

Simulate VPN

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

- Which VPNs are on my network? Where and how are they configured?
- Which VPNs are using congested interfaces?
- Which VPNs will experience congestion under any of a given list of failure scenarios?
- Which failures scenarios cause the worst-case congestion or latency for a VPN?

There are many varieties of VPNs. For example, there are Layer 2 (L2) VPNs and Layer 3 (L3) VPNs, each with different categories within it, and there are vendor-specific VPN implementations. Each VPN type has its own specific configuration and terminology. The Cisco Crosswork Planning VPN model supports a number of these VPN types based on either route-target or full-mesh connectivity.

This section contains the following topics:

- VPN model, on page 2
- VPNs, on page 3
- VPN nodes, on page 5
- Layer 3 VPN example, on page 8
- VPN simulation analysis, on page 11

VPN model

VPN objects

Object	Description	Examples
VPNs	A set of VPN nodes that can exchange data with each other.	• Layer 2 VPN: The VPN represents an individual VPLS containing Virtual Switch Interfaces (VSIs). • Layer 3 VPN: The VPN represents sets of VRFs associated with a set of VPN nodes that forward traffic between themselves. Often, this set of VRFs signifies a single customer or service.
VPN nodes	Connection points in a VPN. They exist on standard nodes, and each node can contain multiple VPN nodes. A VPN node can be in only one VPN.	• Layer 2 VPN: The VPN node represents the VSIs configured on each router. • Layer 3 VPN: The VPN node represents the VRF instances configured on each router.

VPN topology and connectivity

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



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

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

- The two VPN nodes are in the same VPN.
- The demand is in the same service class as the VPN.
- Only for VPNs with RT connectivity, the **RT export** property of one VPN node must match the **RT import** property of another VPN node.

Once demands are associated with the VPN, this configuration simulates the associated access circuits exchanging traffic as if they were on the same LAN.

Note that a demand associated with a VPN can additionally contain other traffic that is for that VPN.

Connectivity	Description
Full Mesh	Full-mesh connectivity is a complete mesh of connections between VPN nodes in a VPN so they can all communicate with each other. This connectivity is typical in a VPLS, where all VSIs identify one another based on a common AGI.
Route Targets (RT)	Route targets model the more complex connectivity used in Layer 3 VPNs, such as hub-and-spoke networks. Here, the VRFs exchange data with one another based on the matching of RT export and RT import properties set for each VPN node. Having an import/export pair does not create bidirectional communication. Rather, traffic flows in the opposite direction of the routed advertisements. For example, if node A's RT import matches node B's RT export, traffic can flow from node A to B. For traffic to flow from node B back to node A, node B must have an RT import that matches an RT export of node A. This combination of matching imported and exported RTs defines which VPN nodes can exchange data. The VPN name identifies the VPN itself.

VPNs

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

- **Name**—Unique name of the VPN.
- **Type**—Type of VPN. Choose from the options: VPWS, VPLS, or L3VPN.
- **Connectivity**—Determines how Cisco Crosswork Planning calculates connectivity and associated demands for VPNs:
  - Full Mesh—Connectivity is between all nodes in the VPN. Cisco Crosswork Planning ignores the RT Import and RT Export properties of the VPN nodes.
  - RT—Connectivity is based on the RT Import and RT Export properties of its VPN nodes.
- **Service class**—Service class associated with this VPN.

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

Create VPNs

You can create new VPNs and then later add VPN nodes to them (see Add VPN nodes to VPNs, on page 7).

Create new VPNs

To create new VPNs, 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** > **Insert** > **VPNs** > **VPN**.

OR

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

The VPNs tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon () and check the **VPNs** check box.

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

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

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

Step 6 Choose the Service class for the VPN.

Step 7 Click **Add**.

Step 8 (Optional) Add VPN nodes to the newly created VPN. For details, see Add VPN nodes to VPNs, on page 7.

VPNs table

The VPNs table lists the VPN properties, its associated service class, traffic, and the number of VPN nodes within that VPN (Table 1: VPNs table columns for normal operation, on page 4). For information on QoS measurements, see Simulate Quality of Service (QoS). For information on the Worst-Case columns not listed here, see Table 3: Simulation analysis columns in the VPNs table, on page 12.



Note Because the traffic and QoS calculations are based on all interfaces within the VPN for the service class specified for that VPN, the plot view might differ from the table. For example, the plot view could show Internet traffic while a VPN carrying voice traffic is selected.



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

Table 1: VPNs table columns for normal operation

Column	Description
Service class	Service class associated with this VPN. All values within the table are associated with this service class.
Num nodes	Number of VPN nodes in this VPN.

Column	Description
Util meas	The maximum measured utilization of all interfaces used by this VPN.
Util sim	The maximum simulated utilization of all interfaces used by this VPN.
Total src traff meas	Total amount of measured source traffic on this VPN.
Total dest traff meas	Total amount of measured destination traffic on this VPN.
QoS violation sim	Maximum QoS violation under normal operations for all simulated traffic for all interfaces used by this VPN. If the number is positive, there is a violation.
QoS violation sim (%)	QoS violation as a percent of the total simulated interface capacity.
QoS violation meas	Maximum QoS violation under normal operations for all measured traffic for all interfaces used by this VPN. If the number is positive, there is a violation.
QoS violation meas (%)	QoS violation as a percent of the total measured interface capacity.
Latency	Maximum latency of all demands used by this VPN.
Tags	User-defined identifiers that makes it easy to group VPNs.

VPNs are not selectable from the network plot; you can only select and filter to VPNs through tables. When selected, all VPN nodes within the VPN are highlighted in the plot (Figure 1: VPN nodes within a VPN, on page 8).

Identify interfaces used by VPNs

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

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

VPN nodes

VPN nodes are defined by the following properties that determine which VPNs the nodes belong to and how the demands are routed.

- **Site**—Name of the site on which the VPN node resides.
- **Node**—Name of the node on which the VPN node resides. This node name corresponds with one in the Nodes table.
- **Type**—The type of VPN. You can choose from the defaults (VPWS, VPLS, or L3VPN), or you can enter a string value to create a new one. Once entered, the new VPN type appears in the drop-down list and is available for other VPN nodes and VPNs.

- **Name**—Name of the VPN node.
- **VPN**—Name of the VPN in which this VPN node resides. The drop-down list shows existing VPNs of the same type set in the Type field. You can create a VPN node without setting its VPN, but without it, the VPN node is not included in simulations as a member of any VPN.

  To simulate RT connectivity, you must set the VPN Connectivity property to RT and then set the RT import and RT export properties on the individual VPN nodes within it.
- **Description**—Description for the VPN node.
- **RT import** and **RT export**—The pairing of RT values identifies which VPN nodes connect with each other. For more information, see VPN topology and connectivity, on page 2.
- (Optional) **RD**—Route Distinguisher (RD) uniquely identifies routes within a VRF as belonging to one VPN or another, thus enabling duplicate routes to be unique within a global routing table.

Create VPN nodes

Procedure

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

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

OR

In the Network Summary panel on the right side, click  in the **VPN nodes** tab.

The VPN nodes tab is available under the **More** tab. If it is not visible, then click the **Show/hide tables** icon () and check the **VPN nodes** check box.

Step 3 Click .

Step 4 In the **Site** and **Node** fields, choose the site in which the VPN node will exist, and choose the node on which the VPN node is being configured.

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

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

Step 7 From the **VPN** drop-down list, choose the VPN to which you are adding this VPN node. If you do not see the VPN that you expect to see, check if you have selected the correct VPN type in the **Type** drop-down list.

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

Step 9 If the Connectivity for the VPN is RT, enter the applicable route targets in the **RT import** and **RT export** fields. All VPN nodes with the same import RT as another VPN node's export RT can receive traffic from that VPN node. Those VPN nodes with the same export RT as another VPN node's import RT can send traffic to that VPN node.

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

Step 11 Click **Add**.

Add VPN nodes to VPNs

To add VPN nodes to VPNs, do the following:

Procedure

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

Step 2 In the Network Summary panel on the right side, select one or more VPN nodes in the **VPN Nodes** table and click .

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

Step 3 In the **VPN** drop-down list, choose the VPN to which you are adding the VPN nodes. If you do not see the VPN that you expect to see, check if you have selected the correct VPN type in the **Type** drop-down list.

Step 4 Click **Save**.

VPN nodes table

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

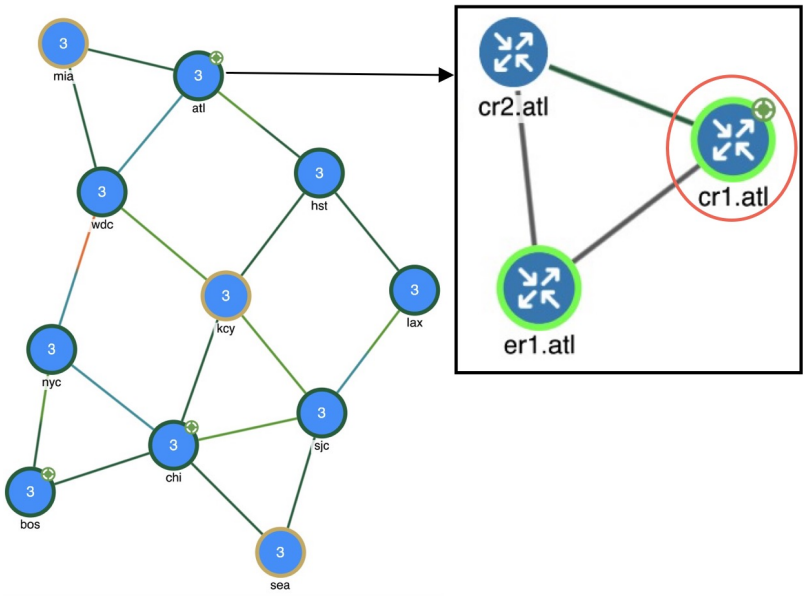
Table 2: VPN nodes table

Column	Description
Total connect	Number of VPN nodes that are connected to this VPN node as defined by the RT Import and RT Export pairings. These may or may not be in the same VPN.
VPN connect	Number of VPN nodes that are connected to this VPN node and are in the same VPN as defined by the VPN column.
Num VPN nodes	Number of nodes in the VPN that this VPN node belongs to as defined by the VPN column. This value is "na" if the VPN node does not belong to a VPN.
Src traff meas	Total amount of measured traffic entering the VPN at this node (source traffic).
Dest traff meas	Total amount of measured traffic leaving the VPN at this node (destination traffic).
Tags	User-defined identifier that makes it easy to group VPN nodes into a single VPN. If you give a VPN node a tag, when you create a VPN later, you can identify its VPN nodes using tags.

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

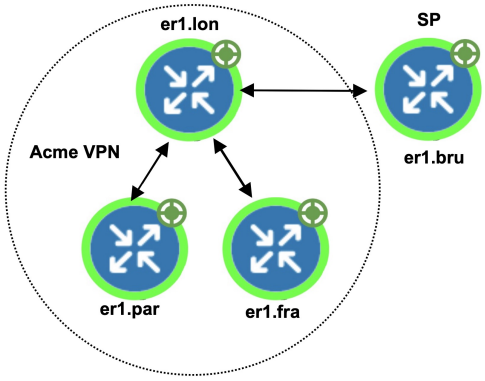
Once selected from the VPN Nodes or VPNs tables, the associated site and the nodes within that site appear with a green circle on it (Figure 1: VPN nodes within a VPN, on page 8).

Figure 1: VPN nodes within a VPN



Layer 3 VPN example

This example illustrates a scenario where the Acme manufacturing company has three offices, but permits the two branch (er1.par and er1.fra) offices to exchange data only with headquarters (er1.lon).

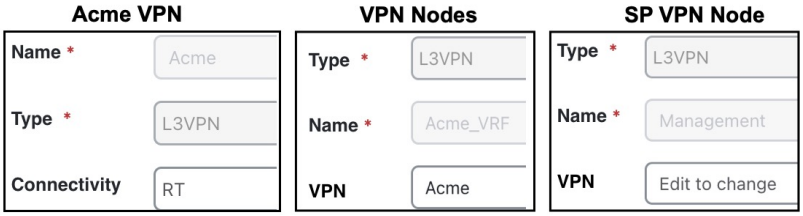


Additionally, headquarters communicates with an SP VPN node (er1.bru) that is not in the Acme VPN. Figure 2: Example RT connectivity and Acme VPN footprint, on page 9 shows the footprint of the Acme VPN and the RTs set for all VPN nodes in this example.

- The VPN is named Acme, and it is set to a Connectivity of RT and a Type of L3VPN.
- In turn, each branch office is set to the Acme VPN, with a Type of L3VPN.
- To exchange data with two other VPN nodes in the Acme VPN, headquarters (er1.lon) imports the offices' exported route targets of 2:1 (er1.par) and 3:1 (er1.fra).
- In turn, headquarters (er1.lon) exports a route target of 1:1.

All three of these other VPN nodes import it (both offices and the SP VPN node).

Because the SP VPN node (er1.bru) is not in the Acme VPN, its communication with er1.lon is not within the context of that VPN.



The VPN footprint in Figure 2: Example RT connectivity and Acme VPN footprint, on page 9 shows that if the circuit between er1.fra and er1.bru becomes congested or fails, the VPN is impacted. However, a failure of the circuit between the two branch offices is not impacted. This failure is illustrated in Figure 3: Example failure between branch offices in the Acme VPN, on page 10, which shows that none of the demands associated with the VPN are rerouted.

Figure 2: Example RT connectivity and Acme VPN footprint

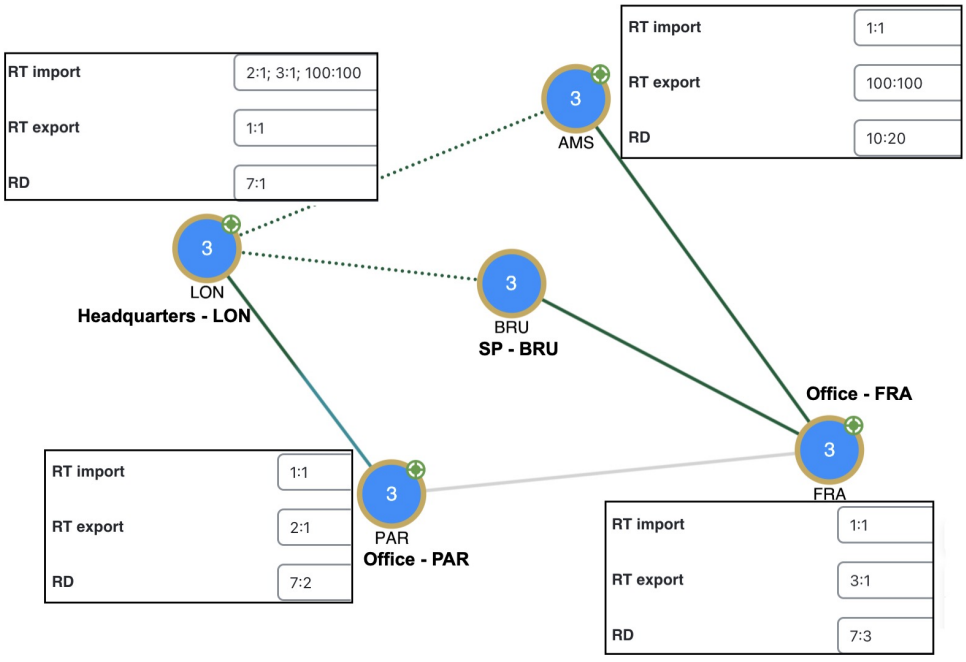
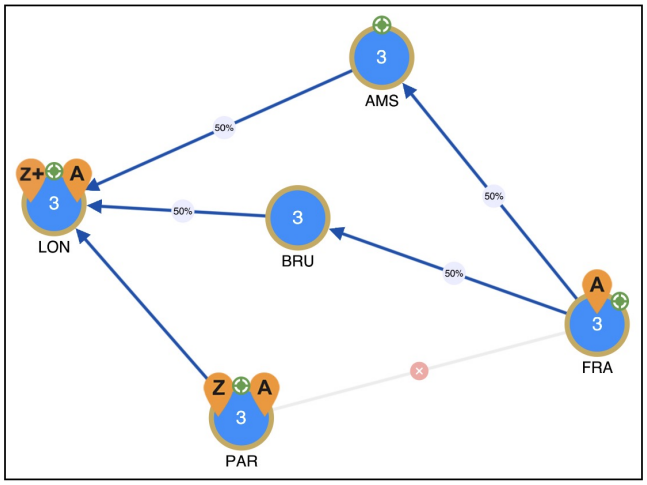


Figure 3: Example failure between branch offices in the Acme VPN



For this example, Figure 4: VPN nodes belonging to Acme VPN, and Acme VPN filtered to demands, on page 11 illustrates the VPN nodes belonging to Acme VPN and the filtering of the Acme VPN to its associated demand traffic. It also shows the calculations of the **Total connect** and **VPN connect** columns in the VPN Nodes table.

- The Total connect for the VPN node residing on er1.lon headquarters is the highest because it exchanges data with three other VPN nodes.

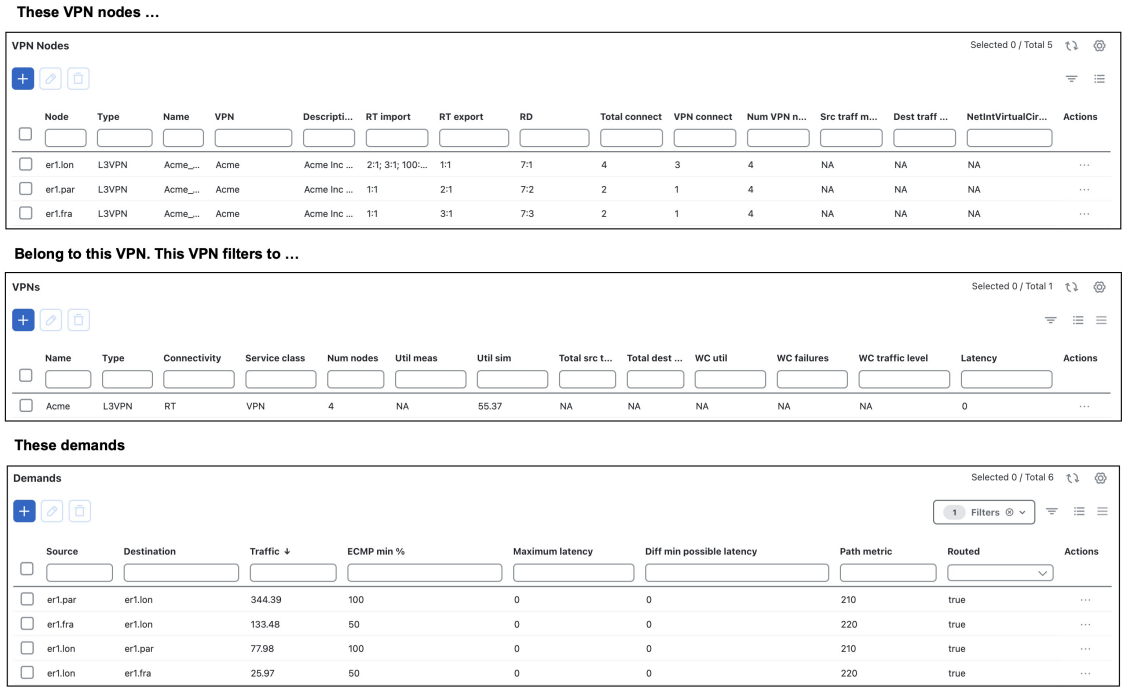
  Each of the offices and the service provider VPN node have 2 in the Total connect column.

- The VPN connect for the VPN node residing on er1.lon headquarters is the highest because it exchanges data with and is in the same VPN as the two offices; all three VPN nodes share the same VPN name.

Each office has 1 in the VPN connect column because it communicates with only one VPN node in the same VPN.

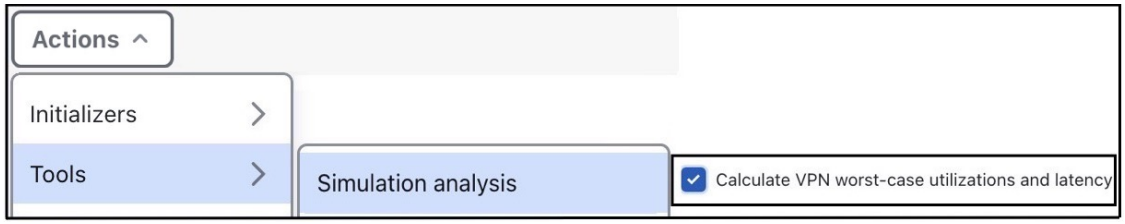
The service provider VPN node (er1.bru) has 0 VPN connects because it does not reside in a defined VPN.

Figure 4: VPN nodes belonging to Acme VPN, and Acme VPN filtered to demands



VPN simulation analysis

When running the **Simulation analysis** tool (from the toolbar, choose **Actions** > **Tools** > **Simulation analysis**), you have the option to record worst-case utilization and latency for VPNs in the VPNs table. You can then select a VPN to fail to its worst-case utilization or worst-case latency using the ••• > **Fail to WC** or **Fail to WC latency** options, respectively.



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

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

Table 3: Simulation analysis columns in the VPNs table

Columns	Description
WC util	Worst-case VPN utilization over all failure scenarios.
WC failures	Failures causing the worst-case utilization of the VPN.
WC traffic level	Traffic level causing the utilization of the interface identified in the WC util column.
WC QoS violation	Highest worst-case QoS violation for all interfaces used by this VPN. A QoS violation is equal to the worst-case traffic minus the worst-case capacity permitted (worst-case QoS bound).
WC QoS violation (%)	Highest worst-case QoS violation for all interfaces in this VPN expressed as a percentage of total capacity.
WC latency	Maximum VPN latency over failure scenarios considered.
WC latency failures	Failures causing the worst-case VPN latency.