



Single Upstream with Two Private Peers

This section discusses the next stage which is commonly encountered after a network operator establishes their first ever peering connection: another party is interested in peering as well, because they see the benefits that can be gained.

To add the second peering, we build on what we learned in the previous section, noting that we have already deployed our address space, set up IBGP, created Route Objects and ROAs, and created policies for our EBGP sessions.

Enabling the second Peer

We already have a network running IBGP, with EBGP operating with the first peer we set up. We may also have set up EBGP with our upstream provider (strongly recommended).

The steps to add the second peer are rather simple, as the hard work has already been done. We can use the same router as for the initial peer, just assign a different interface.

Once we have installed the physical connection, and assigned IP addresses to the point to point link, we are ready to bring up the BGP session with the peer.

We already have our outbound policy decided and configured on the router (the policy is a prefix filter allowing just your address space outbound). But we need to create an inbound policy with this new peer. And that's a simple prefix filter which permits their address space inbound (the same concept as we used for the first peer).

And once the policy is ready and configured, we are ready to bring up the EBGP session. The prefixes learned will be propagated by the IBGP across the network, and all devices connected will now have a direct path to this new peer.

Notes about the new peering

Filtering

As the name suggests, the peering is private. Which means only you see your peers prefixes; and they only see your prefixes. Because your prefix filter only allows your prefixes out, and their prefixes in, nothing else connected to your network (upstream or other private peer) will be able to see their

prefixes, or access their network through you. This demonstrates the importance of always implementing prefix filters, inbound and outbound on EBGp sessions.

[RFC8212](#) discuss default router behaviour, noting the recommendation that no EBGp session should be established with out filters in place.

And [MANRS](#) first action is to “Prevent propagation of incorrect routing information”, and our inbound and outbound filters applied on the peerings ensure that only the correct prefixes are exchanged over the EBGp sessions.

Route Object & ROAs

No change is needed in the Route Object in the IRR, or the ROAs created for the address space. Both indicate the ASN originating the prefix and this does not change when a new peering is established.

Scaling

Adding further private peerings can be done using the same process. However, when many of these peers are in the same location, it is more efficient to create what is known as an Internet Exchange Point, and this will be examined in the next section of the Peering Toolbox.

[Back to 'Establishing Peering' page](#)

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

Permanent link:
https://mail.bgp4all.com/pfs/peering-toolbox/single_upstream_two_private_peer?rev=1652510930

Last update: 2022/05/14 06:48

