


Single Upstream with IXP

This section discusses how we scale multiple peerings with our network, using what is known as an Internet Exchange Point.

Internet Exchange Points are open neutral interconnects where network operators (with their own Internet resources) are able to freely interconnect. An IXP is the most efficient and effective way of scaling interconnections between network operators in any one location.

Lots of information about IXPs is available from many locations, including the Euro-IX website, the IXPDB, as well as in the links noted at the foot of the page.

A diagram showing the typical physical layout of this scenario is shown below:



Participating in an IXP

The section describes how to participate at an Internet Exchange Point. The description is high level as each and every IXP will have their own nuances, variations on the general theme. Discussion with the IXP operator is important to understand their requirements.

We won't discuss why joining an IXP is important - the Value of Peering has already covered why peering is essential for a network operator's business.

Nor will we discuss which IXP to join - there are many factors involved, but common advice is to join

the "local IXP" as that will host network operators with similar common interest, content, and customers, and likely will give the best peering opportunities.

- Joining the IXP
- Physically Connecting to the IXP
- · Connecting to the IXP by Remote Peering
- . Establishing Peering at the IXP

Joining the IXP as a Member

Every Internet Exchange Point will have some form of requirements to join them so you can participate in peering there.

Requirements can be as simple as:

- Agreeing how to access the location, building, datacentre (both for putting connectivity in there, as well as for human access for maintenance work)
- For non-profit member driven IXPs, becoming a member of the IXP
- For commercial IXPs, agreeing and signing a contract of engagement
- Understanding how to establish peering (be it bi-lateral with other members, or via the IXP's Route Server infrastructure)
- Understanding how to use the IXPs member portal (IXP Manager or other).
- Agreeing on any annual cost sharing or fees for the IXP
- Assignment of IP addresses for IXP LAN, and information about Route Servers (if applicable)
- Agreeing basic best practice behaviours

Once the administrative aspects have all been agreed and finalised, we can get on with the task of connecting to the IXP and reaping its benefits.

Physically connecting to the IXP

Physically Connecting to the IXP is covered elsewhere in the Toolbox. That discussion covered both the case where the new member provides their own media to get to the IXP location, or contracts a third party who operates a layer-2 infrastructure connected to the IXP.

Establishing Peering at the IXP

The final part of the process is to establishing peering with other members of the IXP. Most IXPs will offer two methods and we'll look at these now:

- 1. Peering with a Route Server
- 2. Bilateral Peering

Some IXPs will also provide a facility called a Looking Glass. This allows members of the IXP, and often members of the public, to view the BGP table as seen at the IXP.

https://bgp4all.com/pfs/ Printed on 2025/10/17 06:31

2025/10/17 06:31 3/4 Peering at an IXP

Route Server

A Route Server is a device at the IX (there are usually two independent Route Servers) which peers with every member of the IXP. It receives all the routes each member announces to it, and announces all routes it has received to all members.

This is the basic behaviour of a Route Server used in most IXPs around the world. The Route Server is a BGP daemon running on a Linux or FreeBSD virtual machine (most common implementation). The most widely used implementation is BIRD, although some IXPs use FRR, GoBGP, or OpenBGPD.

For a newcomer to peering and BGP in general, setting up a session with the Route Servers at the IXP is the easiest way to get up and running.

Your existing outbound policy applies with the Route Server peering too - you have a prefix-list which only allows your prefixes out to the EBGP peer. Inbound policy, in the basic instance is quite simple: you set up a prefix-list that allows everything, but set up a prefix-limit on the EBGP session to 100% more than the number of routes the Route Server is advertising (which you will find out once you bring the peering up). This protects against any of the peers at the IXP accidentally announcing a large portion of the BGP table via the Route Server. Note that most IXPs will have this protection on their Route Server in any case, but it's a good idea/recommendation that you do this too.

And then establishing the EBGP session is the same as for any private peer, as we saw earlier. There is one point to note though. The Route Server will NOT insert its AS number into the AS path of the routes you will hear from the IX. BGP implementations which conform with the standard require that the first AS in the path is the same as that of the peer AS - so this will cause an issue. You need to turn this feature off. On Cisco, for example, the command is:

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

Thereafter the EBGP session with the Route Server will be established, and connectivity to all the IX peers will be via the IX LAN. (Note that traffic does not go via the Route Server.)

If the IX has two Route Servers, bringing up the EBGP session with the second one will be by the same process - and provides important redundancy should either of the Route Servers go off line (for maintenance or otherwise).

Bilateral Peering

The other type of peering at an IXP is known as Bilateral Peering, and is where one member sets up an EBGP session directly with the other member across the IXP fabric. This type of peering is used by network operators who implement a Selective peering policy.

Establishing a peering with such an operator usually requires initiating contact with them first (via the IXP membership portal or via PeeringDB), agreeing on the peering and any other requirements that either operator may have.

Once this is done, establishing the EBGP session is no different from establishing EBGP with a private peer. Your outbound policy is already known, and the inbound policy only needs to be a prefix filter allowing the prefixes that the other operator said they'd be announcing.

Looking Glass

Most IXPs offer a general service to members and the public called a Route Collector. A Route Collector collects routes from peers for use by a Looking Glass.

A Route Collector is functionally identical to a Route Server but with one very important difference: it does **NOT** send any routes to its peers.

The Looking Glass allows the general Internet public see what prefixes are available at the IXP, and is a very valuable tool for any network infrastructure. The use of a Looking Glass will be covered elsewhere in the Toolbox.

Some IXPs will operate their Looking Glass infrastructure via the Route Servers.

Others IXPs keep the facility separate and will request their members to set up a bi-lateral peering with their IXP's Route Collector - in that sense the configuration is no different from any other bi-lateral peer, but there will be no routes received from the Route Collector itself.

All IXP members are encouraged to peer with the Route Collector - it helps with awareness, promoting peerability (showing the available routes), and equally importantly, with troubleshooting connectivity and reachability issues.

References

This content is sourced from many contributors, including:

- IXP Design Presentation Philip Smith
- Value of Peering Presentation Philip Smith
- BGP Videos Network Startup Resource Center

Back to 'Establishing Peering' page

From:

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

Permanent link:

https://bgp4all.com/pfs/peering-toolbox/single_upstream_ixp?rev=1659337437

Last update: 2022/08/01 07:03



https://bgp4all.com/pfs/ Printed on 2025/10/17 06:31