

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:

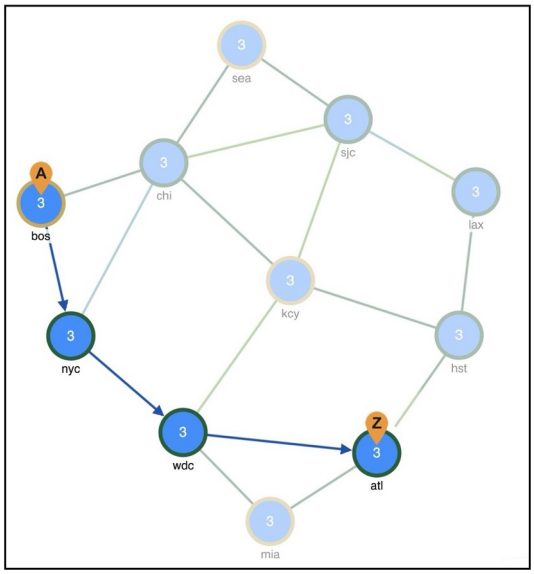
- Demands, on page 1
- How demands can be used?, on page 4
- Create demands and demand meshes, on page 5
- Visualize demands, on page 8
- Demand traffic, on page 13

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 4) 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 1: 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 6. 2. Set the demand traffic manually or using the **Demand Deduction** tool. For details, see Modify demand traffic, on page 14 and Estimate demand traffic using demand deduction, on page 20.
Model future usage in the network	1. Create a demand mesh. For details, see Create demand meshes, on page 6. 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.

Purpose	Suggested Steps to Take
Design networks	1. Create a demand mesh. For details, see Create demand meshes, on page 6. 2. Set the demand traffic using methods described in Simulate traffic using demands, on page 13.
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). It opens in the **Network Design** page.

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

OR

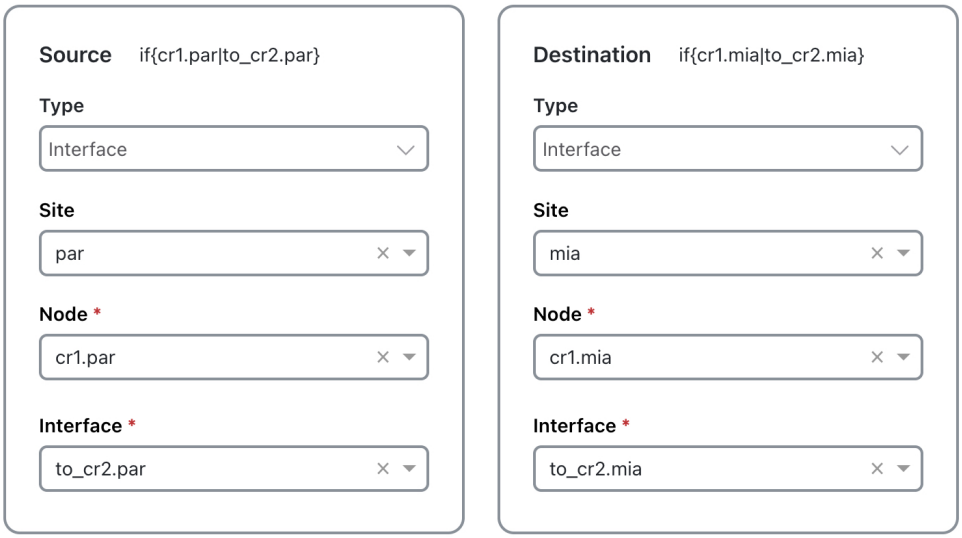
In the Network Summary panel on the right side, click  > **Demands** in the **Demands** tab.

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 2: Source and destination options, on page 6).

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

Figure 2: 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.

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.

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). 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.
- 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.
- 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). 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). 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.

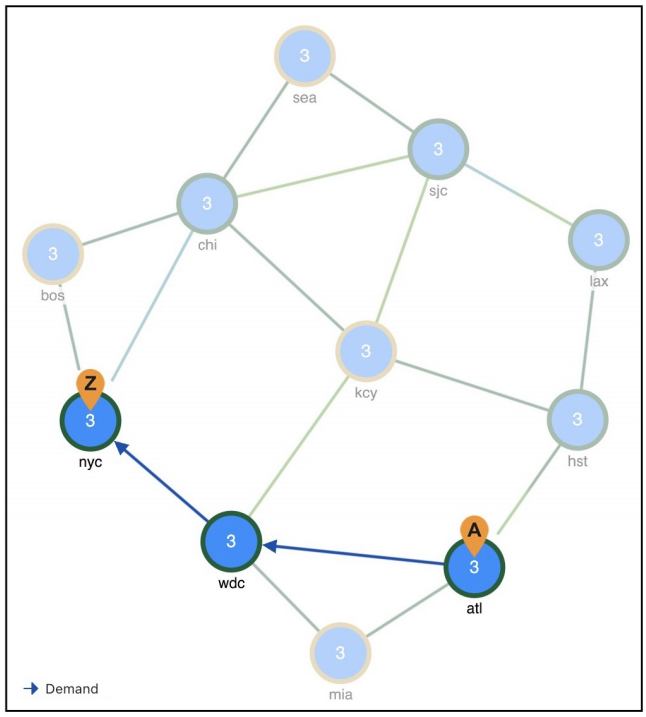
To delete a latency bound, delete the text in this field.

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 3: Demand route



View participating only nodes

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

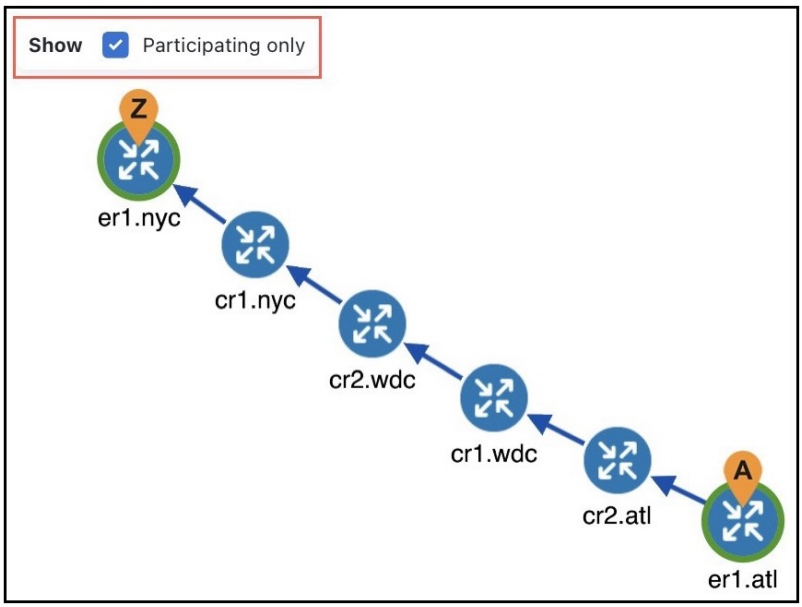
Procedure

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 4: Participating only nodes



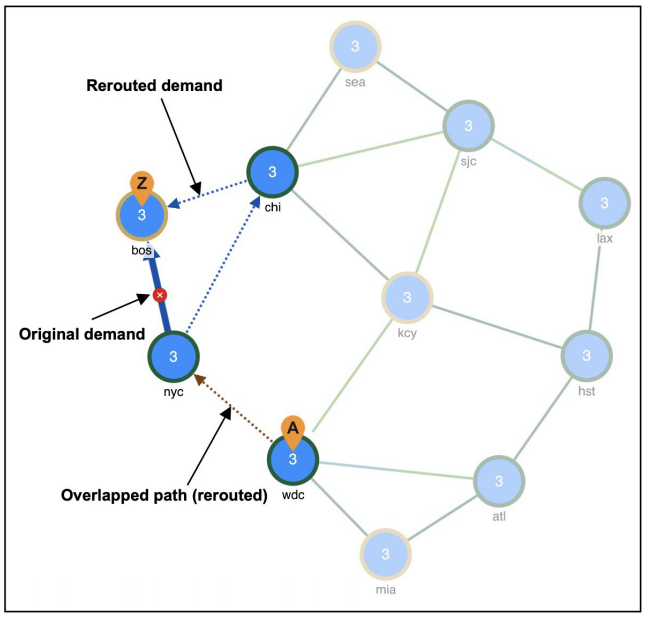
What to do next

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 5: 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.

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 6: Shortest IGP and latency paths, and overlapped path, on page 12.

View shortest latency paths

To view the shortest latency path of a demand:

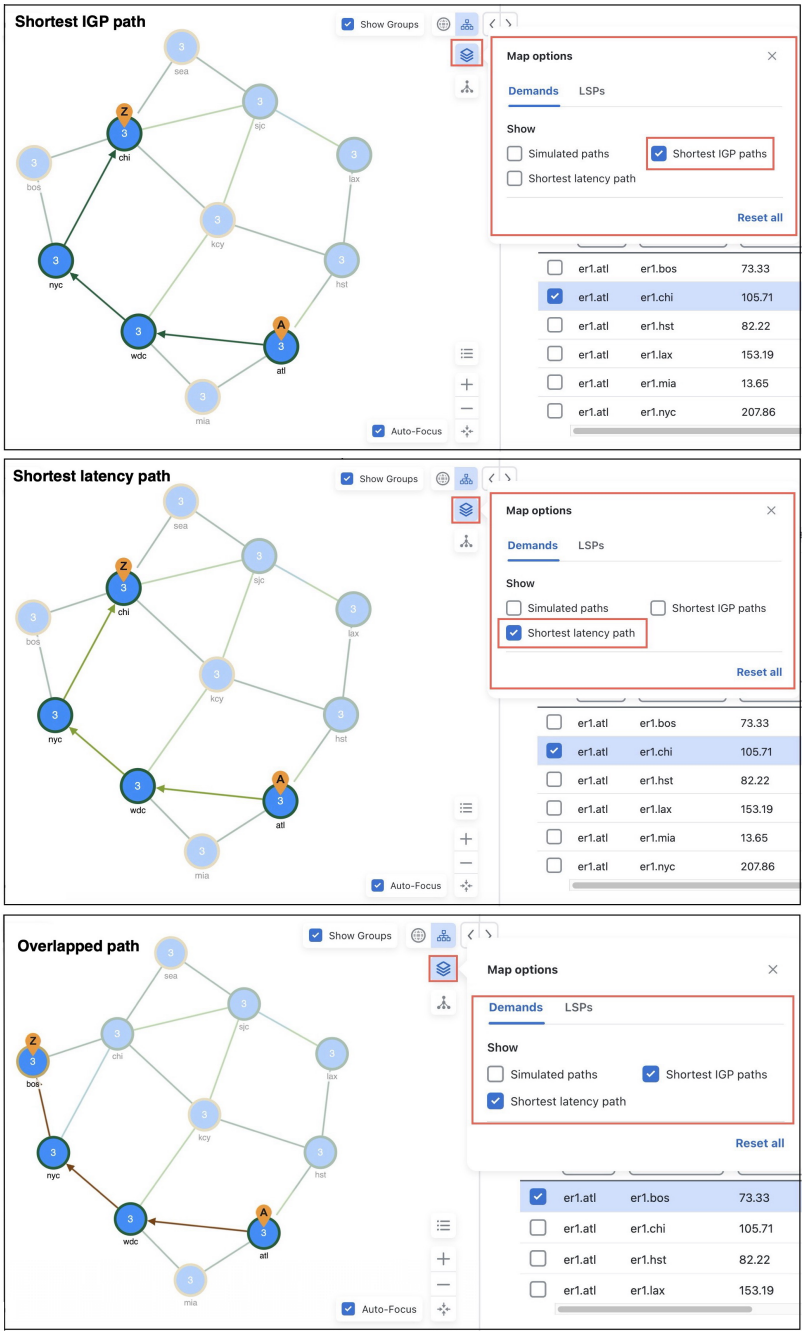
1. Click the **Map options** icon in the network plot ( ) and then check the **Shortest latency path** check box under the **Demands** tab.
2. Select the demand from the Demands table.

The shortest latency path is displayed as a light green arrow in the network plot. See Figure 6: Shortest IGP and latency paths, and overlapped path, on page 12.

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 6: 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

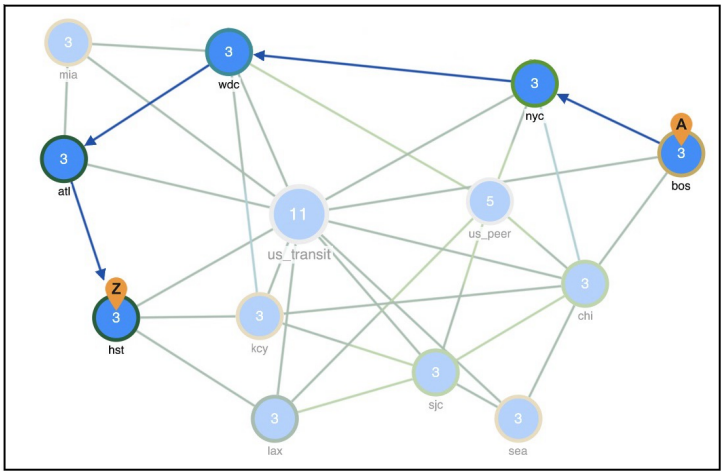
Step 1 Open the plan file (see Open plan files). It opens in the **Network Design** page.

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 7: 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.

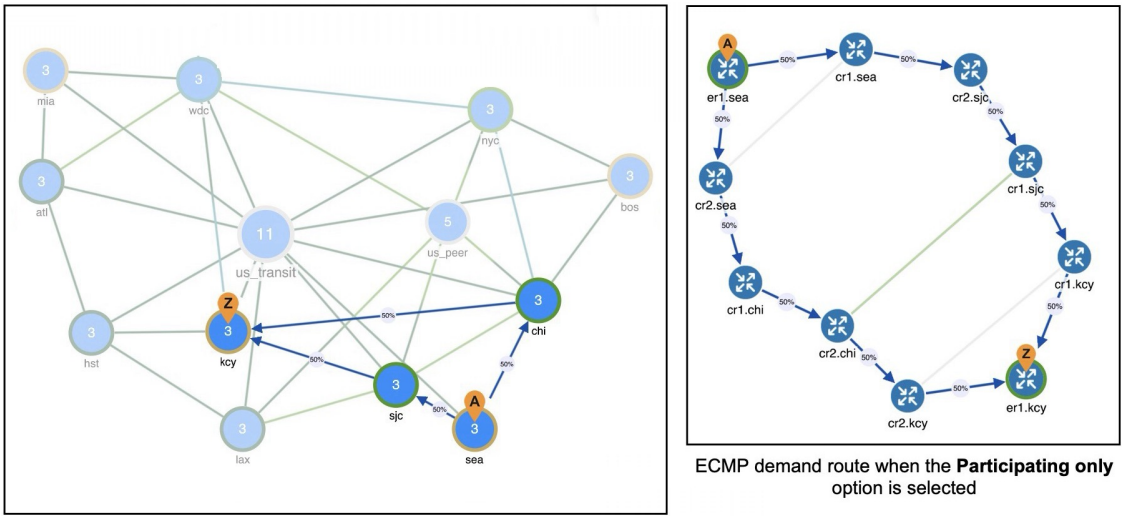
d) Notice that the sum of Delay sim values of all interfaces is equal to the maximum latency of the corresponding Demand. In this case, 23.

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 8: 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 15).
- To set a fixed or relative value, use the **Modify demand traffic** initializer (Modify fixed or relative demand traffic, on page 15).

Modify fixed demand traffic

To modify fixed demand traffic, do the following:

Procedure

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

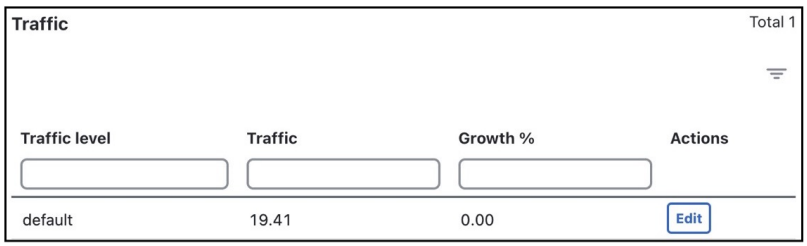
Step 2 In the Network Summary panel on the right side, select one or more demands from the **Demands** table.

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 9: 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). 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 1: Modify demand traffic options, on page 16 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 1: Modify demand traffic options

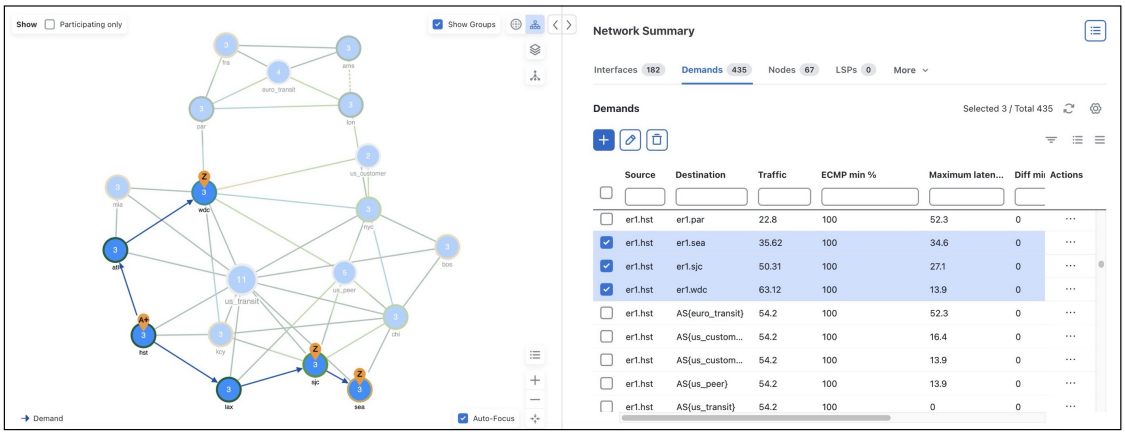
Option	Description
Change traffic by __ %	Change traffic by a specified percentage. Positive percentages add to the traffic, and negative percentages subtract. For example, if the traffic were 1000 Mbps and you entered –10, the traffic would be reduced to 900 Mbps.
Add	• Add __ Mbps in total, proportionally—Add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

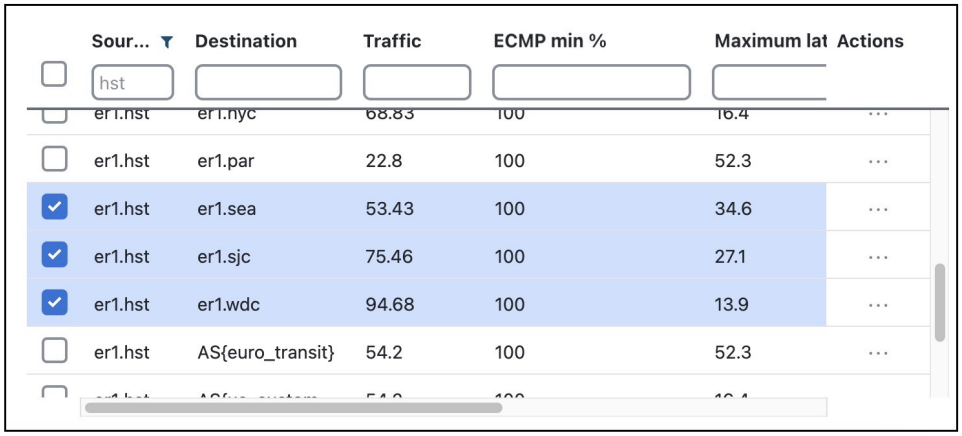
Note the values in the **Traffic** column for the demands. In this example, 35.62, 50.31, and 63.12 Mbps.

Figure 10: 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 11: 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 20.

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.

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 2: Flows table columns, on page 20 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.

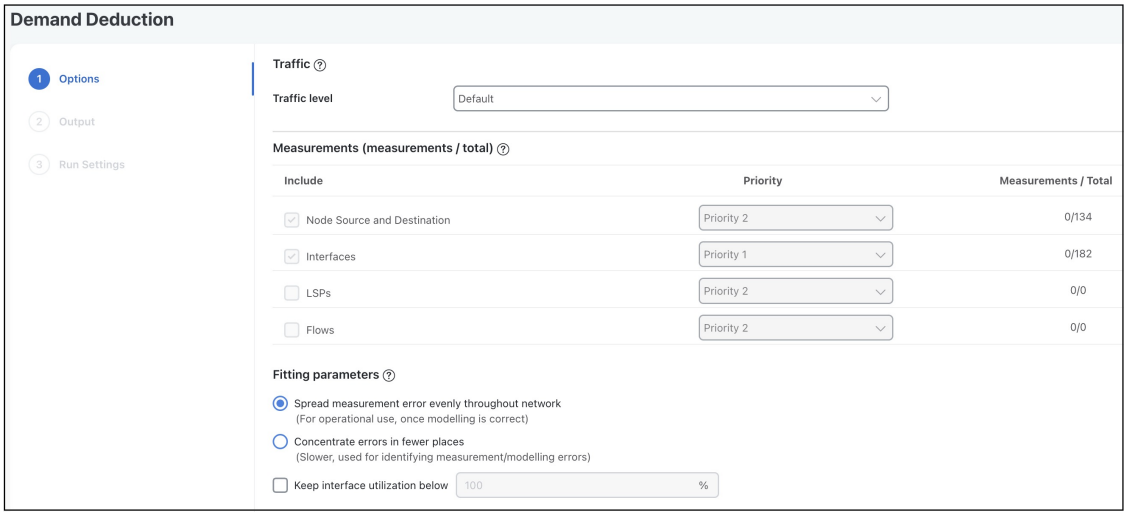
Table 2: Flows table columns

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 19. For information on setting up external endpoint members to be included in Demand deduction calculations, see Simulate Advanced Routing with External Endpoints.



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

Procedure

Step 1 Open the plan file (see Open plan files). 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.

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).

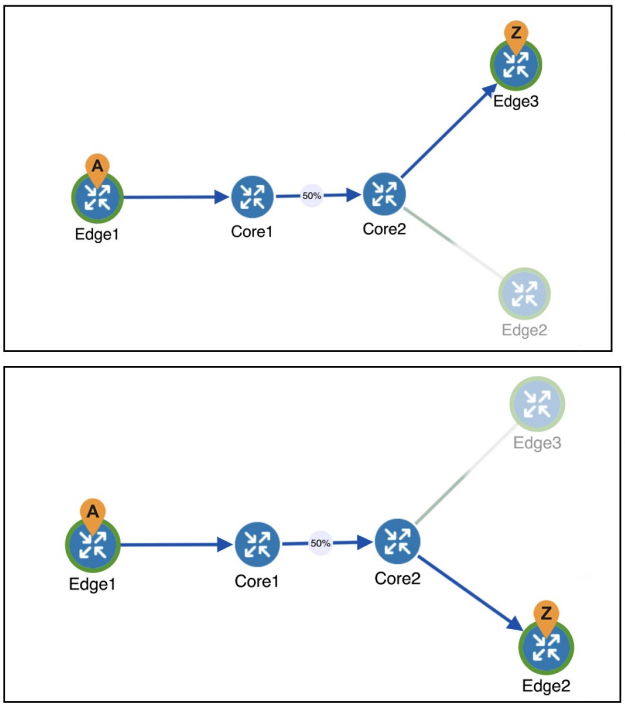
Note
Ensure that you save the plan file before you schedule the job. Any unsaved changes in the plan file are not considered when you run the tool as a scheduled job.

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 12: Network containing two demands and no demand traffic, on page 22 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 12: Network containing two demands and no demand traffic



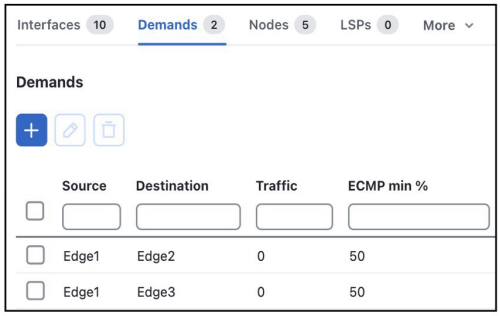


Figure 13: Measured traffic view and interfaces associated with the demands, on page 23 shows the Measured traffic view and the five interfaces associated with the two demands, three of which have measured traffic.

- 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 13: Measured traffic view and interfaces associated with the demands



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 14: Simulated view showing demand traffic, on page 24). 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.

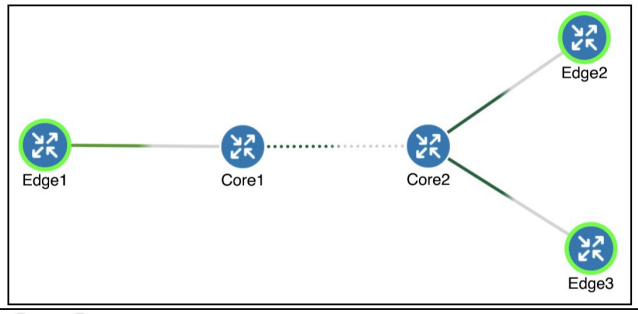
- Edge1 to Core1 has 460 Mbps of simulated traffic.
- Both Core1 to Core2 interfaces have 230 Mbps.
- Core2 to Edge2 and Core2 to Edge3 both have 230 Mbps.

  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 14: Simulated view showing demand traffic



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

Interfaces

Selected 0 / Total 10

Node	Interface	IGP metric	Remote node	Traff sim	Traff meas ↑	Capacity sim	Util sim	Util meas	Abs meas diff	Meas diff/cap (%)	Actions
Core1	to_Core2	1	Core2	230	210	1000	23	21	20	2	···
Core1	to_Core2[1]	1	Core2	230	240	1000	23	24	10	1	···
Edge1	to_Core1	1	Core1	460	470	1000	46	47	10	1	···
Core2	to_Core1[1]	1	Core1	0	NA	1000	0	NA	NA	NA	···
Core2	to_Edge2	1	Edge2	230	NA	1000	23	NA	NA	NA	···
Edge2	to_Core2	1	Core2	0	NA	1000	0	NA	NA	NA	···
Core2	to_Edge3	1	Edge3	230	NA	1000	23	NA	NA	NA	···
Edge3	to_Core2	1	Core2	0	NA	1000	0	NA	NA	NA	···
Core1	to_Edge1	1	Edge1	0	NA	1000	0	NA	NA	NA	···

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.