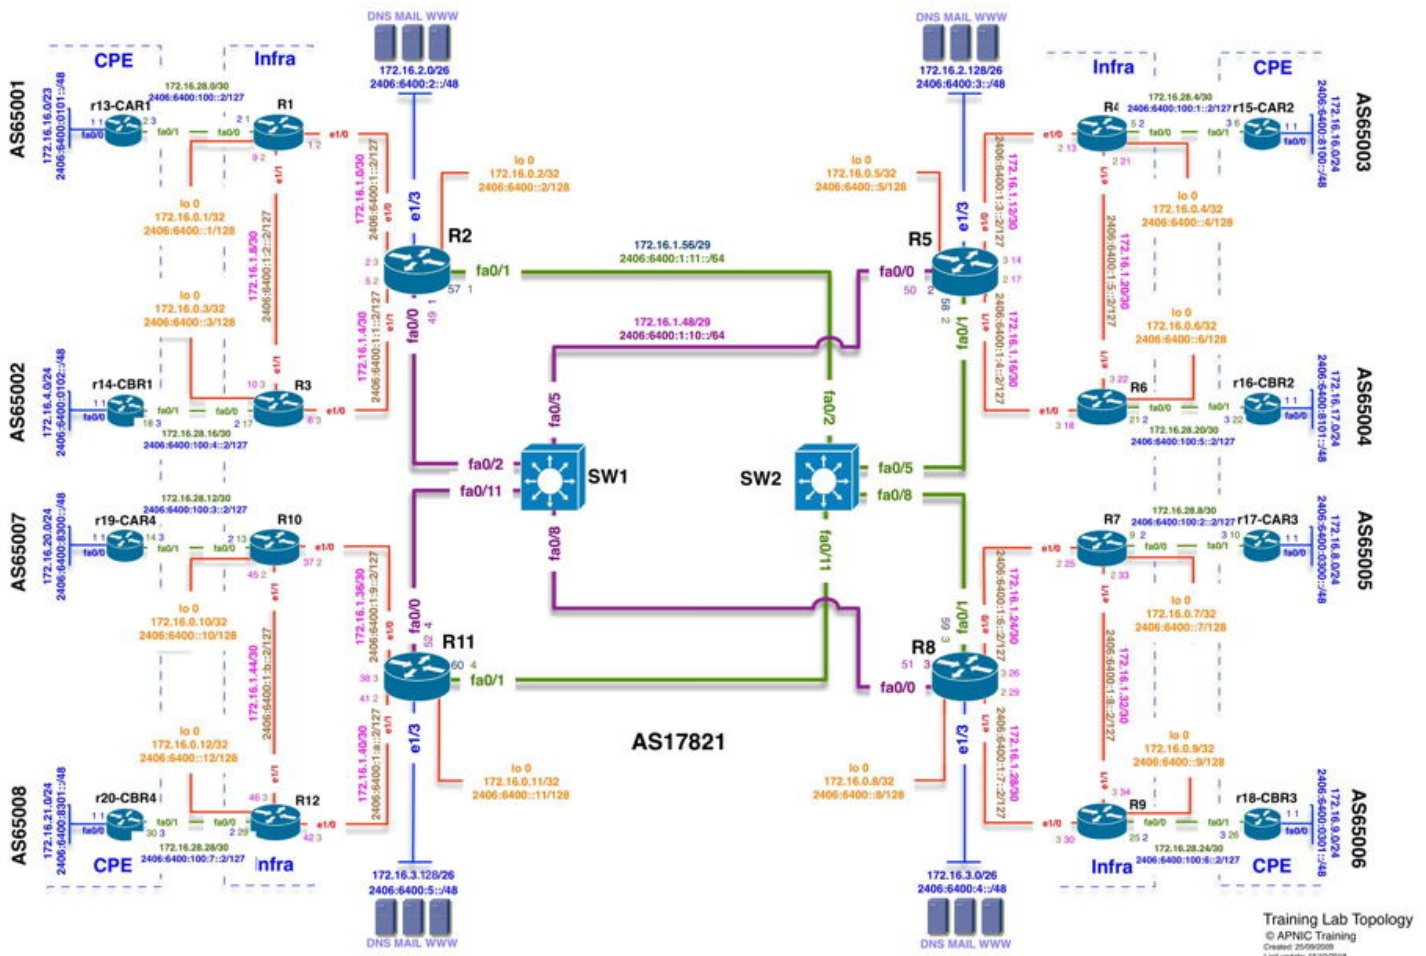




# Module 2 - OSPF(v3) Configuration

## Topology introduction:

- The topology below shows 4 regional networks comprised of a core POP and 2 aggregation POPs (edge routers).
- Edge routers aggregate downstream customers.
- The regional networks are interconnected with redundant transport links.



# Lab Tasks

---

After Module-1, you can only reach the directly connected neighbours (only the directly connected interfaces), but not the other routers in your network (R1-R12).

To be able to reach all routers within your network, we need to deploy interior gateway routing (IGP) protocols like OSPF and IS-IS.

In this lab, we will run OSPFv3 to carry both IPv4 and IPv6 address families. This helps consolidate both address families under a single OSPF process (as opposed to running two separate OSPF processes, OSPFv2/v3 for IPv4 and IPv6).

Take note of the following:

1. To scale, it is advisable to only carry your infrastructure prefixes in IGP (loopbacks, point-to-point addresses and transport links), but not customer routes/prefixes (which includes the customer point-to-point links).
2. The whole network (R1-R12) runs in OSPF area-0 (backbone area).
3. After finishing OSPFv3 configuration we should see the following 26 new prefixes in all infrastructure routers' routing table.

| Loopback                 | Point-to-Point                  | Transport                   |
|--------------------------|---------------------------------|-----------------------------|
| R1 => 2406:6400::1/128   | R1-R2 => 2406:6400:1::2/127     | Purple=>2406:6400:1:10::/64 |
| R2 => 2406:6400::2/128   | R2-R3 => 2406:6400:1:1::2/127   | Green=>2406:6400:1:11::/64  |
| R3 => 2406:6400::3/128   | R1-R3 => 2406:6400:1:2::2/127   |                             |
| R4 => 2406:6400::4/128   | R4-R5 => 2406:6400:1:3::2/127   |                             |
| R5 => 2406:6400::5/128   | R5-R6 => 2406:6400:1:4::2/127   |                             |
| R6 => 2406:6400::6/128   | R4-R6 => 2406:6400:1:5::2/127   |                             |
| R7 => 2406:6400::7/128   | R7-R8 => 2406:6400:1:6::2/127   |                             |
| R8 => 2406:6400::8/128   | R8-R9 => 2406:6400:1:7::2/127   |                             |
| R9 => 2406:6400::9/128   | R7-R9 => 2406:6400:1:8::2/127   |                             |
| R10 => 2406:6400::10/128 | R10-R11 => 2406:6400:1:9::2/127 |                             |
| R11 => 2406:6400::11/128 | R11-R12 => 2406:6400:1:a::2/127 |                             |
| R12 => 2406:6400::12/128 | R10-R12 => 2406:6400:1:b::2/127 |                             |

# Lab Exercise

---

## Step 1 - OSPFv3 Configuration:

Since OSPFv3 messages are sent over IPv6 (sourced using the router's link-local address) for both IPv4 and IPv6, make sure IPv6 is enabled on the router first, and the interfaces have IPv6 addresses configured.

- If there are no active IPv4 addresses configured, we need to manually configure a 32-bit router-ID
- In OSPFv3, `passive-interface default` will be configured. The function is to not create hello message by default to all interface. This approach is very effective if we would like to build OSPF adjacency on the selected interface. Specially if there is an external facing interface and no OSPF routing to them. We will disable `passive-interface` on those interfaces which are running ospfv3.

Example OSPFv3 config on a R1:

```
config t
router ospfv3 17821
address-family ipv4 unicast
router-id 172.16.0.1
passive-interface default
no passive-interface e1/0
no passive-interface e1/1
exit

address-family ipv6 unicast
router-id 172.16.0.1
passive-interface default
no passive-interface e1/0
no passive-interface e1/1
exit
```

To advertise networks (both IPv4 and IPv6) into OSPFv3, you need to enable OSPFv3 directly on the interfaces; there is no “network” command in OSPFv3. For example,

```
interface Loopback 0
ospfv3 17821 ipv4 area 0
ospfv3 17821 ipv6 area 0
exit
```

Since we are using ethernet links (which are broadcast/multi-access interfaces) as point-to-point links, to avoid the election of DR/BDR, we need to explicitly tell OSPF process that they are point-to-point links

On edge POP routers, for example, R4:

```
interface Loopback 0
ospfv3 17821 ipv4 area 0
ospfv3 17821 ipv6 area 0
exit

interface e1/0
ospfv3 17821 ipv4 area 0
ospfv3 17821 ipv6 area 0
ospfv3 network point-to-point
exit

interface e1/1
ospfv3 17821 ipv4 area 0
ospfv3 17821 ipv6 area 0
ospfv3 network point-to-point
exit
```

On Core routers, the transport links are still broadcast, therefore we only configure `point-to-point` on P2P interfaces, not on transport links.

For example, R2:

```
interface Loopback 0
ospfv3 17821 ipv4 area 0
ospfv3 17821 ipv6 area 0
exit

interface e1/0
ospfv3 17821 ipv4 area 0
ospfv3 17821 ipv6 area 0
ospfv3 network point-to-point
exit

interface e1/1
ospfv3 17821 ipv4 area 0
ospfv3 17821 ipv6 area 0
ospfv3 network point-to-point
exit

interface fa0/0
ospfv3 17821 ipv4 area 0
ospfv3 17821 ipv6 area 0
exit

interface fa0/1
ospfv3 17821 ipv4 area 0
ospfv3 17821 ipv6 area 0
```

Please remember to save the configuration.

## Step 2 - Verify OSPFv3 configuration:

`show ospfv3 neighbor` Check OSPFv3 neighbor table

`show ospfv3 database` Check OSPFv3 topology table

`show ip route ospfv3` Check ipv4 routes/prefixes learned through OSPFv3

`show ipv6 route ospf` Check ipv6 routes/prefixes learn through OSPFv3

## Step 3 - Verify Reachability:

Make sure you can reach ( `ping` ) all other routers in the network (loopbacks and point-to-points).

