



IEWB-RS Technology Labs

OSPF

Brian Dennis, CCIE # 2210 (R&S / ISP Dial / Security / Service Provider)
Brian McGahan, CCIE # 8583 (R&S / Service Provider)

Copyright Information

Copyright © 2003 - 2007 Internetwork Expert, Inc. All rights reserved.

The following publication, ***CCIE Routing and Switching Lab Workbook***, was developed by Internetwork Expert, Inc. All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means without the prior written permission of Internetwork Expert, Inc.

Cisco®, Cisco® Systems, CCIE, and Cisco Certified Internetwork Expert, are registered trademarks of Cisco® Systems, Inc. and/or its affiliates in the U.S. and certain countries. All other products and company names are the trademarks, registered trademarks, and service marks of the respective owners. Throughout this manual, Internetwork Expert, Inc. has used its best efforts to distinguish proprietary trademarks from descriptive names by following the capitalization styles used by the manufacturer.

Disclaimer

The following publication, ***CCIE Routing and Switching Lab Workbook***, is designed to assist candidates in the preparation for Cisco Systems' CCIE Routing & Switching Lab exam. While every effort has been made to ensure that all material is as complete and accurate as possible, the enclosed material is presented on an "as is" basis. Neither the authors nor Internetwork Expert, Inc. assume any liability or responsibility to any person or entity with respect to loss or damages incurred from the information contained in this workbook.

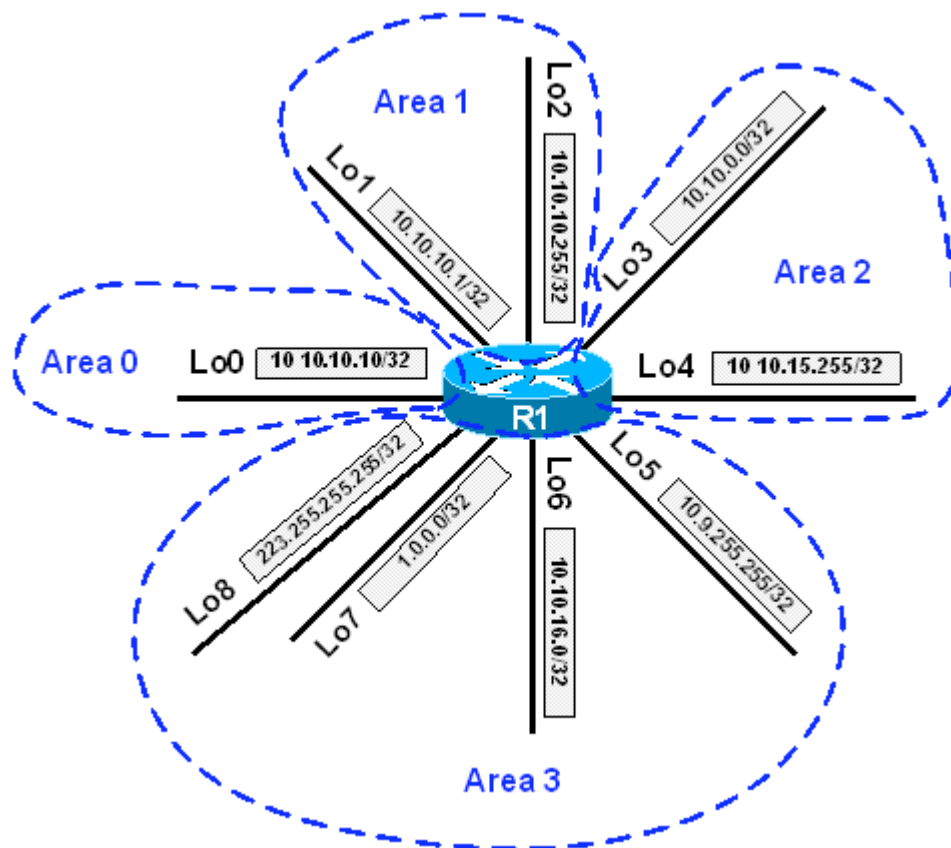
This workbook was developed by Internetwork Expert, Inc. and is an original work of the aforementioned authors. Any similarities between material presented in this workbook and actual CCIE™ lab material is completely coincidental.

UNDERSTANDING THE OSPF NETWORK STATEMENT	1
OSPF DR/BDR ELECTION.....	5
OSPF OVER FRAME RELAY - NON-BROADCAST	9
OSPF OVER FRAME RELAY - BROADCAST	14
OSPF OVER FRAME RELAY - POINT-TO-MULTIPOINT	19
OSPF OVER FRAME RELAY - POINT-TO-MULTIPOINT NON-BROADCAST	25
OSPF OVER FRAME RELAY - POINT-TO-POINT.....	30
OSPF NETWORK TYPE LOOPBACK.....	33
OSPF VIRTUAL LINKS - REPAIRING AREA 0	37
OSPF VIRTUAL LINKS - REPAIRING DISCONTIGUOUS AREAS	43
COMMON MULTI-AREA CONFIGURATION WITH 2 ABRs.....	47
USING TYPE-3 LSA INTER-AREA FILTERING.....	52
TYPE-3 LSA FILTERING WITH NETWORK RANGES.....	54
INGRESS FILTERING WITH DISTRIBUTE-LIST	56
INGRESS FILTERING WITH DISTRIBUTE-LIST AND ROUTE-MAP	58
INGRESS FILTERING WITH ADMINISTRATIVE DISTANCE	61
NSSA AREA TYPE 7TO5 LSA TRANSLATOR ELECTION.....	65
NSSA AREA ABR EXTERNAL PREFIX FILTERING	69
NSSA SUPPRESS FA FEATURE	71
NSSA AREA AND DEFAULT-ROUTE ORIGIANTION AT ABR PART 1	73
NSSA AREA AND DEFAULT-ROUTE ORIGIANTION AT ABR PART 2.....	75
NSSA AREA AND DEFAULT-ROUTE ORIGIANTION AT ASBR	77

Understanding the OSPF Network Statement

Objective: Enable OSPF on R1 with the following area assignments using exactly four network statements

- Interface 10.10.10.10 should be assigned to area 0
- Interfaces 10.10.10.0 - 10.10.10.255 should be assigned to area 1
- Interface 10.10.0.0 - 10.10.15.255 should be assigned to area 2
- All other interfaces (1.0.0.0 - 223.255.255.255) should be assigned to area 3



Directions

- Create interface loopback0 on R1 with the IP address 10.10.10.10/32
- Create interface loopback1 on R1 with the IP address 10.10.10.1/32
- Create interface loopback2 on R1 with the IP address 10.10.10.255/32
- Create interface loopback3 on R1 with the IP address 10.10.0.0/32
- Create interface loopback4 on R1 with the IP address 10.10.15.255/32
- Create interface loopback5 on R1 with the IP address 10.9.255.255/32
- Create interface loopback6 on R1 with the IP address 10.10.16.0/32
- Create interface loopback7 on R1 with the IP address 1.0.0.0/32
- Create interface loopback8 on R1 with the IP address 223.255.255.255/32

- Enable OSPF using a process-id of 1
- Configure a network statement so that the interface 10.10.10.10 is in OSPF area 0
- Configure a network statement so that interfaces in the range of 10.10.10.0 - 10.10.10.255 are in OSPF area 1
- Configure a network statement so that interfaces in the range of 10.10.0.0 - 10.10.15.255 are in OSPF area 2
- Configure a network statement so that interfaces in the range of 1.0.0.0 - 223.255.255.255 are in OSPF area 3

Ask Yourself

- How does the wildcard mask relate to the subnet mask of the interface?
- What happens if multiple OSPF network statements overlap?
- How can I verify which area an interface belongs to?

Final Configuration

```
R1:
interface Loopback0
 ip address 10.10.10.10 255.255.255.255
!
interface Loopback1
 ip address 10.10.10.1 255.255.255.255
!
interface Loopback2
 ip address 10.10.10.255 255.255.255.255
!
interface Loopback3
 ip address 10.10.0.0 255.255.255.255
!
interface Loopback4
 ip address 10.10.15.255 255.255.255.255
!
interface Loopback5
 ip address 10.9.255.255 255.255.255.255
!
interface Loopback6
 ip address 10.10.16.0 255.255.255.255
!
interface Loopback7
 ip address 1.0.0.0 255.255.255.255
!
interface Loopback8
 ip address 223.255.255.255 255.255.255.255
!
router ospf 1
 network 10.10.10.10 0.0.0.0 area 0
 network 10.10.10.0 0.0.0.255 area 1
 network 10.10.0.0 0.0.15.255 area 2
 network 0.0.0.0 255.255.255.255 area 3
```

Verification

```
R1#show ip ospf interface | include (Loopback[0-8]|Area)
Loopback0 is up, line protocol is up
  Internet Address 10.10.10.10/32, Area 0
Loopback2 is up, line protocol is up
  Internet Address 10.10.10.255/32, Area 1
Loopback1 is up, line protocol is up
  Internet Address 10.10.10.1/32, Area 1
Loopback4 is up, line protocol is up
  Internet Address 10.10.15.255/32, Area 2
Loopback3 is up, line protocol is up
  Internet Address 10.10.0.0/32, Area 2
Loopback8 is up, line protocol is up
  Internet Address 223.255.255.255/32, Area 3
Loopback7 is up, line protocol is up
  Internet Address 1.0.0.0/32, Area 3
Loopback6 is up, line protocol is up
  Internet Address 10.10.16.0/32, Area 3
Loopback5 is up, line protocol is up
  Internet Address 10.9.255.255/32, Area 3
```

Breakdown

The first step in configuring OSPF is to enable the routing process globally. This is accomplished by issuing the `router ospf` command, followed by a locally significant process-id in global configuration mode. Once the process has been enabled, the next step in configuring OSPF is to enable the process on an interface by issuing the `network` statement.

Like EIGRP, the OSPF network statement uses a wildcard mask to identify which interface or interfaces are running the protocol. In addition to this, the network statement determines which link state area an interface will belong to. The wildcard mask option of the network statement works like the wildcard mask option of an IP access list, where a binary zero indicates to check the particular bit position, while a binary one indicates to ignore the particular bit position. Usage of the wildcard mask in the network statement does not relate to the IP subnet mask of an interface or its corresponding advertisement, but is simply used to determine which interface or interfaces belong to the area.

This is demonstrated in the above example through an exercise in binary logic. The most specific match for an interface in the OSPF network statement uses a wildcard mask of all zeros. The statement `network 10.10.10.10 0.0.0.0 area 0` dictates that only the interface with the IP address 10.10.10.10 belongs to area 0. The next statement, `network 10.10.10.0 0.0.0.255 area 1`, dictates that any interface with an IP address of 10.10.10.X, where X is any number, belongs to area 1. Next, the statement `network 10.10.0.0 0.0.15.255 area 2` dictates that any interface with an IP address of 10.10.X.Y, where X is a number from 0-15 and Y is any number, belongs to area 2. Lastly, all other interfaces that aren't already assigned to an area are assigned to area 3 with the statement `network 0.0.0.0 255.255.255.255 area 3`.

When multiple OSPF network statement are configured with overlapping matches, as in the above case, the most specific network statement determines which area the interface will exist in. For example the above statement `network 0.0.0.0 255.255.255.255 area 3` dictates that all interfaces will participate in area 3. However, the more specific statement `network 10.10.10.10 0.0.0.0 area 0` dictates that this particular interface will belong to area 0.

To verify what area an interface is running in, issue the `show ip ospf interface` command in privilege level mode.

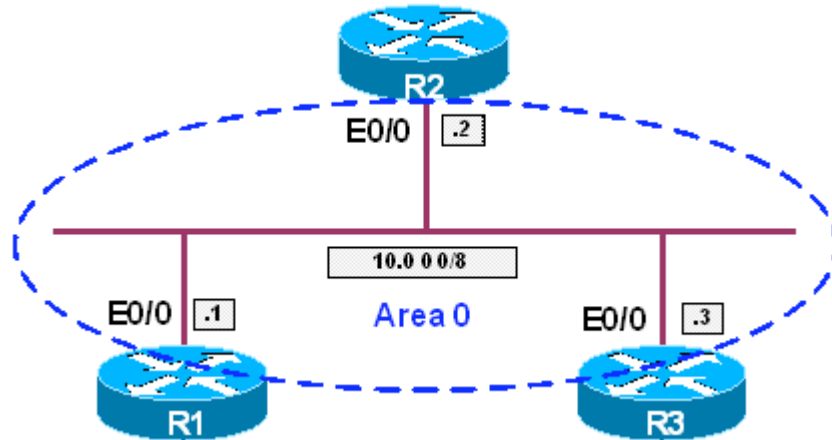


Recommended Reading

[OSPF Commands : network area](#)

OSPF DR/BDR Election

Objective: Configure OSPF on the Ethernet network between R1, R2 and R3 so that R1 is elected the DR, R2 is elected the BDR, and R3 does not participate in the election.



Directions

- Configure R1's interface Ethernet0/0 with the IP address 10.0.0.1/8
- Configure R2's interface Ethernet0/0 with the IP address 10.0.0.2/8
- Configure R3's interface Ethernet0/0 with the IP address 10.0.0.3/8
- Configure R1's interface Ethernet0/0 with an OSPF priority of 100
- Configure R2's interface Ethernet0/0 with an OSPF priority of 50
- Configure R3's interface Ethernet0/0 with an OSPF priority of 0
- Configure these interfaces in OSPF area 0

Ask Yourself

- What is the default OSPF network type on an Ethernet segment?
- Does this network type have a DR/BDR election?
- What is the role of the DR?
- What is the role of the BDR?
- How does the DR/BDR election occur?
- How can the DR/BDR election be influenced?
- What happens when a router with a higher priority than the DR or BDR is added to the segment?

Final Configuration

```
R1:
interface Ethernet0/0
 ip address 10.0.0.1 255.0.0.0
 ip ospf priority 100
!
router ospf 1
 network 10.0.0.1 0.0.0.0 area 0
```



```

R2:
interface Ethernet0/0
 ip address 10.0.0.2 255.0.0.0
 ip ospf priority 50
!
router ospf 1
 network 10.0.0.2 0.0.0.0 area 0

```

```

R3:
interface Ethernet0/0
 ip address 10.0.0.3 255.0.0.0
 ip ospf priority 0
!
router ospf 1
 network 10.0.0.3 0.0.0.0 area 0

```

Verification

```
R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.2	50	FULL/BDR	00:00:38	10.0.0.2	Ethernet0/0
10.0.0.3	0	FULL/DROTHER	00:00:34	10.0.0.3	Ethernet0/0

```
R2#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.3	0	FULL/DROTHER	00:00:33	10.0.0.3	Ethernet0/0
10.0.0.1	100	FULL/DR	00:00:31	10.0.0.1	Ethernet0/0

```
R3#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.2	50	FULL/BDR	00:00:34	10.0.0.2	Ethernet0/0
10.0.0.1	100	FULL/DR	00:00:39	10.0.0.1	Ethernet0/0

```
R1#show ip ospf interface ethernet0/0
```

```

Ethernet0/0 is up, line protocol is up
 Internet Address 10.0.0.1/8, Area 0
 Process ID 1, Router ID 10.0.0.1, Network Type BROADCAST, Cost: 10
 Transmit Delay is 1 sec, State DR, Priority 100
 Designated Router (ID) 10.0.0.1, Interface address 10.0.0.1
 Backup Designated router (ID) 10.0.0.2, Interface address 10.0.0.2

```

Breakdown

The OSPF Designated Router (DR) and Backup Designated Router (BDR) are roles assigned to devices on the OSPF network type broadcast and non-broadcast in order to reduce LSA replication on the segment. When a device on the segment wants to send an LSA into the network, the LSA is sent to DR. The DR is then responsible for sending this LSA back out to all other devices on the segment. The BDR is used to take the place of the DR in the case of a network failure. Determination of which devices will be the DR and BDR on the segment is based on an election process.

This election process is decided by two values, the interface priority and the router-id. The interface priority is a number in the range of 0 to 255, with the default being 1 on all interfaces. An OSPF priority of 0 indicates that a device will not participate in the DR/BDR election, while a priority of 255 indicates that the device is most likely to be elected DR. If there is a tie in the OSPF priority, the device with the highest `router-id` is elected.

OSPF router-id can either be automatically selected or manually configured. For automatic selection the router will choose the highest loopback IP address on the router, or if there are no loopback address, the highest IP address assigned to any other interface. When the router-id is left up to automatic selection it can change if a new interface with a higher IP address is assigned to the router and the OSPF process is restarted. As the router-id must be unique to the entire OSPF domain, it is recommended that it be manually defined with the `router-id` command under the OSPF process.

One of the problems that can occur with the OSPF DR/BDR election relates to the speed at which the OSPF process loads. The OSPF DR/BDR election, unlike the IS-IS DIS election, does not support preemption. This means is that once the DR and BDR for a segment are determined, a new election cannot occur even if a new device with a higher priority or router-id comes on to the segment. Instead, a new election only occurs if the DR or BDR fails. If the DR fails, the BDR assumes the role of the DR, and a new election occurs for the BDR. If the BDR fails, a new election occurs for the BDR.

In the above example the implication of this can be seen if R1 comes on to the segment after R2 and R3 have performed the election. Even though R2's priority is lower than R1, R1 can not preempt R2's DR status if it has already been elected. Therefore the only way to ensure that a particular device is elected the DR in all circumstances is to remove the other routers from the election by setting the OSPF priority to 0. This setup is common in cases where only one device has full layer 2 connectivity to all devices on the segment, such as the hub in a hub-and-spoke NBMA topology, in which case a failure of the hub would cause a loss of physical reachability between all other devices on the network. In such a case it is required to set the spoke devices to have a priority of 0, as the hub is the only device on the segment that can properly perform the duties of the DR.

Once the OSPF election has occurred, issue the `show ip ospf neighbor` and `show ip ospf interface` commands to verify who have been elected the DR and BDR.

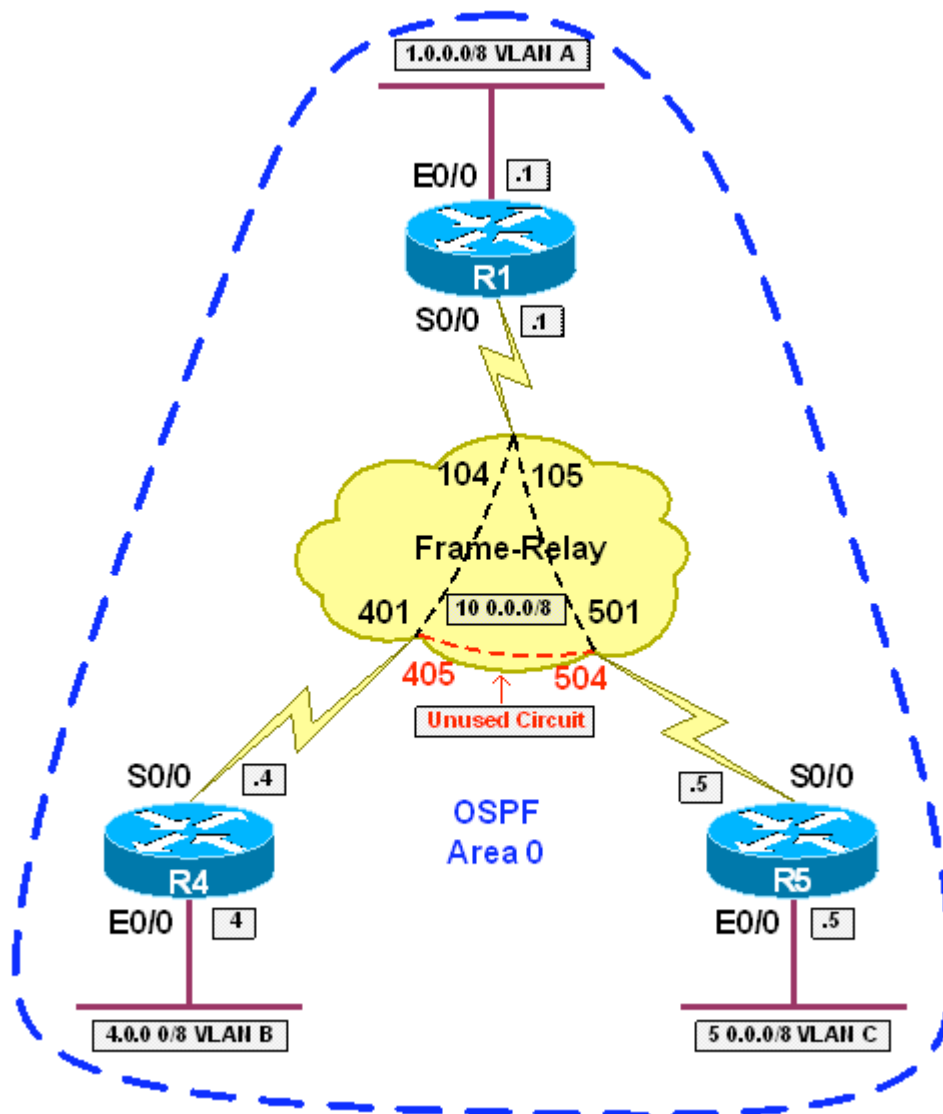


Recommended Reading

[IOS Command Reference: ip ospf priority](#)

OSPF over Frame Relay - Non-Broadcast

Objective: Configure OSPF over the Frame Relay network between R1, R4, and R5 using the OSPF network type non-broadcast to achieve connectivity between VLAN A, VLAN B, and VLAN C



Directions

- Configure R1's Ethernet interface with the IP address of 1.0.0.1/8
- Configure R4's Ethernet interface with the IP address of 4.0.0.4/8
- Configure R5's Ethernet interface with the IP address of 5.0.0.5/8
- Configure R1's main Serial interface with the IP address 10.0.0.1/8
- Configure R4's main Serial interface with the IP address 10.0.0.4/8
- Configure R5's main Serial interface with the IP address 10.0.0.5/8

- Configure Frame Relay encapsulation on R1, R4, and R5's main Serial interfaces
- Disable Frame-Relay Inverse-ARP on R4 and R5's main Serial interfaces
- Configure static layer 3 to layer 2 mapping on R1 to reach 10.0.0.4 via DLCI 104 and 10.0.0.5 via DLCI 105
- Configure static layer 3 to layer 2 mapping on R4 to reach 10.0.0.1 and 10.0.0.5 via DLCI 401
- Configure static layer 3 to layer 2 mapping on R5 to reach 10.0.0.1 and 10.0.0.4 via DLCI 501
- Configure R4 and R5's Serial interfaces with an OSPF priority of 0
- Configure OSPF on R1, R4, and R5 with all interfaces in area 0
- Configure neighbor statements on R1 for R4 and R5 under the OSPF process

Ask Yourself

- What is the default OSPF network type on a multipoint Frame Relay interface?
- Are LSAs for this network type sent as unicast or multicast?
- Does this network type support the DR/BDR election?
- How do these properties affect your configuration?

Final Configuration

```
R1:
interface Ethernet0/0
 ip address 1.0.0.1 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.1 255.0.0.0
 encapsulation frame-relay
 frame-relay map ip 10.0.0.4 104
 frame-relay map ip 10.0.0.5 105
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
 neighbor 10.0.0.4
 neighbor 10.0.0.5

R4:
interface Ethernet0/0
 ip address 4.0.0.4 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.4 255.0.0.0
 encapsulation frame-relay
 ip ospf priority 0
 frame-relay map ip 10.0.0.1 401
 frame-relay map ip 10.0.0.5 401
 no frame-relay inverse-arp
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
```

```

R5:
interface Ethernet0/0
 ip address 5.0.0.5 255.0.0.0
 !
interface Serial0/0
 ip address 10.0.0.5 255.0.0.0
 encapsulation frame-relay
 ip ospf priority 0
 frame-relay map ip 10.0.0.1 501
 frame-relay map ip 10.0.0.4 501
 no frame-relay inverse-arp
 !
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0

```

Verification

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.4	0	FULL/DROTHER	00:01:58	10.0.0.4	serial0/0
10.0.0.5	0	FULL/DROTHER	00:01:49	10.0.0.5	serial0/0

R4#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.1	1	FULL/DR	00:01:47	10.0.0.1	serial0/0

R5#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.1	1	FULL/DR	00:01:56	10.0.0.1	serial0/0

R1#show ip route

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
 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

Gateway of last resort is not set

```

C    1.0.0.0/8 is directly connected, Ethernet0/0
O    4.0.0.0/8 [110/74] via 10.0.0.4, 00:01:49, Serial0/0
O    5.0.0.0/8 [110/74] via 10.0.0.5, 00:01:49, Serial0/0
C    10.0.0.0/8 is directly connected, Serial0/0

```

R4#show ip route

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
 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

Gateway of last resort is not set

```

O    1.0.0.0/8 [110/74] via 10.0.0.1, 00:02:26, Serial0/0

```

```
C 4.0.0.0/8 is directly connected, Ethernet0/0
O 5.0.0.0/8 [110/74] via 10.0.0.5, 00:02:26, Serial0/0
C 10.0.0.0/8 is directly connected, Serial0/0

R5#show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       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

Gateway of last resort is not set

O 1.0.0.0/8 [110/74] via 10.0.0.1, 00:00:11, Serial0/0
O 4.0.0.0/8 [110/74] via 10.0.0.4, 00:00:11, Serial0/0
C 5.0.0.0/8 is directly connected, Ethernet0/0
C 10.0.0.0/8 is directly connected, Serial0/0
```

Breakdown

The default OSPF network type for a multipoint Frame Relay interface is non-broadcast. The OSPF network type non-broadcast transmits OSPF hello packets as unicast, and like the broadcast network type supports the DR/BDR election. As hello packets for this network type are unicast, devices on the segment must be manually defined on the DR and BDR with the `neighbor` statement under the OSPF process. Note that in the above example the `broadcast` keyword is not added to the `frame-relay map` commands. This is due to the fact that broadcast support is not needed since hello packets are sent as unicast, not multicast.

As seen in the previous example, the DR in OSPF is responsible for distributing LSAs to other devices on the segment. In order for this to occur, the DR must have direct layer two connectivity to all devices on the segment. In hub-and-spoke NBMA networks, this implies that the hub, and only the hub, should be candidate for the DR election. This is due to the fact that the hub is the only device that has direct layer two connectivity to all devices on the network.

As seen in the previous example, the OSPF DR election is based on the OSPF priority. An OSPF priority of 0 dictates that the router will not participate in the election, while an OSPF priority of 255 indicates that the router is most likely to be elected the DR. However, recall that the OSPF DR election does not support preemption, and therefore who wins the election can be dependent on the order that devices come on to the segment. This implies that even if the hub has a priority of 255 configured it is possible that another device on the segment could be elected the DR. Therefore, in order to ensure that the hub is always elected the DR, the OSPF priority must be set to 0 on the spoke routers.

To verify that adjacency has been established issue the `show ip ospf neighbor` command. To verify what the OSPF network type is and who the DR/BDR is issue the `show ip ospf interface` command.

Note that in the above `show ip route` output the next-hop values learned via OSPF do not mimic the underlying partially meshed layer 2 topology. In other words, the routing table on R4 says to use the next-hop value of 10.0.0.5 to reach 5.0.0.0, while in reality R4 must send the traffic towards R1 first. This implies that R4 must have direct layer 3 to layer 2 resolution for the address 10.0.0.5 out the Frame Relay interface, while R5 must likewise have direct layer 3 to layer 2 resolution for the address 10.0.0.4 out the Frame Relay interface.

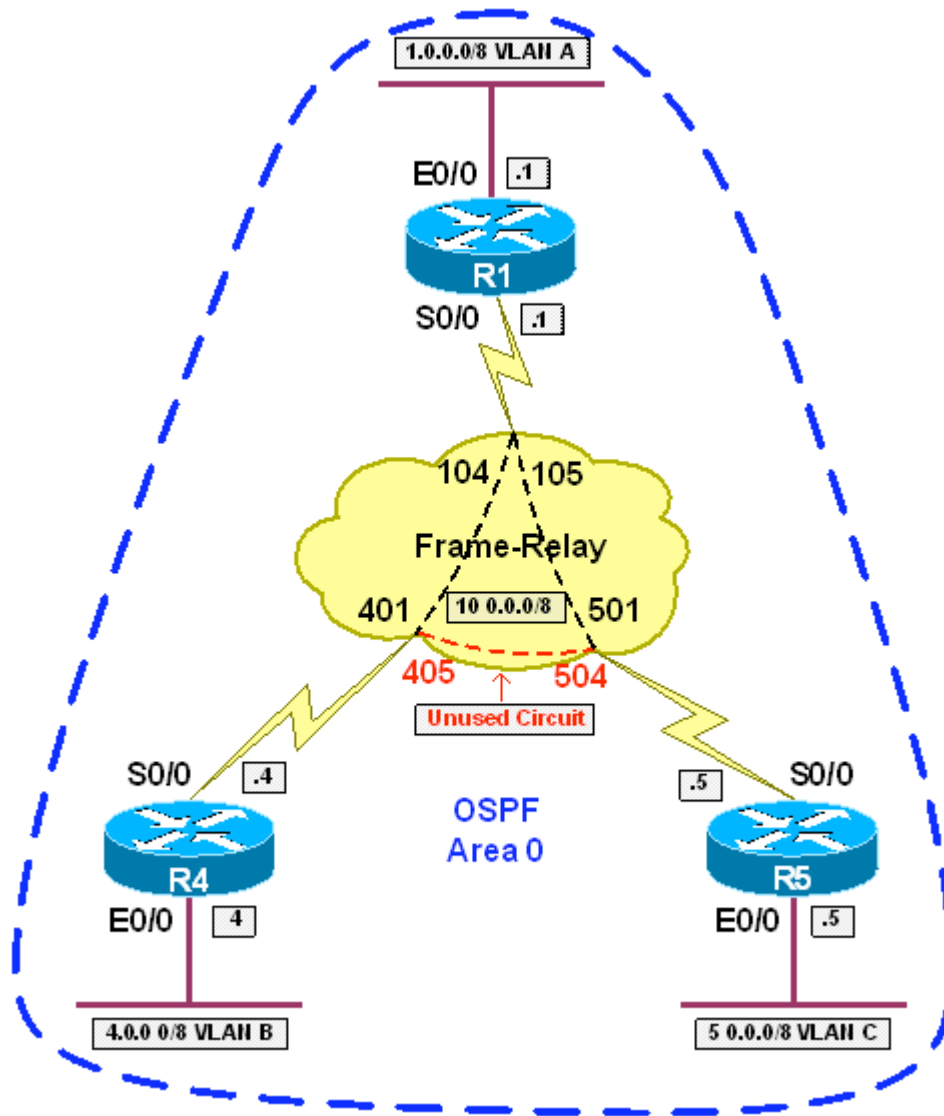


Recommended Reading

[Configuring OSPF: Configuring Your OSPF Network Type](#)

OSPF over Frame Relay - Broadcast

Objective: Configure OSPF over the Frame Relay network between R1, R4, and R5 using the OSPF network type broadcast to achieve connectivity between VLAN A, VLAN B, and VLAN C



Directions

- Configure R1's Ethernet interface with the IP address of 1.0.0.1/8
- Configure R4's Ethernet interface with the IP address of 4.0.0.4/8
- Configure R5's Ethernet interface with the IP address of 5.0.0.5/8
- Configure R1's main serial interface with the IP address 10.0.0.1/8
- Configure R4's main serial interface with the IP address 10.0.0.4/8
- Configure R5's main serial interface with the IP address 10.0.0.5/8
- Configure Frame Relay encapsulation on R1, R4, and R5's main Serial interfaces

- Disable Frame-Relay Inverse-ARP on R4 and R5's main Serial interfaces
- Configure static layer 3 to layer 2 mapping on R1 to reach 10.0.0.4 via DLCI 104 and 10.0.0.5 via DLCI 105
- Configure static layer 3 to layer 2 mapping on R4 to reach 10.0.0.1 and 10.0.0.5 via DLCI 401
- Configure static layer 3 to layer 2 mapping on R5 to reach 10.0.0.1 and 10.0.0.4 via DLCI 501
- Ensure that broadcast and multicast packets can be sent over these circuits
- Configure the OSPF network type broadcast on R1, R4, and R5's Serial interfaces
- Configure R4 and R5's Serial interfaces with an OSPF priority of 0
- Configure OSPF on R1, R4, and R5 with all interfaces in area 0

Ask Yourself

- How does the OSPF network type broadcast differ from the other OSPF network types?
- Are LSAs for this network type sent as unicast or multicast?
- Does this network type support the DR/BDR election?
- How do these properties affect your configuration?

Final Configuration

```
R1:
interface Ethernet0/0
 ip address 1.0.0.1 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.1 255.0.0.0
 encapsulation frame-relay
 ip ospf network broadcast
 frame-relay map ip 10.0.0.4 104 broadcast
 frame-relay map ip 10.0.0.5 105 broadcast
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0

R4:
interface Ethernet0/0
 ip address 4.0.0.4 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.4 255.0.0.0
 encapsulation frame-relay
 ip ospf network broadcast
 ip ospf priority 0
 frame-relay map ip 10.0.0.1 401 broadcast
 frame-relay map ip 10.0.0.5 401
 no frame-relay inverse-arp
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
```

```

R5:
interface Ethernet0/0
 ip address 5.0.0.5 255.0.0.0
 !
interface Serial0/0
 ip address 10.0.0.5 255.0.0.0
 encapsulation frame-relay
 ip ospf network broadcast
 ip ospf priority 0
 frame-relay map ip 10.0.0.1 501 broadcast
 frame-relay map ip 10.0.0.4 501
 no frame-relay inverse-arp
 !
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0

```

Verification

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.4	0	FULL/DROTHER	00:00:33	10.0.0.4	Serial0/0
10.0.0.5	0	FULL/DROTHER	00:00:36	10.0.0.5	Serial0/0

R4#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.1	1	FULL/DR	00:00:39	10.0.0.1	Serial0/0

R5#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.1	1	FULL/DR	00:00:33	10.0.0.1	Serial0/0

R1#show ip route

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
 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

Gateway of last resort is not set

```

C    1.0.0.0/8 is directly connected, Ethernet0/0
O    4.0.0.0/8 [110/74] via 10.0.0.4, 00:01:49, Serial0/0
O    5.0.0.0/8 [110/74] via 10.0.0.5, 00:01:49, Serial0/0
C    10.0.0.0/8 is directly connected, Serial0/0

```

R4#show ip route

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
 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

Gateway of last resort is not set

```
O 1.0.0.0/8 [110/74] via 10.0.0.1, 00:02:26, Serial0/0
C 4.0.0.0/8 is directly connected, Ethernet0/0
O 5.0.0.0/8 [110/74] via 10.0.0.5, 00:02:26, Serial0/0
C 10.0.0.0/8 is directly connected, Serial0/0

R5#show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       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

Gateway of last resort is not set

O 1.0.0.0/8 [110/74] via 10.0.0.1, 00:00:11, Serial0/0
O 4.0.0.0/8 [110/74] via 10.0.0.4, 00:00:11, Serial0/0
C 5.0.0.0/8 is directly connected, Ethernet0/0
C 10.0.0.0/8 is directly connected, Serial0/0
```

Breakdown

The OSPF network type broadcast transmits OSPF hellos packets as multicast, and like the non-broadcast network type supports the DR/BDR election. This network type is default on broadcast medias such as Ethernet and Token-Ring. Note that in the above example the broadcast keyword is added to the `frame-relay map` commands, as opposed to the previous example. This is due to the fact that broadcast support is needed since hello packets are sent as multicast, not unicast. Furthermore, since unicast updates do not need to be configured, the `neighbor` statement is not required on the DR and BDR.

As seen in the previous examples, the DR in OSPF is responsible for distributing LSAs to other devices on the segment. In order for this to occur, the DR must have direct layer two connectivity to all devices on the segment. In hub-and-spoke NBMA networks, this implies that the hub, and only the hub, should be candidate for the DR election. This is due to the fact that the hub is the only device that has direct layer two connectivity to all devices on the network.

As seen in the previous examples, the OSPF DR election is based on the OSPF priority. An OSPF priority of 0 dictates that the router will not participate in the election, while an OSPF priority of 255 indicates that the router is most likely to be elected the DR. However, recall that the OSPF DR election does not support preemption, and therefore who wins the election can be dependent on the order that devices come on to the segment. This implies that even if the hub has a priority of 255 configured it is possible that another device on the segment could be elected the DR. Therefore, in order to ensure that the hub is always elected the DR, the OSPF priority must be set to 0 on the spoke routers.

To verify that adjacency has been established issue the `show ip ospf neighbor` command. To verify what the OSPF network type is and who the DR/BDR is issue the `show ip ospf interface` command.

Note that in the above `show ip route` output the next-hop values learned via OSPF do not mimic the underlying partially meshed layer 2 topology. In other words, the routing table on R4 says to use the next-hop value of 10.0.0.5 to reach 5.0.0.0, while in reality R4 must send the traffic towards R1 first. This implies that R4 must have direct layer 3 to layer 2 resolution for the address 10.0.0.5 out the Frame Relay interface, while R5 must likewise have direct layer 3 to layer 2 resolution for the address 10.0.0.4 out the Frame Relay interface. Short explanation goes here.

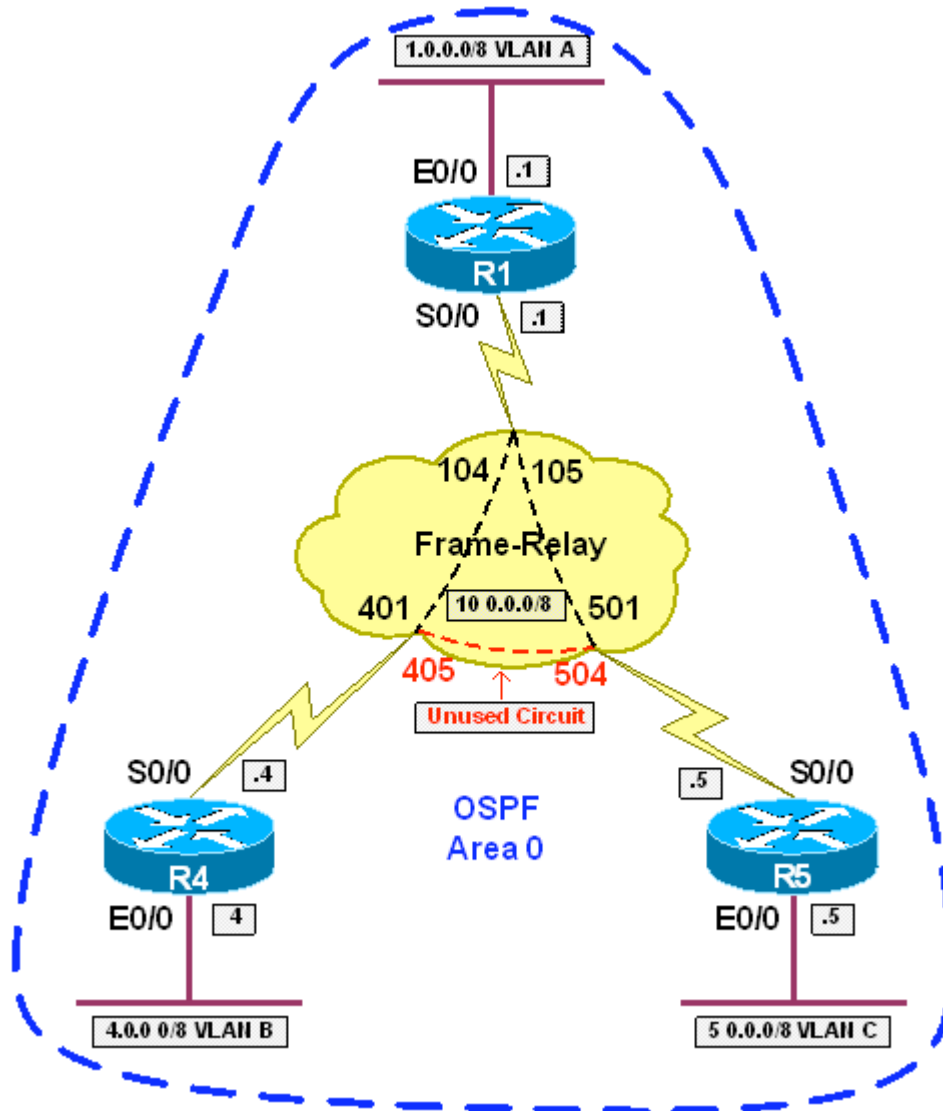


Recommended Reading

[Configuring OSPF: Configuring Your OSPF Network Type](#)

OSPF over Frame Relay - Point-to-Multipoint

Objective: Configure OSPF over the Frame Relay network between R1, R4, and R5 using the OSPF network type point-to-multipoint to achieve connectivity between VLAN A, VLAN B, and VLAN C. R4 and R5 should resolve each other's layer 3 addresses through layer 3 routing instead of static layer 2 mapping



Directions

- Configure R1's Ethernet interface with the IP address of 1.0.0.1/8
- Configure R4's Ethernet interface with the IP address of 4.0.0.4/8
- Configure R5's Ethernet interface with the IP address of 5.0.0.5/8

- Configure R1's main serial interface with the IP address 10.0.0.1/8
- Configure R4's main serial interface with the IP address 10.0.0.4/8
- Configure R5's main serial interface with the IP address 10.0.0.5/8
- Configure Frame Relay encapsulation on R1, R4, and R5's main serial interfaces
- Disable Frame Relay Inverse-ARP on R4 and R5's main serial interfaces
- Configure static layer 3 to layer 2 mapping on R1 to reach 10.0.0.4 via DLCI 104 and 10.0.0.5 via DLCI 105.
- Configure static layer 3 to layer 2 mapping on R4 to reach 10.0.0.1 via DLCI 401
- Configure static layer 3 to layer 2 mapping on R5 to reach 10.0.0.1 via DLCI 501
- Do not configure static layer 3 to layer 2 mappings between R4 and R5
- Ensure that broadcast and multicast traffic can be sent over these circuits
- Configure the OSPF network type on R1, R4, and R5's main serial interfaces as point-to-multipoint
- Configure OSPF on R1, R4, and R5 with all interfaces in area 0

Ask Yourself

- What is the default OSPF network type on the Frame Relay network?
- How does OSPF network type point-to-multipoint differ from the other network types?
- Does OSPF network type point-to-multipoint transmit LSAs as multicast or unicast?
- How does this affect layer 3 to layer 2 resolution?
- How does the next-hop processing of OSPF network type point-to-multipoint differ from the broadcast and non-broadcast network types?
- How does this affect the need for layer 3 to layer 2 resolution?

Final Configuration

```
R1:
interface Ethernet0/0
 ip address 1.0.0.1 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.1 255.0.0.0
 encapsulation frame-relay
 ip ospf network point-to-multipoint
 frame-relay map ip 10.0.0.4 104 broadcast
 frame-relay map ip 10.0.0.5 105 broadcast
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0

R4:
interface Ethernet0/0
 ip address 4.0.0.4 255.0.0.0
!
interface Serial0/0
```

```

ip address 10.0.0.4 255.0.0.0
encapsulation frame-relay
ip ospf network point-to-multipoint
frame-relay map ip 10.0.0.1 401 broadcast
no frame-relay inverse-arp
!
router ospf 1
  network 0.0.0.0 255.255.255.255 area 0

R5:
interface Ethernet0/0
  ip address 5.0.0.5 255.0.0.0
!
interface Serial0/0
  ip address 10.0.0.5 255.0.0.0
  encapsulation frame-relay
  ip ospf network point-to-multipoint
  frame-relay map ip 10.0.0.1 501 broadcast
  no frame-relay inverse-arp
!
router ospf 1
  network 0.0.0.0 255.255.255.255 area 0

```

Verification

R1#show frame-relay map

```

Serial0 (up): ip 10.0.0.4 dlci 104(0x68,0x1880), static,
              broadcast,
              CISCO, status defined, active
Serial0 (up): ip 10.0.0.5 dlci 105(0x69,0x1890), static,
              broadcast,
              CISCO, status defined, active

```

R4#show frame-relay map

```

Serial0 (up): ip 10.0.0.1 dlci 401(0x191,0x6410), static,
              broadcast,
              CISCO, status defined, active

```

R5#show frame-relay map

```

Serial0 (up): ip 10.0.0.1 dlci 501(0x191,0x6410), static,
              broadcast,
              CISCO, status defined, active

```

R1#show ip ospf neighbor

Neighbor ID	Pri	State		Dead Time	Address	Interface
10.0.0.5	1	FULL/	-	00:01:49	10.0.0.5	Serial0/0
10.0.0.4	1	FULL/	-	00:01:49	10.0.0.4	Serial0/0

R4#show ip ospf neighbor

Neighbor ID	Pri	State		Dead Time	Address	Interface
10.0.0.1	1	FULL/	-	00:01:47	10.0.0.1	Serial0/0

R5#show ip ospf neighbor

Neighbor ID	Pri	State		Dead Time	Address	Interface
10.0.0.1	1	FULL/	-	00:01:56	10.0.0.1	Serial0/0

R1#show ip route


```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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
```

Gateway of last resort is not set

```
C 1.0.0.0/8 is directly connected, Ethernet0
O 4.0.0.0/8 [110/74] via 10.0.0.4, 00:55:30, Serial0
O 5.0.0.0/8 [110/74] via 10.0.0.5, 00:55:30, Serial0
 10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C 10.0.0.0/8 is directly connected, Serial0
O 10.0.0.4/32 [110/64] via 10.0.0.4, 00:55:30, Serial0
O 10.0.0.5/32 [110/64] via 10.0.0.5, 00:55:30, Serial0
```

R4#show ip route

```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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
```

Gateway of last resort is not set

```
O 1.0.0.0/8 [110/74] via 10.0.0.1, 02:32:17, Serial0
C 4.0.0.0/8 is directly connected, Ethernet0
O 5.0.0.0/8 [110/138] via 10.0.0.1, 02:32:17, Serial0
 10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C 10.0.0.0/8 is directly connected, Serial0
O 10.0.0.1/32 [110/64] via 10.0.0.1, 02:32:17, Serial0
O 10.0.0.5/32 [110/128] via 10.0.0.1, 02:32:17, Serial0
```

R5#show ip route

```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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
```

Gateway of last resort is not set

```
O 1.0.0.0/8 [110/74] via 10.0.0.1, 00:44:06, Serial0
O 4.0.0.0/8 [110/138] via 10.0.0.1, 00:44:06, Serial0
C 5.0.0.0/8 is directly connected, Ethernet0
 10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C 10.0.0.0/8 is directly connected, Serial0
O 10.0.0.1/32 [110/64] via 10.0.0.1, 00:44:06, Serial0
O 10.0.0.4/32 [110/128] via 10.0.0.1, 00:44:06, Serial0
```

R4#ping 10.0.0.5

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.0.0.5, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 116/116/120 ms

```
R4#ping 5.0.0.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.0.0.5, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/116/120 ms

R1#sh ip ospf interface
Serial0 is up, line protocol is up
 Internet Address 10.0.0.1/8, Area 0
  Process ID 1, Router ID 10.0.0.1, Network Type POINT_TO_MULTIPOINT,
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 10.0.0.4
    Adjacent with neighbor 10.0.0.5
```

Breakdown

Unlike the OSPF network types of broadcast and non-broadcast, OSPF network type point-to-multipoint does not support the DR and BDR election. As with the network types broadcast and point-to-point, point-to-multipoint sends hello packets to the multicast address 224.0.0.5. Therefore, broadcast resolution is required on NBMA networks on which OSPF point-to-multipoint adjacencies need to be formed. As the default OSPF network type of multipoint Frame Relay interfaces is non-broadcast, point-to-multipoint must be manually configured with the interface level command `ip ospf network point-to-multipoint`.

The functional difference between OSPF network type point-to-multipoint versus broadcast and non-broadcast network types is how point-to-multipoint deals with next-hop resolution on a non-broadcast media. OSPF network type point-to-multipoint treats the network as a collection of point-to-point links instead of one flat broadcast network. With the network types non-broadcast and broadcast, OSPF does not understand that the underlying layer 2 topology may not mirror a flat layer 3 network. With OSPF network types broadcast and non-broadcast, next hop values are not modified when updates are transmitted across an NBMA media. This implies that devices on the NBMA cloud require layer 3 to layer 2 resolution for any endpoint injecting routes into the network. With OSPF network type point-to-multipoint, next hop values are modified to the address of the directly connected neighbor when they are advertised across the NBMA cloud. This implies that routers on the NBMA network only need layer 3 to layer 2 resolution for directly connected neighbors when running OSPF network point-to-multipoint.

In the above scenario R1, R4, and R5 are configured in a partially meshed hub-and-spoke topology. In order for traffic to pass from R4 to R5 and vice-versa, it must transit R1. When OSPF network type point-to-multipoint is configured on this topology, the next-hop value of routes learned from R4 is modified to 10.0.0.1 (R1's interface address) when they are advertised out to R5. Likewise

the next-hop value of routes learned from R5 is modified to 10.0.0.1 (R1's interface address) when they are advertised out to R4. This implies that when route recursion for the final destination occurs on R4 or R5, packets sent to the Frame Relay interface will be encapsulated towards 10.0.0.1. Therefore, layer 3 to layer 2 resolution is not required between R4 and R5 as the layer 3 routing domain tells them to use R1 to get to all destinations learned over the Frame Relay network.

If the OSPF network type in this scenario were broadcast or non-broadcast, as seen in previous examples, R4 would use the next-hop value of 10.0.0.5 to reach routes originated by R5, and R5 would use the next-hop value of 10.0.0.4 to reach routes originated by R4. Therefore R4 and R5 would need layer 3 to layer 2 resolution for each others addresses.

In addition to this next-hop modification, OSPF network type point-to-multipoint advertises the endpoints of the point-to-multipoint network as host routes instead of the actual network itself. This is shown in the above example from the `show ip route` output from R1, R4, and R5. This behavior allows non-directly connected devices to use layer 3 routing to reach each other as opposed to layer 3 to layer 2 resolution. This is seen in the above output as R4 has connectivity to R5's interface address 10.0.0.5 without having layer 3 to layer 2 mapping for it. Instead this resolution occurs at layer 3. Note that when this network is advertised out to the rest of the OSPF domain only the host routes are sent. Therefore everyone else in the OSPF network will see only 10.0.0.1/32, 10.0.0.4/32, and 10.0.0.5/32 to represent the network 10.0.0.0/8.

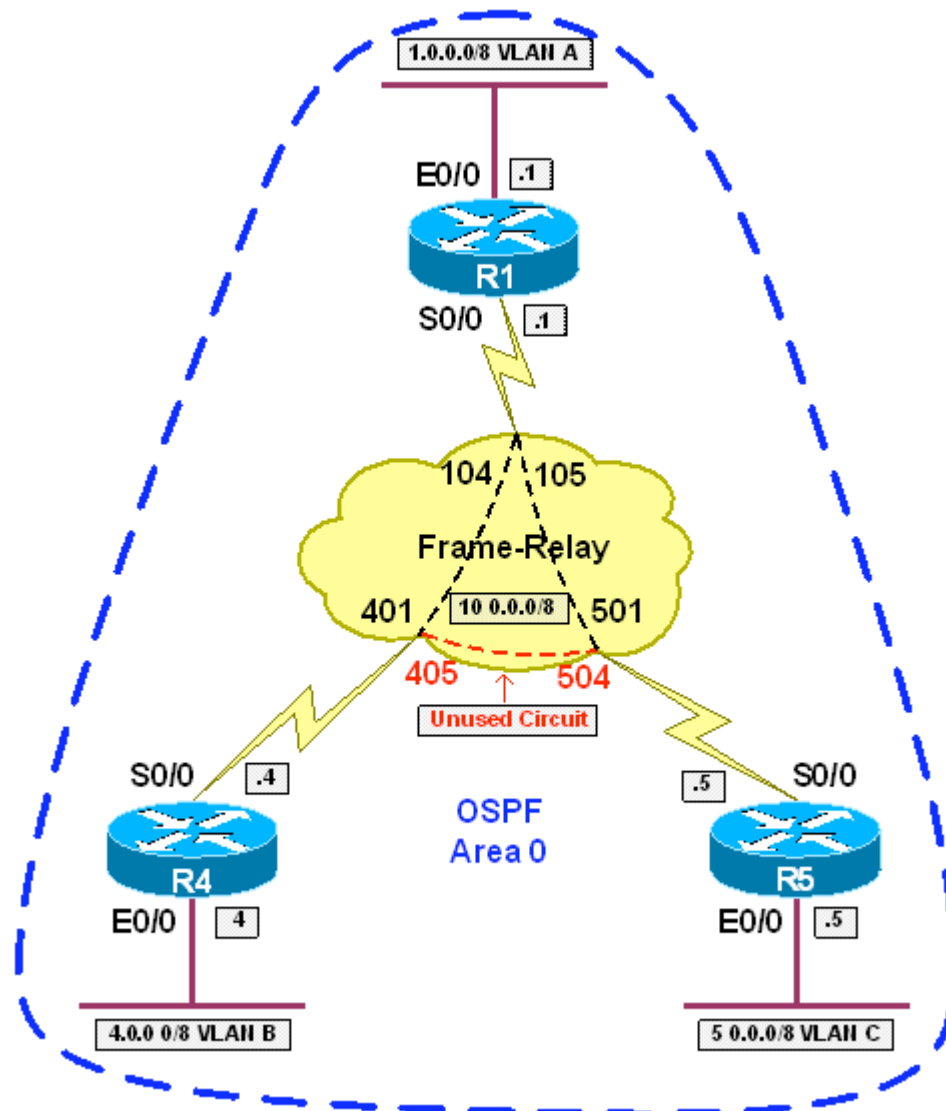


Recommended Reading

[Configuring OSPF: Configuring Your OSPF Network Type](#)

OSPF over Frame Relay - Point-to-Multipoint Non-Broadcast

Objective: Configure OSPF over the Frame Relay network between R1, R4, and R5 using the OSPF network type point-to-multipoint non-broadcast to achieve connectivity between VLAN A, VLAN B, and VLAN C. Do not enable broadcast support on any circuit



Directions

- Configure R1's Ethernet interface with the IP address of 1.0.0.1/8
- Configure R4's Ethernet interface with the IP address of 4.0.0.4/8
- Configure R5's Ethernet interface with the IP address of 5.0.0.5/8
- Configure R1's main serial interface with the IP address 10.0.0.1/8
- Configure R4's main serial interface with the IP address 10.0.0.4/8
- Configure R5's main serial interface with the IP address 10.0.0.5/8

- Configure Frame Relay encapsulation on R1, R4, and R5's main serial interfaces
- Disable Frame Relay Inverse-ARP on R4 and R5's main serial interfaces
- Configure static layer 3 to layer 2 mapping on R1 to reach 10.0.0.4 via DLCI 104 and 10.0.0.5 via DLCI 105
- Configure static layer 3 to layer 2 mapping on R4 to reach 10.0.0.1 via DLCI 401
- Configure static layer 3 to layer 2 mapping on R5 to reach 10.0.0.1 via DLCI 501
- Do not configure static layer 3 to layer 2 mappings between R4 and R5
- Configure the OSPF network type on R1, R4, and R5's main serial interfaces as point-to-multipoint non-broadcast
- Configure OSPF on R1, R4, and R5 with all interfaces in area 0
- Configure neighbor statements on R1 for R4 and R5 under the OSPF process

Ask Yourself

- What is the difference between OSPF network type point-to-multipoint non-broadcast and the other OSPF network types?
- In what way is OSPF network type point-to-multipoint non-broadcast like OSPF network type point-to-multipoint?
- How does this affect layer 3 to layer 2 resolution?
- In what way is OSPF network type point-to-multipoint non-broadcast like OSPF network non-broadcast?
- How does this affect layer 3 to layer 2 resolution?

Final Configuration

```
R1:
interface Ethernet0/0
 ip address 1.0.0.1 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.1 255.0.0.0
 encapsulation frame-relay
 ip ospf network point-to-multipoint non-broadcast
 frame-relay map ip 10.0.0.4 104
 frame-relay map ip 10.0.0.5 105
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
 neighbor 10.0.0.5
 neighbor 10.0.0.4

R4:
interface Ethernet0/0
 ip address 4.0.0.4 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.4 255.0.0.0
 encapsulation frame-relay
```

```

ip ospf network point-to-multipoint non-broadcast
frame-relay map ip 10.0.0.1 401
no frame-relay inverse-arp
!
router ospf 1
network 0.0.0.0 255.255.255.255 area 0

```

R5:

```

interface Ethernet0/0
ip address 5.0.0.5 255.0.0.0
!
interface Serial0/0
ip address 10.0.0.5 255.0.0.0
encapsulation frame-relay
ip ospf network point-to-multipoint non-broadcast
frame-relay map ip 10.0.0.1 501
no frame-relay inverse-arp
!
router ospf 1
network 0.0.0.0 255.255.255.255 area 0

```

Verification

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.5	1	FULL/ -	00:01:49	10.0.0.5	Serial0/0
10.0.0.4	1	FULL/ -	00:01:49	10.0.0.4	Serial0/0

R4#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.1	1	FULL/ -	00:01:47	10.0.0.1	Serial0/0

R5#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.0.0.1	1	FULL/ -	00:01:56	10.0.0.1	Serial0/0

R1#show ip route

```

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
i - IS-IS, 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

```

Gateway of last resort is not set

```

C    1.0.0.0/8 is directly connected, Ethernet0
O    4.0.0.0/8 [110/74] via 10.0.0.4, 00:56:28, Serial0
O    5.0.0.0/8 [110/74] via 10.0.0.5, 00:56:28, Serial0
    10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C    10.0.0.0/8 is directly connected, Serial0
O    10.0.0.4/32 [110/64] via 10.0.0.4, 00:56:28, Serial0
O    10.0.0.5/32 [110/64] via 10.0.0.5, 00:56:28, Serial0

```

R4#show ip route

```

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
i - IS-IS, 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
```

Gateway of last resort is not set

```
O 1.0.0.0/8 [110/74] via 10.0.0.1, 00:00:47, Serial0
C 4.0.0.0/8 is directly connected, Ethernet0
O 5.0.0.0/8 [110/138] via 10.0.0.1, 00:00:47, Serial0
  10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C 10.0.0.0/8 is directly connected, Serial0
O 10.0.0.1/32 [110/64] via 10.0.0.1, 00:00:47, Serial0
O 10.0.0.5/32 [110/128] via 10.0.0.1, 00:00:47, Serial0
```

R5#show ip route

```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
i - IS-IS, 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
```

Gateway of last resort is not set

```
O 1.0.0.0/8 [110/74] via 10.0.0.1, 00:58:35, Serial0
O 4.0.0.0/8 [110/138] via 10.0.0.1, 00:58:35, Serial0
C 5.0.0.0/8 is directly connected, Ethernet0
  10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C 10.0.0.0/8 is directly connected, Serial0
O 10.0.0.1/32 [110/64] via 10.0.0.1, 00:58:35, Serial0
O 10.0.0.4/32 [110/128] via 10.0.0.1, 00:58:35, Serial0
```

R4#ping 10.0.0.5

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.0.0.5, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 116/118/128 ms

R4#ping 5.0.0.5

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 5.0.0.5, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 116/118/128 ms

Breakdown

OSPF network-type point-to-multipoint non-broadcast is similar to the network type point-to-multipoint except that OSPF packets are sent as unicast instead of multicast. This implies that circuits used to form OSPF point-to-multipoint non-broadcast adjacencies do not need to support broadcast or multicast transmission for an OSPF adjacency to form.

Like OSPF network type point-to-multipoint, point-to-multipoint non-broadcast does not support the DR/BDR election, and possesses the same next-hop processing behavior. This implies that layer 3 to layer 2 resolution is only required between directly connected devices on the NBMA network, and that IP routing will be used to establish reachability between devices that are not adjacent at layer 2.




Recommended Reading

[Configuring OSPF: Configuring Your OSPF Network Type](#)

OSPF over Frame Relay - Point-to-Point

Objective: Configure OSPF over the Frame Relay network between R1 and R2 using the OSPF network type point-to-point to achieve connectivity between VLAN A and VLAN B



Directions

- Configure R1's Ethernet interface with the IP address of 1.0.0.1/8
- Configure R2's Ethernet interface with the IP address of 2.0.0.2/8
- Configure R1's Serial interface with the IP address of 10.0.0.1/8
- Configure R2's Serial interface with the IP address of 10.0.0.2/8
- Configure the Frame Relay circuit between R1 and R2 using static layer 3 to layer 2 resolution
- Ensure that broadcast and multicast packets can be sent across the circuit
- Configure the OSPF network type point-to-point on the Serial interfaces of R1 and R2
- Configure OSPF on R1 and R2 with all interfaces in area 0

Ask Yourself

- What is the default OSPF network type on a multipoint Frame Relay interface?
- How does the point-to-point network type differ from the other OSPF network types?
- Are hello packets for this network type sent as unicast or multicast?
- Does this network type support the DR/BDR election? Why or why not?

Final Configuration

```
R1:
interface Ethernet0/0
 ip address 1.0.0.1 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.1 255.0.0.0
 encapsulation frame-relay
 ip ospf network point-to-point
 frame-relay map ip 10.0.0.2 102 broadcast
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
```

R2:

```
interface Ethernet0/0
 ip address 2.0.0.2 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.2 255.0.0.0
 encapsulation frame-relay
 ip ospf network point-to-point
 frame-relay map ip 10.0.0.1 201 broadcast
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
```

```
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
```

www.InternetNetworkExpert.com

```

interface Ethernet0/0
 ip address 2.0.0.2 255.0.0.0
!
interface Serial0/0
 ip address 10.0.0.2 255.0.0.0
 encapsulation frame-relay
 ip ospf network point-to-point
 frame-relay map ip 10.0.0.1 201 broadcast
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0

```

Verification

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface	
10.0.0.2	1	FULL/	-	00:00:33	10.0.0.2	Serial0/0

R2#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface	
10.0.0.1	1	FULL/	-	00:00:29	10.0.0.1	Serial0/0

R2#show ip route

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
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

Gateway of last resort is not set

```

O   1.0.0.0/8 [110/74] via 10.0.0.1, 00:01:08, Serial0
C   2.0.0.0/8 is directly connected, Ethernet0
C   10.0.0.0/8 is directly connected, Serial0

```

R1#show ip route

Codes: C - connected, S - static, I - IGRP, 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, E - EGP
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

Gateway of last resort is not set

```

C   1.0.0.0/8 is directly connected, Ethernet0
O   2.0.0.0/8 [110/74] via 10.0.0.2, 00:00:10, Serial0
C   10.0.0.0/8 is directly connected, Serial0
R1#

```

Breakdown

OSPF network type point-to-point is used for adjacencies between exactly two OSPF neighbors. Network type point-to-point uses multicast for transmission, does not support the DR/BDR election, and is the default network type for point-to-point NBMA medias such as ISDN and point-to-point subinterfaces in Frame Relay and ATM. Since this network type uses multicast for transmission the `broadcast` keyword has been added to the `frame-relay map` statement between R1 and R2.

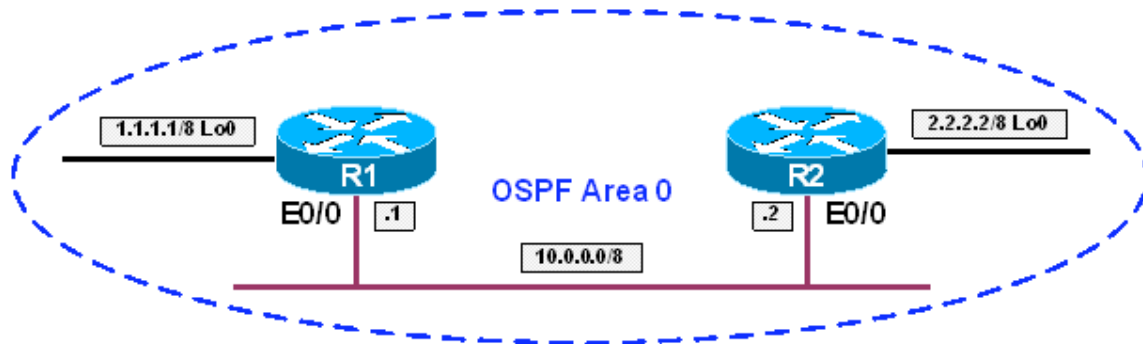


Recommended Reading

[Configuring OSPF: Configuring Your OSPF Network Type](#)

OSPF Network Type Loopback

Objective: Configure OSPF area 0 on the Ethernet and Loopback0 interfaces of R1 and R2. The Loopback networks should appear in the routing table with a mask of /8



Directions

- Configure R1 Ethernet0 with an IP address of 10.0.0.1/8
- Configure R2 Ethernet0 with an IP address of 10.0.0.2/8
- Configure R1 Loopback0 with an IP address of 1.1.1.1/8
- Configure R2 Loopback0 with an IP address of 2.2.2.2/8
- Configure all interfaces on R1 and R2 in OSPF area 0
- Examine the routing tables on R1 and R2 and note the mask of the OSPF routes to the Loopback interfaces
- Configure the OSPF network type point-to-point on R1 and R2's Loopback0 interfaces

Ask Yourself

- What is the default OSPF network type of a Loopback interface?
- How does this differ from the other network types?
- What type of behavior does this interface exhibit in the context of OSPF?
- Can this behavior be disabled?

Final Configuration

```
R1:
interface Loopback0
 ip address 1.1.1.1 255.0.0.0
 ip ospf network point-to-point
!
interface Ethernet0/0
 ip address 10.0.0.1 255.0.0.0
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
```

```

R2:
interface Loopback0
 ip address 2.2.2.2 255.0.0.0
 ip ospf network point-to-point
!
interface Ethernet0
 ip address 10.0.0.2 255.0.0.0
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0

```

Verification

Before point-to-point configuration

```

R1#show ip ospf interface loopback0
Loopback0 is up, line protocol is up
 Internet Address 1.1.1.1/8, Area 0
 Process ID 1, Router ID 1.1.1.1, Network Type LOOPBACK, Cost: 1
 Loopback interface is treated as a stub Host

```

```

R1#show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

```

Gateway of last resort is not set

```

C    1.0.0.0/8 is directly connected, Loopback0
    2.0.0.0/32 is subnetted, 1 subnets
O    2.2.2.2 [110/11] via 10.0.0.2, 00:08:59, Ethernet0
C    10.0.0.0/8 is directly connected, Ethernet0

```

```

R2#show ip route
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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

```

Gateway of last resort is not set

```

    1.0.0.0/32 is subnetted, 1 subnets
O    1.1.1.1 [110/11] via 10.0.0.1, 00:07:28, Ethernet0
C    2.0.0.0/8 is directly connected, Loopback0
C    10.0.0.0/8 is directly connected, Ethernet0

```

After point-to-point configuration

```

R1#show ip ospf interface loopback0
Loopback0 is up, line protocol is up
 Internet Address 1.1.1.1/8, Area 0
 Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 1

```

```
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  oob-resync timeout 40
  No Hellos (Passive interface)
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)
```

R1#show ip route

```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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
```

Gateway of last resort is not set

```
C    1.0.0.0/8 is directly connected, Loopback0
O    2.0.0.0/8 [110/11] via 10.0.0.2, 00:00:53, Ethernet0
C    10.0.0.0/8 is directly connected, Ethernet0
```

R2#show ip route

```
Codes: C - connected, S - static, I - IGRP, 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, E - EGP
       i - IS-IS, 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
```

Gateway of last resort is not set

```
O    1.0.0.0/8 [110/11] via 10.0.0.1, 00:01:47, Ethernet0
C    2.0.0.0/8 is directly connected, Loopback0
C    10.0.0.0/8 is directly connected, Ethernet0
```

Breakdown

The OSPF network type loopback is a special case reserved for Loopback or looped back interfaces. Network type loopback is the default network type on loopback interfaces, and is not manually configurable with the `ip ospf network command`. The loopback network type treats the interface as a stub host and injects the IP address of the interface as host route into the OSPF domain regardless of the actual subnet mask of the interface. Manually setting the OSPF network type to point-to-point disables this special treatment, and causes OSPF to advertise the IP address of the interface with the native subnet mask as configured on the interface.

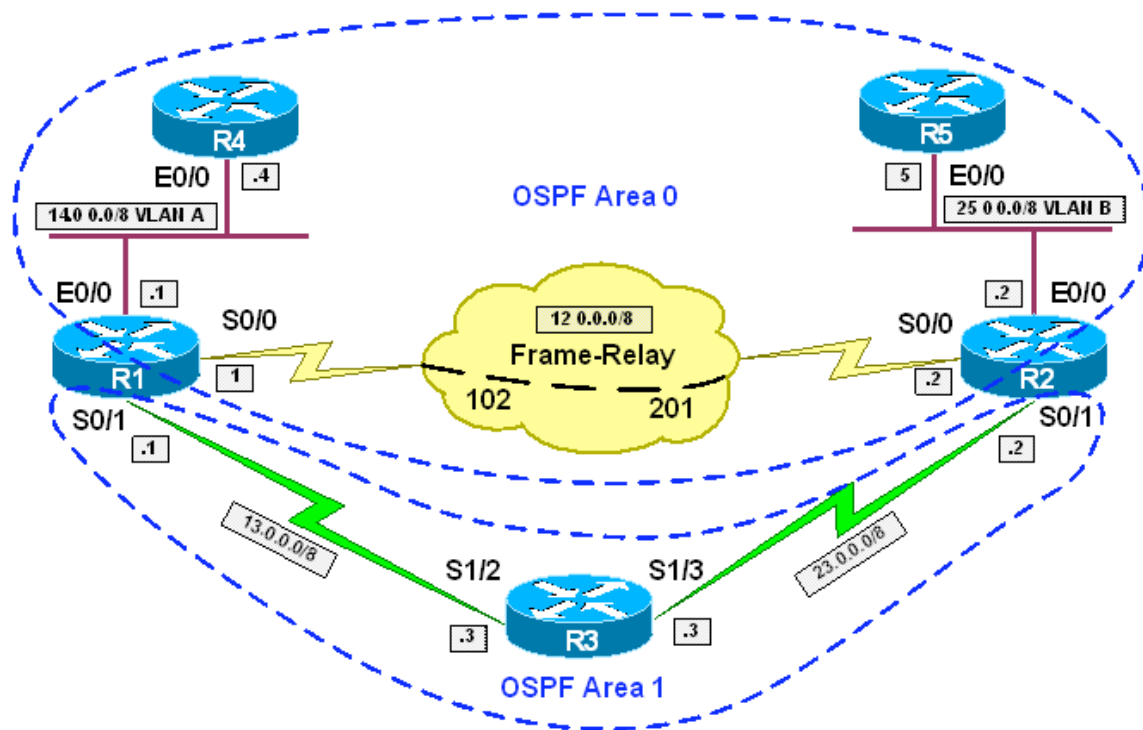


Recommended Reading

[RFC 2328: OSPF version 2 - Section 9.1 Interface states: Loopback](#)

OSPF Virtual Links - Repairing Area 0

Objective: Configure the OSPF domain between R1, R2, R3, R4, and R5 per the diagram. Ensure that connectivity is maintained between R4 and R5 if the link between R1 and R2 is down



Directions

- Configure the IP addressing on R1, R2, R3, R4, and R5 per the above diagram
- Configure the Frame Relay circuit between R1 and R2 using static layer 3 to layer 2 mappings
- Configure OSPF area 0 on VLAN A between R1 and R4
- Configure OSPF area 0 on VLAN B between R2 and R5
- Configure OSPF area 0 on the Frame Relay segment between R1 and R2 using OSPF network type point-to-point
- Configure OSPF area 1 on the Serial links between R1 & R3 and R2 & R3
- Statically configure the OSPF router-ids on R1 and R2 as the IP address of their respective Ethernet interfaces
- Configure an OSPF virtual link between R1 and R2 using area 1 as transit

Ask Yourself

- Why must all areas in OSPF be connected to area 0?

- What happens when an area is not connected to area 0?
- Why must area 0 be contiguous throughout the OSPF domain?
- What happens when area 0 becomes partitioned?
- What are the methods of repairing a partitioned area 0?
- When configuring an OSPF virtual-link, which routers should it terminate on?
- What does the IP address used in the virtual-link statement refer to?
- Must this address be reachable before the virtual-link can be established?

Final Configuration

```
R1:
interface Ethernet0/0
 ip address 14.0.0.1 255.0.0.0
!
interface Serial0/0
 ip address 12.0.0.1 255.0.0.0
 encapsulation frame-relay
 ip ospf network point-to-point
 frame-relay map ip 12.0.0.2 102 broadcast
!
interface Serial0/1
 ip address 13.0.0.1 255.0.0.0
!
router ospf 1
 router-id 14.0.0.1
 area 1 virtual-link 2.0.0.2
 network 12.0.0.1 0.0.0.0 area 0
 network 13.0.0.1 0.0.0.0 area 1
 network 14.0.0.1 0.0.0.0 area 0
```

```
R2:
interface Ethernet0/0
 ip address 25.0.0.2 255.0.0.0
!
interface Serial0/0
 ip address 12.0.0.2 255.0.0.0
 encapsulation frame-relay
 ip ospf network point-to-point
 frame-relay map ip 12.0.0.1 201 broadcast
!
interface Serial0/1
 ip address 23.0.0.2 255.0.0.0
!
router ospf 1
 router-id 25.0.0.2
 area 1 virtual-link 1.0.0.1
 network 12.0.0.2 0.0.0.0 area 0
 network 23.0.0.2 0.0.0.0 area 1
 network 25.0.0.2 0.0.0.0 area 0
```

```
R3:
interface Ethernet0/0
 ip address 3.0.0.3 255.0.0.0
!
interface Serial1/2
 ip address 13.0.0.3 255.0.0.0
 clock rate 64000
!
```

```

interface Serial1/3
 ip address 23.0.0.3 255.0.0.0
 clock rate 64000
 !
router ospf 1
 network 3.0.0.3 0.0.0.0 area 1
 network 13.0.0.3 0.0.0.0 area 1
 network 23.0.0.3 0.0.0.0 area 1

```

R4:

```

interface Ethernet0/0
 ip address 14.0.0.4 255.0.0.0
 !
router ospf 1
 network 14.0.0.4 0.0.0.0 area 0

```

R5:

```

interface Ethernet0/0
 ip address 25.0.0.5 255.0.0.0
 !
router ospf 1
 network 25.0.0.5 0.0.0.0 area 0

```

Verification

Before Virtual Link

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
14.0.0.4	1	FULL/DR	00:00:39	14.0.0.4	Ethernet0/0
25.0.0.2	0	FULL/ -	00:00:39	12.0.0.2	Serial0/0
23.0.0.3	0	FULL/ -	00:00:30	13.0.0.3	Serial0/1

R2#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
25.0.0.5	1	FULL/DR	00:00:38	25.0.0.5	Ethernet0/0
14.0.0.1	0	FULL/ -	00:00:38	12.0.0.1	Serial0/0
23.0.0.3	0	FULL/ -	00:00:34	23.0.0.3	Serial0/1

R3#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
25.0.0.2	0	FULL/ -	00:00:34	23.0.0.2	Serial1/3
14.0.0.1	0	FULL/ -	00:00:36	13.0.0.1	Serial1/2

R4#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
14.0.0.1	1	FULL/BDR	00:00:30	14.0.0.1	Ethernet0/0

R5#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
25.0.0.2	1	FULL/BDR	00:00:31	25.0.0.2	Ethernet0/0

R1#show ip route ospf

```

O   23.0.0.0/8 [110/845] via 13.0.0.3, 00:00:04, Serial0/1
O   25.0.0.0/8 [110/74] via 12.0.0.2, 00:00:40, Serial0/0

```

```
R2#show ip route ospf
O   13.0.0.0/8 [110/845] via 23.0.0.3, 00:00:19, Serial0/1
O   14.0.0.0/8 [110/74] via 12.0.0.1, 00:00:55, Serial0/0

R3#show ip route ospf
O IA 25.0.0.0/8 [110/791] via 23.0.0.2, 00:00:24, Serial1/3
O IA 12.0.0.0/8 [110/845] via 23.0.0.2, 00:00:24, Serial1/3
   [110/845] via 13.0.0.1, 00:00:24, Serial1/2
O IA 14.0.0.0/8 [110/791] via 13.0.0.1, 00:00:24, Serial1/2

R4#show ip route ospf
O IA 23.0.0.0/8 [110/138] via 14.0.0.1, 00:01:04, Ethernet0/0
O   25.0.0.0/8 [110/84] via 14.0.0.1, 00:01:04, Ethernet0/0
O   12.0.0.0/8 [110/74] via 14.0.0.1, 00:01:04, Ethernet0/0
O IA 13.0.0.0/8 [110/74] via 14.0.0.1, 00:01:04, Ethernet0/0

R5#show ip route ospf
O IA 23.0.0.0/8 [110/74] via 25.0.0.2, 00:01:11, Ethernet0/0
O   12.0.0.0/8 [110/74] via 25.0.0.2, 00:01:11, Ethernet0/0
O IA 13.0.0.0/8 [110/138] via 25.0.0.2, 00:01:11, Ethernet0/0
O   14.0.0.0/8 [110/84] via 25.0.0.2, 00:01:11, Ethernet0/0

Before Virtual Link with Frame Relay Down

R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#interface serial0/0
R1(config-if)#shut

%OSPF-5-ADJCHG: Process 1, Nbr 25.0.0.2 on Serial0/0 from FULL to DOWN,
Neighbor Down: Interface down or detached

R1#show ip route ospf
O   23.0.0.0/8 [110/845] via 13.0.0.3, 00:01:51, Serial0/1

R1#ping 25.0.0.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 25.0.0.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#show ip route ospf
O   13.0.0.0/8 [110/845] via 23.0.0.3, 00:01:57, Serial0/1

R2#ping 14.0.0.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 14.0.0.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#show ip route ospf
O IA 25.0.0.0/8 [110/791] via 23.0.0.2, 00:00:37, Serial1/3
O IA 12.0.0.0/8 [110/845] via 23.0.0.2, 00:00:38, Serial1/3
O IA 14.0.0.0/8 [110/791] via 13.0.0.1, 00:00:02, Serial1/2

R4#show ip route ospf
O IA 23.0.0.0/8 [110/855] via 14.0.0.1, 00:00:42, Ethernet0/0
O IA 13.0.0.0/8 [110/74] via 14.0.0.1, 00:00:42, Ethernet0/0
```

```

R5#show ip route ospf
O IA 23.0.0.0/8 [110/74] via 25.0.0.2, 00:00:10, Ethernet0/0
O   12.0.0.0/8 [110/74] via 25.0.0.2, 00:00:10, Ethernet0/0
O IA 13.0.0.0/8 [110/855] via 25.0.0.2, 00:00:10, Ethernet0/0

After Virtual Link with Frame Relay Down

R1#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address      Interface
25.0.0.2         0     FULL/ -         00:00:25   23.0.0.2    OSPF_VL2
14.0.0.4         1     FULL/DR         00:00:36   14.0.0.4    Ethernet0/0
23.0.0.3         0     FULL/ -         00:00:34   13.0.0.3    Serial0/1

R2#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address      Interface
14.0.0.1         0     FULL/ -         00:00:20   13.0.0.1    OSPF_VL2
25.0.0.5         1     FULL/DR         00:00:34   25.0.0.5    Ethernet0/0
23.0.0.3         0     FULL/ -         00:00:37   23.0.0.3    Serial0/1

R1#show ip ospf virtual-links
Virtual Link OSPF_VL2 to router 25.0.0.2 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 1, via interface Serial0/1, Cost of using 845
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
  Adjacency State FULL (Hello suppressed)
  Index 2/3, retransmission queue length 0, number of retransmission 1
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec

R2#show ip ospf virtual-links
Virtual Link OSPF_VL2 to router 14.0.0.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 1, via interface Serial0/1, Cost of using 845
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:00
  Adjacency State FULL (Hello suppressed)
  Index 2/3, retransmission queue length 0, number of retransmission 1
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec

R4#show ip route ospf
O IA 23.0.0.0/8 [110/855] via 14.0.0.1, 00:00:13, Ethernet0/0
O   25.0.0.0/8 [110/865] via 14.0.0.1, 00:00:13, Ethernet0/0
O   12.0.0.0/8 [110/919] via 14.0.0.1, 00:00:13, Ethernet0/0
O IA 13.0.0.0/8 [110/74] via 14.0.0.1, 00:00:13, Ethernet0/0

R5#show ip route ospf
O IA 23.0.0.0/8 [110/74] via 25.0.0.2, 00:00:31, Ethernet0/0
O   12.0.0.0/8 [110/74] via 25.0.0.2, 00:00:31, Ethernet0/0
O IA 13.0.0.0/8 [110/855] via 25.0.0.2, 00:00:31, Ethernet0/0
O   14.0.0.0/8 [110/865] via 25.0.0.2, 00:00:31, Ethernet0/0

```

Breakdown

OSPF area 0 is used as a backbone for the link state domain. In order to properly compute SPF, all routers in OSPF area 0 must have a full copy of the link state database, and must have contiguous adjacencies. This means that a non-backbone area cannot be in the transit path of area 0 adjacency. In the case that area 0 is partitioned and becomes discontinuous, inconsistencies in the OSPF database will appear throughout the network. In this type of circumstance, an OSPF virtual-link can be used to form a logical area 0 adjacency by using a non-backbone area as transit.

An OSPF virtual-link uses a non-backbone transit area to carry LSAs between the Area Border Routers (ABRs) of the discontinuous areas. In order to identify who the intended destination of the area 0 LSAs are, the OSPF router-id is used. Note however that IP reachability need not be established to this address. Instead it is simply a requirement that the routers in the transit path of the non-backbone area know about the router-ids of the ABRs.

In the above case, R1, R2, R4, and R5 make up a contiguous set of area 0 adjacencies. However, when the Frame Relay circuit between R1 and R2 is down, area 0 is partitioned between area 1. The resulting output from `show ip route` on R4 and R5 shows that only partial reachability is maintained when the adjacency between R1 and R2 is down. In addition to this, when the Frame Relay circuit is down, R1 loses connectivity to 25.0.0.2 (R2's router-id), and R2 loses connectivity to 14.0.0.1 (R1's router-id). When the virtual-link is configured between R1 and R2 over area 1, full reachability is established again. Note that R1 and R2 use the configured router-ids of each other to form the virtual-link, and not necessarily the IP address of the interface that is facing the transit area..

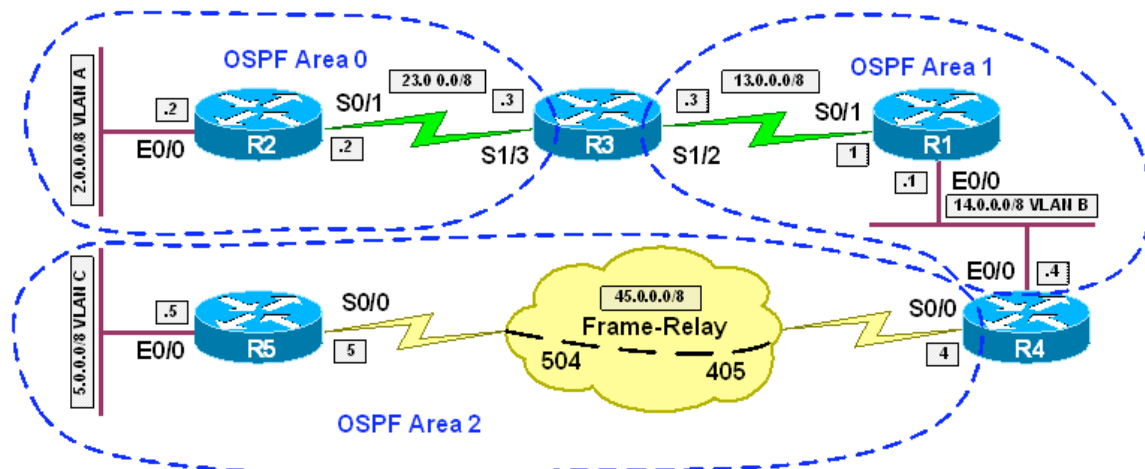


Recommended Reading

[Cisco IOS Command Reference: area virtual-link](#)

OSPF Virtual Links - Repairing Discontiguous Areas

Objective: Configure the OSPF domain per the diagram ensuring complete connectivity



Directions

- Configure the IP addresses of R1, R2, R3, R4, and R5 per the above diagram
- Configure the Frame Relay circuit between R4 and R5 using static layer 3 to layer 2 resolution
- Configure OSPF area 0 on R2's interface Ethernet0/0
- Configure OSPF area 0 on the Serial link between R2 and R3
- Configure OSPF area 1 on the Serial link between R1 and R3
- Configure OSPF area 1 on the Ethernet link between R1 and R4
- Configure OSPF area 2 on the Frame Relay link between R4 and R5 using network type point-to-point
- Configure OSPF area 2 on R5's interface Ethernet0/0
- Configure an OSPF virtual link between R3 and R4 using OSPF area 1 as transit

Ask Yourself

- Why must all areas in OSPF be connected to area 0?
- What happens when an area is not connected to area 0?
- Why must area 0 be contiguous throughout the OSPF domain?
- What happens when area 0 becomes partitioned?
- What are the methods of repairing a partitioned area 0?
- When configuring an OSPF virtual-link, which routers should it terminate on?
- What does the IP address used in the virtual-link statement refer to?
- Must this address be reachable before the virtual-link can be established?

Final Configuration**R1:**

```
interface Ethernet0/0
 ip address 14.0.0.1 255.0.0.0
!
interface Serial0/1
 ip address 13.0.0.1 255.0.0.0
!
router ospf 1
 network 13.0.0.1 0.0.0.0 area 1
 network 14.0.0.1 0.0.0.0 area 1
```

R2:

```
interface Ethernet0/0
 ip address 2.0.0.2 255.0.0.0
!
interface Serial0/1
 ip address 23.0.0.2 255.0.0.0
!
router ospf 1
 network 2.0.0.2 0.0.0.0 area 0
 network 23.0.0.2 0.0.0.0 area 0
```

R3:

```
interface Serial1/2
 ip address 13.0.0.3 255.0.0.0
 clock rate 64000
!
interface Serial1/3
 ip address 23.0.0.3 255.0.0.0
 clock rate 64000
!
router ospf 1
 area 1 virtual-link 45.0.0.4
 network 13.0.0.3 0.0.0.0 area 1
 network 23.0.0.3 0.0.0.0 area 0
```

R4:

```
interface Ethernet0/0
 ip address 14.0.0.4 255.0.0.0
!
interface Serial0/0
 ip address 45.0.0.4 255.0.0.0
 encapsulation frame-relay
 ip ospf network point-to-point
 frame-relay map ip 45.0.0.5 405 broadcast
!
router ospf 1
 area 1 virtual-link 23.0.0.3
 network 14.0.0.4 0.0.0.0 area 1
 network 45.0.0.4 0.0.0.0 area 2
```

R5:

```
interface Ethernet0/0
 ip address 5.0.0.5 255.0.0.0
!
interface Serial0/0
 ip address 45.0.0.5 255.0.0.0
 encapsulation frame-relay
 ip ospf network point-to-point
```

```

frame-relay map ip 45.0.0.4 504 broadcast
!
router ospf 1
network 5.0.0.5 0.0.0.0 area 2
network 45.0.0.5 0.0.0.0 area 2

```

Verification

Before Virtual-Link

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
45.0.0.4	1	FULL/DR	00:00:32	14.0.0.4	Ethernet0/0
23.0.0.3	0	FULL/ -	00:00:38	13.0.0.3	Serial0/1

R1#show ip route ospf

```

O IA 2.0.0.0/8 [110/855] via 13.0.0.3, 00:00:29, Serial0/1
O IA 23.0.0.0/8 [110/845] via 13.0.0.3, 00:00:29, Serial0/1

```

R2#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
23.0.0.3	0	FULL/ -	00:00:36	23.0.0.3	Serial0/1

R2#show ip route ospf

```

O IA 13.0.0.0/8 [110/845] via 23.0.0.3, 00:00:59, Serial0/1
O IA 14.0.0.0/8 [110/855] via 23.0.0.3, 00:00:59, Serial0/1

```

R3#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
23.0.0.2	0	FULL/ -	00:00:36	23.0.0.2	Serial1/3
14.0.0.1	0	FULL/ -	00:00:35	13.0.0.1	Serial1/2

R3#show ip route ospf

```

O 2.0.0.0/8 [110/791] via 23.0.0.2, 00:01:05, Serial1/3
O 14.0.0.0/8 [110/791] via 13.0.0.1, 00:00:55, Serial1/2

```

R4#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
14.0.0.1	1	FULL/BDR	00:00:33	14.0.0.1	Ethernet0/0
45.0.0.5	0	FULL/ -	00:00:34	45.0.0.5	Serial0/0

R4#show ip route ospf

```

O IA 2.0.0.0/8 [110/865] via 14.0.0.1, 00:00:56, Ethernet0/0
O 5.0.0.0/8 [110/74] via 45.0.0.5, 00:00:56, Serial0/0
O IA 23.0.0.0/8 [110/855] via 14.0.0.1, 00:00:56, Ethernet0/0
O 13.0.0.0/8 [110/74] via 14.0.0.1, 00:00:56, Ethernet0/0

```

R5#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
45.0.0.4	0	FULL/ -	00:00:32	45.0.0.4	Serial0/0

R5#show ip route ospf

After Virtual-Link

R1#show ip route ospf

```
O IA 2.0.0.0/8 [110/855] via 13.0.0.3, 00:00:06, Serial0/1
O IA 5.0.0.0/8 [110/84] via 14.0.0.4, 00:00:15, Ethernet0/0
O IA 23.0.0.0/8 [110/845] via 13.0.0.3, 00:00:06, Serial0/1
O IA 45.0.0.0/8 [110/74] via 14.0.0.4, 00:00:15, Ethernet0/0
```

R2#show ip route ospf

```
O IA 5.0.0.0/8 [110/929] via 23.0.0.3, 00:00:18, Serial0/1
O IA 13.0.0.0/8 [110/845] via 23.0.0.3, 00:00:18, Serial0/1
O IA 14.0.0.0/8 [110/855] via 23.0.0.3, 00:00:18, Serial0/1
O IA 45.0.0.0/8 [110/919] via 23.0.0.3, 00:00:18, Serial0/1
```

R3#show ip route ospf

```
O 2.0.0.0/8 [110/791] via 23.0.0.2, 00:00:23, Serial1/3
O IA 5.0.0.0/8 [110/865] via 13.0.0.1, 00:00:23, Serial1/2
O 14.0.0.0/8 [110/791] via 13.0.0.1, 00:00:33, Serial1/2
O IA 45.0.0.0/8 [110/855] via 13.0.0.1, 00:00:23, Serial1/2
```

R3#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
45.0.0.4	0	FULL/ -	-	14.0.0.4	OSPF_VL1
23.0.0.2	0	FULL/ -	00:00:38	23.0.0.2	Serial1/3
14.0.0.1	0	FULL/ -	00:00:37	13.0.0.1	Serial1/2

R3#show ip ospf virtual-links

Virtual Link OSPF_VL1 to router 45.0.0.4 is up

```
Run as demand circuit
DoNotAge LSA allowed.
Transit area 1, via interface Serial1/2, Cost of using 791
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:00
Adjacency State FULL (Hello suppressed)
Index 2/3, retransmission queue length 0, number of retransmission 1
First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
Last retransmission scan length is 1, maximum is 1
Last retransmission scan time is 0 msec, maximum is 0 msec
```

R4#show ip route ospf

```
O 2.0.0.0/8 [110/865] via 14.0.0.1, 00:00:16, Ethernet0/0
O 5.0.0.0/8 [110/74] via 45.0.0.5, 00:00:36, Serial0/0
O 23.0.0.0/8 [110/855] via 14.0.0.1, 00:00:16, Ethernet0/0
O 13.0.0.0/8 [110/74] via 14.0.0.1, 00:00:36, Ethernet0/0
```

R4#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
23.0.0.3	0	FULL/ -	-	13.0.0.3	OSPF_VL1
14.0.0.1	1	FULL/BDR	00:00:32	14.0.0.1	Ethernet0/0
45.0.0.5	0	FULL/ -	00:00:33	45.0.0.5	Serial0/0

R4#show ip ospf virtual-links

Virtual Link OSPF_VL1 to router 23.0.0.3 is up

```
Run as demand circuit
DoNotAge LSA allowed.
Transit area 1, via interface Ethernet0/0, Cost of using 74
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:09
Adjacency State FULL (Hello suppressed)
Index 1/2, retransmission queue length 0, number of retransmission 0
```

```
First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
Last retransmission scan length is 0, maximum is 0
Last retransmission scan time is 0 msec, maximum is 0 msec
```

```
R5#show ip route ospf
```

```
O IA 2.0.0.0/8 [110/929] via 45.0.0.4, 00:00:21, Serial0/0
O IA 23.0.0.0/8 [110/919] via 45.0.0.4, 00:00:21, Serial0/0
O IA 13.0.0.0/8 [110/138] via 45.0.0.4, 00:00:40, Serial0/0
O IA 14.0.0.0/8 [110/74] via 45.0.0.4, 00:00:40, Serial0/0
```

Breakdown

In order to ensure that SPF is computed properly, area 0 is used as a reference point for all other areas in the network. In order for this process to work properly, it is implied that all areas must have a direct adjacency to area 0. In the previous scenario we saw that an OSPF virtual-link can be used to repair a break in area 0. In this scenario a virtual link is used to connect a discontinuous area (area 2) back to area 0.

In the above scenario the virtual link runs between the ABR of the discontinuous area and the area 0 ABR (R4 and R3 respectively) and uses area 1 for transit. Note that the virtual link is configured based on the router-ids of the above-mentioned ABRs involved, not the interfaces that face the transit area in question.

Prior to the virtual-link being configured it is seen through the `show ip route` output that the OSPF databases throughout the network are not synchronized. Once the virtual-link is established, full connectivity is obtained.

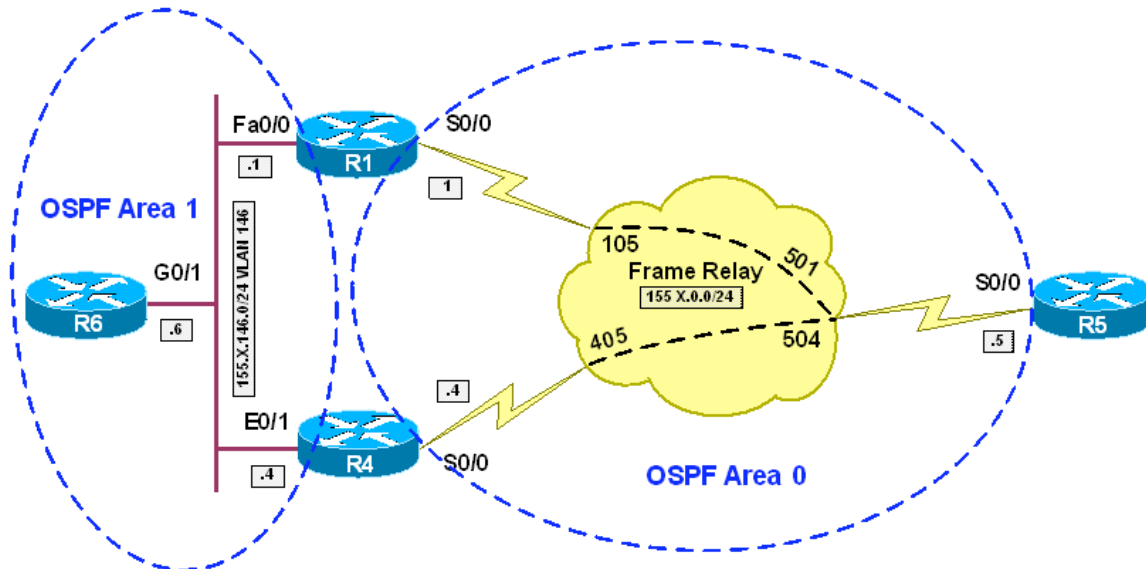


Recommended Reading

[Cisco IOS Command Reference: area virtual-link](#)

Common Multi-Area Configuration with 2 ABRs

Objective: Configure OSPF topology with two Areas and two ABRs



Directions

- Create VLAN 146 on all switches involved in the scenario
- Configure the respective switch ports for R1, R4 and R6 in VLAN 146
- Configure IP addressing on VLAN 146 interfaces as per diagram
- Configure IP addressing on Frame-Relay interfaces. Use only physical FR interfaces, and use static mappings. Use DLCI numbers specified on diagram
- Configure R1 and R4 FR interfaces so that they can reach each other via R5.
- Configure VLAN 146 in OSPF Area 1
- Configure FR interfaces in OSPF Area 0. Use OSPF network type non-broadcast for this task.
- To make sure that R5 is always elected as DR, configure “ip ospf priority 0” on FR interfaces of R4 and R1.
- Configure static neighbors “155.X.0.1” and “155.X.0.4” under OSPF routing process on R5
- Create Loopback0 interfaces on R1, R4, R5, R6 with IP addresses 150.X.Y.Y/24 where Y is router number
- Advertise Loopback 0 interfaces on R1, R4, R6 into OSPF Area 1
- Advertise Loopback 0 interface on R5 into OSPF Area 0

Final Configuration

```
SW1-SW4:
vlan 146

SW1:
interface Fa 0/1
```

```
switchport host
switchport access vlan 146

SW2:
interface Fa 0/6
switchport host
switchport access vlan 146

SW4:
interface Fa 0/4
switchport host
switchport access vlan 146

R1:
interface Fa 0/0
ip address 155.1.146.1 255.255.255.0
no shut
!
interface Serial 0/0
encapsulation frame-relay
no frame-relay inverse-arp
ip address 155.1.0.1 255.255.255.0
frame-relay map ip 155.1.0.5 105
frame-relay map ip 155.1.0.4 105
ip ospf priority 0
no shutdown
!
interface Loopback0
ip address 150.1.1.1 255.255.255.0
!
router ospf 1
router-id 150.1.1.1
network 155.1.146.1 0.0.0.0 area 1
network 150.1.1.1 0.0.0.0 area 1
network 155.1.0.1 0.0.0.0 area 0

R4:
interface Eth 0/1
ip address 155.1.146.4 255.255.255.0
!
interface Serial 0/0
encapsulation frame-relay
no frame-relay inverse-arp
ip address 155.1.0.4 255.255.255.0
frame-relay map ip 155.1.0.5 405
frame-relay map ip 155.1.0.1 405
ip ospf priority 0
!
interface Loopback0
ip address 150.1.4.4 255.255.255.0
!
router ospf 1
router-id 150.1.4.4
network 155.1.146.4 0.0.0.0 area 1
network 150.1.4.4 0.0.0.0 area 1
network 155.1.0.4 0.0.0.0 area 0

R5:
interface Serial 0/0
encapsulation frame-relay
no frame-relay inverse-arp
ip address 155.1.0.5 255.255.255.0
```

```

frame-relay map ip 155.1.0.4 504
frame-relay map ip 155.1.0.1 501
!
interface Loopback0
 ip address 150.1.5.5 255.255.255.0
!
router ospf 1
 router-id 150.1.5.5
 network 150.1.5.5 0.0.0.0 area 0
 network 155.1.0.5 0.0.0.0 area 0
 neighbor 155.1.0.4
 neighbor 155.1.0.1

```

R6:

```

interface Gig 0/1
 ip address 155.1.146.6 255.255.255.0
!
interface Loopback0
 ip address 150.1.6.6 255.255.255.0
!
router ospf 1
 router-id 150.1.6.6
 network 155.1.146.6 0.0.0.0 area 1
 network 150.1.6.6 0.0.0.0 area 1

```

Verification

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:52	155.1.0.5	Serial0/0
150.1.4.4	1	FULL/DROTHER	00:00:32	155.1.146.4	FastEthernet0/0
150.1.6.6	1	FULL/BDR	00:00:33	155.1.146.6	FastEthernet0/0

R4#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:43	155.1.0.5	Serial0/0
150.1.1.1	1	FULL/DR	00:00:37	155.1.146.1	Ethernet0/1
150.1.6.6	1	FULL/BDR	00:00:35	155.1.146.6	Ethernet0/1

R5#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.4.4	0	FULL/DROTHER	00:01:35	155.1.0.4	Serial0/0
150.1.1.1	0	FULL/DROTHER	00:01:57	155.1.0.1	Serial0/0

R6#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.1.1	1	FULL/DR	00:00:30	155.1.146.1	GigabitEthernet0/1
150.1.4.4	1	FULL/DROTHER	00:00:37	155.1.146.4	GigabitEthernet0/1

R5#show ip route ospf

```


155.1.0.0/24 is subnetted, 3 subnets
O IA 155.1.146.0 [110/391] via 155.1.0.1, 00:00:00, Serial0/0
150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA 150.1.6.6/32 [110/392] via 155.1.0.1, 00:00:00, Serial0/0
O IA 150.1.4.4/32 [110/391] via 155.1.0.4, 00:00:00, Serial0/0

```

```
O IA 150.1.1.1/32 [110/391] via 155.1.0.1, 00:00:00, Serial0/0
```

Using Type-3 LSA Inter-Area Filtering

Objective: Filter OSPF summary-prefixes on ABRs



Directions

- Configure routers as per the OSPF scenario “Common Multi-Area Configuration with 2 ABRs”
- Create prefix-list AREA1_FILTER_OUT on R1 and R4. Deny the 150.X.6.6/32 prefix and permit all other
- Apply this filter to prefixes coming out of Area 1 on R1 and R4
- Create prefix-list AREA1_FILTER_IN on R1 and R4. Deny the 150.X.5.5/32 prefix and permit all other
- Apply this filter to prefixes coming into Area 1 on R1 and R4
- Note that keyword “in” means coming into area, and “out” - coming out of the area

Final Configuration**R1 & R4:**

```
no ip prefix-list AREA1_FILTER_OUT
ip prefix-list AREA1_FILTER_OUT deny 150.1.6.6/32
ip prefix-list AREA1_FILTER_OUT permit 0.0.0.0/0 le 32

no ip prefix-list AREA1_FILTER_IN
ip prefix-list AREA1_FILTER_IN deny 150.1.5.5/32
ip prefix-list AREA1_FILTER_IN permit 0.0.0.0/0 le 32

router ospf 1
 area 1 filter-list prefix AREA1_FILTER_OUT out
 area 1 filter-list prefix AREA1_FILTER_IN in
```

Verification

Before filtering has been applied:

R5#sho ip route ospf

```
155.1.0.0/24 is subnetted, 3 subnets
O IA 155.1.146.0 [110/391] via 155.1.0.1, 00:21:06, Serial0/0
150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA 150.1.6.6/32 [110/392] via 155.1.0.1, 00:21:06, Serial0/0
O IA 150.1.4.4/32 [110/391] via 155.1.0.4, 00:21:06, Serial0/0
O IA 150.1.1.1/32 [110/391] via 155.1.0.1, 00:21:06, Serial0/0
```

R6#show ip route ospf

```
155.1.0.0/24 is subnetted, 2 subnets
O IA 155.1.0.0 [110/65] via 155.1.146.1, 00:21:39, GigabitEthernet0/1
150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA 150.1.5.5/32 [110/66] via 155.1.146.1, 00:21:22, GigabitEthernet0/1
O 150.1.4.4/32 [110/2] via 155.1.146.4, 00:21:42, GigabitEthernet0/1
O 150.1.1.1/32 [110/2] via 155.1.146.1, 00:21:42, GigabitEthernet0/1
```

After that:

R6#show ip route ospf


```
155.1.0.0/24 is subnetted, 2 subnets
O IA 155.1.0.0 [110/65] via 155.1.146.1, 00:26:04, GigabitEthernet0/1
150.1.0.0/16 is variably subnetted, 3 subnets, 2 masks
O 150.1.4.4/32 [110/2] via 155.1.146.4, 00:26:06, GigabitEthernet0/1
O 150.1.1.1/32 [110/2] via 155.1.146.1, 00:26:06, GigabitEthernet0/1
```

R5#sho ip route ospf

```
155.1.0.0/24 is subnetted, 3 subnets
O IA 155.1.146.0 [110/391] via 155.1.0.1, 00:25:41, Serial0/0
150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA 150.1.4.4/32 [110/391] via 155.1.0.4, 00:25:41, Serial0/0
O IA 150.1.1.1/32 [110/391] via 155.1.0.1, 00:25:41, Serial0/0
```


Type-3 LSA Filtering with Network Ranges

Objective: Filter OSPF summary-prefixes using network ranges



Directions

- Configure routers as per the OSPF scenario “Common Multi-Area Configuration with 2 ABRs”
- Create network-range 150.1.6.0/24 for Area 1, add keyword “**not-advertise**” to block propagation of the summary
- Create network-range 150.1.5.0/24 for Area 0, add keyword “**not-advertise**” to block propagation of the summary

Final Configuration

```
R1 & R4:
router ospf 1
 area 1 range 150.1.6.0 255.255.255.0 not-adv
 area 0 range 150.1.5.0 255.255.255.0 not-adv
```

Verification

Before summarization has been applied:

```
R5#sho ip route ospf
 155.1.0.0/24 is subnetted, 3 subnets
 O IA   155.1.146.0 [110/391] via 155.1.0.1, 00:21:06, Serial0/0
 150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
 O IA   150.1.6.6/32 [110/392] via 155.1.0.1, 00:21:06, Serial0/0
 O IA   150.1.4.4/32 [110/391] via 155.1.0.4, 00:21:06, Serial0/0
 O IA   150.1.1.1/32 [110/391] via 155.1.0.1, 00:21:06, Serial0/0
```

```
R6#show ip route ospf
 155.1.0.0/24 is subnetted, 2 subnets
```

```
O IA 155.1.0.0 [110/65] via 155.1.146.1, 00:21:39, GigabitEthernet0/1
    150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA 150.1.5.5/32 [110/66] via 155.1.146.1, 00:21:22, GigabitEthernet0/1
O    150.1.4.4/32 [110/2] via 155.1.146.4, 00:21:42, GigabitEthernet0/1
O    150.1.1.1/32 [110/2] via 155.1.146.1, 00:21:42, GigabitEthernet0/1
```

After that:

```
R6#show ip route ospf
```


```
155.1.0.0/24 is subnetted, 2 subnets
O IA 155.1.0.0 [110/65] via 155.1.146.1, 00:26:04, GigabitEthernet0/1
    150.1.0.0/16 is variably subnetted, 3 subnets, 2 masks
O    150.1.4.4/32 [110/2] via 155.1.146.4, 00:26:06, GigabitEthernet0/1
O    150.1.1.1/32 [110/2] via 155.1.146.1, 00:26:06, GigabitEthernet0/1
```

```
R5#sho ip route ospf
```

```
155.1.0.0/24 is subnetted, 3 subnets
O IA 155.1.146.0 [110/391] via 155.1.0.1, 00:25:41, Serial0/0
    150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA 150.1.4.4/32 [110/391] via 155.1.0.4, 00:25:41, Serial0/0
O IA 150.1.1.1/32 [110/391] via 155.1.0.1, 00:25:41, Serial0/0
```

Ingress Filtering with Distribute-List

Objective: Filter OSPF routes entering local IP routing table



Directions

- Configure routers as per the OSPF scenario “Common Multi-Area Configuration with 2 ABRs”
- Create access-list 99 on R1 and R4. Deny the 150.1.6.6/32 prefix and permit everything else with this access-list
- Apply access-list as distribute-list for OSPF processes on R1 and R4
- Note that distribute-list only block route from entering local RIB, and does not stop LSAs propagation

Final Configuration

```
R1 & R4
no access-list 99
access-list 99 deny 150.1.6.6
access-list 99 permit any
!
router ospf 1
  distribute-list 99 in
```

Verification


```
R1#sh ip route ospf
    150.1.0.0/16 is variably subnetted, 3 subnets, 2 masks
O       150.1.5.5/32 [110/65] via 155.1.0.5, 00:01:49, Serial0/0
O       150.1.4.4/32 [110/2] via 155.1.146.4, 00:01:49, FastEthernet0/0

R4#show ip route ospf
    150.1.0.0/16 is variably subnetted, 3 subnets, 2 masks
O       150.1.5.5/32 [110/391] via 155.1.0.5, 00:00:53, Serial0/0
O       150.1.1.1/32 [110/11] via 155.1.146.1, 00:00:53, Ethernet0/1

R5#show ip route ospf
    155.1.0.0/24 is subnetted, 3 subnets
O IA    155.1.146.0 [110/391] via 155.1.0.1, 00:03:46, Serial0/0
    150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA    150.1.6.6/32 [110/392] via 155.1.0.1, 00:01:26, Serial0/0
O IA    150.1.4.4/32 [110/391] via 155.1.0.4, 00:03:59, Serial0/0
O IA    150.1.1.1/32 [110/391] via 155.1.0.1, 00:03:51, Serial0/0
```

Ingress Filtering with Distribute-List and Route-Map

Objective: Filter only external OSPF routes originating at specific router from entering local IP routing table



Directions

- Configure routers as per the OSPF scenario “Common Multi-Area Configuration with 2 ABRs”
- Create static route 160.X.60.0/24 to Null0 on R6. Redistribute this route into OSPF on R6
- Create static route 160.X.50.0/24 to Null0 on R5. Redistribute this route into OSPF on R5
- Create access-list 1 on R1 and R4 and match host 150.1.6.6 (R6’s router-id) with it
- Create route-map FILTER_OSPF on R1 and R4. Deny external routes with this route-map, and permit all other types of routes. Configure route-map to match routes from R6 only, by matching route-source based on access-list 1
- Apply route-map FILTER_OSPF as ingress filter for OSPF processes on R1 and R4.

Final Configuration

```
R1 & R4
access-list 1 permit 150.1.6.6
!
route-map FILTER_OSPF deny 10
 match route-type external type-2
 match ip route-source 1
route-map FILTER_OSPF permit 1000
!
router ospf 1
```

```
    distribute-list route-map FILTER_OSPF in  
  
R5:  
ip route 160.1.50.0 255.255.255.0 null 0  
router ospf 1  
    redistribute static subnets  
  
R6:  
ip route 160.1.60.0 255.255.255.0 null 0  
router ospf 1  
    redistribute static subnets
```

Verification

```
R1#sh ip route ospf
 160.1.0.0/24 is subnetted, 1 subnets
O E2 160.1.50.0 [110/20] via 155.1.0.5, 00:00:31, Serial0/0
 150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O    150.1.6.6/32 [110/2] via 155.1.146.6, 00:00:51, FastEthernet0/0
O    150.1.5.5/32 [110/65] via 155.1.0.5, 00:00:41, Serial0/0
O    150.1.4.4/32 [110/2] via 155.1.146.4, 00:00:51, FastEthernet0/0

R1#sh ip access-lists 1
Standard IP access list 1
 10 permit 150.1.6.6 (5 matches)

R4#sh ip route ospf
 160.1.0.0/24 is subnetted, 1 subnets
O E2 160.1.50.0 [110/20] via 155.1.0.5, 00:00:43, Serial0/0
 150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O    150.1.6.6/32 [110/11] via 155.1.146.6, 00:01:03, Ethernet0/1
O    150.1.5.5/32 [110/391] via 155.1.0.5, 00:00:53, Serial0/0
O    150.1.1.1/32 [110/11] via 155.1.146.1, 00:01:03, Ethernet0/1

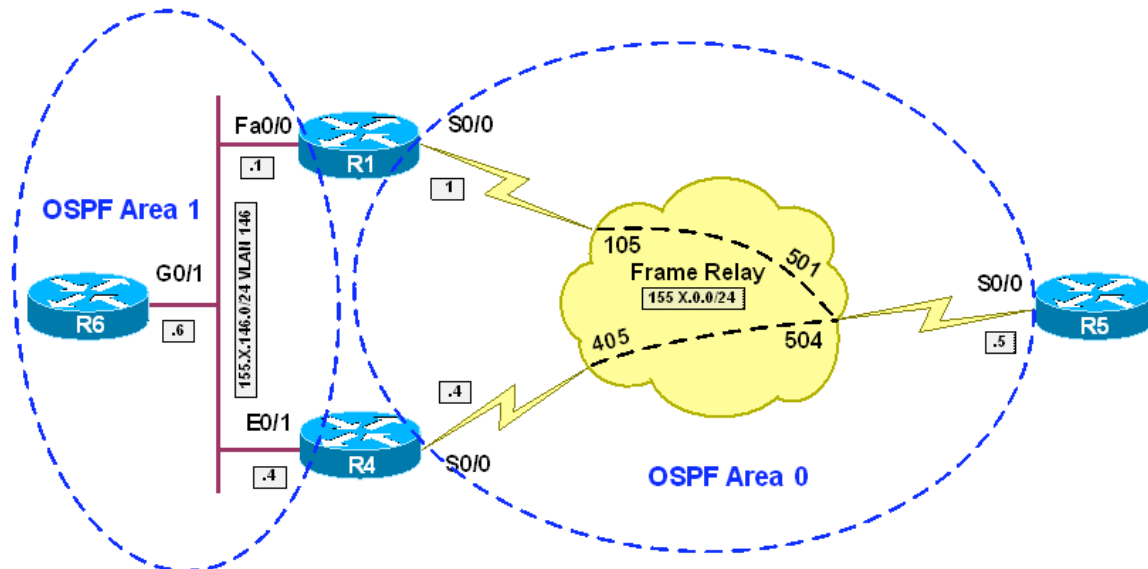
R4#sh ip access-lists 1
Standard IP access list 1
 10 permit 150.1.6.6 (5 matches)

R6#sh ip route ospf
 155.1.0.0/24 is subnetted, 2 subnets
O IA 155.1.0.0 [110/65] via 155.1.146.1, 00:01:11, GigabitEthernet0/1
 160.1.0.0/24 is subnetted, 2 subnets
O E2 160.1.50.0 [110/20] via 155.1.146.1, 00:00:56, GigabitEthernet0/1
 150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA 150.1.5.5/32 [110/66] via 155.1.146.1, 00:01:11, GigabitEthernet0/1
O    150.1.4.4/32 [110/2] via 155.1.146.4, 00:01:11, GigabitEthernet0/1
O    150.1.1.1/32 [110/2] via 155.1.146.1, 00:01:11, GigabitEthernet0/1

R5#sh ip route ospf
 155.1.0.0/24 is subnetted, 3 subnets
O IA 155.1.146.0 [110/391] via 155.1.0.1, 00:01:18, Serial0/0
 160.1.0.0/24 is subnetted, 2 subnets
O E2 160.1.60.0 [110/20] via 155.1.0.1, 00:01:18, Serial0/0
 150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA 150.1.6.6/32 [110/392] via 155.1.0.1, 00:01:18, Serial0/0
O IA 150.1.4.4/32 [110/391] via 155.1.0.4, 00:01:18, Serial0/0
O IA 150.1.1.1/32 [110/391] via 155.1.0.1, 00:01:18, Serial0/0
```

Ingress Filtering with Administrative Distance

Objective: Filter out summary and external routes by changing the administrative distance



Directions

- Configure routers as per the OSPF scenario “Common Multi-Area Configuration with 2 ABRs”
- Create static route 160.X.60.0/24 to Null0 on R6. Redistribute this route into OSPF on R6
- The objective is to filter out prefixes 160.X.60.0/24 and 150.X.6.6/32 on R5
- Note that prefix 160.X.60.0/24 is external prefix originated at R6, and 150.X.6.6/32 is summary prefix originated at R1 and R4
- Create access-list 99 on R5 and match prefix 160.X.60.0 with it
- Create access-list 98 on R5 and match prefix 150.X.6.6 with it
- Under OSPF routing process at R5 change administrative distance for prefixes in access-list 99 originated at R6
- Under OSPF routing process at R5 change administrative distance for prefixes in access-list 98 originated at R1 and R4

Final Configuration

```
R5:
access-list 99 permit 160.1.60.0
access-list 98 permit 150.1.6.6
!
router ospf 1
 distance 255 150.1.6.6 0.0.0.0 99
 distance 255 150.1.1.1 0.0.0.0 98
 distance 255 150.1.4.4 0.0.0.0 98
```

Verification

```
R5#show ip route ospf
      155.1.0.0/24 is subnetted, 3 subnets
O IA   155.1.146.0 [110/391] via 155.1.0.1, 00:03:30, Serial0/0
      150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA   150.1.4.4/32 [110/391] via 155.1.0.4, 00:03:31, Serial0/0
O IA   150.1.1.1/32 [110/391] via 155.1.0.1, 00:03:31, Serial0/0
```

```
R5#show ip ospf database external
```

```
          OSPF Router with ID (150.1.5.5) (Process ID 1)
```

```
          Type-5 AS External Link States
```

```
Routing Bit Set on this LSA
LS age: 984
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 160.1.60.0 (External Network Number )
Advertising Router: 150.1.6.6
LS Seq Number: 80000006
Checksum: 0x71A8
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

```
R5#show ip ospf database summary
```

```
          OSPF Router with ID (150.1.5.5) (Process ID 1)
```

```
          Summary Net Link States (Area 0)
```

```
Routing Bit Set on this LSA
LS age: 1091
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.1.1 (summary Network Number)
Advertising Router: 150.1.1.1
LS Seq Number: 80000001
Checksum: 0x6C9C
Length: 28
```

```
Network Mask: /32
  TOS: 0 Metric: 1
```

```
LS age: 1050
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.1.1 (summary Network Number)
Advertising Router: 150.1.4.4
LS Seq Number: 80000001
Checksum: 0xA94F
Length: 28
Network Mask: /32
  TOS: 0 Metric: 11
```

```
Routing Bit Set on this LSA
LS age: 1053
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.4.4 (summary Network Number)
Advertising Router: 150.1.1.1
LS Seq Number: 80000001
Checksum: 0x37CA
Length: 28
Network Mask: /32
  TOS: 0 Metric: 2
```

```
Routing Bit Set on this LSA
LS age: 690
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.4.4 (summary Network Number)
Advertising Router: 150.1.4.4
LS Seq Number: 8000002D
Checksum: 0xAD23
Length: 28
Network Mask: /32
  TOS: 0 Metric: 1
```

```
Routing Bit Set on this LSA
LS age: 1056
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.6.6 (summary Network Number)
Advertising Router: 150.1.1.1
LS Seq Number: 80000001
Checksum: 0xDF0
Length: 28
Network Mask: /32
  TOS: 0 Metric: 2
```

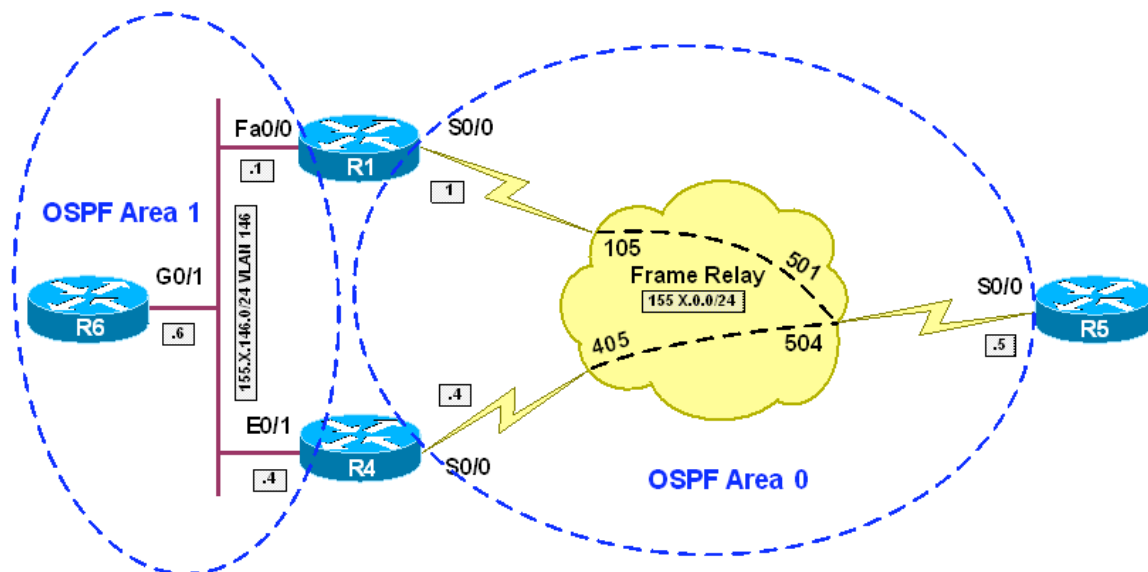
```
Routing Bit Set on this LSA
LS age: 1055
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.6.6 (summary Network Number)
Advertising Router: 150.1.4.4
LS Seq Number: 80000001
Checksum: 0x40AE
Length: 28
Network Mask: /32
  TOS: 0 Metric: 11
```

```
Routing Bit Set on this LSA
LS age: 1058
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 155.1.146.0 (summary Network Number)
Advertising Router: 150.1.1.1
LS Seq Number: 80000001
Checksum: 0xF37F
Length: 28
Network Mask: /24
    TOS: 0  Metric: 1
```

```
LS age: 1058
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 155.1.146.0 (summary Network Number)
Advertising Router: 150.1.4.4
LS Seq Number: 80000001
Checksum: 0x273D
Length: 28
Network Mask: /24
    TOS: 0  Metric: 10
```

NSSA Area Type 7to5 LSA Translator Election

Objective: Influence OSPF NSSA Type 7 to Type 5 translating router election



Directions

- Configure routers as per the OSPF scenario “Common Multi-Area Configuration with 2 ABRs”
- Configure Area 1 as NSSA on R1, R4 and R6
- Add a static route 160.X.60.0/24 to Null0 on R6 and redistribute it into OSPF
- Configure R1 to have higher router-id than R4, so that it may become designated NSSA Area 7->5 translator

Final Configuration

R1,R4,R6:

```
router ospf 1
 area 1 nssa
```

R6:

```
ip route 160.1.60.0 255.255.255.0 Null0
```

```
router ospf 1
 redistribute static subnets
```

R1:

```
router ospf 1
 router-id 150.1.100.100
```

Verification

Before:

R1#show ip route ospf

```
160.1.0.0/24 is subnetted, 1 subnets
O E2 160.1.60.0 [110/20] via 155.1.146.6, 00:00:41, FastEthernet0/0
150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O 150.1.6.6/32 [110/2] via 155.1.146.6, 00:00:41, FastEthernet0/0
O 150.1.5.5/32 [110/65] via 155.1.0.5, 00:01:24, Serial0/0
O 150.1.4.4/32 [110/2] via 155.1.146.4, 00:00:41, FastEthernet0/0
```

R4#show ip route ospf

```
160.1.0.0/24 is subnetted, 1 subnets
O N2 160.1.60.0 [110/20] via 155.1.146.6, 00:01:10, Ethernet0/1
150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O 150.1.6.6/32 [110/11] via 155.1.146.6, 00:01:10, Ethernet0/1
O 150.1.5.5/32 [110/391] via 155.1.0.5, 00:01:25, Serial0/0
O 150.1.1.1/32 [110/11] via 155.1.146.1, 00:01:10, Ethernet0/1
```

R5#show ip route ospf

```
155.1.0.0/24 is subnetted, 3 subnets
O IA 155.1.146.0 [110/391] via 155.1.0.1, 00:01:28, Serial0/0
160.1.0.0/24 is subnetted, 1 subnets
O E2 160.1.60.0 [110/20] via 155.1.0.1, 00:02:06, Serial0/0
150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA 150.1.6.6/32 [110/392] via 155.1.0.1, 00:01:42, Serial0/0
O IA 150.1.4.4/32 [110/391] via 155.1.0.4, 00:04:42, Serial0/0
O IA 150.1.1.1/32 [110/391] via 155.1.0.1, 00:04:36, Serial0/0
```

R5#show ip ospf database external 160.1.60.0

OSPF Router with ID (150.1.5.5) (Process ID 1)

Type-5 AS External Link States

Routing Bit Set on this LSA

LS age: 226

Options: (No TOS-capability, DC)

LS Type: AS External Link

Link State ID: 160.1.60.0 (External Network Number)

Advertising Router: 150.1.4.4

LS Seq Number: 80000001

Checksum: 0xBCC2

Length: 36

Network Mask: /24

Metric Type: 2 (Larger than any link state path)

TOS: 0

Metric: 20

Forward Address: 150.1.6.6

External Route Tag: 0

R1#show ip ospf database external 160.1.60.0

OSPF Router with ID (150.1.1.1) (Process ID 1)

Type-5 AS External Link States

Routing Bit Set on this LSA

LS age: 321

Options: (No TOS-capability, DC)

```

LS Type: AS External Link
Link State ID: 160.1.60.0 (External Network Number )
Advertising Router: 150.1.4.4
LS Seq Number: 80000001
Checksum: 0xBCC2
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 150.1.6.6
    External Route Tag: 0

```

R4#show ip ospf database nssa-external 160.1.60.0

```

                OSPF Router with ID (150.1.4.4) (Process ID 1)

                Type-7 AS External Link States (Area 1)

Routing Bit Set on this LSA
LS age: 376
Options: (No TOS-capability, Type 7/5 translation, DC)
LS Type: AS External Link
Link State ID: 160.1.60.0 (External Network Number )
Advertising Router: 150.1.6.6
LS Seq Number: 80000001
Checksum: 0xE63
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 150.1.6.6
    External Route Tag: 0

```

After:

```

R1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#router ospf 1
R1(config-router)#router-id 150.1.100.100
Reload or use "clear ip ospf process" command, for this to take effect
R1(config-router)#^Z
R1#clear ip ospf process
Reset ALL OSPF processes? [no]: yes

R1#show ip route ospf
    160.1.0.0/24 is subnetted, 1 subnets
O N2  160.1.60.0 [110/20] via 155.1.146.6, 00:00:01, FastEthernet0/0
    150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O     150.1.6.6/32 [110/2] via 155.1.146.6, 00:00:30, FastEthernet0/0
O     150.1.5.5/32 [110/65] via 155.1.0.5, 00:00:01, Serial0/0
O     150.1.4.4/32 [110/2] via 155.1.146.4, 00:00:30, FastEthernet0/0

R4#show ip route ospf
    160.1.0.0/24 is subnetted, 1 subnets
O E2  160.1.60.0 [110/20] via 155.1.146.6, 00:00:08, Ethernet0/1
    150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O     150.1.6.6/32 [110/11] via 155.1.146.6, 00:00:27, Ethernet0/1
O     150.1.5.5/32 [110/391] via 155.1.0.5, 00:00:08, Serial0/0
O     150.1.1.1/32 [110/11] via 155.1.146.1, 00:00:27, Ethernet0/1

```

```
R5#show ip route ospf
 155.1.0.0/24 is subnetted, 3 subnets
O IA   155.1.146.0 [110/391] via 155.1.0.1, 00:00:17, Serial0/0
 160.1.0.0/24 is subnetted, 1 subnets
O E2   160.1.60.0 [110/20] via 155.1.0.1, 00:00:17, Serial0/0
 150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA   150.1.6.6/32 [110/392] via 155.1.0.1, 00:00:17, Serial0/0
O IA   150.1.4.4/32 [110/391] via 155.1.0.4, 00:00:17, Serial0/0
O IA   150.1.1.1/32 [110/391] via 155.1.0.1, 00:00:17, Serial0/0
```

```
R5#show ip ospf database external 160.1.60.0
```


```
      OSPF Router with ID (150.1.5.5) (Process ID 1)
```

```
      Type-5 AS External Link States
```

```
Routing Bit Set on this LSA
LS age: 45
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 160.1.60.0 (External Network Number )
Advertising Router: 150.1.100.100
LS Seq Number: 80000001
Checksum: 0xD7E6
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 150.1.6.6
    External Route Tag: 0
```

NSSA Area ABR External Prefix Filtering

Objective: Configure NSSA ABR to filter external prefixes originated at ASBR



Directions

- Configure routers as per the OSPF scenario “NSSA Area Type 7to5 LSA Translator Election”
- Configure summary address 160.X.6.0/24 on R1 (NSSA translator) with keyword “not-advertise”
- This way you prevent new Type 5 LSA, translated from Type 7 from being propagated

Final Configuration

```
R1:
router ospf 1
summary-address 160.1.6.0 255.255.255.0 not-advertise
```

Verification

Before summary has been configured:

```
R5#show ip route ospf
 155.1.0.0/24 is subnetted, 3 subnets
O IA   155.1.146.0 [110/391] via 155.1.0.1, 00:33:31, Serial0/0
 160.1.0.0/24 is subnetted, 1 subnets
O E2   160.1.60.0 [110/20] via 155.1.0.1, 00:33:31, Serial0/0
 150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA   150.1.6.6/32 [110/392] via 155.1.0.1, 00:33:31, Serial0/0
O IA   150.1.4.4/32 [110/391] via 155.1.0.4, 00:33:31, Serial0/0
O IA   150.1.1.1/32 [110/391] via 155.1.0.1, 00:33:31, Serial0/0

R4#show ip route ospf
```



```

    160.1.0.0/24 is subnetted, 1 subnets
O E2 160.1.60.0 [110/20] via 155.1.146.6, 00:00:09, Ethernet0/1
    150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O    150.1.6.6/32 [110/11] via 155.1.146.6, 00:37:30, Ethernet0/1
O    150.1.5.5/32 [110/391] via 155.1.0.5, 00:37:11, Serial0/0
O    150.1.1.1/32 [110/11] via 155.1.146.1, 00:37:30, Ethernet0/1

```

After:

R5#show ip route ospf

```

    155.1.0.0/24 is subnetted, 3 subnets
O IA 155.1.146.0 [110/391] via 155.1.0.1, 00:34:04, Serial0/0
    150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA 150.1.6.6/32 [110/392] via 155.1.0.1, 00:34:04, Serial0/0
O IA 150.1.4.4/32 [110/391] via 155.1.0.4, 00:34:04, Serial0/0
O IA 150.1.1.1/32 [110/391] via 155.1.0.1, 00:34:04, Serial0/0

```

R1#show ip route ospf

```

    160.1.0.0/24 is subnetted, 1 subnets
O N2 160.1.60.0 [110/20] via 155.1.146.6, 00:34:38, FastEthernet0/0
    150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O    150.1.6.6/32 [110/2] via 155.1.146.6, 00:35:07, FastEthernet0/0
O    150.1.5.5/32 [110/65] via 155.1.0.5, 00:34:38, Serial0/0
O    150.1.4.4/32 [110/2] via 155.1.146.4, 00:35:07, FastEthernet0/0

```

R4#show ip route ospf

```

    160.1.0.0/24 is subnetted, 1 subnets
O N2 160.1.60.0 [110/20] via 155.1.146.6, 00:01:00, Ethernet0/1
    150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O    150.1.6.6/32 [110/11] via 155.1.146.6, 00:35:13, Ethernet0/1
O    150.1.5.5/32 [110/391] via 155.1.0.5, 00:34:53, Serial0/0
O    150.1.1.1/32 [110/11] via 155.1.146.1, 00:35:13, Ethernet0/1

```

R4#show ip ospf database nssa-external 160.1.60.0

OSPF Router with ID (150.1.4.4) (Process ID 1)

Type-7 AS External Link States (Area 1)

Routing Bit Set on this LSA

LS age: 877

Options: (No TOS-capability, Type 7/5 translation, DC)

LS Type: AS External Link

Link State ID: 160.1.60.0 (External Network Number)

Advertising Router: 150.1.6.6

LS Seq Number: 80000002

Checksum: 0xC64

Length: 36

Network Mask: /24

Metric Type: 2 (Larger than any link state path)

TOS: 0


Metric: 20

Forward Address: 150.1.6.6

External Route Tag: 0

NSSA Suppress FA Feature

Objective: Configure NSSA FA-suppress feature to resolve connectivity problem



Directions

- Configure routers as per the OSPF scenario “NSSA Area Type 7to5 LSA Translator Election”
- Configure routers as per scenario “OSPF Type-3 LSA Filtering with Network Ranges”
- Configure NSSA Suppress FA feature on NSSA 7->5 Translating router in order to change FA of translated prefixes

Final Configuration

```
R1:
router ospf 1
area 1 nssa translate type7 suppress-fa
```

Verification

Before Suppress-FA feature has been enabled:

```
R5#show ip route ospf
155.1.0.0/24 is subnetted, 3 subnets
O IA 155.1.146.0 [110/391] via 155.1.0.1, 00:03:43, Serial0/0
150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA 150.1.4.4/32 [110/391] via 155.1.0.4, 00:03:43, Serial0/0
O IA 150.1.1.1/32 [110/391] via 155.1.0.1, 00:03:43, Serial0/0
```

```
R5#show ip os data ext
```

```
OSPF Router with ID (150.1.5.5) (Process ID 1)
```

```

                                Type-5 AS External Link States

LS age: 235
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 160.1.60.0 (External Network Number )
Advertising Router: 150.1.100.100
LS Seq Number: 80000001
Checksum: 0xD7E6
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 150.1.6.6
    External Route Tag: 0

R5#show ip route 150.1.6.6
% Subnet not in table

After that:

R5#show ip route ospf
 155.1.0.0/24 is subnetted, 3 subnets
O IA   155.1.146.0 [110/391] via 155.1.0.1, 00:05:26, Serial0/0
 160.1.0.0/24 is subnetted, 1 subnets
O E2   160.1.60.0 [110/20] via 155.1.0.1, 00:00:04, Serial0/0
 150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA   150.1.4.4/32 [110/391] via 155.1.0.4, 00:05:25, Serial0/0
O IA   150.1.1.1/32 [110/391] via 155.1.0.1, 00:05:26, Serial0/0

R5#show ip ospf data ex


                                OSPF Router with ID (150.1.5.5) (Process ID 1)

                                Type-5 AS External Link States

Routing Bit Set on this LSA
LS age: 13
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 160.1.60.0 (External Network Number )
Advertising Router: 150.1.100.100
LS Seq Number: 80000002
Checksum: 0xAEB2
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

NSSA Area and Default-Route Origination at ABR Part 1

Objective: Configure NSSA ABRs to advertise default-route as Type-3 LSA. Do not assign the default cost to this route.



Directions

- Configure routers as per the OSPF scenario “NSSA Area Type 7to5 LSA Translator Election”
- Configure R1 and R4 as “no-summary” NSSA ABRs
- Configure R1 and R4 to send default route with the cost of 100

Final Configuration

```
R1 & R4:
router ospf 1
 area 1 nssa no-summary
 area 1 default-cost 100
```

Verification

Before:

```
R6#show ip route ospf
 155.1.0.0/24 is subnetted, 2 subnets
O IA   155.1.0.0 [110/65] via 155.1.146.1, 00:46:49, GigabitEthernet0/1
 150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA   150.1.5.5/32 [110/66] via 155.1.146.1, 00:00:15, GigabitEthernet0/1
O      150.1.4.4/32 [110/2] via 155.1.146.4, 00:46:49, GigabitEthernet0/1
O      150.1.1.1/32 [110/2] via 155.1.146.1, 00:46:49, GigabitEthernet0/1
```

After:

```
R6#show ip route ospf
```

```
150.1.0.0/16 is variably subnetted, 3 subnets, 2 masks
O      150.1.4.4/32 [110/2] via 155.1.146.4, 00:50:26, GigabitEthernet0/1
O      150.1.1.1/32 [110/2] via 155.1.146.1, 00:50:26, GigabitEthernet0/1
O*IA 0.0.0.0/0 [110/101] via 155.1.146.4, 00:00:02, GigabitEthernet0/1
      [110/101] via 155.1.146.1, 00:00:02, GigabitEthernet0/1
```

Tear down R1's Link to Area0:

```
R1#conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
R1(config)#inter se 0/0
```

```
R1(config-if)#shut
```

```
R1(config-if)#
```

```
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.5.5 on Serial0/0 from FULL to DOWN,
```


```
Neighbor Down: Interface down or detached
```

```
R6#show ip route ospf
```

```
150.1.0.0/16 is variably subnetted, 3 subnets, 2 masks
O      150.1.4.4/32 [110/2] via 155.1.146.4, 00:00:48, GigabitEthernet0/1
O      150.1.1.1/32 [110/2] via 155.1.146.1, 00:00:48, GigabitEthernet0/1
O*IA 0.0.0.0/0 [110/101] via 155.1.146.4, 00:00:48, GigabitEthernet0/1
```

NSSA Area and Default-Route Origination at ABR Part 2

Objective: Configure NSSA ABRs to advertise default-route as Type-7 LSA. Do not assign the default metric to this route



Directions

- Configure routers as per the OSPF scenario “NSSA Area Type 7to5 LSA Translator Election”
- Configure Area 1 on R1 and R4 as NSSA with “default-information-originate” keyword
- Configure R1 and R4 to send this default route with the metric of 100

Final Configuration

```
R1 & R4:
router ospf 1
 area 1 nssa default-information-originate metric 100
```

Verification

Before default origination has been configured:

```
R6#show ip route ospf
 155.1.0.0/24 is subnetted, 2 subnets
O IA   155.1.0.0 [110/65] via 155.1.146.1, 00:46:49, GigabitEthernet0/1
 150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA   150.1.5.5/32 [110/66] via 155.1.146.1, 00:00:15, GigabitEthernet0/1
O      150.1.4.4/32 [110/2] via 155.1.146.4, 00:46:49, GigabitEthernet0/1
O      150.1.1.1/32 [110/2] via 155.1.146.1, 00:46:49, GigabitEthernet0/1
```

After:

```
R6#show ip route ospf
```

```
155.1.0.0/24 is subnetted, 2 subnets
O IA 155.1.0.0 [110/65] via 155.1.146.1, 00:00:33, GigabitEthernet0/1
150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA 150.1.5.5/32 [110/66] via 155.1.146.1, 00:00:33, GigabitEthernet0/1
O 150.1.4.4/32 [110/2] via 155.1.146.4, 00:03:01, GigabitEthernet0/1
O 150.1.1.1/32 [110/2] via 155.1.146.1, 00:03:01, GigabitEthernet0/1
O*N2 0.0.0.0/0 [110/100] via 155.1.146.4, 00:00:03, GigabitEthernet0/1
[110/100] via 155.1.146.1, 00:00:03, GigabitEthernet0/1
```

Tear down R1's Link to Area0:

R1#**conf t**

Enter configuration commands, one per line. End with CNTL/Z.

R1(config)#**inter se 0/0**

R1(config-if)#**shut**

R1(config-if)#

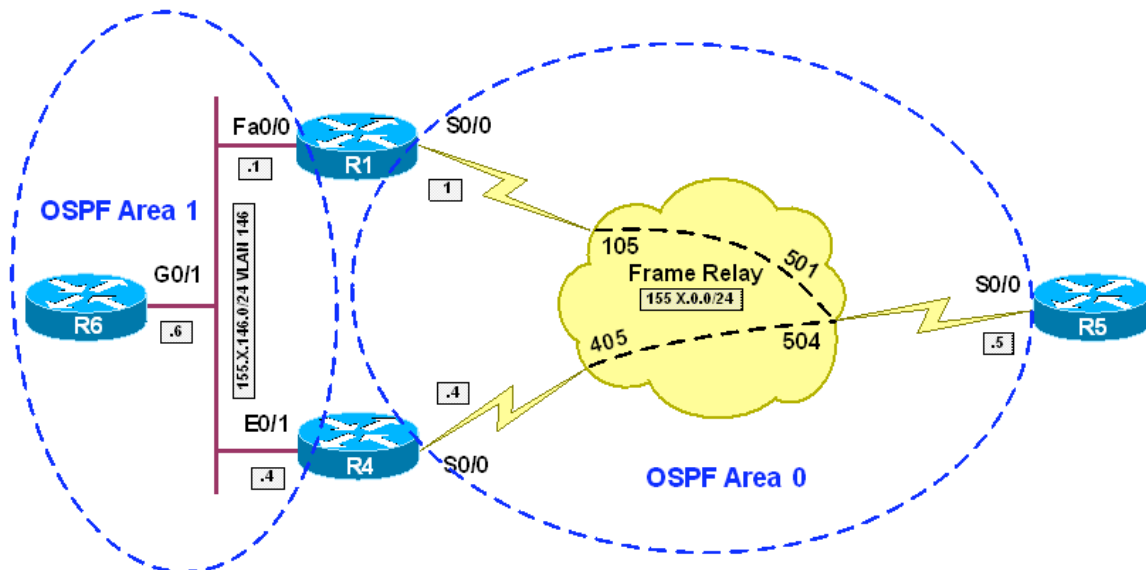
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.5.5 on Serial0/0 from FULL to DOWN,
Neighbor Down: Interface down or detached

R6#**show ip route ospf**

```
155.1.0.0/24 is subnetted, 2 subnets
O IA 155.1.0.0 [110/391] via 155.1.146.4, 00:00:05, GigabitEthernet0/1
150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
O IA 150.1.5.5/32 [110/392] via 155.1.146.4, 00:00:05, GigabitEthernet0/1
O 150.1.4.4/32 [110/2] via 155.1.146.4, 00:03:33, GigabitEthernet0/1
O 150.1.1.1/32 [110/2] via 155.1.146.1, 00:03:33, GigabitEthernet0/1
O*N2 0.0.0.0/0 [110/100] via 155.1.146.4, 00:00:05, GigabitEthernet0/1
```

NSSA Area and Default-Route Origination at ASBR

Objective: Configure NSSA ASBR to advertise default-route as Type-7 LSA



Directions

- Configure routers as per the OSPF scenario “NSSA Area Type 7to5 LSA Translator Election”
- Configure R6 to advertise default-route into NSSA Area (area 1 nssa default-information-originate)
- Configure default static route to Null0 on R6
- Note that you can not use “default-information originate” router command, since it generates Type-5 LSA, which is prohibited in NSSA area

Final Configuration

```
R6:
router ospf 1
 area 1 nssa default-information-originate
!
ip route 0.0.0.0 0.0.0.0 null0
```

Verification

```
R6#show ip route
Codes: 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
```



```

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

    155.1.0.0/24 is subnetted, 2 subnets
C       155.1.146.0 is directly connected, GigabitEthernet0/1
O IA    155.1.0.0 [110/65] via 155.1.146.1, 00:00:10, GigabitEthernet0/1
    160.1.0.0/24 is subnetted, 1 subnets
S       160.1.60.0 is directly connected, Null0
    150.1.0.0/16 is variably subnetted, 4 subnets, 2 masks
C       150.1.6.0/24 is directly connected, Loopback0
O IA    150.1.5.5/32 [110/66] via 155.1.146.1, 00:00:10, GigabitEthernet0/1
O       150.1.4.4/32 [110/2] via 155.1.146.4, 00:00:11, GigabitEthernet0/1
O       150.1.1.1/32 [110/2] via 155.1.146.1, 00:00:11, GigabitEthernet0/1
S*     0.0.0.0/0 is directly connected, Null0

R1#show ip route ospf
    160.1.0.0/24 is subnetted, 1 subnets
O N2    160.1.60.0 [110/20] via 155.1.146.6, 00:02:50, FastEthernet0/0
    150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O IA    150.1.6.0/24 [110/75] via 155.1.0.4, 00:03:01, Serial0/0
O       150.1.6.6/32 [110/2] via 155.1.146.6, 00:03:01, FastEthernet0/0
O       150.1.5.5/32 [110/65] via 155.1.0.5, 00:06:08, Serial0/0
O       150.1.4.4/32 [110/2] via 155.1.146.4, 00:03:01, FastEthernet0/0
O*N2    0.0.0.0/0 [110/1] via 155.1.146.6, 00:00:59, FastEthernet0/0

R4#show ip route ospf
    160.1.0.0/24 is subnetted, 1 subnets
O E2    160.1.60.0 [110/20] via 155.1.146.6, 00:02:57, Ethernet0/1
    150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O       150.1.6.0/24 is a summary, 00:03:08, Null0
O       150.1.6.6/32 [110/11] via 155.1.146.6, 00:03:08, Ethernet0/1
O       150.1.5.5/32 [110/391] via 155.1.0.5, 00:06:17, Serial0/0
O       150.1.1.1/32 [110/11] via 155.1.146.1, 00:03:08, Ethernet0/1
O*E2    0.0.0.0/0 [110/1] via 155.1.146.6, 00:01:06, Ethernet0/1

R5#show ip route ospf
    155.1.0.0/24 is subnetted, 3 subnets
O IA    155.1.146.0 [110/391] via 155.1.0.1, 00:04:24, Serial0/0
    160.1.0.0/24 is subnetted, 1 subnets
O E2    160.1.60.0 [110/20] via 155.1.0.1, 00:04:08, Serial0/0
    150.1.0.0/16 is variably subnetted, 6 subnets, 2 masks
O IA    150.1.6.0/24 [110/401] via 155.1.0.4, 00:04:19, Serial0/0
O IA    150.1.6.6/32 [110/392] via 155.1.0.1, 00:04:19, Serial0/0
O IA    150.1.4.4/32 [110/391] via 155.1.0.4, 00:06:58, Serial0/0
O IA    150.1.1.1/32 [110/391] via 155.1.0.1, 00:06:48, Serial0/0
O*E2    0.0.0.0/0 [110/1] via 155.1.0.1, 00:02:17, Serial0/0

R5#show ip ospf data ex

        OSPF Router with ID (150.1.5.5) (Process ID 1)

                Type-5 AS External Link States

Routing Bit Set on this LSA
LS age: 146
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 0.0.0.0 (External Network Number )
Advertising Router: 150.1.100.100
LS Seq Number: 80000001
Checksum: 0xE3CB
Length: 36
Network Mask: /0

```

```
Metric Type: 2 (Larger than any link state path)
TOS: 0
Metric: 1
Forward Address: 150.1.6.6
External Route Tag: 0
```

```
Routing Bit Set on this LSA
LS age: 257
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 160.1.60.0 (External Network Number )
Advertising Router: 150.1.100.100
LS Seq Number: 80000001
Checksum: 0xD7E6
Length: 36
Network Mask: /24
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 150.1.6.6
    External Route Tag: 0
```

