



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.

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

Connecting to the IXP involves multiple stages, from getting to the location, to physically connecting to the IXP ethernet switch.

Stage One

The first step needed before physically connecting to any IXP is to ensure that there is an available router. An IXP is a large layer 2 network to which many operators are connected, so the aspiring member needs to provide either a router physically at the IXP itself or an available router port on their own network infrastructure.

This router is normally dedicated only for peering connections, whether connecting to public peers at an IXP or private peers. If procuring a separate router is not a possibility, it is possible that an existing router could be used, so long as it meets the appropriate technical requirements for participating in a peering infrastructure (full support of BGP, sufficient controlplane memory and CPU capacity).

If this router will be installed at the IXP location, appropriate arrangements need to be made to procure it and have it delivered, installed, and configured, to coincide with the delivery and commissioning of the physical link back to the aspiring member's network infrastructure.

Stage Two

Next we need to get to the location where the IXP has been established, usually in a datacentre or some independent or neutral data housing facility. These days access is mostly by using fibre optic, which the network operator will arrange with the fibre optic provider. And there are many possibilities here too, depending on the country or region of the world:

- dark fibre (fibre pair for exclusive use for the network operator)
- a wavelength on DWDM network
- fibre bundle installed by the network operator themselves

Where it is not possible to get fibre access, other methods include point-to-point microwave or 802.11-based links, traditional TDM leased lines, or via a third-party's layer-2 network (known as

Remote Peering).

Stage Three

If the new member is providing their own connectivity to the IXP, their next step is getting from the connection media entry point to the IXP switch itself. This is usually done by the building operator (most don't want 3rd parties installing anything inside their premises). The end result is that there will be a fibre or ethernet presentation at or near the IXP switch, usually in a separate patch panel. This usually depends on whether the network operator connecting is going to locate a router at the IXP (usually recommended when the media to get to the IXP location is not fibre optic), or simply connect from their own point of presence directly to the IXP.

If the new member is using a third party's layer-2 service, that layer-2 operator will already be physically present and connected to the IXP, so this step will not apply.

Stage Four

This assumes the new member will install their own router. For this, agreement with either the IXP or the building operator to locate a router is needed (and involve fees to cover the space, power needs, air-conditioning, and remote hands). This router will very likely NOT be installed in the IXP equipment rack but elsewhere in a common user space. The router is usually only a single rack unit tall (it needs one interface to connect to the IXP, and another interface to connect back to the network operator's main network) so the space requirements are minimal.

If the new member is using a third party's layer-2 service, no router needs to be installed at the IXP itself - the new member simply needs to have a router port available on their own infrastructure to connect to the layer-2 operator's infrastructure.

Stage Five

The fifth and final stage of connecting to the IXP is plugging the network operator infrastructure into the IXP itself. Most IXP switches today are fibre optics based, with ports supporting 1Gbps or 10Gbps depending on the fibre optic transceiver (SFP) installed. There are two possibilities here:

1. The IXP membership fee includes providing an SFP for the member to connect to the switch. It's more likely for the IXP to have an inventory of SFPs for their particular brand of switch. In this case the new member only has to provide a suitable SFP for their router - or if their router has no fibre port, the IXP has to use an SFP that supports an RJ45 copper connection instead. The IXP will also patch their switch to the member infrastructure (whether it is a single mode fibre optic patch lead, or Cat6 ethernet cable if fibre is not feasible).
2. The member has to bring their own SFP to connect to the IXP switch. The IXP operator simply provides the switch, and it is up to the member to procure and provide the SFP and the suitable single mode fibre optic patch lead to get from their installation to the IXP switch. The IXP (or building operator) will still do the install though.

Connecting to the IXP by Remote Peering

Connecting to the IXP via Remote Peering is greatly simplified as the layer-2 infrastructure provider will already be present at the IXP.

Stage One

The first step for connecting to an IXP by using a Remote Peering service is to contract with an operator who is already present at the IXP in question to provide the agreed layer-2 capacity to the IXP.

Details of what this service and contract should look like are beyond the scope of the Peering Toolbox. However, it is important to ensure that the committed bandwidth from the layer-2 provider will meet the needs of the new member, and that this capacity can be easily upgraded or downgraded on reasonable notice. The last thing any IXP member will want is a congested IXP connection.

Stage Two

The next step is to ensure that there is an available router. This router is normally dedicated only for peering connections, whether connecting to public peers at an IXP or private peers. If procuring a separate router dedicated to peering is not a possibility, it is possible that an existing router could be used, so long as it meets the appropriate technical requirements for participating in a peering infrastructure (full support of BGP, sufficient controlplane memory and CPU capacity).

Stage Three

Once the new member's router is in place and operational, a free port needs to be connected to the layer-2 infrastructure operator. The most common way is by ethernet whereby the remote IXP LAN is delivered over a particular identified VLAN. This method allows the layer-2 infrastructure operator to provide several different connectivity services to its customers, each allocated their own VLAN.

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. [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.

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 Looking Glass.

A Looking Glass 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 Looking Glass - in that sense the configuration is no different from any other bi-lateral peer, but there will be no routes received from the Looking Glass itself.

All IXP members are encouraged to peer with the Looking Glass - 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=1659247150

Last update: **2022/07/31 05:59**

