cogent COMMUNICATIONS

Optical Internet

Redundancy vs. Capacity

A Cogent Whitepaper

Customers requesting connectivity to redundant Cogent routers frequently expect that their traffic will load-share or load-balance reasonably well in both directions. In the simplest case where the customer has a single router, it is relatively easy for the customer to load-balance his traffic in the outbound direction (i.e. towards Cogent), so long as the customer's router is running BGP with a loopback address, has set "maximum-paths" to a minimum of 2, and so long as the redundant connections to Cogent appear as equal costs routes in his routing table **(Figure 1)**.

Depending on how sensitive a customer's router might be to MEDs, slight differences in incoming MEDs passed to the customer on routes from Router A and Router B may also influence how well the traffic load-balances in the outbound direction. In such instances, a possible workaround is for the customer to reset the incoming MED from Cogent on both sessions to 0. There is no downside to doing this so long as the customer doesn't have multiple connections to Cogent in multiple cities.

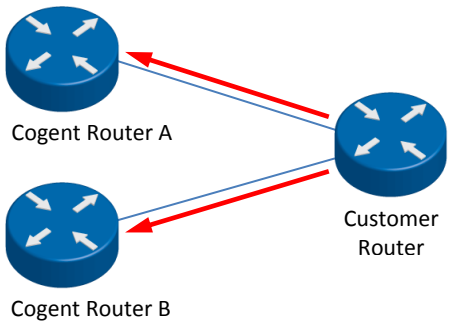

(Figure 1)

In the inbound direction, however, the degree to which traffic will load-share depends to a large extent on where the traffic is coming from, and on whether the source of the traffic is likewise redundantly connected to Cogent. In the following example, we would expect traffic from a content provider to load-share properly as it makes its way towards the customer, if both the source and the endpoint are redundantly connected **(Figure 2)**.

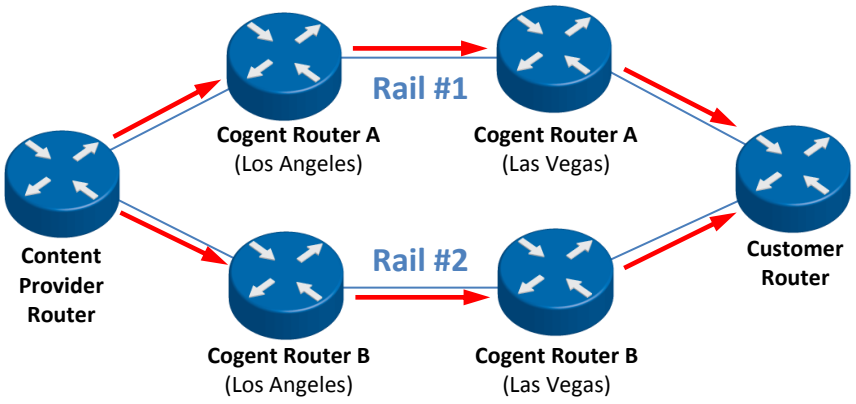

(Figure 2)

If, however, the content provider isn't redundantly connected to Cogent, then there is no way to load-share the traffic across both of the customer's connections to Cogent **(Figure 3)**.

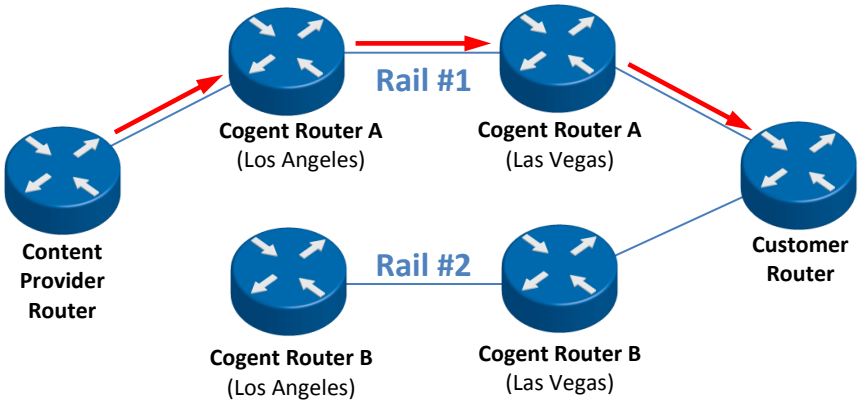

(Figure 3)

In any real world scenario, the customer should expect to pull Internet traffic from a wide variety of geographically dispersed sources, some of which are redundantly connected to Cogent's network, but most of which are not. Given these considerations, it is difficult if not impossible to predict how well inbound load-sharing will work for any given customer.

Customers can, however, influence their inbound traffic somewhat by changing the local preference that is applied to their routes on the Cogent network. Thus if a customer is announcing multiple routes, a simple way to manually balance the traffic would be to split the announcements so that some routes have the default local preference of 130 on their connection to Router A. The remaining routes on that same connection to router A could then be set to a lower local-preference of 125, by sending the community 174:125 on those routes. On the connection to Router B, the exact opposite would be configured – i.e., the first set of routes would be sent with the 174:125 community, while the remaining routes would receive the default local preference of 130. Note that this approach requires some fine tuning and that results will vary.

In the final analysis, customers must therefore decide whether redundancy or capacity is more important. If capacity is most important, then all a customer's connections should ideally terminate on a single Cogent router at any given POP, whereby load-balancing in both directions in easily achievable. In the case where redundancy is more important, the customer would want to avoid filling both connections to Cogent anyway, as doing so runs the risk of dropping significant traffic in the event one of those connections fails. Accepting the premise of such a constraint, the degree of inbound load-sharing across redundant connections becomes insignificant so long as aggregate utilization across both links doesn't exceed 50%.