CCNP ENARSI

Solution Guide Incident 05

IPv4 Preferred Path Drift

Instructor and post-investigation reference

Confirmed cause: R1 Ethernet0/0 OSPFv2 cost was raised from 10 to 100, making the R2 path less preferred.

CLI annotations: GREEN - expected healthy result; YELLOW - warning or interpretation required; RED - fault or failed requirement. Colors are workbook annotations, not native CLI colors. Variable timers and ages may differ. Yellow/red items include a one-line reason.

1. Healthy baseline and isolation

The Incident UNL was derived from the exact healthy save-config startup; its derivation manifest identifies the changed node. The healthy complete-state UNL was independently cold-booted and passed full regression before Incident testing. Each Incident runtime was cleared independently.

2. Exact fault injection

Changed node: R1. The following command sequence creates this incident from the healthy reference. The delivered Incident UNL already contains the resulting startup state.

R1

```
enable
conf t
interface Ethernet0/0
ip ospf cost 100
end
```

3. Intended blast radius

R1 IPv4 path selection to the branch; one next hop is removed from the installed route.

Unrelated services required to stay healthy: All OSPFv2 adjacencies and IPv4 connectivity; all IPv6 neighbors, ECMP and client paths.

4. Reproduced symptom: actual CLI

Actual CLI output - R1

```
R1# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           0   FULL/  -        00:00:39    10.13.13.2      Ethernet0/1
2.2.2.2           0   FULL/  -        00:00:38    10.12.12.2      Ethernet0/0
```

Actual CLI output - R1

```
R1# show ipv6 ospf neighbor
            OSPFv3 Router with ID (1.1.1.1) (Process ID 30)
Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
3.3.3.3           0   FULL/  -        00:00:31    1               Ethernet0/1
2.2.2.2           0   FULL/  -        00:00:31    1               Ethernet0/0
```

Actual CLI output - R1

```
R1# show ip route 192.168.40.0
Routing entry for 192.168.40.0/24
  Known via "ospf 10", distance 110, metric 30, type inter area
  Last update from 10.13.13.2 on Ethernet0/1, 00:03:16 ago
  Routing Descriptor Blocks:
  * 10.13.13.2, from 4.4.4.4, 00:03:16 ago, via Ethernet0/1
      Route metric is 30, traffic share count is 1
```

WARNING - Only one installed path remains where the approved baseline requires two; service may still answer through the survivor.

Actual CLI output - R1

```
R1# show ipv6 route 2001:DB8:40::/64
Routing entry for 2001:DB8:40::/64
  Known via "ospf 30", distance 110, metric 30, type inter area
  Route count is 2/2, share count 0
  Routing paths:
    FE80::A8BB:CCFF:FE00:3000, Ethernet0/1
      Route metric is 30, traffic share count is 1
      From FE80::A8BB:CCFF:FE00:3000
      Last updated 00:03:23 ago
    FE80::A8BB:CCFF:FE00:2000, Ethernet0/0
      Route metric is 30, traffic share count is 1
      From FE80::A8BB:CCFF:FE00:2000
      Last updated 00:03:23 ago
```

Actual CLI output - R1

```
R1# ping 192.168.40.10 source Loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.40.10, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/2 ms
```

Actual CLI output - R1

```
R1# ping 2001:DB8:40::10 source Loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:40::10, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:1::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms
```

5. Diagnostic workflow and reasoning

Compare the failing family with the other family and with the approved two-path route. Then inspect the affected interface or OSPF process on the changed node. The following commands were run on the broken instance in this order.

R1

```
enable
terminal length 0
show ip ospf neighbor
show ip route 192.168.40.0
show ip ospf interface Ethernet0/0
show running-config | section interface Ethernet0/0
```

Actual CLI output - R1

```
R1# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           0   FULL/  -        00:00:38    10.13.13.2      Ethernet0/1
2.2.2.2           0   FULL/  -        00:00:37    10.12.12.2      Ethernet0/0
```

Actual CLI output - R1

```
R1# show ip route 192.168.40.0
Routing entry for 192.168.40.0/24
  Known via "ospf 10", distance 110, metric 30, type inter area
  Last update from 10.13.13.2 on Ethernet0/1, 00:03:18 ago
  Routing Descriptor Blocks:
  * 10.13.13.2, from 4.4.4.4, 00:03:18 ago, via Ethernet0/1
      Route metric is 30, traffic share count is 1
```

WARNING - Only one installed path remains where the approved baseline requires two; service may still answer through the survivor.

Actual CLI output - R1

```
R1# show ip ospf interface Ethernet0/0
Ethernet0/0 is up, line protocol is up
  Internet Address 10.12.12.1/30, Interface ID 2, Area 0
  Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 100
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.2.2
```

Actual CLI output - R1

```
R1# show running-config | section interface Ethernet0/0
interface Ethernet0/0
 description TO_R2
 ip address 10.12.12.1 255.255.255.252
 ip ospf network point-to-point
 ip ospf 10 area 0
 ip ospf cost 100
 ipv6 address 2001:DB8:12::1/64
 ipv6 ospf 30 area 0
 ipv6 ospf network point-to-point
 no passive-interface Ethernet0/0
 no passive-interface Ethernet0/0
```

Confirmed cause: R1 Ethernet0/0 OSPFv2 cost was raised from 10 to 100, making the R2 path less preferred.

6. Minimal corrective action

Change only the identified setting on R1. Allow OSPF to converge before checking route installation.

R1

```
enable
conf t
interface Ethernet0/0
ip ospf cost 10
end
```

7. Restoration and complete regression

The matching live repair was applied to the broken instance. Both address families and every router adjacency were then rechecked; the next excerpts show the decisive restored evidence.

Actual CLI output - R1

```
R1# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           0   FULL/  -        00:00:39    10.13.13.2      Ethernet0/1
2.2.2.2           0   FULL/  -        00:00:38    10.12.12.2      Ethernet0/0
```

Actual CLI output - R1

```
R1# show ipv6 ospf neighbor
            OSPFv3 Router with ID (1.1.1.1) (Process ID 30)
Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
3.3.3.3           0   FULL/  -        00:00:31    1               Ethernet0/1
2.2.2.2           0   FULL/  -        00:00:39    1               Ethernet0/0
```

Actual CLI output - R1

```
R1# ping 192.168.40.10 source Loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.40.10, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

Actual CLI output - R1

```
R1# ping 2001:DB8:40::10 source Loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:40::10, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:1::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

Actual CLI output - R1

```
R1# show ip route 192.168.40.0
Routing entry for 192.168.40.0/24
  Known via "ospf 10", distance 110, metric 30, type inter area
  Last update from 10.12.12.2 on Ethernet0/0, 00:00:56 ago
  Routing Descriptor Blocks:
  * 10.13.13.2, from 4.4.4.4, 00:04:14 ago, via Ethernet0/1
      Route metric is 30, traffic share count is 1
    10.12.12.2, from 4.4.4.4, 00:00:56 ago, via Ethernet0/0
      Route metric is 30, traffic share count is 1
```

Actual CLI output - R1

```
R1# show ipv6 route 2001:DB8:40::/64
Routing entry for 2001:DB8:40::/64
  Known via "ospf 30", distance 110, metric 30, type inter area
  Route count is 2/2, share count 0
  Routing paths:
    FE80::A8BB:CCFF:FE00:3000, Ethernet0/1
      Route metric is 30, traffic share count is 1
      From FE80::A8BB:CCFF:FE00:3000
      Last updated 00:04:21 ago
    FE80::A8BB:CCFF:FE00:2000, Ethernet0/0
      Route metric is 30, traffic share count is 1
      From FE80::A8BB:CCFF:FE00:2000
      Last updated 00:04:21 ago
```

Actual CLI output - H4

```
H4# ping 192.168.40.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.40.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Actual CLI output - H4

```
H4# ping 2001:DB8:40::1 source Ethernet0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:40::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:40::10
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Actual CLI output - H4

```
H4# ping 1.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

Actual CLI output - H4

```
H4# ping 2001:DB8:1::1 source Ethernet0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:1::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:40::10
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

8. Save and handoff

All five nodes were saved only after full regression. The damaged Incident file remains a separate starting state; opening another ticket requires a fresh isolated boot from its own UNL. The release report contains the complete node-by-node console records and hashes.

Reference: Cisco 300-410 ENARSI v1.1 blueprint.