



Segment Routing Optimization

The [SR-TE Optimization, on page 1](#) 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.

The [SR-TE Bandwidth Optimization, on page 7](#) 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.

This section contains the following topics:

- [SR-TE Optimization, on page 1](#)
- [SR-TE Bandwidth Optimization, on page 7](#)
- [Compressing Segment Lists, on page 14](#)

SR-TE Optimization

You can use the SR-TE Optimization tool to design, capacity plan, and manually configure networks to meet the following objectives:



Note SR-TE Optimization 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, L1 links, L1 nodes, 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.

Optimizing SR-TE LSPs

Choose **Tools > SR LSP Optimization > SR-TE Opt**. Then identify which SR LSPs to optimize: all SR LSPs, those selected before opening the dialog box, or those with specific tags. The default is to optimize all selected SR LSPs, and by default, all SR LSPs are considered selected.

- Parameters for optimizing explicit LSP paths are specified in the Bound, Margin, and Constraints areas of the dialog box. If you do not specify a bound, margin, or constraint, WAE Design generates the best possible SR LSP routes for the given path metric.
- A segment list is created or updated for all SR LSP paths that can be optimized using the minimum number of segment list hops possible.
 - If an LSP has a directly associated segment list, the segment list is deleted.
 - If an LSP has no LSP paths, an LSP path with a path option of 1 is created.
 - If an LSP has no segment list, but it has LSP paths, a segment list is created for the active path if it exists and if not, for the lowest path option (where lowest means the smallest value).
 - If an LSP has associated LSP paths and a segment list associated with either the active path or lowest path option, the segment list hops in the segment list are updated.

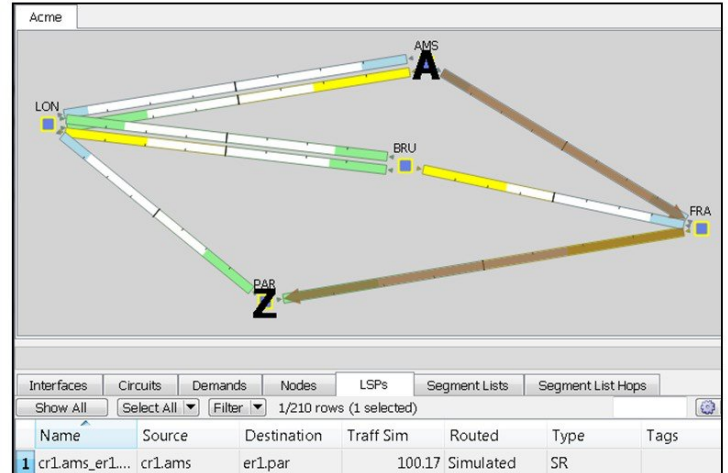
Example: [Figure 1: Example SR LSP Optimization, on page 3](#) shows an SR LSP that normally routes from AMS-FRA-PAR based on IGP metrics. It has no LSP paths. All circuits have a Delay Sim value of 50 except for the AMS-FRA circuit, which has a Delay Sim value of 60. Optimizing this SR LSP based on delay results in rerouting the SR LSP over the shortest achievable delay path, creating an LSP path with a path option of 1, and creating a segment list node hop on the cr.lon node.

- The final hop is either a node hop or interface whose remote node is the destination of the LSP.
- If it is not possible to achieve the lowest possible path without exceeding the specified maximum number of segment list hops, WAE Design provides the best achievable solution within this restriction of maximum permissible hops.
- If it is not possible to avoid all specified nodes for an SR LSP, that SR LSP is not updated.
- If an existing segment list would be better optimized by removing all segment list hops, those hops are removed, though the segment list remains.
- If a previously optimized SR LSP does not need further optimization, it is not changed and its previously set SROpt tag is removed.

- If the source and destination nodes are in different IGP areas, WAE Design creates an explicit hop on the most optimal ABR for each IGP area that the SR LSP traverses.

Figure 1: Example SR LSP Optimization

SR LSP Routed Based on IGP Metrics Before Optimization



SR LSP Route Optimized for Best Path Based on Delay

Minimize Path Metric Delay

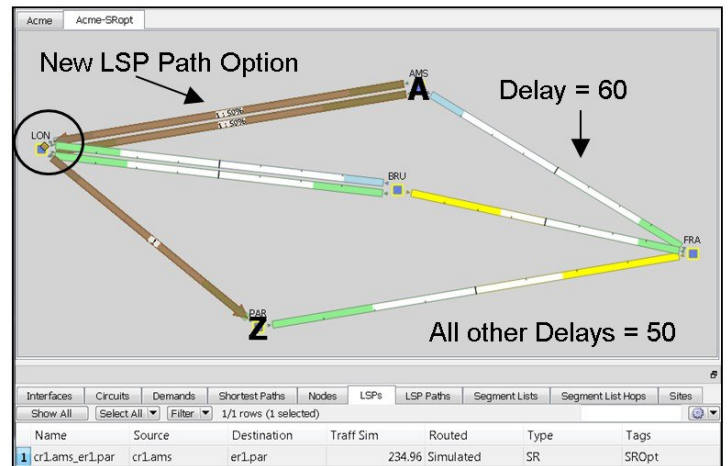
Bound on path length

Fixed

Margin above shortest path

Fixed


Percentage



404671

Optimizing the Path Metric

The Minimize Path Metric selection 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 an interface Properties dialog box, and delays can also be set in a circuit Properties dialog box.



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, WAE Design uses the strictest limitation as the optimization target. If there is no bound or margin specified, WAE Design optimizes LSP paths to the best possible solution (lowest total metrics for the LSP path).

Bound: Maximum Acceptable Path Length

The Fixed Bound entry lets you set the maximum path metric that is acceptable. WAE Design 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 (as in [Figure 1: Example SR LSP Optimization, on page 3](#)), 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 ([Figure 2: Example SR LSPs Optimized with Margins, on page 5](#)). 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) \times $.10$ (percentage margin) = 20 , so all paths greater than 220 must be optimized

Example: If an SR LSP route has a delay of 110 (as in [Figure 1: Example SR LSP Optimization, on page 3](#)), 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 ([Figure 2: Example SR LSPs Optimized with Margins, on page 5](#)).

Fixed Bound = 120

100 (shortest path) + 15 (fixed margin) = 115

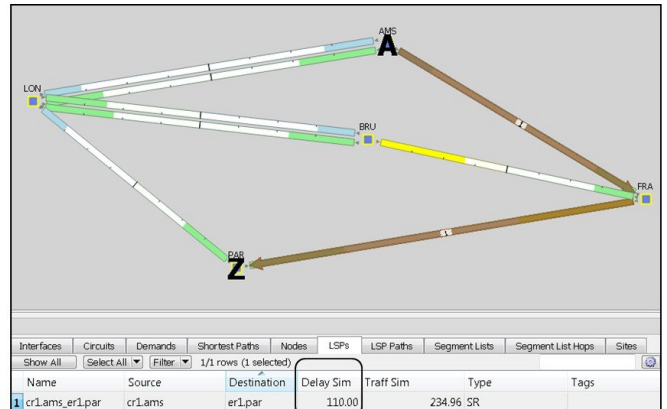
100 (shortest path) \times $.05$ (percentage margin) = 5 , so all paths greater than 105 are optimized because it is the strictest margin

Figure 2: Example SR LSPs Optimized with Margins

Shortest Delay = 100
 SR LSP Delay = 110

SR LSP Is Not Optimized

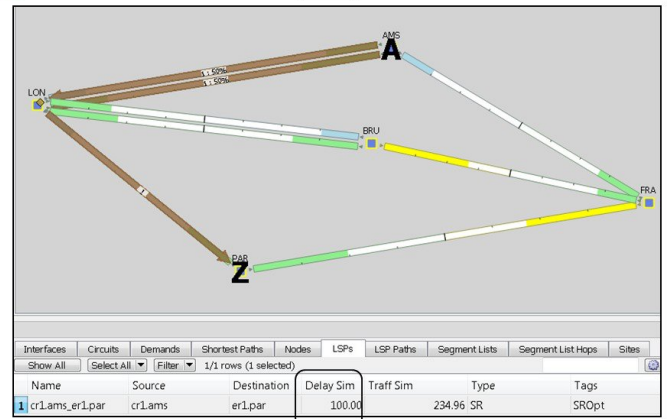
Minimize Path Metric	Delay
Bound on path length	
Fixed	
Margin above shortest path	
Fixed	10
Percentage	



SR LSP Is Optimized

Minimize Path Metric	Delay
Bound on path length	
Fixed	120
Margin above shortest path	
Fixed	15
Percentage	5

Strictest →



404572

Constraints

The constraints let you specify restrictions for the optimizations.


- **Maximum Segment List Hops**—The maximum number of segment list hops that any given segment list can contain after optimization. If no value is specified, WAE Design creates as many hops as needed to optimize the SR LSP.

- Avoid—Do not permit optimized segment lists to route through the selected or tagged objects (nodes, interfaces, L1 links, L1 nodes, or SRLGs). This constraint is useful when modeling dual-plane topologies that route disjoint LSPs.

Example: [Figure 3: Example of Avoiding Tagged Nodes During SR-TE Optimization, on page 6](#) shows an SR LSP from cr1.chi to cr1.mia that is routed based on its shortest IGP route. It also shows the optimized SR LSP path if it were to avoid cr1.okc based on its tag.

- Restrict segment nodes 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.

Figure 3: Example of Avoiding Tagged Nodes During SR-TE Optimization



Using SR Tags and Creating New Plan Files

These options identify how optimized SR LSPs are tagged, and whether to create new plan files with the results:

- Tag updated LSPs with—Create and add tags to all SR LSPs rerouted during optimization. Using tags makes it easier to find the changed SR LSPs that are deployed to the network or to find those SR LSPs that need to be reconfigured in the actual network according to the optimization results. By default, newly optimized LSPs are tagged as *SROpt*.
- New plan for result—This default option creates a new plan with the results of the optimization. Unless a name is specified, WAE Design attaches an *-SROpt* suffix to the current plan filename. If not selected, WAE Design changes the current plan file with the updated information. Saving this plan file then simplifies the process of identifying which LSPs to reconfigure in the network.

Optimization Report

Upon completion, WAE Design 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 **Window > Reports**.

SR-TE Bandwidth Optimization

Table 1: Feature History

Feature	Release Information	Description
SR-TE Bandwidth Optimization Enhancements	Cisco WAE Release 7.4.0	<ul style="list-style-type: none"> The SR-TE Bandwidth Optimization Tool can now be used to specify utilization threshold values for individual interfaces. A new field called Utilization Threshold Tables file allows you to upload a .txt file containing threshold values for specific interfaces in the form of a table. SR-TE Bandwidth Optimization Tool uses these details to reduce utilization on the link to the specified threshold value. Users can now know if they have enough capacity in their network for congestion to be mitigated under different failure scenarios. The SR-TE Bandwidth Optimization Tool in WAE Design is enhanced to include Congestion proximity constraint. This constraint specifies that the new SR LSP source/destination nodes may differ from the demand source/destination nodes and the Source/Destination nodes closer to congestion points are selected.

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 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 SR-TE Bandwidth Optimization tool supports Inter-Area functionality.

Operating Modes

You can run the SR-TE BW Optimization tool in two modes:

- **Analysis**—Performs multiple optimizations, one for each failure scenario within the specified failure set. The results of these optimizations are aggregated in a report.
- **Operation**—Performs an optimization using a specified set of constraints.

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.

Interface Utilization Thresholds

You can specify utilization threshold of an interface using the SR-TE Bandwidth 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.

- **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.

Constraints

Both the analysis and operation modes use the same dialog box for specifying bandwidth constraints:

- Path Selection—When choosing demand routes, the SR-TE BW Optimization tool uses one of two 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, or IGP metrics. All of these properties are configurable from an interface Properties dialog box, and delays can also be set in a circuit Properties dialog box.
 - Create Interface Segment Hops—If checked, the tool considers interface hops when trying to drive down utilization.
- LSPs:
 - Create new LSPs—If checked, new private SR LSPs with optimized routing can be created. If unchecked, new LSPs are not created.
 - Midpoint Optimization 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 Nodes—Nodes selected are not considered as source/destination nodes for SR LSPs. Default is **None**. The nodes selected here are ignored if Midpoint Optimization Mode is set to **Disabled**.
 - 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.
- Rerouting Demands:
 - Traffic Steering Mode—If Individual Demands is selected, each demand is individually carried by a private LSP. If Autoroute is selected, multiple demands can be carried by a non-private Autoroute LSP.
 - Maximum Demand 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 a demand's Properties dialog box.) If the latency of rerouted demands exceeds the configured bound, the tool does not reroute demands on congested interfaces.

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.

Examples

[Figure 4: Network Before Optimization, on page 11](#) shows a network before bandwidth optimization. It has no LSPs, and there is congestion between San Jose (sjc) and Los Angeles (lax), as well as between Los Angeles (lax) and Houston (hst). Sorting the Util Sim column shows the top five highest utilizations with two of them higher than 90%. Notice that the demand from sjc to hst routes directly over these highly used interfaces.

Figure 4: Network Before Optimization



Interfaces					
Node	Interface	Remote Nod	Traff Sim	Capacity Sim	Util Sim
cr1.lax	to_cr2.hst	cr2.hst	2540.09	2488.00	102.09
cr2.sjc	to_cr2.lax	cr2.lax	2270.78	2488.00	91.27
cr2.nyc	to_cr1.chi	cr1.chi	2144.68	2488.00	86.20
cr2.lax	to_cr2.sjc	cr2.sjc	1991.40	2488.00	80.04
cr2.atl	to_cr1.wdc	cr1.wdc	1989.63	2488.00	79.97


4.06234

To optimize the maximum available bandwidth to be below 90% and to enforce latency bounds, the following parameters are set:

- Congestion threshold = 90%.
- Base the rerouted paths on maximizing available bandwidth.
- Enforce latency bounds.

The result of this optimization is that no interfaces have utilization over 90%, and the interfaces with the highest utilization are different interfaces than prior to the optimization. [Figure 5: Network After Reducing Utilization to Below 90%, on page 12](#) shows this redistribution of traffic across the network. The only demand rerouted was the one between sjc and hst, which you can compare to [Figure 4: Network Before Optimization, on page 11](#).

Figure 5: Network After Reducing Utilization to Below 90%



Interfaces					
Node	Interface	Remote Node	Traff Sim	Capacity Sim	Util Sim
cr2.nyc	to_cr1.chi	cr1.chi	2144.68	2488.00	86.20
cr2.kcy	to_cr1.hst	cr1.hst	2124.55	2488.00	85.39
cr1.sjc	to_cr1.kcy	cr1.kcy	2088.27	2488.00	83.93
cr2.lax	to_cr2.sjc	cr2.sjc	1991.40	2488.00	80.04
cr2.atl	to_cr1.wdc	cr1.wdc	1989.63	2488.00	79.97


Number of Rerouted Demands	: 1	

	Before	After
Number of congested interfaces:	2	0
Maximum interface utilization:	102.09%	86.2%

Figure 6: Optimized Network Under Failure, on page 13 shows this same network, but with a failed circuit between Seattle (sea) and Chicago (chi). There are three interfaces with over 100% utilization. Here, the SR-TE BW Optimization tool was rerun using the same parameters as above except the congestion threshold was set to 100. Although congestion cannot be completely mitigated, the tool minimizes congestion as much as it can, leaving only one interface congested (over 100% utilized).

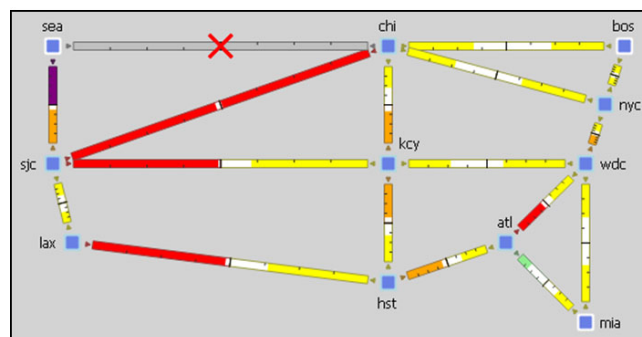
4/8/2017

Figure 6: Optimized Network Under Failure



Util Sim
136.26
129.35
113.98
87.53
86.20
85.39
83.93
80.04

Before Optimization



Util Sim
129.35
99.23
98.46
96.96
96.53
90.21
88.57
87.53

After Optimization

Number of Rerouted Demands	:	4

	Before	After
Number of congested interfaces:	3	1
Maximum interface utilization:	136.26%	129.35%

408298

Optimize Bandwidth

Step 1 Skip this step if you are using tags. You can also skip this step if you want to run the tool on all interfaces, and if you are not selecting demands to fix.

If you are fixing selected demands or running the optimization on selected interfaces, select them before opening the dialog box.

- If you are selecting one or the other, do so from their respective Demands or Interfaces tables.
- If you are selecting both, make your selection in one table, and then press and hold the Alt key (Option key on a Mac) while making selections in another table.

Step 2 Choose **Tools > SR LSP Optimization > SR-TE BW Opt > Operation**.

Step 3 Select whether to optimize bandwidth for all interfaces, selected interfaces, or tagged interfaces.

Step 4 Specify constraints. See [Constraints, on page 5](#).

Step 5 Specify whether to create and add tags to all SR LSPs rerouted during optimization. You can edit the *SRBWopt* default.

- Step 6** Specify whether to create a new plan with the results of the optimization. Unless a name is specified, WAE Design attaches an *-SRBWopt* suffix to the current plan filename. If unchecked, WAE Design changes the current plan file with the updated information.
- Step 7** Click **OK**.

Analyzing 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.

- Step 1** Choose **Tools > SR LSP Optimization > SR-TE BW Opt > Analysis**. By default, the analysis tool presents you with the Basic bandwidth optimization options.
- Step 2** Click the **Failure Set** tab.
- Step 3** Select 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.
- Step 4** (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 5** Click **OK**.
This runs a simulation analysis and creates the optimization report.

Compressing Segment Lists

Some routers might not support long segment lists; however, long segment lists might be required to achieve the desired route.

The Compress Segment Lists tool takes LSPs as input and creates new segment lists that adhere to a maximum length. Any segment lists that exceed the maximum length are replaced by new LSPs, LSP paths, and segment lists. A newly created LSP is set as the final segment list hop of the previous LSP; this process repeats until the initial LSP reaches its destination. The tool also creates new LSPs when the route crosses an AS boundary.

- Step 1** Choose **Tools > SR LSP Optimization > Compress Segment Lists**.
- Step 2** Choose which segment lists to compress: all (the default), those selected before opening the dialog box, or those with specific tags.
- Step 3** In the Maximum Segment List Length field, enter a value of at least 3 hops. The default limit is 9 hops.
- Step 4** Check the **Remove Orphaned Segment Lists** check box if you want to remove segment lists that are not used by other LSP paths or SR FRRs as a result of the optimization.

- Step 5** In the **Tag updated LSPs with** field, enter a tag to add to all segment lists that are compressed during optimization. Using tags makes it easier to find the segment lists that need to be compressed in the actual network according to the optimization results. By default, newly compressed segment lists are tagged as *SLOpt* .
- Step 6** Check or uncheck the **New plan for result** check box. If checked, WAE Design generates a new plan file with the results of the optimization and appends the suffix *SLOpt* to the new plan filename. (You can change the filename.) If unchecked, WAE Design changes the current plan file with the updated information.
- Step 7** Click **OK**.
-

