

IPv6 Module 2 – OSPF Areas

Objective: To migrate from one flat OSPF topology to OSPF areas, plus introduce neighbour authentication and area summarisation in the lab network.

Prerequisite: IPv6 Module 1 and OSPF presentation

Topology:

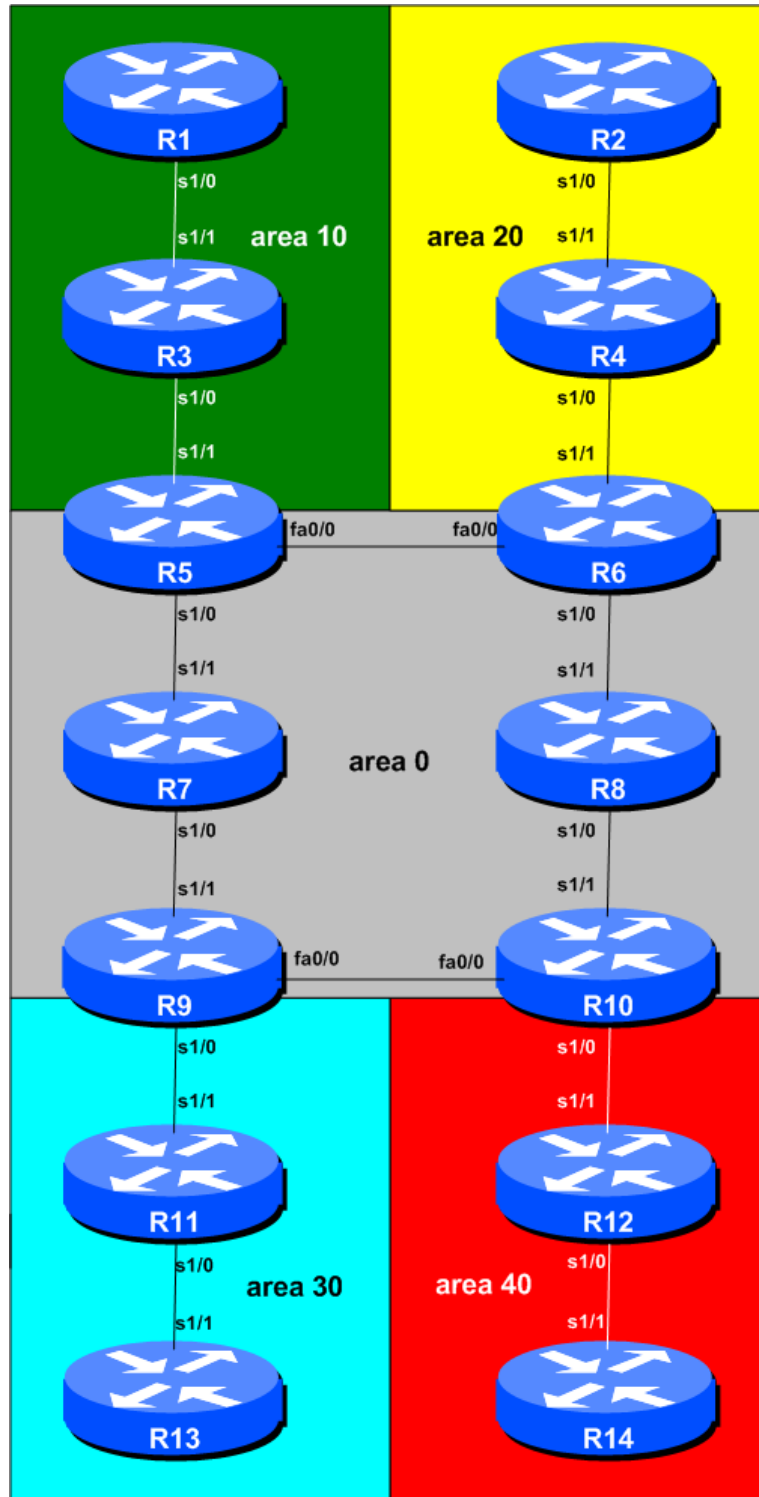


Figure 1 – OSPF Areas

Lab Notes

This module provides an introduction to areas in OSPF as used in ISP networks. The module forms the basis for the Route Reflector Module which follows. It is strongly recommended that the OSPF presentation is reviewed prior to starting this module. Prerequisites for this Module are Module 1, and the OSPF presentation. As before, ask the workshop instructors or refer to the online documentation if there is any doubt.

The topology shown in Figure 1 allows the student to configure the best combination for the study of OSPF areas.

Functional assignments of routers in Figure 1:

- Routers 5 to 10 represent the “core network” and the core interfaces are all in OSPF area 0. In a typical ISP backbone, these routers would carry all the internal link routes known in the ISPs network. Routers 5, 6, 9 and 10 are **Area Border Routers**, whereas Routers 7 and 8 are **Internal Routers**.
- Routers 1 and 3 are completely in OSPF area 10. Router 5 is the boundary between area 0 and area 10, so requires configuration for both areas.
- Routers 2 and 4 are completely in OSPF area 20. Router 6 is the boundary between area 0 and area 20, so requires configuration for both areas.
- Routers 11 and 13 are completely in OSPF area 30. Router 9 is the boundary between area 0 and area 30, so requires configuration for both areas.
- Routers 12 and 14 are completely in OSPF area 40. Router 10 is the boundary between area 0 and area 40, so requires configuration for both areas.

Note: if this module is attempted without completing the IPv4 version, please remember to set the OSPF router-id first.

Lab Exercise

1. **OSPF Migration.** This module is going to show the student how to migrate from a single area OSPF backbone to one using multiple areas. As ISP networks grow, quite often a single OSPF area ends up with an unmanageably large number of routers, resulting in reduced convergence rates impacting network performance. Of course, standard design advice is to start off with areas from day one, but many network designers unfortunately choose to ignore lessons well learnt by their predecessors.

Before starting ensure that the configuration is exactly as it was at the end of Module 1. OSPFv3 adjacencies should all be properly established and iBGP should be running with each neighbour adjacency established and active. Also, shutdown or disconnect interfaces which are not being used for this Module. These include connections between Router1 and Router2, Router1 and Router13, Router2 and Router14, and Router13 and Router14. It makes no sense to have a physical connection between non-zero OSPF areas as traffic between two non-zero areas must always transit Area 0.

Checkpoint #1: *Call the lab assistant and show the router configuration and connectivity.*

2. **Migration strategy.** Review the OSPF for ISPs presentation where one possible migration strategy was discussed. The best strategy is to start at the edge of the network and work inwards.

We want to migrate to using OSPF areas without causing too much downtime in the network – in other words, we want to try and ensure that all the iBGP sessions remain active. To this end, the outermost links of the network are changed from OSPF Area 0 to the sub area first. We want to avoid having any islands of Area0 in the network – this will result in partition of the network and disconnectivity for the “stub” Area 0. The following steps work through the process in detail.

- 3. Step 1: Network Outer Edge.** The first links we will move from Area 0 are the links between Router1 and Router3, Router2 and Router4, Router11 and Router13, and Router12 and Router14. The Router teams at either side of each link **MUST** work together to coordinate the change. For example, Router1 and Router3 should coordinate to change the area on the link between them from Area0 to Area10 at exactly the same time.

Example for Router1:

```
interface serial 1/0
  ipv6 ospf 41 area 10
interface loopback 0
  ipv6 ospf 41 area 10
```

Example for Router3:

```
interface serial 1/1
  ipv6 ospf 41 area 10
interface loopback 0
  ipv6 ospf 41 area 10
```

Once these two teams have changed the area for the link between them (don't forget to do the loopback interface as well!), the OSPF adjacency will re-establish in the new area, and the rest of the network will become visible again. The other teams mentioned above should do a similar thing for their interconnecting links and respective areas.

- 4. Step 2: Network Inner Edge.** Once the outer edge links have been moved to the new areas, we are ready to change the area of the next links in from the outer edge; these are the links between Router3 and Router5, Router4 and Router6, Router11 and Router9, and Router12 and Router10. Again the Router teams at either side of each link **MUST** work together to coordinate the change. For example, Router4 and Router6 should coordinate to change the area on the link between them from Area0 to Area20 at exactly the same time.

Example for Router4:

```
interface serial 1/0
  ipv6 ospf 41 area 20
interface loopback 0
  ipv6 ospf 41 area 20
```

Example for Router6:

```
interface serial 1/1
  ipv6 ospf 41 area 20
```

Once these two teams have changed the area for the link between them, the OSPF adjacency will re-establish in the new area, and the rest of the network will become visible again. The other teams mentioned above should do a similar thing for their interconnecting links and respective areas.

Note that Routers 5, 6, 9 and 10 should only change the Area of the link to their outer neighbour – they should **NOT** change the area of the loopback interface.

5. **Further steps.** Our lab network is not very large or complex, so the migration was quite simple. However, for larger networks, we would repeat the above steps, working from the outside in to the backbone Area, taking great care to ensure that no Area 0 islands are created as part of the process. Such a migration should have negligible impact on the network's operation.
6. **All router teams should check their routing table.** The routing table should look the same as it did in Module 1. All the prefixes in Area 0 will be available in Areas 10 through to 40. Although notice now that several of the routes will be marked as Inter Area (IA) routes. Here is an example of the typical output from show ipv6 route as seen from Router2:

```
IPv6 Routing Table - 29 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S    2001:DB8::/32 [1/0]
     via ::, Null0
OI   2001:DB8::1/128 [110/250]
     via FE80::20C:30FF:FE91:EEA0, Serial1/0
LC   2001:DB8::2/128 [0/0]
     via ::, Loopback0
OI   2001:DB8::3/128 [110/200]
     via FE80::20C:30FF:FE91:EEA0, Serial1/0
O    2001:DB8::4/128 [110/50]
     via FE80::20C:30FF:FE91:EEA0, Serial1/0
OI   2001:DB8::5/128 [110/150]
     via FE80::20C:30FF:FE91:EEA0, Serial1/0
OI   2001:DB8::6/128 [110/100]
     via FE80::20C:30FF:FE91:EEA0, Serial1/0
OI   2001:DB8::7/128 [110/200]
     via FE80::20C:30FF:FE91:EEA0, Serial1/0
OI   2001:DB8::8/128 [110/150]
```

Traces through the network. Once you have analysed the routing table as it is now, and compared with any notes you made in Module 1, try some traceroutes through the network.

Checkpoint #2: Call the lab instructors and show the function of your router. The inter-area OSPF “peerings” will now be up. If you have an area border router, demonstrate the neighbour relationships using “sh ip ospf neigh”. If you are not in an area border, you should now have a more complete routing table.

STOP AND WAIT HERE

7. **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 area will turn on authentication within that area. Routers which are **ABRs** will naturally have to enter configuration to cover all areas the router has interfaces in. This first step will enable each area to support authentication using the *area N authentication* command set.

We have to set a different Security Parameter Index (SPI) for each area. Please use the SPIs as listed below, otherwise each router team will have to negotiate SPIs for each area.

Area0	500	Area30	3000
Area10	1000	Area40	4000
Area20	2000		

An example configuration for Router6 might be:

```
ipv6 router ospf 41
  area 0 authentication ipsec spi 500 md5 0123456789ABCDEF0123456789ABCDEF
  area 20 authentication ipsec spi 2000 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.

Checkpoint #3: Call the lab instructors and demonstrate the OSPF is still working following the addition of neighbour authentication and that the OSPF routing table is still the same as it was previously.

8. Inter-area summarisation. It is fairly important, especially for larger ISP networks, to try and summarise the announcements going from sub-areas into Area 0 as much as is possible. Unfortunately we don't have an extensive lab network here, so it's going to be hard to see much benefit from summarisation, but we will try with the limited resources we have.

Areas 10, 20, 30 and 40 can summarise the point-to-point links they announce into Area 0. Routers 5, 6, 9 and 10, respectively, will introduce this configuration into the network. We will take the example of Area 20 and Router 6 below. Area 20 point to point subnets are 2001:db8:0:6::/64 (Router2-Router4) and 2001:db8:0:7::/64 (Router4-Router6). These can be aggregated into 2001:db8:0:6::/63 (even though we have addressed the links as /127, we have still reserved a /64 for each). The configuration which has to be introduced on Router6 therefore is:

```
ipv6 router ospf 41
 area 20 range 2001:db8:0:6::/63
```

Once the team operating Router6 has done this, check the routing table. Everyone (apart from Area 20 routers) will see that the entries for 2001:db8:0:6::/64 and 2001:db8:0:7::/64 will have been replaced by the aggregate, 2001:db8:0:6::/63.

9. The other three router teams (5, 9 and 10) should do a similar thing for their areas, aggregating the point to point links between their area routers. Once complete, eight /127s in the routing table should have been replaced by four /63 aggregates
10. **Final check.** Use the various “*show ipv6 ospf*” commands to see the OSPF status of the lab network now. Check the routing of IPv6 traffic and the IPv6 routing table. If you are missing any adjacencies, work with your neighbouring routers to work out why, and what might have gone wrong with the neighbour authentication.

Checkpoint #4: Call the lab instructors and show the routing table after the area summarisation was applied.