Single Upstream & Private Peer

Most network operators first encounter with peering is when they have a simple connection to their upstream provider. This connection usually involves the operator having a static default route configured on the link to the upstream, with the upstream pointing a simple route to their customer.

We will now look at how we go about adding a peering with another network. There are two possible cases here:

1. Adding the peering, retaining static route to upstream
2. Adding the peering, introducing BGP with the upstream

Enabling the Peer

Deploying Address Space

If we are not already using our own address space, then we need to consider first how to go about deploying it across our own network (and any customers we might have).

We already have obtained our own address space and ASN for the network. As noted elsewhere, most upstream providers forbid the use of their address space for connecting with any other operator. Renumbering a network is beyond the scope of the Toolbox, but there are many published documents and guides on how to do this. At a high level, the steps are:

1. Create a Route Object for your address space using your upstream provider's ASN
2. Create another Route Object for your address space using your own ASN
3. Create a ROA for your address space using your upstream provider's ASN
4. Create another ROA for your address space using your own ASN
5. Provide a Letter of Authority (LOA) to your upstream provider requesting their peers and transits accept and propagate your address space (some providers request this, so it is good to be prepared)
6. Organise with the upstream (transit provider) for the address space to be announced globally and routed to you
7. Renumber the network (change dynamic pools, and use secondary addressing if needed)
8. Withdraw the old address space
9. Return the old address space to the upstream

Deploying IBGP

Once our own address space is in use for the network, BGP needs to deployed internally (IBGP) across the network (at least for the devices in the core and border of the network). If the network is just of a single router, no IBGP is needed.

There are many online guides on how to deploy IBGP so will not be covered here. The assumption is that an interior routing protocol like OSPF or ISIS is already operating - if it is not, then this will also need to be deployed before IBGP can be deployed. The AS number that BGP requires is the one already obtained from the Regional Internet Registry.

Deploying EBGP with Peer

The final step is to deploy EBGP with the brand new peer. With BGP already running on the border/peering router, all that needs to be done is a session added to talk to the new neighbour.

First we create our BGP policy:

- outbound we announce just our address block, so a prefix filter allowing our address block out to our peer is sufficient here
- inbound we accept just the prefixes our peer will send us, so another prefix filter allowing those in from our peer is sufficient here

And then we set up the EBGP session with our direct peer.

The prefixes learned will be propagated by the IBGP across the network, and all devices connected will now have a direct path to the peer.

BGP with Upstream

Here we go one step further that we did in the previous section, by setting up EBGP with the upstream provider as well. This is preferred as it makes the network truly autonomous and globally visible too.

The address space still needs to be deployed, IBGP needs to be set up across the backbone, and EBGP needs to be set up with the peer. These three steps were described in the previous section. You need to do those first!

Enabling EBGP with the upstream will now make the network operator's ASN visible to the global Internet. There are very distinct steps to enabling this.

Request to use EBGP

We need to discuss with the upstream provider about setting up the EBGP session. They need to know what your intentions are. If they refuse to use EBGP, you can either stay with static routing, or search out a new upstream provider (the latter step is preferred, and usually causes a change of heart of the recalcitrant operator).

Letter of Authority

The next step is to provide them with a Letter of Authority (if required) which requests their upstreams and peers to allow your address space originated by your ASN.

LOAs are not usually required as a ROA should be enough to prove the holder of the address space and the origin ASN. But some operators insist on the LOA, still.

ROA and Route Object

Confirm that the ROA (and Route Object) with your ASN as the origin of your address space is still present in your RIR's database. Note that if you are migrating from a static set up, do **NOT** delete the existing ROA (or Route Object) that declares their ASN as the origin - you still need it for now.

BGP Policy Configuration

Then you need to create two policy statements, one for incoming policy, the other for outgoing policy. The incoming policy statement will likely say that you'll only accept a default route. Most operators will either:

- send a full BGP table, or
- send just a default route

You do not need a full BGP table in this simple case - a default is sufficient.

The outbound policy statement will be the same as the one you use for the peer, namely announce just your address space/block.

Establishing EBGP with Upstream

Once all the previous steps have been completed, you will be ready to set up the EBGP session with your upstream.

Verify that you are getting a default route, and that they are receiving your route from you.

You can then remove the static default route pointing to them.

And during their maintenance window, they can remove the static route for your address space pointing to you. Note that when they do this, they will stop originating your address space from their AS. Which means that the announcement you make to them from your ASN will become globally visible. This change needs to be done once everyone is sure, and the upstream has confirmed, that all their upstreams and peers have updated BGP filters accordingly. It is best that this latter work is all carried out during a maintenance window, and preferably at least 24 to 48 hours after bringing up the EBGP session, to avoid unexpected outages.

Clean up

With EBGP running successfully with the upstream provider, and the address space visible globally (check on RouteViews where multiple paths should be visible), we can now clean up.

The Route Object and ROA created with the upstream provider's ASN can now be deleted, leaving on the ROA and Route Object with your own ASN. This should have no operational impact whatsoever.

Back to 'Establishing Peering' page

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

Permanent link:
https://www.bgp4all.com.au/pfs/peering-toolbox/single_upstream_private_peer?rev=1652509855

Last update: **2022/05/14 06:30**