

C O N F I G U R A T I O N

Technically, we're using OSPFv2 for IPv4, but that only matters because IPv6 uses OSPFv3.

Wildcard—a bitmask controlling address matching much like a netmask. You've seen these before, in access control lists.

- A zero must match; a one means we don't care. For example the address + wildcard pair 123.456.789.123 0.0.0.255 matches anything in the subnet 123.456.789.0 /24
- When you add a network address to its wildcard, you happen to get its broadcast address, the last address in the range
- Any legal netmask can be turned into a wildcard by subtracting it from 255.255.255.255. For example, the wildcard 0.0.0.3 is equivalent to the netmask 255.255.255.252
- In the world beyond OSPF, not all wildcards can be expressed as a netmask. Netmasks must be a contiguous block of ones followed by a contiguous block of zeroes, yielding a single range of addresses with no gaps. In general, wildcards can be any mix of 1s and 0s, yielding matches like "every other IP address"
- In OSPF, wildcards *do* need to express a contiguous block of addresses, so they do need to be a sequence of zeroes followed by a sequence of ones, with no interleaving

Configuring OSPF on each router can be this simple:

```
1 R1(config)# router ospf 10
2 R1(config-router)# network 10.0.0.0 0.0.0.7 area 0
```

Line 1

- Starts an OSPF process on the router
- The process number will be 10. The process number only has meaning on this one router and doesn't need to match the process number on any other router. In fact, you're not even going to care what the number is because in CCNA you'll only ever need one process running on the same router and you'll just call it "the" OSPF process.

Line 2 specifies which interfaces will participate in OSPF

- All interfaces with an IP address between 10.0.0.0 and 10.0.0.7 will participate—*see wildcard logic, above*. This corresponds to the subnet 10.0.0.0 /29
- You can have more lines like this to specify other interface address ranges
- Participating interfaces will have their subnets advertised to other routers through OSPF
- Participating interfaces will be used by OSPF to find neighboring OSPF routers
- These interfaces will be part of area 0

Notice that interfaces are involved in OSPF by matching their addresses as part of the OSPF configuration rather than in the configuration of the interfaces themselves. This is normal for this generation of routing protocols, but differs from the IPv6 version of OSPF that was developed later.

The network statement wildcard doesn't need to line up with the netmask of a network that you want to advertise; it just has to include the address of the interface whose network you want to involve. For simplicity, you can use a broad wildcard to cover several interfaces or even use an exact interface address with a 0.0.0.0 wildcard to clearly specify one interface per network statement.

EXPLICIT ROUTER ID

You can explicitly set the RID to make diagnostic command output easier to read and influence Designated Router (DR) elections on Ethernet subnets. The RID is just a 32-bit number, formatted as if it were an IP address. After you change the number, you'll have to reset the OSPF process [line 4] to make the change take effect.

```
R1
1 R1(config)# router ospf 10
2 R1(config-router)# router-id 0.0.0.1
3 % OSPF: Reload or use "clear ip ospf process" command, for this to take effect
4 R1(config-router)# do clear ip ospf process
5 Reset ALL OSPF processes? [no]: y
6 R1(config-router)#
7 Jan  2 19:43:03.446: %OSPF-5-ADJCHG: Process 10, Nbr 10.2.0.1 on
8 GigabitEthernet0/0 from FULL to DOWN, Neighbor Down: Interface down or detached
9 Jan  2 19:43:03.450: %OSPF-5-ADJCHG: Process 10, Nbr 10.2.0.1 on
10 GigabitEthernet0/0 from LOADING to FULL, Loading Done
```

As mentioned in chapter 7, manually setting the RID is optional. If you don't, the OSPF process will simply use the highest IP address found on a loopback interface, and if there aren't any loopbacks, it'll choose the highest address on a real interface.

PASSIVE INTERFACES

Enabling OSPF on an interface has two effects:

- The network reachable through that interface will be advertised to other OSPF routers
- Advertisements and hellos will be sent and received through the interface and neighbor relationships will be formed

If we only wanted to advertise the network and not recruit neighbors through a given interface, we can declare that interface to be "passive." This can be useful in a few situations:

- The interface leads to an access switch; no routers are reachable out the interface. We can eliminate a slight amount of CPU load and network traffic
- The interface leads to a router that isn't ours, for example, we want to advertise the network addresses on the link leading to our ISP, but we don't want to send OSPF traffic out from our own network to the world at large

Passive interfaces are declared within the OSPF configuration


```
1 R2(config)# router ospf 2
2 R2(config-router)# passive-interface gi0/1
```

To simplify configuration, you have the option of declaring all interfaces passive and then choosing exceptions.

```
1 R2(config)# router ospf 2
2 R2(config-router)# passive-interface default
3 R2(config-router)# no passive-interface gi0/0
```

To see which interfaces are passive without looking at the running config, you'll need to use the full version of the "show ip ospf interface" command and look for a notation that an interface has "No Hellos." The "brief" version of the command lists all OSPF interfaces, passive or not.

V E R I F I C A T I O N



Show IP OSPF Neighbor

This time, "we" are R2. For each neighbor,

- "Neighbor ID" is the RID of the neighbor
- "State" is the neighbor's state. In the following example, line 4 shows that R3 is just beginning to exchange route information. It'll be fully adjacent in a moment, just like R1
- Lines 3 and 4 also show us that R1 is the Designated Router (DR) and R3 is the Backup DR (BDR) on the LAN connecting all 3 routers, leaving us as a "DROTHER." Clearly, this wasn't caused by an election based on router IDs, and we see in the "Pri" column to the left that R1 and R3's interfaces still have the default priority of 1. Therefore, the election must have been skewed by the order in which the already-configured routers were booted or the interfaces enabled—the first 2 routers up would elect the DR and BDR amongst themselves. A new election would be held only if one of them went down.

Serial lines will show "FULL/ -" (a dash) because they don't have DRs and BDRs

- "Address" belongs to the neighbor and is at the other end of the link we share
- "Interface" is ours, leading to the neighbor.

```

1 R2# show ip ospf neighbor
2 Neighbor ID      Pri   State           Dead Time   Address      Interface
3 0.0.0.1          1    FULL/DR         00:00:37   10.0.0.1    GigabitEthernet0/0
4 0.0.0.3          1    EXSTART/BDR     00:00:39   10.0.0.3    GigabitEthernet0/0
  
```

Show IP OSPF Database

This shows the contents of the Link State DataBase (LSDB). The output should be identical on each router within an area. The "Router Link State" section should show an entry for each router in the area. That's all we care about for CCNA. A third section, "Summary Net Link States," is possible if you have networks in more than one area and is shown in 8b "Multi-Area OSPF."

```

1 R1# show ip ospf database
2
3         OSPF Router with ID (0.0.0.1) (Process ID 10)
4
5         Router Link States (Area 0)
6
7 Link ID        ADV Router      Age             Seq#            Checksum Link count
8 0.0.0.1        0.0.0.1        781            0x80000002     0x00A673  1
9 0.0.0.2        0.0.0.2        373            0x80000004     0x00975F  2
10 0.0.0.3        0.0.0.3        754            0x80000002     0x00A62E  3
11
12         Net Link States (Area 0)
13
14 Link ID        ADV Router      Age             Seq#            Checksum
15 10.0.0.1        0.0.0.1        753            0x80000002     0x0045E4
  
```

Show IP OSPF Interface

This command will allow you to clearly see which routers are the DR and BDR for a LAN. In this case (following a reboot after the previous commands), we see that R3 is the DR, R2 is the BDR, and we, R1, are a DROTHER.

```
1 R1# show ip ospf interface Gi0/0
2 GigabitEthernet0/0 is up, line protocol is up
3 Internet Address 10.0.0.1/29, Area 0, Attached via Network Statement
4 Process ID 10, Router ID 0.0.0.1, Network Type BROADCAST, Cost: 1
5 Topology-MTID      Cost      Disabled      Shutdown      Topology Name
6           0          1          no            no            Base
7 Transmit Delay is 1 sec, State DROTHER, Priority 1
8 Designated Router (ID) 0.0.0.3, Interface address 10.0.0.3
9 Backup Designated router (ID) 0.0.0.2, Interface address 10.0.0.2
10 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:00
    Supports Link-local Signaling (LLS)
    Cisco NSF helper support enabled
    IETF NSF helper support enabled
    Index 1/1, flood queue length 0
    Next 0x0(0)/0x0(0)
    Last flood scan length is 0, maximum is 1
    Last flood scan time is 0 msec, maximum is 0 msec
    Neighbor Count is 2, Adjacent neighbor count is 2
        Adjacent with neighbor 0.0.0.2 (Backup Designated Router)
        Adjacent with neighbor 0.0.0.3 (Designated Router)
    Suppress hello for 0 neighbor(s)
```

Show IP Route

Routes learned through ospf [lines 16-18] have an "O" in the left column and an Administrative Distance (AD) of 110 in the square brackets. The second number in square brackets shows the ospf "cost" or "distance" to that network from where we are.

```
1 R1# show ip route
2 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
3         D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
4         N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
5         E1 - OSPF external type 1, E2 - OSPF external type 2
6         i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
7         ia - IS-IS inter area, * - candidate default, U - per-user static route
8         o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
9         + - replicated route, % - next hop override
10
11 Gateway of last resort is not set
12
13     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
14 C       10.0.0.0/29 is directly connected, GigabitEthernet0/0
15 L       10.0.0.1/32 is directly connected, GigabitEthernet0/0
16 O       10.2.0.1/32 [110/2] via 10.0.0.2, 00:11:45, GigabitEthernet0/0
17 O       10.3.0.1/32 [110/2] via 10.0.0.3, 00:18:11, GigabitEthernet0/0
18 O       10.3.1.1/32 [110/2] via 10.0.0.3, 00:18:11, GigabitEthernet0/0
```

Show IP Protocols

This command will tell us about any routing protocols running on the router, not just OSPF.

- Line 4 shows us the process number we chose for OSPF. It corresponds to the configuration command "router ospf 10." We'll need that if we want to change anything later
- Line 7 tells us our RID
- Line 11 recaps our OSPF network statements
- Lines 14-15 literally tell us which OSPF routers we happened to learn routes from. "Gateway" is just an archaic term for router.

```
1 R1# show ip protocols
2 *** IP Routing is NSF aware ***
3
4 Routing Protocol is "ospf 10"
5   Outgoing update filter list for all interfaces is not set
6   Incoming update filter list for all interfaces is not set
7   Router ID 0.0.0.1
8   Number of areas in this router is 1. 1 normal 0 stub 0 nssa
9   Maximum path: 4
10  Routing for Networks:
11     10.0.0.0 0.0.0.7 area 0
12  Routing Information Sources:
13     Gateway          Distance      Last Update
14     0.0.0.2           110          00:16:57
15     0.0.0.3           110          00:23:23
16  Distance: (default is 110)
```

Inclusion in this list is a little counter-intuitive, slightly random, and definitely beyond the scope of CCNA.

- The router needs to have given us a route we didn't already know. Neighbors who tell us nothing or tell us nothing we didn't already know won't appear
- A router doesn't need to be a neighbor to get credit for providing a route. For example, a link between a neighbor and a more distant router could easily be credited to either.
- Because a link between two OSPF routers will be reported by both, each will have a 50/50 chance of being the first to tell us and be listed

TROUBLESHOOTING NEIGHBOR RELATIONSHIPS

A failure to form a neighbor relationship can be caused by all the normal problems at OSI layers 3, 2, and 1, but a tricky test question could easily prevent two routers from becoming neighbors by declaring an interface to be passive in the OSPF configuration. That interface will look perfectly normal in the brief version of the "show ip ospf interface" command, telling you that routes found out that interface are being advertised out other interfaces, but not warning you that the interface is passive, meaning that nothing is advertised out it.

```
1 R2# show ip ospf interface brief
2 Interface    PID    Area      IP Address/Mask    Cost  State Nbrs F/C
3 Gi0/0        2      0         10.12.0.2/24       1     BDR   1/1
4 Gi0/1        2      0         10.23.0.2/24       1     DR    0/0
```

Show IP OSPF Interface

Taking it as an article of faith that you should never trust the "show running config" command to be available during an exam, you'll need to use the full version of the "show ip ospf interface" command [line 12] to see which interfaces are passive.

On lines 8-9, you can also see which routers are the designated router and backup designated router on this segment.

```
1 R2# show ip ospf interface
2 GigabitEthernet0/1 is up, line protocol is up
3   Internet Address 10.23.0.2/24, Area 0, Attached via Network Statement
4   Process ID 2, Router ID 0.0.0.2, Network Type BROADCAST, Cost: 1
5   Topology-MTID    Cost    Disabled    Shutdown    Topology Name
6   0                1      no         no         Base
7   Transmit Delay is 1 sec, State DR, Priority 1
8   Designated Router (ID) 0.0.0.2, Interface address 10.23.0.2
9   No backup designated router on this network
10  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
11    oob-resync timeout 40
12    No Hellos (Passive interface)
13  Supports Link-local Signaling (LLS)
14  Cisco NSF helper support enabled
15  IETF NSF helper support enabled
16  Index 2/2, flood queue length 0
17  Next 0x0(0)/0x0(0)
18  Last flood scan length is 0, maximum is 0
19  Last flood scan time is 0 msec, maximum is 0 msec
20  Neighbor Count is 0, Adjacent neighbor count is 0
21  Suppress hello for 0 neighbor(s)
```