

IPv6 Module 1b – ISIS

Objective: Create a basic physical lab interconnection using IPv6 with one ISIS Area running on top of an existing IPv4 infrastructure.

Prerequisites: IPv4 Lab Module 1, knowledge of Cisco router CLI, and previous hands on experience.

The following will be the common topology used for this supplement.

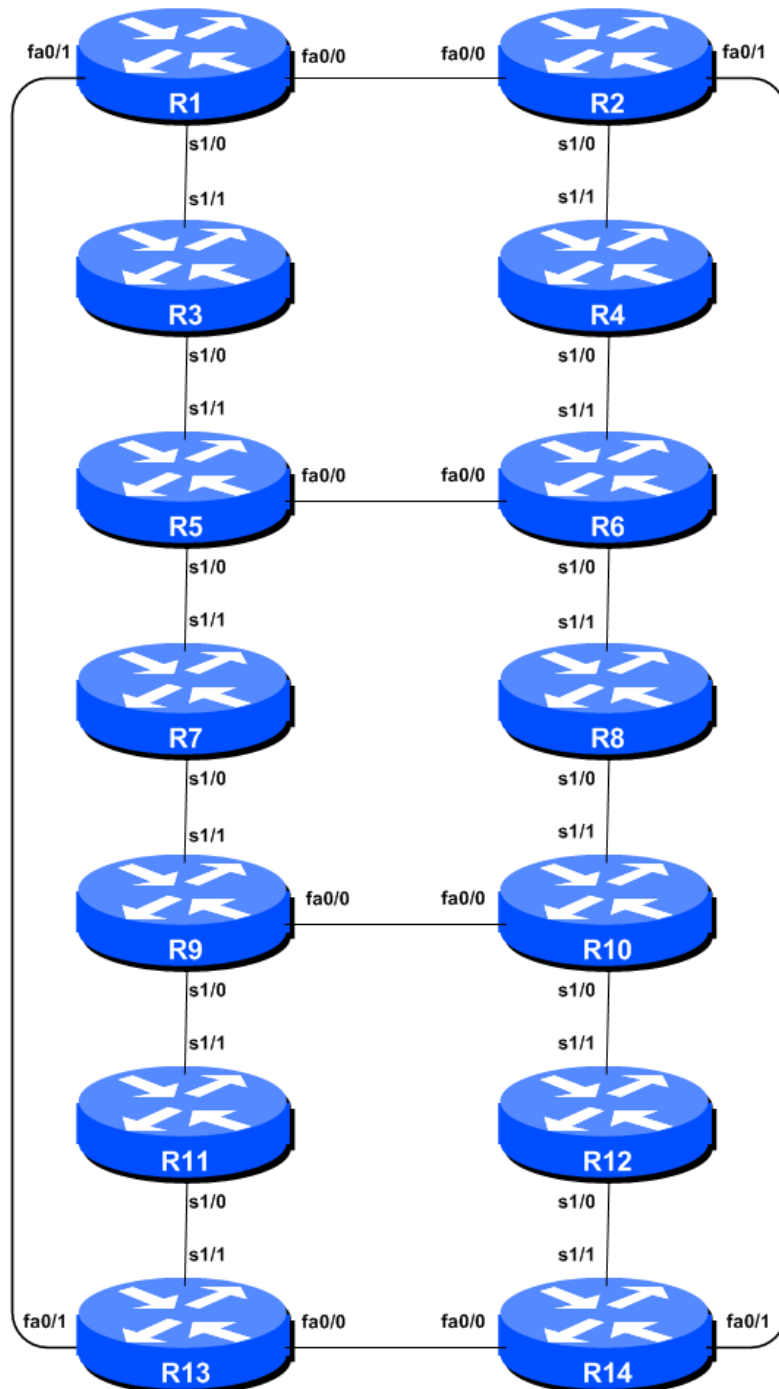


Figure 1 – ISP Lab Basic Configuration

Lab Notes

This Module is intended to supplement Module 1 of the IPv4 Workshop series once that has been completed. The topology and IPv4 configuration should be left exactly as it was at the end of Module 1.

The routers used for this portion of the workshop must support IPv6. This is basically any IP Plus image from 12.2T onwards (IP Plus was renamed to Advanced IP Services for most platforms as from 12.3 mainline). As always, it is best to check the Cisco Feature Navigator www.cisco.com/go/fn to be absolutely sure which images set and platform supports IPv6. Unfortunately IPv6 is not part of the basic IP only or Service Provider IOS images used by most ISPs.

Lab Exercise

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:

```
Router(config)# ipv6 unicast-routing
```

The router is now configured to support IPv6 Unicast (as well as IPv4 Unicast which is the default). Save the configuration.

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:

```
Router(config)# ipv6 cef
```

Nothing will break if IPv6 CEF is not enabled, but more advanced features such as NetFlow will not function without IPv6 CEF being enabled.

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).

```
Router1 (config)# 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. 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 64k possible 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. **Note: This workshop will adopt the recommendations of RFC6164 and use a /127 mask for each point-to-point link – even though the link still has a /64 reserved for it.**

With this revised system, the address plan for IPv6 is greatly 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.** As with the IPv4 portion of this Module we are going to introduce basic concepts of putting together a sensible IPv6 addressing plan for an ISP backbone. The RIRs are typically handing out IPv6 address space in /32 chunks – we assume for the purposes of this lab that our ISP has received a /32. Rather than using public address space, we are going to use 2001:db8::/32, the documentation address for IPv6. In the real world Internet, we would use public address space for our network infrastructure.

The typical way that ISPs split up their allocated address space is to carve it into three pieces. One piece is used for assignments to customers, the second piece is used for infrastructure point-to-point links, and the final piece is used for loopback interface addresses for all their backbone routers. The schematic in

Figure 2 shows what is typically done.




Figure 2 – Dividing allocated block of /32 into Customer, Infrastructure and Loopbacks

Study the address plan which was handed out as an addendum to this workshop module. Notice how the infrastructure addressing uses the first /48 out of the /32 address block. Notice how we have set aside a /64 out of the infrastructure block for the router loopbacks. ISPs tend to document their addressing plans in flat text files or in spreadsheets – Figure 3 below shows an extract from a typical example (using our addressing scheme here).


```

show ipv6 neighbors           : Shows the ipv6 neighbour cache
show ipv6 interface <interface> <number> : Interface status and configuration
show ipv6 interface          : Summary of IP interface status and configuration

```

- 8. Assign IPv6 Addresses to Loopback Interfaces.** While there is no need for a Loopback interface in this lab yet, it is still useful to configure an IPv6 address for it at this time. The loopback will be used for the iBGP peering later on in this lab. Note that OSPF and BGP router IDs are 32 bit integers and in IOS these are derived from the IPv4 address assigned to the Loopback interface (this has potential issues on network devices with no IPv4 address configured).

Q. Why do you think the lack of any IPv4 address on the router would be a problem? Ask the lab instructors to discuss.

As the minimum subnet size possible for IPv6 is a /64, we will assign the first /64 out of our /48 infrastructure block to be used for loopbacks – so we will use 2001:db8:0:0/64 for all the loopbacks. We have 14 routers in our lab – the assigned loopback addresses are:

R1	2001:db8::1/128	R8	2001:db8::8/128
R2	2001:db8::2/128	R9	2001:db8::9/128
R3	2001:db8::3/128	R10	2001:db8::a/128
R4	2001:db8::4/128	R11	2001:db8::b/128
R5	2001:db8::5/128	R12	2001:db8::c/128
R6	2001:db8::6/128	R13	2001:db8::d/128
R7	2001:db8::7/128	R14	2001:db8::e/128

For example, Router Team 1 would assign the following address and mask to the loopback on Router 1:

```

Router1(config)# interface loopback 0
Router1(config-if)# ipv6 address 2001:db8::1/128

```

Q: Why do we use /128 masks for the loopback interface address?

A: There is no physical network attached to the loopback so there can only be one device there. So we only need to assign a /128 mask – it is a waste of address space to use anything else.

Checkpoint #1: call lab assistant to verify the connectivity. Demonstrate that you can ping and telnet to the adjacent routers.

- 9. ISIS within the same AS.** ISIS will have already been set up in the AS for IPv4 – each team should confirm that the IGP is still operational before starting the next steps to add in IPv6 topology support.

- 10. Activating Multi-Topology ISIS.** We also need to activate multi-topology ISIS to roll out ISIS support for IPv6 if the existing network is already using IPv4 ISIS. This allows the IPv6 topology to be incrementally rolled out, very useful during deployment of IPv6. This means that each team can add ISIS IPv6 support without having to coordinate with their neighbouring teams.

```

Router1(config)# router isis workshop
Router1(config-router)# address-family ipv6
Router1(config-router-af)# multi-topology

```

NB. If we do not enable multi-topology, then each team will have to coordinate the enabling of IPv6 ISIS on each interface with their respective neighbouring teams. Failure to do so will result in the ISIS session going down, as there will be a topology mismatch on that interface.

NB. Multi-topology is configurable in IOS 12.3 and 12.4 but it does not work due to a bug which Cisco refuses to fix. The workaround is to use single topology, noting the caveat above, or use 12.2SRE, 12.2SXH, 12.4T, 15.0 or later IOS images.

- 11. Activated Multi-Topology ISIS.** All lab routers **must** have multi-topology ISIS enabled before the lab can proceed to active IPv6 ISIS on each interface. Failure to do so will mean that ISIS will see a topology mismatch and will tear the ISIS adjacency down. This is not a good situation on a live operational network, as it means the iBGP sessions will drop and customers will lose connectivity.

Checkpoint #2: demonstrate that you have enabled multi-topology ISIS to the lab demonstrators.

- 12. Activating ISIS on each interface.** All connected point to point and shared ethernet interfaces need to be configured with IPv6 ISIS. Otherwise, you may not be able to see network advertisements via ISIS from routers two or more hops away.

The example for the Router Team 1 is:

```
Router1(config)# interface fastethernet 0/0
Router1(config-if)# ipv6 router isis workshop
!
Router1(config)# interface fastethernet 0/1
Router1(config-if)# ipv6 router isis workshop
!
Router1(config)# interface serial 1/0
Router1(config-if)# ipv6 router isis workshop
```

Note: the ISIS ID on the interfaces must be matched with the router's ISIS ID.

- 13. ISIS Metrics.** Now each team needs to set the ISIS metric on each physical interface. The default ISIS metric for all interface types is 10. Unlike OSPF in IOS, ISIS has no automatic scheme to convert the interface bandwidth into a metric value. ISPs deploying ISIS have to come up with their own scheme (as in fact many ISPs using OSPF now also do)

In the lab we use metric 2 for the Ethernet interfaces and metric 20 for the Serial interfaces. For example:

```
Router1(config)# interface fastethernet 0/0
Router1(config-if)# isis ipv6 metric 2 level-2
!
Router1(config)# interface fastethernet 0/1
Router1(config-if)# isis ipv6 metric 2 level-2
!
Router1(config)# interface serial 1/0
Router1(config-if)# isis ipv6 metric 20 level-2
```

14. Announcing the Loopback /128. The loopback interface will have already been marked as passive when ISIS was set up for IPv4 routing on the network. Each team should confirm that the passive-interface command is still present for the Loopback.

15. Avoiding Traffic Blackhole on Reboot. When a router restarts after being taken out of service, ISIS will start distribute prefixes as soon as adjacencies are established with its neighbours. In the next part of the workshop lab, we will be introducing iBGP. So if a router restarts, ISIS will start up well before the iBGP mesh is re-established. This will result in the router landing in the transit path for traffic, with out the routing table being completed by BGP. There will not be complete routing information on the router, so any transit traffic (from customer to peer or upstream, or vice-versa) will be either dropped, or resulting in packets bouncing back and forth between adjacent routers. To avoid this problem, we require the router to not announce it is availability until the iBGP mesh is up and running. To do this, we have to provide the following command:

```
Router1(config)#router isis workshop
Router1(config-router)#address-family ipv6
Router1(config-router-af)#set-overload-bit on-startup wait-for-bgp
```

This sets ISIS' overload bit such that all IPv6 routes via this router will be marked as unreachable (very high metric) until iBGP is up and running. Once iBGP is running, the prefixes distributed by ISIS will revert to standard metric values, and the router will pass transit traffic as normal.

16. Ping Test #2. Ping all loopback interfaces in the classroom. This will ensure the ISIS 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 clns neighbor	: see a list of CLNS-IS neighbours that the router sees
show clns interface	: see if ISIS is configured and see the IS type
show isis database	: see ISIS link state database that the router has learned
show isis ipv6 rib	: see IPv6 ISIS routes that the router has learned
show isis topology	: see the ISIS topology as learned by the router

Checkpoint #3: call lab assistant to verify the connectivity. Save the configuration as it is on the router – use a separate worksheet, or the workspace at the end of this Module. You will require this configuration several times throughout the workshop.

17. Traceroute to all routers. Once you can ping all the routers, try tracing routes to all the routers using *trace x:x* command. For example, Router Team 1 would type:

```
Router1# trace 2001:db8::c
```

to trace a route to Router R12. If the trace times out each hop due to unreachable destinations, it is possible to interrupt the *traceroute* using the Cisco break sequence CTRL-^.

Q. Why do some trace paths show multiple IP addresses per hop?

A. If there are more than one equal cost paths, ISIS will “load share” traffic between those paths.

```
Router1>trace 2001:db8::c
```

```
Type escape sequence to abort.  
Tracing the route to 2001:db8::c  
  
 1 2001:db8:0:3::1      4 msec  
   2001:db8:0:2::1      0 msec  
   2001:db8:0:3::1      0 msec  
 2 2001:db8:0:f::1      4 msec  
   2001:db8:0:8::1      4 msec  
   2001:db8:0:f::1      0 msec  
 3 2001:db8:0:13::      4 msec *   4 msec  
Router1>
```

18. Other Features in ISIS. Review the documentation or use command line help by typing ? to see other *show* commands and other ISIS configuration features.

Review Questions

1. What IOS show command(s) will display the router's IPv6 forwarding table?
2. What IOS show command(s) will display the router's IPv6 ISIS database?