

Module 19 – Internet Exchange Points

Objective: To investigate methods for connecting to an Internet Exchange Point.

Prerequisites: Modules 12, 13 and 18, and the Exchange Points Presentation

The following will be the common topology used.




Figure 1 – IXP Configuration

Monday, April 30, 2001

Lab Notes

The purpose of this module is to investigate more deeply the operation of Internet Exchange Points, how to peer at IXPs, and look at some of the recommended configuration practices.

Lab Exercise

- 1. Basic Configuration.** Each router team should configure their router to fit into the network layout depicted in Figure 1. Check all connections. Note that all links are by ethernet.
- 2. Addressing Plan.** These address ranges should be used throughout this module. You are welcome to use your own range within an AS if you desire, just so long as you consult with the teams in other ASes to ensure there is no overlap.

AS100	218.11.0.0/19	AS107	220.10.0.0/19
AS101	218.35.0.0/19	AS108	220.19.0.0/19
AS102	218.76.0.0/19	AS109	220.73.0.0/19
AS103	219.13.0.0/19	AS110	221.19.0.0/19
AS104	219.58.0.0/19	AS111	221.35.0.0/19
AS105	219.64.0.0/19	AS112	221.99.0.0/19
AS106	219.99.0.0/19	AS113	222.11.0.0/19

- 3. Basic Router Setup.** Set up the routers as you would have done in previous modules. That is, basic security, the BGP outline configuration, IOS Essentials, etc. The lab instructor will have connected another router to the exchange point – this is Router15 in the figure. You should set up your routers to synchronise time off that router using NTP. The address range used for the IXP is 220.5.10.0/24 – the management router in this module has an IP address of 220.5.10.254. Each of the ASes is assigned a block of 3 addresses to use on the exchange point LAN. So, for example, AS100 has 220.5.10.1, 220.5.10.2 and 220.5.10.3. AS101 has 220.5.10.4, 220.5.10.5 and 220.5.10.6. And so on.

Checkpoint #1: When you have properly configured your router, and the other routers at the IXP are reachable (i.e. you can ping the other routers), please let the instructor know.

Scenario One – Simple IXP example

The first example is that of very simple IXP. The example only uses prefix lists to configure filtering. eBGP peers should be in peer-groups, soft reconfiguration should be used as in other modules, and

Unicast Reverse Path Forwarding checks should be enabled on the ethernet interface pointing to the IXP.

- 4. Configure the ethernet of each router at the IXP.** The ethernet interfaces connected to the IXP should be configured appropriately for a public connection. Review the IOS Essentials materials and the IXP presentation. The configuration for Router 14 might be:

```
interface ethernet 0
  description Exchange Point LAN
  ip address 220.5.10.40 255.255.255.0
  ip verify unicast reverse-path
  no ip directed-broadcast
  no ip proxy-arp
  no ip redirects
!
```

If you are unclear as to what any of the configuration lines do, please ask the lab instructor.

- 5. Configuring BGP on the routers.** Next, eBGP needs to be set up on the routers. Create a peer-group and apply that peer-group to each eBGP neighbour. A sample configuration for Router13 might be:

```
ip prefix-list myprefixes permit 221.99.0.0/19
ip prefix-list peer100 permit 218.11.0.0/19
..
ip prefix-list peer113 permit 222.11.0.0/19
!
router bgp 112
  network 221.99.0.0 mask 255.255.224.0
  neighbor ixp-peers peer-group
  neighbor ixp-peers soft-reconfiguration in
  neighbor ixp-peers prefix-list myprefixes out
  neighbor <router1> remote-as 100
  neighbor <router1> description Peering with AS100
  neighbor <router1> peer-group ixp-peers
  neighbor <router1> prefix-list peer100 in
..
  neighbor <router14> remote-as 113
  neighbor <router14> description Peering with AS113
  neighbor <router14> peer-group ixp-peers
  neighbor <router14> prefix-list peer113 in
!
```

Monday, April 30, 2001

The configurations for the other routers will be similar to this one.

Note the prefix-lists. There is a per peer inbound prefix-list. Some service providers only filter ASes – that has inherent dangers, and does not prevent against inbound leaking of prefixes incorrectly originated by the peer AS. But only filtering on prefixes doesn't scale, especially in larger IXPs with large participating service providers as they are frequently adding to the prefixes they announce. The Internet Routing Registry is usually used to solve this problem.

- 6. Connectivity Test.** Check connectivity throughout the IXP network. Each router team should be able to see all the other routers at the IXP. When you are satisfied that BGP is working correctly, try running traceroutes to check the paths being followed.

Checkpoint #2: *Once the BGP configuration has been completed, check the routing table and ensure that you have complete reachability over the entire network. If there are any problems, work with the other router teams to resolve those.*




Figure 2 – More Complex IXP

Scenario Two – More complex IXP

This situation above was a simple example of how to configure a basic IXP. However, this doesn't really represent reality in the majority of cases. This next example introduces the situation where an IXP participant provides transit beyond their backbone AS. In this example AS103 and AS107 provide transit to AS200/201 and AS300/301 respectively.

Monday, April 30, 2001

- 7. Reconfigure the network.** All routers should remove any prefix-list filtering and eBGP peering which has been configured. Routers 1 to 8 can retain the remaining configuration, address block, etc. Routers 9 to 14 need to be reconfigured into the new ASes as per the figure. The address ranges are as previously with modifications to cater for AS200, AS201, AS300 and AS301.

AS100	218.11.0.0/19	AS107	220.10.0.0/19
AS101	218.35.0.0/19	AS200	220.19.0.0/19
AS102	218.76.0.0/19	AS201	220.73.0.0/19
AS103	219.13.0.0/19	AS300	221.19.0.0/19
AS104	219.58.0.0/19	AS301	221.35.0.0/19
AS105	219.64.0.0/19		
AS106	219.99.0.0/19		

- 8. Configure the ethernet of each router at the IXP.** The ethernet interfaces connected to the IXP should be configured appropriately for a public connection. This step is similar to step 4 earlier, but applies only to Routers 1 to 8.
- 9. Configuring the edge ASes.** Next the routers in AS200, AS201, AS300 and AS301 should set up eBGP peering with the routers in AS103 and 107 respectively. This has been done several times in previous modules, so no examples given. **HINT:** AS103 and AS107 announce default, only accept customer prefixes, AS200/1, AS300/1 accept default, only announce their prefix.
- 10. Configuring AS103 and AS107.** The routers in AS103 and AS107 should set up OSPF and iBGP between each other. Again, no examples given as this has been covered several times in previous modules.
- 11. Configuring eBGP across the IXP.** As in the previous example, create a peer-group and apply that peer-group to each eBGP neighbour. However, notice that AS103 and AS107 are now providing transit to the exchange point for their customer ASes. The configuration we will use for Routers 4 and 8 is slightly different – these ASes control their announcements using an AS path filter, rather than a prefix list. A sample configuration for Router8 might be:

```
ip prefix-list myprefixes permit 220.10.0.0/19
ip prefix-list peer100 permit 218.11.0.0/19
ip prefix-list peer101 permit 218.35.0.0/19
ip prefix-list peer102 permit 218.76.0.0/19
! peer103 has special inbound policy
ip prefix-list peer104 permit 219.58.0.0/19
ip prefix-list peer105 permit 219.64.0.0/19
ip prefix-list peer106 permit 219.99.0.0/19
```

```

! peer107 has special inbound policy
!
ip as-path access-list 10 permit ^$
ip as-path access-list 10 permit ^300$
ip as-path access-list 10 permit ^301$
!
ip as-path access-list 20 permit ^103$
ip as-path access-list 20 permit ^103_200$
ip as-path access-list 20 permit ^103_201$
!
router bgp 107
 network 220.10.0.0 mask 255.255.224.0
 neighbor ixp-peers peer-group
 neighbor ixp-peers soft-reconfiguration in
 neighbor ixp-peers prefix-list rfc1918-dsua out
 neighbor ixp-peers filter-list 10 out
 neighbor <router1> remote-as 100
 neighbor <router1> description Peering with AS100
 neighbor <router1> peer-group ixp-peers
 neighbor <router1> prefix-list peer100 in
 ..
 neighbor <router4> remote-as 103
 neighbor <router4> description Peering with AS103
 neighbor <router4> peer-group ixp-peers
 neighbor <router4> prefix-list rfc1918-dsua in           ! NOTE THIS
 neighbor <router4> filter-list 20 in
 ..
 neighbor <router7> remote-as 106
 neighbor <router7> description Peering with AS106
 neighbor <router7> peer-group ixp-peers
 neighbor <router7> prefix-list peer106 in
!

```

The configurations for the other routers will be similar to this one.

Note the different configuration for the AS103 peering. Here we are allowing AS103 to send any prefixes, but filtering specifically on AS path. Be aware of the dangers of doing this – AS103 could send prefixes which it isn't meant to – an extra precaution, with extra administrative difficulty, would be to filter on prefixes also.

The comments about the prefix-lists in the earlier example apply to this one too. This is why it is common to find ISPs using the Internet Routing Registry at IXPs – managing the peering at this level can get quite sophisticated, with prefix-lists and filter-lists being rebuilt on a nightly basis.

Monday, April 30, 2001

12. Connectivity Test. Check connectivity throughout the IXP network. Each router team should be able to see all the other routers at the IXP. When you are satisfied that BGP is working correctly, try running traceroutes to check the paths being followed.

13. Checkpoint #2: *Once the BGP configuration has been completed, check the routing table and ensure that you have complete reachability over the entire network. If there are any problems, work with the other router teams to resolve those.*

14. Summary. This module has given examples of configurations used by Internet Service Providers at Internet Exchange Points. They have concentrated on using prefix-lists and as-path filters – more sophisticated configurations are possible by using communities. These examples are left to the reader to consider. If there is time at the end of the workshop, ask the Instructor to test out some other scenarios.

CONFIGURATION NOTES

Documentation is critical! You should record the configuration at each *Checkpoint*, as well as the configuration at the end of the module.