IPv6 Internet Exchange Topology

In this exercise we simulate the construction of an IX.

An IX provides a peering exchange connecting ISPs and other providers together in single common subnet in order to conveniently route traffic between themselves by mutual agreement.

Each IX member network or AS maintains a presence in the IX through a connected router, which acts as a border router to the other networks (ASs) in the IX. The routers are commonly connected to a switch on a common vlan providing them with the peering connectivity.

In this exercise you will:

- Determine the IP addressing scheme for the IX and for your ISP LAN network
- Configure the external interfaces of the routers connecting your ISP to the IX
- Configure an internal LAN for your ISP Configure static routing
- Configure BGP on the routers
- Test this connectivity.

IXP network topology
IX Subnet: 2001:DF0:AA::/48

ISP Allocated blocks

Router 1: 2001:ABC1::/32	Router 2: 2001:ABC2::/32
Router 3: 2001:ABC3::/32	Router 4: 2001:ABC4::/32
Router 5: 2001:ABC5::/32	Router 6: 2001:ABC6::/32
Router 7: 2001:ABC7::/32	Router 8: 2001:ABC8::/32

Allocated AS Numbers

Router 1: ASN 1	Router 1: ASN 2
Router 3: ASN 3	Router 1: ASN 4
Router 5: ASN 5	Router 1: ASN 6
Router 7: ASN 7	Router 1: ASN 8

Exercise 1.1

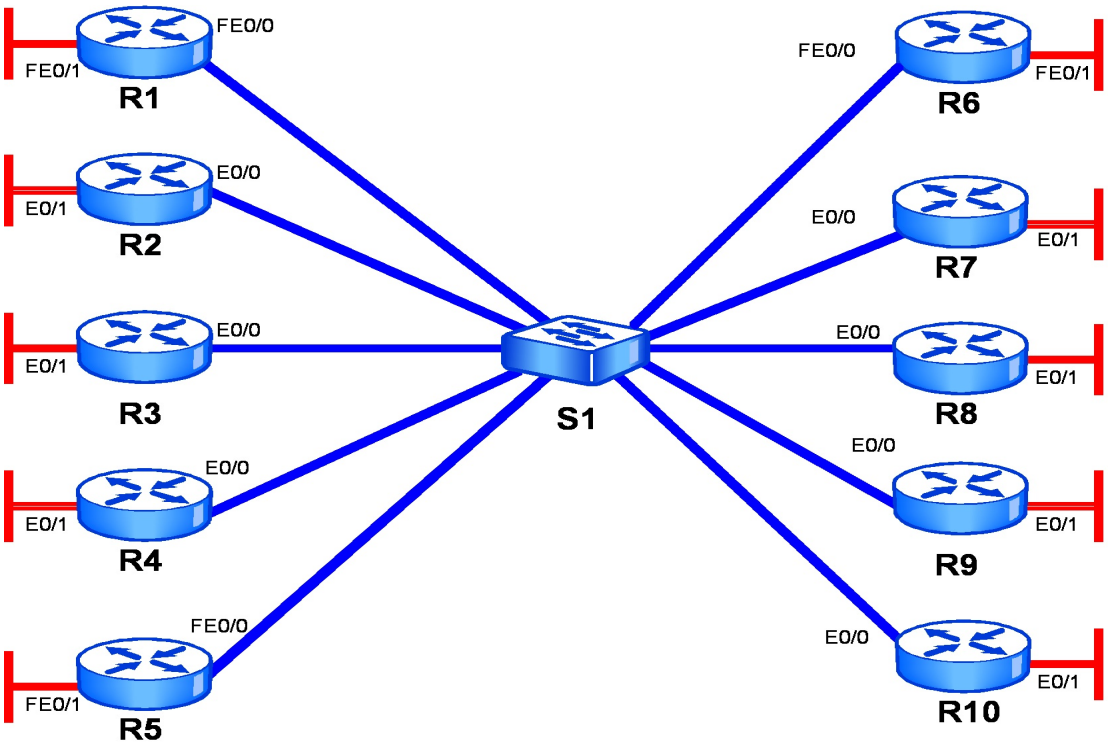

Referring to the diagram above and using the table given below, configure the interface that connects your router to the IX using the appropriate /64 address. These addresses are taken from a /64 taken from the /48 assigned to the IX (2001::0DF0:00AA::/48)

NOTE: All routers interface connecting to the IX should be in the same subnet.

Router 1: 2001:DF0:AA::1/64	Router 2: 2001:DF0:AA::2/64
Router 3: 2001:DF0:AA::3/64	Router 4: 2001:DF0:AA::4/64
Router 5: 2001:DF0:AA::5/64	Router 6: 2001:DF0:AA::6/64
Router 7: 2001:DF0:AA::7/64	Router 8: 2001:DF0:AA::8/64

Exercise 1. 2

Each of the ISPs have been allocated a /32 address. From this block create a /48 subnet for an internal LAN segment of the network. (i.e. assignment to a customer)

Router 1	2001:ABC1:0001::/48	Router 2	2001:ABC2:0002::/48
Router 3	2001:ABC3:0003::/48	Router 4	2001:ABC4:0004::/48
Router 5	2001:ABC5:0005::/48	Router 6	2001:ABC6:0006::/48
Router 7	2001:ABC7:0007::/48	Router 8	2001:ABC8:0008::/48

Exercise 1.3
From the /48 address use the first /64 for the point-to-point connection that will be used to connect the customer LAN. Select the "internal" facing interface (E0/1 or F0/1, refer to the diagram) and configure this interface with the first IP address from that /64.

NOTE: Remember, the E0/0 or F0/0 interfaces will connect to the IX, while the E0/1 or F0/1 interface will be LAN interface of the ISP for its customers.

Router 1 - 2001:ABC1:0001::1/64 **Router 2 - 2001:ABC2:0002::1/64**

Exercise 1.4
Create a static route to every member's network in the IX. This would mean each router would have 6 static route entries to reach destination network if there are 7 routers connected to the IX.

Formula: N-1 = number of static routes needed ("N" is number of routers)

Note: you will create a route to the /32 allocated to that ISP, using the address of the /32 given in the table at the start of this section (p16). As the next hop you will need to use the IP address configured on the IX facing interface of the router for that ISP

Configuring IPv6 address on Cisco routers

Configuring the router interface with an IPv6 address.

The router has several different interfaces, which need to be configured. Assume we are configuring the f0/0 (Fastethernet 0/0)

1. Select the appropriate interface required prior to configuration.

2. Once selected ensure that the interface is up.

   Issue the command "show interface"

   router # show interface f0/0

   Enter configuration mode in order to configure the interface.

   router # **configure terminal**

   Select the interface (interface mode)

   routers(config) # **interface f0/0**
   routers(config-if) #

3. Configure the IP address for the interface selected. :

   f0/0 configuration with a /64 subnet for the interface connected to the IX network

   routers(config-if)#**ipv6 address 2001:DF0:AA::1/64**

   f0/1 configuration with a /64 subnet for the ISP's router interface towards its customer
   (Note: you will of course need to enter interface mode for this interface)

routers(config-if) # **ipv6 address 2001:ABC1:1::1/64**

4. Enable IPv6 on the interface selected.

   routers(config-if) # **ipv6 enable**

5. Exit from the interface configuration and enable IPv6 unicast datagram forwarding by typing the command below in the global mode.

   routers(config) # **ipv6 unicast-routing**

6. To create a static route to a destination network, define the network prefix (ID) and the next hop to be used to reach that network. In our IX example, we define a route to the /32 of each ISP using the IX facing address of the ISP router that connects to the IX as the next hop (or the outbound interface)

   routers(config)#ipv6 route 2001:ABC2::/32 FE 0/0 (outbound interface)

   Or

   routers(config)#ipv6 route 2001:ABC2::/32 2001:DF0:AA::2 (next-hop IP)

7. Save the configuration by issuing the command: "**copy run start**"

8. For testing ping a valid host address of the destination network. Use the IP address configured on the customer-facing interface of the router.

   Example
   routers#ping 2001:DF0:AA::1
   (for interface connected to the IX)

   routers#ping 2001:ABC2:2::1
   (for routers2 interface connected to its customers)

STOP HERE MAKE SURE THAT YOU CAN REACH OTHER ROUTER'S NETWORKS

Exercise 1.5
After testing that the static route configuration is working, we are now ready to configure IXP network to run You will need to configure BGP on each member routers and establish BGP peering between each of the members.

NOTE: Remove the entire static configuration first.

routers(config)#no ipv6 route 2001:ABC1::/32

1. To configure BGP we need to enter “BGP mode” and also define the network prefix (id) that we want to announce. Type “router bgp” with the AS number in the command prompt of the router global.

   routers#configure terminal
   routers(config)#router bgp <ASN>
   routers(config-router)#no auto summary
   routers(config-router)#no synchronization

   Where the AS number is the number of your router

   Note: If the AS that performs BGP does not own the complete classfull network, Cisco recommends that you disable auto-summary using the **no auto-summary** command under **router bgp**. (Reference: http://www.cisco.com/warp/public/459/bgpfaq_5816.shtml#five**Note:** If BGP synchronization is enabled, there must be a match for the prefix in the IP routing table in order for an internal BGP (iBGP) path to be considered a valid path. BGP synchronization is enabled by default in Cisco IOS® Software. If the matching route is learned from an Open Shortest Path First (OSPF) neighbor, its OSPF router ID must match the BGP router ID of the iBGP neighbor. Most users prefer to disable synchronization with use of the no synchronization BGP subcommand. (Reference: http://www.cisco.com/warp/public/459/25.shtml)

2. Configure the peering address of the neighboring AS. Use the point-to-point interface IP address for each router connected to the IX. Route announcements and updates are only send to neighbours.

   NOTE: Each router will have (n-1) neighbours.

   routers(config-router)# neighbor <other ASN interface IP>
   remote-as <other ASN>

   Example: Complete configuration command to configure BGP.

   routers#configure terminal
   routers(config)#router bgp 1
   routers(config-router)#no auto-summary
   routers(config-router)#no synchronization

   routers(config-router)#neighbor 2001:DF0:AA::2 remote-as 2
   (for peering with routers2)

3. BGP uses a router ID to identify BGP-speaking peers. The BGP router ID is 32- bit value that is often represented by an IPv4 address. By default, the Cisco IOS software sets the router ID to the IPv4 address of a loopback interface on the router. If no loopback interface is configured on the router, then the software chooses the highest IPv4 address configured to a physical interface on the router to represent the BGP router ID. When configuring BGP on a router that is enabled only for IPv6 (the router does not have an IPv4 address), you must manually configure the BGP router ID for the router. The BGP router ID, which is represented as a 32-bit value using an IPv4 address syntax, must be unique to the BGP peers of the router.

   http://www.cisco.com/en/US/docs/ios/12_2t/ipv6/SA_bgpv6_ps6350_TSD_Products_Configuration_Guide_Chapter.html)

   If an ID cannot be configure the following error message will be displayed:

   % BGP cannot run because the routers-id is not configured BGP router identifier 0.0.0.0, local AS number 1

   In such a case you will need to configure the bgp router-id

   routers(config-router)#bgp router-id [(32 bit number in ipv4address format]

   Example:

   Router1#configure terminal
   Router1(config)#router bgp 1
   Router1(config-router)#bgp router-id 10.1.1.1

4. From the BGP prompt you need to use the ipv6 address family command, to enter address-family mode in order to configure the peering for each neighbor. This will define the address mode (IPv6) for announcements and then the neighbours who will included the announcements of address updates.

   routers(config-router)# address-family ipv6

   Then activate the peering to each neighbor

   routers(config-router-af)# neighbor (neighbor IPv6 address) activate

   routers(config-router-af)# neighbor 2001::DF0:AA::2 activate

5. Still in the address-family mode, enter the network statement for the networks that will be announced from the AS.

   Example:

   routers(config-router-af)# network ipv6network/prefix size)

   routers(config-router-af)#network 2001:ABC1::/32

6. Exit from the BGP mode and then create a “pull up route” for the /32 ip address block by simply creating a static route to null 0 interface.

   NOTE: For BGP to announce the /32 block it needs to see a valid IGP route match.

   Example:

   routers(config)# ipv6 route ipv6network/prefix size) null0 (metric)

   routers(config)# ipv6 route 2001:ABC1::/32 null0 254

7. Verify that the configuration is working properly by entering the following commands:

   - **show bgp ipv6 unicast summary** (to check the bgp summary table)
   - **sh bgp ipv6 unicast neighbour** (to check neighbour list)
   - **sh bgp ipv6** (to check the routing table for the BGP announcements)
   - **sh ipv6 route** (to check the IPv6 routing table)

STOP HERE MAKE SURE THAT YOU CAN REACH OTHER ROUTER’S NETWORK