Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

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

IPv6 Questions & Answers

Troubleshooting IPv6

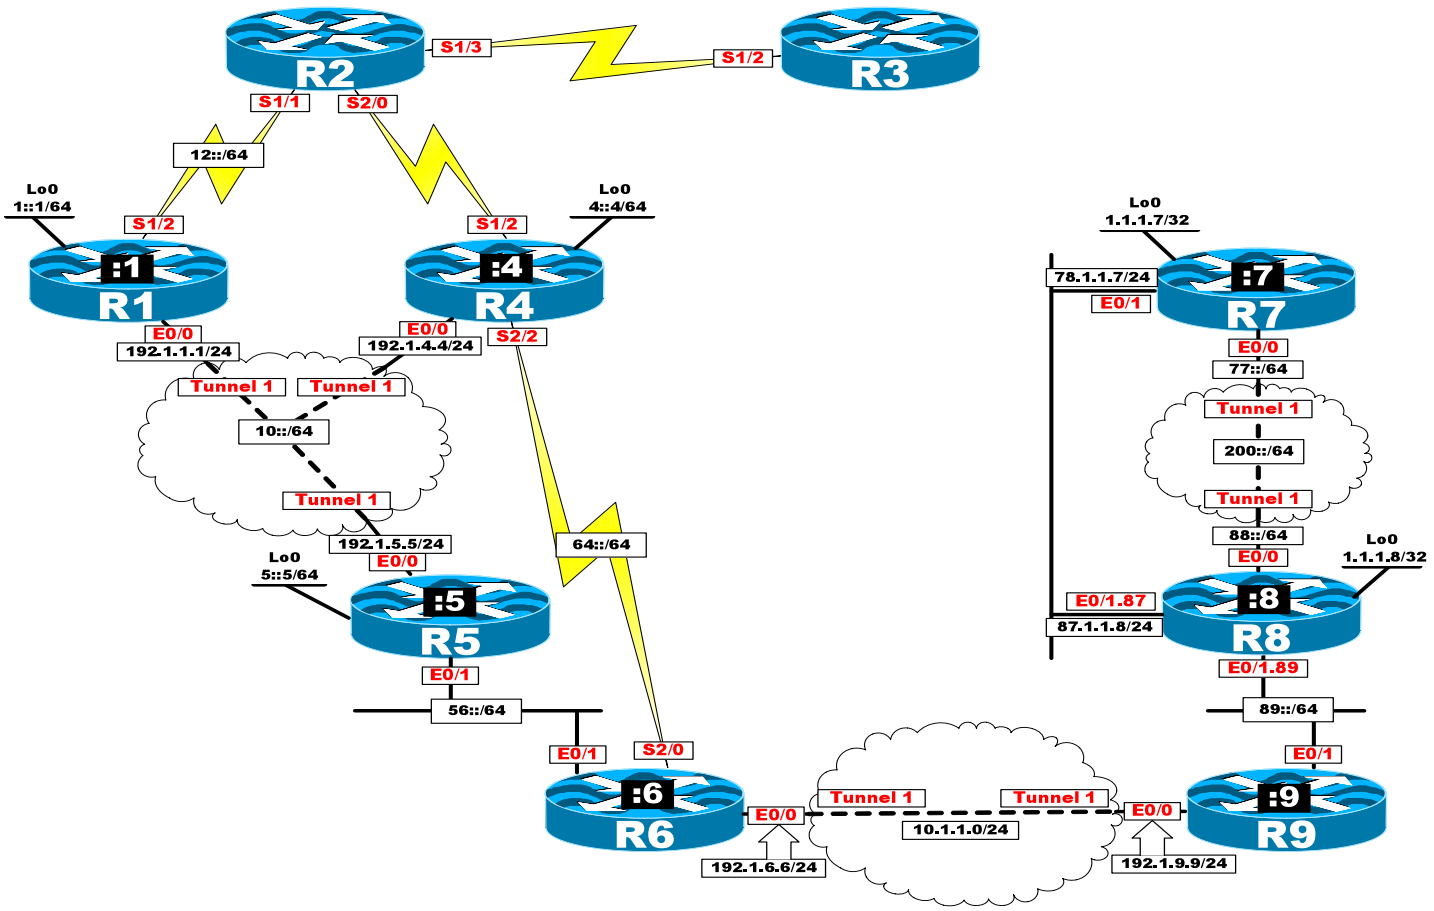

Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "IPv6" →"Lab-1".

Lab Rules:

- **DO NOT remove any command/s unless otherwise stated**

- **DO NOT change the VLAN assignment**
- **You must be VERY specific when resolving these tasks**

Ticket 1

R1 (The PE) has an IPv6 prefix of 20:1:1::/48 and it needs to subnet this prefix to /56 subnets. R2 (The CE) has two sites that are connected through its S1/3 and S2/0 interfaces, but soon this company will grow to 12 sites.
R2 was given subnet one, and subnetted this network to a minimum of 12 subnets. The first subnet one was automatically assigned to its S1/3 interface with the host portion of its IPv6 address as "::23"; R3 should use R2 as its default gateway.
R2's S2/0 interface was configured to automatically assign the next subnet with the host portion of "::24". R4 should use R2 as its default gateway.
R3 and R4 were configured to automatically acquire the network portion of their IPv6 address from R2, but for some strange reason R2 is getting the wrong subnets from R1.
You can remove and reconfigure one command to fix this ticket.

Let's see if R2 acquired an IPv6 address from R1 for its S1/1 interface:

On R2:

```
R2#sh ipv6 inter br s1/1

Serial1/1                  [up/up]
    FE80::2
    12::2
```

Yes, R2 did acquire an IPv6 address from R1 and it's from the correct range, if R2 didn't acquire an IPv6 address you should shut and no shut the interface. Let's verify the IPv6 addresses that R2 assigned to its S1/3 and S2/0 interfaces:

```
R2#show ipv6 inter bri s1/3

Serial1/3                  [up/up]
    FE80::A8BB:CCFF:FE00:2100
    20:1:1:1::23

R2#show ipv6 inter bri s2/0

Serial2/0                  [up/up]
    FE80::A8BB:CCFF:FE00:2100
    20:1:1:2::24
```

Sure enough they are getting an IPv6 address but let's see if the IPv6 addresses are from the correct subnet:

R1 has 20:1:1::/48 and it subneted this network to /56s and it gave the first subnet with /60 to R2. The reason he had to give a /60 to R2 is because the ticket states that R2 will grow to 12 sites, therefore, 12 subnets will require 4 bits and if we add 4 bits to /56, it's going to be /60.

The following diagram shows the bits used by the provider to generate some /56 subnets.

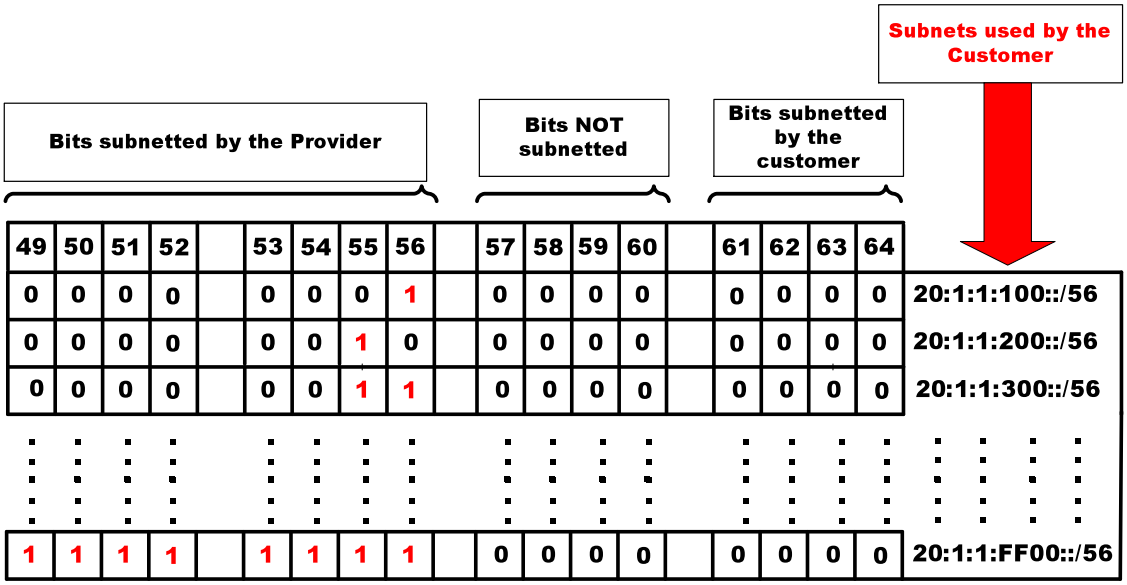

The provider assigned the subnet one to the customer (R2), this is "20:1:1:100::/60". The following diagram shows the customer subnetting the 4 bits to generate up to 16 networks:

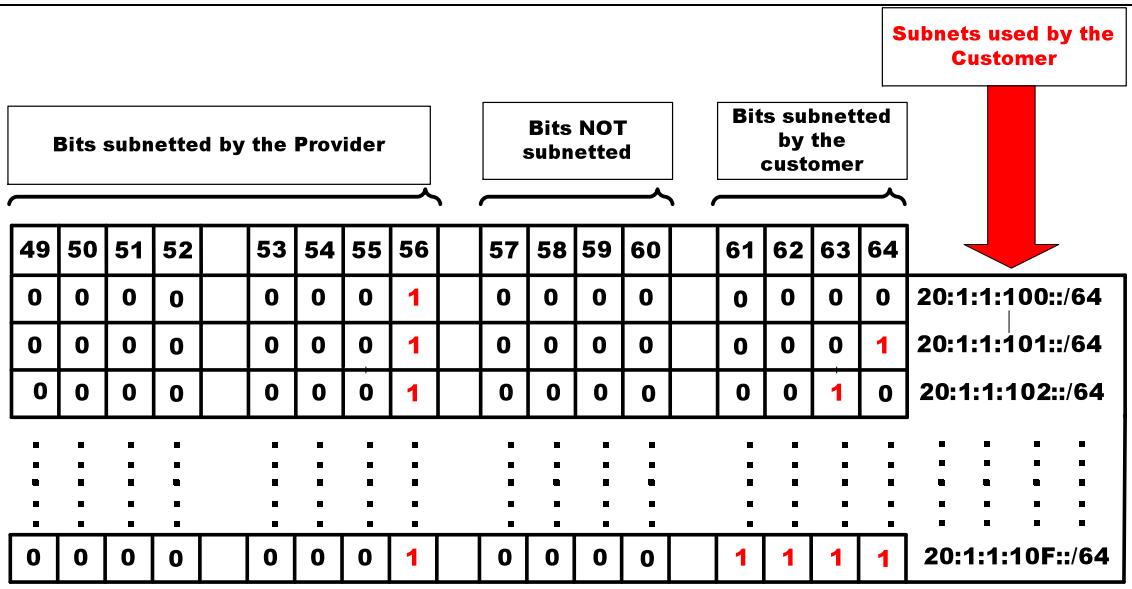

Since R2 is configured to assign the first subnet to its S1/3 interface, it should have been from 20:1:1:101::/60 and NOT 20:1:1:1::/60. This is implemented by the "IPv6 local pool" command. let's see the configuration on R1:

On R1:

```
R1#show run | i ipv6 local pool
ipv6 local pool test 20:1:1::/56 60
```

We can see the problem, let's change this.

```
R1#clear ipv6 dhcp binding *

R1(config)#no ipv6 local pool test 20:1:1::/56 60

R1(config)#ipv6 local pool test 20:1:1:100::/56 60
```

Let's save the configuration and reload R1 and R2, and once the routers are up, we should see the following:

On R2:

```
R2#sh ipv6 general-prefix
```

```
IPv6 Prefix ISP, acquired via DHCP PD
  20:1:1:100::/60 Valid lifetime infinite, preferred lifetime infinite
   Serial1/3 (Address command)
   Serial2/0 (Address command)

R2#sh ipv6 inter bri s1/3

Serial1/3              [up/up]
    FE80::213:C4FF:FE55:F568
    20:1:1:101::23

R2#sh ipv6 inter bri s2/0

Serial2/0              [up/up]
    FE80::213:C4FF:FE55:F569
    20:1:1:102::24
```

If R3 and R4 have not acquired an IPv6 address, you should "shut" and then "no shut" their S1/2 interfaces.

On R3:

```
R3#sh ipv6 int br s1/2

Serial1/2                  [up/up]
    FE80:: A8BB:CCFF:FE00:3100
    20:1:1:101: A8BB:CCFF:FE00:3100
```

On R4:

```
R4#sh ipv6 int br s1/2

Serial1/2                  [up/up]
    FE80:: A8BB:CCFF:FE00:4100
    20:1:1:102: A8BB:CCFF:FE00:4100
```

Ticket 2

R5 is the hub router, and R1 and R4 are the spokes in the DMVPN network. The hub and the spoke routers are configured with static NHRP mapping. These three routers are running IPv6 OSPF on their tunnel 1 and their lo0 interfaces, but for some reason they don't have reachability to each other's lo0 or tunnel IPv6 interfaces.
You should keep the DMVPN configured with static mapping, do not configure additional mapping. DO NOT configure another dynamic routing protocol to fix this ticket.

Let's verify the ticket:

On R5:

```
R5#sh ipv6 route ospf

IPv6 Routing Table - default - 3 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
```

Let's see if we have reachability to the tunnel IPv6 addresses:

```
R5#ping 10::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10::1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#ping 10::4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10::4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's see the tunnel configuration and verify reachability to the tunnel source IP addresses of R1 and R4:

```
R5#ping 192.1.4.4
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.4.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#ping 192.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Do we have a route to these IP addresses?

```
R5#sh ip route | B Gate
Gateway of last resort is not set

      192.1.5.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.1.5.0/24 is directly connected, FastEthernet0/0
L        192.1.5.5/32 is directly connected, FastEthernet0/0
```

Based on the physical diagram, we know that R5's E0/0 interface is connected to SW1's F0/5 interface, let's check the configuration of F0/5 interface on SW1:

On SW1:

```
SW1#sh run int e1/1 | b inter

interface Ethernet1/1
 no switchport
 ip address 192.1.5.10 255.255.255.0
 duplex auto
end
```

This IP address should be the default gateway for R5. Since we are not allowed to run another routing protocol, let's configure two static routes for 192.1.1.1 and 192.1.4.4 IP addresses on R5 and point them to 192.1.5.10.

On R5:

```
R5(config)#ip route 192.1.1.1 255.255.255.255 192.1.5.10
R5(config)#ip route 192.1.4.4 255.255.255.255 192.1.5.10
```

To verify the configuration:

```
R5#ping 192.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Maybe R1 is missing the same configurations, let's verify:

On R1:

```
R1#sh ip route | b Gate
Gateway of last resort is not set

      192.1.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.1.1.0/24 is directly connected, FastEthernet0/0
L        192.1.1.1/32 is directly connected, FastEthernet0/0
```

Yes, it is missing the routes, let's configure two static routes on R1, one for R4 and the second one for R5:

On SW1:

```
SW1#sh run int e0/1 | b inter

interface Ethernet0/1
 no switchport
 ip address 192.1.1.10 255.255.255.0
 duplex auto
end
```

On R1:

```
R1(config)#ip route 192.1.5.5 255.255.255.255 192.1.1.10
R1(config)#ip route 192.1.4.4 255.255.255.255 192.1.1.10
```

To verify the configuration:

```
R1#ping 192.1.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.5.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Since both R1 and R5 are configured correctly, SW1 can be the problem, let's verify:

On SW1:

```
SW1#sh ip route | b Gate

Host                 Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
```

OK, we can see the problem, "IP routing" is not configured, let's configure "IP routing" and verify:

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

To verify the configuration:

```
SW1#sh ip route | B Gate
Gateway of last resort is not set

      192.1.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.1.1.0/24 is directly connected, Ethernet0/1
L        192.1.1.10/32 is directly connected, Ethernet0/1
      192.1.4.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.1.4.0/24 is directly connected, Ethernet1/0
L        192.1.4.10/32 is directly connected, Ethernet1/0
      192.1.5.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.1.5.0/24 is directly connected, Ethernet1/1
L        192.1.5.10/32 is directly connected, Ethernet1/1
      192.1.6.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.1.6.0/24 is directly connected, Ethernet1/2
L        192.1.6.10/32 is directly connected, Ethernet1/2
      192.1.9.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.1.9.0/24 is directly connected, Ethernet2/1
L        192.1.9.10/32 is directly connected, Ethernet2/1
```

On R1:

```
R1#ping 192.1.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.5.5, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
```

Maybe R4 is missing the same configuration that R1 and R5 were missing, let's verify:

On R4:

```
R4#sh ip route | b Gate
Gateway of last resort is not set

      192.1.4.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.1.4.0/24 is directly connected, Ethernet0/0
L        192.1.4.4/32 is directly connected, Ethernet0/0
```

Sure enough, it is missing a route to 192.1.5.5.

On SW1:

```
SW1#sh run int e1/0 | b inter

interface Ethernet1/0
 no switchport
 ip address 192.1.4.10 255.255.255.0
 duplex auto
end
```

On R4:

```
R4(config)#ip route 192.1.1.1 255.255.255.255 192.1.4.10
R4(config)#ip route 192.1.5.5 255.255.255.255 192.1.4.10
```

To verify the configuration:

```
R4#ping 192.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R4#ping 192.1.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Let's go back to R5 and see if this fixed the problem:

On R5:

```
R5#sh ipv6 route ospf

IPv6 Routing Table - default - 3 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
```

Let's verify the IPv6 NHRP mapping:

```
R5#sh ipv6 nhrp

10::1/128 via 10::1
   Tunnel1 created 00:21:56, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
10::4/128 via 10::4
   Tunnel1 created 00:00:57, never expire
   Type: static, Flags:
   NBMA address: 192.1.4.4
```

We can see the problem, since the DMVPN is configured with IPv6 and it is running on top of IPv4, the local router must map next-hop's Link local IPv6 address to the next-hop's NBMA-IP address, and what we see here is a mapping of next-hop's tunnel IPv6 address to its NBMA-IP, let's change this configuration and verify. Let's find out the link local IPv6 addresses of R1 and R4:

On R1:

```
R1#sh ipv6 int br tunne 1

Tunnel1                    [up/up]
    FE80::1
    10::1
```

On R4:

```
R4#sh ipv6 inter bri tunne 1

Tunnel1                    [up/up]
    FE80::4
```

```
    10::4
```

Let's reconfigure R5's nhrp mapping:

On R5:

```
R5(config)#int tunn 1
R5(config-if)#no ipv6 nhrp map 10::4/128 192.1.4.4
R5(config-if)#no ipv6 nhrp map 10::1/128 192.1.1.1
R5(config-if)#ipv6 nhrp map fe80::4/128 192.1.4.4
R5(config-if)#ipv6 nhrp map fe80::1/128 192.1.1.1
```

We should check R1 and R4:

On R1:

```
R1#sh ipv6 nhrp

10::5/128 via 10::1
   Tunnel1 created 00:44:06, never expire
   Type: static, Flags:
   NBMA address: 192.1.5.5

R1(config)#int tunn 1
R1(config-if)#no ipv6 nhrp map 10::5/128 192.1.5.5
R1(config-if)#ipv6 nhrp map fe80::5/128 192.1.5.5
```

To verify the configuration:

```
R1#sh ipv nhrp

FE80::5/128 via FE80::5
   Tunnel1 created 00:00:10, never expire
   Type: static, Flags:
   NBMA address: 192.1.5.5
```

Let's check R4:

On R4:

```
R4#sh ipv6 nhrp

10::5/128 via 10::5
   Tunnel1 created 00:47:35, never expire
```

```
   Type: static, Flags:
   NBMA address: 192.1.4.4

R4(config)#int tunn 1
R4(config-if)#no ipv6 nhrp map 10::5/128 192.1.4.4
R4(config-if)#ipv6 nhrp map fe80::5/128 192.1.5.5
```

To verify the configuration:

```
R4#sh ipv6 nhrp

FE80::5/128 via FE80::5
   Tunnel1 created 00:00:09, never expire
   Type: static, Flags:

   NBMA address: 192.1.5.5
```

Let's verify the configuration:

On R5:

```
R5#sh ipv6 nhrp

FE80::1/128 via FE80::1
   Tunnel1 created 00:09:12, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
FE80::4/128 via FE80::4
   Tunnel1 created 00:09:19, never expire
   Type: static, Flags:
   NBMA address: 192.1.4.4

R5#ping fe80::1

Output Interface: tunnel1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::1, timeout is 2 seconds:
Packet sent with a source address of FE80::5%Tunnel1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms

R5#ping fe80::4

Output Interface: tunnel1
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::4, timeout is 2 seconds:
Packet sent with a source address of FE80::5%Tunnel1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms

R5#sh ipv6 route ospf

IPv6 Routing Table - default - 5 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
```

Let's see if Multicast capability was provided:

```
R5#sh run int tunn 1 | i ipv6 nhrp map

ipv6 nhrp map FE80::1/128 192.1.1.1
ipv6 nhrp map FE80::4/128 192.1.4.4
```

We can see the problem, we have not provided multicast capability, let's fix this problem:

```
R5(config)#int tunn 1
R5(config-if)#ipv6 nhrp map multicast 192.1.1.1
R5(config-if)#ipv6 nhrp map multicast 192.1.4.4
```

Let's check R1:

On R1:

```
R1#sh run int tunn 1 | i ipv6 nhrp map

ipv6 nhrp map FE80::5/128 192.1.5.5

R1(config)#int tunn 1
R1(config-if)#ipv6 nhrp map multicast 192.1.5.5
```

Once the mapping is configured, you should see the following console message:

```
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.5 on Tunnel1 from LOADING to FULL,
Loading Done
```

Let's check R4:

On R4:

```
R4#sh run int tunn 1 | i ipv6 nhrp map

ipv6 nhrp map FE80::5/128 192.1.5.5

R4(config)#int tunn 1
R4(config-if)#ipv6 nhrp map multicast 192.1.5.5
```

To verify the configuration:

On R5:

```
R5#sh ipv6 route ospf

IPv6 Routing Table - default - 5 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::1/128 [110/1000]
     via FE80::1, Tunnel1
```

When OSPF is configured on a tunnel interface, OSPF does not read the "tunnel mode" command and therefore, it sets the OSPF network type to point-to-point, since R5 is the hub router, it should be configured as point-to-multipoint and the timers should be changed to match the spokes, this means that the hello interval should be set to 10 seconds and the dead interval to 40. Let's verify and configure the tunnel interface of R5 to fix this problem:

On R5:

```
R5#sh ipv6 ospf int tunne 1 | i Network

  Network Type POINT_TO_POINT, Cost: 1000

R5(config)#int tunn 1
R5(config-if)#ipv6 ospf network point-to-multipoint
R5(config-if)#ipv6 ospf hello-interval 10
```

You should see the following console messages:

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

%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Tunnel1 from LOADING to FULL,
Loading Done
```

To verify the configuration:

```
R5#sh 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
       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::1/128 [110/1000]
     via FE80::1, Tunnel1
O   4::4/128 [110/1000]
     via FE80::4, Tunnel1

R5#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 = 0/0/4 ms

R5#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 = 4/4/4 ms

R5#ping 10::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10::1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

NOTE: We have reachability to the loopback interfaces because they are exchanged through OSPF using our Linbk local IPv6 address as their next-hop, since the tunnel is configured statically, the routers will not have reachability to tunnel's IPv6 addresses. One way to fix this problem is to change the network type on the spokes to "Point-to-Multipoint", this network type will advertise the local router's tunnel IPv6 address to the adjacent router/s, we can see this by going to the spokes and verifying their routing table:

On R4:

```
R4#sh ipv6 route ospf
```

We can see that R4 has R5's tunnel IPv6 address in its routing table as a host route (/128), this was advertised by R5 because of R5's OSPF network type.

```
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
       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::1/128 [110/2000]
     via FE80::5, Tunnel1
O   5::5/128 [110/1000]
     via FE80::5, Tunnel1
O   10::5/128 [110/1000]
     via FE80::5, Tunnel1
OE2 14::/96 [110/20]
     via FE80::6, Serial2/2
```

On R1:

```
R1#sh ipv6 route ospf

IPv6 Routing Table - default - 11 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   4::4/128 [110/2000]
     via FE80::5, Tunnel1
O   5::5/128 [110/1000]
     via FE80::5, Tunnel1
```

```
O   10::5/128 [110/1000]
     via FE80::5, Tunnel1
```

NOTE: the spoke routers have R5's tunnel IPv6 address because R5 is configured with a "Point-to-Multipoint" OSPF network type, let's reconfigure the same network type on R1 and R4:

On R1 and R4:

```
Rx(config)#int tunn 1
Rx(config-if)#ipv6 ospf network point-to-multipoint
```

Let's change the timers on R5 to match R1 and R4:

On R5:

```
R5(config)#int tunn 1
R5(config-if)#no ipv6 ospf hello-interval 10
```

You should see the following console message:

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

%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Tunnel1 from LOADING to FULL,
Loading Done

R5#sh ipv6 route ospf

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
O   1::1/128 [110/1000]
     via FE80::1, Tunnel1
O   4::4/128 [110/1000]
     via FE80::4, Tunnel1
O   10::1/128 [110/1000]
     via FE80::1, Tunnel1
O   10::4/128 [110/1000]
     via FE80::4, Tunnel1

R5#ping 10::1
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms

R5#ping 10::4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10::4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
```

Ticket 3

R7 is configured to be the NHS, and R8 is configured as a spoke. R7 and R8 are running OSPF veesion 3 on their tunnel and loopback interfaces, but for some reason they don't have reachability to the loopback interface that the other router is advertising. DO NOT run another dynamic routing protocol to fix this problem.

Let's verify the ticket:

On R7:

```
R7#sh ipv route ospf

IPv6 Routing Table - default - 5 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
```

Let's see if we have reachability to the tunnel IPv6 address of R8:

```
R7#ping 200::8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200::8, timeout is 2 seconds:
```

```
% No valid route for destination
Success rate is 0 percent (0/1)
```

Let's verify the configuration of R7 and R8's tunnel interface:

```
R7#sh run int tunn 1 | b inter

interface Tunnel1
 no ip address
 no ip redirects
 ipv6 address FE80::7 link-local
 ipv6 address 200::7/64
 ipv6 nhrp map multicast dynamic
 ipv6 nhrp network-id 777
 ipv6 ospf 1 area 0
 tunnel source Ethernet0/0
 tunnel mode gre multipoint
end
```

On R8:

```
R8#sh run int tunn 1 | B inter

interface Tunnel1
 no ip address
 no ip redirects
 ipv6 address FE80::8 link-local
 ipv6 address 200::8/64
 ipv6 nhrp map 200::7/128 77::7
 ipv6 nhrp network-id 888
 tunnel source Ethernet0/0
 tunnel mode gre multipoint
```

Let's see if these two routers have reachability to the tunnel source IPv6 address of each other:

```
R8#ping 77::7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 77::7, timeout is 2 seconds:

% No valid route for destination
Success rate is 0 percent (0/1)
```

Based on the physical topology, we know that R7 and R8 are connected to the E1/3 and E2/0 interfaces of SW1 respectively, let's verify the configuration of these two interfaces:

On SW1:

```
SW1#sh run int e1/3 | b inter

interface Ethernet1/3
 no switchport
 no ip address
 ipv6 address FE80::10 link-local
 ipv6 address 77::10/64
end

SW1#sh run int e2/0 | b inter

interface Ethernet2/0
 no switchport
 no ip address
 ipv6 address FE80::10 link-local
 ipv6 address 88::10/64
end
```

Let's verify the configuration of the E0/0 interface on R7 and R8:

On R8:

```
R8#sh run int e0/0 | b inter

interface Ethernet0/0
 no ip address
 ipv6 address FE80::8 link-local
 ipv6 address 88::8/64
 ipv6 ospf 1 area 0
end
```

On R7:

```
R7#sh run int e0/0 | b inter

interface Ethernet0/0
 no ip address
 ipv6 address FE80::7 link-local
 ipv6 address 77::7/64
```

Let's verify if R7 can ping its default gateway 77::10

```
R7#ping 77::10
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 77::10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/8 ms

R7#sh ipv6 route

IPv6 Routing Table - default - 5 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
C   7::/64 [0/0]
     via Loopback0, directly connected
L   7::7/128 [0/0]
     via Loopback0, receive
C   77::/64 [0/0]
     via Ethernet0/0, directly connected
L   77::7/128 [0/0]
     via Ethernet0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

We don't see an entry for 88::/64 network in the routing table. Since the ticket states that no additional dynamic routing protocol should be configured, we will configure a static route for 88::/64 network.

```
R7(config)#ipv6 route 88::/64 77::10
```

Let's verify the same on R8:

On R8:

```
R8#sh ipv6 route

IPv6 Routing Table - default - 5 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
C   8::/64 [0/0]
     via Loopback0, directly connected
L   8::8/128 [0/0]
     via Loopback0, receive
C   88::/64 [0/0]
     via Ethernet0/0, directly connected
L   88::8/128 [0/0]
     via Ethernet0/0, receive
C   89::/64 [0/0]
     via Ethernet0/1.89, directly connected
L   89::8/128 [0/0]
     via Ethernet0/1.89, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

We need to configure a static route on R8 for 77::/64 network:

```
R8(config)#ipv6 route 77::/64 88::10
```

To verify the configuration:

```
R8#ping 77::7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 77::7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Do we have "IPv6 unicast-routing" enabled on SW1? Let's verify:

On SW1:

```
SW1#sh run | i ipv6 unicast-routing
SW1#
```

Let's enable "IPv6 unicast-routing" and try again:

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

To verify the configuration:

On R8:

```
R8#ping 77::7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 77::7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

On R7:

```
R7#ping 88::8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 88::8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Let's verify and see if OSPF is running on the correct interfaces:

```
R7#sh ipv6 ospf int br

Interface    PID   Area            Intf ID    Cost  State Nbrs F/C
Lo0          1     0               9          1     LOOP  0/0
Tu1          1     0               8          1000  DOWN  0/0
```

On R8:

```
R8#sh ipv6 ospf int br

Interface    PID   Area            Intf ID    Cost  State Nbrs F/C
Lo0          1     0               9          1     LOOP  0/0
Tu1          1     0               8          1000  DOWN  0/0
```

OSPF is running on the correct interfaces, but the tunnel 1 interface shows up as down. Let's verify the tunnel configuration:

```
R8#sh run int tunn 1 | b inter

interface Tunnel1
 no ip address
 no ip redirects
 ipv6 address FE80::8 link-local
 ipv6 address 200::8/64
 ipv6 nhrp map 200::7/128 77::7
 ipv6 nhrp network-id 888
```

```
 ipv6 ospf 1 area 0
 tunnel source Ethernet0/0
 tunnel mode gre multipoint
end
```

We fixed the tunnel source reachability, but we can see how the tunnel 1 interface is not configured correctly, since R7 is the NHS, R8 should be referencing it in its configuration. We can also see that the IPv6 nhrp mapping is incorrect, and the "IPv6" keyword is missing in the "Tunnel mode" command. Let's fix these problems one at a time:

Let's remove the mapping and reconfigure it correctly:

```
R8(config)#int tunn 1
R8(config-if)#no ipv6 nhrp map 200::7/128 77::7
R8(config-if)#ipv6 nhrp nhs 200::7 nbma 77::7 multicast
```

As we can see the NHS references the tunnel IPv6 address of R7 and it also specifies the NBMA IPv6 address of NHS, and Multicast capability is done on the same command.

When DMVPN's NBMA and Tunnel address is IPv6, we MUST use the "IPv6" keyword on the "Tunnel mode" command or else the tunnel interface will not come up. Let's configure that:

```
R8(config-if)#tunnel mode gre multipoint ipv6
```

You should see the following console message stating that the tunnel interface is up:

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1, changed state to
up
```

Let's verify the configuration:

```
R8#ping 200::7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200::7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check R7's configuration:

On R7:

```
R7#sh run int tunn 1 | b inter
```

```
interface Tunnel1
 no ip address
 no ip redirects
 ipv6 address FE80::7 link-local
 ipv6 address 200::7/64
 ipv6 nhrp map multicast dynamic
 ipv6 nhrp network-id 777
 ipv6 ospf 1 area 0
 tunnel source Ethernet0/0
 tunnel mode gre multipoint
end
```

On R7, the only command missing is the "IPv6" keyword so the tunnel interface can transition into up/up state.

```
R7#sh ipv6 inter br tunne 1

Tunnel1                [up/down]
    FE80::7
    200::7

R7(config)#int tunn 1
R7(config-if)#tunnel mode gre multi ipv6
```

You should see the following console message:

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1, changed state to
up

%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.8 on Tunnel1 from LOADING to FULL,
Loading Done
```

To verify the configuration:

```
R7#ping 200::8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200::8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R7#sh ipv6 route ospf

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
       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   8::8/128 [110/1000]
     via FE80::8, Tunnel1

R7#ping 8::8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8::8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R7#ping 8::8 source lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8::8, timeout is 2 seconds:
Packet sent with a source address of 7::7
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Ticket 4

R8 is configured as a DHCP server on the 89::/64 segment. R9 is configured to acquire the prefix portion of its IPv6 address through the SLAAC process, and its domain-name and DNS-server from the DHCP server. But for some reason this is not working.

Let's verify the ticket:

On R9:

```
R9#sh ipv6 inter br e0/1

Ethernet0/1            [up/up]
    FE80::9
    89::9
```

R9 did acquire an IPv6 address, let's see if it acquired the DNS and domain name from the DHCP server:

```
R9#sh ipv6 dhcp interface
R9#
```

Let's verify the configuration of R9's E0/1 interface:

```
R9#sh run int e0/1 | b inter

interface Ethernet0/1
 no ip address
 ipv6 address FE80::9 link-local
 ipv6 address autoconfig default
end
```

The interface is configured correctly, let's chceck R8's configuration:

On R8:

```
R8#sh run | s ipv6 dhcp pool

ipv6 dhcp pool tst
 address prefix 89::/64
 dns-server 2000:89::89
 domain-name MicronicsTraining.com
```

Let's see if the correct bit is set:

```
R8#sh run int e0/1.89 | b inter

interface Ethernet0/1.89
 encapsulation dot1Q 89
 ipv6 address FE80::8 link-local
 ipv6 address 89::8/64
 ipv6 enable
 ipv6 nd managed-config-flag
 ipv6 dhcp server tst
end
```

We can see the problem, the DHCP server (R8) is configured to assign IPv6 addresses, DNS server and domain name because the "M" bit is set. R8 should be configured with "Other-config-flag" or set the "O" bit so R9 acquires its DNS and domain name through DHCP and its IPv6 address through SLAAC. Let's configure and verify:

On R8:

```
R8(config)#int e0/1.89
R8(config-if)#no ipv6 nd managed-config-flag
R8(config-if)#ipv6 nd other-config-flag

R8(config-subif)#int e0/1
R8(config-if)#shut

%LINK-5-CHANGED: Interface GigabitEthernet0/1, changed state to
administratively down

%LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/1,
changed state to down

R8(config-if)#no shut
```

To verify the configuration:

On R9:

```
R9#sh ipv6 dhcp interface

Ethernet0/1 is in client mode
  Prefix State is IDLE (0)
  Information refresh timer expires in 23:59:49
  Address State is IDLE
  List of known servers:
    Reachable via address: FE80::8
    DUID: 00030001AABBCC008100
    Preference: 0
    Configuration parameters:
      DNS server: 2000:89::89
      Domain name: MicronicsTraining.com
      Information refresh time: 0
  Prefix Rapid-Commit: disabled
  Address Rapid-Commit: disabled

R9#sh ipv6 inter br e0/1

FastEthernet0/1            [up/up]
    FE80::9
    89::9
```

```
R9#ping 89::8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 89::8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/8 ms
```

Ticket 5

R1 is configured as a DHCP server. R6 is configured to acquire the network portion of its IPv6 address from R5 and its domain name and DNS from the DHCP server (R1). But for some reason this is not working.

Let's verify the ticket:

On R6:

```
R6#sh ipv6 inter br e0/1

Ethernet0/1            [up/up]
    FE80::6
```

R6 is configured to acquire the network portion of its IPv6 address from R5. Let's verify the configuration of R5, it should have "IPv6 unicast-routing" and an IPv6 address:

On R5:

```
R5#sh run int e0/1 | b inter

nterface Ethernet0/1
 no ip address
 ipv6 address FE80::5 link-local
 ipv6 nd other-config-flag
 ipv6 nd ra suppress

R5#sh run | i ipv6 unicast

ipv6 unicast-routing
```

IPv6 unicast-routing is configured, but R5's E0/1 interface is not configured with an IPv6 address:

```
R5(config)#int e0/1
R5(config-if)#ipv6 addr 56::5/64
R5(config-if)#shu

%LINK-5-CHANGED: Interface FastEthernet0/1, changed state to
administratively down


%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed
state to down

R5(config-if)#no shut
```

To verify the configuration:

On R6:

```
R6#sh ipv6 int br e0/1

Ethernet0/1            [up/up]
    FE80::6
```

Let's "shut" and "no shut" this interface and verify again:

```
R6(config)#int e0/1
R6(config-if)#shut

%LINK-5-CHANGED: Interface FastEthernet0/1, changed state to
administratively down

%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed
state to down

R6(config-if)#no shut
```

To verify the configuration:

```
R6#sh ipv6 int br e0/1

Ethernet0/1            [up/up]
    FE80::6
```

Let's see, in this case R6 is going through SLAAC operation to acquire the prefix portion of its IPv6 address, in that process, it should see the Router Advertisements (RA) from R5, let's turn on the debug on R6 and

see why R6 is not getting the prefix portion of its IPv6 address:

```
R6#debug ipv6 nd
  ICMP Neighbor Discovery events debugging is on
```

Maybe R5 is not sending them, let's enable this debug on R5:

On R5:

```
R5#debug ipv6 nd
  ICMP Neighbor Discovery events debugging is on
```

These RAs should go out every 3 minutes, let's check R5's configuration:

```
R5#sh run int e0/1 | b inter

interface Ethernet0/1
 no ip address
 ipv6 address FE80::5 link-local
 ipv6 address 56::5/64
 ipv6 nd other-config-flag
 ipv6 nd ra suppress
end
```

We can see the problem, RAs are suppressed. Let's fix this problem:

```
R5(config)#int e0/1
R5(config-if)#no ipv6 nd ra suppress
```

You should see the following debug output:

```
ICMPv6-ND: FastEthernet0/1: ND Feature 1 - on
ICMPv6-ND: Created RA context for FE80::5
ICMPv6-ND: Request to send RA for FE80::5
ICMPv6-ND: Sending RA from FE80::5 to FF02::1 on FastEthernet0/1
ICMPv6-ND:     Other stateful configuration
ICMPv6-ND:     MTU = 1500
ICMPv6-ND:     prefix = 56::/64 onlink autoconfig
ICMPv6-ND:              2592000/604800 (valid/preferred)
```

Since the RAs are sent, let'see if R6 acquired the prefix portion of its IPv6 address:

On R6:

We don't see any debug output on R6, let's verify the switch configuration:

On SW2:

```
SW2#sh run int e1/1 | b inter

interface Ethernet1/1
 switchport access vlan 56
 switchport mode access
end
```

Well, R5 is configured in VLAN 56, let's verify the e1/2 interface on SW2:

```
SW2#sh run int e1/2 | b inter

interface Ethernet1/2
 switchport access vlan 56
 switchport mode access
end
```

SW2's e1/2 interface is configured in the correct VLAN, let's verify R6's e0/1 interface:

On R6:

```
R6#sh run int e0/1 | b inter

interface Ethernet0/1
 no ip address
 ipv6 address dhcp
 ipv6 address FE80::6 link-local
 ipv6 enable
end

R6#sh ipv6 int br e0/1

Ethernet0/1            [up/up]
    FE80::6
```

Let's see the configuration on E0/1 interface of R6:

```
R6#sh run int e0/1 | b inter

interface Ethernet0/1
 no ip address
 ipv6 address dhcp
```

```
 ipv6 address FE80::6 link-local
 ipv6 enable
end
```

R6 is configured to acquire an IPv6 address from a DHCP server and it's not configured for SLAAC operation, let's fix this problem:

```
R6(config)#int e0/1
R6(config-subif)#no ipv6 addr dhcp
R6(config-subif)#ipv6 address autoconfig default
```

Once this is configured, we should see the debug output stating that the local router did acquire its prefix from R5:

To verify the configuration:

```
R6#sh ipv6 inter br e0/1

Ethernet0/1            [up/up]
    FE80::6
    56::6
```

Let's see if R6 acquired its DNS and domain name configuration from the DHCP server:

```
R6#sh ipv6 dhcp inter

Ethernet0/1 is in client mode
  Prefix State is INFORMATION-REQUEST (6)
  Information refresh timer expires in 00:00:02
  Address State is IDLE
  Prefix Rapid-Commit: disabled
  Address Rapid-Commit: disabled
```

R6 is not getting its DNS and domain name parameters from the DHCP server, R1. But in order for R6 to get these information it should be configured as a relay agent, let's verify that configuration:

On R5:

```
R5#sh run int e0/1 | b inter

interface Ethernet0/1
 no ip address
 ipv6 address FE80::5 link-local
 ipv6 address 56::5/64
```

```
 ipv6 nd other-config-flag
end
```

R5 is not configured as a relay agent, let's do that:

```
R5(config)#int e0/1
R5(config-if)#ipv6 dhcp relay destination 10::1 tunnel 1
R5(config-if)#shut

%LINK-5-CHANGED: Interface FastEthernet0/1, changed state to
administratively down

%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed
state to down

R5(config-if)#no shut
```

To verify the configuration:

On R6:

```
R6#sh ipv6 dhcp inter

Ethernet0/1 is in client mode
  Prefix State is IDLE (0)
  Information refresh timer expires in 23:59:25
  Address State is IDLE
  List of known servers:
    Reachable via address: FE80::5
    DUID: 00030001AABBCC001100
    Preference: 0
    Configuration parameters:
      DNS server: 2000:56::56
      Domain name: MicronicsTraining.com
      Information refresh time: 0
  Prefix Rapid-Commit: disabled
  Address Rapid-Commit: disabled
```

On R5, and R6:

```
Rx#u all
All possible debugging has been turned off
```

Ticket 6

R7's E0/1 interface is configured with 78.1.1.7/24 and R8's E0/1.87 is configured with 87.1.1.8/24 IP addresses.
These two routers should form an adjacency and have reachability to each other's loopback 0 interface.
DO NOT change the IP addresses, or use ip unnumbered, or change the network type.
DO NOT configure PPP to accomplish this task.

In this case we can configure OSPFv3, since in OSPFv3 the packets are encapsulated in IPv6, the IP addresses can be from any network. Let's configure this and verify:

On R7:

```
R7(config)#router ospfv3 1
R7(config-router)#address-family ipv6 unicast
R7(config-router-af)#router-id 0.0.0.77

R7#cle ospfv3 pro
Reset selected OSPFv3 processes? [no]: y

R7(config)#int e0/1
R7(config-if)#ipv6 enable
R7(config-if)#ospfv3 1 ipv4 area 0

R7(config)#int lo0
R7(config-if)#ospfv3 1 ipv4 area 0
```

On R8:

```
R8(config)#router ospfv3 1
R8(config-router)#address-family ipv6 unicast
R8(config-router-af)#router-id 0.0.0.88

R8#cle ospfv3 pro
Reset selected OSPFv3 processes? [no]: y

R8(config)#int e0/1.87
R8(config-subif)#ipv6 enable
R8(config-subif)#ospfv3 1 ipv4 area 0
```

Once the above command is entered, you should see the following console message stating that the

adjacency has been established:

```
%OSPFv3-5-ADJCHG: Process 1, IPv4, Nbr 1.1.1.7 on GigabitEthernet0/1.87
from LOADING to FULL, Loading Done

%OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 0.0.0.77 on Tunnel1 from LOADING to
FULL, Loading Done

R8(config)#int lo0
R8(config-if)#ospfv3 1 ipv4 area 0
```

To verify the configuration:

On R7:

```
R7#sh ip route ospfv3 | b Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 2 subnets
O        1.1.1.8 [110/10] via 87.1.1.8, 00:00:45, Ethernet0/1
      87.0.0.0/24 is subnetted, 1 subnets
O        87.1.1.0 is directly connected, 00:00:45, Ethernet0/1

R7#ping 87.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 87.1.1.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R7#ping 1.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Ticket 7

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