

Advanced CCIE Routing & Switching 5.0

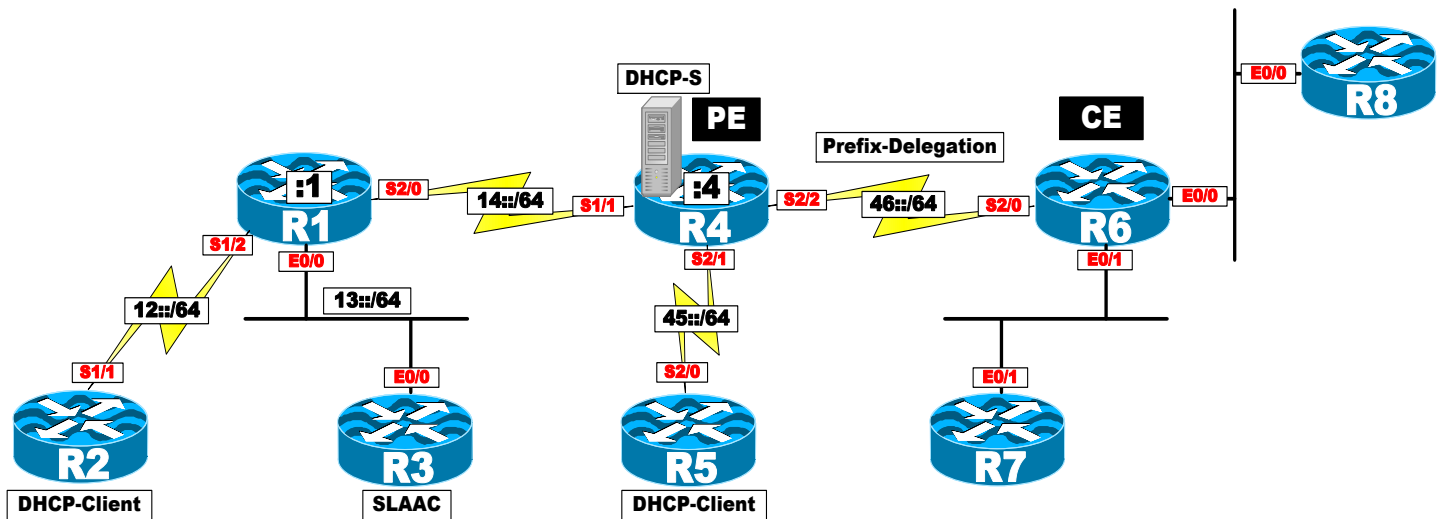
www.MicronicsTraining.com

Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

IPv6

Lab 1

Acquiring an IPv6 Address



Task 1

Configure the Ethernet segment connecting R1 to R3, do not assign an IPv6 address to R3, R3 should acquire the network portion of its IPv6 address through SLAAC process from R1. Use the following mac-addresses for these two routers:

R1 – 0000.1111.1111
 R3 – 0000.3333.3333

Let's configure both E0/0 interfaces of these two routers in VLAN 13 (Any VLAN ID can be used).

On SW1:

```
SW1 (config) #int range e0/1,e0/3
SW1 (config-if-range) #swi
SW1 (config-if-range) #swi mode acc
SW1 (config-if-range) #swi acc v 13
SW1 (config-if-range) #spanning portf
SW1 (config-if-range) #no shut
```

Stateless Address Auto-Configuration (SLAAC) is one method for the IPv6 clients to get the network portion of their IPv6 address. SLAAC provides a very simple process where the clients self-assign an IPv6 address based on the IPv6 prefix.

This process is achieved based on the following:

Host sends a router Solicitation (RS) message.

A router with IPv6 unicast routing enabled will reply with a router Advertisement (RA) message.

The Host takes the first 64 bits of the IPv6 prefix from the router Advertisement message and combines it with the 64 bit EUI-64 address to create a global unicast message.

The host also uses the source IPv6 address, in the IPv6 header, of the router Advertisement message, as its default gateway.

Duplicate Address Detection is performed by IPv6 clients to ensure the uniqueness of the new IPv6 address.

On R1:

In IPv6, unicast routing is disabled by default and in order for R1 to respond to the router Solicitation (RS) messages the unicast routing MUST be enabled:

```
R1 (config) #ipv6 unicast-routing
R1 (config) #int e0/0
R1 (config-if) #mac-address 0000.1111.1111
R1 (config-if) #ipv6 enable
R1 (config-if) #ipv6 address 13::1/64
R1 (config-if) #no shut
```

Let's enable "Debug ipv6 nd":

```
R1#deb ipv nd
  ICMP Neighbor Discovery events debugging is on

ICMPv6-ND: (Ethernet0/0, FE80::200:11FF:FE11:1111) send RA to FF02::1
ICMPv6-ND: (Ethernet0/0, FE80::200:11FF:FE11:1111) Sending RA (1800) to
FF02::1
ICMPv6-ND:   MTU = 1500
ICMPv6-ND:   prefix 13::/64 [LA] 2592000/604800
```

Once R3 is configured, it will take this Link Local IPv6 address and it will use it as its default gateway. It will then look at the prefix field highlighted in the output of the above show command and it will use that as its prefix address.

To verify the configuration:

On R1:

```
R1#show ipv6 interface e0/0 | inc FF
```

```
IPv6 is enabled, link-local address is FE80::200:11FF:FE11:1111
```


```
FF02::1 → All hosts within the local segment
```

```
FF02::2 → All routers within the local segment
```

```
FF02::1:FF00:1 → The Solicited Node Multicast based on the Global unicast IPv6 address
```

```
FF02::1:FF11:1111 → The Solicited Node Multicast based on the Link Local IPv6 address
```

In the output of the above show command we can see that the local router has auto generated a link local address based on the EUI-64 format. The following shows how the EUI-64 format is generated:



Then, the 7th bit in inverted:

0000:11FF:FE11:1111

0000 0000
1234 5678

The 7th bit is inverted, in this case it's 1

0000 0010

The new Link Local Address is:

0200:11FF:FE11:1111

OR

200:11FF:FE11:1111

The host portion of the IPv6 address is taken from the MAC-address of that interface (In Ethernet ONLY), but the MAC address is 48 bits and the node portion of the IPv6 address is 64 bits, so they decided to add "FFFE" which is 16 bits in the middle of the MAC address, then, the flip the most significant 7th bit and add "FE80::" in front of it.

On R3:

```
R3(config)#int e0/0
R3(config-if)#mac-address 0000.3333.3333
R3(config-if)#ipv6 enable
R3(config-if)#ipv6 address autoconfig default
R3(config-if)#no shut
```

The "ipv6 address autoconfig default" command makes R3 a SLAAC client.

On R3:

```
R3#sh ipv int e0/0
Ethernet0/0 is up, line protocol is up
IPv6 is enabled, link-local address is FE80::200:33FF:FE33:3333
No Virtual link-local address(es):
Stateless address autoconfig enabled
Global unicast address(es):
13::200:33FF:FE33:3333, subnet is 13::/64 [EUI/CAL/PRE]
valid lifetime 2591855 preferred lifetime 604655
```

Joined group address(es):

FF02::1

FF02::1:FF33:3333

MTU is 1500 bytes

ICMP error messages limited to one every 100 milliseconds

ICMP redirects are enabled

ICMP unreachable are sent

ND DAD is enabled, number of DAD attempts: 1

ND reachable time is 30000 milliseconds (using 30000)

ND NS retransmit interval is 1000 milliseconds

Default router is FE80::200:11FF:FE11:1111 on Ethernet0/0

R3#**show ipv6 interface brief e0/0**

FastEthernet0/0 [up/up]

FE80::200:33FF:FE33:3333

13::200:33FF:FE33:3333

R3#**sh ipv route ::/0**

Routing entry for ::/0

Known via "ND", distance 2, metric 0

Route count is 1/1, share count 0

Routing paths:

FE80::200:11FF:FE11:1111, Ethernet0/0

Last updated 00:05:55 ago

Task 2

Configure the serial link connecting R4 to R5, do not assign an IPv6 address to the S2/0 interface of R5. R5 should be configured as a DHCP Client acquiring an IPv6 address from R4. R5 should also get its domain name "MicronicsTraining.com" and the DNS server's IPv6 address of 2001:1111::1 from R4.

Let's configure R4's S2/1 interface and also configure the local router as a DHCP server:

On R4:

To work as a DHCP server, unicast-routing MUST be enabled:

```
R4 (config) #ipv6 unicast-routing
```

```
R4 (config) #ipv6 dhcp pool TST
```

Specifies the address range to provide in the pool.

```
R4 (config-dhcpv6) #address prefix 45::/64
```

Provides the DNS server and the domain name option to DHCP clients.

```
R4 (config-dhcpv6) #dns-server 2001:1:1111::1
```

```
R4 (config-dhcpv6) #domain-name MicronicsTraining.com
```

The "IPv6 enable" command MUST BE CONFIGURED.

```
R4 (config) #int s2/1
```

```
R4 (config-if) #ipv6 enable
```

```
R4 (config-if) #ipv6 address 45::4/64
```

The following command configures the DHCP server on the closest interface facing the clients.

```
R4 (config-if) #ipv6 dhcp server TST
```

The following command sets the "M" bit in the ND messages. This enables the hosts to use DHCP for address configuration and all the optional parameters such as domain name and DNS-server.

```
R4 (config-if) #ipv6 nd managed-config-flag
```

```
R4 (config-if) #no shut
```

There are two bits that we can use, the "M" bit and the "O" bits:

The "M" bit:

The "Managed address configuration" flag. When set, it indicates that addresses are available via Dynamic Host Configuration Protocol [DHCPv6]. Clients SHOULD use DHCP to obtain IPv6 addresses.

The "O" bit:

When set, it indicates that the DHCPv6lite is available for autoconfiguration of OTHER (non-address) information.

On R5:

```
R5(config)#int s2/0
R5(config-if)#ipv6 enable
R5(config-if)#ipv6 address dhcp
R5(config-if)#no shut
```

To verify the configuration:

On R5:

```
R5#show ipv6 int br s2/0
```

```
Serial2/0                [up/up]
    FE80::A8BB:CCFF:FE00:500
    45::A437:4578:2315:67BA
```

We can see that the local router acquired an IPv6 address from the DHCP server. How do we display the DHCP optional parameters that the local router acquired from the DHCP server?

```
R5#show ipv6 dhcp interface
```

```
Serial2/0 is in client mode
Prefix State is IDLE
Address State is OPEN
Renew for address will be sent in 11:59:04
List of known servers:
  Reachable via address: FE80::A8BB:CCFF:FE00:400
  DUID: 00030001AABBCC000400
  Preference: 0
Configuration parameters:
  IA NA: IA ID 0x000B0001, T1 43200, T2 69120
  Address: 45::A437:4578:2315:67BA/128
          preferred lifetime 86400, valid lifetime 172800
```


expires at Jun 22 2019 02:38 AM (172744 seconds)

DNS server: 2001:1:1111::1

Domain name: MicronicsTraining.com

Information refresh time: 0

Prefix Rapid-Commit: disabled

Address Rapid-Commit: disabled

On R4:

R4#show ipv6 dhcp binding

Client: FE80::A8BB:CCFF:FE00:500

DUID: 00030001AABBCC000500

Username : unassigned

VRF : default

IA NA: IA ID 0x000B0001, T1 43200, T2 69120

Address: 45::A437:4578:2315:67BA

preferred lifetime 86400, valid lifetime 172800

expires at Jun 22 2019 02:38 AM (172656 seconds)

The output of the above show command reveals that the client's link local IPv6 address is "FE80::A8BB:CCFF:FE00:500" and the IPv6 address assigned to that client is "45::A437:4578:2315:67BA"

Task 3

Configure the serial link connecting R1 to R4 based on the diagram in the beginning of this lab. Configure the serial link connecting R1 to R2 based on the following policy: DO NOT configure an IPv6 address on R2's S1/1 interface, this router should be configured to acquire an IPv6 address from the DHCP Server (R4), R1 should be configured as a DHCP relay agent.

Let's configure the serial interfaces first:

On R1:

R1 (config) #**ipv6 unicast-routing**

```
R1 (config) #int s1/2
R1 (config-if) #ipv6 address 12::1/64
R1 (config-if) #ipv6 enable
R1 (config-if) #no shut
```

```
R1 (config) #int s2/0
R1 (config-if) #ipv6 enable
R1 (config-if) #ipv6 address 14::1/64
R1 (config-if) #no shut
```

On R4:

```
R4 (config) #int s1/1
R4 (config-if) #ipv6 enable
R4 (config-if) #ipv6 address 14::4/64
R4 (config-if) #no shut
```

To test the configuration:

On R4:

```
R4#ping 14::1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 14::1, timeout is 2 seconds:

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

In this case R2 is going to be the DHCP-Client, once configured, R2 will send an RS (Router Solicit), R1 will receive this request and it will relay that request to R4 the DHCP Server. R4 will go through its scopes and it will lease out an IPv6 address from that scope.

R1 will receive the RA (Router Advertise) and it will relay it down to R2, and this is how R2 acquires an IPv6 address. A very simple process that is very similar to IPv4.

Let's configure another scope for the 12::/64 network:

On R4:

```
R4 (config) #ipv6 dhcp pool TST-R2
R4 (config-dhcpv6) #address prefix 12::/64
R4 (config-dhcpv6) #dns-server 2001:12:12:12::12
R4 (config-dhcpv6) #domain-name MicronicsTraining.com
```

Let's apply the pool to the S1/1 interface, which is the closest interface facing the DHCP client:

```
R4 (config-dhcpv6) #int s1/1
R4 (config-if) #ipv6 dhcp server TST-R2
```

NOTE: If the "address prefix 12::/64" is configured in the previous pool (TST), R2 will get two IPv6 address one from the 12::/64 network and the second one from the 45::/64 network.

Let's configure R1, the DHCP relay agent:

On R1:

```
R1 (config) #int s1/2
```

The following command configures the IPv6 address of the DHCP server by using the "destination" keyword. The S2/0 interface configuration is optional in this scenario. The serial2/0 interface has to be configured ONLY if the "destination" keyword references the Link Local ipv6 address of R4.

```
R1 (config-if) #ipv6 dhcp relay destination 14::4 s2/0
R1 (config-if) #ipv6 nd managed-config-flag
```

The "IPv6 unicast-routing" command MUST be configured, since this was configured in Task-1, we don't have to configure it again.

Finally the DHCP client is configured.

On R2:

```
R2 (config) #int s1/1
R2 (config-if) #ipv6 enable
R2 (config-if) #ipv6 address dhcp
R2 (config-if) #no shut
```

To verify the configuration:

On R2:

```
R2#show ipv6 interface bri s1/1
```

```
Serial1/1          [up/up]  
FE80::A8BB:CCFF:FE00:200  
12::D183:8C5:C92D:774
```

```
R2#show ipv6 dhcp interface
```

Serial1/1 is in client mode

```
Prefix State is IDLE  
Address State is OPEN  
Renew for address will be sent in 11:59:30  
List of known servers:
```

```
Reachable via address: FE80::200:11FF:FE11:1111 → The Link Local IPv6
```

```
DUID: 00030001AABBCC000400
```

```
Address of R1
```

```
Preference: 0
```

```
Configuration parameters:
```

```
IA NA: IA ID 0x00080001, T1 43200, T2 69120
```

```
Address: 12::D183:8C5:C92D:774/128
```

```
preferred lifetime 86400, valid lifetime 172800  
expires at Jun 23 2019 06:15 PM (172771 seconds)
```

```
DNS server: 2001:12:12:12::12
```

```
Domain name: MicronicsTraining.com
```

```
Information refresh time: 0
```

```
Prefix Rapid-Commit: disabled
```

```
Address Rapid-Commit: disabled
```

Task 4

Configure R3 to get the DNS and its domain name of “**MicroncisTraining.com**” from the DHCP server but it should continue to use SLAAC for its IPv6 address.

We saw the configuration of the “M” bit in which case the DHCP server gave an IPv6 address plus the optional DHCP parameters such as the DNS and the domain name. With the “O” bit set the DHCP server will ONLY give out the optional DHCP parameters, the client will NOT acquire an IPv6 address from the DHCP server. R3 has already used the SLAAC process to get the network portion of its IPv6 address, let’s configure R1 to accommodate this request.

Let’s see R3’s E0/0 configuration:

On R3:

```
R3#show run int e0/0 | b interface
interface Ethernet0/0
  mac-address 0000.3333.3333
  no ip address
  ipv6 address autoconfig default
  ipv6 enable
end
```

All we need to do is enable the “O” bit and configure the relay configuration on the E0/0 interface of R1:

On R1:

```
R1(config)#int e0/0
R1(config-if)#ipv6 dhcp relay destination 14::4 e0/0
R1(config-if)#ipv6 nd other-config-flag
```

To verify the configuration:

On R3:

```
R3#show ipv6 interface brief e0/0
```

```
Ethernet0/0          [up/up]
  FE80::200:33FF:FE33:3333
  13::200:33FF:FE33:3333
```

```
R3#show ipv6 dhcp interface e0/0
```

```
Ethernet0/0          [up/up]
  FE80::200:33FF:FE33:3333
  13::200:33FF:FE33:3333
```

```
R3#sh ipv dhcp inter e0/0
```

```
Ethernet0/0 is in client mode
Prefix State is IDLE (0)
Information refresh timer expires in 23:59:21
Address State is IDLE
List of known servers:
  Reachable via address: FE80::200:11FF:FE11:1111
  DUID: 00030001AABBCC000400
  Preference: 0
  Configuration parameters:
    DNS server: 2001:12:12:12::12
    Domain name: MicronicsTraining.com
  Information refresh time: 0
Prefix Rapid-Commit: disabled
Address Rapid-Commit: disabled
```

```
R3#sh ipv6 inter e0/0
```

```
Ethernet0/0 is up, line protocol is up
IPv6 is enabled, link-local address is FE80::200:33FF:FE33:3333
No Virtual link-local address(es):
Stateless address autoconfig enabled
Global unicast address(es):
  13::200:33FF:FE33:3333, subnet is 13::/64 [EUI/CAL/PRE]
  valid lifetime 2591866 preferred lifetime 604666
Joined group address(es):
  FF02::1
  FF02::1:FF33:3333
```

```
MTU is 1500 bytes
ICMP error messages limited to one every 100 milliseconds
ICMP redirects are enabled
ICMP unreachable are sent
ND DAD is enabled, number of DAD attempts: 1
ND reachable time is 30000 milliseconds (using 30000)
ND NS retransmit interval is 1000 milliseconds
Default router is FE80::200:11FF:FE11:1111 on Ethernet0/0
```

We can see that R3 received its DNS and domain name information from the DHCP server but it received the network portion of its IPv6 address from R1 through the SLAAC process.

Task 5

Re-configure R5 to acquire its IPv6 address from R4 (The DHCP Server) using two messages instead of four.

The DHCPv6 client can acquire its IPv6 address and optional parameters from a DHCP server in two ways:

Rapid-commit: In this process **ONLY two messages** are exchanged, a Solicit from the client to the server and a reply from the server to the client.

The default: The normal way which is the default, the DHCP client and the DHCP server exchange four DHCP messages and they are: Solicit, Advertise, Request, and Reply.

Before the task is configured, let's enable "debug ipv6 dhcp" and "Default interface S2/0" on R5 and reconfigure a regular or the default four message exchange:

On R5:

```
R5(config)#default inter s2/0
Interface Serial1/4 set to default configuration

R5#debug ipv6 dhcp
IPv6 DHCP debugging is on
```

```
R5 (config) #int s2/0
R5 (config-if) #shut
R5 (config-if) #ipv6 enable
R5 (config-if) #ipv6 address dhcp
R5 (config-if) #no shut
```

Now, let's configure R5 as a DHCP client to exchange four messages before it acquires an IPv6 address and the optional parameters from the DHCP server, R4.

```
IPv6 DHCP: Sending SOLICIT to FF02::1:2 on Serial2/0
IPv6 DHCP: Received ADVERTISE from FE80::A8BB:CCFF:FE00:400 on Serial2/0
IPv6 DHCP: Adding server FE80::A8BB:CCFF:FE00:400
IPv6 DHCP: Sending REQUEST to FF02::1:2 on Serial2/0
IPv6 DHCP: DHCPv6 address changes state from SOLICIT to REQUEST
(ADDR_ADVERTISE_RECEIVED) on Serial2/0
IPv6 DHCP: Received REPLY from FE80::A8BB:CCFF:FE00:400 on Serial2/0
IPv6 DHCP: Processing options
IPv6 DHCP: Adding address 45::A437:4578:2315:67BA/128 to Serial2/0
IPv6 DHCP: T1 set to expire in 43200 seconds
IPv6 DHCP: T2 set to expire in 69120 seconds
IPv6 DHCP: Configuring DNS server 2001:1:1111::1
IPv6 DHCP: Configuring domain name MicronicsTraining.com
```

The “Rapid-commit” option MUST be configured on the DHCP client and the DHCP server. To configure the DHCP server for “rapid-commit”:

On R4:

```
R4 (config) #int s2/1
R4 (config-if) #ipv6 dhcp server TST rapid-commit
```

Let's “Default interface S2/0” on R5 and configure the “Rapid-commit” option and see the difference:

On R5:

```
R5 (config) #default inter s2/0
Interface Serial2/0 set to default configuration
```



```
R5 (config) #int s2/0
R5 (config-if) #ipv6 enable
R5 (config-if) #ipv6 address dhcp rapid-commit
R5 (config-if) #no shut

IPv6 DHCP: Sending SOLICIT to FF02::1:2 on Serial2/0
IPv6 DHCP: Received REPLY from FE80::A8BB:CCFF:FE00:400 on Serial2/0
IPv6 DHCP: Adding server FE80::A8BB:CCFF:FE00:400
IPv6 DHCP: Processing options
IPv6 DHCP: Adding address 45::A437:4578:2315:67BA/128 to Serial2/0
IPv6 DHCP: T1 set to expire in 43200 seconds
IPv6 DHCP: T2 set to expire in 69120 seconds
IPv6 DHCP: Configuring DNS server 2001:1:1111::1
IPv6 DHCP: Configuring domain name MicronicsTraining.com
```

As we can see only two messages were exchanged.

Task 6

Configure R4, R6, R7 and R8 based on the following:

DO NOT assign an IPv6 address to R6, R7 or R8.

Configure the E0/1 interfaces of R6 and R7 in VLAN 67.

Configure the E0/0 interfaces of R6 and R8 in VLAN 68.

Configure the S2/2 interface of R4 with 46::4/64 IPv6 address and 0000.4444.4444 MAC address.

On R4:

```
R4 (config) #int s2/2
R4 (config-if) #ipv6 enable
R4 (config-if) #mac-address 0000.4444.4444
R4 (config-if) #ipv6 addr 46::4/64
R4 (config-if) #no shut
```

On SW2:

```
SW2 (config) #int rang e1/2-3
SW2 (config-if-range) #swi
SW2 (config-if-range) #swi mode acc
SW2 (config-if-range) #swi acc v 168
SW2 (config-if-range) #spannin portf
SW2 (config-if) #no shut
```

On SW1:

```
SW1 (config) #int range e1/2,e2/0
SW1 (config-if-range) #swi mode acc
SW1 (config-if-range) #swi acc v 68
SW1 (config-if-range) #no shut
```

Task 7

ISP-A has an IPv6 prefix of 46:1:1::/48 and it needs to subnet this network to /56 subnets and use subnet one for its client R6.

Company-A (R6) has two sites that are connected through its E0/0 and E0/1 interfaces, but soon this company will have 16 remote sites.

Company-A (R6) should acquire a subnet from ISP-A (R4) as it adds more remote sites.

The first subnet should be automatically assigned to its E0/0 interface with the host portion of its IPv6 address as "::10"; R8 should use R6 as its default gateway.

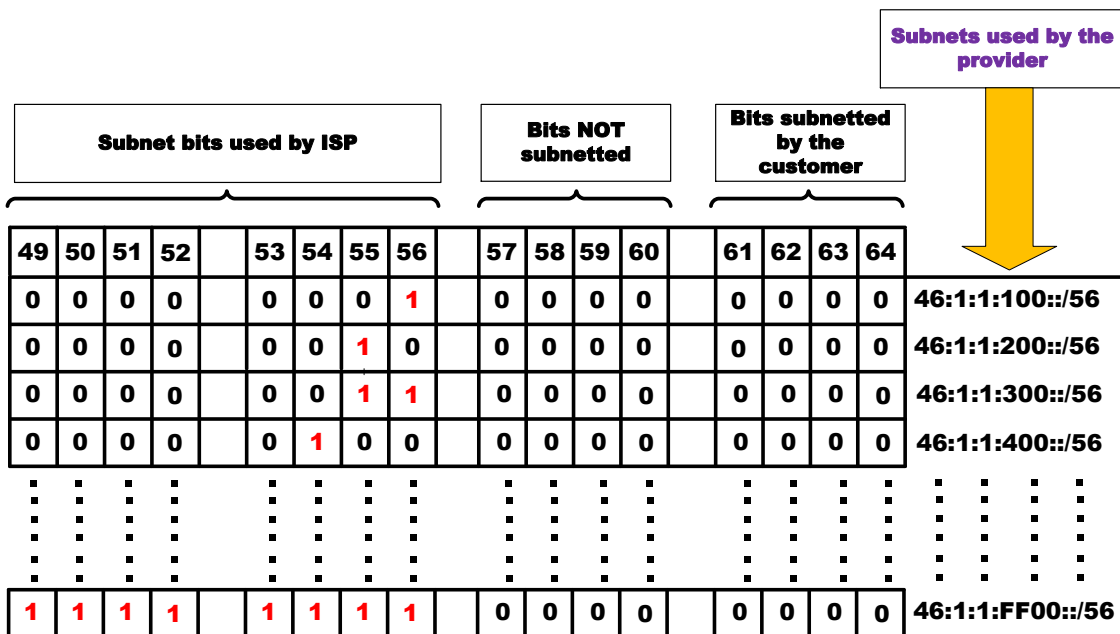
R6's E0/1 interface should be automatically assigned the third subnet with the host portion of "::1". R7 should use R6 as its default gateway.

R7 and R8 should automatically acquire the network portion of their IPv6 address from R6, they should auto generate their host portion of their IPv6 address.

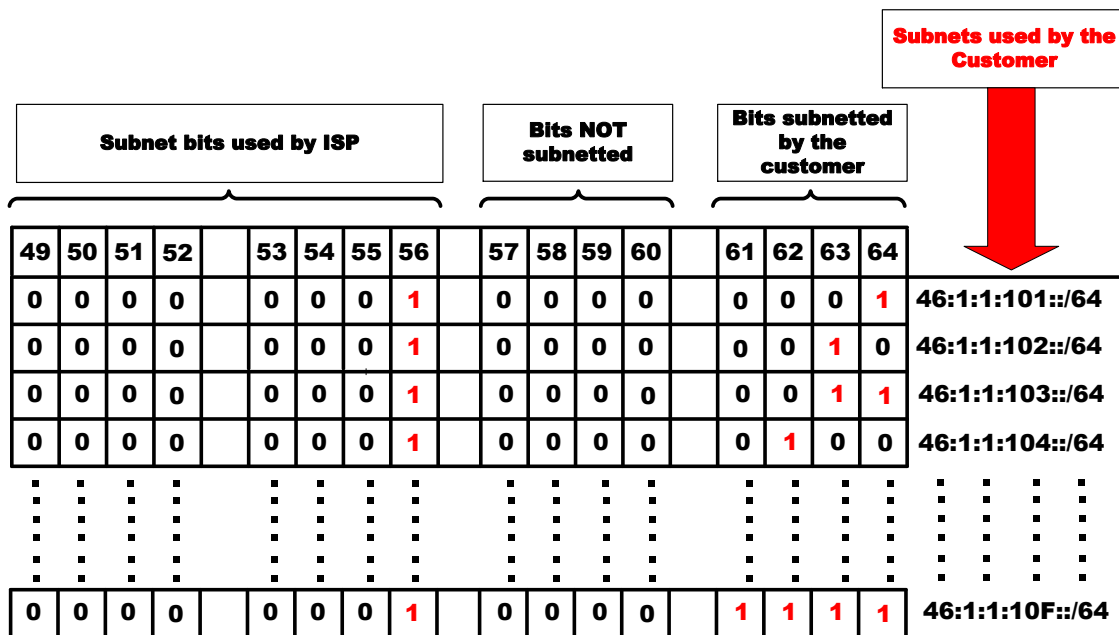
Both R7 and R8 should have reachability to R4's S2/2 IPv6 address.

DO NOT configure any static route or configure static IPv6 address/es to accomplish this task.

The following diagram shows the bits used by the provider to generate more /56 subnets. One of these subnets is given to the customer.



The following diagram shows the 4 bits that are given to the customer. The customer can use these 4 bits to generate 16 networks:



To resolve this task, we need to configure prefix-delegation. The purpose of the prefix delegation mechanism is to delegate prefixes to the CEs automatically. The prefix delegation mechanism typically delegate prefix lengths between /48 to /64. In this topology R4 is the delegating router.

On R4:

The following configures a local pool that instructs the router to hand out /60 addresses but the first 56 bits of the addresses must be 46:1:1:100:

```
R4 (config) #ipv6 local pool test 46:1:1:100::/56 60
```

A regular DHCPv6 pool is configured and it references the local pool and assigns a life time of infinity.

```
R4 (config) #ipv6 dhcp pool ISP
R4 (config-dhcpv6) #prefix-delegation pool test lifetime infinite infinite
```

The pool is referenced on the S2/2 interface of R4 facing R6:

```
R4 (config) #int s2/2
R4 (config-if) #ipv6 dhcp server ISP
```

On R6:

```
R6 (config) #ipv6 unicast-routing
```

This router is typically the CE router, it is acquiring the network portion of its IPv6 address from the PE router (R4) through SLAAC process, and it auto generates the host portion using EUI-64.

```
R6 (config) #int s2/0
R6 (config-if) #ipv6 enable
R6 (config-if) #ipv6 address autoconfig default
```

The “`ipv6 dhcp client pd`” command enables request for prefix delegation through the interface on which this command is configured. If this command is not configured, the local router will NOT generate a PD request. The pool of IPv6 addresses that are received by the local router will be placed in the general cache.

```
R6(config-if)#ipv6 dhcp client pd TST-ISP
R6(config-if)#no shut
```

Before the E0/0 and E0/1 interfaces are configured, let's verify and see if R6 did acquire an IPv6 address through the SLAAC process from R4:

```
R6#sh ipv6 int br s2/0
```

```
Serial2/0 [up/up]
FE80::A8BB:CCFF:FE00:600
46::A8BB:CCFF:FE00:600
```

If R6 fails to get an IPv6 address, you must shut and then no shut its s2/0 interface.

```
R6(config)#int e0/0
R6(config-if)#ipv6 enable
```

The following example shows how to enable IPv6 processing on the interface and configure an address based on the general prefix called "TST-ISP". The "::1" is the first subnet, meaning that you are assigning the first subnet to this interface; so if you wish to assign the second subnet, "::2" should be used. The ":10" is the host portion of the IPv6 address that you wish the local interface to have.

```
R6(config-if)#ipv6 address TST-ISP ::1:0:0:0:10/64
R6(config-if)#no shut
```

```
R6(config)#int e0/1
R6(config-if)#ipv6 enable
R6(config-if)#ipv6 address TST-ISP ::3:0:0:0:1/64
R6(config-if)#no shut
```

To verify the configuration:

On R6:

```
R6#show ipv6 inter bri s2/0
```

```
Serial2/0 [up/up]
FE80::A8BB:CCFF:FE00:600
```

```
46::A8BB:CCFF:FE00:600
```

```
R6#ping 46::4
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 46::4, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms
```

```
R6#show ipv6 inter bri e0/0
```

```
Ethernet0/0 [up/up]
```

```
FE80::A8BB:CCFF:FE00:600
```

```
46:1:1:101::10
```

```
R6#show ipv6 inter bri e0/1
```

```
Ethernet0/1 [up/up]
```

```
FE80::A8BB:CCFF:FE00:610
```

```
46:1:1:103::1
```

```
R6#show ipv6 route
```

```
IPv6 Routing Table - default - 9 entries
```

```
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
```

```
B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
D - EIGRP, EX - EIGRP external, NM - NEMO, ND - neighbor Discovery
```

```
l - LISP
```

```
O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
ND ::/0 [2/0]
```

```
via FE80::A8BB:CCFF:FE00:400, Serial2/0
```

```
NDp 46::/64 [2/0]
```

```
via Serial2/0, directly connected
```

```
L 46::A8BB:CCFF:FE00:600/128 [0/0]
```

```
via Serial2/0, receive
```

```
S 46:1:1:100::/60 [1/0]
```

```
via Null0, directly connected
```

```
C 46:1:1:101::/64 [0/0]
  via Ethernet0/0, directly connected
L 46:1:1:101::10/128 [0/0]
  via Ethernet0/0, receive
C 46:1:1:103::/64 [0/0]
  via Ethernet0/1, directly connected
L 46:1:1:103::1/128 [0/0]
  via Ethernet0/1, receive
L FF00::/8 [0/0]
  via Null0, receive
```

On R7:

```
R7(config)#int e0/1
R7(config-if)#ipv6 enable
R7(config-if)#ipv6 address autoconfig default
R7(config-if)#no shut
```

On R8:

```
R8(config)#int e0/0
R8(config-if)#ipv6 enable
R8(config-if)#ipv6 address autoconfig default
R8(config-if)#no shut
```

To verify the configuration:

On R7:

```
R7#show ipv6 inter bri e0/1
Ethernet0/1 [up/up]

FE80::A8BB:CCFF:FE00:710
46:1:1:103:A8BB:CCFF:FE00:710
```

```
R7#show ipv6 route
```

```
IPv6 Routing Table - default - 4 entries
```

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
ld - LISP dyn-EID, a - Application

ND ::/0 [2/0]

via FE80::A8BB:CCFF:FE00:610, Ethernet0/1

NDp 46:1:1:103::/64 [2/0]

via Ethernet0/1, directly connected

L 46:1:1:103:A8BB:CCFF:FE00:710/128 [0/0]

via Ethernet0/1, receive

L FF00::/8 [0/0]

via Null0, receive

R7#ping 46:1:1:101::10

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 46:1:1:101::10, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/12 ms

R7#ping 46:2::4

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 46:2::4, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

On R8:

R8#show ipv6 inter br e0/0

Ethernet0/0 [up/up]

FE80::A8BB:CCFF:FE00:800

46:1:1:101:A8BB:CCFF:FE00:800

R8#show ipv6 route

IPv6 Routing Table - default - 4 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP

H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea

IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO

ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect

O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2

ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site

ld - LISP dyn-EID, a - Application

ND ::/0 [2/0]

via FE80::A8BB:CCFF:FE00:600, Ethernet0/0

NDp 46:1:1:101::/64 [2/0]

via Ethernet0/0, directly connected

L 46:1:1:101:A8BB:CCFF:FE00:800/128 [0/0]

via Ethernet0/0, receive

L FF00::/8 [0/0]

via Null0, receive

R8#ping 46:1:1:101::10

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 46:1:1:101::10, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/8 ms

R8#ping 46::4

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 46::4, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms

On R4:

R4#show ipv6 local pool

Pool	Prefix	Free	In use
test	46:1:1:100::/56	15	1

NOTE: Since the customer received 46:1:1:100::/60 the customer can have 2⁴ or 16 possible subnets, and since we have used one we still have 15 free.

Task 7

Erase the startup configuration of the routers and reload them before proceeding to the next lab.


Advanced CCIE Routing & Switching v5.0

www.MicronicsTraining.com

Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

DMVPN & IPv6 in IP

Lab 1 – Configuring IPv6 in IPv4 Tunnel



Lab Setup:

To copy and paste the initial configurations, go to “Advanced-init” → “IPv6” → “Lab-1”.

Task 1

Configure OSPFv3 Area 0 on the link connecting R1 to R2 and their lo0 interfaces. The loopback interfaces should be advertised with their correct mask. The OSPF RID should be set to “0.0.0.x” where “x” is the router number. The link local IPv6 addresses should be set to FE80::x, where “x” is the router number.

On R1:

```
R1 (config) #ipv6 unicast-routing
```

```
R1 (config) #ipv6 router ospf 1
R1 (config-rtr) #router-id 0.0.0.1
```

```
R1 (config) #int lo0
R1 (config-if) #ipv6 ospf network point-to-point
R1 (config-if) #ipv6 ospf 1 area 0
```

```
R1 (config) #int s1/2
R1 (config-if) #ipv6 address fe80::1 link
R1 (config-if) #ipv6 ospf 1 area 0
```

On R2:

```
R2 (config) #ipv6 unicast-routing
```

```
R2 (config) #ipv6 router ospf 1
R2 (config-rtr) #router-id 0.0.0.2
```

```
R2 (config) #int lo0
R2 (config-if) #ipv6 ospf network point-to-point
R2 (config-if) #ipv6 ospf 1 area 0
```

```
R2 (config) #int s1/1
R2 (config-if) #ipv6 address fe80::2 link
R2 (config-if) #ipv6 ospf 1 area 0
```

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Serial1/1 from LOADING to FULL, Loading Done
```

```
R2#show ipv6 route ospf
```

```
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       D - EIGRP, EX - EIGRP external, NM - NEMO, ND - neighbor Discovery
       l - LISP
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
O    1::/64 [110/782]
    via FE80::1, Serial1/1
```

To verify the configuration:

On R2:

```
R2#ping 1::1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1::1, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms

```
R2#ping 1::1 source lo0
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1::1, timeout is 2 seconds:

Packet sent with a source address of 2::2

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms

Task 2

Configure OSPFv3 on the link connecting R3 to R4 and advertise their serial and loopback interfaces in Area 0. The loopback interfaces should be advertised with their correct mask.

On R3:

```
R3(config)#ipv6 unicast-routing
```

```
R3(config)#ipv6 router ospf 1
```

```
R3(config-rtr)#router-id 0.0.0.3
```

```
R3(config)#int lo0
```

```
R3(config-if)#ipv6 ospf network point-to-point
```

```
R3(config-if)#ipv6 ospf 1 area 0
```

```
R3(config)#int s2/0
```

```
R3(config-if)#ipv6 address fe80::3 link
```

```
R3(config-if)#ipv6 ospf 1 area 0
```

On R4:

```
R4(config)#ipv6 unicast-routing
```

```
R4(config)#ipv6 router ospf 1
R4(config-rtr)#router-id 0.0.0.4
```

```
R4(config)#int lo0
R4(config-if)#ipv6 ospf network point-to-point
R4(config-if)#ipv6 ospf 1 area 0
```

```
R4(config)#int s1/3
R4(config-if)#ipv6 address fe80::4 link
R4(config-if)#ipv6 ospf 1 area 0
```

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Serial1/3 from LOADING to FULL, Loading Done
```

To verify the configuration:

On R4:

```
R4#show ipv6 route ospf
```

```
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
       lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
```

```
O 3::/64 [110/65]
via FE80::3, Serial1/3
```

```
R4#ping 3::3
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 3::3, timeout is 2 seconds:

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

```
R4#ping 3::3 source lo0
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 3::3, timeout is 2 seconds:

Packet sent with a source address of 4::4

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms

Task 3

Configure a tunnel between R2 and R3 based on the following policy:

- R2's tunnel interface should be configured with an IPv6 address of 23::2/64, and R3's tunnel interface with an IPv6 address of 23::3/64.
- The tunnel mode should be IPv6.
- Configure OSPFv3 Area 0 on tunnel interface. If this configuration is done properly, every router should have an IPv6 NLRI to every other router.

The source address of the tunnel is the ip address assigned to the Serial1/3 interface, but the tunnel's address is an IPv6 address of 23::2/64, in this scenario ip is the transport and the guest protocol is IPv6.

On R2:

```
R2(config)#int tunn 23
R2(config-if)#ipv6 addr 23::2/64
R2(config-if)#ipv6 addr fe80::2 link-local
```

NOTE: The tunnel destination is the ip address of R3's S1/2 interface:

```
R2(config-if)#tunnel source s1/3
R2(config-if)#tunnel destination 23.1.1.3
```

Finally, the tunnel mode is configured to IPv6IP.

```
R2(config-if)#tunnel mode ipv6ip
```

On R3:

```
R3(config)#int tunn 32
R3(config-if)#ipv6 addr 23::3/64
R3(config-if)#ipv6 addr fe80::3 Link-local
R3(config-if)#tunnel source s1/2
```



```
R3(config-if)#tunnel destination 23.1.1.2
R3(config-if)#tunnel mode ipv6ip
```

To verify the configuration:

On R3:

```
R3#ping 23::2
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 23::2, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 36/36/36 ms

Let's run OSPFv3:

```
R3(config)#int tunn 32
R3(config-if)#ipv6 ospf 1 area 0
```

On R2:

```
R2(config)#int tunn 23
R2(config-if)#ipv6 ospf 1 area 0
```

You should see the following console message:

On R2:

```
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Tunnel23 from LOADING to FULL, Loading Done
```

To verify the configuration:

On R1:

```
R1#show ipv6 route ospf
```

IPv6 Routing Table - default - 10 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP

H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea

IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO

ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect

O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2

```
ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
O  2::/64 [110/65]
   via FE80::2, Serial1/2
O  3::/64 [110/1065]
   via FE80::2, Serial1/2
O  4::/64 [110/1129]
   via FE80::2, Serial1/2
O  23::/64 [110/1064]
   via FE80::2, Serial1/2
O  34::/64 [110/1128]
   via FE80::2, Serial1/2
```

R1#ping 2::2

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 2::2, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms

R1#ping 3::3

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 3::3, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 60/63/64 ms

R1#ping 4::4

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 4::4, timeout is 2 seconds:

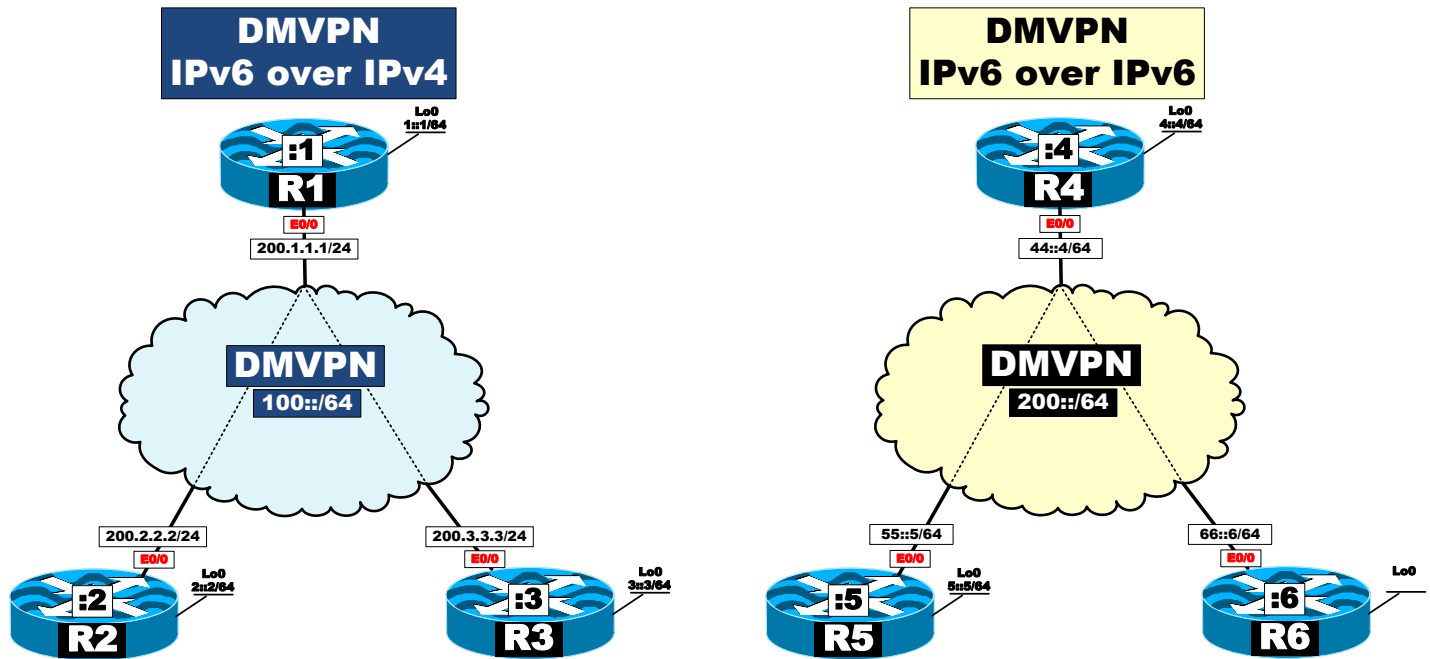
!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 88/88/92 ms

Task 4

Erase the startup configuration of the routers and reload before proceeding to the next lab.

Lab 2 – DMVPN & IPv6



Task 1

SW1 represents the Internet. Configure a static default route on each router pointing to the appropriate interface on SW1. If this configuration is performed correctly, these routers should be able to ping and have reachability to the E0/0 interfaces of all routers in this topology. The switch interface to which the routers are connected to should have a “.10” in the host portion of the IP address for that subnet. Shutdown all the ports on SW1 and set the duplex mode to half. ONLY use the ports in the above topology.

Let's configure SW1's interfaces for these routers. Since SW1 represents the Internet, the IP addresses in the following configuration should be configured as the default gateway on the routers .

On SW1:

```
SW1 (config) #int range e0/0-3,e1/0-3,e2/0-3,e3/0-3,e4/0-3,e5/0-3
```

```
SW1 (config-if-range) #duplex half
SW1 (config-if-range) #no swi
SW1 (config-if-range) #shut

SW1 (config) #int e0/1
SW1 (config-if) #ip addr 200.1.1.10 255.255.255.0
SW1 (config-if) #no shut

SW1 (config) #int e0/2
SW1 (config-if) #ip addr 200.2.2.10 255.255.255.0
SW1 (config-if) #no shut

SW1 (config) #int e0/3
SW1 (config-if) #ip addr 200.3.3.10 255.255.255.0
SW1 (config-if) #no shut
```

Let's NOT forget to enable "ip routing" or else the switch will not be able to route from one subnet to another.

```
SW1 (config) #ip routing
```

Let's configure the routers:

On R1:

```
R1 (config) #int e0/0
R1 (config-if) #ip addr 200.1.1.1 255.255.255.0
R1 (config-if) #no shut

R1 (config) #ip route 0.0.0.0 0.0.0.0 200.1.1.10
```

On R2:

```
R2 (config) #int e0/0
R2 (config-if) #ip addr 200.2.2.2 255.255.255.0
R2 (config-if) #no shut

R2 (config) #ip route 0.0.0.0 0.0.0.0 200.2.2.10
```

On R3:

```
R3 (config) #int e0/0
R3 (config-if) #ip addr 200.3.3.3 255.255.255.0
R3 (config-if) #no shut
```

```
R3(config)#ip route 0.0.0.0 0.0.0.0 200.3.3.10
```

To verify the configuration:

On R1:

```
R1#ping 200.2.2.2
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 200.2.2.2, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

```
R1#ping 200.3.3.3
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 200.3.3.3, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

On R2:

```
R2#ping 200.3.3.3
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 200.3.3.3, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Task 2

Configure DMVPN such that R1 is the HUB. R2, R3, and R4 are configured as the SPOKES. You should use 100::x/64, where “x” is the router number for the tunnel interfaces of these three routers. Configure the link local IPv6 address of these routers to FE80::x, where “x” is the router number. If this configuration is performed correctly, these routers should have reachability to all tunnel end points. You should configure static mappings. The tunnel interfaces should be configured in multipoint manner to accomplish this task.

Since the source of the tunnel interface is an ip address (The NBMA-IP), the transport protocol is ip and IPv6 is configured as a guest protocol.

On R1:

Let's configure the tunnel interface of R1:

```
R1 (config) #int tunn 1
R1 (config-if) #ipv6 address 100::1/64
R1 (config-if) #ipv6 address fe80::1 link-local
```

The tunnel source or the NBMA-IP is based on IP:

```
R1 (config-if) #tunnel source e0/0
R1 (config-if) #tunnel mode gre multipoint
```

Regular NHRP configuration, NHRP is enabled for IPv6:

```
R1 (config-if) #ipv6 nhrp network 111
```

The NHRP mapping is done based on the Link-local IPv6 address of the spokes because in IPv6, the routing protocols use the Link-local IPv6 addresses.

NOTE: In DMVPN, the tunnel-IP is mapped to the NBMA-IP. This is exactly how we configure IP over DMVPN the only difference is that if the tunnel's IPv6 address is mapped to the NBMA-IP address the routing protocols will not establish an adjacency, therefore, instead of using the tunnel IPv6 address, we have used the link-local IPv6 address and mapped it to the NBMA-IP.

```
R1 (config-if) #ipv6 nhrp map fe80::2/128 200.2.2.2
R1 (config-if) #ipv6 nhrp map fe80::3/128 200.3.3.3
```

Finally, Multicast is mapped to the NBMA-IP to provide Multicast capability so we can run OSPF.

```
R1 (config-if) #ipv6 nhrp map multicast 200.2.2.2
R1 (config-if) #ipv6 nhrp map multicast 200.3.3.3
```

On R2:

```
R2 (config) #int tunnel 1
R2 (config-if) #ipv6 address 100::2/64
R2 (config-if) #ipv6 address fe80::2 Link-local

R2 (config-if) #tunnel source e0/0
R2 (config-if) #tunnel mode gre multipoint

R2 (config-if) #ipv6 nhrp network 222
R2 (config-if) #ipv6 nhrp map fe80::3/128 200.3.3.3
```

```
R2(config-if)#ipv6 nhrp map fe80::1/128 200.1.1.1
R2(config-if)#ipv6 nhrp map multicast 200.1.1.1
R2(config-if)#ipv6 nhrp map multicast 200.3.3.3
```

On R3:

```
R3(config)#int tunnel 1
R3(config-if)#ipv6 address 100::3/64
R3(config-if)#ipv6 address fe80::3 link-local
```

```
R3(config-if)#tunnel source e0/0
R3(config-if)#tunnel mode gre multipoint
```

```
R3(config-if)#ipv6 nhrp network 333
R3(config-if)#ipv6 nhrp map fe80::1/128 200.1.1.1
R3(config-if)#ipv6 nhrp map fe80::2/128 200.2.2.2
R3(config-if)#ipv6 nhrp map multicast 200.1.1.1
R3(config-if)#ipv6 nhrp map multicast 200.2.2.2
```

To verify the configuration:

On R1:

```
R1#ping fe80::2
```

Output Interface: **tunnell1**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 100::2, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 0/3/4 ms

```
R1#ping fe80::3
```

Output Interface: **tunnell1**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 100::3, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms

On R2:

```
R2#ping fe80::1
```

```
Output Interface: tunnel1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 100::1, timeout is 2 seconds:
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
```

```
R2#ping fe80::3
```

```
Output Interface: tunnel1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 100::3, timeout is 2 seconds:
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
```

On R3:

```
R3#ping fe80::1
```

```
Output Interface: tunnel1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 100::1, timeout is 2 seconds:
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
```

```
R3#ping fe80::2
```

```
Output Interface: tunnel1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 100::2, timeout is 2 seconds:
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/4 ms
```

Task 3

Configure the loopback0 interfaces of these three routers based on the topology. Configure OSPFv3 Area 0 on the loopback0 and the tunnel 1 interfaces of these three routers. If this configuration is done successfully, these routers should have reachability to the IPv6 addresses of the other routers. The router-id of these routers should be 0.0.0.x, where "x" is the router number. The loopback interfaces should be advertised

with their correct mask. The OSPF network type for the tunnel interfaces should be configured as broadcast.

Let's run OSPFv3 on the tunnel and loopback 0 interfaces of these routers:

On R1:

```
R1 (config) #ipv6 unicast-routing

R1 (config) #ipv6 router ospf 1
R1 (config-rtr) #router-id 0.0.0.1

R1 (config) #int lo0
R1 (config-if) #ipv6 address 1::1/64
R1 (config-if) #ipv6 ospf 1 area 0
R1 (config-if) #ipv6 ospf network point-to-point

R1 (config) #int tunnel 1
R1 (config-if) #ipv6 ospf 1 area 0
R1 (config-if) #ipv6 ospf network broadcast
```

On R2:

```
R2 (config) #ipv6 unicast-routing

R2 (config) #ipv6 router ospf 1
R2 (config-rtr) #router-id 0.0.0.2

R2 (config) #int lo0
R2 (config-if) #ipv6 address 2::2/64
R2 (config-if) #ipv6 ospf 1 area 0

R2 (config-if) #ipv6 ospf network point-to-point

R2 (config) #int tunnel 1
R2 (config-if) #ipv6 ospf 1 area 0
```

The spoke routers should always be configured with a priority of 0 so they do not participate in DR election.

```
R2 (config-if) #ipv6 ospf network broadcast
R2 (config-if) #ipv6 ospf priority 0
```

On R3:

```
R3(config)#ipv6 unicast-routing
```

```
R3(config)#ipv6 router ospf 1  
R3(config-rtr)#router-id 0.0.0.3
```

```
R3(config)#int lo0  
R3(config-if)#ipv6 address 3::3/64  
R3(config-if)#ipv6 ospf network point-to-point  
R3(config-if)#ipv6 ospf 1 area 0
```

```
R3(config)#int tunnel 1  
R3(config-if)#ipv6 ospf 1 area 0  
R3(config-if)#ipv6 ospf network broadcast  
R3(config-if)#ipv6 ospf priority 0
```

On All Routers:

```
Rx#clear ipv6 ospf pro  
Reset ALL OSPF processes? [no]: y
```

To verify the configuration:

On R1:

You should see the following console messages:

```
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Tunnel1 from LOADING to FULL,  
Loading Done
```

```
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Tunnel1 from LOADING to FULL,  
Loading Done
```

```
R1#show ipv6 route ospf
```

```
IPv6 Routing Table - default - 7 entries  
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route  
B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP  
H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea  
IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO  
ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect  
O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2  
ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt  
lr - LISP site-registrations, ld - LISP dyn-eid, a - Application  
O 2::/64 [110/1001]  
via FE80::2, Tunnel1
```

```
O 3::/64 [110/1001]
   via FE80::3, Tunnel1
```

```
R1#ping 2::2 source lo0
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 2::2, timeout is 2 seconds:

Packet sent with a source address of 1::1

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms

```
R1#ping 3::3 source lo0
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 3::3, timeout is 2 seconds:

Packet sent with a source address of 1::1

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms

On R2:

```
R2#show ipv6 route ospf
```

IPv6 Routing Table - default - 7 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP

H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea

IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO

ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect

O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2

ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt

lr - LISP site-registrations, ld - LISP dyn-eid, a - Application

```
O 1::/64 [110/1001]
   via FE80::1, Tunnel1
```

```
O 3::/64 [110/1001]
   via FE80::3, Tunnel1
```

```
R2#ping 1::1 source lo0
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 1::1, timeout is 2 seconds:

Packet sent with a source address of 2::2

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/0 ms

```
R2#ping 3::3 source lo0
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3::3, timeout is 2 seconds:
Packet sent with a source address of 2::2
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/0 ms
```

```
R2#trace
Protocol [ip]: ipv6
Target IPv6 address: 3::3
Source address: 2::2
Insert source routing header? [no]:
Numeric display? [no]:
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Priority [0]:
Port Number [0]:
Type escape sequence to abort.
Tracing the route to 3::3

  1 100::3 2 msec 1 msec 1 msec
```

To summarize:

When configuring IPv6 over DMVPN and the underlying protocol is IPv4, the NHRP mapping should be configured based on the link local IPv6 addresses, if the mapping is configured for the global unicast IPv6 addresses, the routing protocols will not form an adjacency, this is because the routing protocols use the link local IPv6 addresses to form an adjacency.

In the above configuration, the routers will not have reachability to the global unicast IPv6 address of the tunnel interfaces. If this reachability is required, then, an NHRP mapping must be configured for the global unicast IPv6 address of the tunnel interfaces.

Task 4

SW1 represents the Internet. Configure an IPv6 static default route on R4, R5 and R6 pointing to the appropriate interface on SW1. If this configuration is performed correctly, these routers should be able to ping and have reachability to the IPv6 address of the E0/0 interfaces of the other routers in this topology. The switch interface to which the routers are connected to should be configured with a link local IPv6 address of fe80::10 and a global unicast IPv6 address of 44::10, 55::10, and 66::10 for the connections to R4, R5 and R6 respectively.

On SW1:

```
SW1 (config) #ipv6 unicast-routing

SW1 (config) #int e1/0
SW1 (config-if) #no swi
SW1 (config-if) #ipv6 address fe80::10 link
SW1 (config-if) #ipv6 address 44::10/64
SW1 (config-if) #no shut

SW1 (config) #int e1/1
SW1 (config-if) #no swi
SW1 (config-if) #ipv6 address fe80::10 link
SW1 (config-if) #ipv6 address 55::10/64
SW1 (config-if) #no shu

SW1 (config) #int e1/2
SW1 (config-if) #no swi
SW1 (config-if) #ipv6 address fe80::10 link
SW1 (config-if) #ipv6 address 66::10/64
SW1 (config-if) #no shu
```

On R4:

```
R4 (config) #int e0/0
R4 (config-if) #ipv6 enable
R4 (config-if) #ipv6 address 44::4/64
R4 (config-if) #ipv6 address fe80::4 link
R4 (config-if) #no shut

R4 (config) #ipv6 route ::/0 e0/0 fe80::10
```

On R5:

```
R5 (config) #int e0/0
R5 (config-if) #ipv6 enable
R5 (config-if) #ipv6 address 55::5/64
R5 (config-if) #ipv6 address fe80::5 link
R5 (config-if) #no shut

R5 (config) #ipv6 route ::/0 e0/0 fe80::10
```

On R6:

```
R6(config)#int e0/0
R6(config-if)#ipv6 enable
R6(config-if)#ipv6 address fe80::6 link
R6(config-if)#ipv6 address 66::6/64
R6(config-if)#no shut

R6(config)#ipv6 route ::/0 e0/0 fe80::10
```

To verify the configuration:

On R4:

```
R4#ping 55::5
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 55::5, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

```
R4#ping 66::6
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 66::6, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/19 ms

Task 2

Configure the second DMVPN network using the following policy:

- R4, R5, and R6 must use the IPv6 address of their E0/0 interfaces as the source of the tunnel.
- The tunnel mode should be configured as “**GRE Multipoint**”.
- The ipv6 address of the tunnel interfaces should be configured as 200::x/64 where “x” is the router number.
- Configure the IPv6 address of the loopback interfaces based on the topology, these loopback interfaces must be advertised with their correct mask.
- You should run OSPFv3 Area 0 on the tunnel and their Loopback0 interfaces. The loopback interfaces should be advertised with their correct mask. The OSPF network type of the tunnel interface should be configured as Point-to-Multipoint

on the hub and Point-to-Point on the spokes. You must use address-family to accomplish this task.

- R4 should be configured to be the NHS.
- This DMVPN network must be configured in phase 3

NOTE: In this task the source of the tunnel is the IPv6 address of the E0/0 interface, therefore, in this scenario IPv6 is both the transport and the guest protocol.

On R4:

```
R4 (config) #int tunn 1
R4 (config-if) #ipv6 address 200::4/64
R4 (config-if) #ipv6 address fe80::4 link-local
```

The only difference between this configuration and the configuration from the previous MPLS cloud is the “IPv6” keyword at the end of the “tunnel mode” command. If this keyword is not used, the tunnel interface will not come up.

```
R4 (config-if) #tunnel source e0/0
R4 (config-if) #tunnel mode gre multipoint ipv6
```

The NHRP protocol is enabled and Multicast is mapped dynamically.

```
R4 (config-if) #ipv6 nhrp network 444
R4 (config-if) #ipv6 nhrp map multicast dynamic
R4 (config-if) #ipv6 nhrp redirect
```

On R5:

```
R5 (config) #int tunn 1
R5 (config-if) #ipv6 address 200::5/64
R5 (config-if) #ipv6 address fe80::5 link-local
```

```
R5 (config-if) #tunnel source e0/0
R5 (config-if) #tunnel mode gre multipoint ipv6
```

```
R5 (config-if) #ipv6 nhrp network 555
R5 (config-if) #ipv6 nhrp nhs 200::4 nbma 44::4 multicast
R5 (config-if) #ipv6 nhrp shortcut
```

On R6:

```
R6 (config) #int tunn 1
R6 (config-if) #ipv6 address 200::6/64
```

```
R6 (config-if) #ipv6 address fe80::6 link-local
R6 (config-if) #tunnel source e0/0
R6 (config-if) #tunnel mode gre multipoint ipv6

R6 (config-if) #ipv6 nhrp network 666
R6 (config-if) #ipv6 nhrp nhs 200::4 nbma 44::4 multicast
R6 (config-if) #ipv6 nhrp shortcut
```

To verify the configuration:

On R4:

We can see that both the IPv6 address of the tunnel interface, and the link-local IPv6 addresses are dynamically mapped:

```
R7#show ipv6 nhrp

200::5/128 via 200::5
  Tunnel1 created 00:03:23, expire 01:56:36
  Type: dynamic, Flags: unique registered used nhop
  NBMA address: 55::5
200::6/128 via 200::6
  Tunnel1 created 00:00:35, expire 01:59:24
  Type: dynamic, Flags: unique registered used nhop
  NBMA address: 66::6
FE80::5/128 via 200::5
  Tunnel1 created 00:03:23, expire 01:56:36
  Type: dynamic, Flags: unique registered
  NBMA address: 55::5
FE80::6/128 via 200::6
  Tunnel1 created 00:00:35, expire 01:59:24
  Type: dynamic, Flags: unique registered
  NBMA address: 66::6
```

On R5:

```
R5#sh ipv6 nhrp

200::4/128 via 200::4
  Tunnel1 created 00:05:28, never expire
  Type: static, Flags: used
  NBMA address: 44::4
FE80::4/128 via FE80::4
  Tunnel1 created 00:05:28, never expire
```



```
Type: static, Flags: nhs-ll
NBMA address: 44::4
```

On R6:

```
R6#sh ipv6 nhrp
```

```
200::4/128 via 200::4
  Tunnell created 00:03:22, never expire
  Type: static, Flags: used
  NBMA address: 44::4
FE80::4/128 via FE80::4
  Tunnell created 00:03:22, never expire
  Type: static, Flags: nhs-ll
  NBMA address: 44::4
```

On R4:

```
R4#ping 200::5
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200::5, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

```
R4#ping 200::6
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200::6, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Let's run OSPFv3 area 0 on these two routers:

On R4:

```
R4(config)#ipv6 unicast-routing

R4(config)#int lo0
R4(config-if)#ipv6 address 4::4/64
R4(config-if)#ospfv3 network point-to-point

R4(config)#router ospfv3 1
R4(config-router)#address-family ipv6 unicast
R4(config-router-af)#router-id 0.0.0.4
```

```
R4 (config) #int lo0
R4 (config-if) #ospfv3 1 ipv6 area 0
```

```
R4 (config) #int tunnel 1
R4 (config-if) #ospfv3 1 ipv6 area 0
R4 (config-if) #ospfv3 network point-to-multipoint
```

On R5:

```
R5 (config) #ipv6 unicast-routing
```

```
R5 (config) #int lo0
R5 (config-if) #ipv6 address 5::5/64
R5 (config-if) #ospfv3 network point-to-point
```

```
R5 (config) #router ospfv3 1
R5 (config-router) #address-family ipv6 unicast
R5 (config-router-af) #router-id 0.0.0.5
```

```
R5 (config) #int lo0
R5 (config-if) #ospfv3 1 ipv6 area 0
```

```
R5 (config-if) #int tunn 1
R5 (config-if) #ospfv3 1 ipv6 area 0
R5 (config-if) #ospfv3 1 netw point-to-multi
```

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 0.0.0.4 on Tunnel1 from LOADING to FULL, Loading Done
```

On R6:

```
R6 (config) #ipv6 unicast-routing
```

```
R6 (config) #int lo0
R6 (config-if) #ipv6 address 6::6/64
R6 (config-if) #ospfv3 network point-to-point
```

```
R6 (config) #router ospfv3 1
R6 (config-router) #address-family ipv6 unicast
R6 (config-router-af) #router-id 0.0.0.6
```

```
R6 (config) #int lo0
R6 (config-if) #ospfv3 1 ipv6 area 0
```

```
R6(config-if)#int tunn 1
R6(config-if)#ospfv3 1 ipv6 area 0

R6(config-if)#ospfv3 1 netw point-to-multi
```

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 0.0.0.4 on Tunnell from LOADING to FULL, Loading Done
```

To verify the configuration:

On R4:

```
R4#show ipv6 route ospf
```

```
IPv6 Routing Table - default - 12 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
       lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
O  5::/64 [110/1001]
   via FE80::5, Tunnell
O  6::/64 [110/1001]
   via FE80::6, Tunnell
O  200::5/128 [110/1000]
   via FE80::5, Tunnell
O  200::6/128 [110/1000]
   via FE80::6, Tunnell
```

```
R4#ping 5::5
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 5::5, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

```
R4#ping 200::5
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 200::5, timeout is 2 seconds:

```
!!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
R4#ping 6::6
```

```
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 6::6, timeout is 2 seconds:
```

```
!!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
R4#ping 200::6
```

```
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 200::6, timeout is 2 seconds:
```

```
!!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Task 3

Erase the startup configuration of the routers and reload before proceeding to the next lab.


Advanced CCIE Routing & Switching v5.0

www.MicronicsTraining.com

Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

OSPFv3

Lab 1 – Configuring OSPFv3



Task 1

Configure the above topology. DO NOT configure any routing protocol.

On SW1:

```
SW1 (config) #int rang e0/0-3,e1/0-3,e2/0-3,e3/0-3,e4/0-3,e5/0-3
SW1 (config-if-range) #duplex hal
SW1 (config-if-range) #shut

SW1 (config) #int range e0/1-2
SW1 (config-if-range) #swi
SW1 (config-if-range) #swi mode acc
SW1 (config-if-range) #swi acc v 12
SW1 (config-if-range) #spanning portf
SW1 (config-if-range) #no shut
```

On R1:

```
R1 (config) #int e0/0
R1 (config-if) #ipv6 address 12::1/64
R1 (config-if) #ipv6 address fe80::1 link-local
R1 (config-if) #no shut
```

```
R1 (config) #int lo0
R1 (config-if) #ipv6 address 1::1/64
```

On R2:

```
R2 (config) #int e0/0
R2 (config-if) #ipv6 address 12::2/64
R2 (config-if) #ipv6 address fe80::2 link-local
R2 (config-if) #no shut
```

To verify the configuration:

On R2:

```
R2#ping 12::1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 12::1, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/15 ms

```
R2 (config) #int lo0
R2 (config-if) #ipv6 address 2::2/64
```

```
R2 (config-if) #int lo1
R2 (config-if) #ip address 2.2.2.2 255.255.255.0
```

```
R2 (config) #int s2/0
R2 (config-subif) #ip address 24.1.1.2 255.255.255.0
R2 (config-if) #no shu
```

```
R2 (config) #int s1/3
R2 (config-subif) #ipv6 address 23::2/64
R2 (config-subif) #ipv6 address fe80::2 link-local
R2 (config-if) #no shu
```

On R3:

```
R3 (config) #int s1/2
```

```
R3(config-if)#ipv6 addr 23::3/64
R3(config-if)#ipv6 addr fe80::3 link-local
R3(config-if)#no shut
```

```
R3(config)#int lo0
R3(config-if)#ipv6 address 3::3/64
```

To verify the configuration:

On R3:

```
R3#ping 23::2
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 23::2, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms

On R4:

```
R4(config)#int s1/2
R4(config-if)#ip address 24.1.1.4 255.255.255.0
R4(config-if)#no shut
```

```
R4(config)#int lo0
R4(config-if)#ip address 4.4.4.4 255.255.255.0
```

To verify the configuration:

On R4:

```
R4#ping 24.1.1.2
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 24.1.1.2, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/10 ms

Task 2

Configure OSPF and OSPFv3 on the routers in this topology based on the following policy:

- Configure the E0/0 and loopback0 interfaces of R1 and R2 in Area 0.
- Configure R4 and its loopback interface in area 1.
- Configure R2's S2/0 interface in Area 1, and R2's S1/3 should be configured in Area 0.
- ONLY R2 and R4 should accomplish this task using an address-family.
- Configure R3's S1/2 and Lo0 interfaces in Area 0.
- The loopback interfaces should be configured with their correct mask.
- The router ID of these routers should be configured to be 0.0.0.x, where "x" is router number.

On R1:

In IPv4, unicast routing is enabled by default, but this is not the case in IPv6, and therefore it should be statically enabled before any routing protocol can be configured:

```
R1 (config) #ipv6 unicast-routing
```

```
R1 (config) #ipv6 router ospf 1
```

You should see the following console message stating that since OSPFv3 could not see a 32 bit dotted decimal value it could not assign a router-id. Let's statically assign a router-id to this process:

```
%OSPFv3-4-NORTRID: OSPFv3 process 1 could not pick a router-id, please  
configure manually
```

```
R1 (config-rtr) #router-id 0.0.0.1
```

```
R1 (config) #int lo0
```

```
R1 (config-if) #ipv6 ospf 1 area 0
```

```
R1 (config) #int e0/0
```

```
R1 (config-if) #ipv6 ospf 1 area 0
```

To verify the configuration:

```
R1#show ipv6 ospf interface brief
```

Interface	PID	Area	Intf ID	Cost	State	Nbrs	F/C
Lo0	1	0	18	1	LOOP	0/0	
Et0/0	1	0	3	10	WAIT	0/0	

```
R1#show ipv6 ospf inter lo0
```

```
Loopback0 is up, line protocol is up
  Link Local Address FE80::A8BB:CCFF:FE00:100, Interface ID 18
  Area 0, Process ID 1, Instance ID 0, router ID 0.0.0.1
  Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
```

We can see that the loopback0 interface fo R1 is treated as a host, this is the same behavior we saw in OSPFv2. Let's change the network type to Point-to-Point and verify:

```
R1(config)#int lo0
R1(config-if)#ipv6 ospf network point-to-point
```

To verify the configuration:

On R1:

```
R1#show ipv6 ospf int lo0
```

```
Loopback0 is up, line protocol is up
  Link Local Address FE80::A8BB:CCFF:FE00:100, Interface ID 18
  Area 0, Process ID 1, Instance ID 0, router ID 0.0.0.1
  Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Graceful restart helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
```

On R4:

OSPF can be configured in three ways: OSPFv2 configuration, OSPFv3 configuration and OSPFv3 using address-family. The benefit of configuring OSPFv3 using the address-family is the fact that this configuration can handle IPv4 and IPv6 neighbors. As you can see the "IPv6 unicast-routing" **MUST** be enabled.

```
R4 (config) #ipv6 unicast-routing
```

The configuration starts with “router ospfv3” followed by the process id:

```
R4 (config) #router ospfv3 1
```

The router-id is configured directly under the address-family:

```
R4 (config-router) #address-family ipv4 unicast  
R4 (config-router) #router-id 0.0.0.4
```

Running OSPFv3 on the interfaces is done by configuring “OSPFv3” followed by the process-id “1”, followed by “IPv4”, if this was for IPv6, then the “IPv6” keyword should be used, and then the “area” keyword is used ending with the area-id, in this case “1”.

```
R4 (config) #int lo0  
R4 (config-if) #ipv6 enable  
R4 (config-if) #ospfv3 1 ipv4 area 1  
R4 (config-if) #ospfv3 netw point-to-point
```

```
R4 (config) #int s1/2  
R4 (config-if) #ipv6 enable  
R4 (config-if) #ospfv3 1 ipv4 area 1
```

On R2:

```
R2 (config) #ipv6 unicast-routing
```

```
R2 (config) #router ospfv3 1  
R2 (config-router) #address ipv4 unicast  
R2 (config-router-af) #router-id 0.0.0.2
```

```
R2 (config-router-af) #address ipv6 unicast  
R2 (config-router-af) #router-id 0.0.0.2
```

```
R2 (config) #int lo0  
R2 (config-if) #ospfv3 1 ipv6 are 0
```

```
R2 (config-if) #int e0/0  
R2 (config-if) #ospfv3 1 ipv6 are 0
```

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 0.0.0.1 on Ethernet0/0 from  
LOADING to FULL, Loading Done
```

```
R2(config)#int s2/0
R2(config-if)#ospfv3 1 ipv4 area 1
% OSPFv3: IPV6 is not enabled on this interface
```

NOTE: Even though we are running IPv4 on the S2/0 and the Lo1 interfaces of R2, we need to enable IPv6, why?

OSPFv3 can handle IPv4 and IPv6, but to run OSPFv3 on an interface that is running pure IPv4, we must enable IPv6 because all the packets will be encapsulated in IPv6.

```
R2(config-if)#ipv6 enable
R2(config-if)#ospfv3 1 ipv4 area 1
```

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, IPv4, Nbr 0.0.0.4 on Serial2/0 from LOADING
to FULL, Loading Done
```

```
R2(config)#int s1/3
R2(config-if)#ipv6 ena
R2(config-if)#ospfv3 1 ipv6 area 0
```

```
R2(config)#int lo1
R2(config-if)#ipv6 ena
R2(config-if)#ospfv3 1 ipv4 area 1
```

To verify the configuration:

```
R2#sh ipv route ospf
```

```
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
       lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
```

```
O 1::/64 [110/11]
  via FE80::1, Ethernet0/0
```

```
R2#ping 1::1
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1::1, timeout is 2 seconds:
!!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 2/4/5 ms

On R3:

```
R3(config)#ipv6 unicast-routing
```

```
R3(config)#router ospfv3 1
```

```
R3(config-router)#address-family ipv6 unicast
```

```
R3(config-router-af)#router-id 0.0.0.3
```

```
R3(config)#int lo0
```

```
R3(config-if)#ospfv3 netw point-to-point
```

```
R3(config-if)#ospfv3 1 ipv6 area 0
```

```
R3(config)#int s1/2
```

```
R3(config-if)#ospfv3 1 ipv6 area 0
```

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 0.0.0.2 on Serial2/0 from LOADING  
to FULL, Loading Done
```

To verify the configuration:

```
R3#sh ipv6 route ospf
```

```
IPv6 Routing Table - default - 8 entries
```

```
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
```

```
      B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
      H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
```

```
      IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
```

```
      ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
```

```
      O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
      ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
```

```
      lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
```

```
O 1::/64 [110/75]
```

```
      via FE80::2, Serial1/2
```

```
O 2::2/128 [110/64]
```

```
      via FE80::2, Serial1/2
```

```
O 12::/64 [110/74]
```

```
      via FE80::2, Serial1/2
```

The highlighted area in the output of the above show command reveals the fact that R2's Lo0 is NOT advertised with its correct mask. Let's fix this problem and verify:

```
R2(config)#int lo0
```

```
R2 (config-if) #ospfv3 network point-to-point
```

```
R3#sh ipv6 route ospf
```

```
IPv6 Routing Table - default - 8 entries
```

```
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
```

```
      B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
      H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
```

```
      IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
```

```
      ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
```

```
      O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
      ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
```

```
      lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
```

```
O 1::/64 [110/75]
   via FE80::2, Serial1/2
```

```
O 2::/64 [110/65]
   via FE80::2, Serial1/2
```

```
O 12::/64 [110/74]
   via FE80::2, Serial1/2
```

```
R3#ping 2::2
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 2::2, timeout is 2 seconds:
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/9/10 ms
```

```
R3#ping 1::1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1::1, timeout is 2 seconds:
```

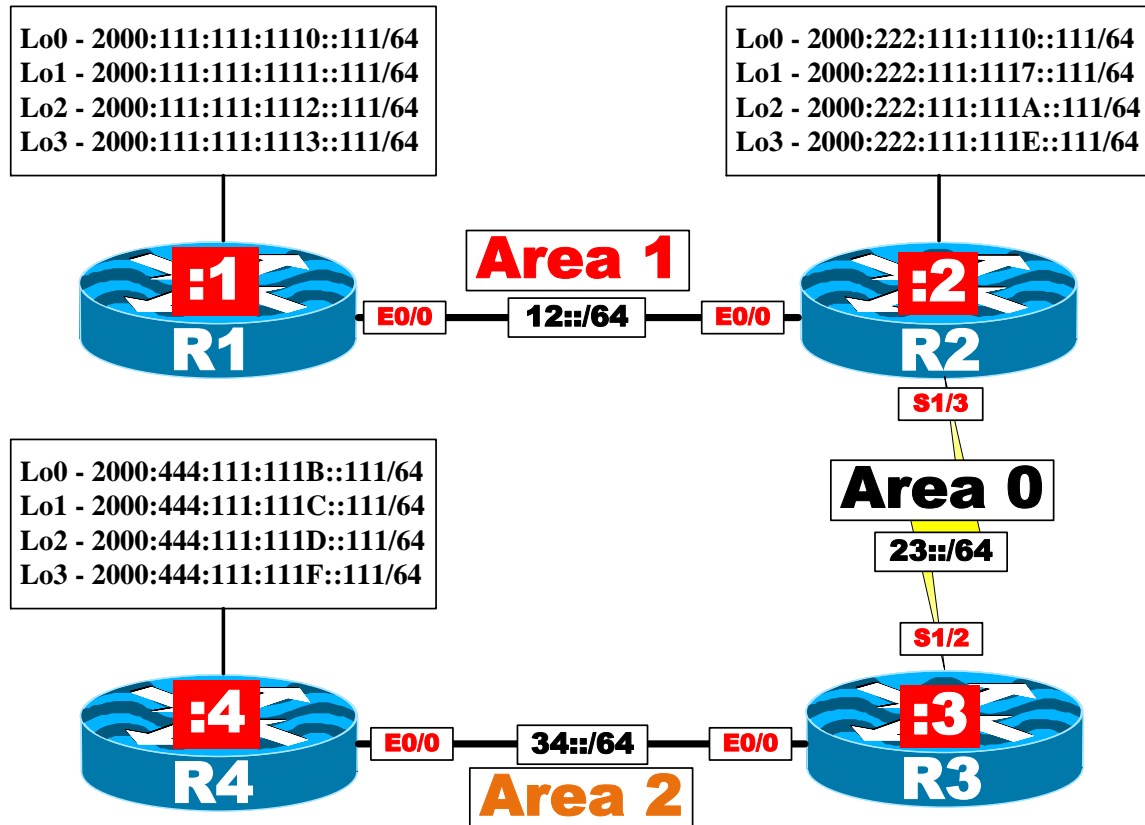
```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 9/9/10 ms
```

Task 3

Erase the startup configuration of the routers and reload before proceeding to the next lab.

Lab 2 – Summarization of Internal and External Networks



Task 1

Configure the above topology and shutdown the unused ports on SW1. DO NOT configure any routing protocol/s. The Link-local IPv6 addresses of these routers should be based on the following:

- R1 – FE80::1
- R2 – FE80::2
- R3 – FE80::3
- R4 – FE80::4

On SW1:

```
SW1(config)#int range e0/0-3,e1/0-3,e2/0-3,e4/0-3,e5/0-3
```

```
SW1 (config-if-range) #duplex half
SW1 (config-if-range) #shut
```

```
SW1 (config) #int range e0/1-2
SW1 (config-if-range) #swi
SW1 (config-if-range) #swi mode acc
SW1 (config-if-range) #swi acc v 12
SW1 (config-if-range) #spanning portf
SW1 (config-if-range) #no shut
```

```
SW1 (config) #int range e0/3,e1/0
SW1 (config-if-range) #swi
SW1 (config-if-range) #swi mode acc
SW1 (config-if-range) #swi acc v 34
SW1 (config-if-range) #spannin portf
SW1 (config-if-range) #no shut
```

On R1:

```
R1 (config) #int e0/0
R1 (config-if) #ipv6 address 12::1/64
R1 (config-if) #ipv6 address FE80::1 Link-Local
R1 (config-if) #no shut
```

```
R1 (config) #int lo0
R1 (config-if) #ipv6 addr 2000:111:111:1110::111/64
```

```
R1 (config) #int lo1
R1 (config-if) #ipv6 addr 2000:111:111:1111::111/64
```

```
R1 (config) #int lo2
R1 (config-if) #ipv6 addr 2000:111:111:1112::111/64
```

```
R1 (config) #int lo3
R1 (config-if) #ipv6 addr 2000:111:111:1113::111/64
```

On R2:

```
R2 (config) #int e0/0
R2 (config-if) #ipv6 address 12::2/64
R2 (config-if) #ipv6 address fe80::2 link-local
R2 (config-if) #no shut
```

```
R2 (config) #int s1/3
```



```
R2 (config-if) #ipv6 address 23::2/64
R2 (config-if) #ipv6 address fe80::2 Link-local
R2 (config-if) #no shut

R2 (config) #int lo0
R2 (config-if) #ipv6 addr 2000:222:111:1110::111/64

R2 (config) #int lo1
R2 (config-if) #ipv6 addr 2000:222:111:1117::111/64

R2 (config) #int lo2
R2 (config-if) #ipv6 addr 2000:222:111:111A::111/64

R2 (config) #int lo3
R2 (config-if) #ipv6 addr 2000:222:111:111E::111/64
```

To verify the configuration:

```
R2#ping 12::1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 12::1, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms

```
R2#show ipv6 neighbors
```

IPv6 Address	Age	Link-layer Addr	State	Interface
12::1	0	aabb.cc00.0100	REACH	Et0/0
FE80::1	0	aabb.cc00.0100	DELAY	Et0/0

On R3:

```
R3 (config) #int s1/2
R3 (config-if) #ipv6 address 23::3/64
R3 (config-if) #ipv6 address fe80::3 link-local
R3 (config-if) #no shut

R3 (config) #int e0/0
R3 (config-if) #ipv6 address 34::3/64
R3 (config-if) #ipv6 address fe80::3 link-local
R3 (config-if) #no shut
```

To verify the configuration:

```
R3#ping 23::3
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 23::3, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 0/2/4 ms

On R4:

```
R4(config)#int e0/0
```

```
R4(config-if)#ipv6 address 34::4/64
```

```
R4(config-if)#ipv6 address fe80::4 link-local
```

```
R4(config-if)#no shut
```

```
R4(config)#int lo0
```

```
R4(config-if)#ipv6 addr 2000:444:111:111B::111/64
```

```
R4(config)#int lo1
```

```
R4(config-if)#ipv6 addr 2000:444:111:111C::111/64
```

```
R4(config)#int lo2
```

```
R4(config-if)#ipv6 addr 2000:444:111:111D::111/64
```

```
R4(config)#int lo3
```

```
R4(config-if)#ipv6 addr 2000:444:111:111F::111/64
```

To verify the configuration:

```
R4#ping 34::3
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 34::3, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/8 ms

```
R4#show ipv6 neighbors
```

IPv6 Address	Age	Link-layer Addr	State	Interface
34::3	0	aabb.cc00.0300	REACH	Et0/0
FE80::3	0	aabb.cc00.0300	DELAY	Et0/0

Task 2

Configure OSPFv3 based on the following requirements:

- Configure OSPFv3 on R1 and run all its directly connected interfaces in area 1. DON'T use Address-family to configure this router. Configure the loopback interfaces with their correct mask. The RID of this router should be set to "0.0.0.1".
- Configure OSPFv3 on R2 using an Address-family. This router should run OSPFv3 area 1 on its E0/0 and OSPFv3 area 0 on its s1/3. The loopback interfaces of this router should be configured in area 0. Configure the loopback interfaces with their correct mask. The RID of this router should be set to "0.0.0.2".
- Configure OSPFv3 on R3 and run its s1/2 interface in area 0, and its E0/0 interface in area 2. DON'T use Address-family to configure this router. The RID of this router should be set to "0.0.0.3".
- Configure OSPFv3 on R4 using an Address-family. This router should run OSPFv3 area 2 on its E0/0 interface. The loopback interfaces of this router should be configured in OSPFv3 routing domain. The RID of this router should be set to "0.0.0.4".

On R1:

```
R1 (config) #ipv6 unicast-routing
```

```
R1 (config) #ipv6 router ospf 1  
R1 (config-rtr) #router-id 0.0.0.1
```

```
R1 (config) #int e0/0  
R1 (config-if) #ipv6 ospf 1 area 1
```

```
R1 (config) #int range Lo0-3  
R1 (config-if-range) #ipv6 ospf 1 area 1  
R1 (config-if-range) #ipv6 ospf network point-to-point
```

To verify the configuration:

```
R1#sh ipv ospf int br
```

Interface	PID	Area	Intf ID	Cost	State	Nbrs	F/C
Lo0	1	1	18	1	P2P	0/0	
Lo1	1	1	19	1	P2P	0/0	
Lo2	1	1	20	1	P2P	0/0	
Lo3	1	1	21	1	P2P	0/0	

```
Et0/0          1      1          3          10      WAIT    0/0
```

On R2:

```
R2 (config) #ipv6 unicast-routing
```

```
R2 (config) #router ospfv3 1
```

```
R2 (config-router) #address-family ipv6 unicast
```

```
R2 (config-router-af) #router-id 0.0.0.2
```

```
R2 (config) #int s1/3
```

```
R2 (config-if) #ospfv3 1 ipv6 area 0
```

```
R2 (config) #int e0/0
```

```
R2 (config-if) #ospfv3 1 ipv6 area 1
```

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 0.0.0.1 on Ethernet0/0 from  
LOADING to FULL, Loading Done
```

```
R2 (config) #int range lo0-3
```

```
R2 (config-if-range) #ospfv3 1 ipv6 area 0
```

```
R2 (config-if-range) #ospfv3 network Point-to-point
```

To verify the configuration:

On R2:

```
R2#show ipv6 route ospf
```

```
IPv6 Routing Table - default - 17 entries
```

```
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
```

```
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
```

```
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
```

```
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
```

```
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
```

```
       ld - LISP dyn-EID, a - Application
```

```
O 2000:111:111:1110::/64 [110/11]
```

```
   via FE80::1, Ethernet0/0
```

```
O 2000:111:111:1111::/64 [110/11]
```

```
   via FE80::1, Ethernet0/0
```

```
O 2000:111:111:1112::/64 [110/11]
```

```
via FE80::1, Ethernet0/0
O 2000:111:111:1113::/64 [110/11]
via FE80::1, Ethernet0/0
```

R2#**show ospfv3 neighbor**

OSPFv3 1 address-family ipv6 (router-id 0.0.0.7)

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
0.0.0.1	1	FULL/DR	00:00:37	3	Ethernet0/0

R2#**show ospfv3 inter br**

Interface	PID	Area	AF	Cost	State	Nbrs	F/C
Lo0	1	0	ipv6	1	P2P	0/0	
Lo1	1	0	ipv6	1	P2P	0/0	
Lo2	1	0	ipv6	1	P2P	0/0	
Lo3	1	0	ipv6	1	P2P	0/0	
Se1/3	1	0	ipv6	64	P2P	0/0	
Et0/0	1	1	ipv6	10	BDR	1/1	

On R3:

R3 (config) #**ipv6 unicast-routing**

R3 (config) #**ipv6 router ospf 1**
R3 (config-rtr) #**router-id 0.0.0.3**

R3 (config) #**int s1/2**
R3 (config-if) #**ipv6 ospf 1 area 0**

R3 (config) #**int e0/0**
R3 (config-if) #**ipv6 ospf 1 area 2**

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Serial1/2 from LOADING to FULL, Loading Done
```

To verify the configuration:

On R3:

R3#**show ipv6 route ospf**

```
IPv6 Routing Table - default - 14 entries
```

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
lr - LISP site-registrations, ld - LISP dyn-eid, a - Application

```
OI 12::/64 [110/74]
    via FE80::2, Serial1/2
OI 2000:111:111:1110::/64 [110/75]
    via FE80::2, Serial1/2
OI 2000:111:111:1111::/64 [110/75]
    via FE80::2, Serial1/2
OI 2000:111:111:1112::/64 [110/75]
    via FE80::2, Serial1/2
OI 2000:111:111:1113::/64 [110/75]
    via FE80::2, Serial1/2
O 2000:222:111:1110::/64 [110/65]
    via FE80::2, Serial1/2
O 2000:222:111:1117::/64 [110/65]
    via FE80::2, Serial1/2
O 2000:222:111:111A::/64 [110/65]
    via FE80::2, Serial1/2
O 2000:222:111:111E::/64 [110/65]
    via FE80::2, Serial1/2
```

On R4:

Since the task does not specify the area in which the loopback interfaces should be configured, the only other way to run them in OSPF is to redistribute them into OSPF routing domain. Let's configure R4 based on the requirements of this task:

```
R4 (config) #ipv6 unicast-routing
```

```
R4 (config-router) #route-map tst
```

```
R4 (config-route-map) #match interface lo0 lo1 lo2 lo3
```

```
R4 (config) #router ospfv3 1
```

```
R4 (config-router) #address-family ipv6 unicast
```

```
R4 (config-router-af) #router-id 0.0.0.4
```

```
R4 (config-router-af) #redistribute connected route-map tst
```

```
R4 (config-if) #int e0/0
```

```
R4 (config-if) #ospfv3 1 ipv6 area 2
```

You should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 0.0.0.3 on Ethernet0/0 from
LOADING to FULL, Loading Done
```

To verify the configuration:

On R4:

```
R4#show ipv6 route ospf
```

```
IPv6 Routing Table - default - 21 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
       lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
OI 12::/64 [110/84]
   via FE80::3, Ethernet0/0
OI 23::/64 [110/74]
   via FE80::3, Ethernet0/0
OI 2000:111:111:1110::/64 [110/85]
   via FE80::3, Ethernet0/0
OI 2000:111:111:1111::/64 [110/85]
   via FE80::3, Ethernet0/0
OI 2000:111:111:1112::/64 [110/85]
   via FE80::3, Ethernet0/0
OI 2000:111:111:1113::/64 [110/85]
   via FE80::3, Ethernet0/0
OI 2000:222:111:1110::/64 [110/75]
   via FE80::3, Ethernet0/0
OI 2000:222:111:1117::/64 [110/75]
   via FE80::3, Ethernet0/0
OI 2000:222:111:111A::/64 [110/75]
   via FE80::3, Ethernet0/0
OI 2000:222:111:111E::/64 [110/75]
   via FE80::3, Ethernet0/0
```

On R2:

```
R2#show ipv route ospf
```

```

IPv6 Routing Table - default - 22 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
       lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
OI  34::/64 [110/74]
    via FE80::3, Serial1/3
O   2000:111:111:1110::/64 [110/11]
    via FE80::1, Ethernet0/0
O   2000:111:111:1111::/64 [110/11]
    via FE80::1, Ethernet0/0
O   2000:111:111:1112::/64 [110/11]
    via FE80::1, Ethernet0/0
O   2000:111:111:1113::/64 [110/11]
    via FE80::1, Ethernet0/0
OE2 2000:444:111:111B::/64 [110/20]
    via FE80::3, Serial1/3
OE2 2000:444:111:111C::/64 [110/20]
    via FE80::3, Serial1/3
OE2 2000:444:111:111D::/64 [110/20]
    via FE80::3, Serial1/3
OE2 2000:444:111:111F::/64 [110/20]
    via FE80::3, Serial1/3

```

Task 3

Summarize the loopback interfaces configured on R1 and R2. These routers should advertise a single route for their loopback interfaces. DON'T configure more than two summary routes to accomplish this task.

In OSPF, summarization can be configured on an ABR and/or an ASBR. Since R2 is the ABR we should configure R2 to summarize the loopbacks advertised on R1.

The process of summarization in IPv6 is identical to IPv4, let's look at the IPv6 addresses on R1:

```

Lo0 - 2000:111:111:1110::111/64
Lo1 - 2000:111:111:1111::111/64
Lo2 - 2000:111:111:1112::111/64
Lo3 - 2000:111:111:1113::111/64

```


We can see that they all start with **2000:111:111:111**, but the last Hex value in the forth Hextet is where they differ. Let's convert the last Hex value of the forth Hextet to binary:

```
0 = 0 0 0 0
1 = 0 0 0 1
2 = 0 0 1 0
3 = 0 0 1 1
```

Let's count the common identical contiguous bits, we should see the following:

```
0 = 0 0 0 0
1 = 0 0 0 1
2 = 0 0 1 0
3 = 0 0 1 1
```

We can see that the last two binary digits is where they differ, therefore, all the bits up to the third binary digit of the forth Hextet are identical, therefore, the network that summarizes the above four networks should be:

```
2000:111:111:1110::/62
```

Basically, we zeroed the last two bits in the forth Hextet, that gave us the "2000:111:111:1110::" and the prefix length specifies the number of identical contiguous bits, in this case we have:

16 bits : 16 bits : 16 bits : 14 bits or /62.

Let's configure OSPFv3 to summarize these networks, but before we summarize let's verify the routing table of R2 and look for these networks:

On R2:

```
R2#show ipv6 route ospf | inc 2000:111:111
```

```
○ 2000:111:111:1110::/64 [110/11]
○ 2000:111:111:1111::/64 [110/11]
○ 2000:111:111:1112::/64 [110/11]
○ 2000:111:111:1113::/64 [110/11]
```

To summarize:

On R2:

```
R2 (config)#router ospfv3 1
```

```
R7 (config-router) #address-family ipv6 unicast
R2 (config-rtr) #area 1 range 2000:111:111:1110::/62
```

To verify the configuration:

On R3:

```
R3#show ipv6 route ospf | inc 2000:111:111
```

```
OI 2000:111:111:1110::/62 [110/75]
```

On R2:

```
R2#show ipv6 route ospf | inc /62|Null
```

```
O 2000:111:111:1110::/62 [110/11]
   via Null0, directly connected
```

We can see that the discard-route is injected to avoid forwarding loops just like OSPFv2. Let's summarize the loopback interfaces configured on R2.

```
Lo0 - 2000:222:111:1110::111/64
Lo1 - 2000:222:111:1117::111/64
Lo2 - 2000:222:111:111A::111/64
Lo3 - 2000:222:111:111E::111/64
```

Once again the last Hex value of the forth Hextet is where they differ, let's go through the same process:

```
0 - 0 0 0 0
7 - 0 1 1 1
A - 1 0 1 0
E - 1 1 1 0
```

In this case the contiguous identical bits stops at the third Hex value of the forth Hextet, therefore, the summary route should be:

```
2000:222:111:1110::
```

And the prefix length should be:

```
16 bits : 16 bits : 16 bits : 12 bits or /60
```

```
R2 (config)#router ospfv3 1
R2 (config-router)#address-family ipv6 unicast
R2 (config-router-af)#area 0 range 2000:222:111:1110::/60
```

To verify the configuration:

On R1:

```
R1#show ipv6 route ospf | inc /60
```

```
OI 2000:222:111:1110::/60 [110/11]
```

```
R1#ping 2000:222:111:1117::111
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 2000:222:111:1117::111, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/4 ms

We need to repeat the same configuration on R3, but before we configure the summary route for R2's Loopback IPv6 addresses let's verify the routing table of R4:

On R4:

```
R4#show ipv route ospf
```

IPv6 Routing Table - default - 18 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP

H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea

IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO

ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect

O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2

ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site

ld - LISP dyn-EID, a - Application

```
OI 12::/64 [110/84]
```

```
via FE80::3, Ethernet0/0
```

```
OI 23::/64 [110/74]
```

```
via FE80::3, Ethernet0/0
```

```
OI 2000:111:111:1110::/62 [110/85]
```

```
via FE80::3, Ethernet0/0
```

```
OI 2000:222:111:1110::/64 [110/75]
```

```
via FE80::3, Ethernet0/0
```

```
OI 2000:222:111:1117::/64 [110/75]
```

```
via FE80::3, Ethernet0/0
OI 2000:222:111:111A::/64 [110/75]
via FE80::3, Ethernet0/0
OI 2000:222:111:111E::/64 [110/75]
via FE80::3, Ethernet0/0
```

On R3:

```
R3(config)#ipv6 router ospf 1
R3(config-rtr)#area 0 range 2000:222:111:1110::/60
```

To verify the configuration:

On R4:

```
R4#show ipv route ospf
```

```
IPv6 Routing Table - default - 15 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
       lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
OI 12::/64 [110/84]
   via FE80::3, Ethernet0/0
OI 23::/64 [110/74]
   via FE80::3, Ethernet0/0
OI 2000:111:111:1110::/62 [110/85]
   via FE80::3, Ethernet0/0
OI 2000:222:111:1110::/60 [110/75]
   via FE80::3, Ethernet0/0
```

```
R4#ping 2000:222:111:1117::111
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2000:222:111:1117::111, timeout is 2
seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 4

Summarize the external routes redistributed on R4. If this is configured correctly, the rest of the routers should see a single summary route for the four networks redistributed into OSPF routing domain.

Let's look at the IPv6 addresses that we are going to summarize:

```
Lo0 - 2000:444:111:111B::111/64
Lo1 - 2000:444:111:111C::111/64
Lo2 - 2000:444:111:111D::111/64
Lo3 - 2000:444:111:111F::111/64
```

Let's configure the last digit of the forth Hextet to binary:

```
B - 1 0 1 1
C - 1 1 0 0
D - 1 1 0 1
F - 1 1 1 1
```

Therefore, the summary route should be 2000:444:111:1118::/61. Let's configure and verify:

On R4:

In OSPFv2, we use "Area range" command for internal routes and "Summary-address" for summarizing external routes. In OSPFv3, we still use the "Area range" command for internal routes, but for external routes instead of "Summary-address" command the "Summary-prefix" command is used.

```
R4(config)#router ospfv3 1
R4(config-router)#address-family ipv6 unicast
R4(config-router-af)#Summary-prefix 2000:888:111:1118::/61
```

To verify the configuration:

On R1:

```
R1#show ipv6 route ospf | inc /61
```

```
OE2 2000:444:111:1118::/61 [110/20]
```

```
R1#ping 2000:444:111:111f::111
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 2000:444:111:111F::111, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms

Task 5

The policy for summarizing external routes has changed, the routers in area 2 should see all specific external routes whereas, the routers in the other areas should see a single summary route for the four external routes.

Let's remove the "Summary-prefix" configured in the previous step:

On R4:

```
R4 (config) #router ospfv3 1
R4 (config-router) #address-family ipv6 unicast
R4 (config-router-af) #no summary-prefix 2000:444:111:1118::/61
```

To verify the configuration:

On R1:

```
R1#show ipv6 route ospf | inc /61
R1#
```

```
R1#show ipv6 route ospf | inc E2
```

```
      O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
OE2 2000:444:111:111B::/64 [110/20]
OE2 2000:444:111:111C::/64 [110/20]
OE2 2000:444:111:111D::/64 [110/20]
OE2 2000:444:111:111F::/64 [110/20]
```

Note: Summarization in OSPFv3 is identical to OSPFv2; there are two routers where the summarization can be configured:

- **The routes are internal:**
If the routes are internal, the summarization can only be configured on the ABRs using the "**Area**

range” command.

➤ **The routes are external:**

The external routes can be summarized on an ASBR. ASBR is a router that originates external routes, these are “E” or “N” routes. Since area 2 is a normal area, the summarization is performed on the ASBR (R4), but if this area is converted into an NSSA area, R3 will receive the “N” routes and it originates “E” routes. Since R3 will originate “E” routes, it becomes an ASBR, therefore, summarization of external routes can be configured on R3:

On R4:

```
R4 (config) #router ospfv3 1
R4 (config-router) #address-family ipv6 unicast
R4 (config-router-af) #area 2 nssa
```

You should see the following console message stating that the adjacency to R3 is down. This is because the area stub flag no longer matches, once R3 is configured as an NSSA area, the area stub flag will match (They will both be in NSSA area) and the adjacency will be reestablished.

```
%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 0.0.0.3 on Ethernet0/0 from FULL
to DOWN, neighbor Down: Adjacency forced to reset
```

On R3:

```
R3 (config) #ipv6 router ospf 1
R3 (config-rtr) #area 2 nssa
R3 (config-rtr) #summary-prefix 2000:444:111:1118::/61
```

You should see the following console message stating that the adjacency with 0.0.0.4 or R4 has transitioned into FULL state:

```
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.4 on Ethernet0/0 from LOADING
to FULL, Loading Done
```

To verify the configuration:

On R1:

```
R1#show ipv6 route ospf | inc /61
```

```
OE2 2000:444:111:1118::/61 [110/20]
```

```
R1#ping 2000:444:111:111f::111
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 2000:444:111:111F::111, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/4 ms

Task 6

None of the routers should have a discard route in their routing table.

In OSPF, the “discard-route/s” are auto-injected on the router that configured the summarization. Let’s verify the existence of the discard routes on R2 and R3:

On R3:

```
R3#show ipv6 route ospf
```

```
IPv6 Routing Table - default - 17 entries
```

```
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
```

```
      B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
      H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
```

```
      IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
```

```
      ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
```

```
      O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
      ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
```

```
      lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
```

```
OI 12::/64 [110/74]
   via FE80::2, Serial1/2
OI 2000:111:111:1110::/62 [110/75]
   via FE80::2, Serial1/2
O 2000:222:111:1110::/60 [110/65]
  via Null0, directly connected
O 2000:222:111:1110::/64 [110/65]
  via FE80::2, Serial1/2
O 2000:222:111:1117::/64 [110/65]
  via FE80::2, Serial1/2
O 2000:222:111:111A::/64 [110/65]
  via FE80::2, Serial1/2
O 2000:222:111:111E::/64 [110/65]
  via FE80::2, Serial1/2
O 2000:444:111:1118::/61 [254/20]
  via Null0, directly connected
ON2 2000:444:111:111B::/64 [110/20]
```



```
via 34::4, Ethernet0/0
ON2 2000:444:111:111C::/64 [110/20]
via 34::4, Ethernet0/0
ON2 2000:444:111:111D::/64 [110/20]
via 34::4, Ethernet0/0
ON2 2000:444:111:111F::/64 [110/20]
via 34::4, Ethernet0/0
```

The output of the above show command reveals that R3 has two discard-routes, the 2000:222:111:1110::/60, which is the discard-route for internal networks within Area 0, and 2000:444:111:1118::/61 which is the discard-route for external networks. Let's remove the discard-routes:

```
R3(config)#router ospfv3 1
R3(config-router)#address ipv6
R3(config-rtr)#no discard-route internal
R3(config-rtr)#no discard-route external
```

To verify this configuration:

On R2:

```
R2#show ipv6 route ospf
```

```
IPv6 Routing Table - default - 21 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
       lr - LISP site-registrations, ld - LISP dyn-eid, a - Application OI
34::/64 [110/74]
via FE80::3, Serial1/3
O 2000:111:111:1110::/62 [110/11]
  via Null0, directly connected
O 2000:111:111:1110::/64 [110/11]
  via FE80::1, Ethernet0/0
O 2000:111:111:1111::/64 [110/11]
  via FE80::1, Ethernet0/0
O 2000:111:111:1112::/64 [110/11]
  via FE80::1, Ethernet0/0
O 2000:111:111:1113::/64 [110/11]
  via FE80::1, Ethernet0/0
O 2000:222:111:1110::/60 [110/1]
```

via Null0, directly connected

```
OE2 2000:444:111:1118::/61 [110/20]
    via FE80::3, Serial1/3
```

Let's remove the discard-route from R2:

```
R2(config)#router ospfv3 1
R2(config-router)#address ipv6
R2(config-router-af)#no discard-route internal
```

To verify this configuration:

```
R2#show ipv6 route ospf
```

```
IPv6 Routing Table - default - 19 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
       lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
OI 34::/64 [110/74]
    via FE80::3, Serial1/3
O 2000:111:111:1110::/64 [110/11]
    via FE80::1, Ethernet0/0
O 2000:111:111:1111::/64 [110/11]
    via FE80::1, Ethernet0/0
O 2000:111:111:1112::/64 [110/11]
    via FE80::1, Ethernet0/0
O 2000:111:111:1113::/64 [110/11]
    via FE80::1, Ethernet0/0
OE2 2000:444:111:1118::/61 [110/20]
    via FE80::3, Serial1/3
```

Task 7

Erase the startup configuration of the routers, config.text and the VLAN.dat of the switches and reload them before proceeding to the next lab.