.

      OR

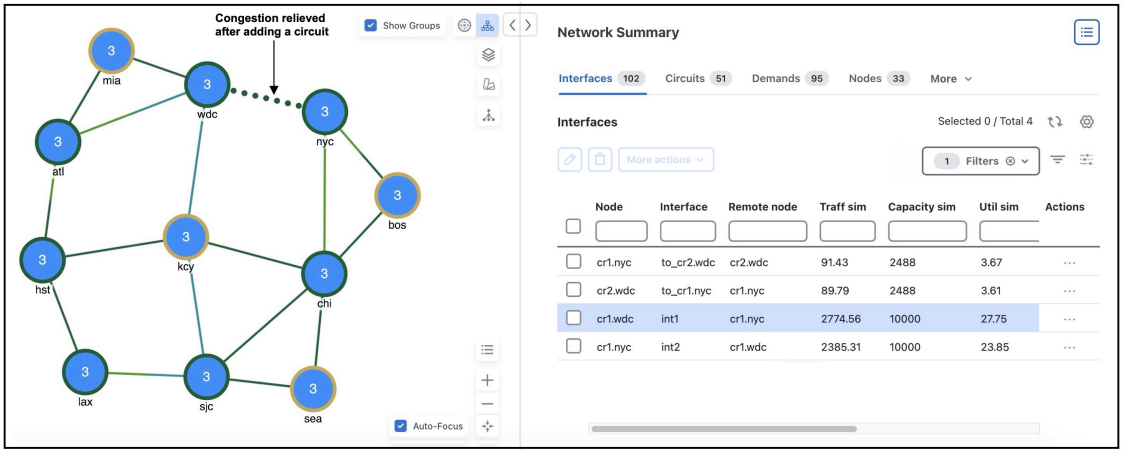
      Go to the **Circuits** table and click +. If the Circuits tab is not visible, then click the **Show/hide tables** icon (☰), select the **Circuits** check box, and click **Apply**.

   b) In the **Add Circuit** page, enter these details:

      - **Circuit name**: Enter the name of the circuit.
      - **Capacity**: Enter the amount of total traffic this circuit can carry. The drop-down list includes the most widely used capacities.
      - **Interface A** and **Interface B** sections: Specify two interfaces that are connected by the circuit. In this case, for Interface A, choose the interface belonging to the site, "wdc". For Interface B, choose the one belonging to "nyc".

   c) Click **Add**.
4. Notice the difference in the traffic utilization color in the link between "wdc" and "nyc". The dark green color indicates that the congestion is relieved. The dotted line indicates the presence of multiple circuits.

   **Figure 37: Network plot after adding a circuit**




If congestion had not been relieved, a larger circuit or different metric could have been analyzed with a simple change to the circuit properties in the UI.

# Perform impact analysis of metric changes

Explains how adjusting interface metrics affects network performance and routing behavior.

Cisco Crosswork Planning lets you update a metric and helps you analyze how it affects the network. For an example of how changing the IGP metric of an interface impacts the utilization, refer to Example: Impact analysis of metric changes on page 126.

To optimize the metric of multiple interfaces, use the **Metric optimization** and **Tactical metric optimization** tools. For more information, refer to Optimize Metrics in the Network Core on page 225.

## Example: Impact analysis of metric changes

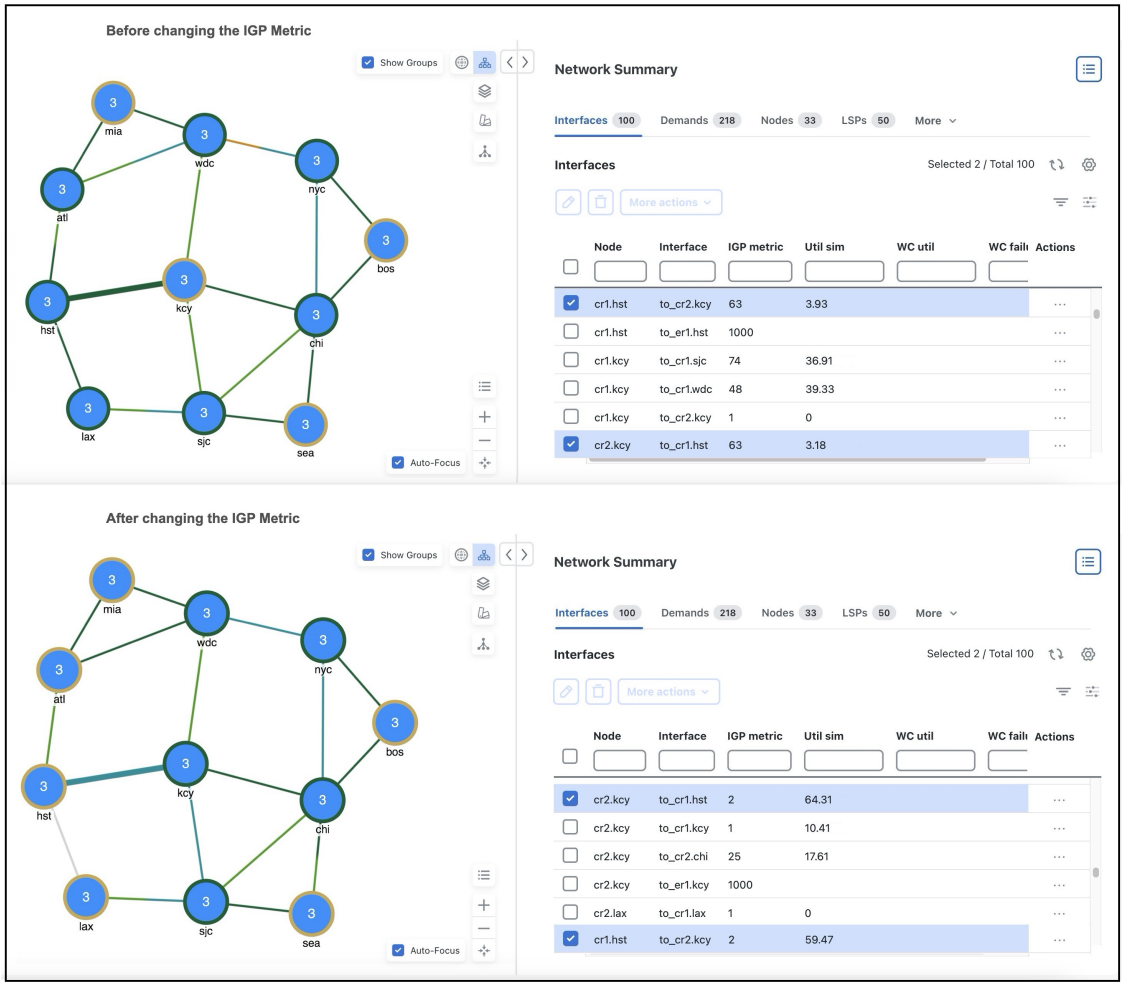
Demonstrates how adjusting the IGP metric value between the "kcy" and "hst" sites alters traffic utilization.

In this example, we see how changing the **IGP metric** value of the interfaces (between "kcy" and "hst" sites) changes the traffic utilization.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Notice that the color of the link between the sites "kcy" and "hst" is dark green. This color indicates that the traffic utilization is between 0% and 30%.

   **Figure 38: Impact of changing the IGP metric**




3. Select this link. In the **Interfaces** table, the interfaces belonging to this circuit are selected.
4. Click .
5. In the **IGP metric** field, change the value from 63 to **2**.
6. Notice that the color of the link between "kcy" and "hst" has changed to blue. This change indicates that the traffic utilization on these interfaces has increased.

# 7 Evaluate Impact of Worst-Case Failures

**Topics:**


- Simulation analysis tool
- Worst-case traffic utilization
- Worst-case demand latency
- Run simulation analysis
- Simulation analysis report
- Failure impact view
- Configure parallelization for simulation analysis


Guides evaluation of worst-case network failures by identifying critical traffic utilization, QoS violations, and latency, running simulation analyses, reviewing reports, leveraging failure visualizations, and configuring parallelization for comprehensive network resilience assessment.

## Simulation analysis tool

Explains how simulation analysis combines failure scenario results to assess network vulnerability to congestion and latency.

The Simulation analysis tool is a Cisco Crosswork Planning tool that

- combines the simulation results of a large set of failure scenarios
- determines how vulnerable a network is to congestion and high latency under different types of failures, and
- enables you to allocate sufficient network capacity for any given failure scenario.

Simulation analysis is run across a set of failure scenarios that include selected objects, such as circuits, and traffic levels. Cisco Crosswork Planning calculates these failure scenarios across all service classes. Each scenario is simulated and produces several forms of analysis. The analysis depends on the network has QoS parameters and the options selected during simulation.

- Worst-case traffic utilization on page 130 on interfaces per service class
- (Optional) VPN worst-case utilization and latency
- (Optional) Worst-case demand latency on page 133
- Failure impact view on page 144, which analyzes the impact that each failed object has on interface utilizations throughout the network

When the simulation completes, a report page appears showing a summary of each analysis and the simulations performed. This information updates every time you run a new simulation.

You can perform simulation analysis under different simulation convergence modes (Fast reroute, IGP and LSP reconvergence, Autobandwidth convergence, and Autobandwidth convergence (including failures)), depending on which stage of the network recovery after failure is being investigated. The default simulation mode is IGP and LSP reconvergence, and except where identified, the documentation describes this simulation mode.

## Worst-case traffic utilization

Outlines worst-case traffic utilization analysis, including failure scenario evaluation, simulation controls, result visualization, and the interpretation of updated table columns for interfaces and circuits, along with guidance on reviewing QoS impacts and VPN-specific calculations.

Worst case traffic utilization is the highest utilization that a particular interface experiences over all the failure sets and traffic levels that you select.

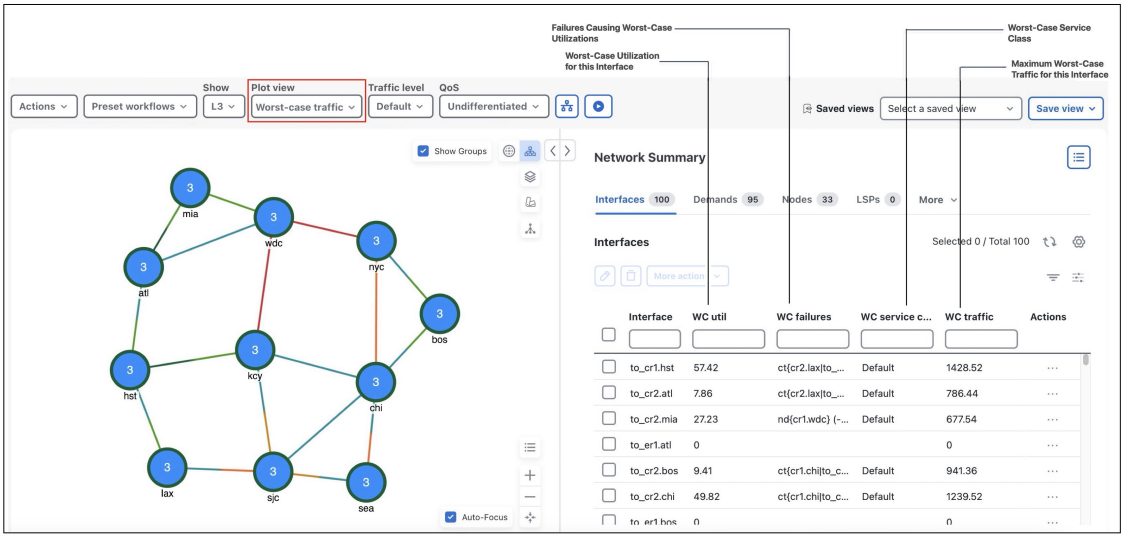
### Analysis behavior

Cisco Crosswork Planning determines which combination of failures would cause this worst-case utilization. The default analysis is to identify up to 10 failures for the worst-case utilization on each interface in the network. Alternatively, you can record failures causing utilizations within a specified percent of the worst-case utilization.

To control the number of threads that Cisco Crosswork Planning processes in parallel when examining failure scenarios, set the **Maximum number of threads** field. For more information on how to run the simulation analysis, refer to Run simulation analysis on page 135.

After finishing the analysis, Cisco Crosswork Planning switches to the Worst-Case traffic view and updates the plot to simultaneously display the worst-case utilization for all interfaces.

**Figure 39: Worst-case traffic utilization**

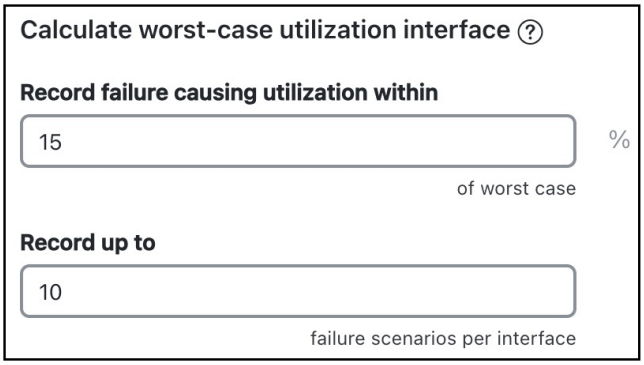



### Updated table columns

In addition to Plot view changes, these columns are also updated in the Interfaces and Circuits tables after the Simulation analysis completes:

- **WC util**: The worst-case utilization for that interface. The worst-case for a circuit is defined to be whichever of the worst cases of the two constituent interfaces results in the larger utilization. Thus, for circuits, this value is the larger of the WC util values for the two interfaces in the circuit.
- **WC traffic**: The actual traffic (Mbps) through the interface under the worst-case scenario.
- **WC traff level**: The traffic level under which this worst-case scenario occurs.
- **WC failure**: List of one or more failures that cause the worst-case failure of the circuit. An easier way to read this list is to select an interface and use ••• > **Fail to WC**.

  If you record failures causing utilizations within a given percent of worst case, this column shows QoS violations as a percent (refer to Worst-case QoS violation on page 132). If the number is positive, then the allotted capacity has been surpassed. If negative, the capacity has not been surpassed. For example, if a circuit has 10,000 Mbps capacity, and if the amount of traffic on it as a result of three different failures is 11,000, 8000, and 4000 Mbps, the utilizations are 10%, -20%, and -60%, respectively, and in descending order.




- **WC service class**: The service class for which this worst-case scenario occurs.

For information on running a Simulation analysis with QoS, refer to Worst-case QoS violation on page 132.

For information on worst-case calculations for VPNs, refer to Simulate VPN on page 197.

## Worst-case QoS violation

Describes how worst-case QoS violations are calculated and identified when QoS parameters are configured.

Cisco Crosswork Planning includes QoS bound (maximum available capacity) as part of the worst-case calculations. If there are no QoS parameters set, then the QoS bound is 100% and violations occur if utilization goes over that 100%. However, if a worst-case policy has been set on a service class or if interface queue parameters have been set, then worst-case QoS violations are calculated. In these instances, Cisco Crosswork Planning identifies the interface with the highest percentage of QoS violation as the worst-case possibility.

### QoS violation parameters

These columns are updated:

- **WC QoS bound**: The worst-case interface capacity available without violating these QoS requirements. This value is based on available capacity, traffic utilization, worst-case policies set on service classes, and interface queue parameters.

  The WC QoS bound (%) column identifies this same value as a percentage of the total capacity.

- **WC QoS violation**: The worst-case traffic minus the worst-case capacity permitted (WC QoS bound). A violation occurs if the QoS capacity allotted through worst-case policies for service classes is exceeded or if QoS capacity allotted through interface queue parameters is exceeded. If the number appearing in the WC QoS violation column is positive, then the allotted capacity has been surpassed. If negative, the capacity has not been surpassed.

  The WC QoS violation (%) column identifies this same value as a percentage of total capacity.

- **WC service class**: The service contributing to the worst-case QoS violation.

To see the cause of worst-case QoS violations, select a circuit and use ••• > **Fail to WC**. The page that appears lists all causes of this circuit's worst-case utilization and its worst-case QoS violations. Choose the worst-case failure to view, and click **Submit**.

| For more information on... | See... |
|---|---|
| • QoS parameters and QoS calculations<br>• Set worst-case policies on service classes<br>• Set interface queue parameters | Simulate Quality of Service (QoS) on page 183 |
| Worst-case QoS calculations for VPNs | Simulate VPN on page 197 |

## Fail circuits to worst-case utilization

Describes how to display worst-case utilization scenarios after running simulation analysis.

Use this task to selectively view each failure scenario that causes the worst-case utilization or worst-case QoS violation for a single interface after running simulation analysis.

If there are multiple possibilities for a worst-case failure, or if there is a range of failures within a percentage of the worst-case failure (listed in the **WC failures** column), the Fail to WC page lists each failure, its worst-case utilization percent, and its QoS violation percent (refer to Figure 40: Failure of a single circuit to its worst-case on page 133).

> **Note**
>
> If you choose an interface, you are actually failing its associated circuit to its worst case.

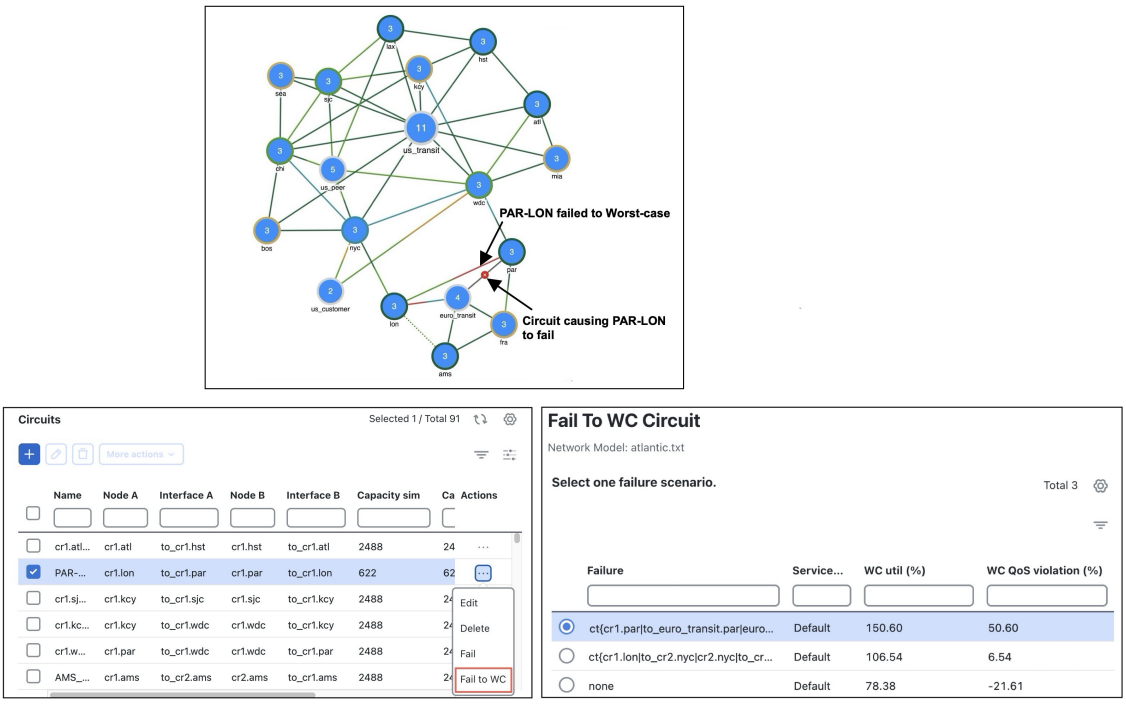
**Procedure**

1. In the **Worst-case traffic** view, select the desired interface or circuit from their respective tables.
2. From the **Actions** column, choose ••• > **Fail to WC**.

   The Fail to WC Interface (or Circuit) page appears.
3. Choose the failure scenario of interest (refer to Figure 40: Failure of a single circuit to its worst-case on page 133).
4. Click **Submit**.

The network plot changes to show this particular failure scenario.

**Figure 40: Failure of a single circuit to its worst-case**




## Worst-case demand latency

Explains the simulation analysis feature that calculates maximum latency for each demand under failure scenarios.

When running Simulation analysis (refer to Run simulation analysis on page 135), you have the option to simulate worst-case latency for each demand in the plan. Cisco Crosswork Planning calculates the maximum latency of each demand under the failure scenarios selected. The result does not depend on service classes or traffic levels because demand routing is independent of these plans. The simulation also records the failures that cause this maximum latency.

#### Worst-case demand latency configuration

These columns in the Demands table are updated when you check the **Calculate demand worst-case latency** check box (under the **Calculate worst-case utilization interface** section) while running the Simulation analysis tool:

- **WC latency**: The highest demand latency over all failure scenarios in the analysis.
- **WC latency failures**: The failures that caused this worst-case latency. Up to 10 failures are identified.

Cisco Crosswork Planning captures the latencies of each demand for each failure case included in Simulation analysis using these two options:

- **Record failures causing demand latency within __ % of worst case**: Records failures causing demand latency within the specified percentage range of the worst-case latency. Default is 0. If you enter 0, only the worst case latency failures are recorded.
- **Record up to __ failure scenarios on demand latency**: Maximum number of failure scenarios to record per demand. Default is 1.

If you record failures causing demand latency within a given percent of worst case, the WC latency failures column shows the WC latency along with failure scenarios.

## Fail demands to worst-case latency

Describes how to fail a demand to its worst-case latency after running simulation analysis.

Failing a demand to its worst-case latency allows you to analyze specific failure scenarios that cause the highest latency.

After running Simulation analysis to calculate demand worst-case latency, you have the option to fail a single demand to its worst-case latency. The failure scenarios causing the worst-case latency are listed in the **WC latency failures** column of the Demands table.

> **Note**
>
> If you have not selected the **Calculate demand worst-case latency** check box while running Simulation analysis and if the Plot view is not **Worst-case traffic**, you will not see the option to fail the demand to its worst-case latency.

### Before you begin

- Select the **Calculate demand worst-case latency** check box during simulation analysis (refer to Run simulation analysis on page 135).
- Ensure the Plot view is set to **Worst-case traffic**.

Complete these steps to fail a demand to its worst-case latency.

#### Procedure

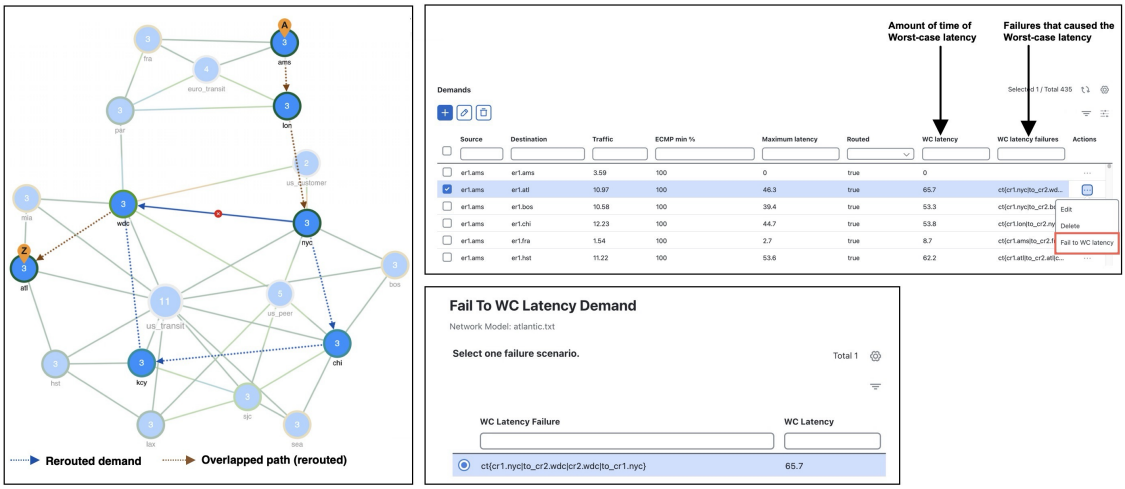
1. In the **Worst-case traffic** view, select the desired demand from the **Demands** table.
2. From the **Actions** column, choose ··· > **Fail to WC latency**.

   The Fail to WC Latency Demand page appears.
3. Select the failure scenario of interest (refer to Figure 41: Example of worst-case demand latency on page 135).

4. Click **Submit**.

The network plot changes to show this particular failure scenario.

**Figure 41: Example of worst-case demand latency**




## Run simulation analysis

Guides you through running simulation analysis to assess worst-case interface and VPN utilization and demand latency.

Use this task to run a simulation analysis to understand how failures affect network utilization, latency, and capacity.

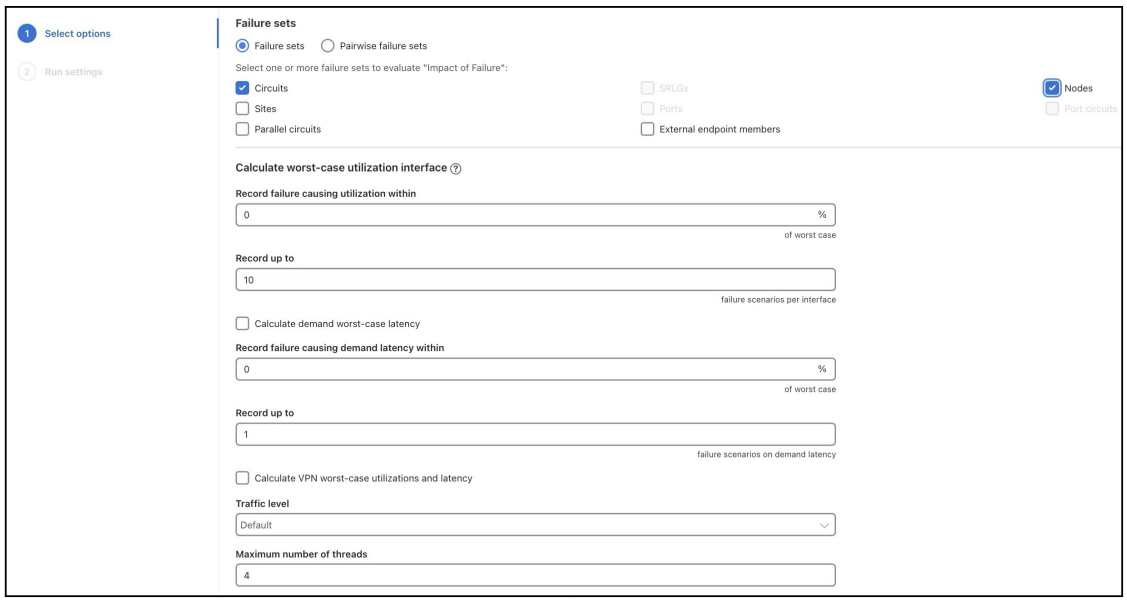
The **Simulation analysis** tool supports four failure analysis options: worst-case utilization on interfaces, worst-case VPN utilization and latency, worst-case demand latency, and failure impact.

> **Note**
>
> Recording worst-case demand latencies or VPN worst-case utilizations increases the time it takes to perform a worst-case analysis.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Tools** > **Simulation analysis**. Alternatively, you can also choose **Preset workflows** > **Evaluate impact of failures**.

   **Figure 42: Simulation analysis options**




3. In the **Failure sets** section, choose one or more failure sets.

   To perform pairwise failure analysis, choose the **Pairwise failure sets** option and select the required objects. For more information, refer to Perform pairwise failure analysis on page 137.

4. Configure these parameters as required:
   a) In the **Record failures causing utilizations within __% of worst case** field, enter 0 to record only worst-case failures, or enter a number to find all failures causing utilizations within that percentage range of the worst-case failure.
   b) In the **Record up to __ failure scenarios per interface** field, enter the maximum number of failure scenarios to record per interface. The default is 10.

      **Example**
      If you record failures causing utilizations within 10% of the worst case, and the worst-case utilization for an interface is 90%, then Cisco Crosswork Planning records failures on this interface resulting in utilization of 81% or higher (90 -(90/10)). In this same scenario, if you record 10 failure scenarios per interface, and there are failures that could cause utilizations of 90%, 85%, 82%, and 76% for an interface, Cisco Crosswork Planning does not record the failure causing 76% utilization.
   c) Select whether or not to record demand worst-case latency calculations using the **Calculate demand worst-case latency** check box.
   d) In the **Record failures causing demand latency within __ % of worst case** field, enter 0 to record only worst case latency failures, or enter a number to find all failures causing demand latency within that percentage range of the worst-case latency.
   e) In the **Record up to __ failure scenarios on demand latency** field, enter the maximum number of failure scenarios to record per demand. The default is 1.

      **Example**
      If you record failures causing demand latency within 10% of the worst case, and the worst-case latency for a demand is 100 ms, then Cisco Cisco Crosswork Planning records failure scenarios which have the latency of 90 ms or higher (100-(100/10)) on this demand. In this same scenario, if you record 5 failure scenarios per demand latency, and there are failures that could cause latency of 92 ms, 95 ms, 98 ms, and 80 ms for a demand, Cisco Cisco Crosswork Planning does not record the failure causing 80 ms demand latency.
   f) Select whether or not to record VPN worst-case utilizations and latencies using the **Calculate VPN worst-case utilizations and latency** check box. For more information, refer to Simulate VPN on page 197.

5. Select one or more traffic levels.

6. Enter the value of maximum number of threads in the **Maximum number of threads** field and click **Next**.
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

   The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

   > **Note**
   >
   > Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. Submit your changes.

You can now use the Worst-case traffic view (from the **Plot view** drop-down, choose **Worst-case traffic**) to analyze the worst-case traffic utilization, worst-case QoS violation, and worst-case latency information. You can fail interfaces and nodes to their worst case, and fail demands to their worst-case latency. You can also use the Failure impact view on page 144 view to identify the circuits that are responsible for worst-case traffic congestion.

## Perform pairwise failure analysis

Explains how to configure the Simulation analysis tool to analyze network models for concurrent failures by selecting pairs of failure sets and choosing from several pairwise failure simulation variations.

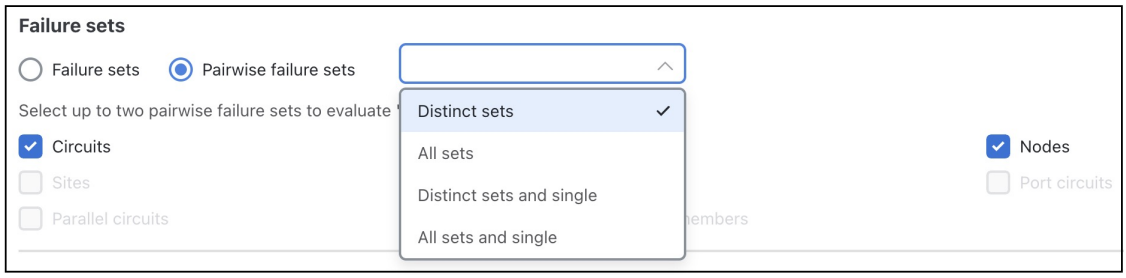
In addition to selecting the failure sets individually, the Simulation analysis tool allows you to select a pair of failure sets. This enables you to analyze a network model for instances when two concurrent failures may cause significant impact. You can select up to two failure sets for pairwise failure analysis and choose from several variations of pairwise failures to simulate.

Complete these steps to perform pairwise failure analysis.

### Procedure

1. In the **Simulation analysis** tool, choose the **Pairwise failure sets** option.
2. Select up to two failure sets.
3. Choose one of the four options from the drop-down list.

**Figure 43: Pairwise failure sets options**



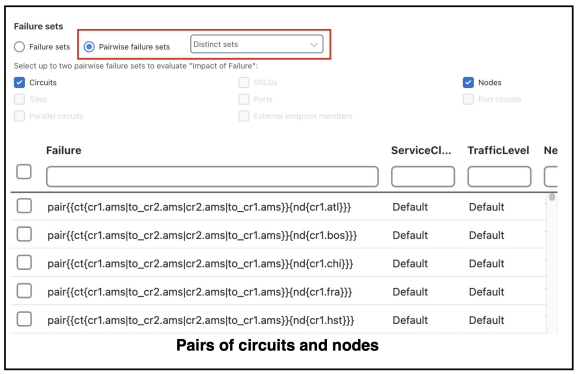

- **Distinct sets**: This option creates pairwise failures using elements from distinct failure sets. For example, if you select circuit and node as failure sets, this option considers the combination of distinct sets (circuit, node) for pairwise failure analysis.
- **All sets**: This option creates pairwise failures using elements from failure sets, including pairs of elements from the same failure set. For example, if you select circuit and node as failure sets, this option considers all combinations for failure analysis, that is, (node, node), (circuit, circuit), and (circuit, node).
- **Distinct sets and single**: This option is similar to **Distinct sets**, but in addition to pairwise failures, it simulates the failure of individual elements in the failure sets. For example, if you select circuit and node as failure sets, this option considers combination of distinct sets and individual elements for failure analysis, that is, (circuit), (node), and (circuit, node).
- **All sets and single**: This option is similar to **All sets**, but in addition to pairwise failures, it simulates the failure of individual elements in the failure sets. For example, if you select circuit and node as failure sets, this option considers all combinations of failure sets and individual elements for failure analysis, that is, (circuit), (node), (node, node), (circuit, circuit), and (circuit, node).

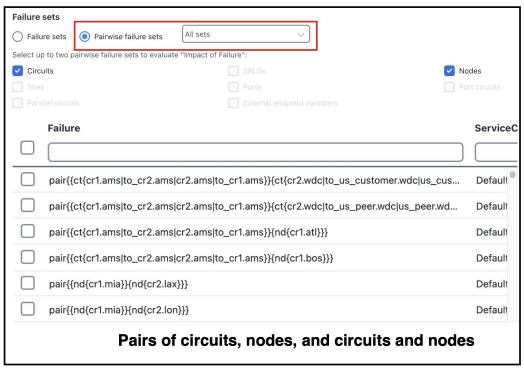
**Note**

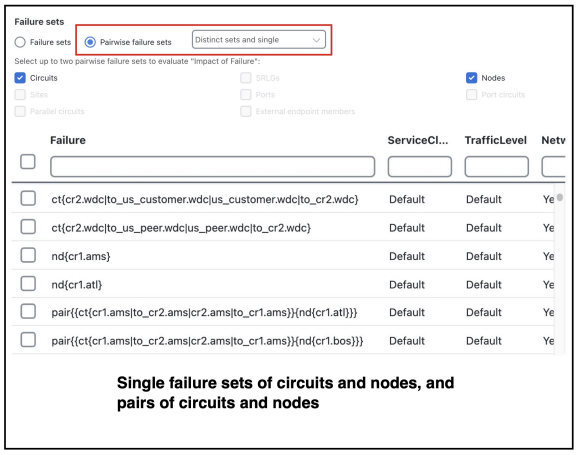
- The memory consumption and performance of the simulation analysis can be significantly impacted as the number of failure scenarios can grow dramatically with pairwise failure sets. For better performance, it is recommended to use the **Distinct sets** or **Distinct sets and single** options, rather than the **All sets** and **All sets and single** options.
- Be cautious when using pairwise failure sets, as they execute significantly more failure scenarios compared to regular failure sets. For instance, in a 100-node scenario, regular failure sets would execute 100 scenarios, whereas pairwise failure sets (node, node) would result in 9,900 scenarios.

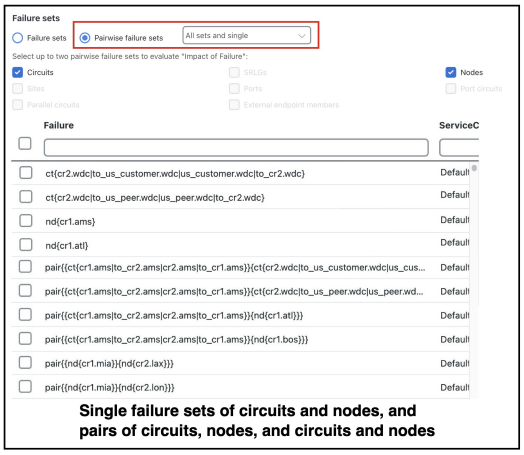
The **Simulations** tab in the generated report provides details of the list of single or paired objects considered in the failure analysis.

**Figure 44: Pairwise failure sets selection and results**




### Pairwise failure analysis example

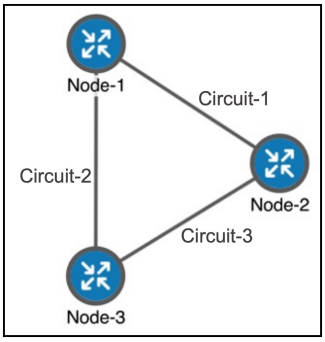
Illustrates how the Simulation analysis tool evaluates pairwise failure scenarios using a network example.

This example demonstrates how the selection of failure sets and pairwise failure options determines which failure scenarios the Simulation analysis tool evaluates.

To illustrate how the pairwise failure options work, consider a simple network topology with:

- Three nodes: Node-1, Node-2, Node-3
- Three circuits (links): Circuit-1, Circuit-2, Circuit-3

**Figure 45: Sample network topology**




This section examines the failure scenarios generated under these two conditions:

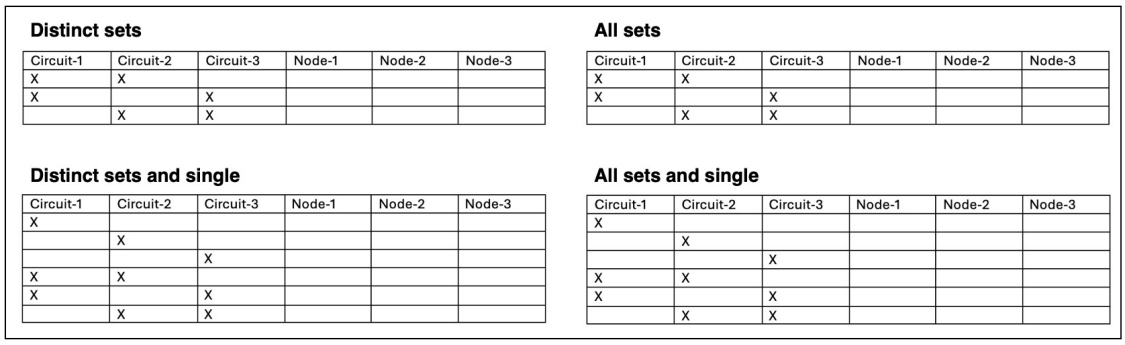
- When only circuits are selected as failure sets
- When both circuits and nodes are selected

**Scenario 1: Pairwise failure analysis with only "Circuits" selected**

When only "Circuits" are chosen as the failure set, the analysis focuses purely on failures involving Circuit-1, Circuit-2, and Circuit-3.

| Pairwise failure option | Action | Generated failure scenarios |
|---|---|---|
| Distinct sets | Simulates the simultaneous failure of two distinct circuits. | (Circuit-1, Circuit-2), (Circuit-1, Circuit-3), (Circuit-2, Circuit-3) |
| All sets | Simulates the simultaneous failure of any two circuits.<br>**Note**<br>The results are same as the "Distinct sets" option because only one failure set was selected. | (Circuit-1, Circuit-2), (Circuit-1, Circuit-3), (Circuit-2, Circuit-3) |
| Distinct sets and single | Simulates individual circuit failures and the simultaneous failure of two distinct circuits. | • Single failures: (Circuit-1), (Circuit-2), (Circuit-3)<br>• Distinct pairs: (Circuit-1, Circuit-2), (Circuit-1, Circuit-3), (Circuit-2, Circuit-3) |
| All sets and single | Simulates individual circuit failures and all possible circuit pair failures.<br>**Note**<br>The results are same as the "Distinct sets and single" option because only one failure set was selected. | • Single failures: (Circuit-1), (Circuit-2), (Circuit-3)<br>• Distinct pairs: (Circuit-1, Circuit-2), (Circuit-1, Circuit-3), (Circuit-2, Circuit-3) |

**Figure 46: Summary of failure scenarios generated**

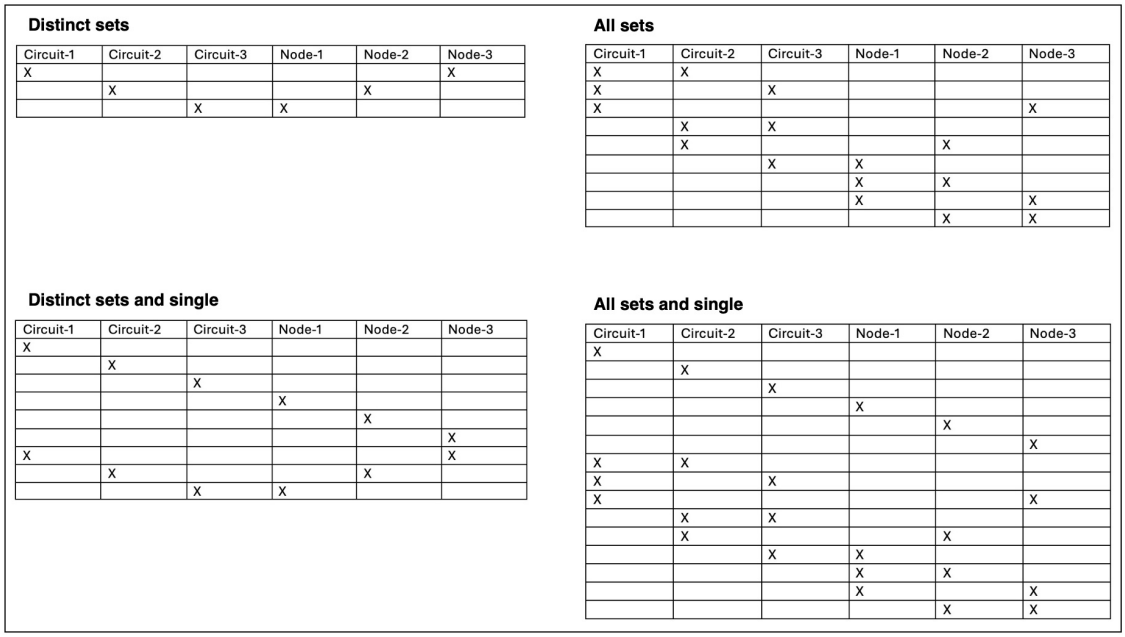


**Scenario 2: Pairwise failure analysis with "Circuits" and "Nodes" selected**

Selecting both "Circuits" (Circuit-1, Circuit-2, Circuit-3) and "Nodes" (Node-1, Node-2, Node-3) as failure sets significantly increases the number and types of failure combinations.

| Pairwise failure option | Action | Generated failure scenarios |
|---|---|---|
| Distinct sets | Simulates the simultaneous failure of distinct circuit-node pairs.<br>In Figure 45: Sample network topology on page 139, Circuit-1 is not connected to Node-3. Similarly, Circuit-2 is not connected to Node-2, and Circuit-3 is not connected to Node-1. Therefore, these pairs are considered as distinct pairs in the failure scenarios. | (Circuit-1, Node-3), (Circuit-2, Node-2), (Circuit-3, Node-1) |
| All sets | Simulates the simultaneous failure of two elements chosen from the combined set of all circuits and all nodes. This includes:<br>• Pairs of circuits<br>• Pairs of nodes<br>• Distinct circuit and node pairs | • Circuit-Circuit pairs: (Circuit-1, Circuit-2), (Circuit-2, Circuit-3), (Circuit-1, Circuit-3)<br>• Node-Node pairs: (Node-1, Node-2), (Node-2, Node-3), (Node-1, Node-3)<br>• Distinct Circuit-Node pairs: (Circuit-1, Node-3), (Circuit-2, Node-2), (Circuit-3, Node-1) |
| Distinct sets and single | Simulates individual failures of circuits and nodes, and the simultaneous failure of distinct circuit-node pairs. | • Single failures: (Circuit-1), (Circuit-2), (Circuit-3), (Node-1), (Node-2), (Node-3)<br>• Distinct Circuit-Node pairs: (Circuit-1, Node-3), (Circuit-2, Node-2), (Circuit-3, Node-1) |
| All sets and single | Simulates individual failures of circuits and nodes, and the simultaneous failure of any two elements from the combined set of all circuits and all nodes.<br>This option generates the largest number of scenarios. | • Single failures: (Circuit-1), (Circuit-2), (Circuit-3), (Node-1), (Node-2), (Node-3)<br>• All pairs: (Circuit-1, Circuit-2), (Circuit-2, Circuit-3), (Circuit-1, Circuit-3), (Node-1, Node-2), (Node-2, Node-3), (Node-1, Node-3), (Circuit-1, Node-3), (Circuit-2, Node-2), (Circuit-3, Node-1) |

**Figure 47: Summary of failure scenarios generated**



**Distinct sets**

| Circuit-1 | Circuit-2 | Circuit-3 | Node-1 | Node-2 | Node-3 |
|---|---|---|---|---|---|
| X | | | | | X |
| | X | | | X | |
| | | X | X | | |

**All sets**

| Circuit-1 | Circuit-2 | Circuit-3 | Node-1 | Node-2 | Node-3 |
|---|---|---|---|---|---|
| X | X | | | | |
| X | | X | | | |
| X | | | | | X |
| | X | X | | | |
| | X | | | X | |
| | | X | X | | |
| | | | X | X | |
| | | | X | | X |
| | | | | X | X |

**Distinct sets and single**

| Circuit-1 | Circuit-2 | Circuit-3 | Node-1 | Node-2 | Node-3 |
|---|---|---|---|---|---|
| X | | | | | |
| | X | | | | |
| | | X | | | |
| | | | X | | |
| | | | | X | |
| | | | | | X |
| X | | | | | X |
| | X | | | X | |
| | | X | X | | |

**All sets and single**

| Circuit-1 | Circuit-2 | Circuit-3 | Node-1 | Node-2 | Node-3 |
|---|---|---|---|---|---|
| X | | | | | |
| | X | | | | |
| | | X | | | |
| | | | X | | |
| | | | | X | |
| | | | | | X |
| X | X | | | | |
| X | | X | | | |
| X | | | | | X |
| | X | X | | | |
| | X | | | X | |
| | | X | X | | |
| | | | X | X | |
| | | | X | | X |
| | | | | X | X |

## Protect objects

Explains how to configure object protection to exclude objects from Simulation analysis failure lists by marking them as Protected.

To exclude an object from the list of those objects failed when performing a Simulation analysis, you can mark it as *Protected* in its Edit page. For example, if you want to run a Simulation analysis only on core nodes, you could first protect all edge nodes.

You can protect nodes, sites, circuits ports, port circuits, external endpoint members, and parallel circuits.




**Note**

If you select an interface, you are actually protecting its associated circuit.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. Select one or more like objects from their respective tables.
3. Click .

   **Note**

   If you are editing a single object, you can also use the ··· > **Edit** option under the **Actions** column.

4. In the **State** field, check the **Protected** check box.

5. Save the changes.

The selected objects are marked as Protected and will be excluded from Simulation analysis failure lists.

## Simulation analysis report

Provides information about simulation analysis reports that are automatically generated each time Simulation analysis is run.

Each time Simulation analysis is run, a report is automatically generated. To view the report, select **Actions > Reports > Generated reports** and then click the **Simulation Analysis** link in the right panel.

- The **Options** tab displays the input parameters used for the simulation analysis.
- The **Summary** tab details the options used in the analysis. It also summarizes the most important problems identified, such as QoS violations and latency bound violations.
- The **Max Utilization** tab shows the impact of failures on maximum utilization in the form of a pie chart.
- The **Simulations** table lists each simulation that was performed in the Simulation analysis.

Note that new reports replace the previous ones.

**Table 8: Simulations table in simulation analysis report**

| Simulation Data Point | Description |
|---|---|
| Failure | Failure scenario used in the analysis. |
| ServiceClass | Service class used in the analysis. |
| TrafficLevel | Traffic level used in the analysis. |
| NetworkBreakpoint | Identifies whether there are network breaks resulting from the failures in this simulation. If multiple network breakpoints occur, the most serious one is listed.<br>• Yes (Total): A break exists that completely partitions the network into two or more disconnected sections.<br>• Yes (AS): A break exists that completely partitions an AS into two or more sections. However, routes exist between the sections of the AS through other ASes.<br>• Yes (OSPF Area 0): A break exists that completely partitions Area 0 of an AS running OSPF. Under OSPF, traffic cannot route between the partitions even if a path is available through non-zero areas in the AS.<br>• No: No break in the network. |
| NumUnroutedDemands | Number of demands that cannot be routed under this failure for any of the reasons identified by the network breakpoint. |
| UnroutedTraffic | Total amount of demand traffic that cannot be routed under this failure for any of the reasons identified by the network breakpoint. |
| MaxUtil | Maximum utilization over all interfaces in this simulation. Utilization is the traffic through the interface as a percentage of the capacity of the interface. |

| Simulation Data Point | Description |
|---|---|
| MaxQoSboundpercent | Worst-case capacity available without violating QoS bounds, expressed as a percentage of the total capacity. |
| NumQoSviolations | Number of times the QoS bound is violated. QoS bounds are set through service class policies and interface queue parameters. |
| LatencyBoundViolations | Number of demands with maximum latency in excess of the latency bound specified for the demand. |
| NumUnRoutedLSPs | Number of unrouted non-Fast Reroute (FRR) LSPs in the analysis. |
| NumUnroutedFRRLSPs | Number of unrouted FRR LSPs in the analysis. |
| NumRoutedCSRSVPLSPs | Number of routed CS-RSVP LSPs in the analysis. |
| NumUnroutedCSRSVPLSPs | Number of unrouted CS-RSVP LSPs in the analysis. |
| NumNonDisjointCSRSVPLSPs | Number of CS-RSVP LSPs with disjointness violations. |
| NumNodeToLinkDSJCSRSVPLSPs | Number of CS-RSVP LSPs falling back from node disjointness to link disjointness. |

## Failure impact view

Describes that failure impact view that colors nodes and circuits according to the maximum utilization level that would occur elsewhere in the network if the node or circuit fails.

A failure impact view is a network plot view that

- colors nodes and circuits according to the maximum utilization level that would be caused elsewhere in the network should the node or circuit fail
- indicates the resulting utilization and severity of the congestion through color representation, and
- becomes available upon running a Simulation analysis.

#### Columns related to failure impact

The Node, Interface, and Circuit tables contain the **Failure impact** and **Failure impact interface** columns. In the Interfaces table, the information describes the failure impact of the circuit containing the interface.

- **Failure impact**: The failure impact of each node or circuit. For example, if the value is 80%, it means that if this node or circuit failed, the resulting traffic utilization on one or more interfaces would exceed 80%.
- **Failure impact interface**: The interface that will experience the highest utilization as a result of the node or circuit going down.

Format = if{Node|Interface}

For example, if{cr2.sjc|to_cr1.kcy} means if the circuit goes down, it will have the greatest traffic impact on the cr2.sjc to cr1.kcy interface.

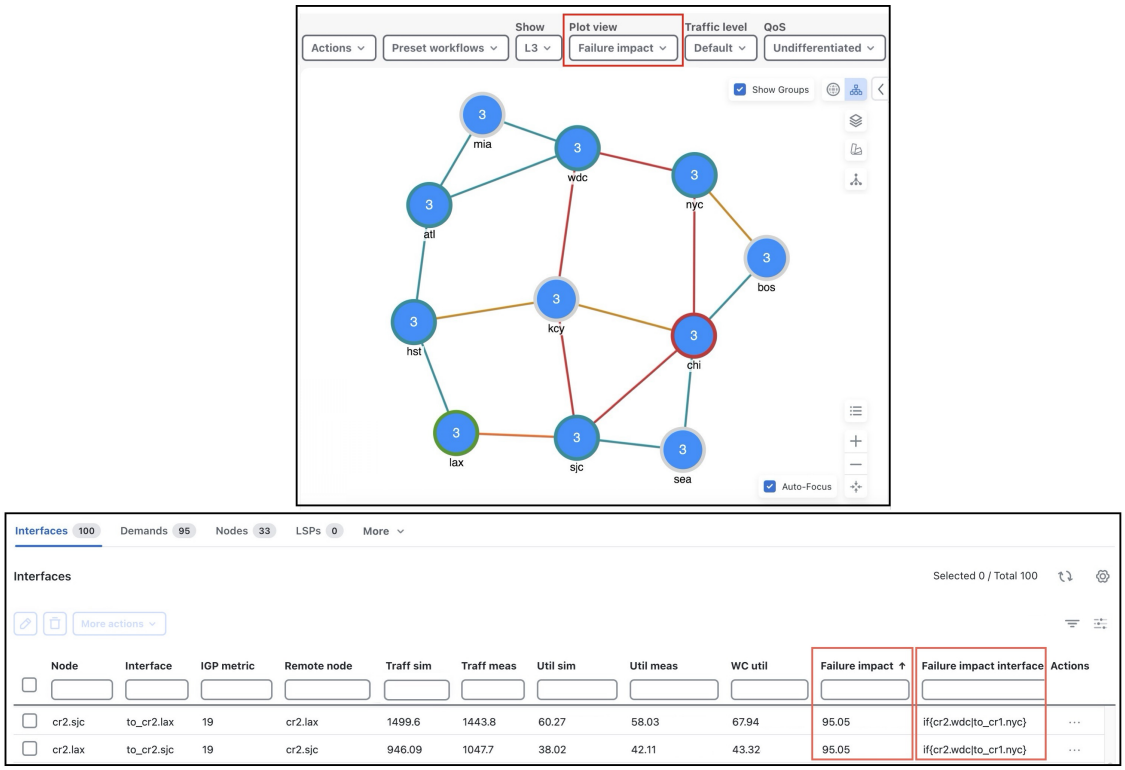
In the network plot, the Site borders show the maximum utilization level that would be caused elsewhere in the network should nodes within it fail or should intra site circuits within it fail.

> **Note**
>
> The Failure impact view only shows the impact of circuit and node failures. It does not show failures of other objects.

**Failure impact view**

In Figure 48: Failure impact view on page 145, a sjc-lax circuit has a utilization of 90-100% and its color representation is orange red. This means that if sjc-lax were to fail, one or more interfaces would react by exceeding a 90% utilization level and correspondingly, would turn orange red in the plot.

**Figure 48: Failure impact view**




# Configure parallelization for simulation analysis

Explains how to configure parallelization to speed up computation and achieve faster results for large network models in the Simulation analysis tool.

Parallelization allows the Simulation Analysis tool to split computation across multiple engines, resulting in faster processing times for large network models with thousands of circuits.

The Simulation Analysis tool supports parallelizing computation to achieve faster results for large network models. For example, if there are 10000 circuits in a network model and 10 different engines are available, the tool can partition the network model into 10 sections with each partition handling 1000 failure scenarios. This creates 10 different result files that are later merged to obtain the final result.

**Procedure**

1. From the main menu, choose **Job Manager**.

2. Click  > **Using CLI**.
3. Choose **Simulation analysis** and click **Next**.
4. Select the network model in which you want to run the simulation analysis and click **Next**.
5. In the **Simulation Analysis input options** field, enter this command.

```
-failure-sets <failure-sets> -num-partitions <number-of-partitions>
-num-threads <number-of-threads> -partition-index <partition-index>
-result-file <result-filename>
```

Parameter definitions:

- **-num-partitions**: Number of partitions of the failure scenarios. Each partition has an associated set of failure scenarios and is identified by an index ranging from 0 up to the number of partitions minus 1. Default is 1.
- **-partition-index**: Simulate the set of failure scenarios belonging to the specified partition. Default is 0.
- **-result-file**: If specified, the simulation analysis report results are written to this file. The file can be a *.txt or a *.db file.

6. To merge the results, invoke the merge_sim_analysis CLI using the Python script. For more information on running tools and initializers using CLI, refer to Run tools or initializers using CLI on page 385.

```
import os
import sys

cmd = "merge_sim_analysis -plan-file {0} -partial-results {1} -out-file
out-plan.txt ".format( sys.argv[1], sys.argv[2])
print(cmd)
os.system(cmd)
```

Parameter definitions:

- **-plan-file**: Input plan file.
- **-out-file**: Output plan file.
- **-partial-results**: Comma separated list of files containing simulation analysis results for each partition. These may be plan files or files generated using the **-result-file** option while running the Simulation analysis CLI tool.
- **-partial-result-paths-file**: File containing list of files, one per line, with simulation analysis results for each partition. These may be plan files or files generated using the **-result-file** option while running the Simulation analysis CLI tool. This is ignored if the **-partial-results** option is specified.

7. In the Job Manager, enter the script arguments while running the script (run_cli_merge.py).

```
run_cli_merge.py input_planfile.pln res_0.txt,res_1.txt,res_2.txt
```

# 8 Evaluate Impact of Traffic Growth

**Topics:**


- Demand groupings
- Growth plans


Outlines the tools and methodologies for forecasting traffic growth, developing and managing demand groupings, utilizing traffic tables, and constructing growth plans to predict and analyze future network demands.

## Demand groupings

Explains how demand groupings aggregate network traffic for reporting and growth planning.

A demand grouping refers to a set of filters (rules) that

- defines a group of demands, and
- provides a convenient way of specifying aggregated traffic in a plan, which can be used as a basis for traffic reports and growth plans.

> **Note**
>
> You can only view the Demand grouping details if they are already present in your plan file. You cannot create, edit, or delete the Demand groupings from the UI.

You can use demand groupings to represent

- total traffic sourced from one specific site
- total VPN traffic, or any defined service class, sourced from one specific site to another, or
- total traffic destined for a particular AS.

## Growth plans

Explains how growth plans predict future network traffic using various forecasting methods and help incorporate performance simulation into planning decisions.

The **Create growth plans** tool lets you generate new plan files containing predictions of future traffic, based on growth forecasts for the current plan.

You can generate the predicted traffic growth in one of three ways:

- Per demand grouping using demand grouping traffic growth rates: refer to Growth plans from demand grouping on page 150.
- Per demand using demand traffic growth rates: refer to Create growth plans from demand growth on page 150.
- Per interface using interface traffic growth rates: refer to Create growth plans from interface growth on page 151.

Regardless of the method, the growth is compounded per period increment identified in the Create growth plans tool.

A series of plan files is created for each growth period, and each plan contains a report on the prediction process used to generate it. Each newly created plan shows the effects of the traffic growth visually in the plot using default traffic utilization colors, as well as in the Traffic column in the Demands table. You can incorporate worst-case performance into planning decisions by automatically running a simulation analysis on the new plans.

### Create growth plans

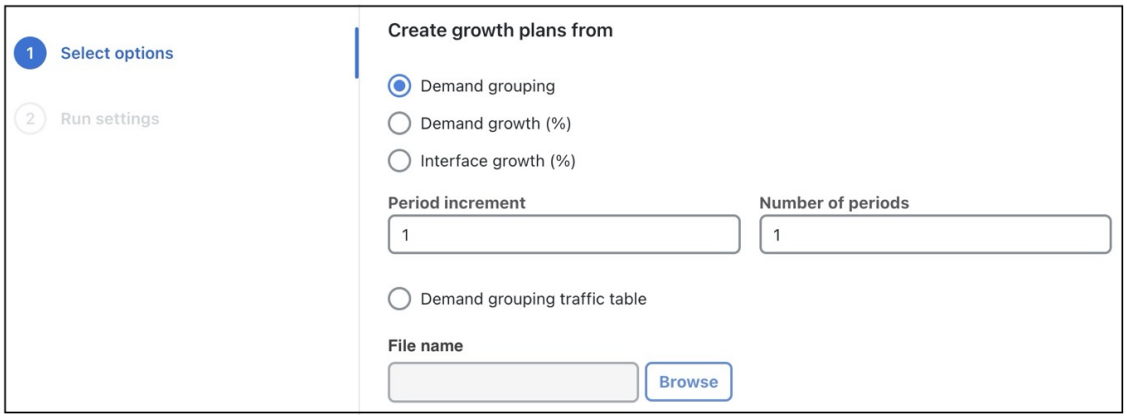
Guides you through creating growth plans to predict network infrastructure needs based on potential traffic increase.

Use this task to create growth plans that predict network infrastructure requirements based on potential traffic increases.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Tools** > **Create growth plans**. Alternatively, you can choose **Preset workflows** > **Evaluate impact of traffic growth**.

   **Figure 49: Create growth plans settings**




3. Specify one of these growth methods:
   - Demand grouping: For information on how to set demand grouping growth based on fixed, incremental, or percents, refer to Growth plans from demand grouping on page 150.
   - Demand growth (%): For information on how to set demand growth, refer to Create growth plans from demand growth on page 150.
   - Interface growth (%): For information on how to set interface growth, refer to Create growth plans from interface growth on page 151.
   - Demand grouping traffic table: Browse to or enter the name of the file containing the Demand grouping traffic table. For information on Demand grouping traffic table, refer to Demand grouping traffic table.
4. Enter the number of plans to create and the increment by which you want to create them. These options are not applicable, if you selected the Demand grouping traffic table option in Step 3.
   - Period increment: Number of percentage increments per period. For example, if a demand is set to grow at 10%, and the period increment is 2, the demand will grow by 21% in the first growth plan.
   - Number of periods: Number of growth plan files to create.

   Plans are named using the period increment. For example, if you enter 4 as the increment and 2 as period, Cisco Crosswork Planning generates two growth plan files: one ending with _4 and one ending with _8.
5. Click **Next**.
6. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.

- Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
- Schedule: Set the time at which you want to run the tool.

The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

> **Note**
>
> Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

7. Submit your changes.

---

The generated growth plan is saved under **Network Models > User space > Other files**. To view the details, download the file (.tar file), extract it, and import the updated plan file to the user space (for more information, refer to Import plan files from the local machine on page 34).

## Growth plans from demand grouping

Explains how network growth plans are constructed by grouping traffic demands and allocating growth forecasts accordingly.

Growth plans for network traffic, such as those produced either by trending analyses, or by financial or sales forecasts, are typically created for traffic aggregates. You can specify these aggregate growth forecasts or trending estimates as demand grouping traffic growth information. The **Create growth plans** tool allocates traffic to demands to match the aggregate traffic growth specified for each demand grouping. If two demand groupings contain a common demand (for example, a grouping of demands sourced from one site, and a grouping of demands destined for another site), this demand traffic is balanced between the requirement to match each of the demand grouping traffic aggregates.

> **Note**
>
> You cannot create, edit, or delete the Demand groupings from the UI. You can only view the Demand grouping details in the UI if it's already present in your plan file.

## Create growth plans from demand growth

Outlines the process for creating growth plans from demand growth by editing the Growth (%) column in the Demands table.

This method grows demand traffic based on the value identified in the **Growth (%)** column of the Demands table in the current network model. This approach is useful when all traffic in the network, all traffic within a given traffic level, or all traffic within a service class is forecast to grow at the same rate. When generating plans, you can choose how many plans to create. Each plan shows growth at the specified rate compared to the previous plan.

Complete these steps to populate the **Growth (%)** column in the Demands table.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. Select one or more demands from the Demands table.
3. Click .
4. In the **Traffic** section, enter the growth rate in the **Growth %** column.
5. Save your changes.

If traffic on a demand is 1000 Mbps, if the growth for that demand is defined as 10%, and if you create two growth plans, the first will show 1100 Mbps and the second will show 1210 Mbps on that same demand, representing 10% growth per period.

## Create growth plans from interface growth

Guides you through creating growth plans based on measured interface traffic, configuring growth rates in the Interfaces table.

This method grows measured traffic based on the **Growth (%)** column in the Interfaces table in the current plan file. It differs from demand-based growth methods because this method grows measured traffic rather than demands (simulated traffic). This method is useful for basic forecasting of the current network state, rather than worst-case failure states, which require simulation. When generating the plans, you can choose how many plans to create, and each plan shows growth at the specified rate compared to the previous plan.

You can also use this method in conjunction with the Demand deduction tool to create demands from the forecast measurements.

Complete these steps to populate the **Growth (%)** column in the Interfaces table.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the toolbar, choose **Actions** > **Edit** > **Traffic** > **Measured traffic**.

   The Measured traffic page opens.




3. For each applicable interface, click the **Edit** button and in the **Growth (%)** column, enter the percent by which you want to increase the interface traffic.
4. Save your changes.

If the traffic on an interface is 525 Mbps, if the growth for that interface is defined as 5%, and you specify three growth plans, the first will show 551.25 Mbps, the second will show 578.81 Mbps, and the last will show 607.75 on the same interface, representing 5% growth per period.

# 9 Perform Capacity Planning

**Topics:**


- Capacity planning optimization
- Perform capacity planning


Outlines strategies and tools for optimizing capacity planning, including generating and interpreting reports to ensure efficient resource utilization.

## Capacity planning optimization

Explains how capacity planning optimizations enable efficient resource allocation to meet changing demands.

As demand grows on your network, you need a way to address the additional traffic and congestion. To alleviate congestion, you can

- upgrade existing circuits by adding more capacity
- augment these circuits with associated port circuits
- add parallel circuits to existing circuits, or
- specify new adjacencies between nodes that were not initially connected.

To perform these planning optimization investigations, Cisco Crosswork Planning includes a **Capacity planning optimization** tool. The goal of a capacity planning optimization is to minimize the installation of additional required network capacity. The capacity planning optimizer operates on Layer 3 network elements, as well as across failure sets, so you can use a combination of elements in a layered fashion to reach an acceptable solution to meet maximum utilization requirements.

## Perform capacity planning

Describes how to configure optimization parameters to relieve network congestion and augment capacity.

Use this task to optimize network capacity by configuring parameters that relieve congestion and augment available bandwidth. This task enables you to set utilization thresholds and select advanced optimization options for circuits and failure sets.

You can set optimization parameters to relieve network congestion and augment capacity. You start by defining a maximum utilization threshold for a specified set of interfaces. Then, you add additional options that operate on Layer 3 and across failure sets.

> **Note**
>
> Some optimizer options are mutually exclusive. For example, if configure the optimizer to create port circuits, it cannot create parallel circuits at the same time. Other options are complementary. For example, you can specify the creation of parallel circuits and new adjacencies together.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.

   From the toolbar, choose **Actions** > **Tools** > **Capacity planning optimization**. Alternatively, you can also choose **Preset workflows** > **Perform capacity planning**.

2. From the **Circuits** panel, select the circuits you want the optimizer to consider. Decide on the Layer 3 optimization options to use. The options are described in Optimization options on page 156.

   You can upgrade all existing circuits or a subset of these circuits. If you want to allow the upgrade of just certain circuits, indicate those circuits to be modified. This is useful if you want to limit your upgrades to a certain geographical location in the network.

   To create new port circuits or parallel circuits, choose one of these options:

- Create new port circuits (LAGs): In this case, the optimizer augments the existing circuits with associated port circuits (LAGs). The optimizer converts non-LAG circuits to LAGs.
- Create new parallel circuits: In this case, the optimizer creates new circuits that are parallel to existing circuits.

> **Note**
>
> - You can specify both the capacity increment and the existing capacity of LAG circuits in tandem.
> - You can apply the IGP metric to new adjacencies, but not to parallel circuits.

3. (Optional) Override the defaults for how upgraded circuits and new objects are tagged and click **Next**.
4. (Optional) In the **Optimization objective** section, choose the options to provide additional capacity planning. Describe the cost for the various elements in your design. In the **Failure sets** section, choose the failure sets, as required. For more information, refer to Advanced optimization options on page 157.
5. (Optional) Specify the maximum number of threads. By default, the optimizer tries to set this value to the optimal number of threads based on the available cores.
6. Click **Next**.
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

     The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

     > **Note**
     >
     > Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:
   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.

- If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

9. Click **Submit** to create the capacity planning optimization reports.

---

Cisco Crosswork Planning routes the traffic and looks at the utilization threshold, and any other optimization parameters you specified. If you have specified any advanced options, the optimizer takes into account whether it makes sense to upgrade existing circuits or set up new adjacencies, and considers capacity increments, cost options, and feasibility limits.

### What to do next

After running an optimization, you can look at the summary report to see what the optimizer did to remove congestion. The key metric to look at is the Total Capacity Added (Mb/s). For more information, refer to Optimization report analysis on page 160.

## Optimization options

Provides Layer 3 optimization options for capacity planning, including utilization thresholds, capacity increments, and circuit or adjacency creation preferences.

This topic describes the Layer 3 options for capacity planning. These parameters allow the optimizer to use your preferences for utilization thresholds, capacity increments, and whether to upgrade existing circuits, create new adjacencies between nodes, create port circuits, or create parallel circuits.

**Table 9: Layer 3 optimization options**

| Option | Description |
|---|---|
| **Interface utilization & capacity** | |
| Maximum interface utilization | Defines the maximum utilization threshold that you want to use for all interfaces in the network. By default, the capacity planning optimizer sets the threshold at 100% congestion. However, for future planning purposes, you might want to set the utilization to a lower number, such as 90%. |
| Capacity increment | Defines the bandwidth increment value (that is, the allowed capacity increment). You can think of this value as the capacity of the ports. The default is 100000 Mb/s. |
| Use capacity of existing LAG members | Augments capacity based on the capacity already defined in the existing LAG Members instead of using the **Capacity increment**. |
| **Circuits** | |
| Upgrade Circuits | You can upgrade all circuits or just a subset of circuits to reduce congestion. To upgrade only certain circuits, select those circuits you want to modify. This is useful if you want to limit your upgrades to a certain geographical location in the network. By default, none of the circuits are selected. |
| Create<br>• Port circuits (LAGs)<br>• Parallel circuits | • The optimizer augments the existing circuits with associated port circuits (LAGs) with additional port circuits. The optimizer converts non-LAG circuits to LAGs.<br>• The optimizer creates parallel circuits to reduce congestion. If you tell the optimizer to create parallel circuits, you cannot tell it to create port circuits and vice versa. |

| Option | Description |
|---|---|
| **Adjacencies** | |
| Restrict new adjacencies between nodes | By default, the optimizer does not create new adjacencies. If you specify a set of candidate nodes for it to use, the optimizer proposes new adjacencies. The optimizer restricts a new adjacency between the specified candidate nodes. For example, you might require that only core nodes be directly connected. In this case, specify only core nodes as your candidate nodes. |
| Maximum number of new adjacencies | Specifies the maximum number of adjacencies to create between candidate nodes. By default, this is an unbounded number. |
| Set IGP metric of new interfaces to<br>• Shortest IGP metric minus decrement<br>• Fixed metric | • Sets the IGP metric of the interfaces corresponding to new adjacencies to be the shortest path minus the value specified in the metric option. This is how the IGP metric of such interfaces is set by default. The IGP metric of new interfaces due to parallel circuit upgrades is the same as the metric of the parallel interfaces.<br>• Defines a fixed IGP metric to use when creating new interfaces. The metric of new interfaces is equal to the value specified in this field. Enter a positive integer as a value in this field. |
| Tag upgraded circuits with | Defines tags for any upgraded circuits. By default, the optimizer tags upgraded circuits with the label *CapacityOpt::Upgraded*. |
| Tag new objects with | Defines tags for any new objects the optimizer creates. By default, the optimizer tags upgraded circuits with the label *CapacityOpt::New*. |

## Advanced optimization options

Provides advanced optimization options for selecting capacity or cost objectives and defining failure scenarios in the optimizer tool.

Use the advanced optimization options to

- choose whether to optimize for capacity or cost, and
- define which failure scenarios to consider.

When you run the optimizer, the tool generates a report on the overall cost, including the cost of the added ports, and recommends the most cost-efficient solution.

**Table 10: Advanced optimization options**

| Option | Description |
|---|---|
| **Optimization objective** | |
| Minimize capacity | Minimizes the capacity that is added to your network. This is the default option. |
| Minimize cost | Select this option to run the optimization with the objective of reducing the overall cost when adding new capacity. Click  to manually enter a cost unit for capacity, L3 port costs, and the feasibility limit. Alternatively, click  to import cost values from a file. |

| Option | Description |
|---|---|
| **Failure sets** | Choose the options that you want the optimizer to consider (Circuits, Nodes, Sites, and so on). Entries appear dimmed if they are not available in your design plan. |
| **Maximum number of threads** | Specify the maximum number of threads. By default, the optimizer tries to set this value to the optimal number of threads based on the available cores. |

## Capacity planning optimization example

Illustrates the impact of capacity planning optimization on network design and interface utilization, using before-and-after scenarios and explaining optimization parameters, adjacency adjustments, and resulting improvements in network performance.

This topic explains how network optimization changes network design and interface utilization, using before-and-after scenarios to illustrate the impact of optimization parameters

Figure 50: Design before optimization on page 158 shows an example of a network design before optimization, where extra demands were placed on the network. As a result, congested interfaces appear in red.

**Figure 50: Design before optimization**

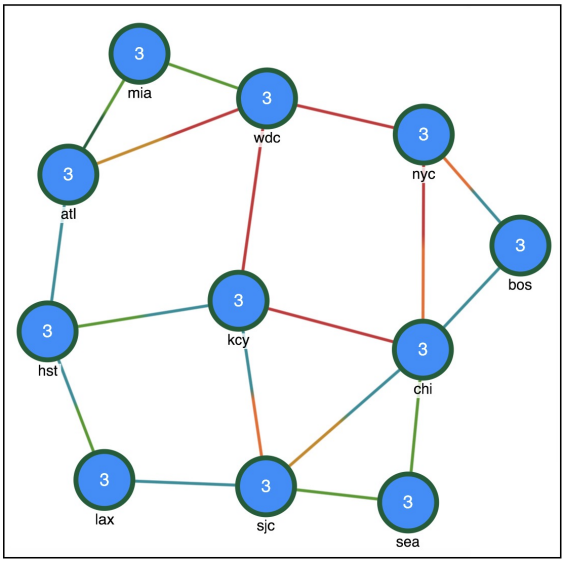



Figure 51: Design after optimization on page 159 shows the design after optimization using these parameters. For the remaining options, the default values are used.

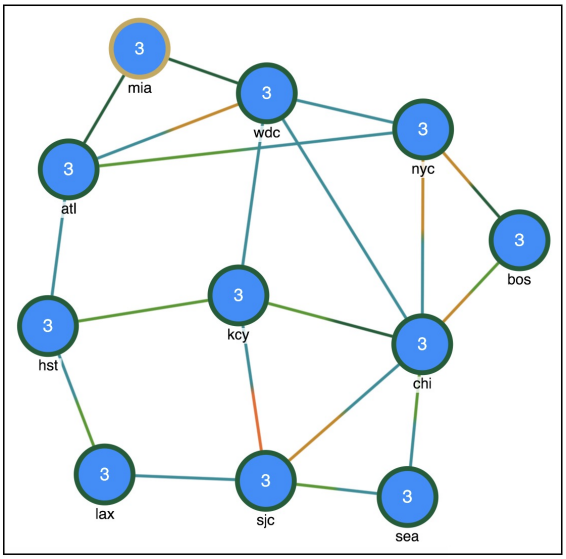
- Maximum interface utilization: 100%
- Capacity increment: 2488 Mbps
- Circuits: Upgrade All circuits and Create Port circuits (LAGs)
- Create new adjacencies: Restrict new adjacencies between all nodes
- Failure sets: Circuits

In this case, the optimizer proposes to set up two new adjacencies:

- One between Atlanta (atl) and New York City (nyc)

- One between Washington, D.C. (wdc) and Chicago (chi)

**Figure 51: Design after optimization**




## Parallel circuits design examples

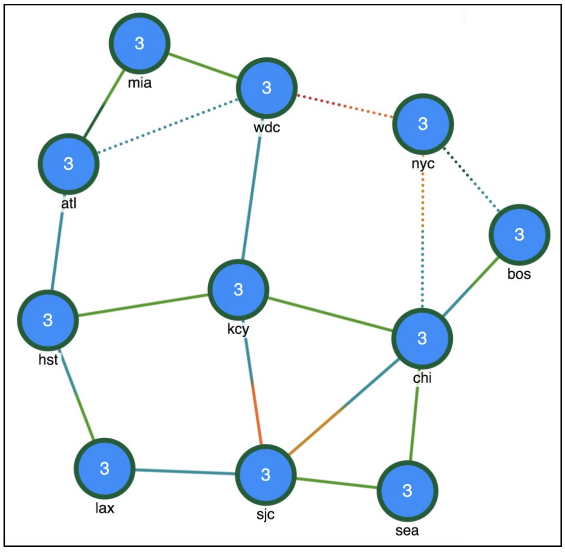
Explains how parallel circuits are added to a network design to increase capacity, using a practical example with multiple city connections.

Figure 52: Adding parallel circuits to increase capacity on page 159 shows the first design in this chapter with parallel circuits added to increase capacity.

In this case, the optimizer proposes four sets of parallel circuits as indicated by the dotted lines:

- One between Atlanta (atl) and Washington, D.C. (wdc)
- One between Washington, D.C. (wdc) and New York City (nyc)
- One between New York City (nyc) and Chicago (chi)
- One between New York City (nyc) and Boston (bos)

**Figure 52: Adding parallel circuits to increase capacity**




## Optimization report analysis

Provides information about accessing and interpreting optimization reports generated by the Capacity planning optimization tool, including available metrics and report navigation.

Each time the Capacity planning optimization tool is run, a report is automatically generated. You can access this information at any time by choosing **Actions** > **Reports** > **Generated reports** and then clicking the **Capacity Planning Optimization** link in the right panel. Note that new reports replace the previous ones.

### Summary tab

The **Summary** tab provides useful metrics at a glance. It provides these details:

- Input parameters used for optimization
- Total capacity added as a result of the optimization
- Number of new adjacencies
- Number of upgraded circuits
- Number of new circuits
- Number of new ports
- Number of new port circuits

### Capacity Upgrades tab

The **Capacity Upgrades** tab provides details on how the optimizer upgraded circuits.

# 10 Simulate IGP Routing Protocol

**Topics:**


- IGP process ID
- OSPF and IS-IS differences
- EIGRP
- Simulate IGP multipath
- Exclude ABR nodes
- Configure IGP simulation
- Multi-topology routing


Describes IGP routing protocol simulation, covering IGP process ID management, OSPF and IS-IS simulation techniques, EIGRP simulation, multipath scenarios, simulation configuration, and multi-topology routing with task-based configuration procedures.

# IGP process ID

Describes an identifier used to differentiate between multiple instances of the same IGP running on a router.

An IGP process ID is an identifier that

- differentiates between multiple instances of the same IGP running on a router, and
- allows you to have separate IGP configurations for different purposes within the same router.

### IGP process ID configuration

In Cisco Crosswork Planning, you can specify IGP process ID to an IGP protocol. Also, you can associate each interface with an IGP process ID.

## Configure an IGP process ID

Describes how to create, delete, or edit an IGP process ID in your network design.

Use this task to configure IGP process IDs to manage interior gateway protocol processes for network routing configurations.

Complete these steps to create, edit, or delete an IGP process ID.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Edit** > **IGP process protocols**.

   The IGP process protocols page opens.
3. Configure an IGP process ID.
   a) To create an IGP process ID, click , enter the IGP process ID, and select the required IGP protocol. Click **Save** to save the IGP process ID.
   b) To edit an IGP process ID, click the **Edit** button of the IGP process ID that you want to edit. Update the values and then click **Save**.
   c) To delete an IGP process ID, select the IGP process ID you want to delete and then click . Alternatively, you can use the **Delete** button.

The IGP process ID is successfully created, edited, or deleted based on your selection.

## Associate interfaces with an IGP process ID

Describes how to configure interfaces to use a specific IGP process ID for routing protocol.

Use this task to associate each interface with an IGP process ID.

### Procedure

1. From the **Interfaces** table, select one or more interfaces for which you want to add the IGP process ID.
2. Click  and then click the **Advanced** tab.

> **Note**
>
> If you are editing a single interface, you can also use the ··· > **Edit** option under the **Actions** column.

3. In the **IGP** panel, enter the value in the **IGP process ID** field.
4. Save the changes.

The selected interfaces are now associated with the specified IGP process ID.

## OSPF and IS-IS differences

Describes the differences between OSPF and IS-IS routing protocol simulations.

OSPF and IS-IS simulations are identical with these exceptions:

- OSPF routing uses OSPF areas, if specified. By default, all interfaces are assigned area zero.
- IS-IS routing uses IS-IS levels, if specified. By default, all interfaces are set to Level 2. Interfaces can belong to Level 1, Level 2, or both. If both, then you can specify an alternate metric for Level 1.

### OSPF area simulation

Describes how OSPF area membership can be specified per interface and how areas are configured for routing simulation.

OSPF area membership can be specified per interface. The two interfaces on each circuit must belong to the same area. Area names can be any string. Area zero, the backbone area, must be denoted by ‘‘0’’, ‘‘0.0.0.0’’, or an empty string. Cisco Crosswork Planning simulates the OSPF area routing configuration in which the areas import Link State Advertisements (LSAs) from the backbone. A demand from a source node to a destination node in a different area will only be routed if it can reach the destination by passing through the source area, directly to area zero, and from there directly to the destination area.

By default, all nodes in a single AS are assumed to belong to a single OSPF area. This list describes how the nodes are assigned to areas:

- If OSPF areas are not defined for interfaces, all nodes are assumed to be in the same area.
- Each interface can be assigned to only one OSPF area. Note that each node can be assigned to one or more areas.
- If a node has an interface in an OSPF area, the node is assigned to that area.
- An Area Border Router (ABR) is a node that belongs to both area 0 and other OSPF areas.

#### Set OSPF area membership

Outline how to configure OSPF area membership for network interfaces.

Use this task to assign OSPF area membership to specific network interfaces.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select one or more interfaces from the **Interfaces** table.

3. Click .

> **Note**
>
> If you are editing a single interface, you can also use the ••• > **Edit** option under the **Actions** column.

4. Click the **Advanced** tab.
5. In the **IGP** panel, enter the required value in the **OSPF area** field.
6. Save the changes.

---

The OSPF area membership is configured for the selected interfaces.

## IS-IS multi-level simulation

Describes how IS-IS interfaces are assigned to different levels and how nodes are organized into areas based on their level assignments.

By default, IS-IS interfaces are assigned to Level 2, though you can assign them to Level 1 or to both Levels 1 and 2. For instructions, refer to Set IS-IS multi-level simulation on page 164.

- The **IGP metric** defines the Level 1 metric for interfaces in Level 1, in Level 2, or in both Level 1 and 2 if these have equal metrics. However, you can change the Level 1 metric.
- If an interface is in both Level 1 and Level 2 with uneven metrics, the IGP metric defines the Level 2 metric and the Level 1 metric defines the Level 1 metric.

The IS-IS level is listed in the **ISIS level** column of the **Interfaces** table, and the Level 1 metric is listed in the **Metric ISIS level 1** column.

### Node assignment rules

By default, all nodes in a single AS are assumed to belong to a single IS-IS level. This list describes how the nodes are assigned to levels:

- Each node can be assigned to one or more levels.
- Two Level 1 nodes are placed in different areas if any route between them (under normal operation) passes through a Level 2 node.
- A node is assigned to a single Level 2 area if there is at least one interface that is Level 2, or both Level 1 and Level 2.
- A node is assigned to one of potentially multiple Level 1 areas if any interface is Level 1, or both Level 1 and Level 2.
- An ABR is a node belonging to both the Level 2 area and another IS-IS area.

### Set IS-IS multi-level simulation

Describes how to configure IGP metrics, IS-IS levels, and Level 1 metrics for network interfaces.

Use this task to configure IGP metrics, IS-IS levels, and Level 1 metrics for network interfaces in your plan file.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select one or more interfaces from the **Interfaces** table.
3. Click .

   **Note**

   If you are editing a single interface, you can also use the ··· > **Edit** option under the **Actions** column.

4. To set the IGP metric, enter its value in the **IGP metric** field.
5. To set IS-IS multi-level simulation, click the **Advanced** tab.

   In the **IGP** panel:

   a) Enter the **IGP process ID**.
   b) Select the level from the **IS-IS level** drop-down list.
   c) In the **Level 1 metric** field, enter the metric value.
6. Save the changes.

The IS-IS multi-level simulation is configured with the specified IGP metrics, IS-IS levels, and Level 1 metrics for the selected interfaces.

## EIGRP

Describes the Enhanced Interior Gateway Routing Protocol and its metric calculation method.

EIGRP is a routing protocol that

- uses a specific formula to derive total feasible distance from a node to a destination subnet, and
- does not use the IGP Metric value for each interface.

### EIGRP metric calculation

EIGRP uses this formula to derive total feasible distance from a node to a destination subnet:

```
path metric to destination = (10,000/(bandwidth) + (delay) * 256
```

The formula components are:

- The *bandwidth* is the minimum interface Capacity value along the path. This is in Mbps. For each interface, if there is no Capacity value, the Capacity Sim value is used instead.
- The *delay* is the sum of the interface delays in 10s of microseconds. This is calculated by taking the sum of the EIGRP Delay values and dividing this sum by 10.

You can set this delay in the **EIGRP delay** field of an interface's Edit window. Here, the value is in microseconds.

If EIGRP delay is not set, Cisco Crosswork Planning uses 10 as the delay in the preceding calculation, which means 10 " 10s of microseconds.". For example, if there are eight interfaces, each with an EIGRP Delay value of 15, the delay used in the calculation is (8 x 15) / 10 = 120 / 10 = 12.

Demands show the EIGRP path metric of the path over which they flow in the **Path metric** column.

## Simulate IGP multipath

Explains how to configure the maximum number of ECMP paths for IGP multipath simulation in network planning.

By default, demands are split equally between all paths from a route of equal distance to the destination, and there is no ECMP limit.

You can specify a maximum number of ECMP paths. In this case, demands transiting through a router are distributed among the available paths up to this maximum ECMP value. Paths are chosen by lowest next-hop IP address, which is the interface IP address for IGP or the destination IP address for LSPs. Paths with no IP address are chosen last.

### Before you begin

Complete these steps to configure ECMP.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

   The Network options page opens.
3. Click the **Protocols** tab.
4. In the **IGP** section, enter the maximum number of ECMPs in the **Max ECMP paths** field. A 0 (zero) means there is no limit.
5. Save the changes.

The ECMP configuration is saved and will be applied to the network model for IGP multipath simulation.

## Exclude ABR nodes

Explains how to configure ABR node exclusions to prevent specific nodes from being used while entering or exiting the domain.

Use this task when you need to exclude ABR nodes while entering or exiting the domain to optimize network performance or maintain routing preferences.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Edit** > **Node ABR exclusions**.

   The Node ABR exclusions page opens. The **Node** column displays the A and Z end. The **Exclude node** column displays the node that is excluded.
3. To add an IGP ABR node exclusion:
   a) Click .
   b) In the **Node** and **Exclude node** sections, specify the details of the node and the node that needs to be excluded, respectively.
   c) Save the changes.

4. To delete an IGP ABR node exclusion:
   a) Select the row that you want to delete.
   b) Click .

The IGP ABR node exclusions are configured according to your specifications.

## Configure IGP simulation

Describes how to configure IGP simulation settings to control route redistribution across IGP processes and manage ECMP path tie breakers.

IGP simulation configuration allows you to define how IGP processes handle route redistribution and path selection.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, click  or choose **Actions > Edit > Network options**.
   The Network options page opens.
3. Click the **Protocols** tab.
4. In the **IGP** section, choose the appropriate protocol from the **Protocol** drop-down list.
5. Click the **Simulation** tab.
6. Select an option for **Redistribute routes across IGP process**.
   The options include:
   - No IGP redistribution
   - Shortest exit
   - Shortest path
7. The **Multiple IGP ABRs** option is disabled by default. This option is helpful if you want to apply tie breakers for ECMP paths based on BGP ID, IP address, and host name respectively.
8. Save the changes.

The IGP simulation configuration is saved and applied to the network plan. The selected route redistribution method and ECMP tie breaker settings will be used for IGP simulation calculations.

## Multi-topology routing

Describes how Cisco Crosswork Planning simulates multi-topology routing with specific rules for interface assignment and routing restrictions.

Cisco Crosswork Planning supports simulation of multi-topology routing. Interfaces can be assigned to one or more IGPs. Demands and LSPs can be assigned to a specific IGP, and will only route through interfaces belonging to that IGP. This multi-topology simulation uses these rules.

- When an interface is common to more than one topology, all IGP properties of the interface (including the metric) must be the same in both topologies.

- The IGPs must be defined to be of the same type, for example, OSPF or IS-IS. This is not a restriction in practice unless either the OSPF topology or the IS-IS topology uses multiple areas.
- Demands and LSPs defined with a specific topology can only route over circuits belonging to that topology. A circuit belongs to a topology if either one of its interfaces does.
- Demands and LSPs with no defined topology have no restrictions on routing. They are routed using the default topology, to which all interfaces belong.

## Configure topologies

Explains how to create or delete network topology configurations in the plan file.

Complete these steps to create or delete topologies.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions > Edit > Topologies**. The Topologies page opens.
3. To create a new topology:
   a) Click .
   b) Enter the topology name.
   c) Click the **Save** button.
4. To delete a topology:
   a) Select the row containing the topology you want to delete.
   b) Click the **Delete** button in the selected row or click .

## Add topologies to objects

Describes how to configure associations between topologies and demands or LSPs to organize network elements for planning and visualization.

Use this task to associate demands or LSPs with topologies. Once you have associated the topologies, you can view them by showing the **Topology** column in the respective tables (use the **Show/hide table columns** icon ().

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select one or more demands or LSPs from their respective tables.
3. Click .
4. Select the desired topology or topologies.
   a) From the **Topology** drop-down list, select the desired topology. You can associate one topology to each demand or LSP.

      In case of LSPs, this drop-down is available under the **Advanced** tab.

   b) Save the changes.

By selecting from a similar **Topology** drop-down list, you can also add a topology to LSP meshes when creating them.

---

The selected demands or LSPs are now associated with the specified topologies and will appear in the Topology column when enabled in the respective tables.

# 11 Simulate BGP Routing

**Topics:**


- Multi AS network modeling for BGP routing
- Internal and external AS types
- Configure ASes
- Route demands between ASes
- Know about BGP routing details
- How BGP routing works


Describes BGP routing simulation concepts, including AS types and configuration, route determination processes, BGP protocol features, and step-by-step procedures for modeling routes between internal and external ASes.

# Multi AS network modeling for BGP routing

Explains how Cisco Crosswork Planning models multi-AS networks and simulates basic BGP routing using high-level peering policies.

Cisco Crosswork Planning models multi-AS networks and simulates basic BGP routing. The system does not directly emulate BGP routing configurations, such as local prefs and MEDs. Rather, it provides a high-level modeling of typical peering policies, such as standard customer, transit, and settlement-free arrangements for service providers. This model lets you evaluate the effects of peering locations and basic policy variations.

You can also extend these high-level models to significantly more complex policy-based routing situations using external endpoints as demand sources and destinations. For information on demands and on external endpoints, refer to Simulate Traffic Flow from Source to Destination Using Demands on page 95 and Simulate Advanced Routing with External Endpoints on page 209.

# Internal and external AS types

Describes the two AS types used in multi-AS network modeling.

To model a multi-AS network, each node is assigned an AS, and each AS is defined as either *internal* or *external*.

### AS type characteristics

A typical multi-AS model in Cisco Crosswork Planning consists of these components:

- A single internal AS representing the full topology of your network.
- Individual peering nodes of neighboring external ASes.
- Peering circuits connecting the internal AS to the nodes in the external ASes.

Generally, there are many external ASes in the network model, but usually only one or a few internal ASes. All nodes in an external AS are typically placed in the same site, although you can place them in any site.

ASNs and their types are defined in the Edit AS page and listed in the **AS** table. Nodes are assigned to ASes in the Edit Node page.

# Configure ASes

Describes how to configure ASes in your network topology to establish routing domains and enable proper BGP communication between different administrative boundaries.

## Create ASes

Explains how to configure ASes in your network design plan by creating empty ASes and defining their properties.

Use this task to create an empty AS.

After creating the AS, you still need to associate nodes with it and create the relationship between this AS and others. Refer to Associate nodes with an AS on page 173 and Edit AS routing policies on page 174.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **AS**.

Alternatively, in the Network Summary panel on the right side, click  in the **AS** tab.

The AS tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon () and check the **AS** check box.

3. Configure the AS properties:
   - **ASN**: AS number, which is a text string that can be a number or name.
   - **Name**: AS name.
   - **Type**: Internal ASes have a full topology. External ASes have a collapsed topology with just border nodes and a virtual node.
   - **External mesh**: When creating a demand mesh, this option tells Cisco Crosswork Planning whether to create external meshes. When one or both ASes are set to Include, Cisco Crosswork Planning creates a mesh between the external ASes (default). If both are set to Exclude, no demands are created.
   - **IGP protocol**: Choose OSPF, ISIS, or EIGRP from the drop-down.
   - **Description**: A text description of the AS.
4. Click **Submit** to create an AS.
5. To change the routing policy, select the AS and click  and then click the **AS Relationships** tab (refer to Edit AS routing policies on page 174).

---

The AS is created with the specified properties.

## Associate nodes with an AS

Describes how to associate network nodes with an AS.

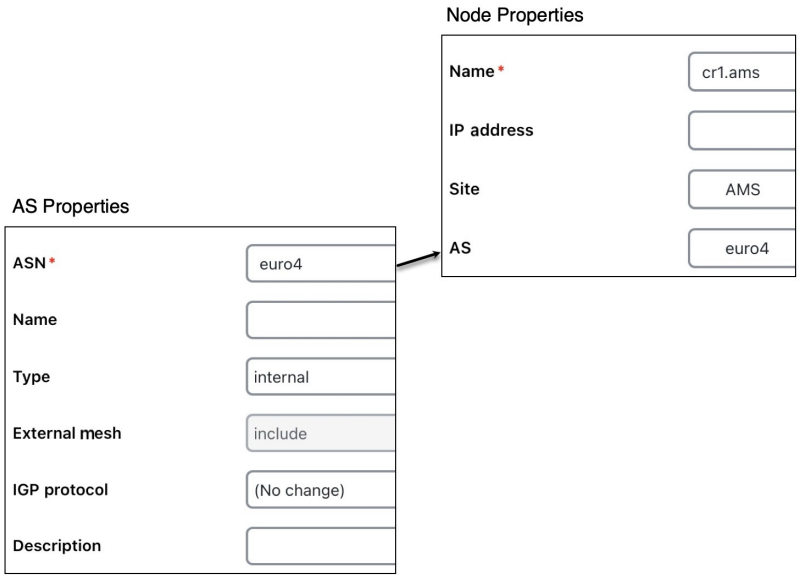
Use this task to associate nodes with an AS to establish proper network topology and routing relationships within your network design.

### Procedure

---

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, choose one or more nodes from the **Nodes** table.
3. Click .

> **Note**
>
> If you are editing a single node, you can also use the ••• > **Edit** option under the **Actions** column.




4. From the **AS** drop-down list, choose the AS to which you want to assign the nodes.
5. Save the changes.

---

The selected nodes are now associated with the specified AS.

## Edit AS routing policies

Explains how to configure AS routing policies to create AS relationships by setting the appropriate routing policy for each AS pair.

Use this task to create AS relationships by setting the routing policy for ASes.

### Procedure

---

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select an AS from the **AS** table.
3. Click or use the ••• > **Edit** option under the **Actions** column.
4. Click the **AS Relationships** tab.
5. Select the AS pair you want to configure. There is a separate line for each direction in the relationship so you can configure them independently.
6. Click , and set the **Routing policy** to **Respect MEDs** or **Shortest Exit**. For more information, refer to Route demands between ASes on page 175.
7. Submit your changes.

---

The AS routing policy is configured for the selected AS pair, establishing the routing relationship according to your specified policy settings.

# Route demands between ASes

Describes the traffic engineering requirements and routing policies that define how data flows should be distributed between different Autonomous Systems to meet performance, capacity, and business objectives.

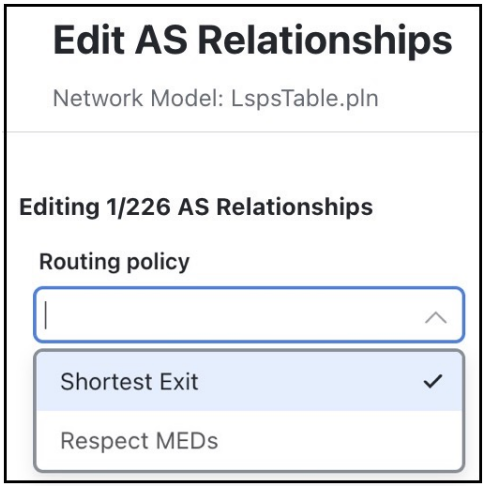
## Route selection between internal ASes

Describes how routing is determined between two connecting internal autonomous systems using IGP protocols and routing policies.

Demands routed within a single AS have a specified source node and destination node where traffic originates and terminates. Demands routed between two connecting internal ASes are specified in the same manner: with a source node in the first AS and a destination node in the second.

### Routing policy configuration

Cisco Crosswork Planning routes within an AS, to and from the border exit point, are determined by the IGP protocols. The selection of border exit point is modeled by the **Routing policy**, which is set to either **Shortest Exit** or **Respect MEDs**. This property is set in the Edit AS Relationships page, which is accessed through the Edit AS page.




The routing policy options include:

- **Shortest Exit**: The border exit node is selected, which is closest to the source node, within the IGP of the source AS. If there is a tie, the exit node with the lowest BGP ID is used.
- **Respect MEDs**: The border exit node is selected, which is closest to the destination node, within the IGP of the destination AS. If there is a tie, the exit node with the lowest BGP ID is chosen.

## Routes between external and internal ASes

Describes how routing policies determine traffic flow control between different autonomous systems in peer, customer, and transit relationships.

This topic lists typical routing configurations that can be constructed by applying different combinations of routing policies for traffic in both directions between two ASes.

### Typical AS routing configurations

Different AS relationship types use specific routing policy combinations:

- In a peer relationship, routing in both directions is Shortest Exit, which means each controls its own border exit points.
- For a customer relationship, the customer determines the border exit points for traffic in both directions.
- For a transit relationship, the transit AS provides paid transit to the internal AS, so the internal AS determines all border exit points.

**Table 11: Typical AS routing configurations**

| Type | Policy to | Policy from |
|---|---|---|
| Peer | Shortest Exit | Shortest Exit |
| Customer | Respect MEDs | Shortest Exit |
| Transit | Shortest Exit | Respect MEDs |

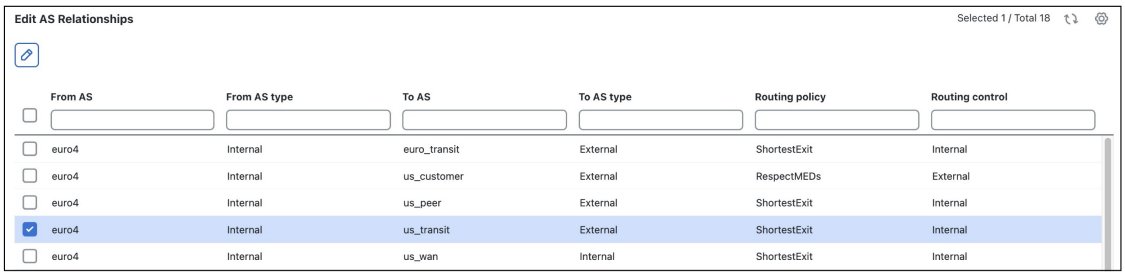
Like traffic routed within an AS, traffic routed between ASes is represented by demands. However, for demands from and/or to external ASes, the external AS is defined AS the source or destination of the demand. Optionally, the specific node in the external AS from which the traffic enters or exits the internal AS is also specified.

Failover between nodes in the external AS can be modeled. For example, if the traffic is sourced from an external AS and if the peering circuit from which traffic is entering the internal AS under normal operation fails, the traffic can enter the internal AS from a different interface or peering node in the same external AS. In the Demands table, the sources and destinations are represented as:

**AS{<ASN>}:if{node_name|interface_name}**

Example: AS{33287}:if{cr01.newyork.ny|POS3/7/0/0}

The AS that controls the routing chooses which peering node to use. If the internal AS controls the routing, then because the topology of the internal AS is known, you can simulate the routing to the peering node. However, because Cisco Crosswork Planning has limited knowledge of the external AS topology, if the external AS controls the routing, you cannot predict how traffic will be distributed among the exit points.




### Determine the AS that controls the routing

The AS that controls the routing is determined by the AS type, direction of the demand, and the Routing policy property AS described in this table.

**Table 12: Determining the AS that controls the routing**

| Direction | Routing Policy | AS with Routing Control |
|---|---|---|
| External AS to Internal AS (Ingress) | Respect MEDs | Internal |
| | Shortest Exit | External |
| Internal AS to External AS (Egress) | Respect MEDs | External |
| | Shortest Exit | Internal |

### Effects of routing policy and routing control

Two ASes can be in one of four different routing relationships to one another, depending on which of the two routing policies is chosen in each direction. Traffic routing behavior depends on which AS has control:

- If traffic is routed to an external AS when it has control and there is no knowledge of its topology, a set of demands is created from the source in the internal AS (or from another external AS), each with a destination set to one of the border nodes in the external AS. This way, any division of traffic between the exit points can be modeled.
- If traffic is routed to an external AS when an internal AS has control, a single demand is created from the source to the AS itself. Cisco Crosswork Planning simulations determine the correct exit point for this single demand based on the source.
- If traffic is to be routed from an external AS when it has control, a demand is created from each node in the external AS to each node in the internal AS.
- If traffic is to be routed from an external AS when an internal AS has control, a demand is created to each node in the internal AS using the external AS AS the source. The demand originates from one or multiple nodes in the external AS, depending on the topology and the metric cost to reach the destination node. For example, a single demand from an external AS to a specific node could be sourced from two different nodes in the external AS, each carrying 50% of the demand traffic.

**Table 13: Effects of routing policy and routing control**

| Direction | Routing Policy | AS with Routing Control | Demand Source or Destination Endpoint in Remote AS | Number of Demands |
|---|---|---|---|---|
| External AS to Internal AS (Ingress) | Respect MEDs | Internal | Entire external AS | One only |
| | Shortest Exit | External | Border nodes | One for each node |
| Internal AS to External AS (Egress) | Respect MEDs | External | Border nodes | One for each node |
| | Shortest Exit | Internal | Entire external AS | One only |

## External meshes

Describes how external meshes control BGP route advertisement restrictions between external autonomous systems in Cisco Crosswork Planning.

An external mesh is a network configuration that

- consists of two or more external ASes with a **Type** property of **external**, and
- enables or restricts the advertisement of BGP routes between external ASes.

### External mesh property configuration

An internal AS typically restricts advertisement of BGP routes for some external ASes to other external ASes. For example, destinations reachable through the transit network would not be advertised to a peer, or vice versa. In Cisco Crosswork Planning, these restrictions are represented by the absence of demands between the two external ASes.

Each AS has a property called **External mesh**, which Cisco Crosswork Planning uses when inserting demand meshes into a plan. Demands are created for external ASes only if one or both ASes have **External mesh** set to **include**. If both ASes are set to **exclude**, no demands are created for the external AS. For example, in Figure 53: External mesh control on page 178 the peer and transit ASes are both set to Exclude, so no demands are created between those ASes. All other external AS demands are included in the demand mesh. Table 14: External mesh settings for common AS relationships on page 178 shows the External Mesh settings for common AS relationships.

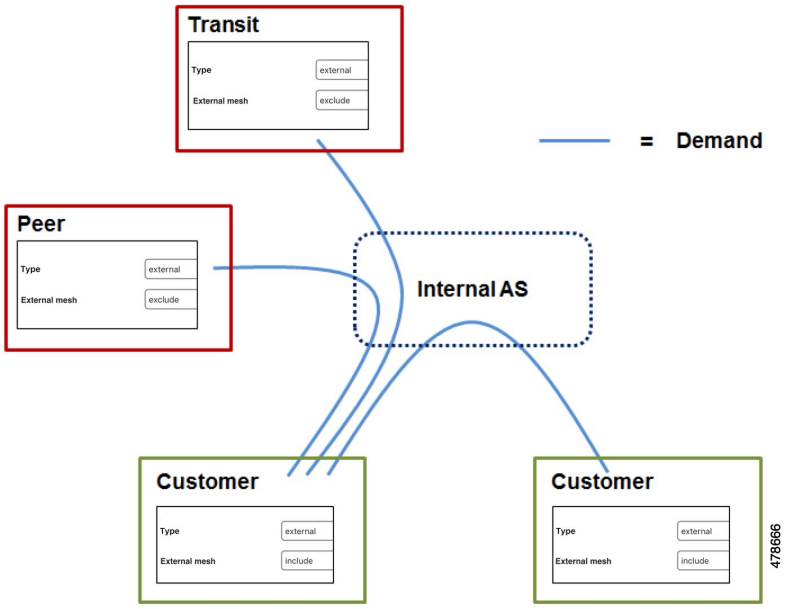
This table shows the External Mesh settings for common AS relationships.

**Table 14: External mesh settings for common AS relationships**

| Relationship | External Mesh Setting | Result |
|---|---|---|
| Peer | Exclude | Demands created to/from customers only |
| Customer | Include | Demands created to/from all external ASes |
| Transit | Exclude | Demands created to/from customers only |

For internal ASes, the **External mesh** property is ignored. More complex route advertisement policies cannot be represented by these simple External mesh settings. In this case, demand mesh creation must be performed in several steps, possibly using a script.

**Figure 53: External mesh control**




# Know about BGP routing details

Describes the BGP routing mechanisms, including path selection algorithms, route advertisements, and inter-autonomous system communication for Internet-scale network routing decisions.

## BGP multihop

Explains the concept of BGP multihops and how pseudonode modeling helps represent connections across autonomous systems.

Cisco Crosswork Planning automatically constructs BGP pseudonodes where necessary when BGP multihops are detected.

Cisco Crosswork Planning models the nodes in external ASes that are directly connected, for example, through eBGP, to nodes in internal ASes. One exception is that you can model BGP multihops by setting the node **Type** property to **psn** (pseudonode), such as might occur at a peering exchange. This pseudonode can represent the switch that connects a number of external AS nodes to the same internal AS node. In this instance, multiple external AS nodes are connected by circuits to a BGP psn node, and this node is connected to a node in the internal AS.

> **Note**
>
> In all cases, eBGP multipaths across parallel border circuits is assumed.

## Configure BGP load balancing

Explains how to configure BGP multipath options globally to enable load balancing to external ASes using EBGP multipaths or EBGP multihops.

Use this task to set BGP multipath options globally.

BGP load balancing to an external AS uses eBGP multipaths or eBGP multihops. Cisco Crosswork Planning models these two eBGP load balancing designs in the same manner, though in the UI they are identified only as multipaths. BGP multipath options are disabled by default.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

   The Network options page opens.
3. Click the **Protocols** tab.
4. In the **BGP** section, for each BGP multipath option that you want enabled, choose **Enabled** from the drop-down list. By default, all these options are disabled.
   - EBGP multipath: Turns on eBGP multipath within the internal ASes. Demand routings through the internal AS to an external AS are divided among external routes with equal-cost BGP exit routes.
   - EBGP multipath incoming: Turns on eBGP multipath in all external ASes. Demand routings from external ASes to an internal AS are divided among external routes with equal-cost BGP exit routes.
   - IBGP multipath: Turns on iBGP multipath within the internal ASes. Demand routes through an internal AS to an external AS are divided among internal paths to equal-cost BGP exit routes.
5. Save the changes.

BGP multipath options are now enabled globally, allowing for load balancing across multiple equal-cost BGP paths.

### BGP next hop

Describes the BGP next hop configurations and how they affect IGP metric calculations in path selection.

In networks, there are two common configurations for the BGP next-hop IGP metric used in the path selection. One is to set the next-hop self on the iBGP peers (next-hop self = on). The other is to configure IGP metrics on eBGP interfaces, and to inject the interface prefix into the IGP database by setting the interface to be a passive IGP interface (next-hop self = off).

Cisco Crosswork Planning does not have an explicit next-hop self setting, so it simulates paths as if next-hop self is off. That is, the IGP metric of the egress peering interface is included in the IGP distance to the peering router and is used in the iBGP path selection. However, next-hop self to an external AS can effectively be simulated by setting the metrics on all egress interfaces to that external AS to 0. You can set the IGP metric in either the Edit Interface or Edit Circuit page.

## How BGP routing works

Describes how BGP routing operates in AS routing simulations using demands and specified protocols for particular failure scenarios and traffic levels.

As with all Cisco Crosswork Planning simulations, AS routing uses demands. An IP simulation for a particular failure scenario and traffic level performs these steps.

#### Summary

The key components involved in the BGP routing process are:

- Demands and LSPs
- BGP protocols

#### Workflow

These stages describe how BGP routing operates:

1. Demands are routed using the established LSPs (if applicable) and using the specified BGP protocols given the specified failure scenarios.
2. Interface utilizations are calculated from the demand traffic using the specified traffic level.

   Cisco Crosswork Planning allows routes to be calculated between selected nodes even if no demands are present. In this case, only the first step applies.

   BGP demands do not failover between external ASes. That is, all traffic to or from an external AS behaves the same under peering failures to an external AS. You can change this default behavior using external endpoints to simulate specific external AS nodes where traffic goes in and out of the network, as well as set priorities so that if one traffic source or destination goes down, the traffic can still be sourced from or delivered to another external AS node.

# 12 Simulate Quality of Service (QoS)

**Topics:**


- Quality of Service (QoS)
- QoS requirements
- QoS bound and QoS violations
- Queue and service class information
- QoS bound and QoS violation columns
- Create queues
- Create service classes
- Create policy groups for interfaces
- Create or edit service class policies
- Apply QoS settings globally


Explains Quality of Service simulation concepts, requirements, bound calculations, queue and service class management, policy configuration, global settings, and violation indicators to optimize network performance and resource allocation.

# Quality of Service (QoS)

Explains how QoS features enable networks to provide preferential treatment to critical traffic and ensure service levels without over-provisioning.

Quality of Service (QoS) is a network performance mechanism that

- ensures high-quality performance for critical applications
- allows preferential treatment for critical applications
- differentiates traffic based on user or application requirements, and
- ensures that service levels are met without reactively expanding or over-provisioning the network.

QoS features are available for undifferentiated traffic, for service classes, and for interface queues.

- Undifferentiated traffic: Aggregate traffic on an interface.
- Service class: A user-defined classification of traffic that is not discovered by Cisco Crosswork Planning. Examples include voice, video, and data. Service classes apply to the entire network.
- Interface Queue: In live networks, traffic waits in conceptual lines (queues) and then is forwarded over an interface on a per-queue basis according to QoS parameters. Similarly, in Cisco Crosswork Planning, each queue has a set of user-defined QoS parameters (interface queue properties) that specify how these queues are prioritized and what percentage of traffic they carry. An interface contains zero or more queues that are discoverable by Cisco Crosswork Planning. You can also manually create and configure them. The traffic per queue is also discovered.

# QoS requirements

Outlines QoS requirements, including the definition and implementation of traffic policies for service classes and the configuration of interface queue properties such as priority, weight, and police limit.

In Cisco Crosswork Planning, QoS requirements are defined by policies and interface queue properties.

### Policy

A policy defines the maximum percentage of traffic capacity that can be used by either a service class or by undifferentiated traffic. There are two policies: one for normal operation and one for worst-case scenarios. Policies set on service classes do not affect QoS requirements of any other service class, nor do they affect live network behavior.

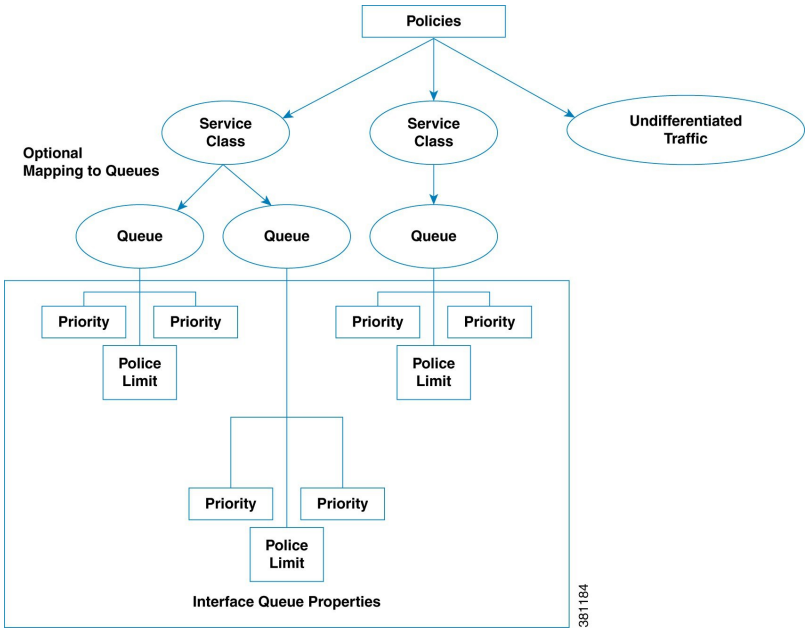
### Interface queue properties

These parameters affect routing behavior in a live network. In Cisco Crosswork Planning, the interface queue properties are priority, weight, and police limit. To set these properties, refer to Edit interface queue properties on page 192.

- **Priority**: Identifies the precedence of the queue. For example, traffic in a priority 1 queue is routed before traffic in a priority 2 queue. Queues with the same priority share capacity evenly based on weighted round robin (WRR) calculations. You can change this behavior using the weight and police limit parameters. There are an unlimited number of priorities; however, most networks only use no more than three. By default, queues do not have priorities.
- **Weight**: Identifies the percentage of preference given to queues of an equal priority level, which enables the network to fairly distribute the load among available resources. For example, if 10 Gbps were passing through a 10GbE interface on two priority 1 queues, by default 5 Gbps would pass through each queue. However, if you set the weight of one queue to 75% and the other to 25%, the distribution would be 7 Gbps and 2.5 Gbps, respectively. By default, all queues have a weight of 100%.

- **Police limit**: Identifies the maximum percentage of available capacity permitted through a queue of a given priority level, thereby preventing traffic from higher priority queues from starving lower priority queues. For example, if the interface is a 20GbE and a priority 1 queue has a police limit of 40%, then only 8 Gbps of interface traffic can go through this queue. By default, all queues have a police limit of 100%. To see examples of this *starvation*, refer to the examples in Policies and QoS bound calculations on page 186, where you can see that lower priority queues received zero traffic due to priority settings.

**Figure 54: Policies and interface queue parameters**




## QoS bound and QoS violations

Introduces the concepts of QoS bound and QoS violations, explaining how policies and interface queue properties define capacity limits, how violations are detected, and how these calculations facilitate network service planning and monitoring in the Network Summary table.

Cisco Crosswork Planning uses the concepts of **QoS bound** and **QoS violation** to identify whether QoS parameters are being met or surpassed. This helps you plan for service requirements across the network. Policies and queue properties determine the QoS bound calculation. In turn, this calculation determines whether there is a violation.

### QoS bound

The QoS bound is the maximum interface capacity available without violating these QoS requirements. Separate QoS bounds are calculated for both policy and interface queue properties.

| QoS bound for... | Calculated based on... |
|---|---|
| Interface queues | Combination of interface queue properties, or in live networks, it is the capacity percentage that is discovered. |
| Service class | Policy |
| Service class mapped to queues | The lower of these two calculations is used:<br>• Policy for service class<br>• Queue properties for queues |

| QoS bound for... | Calculated based on... |
|---|---|
| Undifferentiated traffic | Policy |

In the Network Summary table, these columns convey the QoS bound information: QoS bound meas, QoS bound meas (%), QoS bound sim, and QoS bound sim (%).

QoS bound calculations involve determining how to increase traffic on the queues until it reaches the maximum allowable amount. This capacity, or the reason the traffic cannot be increased further, is defined both by the QoS parameters and the amount of traffic. For example, when traffic arrives at Queue X, Cisco Crosswork Planning fixes the traffic on all other queues and then determines how it can increase the traffic on Queue X until some other traffic blocks it.

For those queues that do not reach full capacity, unused queue capacity is made available for other queues.

#### QoS violation

QoS violation is the total traffic minus the capacity permitted for the queue (QoS bound). A violation occurs if the maximum QoS capacity allotted through policies and interface queue properties is exceeded. If the number appearing in the **QoS violation** column is positive, the allotted capacity has been surpassed. If the number is negative, the allotted capacity has not been reached.

## Policies and QoS bound calculations

Lists example policy configurations, their resulting QoS bounds, and potential violations for network interfaces.

If you do not set other QoS parameters such as Priority, Weight, or Police limit in the Interface queue properties, the QoS bound equals the policy that you configure.

This table provides examples of how interface capacity, traffic type, and policy settings affect QoS bound calculations and whether violations occur.

**Table 15: Example policies and QoS bound calculations**

| Example configuration | QoS bound | QoS violation (Positive # = Violation) |
|---|---|---|
| Interface capacity = 10,000 Mbps<br>Undifferentiated traffic = 5000 Mbps<br>Normal operation policy = 60% | 6000 Mbps (60%) | -1000 Mbps (-10%)<br>Because this number is negative, there is no capacity violation. |
| Interface capacity = 10,000 Mbps<br>Undifferentiated traffic = 8000 Mbps<br>Normal operation policy = 60% | 6000 Mbps (60%) | 2000 Mbps (20%)<br>Because this number is positive, there is a capacity violation. |
| Interface capacity = 10,000 Mbps<br>Voice traffic = 6000 Mbps<br>Video traffic = 2000 Mbps<br>Voice normal operation policy = 90%<br>Video normal operation policy = 60% | Voice = 9000 Mbps (90%)<br>Video = 6000 Mbps (60%) | Voice = -3000 Mbps (30%)<br>Video = -4000 Mbps (40%) |

## Interface queue properties and QoS bound calculations

Provides rules and attributes for determining QoS bounds of interface queues, including priority, weight, and police limits.

Cisco Crosswork Planning simultaneously calculates QoS bound for each queue in the interface. Cisco Crosswork Planning uses the Interface queue parameters (Priority, Weight, and Police limit) and measured or simulated traffic for all queues in the interface. Priority is always considered first. If there are queues of equal Priority, then Weight is applied next.

- Queues with priority 1 share all available interface capacity. Their weight and police limits further refine how much each priority 1 queue can use (their QoS bound). Each priority 1 queue can borrow available capacity from other priority 1 queues up to the limit of their QoS bound.
- The available capacity for priority 2 queues is the total interface capacity minus all capacity consumed by priority 1 queues. The process then begins again for all priority 2 queues. Their weight and police limits determine their QoS bound, and priority 2 queues can borrow capacity from each other up to the limits set by the QoS bound.
- This process continues for each successive priority level. Traffic that is outside any QoS bound is dropped to the lowest priority of all traffic on the interface.

For discovered networks with measured traffic, if no Cisco Crosswork Planning QoS parameters are set, the QoS bound is based on whatever capacity percentages the live network has for each queue.

**Priority**

Lists how QoS bounds are calculated for priority queues and provides examples.

If policies are not set that further affect the QoS bound, a queue's QoS bound is calculated using this pattern:

- Priority 1 QoS bound = 100% of the interface capacity.
- Priority 2 QoS bound = Total interface capacity – amount of traffic consumed by priority 1 queues.
- Priority 3 QoS bound = Total interface capacity – amount of traffic consumed by (priority 1 + priority 2 queues).
- QoS bound for each succeeding priority follows this same pattern where the traffic consumed by all higher priority queues is subtracted from the total interface capacity.

**Table 16: Examples of priority QoS bound calculations**

| Example configuration | QoS bound | QoS violation (Positive # = Violation) | QoS bound calculations |
|---|---|---|---|
| Interface capacity = 20,000 Mbps<br>EF traffic = 6000 Mbps; priority = 1<br>BE traffic = 3000 Mbps; no priority set | EF = 20,000 Mbps<br>BE = 14,000 Mbps | EF = -14,000 Mbps<br>BE = -11,000 Mbps | EF = Total interface capacity because it is the only priority 1 queue<br>BE = 20,000 (interface capacity) – 6000 (consumed by higher priority queues) |
| Interface capacity = 10,000 Mbps<br>EF Traffic = 6000 Mbps; priority = 1<br>BE Traffic = 5000 Mbps; priority = 2 | EF = 10,000 Mbps<br>BE = 4000 Mbps | EF = -4000 Mbps<br>BE = 1000 Mbps | EF = Total interface capacity because it is the only priority 1 queue<br>BE = 10,000 (interface capacity) – 6000 (consumed by higher priority queues |

**Weight**

Provides details and examples on how queue weights determine forwarding precedence and how QoS bounds are calculated.

The weight identifies the forwarding precedence for queues of equal priority. If weights for queues of the same priority do not add up to 100%, weights are converted proportionally so they do add up to 100%.

**Table 17: Examples of weight QoS bound calculations**

| Example configuration | QoS bound | QoS violation (Positive # = Violation) | QoS bound calculations |
|---|---|---|---|
| Interface capacity = 10,000 Mbps<br>AF1 traffic = 3000 Mbps; priority = 1; weight = 100%<br>AF2 traffic = 6000 Mbps; priority = 1; weight = 100% | AF1 = 5000 Mbps<br>AF2 = 7000 Mbps | AF1 = -2000 Mbps<br>AF2 = -1000 Mbps | AF1 = Half of capacity for priority 1 queues because both queues have equal weights<br>AF2 = 5000 (half of capacity) + 2000 (unused AF1 capacity) |
| Interface capacity = 10,000 Mbps<br>AF1 = 5000 Mbps; priority = 1; weight = 60%<br>AF2 traffic = 6000 Mbps; priority = 1; weight = 40% | AF1 = 6000 Mbps<br>AF2 = 5000 Mbps | AF1 = -1000 Mbps<br>AF2 = 1000 Mbps | AF1 = 60% of capacity for all priority 1 queues<br>AF2= 10,000 (interface capacity) – 5000 (consumed by AF1 queue) |

### Police limits

Provides an overview of QoS police limits and demonstrates example calculations for queue management.

Priority 1 queues have 100% of the interface traffic, and thus starve out the remaining queues. To prevent this queue starvation, use police limits to configure how much of the maximum percentage should be available for a given priority level.

**Table 18: Examples of police limit QoS bound calculations**

| Example configuration | QoS bound | QoS violation (Positive # = Violation) | QoS bound calculations |
|---|---|---|---|
| Interface capacity = 10,000 Mbps<br>EF traffic = 1000 Mbps; priority = 1; police limit = 50%<br>BE traffic = 2000 Mbps; priority = 2 | EF = 5000 Mbps<br>BE = 9000 Mbps | EF = -4000 Mbps<br>BE = -7000 Mbps | EF = 50% of total interface capacity<br>BE = 10,000 (interface capacity) – 1000 (capacity consumed by EF) |
| Interface capacity = 10,000 Mbps<br>EF traffic = 1000 Mbps; priority = 1; police limit = 5%<br>BE traffic = 2000 Mbps; priority = 2 | F = 500 Mbps<br>BE = 9500 Mbps | EF = 500 Mbps<br>BE = -7500 Mbps | EF = 5% of total interface capacity<br>BE = 10,000 (interface capacity) – 500 (capacity consumed by EF) |

| Example configuration | QoS bound | QoS violation (Positive # = Violation) | QoS bound calculations |
|---|---|---|---|
| Interface capacity = 10,000 Mbps<br>EF = 3000 Mbps; priority = 1; police limit = 20%<br>AF1 traffic = 4000 Mbps; priority = 2; police limit = 75%<br>AF2 traffic = 2500 Mbps; priority = 2; police limit 25% | EF = 2000 Mbps<br>AF1 = 6000 Mbps<br>AF2 = 4000 Mbps | EF = 1000 Mbps<br>AF1 = -2000 Mbps<br>AF2 = -1500 Mbps | EF = 20% of total interface capacity<br>AF1 = 75% of (10,000 [interface capacity] – 2000 [capacity consumed by EF])<br>AF2 = 10,000 (interface capacity) – 2000 (capacity consumed by EF) – 4000 (capacity consumed by AF1) |

## Interface QoS bound calculations using multiple QoS parameters

Provides examples and formulas for calculating QoS bounds for interface queues using priority, weight, and police limits.

Cisco Crosswork Planning calculates a QoS bound for interface queues based on three parameters–priority, weight, and police limit–when all are configured.

**Table 19: Examples of interface QoS bound calculations using multiple QoS parameters**

| Example configuration | QoS bound | QoS violation (Positive #= Violation) | QoS bound calculation |
|---|---|---|---|
| Interface capacity = 10,000 Mbps<br>EF = 3000 Mbps; priority = 1; police limit = 20%<br>AF1 traffic = 4000 Mbps; priority = 2; weight = 75%<br>AF2 traffic = 2500 Mbps; priority = 2; weight = 25% | EF = 2000 Mbps<br>AF1 = 6000 Mbps<br>AF2 = 4000 Mbps | EF = 1000 Mbps<br>AF1 = -2000 Mbps<br>AF2 = -1500 Mbps | EF = 20% of total interface capacity<br>AF1 = Maximum of these two values.<br>• 75% of (10,000 [interface capacity] – 2000 [capacity consumed by EF])<br>• 8000 (available capacity) – 2500 (AF2 traffic)<br>AF2 = Maximum of these two values.<br>• 25% of (10,000 [interface capacity] – 2000 [capacity consumed by EF])<br>• 8000 (available capacity) – 4000 (AF1 traffic) |

### Service class QoS bound calculations using multiple QoS parameters

Provides details on how QoS bounds are calculated for service classes and queues using multiple QoS parameters.

When service classes have defined policies and are mapped to queues, Cisco Crosswork Planning calculates a QoS bound for both. The system then uses the lowest value of the two so as to enforce restrictions in the strictest possible manner.

**Example calculation**

- Interface capacity = 10,000 Mbps
- QoS bound for service class = 50%, or 5000 Mbps based on policy
- QoS bound for EF queue = 7500 Mbps based on combined parameters of priority, weight, and police limit

The QoS bound for this service class is 5000 Mbps because the policy QoS bound calculation is lower.

## Queue and service class information

Provides instructions for viewing queue and service class data in interface tables and visualizations.

This table lists the options to view different types of queue and service class information.

**Table 20: Queue and service class information**

| To view | Show or select |
|---|---|
| Queue information | Show the Interface Queues table. Select the queue from the **QoS** drop-down list in the visualization toolbar. Both the network plot and the **Traff meas** and **Traff sim** columns display traffic data specific to the queue type selected. |
| Per-queue traffic in the Interfaces table | Select the queue from the **QoS** drop-down list. Both the network plot and the **Traff meas** and **Traff sim** columns display traffic data specific to the queue type selected. |
| Service class demands | Show the **Service class** column in the Demands table. |
| Per-service-class traffic | Select the service class from the **QoS** drop-down list. Both the network plot and the **Traff meas** and **Traff sim** columns in the Interfaces table display data specific to the service class selected. |

## QoS bound and QoS violation columns

Lists and describes the available columns for monitoring QoS bounds and violations in network interface, circuit, and queue tables.

This table lists the available column options to display numeric values of the QoS bound calculations. For information on QoS values as they relate to VPNs, refer to Simulate VPN on page 197.

**Table 21: QoS bounds and QoS violations**

| To view | Show this column in Interfaces, Circuits, or Interface queues table |
|---|---|
| **Measured Traffic** | |
| Maximum capacity before a QoS bound is violated under normal operations | QoS bound meas |
| QoS bound as a percentage of total interface capacity | QoS bound meas (%) |
| QoS violations under normal operations; if the number is positive, there is a violation | QoS violation meas |
| QoS violation as a percent of the total interface capacity | QoS violation meas (%) |
| **Simulated Traffic** | |
| Maximum capacity before a QoS bound is violated under normal operations | QoS bound sim |
| QoS bound as a percentage of total interface capacity | QoS bound sim (%) |
| QoS violations under normal operations; if the number is positive, there is a violation | QoS violation sim |
| QoS violation as a percent of the total interface capacity | QoS violation sim (%) |
| **Worst-Case Traffic** | |
| Maximum capacity before a QoS bound is violated under worst-case operations | WC QoS bound |
| WC QoS bound as a percentage of total interface capacity | WC QoS bound (%) |
| QoS violations under worst-case operations; if the number is positive, there is a violation | WC QoS violation |
| WC QoS violation as a percent of the total interface capacity | WC QoS violation (%) |
| Service class causing the worst-case utilization | WC service class |

## Create queues

Guides you through creating new interface queues by manually adding them and configuring queue properties.

Use this task to create new interface queues.

Cisco Crosswork Planning discovers queues. However, you can manually add them. Once discovered or created, queues appear in the Interface Queues table.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **Others** > **Interface Queues**.

   Alternatively, in the Network Summary panel, click  in the **Interface queues** tab.

The **Interface queues** tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **Interface queues** check box.

3. Select the required interfaces.
4. Click **Next**.
5. Enter the queue name.
6. Optionally enter the queue properties of priority, weight, and police limit. For information on how these queue properties behave, refer to Interface queue properties and QoS bound calculations on page 186.
7. Submit your changes.

---

The new queue appears as an option in the **QoS** drop-down list in the toolbar.

**What to do next**

Map a service class to the queue. For instructions, refer to Map service classes to queues on page 193.

### Edit interface queue properties

Provides instructions for editing interface queue properties, including accessing the Interface queues table and modifying QoS parameters.

Complete these steps to edit the QoS parameters using Interface queues properties.

**Procedure**

---

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select the interface queues from the **Interface queues** table.

   If the **Interface queues** tab is not visible, then click the **Show/hide tables** icon ( ), select the **Interface queues** check box, and click **Apply**.
3. Click .
4. Update one or more QoS fields: Priority, Weight, and Police limit. This enables you to create the desired QoS requirement.
5. Save your changes.

---

## Create service classes

Provides instructions for creating a new service class, including naming, mapping to queues, and saving the configuration within the QoS management interface.

Use this task to create new service classes.

**Procedure**

---

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, select **Manage QoS** from the **QoS** drop-down list, or use the **Actions** > **Edit** > **Manage QoS** option.
3. In the **Service Classes** section, click .

An empty row appears.

4. Enter a unique name under the **Name** column.
5. (Optional) If queues exist and if you want to map this new service class to one or more queues, select them from the list.

   If queues do not exist, but you want them, you must manually create the queues and then return to this dialog box to select them. For more information, see Map service classes to queues on page 193.
6. Save the changes.

The newly created service class appears under this section and in the **QoS** drop-down list in the toolbar.

### Map service classes to queues

Provides instructions for mapping service classes to interface queues.

Use this task to map service classes to interface queues.

#### Before you begin

To map service classes to queues, those queues must first exist either because they were discovered by Cisco Crosswork Planning or because you manually added them. For more information on creating queues, refer to Create queues on page 191.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, select **Manage QoS** from the **QoS** drop-down list, or use the **Actions** > **Edit** > **Manage QoS** option.
3. In the Service Classes section, select the service class.
4. Click **Edit**.
5. In the **Mapped to queues** column, select one or more queues.
6. Save your changes.
7. Repeat for each service class to which you are mapping queues.

## Create policy groups for interfaces

Describes how to configure a policy group for network interfaces to manage QoS settings collectively.

Creating policy groups for interfaces allows you to apply consistent QoS settings across multiple interfaces from the Manage QoS page.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select one or more interfaces from the **Interfaces** table.
3. Click .
4. Click the **Advanced** tab.
5. Expand the **QoS** panel.

- To add this interface to an existing policy group, select it from the drop-down list.
- To add a new policy group, type the name of the new group in the drop-down field and click **Add** "*PoliyGroupName*".

6. Save your changes.
7. To assign a service class to this policy group, select this policy group while configuring the Service Class Policy. For more information, refer to Create or edit service class policies on page 194.

## Create or edit service class policies

Provides instructions for creating or editing service class policies, including configuring policies for undifferentiated traffic and service classes.

This topic describes how to create or edit service class policies.

You can configure policies for undifferentiated traffic and for service classes.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, select **Manage QoS** from the **QoS** drop-down list, or use the **Actions** > **Edit** > **Manage QoS** option.
3. To click a new service class policy:
   a) In the **Service Class Policies** section, click ⊞.

      The New Service Class Policy page opens.

   b) If creating a policy for undifferentiated traffic, select **Undifferentiated traffic**. If creating a policy for an existing service class, select **Service class**. Then, select the service class from the **Service Class** drop-down list.
   c) To apply this service class mapping to a group of interfaces, select from or enter the name in the **Interface policy group** drop-down list. You can enter a name that does not exist and a policy group will be created with the name you entered. For details on creating policy groups for interfaces, refer to Create policy groups for interfaces on page 193.
   d) In **Normal operation (in %)**, enter the percentage of bandwidth capacity that you do not want this interface (or group of interfaces) to exceed for this traffic or service class under normal conditions.
   e) In **Worst-case (in %)**, enter the percentage of bandwidth capacity that you do not want this interface (or group of interfaces) to exceed for this traffic or service under worst-case operating conditions.
   f) Click **Submit**.
4. To edit an existing service class policy:
   a) In the **Service Class Policies** section, select the service class row you want to edit.
   b) In the **Actions** column, click ••• > **Edit**.

      The Edit Service Class Policy page opens.

   c) Enter the details as described in Step 3.
   d) Click **Save**.

## Apply QoS settings globally

Guides you through applying QoS settings globally, including selecting service classes or queues and updating your network plan to reflect changes in network summary tables and plot displays.

Selecting QoS settings and the traffic level at the global level determines the network summary table calculations and what is displayed in the network plot.

### Before you begin

Ensure that the required queue or the service class is available.

Complete these steps to apply the QoS parameters globally for an opened plan file.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, select **Manage QoS** from the **QoS** drop-down list, or use the **Actions** > **Edit** > **Manage QoS** option.
3. Select the service class or queue from the **QoS** drop-down list.

The selected QoS setting applies globally to your network plan. The network plot and the network summary table display data based on the selected QoS parameter.

## How QoS violations are indicated in the topology

Outlines how QoS violations are visually indicated in the topology.

### Summary

Cisco Crosswork Planning visually signals when network interfaces exceed QoS bounds. This feature helps you quickly identify and address potential performance issues.

The key components involved in the process are:

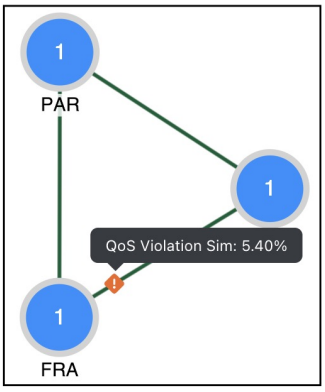
- Topology plot: Displays network interfaces and their current status.
- QoS violation icon ( ): Marks interfaces that violate QoS thresholds.
- Severity text: Provides details about the violation.

### Workflow

The process includes these stages:

1. The system monitors interface utilization and QoS thresholds.
2. When an interface exceeds its QoS bound, a violation icon appears on the topology plot.
3. The severity text provides details about the violation.

**Figure 55: QoS violation indication**




**Note**

Cisco Crosswork Planning indicates when an interface violates a QoS bound in the network plot, but does not display the QoS bound values on the interfaces.

# 13 Simulate VPN

**Topics:**


- VPN model
- VPNs
- VPN node properties
- Layer 3 VPN example
- VPN simulation analysis


Outlines VPN simulation concepts, including VPN modeling, object and topology management, VPN creation and interface identification, and analysis techniques to support comprehensive VPN deployment and validation.

# VPN model

Describes the virtual private network (VPN) architecture and configuration framework used to establish secure, encrypted connections between remote sites or users and the corporate network over public internet infrastructure.

A VPN model is a representation of a virtual subnetwork within the overall network architecture. Viewing and simulating VPN within Cisco Crosswork Planning helps with many network tasks and answers these questions:

- Which VPNs are on my network? Where and how are they configured?
- Which VPNs are using congested interfaces?
- Which VPNs will experience congestion under any of a given list of failure scenarios?
- Which failure scenarios cause the worst-case congestion or latency for a VPN?

There are many varieties of VPNs. For example, there are Layer 2 (L2) VPNs and Layer 3 (L3) VPNs, each with different categories within it, as well as vendor-specific VPN implementations. Each VPN type has its own specific configuration and terminology. The Cisco Crosswork Planning VPN model supports a number of these VPN types based on either route-target or full-mesh connectivity.

## VPN objects

Lists VPN objects and their descriptions with examples for both Layer 2 and Layer 3 VPN configurations.

This section lists the VPN objects and their descriptions with examples for both Layer 2 and Layer 3 VPN configurations.

| Object | Description | Examples |
|---|---|---|
| VPNs | A set of VPN nodes that can exchange data with each other. | • Layer 2 VPN: The VPN represents an individual VPLS containing Virtual Switch Interfaces (VSIs).<br>• Layer 3 VPN: The VPN represents sets of VRFs associated with a set of VPN nodes that forward traffic between themselves. This set of VRFs signifies a single customer or service. |
| VPN nodes | Connection points in a VPN. They exist on standard nodes, and each node can contain multiple VPN nodes. A VPN node can be in only one VPN. | • Layer 2 VPN: The VPN node represents the VSIs configured on each router.<br>• Layer 3 VPN: The VPN node represents the VRF instances configured on each router. |

## VPN topology and connectivity

Describes how VPN topology route connections are established and how connectivity properties affect traffic flow calculations.

VPN topology and connectivity is a network configuration that

- are established through Route Targets (RTs) or through a full mesh of VPN nodes

- enables Cisco Crosswork Planning to calculate which demands between VPN nodes carry traffic for a particular VPN, and
- allows calculation of VPN vulnerability to certain failure and congestion scenarios.

A demand is associated with a VPN, meaning it carries traffic for that VPN, if these conditions are true:

- Two VPN nodes are in the same VPN.
- Demand is in the same service class as the VPN.
- Only for VPNs with RT connectivity, the **RT export** property of one VPN node must match the **RT import** property of another VPN node.

After demands are associated with the VPN, this configuration simulates the associated access circuits exchanging traffic as if they were on the same LAN. Note that a demand associated with a VPN can additionally contain other traffic that is for that VPN.

#### VPN connectivity types

The **Connectivity** property is set in the Add/Edit VPN page.




This table compares connectivity types and their characteristics.

| Connectivity | Description |
|---|---|
| **RT** | Route targets model the more complex connectivity used in Layer 3 VPNs, such as hub-and-spoke networks. Here, the VRFs exchange data with one another based on the matching of RT export and RT import properties set for each VPN node. However, having an import/export pair does not create bidirectional communication. Instead, traffic flows in the opposite direction of the routed advertisements.<br><br>For example, if node A's RT import matches node B's RT export, traffic can flow from node A to B. For traffic to flow from node B back to node A, node B must have an RT import that matches an RT export of node A. This combination of matching imported and exported RTs defines which VPN nodes can exchange data. The VPN name identifies the VPN itself. |
| **Full Mesh** | Full mesh connectivity refers to a complete mesh of connections between VPN nodes in a VPN where they can all communicate with each other. This connectivity is typical in a VPLS, where all VSIs identify one another based on a common AGI. |

## VPNs

Explains how VPNs create secure connections across public networks by defining key properties and routing approaches.

A VPN is a network construct that

- consists of a set of VPN nodes that can exchange data within it
- has key properties that uniquely identify the VPN, and

- defines how the traffic within the VPN is routed.

#### VPN properties

These are the key properties of the VPNs:

- **Name**: Unique name of the VPN.
- **Type**: Type of VPN. The options are VPWS, VPLS, or L3VPN.
- **Connectivity**: Determines how Cisco Crosswork Planning calculates connectivity and associated demands for VPNs. The options are:
  - **Full Mesh**: Connectivity is between all nodes in the VPN. Cisco Crosswork Planning ignores the RT import and RT export properties of the VPN nodes.
  - **RT**: Connectivity is based on the RT import and RT export properties of its VPN nodes.
- **Service class**: Service class associated with the VPN.

Once the VPN is created, it appears in the **VPN** drop-down list of VPN nodes.

## Create VPNs

Describes how to create new VPNs in the system, which can then be configured with VPN nodes as needed for your network topology.

You can create new VPNs and then later add VPN nodes to them (refer to Add VPN nodes to VPNs on page 203).

#### Create new VPNs

Explains how to configure new VPN connections with customized settings and parameters.

Use this task when you need to establish new VPN connections in your network design plan. You can create different types of VPNs including L3VPN, VPLS, and VPWS with various connectivity and service configurations.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **VPNs** > **VPN**.

   Alternatively, in the Network Summary panel on the right side, click  in the **VPNs** tab.

   The VPNs tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon () and check the **VPNs** check box.
3. In the **Name** field, enter a unique name for the VPN.
4. From the **Type** drop-down list, choose a VPN type.

   The options are L3VPN, VPLS, and VPWS.
5. Select the Connectivity type. The options are RT or Full Mesh.
6. Select the Service class for the VPN.
7. Click **Add**.
8. (Optional) Add VPN nodes to the newly created VPN. For more information, refer to Add VPN nodes to VPNs on page 203.

The new VPN is created with the specified configuration settings and appears in your plan file.

## VPNs table

Lists the VPN properties, associated service class, traffic, and number of VPN nodes within each VPN.

The VPNs table lists the VPN properties, its associated service class, traffic, and the number of VPN nodes within that VPN (Table 22: VPNs table columns on page 201). For information on QoS measurements, refer to Simulate Quality of Service (QoS) on page 183. For information on the Worst-Case columns not listed here, refer to Table 24: Updated columns in the VPNs table on page 208.

**Note**

- Because the traffic and QoS calculations are based on all interfaces within the VPN for the service class specified for that VPN, the plot view might differ from the table. For example, the plot view could show Internet traffic while a VPN carrying voice traffic is selected.
- All traffic and QoS violations are based on traffic carried by all interfaces used by the VPN for the service class defined for that VPN.

**Table 22: VPNs table columns**

| Column | Description |
|---|---|
| Service class | Service class associated with this VPN. All values within the table are associated with this service class. |
| Num nodes | Number of VPN nodes in this VPN. |
| Util meas | The maximum measured utilization of all interfaces used by this VPN. |
| Util sim | The maximum simulated utilization of all interfaces used by this VPN. |
| Total src traff meas | Total amount of measured source traffic on this VPN. |
| Total dest traff meas | Total amount of measured destination traffic on this VPN. |
| QoS violation sim | Maximum QoS violation under normal operations for all simulated traffic for all interfaces used by this VPN. If the number is positive, there is a violation. |
| QoS violation sim (%) | QoS violation as a percent of the total simulated interface capacity. |
| QoS violation meas | Maximum QoS violation under normal operations for all measured traffic for all interfaces used by this VPN. If the number is positive, there is a violation. |
| QoS violation meas (%) | QoS violation as a percent of the total measured interface capacity. |
| Latency | Maximum latency of all demands used by this VPN. |
| Tags | User-defined identifiers that makes it easy to group VPNs. |

VPNs are not selectable from the network plot. You can only select and filter VPNs only through tables. When selected, all VPN nodes within the VPN are highlighted in the plot (Figure 56: VPN nodes within a VPN on page 204).

### Identify interfaces used by VPNs

Explains how to view which interfaces are associated with a VPN by using the filter options.

Use this task to identify which interfaces are associated with a specific VPN and visualize the VPN in the network plot.

**Procedure**

1. Select the required VPN from the **VPNs** table.
2. Click > **Filter to interfaces**.

   **Note**

   Utilization measurements might be different between the tables because the VPN table calculates measurements only for the service class associated with that VPN.

3. Select the filtered interfaces to see the VPN outlined in the network plot.

The system displays the interfaces associated with the selected VPN. When you select the filtered interfaces, the VPN appears outlined in the network plot for visual representation.

## VPN node properties

Lists the properties that define VPN nodes and determine which VPNs the nodes belong to and how the demands are routed.

VPN nodes are defined by the properties that determine which VPNs the nodes belong to and how the demands are routed.

- **Site**: Name of the site on which the VPN node resides.
- **Node**: Name of the node on which the VPN node resides. This node name corresponds to the one in the Nodes table.
- **Type**: The type of VPN. The options include VPWS, VPLS, or L3VPN. Alternatively, you can enter a string value to create a new one. Once entered, the new VPN type appears in the drop-down list and is available for other VPN nodes and VPNs.
- **Name**: Name of the VPN node.
- **VPN**: Name of the VPN in which this VPN node resides. The drop-down list shows existing VPNs of the same type set in the Type field. You can create a VPN node without setting its VPN, but without it, the VPN node is not included in simulations as a member of any VPN.
- **Description**: Description for the VPN node.

  To simulate RT connectivity, you must set the VPN Connectivity property to RT and then set the RT import and RT export properties on the individual VPN nodes within it.
- **RT import** and **RT export**: The pairing of RT values identifies which VPN nodes connect with each other. For more information, refer to VPN topology and connectivity on page 198.
- **RD**: Route Distinguisher (RD) uniquely identifies routes within a VRF as belonging to one VPN or another, thus enabling duplicate routes to be unique within a global routing table.

## Create VPN nodes

Describes how to configure VPN nodes within a network design plan to establish VPN connectivity between sites and nodes.

Use this task to create VPN nodes.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **VPNs** > **VPN node**.

   Alternatively, in the Network Summary panel on the right side, click  in the **VPN nodes** tab.

   The VPN nodes tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon () and check the **VPN nodes** check box.
3. Click .
4. Enter the VPN node configuration details.
   a) In the **Site** and **Node** fields, choose the site in which the VPN node will exist, and choose the node on which the VPN node is being configured.
   b) From the **Type** drop-down list, choose a VPN type. The options are: L3VPN, VPLS, and VPWS.
   c) In the **Name** field, enter the name of the VPN node, which does not have to be unique.
   d) From the **VPN** drop-down list, choose the VPN to which you are adding this VPN node. If you do not see the VPN that you expect to see, check if you have selected the correct VPN type in the **Type** drop-down list.
   e) Optionally, enter a description that identifies the VPN node. For example, a customer name might be helpful.
   f) If the Connectivity for the VPN is RT, enter the applicable route targets in the **RT import** and **RT export** fields. All VPN nodes with the same RT import as another VPN node's RT export can receive traffic from that VPN node. Those VPN nodes with the same RT export as another VPN node's RT import can send traffic to that VPN node.
   g) Optionally, in the **RD** field, enter a route distinguisher.
5. Click **Add** to create the VPN node.

The VPN node is created and appears in the VPN nodes table. The VPN node is now associated with the specified site, node, and VPN instance with the configured connectivity settings.

## Add VPN nodes to VPNs

Explains how to configure VPN nodes by adding them to existing VPNs in the Network Design page.

Use this task to add VPN nodes to a specific VPN.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, select one or more VPN nodes in the **VPN nodes** table and click .

> **Note**
>
> If you are editing a single VPN node, you can also use the ••• > **Edit** option under the **Actions** column.

3. In the **VPN** drop-down list, choose the VPN to which you are adding the VPN nodes. If you do not see the VPN that you expect to see, check if you have selected the correct VPN type in the **Type** drop-down list.
4. Save the changes.

---

The VPN nodes are successfully added to the selected VPN and the configuration is saved.

## VPN nodes table

Lists the VPN node properties, as well as columns that identify the VPN nodes' relationship within the VPN and its traffic.

The VPN nodes table lists the VPN node properties. It also includes columns that show each VPN nodes' relationship within the VPN and its traffic.

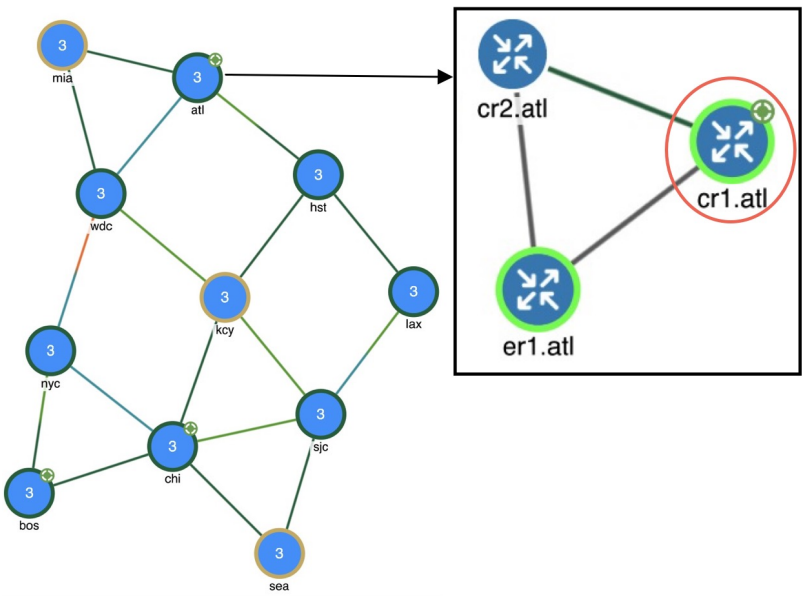
**Table 23: VPN nodes table**

| Column | Description |
|---|---|
| Total connect | Number of VPN nodes that are connected to this VPN node as defined by the RT Import and RT Export pairings. These may or may not be in the same VPN. |
| VPN connect | Number of VPN nodes that are connected to this VPN node and are in the same VPN as defined by the VPN column. |
| Num VPN nodes | Number of nodes in the VPN that this VPN node belongs to as defined by the VPN column. This value is "na" if the VPN node does not belong to a VPN. |
| Src traff meas | Total amount of measured traffic entering the VPN at this node (source traffic). |
| Dest traff meas | Total amount of measured traffic leaving the VPN at this node (destination traffic). |
| Tags | A user-defined identifier that makes it easy to group VPN nodes into a single VPN. If you give a VPN node a tag, when you create a VPN later, you can identify its VPN nodes using tags. |

Once selected from the VPN nodes or VPNs table, the associated site and the nodes within that site appear with a green circle on it (Figure 56: VPN nodes within a VPN on page 204).

> **Note**
>
> VPN nodes are not selectable from the network plot. You can only select and filter to them through tables.

**Figure 56: VPN nodes within a VPN**




## Layer 3 VPN example

Describes a scenario where the Acme manufacturing company has three offices, with branch offices exchanging data only with headquarters.

This example illustrates a scenario where the Acme manufacturing company has three offices, but permits the two branch (er1.par and er1.fra) offices to exchange data only with headquarters (er1.lon). Additionally, the headquarters communicates with an SP VPN node (er1.bru) that is not in the Acme VPN.

**VPN configuration details**

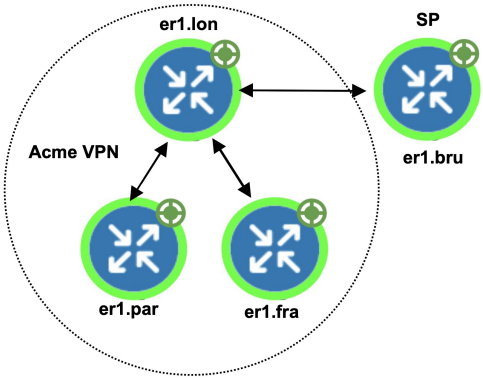
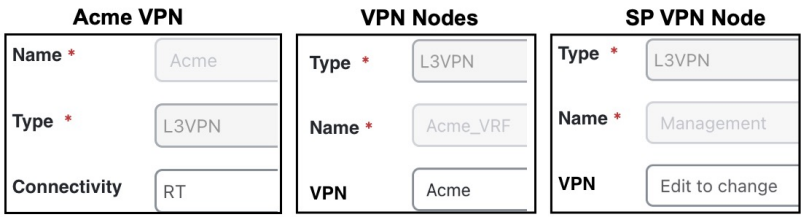



Figure 57: Example RT connectivity and Acme VPN footprint on page 206 shows the footprint of the Acme VPN and the RTs set for all VPN nodes in this example.

- The VPN is named Acme, and it is set to a Connectivity of RT and a Type of L3VPN.
- In turn, each branch office is set to the Acme VPN with a Type of L3VPN.
- To exchange data with two other VPN nodes in the Acme VPN, headquarters (er1.lon) imports the offices' exported route targets of 2:1 (er1.par) and 3:1 (er1.fra).
- In turn, headquarters (er1.lon) exports a route target of 1:1. All three other VPN nodes import it (both offices and the SP VPN node).
- Because the SP VPN node (er1.bru) is not in the Acme VPN, its communication with er1.lon is not within the context of the Acme VPN.




The VPN footprint in Figure 57: Example RT connectivity and Acme VPN footprint on page 206 shows that if the circuit between er1.fra and er1.bru becomes congested or fails, the VPN is impacted. However, a failure of the circuit between the two branch offices is not impacted. This failure is illustrated in Figure 58: Example failure between branch offices in the Acme VPN on page 206, which shows that none of the demands associated with the VPN are rerouted.

**Figure 57: Example RT connectivity and Acme VPN footprint**

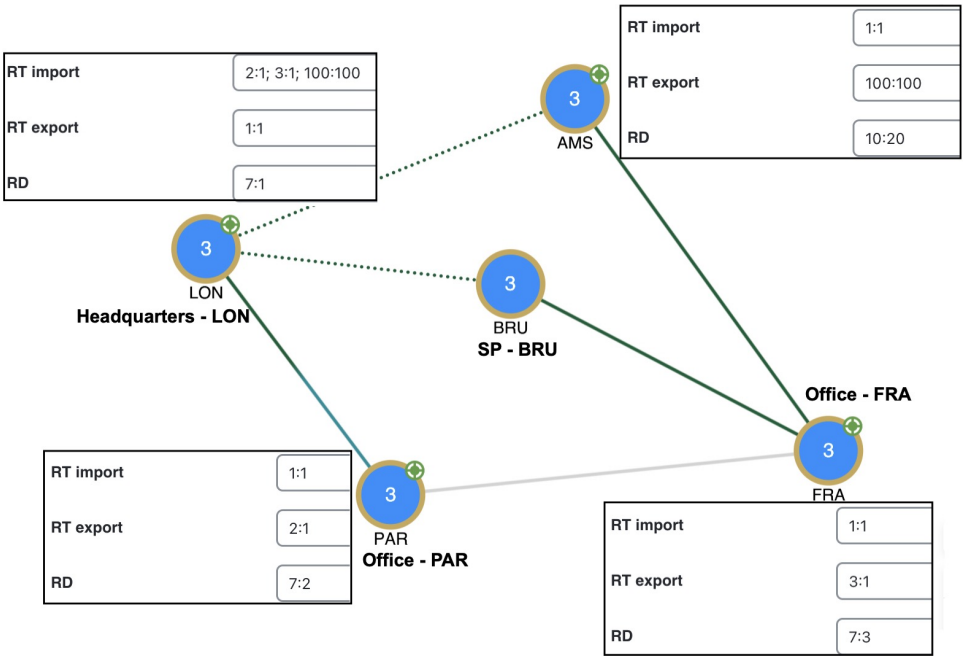
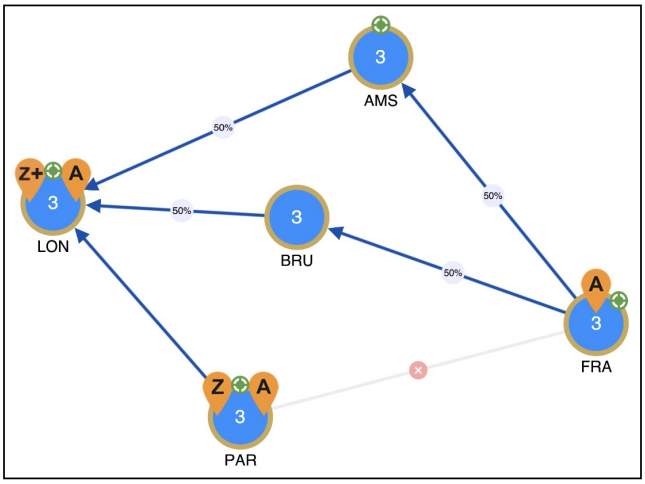



**Figure 58: Example failure between branch offices in the Acme VPN**



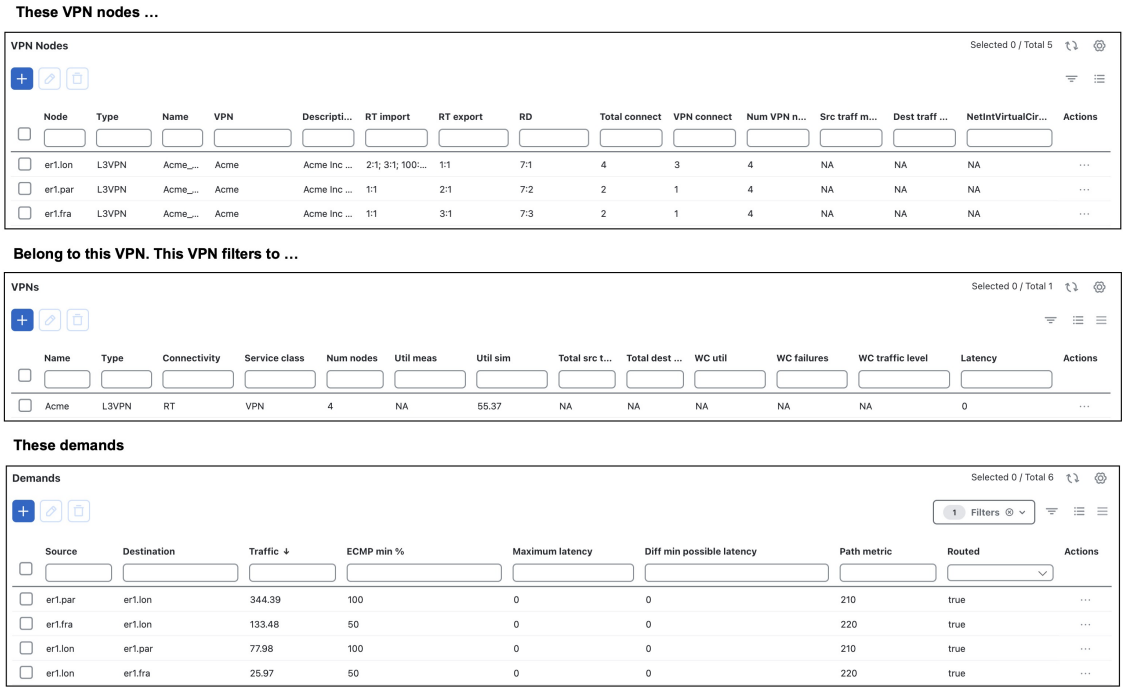

For this example, Figure 59: VPN nodes belonging to Acme VPN, and Acme VPN filtered to demands on page 207 illustrates the VPN nodes belonging to Acme VPN and the filtering of the Acme VPN to its associated demand traffic. It also shows the calculations of the **Total connect** and **VPN connect** columns in the VPN Nodes table.

- The **Total connect** for the VPN node residing on er1.lon headquarters is the highest because it exchanges data with three other VPN nodes. Each of the offices and the service provider VPN node have 2 in the Total connect column.
- The **VPN connect** for the VPN node at the er1.lon headquarters is the highest because it exchanges data with and is in the same VPN as the two offices. All three VPN nodes share the same VPN name.

  Each office has 1 in the VPN connect column because it communicates with only one VPN node in the same VPN.

  The service provider VPN node (er1.bru) has 0 VPN connects because it is not part of the Acme VPN.

**Figure 59: VPN nodes belonging to Acme VPN, and Acme VPN filtered to demands**

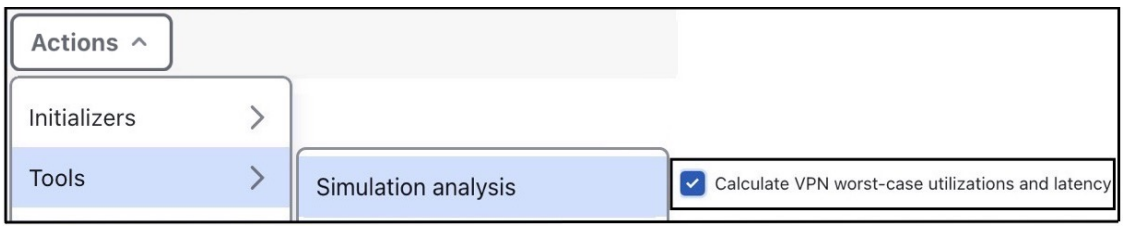



## VPN simulation analysis

Lists the columns updated in the VPNs table after running simulation analysis and provides options for VPN failure testing.

When running the **Simulation analysis** tool, you can record worst-case utilization and latency for VPNs. Then, you can select a VPN to fail to its worst-case utilization or worst-case latency using the ••• > **Fail to WC** or **Fail to WC latency** options, respectively.

> **Note**
>
> All calculations are based on traffic carried by all interfaces used by the VPN for the service class defined for that VPN.



### Columns updated in the VPNs table after simulation analysis

These columns are updated in the **VPNs** table after completion of Simulation analysis.

**Table 24: Updated columns in the VPNs table**

| Column | Description |
|---|---|
| WC util | Worst-case VPN utilization over all failure scenarios. |
| WC failures | Failures causing the worst-case utilization of the VPN. |
| WC traffic level | Traffic level causing the utilization of the interface identified in the **WC util** column. |
| WC QoS violation | Highest worst-case QoS violation for all interfaces used by this VPN. A QoS violation is equal to the worst-case traffic minus the worst-case capacity permitted (worst-case QoS bound). |
| WC QoS violation (%) | Highest worst-case QoS violation for all interfaces in this VPN expressed as a percentage of total capacity. |
| WC latency | Maximum VPN latency over failure scenarios considered. |
| WC latency failures | Failures causing the worst-case VPN latency. |

# 14 Simulate Advanced Routing with External Endpoints

**Topics:**


- External endpoints
- Create external endpoints and add members
- Simulate routing with external endpoints


Explains advanced routing strategies using external endpoints, demonstrating configuration, simulation, member specification, and traffic distribution to enhance network routing flexibility and evaluate use cases for integrating external resources.

# External endpoints

Explains how external endpoints act as identified entry and exit points for network traffic and how they are used for advanced routing scenarios.

An external endpoint is a Cisco Crosswork Planning object that

- identifies specific entry (source) or exit (destination) points for demands, and
- consists of one or more members that are defined as nodes, external ASes, or a combination of an external AS and external node.

By setting a demand's source or destination to an external endpoint, you can simulate traffic going from multiple sources to a single destination, from a single source to multiple destinations, or multiple sources going to multiple destinations. Because of this flexibility, they are useful for specifying secondary entry and exit points in the event of failures.

## Use cases for routing with external endpoints

Lists common scenarios and benefits of routing with external endpoints in network modeling and simulation.

To model basic IGP routing, demands are sourced or destined for nodes within the topology. To model basic inter-AS routing, the sources and destinations are neighboring external ASes, or a combination of the external AS and the peering node in that AS. However, more complex routing situations require the use of external endpoints as the source or destination. External endpoints can contain multiple member nodes and ASes, and you can specify when traffic enters or exits from each of them individually. This allows you to simulate routing within and between ASes where multiple traffic entry and exit points are used simultaneously. You can also prioritize where the traffic fails over to other nodes and ASes.

Use cases for both IGP and inter-AS routing with external endpoints include:

- Simulating content caching failovers for in-network source of demands that are backed up by another in-network source. If connectivity is lost to the first, traffic is sourced from the second.
- Simulating edge routing with a single entry point into the network edge and a specific failover point. Alternatively, you could model multiple entry points, depending on which is closest to the destination.
- Simulating complex BGP routing policies from a transit provider. For example, you can specify a transit entry location and failover location per destination.
- Simulating failover between peering ASes; for example, from one single-connection transit provider to another.

# Create external endpoints and add members

Explains how to configure external endpoints and add their members within your network design plan.

Use this task to enable traffic routing and demand management by defining external endpoints and their members in a network design plan.

The recommended method of creating external endpoint members is to do so while creating the associated external endpoint.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **Demands** > **External endpoints**.

   Alternatively, in the Network Summary panel, click  in the **External endpoints** tab.

The External endpoints tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **External endpoints** check box.

3. In the **Name** field, enter a unique external endpoint name.
4. To add a new member, click . To edit an existing member, select the member from the list and click .

   An External Endpoint Member page opens (refer to Figure 60: External endpoint member settings on page 211). Specify the parameters described in External endpoint member attributes on page 211 and click **Add**.
5. Click **Add** in the Add External Endpoint page.

## External endpoint member attributes

Lists the attributes used to define and prioritize external endpoint members, and describes options for traffic distribution.

Each member is assigned properties that prioritize traffic entering and exiting that member, in which order to fail over to another member, and how to distribute traffic for members of equal priority. These properties are set when you create the member. All members appear in the External endpoint members table.

**Figure 60: External endpoint member settings**




- **Member endpoint**: Defines whether the member is a node, AS, or external node through an AS.
  - For nodes, choose the site and node name.
  - For interfaces, choose the site, node, and interface. Using an interface lets you specify the exact interface on which the demand traffic is going into or out of a node.
  - For ASes, choose the AS name, and select either the “(None)” option or a node name within the AS. The “(None)” option routes traffic evenly across the entire AS.
- **Priority**: The sequential order in which external endpoint members are used in simulations if failures occur.
- **Equal priority routing**: If members have the same priority, this property identifies how the traffic is distributed. Select from these options:
  - **Shortest Path**: Use the member that results in the use of shortest path between source and destination.

> **Note**
>
> Members within an external endpoint that have the same priority must either all be Shortest Path or none of them be Shortest Path.

If a member is using the shortest path, the **Traffic balance (%)** field is displayed as disabled.

- **Fix Traffic**: Set the traffic across members of equal priority as defined in the **Traffic balance (%)** field.
- **Deduce Traffic**: Behaves like Fix Traffic by setting traffic across members of equal priority as defined in the Traffic balance (%) field. However, upon running the Demand deduction tool, the Traffic balance (%) field is updated based on the measured traffic in the network. Note that Demand deduction only estimates the traffic balances for external endpoints with a priority that is in use in the current no-failure simulation. Therefore, Deduce Traffic is typically set to Priority 1.

## Simulate routing with external endpoints

Describes how routing simulations select endpoint members by priority and availability, redistribute traffic during failures, and apply BGP routing policies for external AS members.

> **Note**
>
> Although this section describes demands as being sourced from an external endpoint, the same methodology holds true if a demand's destination is an external endpoint.

### Summary

The simulation of routing with external endpoints involves selecting members based on priority and availability, redistributing traffic upon failures, and applying BGP routing policies for external AS members.

The key components involved in the process are:

- External endpoint members: Nodes, ASes, or external nodes that serve as traffic sources or destinations.
- Priority levels: Determine the order in which members are used.
- BGP routing policy: Governs routing when members are external ASes.

### Workflow

The process involves these stages:

1. Members with the highest priority (lowest number) are used as the demand's source. For example, if the external endpoint has two members with a priority of 1, the demand is sourced from both members provided they are available.
2. If some members are not available, their traffic is evenly redistributed to the other top priority members.
3. If none of the top priority members are available to source the traffic but there are next-priority members available, Step 1 is repeated for the next priority external endpoint members.
4. If all members with the same priority fail, the traffic is routed according to the next priority in the sequence.
5. If a failure occurs that does not affect the external endpoint member's ability to send or receive traffic, then traffic is rerouted as usual without a need to use the additional members.

6. If an external endpoint member is an external AS, with or without a node specified, the routing for that member is based on the BGP routing policy set by the AS relationships. The distribution of traffic between external endpoints with the same priority is the same as that for node members.

## Traffic distribution

Lists the rules, properties, and failure-handling logic used for distributing traffic among demand members based on priority, path, and percentage settings.

The traffic distribution through these demands is based on the **Equal priority routing** property, and if applicable, the **Traffic balance (%)** property is used to define the external endpoint members.




- If there is only one member and it is defined as **Shortest Path**, the demand takes the shortest path as defined by the IGP metrics.

  If the **Traffic balance (%)** values are empty for all routable demands, traffic is routed and equally load balanced across the demands with the shortest IGP paths. In the case of multiple internal ASes, the shortest IGP route is the shortest route in the first AS the demand enters.

- If multiple members of the same priority are set to **Shortest Path**, the demand takes the path with the shortest IGP path. If all interfaces between the source members and the destination have the same shortest IGP paths, then the traffic is load balanced equally across them.

- If one or more members of the same priority have their **Equal priority routing** property set to **Fix Traffic** or **Deduce Traffic**, the demand traffic is split according to each member's **Traffic balance (%)** value.

  - If the Traffic balance percentages across sources with the same priority sum to less than 100%, the overall demand traffic is decreased to that percentage.

  - If an external endpoint member of the same priority fails, the traffic on the remaining members increases in proportion, so that the same amount of traffic is still routed.

### Example

Consider Node A has failed. Nodes B, C, and D each with priority 2 and Fix Traffic type. Their traffic balances are 20%, 20%, and 40%, respectively.

| Member | Priority | Traffic balance % | Type |
|---|---|---|---|
| A | 1 | NA | ShortestPath |
| B | 2 | 20 | FixTraffic |
| C | 2 | 20 | FixTraffic |
| D | 2 | 40 | FixTraffic |
| E | 3 | NA | ShortestPath |

The demand has 1000 Mbps of traffic.

- With Node A failed, the demand routes 200 Mbps traffic through node B, 200 Mbps through C, and 400 Mbps through D, totaling 800 Mbps.

- If Node D fails, the demand routes 400 Mbps traffic through Node B and 400 through C. If Node B fails too, the entire 800 Mbps is routed through C.
- If all three priority 2 members fail, 1000 Mbps is routed through node E, which is the priority 3 member.

# 15 Simulate Multicast

**Topics:**


- Source-Specific Multicast (SSM)
- Discovered versus simulated multicast flows
- Configure multicast flow hops for simulations
- Configure Cisco next hop for multicast flows
- Multicast demands
- View multicast flows
- Create demands for multicast flows


Describes multicast simulation capabilities, covering SSM parameter representation, comparison of discovered and simulated multicast flows, configuration of flow hops and next hop, demand creation and simulation, and tools to view and analyze multicast flows.

## Source-Specific Multicast (SSM)

Explains how source-specific multicast delivers packets from a specified source to receivers for improved efficiency and security.

A Source-Specific Multicast (SSM) is a multicast forwarding method that

- delivers multicast packets to receivers only from a specified source address requested by the receiver
- reduces resource requirements and improves security by restricting sources, and
- is specified by an (S,G) parameter per multicast flow.

The (S,G) pair is labeled using a dot-decimal notation similar to IP addresses (for example, 1.1.1.1, 2.2.2.2). Each (S,G) pair is the name of a multicast flow, and each flow is listed in the **Multicast flows** table.

## Discovered versus simulated multicast flows

Lists and compares how discovered multicast flows and simulated multicast flows are handled within Cisco Crosswork Planning, including creation, pathing, and applicable tables.

How you work with multicast flows within Cisco Crosswork Planning differs depending on whether Cisco Crosswork Planning discovered the multicast flow or you are simulating it (Table 25: Discovered versus simulated multicast flows on page 216).

**Table 25: Discovered versus simulated multicast flows**

| | Discovered multicast flow | Simulated multicast flow |
|---|---|---|
| **Creation** | Cisco Crosswork Planning discovers multicast flows–(S,G) pairs)–and multicast traffic is derived from the Multicast Flow Traffic table. | First, manually create the multicast flow, defining both the source (S) and the destinations (G). Next, create a demand that links the source to these multicast destinations, allowing the routing to be simulated through the network. |
| **Hops** | Each discovered multicast flow includes multicast flow hops, which are node-interface combinations through which the multicast path flows. | When you create a demand, it determines the path to take. |
| **External hops** | Cisco Crosswork Planning discovers multicast flow hops on interfaces that are external to the plan file. These are interfaces from a plan node to an external node, or from an external node to a plan node. | When you create a demand, it determines the path to take, but multicast external flow hops are not identified. |
| **Destinations** | Cisco Crosswork Planning does not identify a list of multicast destinations for each flow. | Specify destinations (nodes, interfaces, external ASes, or external endpoints) when you create multicast flows. When creating multicast demands, specify these as multicast destinations. |
| **Applicable tables** | • Multicast flows<br>• Multicast flow hops<br>• Multicast flow external hops | • Multicast flows<br>• Multicast flow destinations<br>• Demands |

## Configure multicast flow hops for simulations

Describes how to configure simulation settings in a plan file to use or disregard multicast flow hops for accurate planning and routing analysis.

Cisco Crosswork Planning uses the network state in multicast simulation. Use this task to set simulations to consider multicast flow hops.

If a plan file contains multicast information, it includes the current hops taken by multicast flows. You can specify that Cisco Crosswork Planning multicast simulations should use these flow hops if possible. This is useful, for example, when calculating incremental routing changes on the current network state, such as those caused by a failure. For planning purposes, when the current state is not relevant, you can change this behavior to disregard the multicast flow hops.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.
3. Click the **Simulation** tab.
4. To use or disregard multicast flow hops in simulation, check or uncheck **Use multicast flow hops**, and click **Save**.

## Configure Cisco next hop for multicast flows

Explains how to configure your plan to use Cisco next hop selection method for multicast traffic flows.

Use this task to set Cisco Crosswork Planning to use Cisco next hops.

For manually inserted demands, all SSM demand traffic for the (S,G) pair (multicast flow) goes through any interface that is traversed by that demand.

However, you can set Cisco Crosswork Planning to use Cisco next hops, which are calculated using a hash on S and G for a multicast flow (S,G). The hash calculation is different in IOS and IOS XR. The default behavior is that of IOS. The IOS XR hash is used on all nodes whose OS field starts with IOS XR.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

   The Network options page opens.
3. Click the **Protocols** tab.
4. Under the **Multicast** section, from the **Multipath selection** drop-down list, choose **CiscoNextHop**.
5. Save your changes.

## Multicast demands

Explains the concept of multicast demands and how they relate to discovered and simulated multicast flows.

A multicast demand is a network traffic requirement that

- specifies how multicast flows are to be handled within a network simulation
- can be inserted manually for discovered multicast flows, and
- must be created manually for simulated multicast flows.

## Simulated multicast demands

Explains how simulated multicast demands are routed and modeled, including interface and external endpoint behaviors.

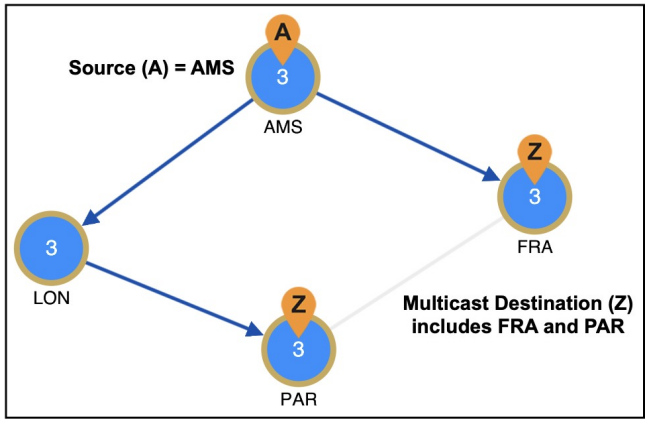
For manually inserted demands, all SSM demand traffic for the (S,G) pair (multicast flow) goes through any interface that is traversed by that demand. SSM demands are routed to take the shortest path from the destination to the source, rather than from the source to the destination, as unicast demands do. By default, multicast multipath is disabled. If two paths of equal cost exit from a node on the route back to the source, the path is chosen based on:

- The remote interface with the highest IP address is used.
- If IP addresses are not available, the router name with the lowest lexicographical name is used.

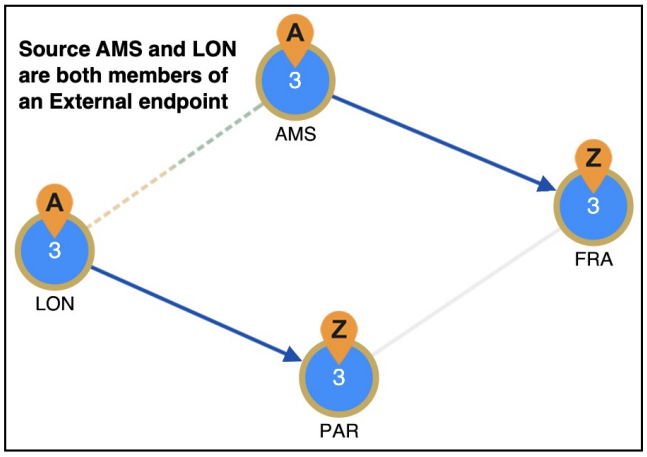
However, you can set Cisco Crosswork Planning to use the Cisco next-hop multicast, multipath selection method. For more information, refer to Configure Cisco next hop for multicast flows on page 217.

The sources of these multicast demands can be nodes, interfaces, external ASes, or external endpoints. Using interfaces lets you specify the exact interface on which demand traffic is entering. Using external endpoints lets you model situations where the multicast source is outside of the nodes in the plan, and there is more than one possible entry point of the traffic flow into and through the interfaces in the plan.

In this example, the selected demand has a node source of AMS, and a multicast destination containing two nodes: FRA and PAR. The source node (S) is marked by A, and the destinations in the receiver group (G) are marked by Z.




You can model a source that is outside the plan using external endpoints. In this example, the demand's source is an external endpoint that contains AMS and LON as members. The multicast destination includes both FRA and PAR.




## View multicast flows

Lists the types of multicast flows and related information available for viewing in the user interface.

To view sources and destinations of multicast demands, select the demand from its table. The source is identified in the plot by A and the destinations are identified by Z.

> **Note**
>
> You can only view the Multicast flow details in the UI if it's already present in your plan file. You cannot create, edit, or delete the multicast flows from the UI.

| To view... | Show this table |
|---|---|
| Multicast flows, including both source and receiver names | Multicast flows |
| Discovered multicast flow hops, including (S,G) name, hop node, and hop interface | Multicast flow hops |

| To view... | Show this table |
|---|---|
| Discovered multicast flow external hops (destinations that are inferred as external to the plan), including (S,G) name, direction of the outbound interface | Multicast flow external hops |
| User-created multicast flow destinations, including (S,G) name and destination node | Multicast flow destinations |
| Multicast demands | Demands |

## Create demands for multicast flows

Explains how to configure a demand to simulate multicast flows in a network plan.

Use this task to create demands for multicast flows.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36) with Multicast details. The plan file opens in the Network Design page.
2. From the toolbar, choose **Actions** > **Insert** > **Demands** > **Demand**.

   Alternatively, in the Network Summary panel, click + > **Demands** in the **Demands** table.
3. In the **Name** field, enter the name of the demand.

4. In the **Source** section, define the source (S) of the multicast flow.
   a) From the **Type** list, select the source as a node, interface, external AS, or external endpoint.
   b) For node sources, select both the site and the node.

      For interface sources, select the site, node, and the interface.

      For external ASes, select both the external AS and the node within the plan file through which this external traffic flows.

      For external endpoints, select its name.

5. In the **Destination** section, define the receiver group (G) of the multicast flow.
   a) From the **Type** list, choose **Multicast destination**.
   b) From the **(G) Receiver** list, choose the receiver that identifies the simulated multicast flow (S,G).
6. Optionally, complete all other fields as needed and then click **Add**.

# Part III

## Traffic Engineering and Optimization

**Topics:**


- Optimize Metrics in the Network Core
- Configure MPLS Routing
- Optimize LSPs
- Configure RSVP-TE Routing
- Optimize RSVP-TE Routing
- Perform Explicit and Tactical Explicit RSVP LSP Optimization
- Configure Segment Routing
- Optimize Segment Routing

# 16 Optimize Metrics in the Network Core

**Topics:**


- Metric optimization tools
- Core and edge distinctions
- Perform metric optimization
- Tactical metric optimization
- Metric optimization reports


Outlines strategies for performing metric optimization tasks, understanding tactical optimization concepts, and accessing metric optimization reports within the network core.

## Metric optimization tools

Introduces metric optimization tools for improving network core metrics, maximizing throughput across routes under various failure scenarios, and resolving local congestion with minimal changes using tactical metric optimization.

**Metric optimization** is a Cisco Crosswork Planning tool that

- optimizes metrics in the network core, and
- replaces the IGP metrics of core interfaces in a network plan with new metrics chosen to maximize the throughput achievable by the resultant routes, under a specified set of failure scenarios.

The **Tactical metric optimization** tool lets you quickly fix local congestion by doing as few modifications as possible on selected interfaces to push utilizations under a given level.

> **Note**
>
> The core of the network must be connected, and all nodes must be contained in one AS. That is, there must be a path between any two core nodes that passes only through other core nodes in the same AS, and not through edge nodes. If this is not the case, Metric optimization signals an error.

## Core and edge distinctions

Explains how nodes and circuits in a network model are classified as core or edge and describes the impact of this distinction on metric optimization policies.

Cisco Crosswork Planning network models allow nodes to be defined as either core nodes or edge nodes. Note that these definitions are not related to any explicit router configurations. A similar classification of circuits is determined implicitly by the nodes they connect. If a circuit is attached to one or more edge nodes, it is defined as an edge circuit. Otherwise, it is a core circuit. Edge metrics and core metrics mean metrics of edge and core circuits, respectively.

This core and edge distinction helps define certain policies followed by Cisco Crosswork Planning tools. In Metric optimization, these are:

- Only core metrics are changed.
- The objective is to minimize utilization in the core. Utilization in the edge is not taken into account.
- The routes selected by the metrics are restricted by the tool so that any route between a source and a destination will enter and leave the core at most once. That is, *edge leakage* of routes is prevented.

A consequence of these policies is that larger edge metrics in the plan file to be optimized may result in better optimized core routes. With larger given edge metrics, the Metric optimization tool has more flexibility in how it can set the core metrics without leakage of routes into the edge.

## Perform metric optimization

Describes how to optimize network metrics and balance performance under normal and failure scenarios using the Metric optimization tool.

There are many trade-offs in metric optimization. The main trade-off is between performance under normal operation versus performance under failure. Metrics can be selected to do well under the former while ignoring the latter, or to try to balance performance between the two.

Metric optimization optimizes undifferentiated traffic, but does not optimize traffic per service class or per interface queue. If there are policies set on the undifferentiated traffic, Cisco Crosswork Planning tries to minimize both the interface utilization and the policy violations.

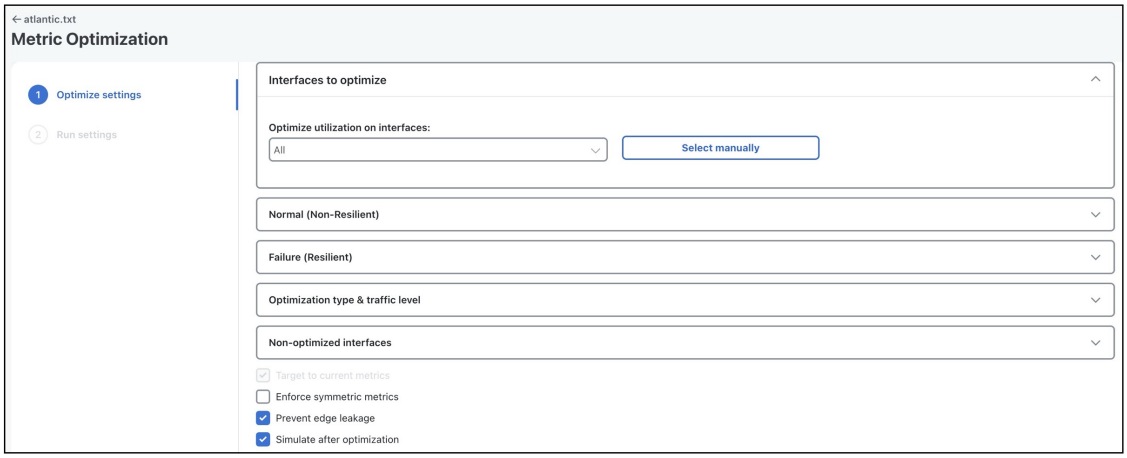
Complete these steps to run the Metric optimization tool.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). The plan file opens in the **Network Design** page.

   From the toolbar, choose **Actions** > **Tools** > **Metric optimization**.

   **Figure 61: Metric optimization options**




2. In the **Interfaces to optimize** panel, select the required interfaces to optimize. Use the drop-down list to select all or the Core interfaces, or the **Select manually** button to select the required interfaces manually.
3. Expand the **Normal (Non-Resilient)** section and choose the optimization options to use. Refer to Metric optimization options on page 228 for field descriptions.
4. Expand the **Failure (Resilient)** section and decide on the options to use. Refer to Metric optimization options on page 228 for field descriptions.

   For planning purposes, the most benefit from metric optimization is that it is able to optimize routing over a large number of failure scenarios. This capability controls that optimization, as well as controls the trade-off between normal and failure optimizations.

5. Specify the remaining options to fine-tune the optimization. Click **Next**.

   For detailed field descriptions, refer to Metric optimization options on page 228.

6. (Optional) Specify the maximum number of threads. By default, the optimizer tries to set this value to the optimal number of threads based on the available cores.
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:

   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.

- Schedule: Set the time at which you want to run the tool.

The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

> **Note**
>
> Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:

   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
   - If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

9. Submit your changes.

---

The Metric optimization tool executes with your selected preferences, optimizing network metrics for both regular operation and failure scenarios.

## Metric optimization options

Lists the available metric optimization options and their configurable fields to help you customize network optimization behaviors.

This table lists the available metric optimization options and their descriptions.

**Table 26: Metric optimization options**

| Field | Description |
|---|---|
| **Interfaces to optimize** | |
| Optimize utilization on interfaces | By default, Cisco Crosswork Planning optimizes all interfaces. You can restrict the optimized interfaces using this option. Choose the **Optimize utilization on interfaces** value as Core, or select the interfaces manually using **Select manually**. |
| **Normal (Non-Resilient) options** | |
| Minimize maximum interface utilization | This default is usually the primary objective, and concentrates on reducing the utilization of the interface with the highest utilization. |

| Field | Description |
|---|---|
| Minimize # of interfaces with utilization > __ % | Often a small number of interfaces are bottlenecks in the network, and their utilization cannot be reduced through rerouting. However, it is still desirable to reduce the utilization of other, under-utilized interfaces. This option tries to reduce all interface utilization to below a specified utilization level, or, if this is not possible, to minimize the number of interfaces over this level. |
| Minimize average latency | If after the optimizations flexibility is available in the choice of metrics, Cisco Crosswork Planning uses this flexibility to minimize the average latency of the routes across the network. This is not user-configurable and is always on. |
| Enforce latency bounds | If checked, the choice of paths for each demand is restricted, if possible, to those demands with latency below the latency bound. |
| **Failure (Resilient) options** | |
| Minimize maximum interface utilization | If checked, Cisco Crosswork Planning selects metrics to minimize the maximum interface utilization over all interfaces and all failure scenarios. If unchecked, only normal operation is optimized, and the remainder of the failure section is ignored. |
| Minimize # of interfaces with utilization > __ % | The number of interfaces with utilization (under any failure scenario) above a specified percentage is minimized. |
| Enforce no-failure utilization bound | If unchecked, maximum failure utilization is minimized with no consideration for utilization under normal operation. It is often convenient to check this option and place a bound on the utilization under normal operation, which Metric Optimization then satisfies, if possible. |
| Failure sets | This option determines the failure scenarios over which optimization is performed. You can select any combination of failures: All unprotected circuits, SRLGs, nodes, sites, ports, port circuits, parallel circuits, and external endpoint members. |
| **Optimization type & traffic level options** | |
| Optimization type | Cisco Crosswork Planning can perform global or incremental metric optimizations.<br>• Global: Creates optimal routes from scratch, and targets the old metrics afterward. Because Cisco Crosswork Planning can choose routes globally, this option tends to perform better under some circumstances, particularly when the target routes are poorly chosen. However, global optimization typically takes longer than incremental optimization.<br>• Incremental: Uses the current routes as a basis from which improvements can be made. The target routes are not changed as radically as when using global optimizations, and therefore the number of metric changes from the metric targets is typically much smaller when using the incremental option. |
| Set metrics on interfaces | By default, Cisco Crosswork Planning sets or changes metrics on any of the core interfaces in the network to achieve its optimization goals. Alternatively, as provided by this option, Metric optimization restricts metric changes to the these interfaces: all, core, interfaces selected manually. |
| Traffic level | The traffic level over which the optimization is performed. |

| Field | Description |
|---|---|
| **Non-optimized Interfaces options** | |
| Non-optimized Interfaces | For interfaces that are not optimized, you can set them as:<br>• Ignore<br>• Minimize number of interfaces with util > __ %<br>• Minimize number of interfaces with util > current + __ % |
| **Other options** | |
| Target to current metrics | There are many sets of metrics that will result in equivalent sets of routes. (Doubling all metrics, for example, does not change routes.) Cisco Crosswork Planning selects between these sets of metrics by metric targeting. Metrics can be selected to match as closely as possible the current metrics in the network. This simplifies the changeover from one set of metrics to the next. |
| Enforce symmetric metrics | This option sets the two metrics on each interface of each circuit equal to one another. This constraint can reduce optimization performance. However, it ensures that demands from A (source) to Z (destination) and from Z to A (in the core) take the same path, and so the chances of a session between A and Z being interrupted by a failure is minimized. |
| Prevent edge leakage | This default option restricts how Cisco Crosswork Planning can change demand routes. If a demand starts by routing only through the core, or from the edge through the core to the edge, this option prevents the routing from being changed so that it routes through the edge in the middle of its path. This is usually a desired routing policy restriction. |
| Simulate after optimization | This default option performs a simulation analysis after optimizing metrics. The simulation analysis allows a detailed view, for example, of the utilization under failure resulting from the new metric settings. |

## Tactical metric optimization

Explains tactical metric optimization and its role as a streamlined version of broader metric optimization tools.

The **Tactical metric optimization** is a tool in Cisco Crosswork Planning that

- is a reduced version of the Metric optimization tool
- minimizes the number of changes and runs faster, and
- offers a subset of Metric optimization tool options.

The tool tries to find the most optimal solution it can. In general, the longer it runs, the better the result. The **Time limit** option tells the tool to stop looking for better solutions after the specified amount of time and use the best solution it has found so far.

To access this tool, choose **Actions** > **Tools** > **Tactical metric optimization** from the toolbar.

# Metric optimization reports

Lists the contents and attributes available in Metric optimization reports, including report options, warning messages, routing statistics, and metric details.

Each time Metric optimization is run, a Metric-Opt report is automatically generated. You can access this information at any time by choosing **Actions > Reports > Generated reports** and then clicking the **Design History** link in the right panel.

The Metric-Opt report contains these sections:

- **Options**: Summarizes the options used to call Metric optimization. Some of these options (for example, the Output File and Report File) are only relevant when using the CLI tool.
- **Number of demands**: Total number of demands in the plan.
- **Objectives**: Summarizes the optimization objectives in order of priority.

### Warnings

Warnings detected by Metric optimization might lead to misleading or undesirable results. These are some of the warning message examples:

- **General warnings**
  - No Improvement in Optimization

    **No routing improvement found: metrics unchanged**

    Metric optimization was unable to improve on the metrics in the network. If the incremental option was specified and latency bounds were enforced, this could be because the target metrics violated latency bounds in ways that incremental optimization could not fix.
  - Optimization Exit Diagnostics

    **Optimization performance may be limited by low edge metrics**

    During optimization, some desirable routes could not be selected because the given edge metrics would have caused edge leakage.
  - Optimization Constraints Violated

    **<n> \{intra. inter\site core metrics exceed \crlf \{intra, inter\site-metric \{upper, lower\-bound**

    One or more intrasite or intersite core metric bounds have been violated.
  - Maximum normal utilization exceeds specified bound

    The original plan had normal utilization above the specified bound, and Metric optimization was unable to reduce the normal utilization below this bound. If this occurs, Metric optimization tries to set normal utilization as low as possible. No worst-case optimization is performed because reducing utilization under the normal scenario is the highest priority.

- **Constraint and routing warnings**:
  - **<n> demands with non-zero bandwidth exceed latency bounds**

    Because enforcing latency bounds is the highest optimization priority, this situation should not occur unless a latency bound is less than the shortest possible latency for a demand.
  - **<n> demands with zero bandwidth exceed latency bounds**

    Metric optimization does not always enforce latency bounds for zero bandwidth demands. These warnings can occasionally occur even if it is possible to find lower latency paths for the signaled demands.

- Unrouted Demands

  **<n> unroutable demands under normal operation**

  There are <n> demands for which no route exists between source and destination, even under the No Failure scenario.

- **<n> unroutable demands under <m> circuit failure scenarios**

  There are <n> demands for which no route exists between source and destination under <m> circuit failure scenarios. (Different demands might be unroutable under different circuit failures.)

- **<n> unroutable demands under <m> SRLG failure scenarios**

  There are <n> demands for which no route exists between source and destination under <m> SRLG failure scenarios. (Different demands might be unroutable under different SRLG failures.)

- **<n> unroutable demands under <m> source/dest node failure scenarios**

  There are <n> demands whose source and/or destination nodes fail under <m> failure scenarios. These demands clearly cannot be routed in these circumstances. (Different demands might be unroutable under different node failures.)

- **<n> unroutable demands under <m> non-source/dest node failure scenarios**

  There are <n> demands for which no route exists between source and destination under <m> node failure scenarios, where the nodes that fail are intermediate nodes in the path of the demands, and are not source/destination nodes of the demands. (Different demands might be unrouted under different node failures.)

### Routing summary

Summary statistics are displayed of the routes in network, both before (in parentheses) and after the optimization.

- Core/Edge Max Utilization: The maximum percentage utilization over any interface in the core or edge network. The normal utilization is under the normal (no-failure) scenario, and the worst-case utilization is maximized over all failure scenarios in the failure set.
- Latency: Median and average latencies over all routed demands, both in milliseconds and as a percentage of the latency of the smallest latency route possible.

  The latency of each demand's routing is calculated as a percentage of the latency of the shortest possible route, by latency, for that demand. Under the percentage column, the median and average of these percentages are displayed.
- Num of demands routed away from shortest path: The number of routes (out of all the demands routed) that do not follow the shortest path, by latency. This statistic is an indicator of how much the routes have been affected by utilization minimization as the primary criterion for optimization.
- Num of demand routes exceeding latency bounds: The number of routes (out of all the demands routed) that exceed the latency bounds.

### Metrics

The metric targeting choice is shown, together with the number of metrics different from the target metrics. Lists of interfaces for which metrics differ from the target metrics are displayed, together with both the target metric and the optimized metric. The metrics are separated into those on intrasite and intersite interfaces.

# 17 Configure MPLS Routing

**Topics:**


- MPLS routing
- Supported LSP types
- Create and visualize LSPs and LSP paths
- LSP paths
- Route demands through LSPs
- Global simulation parameters
- LSP simulation diagnostics


Explains MPLS routing concepts, supported LSP types, LSP creation and path configuration, and troubleshooting procedures essential for deploying and maintaining MPLS networks.

## MPLS routing

Explains how MPLS routings handle LSPs, including their establishment, rerouting procedures, and demand management for effective traffic engineering.

All LSPs other than SR (segment routed) LSPs are routed like RSVP LSPs. For MPLS simulation information specific to these types of LSPs, refer to Configure RSVP-TE Routing on page 263 and Configure Segment Routing on page 337.

- LSPs are established under normal operation: that is, with failures not taken into account.
- Any LSPs that are affected by the failures are rerouted. Depending on the LSP path settings, reroutes might involve moving to a secondary path, dynamically rerouting the LSPs, or rerouting based on a segment list.
- Demands are routed using the established LSPs using the specified IGP protocols given the specified failure scenarios.
- LSP utilizations are calculated from the demand traffic using the specified traffic level.

## Supported LSP types

Describes the LSP types supported by Cisco Crosswork Planning to facilitate path management. These types enable users to identify and configure specific routing mechanisms within the network infrastructure.

Cisco Crosswork Planning supports these LSP types:

- SR LSPs: Segment Routing LSPs, which do not use RSVP for routing. You can create SR LSPs using the Cisco Crosswork Planning UI. They are identified with a **Type** property of **SR**. For more information, refer to Configure Segment Routing on page 337.
- RSVP LSPs: LSPs established through RSVP. These are commonly known as MPLS TE tunnels. Cisco Crosswork Planning discovers RSVP LSPs, and you can also create them using the Cisco Crosswork Planning UI. They are identified with a **Type** property of **RSVP**. For more information, refer to Configure RSVP-TE Routing on page 263.

> **Note**
>
> Cisco Crosswork Planning does not model LDP tunnels as LSPs.

## Create and visualize LSPs and LSP paths

Describes how to create and visualize MPLS LSPs within the network design interface.

To simulate MPLS LSPs in the network, you must first set up the LSPs to be routed. Use this task to create and visualize LSPs from the UI.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36 ). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP**.

   Alternatively, in the Network Summary panel on the right side, click  > **LSPs** in the **LSPs** tab.

The LSPs tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **LSPs** check box.

3. Choose the **Type**, which determines whether this is an RSVP LSP or an SR LSP.
4. In the **Name** field, enter the name of the LSP.
5. Select the **Active** or **FRR enabled** check box.
6. Expand the **Source & destination** panel. Choose the appropriate source and destination site and node details.
7. (Optional) Expand the other panels (Routing, CSPF, and Other) and enter the relevant parameters.
8. Save your changes.
9. To visualize the LSP in the network plot, select the LSP from the LSPs table. They appear in the plot as purple arrows.

**Figure 62: LSP visualization**




- To add a mesh of LSPs between the selected nodes, use any of these options:
  - From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP mesh**.
  - In the Network Summary panel on the right side, click + > **LSP mesh** in the **LSPs** tab.
- To filter to related information such as related interfaces, source and destination nodes, demands, or LSP paths, select the LSP, click , and choose the appropriate option.

## LSP paths

Describes the configuration and visualization properties of LSP paths within a network design.

LSPs can be assigned one or more LSP paths. Like LSPs, LSP paths have properties that vary depending on whether the path is for an RSVP LSP or SR LSP. If these properties are omitted, then they are inherited from the LSP. If these properties are set in the LSP path, they override the LSP settings. For information on these properties, refer to Configure RSVP-TE Routing on page 263 and Configure Segment Routing on page 337.

### Path options and active path

Describes the configuration of LSP path options and the selection of active paths to manage network routing behavior.

A path option is a configuration parameter that

- assigns a numeric value to each LSP path
- determines the sequence in which paths are established, and
- enables prioritization of network routes based on operational requirements.

Each LSP path includes a **Path option** property, which you can set when adding or editing an LSP path. The system attempts to establish LSP paths in order of increasing path option value: path option 1 is tried first, followed by path option 2, and so on. The LSP is routed using the first path that can be successfully established.

**Figure 63: Path option**




Alternatively, you can specify which LSP path to use by setting the **Active path** field under the **Routing** section (Add/Edit LSP page). Selecting an active path allows you to override the default path option order and explicitly control which path is used for routing.

## Create LSP paths

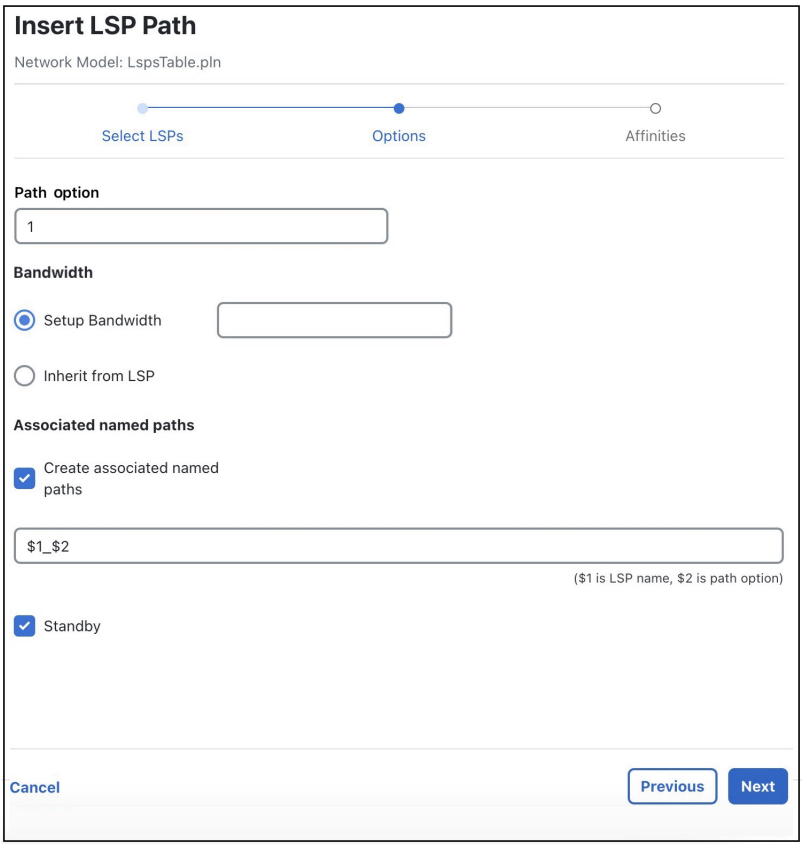
Outlines the steps to configure LSP paths within a network plan file to define activation order and path properties.

Use this task to create LSP paths.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the **Network Summary** panel on the right side, from the **LSPs** table, choose the LSPs to which you are adding LSP paths.
3. From the toolbar, choose **Actions > Insert > LSPs > LSP Paths**, or click + in the **LSP Paths** table.

   The Insert LSP Path window appears.
4. If you are satisfied with the LSPs that you selected in Step 2, click **Next**. Make changes to the selection, if required.
5. In the **Path option** field, enter the order in which the LSP path is activated. Lower numbers have preference over higher numbers.



6. If you are creating an LSP path from SR LSPs, proceed to Step 7. If this is an RSVP LSP path, optionally set these properties. For information on these properties, refer to Configure RSVP-TE Routing on page 263.
   a) To set the bandwidth, specify the **Setup Bandwidth** for the LSP path, or specify that it should inherit the bandwidth from the LSP.
   b) To create associated named paths, check this option and complete the name using $1 for the LSP name and $2 for the path option.
   c) If this is a standby LSP path, check **Standby**, which ensures the paths are always active.
7. Click **Next**.
8. Make the required changes in the **Affinities** page to associate the LSP path with affinities.
   a. For each rule (Include, Include any, Exclude), choose whether the LSP path should inherit the LSP affinity rule or whether it should be based on rules defined in the table.
   b. If you chose to use the table, choose the rule for each affinity you are associating with the LSP path.
9. Submit your changes.

---

### Path latency calculations

Provides the rules used to calculate the shortest latency paths for LSPs and demands in Cisco Crosswork Planning.

Cisco Crosswork Planning calculates shortest latency paths for LSPs and demands using these rules:

- By default, Cisco Crosswork Planning uses the shortest path from the source to the destination of the LSP or demand, where the weight on each interface is the delay value for that interface.
- If the delay for any interface is zero (as it is by default), a small (0.00001 ms) delay is used instead. This means that if all the delays for the interfaces in the plan are zero, the path with the smallest number of hops is selected.
- If two paths have the same latency, paths with the smaller number of hops are preferred.

## Route demands through LSPs

Explains how to route demands through intra-area, inter-area, and specific LSPs, including demand routing order, explicit versus dynamic inter-area routing, creating private demands and LSPs, deleting demands, and configuring load sharing.

### Routing through intra-area LSPs

Introduces intra-area LSP routing as IGP shortcuts, explains metric options for autoroute and forwarding adjacency LSPs, and outlines how to configure routing settings and metrics for these LSPs.

Intra-area LSPs are modeled as IGP shortcuts. That is, the source of the demand using the LSP can be a node other than the ingress node into the IGP, and the destination node of the LSP can be a node other than the egress node from the IGP. A demand through intra-area LSPs does not require the LSP to traverse the full demand path.

Each intra-area LSP has a metric that helps determine which traffic is routed through the LSP. By default, autoroute LSPs have a metric equal to the shortest IGP distance from the source to the destination. However, you can configure static metrics for these LSPs that override this default. Static metrics can also be defined as relative to the shortest IGP distance, but these relative metrics are not currently supported in Cisco Crosswork Planning.

#### Configuration options

Forwarding adjacency (FA) LSPs always have a metric. If not specified, the default is 10. Forwarding adjacency LSP metrics are injected into the IGP so that nodes other than the source node of the LSP are aware of the path length through the forwarding adjacency LSP, and use it in their shortest path calculations.

To edit Autoroute and Forwarding adjacency settings, choose one or more LSPs from the LSP table, click  to open the **Edit** page, and edit the values in the **Routing** section.

## Routing demands through inter-area LSPs

Outlines the routing of demands through inter-area LSPs, including endpoint requirements, routing protocols, constraints for LSP types, limitations on fast reroute and IGP shortcuts, and necessary privacy conditions for successful traffic routing.

Since it is not possible to define metrics distances across ISP areas, there can be no well-defined metric for inter-area LSPs. Therefore, in Cisco Crosswork Planning, a demand only routes through an inter-area LSP if the demand's endpoints are nodes matching the source and destination of the LSP, interfaces on these nodes, or external ASes whose ingress and egress points through the IGP are these nodes. This requirement is true regardless of the network option settings.

### Routing inter-area LSPs

In Cisco Crosswork Planning, all nodes in a single AS are assumed to belong to a single IGP. If the plan file contains more than one AS, all IGPs defined in these ASes are of the same type. For information on how nodes are assigned to areas or levels, refer to Simulate Traffic Flow from Source to Destination Using Demands on page 95.

An inter-area LSP is an LSP whose source node and destination node have no areas in common. If available, inter-area LSPs follow actual paths, regardless of whether doing so violates the required order of routing through areas. For example, if following an actual path, an inter-area can enter and leave an OSPF area 0 more than once.

To understand the protocols that determine routing order when setting up inter-area LSPs, refer to Order of routing through areas on page 238. Other factors that determine how the inter-area LSPs are routed include whether the LSP type is RSVP or SR, and whether you have selected to require explicit hops at ABRs. For more information, refer to Explicit and dynamic inter-area LSP routing on page 239.

**Note**

- Inter-area Fast Reroute LSPs and inter-area IGP shortcut LSPs are not supported. If the source and destination nodes of a Fast Route LSP are in different areas, the LSP is not routed.
- To successfully route traffic, demands must also match the privacy requirements. Autoroute, Forwarding adjacency (FA) properties, and LSP metrics are ignored. For more information on privacy, refer to Routing demands through private LSPs on page 239.

### Order of routing through areas

Explains how routing protocols determine the order of traversing backbone and non-backbone areas when setting up inter-area LSPs.

Regardless of the LSP type and whether explicit hops are required, inter-area LSPs route through the backbone areas, where "backbone" means area 0 for OSPF or the Level 2 area for IS-IS.

- If there are three or more areas, backbone areas must be between the source and destination nodes. Typically, there are no more than three areas and therefore, the backbone area must be in the middle. For example, an OSPF inter-area LSP would route from area 1 to area 0 (backbone) and then to area 2.
- If there are only two areas, there can only be one backbone area, and either the source or the destination node must be in the backbone area.

#### Explicit and dynamic inter-area LSP routing

Explains the two approaches to routing inter-area LSPs: explicit hop-based and dynamic methods, and describes their simulation and planning benefits.

An OSPF ABR is a node that belongs to both area 0 and other OSPF areas. An IS-IS ABR is a node that belongs to both the Level 2 area and another IS-IS level.

There are two modes of routing inter-area LSPs:

- Explicit hop mode: Requires explicit hops to be set on ABR nodes. This approach simulates actual router behavior where ABR explicit hops are required.
- Dynamic mode: Does not require explicit hops at ABR nodes and allows routing an LSP dynamically across multiple areas. Though this does not simulate actual router behavior, it is useful for planning inter-area LSP routes.

#### Configuration option

These modes are specified using the network option **LSP routing requires ABR explicit hops** in the **Label Switched Paths** section. If this option is selected, inter-area LSPs are routed based on explicit hops set on the ABR nodes.

- An inter-area RSVP LSP must contain a named path, and the named path must contain explicit hops at ABRs for each required area crossing.
- An inter-area SR LSP must contain a segment list, and the segment list must contain explicit node hops at ABRs for each required area crossing.

If this option is not selected, inter-area LSPs are routed dynamically and explicit hops at ABRs are not required. To leave one area and enter another, the inter-area LSP routes to the closest ABR in the current area that also borders the area it is entering.

## Routing demands through private LSPs

Explains what private LSPs are, their key attributes, and how to configure and identify them.

A private LSP is a special LSP that

- carries traffic dedicated to specific demands only
- simulates the operation of MPLS Layer 2 VPNs, providing an exclusive route for the associated demands, and
- interrupts all traffic associated with the LSP if it fails.

#### Configuration and identification

To configure LSPs that simulate Layer 2 VPNs, use one of these two methods:

- Create dedicated demands for existing LSPs. The created demands will match the LSPs in source and destination. For more information, refer to Create private demands for existing LSPs on page 240.
- Create private LSPs from existing demands. The created LSPs will match the existing demands in source and destination. For more information, refer to Create private LSPs for demands on page 241.

System indicators for private LSPs include:

- LSPs table: The **Private** column is set to true.
- Demands table columns: The **Private LSP name** and **Private LSP source** columns are set to true.

### Create private demands for existing LSPs

Explains the steps to generate private demands for existing LSPs within the network model and secure traffic flow.

Use this task to assign specific traffic parameters to LSPs and mark the demands as private.

#### Before you begin

Ensure that LSPs currently exist in the network model.

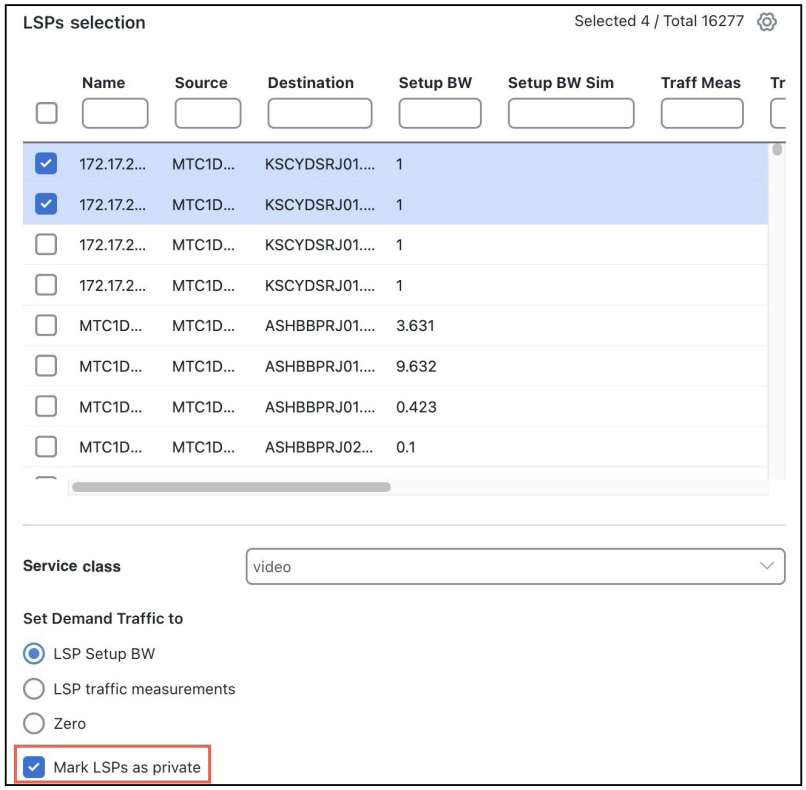
#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **LSPs** > **Demands for LSPs**.

   Alternatively, in the **Network Summary** panel on the right side, click + > **Demands for LSPs** in the **LSPs** tab.

   The **LSPs** tab may be available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **LSPs** check box.

   **Figure 64: Create demand for LSPs page**




3. From the LSPs list, select the LSPs for which you want to create demands.
4. Select the service class to which these LSPs belong.
5. Set the demand traffic to equal the LSP setup bandwidth, the LSP traffic measurements, or zero.
6. Check the **Mark LSPs as private** check box.
7. Submit your changes.

   The newly created demands are highlighted in the **Demands** table.

### Create private LSPs for demands

Outlines the steps to create private LSPs for existing demands within a plan file.

Use this task to add private LSPs to selected demands within your plan file to control bandwidth assignments and mark paths as private.

#### Before you begin

Ensure that demands exist in the plan file.

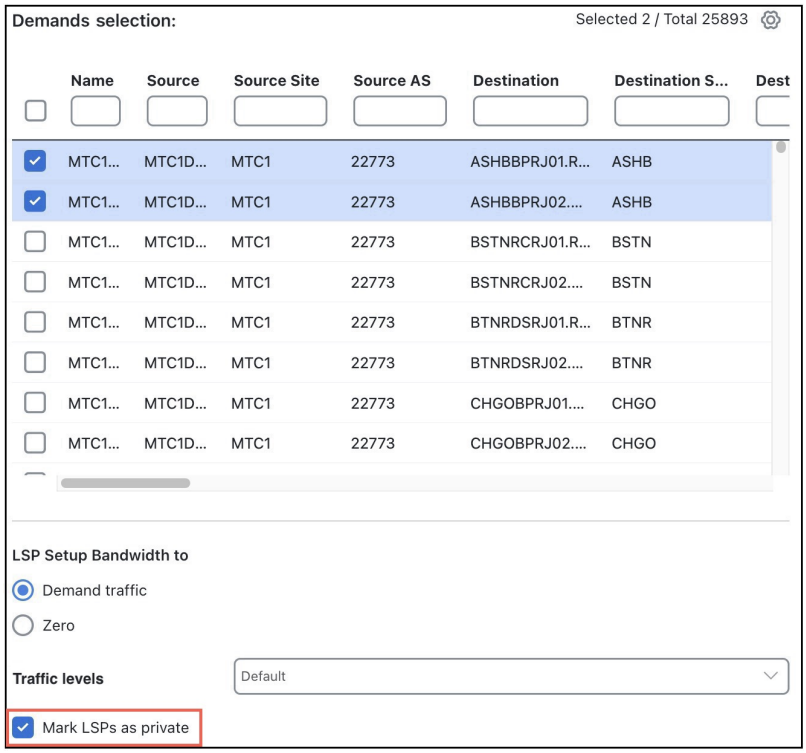
#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSPs for demands**.

   Alternatively, in the Network Summary panel on the right side, click + > **LSPs for demands** in the **LSPs** tab.

   The **LSPs** tab may be available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **LSPs** check box.

   **Figure 65: Demands selection page**




3. In the list of demands, select the demands for which you want to create LSPs.
4. Set the bandwidth traffic to a specific demand traffic or to zero.
5. Check the **Mark LSPs as private** check box.
6. Submit your changes.

The newly created LSPs are highlighted in the **LSPs** table.

### Configure automatic demand deletion for private LSPs

Describes the steps to configure settings so that when you delete a private LSP, the related demand is also deleted automatically.

Use this task to delete demands when associated private LSPs are removed from the plan file.

By default, when you delete a private LSP, the related demand remains in the system.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

   The Network options page opens.
3. Click the **Advanced** tab.
4. In the **Demands** section, ensure that these check boxes are selected.
   - **Unrouted if the private LSPs are unrouted**
   - **Removed if the private LSPs are removed**
5. Save your changes.

## Load sharing between LSPs

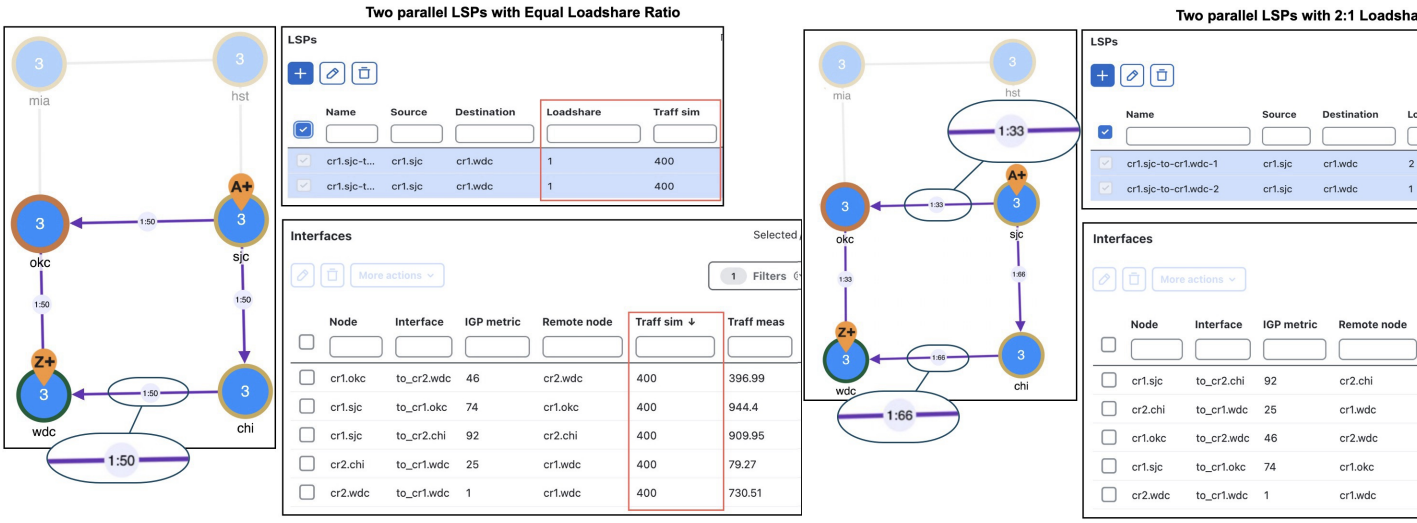
Describes how traffic is distributed across multiple LSPs with the same source and destination using the LSP Loadshare property. This mechanism allows for proportional traffic routing between parallel paths.

Two or more LSPs with the same source and destination (and metrics, if these are defined) loadshare traffic between them. How the load is shared between the LSPs is determined by the LSP Loadshare property. By default, LSPs have a Loadshare property of 1, and thus route traffic between them in equal proportions. Changing the Loadshare value changes the distribution of LSP traffic and interface traffic in proportion to these values.

### Loadshare ratio examples

If two LSPs are parallel, and one has a Loadshare property of 2 and one has a Loadshare property of 1, there will be a 2:1 ratio of traffic shared between them. The top half of Figure 67: Examples of two parallel LSPs with equal loadsharing and 2:1 loadsharing shows an example of two parallel LSPs that are routed using strict explicit paths. Each has a Loadshare value of 1, which means the traffic is routed using a 1:1 loadshare ratio so that each LSP carries 50% of the traffic. In contrast, the lower half shows the same parallel LSPs with a 2:1 ratio. That is, one LSP has a Loadshare value of 2, and one has a Loadshare property value of 1. The LSP with a Loadshare value of 2 carries 66% of the traffic, while the other carries 33%.

**Figure 66: Examples of two parallel LSPs with equal loadsharing and 2:1 loadsharing**




Each LSP has a path showing the percentage of traffic that crosses the LSP as a result of the Loadshare value for that LSP. The format is *<path option>:<loadshare percentage>*. Although loadshare applies only to path option 1, the LSP itself might be routed along a different path.

Note that this percentage may or may not be the same as the Loadshare value in the LSPs table. This is because the Loadshare value is relative to the other Loadshare values for parallel LSPs. For this reason, if manually editing the Loadshare property, we recommend that you consider the entire set of parallel LSPs. Note also that because an LSP can be in multiple sets of parallel LSPs, the percentage on its path can differ.

To optimize these Loadshare values, use the LSP Loadshare Optimization tool. For information, refer to LSP loadshare optimization on page 254.

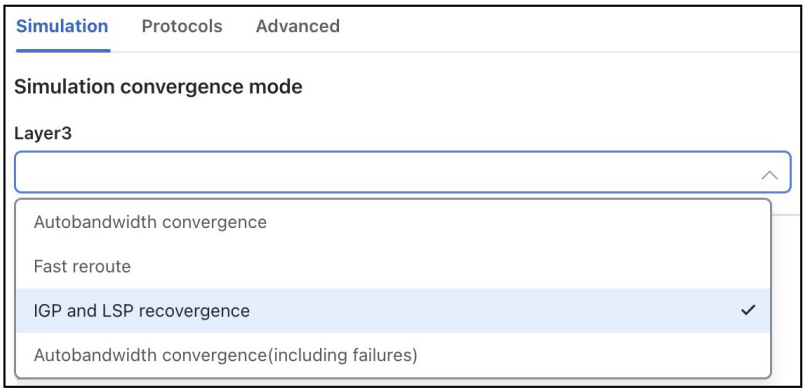
## Global simulation parameters

Describes how global simulation parameters define network behavior during failure scenarios and impact route and optimization simulations.

Cisco Crosswork Planning lets you set global parameters that affect how LSPs are routed or rerouted.

To access these options, click  or choose **Actions** > **Edit** > **Network options** in the toolbar, then click the **Simulation** tab.

**Figure 67: Global simulation parameters**




By default, Cisco Crosswork Planning simulates the state of the network once it has fully responded to a failure. Specifically, this is the network state after LSPs have re-established new routes around the failure, and the IGP has fully reconverged. This is called the **IGP and LSP reconvergence** simulation mode.

Other simulation modes include **Fast reroute** (FRR), **Autobandwidth convergence**, and **Autobandwidth convergence (including failures)**. For more information on these modes, refer to Configure RSVP-TE Routing on page 263.

#### Optimization tool behavior

The optimization tools only work in *IGP and LSP reconvergence* mode. If you attempt to run an optimization tool in a different convergence mode, the system prompts you to confirm; if you proceed, the simulation automatically switches to *IGP and LSP reconvergence* mode.

### Set LSP establishment order

Outlines the steps to modify the sequence in which LSPs are established by changing a random seed value.

Use this task for testing different routing outcomes based on varying establishment orders.

Cisco Crosswork Planning establishes LSPs in the order in which they appear in the plan. The routing of a specific LSP might depend on previously established LSP routes. You can modify this order by changing a random seed. Cisco Crosswork Planning then establishes LSPs in a random order that is determined by this number. Although you cannot predict the order based on the number, if you use the same number multiple times, Cisco Crosswork Planning establishes the LSPs in the same order each time. For example, varying the LSP establishment order lets you check whether certain orders result in higher utilizations.

#### Procedure

1. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

   The Network options page opens.
2. In the **Label switched paths** section under the **Simulation** tab, enter a number in the **LSP establishment order seed** field. The default is 0.
3. Save your changes.

## LSP simulation diagnostics

Provides details about the LSP simulation diagnostic tool, including its features and behaviors for analyzing simulated routing and LSP path issues.

To help troubleshoot RSVP LSP and LSP path simulations, Cisco Crosswork Planning analyzes simulation routes and provides reasons for certain types of routing behavior. The simulation diagnostics tool identifies why an LSP is not routed, why an LSP is routed away from its actual path, and why an LSP is not following the shortest TE path.

The diagnostic reports assume that all other LSPs, except for the one being tested, have been routed. That is, Cisco Crosswork Planning calculates whether an LSP can route on the actual path after all other routed LSPs have had their bandwidth reserved.

Note that running a report overwrites the previous report of the same type. For example, an LSPs diagnostic report would overwrite the previous LSPs report, but not the LSP path diagnostic report.

### Run LSP simulation diagnostics

Outlines the steps to execute diagnostic simulations for LSPs or LSP paths to identify routing issues.

Use this task to run simulation diagnostics.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Tools** > **Diagnostics** > **LSP simulation** or **LSP path simulation**.
3. Select the LSPs or the LSP paths you want to optimize.
4. Click **Submit**.

   A diagnostic report is automatically generated. You can access this report at any time by choosing **Actions > Reports > Generated reports**, and then clicking the **LSP Routing Diagnostics** or **LSP Path Routing Diagnostics** link in the right panel.

## Simulation diagnostics for LSP routing failures

Provides details about the diagnostic information to assist in troubleshooting LSP and LSP path routing issues within the network simulation environment.

Simulation diagnostics provide reasons that explain why an LSP or LSP path fails to route during simulation. This information helps identify specific reasons for routing failures during simulation.

The table lists the reasons and a brief description to assist in troubleshooting:

**Table 27: LSP simulation diagnostic reasons**

| Reason | Description |
|---|---|
| Affinities | The affinity settings prevent the LSP or LSP path from routing. |
| Available BW | There is insufficient bandwidth to route the LSP or LSP path. |
| Explicit Hops | The route is determined by a named path that contains one or more explicit hops. |
| Hop Limit | The hop limit is too low. |
| Invalid Actual Path | The actual path is invalid and cannot be interpreted as an LSP route. |
| No Actual Path | The LSP does not have an actual path. |
| No Attempt | The LSP path is not routed because it is not a standby path and it is not the lowest routable path option for the LSP. |
| No Destination | There is no destination defined. The collected network might not contain the destination node of the LSP. |
| Simulation Option | The simulation option to use actual paths is not enabled. |
| TE not enabled | There are interfaces that the LSP or LSP path traverses that are not TE enabled. |
| Topology | The source and destination are disconnected. |

# 18 Optimize LSPs

**Topics:**


- LSP disjoint path optimization
- LSP loadshare optimization
- LSP setup bandwidth optimization


Describes tools and procedures for optimizing LSP operation, including disjoint path optimization, loadshare configuration, and setup bandwidth management.

# LSP disjoint path optimization

Explains how LSP disjoint path optimization enables the creation and optimization of disjoint LSP routes to ensure that LSPs do not share routing objects, such as interfaces or nodes.

A LSP disjoint path optimization is a tool in Cisco Crosswork Planning that

- creates disjoint LSP paths for RSVP LSPs and SR LSPs, and
- optimizes these paths based on user-specified constraints.

LSPs and LSP paths are *disjoint* when they do not route over common objects, such as interfaces, or nodes. Using disjoint LSPs ensures service continuity and high resiliency in the event of network failures. For example, this tool lets you route LSPs to have disjoint primary and secondary paths that are optimized to use the lowest delay metric possible.

If it is not possible to achieve the optimization as defined by the routing selection, path requirements, and constraints, the tool provides the best disjoint paths and optimization possible.

**Note**

- The LSP disjoint path optimization tool supports both Inter-Area and Inter-AS functionalities.
- The optimizer can be applied to both RSVP LSPs and SR LSPs. However, only one type of LSPs can be optimized at a time.

To access the LSP disjoint path optimization tool, from the toolbar, choose **Actions** > **Tools** > **LSP optimization** > **LSP disjoint path optimization**.

When the optimization completes, the system tags the LSPs with *DSJOpt* and generates a report containing the optimization results.

## Disjoint path optimization inputs

Explains optimization input options, including disjoint routing selection behaviors, routing constraints, and available configuration parameters for rerouting LSP paths based on explicit hops, segment lists, and optimization requirements.

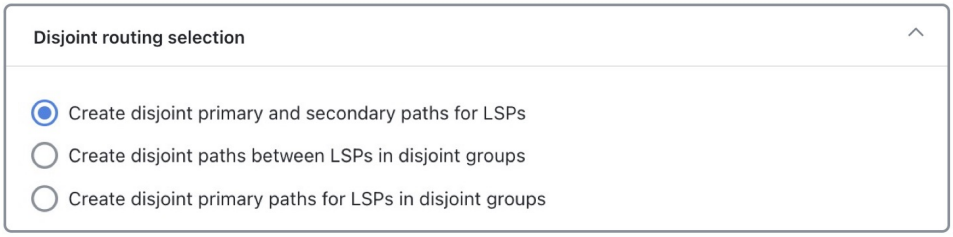
### Disjoint routing selection options

Lists and describes the behaviors and configurable options for disjoint routing selection in Cisco Crosswork Planning.

The **Disjoint routing selection** section provides several options for configuring how LSP routes are managed in the system.

#### Configurable disjoint routing options

**Figure 68: Disjoint routing selection options**



The routing options include:

- **Create disjoint primary and secondary paths for LSPs**: For all LSPs, whether they are in a disjoint group or not, routes all LSP paths so that they are disjoint from all other paths belonging to that LSP. This disjointness extends beyond primary and secondary paths to include all other path options (for example, tertiary).
- **Create disjoint paths between LSPs in disjoint groups**: For all LSPs that are in disjoint groups, routes all LSP paths so that they are disjoint from all other paths belonging to LSPs in that disjoint group. This disjointness extends beyond primary and secondary paths to include all other path options (for example, tertiary).

  For example, all LSPs in disjoint group East are rerouted to be disjoint from each other. All LSPs in disjoint group Southeast are rerouted to be disjoint from each other. However, LSP paths in the East group are not rerouted to be disjoint from those in the Southeast group.
- **Create disjoint primary paths for LSPs in disjoint groups**: For all LSPs that are in disjoint groups, reroutes only their primary paths so that they are disjoint from each other.

#### Additional behaviors in Cisco Crosswork Planning

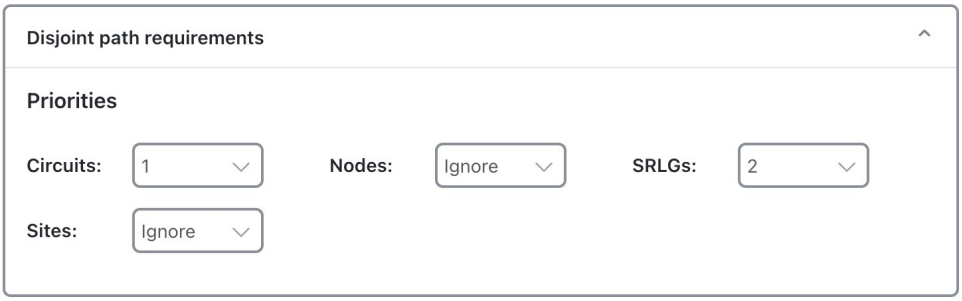
- No new LSP paths are created. Only existing LSP paths are rerouted.
- Explicit hops are modified or created for RSVP LSPs.
- Segment lists are created only for LSP paths, and only LSP path segment lists are updated. If a segment list is associated with an LSP (rather than an LSP path), that LSP segment list is removed.

### Disjoint path requirements

Lists the requirements and priority system used for achieving disjoint paths in the network to ensure optimal separation and resilience.

The disjoint path requirements identify the priority for creating disjointness across a path. Disjointness priorities **1**, **2**, **3**, and **Ignore** are available for circuits, SRLGs, nodes, and sites. The tool tries to create disjointness for all objects that have a priority set other than **Ignore**. If full disjointness cannot be achieved, the tool prioritizes disjointness based on these values.

**Figure 69: Disjoint path requirements**




For example, in Figure 69: Disjoint path requirements on page 249, Circuits have a priority of 1, SRLGs have a priority of 2, and the other objects are ignored. If the tool cannot achieve full disjointness across both circuits and SRLGs, it prioritizes the disjointness of circuits over SRLGs.

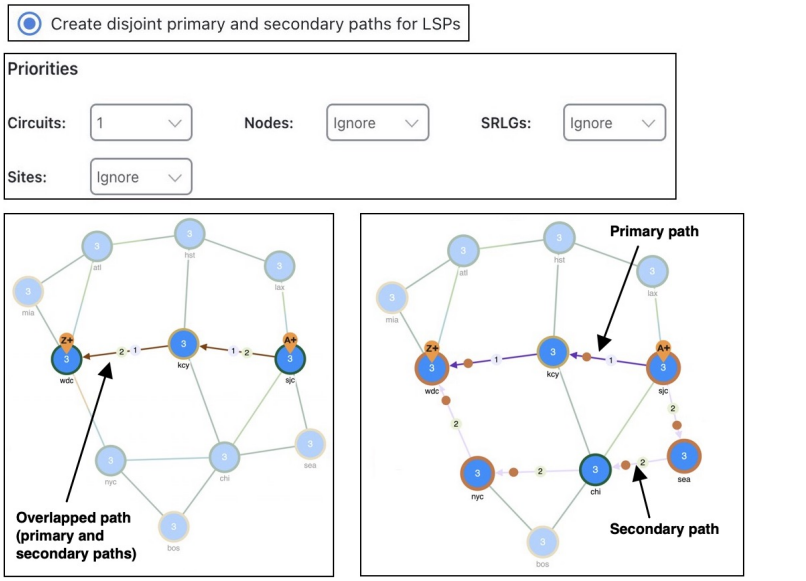
### RSVP LSP path comparison example

Describes the key attributes and differences between primary and secondary RSVP LSP paths when disjoint circuit requirements are applied.

This example compares the characteristics of primary and secondary RSVP LSP paths based on disjoint circuit settings. The LSP has a primary and secondary LSP paths that use the same route from cr2.sjc to cr2.wdc. The LSP is not a member of a disjoint group.

The Figure 70: Primary and secondary paths based on disjoint circuit requirements on page 250 shows the explicit interface hops set for the secondary LSP path that force it to take a different route from cr2.sjc to cr2.wdc. As the disjoint path requirement is only circuits, the primary and secondary paths route across different circuits, as indicated in the resulting report.

**Figure 70: Primary and secondary paths based on disjoint circuit requirements**



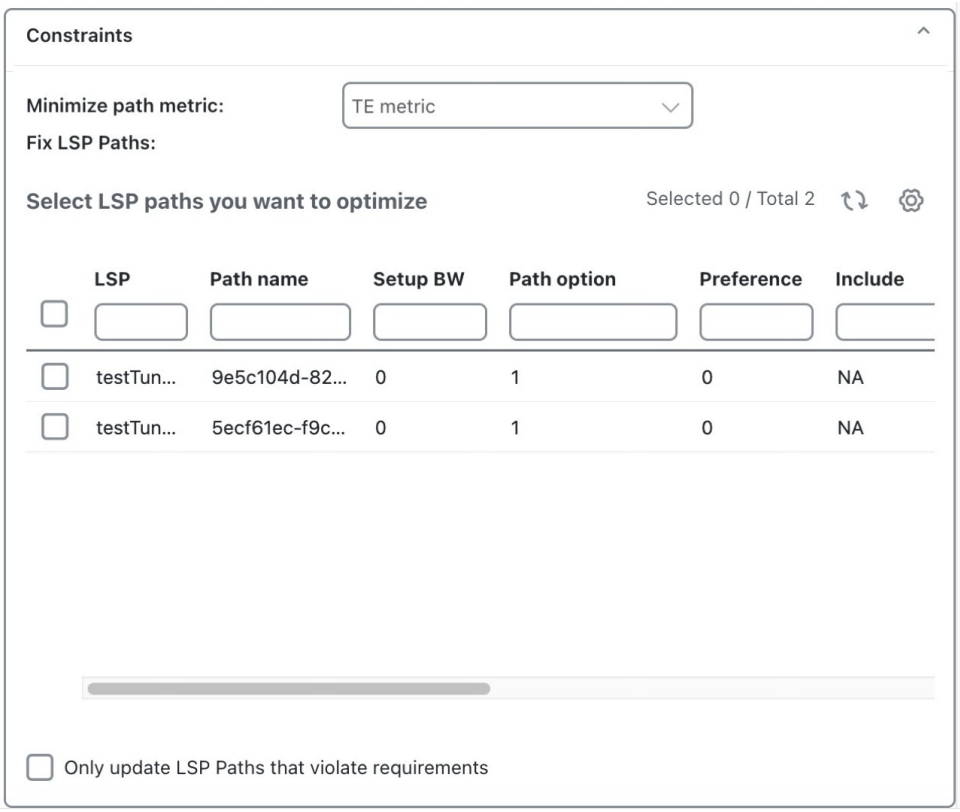

| LSPNAME | LSPSOURCE | COMMONCIRCUITSBEFORE | COMMONCIRCUITSAFTER |
|---|---|---|---|
| lsp_cr2.sjc | cr2.sjc | 4 | 0 |

### Path constraint options

Lists the available constraint options for path optimization and management, including recommendations for their effective use.

The **Constraints** section provides the these options for path management:

- **Minimize path metric**: Paths are optimized to minimize the sum of the metrics along the path with respect to delay, TE, or IGP metrics. These properties are configurable from the Add/Edit Interface page, and delays can also be set in the Add/Edit Circuit page.
- **Fix LSP Paths**: Selected or tagged LSP paths are not rerouted. This constraint is useful when you have previously optimized specific LSPs within the network and want to maintain their routes.
- **Only update LSP Paths that violate requirements**: Paths are modified only if they violate the requirements specified in the area of the Disjoint path requirements on page 249 page.




**Disjoint route optimization for SR LSPs**

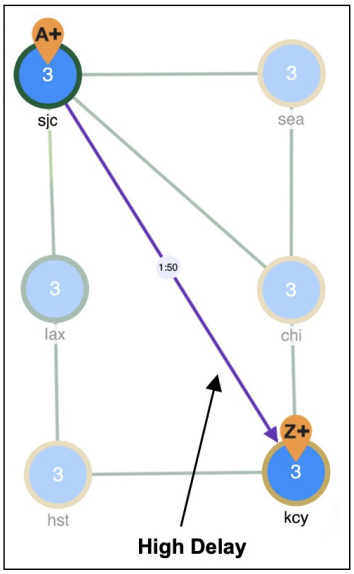
Provides an example showing how disjoint route optimization can improve path selection for two SR LSPs sharing the same source and destination.

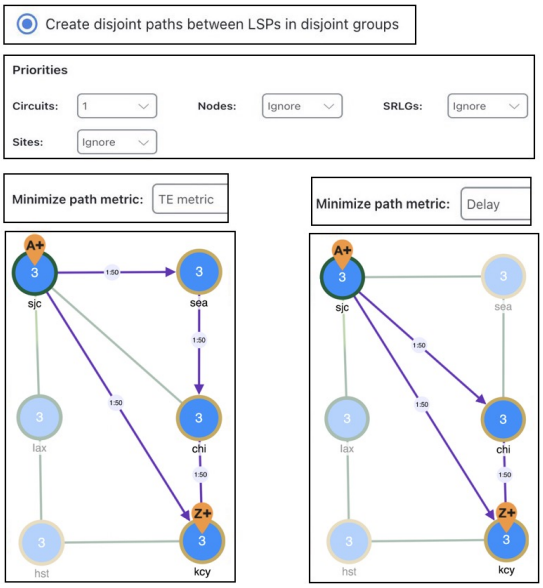
This example shows how disjoint routes can be optimized for two SR LSPs using the same source (sjc) and destination (kcy).

In this example:

- Both LSPs belong to the same disjoint group, and both LSPs have an LSP path.
- The circuit between sjc and kcy has significantly higher delay than the other circuits.
- The only disjoint path requirement selected is circuits.
- Figure 71: Routing of SR LSPs on page 251 shows this:
  - Before using the LSP Disjoint Path Optimization tool, both LSP paths use the same route. There are no segment list hops.
  - When you select the option to create disjoint paths between LSPs in the same disjoint group and use the TE metric for the shortest path calculation, two disjoint LSP path routes are created. Both routes include a segment list node hop (as indicated by the orange circle around the node) on the destination node. One route also includes a segment list node hop on sea, which ensures a different route is used.
  - When you select the same disjoint option and use the Delay metric for the shortest path calculation, one LSP is moved away from the high-delay sjc-kcy circuit. Traversing that circuit does not provide the shortest latency path.

**Figure 71: Routing of SR LSPs**




### Create disjoint groups for LSPs

Explains the steps to configure disjoint groups to enhance LSP redundancy and optimize path selection.

Use this task to assign LSPs to disjoint groups, to define separate routing options and priorities.

Disjoint groups are used to ensure that LSPs within the same group choose alternate routes. This setup minimizes path overlap, so each LSP uses a different route based on assigned priorities.

#### Before you begin

- The network model must already contain the primary LSP paths. If using the option to create disjoint primary and secondary paths, it must also at least contain secondary LSP paths, though it can contain other path options, such as tertiary.
- If you are creating disjoint paths between LSPs in the same disjoint group, you must first add the LSPs to the disjoint groups.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the Network Summary panel on the right side, choose one or more LSPs from the LSPs table.
3. Click .
4. Click the **Advanced** tab.
5. Expand the **Explicit route selection** section and enter the Disjoint group name.
6. Assign priorities to LSPs within these groups, if needed. Higher priority LSPs are assigned shorter routes based on the selected metric. The higher the number, the lower the priority.

   **Example**
   There are two LSPs in the same disjoint group, each with a different disjoint priority. Run the LSP disjoint path optimization tool to create disjoint paths for LSPs in the same disjoint group using TE metrics as a constraint. The LSP with a disjoint priority of 1 routes using the lowest TE metric, and the LSP with a disjoint priority of 2 routes using the second lowest TE metric.
7. Save your changes.

### Run the LSP disjoint path optimization tool

Explains the steps to configure and run the LSP disjoint path optimization tool to optimize disjoint paths for selected LSPs in your network design.

Use this task to run the LSP disjoint path optimization tool to ensure that your LSPs use disjoint paths. This reduces the risk from single points of failure and enforces compliance with network path requirements.

#### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.

   From the toolbar, choose **Actions** > **Tools** > **LSP optimization** > **LSP disjoint path optimization**.

   Alternatively, choose **Preset workflows** > **Perform optimization**, select **LSP Optimization** as the optimization type, select **LSP disjoint path optimization** from the drop-down list, and click **Launch**.

2. Select the LSPs for which you want to optimize the disjoint paths.
3. In the **Disjoint routing selection** section, select how to route the disjoint paths. For more information, refer to Disjoint routing selection options on page 248.
4. In the **Disjoint path requirements** section, select the disjoint path requirements and priorities. For more information, refer to Disjoint path requirements on page 249.
5. In the **Constraints** section, select the constraints. For more information, refer to Path constraint options on page 250.
6. (Optional) In the **Tag updated LSP Paths with** field on the **Run Settings** page, override the defaults for how LSP paths are tagged (*DSJopt*).
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:

   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

   The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

   > **Note**
   >
   > Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:

- If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
- If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

9. Submit your changes.

---

**What to do next**

Refer to LSP disjoint path optimization report on page 254.

**LSP disjoint path optimization report**

Lists the sections and attributes included in a LSP Disjoint Path Optimization report generated by the optimization tool, including how to access the report and summarized metrics.

The LSP Disjoint Path Optimization report provides a summary of the optimization results.

Each time the optimization tool is run, a report is automatically generated. You can access this information at any time by choosing **Actions > Reports > Generated reports** and then clicking the **LSP Disjoint Path Optimization** link in the right panel.

The resulting report summarizes the number of LSPs and number of updated LSP paths. Depending on the disjoint option selected, the report summarizes the uniquely distinguishing attributes, such as LSP name and disjoint group name.

The LSP Disjointness, and Path Disjointness areas all list the number of common objects (selected as disjoint path requirements) before and after the optimization, as well as disjointness violations based on these requirements before and after the optimization.

## LSP loadshare optimization

Explains how LSP loadshare optimization automates the balancing of traffic across parallel LSPs and helps prevent congestion in the network.

A LSP loadshare optimization is a Cisco Crosswork Planning tool that

- automates the process of finding and setting the most favorable loadshare ratios across parallel LSPs
- includes only interfaces that use parallel LSPs in the optimization process, and
- allows you to limit parallel LSP interfaces to only those on which you want to drive down the maximum utilization.

After completing the operation, the system tags the LSPs with *LSPLoadshare* and generates a report. The report includes:

- The number of LSPs and interfaces affected
- The number of bins used
- The resulting maximum interface utilization
- The number of interfaces exceeding the specified utilization threshold

### Run LSP loadshare optimization

Explains the steps to configure loadsharing to balance traffic across parallel LSPs and optimize interface utilization.

Use this task to optimize LSP loadshare settings to improve traffic distribution and minimize overutilization in your network.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.

   From the toolbar, choose **Actions** > **Tools** > **LSP optimization** > **LSP loadshare optimization**.

   Alternatively, choose **Preset workflows** > **Perform optimization**, select **LSP Optimization** as the optimization type, select **LSP loadshare optimization** from the drop-down list, and click **Launch**.

2. Choose the LSPs on which you want to run the optimizer and click **Next**.

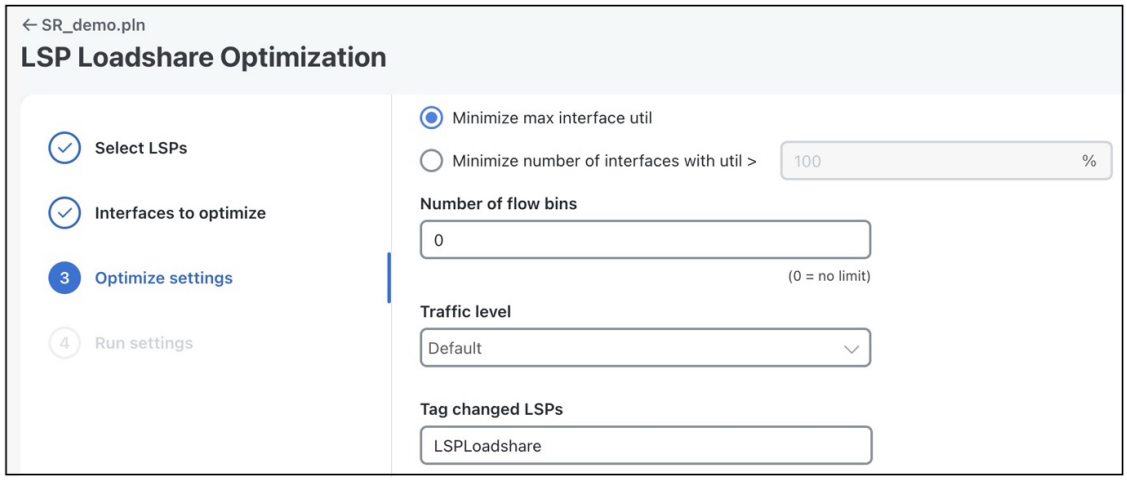
   It sets the loadshare setting between the selected parallel LSPs in the network. Parallel LSPs are those with the same source and destination nodes. By default, none of the LSPs are selected. However, you can also choose a subset or all the LSPs to optimize for load sharing purposes.

   The optimizer does not set a loadshare value of 0 on LSPs.

3. Choose the interfaces on which you want to drive down utilization and click **Next**.
   - If you select all the interfaces, the optimizer uses all interfaces in the network model.
   - Choose only those interfaces of interest on which to drive down the maximum utilization.

4. Specify the relevant optimization settings. For field descriptions, refer to LSP loadshare optimization settings on page 256.

   **Figure 72: LSP loadshare optimization options**




5. Select which traffic level to optimize.
6. (Optional) In the **Tag changed LSPs** field, override the defaults for how LSPs are tagged (*LSPLoadshare*).
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

> **Note**
>
> Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:

   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
   - If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

9. Submit your changes.

---

### LSP loadshare optimization settings

Lists the LSP loadshare optimization parameters and explains how each affects traffic distribution and network usage across LSPs.

This table describes the available parameters for LSP loadshare optimization which controls how traffic is distributed among LSPs.

**Table 28: LSP loadshare optimization settings**

| Field | Description |
|---|---|
| Minimize max interface util | Minimizes the maximum interface utilization over all interfaces on the LSP routes. Cisco Crosswork Planning tries to minimize the number of loadshare parameter changes required to do so. |
| Minimize number of interfaces with util >___% | Minimizes the number of interfaces that have a utilization greater than the specified value. This is a looser constraint than minimizing the maximum interface utilization. Therefore, Cisco Crosswork Planning has the opportunity to modify as few loadshare values as possible, thus reducing the amount of reconfiguration required. |
| Number of flow bins | Routers typically cannot divide flows arbitrarily between parallel LSPs, but instead allocate them to a fixed number of “bins” of approximately equal size. The bins are then divided between the parallel LSPs. This option lets you specify the total number of bins, which in turn places a constraint on the traffic division. |
| Traffic level | Specifies the traffic level to use in the optimization. |
| Tag changed LSPs | Specifies a tag for any changed LSPs. By default, the optimizer tags upgraded circuits with the label *LSPLoadshare*. |

### Minimization strategies for interface utilization

Provides details and outcomes on applying minimization strategies to optimize interface utilization within a sample network plan.

This is an example to apply minimization strategies and adjust Loadshare parameters to a sample plan file. The tags are applied on the optimized LSPs only, resulting in lower interface utilization and an updated plan file for further analysis.

The base plan for this example, *Acme_Network*, contains two sets of parallel LSPs, each with a Loadshare value of 1 and each using strict hops on named paths to reach its destination. All of these LSPs have a Traff sim of 3000 Mbps. The interfaces to which these LSPs filter all have a Traff sim value of 3000 Mbps, except for sjc-to-okc, which has 6000 Mbps of simulated traffic (Figure 73: Example Acme network before LSP loadshare optimization on page 257).

Using the *Acme_Network* plan file, if you minimize the maximum interface utilization across all LSPs, all four Loadshare parameters change, and the maximum interface utilization is 40% (Figure 74: Example Acme network after minimizing maximum interface utilization on page 257). The default LSP tags are used, thus naming the tags as *LSPLoadshare*. If you select the **Display results in a new plan file** option, then a new plan file *Plan-file-1.pln* is created by default.

**Figure 73: Example Acme network before LSP loadshare optimization**

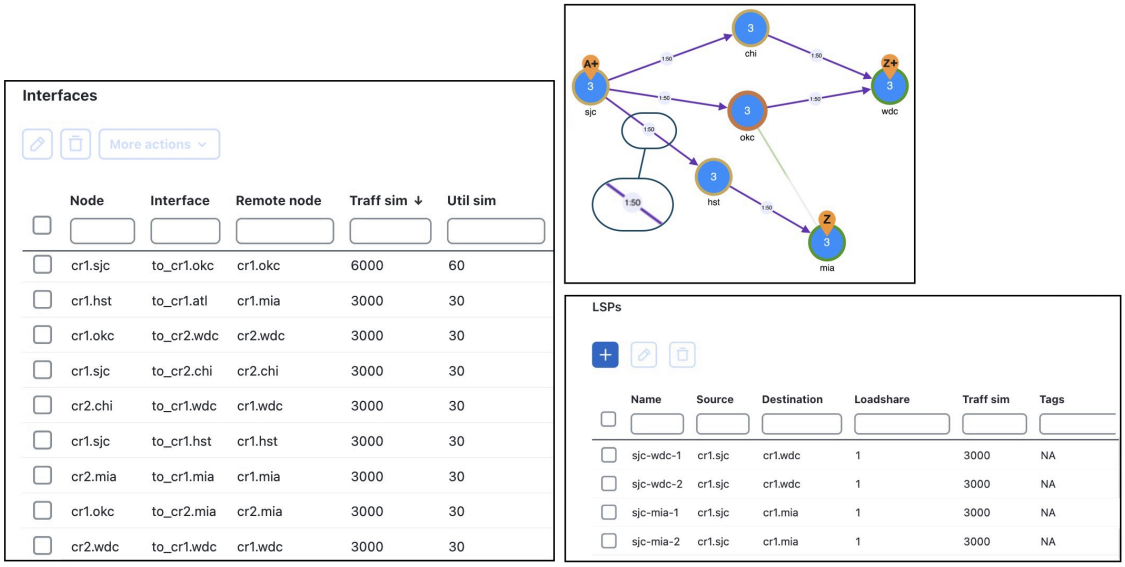
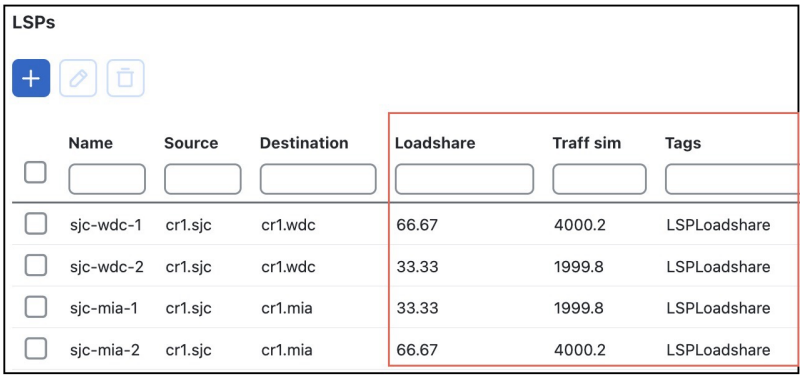
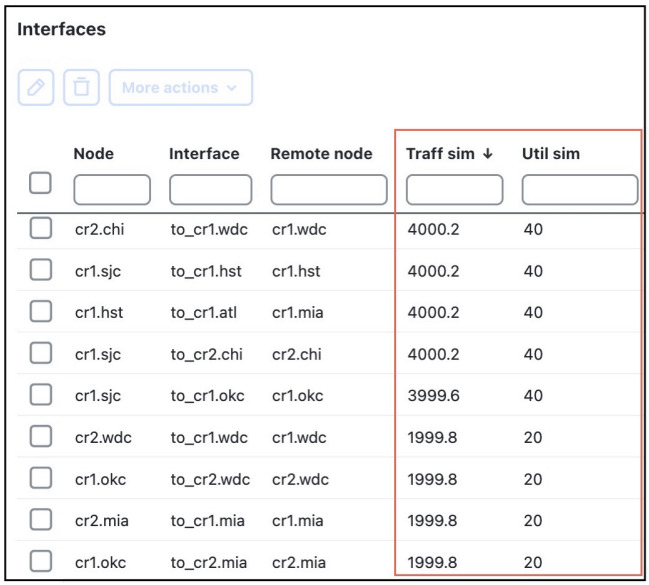



**Figure 74: Example Acme network after minimizing maximum interface utilization**




### Traffic allocation minimization for LSP bins

Provides a practical example of optimizing traffic allocation across two LSPs using bins, illustrating how network nodes approach fractional traffic splits.

This example describes how traffic allocation minimization uses bins to approximate optimal splits when a node must route traffic over multiple LSPs.

In this example, consider a node that uses a maximum of 32 bins, and the optimal traffic allocation is 45% of the traffic through LSP A and 55% of the traffic through LSP B. The node that sources these two LSPs, however, cannot do this exact split.

The optimization divides the traffic into 32 bins, each with the same amount of traffic in them. So, each bin has 3.125% of the traffic (100% of the traffic divided by 32). The optimization also determines how the node should split the traffic. In this case, the split is to give 43.75% (which is 14 bins of 3.125% each) to LSP A and 56.25% (which is 18 bins of 3.125% each) to LSP B. Thus, it optimizes and distributes 32 (14+18) bins of traffic. This is as close as possible to the optimal 45%/55% split using 32 bins.

## LSP setup bandwidth optimization

Explains how optimizing LSP setup bandwidth improves network scalability, load balancing, and resilience during failures.

As a network operator, you may need to periodically update LSPs as traffic changes in your network. You may have some fixed rules around the maximum setup bandwidth for LSPs. Restrictions on the setup bandwidth per LSP affect the number of LSPs in the network. As the setup bandwidth increases, the number of LSPs that you have to manage in your network

is reduced. However, the risk that certain LSPs cannot be routed also increases. As the setup bandwidth decreases, more alternative paths can be found, allowing for better load balancing. This holds true both in the fail-free and under failure cases.

To address these LSP setup bandwidth requirements, Cisco Crosswork Planning includes an **LSP setup BW optimization** tool (**Actions > Tools > RSVP LSP Optimization > LSP setup BW optimization**).

## Run LSP setup bandwidth optimization

Describes the configuration of LSPs to meet specific setup bandwidth requirements using the LSP setup bandwidth optimization tool.

Use this task to optimize the number and configuration of LSPs in your network to match set bandwidth goals.

The **LSP setup BW optimization** tool lets you add or remove LSPs based on set bandwidth requirements that you specify. You can select the LSPs that you want the optimizer to consider. The optimizer then separates these LSPs into different groups. The optimizer defines an LSP group based on the fact that the LSPs within the group share common source and destination nodes. You can also specify additional custom groupings to get a finer granularity on the setup bandwidth.

The optimizer lets you create LSPs by specifying:

- The number of LSPs to create for each LSP group.
- The maximum setup bandwidth per LSP. In this case, the minimum number of LSPs is created for each group in order to meet this requirement.

Additionally, the optimizer lets you remove LSPs by specifying:

- The number of LSPs to remove for each group.
- The minimum setup bandwidth per LSP. In this case, the minimum number of LSPs is removed from each group in order to meet this requirement.

You can then perform different operations on each group to test different design scenarios.

> **Note**
>
> The sum of the setup bandwidth of the LSPs per group is not modified by adding or removing LSPs. The setup bandwidth is evenly redistributed among the LSPs within each LSP group.

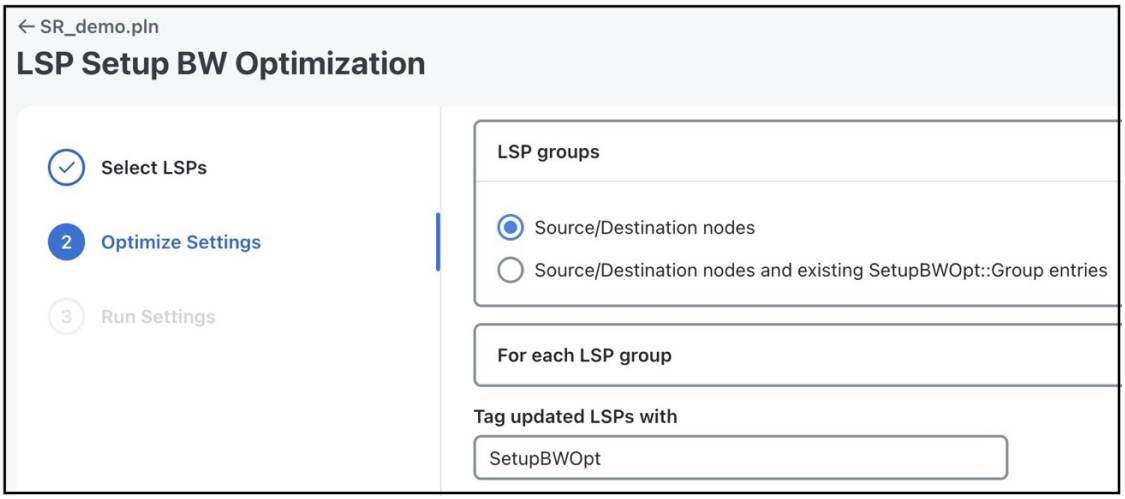
### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose any of the these options:

   From the toolbar, choose **Actions** > **Tools** > **RSVP LSP optimization** > **LSP setup BW optimization**.

   Alternatively, you can also choose **Preset workflows** > **Perform optimization**, select **RSVP LSP Optimization** as the optimization type, select **LSP setup BW optimization** from the drop-down list, and click **Launch**.
3. Choose the LSPs you want the optimizer to consider. By default, none of the LSPs are selected.
4. In the **LSP groups** section, choose the manner in which LSPs are separated into LSP groups:

   **Figure 75: LSP setup bw optimization options**




- **Source/Destination nodes**: LSPs with a common source and destination belong in the same group.
- **Source/Destination nodes and existing SetupBWOpt::Group entries**: LSPs with a common source and destination and common SetupBWOpt::Group entries belong in the same group. This option is useful when you want to group LSPs based on different service classes.

5. In the **For each LSP group** section, provide your setup bandwidth requirements for each LSP group:
   - **Create x LSPs**: Enter a positive integer. The optimizer creates a certain number of LSPs per group; for example, 1.
   - **Create the minimum number of LSPs so that setup BW <= x Mbps**: Specify the minimum number of LSPs to meet your setup bandwidth rule. For example, the bandwidth rule could be 10000 Mbps.
   - **Remove x LSPs**: Enter a positive integer. The optimizer removes a certain number of LSPs per group; for example, 1.
   - **Remove the minimum number of LSPs so that setup BW >= x Mbps**: Specify the removal of the minimum number of LSPs to meet your setup bandwidth rule; for example, 10000 Mbps.
6. (Optional) In the **Tag updated LSPs with** field, override the defaults for how LSPs are tagged ( *SetupBWOpt*).
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

     The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

> **Note**
>
> Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:

   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
   - If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

9. Submit your changes.

---

### What to do next

Refer to the LSP setup BW optimization report on page 261.

### LSP setup BW optimization report

Lists the key attributes available in the LSP Setup BW Optimization report, including summary statistics and details about LSP groups, to help users analyze optimization results.

The LSP Setup BW Optimization tool automatically generates a report every time it is run. You can access this information at any time by choosing **Actions > Reports > Generated reports** and then clicking the **LSP Setup BW Optimization** link in the right panel.

The report contains these two tabs:

#### Summary

The **Summary** tab summarizes the total number of

- LSPs you selected
- LSP groups identified
- LSPs created
- LSPs removed

#### LSP groups

The **LSP groups** tab provides the details on how the optimizer operated on LSPs:

- **Group**: Name of the LSP group as specified by SetupBWOpt::Group. If this group is not specified, the Group column is empty
- **LSP Source**: Source node of the LSP group
- **LSP Destination**: Destination node of the LSP group

- **Total Setup BW**: Sum of the setup bandwidth values of the LSPs within the LSP group
- **LSP Number Before**: Number of LSPs within the LSP group before running the optimization
- **LSP Number After**: Number of LSPs within the LSP group after the optimization
- **LSP Setup BW Before**: Average setup bandwidth within the LSP group before the optimization
- **LSP Setup BW After**: Setup bandwidth within the LSP group after the optimization

# 19 Configure RSVP-TE Routing

**Topics:**


- Dynamic LSP routing and CSPF
- RSVP LSP path attributes and behavior
- Named paths and explicit LSP routing
- Actual paths
- Affinities
- Global simulation parameters
- Fast Reroute (FRR) simulations
- Autobandwidth-enabled LSPs
- Circuit Style RSVP LSPs
- Advanced RSVP-TE LSP simulations
- P2MP LSPs
- Unresolved LSP destinations and hops


Describes how Cisco Crosswork Planning simulates RSVP-TE routing. Unless otherwise noted, *LSPs* refer to RSVP-TE LSPs. Note that Cisco Crosswork Planning does not distinguish between LDP and RSVP-TE LSPs.

## Dynamic LSP routing and CSPF

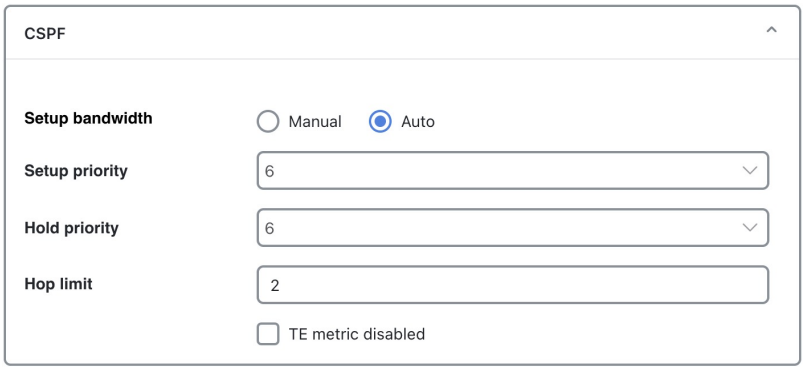
Explains how dynamic RSVP LSP routing uses CSPF, describes related MPLS interface properties, LSP reservation views, and setup bandwidth initialization.

If an RSVP LSP contains no LSP paths (also called *MPLS TE tunnel paths*), it is routed dynamically using Constrained Shortest Path First (CSPF). The weights used for the CSPF calculation are the TE metrics per interface. If the TE metric is not configured for an interface, then IGP metric is used.

If there are equal-cost routes, CSPF uses these selection criteria in order:

- Bandwidth available: The route with the greatest reservable bandwidth is chosen.
- Hop count: If bandwidth is equal, the route with the fewest hops is chosen.
- Random: If neither bandwidth nor hop count differentiate the routes, a route is chosen at random.

CSPF properties are set in the Edit LSP page and are viewable from the LSPs table. Note that these properties are available only for RSVP LSPs.




- **Setup bandwidth** (Manual): The amount of traffic (in Mbps) the source node requests for this LSP. The requested bandwidth is available from the reservable bandwidth of each interface in the path.
- **Setup bandwidth** (Auto): Automatically updates the Setup BW Sim value when using the Auto bandwidth convergence mode.
- **Setup priority**: Sets the priority for allocating reservable bandwidth. This is the order in which LSPs are signaled. The lower the number, the higher the priority.
- **Hold priority**: Determines which LSPs can preempt others on the shortest path. This is typically set for a particular service or traffic type. The lower the number, the higher the priority.
- **Hop limit**: Specifies the maximum number of hops permitted in the LSP route. If no path is available with this number of hops or fewer, the LSP is not routed.
- **TE metric disabled**: Uses IGP metrics for the LSP route if checked. If unchecked (default), the LSP uses TE metrics.

### Configure interface MPLS properties

Provides instructions for configuring interface MPLS properties that affect affinities, reservable bandwidth, FRR behavior, metrics, and RSVP LSP routing simulations.

Use this task to update the MPLS properties on one or more interfaces or circuits to control reserved bandwidth, traffic engineering metrics, and LSP routing behavior.

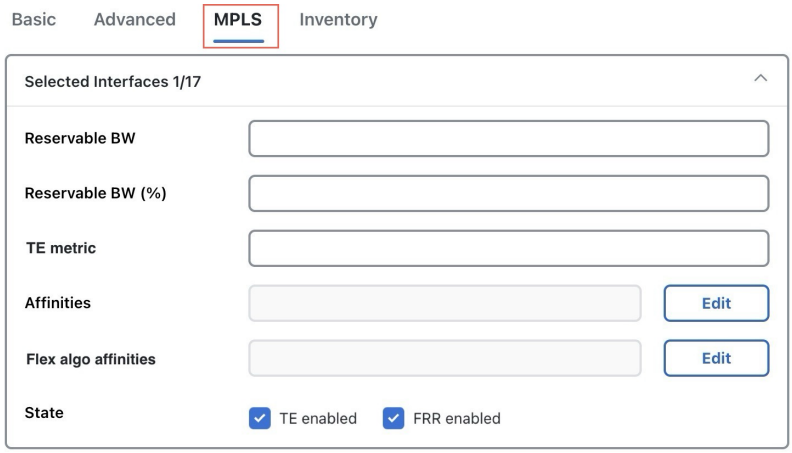
MPLS properties are set in the MPLS tab of the Edit interface or circuit page, and are viewable in the Interfaces or Circuits table.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. Select one or more interfaces or circuits from their respective tables.
3. Click .

> **Note**
>
> If you are editing a single interface or circuit, you can also use the ··· > **Edit** option.

**Figure 76: Interface MPLS properties**




4. Set these properties:
   - **Reservable BW**: Shows the amount of bandwidth that can be reserved by LSPs routed over the interface.
   - **Reservable BW (%)**: Shows the percentage of bandwidth that can be reserved by LSPs routed over the interface.
   - **TE metric**: Determines which path an LSP takes.
   - **Affinities**: Lists which affinities are set. For more information, refer to Affinities on page 274.
   - **Flex algo affinities**: Lists the FlexAlgo affinities assigned to the interface (MPLS).
   - **State**:
     - **TE enabled**: Identifies whether an LSP can be routed over the interface. If checked, the value is set to True, and TE metrics are used in routing the LSP, provided the LSP's **TE metric disabled** field is unchecked.
     - **FRR enabled**: Designated interface to be avoided by the FRR LSP. Only interfaces with this property set are used by the FRR LSPs initializer when creating FRR LSPs. This property does not affect FRR LSPs that are manually created. For more information, refer to Fast Reroute (FRR) simulations on page 280.
5. Save your changes.

These columns are updated in the Interfaces or Circuits tables:

- **TE metric sim**: Derived column that shows the effective TE metric. If the TE Metric is empty, this is set to the IGP metric.
- **Resv BW sim**: Derived column.
  - If a Reservable BW value is entered, this is copied to the Resv BW sim column.
  - If Reservable BW and Reservable BW (%) are “na”, the Resv BW sim is copied from the Capacity sim column.
  - If Reservable BW is “na”, but Reservable BW (%) has a value, the Resv BW sim value is derived by this formula:

    ```
    Capacity sim * (Reservable BW (%) / 100)
    ```

    If needed, you can change this behavior so that reserved bandwidth constraints are not used during simulations, which could be useful for planning purposes. For information, refer to Ignore reservable bandwidths for capacity planning on page 306.
- To determine how much of the interface traffic is contained in LSPs, show these columns in the Interfaces table. The total of the two equals **Traff sim**.
  - **Traff sim LSP**: The total amount of LSP traffic on the interface.
  - **Traff sim non LSP**: The total amount of non-LSP traffic on the interface.

  To determine which interfaces carry LSP traffic, sort the columns. To determine the demands associated with one or more selected interfaces based on whether or not they have LSP traffic, select them and choose > **Filter to demands > Through all interfaces**. Then, choose > **Filter to LSPs**.
- **Loadshare**: Show this column in the LSPs table. It indicates the value for redistributing either traffic across LSPs based on the ratio of this value with other LSPs that are parallel (have the same source and destination).

## View LSP reservations

Guides you through viewing LSP reservations on interfaces and comparing simulated traffic with LSP reservation utilization in the network plot.

Use this task to view LSP reservations on interfaces and compare reservation utilization with simulated traffic in the network plot.

In the LSP reservations mode, interfaces are colored according to the percentage of bandwidth (in comparison to what is reserved for the LSP).

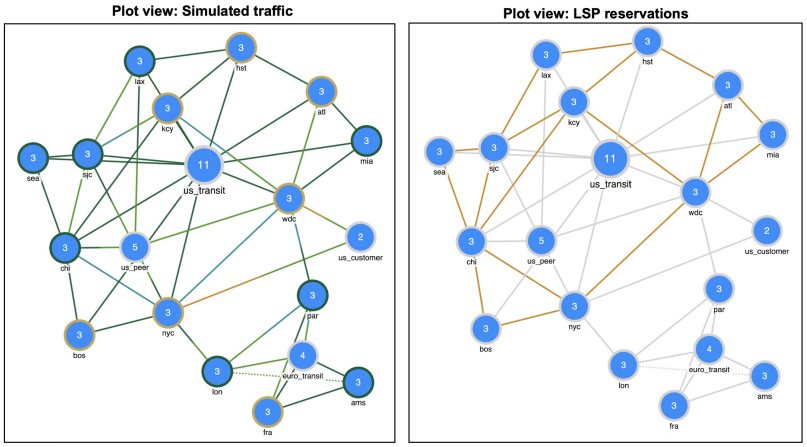
### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the **Plot view** drop-down list in the toolbar, choose **LSP reservations**.

The network plot updates, displaying interfaces color-coded by the percentage of bandwidth reserved for LSPs. The utilization colors are based on the LSP util column in the Interfaces table. LSP util is the sum of all setup bandwidths for all LSPs through the interface, as a percentage of the Resv BW sim.

Figure 77: Simulated traffic versus LSP reservations on page 266 shows how the network plot changes based on the Plot view options you select, such as Simulated traffic or LSP reservations.

**Figure 77: Simulated traffic versus LSP reservations**




## Calculate LSP setup bandwidth

Provides instructions for running the LSP setup bandwidth initializer to calculate setup bandwidth values for selected RSVP-TE LSPs and traffic levels.

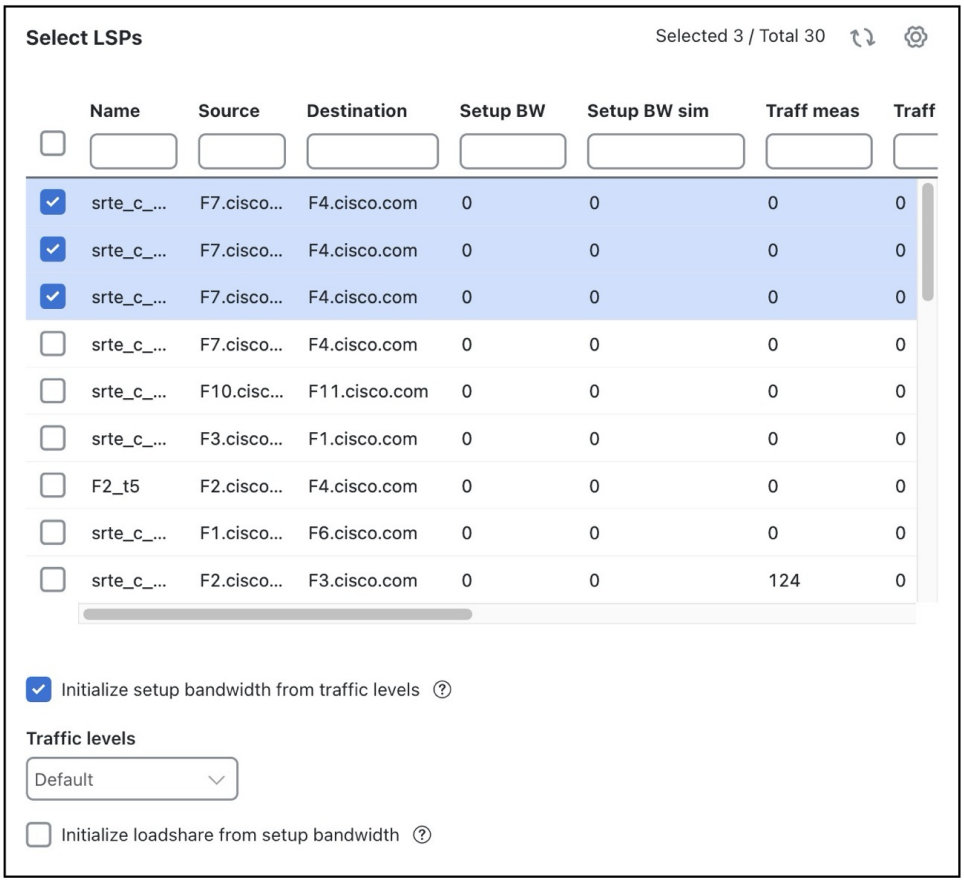
Use this task to calculate setup bandwidth values for selected RSVP-TE LSPs by running the LSP setup bandwidth initializer.

The LSP setup bandwidth initializer calculates the setup bandwidth for existing RSVP-TE LSPs. The setup bandwidth is set to the maximum amount of demand traffic passing through the LSP for the traffic levels selected.

If using only one traffic level, the best practice is to use the **Autobandwidth convergence** mode ( > **Simulation** > **Simulation convergence mode**). For more information, refer to Autobandwidth-enabled LSPs on page 289.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the toolbar, choose **Actions** > **Initializers** > **LSP setup bandwidth**.




3. Select the LSPs you want to optimize.
4. Choose whether to set the setup bandwidth value to be equal to the maximum amount of traffic passing through each tunnel across selected traffic levels. Use the **Initialize setup bandwidth from traffic levels** check box for this purpose.
5. Choose which traffic levels to use for demand traffic.
6. Choose whether to set the load share values to be equal to the setup bandwidth. To set, check the **Initialize loadshare from setup bandwidth** check box.
7. Submit your changes.

## RSVP LSP path attributes and behavior

Lists and explains the key attributes of RSVP LSP paths and describes how their configuration influences route selection and simulation scenarios.

Like LSPs, LSP paths have CSPF properties, such as **Setup priority** and **Hop limit**. For a description of LSP properties, refer to Dynamic LSP routing and CSPF on page 264.

### Key points

- If the LSP path does not define properties like setup priority or hop limit, it inherits them from the LSP. If these properties are set in the LSP path, they override the LSP settings.
- If operating in the **Autobandwidth convergence** simulation mode and if the LSP's **Setup bandwidth** is set to **Auto**, the LSP path Setup BW property (like the LSP Setup BW) is ignored, and the LSP's Setup BW sim value is calculated and used.

- If an LSP path does not have an associated named path, it is routed dynamically. However, you can use named paths to fully or partially describe a path from the LSP source to destination.

> **Note**
>
> This chapter refers to LSPs and LSP paths as *dynamic*, *explicit*, or *loose explicit*, depending on whether their associated named paths have no defined hops, all hops defined, or only some hops defined.

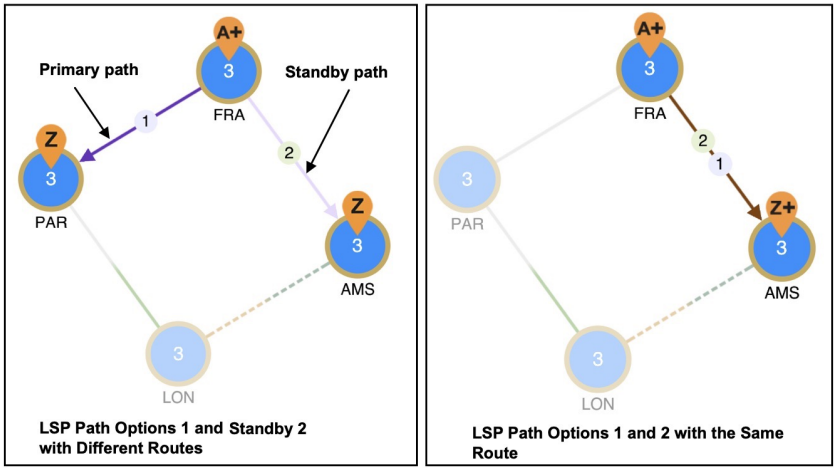
## Standby paths

Details how standby paths provide alternate RSVP LSP paths and describes how they support routing continuity during failure simulations.

A standby path is an LSP path

- establishes alternative LSP paths for traffic in RSVP-TE networks, and
- reserves setup bandwidth for these alternate paths to ensure immediate availability during failure.

Standby paths are always established even if the associated LSP is routed using a path with a different path option.




## How LSP paths respond to failures

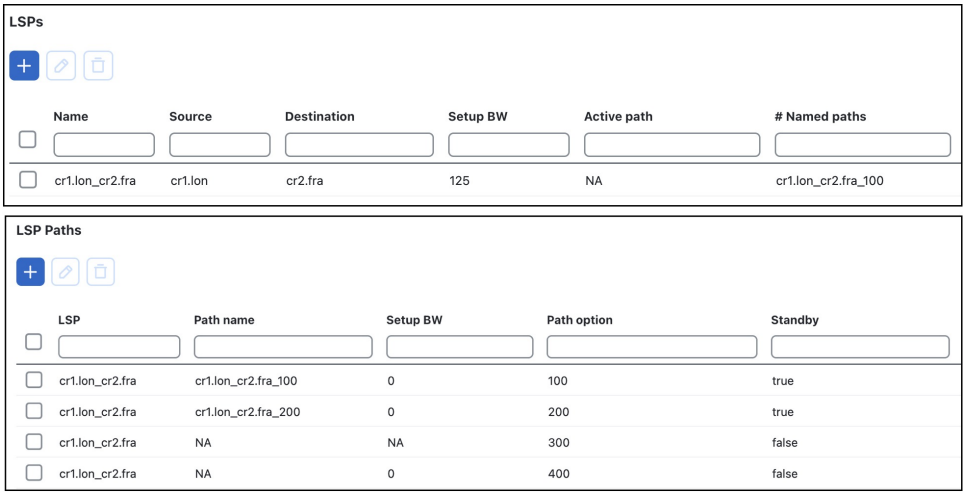
Outlines a failure scenario that shows how RSVP LSP paths are evaluated and selected when links or paths become unavailable.

### Summary

Figure 78: Example RSVP LSP and associated LSP paths on page 269 shows an RSVP LSP (cr1.lon_cr2.fra) that uses four LSP paths. The cr1.lon_cr2.fra LSP has a primary and standby secondary path options (100 and 200), each with a setup bandwidth of 0. They are both established, using their defined named paths. Both primary and secondary LSP paths have associated named paths, which appear in the **Path name** column. The other two paths do not have associated named paths, and so are dynamically routed.

### Workflow

**Figure 78: Example RSVP LSP and associated LSP paths**




This example shows how LSP responds to failures:

1. If the primary path (cr1.lon_cr2.fra_100) fails, the secondary path (cr1.lon_cr2.fra_200) is used.
2. If both paths with associated name paths fail, a dynamic path is established with a reserved bandwidth of 125 Mbps, which it inherits from the LSP. This path is not set up initially because it is not a standby.
3. If all paths fail, the path of last resort is a dynamic path with zero bandwidth. This ensures that there is always an LSP up so that services relying on the LSP are always available.
4. The simulation changes if the LSP Active path has been set. For example, if cr1.lon_cr2.fra_200 is the active path (that is, 200 appears in the LSP's Active path column), Cisco Crosswork Planning uses cr1.lon_cr2.fra_200 first, and then goes through the previous sequence.

## Named paths and explicit LSP routing

Introduces named paths for explicit RSVP LSP routing and outlines hop examples plus procedures to create and edit named paths.

Named paths specify a route though the network using an ordered list of adjacencies. The route is defined by named path hops, which can be nodes or interfaces. Each hop type specifies whether the route should be strict, loose, or excluded.

- Strict: The LSP must reach the named path hop directly from the previous one, with no intermediate interfaces.
- Loose: The LSP is routed to this hop from the previous hop using CSPF. Intermediate interfaces can be used.
- Exclude: The node or interface is excluded from the LSP path. This hop type cannot be combined with strict or loose hops in the same named path.

The named paths and their hops are listed and selectable from the Named paths and Named path hops tables. The advantage of having separate tables is that you can have named paths that have missing or unresolved hops. Also, the path name can be reserved even if it is not part of the Cisco Crosswork Planning simulation. The hop type is identified visually on the plot, as well as listed in the **Type** column of the Named path hops table.

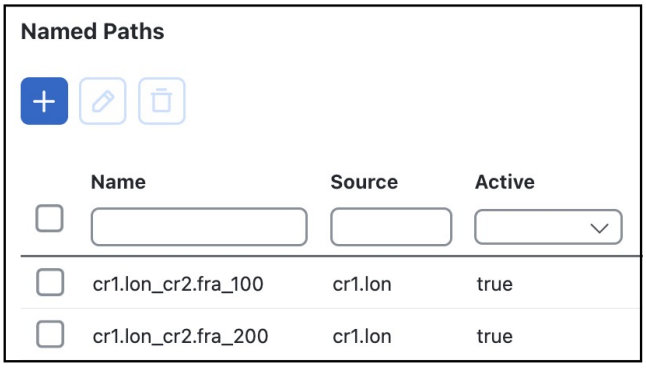
### Named path hops example

Details an example of named path hops and shows how ordered adjacencies guide explicit routing for RSVP LSP paths through the network.

Figure 79: Example named paths the cr1.lon_cr2.fra LSP on page 271 extends the earlier example (Figure 78: Example RSVP LSP and associated LSP paths on page 269), showing two named paths for the cr1.lon_cr2.fra LSP. This shows

that the first two LSP paths each have a named path. The naming convention is *<LSP_Name>_<Path_Option>*, which in this case is cr1.lon_cr2.fra_100 and cr1.lon_cr2.fra_200.

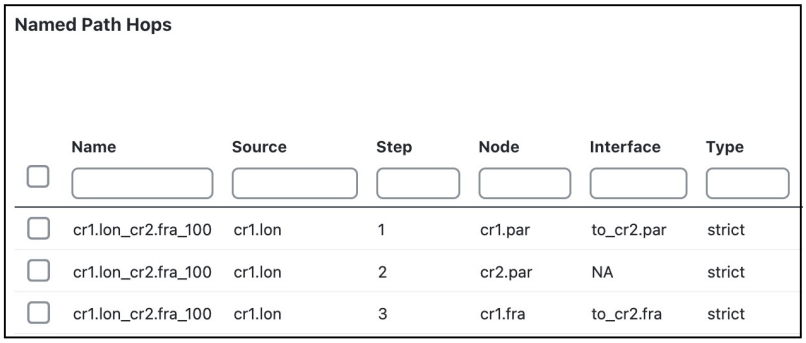
**Figure 79: Example named paths the cr1.lon_cr2.fra LSP**




The hops for each named path are defined in a separate table called **Named Path Hops**. Figure 80: Example named path hops for cr1.lon_cr2.fra_100 on page 271 shows the named path hops for the cr1.lon_cr2.fra_100 named path. The **Step** column shows the hop order.

- The first hop is an interface hop: a strict hop on the interface from cr1.par to cr2.par.
- The second is a node hop, also strict, to cr2.par.
- The third hop is an interface hop: a strict hop on the interface from cr1.fra to cr2.fra.
- Hops for cr1.lon_cr2.fra_200 are defined similarly.

**Figure 80: Example named path hops for cr1.lon_cr2.fra_100**




## Create named paths and their hops

Provides instructions for creating named paths and defining ordered hops so RSVP LSPs can follow explicit routes through the network.

Use this task to create named paths and define ordered hops for explicit RSVP LSP routing through the network.

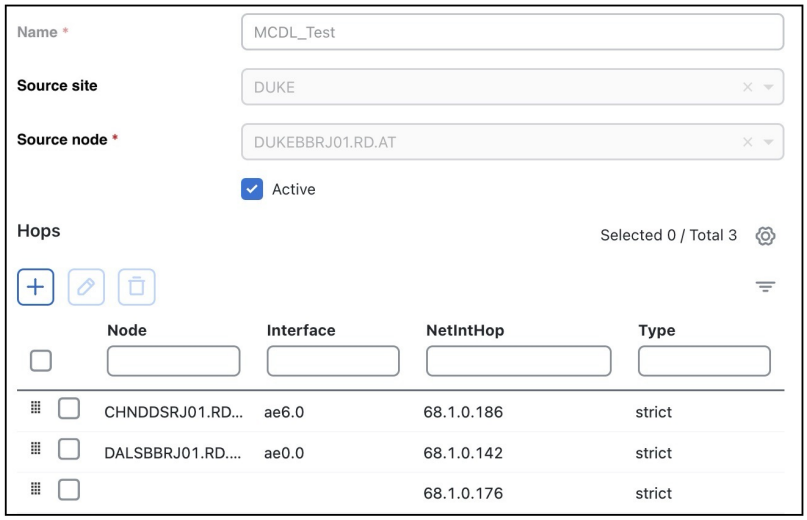
### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. Create named paths using any of these options:
   - From the toolbar, choose **Actions** > **Insert** > **LSPs** > **Named path**.
   - In the Network Summary panel on the right side, click + in the **Named paths** tab.

The Named paths tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **Named paths** check box.

- When creating LSP paths, check the **Create associated named paths** check box. Using this option does not create hops.

3. Enter these information:
   - **Name**: Enter the name of the named path.
   - **Source site** and **Source node**: Use these options to change the source site or source node.
   - **Active** check box: Use this option to make the named path active.




4. Configure the hop options:
   - To add a new hop, click .
   - To edit an existing hop, select it from the Hops list and then click . The Edit page opens for named path hops.
   - To delete an existing hop, choose it from the Hops list, click .

5. To continue creating or editing the named path hop, use these options.
   - Choose the site, node, and interface. For node hops, do not choose an interface.
   - Choose the **Type** as Loose, Strict, or Exclude. For more information on these options, refer to Named paths and explicit LSP routing on page 270.

6. Click **Add**.

## Edit named paths and their hops

Provides instructions for editing existing named paths, including adding, modifying, or deleting hops that define explicit RSVP LSP routes through the network.

Once the named paths are created, you can create, edit, or delete their named path hops. Note that discovered named paths can contain *unresolved* hops, which are nodes and interfaces that are not in the plan file. For information about resolving named paths, refer to Unresolved LSP destinations and hops on page 316.

The recommended method of editing the named path hop type is to access it directly from the Named paths table. This is the most efficient way of viewing the entire path.

**Procedure**

1. Select one or more named paths from the Named paths table.
2. Click .

> **Note**
>
> If you are editing a single named path, you can also use the ••• > **Edit** option.

3. In the **Hop** section, edit the Named path hop details, as required. For more information, refer to Create named paths and their hops on page 271.
4. Save your changes.

## Actual paths

Describes how actual paths are used in LSP routing simulations and how their states affect overall MPLS routing predictions.

Actual path is a path that

- represents the real routes taken by routed LSP paths
- is read from the live network, per LSP path, and
- reflects current LSP path states, including successful and failed routings.

RSVP-TE routing is not always completely predictable because it depends on the order in which LSPs are established and reserve their bandwidth on the network. Cisco Crosswork Planning uses actual paths in LSP routing simulations to match the current LSP routing state of a network as accurately as possible.

When an LSP path is routed, the actual route is attempted first. If this routing fails due to CSPF constraints, such as not enough reservable bandwidth or affinity restrictions, Cisco Crosswork Planning reverts to the standard CSPF routing algorithm.

**Example**

Cisco Crosswork Planning routes all LSPs in its simulation according to the actual paths read from the network, and the routings completely match the network. If a circuit fails, only the LSPs through this circuit are rerouted. Because they cannot be established along their actual paths, standard CSPF is used. The result is a prediction of incremental routing changes that would occur on the current network if such a failure were to occur.

An LSP or an LSP path might have a corresponding actual path. An LSP path with no actual path inherits its LSP actual path, if available.

Two columns in the LSPs and LSP Paths tables are useful to refer to the result of actual paths on LSP routing.

| Column | Description |
|---|---|
| **Actual** column in the **LSP Paths** table | The simulation can only use actual paths if they are resolved. If the network discovery was incomplete, this might not be possible.<br>• true = the simulation converted the actual path hops into a complete path from source to destination of the LSP<br>• false = the simulation did not convert the actual path hops into a complete path<br>• NA = not applicable; no actual path was available. |
| **Routed** column in the **LSPs** table | • Actual = The LSP follows an actual path.<br>• Simulated = The LSP does not follow an actual path.<br>• Unrouted = Routing was not possible. |

### Deactivate actual paths for simulations

Provides instructions for deactivating actual paths so MPLS simulations use configured LSP active path settings instead of discovered routes during routing.

Normally, the simulation routes LSPs using actual paths based on current network state and LSP active path settings. For planning purposes, when state is not relevant, you can change this behavior to disregard actual paths.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

   The Network options page opens.
3. Click the **Simulation** tab.
4. To use or disregard actual paths and active paths in MPLS simulation, check or uncheck **Use LSP actual paths, active paths** check box.
5. Save your changes.

## Affinities

Introduces affinity-based RSVP LSP routing and outlines affinity creation, interface assignment, association rules, routing examples, and LSP mesh usage for simulations.

An affinity is a mechanism that

- enables path diversity for LSPs, and
- allows assignment of attributes to physical network circuits and LSPs to implement a variety of routing policies.

For example, you can use affinities to restrict certain traffic to specific topological regions, or you can force primary and backup LSP paths onto different routes so they do not simultaneously fail.

Cisco Crosswork Planning supports an unlimited number of 64-bit affinities, each of which is defined by a number and an optional name. Named affinities are sometimes called *administrative groups* or *link coloring*.

The default plan files have 0-31 unnamed and unassigned affinities.

## Create and edit affinities

Provides instructions for creating and editing affinities that can later be assigned to interfaces and associated with RSVP LSPs for routing.

Use this task to create new affinities or edit existing affinities used in RSVP LSP routing policies.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the toolbar, choose **Actions** > **Edit** > **Admin groups**.

   The Admin groups page opens.
3. To create a new affinity:
   a) Click **Add**.
   b) In the **Affinity** and **Name** fields, enter the affinity number and affinity name.

      Affinity numbers must be unique. Affinity names are optional and must also be unique.

   c) Save your changes.
4. To modify an existing affinity:
   a) Select the affinity you want to edit.
   b) Click **Edit**.
   c) In the **Affinity** and **Name** fields, enter the affinity number and affinity name.

      Affinity numbers must be unique. Affinity names are optional and must also be unique.

   d) Save your changes.

## Assign affinities to interfaces

Provides instructions for assigning affinities to interfaces.

Use this task to assign affinities to interfaces so RSVP LSP routing can include or avoid selected network resources during simulations.

To use affinities, you must assign them to interfaces in a way that encourages favorable routes and discourages others. For example, continental paths can have one affinity and international paths can have another.



**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the Interfaces table, select one or more interfaces to which you are assigning the affinities.
3. Click .

   **Note**

   If you are editing a single interface, you can also use the ··· > **Edit** option.

4. Click the **MPLS** tab.
   a) Click the **Edit** button next to the **Affinities** field.
   b) In the **Include** column, check the boxes for the affinities that you want to associate with the selected circuits.
   c) Click **Save**.
5. Save your changes.

You can also use these steps to assign affinities to interfaces through the Edit Circuit page. However, you must choose and assign the affinity twice: once to each interface.

## LSP affinity associations

Details how LSP affinity associations use interface affinity settings to include or exclude resources during RSVP LSP routing path selection.

For affinities to influence LSP routing, you must associate LSPs with the affinities assigned to interfaces that you want to include or exclude in the routing process. If no affinity is specified, the LSP can be routed through any path. Explicit routes can also have affinities for hops with a Loose relationship (refer to Named paths and explicit LSP routing on page 270).

### Inclusion and exclusion rules

You can choose any of these rules while associating LSPs with affinities:

- **Include**: Only use interfaces that include all affinities.
- **Include any**: Use interfaces that include at least one affinity.
- **Exclude**: Do not use interfaces with this affinity.

LSP paths inherit the LSP affinities. However, you can edit LSP path affinities to have priority over the LSP affinities.

### Example of affinity-based LSP routing

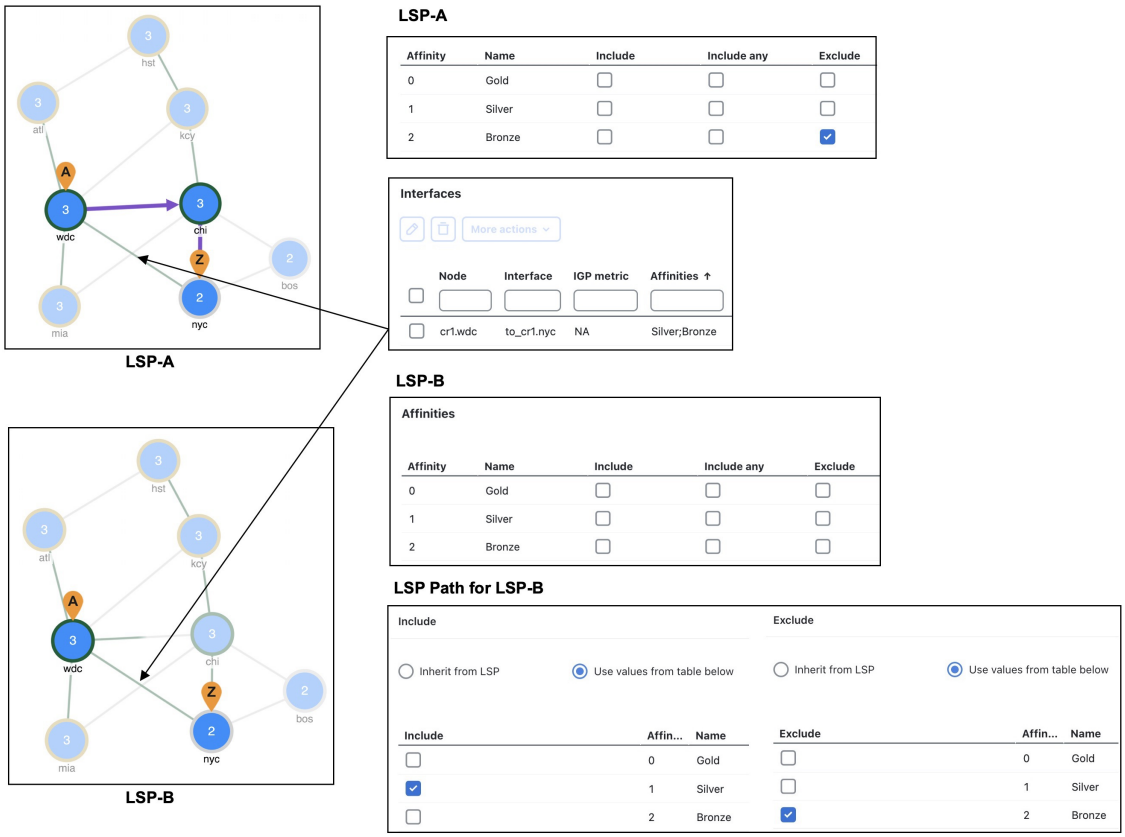
Details how affinities affect the routing of two LSPs.

This example demonstrates how affinities affect the routing of two LSPs with the same source and destination nodes.

- The to_cr1.nyc interface between wdc and nyc is assigned to both affinities 1 (Silver) and 2 (Bronze). No other interfaces are assigned affinities.
- LSP A is configured to exclude all interfaces assigned to affinity 2. While it cannot take the shortest path using the to_cr1.nyc interface, it can route around it.

- LSP B is not associated with any affinities. Its LSP path is configured to include any interfaces assigned to affinity 1 (Silver) and exclude those assigned to affinity 2 (Bronze). Since the to_cr1.nyc interface is assigned to both affinities, LSP B cannot route.

**Figure 81: Affinity-based LSP routing example**




## Assign LSPs to affinities

Provides instructions for associating LSPs with affinities.

Use this task to associate LSPs with affinities so routing can include or exclude interfaces with matching affinity settings.

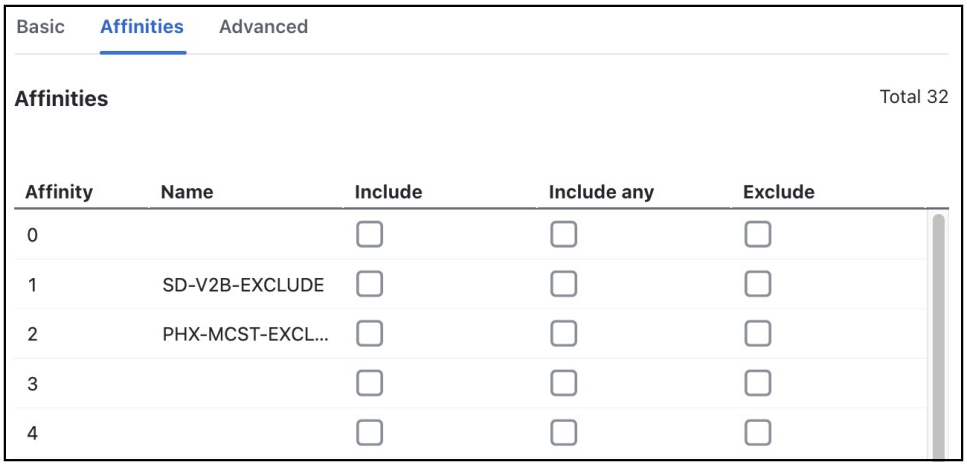
### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the **LSPs** table, choose one or more LSPs to which you are associating the affinities.
3. Click .

> **Note**
>
> If you are editing a single LSP, you can also use the ••• > **Edit** option under the **Actions** column.

4. Click the **Affinities** tab.

**Figure 82: LSP affinities**




5. Choose the inclusion or exclusion rule (Include, Include any, or Exclude) for the affinity you are associating with the selected LSPs.
6. Save your changes.

## Assign LSP paths to affinities

Provides instructions for associating LSP paths with affinities.

Use this task to associate LSP paths with affinities.

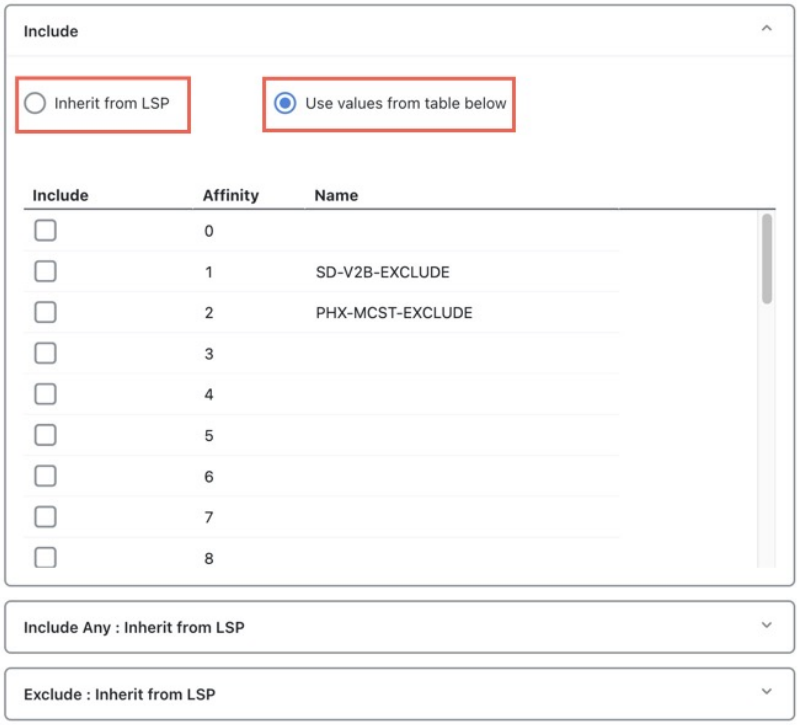
### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the LSP paths table, choose one or more LSP paths to which you are associating the affinities.
3. Click .

> **Note**
>
> If you are editing a single LSP path, you can also use the ••• > **Edit** option under the **Actions** column.

4. In the **Affinities** section, click the **Edit** button.
5. Choose the inclusion or exclusion rule (Include, Include any, or Exclude) for the affinity you are associating with the selected LSP paths. You can choose the LSP paths to inherit the LSP affinities or choose the values from the table. After making the selections, click **Save**.

**Figure 83: LSP path affinities**




6. Save your changes.

### Assign affinities when creating LSP meshes

Provides instructions for assigning affinities while creating LSP meshes.

Use this task to assign affinities while creating LSP meshes so that the generated LSPs use the intended routing constraints.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP mesh**.

   Alternatively, in the Network Summary panel on the right side, click + > **LSP mesh** in the **LSPs** tab.
3. After selecting the required nodes and options, go to the **Affinities** page.
4. In the **Affinities** section, click **Choose affinities**.
5. Choose the inclusion or exclusion rule for each affinity you are associating with all LSPs in the mesh, and click **Add**.
6. Submit your changes.

## Global simulation parameters

Outlines global simulation parameters that affect RSVP LSP routing, rerouting, bandwidth handling, and simulation behavior across MPLS planning workflows and analyses.

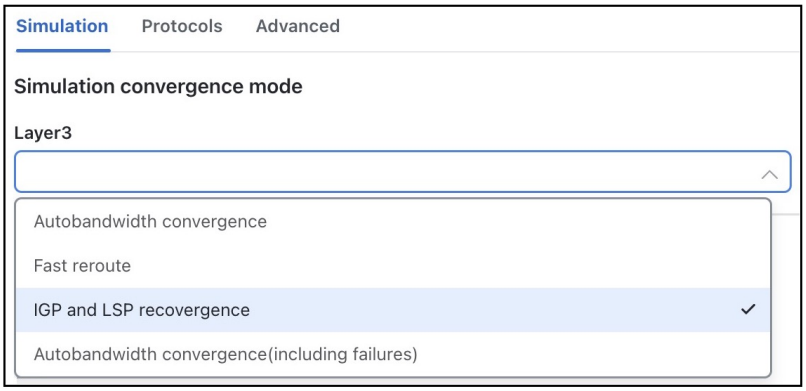
Cisco Crosswork Planning lets you set global parameters that affect how LSPs are routed or rerouted.

To access these options:

1. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.
2. Click the **Simulation** tab.

### Simulation modes

Cisco Crosswork Planning supports four simulation modes for RSVP-TE LSPs. These are listed as options in the **Layer 3** drop-down list in the **Simulation convergence mode** section.




- **Autobandwidth convergence**: In this mode, the network is simulated after the traffic (Traff sim) has been rerouted to other LSPs, but before the Setup BW sim values have been reset. Refer to Autobandwidth-enabled LSPs on page 289.
- **Fast reroute**: This mode uses FRR LSPs, which is a common failure restoration mechanism used by RSVP-TE LSPs. In a live network, FRR restoration typically occurs within milliseconds. Refer to Fast Reroute (FRR) simulations on page 280.
- **IGP and LSP reconvergence**: By default, Cisco Crosswork Planning simulates the state of the network once it has fully responded to a failure. Specifically, this is the network state after LSPs have re-established new routes around the failure, and the IGP has fully reconverged. The optimization tools only work in IGP and LSP reconvergence mode.
- **Autobandwidth convergence (including failures)**: In this mode, the network is simulated after the Setup BW sim values have been reset. Refer to Autobandwidth-enabled LSPs on page 289.

## Fast Reroute (FRR) simulations

Explains how Fast Reroute (FRR) simulations work, including FRR fundamentals, routing during failures, protection mechanisms, and simulation outcomes.

A Fast Reroute (FRR) simulation is a network analysis technique that

- models the behavior of Label Switched Paths (LSPs) under failure scenarios
- evaluates how locally protected LSPs reroute traffic using pre-signaled FRR LSPs or bypass tunnels, and
- measures restoration performance, often aiming for a 50 millisecond recovery.

Cisco Crosswork Planning simulates a Fast Reroute convergence mode using FRR LSPs. If a failure occurs in the route of a protected RSVP-TE LSP, a presignaled FRR LSP (or bypass tunnel) forwards traffic locally around the point of failure, typically from the node just before the failure to a node downstream. This restoration mechanism gives the source node (head-end) of the FRR LSP time to re-establish an alternate route around the failure. This period is sometimes described as the *50 millisecond* behavior of the network because FRR restoration ideally becomes effective within approximately 50 milliseconds of a failure.

Routing of LSPs and demands in FRR simulations differs from full convergence simulations:

- The source node of the protected LSP does not reroute the traffic. If a failure occurs on a node or circuit (link) within the LSP path, the routing occurs as described in these steps:
  - If an FRR LSP exists to protect a failed node or circuit, the protected LSP continues to be routed, but its new path includes the FRR LSP route that redirects the traffic around the failure. This traffic enters the FRR LSP's source node just prior to the failure and rejoins the protected LSP path after the failure at the FRR LSP's destination node.
  - If the LSP is not protected by an FRR LSP, then the LSP traffic is not routed.
- Since there is no IGP reconvergence, demands through unrouted LSPs are unrouted. Additionally, demands through failures outside of these LSPs are not routed.

## FRR fundamentals

Outlines key FRR properties, terms, network plot indicators, and failure visualization details used in RSVP-TE simulations for analyzing protected LSPs.

This section summarizes key properties and terms, and how FRR LSPs are visualized in the network plot.

**Objects and properties**

> **Note**
>
> FRR LSPs appear in the same table as the regular LSPs and are identified in the **Metric type** column as **FRR**.

**Table 29: Objects and properties**

| | | | |
|---|---|---|---|
| LSP | FRR LSP | FRR interface | Egress interface to be used by the FRR LSP. |
| | | FRR type | Type of FRR protection link or Node |
| | Active Egress Any | FRR detour | Mark an LSP to be protected by FRR LSPs |

| | | | |
|---|---|---|---|
| | na | FRR | Mas<br>an<br>to<br>be<br>by<br>FRR<br>LSPs<br>Thi<br>is<br>if<br>an<br>is<br>to<br>be<br>us<br>wh<br>usi<br>the<br>FRR<br>LSPs |

**Terms**

- Protected interface: An interface that is avoided when using FRR LSPs.
- Protected LSP: An LSP that would be rerouted through FRR LSPs in the event of failures. In Cisco Crosswork Planning, this requires the Fast Reroute simulation mode. To turn this mode on, refer to Run fast reroute simulation on page 288.
- Protected SRLG: An SRLG that contains a protected interface.

**Visualization**

An unused FRR LSP path highlights in purple when it is selected from the LSPs table. When routed under failure, FRR LSPs appear as a dotted line.

## FRR LSP routing

Details how FRR LSP source and destination choices support link protection and node protection in RSVP-TE simulations for failure analysis.

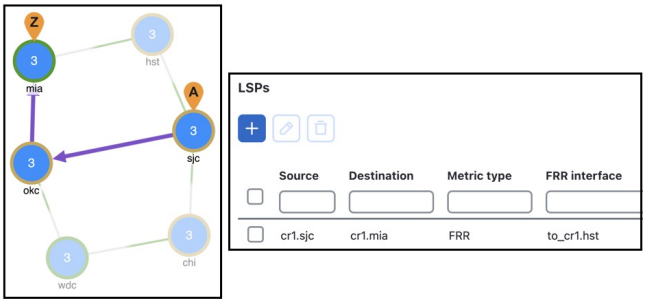
The source and destination of an FRR LSP can be any pair of nodes, although in practice the source and destination nodes are usually one hop (link protection) or two hops (node protection) apart. The source node for the FRR LSP has a designated interface that is configured so that it is avoided during Fast Reroute simulations. This is the *protected interface*.

This is configured using the **FRR interface** property, which is set in the Edit LSP page when manually creating the FRR LSP, or it is automatically created when running the FRR LSPs initializer. The initializer derives it from interface's FRR enabled property.

### Example

This image shows an FRR LSP sourced at sjc, destined for mia, and whose protected interface (FRR interface) is to_cr1.hst. Therefore, the FRR LSP path routes towards okc and away from hst. If this FRR interface property were to_cr1.okc instead, the FRR LSP path would route sjc-hst-mia, away from okc.

**Figure 84: FRR LSP routing example**




Like other LSPs, FRR LSPs can use actual and named paths. Unlike other LSPs, they do not use setup bandwidth when routing and they only use one LSP path.

## Types of FRR LSP protection

Lists and explains the types of LSP protection supported in Cisco Crosswork Planning, including link, node, and SRLG protection.

Cisco Crosswork Planning supports three main types of protection for FRR LSPs: link protection, node protection, and SRLG protection.

### Link and node protection

The type of link and node protection are displayed in the **FRR type** column of the LSPs table.

- Link-protection LSPs: Protect an LSP from a circuit failure by routing an FRR LSP around the failed circuit. The source node of the FRR LSP is the local node of the failed circuit, and if configured correctly, the destination of the FRR LSP is the remote node of the failed circuit.
- Node-protection LSPs: Protect an LSP from a node failure by routing an FRR LSP from the node one hop prior to the failed node in the protected LSP path to the node one hop after the failed node.

### SRLG protection

In a network, if interfaces are sourced from the same node (router) and are configured as an SRLG, an FRR LSP protecting one of these interfaces can be configured to reroute around that interface. If possible, the FRR LSP avoids all other interfaces in the SRLG. In Cisco Crosswork Planning, FRR LSPs can be configured to avoid all objects defined in the SRLG, not just interfaces on the source node of the FRR LSP.

## Configure routing around failures

Provides instructions for configuring routing around failures so protected RSVP LSPs can use FRR LSPs during simulations and analysis workflows.

When one or more failures occur, these decisions determine which FRR LSP to take:

### Before you begin

- LSPs must be configured to be protected by FRR LSPs. For information about setting up FRR LSPs, refer to Set up fast reroute simulations on page 285.

- Any interfaces along the FRR LSP path must be configured to be protected by FRR LSPs. For information about protecting interfaces, refer to Mark LSPs for protection on page 285 or the FRR LSPs initializer on page 286 (which automates the process of setting these).

**Procedure**

1. For each such interface on the LSP path, Cisco Crosswork Planning checks for an FRR LSP protecting that interface whose destination is a node further down the path of the LSP.
2. If more than one of these FRR LSPs exist, Cisco Crosswork Planning chooses the one that is farthest down that path of the protected LSP, thus simulating standard FRR behavior that prefers node protection over link protection.

   If more than one eligible FRR LSP has the same destination farthest down the protected LSP path, the FRR LSP is chosen arbitrarily among them.
3. If more failed circuits or nodes occur after this destination, further FRR LSPs are chosen in the same manner.

## Set up fast reroute simulations

Provides instructions for setting up FRR simulations, including protection properties, SRLG identification, FRR LSP initialization, and simulation execution workflow requirements.

Use this task to set up FRR simulations by configuring protection properties, SRLGs, FRR LSP initialization, and simulation execution.

**Procedure**

1. Set appropriate properties on LSPs to mark them for protection. Refer to Mark LSPs for protection on page 285.
2. If protecting SRLGs, associate a node with the SRLG in order to establish the protected interface within it. Refer to Identify SRLGs for protection on page 286.
3. If you are using the initializer to create FRR LSPs, identify which interfaces to protect. Refer to FRR LSPs initializer on page 286.
4. Create the FRR LSPs using the FRR LSPs initializer or manually. Refer to FRR LSPs initializer on page 286 or Create FRR LSPs manually on page 288.

### Mark LSPs for protection

Provides instructions for marking LSPs for FRR protection.

Use this task to mark LSPs for FRR protection by setting properties that identify protected interfaces, nodes, or SRLG resources.

**Procedure**

1. Select the required LSPs from the LSPs table and click .
2. In the **LSPs details** section, check the **FRR enabled** check box.

   This appears as "true" in the **FRR enabled** column of the **LSPs** table. If this is "false", the LSP is not protected.
3. In the **Routing** section, set the **Routing type** as **Autoroute** or **FA** (Forwarding Adjacency).

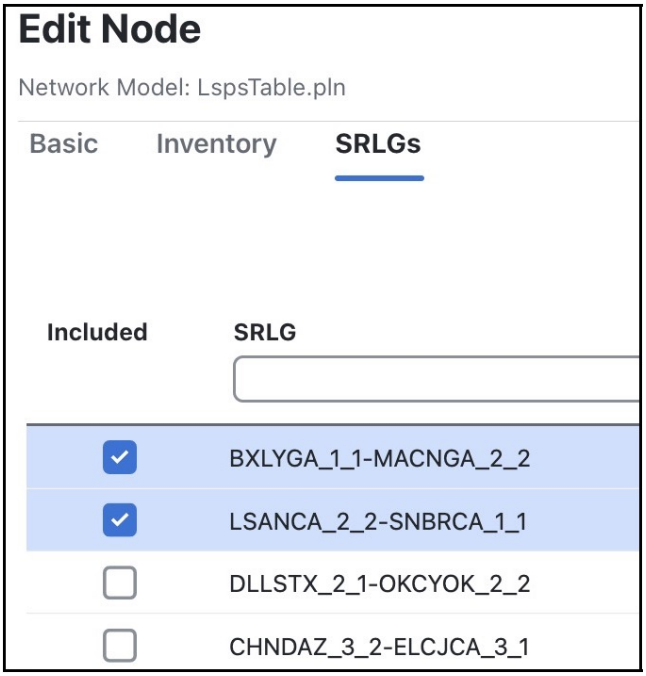
The type appears as "autoroute" or “FA” in the **Metric type** column of the LSPs table.

---

**Identify SRLGs for protection**

Provides instructions for identifying SRLGs so FRR simulations can avoid shared-risk resources when protecting circuits and related LSP path routing.

To simulate a router’s “exclude protect” configuration, associate the source node of the protected circuit with the SRLG. If that circuit fails, the FRR LSP reroutes around all circuits within that SRLG.




**Procedure**

---

1. Select the source node of a circuit within an SRLG that you want to protect. Click or choose ··· > **Edit**.
2. Click the **SRLGs** tab.
3. Choose one or more SRLGs with which you want to associate this node.
4. Save your changes.

---

## FRR LSPs initializer

Details how the FRR LSPs initializer creates protection LSPs when eligible interfaces, protected LSPs, and simulation settings are configured correctly.

The **FRR LSPs** initializer automatically creates link-protection or node-protection FRR LSPs when these conditions are met:

- One or more LSPs have their **FRR enabled** property set (refer to Mark LSPs for protection on page 285).

- One or more interfaces on the path of these LSPs have their **FRR enabled** property set. This tells the FRR LSPs initializer how to create the FRR LSPs so that in the event of a failure, the interface would be avoided when the LSP is being rerouted.
  - This property is set in the **MPLS** tab of the Edit Interface window. The **FRR enabled** column in the Interfaces table shows "true" if the interface is included or "false" if it is not.
  - After FRR LSPs are created, the interface is listed in the **FRR interface** column in the LSPs table.

The initializer creates the FRR LSPs based on these FRR enabled properties and based on whether you choose to create link or node protection.

- Link protection: Cisco Crosswork Planning creates FRR LSPs for each pair of next-hop nodes (two connected nodes) in the protected LSP path where the egress interface of the first hop has its FRR enabled property set.
- Node protection: Cisco Crosswork Planning creates FRR LSPs between each set of next next-hop nodes (three connected nodes) in the protected LSP path where the egress interface of the first hop has its FRR enabled property set.

These new FRR LSPs are named FRR_*<source>* _*<destination>* _*<postfix>*, where *postfix* is optional and set in the initializer's page.

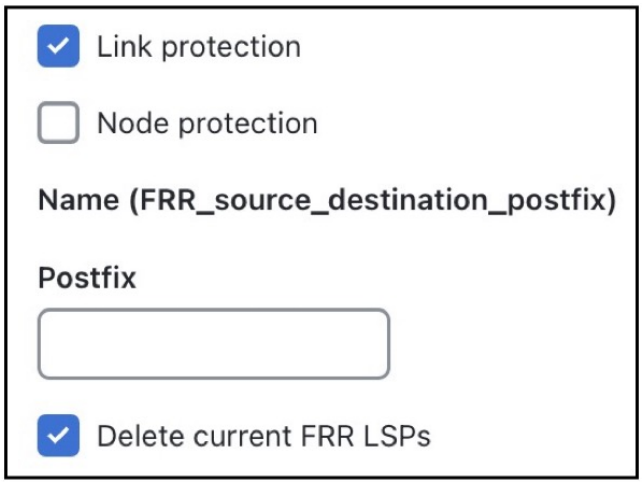
### Run the FRR LSPs initializer

Provides instructions for running the FRR LSPs initializer to create link-protection or node-protection FRR LSPs for eligible protected LSPs in simulations.

Use this task to run the FRR LSPs initializer and create protection LSPs for eligible protected LSPs.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the toolbar, choose **Actions** > **Initializers** > **FRR LSPs**.
3. Select the nodes you want to optimize. If you do not choose any nodes, the initializer uses all nodes in the plan file. Click **Next**.
4. Choose whether to create FRR LSPs that protect links (circuits), nodes, or both.




5. (Optional) Enter a postfix to add to the end of the names of the newly created FRR LSPs.
6. (Optional) Uncheck the **Delete current FRR LSPs** check box to retain the existing FRR LSPs.
7. Submit your changes.

## Create FRR LSPs manually

Provides instructions for manually creating FRR LSPs and setting the properties required for link protection or node protection during simulations.

Use this task to create FRR LSPs manually and set the properties required for link or node protection.

This section describes the required properties for an FRR LSP. It does not describe all the options available when creating an LSP.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. Create an LSP using one of these options:
   - In the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP**.
   - In the Network Summary panel, click  > **LSPs** in the LSPs page.
   - You can also choose an existing LSP to modify to be an FRR LSP. If doing so, click  or choose ··· > **Edit**.
3. Choose the source site and node of the FRR LSP.
4. Choose the destination site and node, or enter a NetInt destination IP address. For information about NetInt destination IP addresses, refer to Unresolved LSP destinations and hops on page 316.
5. In the **Routing** area, choose **FRR** as the Routing type.
6. Specify the interface that is to be avoided during routing by choosing it from the **FRR interface** drop-down list or by entering its IP address in the **NetInt FRRInterface** field.
7. Save your changes.

## Run fast reroute simulation

Provides instructions for running Fast Reroute simulations to evaluate how FRR LSPs route traffic around selected failures in RSVP-TE networks.

FRR LSPs are routed only when simulating failures in the Fast Reroute simulation mode. Complete these steps to enable the Fast Reroute simulation mode.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

   The Network options page opens.
3. Click the **Simulation** tab.
4. In the **Simulation convergence mode** section, choose **Fast reroute** from the Layer 3 drop-down list.
5. Save your changes.

# Autobandwidth-enabled LSPs

Introduces autobandwidth-enabled LSP simulations and outlines how convergence behavior changes with traffic levels and failure conditions during RSVP-TE routing analysis.

In a network, autobandwidth-enabled LSPs reset their setup bandwidth periodically based on the configured autobandwidth timers, and establish new routes if necessary. In networks with constrained bandwidth, changes in LSP routes can make other LSPs unroutable, which can then affect the amount of traffic in each LSP, the setup bandwidth settings, and the routes taken, without ever converging.

While Cisco Crosswork Planning does not simulate this instability, it does simulate autobandwidth-enabled LSPs. In this simulation mode, internally Cisco Crosswork Planning first routes the autobandwidth-enabled LSPs with a setup bandwidth of zero. Then, demands are routed to determine the simulated traffic (Traff sim) through each LSP. The result in the plan file and UI is that the Traff sim value is copied to the Setup BW sim for each autobandwidth-enabled LSP, and these LSPs are then routed using their new Setup BW sim value. Note that if some LSPs cannot be routed, the Traff sim and Setup BW sim values resulting from this process might not match one another for all autobandwidth-enabled LSPs.

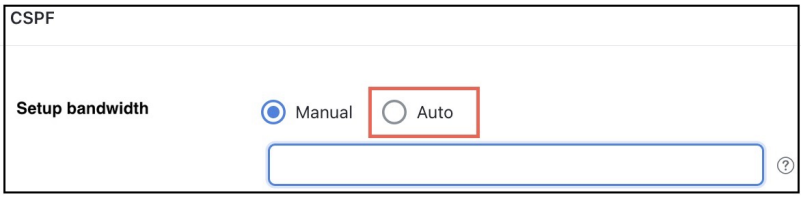
### Autobandwidth simulation modes

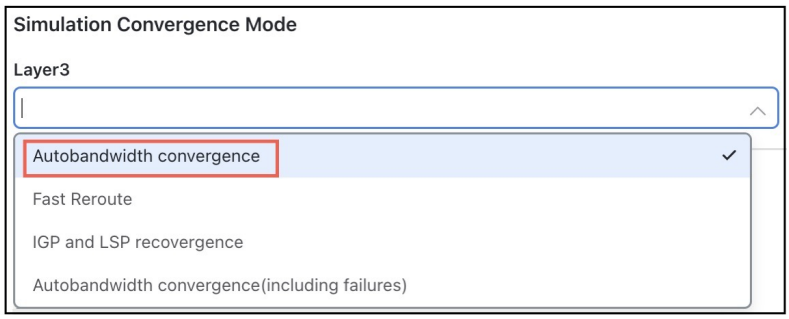
There are two autobandwidth simulation modes available. Before simulating failures or running Simulation analysis on autobandwidth-enabled LSPs, choose the appropriate mode depending on which state you need to simulate.

> **Note**
>
> These simulation modes are only available in plan files with a single traffic level.

- **Autobandwidth convergence**: In this mode, the network is simulated after the traffic (Traff sim) has been rerouted to other LSPs, but before the Setup BW sim values have been reset. This simulates the immediate, less optimal response to failures.
- **Autobandwidth convergence (including failures)**: In this mode, the network is simulated after the Setup BW sim values have been reset. This simulates the longer, more optimal response to failures.

These autobandwidth simulations require that the LSPs have their Autobandwidth property set to true, and that one of the two autobandwidth modes be selected from the Network options page (in the toolbar, click or choose **Actions** > **Edit** > **Network options**).



## Autobandwidth convergence example with and without failures

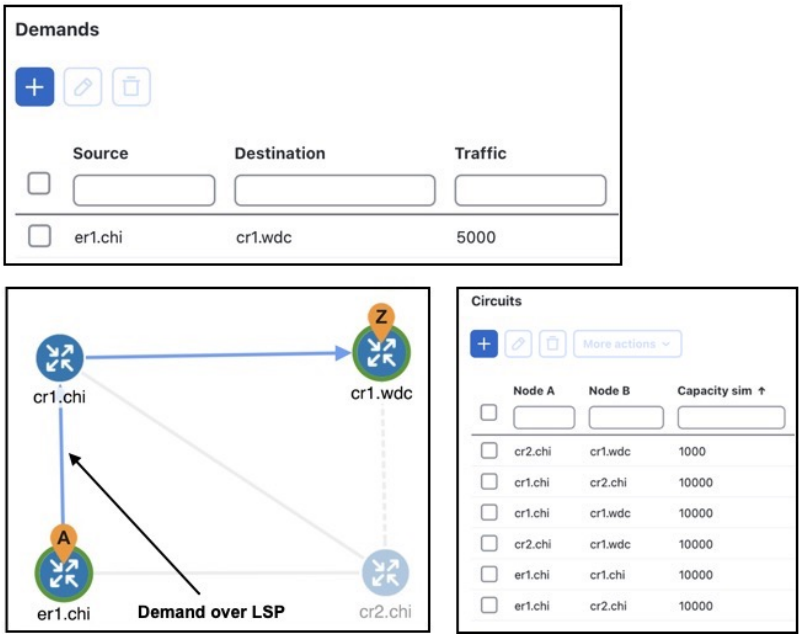
Provides an example comparing LSP setup bandwidth behavior under autobandwidth convergence with and without network failures.

This example describes how autobandwidth convergence behaves in a sample MPLS network under normal operation and when failures occur. The scenario highlights key network parameters, routing decisions, and bandwidth allocation across two LSPs.

### Network parameters

The network in this example has these parameters:

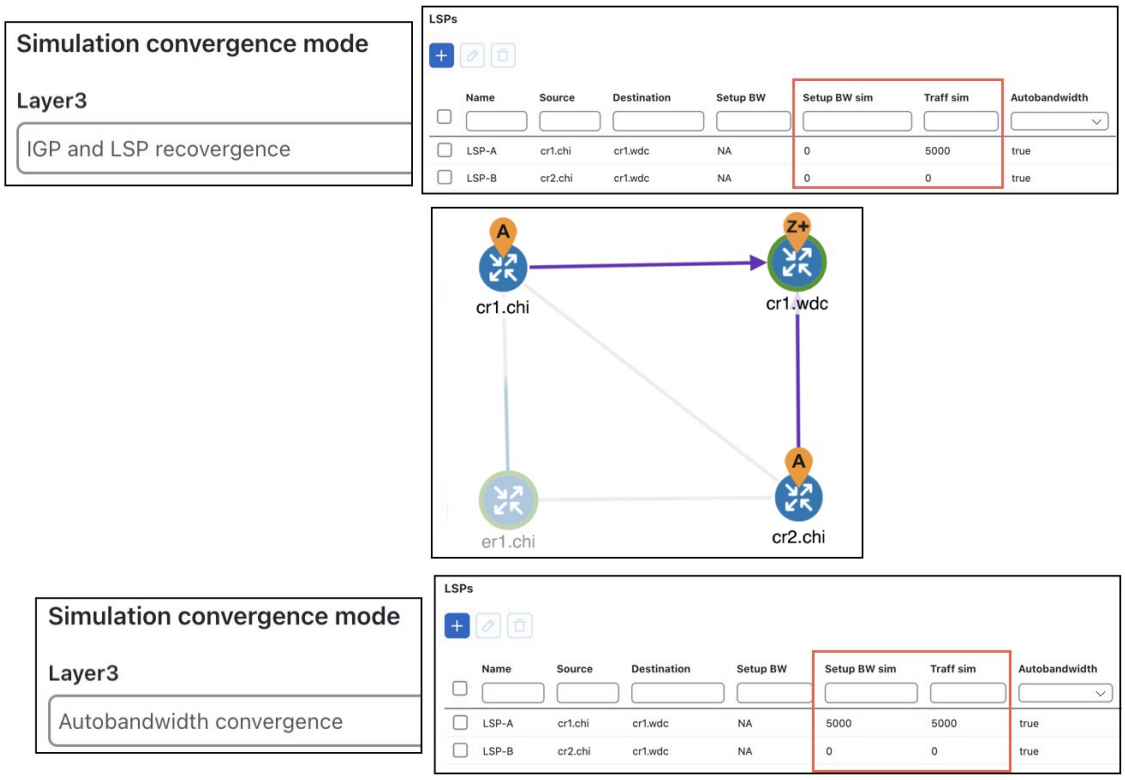
- Demand: One demand originates from er1.chi and terminates at cr1.wdc.
- Two LSPs: LSP-A from cr1.chi to cr1.wdc and LSP-B from cr2.chi to cr1.wdc.
- Setup bandwidth: Neither LSPs has a Setup BW.
- Circuit capacities: All circuits have a 10,000 Mbps capacity except for one (between cr2.chi and cr1.wdc), which is only 1000 Mbps.




### Scenario comparisons

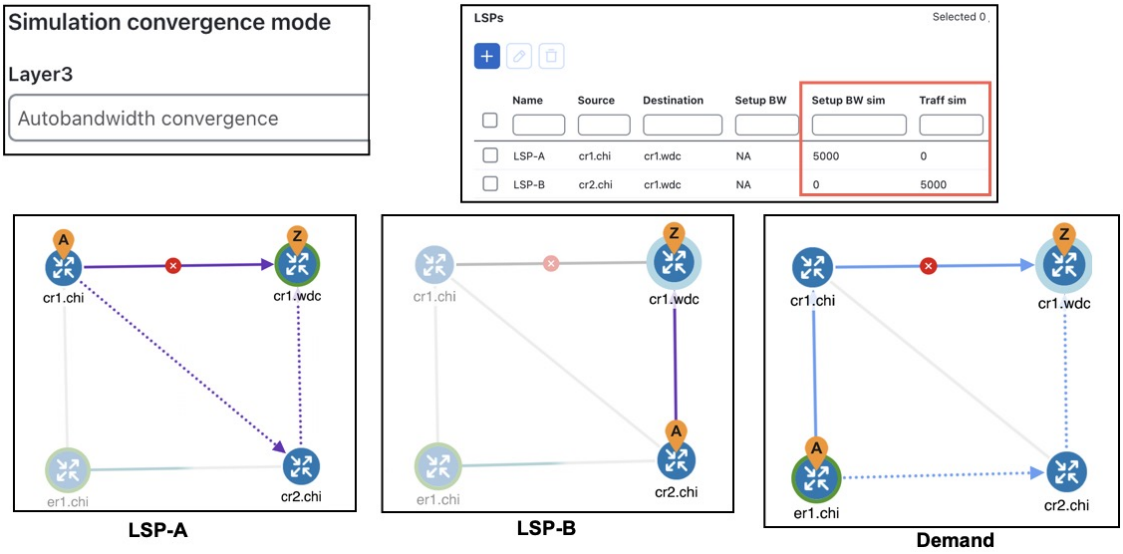
- **Normal operation (no failures):** Figure 85: Example LSP routing using IGP and LSP reconvergence and using autobandwidth convergence on page 290 shows that in the IGP and LSP Reconvergence mode, the Setup BW sim value for both LSPs is 0.
  - In the Autobandwidth convergence mode, the Setup BW sim value for LSP-A is copied from its Traff sim value, which is 5000.
  - There are no failures, and each LSP follows the same route in both convergence modes.

**Figure 85: Example LSP routing using IGP and LSP reconvergence and using autobandwidth convergence**



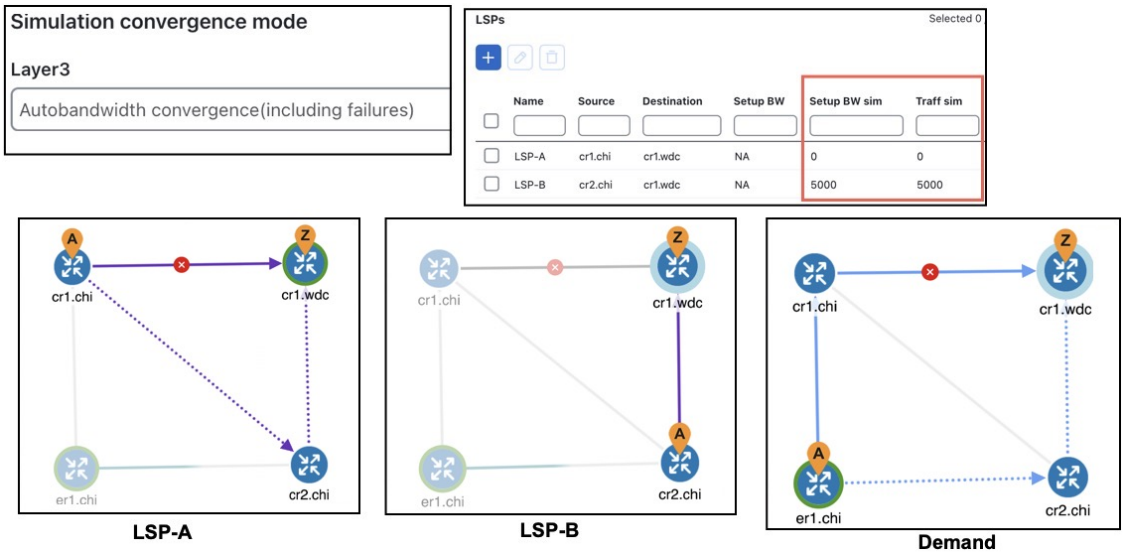

- **Failure event (autobandwidth convergence mode)**: Figure 86: Example autobandwidth convergence with failures on page 291 shows the Autobandwidth convergence mode when there is a failure.
  - Failing the circuit between cr1.chi and cr1.wdc causes LSP-A to reroute through cr2.chi.
  - The shortest path from er1.chi through cr1.chi is now longer than the shortest path from er1.chi through cr2.chi, and so the demand moves over to LSP-B. This is shown in the LSPs table where the traffic (Traff Sim) of LSP-A is moved to LSP-B.
  - Because the smaller of the two circuits is congested, LSP-A uses the larger circuit, which has enough capacity to carry the traffic.
  - LSP-B continues to route on the lower capacity circuit, thus causing congestion.
  - The Setup BW sim value is not updated for either LSP because the Autobandwidth convergence mode does not update setup bandwidth as a result of traffic shifts due to failure.

**Figure 86: Example autobandwidth convergence with failures**




- **Failure event (autobandwidth convergence including failures mode):** Figure 87: Example autobandwidth convergence including failures on page 292 shows the same failure in Autobandwidth convergence (including failures) mode.
  - The LSP-B Setup BW sim is updated to the Traff sim value, forcing it to take the larger circuit to reach its destination. Note that the plot shows both the normal (solid) and failure (dotted) route for LSP-B, indicating that it has rerouted under the failure, even though the failure did not interrupt its path.
  - The demand continues to route on LSP-B, but because the LSP has found a path with sufficient capacity, there is no more congestion.

**Figure 87: Example autobandwidth convergence including failures**




## Circuit Style RSVP LSPs

Introduces CS-RSVP LSP simulations and outlines paths, restore calculation, biased routing, creation, optimization, diagnostics, reporting, examples, and visualization workflows for planning.

Circuit Style RSVP (CS-RSVP) LSPs are logical entities that bundle two uni-directional RSVP LSPs with the same endpoints to form bi-directional RSVP LSPs. This allows traffic to consistently travel in both directions between the endpoints. Cisco Crosswork Planning supports creating and visualizing CS-RSVP LSPs.

## CS-RSVP LSP key features

Outlines key CS-RSVP LSP features and requirements.

These are the key features of CS-RSVP LSPs:

- **Persistent end-to-end traffic engineered paths**: provides predictable and identical latency in both directions.
- **Strict bandwidth commitment**: ensures no impact on the SLA due to changing network load from other services.
- **End-to-end protection and restoration mechanism**: implements robust protection and restoration mechanisms.
- **Monitoring and maintenance of path integrity**: continuously monitors and maintains the integrity of the paths.
- **Data plane resilience**: ensures the data plane remains operational even if the control plane is down.

## CS-RSVP LSP paths

Details the working, protect, and restore paths used by CS-RSVP LSPs to support redundancy and path simulation in both directions.

To improve redundancy, a CS-RSVP LSP has four paths: Working (W), Protect (P), and the corresponding restore paths, named WR and PR respectively.

The paths are categorized as W, P, and restore paths according to the "Path option" field in the LSP paths table. From the lowest to the highest path option, they are classified as W, P, WR, and PR, respectively.

- Working path (W): This is the path with the lowest path option value.
- Protect path (P): This is the path with the second lowest path option value.
- Restore paths (WR and PR): These paths are calculated dynamically when the corresponding W or P is down, considering disjointness constraints. By default, node-disjoint is considered. If node-disjointness is not possible, it automatically falls back to link-disjointness. If an SRLG-disjoint is specified, it takes precedence over node and link disjointness. If the system cannot find a disjointed restore path, it keeps the corresponding CS-RSVP LSP operational without backup path protection.

  Restore paths are calculated based on the "Protected by" field of the LSP Path. For more information, refer to Restore path calculation on page 293.

Cisco Crosswork Planning ensures that both W and P paths have a fully explicit path defined in the named path.

### Restore path calculation

Details restore path calculation for CS-RSVP LSPs and explains how restore paths protect working and protect paths during routing.

A restore path calculation is a method in CS-RSVP LSP configuration that

- determines how traffic is rerouted when a primary LSP path fails
- identifies which LSPs serve as working, protect, and restore paths under different failure scenarios, and
- uses attributes like **Protected by** and **Routed** to control path behavior.

The idea of a restore path is to protect the path that carries traffic. This is defined in the **Protected by** field of the LSP path. This field specifies which alternative LSPs can take over if the primary (working) path fails. The "Routed" field shows whether a specific LSP is actively routing traffic.

### Example

Figure 88: Restore path calculation example on page 294 shows the path configuration of the four LSP paths associated with a CS-RSVP LSP. Consider LSP path with path option 1 as LSP-1, Path option 2 as LSP-2, Path option 3 as LSP-3, and Path option 4 as LSP-4. Observe that LSP-1 and LSP-2 have the **Routed** option as "true", and LSP-3 and LSP-4 have this option as "false".

These steps describe how the restore path is calculated:

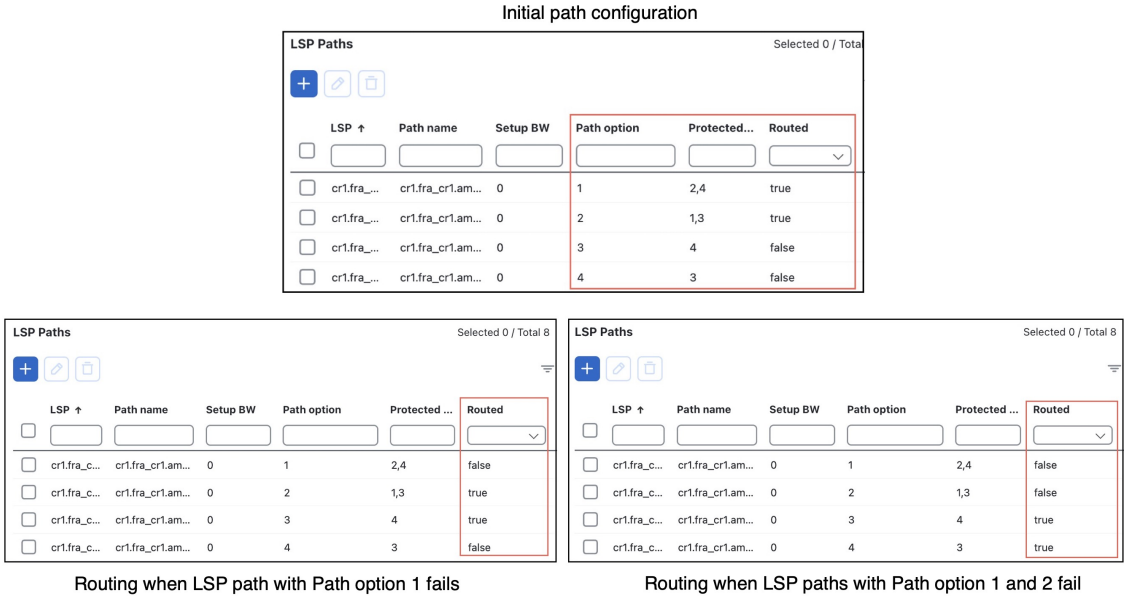
1. If LSP-1 fails, the traffic shifts to LSP-2 and LSP-3 becomes the Protect path. This happens because the **Protected by** field for LSP-1 is set to **2, 4**, indicating that when LSP-1 goes down, LSP-2 becomes the Working path and LSP-4 becomes the restore path. Since LSP-2 has the **Protected by** field as **1, 3** and since LSP-1 is already down, the LSP-3 becomes the Protect path.

   Notice that the **Routed** column value for LSP-1 changes to "false", and LSP-2 and LSP-3 have this value as "true".

2. Now, if LSP-2 goes down, traffic shifts to LSP-3 and LSP-4 becomes the Protect path. This happens because the LSP-2 has **1, 3** as the **Protected by** option. Since LSP-1 is down as indicated in Step 1, LSP-3 becomes the Working path. LSP-4 becomes the Protect path because LSP-3 has the **Protected by** field as 4.

   Observe that the **Routed** column value for LSP-2 changes to "false", and LSP-3 and LSP-4 have this value as "true".

**Figure 88: Restore path calculation example**




## How CS-RSVP LSP path simulation works

Outlines how CS-RSVP LSP path simulation checks explicit paths, disjointness, co-routing, and restore path requirements before determining routability for each path.

### Summary

Cisco Crosswork Planning performs these checks to determine if the path is routable:

### Workflow

1. The explicit paths go through the same hops in both directions. This check applies only to restore paths routed dynamically.

2. The explicit path can reach the endpoints in both directions. Note that the configuration check is not performed.
3. The path has enough bandwidth to satisfy the setup bandwidth requirement. Note that asymmetric configuration is not checked.
4. If W or P is down, Cisco Crosswork Planning calculates the corresponding restore path. If the **Biased routing** option in the Network Options page is enabled, the restore path is calculated by reusing the existing reserved resources. If **Biased routing** is disabled, the restore paths are calculated regularly. For more information, refer to Configure biased routing on page 295.
5. When routing an LSP path, the route follows the hops defined in the named path. W and P have fully explicit path listed in the named path. If the named path is not explicit or unroutable, the LSP path becomes unrouted. Paths are found dynamically only for restore paths.
6. When routing the restore path, node disjointness and link disjointness from the existing path carrying traffic are considered by default. If node disjointness is not possible, then link disjointness is considered. SRLG disjointness is considered if SRLG disjointness is configured at the LSP level.

#### Configure biased routing

Provides instructions for setting biased routing options.

Use this task to set biased routing options that influence CS-RSVP LSP restore path calculation during simulations.

In normal routing, restore path is not required to follow the corresponding W or P path as much as possible. The restore path is routed regularly with the constraint of disjointness.

In Biased Routing, the restore paths (WR and PR) reuse the corresponding W and P paths as much as possible, while ensuring that the total metric of the restore path does not exceed the maximum metric constraint. If biased routing is disabled, the restore paths are calculated regularly.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. In the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

   The Network options page opens.
3. Under the **Simulation** tab, in the **Metric constraint** section, enable biased routing.
   a) Check the **Biased routing** check box.
   b) In the **Maximum metric constraint** field, enter the maximum allowable metric for the restore path. It can be set to an extremely large value if this constraint is not required.
4. Save your changes.

## Create CS-RSVP LSPs

Provides instructions for creating CS-RSVP LSPs by configuring participating RSVP LSPs, co-routing settings, association IDs, and path roles for simulations.

Use this task to create CS-RSVP LSPs by configuring participating RSVP LSPs, co-routing settings, association IDs, and path roles.

You can create CS-RSVP LSPs in the Cisco Crosswork Planning Design application either manually or with the **CS-RSVP LSP optimizer** tool.

This task describes the manual method. For the optimizer method, refer to Create CS-RSVP LSPs using the optimizer tool on page 296.

### Before you begin

In the Add/Edit LSP page of the two RSVP LSPs participating in the CS-RSVP LSP, ensure that:

- The **Co-routed** option is enabled. This option is available in the **CSPF** section of the Add/Edit LSP page.
- They have the same Association ID.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **LSPs** > **CS-RSVP LSP**.

   Alternatively, in the Network Summary panel on the right side, click  in the **CS-RSVP LSPs** tab.

   The CS-RSVP LSPs tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon () and check the **CS-RSVP LSPs** check box.
3. In the **Node A** and **Node B** sections, choose the endpoints of the CS-RSVP LSP.
4. Enter the details in these fields.
   - **Association ID**: This field is used to associate two RSVP LSPs. This is a required field.
   - **Source address**: This field represents the source address for the bidirectional association. This is an optional field. If you do not enter a value, it defaults to an empty value.
   - **Global ID**: This field specifies the global ID for the association's global source. This is an optional field. If you do not enter a value, it defaults to an empty value.
   - **Routed**: This field indicates whether the CS-RSVP LSP is routable. It derives its value based on specific criteria and cannot be manually set.
5. Click **Add**.

## CS-RSVP LSP optimizer

Details how the CS-RSVP LSP optimizer creates, finds, and optimizes CS-RSVP LSPs based on endpoint and association settings for simulations.

**CS-RSVP LSP optimizer** is a Cisco Crosswork Planning tool that

- helps to create, find, and optimize CS-RSVP LSPs based on specified endpoint pairs, association ID, source address, and global ID
- optimizes both existing and newly created CS-RSVP LSPs together, and
- allows you to choose to optimize only the existing LSPs or only the newly created ones

### Create CS-RSVP LSPs using the optimizer tool

Provides instructions for using the optimizer tool to create CS-RSVP LSPs from endpoint, association, source address, and global ID settings.

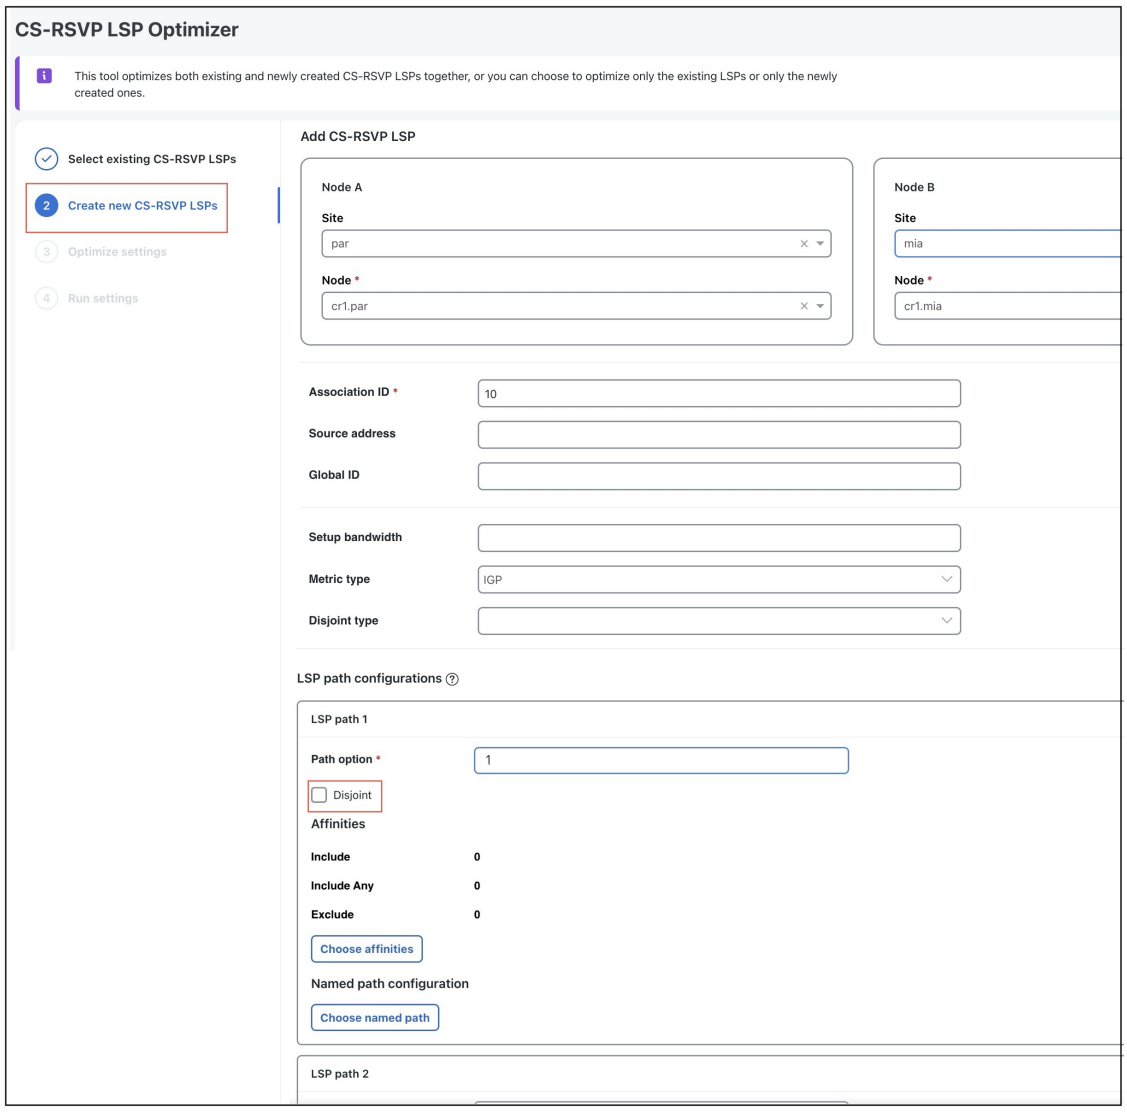
If CS-RSVP LSPs do not exist, the optimizer creates the corresponding LSPs based on the given configurations. The configuration process involves specifying several key parameters as described in this procedure.

The tool creates two LSP paths as W and P for each LSP. It creates the hops of each path as the named path of the corresponding LSP path. In addition to W and P paths, it also creates two additional paths, WR and PR. These additional paths do not have named paths and are associated with their respective W and P paths.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Tools** > **CS-RSVP LSP optimizer**.
3. (Optional) If you want to optimize the existing CS-RSVP LSPs, select them and click **Next**.
4. Click **Create** in the Create new CS-RSVP LSPs page.
5. Enter the details in the required fields. In the **LSP path configurations** section, define the configurations for the four LSP paths: W, P, WR, and PR. For field descriptions, refer to Table 30: CS-RSVP LSP configuration parameters on page 300.

   **Figure 89: Create new CS-RSVP LSPs page**




6. Click **Save**. Then, select the newly created CS-RSVP LSP and click **Next**.

   To create additional CS-RSVP LSPs, click .

7. Select the interfaces you want to exclude from the route calculation of CS-RSVP LSPs.
8. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

   The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

   > **Note**
   >
   > Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

9. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:
   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
   - If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

---

Observe these changes:

- The newly created CS-RSVP LSPs are available on the CS-RSVP LSPs table with *NewCSRSVPLSPs* as tags.
- The corresponding LSPs are available on the LSPs table.
- The corresponding LSP paths are available on the LSP paths table. Note that the LSP paths are created based on the Path option values you entered and the **Protected by** field is automatically calculated.

### Optimize existing CS-RSVP LSPs with the optimizer tool

Provides instructions for using the optimizer tool to optimize existing CS-RSVP LSPs based on configured requirements and current LSP associations.

If CS-RSVP LSPs already exist, the optimizer can optimize them based on the current configurations and requirements. Use this task to optimize the existing CS-RSVP LSPs using the CS-RSVP LSP optimizer tool.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36) with CS-RSVP LSPs. It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Tools** > **CS-RSVP LSP optimizer**.
3. Select the existing CS-RSVP LSPs you want to optimize and click **Next**.
4. (Optional) If you want to create new CS-RSVP LSPs, click **Create** and configure the required parameters. For more information, refer to Create CS-RSVP LSPs using the optimizer tool on page 296.
5. Select the interfaces you want to exclude from the route calculation of CS-RSVP LSPs.
6. Click **Next**.
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

   The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

   > **Note**
   >
   > Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:
   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
   - If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

Observe that the selected CS-RSVP LSPs are updated and have the tags*UpdatedCSRSVPLSPs*.

**CS-RSVP LSP configuration parameters**

Details the CS-RSVP LSP configuration parameters used by the optimizer tool, including endpoint, path, bandwidth, and association settings for creation.

This table lists the parameters to configure when creating CS-RSVP LSPs using the optimizer tool.

**Table 30: CS-RSVP LSP configuration parameters**

| Field | Description |
|---|---|
| Node A and Node B | These fields indicate the endpoints of the LSPs corresponding to the CS-RSVP LSP. |
| Association ID | This field is used to associate two RSVP LSPs. This is a required field. |
| Source address | This field represents the source address for the bidirectional association. This is an optional field. If you do not enter a value, it defaults to an empty value. |
| Global ID | This field specifies the global ID for the association's global source. This is an optional field. If you do not enter a value, it defaults to an empty value. |
| Setup bandwidth | This field indicates the amount of traffic the source nodes requests for this CS-RSVP LSP in Mbps. This is an optional field. |
| Metric type | This field indicates the metric type for route calculation. Choose either IGP or TE as per your requirement. |
| Disjoint type | This field indicates the disjoint type to consider. Choose **SRLGs** to enable SRLG disjointness in addition to the default node and link disjointness. |
| **LSP path configurations** | |
| Path option | This field specifies the path option of the LSP paths corresponding to the CS-RSVP LSP. The path with the lowest Path option value is assigned as the W path, the next lowest as P, the third lowest as WR, and the fourth lowest as the PR. Ensure that each LSP path has a unique Path option value.<br>For example, if the path options are 1, 2, 3, and 4, then the LSP path with Path option 1 is designated as W, path option 2 as P, path option 3 as WR, and path option 4 as PR. |
| Disjoint | This check box indicates whether to consider this LSP path for disjointness with any other LSP path that has this check box enabled. Ensure that exactly two LSP paths are configured as disjoint. |
| Affinities | This is an optional field. Click **Choose affinities** to assign affinities to the LSP paths. |
| Named path configuration | This is an optional field. Click **Choose named path** to assign named path configurations to the LSP paths. |

## CS-RSVP LSP simulation diagnostics

Outlines the CS-RSVP LSP simulation diagnostics checks for configuration consistency, disjointness, path requirements, and restoration behavior before simulation analysis runs.

The **CS-RSVP LSP simulation diagnostics** tool in Cisco Crosswork Planning performs configuration and disjointness checks on the CS-RSVP LSPs.

### Configuration check

### Disjointness check

The configuration check ensures that these properties of a given CS-RSVP LSP are identical on the associated RSVP LSPs:

- Association ID, Source address, and Global ID
- Setup bandwidth
- LSP type and LSP path count
- Protected by and Dynamic field configurations
- Hold priority and setup priority
- Named path, ensuring that Working and Protect paths have named paths
- Source and Destination
- Path options on LSP paths

The disjointness check ensures that all CS-RSVP LSPs in the network model are disjoint, meaning they do not share certain resources. This helps to maintain redundancy and avoid single points of failure.

The tool summarizes these violations:

- SRLG disjointness violation (only for CS-RSVP LSP with SRLG disjointness configuration)
- Node disjointness violation
- Link disjointness violation
- Node to link disjointness fallback

### Run the CS-RSVP LSP simulation diagnostics tool

Provides instructions for running the CS-RSVP LSP simulation diagnostics tool to identify configuration, association, and path issues before optimization workflows.

Use this task to run the CS-RSVP LSP simulation diagnostics tool and identify configuration, association, and path issues.

### Procedure

1. Open the plan file (refer to Open plan files on page 36) with CS-RSVP LSPs. It opens in the Network Design page.
2. From the toolbar, choose **Actions** > **Tools** > **Diagnostics** > **CS-RSVP LSP simulation**.
3. Click **Proceed** in the information dialog box.

A report is generated indicating the details of the failures.

### CS-RSVP LSP simulation diagnostics report

Details the CS-RSVP LSP simulation diagnostics report and explains how generated results identify configuration and path validation issues for review.

Each time the CS-RSVP LSP simulation diagnostics tool is run, a report is automatically generated. You can access this information at any time by choosing **Actions** > **Reports** > **Generated reports** and then clicking the **CS-RSVP LSP Simulation** link in the right panel. Note that new reports replace the previous ones.

The **Summary** tab provides information on the number of failed CS-RSVP LSPs.

The **CS-RSVL LSP simulation** tab provides failure details. Expand the **FAILUREMESSAGE** column to view the failure messages.

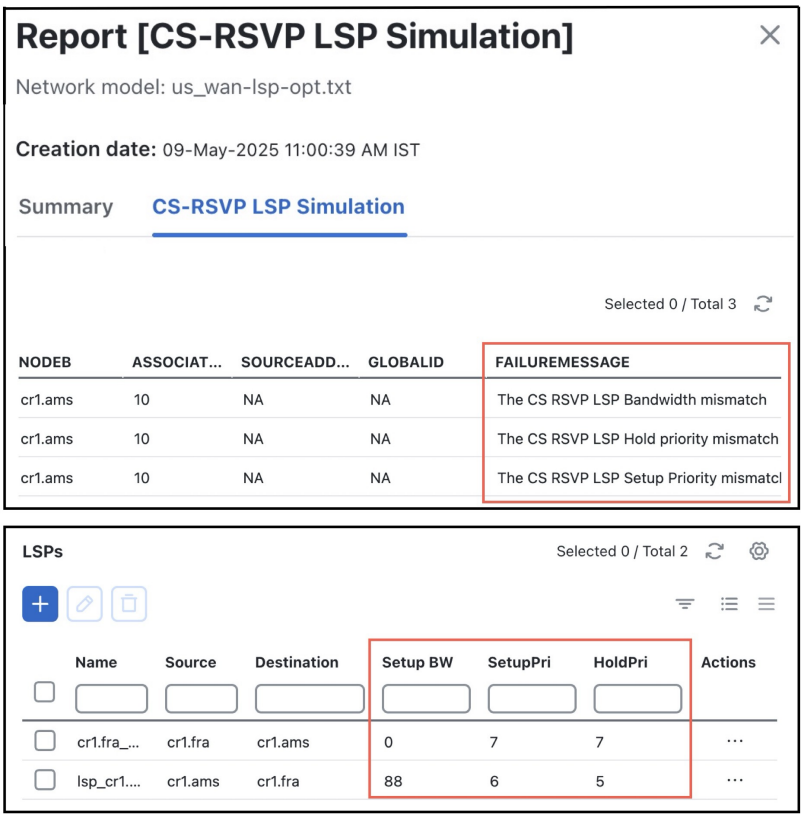
### CS-RSVP LSP simulation diagnostics example

Details an example of CS-RSVP LSP diagnostics that shows how configuration differences and disjointness issues appear in diagnostics analysis reports.

The plan file has a CS-RSVP LSP with configuration issues, such as associated LSPs having different Setup priority, Hold priority, and Setup bandwidth. After running the CS-RSVP LSP simulation diagnostics tool, the generated report shows this information in the **FAILUREMESSAGE** column.

By filtering from the CS-RSVP LSPs table to the LSPs table, you can refer to that the values for these fields do not match.

**Figure 90: CS-RSVP LSP simulation diagnostics report and mismatching LSPs**




## CS-RSVP LSP visualization

Outlines CS-RSVP LSP visualization relationships and explains how selected CS-RSVP LSPs, associated LSPs, and LSP paths appear in network views.

CS-RSVP LSPs are associated with LSPs and LSP paths. This section describes how to visualize:

- CS-RSVP LSPs
- LSPs associated with CS-RSVP LSPs
- LSP paths associated with CS-RSVP LSPs

### Visualize CS-RSVP LSPs

Provides instructions for visualizing selected CS-RSVP LSPs in the network plot.

Use this task to visualize selected CS-RSVP LSPs in the network plot.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36) with CS-RSVP LSPs. It opens in the Network Design page.
2. In the Network Summary panel on the right side, navigate to the **CS-RSVP LSPs** tab.

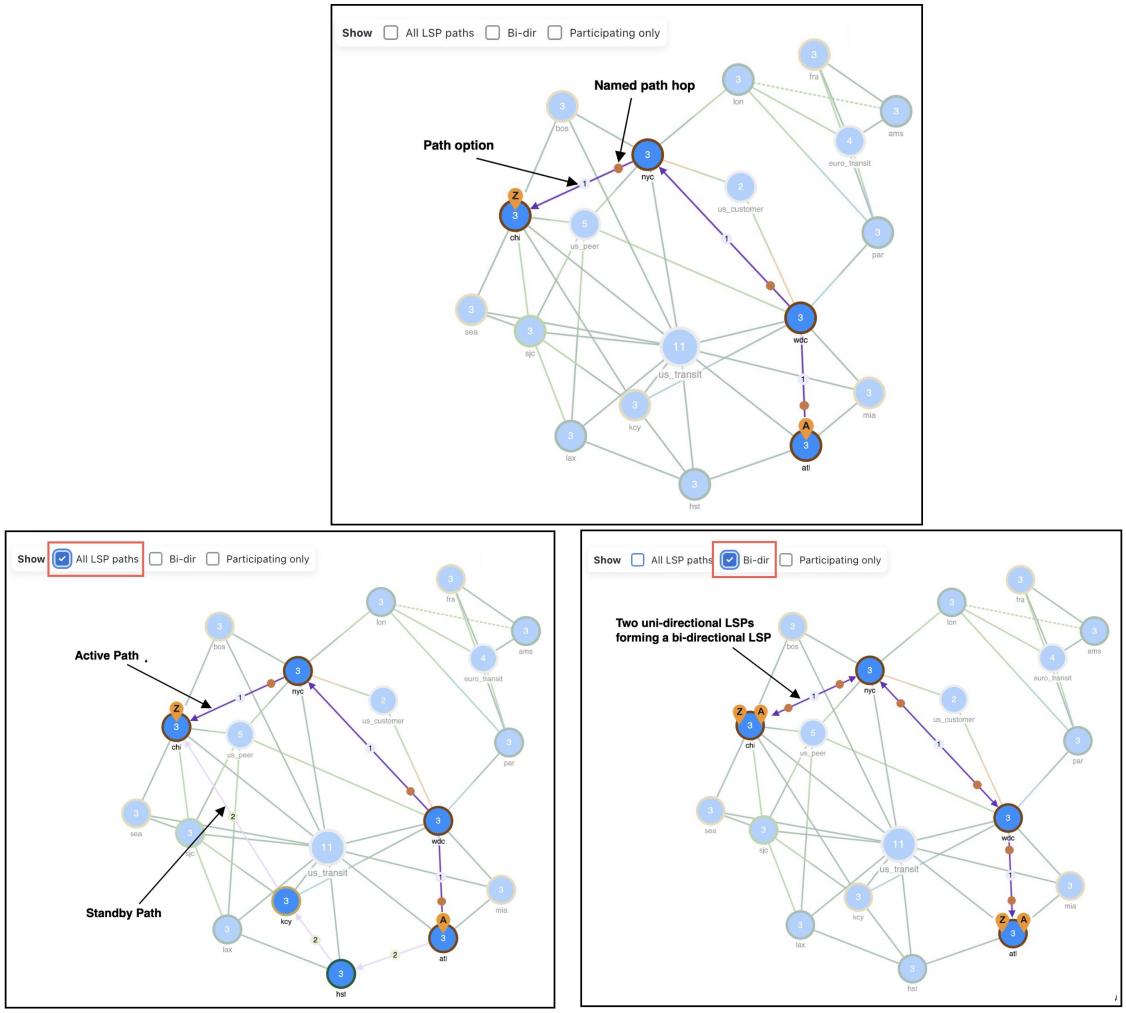
   The CS-RSVP LSPs tab is available under the **More** tab. If it is not visible, then click the Show/hide tables icon ( ) and check the **CS-RSVP LSPs** check box.

   > **Note**
   >
   > You can select only one CS-RSVP LSP at a time to view on the network plot.

3. Select the required CS-RSVP LSP from the list.
4. Observe these in the network plot on the left side:
   - Both LSPs associated with the CS-RSVP LSP are plotted. The purple arrow indicates the LSP path is active (carrying traffic). The number on the path (here 1) is the path option. For example, refer to the first image in Figure 91: CS-RSVP LSP path visualization on page 303.

     A dim purple arrow indicates the standby path, which is an LSP path that is not a primary path and is set as a Standby path. To view the standby paths, check the **All LSP paths** check box in the network plot. For example, refer to the second image in Figure 91: CS-RSVP LSP path visualization on page 303.
   - The bi-directional LSPs are plotted when you check the **Bi-dir** check box in the network plot. For example, refer to the third image in Figure 91: CS-RSVP LSP path visualization on page 303.

**Figure 91: CS-RSVP LSP path visualization**




## Visualize LSPs associated with CS-RSVP LSPs

Provides instructions for visualizing RSVP LSPs associated with selected CS-RSVP LSPs in the network view for relationship analysis and verification.

Use this task to visualize RSVP LSPs associated with selected CS-RSVP LSPs.

### Procedure

1. Open the plan file (refer to Open plan files on page 36) with CS-RSVP LSPs. It opens in the Network Design page.
2. In the Network Summary panel on the right side, navigate to the **CS-RSVP LSPs** tab and select the required CS-RSVP LSP from the list.

   **Note**

   You can select only one row at a time.

3. Choose > **Filter to LSPs**. The LSPs table opens, displaying the LSPs associated with the selected CS-RSVP LSP.

4. Select the LSPs.

> **Note**
>
> You can select up to three LSPs at a time.

5. Observe these in the network plot on the left side:
   - All LSP paths which belong to the LSPs and are routed are plotted.
   - The active path is distinguishable from the other alternate or standby paths.
   - If the Working (W) path fails and traffic is rerouted, both the no-fail path and fail (rerouted) path are plotted.

---

### Visualize LSP paths associated with CS-RSVP LSPs

Provides instructions for visualizing LSP paths associated with selected CS-RSVP LSPs in the network view for path analysis and verification.

Use this task to visualize LSP paths associated with selected CS-RSVP LSPs.

#### Procedure

---

1. Open the plan file (refer to Open plan files on page 36) with CS-RSVP LSPs. It opens in the Network Design page.
2. In the Network Summary panel on the right side, navigate to the **CS-RSVP LSPs** tab and select the required CS-RSVP LSP from the list.

> **Note**
>
> You can select only one row at a time.

3. Choose > **Filter to LSP paths**. The LSP Paths table opens, displaying the LSP paths associated with the selected CS-RSVP LSP.
4. Select the LSP paths to visualize in the network plot.

> **Note**
>
> You can select up to three LSP paths at a time.

---

# Advanced RSVP-TE LSP simulations

## Ignore reservable bandwidths for capacity planning

Provides instructions for ignoring reservable bandwidth constraints when capacity-planning simulations require traffic levels beyond current reservations and available bandwidth limits.

When using Cisco Crosswork Planning for capacity planning, it is common to increase the demand traffic to a level expected at a future date. Various simulations are then performed, such as failures and worst-case analysis, to determine which circuits will likely experience over-utilization.

In MPLS networks using RSVP-TE LSPs, a setup bandwidth is configured for each LSP. If the demand traffic is increased, LSP setup bandwidths are usually increased correspondingly (for example, by using the LSP Setup Bandwidth initializer). When these LSP setup bandwidths become too large, some LSPs might not be routable, thus rendering an improbable view of the future network. When this occurs, you have the option to temporarily remove the reservable bandwidth constraints for a more realistic view into future capacity requirements. While this is not a routing method that actually exists on routers, it is a convenient Cisco Crosswork Planning simulation that can assist you when planning networks.

When you enable this routing mode, Cisco Crosswork Planning routes LSPs in the order they appear in the plan file LSPs table. Note that static order of the LSPs table in the plan file likely differs from the order of LSPs in the GUI, which changes frequently due to user interaction.

This operates only on LSPs that do not have autobandwidth enabled.

### Procedure

1. Route LSPs as usual. For information, refer to Dynamic LSP routing and CSPF on page 264.
2. If an LSP is not routable (call it “LSP-A”), try to route it with zero setup bandwidth.
   - If LSP-A cannot be routed, leave it unrouted and try to route the next LSP in the LSPs table. If this LSP cannot be routed, try to route the next LSP listed in the LSPs table in the plan file.
   - If LSP-A can be routed, then there will be one or more interfaces on its route whose reservable bandwidth is exceeded. These interfaces are marked as having unlimited reservable bandwidth during the simulation of all subsequent LSPs, thus allowing for the routing of LSP-A and other LSPs that would otherwise exceed the reservable bandwidth on these interfaces. Note that actual Resv BW and Resv BW sim values in the Interfaces table do not change.

   On completion of the simulation, a number of interfaces will have their reserved bandwidth greater than their configured reservable bandwidth. The circuits containing these interfaces are candidates for capacity expansion.

### Ignore reservable bandwidth constraints

Provides instructions for configuring routing simulations to ignore reservable bandwidth constraints during RSVP LSP path calculation and capacity analysis workflows.

Use this task to configure routing simulations to ignore reservable bandwidth constraints during routing simulations.

### Procedure

1. From the toolbar, click  or choose **Actions** > **Edit** > **Network options**.

   The Network options page opens.

2. Click the **Simulation** tab.
3. In the **Label switched paths** section, check the **LSP capacity planning mode** check box. This ignores the reservable bandwidth constraints.
4. Save your changes.

---

### Reservable bandwidth capacity planning example

Provides an example that illustrates how reservable bandwidth constraints affect RSVP LSP routing and rerouting during simulation analysis workflows.

This example demonstrates RSVP LSP routing in a simple network and examines how reservable bandwidth influences routing outcomes during simulation analysis.

#### Network configuration

In this example, there are four LSPs, each with a setup bandwidth (Setup BW) of 600 Mbps. All interfaces have a reservable bandwidth (Resv BW) of 1000 Mbps.

- LSP details
  - LSP-1, LSP-2, and LSP-3 go from LON to FRA.
  - LSP-4 goes from LON to ROM.
- IGP metrics
  - Interfaces between PAR and FRA, and interfaces between FRA and ROM have an IGP metric of 2.
  - All other interfaces have an IGP metric of 1.

Figure 92: Comparison of reservable bandwidth being used vs not being used on page 307 shows the LSP reservations view of this simple network in both simulations modes: the usual where reservable bandwidths are observed and the capacity planning version wherein they are ignored. In this latter mode, the two red circuits indicate they are being over utilized and likely need their capacity increased. (For information about the LSP Reservations view, refer to View LSP reservations on page 266.)

**Figure 92: Comparison of reservable bandwidth being used vs not being used**

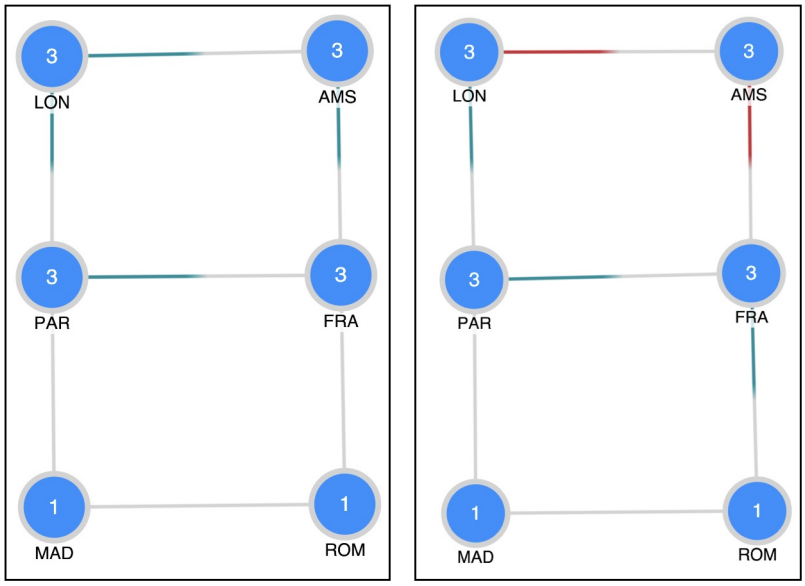
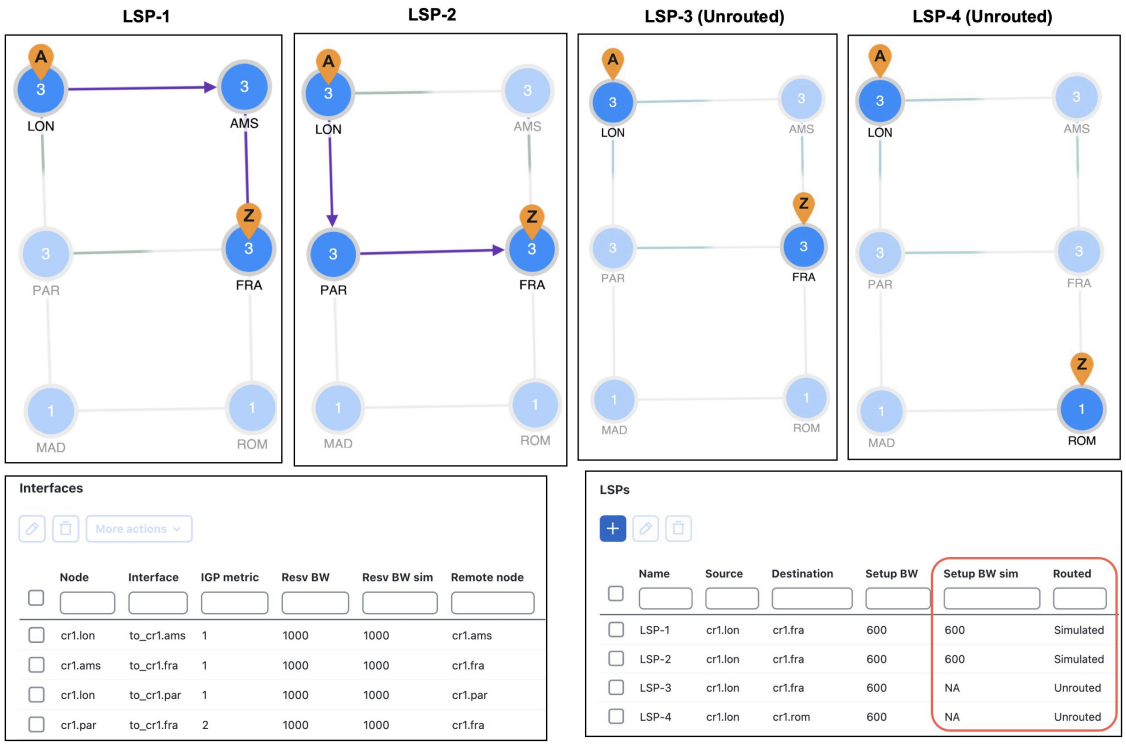



Figure 93: Simulated routing using reservable bandwidth constraints on page 308 demonstrates routing simulations where reservable bandwidth constraints are applied (**LSP capacity planning mode** unchecked), which is the default behavior.

- LSP-1 routes across LON-AMS-FRA, leaving circuits LON-AMS and AMS-FRA each with only 400 Mbps of available bandwidth.
- LSP-2 routes across LON-PAR-FRA, leaving circuits LON-PAR and PAR-FRA each with only 400 Mbps of available bandwidth.
- LSP-3 cannot be routed because there is not enough available bandwidth on LON-AMS or LON-PAR. Its Setup BW sim value in the LSPs table is set to “NA” and the Routed column is set to “Unrouted.”
- LSP-4 cannot be routed, again because there is not enough available bandwidth. Its Setup BW Sim value in the LSPs table is set to “na” and the Routed column is set to “Unrouted.”

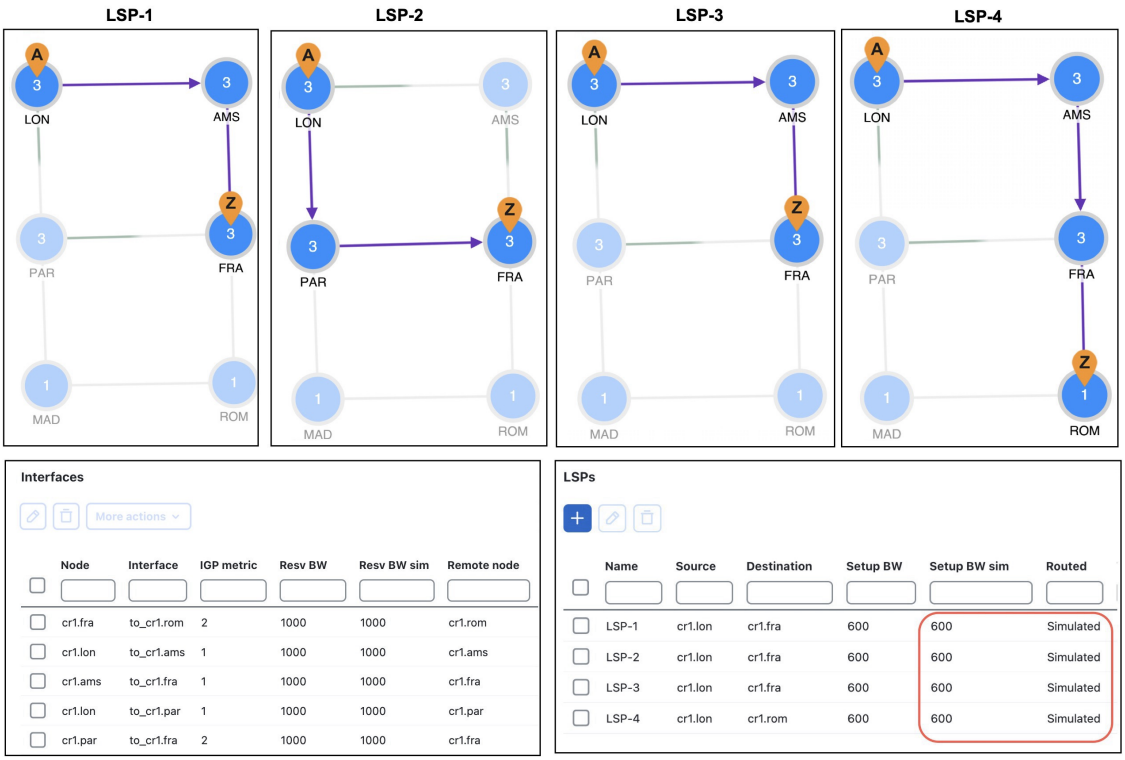
**Figure 93: Simulated routing using reservable bandwidth constraints**




In Figure 94: Simulated routing without using reservable bandwidth constraints on page 309, the use of reservable bandwidth constraints has been disabled (**LSP capacity planning mode** checked). The LSP routes are:

- As with the default behavior, LSP-1 routes across LON-AMS-FRA and LSP-2 routes across LON-PAR-FRA.
- Once it is determined that there is not enough reservable bandwidth for LSP-3 to route, it takes the shortest path, LON-AMS-FRA. The reservable bandwidth on interfaces LON-to-AMS and AMS-to-FRA is ignored for the remainder of this simulation. LSP-3's Setup BW sim value in the LSPs table is set to 600.
- LSP-4 routes across LON-AMS-FRA-ROM because these reservable bandwidth constraints have been removed and because it is the only route available. Even though LON-PAR-MAD-ROM is the shortest path, there is not enough reservable bandwidth on LON-to-PAR or PAR-to-FRA. LSP-4's Setup BW sim value in the LSPs table is set to 600.

**Figure 94: Simulated routing without using reservable bandwidth constraints**




If, however, LSP-4 were routed before LSP-3, then LSP-4 would route across its shortest path of LON-PAR-MAD-ROM, and the LON-to-PAR interface would have its reservable bandwidth constraint ignored. With this routing, LSP-3 would still route LON-AMS-FRA, and the interfaces LON-to-AMS and AMS-to-FRA would still be set to remove their reservable bandwidth constraint.

## P2MP LSPs

Introduces P2MP LSP simulations and outlines bandwidth behavior, demand routing, creation requirements, sub-LSP workflows, demand creation, and deletion behavior rules.

Point-to-multipoint (P2MP) LSPs offer an MPLS alternative to IP multicast for setting up multiple paths to multiple locations from a single source. In live networks, the signal is sent once across the network and packets are replicated only at the relevant or designated branching MPLS nodes.

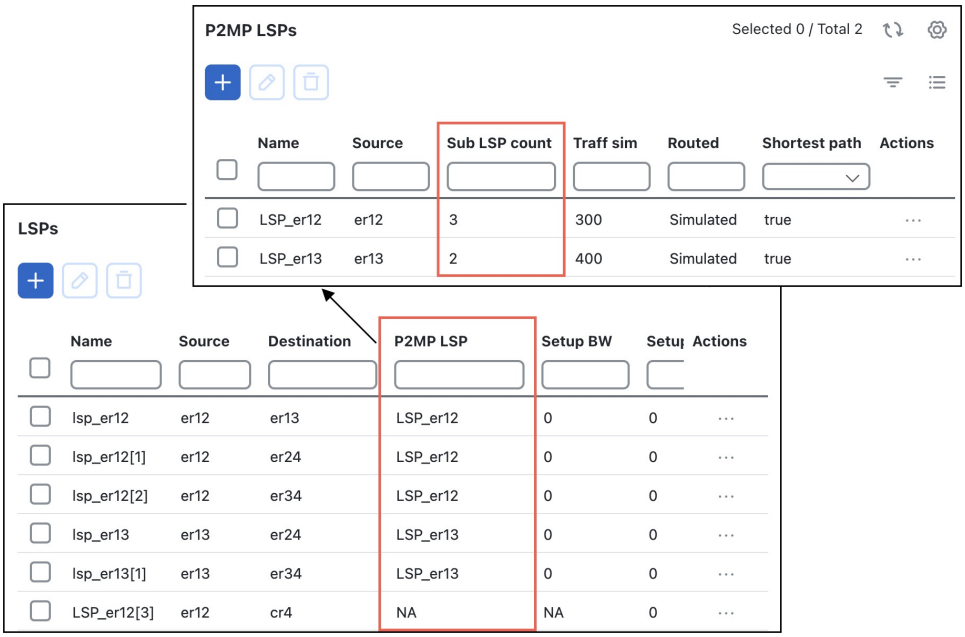
A P2MP LSP consists of two or more sub-LSPs that have the same source node. Most of Cisco Crosswork Planning refers to sub-LSPs as LSPs, and there is no distinction between the two.

To view the P2MP LSPs or their associated sub-LSPs, choose them from the P2MP LSPs and LSPs tables. To determine if an LSP belongs to a P2MP LSP, you can show the **P2MP LSP** column in the LSPs table.

### Example

In Figure 95: P2MP LSPs and associated sub-LSPs on page 310, there are two P2MP LSPs, LSP_er12 and LSP_er13. LSP_er12 and LSP_er13 have "er12" and "er13" as their source nodes, respectively. In the LSPs table, the first three LSPs have er12 as the source node and are sub-LSPs of LSP_er12 P2MP LSP (as indicated by the **P2MP LSP** column). Similarly, the LSPs with er13 as the source node are the sub-LSPs of LSP_er13 P2MP LSP. The **Sub LSP count** in the P2MP LSPs table indicates the number of sub-LSPs of each P2MP LSP.

**Figure 95: P2MP LSPs and associated sub-LSPs**




## P2MP LSP bandwidth

Details how P2MP LSP bandwidth is shared across sub-LSPs and common circuits during P2MP routing simulations and capacity analysis workflows.

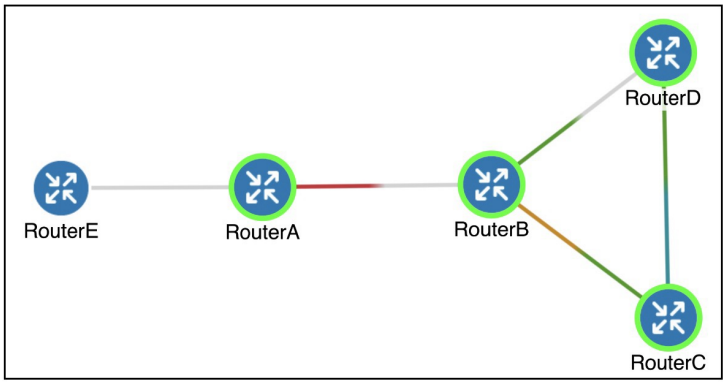
All sub-LSPs within a given P2MP LSP share bandwidth and therefore, common circuits between these sub-LSPs have the same bandwidth, rather than an aggregate. Common practice is to set all sub-LSPs within a P2MP LSP to the same bandwidth. Otherwise, the bandwidth for the shared circuit is the highest bandwidth set.

To refer to the reserved LSP bandwidth, show the LSP Resv column in the Interfaces table. You can also view the associated traffic by choosing **LSP Reservations** in the Network Plot menu in the Visualization toolbar. For more information, refer to View LSP reservations on page 266.

### Example

Example: If A-C LSP and A-D LSP each have a bandwidth 400 Mbps, the A-B circuit has an aggregate bandwidth of 800 Mbps (Figure 96: Bandwidth behavior without a P2MP LSP on page 311). However, if A-C and A-D are sub-LSPs within a P2MP LSP, the A-B circuit has a bandwidth of 400 Mbps (Figure 97: Bandwidth behavior with a P2MP LSP on page 312).

**Figure 96: Bandwidth behavior without a P2MP LSP**

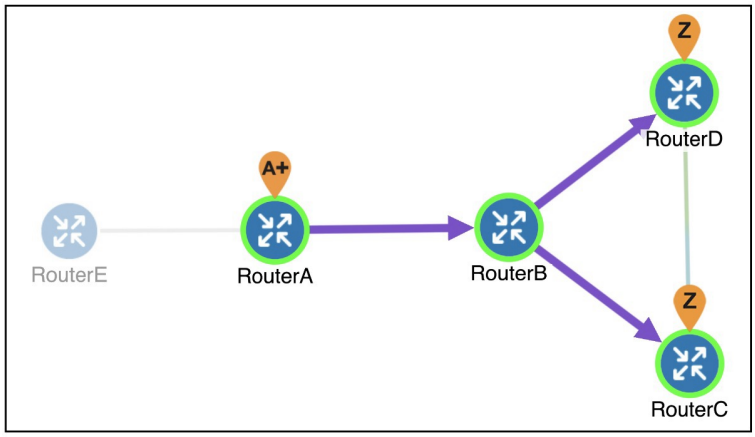



Interfaces

More actions

| Node | Interface | Remote node | Traff meas | Util meas | LSP resv ↓ |
|---|---|---|---|---|---|
| RouterA | to_RouterB | RouterB | 200 | 20 | 800 |
| RouterB | to_RouterD | RouterD | 650 | 65 | 400 |
| RouterB | to_RouterC | RouterC | 500 | 50 | 400 |

**Figure 97: Bandwidth behavior with a P2MP LSP**




Interfaces 56 | Demands 1 | Nodes 23 | LSPs 4 | More

Interfaces

More actions

| Node | Interface | Remote node | Traff meas | Util meas | LSP resv ↓ |
|---|---|---|---|---|---|
| RouterB | to_RouterD | RouterD | 650 | 65 | 400 |
| RouterB | to_RouterC | RouterC | 500 | 50 | 400 |
| RouterA | to_RouterB | RouterB | 200 | 20 | 400 |

## P2MP LSP demands

Details how demands route through P2MP LSPs and how source, destination, and sub-LSP requirements affect routing in simulations and analysis.

Demands routed through P2MP LSPs must have the same source as the P2MP LSP; the demand then terminates at each of the P2MP LSP destinations. No other demands can go through the P2MP LSP except for the one created for it. If the P2MP LSP fails, its demand cannot be routed. For information on creating these demands, refer to Create demands for P2MP LSPs on page 314.

Demands must be private to the P2MP LSP. This privacy is automatically created when you insert a demand for a P2MP LSP.

## Requirements for P2MP LSPs and sub-LSPs

Outlines requirements and guidelines for P2MP LSPs and sub-LSPs, including source, destination, bandwidth, and association rules used in simulation workflows.

Rules and guidelines:

- All sub-LSPs within a given P2MP LSP must have the same source site and source node.
- Because sub-LSPs within a P2MP LSP share setup bandwidth, typically bandwidth is set to the same value for each sub-LSP.

## Create P2MP LSPs

Provides instructions for creating P2MP LSPs by selecting sources, destinations, bandwidth settings, and related routing properties for simulations and analysis.

Use this task to create P2MP LSPs by selecting sources, destinations, bandwidth settings, and routing properties.

### Procedure

1. Create an empty P2MP LSP.
   a) From the toolbar, choose **Actions** > **Insert** > **LSPs** > **P2mp LSP**, or click  in the **P2MP LSPs** table.
   b) Enter a unique P2MP LSP name.
   c) Choose the source site and node. This selection is the root for all sub-LSPs contained in this P2MP LSP.
   d) (Optional) Enter or choose a disjoint group and disjoint priority.
2. If you have not done so, create sub-LSPs for the P2MP LSP. For more information, refer to Create sub-LSPs on page 313.
3. Associate the sub-LSPs to the P2MP LSP. You can use these steps individually or you can choose multiple LSPs and edit them collectively. If you are editing multiple sub-LSPs, verify that all of them have the same source site and node as the P2MP LSP.
   a) Select the sub-LSPs from the LSPs table and click .
   b) In the **Sub-LSP of P2MP LSP** list under the **Other** section at the bottom, choose the P2MP LSP to which you want these LSPs to belong.
   c) Click **Save**.

### Create sub-LSPs

Provides instructions for creating sub-LSPs under a P2MP LSP and configuring their destination and bandwidth settings for routing simulations and analysis.

Use this task to create sub-LSPs under a P2MP LSP and configure their destination and bandwidth settings.

**Procedure**

1. From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP** or click > **LSPs** in the **LSPs** table.
2. Choose the site and node that will be the source of the P2MP LSP. All sub-LSPs in a P2MP LSP must have the same source site and source node.
3. Choose the site and node that are the final destination for the sub-LSP.
4. Complete the remaining optional fields and tabs, as needed. Because sub-LSPs share bandwidth, typically the CSPF Setup bandwidth field is set to the same value for each sub-LSP in the P2MP LSP.
5. If you have already created the P2MP LSP to which this sub-LSP will belong, choose that P2MP LSP from the Sub-LSP of P2MP LSP list (at the bottom). For creating P2MP LSPs, refer to Create P2MP LSPs on page 313.
6. Save your changes.

### Create a mesh of sub-LSPs

Provides instructions for creating a mesh of sub-LSPs from selected source and destination nodes for P2MP routing simulations and demand planning.

Use this task to create a mesh of sub-LSPs from selected source and destination nodes for P2MP routing simulations.

**Procedure**

1. In the Nodes table, choose the node you will be designating as the source for the P2MP LSP. All sub-LSPs in a P2MP LSP must have the same source node. A P2MP LSP cannot have multiple source nodes.
2. From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP mesh** or click > **LSP mesh** in the **LSPs** table.
3. In the **Destination nodes** section, uncheck these check boxes and click **Next**.
   - **Same destinations as sources**
   - **Create LSPs from destination to source**
4. Complete the remaining optional fields as needed and click **Next**.
5. To set the bandwidth to be the same for all sub-LSPs in this mesh, choose **Fixed Value** in the **Bandwidth** drop-down list. Enter that value in the **Fixed value** field (in Mbps).

   To use the auto bandwidth feature so that these LSPs have their Setup BW sim value automatically calculated and used in auto bandwidth simulations, choose **Auto Bandwidth** in the **Bandwidth** drop-down list.
6. Complete the remaining optional fields as needed and click **Submit**.

### Create demands for P2MP LSPs

Provides instructions for creating demands that route traffic through P2MP LSPs and their associated sub-LSPs during simulations and analysis workflows.

To simulate traffic routing through a P2MP LSP, create a demand for it.

**Procedure**

1. If you want to create demands for P2MP LSPs, choose one or more P2MP LSPs from the P2MP LSP table.

2. From the toolbar, choose **Actions** > **Insert** > **LSPs** > **Demands for P2MP LSPs** or click > **Demands for P2MP LSPs** in the **LSPs** table.
3. Choose the P2MP LSPs through which you are routing traffic.
4. (Optional) Choose a service class for the newly created demand.
5. Choose the desired bandwidth: **Minimum sub-LSP Setup BW** or **Zero**.
6. Submit your changes.

## Deletion behavior for P2MP LSPs and sub-LSPs

Outlines deletion behavior for P2MP LSPs and sub-LSPs, including impacts to demands and recommended disassociation steps before cleanup operations proceed.

Deleting a P2MP LSP will invalidate the simulation and remove all demands that use the P2MP LSP as a private LSP. If you do not want all the sub-LSPs deleted, disassociate them first.

### Disassociate sub-LSPs from their P2MP LSPs

Provides instructions for disassociating sub-LSPs from P2MP LSPs.

Use this task to disassociate sub-LSPs from P2MP LSPs so selected sub-LSPs can be preserved during deletion.

#### Procedure

1. In the **LSPs** table, choose one or more sub-LSPs.
2. Click .

> **Note**
>
> If you are editing a single sub-LSP, you can also use the ••• > **Edit** option under the **Actions** column.

3. In the Sub-LSP of P2MP LSP list, perform one of these actions:
   - Choose the empty row to disassociate the sub-LSP and change it to an LSP.
   - Choose a different P2MP LSP to associate the sub-LSP with a different P2MP LSP.
4. Click **OK**.

### Delete sub-LSPs or P2MP LSPs

Provides instructions for deleting sub-LSPs or P2MP LSPs and choosing the cleanup behavior for associated objects in safe maintenance workflows.

Use this task to delete sub-LSPs or P2MP LSPs and choose the cleanup behavior for associated objects.

#### Procedure

1. To delete sub-LSPs but keep the P2MP LSP:

a) Select one or more sub-LSPs in the **LSPs** table.
b) Click . If you are deleting a single sub-LSP, you can also use the ··· > **Delete** option under the **Actions** column.
c) Click **Delete** in the confirmation dialog box to continue.

2. To delete the P2MP LSP:
a) Select one or more P2MP LSPs in the **P2MP LSPs** table.
b) Click . If you are deleting a single P2MP LSP, you can also use the ··· > **Delete** option under the **Actions** column.
c) Click **Delete** in the confirmation dialog box to continue.

## Unresolved LSP destinations and hops

Details how unresolved LSP destinations and hops appear after discovery and how they affect RSVP-TE LSP configuration review and cleanup.

RSVP-TE LSP configurations and state are read using network discovery from the source node of the LSP, through SNMP, the Parse config tool, or other methods. These configurations might reference nodes or interfaces that do not exist in the plan file, so they are referred to as being unresolved. For example, the LSP destination node, read from the configuration on the source, might not be in the plan file.

There are several reasons for these differences:

- The plan file was modified to contain fewer nodes than are in the IGP.
- Cisco Crosswork Planning cannot read all the nodes, or cannot obtain the IP addresses of all the nodes.
- The LSPs themselves are not configured correctly.

References that might not be resolved include:

- Destination nodes of the LSPs.
- Hops in named paths configured on the source node.
- Hops in the actual paths read from the source node.

Cisco Crosswork Planning tries to resolve as many of these references as possible. Tables and columns are updated:

- If the LSP destinations are resolved, they are updated in the Destination column of the LSPs table. If they are not resolved, the column remains empty. The original IP address remains in the NetInt Destination column regardless of whether the destination is resolved or not.
- If the hops are resolved, they are updated in the Node and Interface columns in the Named Path Hops and Actual Path Hops tables. If they are not resolved, the columns remain empty. The original IP address remains in the NetIntHop column regardless of whether the destination is resolved.

Cisco Crosswork Planning does not make further attempts to resolve these references unless you explicitly make the request. There are special cases in which this might be useful. For example, if additional nodes are added to a network from a second network discovery procedure and if LSPs in the original network have unresolved references to these nodes, it is useful to resolve the LSPs again.

# 20 Optimize RSVP-TE Routing

**Topics:**


- RSVP-TE optimization
- RSVP-TE optimization input properties
- Post optimization options
- Run RSVP-TE optimization
- Optimization output


Guides the optimization of RSVP-TE routing by outlining input property configuration, optimization parameters, disjoint groups, and avoidance constraints.

## RSVP-TE optimization

Explains how RSVP-TE optimization centralizes path selection to minimize distance, reduce congestion, and improve traffic engineering in networks.

**RSVP-TE optimization** is a Cisco Crosswork Planning that

- creates explicitly routed LSP paths to optimize route properties such as distance, TE metrics, utilization, and disjointness
- centralizes routing decisions within itself instead of distributing across head-end routers in the network, and
- delivers faster and more predictable performance than the explicit LSP optimization tool.

The main optimization goal is to minimize the path distance, which can be based on TE metrics, delay, or a user-defined value, while not exceeding a *bandwidth bound* defined for each interface so as to avoid congestion. Unlike RSVP-TE, LSPs are routed even if they cannot avoid exceeding bandwidth bounds for certain interfaces, though these bound violations are minimized. Additionally, you can use this tool for reoptimizing and reconfiguring RSVP-TE LSPs after events, such as network failures and subsequent route reconvergence.

You can specify different treatment for different LSPs: whether the LSPs should be rerouted to shorten the path distance, only as necessary to reduce congestion, or not optimized at all. These three LSP groups are called Opt, Fit, and Fix. Using these input parameters allows for a variety of use cases, such as:

- Global optimization: Minimize the path length of all LSPs subject to specified interface bandwidth constraints.
- New LSP optimization: Minimize the path length of newly created LSPs without exceeding bandwidth bounds specified for the interfaces. Existing LSPs are kept on their current paths, being moved only if necessary to prevent new LSPs from violating bandwidth bound constraints.
- Tactical congestion mitigation: Change the routing of as few LSPs as possible to bring bandwidth on all interfaces below the specified bandwidth bound.

The tool is beneficial for offline planning and configuration of an explicitly routed RSVP-TE network. It is also useful for providing routing decisions for a centralized, reactive network control application, such as the Cisco Crosswork Planning platform.

> **Note**
>
> RSVP-TE optimization tool supports both Inter-Area and Inter-AS functionalities.

## RSVP-TE optimization input properties

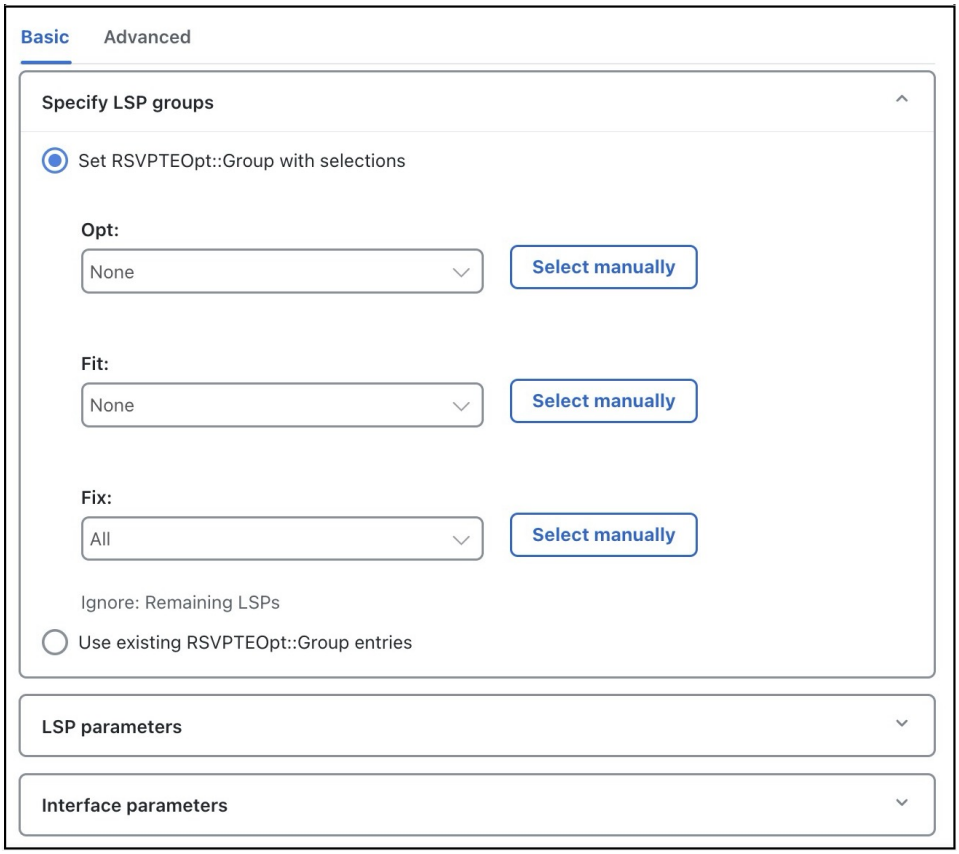
Lists and describes the LSP and interface input properties required by the RSVP-TE optimization tool.

The RSVP-TE optimization tool (found under **Actions** > **Tools** > **RSVP LSP optimization** > **RSVP-TE optimization**) uses the LSP and interface input properties listed in Table 31: RSVPTEOpt input parameters on page 320 to execute its optimization calculations.

These properties, which are specified in user columns in the LSPs and Interfaces tables, can be created in one of the two ways:

- Automatic creation: Use the tool to automatically create the properties if they do not exist. The selections made in the **Specify LSP groups**, **LSP parameters**, and **Interface parameters** areas populate the input properties.

  **Figure 98: RSVP-TE optimization options**




Example 1: If you choose LSPs to optimize, to fit, and to fix (in the **Specify LSP groups** section), the tool creates the RSVPTEOpt::Group column in the LSPs table and sets this property to the appropriate values for each LSP (Opt, Fit, Fix). Existing properties are updated if values change.

Example 2: If you select to set the interface bandwidth bound based on the Resv BW property (in the **Interface parameters** section), the tool creates the RSVPTEOpt::BWBound column in the Interfaces table and for each interface, copies the **Resv BW** value to it. Existing properties are updated if the new value changes.

- Manual creation: Before running the tool, manually create the input properties as user columns in the LSPs and Interfaces tables. When running the tool, you must then select the **Use existing RSVPTEOpt:: property_name** option. This is useful if you want to configure specific values rather than using existing properties and values. For information on creating user-defined columns, refer to Table 31: RSVPTEOpt input parameters on page 320.

> **Note**
>
> As the resulting network model can potentially reset properties that can be used as input to optimizations, it is a best practice to run the optimization only on plan files or LSPs that have not yet been optimized. For example, by default, Cisco Crosswork Planning resets the LSP setup BW property to 0 after optimizing LSPs, which could impact subsequent optimizations.

**RSVPTEOpt input parameters**
**Table 31: RSVPTEOpt input parameters**

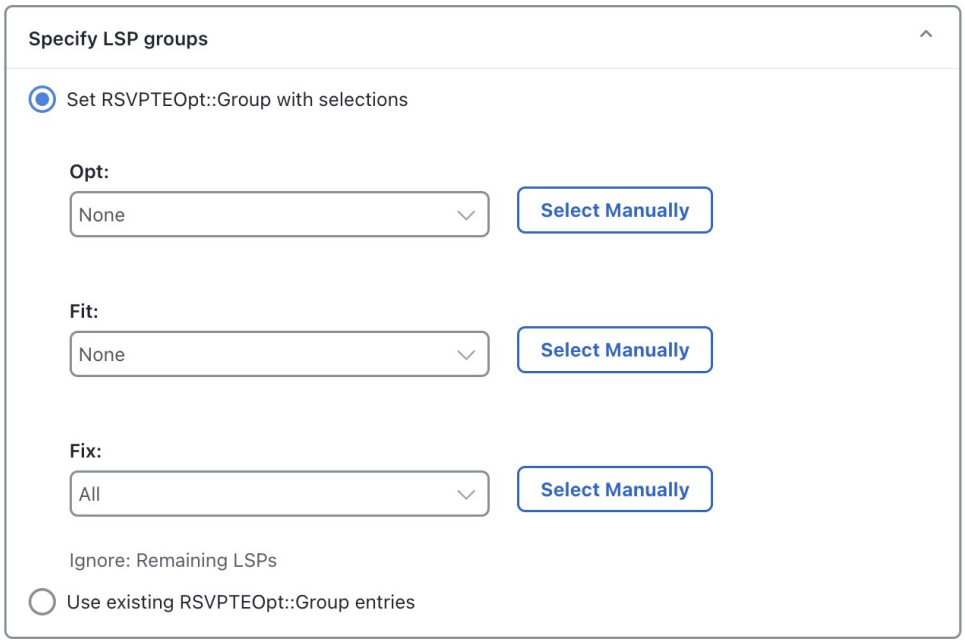
| Table | Input property | Description and values used to create the Input property |
|---|---|---|
| LSPs | RSVPTEOpt::Group | LSP optimization group.<br>The input property is derived based on your entries in the **Specify LSP groups** area unless you select to use the existing RSVPTEOpt::Group property. Values are Opt, Fit, or Fix. If the value is not defined, it is set to None.<br>For more information, refer to LSP group options on page 320. |
| | RSVPTEOpt::BWReq | The bandwidth required by each LSP that is being optimized.<br>The input property is derived from the Setup BW, Traff meas, or Traff sim properties in the **LSP properties** section, unless you specify to use the existing RSVPTEOpt::BWReq value.<br>For more information, refer to Optimization parameters on page 321. |
| Interfaces | RSVPTEOpt::BWBound | Sum total of LSP bandwidth routed over the interface should not exceed this value.<br>The input property is derived from the Resv BW, Capacity, or Capacity sim properties in the **Interface properties** section, unless you specify to use the existing RSVPTEOpt::BWBound value.<br>For more information, refer to Optimization parameters on page 321. |
| | RSVPTEOpt::Metric | Metric value used in the shortest path calculations.<br>The input property is derived from either the TE metric or Delay property in the **Interface properties** section, unless you specify to use the existing RSVPTEOpt::Metric value.<br>For more information, refer to Optimization parameters on page 321. |

## LSP group options

Provides details about LSP group selection options and how these options determine optimization treatment.

The LSP group options determine the optimization treatment of LSPs using the RSVP-TE optimization tool. To determine how and whether LSPs are optimized, select them from the **Specify LSP groups** section. If an LSP belongs to multiple groups, its optimization is determined by the first group in which it appears.

**Figure 99: LSP groups**




Select the **Set RSVPTEOpt::Group with selections** radio button and set these properties:

- **Opt**: Explicitly routes or reroutes the LSPs to minimize their path length, while respecting the bandwidth bounds specified for the interfaces. The RSVP-TE optimization tool changes the LSP path as much as needed to find the optimal route. First, it tries to set the LSP to its shortest path, which is defined by the interface Metric parameter. The LSP is moved away from the shortest path only if needed due to congestion.

  Setting all LSPs to this option is useful when trying to achieve global optimization on all LSPs.

- **Fit**: Reroutes the LSPs only as necessary to respect specified bandwidth bounds on the interfaces, while accommodating the Opt group. If there is currently no explicit route or if the route is incomplete, treats like Opt LSPs.

  This option is useful when attempting tactical congestion mitigation, where it is desirable to move as few LSPs as possible to remove existing congestion.

  A combination of Opt and Fit, or a combination of Opt and Fix is useful for targeting the optimization of new LSPs.

- **Fix**: Does not reroute the LSPs, but includes them in the optimization calculations.

#### Selection behavior

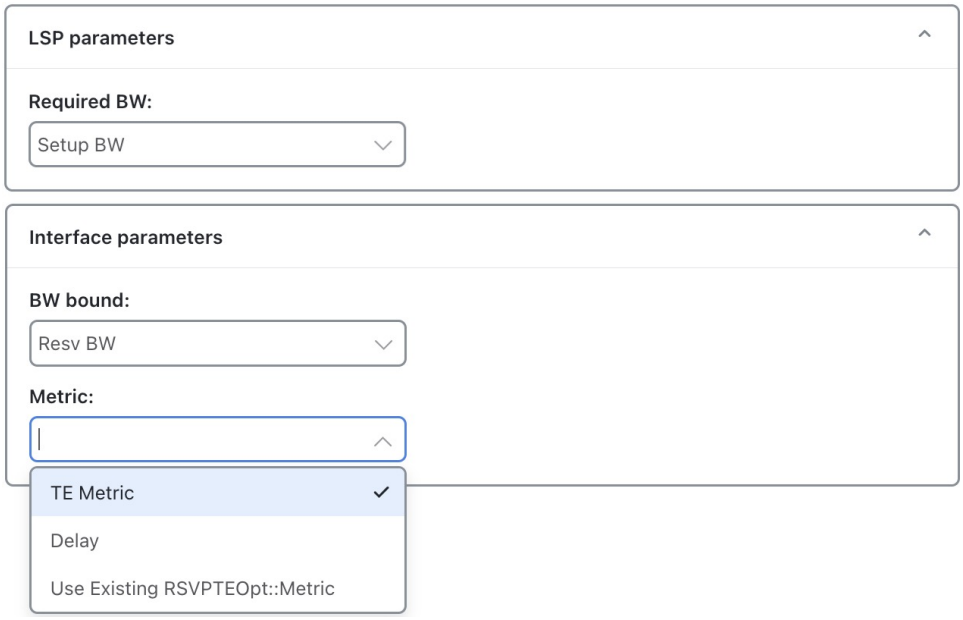
LSPs that are not selected are ignored. They maintain their original configuration and are not considered in optimization calculations. This is useful for network regions that do not interfere with optimizations. Note that if measured interface traffic is used in optimizations, traffic contributions from the ignored LSPs to this measured traffic may affect the results.

You can set each of the **Set RSVPTEOpt::Group with selections** to None, All, or manually selected. This configuration allows you to assign specific LSPs to different optimization treatments.

## Optimization parameters

Lists the parameters for LSP group selection and describes their effects on bandwidth optimization and path calculation.

Optimization calculations are based on two sets of parameters that define how much bandwidth the LSPs require, how much bandwidth the interfaces have available for these LSPs, and which metric to use for shortest path.




The **LSP parameters** and **Interface parameters** sections offer these options:

**Table 32: Optimization input parameters**

| Parameter | Description | Options |
|---|---|---|
| Required BW | This LSP parameter specifies how to determine the bandwidth required for primary LSP paths. | • Setup BW<br>• Traff meas<br>• Traff sim |
| BW bound | This interface parameter indicates the bandwidth available on interfaces to carry LSPs. The system tries to route LSPs without surpassing this bound, but exceeds it if necessary. | • Resv BW<br>• Capacity<br>• Capacity sim |
| Metric | This interface parameter defines the metric used for shortest path calculations. | • TE metric<br>• Delay |

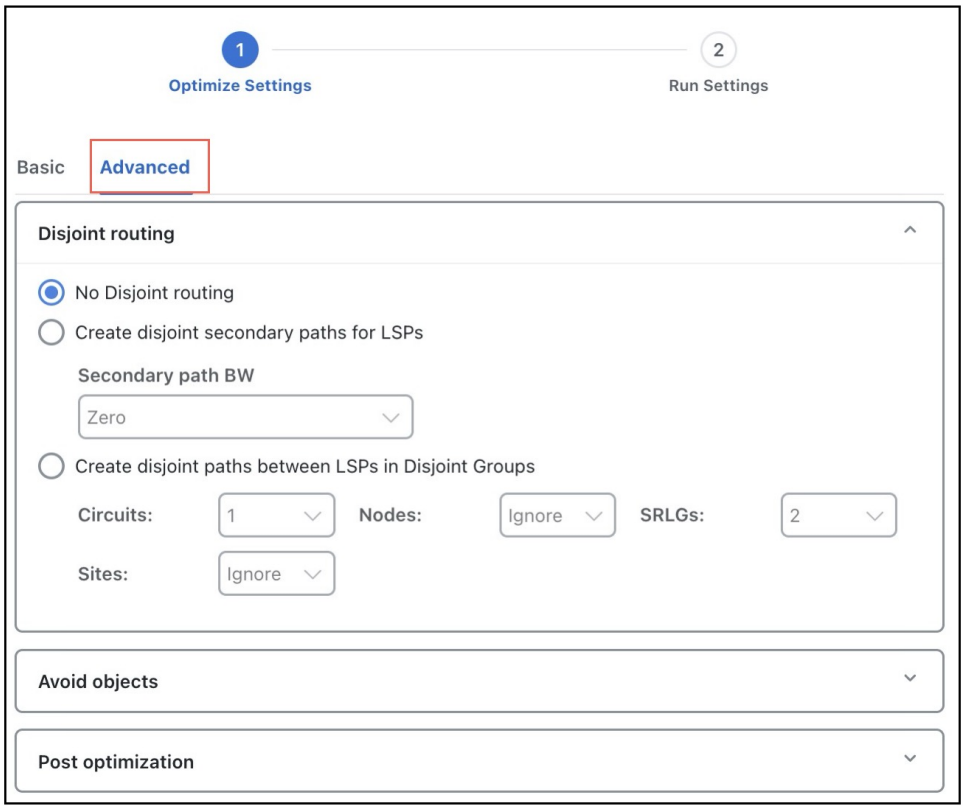
## Disjoint groups

Summarizes the configuration options and attributes of disjoint groups for LSP path setup.

Two LSP paths are *disjoint* if they do not route over common objects. These objects are configurable and can be circuits, nodes, sites, or SRLGs.

In the **Advanced** tab, you can create disjoint paths or use existing disjoint paths.

**Figure 100: RSVP-TE LSP optimization advanced options**




These are the available disjoint routing options:

- **No disjoint routing**: Does not create disjoint secondary paths and does not enforce any disjointness on primary paths.
- **Create disjoint secondary paths for LSPs**: Creates secondary LSP paths that are disjoint from the primary LSP paths.

  Additionally, specifies whether to include secondary paths in the optimization. If set to Zero, their bandwidths are set to 0 and are not included. Otherwise, their required bandwidths are set to be the same as the primary paths, and they are included in the optimization.
- **Create disjoint paths between LSPs in Disjoint Groups**: Creates disjoint paths between LSPs that are in disjoint groups.

## Avoidance constraints

Lists the types of avoidance constraints available for LSP optimization and describes their impact on path selection.

Avoidance constraints specify network elements that should be excluded from the optimization of LSPs.

The **Avoid objects** section in the **Advanced** tab (Figure 100: RSVP-TE LSP optimization advanced options on page 322) lets you select Nodes, Interfaces, and SRLGs to be avoided when optimizing LSPs.

The available object types are:

- Nodes: LSPs are created to route away from the Nodes selected. The default value is None.
- Interfaces: LSPs are created to route away from the Interfaces selected. The default value is None.
- SRLGs: LSPs are created to route away from the SRLGs selected. The default value is None.

## Post optimization options

Lists the available options for configuring Setup BW reset for modified LSPs in the **Advanced** tab after optimization.

The **Post optimization** section in the **Advanced** tab (refer to Figure 100: RSVP-TE LSP optimization advanced options on page 322) lets you identify how to reset the Setup BW for modified LSPs.

The available options are:

- Reset Setup BW to 0.
- Reset Setup BW to be the same as the resulting RSVPTEOpt::BWReq value.
- Keep the Setup BW the same.

## Run RSVP-TE optimization

Describes how to run the RSVP-TE optimization tool to improve LSP routing and bandwidth allocation in your network.

Use this task optimize RSVP-TE LSPs and interfaces to enhance routing and ensure bandwidth requirements are met.

### Before you begin

- If this is the first time using the tool and if you want to use your own preconfigured RSVPTEOpt properties, you must manually edit the LSPs and Interfaces tables to create them:
  - LSP RSVPTEOpt::Group: Place LSPs into the Opt, Fit, or Fix groups. Any remaining LSPs are placed in the Ignore group.
  - LSP RSVPTEOpt::RequiredBW: Define an exact amount for the sum total of LSP bandwidth that is using the interface.
  - Interface RSVPTEOpt::BWBound: Define an exact amount of bandwidth that the interface can hold.
  - Interface RSVPTEOpt::Metric: Define an exact amount to use in shortest path calculations.
- If you are creating disjoint paths between LSPs in the same disjoint group, you must first add the LSPs to the disjoint groups in the Edit LSPs page. Here, you can also assign priorities to LSPs within these groups. Higher priority LSPs are assigned shorter routes based on the metric being used for shortest path calculations. The higher the number, the lower the priority.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Tools** > **RSVP LSP optimization** > **RSVP-TE optimization**.

   Alternatively, choose **Preset workflows** > **Perform optimization**, select **RSVP LSP Optimization** as the optimization type, select **RSVP-TE optimization** from the drop-down list, and click **Launch**.
3. In the **Specify LSP groups** section, make selections for optimizing LSPs based on Opt, Fit, or Fix LSP groups, or based on those defined in the RSVPTEOpt::Group property. For more information, refer to LSP group options on page 320.
4. Expand the **LSP parameters** panel. In the **Required BW** field, set bandwidth requirements for the primary LSP paths. For more information, refer to Optimization parameters on page 321.
5. Expand the **Interface parameters** panel. In the **BW bound** field, set how much bandwidth the interfaces can carry. In the **Metric** field, specify on which property to base the shortest path. For more information, refer to Optimization parameters on page 321.

6. Click **Next**.
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

   The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

   > **Note**
   >
   > Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:
   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
   - If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

9. Submit your changes.

---

**What to do next**

Refer to Reports.

## Optimization output

Lists and describes the generated optimization output by the optimization tool.

The RSVP-TE optimization tool produces several key outputs, including specific properties, detailed reports, and updated LSP configurations designed to maximize network efficiency.

### Properties created

In addition to creating the input parameters as described in Table 31: RSVPTEOpt input parameters on page 320, the tool generates the properties listed in Table 33: RSVPTEOpt parameters created on page 326. These properties provide insights into the optimizations.

**Table 33: RSVPTEOpt parameters created**

| Table | Property | Description |
|---|---|---|
| LSPs | RSVPTEOpt::Action | Identifies whether the LSP's path was updated (optimized). |
| | RSVPTEOpt::PathMetric | Sum of interface metrics on the LSP path. |
| | RSVPTEOpt::ShortestPath | Identifies whether the LSP is now taking the shortest path. |
| | RSVPTEOpt::ShortestPathMetric | Sum of interface metrics on the shortest LSP path. |
| Interfaces | RSVPTEOpt::BWTotal | Sum of the bandwidth required for all LSPs routed on the interface. |
| | RSVPTEOpt::BWBoundExceeded | Identifies whether the bandwidth bound of the interface was exceeded. |

### Reports

Each time the optimizer is run, a report is automatically generated. You can access this information by choosing **Actions** > **Reports** > **Generated reports**. Note that new reports replace any previously generated reports.

The report summarizes the optimization results, including:

- The total number of LSPs and how many were optimized
- How the optimizations were performed
- Handling of disjoint priorities
- The number of paths that exceeded the interface's bandwidth limit

### Reconfigured optimized LSPs

The optimized LSPs are reconfigured with these parameters:

- Metric type: Autoroute
- Metric and Hop limit: NA
- Setup BW: 0, the same value as the resulting RSVPTEOpt::BWReq, or unchanged, depending on how you configured the Post Optimization option in the Advanced tab.
- Primary paths Standby: T (true)
- Status: Active
- FRR enabled: T (true)
- LSP entry in the **Actual paths** table is removed.

# 21 Perform Explicit and Tactical Explicit RSVP LSP Optimization

**Topics:**


- Explicit RSVP LSP optimization tool
- Run Explicit RSVP LSP optimization
- Tactical explicit RSVP-TE LSP optimization


Outlines the principles, usage, and configuration of explicit and tactical RSVP-TE LSP optimization, covering execution methods and available options for efficient traffic engineering.

## Explicit RSVP LSP optimization tool

Explains how tools optimize explicit RSVP LSPs to minimize network congestion and improve path reliability.

The **Explicit RSVP LSP optimization** tool is a Cisco Crosswork Planning tool that

- minimizes congestion by optimizing the placement of primary and secondary paths for selected RSVP LSPs
- reduces the utilization across primary paths under normal operation, and
- creates disjoint secondary paths so that a single failure cannot disrupt both paths simultaneously.

By default, the tool optimizes utilization on all interfaces using all RSVP LSPs on those interfaces. It removes CSPF constraints, such as affinities and hop limits. For example, Cisco Crosswork Planning sets the setup bandwidth to 0, thus providing the greatest flexibility when setting these explicit paths.

Upon completion, Cisco Crosswork Planning writes a report containing the results of the optimization. To access the report, choose **Actions > Reports > Generated reports**.

## Run Explicit RSVP LSP optimization

Outlines the steps to configure objectives and advanced settings for RSVP LSP optimization.

Use this task to run the Explicit RSVP LSP optimization tool.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.

   From the toolbar, choose **Actions** > **Tools** > **RSVP LSP optimization** > **Explicit optimization**.

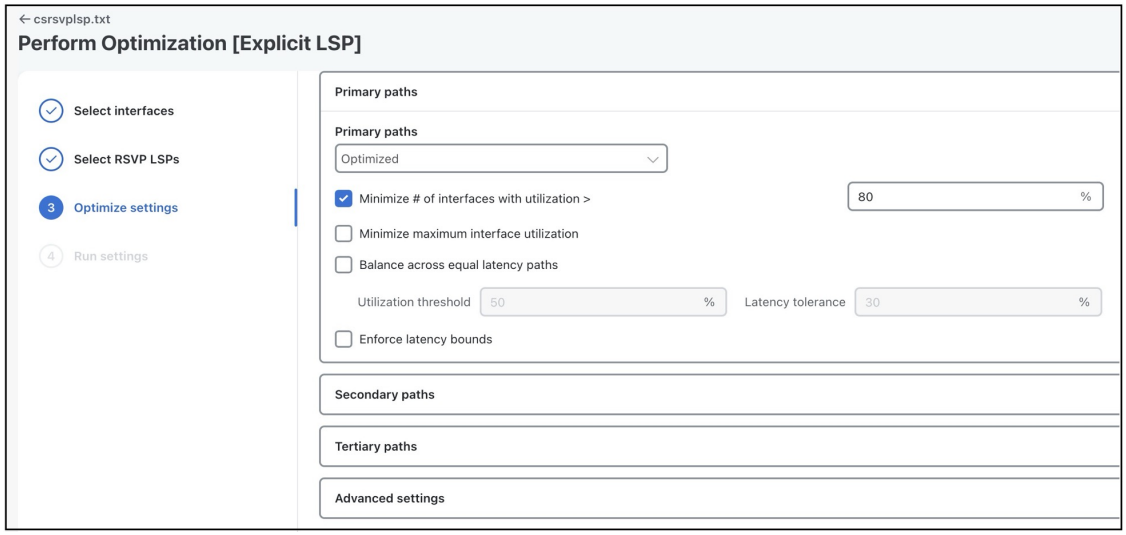
   Alternatively, choose **Preset workflows** > **Perform optimization**, select **RSVP LSP Optimization** as the optimization type, select **Explicit optimization** from the drop-down list, and click **Launch**.

2. Select the interfaces you want to optimize and click **Next**.

   By default, none of the interfaces are selected. You can, however, optimize a limited number of interfaces or specify a set of RSVP LSPs by preselecting them.

3. Select the RSVP LSPs you want to optimize and click **Next**.
4. Specify objectives for primary, secondary, and tertiary paths. For field descriptions, refer to Table 1.

   **Figure 101: Explicit RSVP-TE LSP optimization options**




5. (Optional) Specify the required advanced settings. For field descriptions, refer to Table 1.
6. Click **Next**.
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

   The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

   **Note**

   Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:
   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
   - If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

9. Click **Submit**.

## Explicit RSVP LSP optimization options

Lists the available options for optimizing RSVP-TE LSPs, including detailed settings for primary, secondary, and tertiary paths, interface controls, and advanced configuration.

This table describes the options that are available while running the Explicit and Tactical explicit RSVP LSP optimization tools.

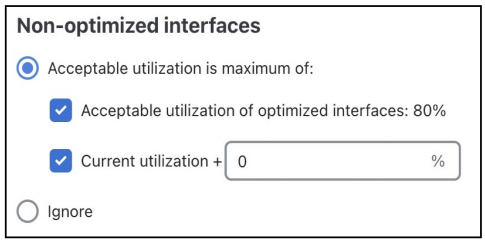
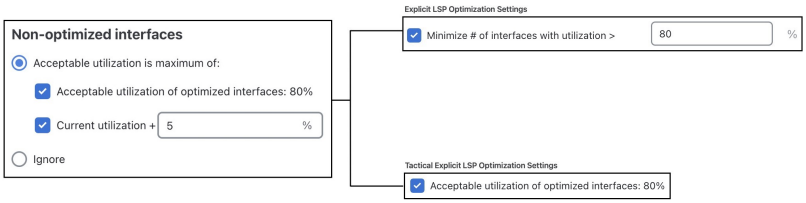
**Table 34: Explicit RSVP LSP optimization options**

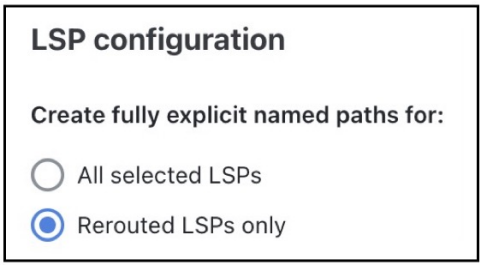
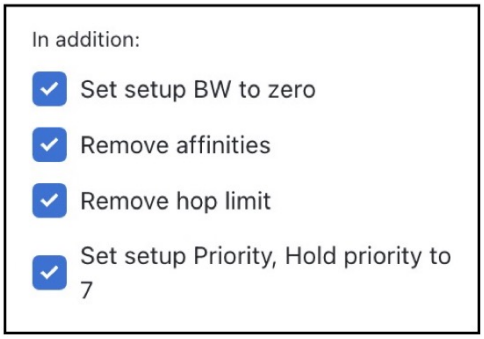
| Field | Description |
|---|---|
| **Primary path options** | |
| Primary paths | Define whether to reroute primary RSVP LSP paths:<br>• **Optimized**: Create optimized explicit primary paths based on the objectives selected.<br>• **Keep**: Route RSVP LSPs along the existing primary paths.<br>Primary paths are optimized using three objectives in order of priority. The first two objectives move RSVP LSPs away from their shortest latency path in an attempt to reduce the utilizations of the most highly utilized interfaces in the network. |
| Minimize # of interfaces with utilization > ___ % | Specify a percentage and minimize the number of selected interfaces with utilizations over that percentage under normal operation (default). |
| Minimize maximum interface utilization | Route primary paths so that the maximum interface utilization over all selected interfaces is minimized under normal operation. |
| Balance across equal latency paths | Balance utilizations over lower utilized interfaces. For example, use this option to balance utilizations on parallel interfaces between two nodes if the two interfaces have the same latency.<br>• **Utilization threshold**: Keep the number of interfaces with utilization greater than this value as low as possible without significantly increasing the latencies of the primary paths.<br>• **Latency tolerance**: Permit this percentage of additional latency. |
| Enforce latency bounds | Enforce latency bounds that can be specified for some or all demands in the plan file. If checked, this option takes precedence over all of the preceding objectives. |
| **Secondary path options** | |

| Field | Description |
|---|---|
| Secondary paths | Define whether and how to route secondary paths. These are the options:<br>• **Optimized**: Create optimized explicit secondary paths based on the objectives selected.<br>• **Dynamic**: Route secondary paths dynamically. No explicit hops will be created for the path.<br>• **None**: No secondary paths are created; existing paths are removed.<br>For optimized secondary paths, the objectives are used in order of priority listed. |
| Hot standby | Set the secondary path to be a *hot* standby, which means it is established at the same time as the primary path, rather than after the primary path fails. |
| 1. Maximize primary/secondary path disjointness with respect to: | Define primary and secondary paths for each RSVP LSP that are disjoint with respect to circuit, SRLG, and nodes, depending on what is selected.<br>• **Circuits**: No circuit is used by both the primary and secondary paths (default = 1).<br>• **SRLGs**: No SRLG is used by both the primary and secondary paths (default = 2).<br>• **Nodes**: No node is used by both the primary and secondary paths (default = 3).<br>• **Traffic disjointness only**: A path is acceptable even if it uses similar circuits, SRLGs, or nodes as other paths provided there is no traffic routed over the RSVP LSP when failures occur.<br>You can specify the degree of disjointness of primary and secondary paths. The lower the number, the higher the disjointness priority. For example, if it is important that the paths are node disjoint and the SRLG disjointness are less important, you can change the setting to circuits 1, nodes 2, and SRLGs 3.<br>Note that the network topology sometimes makes it impossible to fulfill all the selected disjointness requirements. In this case, paths are selected that are maximally disjoint. That is, they are disjoint for as many circuits, SRLGs, and nodes as possible. |
| 2. Minimize # of interfaces with utilization > ___ % | Options 2 and 3 both operate on the selected failure scenarios (circuits, SRLGs, and nodes) listed in the “Failures to consider” options listed under 3. These choices are the failure scenarios over which the simulation is performed at the end of the optimization. Note that this selection of failure scenarios is distinct from the failure scenarios selected for the disjointness objective 1.<br>Use option 2 to minimize the number of interfaces with utilizations over the specified percentage across all selected failure scenarios. |

| Field | Description |
|---|---|
| 3. Minimize maximum interface utilization | Use option 3 to minimize the maximum interface utilization over all interfaces and over all selected failure scenarios. |
| Failures to consider | Choose the objects to consider. |
| **Tertiary path options** | |
| Tertiary paths | Define whether to create tertiary paths:<br>• **Dynamic**: Create dynamic tertiary paths.<br>• **None**: No tertiary paths are created. |
| Hot standby | Set the tertiary path to be a hot standby path, which means it is brought up with the primary path, rather than after the primary path fails. |

**Table 35: Advanced explicit RSVP LSP optimization options**

| Field | Description |
|---|---|
| Non-optimized interfaces | You can specify whether to ignore non-optimized interfaces or set the acceptable level of utilization for them.<br><br>If setting an acceptable utilization level and if both options are selected, Cisco Crosswork Planning uses the higher of the two. These settings are calculated on a per-interface basis.<br><br>• **Acceptable utilization of optimized interfaces: ___%**: This value is the same as the utilization threshold set in the Primary path area (**Minimize # of interfaces with utilization > ___%**, where the default is 80). To change this value, you must change it in the Primary path area.<br><br>If using the Tactical explicit RSVP-TE LSP optimization tool, this field is equivalent to, and thus only changeable in, the Acceptable utilization ___% field.<br><br>• **Current utilization + ___%**: Current utilization of non-optimized interfaces plus the added percentage.<br><br><br><br><br>Example: There are two non-optimized interfaces: cr1.chi_cr1.mia has a utilization of 60% and cr2.sjc_cr2.okc has a utilization of 78%.<br><br>The acceptable utilization settings for non-optimized interfaces are:<br><br>• The primary path utilization threshold is 80%.<br><br>• The current utilization has 5% added to it.<br><br><br><br><br>Result: The maximum utilization for each interface is individually calculated. The acceptable level of utilization for cr1.chi_cr1.mia is 80%, and the acceptable level of utilization for cr2.sjc_cr2.okc is 83% (78 + 5). |

| Field | Description |
|---|---|
| LSP configuration | By default, Cisco Crosswork Planning creates named paths across the RSVP LSPs that were selected if a reroute is required to achieve the optimization objectives. For example, if a selected dynamically routed RSVP LSP has an acceptable route, Cisco Crosswork Planning does not create a named path for it.<br><br>To change this default, click **All selected LSPs**. Cisco Crosswork Planning then creates and routes fully explicit named paths for all selected RSVP LSPs.<br><br><br><br><br>By default, Cisco Crosswork Planning sets the Setup BW to zero, providing the greatest flexibility when creating explicit routes. As well, all affinities and hop limits are removed, and the setup and hold priorities are set to 7. These changes apply only to the RSVP LSPs with newly created or changed explicit named paths.<br><br>You can turn off these defaults, and you can also set them individually after the optimization is performed using the Edit LSP Path window. If these defaults are turned off, the original parameters are preserved.<br><br><br> |
| Traffic level | The traffic level used in the utilization calculations and optimizations. For information on traffic levels, refer to Simulate Traffic Flow from Source to Destination Using Demands on page 95. |
| Rerouting preference | By default, preferences on which RSVP LSPs to reroute are not based on traffic volume. You can select high-traffic or low-traffic options to sequentially give priority to RSVP LSPs with higher and lower traffic. |

## Tactical explicit RSVP-TE LSP optimization

Describes how the Tactical explicit optimization tool reduces congestion in targeted network areas by minimizing RSVP LSP reconfigurations. Provides context for users seeking efficient path optimization with limited changes.

The **Tactical explicit optimization** tool is a reduced version of the Explicit RSVP LSP optimization tool that optimizes only primary paths using the minimum number of path changes required to bring utilization below an acceptable level. It

is useful when you need to reduce congestion in a specific area of the network with a limited number of RSVP LSP reconfigurations.

If you need to target specific interfaces or RSVP LSPs to confine the optimization to problem areas, select those interfaces or RSVP LSPs first. Then, from the toolbar, choose any of these options:

- **Actions** > **Tools** > **RSVP LSP optimization** > **Tactical explicit optimization**
- **Preset workflows** > **Perform optimization**, select **RSVP LSP Optimization** as the optimization type, choose **Tactical explicit optimization** from the drop-down list, and click **Launch**

# 22 Configure Segment Routing

**Topics:**


- Segment Routing LSPs
- SR LSP segment types
- SR LSP paths
- SR LSP routing rules
- SR LSPs and their LSP paths
- Flexible Algorithms
- Create segment lists
- Create SIDs
- SR-TE protection


Explains Segment Routing configuration, including segment list concepts, SR LSP paths, Flex Algorithm creation, and SR-TE protection features to support scalable, flexible, and resilient SR deployments.

## Segment Routing LSPs

Explains how segment routing LSPs operate, how they differ from RSVP LSPs, and how segment lists guide traffic through networks.

By default, Cisco Crosswork Planning creates and routes RSVP LSPs. However, you can create Segment Routing (SR) LSPs by changing the LSP **Type** property to **SR**.

SR LSPs use autoroute, forwarding adjacencies, and class-based forwarding, just like RSVP LSPs. Once the packet enters the tunnel, the different routing mechanisms determine the path. An RSVP LSP is first established along a path, with each node on the path accepting the reservation request and maintaining the reservation state. In contrast, SR LSPs rely on a segment list that is created at the LSP's source node (head-end). This segment list, which is stored in the packet (demand), directs the traffic over a series of configured node segments and interface segments.

> **Note**
>
> Cisco Crosswork Planning uses the terms *SR LSPs* that use *segment lists*, which contain *segment list hops*. In Cisco routing terminology, these terms are *SR tunnels* that use *segment ID (SID) lists*, which contain *SIDs*.

## SR LSP segment types

Describes the different types of segments found in an SR LSP and how they are used in constructing segment lists.

The segment list of an SR LSP contains one or more segment list hops, which can be nodes, interfaces, SR LSPs, or anycast groups.

To view a segment list, select the SR LSP path from the **LSP paths** table and choose > **Filter to segment lists**. You can then select the segment list and choose > **Filter to segment list hops** to view the corresponding segment list hops.

For more information on creating segment lists and their hops, refer to Create segment lists on page 345.

> **Note**
>
> The process of creating and optimizing segment lists and their segment list hops can be automated using the SR-TE optimization tool. For more information, refer to SR-TE optimization on page 354. You can also create SR LSPs and their hops by optimizing bandwidth. For more information, refer to SR-TE bandwidth optimizations on page 358.

### Node segment list hops

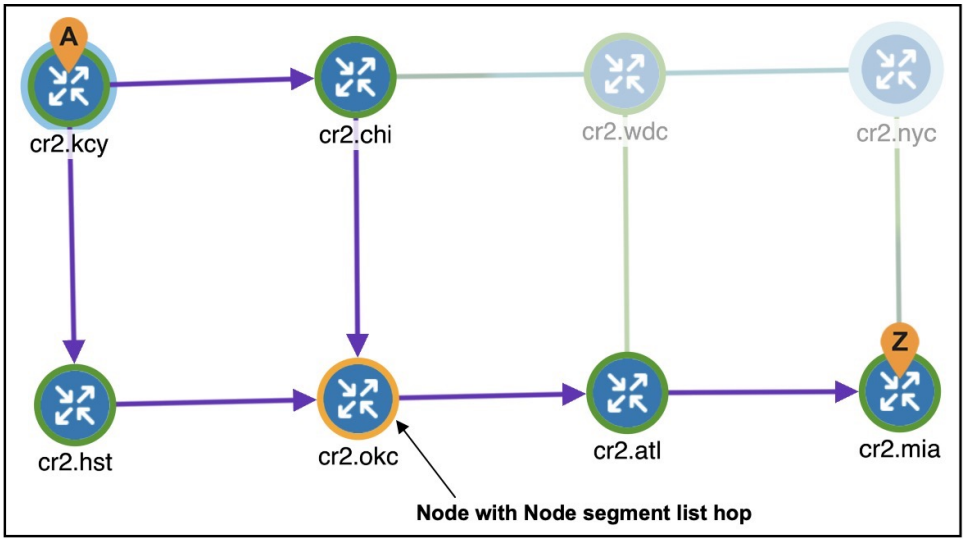
Describes how node segment list hops are used to control the path of an SR LSP, ensuring traffic follows the shortest path through specified nodes and enabling equal-cost multipath routing when metrics are equal.

When using node segment list hops, the SR LSP takes the shortest path to the specified node.

### Example

Figure 102: Example node segment list hop on page 339 shows an example SR LSP from the cr2.kcy node to the cr2.mia node using one node segment list hop cr2.okc (as indicated by orange circle around the node). Since the IGP metrics are equal, the traffic is routed using ECMP to reach cr2.okc.

**Figure 102: Example node segment list hop**




## Interface segment list hops

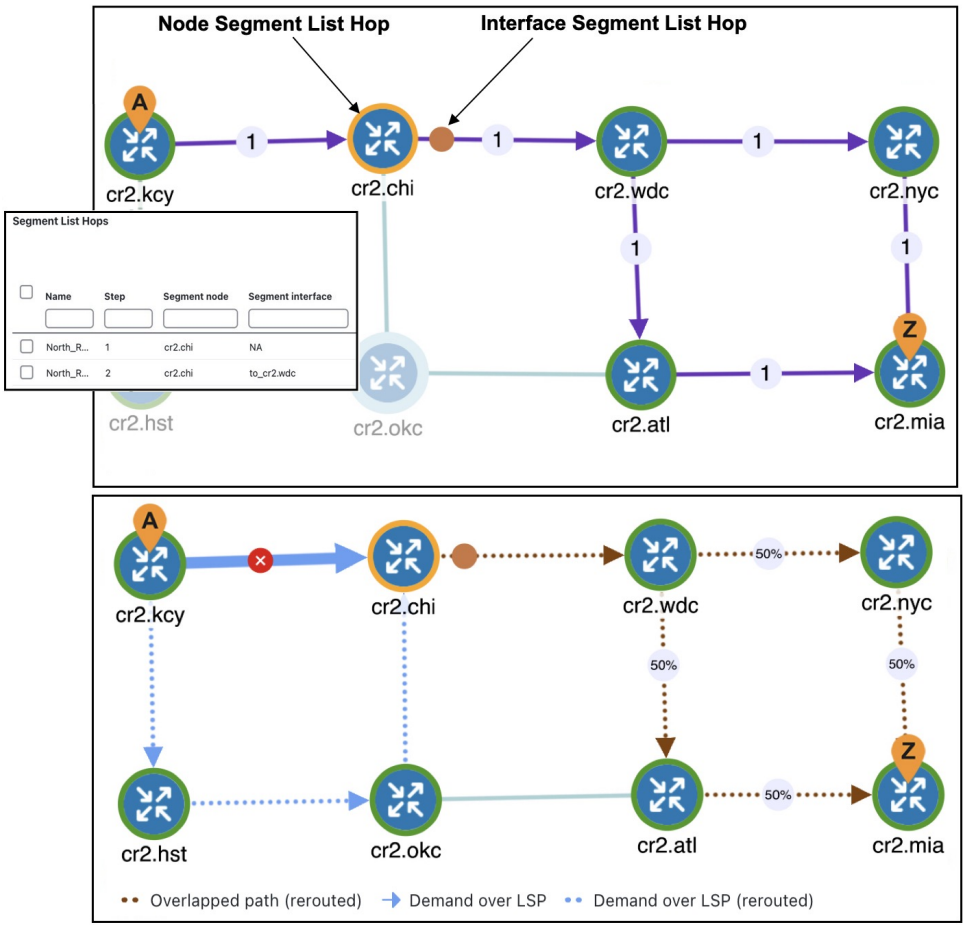
Explains how interface segment list hops are used in segment lists and the differences between standard routers and advanced platforms.

On routers, interfaces can only be used in segment lists if the previous segment list hop is the node containing the interface, or if the interface is the egress of the source node. There is no restriction that an interface segment list hop be local. There is no requirement that the segment list contain a preceding segment to the node containing the interface.

### Example

Figure 103: Example interface segment list hop on page 339 shows an example of an interface segment list hop from cr2.chi to cr2.wdc. To ensure the path reaches cr2.chi (the node the packet must reach to use this interface segment list hop), a node segment list hop is configured first to reach chi. This figure also shows how a demand through an SR LSP reroutes when a node fails. The demand uses IGP to route around the failure and back to the segment, if possible.

**Figure 103: Example interface segment list hop**




## LSP segment list hops

Explains the purpose and restrictions of LSP segment list hops and outlines how they appear for analysis in the system interface.

SR LSPs can use other SR LSPs as segment list hops, and in turn, these SR LSPs segment list hops can contain other SR LSPs as segment list hops.

### Governing rules

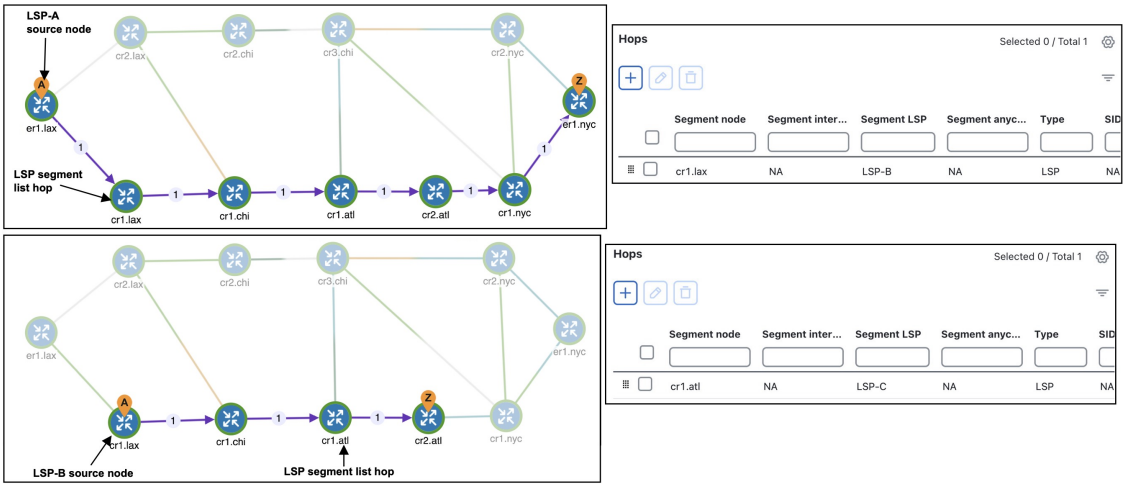
Two rules govern LSP segment list hops:

- The segment list hop must be an SR LSP, not an RSVP LSP.
- An SR LSP referenced as a segment list hop cannot reference (directly or indirectly) the SR LSP that contains it. For example, if the segment for SR LSP A contains SR LSP B as a segment list hop, SR LSP B cannot contain SR LSP A as a segment list hop.

### Example

Figure 104: Example LSP segment list hop on page 340 shows that LSP A (from er1.lax to er1.nyc) contains an LSP B segment list hop, whose source node is cr1.lax. By deselecting LSP A and then selecting LSP B, you see that LSP B also contains an LSP segment list hop, which is LSP C whose source node is cr1.atl.

**Figure 104: Example LSP segment list hop**




### How to view LSP segment list hops

To view the LSP segment list hops:

1. Select the LSP and choose > **Filter to LSP paths**. The LSP Paths page appears.
2. Select the LSP paths and choose > **Filter to segment lists**. The Segment Lists page appears.
3. Select the segment list and click > **Filter to segment list hops** to see the relevant segment list hops.

## Anycast group segment list hops

Explains how Anycast group segment list hops enable Segment Routing LSPs to flexibly select the shortest path among multiple group members.

An SR LSP routes through the node in an anycast group segment list hop that has the shortest path to it. In case of a tie between multiple nodes in an anycast group, ECMP is applied. This mechanism lets you impose routing restrictions in terms of potential intermediate segment destinations. It also lets the SR LSP choose among the possible next hops (next segment list hops), potentially reducing latency and improving load balancing.

### Example

Figure 105: Example anycast group segment list hop on page 341 shows an example of anycast group nodes cr2.chi and cr1.chi. In one instance, the IGP metrics to those nodes have equal IGP metrics, so the SR LSP uses ECMP to route through both of them. When the IGP metric increases from cr2.lax to cr1.chi, the SR LSP routes only through cr2.chi because it has the shortest path. However, as Figure 106: Example anycast group rerouting around failure on page 342 shows, if a failure prevents the shortest path from being used, the SR LSP can still route through the anycast group using the next-highest cost path.

**Figure 105: Example anycast group segment list hop**

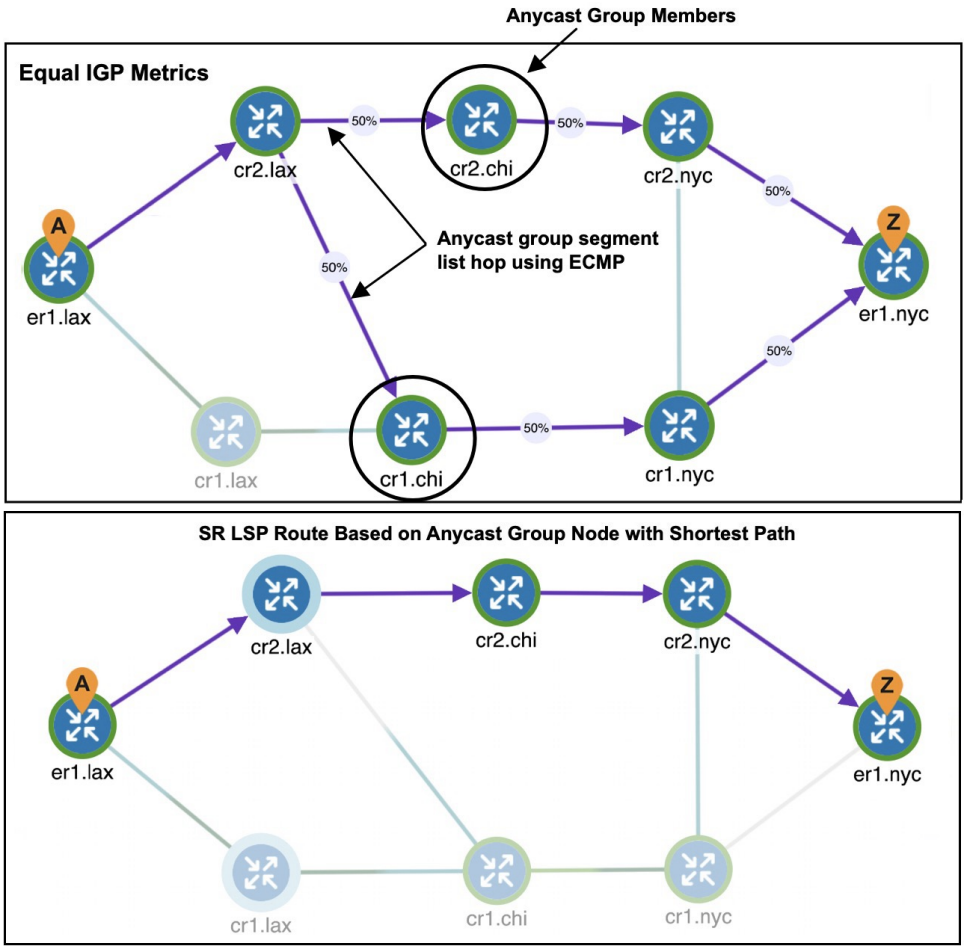
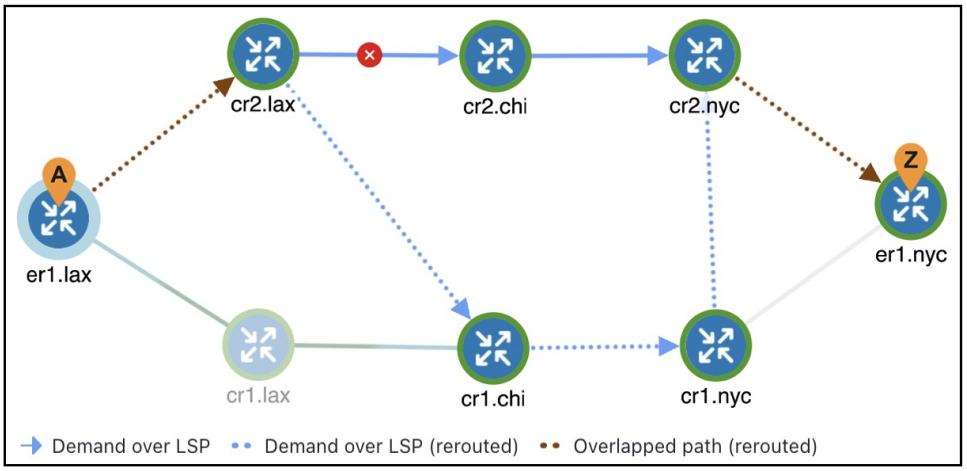



**Figure 106: Example anycast group rerouting around failure**




## SR LSP paths

Details how SR LSP paths function within Segment Routing to provide flexible routing options.

SR LSPs can have multiple LSP paths, where each LSP path has its own segment list. There is no distinction between standby and non-standby LSP paths for SR LSPs.

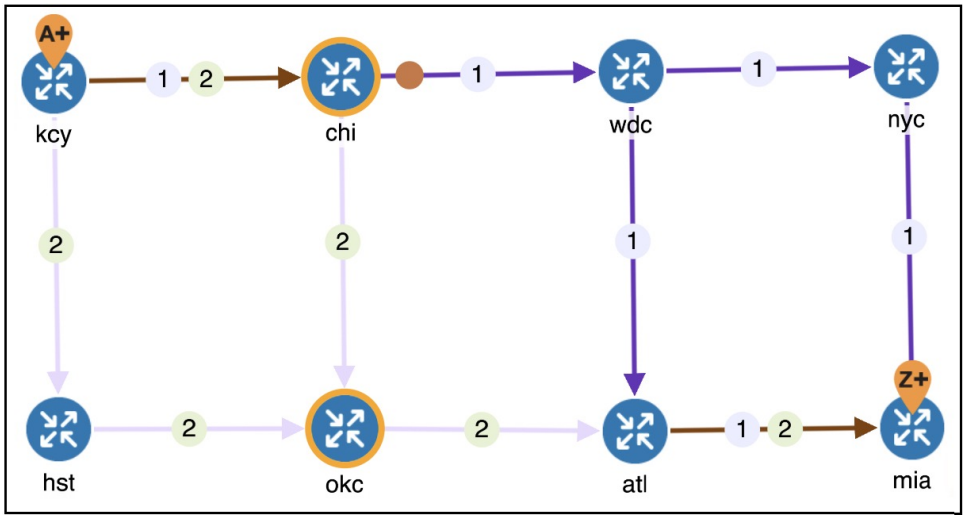
**Note**

If both the SR LSP and its SR LSP paths have segments, the segments on the LSP paths override those on the SR LSP.

**Example**

Figure 107: Example SR LSP on page 343 shows an SR LSP from kcy to mia that has a primary and a secondary path. The primary path is configured with a chi node segment list hop and chi-to-wdc interface segment list hop, and the secondary path is configured with a node segment list hop in okc. This can be seen by selecting the LSP paths individually from the LSP Paths table.

**Figure 107: Example SR LSP**




## SR LSP routing rules

Provides rules and behaviors for how SR LSPs route demands using segment lists, path selection, and node reachability to achieve efficient, flexible traffic routing.

These rules determine how SR LSPs route demands through a network, assuming the common LSP mechanisms such as autoroute and forwarding adjacencies have determined that a demand will enter an SR LSP. If the destination node or segment list hops are not reachable, the SR LSP is not established.

**Note**

The process of creating final segment list hops can be automated. For more information, refer to Simulate Traffic Flow from Source to Destination Using Demands on page 95.

- If the SR LSP does not have a segment list, the demand is routed via IGP to the destination node.

  If the SR LSP has a segment list, the demand is routed using the segment list hops within the segment list in the sequential order in which they appear.

- If it is using a node segment list hop, the demand takes the shortest path to that node.

If it is using an interface segment list hop, the demand takes the shortest path of the source node of that interface, and then routes to that interface's destination node.

If it is using an anycast group segment list hop, the demand routes through the node with the shortest path. In case of a tie between multiple nodes in an anycast group, the SR LSP uses ECMP to route.

- If you select both a node and an anycast group containing that node when creating the segment list hop, the node segment has precedence.
- If the SR LSP contains LSP paths, the demand is routed on the LSP path with the lowest path option for which the demand can be routed from source to destination. Demands for SR LSP paths route in the same manner as described above for SR LSPs.
- If the SR LSP contains multiple SR candidate paths, the demand is routed on the LSP path with the highest preference option for which the demand can be routed from source to destination. Demands for SR LSP paths route in the same manner as described above for SR LSPs.
- If an SR LSP path is configured as 'Dynamic', the LSP path is simulated based on the specified 'Metric type'. This is applicable for IGP, TE, and Delay metric types. As ECMP is not allowed along the path, interfaces with the lowest IP address are preferred. If there is a segment list configured, the LSP path is not dynamically allocated.
- If the demand cannot be routed on the segments defined in the SR LSP or its LSP paths (for example, because of a failure in a node segment list hop), the demand is routed to the destination using the IGP shortest path.
- If a hop (nodes, interfaces, LSPs, or anycast groups) contains a segment ID (SID), the segment list hop inherits that SID.

### Inter-AS LSP routing

Explains how SR LSPs operate across Autonomous System (AS) boundaries, including required hop configurations and demand requirements.

An inter-AS LSP routing is a network routing mechanism that enables SR LSPs to traverse multiple ASes.

To route SR LSPs with source and destination nodes in different ASes, these hop configurations must exist:

- A node segment list hop on the border of the AS that the LSP is exiting. This node is the exit node.
- A node segment list hop on the first node traversed in another AS. This node is the entry node. If there are multiple peering interfaces, their metrics determine which one the LSP uses.
- An interface segment list hop that connects the exit node of one AS to an entry node in another AS.

To use demands with inter-AS SR LSPs, these two conditions must be met:

- The demand must be private to the SR LSP. For information on creating private LSPs, refer to Configure MPLS Routing on page 233.
- The demand must have the same source and destination as the SR LSP.

## SR LSPs and their LSP paths

Describes the attributes and creation process of SR LSPs and their associated LSP paths.

SR LSPs and SR LSP meshes follow the same creation method as RSVP-TE LSPs but require the **Type** property to be set to **SR**. For more information, refer to Create and visualize LSPs.

The SR LSP has a **Color** property which is a 32-bit numerical value. Once created, the next step is to either create the segment list or create LSP paths and their segment lists.

Likewise, SR LSP paths are created in the same way as RSVP-TE LSP paths (refer to Create LSP paths). After creating the SR LSP paths, you can set the **Protocol origin** (Edit LSP Path page) to identify the component or protocol that originates the SR LSP path. Choose from Local, PCE initiated, and BGP initiated.

The **LSP Path Weighted Segment Lists** table provides details of the segment lists associated with the SR LSPs, with the **SLName** column specifying the segment list names.

For more information on LSPs and LSP paths, refer to Configure MPLS Routing on page 233.

# Flexible Algorithms

Describes the role of flexible algorithms in enhancing network routing flexibility with tailored path selection.

Flexible Algorithm (Flex Algo) is a mechanism that

- allows operators to customize and compute the IGP shortest path according to their needs and constraints (specific metrics and link properties), and
- supports user-defined constraints rather than relying on standard algorithm definitions.

For example, Flexible Algorithm can confine the path to a particular plane for networks with multiple logical planes. Since the meaning of the algorithm is not defined by any standard but is defined by the user, it is called a Flexible Algorithm.

## Application-Specific Link Attribute (ASLA)

Introduces ASLA as a link attribute for Flexible Algorithm routing, outlines its role in advertising application-specific metrics.

Application-Specific Link Attribute (ASLA) is a link attribute used in Flexible Algorithm routing that

- advertises application-specific link metrics such as TE metrics and delay metrics for use in algorithm path calculations
- enables precise and flexible traffic engineering by carrying detailed link metrics and constraints, and
- corresponds to the Flexible Algorithm metrics on the interfaces.

### Supported ASLA metrics in Cisco Crosswork Planning

Cisco Crosswork Planning supports ASLA delay and ASLA TE metrics in the Flexible Algorithm calculations. These metrics replace the standard delay and TE metrics during route calculations when using Flexible Algorithm. If these metrics are not available on an interface, Flexible Algorithm excludes that interface from its calculations.

These ASLA attributes are available on the Add or Edit Interface page. You can find them in the **ASLA** section under the **Advanced** tab.

- ASLA delay: A link attribute representing the one-way transmission latency used in Flexible Algorithm path calculations. This attribute corresponds to the interface's Flex-algo metric Min latency.
- ASLA TE metric: A link attribute representing the Traffic Engineering metric used for path selection in Flexible Algorithm routing. This attribute corresponds to the interface's Flex-algo metric TE.
- SRv6 ASLA delay: ASLA delay for SRv6 interfaces.
- SRv6 ASLA TE metric: ASLA TE metric for SRv6 interfaces.

# Create segment lists

Explains the steps to configure and manage segment lists to define the allowed paths for LSPs in your network model.

Use this task to set up segment lists to control and optimize routing of explicit LSPs by specifying the sequence of hops nodes must traverse.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the toolbar, choose **Actions** > **Insert** > **LSPs** > **Segment list**.

   Alternatively, in the Network Summary panel on the right side, click  in the Segment lists page.

   The Segment lists tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon () and check the **Segment lists** check box.
3. In the **Name** field, enter the name of the segment list.
4. From the **Site** and **Node** drop-down lists, choose the relevant site and node.
5. In the **Hops** section, add the segment list hops and rearrange them in the correct order. Click  and drag the entries to reorder.
   - To add a new segment list hop, click  and enter the details.
   - To edit an existing segment list hop, select it from the list and then click .
   - To delete an existing segment list hop, select it from the list, click .
6. To continue creating or editing a segment, use these options and then click **Add**. Choose the site, node, interface, or anycast group.

   > **Note**
   >
   > Interface segment list hops must be the egress interface of a node segment list hop.

   - If the node or source node (for interfaces and LSPs) is within a site, first choose the site.
   - For node segment list hops, do not choose an interface.
   - For interfaces and LSPs segment list hops, first choose the source node.
   - For anycast group segment list hops, do not choose a site, node, or interface.

   After the segment list is created, another way to manage segments is to double-click one from the Segment List Hops table. From there you can change the site, node, and interface, but not the sequence.

## Create anycast groups

Outlines the steps to configure and group nodes as anycast groups for optimized network operation.

Use this task to create anycast groups.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.

2. From the toolbar, choose **Actions** > **Insert** > **Others** > **Anycast group**.

   Alternatively, in the Network Summary panel on the right side, click + in the **Anycast groups** tab.

   The Anycast groups tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon (≡) and check the **Anycast groups** check box.

3. In the **Group name** field, enter a unique name to identify the anycast group.
4. For each node you want to include in the anycast group, check the check boxes in the **Included** column.
5. Click **Add**.

## Create SIDs

Outlines procedures for creating and configuring SIDs, including node SIDs, SRv6 node SIDs, interface SIDs, SRv6 interface SIDs, and flex algorithms to enable segment routing and optimize network path selection.

Use these procedures to create new SIDs.

### Create a node SID

Details the steps to configure node SIDs for routers to enable segment routing and efficient path selection.

Use this task to assign node SIDs to routers within a network to uniquely identify routers, enabling packets to follow the shortest or engineered path through the network.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the Network Design page.
2. From the toolbar, choose **Actions** > **Insert** > **SIDs** > **Node SID**.

   Alternatively, in the Network Summary panel on the right side, click + in the **Node SIDs** table.

   The Node SIDs tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon (≡) and check the **Node SIDs** check box.

3. Select the values for **Site**, **Node**, and **Type** fields.
4. Enter values for **SID** and **SR algorithm**.
5. Use the check box to select **Protected** or **IPv6**, as necessary.
6. Save your changes.

### Create an SRv6 node SID

Outlines the steps to configure a SRv6 node SID in a network design plan.

Use this task to define an SRv6 node SID within a network plan and associate it with a chosen site and node.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **SIDs** > **SRv6 node SID**.

Alternatively, in the **Network Summary** panel on the right side, click + in the **SRv6 node SIDs** table.

The SRv6 node SIDs tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **SRv6 node SIDs** check box.

3. Select values for **Site** and **Node** from the drop-down list.
4. Enter values for **SID** and **SR algorithm**.
5. Use the check box to select **Protected** as necessary.
6. Save your changes.

## Create an interface SID

Defines the steps to configure an interface SID in your network plan.

Use this task to associate a SID with a network interface to enable segment routing features in your network design.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **SIDs** > **Interface SID**.

   Alternatively, in the Network Summary panel on the right side, click + in the **Interface SIDs** table.

   The Interface SIDs tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **Interface SIDs** check box.
3. Select values for **Site**, **Node**, **Interface**, and **Type** fields.
4. Enter the value for **SID**.
5. Use the check box to select **Protected** or **IPv6**, as necessary.
6. Save your changes.

## Create a SRv6 interface SID

Provides instructions to add a SRv6 interface SID in a network plan.

Use this task to set up SRv6 interface SIDs for segment routing in your network.

### Procedure

1. From the toolbar, choose **Actions** > **Insert** > **SIDs** > **SRv6 interface SID**.

   Alternatively, in the **Network Summary** panel on the right side, click + in the **SRv6 interface SIDs** table.

   The SRv6 interface SIDs tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon ( ) and check the **SRv6 interface SIDs** check box.
2. Select values for **Site**, **Node**, and **Interface** fields.
3. Enter values for **SID** and **SR algorithm**.
4. Use the check box to select **Protected**, as necessary.

5. Save your changes.

### Create flex algorithm

Outlines the steps to configure flex algorithm to optimize routing within a network plan.

Use this task to add and configure flex algorithm to apply custom routing behaviors and improve traffic distribution in your network.

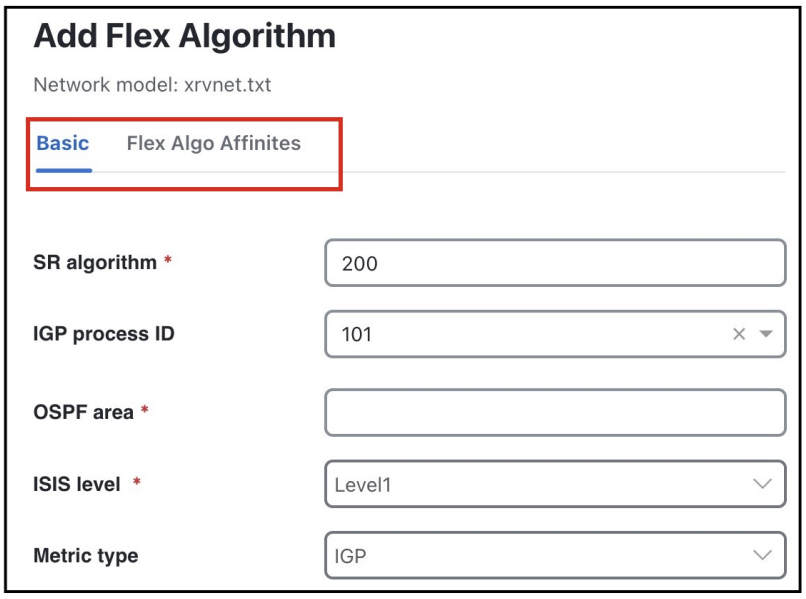
**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose **Actions** > **Insert** > **SIDs** > **Flex algorithm**.

   Alternatively, in the Network Summary panel on the right side, click + in the **Flex algorithms** table.

   The Flex algorithms tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon (☰) and check the **Flex algorithms** check box.

   **Figure 108: Add flex algorithm page**




3. In the **Basic** tab, enter the value for **SR algorithm**.
4. Select values for **IGP process ID**, **OSPF area**, **ISIS level**, and **Metric type**.
5. Click the **Flex Algo Affinities** tab and choose the inclusion or exclusion rule for each affinity.
6. Submit your changes.

## SR-TE protection

Provides information on the attributes of SR-TE protection mechanism, including how they maintain network reliability and ensure traffic continuity.

If the network contains a configurable topology-independent loop-free alternate (TI-LFA), you can simulate SR-TE tunnels before and after convergence.

To see the simulation before convergence, you must be in **Fast reroute** mode (click  or **Actions** > **Edit** > **Network options** and choose **Fast reroute** in the **Simulation convergence mode** section).

In the **IGP and LSP Reconvergence** mode, you can only see the after-convergence routes.

## SR-TE tunnel simulation before convergence

Explains how SR-TE tunnel path simulations model tunnel routing behavior before convergence.

In a normal path state, Cisco Crosswork Planning determines the no-failure route from segment list hop to segment list hop. Between segment list hops, Cisco Crosswork Planning follows the IGP shortest path to the next segment list hop. At each interface along the path, if the connected circuit is down, Cisco Crosswork Planning determines whether the next segment list hop is an interface or a node, and whether the SR tunnel is routable.

For each SR tunnel in a network model, and for each SID in the path, Cisco Crosswork Planning determines the no-failure route.

- If SR FRR is enabled for the interface in the Interfaces table, Cisco Crosswork Planning configures the *protect path state*:
  - Cisco Crosswork Planning derives the post-convergence path from the local node to the next SID hop with the subject circuit removed from the topology.
  - The next SID hop depends on the SID hop type. For an interface SID, the next hop is the remote node of that interface. For a node SID, the next hop is that node.
  - Cisco Crosswork Planning follows the derived after-convergence path hop by hop until it reaches the next SID hop. (If a failure occurs along the way, the SR FRR and the SR LSP do not route.)
  - When Cisco Crosswork Planning reaches the next SID hop, it returns to the normal path state and resumes following the no-fail path to the SID hop.
- If SR FRR is not enabled for the interface in the Interfaces table, the SR tunnel does not route.

## Simulate SR-TE tunnels after convergence

Outlines the simulation of SR-TE tunnels after network convergence.

Cisco Crosswork Planning routes to each segment list hop through the IGP shortest path under all failure conditions.

> **Note**
>
> An SR LSP might be routable before convergence through a TI-LFA, but unroutable after convergence.

## Constraints for SR FRR and TI-LFA

Lists the main constraints that apply to SR FRR protection and TI-LFA path computation.

These constraints define the behavior of protection features and path calculations.

- In the Interfaces table, if the **SR FRR enabled** column displays "true" for a given interface, all SR LSPs with a node or interface SID as the next hop are protected over that interface. If the SRR FRR enabled column displays "false", all SR LSP SIDs are unprotected over that interface.
- Protected versus unprotected adjacency SIDs are not supported. All adjacency SIDs are considered protected.

- Cisco Crosswork Planning does not impose a limit on the TI-LFA path's label stack depth.
- Cisco Crosswork Planning does not explicitly derive P and Q nodes, but instead derives a protect path that aligns with the after-convergence shortest path from the PLR to the next segment list hop.
- If any part of the TI-LFA path encounters a failure before reaching the next segment list hop, the TI-LFA path fails and the SR LSP does not route.

# 23 Optimize Segment Routing

**Topics:**


- SR LSP optimization tools
- SR-TE optimization
- SR-TE bandwidth optimizations


Explains SR-TE and bandwidth optimization methods and reporting for advanced segment routing performance and traffic engineering.

# SR LSP optimization tools

Outlines SR LSP optimization tools that minimize metrics and hops for SR LSPs and set routes based on traffic engineering, alongside bandwidth optimization for congestion mitigation and post-failure reoptimization.

The **SR-TE optimization** is a Cisco Crosswork Planning tool

- creates or updates segment lists to minimize the sum total of user-specified metrics for selected SR LSPs, including inter-AS LSPs, using the fewest number of segment list hops as possible
- sets the routes based on traffic engineering criteria, and
- optimizes a sequence of node or adjacency hops for the SR LSP paths to follow.

The **SR-TE bandwidth optimization** tool lets you reduce the traffic utilization of selected interfaces to below a specified threshold. This ability to mitigate congestion is useful when planning for increased traffic and when determining if service-level agreements can be met in the event of congestion. Additionally, you can use this tool for reoptimizing bandwidth after events, such as network failures and subsequent route reconvergence.

# SR-TE optimization

Explains how SR-TE optimizations enable advanced traffic engineering, flexible path calculation, and network reliability for modern IP and MPLS infrastructures.

You can use the **SR-TE optimization** tool (**Actions** > **Tools** > **SR LSP optimization** > **SR-TE optimization** to design, capacity plan, and manually configure networks to meet the objectives. This tool supports both Inter-Area and Inter-AS functionalities.

This tool helps you design networks for objectives such as:

- TE Metric or Delay Minimization: Minimize distance between hops with respect to metrics other than IGP metrics. These can be either TE metrics configured on interfaces (which can be set proportional to circuit latency) or latencies (delays) for each circuit. An example application is a differentiated service in which latency-sensitive network traffic is routed on shortest latency paths, while the bulk of the traffic routes over cost-optimized paths.
- Avoidance: Create or optimize segment lists so they avoid routing through specified objects (nodes, interfaces, or SRLGs). An example application is routing pairs of LSPs, each over a different plane in a dual-plane network. The same traffic is routed over both LSPs simultaneously, thus improving availability.

  While there is an option to specify a maximum number of hops, doing so might not achieve the lowest possible latency. In this case, the best achievable solution is provided.

  You can additionally avoid unnecessary LSP churn by specifying boundaries (bounds) on the path length and margins within which the shortest path must be optimized.

> **Note**
>
> Unless qualified with “TE” or “IGP,” the term *metric* in this chapter applies to IGP metric, TE metric, or delay.

## SR-TE optimization input properties

Describes the configurable input parameters for SR-TE optimization including path metric selection methods, bounds and margins, and allowable path length settings.

### Path metric optimization options

Describes the configuration options for optimizing the path metric for SR LSPs.

The **Minimize path metric** section defines whether to optimize SR LSPs based on interface IGP metrics, interface TE metrics, or circuit delays. This minimization is for the sum of the metrics along the path. These properties are configurable from the Edit Interface page, and delays can also be set in the Edit Circuit page.

> **Note**
>
> For inter-AS SR LSPs, these metrics are calculated end-to-end for the LSP, not per AS.

**Figure 109: Minimize path metric options**

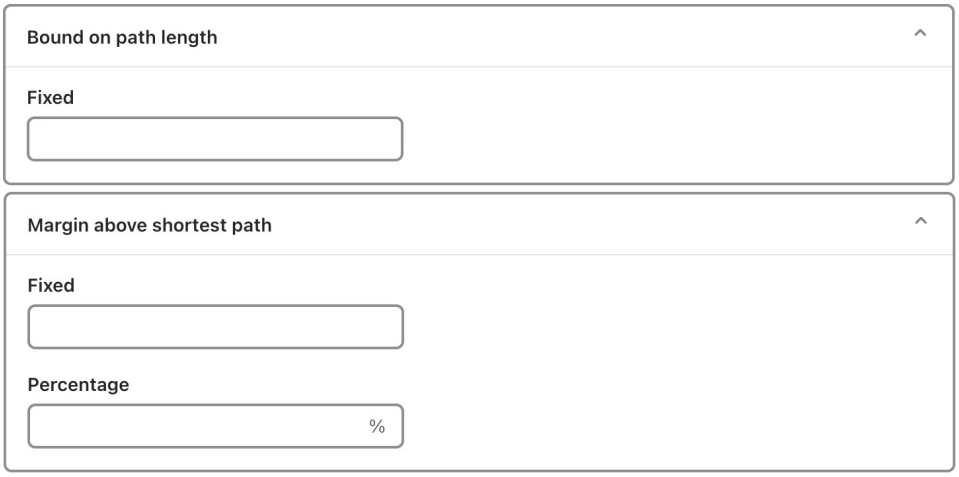



### Bound and margin options

Explains the bounds and margins optimization parameters for LSP path selection in Cisco Crosswork Planning.

Bounds and margins identify which paths to optimize, as well as when to stop optimizing a given path. If you enter a value for more than one restriction, Cisco Crosswork Planning uses the strictest limitation as the optimization target. If there is no bound or margin specified, the system optimizes LSP paths to the best possible solution (lowest total metrics for the LSP path).

**Figure 110: Bound and Margin panels**




### Bound: Maximum acceptable path length

The **Fixed** Bound entry in the **Bound on path length** section lets you set the maximum path metric that is acceptable. Cisco Crosswork Planning tries to optimize LSP paths with metrics that exceed this bound. If a solution adhering to this bound cannot be found, the best possible solution is provided and bound violations are listed in the report. LSP paths that are less than or equal to this bound are not optimized.

Example: If you select to optimize LSP paths based on TE metrics, the value entered is 50, and the sum of TE metrics on the LSP path is 51, that LSP path is optimized.

Enter a number based on the selected path metric. A TE metric is a property values whose total sum for the LSP path cannot be exceeded. The delay is also a property value, but it is in milliseconds (ms). If you enter "50" and you have selected delay as the metric to optimize, this represents 50 ms as the maximum acceptable delay for the LSP path.

#### Margin: Maximum acceptable metric above shortest path

The Margin entries let you identify the acceptable deviation above the shortest achievable path metric. Any existing LSP path with a metric that is less than or equal to the shortest path metric plus the margin is not optimized.

- **Fixed**: The amount by which a metric must be surpassed before it is optimized.

  Example: If an SR LSP route has a delay of 110, a fixed margin set to 10, and the shortest achievable delay path is 100, the current SR LSP is within the margin and will not be updated. If the fixed margin is set to 9, the SR LSP would be optimized.

  100 (shortest path) + 10 (fixed margin) = 110, so all paths greater than 110 are optimized.

- **Percentage**: The amount by which a metric must be surpassed, expressed as a percentage of the shortest path, before it is optimized.
  - Example 1: If an existing SR LSP route has a TE metric of 210, a percentage margin set to 10%, and the shortest achievable TE metric path is 200, the current SR LSP is within the margin and will not be updated. If the current SR LSP had a metric of 225, it would be optimized.

    200 (shortest path) x .10 (percentage margin) = 20, so all paths greater than 220 must be optimized
  - Example 2: If an SR LSP route has a delay of 110, a fixed bound is set to 120, a fixed margin is set to 15, a percentage margin is set to 5%, and the shortest achievable delay path is 100, the strictest of these restrictions takes precedence, and the SR LSP is optimized.

    Fixed Bound = 120

    100 (shortest path) + 15 (fixed margin) = 115

    100 (shortest path) x .05 (percentage margin) = 5, so all paths greater than 105 are optimized because it is the strictest margin.

#### Optimization constraints

Lists the constraints that restrict available optimization options.

These constraints specify restrictions for the optimizations:

- **Maximum segment list hops per SR LSP**: The maximum number of segment list hops that any given segment list can contain after optimization. If no value is specified, Cisco Crosswork Planning creates as many hops as needed to optimize the SR LSP.
- **Avoid**: Do not permit optimized segment lists to route through the selected objects (nodes, interfaces, or SRLGs). This constraint is useful when modeling dual-plane topologies that route disjoint LSPs.
- **Restrict segment node to core nodes**: Segment list node hops must be core nodes (nodes that have their Function property set to "core"), and the local node of segment list interface hops must be a core node. An SR LSP could still route using edge nodes if they are not used as hops.

## Run SR-TE optimization

Describes the steps to execute SR-TE optimization on a network plan to optimize LSPs for improved performance and resource utilization.

Complete these steps to perform SR-TE optimization.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose any of these options:

   From the toolbar, choose **Actions** > **Tools** > **SR LSP optimization** > **SR-TE optimization**.

   Alternatively, choose **Preset workflows** > **Perform optimization**, select **SR LSP Optimization** as the optimization type, select **SR-TE optimization** from the drop-down list, and click **Launch**.
3. Choose the LSPs you want the optimizer to consider and click **Next**.
4. Specify these parameters:
   a) In the **Minimize path metric** section, choose whether to optimize SR LSPs based on interface IGP metrics, interface TE metrics, or circuit delays. For more information, refer to Path metric optimization options on page 355.
   b) In the **Bound on path length** and **Margin above shortest path** sections, specify the values as per your requirement. For more information, refer to Bound and margin options on page 355.
   c) In the **Constraints** section, specify any restrictions for optimization. For more information, refer to Optimization constraints on page 356.
5. (Optional) In the **Tag updated LSPs with** field, override the defaults for how LSPs are tagged (*SROpt*).
6. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:

   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

   The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

   > **Note**
   >
   > Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

7. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:

   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.

- If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

8. Submit your changes.

---

**What to do next**

Upon completion, a report is generated containing the results of the optimization, as well as verifications that the results met the criteria for avoiding nodes and meeting the specified bounds. To view this information, choose **Actions** > **Reports** > **Generated report** and click the **Segment Route TE Optimization** report link in the right panel.

# SR-TE bandwidth optimizations

Describes how SR-TE bandwidth optimization manages network congestion by rerouting demands and using efficient label-switched paths under defined constraints.

The goal when reducing traffic is to reroute as few demands as possible. This is achieved by creating private SR LSPs for demands. The **SR-TE BW optimization** tool also creates LSP paths and segment lists for these LSP paths with the fewest node or interface segment list hops possible, with a maximum of three hops. The final hop is either a node hop or interface whose remote node is the destination of the LSP.

The **SR-TE BW optimization** tool operates on a set of constraints that determine which paths are selected, whether demands and LSPs can be rerouted, and whether demand latency bounds can be surpassed. If the traffic cannot be reduced to below the threshold given these constraints, private SR LSPs are still created to route demands to reduce congestion as much as possible. If interface utilization surpasses the threshold prior to the optimization, it does not go higher. Utilization of non-congested interfaces might increase, though these will not surpass the threshold.

The tool reroutes existing SR LSPs only if they are private to a demand.

> **Note**
>
> The SR-TE BW optimization tool supports Inter-Area functionality.

## Operating modes for SR-TE bandwidth optimization

Lists and describes the available operating modes in the SR-TE BW optimization tool.

The SR-TE BW optimization tool runs in two modes:

- **Analysis**: The **SR-TE BW optimization analysis** tool performs multiple optimizations, one for each failure scenario within the specified failure set. The results of these optimizations are aggregated in a report. For more information on how to run this tool, refer to Analyze the congestion under different failure sets on page 363.
- **Operation**: The **SR-TE BW optimization operation** tool performs an optimization using a specified set of constraints. For more information on how to run this tool, refer to Optimize SR-TE bandwidth on page 362.

An output plan file is generated with the optimization results. Only a single state of the network is considered. For example, if no objects are failed in the input plan file, the optimization is performed for normal operation. If a circuit is failed in the input plan file, the optimization is performed considering this particular failure scenario.

## SR-TE bandwidth optimization input parameters

Lists the sections that describe the key input options and parameters that must be specified to optimize bandwidth for SR-TE.

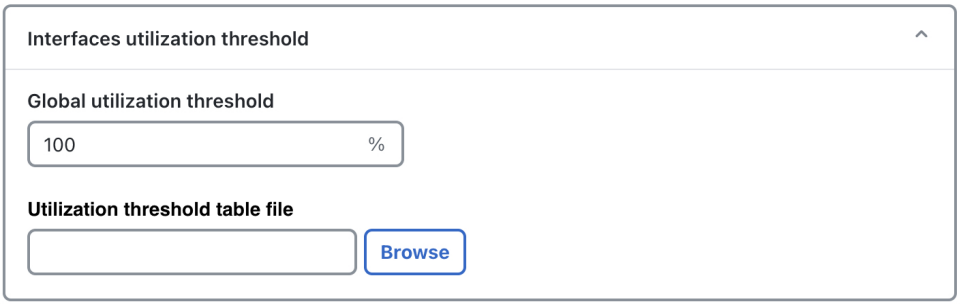
To perform the SR-TE bandwidth optimization, you must specify key input options that determine how traffic is managed, rerouted, and constrained within the network.

### Interface utilization thresholds

Lists supported interface utilization threshold configuration modes, usage guidance, and examples for bandwidth optimization.

Interface utilization thresholds define the maximum allowed usage level for network interfaces. You can specify utilization threshold of an interface using the SR-TE BW optimization tool. If the utilization for an interface is larger than this threshold value, then that interface is considered to be congested, and the tool reduces traffic on the link.

**Figure 111: Interface utilization threshold panel**




- **Global utilization threshold**: Use this option to specify the global threshold value which applies to all interfaces on the network.
- **Utilization threshold tables fil**e: Use this option to upload Utilization threshold tables file which is a .txt file. This file contains the threshold value for a specific interface.

  The file contains a table called `<InterfaceThresholds>` which has the threshold values defined and separated using tabs. The table has three columns: Node, Interface, and Threshold.

  Example: Refer to this sample table:

  ```
  <InterfaceThresholds>
  Node Interface Threshold
  cr1.nyc to_cr2.wdc 70
  ```

  This table specifies that utilization on interface to_cr2.wdc from cr1.nyc must be less than 70 percent.

  Use the **Browse** button to upload the **Utilization Threshold Tables** file. Uploading this file is optional, but when the interface threshold is specified using this file, the tool aims to reduce the utilization on this link to the specified threshold.

### Rerouting demands

Explains the available options and attributes for rerouting demands in analysis and operation modes. Provides details on how each setting affects traffic distribution, demand splitting, path constraints, and latency enforcement.

The analysis and operation modes provide several configurable options for rerouting demands within the network. You can use these options determine how demands are selected, routed, and whether specific constraints such as latency or previously optimized paths are enforced.

**Figure 112: Rerouting demands section**



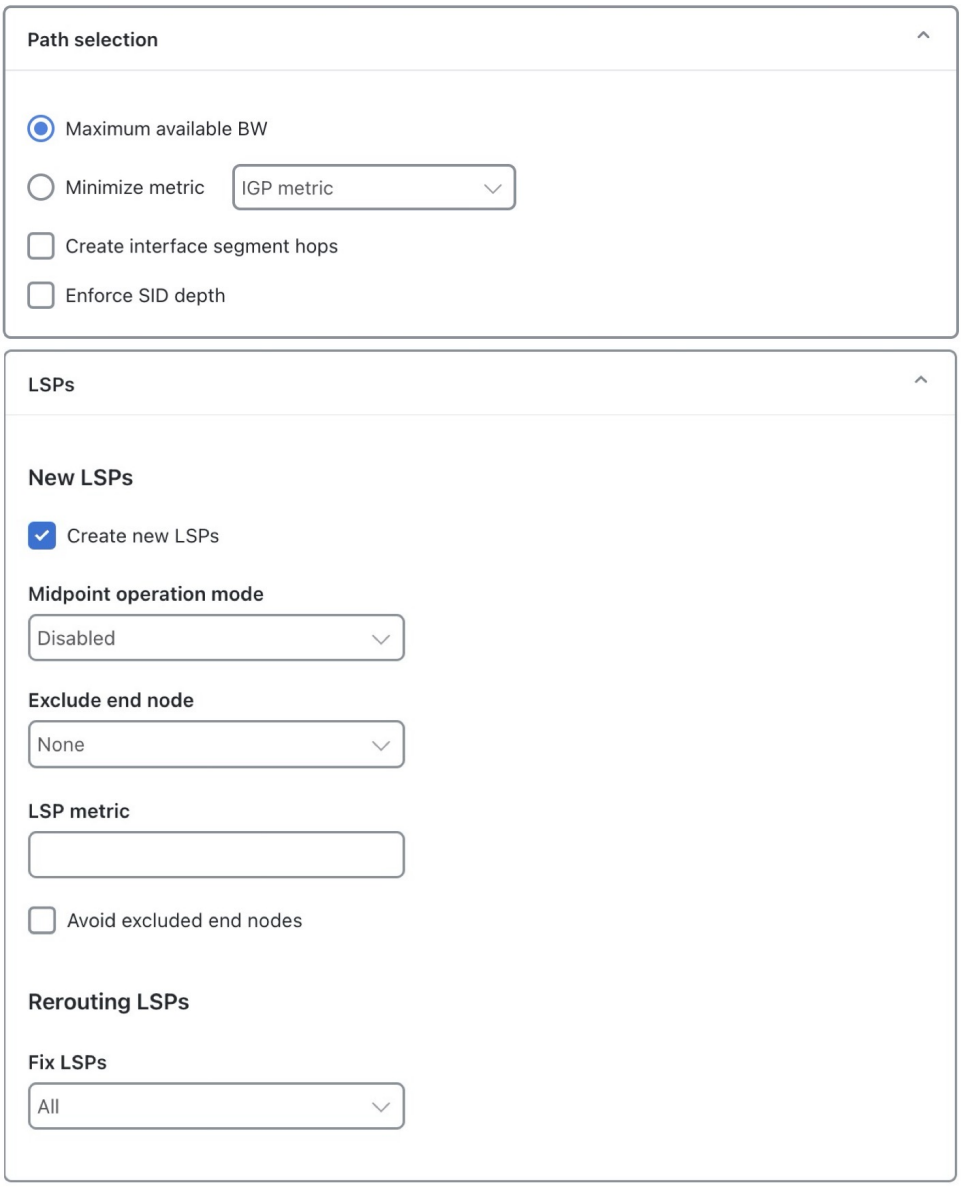

- **Traffic steering mode**: If you choose **Individual Demands**, each demand is individually carried by a private LSP. If you choose **Autoroute**, multiple demands can be carried by a non-private Autoroute LSP.
- **Maximum demands split**: Use this option to split demands into smaller demands. If *n* is the value entered, the demands are split into '*n*+1' smaller demands. The default value is 0.
- **Fix demands**: Selected or tagged demands are not rerouted. This constraint is useful, for example, when you have previously optimized specific LSPs within the network and want to maintain the route of one or more existing demands.
- **Enforce latency bounds**: If checked, demands cannot surpass their configured latency bound. This property is set in the Edit Demand window. If the latency of rerouted demands exceeds the configured bound, the tool does not reroute demands on congested interfaces.

### Bandwidth constraint options

Lists the bandwidth constraint options, attributes, and selection criteria available when rerouting demands in analysis and operation modes.

Both analysis and operation modes use these options for specifying bandwidth constraints.

**Figure 113: Constraints section**




- **Path Selection**: When choosing demand routes, the SR-TE BW Optimization tool uses one of these criteria:
  - **Maximize available BW**: Paths are optimized to achieve the highest available bandwidth on interfaces (Capacity Sim - Traff Sim).
  - **Minimize metric**: Paths are optimized to minimize the sum of the metrics along the path with respect to Delay, TE Metric, or IGP Metric. These properties are configurable from an interface Properties window, and delays can also be set in a circuit Properties window.
  - **Create interface segment hops**: If checked, the tool considers interface hops when trying to drive down utilization.
  - **Enforce SID depth**: If checked, the number of hops in the optimized segment lists are bound by the Max SID Depth defined at the source of the lists.
- **LSPs**:
  - **Create new LSPs**: If checked, new private SR LSPs with optimized routing can be created. If unchecked, new LSPs are not created.
    - **Midpoint operation mode**

    - **Disabled**: Default mode indicating that the new SR LSP source/destination nodes must match demand source/destination nodes.
    - **Demand endpoint proximity**: Specifies that the new SR LSP source/destination nodes may differ from the demand source/destination nodes. Source/Destination nodes closer to demand endpoints are selected.
    - **Congestion proximity**: Specifies that the new SR LSP source/destination nodes may differ from the demand source/destination nodes. Source/Destination nodes closer to congestion points are selected.
  - **Exclude end node**: Nodes selected are not considered as source/destinations nodes for SR LSPs. Default value is **None**. The nodes selected here are ignored if **Midpoint operation mode** is set to **Disabled**.
  - **LSP metric**: When specified, LSPs created by the tool will have their **Metric Value** set to the specified value.
  - **Avoid excluded end nodes**: If the option is selected, the specified nodes will not be included in the newly created LSPs.
- **Rerouting LSPs**
  - **Fix LSPs**: Controls whether or not LSP routes can be modified. This constraint is useful if you want to reroute existing LSPs to mitigate congestion.

## Optimize SR-TE bandwidth

Explains the steps to optimize bandwidth for network interfaces using SR-TE bandwidth optimization operations.

Bandwidth optimization helps ensure efficient use of available network resources. Complete these steps to configure the optimization process, specify the thresholds, constraints, and segment hop requirements.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.

   From the toolbar, choose **Actions** > **Tools** > **SR LSP optimization** > **SR-TE BW optimization operation**.

   Alternatively, choose **Preset workflows** > **Perform optimization**, select **SR LSP Optimization** as the optimization type, select **SR-TE BW optimization operation** from the drop-down list, and click **Launch**.
2. Select whether to optimize bandwidth for all interfaces or selected interfaces.
3. Specify the utilization threshold of an interface and the routing demand options. For more information, refer to Interface utilization thresholds on page 359 and Rerouting demands on page 359.
4. (Optional) In the **Tag updated LSPs with** field, override the defaults for how LSPs are tagged (*SRBWOpt*).
5. Specify constraints. For more information, refer to Bandwidth constraint options on page 360.
6. Select the required segment hop constraints from the **Node segment hop** and **Interface segment hop** drop-down lists, and click **Next**.
7. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:
   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.

- Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
- Schedule: Set the time at which you want to run the tool.

The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

> **Note**
>
> Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

8. (Optional) If you want to display the result in a new plan file, specify a name for the new plan file in the **Display results** section.

   In the previous step:

   - If you have selected to run the task immediately, by default, the changes are applied on the current plan file. If you want to display the results in a new file, select the **Display results in a new plan file** check box and enter the name of the new plan file.
   - If you have scheduled the task to run at a later time, by default, the results are displayed in the *Plan-file-1*. Update the name, if required.

9. Submit your changes.

## Analyze the congestion under different failure sets

Describes the steps to evaluate and mitigate network congestion by analyzing optimization results across multiple failure scenarios in your design plan.

Use this task to perform multiple optimizations, one for each failure scenario within a specified failure set. The analysis aggregates results, helping you identify configurations that minimize congestion under different failure scenarios.

### Procedure

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.
2. From the toolbar, choose any of these these options:

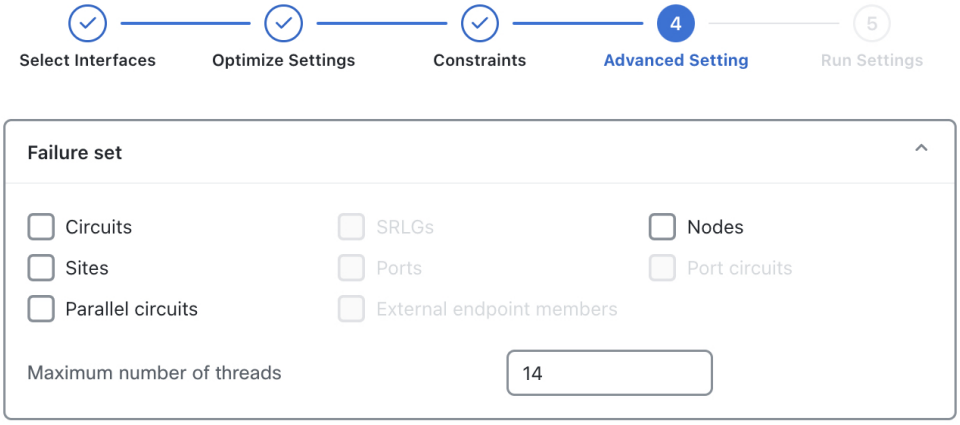
   From the toolbar, choose **Actions** > **Tools** > **SR LSP optimization** > **SR-TE BW optimization analysis**.

   Alternatively, choose **Preset workflows** > **Perform optimization**, select **SR LSP Optimization** as the optimization type, select **SR-TE BW optimization analysis** from the drop-down list, and click **Launch**.

3. By default, the analysis tool presents you with the bandwidth optimization options. Configure these options, as required. For more information, refer to Optimize SR-TE bandwidth on page 362.
4. Click **Next** to go to the **Advanced Setting** page. Click the **Failure set** section.
5. Select the failure sets that you want the optimizer to consider (circuits, nodes, sites, and so on).

Entries appear dimmed if they are not available in your design plan.

**Figure 114: Failure sets panel**




6. (Optional) Specify the maximum number of threads.

   By default, the optimizer tries to set this value to the optimal number of threads based on the available cores.

7. Select the required segment hop constraints from the **Node segment hop** and **Interface segment hop** drop-down lists, and click **Next**.
8. On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose one of these **Execute** options:

   - Now: Choose this option to execute the job immediately. The tool runs and changes are applied to the network model immediately. A summary report appears. You can access the report any time later using **Actions** > **Reports** > **Generated reports** option.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. Set these options:
     - Priority: Select the priority of the task.
     - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
     - Schedule: Set the time at which you want to run the tool.

   The tool runs at the scheduled time using the selected engine profile. You can track the status of the job at any time using the Job Manager page (from the main menu, choose **Job Manager**). Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

   **Note**

   Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

9. Submit your changes.

## Bandwidth optimization report

Summarizes the results and key network statistics captured after running the SR-TE bandwidth optimization tool.

After completion of a basic optimization, the SR-TE BW optimization tool generates a summary report to identify key changes due to the optimization.

The tool tags the LSPs with *SRBWopt* and generates a new plan file with an *-SRBWopt* suffix. This plan file opens, showing the LSPs table that is filtered to these rerouted (and newly tagged) LSPs. Saving this plan file then simplifies the process of identifying which LSPs to reconfigure in the network. Both the tags and the new plan file names are editable.

After running the bandwidth optimization, the optimizer creates a summary report that includes:

- Number of interfaces that exceeded the utilization both before and after the optimization
- Maximum utilization both before and after the optimization
- Number of demands that were rerouted
- Number of LSPs that were created and rerouted
- Average segment list length
- Number of latency bound violations both before and after the optimization

# Part IV

# Access Reports, Jobs, Add-ons, Patch Files, and Changeover Tool

**Topics:**


- Access Reports
- View Jobs
- Manage Add-ons
- Update Configuration from One File to Another
- Create and Use Patch Files
- Export Routes

# 24 Access Reports

**Topics:**


- Reports in Cisco Crosswork Planning
- Plan comparison reports
- Traffic comparison reports
- View a generated report


Explains the types of reports available, including plan and traffic comparisons, and guides you on how to view and interpret report data for informed decision-making.

# Reports in Cisco Crosswork Planning

Explains about reports in Cisco Crosswork Planning that enable you to compare network topology, configuration, and traffic information between plan files.

Cisco Crosswork Planning lets you generate reports that compare IP/MPLS topology and configuration information, demand routings, LSP routings, or traffic between two plan files. You can use these reports to:

- Plan for upgrades: Compare plan files before and after topology and configuration changes are made to a current network. For example, you can compare the original plan file to a proposed new plan in which circuits or nodes are added or upgraded.
- Mitigate congestion: Compare plans before and after making configuration changes to mitigate congestion due to a failure or planned maintenance, such as metric changes or LSP additions and reroutings.
- Validate simulations: Compare the simulated traffic under failure in one plan to the measured traffic after a failure has occurred in another plan to validate the accuracy of the simulation.
- Audit plans: Review any changes between two plan files.
- Identify demand and LSP reroutes: Compare plans to determine which demands or LSPs experienced routing changes, for example as the result of an interface metric change or capacity upgrades.

Upon comparing plan files, a report of the results opens automatically. You can access this report later by choosing **Actions > Reports**.

# Plan comparison reports

Explains how plan comparison reports enable you to compare objects between two plan files to identify differences and similarities in network configurations.

A plan comparison report is a report that

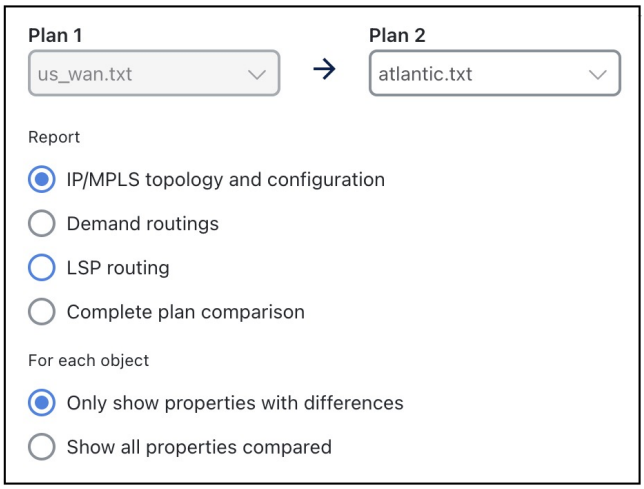
- lets you compare objects between two plan files
- helps to determine which objects are present in one plan but not in the other, and
- identifies objects present in both plans and highlights differences in their properties.

### Types of plan comparison reports

Plan comparison reports include several types that focus on different aspects of network topology and configuration.

- IP/MPLS topology and configuration: Compares objects describing the topology (such as nodes, circuits, and interfaces), configured objects (such as LSPs and LSP paths), and related properties (such as IGP metrics and capacity). For a complete list of compared properties, refer to Table 36: Plans comparison report: IP/MPLS topology and configuration on page 372.
- Demand routings: Compares demand paths and the properties that indicate routing changes, such as Path Length and Max Latency. For a complete list of compared properties, refer to Table 37: Plans comparison report: Demand routings on page 375.
- LSP routing: Compares properties of LSPs that indicate routing changes, such as TE Path Metric and Active Path Sim.
- Complete plan comparison: Compares all tables in the table schema, except for internal simulation caching tables.

**Figure 115: Compare Plans page**




## Create plan file comparison reports

Describes how to configure plan file comparison reports to analyze differences between two plan files. This helps you verify changes and understand network modifications efficiently.

This topic describes how to generate comparison reports between two plan files, allowing you to identify differences in topology, configuration, and routing.

Use this task when you need to compare two plan files to review changes or validate network updates.

### Before you begin

Ensure you have access to the plan files you want to compare.

#### Procedure

1. Open the two plan files you want to compare. For more information, refer to Open plan files on page 36.
2. In either of these two plan files, choose **Actions** > **Reports** > **Compare plans** from the toolbar.
3. Select the plan file to which you are comparing the open plan.
4. Select the type of comparisons you want to generate:
   - IP/MPLS topology and configuration
   - Demand routings
   - LSP routing
   - Complete plan files
5. Select whether to show only properties with differences or whether to show all properties compared.
6. Click **Compare**.

The comparison report displays the differences and similarities between the selected plan files, helping you review network changes effectively.

## Report columns

Lists the types of columns in each Cisco Crosswork Planning table and describes their roles in plan comparison reports.

Each Cisco Crosswork Planning table contains key columns that uniquely identify objects. In the Interfaces table, for example, the Node and Interface columns represent the name of the node containing the interface and the name of the interface itself. In a Plan Comparison report, objects in different plans are matched only if their key columns are identical. Key columns determine whether an object is present in both plans or only one.

The Plan Comparison report displays three other types of columns. Table 36: Plans comparison report: IP/MPLS topology and configuration on page 372, Table 37: Plans comparison report: Demand routings on page 375, and Table 38: Plans comparison report: LSP routings on page 376 list the properties reported upon and their associated column type.

- Information only (Info): There are no comparisons made. Information is reported for the plan file from which you are running the report (Plan 1).
- Differences (Diff): Specific properties are compared for each object that has matching key columns across the two plan files. Objects are then identified as belonging in either Plan 1 only (file from which you are running the report), in Plan 2 only (plan file to which you are comparing), or in both Plan 1 and Plan 2. If the object exists in both plan files, but has different properties, a Diff column shows T (true) if there are any differences or F (false) if there are none.

  Differences are based on the current state of the open plan files. For example, if a circuit has failed, the demand routings change to route around the failure.
- Summary Differences (Summ Diff): These are not columns within Cisco Crosswork Planning tables. They are T/F (true or false) values, depending on some differences in the table objects that are not represented in the table columns. For example, if a common LSP path is configured differently between plans, this difference is represented in a summary difference column in the LSP section. If demand routings differ, this is represented in a summary difference column in the Demands section.

> **Note**
>
> For easy reference, these tables list columns in the order of type first: Key, Info, Diff, and Summ Diff. Properties within each column type are alphabetical.

**Table 36: Plans comparison report: IP/MPLS topology and configuration**

| Compared Table | Compared Columns | Comparison Column Type | Summary Difference Description |
|---|---|---|---|
| Interfaces | Name | Key | |
| | Node | Key | |
| | Remote Node | Info | |
| | Affinities | Diff | |
| | Area | Diff | |
| | Capacity | Diff | |
| | Circuit | Diff | |
| | Description | Diff | |
| | FRR Protect | Diff | |
| | IGP Metric | Diff | |

| Compared Table | Compared Columns | Comparison Column Type | Summary Difference Description |
|---|---|---|---|
| | IP Address | Diff | |
| | PC Min BW | Diff | |
| | PC Min Links | Diff | |
| | Resv BW | Diff | |
| | TE Metric | Diff | |
| Circuits | InterfaceA | Key | |
| | InterfaceB | Key | |
| | NodeA | Key | |
| | NodeB | Key | |
| | Name | Info | |
| | Active | Diff | |
| | Capacity | Diff | |
| | Delay | Diff | |
| Nodes | Name | Key | |
| | Active | Diff | |
| | AS | Diff | |
| | BGP ID | Diff | |
| | ECMP Max | Diff | |
| | IP Address | Diff | |
| | IP Manage | Diff | |
| | Model | Diff | |
| | OS | Diff | |
| | Vendor | Diff | |
| LSPs | Node | Key | |
| | Source | Key | |
| | Active | Diff | |
| | Destination | Diff | |
| | Exclude | Diff | |
| | FRR Link Protect | Diff | |

| Compared Table | Compared Columns | Comparison Column Type | Summary Difference Description |
|---|---|---|---|
| | Hold Pri | Diff | |
| | Hop Limit | Diff | |
| | Include | Diff | |
| | Include Any | Diff | |
| | Load Share | Diff | |
| | Metric | Diff | |
| | Metric Type | Diff | |
| | Setup BW | Diff | |
| | Setup Pri | Diff | |
| | Unresolved Destination | Diff | |
| | LSPs Path Diff | Summ Diff | T (true) if there are any differences in the LSP Paths table for this LSP. If two LSPs have differences in their named path hops or LSP paths, these differences are propagated to this column. |
| LSP Paths | Node | Key | |
| | Path Option | Key | |
| | Source | Key | |
| | Exclude | Diff | |
| | Hold Pri | Diff | |
| | Hop Limit | Diff | |
| | Include | Diff | |
| | Include Any | Diff | |
| | Path Name | Diff | |
| | Setup BW | Diff | |
| | Setup Pri | Diff | |
| | Standby | Diff | |
| | Named Path Diff | Summ Diff | T (true) if there are any differences in the Named Paths table for this LSP path. |
| Named Paths | Active | Key | |
| | Name | Key | |

| Compared Table | Compared Columns | Comparison Column Type | Summary Difference Description |
|---|---|---|---|
| | Source | Key | |
| | Named Path Hops Diff | Summ Diff | T (true) if there are any differences in the Named Path Hops table for this named LSP path. |
| Named Path Hops | Name | Key | |
| | Source | Key | |
| | Step | Key | |
| | Interface | Diff | |
| | IP Address | Diff | |
| | Node | Diff | |
| | Type | Diff | |
| | Unresolved Hop | Diff | |

**Table 37: Plans comparison report: Demand routings**

| Compared Table | Compared Columns | Comparison Column Type | Summary Difference Description |
|---|---|---|---|
| Demands | Destination | Key | |
| | Name | Key | |
| | Service Class | Key | |
| | Source | Key | |
| | Destination Site | Info | |
| | Source Site | Info | |
| | Active | Diff | |
| | ECMP Min % | Diff | |
| | Max Latency | Diff | |
| | Path Metric | Diff | |
| | Path Diff | Summ Diff | T (true) if there are any differences in routing of the demand. |

**Table 38: Plans comparison report: LSP routings**

| Compared Table | Compared Columns | Comparison Column Type | Summary Difference Description |
|---|---|---|---|
| LSPs | Name | Key | |
| | Source | Key | |
| | Active Path Sim | Diff | |
| | Destination | Diff | |
| | Routed | Diff | |
| | Shortest TE Path | Diff | |
| | TE Path Metric | Diff | |
| | Actual Path Diff | Summ Diff | T (true) if there are any differences in routing of the actual LSP path. |
| | Simulated Path Diff | Summ Diff | T (true) if there are any differences in routing of the LSP. |

## Report summary

Explains how the report summarizes object counts, property differences, and display content based on selected report options.

The report **Summary** section is a report component that

- shows the number of objects in each plan file
- displays the number of objects in both plan files, and
- indicates the number of objects with property differences.

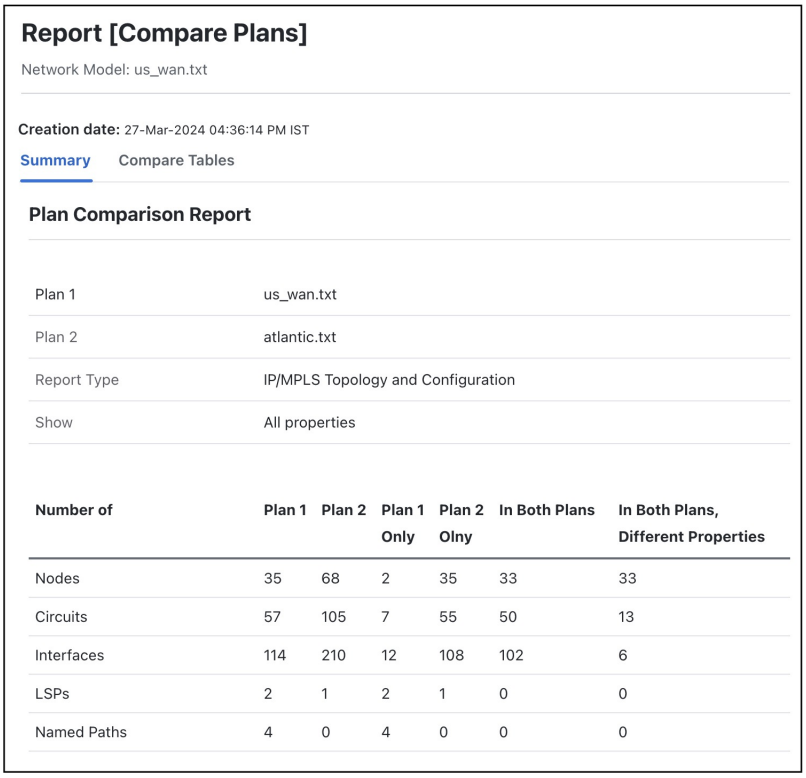
### Report section details

An individual section is generated for each object that appears in only one plan file, and if there are differences between properties, a section is generated showing these differences. Up to three sections may be generated per compared table: one for objects appearing only in Plan 1, one for objects appearing only in Plan 2, and one listing objects existing in both plan files, with differences noted as "true" or "false".

Other than Summary, the report sections and columns (properties) that appear depend on the options selected when generating the report.

- If you select to show only properties with differences, only those columns with differences in properties appear in the report. If only one value is different, all objects are listed (not just the one with the different property).
- If you select to show all properties compared, all compared properties appear in the report whether there are differences between the two plan files or not.

**Figure 116: Plan comparison report summary example**



Report [Compare Plans]

Network Model: us_wan.txt

**Creation date:** 27-Mar-2024 04:36:14 PM IST

Summary Compare Tables

**Plan Comparison Report**

| | |
|---|---|
| Plan 1 | us_wan.txt |
| Plan 2 | atlantic.txt |
| Report Type | IP/MPLS Topology and Configuration |
| Show | All properties |

| Number of | Plan 1 | Plan 2 | Plan 1 Only | Plan 2 Olny | In Both Plans | In Both Plans, Different Properties |
|---|---|---|---|---|---|---|
| Nodes | 35 | 68 | 2 | 35 | 33 | 33 |
| Circuits | 57 | 105 | 7 | 55 | 50 | 13 |
| Interfaces | 114 | 210 | 12 | 108 | 102 | 6 |
| LSPs | 2 | 1 | 2 | 1 | 0 | 0 |
| Named Paths | 4 | 0 | 4 | 0 | 0 | 0 |

**Report section example**

Nodes A, B, and C are in Plan 1, and B, C, and D are in Plan 2, respectively. Nodes B and C have matching key columns. The B nodes have identical properties, but the C nodes have different IP addresses. Therefore, in the Summary section, the **In Both Plans** column shows 2, and the Different Properties column shows 1.

## Traffic comparison reports

Explains how traffic comparison reports enable you to compare traffic values within a single plan or between two plan files, highlighting differences across network objects and traffic columns.

A traffic comparison report is a report that

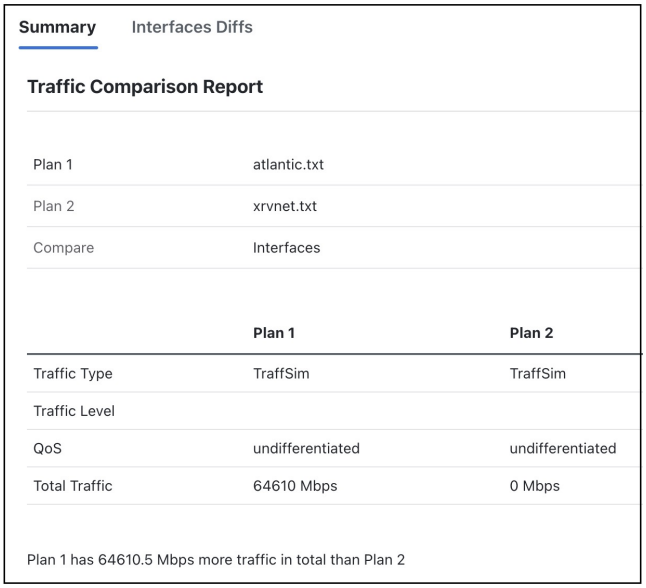
- compares traffic values within a single plan or between two plan files, and
- allows comparison of traffic on one object at a time, such as nodes, interfaces, circuits, demands, LSPs, multicast flows, and flows.

Each object has a different set of comparisons from which to choose. For example, interface traffic and capacity can be compared, and LSP traffic and setup bandwidth can be compared. Another example would be to compare the simulated interface traffic in a plan file before a failure to the measured interface traffic of a plan file in which a failure occurred. Such a comparison lets you determine which interfaces show the greatest differences.

For each plan file compared, the current traffic level and either the selected service class or queue are used. They can be different in each plan.

The Summary section displays a high-level overview of the comparison. Additionally, a section detailing the difference on the selected object and traffic columns is generated and named accordingly.

**Figure 117: Traffic comparison report**




## Create traffic comparison reports

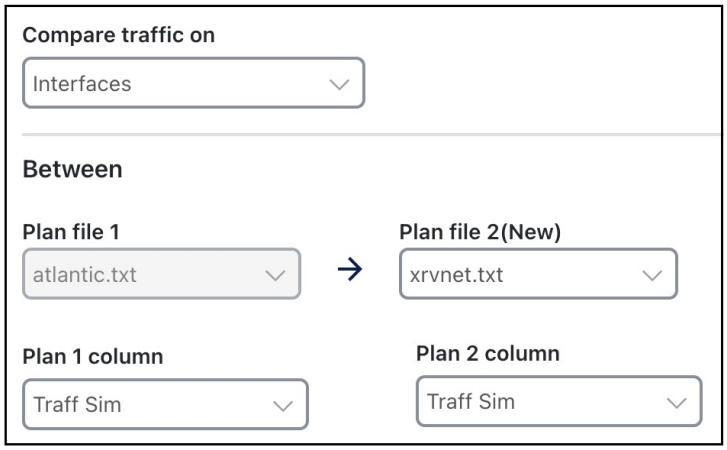
Describes how to configure traffic comparison reports to analyze and compare network traffic between different plan files. This task helps you identify differences in traffic patterns across various objects such as interfaces, circuits, and nodes.

This topic describes how to create traffic comparison reports to evaluate and analyze network traffic across multiple plan files. This enables you to identify changes and trends in traffic for specific objects.

Traffic comparison reports are useful for understanding how network traffic varies between different plan files. You can compare traffic for objects such as interfaces, circuits, and nodes to gain insights into network performance and changes.

### Procedure

1. If comparing traffic from the open plan file to another, open the other plan file.




2. From the toolbar, choose **Actions** > **Reports** > **Compare traffic**.

   The currently opened plan file is identified as Plan 1 in the Compare Traffic window that opens.

3. From the **Compare traffic on** drop-down list, select the object type on which to compare traffic (interfaces, circuits, nodes, and so on).
4. Select the traffic column in the current plan that you want to compare.

5. To change the plan file 2, select a different plan file from the Plan file 2 drop-down.
6. Select the traffic column being compared for Plan 2.
7. Click **Compare**.

The traffic comparison report displays differences in traffic between the selected plan files for the chosen object types and columns. You can review the results to identify changes and trends.

## View a generated report

Verify that you can view generated reports in the application. Configure the report viewing process to access report sections and save reports for later use.

This topic describes how to access and view generated reports within the application. You can examine various report sections and save reports for future reference.

Reports open automatically after generation. If you close a report, you can reopen it from the application toolbar. Saving the plan file preserves generated reports, but rerunning a report will overwrite the previous version.

### Before you begin

Ensure that a report has been generated in the application.

#### Procedure

1. To reopen a report after closing it, choose **Actions** > **Reports** > **Generated reports** from the toolbar

   All reports generated on the plan file appear in a single window and are selectable.

2. To save reports for later use, you must save the plan file.

   Choose **Actions** > **File** > **Save** or **Save as** from the toolbar.

# 25 View Jobs

**Topics:**


- Job manager
- View job details
- Access output plan files from job manager
- Run tools or initializers using CLI
- Run external scripts


Outlines how to manage and review jobs, access output files, and execute tools or scripts through the CLI or external scripts.

# Job manager

Explains how job manager enables scheduling, prioritizing, and managing background jobs in Cisco Crosswork Planning, including execution options and job control features.

A job manager is a background job management tool that

- enables you to schedule tools and initializers to run as background jobs, also known as asynchronous jobs, and
- offers options to execute jobs using CLI and custom scripts.

The asynchronous jobs execute in the background on design engine instances configured during deployment.

The Job manager provides these options for background jobs:

- Set a priority before submission
- Schedule jobs to run at a required time
- Cancel jobs before they move to the Running state

## User role permissions for job manager

Lists user role permissions for Job Manager, including actions available to regular and admin users.

This topic describes the permissions assigned to different user roles in the Job Manager page.

- Any user can view the jobs submitted by any other user.
- Any user can view the results of the jobs submitted by any other user.
- Any user can cancel (abort) the jobs submitted only by themselves.
- Only admin users can cancel (abort) the jobs submitted by any user.

# View job details

Describes how to verify the details of submitted background jobs using the Job Manager. This allows you to monitor job status, review output files, and access comprehensive job information efficiently.

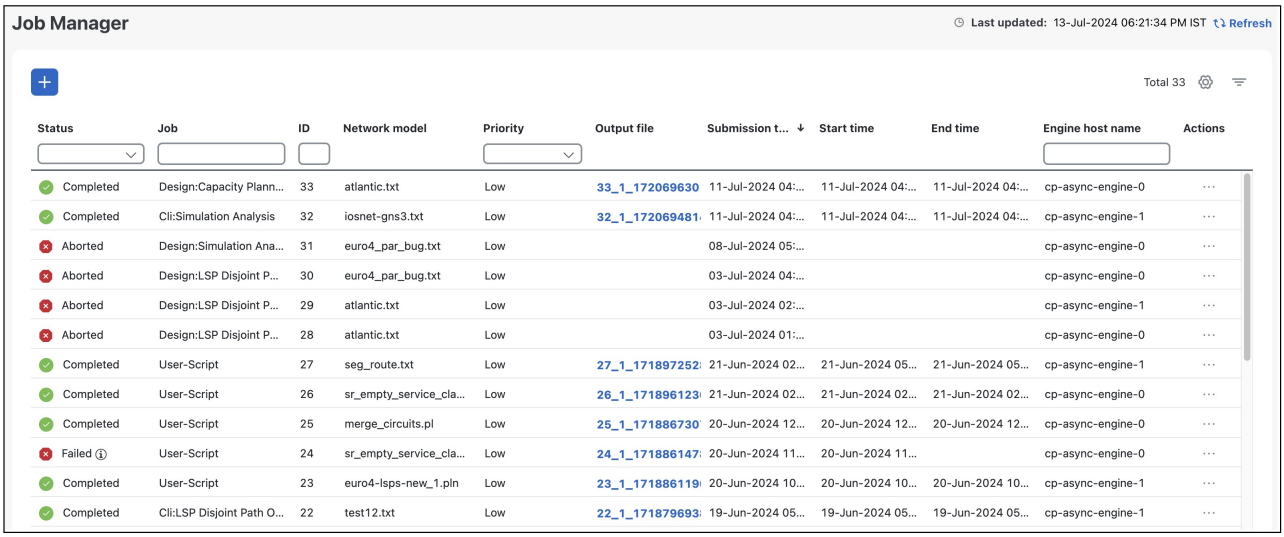
Complete these steps to see the details of the submitted background jobs.

**Procedure**

1. From the main menu, choose **Job Manager**.

   The Job Manager page opens, displaying the list of all jobs submitted as background jobs.

   **Figure 118: Job manager page**




The jobs are displayed in descending order of creation time, with the most recent job shown first. To sort the data in the table, click a column heading. Click the column heading again to toggle between ascending and descending sort order.

Click to show or hide the columns. Check the check boxes for the columns you want to show and uncheck the ones you want to hide.

Use to toggle the display of the floating filters at the top of each column. Using this filter you can set the filter criteria on one or more columns in the table. To clear all the filters, click the **X** icon in the Filters field that appears above the table.

2. The **Status** column shows the types of states: Completed, Scheduled, Queued, Running, Failed, and Aborted. For any failed job, click shown next to the error for more information. For more information on all the columns, refer to Table 39: Job manager column details on page 383.
3. The **Output file** column lists the tar file that is generated after the job has been successfully completed. Click the name of the output file to view the output plan file details. You can either download this plan file or import to the user space. For more information, refer to Access output plan files from job manager on page 384.
4. Click ••• > **View details** under the **Actions** column to view a comprehensive summary of the details of the submitted jobs.

   If a job is in Scheduled, Queued, or Running state, you can use ••• > **Cancel** under the **Actions** column to abort the job.

## Job manager column details

Explains each column displayed on the Job Manager page, including their purpose and available options.

This topic describes each column displayed on the Job Manager page, including their purpose and available options.

**Table 39: Job manager column details**

| Column | Description |
|---|---|
| Status | Indicates the status of the submitted background job. The available status types are: Completed, Scheduled, Queued, Running, Failed, and Aborted. |

| Column | Description |
|---|---|
| Job | Indicates the type of the job.<br>If a job is submitted via<br>• tools or initializers in the **Network Design** page, the Job column displays "Design: *Tool/initializer name*"<br>• add-ons, the Job column displays "Add-on: *Add-on name*"<br>• **Using CLI** option in the **Job Manager** page, the Job column displays "Cli: *Tool/initializer name*", and<br>• **Using Script** option in the **Job Manager** page, the Job column displays "User-Script". |
| ID | Indicates the Job ID. When a job is submitted as a background job, a Job ID is created, which is displayed in this column. |
| Network model | Indicates the network model in which the job was submitted. |
| Priority | Indicates the priority of the job. When a job is submitted as a background job, you have the option to assign priority, which is displayed in this column. The available options are: High, Medium, and Low. |
| Options | Indicates the configuration options used while submitting the job. This column is hidden by default. |
| Output file | Indicates the file that is generated after the job has been successfully completed. Click this file to access the updated network model. |
| Submission time | Indicates the timestamp at which the job was submitted. |
| Start time | Indicates the timestamp at which the job execution began. |
| End time | Indicates the timestamp at which the job execution was completed successfully. |
| Engine host name | Indicates the engine name used for the job execution. |
| Schedule | Indicates the timestamp at which the job was scheduled to begin. This column is hidden by default. |
| Submitted by | Indicates the name of the system or user that submitted the job. For scheduled jobs, this column displays "system" as the job is submitted by the system at the scheduled time. For all other jobs, the "Submitted by" and "Created by" columns display the same information. This column is hidden by default. |
| Created by | Indicates the user who submitted the job. This column is hidden by default. |
| Actions | Displays the comprehensive summary of the submitted jobs when you use the ••• > **View details** option under this column. You can also use ••• > **Cancel** to abort the job if it is in the Scheduled, Queued, or Running state. |

## Access output plan files from job manager

Explains how to access the output plan files from the Job Manager page for further analysis or visualization.

Access output plan files from Job Manager to retrieve or download files generated by background jobs. This is useful for reviewing results, importing plan files into your user space, or downloading files for local analysis.

Use this task when you need to view, download, or import output plan files generated by jobs in the Job Manager.

**Procedure**

1. From the main menu, choose **Job Manager**.

   The Job Manager page opens and displays all background jobs.

2. Click the .tar file in the **Output file** column to access the generated output files. To view the contents of a specific file, click on its file name.

   If the file is in a human-readable format, its contents display correctly. For file types such as a .pln file, which are not human-readable, the contents may appear as garbled text. In such cases, download the file and open it with an appropriate application to view its contents properly.

3. To import the output plan file into your user space:
   a) Click the name of the output file.

      The **Extracted result** page opens and displays all output files and logs.

   b) Locate the output plan file and click ••• > **Export to userspace**.

   The plan file is imported into your user space and becomes available for visualization.
4. To download the file to your local machine:
   a) Click the name of the output file.

      The **Extracted result** page opens and displays all output files and logs.

   b) Locate the required file and click ••• > **Download**.

   The selected file is downloaded to your local machine.
5. To download the complete output results, click ••• > **Download output file** under the **Actions** column.
   The tar file is downloaded to your local machine. Extract the file and import the required plan file into your user space. For more information, refer to Import plan files from the local machine on page 34.

After completing these steps, you can access, download, or import output plan files from Job Manager. The files are available for review, visualization, or further processing as needed.

## Run tools or initializers using CLI

Explains how to configure tools or initializers using CLI to automate network tasks.

In addition to running the tools and initializers from the **Network Design** page, Cisco Crosswork Planning supports running them using CLI.

Complete these steps to run the tools or initializers using CLI:

**Procedure**

1. From the main menu, choose **Job Manager**.

   The Job Manager page opens, displaying the list of all jobs submitted as background jobs.

2. Click + > **Using CLI**.

A list of all available tools and initializers appears.

3. Select the required tool or initializer and click **Next**.

   A list of all available plan files appears.

4. Select the plan file in which you want to run the tool or initializer, and click **Next**.

   The number of plan files you can select varies based on the tool or initializer chosen in the previous step. An error message appears at the top when the allowed number is exceeded.

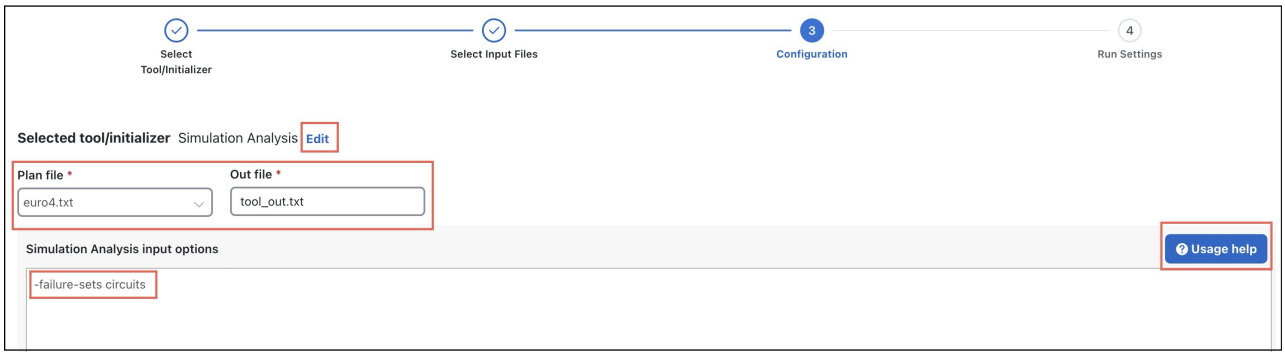
5. Enter the input configuration options in the text field, and click **Next**.

   For help with this, click **Usage Help**. The Usage Help page displays the details of the configuration parameters that you can use in a specific tool or initializer. It also displays syntax, required configuration options, optional configuration options, and an example command.

   > **Note**
   >
   > You can enter only the Optional options. You need not enter the tool name and the Required options.

   On this page, you can also update the tool or initializer, and name of the output file.

   - To change the tool or initializer, click **Edit** next to the tool name that you already selected.
   - To update the name of the output file, enter the new name in the **Output file** field. By default, *tool_out.txt* is used as the file name.




6. On the **Run Settings** page, choose whether to execute the task immediately or schedule it for a later time. Choose from these **Execute** options:

   - Now: Choose this option to run the tool/initializer immediately using the selected engine profile.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. The tool runs at the scheduled time and using the selected engine profile.

   Based on the selection, set these options:

   - Priority: Select the priority of the task.
   - Engine profiles: Select the engine profile as needed. This section lists all the available asynchronous engine profiles.
   - Schedule: Set the time at which you want to run the tool. This option is available only for scheduled jobs.

7. Track the status of the job on the Job Manager page. The job name is prefixed with "Cli:" in the **Job** column. Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

The tool or initializer runs as configured, and the output plan file is available for import and visualization in user space.

## Run external scripts

Outlines how to configure external scripts to automate network model operations using the OPM and Design RPC Python libraries. This enables you to run customized scripts and CLI tools efficiently in your environment.

In Cisco Crosswork Planning, OPM Python library (previously known as OPM API) and Design RPC Python libraries (previously known as Design RPC APIs) are supported via scripts. You can upload customized scripts that are created using these APIs. You can also use the scripts to run the CLI tools.

OPM Python libraries provide a powerful Python API to manipulate network models. It lets you operate on the network without having to worry about device-specific properties. Even if the underlying routers are replaced by routers from a different vendor, the API calls remain exactly the same.

For a sample Python script, refer to Example: Run external scripts on page 388.

Complete these steps to run the external scripts.

**Procedure**

1. From the main menu, choose **Job Manager**.

   The Job Manager page opens, displaying the list of all jobs submitted as background jobs.

2. Click  > **Using script**.

   A list of scripts available in the user space appears.

3. Select the required script and click **Next** to continue.

   If the required script is not available in the list, use the **Upload script** option to upload the script. The script name appears under Available Scripts. You can also find the script file under the **Network Models > My user space > All files** section.

4. Select the network models on which you want to run the script, and click **Next**.
5. Enter the input configuration options in the text field using this format and click **Next**.

   ```
   script_name arg1 arg2
   ```

   Use the **Edit** button to change the script file, if required.

6. On the **Run Settings** page, choose whether to execute the task immediately or schedule it for a later time. You can also set priority for the jobs. Choose from these **Execute** options:

   - Now: Choose this option to run the script immediately.
   - As a scheduled job: Choose this option to execute the task as an asynchronous job. If you choose this option, set the time at which you want to run the script. The script runs at the scheduled time.

7. Track the status of the job on the Job Manager page. The **Job** column displays "User-Script" for the script jobs. Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.

---

## Example: Run external scripts

Provides an example of how to use external scripts in Cisco Crosswork Planning to append interface descriptions.

This topic explains how to use an external Python script with Cisco Crosswork Planning to append a description to every interface in the network, demonstrating the process and command usage.

The sample Python script (ext_exe_eg.py) appends a description to every interface in the network with "My IGP metric is *<value>*."

Contents of **ext_exe_eg.py**:

```
import sys
from com.cisco.wae.opm.network import Network

src  = sys.argv[1]
dest = sys.argv[2]

srcNet =  Network(src)

for node in srcNet.model.nodes:
    cnt = 1
    for iface in node.interfaces:
        iface.description = 'My IGP metric is ' + str(iface.igp_metric)
        cnt = cnt + 1

srcNet.write(dest)
```

Upload this script and run it from Job Manager using the steps mentioned in Run external scripts on page 387. Use this command:

```
ext_exe_eg.py input-plan.pln out-plan.pln
```

Once the job is completed successfully, download the output file (.tar file) from the Job Manager page, extract it, and import the updated plan file into the user space to access it.

# 26 Manage Add-ons

**Topics:**


- Add-ons
- Create add-ons
- Run an add-on on a plan file
- Import an add-on
- Export an add-on
- Edit an add-on
- Sample add-on


Outlines add-on management processes, including configuration requirements, user permissions, creation and editing procedures, import and export operations, and provides guidance for running add-ons and utilizing example configurations.

# Add-ons

Introduces add-on concepts, covering configuration requirements, limitations, and necessary user role permissions for effective add-on management.

An add-on is a customizable Cisco Crosswork Planning component that

- extends the Cisco Crosswork Planning Design UI functionality, enabling you to meet your specific needs
- allows you to run the scripts like any other tools and view generated output plan files and reports, and
- enhances Cisco Crosswork Planning with additional features and integrations that are not included in the base installation.

Add-ons can invoke Python executable scripts that can modify the network models. These scripts are executed as an asynchronous job and you can view the results in Cisco Crosswork Planning Design.

## Requirements and limitations for add-on configuration

Outlines requirements and limitations for configuring add-ons.

Ensure you adhere to these requirements and limitations while configuring add-ons.

- You can enable a maximum of 15 add-ons at a time for each user.
- Only Python scripts are supported for add-ons.
- You must migrate the existing Cisco WAE add-on scripts to Python3.
- Add-on scripts must be compatible with the provided Python libraries. Installing additional libraries is not supported.
- All add-on scripts run using the design_api_python binary. Executing them within other Python wrapper is not supported.
- You can import add-ons from other systems, such as Cisco WAE, if they meet the requirements described.

## User role permissions for add-ons

Outlines user role permissions for managing add-ons.

On the **Manage Add-ons** page of Cisco Crosswork Planning:

- Any user can create, view, and enable any add-on.
- Only the creator of the add-on or a user with the "CW Planning Design Settings APIs" role permission can edit or delete it.

This table shows which actions users can perform on add-ons.

| Action | Any user | Add-on creator | User with the "CW Planning Design Settings APIs" role permission |
|---|---|---|---|
| Create add-ons | ✓ | ✓ | ✓ |
| View add-ons | ✓ | ✓ | ✓ |
| Enable add-ons | ✓ | ✓ | ✓ |

| Action | Any user | Add-on creator | User with the "CW Planning Design Settings APIs" role permission |
|---|---|---|---|
| Edit add-ons | ✗ | ✓ | ✓ |
| Delete add-ons | ✗ | ✓ | ✓ |

## Create add-ons

Provides step-by-step instructions for creating a customizable add-on in Cisco Crosswork Planning Design including uploading scripts, configuring UI fields with the Form Builder, validating requirements, and finalizing the add-on for use within the system.

This topic describes how to add a new customizable add-on to automate or enhance tasks in Cisco Crosswork Planning Design.

Use this process to create an add-on for use within the system, supplying your own script and customizing the UI.

### Before you begin

Review the requirements in Requirements and limitations for add-on configuration on page 390.

#### Procedure

1. Open the Manage Add-ons page using one of these options:
   - On the Network Design page, click **Manage add-ons**.
   - From the toolbar, select **Actions** > **Add-ons** > **Manage add-ons**.
2. Click + to open the Add-on Settings page.
3. Enter the basic details.
   a) Enter the name of the add-on.
   b) In **File**, upload the compressed file containing your add-on script.
   c) In **Specify executable**, enter the name of the add-on script file within the uploaded zip folder that will be executed when the add-on runs.
   d) Click **Proceed to form builder**.

   The Form Builder page opens.
4. (Optional) To make any changes to the add-on name, executable name, or description, click **Edit** in the **Add-on settings** section.

> **Note**
>
> You cannot upload a different add-on file using this option.

5. On the **Form Builder** page, drag the required component from the left pane to the main area. Configure the steps and input controls for your add-on's user interface. For descriptions of all the fields, refer to Add-on controls on page 392.
   a) Use the **Step** element to create new pages. Click to rename the label of the page.
   b) Use the **Section** element to create new sections in a page. Click to rename the label of your section.
   c) Use the elements under the **Fields** section to configure the required fields. Click to specify name, label, default values, or limitations for your fields.
6. Ensure that each configured field meets the minimum required settings. The fields must display . If you have not met the requirements, an alert ( ) appears. In this situation, click and specify the required fields.
7. Click **Preview** to preview the add-on configuration. Verify all the details you configured.
8. When you are satisfied with the changes, click **Back to editor** at the top to go back to the Form Builder page.
9. Click **Create add-on** to create the add-on.

---

The new add-on appears on the Manage Add-ons page and is ready for use.

## Add-on controls

Outlines the fields available for configuring add-ons.

This topic describes the fields you can use while configuring add-ons.

**Table 40: Add-on controls**

| Field | Description |
|---|---|
| Step (new page) | Creates new pages. |
| Section | Groups fields together that are created within a step (page). |

| Field | Description |
|---|---|
| Field | Type of field.<br>• Text Box: Text field.<br>• Number: Numeric value. Supported format includes both whole and decimal numbers.<br>• Dropdown: Selection of items to appear in a drop-down menu.<br>• Check Box: Toggle entry. If selected, it is true and if not selected, it is false.<br>• Table Selection: Select objects in a specified table. The selected objects are written to a file and then passed to the script.<br>For example, if you select Circuits in the **Select Table** field, you need to select the required circuit objects while running the add-on.<br>• Table Entry: Select a table column in a specified table.<br>• File: Select a file. You can browse with a browse button to a different directory. The complete path is passed to the script.<br>• Password: Text string. Characters entered are hidden. This field is used to pass sensitive text to script. |

## Run an add-on on a plan file

Guides you through enabling, configuring, and running an add-on on a plan file, including job scheduling, output import, and accessing generated reports to visualize and review changes.

This topic describes how to execute a selected add-on on a plan file to generate required outputs or reports.

### Before you begin

- Ensure the required add-on is available and enabled on the Manage Add-ons page.
- You can enable a maximum of 15 add-ons for each user.

#### Procedure

1. Enable the required add-on from the Manage Add-ons page.
   a) Select the add-on on the Manage Add-ons page.
   b) Click **More actions** > **Enable add-on**.

   The enabled add-on appears under the **Actions** > **Add-ons** menu on the visualization toolbar in the Network Design page.
2. Select the enabled add-on from the **Actions** > **Add-ons** menu on the Network Design page.
3. Enter the required input fields as defined by the add-on controls.
4. On the **Run Settings** page, enter a job name and schedule when to run the add-on.

   Optionally, set these options:
   - Priority: Select the priority of the task.
   - Engine profiles: Select the engine profile as needed. This section lists all available asynchronous engine profiles.

- Schedule: Set the time at which you want to run the tool.

> **Note**
>
> Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are ignored when you run the tool as a scheduled job.

5. Submit the job.

   The add-on runs at the scheduled time using the selected engine profile. You can track the job status on the Job Manager page (from the main menu, choose **Job Manager**).

6. Once the job completes, import the output plan file into user space to visualize it. For more information, refer to Access output plan files from job manager on page 384.
7. (Optional) Access the generated report to view the changes caused by the add-on.
   a) Open the newly imported plan file in the Network Design page.
   b) From the toolbar, click **Actions** > **Reports** > **Generated reports**.
   c) Click the name of the add-on to view the summary and details of the changes.

## Import an add-on

Provides step-by-step instructions for importing existing add-ons from other systems into the Cisco Crosswork Planning Design application.

This topic describes how to import an existing add-on to your Cisco Crosswork Planning Design application.

Use this procedure when you have an add-on that was created in a different Cisco Crosswork Planning system or the Cisco WAE Design application.

### Before you begin

- Ensure that the add-on package is a .tar file.
- Ensure that the add-on package contains the addon.json file if the add-on was created with Cisco Crosswork Planning, or the addon.txt file if it was created with Cisco WAE.

### Procedure

1. Open the Manage Add-ons page using one of these options:
   - On the Network Design page, click **Manage add-ons** at the top right.
   - From the toolbar, select **Actions** > **Add-ons** > **Manage add-ons**.
2. Click  to open the Add-on Settings page.
3. Click **Browse** and select the add-on package file.
4. Click **Generate form** to import the add-on.
5. Configure or update any fields or options as needed. For descriptions of all the fields, refer to Add-on controls on page 392.

6. Click **Create add-on**.

The imported add-on is available on the Manage Add-ons page and is ready for use.

## Export an add-on

Provides instructions for exporting add-ons and their configurations from the Cisco Crosswork Planning Design application, enabling you to move add-ons between instances for reuse or migration.

This topic describes how to export add-ons and their configurations from the Cisco Crosswork Planning Design application.

Use this procedure to move add-ons between Cisco Crosswork Planning instances. First, export the add-on, then import it into another instance.

### Procedure

1. Open the Manage Add-ons page using one of these options:
   - On the Network Design page, click **Manage add-ons** at the top right.
   - From the toolbar, select **Actions** > **Add-ons** > **Manage add-ons**.
2. Select the add-on you want to export and click **More actions** > **Export**.
3. Save the package containing the add-on configurations and scripts.

The selected add-on is exported to your local machine and is available for use.

### What to do next

To use this add-on in another Cisco Crosswork Planning instance, log in to that instance and use the steps in Import an add-on on page 394.

## Edit an add-on

Guides you through modifying an existing add-on, including updating the add-on file, adjusting user interface controls and properties, previewing changes, and saving the updated configuration in the management interface.

This topic describes how to modify user interface or input controls of an existing add-on.

### Procedure

1. Open the Manage Add-ons page using one of these options:
   - On the Network Design page, click **Manage add-ons** at the top right.
   - From the toolbar, select **Actions** > **Add-ons** > **Manage add-ons**.
2. To use a different add-on file, select the add-on you want to edit and click **More actions** > **Update add-on file**. Then, upload the new add-on file.
3. To modify any UI controls, select the add-on you want to edit and click **Edit**.
   a) Update UI controls, default values, or other UI-related properties as needed. For descriptions of all the fields, refer to Add-on controls on page 392.

b) Ensure that each configured field meets the minimum required settings. The fields must display . If you have not met the requirements, an alert ( ) appears. In this situation, click and specify the required fields.
c) Click **Preview** to preview the add-on configuration. Verify all the details you configured.
d) When you are satisfied with the changes, click **Back to editor** at the top to go back to the Form Builder page.
e) Click **Save add-on** to save the changes.

---

## Sample add-on

Introduces a sample add-on.

Cisco Crosswork Planning includes a Sample Addon on the Manage Add-ons page. This add-on updates the description of the selected interfaces to "sample description". The Sample Add-on is provided to serve as a practical example, allowing you to explore and better understand how add-ons function within Cisco Crosswork Planning. By experimenting with the Sample Add-on, you can safely familiarize themselves with the process of configuring and managing add-ons.

When you run this add-on, Cisco Crosswork Planning performs these actions.

- An output plan file named "sample-addon-output-planfile.txt" is generated.
- In the newly created plan file, the description of each interface you selected while executing the add-on is updated to "sample description".
- A report named "Sample Addon" is generated. This report includes two sections: Summary and UpdatedInterfaces.
  - The **Summary** section displays the number of interfaces updated with a new description.
  - The **UpdatedInterfaces** section lists the interfaces whose descriptions have been updated.

# 27 Update Configuration from One File to Another

**Topics:**


- Changeover tool
- Run the changeover tool
- Report analysis


Guides you through updating configurations using the Changeover tool and explains how to analyze generated reports to ensure successful migration.

## Changeover tool

Outlines the function of Changeover tool, which guides network transitions by providing sequenced routing configuration changes and minimizing congestion during interim states.

The **Changeover** tool is a Cisco Crosswork Planning tool that

- provides a step-by-step sequence of routing configuration changes to move a network safely from an initial configuration to a pre-specified final configuration, and

  optimizes the order of changes to minimize network congestion during intermediate configurations.

Cisco Crosswork Planning lets you model the topology, routing, and utilization of a currently operational network. It enables exploration of modifications, such as adjusting interface metrics or explicit LSP routings to obtain a new routing configuration.

Supported configuration changes between initial and final plans include:

- Changes to interface metrics
- Changes to LSP configurations
- Changes of circuits and nodes between active and inactive states

An individual step in the changeover sequence consists of:

- A single metric change on a specific interface
- LSP configuration changes on all LSPs sourced from a specific node
- The activation or inactivation of a specific circuit or node

## Run the changeover tool

Explains how to configure the Changeover tool to compare initial and final plan files and execute changeover operations. This process helps you analyze network changes and monitor interface utilization during the changeover sequence.

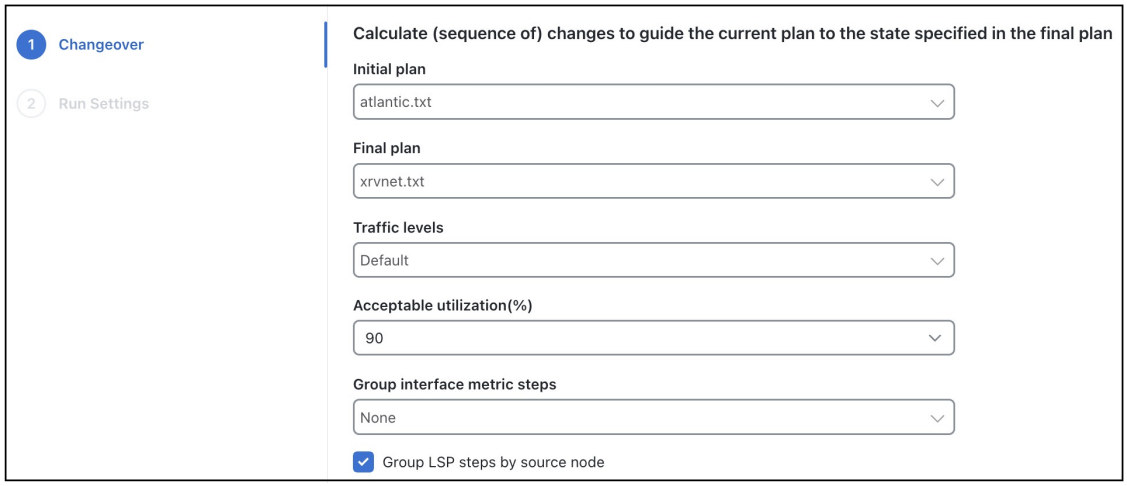
Complete these steps to run the Changeover tool.

Use the Changeover tool to analyze differences between two plan files. This tool is useful when you need to evaluate the impact of network changes and ensure optimal utilization during the changeover sequence.

**Procedure**

1. Open both the initial and final plan files (refer to Open plan files on page 36), and choose the desired plan to view.
2. From the toolbar, choose **Actions** > **Tools** > **Changeover**.

   **Figure 119: Changeover options**




3. Decide on the changeover options to use. For field descriptions, refer to Table 41: Changeover options on page 399
4. Click **Next**.
5. On the Run Settings page, under **Execute**, choose either of these options:
   - **Now**: Choose this option to run the tool immediately. Upon completion, a report of the results opens automatically.
   - **As a scheduled job**: Choose this option to run the tool at a scheduled time. Set these options:
     - **Priority**: Select the priority of the task.
     - **Engine profiles**: Select the engine profile as per your requirement. This section lists all the available asynchronous engine profiles.
     - **Schedule**: Set the time at which you want to run the tool. This option is available only for scheduled jobs.

   The job runs at the scheduled time and using the selected engine profile. You can track the status of the job from the **Job Manager** page (from the main menu, choose **Job Manager**).

   **Note**

   Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

**Table 41: Changeover options**

| Field | Description |
|---|---|
| Initial plan | Name of the initial plan, selected from opened plan files. |
| Final plan | Name of the final plan, selected from opened plan files. |
| Traffic levels | Changeover monitors maximum interface utilization levels for all steps in the sequence. The utilizations are calculated using this traffic level. |

| Field | Description |
| --- | --- |
| Acceptable utilization (%) | The percentage of maximum acceptable utilization of any interface at any step during the changeover sequence. Changeover tries to keep utilizations below this level, although it is not always possible. Often the utilization spikes upward during the last few steps of the sequence, if, for example, new ECMP paths must be put in place. Cisco Crosswork Planning tries to keep the number of high-utilization steps to a minimum. |
| Group interface metric steps | The Changeover tool allows you to group Interface Metric changes together into steps. You can choose from these options:<br>• None: Each Interface Metric change is treated as a single step.<br>• Parallel: Parallel Interfaces are grouped into a single step.<br>• Source Node: Interfaces sourced from the same node are grouped into a single step if there is no impact on maximum utilization. |
| Group LSP steps by source node | Specify whether to group the LSP steps by source node. |

## Report analysis

Provides information about the sections included in the generated report and how to access it from the user interface.

To access the report created by the Changeover tool, choose **Actions** > **Reports** > **Generated reports** link from the right panel.

The report contains several sections that help you review the options used, actions taken, and interface utilizations during the changeover process.

- **Summary**: Contains a list of the options used to run the changeover. It also contains a summary of the differences between the initial and final plans, the number of steps taken, and the number of steps resulting in intermediate configurations with utilizations in excess of the acceptable utilization.
- **Steps**: Contains details of the actions taken at each step while running the Changeover tool.
- **Utilizations**: Contains all utilizations for all interfaces in the network, for each step in the changeover sequence.

# 28 Create and Use Patch Files

**Topics:**


- Patch files
- Create patch files
- Apply patch files
- View or edit patch files


Explains the concepts, creation, application, and management of patch files, guiding users through understanding patch files, generating new patches, applying patches, and viewing or editing existing patch files.

# Patch files

Explains what patch files are, how they differ from plan files, and how they are used in Cisco Crosswork Planning.

A patch file is a file that

- represents differences (or patches) between plan files
- enables these differences to be applied to other plan files, and
- offers a compact, efficient way to manage updates or changes across plans.

A common Cisco Crosswork Planning use case is to create a plan file for deployment. A typical workflow for doing this includes these steps:

1. Open a plan file from the Cisco Crosswork Planning UI (refer to Open plan files on page 36).
2. Duplicate the plan file using ••• > **Make a copy**.
3. Make LSP modifications as needed in the duplicated plan file.
4. Open both plan files and create a patch file using **Actions** > **Tools** > **Patches** > **Create patch**.
5. Deploy the patch file to the network using **Actions** > **Tools** > **Patches** > **Apply patch**.

# Create patch files

Explains how to configure patch files to identify differences between two plan files and apply targeted changes.

Create patch files to automate the process of updating one plan file to match another. This process efficiently identifies additions, deletions, or modifications between the files.

The patch file created identifies how to get to the plan file identified in the **To file** field from the plan file identified in the **From file** field. These patch files only include differences in new, modified, and deleted objects that are supported by the Cisco Crosswork Planning YANG model.

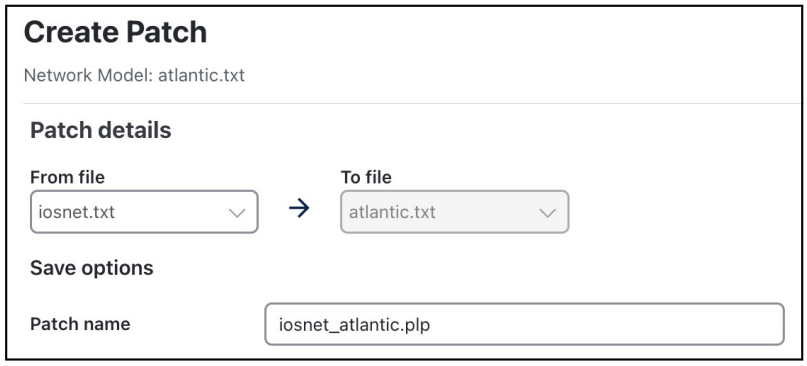
Examples:

- If From file iosnet.txt contains 33 LSPs and To file atlantic.txt contains 23 LSPs, the patch file identifies which 10 LSPs to delete from plan iosnet.txt to create plan atlantic.txt.
- If From file iosnet.txt contains 10 LSPs and To file atlantic.txt contains 15 LSPs, the patch file identifies which 5 LSPs to add to plan iosnet.txt to create plan atlantic.txt.
- If both plan files contain the same LSP, but in From file iosnet.txt the LSP has a SetupBW property equal to 100 and in the To file atlantic.txt it does not, the patch file identifies the need to modify the SetupBW property of this LSP in plan iosnet.txt so that it matches plan atlantic.txt.

Complete these steps to create patch files that identify differences between two plan files.

### Procedure

1. Open the two plan files that you are using for the patch file creation. For more information, refer to Open plan files on page 36.
2. From the plan file you want to use as the target (the "To file"), choose **Actions** > **Tools** > **Patches** > **Create patch** from the toolbar.




3. From the **From file** drop-down list, choose the plan file from which you are determining required changes to achieve the plan file identified in the "To file" field.
4. Review the default patch filename, which combines the From and To plan filenames. Change the filename if required. Ensure it ends with a .plp extension. To prevent overwriting a patch file, ensure the filename is unique within its location.
5. Click **Create**.

The patch file is created successfully. The generated patch file is saved under **Network Models** > **User space** > **Other files**. Use the ••• button under the **Actions** column to open, download, delete, or view details of the patch file.

### What to do next

Use the **Edit patch** button to edit the patch file. Use the **Download** button to download the file to your local machine.

## Apply patch files

Explains how to apply patch files to plan files to update or modify their contents as needed.

Use this task to make targeted changes or corrections to existing plan files using patch files.

You may need to update plan files with patch files to fix errors, add enhancements, or synchronize the design with new requirements.

### Procedure

1. Open the plan file in the **Network Design** page. For more information, refer to Open plan files on page 36.
2. From the toolbar, choose **Actions** > **Tools** > **Patches** > **Apply patch**.
3. Click **Browse** to browse to the location where the patch file resides. Choose from these options:
   - **From user space**: Choose this option to select the patch file saved in your user space.
   - **From local**: Choose this option to select the patch files saved in your local system. Click **Browse**, choose the patch file located in your local system, and click **Import**.

     Note that the patch file is imported into the user space as well.
4. Specify whether to stop or continue in case of error.
5. To verify the patch file's contents before applying it, click the **View details** link. After viewing the patch, click **Close**. Use the **Edit patch** button to edit the patch file.
6. Click **Run** in the Apply Patch page.

The patch file is applied to the selected plan file, updating its contents as specified.

## View or edit patch files

Describes how to configure and review patch files before applying or deploying them to ensure they meet your requirements.

Viewing the patch files lets you ensure they meet your needs before applying or deploying them.

If the patch files are available under **Network Models** > **User space** > **Other files**, then use ••• > **Open** to view the patch file details.

### Before you begin

Confirm patch files are available in either the user space or on your local system.

#### Procedure

1. Open the plan file in the **Network Design** page. For more information, refer to Open plan files on page 36.
2. From the toolbar, choose **Actions** > **Tools** > **Patches** > **View patch**.
3. Browse to the location where the patch file resides. Choose from these options:
   - **From user space**: Lists the patch files saved in the User Space. Select the required patch file.
   - **From local**: Click **Browse** and choose the patch file located in your local system. Then, click **Import**.
4. Click **View details** to examine the patch file contents.
5. After reviewing the patch, take one of these actions:
   - **Edit patch**: Use this option to edit the patch file. The file becomes editable, displaying the contents of the patch as XML text. You can save or discard the patch text edits. The tool warns you if you try to save invalid XML syntax. Other standard text editing (cut, copy, paste, and so on) is also supported.
   - **Download**: Use this option to download the patch file to your local system.
   - **Close**: Use this option to close the window without deploying the patch.

# 29 Export Routes

**Topics:**


- Export routes in Cisco Crosswork Planning
- Route tables
- Output table column details
- Export routes for selected objects


Outlines export route capabilities, covering the overview of export routes, route table structures, table output details, and step-by-step procedures for exporting routes using the UI for efficient network planning and management.

## Export routes in Cisco Crosswork Planning

Outlines how to export routes to a text-editable format for route analysis.

The ability to export routes facilitates the exchange of information in a text editable format. It also supplements the analysis of routes and plan file objects. For example, you could export demand routes using link to determine which links are used by demands. Since you can export routes that changed due to failure states, you can fine-tune your analysis of how failures impact the network.

## Route tables

Explains the concept of route tables and their role in exporting network routing information.

A route table is a table created when exporting routes using Cisco Crosswork Planning UI or CLI job in Job Manager that

- contains routing information for specific objects such as demands or LSPs
- output to a .txt file containing only the table data, and
- is accessible from **My user space** > **All files**.

The types of route tables include:

- <DemandHops>
- <LSPHops>
- <LSPPathHops>
- <ShortestIGPPathHops>
- <ShortestLatencyPathHops>
- <ShortestTEPathHops>
- <CircuitStyleLSPPathHops>

Optionally, you can fail objects before exporting their routes and then export only the routes that changed due to failure.

## Output table column details

Lists and describes the key columns included in exported route tables, including how each column uniquely identifies and organizes route hop information.

The key columns in the exported tables are a combination of the key column identified for its object type plus a Step column. The Step column is an integer that identifies the sequence of this hop in the route path. For example, the <Demands> table includes the key columns Name, Source, Destination, and ServiceClass. Therefore, the <DemandHops> table has four columns to uniquely identify the demand, and a Step column to identify the hop sequence.

Each exported hops table includes these columns:

- **UnresolvedHop**: Identifies the hop type and complete definition. For example, if{cr1.atl|to_cr1.hst} means this hop is an interface with a source node of cr1.atl and an egress interface of to_cr1.hst.
- **PreviousStep**: Identifies the previous step or steps in the route. As such, this identifies the previous hop in the route. For example, if you are on the row corresponding to Step 4, and the PreviousStep is 2, then the previous hop is contained in the row corresponding to Step 2. If the PreviousStep column is empty, then this is the first hop in the route.

- **NextStep**: Identifies the next step in that table. As such, this identifies the next hop in the route. If the NextStep column is empty, then this is the last hop in the route.
- **TrafficProportion**: Identifies the proportion of traffic on this hop compared to the other hops in the route. Although this value is usually 1, in case of ECMP routing, it could be less than 1.

> **Note**
>
> Demand routes may be split due to ECMP routing or multicast. Therefore, demand hops might not be listed sequentially. A hop could have two next hops because the route and traffic split into two.

### Export route table example

Outlines an example export of route tables using the LSPHops table.

This table is an example LSPHops table that was created using active LSPs with interface hops. The cr2.par_cr1.fra LSP has no interface hops in its route. In the table, it has only one step and no previous or next steps. The cr1.ams_cr1.par LSP has three interface hops (steps 1, 2, and 3) between source (cr1.ams) and destination (cr1.par).

**Table 42: Example LSPHops table**

| Name | Source | Destination | Step | UnresolvedHop | Previous Step | Next Step | Traffic Proportion |
|---|---|---|---|---|---|---|---|
| cr2.par_cr1.fra | cr2.par | cr1.fra | 1 | if{cr2.par\|{to_cr1.fra}} | | | 1 |
| cr1.ams_cr1.par | cr1.ams | cr1.par | 1 | if{cr1.ams \| to_cr2.lon}} | | 2 | 1 |
| cr1.ams_cr1.par | cr1.ams | cr1.par | 2 | if{cr2.lon\|{to_cr1.lon}} | 1 | 3 | 1 |
| cr1.ams_cr1.par | cr1.ams | cr1.par | 3 | if{cr1.lon\|{to_cr1.par}} | 2 | | 1 |

## Export routes for selected objects

Describes how to configure options and export routes for network objects using the user interface.

Use this task to export route information for selected objects from your plan file to a hops table.

Hops tables are created for the object selected in the **Export Routes** page. For example, if you select Demands, Cisco Crosswork Planning creates a <DemandHops> table containing all demands in the plan file.

If you want to export routes only for objects whose route changed due to failure, first fail all relevant objects. Then export the routes and select the **Only export routes changed by failure state** option. Alternatively, you can run the Simulation analysis tool and then fail one or more objects to the worst-case failure. For more information, refer to Evaluate Impact of Worst-Case Failures on page 129.

You can also use the **Export routes** CLI tool to export routes to the hops table. For more information on using the CLI tools, refer to Run tools or initializers using CLI on page 385.

**Procedure**

1. Open the plan file (refer to Open plan files on page 36). It opens in the **Network Design** page.

2. From the toolbar, choose **Actions** > **File** > **Export** > **Routes**.

   The Export Routes page opens.

3. From the **Object** drop-down list, select the object for which you want to export routes.
4. From the **Hop type** drop-down list, select the hop type that you want to export.
5. If exporting only routes that changed due to a failure, select the **Only export routes changed by failure state** option.
6. In the **File name** field, enter the file name for the output file. By default, the *planfile.object* format is used as the file name.
7. Click **Submit**.

   The exported file is saved under **Network Design** > **My user space** > **All files**.

8. Navigate to **Network Design** > **My user space** > **All files** and download the exported file to your local machine.
   - Click the name of the exported file, or
   - Use the ··· > **Download** option under the **Actions** column.

# A Cross Table Filters

**Topics:**


- Supported cross table filters


Outlines the cross table filters supported in Cisco Crosswork Planning.

## Supported cross table filters

Lists all available cross table filter options in Cisco Crosswork Planning and the software version in which each filter option was introduced.

This section lists all the cross table filters supported in Cisco Crosswork Planning. To access the cross table filters, select the objects from their corresponding Network Summary tables and click the **Filter options** icon ( ).

**Table 43: Cross table filters in Cisco Crosswork Planning**

| Network summary table | Cross table filter option | Version introduced |
|---|---|---|
| Interfaces | Filter to circuits | 7.0.0 |
| | Filter to nodes > Local | 7.0.0 |
| | Filter to nodes > Remote | 7.0.0 |
| | Filter to nodes > Both | 7.0.0 |
| | Filter to ports | 7.0.0 |
| | Filter to demands > Through all interfaces | 7.0.0 |
| | Filter to LSPs | 7.0.0 |
| | Filter to interface queues | 7.1.0 |
| | Filter to CS-RSVP LSPs | 7.1.0 |
| | Filter to VPNs | 7.2.0 |

| Network summary table | Cross table filter option | Version introduced |
|---|---|---|
| Nodes | Filter to demands > Between selected nodes | 7.0.0 |
| | Filter to demands > Through all selected nodes | 7.0.0 |
| | Filter to demands > Selected nodes as source | 7.2.0 |
| | Filter to demands > Selected nodes as destination | 7.2.0 |
| | Filter to demands > Through any selected node | 7.2.0 |
| | Filter to LSPs > Between selected nodes | 7.0.0 |
| | Filter to LSPs > Through selected nodes | 7.0.0 |
| | Filter to LSPs > Selected nodes as source | 7.2.0 |
| | Filter to LSPs > Selected nodes as destination | 7.2.0 |
| | Filter to flex algorithms | 7.1.0 |
| | Filter to node SIDs | 7.1.0 |
| | Filter to SRv6 node SIDs | 7.1.0 |
| | Filter to CS-RSVP LSPs | 7.1.0 |
| | Filter to sites | 7.2.0 |
| | Filter to interfaces > From | 7.2.0 |
| | Filter to interfaces > To | 7.2.0 |
| | Filter to interfaces > From or to | 7.2.0 |
| | Filter to interfaces > From and to | 7.2.0 |
| | Filter to ports | 7.2.0 |
| Demands | Filter to nodes > Source | 7.0.0 |
| | Filter to nodes > Destination | 7.0.0 |
| | Filter to nodes > Both | 7.0.0 |
| | Filter to LSPs | 7.0.0 |
| | Filter to interfaces | 7.0.0 |
| | Filter to CS-RSVP LSPs | 7.1.0 |
| | Filter to external endpoints > Source | 7.2.0 |
| | Filter to external endpoints > Destination | 7.2.0 |
| | Filter to external endpoints > Both | 7.2.0 |
| | Filter to VPNs | 7.2.0 |

| Network summary table | Cross table filter option | Version introduced |
|---|---|---|
| Circuits | Filter to interfaces | 7.0.0 |
| | Filter to nodes | 7.2.0 |
| | Filter to SRLGs | 7.2.0 |
| LSPs | Filter to demands | 7.0.0 |
| | Filter to nodes > Source | 7.0.0 |
| | Filter to nodes > Destination | 7.0.0 |
| | Filter to nodes > Both | 7.0.0 |
| | Filter to LSP paths | 7.0.0 |
| | Filter to interfaces | 7.0.0 |
| | Filter to actual path hops | 7.0.0 |
| | Filter to CS-RSVP LSPs | 7.1.0 |
| LSP Paths | Filter to LSPs | 7.0.0 |
| | Filter to named paths | 7.0.0 |
| | Filter to segment lists | 7.0.0 |
| | Filter to interfaces | 7.0.0 |
| | Filter to CS-RSVP LSPs | 7.1.0 |
| | Filter to flex algorithms | 7.2.0 |
| Ports | Filter to ports circuits | 7.0.0 |
| | Filter to interfaces | 7.0.0 |
| Port Circuits | Filter to circuits | 7.1.0 |
| | Filter to ports | 7.1.0 |
| SRLGs | Filter to nodes | 7.0.0 |
| | Filter to interfaces | 7.0.0 |
| | Filter to ports | 7.2.0 |
| | Filter to port circuits | 7.2.0 |
| | Filter to circuits | 7.2.0 |
| | Filter to configured nodes | 7.2.0 |
| External Endpoints | Filter to demands > Source | 7.1.0 |
| | Filter to demands > Destination | 7.1.0 |
| | Filter to external endpoint members | 7.1.0 |

| Network summary table | Cross table filter option | Version introduced |
|---|---|---|
| External Endpoint Members | Filter to nodes | 7.0.0 |
| | Filter to interfaces | 7.0.0 |
| | Filter to AS | 7.1.0 |
| VPNs | Filter to VPN nodes | 7.1.0 |
| | Filter to interfaces | 7.0.0 |
| | Filter to demands | 7.0.0 |
| VPN Nodes | Filter to VPNs | 7.1.0 |
| | Filter to nodes | 7.1.0 |
| | Filter to demands > From | 7.2.0 |
| | Filter to demands > To | 7.2.0 |
| | Filter to demands > Between | 7.2.0 |
| Sites | Filter to contained nodes > Direct | 7.1.0 |
| | Filter to contained nodes > All | 7.1.0 |
| | Filter to interfaces > In | 7.1.0 |
| | Filter to interfaces > From | 7.1.0 |
| | Filter to interfaces > To | 7.1.0 |
| | Filter to interfaces > From or To | 7.1.0 |
| | Filter to interfaces > From and To | 7.1.0 |
| Interface Queues | Filter to interfaces | 7.1.0 |
| Segment Lists | Filter to LSP paths | 7.0.0 |
| | Filter to segment list hops | 7.1.0 |
| Segment List Hops | Filter to node SIDs | 7.1.0 |
| | Filter to interface SIDs | 7.1.0 |
| | Filter to SRv6 node SIDs | 7.1.0 |
| | Filter to SRv6 interface SIDs | 7.1.0 |
| | Filter to segment lists | 7.1.0 |
| Named Paths | Filter to named path hops | 7.1.0 |
| Named Path Hops | Filter to interfaces | 7.1.0 |
| | Filter to nodes | 7.1.0 |

| Network summary table | Cross table filter option | Version introduced |
|---|---|---|
| Flex Algorithms | Filter to nodes | 7.1.0 |
| | Filter to interfaces | 7.1.0 |
| | Filter to node SIDs | 7.1.0 |
| | Filter to SRv6 node SIDs | 7.1.0 |
| | Filter to LSP paths | 7.2.0 |
| Node SIDs | Filter to nodes | 7.1.0 |
| | Filter to segment list hops | 7.1.0 |
| | Filter to flex algorithms | 7.1.0 |
| | Filter to node prefix loopbacks | 7.1.0 |
| SRv6 Node SIDs | Filter to nodes | 7.1.0 |
| | Filter to segment list hops | 7.1.0 |
| | Filter to flex algorithms | 7.1.0 |
| | Filter to node prefix loopbacks | 7.1.0 |
| Interface SIDs | Filter to interfaces | 7.1.0 |
| | Filter to segment list hops | 7.1.0 |
| SRv6 Interface SIDs | Filter to interfaces | 7.1.0 |
| | Filter to segment list hops | 7.1.0 |
| | Filter to flex algorithms | 7.1.0 |
| CS-RSVP LSPs | Filter to LSPs | 7.1.0 |
| | Filter to LSP paths | 7.1.0 |
| | Filter to interfaces | 7.1.0 |
| | Filter to nodes | 7.1.0 |
| | Filter to demands | 7.1.0 |
| AS | Filter to nodes | 7.1.0 |

# B Legend for Plan Objects

**Topics:**


- Interfaces and Circuits
- Nodes
- Sites
- Demands
- LSPs


Outlines graphical representations and legends for plan objects, including interfaces, circuits, nodes, sites, demands, and LSPs, to help users interpret plan diagrams and symbols within the network management interface.

## Interfaces and Circuits

Lists and describes the visual cues used to represent interfaces and circuits in network plots.

This section explains how the interfaces and circuits are represented in the network plot.

**Table 44: Interface and Circuit visualization**

| Graphic | Description |
|---|---|
|  | A solid line indicates a link between two devices. The color of the line varies depending on the outbound traffic utilization. To learn how link colors correspond to traffic utilization, refer to Traffic utilization colors on page 55.<br>When an interface or circuit is selected, the width of the link doubles ().<br>When there are more than one link between two devices, a dotted line appears (). |
|  | A red circle with a white X on a circuit indicates the circuit has failed. |
|  | A red circle with a white down arrow on a circuit indicates that the circuit is not operational because of one of these reasons:<br>• A node connected to the circuit has failed.<br>• A site that contains a node connected to the circuit has failed.<br>• The circuit is a LAG (port channel) that does not meet its operational requirements (minimum number of available ports or minimum capacity, whichever occurs first) due to failure of some of its ports or port circuits.<br>• An SRLG that contains the circuit or one of the required objects for circuit operation has failed. |
|  | A gray circle with a white X on a circuit indicates the circuit is inactive. |
|  | A gray circle with a white down arrow on a circuit indicates that the circuit is not operational because of one of these reasons:<br>• A node connected to the circuit is inactive.<br>• A site containing a node connected to the circuit is inactive.<br>• The circuit is a LAG that does not meet its operational requirements (minimum number of available ports or minimum capacity, whichever occurs first) due to some of its ports or port circuits being inactive.<br>• An SRLG is inactive, containing the circuit or one of the above objects that are required for circuit operation. |

## Nodes

Lists and describes the visual cues used to represent nodes in network plots.

This section explains how the nodes are represented in the network plot.

**Table 45: Node visualization**

| Graphic | Description |
|---|---|
|  | A blue icon indicates a node. A light blue outline indicates that the node is selected. |
|  | A green border surrounds the node when:<br>• Simulated traffic view: The node is the source or destination of simulated traffic that is greater than zero.<br>• Measured traffic view: The node is the source or destination of measured traffic that is greater than zero.<br>• LSP reservation view: The node is the source or destination of LSPs with non-zero bandwidth reservations. |
|  | A light gray border surrounds the node when:<br>• Simulated traffic view: The node is not the source or destination of simulated traffic.<br>• Measured traffic view: The node has ingress and egress traffic of zero.<br>• LSP reservation view: The node is the source or destination of LSPs with zero bandwidth reservations. |
|  | A dark gray border surrounds the node when:<br>• Measured traffic view: The node has no ingress or egress traffic.<br>• Worst-case traffic view: Nodes always have gray borders.<br>• Failure impact view: The node is not included in the failures considered, or failing the node would not increase the traffic utilization on any circuit. |
|  | In the Failure impact view, the node border color represents the failure impact of a node. This is defined as the highest utilization experienced by any circuit if that node were to fail. |
|  | A red circle with a white X on the node indicates the node has failed. |
|  | A red circle with a white down arrow on the node indicates the node is not operational because it is contained in a failed site or in a failed SRLG. |
|  | A gray colored node indicates the node is not active. |
|  | A gray circle with a white down arrow on the node indicates the node is not operational because it is contained in an inactive site or an SRLG. |
|  | A green icon indicates that one or more selected VPN nodes exist on the node. |

## Sites

Lists and describes the visual cues used to represent sites in network plots.

This section explains how the sites are represented in the network plot.

Site utilization colors and symbols are based on intrasite objects. Intrasite objects include child sites and the objects contained within them.

**Table 46: Site visualization**

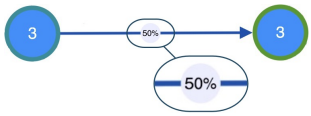
| Graphic | Description |
|---|---|
|  | A blue circle indicates a site. The number inside the circle indicates the number of nodes (devices) in the site.<br>A light blue border indicates that the site is selected. |
|  | • Measured traffic and Simulated traffic views: The border color is the maximum traffic utilization color of all intrasite interfaces within the site.<br>• Worst-case traffic view: The border color is the highest utilization that the interfaces within it experience over all the failure sets selected for Simulation Analysis.<br>• Failure impact view: The border color is the maximum utilization level that would be caused elsewhere in the network should anode or circuit contained within it fail. |
|  | A gray border surrounds the site if all intrasite interfaces have zero utilization. |
|  | A red circle with a white X mark on the site indicates that the site has failed. |
|  | A red circle with a white down arrow on the site indicates:<br>• One or more intrasite objects have failed.<br>• One or more nodes or circuits contained within the site are not operational because they are contained in a failed SRLG. |
|  | A gray colored site indicates:<br>• The site is inactive.<br>• One or more intrasite objects are inactive.<br>• One or more nodes or circuits contained within the site are not operational because they are contained in an inactive SRLG. |
|  | A green icon as shown in the image indicates one or more selected VPN nodes exist on nodes within the site. |

## Demands

Lists and describes the visual cues used to represent demands in network plots.

This section explains how the demands are represented in the network plot.

**Table 47: Demand visualization**

| Graphic | Description |
|---|---|
|  | A dark blue arrow indicates a demand routed through an interface or a circuit. “A” and “Z” mark the source and destination site/node of a demand. |
|  | A dotted dark blue arrow indicates the rerouted demand under a failure. |
|  | Percentages show the proportion of a split demand (for example from ECMP) through an interface. |
|  | A light blue arrow indicates a demand routed over an LSP.<br>A dotted light blue arrow () indicates the rerouted demand over an LSP. |
|  | A dark green arrow indicates the shortest IGP path of a demand.<br>A dotted dark green arrow () shows the rerouted shortest IGP path of a demand during a failure. |
|  | A light green arrow indicates the shortest latency path of a demand.<br>A dotted light green arrow () shows the rerouted shortest latency path of a demand during a failure. |
|  | A dark brown arrow indicates the aggregated path of a demand. This appears when different types of demands, such as the shortest IGP path or shortest latency path, share the same demand route.<br>A dotted brown arrow () shows the rerouted aggregated path of a demand during a failure. |

## LSPs

Lists and describes the visual cues used to represent LSPs in network plots.

This section explains how the LSPs are represented in the network plot.

**Table 48: LSP visualization**

| Graphic | Description |
|---|---|
|  | A purple arrow indicates how LSP traffic is routed. “A” and “Z” mark the source and destination site/node of an LSP. |

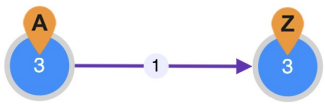
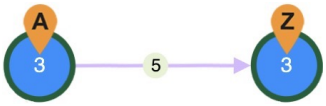
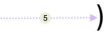
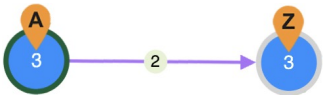
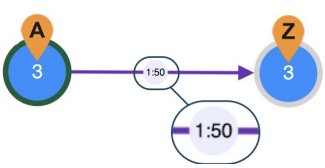
| Graphic | Description |
|---|---|
|  | A dotted purple arrow indicates the rerouting of LSP traffic when a failure occurs in Convergence mode. |
|  | A light green arrow indicates the shortest latency path of an LSP.<br>A dotted light green arrow () indicates the rerouted shortest latency path of an LSP during a failure. |
|  | A gray arrow indicates the shortest TE path of an LSP.<br>A dotted light gray arrow () indicates the rerouted shortest TE path of an LSP during a failure. |
|  | A brown arrow indicates the aggregated path of an LSP. This appears when different types of LSPs, such as the shortest IGP path, shortest TE path, or shortest latency path, share the same LSP route.<br>A dotted brown arrow () indicates the rerouted aggregated path of an LSP during a failure. |

## LSP paths

Lists and describes the visual cues used to represent LSP paths in network plots.

This section explains how the LSP paths are represented in the network plot.

**Table 49: LSP path visualization**

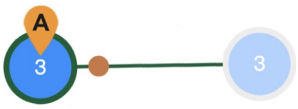
| Graphic | Description |
|---|---|
|  | A purple arrow indicates the LSP primary path that is active (carrying traffic). “A” and “Z” mark the source and destination site/node of an LSP path. This is the path with lowest path option (for RSVP LSPs) or the highest preference (for SR LSPs). The number in the circle (here 1) is the path option.<br>A dotted purple arrow () indicates the rerouted primary path of an LSP during a failure. |
|  | A very dim purple arrow indicates the standby LSP path which is not a primary path and is defined as a Standby path. The number in the circle (here 5) is the path option.<br>A dotted arrow () shows the rerouted standby path of an LSP during a failure. |
|  | A light purple arrow indicates the alternative LSP path which is not a primary path of an LSP. The number in the circle (here 2) is the path option.<br>A dotted light purple arrow () shows the rerouted alternate path of an LSP during a failure. |
|  | The value in the circle (for example, 1:50) helps to identify the loadshare percentage that this LSP path is carrying for the selected LSP. In this example, 1 indicates the LSP path option and 50 indicates that the LSP loadshare percentage is 50%.<br>The maximum number of characters that can be displayed in this circle is 4. Therefore, it always shows the truncated value if the number of characters exceeds 4. For example, 1:50 is shown for 1:50% and 12:5 is shown for 12:50%. You can click the badge to see the complete value. |

## LSP named paths

Lists and describes the visual cues used to represent LSP named paths in network plots.

This section explains how the LSP named paths are represented in the network plot.

**Table 50: LSP named path visualization**

| Graphic | Description |
|---|---|
|  | A brown circle on an interface denotes it as a named path hop. “A” marks the source of a named path. |
|  | A brown outline on a site indicates that the site contains a node or interface that is a named path hop. |
|  | A brown outline on a node denotes it as a named path hop. |

## LSP actual paths

Lists and describes the visual cues used to represent LSP actual paths in network plots.

This section explains how the LSP actual paths are represented in the network plot.

**Table 51: LSP actual path visualization**

| Graphic | Description |
|---|---|
|  | A dark yellow diamond on an interface denotes an LSP actual path hop on the interface. |
|  | A dark yellow outline on a site denotes an actual path hop on a node or interface within the site. |
|  | A dark yellow circle on a node denotes an actual path hop on the node. |

## SR LSP paths

Lists and describes the visual cues used to represent SR LSP paths in network plots.

This section explains how the SR LSP paths are represented in the network plot.

**Table 52: SR LSP path visualization**

| Graphic | Description |
| --- | --- |
|  | A brown circle on an interface denotes an interface segment in an SR LSP. |
|   | A brown circle around a site indicates a site containing a node segment or an interface segment. |
|  | A brown circle around a node denotes the node segment in an SR LSP. |