Implementing a Scalable, Multiarea Network, OSPF-Based Solution

Interconnecting Cisco Networking Devices, Part 2 (ICND2) v2.0

© 2013 Cisco Systems, Inc.

OSPF Overview

Implementing a Scalable, Multiarea Network, OSPF-Based Solution

© 2013 Cisco Systems, Inc.

Link-State Routing Protocol Overview

Link-state routing protocols such as OSPF have several advantages when compared to distance vector routing protocols:

- Link-state protocols are more scalable.
- Each router has a full picture of a topology.
- Updates are sent when a topology change occurs and are reflooded periodically.
- Link-state protocols respond quickly to topology changes.
- More information is communicated between routers.

© 2013 Cisco Systems, Inc.

Link-State Routing Protocol Data Structures

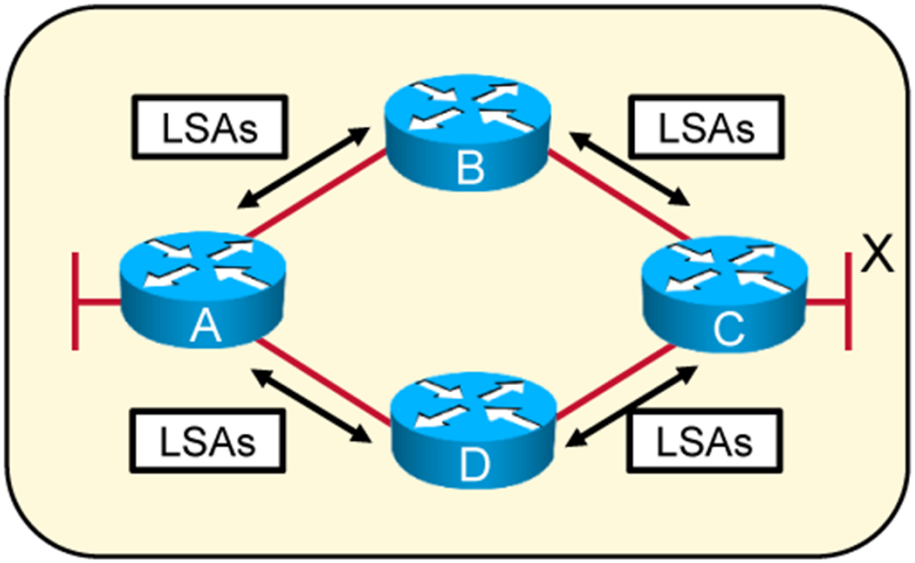

A: Neighbors Table
B
D

LSA Flooding

A: Topology Table
Router A Network 2 transit networks
Router D Network 2 transit networks
…

Dijkstra SPF

A: Routing Table
Network X → D
…

© 2013 Cisco Systems, Inc.

OSPF Metric

- OSPF uses path cost as a metric.
- By default, the cost is calculated based on the interface bandwidth.
- Cost = Reference Bandwidth / Interface Bandwidth, where reference bandwidth is 100 Mb/s.
- Path cost is a cumulated cost of all links on the path to destinations.

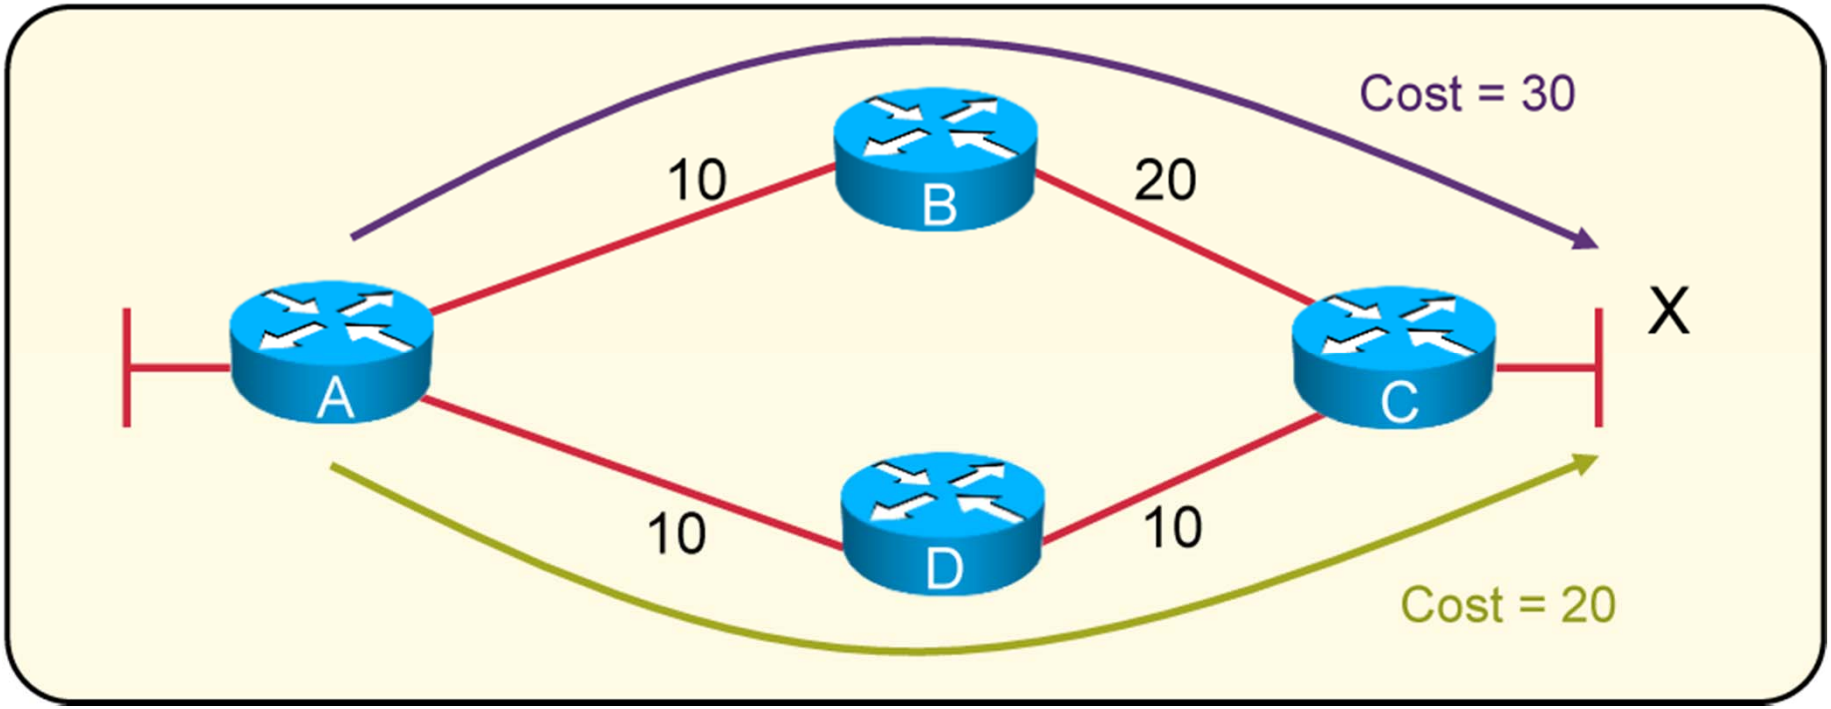

© 2013 Cisco Systems, Inc.

Establishing OSPF Neighbor Adjacencies

- OSPF routers first establish adjacencies.
- Hello packets are periodically sent to multicast address 224.0.0.5.
- Routers must agree on certain information (*) inside the hello packet before adjacency can be established.

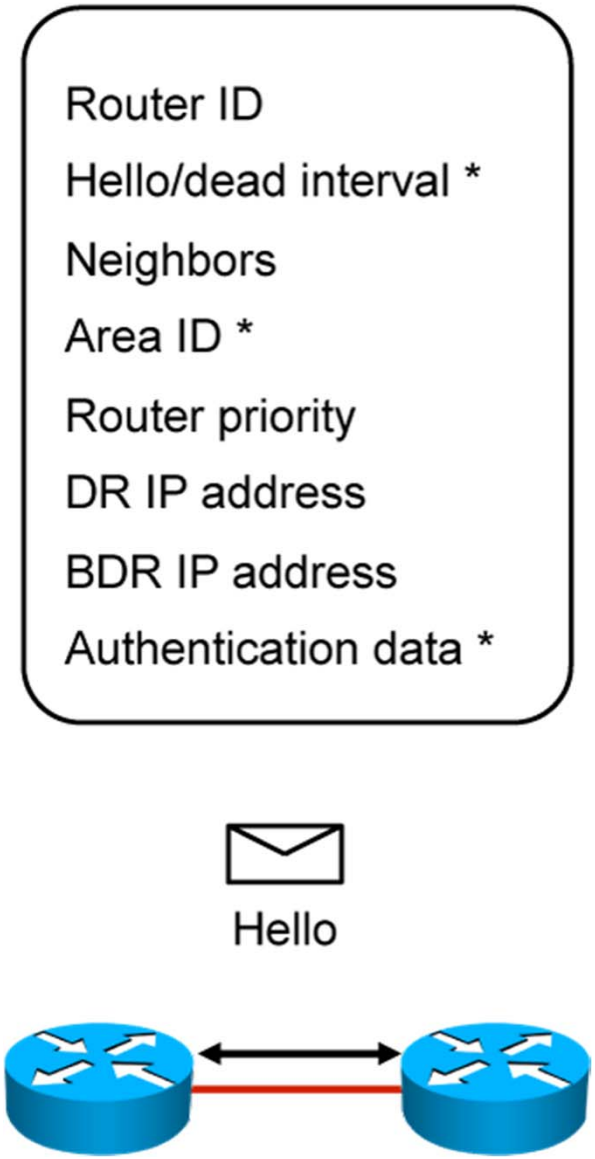

© 2013 Cisco Systems, Inc.

Building a Link-State Database

Four types of OSPF packets are involved in building a link-state database.

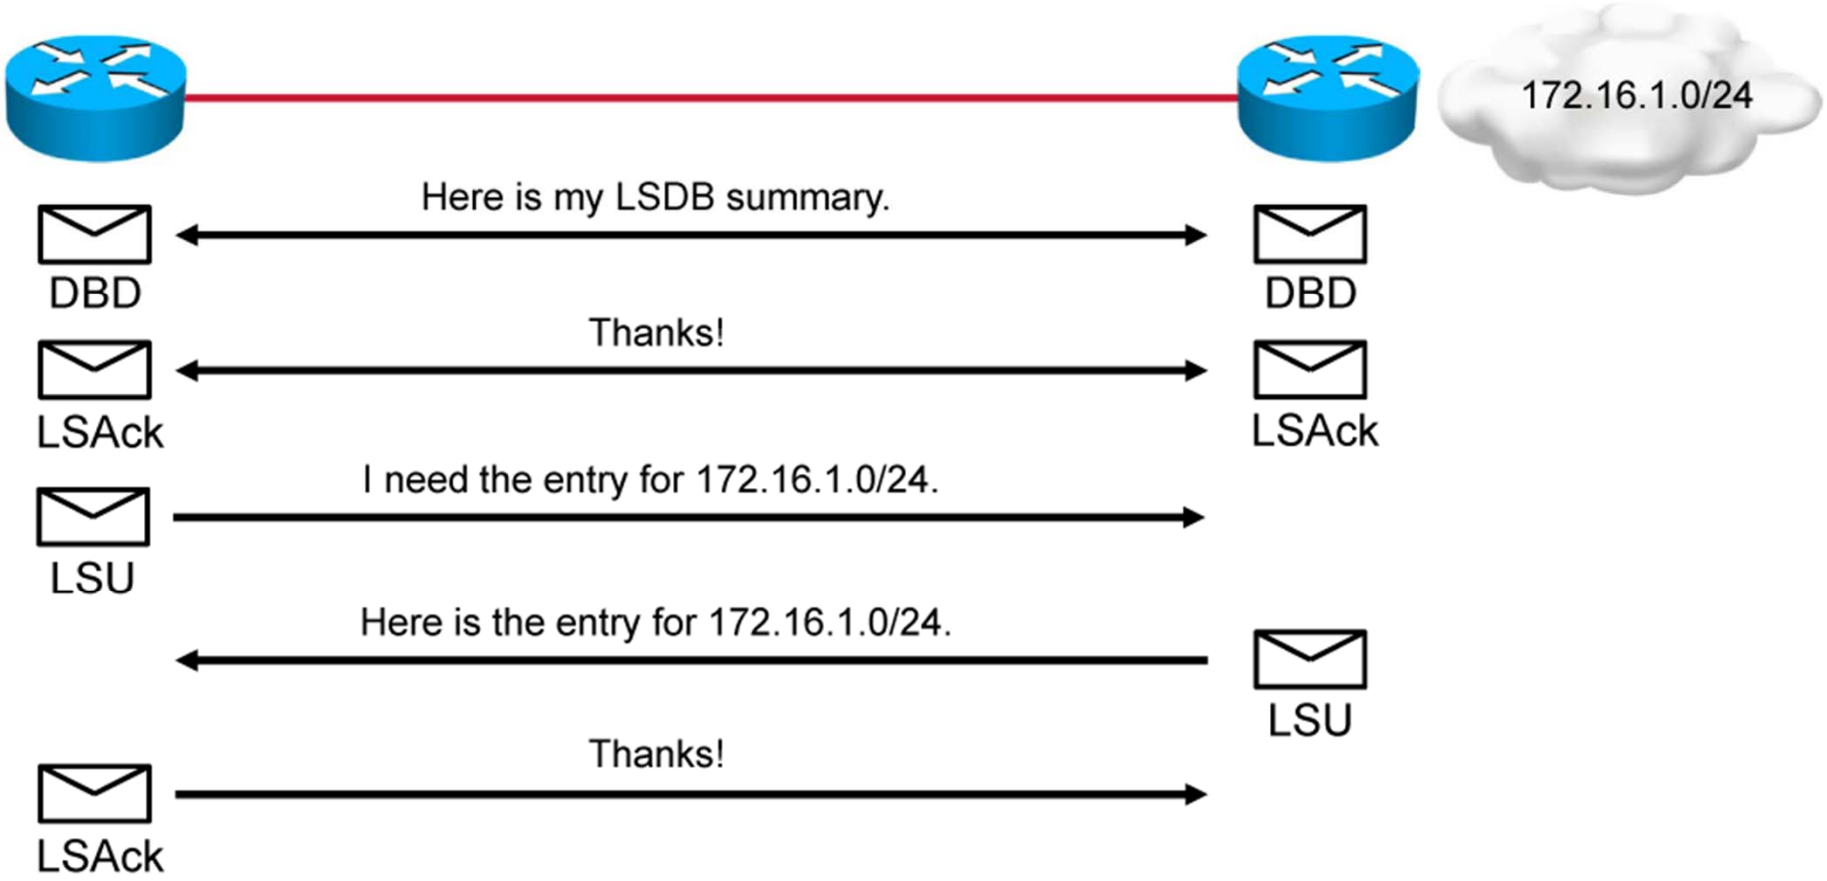

© 2013 Cisco Systems, Inc.

OSPF Area Structure

- OSPF supports a hierarchical network structure.
- The two-level hierarchy consists of the following:
  - Backbone area (Area 0)
  - Normal areas

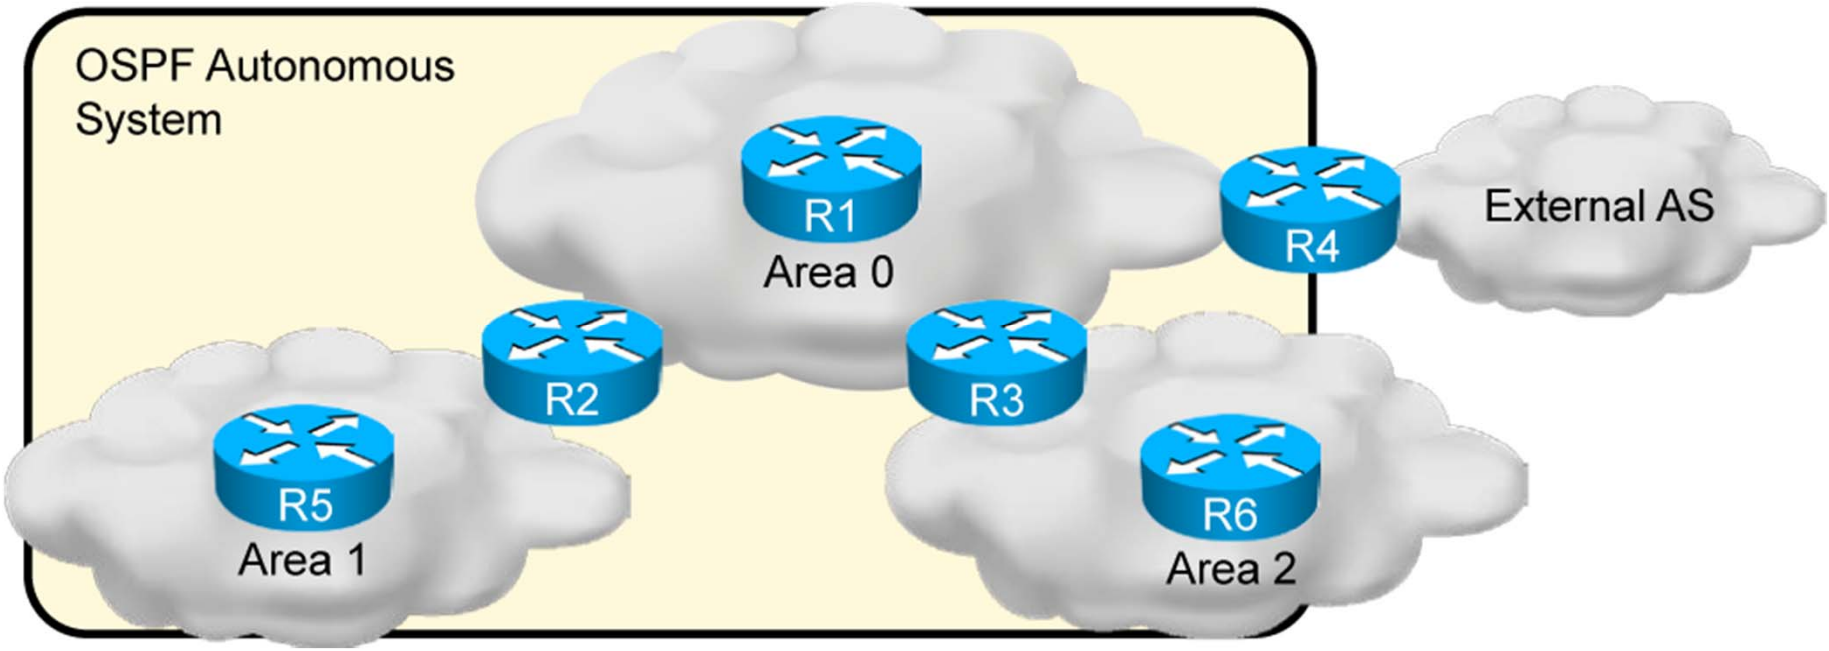

© 2013 Cisco Systems, Inc.

OSPF Area Structure (Cont.)

- Backbone routers
- Internal routers
- ABRs
- ASBRs

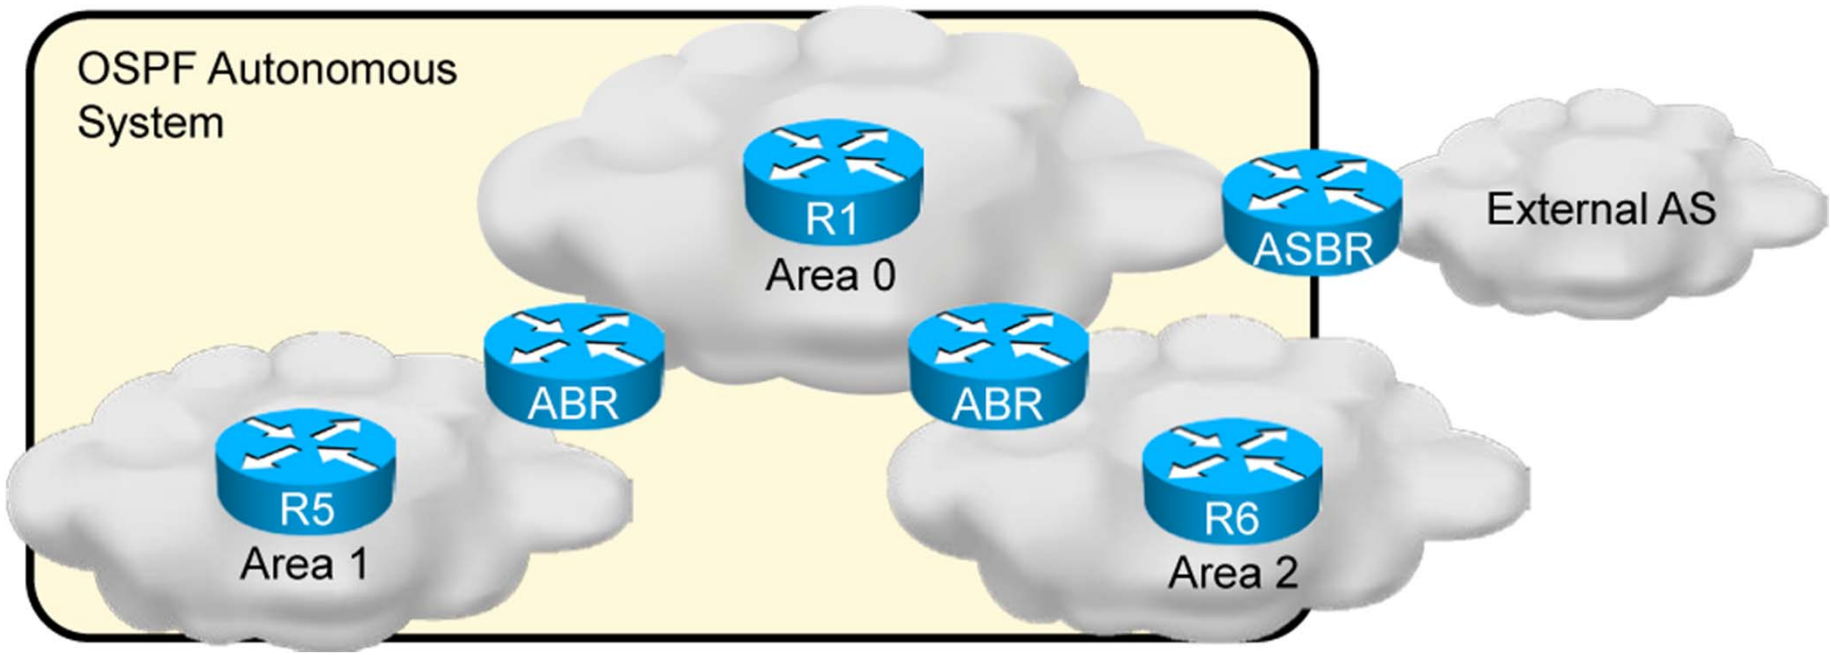

© 2013 Cisco Systems, Inc.

Basic LSA Types

LSA Type	Description
1	Router LSA
2	Network LSA
3	Summary LSA
4	ASBR summary LSA
5	Autonomous system LSA

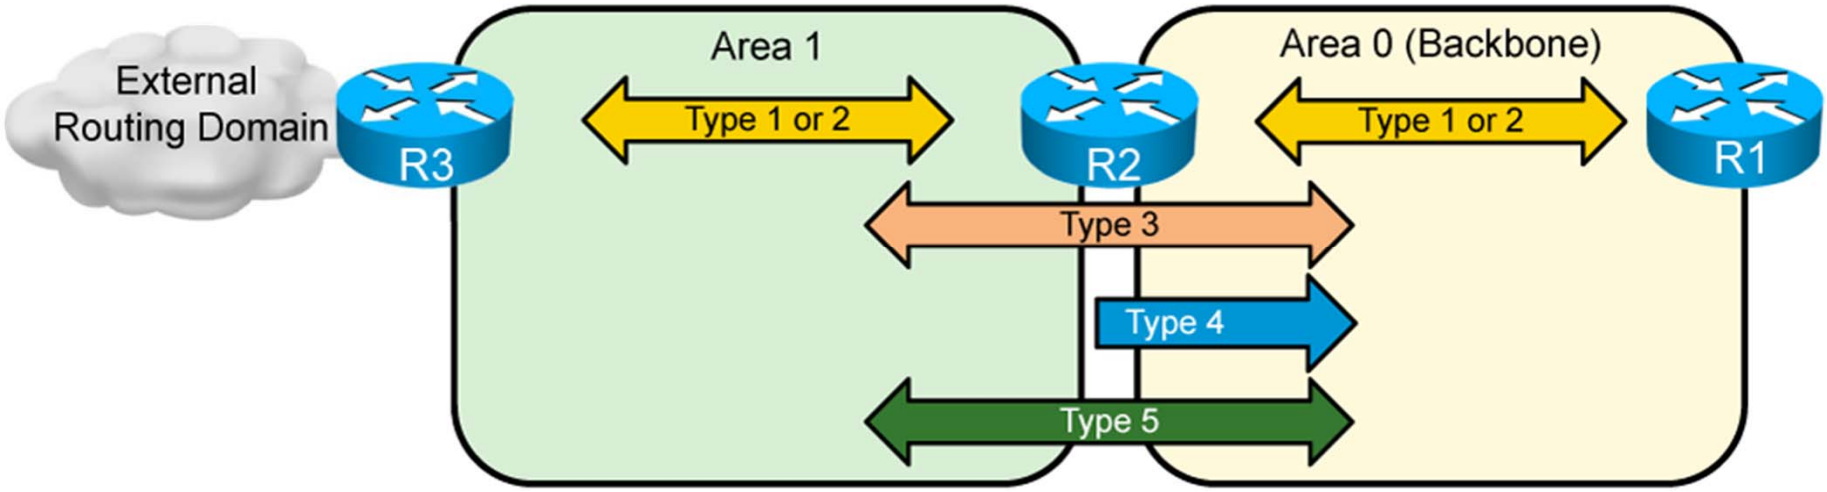

© 2013 Cisco Systems, Inc.

Summary

- Link-state routing protocols such as OSPF are more scalable and converge faster than distance-vector routing protocols.
- OSPF uses three data structures: neighbors table, topology table, and routing table.
- OSPF uses cost as a metric. The cost of an interface is inversely proportional to the bandwidth of the interface. Smaller cost indicates a better path than higher cost.
- OSPF routers first establish adjacency using hello packets.
- OSPF routers synchronize the LSDB by using the exchange protocol, which utilizes four OSPF packet types.
- OSPF supports a two-tier hierarchical network architecture.

© 2013 Cisco Systems, Inc.

© 2013 Cisco Systems, Inc.

Multiarea OSPF Implementation

Implementing a Scalable, Multiarea Network, OSPF-Based Solution

© 2013 Cisco Systems, Inc.

Single-Area vs. Multiarea OSPF

- Single-area OSPF:
  - Many LSAs processed on every router
  - Large routing tables
- Multiarea OSPF:
  - LSA processing confined to an area
  - Smaller routing tables if summarization is used

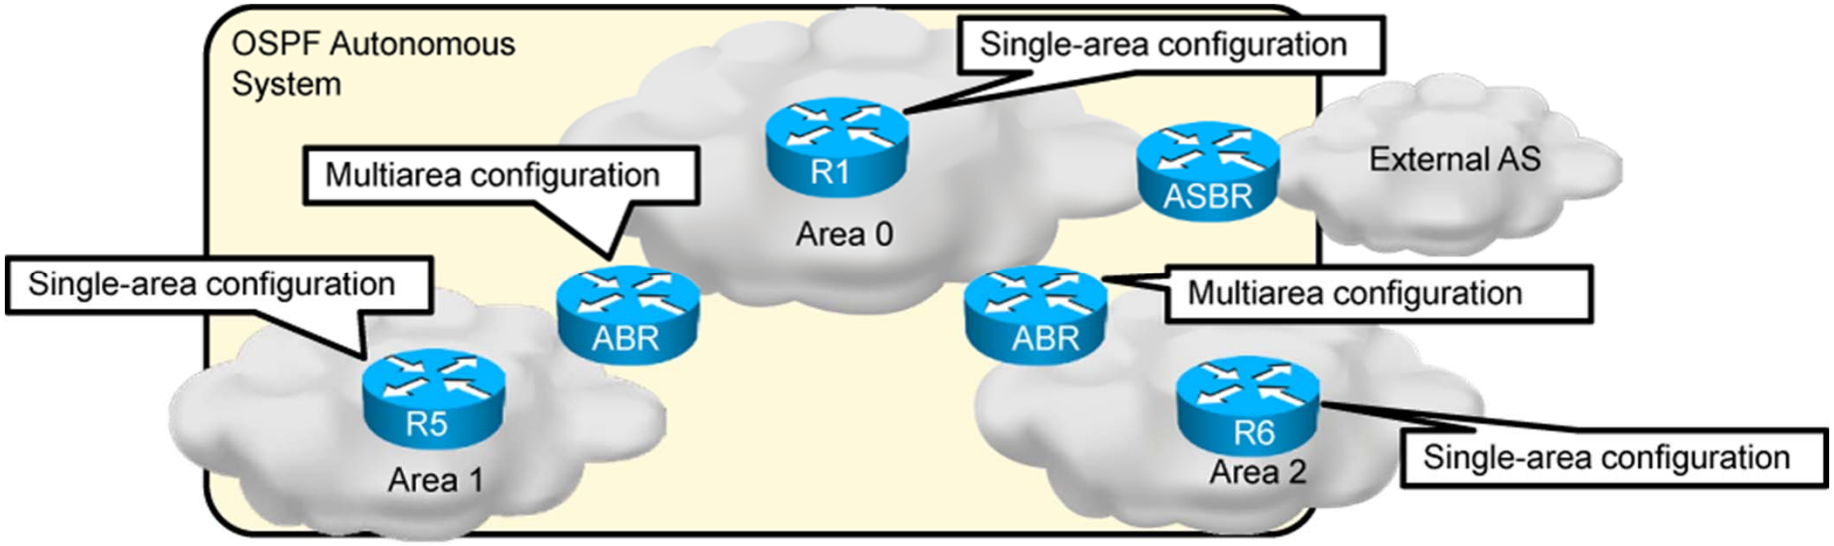

© 2013 Cisco Systems, Inc.

Planning for the Implementation of OSPF

- Assess the requirements and options:
  - Verify IP addressing.
  - Verify network topology.
- Define ABRs and ASBRs.
- Create an implementation plan.
- Configure OSPF.

© 2013 Cisco Systems, Inc.

Multiarea OSPF Configuration

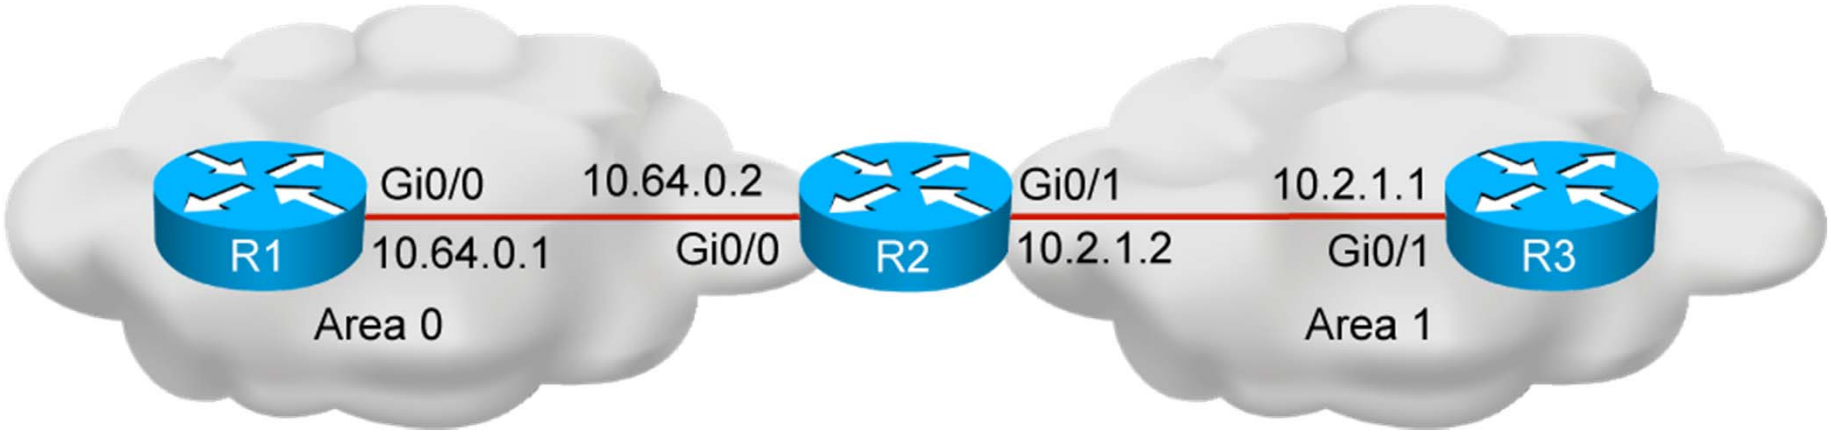


```
R1(config)#interface GigabitEthernet0/0
R1(config-if)#ip address 10.64.0.1 255.255.255.0
R1(config-if)#ip ospf cost 10
R1(config-if)#exit
R1(config)#router ospf 1
R1(config-router)#network 10.64.0.0  0.0.0.255 area 0
```

OSPF configuration on R1

© 2013 Cisco Systems, Inc.

Multiarea OSPF Configuration (Cont.)

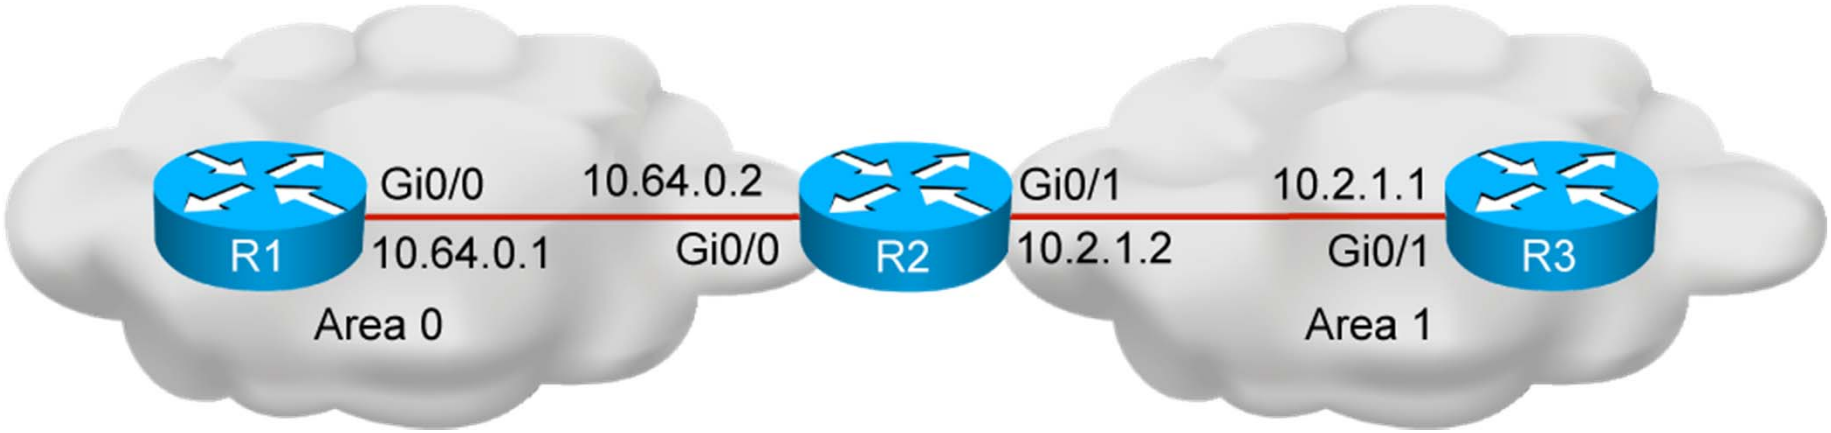


```
R2(config)#interface GigabitEthernet0/0
R2(config-if)#ip address 10.64.0.2  255.255.255.0
R2(config-if)#exit
R2(config)#interface GigabitEthernet0/1
R2(config-if)#ip address 10.2.1.2  255.255.255.0
R2(config-if)#exit
R2(config)#router ospf 1
R2(config-router)#network 10.64.0.0  0.0.0.255 area 0
R2(config-router)#network 10.2.1.2  0.0.0.0 area 1
```

OSPF configuration on R2

© 2013 Cisco Systems, Inc.

Multiarea OSPF Verification

```
R2#show ip ospf neighbor

Neighbor ID  Pri   State           Dead Time   Address         Interface
10.64.0.1      1   FULL/DR         00:00:31    10.64.0.1 GigabitEthernet0/0
10.2.1.1       1   FULL/DR         00:00:51    10.2.1.1  GigabitEthernet0/1
```

- Verify OSPF neighbors. R2 has two neighbors in this example.

```
R2#show ip ospf interface
GigabitEthernet0/0 is up, line protocol is up
  Internet Address 10.64.0.2/24, Area 0
  Process ID 1, Router ID 10.64.0.2, Network Type BROADCAST, Cost: 1
  <output omitted>
GigabitEthernet0/1 is up, line protocol is up
  Internet Address 10.2.1.2/24, Area 1
  Process ID 1, Router ID  10.64.0.2/24, Network Type BROADCAST, Cost: 1
  <output omitted>
```

- Verify OSPF-enabled interfaces. R2 has two interfaces in this example: GigabitEthernet0/0 in Area 0 and GigabitEthernet0/1 in Area 1.

© 2013 Cisco Systems, Inc.

Multiarea OSPF Verification (Cont.)

```
R2#show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 1"
<output omitted>
  Routing for Networks:
    10.64.0.0 0.0.0.255 area 0
    10.2.1.2 0.0.0.0 area 1
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.64.0.1       110           00:06:07
    10.2.1.1        110           00:06:07
   <output omitted>
```

- Verify for which networks that R2 is routing.

© 2013 Cisco Systems, Inc.

Multiarea OSPF Verification (Cont.)

```
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 not set

      10.0.0.0/24 is subnetted, 1 subnets
O IA     10.2.1.0 [110/11 via 10.64.0.2, 00:46:20, GigabitEthernet0/0
<output omitted>
```

- Verify the routing table.

© 2013 Cisco Systems, Inc.

Summary

- Multiarea OSPF design enables segmentation of a network to limit the propagation of LSAs inside an area and to make routing tables smaller (if utilizing summarization).
- The type of OSPF implementation depends on requirements and existing topology.
- Multiarea OSPF configuration is similar to single-area configuration. Interfaces are assigned to areas using the **network** command.
- Multiarea OSPF verification is similar to single-area verification. Interarea routes are marked with the "IA" code in the routing table.

© 2013 Cisco Systems, Inc.

© 2013 Cisco Systems, Inc.

Troubleshooting Multiarea OSPF

Implementing a Scalable, Multiarea Network, OSPF-Based Solution

© 2013 Cisco Systems, Inc.

OSPF Neighbor States

OSPF routers go through different OSPF states:

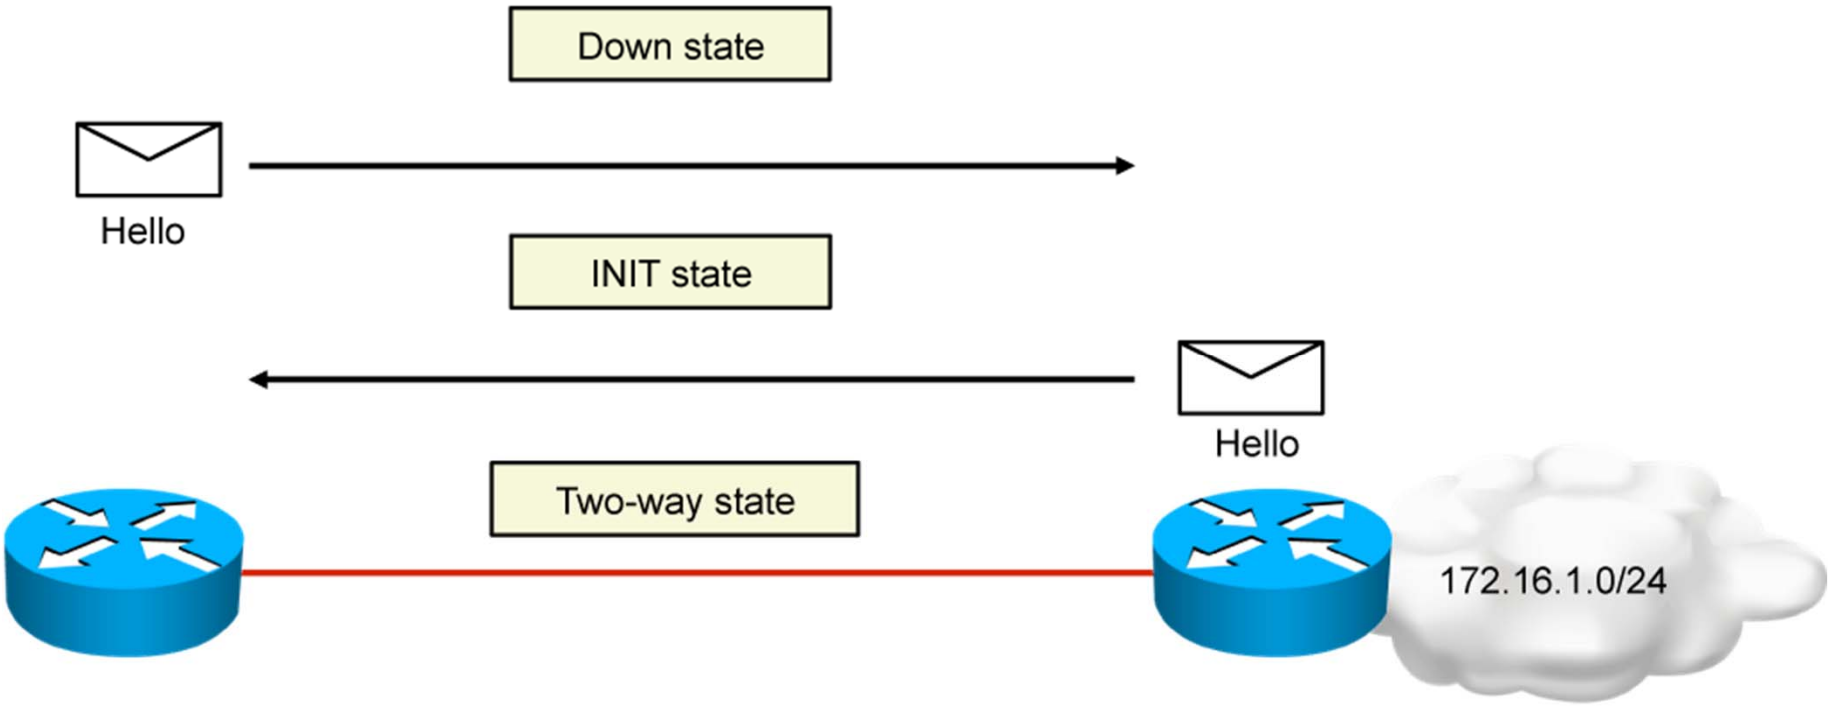

© 2013 Cisco Systems, Inc.

OSPF Neighbor States (Cont.)

All states except two-way and full are transitory, and routers should not remain in these states for extended periods of time.

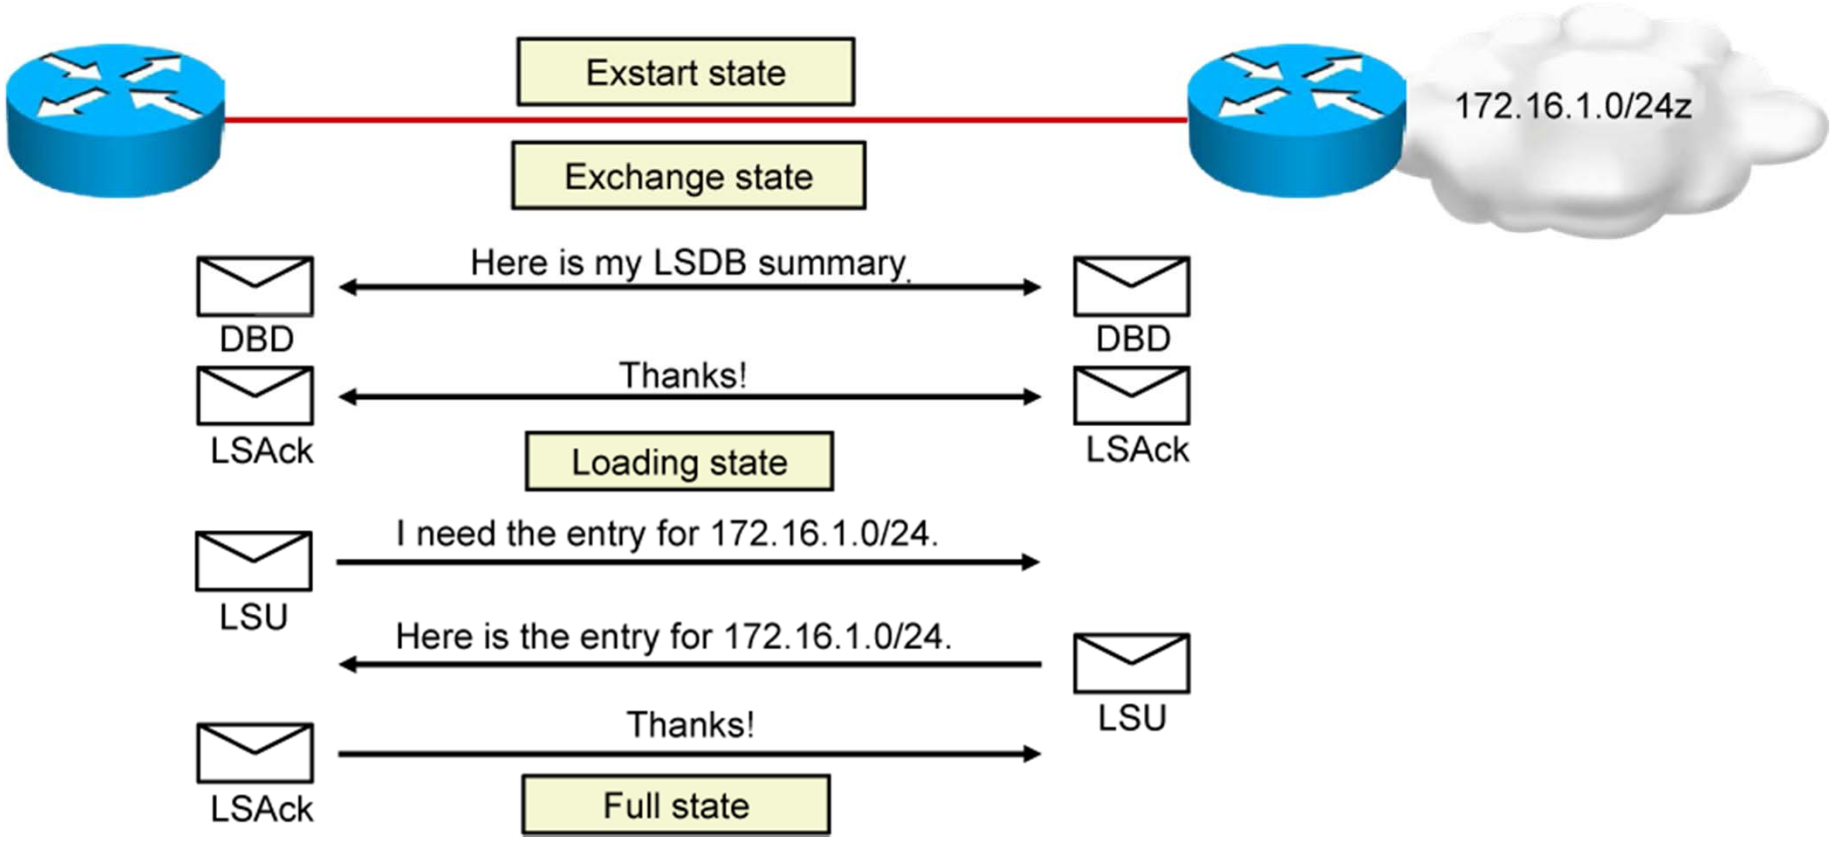

© 2013 Cisco Systems, Inc.

Components of Troubleshooting OSPF

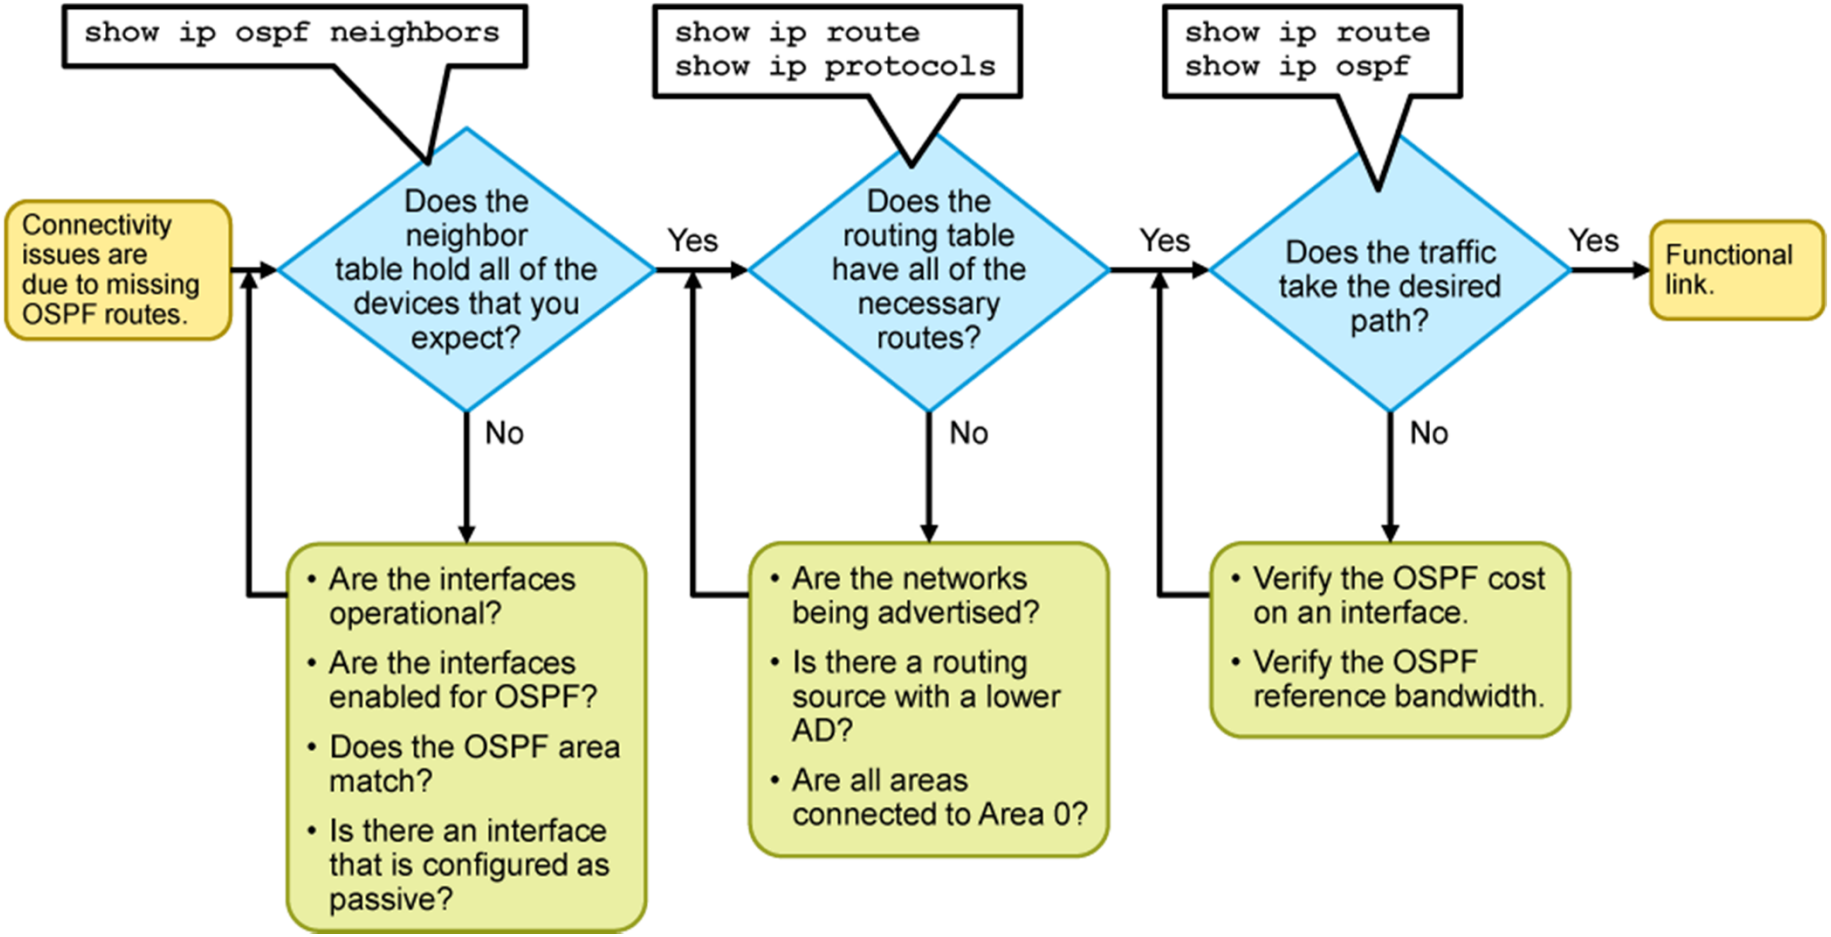

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues

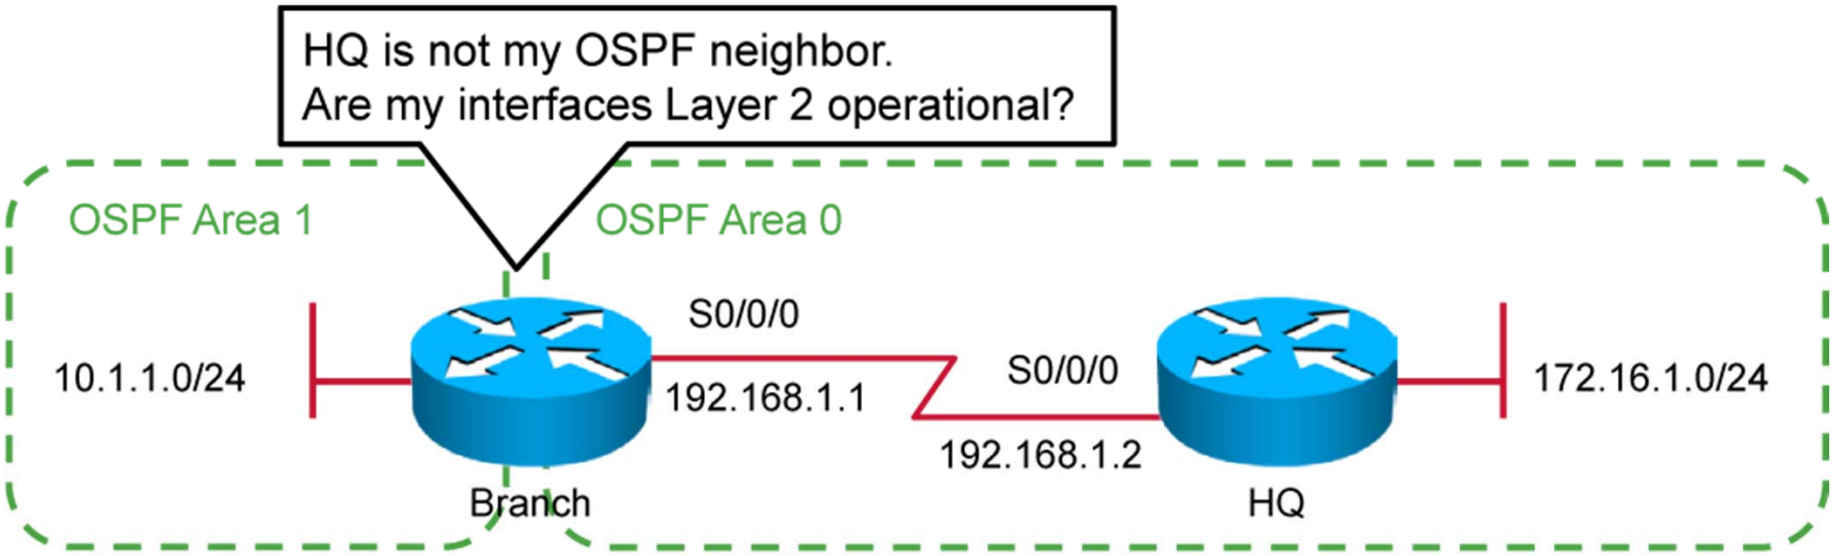


```
Branch#show ip interface brief
Interface                  IP-Address      OK? Method Status  Protocol
GigabitEthernet0/0         10.1.1.1        YES manual up       up
Serial0/0/0                192.168.1.1     YES manual up       up
<output omitted>
```

Verify if Serial 0/0/0 on the Branch router is Layer 2 operational.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)

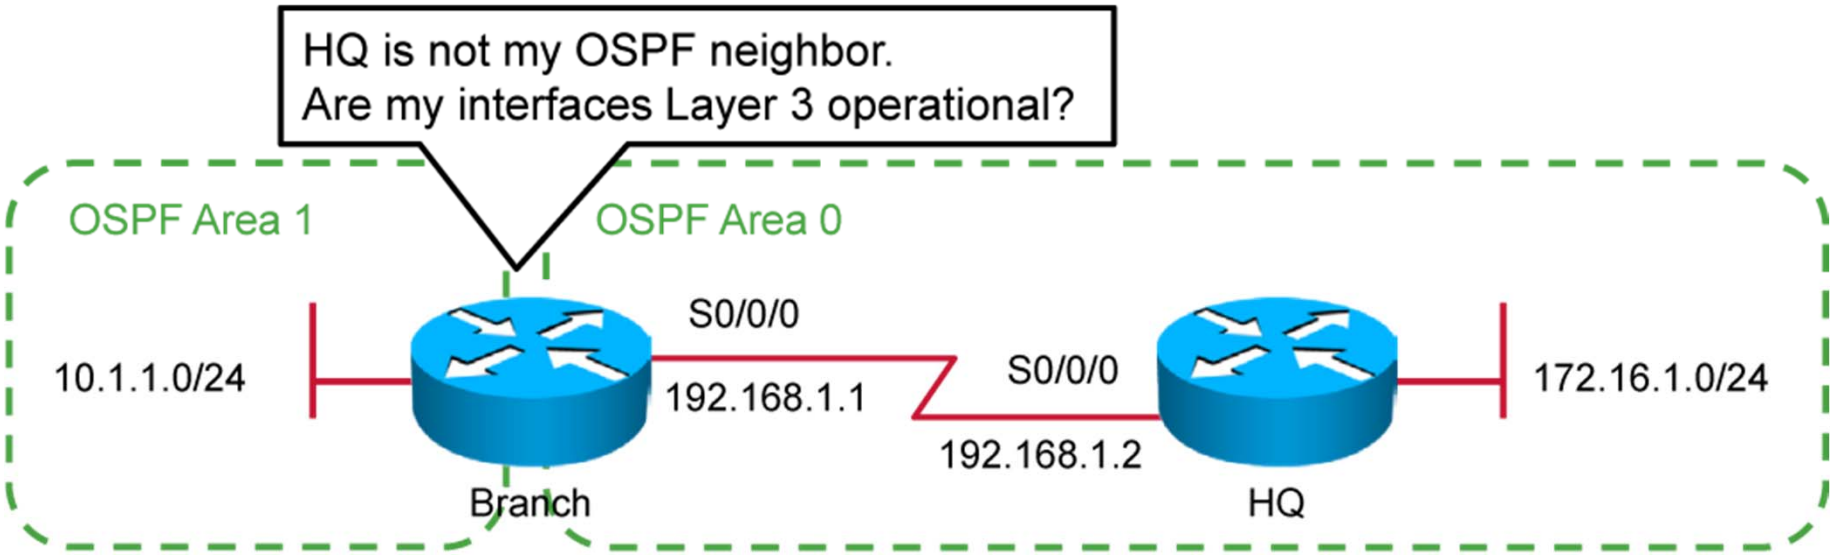


```
Branch#ping 192.168.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Verify Layer 3 connectivity between the Branch and HQ routers.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)

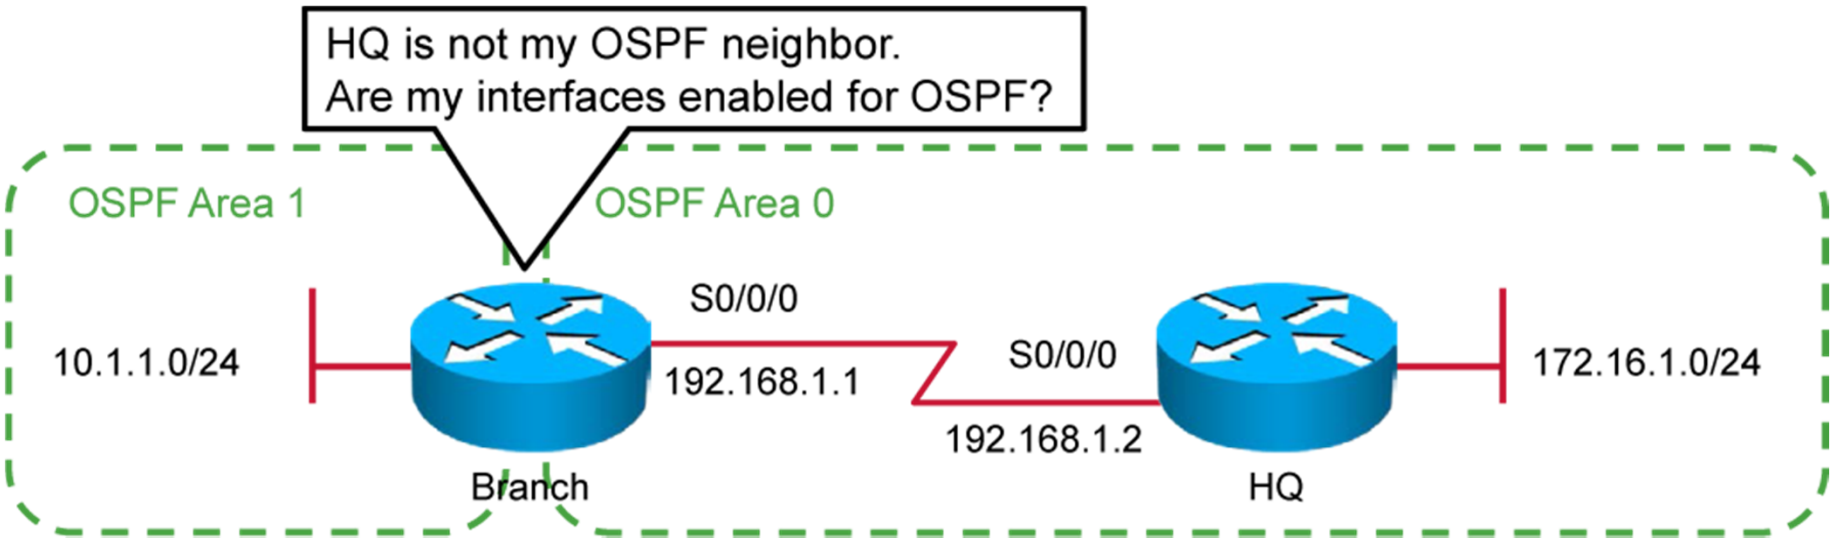


```
Branch#show ip ospf interface
Serial0/0/0 is up, line protocol is up
  Internet Address 192.168.1.1/24, Area 0
  Process ID 1, Router ID 209.165.201.1, Network Type POINT_TO_POINT, Cost:
64
 <output omitted>
```

If the interfaces on both routers are not enabled for OSPF, the adjacency will not form.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)

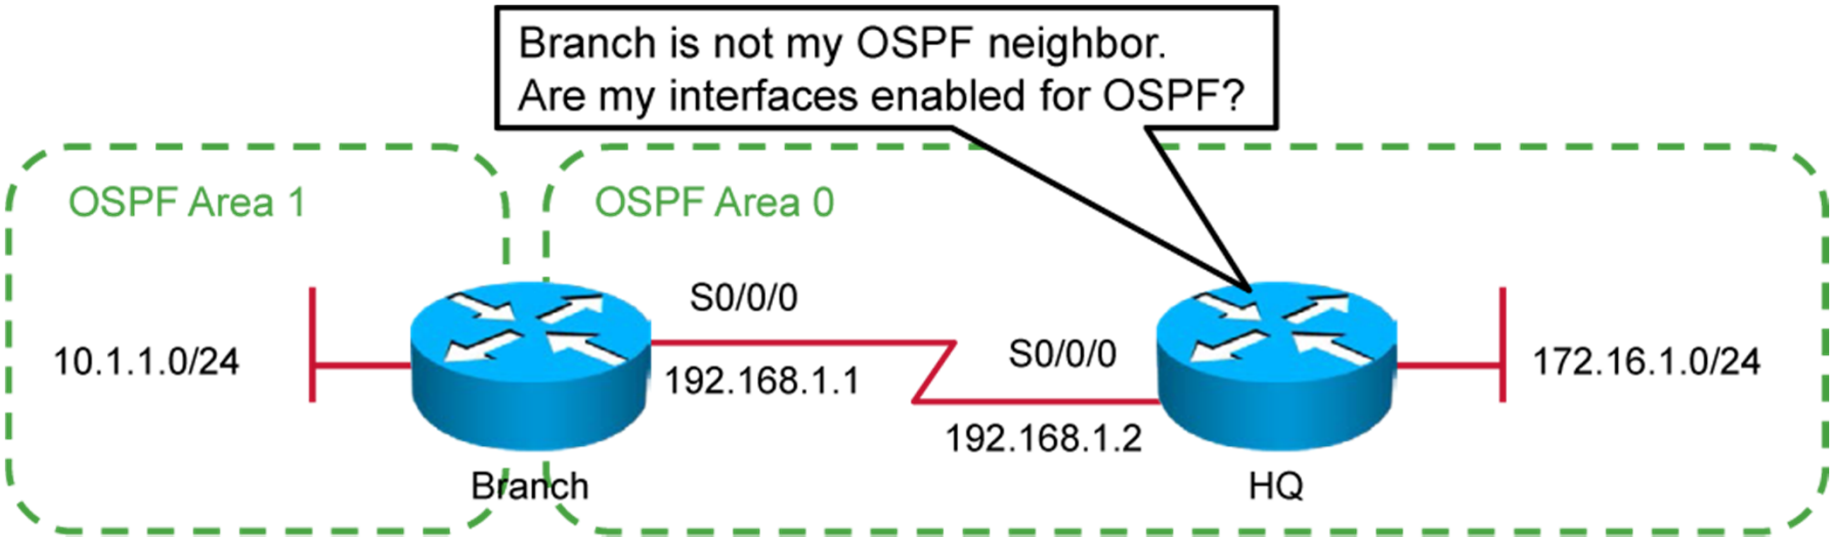


```
HQ#show ip ospf interface
Serial0/0/0 is up, line protocol is up
  Internet Address 192.168.1.2/24, Area 0
  Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 64
<output omitted>
```

If the interfaces on both routers are not enabled for OSPF, the adjacency will not form.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)

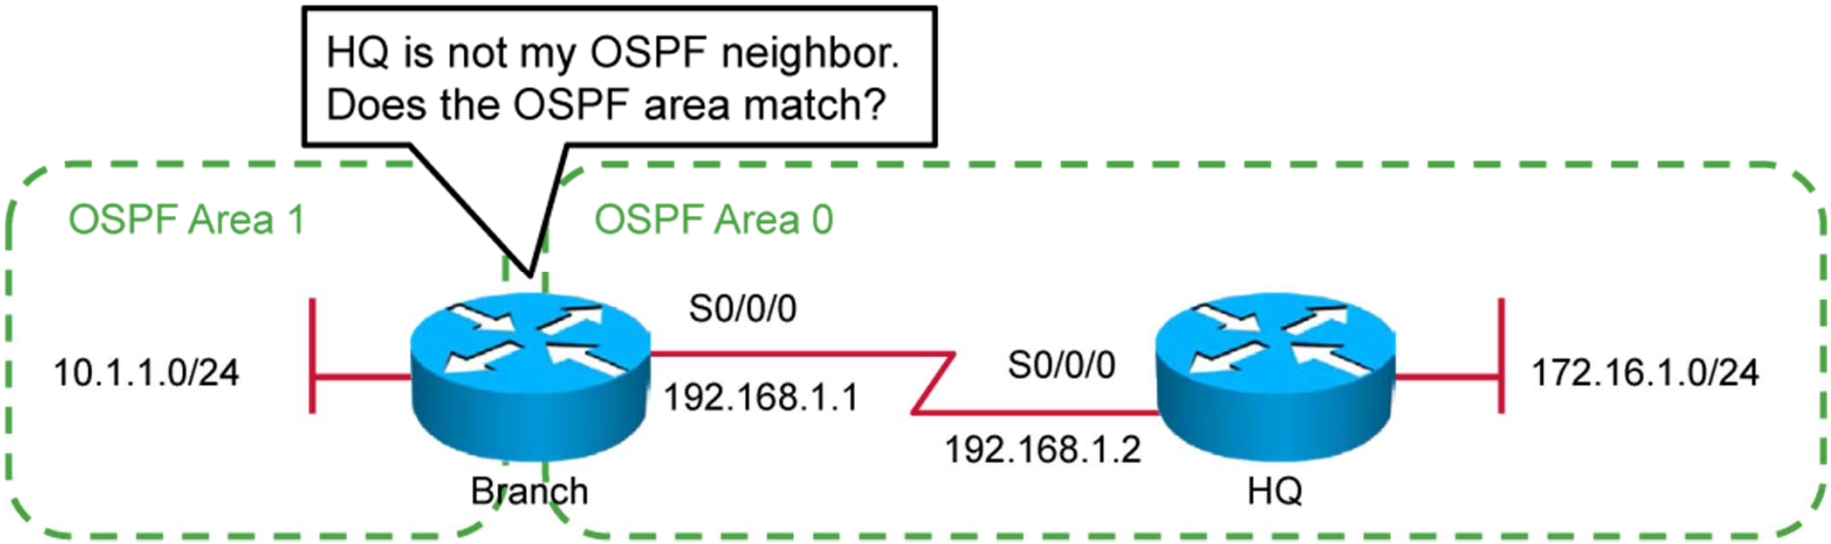


```
Branch#show ip protocols
Routing Protocol is "ospf 1"
<output omitted>
  Maximum path: 4
  Routing for Networks:
    10.1.1.0 0.0.0.255 area 1
    192.168.1.0 0.0.0.255 area 0
 <output omitted>
```

If the OSPF area does not match on both ends, the adjacency will not form.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)

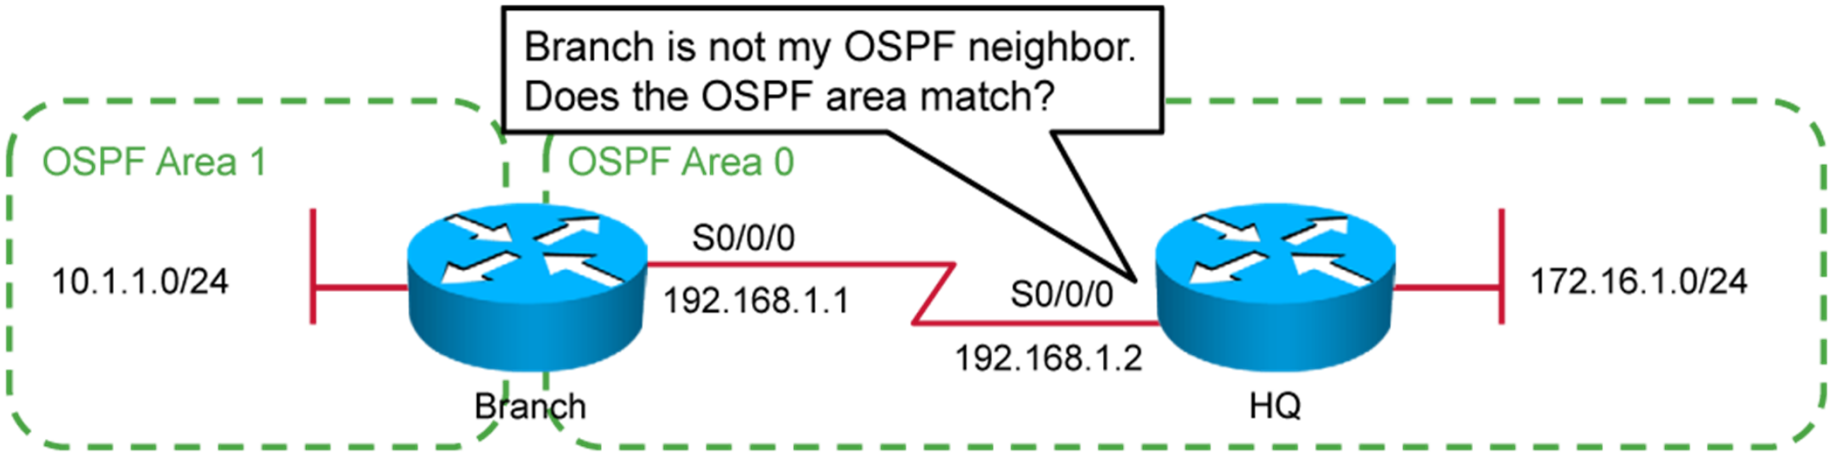


```
HQ#show ip protocols
Routing Protocol is "ospf 1"
<output omitted>
  Routing for Networks:
    172.16.1.0 0.0.0.255 area 0
    192.168.1.0 0.0.0.255 area 0
 <output omitted>
```

If the OSPF area does not match on both ends, the adjacency will not form.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)

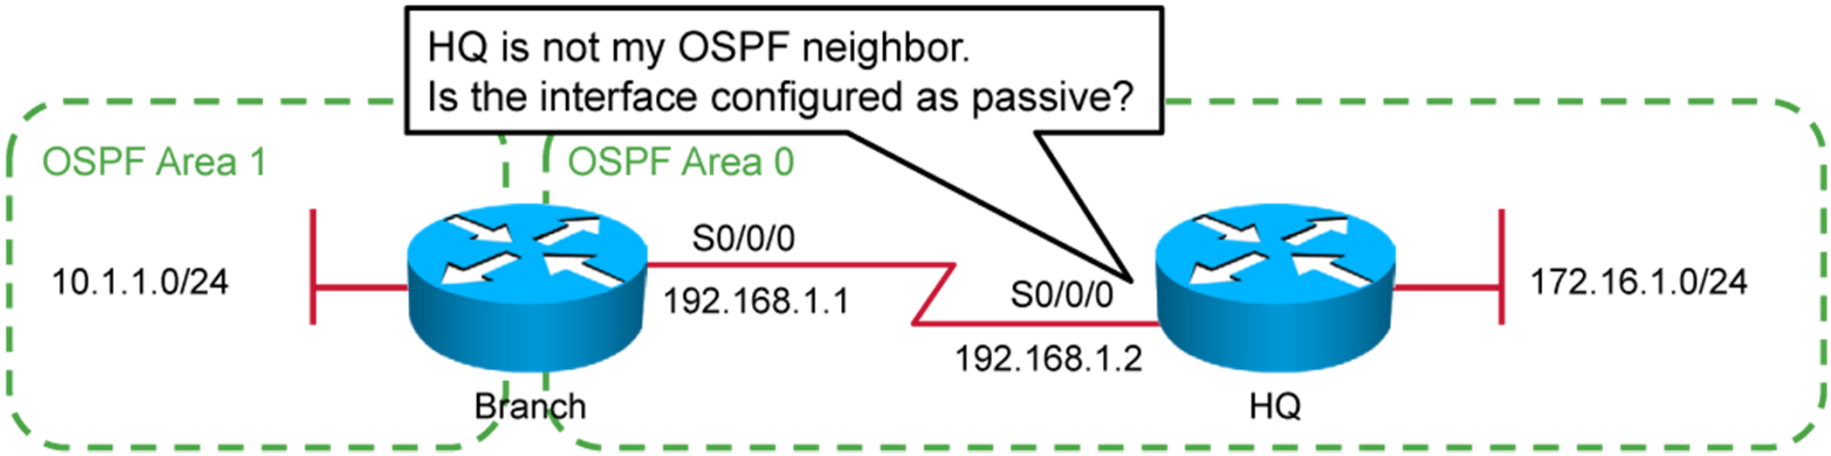


```
Branch#show ip protocols
Routing Protocol is "ospf 1"
<output omitted>
  Routing for Networks:
    10.1.1.0 0.0.0.255 area 1
    192.168.1.0 0.0.0.255 area 0
 <output omitted>
```

Check if the interface toward the HQ router is configured as passive.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)

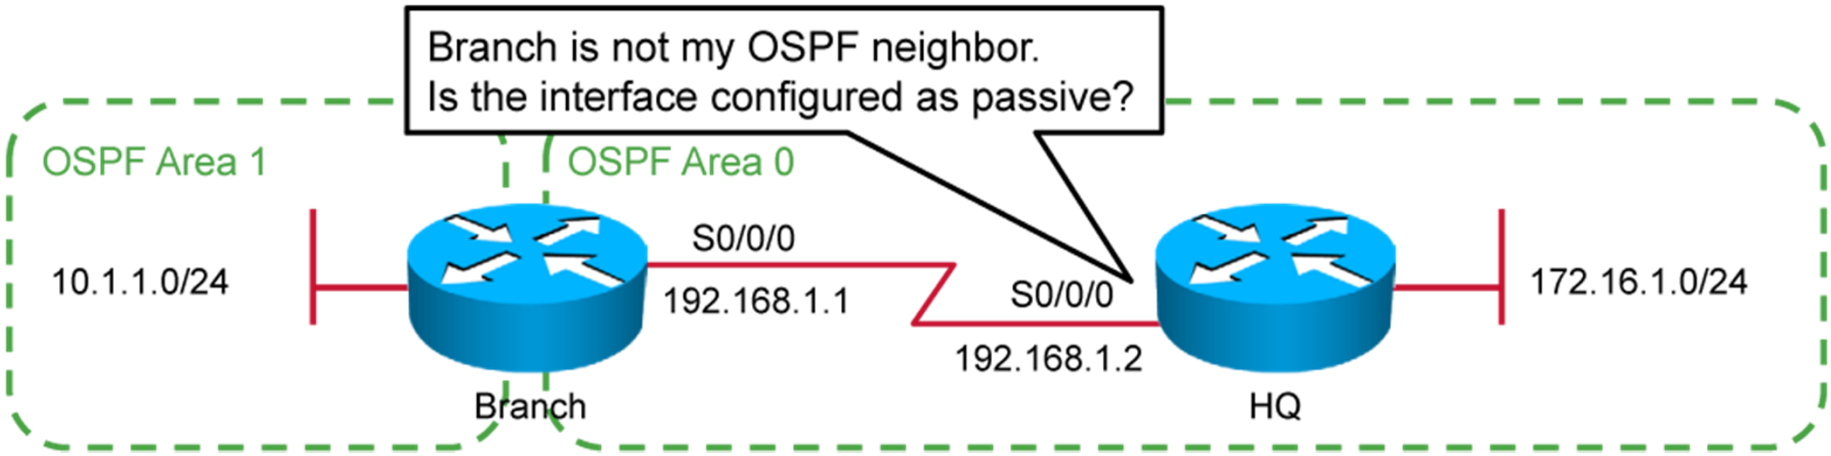


```
HQ#show ip protocols
Routing Protocol is "ospf 1"
  <output omitted>
  Routing for Networks:
    172.16.1.0 0.0.0.255 area 0
    192.168.1.0 0.0.0.255 area 0
  Passive Interface(s):
    Serial0/0/0
 <output omitted>
```

HQ has the interface toward the Branch router configured as passive. This is why the two routers are not forming an adjacency.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Routing Table Issues

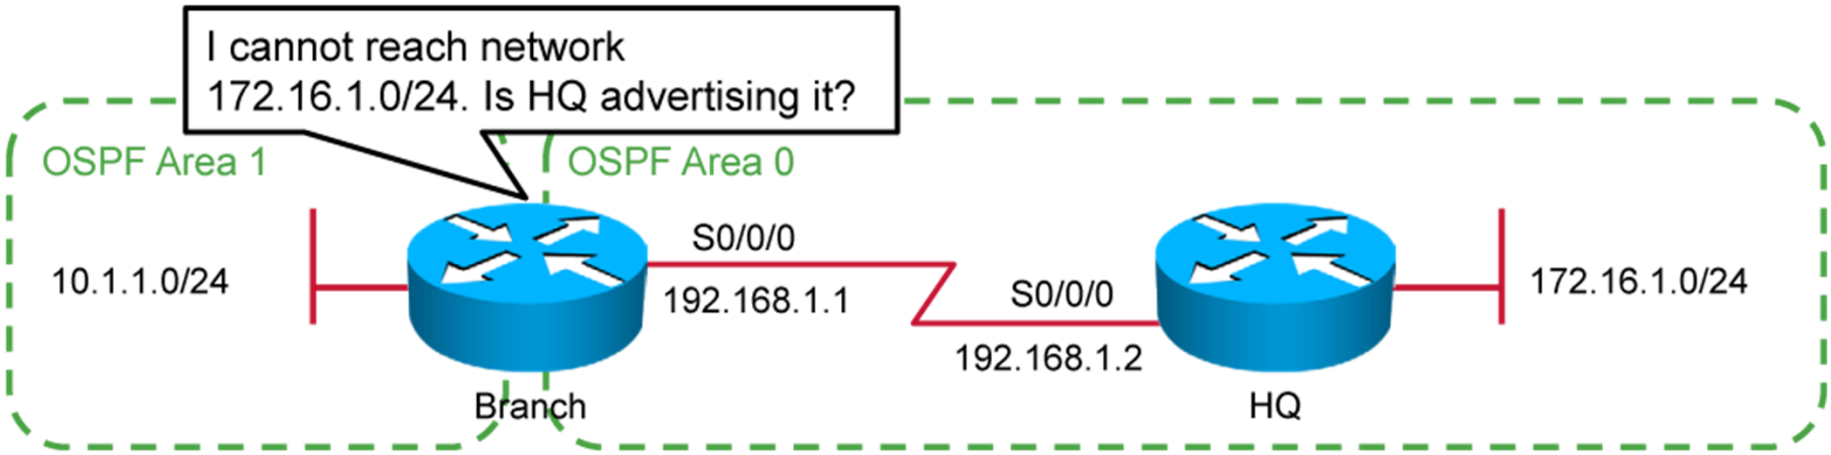


```
HQ#show ip protocols
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    192.168.1.0 0.0.0.255 area 0
<output omitted>
```

HQ has to advertise the network in order to reach it from the Branch router.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Routing Table Issues (Cont.)

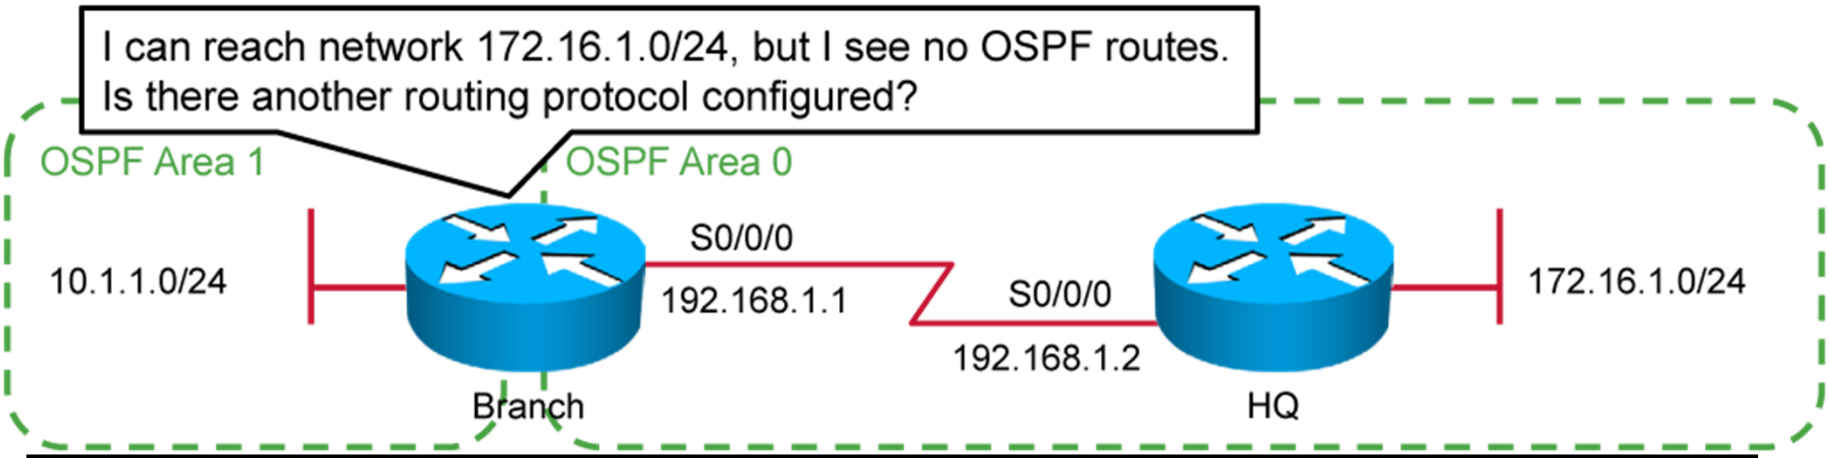


```
Branch#show ip route 172.16.1.0
Routing entry for 172.16.1.0/24
  Known via "eigrp 1", distance 90, metric 2297856, type internal
  Redistributing via eigrp 1
  Last update from 192.168.1.2 on Serial0/0/0, 00:00:39 ago
  Routing Descriptor Blocks:
  * 192.168.1.2, from 192.168.1.2, 00:00:39 ago, via Serial0/0/0
      Route metric is 2297856, traffic share count is 1
      Total delay is 25000 microseconds, minimum bandwidth is 1544 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

If several routing protocols are configured on routers, administrative distance will decide which protocol will be used.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Path Selection

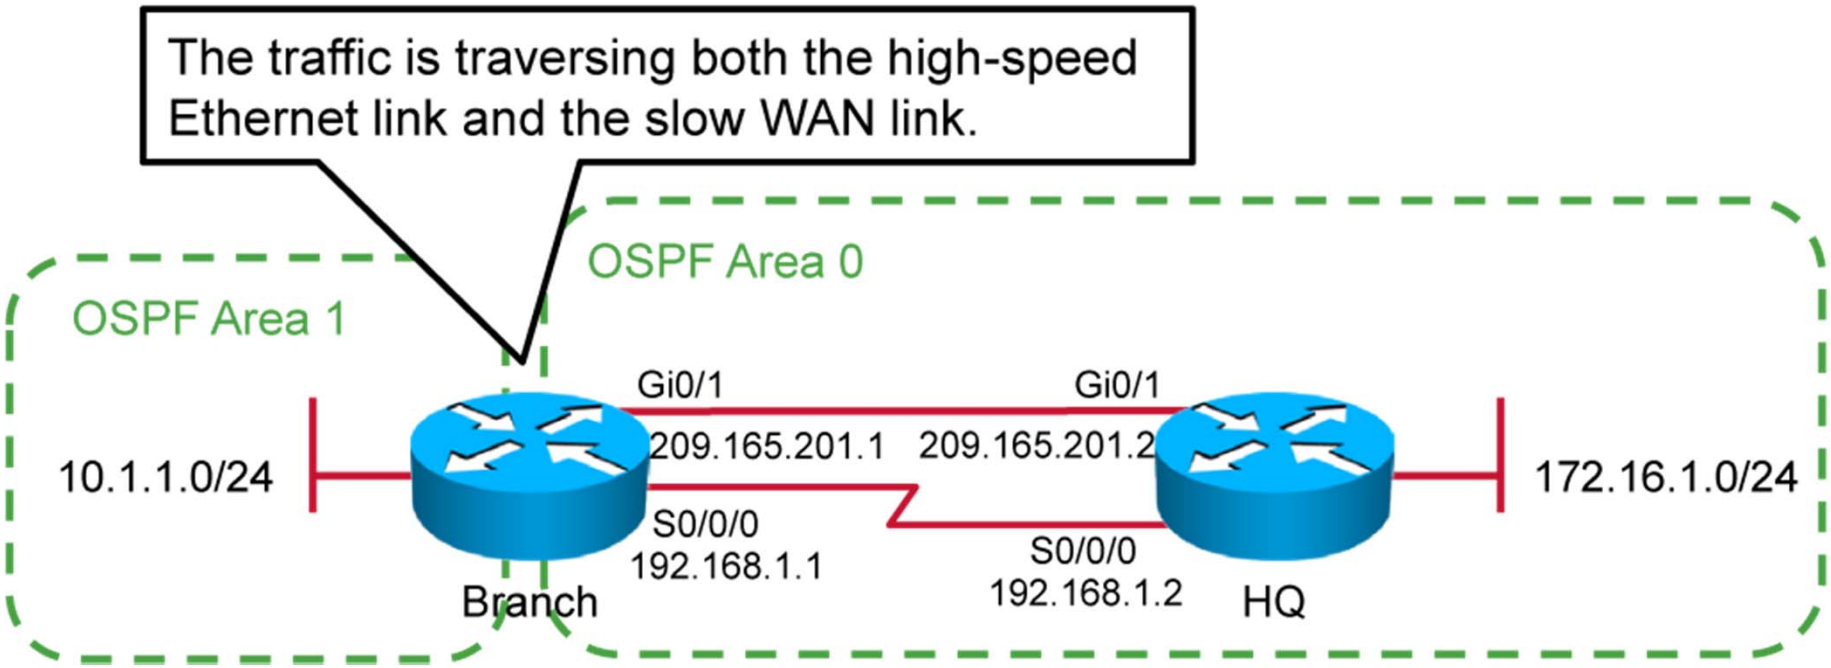


```
Branch#show ip route ospf
     172.16.0.0/24 is subnetted, 1 subnets
O       172.16.1.0 [110/2] via 209.165.201.2, 00:04:34, GigabitEthernet0/1
                   [110/2] via 192.168.1.2, 00:00:32, Serial0/0/0
```

Traffic is traversing both links because interfaces have the same cost.

© 2013 Cisco Systems, Inc.

Summary

- OSPF routers traverse different OSPF states when adjacencies are being established.
- When troubleshooting connectivity problems due to OSPF, you should first verify OSPF neighbors.
- If OSPF areas on two routers do not match, an adjacency will not form.
- OSPF routing issues could stem from the fact that networks are not advertised, or there could be an ACL that is blocking a routing advertisement.
- When you have OSPF path selection issues, you have to make sure that traffic takes the desired path through the network by manipulating the OSPF cost.

© 2013 Cisco Systems, Inc.

© 2013 Cisco Systems, Inc.

Examining OSPFv3

Implementing a Scalable, Multiarea Network, OSPF-Based Solution

© 2013 Cisco Systems, Inc.

OSPFv3 Key Characteristics

- OSPFv3 is an implementation of the OSPF routing protocol for IPv6.
- OSPFv2 (for IPv4) and OSPFv3 (for IPv6) run independently on the router.
- OSPFv3 has the same key capabilities as OSPFv2 for IPv4 networks:
  - Multiarea network design with ABRs that segment the network
  - SPF algorithm for optimal path calculation

© 2013 Cisco Systems, Inc.

OSPFv3 Key Characteristics (Cont.)

- The router ID is a 32-bit number that is based on the IPv4 address of the router. If there is no IPv4 address that is present on the router, you are prompted to configure it using the **router-id** command.
- Adjacencies and next-hop attributes use link-local addresses.
- IPv6 is used for transport of the LSA.
- OSPFv3 is enabled per link, not per network.
- OSPFv3 communicates using IPv6 multicast addresses.

© 2013 Cisco Systems, Inc.

OSPFv3 Configuration

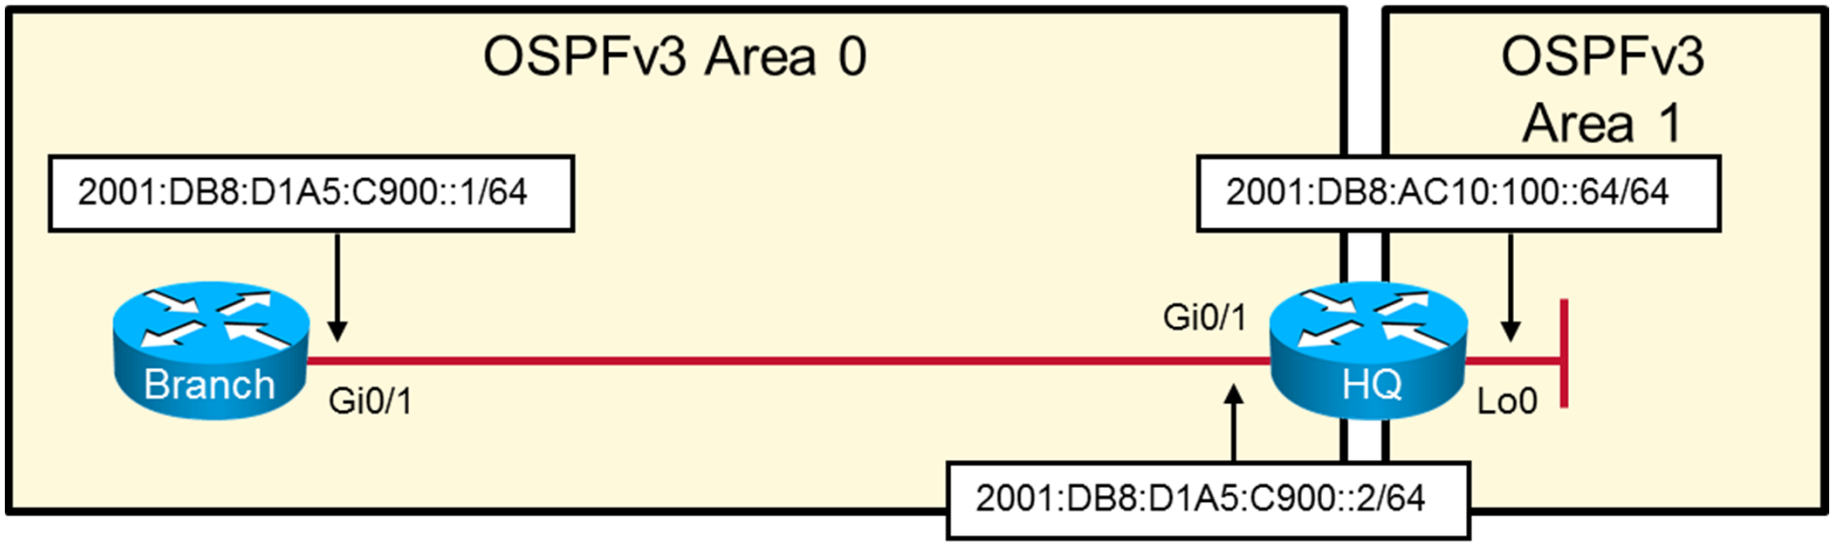


```
Branch(config)#ipv6 router ospf 99
Branch(config-rtr)#router-id 1.1.1.1
Branch(config-rtr)#exit
Branch(config)#interface GigabitEthernet0/1
Branch(config-if)#ipv6 address 2001:DB8:D1A5:C900::1/64
Branch(config-if)#ipv6 ospf 99 area 0
```

- OSPFv3 configuration on the Branch router

© 2013 Cisco Systems, Inc.

OSPFv3 Configuration (Cont.)

```
HQ(config)#ipv6 router ospf 99
HQ(config-rtr)#router-id 2.2.2.2
HQ(config-rtr)#exit
HQ(config)#interface Loopback0
HQ(config-if)#ipv6 address 2001:DB8:AC10:100::64/64
HQ(config-if)#ipv6 ospf 99 area 0.0.0.1
HQ(config-if)#exit
HQ(config)#interface GigabitEthernet0/1
HQ(config-if)#ipv6 address 2001:DB8:D1A5:C900::2/64
HQ(config-if)#ipv6 ospf 99 area 0.0.0.0
```

- OSPFv3 configuration on the HQ router

© 2013 Cisco Systems, Inc.

Cisco OSPFv3 Verification

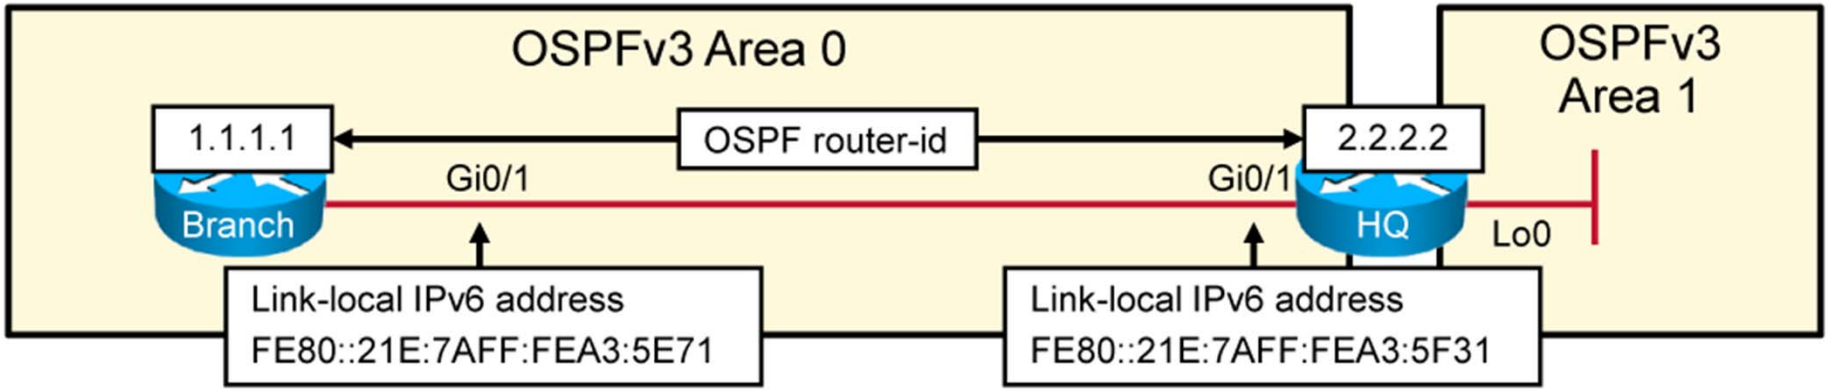


```
Branch#show ipv6 ospf interface
GigabitEthernet0/1 is up, line protocol is up
  Link Local Address FE80::21E:7AFF:FEA3:5E71, Interface ID 5
  Area 0.0.0.0, Process ID 99, Instance ID 0, Router ID 1.1.1.1
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.1, local address FE80::21E:7AFF:FEA3:5E71
  Backup Designated router (ID) 2.2.2.2, local address
FE80::21E:7AFF:FEA3:5F31
  <output omitted>
    Adjacent with neighbor 2.2.2.2  (Backup Designated Router)
  <output omitted>
```

- OSPFv3 interface information on the Branch router

© 2013 Cisco Systems, Inc.

Cisco OSPFv3 Verification (Cont.)

```
Branch#show ipv6 ospf
 Routing Process "ospfv3 99" with ID 1.1.1.1
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
<output omitted>
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 Reference bandwidth unit is 100 mbps
    Area BACKBONE(0.0.0.0)
        Number of interfaces in this area is 1
<output omitted>
```

- OSPFv3 general information on the Branch router

© 2013 Cisco Systems, Inc.

Cisco OSPFv3 Verification (Cont.)

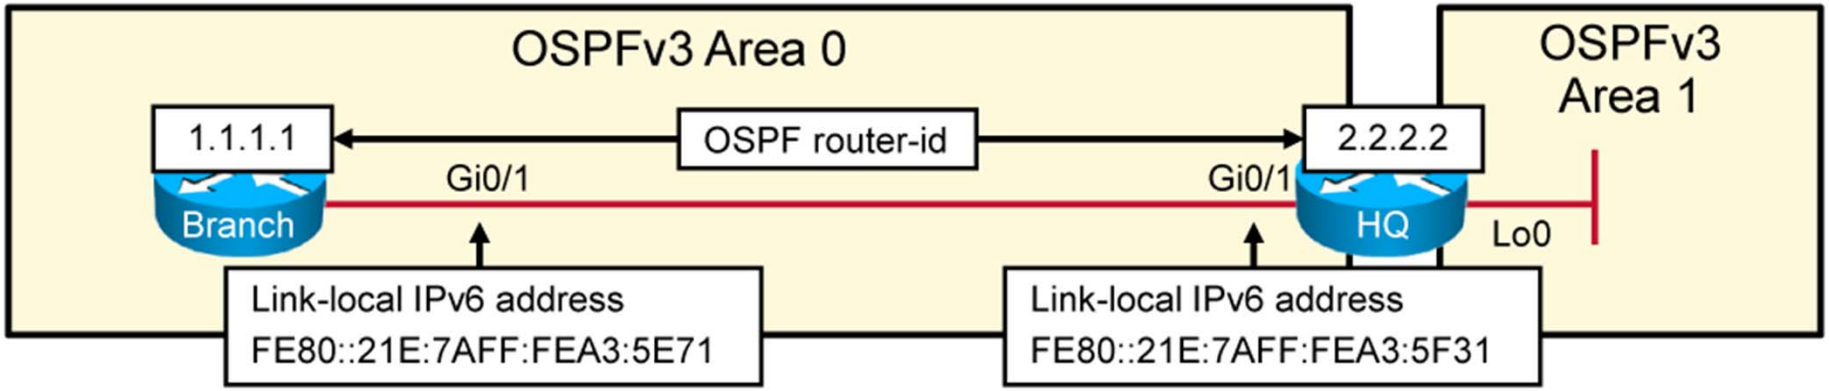


```
Branch#show ipv6 ospf neighbor

            OSPFv3 Router with ID (1.1.1.1) (Process ID 99)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
2.2.2.2           0   FULL/  -        00:00:37    6               Gi0/1
```

- OSPFv3 neighbor information on the Branch router

© 2013 Cisco Systems, Inc.

Cisco OSPFv3 Verification (Cont.)

```
Branch#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, 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
OI  2001:DB8:AC10:100::64/64 [110/64]
     via FE80::21E:7AFF:FEA3:5F31, GigabitEthernet0/1
```

- OSPFv3 route information on the Branch router

© 2013 Cisco Systems, Inc.

Summary

- OSPFv3 for IPv6 supports the same basic mechanisms that OSPFv2 for IPv4 supports, including the use of areas to provide network segmentation and LSAs to exchange routing updates.
- OSPFv3 is configured per-interface on Cisco routers.
- To verify OSPFv3 configuration, use commands that are similar to those that are used for OSPFv2. Instead of the keyword **ip,** use the keywords **ipv6**: **show ipv6 ospf interface** and **show ipv6 ospf**.

© 2013 Cisco Systems, Inc.

© 2013 Cisco Systems, Inc.

Module Summary

- OSPF uses three data structures: neighbors table, topology table, and routing table.
- Multiarea OSPF design enables segmentation of a network to limit the propagation of LSAs inside an area and to make routing tables smaller.
- When troubleshooting connectivity problems due to OSPF, you should first verify OSPF neighbors.
- OSPFv3 for IPv6 supports the same basic mechanisms that OSPFv2 for IPv4 does, including the use of areas to provide network segmentation and LSAs to exchange routing updates.

© 2013 Cisco Systems, Inc.

© 2013 Cisco Systems, Inc.