OSPF for CCNP ENCOR

The coverage follows ENCOR v1.2: OSPFv2/v3, neighbors, point-to-point and broadcast networks, passive interfaces, multiple normal areas, summarization, filtering, and area types. Cisco exam blueprint.

OSPF builds a topology map and calculates paths through it. Open Shortest Path First is a link-state Interior Gateway Protocol: a routing protocol used within an organization's routing domain. Routers describe their links using **Link-State Advertisements (LSAs)**. They store those advertisements in a **Link-State Database (LSDB)** and use the **Shortest Path First (SPF)** algorithm to calculate paths.

Keep these three views separate:

View	What it contains	Cisco command
Neighbor table	Other OSPF routers discovered on connected links	`show ip ospf neighbor`
LSDB	Advertisements describing topology and reachable prefixes	`show ip ospf database`
Routing table	Selected routes used for forwarding	`show ip route ospf`

Routers in the same area converge on the same area topology information. **Their routing tables differ because each calculates paths from its own position.** OSPFv2 carries IPv4 routing information and runs directly over IP protocol **89**. Its multicast groups are **224.0.0.5** for all OSPF routers and **224.0.0.6** for DR/BDR routers. OSPFv2 specification.

Areas limit how much topology each router must understand. Within an area, routers know the detailed topology. Between areas, they learn destination prefixes and costs through an **Area Border Router (ABR)**.

We will use this topology throughout the lab. Loopbacks simulate destination networks, so you need only three routers.

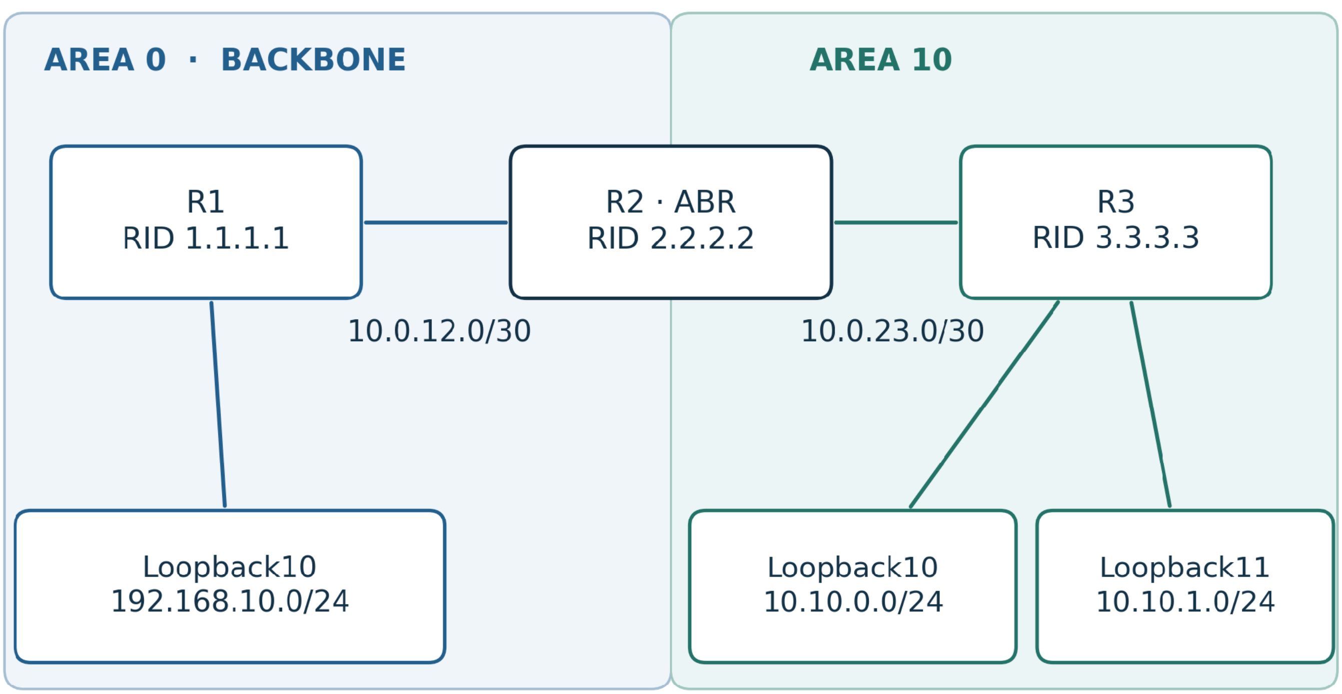

R2 connects Area 10 to **Area 0, the backbone**. A normal multiarea design connects every nonbackbone area to Area 0 and keeps the backbone contiguous. The two ends of a link belong to the same area. R2 uses one OSPF process for both areas; it does not redistribute between them.

An **internal router** has all its OSPF interfaces in one area. A **backbone router** has an interface in Area 0. An **Autonomous System Boundary Router (ASBR)** introduces external routing information into OSPF, for example by redistributing routes from another protocol. A router can have more than one role. Cisco: interarea routing.

The LSA type tells you what information you are looking at. These are the main OSPFv2 types; DR and NSSA are explained below:

Type	Name	What it describes	Created by / scope
1	Router	A router's links and costs within an area	Every router; stays in that area
2	Network	Routers attached to a shared transit network	DR; stays in that area
3	Summary	A prefix reachable through another area	ABR; advertised into another area
4	ASBR Summary	How to reach an ASBR in another area	ABR; advertised where needed
5	AS External	An external prefix	ASBR; floods through normal areas
7	NSSA External	An external prefix introduced inside an NSSA	NSSA ASBR; stays in that NSSA

Eligible Type 7 advertisements are translated into Type 5 by an NSSA ABR. **A Type 3 "Summary LSA" can describe a single subnet.** Its name does not mean aggregation has been configured. In our lab, R2 advertises R3's two /24s into Area 0 using Type 3 LSAs. Cisco: interarea LSAs, external LSAs.

Neighbors discover each other before they exchange the database. The five OSPF packet types have distinct jobs:

Packet	Job
Hello	Discover neighbors and check that they remain reachable
Database Description — DBD	Describe database contents using LSA headers
Link-State Request — LSR	Request particular LSAs
Link-State Update — LSU	Carry one or more complete LSAs
Link-State Acknowledgment — LSAck	Acknowledge received LSAs

An **LSA is an advertisement inside the protocol; an LSU is a packet that carries advertisements**. After initial synchronization, routers flood changes and periodically refresh their own LSAs; they do not repeatedly send their entire routing tables. OSPFv2 specification.

The neighbor states describe progress:

State	Meaning
Down	No valid Hello has been received from the neighbor
Init	A Hello arrived, but it does not list our router ID
2-Way	The neighbor's Hello lists our router ID: communication works both ways
ExStart	Negotiate DBD exchange roles and the initial sequence number
Exchange	Exchange DBDs and compare database contents

State	Meaning
Loading	Request and receive missing or newer LSAs
Full	Synchronization is complete for this adjacency

`Attempt` is an additional state associated with manually configured NBMA neighbors. For ENCOR's Ethernet and point-to-point examples, focus on the states above. [Cisco: neighbor states].

Full adjacency requires compatible settings on the link. Check these when a neighbor will not form:

Setting	Requirement
Interfaces	Up, OSPF enabled, and not passive on the transit link
Area	Same area ID and compatible normal/stub/NSSA settings
Hello / Dead intervals	Match; normally **10 / 40 seconds** on broadcast and point-to-point links
Authentication	Compatible method and matching credentials, when enabled
IPv4 addressing	Correct addressing; same subnet/mask on a broadcast segment
Router IDs	Unique
Network type	Configure consistently at both ends
MTU	Compatible for database exchange; a mismatch commonly stalls ExStart/Exchange
Packet delivery	OSPF protocol 89 must pass between peers

The **process ID is locally significant**: `router ospf 10` on one router can neighbor `router ospf 50` on another. Costs and DR priorities also need not match. Network type itself is not a Hello field; mismatches can sometimes reach Full while producing broken routing. [Cisco: neighbor troubleshooting].

Broadcast networks elect a DR and BDR; point-to-point networks do not. Ethernet defaults to OSPF's broadcast network type, even when only two routers are connected.

Behavior	Broadcast	Point-to-point
Typical use	Several routers sharing a VLAN	A link with exactly two OSPF routers
DR / BDR election	Yes	No
Full adjacencies	Each router with the DR and BDR	With the other router
DROTHER-to-DROTHER relationship	Normally remains at 2-Way	Not applicable: no DR roles
Type 2 LSA for the link	Yes, when the DR has a Full neighbor	No

The **Designated Router (DR)** coordinates flooding on that segment. The **Backup Designated Router (BDR)** maintains adjacencies and takes over if the DR fails. Other routers are called **DROTHERs**.

This diagram shows OSPF relationships among four routers on the **same VLAN**, rather than physical cabling:

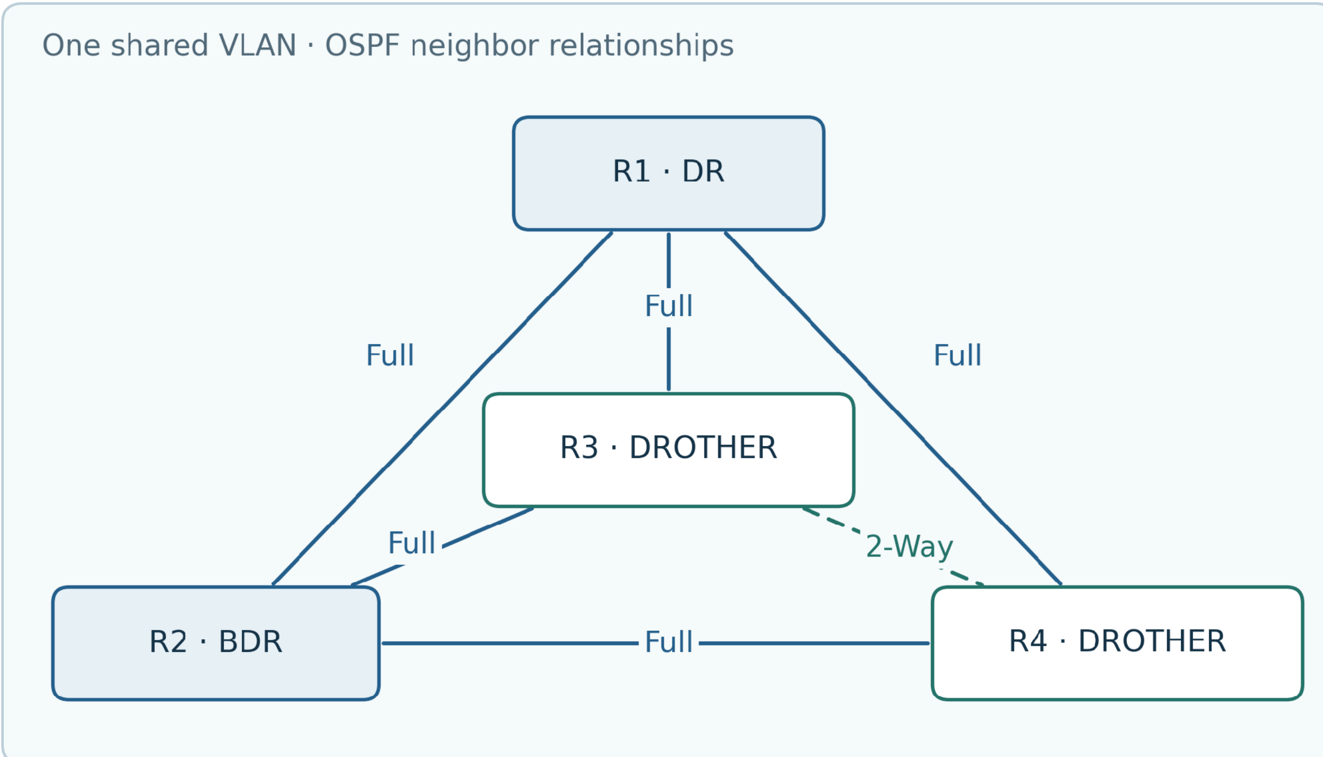

Election preference is highest interface priority, then highest router ID. Default priority is **1**; **0** makes a router ineligible for DR and BDR. Existing DR/BDR roles are **not preempted** by a newcomer with a higher priority. Configure priorities before the initial election:

```
interface GigabitEthernet0/0
 ip ospf priority 200
```

DR is a **per-segment role**. It does not make that router the default gateway or require user traffic to traverse it. `FULL/DR` in neighbor output describes the **neighbor's** role. `2WAY/DROTHER` between two DROTHERs is healthy. [Cisco: interpreting neighbor output](#).

OSPF selects paths by cost. Cisco normally derives interface cost from:

$$Interface\ cost\ =\ \frac{reference\ bandwidth}{interface\ bandwidth}$$

Cost is an integer with a minimum of 1. The default reference bandwidth is **100 Mbps**, so 100 Mbps, 1 Gbps, and 10 Gbps interfaces all normally receive cost 1. Set a consistent reference bandwidth across routers to distinguish faster links:

```
router ospf 10
 auto-cost reference-bandwidth 10000
```

`10000` is in **Mbps**. With that reference, 100 Mbps costs 100, 1 Gbps costs 10, and 10 Gbps costs 1. An explicit `ip ospf cost 25` overrides the calculated cost on that interface. OSPF does not measure current congestion to choose the path. [Cisco: OSPF configuration guide](#).

For an internal route, add the outgoing interface costs toward the destination, including the advertising router's cost to the destination network. In this separate example, every link is in one area and arrows show costs toward the LAN:

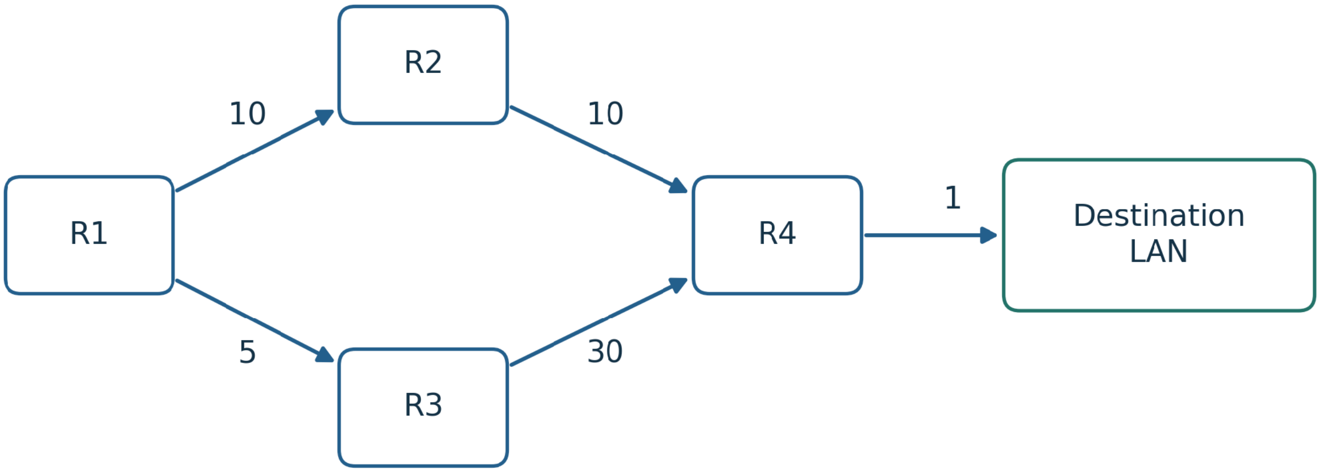

Path from R1	Calculation	Result
Through R2	10 + 10 + 1	**21 — selected**
Through R3	5 + 30 + 1	36

If both paths have the same eligible cost, OSPF can install both for **equal-cost multipath (ECMP)**, subject to the configured path limit.

Within one OSPF process, route type matters before comparing costs for the same prefix. An intra-area route (`O`) is preferred over an interarea route (`O IA`), which is preferred over an external route. For example, `O` with cost 100 beats `O IA` with cost 20 for the same prefix.

External route	How its metric is evaluated
`O E1`	External metric plus internal cost to the external exit
`O E2`	External metric first; internal cost breaks equal-external-metric ties
`O N1` / `O N2`	NSSA counterparts of E1 / E2

E1 is preferred over E2 for the same external prefix. Administrative distance is a separate comparison between routing sources; Cisco's default for OSPF is **110**. Once routes are installed, packet forwarding uses **longest-prefix match**. Do not confuse that with choosing between candidates for the same prefix. Cisco: OSPF route preference.

ENCOR also expects you to distinguish OSPF from EIGRP. The comparison is about how they calculate and maintain routes:

Concept	OSPF	EIGRP
Model	Link state	Advanced distance vector
Path calculation	SPF over the area topology	DUAL using neighbor-reported route information
Default metric basis	Interface cost, usually derived from bandwidth	Minimum path bandwidth and accumulated delay
Load balancing	Equal-cost paths	Equal-cost, plus eligible unequal-cost paths using `variance`
Hierarchy	Backbone and other areas	No OSPF-style area hierarchy

Concept	OSPF	EIGRP
Response to a failure	Flood the change and recalculate affected routes	Use a feasible successor if available; otherwise perform a diffusing computation

An EIGRP feasible successor is a prequalified loop-free alternative. OSPF does not use that terminology or selection mechanism. Cisco: EIGRP operation.

Now configure the three-router topology. These examples use Cisco IOS/IOS XE interface names; substitute yours if needed. Enter each block from privileged EXEC mode. The displayed results later are expected results, not captures from a running emulator.

Router	Interface	IPv4 address	IPv6 address, used later	Area
R1	G0/0 → R2	10.0.12.1/30	2001:db8:12::1/64	0
R1	Loopback10	192.168.10.1/24	—	0
R2	G0/0 → R1	10.0.12.2/30	2001:db8:12::2/64	0
R2	G0/1 → R3	10.0.23.1/30	2001:db8:23::1/64	10
R3	G0/0 → R2	10.0.23.2/30	2001:db8:23::2/64	10
R3	Loopback10	10.10.0.1/24	2001:db8:100::1/64	10
R3	Loopback11	10.10.1.1/24	2001:db8:100:1::1/64	10

We assign transit interfaces cost **10** and destination loopbacks cost **1**, so the route calculations are predictable across lab platforms.

On **R1**:

```
configure terminal
hostname R1
router ospf 10
 router-id 1.1.1.1
 auto-cost reference-bandwidth 10000
 passive-interface default
 no passive-interface GigabitEthernet0/0
exit
interface GigabitEthernet0/0
 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
exit
interface Loopback10
 ip address 192.168.10.1 255.255.255.0
 ip ospf 10 area 0
 ip ospf network point-to-point
 ip ospf cost 1
end
```

On **R2**:

```
configure terminal
hostname R2
router ospf 10
 router-id 2.2.2.2
 auto-cost reference-bandwidth 10000
 passive-interface default
 no passive-interface GigabitEthernet0/0
 no passive-interface GigabitEthernet0/1
exit
```

```
interface GigabitEthernet0/0
 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
exit
interface GigabitEthernet0/1
 ip address 10.0.23.1 255.255.255.252
 ip ospf 10 area 10
 ip ospf network point-to-point
 ip ospf cost 10
 no shutdown
end
```

On **R3**:

```
configure terminal
hostname R3
router ospf 10
 router-id 3.3.3.3
 auto-cost reference-bandwidth 10000
 passive-interface default
 no passive-interface GigabitEthernet0/0
exit
interface GigabitEthernet0/0
 ip address 10.0.23.2 255.255.255.252
 ip ospf 10 area 10
 ip ospf network point-to-point
 ip ospf cost 10
 no shutdown
exit
interface Loopback10
 ip address 10.10.0.1 255.255.255.0
 ip ospf 10 area 10
 ip ospf network point-to-point
 ip ospf cost 1
exit
interface Loopback11
 ip address 10.10.1.1 255.255.255.0
 ip ospf 10 area 10
 ip ospf network point-to-point
 ip ospf cost 1
end
```

Each important command has a specific purpose:

Command	Effect
`router-id 1.1.1.1`	Sets a stable, unique OSPF identifier; it need not be an assigned interface address
`ip ospf 10 area 0`	Enables OSPF process 10 on this interface and assigns it to Area 0
`passive-interface default`	Prevents Hellos and neighbor formation by default; enabled connected prefixes are still advertised
`no passive-interface G0/0`	Allows OSPF neighbors on this transit interface
`ip ospf network point-to-point`	Removes DR/BDR election on the two-router transit links
The same command on the loopbacks	Makes OSPFv2 advertise the configured /24 instead of its default loopback /32

Without an explicit router ID, Cisco normally selects the highest IPv4 loopback address, otherwise the highest IPv4 address on an active interface, when the process chooses its ID. Set it before bringing up OSPF; changing an active router ID may require restarting the process.

You will also encounter this alternative interface-selection syntax:

```
configure terminal
router ospf 10
 network 10.0.12.0 0.0.0.3 area 0
end
```

It enables OSPF on interfaces whose addresses match the statement. The wildcard uses **0 = must match; 1 = ignore**. It does not create a summary route or directly announce an arbitrary network. Our lab uses interface commands, so this alternative is unnecessary. Cisco: enabling OSPF on interfaces.

Verify the neighbors, then the routes, then forwarding. R2 should have two Full neighbors:

```
R2# show ip ospf neighbor

Neighbor ID     Pri   State       Dead Time    Address        Interface
1.1.1.1           0   FULL/-      00:00:3x     10.0.12.1     GigabitEthernet0/0
3.3.3.3           0   FULL/-      00:00:3x     10.0.23.2     GigabitEthernet0/1
```

`FULL/-` means Full adjacency without a DR/BDR role. R1 should learn:

```
O IA 10.10.0.0/24 [110/21] via 10.0.12.2, GigabitEthernet0/0
O IA 10.10.1.0/24 [110/21] via 10.0.12.2, GigabitEthernet0/0
```

`110` is administrative distance. **21 = R1's outgoing cost 10 + R2's outgoing cost 10 + R3's destination cost 1.** They are `O IA` because R1 reaches them through another area. R2 sees these same prefixes as `O` because it participates directly in Area 10.

From R1, verify end-to-end routing with an explicit source:

```
ping 10.10.0.1 source 192.168.10.1
ping 10.10.1.1 source 192.168.10.1
```

These pings also require R3 to have a return route to `192.168.10.1`.

Summarization replaces several advertised prefixes with one covering prefix. Our two Area 10 networks are contiguous and aligned:

Component prefixes	Aggregate	Mask
10.10.0.0/24 and 10.10.1.0/24	**10.10.0.0/23**	255.255.254.0

On **R2, the ABR**, configure the range under the area that contains those networks:

```
configure terminal
router ospf 10
 area 10 range 10.10.0.0 255.255.254.0
end
```

R1 now learns one aggregate in place of those two specific routes:

```
O IA 10.10.0.0/23 [110/21] via 10.0.12.2, GigabitEthernet0/0
```

R2 retains the specific routes in Area 10. The default aggregate metric is the **highest component cost at the ABR**, here 11; R1 adds 10 to reach R2. The summary remains advertised while at least one qualifying component exists. Cisco also normally installs a discard route to **Null0** on the ABR: traffic matching the aggregate without a more specific route is dropped, preventing a loop through a less specific route. Cisco: area-range and discard-route commands.

OSPF has two main aggregation locations: **ABR `area range` for interarea routes**, and **ASBR `summary-address` for external routes**. It does not automatically summarize an area merely because that area exists.

Filtering advertisements and filtering local routes have different effects. To stop R2 advertising `10.10.1.0/24` from Area 10 into other areas, remove the previous summary and apply an ABR prefix filter:

```
configure terminal
ip prefix-list AREA10-EXPORT seq 5 deny 10.10.1.0/24
ip prefix-list AREA10-EXPORT seq 10 permit 0.0.0.0/0 le 32
router ospf 10
 no area 10 range 10.10.0.0 255.255.254.0
 area 10 filter-list prefix AREA10-EXPORT out
end
```

The first entry denies exactly that /24. The second permits everything else; without it, the prefix list's implicit deny would block the other prefixes too.

ABR command	Direction relative to Area 10
`area 10 filter-list prefix NAME in`	Other areas → Area 10
`area 10 filter-list prefix NAME out`	Area 10 → other areas

R1 keeps `10.10.0.0/24` and loses `10.10.1.0/24`. R2 still knows both inside Area 10. This controls **Type 3 advertisements**; it is not a traffic ACL. A covering summary or default route could still provide reachability, which is why we removed the summary for this exercise. Cisco: ABR filtering command.

By comparison, an inbound `distribute-list` filters installation into the **local routing table**. The LSAs remain in the database and continue flooding. Thus, "present in `show ip ospf database`" does not guarantee "present in `show ip route`." Cisco: local OSPF route filtering.

Area types control which information an area receives. All these area types retain their own internal topology information:

Area type	Specific interarea routes	Type 5 external routes	Local external redistribution
Normal	Yes	Yes	Yes
Stub	Yes	No	No
Totally stubby	Replaced by an ABR default route	No	No
NSSA	Yes	No	Yes, using Type 7
Totally NSSA	Replaced by an ABR default route	No	Yes, using Type 7

Stub areas also receive an ABR-generated default route. Do not assume a plain NSSA automatically receives one; default origination has separate rules. Area 0 remains a normal area. Cisco's "totally" variants use `no-summary` on the ABR. Cisco: area-type behavior.

For example, to make our Area 10 stub, configure `area 10 stub` under `router ospf 10` on **both R2 and R3**. For totally stubby, use `area 10 stub no-summary` on **R2** and retain `area 10 stub` on **R3**. Practice these as separate variations after restoring the normal-area lab.

A normal area needs explicit default-route origination. If an edge router already has a usable default route, this advertises it into OSPF:

```
configure terminal
router ospf 10
 default-information originate
end
```

The ordinary command requires a default route in that router's routing table. Adding `always` advertises a default even without one. An advertised default does not itself create Internet connectivity. The default external metric type is E2. [Cisco: default-route origination](#).

OSPFv3 uses the same core ideas with different addressing and commands. Areas, SPF, cost, neighbor states, DR/BDR, and passive interfaces still apply.

Detail	OSPFv2	OSPFv3
Routing information	IPv4	IPv6; address-family extensions also support IPv4
Router ID	32-bit, written like an IPv4 address	Still 32-bit; set explicitly even in an IPv6-only network
Normal neighbor communication	IPv4	IPv6 link-local addresses
Multicast groups	224.0.0.5 / 224.0.0.6	FF02::5 / FF02::6
Protocol number	89	89, carried over IPv6
Interface activation	`ip ospf 10 area 0`	`ospfv3 10 ipv6 area 0`
Prefix information	Included in several topology LSAs	Separated from the Router and Network LSAs

OSPFv3 adds **Link-LSAs** and **Intra-Area-Prefix-LSAs**. Interarea prefixes use **Inter-Area-Prefix-LSAs**, the counterpart of OSPFv2 Type 3. OSPFv3 is not automatically encrypted; authentication requires configuration and supported mechanisms. [OSPFv3 specification](#).

Add this IPv6 configuration to the **existing three-router lab**. It uses a separate OSPFv3 process alongside OSPFv2; the shared number 10 does not join the two protocols.

On **R1**:

```
configure terminal
ipv6 unicast-routing
router ospfv3 10
 router-id 1.1.1.1
 passive-interface default
 no passive-interface GigabitEthernet0/0
 address-family ipv6 unicast
 exit-address-family
exit
interface GigabitEthernet0/0
 ipv6 address 2001:db8:12::1/64
 ospfv3 10 ipv6 area 0
 ospfv3 network point-to-point
 ospfv3 cost 10
end
```

On **R2**:

```
configure terminal
ipv6 unicast-routing
router ospfv3 10
 router-id 2.2.2.2
 passive-interface default
 no passive-interface GigabitEthernet0/0
 no passive-interface GigabitEthernet0/1
 address-family ipv6 unicast
 exit-address-family
exit
```

```
interface GigabitEthernet0/0
 ipv6 address 2001:db8:12::2/64
 ospfv3 10 ipv6 area 0
 ospfv3 network point-to-point
 ospfv3 cost 10
exit
interface GigabitEthernet0/1
 ipv6 address 2001:db8:23::1/64
 ospfv3 10 ipv6 area 10
 ospfv3 network point-to-point
 ospfv3 cost 10
end
```

On **R3**:

```
configure terminal
ipv6 unicast-routing
router ospfv3 10
 router-id 3.3.3.3
 passive-interface default
 no passive-interface GigabitEthernet0/0
 address-family ipv6 unicast
 exit-address-family
exit
interface GigabitEthernet0/0
 ipv6 address 2001:db8:23::2/64
 ospfv3 10 ipv6 area 10
 ospfv3 network point-to-point
 ospfv3 cost 10
exit
interface Loopback10
 ipv6 address 2001:db8:100::1/64
 ospfv3 10 ipv6 area 10
 ospfv3 network point-to-point
 ospfv3 cost 1
exit
interface Loopback11
 ipv6 address 2001:db8:100:1::1/64
 ospfv3 10 ipv6 area 10
 ospfv3 network point-to-point
 ospfv3 cost 1
end
```

`ipv6 unicast-routing` enables IPv6 forwarding. IPv6 addresses also enable link-local addressing on these interfaces. As with our IPv4 simulated LANs, the loopbacks use point-to-point network type so the configured prefixes can be advertised instead of loopback host routes.

Verify with:

```
show ospfv3 ipv6 neighbor
show ospfv3 interface brief
show ospfv3 database
show ipv6 route ospf
```

R2 should have two Full neighbors. R1 should learn R3's two /64s as **`OI` interarea routes**, with an IPv6 link-local next hop. From R1:

```
ping ipv6 2001:db8:100::1 source 2001:db8:12::1
```

Older IOS materials use `ipv6 router ospf 10` and `ipv6 ospf 10 area 0`. Recognize that legacy form; these examples use the modern address-family form. [Cisco Press: OSPFv3 configuration and verification].

IPv6 summarization uses the same ABR concept. On **R2**:

```
configure terminal
router ospfv3 10
 address-family ipv6 unicast
  area 10 range 2001:db8:100::/63
 exit-address-family
end
```

The /63 combines `2001:db8:100::/64` and `2001:db8:100:1::/64`. R1 should learn the aggregate instead of those two specific prefixes. [Cisco: OSPFv3 address families and summarization](#).

For a **separate IPv6 filtering exercise**, remove that range on R2 and filter the second /64:

```
configure terminal
ipv6 prefix-list AREA10-V6-EXPORT seq 5 deny 2001:db8:100:1::/64
ipv6 prefix-list AREA10-V6-EXPORT seq 10 permit ::/0 le 128
router ospfv3 10
 address-family ipv6 unicast
  no area 10 range 2001:db8:100::/63
  area 10 filter-list prefix AREA10-V6-EXPORT out
 exit-address-family
end
```

R1 should retain the first /64 and lose the second. IPv4 prefix lists and IPv6 prefix lists are separate, as are the policies in the two routing processes. [Cisco: OSPFv3 ABR filtering](#).

Troubleshoot the stage that is failing. Start with interface state, then neighbors, then the database, then the routing table and forwarding:

Symptom	Most useful checks
No neighbor	Interface up/up, correct VLAN/subnet, OSPF enabled, passive state, area, timers, authentication, protocol 89 delivery
Stuck in Init	One-way Hello delivery; our router ID is absent from the peer's Hello
Stuck in 2-Way	Expected between DROTHERs; otherwise inspect network type and DR eligibility
Stuck in ExStart / Exchange	MTU mismatch, duplicate router ID, packet loss, or DBD negotiation problems
Stuck in Loading	Missing LSAs, retransmissions, or packet loss
Full, but route absent	Prefix enabled at origin, ABR/Area 0 connectivity, area filter, local distribute-list, or a preferred competing route
Route present, but ping fails	Return route, ACL, destination state, source address, and forwarding path

An MTU mismatch is checked during database exchange, so Hellos can succeed first. Fix the underlying mismatch instead of treating `ip ospf mtu-ignore` as a universal solution. [Cisco: ExStart/Exchange troubleshooting](#).

Command	What to inspect
`show ip ospf interface brief`	Interfaces, areas, costs, and neighbor counts
`show ip ospf interface G0/0`	Network type, timers, DR/BDR, and passive status
`show ip ospf neighbor`	Neighbor IDs and adjacency states
`show ip ospf database`	LSAs available to OSPF
`show ip ospf database summary`	OSPFv2 Type 3 advertisements
`show ip route 10.10.1.1`	The route actually selected for that destination
`show ip protocols`	Routing-process settings and filtering

To test your understanding, make a change and predict the result before checking it. Use the base lab and restore each change before proceeding:

1. Make R2's G0/1 passive. Its adjacency with R3 should disappear; enabling OSPF on the interface alone does not override passive mode.
2. Change R1's G0/0 cost from 10 to 50. Its unsummarized routes to R3's /24s should rise from **21 to 61**.
3. Remove point-to-point network type from both ends of R1–R2. Ethernet should use broadcast behavior and elect a DR/BDR.
4. Configure the /23 summary on R2. R1 should replace the two /24 advertisements with one /23 route.
5. Remove the summary and filter `10.10.1.0/24`. R1 should lose that route while R2 retains it.
6. Repeat summarization and filtering with the IPv6 prefixes, then verify a sourced IPv6 ping.