

AFIX Technical Workshop: Session 5

Preparation for a real IXP

Contents

Contents	1
Each ISP gets an ASN from AfriNIC	1
Fees.....	2
IXP operator to obtain IP addresses	2
Each ISP installs a router at XP	2
Configure iBGP mesh.....	3
Import internal routes to iBGP	3
Configure interface to XP switch	3
eBGP to each peer	3
Database of peers	3
Conclusion	4

Overview

This session covers the steps that you will need to perform in order to prepare for your real IXP installation.

1. Each ISP gets an ASN from AfriNIC

Although it may be technically possible to omit this step (and use private ASN numbers or something else entirely), we strongly recommend that participants at an exchange point get their own globally routable Autonomous System Number ("AS number" or "ASN").

AfriNIC is the Regional Internet Registry (RIR) responsible for the distribution of Internet resources in Africa and Indian Ocean service region. Most organisations request these resources, i.e. IP addresses and Autonomous System Numbers (ASN), directly from their upstream service provider. However there are circumstances when this is not appropriate, such as when a national IXP is being established.

When this situation occurs, an organisation must initially apply to become a member of AfriNIC. There are three categories of membership available:

- Member only (this is only for those who want to participate and vote on policy making, not for those who simply require resources, and as such would not be applicable in this instance).
- Local Internet Registry (LIR status is for those ISPs that are using more than /24 addresses and are growing and needing to allocate addresses to their customers).
- End User (this is if you do not yet fit into the LIR category and just want an ASN and also the fastest route to begin with).

After you decide what type of member you wish to be you will need to complete three forms. Keep in mind that the forms are read and processed by Afrinic support staff (that is, by people, not by a

computer) and you need not know every answer to every question. You must complete these forms to the best of your ability.

Download and complete the following:

1. The membership application form:

www.afrinic.net/forms/affrm-ma200501.pdf

2. The Registry Services Agreement ("RSA") form

www.afrinic.net/corporate/aflgl-rsa200407.pdf

Then fax the completed forms to AfriNIC in Mauritius at fax number: +230 466 6758.

If you don't hear from them in a week try new-member@afrinic.net.

3. You should then have an invoice. Once you have paid, complete the ASN form:

www.afrinic.net/forms/affrm-asn200501.htm

and email it to hostmaster@afrinic.net

Note that you should complete the AfriNIC membership process (including payment) **before** you submit the request for an AS number.

Fees

The setup fee for an ASN is US\$400 once off. The annual recurring fee is a membership fee and is dependent on the type of membership. For end users it will be US\$50 and for LIRs it depends on the size of address allocation.

2. The IXP operator obtains IP addresses

The IXP operator must go through the same process as each ISP, i.e. apply for AFRINIC membership and sign the RSA agreement.

Once membership has been approved then the IPv4 First Allocation Request Form must be completed:

www.afrinic.net/forms/affrm-v4fst200501.htm

and sent by email to hostmaster@afrinic.net

Due to the high cost of a micro allocation from AfriNIC (US\$2500), AFIX is exploring the services provided by www.xp.net specifically for IXPs.

3. Each ISP installs a router at the IXP

Once the IXP operator has obtained the address space to be used at the IXP, each ISP that connects at the IXP must be allocated an IP address by the IXP operator for each router that will be connected to the IXP.

Each ISP must procure a router and a line to connect their headquarters to the IXP.

4. Each ISP must establish its interior routing

Each ISP must use appropriate interior routing mechanisms to connect its exchange point router to the rest of its network. This will probably involve OSPF (or some other interior routing protocol), or static routes.

In almost all cases, each ISP will also establish iBGP neighbour relationships between its exchange point router and other routers in its network. The procedure for deciding which routers need to use iBGP is:

- Identify all routers in your network that will use eBGP (for example, to speak to your upstream provider and to speak to your peers at an exchange point). Also identify all routers that have non-BGP routes pointing to destinations outside your control (such as a default route to your upstream

provider). In the simplest likely case, there will be two such routers: a router at your headquarters ("HQ") that connects to your upstream provider, and the new router at the exchange point ("XP").

- Identify all routers in your network that could possibly be involved in forwarding packets from one of the routers identified in (1) above to another such router. This must also include routers on backup paths that might be used only if a primary network link fails. In the simplest case, there might not be any such routers (for example, if there is a leased line directly between your HQ router and your XP router).
- All the routers identified in (1) and (2) above must use iBGP.

Unless you are using iBGP route reflectors (which is outside the scope of this workshop), every router that uses iBGP must speak iBGP to every other such router. This is called "full mesh" iBGP. Each router that is part of the iBGP mesh will have a loopback interface, and will have an iBGP neighbour relationship configured between its own loopback interface and the loopback interfaces of every other router that is also part of the iBGP mesh. In the simplest case of only two routers, they simply need to speak iBGP to each other.

Configure iBGP mesh

- Allocate loopback addresses (central database).
- Configure loopback interfaces (on each router).
- Configure BGP AS number.
- Use BGP "neighbor" commands for full mesh iBGP.

Import internal routes to iBGP

- In a small network, it's reasonable to do this only on your HQ router.
- Use a BGP "network" command for each prefix.
- If prefix is an aggregate, use a static pullup route ("ip route <network> <netmask> Null0 254").

Configure interface to IXP switch

- Address from IXP operator.
- Get an IP address from IXP operator.
- Configure IP address on exchange point interface.

5. Each ISP must establish its external routing

Configure eBGP to each peer as per the exercises, ensuring that you remember each of the following:

- Filters
- "Neighbour" statements
- eBGP neighbours usually use the IP address associated with the physical exchange point interface (e.g. ethernet), not loopback interfaces.

6. Database of peers

The IXP operator and each ISP that connects at the IXP should maintain a database of all the peers that includes the following information:

- Name
- AS No.
- Router IP address
- Contact name, phone, email
- After hours procedures

- List of network/netmask

Conclusion

If there are any problems in the construction of your IXP please consult with AFIX or use the AfriNIC discussion list. Please refer to either web site afix.afrispa.org or afrinic.net for further information.