IPv6 Module 4 – OSPF to ISIS for IPv6

Objective: To migrate the OSPF version of Module 1 (running IPv4) to using ISIS as part of an IPv6 migration strategy. OSPF will be completely removed once the migration is completed. IPv6 Module 1c should be completed after this one.

Prerequisites: IPv4 Module1a and 1c (OSPF/iBGP).

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

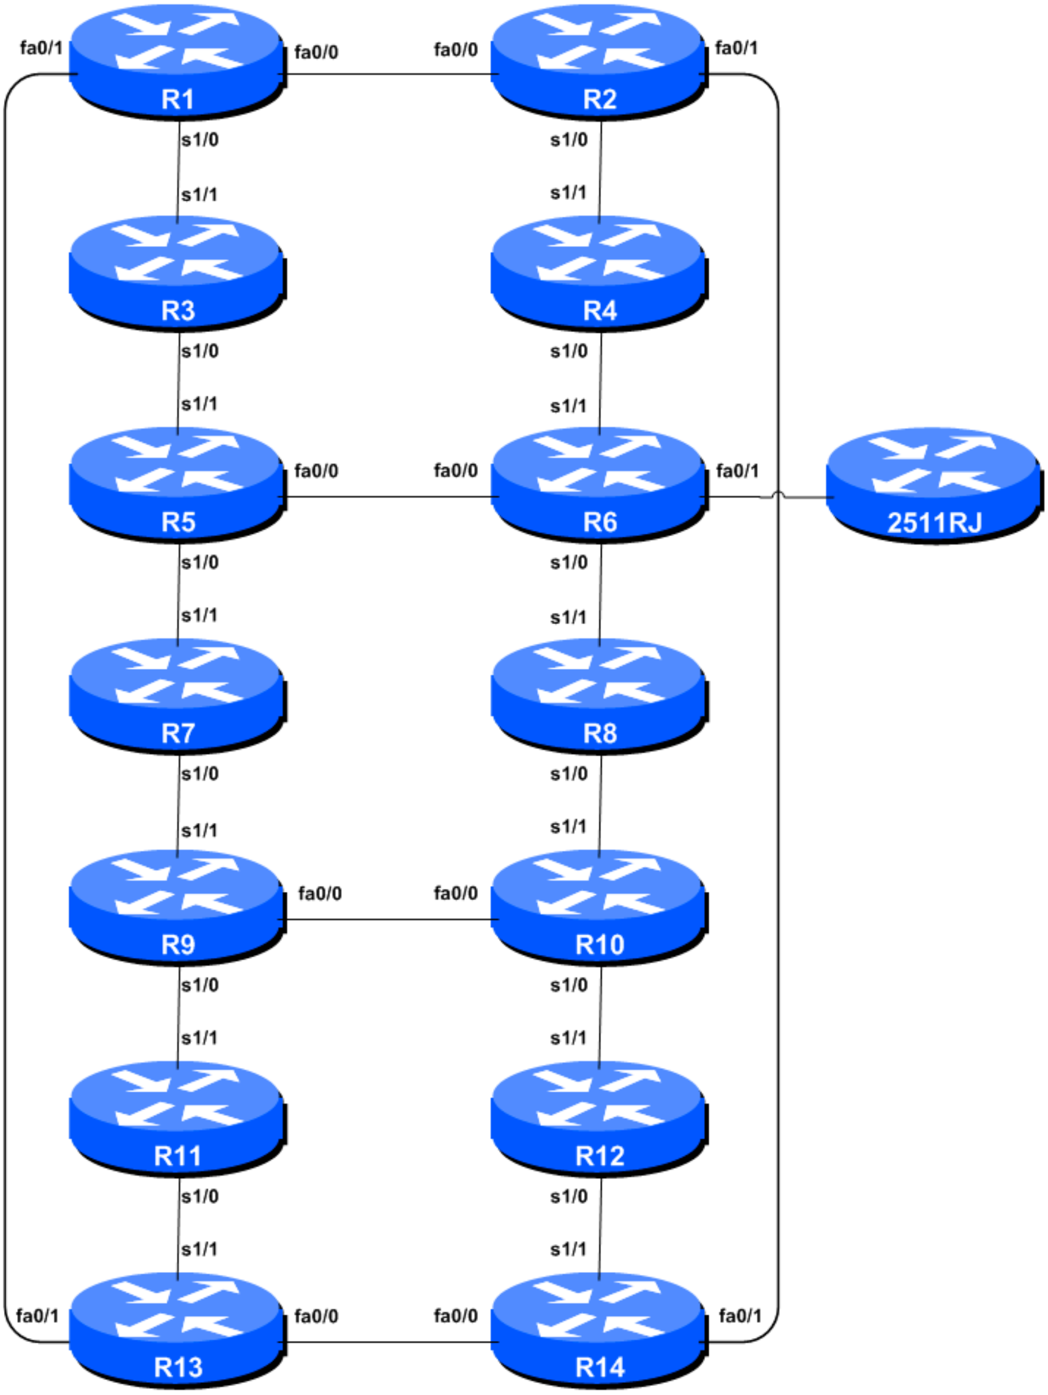

Figure 1 – ISP Lab Basic Configuration

Lab Notes

This Module is intended to migrate Module 1 of the IPv4 Workshop from OSPF to ISIS as part of a strategy to deploy IPv6 across the network. The topology and network 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.

Note: these labs assume that the routers used are using a minimum of IOS 12.4 mainline. Syntax predating IOS 12.4 is discussed in the optional sections throughout the workshop.

Lab Exercise

1. **Introduction.** This module works through a typical migration strategy an ISP would follow as they deploy IPv6 across their network. At the start of the module, they are simply running OSPF as their IGP, with iBGP providing the transit network overlay. By the end of the module, they will have added IPv6 dual stack, deployed ISIS to carry both IPv4 and IPv6 prefixes, and removed OSPF from their infrastructure. All without impacting the iBGP or the carrying of traffic across the backbone. The steps and process described here show a seamless deployment.

2. **ISIS rollout.** ISIS has administrative distance of 115. OSPF has administrative distance of 110. Therefore any routes carried within OSPF and within ISIS will see the OSPF versions taking priority. Which means that ISIS can deployed over an ISP backbone without impacting the operation of the network.

3. **ISIS with one area and one level (level-2) within the same AS.** Each router team should enable ISIS on their router, and use *workshop* as the ISIS ID in the configuration. In this module, we use level-2 in one area (*49.0001*) and use wide metrics (IOS default is the historical narrow metric and is not considered good practice). The NET should be *49.0001.x.x.x.x.00*, where *x.x.x.x* represents the router loopback IP address. For example, the loopback for Router1 is 10.0.15.241 which will make the NSAP address *49.0001.0100.0001.5241.00.*

```
Router1(config)# router isis workshop
Router1(config-router)#net 49.0001.0100.0001.5241.00
Router1(config-router)#is-type level-2-only
```

Q: Why do you have `is-type level-2-only` configured? Write your answer here:

Hint: A nice trick for converting the loopback interface address into the NSAP address is to take the loopback address and put the missing leading zeroes in. For example, Router 5 loopback address is 10.0.15.245; this is rewritten to 010.000.015.245 putting in the missing zeroes. Then rather than having the dot after every third character, move it to be after every fourth character. So 010.000.015.245 becomes 0100.0001.5245.

4. **Setting Wide Metrics.** We also set the metric-style to wide. ISIS supports two types of metric, narrow (historic now and not suitable for modern networks) and wide. IOS still defaults to narrow metrics, so we need to enter explicit configuration to change this to wide.

```
Router1(config)# router isis workshop
Router1(config-router)#metric-style wide level-2
```

5. **Activating ISIS on each interface.** Now that the ISIS process is configured, all connected point to point and shared ethernet interfaces need to be configured with ISIS. Else, you will not be able to see network advertisements via ISIS from routers two or more hops away. Here is an example configuration as would be used on Router1:

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

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

6. **ISIS Circuit Type and ISIS Metrics.** Now each team needs to set the circuit type and ISIS metric on each physical interface.

   The default circuit type is level-1-2 even though the router has been defined to be a level-2-only router.

   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 will use metric 2 for the Ethernet interfaces and metric 20 for the Serial interfaces.

   Combining the above, gives the example:

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

7. **Announcing the Loopback /32.** We do not need to set up ISIS adjacencies on the loopback interface as there are no neighbours there, so we mark it as passive:

```
Router1(config)# router isis workshop
Router1(config-router)# passive-interface Loopback0
```

Note that this will tell ISIS to install the loopback interface address in the ISIS RIB. We do NOT need to add an `ip router isis` statement onto the loopback interface itself. This is different from the required OSPF configuration, and often catches many engineers out, especially those who are learning ISIS after gaining experience with OSPF.

8. **ISIS Adjacencies.** Enable logging of ISIS adjacency changes. This is so that a notification is generated every time the state of a CLNS neighbor changes, and is useful for debugging purposes.

   (**Note:** From IOS 12.4 onwards, *log-adjacency-changes* is activated by default when ISIS is first configured.)

```
        Router1(config)#router isis workshop
        Router1(config-router)#log-adjacency-changes
```

9. **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)#set-overload-bit on-startup wait-for-bgp
```

   This sets ISIS' overload bit such that all 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.

10. **Ping Test #2.** Ping all loopback interfaces. 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 ip 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 rib`	: see ISIS local RIB that the router has learned
`show isis topology`	: see ISIS local RIB that the router has learned

Checkpoint #1: *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.*

STOP AND WAIT HERE

11. **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.

12. **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.

13. **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
    ```

14. **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. **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.

15. **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.

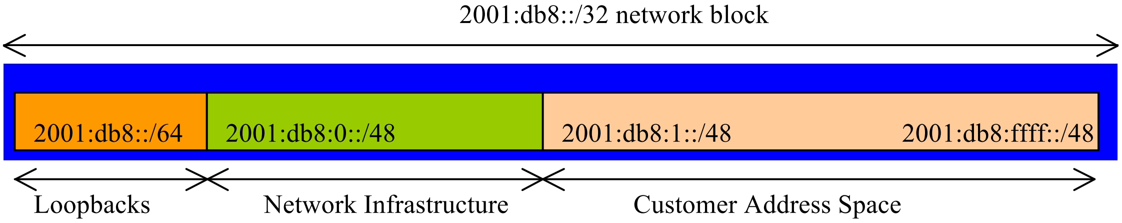

Figure 2 – Dividing allocated block of /32 into Customer, Infrastructure and Loopbacks

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 a side on /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).

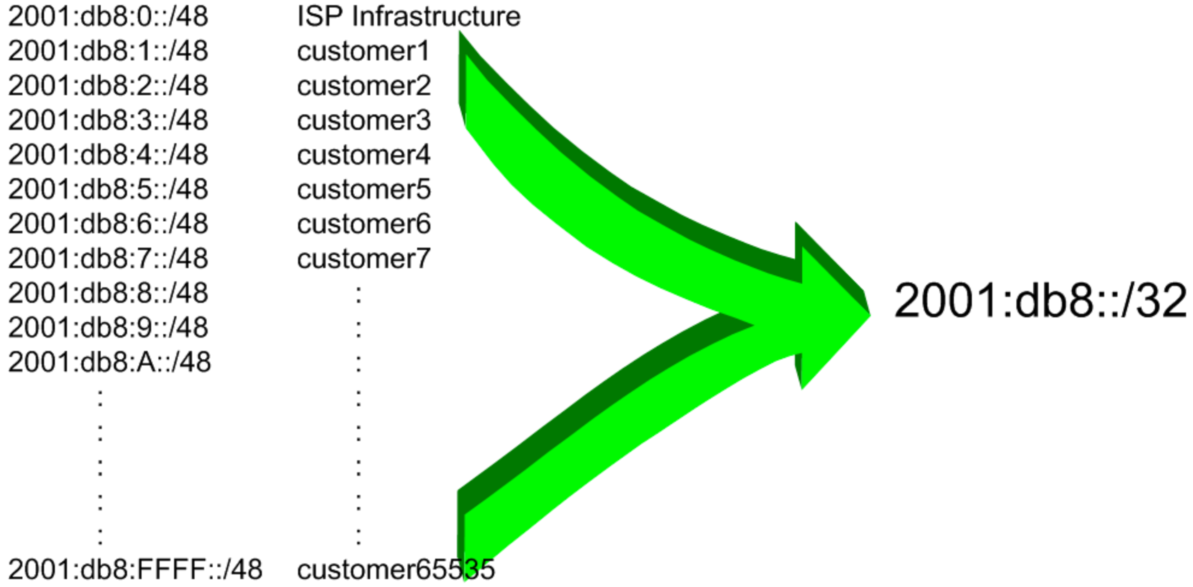

2001:db8:0::/48	ISP Infrastructure
2001:db8:1::/48	customer1
2001:db8:2::/48	customer2
2001:db8:3::/48	customer3
2001:db8:4::/48	customer4
2001:db8:5::/48	customer5
2001:db8:6::/48	customer6
2001:db8:7::/48	customer7
2001:db8:8::/48	:
2001:db8:9::/48	:
2001:db8:A::/48	:
:	:
:	:
:	:
:	:
:	:
2001:db8:FFFF::/48	customer65535

Figure 3 - Extract from an ISP Addressing Plan

15. **Back-to-Back Serial Connections.** Each team now needs to assign IPv6 addresses to the serial connections between the routers. See the addressing plan in the Appendices for the recommended addressing plan.

    Note: this lab will **not** use EUI-64 interface addressing, but instead will assign absolute addressing to each interface. The latter is much easier to manage, easier to handle for managing point-to-point peers and neighbour relationships.

A sample configuration might look like:

```
Router2(config)# interface serial 1/0
Router2(config-if)# ipv6 address 2001:db8:0:6::/127
```

Q: What network mask should be used on all IPv6 enabled interfaces?
A: The network mask should be /127. This is the subnet size used for all point-to-point links as recommended in RFC6164. We still reserve the entire /64 for this point-to-point link though, allowing simpler operational scalability should future changes be required.

Note: As discussed in the IPv6 presentation, ISPs are also using /126 and /120 as the subnet mask for point-to-point link addresses. We could have done this in the workshop as well, but chose to follow RFC6164 recommendations. We could also have numbered all our point-to-point links out of a single /64. However, this could mean potential problems in the future where developments in the IPv6 standard call on special uses for some of the bits between /65 and /128.

16. **Ethernet Connections.** As for the previous step, assign IPv6 addresses to the Ethernet point-to-point connections.

17. **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. Use the following commands to troubleshoot the connection:

```
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
```

18. **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 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.

19. **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.

20. **Activating ISIS on each interface.** The ISIS process is already configured from our IPv4 deployment earlier on in this module. We now need to configure all connected point-to-point and shared ethernet interfaces with ISIS.

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.

21. **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)

As for the IPv4 version earlier in this 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
```

22. Avoiding Traffic Blackhole on Reboot. As for the IPv4 version of this lab earlier, we want to configure ISIS so that the router does not blackhole transit traffic on restart. 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.

23. 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 rib`	: see ISIS local RIB that the router has learned
`show isis topology`	: see ISIS local RIB that the router has learned

Checkpoint #2: *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.*

STOP AND WAIT HERE

24. Switching to ISIS. The next step is a significant one. This is where we move the lab from using OSPF as the IGP to using ISIS as the IGP. Before doing so, review the OSPF to ISIS transition presentation, and make sure that all the points there have been met.

25. Changing OSPFv2 distance. All lab participants should now change the distance of OSPF on their router from the default of 110 to something higher than 115. We will choose a distance of 150. Set the distance for OSPFv2. To do this, we use the command as in the example below:

```
Router4(config)# router ospf 41
Router4(config-router)# distance 150
```

Once all the network is running with OSPF distance of 150, you should notice that ISIS is now carrying the best IGP path – and that there are no OSPF routes left in the routing table. Check the routing table by doing:

`show ip route`	: see if there is an IPv4 route for the intended destination
`show ipv6 route`	: see if there is an IPv6 route for the intended destination

You should see that the backbone IPv4 and IPv6 routes are known via ISIS, and there will be no prefixes left in OSPF. If there are prefixes left in OSPF, troubleshoot to find out what has been omitted from the ISIS configuration.

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.*

STOP AND WAIT HERE

26. Removing OSPF. Now that OSPF is no longer being used for backbone routes, we can remove OSPF from the network. To do this, each participant should do:

```
Router4(config)# no router ospf 41
```

Depending on IOS versions, there may also be some interface commands for OSPF which need to be removed. Once complete, check that there are no OSPF commands left in the router configuration.

```
Router4# sh run | include ospf
```

is a helpful command that will show you all configuration that still includes ospf in the line.

27. Successful migration. A successful migration has now been completed. Note that the iBGP for both IPv4 will not have been impacted, and that the network should have carried on operating as normal. The lab teams can now move forwards and deploy iBGP for IPv6 across the network (as per IPv6 Module 1c).

Checkpoint #4: *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.*

Review Questions

1. What is the administrative distance of OSPF?

2. What is the administrative distance of ISIS?

3. How does the administrative distance help in the migration from one IGP to another?