IPv6 Module 11 – Advanced Router Configuration

Objective: Create a basic physical lab interconnection with two autonomous systems. Each AS should use OSPF, iBGP and eBGP appropriately to construct a working network.

Prerequisites: Basic ISP Workshop (at least Modules 1 and 6 to 8)

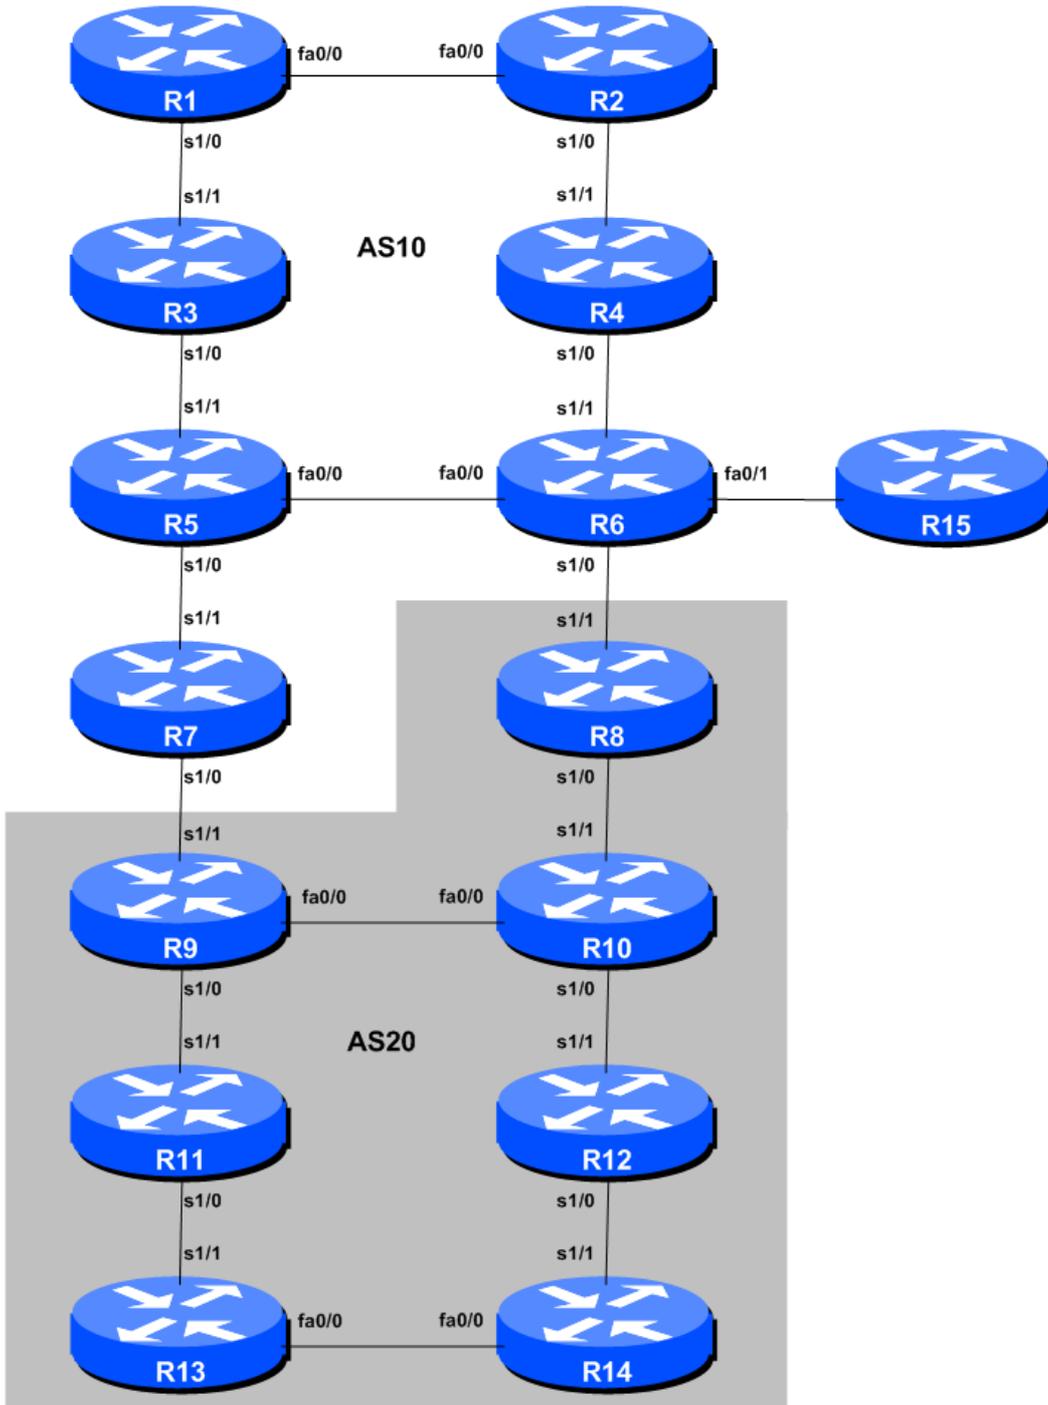The following will be the common topology used.

Figure 1 – ISP Lab Basic Configuration

Tuesday, July 29, 2014

Lab Notes

The purpose of this module is to construct the workshop lab and serve as a reminder of the basic principles of building a network, introducing an IGP, properly function iBGP, and the basics of eBGP:

- After the **physical design** is established, the connections between the hardware should be built and verified.

- Next, the routers should have the **base configuration** installed, and basic but sufficient security should be set up. Note that Router15 is the Workshop Instructor's router and it will be used at various instances throughout the workshop.

- Next the **basic IP connectivity** be tested and proven. This means assigning IP addresses on all links which are to be used, and testing the links to the neighbouring devices.

- Only once one router can see its neighbour does it make sense to start configuring routing protocols. And **start with the IGP** (OSPF is chosen for this workshop). There is no purpose to building BGP while the chosen IGP (in this case OSPF) is not functioning properly. BGP relies on OSPF to find its neighbours and next hops, and an improperly or non-functioning OSPF will result in much time wasted attempting to debug routing problems.

- Once the IGP is functioning properly, the **BGP configuration** can be started, first internal BGP, then external BGP.

- Finally, **documentation**. Documentation is often overlooked or forgotten. It is an ongoing process in this workshop. If the instructor asks you to document something, either on the whiteboard in the class, or at the back of this booklet, it is in your best interests to do so. There can never be too much documentation, and documentation at the time of network design and construction can usually saves much frustration at a future date or event.

Lab Exercise

The following list is typical for what needs to be done to bring up the lab configuration:

1. **Enable IPv6.** Cisco routers with an IOS supporting IPv6 currently do not ship with IPv6 enabled by default. This needs to be done before any of the following exercises can be completed. To do this, use the following command:

    ```
    ipv6 unicast-routing
    ```

2. **Enable IPv6 CEF.** Unlike IPv4, CEFv6 is not enabled by default. So we now need to enable IPv6 CEF also, using the following command:

    ```
    ipv6 cef
    ```

3. **Disable IPv6 Source Routing.** Unless you really believe there is a need for it, source routing should be disabled. This option, enabled by default, allows the router to process packets with source routing header options. This feature is a well-known security risk as it allows remote sites to send packets with different source address through the network (this was useful for

troubleshooting networks from different locations on the Internet, but in recent years has been widely abused for miscreant activities on the Internet).

```
no ipv6 source-route
```

4. **IPv6 Addressing Plans.** Addressing plans in IPv6 are somewhat different from what has been considered the norm for IPv4. The IPv4 system is based around the RIRs allocating address space to an LIR (an ISP who is a member of the RIR) based on the needs of the ISP; that allocation is intended to be sufficient for a year of operation without returning to the RIR. The ISP is expected to implement a similar process towards their customers – so assigning address space according to the needs of the customer.

The system changes a little for IPv6. While the RIRs still allocate address space to their membership according to their membership needs, the justification to receive an IPv6 allocation is somewhat lighter than it is for IPv4. If the ISP can demonstrate a plan to connect at least 200 customers to the Internet using IPv6, they will receive an allocation. However, a bigger advantage starts with the customer assignments made by the ISP – the ISP simply has to assign a /48 to each of their customers. This is the minimum assignment for any site/customer – within this /48 there are a possible 64k subnets, deemed sufficient for all but the largest networks around these days. Within this /48, the smallest unit which can be assigned is a /64 – so every LAN and point to point link receives a /64.

With this revised system, the address plan for IPv6 is hugely simplified. ISPs assign a single /48 for their network infrastructure, and the remainder of their /32 block is used for customer assignments. This workshop assumes this principle, as will be seen in the following steps.

5. **IPv6 Addresses.** Each AS is assigned a block of IPv6 addresses.

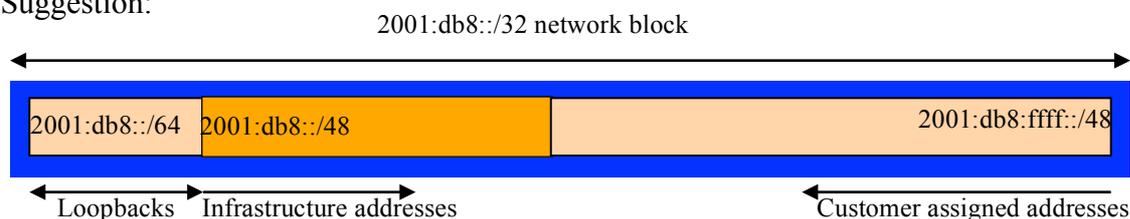**AS10   2001:db8::/32   AS20   2001:db9::/32**

Decide among your team what the addressing plan for you AS should be.
Hint One: point-to-point links only require /64s and should be addressed as /127s.
Hint Two: loopbacks only require a /128 host address.
Hint Three: number your backbone sequentially, preferably from the start of the range, using the first /48 for your infrastructure.

Suggestion:

2001:db8::/32 network block

2001:db8::/64   2001:db8::/48   2001:db8:ffff::/48

Loopbacks   Infrastructure addresses   Customer assigned addresses

Note: When the IPv6 addresses are assigned, they **MUST** be annotated on the **WHITE-BOARD** at the front of the workshop room. A large network map will have been drawn on the white-board – all the IP address assignments need to be annotated there so that other Router Teams can document and understand the links and routing in this and future modules.

6. **Ping Test #1.** Ping all physically connected subnets of the neighbouring routers. If the physically connected subnets are unreachable, consult with your neighbouring teams as to what might be wrong. Don't ignore the problem – it may not go away.

7. **Create vty filters.** Set up filters on the vty interfaces restricting vty access to your router to those addresses you would like to permit. For the purpose of this lab, even though it is not connected to the Internet, set up filters so that only the address space in the lab has access.

```
ipv6 access-list ipv6-vty
 remark VTY filters
 permit ipv6 2001:db8::/32 any
 permit ipv6 2001:db9::/32 any
 deny   ipv6 any any log
!
line vty 0 4
 ipv6 access-class ipv6-vty in
```

8. **OSPF within the same AS.** Each router Team should enable OSPF on their router. The OSPF process identifier should be the same as your AS number. All the routers in one AS will be in OSPF Area 0. Those interfaces over which we expect to form OSPF adjacencies should be marked as active interfaces. Don't forget to mark the loopback interface as a passive interface – for example:

```
ipv6 router ospf 41
 passive-interface default
 no passive-interface serial 1/0
 log-adjacency-changes
!
interface loopback 0
 ipv6 ospf 41 area 0
!
interface serial 1/0
 ipv6 ospf 41 area 0
!
```

Note that the loopback interface has no "network" attached to it, so there is no way it can be connected to another device in such a way as to form an OSPF adjacency.

9. **DMZ between AS10 and AS20.** OSPF must **NOT** run on the demarcation links between AS10 and AS20. So Routers 6, 7, 8 and 9 must **NOT** configure the serial links between each other to be non-passive. This is a very important point, and a mistake frequently made by many ISPs. Also, do not put a network statement on external facing interfaces – again see iBGP discussion at step 14. Router 6 should have two adjacencies only – with Router 4 and 5. Router 7 should have one adjacency only – with Router 9. And so on.

10. **Intra Area Authentication.** OSPF supports router authentication within areas. This is quite important inside ISP networks to prevent the introduction of improperly configured or unintended equipment.

    Each team will turn on authentication for area 0 using the *area N authentication* command. An example configuration for Router6 might be:

    ```
    ipv6 router ospf 41
     area 0 authentication ipsec spi 256 md5 0123456789ABCDEF0123456789ABCDEF
    ```

```
    !
```

Notice now that the OSPF adjacencies do not come up unless the neighbouring router has also entered the same configuration and key. Notice also how the OSPF adjacencies timed out after the configuration was entered – the adjacent router still had not set up OSPF neighbour authentication. Once they also have done this, the adjacency is re-established.

Also, if we want to configure neighbour authentication per area but disable it on a per link basis, we need to use the `authentication null` option. The following example shows how to disable neighbour authentication on one specific interface, hssi 5/0.

```
interface hssi 5/0
 ipv6 ospf authentication null
 ipv6 ospf 1 area 10
```

Note: Wherever an OSPF session is configured from now on in the workshop, all Router Teams MUST use passwords on these OSPF sessions.

11. **Ping Test #2.** Ping all loopback interfaces in your AS. They should all respond. This will ensure the OSPF IGP is connected End-to-End. If there are problems, use the following commands to help determine the problem:

```
show ipv6 route           : see if there is a route for the intended destination
show ipv6 ospf            : see general OSPF information
show ipv6 ospf interface  : Check if OSPF is enabled on all intended interface
show ipv6 ospf neighbor   : see a list of OSPF neighbours that the router sees
```

Checkpoint #2: *call lab assistant to verify the connectivity. Save the configuration as it is on the router either on the worksheet on the end of this hand out, or own your own laptop, or on the classroom tftp server if it is available.*

12. **BGP Distance.** Before we set up iBGP with our neighbours in our AS, we need to do some basic preparation on the router. The IOS defaults are not optimised for Service Provider networks, so before we bring up BGP sessions, we should set the defaults that we require.

The default distance for eBGP is 20, the default distance for iBGP is 200, and the default distance for ISIS is 115. This means that there is a potential for a prefix learned by eBGP to override the identical prefix carried by ISIS. Recall from the Routing presentation that there is a distinct separation between BGP and ISIS processes – prefixes present in ISIS will never be found in BGP, and vice-versa. To protect against accidents[1], the eBGP distance is set to 200 also. The command to do this is the bgp distance subcommand, syntax is:

```
distance bgp <external-routes> <internal-routes> <local-routes>
```

Note: This should be included in all future BGP configurations in this workshop. Set the BGP protocol distance so that BGP is always less preferred than any IGP. So:

---

[1]   There have been several incidents in the past where denial of service attacks on ISP networks have been successful because ISPs have omitted basic routing protocol security. Setting the BGP distances to be greater than any IGP is one of the mitigation methods available.

```
router bgp 10
 address-family ipv6
  distance bgp 200 200 200
```

13. **Passwords on BGP sessions.** It is now considered very good practice to use passwords on the BGP sessions on the router. When BGP is set up in the next step, don't forget to include a password on the BGP peering.

    The password used for this module will be *cisco* – obviously on a real operational network operators will use a password which follows their normal password rules, and not something which is easily guessable. An example configuration might be:

    ```
    router bgp 10
     address-family ipv6
      neighbor 2001:db8::2 password cisco
    ```

    Note: Passwords should be included in all future BGP configurations in this workshop.

14. **Configuring next-hop-self on iBGP Neighbours.** So that BGP has a valid next-hop for external destinations, we introduce the `next-hop-self` BGP configuration. This changes the iBGP default by replacing the next-hop address for external sites from that of the external neighbour address to the loopback address of the local router. The local router knows how to get to the external destinations because it is connected to the LAN that leads there – the rest of the network internal to the AS is told simply to go via this router. Note that because we do this, we no longer need to quote the external point to point link in our OSPF (or ISIS) configuration – see steps 7 & 9 earlier. For example:

    ```
    router bgp 10
     address-family ipv6
      neighbor 2001:db8::2 next-hop-self
    ```

    Note that the use of *next-hop-self* on all iBGP sessions is considered industry best practice, and its use from now on in the workshop is **strongly recommended**.

15. **Configuring iBGP Neighbours.** Configure iBGP peers within each autonomous system. Use a full iBGP mesh. Don't forget that iBGP peering is configured to be between the loopback interfaces on the routers. Also, it is good practice to use a peer-group – and because we are using IPv6, put in a "v6" to differentiate it from the equivalent IPv4 peer-group. For example:

    ```
    router bgp 10
     no bgp default ipv4-unicast
     address-family ipv6
      neighbor v6-ibgp-peers peer-group
      neighbor v6-ibgp-peers remote-as 10
      neighbor v6-ibgp-peers description iBGP v6 peergroup
      neighbor v6-ibgp-peers update-source loopback 0
      neighbor v6-ibgp-peers next-hop-self
      neighbor v6-ibgp-peers password cisco
      neighbor v6-ibgp-peers send-community
      neighbor 2001:db8::2 peer-group v6-ibgp-peers
      neighbor 2001:db8::3 peer-group v6-ibgp-peers
      neighbor 2001:db8::4 peer-group v6-ibgp-peers
     ..etc..
    ```

Use *show bgp ipv6 unicast summary* to check the status of the iBGP neighbour connections. If the iBGP session is not up and/or no updates are being sent, work with the Router Team for that neighbour connection to troubleshoot the problem. Note: get into the habit of using peer-groups and configuring them fully, including the "send-community" directive. This workshop makes extensive use of communities, and making them part of your configuration is good practice.

Note: Router6 should also include the network connecting to Router15 in the iBGP configuration. This is so that the network connected to Router15 can be accessed – it has the DNS server and NTP server located on it.

16. **Add Prefixes to BGP.** Each Router Team will advertise the CIDR block assigned to them via BGP. AS10 would advertise 2001:db8::/32 and AS20 would advertise 2001:db9::/32:

```
router bgp 10
 bgp log-neighbor-changes
 !
 address-family ipv6
  no synchronization
  network 2001:db8::/32
 !
ipv6 route 2001:db8::/32 null0
```

Don't forget the static route to Null0. This ensures that the prefix has an entry in the routing table, and therefore will appear in the BGP table. Synchronisation is disabled by default in IOS, but it's still good practice to remember to include the command. (Note that a distance of 250 could be applied to the static route to ensure that routing protocols announcing this exact prefix will override the static (if this is required/desired).)

Checkpoint #3: call the lab assistant to verify the connectivity.

17. **Configure eBGP peering.** Now that iBGP is functioning, it is time to configure eBGP. External BGP will be set up between AS10 and AS20, specifically between Routers 6 and 8, and Routers 7 and 9 only. The remaining lab teams should monitor the BGP table they see on their routers.

Firstly, agree on what IP addresses should be used for the point to point links between the ASes. Put the /64 networks used for the DMZ links into OSPF (network statement and passive interface). Then configure eBGP between the router pairs, for example:

```
router bgp 10
 address-family ipv6
  neighbor 2001:db9:0:1:: remote-as 20
  neighbor 2001:db9:0:1:: password cisco
  neighbor 2001:db9:0:1:: description eBGP with RouterXX
```

Use the BGP show commands to ensure that you are receiving prefixes from your neighbouring AS.

18. **eBGP filters.** The final step of setting up the eBGP sessions is to apply filters on the BGP sessions. The teams operating Routers 6, 7, 8 and 9 should now set up filters to allow their

neighbour AS prefixes into their network; and set up filters to allow their own AS prefixes out of their network. For example:

```
router bgp 10
 address-family ipv6
  neighbor 2001:db8:0:1:: prefix-list AS20v6peer-in in
  neighbor 2001:db8:0:1:: prefix-list AS20v6peer-out out
!
ipv6 prefix-list AS20v6peer-in permit 2001:db9::/32
ipv6 prefix-list AS20v6peer-out permit 2001:db8::/32
!
```

This only allows 2001:db9::/32 in from an AS20 eBGP neighbour, and allows 2001:db8::/32 out to an AS20 eBGP neighbour.

19. **BGP Communities on eBGP sessions.** Network operators only send communities to their eBGP peers if they have a policy requirement to do so. Cisco IOS also does not send communities on the eBGP sessions by default, so we will now add this to eBGP peerings. The teams operating Routers 6, 7, 8 and 9 should add the "send-community" directive to their eBGP peerings. For example:

```
router bgp 10
 address-family ipv6
  neighbor 2001:db8:0:1:: send-community
!
```

20. **Check the network paths and the routing table.** Run traceroutes between your router and other routers in the classroom. Ensure that all routers are reachable. If any are not, work with the other router teams to establish what might be wrong. Make sure that you can see Router15. The lab instructor will have written the addresses and network up on the whiteboard. (The network is 2001:dba::/64, the address of Router6 on that LAN is 2001:dba::254, and the address of Router15 is 2001:dba::1.)

21. **Saving the configuration. For software releases from 12.0 onwards**, the commands to save the configuration are of the format *copy <source> <destination>* where the source and destinations can be any of the following options: *ftp, lex, null, nvram, rcp, running-config, startup-config, system, tftp.* To save the configuration to the TFTP server, use the "*copy system:/running-config tftp:*" command sequence. If the TFTP server is unreachable, "."s followed by an error message will be displayed rather than "!"s. (Note that the "*write net*" command of earlier releases is still supported but may be removed at a future release.)

An example of saving the configuration for Router 1 might be:

```
Router1#copy system:running-config tftp:
Address or name of remote host[]? 192.168.1.4
Destination filename [running-config]? router1-confg
!!
2259 bytes copied in 2.920 secs (1129 bytes/sec)
Router1#
```

Checkpoint #4: call the lab assistant to verify the connectivity.

22. Summary. This module has covered most of the fundamental configuration topics required to add IPv6 routing support to an existing ISP network. It has covered enabling IPv6, OSPF configuration, iBGP configuration, and finally simple eBGP configuration. No routing policy has been implemented. **Each Router team is strongly recommended to make a copy of their configuration as most of the configuration concepts will be required throughout the remainder of the workshop.**