

Cisco Crosswork Planning Design 7.1 User Guide

First Published: 2025-06-25

Last Modified: 2025-10-17

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

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)

© 2025 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1 **Overview** 1

Core capabilities of Cisco Crosswork Planning 1

Cisco Crosswork Planning Design use cases 2

System overview 2

CHAPTER 2 **Get Started** 5

Key concepts 5

Log in to and log out of Cisco Crosswork Planning 7

Dashboard 7

Custom dashboards 8

Create custom dashboards 8

Rename custom dashboards 9

Delete custom dashboards 9

Customize dashlets 10

Reset the General dashboard to the default view 11

Copy dashboards 11

Design engine instances 11

Resource allocation for design engine instances 11

Create design engine services 12

Reset engine sessions 13

Delete design engine instances 14

Recreate design engine services 14

PART I **Visualize Network Models** 17

CHAPTER 3 **Visualize Network Models** 19

Create or import network models 19
Import plan files from the local machine 19
Import plan files from archive 20
Create new network models manually 21
Open plan files 22
Delete plan files 23
Important considerations for plan file deletion 23
Download plan files 24
Key features of the Network Models page 24
Key features of the Network Design page 27
Select design engine profiles 30
Assign sites to nodes 30
Site assignment rules 31
Simple mapping rules 31
Node-to-Site mapping rules 32
Run the Assign sites to nodes initializer 33
Assign geographic locations to sites 34
Initialize site location 35
Network topology 36
Network plots 36
Site plots 38
Graphical network topology views 38
Grouping of nodes 38
Representation of plan objects 39
Traffic utilization 39
Example: Determine traffic utilization 40
Identify most highly utilized interface 42
Network summary tables 42
Common tables 42
Work with tables and object selections 43
Show or hide tables or columns 44
Search and filter in tables 44
Sort columns in tables 45
Save table views 45

CHAPTER 4 **Understand Plan Objects** 47

Nodes and Sites 47

Parent sites and contained objects 48

Delete sites 48

PSN nodes 48

Create nodes and sites 49

Create nodes 49

Create sites 49

Merge nodes 50

Circuits and interfaces 51

Create circuits and interfaces 51

Merge circuits 51

SRLGs 52

Create SRLGs 52

Create SRLGs for circuits only 52

Ports, port circuits, and LAGs 53

Create ports 54

Create port circuits 54

Create LAGs 55

Set LAG simulation properties 55

Key operations you can perform on objects 55

Create objects 56

Edit objects 56

Delete objects 57

PART II **Simulate Your Network** 59

CHAPTER 5 **Simulation Overview** 61

Use cases of network simulation 61

Auto-resimulation 62

States of plan objects 63

Failed state 63

Fail and recover objects 65

Protect circuits from SRLG failures 65
Active state 66
Set objects to Active or Inactive state 67
Operational state 67
Simulated capacity 68
Simulated Delay 69

CHAPTER 6 **Simulate Traffic Flow from Source to Destination Using Demands 71**
Demands 71
Demand sources and destinations 73
Demand meshes 74
Demand latency bounds 74
How demands can be used? 74
Create demands and demand meshes 75
Create demands 75
Create demand meshes 76
Create demands for LSPs 77
Set demand latency bounds 78
Visualize demands 78
View participating only nodes 79
Traffic rerouting 80
View shortest IGP and latency paths, and overlapped paths 81
Demand traffic 83
Simulate traffic using demands 83
Modify demand traffic 84
Modify fixed demand traffic 85
Modify fixed or relative demand traffic 85
Understand demand deduction 87
Differences in measured and simulated traffic 88
Minimize differences between measured and simulated traffic 89
Flow measurements in demand deduction 89
Estimate demand traffic using demand deduction 90
Demand deduction example 92

CHAPTER 7 **Perform What-If Analysis** 95

Examine failure scenarios 95

Perform failure analysis 95

Example: Failure analysis 96

Perform impact analysis of topology changes 98

Example: Impact analysis of topology changes 99

Perform impact analysis of metric changes 100

Example: Impact analysis of metric changes 100

CHAPTER 8 **Evaluate Impact of Worst-Case Failures** 103

Worst-case traffic utilization 104

Worst-case QoS violations 105

Fail circuits to worst-case utilization 106

Worst-case demand latency 107

Fail demands to worst-case latency 108

Run simulation analysis 109

Pairwise failure analysis 111

Pairwise failure analysis example 112

Protect objects 115

Analyze simulation analysis reports 116

Failure impact view 118

Parallelization 119

CHAPTER 9 **Evaluate Impact of Traffic Growth** 121

Demand Groupings 121

Growth plans 122

Create growth plans 122

Growth plans from demand grouping 124

Create growth plans from demand growth 124

Create growth plans from interface growth 125

CHAPTER 10 **Perform Capacity Planning** 127

Optimize capacity 127

Example 129
Optimization options 131
Advanced optimization options 132
Parallel circuits design example 133
Analyze optimization reports 134

CHAPTER 11 **Simulate IGP Routing Protocol** 135
IGP process ID 135
Configure IGP process ID 135
Associate interfaces with an IGP process ID 136
OSPF and IS-IS 136
OSPF area simulation 137
Set OSPF area membership 137
IS-IS multi-level simulation 137
Set IS-IS multi-level simulation 138
EIGRP 139
Simulate IGP multipath 139
Exclude ABR nodes 140
Configure IGP simulation 140
Simulate multi-topology routing 141
Configure topologies 141
Add topologies to objects 142

CHAPTER 12 **Simulate BGP Routing** 143
Internal and external AS types 143
Configure ASes 144
Create ASes 144
Associate nodes with an AS 144
Edit AS routing policies 145
Route demands between ASes 146
Determine routes between internal ASes 146
Determine routes between external and internal ASes 147
Configure external meshes 149
Know about BGP routing details 150

BGP multihop 150
BGP load balancing 150
BGP next hop 151
BGP routing 151

CHAPTER 13 **Simulate Quality of Service (QoS)** 153
QoS requirements 153
Edit QoS requirements 155
QoS bound and QoS violations 155
Policies and QoS bound calculations 156
Interface queue properties and QoS bound calculations 156
Priority 157
Weight 158
Police limits 158
Interface QoS bound calculations using multiple QoS parameters 159
Service class QoS bound calculations using multiple QoS parameters 160
View QoS bounds and QoS violations 160
Configure queues and service classes 161
Create service classes 161
Create queues 162
View queue and service class information 162
Create service class policies 163

CHAPTER 14 **Simulate VPN** 165
VPN model 166
VPN objects 166
VPN topology and connectivity 166
VPNs 167
Create VPNs 167
Create new VPNs 167
VPNs table 168
Identify interfaces used by VPNs 169
VPN nodes 169
Create VPN nodes 170

Add VPN nodes to VPNs 171
VPN nodes table 171
Layer 3 VPN example 172
VPN simulation analysis 175

CHAPTER 15 **Simulate Advanced Routing with External Endpoints** 177
Routing with external endpoints 177
Create external endpoints and their members 178
Specify external endpoint members 178
Simulate routing 180
Traffic distribution 180

CHAPTER 16 **Simulate Multicast** 183
Representation of SSM parameters 183
Discovered versus simulated multicast flows 183
Set global multicast simulation parameters 184
Flow hops 184
Cisco next hop 185
Multicast demands 185
Simulated multicast demands 185
Multicast flows 186
View multicast flows 186
Create demands for multicast flows 187

PART III **Traffic Engineering and Optimization** 189

CHAPTER 17 **Optimize Metrics in the Network Core** 191
Core versus edge 191
Perform metric optimization 192
Tactical metric optimization 196
Metric optimization reports 196

CHAPTER 18 **Configure MPLS Routing** 199
Supported LSP types 199

Create and visualize LSPs 200
LSP paths 201
Path options and active path 201
Create LSP paths 201
Path latency calculations 202
Route demands through LSPs 203
Route demands through intra-area LSPs 203
Inter-area LSPs 203
Routing inter-area LSPs 203
Route demands through specific LSPs (Private LSPs) 204
Create private demands for existing LSPs 205
Create private LSPs for demands 206
Delete demands when deleting private LSPs 207
Load sharing between LSPs 208
Set global simulation parameters 209
Set LSP establishment order 210
Troubleshoot LSP simulation 210
Run LSP simulation diagnostics 211
Use simulation diagnostics to troubleshoot 211

CHAPTER 19 **Optimize LSPs 213**
Optimize disjoint LSP paths 213
Specify optimization inputs 214
Disjoint routing selection 214
Disjoint path requirements 214
Constraints 216
Create disjoint groups 217
Run the LSP disjoint path optimization tool 218
Disjoint report 219
Optimize LSP loadshare 219
Run LSP loadshare optimization 220
Minimization example 222
Bin example 224
Optimize LSP setup bandwidth 224

Run LSP setup bandwidth optimization 225
LSP setup BW optimization report 227

CHAPTER 20 **Configure RSVP-TE Routing** 229
Dynamic LSP routing and CSPF 229
Set interface MPLS properties 230
View LSP reservations 232
Calculate LSP setup bandwidth 232
RSVP LSP paths 234
Hot standby paths 234
Example response to failures 234
Named paths and explicit LSP routing 235
Named path hops example 236
Create named paths and their hops 237
Edit named paths and their hops 238
Actual paths 238
Deactivate actual paths for simulations 239
Affinities 240
Create and edit affinities 240
Assign affinities to interfaces 241
Associate LSPs with affinities 241
Inclusion and exclusion rules 241
Example 242
Assign LSPs to affinities 243
Assign LSP paths to affinities 243
Assign affinities when creating LSP meshes 244
Set global simulation parameters 245
Fast Reroute simulations 246
FRR fundamentals 247
FRR LSP routing 247
Link and node protection 248
SRLG protection 248
Routing around Failures 248
Set up fast reroute simulations 249

Run fast reroute simulation 252
Simulate autobandwidth-enabled LSPs 252
Example autobandwidth convergence without and with failures 253
Circuit Style RSVP LSPs 256
Key features of CS-RSVP LSPs 256
CS-RSVP LSP paths 256
Restore path calculation 257
CS-RSVP LSP path simulation 258
Biased routing 258
Create CS-RSVP LSPs 259
CS-RSVP LSP optimizer 260
Create CS-RSVP LSPs using the optimizer tool 260
Optimize the existing CS-RSVP LSPs using the optimizer tool 262
CS-RSVP LSP configuration parameters 263
CS-RSVP LSP simulation diagnostics 264
Run the CS-RSVP LSP simulation diagnostics tool 265
CS-RSVP LSP simulation diagnostics report 265
Example 266
Visualize CS-RSVP LSPs, and associated LSPs and LSP paths 266
Visualize CS-RSVP LSPs 266
Visualize LSPs associated with CS-RSVP LSPs 268
Visualize LSP paths associated with CS-RSVP LSPs 269
Advanced RSVP-TE LSP simulations 269
Ignore reservable bandwidths for capacity planning 269
Ignore reservable bandwidth constraints 270
Example 271
P2MP LSPs 273
P2MP LSP bandwidth 274
P2MP LSP demands 276
Create P2MP LSPs and sub-LSPs 276
Create Demands for P2MP LSPs 278
Delete P2MP LSPs and sub-LSPs 279
Unresolved LSP destinations and hops 280

CHAPTER 21 **Optimize RSVP-TE Routing** 281

Specify inputs for optimization 282

Select LSP groups 283

Set optimization parameters 285

Set disjoint groups 285

Setting avoidance constraints 286

Set post-optimization parameters 287

Run RSVP-TE optimization 287

Analyze optimization output 289

Properties created 289

Reports 289

Optimized LSPs reconfigured 289

CHAPTER 22 **Perform Explicit and Tactical RSVP-TE LSP Optimization** 291

Run Explicit RSVP-TE LSP optimization 291

Optimization options 293

Tactical explicit RSVP-TE LSP optimization 299

CHAPTER 23 **Configure Segment Routing** 301

SR LSP segment types 301

Node segment list hops 302

Interface segment list hops 302

LSP segment list hops 303

Anycast group segment list hops 304

SR LSP paths 306

SR LSP routing 306

Inter-AS LSP routing 307

SR LSPs and their LSP paths 308

Create segment lists 308

Create anycast groups 309

Create SIDs 309

Create node SID 309

Create SRV6 node SID 310

Create interface SID 310
Create SRV6 interface SID 311
Create flex algorithm 311
SR-TE protection 312
Simulate SR-TE tunnels before convergence 313
Simulate SR-TE tunnels after convergence 313
Constraints 313

CHAPTER 24 **Optimize Segment Routing** 315
Optimize SR-TE 315
Specify inputs for SR-TE optimization 316
Optimize the path metric 316
Bound and margin 316
Constraints 317
Run SR-TE optimization 318
Optimization report 319
Optimize and analyze SR-TE bandwidth 319
Select operating modes for SR-TE bandwidth optimization 320
Specify inputs for SR-TE bandwidth optimization 320
Interface utilization thresholds 320
Rerouting demands 321
Constraints 321
Optimize bandwidth 323
Analyze congestion under different failure sets 325
Bandwidth optimization report 326

PART IV **Access Reports, Jobs, Patch Files, and Changeover Tool** 327

CHAPTER 25 **Access Reports** 329
Plan comparison reports 329
Create plan file comparison reports 330
Report columns 331
Report sections 335
Traffic comparison reports 336

Create traffic comparison reports 337
View reports 338

CHAPTER 26 View Jobs 339
Job Manager 339
User role permissions for Job Manager 339
View job details 340
Run tools or initializers using CLI 342
Run external scripts 343
Example: Run external scripts 344

CHAPTER 27 Update Configuration from One File to Another 347
Run the Changeover tool 347
Analyze reports 349

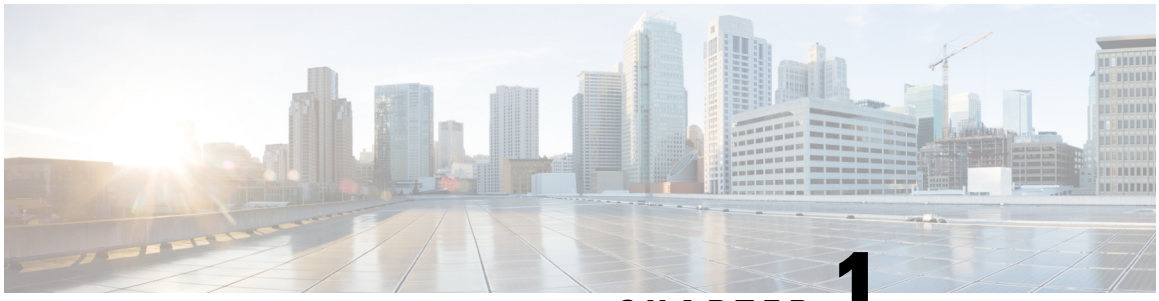
CHAPTER 28 Create and Use Patch Files 351
Create patch files 351
Apply patch files 352
View or edit patch files 353

CHAPTER 29 Export Routes 355
Route tables 355
Output table column details 355
Example: Export route tables 356
Export routes 356

APPENDIX A Cross Table Filters 359
Supported cross table filters 359

APPENDIX B Legend for Plan Objects 365
Interfaces and Circuits 365
Nodes 366
Sites 368
Demands 369

LSPs 370
LSP paths 370
LSP named paths 371
LSP actual paths 372
SR LSP paths 372



CHAPTER 1

Overview

This document offers guidance on how you can use the Cisco Crosswork Planning Design application to

- model, simulate, and analyze failures, design changes, and impact of traffic growth
- optimize your network for maximum efficiency
- visualize the health of your network, and
- generate reports that compare traffic

This chapter contains the following topics:

- Core capabilities of Cisco Crosswork Planning , on page 1
- Cisco Crosswork Planning Design use cases, on page 2
- System overview, on page 2

Core capabilities of Cisco Crosswork Planning

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 to network performance of changes in the network topology, which is 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.

Cisco Crosswork Planning Design use cases

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:

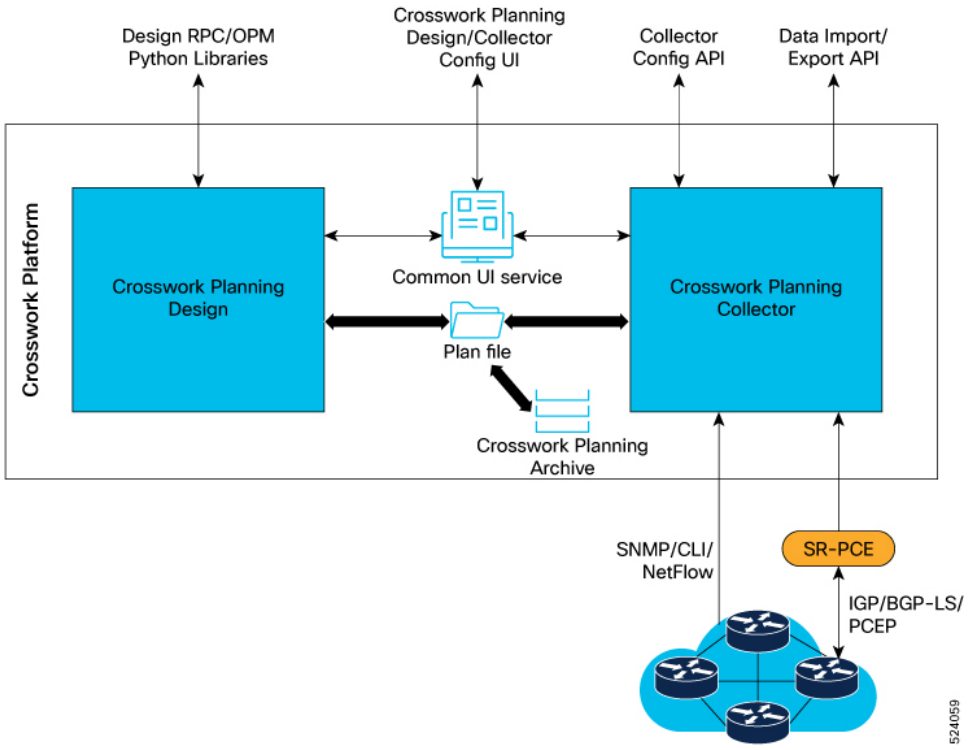
- How can I visualize the current health of the network?
- What caused the outage experienced earlier today?
- How would a current failure affect the network at peak time tomorrow?
- Which VPN customers will be affected by a planned outage or by an unplanned failure?

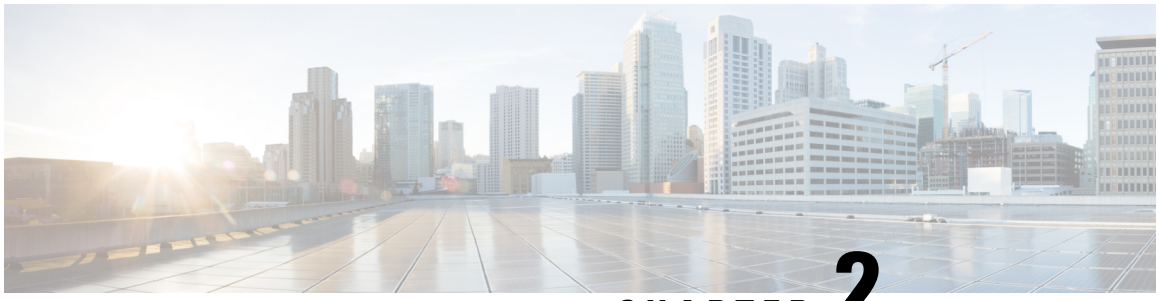
System overview

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: System overview





CHAPTER 2

Get Started

- Key concepts, on page 5
- Log in to and log out of Cisco Crosswork Planning , on page 7
- Dashboard, on page 7
- Design engine instances, on page 11

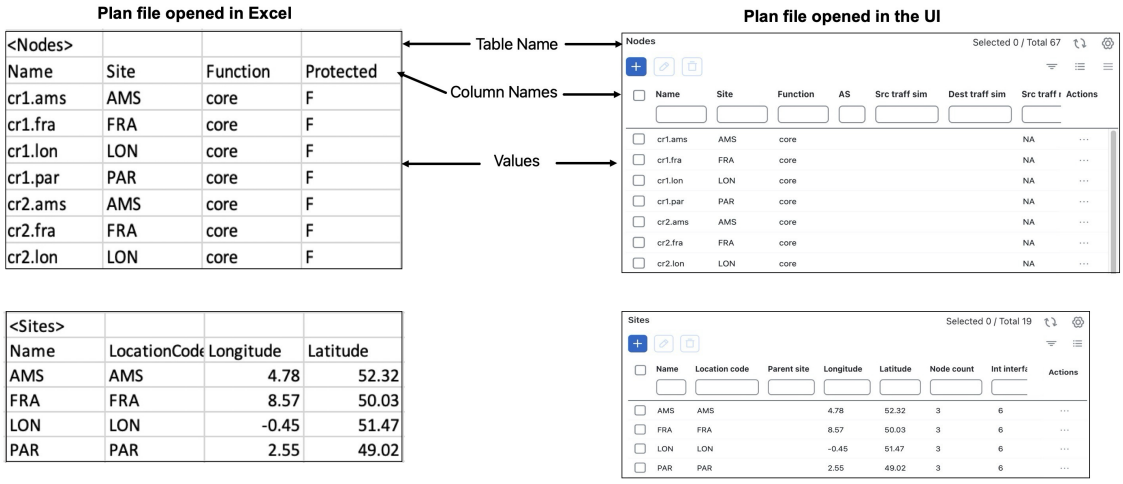
Key concepts

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 this 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, see Understand Plan Objects, on page 47.

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. You can open a maximum of three plan files simultaneously. Each plan file is maintained individually as a network model, with only one network model active at any given time. You can switch the active network model among all opened network models and all activities are applied to the active network model only.

- **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 the space it operates in user space or job space, it operates in *synchronous mode* 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 the other user activities.

In asynchronous mode, the engine processes every 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, see Design engine instances, on page 11.

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, see Create and Use Patch Files, on page 351.

Log in to and log out of Cisco Crosswork Planning

Cisco Crosswork Planning is a browser-based application. For details on supported browser versions, see the *"Supported web browsers" section in the Cisco Crosswork Planning 7.1 Installation Guide*.

After installing Cisco Crosswork Planning, follow these steps to access the Cisco Crosswork Planning UI.

Procedure

Step 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. When this happens, follow the prompts to add a security exception and download the self-signed certificate from the server. After you do this, the browser accepts the Cisco Crosswork Planning server as a trusted site in all subsequent logins.

Step 2 Enter your username and password. The default administrator username and password is **admin**. This account is created automatically at installation. You must change the initial password for this account during installation verification. We strongly recommend that you keep the default administrator credential secure, and never use it for routine logins. Instead, create new user roles with appropriate privileges and assign new users to those roles. At least one of the users you create should be assigned the "administrator" role.

Step 3 Click **Login**.

Step 4 To log out, at the top right of the main window, click  > **Logout**.

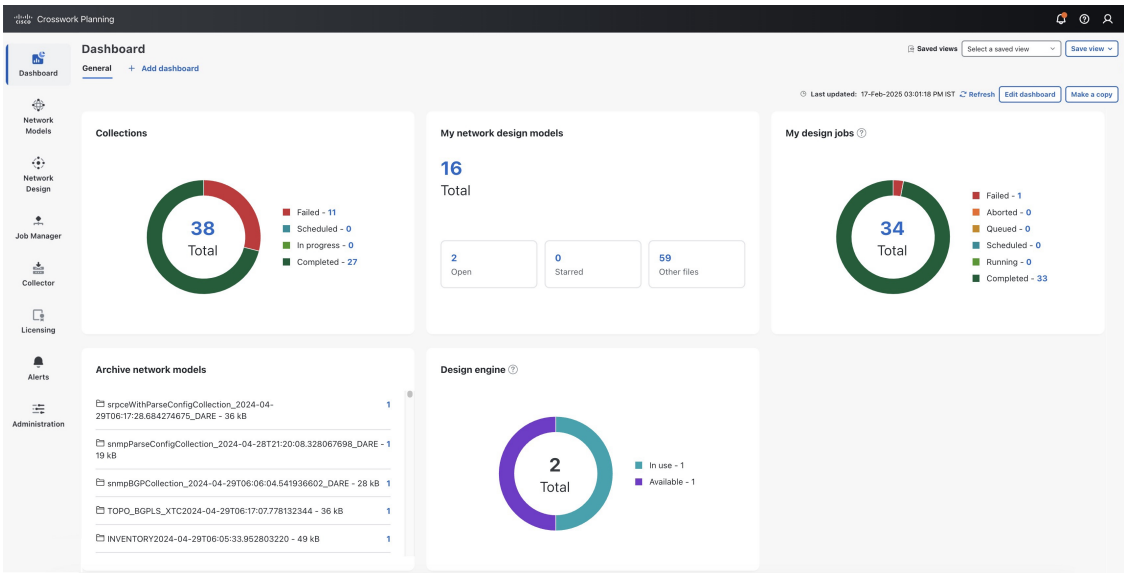
Note
Logging out does not close the plan file you are working on. It remains open.

Dashboard

The **Dashboard** page provides a quick operational summary of Cisco Crosswork Planning. This page consists of various dashlets, which vary based on the Cisco Crosswork Planning application installed.

For example, 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 8, 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. For details, see Customize dashlets, on page 10 .

Custom dashboards

In addition to the default **General** dashboard, you can create custom dashboards to include dashlets of your choice.

Create custom dashboards

Follow these steps to create custom dashboards.

Procedure

Step 1 Click the **+ Add dashboard** at the top right.

Step 2 Click **+ Add dashlet**. The Add Dashlets window appears.

Step 3 From the list, select the dashlet you want and click **+ Add**.

Step 4 Enter the name for the dashboard by clicking the edit icon next to the dashboard title at the top left. Ensure that the title is between 5 and 50 characters long.

Step 5 Click **Save** in the Rename Dashboard dialog box.

Step 6 Click **Save** to save the dashboard.

Rename custom dashboards

Follow these steps to rename the custom dashboards.



Note You cannot rename the default **General** dashboard.

Procedure

Step 1 Go to the required custom dashboard and click **Edit dashboard** at the top right.

Step 2 Enter the new name for the dashboard by clicking the edit icon next to the dashboard title at the top left. Ensure that the title is between 5 and 50 characters long.

Step 3 Click **Save** in the Rename Dashboard dialog box.

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

Delete custom dashboards

Follow these steps to delete the custom dashboards.



Note You cannot delete the default **General** dashboard.

Procedure

Step 1 Go to the required dashboard and click **Edit dashboard** at the top right.

Step 2 Click **Delete dashboard** at the top right.

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

Customize dashlets

By default, all the available dashlets appear in the default **General** dashboard. Using the **Edit dashboard** button 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

Step 1 Go to the required dashboard and click **Edit dashboard** at the top right. The page becomes editable.

Step 2 To rearrange dashlets:

a) Click ⁘ in the dashlet you want to move, and drag and drop to the position you want.
b) Click **Save**.

Step 3 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) Click **Save**.

Step 4 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) Click **Save**.

Step 5 To copy dashlets:

a) Click ••• > **Copy** in the dashlet you want to copy.

A copy of the dashlet appears in the dashboard.

b) Click **Save**.

Step 6 To remove dashlets:

a) Click ••• > **Delete** in the dashlet you want to delete.

The selected dashlet is removed from the dashboard.

b) Click **Save**.

Reset the General dashboard to the default view

Follow these steps to reset the General dashboard to the default view.

Note You cannot perform this procedure on custom dashboards.

Procedure

Step 1 On the **General** dashboard, click **Edit dashboard** at the top right.

Step 2 Click the **Reset to default** button at the bottom of the page.

The Reset Dashboard confirmation window appears.

Step 3 Click **Reset** in the confirmation window.

Copy dashboards

To make a copy of the dashboard, go to the required dashboard and click **Make a copy** at the top.

Design engine instances

For an overview of design engines, see Key concepts, on page 5.

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, and 8 vCPU (8000m vCPU) and 24 Gi of memory for extra large (XLarge) profiles.

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.

Resource allocation for design engine instances

In large scale deployments, you may need to complete simulation analysis or capacity planning faster, or perform tasks on huge plan files. In such cases, you need to allocate higher vCPU and memory. Cisco Crosswork Planning allows you to allocate the vCPU and memory values to engine instances while creating them.

The maximum vCPU and memory you can allocate to an instance depend on the system's total resources. For example, if your system has 4 vCPUs and 64 Gi of memory, you can allocate a maximum of 4000m vCPU (that is, 4 vCPU) and 64 Gi of memory for an engine instance. To find the system's resource information, go to the **Administration** > **Crosswork Manager** > **Crosswork summary** > **System summary** page. The resource details allocated to the design engine instances are listed under the **Engine profile** column of the Design Engine Instances page.

Create design engine services

Follow 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 details on managing user roles and permissions, see the *"Manage Users" section of the Cisco Crosswork Planning 7.1 Collection Setup and Administration* document.

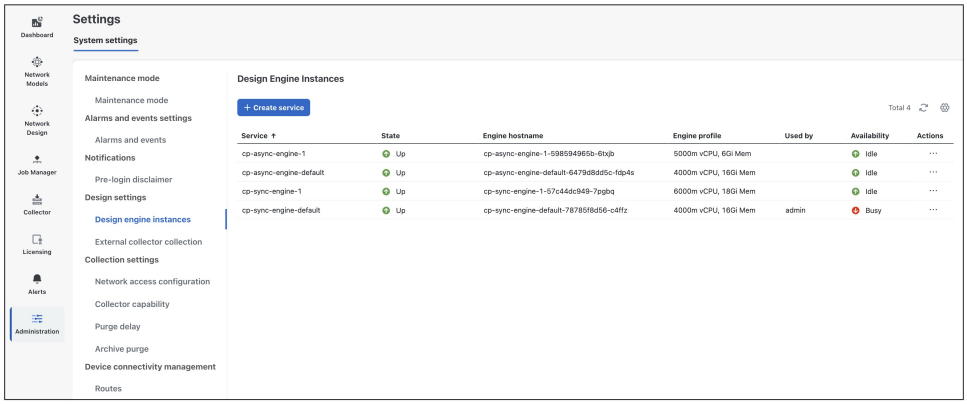
Note You can allocate a maximum of 10 **design engines** for both synchronous and asynchronous jobs combined.

Procedure

Step 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



Step 2 Click **Create Service**.

The Create Service page opens.

Step 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.

Step 4 Enter the **Memory** in Gi (gibibyte) as per your requirement. Ensure that the entered value does not exceed the system's memory.

Step 5 Choose whether to create synchronous or asynchronous instances.

Step 6 Click **Save**.

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, see Select design engine profiles, on page 30.
- The newly created asynchronous engine profile appears under the **Engine Profile** section on the **Run Settings** page while running the tools as scheduled jobs.

Reset engine sessions

Follow these steps to reset engine sessions.

Before you begin

Ensure you have the "CW Planning Design Setting APIs" Role permission (**Administration** > **Users and Roles** > **Users**). For more details on managing user roles and permissions, see the *"Manage Users" section of the Cisco Crosswork Planning 7.1 Collection Setup and Administration* document.

Procedure

Step 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.

Step 2 Under the **Actions** column, click ••• > **Reset session** for the engine instance whose session you want to reset.

Step 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 profiles, the opened plan files are closed.
- For async engine profiles, the scheduled jobs are aborted.

Delete design engine instances

Follow these steps to delete the dynamically 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 details on managing user roles and permissions, see the *"Manage Users" section of the Cisco Crosswork Planning 7.1 Collection Setup and Administration* document.

Note You cannot delete the default design engine instances.

Procedure

Step 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.

Step 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.

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

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

Recreate design engine services

The dynamically created design engine services 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.

Follow these steps to recreate them.

Procedure

Step 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.

Step 2 Under the **Actions** column, click ••• > **Recreate** for the service you want to recreate.

Step 3 Click **Recreate** in the confirmation dialog box.

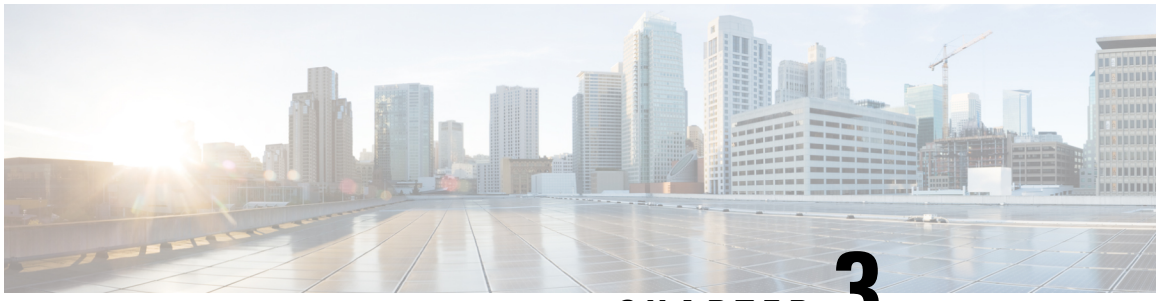
The selected design engine instance is recreated.



PART I

Visualize Network Models

- Visualize Network Models, on page 19
- Understand Plan Objects, on page 47



CHAPTER 3

Visualize Network Models

This chapter contains the following topics:

- Create or import network models, on page 19
- Open plan files, on page 22
- Delete plan files, on page 23
- Download plan files, on page 24
- Key features of the Network Models page, on page 24
- Key features of the Network Design page, on page 27
- Select design engine profiles, on page 30
- Assign sites to nodes, on page 30
- Assign geographic locations to sites, on page 34
- Network topology, on page 36
- Traffic utilization, on page 39
- Network summary tables, on page 42

Create or import network models

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

This section explains how to import the plan files into the user space from local machine.

Before you begin

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

Procedure

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

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

Step 2 Click the **Import plan file** button.

The Import Plan File window appears.

Step 3 Click **Browse** and choose the plan file that you want to import.

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

Follow 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 the available plan files, then it is deleted from this list as well.

Step 5 Click **Import**.

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

Import plan files from archive

This section explains how to import plan files into the 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 details, see the "Scenario 2: When the Cisco Crosswork Planning Design and Collector Applications are Installed on Different Machines" section in the *Cisco Crosswork Planning 7.1 Collection Setup and Administration*.

Procedure

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

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

Step 2 On the left pane, under **Local archive** or **Remote archive**, list of archived collections are displayed. Select the required collection name from the list. The right panel displays the list of plan files created under this collection at various scheduled times. Use the **Last updated** column to know the time at which the plan file was created.

Step 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 window appears.

Step 4 (Optional) In the **Save as** field, enter a new name for the plan file.

Step 5 (Optional) Select the required tags from the list (if available) or create new tags.

Follow 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 the available plan files, then it is deleted from this list as well.

Step 6 Click **Save**.

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

Create new network models manually

This section explains how to create new network models in the user space manually.

Procedure

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

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

Step 2 Click the **Create new** button.

The New Network Model window opens.

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

Note
Make sure to include the .txt or .pln extension while entering the plan file name.

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

Follow 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 the available plan files, then it is deleted from this list as well.

Step 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 details, see Create objects, on page 56.
- Open the plan file in the **Network Design** page (see Open plan files, on page 22) to perform any actions, as per your requirement.

Open plan files

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 19.

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 wherein you can perform various actions on them, as required.

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.

To open from the Network Models page:	To open from the Network Design page:
1. From the main menu, choose **Network Models**. By default, **User space** > **My network models** is selected. A list of available plan files is displayed. 2. Open the required plan file in either of the following ways: • Click the name of the required plan file. OR • Under the **Actions** column, click ••• > **Open** for the plan file you want to open. The plan file opens in the **Network Design** page.	1. From the main menu, choose **Network Design**. 2. 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. 3. Open the required plan file in either of the following ways: • Click the name of the required plan file. OR • Under the **Actions** column, click ••• > **Open** for the plan file you want to open. The plan file opens in the **Network Design** page.

After the plan file opens:

- 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 from 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

This section explains how to delete plan files from the user spaces.

Important considerations for plan file deletion

Note these points while deleting the plan files:

- You can delete up to 10 files at a time.
- Only admin users can delete the files which are available in the other user spaces.
- You cannot delete the plan files that are currently open.
- Deleting the plan files will result in the deletion of the associated reports.

Procedure

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

Step 2 From the user space, select the plan files that you want to delete and click the icon at the top.

If you are deleting a single file, then you can use the ••• > **Delete** option under the **Actions** column.

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

Download plan files

This section explains how to download plan files.

You can download up to 10 files at a time.

Procedure

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

Step 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, then you can use the ••• > **Download** option under the **Actions** column.

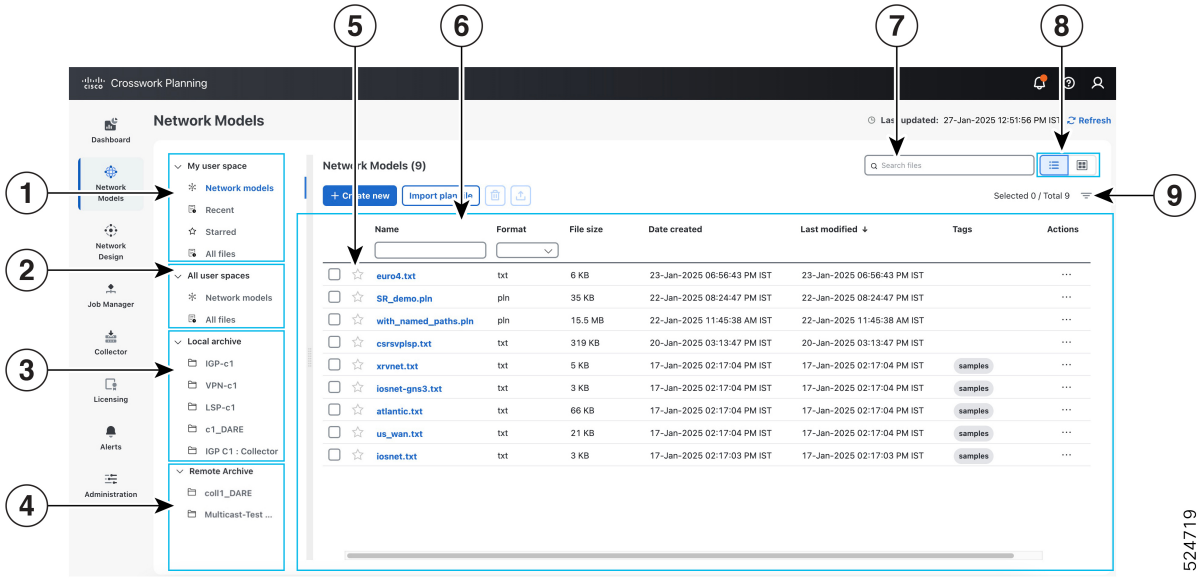
The plan files are downloaded to your local machine. When downloading more than one plan file, the files are bundled together in a single .tar file.

Key features of the Network Models page

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

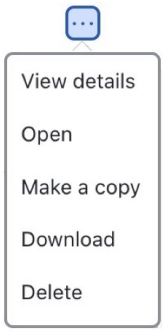
The following image and the table describe the various actions you can perform in the **Network Models** page:

Figure 5: Network models user interface



524719

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, the users cannot view the plan files used by the other users. By default, My user space contains few sample plan files. There are four sub-sections under this section: • **Network models**—Lists all the plan files accessible to the logged-in user. **Note** 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. • **Recent**—Displays the logged-in user's recently accessed plan files. • **Starred**—Displays the plan files that are marked as favorites by the logged-in user. • **All files**—In addition to plan files, Cisco Crosswork Planning supports 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.
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.

Callout No.	Description
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. Use the ••• icon under the **Actions** column to • view the details of the plan file. **Note** Use this option to add or delete tags from the plan files. • open the plan file in the Network Design page • make a copy of the plan file in the user space • download the plan file to your local machine, and • delete the plan file from the user space. ***Figure 6: Actions menu*** 
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 inside the user space can be either 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. 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

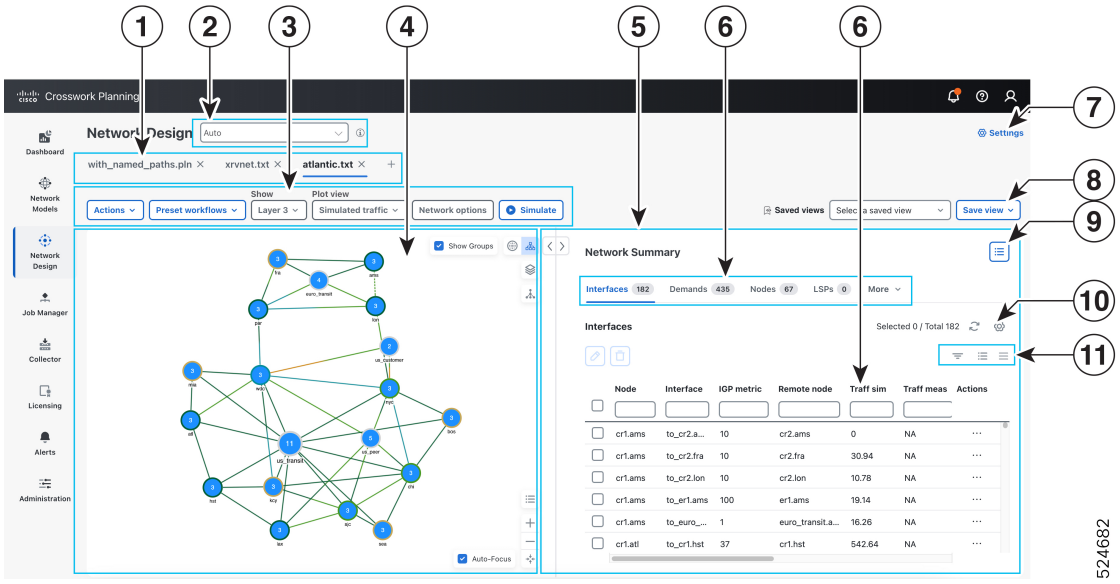
When you click the plan file name in the Network Models page, it opens in the **Network Design** page. This page displays the network plot which shows sites that appear to be connected by circuits, though they actually connect to nodes (routers) within each site. The nodes within the sites are connected with intra site 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

The following image and the table describe the various actions that you can perform in the Network Design page:

Figure 7: Network Design user interface



Callout No.	Description
1	**Plan file tabs**—Identifies which plan files are currently open. At a time, you can open three plan files. 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 details, see Select design engine profiles, on page 30.

Callout No.	Description

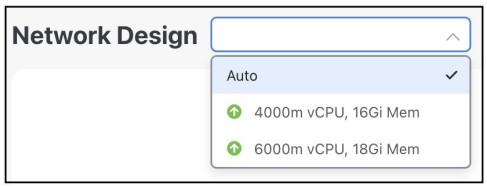
3	**Toolbar**—Provides options for network modeling, simulation, and optimization.
	• **Actions**—Contains options related to tools, initializers, and reports. Use the **Actions** > **Insert** option to create various plan objects.
	• **Preset workflows**—Includes options that allow you to navigate to the simulation and optimization tools quickly.
	  • **Evaluate impact of failures**—Helps you navigate to the Simulation analysis tool.
	  • **Evaluate impact of traffic growth**—Helps you navigate to the Create growth plans tool.
	  • **Perform capacity planning**—Helps you navigate to the Capacity planning optimization tool.
	  • **Perform optimization**—Helps you navigate to the various LSP, SR LSP, and RSVP LSP optimization tools.
	• **Show Layer 3**—Indicates that the UI is currently displaying Layer 3 (L3) networks.
	• **Plot view**—Lets you choose the type of traffic displayed in the plot. For details, see Network plots, on page 36.
	• **Network options**—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.
	• **Simulate icon** ()—Triggers resimulation. By default, any change that voids the current simulation does not automatically trigger resimulation. For example, changes in topology trigger requires a resimulation. For details, see Auto-resimulation, on page 62.

Callout No.	Description
4	**Network plot**—Displays a visual representation of the network topology. The color on each interface represents the percentage of traffic utilization on that interface (see Traffic utilization, on page 39). When you select the objects in the Network Summary tables, such as demands or interfaces, they are highlighted in the plot. • and icons—There are two views available to visualize the network topology graphically: Schematic view and Geographical view. For details, see Graphical network topology views, on page 38. • **Show Groups**—In a schematic view, use this check box to group or ungroup the nodes. For details, see Grouping of nodes, on page 38. • **Auto-Focus**—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. • The legend icon ()—Use this icon to display a legend of icons and line types in the plot, and their meanings. • and icons—Use these icons to gradually zoom in and zoom out the network topology. • The Auto-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 details, see Network summary tables, on page 42.
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, on page 42 are the most commonly used and are the defaults. To bring a table to the forefront and make it the active table, 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. The tables display a series of rows and columns, where the rows are objects and the columns are properties.
7	**Settings button**—Use the Settings button at the top right corner to enable/disable the Auto-resimulate option. For more information, see Auto-resimulation, on page 62.
8	**Save views**—Use the **Save view** button to save views of the commonly used tables. For more details, see Save table views, on page 45.
9	**Show/hide tables**—Use to show or hide one or more object tables from the Network Summary panel. For details, see Show or hide tables or columns, on page 44.
10	**Show/hide table columns**—Use to show or hide one or more columns from the object tables. For details, see Show or hide tables or columns, on page 44.

Callout No.	Description
11	**Filter icons**—There are three filter options available in Cisco Crosswork Planning: Floating filter, Advanced filter, and Cross table filter. For details, see Search and filter in tables, on page 44.

Select design engine profiles

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, see Design engine instances, on page 11.



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

Before you begin

Ensure that the required engine profiles are available. For details on creating design engine instances, see Create design engine services, on page 12

Procedure

Step 1 From the main menu, choose **Network Design**.

Step 2 Close any open plan files.

Step 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.

Step 4 Open the required plan files and start working on them. They will now use the selected engine profile for all synchronous tasks.

Assign sites to nodes

The **Assign sites to nodes** initializer helps to reorganize your network by changing the node-to-site mappings. 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 would typically place all nodes in the same geographic location or point-of-presence (PoP) into a single site.

Using this initializer, you can:

- 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, which uses rules to map nodes to sites based on regular expression substitutions. Because the table is stored in the plan file (as a <NodeSiteMappingRules> table), you can maintain and reuse it.

Site assignment rules

This section explains how the **Assign sites to nodes** initializer assigns sites to nodes.

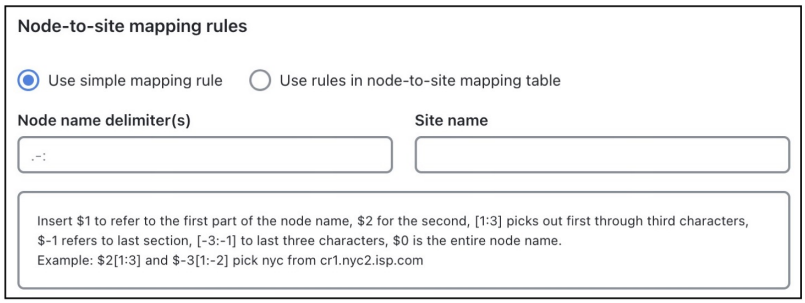
If a node is ...	Then ...
in an external AS	Cisco Crosswork Planning assigns it to a site named after the AS name. If the AS name does not exist, Cisco Crosswork Planning assigns the node to a site named after its ASN.
not in an external AS and is not a pseudo-node (PSN)	Cisco Crosswork Planning assigns nodes to sites based on the Simple mapping rules, on page 31 or on the Node-to-Site mapping rules, on page 32, 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 in the **Assign sites to nodes** window.

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 are a period, a hyphen, and a colon (.-:). New site names are based on 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 the following 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 in the **Assign Sites to Nodes** window.
2. Click  and enter the details.

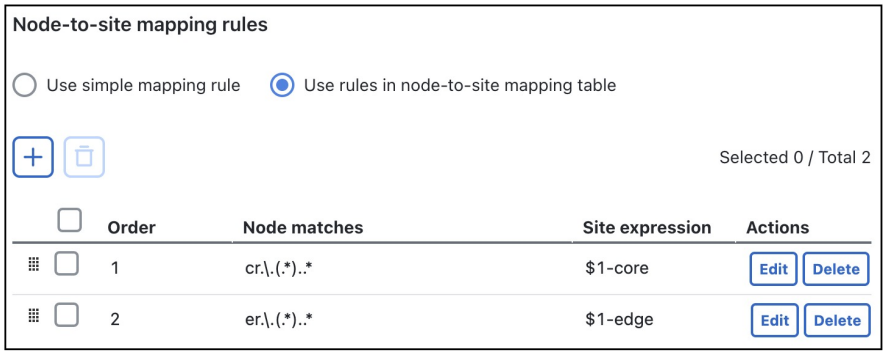
Column	Description
Order	Identifies the order in which rules are applied.
Node matches	Regular expression matching the node names.
Site expression	Site name expression, which can use references in the Node matches rule.

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 8: Node-to-site mapping example, on page 32 matches cr1.lax into lax-core and er1.lax into lax-edge.

Figure 8: 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, but also maps node cr1.okc.isp.net to site okc.
..(.)\.(.*)..*	$2-$1	Map node cr1.par.isp.net to site par-1.

Run the Assign sites to nodes initializer

This section explains how to run the Assign sites to nodes initializer.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 Choose **Actions** > **Initializers** > **Assign sites to nodes**.

Step 3 Choose one or more nodes for which you want to assign sites.

Step 4 Click **Next**.

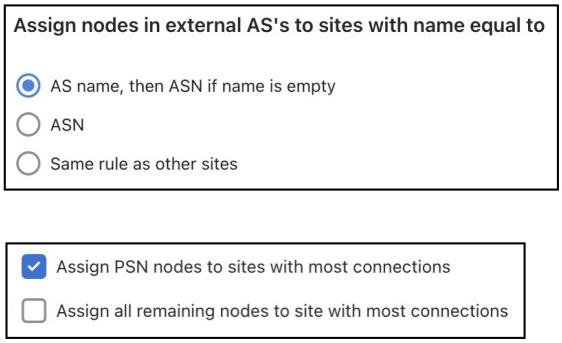
Step 5 Choose which method of node-to-site mappings to use.

- 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 the **Edit** button.
  - To delete an existing rule, click the **Delete** button. To delete multiple rules, select the rules and click .

Step 6 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.

Step 7 Choose the options as follows.

Option	Description

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

Step 8 Click **Next** and verify the assignments.

Step 9 When you are satisfied with the mappings, click **Submit**.

Assign geographic locations to sites

Cisco Crosswork Planning includes a database of worldwide city names and airport codes that identify the longitude and latitude of major cities. Use the **Actions** > **Initializers** > **Assign locations to sites** initializer that accesses this database to let 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.

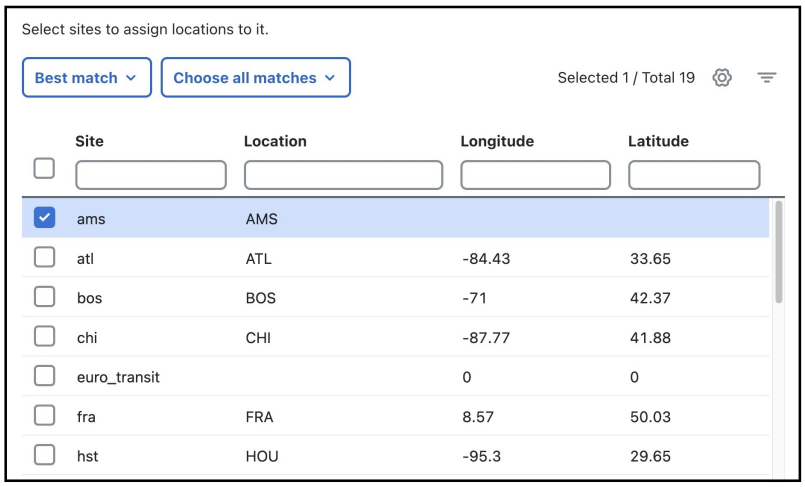
Follow these steps to assign multiple sites to locations.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Initializers** > **Assign locations to sites**.

Figure 9: Assign locations to sites



Step 3 Set the longitude and latitude values using one of the following methods. In each of these two methods, Cisco Crosswork Planning fills in the fields with the most closely associated airport or city code in its database, thereby assigning 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.

Step 4 Click **Save** to accept the newly assigned locations.

Initialize site location

As an alternative to using the **Assign locations to sites** initializer, you can set a location from the Edit window. For best results, use airport codes.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

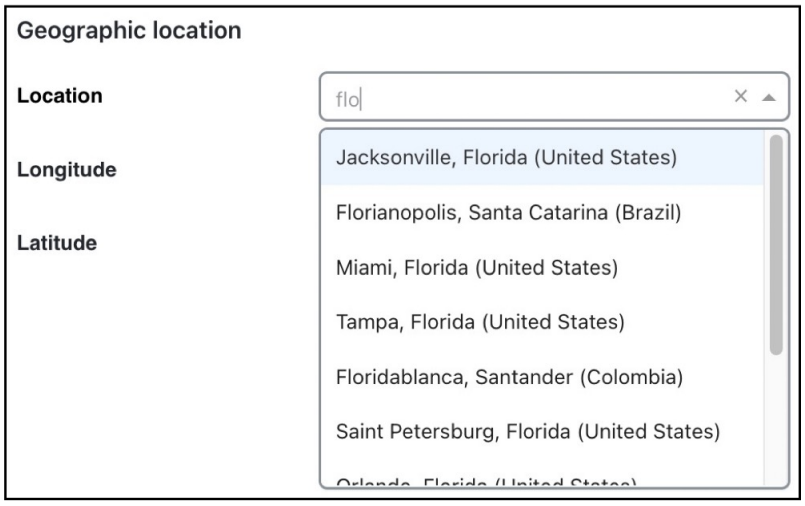
Step 2 In the Network Summary panel on the right side, select one or more sites from the **Sites** table.

Step 3 Click .

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

Step 4 In the **Location** field, enter a geographic name, such as a city name. Cisco Crosswork Planning fills in the **Location** fields with the most closely associated airport or city code in its database. Choose the location.

Figure 10: Choosing geographic location



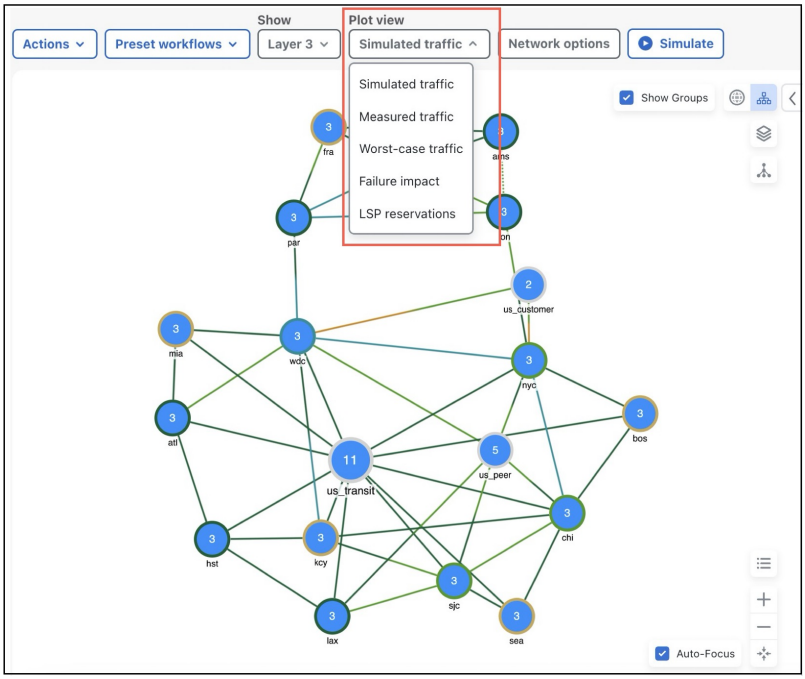
Step 5 Click **Save** to accept the newly assigned location.

Network topology

Network plots

The network plot, also referred to simply as the *plot*, is the main area showing the network topology. It can contain both sites and nodes. From the main menu, choose **Network Design** to view the network plots. Although Cisco Crosswork Planning shows only the sites in the network plot and tables, you can view those sites in more depth.

Figure 11: Network plot views



There are multiple network plot views available. The default view is **Simulated traffic**. To change the view, select it from the **Plot view** drop-down list.

- **Simulated traffic**—View of simulated traffic that can be used for capacity planning, what-if analysis, and failure planning. Circuits are sized proportional to their capacity. Interfaces are colored according to the percentage of simulated traffic (in comparison to available capacity as 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 Simulate in the toolbar. For details on updating the auto-resimulation setting, see Auto-resimulation, on page 62.

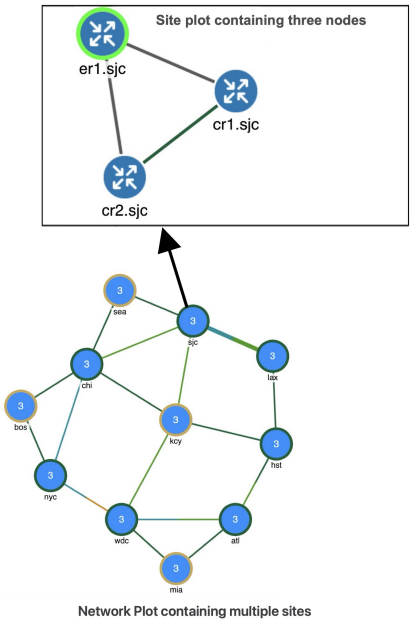
- **Measured traffic**—View of traffic as measured from live views of the network. Circuits are sized proportional to their capacity. Interfaces are colored according to the percentage of measured traffic (in comparison to available capacity as defined by the size of the interface or the QoS bound).
- **Worst-case traffic**—Interfaces are colored to show the maximum utilization of the interface over all failure scenarios as defined by the most recently run simulation analysis. For more information, see Evaluate Impact of Worst-Case Failures, on page 103.
- **Failure impact**—Interfaces and nodes are colored to indicate how a failure of that interface or node would impact other interfaces and nodes. For more information, see Evaluate Impact of Worst-Case Failures, on page 103.
- **LSP reservations**—Circuits are sized proportional 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, see Configure RSVP-TE Routing, on page 229.

Site plots

The site plot can contain other sites or nodes and their connecting interfaces, including external interfaces that are labeled with their destination node.

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

Figure 12: Site with three nodes



Graphical network topology views

There are two views available to visualize the network topology graphically.

- **Schematic View**: Click in the network plot to visualize the topology in a schematic view. This view shows the network topology, positioned according to an automatic layout algorithm, ignoring their geographical location. 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 persisted across user sessions.
- **Geograpic View**: Click in the network plot to visualize the topology in a geographic view. This view shows network topology, superimposed on a map of the world. Each device location on the map reflects the device's GPS coordinates (longitude and latitude) as defined in the device inventory. If the GPS coordinates are not available, then this view will be unavailable and nodes will be placed in a pre-determined location.

Grouping of nodes

In the schematic 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. When a site is clicked, the topology schematic map is replaced with a new map that will only show the nodes that are part of the site and the links that interconnect these nodes.

In the geographical view ( ), the nodes are grouped and ungrouped automatically based on the geographical location.

Representation of plan objects

For detailed information on how the plan objects are represented in the network plot, click the icon in the network plot.

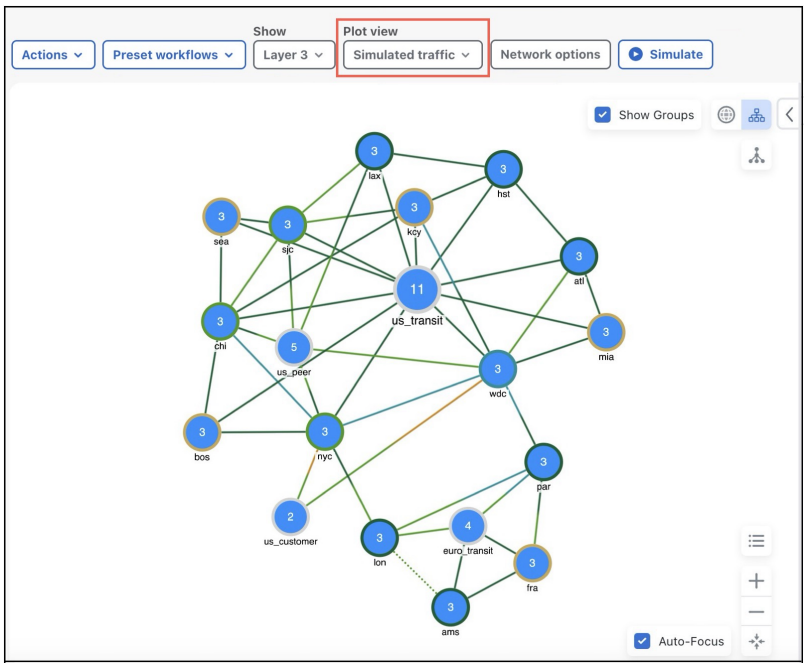
Traffic utilization

The links between nodes are colored according to the simulated, measured, or worst case utilization. Each link may have two different colors based on the utilization of each unidirectional interface that the link represents. The color fill of an interface shows the bandwidth utilization for traffic leaving that interface in proportion to the interface capacity. That is, it shows the *traffic utilization*. Rows that are selected in the tables, such as demands or interfaces, are highlighted in the plot.

Note You can observe the colored links only when you trigger the resimulation. To trigger the resimulation, click Simulate in the toolbar. By default, the auto resimulation is off in Cisco Crosswork Planning. To change this default setting, see Auto-resimulation, on page 62.

Following table lists the mapping between the color of the links and the traffic utilization.

Traffic Utilization	Color	Example
0%	Light gray	
< 30%	Dark green	
30 to 50%	Light green	
50 to 80%	Blue	
80 to 90%	Light orange	
90 to 100%	Orange red	
Above 100%	Red	
No traffic	Dark gray	

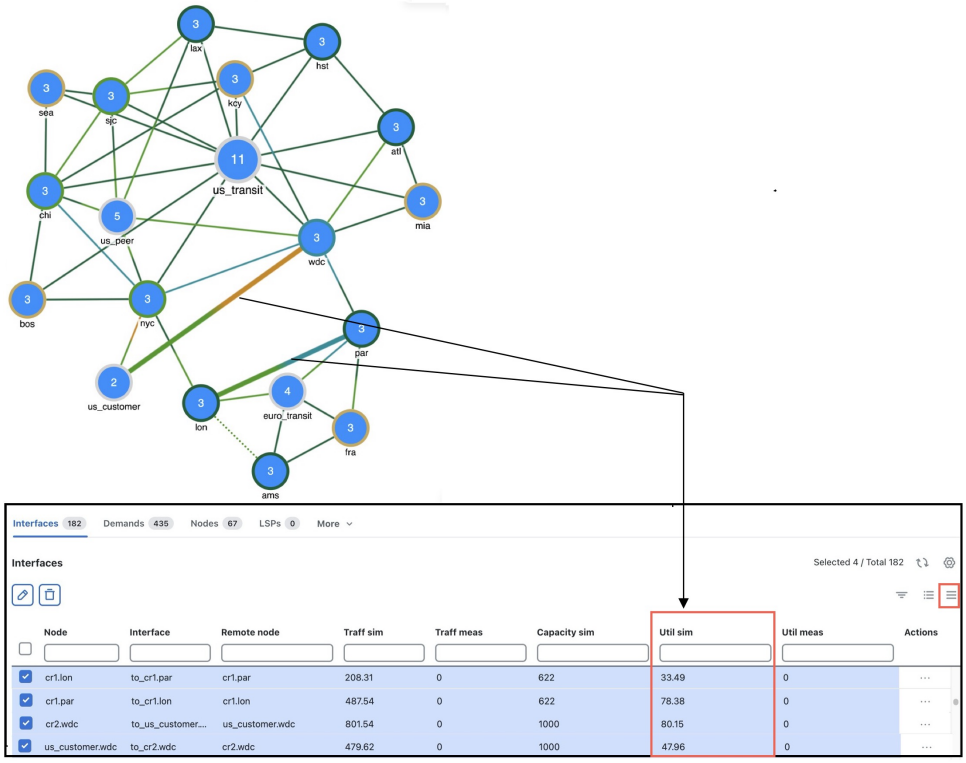


All parallel links between two nodes to be clustered to form a single link are shown as dotted lines. When a clustered link is selected, the network summary table on the right side shows the constituent links and provides a way to select each one of them separately.

Note All traffic is displayed in Mbps. In the network plot, the traffic utilization colors represent outbound traffic.

Example: Determine traffic utilization

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



Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 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.

Step 3 In the Network Summary panel on the right side, choose ☰ > **Filter to circuits**. The Circuits table opens and shows only the selected circuits.

Step 4 Select the circuits and choose ☰ > **Filter to interfaces**. The Interfaces table opens and shows only the interfaces included in the selected circuits.

Step 5 Notice the utilization values in the **Util sim** column.

- For "to_cr1.par" interface, the Util sim value is 33.49%, which is between the 30-50% level denoted by light green. The interface is filled with green that extends to fill almost half of this interface. The other half of the circuit ("to_cr1.lon") displays the traffic utilization for the reverse direction, which is 78.38% for this interface, which is between 50-80% level denoted by Blue.
- For "to_us_customer" interface, the Util sim value is 80.15%, which is between the 80-90% level denoted by light orange. The other half of the circuit ("to_cr2.wdc") displays the traffic utilization for the reverse direction, which is 47.96% for this interface, which is between 30-50% level denoted by light green.

Step 6 To show all the interfaces, clear the applied filter.

Identify most highly utilized interface

Identifying the most utilized interfaces is useful when analyzing large, complex networks.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 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.

Step 3 Notice that the most utilized interface is at the top row. Select this top row. The interface highlighted in the plot is the most highly utilized interface.

Step 4 Deselect the row to deselect the interfaces.

Network summary tables

The Network Summary tables are displayed on the right side of the Network Design page. They display a series of rows and columns, where the rows are objects and the columns are properties (Figure 13: Common table functions, on page 43). 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>.

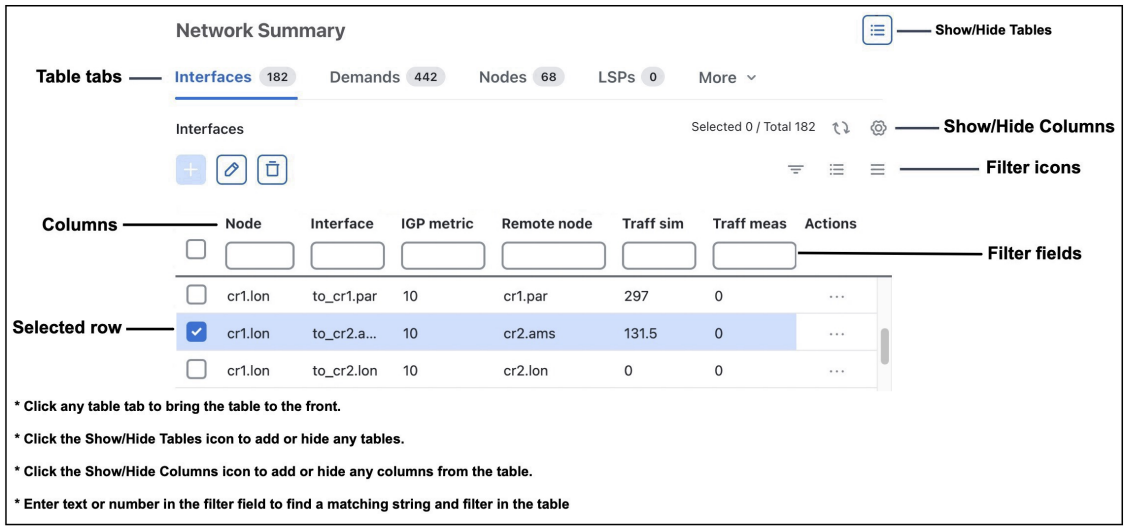
Common tables

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

To bring a table to the forefront and make it the active table, click its associated tab. For example, to open the LSPs table, click the **LSPs** tab.

- 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 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 13: Common table functions



Work with tables and object selections

This section familiarizes you with the basic dynamics between tables and objects, such as nodes, sites, circuits, interfaces, and ports. When you select an object in a table, it is automatically selected in the network plot. Conversely, whatever you select in the network plot is selected in the tables.

Task: 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

Step 1 If the plan file is not open, open it using the steps mentioned in Open plan files, on page 22.

Step 2 Notice that the Interfaces tab is selected in the Network Summary panel on the right side. Click the **Sites** tab (click **More** > **Sites**) to show the Sites table.

Step 3 From the Sites table on the right side, select various sites, and each time notice that the circle color around the selected site changes to Grey color.

Step 4 Show the Circuits table (which is not in the default table tabs) using the **Show/hide tables** icon ( ). For details, see Show or hide tables or columns, on page 44.

Step 5 In the Circuits table, select the circuits of your choice. Notice that the selected circuit rows are highlighted as are the corresponding circuits in the network plot. Conversely, if you select the circuits in the network plot, the corresponding rows are highlighted in the table.

If 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 the **Cross table filter** icon ( ). Then, choose **Filter to interfaces**. The Interfaces table shows the four interfaces.

Step 6 In all the tables, check the selection check box in the header row to select all the objects. The selected objects are highlighted in the network plot. Uncheck this check box to deselect all the objects.



Show or hide tables or columns

Table 1: Show or Hide Tables or Columns

To...	Do This...
Show one or more tables	1. In the Network Summary panel on the right side, click the **Show/hide tables** icon ( ). 2. 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. 3. Click **Apply**.
Show or hide one or more columns	1. In the Network Summary panel on the right side, select the required Object tab. 2. Click the **Show/hide table columns** icon ( ). 3. 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.

Search and filter in tables

Use the following filter options to search and filter the values in tables:

- **Floating filter**—Click the icon 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 appear above the table.

- **Advanced filter**—Use the icon to filter the table to a list of specific search values:

  1. Choose from the drop-down lists and enter the value.
  2. If you are filtering by more than one criterion, you can optionally use these controls:
     - All (AND)—(Default) Filters to only those rows that match all the specified criteria (default).

- Any (OR)—Filters to rows that match any one of the criteria.

- **Cross table filter**—Use the icon to filter the selection to associated objects. For example, you might want to filter a circuit to associated interfaces. To use this option, select the object from the Network Summary table, and choose **Filter** > **Filter to** *X*, where *X* is a name of a viable object.

Sort columns in tables

Click the column headings to toggle the columns to sort in ascending or descending order. Notice the upward and downward arrows in the column headings. The upward arrow indicates that the columns are sorted in the ascending order, while the downward arrow indicates the descending order.

Save table views

The network models contain multiple object tables. In Cisco Crosswork Planning, you can save the views of the most commonly used tables. When the saved view is applied, the network model displays only the tables that are saved in the view.

Procedure

Step 1 If the plan file is not open, open it using the steps mentioned in Open plan files, on page 22.

Step 2 Use to show or hide the tables as per your requirement. For details, see Show or hide tables or columns, on page 44.

Step 3 Click Save view > **Save view** in the top right corner.

The Save View As window appears.

Step 4 Enter the name and click **Save**.

The newly created view appears in the **Saved views** list.

Step 5 **Apply Views:**

a) Click the **Saved views** drop-down list. List of all the saved views are displayed.
b) Click the thumbnail of the view (in the card layout) or the view name (in the list layout) to apply the view on the network model.

Step 6 **Rename Views:**

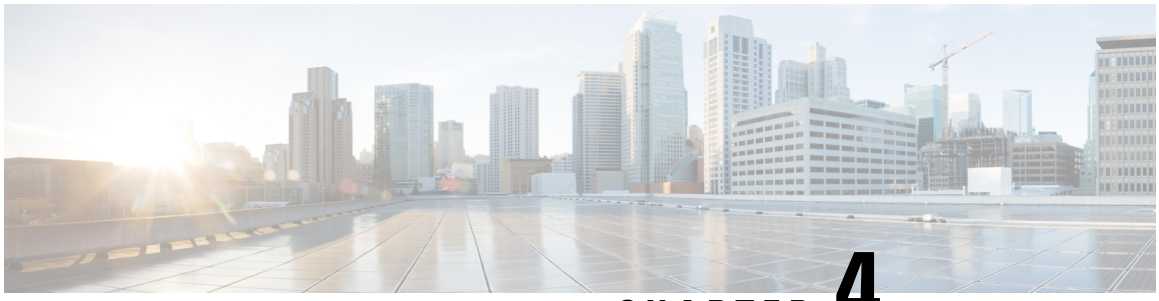
When a view is applied on a network model, you can rename it.

a) Ensure that the view is applied (see Step 5).
b) Click Save view > **Rename view**. The Rename View window appears.
c) Enter a new name for the view and click **Save**.

Step 7 **Delete Views:**

a) Click Save view > **Manage views**. List of all the saved views are displayed.
b) Click the icon in the view that you want to delete.

c) Click **Delete** in the confirmation window.



CHAPTER 4

Understand Plan Objects

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.

This section describes these basic objects and their relationships, as well as how to create, edit, and delete them.

- Nodes and Sites, on page 47
- Circuits and interfaces, on page 51
- SRLGs, on page 52
- Ports, port circuits, and LAGs, on page 53
- Key operations you can perform on objects, on page 55

Nodes and Sites

Both *node* and *site* are Cisco Crosswork Planning terminology.

- Node—A device in the network, which can be one of three types: physical, PSN (pseudonode), 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 pseudo-node (PSN) is typically used to represent 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 collection of nodes and/or other sites that potentially form 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.

- Nodes are shown as blue router icons ( ). The border color is an indication of 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 are shown as blue circles ( ). 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

Unless a site is empty, there is a hierarchy of the sites and nodes that a site contains. A site can be both a parent and a child site. If a site contains another site, it is a *parent* site. Any nodes, sites, or circuits within a parent site are *contained* (*nested*) objects. The contained nodes and sites are also called *children*. Often the child nodes and sites are geographically co-located. For example, a site might be a PoP where the routers reside.

- The site's **Parent site** property defines whether it is nested within another site. If it is empty, the site is not nested.
- 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. Similarly, this is true for each child site plot within it.
- Selecting a site from the network plot does not select the sites or nodes under it.

Delete sites

Selecting a site from the network plot does not select the sites or nodes under it except when deleting the site. In this case, all objects within a site are selected for deletion. However, in the confirmation that appears, you have the option to keep the contained sites and nodes. If you do, then the objects that are contained directly within it are moved to be on the same level as the site that is removed. The other, more deeply nested objects maintain their parent relationships.

PSN nodes

Cisco Crosswork Planning network models can contain nodes of with a **Type** property of "psn". These nodes represent pseudonodes (PSNs), which are used to model LANs or switches that connect more than two routers. They are used in two situations: for IGP modeling and BGP peer modeling.

In an IGP network, 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 whereby one of the routers on that LAN is the designated router (DR) for OSPF or the designated intermediate system (DIS) for IS-IS. The PSN node is named after 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.

Few points to consider when working with nodes that have **PSN** as a **Type** property are:

- Two PSNs cannot be connected by a circuit.
- If a PSN node is created by Cisco Crosswork Planning, “psn” is prepended 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. This is possible in manual demand creation, but 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.

Create nodes and sites

Create nodes

To create nodes, follow the steps in Create objects, on page 56, where *Object* is **Node**.

These are some of the frequently used fields and their descriptions.

- Name—Required unique name for the node.
- IP address—Often the loopback address used for the router ID.
- Site—Name of the site in which the node exists. If left empty, the node resides in the network plot. This offers a convenient way to create a site while 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—Name of the AS in which this node resides, which identifies its routing policy. This can be left empty if no BGP is being simulated.
- BGP ID—IP address that is used for BGP.
- Function—Identifies 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—Geographic location of the node within the network plot. These values are relevant when using geographic backgrounds.

Create sites

To create sites, follow the steps in Create objects, on page 56, where *Object* is **Site**.

These are some of the frequently used fields and their descriptions.

- Name—Unique name for the site.

- Display name—Site name that appears in the plot. If this field is empty, the Name entry is used.
- Parent site—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—Geographic location of the site within the network plot. These values are relevant when using geographic backgrounds.

Merge nodes

Real network topologies often have a number of nodes, typically edge nodes, connected to the network in the same way. For example, they might all be connected to the same core node or 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 that merging nodes changes the plan itself, 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 node merge effects are as follows:

- 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.

Follow these steps to merge nodes.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions > Initializers > Merge nodes**.

Step 3 Select the nodes that you want to merge in the Merge Nodes wizard. If you do not select any nodes, Cisco Crosswork Planning merges all nodes.

Step 4 Click **Next**.

Step 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.

Step 6 Click **Next**.

Step 7 Preview the list of effects of node merge. If these are acceptable, then click **Merge**.

Circuits and interfaces

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, see Ports, port circuits, and LAGs, on page 53.

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.

Create circuits and interfaces

To create circuits, follow the steps in Create objects, on page 56, where *Object* is **Circuit**. As a result of creating a circuit, two interfaces are also created.

Following are some of the frequently used fields:

- Capacity—The amount of total traffic this circuit can carry. The drop-down list has a selection of the most widely used capacities.
- SRLGs—If you want this circuit to belong to an SRLG, select it from this list or create a new one by clicking **Edit**. See Create SRLGs for circuits only, on page 52.
- Parallel group name—To include this circuit in a new or existing parallel grouping, enter its name.
- Interface A and B—You must specify two interfaces that are connected by the circuit.

Merge circuits

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

Merging circuits changes the network model itself, not just the visual representation.

The effects of circuit merge are as follows:

- 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.

Follow these steps to merge circuits.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.
Step 2 From the toolbar, choose **Actions > Initializers > Merge circuits**.
Step 3 Select the circuits that have the same source and destination endpoints.
Step 4 Preview the list of effects of circuit merge. If these are acceptable, then click **Submit**.

SRLGs

An SRLG is a group of objects that might all fail due to a common cause. For example, an SRLG could contain all the circuits whose interfaces belong to a common line card.

Create SRLGs

Follow these steps to create SRLGs.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.
Step 2 Create SRLGs by following the steps in Create objects, on page 56, where *Object* is **SRLG**.
Step 3 Enter the SRLG name.
Step 4 From the **Object type** drop-down list, choose the type of object that you want to include in the SRLG.
Step 5 For each object you want to include in the SRLG, check the check box under the **Included** column.
Step 6 Click **Add**.

Create SRLGs for circuits only

Follow these steps to create SRLGs for circuits.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 Open the Add/Edit Circuits window using any of these methods.

- If you are creating a new circuit, follow the steps in Create objects, on page 56, where *Object* is **Circuit**.
- If you are creating SRLGs for selected circuits, select one or more circuits from the **Circuits** tab. Then, click .

Step 3 Click the **Edit** button associated with the **SRLGs** field.



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

Associate Circuits with Existing SRLGs	Create New SRLG
a. For each SRLG in which you want to include the selected circuits, check the check box. **b.** Click **Save**.	**a.** Click . **b.** Enter the new SRLG name, and click **Save**.

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

Ports, port circuits, and LAGs

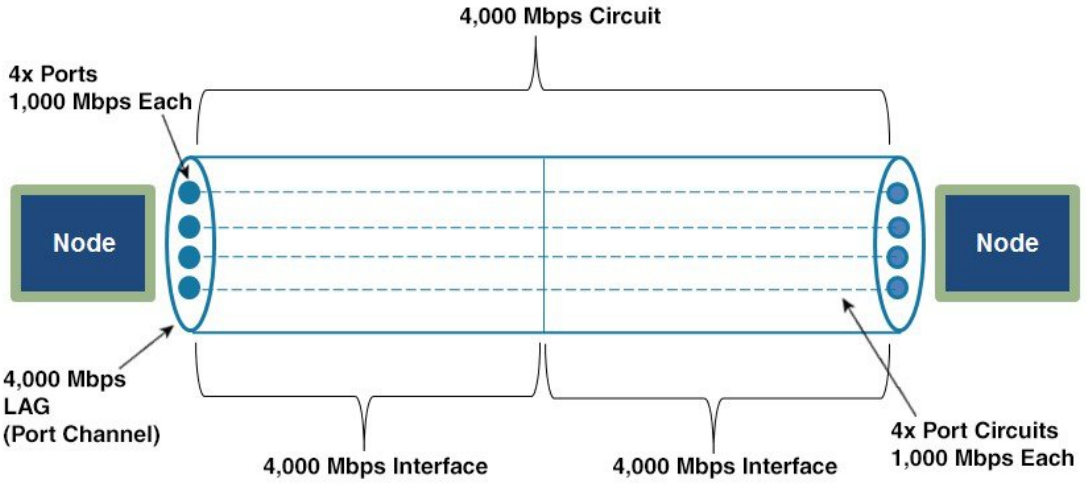
In Cisco Crosswork Planning, a port is a physical interface. You can model link aggregation groups (LAGs) and port channels using Cisco Crosswork Planning port and port circuits.

A LAG is a group of physical ports that are bundled into a single logical interface. A LAG is also known as *bundling* or *trunking*.

By default, each logical interface listed in the Interfaces table corresponds to a single physical port, and 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.

A port circuit is a connection between two ports. However, ports are not required to be connected to other ports by port circuits.

Figure 14: Ports, port circuits, and LAGs



Create ports

To create ports, follow the steps in Create objects, on page 56, where *Object* is **Port**.

These are some of the frequently used fields and their descriptions.

- Name—Required 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 amount of total traffic this port can carry. The drop-down list has a selection of the most widely used capacities.

Create port circuits

A port circuit specifies a pair of connected ports.

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

  To create ports and map them to their interfaces, see Create ports, on page 54.

  To create port circuits, follow the steps in Create objects, on page 56, where *Object* is **Port circuit**.

  These are the required fields and their descriptions.

  - Site and Node—Site and node on which the port exists.
  - Port—Name of the port.

- Capacity—The amount of total traffic this port circuit can carry. The drop-down list has a selection of the most widely used capacities.

Create LAGs

Follow 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 (see Create ports, on page 54).

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, from the **Ports** table, select all ports that belong to the LAG. You can quickly do this by filtering to their common interface.

Step 3 Select one of the ports and click  to open the Edit window.

Step 4 Select the interface that contains these ports.

Step 5 Click **Save**.

Set LAG simulation properties

Each interface is considered to be a LAG (port channel). You can configure LAG properties so that if it loses too much capacity due to non-operational ports, the entire LAG is taken down.

Follow these steps to configure LAG properties.

Procedure

Step 1 Open the Edit window of an interface or circuit (see Edit objects, on page 56).

Step 2 Click the **Advanced** tab.

Step 3 In the **Port channel** area, set one or both of the following parameters for one or both interfaces on the circuit:

a) In the **Min number of ports** field, enter the minimum number of ports that must be active and operating for the LAG circuit to be up.

b) In the **Min capacity** field, enter the minimum capacity that must be available for the LAG circuit to be up.

Step 4 Click **Save**.

Key operations you can perform on objects

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

Create objects

Follow these steps to add objects to the network model.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 Create objects in either of the following ways:

- From the toolbar, choose **Actions** > **Insert** > *Object*.
- In the Network Summary panel on the right side, 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.

  For reference, see Work with tables and object selections, on page 43.

Step 3 Enter the required details. The properties differ for each object.

Step 4 Click **Add**.

Edit objects

Most objects in the plot and in the associated tables have a set of properties that you can manage using the Edit window. These are the properties that Cisco Crosswork Planning uses to define and simulate an object.

Follow these steps to edit objects in the network model.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, navigate to the desired *Object* tab, and edit the object properties.

Step 3 **To edit a single object**:

Use either of the following ways:

- Select the required object and click the **Edit** icon ().

  OR

- In the **Actions** column, click ••• > **Edit** for the object you want to edit the properties.

Step 4 **To edit multiple objects (bulk edit)**:

Select the required objects and click the **Edit** icon ().

Note
In case of a bulk edit operation, you can select up to 50 objects at a time.

Step 5 Edit the properties, as required.

Step 6 Click **Save**.

Delete objects

Follow these steps to delete objects from the network model.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, navigate to the desired *Object* tab, and delete the objects.

Step 3 **To delete a single object**:

Use either of the following ways:

- Select the required object and click the **Delete** icon ( ).

  OR

- In the **Actions** column, click ••• > **Delete** for the object you want to delete.

Step 4 **To delete multiple objects (bulk delete)**:

Select the required objects and click the **Delete** icon ( ).

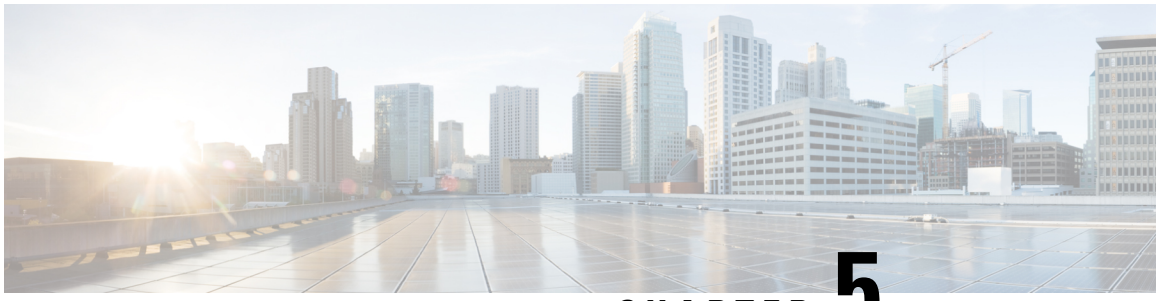
The object deletion successful message appears.



PART II

Simulate Your Network

- Simulation Overview, on page 61
- Simulate Traffic Flow from Source to Destination Using Demands, on page 71
- Perform What-If Analysis, on page 95
- Evaluate Impact of Worst-Case Failures, on page 103
- Evaluate Impact of Traffic Growth, on page 121
- Perform Capacity Planning, on page 127
- Simulate IGP Routing Protocol, on page 135
- Simulate BGP Routing, on page 143
- Simulate Quality of Service (QoS), on page 153
- Simulate VPN, on page 165
- Simulate Advanced Routing with External Endpoints, on page 177
- Simulate Multicast, on page 183



CHAPTER 5

Simulation Overview

Cisco Crosswork Planning network simulations calculate 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. A number of protocols and models are supported, including IGP, MPLS RSVP-TE, BGP, QoS, VPNs, and Multicast.

This chapter focuses on the general features of Cisco Crosswork Planning simulations. The individual protocols and models are described in their respective chapters.

This section contains the following topics:

- Use cases of network simulation, on page 61
- Auto-resimulation, on page 62
- States of plan objects, on page 63
- Simulated capacity, on page 68
- Simulated Delay, on page 69

Use cases of network simulation

Network simulation is a concept that allows for predictive analysis of network changes. It

- enables what-if analysis to predict the outcomes of changes in the network model
- supports capacity planning through simulations, and
- helps forecasting by projecting growth percentages on demands.

Examples

- **What-if analysis**—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?

  For details, refer to Perform What-If Analysis, on page 95.

- **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 the "worst-case" scenarios that highlights areas most at risk of congestion. You will also get a "failure impact" view, detailing the failures that cause the worst case. For details, refer to Evaluate Impact of Worst-Case Failures, on page 103.
- **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 details, refer to Evaluate Impact of Traffic Growth, on page 121.

Auto-resimulation

Auto-resimulation is a feature that automatically triggers resimulation when any changes affect routing in a network. Examples of such changes include:

- 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

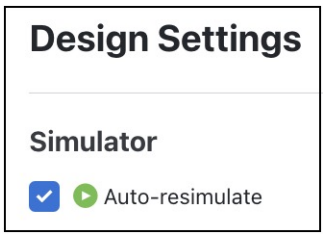
Auto-resimulation is disabled by default in newly opened plan files.

Enable auto-resimulation by default

Follow these steps to enable auto-resimulation by default.

1. Click Settings at the top right corner.

   Figure 15: Auto resimulation settings

2. Enable the **Auto-resimulate** check box.

   After updating this setting, it applies to the specific plan file. Each time you open this particular file, the same setting is used.

Run simulation manually

Additionally, 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 Simulate icon in the **Network Design** page.

States of plan objects

The state of a plan object is a condition that affects the simulation and determines whether an 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" means the object is in that state and "false" means it is not. Note that the "true" or "false" for Active state in the Interfaces table is reflective of the associated circuit. The plot shows graphical representations of these states with either a white cross or a down arrow 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

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 plot immediately displays where the traffic increases as a result (Figure 16: Failed circuit, on page 64)
- the **Util sim** column in the Interfaces table reflects the traffic changes, and
- demands are then rerouted around the failure (Figure 17: Reroute of demand around a failed circuit, on page 64).

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 16: Failed circuit

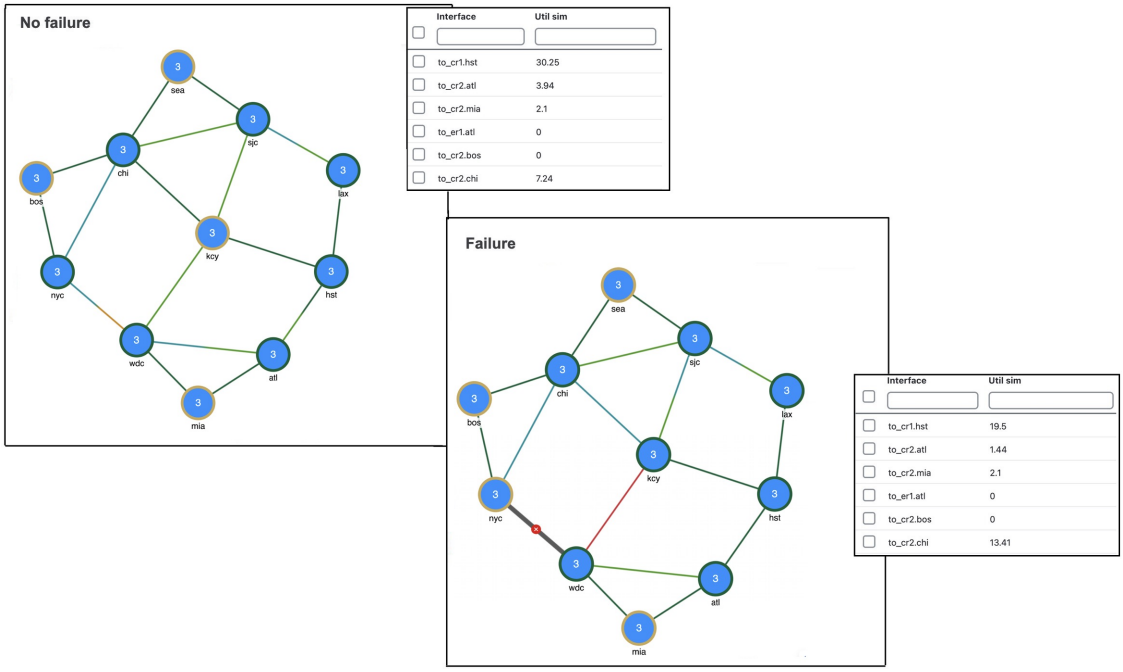
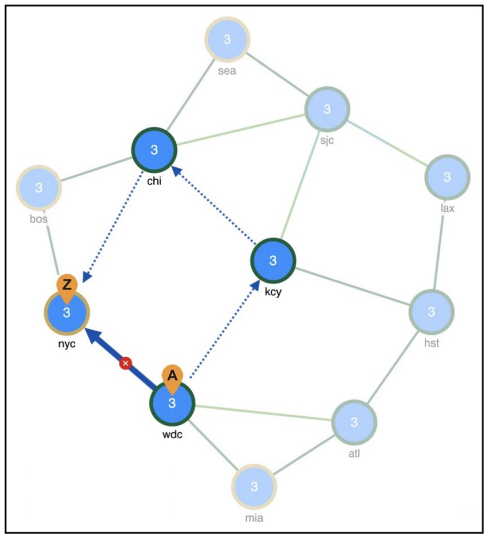


Figure 17: Reroute of demand around a failed circuit



Disabling demand rerouting

To specify that a demand should not reroute around failures, uncheck the **Reroutable** check box in the demand's Edit window. 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 and recover objects

Follow these steps to fail or recover objects.

Before you begin

Make a note of the list of objects that can fail by referring to the States of plan objects section.

Procedure

Step 1 Select an object from its respective table.

Step 2 Under the **Actions** column, click the ••• > **Fail** option.

In the network plot, notice a red circle with a white X on the object you failed (for example, see Figure 16: Failed circuit, on page 64).

Step 3 After failing an object, the menu option changes to **Recover**. Use this option to recover the failed objects.

Protect circuits from SRLG failures

You can protect circuits from being included in SRLG failures and SRLG worst-case analysis, though 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, see Optimize RSVP-TE Routing, on page 281.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 Set the **Protected** property for circuits.

a) In the Network Summary panel on the right side, 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) Check the **Protected** check box, which is an option in the **State** field.
d) Click **Save**.

Step 3 Set the **Network options** property for protecting circuits included in SRLGs.

a) Click **Network options** 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) Click **Save**.

Active state

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

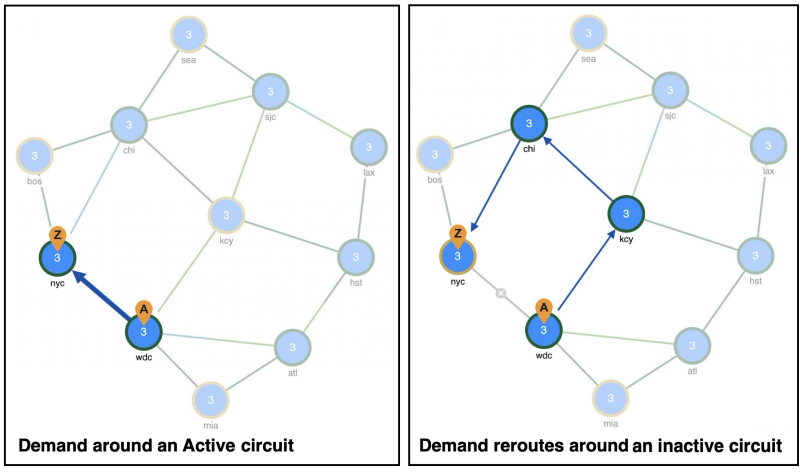
An object can be inactive because:

- 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 in the original plan that was discovered.

You can simultaneously change the active state of one or more objects. If you change the active state of an interface, you are actually changing 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 18: Inactive circuit, on page 66).

Figure 18: 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

Follow 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

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, select one or more like objects from their respective tables.

Step 3 Click .

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

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

A check mark means the object is active; uncheck the box to make it inactive.

Step 5 Click **Save**.

Operational state

The operational state identifies whether the object is functioning. You cannot set an operational state; rather, 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 mirrors the state of those objects.

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

Simulated capacity refers to the calculated capacity of an object, taking into account network state, including any failures that may reduce capacity. This 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

When referencing the Capacity sim value, there are a few rules to note regarding its 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.

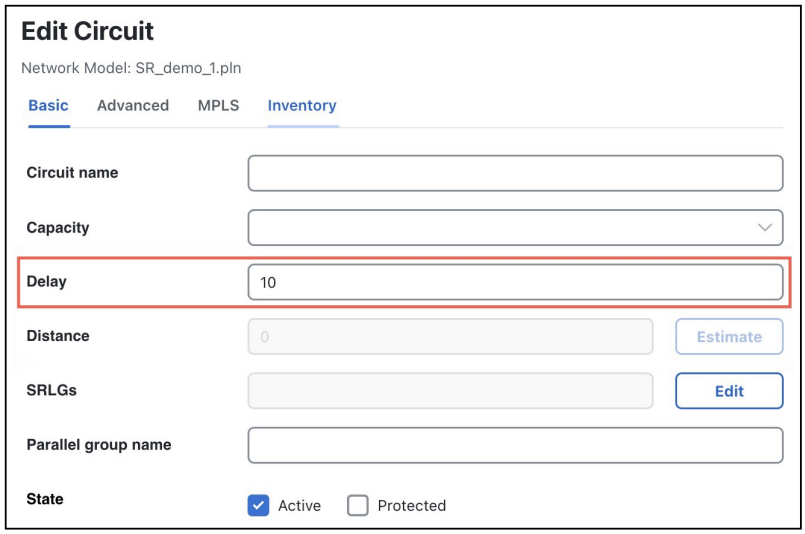
Note If a pair of ports is considered in Capacity sim calculations, both must be operational to be considered.

Simulated Delay

Simulated delay is the derived latency value, specifically within the context of L3 circuit delay calculations. This is indicated by the **Delay sim** column.

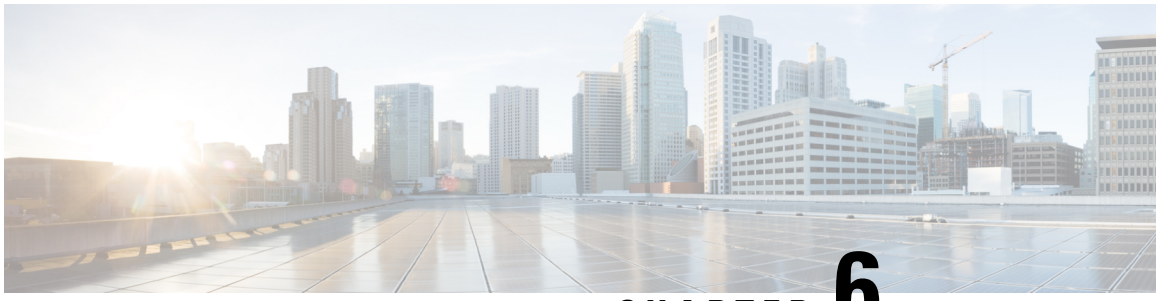
Delay is a property that can be set in the Circuit's Edit window.

Figure 19: Edit Circuit window



All Cisco Crosswork Planning delay calculations that use the L3 circuit delay, such as Metric optimization, use the Delay sim value.

Columns in the Circuits table	Description
Delay	One-way transmission latency over the circuit in milliseconds (ms).
Delay sim	(Derived) Entering a circuit Delay value copies it to the Delay sim column.



CHAPTER 6

Simulate Traffic Flow from Source to Destination Using Demands

Cisco Crosswork Planning uses *demands* to describe the source and destination of a potential traffic flow across a network. 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. 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.

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 the demands is essential for accurate traffic trending and traffic growth predictions.

This chapter describes demands, including how they are created and how their traffic can be estimated.

This section contains the following topics:

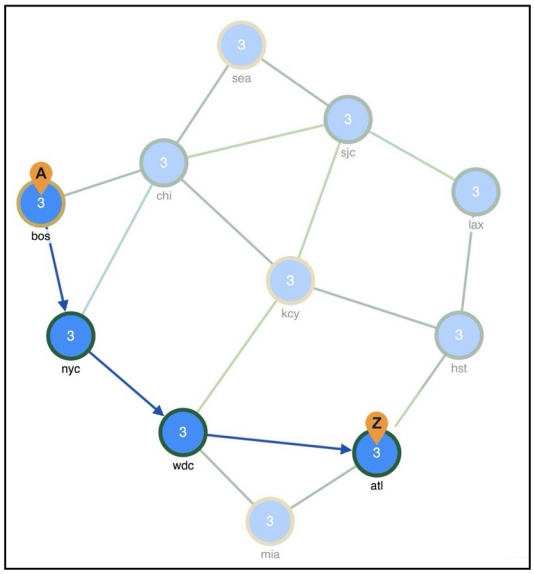
- Demands, on page 71
- How demands can be used?, on page 74
- Create demands and demand meshes, on page 75
- Visualize demands, on page 78
- Demand traffic, on page 83

Demands

A demand in Cisco Crosswork Planning describes the source and destination of a potential traffic flow across a network. Since demands determine how traffic is routed through the simulated model, creating realistic demands and demand meshes (see Demand meshes, on page 74) 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.

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

Figure 20: Demand route



Each demand is comprised of unique properties (keys) that define it, other properties, and traffic. The following list summarizes these. For a complete list of properties, refer to the available columns in the Demands table.

Unique Properties (Keys)	Each demand is defined by a unique combination of these four properties:
	Name—By default, this is blank.
	Source—Nodes, interfaces, external ASes, or external endpoints.
	Destination—Nodes, interfaces, external ASes, external endpoints, or multicast destinations.

Commonly Used Properties	Service class—User-defined classification of traffic, such as for voice or video.
	Latency bound—Policy that sets the maximum permissible latency on a demand under normal operation. This property is used by Cisco Crosswork Planning traffic engineering tools.
	Topology—Demands can be assigned to a specific IGP, and only route through interfaces that belong to that IGP.
	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. 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.
	Active—Only active demands are routed during simulations.
	Reroutable—Enable/disable the routing of demands around failures. Turning off reroutes around failures might be useful.
	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, so you must add the simulated traffic to them. Demand traffic belongs to the service class of the demand.

Demand sources and destinations

When creating sources and destinations, follow these recommendations:

- For internal routing, use nodes.
- For external ASes, use a combination of ASes, nodes, and interfaces. Using interfaces lets you specify the exact interface on which the demand traffic is going into or out of a node.
- For more complex routing where multiple sources or destinations (and multiple failover scenarios) are required, use external endpoints.
- For multicast routing, use multicast destinations.

If multiple interfaces are attached to a node and if a demand is sourced to or destined for that node, the 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, the source is the inbound interface. If using an interface as the destination of a demand, the destination is the outbound interface.

Demand meshes

Demand meshes are a time-efficient way of creating numerous demands for all or part of the 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

Each demand can have a latency bound, which is a policy that sets the maximum permissible latency on a demand under normal operation. These can then be used to guide the route selection of the traffic engineering tools. The Simulation analysis tool can use these values to determine if latency bounds are violated when worse-case failures occur.

The Demands table has several Latency columns. Key ones are as follows:

- 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.

How demands can be used?

You can use demands for the following purposes.

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 details, see Create demand meshes, on page 76. 2. Set the demand traffic manually or using the **Demand Deduction** tool. For details, see Modify demand traffic, on page 84 and Estimate demand traffic using demand deduction, on page 90.
Model future usage in the network	1. Create a demand mesh. For details, see Create demand meshes, on page 76. 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 details, see Evaluate Impact of Traffic Growth, on page 121.

Purpose	Suggested Steps to Take
Design networks	1. Create a demand mesh. For details, see Create demand meshes, on page 76. 2. Set the demand traffic using methods described in Simulate traffic using demands, on page 83.
Analyze existing plans	Use a variety of Cisco Crosswork Planning tools that rely on demand traffic.

Create demands and demand meshes

Create demands

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

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **Demands** > **Demand**.

OR

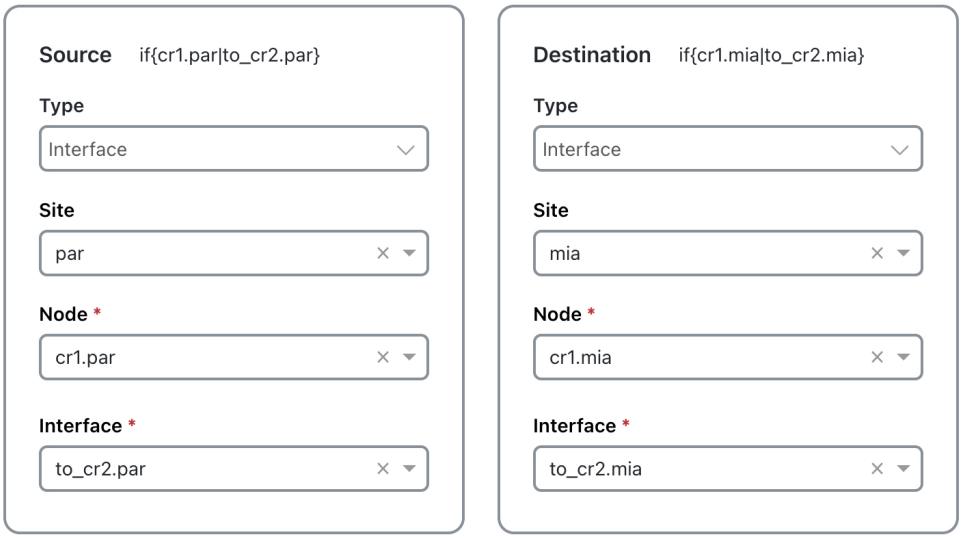
In the Network Summary panel on the right side, click  > **Demands** in the **Demands** tab.

Step 3 Enter a demand name.

Step 4 Specify the source as a node, interface, external AS, or external endpoint. Choose the other Source details, as required (see Figure 21: Source and destination options, on page 76).

Step 5 Specify the destination as a node, interface, external AS, external endpoint, or multicast destination. Choose the other Destination details, as required.

Figure 21: Source and destination options



Step 6 Choose the service class. If there are no service classes, create a new service class and select it. For instructions on creating a new service class, see Create service classes, on page 161.

Step 7 Enter a value for the latency bound.

Step 8 Choose a topology to restrict the demand routes only to interfaces or LSPs belonging to that topology. The default is unrestricted routing.

Step 9 Retain the **Active** default to include the demand in Cisco Crosswork Planning simulations, or uncheck **Active** to exclude this demand from simulations.

Step 10 For the default traffic level, click the **Edit** button to specify the amount of traffic or leave it empty for the Demand deduction tool to complete.

Step 11 Click the **Edit** button to specify the amount of growth rate you want to use for forecasting purposes. For more information, see Evaluate Impact of Traffic Growth, on page 121.

Step 12 Click **Add**.

Create demand meshes

Follow these steps to create demand meshes.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **Demands** > **Demand mesh**.

OR

In the Network Summary panel on the right side, click the **Demands** tab, and then click > **Demand mesh**.

Step 3 In the **Demand mesh details** panel:

a) Enter a demand name. The default is to have no name to prevent large numbers of demands using the same name from being created. The names are useful if needing to identify a specific area of the network, such as a VPN. However, not using demand names helps ensure you do not create a large number of demands that all have the same name.
b) Choose a service class.
c) Choose a topology.

Step 4 In the **Source** panel:

Select one or more sources from the Source check boxes. Options include using the Nodes, External ASes, and External Endpoints. By default, all the options are selected. Also, all available nodes, external ASes, and external endpoints are selected. Use the **Edit** button next to each option to select only the required nodes, external ASes, or external endpoints.

Step 5 In the **Destination** panel:

a) If you want to create demands to destinations other than what has been selected as the source, check the **Specify separate set of destinations** check box and choose the other required details.
b) Uncheck the **Also create demands from destination to source** check box if you want the demands created in only one direction. This applies only if you have selected a different set of destinations.

Step 6 For any of the following options, 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, see Simulate BGP Routing, on page 143.
- 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 information on AS relationships and routing policies, see Simulate BGP Routing, on page 143.
- 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).

Step 7 Click **Save**.

Create demands for LSPs

To create demands for LSPs, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.
Step 2 In the Network Summary panel on the right side, click the **LSPs** tab.
Step 3 Click  and choose **Demands for LSPs**.
Step 4 Choose the LSPs over which you want to run the demands.
Step 5 Choose the service class for the resulting demands.

Step 6 Choose 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.

Step 7 To remove the restriction of setting these demands to only these LSPs, uncheck **Mark LSPs as private**. Otherwise, the default is to restrict these LSPs so that they use only the resulting demands.

Step 8 Click **Submit**.

Set demand latency bounds

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

To set the demand latency bounds, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, select one or more demands from the **Demands** table.

Step 3 Click  and choose **Demands**.

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

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

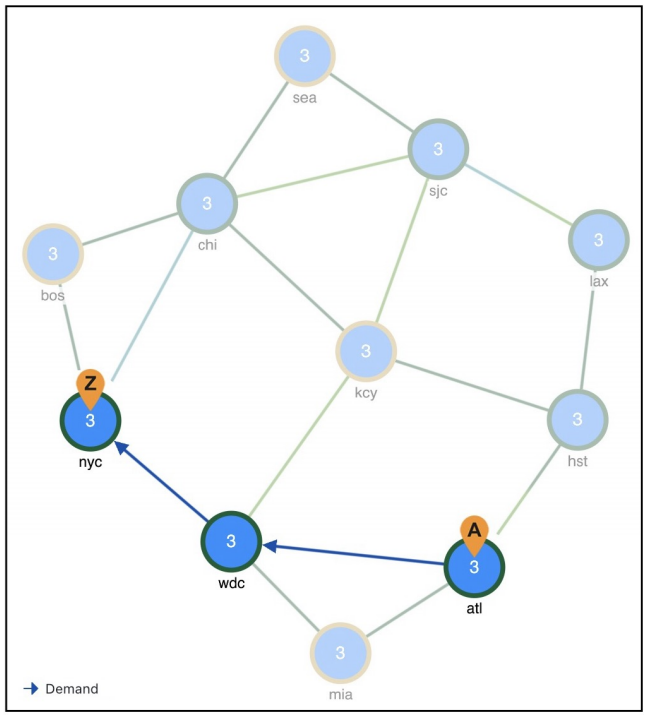
To delete a latency bound, delete the text in this field.

Step 5 Click **Save**.

Visualize demands

To view demand paths in the network plot, select them in the Demands table. They appear as blue arrows in the network plot. An “A” labels the source, and a “Z” labels the destination. If sites are nested, these “A” and “Z” labels appear in all relevant child sites. You can select up to three demands to view in the network plot. When you hover over any selected demand in the Network Summary table, it is highlighted in the network plot.

Figure 22: Demand route



View participating only nodes

Follow these steps to view only the nodes associated with the selected demand.

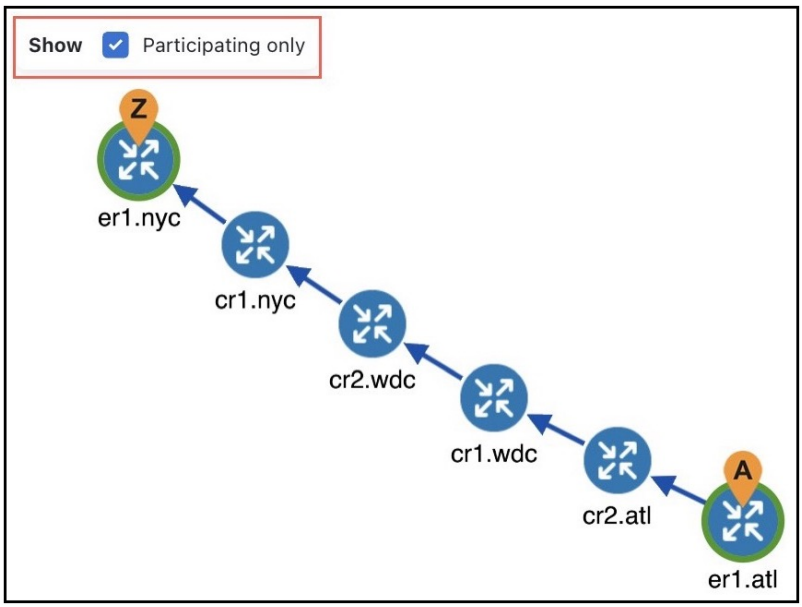
Procedure

Step 1 Open the plan file (see Open Plan Files). It opens in the **Network Design** page.

Step 2 Select the demand from the **Demands** table.

Step 3 Check the **Participating only** check box in the network plot.

Figure 23: Participating only nodes



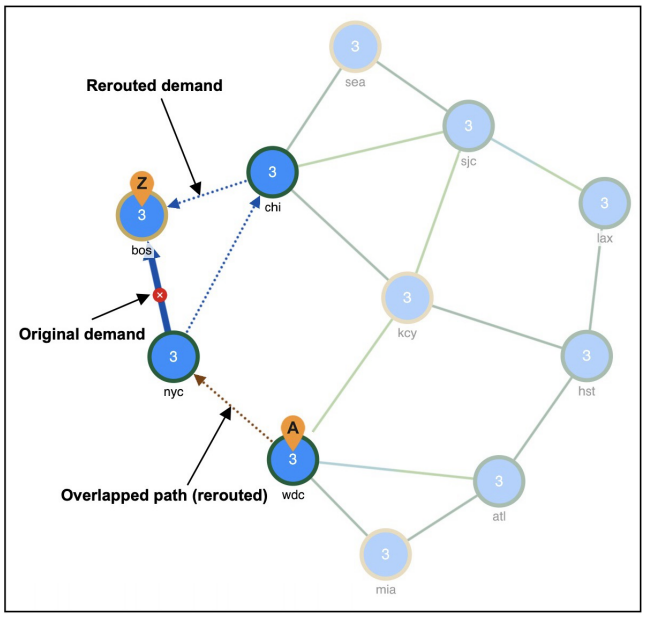
What to do next

To return to the previous view, uncheck the **Participating only** check box or check the **Show Groups** check box in the network plot.

Traffic rerouting

Demands are typically used to show how traffic reroutes around failures. A blue dotted line represents the rerouted demand. 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

For a description of shortest IGP paths and shortest latency paths, see Simulate IGP Routing Protocol, on page 135.

View shortest IGP paths

To view the shortest IGP path of a demand:

1. Click the **Map options** icon in the network plot ( ) and then check the **Shortest IGP paths** check box under the **Demands** tab.

   By default, **Simulated paths** option is selected, which displays the demand as a blue arrow.

2. Select the demand from the Demands table.

The shortest IGP path is displayed as a dark green arrow in the network plot. See Figure 25: Shortest IGP and latency paths, and overlapped path, on page 82.

View shortest latency paths

To view the shortest latency path of a demand:

1. Click the **Map options** icon in the network plot ( ) and then check the **Shortest latency path** check box under the **Demands** tab.

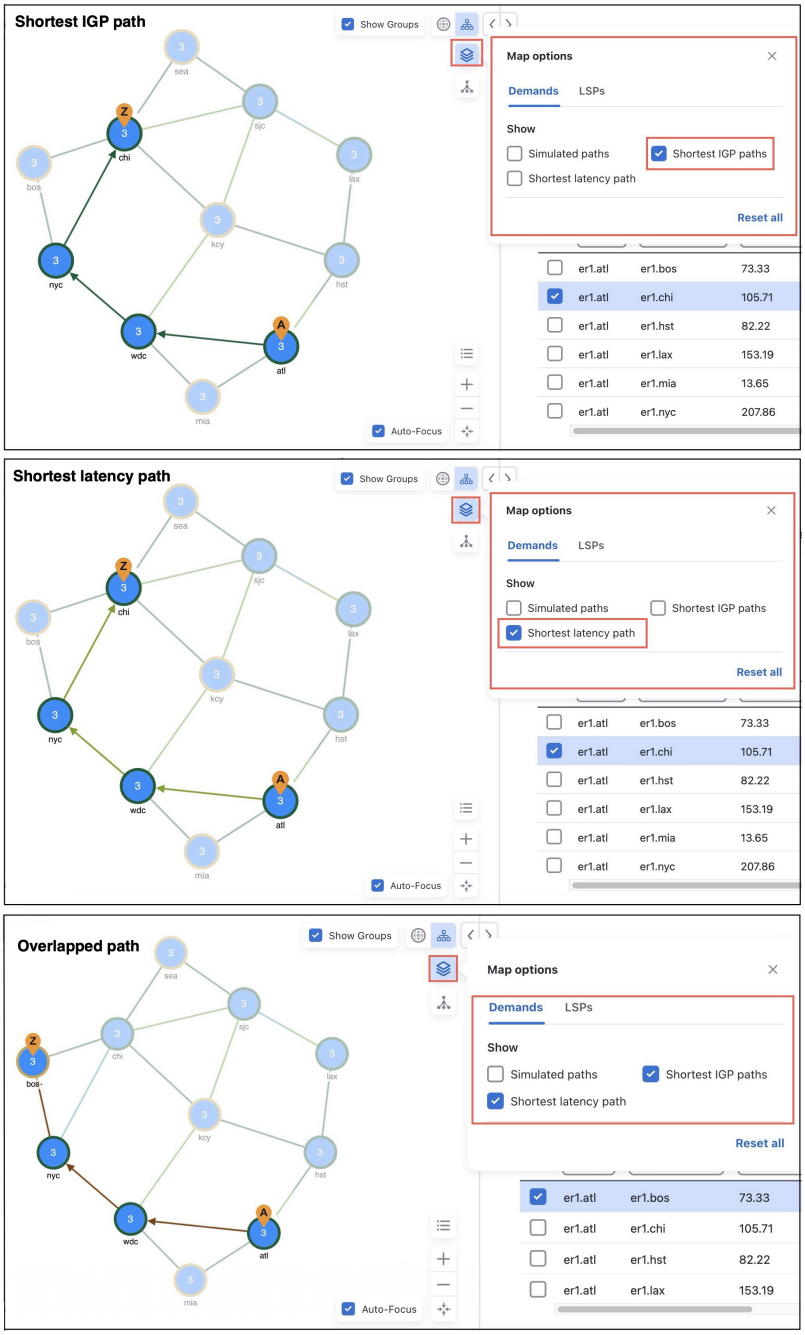
2. Select the demand from the Demands table.

The shortest latency path is displayed as a light green arrow in the network plot. See Figure 25: Shortest IGP and latency paths, and overlapped path, on page 82.

View overlapped paths

If you choose multiple options in the **Map options** icon and if the route overlaps with another route, then a brown arrow is displayed.

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



Demand traffic

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 most complex and powerful method of adding demand traffic is the Demand deduction tool, which estimates demand traffic from measured traffic values.

Simulate traffic using demands

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 traffic *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.

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.

Procedure

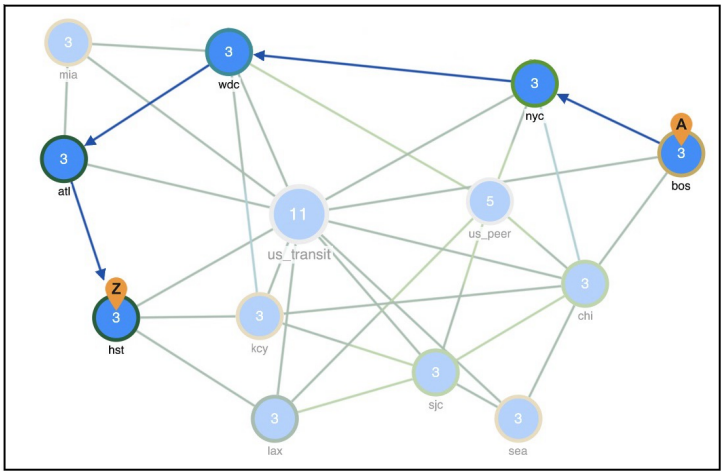
Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, click the **Demands** tab to show the Demands table.

Step 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



Step 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

Step 5 Click the **Service class** column heading to sort demands by service class.

Step 6 Read the **Traffic** column values to determine how much traffic each demand is attempting to route.

Step 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.

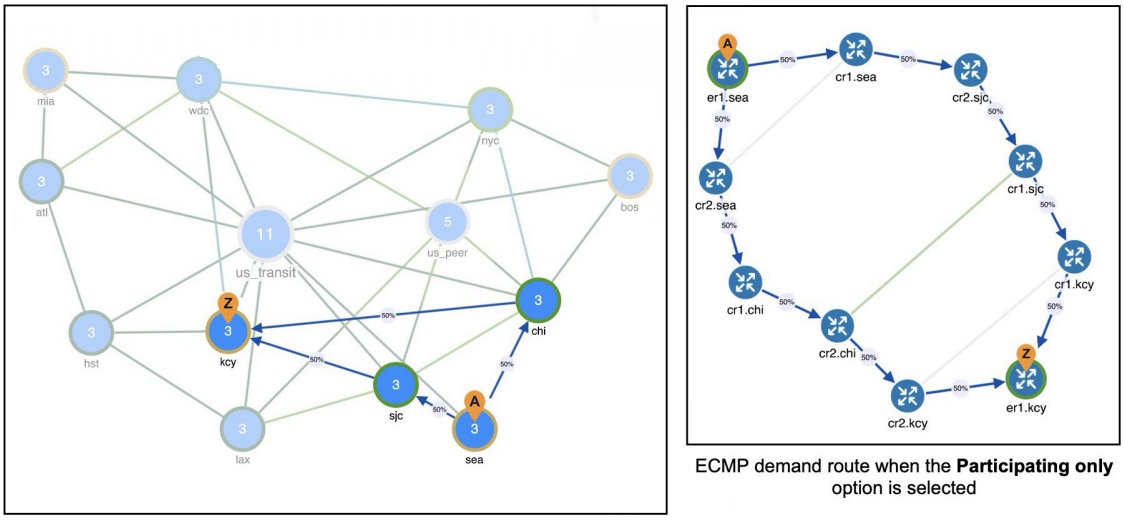
d) Notice that the sum of Delay sim values of all interfaces is equal to the maximum latency of the corresponding Demand. In this case, 23.

Clear the filter once you have completed the above steps.

Step 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

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 the following relative values.

- To set a fixed value, use the Edit Demands window (Modify fixed demand traffic, on page 85).
- To set a fixed or relative value, use the **Modify demand traffic** initializer (Modify fixed or relative demand traffic, on page 85).

Modify fixed demand traffic

To modify fixed demand traffic, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

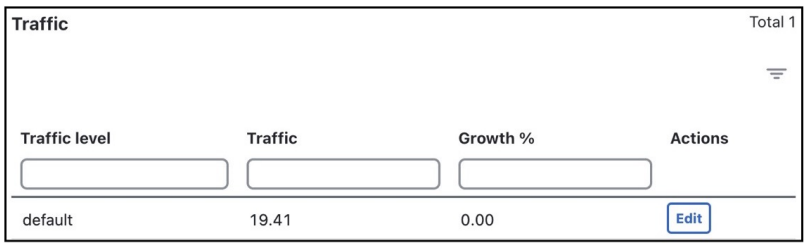
Step 2 In the Network Summary panel on the right side, select one or more demands from the **Demands** table.

Step 3 Click  and choose the **Demands** option.

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

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

Figure 28: Modify demand traffic



Step 5 Enter the desired amount of simulated traffic in the **Traffic** field and click **Save**.

Step 6 Click **Save** in the Edit Demand window.

Modify fixed or relative demand traffic

To modify the demand traffic to a fixed or relative value using the **Modify demand traffic** initializer, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, click **Actions > Initializers > Modify demand traffic** to open the Modify Demand Traffic page.

Step 3 Select the demands for which you want to modify the traffic. By default, all demands are selected. Deselect all and select the required demands.

Step 4 Click **Next**.

Step 5 Choose one of the options identified in Table 2: Modify demand traffic options, on page 86 and choose a relevant value.

Step 6 Click **Submit**.

Following options are the available in the Modify Demand Traffic page. Except for the percentage option, all values are in Mbps.

Table 2: Modify demand traffic options

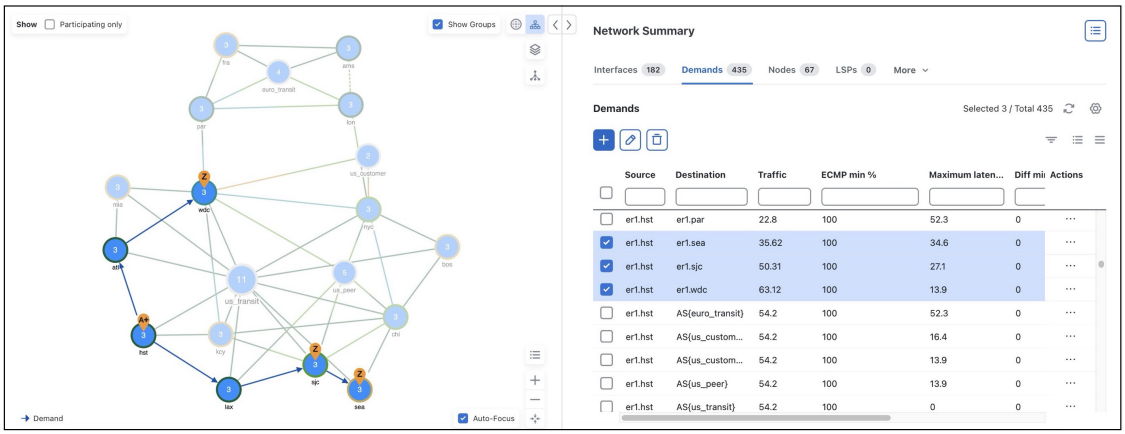
Option	Description
Change traffic by __ %	Change traffic by a specified percentage. Positive percentages add to the traffic, and negative percentages subtract. For example, if the traffic were 1000 Mbps and you entered –10, the traffic would be reduced to 900 Mbps.
Add	• Add __ Mbps in total, proportionally—Add 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. • 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. • 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. • 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.

Example: Modify demand traffic

In the following example, we increase the demand traffic of the selected demands by 50%.

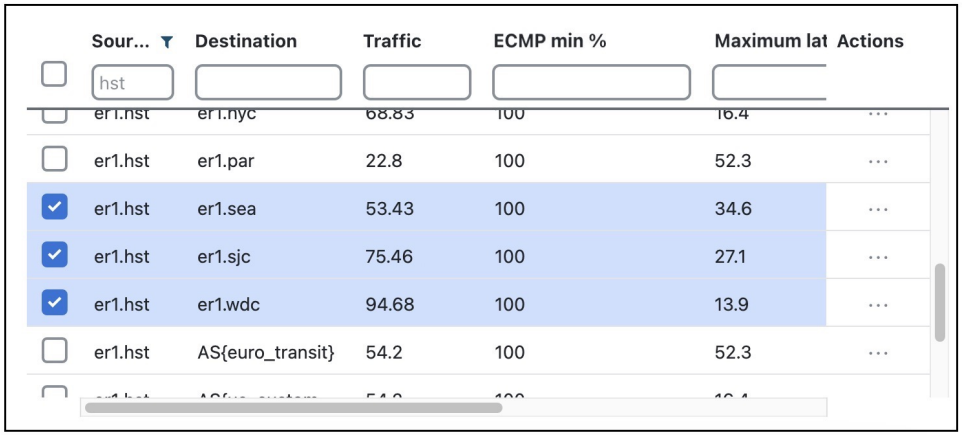
Note the values in the **Traffic** column for the demands. In this example, 35.62, 50.31, and 63.12 Mbps.

Figure 29: Modify demand traffic



To increase the demand traffic by 50%, enter **50** in the **Change traffic by ___ %** field. Note 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%



Understand demand deduction

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. For details, see Estimate demand traffic using demand deduction, on page 90.

The accuracy and usefulness of the results depend on many factors, including how much measured traffic is available, and of what type. 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, predicting the overall utilization after a failure, a topology change, or a metric change, can be very accurate 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 flows 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. See Specify external endpoint members, on page 178.

Differences in measured and simulated traffic

Demand deduction relies on accurate topologies, demand meshes, and traffic measurements. These can affect the results of the traffic simulated in the demands and cause simulated traffic to differ from the measured traffic, thus affecting the accuracy of Cisco Crosswork Planning simulations. 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.

If these columns show large values, one of the following 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—There are typically many more demands in a network than measurements, and 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.

  Additionally, 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 nondeterministic. 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 and 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.

Minimize differences between measured and simulated traffic

Demand deduction estimates demands that predict interface utilizations under incremental changes to the topology, for example failures, metric changes, or design changes, such as adding a new express route. If interface measurements alone are available, you might choose to fine-tune the Demand deduction calculations to get better results, such as for site-to-site traffic. To enhance the accuracy of Demand deduction results, consider the following 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.
- In the Demand Deduction window, always 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 the flow measurements are inconsistent and the interface measurements are very consistent, then interfaces should be set 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, in the Demand Deduction window, 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 in demand deduction



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.

Besides node, interface, and LSP traffic measurements, Demand deduction allows more general flow measurements to be used. These flow measurements can be flows from (or through) a specified node, to (or through) another node. Measurements can also be 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.

Flow measurements are entered in the plan file in the <Flows> table, and appear in the UI in the Flows table. Table 3: Flows table columns, on page 90 lists some of the more useful columns in the Flows table. Note that exactly which traffic is included is defined in the **From type** and **To type** columns.

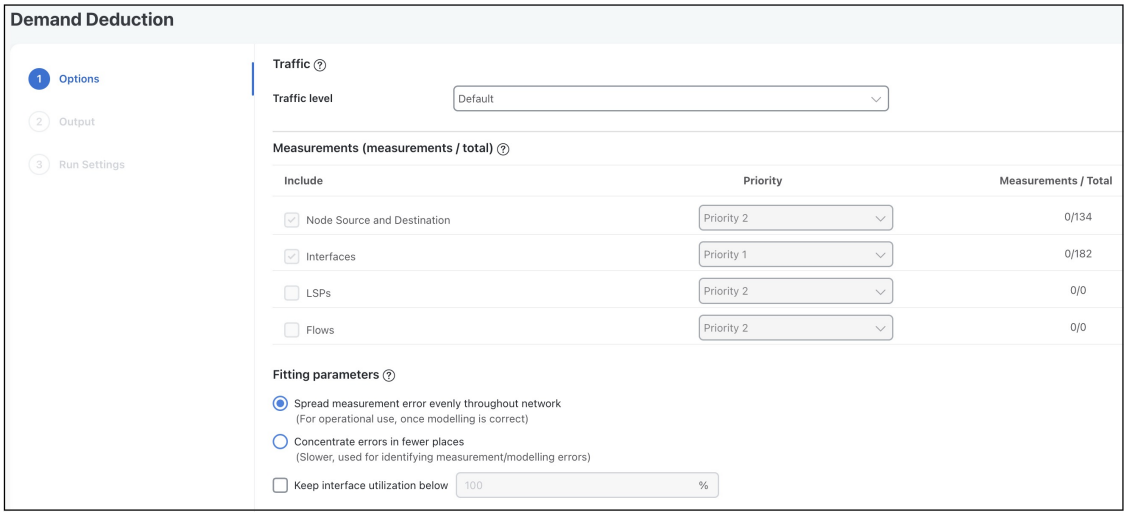
Table 3: Flows table columns

Column	Description
From	Specifies the source node.
From type	• Source—Traffic originating at the From node is included in the flow. • Interior—Traffic is included that passes through the From node, entering that node from another node in the same AS • 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. • Interior—Traffic that passes through the To node to another node in the same AS. • Border—Traffic that passes through the To node to another node in a different AS.
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

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, see Minimize differences between measured and simulated traffic, on page 89. For information on setting up external endpoint members to be included in Demand deduction calculations, see Simulate Advanced Routing with External Endpoints, on page 177.



To estimate the demand traffic using the Demand deduction tool, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions > Tools > Demand deduction**.

Step 3 (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.

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, see Evaluate Impact of Traffic Growth, on page 121.

Step 4 Identify one or more types of measurements used in the calculations: Nodes (source and destination), Interfaces, LSPs, and Flows.

Step 5 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.

Step 6 Choose the required **Fitting parameters**.

Step 7 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.

Step 8 Click **Next**.

Step 9 Choose 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.

Step 10 Determine whether to fix multicast demands. If selected, the multicast demands are fixed at their current traffic value.

Step 11 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.

Step 12 Click **Next**.

Step 13 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

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.

Step 14 Click **Submit**. The Demand deduction tool calculates the simulated traffic and lists the results in a Demand Deduction report.

Demand deduction example

This example demonstrates results when using the Demand Deduction tool on a simple network. Figure 31: Network containing two demands and no demand traffic, on page 92 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 because these demands do not yet contain traffic.

Figure 31: Network containing two demands and no demand traffic

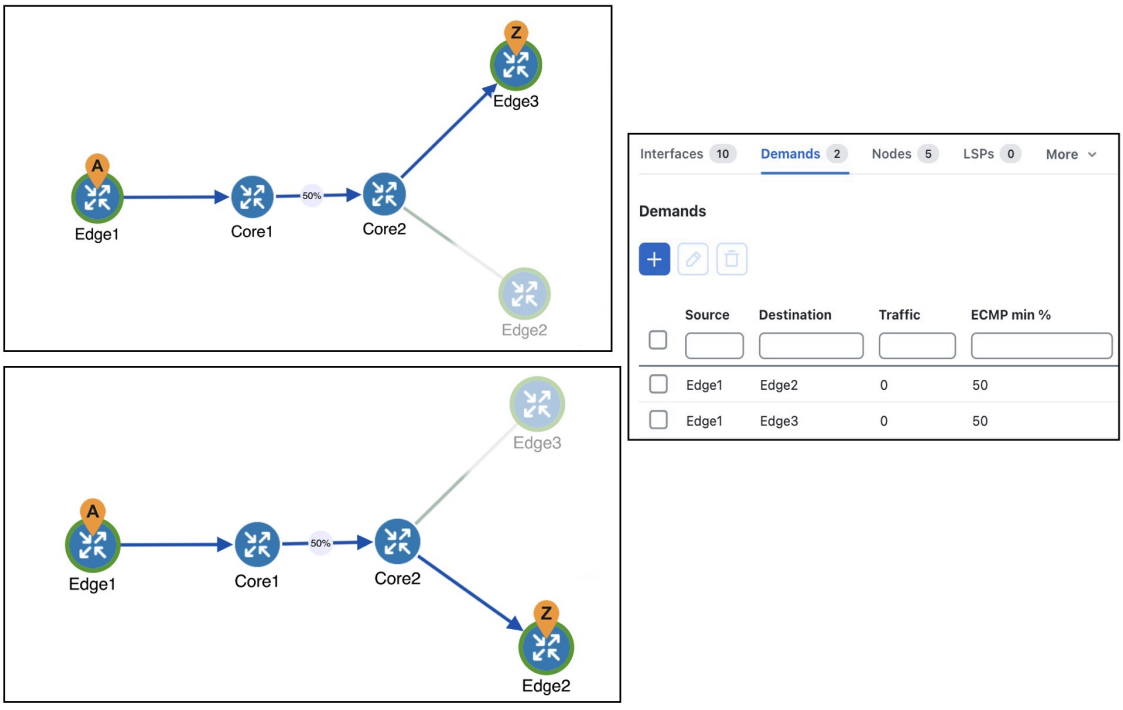
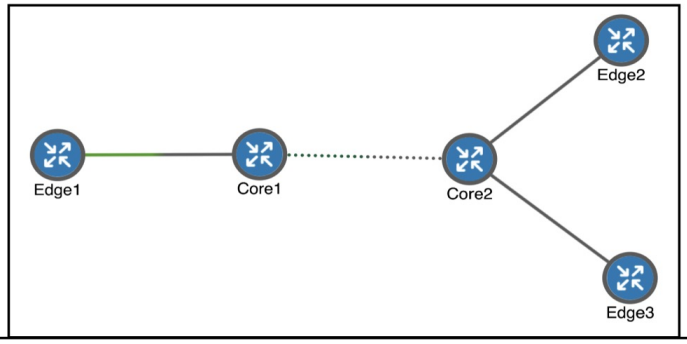


Figure 32: Measured traffic view and interfaces associated with the demands, on page 93 shows the Measured traffic view and the five interfaces associated with the two demands, three of which have measured traffic.

- Edge1 to Core1 has 470 Mbps of measured traffic.
- One Core1 to Core2 interface has 210 Mbps, while the other has 240 Mbps, for a total of 450 Mbps. This unequal split is due to 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



Interfaces 10 | Demands 2 | Nodes 5 | LSPs 0 | More

Interfaces

Node	Interface	IGP metric	Remote node	Traff sim	Traff meas	Capacity sim	Util sim	Util meas
Edge1	to_Core1	1	Core1	0	470	1000	0	47
Core1	to_Edge1	1	Edge1	0	NA	1000	0	NA
Core1	to_Core2	1	Core2	0	210	1000	0	21
Core2	to_Core1	1	Core1	0	NA	1000	0	NA
Core1	to_Core2[1]	1	Core2	0	240	1000	0	24

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 94). 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.

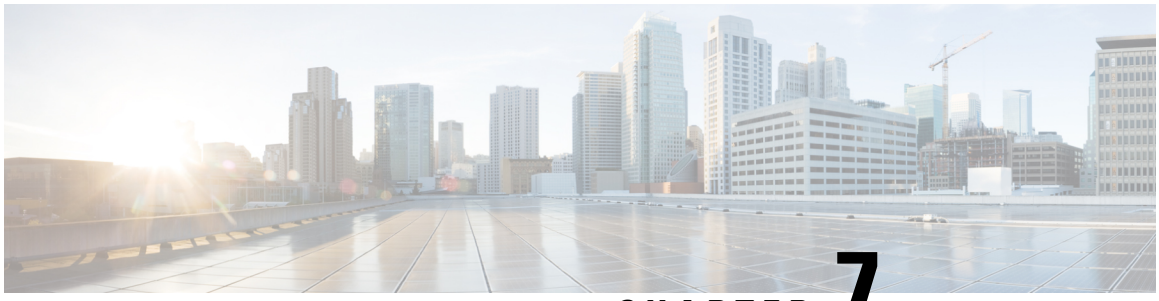
- Edge1 to Core1 has a difference of 10 Mbps, or 1%.
- One Core1 to Core2 has a difference of 20 Mbps, or 2%, while the other has a difference of 10 Mbps, or 1%.

- 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



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.



CHAPTER 7

Perform What-If Analysis

This section contains the following topics:

- Examine failure scenarios, on page 95
- Perform failure analysis, on page 95
- Perform impact analysis of topology changes, on page 98
- Perform impact analysis of metric changes, on page 100

Examine failure scenarios

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, rather than a huge upgrade, a planner may consider whether adding a single link between two sites can mitigate congestion over 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

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 utilizations are well under 100% under normal operation. But what happens under failure?

Cisco Crosswork Planning lets you perform failure analysis, which can be useful in long-term planning and short-term preparation. It simulates reroutings and changes in traffic utilization under individual failure events. For an example of how failure analysis is performed in Cisco Crosswork Planning, see Example: Failure analysis, on page 96.

Follow these steps to fail or recover objects.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 Select an object from its respective table. To see the complete list of objects that can be failed, see the list mentioned in States of plan objects.

Step 3 Choose the ••• > **Fail** option under the **Actions** column.

Once failed, the menu option changes to ••• > **Recover**. Use this option to recover the objects.

Step 4 Notice the difference in traffic utilization values and how the demand has been rerouted in the network plot.

Example: Failure analysis

In this example, we determine the effects on the network if a single interface fails.

Procedure

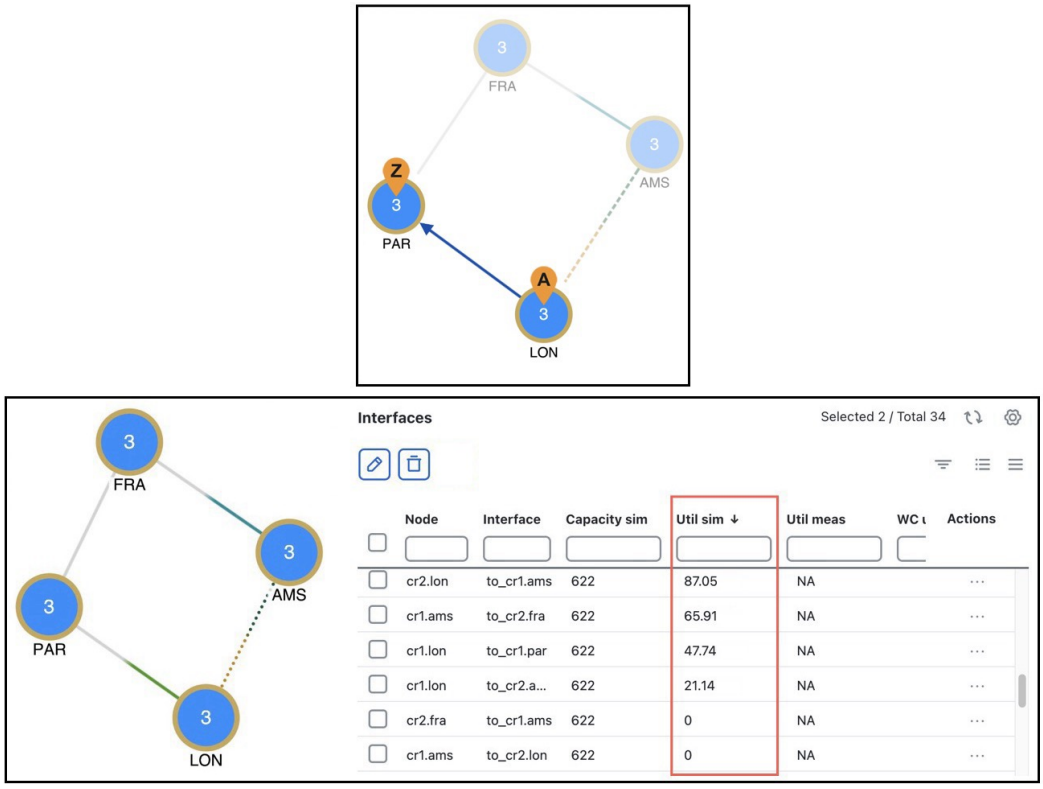
Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the network plot (see Figure 34: Before failing the interface, on page 97), notice 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.

Step 3 In the Interfaces table, select the cr1.lon-cr1.par interface (LON-PAR circuit) and choose ☰ > **Filter to demands** > **Through all interfaces**.

Step 4 Click the demand to see its route. The solid violet arrow indicates the demand path.

Figure 34: Before failing the interface



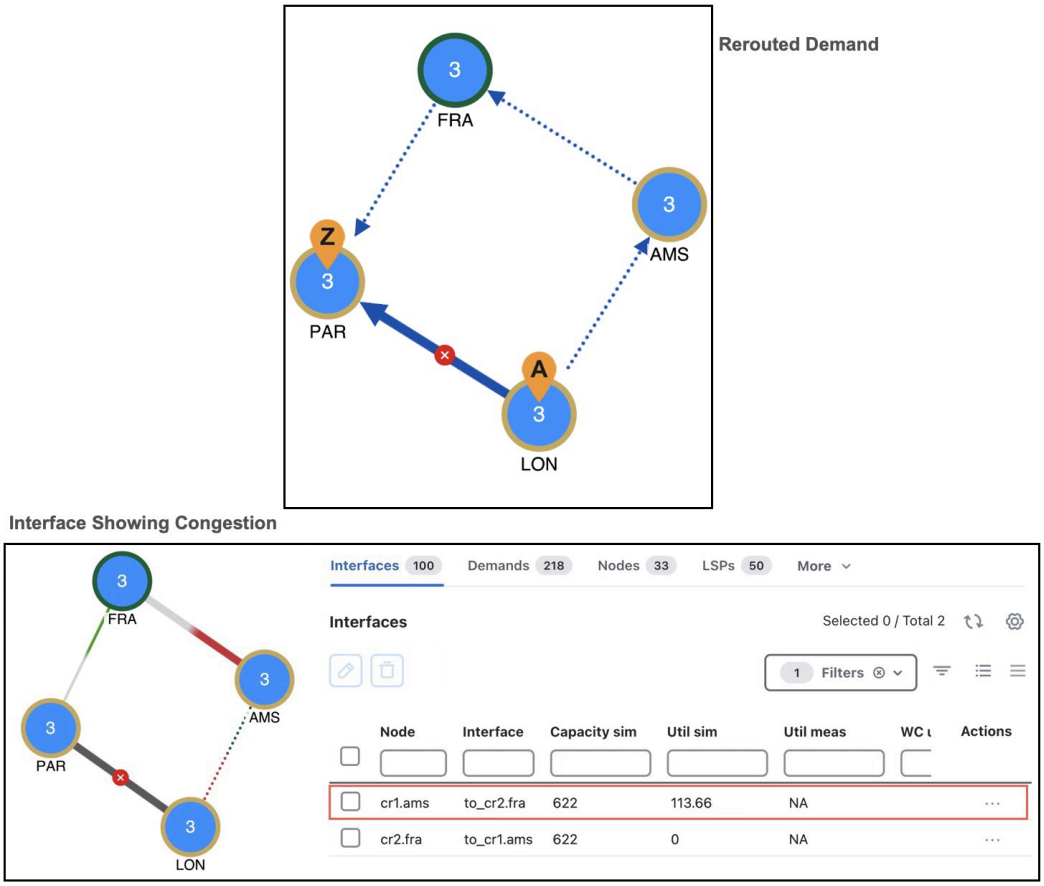
Step 5 In the interfaces table, select the cr1.lon-cr1.par interface and choose ••• > **Fail** option under the **Actions** column.

Step 6 Notice how the demand has been rerouted (see Figure 35: After failing the interface, on page 98). The dotted line shows the current demand path under failure.

The red cross on the link indicates a failure.

Step 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: After failing the interface



Step 8 To restore the failed circuit/interface, select it and choose ••• > **Recover** option under the **Actions** column.

Perform impact analysis of topology changes

A frequent question in the planning process is how to test the value of a minor change that may obviate the need for a larger one. For instance, rather than a huge upgrade, a planner may consider whether adding a single link between two sites can mitigate congestion over part of the network.

Being able to emulate and predict the impact of these changes promotes Service Level Agreement (SLA) adherence and staff efficiency. Making a topology change could impact traffic flows, congestion, and latency. Therefore, knowing the effect such changes would have is valuable for operators, planners, and designers alike, 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, see Example: Impact analysis of topology changes, on page 99.

Example: Impact analysis of topology changes

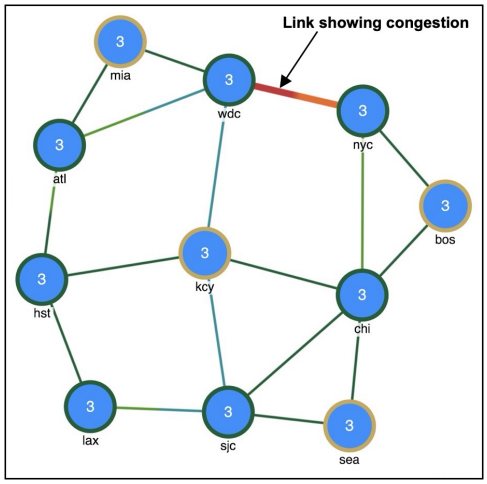
In this example, we see how adding a circuit between two sites ("wdc" and "nyc") can help in relieving the congestion.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 Notice that the link between the sites "wdc" and "nyc" is congested, as indicated by the red interface.

Figure 36: Before adding a circuit



Step 3 Add a circuit beween "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**.

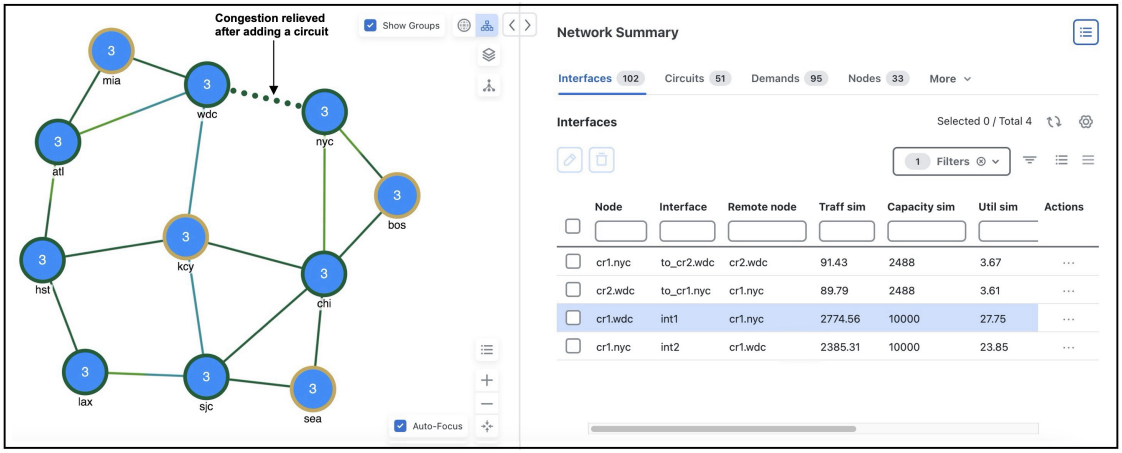
b) In the **Add Circuit** window, enter the following details:

- **Circuit name**—Enter the name of the circuit.
- **Capacity**—Enter the amount of total traffic this circuit can carry. The drop-down list has a selection of the most widely used capacities.
- **Interface A** and **Interface B** panels—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**.

Step 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: 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

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, see Example: Impact analysis of metric changes, on page 100.

To optimize the metric of multiple interfaces, use the **Metric optimization** and **Tactical metric optimization** tools. For details, see Optimize Metrics in the Network Core, on page 191.

Example: Impact analysis of metric changes

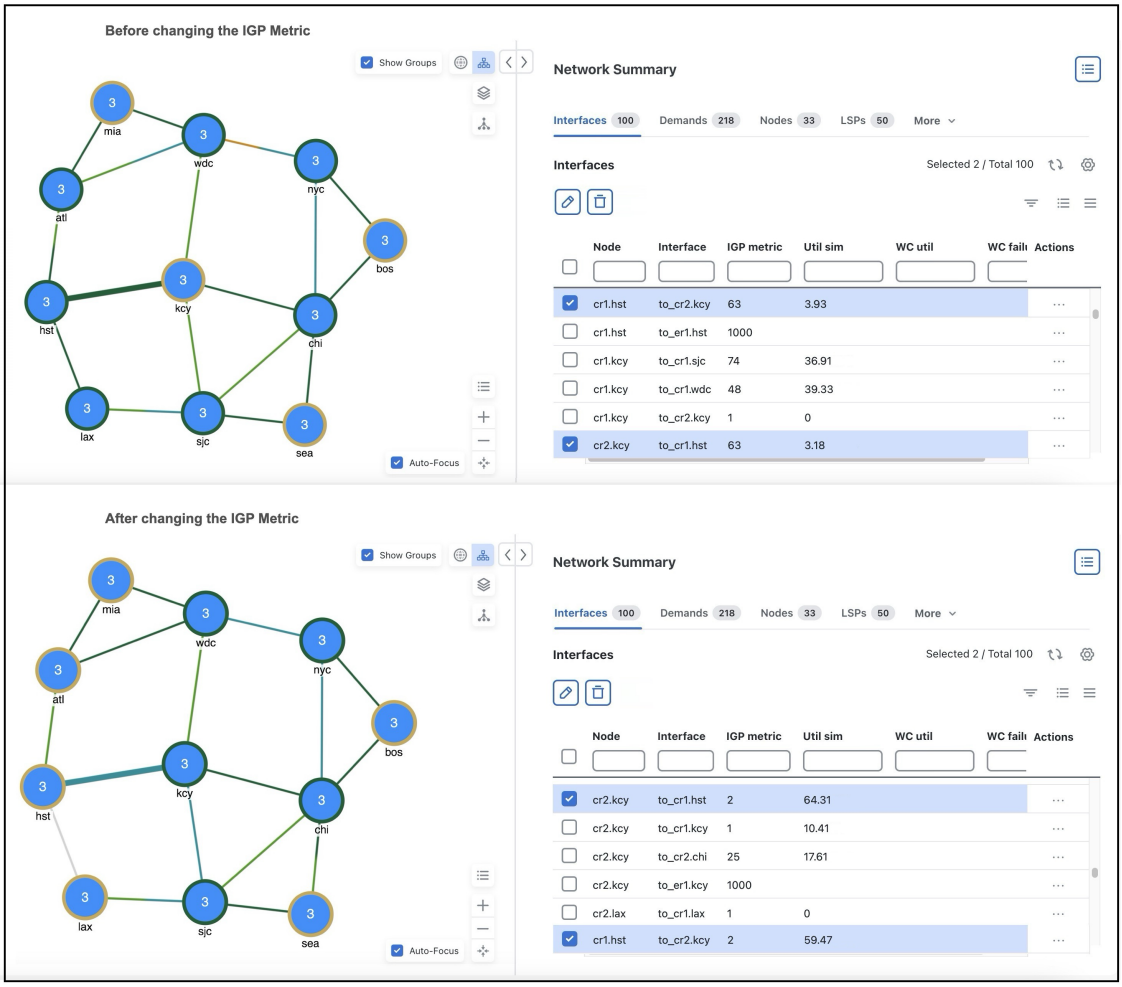
In this example, we see how changing the **IGP metric** value of the interfaces (between "kcy" and "hst" sites) changes the traffic utilization.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 Notice that the color of the link between the sites "kcy" and "hst" is dark green, which indicates that the traffic utilization is between 0-30%.

Figure 38: Impact of changing the IGP metric

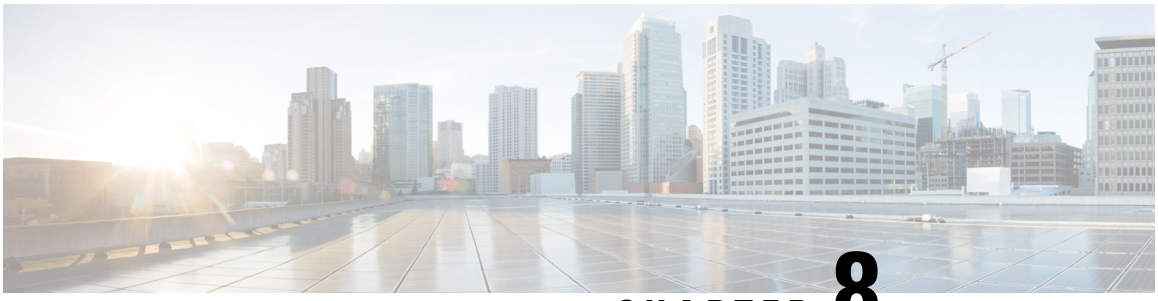


Step 3 Select this link. In the **Interfaces** table, the interfaces belonging to this circuit is selected.

Step 4 Click .

Step 5 In the **IGP metric** field, change the value from 63 to **2**.

Step 6 Notice the difference in the color of the link between "kcy" and "hst". The blue color indicates that the traffic utilization on these interfaces has increased.



CHAPTER 8

Evaluate Impact of Worst-Case Failures

The **Simulation analysis** tool combines the simulation results of a large set of failure scenarios. These results are useful for determining how vulnerable a network is to congestion and high latencies under failures, thus allowing you to plan sufficient 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 results in the following available analyses, which vary depending on whether the network has QoS parameters and depending on which options are selected when running the simulation.

- Worst-case traffic utilization, on page 104 on interfaces per service class
- (Optional) VPN worst-case utilization and latency
- (Optional) Worst-case demand latency, on page 107
- Failure impact view, on page 118, which analyzes the impact that each failed object has on interface utilizations throughout the network

Upon completion, a report window opens with a summary of each analysis, along with the list of simulations performed. Each time you run a simulation, this information is updated (replaced).

Simulation analysis can be performed 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.

This section contains the following topics:

- Worst-case traffic utilization, on page 104
- Worst-case demand latency, on page 107
- Run simulation analysis, on page 109
- Analyze simulation analysis reports, on page 116
- Failure impact view, on page 118
- Parallelization, on page 119

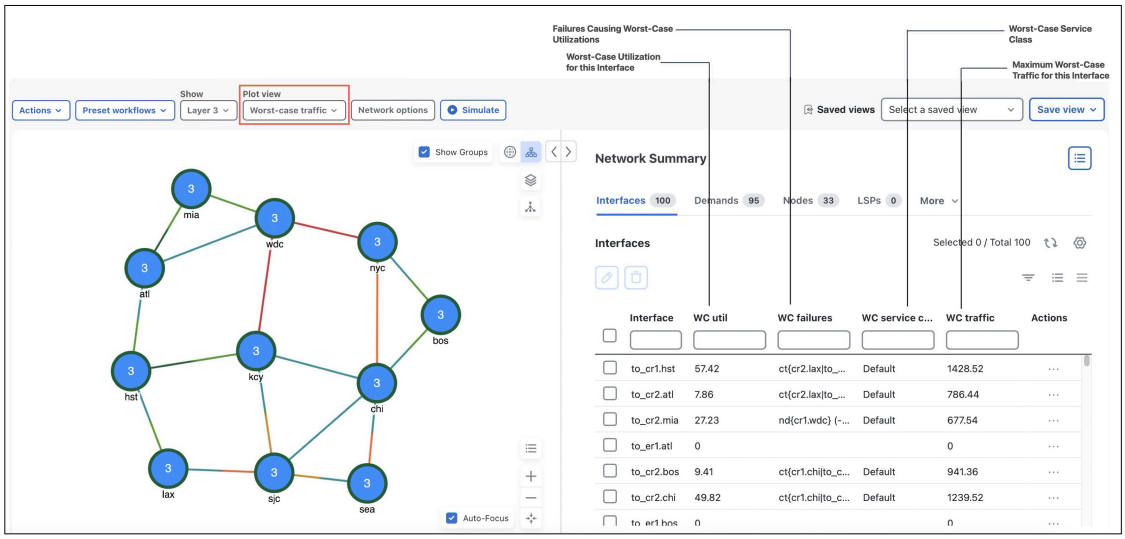
Worst-case traffic utilization

Worst case is the highest utilization that a particular interface experiences over all the failure sets and traffic levels that you selected. 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 details on how to run the Simulation analysis, see Run simulation analysis, on page 109.

Upon 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 for all interfaces

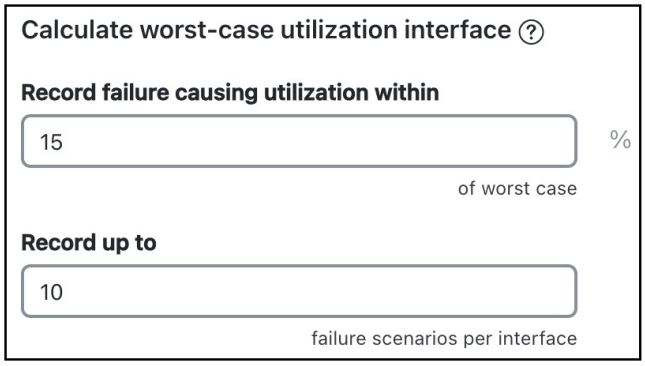


In addition to Plot view changes, the following columns are also updated in the Interfaces and Circuits tables upon finishing the Simulation analysis:

- 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. (See Worst-case QoS violations, on page 105.) 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, see Worst-case QoS violations, on page 105.

For information on worst-case calculations for VPNs, see Simulate VPN, on page 165.

Worst-case QoS violations

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.

The following columns are updated accordingly.

- 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.

  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**.

- WC service class—The service contributing to the worst-case QoS violation.

For more information on...	See...
• QoS parameters and QoS calculations • Set worst-case policies on service classes • Set interface queue parameters	Simulate Quality of Service (QoS), on page 153
Worst-case QoS calculations for VPNs	Simulate VPN, on page 165

Fail circuits to worst-case utilization

After running Simulation analysis (see Run simulation analysis, on page 109), you have the option to selectively view each failure scenario that causes the worst-case utilization or worst-case QoS violation for a single interface.

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 (see Figure 40: Failure of single circuit to its worst-case, on page 107).

Note If you choose an interface, you are actually failing its associated circuit to its worst case.

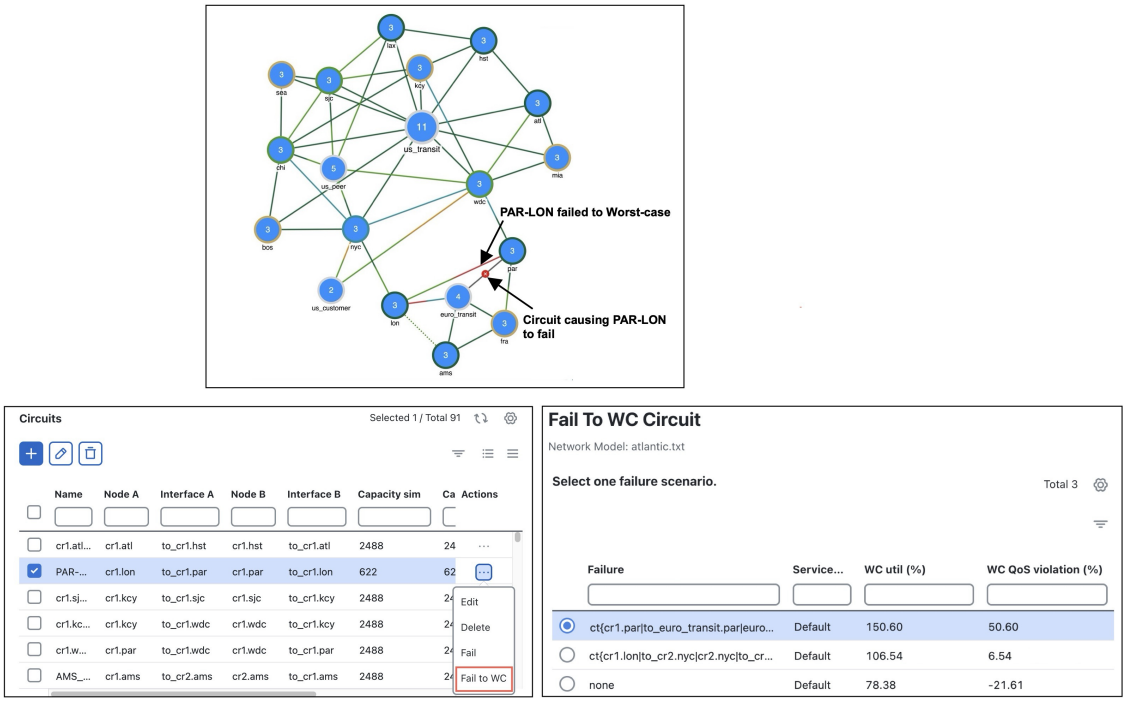
To fail the interface/circuit to its worst-case utilization, do the following:

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 (see Figure 40: Failure of single circuit to its worst-case, on page 107).
4. Click **Submit**.

The network plot changes to show this particular failure scenario (see Figure 40: Failure of single circuit to its worst-case, on page 107).

Figure 40: Failure of single circuit to its worst-case



Worst-case demand latency

When running Simulation analysis (see Run simulation analysis, on page 109), 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.

The following 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 the following 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

After running Simulation analysis (see Run simulation analysis, on page 109) to calculate demand worst-case latency, you have the option to fail a single demand to its worst-case latency.

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.

The failure scenarios causing the worst-case latency are listed in the **WC latency failures** column of the Demands table.

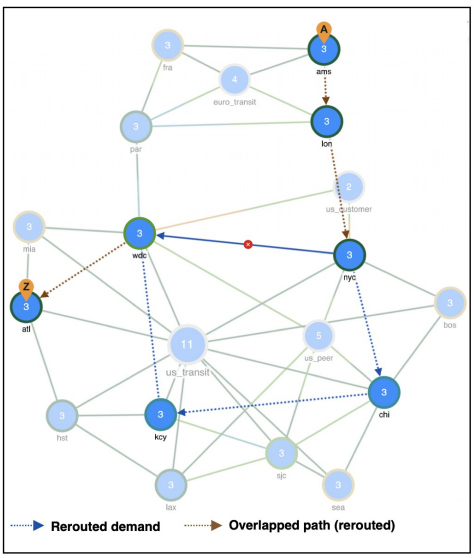
To fail a demand to its worst-case latency, do the following:

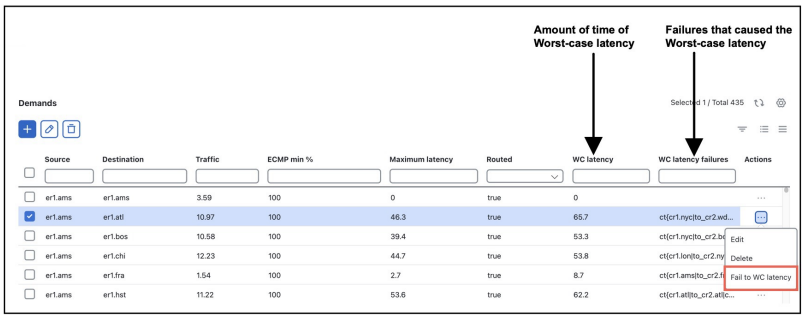
1. In the **Worst-case traffic** view, select the desired demand from the **Demands** table.
2. From the **Actions** column, choose ••• > **Fail to WC latency**.

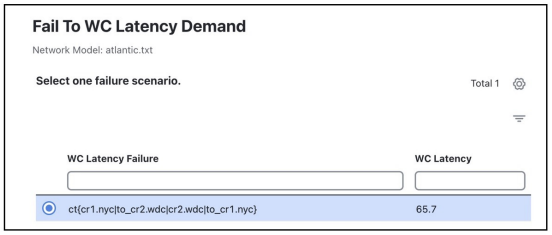
   The Fail to WC Latency Demand page appears.
3. Choose the failure scenario of interest (see Figure 41: Example of worst-case demand latency, on page 108).
4. Click **Submit**.

The network plot changes to show this particular failure scenario (see Figure 41: Example of worst-case demand latency, on page 108).

Figure 41: Example of worst-case demand latency







Run simulation analysis

The Simulation analysis tool is the basis of 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

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

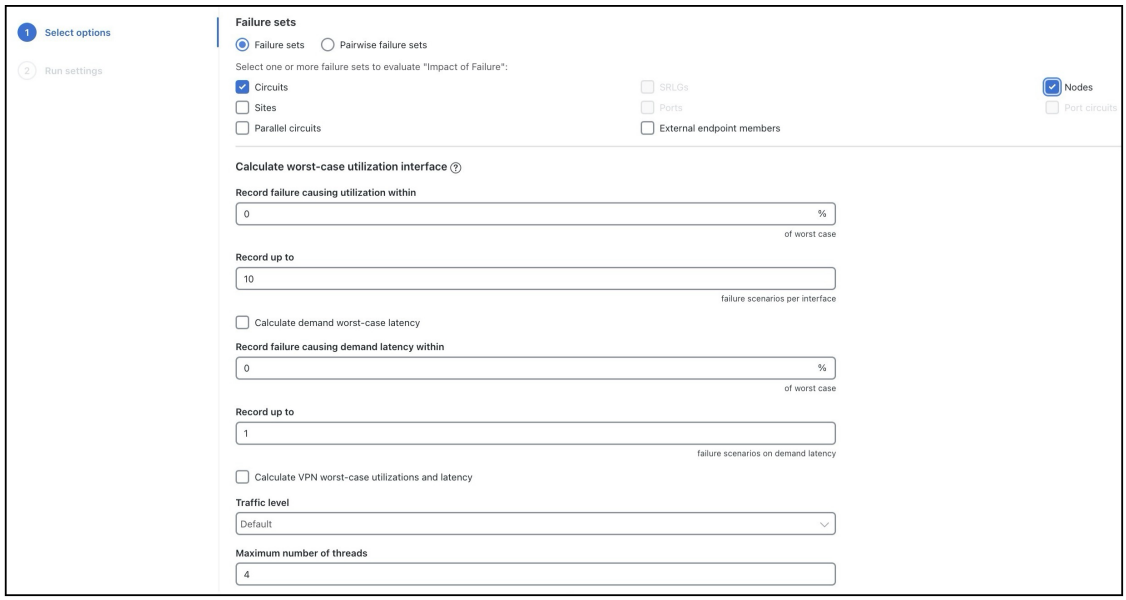
Step 2 From the toolbar, choose any of the following options:

- **Preset workflows > Evaluate impact of failures**

  OR

- **Actions > Tools > Simulation analysis**

Figure 42: Configure simulation analysis



Step 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 details, see Pairwise failure analysis, on page 111.

Step 4 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.

Step 5 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 if 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 if 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.

Step 6 Select whether or not to record demand worst-case latency calculations using the **Calculate demand worst-case latency** check box.

Step 7 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.

Step 8 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 if 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 if 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.

Step 9 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, see Simulate VPN, on page 165.

Step 10 Enter the value of maximum number of threads in the **Maximum number of threads** field.

Step 11 Click **Next**.

Step 12 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

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.

Step 13 Click **Submit**.

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. Here you can also 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 118 view to identify the circuits that are responsible for worst-case traffic congestion.

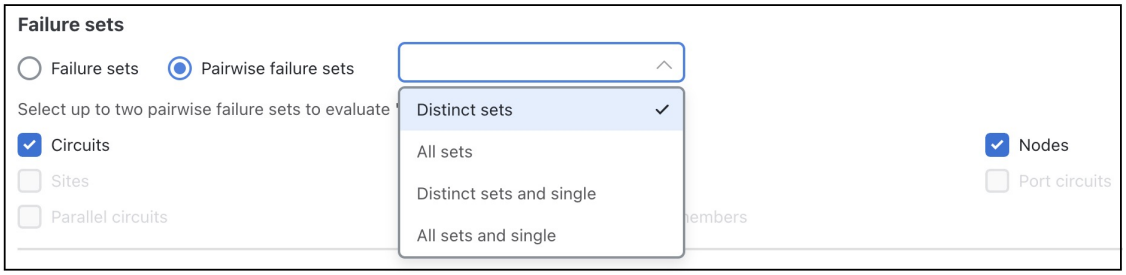
Pairwise failure analysis

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 analyse 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.

To perform pairwise failure analysis:

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



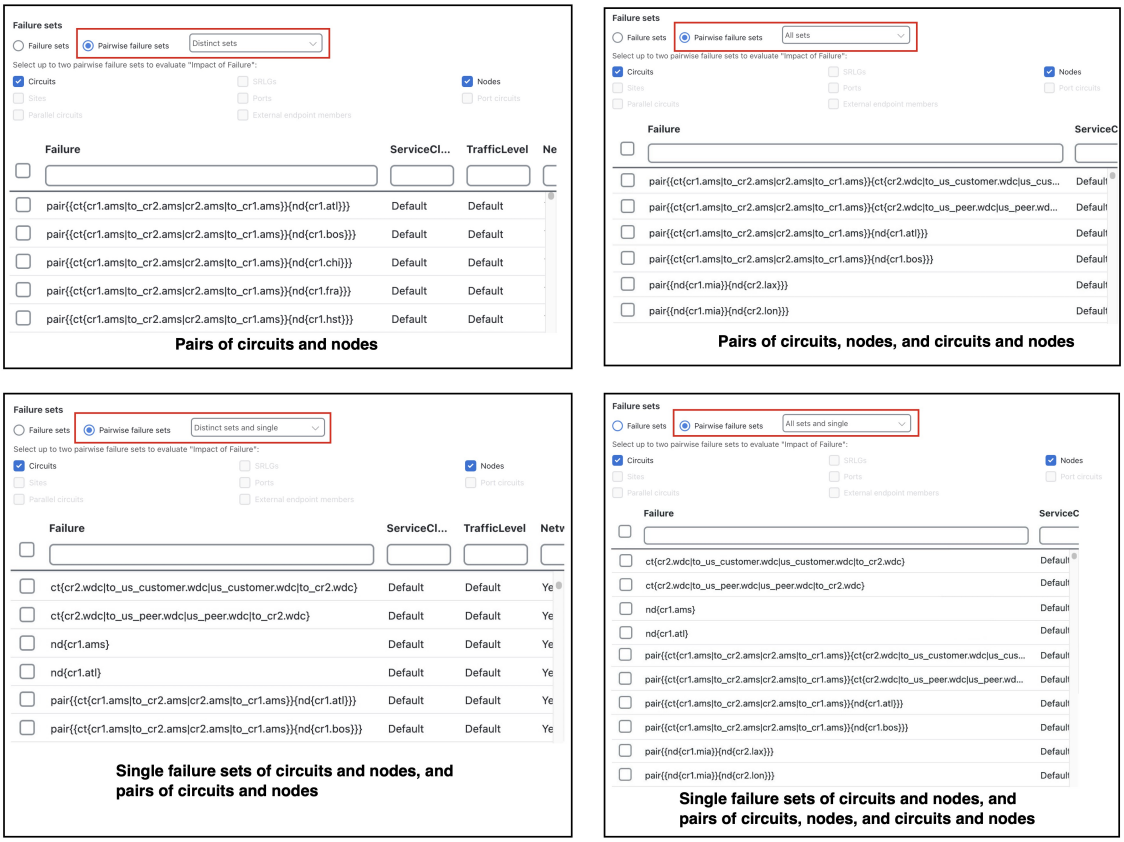
- **Distinct sets**: This option creates pairwise failures using elements from distinct failure sets. For example, if you select circuit and node as failure sets, this option considers the combination of distinct sets (circuit, node) for pairwise failure analysis.
- **All sets**: This option creates 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).

Important

- 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, on page 112 shows the example of the failure sets considered in all the above scenarios.

Figure 44: Pairwise failure sets selection and results



Pairwise failure analysis example

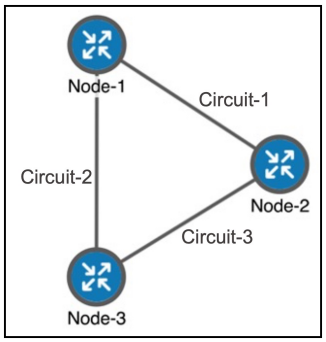
This example demonstrates how the selection of failure sets and pairwise failure options determines which failure scenarios the Simulation analysis tool evaluates.

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:

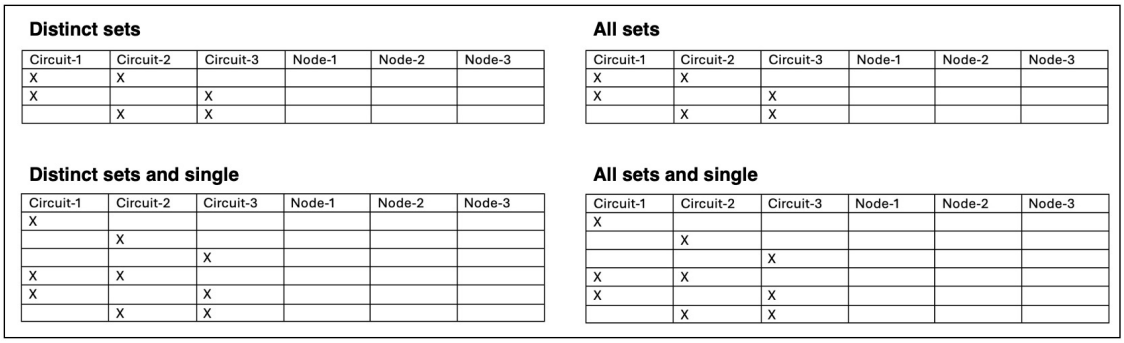
- When only circuits are selected as failure sets
- When both circuits and nodes are selected

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. **Note** 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) • 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. **Note** 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) • Distinct pairs: (Circuit-1, Circuit-2), (Circuit-1, Circuit-3), (Circuit-2, Circuit-3)

Figure 46: 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			

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			

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			

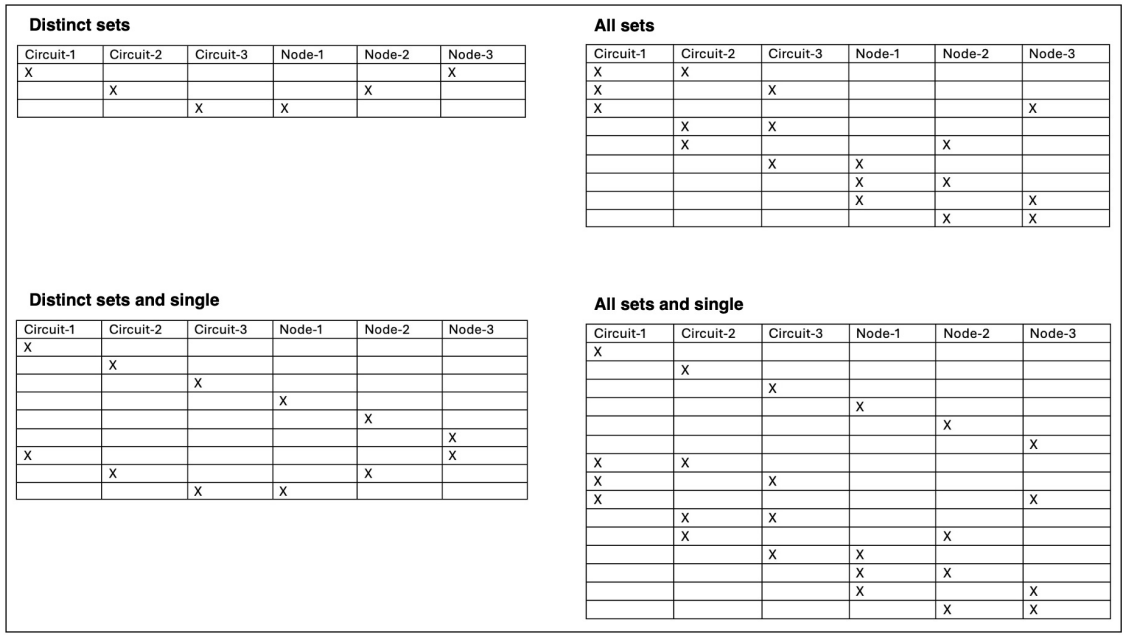
Scenario 2: Pairwise failure analysis with "Circuits" and "Nodes" selected

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. In Figure 45: Sample network topology, on page 113, 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: • Pairs of circuits • Pairs of nodes • Distinct circuit and node pairs	• Circuit-Circuit pairs: (Circuit-1, Circuit-2), (Circuit-2, Circuit-3), (Circuit-1, Circuit-3) • Node-Node pairs: (Node-1, Node-2), (Node-2, Node-3), (Node-1, Node-3) • 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) • Distinct Circuit-Node pairs: (Circuit-1, Node-3), (Circuit-2, Node-2), (Circuit-3, Node-1)

Pairwise failure option	Action	Generated failure scenarios
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. This option generates the largest number of scenarios.	• Single failures: (Circuit-1), (Circuit-2), (Circuit-3), (Node-1), (Node-2), (Node-3) • 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



Protect objects

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 window. 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

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 Select one or more like objects from their respective tables.

Step 3 Click .

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

Step 4 In the **State** field, check the **Protected** check box.

Step 5 Click **Save**.

Analyze simulation analysis reports

Each time Simulation analysis 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 **Simulation Analysis** link in the right panel. Note that new reports replace the previous ones.

The **Options** tab displays the input parameters used for the simulation analysis.

The **Summary** tab details the options used in the analysis and 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 (Table 4: Simulations table in simulation analysis report, on page 116).

Table 4: 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.

Simulation Data Point	Description
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. • Yes (Total)—A break exists that completely partitions the network into two or more disconnected sections. • 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. • 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. • 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.
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

The **Failure impact** view is available upon running a Simulation analysis (Figure 48: Example failure impact, on page 119). The plot in this view colors the nodes and circuits according to the maximum utilization level that would be caused elsewhere in the network should the node or circuit fail. The color indicates the resulting utilization and severity of the congestion.

Example: In the Failure impact view, 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.

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}

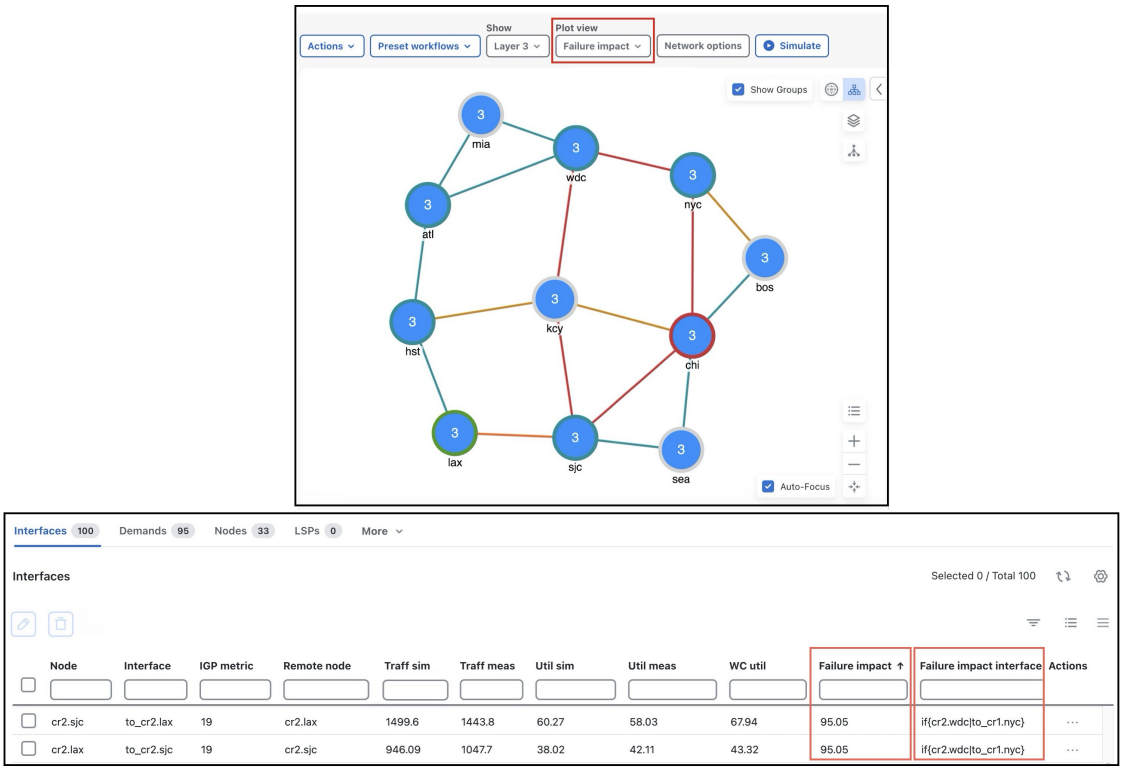
Example: if{cr2.sjc|to_cr1.kcy} means if the circuit goes down, it will have the greatest traffic impact on the cr2.sjc to cr1.kcy interface.

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.

Figure 48: Example failure impact



Parallelization

Simulation Analysis tool allows parallelizing the computaion and hence arrive at a faster result for large network models.

Example: If there are 10000 circuits in a network model, and there are 10 different engines available, the tool can be used to break up the network model into 10 partitions with each partition handling 1000 failure scenarios. This results in 10 different result files. The results from each of these independent runs are merged together to obtain the final result.

Execute parallelization:

Use the CLI tool to execute parallelization. For details, see Run tools or initializers using CLI, on page 342.

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, use the following:

```
-failure-sets <failure-sets> -num-partitions <number-of-partitions> -num-threads
<number-of-threads> -partition-index <partition-index> -result-file <result-filename>
```

where

- **-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. Can be *.txt or *.db file.

For more information on running tools and initializers using CLI, see Run tools or initializers using CLI, on page 342.

Merge results:

To merge the results, invoke the merge_sim_analysis CLI using the Python script. For details on running the scripts, see Run external scripts, on page 343.

Sample script (run_cli_merge.py):

```
import os
import sys

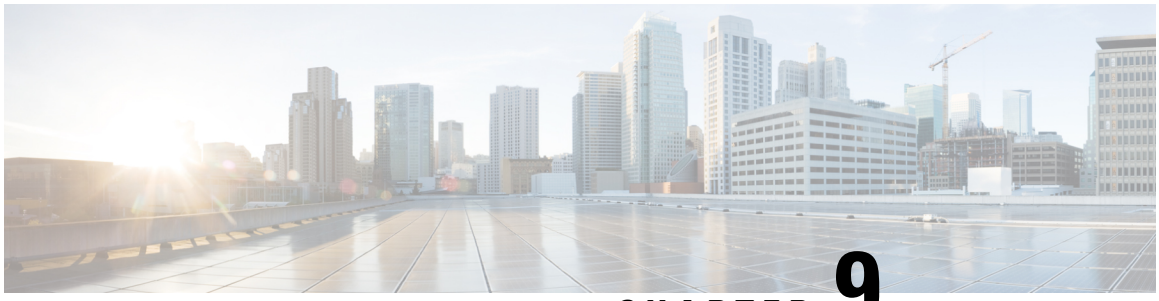
cmd = "merge_sim_analysis -plan-file {0} -partial-results {1} -out-file out-plan.txt ".format(
 sys.argv[1], sys.argv[2])
print(cmd)
os.system(cmd)
```

where

- **-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.

In the Job Manager, enter the following 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
```



CHAPTER 9

Evaluate Impact of Traffic Growth

Cisco Crosswork Planning traffic forecasting tool (**Create growth plans**) lets you create plans containing estimates of future traffic based on projected growth of current measured or simulated traffic. From these new plans you can determine the impact of the new traffic on the network and plan for upgrades. The traffic growth estimates used in this process can be manually entered, for example, based on knowledge of new services being introduced. After creating these growth plans, you can analyze their capacity requirements using Simulation analysis, and use the UI to upgrade the network to meet these requirements.

One example use of growth plan tool would be to predict the effects of a certain percentage of increased traffic for each quarter of the coming year. You can generate a separate plan for each of these traffic predictions to forecast the impact of each quarter's anticipated growth.

This section contains the following topics:

- Demand Groupings, on page 121
- Growth plans, on page 122

Demand Groupings



Important 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.

Demand groupings define a group of demands. They provide a convenient way of specifying aggregated traffic in a plan, which can be used as a basis for traffic reports and growth plans. For example, you can use demand groupings to represent the following.

- Total traffic sourced from one specific site
- Total VPN traffic, or any defined service class, sourced from one specific site to another
- Total traffic destined for a particular AS

Growth plans

The **Create growth plans** tool (**Actions** > **Tools** > **Create growth plans**) lets you generate new plan files containing predictions of future traffic, based on growth forecasts for the current plan. The predicted traffic growth can be generated in one of three ways.

- Per demand grouping using demand grouping traffic growth rates—See Growth plans from demand grouping, on page 124.
- Per demand using demand traffic growth rates—See Create growth plans from demand growth, on page 124.
- Per interface using interface traffic growth rates—See Create growth plans from interface growth, on page 125.

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. By automatically running a simulation analysis on the new plans, you can incorporate worst-case performance into planning decisions.

Create growth plans

To create growth plans, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

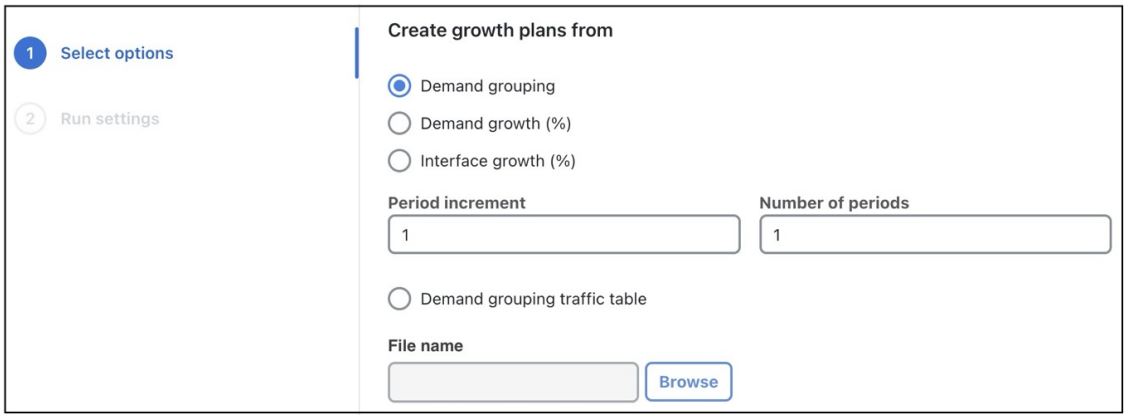
Step 2 From the toolbar, choose any of the following options:

- **Preset workflows > Evaluate impact of traffic growth**

  OR

- **Actions > Tools > Create growth plans**

Figure 49: Create Growth Plans Settings



Step 3 Specify one of the following growth methods:

- Demand grouping—For information on how to set demand grouping growth based on fixed, incremental, or percents, see Growth plans from demand grouping, on page 124.
- Demand growth (%)—For information on how to set demand growth, see Create growth plans from demand growth, on page 124.
- Interface growth (%)—For information on how to set interface growth, see Create growth plans from interface growth, on page 125.
- Demand grouping traffic table—Browse to or enter the name of the file containing the Demand grouping traffic table.

Step 4 In the following fields, 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.

Step 5 Click **Next**.

Step 6 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

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.

Step 7 Click **Submit**.

Step 8 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 details, see Import plan files from the local machine, on page 19).

Growth plans from demand grouping

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.

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.

Create growth plans from demand growth

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 the plans, you have the option of how many plans to create, and each plan will show growth at the same specified rate relative to the previous plan.

Follow these steps to populate the **Growth (%)** column in the **Demands** table.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the Network Design page.

Step 2 Select one or more demands from the **Demands** table.

Step 3 Click .

Step 4 In the **Traffic** section, enter the growth rate in the **Growth %** column.

Step 5 Click **Save**.

Example

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

This method grows measured traffic based on the **Growth (%)** column in the **Interfaces** table in the current network model. It differs from demand-based growth methods because measured traffic, rather than demands (simulated traffic), is grown. This is useful for basic forecasting of the current network state, rather than worst-case failure states, which require simulation. When generating the plans, you have the option of how many plans to create, and each plan will show growth at the same specified rate relative to the previous plan.

You can also use this method in conjunction with the Demand deduction tool to create demands from the forecast measurements.

You can populate the **Growth (%)** column in the **Interfaces** table using the following steps:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the Network Design page.

Step 2 From the toolbar, choose **Actions** > **Edit** > **Traffic** > **Measured traffic**.

The Measured traffic page opens.

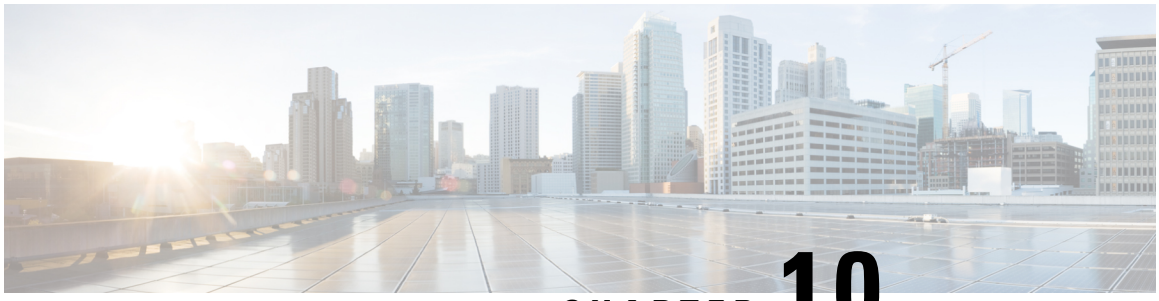


Step 3 For each applicable interface, click the **Edit** button and in the **Growth (%)** column, enter the percent by which you want to increase the interface traffic.

Step 4 Click the **Save** button.

Example

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.



CHAPTER 10

Perform Capacity Planning

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 to them
- Augment these circuits with associated port circuits
- Add parallel circuits to existing circuits
- 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 addition of any required capacity to be installed on the network. 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.

This section contains the following topics:

- Optimize capacity, on page 127
- Analyze optimization reports, on page 134

Optimize capacity

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, and then layering additional options that operate on Layer 3 and across failure sets.

Note Some of the optimizer options are mutually exclusive; for example, if you tell 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

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose any of the following options:

- **Preset workflows > Perform capacity planning**

  OR

- **Actions > Tools > Capacity planning optimization**

Step 3 From the **Circuits** panel, choose 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 131.

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 the following options:

- Create new port circuits (LAGs): In this case, the optimizer augments the existing circuits with associated port circuits (LAGs) with additional port circuits. 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.

Note
You can apply the IGP metric to new adjacencies, but not to parallel circuits.

Step 4 (Optional) Override the defaults for how upgraded circuits and new objects are tagged.

Step 5 Click **Next**.

Step 6 (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, see Advanced optimization options, on page 132.

Step 7 (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.

Step 8 Click **Next**.

Step 9 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

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.

Step 10 (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.

Step 11 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.

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 details, see Analyze optimization reports, on page 134.

Example

Figure 50: Design before optimization, on page 130 shows an example of a network design before optimization, where extra demands were placed on the network. As a result, you can see congested interfaces in red.

Figure 50: Design before optimization

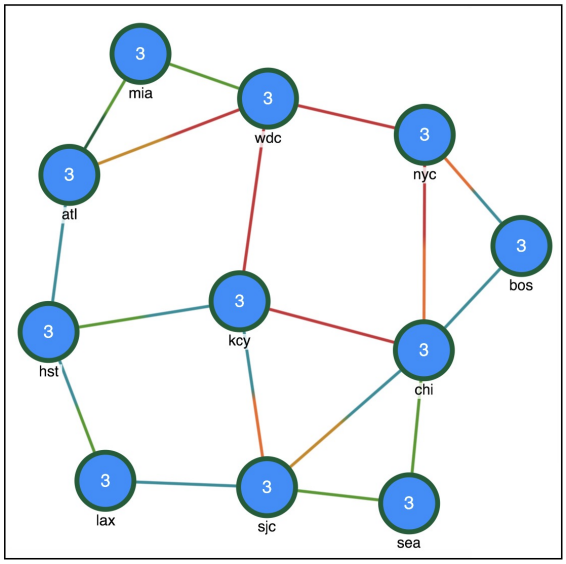


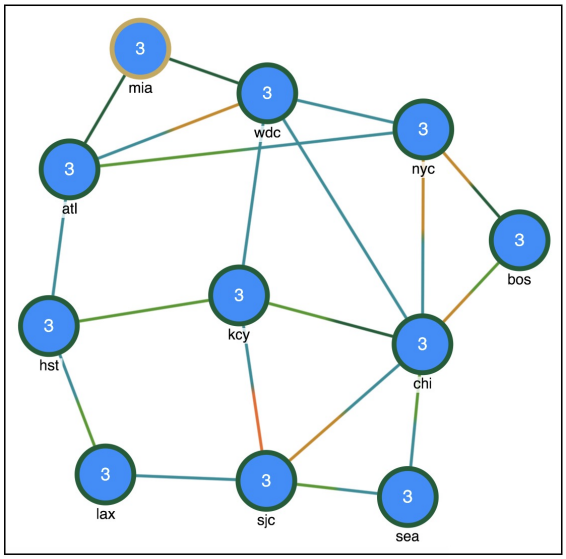
Figure 51: Design after optimization, on page 131 shows the design after optimization using the following parameters (For the remaining options, the default values are used.):

- 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



Optimization options

You have several Layer 3 options for capacity planning. These parameters tell the optimizer what your preferences are in terms of utilization thresholds, capacity increments, and whether to upgrade existing circuits, create new adjacencies between nodes, create port circuits, or create parallel circuits.

Table 5: 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. If you want to allow the upgrade of just certain circuits, select those circuits to be modified. 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.

Option	Description
Create • Port circuits (LAGs) • 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.
	• 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.
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 • Shortest IGP metric minus decrement • 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.
	• 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

Use the advanced optimization options to:

- Choose whether to optimize for capacity or cost.
- Define what 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 6: Advanced optimization options

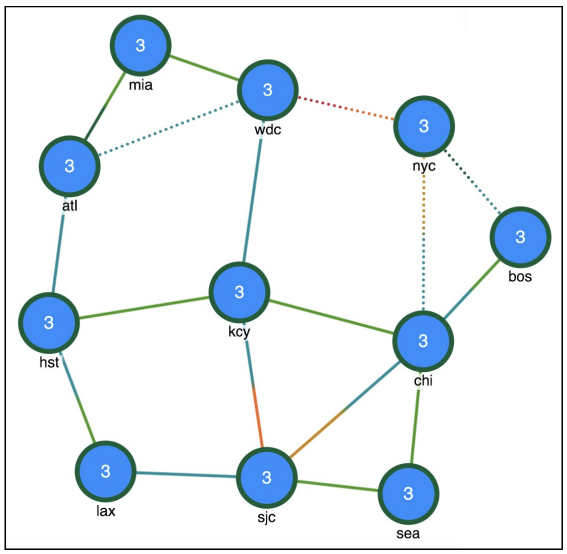
Option	Description
Optimization objective	
Minimize capacity	Minimizes the capacity that is added to your network. This is the default option.
Minimize cost	Select this option to run the optimization with the objective of reducing the overall cost when adding new capacity. Click to manually enter a cost unit for capacity, L3 port costs, and the feasibility limit. Alternately, click to import cost values from a file.
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.

Parallel circuits design example

Figure 52: Adding parallel circuits to increase capacity, on page 134 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



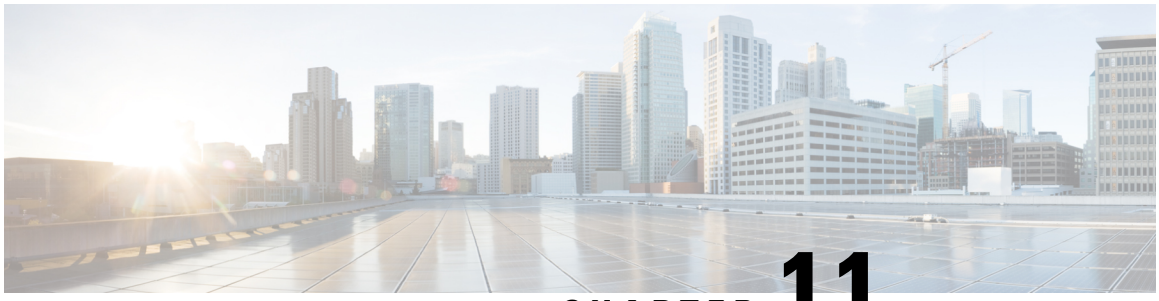
Analyze optimization reports

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.

The **Summary** tab provides useful metrics at a glance. It provides details on the:

- 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

The **Capacity Upgrades** tab provides details on how the optimizer upgraded circuits.



CHAPTER 11

Simulate IGP Routing Protocol

This chapter describes options for simulating an IGP. Cisco Crosswork Planning can simulate IS-IS, OSPF, and EIGRP, as well as supports multi-topology routing for IS-IS and OSPF IGPs.

This section contains the following topics:

- IGP process ID, on page 135
- OSPF and IS-IS, on page 136
- EIGRP, on page 139
- Simulate IGP multipath, on page 139
- Exclude ABR nodes, on page 140
- Configure IGP simulation, on page 140
- Simulate multi-topology routing, on page 141

IGP process ID

An **IGP process ID** is used to differentiate between multiple instances of the same IGP running on a router. This allows you to have separate IGP configurations for different purposes within the same router.

In Cisco Crosswork Planning, you can specify IGP process ID to an IGP protocol. Also, each interface can be associated with an IGP process ID.

Configure IGP process ID

Follow these steps to create, delete, or edit an IGP process ID.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Edit** > **IGP process protocols**.

The IGP process protocols page opens.

Step 3 Configure IGP process ID.

Figure 53: Configure IGP process ID



a) To create IGP process ID, click , enter the name, 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 that you want to delete, and then click or use the **Delete** button.

Associate interfaces with an IGP process ID

Follow these steps to associate each interface with an IGP process ID.

Procedure

Step 1 From the **Interfaces** table, select one or more interfaces for which you want to add the IGP process ID.

Step 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.

Step 3 In the **IGP** panel, enter the value in the **IGP process ID** field.

Step 4 Click **Save**.

OSPF and IS-IS

OSPF and IS-IS simulations are identical with the following 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 an alternate metric for Level 1 can also be specified.

OSPF area simulation

OSPF area membership can be specified per interface. For instructions, see Set OSPF area membership, on page 137.

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. Nodes are assigned to areas, as follows:

- 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

To specify OSPF area membership, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, select one or more interfaces from the **Interfaces** table.

Step 3 Click .

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

Step 4 Click the **Advanced** tab.

Step 5 In the **IGP** panel, enter the required value in the **OSPF area** field.

Step 6 Click **Save**.

IS-IS multi-level simulation

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, see Set IS-IS multi-level simulation, on page 138.

- 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.

By default, all nodes in a single AS are assumed to belong to a single IS-IS level. Nodes are assigned to levels, as follows:

- 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

To set IGP metrics, IS-IS levels, and Level 1 metrics, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, select one or more interfaces from the **Interfaces** table.

Step 3 Click .

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

Step 4 To set the IGP metric, enter its value in the **IGP metric** field.

Step 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.

Step 6 Click **Save**.

EIGRP

The IGP Metric value for each interface is not used when EIGRP routing is selected. Instead, EIGRP uses the following formula to derive total feasible distance (total cost) from a node to a destination subnet.

```
path metric to destination = (10,000/(bandwidth) + (delay) * 256
```

- 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

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.

Follow these steps to configure ECMP.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the toolbar, click **Network options** or choose **Actions > Edit > Network options**.

The Network options page opens.

Step 3 Click the **Protocols** tab.

Step 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.

Step 5 Click **Save**.

Exclude ABR nodes

To exclude the nodes while entering or exiting the domain, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 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.

Step 3 To add a node that has to be excluded:

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) Click **Save**.

Step 4 To delete an ABR node:

a) Select the row that you want to delete.
b) Click .

Configure IGP simulation

Follow these steps to configure IGP simulation.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the toolbar, click **Network options** or choose **Actions > Edit > Network options**.

The Network options page opens.

Step 3 Click the **Protocols** tab.

Step 4 Under the **IGP** section, choose the appropriate protocol from the **Protocol** drop-down list.

Step 5 Click the **Simulation** tab.

Step 6 Select an option for **Redistribute routes across IGP process**.

The options include:

- No IGP redistribution
- Shortest exit

- Shortest path

Step 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.

Step 8 Click **Save**.

Simulate multi-topology routing

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

To create, rename, or delete topologies, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions > Edit > Topologies**. The Topologies page opens.

Step 3 To create a new topology:

a) Click .
b) Enter the topology name.
c) Click the **Save** button.

Step 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

To associate demands or LSPs with topologies, do the following. 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 ( ).

Before you begin

Create the topologies. For details, see .

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, select one or more demands or LSPs from their respective tables.

Step 3 Click .

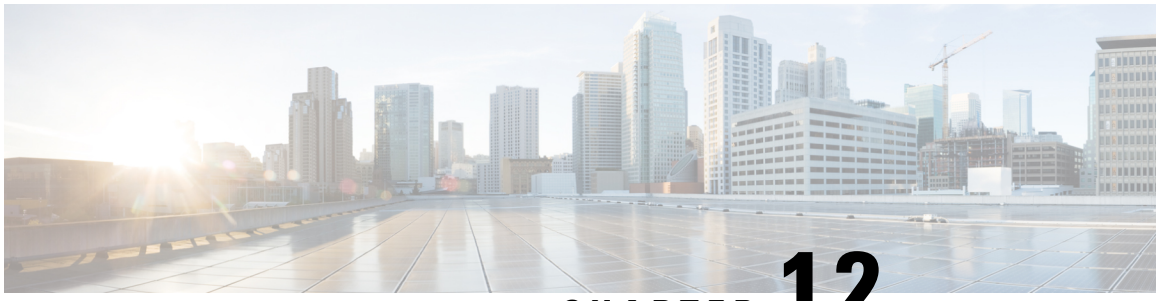
Step 4 Select the desired topology or topologies, as follows.

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) Click **Save**.

By selecting from a similar **Topology** drop-down list, you can also add a topology to LSP meshes when creating them.



CHAPTER 12

Simulate BGP Routing

This chapter describes how Cisco Crosswork Planning models multi-AS networks and simulates basic BGP routing. Cisco Crosswork Planning 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 quickly and easily evaluate the effects of peering locations and basic policy variations.

Additionally, you can 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, see Simulate Traffic Flow from Source to Destination Using Demands, on page 71 and Simulate Advanced Routing with External Endpoints, on page 177.

This section contains the following topics:

- Internal and external AS types, on page 143
- Configure ASes, on page 144
- Route demands between ASes, on page 146
- Know about BGP routing details, on page 150
- BGP routing, on page 151

Internal and external AS types

To model a multi-AS network, each node is assigned an AS, and each AS is defined as either *internal* or *external*. A typical multi-AS model in Cisco Crosswork Planning consists of the following:

- 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 AS Properties window and listed in the AS table. Nodes are assigned to ASes in the Node Properties window.

Configure ASes

Create ASes

Follow these steps 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. See Associate nodes with an AS, on page 144 and Edit AS routing policies, on page 145.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **AS**.

OR

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.

Step 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.

Step 4 Click **Submit** to create an AS.

Step 5 To change the routing policy, select the AS and click  and then choose the **AS Relationships** tab (see Edit AS routing policies, on page 145).

Associate nodes with an AS

To associate nodes with an AS, do the following:

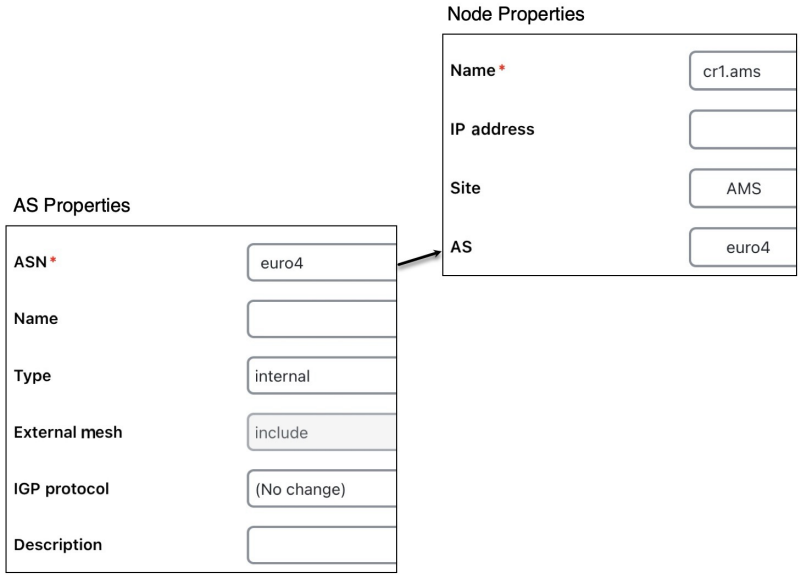
Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, choose one or more nodes from the **Nodes** table.

Step 3 Click .

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



Step 4 From the **AS** drop-down list, choose the AS to which you want to assign the nodes.

Step 5 Click **Save**.

Edit AS routing policies

To create AS relationships, set the routing policy.

Procedure

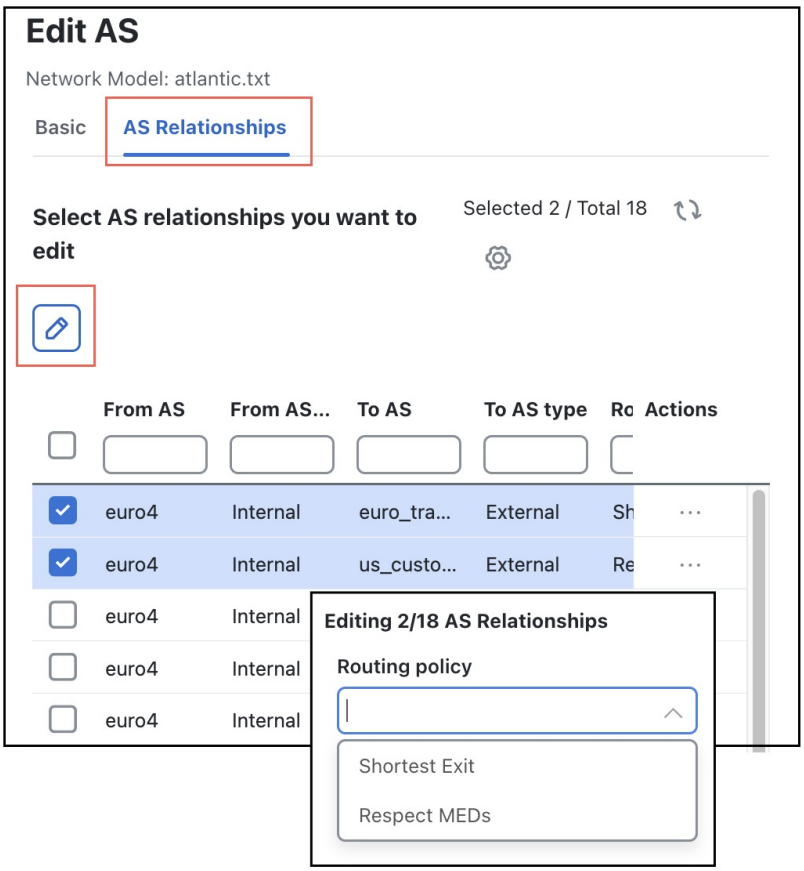
Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, choose an AS from the **AS** table.

Step 3 Click  or use the ••• > **Edit** option under the **Actions** column.

Step 4 Click the **AS Relationships** tab.

Figure 54: Edit AS relationships



Step 5 Choose the AS pair that you want to configure. There is a separate line for each direction in the relationship so you can configure them independently.

Step 6 Click the icon, set the **Routing policy** to **Respect MEDs** or **Shortest Exit**. For details, see Route demands between ASes, on page 146.

Step 7 Click **Submit**.

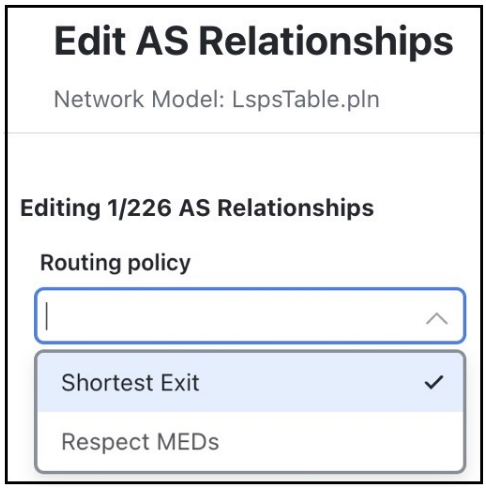
Route demands between ASes

Determine routes between internal ASes

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.

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 window, which is accessed through the Edit AS window.



- **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.

Determine routes between external and internal ASes

Table 7: Typical AS routing configurations, on page 147 lists typical routing configurations that can be constructed by applying different combinations of routing policies for traffic in both directions between two ASes.

- 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 7: 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.

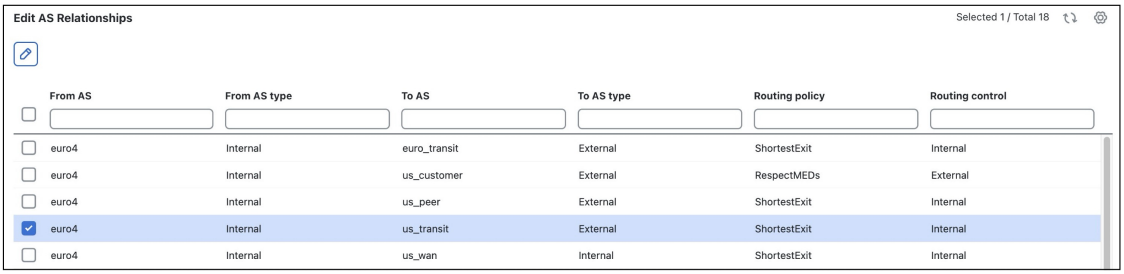
Failover between nodes in the external AS can be modeled. For example, if the traffic is sourced from an external AS and if the peering circuit from which traffic is entering the internal AS under normal operation fails, the traffic can enter the internal AS from a different interface or peering node in the same external AS. In the Demands table, the sources and destinations are represented as follows.

AS{<ASN>}:if{node_name|interface_name}

Example: AS{33287}:if{cr01.newyork.ny|POS3/7/0/0}

For more detailed information on demand sources and destinations, see Simulate Traffic Flow from Source to Destination Using Demands, on page 71.

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.



The AS that controls the routing is determined by the AS type, direction of the demand, and the Routing policy property as described in Table 8: Determining the AS that controls the routing, on page 148.

Table 8: 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

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 (Table 9: Effects of routing policy and routing control , on page 149).

- 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 9: 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

Configure external meshes

An external mesh consists of two or more external ASes with a **Type** property of **external**. 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 55: External mesh control, on page 150 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 10: External mesh settings for common AS relationships, on page 149 shows the External Mesh settings for common AS relationships.

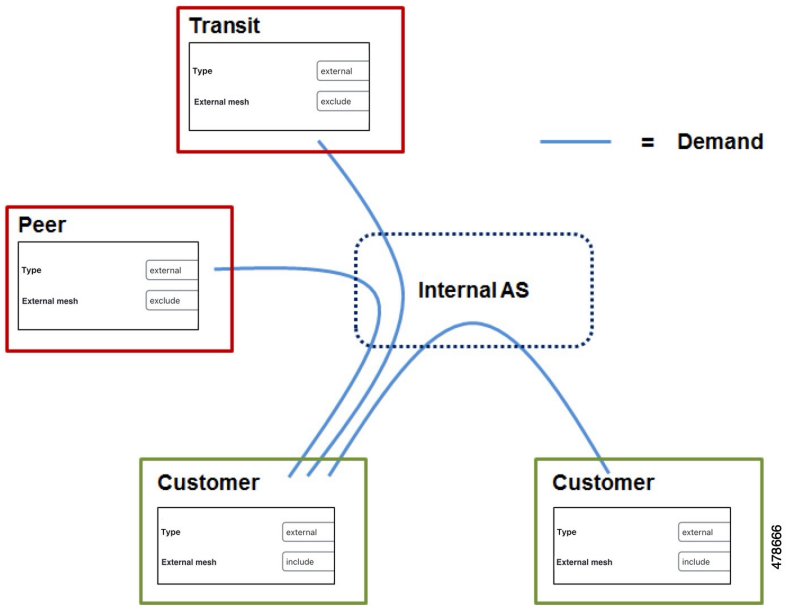
The External mesh property is set in the Edit AS window.

Table 10: External mesh settings for common AS relationships

Relationship	External Mesh Setting	Result
Peer	Exclude	Demands permitted to/from customers only
Customer	Include	Demands permitted to/from all external ASes
Transit	Exclude	Demands permitted 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 55: External mesh control



Know about BGP routing details

BGP multihop

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.

BGP load balancing

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.

Follow these steps to set BGP multipath options globally.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the toolbar, click **Network options** or choose **Actions** > **Edit** > **Network options**.

The Network options page opens.

Step 3 Click the **Protocols** tab.

Step 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.

Step 5 Click **Save**.

BGP next hop

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 window.

BGP routing

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.

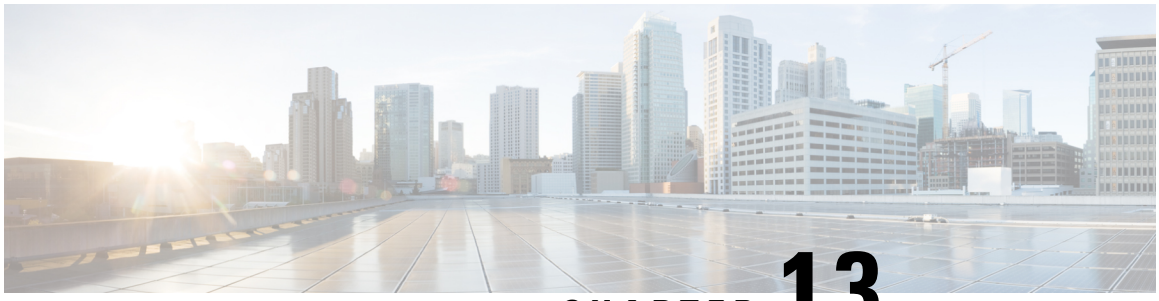
Procedure

Step 1 Demands are routed using the established LSPs (if applicable) and using the specified BGP protocols given the specified failure scenarios.

Step 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.



CHAPTER 13

Simulate Quality of Service (QoS)

Quality of Service (QoS) is a means of ensuring high-quality performance for critical applications. The concept is that because requirements of some users and services are more critical than others, some traffic requires preferential treatment.

Using Cisco Crosswork Planning QoS features, you can ensure 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.
- 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.

Note Cisco Crosswork Planning supports only Default queue. Also, you do not have the option to map the service classes to queues.

This section contains the following topics:

- QoS requirements, on page 153
- QoS bound and QoS violations, on page 155
- Configure queues and service classes, on page 161
- Create service class policies, on page 163

QoS requirements

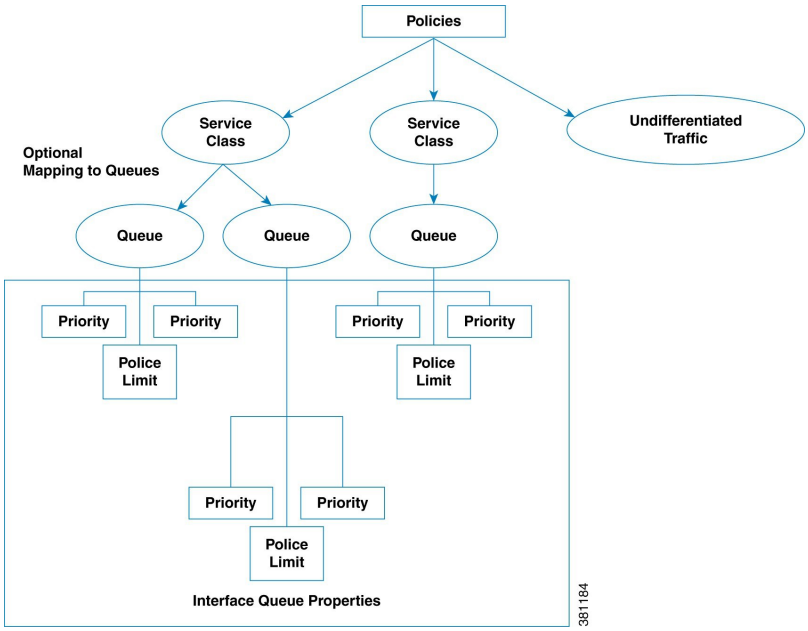
In Cisco Crosswork Planning, QoS requirements are defined by policies and interface queue properties.

- Policy—Maximum percentage of traffic capacity that can be utilized 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 would this parameter have any effect on live network behavior.

- Interface queue properties—Configured parameters that would 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, see Edit QoS requirements, on page 155.
  - The **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 evenly share the capacity based on weighted-round robin (WRR) calculations. You can change this behavior using the weight and police limit parameters. There are unlimited number of priorities, though most networks only use no more than three. By default, queues do not have priorities.
  - The **Weight** is 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%.
  - The **Police limit** is 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 156, where you can see that lower priority queues received zero traffic due to priority settings.

Figure 56: Policies and interface queue parameters



Edit QoS requirements

To edit the QoS parameters using Interface queues properties, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 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**.

Step 3 Click .

Step 4 Update one or more QoS fields (Priority, Weight, and Police limit) to create the desired QoS requirement.

Step 5 Click **Save**.

QoS bound and QoS violations

Cisco Crosswork Planning uses the concepts of **QoS bound** and **QoS violation** as a way of identifying whether QoS parameters are being met or surpassed, thus better enabling you to 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—Maximum interface capacity available without violating these QoS requirements. A separate QoS bound is 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: • Policy for service class • Queue properties for queues
Undifferentiated traffic	Policy

In the Network Summary table, the columns that convey the QoS bound information are: QoS bound meas, QoS bound meas (%), QoS bound sim, and QoS bound sim (%).

QoS bound calculations are a set of decisions being made to determine how to raise traffic on the queues until that traffic cannot be raised any further. This capacity, or the reason the traffic cannot be raised 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 raise 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—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

If no other QoS parameters are set via the Interface queue properties of Priority, Weight, and Police limit, the QoS bound is equivalent to the policy set.

Table 11: Example policies and QoS bound calculations

Example configuration	QoS bound	QoS violation (Positive # = Violation)
Interface capacity = 10,000 Mbps Undifferentiated traffic = 5000 Mbps Normal operation policy = 60%	6000 Mbps (60%)	–1000 Mbps (–10%) Because this number is negative, there is no capacity violation.
Interface capacity = 10,000 Mbps Undifferentiated traffic = 8000 Mbps Normal operation policy = 60%	6000 Mbps (60%)	2000 Mbps (20%) Because this number is positive, there is a capacity violation.
Interface capacity = 10,000 Mbps Voice traffic = 6000 Mbps Video traffic = 2000 Mbps Voice normal operation policy = 90% Video normal operation policy = 60%	Voice = 9000 Mbps (90%) Video = 6000 Mbps (60%)	Voice = –3000 Mbps (30%) Video = –4000 Mbps (40%)

Interface queue properties and QoS bound calculations

Cisco Crosswork Planning simultaneously calculates QoS bound for each queue in the interface. In doing so, Cisco Crosswork Planning uses the Interface queue parameters (Priority, Weight, and Police limit) and the traffic measured or simulated 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

Provided policies are not set that further affect the QoS bound, a queue's QoS bound is calculated as follows:

- 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 12: Examples of priority QoS bound calculations

Example configuration	QoS bound	QoS violation (Positive # = Violation)	QoS bound calculations
Interface capacity = 20,000 Mbps EF traffic = 6000 Mbps; priority = 1 BE traffic = 3000 Mbps; no priority set	EF = 20,000 Mbps BE = 14,000 Mbps	EF = –14,000 Mbps BE = –11,000 Mbps	EF = Total interface capacity because it is the only priority 1 queue BE = 20,000 (interface capacity) – 6000 (consumed by higher priority queues)
Interface capacity = 10,000 Mbps EF Traffic = 6000 Mbps; priority = 1 BE Traffic = 5000 Mbps; priority = 2	EF = 10,000 Mbps BE = 4000 Mbps	EF = –4000 Mbps BE = 1000 Mbps	EF = Total interface capacity because it is the only priority 1 queue BE = 10,000 (interface capacity) – 6000 (consumed by higher priority queues

Weight

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 13: Examples of weight QoS bound calculations

Example configuration	QoS bound	QoS violation (Positive # = Violation)	QoS bound calculations
Interface capacity = 10,000 Mbps AF1 traffic = 3000 Mbps; priority = 1; weight = 100% AF2 traffic = 6000 Mbps; priority = 1; weight = 100%	AF1 = 5000 Mbps AF2 = 7000 Mbps	AF1 = –2000 Mbps AF2 = –1000 Mbps	AF1 = Half of capacity for priority 1 queues because both queues have equal weights AF2 = 5000 (half of capacity) + 2000 (unused AF1 capacity)
Interface capacity = 10,000 Mbps AF1 = 5000 Mbps; priority = 1; weight = 60% AF2 traffic = 6000 Mbps; priority = 1; weight = 40%	AF1 = 6000 Mbps AF2 = 5000 Mbps	AF1 = –1000 Mbps AF2 = 1000 Mbps	AF1 = 60% of capacity for all priority 1 queues AF2= 10,000 (interface capacity) – 5000 (consumed by AF1 queue)

Police limits

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 14: Examples of police limit QoS bound calculations

Example configuration	QoS bound	QoS violation (Positive # = Violation)	QoS bound calculations
Interface capacity = 10,000 Mbps EF traffic = 1000 Mbps; priority = 1; police limit = 50% BE traffic = 2000 Mbps; priority = 2	EF = 5000 Mbps BE = 9000 Mbps	EF = –4000 Mbps BE = –7000 Mbps	EF = 50% of total interface capacity BE = 10,000 (interface capacity) – 1000 (capacity consumed by EF)

Example configuration	QoS bound	QoS violation (Positive # = Violation)	QoS bound calculations
Interface capacity = 10,000 Mbps EF traffic = 1000 Mbps; priority = 1; police limit = 5% BE traffic = 2000 Mbps; priority = 2	F = 500 Mbps BE = 9500 Mbps	EF = 500 Mbps BE = –7500 Mbps	EF = 5% of total interface capacity BE = 10,000 (interface capacity) – 500 (capacity consumed by EF)
Interface capacity = 10,000 Mbps EF = 3000 Mbps; priority = 1; police limit = 20% AF1 traffic = 4000 Mbps; priority = 2; police limit = 75% AF2 traffic = 2500 Mbps; priority = 2; police limit 25%	EF = 2000 Mbps AF1 = 6000 Mbps AF2 = 4000 Mbps	EF = 1000 Mbps AF1 = –2000 Mbps AF2 = –1500 Mbps	EF = 20% of total interface capacity AF1 = 75% of (10,000 [interface capacity] – 2000 [capacity consumed by EF]) AF2 = 10,000 (interface capacity) – 2000 (capacity consumed by EF) – 4000 (capacity consumed by AF1)

Interface QoS bound calculations using multiple QoS parameters

Cisco Crosswork Planning calculates a QoS bound for interface queues based on all three parameters if they are all configured: priority, weight, and police limits.

Table 15: 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 EF = 3000 Mbps; priority = 1; police limit = 20% AF1 traffic = 4000 Mbps; priority = 2; weight = 75% AF2 traffic = 2500 Mbps; priority = 2; weight = 25%	EF = 2000 Mbps AF1 = 6000 Mbps AF2 = 4000 Mbps	EF = 1000 Mbps AF1 = –2000 Mbps AF2 = –1500 Mbps	EF = 20% of total interface capacity AF1 = Maximum of these two values. • 75% of (10,000 [interface capacity] – 2000 [capacity consumed by EF]) • 8000 (available capacity) – 2500 (AF2 traffic) AF2 = Maximum of these two values. • 25% of (10,000 [interface capacity] – 2000 [capacity consumed by EF]) • 8000 (available capacity) – 4000 (AF1 traffic)

Service class QoS bound calculations using multiple QoS parameters

If service classes have policies and they are mapped to queues, Cisco Crosswork Planning calculates a QoS bound for both. Cisco Crosswork Planning then uses the lowest value of the two so as to enforce restrictions in the strictest possible manner.

Example:

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.

View QoS bounds and QoS violations

Table 16: QoS bounds and QoS violations , on page 161 lists the available column options to display numeric values of the QoS bound calculations. For information on QoS values as they relate to VPNs, see Simulate VPN, on page 165.

Table 16: 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

Configure queues and service classes

Create service classes

Follow these steps to create service classes.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Edit** > **Manage QoS**.

Step 3 Under the **Service Classes** section, click .

An empty row appears.

Step 4 Enter a unique name under the **Name** column and click the **Save** button.

The newly created service class appears under this section.

Create queues

Cisco Crosswork Planning discovers queues. However, you can manually add them. Once discovered or created, queues appear in the Interface Queues table.

To create queues, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **Others** > **Interface Queues**.

OR

In the Network Summary panel on the right side, 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.

Step 3 Select the required interfaces.

Step 4 Click **Next**.

Step 5 Enter the queue name.

Step 6 (Optional) Enter the queue properties of priority, weight, and police limit. For information on how these queue properties behave, see Interface queue properties and QoS bound calculations, on page 156.

Step 7 Click **Submit**.

View queue and service class information

Table 17: Queue and service class information

To view	Show or select
Queue information	Show the **Queue** column in the **Interface queues** table (click , select the **Queue** check box, and click **Apply**).
Per-queue traffic in the Interfaces table	The **Traff meas** and **Traff sim** columns display data specific to the queue type selected.

To view	Show or select
Service class demands	Show the **Service class** column in the **Demands** table (click ⚙, select the **Service class** check box, and click **Apply**).

Create service class policies

You can configure policies for undifferentiated traffic and for service classes.

Follow these steps to configure service class policies.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Edit** > **Manage QoS**.

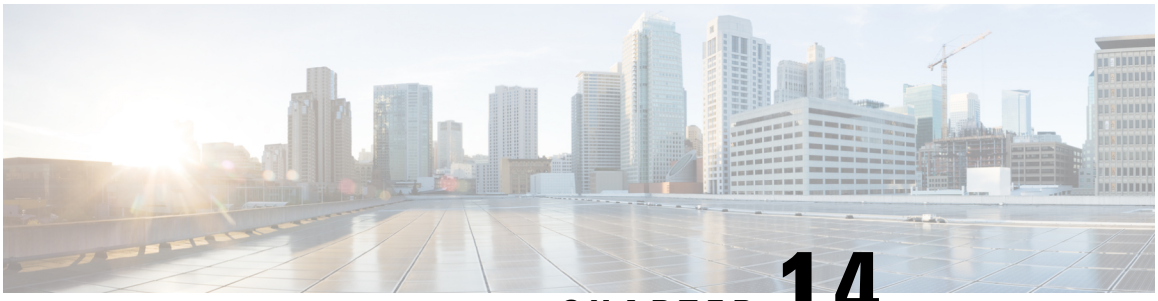
Step 3 Under the **Service Class Policies** section, click ⊕.

a) If creating a policy for undifferentiated traffic, select the **Undifferentiated Traffic** option. If creating a policy for an existing service class, select the **Service Class** option. Then, select the service class from the **Service Class** drop-down list.

b) In the **Normal operation (in %)** field, 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.

c) In the **Worst-case (in %)** field, 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.

Step 4 Click **Save**. These values now appear in the Manage QoS window.



CHAPTER 14

Simulate VPN

The Cisco Crosswork Planning Virtual Private Network (VPN) model is a representation of a virtual subnetwork within the network model. Viewing and simulating VPN within Cisco Crosswork Planning helps many network tasks and can answer questions, such as:

- 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 failures 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, and there are 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.

This section contains the following topics:

- VPN model, on page 166
- VPNs, on page 167
- VPN nodes, on page 169
- Layer 3 VPN example, on page 172
- VPN simulation analysis, on page 175

VPN model

VPN objects

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). • Layer 3 VPN: The VPN represents sets of VRFs associated with a set of VPN nodes that forward traffic between themselves. Often, 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. • Layer 3 VPN: The VPN node represents the VRF instances configured on each router.

VPN topology and connectivity

Cisco Crosswork Planning VPN topology route connections are established through Route Targets (RTs) or through a full mesh of VPN nodes. The **Connectivity** property is set in the Add/Edit VPN window.



Knowing a VPN's topology and connectivity lets Cisco Crosswork Planning calculate which demands between VPN nodes carry traffic for a particular VPN, and thus which interfaces carry traffic for that VPN. In turn, Cisco Crosswork Planning can calculate the vulnerability of a VPN to certain failure and congestion scenarios.

A demand is associated with a VPN, meaning it carries traffic for that VPN, if the following is true:

- The two VPN nodes are in the same VPN.
- The 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.

Once 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.

Connectivity	Description
Full Mesh	Full-mesh connectivity is a complete mesh of connections between VPN nodes in a VPN so 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.
Route Targets (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. Having an import/export pair does not create bidirectional communication. Rather, traffic flows in the opposite direction of the routed advertisements. 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.

VPNs

Each VPN consists of a set of VPN nodes that can exchange data within it. VPNs have the following key properties that uniquely identify them and define how the traffic within them is routed.

- **Name**—Unique name of the VPN.
- **Type**—Type of VPN. Choose from the options: VPWS, VPLS, or L3VPN.
- **Connectivity**—Determines how Cisco Crosswork Planning calculates connectivity and associated demands for VPNs:
  - 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 this VPN.

Once the VPN is created, it appears in the **VPN** drop-down list of VPN nodes.

Create VPNs

You can create new VPNs and then later add VPN nodes to them (see Add VPN nodes to VPNs, on page 171).

Create new VPNs

To create new VPNs, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **VPNs** > **VPN**.

OR

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.

Step 3 In the **Name** field, enter a unique name for the VPN.

Step 4 From the **Type** drop-down list, choose a VPN type. The options are: L3VPN, VPLS, and VPWS.

Step 5 Choose the Connectivity type: RT or Full Mesh.

Step 6 Choose the Service class for the VPN.

Step 7 Click **Add**.

Step 8 (Optional) Add VPN nodes to the newly created VPN. For details, see Add VPN nodes to VPNs, on page 171.

VPNs table

The VPNs table lists the VPN properties, its associated service class, traffic, and the number of VPN nodes within that VPN (Table 18: VPNs table columns for normal operation, on page 168). For information on QoS measurements, see Simulate Quality of Service (QoS), on page 153. For information on the Worst-Case columns not listed here, see Table 20: Simulation analysis columns in the VPNs table, on page 176.



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.



Note All traffic and QoS violations are based on traffic carried on all interfaces used by the VPN for the service class defined for that VPN.

Table 18: VPNs table columns for normal operation

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.

Column	Description
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 to VPNs through tables. When selected, all VPN nodes within the VPN are highlighted in the plot (Figure 57: VPN nodes within a VPN, on page 172).

Identify interfaces used by VPNs

To view which interfaces are associated with a VPN, select the VPN, click ☰, and choose **Filter to interfaces**. If you then choose all of these filtered interfaces, you can see the VPN outlined in the network plot.

Note Utilization measurements might be different between the tables because the VPN table calculates measurements only for the service class associated with that VPN.

VPN nodes

VPN nodes are defined by the following 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 with one in the Nodes table.
- **Type**—The type of VPN. You can choose from the defaults (VPWS, VPLS, or L3VPN), or 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.

  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.
- **Description**—Description for the VPN node.
- **RT import** and **RT export**—The pairing of RT values identifies which VPN nodes connect with each other. For more information, see VPN topology and connectivity, on page 166.
- (Optional) **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

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **VPNs** > **VPN node**.

OR

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.

Step 3 Click .

Step 4 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.

Step 5 From the **Type** drop-down list, choose a VPN type. The options are: L3VPN, VPLS, and VPWS.

Step 6 In the **Name** field, enter the name of the VPN node, which does not have to be unique.

Step 7 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.

Step 8 (Optional) Enter a description that identifies the VPN node, for example, a customer name might be helpful.

Step 9 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 import RT as another VPN node's export RT can receive traffic from that VPN node. Those VPN nodes with the same export RT as another VPN node's import RT can send traffic to that VPN node.

Step 10 (Optional) In the **RD** field, enter a route distinguisher.

Step 11 Click **Add**.

Add VPN nodes to VPNs

To add VPN nodes to VPNs, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 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.

Step 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.

Step 4 Click **Save**.

VPN nodes table

The VPN Nodes table lists the VPN node properties, as well as columns that identify the VPN nodes' relationship within the VPN and its traffic.

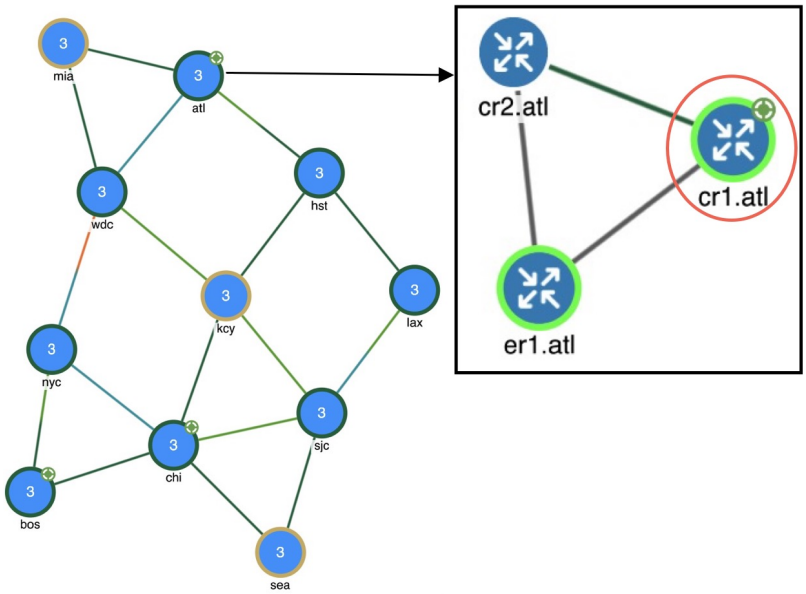
Table 19: VPN nodes table

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	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.

VPN nodes are not selectable from the network plot; you can only select and filter to them through tables.

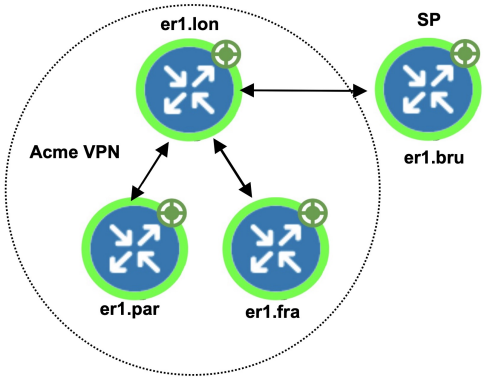
Once selected from the VPN Nodes or VPNs tables, the associated site and the nodes within that site appear with a green circle on it (Figure 57: VPN nodes within a VPN, on page 172).

Figure 57: VPN nodes within a VPN



Layer 3 VPN example

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).

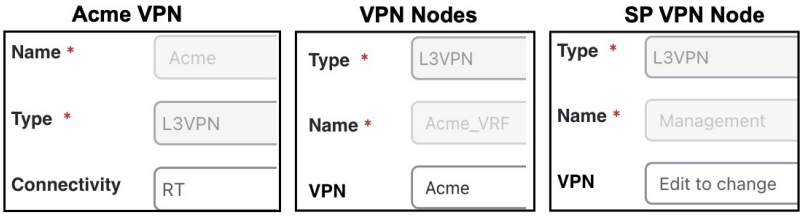


Additionally, headquarters communicates with an SP VPN node (er1.bru) that is not in the Acme VPN. Figure 58: Example RT connectivity and Acme VPN footprint, on page 173 shows the footprint of the Acme VPN and the RTs set for all VPN nodes in this example.

- 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 of these 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 that VPN.



The VPN footprint in Figure 58: Example RT connectivity and Acme VPN footprint, on page 173 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 59: Example failure between branch offices in the Acme VPN, on page 174, which shows that none of the demands associated with the VPN are rerouted.

Figure 58: Example RT connectivity and Acme VPN footprint

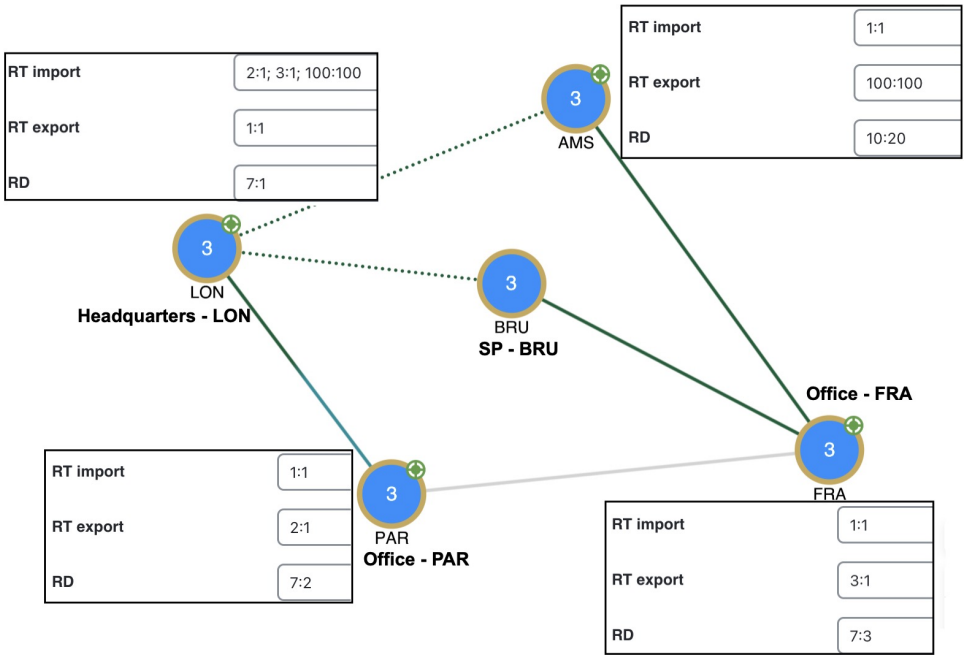
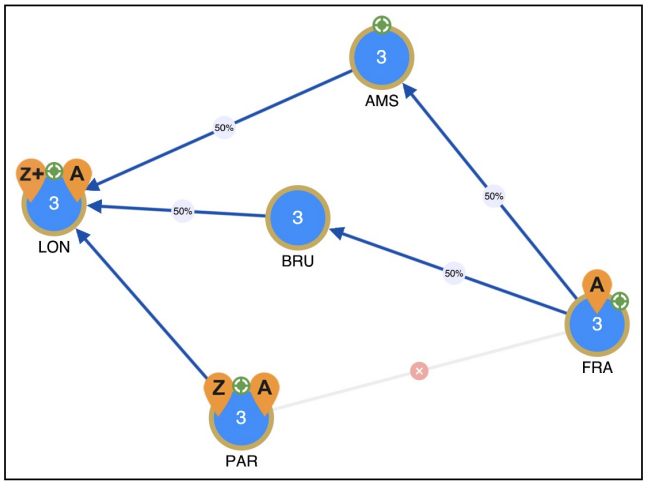


Figure 59: Example failure between branch offices in the Acme VPN



For this example, Figure 60: VPN nodes belonging to Acme VPN, and Acme VPN filtered to demands, on page 175 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.

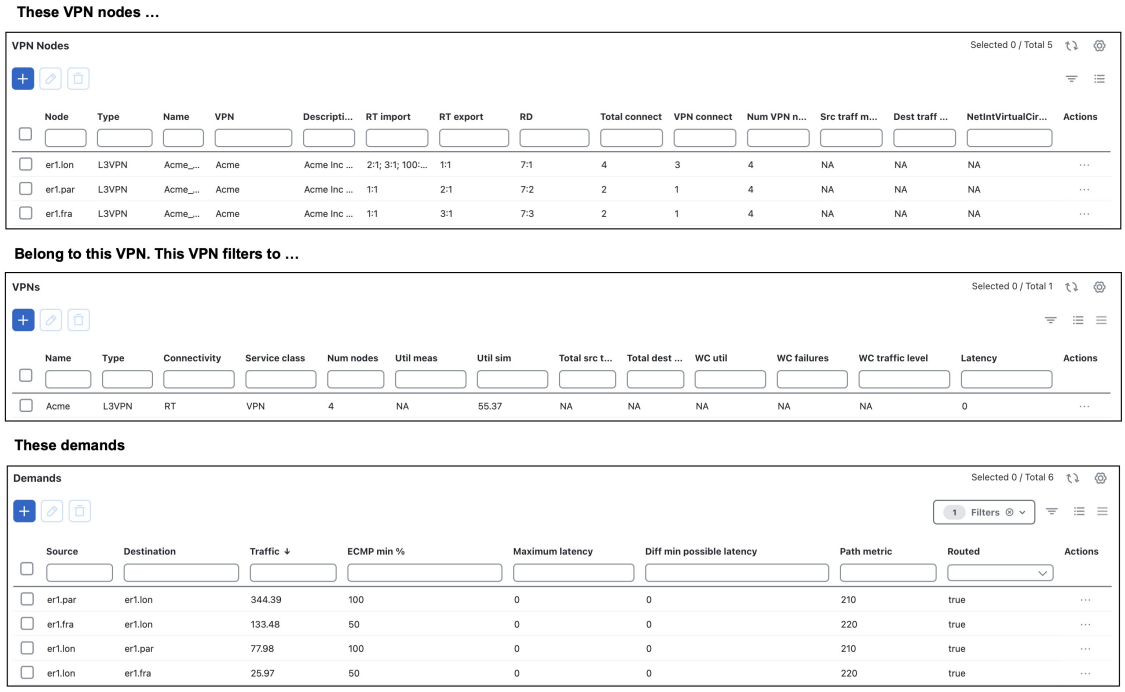
  Each of the offices and the service provider VPN node have 2 in the Total connect column.

- The VPN connect for the VPN node residing on 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.

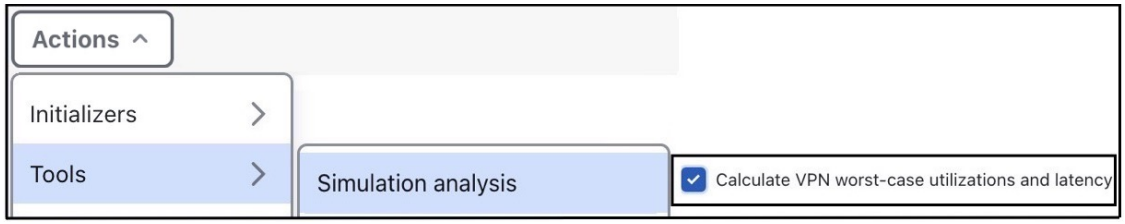
The service provider VPN node (er1.bru) has 0 VPN connects because it does not reside in a defined VPN.

Figure 60: VPN nodes belonging to Acme VPN, and Acme VPN filtered to demands



VPN simulation analysis

When running the **Simulation analysis** tool (from the toolbar, choose **Actions** > **Tools** > **Simulation analysis**), you have the option to record worst-case utilization and latency for VPNs in the VPNs table. You can then 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.

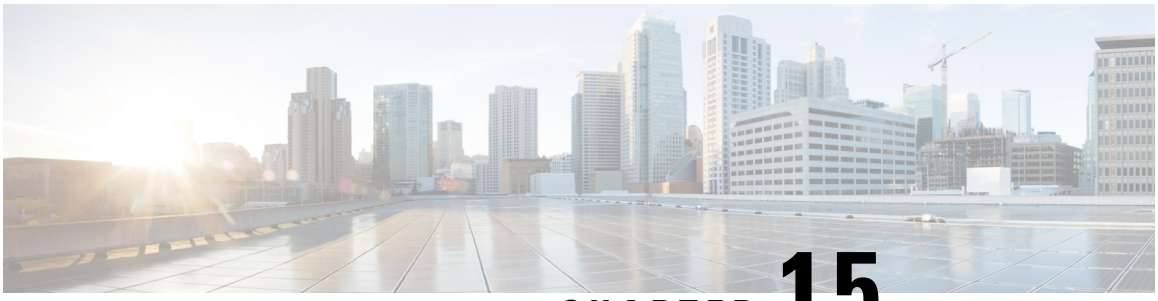


Note All calculations are based on traffic carried on all interfaces used by the VPN for the service class defined for that VPN.

The following columns are updated in the **VPNs** table upon finishing the Simulation analysis:

Table 20: Simulation analysis columns in the VPNs table

Columns	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.



CHAPTER 15

Simulate Advanced Routing with External Endpoints

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.

There are numerous use cases both for IGP and inter-AS routing:

- Simulate 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.
- Simulate 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.
- Simulate complex BGP routing policies from a transit provider. For example, you can specify a transit entry location and failover location per destination.
- Simulate failover between peering ASes; for example, from one single-connection transit provider to another.

This section contains the following topics:

- Routing with external endpoints, on page 177
- Create external endpoints and their members, on page 178
- Simulate routing, on page 180

Routing with external endpoints

An external endpoint is a Cisco Crosswork Planning object that identifies specific entry (source) or exit (destination) points for demands. These are identified in the External Endpoints table by a name.

Each external endpoint 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.

Create external endpoints and their members

The recommended method of creating external endpoint members is to do so while creating the associated external endpoint, as follows:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **Demands** > **External endpoints**.

OR

In the Network Summary panel on the right side, 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.

Step 3 In the **Name** field, enter a unique external endpoint name.

Step 4 To add a new member, click . To edit an existing member, select it from the Members table and click .

An External Endpoint Member window opens (see Figure 61: External endpoint members, on page 179). Specify the parameters as described in Specify external endpoint members, on page 178 and click **Add**.

Step 5 Click **Add** in the Add External Endpoint window.

Specify external endpoint members

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, and all members are listed in the **External endpoint members** table. For details on creating External endpoint members, see Create external endpoints and their members, on page 178.

Figure 61: External endpoint members



- **Member endpoint**—Defines whether the member is a node, AS, or external node via 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 node name within the AS. The "(None)" option routes traffic evenly throughout the entire AS.
- **Priority**—The sequential order in which external endpoint members are used in the simulations should failures occur.
- **Equal priority routing**—If members have the same priority, this property identifies how the traffic is distributed. Choose from the following options, as required:
  - **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 the same as Fix Traffic in that it sets 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. Thus, Deduce Traffic is usually set to Priority 1.

Simulate routing

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.

If a demand's source is defined as an external endpoint, the following selection of external endpoint members ensues.

Procedure

Step 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.

If one or more of the members are not available, the traffic from the unavailable members is evenly redistributed to the other top priority members.

Step 2 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. Only if all members with the same priority fail does the traffic get routed according to the next priority in the sequence.

Note that 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.

If an external endpoint member is an external AS, with or without a node specified, then the routing from or to that member is determined by the BGP routing policy determined 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

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.

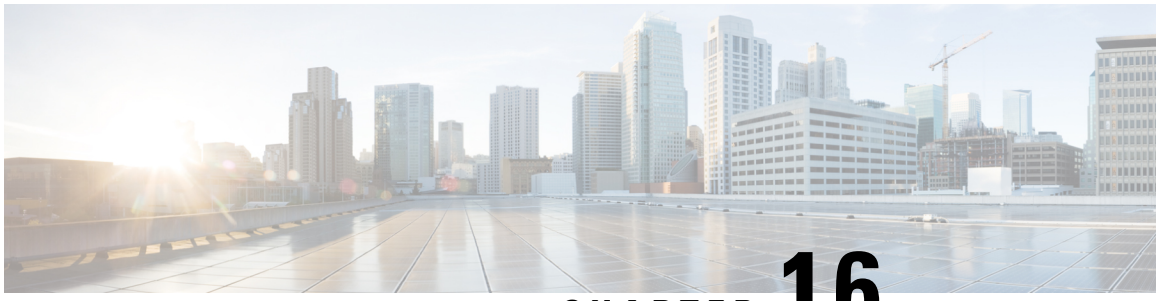
Of the routable demands, if the **Traffic balance (%)** values are all empty, the traffic is routed and equally load balanced across the demands with the shortest IGP paths. Note in the case of multiple internal ASes, the shortest IGP route is the shortest route in the first AS the demand enters.

- 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: Node A failed. Nodes B, C, and D each have a priority of 2 and are each a Fix Traffic type. Their traffic balances are 20%, 20%, and 40%, respectively. The demand has 1000 Mbps of traffic.

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

- Because 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.



CHAPTER 16

Simulate Multicast

Cisco Crosswork Planning supports Source-Specific Multicast (SSM), which is a method of delivering multicast packets to receivers only from a specified source address that is requested by the receiver. By limiting the source, SSM reduces resource requirements and improves security.

This section contains the following topics:

- Representation of SSM parameters , on page 183
- Discovered versus simulated multicast flows, on page 183
- Set global multicast simulation parameters, on page 184
- Multicast demands, on page 185
- Multicast flows, on page 186

Representation of SSM parameters

SSM 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

How you work with multicast flows within Cisco Crosswork Planning differs depending on whether the multicast flow was discovered by Cisco Crosswork Planning or whether you are simulating it (Table 21: Discovered versus simulated multicast flows , on page 183).

Table 21: 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.

	Discovered multicast flow	Simulated multicast flow
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 • Multicast Flow Hops • Multicast Flow External Hops	• Multicast Flows • Multicast Flow Destinations • Demands

Set global multicast simulation parameters

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 follow 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.

Cisco Crosswork Planning uses the network state in its multicast simulation. You can set simulations to take into account multicast flow hops, as follows.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the toolbar, click **Network options** or choose **Actions > Edit > Network options**.

Step 3 Click the **Simulation** tab.

Step 4 To use or disregard multicast flow hops in simulation, check or uncheck **Use multicast flow hops**, and click **Save**.

Cisco next hop

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.

Follow these steps to set Cisco Crosswork Planning to use Cisco next hops.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the toolbar, click **Network options** or choose **Actions > Edit > Network options**.

The Network options page opens.

Step 3 Click the **Protocols** tab.

Step 4 Under the **Multicast** section, from the **Multipath selection** drop-down list, choose **CiscoNextHop**.

Step 5 Click **Save**.

Multicast demands

For discovered multicast flows, you can insert demands manually. For details, see Create demands for multicast flows, on page 187.

For simulated multicast flows, create demands manually. For information on demands in general, see Simulate Traffic Flow from Source to Destination Using Demands, on page 71.

Simulated multicast demands

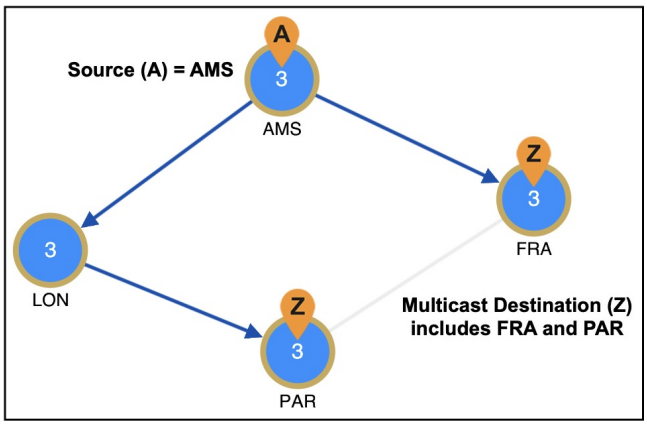
For manually inserted demands, all SSM demand traffic for the (S,G) pair (multicast flow) goes through any interface that is traversed by that demand. SSM demands are routed to take the shortest path from the destination to the source, rather than from the source to the destination, as unicast demands 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.

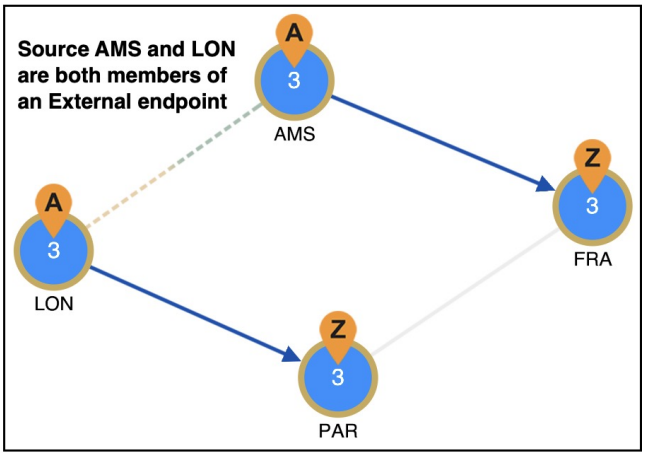
However, you can set Cisco Crosswork Planning to use the Cisco next-hop multicast, multipath selection method. For information, see Cisco next hop, on page 185.

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.



Multicast flows

View multicast flows

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
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

Follow these steps to create demands for multicast flows.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22) with Multicast details. The plan file opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **Demands** > **Demand**.

OR

In the Network Summary panel on the right side, click  > **Demands** in the **Demands** table.

Step 3 In the **Name** field, enter the name of the demand.

Step 4 In the **Source** area, 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.

Step 5 In the **Destination** area, 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).

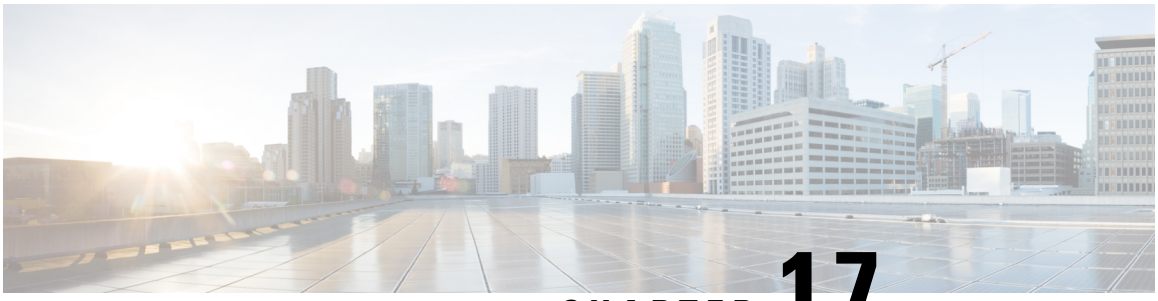
Step 6 (Optional) Complete all other fields as needed, and then click **Add**.



PART III

Traffic Engineering and Optimization

- Optimize Metrics in the Network Core, on page 191
- Configure MPLS Routing, on page 199
- Optimize LSPs, on page 213
- Configure RSVP-TE Routing, on page 229
- Optimize RSVP-TE Routing, on page 281
- Perform Explicit and Tactical RSVP-TE LSP Optimization, on page 291
- Configure Segment Routing, on page 301
- Optimize Segment Routing, on page 315



CHAPTER 17

Optimize Metrics in the Network Core

The **Metric optimization** tool optimizes metrics in the network core. It 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.

This section contains the following topics:

- Core versus edge, on page 191
- Perform metric optimization, on page 192
- Tactical metric optimization, on page 196
- Metric optimization reports, on page 196

Core versus edge

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 as follows:

- 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

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.

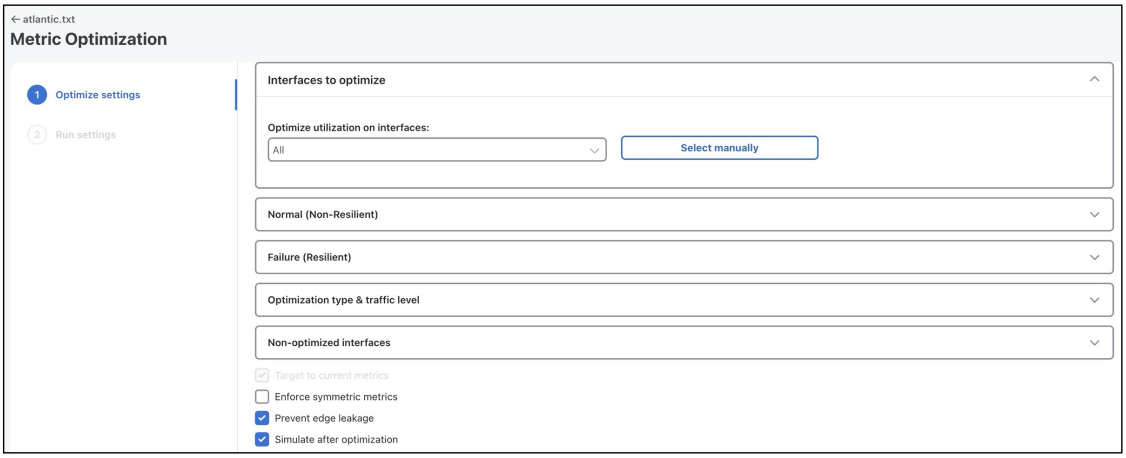
To run the Metric optimization tool, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). The plan file opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions > Tools > Metric optimization**.

Figure 62: Metric optimization options



Step 3 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.

Step 4 Expand the **Normal (Non-Resilient)** section and choose the optimization options to use. See Table 22: Metric optimization options, on page 193 for field descriptions.

Step 5 Expand the **Failure (Resilient)** section and decide on the options to use. See Table 22: Metric optimization options, on page 193 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.

Step 6 Specify the remaining options to fine-tune the optimization. See Table 22: Metric optimization options, on page 193 for field descriptions.

Step 7 Click **Next**.

Step 8 (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.

Step 9 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

  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.

Step 10 (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.

Step 11 Click **Submit**.

Table 22: Metric optimization options

Field	Description
Interfaces to optimize	

Field	Description
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.
Minimize # of interfaces with utilization > __ %	Often a small number of interfaces are bottlenecks in the network, and their utilizations 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 utilizations 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	

Field	Description
Optimization type	Cisco Crosswork Planning can perform global or incremental metric optimizations. • 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. • 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 following interfaces: all, core, interfaces selected manually.
Traffic level	The traffic level over which the optimization is performed.
Non-optimized Interfaces options	
Non-optimized Interfaces	For those interfaces that are not optimized, you can set the following: • Ignore • Minimize number of interfaces with util > __ % • 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.

Field	Description
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 utilizations under failure resulting from the new metric settings.

What to do next

See Metric optimization reports, on page 196.

Tactical metric optimization

The **Tactical metric optimization** tool (**Actions > Tools > Tactical metric optimization**) is a reduced version of the Metric optimization tool that minimizes the number of changes and runs faster. The options for this tool are a subset of those for the Metric optimization tools.

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. If the option is not set, the tool continues to run until it runs out of solutions to explore, which can take a long time. If you need all available solutions to be explored, use the **Metric optimization** tool.

Metric optimization reports

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.

A report is majorly categorized into the following sections:

Metric optimization report

Following information is available in the **Metric-Opt Report** section:

- 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. Following are some of the warning message examples:

- 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 the 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.

- **<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

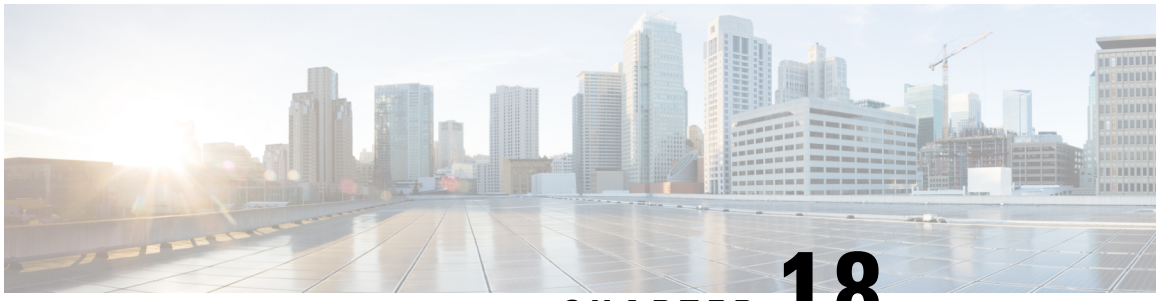
Summary statistics are displayed of the routes in network, both before (in parentheses) and after the optimization.

- 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.



CHAPTER 18

Configure MPLS Routing

This chapter describes how Cisco Crosswork Planning configures MPLS routing. All LSPs other than SR (segment routed) LSPs are routed like RSVP LSPs. For MPLS simulation information specific to these types of LSPs, see Configure RSVP-TE Routing, on page 229 and Configure Segment Routing, on page 301.

- 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.

This section contains the following topics:

- Supported LSP types, on page 199
- Create and visualize LSPs, on page 200
- LSP paths, on page 201
- Route demands through LSPs, on page 203
- Set global simulation parameters, on page 209
- Troubleshoot LSP simulation, on page 210

Supported LSP types

Cisco Crosswork Planning supports the following 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, see Configure Segment Routing, on page 301.
- RSVP LSPs—LSPs that are 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, see Configure RSVP-TE Routing, on page 229.

Note Cisco Crosswork Planning does not model LDP tunnels as LSPs.

Create and visualize LSPs

To simulate MPLS LSPs in the network, you must first set up the LSPs to be routed. Follow these steps to create and visualize LSPs in the Cisco Crosswork Planning UI.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **LSPs** > **LSP**

OR

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.

Step 3 Choose the **Type**, which determines whether this is an RSVP LSP or an SR LSP.

Step 4 In the **Name** field, enter the name of the LSP.

Step 5 Select the **Active** or **FRR enabled** check box.

Step 6 Expand the **Source & destination** panel. Choose the appropriate source and destination site and node details.

Step 7 (Optional) Expand the other panels (Routing, CSPF, and Other) and enter the relevant parameters.

Step 8 Click **Save**.

Step 9 To visualize the LSP in the network plot, select the LSP from the LSPs table. They appear in the plot as a purple arrow.

Figure 63: LSP visualization



You can also add a mesh of LSPs between the selected nodes using any of the following 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.

You can filter to related information, such as related interfaces, source and destination nodes, demands, or LSP paths. To do this, select the LSP, click the **Cross table filter** icon ( ), and choose the appropriate option.

LSP paths

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, see Configure RSVP-TE Routing, on page 229 and Configure Segment Routing, on page 301.

Path options and active path

Each LSP path has a **Path option** property in the Add/Edit LSP Path window. The LSP is routed using the first LSP path that can successfully be established. LSP paths are established in increasing order of their path option, where path option 1 is established first.



Alternatively, in the **Active path** field under the **Routing** section (Add/Edit LSP page), you can enter which LSP path to use.

Create LSP paths

To create LSP paths, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

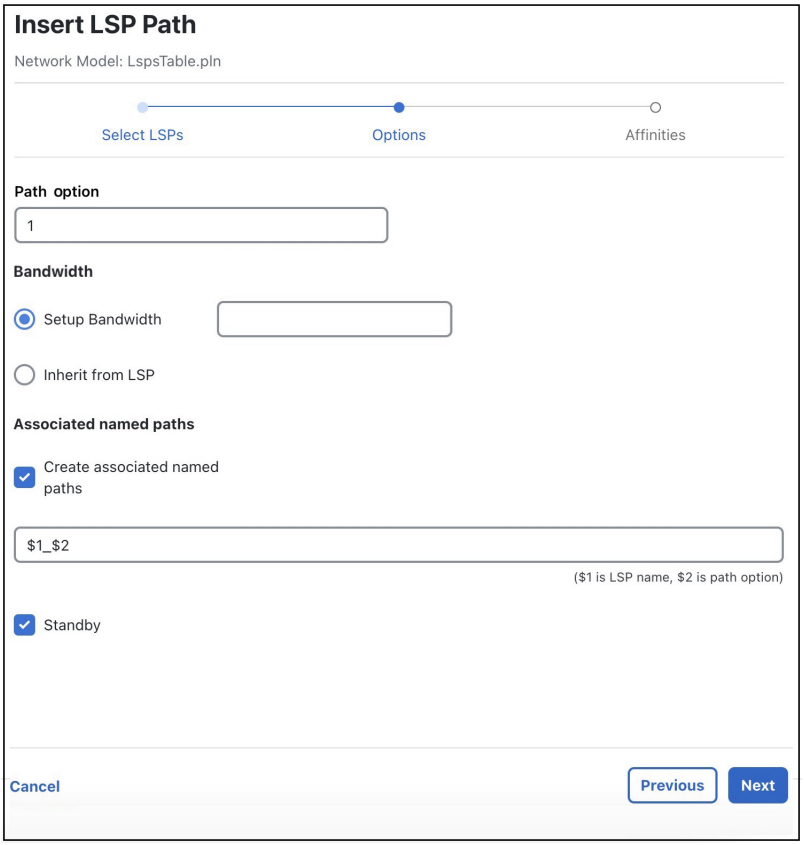
Step 2 In the Network Summary panel on the right side, from the **LSPs** table, choose the LSPs to which you are adding LSP paths.

Step 3 From the toolbar, choose **Actions > Insert > LSPs > LSP Paths**, or click  in the **LSP Paths** table.

The Insert LSP Path window appears.

Step 4 If you are satisfied with the LSPs that you selected in Step 2, click **Next**. Make changes to the selection, if required.

Step 5 In the **Path option** field, enter the order in which the LSP path is activated. Lower numbers have preference over higher numbers.



Step 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, see Configure RSVP-TE Routing, on page 229.

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.

Step 7 Click **Next**.

Step 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 below these options.

b. If you chose to use the table, choose the rule for each affinity you are associating with the LSP path.

Step 9 Click **Submit**.

Path latency calculations

Cisco Crosswork Planning calculates shortest latency paths for LSPs and demands as follows:

- 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

Route demands through intra-area 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.

Forwarding adjacency 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 (FA) settings, choose one or more LSPs from the LSP table, click  to open the Edit window, and edit the values in the **Routing** section.

Inter-area LSPs

Because 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.

To be routed, these demands must also match the privacy requirements. Autoroute, Forwarding adjacency (FA) properties, and LSP metrics are ignored. For information on privacy, see Route demands through specific LSPs (Private LSPs), on page 204.

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, see Simulate Traffic Flow from Source to Destination Using Demands, on page 71.

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.

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. (See Explicit versus dynamic inter-area LSP routing, on page 204.)

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.

Order of routing through areas

Regardless of the LSP type and whether explicit hops are required, inter-area LSPs route through the backbone areas, as follows, 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 versus dynamic inter-area LSP routing

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. One requires that explicit hops be set on ABR nodes. This mode correctly simulates actual router behavior where ABR explicit hops are required. The other mode does not require explicit hops at ABR nodes and can route an LSP fully dynamically across multiple areas. While this mode does not simulate actual router behavior, it is useful for planning inter-area LSP routes. These modes are specified using the network option labeled **LSP routing requires ABR explicit hops** under 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.

Route demands through specific LSPs (Private LSPs)

Cisco Crosswork Planning provides two ways to route selected traffic demands through specific LSPs. One way is to dedicate the traffic for specific demands to a private LSP, which is a special LSP that carries those

demands only. This type of LSP models an MPLS Layer 2 VPN, providing an exclusive route for the associated demands. If the LSP goes down, all traffic associated with the LSP is interrupted.

To configure LSPs that simulate Layer 2 VPNs, use one of these two tools.

- One tool creates dedicated demands for existing LSPs. The created demands will match the LSPs in source and destination. For details, see Create private demands for existing LSPs, on page 205.
- One tool creates private LSPs from existing demands. The created LSPs will match the existing demands in source and destination. For details, see Create private LSPs for demands, on page 206.

The LSP **Private** column is set to true in the LSPs table, and the **Private LSP Name** and **Private LSP Source** columns in the Demands table are set.

Create private demands for existing LSPs

Before you begin

Ensure that LSPs currently exist in the network model.

To create private demands for existing LSPs, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

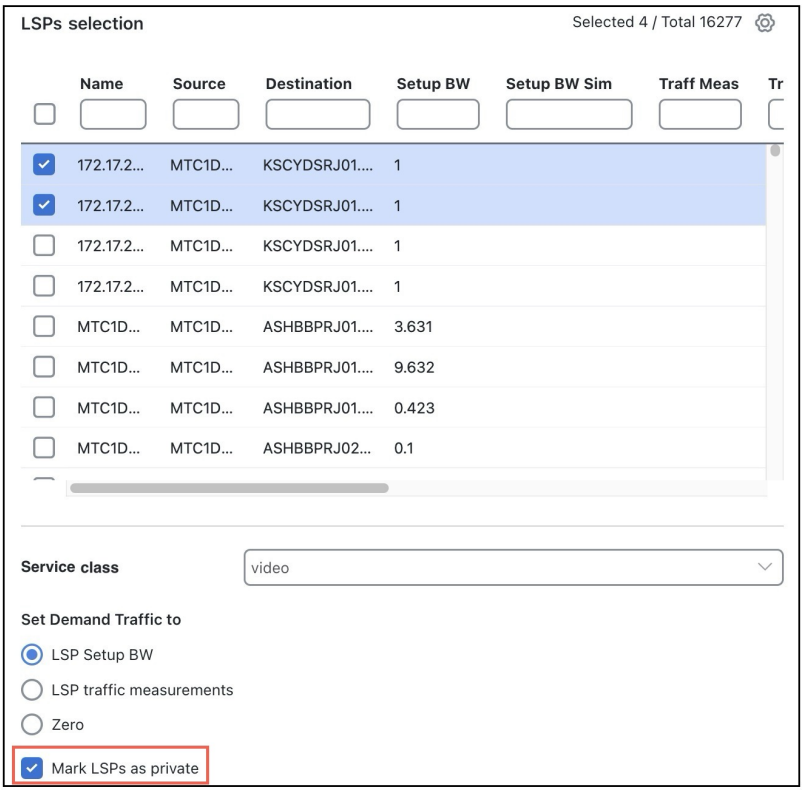
Step 2 From the toolbar, choose **Actions > Insert > LSPs > Demands for LSPs**.

OR

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



Step 3 From the LSPs list, select the LSPs for which you want to create demands.

Step 4 Select the service class to which these LSPs belong.

Step 5 Set the demand traffic to equal the LSP setup bandwidth, the LSP traffic measurements, or zero.

Step 6 Check the **Mark LSPs as private** check box.

Step 7 Click **Submit**. The newly created demands are highlighted in the Demands table.

Create private LSPs for demands

Before you begin

Ensure that demands currently exist in the plan file. See Simulate Traffic Flow from Source to Destination Using Demands, on page 71.

To create private LSPs for demands, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

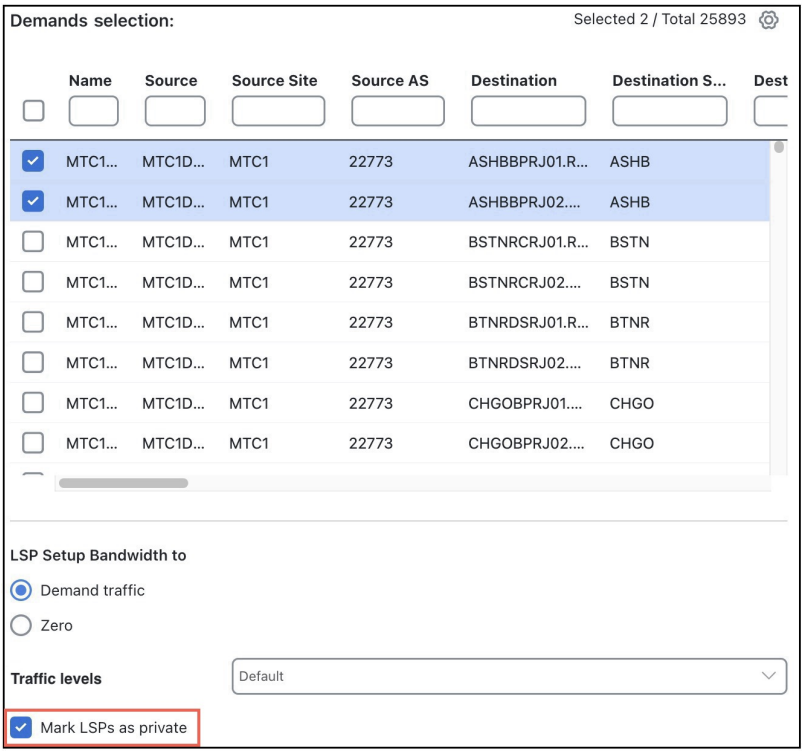
Step 2 From the toolbar, choose **Actions > Insert > LSPs > LSPs for demands**.

OR

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



Step 3 In the list of demands, select the demands for which you want to create LSPs.

Step 4 Set the bandwidth traffic to a specific demand traffic or to zero.

Step 5 Check the **Mark LSPs as private** check box.

Step 6 Click **Submit**. The newly created LSPs are highlighted in the LSPs table.

Delete demands when deleting private LSPs

You can choose to delete a demand when deleting a Private LSP. By default, when a private LSP is deleted, the corresponding demand is not deleted.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the toolbar, click **Network options** or choose **Actions > Edit > Network options**.

The Network options page opens.

Step 3 Click the **Advanced** tab.

Step 4 In the **Demands** section, ensure that these check boxes are checked.

- Unrouted if the private LSPs are unrouted
- Removed if the private LSPs are removed

Step 5 Click **Save**.

Load sharing between LSPs

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.

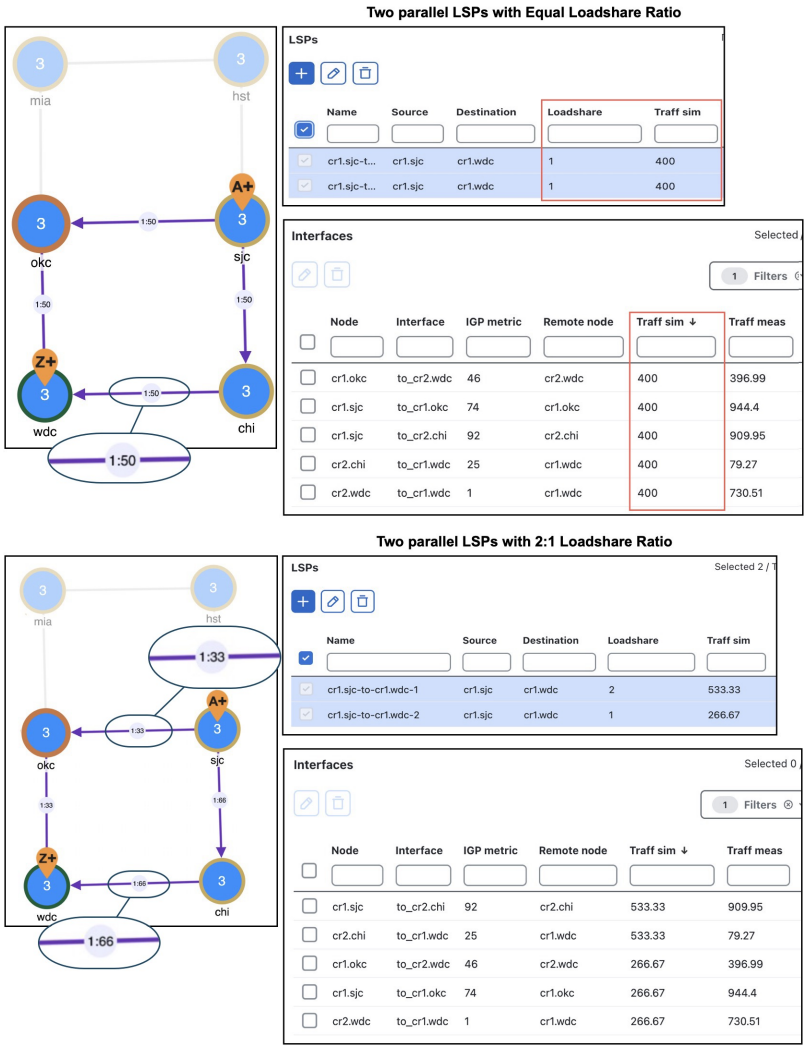
Example: If two LSPs are parallel, and one has a Loadshare property of 2 and one has a Loadshare property of 1, there will be a 2-to-1 ratio of traffic shared between them. The top half of Figure 66: Examples of two parallel LSPs with equal loadsharing and 2:1 loadsharing, on page 209 shows an example 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%.

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.

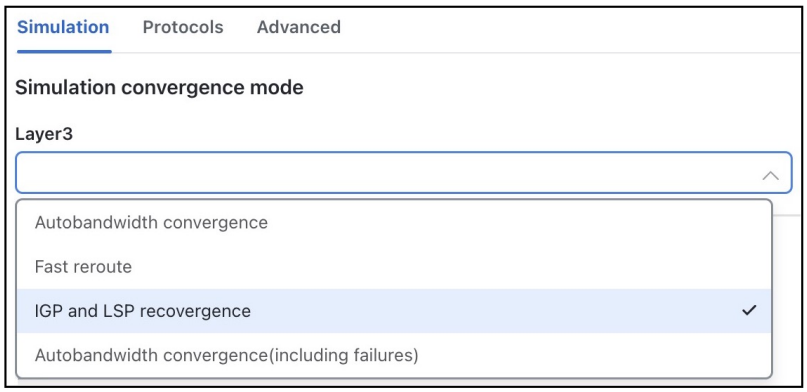
To optimize these Loadshare values, use the LSP Loadshare Optimization tool. For information, see Optimize LSP loadshare, on page 219 .

Figure 66: Examples of two parallel LSPs with equal loadsharing and 2:1 loadsharing



Set global simulation parameters

Cisco Crosswork Planning lets you set global parameters that affect how LSPs are routed or rerouted. To access these options, click **Network options** or choose **Actions** > **Edit** > **Network options** in the toolbar. Then, click the **Simulation** tab.



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 information, see Configure RSVP-TE Routing, on page 229.

The Optimization tools only work in IGP and LSP reconvergence mode. If you try to run one while in a different convergence mode, you are prompted whether to continue, and if you do, the simulation changes to IGP and LSP reconvergence mode.

Set LSP establishment order

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. Varying the LSP establishment order lets you check, for example, whether certain orders result in higher utilizations.

Procedure

Step 1 In the toolbar, click **Network options** or choose **Actions > Edit > Network options**.

The Network options page opens.

Step 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.

Step 3 Click **Save**.

Troubleshoot LSP simulation

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.

These routing diagnostics 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

To run LSP simulation diagnostics, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions > Tools > Diagnostics > LSP simulation** or **LSP path simulation**.

Step 3 Select the LSPs or the LSP paths you want to optimize.

Step 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.

What to do next

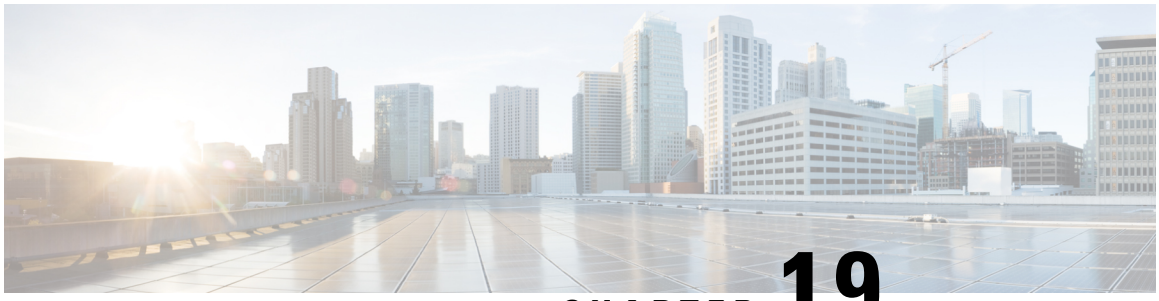
See Use simulation diagnostics to troubleshoot, on page 211.

Use simulation diagnostics to troubleshoot

The following information is provided to help you troubleshoot LSP and LSP path issues.

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.

Reason	Description
No Destination	There is no destination defined. The collected network might not contain the destination node of the LSP.
Simulation Option	The simulation option for following 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.



CHAPTER 19

Optimize LSPs

This section contains the following topics:

- Optimize disjoint LSP paths, on page 213
- Optimize LSP loadshare, on page 219
- Optimize LSP setup bandwidth, on page 224

Optimize disjoint LSP paths

LSPs and LSP paths are *disjoint* if they do not route over common objects, such as interfaces, or nodes. The **LSP disjoint path optimization** tool (**Actions > Tools > LSP optimization > LSP disjoint path optimization**) creates disjoint LSP paths for RSVP LSPs and SR LSPs, and optimizes these paths based on user-specified constraints.

Note The LSP disjoint path optimization tool supports both Inter-Area and Inter-AS functionalities.

One common use case for this optimization is that disjoint LSPs can ensure that the services are highly resilient when network failures occur. 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.

Upon completion, by default, Cisco Crosswork Planning tags the LSPs with *DSJOpt* and writes a report containing the results of the optimization.

Note While the optimizer applies to both RSVP LSPs and SR LSPs, only one of these types of LSPs can be optimized at a time.

Specify optimization inputs

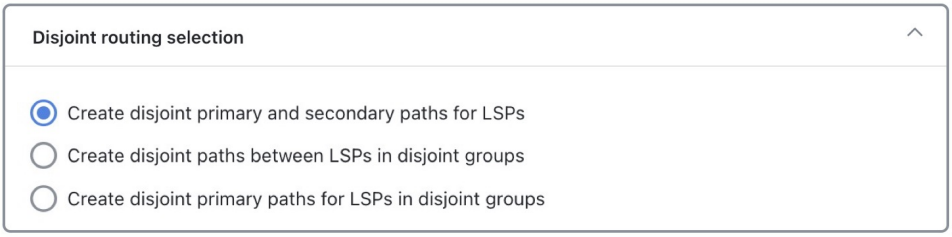
Disjoint routing selection

In Cisco Crosswork Planning, only existing LSP paths are rerouted. New LSP paths are not created.

- Explicit hops are modified or created for RSVP LSPs.
- Segment list hops are modified or created for SR LSPs. The final hop is either a node hop or interface whose remote node is the destination of the LSP.

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.

Figure 67: Disjoint routing selection options



Routing options include:

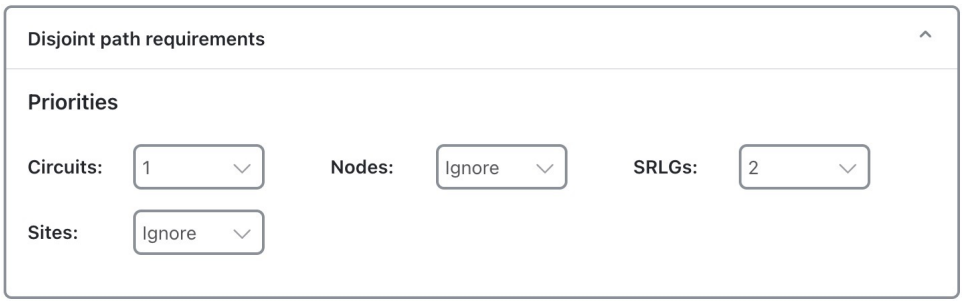
- **Create disjoint primary and secondary paths for LSPs**—For all LSPs, whether they are in a disjoint group or not, routes all LSP paths so that they are disjoint from all other paths belonging to that LSP. This disjointness extends beyond primary and secondary paths to include all other path options (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).

  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.

Disjoint path requirements

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 68: Disjoint path requirements



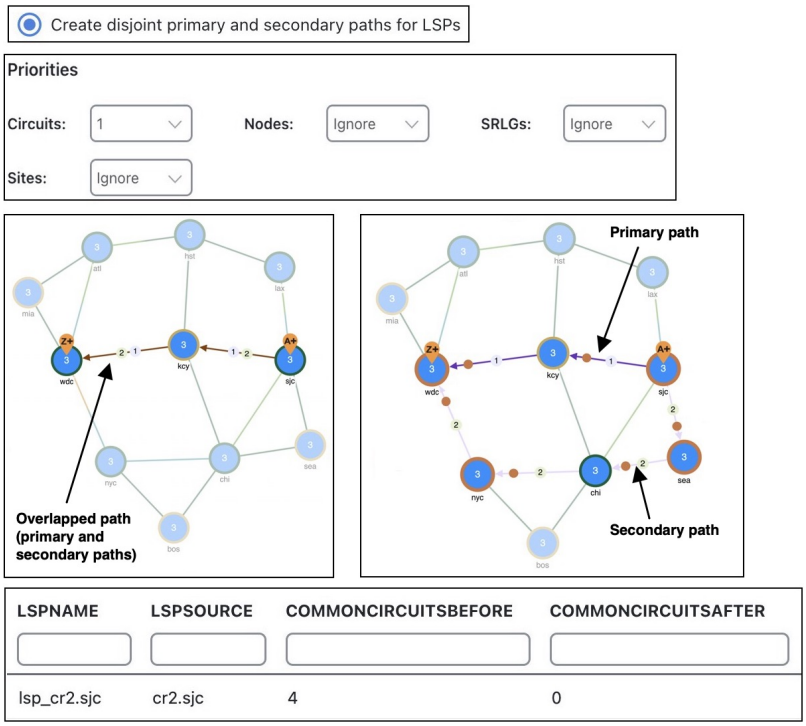
For example, in Figure 68: Disjoint path requirements, on page 215, Circuits have a priority of 1, SRLGs have a priority of 2, and the other objects are ignored. If the tool cannot achieve full disjointness across both circuits and SRLGs, it prioritizes the disjointness of circuits over SRLGs.

Example

This example shows how disjoint routes can be created for primary and secondary RSVP LSP paths, and how those routes differ, depending on the path requirements set. 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 69: Primary and secondary paths based on disjoint circuit requirements, on page 215 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. Because the disjoint path requirement is only circuits, the primary and secondary paths route across different circuits, as indicated in the resulting report.

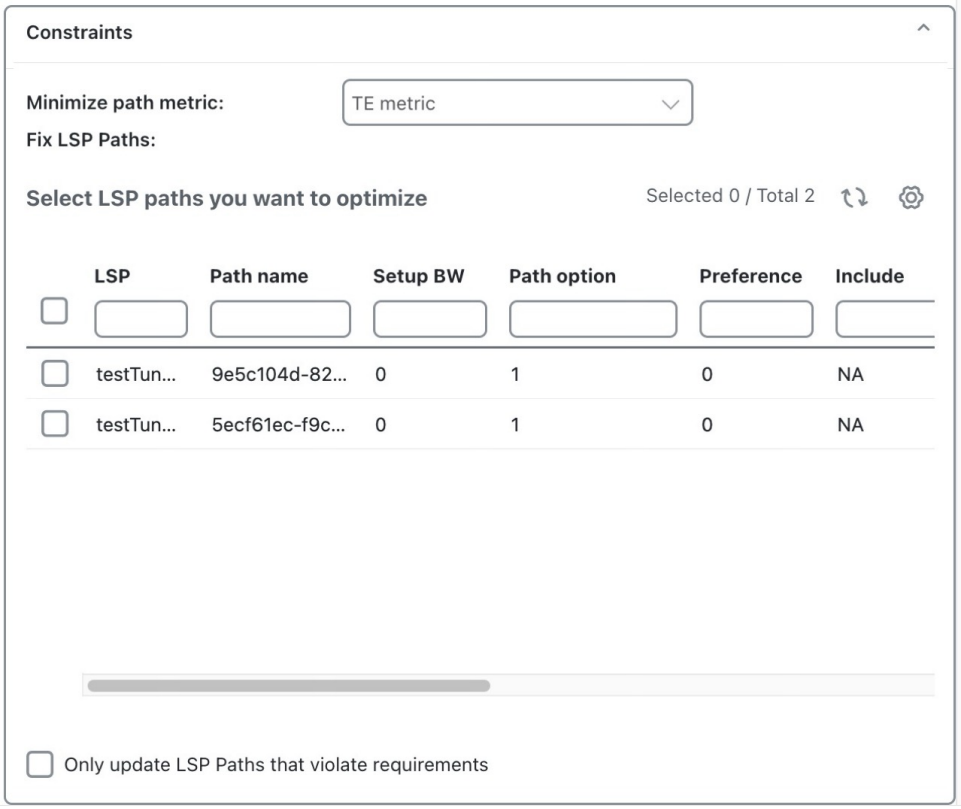
Figure 69: Primary and secondary paths based on disjoint circuit requirements



Constraints

Following options are available in the **Constraints** section:

- **Minimize path metric**—Paths are optimized to minimize the sum of the metrics along the path with respect to delay, TE, or IGP metrics. All of these properties are configurable from an interface Properties window, and delays can also be set in a circuit Properties window.
- **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 214 window.



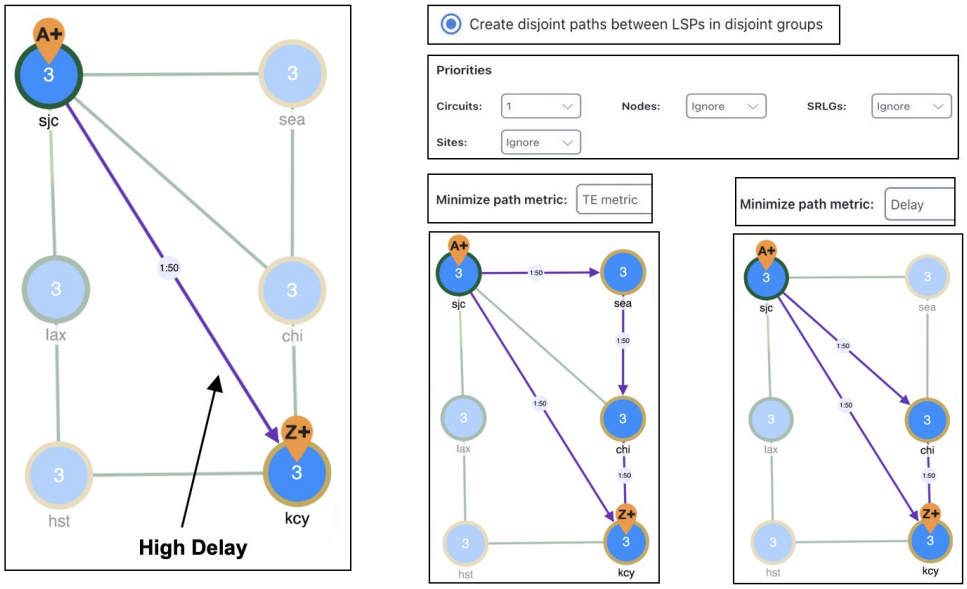
Example

This example shows how disjoint routes can be optimized for two SR LSPs using the same source (sjc) and destination (kcy).

- 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 70: Example routing of SR LSPs, on page 217 shows the following:
  - Before using the LSP Disjoint Path Optimization tool, both LSP paths use the same route. There are no segment list hops.

- By selecting the option to create disjoint paths between LSPs in the same disjoint group and using the TE metric for the shortest path calculation, two disjoint LSP path routes are created. Both have a segment list node hop (as indicated by the orange circle around the node) on the destination node. One has an additional segment list node hop on sea to force a different route.
- By selecting the same disjoint option and using the Delay metric for the shortest path calculation, one LSP is moved away from the high-delay sjc-kcy circuit because traversing that circuit is not the shortest latency path.

Figure 70: Example routing of SR LSPs



Create disjoint groups

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 creating disjoint paths between LSPs in the same disjoint group, the LSPs must first be added to the disjoint groups.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the Network Summary panel on the right side, choose one or more LSPs from the **LSPs** table.

Step 3 Click .

Step 4 Click the **Advanced** tab.

Step 5 Expand the **Explicit route selection** panel and enter the Disjoint group name.

Step 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.

Step 7 Click **Save**.

Run the LSP disjoint path optimization tool

To run the LSP disjoint path optimization tool, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose any of the following options:

- **Actions > Tools > LSP optimization > LSP disjoint path optimization**

  OR

- **Preset workflows > Perform optimization**, select **LSP Optimization** as the optimization type, choose **LSP disjoint path optimization** from the drop-down list, and click **Launch**.

Step 3 Select the LSPs for which you want to optimize the disjoint paths.

Step 4 Click **Next**.

Step 5 In the **Disjoint routing selection** section, select how to route the disjoint paths. For details, see Disjoint routing selection, on page 214.

Step 6 In the **Disjoint path requirements** section, select the disjoint path requirements and priorities. For details, see Disjoint path requirements, on page 214.

Step 7 In the **Constraints** section, select the constraints. For details, see Constraints, on page 216.

Step 8 Click **Next**.

Step 9 (Optional) In the **Tag updated LSP Paths with** field on the **Run Settings** page, override the defaults for how LSP paths are tagged (*DSJopt*).

Step 10 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

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.

Step 11 (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.

Step 12 Click **Submit**.

What to do next

See Disjoint report, on page 219.

Disjoint report

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. Note that new reports replace the previous ones.

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.

Optimize LSP loadshare

The **LSP loadshare optimization** tool automates the process of finding and setting the most favorable loadshare ratios across parallel LSPs to balance traffic and avoid congestion. The optimizer only includes interfaces that use parallel LSPs in the optimization. You can further limit parallel LSP interfaces to only those on which you want to drive down the maximum utilization.

Upon completion, Cisco Crosswork Planning tags the LSPs with *LSPLoadshare* and generates a report identifying how many LSPs and interfaces were affected, the number of bins used, the resulting maximum interface utilization, and if applicable, the number of interfaces with utilization over the specified threshold.

Run LSP loadshare optimization

To run the LSP loadshare optimization tool, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose any of the following options:

- **Actions > Tools > LSP optimization > LSP loadshare optimization**

  OR

- **Preset workflows > Perform optimization**, select **LSP Optimization** as the optimization type, choose **LSP loadshare optimization** from the drop-down list, and click **Launch**.

Step 3 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.

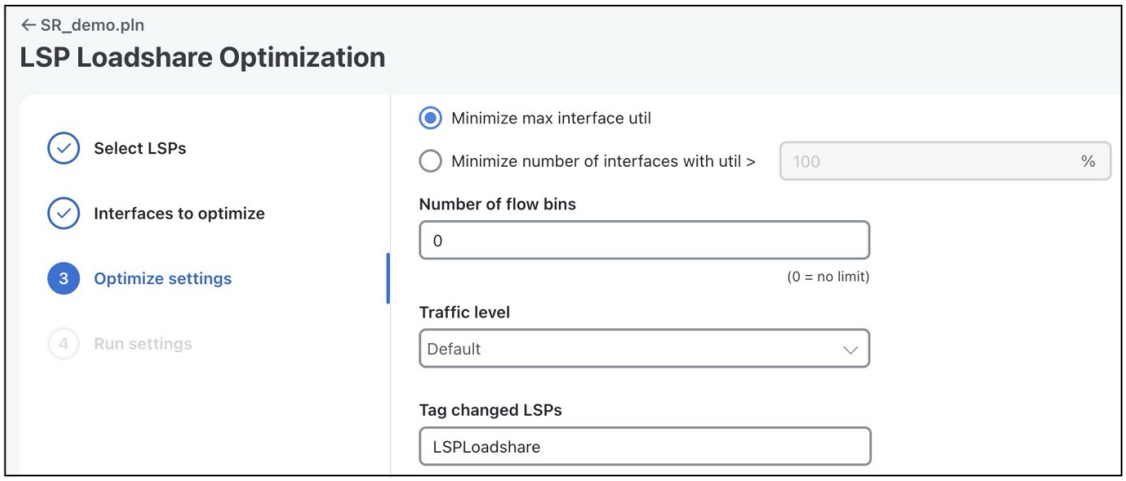
The optimizer does not set a loadshare value of 0 on LSPs.

Step 4 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.

Step 5 Specify the relevant optimization settings. For field descriptions, see Table 23: LSP loadshare optimization settings, on page 222.

Figure 71: LSP loadshare optimization options



Step 6 (Optional) In the **Tag changed LSPs** field, override the defaults for how LSPs are tagged (*LSPLoadshare*).

Step 7 Click **Next**.

Step 8 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

  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.

Step 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.

Step 10 Click **Submit**.

Table 23: LSP loadshare optimization settings

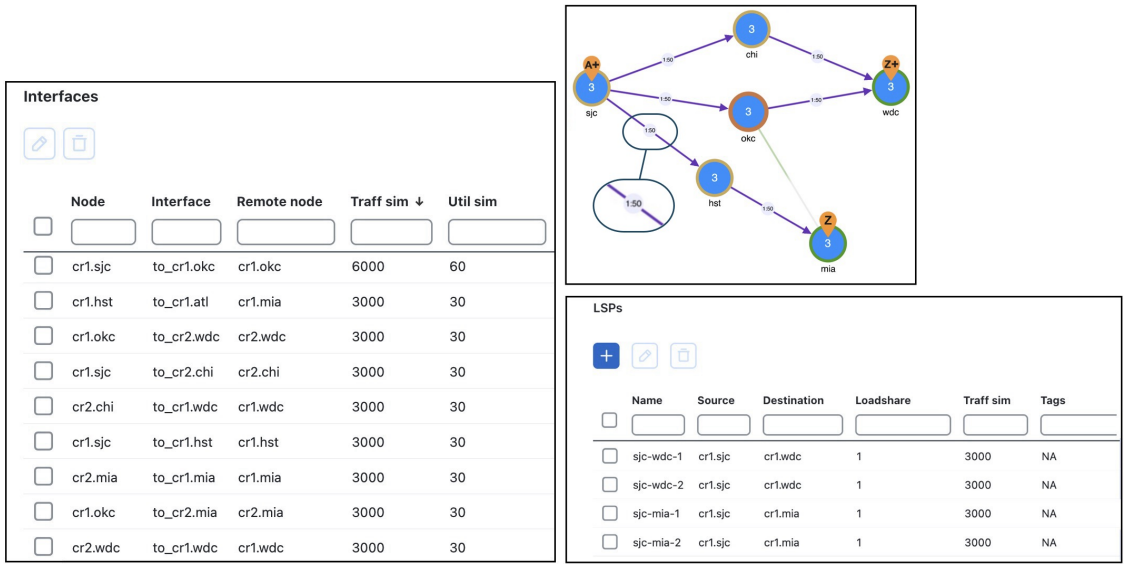
Field	Description
Minimize max interface util	Minimizes the maximum interface utilization over all interfaces on the LSP routes. Cisco Crosswork Planning tries to minimize the number of loadshare parameter changes required to do so.
Minimize number of interfaces with util >___%	Minimizes the number of 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 turns 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 example

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 72: Example Acme network before LSP loadshare optimization, on page 223).

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 73: Example Acme network after minimizing maximum interface utilization, on page 224). 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 72: Example Acme network before LSP loadshare optimization



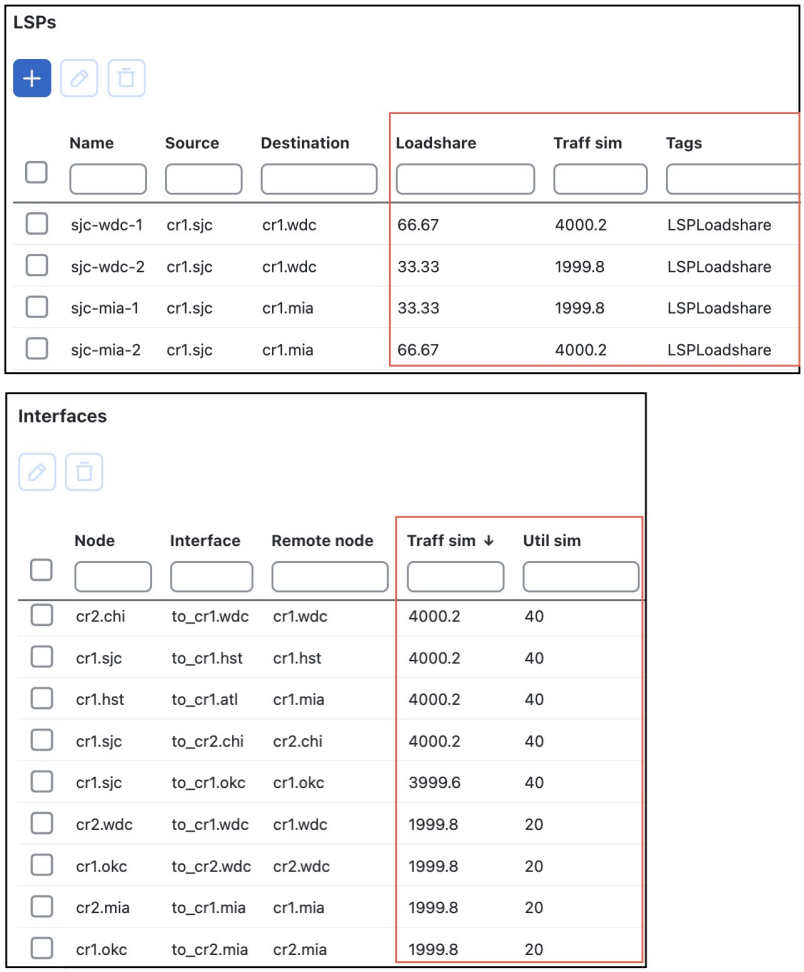
Interfaces

Node	Interface	Remote node	Traff sim ↓	Util sim
cr1.sjc	to_cr1.okc	cr1.okc	6000	60
cr1.hst	to_cr1.atl	cr1.mia	3000	30
cr1.okc	to_cr2.wdc	cr2.wdc	3000	30
cr1.sjc	to_cr2.chi	cr2.chi	3000	30
cr2.chi	to_cr1.wdc	cr1.wdc	3000	30
cr1.sjc	to_cr1.hst	cr1.hst	3000	30
cr2.mia	to_cr1.mia	cr1.mia	3000	30
cr1.okc	to_cr2.mia	cr2.mia	3000	30
cr2.wdc	to_cr1.wdc	cr1.wdc	3000	30

LSPs

Name	Source	Destination	Loadshare	Traff sim	Tags
sjc-wdc-1	cr1.sjc	cr1.wdc	1	3000	NA
sjc-wdc-2	cr1.sjc	cr1.wdc	1	3000	NA
sjc-mia-1	cr1.sjc	cr1.mia	1	3000	NA
sjc-mia-2	cr1.sjc	cr1.mia	1	3000	NA

Figure 73: Example Acme network after minimizing maximum interface utilization



Bin example

In this example, a node 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. Thus, 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.

Optimize LSP setup bandwidth

As a network operator, you might need to periodically update LSPs as traffic changes in your network. You might 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 failures 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

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

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose any of the following options:

- **Actions > Tools > RSVP LSP optimization > LSP setup BW optimization**

  OR

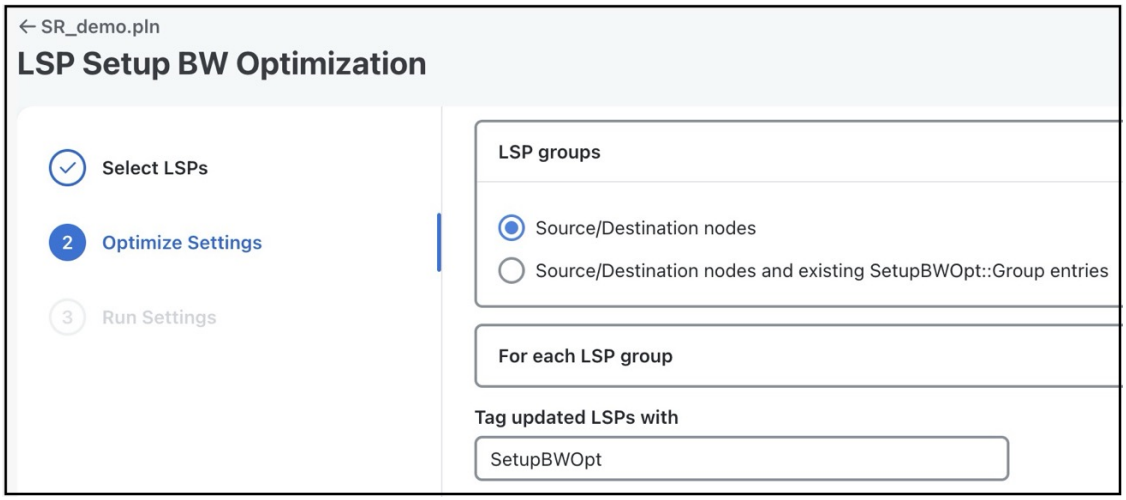
- **Preset workflows > Perform optimization**, select **RSVP LSP Optimization** as the optimization type, choose **LSP setup BW optimization** from the drop-down list, and click **Launch**.

Step 3 Choose the LSPs you want the optimizer to consider. By default, none of the LSPs are selected.

Step 4 Click **Next**.

Step 5 In the **LSP groups** section, choose the manner in which LSPs are separated into LSP groups:

Figure 74: 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.

Step 6 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.

Step 7 (Optional) In the **Tag updated LSPs with** field, override the defaults for how LSPs are tagged (*SetupBWOpt*).

Step 8 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

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.

Step 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.

Step 10 Click **Next**.

What to do next

See LSP setup BW optimization report, on page 227.

LSP setup BW optimization report

Each time the LSP setup BW 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 Setup BW Optimization** link in the right panel.

The report contains the following two tabs:

Summary

The **Summary** tab in the report summarizes the total number of:

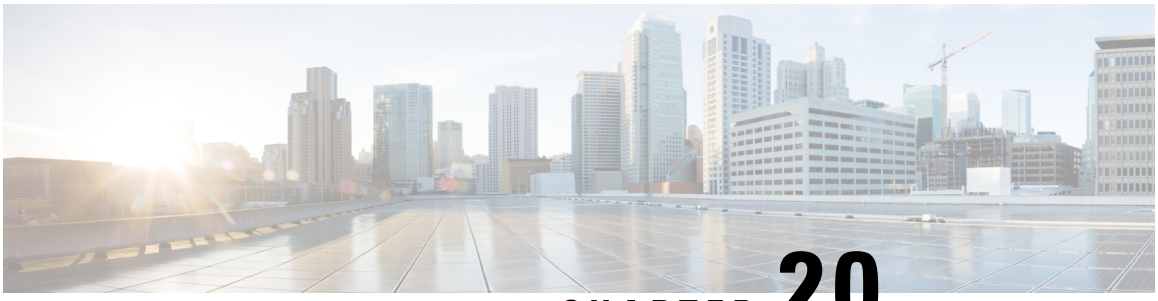
- LSPs you selected.
- LSP groups identified.
- LSPs created.
- LSPs removed.

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.



CHAPTER 20

Configure RSVP-TE Routing

This chapter 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.

This section contains the following topics:

- Dynamic LSP routing and CSPF, on page 229
- RSVP LSP paths , on page 234
- Named paths and explicit LSP routing, on page 235
- Actual paths, on page 238
- Affinities , on page 240
- Set global simulation parameters, on page 245
- Circuit Style RSVP LSPs, on page 256
- Advanced RSVP-TE LSP simulations, on page 269

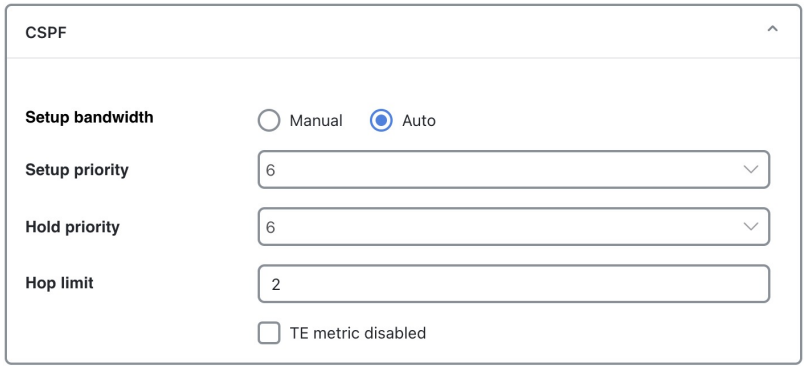
Dynamic LSP routing and CSPF

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, the following selection criteria apply in this order:

- Bandwidth available—The route with the greatest reservable bandwidth is chosen.
- Hop count—The route with the fewest hops is chosen.
- Random—If neither of the above criteria can be used, the route is randomly chosen.

CSPF properties are set in the **Edit LSP** window and are viewable from the LSPs table. Note that these properties are available only for RSVP LSPs.



- Setup bandwidth (Manual)—The amount of traffic the source node requests for this LSP in Mbps. The requested bandwidth is available from the reservable bandwidth of each interface in the path.
- Setup bandwidth (Auto)—Dynamically update the Setup BW Sim value when using the Auto bandwidth convergence mode.
- Setup priority—A priority for allocating reservable bandwidth. This is the order in which the LSPs are signaled. The lower the number, the higher the priority.
- Hold priority—Priority that can preempt LSPs that are 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—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—If checked, the LSP is routed using IGP metrics. If unchecked (default), the LSP is routed using TE metrics.

Set interface MPLS properties

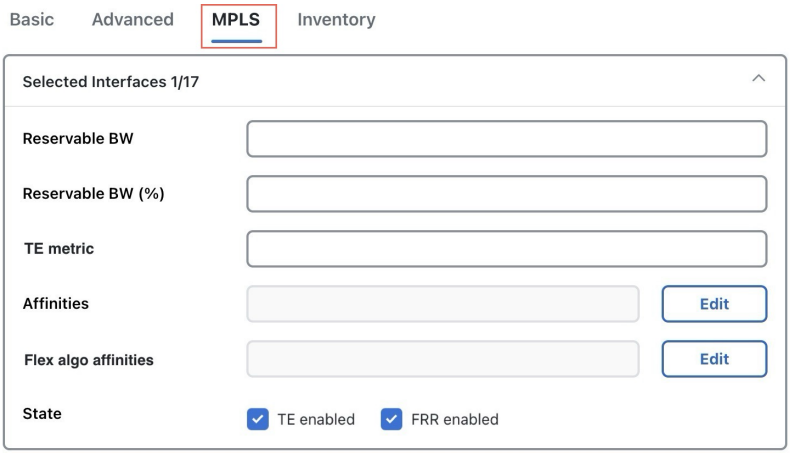
MPLS properties are set in the **MPLS** tab of the interface or circuit properties window and are viewable in the Interfaces or Circuits table.

1. Open the plan file (see Open plan files, on page 22). 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 under the **Actions** column.

4. Click the **MPLS** tab.

Figure 75: Interface MPLS properties



Set the following properties, as required:

- 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, see Affinities , on page 240.
- 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, see Fast Reroute simulations, on page 246.

Following 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, see Ignore reservable bandwidths for capacity planning, on page 269.

- 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.

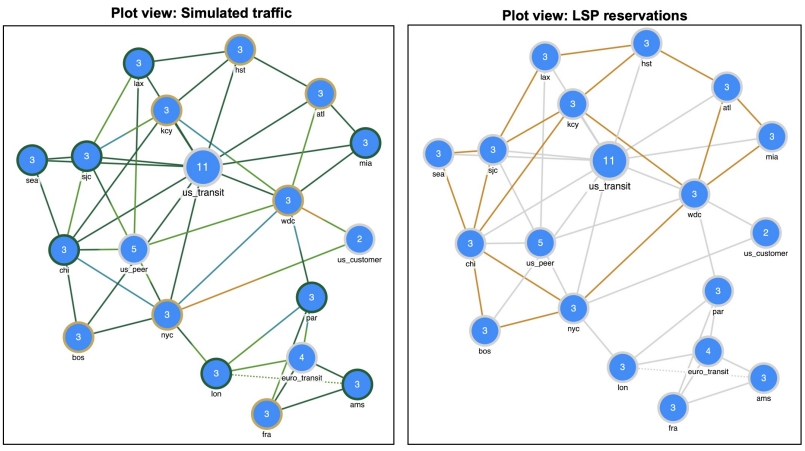
  To determine which interfaces carry LSP traffic, sort the columns. To determine the demands associated with one or more selected interfaces based on whether or not they have LSP traffic, select them and choose ☰ > **Filter to demands > Through all interfaces**. Then, choose ☰ > **Filter to LSPs**.
- 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

To view the LSP reservations on the interfaces, choose **LSP reservations** from the **Plot view** drop-down list in the toolbar. The utilization colors are based on the **LSP util** column in the Interfaces table. The LSP util is the sum of all setup bandwidths for all LSPs through the interface, as a percentage of the Resv BW sim.

Figure 76: Simulated traffic versus LSP reservations, on page 232 shows how the network plot varies based on the Plot view option you choose, Simulated traffic or LSP reservations.

Figure 76: Simulated traffic versus LSP reservations



Calculate LSP setup bandwidth

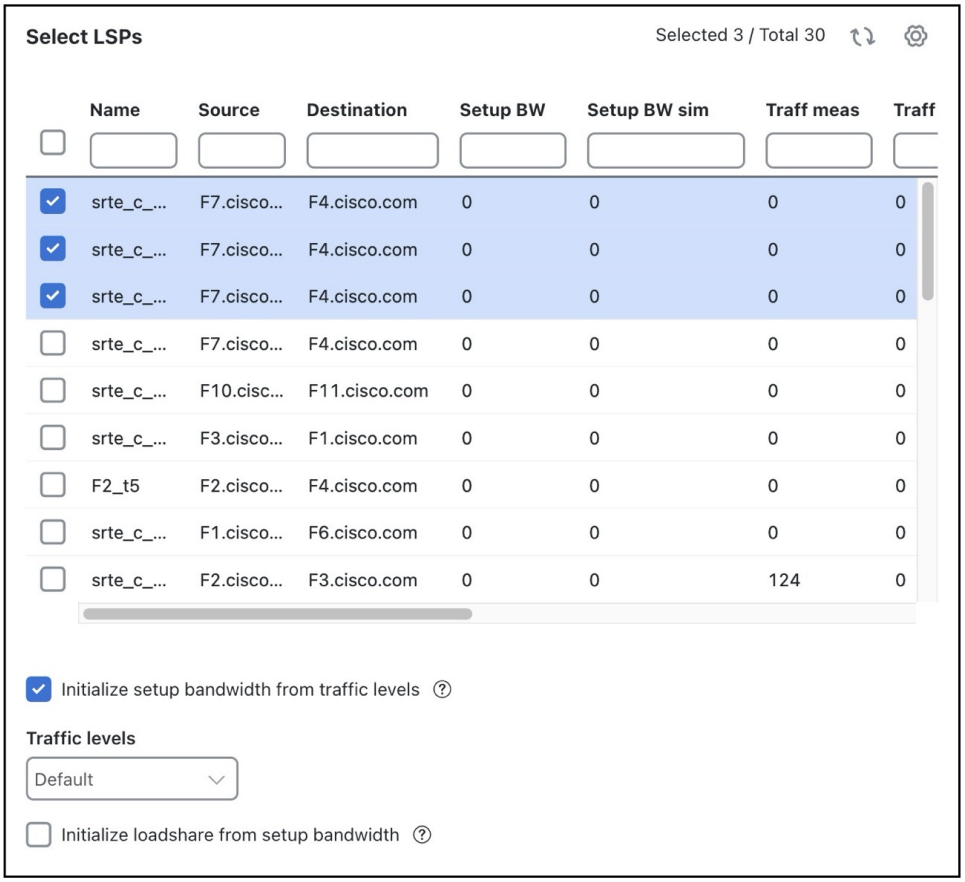
The **LSP setup bandwidth** initializer calculates the setup bandwidth for existing RSVP-TE LSPs. The setup bandwidth is set to the maximum amount of demand traffic passing through the LSP for the traffic levels selected.

The best practice is to use the **Autobandwidth convergence** mode (**Network options** > **Simulation** > **Simulation convergence mode**). For more information, see Simulate autobandwidth-enabled LSPs, on page 252

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Initializers** > **LSP setup bandwidth**.



Step 3 Select the LSPs you want to optimize.

Step 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.

Step 5 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.

Step 6 Click **Submit**.

RSVP LSP paths

Like LSPs, LSP paths have CSPF properties, such as **Setup priority** and **Hop limit**. For a description of LSP properties, see Dynamic LSP routing and CSPF, on page 229.

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.

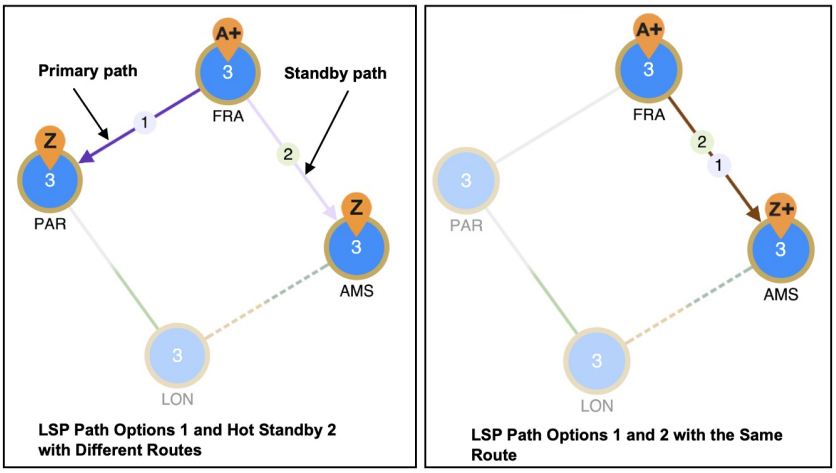
Note If operating in the **Autobandwidth convergence** simulation mode and if the LSP's **Setup bandwidth** is set to **Auto**, the LSP path Setup BW property (like the LSP Setup BW) is ignored, and the LSP's Setup BW sim value is calculated and used.

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.

Hot standby paths

An LSP path can be defined as a "hot" standby path. Standby paths are always established even if the associated LSP is routed using a path with a different path option. For RSVP LSPs, setup bandwidth, if any, is reserved for them. These standby LSP paths are then immediately available should the currently routed path become unavailable.

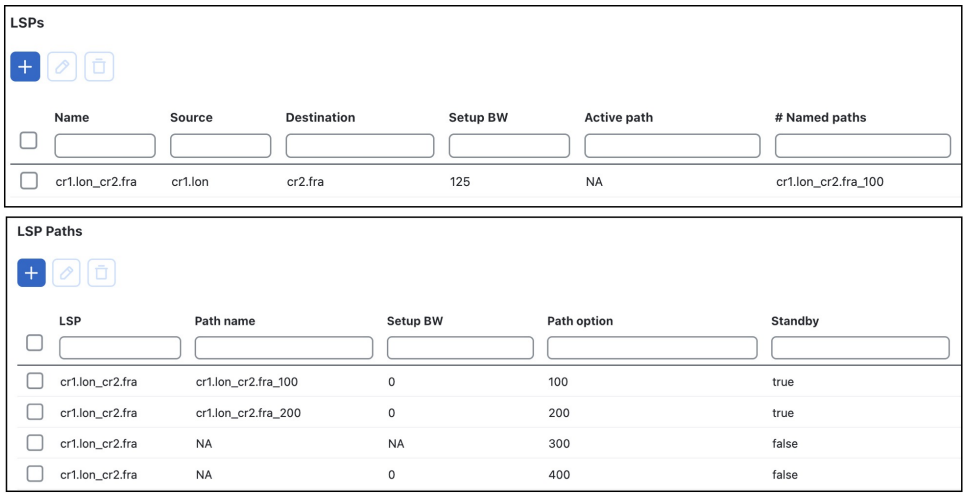


Example response to failures

Figure 77: Example RSVP LSP and associated LSP paths, on page 235 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.

Figure 77: Example RSVP LSP and associated LSP paths



This example LSP responds to failures as follows:

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

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, and each hop type specifies whether the route should be strict, loose, or excluded. Named path hops can be nodes or interfaces. The hop type is identified visually on the plot, as well as listed in the **Type** column of the **Named path hops** table.

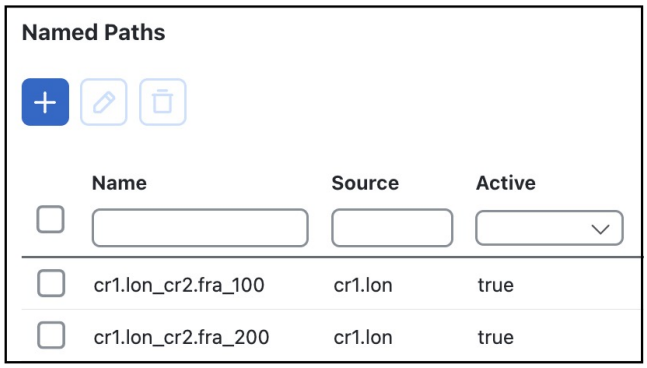
- Strict—The LSP must reach the named path hop directly from the previous one, with no intermediate interfaces.
- Loose—The LSP is routed to this hop from the previous hop using CSPF. Intermediate interfaces can be used.
- Exclude—The node or interface is excluded from the LSP path. This hop type cannot 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.

Named path hops example

Figure 78: Example Named Paths the cr1.lon_cr2.fra LSP, on page 236 extends the earlier example (Figure 77: Example RSVP LSP and associated LSP paths, on page 235), 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 78: Example Named Paths the cr1.lon_cr2.fra LSP

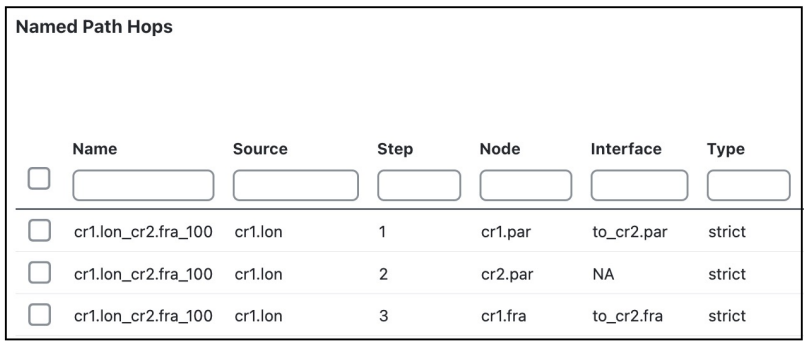


The hops for each named path are defined in a separate table called **Named Path Hops**. Figure 79: Example named path hops for cr1.lon_cr2.fra_100, on page 236 shows the named path hops for the cr1.lon_cr2.fra_100 named path. The **Step** column shows the hop order.

- 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 79: Example named path hops for cr1.lon_cr2.fra_100



Create named paths and their hops

To create Named paths and their hops, do the following:

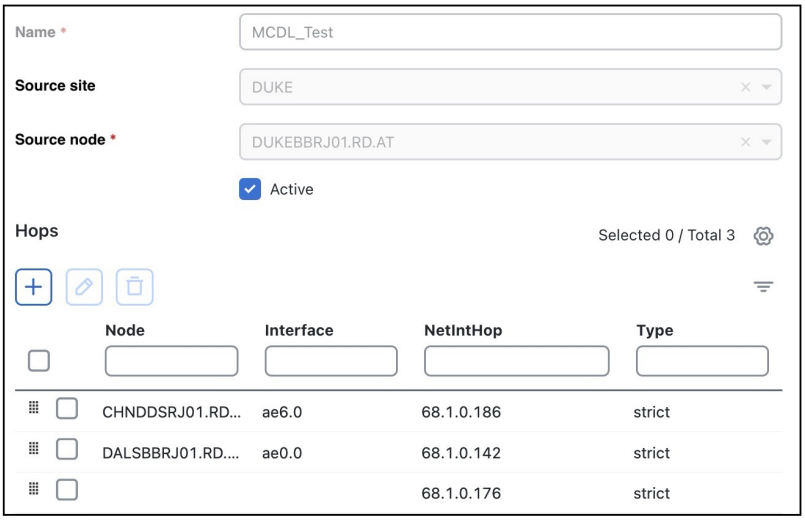
Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 Named paths can be created in the Cisco Crosswork Planning UI in using the following two methods:

- When creating LSP paths, check the **Create associated named paths** check box. Using this option does not create hops.
- 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.



Step 3 Named path options:

- Enter the name in the **Name** field.
- To change the source site or source nodes, choose them from the drop-down lists.
- To activate or deactivate, check **Active**.

Step 4 Hop options:

- To add a new hop, click . Go to Step 5.
- To edit an existing hop, select it from the Hops list and then click . The Edit window opens for named path hops. Go to Step 5.

- To delete an existing hop, choose it from the Hops list, click .

Step 5 To continue creating or editing the named path hop, use the following options.

- As needed, choose the site, node, and interface. For node hops, do not choose an interface.
- Choose the **Type** as Loose, Strict, or Exclude. See Named paths and explicit LSP routing, on page 235 for description.

Step 6 Click **Add**.

Edit named paths and their hops

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 on resolving named paths, see Unresolved LSP destinations and hops, on page 280.

The recommended method of editing the named path hop type is to access it directly from the Named paths table as described in the following steps. This is the most efficient way of viewing the entire path.

Procedure

Step 1 Select one or more named paths from the **Named paths** table.

Step 2 Click .

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

Step 3 In the **Hop** section, edit the Named path hop details, as required. For reference, see Create named paths and their hops, on page 237.

Step 4 Click **Save**.

Actual paths

In Cisco Crosswork Planning, actual paths are the actual route that routed LSP paths are taking. They are read from the live network, per LSP path. 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 see 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. • true = the simulation converted the actual path hops into a complete path from source to destination of the LSP • false = the simulation did not convert the actual path hops into a complete path • NA = not applicable; no actual path was available.
Routed column in the **LSPs** table	• Actual = The LSP follows an actual path. • Simulated = The LSP does not follow an actual path. • Unrouted = Routing was not possible.

Deactivate actual paths for simulations

Cisco Crosswork Planning uses network state in its MPLS simulation, routing LSPs on actual paths where possible and using LSP active path settings. For planning purposes, when state is not relevant, you can change this behavior to disregard actual paths.

Follow these steps to disregard actual paths for simulations.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the toolbar, click **Network options** or choose **Actions > Edit > Network options**.

The Network options page opens.

Step 3 Click the **Simulation** tab.

Step 4 To use or disregard actual paths and active paths in MPLS simulation, check or uncheck **Use LSP actual paths, active paths** check box.

Step 5 Click **Save**.

Affinities

Affinities provide a mechanism for implementing LSP path diversity. By assigning affinities to physical circuits and associating LSPs with affinities, you can 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 network models have 0-31 unnamed and unassigned affinities.

Workflow:

1. Create and edit affinities, on page 240.
2. Assign affinities to interfaces , on page 241.
3. Associate LSPs with affinities , on page 241.

Create and edit affinities

To create or edit affinities, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Edit** > **Admin groups**.

The Admin groups page opens.

Step 3 To create a new affinity:

a) Click .

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) Click **Save**.

Step 4 To modify an existing affinity:

a) Select the affinity you want to edit.

b) Click .

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) Click **Save**.

Assign affinities to interfaces

To use affinities, you must assign them to interfaces in a way that encourages favorable routes and discourages others. For example, continental paths might have one affinity and international paths have another.



To assign affinities to interfaces, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the **Interfaces** table, select one or more interfaces to which you are assigning the affinities.

Step 3 Click .

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

Step 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**.

Step 5 Click **Save**.

You can also follow the above steps to assign affinities to interfaces through a Edit Circuit window. However, you must choose and assign the affinity twice: once to each interface.

Associate LSPs with affinities

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 (see Named paths and explicit LSP routing, on page 235).

Inclusion and exclusion rules

You can choose any of the following 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.

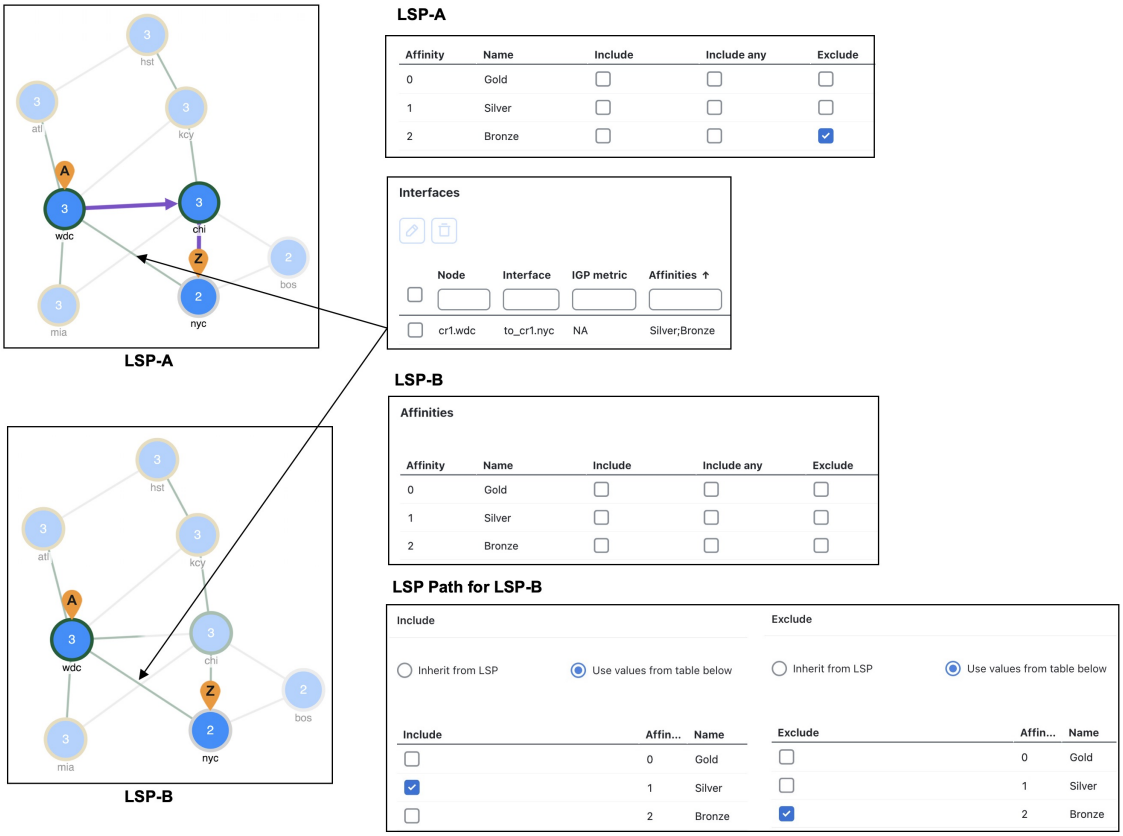
LSP paths inherit the LSP affinities. However, you can edit LSP path affinities to have priority over the LSP affinities.

Example

This example (Figure 80: Affinities example, on page 242) 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 to exclude interfaces assigned to affinity 2 (Bronze). Because the to_cr1.nyc interface is assigned to both of these affinities, LSP B cannot route.

Figure 80: Affinities example



Assign LSPs to affinities

To assign LSPs to affinities, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

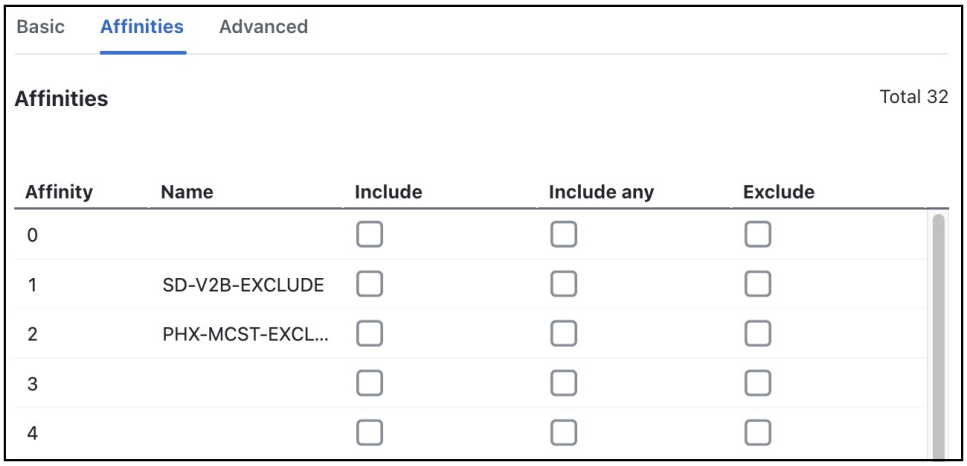
Step 2 From the **LSPs** table, choose one or more LSPs to which you are associating the affinities.

Step 3 Click .

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

Step 4 Click the **Affinities** tab.

Figure 81: LSP affinities



Step 5 Choose the inclusion or exclusion rule (Include, Include any, or Exclude) for the affinity you are associating with the selected LSPs.

Step 6 Click **Save**.

Assign LSP paths to affinities

To assign LSP paths to affinities, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the **LSP paths** table, choose one or more LSP paths to which you are associating the affinities.

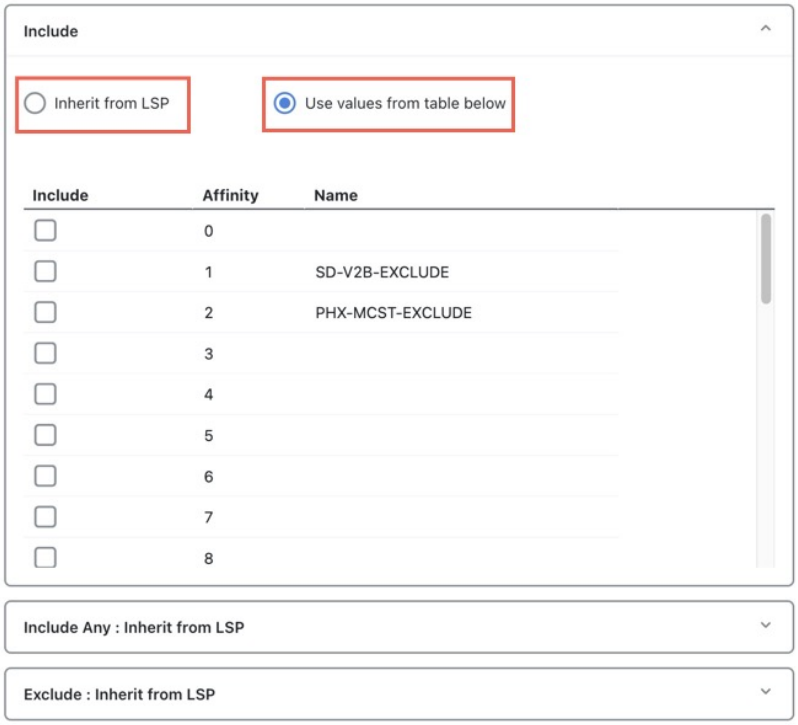
Step 3 Click .

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

Step 4 Click the **Edit** button in the **Affinities** section.

Step 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 82: LSP path affinities



Step 6 Click **Save** in the Edit window to save and exit.

Assign affinities when creating LSP meshes

You can associate affinities with LSPs when creating a new LSP mesh.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

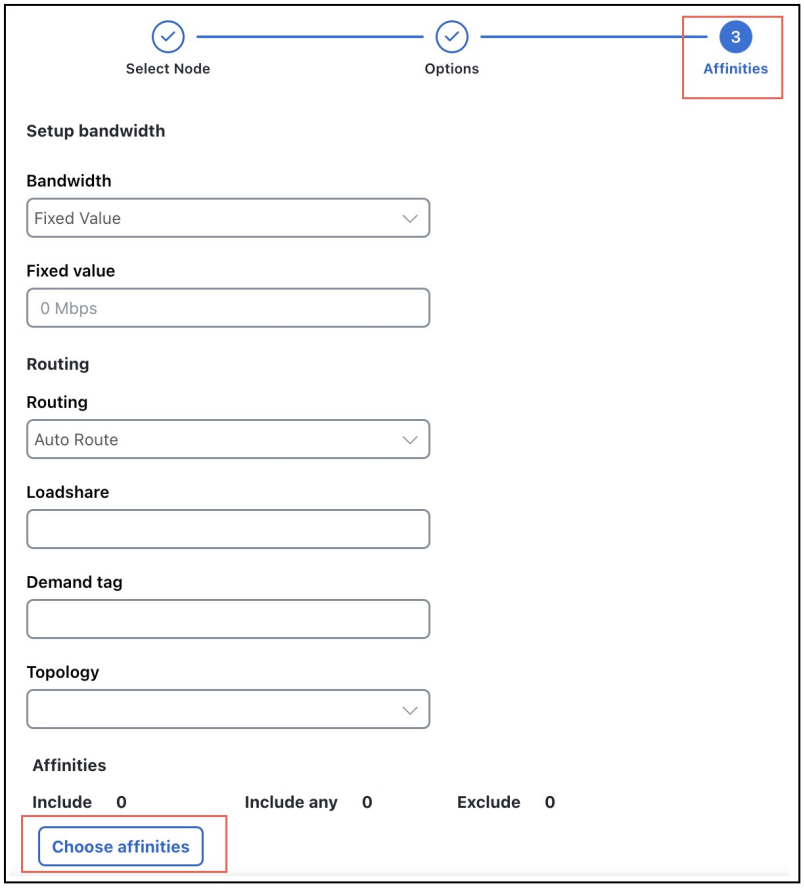
Step 2 From the toolbar, choose **Actions > Insert > LSPs > LSP mesh**.

OR

In the Network Summary panel on the right side, click + > **LSP mesh** in the **LSPs** tab.

Step 3 After selecting the required nodes and options, go to the **Affinities** page.

Figure 83: Insert LSP Mesh window



Step 4 Click **Choose affinities** under the **Affinities** section.

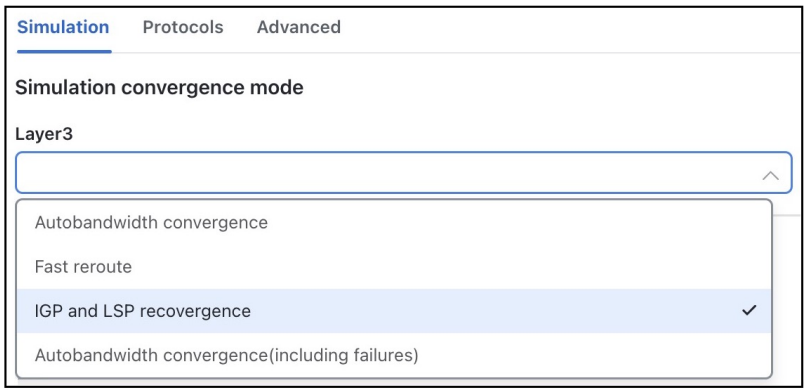
Step 5 Choose the inclusion or exclusion rule for each affinity you are associating with all LSPs in the mesh, and click **Add**.

Step 6 Click **Submit** to save and exit.

Set global simulation parameters

Cisco Crosswork Planning lets you set global parameters that affect how LSPs are routed or rerouted. To access these options, in the toolbar, click **Network options** or choose **Actions** > **Edit** > **Network options**. Then, click the **Simulation** tab.

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. See Simulate autobandwidth-enabled LSPs, on page 252.
- **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. See Fast Reroute simulations, on page 246.
- **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. See Simulate autobandwidth-enabled LSPs, on page 252.

Fast Reroute simulations

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 differ in FRR simulations compared to full convergence simulations, as follows:

- 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 following routing occurs:
  - 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

This section summarizes key properties and terms, and how FRR LSPs are visualized in the network plot. More details about these are further explained throughout the following sections. Note that FRR LSPs appear in the same table as the regular LSPs and are identified in the **Metric type** column as **FRR**.

Objects and properties

Object	Type	Column	Description
LSP	FRR LSP	FRR interface	Designated interface to be avoided by the FRR LSP.
		FRR type	Type of FRR protection: Link or Node.
	Autoroute Forwarding Adjacency	FRR enabled	Marks an LSP to be protected by FRR LSPs.
Interface	na	FRR enabled	Marks an interface to be protected by FRR LSPs. This property is required if an interface is to be used when running the FRR LSPs initializer.

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, see Run fast reroute simulation, on page 252.
- 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.

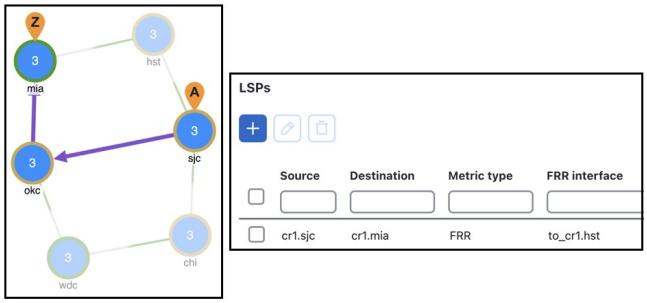
FRR LSP routing

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 window 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 figure 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.



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.

Link and node protection

Cisco Crosswork Planning supports both link (circuit) and node protection FRR LSPs. The type of protection shows 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.

Routing around Failures

Before you begin

- LSPs must be configured to be protected by FRR LSPs. For information on setting up FRR LSPs, see Set up fast reroute simulations, on page 249.
- Any interfaces along the FRR LSP path must be configured to be protected by FRR LSPs. For information on protecting interfaces, see Mark LSPs for protection, on page 249 or the FRR LSPs initializer, on page 250 (which automates the process of setting these).

Assuming one or more failures occurred, the following decisions determine which FRR LSP to take:

Procedure

Step 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.

Step 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.

Step 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

To set up FRR simulations, do the following:

Procedure

Step 1 Set appropriate properties on LSPs to mark them for protection. See Mark LSPs for protection, on page 249.

Step 2 If protecting SRLGs, associate a node with the SRLG in order to establish the protected interface within it. See Identify SRLGs for Protection, on page 249.

Step 3 If you are using the initializer to create FRR LSPs, identify which interfaces to protect. See FRR LSPs initializer, on page 250.

Step 4 Create the FRR LSPs using the FRR LSPs initializer or manually. See FRR LSPs initializer, on page 250 or Create FRR LSPs Manually, on page 251.

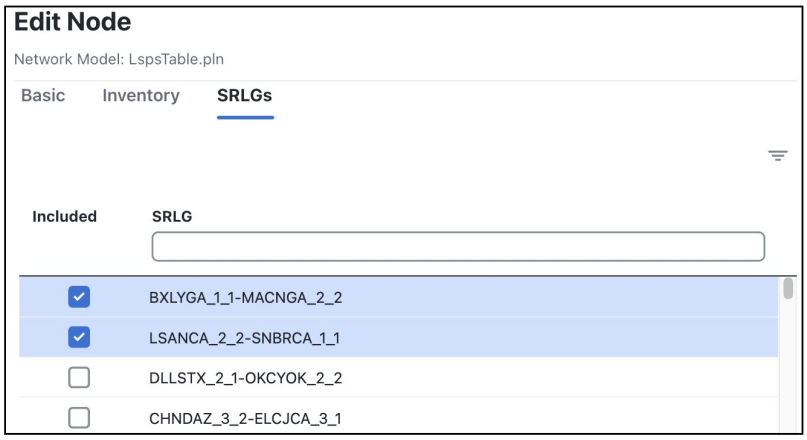
Mark LSPs for protection

To mark the LSPs for Fast Reroute protection, you must set the following properties in the Edit LSP window.

- 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.
- 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

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

Step 1 Select the source node of a circuit within an SRLG that you want to protect. Click or choose ··· > **Edit**.

Step 2 Click the **SRLGs** tab.

Step 3 Choose one or more SRLGs with which you want to associate this node, and click **Save**.

FRR LSPs initializer

The **FRR LSPs** initializer automatically creates link-protection or node-protection FRR LSPs when the following conditions are met:

- One or more LSPs have their **FRR enabled** property set (see Mark LSPs for protection, on page 249).
- 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 window.

Run the FRR LSPs Initializer

Follow these steps to run the FRR LSPs initializer.

Procedure

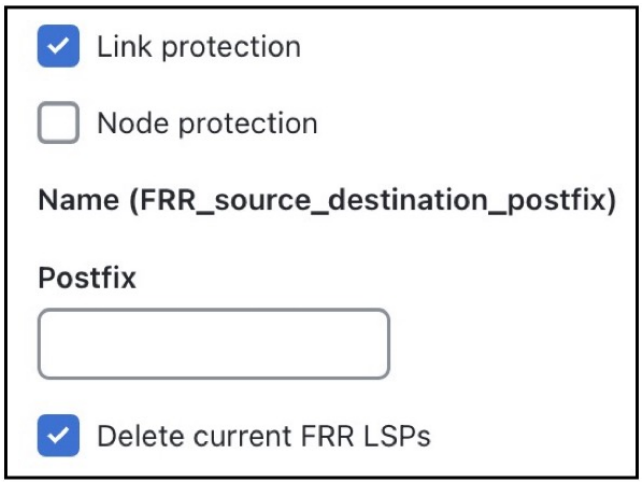
Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions > Initializers > FRR LSPs**.

Step 3 Select the nodes that you want to optimize. If you do not choose any nodes, the initializer uses all nodes in the plan file.

Step 4 Click **Next**.

Step 5 Choose whether to create FRR LSPs that protect links (circuits), nodes, or both.



Step 6 (Optional) Enter a postfix to add to the end of the names of the newly created FRR LSPs.

Step 7 (Optional) Uncheck the **Delete current FRR LSPs** check box if you want to retain the existing FRR LSPs.

Step 8 Click **Submit**.

Create FRR LSPs Manually

This section describes the required properties for an FRR LSP. It does not describe all the options available when creating an LSP.

Procedure

Step 1 Choose **Actions > Insert > LSPs > LSP**, or click > **LSPs** in the **LSPs** page. Alternatively, you can choose an existing LSP to modify to be an FRR LSP. If doing so, click or choose > **Edit**.

Step 2 Choose the source site and node of the FRR LSP.

Step 3 Choose the destination site and node, or enter a NetInt destination IP address. For information on NetInt destination IP addresses, see Unresolved LSP destinations and hops, on page 280.

Step 4 In the **Routing** area, choose **FRR** as the Routing type.

Step 5 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.

Step 6 Click **Save**.

Run fast reroute simulation

FRR LSPs are routed only when simulating failures in the Fast Reroute simulation mode. To enable the Fast Reroute simulation mode, follow these steps.

Procedure

Step 1 In the toolbar, click **Network options** or choose **Actions** > **Edit** > **Network options**.

The Network options page opens.

Step 2 Click the **Simulation** tab.

Step 3 In the **Simulation convergence mode** section, choose **Fast reroute** from the Layer 3 drop-down list.

Step 4 Click **Save**.

Simulate autobandwidth-enabled LSPs

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.

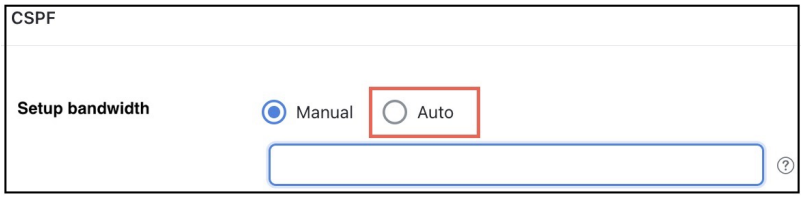
While Cisco Crosswork Planning does not simulate this instability, it does simulate autobandwidth-enabled LSPs. In this simulation mode, internally Cisco Crosswork Planning first routes the autobandwidth-enabled LSPs with a setup bandwidth of zero. Then, demands are routed to determine the simulated 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.

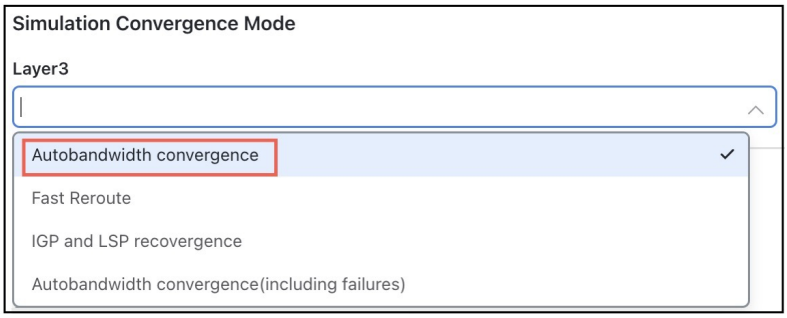
There are two autobandwidth simulation modes available. Before simulating failures or running Simulation analysis on autobandwidth-enabled LSPs, choose the appropriate mode depending on which state you need to simulate.

- In the **Autobandwidth convergence** 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.
- In the **Autobandwidth convergence (including failures)** mode, the network is simulated after the Setup BW Sim values have been reset. This simulates the longer, more optimal response to failures.

Note The **Autobandwidth convergence** and **Autobandwidth convergence (including failures)** simulation modes are only available in plan files with a single traffic level.

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 **Network options** or choose **Actions** > **Edit** > **Network options**).

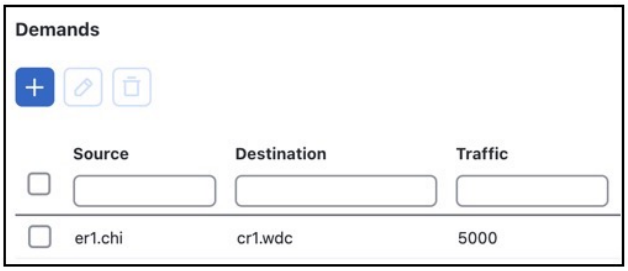


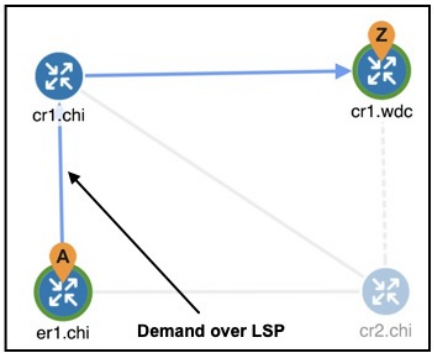


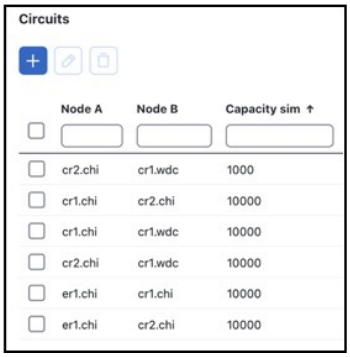
Example autobandwidth convergence without and with failures

The network in this example has the following parameters:

- One demand from er1.chi to cr1.wdc.
- Two LSPs: LSP-A from cr1.chi to cr1.wdc and LSP-B from cr2.chi to cr1.wdc.
- Neither of these LSPs has a Setup BW.
- All circuits have a 10,000 Mbps capacity except for one (between cr2.chi and cr1.wdc), which is only 1000 Mbps.







- Figure 84: Example LSP routing using IGP and LSP reconvergence and using autobandwidth convergence, on page 255 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 autobandwidth convergence with failures, on page 255 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 including failures, on page 256 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 84: Example LSP routing using IGP and LSP reconvergence and using autobandwidth convergence

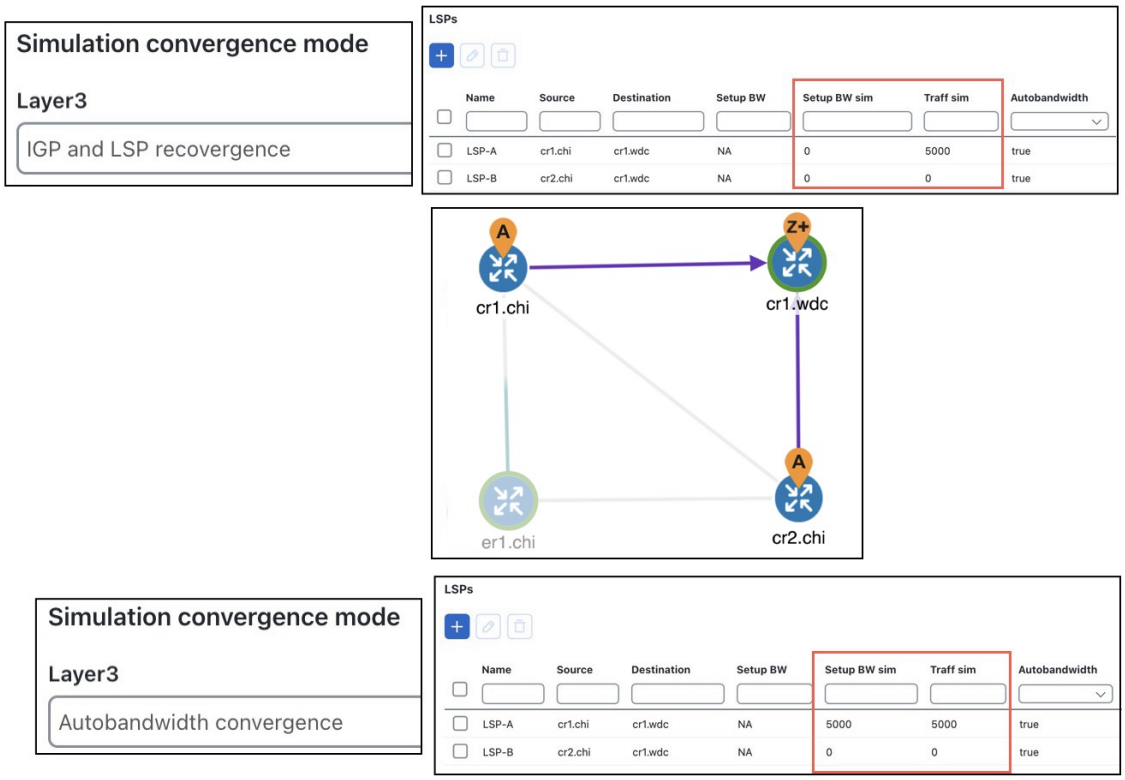


Figure 85: Example autobandwidth convergence with failures

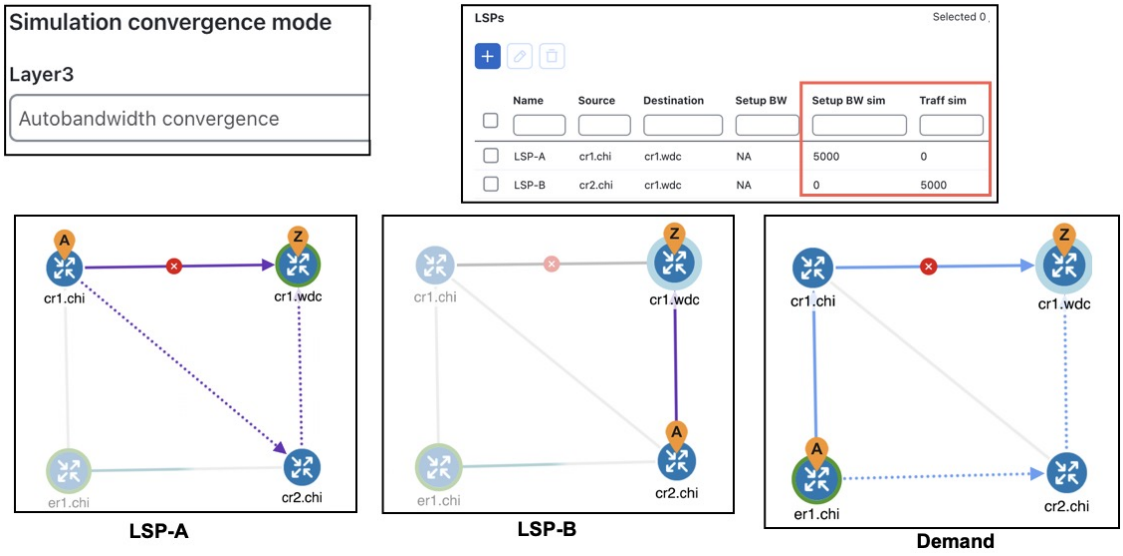
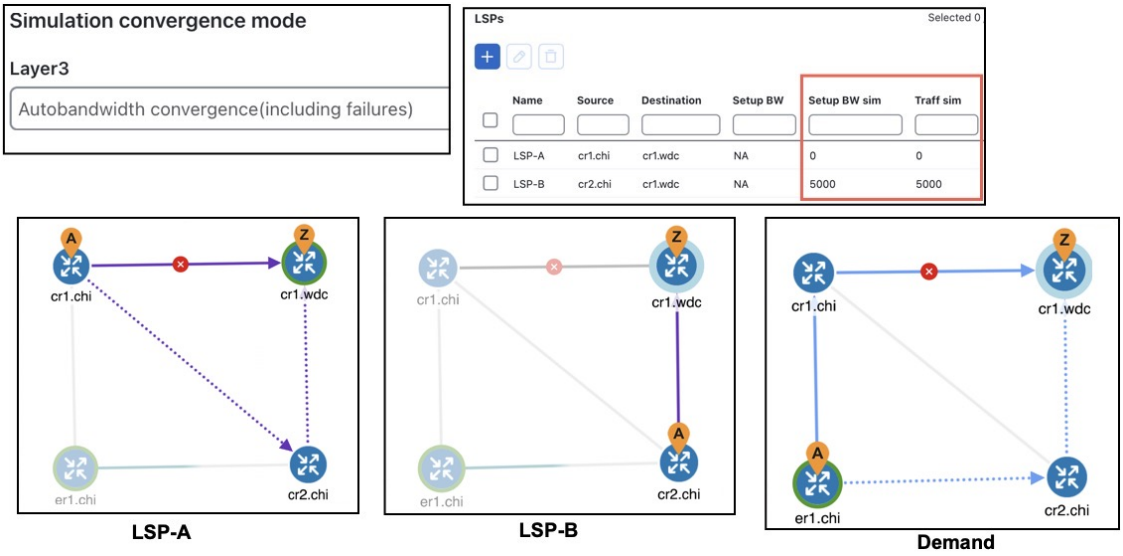


Figure 86: Example autobandwidth convergence including failures



Circuit Style RSVP LSPs

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.

Key features of CS-RSVP LSPs

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

To improve redundancy, a CS-RSVP LSP has four paths: Working (W), Protect (P), 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 details, see Restore path calculation, on page 257.

Cisco Crosswork Planning ensures that both W and P paths have a fully explicit path defined in the named path.

Restore path calculation

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.

Example

Figure 87: Restore path calculation example, on page 258 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".

The restore path is calculated as follows:

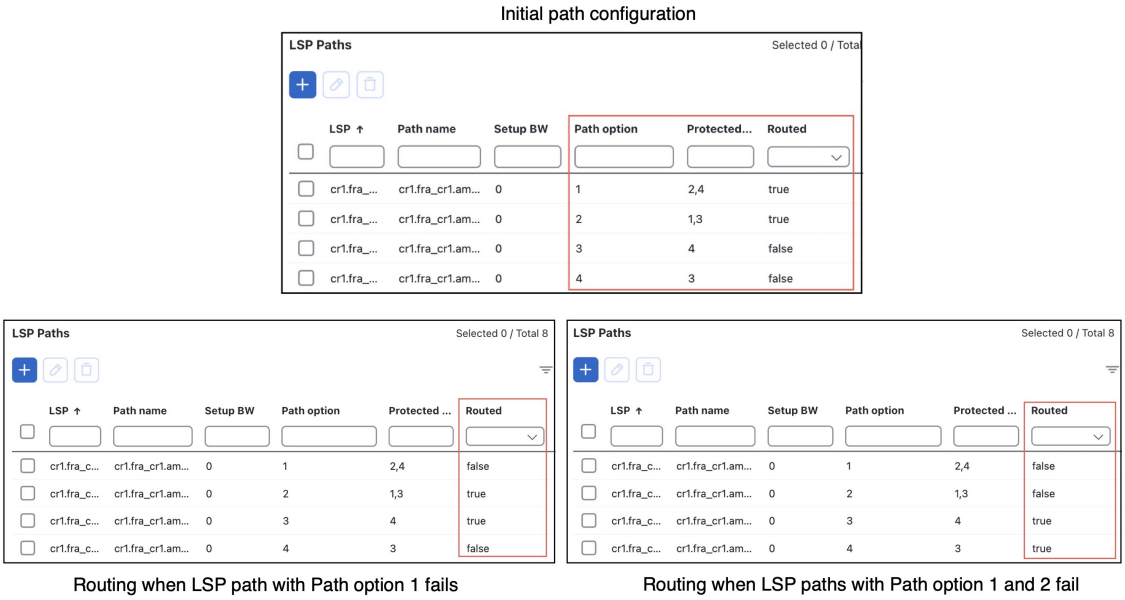
1. If LSP-1 fails, the traffic shifts to LSP-2 and LSP-3 becomes the Protect path. This happens because the **Protected by** field for LSP-1 is set to **2, 4**, indicating that when LSP-1 goes down, LSP-2 becomes the Working path and LSP-4 becomes the restore path. Since LSP-2 has the **Protected by** 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 87: Restore path calculation example



CS-RSVP LSP path simulation

Cisco Crosswork Planning performs these checks to determine if the path is routable:

- The explicit paths go through the same hops in both directions. This check applies only to restore paths routed dynamically.
- The explicit path can reach the endpoints in both directions. Note that the configuration check is not performed.
- The path has enough bandwidth to satisfy the setup bandwidth requirement. Note that asymmetric configuration is not checked.
- 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, see Biased routing, on page 258.
- 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 un-routable, the LSP path becomes unrouted. Paths are found dynamically only for restore paths.
- 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.

Biased routing

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.

Set the biased routing options

Follow these steps set the biased routing options.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 In the toolbar, click **Network options** or choose **Actions** > **Edit** > **Network options**.

The Network options page opens.

Step 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.

Step 4 Click **Save**.

Create CS-RSVP LSPs

You can create CS-RSVP LSPs in the Cisco Crosswork Planning Design application either manually or using the **CS-RSVP LSP optimizer** tool.

This section explains how to create CS-RSVP LSPs manually. For details on using the CS-RSVP LSP optimizer tool, see Create CS-RSVP LSPs using the optimizer tool, on page 260.

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

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **LSPs** > **CS-RSVP LSP**.

OR

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.

Step 3 In the **Node A** and **Node B** sections, choose the endpoints of the CS-RSVP LSP.

Step 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.

Step 5 Click **Add**.

CS-RSVP LSP optimizer

The **CS-RSVP LSP optimizer** tool helps to create, find, and optimize CS-RSVP LSPs based on specified endpoint pairs, association ID, source address, and global ID. It optimizes both existing and newly created CS-RSVP LSPs together, or you can choose to optimize only the existing LSPs or only the newly created ones.

Create CS-RSVP LSPs using the optimizer tool

If the 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

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

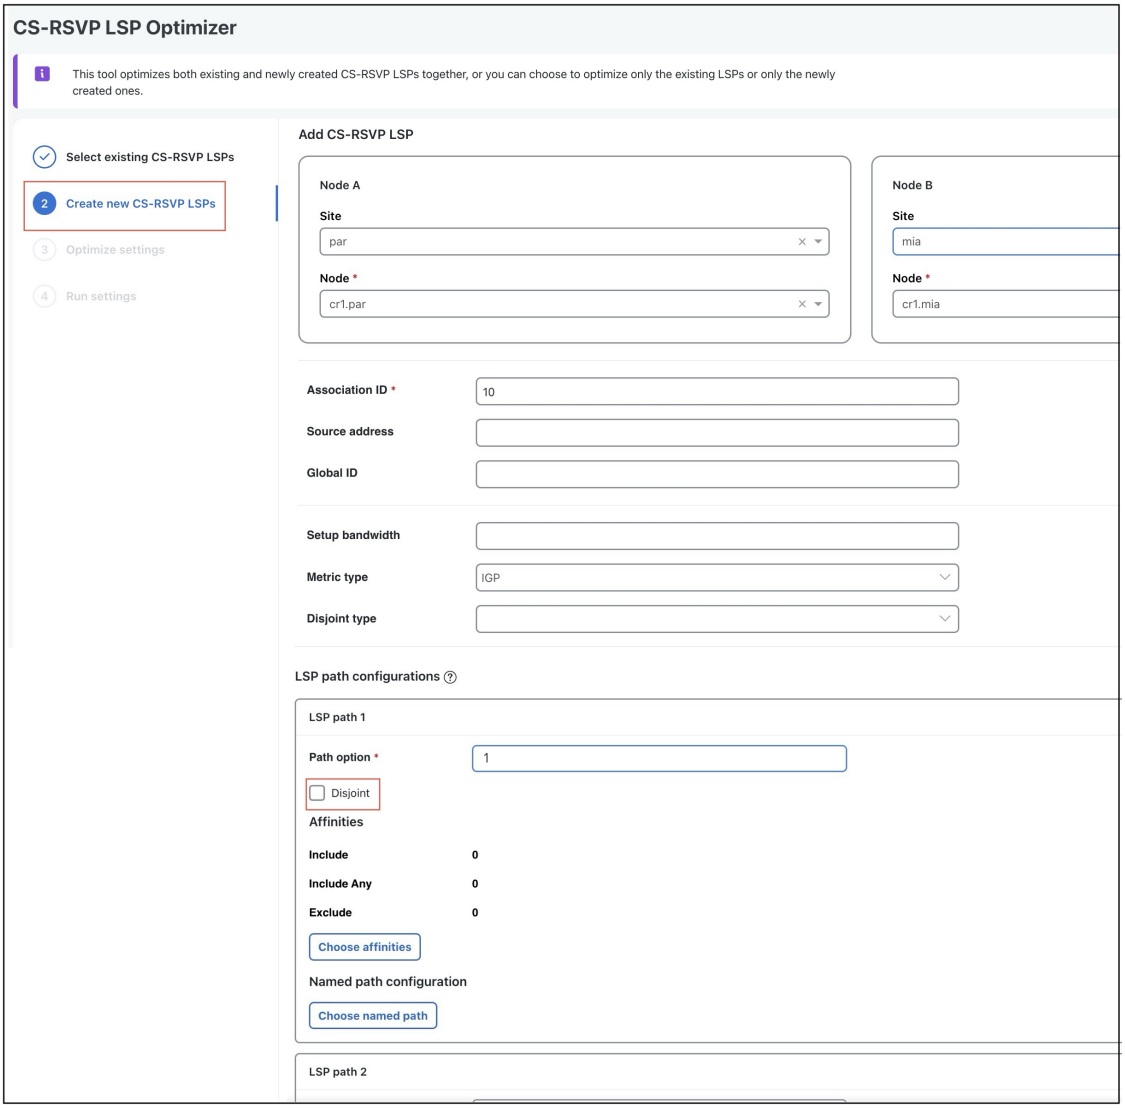
Step 2 From the toolbar, choose **Actions** > **Tools** > **CS-RSVP LSP optimizer**.

Step 3 (Optional) If you want to optimize the existing CS-RSVP LSPs, select them and click **Next**.

Step 4 Click **Create** in the Create new CS-RSVP LSPs page.

Step 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, see Table 24: CS-RSVP LSP configuration parameters, on page 263.

Figure 88: Create new CS-RSVP LSPs page



Step 6 Click **Save**. Then, select the newly created CS-RSVP LSP and click **Next**.

To create additional CS-RSVP LSPs, click .

Step 7 Select the interfaces you want to exclude from the route calculation of CS-RSVP LSPs.

Step 8 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

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.

Step 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 that:

- 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 the existing CS-RSVP LSPs using the optimizer tool

If the CS-RSVP LSPs already exist, the optimizer can optimize them based on the current configurations and requirements. Follow these steps to optimize the existing CS-RSVP LSPs using the **CS-RSVP LSP optimizer** tool.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22) with CS-RSVP LSPs. It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Tools** > **CS-RSVP LSP optimizer**.

Step 3 Select the existing CS-RSVP LSPs you want to optimize and click **Next**.

Step 4 (Optional) If you want to create new CS-RSVP LSPs, click **Create** and configure the required parameters. For details, see Create CS-RSVP LSPs using the optimizer tool, on page 260.

Step 5 Select the interfaces you want to exclude from the route calculation of CS-RSVP LSPs.

Step 6 Click **Next**.

Step 7 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

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.

Step 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

This table lists the parameters to configure when creating CS-RSVP LSPs using the optimizer tool.

Table 24: 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.

Field	Description
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. 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

The **CS-RSVP LSP simulation diagnostics** tool in Cisco Crosswork Planning performs configuration and disjointness checks on the CS-RSVP LSPs.

Configuration 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

Disjointness check

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 three 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

Follow these steps to use the **CS-RSVP LSP simulation diagnostics** tool.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22) with CS-RSVP LSPs. It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Tools** > **Diagnostics** > **CS-RSVP LSP simulation**.

Step 3 Click **Proceed** in the information dialog box.

A report is generated indicating the details of the failures.

CS-RSVP LSP simulation diagnostics report

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.

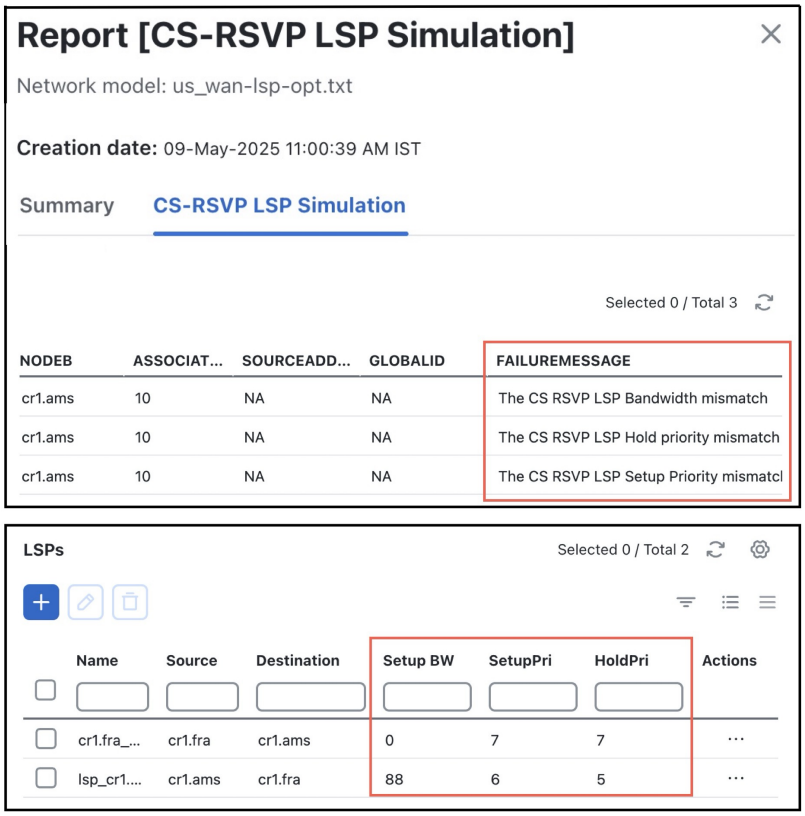
The **CS-RSVL LSP simulation** tab provides failure details. Expand the **FAILUREMESSAGE** column to view the failure messages.

Example

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 see that the values for these fields do not match.

Figure 89: CS-RSVP LSP simulation diagnostics report and mismatching LSPs



Visualize CS-RSVP LSPs, and associated LSPs and LSP paths

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, and
- LSP paths associated with CS-RSVP LSPs.

Visualize CS-RSVP LSPs

This section explains how to visualize CS-RSVP LSPs in the UI.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22) with CS-RSVP LSPs. It opens in the **Network Design** page.

Step 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.

Step 3 Select the required CS-RSVP LSP from the list.

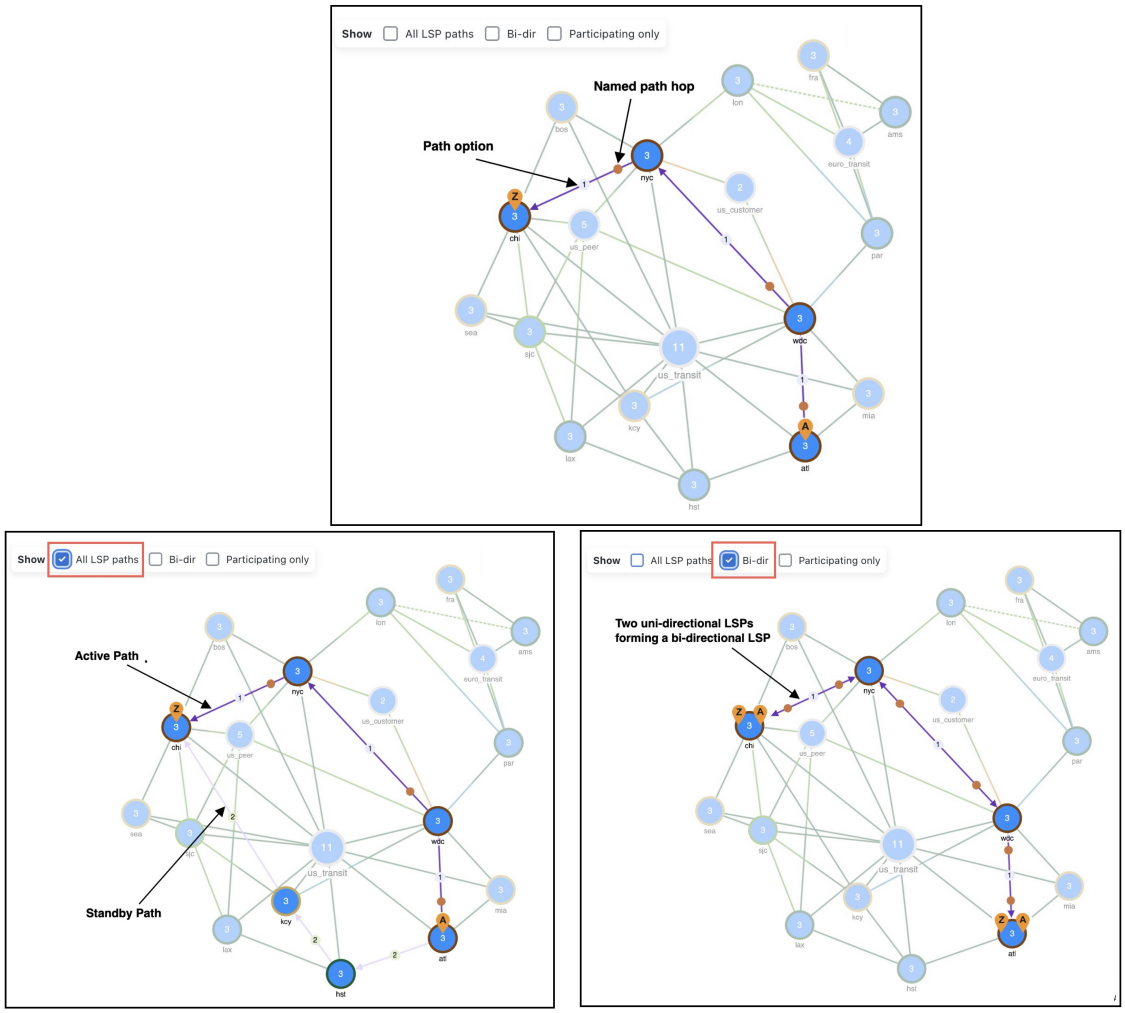
Step 4 Observe these in the network plot on the left side:

- 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, see the first image in Figure 90: CS-RSVP LSP path visualization, on page 268.

  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, see the second image in Figure 90: CS-RSVP LSP path visualization, on page 268.

- The bi-directional LSPs are plotted when you check the **Bi-dir** check box in the network plot. For example, see the third image in Figure 90: CS-RSVP LSP path visualization, on page 268.

Figure 90: CS-RSVP LSP path visualization



Visualize LSPs associated with CS-RSVP LSPs

This section explains how to visualize the LSPs associated with the selected CS-RSVP LSPs.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22) with CS-RSVP LSPs. It opens in the **Network Design** page.

Step 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.

Step 3 Choose ☰ > **Filter to LSPs**. The LSPs table opens, displaying the LSPs associated with the selected CS-RSVP LSP.

Step 4 Select the LSPs.

Note
You can select up to three LSPs at a time.

Step 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

This section explains how to visualize the LSP paths associated with the selected CS-RSVP LSPs.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22) with CS-RSVP LSPs. It opens in the **Network Design** page.

Step 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.

Step 3 Choose ☰ > **Filter to LSP paths**. The LSP Paths table opens, displaying the LSP paths associated with the selected CS-RSVP LSP.

Step 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

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, as described in the following steps. 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

Step 1 Route LSPs as usual. (For information, see Dynamic LSP routing and CSPF, on page 229.)

Step 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

Follow these steps to ignore reservable bandwidth constraints in routing simulations.

Procedure

Step 1 From the toolbar, click **Network options** or choose **Actions > Edit > Network options**.

The Network options page opens.

Step 2 Click the **Simulation** tab.

Step 3 In the **Label switched paths** section, check the **LSP capacity planning mode** check box. This ignores the reservable bandwidth constraints.

Step 4 Click **Save**.

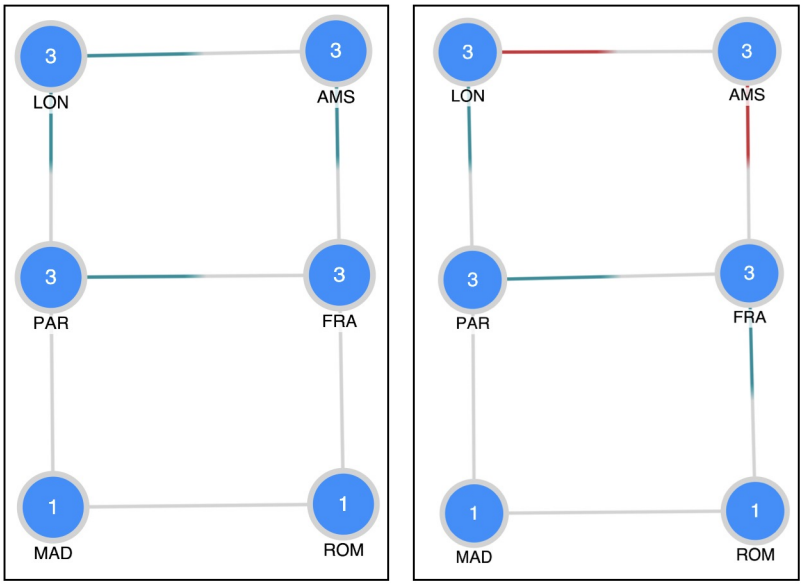
Example

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.

- LSPs
  - 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 91: Comparison of reservable bandwidth being used vs not being used, on page 271 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 on the LSP Reservations view, see View LSP reservations, on page 232.)

Figure 91: Comparison of reservable bandwidth being used vs not being used



Reservable Bandwidth Mode

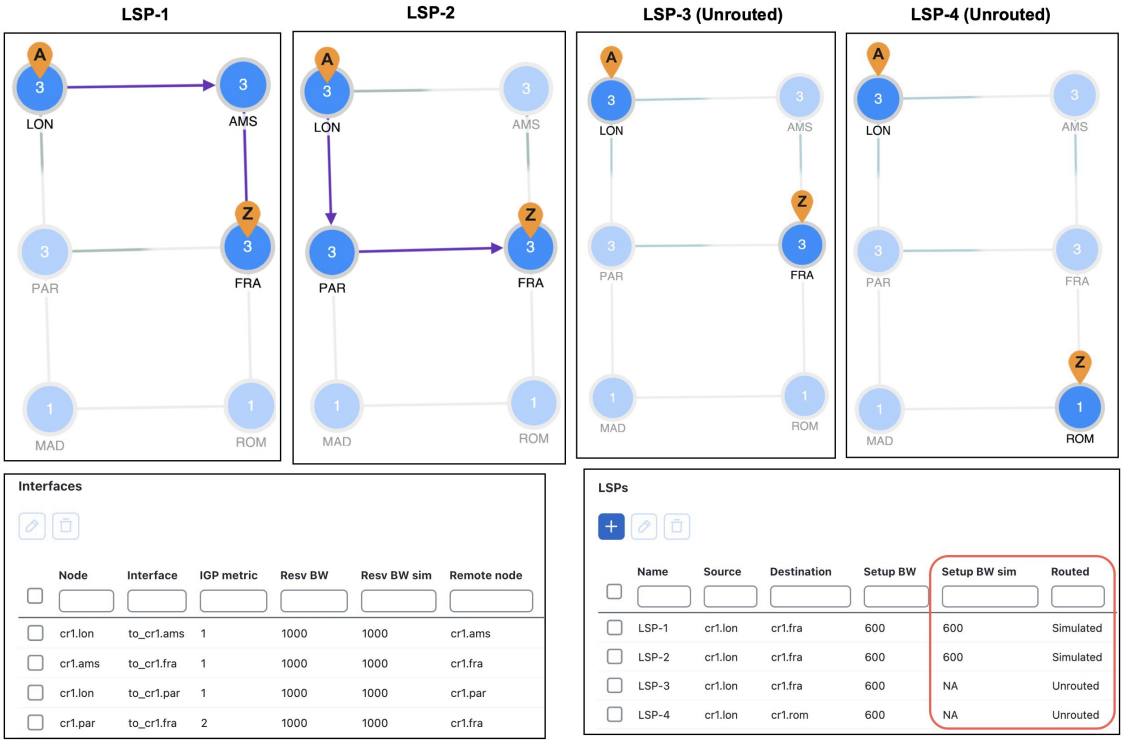
Capacity Planning Mode where Reservable Bandwidth is Ignored

Figure 92: Simulated routing using reservable bandwidth constraints, on page 272 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.”

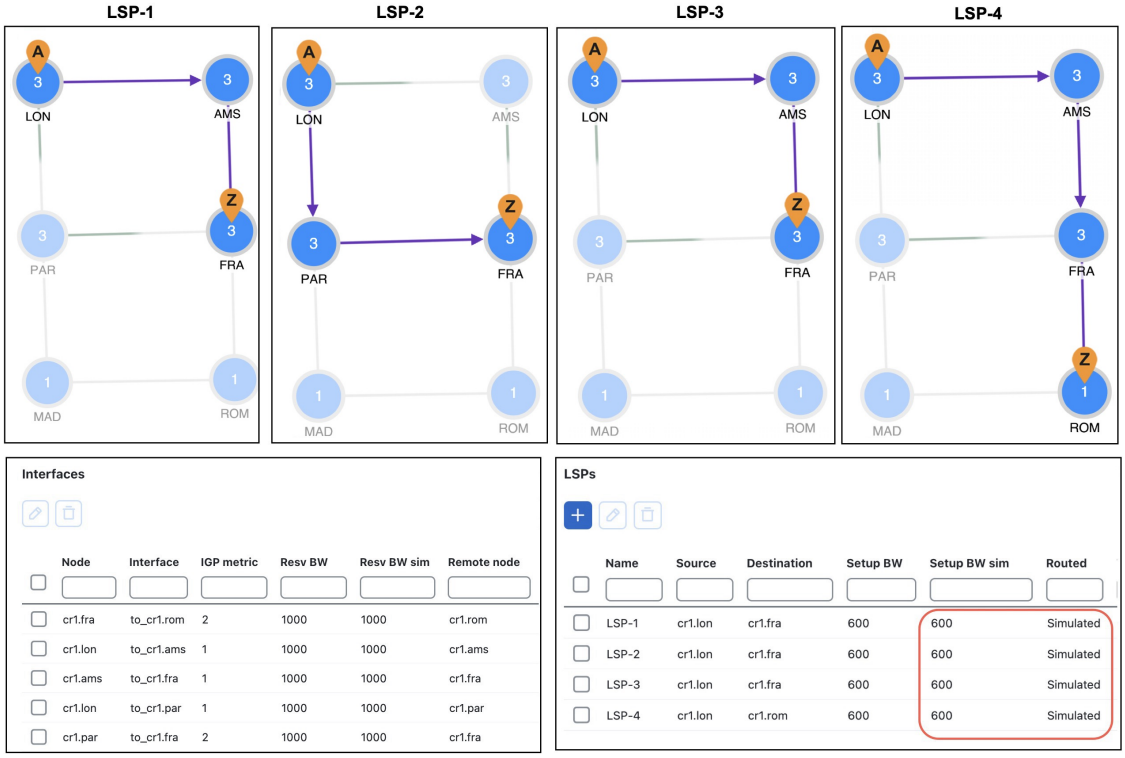
Figure 92: Simulated routing using reservable bandwidth constraints



In Figure 93: Simulated routing without using reservable bandwidth constraints, on page 273, the use of reservable bandwidth constraints has been disabled (**LSP capacity planning mode** checked). The LSP routes are now as follows.

- 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 93: 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

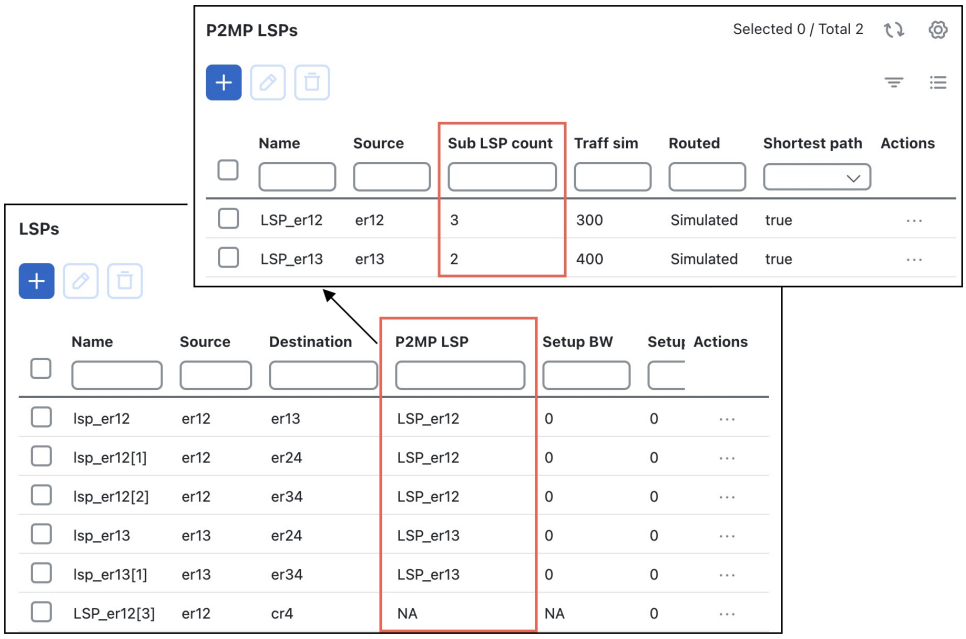
Point-to-multipoint (P2MP) LSPs offer an MPLS alternative to IP multicast for setting up multiple paths to multiple locations from a single source. In live networks, the signal is sent once across the network and packets are replicated only at the relevant or designated branching MPLS nodes.

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 94: P2MP LSPs and associated sub-LSPs, on page 274, 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 94: P2MP LSPs and associated sub-LSPs



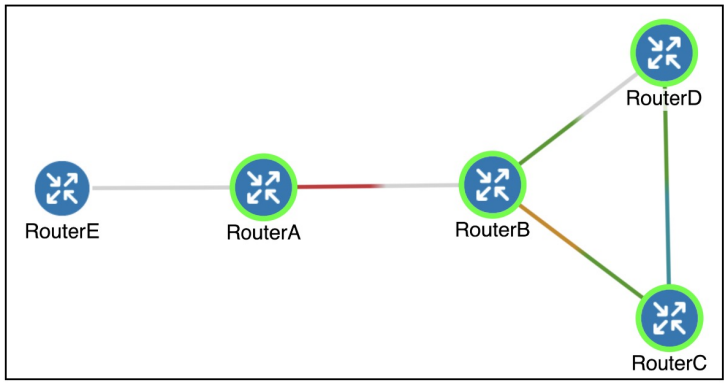
P2MP LSP bandwidth

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 see 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, see View LSP reservations, on page 232.

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 95: Bandwidth behavior without a P2MP LSP, on page 275). 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 96: Bandwidth behavior with a P2MP LSP, on page 276).

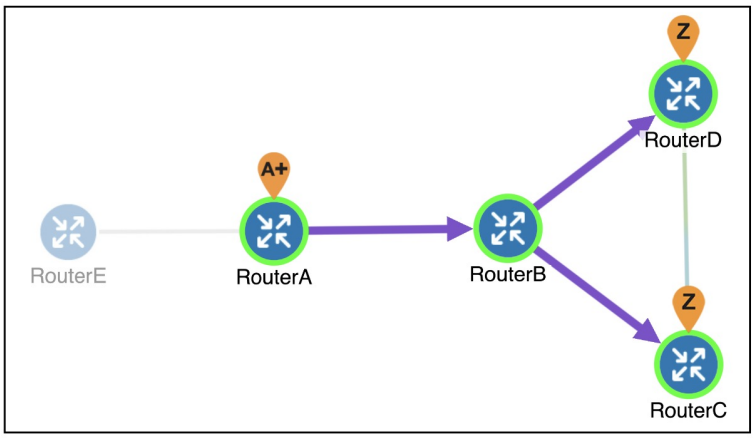
Figure 95: Bandwidth behavior without a P2MP LSP



Interfaces

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 96: Bandwidth behavior with a P2MP LSP



Interfaces 56 | Demands 1 | Nodes 23 | LSPs 4 | More

Interfaces

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

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. (See Create Demands for P2MP LSPs, on page 278 for information on creating these demands.)

Demands must be private to the P2MP LSP. This privacy is automatically created when you insert a demand for a P2MP LSP.

Create P2MP LSPs and sub-LSPs

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

Follow these steps to create P2MP LSPs.

Procedure

Step 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.

Step 2 If you have not done so, create sub-LSPs for the P2MP LSP. For details, see Create Sub-LSPs, on page 277.

Step 3 Associate the sub-LSPs to the P2MP LSP. You can follow these steps individually, or you can choose up to three 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

Follow these steps to create sub-LSPs.

Procedure

Step 1 From the toolbar, choose **Actions > Insert > LSPs > LSP**, or click  > **LSPs** in the **LSPs** table.

Step 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.

Step 3 Choose the site and node that are the final destination for the sub-LSP.

Step 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.

Step 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, see Create P2MP LSPs, on page 276.

Step 6 Click **Save**.

Create a mesh of sub-LSPs

Procedure

Step 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.

Step 2 From the toolbar, choose **Actions > Insert > LSPs > LSP mesh**, or click > **LSP mesh** in the **LSPs** table.

Step 3 In the **Destination nodes** section, uncheck these check boxes and click **Next**.

- **Same destinations as sources**
- **Create LSPs from destination to source**

Step 4 Complete the remaining optional fields as needed and click **Next**.

Step 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.

Step 6 Complete the remaining optional fields as needed and click **Submit**.

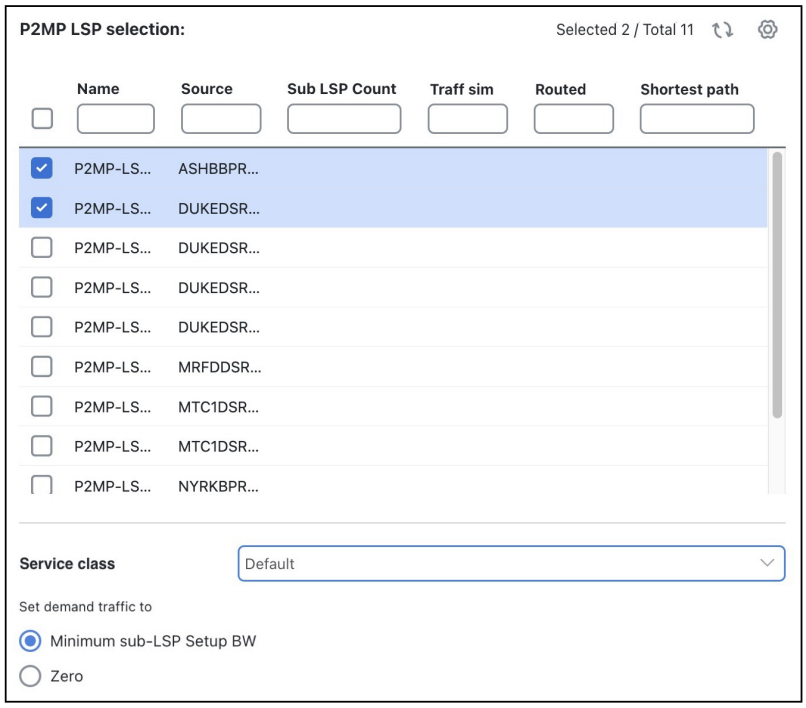
Create Demands for P2MP LSPs

To simulate traffic routing through a P2MP LSP, create a demand for it.

Procedure

Step 1 If you want to create demands for P2MP LSPs, choose one or more P2MP LSPs from the P2MP LSP table.

Step 2 From the toolbar, choose **Actions > Insert > LSPs > Demands for P2MP LSPs**, or click > **Demands for P2MP LSPs** in the **LSPs** table.



Step 3 Choose the P2MP LSPs through which you are routing traffic.

Step 4 (Optional) Choose a service class for the newly created demand.

Step 5 Choose the desired bandwidth: **Minimum sub-LSP Setup BW** or **Zero**.

Step 6 Click **Submit**.

Delete P2MP LSPs and sub-LSPs

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

Procedure

Step 1 In the **LSPs** table, choose one or more sub-LSPs.

Step 2 Click .

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

Step 3 In the Sub-LSP of P2MP LSP list, do one of the following:

- 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.

Step 4 Click **OK**.

Delete sub-LSPs or P2MPs

Delete sub-LSPs but keep P2MP LSP	Delete P2MP LSP
1. Select one or more sub-LSPs in the **LSPs** table. 2. Click . If you are deleting a single sub-LSP, you can also use the ••• > **Delete** option under the **Actions** column. 3. Click **Delete** in the confirmation dialog box to continue.	1. Select one or more P2MP LSPs in the **P2MP LSPs** table. 2. Click . If you are deleting a single P2MP LSP, you can also use the ••• > **Delete** option under the **Actions** column. 3. Click **Delete** in the confirmation dialog box to continue.

Unresolved LSP destinations and hops

RSVP-TE LSP configurations and state are read using network discovery from the source node of the LSP, through SNMP, the Parse Configs 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.

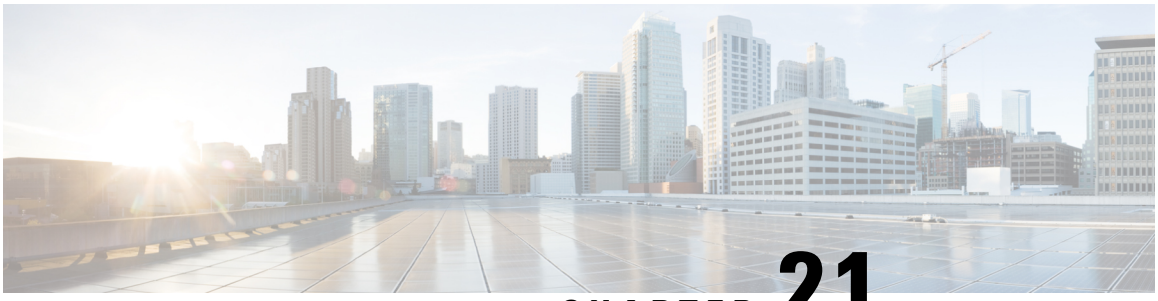
References that might not be resolved are as follows:

- 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 as follows:

- 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.



CHAPTER 21

Optimize RSVP-TE Routing

The **RSVP-TE optimization** tool creates explicitly routed LSP paths to optimize route properties such as distance, TE metrics, utilization, and disjointness. The routing decisions are centralized within the tool, in contrast to traditional RSVP-TE routing, in which routing decisions are distributed across the head-end routers in the network. Although it has fewer parameters, this tool is faster and has a more predictable run-time than the Explicit LSP Optimization tool and is the recommended method for explicit RSVP-TE LSP path optimization.

The primary optimization goal is to minimize path distance (which can be based on interface TE metrics, delay, or 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.

As such, 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 can be implemented on the Cisco Crosswork Planning platform.

Note RSVP-TE optimization tool supports both Inter-Area and Inter-AS functionalities.

This section contains the following topics:

- Specify inputs for optimization, on page 282
- Run RSVP-TE optimization, on page 287

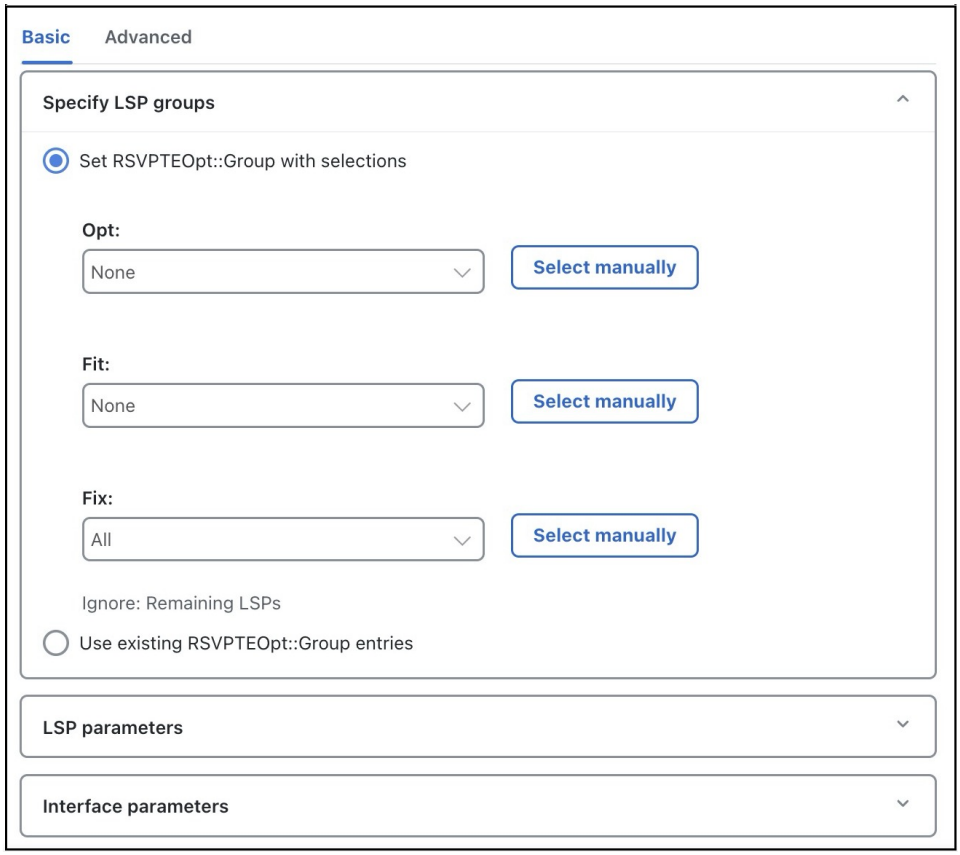
• Analyze optimization output, on page 289

Specify inputs for optimization

The **RSVP-TE optimization** tool (**Actions > Tools > RSVP LSP optimization > RSVP-TE optimization**) uses the LSP and interface input properties listed in Table 25: RSVPTEOpt input parameters, on page 283 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:

- Use the tool to automatically create (initialize) 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 97: RSVP-TE optimization options



Example: 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). If the property already exists, it is updated with the new value if it has changed.

Example: 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. If the property already exists, it is updated with the new bandwidth bound value if it has changed.

- 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, see Specify inputs for optimization, on page 282.

Note Because 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, and this same property could be used as input to other optimizations.

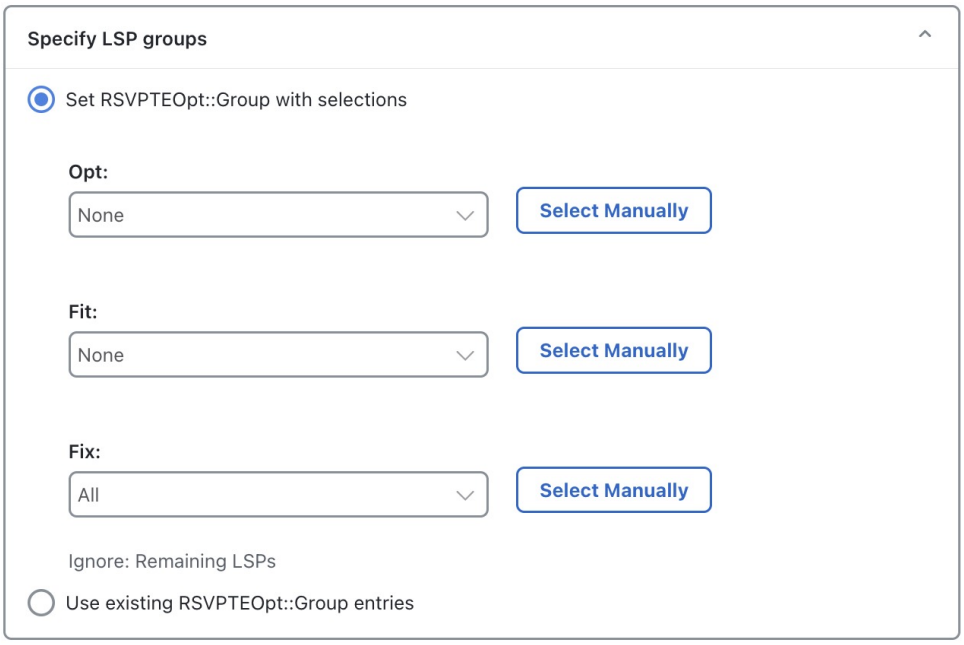
Table 25: RSVPTEOpt input parameters

Table	Input Property	Description and Values Used to Create the Input Property
LSPs	RSVPTEOpt::Group	LSP optimization group. 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. For details, see Select LSP groups, on page 283.
	RSVPTEOpt::BWReq	The bandwidth required by each LSP that is being optimized. 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. For details, see Set optimization parameters, on page 285.
Interfaces	RSVPTEOpt::BWBound	Sum total of LSP bandwidth routed over the interface should not exceed this value. 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. For details, see Set optimization parameters, on page 285.
	RSVPTEOpt::Metric	Metric value used in the shortest path calculations. 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. For details, see Set optimization parameters, on page 285.

Select LSP groups

To determine how and whether LSPs are optimized, select them from the **Specify LSP groups** area. If an LSP is in multiple groups, the first group in which it resides is how the LSP is treated.

Figure 98: Select LSP groups



Select the **Set RSVPTEOpt::Group with selections** radio button and set the following properties, as required:

- **Opt**—Explicitly route or reroute these 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**—Reroute these 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, treat 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.

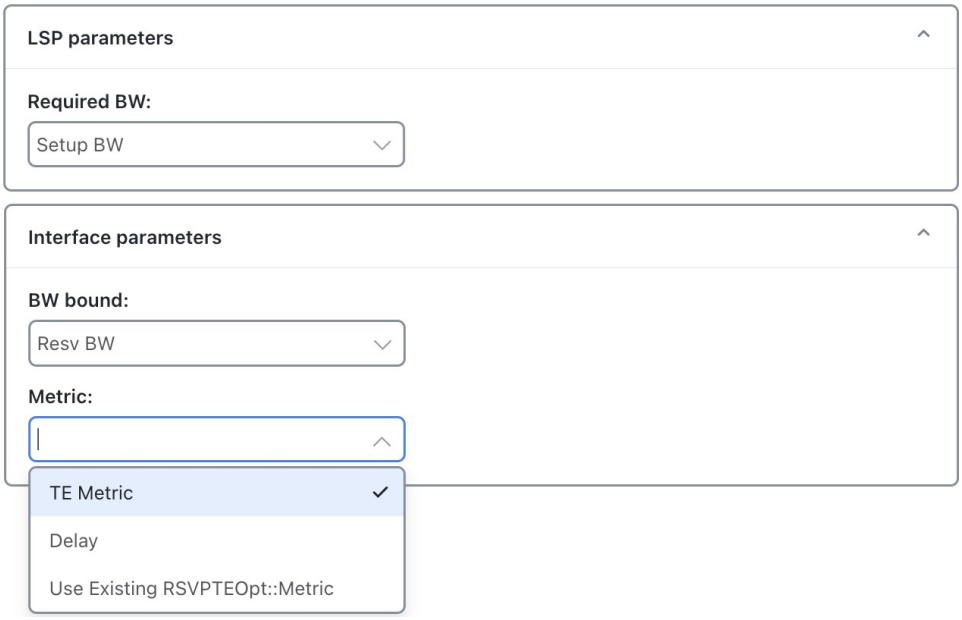
- **Fix**—Do not reroute these LSPs, but consider them in the optimization calculations.

LSPs that are not selected are ignored, and as such, maintain their original configuration and are not considered in the optimization calculations. This can be useful if you have network regions that do not interfere with optimizations. Note that if measured interface traffic is used in the optimizations, traffic contributions from the ignored LSPs to this measured traffic might affect the results.

Because you can set each of the **Set RSVPTEOpt::Group with selections** selections to None, All, or manually selected, you can section off specific LSPs for one optimization treatment or another.

Set optimization parameters

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. See Table 25: RSVPTEOpt input parameters, on page 283 for selection options.



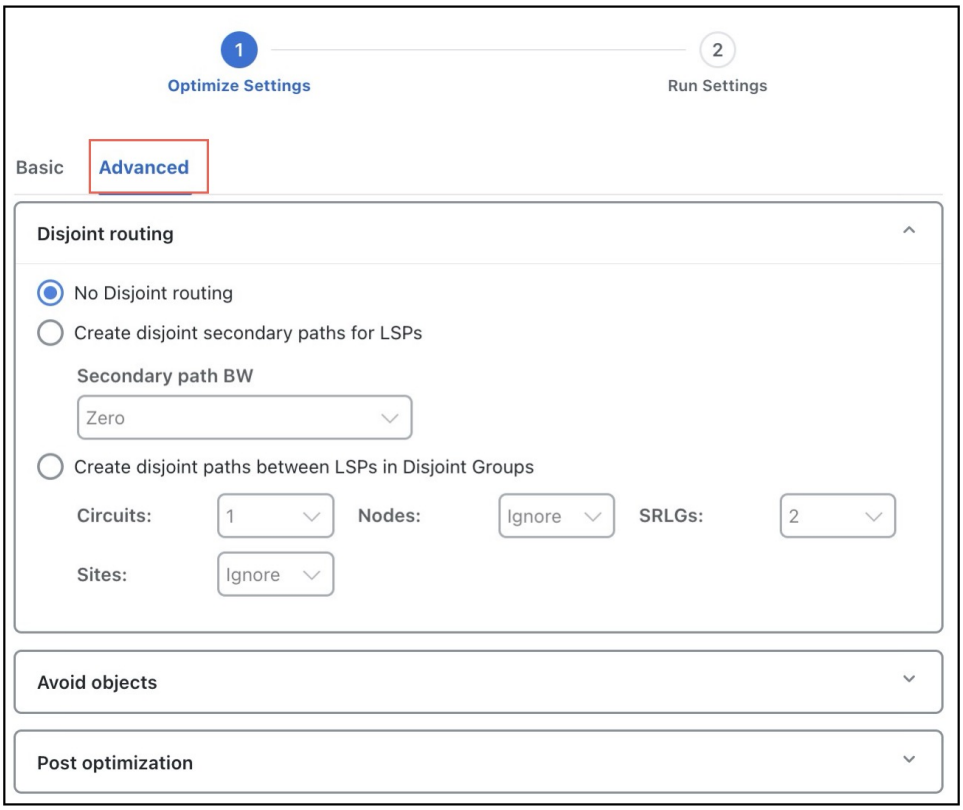
Expand the **LSP parameters** and **Interface parameters** panels to specify the following options:

- **LSP parameters**:
  - Required BW—Specify how to determine the bandwidth required for the primary LSP paths. The options are: Setup BW, Traff meas, and Traff sim.
- **Interface parameters**:
  - BW bound—Specify how much bandwidth the interfaces have available to carry LSPs. Cisco Crosswork Planning tries to route LSPs without surpassing this bound, but exceeds it if necessary. The options are: Resv BW, Capacity, and Capacity sim.
  - Metric—Specify which metric to use for shortest path calculations. The options are: TE Metric and Delay.

Set disjoint groups

In the **Advanced** tab, you can create disjoint paths or use existing disjoint paths as follows. Note that 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.

Figure 99: RSVP-TE LSP optimization advanced options



- **No disjoint routing**—Do not create disjoint secondary paths and do not enforce any disjointness on primary paths.
- **Create disjoint secondary paths for LSPs**—Create secondary LSP paths that are disjoint from the primary LSP paths.

  Additionally, specify 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**—Create disjoint paths between LSPs that are in disjoint groups.

Setting avoidance constraints

The **Avoid objects** section in the **Advanced** tab (see Figure 99: RSVP-TE LSP optimization advanced options, on page 286) lets you to select Nodes, Interfaces, and SRLGs to be avoided when optimizing LSPs.

- Nodes—LSPs are created to route away from the Nodes selected. Default is None.
- Interfaces —LSPs are created to route away from the Interfaces selected. Default is None.
- SRLGs—LSPs are created to route away from the SRLGs selected. Default is None.

Set post-optimization parameters

The **Post optimization** section in the **Advanced** tab (see Figure 99: RSVP-TE LSP optimization advanced options, on page 286) lets you identify how to reset the Setup BW for modified LSPs. The available options are as follows:

- 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

To run the RSVP-TE LSP optimization tool, do the following:

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 creating disjoint paths between LSPs in the same disjoint group, the LSPs must first be added to the disjoint groups in the Edit LSP window. 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

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose any of the following options:

- **Actions > Tools > RSVP LSP optimization > RSVP-TE optimization**

  OR

- **Preset workflows > Perform optimization**, select **RSVP LSP Optimization** as the optimization type, choose **RSVP-TE optimization** from the drop-down list, and click **Launch**.

Step 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, see Select LSP groups, on page 283.

Step 4 Expand the **LSP parameters** panel. In the **Required BW** field, set bandwidth requirements for the primary LSP paths. For more information, see Set optimization parameters, on page 285.

Step 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, see Set optimization parameters, on page 285.

Step 6 Click **Next**.

Step 7 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

  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.

Step 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.

Step 9 Click **Submit**.

What to do next

See Reports, on page 289.

Analyze optimization output

Properties created

In addition to creating the input parameters, as described in Table 25: RSVPTEOpt input parameters, on page 283, the tool generates the properties listed in Table 26: RSVPTEOpt parameters created, on page 289. These properties provide further insight into the optimizations.

Table 26: 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

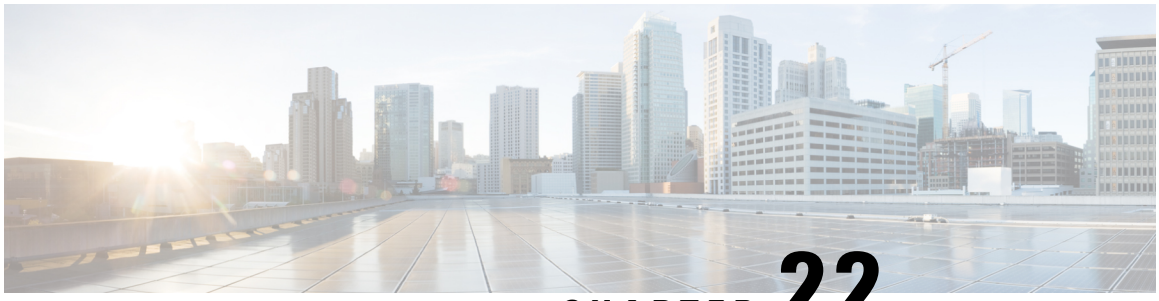
Each time the optimizer is run, a report is automatically generated. You can access this information at any time by choosing **Actions > Reports > Generated reports**. Note that new reports replace the previous ones.

The report summarizes the results of the optimization, such as how many LSPs there are total compared to how many were optimized, as well as how they were optimized. It also summarizes how disjoint priorities were handled and the number of paths that exceed the interface's bandwidth bound.

Optimized LSPs reconfigured

The optimized LSPs are reconfigured with the following 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



CHAPTER 22

Perform Explicit and Tactical RSVP-TE LSP Optimization

The **Explicit RSVP-TE LSP optimization** tool minimizes congestion by optimizing the placement of primary and secondary paths for selected RSVP LSPs. By default, Cisco Crosswork Planning minimizes the utilization across primary paths under normal operation and creates disjoint secondary paths so that a single failure cannot disrupt both paths simultaneously.

The default is to optimize utilization on all interfaces using all RSVP LSPs on those interfaces. Another default is to remove 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 this information later, choose **Actions > Reports > Generated reports**.

This section contains the following topics:

- Run Explicit RSVP-TE LSP optimization, on page 291
- Tactical explicit RSVP-TE LSP optimization, on page 299

Run Explicit RSVP-TE LSP optimization

To run the Explicit RSVP-TE LSP optimization tool, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose any of the following options:

- **Actions > Tools > RSVP LSP optimization > Explicit optimization**

  OR

- **Preset workflows > Perform optimization**, select **RSVP LSP Optimization** as the optimization type, choose **Explicit optimization** from the drop-down list, and click **Launch**.

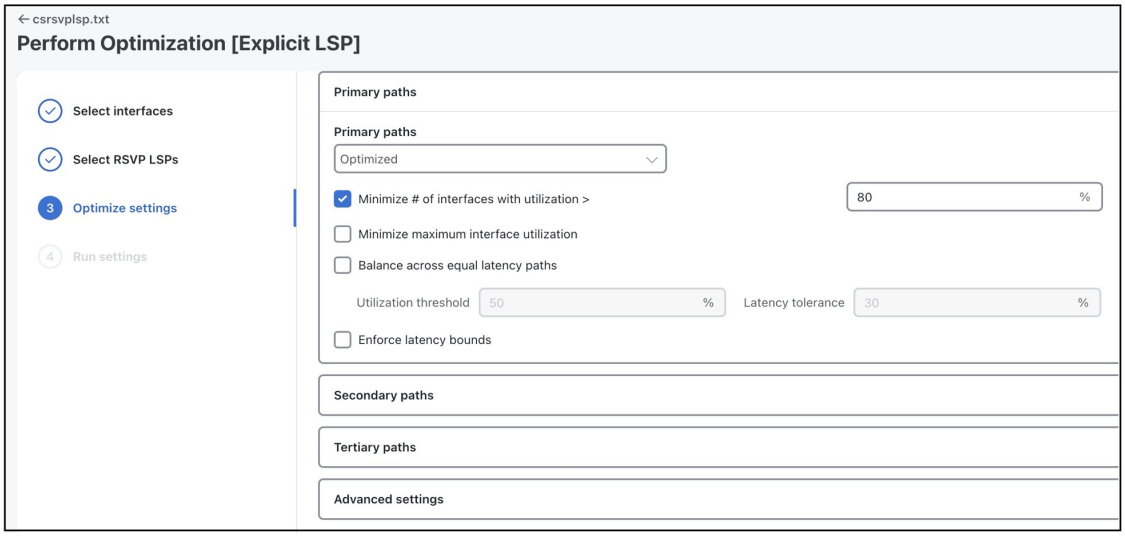
Step 3 Select the interfaces you want to optimize and click **Next**.

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.

Step 4 Select the RSVP LSPs you want to optimize and click **Next**.

Step 5 Specify objectives for primary, secondary, and tertiary paths. For field descriptions, see Table 27: Explicit RSVP-TE LSP optimization options , on page 293.

Figure 100: Explicit RSVP-TE LSP optimization options



Step 6 (Optional) Specify the required advanced settings. For field descriptions, see Table 27: Explicit RSVP-TE LSP optimization options , on page 293.

Step 7 Click **Next**.

Step 8 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

  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.

Step 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.

Step 10 Click **Submit**.

Optimization options

Following table describes the options that are available while running the Explicit and Tactical explicit RSVP-TE LSP optimization tools (**Actions** > **Tools** > **RSVP LSP Optimization** > **Explicit optimization** or **Tactical explicit optimization**):

Table 27: Explicit RSVP-TE LSP optimization options

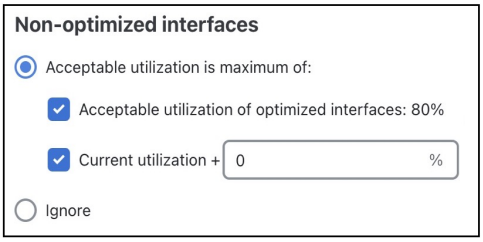
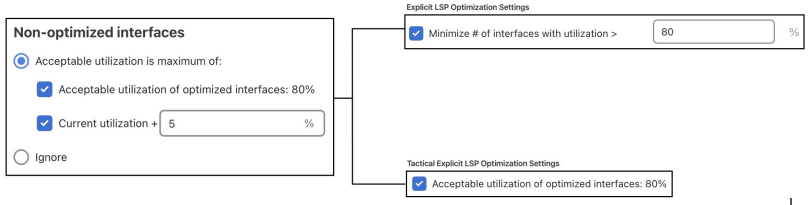
Field	Description
Primary path options	
Primary paths	Define whether to reroute primary RSVP LSP paths: • **Optimized**—Create optimized explicit primary paths based on the objectives selected. • **Keep**—Route RSVP LSPs along the existing primary paths. Primary paths are optimized using the following 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.

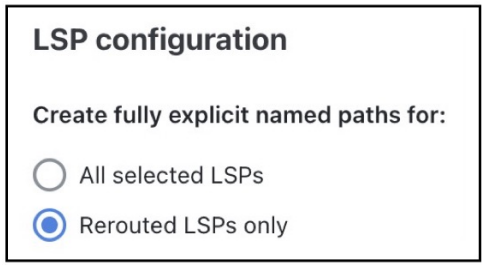
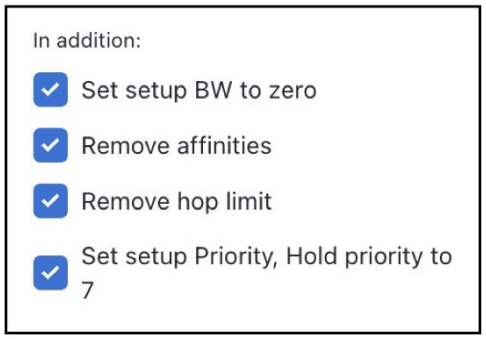
Field	Description
Balance across equal latency paths	Balance utilizations over lower utilized interfaces. For example, use this option to balance utilizations on parallel interfaces between two nodes if the two interfaces have the same latency. • **Utilization threshold**—Keep the number of interfaces with utilization greater than this value as low as possible without significantly increasing the latencies of the primary paths. • **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	
Secondary paths	Define whether and how to route secondary paths. Following are the options: • **Optimized**—Create optimized explicit secondary paths based on the objectives selected. • **Dynamic**—Route secondary paths dynamically. No explicit hops will be created for the path. • **None**—No secondary paths are created; existing paths are removed. 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.

Field	Description
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. • **Circuits**—No circuit is used by both the primary and secondary paths (default = 1). • **SRLGs**—No SRLG is used by both the primary and secondary paths (default = 2). • **Nodes**—No node is used by both the primary and secondary paths (default = 3). • **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. 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. 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. Use option 2 to minimize the number of interfaces with utilizations over the specified percentage across all selected failure scenarios.
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	

Field	Description
Tertiary paths	Define whether to create tertiary paths: • **Dynamic**—Create dynamic tertiary paths. • **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 28: Advanced explicit RSVP-TE 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. 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. • **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. If using the Tactical explicit RSVP-TE LSP optimization tool, this field is equivalent to, and thus only changeable in, the Acceptable utilization ___% field. • **Current utilization + ___%**—Current utilization of non-optimized interfaces plus the added percentage.  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%. The acceptable utilization settings for non-optimized interfaces are: • The primary path utilization threshold is 80%. • The current utilization has 5% added to it.  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. To change this default, click **All selected LSPs**. Cisco Crosswork Planning then creates and routes fully explicit named paths for all selected RSVP LSPs. 
	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. 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. 
Traffic level	The traffic level used in the utilization calculations and optimizations. For information on traffic levels, see Simulate Traffic Flow from Source to Destination Using Demands, on page 71.
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

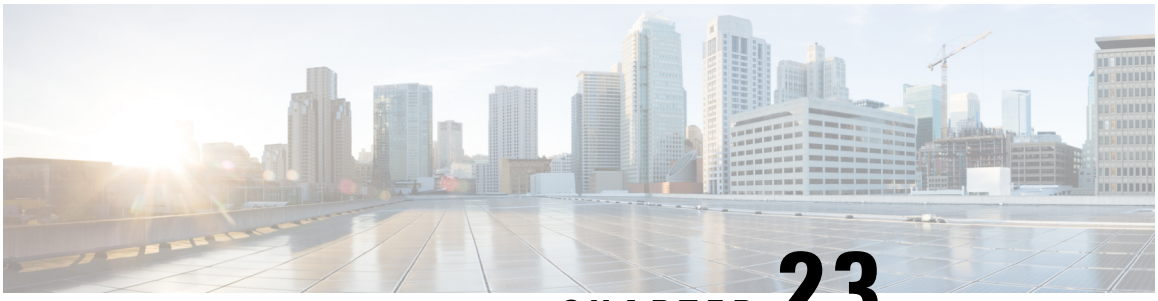
The **Tactical explicit optimization** tool is a reduced version of the Explicit RSVP-TE LSP optimization tool. This tool optimizes only primary paths using the minimum number of path changes required to bring utilizations below an acceptable level. It is useful when you need to reduce congestion in 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 the following options:

- **Actions > Tools > RSVP LSP optimization > Tactical explicit optimization**

  OR

- **Preset workflows > Perform optimization**, select **RSVP LSP Optimization** as the optimization type, choose **Tactical explicit optimization** from the drop-down list, and click **Launch**.



CHAPTER 23

Configure Segment Routing

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**. These 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. SR LSPs use IP routing (and thus, potentially ECMP) between 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*.

This section contains the following topics:

- SR LSP segment types, on page 301
- SR LSP paths, on page 306
- SR LSP routing, on page 306
- SR LSPs and their LSP paths, on page 308
- Create segment lists, on page 308
- Create SIDs, on page 309
- SR-TE protection, on page 312

SR LSP segment types

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 and choose ☰ > **Filter to segment lists**.

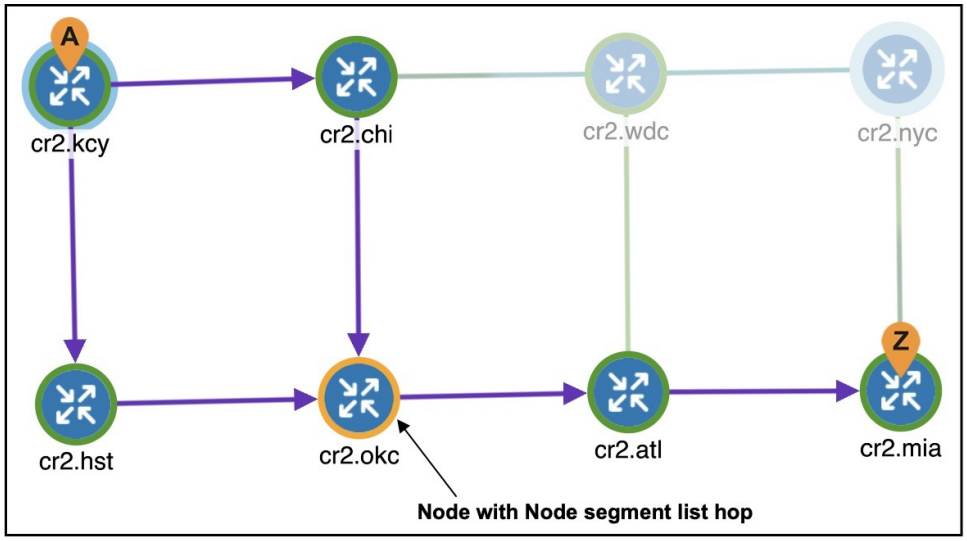
For details on creating segment lists and their hops, see Create segment lists, on page 308.

Note The process of creating and optimizing segment lists and their segment list hops can be automated using the SR-TE optimization tool. For information, see Optimize SR-TE, on page 315. You can also create SR LSPs and their hops by optimizing bandwidth. For information, see Optimize and analyze SR-TE bandwidth, on page 319.

Node segment list hops

When using node segment list hops, the SR LSP takes the shortest path to the specified node. Figure 101: Example node segment list hop, on page 302 shows an example SR LSP from cr2.kcy node to cr2.mia node using one node segment list hop, which is cr2.okc (as indicated by orange circle around the node). Because the IGP metrics are equal, the traffic is routed using ECMP to reach okc.

Figure 101: Example node segment list hop

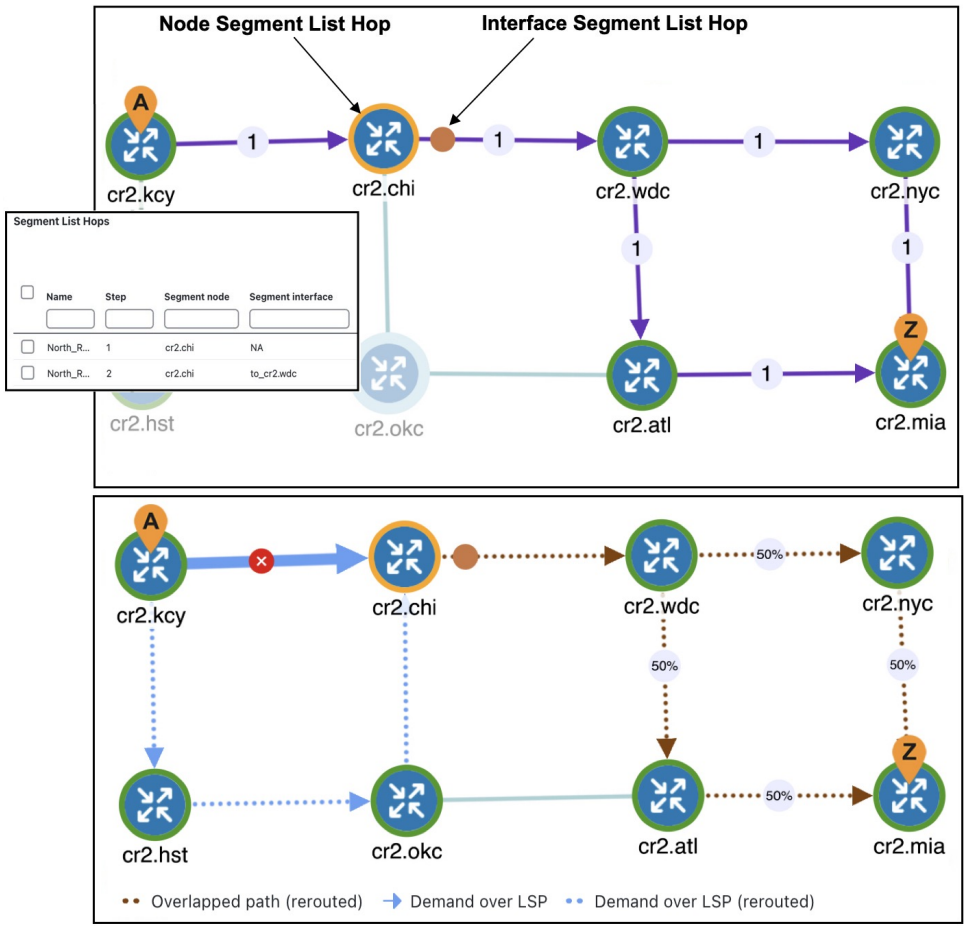


Interface segment list hops

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. In Cisco Crosswork Planning, 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.

Figure 102: Example interface segment list hop, on page 303 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 102: Example interface segment list hop



LSP segment list hops

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. Two rules apply:

- The LSP segment list hop must be an SR LSP, not RSVP LSP.
- An LSP segment list hop cannot reference another SR LSP that directly or indirectly references 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.

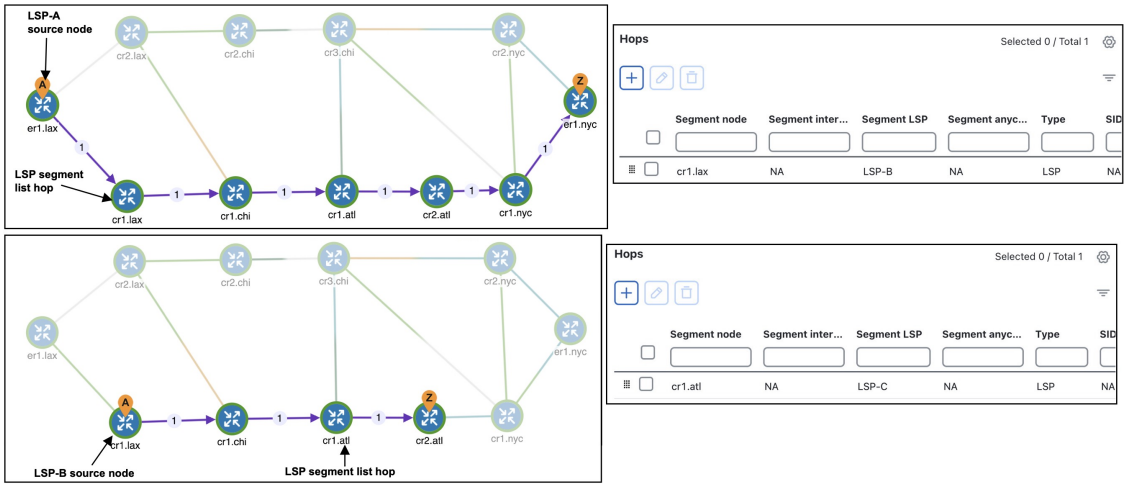
Figure 103: Example LSP segment list hop, on page 304 shows that LSP A (from er1.lax to er1.nyc) contains an LSP B segment list hop, whose source node is cr1.lax. By deselecting LSP A and then selecting LSP B, you see that LSP B also contains an LSP segment list hop, which is LSP C whose source node is cr1.atl.

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  to see the relevant segment list hops.

Figure 103: Example LSP segment list hop

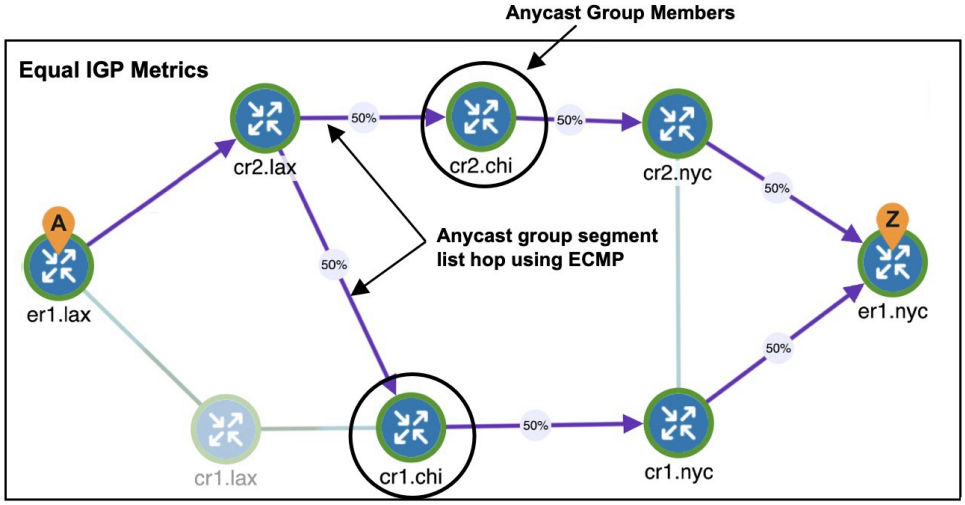


Anycast group segment list hops

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.

Figure 104: Example anycast group segment list hop, on page 305 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 105: Example anycast group rerouting around failure, on page 305 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 104: Example anycast group segment list hop



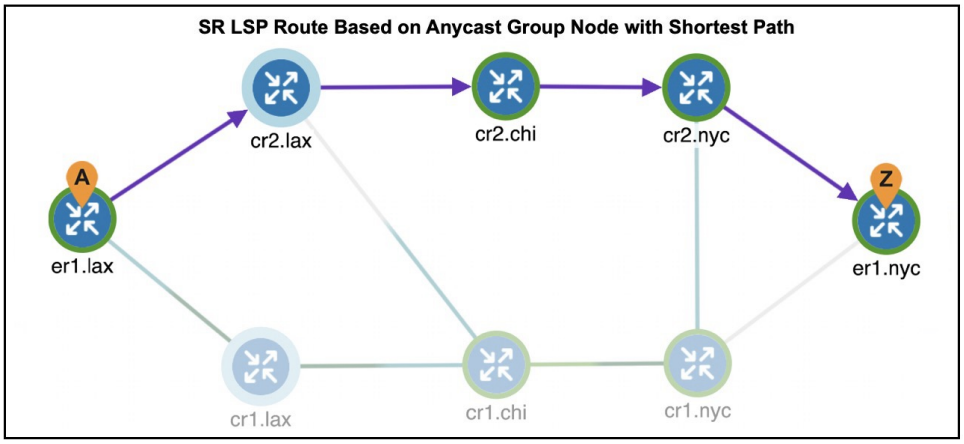
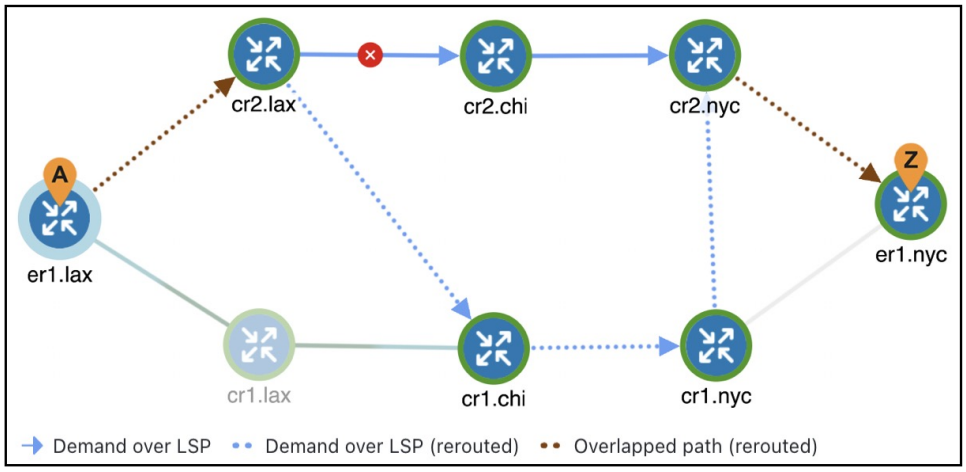


Figure 105: Example anycast group rerouting around failure



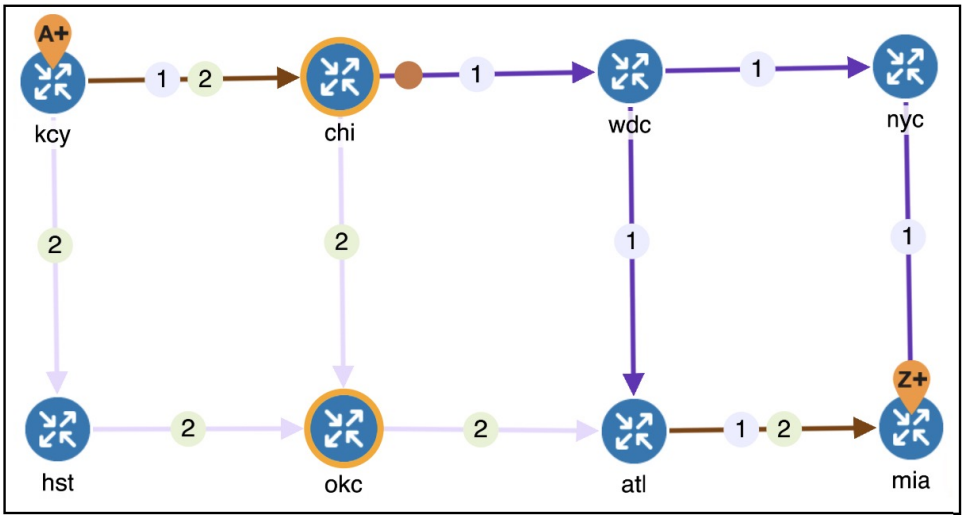
SR LSP paths

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.

Figure 106: Example SR LSP, on page 306 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 106: Example SR LSP



SR LSP routing

Assuming the common LSP mechanisms such as autoroute and forwarding adjacencies have determined that a demand will enter an SR LSP, the LSP traffic is routed as follows. 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 using the Simulate Traffic Flow from Source to Destination Using Demands, on page 71 tool.

- 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

SR LSPs with source and destination nodes in different ASes can route provided the following hop configurations exist:

- There is a node segment list hop on the border of the AS that the LSP is exiting. This node is the exit node.
- One of these hops exists:
  - A node segment list hop exists 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 connects the exit node of one AS to an entry node in another AS.

Using demands with inter-AS SR LSPs requires that two conditions be met:

- The demand must be private to the SR LSP. For information on creating private LSPs, see Configure MPLS Routing, on page 199.
- The demand must have the same source and destination as the SR LSP.

SR LSPs and their LSP paths

SR LSPs and mesh SR LSPs are created the same way RSVP-TE LSPs are created. The difference is that you must set the **Type** property to **SR**. For more information, see Create and visualize LSPs, on page 200.

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 (see Create LSP paths, on page 201). After creating the SR LSP paths, you can set the **Protocol origin** (Edit LSP Path window) 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, see Configure MPLS Routing, on page 199.

Create segment lists

To create segment lists, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions > Insert > LSPs > Segment list**.

OR

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.

Step 3 In the **Name** field, enter the name of the segment list.

Step 4 From the **Site** and **Node** drop-down lists, choose the relevant site and node.

Step 5 In the **Hops** section, add the segment list hops and rearrange them in the order they should be followed. 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 .

Step 6 To continue creating or editing a segment, use the following options and then click **Add**. As needed, 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

To create Anycast groups, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **Others** > **Anycast group**.

OR

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.

Step 3 In the **Group name** field, enter a unique name to identify the anycast group.

Step 4 For each node you want to include in the anycast group, check the check boxes in the **Included** column.

Step 5 Click **Add**.

Create SIDs

Use the following procedures to create new SIDs.

Create node SID

To create a node SID, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **SIDs** > **Node SID**.

OR

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.

Step 3 Select the values for **Site**, **Node**, and **Type** fields.

Step 4 Enter values for **SID** and **SR algorithm**.

Step 5 Use the check box to select **Protected** or **IPv6**, as necessary.

Step 6 Click **Save**.

Create SRV6 node SID

To create SRv6 node SIDs, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **SIDs** > **SRv6 node SID**.

OR

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.

Step 3 Select values for **Site** and **Node** from the drop-down.

Step 4 Enter values for **SID** and **SR algorithm**.

Step 5 Use the check box to select **Protected** as necessary.

Step 6 Click **Save**.

Create interface SID

To create interface SIDs, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **Insert** > **SIDs** > **Interface SID**.

OR

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.

Step 3 Select values for **Site**, **Node**, **Interface**, and **Type** fields.

Step 4 Enter the value for **SID**.

Step 5 Use the check box to select **Protected** or **IPv6**, as necessary.

Step 6 Click **Save**.

Create SRV6 interface SID

To create an SRv6 interface SID, do the following:

Procedure

Step 1 From the toolbar, choose **Actions** > **Insert** > **SIDs** > **SRv6 interface SID**.

OR

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.

Step 2 Select values for **Site**, **Node**, and **Interface** fields.

Step 3 Enter values for **SID** and **SR algorithm**.

Step 4 Use the check box to select **Protected**, as necessary.

Step 5 Click **Save**.

Create flex algorithm

To create Flex Algorithm, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

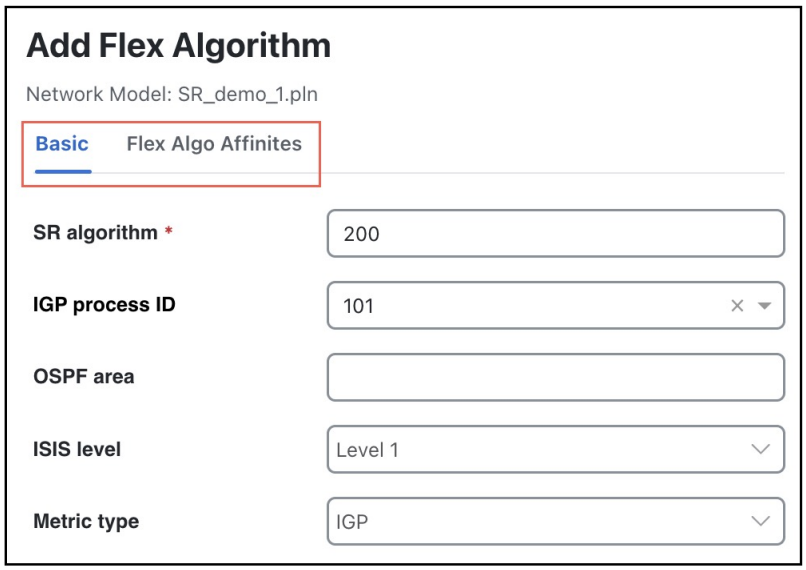
Step 2 From the toolbar, choose **Actions** > **Insert** > **SIDs** > **Flex algorithm**.

OR

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 107: Add flex algorithm window



Step 3 In the **Basic** tab, enter the value for **SR algorithm**.

Step 4 Select values for **IGP process ID**, **OSPF area**, **ISIS level**, and **Metric type**.

Step 5 Click the **Flex Algo Affinities** tab and choose the inclusion or exclusion rule for each affinity.

Step 6 Click **Submit**.

SR-TE protection

If the network contains a configurable topology-independent loop-free alternate (TI-LFA), you can simulate SR-TE tunnels before and after convergence.



Note To see the simulation before convergence, you must be in **Fast reroute** mode (click **Network options** or **Actions** > **Edit** > **Network options**; in the **Simulation convergence mode** area, choose **Fast reroute**). In the IGP and LSP Reconvergence mode, you can only see the after-convergence routes.

Simulate SR-TE tunnels before convergence

For each SR tunnel in a network model, and for each SID in the path, Cisco Crosswork Planning determines the no-failure route.

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.

- 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

Cisco Crosswork Planning routes to each segment list hop through the IGP shortest path under all failure conditions.

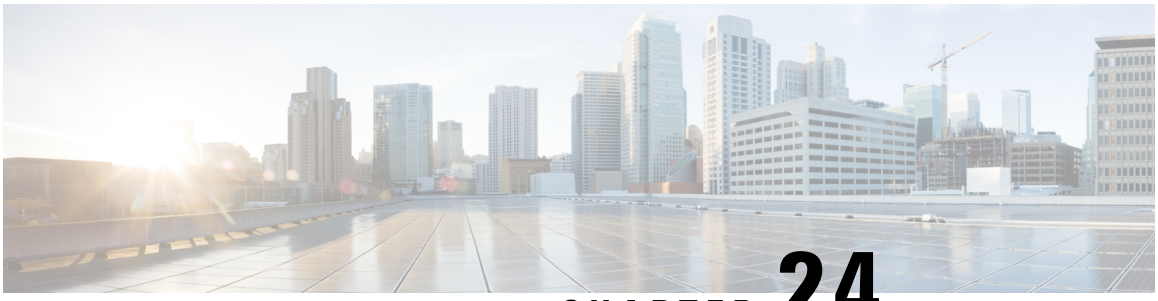


Note An SR LSP might be routable before convergence through a TI-LFA, but unroutable after convergence.

Constraints

- 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.



CHAPTER 24

Optimize Segment Routing

The **SR-TE optimization** 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. The tool sets the routes based on traffic engineering criteria and optimizes a sequence of node or adjacency hops for the SR LSP paths to follow. For more information, see Optimize SR-TE, on page 315.

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. For more information, see Optimize and analyze SR-TE bandwidth, on page 319.

This section contains the following topics:

- Optimize SR-TE, on page 315
- Optimize and analyze SR-TE bandwidth, on page 319

Optimize SR-TE

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 following objectives. This tool supports both Inter-Area and Inter-AS functionalities.

- 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.

Specify inputs for SR-TE optimization

Optimize the path metric

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. Note that for inter-AS SR LSPs, these metrics are calculated end-to-end for the LSP, not per AS. These properties are configurable from the Edit Interface window, and delays can also be set in the Edit Circuit window.

Figure 108: Minimize path metric options



Bound and margin

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, Cisco Crosswork Planning optimizes LSP paths to the best possible solution (lowest total metrics for the LSP path).

Figure 109: 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: 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: 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.

Constraints

The constraints let you 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

To run the SR-TE optimization tool, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose any of the following options:

- **Actions > Tools > SR LSP optimization > SR-TE optimization**

  OR

- **Preset workflows > Perform optimization**, select **SR LSP Optimization** as the optimization type, choose **SR-TE optimization** from the drop-down list, and click **Launch**.

Step 3 Choose the LSPs you want the optimizer to consider.

Step 4 Click **Next**.

Step 5 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, see Optimize the path metric, on page 316.

Step 6 In the **Bound on path length** and **Margin above shortest path** sections, specify the values as per your requirement. For more information, see Bound and margin, on page 316.

Step 7 In the **Constraints** section, specify any restrictions for optimization. For more information, see Constraints, on page 317.

Step 8 Click **Next**.

Step 9 (Optional) In the **Tag updated LSPs with** field, override the defaults for how LSPs are tagged (*SROpt*).

Step 10 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

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.

Step 11 (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.

Step 12 Click **Submit**.

Optimization report

Upon completion, Cisco Crosswork Planning writes a report 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 access this information later, choose **Actions > Reports > Generated reports** and click the **Segment Route TE Optimization** report link in the right panel.

Optimize and analyze SR-TE bandwidth

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.

Select operating modes for SR-TE bandwidth optimization

You can run the SR-TE BW Optimization tool 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 details on how to run this tool, see Analyze congestion under different failure sets, on page 325.
- Operation—The **SR-TE BW optimization operation** tool performs an optimization using a specified set of constraints. For details on how to run this tool, see Optimize bandwidth, on page 323.

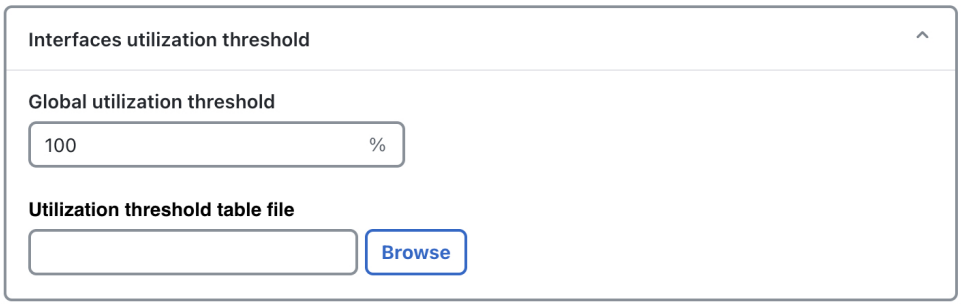
An output plan file is produced with the optimization results. Only a single state of the network is considered. For example, if no objects are failed in the input plan file, the optimization is performed for normal operation. If a circuit is failed in the input plan file, the optimization is performed considering this particular failure scenario.

Specify inputs for SR-TE bandwidth optimization

Interface utilization thresholds

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 110: Interface utilization thresholds panel



- Global utilization threshold—Use this field to specify the global threshold value which applies to all interfaces on the network.
- Utilization threshold tables file—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 namely Node, Interface, and Threshold.

  A sample table is shown below:

  ```
  <InterfaceThresholds>
  Node Interface Threshold
  cr1.nyc to_cr2.wdc 70
  ```

  This above sample 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

Both the analysis and operation modes use the following options for rerouting demands:

Figure 111: Rerouting demands panel

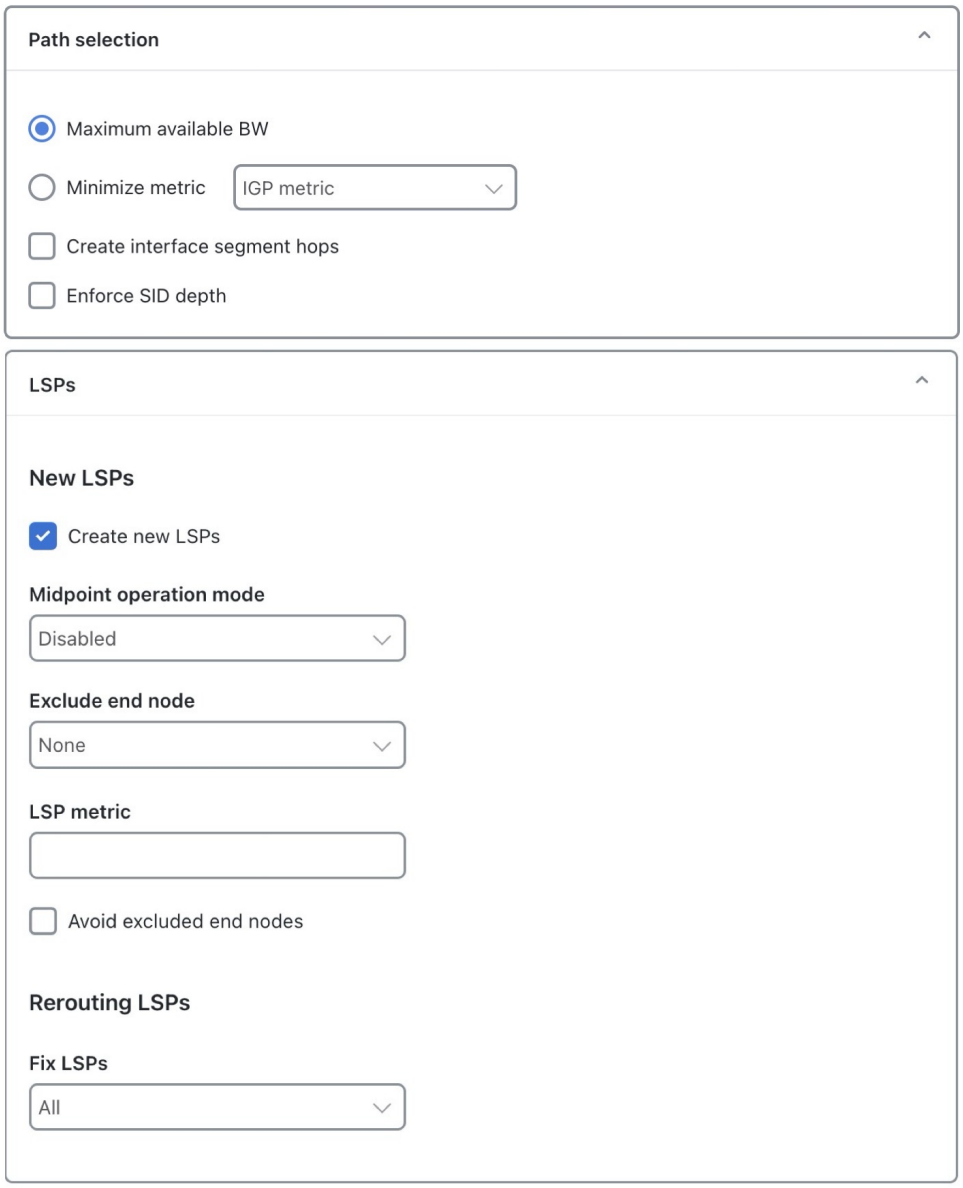


- Traffic steering mode—If you choose **Individual Demands**, each demand is individually carried by a private LSP. If you choose **Autoroute**, multiple demands can be carried by a non-private Autoroute LSP.
- Maximum demands split—Use this option to split demands into smaller demands. If *n* is the 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.

Constraints

Both the analysis and operation modes use the following options for specifying bandwidth constraints:

Figure 112: Constraints section



- **Path Selection**—When choosing demand routes, the SR-TE BW Optimization tool uses one of the following 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—Choose:
      - **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 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 must 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 bandwidth

To run the SR-TE BW optimization operation tool, do the following:

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose any of the following options:

- **Actions > Tools > SR LSP optimization > SR-TE BW optimization operation**

  OR

- **Preset workflows > Perform optimization**, select **SR LSP Optimization** as the optimization type, choose **SR-TE BW optimization operation** from the drop-down list, and click **Launch**.

Step 3 Select whether to optimize bandwidth for all interfaces or selected interfaces.

Step 4 Click **Next**.

Step 5 Specify the utilization threshold of an interface and the routing demand options. For details, see Interface utilization thresholds, on page 320 and Rerouting demands, on page 321.

Step 6 (Optional) In the **Tag updated LSPs with** field, override the defaults for how LSPs are tagged (*SRBWOpt*).

Step 7 Click **Next**.

Step 8 Specify constraints. See Constraints, on page 317.

Step 9 Click **Next**.

Step 10 Select the required segment hop constraints from the **Node segment hop** and **Interface segment hop** drop-down lists, and click **Next**.

Step 11 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

  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.

Step 12 (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.

Step 13 Click **Submit**.

Analyze congestion under different failure sets

You can perform multiple optimizations, one for each failure scenario within a specified failure set. The results of these optimizations are aggregated in a report. This way, you can evaluate the ability of the optimizer to mitigate congestion under different failure scenarios.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose any of the following options:

- **Actions > Tools > SR LSP optimization > SR-TE BW optimization analysis**

  OR

- **Preset workflows > Perform optimization**, select **SR LSP Optimization** as the optimization type, choose **SR-TE BW optimization analysis** from the drop-down list, and click **Launch**.

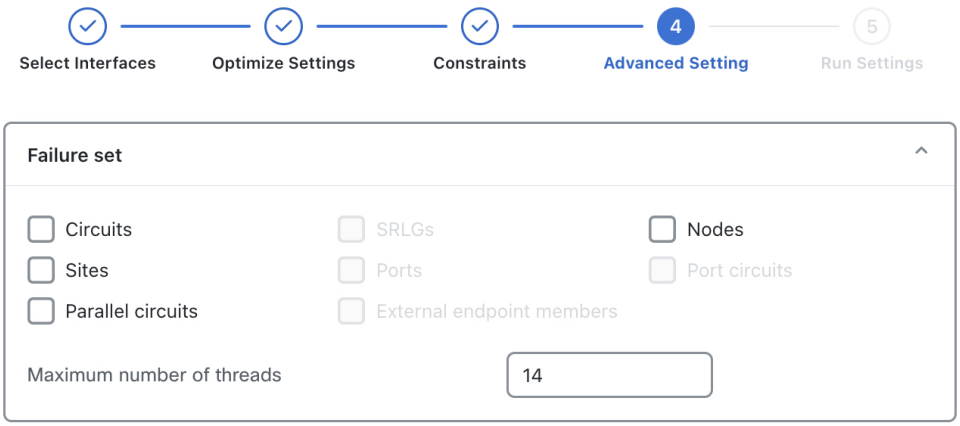
Step 3 By default, the analysis tool presents you with the bandwidth optimization options. Configure these options, as required. For details. see Optimize bandwidth, on page 323.

Step 4 Click **Next** to go to the **Advanced Setting** page. Click the **Failure set** panel.

Step 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 113: Failure sets panel



Step 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.

Step 7 Select the required segment hop constraints from the **Node segment hop** and **Interface segment hop** drop-down lists, and click **Next**.

Step 8 On the **Run Settings** page, choose whether to execute the task now or schedule it for a later time. Choose from the following **Execute** options:

- Now—Choose this option to execute the job immediately. The tool is run and changes are applied on the network model immediately. Also, a summary report is displayed. 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 per your requirement. 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 and using the selected engine profile. You can track the status of the job at any time using the Job Manager window (from the main menu, choose **Job Manager**). Once the job is completed, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

  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.

Step 9 Click **Submit**.

This runs a simulation analysis and creates the optimization report.

Bandwidth optimization report

Upon 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 filenames are editable.

After running the bandwidth optimization, the optimizer creates a summary report of the:

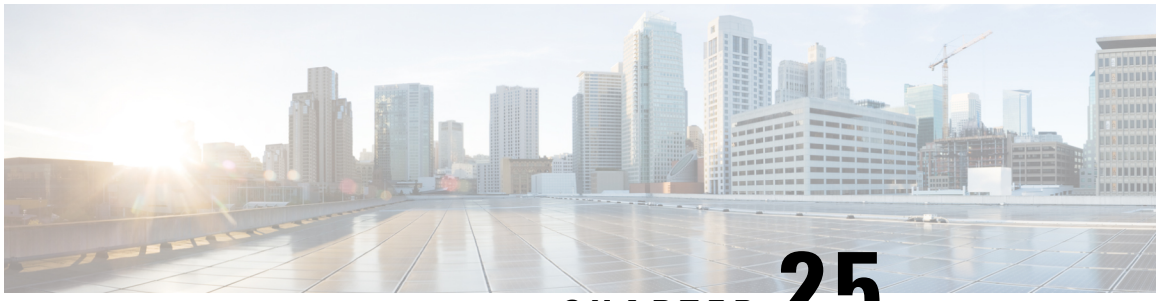
- Number of interfaces that exceeded the utilization both before and after the optimization.
- Maximum utilization both before and after the optimization.
- Number of demands that were rerouted.
- Number of LSPs that were created and rerouted.
- Average segment list length.
- Number of latency bound violations both before and after the optimization.



PART IV

Access Reports, Jobs, Patch Files, and Changeover Tool

- Access Reports, on page 329
- View Jobs, on page 339
- Update Configuration from One File to Another, on page 347
- Create and Use Patch Files, on page 351
- Export Routes, on page 355



CHAPTER 25

Access Reports

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**.

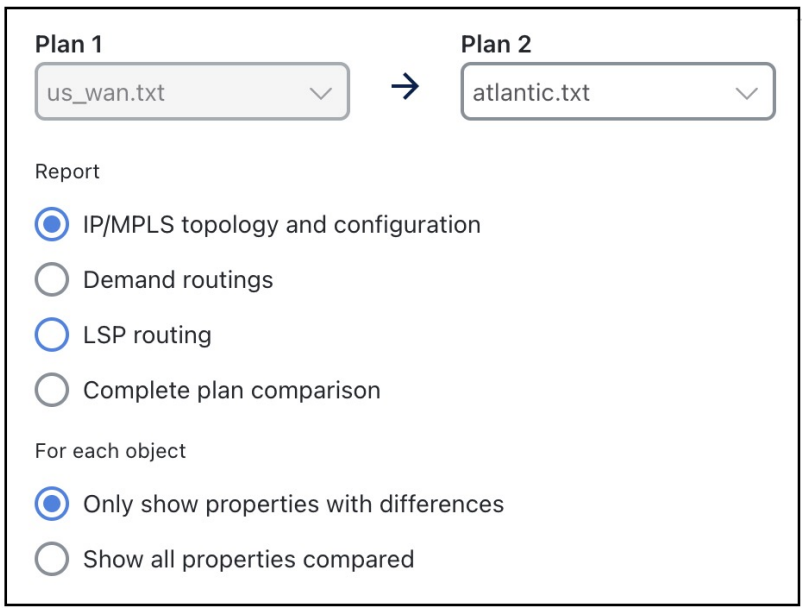
This section contains the following topics:

- Plan comparison reports, on page 329
- Traffic comparison reports, on page 336
- View reports, on page 338

Plan comparison reports

Comparison reports let you compare objects between two plan files to determine:

- Which objects are present in one plan, but not in the other.
- Which objects are present in both plans and which, if any, have different properties.



You can run four types of plan comparison reports:

- IP/MPLS topology and configuration—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, see Table 29: Plans comparison report: IP/MPLS topology and configuration , on page 331.
- Demand routings—Demand paths and the properties that indicate routing changes, such as Path Length and Max Latency. For a complete list of compared properties, see Table 30: Plans Comparison Report: Demand Routings , on page 334.
- LSP routing—Properties of LSPs that indicate routing changes, such as TE Path Metric and Active Path Sim.
- Complete plan comparison—All tables in the table schema are compared, except for internal simulation caching tables.

Create plan file comparison reports

To create plan file comparison reports, do the following:

Procedure

Step 1 Open the two plan files you want to compare. For details, see Open plan files, on page 22.

Step 2 In either of these two plan files, choose **Actions > Reports > Compare plans** from the toolbar.

Step 3 Choose the plan file to which you are comparing the open plan.

Step 4 Choose the type of comparisons you want to generate:

- IP/MPLS topology and configuration
- Demand routings

- LSP routing
- Complete plan files

Step 5 Choose whether to show only properties with differences or whether to show all properties compared.

Step 6 Click **Compare**.

Report columns

Each Cisco Crosswork Planning table has key columns that uniquely identify objects. For example, the key columns in the Interfaces table are the Node and Interface columns, which represent the name of the node containing the interface and the name of the interface itself. In a Plan Comparison report, an object in one plan is identified with an object in another plan only if the key columns of the two objects match. That is, key columns determine if an object is present in both plans or in only one of the plans.

The Plan Comparison report displays three other types of columns. Table 29: Plans comparison report: IP/MPLS topology and configuration , on page 331, Table 30: Plans Comparison Report: Demand Routings , on page 334, and Table 31: Plans Comparison Report: LSP Routings , on page 335 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, the following tables list columns in the order of type first: Key, Info, Diff, and Summ Diff. Properties within each column type are alphabetical.

Table 29: Plans comparison report: IP/MPLS topology and configuration

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
Interfaces	Name	Key	
	Node	Key	

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
	Remote Node	Info	
	Affinities	Diff	
	Area	Diff	
	Capacity	Diff	
	Circuit	Diff	
	Description	Diff	
	FRR Protect	Diff	
	IGP Metric	Diff	
	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	

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
	Model	Diff	
	OS	Diff	
	Vendor	Diff	
LSPs	Node	Key	
	Source	Key	
	Active	Diff	
	Destination	Diff	
	Exclude	Diff	
	FRR Link Protect	Diff	
	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	

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
	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	
	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 30: 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	

Compared Table	Compared Columns	Comparison Column Type	Summary Difference Description
	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 31: 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.

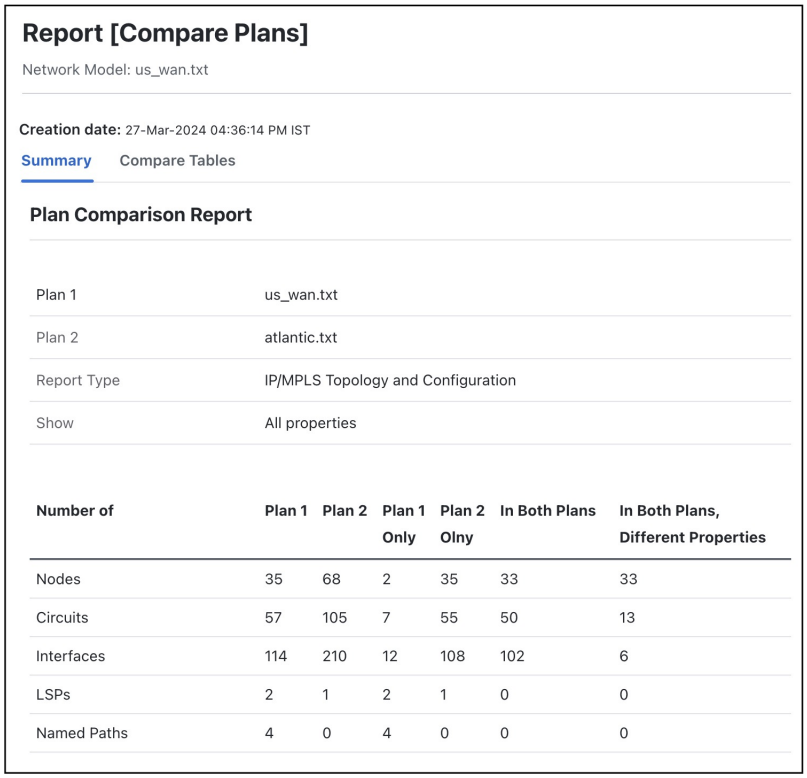
Report sections

The Summary section shows the number of objects in each plan file, the number of objects in both plan files, and the number of objects with property differences (Figure 114: Example plan comparison report summary, on page 336).

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.

An individual section is generated for each object that appears in only one plan file and if there are differences between the properties, a section is generated showing these differences. Thus, there is the potential of generating one to three sections 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, along with their differences noted with "true" or "false".

Figure 114: Example plan comparison report summary



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

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.

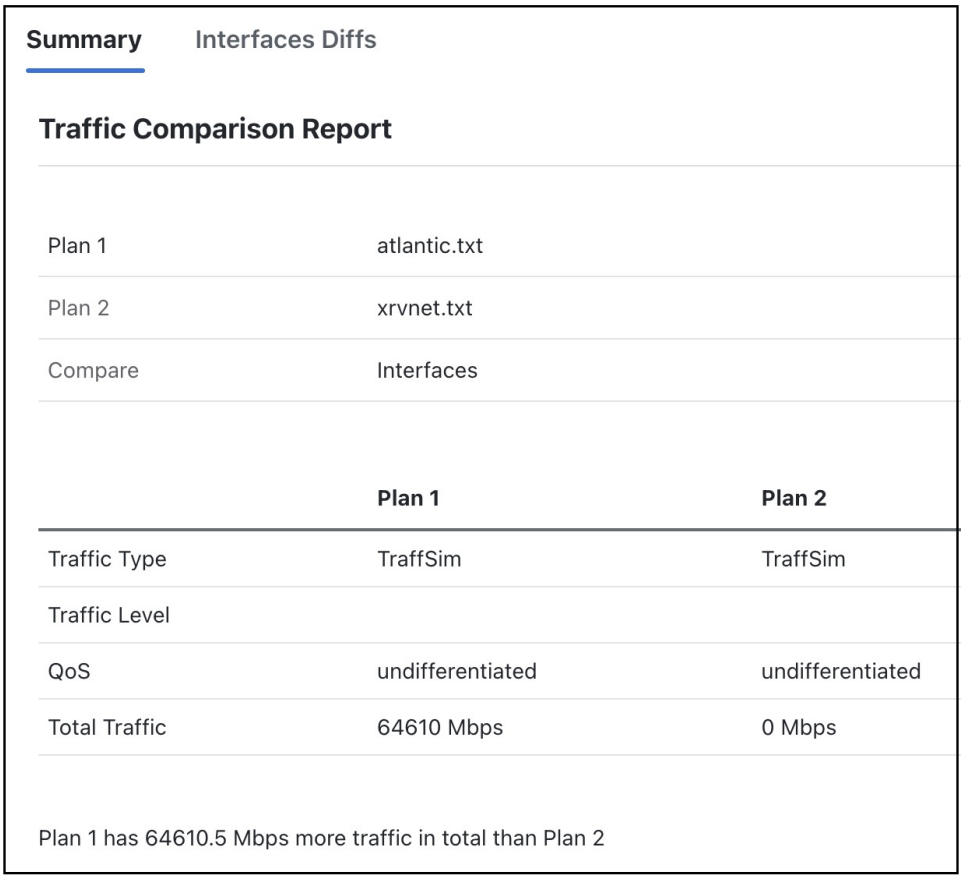
Traffic comparison reports

Traffic Comparison reports compare traffic values within a single plan or between two plan files. You can compare traffic on one object at a time (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 shows a high-level summary of what is being compared. Additionally, a section of the difference on the selected object and traffic columns is generated and named accordingly.



Summary Interfaces Diffs

Traffic Comparison Report

Plan 1	atlantic.txt
Plan 2	xrvnet.txt
Compare	Interfaces

	Plan 1	Plan 2
Traffic Type	TraffSim	TraffSim
Traffic Level		
QoS	undifferentiated	undifferentiated
Total Traffic	64610 Mbps	0 Mbps

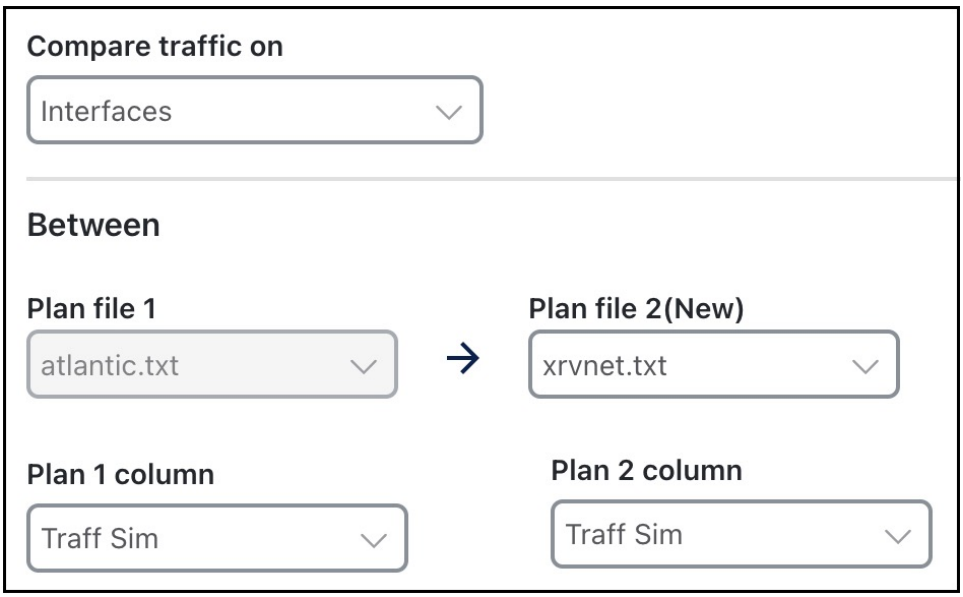
Plan 1 has 64610.5 Mbps more traffic in total than Plan 2

Create traffic comparison reports

To create traffic comparison reports, do the following:

Procedure

Step 1 If comparing traffic from the open plan file to another, open the other plan file.



Step 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.

Step 3 From the **Compare traffic on** drop-down list, choose the object type on which to compare traffic (interfaces, circuits, nodes, and so on).

Step 4 Choose the traffic column in the current plan that you want to compare.

Step 5 To change the plan file 2, choose a different plan file from the Plan file 2 drop-down.

Step 6 Choose the traffic column being compared for Plan 2.

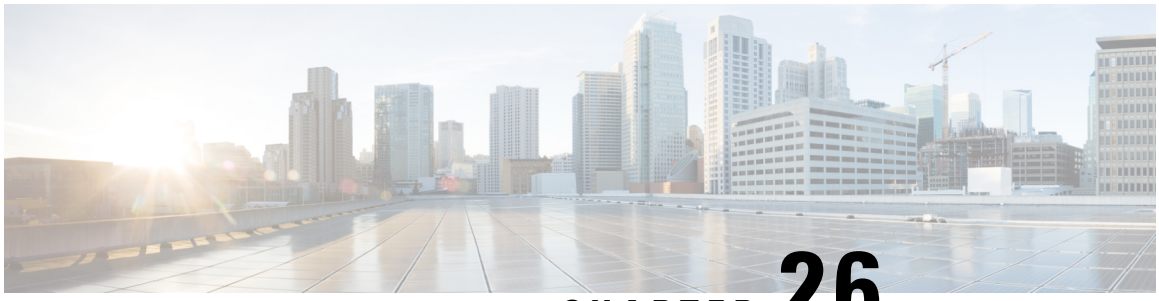
Step 7 Click **Compare**.

View reports

Upon generating a report, it opens automatically. From there, you can view the various sections of the report.

To view a report after having closed it, choose **Actions > Reports > Generated reports** from the toolbar. All reports generated on the network model appear and are selectable from one report window.

If you want to save these reports for later use, you must save the network model (use **Actions** > **File** > **Save** or **Save as**). However, running the report again overwrites the report.



CHAPTER 26

View Jobs

This section contains the following topics:

- Job Manager, on page 339
- View job details, on page 340
- Run tools or initializers using CLI, on page 342
- Run external scripts, on page 343

Job Manager

In Cisco Crosswork Planning, you can schedule tools and initializers to run as background jobs, also known as asynchronous jobs. These jobs execute in the background on the design engine instances configured during deployment. The **Job Manager** page provides details of these jobs.

With background jobs, you can

- set a priority before it is submitted
- schedule to run at a required time, or
- cancel before it moves to the Running state

The **Job Manager** page also provides options to execute jobs using CLI and custom scripts. For details, see Run tools or initializers using CLI, on page 342 and Run external scripts, on page 343.

User role permissions for Job Manager

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, and
- only admin users can cancel (abort) the jobs submitted by any user.

View job details

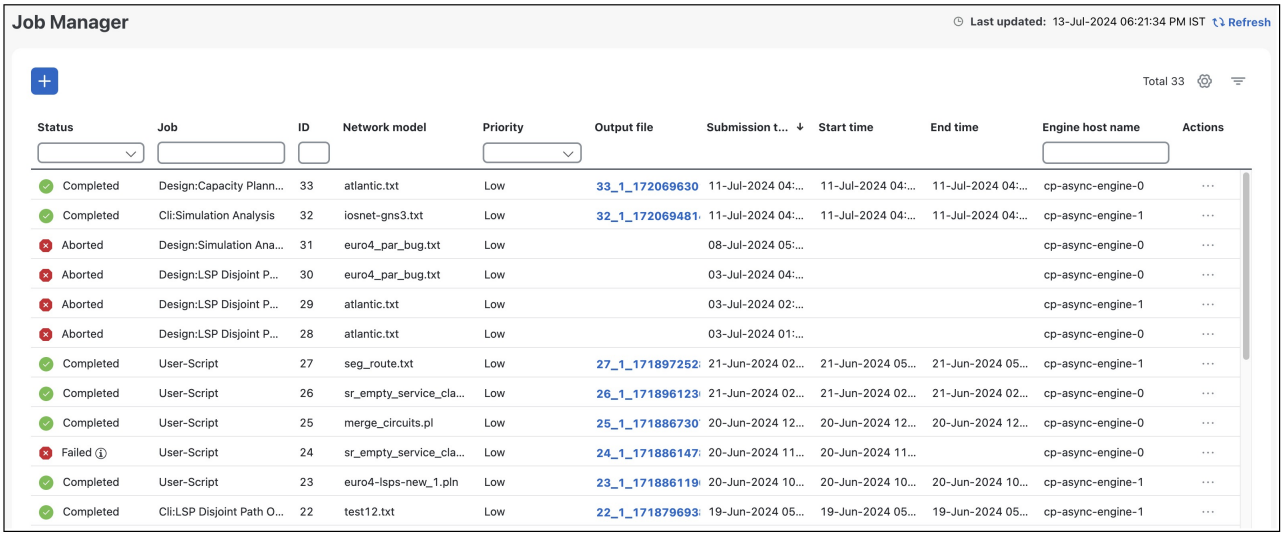
Follow these steps to see the details of the submitted background jobs.

Procedure

Step 1 From the main menu, choose **Job Manager**.

The Job Manager page opens, displaying the list of all jobs submitted as background jobs.

Figure 115: 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.

Step 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 details on all the columns, see Table 32: Job Manager column details, on page 341.

Step 3 The **Output file** column shows the .tar file that is generated after the job has been successfully completed. Click the name of the output file to download it to your local machine and extract the file to view the updated network model. Import this updated network model into the user space to access it (for details, see Import plan files from the local machine, on page 19).

Step 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** to abort the job.

Table 32: 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.
Job	Indicates the type of the job. If a job is submitted via • tools or initializers in the **Network Design** page, the Job column displays "Design: *Tool/initializer name*" • **Using CLI** option in the **Job Manager** page, the Job column displays "Cli: *Tool/initializer name*", and • **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. Download 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.

Column	Description
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.

Run tools or initializers using CLI

In addition to running the tools and initializers from the **Network Design** page, Cisco Crosswork Planning supports running them using CLI.

Follow these steps to run the tools or initializers using CLI.

Procedure

Step 1 From the main menu, choose **Job Manager**.

The Job Manager page opens, displaying the list of all jobs submitted as background jobs.

Step 2 Click + > **Using CLI**.

A list of all available tools and initializers appears.

Step 3 Select the required tool or initializer and click **Next**.

A list of all available network models appears.

Step 4 Select the network models in which you want to run the tool or initializer, and click **Next**.

The number of network models 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.

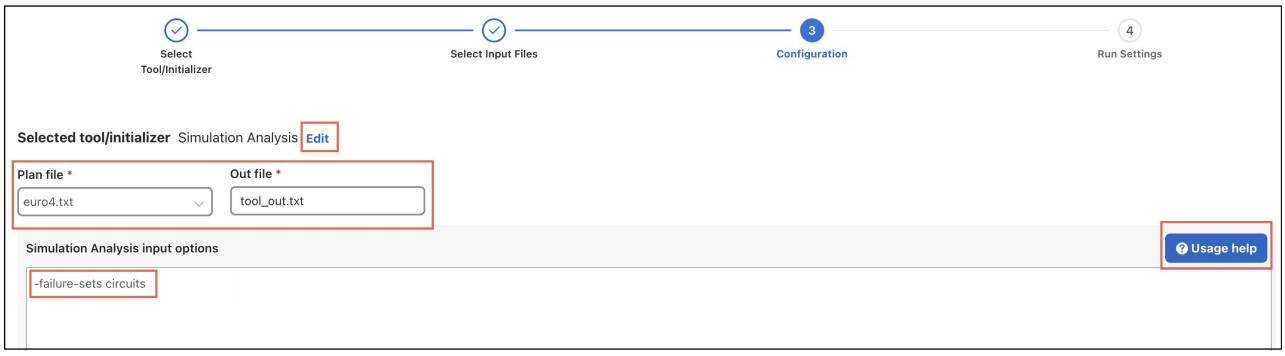
Step 5 Enter the input configuration options in the text field, and click **Next**.

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.



Step 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 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.

Step 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 is completed successfully, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

Run external scripts

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, see Example: Run external scripts, on page 344.

Follow these steps to run the external scripts.

Procedure

Step 1 From the main menu, choose **Job Manager**.

The Job Manager page opens, displaying the list of all jobs submitted as background jobs.

Step 2 Click  > **Using script**.

A list of scripts available in the user space appears.

Step 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.

Step 4 Select the network models on which you want to run the script, and click **Next**.

Step 5 Enter the input configuration options in the text field using the following format, and click **Next**.

```
script_name arg1 arg2
```

Use the **Edit** button to change the script file, if required.

Step 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 the following **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.

Step 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 is completed successfully, download the output file (.tar file), extract it, and import the updated plan file into the user space to access it (for details, see Import plan files from the local machine, on page 19).

Example: Run external scripts

This example describes how to use the external scripts in Cisco Crosswork Planning. 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)

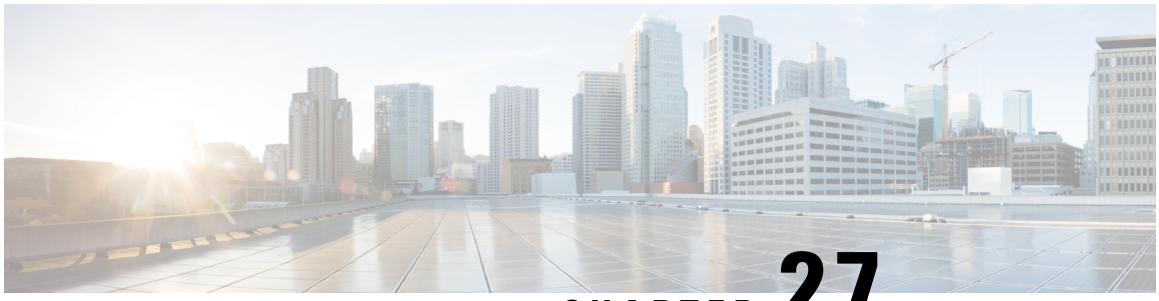
for node in srcNet.model.nodes:
    cnt = 1
    for iface in node.interfaces:
        iface.description = 'My IGP metric is ' + str(iface.igp_metric)
        cnt = cnt + 1

srcNet.write(dest)
```

Upload this script and run it from Job Manager using the steps mentioned in Run external scripts, on page 343. Use the following 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.



CHAPTER 27

Update Configuration from One File to Another

Cisco Crosswork Planning lets you model the topology, routings, and utilizations of a currently operational network. It also allows exploration of modifications to that network. For example, interface metrics can be modified or explicit LSP routings changed to obtain a new routing configuration.

The **Changeover** tool provides a step-by-step sequence of routing configuration changes to move a network safely from an initial configuration to a prespecified final configuration. Cisco Crosswork Planning selects the order of these changes so that there is as little congestion as possible in the network during intermediate configurations, and so that this congestion lasts for as few intermediate steps as possible.

Only certain configuration changes are allowed between the initial and final plan:

- Changes to interface metrics
- Changes to LSP configurations
- Changes of circuits and node from the active to the inactive state, and vice versa

An individual step in the changeover sequence consists of one of the following:

- 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

This section contains the following topics:

- Run the Changeover tool, on page 347
- Analyze reports, on page 349

Run the Changeover tool

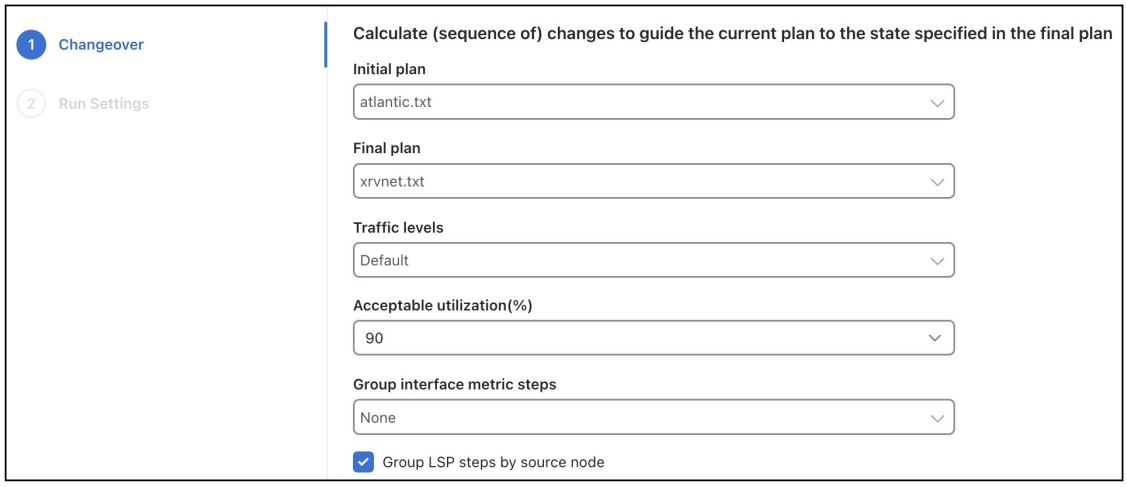
Follow these steps to run the changeover tool.

Procedure

Step 1 Open both the initial and final plan files (see Open plan files, on page 22), and choose the desired plan to view.

Step 2 From the toolbar, choose **Actions** > **Tools** > **Changeover**.

Figure 116: Changeover options



Step 3 Decide on the changeover options to use. For field descriptions, see Table 33: Changeover options , on page 348.

Step 4 Click **Next**.

Step 5 On the Run Settings page, under **Execute**, choose either of the following 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 33: 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.

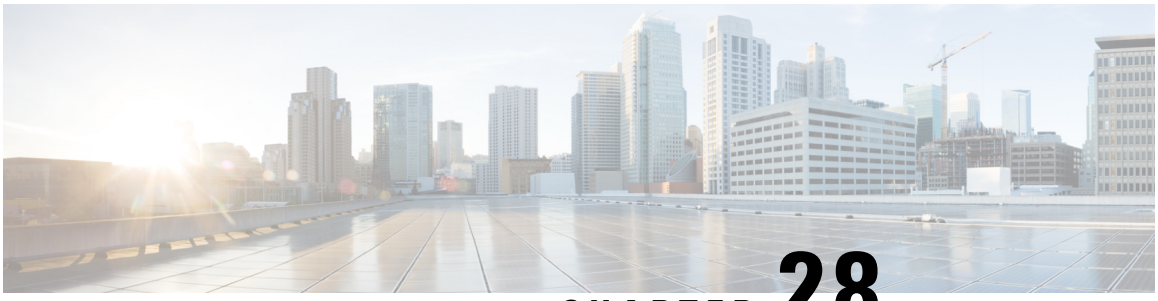
Field	Description
Acceptable utilization (%)	The percentage of maximum acceptable utilization of any interface at any step during the changeover sequence. Changeover tries to keep utilizations below this level, although it is not always possible. Often the utilization spikes upward during the last few 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 the following options: • None—Each Interface Metric change is treated as a single step. • Parallel—Parallel Interfaces are grouped into a single step. • 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.

Analyze reports

To access the report created by Changeover, choose **Actions > Reports > Generated reports** and then click the **Changeover** link from the right panel.

Changeover creates a report that contains the following sections:

- **Summary**: This tab 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**: This tab contains details of the actions taken at each step while running the Changeover tool.
- **Utilizations**: This tab contains all utilizations for all interfaces in the network, for each step in the changeover sequence.



CHAPTER 28

Create and Use Patch Files

A patch file is a compact way to represent differences between plan files. These differences or *patches* can be applied to other plan files or deployed to the network.

You can create, apply, view, and edit patches by choosing **Actions > Tools > Patches** from the toolbar.

A common Cisco Crosswork Planning use case is to create a plan file for deployment. One potential workflow for doing this is the following:

1. Open a plan file from the Cisco Crosswork Planning UI (see Open plan files, on page 22).
2. Duplicate the plan file by using the ••• > **Make a copy** option.
3. Make LSP modifications as needed in the duplicated plan file.
4. Open both plan files, and create a patch file (**Actions > Tools > Patches > Create patch**).
5. Deploy the patch file to the network (**Actions > Tools > Patches > Apply patch**).

This section contains the following topics:

- Create patch files, on page 351
- Apply patch files, on page 352
- View or edit patch files, on page 353

Create patch files

The patch created identifies how to get from the plan file identified in the **From file** field to the plan file identified in the **To file** field. These patch files only include differences in new, modified, and deleted objects that are supported by the Cisco Crosswork Planning YANG model.

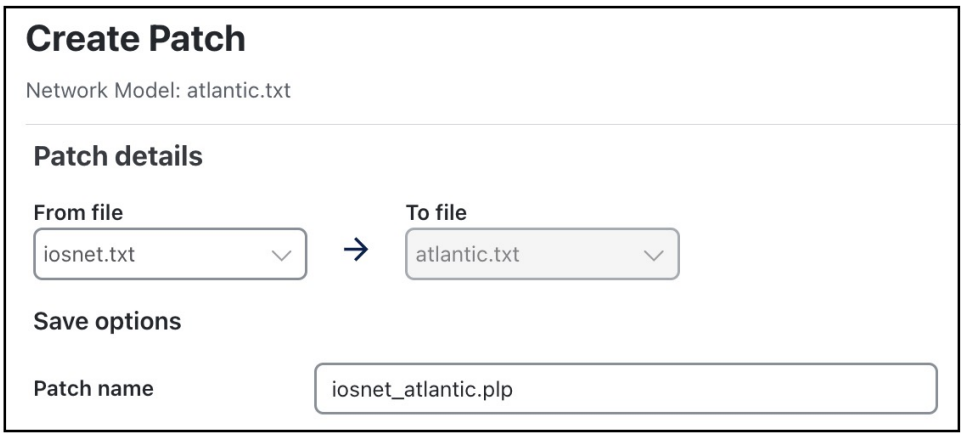
Examples:

- 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, whereas 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.

Procedure

Step 1 Open the two plan files that you are using for the patch file creation. For details, see Open plan files, on page 22.

Step 2 From the plan file used as the To file, which is the plan file you want to achieve using the patch file, choose **Actions** > **Tools** > **Patches** > **Create patch** from the toolbar.



Step 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.

Step 4 The default patch filename combines the From and To plan filenames. If needed, change the filename. It must have a .plp extension. Unless you want the patch file to be overwritten, the patch filename must be unique to its location.

Step 5 Click **Create**.

The patch file is created successfully. Use the **Edit patch** button to edit the patch file. Use the **Download** button to download the file to your local machine.

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.

Apply patch files

To apply patch files to the plan files, do the following:

Procedure

Step 1 Open the plan file in the **Network Design** page. For details, see Open plan files, on page 22.

Step 2 From the toolbar, choose **Actions** > **Tools** > **Patches** > **Apply patch**.

Step 3 Click **Browse** to browse to the location where the patch file resides. Choose from the following 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.

Step 4 Specify whether you want to stop or continue in case of error.

Step 5 (Optional) 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.

Step 6 Click **Run** in the Apply Patch page.

View or edit patch files

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 the ••• > **Open** option to view the patch file details.

You can also use the following steps to view the patch files:

Procedure

Step 1 Open the plan file in the **Network Design** page. For details, see Open plan files, on page 22.

Step 2 From the toolbar, choose **Actions > Tools > Patches > View patch**.

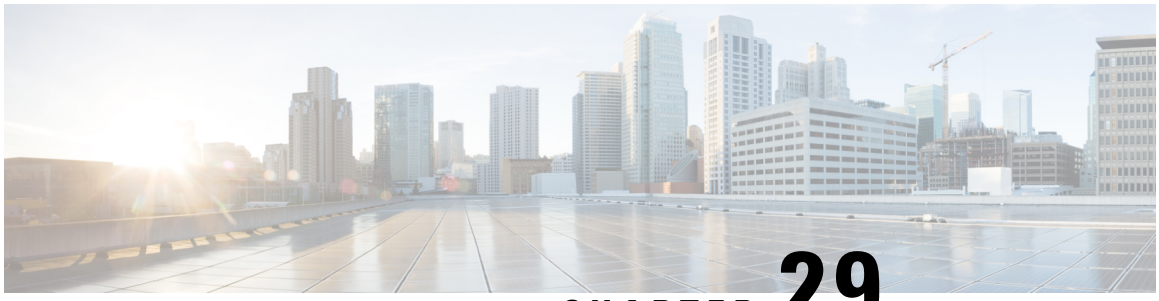
Step 3 Browse to the location where the patch file resides. Choose from the following options:

- 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**.

Step 4 Click the **View details** link.

Step 5 View the patch and then click one of the following:

- **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.



CHAPTER 29

Export Routes

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.

This section contains these topics:

- Route tables, on page 355
- Output table column details, on page 355
- Export routes, on page 356

Route tables

You can export routes using the Cisco Crosswork Planning UI or as a CLI job in Job Manager.

This tool creates one of these tables, depending on what you are exporting. The table generated is output to a .txt file containing only this table. This exported route file is accessible from **My user space > All files**.

- <DemandHops>
- <LSPHops>
- <LSPPathHops>
- <ShortestIGPPathHops>
- <ShortestLatencyPathHops>
- <ShortestTEPathHops>
- <CircuitStyleLSPPathHops>

Optionally, you can fail objects prior to exporting their routes and then export only the routes that changed due to that failed state.

Output table column details

The key columns in the exported table 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 has key columns of Name, Source, Destination, and ServiceClass. Therefore, the <DemandHops> table has these four columns to uniquely identify the demand, plus a Step column to identify the hop sequence.

Each exported hops table includes these columns to specify details of the hop and its position with respect to other hops in the route.

- 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 Routes of demands might be split to account for ECMP routing or multicast. Therefore, demand hops might not be listed sequentially as a hop could be followed by two next hops as the route and traffic split into two.

Example: Export route tables

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 34: 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

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 those objects whose route changed due failure state, fail all relevant objects before exporting 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, see Evaluate Impact of Worst-Case Failures, on page 103.

You can also use the **Export routes** CLI tool to export routes to the hops table. For details on using the CLI tools, see Run tools or initializers using CLI, on page 342.

Procedure

Step 1 Open the plan file (see Open plan files, on page 22). It opens in the **Network Design** page.

Step 2 From the toolbar, choose **Actions** > **File** > **Export** > **Routes**.

The Export Routes page opens.

Step 3 From the **Object** drop-down list, select the object for which you want to export routes.

Step 4 From the **Hop type** drop-down list, select the hop type that you want to export.

Step 5 If exporting only routes that changed due to a failure, select the **Only export routes changed by failure state** option.

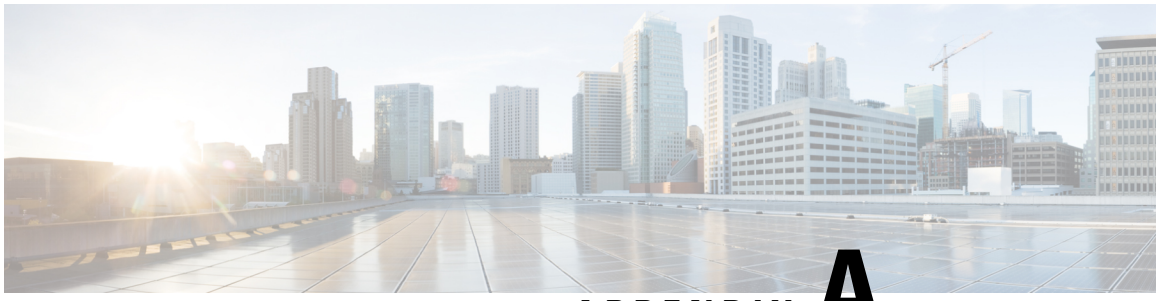
Step 6 In the **File name** field, enter the file name for the output file. By default, the *planfile.object* format is used as the file name.

Step 7 Click **Submit**.

The exported file is saved under **Network Design** > **My user space** > **All files**.

Step 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.



APPENDIX A

Cross Table Filters

This section contains the following topics:

- Supported cross table filters, on page 359

Supported cross table filters

Table 35: Cross table filters in Cisco Crosswork Planning, on page 359 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 **Cross table filter** icon ( ).

Table 35: Cross table filters in Cisco Crosswork Planning

Network Summary table	Cross table filter option	Version introduced
Interfaces	Filter to circuits	7.0
	Filter to nodes > Local	7.0
	Filter to nodes > Remote	7.0
	Filter to nodes > Both	7.0
	Filter to ports	7.0
	Filter to demands > Through all interfaces	7.0
	Filter to LSPs	7.0
	Filter to interface queues	7.1
	Filter to CS-RSVP LSPs	7.1

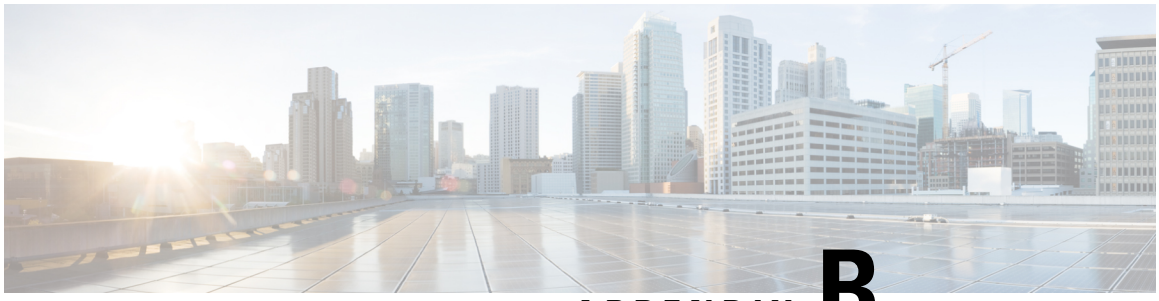
Network Summary table	Cross table filter option	Version introduced
Nodes	Filter to demands > Between selected nodes	7.0
	Filter to demands > Through all selected nodes	7.0
	Filter to LSPs > Between selected nodes	7.0
	Filter to LSPs > Through selected nodes	7.0
	Filter to flex algorithms	7.1
	Filter to node SIDs	7.1
	Filter to SRv6 node SIDs	7.1
	Filter to CS-RSVP LSPs	7.1
Demands	Filter to nodes > Source	7.0
	Filter to nodes > Destination	7.0
	Filter to nodes > Both	7.0
	Filter to LSPs	7.0
	Filter to interfaces	7.0
	Filter to CS-RSVP LSPs	7.1
Circuits	Filter to interfaces	7.0
LSPs	Filter to demands	7.0
	Filter to nodes > Source	7.0
	Filter to nodes > Destination	7.0
	Filter to nodes > Both	7.0
	Filter to LSP paths	7.0
	Filter to interfaces	7.0
	Filter to actual path hops	7.0
	Filter to CS-RSVP LSPs	7.1
LSP Paths	Filter to LSPs	7.0
	Filter to named paths	7.0
	Filter to segment lists	7.0
	Filter to interfaces	7.0
	Filter to CS-RSVP LSPs	7.1

Network Summary table	Cross table filter option	Version introduced
Ports	Filter to ports circuits	7.0
	Filter to interfaces	7.0
Port Circuits	Filter to circuits	7.1
	Filter to ports	7.1
SRLGs	Filter to nodes	7.0
	Filter to interfaces	7.0
External Endpoints	Filter to demands > Source	7.1
	Filter to demands > Destination	7.1
	Filter to external endpoint members	7.1
External Endpoint Members	Filter to nodes	7.0
	Filter to interfaces	7.0
	Filter to AS	7.1
VPNs	Filter to VPN nodes	7.1
	Filter to interfaces	7.0
	Filter to demands	7.0
VPN Nodes	Filter to VPNs	7.1
	Filter to nodes	7.1
Sites	Filter to contained nodes > Direct	7.1
	Filter to contained nodes > All	7.1
	Filter to interfaces > In	7.1
	Filter to interfaces > From	7.1
	Filter to interfaces > To	7.1
	Filter to interfaces > From or To	7.1
	Filter to interfaces > From and To	7.1
Interface Queues	Filter to interfaces	7.1
Segment Lists	Filter to LSP paths	7.0
	Filter to segment list hops	7.1

Network Summary table	Cross table filter option	Version introduced
Segment List Hops	Filter to node SIDs	7.1
	Filter to interface SIDs	7.1
	Filter to SRv6 node SIDs	7.1
	Filter to SRv6 interface SIDs	7.1
	Filter to segment lists	7.1
Named Paths	Filter to named path hops	7.1
Named Path Hops	Filter to interfaces	7.1
	Filter to nodes	7.1
Flex Algorithms	Filter to nodes	7.1
	Filter to interfaces	7.1
	Filter to node SIDs	7.1
	Filter to SRv6 node SIDs	7.1
Node SIDs	Filter to nodes	7.1
	Filter to segment list hops	7.1
	Filter to flex algorithms	7.1
	Filter to node prefix loopbacks	7.1
SRv6 Node SIDs	Filter to nodes	7.1
	Filter to segment list hops	7.1
	Filter to flex algorithms	7.1
	Filter to node prefix loopbacks	7.1
Interface SIDs	Filter to interfaces	7.1
	Filter to segment list hops	7.1
SRv6 Interface SIDs	Filter to interfaces	7.1
	Filter to segment list hops	7.1
	Filter to flex algorithms	7.1

Network Summary table	Cross table filter option	Version introduced
CS-RSVP LSPs	Filter to LSPs	7.1
	Filter to LSP paths	7.1
	Filter to interfaces	7.1
	Filter to nodes	7.1
	Filter to demands	7.1
AS	Filter to nodes	7.1

Cross Table Filters



APPENDIX B

Legend for Plan Objects

This chapter describes the representation of plan objects within the network plot in the Cisco Crosswork Planning Design UI. The objects, such as nodes and circuits, are visually depicted to facilitate network planning and design.

Plan objects	Description
Interfaces and circuits	Interfaces and Circuits, on page 365
Nodes	Nodes, on page 366
Sites	Sites, on page 368
Demands	Demands, on page 369
LSPs	LSPs, on page 370
	LSP paths, on page 370
	LSP named paths, on page 371
	LSP actual paths, on page 372
	SR LSP paths, on page 372

- Interfaces and Circuits, on page 365
- Nodes, on page 366
- Sites, on page 368
- Demands, on page 369
- LSPs, on page 370

Interfaces and Circuits

This section explains how the interfaces and circuits are represented in the network plot.

Table 36: 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. For details on mapping between the color of the links and the traffic utilization, see Traffic utilization, on page 39. When an interface or circuit is selected, the width of the link doubles (). 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 the circuit is not operational because of one of the these reasons: • A node connected to the circuit has failed. • A site that contains a node connected to the circuit has failed. • The circuit is a LAG (port channel) that does not meet its operational requirements (minimum number of available ports or minimum capacity, whichever is surpassed first) due to failure of some of its ports or port circuits. • An SRLG that contains the circuit or contains one of the above objects that are required 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 the circuit is not operational because of one of the following reasons: • A node connected to the circuit is inactive. • A site containing a node connected to the circuit is inactive. • The circuit is a LAG that does not meet its operational requirements (minimum number of available ports or minimum capacity, whichever is surpassed first) due to some of its ports or port circuits being inactive. • An SRLG is inactive, containing the circuit or one of the above objects that are required for circuit operation.

Nodes

This section explains how the nodes are represented in the network plot.

Table 37: Node visualization

Graphic	Description
	A blue icon as shown in the image indicates a node. Light blue outline indicates the node is selected.
	A green border surrounds the node when: • Simulated traffic view—The node is the source or destination of simulated traffic that is greater than 0. • Measured traffic view—The node is the source or destination of measured traffic that is greater than 0. • 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: • Simulated traffic view—The node is not the source or destination of simulated traffic. • Measured traffic view—The node has ingress and egress traffic of 0 (zero). • LSP reservation view—The node is the source or destination of LSPs with 0 (zero) bandwidth reservations.
	A dark gray border surrounds the node when: • Measured traffic view—The node has no ingress or egress traffic. • Worst-case traffic view—Nodes always have gray borders. • 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.

Graphic	Description
	A green icon as shown in the image indicates one or more selected VPN nodes exist on the node.

Sites

Site utilization colors and symbols are based on intrasite objects. Note that intrasite objects include child sites and the objects contained within them.

This section explains how the sites are represented in the network plot.

Table 38: Site visualization

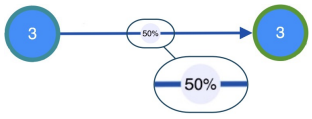
Graphic	Description
	A blue circle indicates a site. The number inside the circle indicates the number of nodes (devices) in the site. Light blue border indicates the site is selected.
	• Measured traffic and Simulated traffic views—The border color is the maximum traffic utilization color of all intrasite interfaces within the site. • 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. • 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 the site has failed.
	A red circle with a white down arrow on the site indicates: • One or more intrasite objects have failed. • One or more nodes or circuits contained within the site are not operational because they are contained in a failed SRLG.

Graphic	Description
	A gray colored site indicates: • The site is inactive. • One or more intrasite objects are inactive. • 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

This section explains how the demands are represented in the network plot.

Table 39: 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. A dotted light blue arrow () indicates the rerouted demand over an LSP.
	A dark green arrow indicates the shortest IGP path of a demand. A dotted dark green arrow () shows the rerouted shorted IGP path of a demand during a failure.
	A light green arrow indicates the shortest latency path of a demand. A dotted light green arrow () shows the rerouted shorted latency path of a demand during a failure.

Graphic	Description
	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. A dotted brown arrow () shows the rerouted aggregated path of a demand during a failure.

LSPs

This section explains how the LSPs are represented in the network plot.

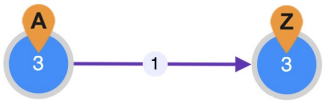
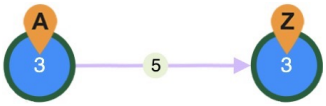
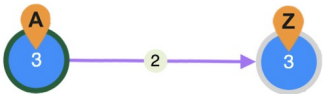
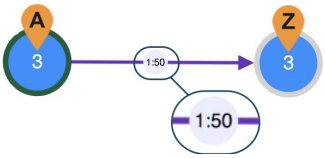
Table 40: LSP visualization

Graphic	Description
	A purple arrows indicates how LSP traffic is routed. “A” and “Z” mark the source and destination site/node of an LSP.
	A dotted purple 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. A dotted light green arrow () indicates the rerouted shorted latency path of an LSP during a failure.
	A gray arrow indicates the shortest TE path of an LSP. A dotted light gray arrow () indicates the rerouted shorted 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. A dotted brown arrow () indicates the rerouted aggregated path of an LSP during a failure.

LSP paths

This section explains how the LSP paths are represented in the network plot.

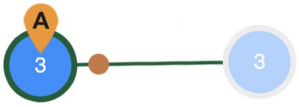
Table 41: LSP path visualization

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. 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. 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. A dotted light purple arrow () shows the rerouted alternate path of an LSP during a failure.
	The value in the circle (here 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%. 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

This section explains how the LSP named paths are represented in the network plot.

Table 42: 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

This section explains how the LSP actual paths are represented in the network plot.

Table 43: 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

This section explains how the SR LSP paths are represented in the network plot.

Table 44: 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.