IXP Route Server

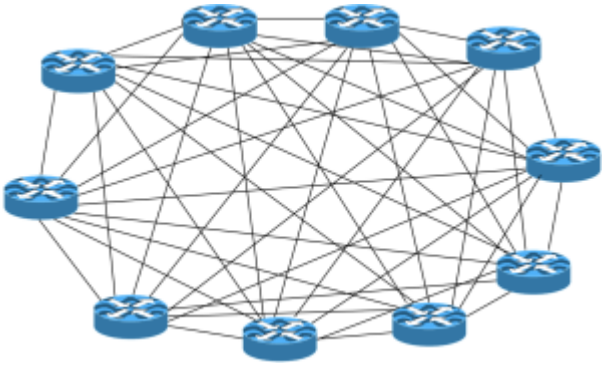
A Route Server has all the features of a Route Collector, but also announces routes to participating IXP members according to their routing policy definitions.

They are implemented using the same specification as for a Route Collector.

Purpose

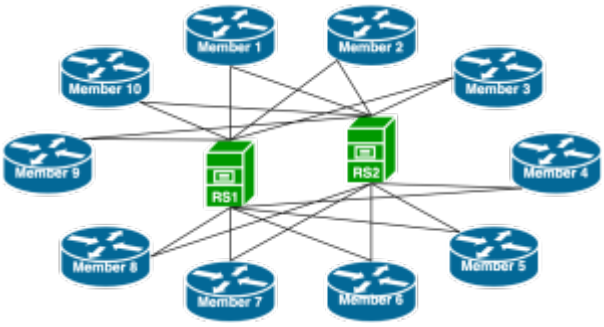
The main function of a Route Server is to scale the route distribution at Internet Exchange Points. Rather than each participant at an IXP having to peer with every other participant, they only peer with the IXP's Route Server(s), thereby only needing to maintain the BGP session(s) to the Route Server(s).

Without a Route Server, the BGP mesh at an IXP would look similar to the following diagram:



basically an n-squared mesh.

Introducing a Route Server results in the BGP sessions for each member being with the Route Route Server, meaning that the number of BGP sessions goes up linearly per member (from Route Server point of view).



with each member only having to maintain their BGP sessions with the two Route Servers.

Routing & Forwarding

It is important to note that the Route Server is not involved in packet forwarding at the IXP, nor does it carry any of the IXP's traffic. The Route Server takes advantage of a BGP feature known as “third

party next-hop" where by the IP address of the "next router" (where the prefix was heard from) is unchanged when it is passed on to the other peers on the same LAN.

The Route Server also does not insert its Autonomous System Number into the AS Path. This means that, to all intents and purposes, it appears as though the IXP members are peering directly with each other, as their AS Numbers appear immediately adjacent in the AS Path. BGP implementations adhering to the BGP specification as per RFC4271 will need to disable the AS adjacency check for any session with a Route Server - BGP requires that the peer AS must be the first AS in the AS_PATH, something that a Route Server does not do.

On Cisco IOS, for example, the configuration commands to do this are:

```
router bgp <ASN>
 no bgp enforce-first-as
```

Note that on Cisco IOS there is no ability to set this per neighbour, something which needs to be considered if the peering router used has multiple private and public peers connected to it.

Back to "Peering Technologies" page

From:
https://www.bgp4all.com.au/pfs/ - **Philip Smith's Internet Development Site**

Permanent link:
https://www.bgp4all.com.au/pfs/peering-toolbox/route_server?rev=1659265894

Last update: **2022/07/31 11:11**