

Lab 5 – Migrating Dual Stack OSPF to ISIS

Objective: To migrate the OSPF version of Labs 1 to 4 (running IPv4 and IPv6) to using ISIS. OSPF will be completely removed once the migration is completed.

Prerequisites: Labs 1 to 4 for IPv4 and IPv6.

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

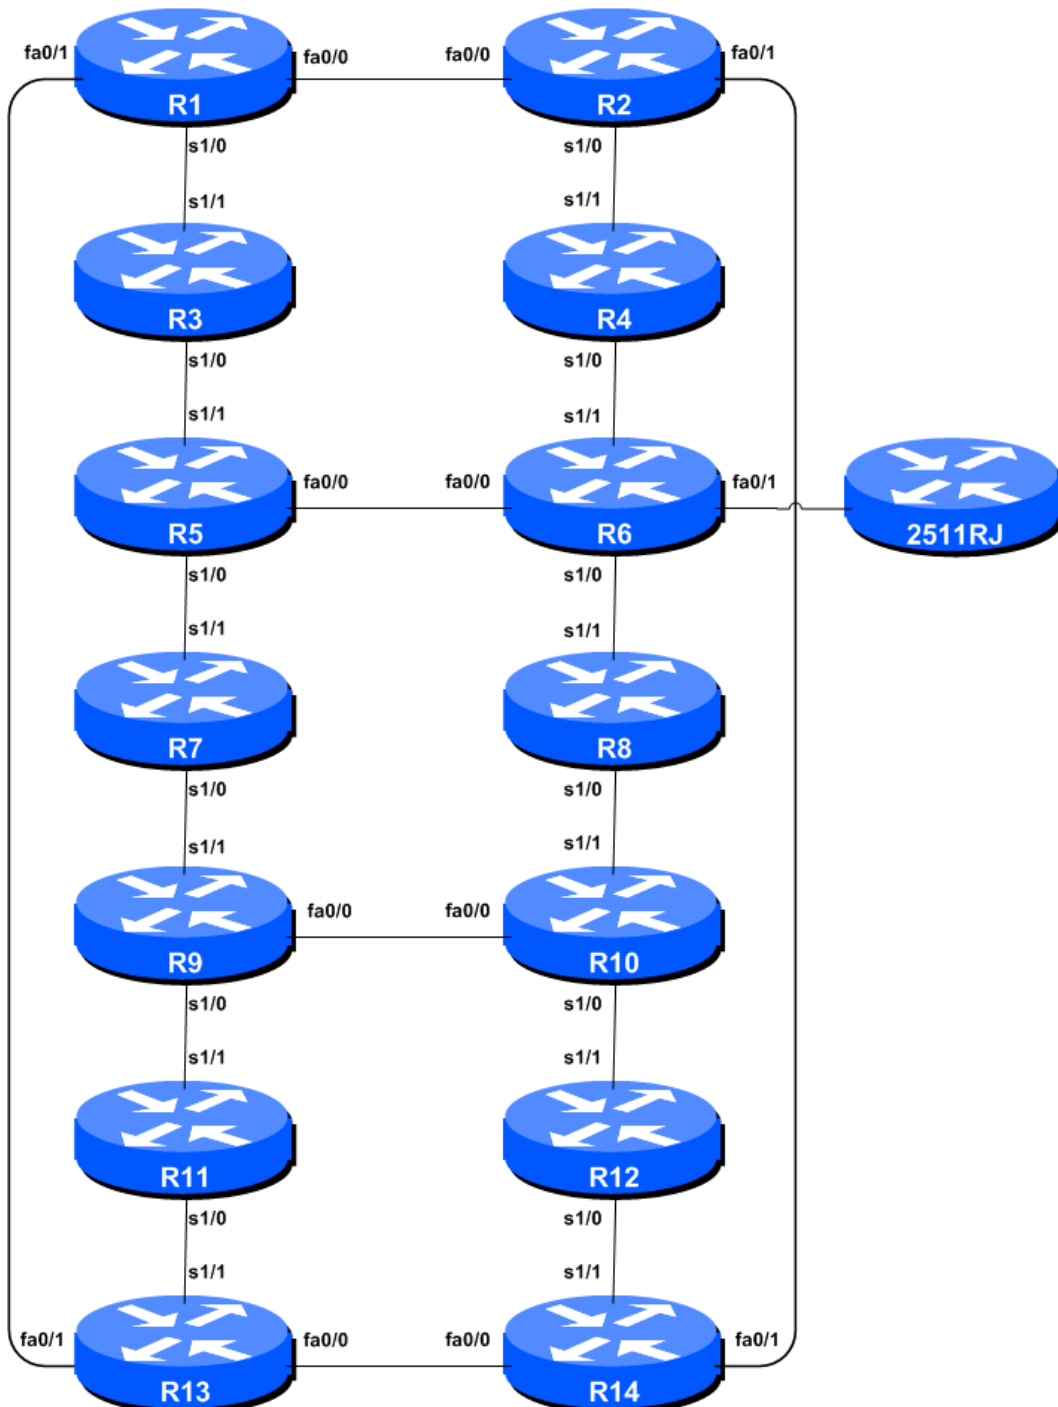


Figure 1 – ISP Lab Basic Configuration

Lab Notes

This Lab is intended to migrate Labs 1 to 4 of the IPv4/IPv6 Workshop from OSPF to ISIS once that lab has been completed. The topology and network configuration should be left exactly as it was at the end of Lab 4.

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. Preparation.** The protocol distance for ISIS is 115 (refer to the Routing Presentations) while the distance for OSPF is 110. Therefore ISIS can be rolled out on an existing network already using OSPF without any impact on the network's operation. The following steps guide the lab participant through the process of migrating from dual stack OSPF (using OSPFv2 for IPv4 and OSPFv3 for IPv6) to the multi-protocol ISIS.
- 2. ISIS within the same AS.** Each router Team should enable ISIS for IPv4 and IPv6 on their router. The ISIS ID should be *workshop* (see example). We use level-2 in one area (*49.0001*). The NET should be *49.0001.x.x.x.x.00*, where *x.x.x.x* represents the router loopback IPv4 address. For example, the loopback for Router1 should be *49.0001.0100.0001.5241.00*. Note that we do not set up CLNS adjacencies on the loopback interface so we mark this as a passive interface. You might have to activate ISIS on one interface first before you can mark any interface as passive.

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

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.

- 3. Setting Wide Metrics.** We also set the metric-style to wide. ISIS supports two types of metric, narrow (historic now) 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
```

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

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

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

- 6. Announcing the Loopback /128.** 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 `ipv6 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.

- 7. 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:

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

- 8. 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 commands for both IPv4 and IPv6 address families:

```
Router1(config)#router isis workshop
Router1(config-router)#set-overload-bit on-startup wait-for-bgp
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 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.

- 9. 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:

<code>show ip route</code>	: see if there is an IPv4 route for the intended destination
<code>show ipv6 route</code>	: see if there is an IPv6 route for the intended destination
<code>show clns neighbor</code>	: see a list of CLNS-IS neighbours that the router sees
<code>show clns interface</code>	: see if ISIS is configured and see the IS type
<code>show isis database</code>	: see ISIS link state database that the router has learned
<code>show isis rib</code>	: see ISIS local RIB that the router has learned
<code>show isis ipv6 rib</code>	: see ISIS local RIB that the router has learned
<code>show isis topology</code>	: 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.

STOP AND WAIT HERE

- 10. 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. We will only do IPv4 for now, rather than trying to convert both IPv4 and IPv6 routing. 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 routes are known via ISIS, and the backbone IPv6 routes are still known via OSPFv3.

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.

STOP AND WAIT HERE

11. Changing OSPFv3 distance. All lab participants should now change the distance of OSPFv3 on their router from the default of 110 to something higher than 115. We will choose a distance of 150. Set the distance for OSPFv3. To do this, we use the command as in the example below:

```
Router4(config)# ipv6 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 for IPv6 routes as well – and that there are no OSPFv3 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.

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

12. 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
Router4(config)# no ipv6 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.

13. Successful migration. A successful migration has now been completed. Note that the iBGP for both IPv4 and IPv6 will not have been impacted, and that the network should have carried on operating as normal.

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 ISIS default distance?
2. What is the OSPF default distance?