OSPF – Open Shotest Path First

Routing Configuration

Cisco | Networking Academy®
Mind Wide Open™

OSPF – Open Shotest Path First

Single-Area OSPF

Characteristics of OSPF

Configuring Single-area OSPFv2

Configuring Single-area OSPFv3

Advanced Single-Area OSPF Configurations

Multiarea OSPF Operation

Configuring Multiarea OSPF

Single-Area OSPF

Cisco | Networking Academy®
Mind Wide Open™

Characteristics of OSPF

Cisco | Networking Academy®
Mind Wide Open™

Evolution of OSPF

Interior Gateway Protocols

	Interior Gateway Protocols				Exterior Gateway Protocols
	Distance Vector		Link-State		Path Vector
IPv4	RIPv2	EIGRP	OSPFv2	IS-IS	BGP-4
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	BGP-MP

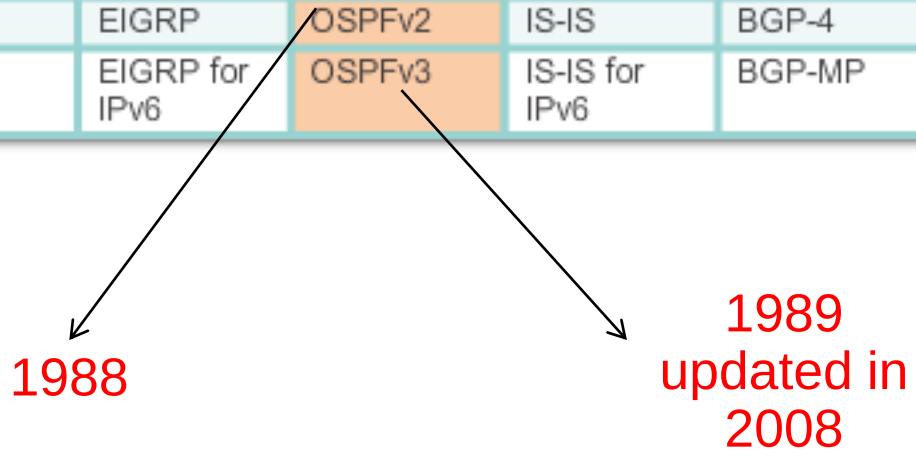

Open Shortest Path First

Features of OSPF

Open Shortest Path First

Components of OSPF

OSPF Data Structures

Database	Table	Description
Adjacency Database	Neighbor Table	• List of all neighbor routers to which a router has established bidirectional communication. • This table is unique for each router. • Can be viewed using the **`show ip ospf neighbor`** command.
Link-state Database (LSDB)	Topology Table	• Lists information about all other routers in the network. • The database shows the network topology. • All routers within an area have identical LSDB. • Can be viewed using the **`show ip ospf database`** command.
Forwarding Database	Routing Table	• List of routes generated when an algorithm is run on the link-state database. • Each router's routing table is unique and contains information on how and where to send packets to other routers. • Can be viewed using the **`show ip route`** command.

Components of OSPF

OSPF Routers Exchange Packets - These packets are used to discover neighboring routers and also to exchange routing information to maintain accurate information about the network.

Open Shortest Path First

Link-State Operation

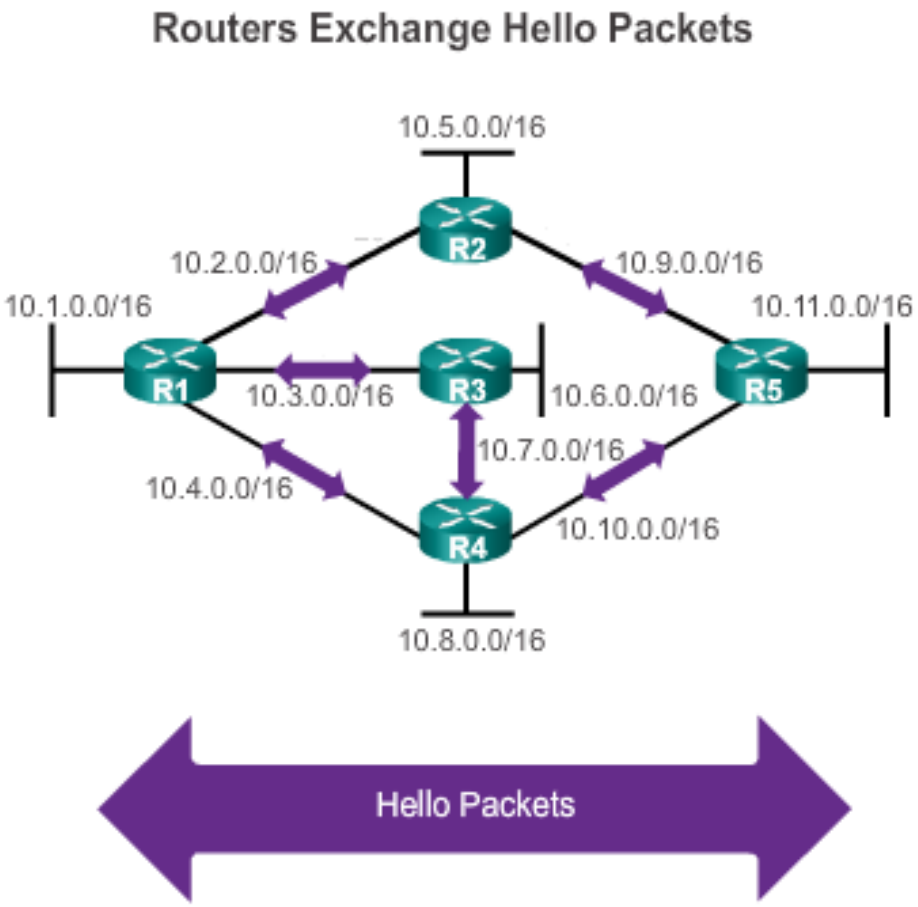

If a neighbor is present, the OSPF-enabled router attempts to establish a neighbor adjacency with that neighbor

Link-State Operation

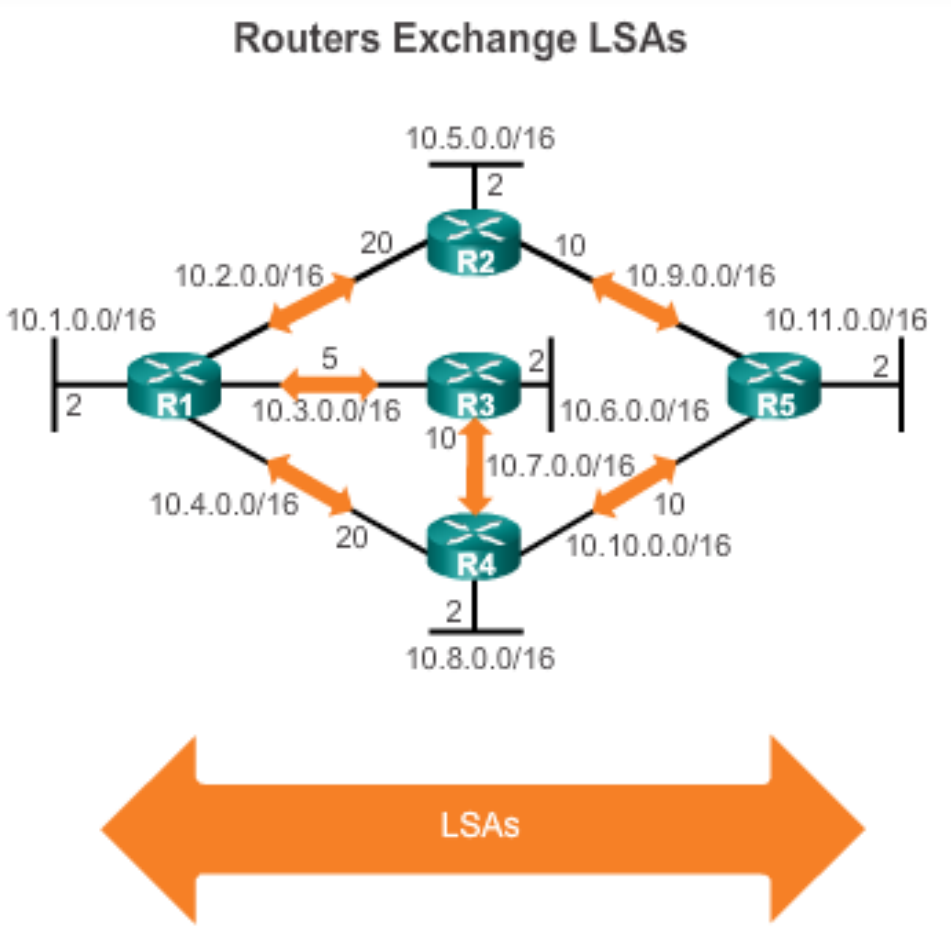

- LSAs contain the state and cost of each directly connected link.
- Routers flood their LSAs to adjacent neighbors.
- Adjacent neighbors receiving the LSA immediately flood the LSA to other directly connected neighbors, until all routers in the area have all LSAs.

Link-State Operation

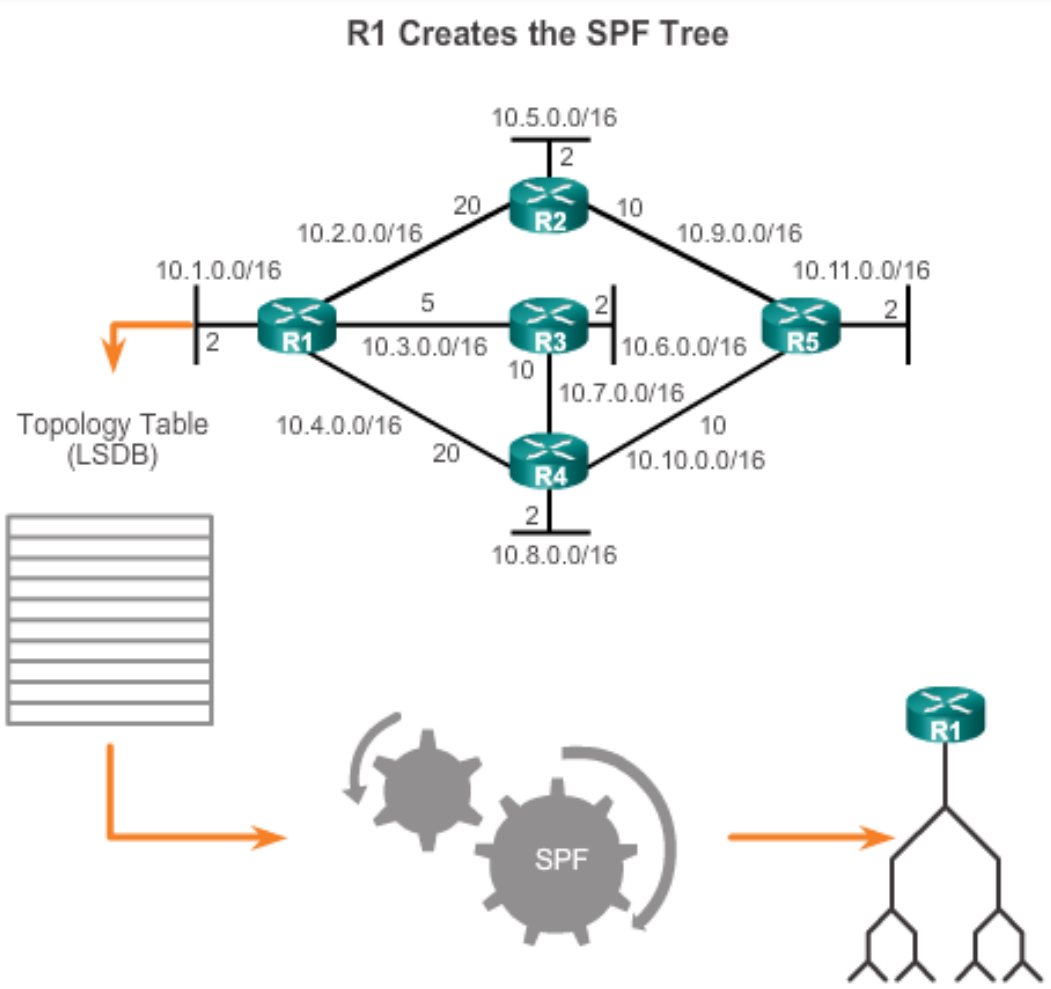

- Build the topology table based on the received LSAs.
- This database eventually holds all the information about the topology of the network.
- Execute the SPF Algorithm.

Link-State Operation

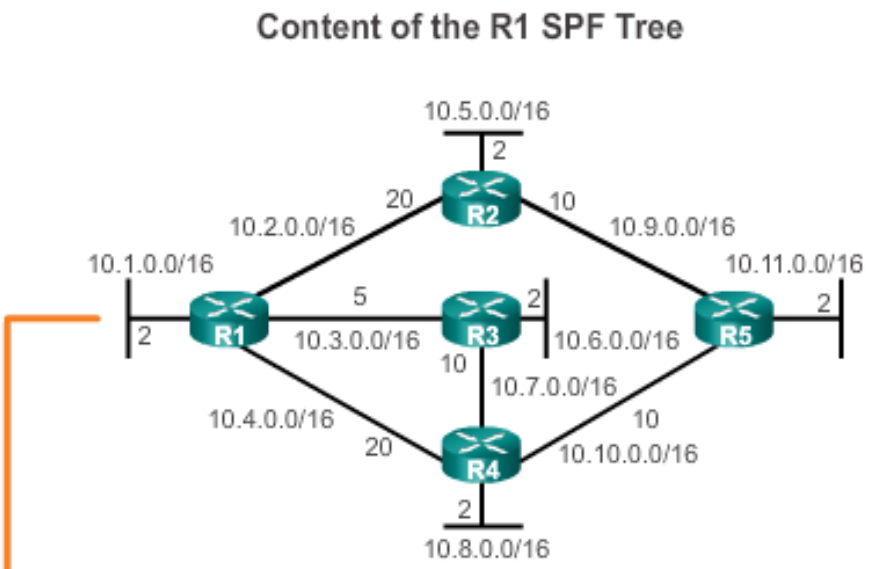

Destination	Shortest Path	Cost
10.5.0.0/16	R1 → R2	22
10.6.0.0/16	R1 → R3	7
10.7.0.0/16	R1 → R3	15
10.8.0.0/16	R1 → R3 → R4	17
10.9.0.0/16	R1 → R2	30
10.10.0.0/16	R1 → R3 → R4	25
10.11.0.0/16	R1 → R3 → R4 → R5	27

From the SPF tree, the best paths are inserted into the routing table.

Open Shortest Path First

Single-area and Multiarea OSPF

Single-Area OSPF

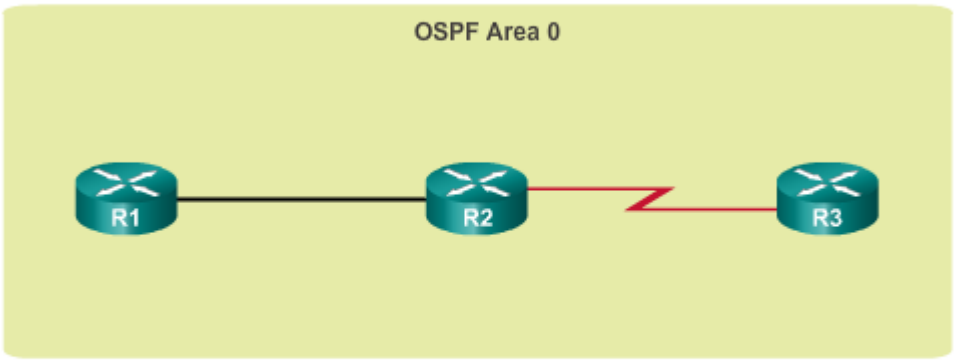

- Area 0 is also called the backbone area.
- Single-area OSPF is useful in smaller networks with few routers.

Multiarea OSPF

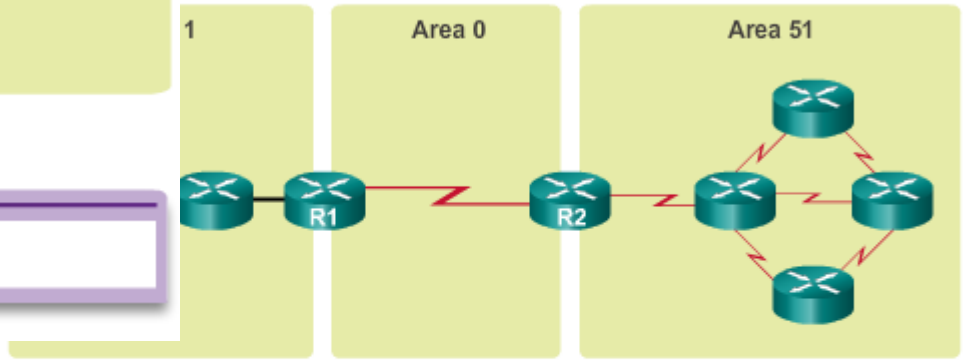

- Implemented using a two-layer area hierarchy as all areas must connect to the backbone area (area 0).
- Interconnecting routers are called Area Border Routers (ABR).
- Useful in larger network deployments to reduce processing and memory overhead.

Single-area and Multiarea OSPF

Link Change Impacts Local Area Only

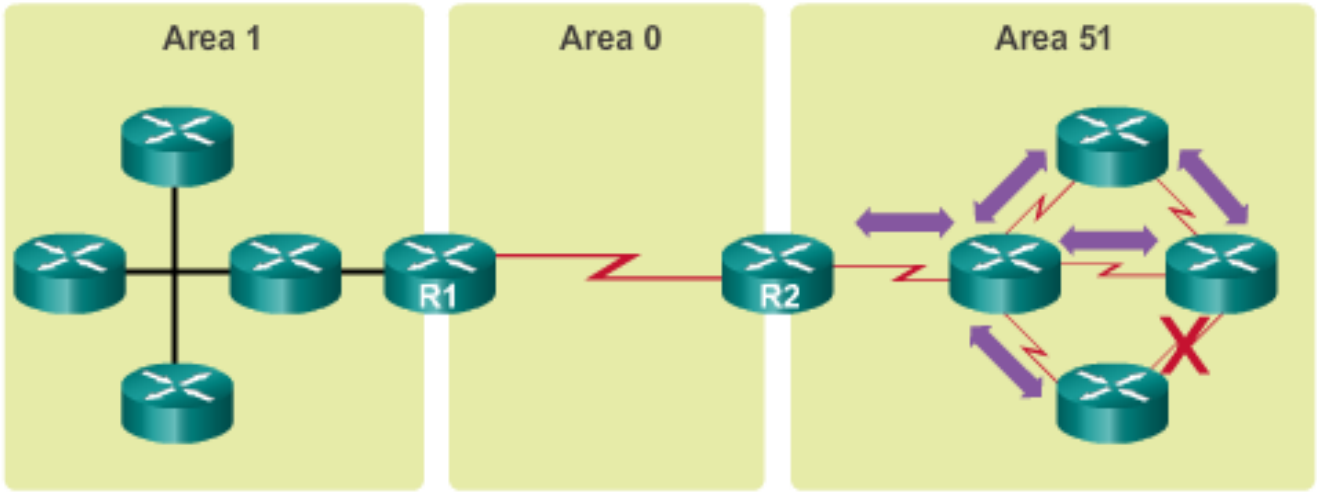

- Link failure affects the local area only (area 51).
- The ABR (R2) isolates the fault to area 51 only.
- Routers in areas 0 and 1 do not need the run the SPF algorithm.

Encapsulating OSPF Messages

OSPF IPv4 Header Fields

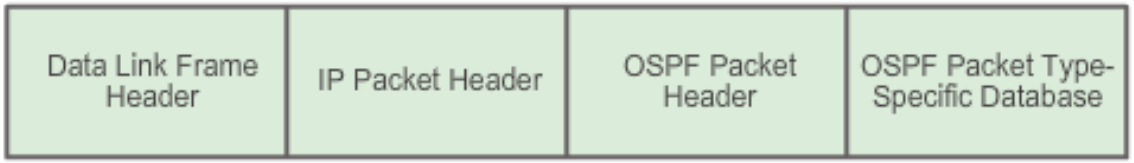

Data Link Frame (Ethernet Fields shown here)
MAC Destination Address = Multicast: 01-00-5E-00-00-05 or 01-00-5E-00-00-06
MAC Source Address = Address of sending interface

IP Packet
IP Source Address = Address of sending interface
IP Destination Address = Multicast: 224.0.0.5 or 224.0.0.6
Protocol field = 89 for OSPF

OSPF Packet Header
Type code for OSPF Packet type
Router ID and Area Id

OSPF Packet types
0x01 Hello
0x02 Database Description (DD)
0X03 Link State Request
0X04 Link State Update
0X05 Link State Acknowledgment

OSPF Messages

Types of OSPF Packets

OSPF Packet Descriptions

Type	Packet Name	Description
1	Hello	Discovers neighbors and builds adjacencies between them
2	Database Description (DBD)	Checks for database synchronization between routers
3	Link-State Request (LSR)	Requests specific link-state records from router to router
4	Link-State Update (LSU)	Sends specifically requested link-state records
5	Link-State Acknowledgment (LSAck)	Acknowledges the other packet types

Hello Packet

OSPF Type 1 packet = Hello packet

- Discover OSPF neighbors and establish neighbor adjacencies
- Advertise parameters on which two routers must agree to become neighbors
- Elect the Designated Router (DR) and Backup Designated Router (BDR) on multiaccess networks like Ethernet and Frame Relay

OSPF Messages

Hello Packet

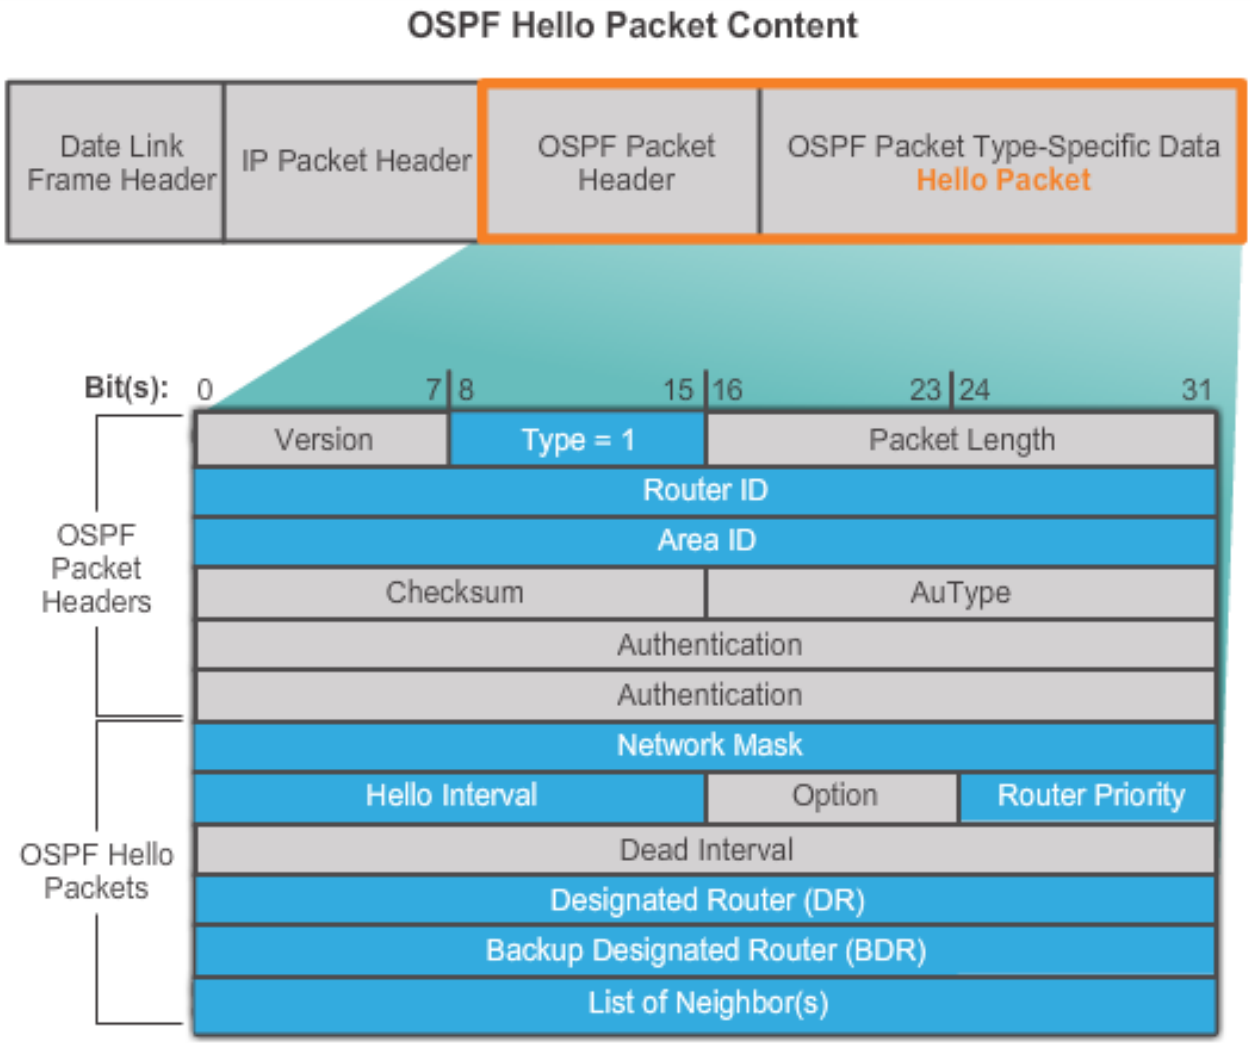

Hello Packet Intervals

OSPF Hello packets are transmitted

- To 224.0.0.5 in IPv4 and FF02::5 in IPv6 (all OSPF routers)
- Every 10 seconds (default on multiaccess and point-to-point networks)
- Every 30 seconds (default on non-broadcast multiaccess [NBMA] networks)
- Dead interval is the period that the router waits to receive a Hello packet before declaring the neighbor down
- Router floods the LSDB with information about down neighbors out all OSPF enabled interfaces
- Cisco's default is 4 times the Hello interval

OSPF Messages

Link-State Updates

LSUs Contain LSAs

Type	Packet Name	Description
1	Hello	Discovers neighbors and builds adjacencies between them
2	DBD	Checks for database synchronization between router
3	LSR	Requests specific link-state records from router to router
4	LSU	Sends specifically requested link-state records
5	LSAck	Acknowledges the other packet types

- An LSU contains one or more LSAs.
- LSAs contain route information for destination networks.

LSA Type	Description
1	Router LSAs
2	Network LSAs
3 or 4	Summary LSAs
5	Autonomous System External LSAs
6	Multicast OSPF LSAs
7	Defined for Not-So-Stubby Areas
8	External Attributes LSA for Border Gateway Protocol (BGP)
9,10,11	Opaque LSAs

OSPF Operational States

When an OSPF router is initially connected to a network, it attempts to:

- Create adjacencies with neighbors
- Exchange routing information
- Calculate the best routes
- Reach convergence
- OSPF progresses through several states while attempting to reach convergence.

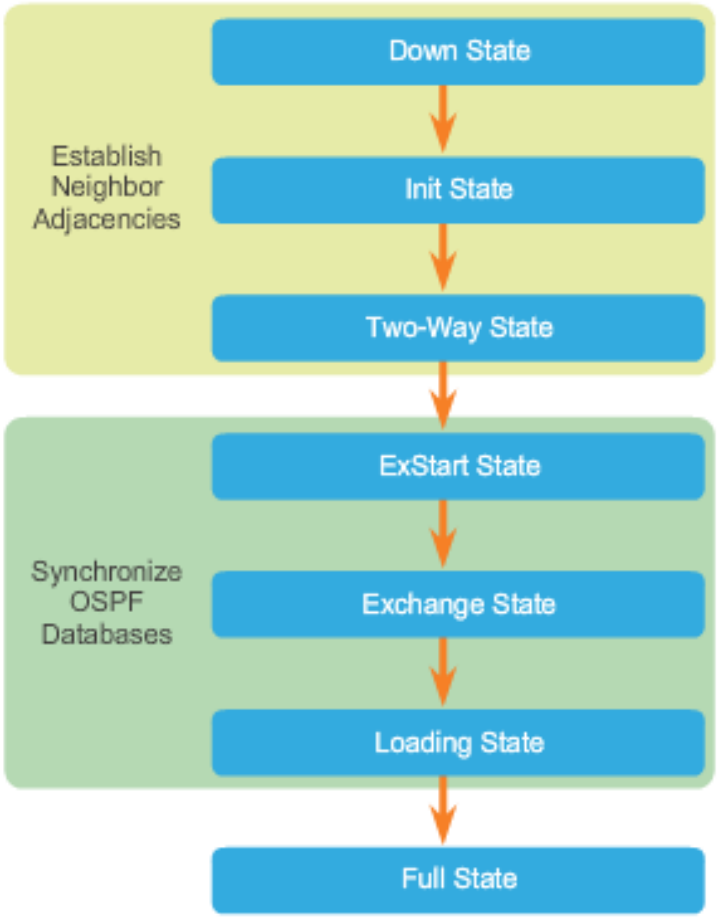

OSPF Operation

Establish Neighbor Adjacencies

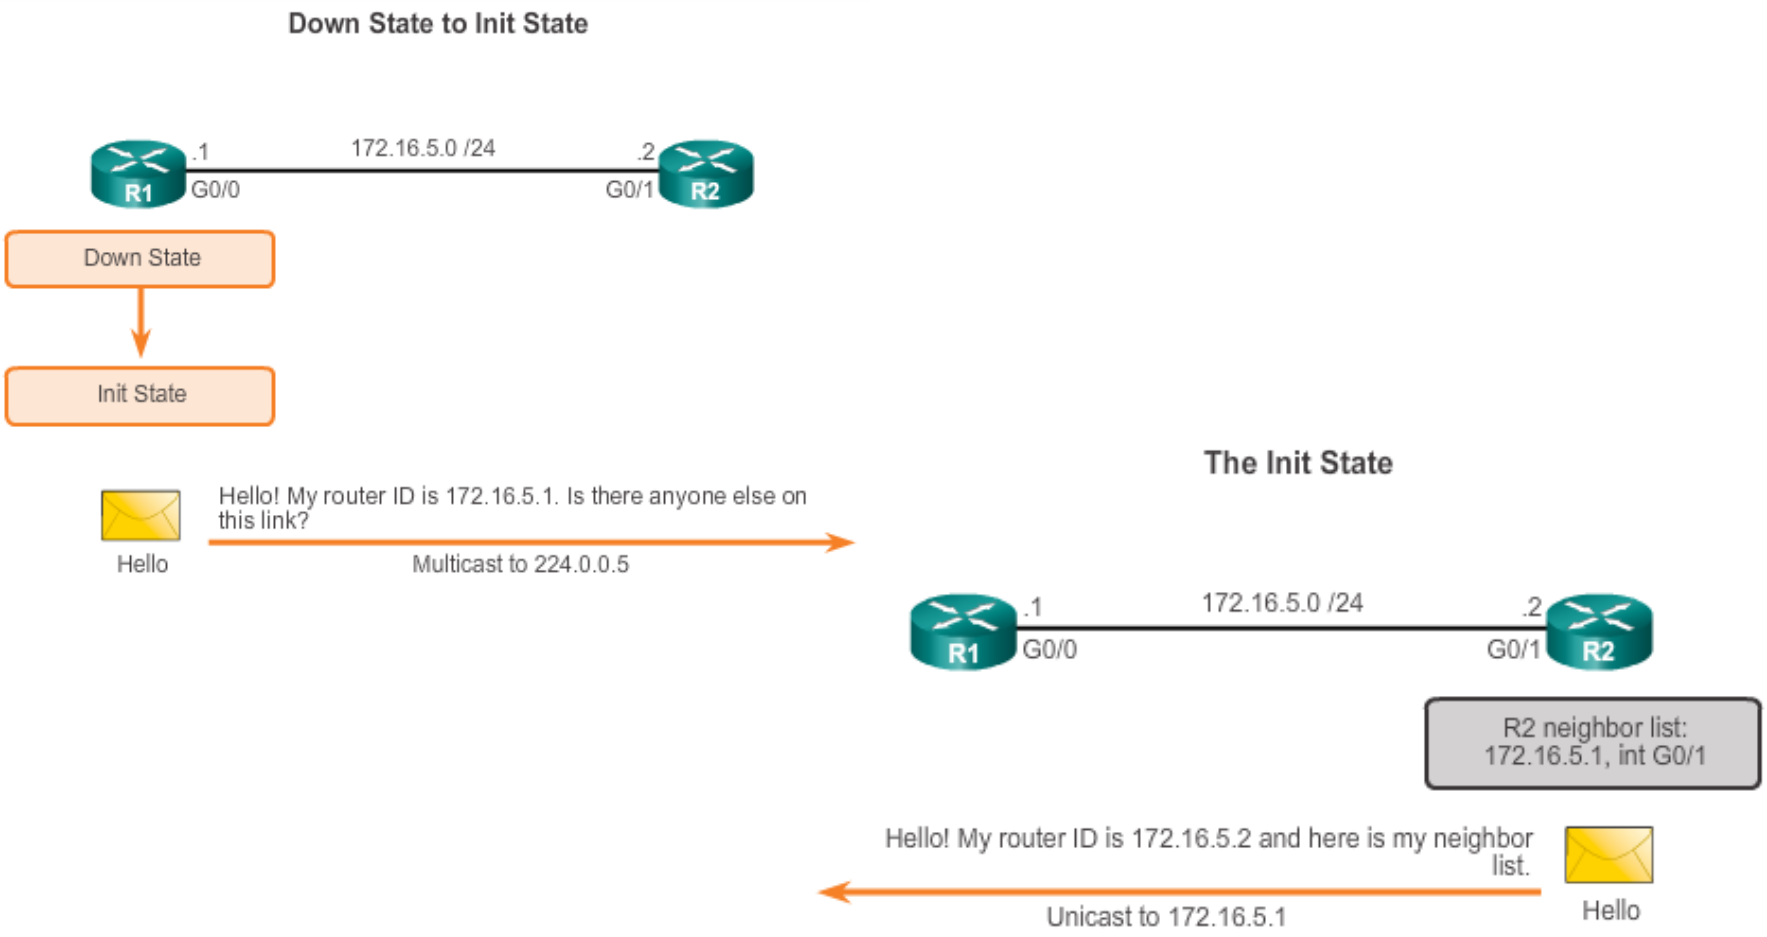

Establish Neighbor Adjacencies

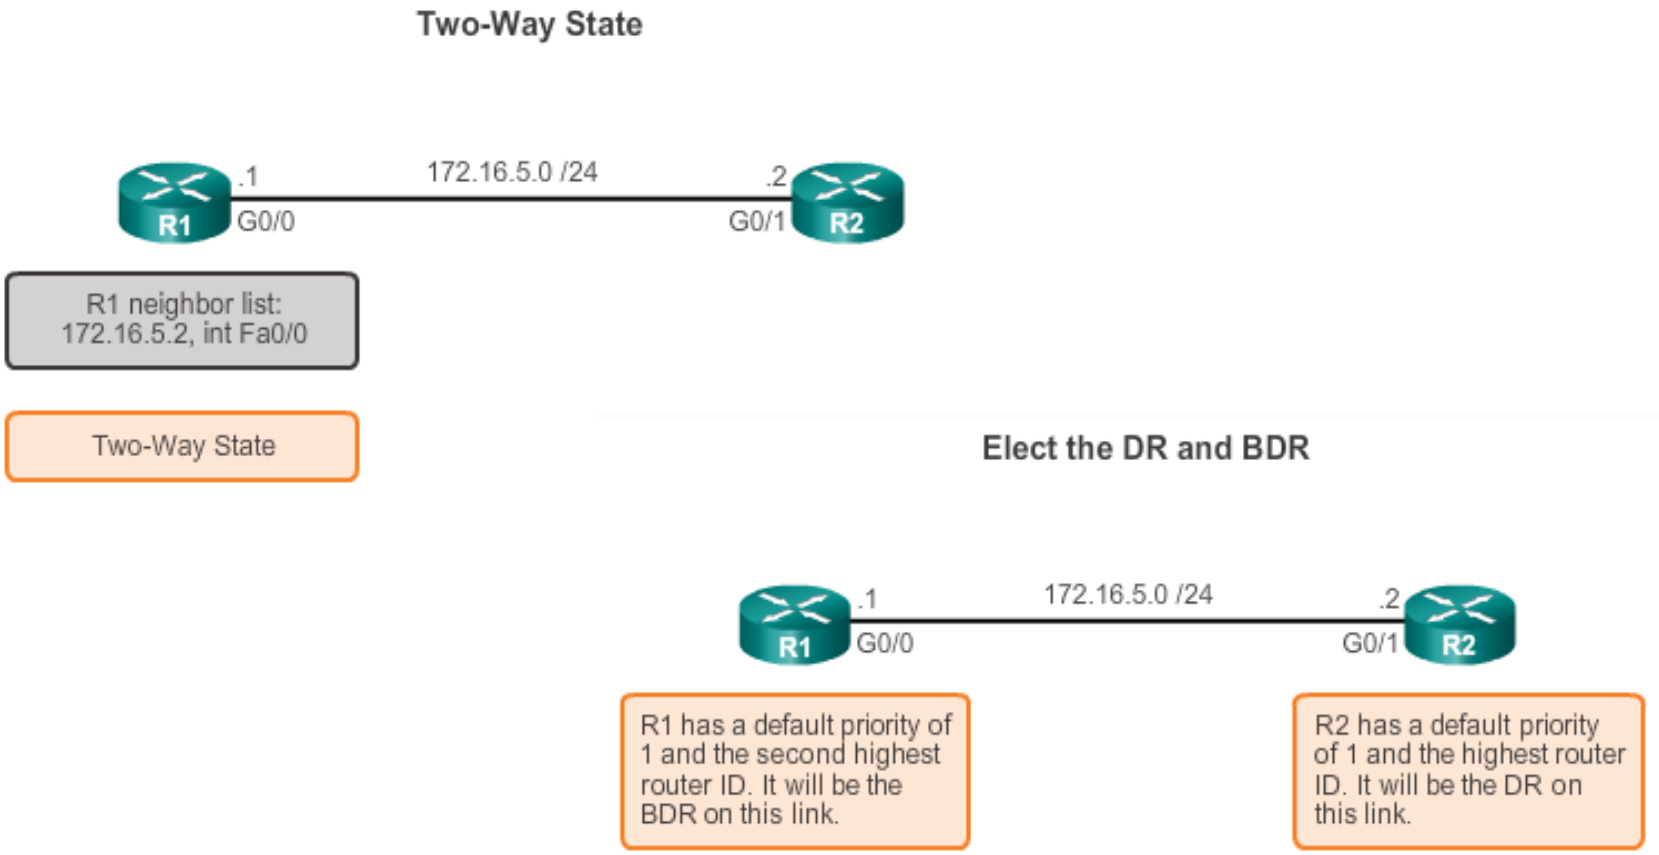

DR and BDR election only occurs on multi-access networks such as Ethernet LANs.

OSPF Operation

OSPF DR and BDR

Creating Adjacencies With Every Neighbor

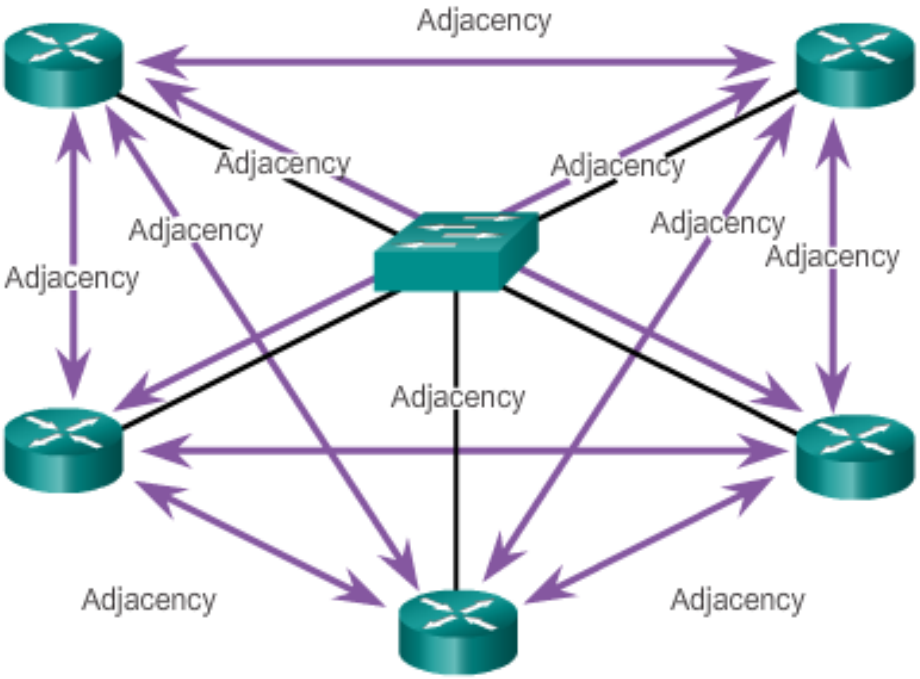

Number of Adjacencies=n(n-1)/2
n=number of routers
Example:5 routers (5-1)/2=10 adjacencies

Synchronizing OSPF Database

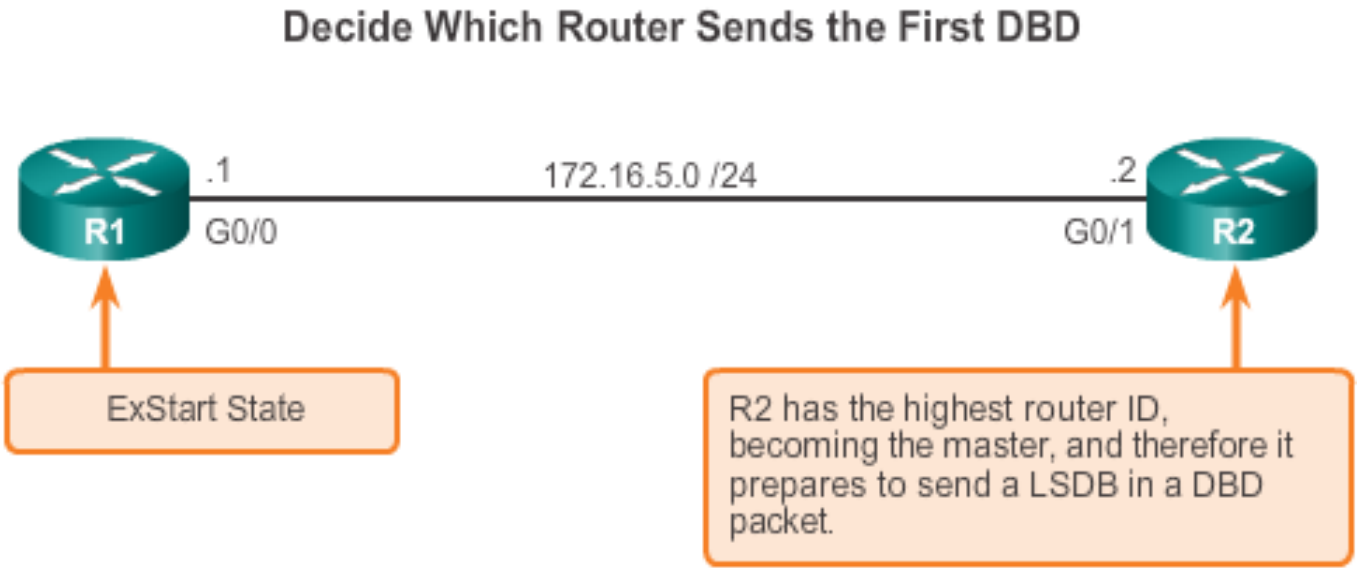

Synchronizing OSPF Database

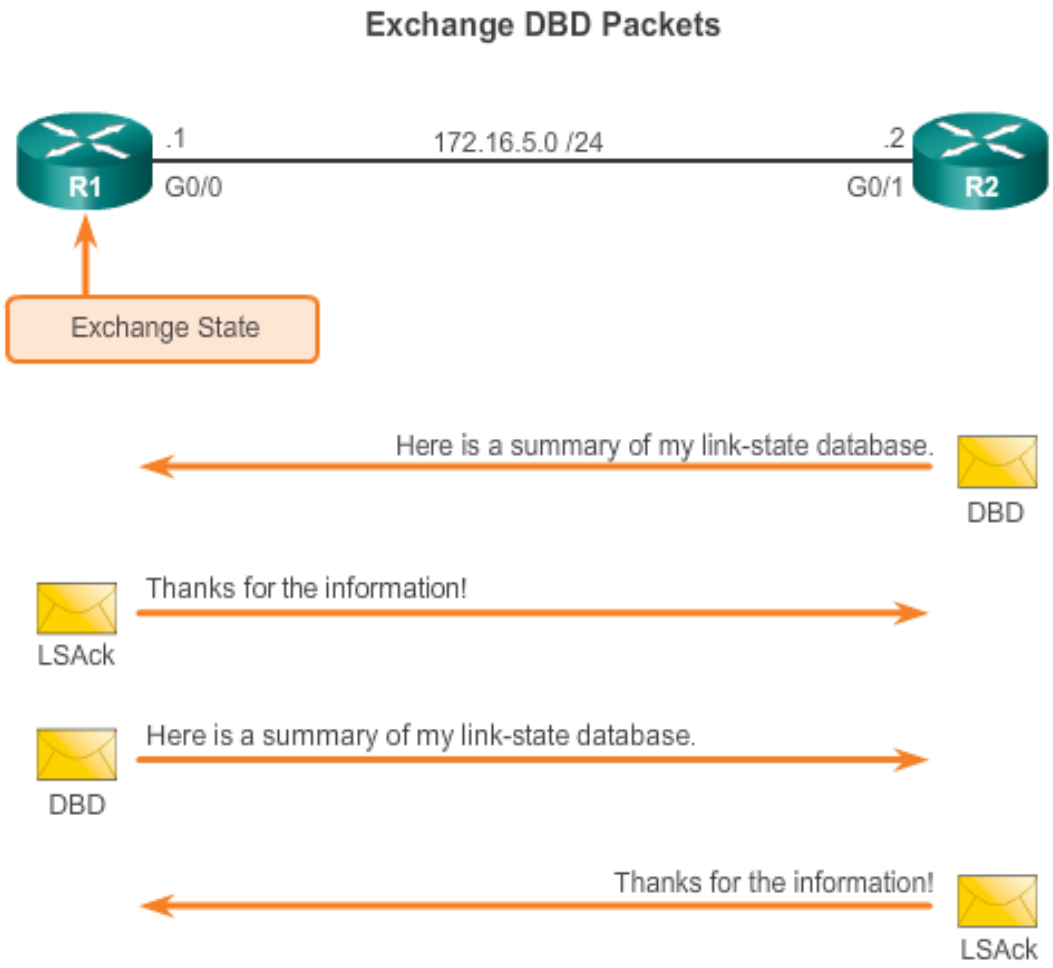

Trivia

- Who Invented the SPF algorithm?

Configuring Single-area OSPFv2

Cisco | Networking Academy®
Mind Wide Open™

OSPF Router ID

OSPF Network Topology

Entering Router OSPF Configuration Mode on R1

```
R1(config)# router ospf 10
R1(config-router)# ?
Router configuration commands:
  auto-cost               Calculate OSPF interface cost
                          according to bandwidth
  network                 Enable routing on an IP network
  no                      Negate a command or set its defaults
  passive-interface       Suppress routing updates on an
                          interface
  priority                OSPF topology priority
  router-id               router-id for this OSPF process
```

Note: Output has been altered to display only the commands that will be used in this chapter.

OSPF Router ID

Router IDs

```
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
% OSPF: Reload or use "clear ip ospf process" command, for
this to take effect
R1(config-router)# end
R1#
*Mar 25 19:46:09.711: %SYS-5-CONFIG_I: Configured from
console by console
```

```
R1(config)# interface loopback 0
R1(config-if)# ip address 1.1.1.1 255.255.255.255
R1(config-if)# end
R1#
```

Router ID Order of Precedence

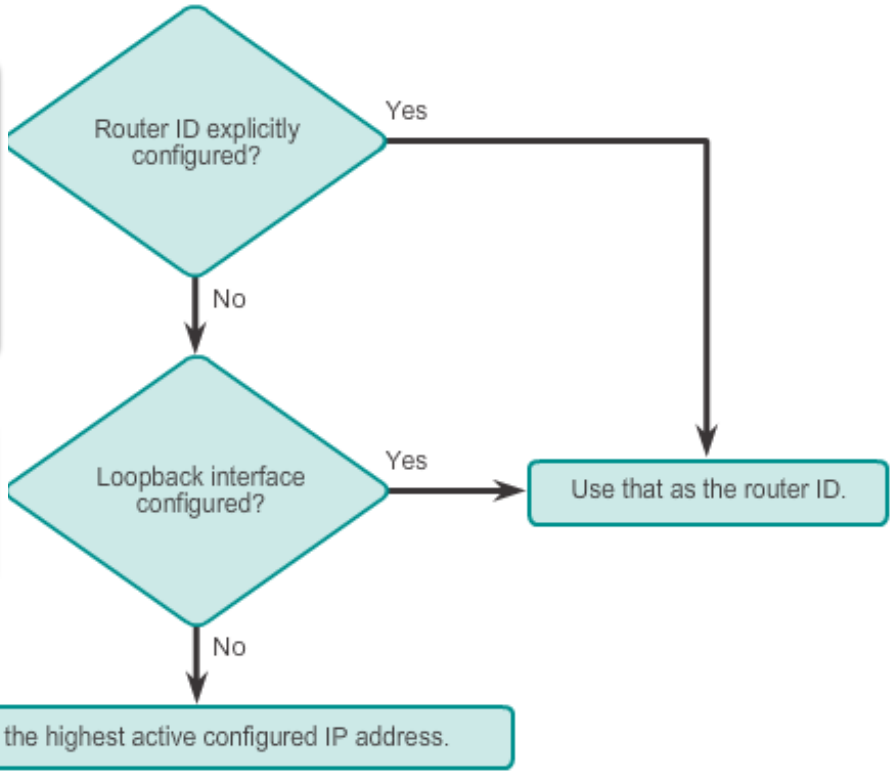

Clearing the OSPF Process

```
R1# clear ip ospf process
Reset ALL OSPF processes? [no]: y
R1#
*Mar 25 19:46:22.423: %OSPF-5-ADJCHG: Process 10, Nbr
3.3.3.3 on Serial0/0/1 from FULL to DOWN, Neighbor Down:
Interface down or detached
*Mar 25 19:46:22.423: %OSPF-5-ADJCHG: Process 10, Nbr
2.2.2.2 on Serial0/0/0 from FULL to DOWN, Neighbor Down:
Interface down or detached
```

The network Command

Assigning Interfaces to an OSPF Area

```
R1(config)# router ospf 10
R1(config-router)# network 172.16.1.0 0.0.0.255 area 0
R1(config-router)# network 172.16.3.0 0.0.0.3 area 0
R1(config-router)# network 192.168.10.4 0.0.0.3 area 0
R1(config-router)#
R1#
```

Assigning Interfaces to an OSPF Area with a Quad Zero

```
R1(config)# router ospf 10
R1(config-router)# network 172.16.1.1 0.0.0.0 area 0
R1(config-router)# network 172.16.3.1 0.0.0.0 area 0
R1(config-router)# network 192.168.10.5 0.0.0.0 area 0
R1(config-router)#
R1#
```

Configuring Passive Interfaces

Configuring a Passive Interface on R1

```
R1(config)# router ospf 10
R1(config-router)# passive-interface GigabitEthernet 0/0
R1(config-router)# end
R1#
```

Use the **passive-interface** router configuration mode command to prevent the transmission of routing messages through a router interface, but still allow that network to be advertised to other routers.

OSPF Metric = Cost

Cost = *reference bandwidth* / *interface bandwidth*
(default reference bandwidth is 10^8)
Cost = 100,000,000 bps / *interface bandwidth in bps*

Default Cisco OSPF Cost Values

Interface Type	Reference Bandwidth in bps		Default Bandwidth in bps	Cost
Gigabit Ethernet 10 Gbps	100,000,000	÷	10,000,000,000	1
Gigabit Ethernet 1 Gbps	100,000,000	÷	1,000,000,000	1
Fast Ethernet 100 Mbps	100,000,000	÷	100,000,000	1
Ethernet 10 Mbps	100,000,000	÷	10,000,000	10
Serial 1.544 Mbps	100,000,000	÷	1,544,000	64
Serial 128 kbps	100,000,000	÷	128,000	781
Serial 64 kbps	100,000,000	÷	64,000	1562

Same Cos
due to
reference
bandwidth

OSPF Accumulates Costs

Cost of an OSPF route is the accumulated value from one router to the destination network

```
R1# show ip route | include 172.16.2.0
O        172.16.2.0/24 [110/65] via 172.16.3.2, 03:39:07,
         Serial0/0/0
R1#
R1# show ip route 172.16.2.0
Routing entry for 172.16.2.0/24
  Known via "ospf 10", distance 110, metric 65, type intra
  area
  Last update from 172.16.3.2 on Serial0/0/0, 03:39:15 ago
  Routing Descriptor Blocks:
  * 172.16.3.2, from 2.2.2.2, 03:39:15 ago, via Serial0/0/0
      Route metric is 65, traffic share count is 1
R1#
```

OSPF Cost

Adjusting the Reference Bandwidth

- Use the **command - auto-cost reference-bandwidth**
- Must be configured on every router in the OSPF domain
- Notice that the value is expressed in Mb/s:
  Gigabit Ethernet - auto-cost reference-bandwidth 1000
  10 Gigabit Ethernet - auto-cost reference-bandwidth 10000

Verifying the S0/0/0 Link Cost

```
R1# show ip ospf interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
 Internet Address 172.16.3.1/30,Area 0,Attached via Network Statement
 Process ID 10,Router ID 1.1.1.1,Network Type POINT_TO_POINT,Cost:647
 Topology-MTID    Cost    Disabled    Shutdown      Topol
        0           647       no          no            E
 Transmit Delay is 1 sec, State POINT_TO_POINT
 Timer intervals configured, Hello 10, Dead 40, Wait 40,
   oob-resync timeout 40
   Hello due in 00:00:01
 Supports Link-local Signaling (LLS)
 Cisco NSF helper support enabled
 IETF NSF helper support enabled
 Index 3/3, flood queue length 0
 Next 0x0(0)/0x0(0)
 Last flood scan length is 1, maximum is 1
 Last flood scan time is 0 msec, maximum is 0 msec
 Neighbor Count is 1, Adjacent neighbor count is 1
   Adjacent with neighbor 2.2.2.2
 Suppress hello for 0 neighbor(s)
R1#
```

Verifying the Metric to the R2 LAN

```
R1# show ip route | include 172.16.2.0
O       172.16.2.0/24 [110/648] via 172.16.3.2, 00:06:03, Serial0/0/0
R1#
R1# show ip route 172.16.2.0
Routing entry for 172.16.2.0/24
  Known via "ospf 10", distance 110, metric 648, type intra area
  Last update from 172.16.3.2 on Serial0/0/0, 00:06:17 ago
  Routing Descriptor Blocks:
  * 172.16.3.2, from 2.2.2.2, 00:06:17 ago, via Serial0/0/0
      Route metric is 648, traffic share count is 1
R1#
R1#
```

Default Interface Bandwidths

On Cisco routers, the default bandwidth on most serial interfaces is set to 1.544 Mb/s

Verifying the Default Bandwidth Settings of R1 Serial 0/0/0

```
R1# show interfaces serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is WIC MBRD Serial
  Description: Link to R2
  Internet address is 172.16.3.1/30
  MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation HDLC, loopback not set
  Keepalive set (10 sec)
  Last input 00:00:05, output 00:00:03, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total
```

OSPF Cost

Adjusting the Interface Bandwidths

Adjusting the R1 Serial 0/0/1 Interface

```
R1(config)# int s0/0/1
R1(config-if)# bandwidth 64
R1(config-if)# end
R1#
*Mar 27 10:10:07.735: %SYS-5-CONFIG_I: Configured from console by c
R1#
R1# show interfaces serial 0/0/1 | include BW
  MTU 1500 bytes, BW 64 Kbit/sec, DLY 20000 usec,
R1#
R1# show ip ospf interface serial 0/0/1 | include Cost:
  Process ID 10, Router ID 1.1.1.1, Network Type
  POINT_TO_POINT, Cost: 15625
R1#
```

Manually Setting the OSPF Cost

Both the **bandwidth** interface command and the **ip ospf cost** interface command achieve the same result, which is to provide an accurate value for use by OSPF in determining the best route.

```
R1(config)# int s0/0/1
R1(config-if)# no bandwidth 64
R1(config-if)# ip ospf cost 15625
R1(config-if)# end
R1#
R1# show interface serial 0/0/1 | include BW
    MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
R1#
R1# show ip ospf interface serial 0/0/1 | include Cost:
    Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT,
    Cost: 15625
R1#
```

Verify OSPF Neighbors

Verify that the router has formed an adjacency with its neighboring routers

```
R1# show ip ospf neighbor

Neighbor ID  Pri  State    Dead Time Address        Interface
3.3.3.3      0    FULL/-   00:00:37  192.168.10.6   Serial0/0/1
2.2.2.2      0    FULL/-   00:00:30  172.16.3.2     Serial0/0/0
R1#
```

Verify OSPF Protocol Settings

Verifying R1's OSPF Neighbors

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 10"
  Outgoing update filter list for all interfaces is not
  set
  Incoming update filter list for all interfaces is not
  set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0
  nssa
  Maximum path: 4
  Routing for Networks:
    172.16.1.0 0.0.0.255 area 0
    172.16.3.0 0.0.0.3 area 0
    192.168.10.4 0.0.0.3 area 0
  Routing Information Sources:
    Gateway         Distance      Last Update
    2.2.2.2              110      00:17:18
    3.3.3.3              110      00:14:49
  Distance: (default is 110)

R1#
```

Verify OSPF Interface Settings

Verifying R1's OSPF Interfaces

```
R1# show ip ospf interface brief
Interface   PID   Area   IP Address/Mask   Cost   State   Nbrs F/C
Se0/0/1     10    0      192.168.10.5/30   15625  P2P     1/1
Se0/0/0     10    0      172.16.3.1/30     647    P2P     1/1
Gi0/0       10    0      172.16.1.1/24     1      DR      0/0
R1#
```

Trivia

- What is the RFC number of OSPFv2?
- What is the RFC number of OSPFv3?

Configuring Single-area OSPFv3

Cisco | Networking Academy®
Mind Wide Open™

OSPFv2 vs. OSPFv3

OSPFv3

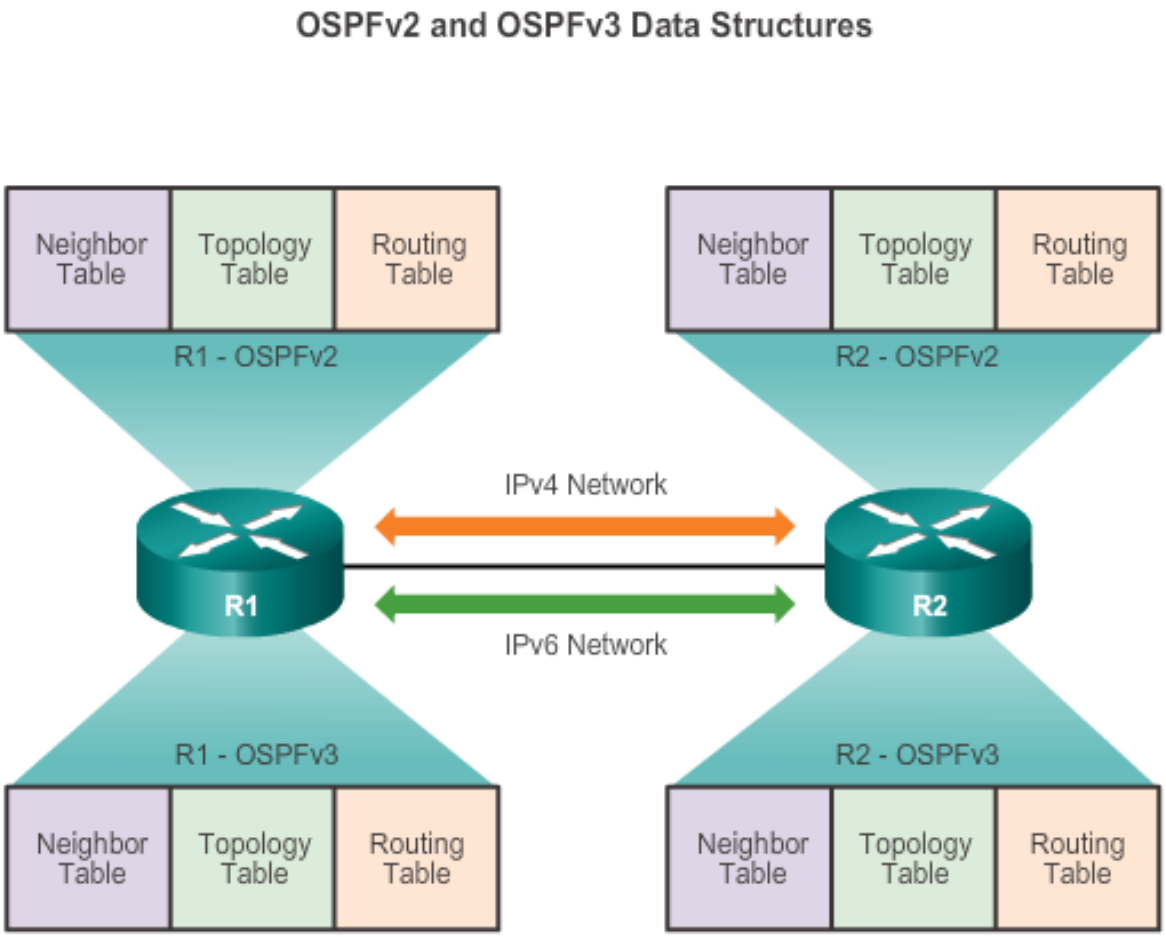

Similarities Between OSPFv2 to OSPFv3

OSPFv2 and OSPFv3	
Link-State	Yes
Routing Algorithm	SPF
Metric	Cost
Areas	Supports the same two-level hierarchy
Packet Types	Same Hello, DBD, LSR, LSU and LSAck packets
Neighbor Discovery	Transitions through the same states using Hello packets
DR and BDR	Function and election process is the same
Router ID	32-bit router ID: determined by the same process in both protocols

OSPFv2 vs. OSPFv3

Differences Between OSPFv2 to OSPFv3

	OSPFv2	OSPFv3
Advertises	IPv4 networks	IPv6 prefixes
Source Address	IPv4 source address	IPv6 link-local address
Destination Address	Choice of: • Neighbor IPv4 unicast address • 224.0.0.5 all-OSPF-routers multicast address • 224.0.0.6 DR/BDR multicast address	Choice of: • Neighbor IPv6 link-local address • FF02::5 all-OSPFv3-routers multicast address • FF02::6 DR/BDR multicast address
Advertise Networks	Configured using the **`network`** router configuration command	Configured using the **`ipv6 ospf process-id area-id`** interface configuration command
IP Unicast Routing	IPv4 unicast routing is enabled by default.	IPv6 unicast forwarding is not enabled by default. The **`ipv6 unicast-routing`** global configuration command must be configured.
Authentication	Plain text and MD5	IPv6 authentication

OSPFv2 vs. OSPFv3

Link-Local Addresses

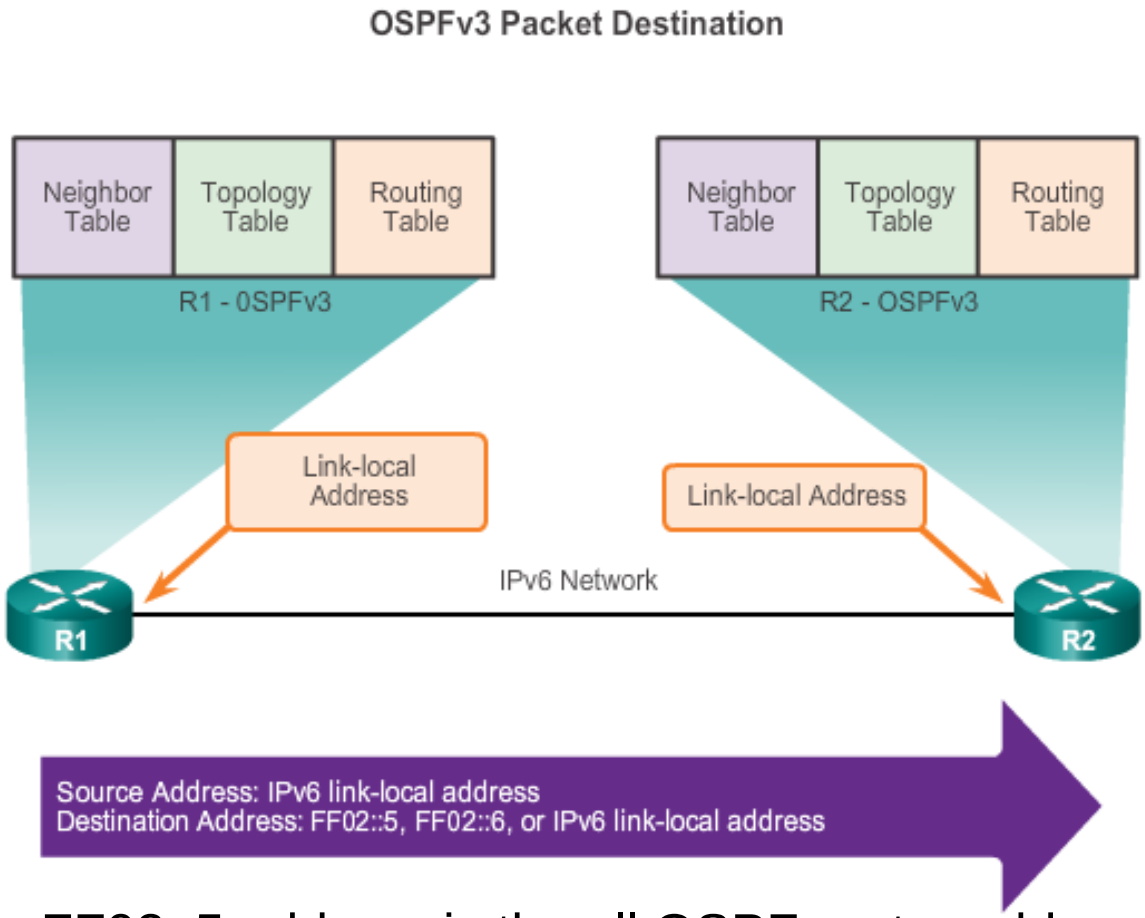

FF02::5 address is the all OSPF router address
FF02::6 is the DR/BDR multicast address

Configuring OSFPv3

OSPFv3 Network Topology

Configuring Global-Unicast Addresses on R1

```
R1(config)# ipv6 unicast-routing
R1(config)#
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# description R1 LAN
R1(config-if)# ipv6 address 2001:DB8:CAFE:1::1/64
R1(config-if)# no shut
R1(config-if)#
R1(config-if)# interface Serial0/0/0
R1(config-if)# description Link to R2
R1(config-if)# ipv6 address 2001:DB8:CAFE:A001::1/64
R1(config-if)# clock rate 128000
R1(config-if)# no shut
R1(config-if)#
R1(config-if)# interface Serial0/0/1
R1(config-if)# description Link to R3
R1(config-if)# ipv6 address 2001:DB8:CAFE:A003::1/64
R1(config-if)# no shut
R1(config-if)# end
R1#
```

Link-Local Addresses

```
R1# show ipv6 interface brief
Em0/0                      [administratively down/down]
    unassigned
GigabitEthernet0/0         [up/up]
    FE80::32F7:DFF:FEA3:DA0
    2001:DB8:CAFE:1::1
GigabitEthernet0/1         [administratively down/down]
    unassigned
Serial0/0/0                [up/up]
    FE80::32F7:DFF:FEA3:DA0
    2001:DB8:CAFE:A001::1
Serial0/0/1                [up/up]
    FE80::32F7:DFF:FEA3:DA0
    2001:DB8:CAFE:A003::1
R1#
```

- Link-local addresses are automatically created when an IPv6 global unicast address is assigned to the interface (required).
- Global unicast addresses are not required.
- Cisco routers create the link-local address using FE80::/10 prefix and the EUI-64 process unless the router is configured manually,
- EUI-64 involves using the 48-bit Ethernet MAC address, inserting FFFE in the middle and flipping the seventh bit. For serial interfaces, Cisco uses the MAC address of an Ethernet interface.
- Notice in the figure that all three interfaces are using the same link-local address.

Assigning Link-Local Addresses

```
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# exit
R1(config)# interface Serial0/0/0
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# exit
R1(config)# interface Serial0/0/1
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)#
```

Configuring the link-local address provides the ability to create an address that is recognizable and easier to remember

```
R1# show ipv6 interface brief
Em0/0                  [administratively down/down]
    unassigned
GigabitEthernet0/0     [up/up]
    FE80::1
    2001:DB8:CAFE:1::1
GigabitEthernet0/1     [administratively down/down]
    unassigned
Serial0/0/0            [up/up]
    FE80::1
    2001:DB8:CAFE:A001::1
Serial0/0/1            [up/up]
    FE80::1
    2001:DB8:CAFE:A003::1
R1#
```

Configuring the OSPFv3 Router ID

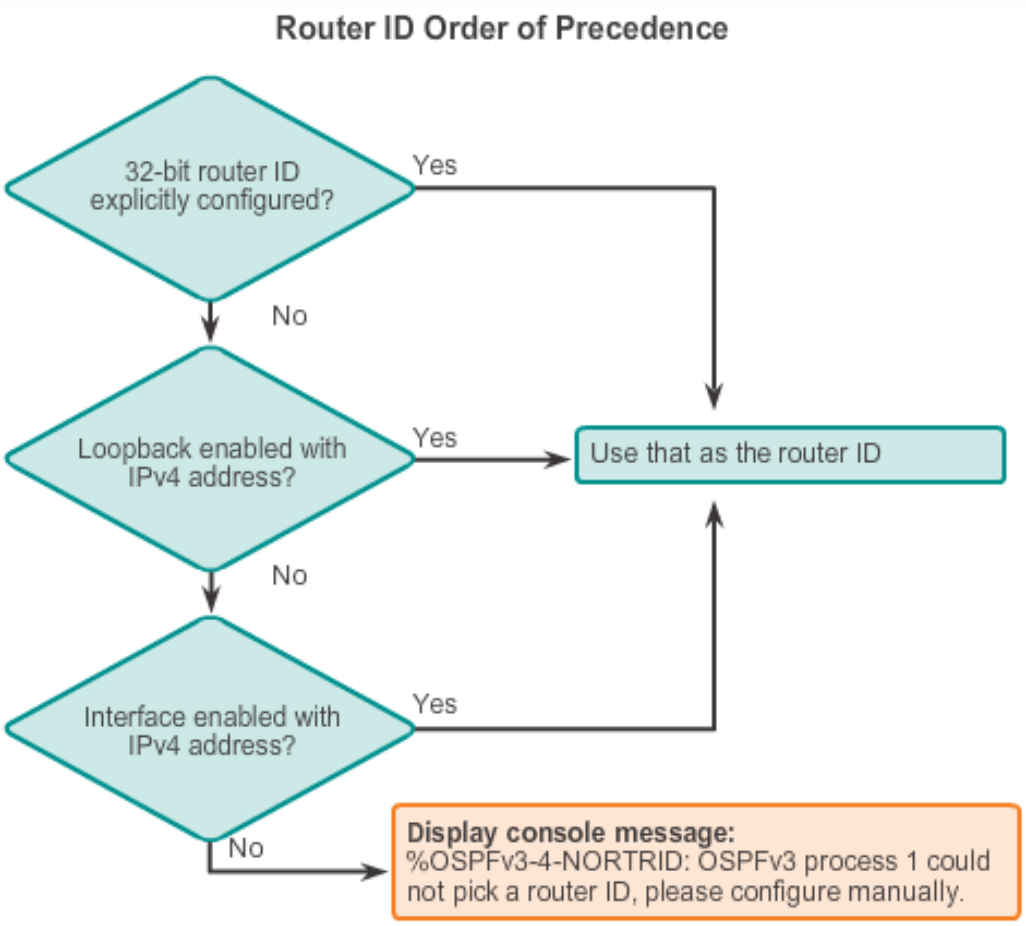

Configuring the OSPFv3 Router ID

Assigning a Router ID to R1

```
R1(config)# ipv6 router ospf 10
R1(config-rtr)#
*Mar 29 11:21:53.739: %OSPFv3-4-NORTRID: Process OSPFv3-1-
IPv6 could not pick a router-id, please configure manually
R1(config-rtr)#
R1(config-rtr)# router-id 1.1.1.1
R1(config-rtr)#
R1(config-rtr)# auto-cost reference-bandwidth 1000
% OSPFv3-1-IPv6: Reference bandwidth is changed. Please
ensure reference bandwidth is consistent across all routers.
R1(config-rtr)#
R1(config-rtr)# end
R1#
R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 1.1.1.1
  Number of areas: 0 normal, 0 stub, 0 nssa
  Redistribution:
    None
R1#
```

Configuring OSFPv3

Modifying an OSPFv3 Router ID

```
R1(config)# ipv6 router ospf 10
R1(config-rtr)# router-id 1.1.1.1
R1(config-rtr)# end
R1#
```

```
R1# clear ipv6 ospf process
Reset selected OSPFv3 processes? [no]: y
R1#
R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 1.1.1.1
  Number of areas: 0 normal, 0 stub, 0 nssa
  Redistribution:
    None
R1#
```

Enabling OSPFv3 on Interfaces

Instead of using the **network** router configuration mode command to specify matching interface addresses, OSPFv3 is configured directly on the interface.

```
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)#
R1(config-if)# interface Serial0/0/0
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)#
R1(config-if)# interface Serial0/0/1
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)#
R1(config-if)# end
R1#
R1# show ipv6 ospf interfaces brief
Interface  PID   Area      Intf ID  Cost   State  Nbrs F/C
Se0/0/1    10      0         7      15625  P2P    0/0
Se0/0/0    10      0         6      647    P2P    0/0
Gi0/0      10      0         3       1     WAIT   0/0
R1#
```

Verify OSPFv3

Verify OSPFv3 Neighbors/Protocol Settings

```
R1# show ipv6 ospf neighbor

OSPFv3 Router with ID (1.1.1.1) (Process ID 10)

Neighbor ID     Pri   State           Dead Time   Interface ID  Interface
3.3.3.3           0   FULL/  -        00:00:39    6             Serial0/0/1
2.2.2.2           0   FULL/  -        00:00:36    6             Serial0/0/0
R1#
```

```
R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 1.1.1.1
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    Serial0/0/1
    Serial0/0/0
    GigabitEthernet0/0
  Redistribution:
    None
R1#
```

Verify OSPFv3 Interfaces

```
R1# show ipv6 ospf interface brief
Interface    PID   Area            Intf ID    Cost  State Nbrs F/C
Se0/0/1      10    0               7          15625 P2P   1/1
Se0/0/0      10    0               6          647   P2P   1/1
Gi0/0        10    0               3          1     DR    0/0
R1#
```

Verify OSPFv3

Verify IPv6 Routing Table

```
R1# show ipv6 route ospf
IPv6 Routing Table - default - 10 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user
Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS
summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND
Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter,
OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF
NSSA ext 2
O   2001:DB8:CAFE:2::/64 [110/657]
     via FE80::2, Serial0/0/0
O   2001:DB8:CAFE:3::/64 [110/1304]
     via FE80::2, Serial0/0/0
O   2001:DB8:CAFE:A002::/64 [110/1294]
     via FE80::2, Serial0/0/0
R1#
```

Trivia

- Do we have any ISPs supporting IPv6?

Advanced Single-Area OSPF Configurations

Cisco | Networking Academy®
Mind Wide Open™

Configuring Single-Area OSPF

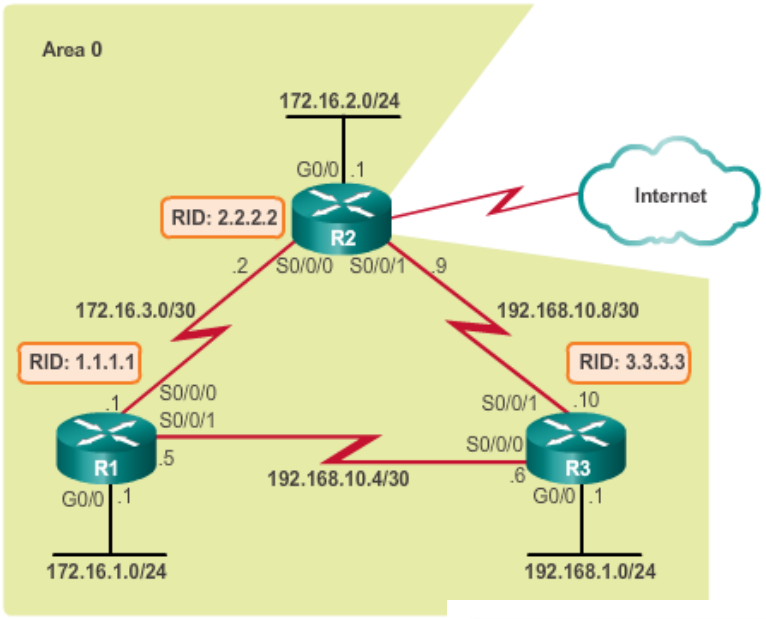


```
R1(config)# interface GigabitEthernet0/0
R1(config-if)# bandwidth 1000000
R1(config-if)# exit
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
R1(config-router)# auto-cost reference-bandwidth 1000
% OSPF: Reference bandwidth is changed.
        Please ensure reference bandwidth is consistent
across all routers.
R1(config-router)# network 172.16.1.0 0.0.0.255 area 0
R1(config-router)# network 172.16.3.0 0.0.0.3 area 0
R1(config-router)# network 192.168.10.4 0.0.0.3 area 0
R1(config-router)#
R1(config-router)# passive-interface g0/0
R1(config-router)#
```

```
R2(config)# interface GigabitEthernet0/0
R2(config-if)# bandwidth 1000000
R2(config-if)# exit
R2(config)# router ospf 10
R2(config-router)# router-id 2.2.2.2
R2(config-router)# auto-cost reference-bandwidth 1000
% OSPF: Reference bandwidth is changed.
        Please ensure reference bandwidth is consistent
across all routers.
R2(config-router)# network 172.16.2.1 0.0.0.0 area 0
R2(config-router)# network 172.16.3.2 0.0.0.0 area 0
R2(config-router)# network 192.168.10.9 0.0.0.0 area 0
R2(config-router)#
R2(config-router)# passive-interface g0/0
R2(config-router)#
```

Routing in the Distribution and Core Layers

Verifying Single-Area OSPF

```
R1# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           0   FULL/  -        00:00:32    192.168.10.6    Serial0/0/1
2.2.2.2           0   FULL/  -        00:00:38    172.16.3.2      Serial0/0/0
R1#
```

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    172.16.1.0 0.0.0.255 area 0
    172.16.3.0 0.0.0.3 area 0
    192.168.10.4 0.0.0.3 area 0
  Passive Interface(s):
    GigabitEthernet0/0
  Routing Information Sources:
    Gateway         Distance      Last Update
    3.3.3.3              110      00:12:14
    2.2.2.2              110      00:12:46
  Distance: (default is 110)

R1#v
```

Verifying Single-Area OSPF (cont.)

```
R1# show ip ospf
 Routing Process "ospf 10" with ID 1.1.1.1
 Start time: 00:06:18.952, Time elapsed: 00:39:56.400

 <Output omitted>

 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 Number of areas transit capable is 0
 External flood list length 0
 IETF NSF helper support enabled
 Cisco NSF helper support enabled
 Reference bandwidth unit is 1000 mbps
 Area BACKBONE(0)
        Number of interfaces in this area is 3
Area has no authentication
SPF algorithm last executed 00:15:21.436 ago
SPF algorithm executed 6 times
Area ranges are
Number of LSA 3. Checksum Sum 0x023523
Number of opaque link LSA 0. Checksum Sum 0x000000
Number of DCbitless LSA 0
Number of indication LSA 0
Number of DoNotAge LSA 0
Flood list length 0

R1#
```

Verifying Single-Area OSPF (cont.)

```
R1# show ip ospf interface
GigabitEthernet0/0 is up, line protocol is up
  Internet Address 172.16.1.1/24, Area 0, Attached via Network
Statement
  Process ID 10, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1         no          no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.1, Interface address 172.16.1.1
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
Retransmit 5
    oob-resync timeout 40
    No Hellos (Passive interface)
  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 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
Serial0/0/1 is up, line protocol is up
  Internet Address 192.168.10.5/30, Area 0, Attached via Network
Statement
  Process ID 10, Router ID 1.1.1.1, Network Type POINT_TO_POINT,
Cost: 647

 <Output omitted>
```

```
R1# show ip ospf interface brief
Interface    PID   Area    IP Address/Mask    Cost  State Nbrs F/C
Gi0/0        10    0       172.16.1.1/24      1     DR    0/0
Se0/0/1      10    0       192.168.10.5/30    647   P2P   1/1
Se0/0/0      10    0       172.16.3.1/30      647   P2P   1/1
R1#
```

Configuring Single-Area OSPFv3

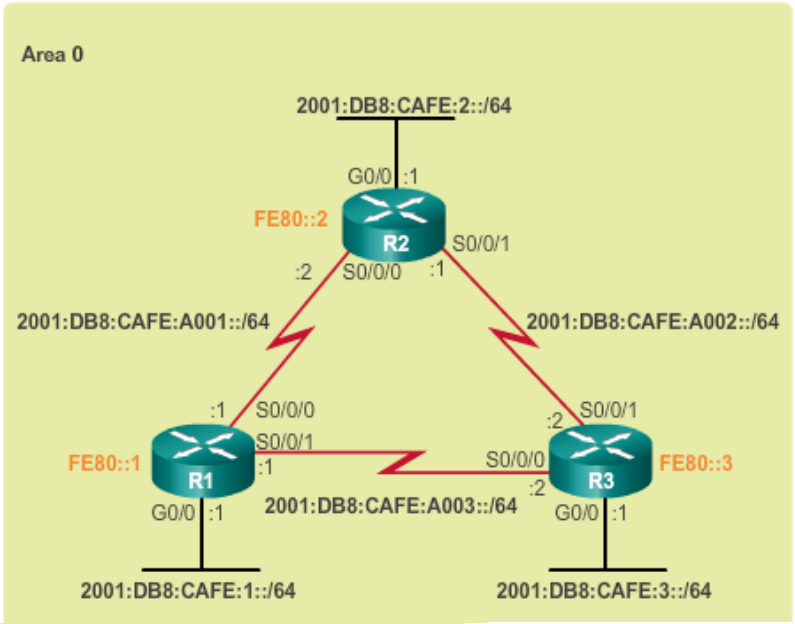


```
R1(config)# ipv6 router ospf 10
R1(config-rtr)# router-id 1.1.1.1
R1(config-rtr)# auto-cost reference-bandwidth 1000
% OSPFv3-10-IPv6:  Reference bandwidth is changed.
        Please ensure reference bandwidth is consistent across all
routers.
R1(config-rtr)#
R1(config-rtr)# interface GigabitEthernet 0/0
R1(config-if)# bandwidth 1000000
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)#
R1(config-if)# interface Serial0/0/0
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)#
R1(config-if)# interface Serial0/0/1
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# end
R1#
```

```
R2(config)# ipv6 router ospf 10
R2(config-rtr)# router-id 2.2.2.2
R2(config-rtr)# auto-cost reference-bandwidth 1000
% OSPFv3-10-IPv6:  Reference bandwidth is changed.
        Please ensure reference bandwidth is consistent across all
routers.
R2(config-rtr)#
R2(config-rtr)# interface GigabitEthernet 0/0
R2(config-if)# bandwidth 1000000
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)#
R2(config-if)# interface Serial0/0/0
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)#
R2(config-if)# interface Serial0/0/1
R2(config-if)# ipv6 ospf 10 area 0
R2(config-if)# end
R2#
```

Routing in the Distribution and Core Layers

Verifying Single-Area OSPFv3

```
R1# show ipv6 ospf neighbor

            OSPFv3 Router with ID (1.1.1.1) (Process ID 10)

Neighbor ID Pri State       Dead Time Interface ID Interface
3.3.3.3     0  FULL/  -     00:00:31  6            Serial0/0/1
2.2.2.2     0  FULL/  -     00:00:37  6            Serial0/0/0
2.2.2.2     1  FULL/BDR     00:00:38  3            GigabitEthernet0/0
3.3.3.3     1  FULL/DROTHER 00:00:32  3            GigabitEthernet0/0
R1#
```

```
R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 1.1.1.1
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    Serial0/0/1
    Serial0/0/0
    GigabitEthernet0/0
  Redistribution:
    None
R1#
```

Verifying Single-Area OSPFv3 (cont.)

```
R1# show ipv6 route ospf
IPv6 Routing Table - default - 10 entries
Codes: C - Connected, L - Local, S - Static,
       U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary,
       D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix,
       DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter,
       OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1,
       ON2 - OSPF NSSA ext 2
O   2001:DB8:CAFE:2::/64 [110/1]
     via GigabitEthernet0/0, directly connected
O   2001:DB8:CAFE:3::/64 [110/1]
     via GigabitEthernet0/0, directly connected
O   2001:DB8:CAFE:A002::/64 [110/648]
     via FE80::2, GigabitEthernet0/0
     via FE80::3, GigabitEthernet0/0
R1#
```

```
R1# show ipv6 ospf interface brief
Interface    PID   Area            Intf ID    Cost  State Nbrs F/C
Se0/0/1      10    0               7          647   P2P   1/1
Se0/0/0      10    0               6          647   P2P   1/1
Gi0/0        10    0               3          1     DR    2/2
R1#
```

OSPF Network Types

- **Point-to-point** – Two routers interconnected over a common link. Often the configuration in WAN links.
- **Broadcast Multiaccess** – Multiple routers interconnected over an Ethernet network.
- **Non-broadcast Multiaccess (NBMA)** – Multiple routers interconnected in a network that does not allow broadcasts, such as Frame Relay.
- **Point-to-multipoint** – Multiple routers interconnected in a hub-and-spoke topology over an NBMA network.
- **Virtual links** – Special OSPF network used to interconnect distant OSPF areas to the backbone area.

Challenges in Multiaccess Networks

Multiaccess networks can create two challenges for OSPF:

- **Creation of multiple adjacencies** – creating adjacencies with multiple routers would lead to an excessive number of LSAs being exchanged.
- **Extensive flooding of LSAs** – Link-state routers flood the network when OSPF is initialized or when there is a change.

- Formula used to calculate the number of required adjacencies n(n-1)/2
- A topology of 4 routers would result in 4(4-1)/2 = 6

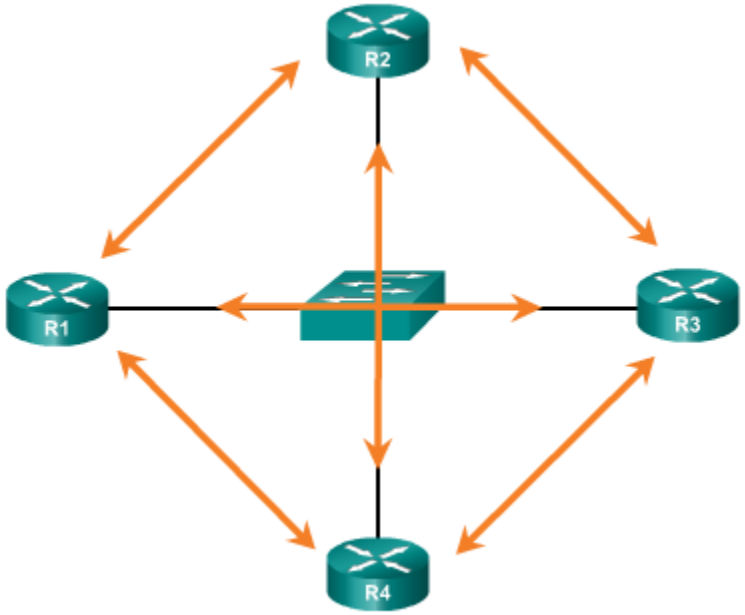

OSPF Designated Router

- The designated router (DR) is the solution to managing adjacencies and flooding of LSAs on a multiaccess network.
- The backup designated router (BDR) is elected in case the DR fails.
- All other non-DR and non-BDR routers become DROTHERs. DROTHERs only form adjacencies with the DR and BDR.
- DROTHERs only send their LSAs to the DR and BDR using the multicast address 224.0.0.6.
- DR uses the multicast address 224.0.0.5 to send LSAs to all other routers. DR only router flooding LSAs.
- DR/BDR Elections only necessary on multiaccess networks.

OSPF Designated Router (cont.)

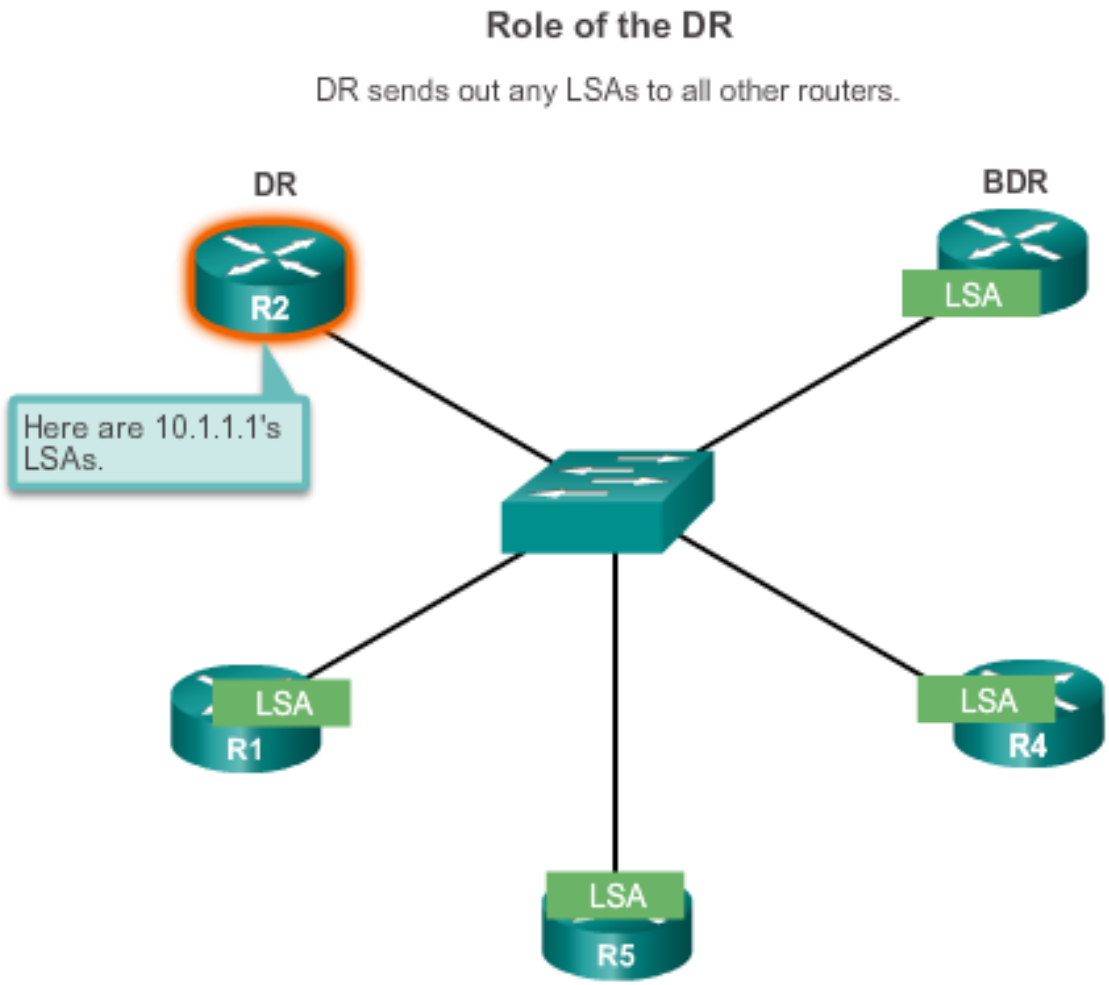

OSPF in Multiaccess Networks

Verifying DR/BDR Roles

Verifying the Role of R1

```
R1# show ip ospf interface GigabitEthernet 0/0
GigabitEthernet0/0 is up, line protocol is up
 Internet Address 192.168.1.1/28,Area 0,Attached via Network Statement
 Process ID 10, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
 Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1        no          no            Base
1 Transmit Delay is 1 sec, State DROTHER, Priority 1
 Designated Router (ID) 3.3.3.3, Interface address 192.168.1.3
2 Backup Designated router (ID) 2.2.2.2, Interface address 192.168.1.2
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
   oob-resync timeout 40
   Hello due in 00:00:06
 Supports Link-local Signaling (LLS)
 Cisco NSF helper support enabled
 IETF NSF helper support enabled
 Index 2/2, flood queue length 0
 Next 0x0(0)/0x0(0)
 Last flood scan length is 1, maximum is 2
 Last flood scan time is 0 msec, maximum is 0 msec
 Neighbor Count is 2, Adjacent neighbor count is 2
3   Adjacent with neighbor 2.2.2.2  (Backup Designated Router)
   Adjacent with neighbor 3.3.3.3  (Designated Router)
 Suppress hello for 0 neighbor(s)
R1#
```

Verifying DR/BDR Adjacencies

State of neighbors in multiaccess networks can be:

- **FULL/DROTHER** – This is a DR or BDR router that is fully adjacent with a non-DR or BDR router.
- **FULL/DR** – The router is fully adjacent with the indicated DR neighbor.
- **FULL/BDR** – The router is fully adjacent with the indicated BDR neighbor.
- **2-WAY/DROTHER** – The non-DR or BDR router has a neighbor adjacency with another non-DR or BDR router.

```
R1# show ip ospf neighbor

Neighbor ID Pri State       Dead Time  Address     Interface
1 2.2.2.2     1  FULL/BDR     00:00:36  192.168.1.2 GigabitEthernet0/0
2 3.3.3.3     1  FULL/DR      0:00:35   192.168.1.3 GigabitEthernet0/0
R1#
```

Default DR/BDR Election Process

- The router with the highest interface priority is elected as the DR.
- The router with the second highest interface priority is elected as the BDR.
- Priority can be configured between 0-255. (Priority of 0 - router cannot become the DR. 0
- If interface priorities are equal, then the router with highest router ID is elected DR and second highest the BDR
- Three ways to determine router ID:
  - Router ID can be manually configured.
  - If not configured, the ID determined by the highest loopback IP address.
  - If no loopbacks, the ID is determined by the highest active IPv4 address.
- In an IPv6 network, the router ID must be configured manually.

DR/BDR Election Process

DR remains the DR until one of the following occurs:

- The DR fails.
- The OSPF process on the DR fails or is stopped.
- The multiaccess interface on the DR fails or is shutdown.

If the DR fails, the BDR is automatically promoted to DR.

- There is then a new BDR election and the DROTHER with the higher priority or router ID is elected as the new BDR.

The OSPF Priority

- Instead of setting the router ID on all routers, it is better to control the election by setting interface priorities.
  - To change the priority, use one of the following commands:

    `ip ospf priority` *value* (OSPFv2 interface command)

    `ipv6 ospf priority` *value* (OSPFv3 interface command)
- To begin another OSPF election, use one of the following methods:
  - Shutdown the router interfaces and then re-enable them starting with the DR, then the BDR, and then all other routers.
  - Reset the OSPF process using the **`clear ip ospf process`** privileged EXEC mode command on all routers.

```
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ip ospf priority 255
R1(config-if)# end
R1#
```

Propagating a Default Static Route in OSPFv2

The router connected to the Internet that is used to propagate a default route is often called the edge, entrance or gateway router. In an OSPF network, it may also be call the autonomous system boundary router (ASBR).

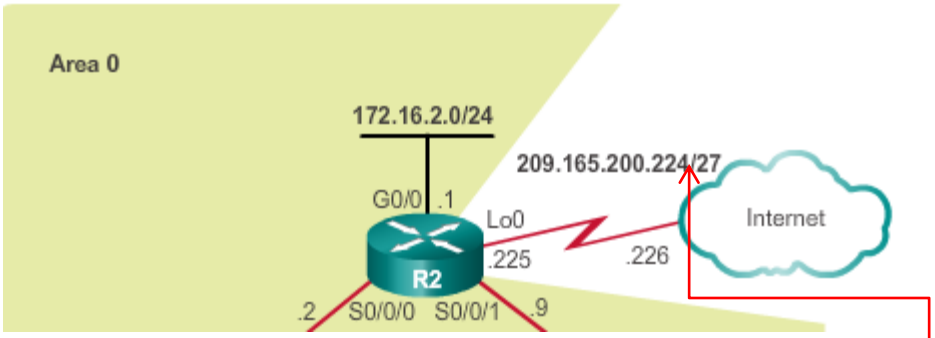


```
R2(config)# ip route 0.0.0.0 0.0.0.0 209.165.200.226
R2(config)#
R2(config)# router ospf 10
R2(config-router)# default-information originate
R2(config-router)# end
R2#
```

Default Route Propagation

Verifying the Propagated Default Route

```
R2# show ip route | begin Gateway

Gateway of last resort is 209.165.200.226 to network
0.0.0.0

S* 0.0.0.0/0 [1/0] via 209.165.200.226, Loopback0
   172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
O  172.16.1.0/24 [110/65] via 172.16.3.1, 00:01:44,
   Serial0/0/0
C  172.16.2.0/24 is directly connected, GigabitEthernet0/0
L  172.16.2.1/32 is directly connected, GigabitEthernet0/0
C  172.16.3.0/30 is directly connected, Serial0/0/0
L  172.16.3.2/32 is directly connected, Serial0/0/0
O  192.168.1.0/24 [110/65] via 192.168.10.10, 00:01:12,
   Serial0/0/1
   192.168.10.0/24 is variably subnetted, 3 subnets, 2
    masks
O  192.168.10.4/30 [110/128] via 192.168.10.10, 00:01:12,
   Serial0/0/1
            [110/128] via 172.16.3.1, 00:01:12, Serial0/0/0
C  192.168.10.8/30 is directly connected, Serial0/0/1
L  192.168.10.9/32 is directly connected, Serial0/0/1
   209.165.200.0/24 is variably subnetted, 2 subnets, 2
    masks
```

Propagating a Default Static Route in OSPFv3

Enabling OSPFv3 on the R1 Interfaces

```
R2(config)# ipv6 route ::/0 2001:DB8:FEED:1::2
R2(config)#
R2(config)# ipv6 router ospf 10
R2(config-rtr)# default-information originate
R2(config-rtr)# end
R2#
*Apr 10 11:36:21.995: %SYS-5-CONFIG_I: Configured from console by
console
R2#
```

Verifying the propagated IPv6 default Route

```
R2# show ipv6 route static
IPv6 Routing Table - default - 12 entries
Codes:C -Connected, L - Local, S - Static, U - Per-user Static route
      B -BGP, R - RIP, H - NHRP, I1 - ISIS L1
      I2 -ISIS L2, IA - ISIS interarea, IS-ISIS summary,D-EIGRP
      EX -EIGRP external, ND-ND Default,NDp-ND Prefix,DCE-Destination
      NDr -Redirect, O - OSPF Intra,OI-OSPF Inter,OE1-OSPF ext 1
      OE2 -OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S   ::/0 [1/0]
     via 2001:DB8:FEED:1::2, Loopback0
R2#
```

Fine-tuning OSPF Interfaces

OSPF Hello and Dead Intervals

OSPF Hello and Dead intervals must match, or a neighbor adjacency will not occur.

Verifying the OSPF Intervals on R1

```
R1# show ip ospf interface serial 0/0/0 | include Timer
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
Retransmit 5
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
Retransmit 5
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
Retransmit 5
R1#
```

Verifying OSPF Timer Activity

```
R1# show ip ospf neighbor

Neighbor ID  Pri   State     Dead Time   Address         Interface
3.3.3.3        0   FULL/-    00:00:35    192.168.10.6    Serial0/0/1
2.2.2.2        0   FULL/-    00:00:33    172.16.3.2      Serial0/0/0
R1#
```

Fine-tuning OSPF Interfaces

Modifying OSPF Intervals

- Modifying OSPFv2 Intervals

```
R1(config)# interface serial 0/0/0
R1(config-if)# ip ospf hello-interval 5
R1(config-if)# ip ospf dead-interval 20
R1(config-if)# end
R1#
```

- Modifying OSPFv3 Intervals

```
R1(config)# interface serial 0/0/0
R1(config-if)# ipv6 ospf hello-interval 5
R1(config-if)# ipv6 ospf dead-interval 20
R1(config-if)# end
R1#
```

- Verifying the OSPFv3 interface intervals

```
R2# show ipv6 ospf interface s0/0/0 | include Timer
  Timer intervals configured, Hello 5, Dead 20, Wait 20,
Retransmit 5
R2#
R2# show ipv6 ospf neighbor

            OSPFv3 Router with ID (2.2.2.2) (Process ID 10)

Neighbor ID  Pri   State   Dead Time   Interface ID  Interface
3.3.3.3        0   FULL/-  00:00:38    7             Serial0/0/1
1.1.1.1        0   FULL/-  00:00:19    6             Serial0/0/0
R2#
```

Secure Routing Updates

- When neighbor authentication has been configured on a router, the router authenticates the source of each routing update packet that it receives.
- An authenticating key that is known to both the sending and the receiving route is exchanged.
- OSPF supports three types of authentication:
  - **Null** – no authentication.
  - **Simple password authentication** – the password in the update is sent in plaintext over the network (outdated method).
  - **MD5 authentication** – Most secure and recommended method of authentication. Password is calculated using the MD5 algorithm.

Secure OSPF

MD5 Authentication

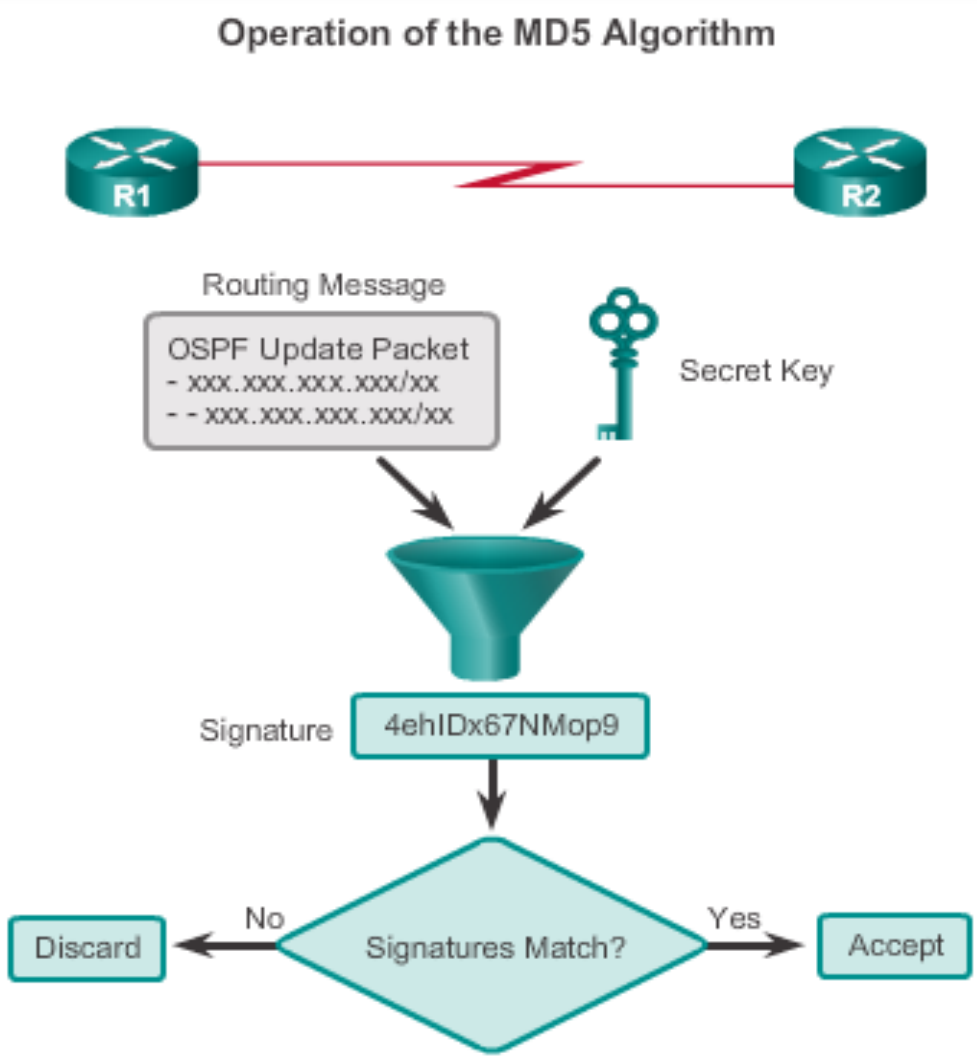

Secure OSPF

Configuring OSPF MD5 Authentication

- MD5 authentication can be enabled globally for all interfaces or on a per-interface basis.
- To enable OSPF MD5 authentication globally, configure:
  - **`ip ospf message-digest-key`** *key* **`md5`** *password* (interface configuration command)
  - **`area`** *area-id* **`authentication message-digest`** (router configuration command)
- To enable MD5 authentication on a per-interface basis, configure:
  - **`ip ospf message-digest-key`** *key* **`md5`** *password* (interface configuration command)
  - **`ip ospf authentication message-digest`** (interface configuration command)

Secure OSPF

OSPF MD5 Authentication Example

```
R1(config)# router ospf 10
R1(config-router)# area 0 authentication message-digest
R1(config-router)# exit
R1(config)#
*Apr  8 09:58:09.899: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2
on Serial0/0/0 from FULL to DOWN, Neighbor Down: Dead timer
expired
R1(config)#
*Apr  8 09:58:28.627: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3
on Serial0/0/1 from FULL to DOWN, Neighbor Down: Dead timer
expired
R1(config)#
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ip ospf message-digest-key 1 md5 CISCO-123
R1(config-if)# exit
R1(config)#
R1(config)# interface Serial 0/0/0
R1(config-if)# ip ospf message-digest-key 1 md5 CISCO-123
R1(config-if)# exit
R1(config)#
R1(config)# interface Serial 0/0/1
R1(config-if)# ip ospf message-digest-key 1 md5 CISCO-123
R1(config-if)#
```

continued

OSPF MD5 Authentication Example (cont.)

```
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ip ospf message-digest-key 1 md5 CISCO-123
R1(config-if)# ip ospf authentication message-digest
R1(config-if)# exit
R1(config)#
R1(config)# interface Serial 0/0/0
R1(config-if)# ip ospf message-digest-key 1 md5 CISCO-123
R1(config-if)# ip ospf authentication message-digest
R1(config-if)# exit
R1(config)#
R1(config)# interface Serial 0/0/1
R1(config-if)# ip ospf message-digest-key 1 md5 CISCO-123
R1(config-if)# ip ospf authentication message-digest
R1(config-if)# exit
R1(config)#
*Apr  8 10:20:10.647: %OSPF-5-ADJCHG: Process 10, Nbr 2.2.2.2
on Serial0/0/0 from FULL to DOWN, Neighbor Down: Dead timer
expired
R1(config)#
*Apr  8 10:20:50.007: %OSPF-5-ADJCHG: Process 10, Nbr 3.3.3.3
on Serial0/0/1 from FULL to DOWN, Neighbor Down: Dead timer
expired
R1(config)#
```

Secure OSPF

Verifying OSPF MD5 Authentication

```
R1# show ip ospf interface serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Internet Address 172.16.3.1/30, Area 0, Attached via
Network Statement
  Process ID 10, Router ID 1.1.1.1, Network Type
POINT_TO_POINT, Cost: 64
Topology-MTID   Cost  Disabled  Shutdown      Topology Name
      0          64      no        no             Base
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 5, Dead 20, Wait 20,
Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:02
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 2.2.2.2
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    Youngest key id is 1
R1#
R1# show ip ospf interface | include Message
  Message digest authentication enabled
  Message digest authentication enabled
  Message digest authentication enabled
R1#
```

Secure OSPF

Verifying OSPF MD5 Authentication (cont.)

Verify the Routing Table on R1

```
R1# show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP,
       M - mobile, B - BGP, D - EIGRP,
       EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1
       E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1,
       L2 - IS-IS level-2, ia - IS-IS inter area,
       * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route
       H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 172.16.3.2 to network 0.0.0.0

O*E2  0.0.0.0/0 [110/1] via 172.16.3.2, 00:33:17, Serial0/0/0
      172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
O        172.16.2.0/24 [110/65] via 172.16.3.2, 00:33:17, Serial0/0/0
O     192.168.1.0/24 [110/65] via 192.168.10.6, 00:30:43, Serial0/0/1
      192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
O        192.168.10.8/30 [110/128] via 192.168.10.6, 00:30:43, Serial0/0/1
                         [110/128] via 172.16.3.2, 00:33:17, Serial0/0/0
R1#
```

OSPF
Trivia

- Who designed MD5?

Multiarea OSPF Operation

Cisco | Networking Academy®
Mind Wide Open™

Single-Area OSPF

Single-area OSPF is useful in smaller networks. If an area becomes too big, the following issues must be addressed:

- Large routing table (no summarization by default)
- Large link-state database (LSDB)
- Frequent SPF algorithm calculations

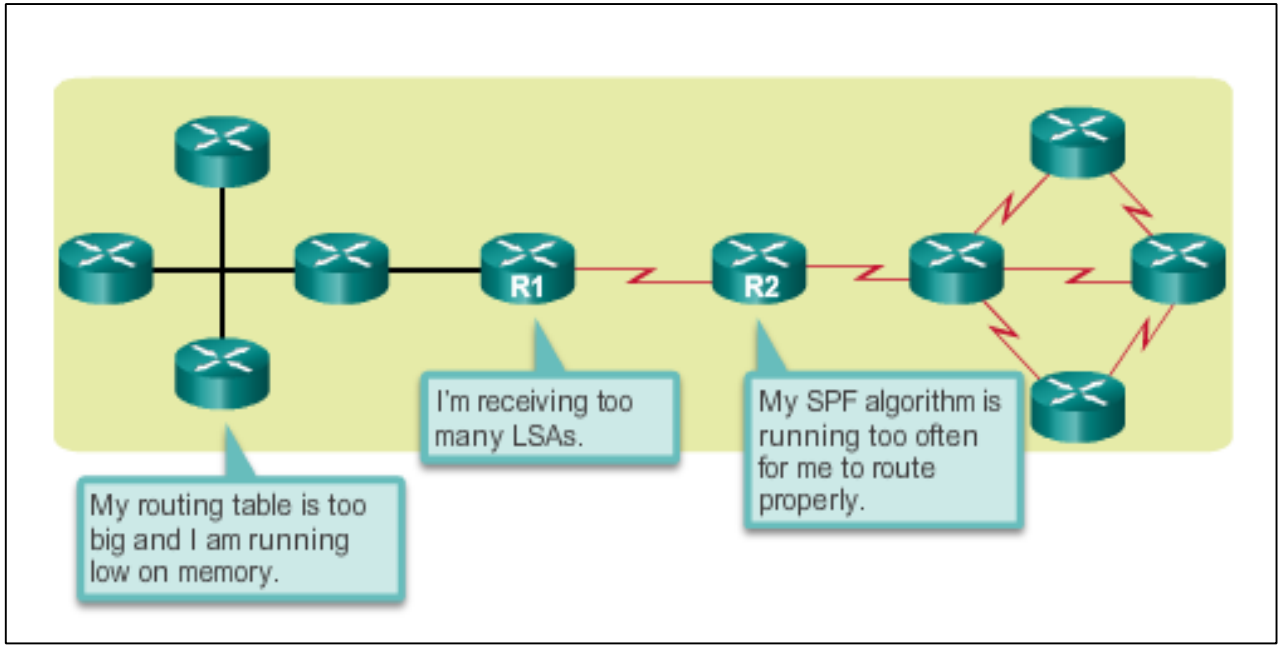

Why Multiarea OSPF?

Multiarea OSPF

Multiarea OSPF requires a hierarchical network design and the main area is called the backbone area, or area 0, and all other areas must connect to the backbone area.

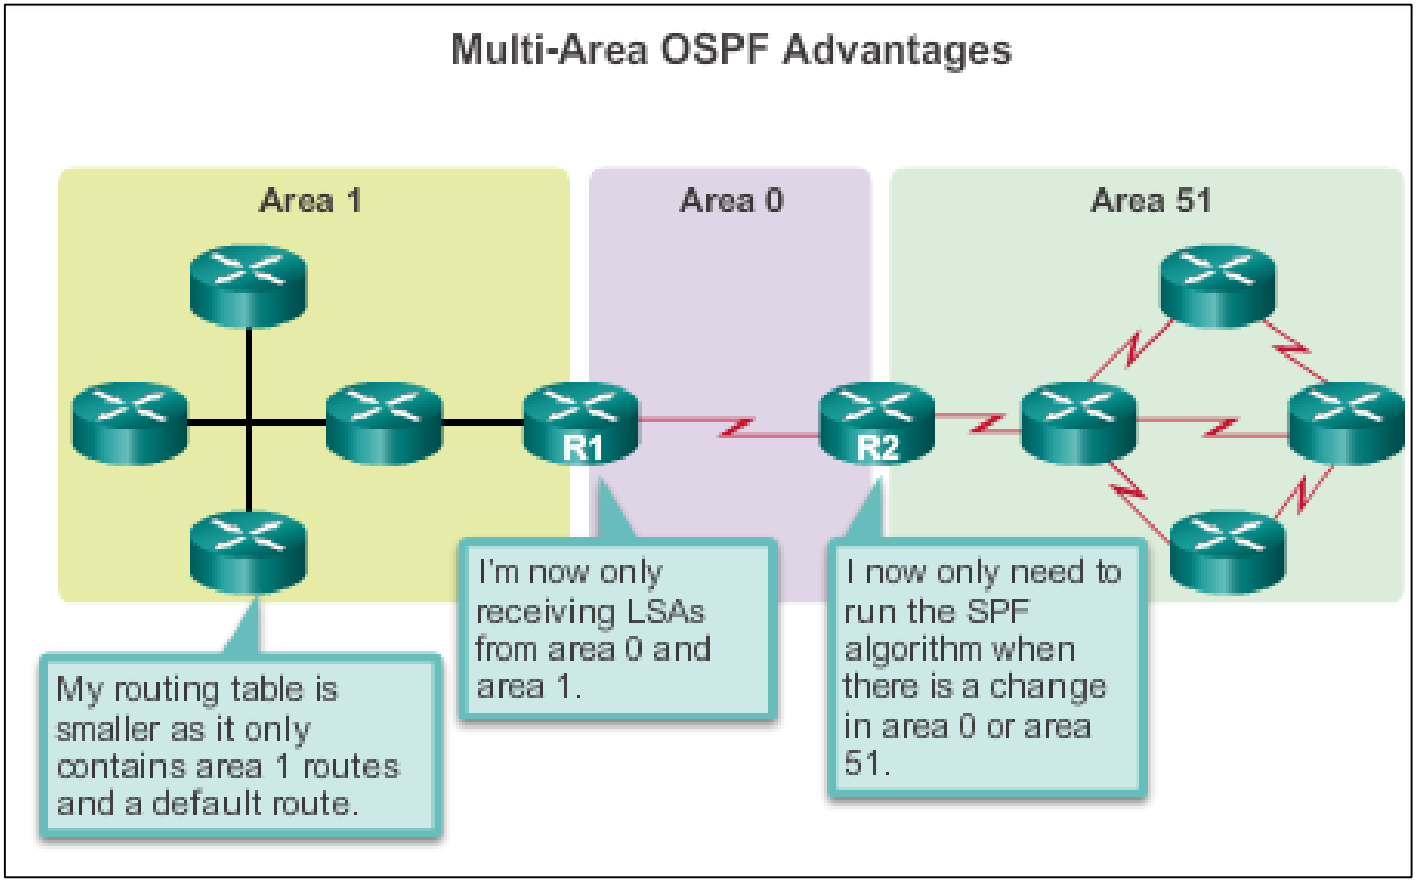

OSPF Two-Layer Area Hierarchy

Multiarea OSPF is implemented in a two-layer area hierarchy:

- **Backbone (transit) area**
  - Area whose primary function is the fast and efficient movement of IP packets.
  - Interconnects with other OSPF area types.
  - Called OSPF area 0, to which all other areas directly connect.
- **Regular (nonbackbone) area**
  - Connects users and resources.
  - A regular area does not allow traffic from another area to use its links to reach other areas.

Types of OSPF Routers

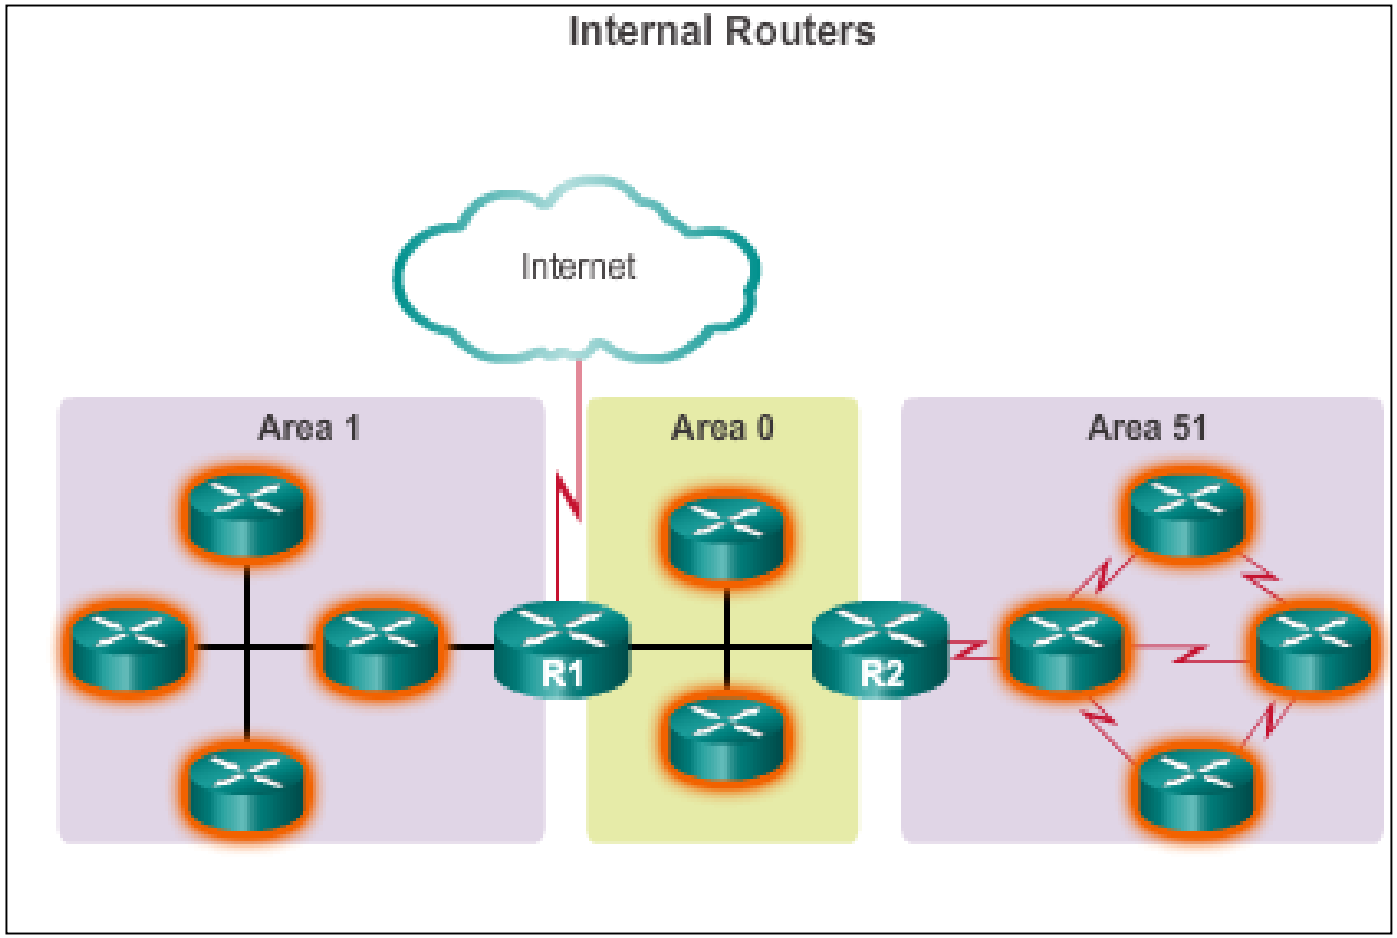

Why Multiarea OSPF?

Types of OSPF Routers (cont.)

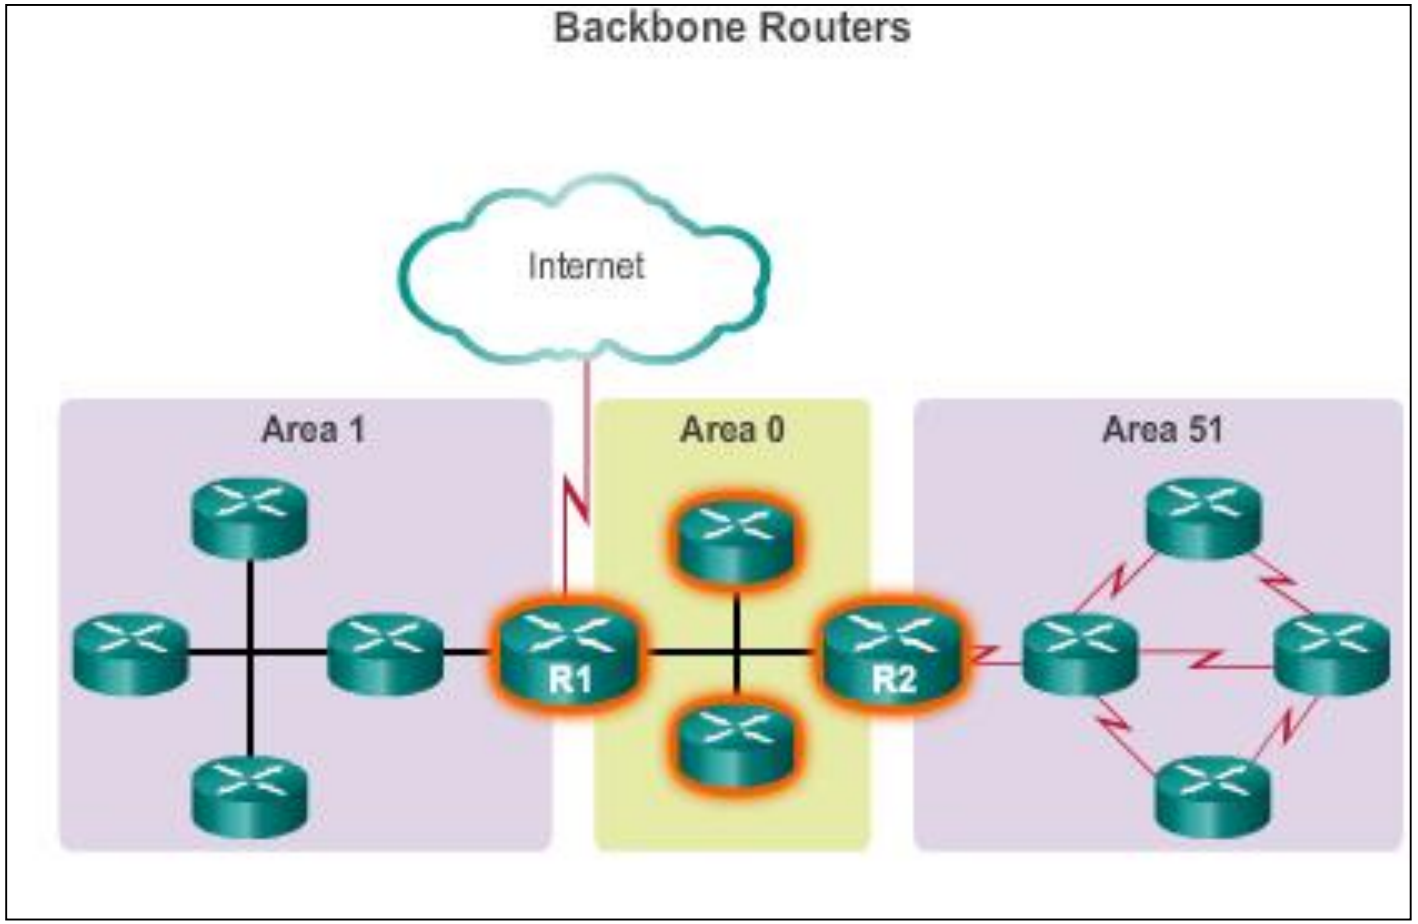

Why Multiarea OSPF?

Types of OSPF Routers (cont.)

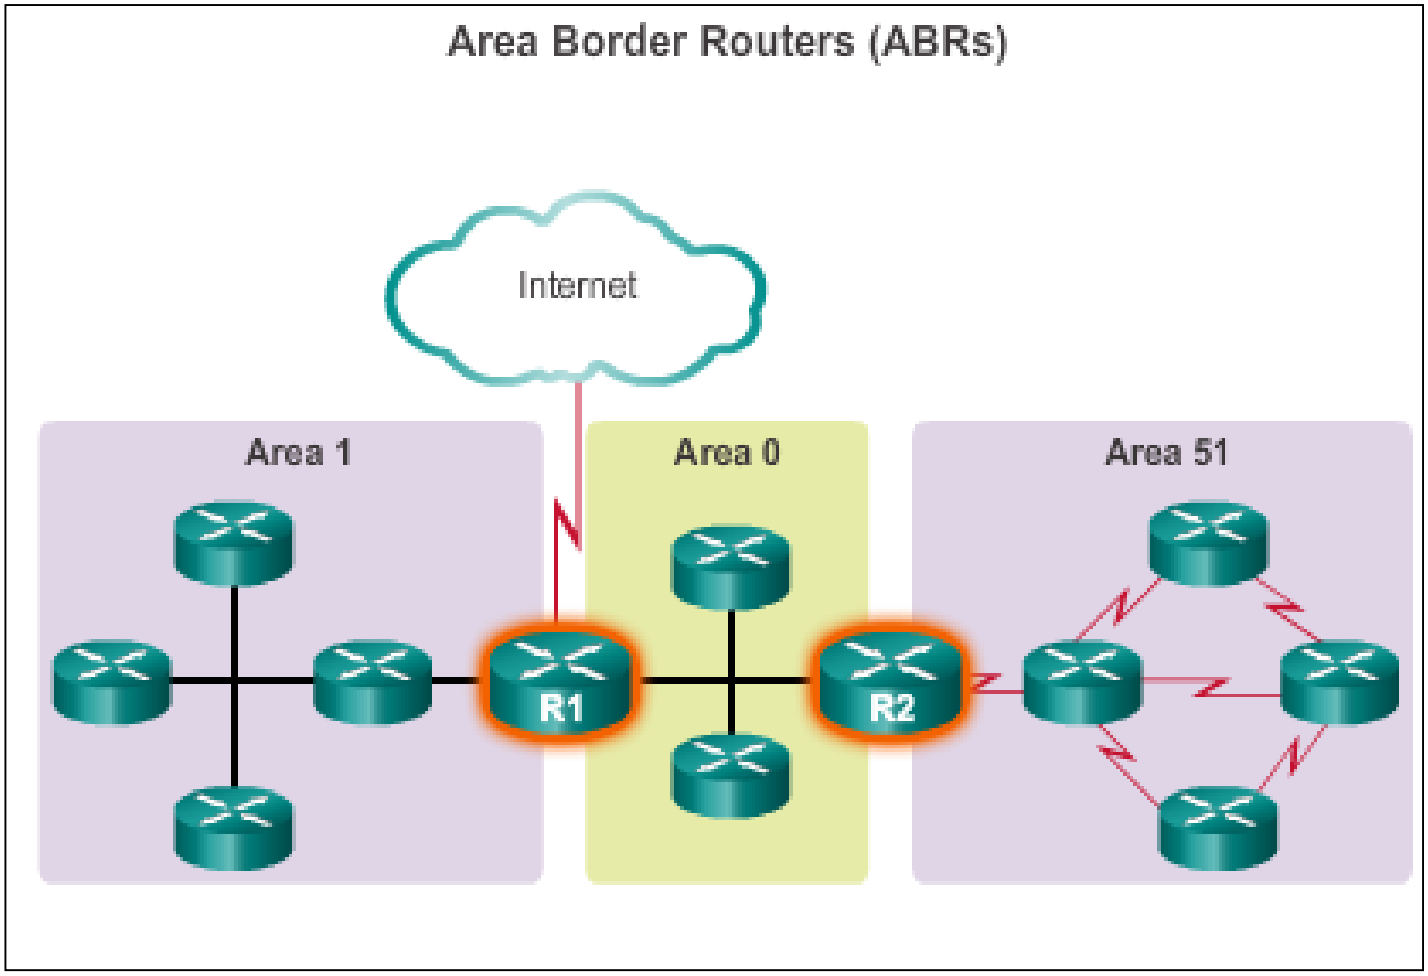

Why Multiarea OSPF?

Types of OSPF Routers (cont.)

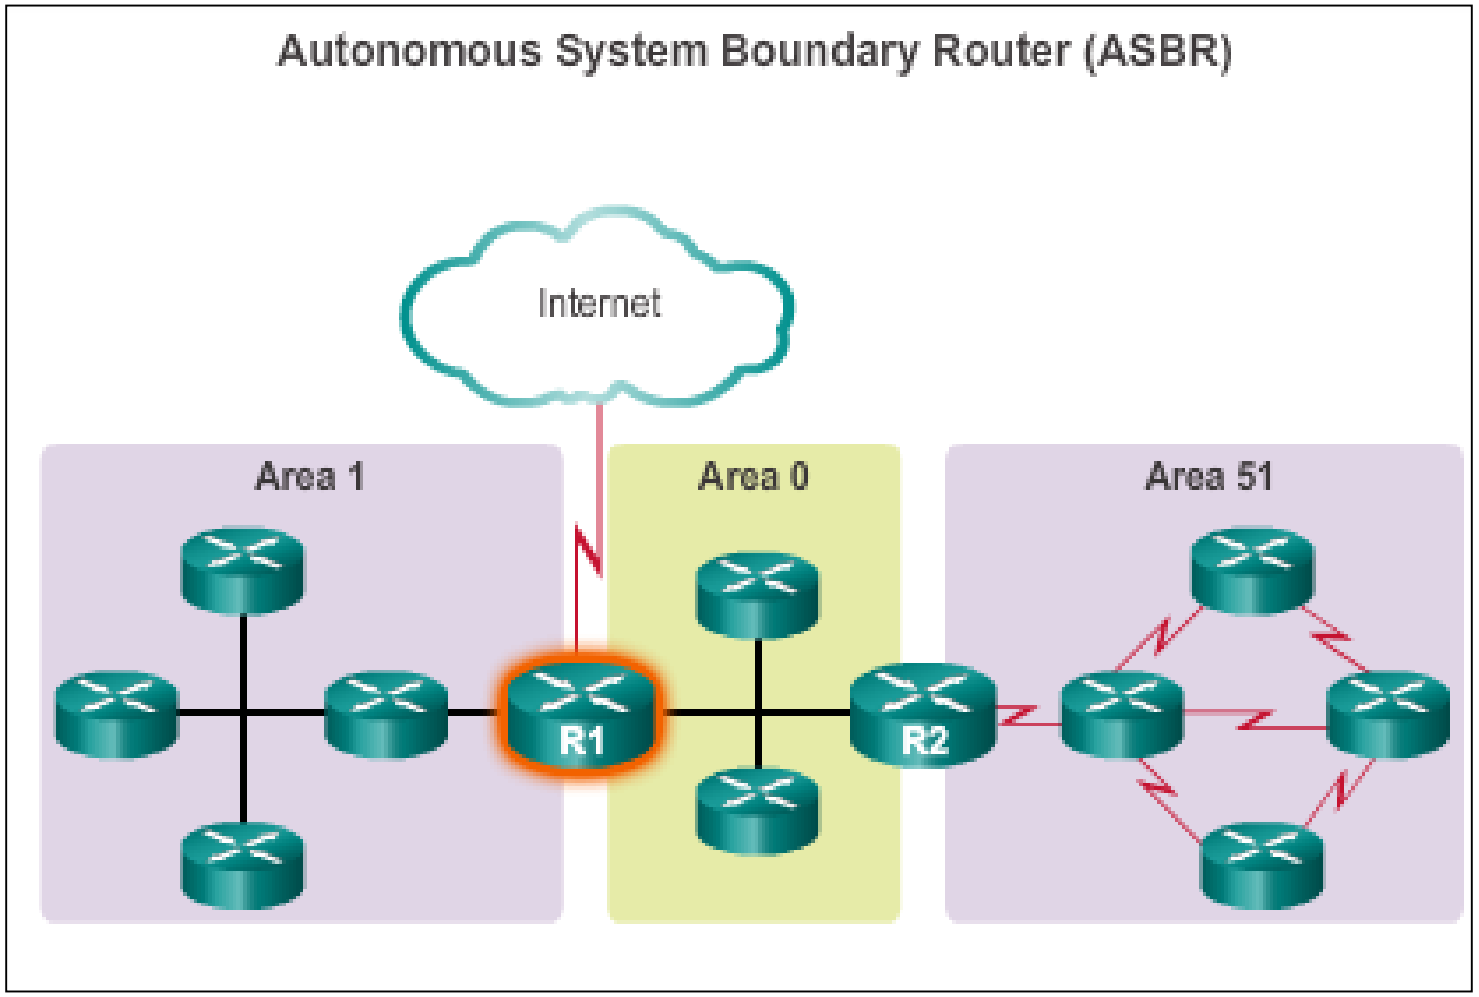

Multiarea OSPF LSA Operation

OSPF LSA Types

LSA Type	Description
1	Router LSA
2	Network LSA
3 and 4	Summary LSAs
5	AS External LSA
6	Multicast OSPF LSA
7	Defined for NSSAs
8	External Attributes LSA for Border Gateway Protocol (BGP)
9, 10, or 11	Opaque LSAs

Most common and covered in this course – 1 thru 5

Multiarea OSPF LSA Operation

OSPF LSA Type 1

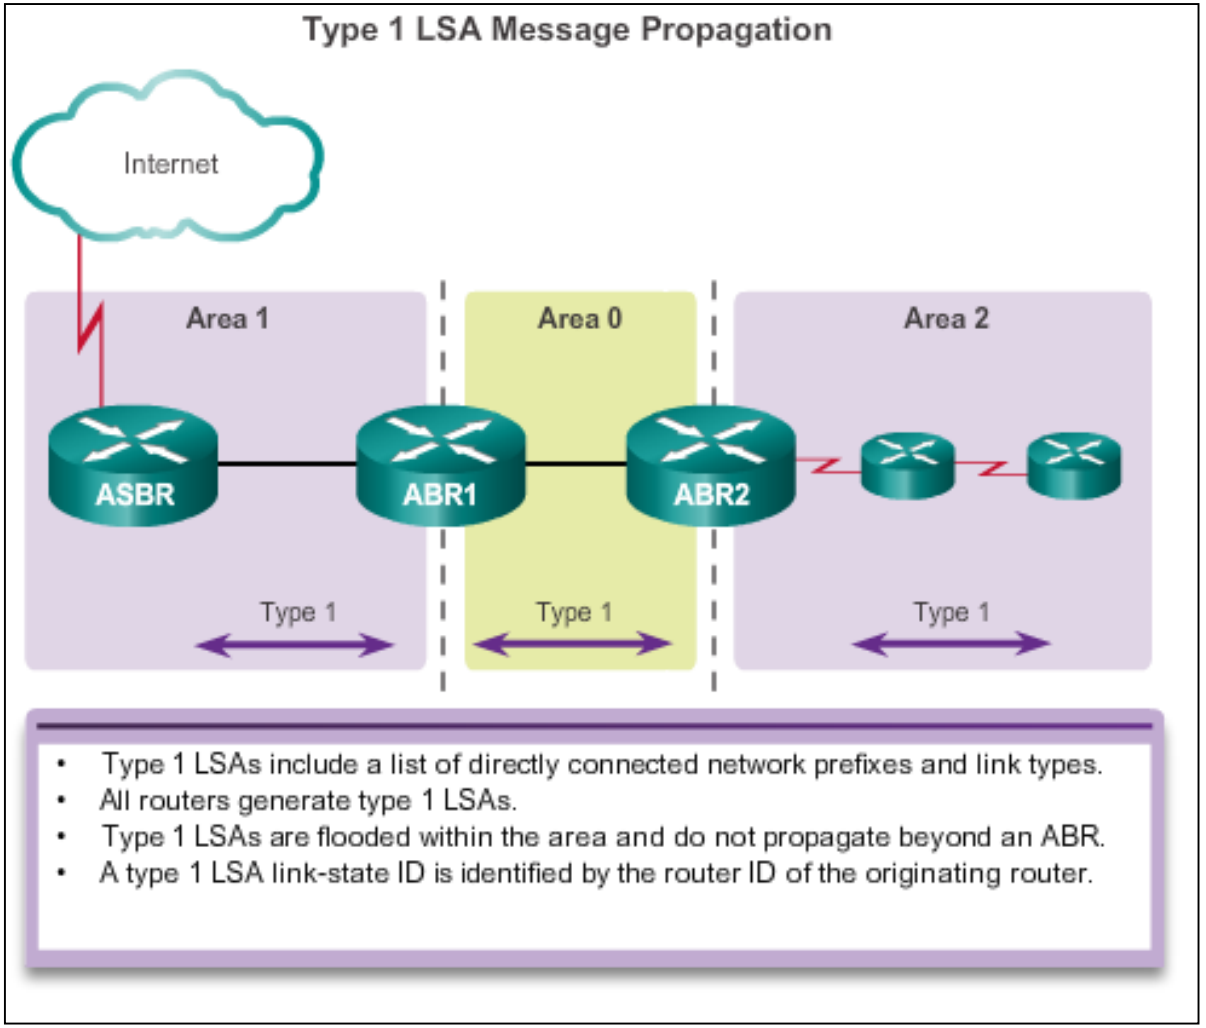

Multiarea OSPF LSA Operation

OSPF LSA Type 2

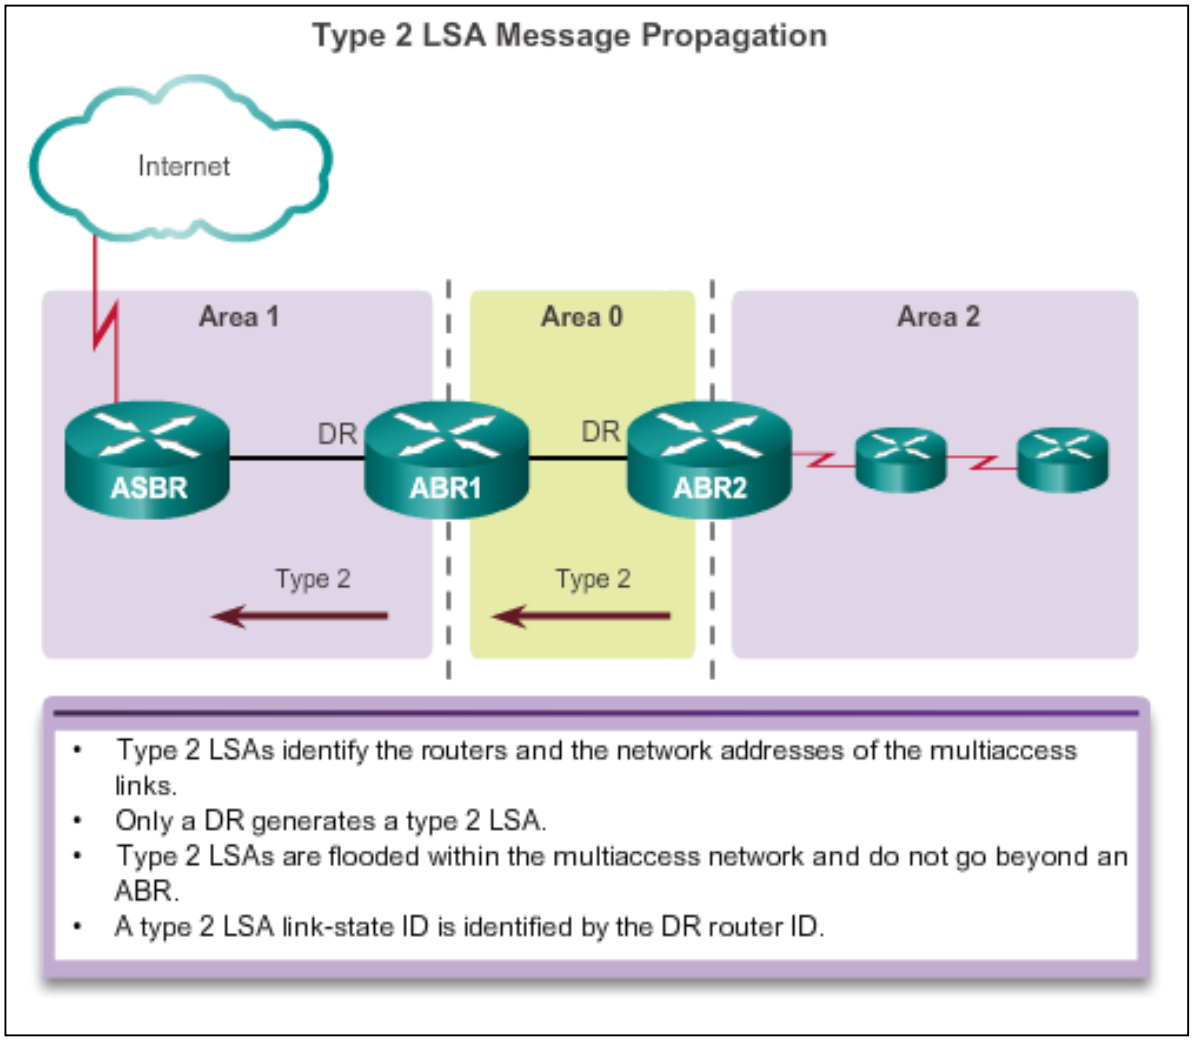

- Type 2 LSAs identify the routers and the network addresses of the multiaccess links.
- Only a DR generates a type 2 LSA.
- Type 2 LSAs are flooded within the multiaccess network and do not go beyond an ABR.
- A type 2 LSA link-state ID is identified by the DR router ID.

Multiarea OSPF LSA Operation

OSPF LSA Type 3

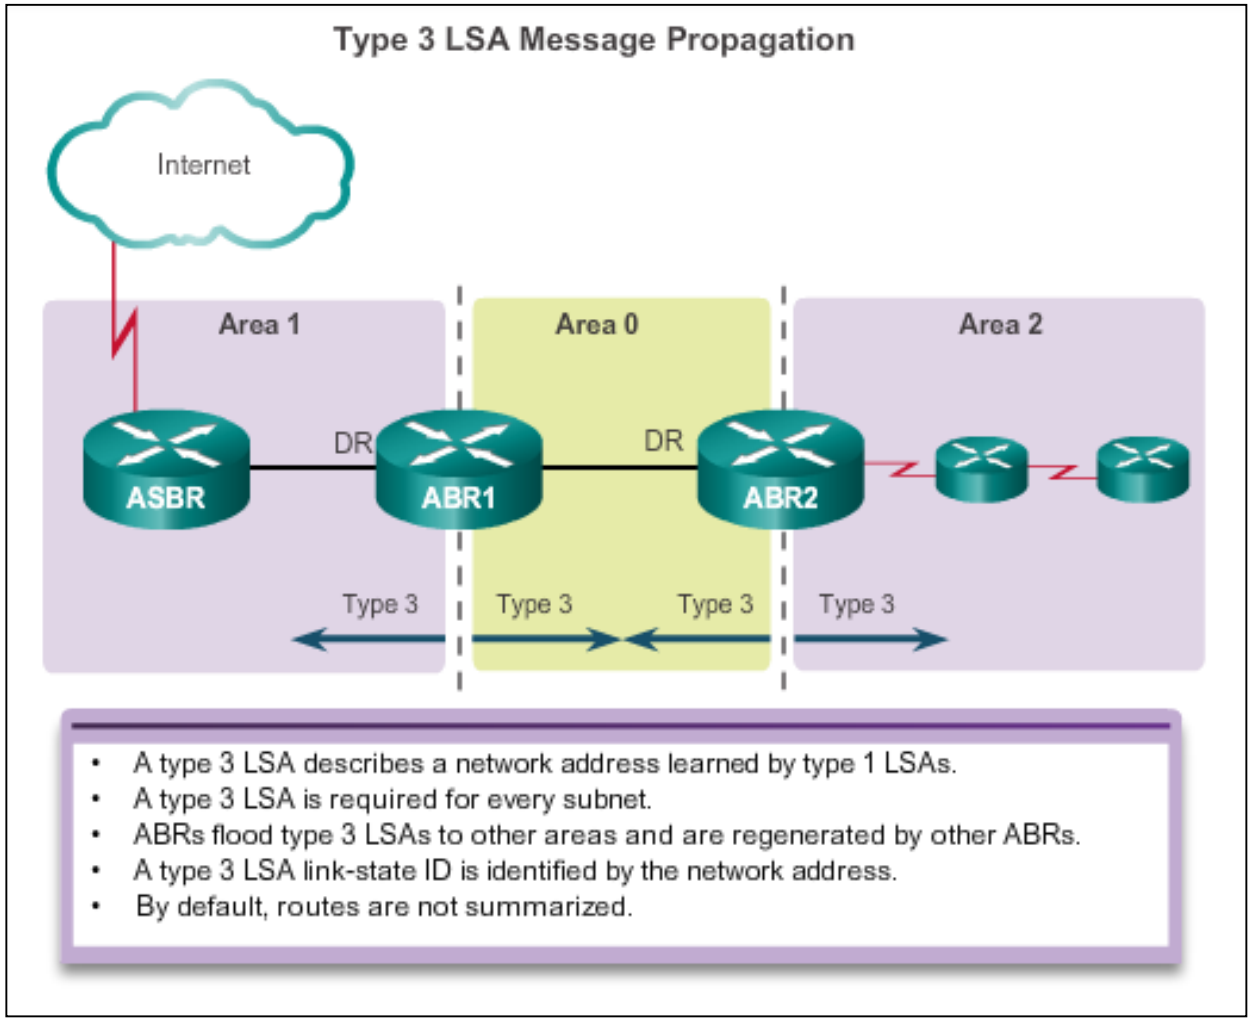

Multiarea OSPF LSA Operation

OSPF LSA Type 4

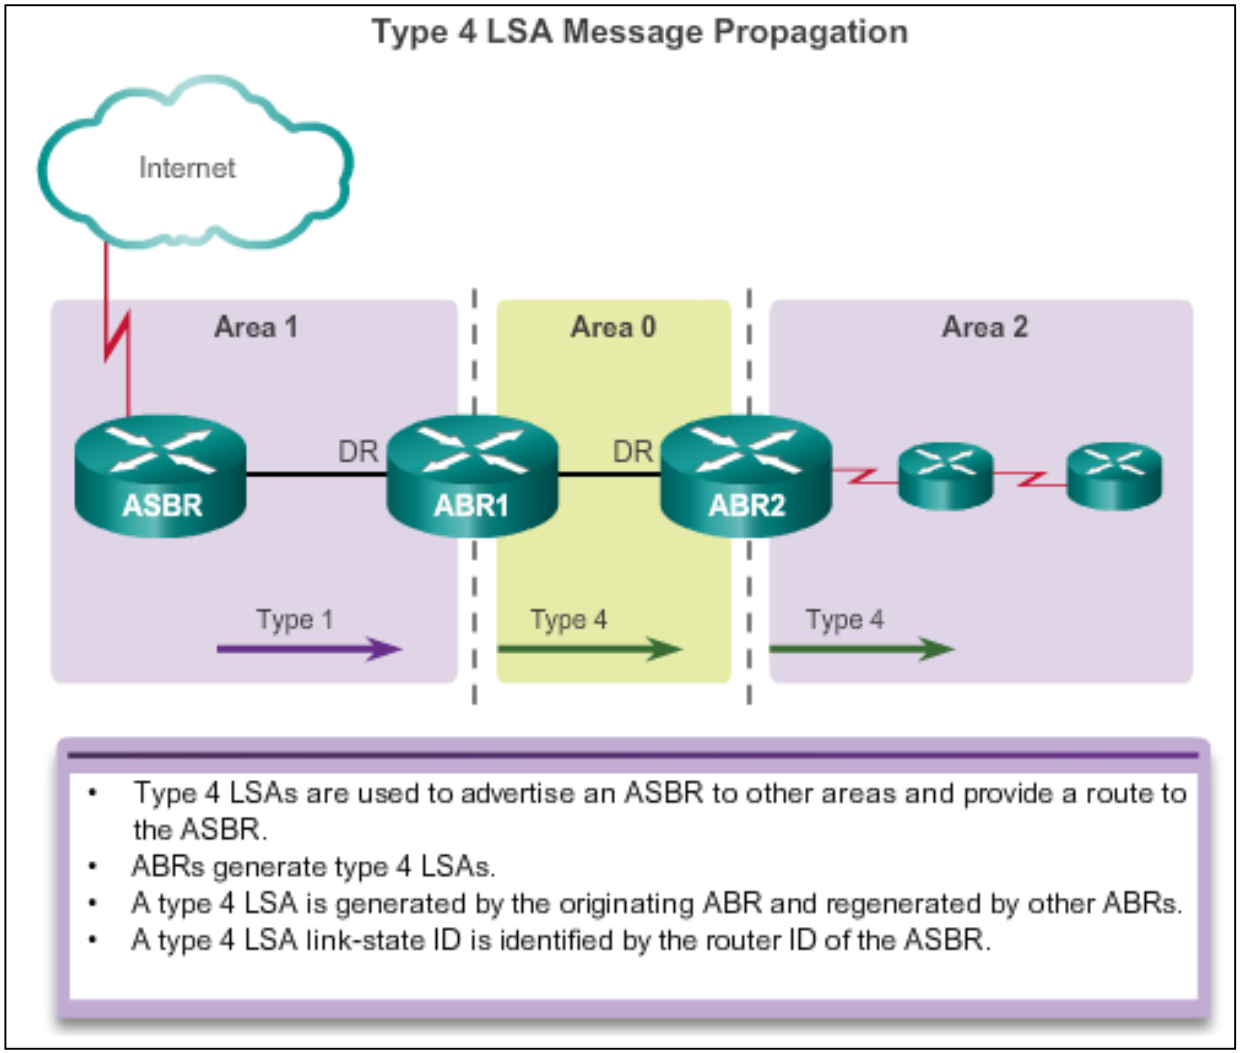

Multiarea OSPF LSA Operation

OSPF LSA Type 5

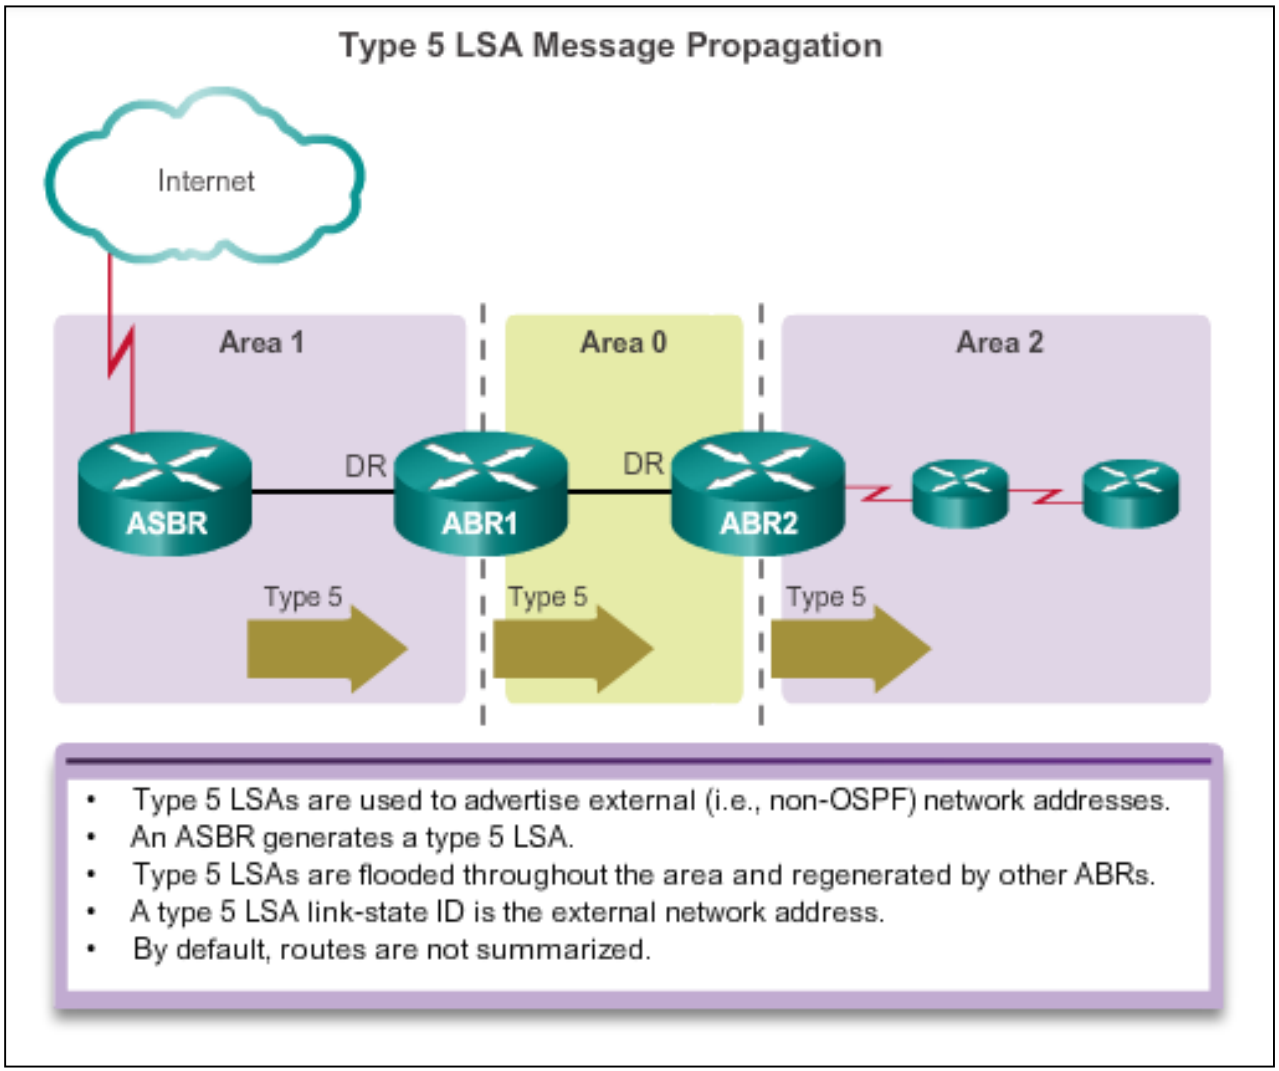

OSPF Routing Table Entries

- **O** – Router (type 1) and network (type 2) LSAs describe the details within an area (the route is intra-area).
- **O IA** – Summary LSAs appear in the routing table as IA (interarea routes)
- **O E1** or **OE 2** – External LSAs external type 1 (E1) or external type 2 (E2) routes

Router and Network Routing Table Entries

```
R1# show ip route
Codes:L - local, C-connected, S-static, R-RIP, M-mobile, B-BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su-IS-IS summary, L1-IS-IS level-1, L2-IS-IS level-2
      ia - IS-IS inter area,*-candidate default,U-per-user static route
      o - ODR, P-periodic downloaded static route, H-NHRP, l-LISP
      + - replicated route, % - next hop override

Gateway of last resort is 192.168.10.2 to network 0.0.0.0

O*E2 0.0.0.0/0 [110/1] via 192.168.10.2, 00:00:19, Serial0/0/0
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C       10.1.1.0/24 is directly connected, GigabitEthernet0/0
L       10.1.1.1/32 is directly connected, GigabitEthernet0/0
C       10.1.2.0/24 is directly connected, GigabitEthernet0/1
L       10.1.2.1/32 is directly connected, GigabitEthernet0/1
O       10.2.1.0/24 [110/648] via 192.168.10.2, 00:04:34, Serial0/0/0
O IA 192.168.1.0/24 [110/1295] via 192.168.10.2, 00:01:48,Serial0/0/0
O IA 192.168.2.0/24 [110/1295] via 192.168.10.2, 00:01:48,Serial0/0/0
     192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
C       192.168.10.0/30 is directly connected, Serial0/0/0
L       192.168.10.1/32 is directly connected, Serial0/0/0
O       192.168.10.4/30 [110/1294] via 192.168.10.2, 00:01:55,Serial0/0/0
R1#
```

OSPF Routing Table Entries (cont.)

- **O** – Router (type 1) and network (type 2) LSAs describe the details within an area (the route is intra-area)
- **O IA** – Summary LSAs appear in the routing table as IA (interarea routes)
- **O E1** or **OE 2** – External LSAs external type 1 (E1) or external type 2 (E2) routes

OSPFv3 Routing Table Entries

```
R1# show ipv6 route
IPv6 Routing Table - default - 9 entries
Codes:C - Connected, L - Local, S - Static, U-Per-user Static route
      B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
      I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
      EX - EIGRP external, ND-ND Default,NDp-ND Prefix,DCE-Destination
      NDr - Redirect, O-OSPF Intra, OI-OSPF Inter, OE1-OSPF ext 1
      OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
OE2 ::/0 [110/1], tag 10
     via FE80::2, Serial0/0/0
C   2001:DB8:CAFE:1::/64 [0/0]
     via GigabitEthernet0/0, directly connected
L   2001:DB8:CAFE:1::1/128 [0/0]
     via GigabitEthernet0/0, receive
O   2001:DB8:CAFE:2::/64 [110/648]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:3::/64 [110/1295]
     via FE80::2, Serial0/0/0
C   2001:DB8:CAFE:A001::/64 [0/0]
     via Serial0/0/0, directly connected
L   2001:DB8:CAFE:A001::1/128 [0/0]
     via Serial0/0/0, receive
O   2001:DB8:CAFE:A002::/64 [110/1294]
     via FE80::2, Serial0/0/0
L   FF00::/8 [0/0]
     via Null0, receive
R1#
```

OSPF Route Calculation

1. All routers calculate the best paths to destinations within their area (intra-area) and add these entries to the routing table.
2. All routers calculate the best paths to the other areas within the internetwork (interarea) or type 3 and type 4 LSAs.
3. All routers calculate the best paths to the external autonomous system (type 5) destinations. These are noted with either an O E1 or an O E2 route designator.

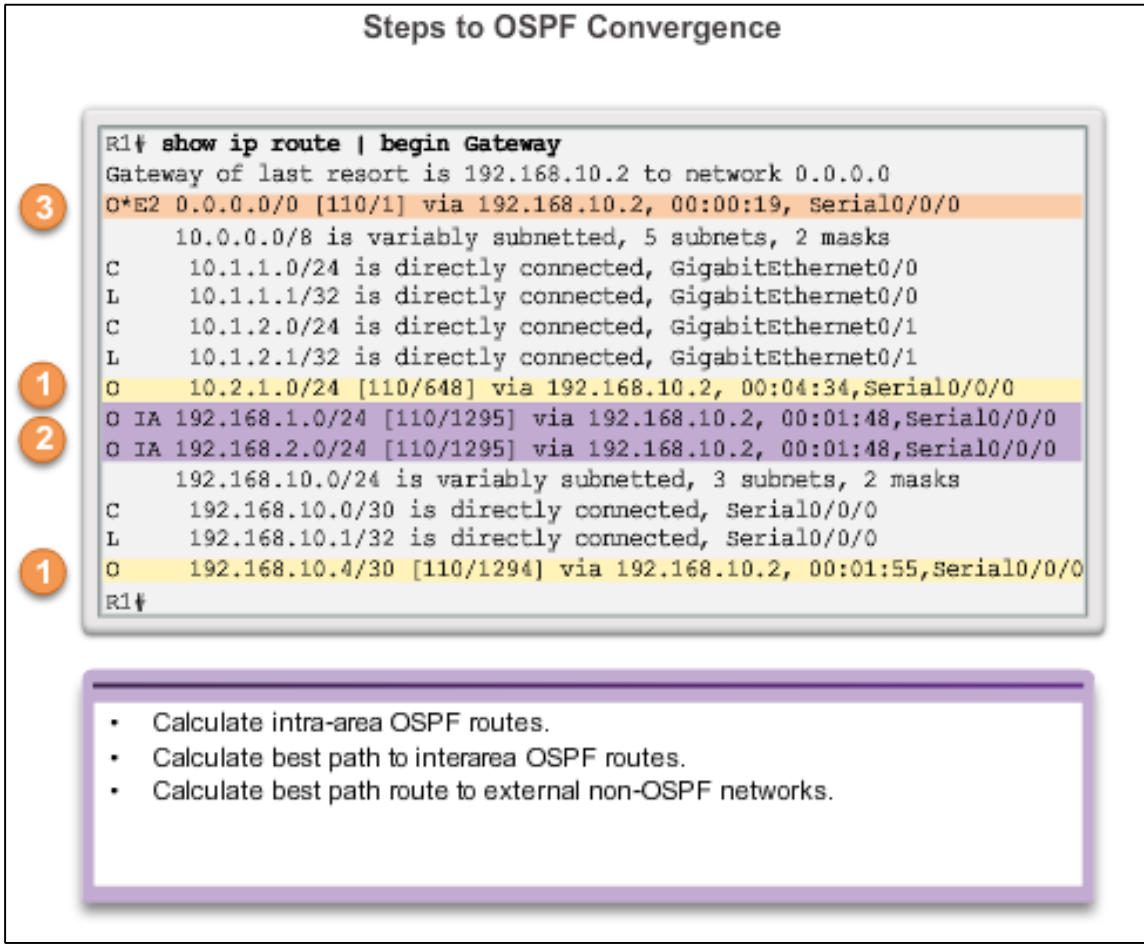

OSPF
Trivia

- How many bits are used to create Areas?

Configuring Multiarea OSPF

Cisco | Networking Academy®
Mind Wide Open™

Implementing Multiarea OSPF

Implementation Plan Steps

1. Gather the network requirements and parameters.
2. Define the OSPF parameters.
3. Configure OSPF.
4. Verify OSPF.

Configuring Multiarea OSPF

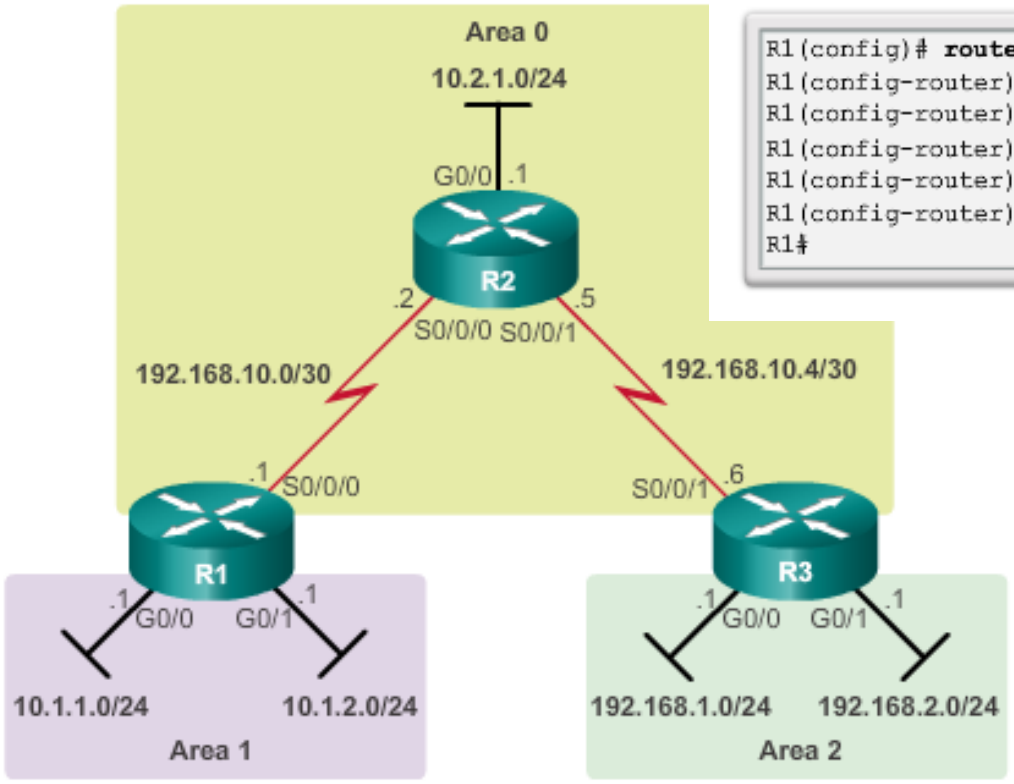


```
R1(config)# router ospf 10
R1(config-router)# router-id 1.1.1.1
R1(config-router)# network 10.1.1.1 0.0.0.0 area 1
R1(config-router)# network 10.1.2.1 0.0.0.0 area 1
R1(config-router)# network 192.168.10.1 0.0.0.0 area 0
R1(config-router)# end
R1#
```

Configuring Multiarea OSPFv3

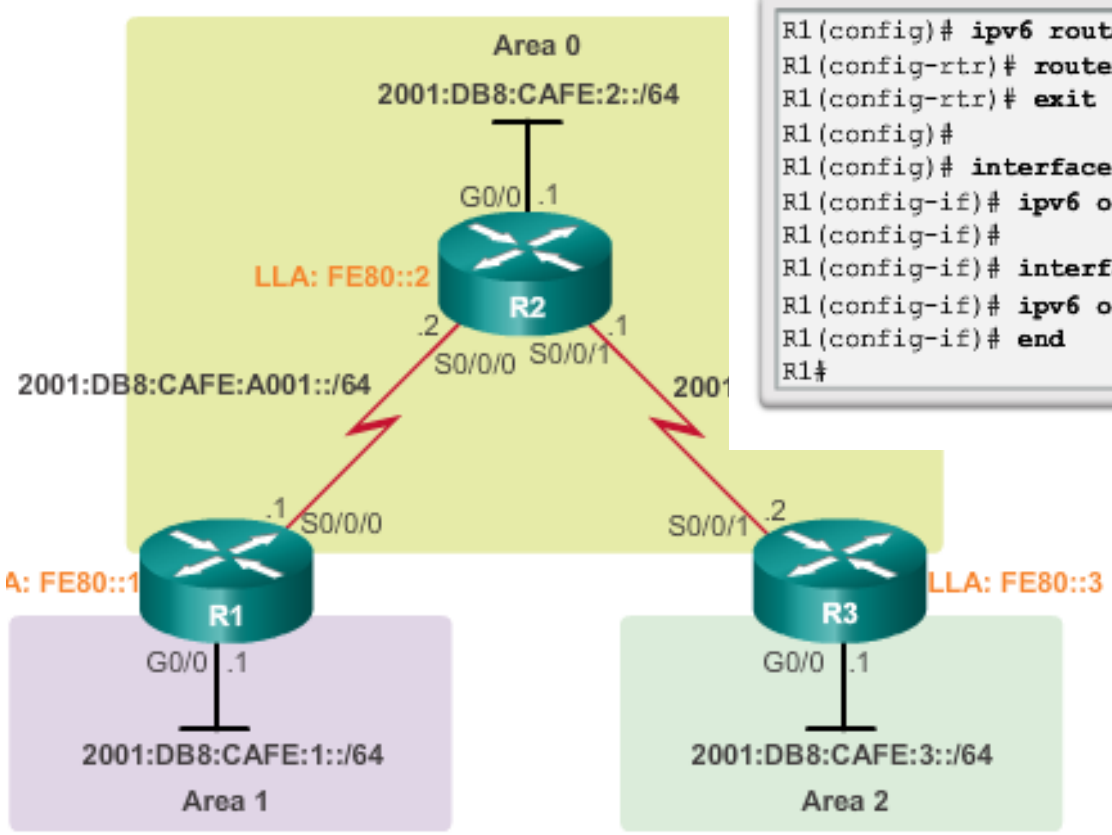


```
R1(config)# ipv6 router ospf 10
R1(config-rtr)# router-id 1.1.1.1
R1(config-rtr)# exit
R1(config)#
R1(config)# interface GigabitEthernet 0/0
R1(config-if)# ipv6 ospf 10 area 1
R1(config-if)#
R1(config-if)# interface Serial0/0/0
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# end
R1#
```

OSPF Route Summarization

- R1 forwards a summary LSA to the core router C1.
- C1, in turn, forwards the summary LSA to R2 and R3.
- R2 and R3 then forward it to their respective internal routers.

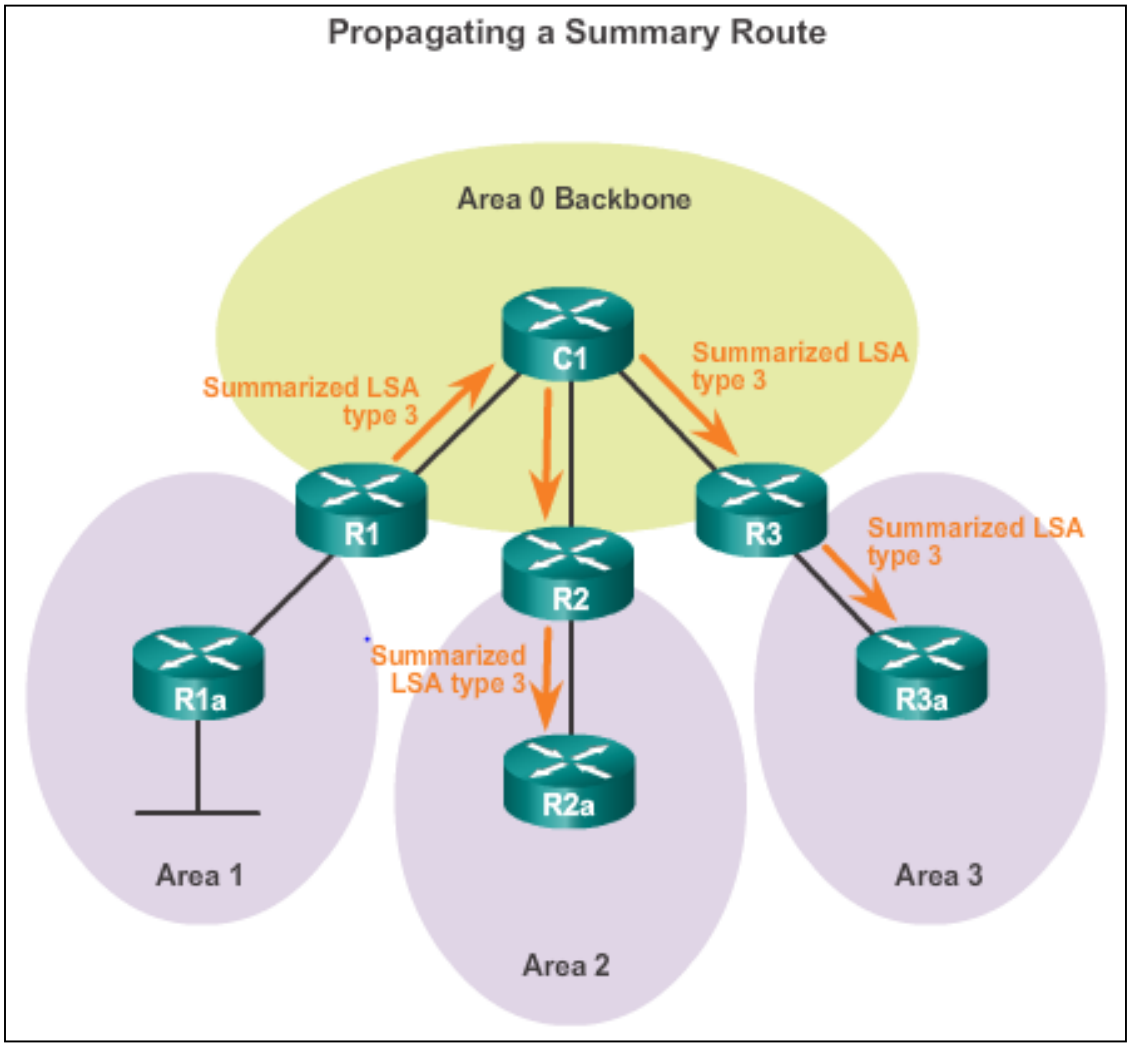

OSPF Route Summarization

Interarea and External Route Summarization

Occurs on ABRs and applies to routes from within each area

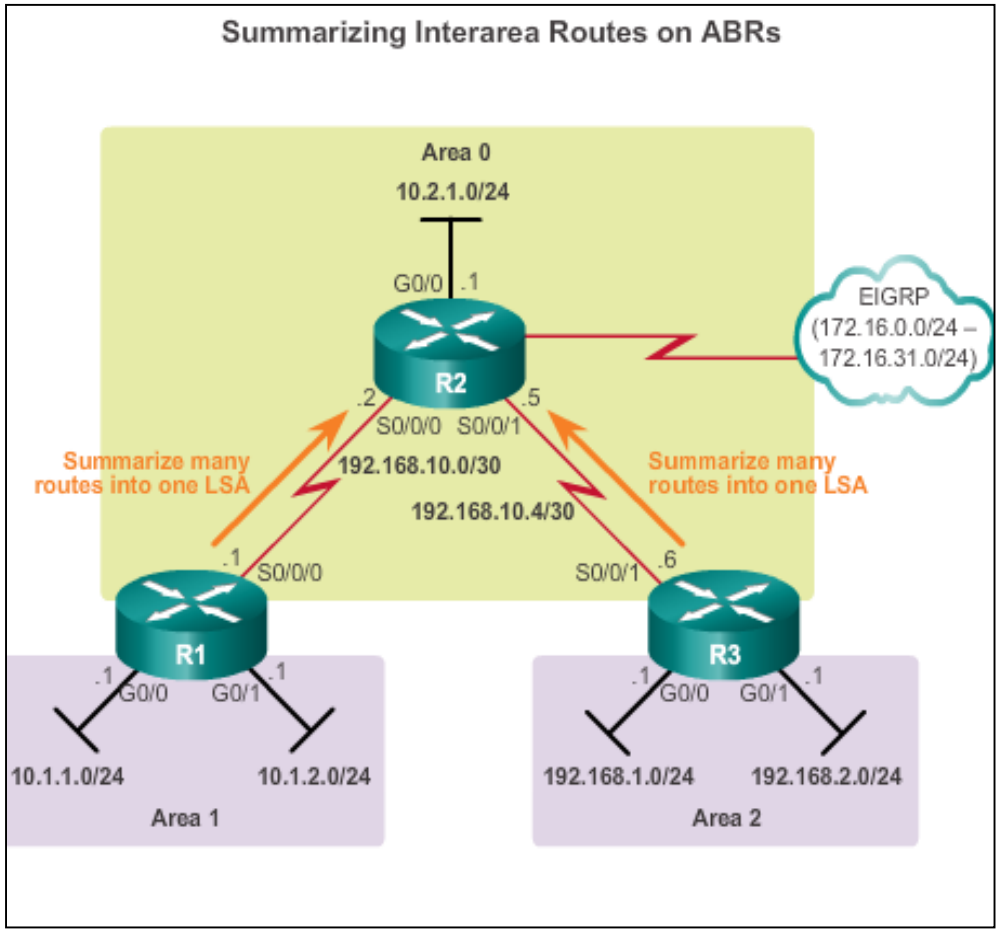

Interarea and External Route Summarization (cont.)

Specific to external routes that are injected into OSPF via route redistribution; ASBRs summarize external routes

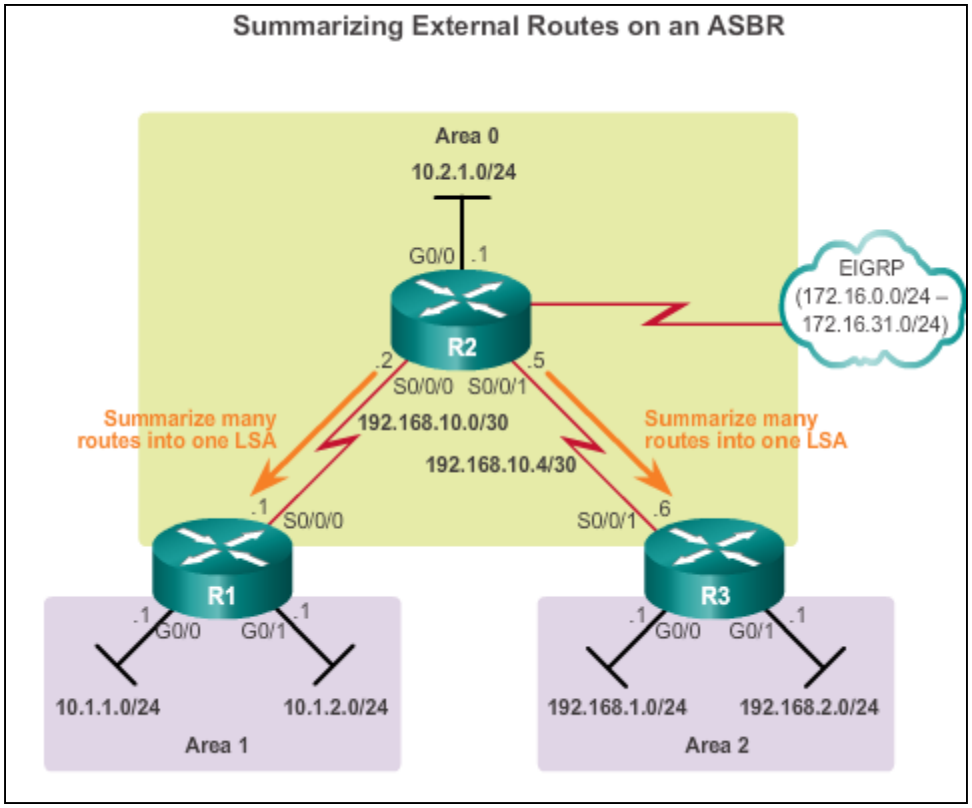

Interarea Route Summarization

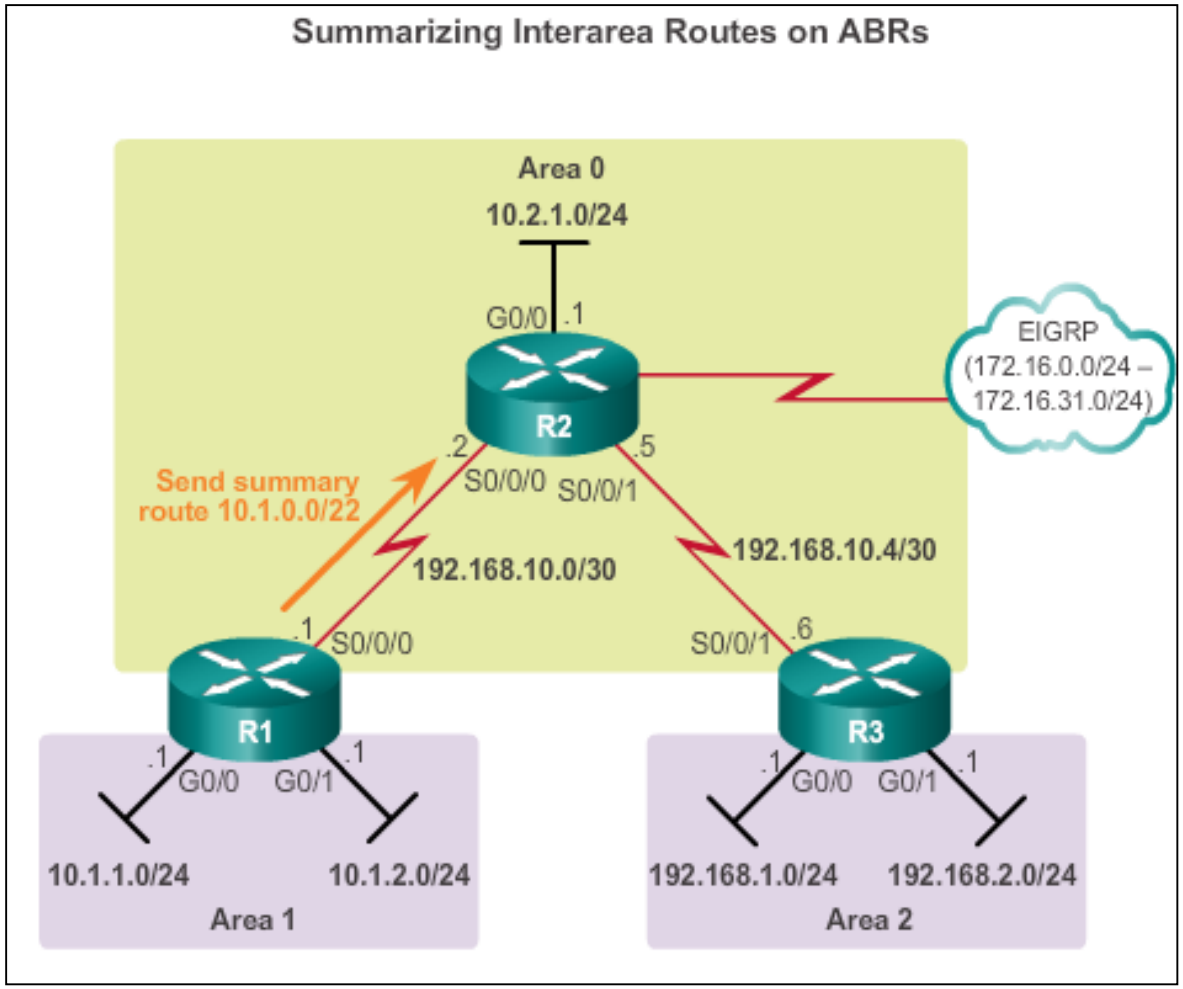

Interarea Route Summarization (cont.)

Verify the R1 Routing Table Before Summarization

```
R1# show ip route ospf | begin Gateway
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O        10.2.1.0/24 [110/648] via 192.168.10.2, 00:00:49,
         Serial0/0/0
O IA  192.168.1.0/24 [110/1295] via 192.168.10.2, 00:00:49,
      Serial0/0/0
O IA  192.168.2.0/24 [110/1295] via 192.168.10.2, 00:00:49,
      Serial0/0/0
      192.168.10.0/24 is variably subnetted, 3
      masks
O        192.168.10.4/30 [110/1294] via 192.16
         00:00:49, Serial0/0/0
R1#
```

Verify the R3 Routing Table Before Summarization

```
R3# show ip route ospf | begin Gateway
Gateway of last resort is not set

      10.0.0.0/24 is subnetted, 3 subnets
O IA     10.1.1.0 [110/1295] via 192.168.10.5, 00:27:14, Serial0/0/1
O IA     10.1.2.0 [110/1295] via 192.168.10.5, 00:27:14, Serial0/0/1
O        10.2.1.0 [110/648] via 192.168.10.5, 00:27:57,  Serial0/0/1
      192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks
O        192.168.10.0/30 [110/1294] via 192.168.10.5, 00:27:57,
         Serial0/0/1
R3#
```

OSPF Route Summarization

Calculating the Summary Route

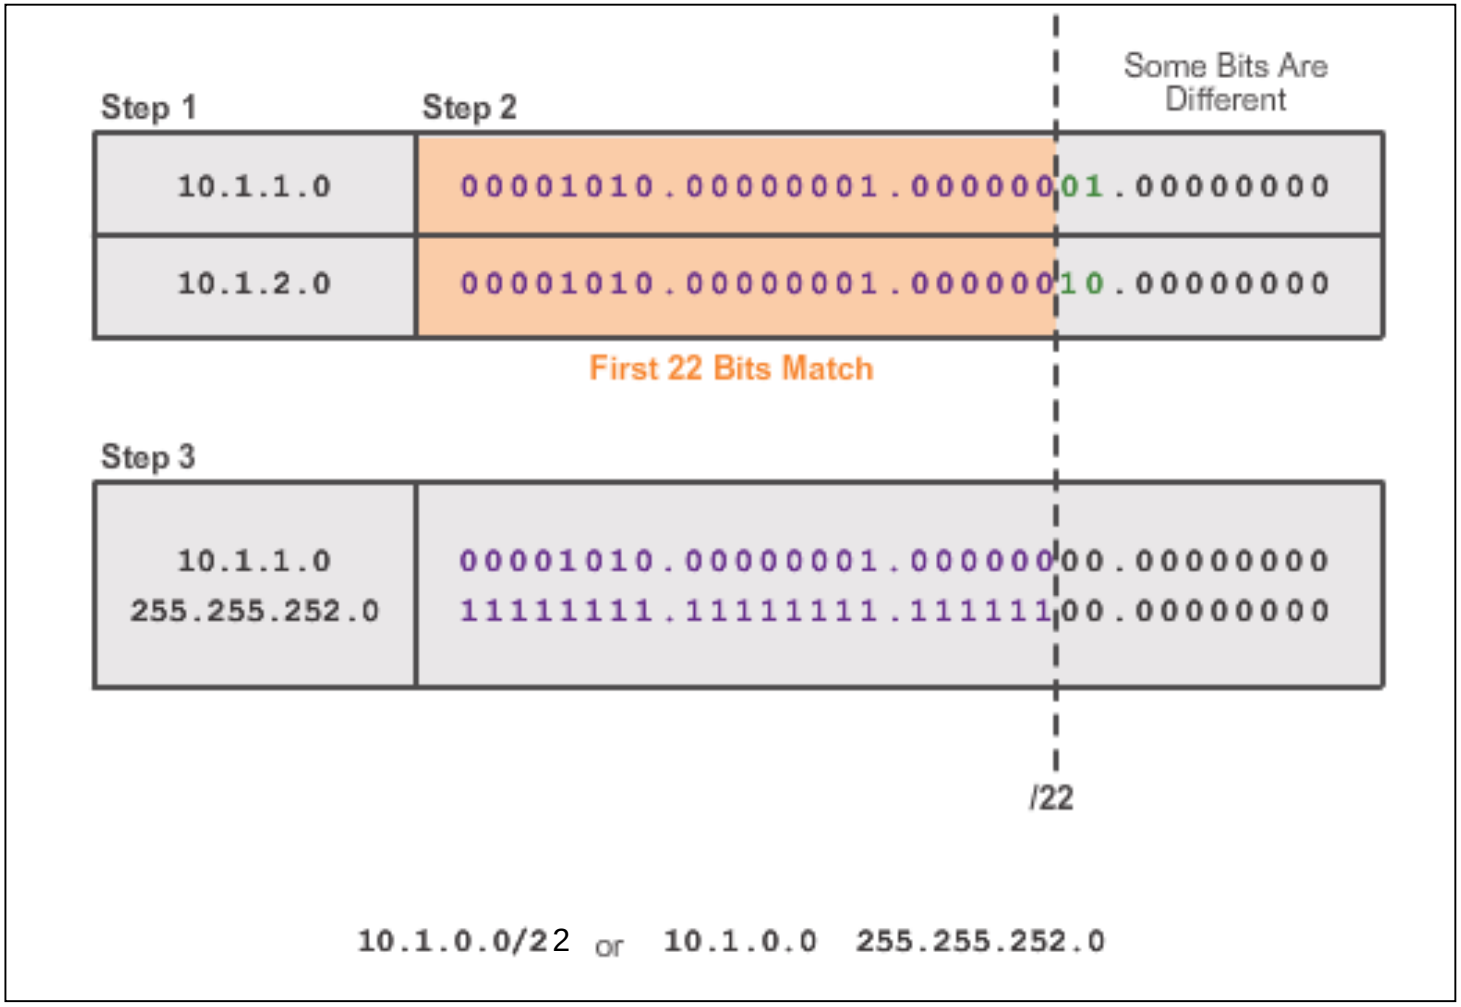

OSPF Route Summarization

Configuring Interarea Route Summarization

```
R1(config)# router ospf 10
R1(config-router)# area 1 range 10.1.0.0 255.255.252.0
R1(config-router)#
```

R1

```
R1# show ip route ospf | begin Gateway
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
O        10.1.0.0/22 is a summary, 00:00:09, Null0
O        10.2.1.0/24 [110/648] via 192.168.10.2, 00:00:09,
Serial0/0/0
O IA  192.168.1.0/24 [110/1295] via 192.168.10.2, 00:00:09,
Serial0/0/0
O IA  192.168.2.0/24 [110/1295] via 192.168.10.2, 00:0
Serial0/0/0
      192.168.10.0/24 is variably subnetted, 3 subnets
masks
O        192.168.10.4/30 [110/1294] via 192.168.10.2,
00:00:09, Serial0/0/0
R1#
```

R3

```
R3# show ip route ospf | begin Gateway
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
O IA     10.1.0.0/22 [110/1295] via 192.168.10.5, 00:00:06,
Serial0/0/1
O        10.2.1.0/24 [110/648] via 192.168.10.5, 00:29:23,
Serial0/0/1
      192.168.10.0/24 is variably subnetted, 3 subnets, 2
masks
O        192.168.10.0/30 [110/1294] via 192.168.10.5,
00:29:23, Serial0/0/1
R3#
```

Verifying Multiarea OSPF

The same verification commands are used to verify single-area OSPF and can be used to verify multiarea OSPF:

- **`show ip ospf neighbor`**
- **`show ip ospf`**
- **`show ip ospf interface`**

Commands specific to multiarea information include:

- **`show ip protocols`**
- **`show ip ospf interface brief`**
- **`show ip route ospf`**
- **`show ip ospf database`**

Note: For OSPFv3, substitute **`ip`** with **`ipv6`**.

Verifying Multiarea OSPF

Verifying General Multiarea OSPF Settings

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  It is an area border router
  Number of areas in this router is 2. 2 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.1.1.1 0.0.0.0 area 1
    10.1.2.1 0.0.0.0 area 1
    192.168.10.1 0.0.0.0 area 0
  Routing Information Sources:
    Gateway         Distance      Last Update
    3.3.3.3              110      02:20:36
    2.2.2.2              110      02:20:39
  Distance: (default is 110)

R1#
```

```
R1# show ip ospf interface brief
Interface  PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Se0/0/0    10    0               192.168.10.1/30    64    P2P   1/1
Gi0/1      10    1               10.1.2.1/24        1     DR    0/0
Gi0/0      10    1               10.1.1.1/24        1     DR    0/0
R1#
```

Verifying Multiarea OSPF

Verify the OSPF Routes

```
R1# show ip route ospf | begin Gateway
Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O       10.2.1.0/24 [110/648] via 192.168.10.2, 00:26:03,
                                                Serial0/0/0
O IA 192.168.1.0/24 [110/1295] via 192.168.10.2, 00:26:03,
                                                Serial0/0/0
O IA 192.168.2.0/24 [110/1295] via 192.168.10.2, 00:26:03,
                                                Serial0/0/0
     192.168.10.0/24 is variably subnetted,3 subnets,2 masks
O       192.168.10.4/30 [110/1294] via 192.168.10.2, 00:26:03,
                                                Serial0/0/0

R1#
```

Verifying Multiarea OSPF

Verifying the Multiarea OSPF LSDB

Verifying the OSPF LSDB on R1

```
R1# show ip ospf database
            OSPF Router with ID (1.1.1.1) (Process ID 10)

                Router Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum Link count
1.1.1.1         1.1.1.1         725         0x80000005 0x00F9B0 2
2.2.2.2         2.2.2.2         695         0x80000007 0x003DB1 5
3.3.3.3         3.3.3.3         681         0x80000005 0x00FF91 2
              Summary Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.0        1.1.1.1         725         0x80000006 0x00D155
10.1.2.0        1.1.1.1         725         0x80000005 0x00C85E
192.168.1.0     3.3.3.3         681         0x80000006 0x00724E
192.168.2.0     3.3.3.3         681         0x80000005 0x006957

                Router Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum Link count
1.1.1.1         1.1.1.1         725         0x80000006 0x007D7C 2
              Summary Net Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum
10.2.1.0        1.1.1.1         725         0x80000005 0x004A9C
192.168.1.0     1.1.1.1         725         0x80000005 0x00B593
192.168.2.0     1.1.1.1         725         0x80000005 0x00AA9D
192.168.10.0    1.1.1.1         725         0x80000005 0x00B3D0
192.168.10.4    1.1.1.1         725         0x80000005 0x000E32
R1#
```

Verifying Multiarea OSPFv3

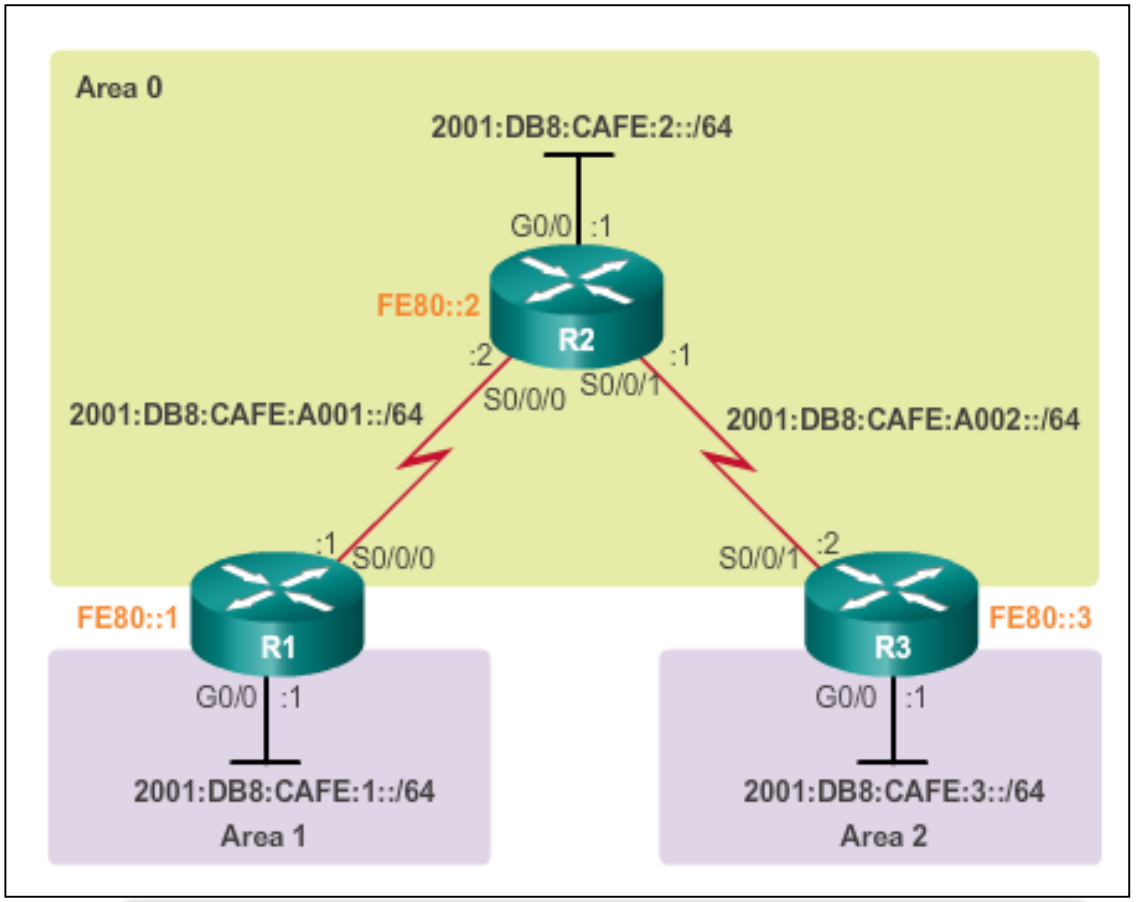


```
R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 1.1.1.1
  Area border router
  Number of areas: 2 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    Serial0/0/0
  Interfaces (Area 1):
    GigabitEthernet0/0
  Redistribution:
    None
R1#
```

```
R1# show ipv6 ospf interface brief
Interface    PID   Area            Intf ID    Cost  State Nbrs F/C
Se0/0/0      10    0               6          647   P2P   1/1
Gi0/0        10    1               3          1     DR    0/0
R1#
```

Verifying Multiarea OSPFv3 (cont.)

```
R1# show ipv6 route ospf
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static
route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D -
       EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE -
Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF
       ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
O   2001:DB8:CAFE:2::/64 [110/648]
     via FE80::2, Serial0/0/0
OI  2001:DB8:CAFE:3::/64 [110/1295]
     via FE80::2, Serial0/0/0
O   2001:DB8:CAFE:A002::/64 [110/1294]
     via FE80::2, Serial0/0/0
R1#
```

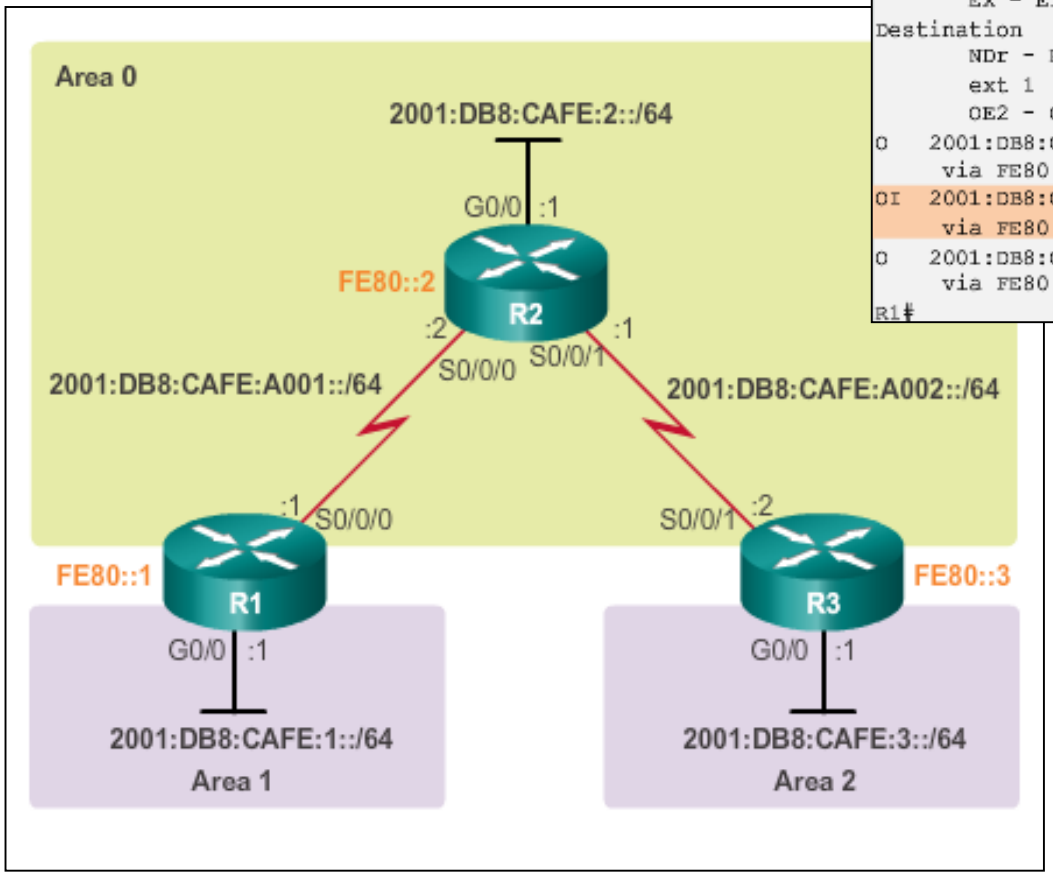

Verifying Multiarea OSPFv3 (cont.)

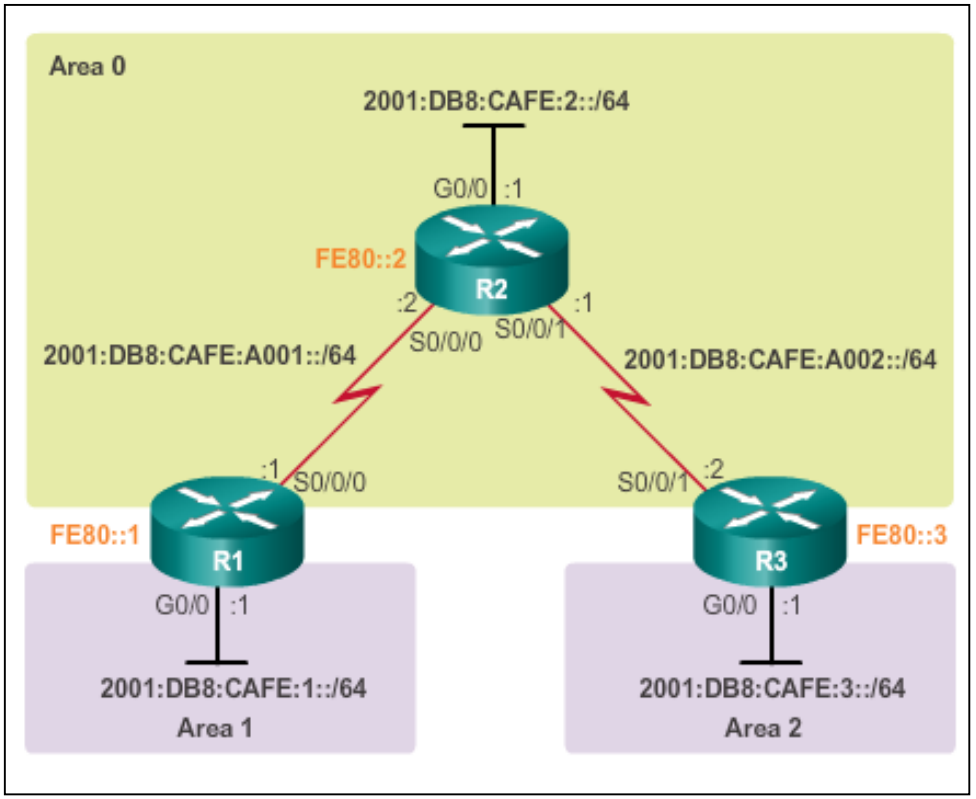


```
R1# show ipv6 ospf database

        OSPFv3 Router with ID (1.1.1.1) (Process ID 10)

                Router Link States (Area 0)

ADV Router    Age     Seq#       Fragment ID  Link count Bits
 1.1.1.1      1617    0x80000002 0            1          B
 2.2.2.2      1484    0x80000002 0            2          None
 3.3.3.3      14      0x80000001 0            1          B

             Inter Area Prefix Link States (Area 0)

ADV Router    Age      Seq#       Prefix
 1.1.1.1      1833     0x80000001 2001:DB8:CAFE:1::/64
 3.3.3.3      1476     0x80000001 2001:DB8:CAFE:3::/64

             Link (Type-8) Link States (Area 0)

ADV Router    Age      Seq#       Link ID    Interface
 1.1.1.1      1843     0x80000001 6          Se0/0/0
 2.2.2.2      1619     0x80000001 6          Se0/0/0

             Intra Area Prefix Link States (Area 0)
```

Trivia

- Did Chuck Norris in fact invent the internet?