

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:



where each line represents a BGP session between two IXP members. This is basically an n-squared mesh.

Introducing a Route Server results in the BGP sessions for each member being with the 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](https://tools.ietf.org/html/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.

Route Server Uses

Apart from helping scale the BGP mesh at larger IXPs, Route Servers have other value add considerations for Internet Exchange Points.

Bogon Filtering

Internet Exchange Points often use their Route Servers to implement bogon prefix filtering and bogon

AS filtering. While members of IXPs will normally do this if they are following [MANRS](#) principles, it is often helpful for the IXP to also implement this via the Route Server in case the member configuration is omitted.

Bogon prefix filtering means dropping all prefixes which are not used in the public Internet. This includes private, reserved, and unassigned address space.

Bogon AS filtering means dropping all prefixes originated by Autonomous Systems using private, reserved, or unassigned ASNs.

Route Origin Validation

Route Servers are also used nowadays to implement Route Origin Validation, also noted in the [MANRS](#) best practices. This ensures that any prefix with invalid Route Origin Authorisations are dropped by the Route Server, so that misconfigurations or route hijacks are not propagated between members across the IXP fabric.

Route Servers and Peering Policies

Operators who have an Open Peering policy are typically those who would peer at a Route Server. There is no operational relationship between the peers when they use a Route Server, as each peer only sets up their BGP session with the Route Server.

Operators who have a Selective Peering policy will typically not use the IXPs Route Server unless the Route Server has been set up to allow per-member policy. Operators with Selective Peering policies normally want to have a direct operational relationship with their peer, and would typically set up a Bi-Lateral peering with them.

Implementations

There are several Route Server implementations, which all run on Linux or BSD based Unix implementations:

- [BIRD](#) - the standard & works best
- [GoBGP](#)
- [OpenBGPD](#)
- [FRR](#) (fork of Quagga)
- [Quagga](#) (LINX fork)

The majority of Internet Exchange Points use BIRD, and it integrates very easily with [IXP Manager](#), the tool used by many IXPs to manage the operation of their peering infrastructure.

At least Cisco and Juniper have Route Server implementations in their router software, but they are generally not recommended for use today (in the author's experience).

On Cisco, from IOS 15.2 and IOS-XE 3.7, the BGP configuration command to enable a neighbour as a Route Server Client is:

```
router bgp <ASN>  
  neighbor <addr> route-server-client
```

and on Juniper as from JunOS 17.4, the BGP configuration command to enable a neighbour as a Route Server Client is:

```
set routing-instances protocols bgp group <name> route-server-client
```

[Back to "Peering Technologies" page](#)

From:

<https://bgp4all.com/pfs/> - **Philip Smith's Internet Development Site**

Permanent link:

https://bgp4all.com/pfs/peering-toolbox/route_server

Last update: **2022/08/26 10:04**

