

Simulate BGP Routing

This chapter describes how Cisco Crosswork Planning models multi-AS networks and simulates basic BGP routing. Cisco Crosswork Planning does not directly emulate BGP routing configurations, such as local prefs and MEDs. Rather, it provides a high-level modeling of typical peering policies, such as standard customer, transit, and settlement-free arrangements for service providers. This model lets you quickly and easily evaluate the effects of peering locations and basic policy variations.

Additionally, you can extend these high-level models to significantly more complex policy-based routing situations using *external endpoints* as demand sources and destinations. For information on demands and on external endpoints, see Simulate Traffic Flow from Source to Destination Using Demands and Simulate Advanced Routing with External Endpoints.

This section contains the following topics:

- Internal and external AS types, on page 1
- Configure ASes, on page 2
- Route demands between ASes, on page 4
- Know about BGP routing details, on page 8
- BGP routing, on page 9

Internal and external AS types

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

- A single internal AS representing the full topology of your network.
- Individual peering nodes of neighboring external ASes.
- Peering circuits connecting the internal AS to the nodes in the external ASes.

Generally, there are many external ASes in the network model, but usually only one or a few internal ASes. All nodes in an external AS are typically placed in the same site, although you can place them in any site.

ASNs and their types are defined in the AS Properties window and listed in the AS table. Nodes are assigned to ASes in the Node Properties window.

Configure ASes

Create ASes

Follow these steps to create an empty AS. After creating the AS, you still need to associate nodes with it and create the relationship between this AS and others. See Associate nodes with an AS, on page 2 and Edit AS routing policies, on page 3.

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** > **AS**.

OR

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

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

Step 3 Configure the AS properties:

- **ASN**—AS number, which is a text string that can be a number or name.
- **Name**—AS name.
- **Type**—Internal ASes have a full topology. External ASes have a collapsed topology with just border nodes and a virtual node.
- **External mesh**—When creating a demand mesh, this option tells Cisco Crosswork Planning whether to create external meshes. When one or both ASes are set to Include, Cisco Crosswork Planning creates a mesh between the external ASes (default). If both are set to Exclude, no demands are created.
- **IGP protocol**—Choose OSPF, ISIS, or EIGRP from the drop-down.
- **Description**—A text description of the AS.

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

Step 5 To change the routing policy, select the AS and click  and then choose the **AS Relationships** tab (see Edit AS routing policies, on page 3).

Associate nodes with an AS

To associate nodes with an AS, do the following:

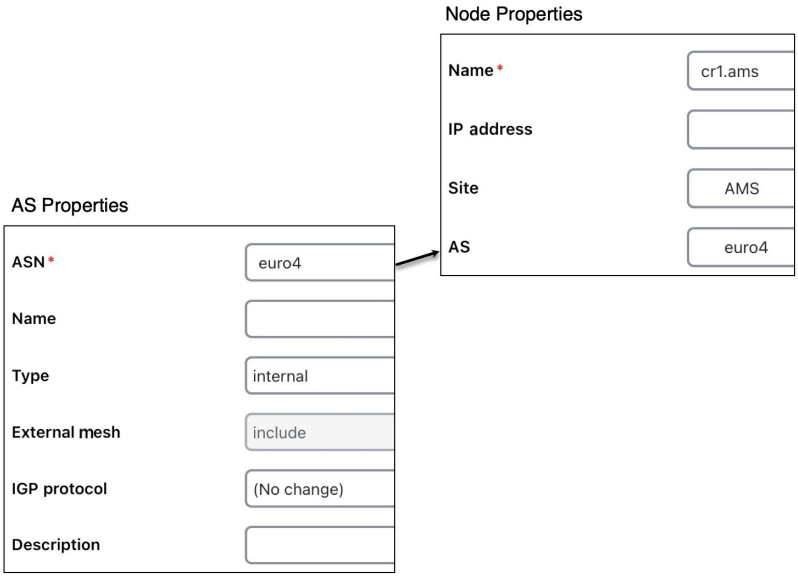
Procedure

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

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

Step 3 Click .

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



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

Step 5 Click **Save**.

Edit AS routing policies

To create AS relationships, set the routing policy.

Procedure

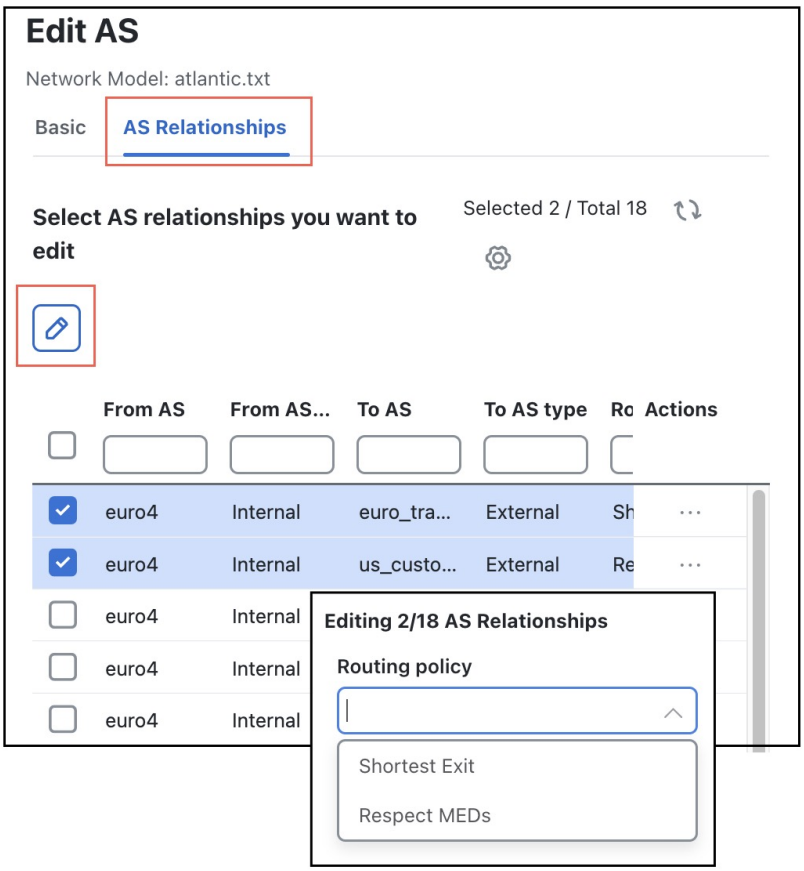
Step 1 Open the plan file (see Open plan files). It opens in the **Network Design** page.

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

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

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

Figure 1: Edit AS relationships



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

Step 6 Click the icon, set the **Routing policy** to **Respect MEDs** or **Shortest Exit**. For details, see Route demands between ASes, on page 4.

Step 7 Click **Submit**.

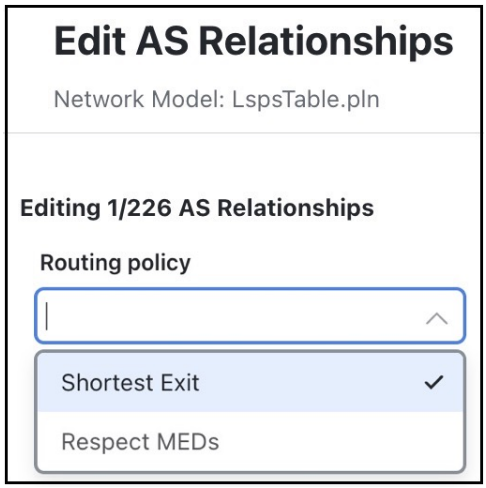
Route demands between ASes

Determine routes between internal ASes

Demands routed within a single AS have a specified source node and destination node where traffic originates and terminates. Demands routed between two connecting internal ASes are specified in the same manner: with a source node in the first AS and a destination node in the second.

Cisco Crosswork Planning routes within an AS, to and from the border exit point, are determined by the IGP protocols. The selection of border exit point is modeled by the **Routing policy**, which is set to either Shortest

Exit or Respect MEDs. This property is set in the Edit AS Relationships window, which is accessed through the Edit AS window.



- **Shortest Exit**—The border exit node is selected, which is closest to the source node, within the IGP of the source AS. If there is a tie, the exit node with the lowest BGP ID is used.
- **Respect MEDs**—The border exit node is selected, which is closest to the destination node, within the IGP of the destination AS. If there is a tie, the exit node with the lowest BGP ID is chosen.

Determine routes between external and internal ASes

Table 1: Typical AS routing configurations, on page 5 lists typical routing configurations that can be constructed by applying different combinations of routing policies for traffic in both directions between two ASes.

- In a peer relationship, routing in both directions is Shortest Exit, which means each controls its own border exit points.
- For a customer relationship, the customer determines the border exit points for traffic in both directions.
- For a transit relationship, the transit AS provides paid transit to the internal AS, so the internal AS determines all border exit points.

Table 1: Typical AS routing configurations

Type	Policy to	Policy from
Peer	Shortest Exit	Shortest Exit
Customer	Respect MEDs	Shortest Exit
Transit	Shortest Exit	Respect MEDs

Like traffic routed within an AS, traffic routed between ASes is represented by demands. However, for demands from and/or to external ASes, the external AS is defined as the source or destination of the demand. Optionally, the specific node in the external AS from which the traffic enters or exits the internal AS is also specified.

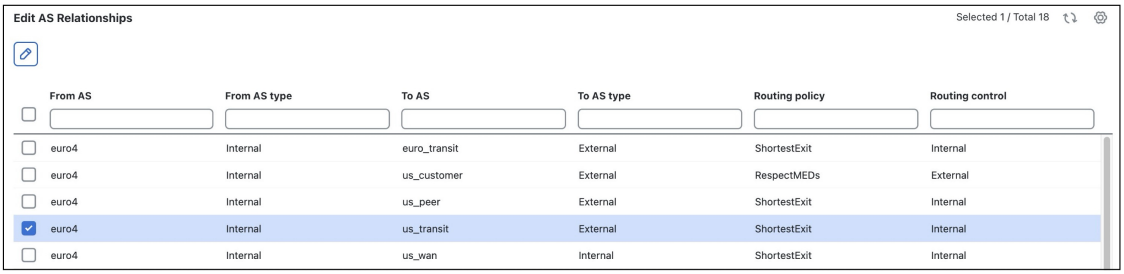
Failover between nodes in the external AS can be modeled. For example, if the traffic is sourced from an external AS and if the peering circuit from which traffic is entering the internal AS under normal operation fails, the traffic can enter the internal AS from a different interface or peering node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

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

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

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

The AS that controls the routing chooses which peering node to use. If the internal AS controls the routing, then because the topology of the internal AS is known, you can simulate the routing to the peering node. However, because Cisco Crosswork Planning has limited knowledge of the external AS topology, if the external AS controls the routing, you cannot predict how traffic will be distributed among the exit points.



The AS that controls the routing is determined by the AS type, direction of the demand, and the Routing policy property as described in Table 2: Determining the AS that controls the routing, on page 6.

Table 2: Determining the AS that controls the routing

Direction	Routing Policy	AS with Routing Control
External AS to Internal AS (Ingress)	Respect MEDs	Internal
	Shortest Exit	External
Internal AS to External AS (Egress)	Respect MEDs	External
	Shortest Exit	Internal

Two ASes can be in one of four different routing relationships to one another, depending on which of the two routing policies is chosen in each direction (Table 3: Effects of routing policy and routing control , on page 7).

- If traffic is routed to an external AS when it has control and there is no knowledge of its topology, a set of demands is created from the source in the internal AS (or from another external AS), each with a destination set to one of the border nodes in the external AS. This way, any division of traffic between the exit points can be modeled.
- If traffic is routed to an external AS when an internal AS has control, a single demand is created from the source to the AS itself. Cisco Crosswork Planning simulations determine the correct exit point for this single demand based on the source.

- If traffic is to be routed from an external AS when it has control, a demand is created from each node in the external AS to each node in the internal AS.
- If traffic is to be routed from an external AS when an internal AS has control, a demand is created to each node in the internal AS using the external AS as the source. The demand originates from one or multiple nodes in the external AS, depending on the topology and the metric cost to reach the destination node. For example, a single demand from an external AS to a specific node could be sourced from two different nodes in the external AS, each carrying 50% of the demand traffic.

Table 3: Effects of routing policy and routing control

Direction	Routing Policy	AS with Routing Control	Demand Source or Destination Endpoint in Remote AS	Number of Demands
External AS to Internal AS (Ingress)	Respect MEDs	Internal	Entire external AS	One only
	Shortest Exit	External	Border nodes	One for each node
Internal AS to External AS (Egress)	Respect MEDs	External	Border nodes	One for each node
	Shortest Exit	Internal	Entire external AS	One only

Configure external meshes

An external mesh consists of two or more external ASes with a **Type** property of **external**. An internal AS typically restricts advertisement of BGP routes for some external ASes to other external ASes. For example, destinations reachable through the transit network would not be advertised to a peer, or vice versa. In Cisco Crosswork Planning, these restrictions are represented by the absence of demands between the two external ASes.

Each AS has a property called **External mesh**, which Cisco Crosswork Planning uses when inserting demand meshes into a plan. Demands are created for external ASes only if one or both ASes have **External mesh** set to **include**. If both ASes are set to **exclude**, no demands are created for the external AS. For example, in Figure 2: External mesh control, on page 8 the peer and transit ASes are both set to Exclude, so no demands are created between those ASes. All other external AS demands are included in the demand mesh. Table 4: External mesh settings for common AS relationships, on page 7 shows the External Mesh settings for common AS relationships.

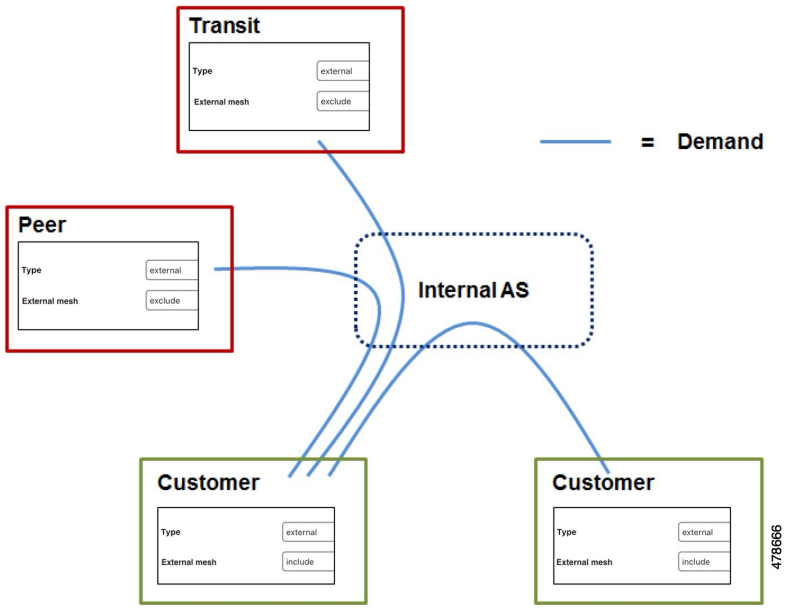
The External mesh property is set in the Edit AS window.

Table 4: External mesh settings for common AS relationships

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

For internal ASes, the **External mesh** property is ignored. More complex route advertisement policies cannot be represented by these simple External mesh settings. In this case, demand mesh creation must be performed in several steps, possibly using a script.

Figure 2: External mesh control



Know about BGP routing details

BGP multihop

Cisco Crosswork Planning automatically constructs BGP pseudonodes where necessary when BGP multihops are detected.

Cisco Crosswork Planning models the nodes in external ASes that are directly connected, for example, through eBGP, to nodes in internal ASes. One exception is that you can model BGP multihops by setting the node **Type** property to **psn** (pseudonode), such as might occur at a peering exchange. This pseudonode can represent the switch that connects a number of external AS nodes to the same internal AS node. In this instance, multiple external AS nodes are connected by circuits to a BGP psn node, and this node is connected to a node in the internal AS.

Note In all cases, eBGP multipaths across parallel border circuits is assumed.

BGP load balancing

BGP load balancing to an external AS uses eBGP multipaths or eBGP multihops. Cisco Crosswork Planning models these two eBGP load balancing designs in the same manner, though in the UI they are identified only as multipaths. BGP multipath options are disabled by default.

Follow these steps to set BGP multipath options globally.

Procedure

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

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

The Network options page opens.

Step 3 Click the **Protocols** tab.

Step 4 In the **BGP** section, for each BGP multipath option that you want enabled, choose **Enabled** from the drop-down list. By default, all these options are disabled.

- **EBGP multipath**: Turns on eBGP multipath within the internal ASes. Demand routings through the internal AS to an external AS are divided among external routes with equal-cost BGP exit routes.
- **EBGP multipath incoming**: Turns on eBGP multipath in all external ASes. Demand routings from external ASes to an internal AS are divided among external routes with equal-cost BGP exit routes.
- **IBGP multipath**: Turns on iBGP multipath within the internal ASes. Demand routes through an internal AS to an external AS are divided among internal paths to equal-cost BGP exit routes.

Step 5 Click **Save**.

BGP next hop

In networks, there are two common configurations for the BGP next-hop IGP metric used in the path selection. One is to set the next-hop self on the iBGP peers (next-hop self = on). The other is to configure IGP metrics on eBGP interfaces, and to inject the interface prefix into the IGP database by setting the interface to be a passive IGP interface (next-hop self = off).

Cisco Crosswork Planning does not have an explicit next-hop self setting, so it simulates paths as if next-hop self is off. That is, the IGP metric of the egress peering interface is included in the IGP distance to the peering router and is used in the iBGP path selection. However, next-hop self to an external AS can effectively be simulated by setting the metrics on all egress interfaces to that external AS to 0. You can set the IGP metric in either the Edit Interface or Edit Circuit window.

BGP routing

As with all Cisco Crosswork Planning simulations, AS routing uses demands. An IP simulation for a particular failure scenario and traffic level performs these steps.

Procedure

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

Step 2 Interface utilizations are calculated from the demand traffic using the specified traffic level.

Cisco Crosswork Planning allows routes to be calculated between selected nodes even if no demands are present. In this case, only the first step applies.

BGP demands do not failover between external ASes. That is, all traffic to or from an external AS behaves the same under peering failures to an external AS. You can change this default behavior using external endpoints to simulate specific external AS nodes where traffic goes in and out of the network, as well as set priorities so that if one traffic source or destination goes down, the traffic can still be sourced from or delivered to another external AS node.