

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.

- Identify Worst-Case Traffic Utilization, on page 2 on interfaces per service class
- (Optional) VPN worst-case utilization and latency
- (Optional) Identify Worst-Case Demand Latency, on page 5
- Visualize Network in Failure Impact View, on page 11, 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:

- Identify Worst-Case Traffic Utilization, on page 2
- Identify Worst-Case Demand Latency, on page 5
- Run Simulation Analysis, on page 7
- Analyze Simulation Analysis Reports, on page 10
- Visualize Network in Failure Impact View, on page 11
- Parallelization, on page 13

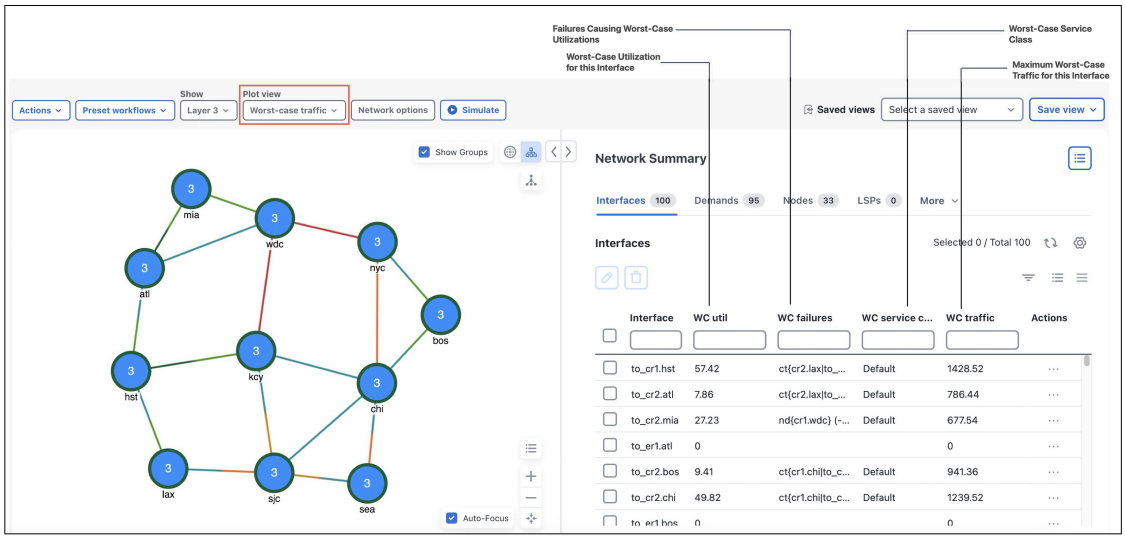
Identify Worst-Case Traffic Utilization

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 7.

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 1: Worst-Case Traffic Utilization for All Interfaces

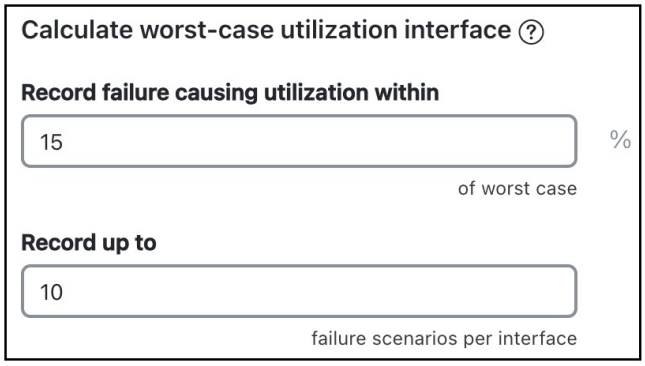


In addition to Plot view changes, the following columns are also updated in the Interfaces and Circuits tables upon finishing the Simulation analysis:

- 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 Identify Worst-Case QoS Violations, on page 3.) 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 Identify Worst-Case QoS Violations, on page 3.

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

Identify 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)
Worst-case QoS calculations for VPNs	Simulate VPN

Fail Circuits to Worst-Case Utilization

After running Simulation analysis (see Run Simulation Analysis, on page 7), 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 2: Failure of Single Circuit to Its Worst-Case, on page 5).

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

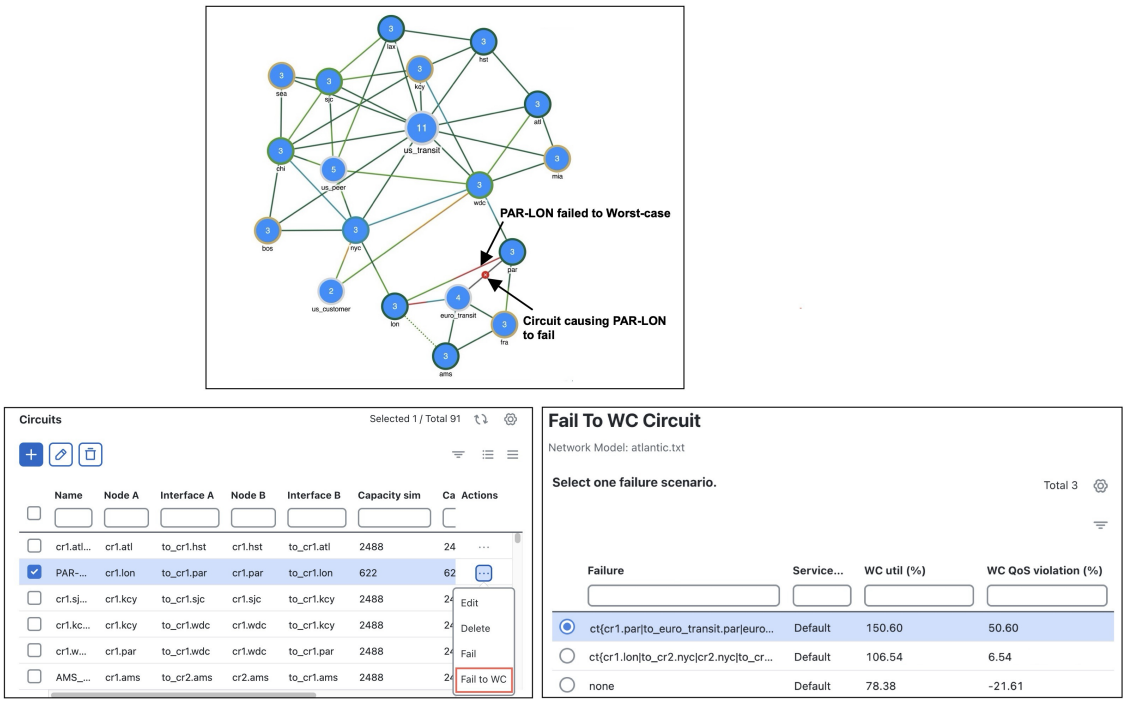
To fail the interface/circuit to its worst-case utilization, do the following:

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 2: Failure of Single Circuit to Its Worst-Case, on page 5).
4. Click **Submit**.

The network plot changes to show this particular failure scenario (see Figure 2: Failure of Single Circuit to Its Worst-Case, on page 5).

Figure 2: Failure of Single Circuit to Its Worst-Case



Identify Worst-Case Demand Latency

When running Simulation analysis (see Run Simulation Analysis, on page 7), 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 7) 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.

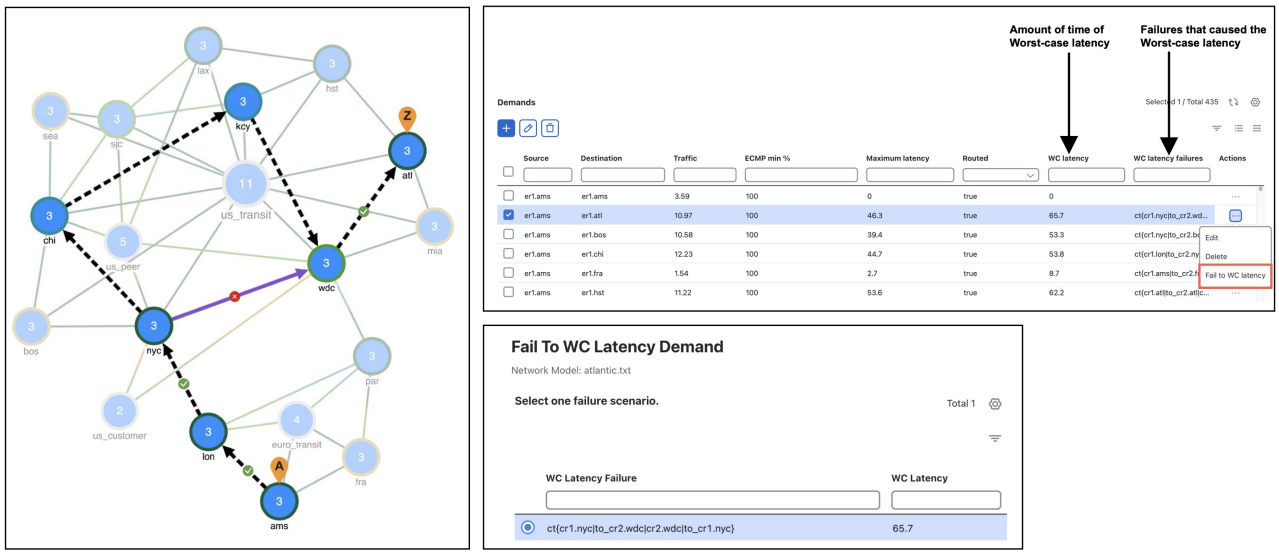
To fail a demand to its worst-case latency, do the following:

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 3: Example of Worst-Case Demand Latency, on page 6).
4. Click **Submit**.

The network plot changes to show this particular failure scenario (see Figure 3: Example of Worst-Case Demand Latency, on page 6).

Figure 3: 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). The plan file opens in the **Network Design** page.

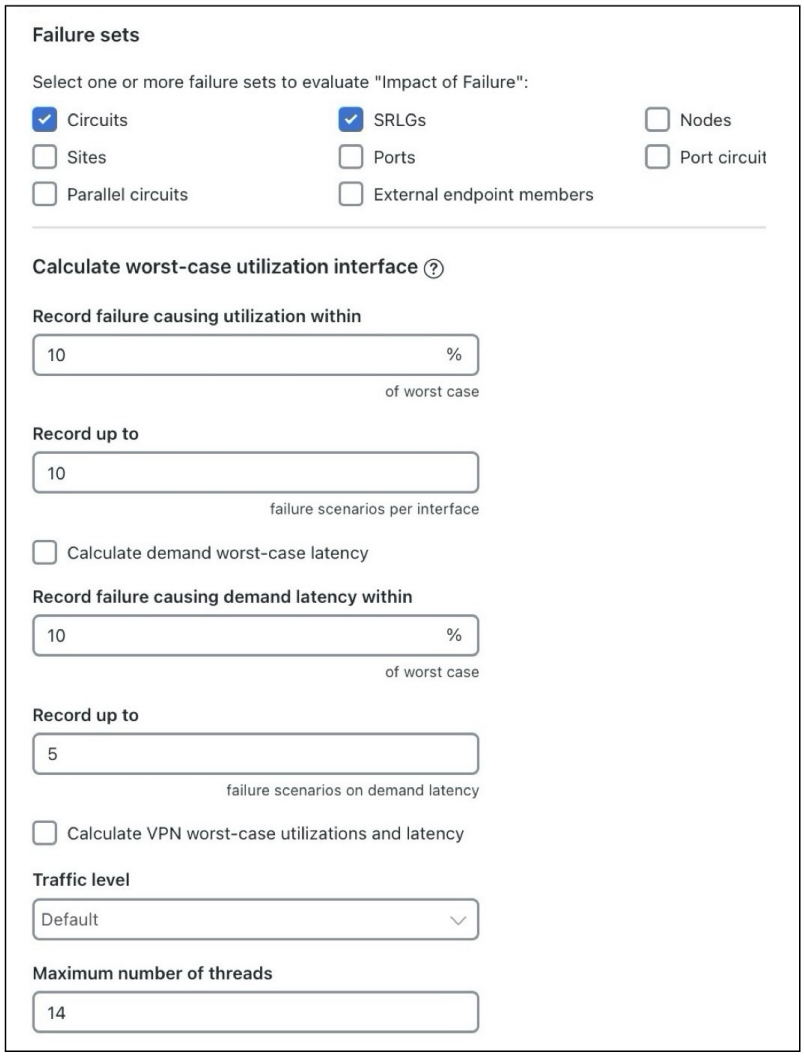
Step 2 From the toolbar, choose any of the following options:

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

  OR

- **Actions > Tools > Simulation analysis**

Figure 4: Configure Simulation Analysis



Step 3 In the **Failure sets** section, choose one or more failure sets.

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.

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. If you choose this option, select the priority of the task and set the time at which you want to run the tool. The tool runs at the scheduled time. 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).

  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 Visualize Network in Failure Impact View, on page 11 view to identify the circuits that are responsible for worst-case traffic congestion.

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). The plan file 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 1: Simulations Table in Simulation Analysis Report, on page 10).

Table 1: Simulations Table in Simulation Analysis Report

Simulation Data Point	Description
Failure	Failure scenario used in the analysis.
Service class	Service class used in the analysis.
Traffic level	Traffic level used in the analysis.

Simulation Data Point	Description
Network breakpoint	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.
Num unrouted demands	Number of demands that cannot be routed under this failure for any of the reasons identified by the network breakpoint.
Unrouted traffic	Total amount of demand traffic that cannot be routed under this failure for any of the reasons identified by the network breakpoint.
Max util	Maximum utilization over all interfaces in this simulation. Utilization is the traffic through the interface as a percentage of the capacity of the interface.
Max QoS bound percent	Worst-case capacity available without violating QoS bounds, expressed as a percentage of the total capacity.
Num QoS violations	Number of times the QoS bound is violated. QoS bounds are set through service class policies and interface queue parameters.
Latency bound violations	Number of demands with maximum latency in excess of the latency bound specified for the demand.
Num unrouted LSPs	Number of unrouted non-Fast Reroute (FRR) LSPs in the analysis.
Num unrouted FRR LSPs	Number of unrouted FRR LSPs in the analysis.

Visualize Network in Failure Impact View

The **Failure impact** view is available upon running a Simulation analysis (Figure 5: Example Failure Impact, on page 12). 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.

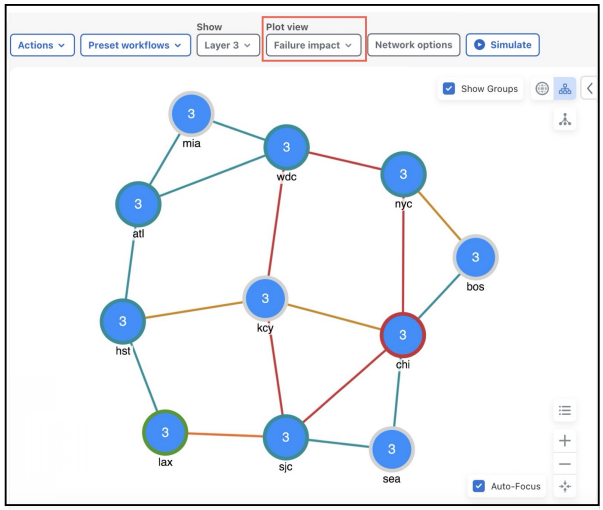
Format = if{Node|Interface}

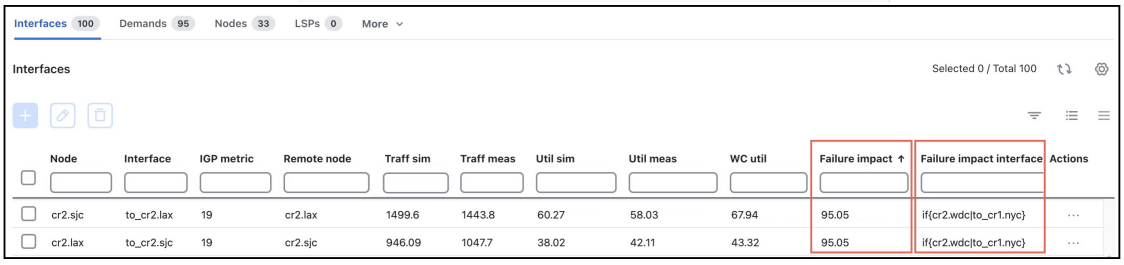
Example: if{cr2.sjc|to_cr1.kcy} means if the circuit goes down, it will have the greatest traffic impact on the cr2.sjc to cr1.kcy interface.

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 5: 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.

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.

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.

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
```