Indiana University School of Informatics and Computing

HATS Research Group

SDN: IXP Case Study

IXP USE CASE

An Internet Exchange Point (IXP) is, essentially, a switch-that is, Layer 2 infrastructure to which many ISPs can connect, and over which those ISPs can establish Peering (or, possibly, transit) connections.

The larger IXPs operate 1+1 redundant switching infrastructure, in some cases supporting a virtual Layer 2 mesh so that failures of the underlying infrastructure are invisible. For large flows and connections, some IXPs provide direct Layer 1 connections between ISPs and some use Layer 1 switches to switch connections between redundant devices.

An IXP is fairly straightforward and should change only as connections are added or removed. For the largest IXPs the challenges are traffic volumes and reliability requirements. Nevertheless, centralising the configuration and control of the switching layer could make the job of running an IXP easier. The virtualisation of the switch infrastructure may be achieved more easily and less expensively with a new (SDN) control plane over an (OpenFlow) data plane.

At an IXP it is up to the individual ISPs to establish and maintain eBGP sessions, to exchange routes and traffic. The IXP itself is not involved in this process, though the IXP must provide adequate capacity-when the ISPs ask for it-and it is in the IXP's reputational and commercial interest to ensure that all its clients maintain adequate capacity.

However, most IXPs also provide a Route Server. To peer with others at the IXP an ISP can connect to the Route Server, so that a single eBGP session between the ISP and the Route Server replaces many connections between the ISP and all the other ISPs at the exchange. This replaces a full mesh of individual eBGP sessions by a hub and spoke arrangement, where the Route Server is, effectively, a proxy for all its clients. The principal advantage of the Route Server is that a new ISP joining the exchange does not have to persuade all existing ISPs at the exchange to establish a new eBGP session-where the marginal cost to each of the existing ISPs can mean that the task lingers at the bottom of the list.

Many ISPs connect to the Route Server and will peer with any and all other ISPs. Some ISPs connect to the Route Server but wish to pick and choose their peering partners. So, the Route Server most provide a "peering matrix" function, which provides, at a minimum, each ISP with ability to deny or permit announcements to other ISPs. Because the Route Server exchanges routes using BGP, where it has more than one route for a given prefix the Route Server must select the best one on behalf of each client. So, each client is delegating some policy to the Route Server-though current Route Servers do not allow the client much, if any, control over the route selection made on its behalf. This is an area where opening up the eBGP routeing layer may allow the Route Server and its clients to cooperate more closely, or to do away with the Route Server altogether (either by automating the process of setting up a new peering connection, or by providing a form of broadcast for announcements).

Some IXPs configure their Route Servers to filter incoming route announcements (that is, announcements coming from each client) to allow only prefixes which are known to be valid. This filtering may be configured from routeing policy published in an RIR. It is not, strictly speaking, the IXP's business what routes its clients announce to each other. However, some feel they should ensure that only valid routes are announced by their Route Servers.

In a "Software Defined IXP" we may see a full integration from the ISP policy down to the packet forwarding, so that:

- routes announced (broadcast) by the Route Server are checked for validity against the originating clients' policies.
- packets forwarded are checked for validity against the routes announced.

The effect of which is to ensure that nothing crosses the exchange which should not cross it.

The control plane of the IXP is entirely a matter for the IXP. So, this might be secured simply by ensuring that no third part has access to the control and command network(s). But the control plane would be more secure if it were secured as if third parties could access it. There is not much which is novel here.

The use of Route Servers at exchanges is an interesting example of what may be achieved by the SDN approach. First, providing a better and less expensive way to build and manage an IXP. Second, by allowing the control plane to be extended-once it is unbundled from the data plane, and given fine grained control over the forwarding plane-to improve the working of the network: in this case by ensuring the IXP only carries the routes and traffic it is intended to carry.

Further, given the ability of both the IXP and the clients to extend what the Route Server does and how clients interact with it, it would be possible for each client to apply its own policy to the routes available at the IXP. This implies some mechanism for each client to influence "its part" of the Route Server, in much the same was as a Virtual Private Network client needs to be able to manage "their part" of the host ISP's network.