



Advanced Routing with External Endpoints

To model basic IGP routing, demands are sourced or destined for nodes within the topology. To model basic inter-AS routing, the sources and destinations are neighboring external ASes, or a combination of the external AS and the peering node in that AS. However, more complex routing situations require the use of *external endpoints* as the source or destination. External endpoints can contain multiple member nodes and ASes, and you can specify when traffic enters or exits from each of them individually. This allows you to simulate routing within and between ASes where multiple traffic entry and exit points are used simultaneously. You can also prioritize where the traffic fails over to other nodes and ASes.

There are numerous use cases both for IGP and inter-AS routing:

- Simulate content caching failovers for in-network source of demands that are backed up by another in-network source. If connectivity is lost to the first, traffic is sourced from the second.
- Simulate edge routing with a single entry point into the network edge and a specific failover point. Alternatively, you could model multiple entry points, depending on which is closest to the destination.
- Simulate complex BGP routing policies from a transit provider. For example, you can specify a transit entry location and failover location per destination.
- Simulate failover between peering ASes; for example, from one single-connection transit provider to another.

This section contains the following topics:

- [Routing with External Endpoints, on page 1](#)
- [States, on page 5](#)
- [In-Network Routing Examples, on page 6](#)
- [Intra-AS Routing Examples, on page 9](#)
- [Creating External Endpoints and Their Members, on page 12](#)

Routing with External Endpoints

An external endpoint is a WAE Design object that identifies specific entry (source) or exit (destination) points for demands. These are identified in the External Endpoints table by a name and an optional tag.

Each external endpoint consists of one or more members that are defined as nodes, external ASes, or a combination of an external AS and external node. By setting a demand's source or destination to an external endpoint, you can simulate traffic going from multiple sources to a single destination, from a single source

to multiple destinations, or multiple sources going to multiple destinations. Because of this flexibility, they are useful for specifying secondary entry and exit points in the event of failures.

External Endpoint Members

Each member is assigned properties that prioritize traffic entering and exiting that member, in what order to fail over to another member, and how to distribute traffic for members of equal priority. These properties are set when you create the member, and all members are listed in the External Endpoint Members table.

Member Endpoint: Node

Site: chi

Node: cr2.chi

Priority (1,2,3...) : 1

Equal Priority Routing: Fix Traffic

Traffic Balance (%): 50

380816

- Member Endpoint—Defines whether the member is a node, AS, or external node via an AS.
- Priority—The sequential order in which external endpoint members are used in the simulations should failures occur.
- Equal priority routing—If members have the same priority, this property identifies how the traffic is distributed.



Note Members within an external endpoint that have the same priority must either all be Shortest Path or none of them be Shortest Path.

- Shortest Path—Use the member that results in the use of shortest path between source and destination.

If a member is using a shortest path, the Traffic Balance (%) column in the External Endpoint Members table shows “na.”

- Fix Traffic—Set the traffic across members of equal priority as defined in the Traffic Balance (%) field.
- Deduce Traffic—Behaves the same as Fix Traffic in that it sets traffic across members of equal priority as defined in the Traffic Balance (%) field. However, upon running Demand Deduction, the Traffic Balance (%) field is updated based on the measured traffic in the network. Note that Demand Deduction only estimates the traffic balances for external endpoints with a priority that is in use in the current no-failure simulation. Thus, Deduce Traffic is usually set to Priority 1.

Routing Simulations



Note Although this section describes demands as being sourced from an external endpoint, the same methodology holds true if a demand's destination is an external endpoint.

If a demand's source is defined as an external endpoint, the following selection of external endpoint members ensues.

Priority (1,2,3...) : 1

Step 1 Members with the highest priority (lowest number) are used as the demand's source. For example, if the external endpoint has two members with a priority of 1, the demand is sourced from both members provided they are available.

If one or more of the members are not available, the traffic from the unavailable members is evenly redistributed to the other top priority members.

Step 2 If none of the top priority members are available to source the traffic but there are next-priority members available, Step 1 is repeated for the next priority external endpoint members. Only if all members with the same priority fail does the traffic get routed according to the next priority in the sequence.

Note that if a failure occurs that does not affect the external endpoint member's ability to send or receive traffic, then traffic is rerouted as usual without a need to use the additional members.

If an external endpoint member is an external AS, with or without a node specified, then the routing from or to that member is determined by the BGP routing policy determined by the AS relationships. The distribution of traffic between external endpoints with the same priority is the same as for node members.

Member Endpoint: External AS
 AS: euro_transit
 Node: (unspecified)

Traffic Distribution

The traffic distribution through these demands is based on the Equal Priority Routing property, and if applicable, the Traffic Balance (%) property used to define the external endpoint members.

Equal Priority Routing: Shortest Path
 Traffic Balance (%): Shortest Path
 Fix Traffic
 Deduce Traffic

- If there is only one member and it is defined as Shortest Path, the demand takes the shortest path as defined by the IGP metrics.

Of the routable demands, if the Traffic Balance (%) values are all empty, the traffic is routed and equally load balanced across the demands with the shortest IGP paths. Note in the case of multiple internal ASes, the shortest IGP route is the shortest route in the first AS the demand enters.


- If multiple members of the same priority are set to Shortest Path, the demand takes the path with the shortest IGP path. If all interfaces between the source members and the destination have the same shortest IGP paths, then the traffic is load balanced equally across them.
- If one or more members of the same priority have their Equal Priority Routing property set to Fix Traffic or Deduce Traffic, the demand traffic is split according to each member's Traffic Balance (%) value.
 - If the Traffic Balance percentages across sources with the same priority sum to less than 100%, the overall demand traffic is decreased to that percentage.
 - If an external endpoint member of the same priority fails, the traffic on the remaining members increases in proportion, so that the same amount of traffic is still routed.

Example: Node A failed. Nodes B, C, and D each have a priority of 2 and are each a Fix Traffic type. Their traffic balances are 20%, 20%, and 40%, respectively. The demand has 1000 Mbps of traffic.

Member	Priority	Type	Traffic Balance (%)
A	1	ShortestPath	na
C	2	FixTraffic	20.00
B	2	FixTraffic	20.00
D	2	FixTraffic	40.00
E	3	ShortestPath	na

- Because Node A failed, the demand routes 200 Mbps traffic through node B, 200 Mbps through C, and 400 Mbps through D, totaling 800 Mbps ([Figure 1: Example Effect of Failures When Using External Endpoints, on page 5](#)).
- If Node D fails, the demand routes 400 Mbps traffic through Node B and 400 through C. If Node B fails too, the entire 800 Mbps is routed through C ([Figure 1: Example Effect of Failures When Using External Endpoints, on page 5](#)).
- If all three priority 2 members fail, 1000 Mbps is routed through node E, which is the priority 3 member ([Figure 1: Example Effect of Failures When Using External Endpoints, on page 5](#)).

Figure 1: Example Effect of Failures When Using External Endpoints







States

External endpoint members can be failed, set to an inactive state, and included in Simulation Analysis to calculate a worst-case failure analysis. Additionally, they can be included in SRLGs, which in turn can also be failed, set to inactive, and included in Simulation Analysis. [Table 1: L1 Link Failure and Inactive States](#), on page 6 shows the icons used for these states.

If the external endpoint member that you fail or set to inactive is an AS, entry through any of its nodes is not possible. The same is true if an SRLG is failed or set to inactive and if it contains an external endpoint member that is an AS.

Table 1: L1 Link Failure and Inactive States

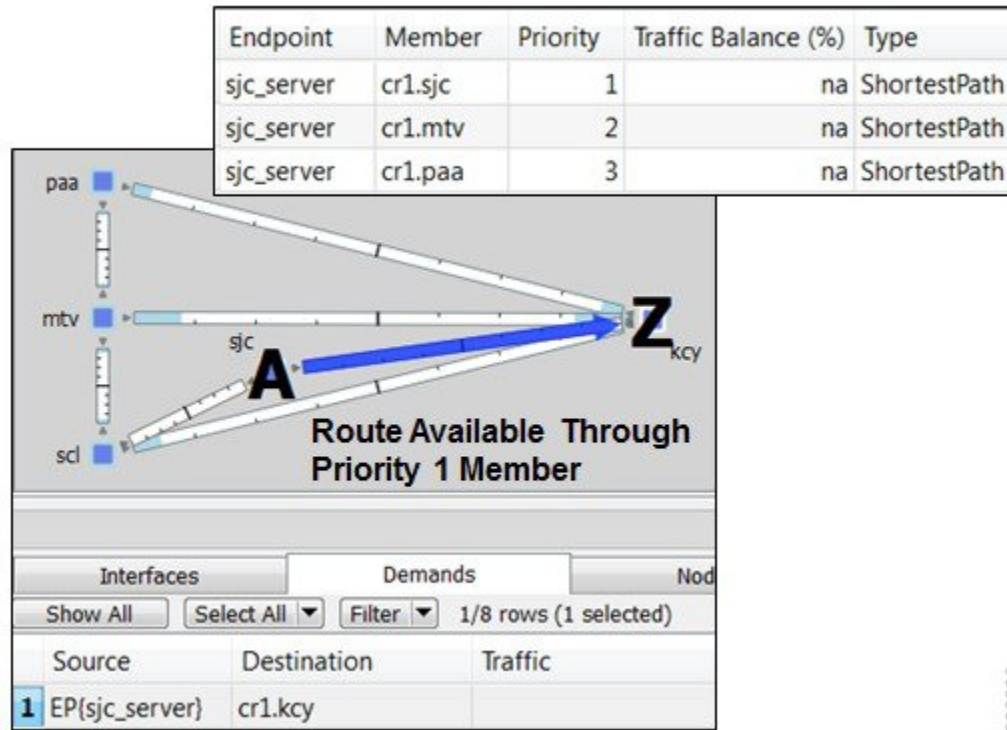
External Endpoint Member	Description
If the external endpoint member is a node, this icon appears under that node in the plot. If the member is an external AS or an external AS via a node, this icon appears under all nodes in the AS.	
	Failed external endpoint member.
	External endpoint member is not operational because it is contained in a failed SRLG.
	External endpoint member is not active.
	External endpoint member is not operational because it is contained in an SRLG that is not active.

In-Network Routing Examples

Content from Multiple Data Centers

This example demonstrates servers that cache content for customers in Kansas City (cr1.kcy). The primary server is located in San Jose (cr1.sjc), with backup servers in Mountain View (cr1.mtv) and Palo Alto (cr1.paa). Each of the servers is a member of an external endpoint (sjc_server) that is used as the source of the demand to cr1.kcy.

Figure 2:




383288

- Member cr1.sjc has a priority of 1, and it takes the shortest path.
- Member cr1.mtv has a priority of 2, and it takes the shortest path.
- Member cr1.paa has a priority of 3, and it takes the shortest path.


Figure 3: Example Failover for Multiple Data Centers, on page 8 shows how the failure scenario works.

- When the priority 1 server node fails (cr1.sjc), content continues to be delivered using cr1.mtv as the source.
- If cr1.mtv also fails, cr1.paa sources the traffic to cr1.kcy.

Figure 3: Example Failover for Multiple Data Centers



If priority 1 member fails, traffic is sourced from priority 2



If both priority 1 and priority 2 members fail, traffic is routed on the priority 3 member

380796

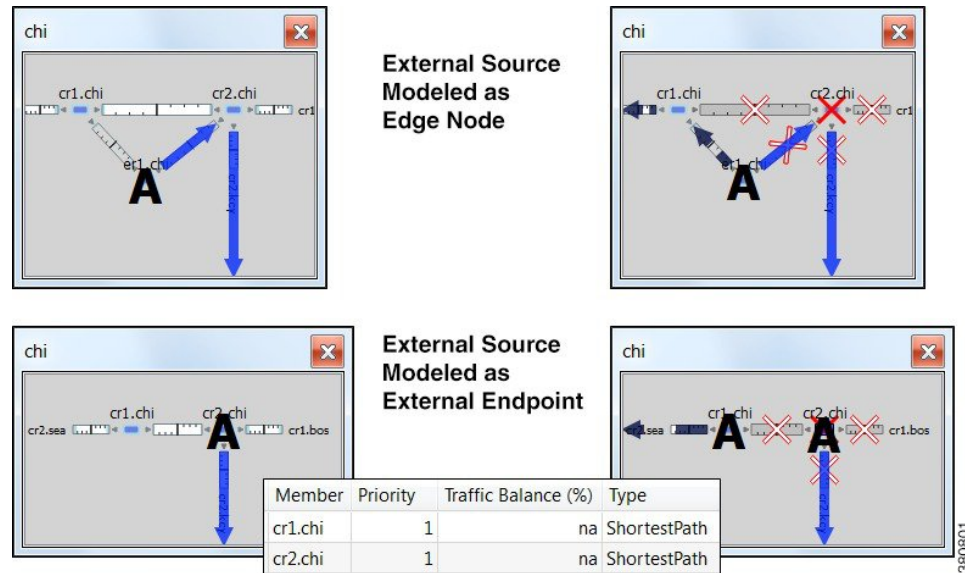
External Source

One method of modeling an external demand source with multiple entry points is to construct an edge node, connecting it to these entry points, and giving the interfaces equal metrics. Then source the demand from this edge node.

Alternatively, you can construct these multiple entry points using a single external endpoint as the demand source, thus reducing the complexity of the topology.

In this example, a demand has an external source with two entry points that come into the network at Chicago. [External Source, on page 8](#) demonstrates two ways to model this source. One uses additional topology (edge node) to source the demand. The other uses a single external endpoint that has two members (cr1.chi and cr2.chi), both set to priority 1 with shortest path traffic.

Figure 4: Example External Demand Source



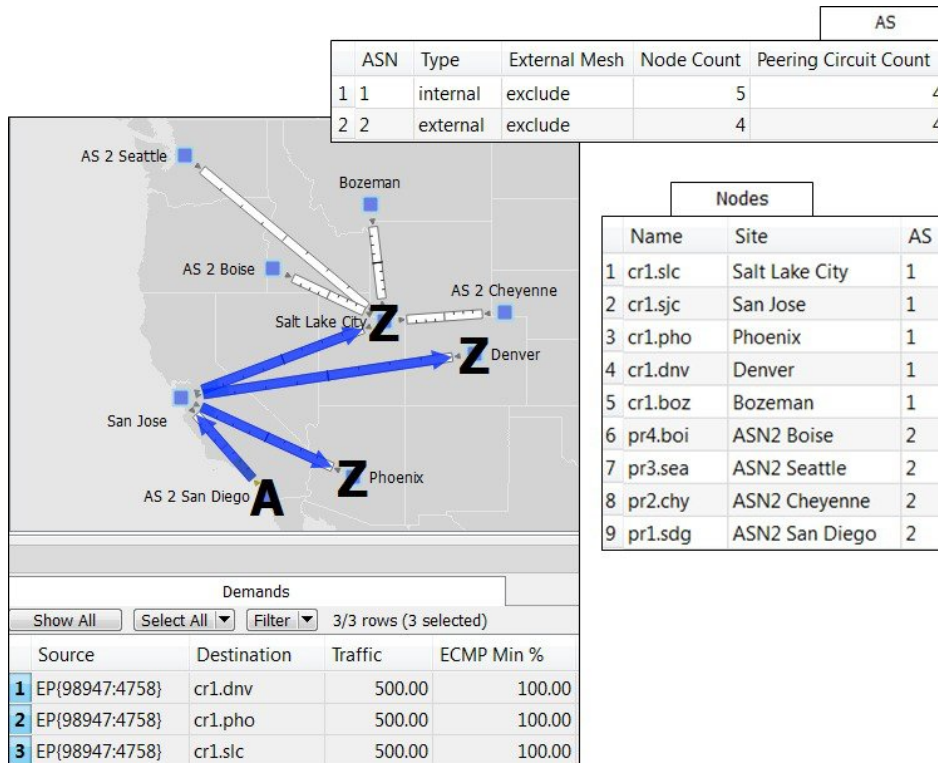
Intra-AS Routing Examples

Failover to Nodes in External ASes

This example demonstrates a content provider (Acme) who sources most incoming Internet traffic through a transit provider, and this traffic is then distributed to Acme's customers. Acme has identified specific policies for each market using communities, and these policies determine, per location, which transit provider peering points are used for primary routing and failover. The network is set up as follows (Figure 5: Example Transit Peer AS 2 San Diego Sends Traffic to Customers in an Internal AS, on page 10):

- The San Diego, Boise, Cheyenne, and Seattle nodes each belong to the same transit provider, AS 2.
- The remaining nodes belong to Acme's internal AS (AS 1), and these are Acme's customers.
- The San Diego node (pr1.sdg) in AS 2 is sending traffic to three key customers located in Phoenix (cr1.pho), Salt Lake City (cr1.slc), and Denver (cr1.dnv).

Figure 5: Example Transit Peer AS 2 San Diego Sends Traffic to Customers in an Internal AS



380807

In WAE Design, Acme uses an external endpoint named after community 98947:4758. It contains the following members, all of which are in the same AS 2 transit provider:


- pr1.sdg set to priority 1 with shortest path traffic
- pr2.chy set to priority 2 with shortest path traffic
- AS 2, which means the remaining nodes in the AS, set to priority 3 with shortest path traffic

As [Figure 5: Example Transit Peer AS 2 San Diego Sends Traffic to Customers in an Internal AS, on page 10](#) shows, multiple demands are created using this external endpoint (98947:4758) as the source, and each sending traffic to a different customer destination. [Figure 6: Example Failover Within the Same External AS, on page 11](#) focuses on the traffic destined for Denver (cr1.dnv), showing that the failover scenario behaves as follows:

- If a failure prevents the peer pr1.sdg node from sending traffic to cr1.pho, a failover occurs to pr2.chy within that same peer to avoid a service disruption.
- If failures prevent pr1.sdg or pr2.chy from sending traffic to cr1.dnv, the remaining nodes in the AS source the traffic. Although this is set to use the shortest path, because there are multiple paths, the traffic is distributed evenly.

The same result could have been achieved by setting the priority 1 to an interface going from San Jose to pr1.sdg and by setting the priority 2 to an interface going from Salt Lake City to pr2.chy.

Figure 6: Example Failover Within the Same External AS



Failover from an External AS to Another

In this example, a small ISP (AcmeNet) buys transit from two other service providers. AcmeNet peers with AS 2 in Phoenix and AS 3 in Denver. AcmeNet has a customer base in Phoenix (cr1.pho). The customer has a preferred access to the Internet through San Francisco (pr1.sfo) in AS 2 (Figure 7: Example Transit Peer ASN2 San Francisco Sends Traffic to Customers in an Internal AS, on page 12), and the secondary access is through Denver (pr2.dnv) in AS 3 (Figure 8: Example Failover to a Different External AS, on page 12). In WAE Design, AcmeNet simulates this network by creating an external endpoint (Internet) containing two different external ASes as members.

- AS 2 is priority 1 with shortest path
- AS 3 is priority 2 with shortest path

The demand is sourced from this external endpoint named Internet.

Figure 7: Example Transit Peer ASN2 San Francisco Sends Traffic to Customers in an Internal AS


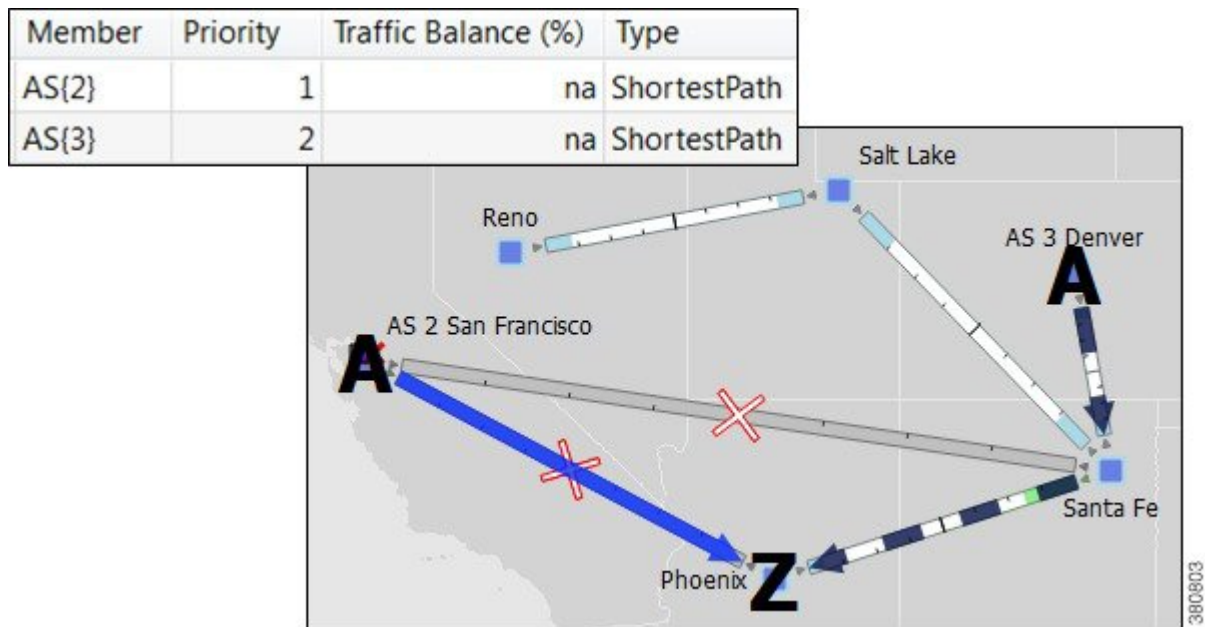


Figure 8: Example Failover to a Different External AS



Creating External Endpoints and Their Members

The recommended method of creating external endpoint members is to do so while creating the associated external endpoint, as follows.

Step 1 Right-click in an empty area and choose **New > Demands > External Endpoint**, or choose **Insert > Demands > External Endpoint**. The External Endpoint Properties dialog box opens.

To edit an existing external endpoint, double-click it from the External Endpoints table.

Step 2 Enter a unique external endpoint name in the Name field.

Step 3 To add a new member, click **New**.

To edit an existing member, select it from the Members table and click **Edit**.

An external endpoint member dialog box opens.

- a) Specify the member type as **Node**, **Interface**, or **External AS**.
 - For nodes, choose the site and node name.
 - For interfaces, choose the site, node, and interface. Using an interface lets you specify the exact interface on which the demand traffic is going into or out of a node.
 - For ASes, choose the AS name, and select either the “unspecified” option or node name within the AS. The “unspecified” option routes traffic evenly throughout the entire AS.
- b) Enter a priority number to determine the sequence in which the member is used with other members in this external endpoint. Multiple members can have the same priority.
- c) Select whether to use a shortest path route, or to direct a specified percentage of traffic across members that have the same priority (Fix Traffic or Deduce Traffic).
- d) If balancing traffic, enter the percentage of traffic this member should route.

If you set Deduce Traffic, this value is updated by Demand Deduction provided there are no members with a higher priority that are in use under normal (non-failure) operation.
- e) Click **OK**.

Step 4 Click **OK** in the External Endpoint Properties dialog box.

