CCNA 200-301 v1.1

MODULE 3 — IP CONNECTIVITY

Complete Hands-on Training Labs

Topology • Sequential Device Configuration • Verification Output • Troubleshooting

NETWORK GUIDE ACADEMY

Instructor / Student Lab Guide

Scope: CCNA v1.1 Module 3 — IP Connectivity objectives 3.1 through 3.5. The labs intentionally use one configuration path per device and then move forward to verification and troubleshooting.

Module 3 Lab Map

- 3.1 Routing table and forwarding-plane interpretation — route codes, prefix/mask, next hop, administrative distance, metric, gateway of last resort and CEF.
- 3.2 Router forwarding decision — longest prefix match first, administrative distance second; OSPF metric selection is validated in Lab 3.4.
- 3.3A IPv4 static routing — default route, network route, host route and floating static backup.
- 3.3B IPv6 static routing — default route, network route, host route and floating static backup.
- 3.4 Single-area OSPFv2 — neighbor adjacency, point-to-point network type, broadcast DR/BDR election, router ID and route cost.
- 3.5 First-hop redundancy — HSRP, VRRP and GLBP using one dual-distribution topology with failover verification.

How to use the detailed device configurations

Every lab follows one forward sequence: topology and addressing → Device 1 complete configuration and verification → Device 2 / remaining devices → end-to-end validation → controlled fault test → recovery verification. A command that already appears in the full device configuration is not repeated later unless the later command intentionally changes state for troubleshooting.

Verification output note

CLI output is representative Cisco IOS / IOS XE output. Exact interface hardware strings, uptime, MAC addresses, timers and minor formatting vary by platform and software release. Verify the fields and state identified in the What to verify section rather than memorizing every character of the sample output.

Recommended lab platforms

- Cisco Packet Tracer for routing-table reading, IPv4/IPv6 static routes, basic OSPF and HSRP where supported.
- Cisco Modeling Labs (CML), EVE-NG or GNS3 with a compatible IOS/IOS XE image for complete OSPF, VRRP and GLBP behavior.
- Physical IOS/IOS XE routers and multilayer switches when available. Adapt GigabitEthernet numbering to the hardware in use.

Standard verification workflow

1. Confirm every routed interface is up/up and has the intended IPv4/IPv6 address.
2. Read the routing table before testing traffic. Confirm the selected route, next hop, AD/metric and default route.
3. Use route-specific show commands and CEF/forwarding information to prove the router will send the packet where expected.
4. Run ping and traceroute only after the control-plane state is correct.
5. Introduce one controlled failure, predict the new route or FHRP state, then verify that prediction.
6. Restore the failed component and prove the original state returns.

Lab 3.1 — Interpret the Routing Table and CEF Forwarding Entry

CCNA objective: Interpret routing-table components including routing protocol code, prefix, network mask, next hop, administrative distance, metric and gateway of last resort.

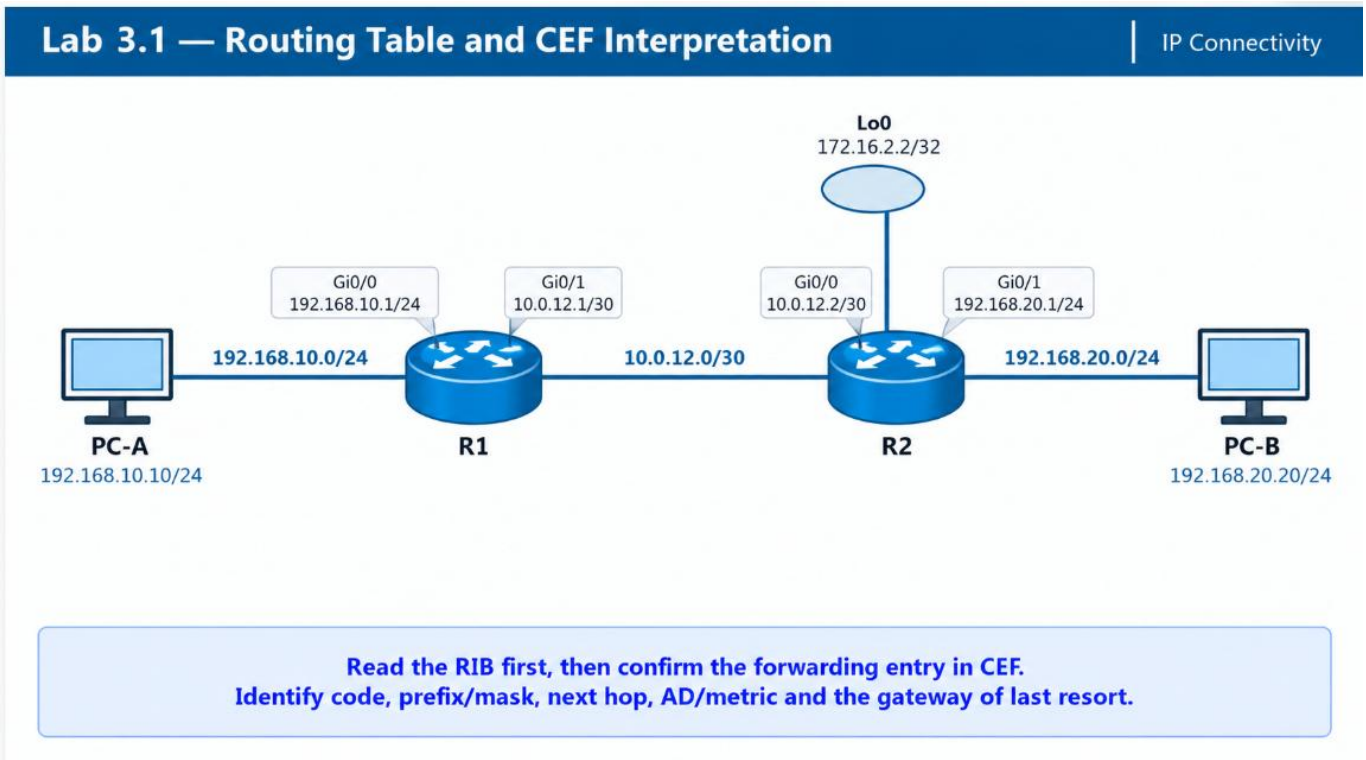

Figure 3.1: Routing Table and CEF Interpretation reference topology

Lab objectives

- Build a small routed network with connected, local, static, host and default routes.
- Read the route code, prefix length/mask, next hop and outgoing interface.
- Interpret the [administrative-distance/metric] values on a static route.
- Identify the gateway of last resort and the S* default route.
- Compare the routing table (RIB) with the Cisco Express Forwarding entry used by the data plane.

Addressing / interface plan

Device	Interface	Address / Mask	Purpose
PC-A	NIC	192.168.10.10/24	User LAN; GW 192.168.10.1
R1	Gi0/0/0	192.168.10.1/24	PC-A LAN
R1	Gi0/0/1	10.0.12.1/30	Transit to R2
R2	Gi0/0/0	10.0.12.2/30	Transit to R1
R2	Gi0/0/1	192.168.20.1/24	PC-B LAN
R2	Lo0	172.16.2.2/32	Host-route target
PC-B	NIC	192.168.20.20/24	Remote LAN; GW 192.168.20.1

Execution sequence

Configure each device once in the order shown below. After the configuration is complete, use the verification output and the What to verify bullets to prove the expected state. Configuration is not repeated later; subsequent commands are only verification or controlled fault tests.

Lab execution — device-by-device configuration and verification

R1 — Complete lab configuration

```
enable
configure terminal
hostname R1
!
interface GigabitEthernet0/0/0
 description LAN_PC_A
 ip address 192.168.10.1 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/0/1
 description TRANSIT_TO_R2
 ip address 10.0.12.1 255.255.255.252
 no shutdown
!
ip route 192.168.20.0 255.255.255.0 10.0.12.2
ip route 172.16.2.2 255.255.255.255 10.0.12.2
ip route 0.0.0.0 0.0.0.0 10.0.12.2
end
write memory
```

R1 — Verification output

```
R1# show ip interface brief
Interface              IP-Address      OK? Method Status                Protocol
GigabitEthernet0/0/0   192.168.10.1    YES manual up                    up
GigabitEthernet0/0/1   10.0.12.1       YES manual up                    up

R1# show ip route
Gateway of last resort is 10.0.12.2 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 10.0.12.2
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.0.12.0/30 is directly connected, GigabitEthernet0/0/1
L        10.0.12.1/32 is directly connected, GigabitEthernet0/0/1
S     172.16.2.2/32 [1/0] via 10.0.12.2
      192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.10.0/24 is directly connected, GigabitEthernet0/0/0
L        192.168.10.1/32 is directly connected, GigabitEthernet0/0/0
S     192.168.20.0/24 [1/0] via 10.0.12.2

R1# show ip route 192.168.20.20
Routing entry for 192.168.20.0/24
  Known via "static", distance 1, metric 0
  Routing Descriptor Blocks:
  * 10.0.12.2
      Route metric is 0, traffic share count is 1

R1# show ip cef 192.168.20.20
192.168.20.0/24
  nexthop 10.0.12.2 GigabitEthernet0/0/1
```

What to verify

- Gi0/0/0 and Gi0/0/1 are up/up. A route cannot use a next hop reached through a failed interface.
- C identifies the directly connected network and L identifies R1's own /32 interface address.
- S identifies a static route. S* identifies the candidate default route.
- The value [1/0] means administrative distance 1 and metric 0 for the static route.
- The gateway of last resort is 10.0.12.2 because the default route points to R2.
- The CEF entry forwards 192.168.20.20 to next hop 10.0.12.2 out Gi0/0/1, matching the RIB.

R2 — Complete lab configuration

```
enable
configure terminal
hostname R2
!
interface GigabitEthernet0/0/0
 description TRANSIT_TO_R1
 ip address 10.0.12.2 255.255.255.252
 no shutdown
!
interface GigabitEthernet0/0/1
 description LAN_PC_B
 ip address 192.168.20.1 255.255.255.0
 no shutdown
!
interface Loopback0
 description HOST_ROUTE_TARGET
 ip address 172.16.2.2 255.255.255.255
!
ip route 192.168.10.0 255.255.255.0 10.0.12.1
end
write memory
```

R2 — Verification output

```
R2# show ip route static
S    192.168.10.0/24 [1/0] via 10.0.12.1

R2# ping 192.168.10.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.10.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5)
```

What to verify

- R2 has a return route to the PC-A LAN through R1. Without the return route, forward traffic may arrive but replies fail.
- The successful ping proves two-way routing after PC-A is configured with default gateway 192.168.10.1.

End-device addressing

- PC-A: 192.168.10.10/24, default gateway 192.168.10.1.
- PC-B: 192.168.20.20/24, default gateway 192.168.20.1.

End-to-end verification

```
PC-A> ping 192.168.20.20
Reply from 192.168.20.20: bytes=32 time<1ms TTL=126

PC-A> tracert 192.168.20.20
  1  192.168.10.1
  2  10.0.12.2
  3  192.168.20.20
```

What to verify

- PC-A reaches PC-B through R1 and R2.
- The first router hop is the PC default gateway; the next routed hop is R2.
- A successful ping confirms reachability, while the routing-table and CEF checks prove why it succeeds.

Lab 3.2 — Determine the Router Forwarding Decision

CCNA objective: Determine how a router makes a forwarding decision by default using longest prefix match, administrative distance and routing-protocol metric.

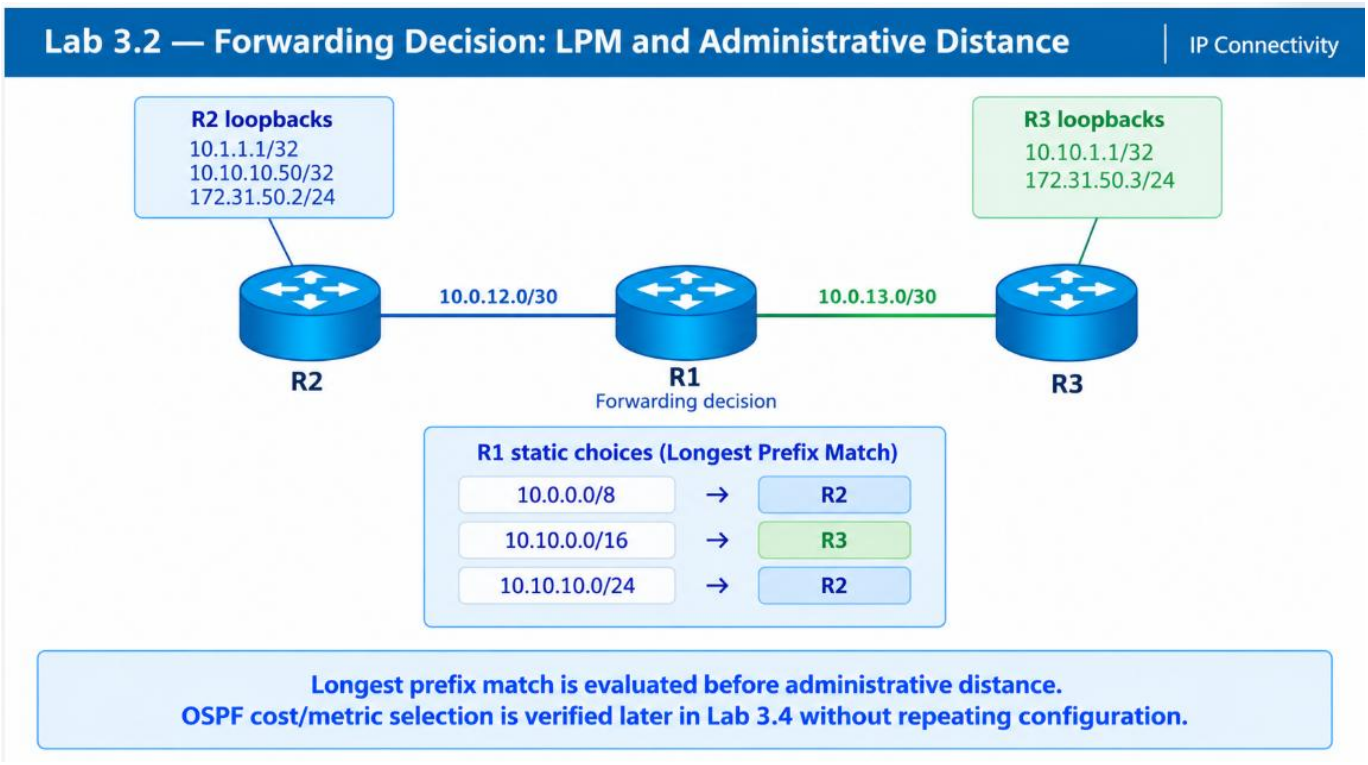

Figure 3.2: Longest Prefix Match and Administrative Distance reference topology

Lab objectives

- Prove that the most-specific route wins even when a less-specific route points to a different next hop.
- Create two routes to the same prefix with different administrative distances and observe which route enters the RIB.
- Fail the preferred next hop and verify the higher-AD route becomes active.
- Carry the metric portion of the forwarding decision forward to Lab 3.4, where two OSPF paths are compared by cost.

Addressing / interface plan

Device	Interface	Address / Mask	Purpose
R1	Gi0/0/0	10.0.12.1/30	To R2
R1	Gi0/0/1	10.0.13.1/30	To R3
R1	Lo0	192.0.2.1/32	Ping source
R2	Gi0/0/0	10.0.12.2/30	To R1
R2	Lo10	10.10.10.50/32	Matches /24 route
R2	Lo20	10.1.1.1/32	Matches only /8 route
R2	Lo50	172.31.50.2/24	Preferred AD test network
R3	Gi0/0/0	10.0.13.2/30	To R1
R3	Lo10	10.10.1.1/32	Matches /16 route
R3	Lo50	172.31.50.3/24	Backup AD test network

Execution sequence

Configure each device once in the order shown below. After the configuration is complete, use the verification output and the What to verify bullets to prove the expected state. Configuration is not repeated later; subsequent commands are only verification or controlled fault tests.

Lab execution — device-by-device configuration and verification

R1 — Complete lab configuration

```
enable
configure terminal
hostname R1
!
interface GigabitEthernet0/0/0
 description TO_R2
 ip address 10.0.12.1 255.255.255.252
 no shutdown
!
interface GigabitEthernet0/0/1
 description TO_R3
 ip address 10.0.13.1 255.255.255.252
 no shutdown
!
interface Loopback0
 ip address 192.0.2.1 255.255.255.255
!
ip route 10.0.0.0 255.0.0.0 10.0.12.2
ip route 10.10.0.0 255.255.0.0 10.0.13.2
ip route 10.10.10.0 255.255.255.0 10.0.12.2
!
ip route 172.31.50.0 255.255.255.0 10.0.12.2 1
ip route 172.31.50.0 255.255.255.0 10.0.13.2 5
end
write memory
```

R1 — Verification output

```
R1# show ip route 10.1.1.1
Routing entry for 10.0.0.0/8
  Known via "static", distance 1, metric 0
  * 10.0.12.2

R1# show ip route 10.10.1.1
Routing entry for 10.10.0.0/16
  Known via "static", distance 1, metric 0
  * 10.0.13.2

R1# show ip route 10.10.10.50
Routing entry for 10.10.10.0/24
  Known via "static", distance 1, metric 0
  * 10.0.12.2

R1# show ip route 172.31.50.0
Routing entry for 172.31.50.0/24
  Known via "static", distance 1, metric 0
  * 10.0.12.2
```

What to verify

- 10.1.1.1 matches the /8 route because no longer matching route exists.
- 10.10.1.1 matches /16 instead of /8 because /16 is more specific.

- 10.10.10.50 matches /24 instead of /16 or /8 because /24 is the longest prefix match.
- For 172.31.50.0/24, both configured routes have the same prefix length, so administrative distance decides. AD 1 through R2 is installed; AD 5 through R3 stays out of the RIB.

R2 — Complete lab configuration

```
enable
configure terminal
hostname R2
!
interface GigabitEthernet0/0/0
 ip address 10.0.12.2 255.255.255.252
 no shutdown
!
interface Loopback10
 ip address 10.10.10.50 255.255.255.255
interface Loopback20
 ip address 10.1.1.1 255.255.255.255
interface Loopback50
 ip address 172.31.50.2 255.255.255.0
!
ip route 192.0.2.1 255.255.255.255 10.0.12.1
end
write memory
```

R2 — Verification output

```
R2# show ip interface brief | include up
GigabitEthernet0/0/0   10.0.12.2      YES manual up                    up
Loopback10             10.10.10.50    YES manual up                    up
Loopback20             10.1.1.1       YES manual up                    up
Loopback50             172.31.50.2    YES manual up                    up

R2# ping 192.0.2.1 source Loopback10
!!!!!
```

What to verify

- The /8 and /24 test destinations are active loopbacks on R2.
- The return host route to R1's loopback allows sourced pings to complete.

R3 — Complete lab configuration

```
enable
configure terminal
hostname R3
!
interface GigabitEthernet0/0/0
 ip address 10.0.13.2 255.255.255.252
 no shutdown
!
interface Loopback10
 ip address 10.10.1.1 255.255.255.255
interface Loopback50
 ip address 172.31.50.3 255.255.255.0
!
ip route 192.0.2.1 255.255.255.255 10.0.13.1
end
write memory
```

R3 — Verification output

```
R3# show ip interface brief | include up
GigabitEthernet0/0/0   10.0.13.2      YES manual up                    up
Loopback10             10.10.1.1      YES manual up                    up
Loopback50             172.31.50.3    YES manual up                    up

R3# ping 192.0.2.1 source Loopback10
!!!!!
```

What to verify

- The /16 destination 10.10.1.1 is reachable through R3.
- R3 also owns the backup 172.31.50.0/24 path that becomes active only after the preferred R2 next hop is unavailable.

Controlled failure — administrative-distance fallback

```
R1# configure terminal
R1(config)# interface GigabitEthernet0/0/0
R1(config-if)# shutdown
R1(config-if)# end

R1# show ip route 172.31.50.0
Routing entry for 172.31.50.0/24
  Known via "static", distance 5, metric 0
  * 10.0.13.2
```

What to verify

- The primary next hop 10.0.12.2 is no longer reachable because the R1–R2 interface is down.
- The same /24 route is now installed through R3 with administrative distance 5.
- Restore the link with no shutdown and verify the AD 1 route through R2 returns.

Lab 3.3A — Configure and Verify IPv4 Static Routing

CCNA objective: Configure and verify IPv4 static routes including default, network, host and floating static routes.

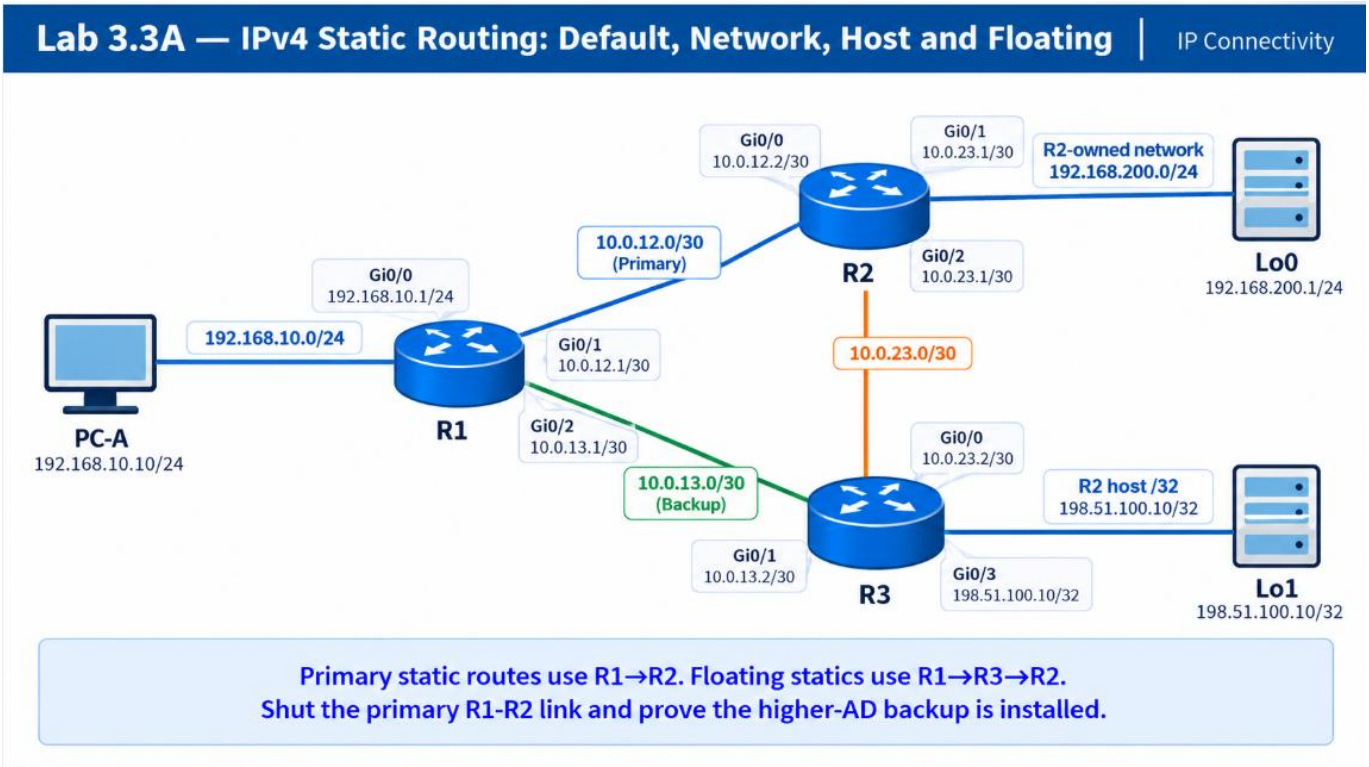

Figure 3.3A: IPv4 Static Routing with a Primary and Floating Backup Path

Lab objectives

- Configure a network static route to an R2-owned loopback network.
- Configure a /32 host route to a specific R2 loopback address.
- Configure a default route from R1 toward R2.
- Configure higher-administrative-distance floating static routes through R3.
- Fail the directly connected primary path and prove the floating routes enter the routing table.

Addressing / interface plan

Device	Interface	Address / Mask	Purpose
PC-A	NIC	192.168.10.10/24	GW 192.168.10.1
R1	Gi0/0/0	192.168.10.1/24	LAN
R1	Gi0/0/1	10.0.12.1/30	Primary to R2
R1	Gi0/0/2	10.0.13.1/30	Backup to R3
R2	Gi0/0/0	10.0.12.2/30	Primary to R1
R2	Gi0/0/1	10.0.23.1/30	To R3
R2	Lo0	192.168.200.1/24	Network-route target
R2	Lo1	198.51.100.10/32	Host-route target
R3	Gi0/0/0	10.0.13.2/30	To R1
R3	Gi0/0/1	10.0.23.2/30	To R2

Execution sequence

Configure each device once in the order shown below. After the configuration is complete, use the verification output and the What to verify bullets to prove the expected state. Configuration is not repeated later; subsequent commands are only verification or controlled fault tests.

Lab execution — device-by-device configuration and verification

R1 — Complete lab configuration

```
enable
configure terminal
hostname R1
!
interface GigabitEthernet0/0/0
 description LAN_PC_A
 ip address 192.168.10.1 255.255.255.0
 no shutdown
interface GigabitEthernet0/0/1
 description PRIMARY_TO_R2
 ip address 10.0.12.1 255.255.255.252
 no shutdown
interface GigabitEthernet0/0/2
 description BACKUP_TO_R3
 ip address 10.0.13.1 255.255.255.252
 no shutdown
!
ip route 192.168.200.0 255.255.255.0 10.0.12.2
ip route 192.168.200.0 255.255.255.0 10.0.13.2 10
ip route 198.51.100.10 255.255.255.255 10.0.12.2
ip route 198.51.100.10 255.255.255.255 10.0.13.2 10
ip route 0.0.0.0 0.0.0.0 10.0.12.2
end
write memory
```

R1 — Verification output

```
R1# show ip route static
S*    0.0.0.0/0 [1/0] via 10.0.12.2
S     192.168.200.0/24 [1/0] via 10.0.12.2
S     198.51.100.10/32 [1/0] via 10.0.12.2

R1# show running-config | include ^ip route
ip route 192.168.200.0 255.255.255.0 10.0.12.2
ip route 192.168.200.0 255.255.255.0 10.0.13.2 10
ip route 198.51.100.10 255.255.255.255 10.0.12.2
ip route 198.51.100.10 255.255.255.255 10.0.13.2 10
ip route 0.0.0.0 0.0.0.0 10.0.12.2
```

What to verify

- The default route is marked S* and points to R2.
- The network route 192.168.200.0/24 and host route 198.51.100.10/32 use the primary next hop 10.0.12.2.
- The floating routes are present in running configuration but not in the RIB while the AD 1 primary routes are valid.
- The floating routes use AD 10, which is worse than the normal static AD 1.

R2 — Complete lab configuration

```
enable
configure terminal
hostname R2
!
interface GigabitEthernet0/0/0
 description PRIMARY_TO_R1
 ip address 10.0.12.2 255.255.255.252
 no shutdown
interface GigabitEthernet0/0/1
 description TO_R3
 ip address 10.0.23.1 255.255.255.252
 no shutdown
interface Loopback0
 description NETWORK_ROUTE_TARGET
 ip address 192.168.200.1 255.255.255.0
interface Loopback1
 description HOST_ROUTE_TARGET
 ip address 198.51.100.10 255.255.255.255
!
ip route 192.168.10.0 255.255.255.0 10.0.12.1
ip route 192.168.10.0 255.255.255.0 10.0.23.2 10
end
write memory
```

R2 — Verification output

```
R2# show ip route static
S    192.168.10.0/24 [1/0] via 10.0.12.1

R2# ping 192.168.10.10 source Loopback0
!!!!!
```

What to verify

- R2 uses the direct R1 path as the preferred return route to PC-A.
- The AD 10 return route through R3 is configured but inactive while 10.0.12.1 is directly reachable.
- The sourced ping confirms return routing from an R2-owned destination network.

R3 — Complete lab configuration

```
enable
configure terminal
hostname R3
!
interface GigabitEthernet0/0/0
 description TO_R1
 ip address 10.0.13.2 255.255.255.252
 no shutdown
interface GigabitEthernet0/0/1
 description TO_R2
 ip address 10.0.23.2 255.255.255.252
 no shutdown
!
ip route 192.168.200.0 255.255.255.0 10.0.23.1
ip route 198.51.100.10 255.255.255.255 10.0.23.1
ip route 192.168.10.0 255.255.255.0 10.0.13.1
end
write memory
```

R3 — Verification output

```
R3# show ip route static
S    192.168.10.0/24 [1/0] via 10.0.13.1
S    192.168.200.0/24 [1/0] via 10.0.23.1
S    198.51.100.10/32 [1/0] via 10.0.23.1
```

What to verify

- R3 knows how to reach the R2-owned network and host when it is used as the backup transit router.
- R3 also has a direct return static route to the PC-A LAN through R1.

End-device addressing

- PC-A: 192.168.10.10/24, default gateway 192.168.10.1.

Normal-path end-to-end verification

```
PC-A> ping 192.168.200.1
Reply from 192.168.200.1: bytes=32 time<1ms TTL=254

PC-A> ping 198.51.100.10
Reply from 198.51.100.10: bytes=32 time<1ms TTL=254

R1# traceroute 198.51.100.10
  1 10.0.12.2
  2 198.51.100.10
```

What to verify

- Both the network target and host target are reachable.
- The normal path from R1 uses 10.0.12.2 directly, not R3.

Controlled failure — activate the floating static routes

```
R1# configure terminal
R1(config)# interface GigabitEthernet0/0/1
R1(config-if)# shutdown
R1(config-if)# end

R1# show ip route static
S    192.168.200.0/24 [10/0] via 10.0.13.2
S    198.51.100.10/32 [10/0] via 10.0.13.2

R1# traceroute 198.51.100.10
  1 10.0.13.2
  2 10.0.23.1
  3 198.51.100.10
```

What to verify

- The primary next hop 10.0.12.2 disappears when Gi0/0/1 is shut.
- The floating network and host routes become active with AD 10 via R3.
- The traceroute now crosses R3 and then R2.
- The unprotected default route disappears because only the network and host routes were given floating backups.
- Restore Gi0/0/1 with no shutdown and confirm AD 1 routes return.

Lab 3.3B — Configure and Verify IPv6 Static Routing

CCNA objective: Configure and verify IPv6 static routes including default, network, host and floating static routes.

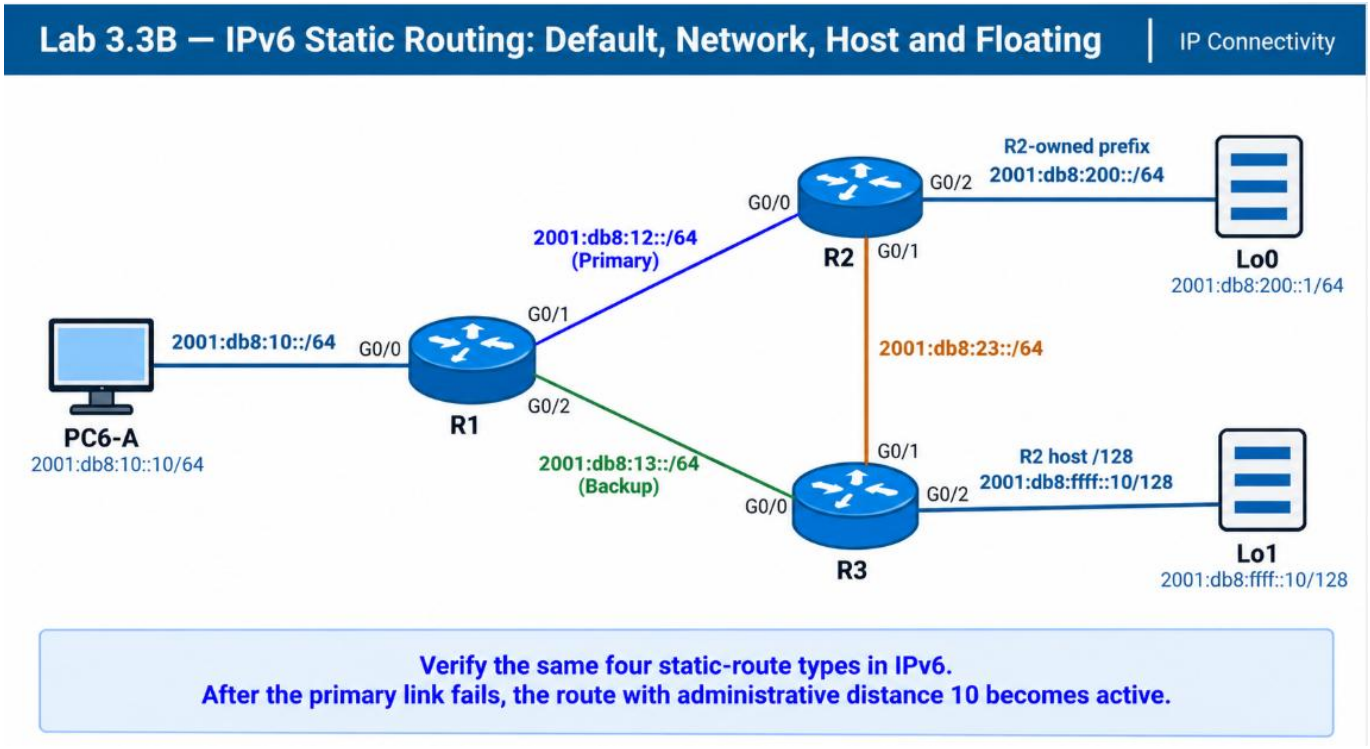

Figure 3.3B: IPv6 Static Routing with a Primary and Floating Backup Path

Lab objectives

- Enable IPv6 routing and configure global-unicast addresses on each routed link.
- Configure an IPv6 network route and /128 host route through the primary R1–R2 path.
- Configure ::/0 as an IPv6 default route.
- Configure floating IPv6 static routes through R3 with AD 10.
- Fail the primary link and verify the higher-AD routes become active.

Addressing / interface plan

Device	Interface	Address / Mask	Purpose
PC6-A	NIC	2001:db8:10::10/64	GW 2001:db8:10::1
R1	Gi0/0/0	2001:db8:10::1/64	LAN
R1	Gi0/0/1	2001:db8:12::1/64	Primary to R2
R1	Gi0/0/2	2001:db8:13::1/64	Backup to R3
R2	Gi0/0/0	2001:db8:12::2/64	Primary to R1
R2	Gi0/0/1	2001:db8:23::1/64	To R3
R2	Lo0	2001:db8:200::1/64	Network-route target
R2	Lo1	2001:db8:ffff::10/128	Host-route target
R3	Gi0/0/0	2001:db8:13::2/64	To R1
R3	Gi0/0/1	2001:db8:23::2/64	To R2

Execution sequence

Configure each device once in the order shown below. After the configuration is complete, use the verification output and the What to verify bullets to prove the expected state. Configuration is not repeated later; subsequent commands are only verification or controlled fault tests.

Lab execution — device-by-device configuration and verification

R1 — Complete lab configuration

```
enable
configure terminal
hostname R1
ipv6 unicast-routing
!
interface GigabitEthernet0/0/0
 ipv6 address 2001:DB8:10::1/64
 no shutdown
interface GigabitEthernet0/0/1
 ipv6 address 2001:DB8:12::1/64
 no shutdown
interface GigabitEthernet0/0/2
 ipv6 address 2001:DB8:13::1/64
 no shutdown
!
ipv6 route 2001:DB8:200::/64 2001:DB8:12::2
ipv6 route 2001:DB8:200::/64 2001:DB8:13::2 10
ipv6 route 2001:DB8:FFFF::10/128 2001:DB8:12::2
ipv6 route 2001:DB8:FFFF::10/128 2001:DB8:13::2 10
ipv6 route ::/0 2001:DB8:12::2
end
write memory
```

R1 — Verification output

```
R1# show ipv6 route static
S   ::/0 [1/0]
     via 2001:DB8:12::2
S   2001:DB8:200::/64 [1/0]
     via 2001:DB8:12::2
S   2001:DB8:FFFF::10/128 [1/0]
     via 2001:DB8:12::2

R1# show running-config | include ^ipv6 route
ipv6 route 2001:DB8:200::/64 2001:DB8:12::2
ipv6 route 2001:DB8:200::/64 2001:DB8:13::2 10
ipv6 route 2001:DB8:FFFF::10/128 2001:DB8:12::2
ipv6 route 2001:DB8:FFFF::10/128 2001:DB8:13::2 10
ipv6 route ::/0 2001:DB8:12::2
```

What to verify

- The active IPv6 static routes use the primary next hop 2001:db8:12::2 with AD 1.
- The /64 network route, /128 host route and ::/0 default route are all visible.
- The AD 10 floating routes are configured but do not enter the IPv6 routing table until the preferred route is invalid.

R2 — Complete lab configuration

```
enable
configure terminal
hostname R2
ipv6 unicast-routing
!
interface GigabitEthernet0/0/0
 ipv6 address 2001:DB8:12::2/64
 no shutdown
interface GigabitEthernet0/0/1
 ipv6 address 2001:DB8:23::1/64
 no shutdown
interface Loopback0
 ipv6 address 2001:DB8:200::1/64
interface Loopback1
 ipv6 address 2001:DB8:FFFF::10/128
!
ipv6 route 2001:DB8:10::/64 2001:DB8:12::1
ipv6 route 2001:DB8:10::/64 2001:DB8:23::2 10
end
write memory
```

R2 — Verification output

```
R2# show ipv6 route static
S   2001:DB8:10::/64 [1/0]
     via 2001:DB8:12::1

R2# ping 2001:DB8:10::10 source 2001:DB8:200::1
!!!!!
```

What to verify

- R2 has a preferred return route to the PC6-A LAN through the direct R1 link.
- The floating AD 10 return path through R3 remains inactive until the primary adjacency fails.
- A sourced IPv6 ping confirms two-way routing.

R3 — Complete lab configuration

```
enable
configure terminal
hostname R3
ipv6 unicast-routing
!
interface GigabitEthernet0/0/0
 ipv6 address 2001:DB8:13::2/64
 no shutdown
interface GigabitEthernet0/0/1
 ipv6 address 2001:DB8:23::2/64
 no shutdown
!
ipv6 route 2001:DB8:200::/64 2001:DB8:23::1
ipv6 route 2001:DB8:FFFF::10/128 2001:DB8:23::1
ipv6 route 2001:DB8:10::/64 2001:DB8:13::1
end
write memory
```

R3 — Verification output

```
R3# show ipv6 route static
S   2001:DB8:10::/64 [1/0]
     via 2001:DB8:13::1
S   2001:DB8:200::/64 [1/0]
     via 2001:DB8:23::1
S   2001:DB8:FFFF::10/128 [1/0]
     via 2001:DB8:23::1
```

What to verify

- R3 has complete IPv6 transit reachability between R1's LAN and R2's test prefixes.
- These routes are required before R3 can carry traffic after R1 switches to its floating static routes.

End-device addressing

- PC6-A: 2001:db8:10::10/64, default gateway 2001:db8:10::1.

Normal-path IPv6 verification

```
PC6-A> ping 2001:db8:200::1
Reply from 2001:DB8:200::1: time<1ms

PC6-A> ping 2001:db8:ffff::10
Reply from 2001:DB8:FFFF::10: time<1ms

R1# traceroute ipv6 2001:DB8:FFFF::10
  1 2001:DB8:12::2
  2 2001:DB8:FFFF::10
```

What to verify

- The network-prefix and host-prefix destinations are reachable through R2.
- The first routed IPv6 hop from R1 is 2001:db8:12::2 on the primary link.

Controlled failure — activate the IPv6 floating routes

```
R1# configure terminal
R1(config)# interface GigabitEthernet0/0/1
R1(config-if)# shutdown
R1(config-if)# end

R1# show ipv6 route static
S   2001:DB8:200::/64 [10/0]
     via 2001:DB8:13::2
S   2001:DB8:FFFF::10/128 [10/0]
     via 2001:DB8:13::2

R1# traceroute ipv6 2001:DB8:FFFF::10
  1 2001:DB8:13::2
  2 2001:DB8:23::1
  3 2001:DB8:FFFF::10
```

What to verify

- The primary next hop is removed when the R1–R2 link goes down.
- The same destination prefixes are installed via R3 with AD 10.
- The IPv6 traceroute now crosses R3 before reaching R2.
- The default route ::/0 disappears because no floating default was configured.

- Restore Gi0/0/1 and verify the AD 1 primary routes return.

Lab 3.4 — Configure and Verify Single-Area OSPFv2

CCNA objective: Configure and verify single-area OSPFv2 neighbor adjacencies, point-to-point links, broadcast DR/BDR election and router IDs.

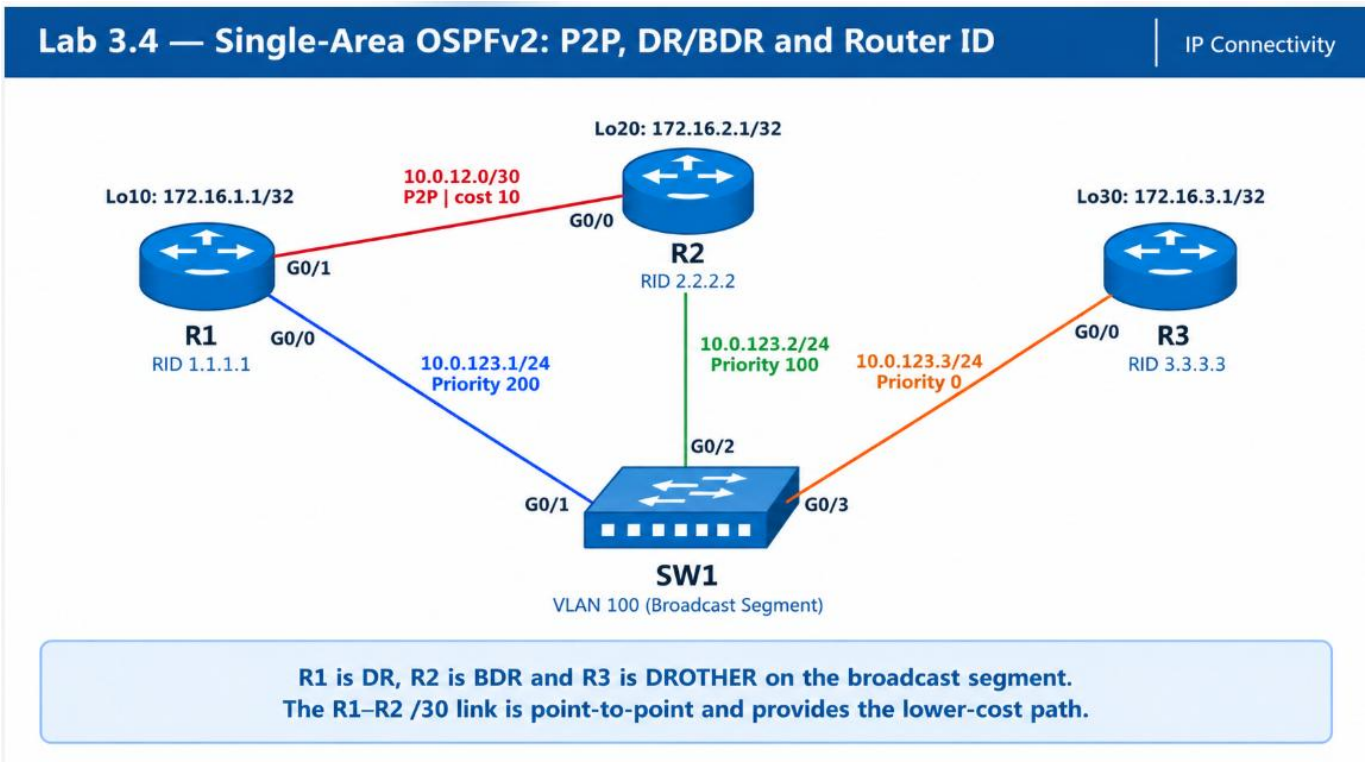

Figure 3.4: Single-Area OSPFv2 Broadcast and Point-to-Point reference topology

Lab objectives

- Form OSPF neighbors across a shared Ethernet broadcast network and a separate point-to-point link.
- Set explicit router IDs so election and verification are deterministic.
- Control the broadcast DR/BDR election with OSPF interface priority.
- Use ip ospf network point-to-point on the R1–R2 /30 link.
- Use OSPF interface cost to create a preferred path and verify the metric shown in the routing table.
- Fail the lower-cost point-to-point link and verify OSPF reconverges to the higher-cost broadcast path.

Addressing / interface plan

Device	Interface	Address / Mask	Purpose
SW1	Gi0/1-3	VLAN 100	OSPF broadcast segment
R1	Gi0/0/0	10.0.123.1/24	Broadcast; priority 200; cost 20
R1	Gi0/0/1	10.0.12.1/30	P2P to R2; cost 10
R1	Lo0	1.1.1.1/32	Router ID
R1	Lo10	172.16.1.1/32	Advertised stub
R2	Gi0/0/0	10.0.123.2/24	Broadcast; priority 100; cost 20
R2	Gi0/0/1	10.0.12.2/30	P2P to R1; cost 10
R2	Lo0	2.2.2.2/32	Router ID
R2	Lo20	172.16.2.1/32	Advertised stub
R3	Gi0/0/0	10.0.123.3/24	Broadcast; priority 0; cost 20
R3	Lo0	3.3.3.3/32	Router ID

Device	Interface	Address / Mask	Purpose
R3	Lo30	172.16.3.1/32	Advertised stub

Execution sequence

Configure each device once in the order shown below. After the configuration is complete, use the verification output and the What to verify bullets to prove the expected state. Configuration is not repeated later; subsequent commands are only verification or controlled fault tests.

Lab execution — device-by-device configuration and verification

SW1 — Complete lab configuration

```
enable
configure terminal
hostname SW1
!
vlan 100
 name OSPF_TRANSIT
!
interface range GigabitEthernet0/1 - 3
 description OSPF_ROUTER_TRANSIT
 switchport mode access
 switchport access vlan 100
 no shutdown
end
write memory
```

SW1 — Verification output

```
SW1# show vlan brief
VLAN Name                      Status    Ports
100  OSPF_TRANSIT                 active    Gi0/1, Gi0/2, Gi0/3

SW1# show interfaces status | include Gi0/
Gi0/1  connected  100  a-full a-1000
Gi0/2  connected  100  a-full a-1000
Gi0/3  connected  100  a-full a-1000
```

What to verify

- All three router-facing switchports are active in the same VLAN 100 broadcast domain.
- The broadcast segment is Layer 2 only; OSPF runs on the three router interfaces.

R1 — Complete lab configuration

```
enable
configure terminal
hostname R1
!
interface Loopback0
 ip address 1.1.1.1 255.255.255.255
 ip ospf 10 area 0
interface Loopback10
 ip address 172.16.1.1 255.255.255.255
 ip ospf 10 area 0
!
interface GigabitEthernet0/0/0
 description OSPF_BROADCAST_TO_SW1
 ip address 10.0.123.1 255.255.255.0
 ip ospf 10 area 0
 ip ospf priority 200
 ip ospf cost 20
 no shutdown
!
interface GigabitEthernet0/0/1
 description OSPF_P2P_TO_R2
 ip address 10.0.12.1 255.255.255.252
 ip ospf 10 area 0
 ip ospf network point-to-point
 ip ospf cost 10
 no shutdown
!
router ospf 10
 router-id 1.1.1.1
 passive-interface Loopback0
 passive-interface Loopback10
end
write memory
```

R1 — Verification output

```
R1# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2         100   FULL/BDR        00:00:35    10.0.123.2      Gi0/0/0
3.3.3.3           0   FULL/DROTHER    00:00:36    10.0.123.3      Gi0/0/0
2.2.2.2           0   FULL/-          00:00:37    10.0.12.2       Gi0/0/1

R1# show ip ospf interface GigabitEthernet0/0/0 | include Internet|Process|State|Designated|Backup
  Internet Address 10.0.123.1/24, Interface ID 2, Area 0
  Process ID 10, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 20
  State DR, Priority 200
  Designated Router (ID) 1.1.1.1, Interface address 10.0.123.1
  Backup Designated router (ID) 2.2.2.2, Interface address 10.0.123.2

R1# show ip ospf interface GigabitEthernet0/0/1 | include Network Type|State|Cost
  Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 10
  State POINT_TO_POINT

R1# show ip route ospf
O     2.2.2.2/32 [110/11] via 10.0.12.2, GigabitEthernet0/0/1
O     172.16.2.1/32 [110/11] via 10.0.12.2, GigabitEthernet0/0/1
O     3.3.3.3/32 [110/21] via 10.0.123.3, GigabitEthernet0/0/0
O     172.16.3.1/32 [110/21] via 10.0.123.3, GigabitEthernet0/0/0
```

What to verify

- R1 has two separate adjacencies to R2: one on the broadcast segment and one on the point-to-point link.

- R1 is DR on VLAN 100 because it has the highest OSPF priority, 200.
- R2 is BDR and R3 is DROTHER. Priority 0 prevents R3 from becoming DR or BDR.
- The R1–R2 /30 interface reports Network Type POINT_TO_POINT and has no DR/BDR role.
- OSPF routes display administrative distance 110. The metric to R2 loopbacks is 11 through the cost-10 P2P link; the metric to R3 loopbacks is 21 through the cost-20 broadcast interface plus the loopback cost.
- This is the routing-protocol metric decision referenced in Lab 3.2: when the prefix and route source are the same protocol, the lower OSPF cost is preferred.

R2 — Complete lab configuration

```
enable
configure terminal
hostname R2
!
interface Loopback0
 ip address 2.2.2.2 255.255.255.255
 ip ospf 10 area 0
interface Loopback20
 ip address 172.16.2.1 255.255.255.255
 ip ospf 10 area 0
!
interface GigabitEthernet0/0/0
 ip address 10.0.123.2 255.255.255.0
 ip ospf 10 area 0
 ip ospf priority 100
 ip ospf cost 20
 no shutdown
!
interface GigabitEthernet0/0/1
 ip address 10.0.12.2 255.255.255.252
 ip ospf 10 area 0
 ip ospf network point-to-point
 ip ospf cost 10
 no shutdown
!
router ospf 10
 router-id 2.2.2.2
 passive-interface Loopback0
 passive-interface Loopback20
end
write memory
```

R2 — Verification output

```
R2# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.1         200   FULL/DR         00:00:35    10.0.123.1      Gi0/0/0
3.3.3.3           0   FULL/DROTHER    00:00:36    10.0.123.3      Gi0/0/0
1.1.1.1           0   FULL/-          00:00:37    10.0.12.1       Gi0/0/1

R2# show ip ospf interface GigabitEthernet0/0/0 | include State|Designated|Backup
  State BDR, Priority 100
  Designated Router (ID) 1.1.1.1, Interface address 10.0.123.1
  Backup Designated router (ID) 2.2.2.2, Interface address 10.0.123.2
```

What to verify

- R2 is BDR on the shared Ethernet segment.
- R2 forms FULL adjacency with R1 and R3 on the broadcast segment and a second FULL/- adjacency with R1 on the P2P link.

- The explicit router ID is 2.2.2.2.

R3 — Complete lab configuration

```
enable
configure terminal
hostname R3
!
interface Loopback0
 ip address 3.3.3.3 255.255.255.255
 ip ospf 10 area 0
interface Loopback30
 ip address 172.16.3.1 255.255.255.255
 ip ospf 10 area 0
!
interface GigabitEthernet0/0/0
 ip address 10.0.123.3 255.255.255.0
 ip ospf 10 area 0
 ip ospf priority 0
 ip ospf cost 20
 no shutdown
!
router ospf 10
 router-id 3.3.3.3
 passive-interface Loopback0
 passive-interface Loopback30
end
write memory
```

R3 — Verification output

```
R3# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.1         200   FULL/DR         00:00:35    10.0.123.1      Gi0/0/0
2.2.2.2         100   FULL/BDR        00:00:36    10.0.123.2      Gi0/0/0

R3# show ip ospf interface GigabitEthernet0/0/0 | include State|Priority
  State DROTHER, Priority 0
```

What to verify

- R3 is DROTHER because its priority is 0 and it cannot participate in DR/BDR election.
- A DROTHER forms FULL adjacency with the DR and BDR. It does not need a FULL adjacency with another DROTHER.
- The explicit router ID is 3.3.3.3.

OSPF database and reachability verification

```
R1# show ip ospf database router | include Link ID|1.1.1.1|2.2.2.2|3.3.3.3
Link ID         ADV Router      Age         Seq#       Checksum Link count
1.1.1.1         1.1.1.1         534         0x8000000A 0xA1B2   5
2.2.2.2         2.2.2.2         527         0x80000009 0xB2C3   5
3.3.3.3         3.3.3.3         515         0x80000007 0xC3D4   3

R1# ping 172.16.3.1 source Loopback10
!!!!!
```

What to verify

- Router LSAs exist for all three router IDs in area 0.
- R1 can reach R3's advertised loopback using an OSPF-learned route.

- The OSPF control plane (neighbors/LSDB) and data plane (routing/forwarding) are both operational.

Controlled failure — verify OSPF metric-based reconvergence

```
R1# configure terminal
R1(config)# interface GigabitEthernet0/0/1
R1(config-if)# shutdown
R1(config-if)# end

R1# show ip route 172.16.2.1
Routing entry for 172.16.2.1/32
  Known via "ospf 10", distance 110, metric 21
  * 10.0.123.2, from 2.2.2.2, GigabitEthernet0/0/0

R1# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2         100   FULL/BDR        00:00:34    10.0.123.2      Gi0/0/0
3.3.3.3           0   FULL/DROTHER    00:00:35    10.0.123.3      Gi0/0/0
```

What to verify

- The P2P adjacency to R2 disappears after Gi0/0/1 is shut.
- The same R2 loopback prefix remains reachable through the broadcast segment.
- The OSPF metric changes from 11 to 21 because the remaining path has a higher interface cost.
- Restore Gi0/0/1 and verify the lower-cost metric 11 path returns.

Lab 3.5 — First-Hop Redundancy with HSRP, VRRP and GLBP

CCNA objective: Describe the purpose, functions and concepts of first-hop redundancy protocols. This hands-on lab uses three common FHRP implementations so the virtual-gateway behavior can be observed directly.

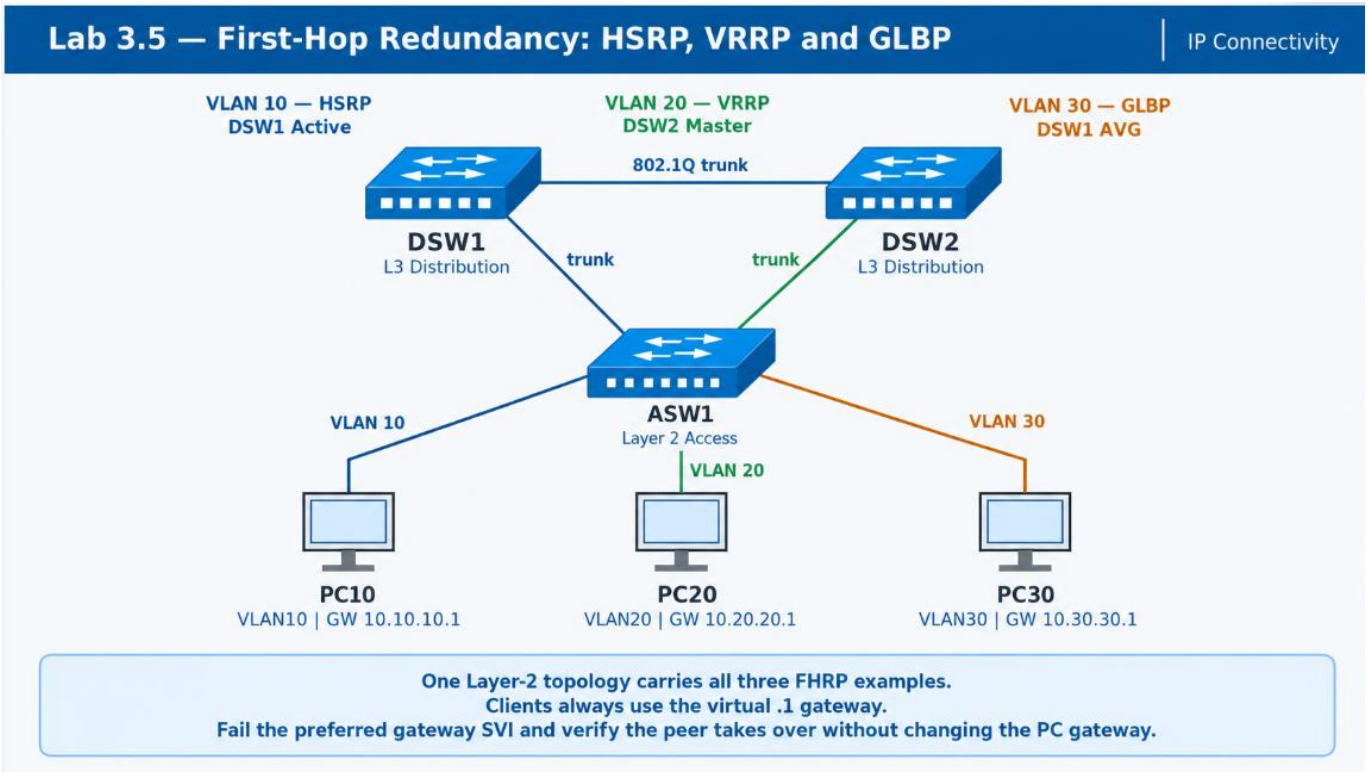

Figure 3.5: Dual-Distribution First-Hop Redundancy reference topology

Lab objectives

- Build one redundant Layer-2 access topology and reuse it for three FHRP VLANs without repeating the base configuration.
- Configure HSRP in VLAN 10 with DSW1 as Active and DSW2 as Standby.
- Configure VRRP in VLAN 20 with DSW2 as Master and DSW1 as Backup.
- Configure GLBP in VLAN 30 with DSW1 as the Active Virtual Gateway (AVG).
- Use the virtual IP as the client default gateway and verify failover after the preferred gateway SVI is shut.

Addressing / interface plan

Device	Interface	Address / Mask	Purpose
DSW1	Vlan10	10.10.10.2/24	HSRP Active; VIP 10.10.10.1
DSW2	Vlan10	10.10.10.3/24	HSRP Standby
DSW1	Vlan20	10.20.20.2/24	VRRP Backup
DSW2	Vlan20	10.20.20.3/24	VRRP Master; VIP 10.20.20.1
DSW1	Vlan30	10.30.30.2/24	GLBP AVG; VIP 10.30.30.1
DSW2	Vlan30	10.30.30.3/24	GLBP Standby AVG
PC10	NIC	10.10.10.10/24	GW 10.10.10.1
PC20	NIC	10.20.20.10/24	GW 10.20.20.1
PC30	NIC	10.30.30.10/24	GW 10.30.30.1

Execution sequence

Configure each device once in the order shown below. After the configuration is complete, use the verification output and the What to verify bullets to prove the expected state. Configuration is not repeated later; subsequent commands are only verification or controlled fault tests.

Lab execution — device-by-device configuration and verification

DSW1 — Complete lab configuration

```
enable
configure terminal
hostname DSW1
ip routing
!
vlan 10
 name HSRP_USERS
vlan 20
 name VRRP_USERS
vlan 30
 name GLBP_USERS
!
interface GigabitEthernet1/0/1
 description TRUNK_TO_DSW2
 switchport mode trunk
 switchport trunk allowed vlan 10,20,30
 no shutdown
interface GigabitEthernet1/0/2
 description TRUNK_TO_ASW1
 switchport mode trunk
 switchport trunk allowed vlan 10,20,30
 no shutdown
!
interface Vlan10
 ip address 10.10.10.2 255.255.255.0
 standby 10 ip 10.10.10.1
 standby 10 priority 110
 standby 10 preempt
 no shutdown
!
interface Vlan20
 ip address 10.20.20.2 255.255.255.0
 vrrp 20 ip 10.20.20.1
 vrrp 20 priority 100
 vrrp 20 preempt
 no shutdown
!
interface Vlan30
 ip address 10.30.30.2 255.255.255.0
 glbp 30 ip 10.30.30.1
 glbp 30 priority 120
 glbp 30 preempt
 no shutdown
end
write memory
```

DSW1 — Verification output

```
DSW1# show standby brief
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Vl10        10   110 P Active  local           10.10.10.3      10.10.10.1

DSW1# show vrrp Vlan20 20
Vlan20 - Group 20
  State is Backup
  Virtual IP address is 10.20.20.1
  Master Router is 10.20.20.3, priority is 120
  Preemption enabled

DSW1# show glbp Vlan30 30
Vlan30 - Group 30
  State is Active
  Virtual IP address is 10.30.30.1
  Active virtual gateway is local
  Standby virtual gateway is 10.30.30.3
  Load balancing: round-robin
```

What to verify

- HSRP VLAN 10 state is Active because priority 110 is higher than DSW2's default 100.
- VRRP VLAN 20 state is Backup because DSW2 has the higher VRRP priority.
- GLBP VLAN 30 state is Active/AVG on DSW1 because priority 120 is higher than DSW2.
- Each FHRP uses the .1 virtual IP as the stable client default gateway.

DSW2 — Complete lab configuration

```
enable
configure terminal
hostname DSW2
ip routing
!
vlan 10
 name HSRP_USERS
vlan 20
 name VRRP_USERS
vlan 30
 name GLBP_USERS
!
interface GigabitEthernet1/0/1
 description TRUNK_TO_DSW1
 switchport mode trunk
 switchport trunk allowed vlan 10,20,30
 no shutdown
interface GigabitEthernet1/0/2
 description TRUNK_TO_ASW1
 switchport mode trunk
 switchport trunk allowed vlan 10,20,30
 no shutdown
!
interface Vlan10
 ip address 10.10.10.3 255.255.255.0
 standby 10 ip 10.10.10.1
 standby 10 priority 100
 standby 10 preempt
 no shutdown
!
interface Vlan20
 ip address 10.20.20.3 255.255.255.0
 vrrp 20 ip 10.20.20.1
 vrrp 20 priority 120
 vrrp 20 preempt
 no shutdown
!
interface Vlan30
 ip address 10.30.30.3 255.255.255.0
 glbp 30 ip 10.30.30.1
 glbp 30 priority 100
 glbp 30 preempt
 no shutdown
end
write memory
```

DSW2 — Verification output

```
DSW2# show standby brief
Interface   Grp  Pri P State    Active          Standby    Virtual IP
Vl10        10   100 P Standby  10.10.10.2      local      10.10.10.1

DSW2# show vrrp Vlan20 20
Vlan20 - Group 20
  State is Master
  Virtual IP address is 10.20.20.1
  Master Router is 10.20.20.3 (local), priority is 120
  Preemption enabled

DSW2# show glbp Vlan30 30
Vlan30 - Group 30
  State is Standby
  Virtual IP address is 10.30.30.1
  Active virtual gateway is 10.30.30.2
  Standby virtual gateway is local
```

What to verify

- DSW2 is HSRP Standby for VLAN 10 and is ready to take the VIP if DSW1 fails.
- DSW2 is VRRP Master for VLAN 20 because it has priority 120.
- DSW2 is the GLBP standby virtual gateway for VLAN 30.

ASW1 — Complete lab configuration

```
enable
configure terminal
hostname ASW1
!
vlan 10
 name HSRP_USERS
vlan 20
 name VRRP_USERS
vlan 30
 name GLBP_USERS
!
interface GigabitEthernet0/1
 description TRUNK_TO_DSW1
 switchport mode trunk
 switchport trunk allowed vlan 10,20,30
 no shutdown
interface GigabitEthernet0/2
 description TRUNK_TO_DSW2
 switchport mode trunk
 switchport trunk allowed vlan 10,20,30
 no shutdown
!
interface FastEthernet0/10
 description PC10_HSRP
 switchport mode access
 switchport access vlan 10
 spanning-tree portfast
 no shutdown
interface FastEthernet0/20
 description PC20_VRRP
 switchport mode access
 switchport access vlan 20
 spanning-tree portfast
 no shutdown
interface FastEthernet0/30
 description PC30_GLBP
 switchport mode access
 switchport access vlan 30
 spanning-tree portfast
 no shutdown
end
write memory
```

ASW1 — Verification output

```
ASW1# show interfaces trunk
Port      Mode  Encapsulation Status    Native vlan
Gi0/1     on    802.1q        trunking  1
Gi0/2     on    802.1q        trunking  1

Port      Vlans allowed on trunk
Gi0/1     10,20,30
Gi0/2     10,20,30

ASW1# show vlan brief | include 10|20|30
10   HSRP_USERS   active    Fa0/10
20   VRRP_USERS   active    Fa0/20
30   GLBP_USERS   active    Fa0/30
```

What to verify

- Both distribution uplinks carry VLANs 10, 20 and 30.

- Each test PC is assigned to the VLAN that matches its FHRP demonstration.
- STP may place one redundant link in a non-forwarding state per VLAN; the FHRP peers must still share the same Layer-2 VLAN through the remaining forwarding topology.

End-device addressing

- PC10: 10.10.10.10/24, default gateway 10.10.10.1.
- PC20: 10.20.20.10/24, default gateway 10.20.20.1.
- PC30: 10.30.30.10/24, default gateway 10.30.30.1.

Client virtual-gateway verification

```
PC10> ping 10.10.10.1
Reply from 10.10.10.1: bytes=32 time<1ms

PC20> ping 10.20.20.1
Reply from 10.20.20.1: bytes=32 time<1ms

PC30> ping 10.30.30.1
Reply from 10.30.30.1: bytes=32 time<1ms
```

What to verify

- Each PC reaches the virtual IP, not the physical SVI address, as its default gateway.
- The virtual IP remains constant when the active/master gateway changes.

Controlled failover 1 — HSRP VLAN 10

```
DSW1# configure terminal
DSW1(config)# interface Vlan10
DSW1(config-if)# shutdown
DSW1(config-if)# end

DSW2# show standby brief
Interface   Grp  Pri P State   Active  Standby   Virtual IP
Vl10        10   100 P Active  local   unknown   10.10.10.1
```

What to verify

- DSW2 moves from Standby to Active and owns virtual IP 10.10.10.1.
- PC10 continues using the same default gateway address.
- Restore DSW1 Vlan10 with no shutdown; because preemption is enabled and DSW1 has higher priority, DSW1 becomes Active again.

Controlled failover 2 — VRRP VLAN 20

```
DSW2# configure terminal
DSW2(config)# interface Vlan20
DSW2(config-if)# shutdown
DSW2(config-if)# end

DSW1# show vrrp Vlan20 20
Vlan20 - Group 20
  State is Master
  Virtual IP address is 10.20.20.1
  Master Router is 10.20.20.2 (local), priority is 100
```

What to verify

- DSW1 changes from Backup to Master and owns the same virtual IP 10.20.20.1.
- PC20 does not change its gateway configuration.
- Restore DSW2 Vlan20; with higher priority and preemption, DSW2 returns to Master.

Controlled failover 3 — GLBP VLAN 30

```
DSW1# configure terminal
DSW1(config)# interface Vlan30
DSW1(config-if)# shutdown
DSW1(config-if)# end

DSW2# show glbp Vlan30 30
Vlan30 - Group 30
  State is Active
  Virtual IP address is 10.30.30.1
  Active virtual gateway is local
```

What to verify

- DSW2 takes the GLBP Active Virtual Gateway role for virtual IP 10.30.30.1.
- GLBP is designed to provide gateway redundancy and can also distribute hosts across virtual forwarders when multiple forwarders are available.
- Restore DSW1 Vlan30 and verify the preferred AVG state returns according to priority and preemption.

Module 3 — Final Verification Checklist

Use this checklist after completing all labs. Each item should be supported by show-command output, not only by a successful ping.

- Routing table: correctly identify C, L, S, S* and O route codes and the destination prefix/mask.
- Route entry: identify the selected next hop, outgoing interface, administrative distance and metric.
- Default routing: locate the gateway of last resort / default route and explain when it is used.
- Forwarding decision: prove longest prefix match wins before administrative distance is compared.
- Administrative distance: demonstrate a higher-AD floating static route becoming active after the primary next hop fails.
- IPv4 static routing: verify default, network, host and floating static routes.
- IPv6 static routing: verify ::/0, network, /128 host and floating static routes.
- OSPF adjacency: neighbors reach FULL state on the expected interfaces.
- OSPF broadcast: identify DR, BDR and DROTHER states and explain the effect of interface priority.
- OSPF point-to-point: verify Network Type POINT_TO_POINT and absence of DR/BDR election on that link.
- OSPF metric: prove a lower-cost route is selected and a higher-cost route is used after failure.
- FHRP: verify the virtual IP, active/master/standby state and successful gateway failover for HSRP, VRRP and GLBP.

Core verification commands

```
show ip interface brief
show ip route
show ip route <destination>
show ip route static
show ip cef <destination>
show ipv6 interface brief
show ipv6 route
show ipv6 route static
show ip ospf neighbor
show ip ospf interface <interface>
show ip ospf database
show ip route ospf
show standby brief
show vrrp
show glbp
ping <destination>
traceroute <destination>
traceroute ipv6 <destination>
```

Instructor note

When a simulator or image does not implement VRRP or GLBP, complete the HSRP portion on that platform and run the VRRP/GLBP sections on CML, EVE-NG, GNS3 or supported physical IOS/IOS XE equipment. Do not replace the missing protocol with repeated HSRP configuration; the purpose of this guide is to preserve one clear sequence for each protocol example.