

***Copyright Information***

---

Copyright © 2008 Internetwork Expert, Inc. All rights reserved.

The following publication, CCIE R&S Lab Workbook Volume I Version 5.0, 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 R&S Lab Workbook Volume I Version 5.0, 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.

# OSPF

 **Note**

Load the *Initial OSPF* initial configurations prior to starting.

## 6.1 OSPF over Broadcast Media

- Enable OSPF on all devices using Process-ID 1.
- On R1 configure OSPF area 1 on the link to VLAN 146.
- On R2 configure OSPF area 51 on the link to VLAN 22.
- On R3 configure OSPF area 2 on the link to SW1.
- On R4 configure OSPF area 1 on the link to VLAN 146.
- On R5 configure OSPF area 3 on the links to VLANs 5 and 58.
- On R6 configure OSPF area 1 on the link to VLAN 146, and area 2 on the link to VLAN 67.
- Do not use the **network** statement under the OSPF process on R1, R2, R3, R4, R5, or R6.
- On SW1 configure OSPF area 2 on all interfaces with an IP address assigned using one single **network** statement.
- On SW2 configure OSPF area 3 on interfaces with IP addresses in the range of 155.X.0.0 – 155.X.127.255 using one single **network** statement.
- On SW3 configure OSPF area 2 on interfaces with IP addresses in the range of 155.X.9.0 – 155.X.9.255 and 155.X.79.0 – 155.X.79.255 using two **network** statements.
- On SW4 configure OSPF area 3 only on interfaces with the exact IP addresses 155.X.10.10 and 155.X.108.10.
- Note any reachability problems throughout the OSPF topology.

## 6.2 OSPF over Non-Broadcast Media

- Configure OSPF area 0 on R1, R2, R3, R4, and R5's connections to the Frame Relay network.
- Without modifying the OSPF network type on any of these devices ensure full reachability is obtained throughout the OSPF domain.

## 6.3 OSPF DR/BDR Election Manipulation

- Configure the network so that R6 is elected the OSPF Designated Router for VLANs 67 and 146.
- If R6 goes down R1 should take over the DR status for VLAN 146.
- When R6 comes back up it should become the BDR.
- Modify the DR/BDR election on the Frame Relay network to ensure that if R5's connection to the Frame Relay network goes down and comes back up, full reachability is still maintained.

## 6.4 OSPF Network Point-to-Point

- Configure OSPF area 0 on the point-to-point Serial link between R4 & R5.
- Configure OSPF area 4 on the point-to-point Serial link between R1 & R3.
- Configure OSPF area 5 on the point-to-point Serial link between R2 & R3.

## 6.5 OSPF Network Point-to-Multipoint

- Remove all `frame-relay map` statements on R1, R2, R3, and R4 with the exception of the mapping to R5, and note any changes in reachability throughout the OSPF domain.
- Modify the Frame Relay network's OSPF configuration to use network type point-to-multipoint, and note any changes in the OSPF database, the routing table, and in reachability throughout the network.

## 6.6 OSPF Network Point-to-Multipoint Non-Broadcast

- Remove the `broadcast` keyword from the `frame-relay map` statements of R1, R2, R3, R4, and R5, and note any changes in reachability throughout the OSPF domain.
- Modify the Frame Relay network's OSPF configuration to use network type point-to-multipoint non-broadcast, and note any changes in the OSPF database, the routing table, and in reachability throughout the network.

# OSPF Solutions

## 6.1 OSPF over Broadcast Media

- Enable OSPF on all devices using Process-ID 1.
- On R1 configure OSPF area 1 on the link to VLAN 146.
- On R2 configure OSPF area 51 on the link to VLAN 22.
- On R3 configure OSPF area 2 on the link to SW1.
- On R4 configure OSPF area 1 on the link to VLAN 146.
- On R5 configure OSPF area 3 on the links to VLANs 5 and 58.
- On R6 configure OSPF area 1 on the link to VLAN 146, and area 2 on the link to VLAN 67.
- Do not use the `network` statement under the OSPF process on R1, R2, R3, R4, R5, or R6.
- On SW1 configure OSPF area 2 on all interfaces with an IP address assigned using one single `network` statement.
- On SW2 configure OSPF area 3 on interfaces with IP addresses in the range of 155.X.0.0 – 155.X.127.255 using one single `network` statement.
- On SW3 configure OSPF area 2 on interfaces with IP addresses in the range of 155.X.9.0 – 155.X.9.255 and 155.X.79.0 – 155.X.79.255 using two `network` statements.
- On SW4 configure OSPF area 3 only on interfaces with the exact IP addresses 155.X.10.10 and 155.X.108.10.
- Note any reachability problems throughout the OSPF topology.

### Configuration

---

```
R1:
interface FastEthernet0/0
 ip ospf 1 area 1
```

```
R2:
interface FastEthernet0/0
 ip ospf 1 area 51
```

```
R3:
interface FastEthernet0/0
 ip ospf 1 area 2
!
router ospf 1
 router-id 150.1.3.3
```

```
R4:
interface FastEthernet0/1
 ip ospf 1 area 1
```

```
R5:
interface FastEthernet0/0
 ip ospf 1 area 3
!
interface FastEthernet0/1
 ip ospf 1 area 3

R6:
interface FastEthernet0/0.67
 ip ospf 1 area 2
!
interface FastEthernet0/0.146
 ip ospf 1 area 1

SW1:
router ospf 1
 network 0.0.0.0 255.255.255.255 area 2

SW2:
router ospf 1
 network 155.1.0.0 0.0.127.255 area 3

SW3:
router ospf 1
 router-id 150.1.9.9
 network 155.1.9.0 0.0.0.255 area 2
 network 155.1.79.0 0.0.0.255 area 2

SW4:
router ospf 1
 network 155.1.10.10 0.0.0.0 area 3
 network 155.1.108.10 0.0.0.0 area 3
```

## Verification

### Note

As of IOS 12.4 there are two ways to enable the OSPF process on an interface. The legacy **network** statement under the OSPF process, and the interface level command **ip ospf [process-id] area [area-id]**. Both accomplish the same thing with one minor exception. If an interface is IP unnumbered, and there is a **network** statement that matches the IP address of the primary interface, both the primary interface and the unnumbered interface will have OSPF enabled on them in the designated area.

Despite common confusion, the **network** statement in OSPF, just like the network statement under the EIGRP process, is not used to *originate* a network advertisement, instead it simply enables the OSPF process on the interface. If multiple network statements overlap the same interface, the most specific match based on the wildcard wins.

In this particular example SW1 enables the OSPF process on all interfaces with the **network 0.0.0.0 255.255.255.255 area 2** command. This means that all interfaces with an IP address assigned will be placed into area 2. This does *not* mean, however, that the network 0.0.0.0/0 itself will be advertised.

Likewise on SW2 the **network 155.1.0.0 0.0.127.255 area 3** command means that any address with the first 17 contiguous bits match the address 155.1.0.0 will be placed into area 3. This does not mean that the network 155.1.0.0/17 will be originated.

The most specific match is seen on SW4, with the **network 155.1.10.10 0.0.0.0 area 3** command. This means that only the interface with the exact IP address 155.1.10.10 will be placed into area 3. If the interfaces 155.1.10.9 or 155.1.10.11, or any other variation, exist on the device, they will not be placed into area 3.

Once the **network** statement or the **ip ospf** statement are configured this can be quickly verified with the **show ip ospf interface brief** command. Note that in the below output there is no functional difference seen between R1 – R6 who used the interface level command to enable OSPF vs. SW1 – SW4 who used the network statement.

```
Rack1R1#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Fa0/0      1    1             155.1.146.1/24  1     DROTH 2/2

Rack1R2#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Fa0/0      1    51            192.10.1.2/24   1     BDR   1/1

Rack1R3#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Fa0/0      1    2             155.1.37.3/24   1     DR    1/1

Rack1R4#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Fa0/1      1    1             155.1.146.4/24  1     DR    2/2

Rack1R5#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Fa0/1      1    3             155.1.5.5/24    1     DR    0/0
Fa0/0      1    3             155.1.58.5/24   1     BDR   1/1

Rack1R6#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Fa0/0.146  1    1             155.1.146.6/24  1     BDR   2/2
Fa0/0.67   1    2             155.1.67.6/24   1     DR    1/1

Rack1SW1#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Lo0        1    2             150.1.7.7/24    1     LOOP  0/0
Fa0/3      1    2             155.1.37.7/24   1     BDR   1/1
Vl179     1    2             155.1.79.7/24   1     BDR   1/1
Vl167     1    2             155.1.67.7/24   1     BDR   1/1
Vl17      1    2             155.1.7.7/24    1     DR    0/0

Rack1SW2#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Po1        1    3             155.1.108.8/24  1     BDR   1/1
Vl158     1    3             155.1.58.8/24   1     DR    1/1
Vl18      1    3             155.1.8.8/24    1     DR    0/0

Rack1SW3#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Vl179     1    2             155.1.79.9/24   1     DR    1/1
Vl19      1    2             155.1.9.9/24    1     DR    0/0

Rack1SW4#show ip ospf interface brief
Interface  PID  Area          IP Address/Mask  Cost  State Nbrs F/C
Po1        1    3             155.1.108.10/24 1     DR    1/1
Vl110     1    3             155.1.10.10/24  1     DR    0/0
```



Once it is verified that the interfaces are configured in the correct areas, the next verification is to check the adjacency state of the OSPF neighbors with the `show ip ospf neighbor` command.

At this point it is important to note the attributes that must be common along with those that must be unique in order for adjacency to establish. The common attributes that must match are the area, timers, authentication, stub flags, MTU, and compatible network types. The attributes that must be unique are the interface IP address and the router-id. The router-id is chosen first based on the process level `router-id` command, second based on the highest active Loopback IP interface, and lastly on the highest active non-Loopback interface IP address.

In this design R3 and SW3 are configured with the duplicate Loopback IP address 222.255.255.255/32. Since this is the highest Loopback IP address it is chosen as the router-id. Since these devices are in the same area, and share the same router-id, both adjacency problems and SPF problems become apparent. This can be seen by log messages on R3 and SW3 along with SPF and routing table calculation problems on SW1.

```
Rack1SW1#debug ip routing
IP routing debugging is on
RT: add 155.1.108.0/24 via 155.1.37.3, ospf metric [110/784]
RT: SET_LAST_RDB for 155.1.146.0/24
    NEW rdb: via 155.1.37.3

RT: del 155.1.108.0/24 via 155.1.37.3, ospf metric [110/784]
RT: delete subnet route to 155.1.108.0/24
```

Since the LSA origination is based on the router-id, SW1 thinks it has two interfaces connecting to the same device. The result of this is that SW1 constantly recalculates SPF forcing it to add and remove prefixes from the routing table because it sees two neighbors with the same router-id advertising different information.

From R3 and SW3's perspective they see each others Router LSA (LSA 1) with the same router-id as the local device, and cannot compute SPF. This can be clearly seen from their log messages to the console.

```
Rack1R3#
%OSPF-4-DUP_RTRID_AREA: Detected router with duplicate router ID
222.255.255.255 in area 2
```

```
Rack1SW3#
%OSPF-4-DUP_RTRID_AREA: Detected router with duplicate router ID
222.255.255.255 in area 2
```

To fix this design problem the router-id needs to be changed to something unique. This can be accomplished by changing or removing IP addresses on the devices, or by issuing the router-id command under the OSPF process. In this solution the router-id command is used to match the unique Loopback0 interfaces of R3 and SW1.

```
Rack1SW3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1SW3(config)#router ospf 1
Rack1SW3(config-router)#router-id 150.1.9.9
Reload or use "clear ip ospf process" command, for this to take effect
Rack1SW3(config-router)#end
Rack1SW3#
%SYS-5-CONFIG_I: Configured from console by console
Rack1SW3#clear ip ospf process
Reset ALL OSPF processes? [no]: yes
Rack1SW3#
```

Once modified on both R3 and SW3, SPF calculation inside area 2 succeeds, and proper routing information can be installed.

```
Rack1R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.6.6	1	FULL/BDR	00:00:37	155.1.146.6	FastEthernet0/0
223.255.255.255	1	FULL/DR	00:00:36	155.1.146.4	FastEthernet0/0

```
Rack1R2#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.10.1.254	1	FULL/DR	00:00:36	192.10.1.254	FastEthernet0/0

```
Rack1R3#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.7.7	1	FULL/BDR	00:00:32	155.1.37.7	FastEthernet0/0

```
Rack1R4#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.1.1	1	FULL/DROTHER	00:00:37	155.1.146.1	FastEthernet0/1
150.1.6.6	1	FULL/BDR	00:00:37	155.1.146.6	FastEthernet0/1

```
Rack1R5#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.8.8	1	FULL/DR	00:00:32	155.1.58.8	FastEthernet0/0

```
Rack1R6#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.1.1	1	FULL/DROTHER	00:00:39	155.1.146.1	FastEthernet0/0.146
223.255.255.255	1	FULL/DR	00:00:39	155.1.146.4	FastEthernet0/0.146
150.1.7.7	1	FULL/BDR	00:00:36	155.1.67.7	FastEthernet0/0.67

```
Rack1SW1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.3.3	1	FULL/DR	00:00:32	155.1.37.3	FastEthernet0/3
150.1.9.9	1	FULL/DR	00:00:38	155.1.79.9	Vlan79
150.1.6.6	1	FULL/DR	00:00:32	155.1.67.6	Vlan67

```
Rack1SW2#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
223.255.255.255	1	FULL/DR	00:00:30	155.1.108.10	Port-channel1
150.1.5.5	1	FULL/BDR	00:00:38	155.1.58.5	Vlan58

```
Rack1SW3#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.7.7	1	FULL/BDR	00:00:37	155.1.79.7	Vlan79

```
Rack1SW4#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.8.8	1	FULL/BDR	00:00:38	155.1.108.8	Port-channel1

Once adjacency has been established a fundamental underlying design issue still exists in the network. At this point only areas 1, 2, 3, and 51 are configured. The backbone area 0 is not configured on any links. This implies that the devices can route within their own area (Intra-Area), but not between areas (Inter-Area). This is due to the fact that the Area Border Router (ABR) who connects one area to area zero is responsible for generating the Network Summary LSA (LSA 3) that describes the Inter-Area routes. The result of this can be seen by viewing the routing table.

R1 and R4 do not see any routes installed because the only link advertised in area 1 is their connected VLAN 146 interface.

```
Rack1R1#show ip route ospf
```

```
Rack1R4#show ip route ospf
```

R2 sees area 51 routes from its adjacent neighbor, BB2.

```
Rack1R2#show ip route ospf
```

```
51.0.0.0/32 is subnetted, 1 subnets
O E2 51.51.51.51 [110/20] via 192.10.1.254, 00:08:46,
```

R3, R6, SW1, and SW3 know about the routes in area 2, but not the routes in area 1.

```
Rack1R3#show ip route ospf
 155.1.0.0/24 is subnetted, 8 subnets
O   155.1.9.0 [110/3] via 155.1.37.7, 00:01:09, FastEthernet0/0
O   155.1.7.0 [110/2] via 155.1.37.7, 00:01:09, FastEthernet0/0
O   155.1.79.0 [110/2] via 155.1.37.7, 00:01:09, FastEthernet0/0
O   155.1.67.0 [110/2] via 155.1.37.7, 00:01:09, FastEthernet0/0
 150.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
O   150.1.7.7/32 [110/2] via 155.1.37.7, 00:01:10, FastEthernet0/0
```

```
Rack1R6#show ip route ospf
 155.1.0.0/24 is subnetted, 6 subnets
O   155.1.9.0 [110/3] via 155.1.67.7, 00:01:28, FastEthernet0/0.67
O   155.1.7.0 [110/2] via 155.1.67.7, 00:01:28, FastEthernet0/0.67
O   155.1.37.0 [110/2] via 155.1.67.7, 00:01:28, FastEthernet0/0.67
O   155.1.79.0 [110/2] via 155.1.67.7, 00:01:28, FastEthernet0/0.67
 150.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
O   150.1.7.7/32 [110/2] via 155.1.67.7, 00:01:28,
```

```
Rack1SW1#show ip route ospf
 155.1.0.0/24 is subnetted, 5 subnets
O   155.1.9.0 [110/2] via 155.1.79.9, 00:01:35, Vlan79
```

```
Rack1SW3#show ip route ospf
 155.1.0.0/24 is subnetted, 5 subnets
O   155.1.7.0 [110/2] via 155.1.79.7, 00:01:48, Vlan79
O   155.1.37.0 [110/2] via 155.1.79.7, 00:01:48, Vlan79
O   155.1.67.0 [110/2] via 155.1.79.7, 00:01:48, Vlan79
 150.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
O   150.1.7.7/32 [110/2] via 155.1.79.7, 00:01:49, Vlan79
```

R5, SW2, and SW4 know about the routes in area 3.

```
Rack1R5#show ip route ospf
 155.1.0.0/24 is subnetted, 7 subnets
O   155.1.10.0 [110/3] via 155.1.58.8, 00:03:27, FastEthernet0/0
O   155.1.8.0 [110/2] via 155.1.58.8, 00:03:27, FastEthernet0/0
O   155.1.108.0 [110/2] via 155.1.58.8, 00:03:27, FastEthernet0/0
```

```
Rack1SW2#show ip route ospf
 155.1.0.0/24 is subnetted, 5 subnets
O   155.1.10.0 [110/2] via 155.1.108.10, 00:03:41, Port-channel1
O   155.1.5.0 [110/2] via 155.1.58.5, 00:03:41, Vlan58
```

```
Rack1SW4#show ip route ospf
 155.1.0.0/24 is subnetted, 5 subnets
O   155.1.8.0 [110/2] via 155.1.108.8, 00:03:53, Port-channel1
O   155.1.5.0 [110/3] via 155.1.108.8, 00:03:53, Port-channel1
O   155.1.58.0 [110/2] via 155.1.108.8, 00:03:53, Port-channel1
```

Intra-area routing is successful but inter-area routing fails.

```
Rack1SW3#ping 155.1.67.6
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 155.1.67.6, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

```
Rack1SW3#debug ip packet
```

```
IP packet debugging is on
```

```
Rack1SW3#ping 155.1.146.6
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 155.1.146.6, timeout is 2 seconds:
```

```
IP: s=155.1.9.9 (local), d=155.1.146.6, len 100, unroutable.
```

```
IP: s=155.1.9.9 (local), d=155.1.146.6, len 100, unroutable.
```

```
IP: s=155.1.9.9 (local), d=155.1.146.6, len 100, unroutable.
```

```
IP: s=155.1.9.9 (local), d=155.1.146.6, len 100, unroutable.
```

```
IP: s=155.1.9.9 (local), d=155.1.146.6, len 100, unroutable.
```

```
Success rate is 0 percent (0/5)
```

## 6.2 OSPF over Non-Broadcast Media

- Configure OSPF area 0 on R1, R2, R3, R4, and R5's connections to the Frame Relay network.
- Without modifying the OSPF network type on any of these devices ensure full reachability is obtained throughout the OSPF domain.

### Configuration

---

```
R1:
interface Serial0/0
 ip ospf 1 area 0
 frame-relay map ip 155.1.0.2 105
 frame-relay map ip 155.1.0.3 105
 frame-relay map ip 155.1.0.4 105
```

```
R2:
interface Serial0/0
 ip ospf 1 area 0
 frame-relay map ip 155.1.0.1 205
 frame-relay map ip 155.1.0.3 205
 frame-relay map ip 155.1.0.4 205
```

```
R3:
interface Serial1/0
 ip ospf 1 area 0
 frame-relay map ip 155.1.0.1 305
 frame-relay map ip 155.1.0.2 305
 frame-relay map ip 155.1.0.4 305
```

```
R4:
interface Serial0/0
 ip ospf 1 area 0
 frame-relay map ip 155.1.0.1 405
 frame-relay map ip 155.1.0.2 405
 frame-relay map ip 155.1.0.3 405
```

```
R5:
interface Serial0/1
 !
router ospf 1
 neighbor 155.1.0.3
 neighbor 155.1.0.1
 neighbor 155.1.0.2
 neighbor 155.1.0.4
```

**Verification** **Note**

In this solution OSPF is enabled on the interfaces with the interface level **ip ospf** command. The **network** statement could have also been used under the OSPF process. The first goal of the section is simply to ensure that the **show ip ospf interface** output indicates that the Frame Relay interfaces are running OSPF in area 0.

```
Rack1R1#show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Se0/0	1	0	155.1.0.1/24	64	BDR	1/1	
Fa0/0	1	1	155.1.146.1/24	1	DROTH	2/2	

```
Rack1R2#show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Se0/0	1	0	155.1.0.2/24	64	BDR	1/1	
Fa0/0	1	51	192.10.1.2/24	1	BDR	1/1	

```
Rack1R3#show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Se1/0	1	0	155.1.0.3/24	781	BDR	1/1	
Fa0/0	1	2	155.1.37.3/24	1	DR	1/1	

```
Rack1R4#show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Se0/0	1	0	155.1.0.4/24	64	BDR	1/1	
Fa0/1	1	1	155.1.146.4/24	1	DR	2/2	

```
Rack1R5#show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Se0/0	1	0	155.1.0.5/24	64	DR	4/4	
Fa0/1	1	3	155.1.5.5/24	1	DR	0/0	
Fa0/0	1	3	155.1.58.5/24	1	BDR	1/1	

The next consideration for establishing reachability in this design is the difference between the default OSPF network types on NBMA Frame Relay interfaces versus broadcast Ethernet interfaces. As seen in the **show ip ospf interface** output the default network type for a multipoint Frame Relay interface is **NON\_BROADCAST**.

```
Rack1R5#show ip ospf interface Serial0/0
Serial0/0 is up, line protocol is up
  Internet Address 155.1.0.5/24, Area 0
  Process ID 1, Router ID 150.1.5.5, Network Type NON_BROADCAST, Cost: 64
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 150.1.5.5, Interface address 155.1.0.5
  Backup Designated router (ID) 223.255.255.255, Interface address 155.1.0.4
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:16
  Supports Link-local Signaling (LLS)
  Index 1/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 4
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 4, Adjacent neighbor count is 4
    Adjacent with neighbor 150.1.3.3
    Adjacent with neighbor 150.1.1.1
    Adjacent with neighbor 150.1.2.2
    Adjacent with neighbor 223.255.255.255 (Backup Designated Router)
```

```
Rack1R1#show ip ospf interface Serial0/0
Serial0/0 is up, line protocol is up
  Internet Address 155.1.0.1/24, Area 0
  Process ID 1, Router ID 150.1.1.1, Network Type NON_BROADCAST, Cost: 64
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 150.1.5.5, Interface address 155.1.0.5
  Backup Designated router (ID) 150.1.1.1, Interface address 155.1.0.1
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:12
  Supports Link-local Signaling (LLS)
  Index 1/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.5.5 (Designated Router)
  Suppress hello for 0 neighbor(s)
```



The Non-Broadcast network type means that there will be a DR/BDR election, and that hellos are exchanged as unicast. In order to unicast OSPF hellos the **neighbor** statement needs to be configured under the OSPF process of the DR. Once the DROTHERs and/or BDR hear the unicast hellos from the DR, they will automatically respond back with their own unicast hellos. This implies that the **neighbor** statement can be configured everywhere, but is only required on the DR. Once R5 is configured with the neighbor statement the `show ip ospf neighbor` output should be checked to verify adjacency.

```
Rack1R5#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.3.3	1	FULL/DROTHER	00:01:34	155.1.0.3	Serial0/0
150.1.1.1	1	FULL/DROTHER	00:01:48	155.1.0.1	Serial0/0
150.1.2.2	1	FULL/DROTHER	00:01:59	155.1.0.2	Serial0/0
223.255.255.255	1	FULL/BDR	00:01:45	155.1.0.4	Serial0/0
150.1.8.8	1	FULL/DR	00:00:33	155.1.58.8	FastEthernet0/0

Once adjacency is established in area 0, Inter-Area routing advertisements can be propagated throughout the entire topology. This is due to the fact that R1, R2, R3, R4, and R5 are now ABRs, and can originate the Network Summary LSA (LSA 3) describing Inter-Area routes to the other neighbors in their attached areas.

From the `show ip ospf database` output on R5 the *Summary Net link States (Area 0)* shows ABRs that are advertising information from other areas into area 0. For example R2 is advertising the link 192.10.1.0/24 from area 51 into area 0, and can be seen as the last entry in this category.

```
Rack1R5#show ip ospf database
```

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

```
Router Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.1.1	150.1.1.1	1489	0x80000003	0x007B03	1
150.1.2.2	150.1.2.2	1670	0x80000003	0x006712	1
150.1.3.3	150.1.3.3	1452	0x80000003	0x00ECB7	1
150.1.5.5	150.1.5.5	1138	0x80000004	0x002940	1
223.255.255.255	223.255.255.255	982	0x80000006	0x008962	1

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

Link ID	ADV Router	Age	Seq#	Checksum
155.1.0.5	150.1.5.5	981	0x80000002	0x001283

## Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
150.1.7.7	150.1.3.3	1704	0x80000002	0x00DB1B
155.1.5.0	150.1.5.5	1709	0x80000002	0x00D225
155.1.7.0	150.1.3.3	1704	0x80000002	0x00E018
155.1.8.0	150.1.5.5	1709	0x80000002	0x00BB38
155.1.9.0	150.1.3.3	1706	0x80000002	0x00D421
155.1.10.0	150.1.5.5	1710	0x80000002	0x00AF41
155.1.37.0	150.1.3.3	1706	0x80000002	0x008B50
155.1.58.0	150.1.5.5	1710	0x80000002	0x008939
155.1.67.0	150.1.3.3	1706	0x80000002	0x004A72
155.1.79.0	150.1.3.3	1706	0x80000002	0x00C5EA
155.1.108.0	150.1.5.5	1710	0x80000002	0x006B24
155.1.146.0	150.1.1.1	1739	0x80000002	0x00F180
155.1.146.0	223.255.255.255	983	0x80000003	0x0071B9
192.10.1.0	150.1.2.2	1672	0x80000002	0x00D6FC

## Summary ASB Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
192.10.1.254	150.1.2.2	1672	0x80000002	0x00D201

## Router Link States (Area 3)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.5.5	150.1.5.5	1710	0x8000002B	0x004017	2
150.1.8.8	150.1.8.8	219	0x80000031	0x007D8E	3
223.255.255.255	223.255.255.255	443	0x8000002C	0x00A0C9	2

## Net Link States (Area 3)

Link ID	ADV Router	Age	Seq#	Checksum
155.1.58.8	150.1.8.8	219	0x80000029	0x009AA7
155.1.108.10	223.255.255.255	443	0x80000029	0x00ECAA

Since R5 is the ABR connecting area 0 to area 3, R5 re-originates all Inter-Area routes coming from other neighbors in area 0 as LSA type 3 routes into area 3. This is seen under the *Summary Net Link States (Area 3)* field.

## Summary Net Link States (Area 3)

Link ID	ADV Router	Age	Seq#	Checksum
150.1.7.7	150.1.5.5	1135	0x80000001	0x00466D
155.1.0.0	150.1.5.5	1712	0x80000002	0x00823B
155.1.7.0	150.1.5.5	1135	0x80000001	0x004B6A
155.1.9.0	150.1.5.5	1135	0x80000001	0x003F73
155.1.37.0	150.1.5.5	1135	0x80000001	0x00F5A2
155.1.67.0	150.1.5.5	1135	0x80000001	0x00B4C4
155.1.79.0	150.1.5.5	1135	0x80000001	0x00303D
155.1.146.0	150.1.5.5	1451	0x80000002	0x0040E9
192.10.1.0	150.1.5.5	1135	0x80000001	0x00345A

## Summary ASB Link States (Area 3)

Link ID	ADV Router	Age	Seq#	Checksum
192.10.1.254	150.1.5.5	1136	0x80000001	0x00305E

## Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
51.51.51.51	192.10.1.254	1706	0x80000052	0x00F9BD	0

The result of the Inter-Area routing advertisements can be seen in the routing table via the `show ip route ospf` output. These new routes are denoted as `O IA` for OSPF Inter-Area. The next important design issue in this example is how the DR processes routing advertisements on the NBMA segment.

First, if we view the `show ip ospf neighbor` output on R1 – R5 we can see that the spokes of the Frame Relay network only form adjacency with R5, the DR. There is no direct adjacency between R1 and R2, because there is no direct layer 2 Frame Relay PVC between them. If you reached this step and R5 is *not* the DR for the segment, an additional problem will be seen that is covered in the next section regarding the DR/BDR election process. However for the sake of this example let's assume that R5 is the DR.

```
Rack1R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:37	155.1.0.5	Serial0/0
150.1.6.6	1	FULL/BDR	00:00:29	155.1.146.6	FastEthernet0/0
223.255.255.255	1	FULL/DR	00:00:39	155.1.146.4	FastEthernet0/0

```
Rack1R2#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:30	155.1.0.5	Serial0/0
192.10.1.254	1	FULL/DR	00:00:33	192.10.1.254	FastEthernet0/0

```
Rack1R3#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:54	155.1.0.5	Serial1/0
150.1.7.7	1	FULL/BDR	00:00:30	155.1.37.7	FastEthernet0/0

```
Rack1R4#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:50	155.1.0.5	Serial0/0
150.1.1.1	1	FULL/DROTHER	00:00:32	155.1.146.1	FastEthernet0/1
150.1.6.6	1	FULL/BDR	00:00:32	155.1.146.6	FastEthernet0/1

```
Rack1R5#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.3.3	1	FULL/DROTHER	00:01:34	155.1.0.3	Serial0/0
150.1.1.1	1	FULL/DROTHER	00:01:48	155.1.0.1	Serial0/0
150.1.2.2	1	FULL/DROTHER	00:01:59	155.1.0.2	Serial0/0
223.255.255.255	1	FULL/BDR	00:01:45	155.1.0.4	Serial0/0
150.1.8.8	1	FULL/DR	00:00:33	155.1.58.8	FastEthernet0/0

The DR state on the segment means that R5 is responsible for replicating LSA information between its adjacent neighbors. For example R3 sends R5 the LSA type-3 route 155.1.9.0/24 describing SW3's link to VLAN 9. When R2 learns this on the Frame Relay network it comes from the DR, R5, but the next-hop value of the route is 155.1.0.3, the originator of the LSA type-3 advertisement. However when R2 learns the route to 155.1.10.0/24, the LSA type-3 originator is R5, so the next-hop value towards this route is also R5. This is seen in the below `show ip route ospf` output.

```
Rack1R2#show ip route ospf
 51.0.0.0/32 is subnetted, 1 subnets
O E2   51.51.51.51 [110/20] via 192.10.1.254, 00:04:23, FastEthernet0/0
 155.1.0.0/24 is subnetted, 13 subnets
O IA   155.1.146.0 [110/65] via 155.1.0.4, 00:04:23, Serial0/0
       [110/65] via 155.1.0.1, 00:04:23, Serial0/0
O IA   155.1.10.0 [110/67] via 155.1.0.5, 00:04:23, Serial0/0
O IA   155.1.8.0 [110/66] via 155.1.0.5, 00:04:23, Serial0/0
O IA   155.1.9.0 [110/67] via 155.1.0.3, 00:04:23, Serial0/0
O IA   155.1.7.0 [110/66] via 155.1.0.3, 00:04:23, Serial0/0
O IA   155.1.5.0 [110/65] via 155.1.0.5, 00:04:23, Serial0/0
O IA   155.1.58.0 [110/65] via 155.1.0.5, 00:04:23, Serial0/0
O IA   155.1.37.0 [110/65] via 155.1.0.3, 00:04:23, Serial0/0
O IA   155.1.79.0 [110/66] via 155.1.0.3, 00:04:23, Serial0/0
O IA   155.1.67.0 [110/66] via 155.1.0.3, 00:04:23, Serial0/0
O IA   155.1.108.0 [110/66] via 155.1.0.5, 00:04:23, Serial0/0
 150.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
O IA   150.1.7.7/32 [110/66] via 155.1.0.3, 00:04:24, Serial0/0
```

This is due to the fact that the DR passes the routes along, but it does not modify any of the routing lookup attributes. The result of this behavior is seen when route recursion is performed to the final destination. From the *IP Routing* section of this publication recall that when a routing lookup is done the router also needs to perform layer 3 to layer 2 mapping for the next-hop value on the link. Let's look at what happens when R2 tries to send traffic to 155.1.10.10.

First R2 finds the longest match to 155.1.10.10, which is 155.1.10.0/24 via the next-hop 155.1.0.5

```
Rack1R2#show ip route 155.1.10.10
Routing entry for 155.1.10.0/24
  Known via "ospf 1", distance 110, metric 67, type inter area
  Last update from 155.1.0.5 on Serial0/0, 00:05:20 ago
  Routing Descriptor Blocks:
  * 155.1.0.5, from 150.1.5.5, 00:05:20 ago, via Serial0/0
    Route metric is 67, traffic share count is 1
```

R2 then needs to do another recursive lookup to find out how to forward towards 155.1.0.5. This is seen via the match 155.1.0.0/24 out Serial0/0.

```
Rack1R2#show ip route 155.1.0.5
Routing entry for 155.1.0.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
    * directly connected, via Serial0/0
      Route metric is 0, traffic share count is 1
```

Since Serial0/0 is a multipoint Frame Relay interface, R2 must now find out which Frame Relay DLCI is associated with the next-hop 155.1.0.5. From the Frame Relay mapping table R2 sees that DLCI 205 is used.

```
Rack1R2#show frame-relay map | include 155.1.0.5
Serial0/0 (up): ip 155.1.0.5 dlci 205(0xCD,0x30D0), static,
```

In this case the layer 3 routing lookup is successful, and the layer 2 resolution is successful. The result is a successful ICMP PING to the destination.

```
Rack1R2#ping 155.1.10.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 155.1.10.10, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/61 ms
```

Now let's look at what R2 does when it tries to send traffic towards 155.1.9.9. The routing lookup for this destination says the longest match is 155.1.9.0/24 via 155.1.0.3. Even though the route came from R5, the next-hop is still the originator, R3, because the DR does not update this.

```
Rack1R2#show ip route 155.1.9.9
Routing entry for 155.1.9.0/24
  Known via "ospf 1", distance 110, metric 67, type inter area
  Last update from 155.1.0.3 on Serial0/0, 00:05:42 ago
  Routing Descriptor Blocks:
    * 155.1.0.3, from 150.1.3.3, 00:05:42 ago, via Serial0/0
      Route metric is 67, traffic share count is 1
```

R2's recursive lookup for 155.1.0.3 says that it is directly connected via the multipoint Frame Relay interface, Serial0/0.

```
Rack1R2#show ip route 155.1.0.3
Routing entry for 155.1.0.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
    * directly connected, via Serial0/0
      Route metric is 0, traffic share count is 1
```

R2 now tries to find the Frame Relay PVC that is used to get towards 155.1.0.3. Since there is no mapping configured, and since R2 and R3 cannot use Inverse-ARP to resolve each others non-directly connected interfaces, this lookup fails.

```
Rack1R2#show frame-relay map | include 155.1.0.3
```

The result of this layer 2 failure can be seen from both the **debug ip packet** and **debug frame-relay packet** output.

```
Rack1R2#debug ip packet
IP packet debugging is on
```

```
Rack1R2#debug frame-relay packet
Frame Relay packet debugging is on
```

```
Rack1R2#ping 155.1.9.9
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 155.1.9.9, timeout is 2 seconds:
```

```
IP: tableid=0, s=155.1.0.2 (local), d=155.1.9.9 (Serial0/0), routed via RIB
IP: s=155.1.0.2 (local), d=155.1.9.9 (Serial0/0), len 100, sending
Serial0/0:Encaps failed--no map entry link 7(IP)
IP: s=155.1.0.2 (local), d=155.1.9.9 (Serial0/0), len 100, encapsulation failed
<output omitted>
Success rate is 0 percent (0/5)
```

The first portion of this debug output is from the **debug ip packet**, and says that the destination 155.1.9.9 is out Serial0/0 and is *routed via RIB*. This means that the layer 3 routing lookup was successful. A failure in the routing lookup was seen in the previous task, and would show the *unroutable* output in this debug.

The next portion is the routing process sending the packet to the layer 2 process. The **debug frame-relay packet** output says that *Encaps failed* and that there is *no map entry link*. **Debug ip packet** says that *encapsulation failed*. These two outputs mean that the router does not know which layer 2 PVC to use when building the frame towards the next-hop in the network. Although this technically is not an issue directly related to OSPF, it is part of the overall network design that must be solved.

There are basically two ways to solve this design problem, either change the next-hop to something the router *does* know how to resolve, or allow the router to properly resolve the current next-hop. In this solution the current next-hop is resolved by adding additional **frame-relay map** statements. Changing the next-hop value is explored in the next few examples with the point-to-multipoint OSPF network type.

By adding additional mappings for R2 to reach R3 via R5, and for R3 to reach R2 via R5, layer 2 encapsulation is successful, and end-to-end transport is achieved.

```
Rack1R2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R2(config)#interface Serial0/0
Rack1R2(config-if)#frame-relay map ip 155.1.0.3 205
Rack1R2(config-if)#end
Rack1R2#

Rack1R3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R3(config)#interface Serial1/0
Rack1R3(config-if)#frame-relay map ip 155.1.0.2 305
Rack1R3(config-if)#end
Rack1R3#

Rack1R2#ping 155.1.9.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 155.1.9.9, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/117/125 ms
```

Based on this it is required that all spokes of the Frame Relay network, R1, R2, R3, and R4, must have **frame-relay map** statements for all other spokes via the DLCI to R5. Only once these additional statements are configured will there be end-to-end transport throughout the entire OSPF network.

### 6.3 OSPF DR/BDR Election Manipulation

- Configure the network so that R6 is elected the OSPF Designated Router for VLANs 67 and 146.
- If R6 goes down R1 should take over the DR status for VLAN 146.
- When R6 comes back up it should become the BDR.
- Modify the DR/BDR election on the Frame Relay network to ensure that if R5's connection to the Frame Relay network goes down and comes back up, full reachability is still maintained.

#### **Configuration**

---

```
R1:
interface Serial0/0
 ip ospf priority 0

R2:
interface Serial0/0
 ip ospf priority 0

R3:
interface Serial1/0
 ip ospf priority 0

R4:
interface Serial0/0
 ip ospf priority 0
!
interface FastEthernet0/1
 ip ospf priority 0

R6:
interface FastEthernet0/0.67
 ip ospf priority 255
!
interface FastEthernet0/0.146
 ip ospf priority 255
```



## Verification

### Note

The OSPF DR/BDR election is determined based on the interface level OSPF priority value along with the router-id. The device with the highest priority is elected the DR, and the device with the second highest priority is elected the BDR. If there is a tie in the priority value the device with the higher router-id is more likely to be elected the DR or BDR. However, in certain designs the election may become unpredictable.

This unpredictability is due to the fact that the OSPF DR/BDR election does not support preemption like IS-IS does for its DIS election. Preemption means that if a new device comes onto the segment with a higher priority or router-id, it can take the DR/BDR status away from the current device. Since OSPF does not support this, new devices must wait for a failure of the DR or BDR before the next election occurs. Additionally in certain cases the order in which the routers load their OSPF process can influence the election.

In the below case on the LAN segment of VLAN 146, R1, R4, and R6 compete for the DR/BDR election. Since R4 has the Loopback interface 223.255.255.255/32 configured it is most likely to be elected the DR by default. This is based on the fact that the default interface priority is 1, and therefore if all three routers load the OSPF process at the same time R4 wins the election. R6 with the router-id 150.1.6.6 is next in line, with R1 having the router-id 150.1.1.1 being last.

By changing R6's interface level priority to 255 it is most likely to be elected the DR, and by changing R4's priority to 0 it means it can *never* be elected the DR or BDR. The step-by-step election can be tracked as follows.

R1's `show ip ospf neighbor` and `show ip ospf interface` output indicate that R6 is the DR with a priority of 255, R1 is the BDR with a priority of 1, and R4 is a DROTHER since it has a priority of 0.

```
Rack1R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:58	155.1.0.5	Serial10/0
150.1.6.6	255	FULL/DR	00:00:39	155.1.146.6	FastEthernet0/0
223.255.255.255	0	FULL/DROTHER	00:00:39	155.1.146.4	FastEthernet0/0

```

Rack1R1#show ip ospf interface FastEthernet0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 155.1.146.1/24, Area 1
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 1
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 150.1.6.6, Interface address 155.1.146.6
  Backup Designated router (ID) 150.1.1.1, Interface address 155.1.146.1
  Old designated Router (ID) 223.255.255.255, Interface address 155.1.146.4
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  oob-resync timeout 40
  Hello due in 00:00:01
  Supports Link-local Signaling (LLS)
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 13
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 150.1.6.6 (Designated Router)
    Adjacent with neighbor 223.255.255.255
  Suppress hello for 0 neighbor(s)

```

R6's link to VLAN 146 goes down. After R1 and R4's dead timer expires a new DR/BDR election occurs.

```

Rack1R6#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R6(config)#interface FastEthernet0/0.146
Rack1R6(config-subif)#shutdown

```

When R1 detects this event it asks all other neighbors on the segment to perform a new election. Since R4's priority is 0 it does not send a response back to R1. The result is that R1 is promoted from the BDR status to the DR, but there is no new BDR elected.

```

Rack1R1#debug ip ospf adj
OSPF adjacency events debugging is on

Rack1R1#
OSPF: Rcv LS UPD from 223.255.255.255 on FastEthernet0/0 length 64 LSA count 1
OSPF: Rcv LS UPD from 223.255.255.255 on FastEthernet0/0 length 64 LSA count 1
OSPF: 150.1.6.6 address 155.1.146.6 on FastEthernet0/0 is dead
OSPF: 150.1.6.6 address 155.1.146.6 on FastEthernet0/0 is dead, state DOWN
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.6.6 on FastEthernet0/0 from FULL to DOWN,
Neighbor Down: Dead timer expired
OSPF: Neighbor change Event on interface FastEthernet0/0
OSPF: DR/BDR election on FastEthernet0/0
OSPF: Elect BDR 150.1.1.1
OSPF: Elect DR 150.1.1.1
OSPF: Elect BDR 0.0.0.0
OSPF: Elect DR 150.1.1.1
DR: 150.1.1.1 (Id) BDR: none
OSPF: Remember old DR 150.1.6.6 (id)
OSPF: Build router LSA for area 1, router ID 150.1.1.1, seq 0x80000034
OSPF: Build network LSA for FastEthernet0/0, router ID 150.1.1.1
OSPF: Build network LSA for FastEthernet0/0, router ID 150.1.1.1

```

This can also be verified by the `show ip ospf neighbor` or `show ip ospf interface` output on R1 or R4. R1 sees R4 as a neighbor with a priority of 0, therefore it is a DROTHER.

```
Rack1R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:57	155.1.0.5	Serial0/0
223.255.255.255	0	FULL/DROTHER	00:00:39	155.1.146.4	FastEthernet0/0

R4 indicates that there is a DR on the segment with the RID 150.1.1.1, but no BDR.

```
Rack1R4#show ip ospf interface FastEthernet0/1
```

```
FastEthernet0/1 is up, line protocol is up
  Internet Address 155.1.146.4/24, Area 1
  Process ID 1, Router ID 223.255.255.255, Network Type BROADCAST, Cost: 1
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State DROTHER, Priority 0
  Designated Router (ID) 150.1.1.1, Interface address 155.1.146.1
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:06
  Supports Link-local Signaling (LLS)
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 14
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.1.1 (Designated Router)
```

When R6's connection to VLAN 146 comes back up R1 learns about the neighbor and a new election occurs. However since there is no preemption for the DR election, R6 can only be elected the BDR, even though its priority is higher than R1.

```
Rack1R6(config-subif)#no shutdown
Rack1R6(config-subif)#
```

```
Rack1R1#
OSPF: Rcv DBD from 150.1.6.6 on FastEthernet0/0 seq 0x160C opt 0x52 flag 0x7
len 32 mtu 1500 state INIT
OSPF: 2 Way Communication to 150.1.6.6 on FastEthernet0/0, state 2WAY
OSPF: Neighbor change Event on interface FastEthernet0/0
OSPF: DR/BDR election on FastEthernet0/0
OSPF: Elect BDR 150.1.6.6
OSPF: Elect DR 150.1.1.1
      DR: 150.1.1.1 (Id)   BDR: 150.1.6.6 (Id)
<output omitted>
```

```
Rack1R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:40	155.1.0.5	Serial0/0
150.1.6.6	255	FULL/BDR	00:00:30	155.1.146.6	FastEthernet0/0
223.255.255.255	0	FULL/DROTHER	00:00:32	155.1.146.4	FastEthernet0/0

```
Rack1R1#show ip ospf interface FastEthernet0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 155.1.146.1/24, Area 1
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 1
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 150.1.1.1, Interface address 155.1.146.1
  Backup Designated router (ID) 150.1.6.6, Interface address 155.1.146.6
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:00
  Supports Link-local Signaling (LLS)
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 13
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 2, Adjacent neighbor count is 2
    Adjacent with neighbor 150.1.6.6 (Backup Designated Router)
    Adjacent with neighbor 223.255.255.255
  Suppress hello for 0 neighbor(s)
```

For LAN segments such as VLAN 146 it technically does not matter which device is the DR or the BDR, because everyone has direct layer 2 connectivity with each other. The only design issue of the DR/BDR placement in this case is a function of the memory and CPU resources of the DR/BDR and how many neighbors are on the segment. For partially meshed NBMA networks however, a larger design issue becomes evident.

On the Frame Relay network between R1, R2, R3, R4, and R5, R5 was elected as the DR due to the order of operations in which OSPF was configured. First OSPF was enabled on R5's link, then R1, R2, R3, and R4 in that order. Since R5's OSPF process loaded first it elected itself as the DR. The neighbor with the next highest router-id, R4 with 223.255.255.255/32, was then elected the BDR. This can be seen from the `show ip ospf interface` output on R5. All routers still have the default interface priority of 1.

```
Rack1R5#show ip ospf interface Serial0/0
Serial0/0 is up, line protocol is up
  Internet Address 155.1.0.5/24, Area 0
  Process ID 1, Router ID 150.1.5.5, Network Type NON_BROADCAST, Cost: 64
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 150.1.5.5, Interface address 155.1.0.5
  Backup Designated router (ID) 223.255.255.255, Interface address 155.1.0.4
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:02
  Supports Link-local Signaling (LLS)
  Index 1/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 4
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 4, Adjacent neighbor count is 4
    Adjacent with neighbor 150.1.3.3
    Adjacent with neighbor 150.1.1.1
    Adjacent with neighbor 150.1.2.2
    Adjacent with neighbor 223.255.255.255 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
```

When R5's link to the Frame Relay network goes down all spokes discover this after the dead timer of 120 seconds expires. A new election then occurs, and each of these routers elect themselves as the DR, and no one as the BDR.

```
Rack1R5#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R5(config)#interface Serial0/0
Rack1R5(config-if)#shutdown
Rack1R5(config-if)#

Rack1R1#debug ip ospf adj
OSPF adjacency events debugging is on
OSPF: 150.1.5.5 address 155.1.0.5 on Serial0/0 is dead
OSPF: 150.1.5.5 address 155.1.0.5 on Serial0/0 is dead, state DOWN
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.5.5 on Serial0/0 from FULL to DOWN,
Neighbor Down: Dead timer expired
OSPF: Neighbor change Event on interface Serial0/0
OSPF: DR/BDR election on Serial0/0
OSPF: Elect BDR 150.1.1.1
OSPF: Elect DR 150.1.1.1
OSPF: Elect BDR 0.0.0.0
OSPF: Elect DR 150.1.1.1
DR: 150.1.1.1 (Id) BDR: none

Rack1R4#debug ip ospf adj
OSPF adjacency events debugging is on
OSPF: 150.1.5.5 address 155.1.0.5 on Serial0/0 is dead
OSPF: 150.1.5.5 address 155.1.0.5 on Serial0/0 is dead, state DOWN
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.5.5 on Serial0/0 from FULL to DOWN,
Neighbor Down: Dead timer expired
Rack1R4#
OSPF: Neighbor change Event on interface Serial0/0
OSPF: DR/BDR election on Serial0/0
OSPF: Elect BDR 223.255.255.255
OSPF: Elect DR 223.255.255.255
OSPF: Elect BDR 0.0.0.0
OSPF: Elect DR 223.255.255.255
DR: 223.255.255.255 (Id) BDR: none
```

The problem in this design occurs when R5 tries to come back onto the segment.

```
Rack1R5#debug ip ospf adj
OSPF adjacency events debugging is on
Rack1R5#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R5(config)#interface Serial0/0
Rack1R5(config-if)#no shutdown
Rack1R5(config-if)#
```

R5 sends unicast updates to the neighbors configured under the OSPF process.

```
OSPF: Starting 0.0.0.0 address 155.1.0.3 on Serial0/0
OSPF: Starting 0.0.0.0 address 155.1.0.1 on Serial0/0
OSPF: Starting 0.0.0.0 address 155.1.0.2 on Serial0/0
OSPF: Starting 0.0.0.0 address 155.1.0.4 on Serial0/0
<output omitted>
OSPF: 2 Way Communication to 150.1.3.3 on Serial0/0, state 2WAY
OSPF: 2 Way Communication to 150.1.1.1 on Serial0/0, state 2WAY
OSPF: 2 Way Communication to 150.1.2.2 on Serial0/0, state 2WAY
OSPF: 2 Way Communication to 223.255.255.255 on Serial0/0, state 2WAY
```

R5 finds four neighbors, and sees that R4 has the highest router-id of 223.255.255.255, and is the DR. The next highest router-id is 150.1.5.5, R5, and it elects itself as the BDR. R1, R2, and R3 revert to DROTHERS.

```
OSPF: Neighbor change Event on interface Serial0/0
OSPF: DR/BDR election on Serial0/0
OSPF: Elect BDR 0.0.0.0
OSPF: Elect DR 223.255.255.255
OSPF: Elect BDR 150.1.5.5
OSPF: Elect DR 223.255.255.255
DR: 223.255.255.255 (Id) BDR: 150.1.5.5 (Id)
<output omitted>
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.3.3 on Serial0/0 from LOADING to FULL,
Loading Done
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.1.1 on Serial0/0 from LOADING to FULL,
Loading Done
%OSPF-5-ADJCHG: Process 1, Nbr 223.255.255.255 on Serial0/0 from LOADING to
FULL, Loading Done
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.2.2 on Serial0/0 from LOADING to FULL,
Loading Done
<output omitted>

Rack1R5#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address      Interface
150.1.3.3        1     FULL/DROTHER    00:01:39   155.1.0.3    Serial0/0
150.1.1.1        1     FULL/DROTHER    00:01:44   155.1.0.1    Serial0/0
150.1.2.2        1     FULL/DROTHER    00:01:55   155.1.0.2    Serial0/0
223.255.255.255  1     FULL/DR         00:01:41   155.1.0.4    Serial0/0
150.1.8.8        1     FULL/DR         00:00:36   155.1.58.8   FastEthernet0/0
```

Due to the desired function of the DR, the network design is now broken. For example in order for R2 to advertise the network 192.10.1.0/24, this update must be sent to the DR before being sent to the rest of the neighbors on the segment. However R2 does not have a direct PVC on the Frame Relay network connecting to R4 who is the DR. The result is that LSA replication inside area 0 is incomplete, and different routers end up with different views of the topology. Based on this inconsistency of information SPF is not uniform, and some routers have reachability to some segments, with others do not. The result of this can be seen in either the OSPF database or the routing tables of the devices.

R1 thinks that R5 is the DR, but R5 thinks that R4 is the DR. LSAs that R1 directs to R5 are dropped, and LSAs from R4 that should be replicated back to R1 cannot be. The result is that R1 has no routes installed in the routing table.

```
Rack1R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:47	155.1.0.5	Serial0/0
150.1.6.6	255	FULL/BDR	00:00:37	155.1.146.6	FastEthernet0/0
223.255.255.255	0	FULL/DROTHER	00:00:39	155.1.146.4	FastEthernet0/0

```
Rack1R1#show ip route ospf
```

R2 has the same result as R1.

```
Rack1R2#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:40	155.1.0.5	Serial0/0
192.10.1.254	1	FULL/DR	00:00:37	192.10.1.254	FastEthernet0/0

```
Rack1R2#show ip route ospf
```

```
51.0.0.0/32 is subnetted, 1 subnets
O E2 51.51.51.51 [110/20] via 192.10.1.254, 00:01:13, FastEthernet0/0
```

R3 knows about area 2 routes, but cannot install Inter-Area routes that should come from other area 0 routers.

```
Rack1R3#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	1	FULL/DR	00:01:31	155.1.0.5	Serial1/0
150.1.7.7	1	FULL/BDR	00:00:33	155.1.37.7	FastEthernet0/0

```
Rack1R3#show ip route ospf
```

```
155.1.0.0/24 is subnetted, 8 subnets
O 155.1.9.0 [110/3] via 155.1.37.7, 04:48:25, FastEthernet0/0
O 155.1.7.0 [110/2] via 155.1.37.7, 04:48:25, FastEthernet0/0
O 155.1.79.0 [110/2] via 155.1.37.7, 04:48:25, FastEthernet0/0
O 155.1.67.0 [110/2] via 155.1.37.7, 04:48:25, FastEthernet0/0
150.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
O 150.1.7.7/32 [110/2] via 155.1.37.7, 04:48:25, FastEthernet0/0
```



The process continues as such depending on where the OSPF topology is viewed from. Therefore at this point in the configuration the network design works in some cases, but not all cases. In order to ensure that the network works in *all* cases, we must guarantee that R5 is always elected the DR. This is due to the fact that R5 is the only neighbor who can form a direct adjacency with the other devices on the Frame Relay segment, because the way the layer 2 circuits are provisioned.

The logic of the solution in this case is somewhat backwards due to how OSPF deals with preemption. Even if R5 is configured with a priority value of 255, it cannot preempt whichever router elected itself as the DR once R5 went down. The only thing R5 can do is re-elect itself as the BDR. Therefore instead of ensuring that R5 *is* elected the DR, we need to ensure that R1, R2, R3, and R4 *are not* elected the DR. This is accomplished by configuring the OSPF priority value as 0 at the interface level. Since priority 0 means they will not participate in the election, R5 is elected the DR and no one is elected the BDR. If R5 goes down Inter-Area routing is lost, however when R5 comes back it is fully restored.

Verification of this can be seen from the `show ip ospf neighbor` output on R5. With the remote devices configured with priority 0 R5 is elected the DR, while all other devices revert back to DROTHERs.

```
Rack1R5#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.3.3	0	FULL/DROTHER	00:01:56	155.1.0.3	Serial0/0
150.1.1.1	0	FULL/DROTHER	00:01:41	155.1.0.1	Serial0/0
150.1.2.2	0	FULL/DROTHER	00:01:51	155.1.0.2	Serial0/0
223.255.255.255	0	FULL/DROTHER	00:01:38	155.1.0.4	Serial0/0
150.1.8.8	1	FULL/DR	00:00:33	155.1.58.8	FastEthernet0/0