

***Copyright Information***

---

Copyright © 2009 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.

# Table of Contents

OSPF .....	1
6.1 OSPF over Broadcast Media .....	1
6.2 OSPF over Non-Broadcast Media .....	2
6.3 OSPF DR/BDR Election Manipulation .....	2
6.4 OSPF Network Point-to-Point .....	2
6.5 OSPF Network Point-to-Multipoint .....	2
6.6 OSPF Network Point-to-Multipoint Non-Broadcast .....	2
6.7 OSPF Network Loopback .....	3
6.8 OSPF Path Selection with Auto-Cost .....	3
6.9 OSPF Path Selection with Cost .....	3
6.10 OSPF Path Selection with Bandwidth .....	3
6.11 OSPF Path Selection with Per-Neighbor Cost .....	3
6.12 Repairing Discontiguous OSPF Areas with Virtual-Links .....	3
6.13 OSPF Path Selection with Non-Backbone Transit Areas .....	4
6.14 OSPF Path Selection with Virtual-Links .....	4
6.15 OSPF Demand Circuit .....	4
6.16 OSPF Flooding Reduction .....	4
6.17 OSPF Clear Text Authentication .....	4
6.18 OSPF MD5 Authentication .....	4
6.19 OSPF Null Authentication .....	5
6.20 OSPF MD5 Authentication with Multiple Keys .....	5
6.21 OSPF Internal Summarization .....	5
6.22 OSPF Path Selection with Summarization .....	5
6.23 OSPF External Summarization .....	5
6.24 OSPF Stub Areas .....	6
6.25 OSPF Totally Stubby Areas .....	6
6.26 OSPF Not-So-Stubby Areas .....	6
6.27 OSPF Not-So-Stubby Areas and Default Routing .....	6
6.28 OSPF Not-So-Totally-Stubby Areas .....	6
6.29 OSPF Stub Areas with Multiple Exit Points .....	7
6.30 OSPF NSSA Type-7 to Type-5 Translator Election .....	7
6.31 OSPF NSSA Redistribution Filtering .....	7
6.32 OSPF LSA Type-3 Filtering .....	7
6.33 OSPF Forwarding Address Suppression .....	7
6.34 OSPF Default Routing .....	8
6.35 OSPF Conditional Default Routing .....	8
6.36 OSPF Reliable Conditional Default Routing .....	8
6.37 OSPF Filtering with Distribute-Lists .....	9
6.38 OSPF Summarization and Discard Routes .....	9
6.39 OSPF Filtering with Administrative Distance .....	9
6.40 OSPF Filtering with Route-Maps .....	9
6.41 OSPF NSSA ABR External Prefix Filtering .....	9
6.42 OSPF Database Filtering .....	9

6.43	OSPF Stub Router Advertisement.....	9
6.44	OSPF Interface Timers.....	10
6.45	OSPF Global Timers .....	10
6.46	OSPF Resource Limiting .....	11
6.47	Miscellaneous OSPF Features.....	11
OSPF Solutions.....		13
6.1	OSPF over Broadcast Media.....	13
6.2	OSPF over Non-Broadcast Media .....	22
6.3	OSPF DR/BDR Election Manipulation .....	32
6.4	OSPF Network Point-to-Point.....	44
6.5	OSPF Network Point-to-Multipoint.....	47
6.6	OSPF Network Point-to-Multipoint Non-Broadcast.....	55
6.7	OSPF Network Loopback .....	59
6.8	OSPF Path Selection with Auto-Cost .....	62
6.9	OSPF Path Selection with Cost.....	65
6.10	OSPF Path Selection with Bandwidth.....	75
6.11	OSPF Path Selection with Per-Neighbor Cost.....	81
6.12	Repairing Discontiguous OSPF Areas with Virtual-Links.....	84
6.13	OSPF Path Selection with Non-Backbone Transit Areas .....	96
6.14	OSPF Path Selection with Virtual-Links.....	106
6.15	OSPF Demand Circuit.....	112
6.16	OSPF Flooding Reduction.....	114
6.17	OSPF Clear Text Authentication .....	116
6.18	OSPF MD5 Authentication .....	119
6.19	OSPF Null Authentication.....	123
6.20	OSPF MD5 Authentication with Multiple Keys.....	124
6.21	OSPF Internal Summarization .....	126
6.22	OSPF Path Selection with Summarization .....	128
6.23	OSPF External Summarization.....	132
6.24	OSPF Stub Areas.....	143
6.25	OSPF Totally Stubby Areas.....	150
6.26	OSPF Not-So-Stubby Areas.....	153
6.27	OSPF Not-So-Stubby Areas and Default Routing .....	159
6.28	OSPF Not-So-Totally-Stubby Areas .....	160
6.29	OSPF Stub Areas with Multiple Exit Points .....	162
6.30	OSPF NSSA Type-7 to Type-5 Translator Election .....	167
6.31	OSPF NSSA Redistribution Filtering .....	175
6.32	OSPF LSA Type-3 Filtering .....	179
6.33	OSPF Forwarding Address Suppression.....	183
6.34	OSPF Default Routing.....	190
6.35	OSPF Conditional Default Routing .....	193
6.36	OSPF Reliable Conditional Default Routing .....	194
6.37	OSPF Filtering with Distribute-Lists .....	198
6.38	OSPF Summarization and Discard Routes .....	200
6.39	OSPF Filtering with Administrative Distance .....	203

---

6.40	OSPF Filtering with Route-Maps .....	205
6.41	OSPF NSSA ABR External Prefix Filtering.....	207
6.42	OSPF Database Filtering.....	209
6.43	OSPF Stub Router Advertisement.....	211
6.44	OSPF Interface Timers.....	214
6.45	OSPF Global Timers .....	216
6.46	OSPF Resource Limiting .....	218
6.47	Miscellaneous OSPF Features.....	219



# 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.



## 6.7 OSPF Network Loopback

- Advertise the Loopback0 interfaces of R1, R2, R3, and R4 into OSPF area 0.
- Advertise the Loopback0 interface of R6 into OSPF area 1.
- Advertise the Loopback0 interfaces of SW1, and SW3 into OSPF area 2.
- Advertise the Loopback0 interfaces of R5, SW2, and SW4 into OSPF area 3.
- Modify the network type of these interfaces so that the links are not advertised as host routes.

## 6.8 OSPF Path Selection with Auto-Cost

- Modify the global OSPF cost calculation of all devices so that a TenGigabitEthernet interface has a cost of 3, and an OC-3 link has a cost of 193.

## 6.9 OSPF Path Selection with Cost

- Using the interface level `ip ospf cost` command modify the OSPF domain so that traffic from SW4 going to the Loopback0 network of R6 uses the point-to-point Serial link between R4 and R5.

## 6.10 OSPF Path Selection with Bandwidth

- Using the interface level `bandwidth` keyword modify the OSPF domain so that traffic from R6 to the Loopback0 network of SW2 is first sent to R1.

## 6.11 OSPF Path Selection with Per-Neighbor Cost

- The Frame Relay PVC between R1 & R5 has been provisioned with 512Kbps of bandwidth by the service provider, while the PVC between R4 & R5 has been provisioned at 768Kbps. Configure per-neighbor cost values on R5 to reflect the underlying bandwidth of the Frame Relay circuits connecting it to R1 and R4.

## 6.12 Repairing Discontiguous OSPF Areas with Virtual-Links

- Configure the network so that full reachability is maintained if R3's connection to SW1 goes down.
- Do not apply any configuration onto R4 to solve this task.

### 6.13 OSPF Path Selection with Non-Backbone Transit Areas

- Configure R4's Frame Relay and point-to-point Serial links connecting to R5 to have an OSPF cost of 1.
- Modify the SPF calculation in the OSPF domain so that R6 cannot use R4 to reach area 0 by transiting area 1.
- Verify this by ensuring that traffic from R6 going to the Loopback0 network of SW2 is still first sent to R1, and then directly to R5.

### 6.14 OSPF Path Selection with Virtual-Links

- Configure the OSPF domain so that traffic from SW3 going to OSPF area 51 transits the point-to-point Serial link between R3 and R2.

### 6.15 OSPF Demand Circuit

- The point-to-point between R4 and R5 is a low-speed Serial link. Configure the OSPF demand circuit feature on this link to reduce periodic OSPF hello transmission and paranoid update flooding.

### 6.16 OSPF Flooding Reduction

- Configure R5, SW2, and SW4 so that links in OSPF area 3 do not participate in periodic paranoid update LSA flooding.

### 6.17 OSPF Clear Text Authentication

- Configure clear-text OSPF authentication for all adjacencies in area 2 using the password CLEARKEY.
- SW1 should enable authentication on all interfaces in area 2 with one single command.
- R3, R6, and SW3 should only enable authentication on their links connecting to SW1.

### 6.18 OSPF MD5 Authentication

- Configure MD5 based OSPF authentication for all adjacencies in area 0, including the Virtual-Links, using the password MD5KEY.
- R1 should enable MD5 authentication on all interfaces in area 0 with a single command.
- All other devices in area 0 should enable MD5 authentication on a per interface basis.

### 6.19 OSPF Null Authentication

- Modify SW1 so that if a new router is added to area 2 on VLAN 7 it does not require authentication.

### 6.20 OSPF MD5 Authentication with Multiple Keys

- Enable MD5 authentication in OSPF area 1.
- R1 should authenticate R6 using the password R1R6KEY.
- R4 should authenticate R6 using the password R4R6KEY.

### 6.21 OSPF Internal Summarization

- Configure R5 to advertise a summary route for the VLAN 8 and 10 prefixes as they are sent into area 0.
- This summary should be as specific as possible while still encompassing all addresses in both subnets.

### 6.22 OSPF Path Selection with Summarization

- Configure R3 to originate the summary route 155.X.6.0/23 in order to force traffic destined to VLAN 7 to transit R6.
- If R6's connection to VLAN 146 is down, traffic for VLAN 7 should transit R3.

### 6.23 OSPF External Summarization

- Redistribute between RIPv2 and OSPF on R4.
- Redistribute between EIGRP and OSPF on R6.
- R4 should advertise a single summary route into OSPF for the prefixes learned from BB3 that has a cost of 50.
- R6 should advertise a single summary route into OSPF for the prefixes learned from BB1 that has a cost of 100, and includes the cost needed to reach R6.

 **Note**

Revert all devices to the *Full OSPF Routing* initial configurations before continuing.

## 6.24 OSPF Stub Areas

- Configure OSPF area 3 so that R5 filters External routes out as they are sent from area 0 to area 3.
- Devices in area 3 should still have reachability to routes external to the OSPF domain.

## 6.25 OSPF Totally Stubby Areas

- Configure OSPF area 3 so that R5 filters Inter-Area and External routes as they are sent from area 0 to area 3.
- Devices in area 3 should still have reachability to routes external to the OSPF domain.

## 6.26 OSPF Not-So-Stubby Areas

- Shutdown R6's Loopback0 interface and connection to VLAN 146.
- Configure OSPF area 2 so that R3 filters External routes out as they are sent from area 0 to area 2.
- R6 should still be allowed to redistribute EIGRP into OSPF.
- Note any reachability problems throughout the domain.

## 6.27 OSPF Not-So-Stubby Areas and Default Routing

- Configure R3 to advertise a default route into area 2.
- This default route should have a cost of 500.

## 6.28 OSPF Not-So-Totally-Stubby Areas

- Remove the previous default routing configuration on R3.
- Modify the area 2 configuration so that R3 filters Inter-Area and External routes as they are sent from area 0 to area 2.
- R6 should still be allowed to redistribute EIGRP into OSPF.

## 6.29 OSPF Stub Areas with Multiple Exit Points

- Re-enable R6's Loopback0 interface and connection to VLAN 146.
- Modify area 2 in such a way that traffic from devices in area 2 going to area 3 uses R6 as the exit point, while traffic from area 2 going to the RIP routes learned from BB3 uses R3.
- If R3 loses its connection to area 2, traffic for External destinations should be rerouted to R6.
- If R6 loses its connection to area 2, traffic for Inter-Area destinations should be rerouted to R3.
- Do not modify the cost of any links in area 2 to accomplish this.

## 6.30 OSPF NSSA Type-7 to Type-5 Translator Election

- Create a new Loopback9 interface on SW3 with the IP address 9.9.9.9/32, and redistribute it into OSPF.
- Modify the OSPF domain so that only R3 advertises this route into area 0, but traffic from R5 going to this destination transits R6.
- If R6 is down this traffic should be rerouted to R3.

## 6.31 OSPF NSSA Redistribution Filtering

- Create new Loopback interfaces on R5 and SW4 with IP addresses 5.5.5.5/32 and 10.10.10.10/32 respectively, and redistribute them into the OSPF domain.
- Configure area 3 as an NSSA so R5 blocks all LSA types 3, 4, and 5 and replaces them with a default route.
- Modify area 3 so that R5's redistributed Loopback is injected into area 0 as LSA Type-5, but is not injected into area 3 as LSA Type-7.

## 6.32 OSPF LSA Type-3 Filtering

- Configure LSA Type-3 Filtering on R5 so that devices in area 0 do not have reachability information about the layer 3 EtherChannel link between SW2 and SW4 or SW4's Loopback0 interface.
- Configure LSA Type-3 Filtering on R2 so that devices in area 51 do not have reachability information about the Loopback0 interface of R5; this should not affect devices in area 5.

## 6.33 OSPF Forwarding Address Suppression

- Modify R5's NSSA configuration so that devices outside of area 3 maintain connectivity to SW4's redistributed Loopback 10.10.10.10/32.

 **Note**

Revert all devices to the *Full OSPF Routing* initial configurations before continuing.

### 6.34 OSPF Default Routing

- Configure R6 with a static default route pointing to BB1.
- Configure R6 to originate an external type-2 default route with a metric of 60 into OSPF as long as the static default route is installed in the routing table.
- Do not use a route-map to accomplish this.
- Configure R4 to originate an external type-1 default route with a metric of 40 regardless of whether it has a default route already installed in the routing table.

### 6.35 OSPF Conditional Default Routing

- Remove the previous default route advertisement on R6.
- Configure R6 to originate a default route into OSPF, but only if the local Serial interface connecting to BB1 is up.

### 6.36 OSPF Reliable Conditional Default Routing

- Remove the previous default route advertisement on R4.
- Configure an IP SLA instance on R4 to check ICMP reachability to BB3 every five seconds.
- R4 should advertise a default route into OSPF, but only if the SLA monitoring reports its status as OK.

### 6.37 OSPF Filtering with Distribute-Lists

- Configure distribute-list filtering on R5, SW2, and SW4 so that these devices do not install routes to the Loopback0 networks of R1 and R2.

### 6.38 OSPF Summarization and Discard Routes

- Configure R5 to advertise the summary 150.X.0.0/22 into area 3.
- Ensure that R5, SW2, and SW4 can still reach the Loopback0 networks of R1 and R2, even with the previously configured distribute-list filter applied.

### 6.39 OSPF Filtering with Administrative Distance

- Configure administrative distance filtering on R5 so that traffic going towards the VLAN 67 network is sent towards R3.

### 6.40 OSPF Filtering with Route-Maps

- Configure route-map filtering on R6 so that traffic going towards the Loopback0 network of R3 is sent towards R1.

### 6.41 OSPF NSSA ABR External Prefix Filtering

- Disable R6's link to VLAN 146 and Loopback0.
- Configure area 2 as an NSSA.
- Configure summarization on R3 so that devices outside of area 2 do not have a route to the network 200.0.0.0/24.
- This filter should not affect any other prefixes.

### 6.42 OSPF Database Filtering

- Configure SW1 so that SW3 cannot learn any OSPF routes from SW1, but SW1 can still learn OSPF routes from SW3.
- Configure R5 so that R2 cannot learn from R5, but R5 can still learn OSPF routes from R2.

### 6.43 OSPF Stub Router Advertisement

- Configure R4 to advertise the maximum metric value to all neighbors inside its Type-1 Router LSA.

## 6.44 OSPF Interface Timers

- Configure R5 to send OSPF hello packets out its link to the Frame Relay network every five seconds, and wait for seven seconds before declaring a neighbor down.
- Configure R4 to send OSPF hello packets every 250ms on the point-to-point link to R5.

## 6.45 OSPF Global Timers

- Modify R4 and R5's OSPF timers as follows:
  - Configure SPF throttling to start new re-calculation at least 100ms after a new LSA arrives. The second SPF calculation should occur in no less than 1s after the first one, and the maximum wait time should be no more than 10 seconds.
  - Configure LSA pacing to wait at least 50ms between consecutive link-state updates. LSA retransmissions should be paced at least 75ms apart.
  - Configure LSA throttling to generate subsequent LSAs after 10ms, to wait at least 4 seconds to generate the next LSA, and no more than 6 seconds between generation of the same LSA.
  - Configure LSA arrival throttling to wait 2 seconds between reception of the same LSA from a neighbor.
  - of the re-calculation at least 100ms after a new LSA arrives. The second SPF calculation should occur in no less than 1s after the first one, and the maximum wait time should be no more than 10 seconds.
- Configure R4 and R5 to assume that LSA transmission takes 2 seconds on the point-to-point link between them. LSA retransmission should occur if an acknowledgement is not received within 10 seconds over this link.



## 6.46 OSPF Resource Limiting

- Configure R4 so that no more than 5000 LSAs can exist in the database.
- No more than 500 of these routes should be originated through redistribution.
- R4's OSPF process should be allowed to use up to 20% CPU utilization before allowing interrupts from higher priority processes.

## 6.47 Miscellaneous OSPF Features

- Configure SW4 so that it does not generate a log message upon receipt of a Type-6 LSA advertisement.
- Configure SW4 so that it does not take the MTU value into account when establishing adjacencies on its Port-Channel1 interface.
- Configure SW4 to reflect the following output:

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
Rack1SW2	1	FULL/BDR	00:00:38	155.1.108.8	Port-channel1



# 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:
ip routing
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 2

SW2:
ip routing
!
router ospf 1
 network 155.1.0.0 0.0.127.255 area 3

SW3:
ip routing
!
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:
ip routing
!
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 Serial0/0
 ip ospf 1 area 0
 !
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

Once the design is complete it can also be verified by the `show ip ospf database` output. Only the DR on an OSPF segment originates the Network LSA (LSA 2). The Network LSA is used to describe all the neighbors that the DR on the segment is adjacent with to the BDR and DROTHERs on that link.

From SW3's below output the *Link ID* under *Net Link States (Area 2)* is the interface IP address of the DR, and the *ADV Router* is the router-id of the DR. Based on this we can see that R3 is the DR for the link from R3 to SW1, R6 is the DR for VLAN 67, and SW3 is the DR for VLAN 79. SW3 does not see Network LSAs for other areas, as they are only needed for Intra-Area SPF calculations.

```
Rack1SW3#show ip ospf database
```

```
OSPF Router with ID (150.1.9.9) (Process ID 1)
```

```
Router Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.3.3	150.1.3.3	660	0x80000051	0x0061BB	1
150.1.6.6	150.1.6.6	1019	0x80000050	0x00D8F7	1
150.1.7.7	150.1.7.7	894	0x80000055	0x0005C0	5
150.1.9.9	150.1.9.9	670	0x8000004F	0x00B13F	2

```
Net Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.37.3	150.1.3.3	660	0x8000004E	0x0082C4
155.1.67.6	150.1.6.6	1019	0x8000004E	0x002BEE
155.1.79.9	150.1.9.9	671	0x8000004E	0x009A64

```
<output omitted>
```



R1 participates in both area 0 and area 1 SPF, so it sees the Network LSA coming from R5 for the Frame Relay segment, and R6 for VLAN 146. This means that R5 and R6 are the DRs for those segments respectively.

```
Rack1R1#show ip ospf database
```

```
OSPF Router with ID (150.1.1.1) (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	212	0x80000030	0x002130	1
150.1.2.2	150.1.2.2	219	0x80000032	0x000941	1
150.1.3.3	150.1.3.3	213	0x80000032	0x00DBEA	2
150.1.5.5	150.1.5.5	218	0x80000040	0x009F76	3
223.255.255.255	223.255.255.255	219	0x80000037	0x00568C	3

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

Link ID	ADV Router	Age	Seq#	Checksum
155.1.0.5	150.1.5.5	213	0x80000002	0x00DEB6

```
<output omitted>
```

```
Net Link States (Area 1)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.146.6	150.1.6.6	294	0x80000001	0x00E05F

```
<output omitted>
```

SW4 only sees Network LSAs for area 3 because it does not participate in other areas. The below output indicates that SW2 is the DR for the VLAN 58 segment and SW4 is the DR for the layer 3 EtherChannel connecting to SW2.

```
Rack1SW4#show ip ospf database
```

```
OSPF Router with ID (223.255.255.255) (Process ID 1)
```

```
Router Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.5.5	150.1.5.5	1856	0x80000052	0x00F13E	2
150.1.8.8	150.1.8.8	1660	0x80000056	0x0033B3	3
223.255.255.255	223.255.255.255	1506	0x80000051	0x0056EE	2

```
Net Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.58.8	150.1.8.8	1660	0x8000004E	0x0050CC
155.1.108.10	223.255.255.255	1506	0x8000004E	0x00A2CF

```
<output omitted>
```

## 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.

### **Configuration**

---

```
R1:
interface Serial0/1
 ip ospf 1 area 4
```

```
R2:
interface Serial0/1
 ip ospf 1 area 5
```

```
R3:
interface Serial1/2
 ip ospf 1 area 4
!
interface Serial1/3
 ip ospf 1 area 5
```

```
R4:
interface Serial0/1
 ip ospf 1 area 0
```

```
R5:
interface Serial0/1
 ip ospf 1 area 0
```

## Verification

### Note

OSPF network point-to-point is the default option for point-to-point interfaces such as HDLC, PPP, or point-to-point NBMA subinterfaces. It uses multicast hellos, does not use the DR/BDR election, and only supports the adjacency of exactly two neighbors on a segment. No special design considerations need to be taken into account for point-to-point OSPF interfaces.

```
Rack1R4#show ip ospf interface Serial0/1
Serial0/1 is up, line protocol is up
  Internet Address 155.1.45.4/24, Area 0
  Process ID 1, Router ID 223.255.255.255, Network Type POINT_TO_POINT, Cost: 64
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:02
  Supports Link-local Signaling (LLS)
  Index 2/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.5.5
  Suppress hello for 0 neighbor(s)
```

```
Rack1R5#show ip ospf interface Serial0/1
Serial0/1 is up, line protocol is up
  Internet Address 155.1.45.5/24, Area 0
  Process ID 1, Router ID 150.1.5.5, Network Type POINT_TO_POINT, Cost: 64
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:09
  Supports Link-local Signaling (LLS)
  Index 2/4, 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 1, Adjacent neighbor count is 1
    Adjacent with neighbor 223.255.255.255
  Suppress hello for 0 neighbor(s)
```

The null output after the slash in the *State* column in the `show ip ospf neighbor` output indicates that no DR or BDR is elected for network type point-to-point.

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	0	FULL/ -	00:00:37	155.1.45.5	Serial0/1
150.1.5.5	1	FULL/DR	00:01:50	155.1.0.5	Serial0/0
150.1.1.1	1	FULL/DR	00:00:32	155.1.146.1	FastEthernet0/1
150.1.6.6	255	FULL/BDR	00:00:30	155.1.146.6	FastEthernet0/1

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
223.255.255.255	0	FULL/ -	00:00:38	155.1.45.4	Serial0/1
150.1.3.3	0	FULL/DROTHER	00:01:35	155.1.0.3	Serial0/0
150.1.1.1	0	FULL/DROTHER	00:01:50	155.1.0.1	Serial0/0
150.1.2.2	0	FULL/DROTHER	00:01:30	155.1.0.2	Serial0/0
223.255.255.255	0	FULL/DROTHER	00:01:47	155.1.0.4	Serial0/0
150.1.8.8	1	FULL/DR	00:00:31	155.1.58.8	FastEthernet0/0

## 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.

### Configuration

---

```
R1:
interface Serial0/0
 ip ospf network point-to-multipoint
 no frame-relay map ip 155.1.0.2 105
 no frame-relay map ip 155.1.0.3 105
 no frame-relay map ip 155.1.0.4 105
 frame-relay map ip 155.1.0.5 105 broadcast
```

```
R2:
interface Serial0/0
 ip ospf network point-to-multipoint
 no frame-relay map ip 155.1.0.1 205
 no frame-relay map ip 155.1.0.3 205
 no frame-relay map ip 155.1.0.4 205
 frame-relay map ip 155.1.0.5 205 broadcast
```

```
R3:
interface Serial1/0
 ip ospf network point-to-multipoint
 no frame-relay map ip 155.1.0.1 305
 no frame-relay map ip 155.1.0.2 305
 no frame-relay map ip 155.1.0.4 305
 frame-relay map ip 155.1.0.5 305 broadcast
```

```
R4:
interface Serial0/0
 ip ospf network point-to-multipoint
 no frame-relay map ip 155.1.0.1 405
 no frame-relay map ip 155.1.0.2 405
 no frame-relay map ip 155.1.0.3 405
 frame-relay map ip 155.1.0.5 405 broadcast
```

```
R5:
interface Serial0/0
 ip ospf network point-to-multipoint
 frame-relay map ip 155.1.0.1 501 broadcast
 frame-relay map ip 155.1.0.2 502 broadcast
 frame-relay map ip 155.1.0.3 503 broadcast
 frame-relay map ip 155.1.0.4 504 broadcast
```

## Verification

### Note

OSPF network type point-to-multipoint is specifically designed to solve reachability problems in partially meshed NBMA network designs. Like network type point-to-point, it sends hellos as multicasts and does not support the DR/BDR election. Unlike point-to-point however, multiple adjacencies on a single interface are supported.

In this particular scenario the first design issue that must be solved is allowing the transport of multicast OSPF packets over the Frame Relay circuits. Since the initial configurations exclude the broadcast keyword from the **frame-relay map** statements configured, the **debug ip packet** and **debug frame-relay packet** outputs indicate that the encapsulation is failing when the OSPF multicast 224.0.0.5 is sent out the Frame Relay circuit. Note that this was not an issue when using network type non-broadcast, because non-broadcast uses unicast hellos, and therefore does not require multicast transport.

```
Rack1R5#debug ip packet
IP packet debugging is on
```

```
Rack1R5#debug frame-relay packet
Frame Relay packet debugging is on
```

```
IP: s=155.1.45.5 (local), d=224.0.0.5 (Serial0/1), len 80, sending
broad/multicast
IP: s=155.1.5.5 (local), d=224.0.0.5 (FastEthernet0/1), len 76, sending
broad/multicast
IP: s=155.1.58.5 (local), d=224.0.0.5 (FastEthernet0/0), len 80, sending
broad/multicast
IP: s=155.1.0.5 (local), d=224.0.0.5 (Serial0/0), len 76, sending
broad/multicast
Serial0/0: broadcast search
Serial0/0:encaps failed on broadcast for link 7(IP)
IP: s=155.1.0.5 (local), d=224.0.0.5 (Serial0/0), len 76, encapsulation failed
```

To resolve this, the **broadcast** keyword must be added on the Frame Relay mappings from R5 to R1, R2, R3, & R4, and from R1, R2, R3, & R4 back to R5.

```
Rack1R5#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R5(config)#interface Serial0/0
Rack1R5(config-if)# frame-relay map ip 155.1.0.1 501 broadcast
Rack1R5(config-if)# frame-relay map ip 155.1.0.2 502 broadcast
Rack1R5(config-if)# frame-relay map ip 155.1.0.3 503 broadcast
Rack1R5(config-if)# frame-relay map ip 155.1.0.4 504 broadcast
%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 150.1.2.2 on Serial0/0 from LOADING to FULL,
Loading Done
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.4.4 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
```

Once adjacency is established the **show ip ospf neighbor** output indicates that there is no DR or BDR for the segment with the null output in the *State* field.

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.1.1	0	FULL/ -	00:01:59	155.1.0.1	Serial0/0
150.1.3.3	0	FULL/ -	00:01:59	155.1.0.3	Serial0/0
150.1.2.2	0	FULL/ -	00:01:59	155.1.0.2	Serial0/0
223.255.255.255	0	FULL/ -	00:01:59	155.1.0.4	Serial0/0
223.255.255.255	0	FULL/ -	00:00:36	155.1.45.4	Serial0/1
150.1.8.8	1	FULL/DR	00:00:35	155.1.58.8	FastEthernet0/0

The spokes of the network only form adjacency with R5, and not each other, because layer 2 PVCs only exist between the hub and the spokes.

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	0	FULL/ -	00:01:54	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
150.1.3.3	0	FULL/ -	00:00:31	155.1.13.3	Serial0/1

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	0	FULL/ -	00:01:45	155.1.0.5	Serial0/0
150.1.3.3	0	FULL/ -	00:00:36	155.1.23.3	Serial0/1
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
150.1.5.5	0	FULL/ -	00:01:33	155.1.0.5	Serial1/0
150.1.7.7	1	FULL/BDR	00:00:32	155.1.37.7	FastEthernet0/0
150.1.1.1	0	FULL/ -	00:00:34	155.1.13.1	Serial1/2
150.1.2.2	0	FULL/ -	00:00:32	155.1.23.2	Serial1/3

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	0	FULL/ -	00:01:39	155.1.0.5	Serial0/0
150.1.5.5	0	FULL/ -	00:00:33	155.1.45.5	Serial0/1
150.1.1.1	1	FULL/DR	00:00:37	155.1.146.1	FastEthernet0/1
150.1.6.6	255	FULL/BDR	00:00:35	155.1.146.6	FastEthernet0/1

The output from the **show ip ospf database** in area 0 also reinforces the fact that there is no DR or BDR for the segment. Recall from the DR/BDR election section that the Network LSA (LSA 2) is used to describe the DR and its attached neighbors. Since there are no DRs for any transit links in area 0, the *Net Link States (Area 0)* section does not appear in the below output.

```
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	989	0x8000002E	0x004660	2
150.1.2.2	150.1.2.2	1386	0x80000030	0x00445A	2
150.1.3.3	150.1.3.3	1065	0x80000030	0x00DFE8	2
150.1.5.5	150.1.5.5	1119	0x8000003D	0x00EB9B	7
223.255.255.255	223.255.255.255	1522	0x80000033	0x001F18	4

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

Link ID	ADV Router	Age	Seq#	Checksum
150.1.7.7	150.1.3.3	322	0x80000028	0x008F41
155.1.5.0	150.1.5.5	1119	0x80000028	0x00864B
155.1.7.0	150.1.3.3	322	0x80000028	0x00943E
155.1.8.0	150.1.5.5	1119	0x80000028	0x006F5E
155.1.9.0	150.1.3.3	322	0x80000028	0x008847
155.1.10.0	150.1.5.5	1119	0x80000028	0x006367
155.1.13.0	150.1.1.1	1243	0x8000001D	0x00F0AC
155.1.13.0	150.1.3.3	1823	0x8000001D	0x00F1D7
155.1.23.0	150.1.2.2	1876	0x8000001D	0x00751C
155.1.23.0	150.1.3.3	1823	0x8000001D	0x00833C
155.1.37.0	150.1.3.3	323	0x80000028	0x003F76
155.1.58.0	150.1.5.5	1120	0x80000028	0x003D5F
155.1.67.0	150.1.3.3	323	0x80000028	0x00FD98
155.1.79.0	150.1.3.3	323	0x80000028	0x007911
155.1.108.0	150.1.5.5	1120	0x80000028	0x001F4A
155.1.146.0	150.1.1.1	2008	0x80000027	0x00A7A5
155.1.146.0	223.255.255.255	1523	0x80000028	0x0027DE
192.10.1.0	150.1.2.2	130	0x80000028	0x008A23

```
<output omitted>
```



Once LSA replication is complete the next change that should be evident is how the routing table is processed on the Frame Relay segment between the neighbors. For example take the following view of the network from R2's perspective.

```
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, 00:02:33, FastEthernet0/0
 155.1.0.0/16 is variably subnetted, 19 subnets, 2 masks
O IA   155.1.146.0/24 [110/129] via 155.1.0.5, 00:02:33, Serial0/0
O IA   155.1.10.0/24 [110/67] via 155.1.0.5, 00:02:33, Serial0/0
O IA   155.1.8.0/24 [110/66] via 155.1.0.5, 00:02:33, Serial0/0
O IA   155.1.9.0/24 [110/131] via 155.1.0.5, 00:02:33, Serial0/0
O IA   155.1.13.0/24 [110/192] via 155.1.0.5, 00:02:33, Serial0/0
O      155.1.0.3/32 [110/128] via 155.1.0.5, 00:02:33, Serial0/0
O      155.1.0.1/32 [110/128] via 155.1.0.5, 00:02:33, Serial0/0
O IA   155.1.7.0/24 [110/130] via 155.1.0.5, 00:02:33, Serial0/0
O      155.1.0.4/32 [110/128] via 155.1.0.5, 00:02:33, Serial0/0
O IA   155.1.5.0/24 [110/65] via 155.1.0.5, 00:02:33, Serial0/0
O      155.1.0.5/32 [110/64] via 155.1.0.5, 00:02:33, Serial0/0
O IA   155.1.58.0/24 [110/65] via 155.1.0.5, 00:02:33, Serial0/0
O      155.1.45.0/24 [110/128] via 155.1.0.5, 00:02:34, Serial0/0
O IA   155.1.37.0/24 [110/129] via 155.1.0.5, 00:02:34, Serial0/0
O IA   155.1.79.0/24 [110/130] via 155.1.0.5, 00:02:34, Serial0/0
O IA   155.1.67.0/24 [110/130] via 155.1.0.5, 00:02:34, Serial0/0
O IA   155.1.108.0/24 [110/66] via 155.1.0.5, 00:02:34, Serial0/0
 150.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
O IA   150.1.7.7/32 [110/130] via 155.1.0.5, 00:02:34, Serial0/0
```

For all routes that are learned from the hub of the Frame Relay network the next-hop value has been updated to the interface IP address of the hub. Recall from the previous design with the non-broadcast network type that the DR did not update the next-hop. The result of this is that when route recursion is performed for traffic that needs to be routed between the spokes, *layer 3 to layer 2 resolution now occurs for the hub, not the spoke.*

Specifically we can see this on R2 for its lookup to reach the VLAN 9 segment. R2's longest match for 155.1.9.9 is 155.1.9.0/24. The evident change here is that the next-hop value is not the LSA 3 originator, who is R3, but the hub, R5. Note the field *from 150.1.3.3* still tells R2 that R3 is the LSA 3 originator.

```
Rack1R2#show ip route 155.1.9.9
Routing entry for 155.1.9.0/24
  Known via "ospf 1", distance 110, metric 131, type inter area
  Last update from 155.1.0.5 on Serial0/0, 00:02:48 ago
  Routing Descriptor Blocks:
  * 155.1.0.5, from 150.1.3.3, 00:02:48 ago, via Serial0/0
    Route metric is 131, traffic share count is 1
```

R2 now does a recursive lookup on 155.1.0.5, and it points to an OSPF host route for 155.1.0.5/32 that is learned from R5 via Serial0/0.

```
Rack1R2#show ip route 155.1.0.5
Routing entry for 155.1.0.5/32
  Known via "ospf 1", distance 110, metric 64, type intra area
  Last update from 155.1.0.5 on Serial0/0, 00:02:59 ago
  Routing Descriptor Blocks:
  * 155.1.0.5, from 150.1.5.5, 00:02:59 ago, via Serial0/0
    Route metric is 64, traffic share count is 1
```

This implies that R2 must now resolve 155.1.0.5 at the Frame Relay layer. Since there is a static **frame-relay map** statement pointing towards 155.1.0.5, R2 knows to use PVC 205 when the layer 2 frame is built.

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

The end result is end-to-end reachability.

```
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 = 112/116/124 ms
```

The design implication of this change is that the spokes of the NBMA network no longer need to maintain detailed mapping tables for all the other spokes. For example take the Frame Relay interface configuration of R2 below.

```
Rack1R2#show run interface Serial0/0
Building configuration...

Current configuration : 237 bytes
!
interface Serial0/0
 ip address 155.1.0.2 255.255.255.0
 encapsulation frame-relay
 ip ospf network point-to-multipoint
 ip ospf priority 0
 ip ospf 1 area 0
 frame-relay map ip 155.1.0.5 205 broadcast
 no frame-relay inverse-arp
end
```

R2 has a **frame-relay map** statement pointing to R5, but not R1, R3, or R4. When R2 does a layer 3 routing lookup on R1's Frame Relay address, 155.1.0.1, it shows that it is learned from R5 as a host route. In the previous non-broadcast design R2 would have seen this as directly connected. In this case R2 now says that 155.1.0.1 is via 155.1.0.5, and 155.1.0.5 is via Serial0/0.

```
Rack1R2#show ip route 155.1.0.1
Routing entry for 155.1.0.1/32
  Known via "ospf 1", distance 110, metric 128, type intra area
  Last update from 155.1.0.5 on Serial0/0, 00:08:18 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.1.1, 00:08:18 ago, via Serial0/0
      Route metric is 128, traffic share count is 1
```

```
Rack1R2#show ip route 155.1.0.5
Routing entry for 155.1.0.5/32
  Known via "ospf 1", distance 110, metric 64, type intra area
  Last update from 155.1.0.5 on Serial0/0, 00:08:21 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.5.5, 00:08:21 ago, via Serial0/0
      Route metric is 64, traffic share count is 1
```

This implies that as long as R2 has a mapping for 155.1.0.5 in the Frame Relay table, encapsulation is successful.

```
Rack1R2#ping 155.1.0.1
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 155.1.0.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/115/120 ms
```

Likewise R1 says that 155.1.0.2 is via 155.1.0.5, and 155.1.0.5 is via Serial0/0. Since R1 has a mapping for 155.1.0.5 the traffic flow is successful.

```
Rack1R1#show ip route 155.1.0.2
Routing entry for 155.1.0.2/32
  Known via "ospf 1", distance 110, metric 128, type intra area
  Last update from 155.1.0.5 on Serial0/0, 00:11:52 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.2.2, 00:11:52 ago, via Serial0/0
      Route metric is 128, traffic share count is 1
```

```
Rack1R1#show ip route 155.1.0.5
Routing entry for 155.1.0.5/32
  Known via "ospf 1", distance 110, metric 64, type intra area
  Last update from 155.1.0.5 on Serial0/0, 00:11:55 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.5.5, 00:11:55 ago, via Serial0/0
      Route metric is 64, traffic share count is 1
```

```
Rack1R1#ping 155.1.0.2
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 155.1.0.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/116/125 ms
```

From views outside of the connected segment the point-to-multipoint network is only seen as a collection of endpoints, not a transit segment itself. For example take the following output on SW4 who is in area 3.

```
Rack1SW4#show ip route 155.1.0.0 255.255.255.0
% Subnet not in table

Rack1SW4#show ip route | include 155.1.0.
  155.1.0.0/16 is variably subnetted, 19 subnets, 2 masks
O IA   155.1.0.2/32 [110/66] via 155.1.108.8, 00:15:21, Port-channel1
O IA   155.1.0.3/32 [110/66] via 155.1.108.8, 00:15:21, Port-channel1
O IA   155.1.0.1/32 [110/66] via 155.1.108.8, 00:15:21, Port-channel1
O IA   155.1.0.4/32 [110/66] via 155.1.108.8, 00:15:31, Port-channel1
O IA   155.1.0.5/32 [110/2]  via 155.1.108.8, 00:15:31, Port-channel1
```

SW4 does not have an exact match for the prefix 155.1.0.0/24, which is the actual subnet assignment of the Frame Relay link. Instead it knows that there are five endpoints on the network, 155.1.0.1, 155.1.0.2, 155.1.0.3, 155.1.0.4, and 155.1.0.5. This is the normal and desirable behavior for OSPF network type point-to-multipoint per the RFC standard.

## 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.

### **Configuration**

---

```
R1:
interface Serial0/0
 ip ospf network point-to-multipoint non-broadcast
 no frame-relay map ip 155.1.0.5 105 broadcast
 frame-relay map ip 155.1.0.5 105

R2:
interface Serial0/0
 ip ospf network point-to-multipoint non-broadcast
 no frame-relay map ip 155.1.0.5 205 broadcast
 frame-relay map ip 155.1.0.5 205

R3:
interface Serial1/0
 ip ospf network point-to-multipoint non-broadcast
 no frame-relay map ip 155.1.0.5 305 broadcast
 frame-relay map ip 155.1.0.5 305

R4:
interface Serial0/0
 ip ospf network point-to-multipoint non-broadcast
 no frame-relay map ip 155.1.0.5 405 broadcast
 frame-relay map ip 155.1.0.5 405

R5:
interface Serial0/0
 ip ospf network point-to-multipoint non-broadcast
 no frame-relay map ip 155.1.0.1 501 broadcast
 no frame-relay map ip 155.1.0.2 502 broadcast
 no frame-relay map ip 155.1.0.3 503 broadcast
 no frame-relay map ip 155.1.0.4 504 broadcast
 frame-relay map ip 155.1.0.1 501
 frame-relay map ip 155.1.0.2 502
 frame-relay map ip 155.1.0.3 503
 frame-relay map ip 155.1.0.4 504
!
router ospf 1
 neighbor 155.1.0.1
 neighbor 155.1.0.2
 neighbor 155.1.0.3
 neighbor 155.1.0.4
```

**Verification** **Note**

OSPF network type point-to-multipoint non-broadcast is essentially the same as network type point-to-multipoint, with one exception. Point-to-multipoint network type uses multicast hellos, while point-to-multipoint non-broadcast uses unicast hellos. Both do not support the DR/BDR election, automatically update the next-hop value of routes learned on partially meshed networks to the directly connected neighbor, and advertise the network as a set of endpoints instead of a transit network.

The `show ip ospf neighbor` output is identical between the two. In this case we can see that the spokes are adjacent with R5, the hub, and the hub is adjacent with the spokes. The null field under State indicates that no DR/BDR election has occurred.

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	0	FULL/ -	00:01:50	155.1.0.5	Serial0/0
150.1.6.6	255	FULL/DR	00:00:37	155.1.146.6	FastEthernet0/0
223.255.255.255	0	FULL/DROTHER	00:00:30	155.1.146.4	FastEthernet0/0
150.1.3.3	0	FULL/ -	00:00:37	155.1.13.3	Serial0/1

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	0	FULL/ -	00:01:44	155.1.0.5	Serial0/0
150.1.3.3	0	FULL/ -	00:00:35	155.1.23.3	Serial0/1
192.10.1.254	1	FULL/DR	00:00:39	192.10.1.254	FastEthernet0/0

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	0	FULL/ -	00:01:34	155.1.0.5	Serial1/0
150.1.7.7	1	FULL/BDR	00:00:39	155.1.37.7	FastEthernet0/0
150.1.1.1	0	FULL/ -	00:00:38	155.1.13.1	Serial1/2
150.1.2.2	0	FULL/ -	00:00:31	155.1.23.2	Serial1/3

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	0	FULL/ -	00:01:55	155.1.0.5	Serial0/0
150.1.5.5	0	FULL/ -	00:00:36	155.1.45.5	Serial0/1
150.1.1.1	1	FULL/BDR	00:00:33	155.1.146.1	FastEthernet0/1
150.1.6.6	255	FULL/DR	00:00:36	155.1.146.6	FastEthernet0/1

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.1.1	0	FULL/ -	00:01:48	155.1.0.1	Serial0/0
150.1.3.3	0	FULL/ -	00:01:48	155.1.0.3	Serial0/0
150.1.2.2	0	FULL/ -	00:01:48	155.1.0.2	Serial0/0
223.255.255.255	0	FULL/ -	00:01:48	155.1.0.4	Serial0/0
223.255.255.255	0	FULL/ -	00:00:32	155.1.45.4	Serial0/1
150.1.8.8	1	FULL/DR	00:00:30	155.1.58.8	FastEthernet0/0

The routing table processing is the same between both these types. For example R2 sees the network 155.1.7.0/24 as reachable through R5, not through the originator, R3. This means that R2 does not need a **frame-relay map** statement to R3, it just needs one towards R5.

```
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, 00:06:24, FastEthernet0/0
 155.1.0.0/16 is variably subnetted, 19 subnets, 2 masks
O IA   155.1.146.0/24 [110/129] via 155.1.0.5, 00:06:34, Serial0/0
O IA   155.1.10.0/24 [110/67] via 155.1.0.5, 00:06:34, Serial0/0
O IA   155.1.8.0/24 [110/66] via 155.1.0.5, 00:06:34, Serial0/0
O IA   155.1.9.0/24 [110/131] via 155.1.0.5, 00:06:34, Serial0/0
O IA   155.1.13.0/24 [110/192] via 155.1.0.5, 00:06:34, Serial0/0
O      155.1.0.3/32 [110/128] via 155.1.0.5, 00:06:34, Serial0/0
O      155.1.0.1/32 [110/128] via 155.1.0.5, 00:06:34, Serial0/0
O IA   155.1.7.0/24 [110/130] via 155.1.0.5, 00:06:34, Serial0/0
O      155.1.0.4/32 [110/128] via 155.1.0.5, 00:06:34, Serial0/0
O IA   155.1.5.0/24 [110/65] via 155.1.0.5, 00:06:35, Serial0/0
O      155.1.0.5/32 [110/64] via 155.1.0.5, 00:06:35, Serial0/0
O IA   155.1.58.0/24 [110/65] via 155.1.0.5, 00:06:35, Serial0/0
O      155.1.45.0/24 [110/128] via 155.1.0.5, 00:06:35, Serial0/0
O IA   155.1.37.0/24 [110/129] via 155.1.0.5, 00:06:35, Serial0/0
O IA   155.1.79.0/24 [110/130] via 155.1.0.5, 00:06:35, Serial0/0
O IA   155.1.67.0/24 [110/130] via 155.1.0.5, 00:06:35, Serial0/0
O IA   155.1.108.0/24 [110/66] via 155.1.0.5, 00:06:35, Serial0/0
 150.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
O IA   150.1.7.7/32 [110/130] via 155.1.0.5, 00:06:36, Serial0/0

```

Like point-to-multipoint, when neighbors outside area 0 see this segment it appears as a collection of endpoints, not the transit subnet itself.

```
Rack1SW4#show ip route ospf | include 155.1.0.
```

```

 155.1.0.0/16 is variably subnetted, 19 subnets, 2 masks
O IA   155.1.0.2/32 [110/66] via 155.1.108.8, 00:07:09, Port-channel1
O IA   155.1.0.3/32 [110/66] via 155.1.108.8, 00:07:09, Port-channel1
O IA   155.1.0.1/32 [110/66] via 155.1.108.8, 00:07:09, Port-channel1
O IA   155.1.0.4/32 [110/66] via 155.1.108.8, 04:15:58, Port-channel1
O IA   155.1.0.5/32 [110/2] via 155.1.108.8, 04:15:58, Port-channel1

```

The difference between the two types can be seen in the `debug ip packet` output. R5 sends multicast hellos to 224.0.0.5 out the other LAN interfaces which use network type broadcast, along with the point-to-point Serial link to R4 that uses network type point-to-point. Out the Frame Relay link running point-to-multipoint non-broadcast R5 sends unicast hellos. This implies that the neighbor statement needs to be configured under the OSPF process, like the non-broadcast network type, to tell the router which devices to send hellos to.

```
Rack1R5#debug ip packet detail
IP packet debugging is on (detailed)
IP: s=155.1.45.5 (local), d=224.0.0.5 (Serial0/1), len 80, sending
broad/multicast, proto=89
IP: s=155.1.5.5 (local), d=224.0.0.5 (FastEthernet0/1), len 76, sending
broad/multicast, proto=89
IP: s=155.1.58.5 (local), d=224.0.0.5 (FastEthernet0/0), len 80, sending
broad/multicast, proto=89
IP: s=155.1.0.5 (local), d=155.1.0.1 (Serial0/0), len 76, sending, proto=89
IP: s=155.1.0.5 (local), d=155.1.0.3 (Serial0/0), len 76, sending, proto=89
IP: s=155.1.0.5 (local), d=155.1.0.2 (Serial0/0), len 76, sending, proto=89
IP: s=155.1.0.5 (local), d=155.1.0.4 (Serial0/0), len 76, sending, proto=89
IP: s=155.1.0.1 (Serial0/0), d=155.1.0.5, len 76, rcvd 0, proto=89
IP: s=155.1.0.4 (Serial0/0), d=155.1.0.5, len 76, rcvd 0, proto=89
IP: s=155.1.0.3 (Serial0/0), d=155.1.0.5, len 104, rcvd 0, proto=89
IP: s=155.1.0.2 (Serial0/0), d=155.1.0.5, len 104, rcvd 0, proto=89
```



## 6.7 OSPF Network Loopback

- Advertise the Loopback0 interfaces of R1, R2, R3, and R4 into OSPF area 0.
- Advertise the Loopback0 interface of R6 into OSPF area 1.
- Advertise the Loopback0 interfaces of SW1, and SW3 into OSPF area 2.
- Advertise the Loopback0 interfaces of R5, SW2, and SW4 into OSPF area 3.
- Modify the network type of these interfaces so that the links are not advertised as host routes.

### Configuration

---

```
R1:
interface Loopback0
 ip ospf 1 area 0
 ip ospf network point-to-point
```

```
R2:
interface Loopback0
 ip ospf 1 area 0
 ip ospf network point-to-point
```

```
R3:
interface Loopback0
 ip ospf 1 area 0
 ip ospf network point-to-point
```

```
R4:
interface Loopback0
 ip ospf 1 area 0
 ip ospf network point-to-point
```

```
R5:
interface Loopback0
 ip ospf 1 area 3
 ip ospf network point-to-point
```

```
R6:
interface Loopback0
 ip ospf 1 area 1
 ip ospf network point-to-point
```

```
SW1:
interface Loopback0
 ip ospf network point-to-point
```

```
SW2:
interface Loopback0
 ip ospf network point-to-point
!
router ospf 1
 network 150.1.8.8 0.0.0.0 area 3
```

```
SW3:
interface Loopback0
 ip ospf network point-to-point
!
router ospf 1
 network 150.1.9.9 0.0.0.0 area 2
```

```
SW4:
interface Loopback0
 ip ospf network point-to-point
!
router ospf 1
 network 150.1.10.10 0.0.0.0 area 3
```

### Verification

#### Note

OSPF network type Loopback is a special case that is used for Loopback and looped back interfaces. Similar to point-to-multipoint an interface running network type Loopback is advertised as a stub endpoint instead of a transit subnet. The result of this network type that we see in our configuration is that when the Loopback0 interfaces of these devices are advertised into OSPF they appear in the routing table as /32 host routes, instead of the actual subnet mask of /24.

```
Rack1R1#show ip route 150.1.10.0 255.255.255.0
% Subnet not in table
```

```
Rack1R1#show ip route 150.1.10.10
Routing entry for 150.1.10.10/32
  Known via "ospf 1", distance 110, metric 67, type inter area
  Last update from 155.1.0.5 on Serial0/0, 00:00:07 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.5.5, 00:00:07 ago, via Serial0/0
      Route metric is 67, traffic share count is 1
```

```
Rack1SW4#show ip ospf interface Loopback0
Loopback0 is up, line protocol is up
  Internet Address 150.1.10.10/24, Area 3
  Process ID 1, Router ID 223.255.255.255, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
```

If the network type of the interface is changed to point-to-point the host route advertisement is withdrawn and the interface is advertised with its native mask.

```
Rack1SW4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1SW4(config)#interface Loopback0
Rack1SW4(config-if)#ip ospf network point-to-point
Rack1SW4(config-if)#end
Rack1SW4#

Rack1R1#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, 00:01:39 ago
  Routing Descriptor Blocks:
  * 155.1.0.5, from 150.1.5.5, 00:01:39 ago, via Serial0/0
    Route metric is 67, traffic share count is 1
```

## 6.8 OSPF Path Selection with Auto-Cost

- Modify the global OSPF cost calculation of all devices so that a TenGigabitEthernet interface has a cost of 3, and an OC-3 link has a cost of 193.

### Configuration

---

```
R1 - R6, SW1 - SW4:  
router ospf 1  
  auto-cost reference-bandwidth 30000
```

### Verification

---

#### Note

OSPF automatic cost calculation is an inverse function of the bandwidth of an interface, meaning the higher the bandwidth value of an interface, the lower the cost value. Specifically the formula for this calculation is:

$$\text{Interface Cost} = \text{Reference Bandwidth} / \text{Interface Bandwidth}$$

The issue with this calculation in current practical network designs is that the default reference bandwidth value used is 100Mbps. This means that a 100Mbps FastEthernet interface will have a cost of 1, and furthermore all interfaces with higher bandwidth values are also 1. The result of this is that OSPF routers in the topology can not make an accurate path calculation when comparing an OC-48 POS interface versus a GigabitEthernet interface, as both interfaces revert to a cost of 1.

In order to resolve this the reference bandwidth value is typically modified to allow these higher bandwidth interfaces to have separate, more granular cost values. In this particular case a TenGigabitEthernet interface, or an OC-192 POS interface, are asked to have a cost value of 3, while a 155Mbps OC-3 POS/ATM interface is asked to have a cost value of 193.

For tasks like this it is not necessary to memorize the formula for cost calculation, but instead just to know how to derive it using various show outputs on the IOS CLI. Since the OSPF cost is derived from the bandwidth, changing the interface **bandwidth** value should change the cost. Without auto-cost modification, as seen below, R1 sees a 10Gbps interface with a cost of 1.

```
Rack1R1#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R1(config)#interface FastEthernet0/0
Rack1R1(config-if)#bandwidth ?
  <1-10000000> Bandwidth in kilobits
  inherit      Specify that bandwidth is inherited
  receive      Specify receive-side bandwidth

Rack1R1(config-if)#bandwidth 10000000
Rack1R1(config-if)#do show ip ospf interface FastEthernet0/0 | include Cost
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 1
```

Changing the reference bandwidth to 100,000 has the following result:

```
Rack1R1(config-if)#router ospf 1
Rack1R1(config-router)#auto-cost reference-bandwidth 100000
Rack1R1(config-router)#do show ip ospf interface FastEthernet0/0 | include Cost
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 10
```

Since the desired result is a cost of 3 for 10Gbps, the reference bandwidth 100,000 is too high. By process of elimination the correct value can be inferred.

```
Rack1R1(config-router)#auto-cost reference-bandwidth 50000
Rack1R1(config-router)#do show ip ospf interface FastEthernet0/0 | include Cost
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 5
Rack1R1(config-router)#auto-cost reference-bandwidth 40000
Rack1R1(config-router)#do show ip ospf interface FastEthernet0/0 | include Cost
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 4
Rack1R1(config-router)#auto-cost reference-bandwidth 30000
Rack1R1(config-router)#do show ip ospf interface FastEthernet0/0 | include Cost
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 3
```

A reference bandwidth of 30,000 results in a cost of 3 (30000 reference bw/ 10000 interface bw = 3). This calculation then stays true for 155Mbps OC-3.

```
Rack1R1(config-router)#interface FastEthernet0/0
Rack1R1(config-if)#bandwidth 10000000
Rack1R1(config-if)#do show ip ospf interface FastEthernet0/0 | include Cost
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 3
Rack1R1(config-if)#bandwidth 155000
Rack1R1(config-if)#do show ip ospf interface FastEthernet0/0 | include Cost
  Process ID 1, Router ID 150.1.1.1, Network Type BROADCAST, Cost: 193
```

### **Pitfall**

Technically OSPF auto-cost reference-bandwidth does not need to match for neighbors to form adjacency. However if different neighbors throughout the topology calculate SPF based on conflicting cost calculations, loops can occur. For this reason it is always recommended to set the auto-cost value consistent throughout the entire OSPF domain.

## 6.9 OSPF Path Selection with Cost

- Using the interface level `ip ospf cost` command modify the OSPF domain so that traffic from SW4 going to the Loopback0 network of R6 uses the point-to-point Serial link between R4 and R5.

### Configuration

```
R5:
interface Serial0/1
 ip ospf cost 10000
```

### Verification

#### Note

In order to understand how OSPF path selection can be modified for traffic engineering purposes, it is important to understand why the current path is being selected. OSPF is sometimes described as being partially link-state, and partially distance vector. This is due to the fact that SPF calculations are only performed for Intra-Area routing, while the Area Border Routers' (ABRs) advertised information is trusted for Inter-Area and External calculations.

In this particular scenario the goal is to modify the traffic flow so that packets coming from SW4 going to the network 150.1.6.0/24 use the point-to-point Serial link between R4 and R5. By visually mapping out the traffic flow we can see that SW4 has only one possible path to reach this network in the first place, the EtherChannel link in area 3 to SW2. SW2 likewise has only one path, the VLAN 58 link in area 3 to R5. R5 has multiple possible paths, the Frame Relay link to R1, the Frame Relay link to R4, and the point-to-point Serial link to R5. Since we are trying to affect the decision process of R5 this means that any cost changes in area 3 will not affect the traffic flow. Only calculation changes on or upstream from R5 can affect this path.

To reinforce this let's first find out how, and more importantly why, SW4 is choosing its current path. This can be determined by finding the longest match R6's Loopback0 interface via the `show ip route` output.

```
Rack1SW4#show ip route 150.1.6.6
Routing entry for 150.1.6.0/24
  Known via "ospf 1", distance 110, metric 19911, type inter area
  Last update from 155.1.108.8 on Port-channel1, 00:03:12 ago
  Routing Descriptor Blocks:
  * 155.1.108.8, from 150.1.5.5, 00:03:12 ago, via Port-channel1
    Route metric is 19911, traffic share count is 1
```

SW4 says that the longest match to 150.1.6.6 is an OSPF Inter-Area route to 150.1.6.0/24 via the next-hop 155.1.108.8 (SW2). Note that after the next-hop value this output shows that the route is *from 150.1.5.5*. This means that R5 is the ABR originating the Network Summary LSA (LSA 3) from area 0 into area 3. By viewing the detailed information about this LSA 3 we can see what R5's cost to the destination is. Specifically the syntax for this is **show ip ospf database summary 150.1.6.0**, where 150.1.6.0 is the link-state ID (the network without the mask) for the destination.

```
Rack1SW4#show ip ospf database summary 150.1.6.0

      OSPF Router with ID (223.255.255.255) (Process ID 1)

      Summary Net Link States (Area 3)

Routing Bit Set on this LSA
LS age: 212
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.6.0 (summary Network Number)
Advertising Router: 150.1.5.5
LS Seq Number: 80000004
Checksum: 0x7129
Length: 28
Network Mask: /24
      TOS: 0 Metric: 19731
```

The above output indicates that SW4 is using the ABR 150.1.5.5 as the exit point to reach 150.1.6.0/24. Note that the metric value of 19731 in this output is not SW4's metric to the destination, but the ABR's (R5's) metric to the destination. In order for SW4 to actually route to this destination it must now run SPF on the advertising router 150.1.5.5 (R5's RID).

This is why OSPF is considered partially distance vector and partially link-state. SW4 does not know the detailed information about the path selection that is occurring behind R5. SW4's trusts R5's calculation for the prefix, which is the distance vector behavior, and then calculates the shortest Intra-Area path to reach R5, which is the link-state behavior. The advantage of this design inherent to OSPF is that SW4 does not need to run SPF on the entire OSPF topology, which immensely increases scalability. Now let's find out how SW4 routes to get to R5.



OSPF's Intra-Area SPF determines this by finding out who SW4 is directly adjacent with, and what the cost values to these neighbors are. Once this is determined, SW4 asks these neighbors who they are adjacent with, and what their costs are to their adjacent neighbors. The process continues over and over until SW4 finds the shortest path to the ABR 150.1.5.5. This Intra-Area metric is then added to the LSA 3 metric that R5 is advertising, and results in the end-to-end metric value that is actually installed in the routing table. We can verify this by first finding out who SW4 is adjacent with by viewing its Router LSA (LSA 1). Specifically this is accomplished with the `show ip ospf database router 223.255.255.255` syntax, where 223.255.255.255 is SW4's OSPF Router-ID.

```
Rack1SW4#show ip ospf database router 223.255.255.255
```

```
OSPF Router with ID (223.255.255.255) (Process ID 1)
```

```
Router Link States (Area 3)
```

```
LS age: 699
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 223.255.255.255
Advertising Router: 223.255.255.255
LS Seq Number: 8000005E
Checksum: 0xBA0E
Length: 60
Number of Links: 3
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 150.1.10.10
(Link Data) Network Mask: 255.255.255.255
Number of TOS metrics: 0
TOS 0 Metrics: 1
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.108.10
(Link Data) Router Interface address: 155.1.108.10
Number of TOS metrics: 0
TOS 0 Metrics: 150
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 155.1.10.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 30
```

SW4 sees that it is adjacent with the *Designated Router address: 155.1.108.10*, and that the cost is 150. Although technically the DR for this segment is SW4 itself, OSPF (like IS-IS) treats the DR as a separate "pseudonode" which it must consult first before computing SPF. This is because the EtherChannel link connecting to SW2 runs OSPF network type broadcast, which requires the DR/BDR election. Therefore SW4 must ask the DR who it is adjacent with, and what its cost calculation to its adjacent neighbors is.

This information can be viewed by checking the DR's Network LSA (LSA 2) via the **show ip ospf database network 155.1.108.10**. From the below information we can see that SW4 finds out that the DR for the segment is adjacent with 150.1.8.8 (SW2's RID).

```
Rack1SW4#show ip ospf database network 155.1.108.10

          OSPF Router with ID (223.255.255.255) (Process ID 1)

              Net Link States (Area 3)

Routing Bit Set on this LSA
LS age: 641
Options: (No TOS-capability, DC)
LS Type: Network Links
Link State ID: 155.1.108.10 (address of Designated Router)
Advertising Router: 223.255.255.255
LS Seq Number: 80000002
Checksum: 0x3B83
Length: 32
Network Mask: /24
    Attached Router: 223.255.255.255
    Attached Router: 150.1.8.8
```

Based on this SW4 infers that it should now ask the router 150.1.8.8 who it is adjacent with. We can view this information on SW4 with the **show ip ospf database router 150.1.8.8** or on SW2 with the **show ip ospf database router 150.1.8.8 self-originate**.

```
Rack1SW4#show ip ospf database router 150.1.8.8

          OSPF Router with ID (223.255.255.255) (Process ID 1)

              Router Link States (Area 3)

LS age: 428
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.8.8
Advertising Router: 150.1.8.8
LS Seq Number: 80000005
Checksum: 0x7446
Length: 72
Number of Links: 4

Link connected to: a Stub Network
(Link ID) Network/subnet number: 150.1.8.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
    TOS 0 Metrics: 1

Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.108.10
(Link Data) Router Interface address: 155.1.108.8
Number of TOS metrics: 0
    TOS 0 Metrics: 150
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.58.8
(Link Data) Router Interface address: 155.1.58.8
Number of TOS metrics: 0
TOS 0 Metrics: 30
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 155.1.8.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 30
```

The above output means that SW2 is adjacent with the DR 155.1.58.8 (SW2), and that the cost to the DR is 30. SW2 must now ask the DR who it is adjacent with.

```
Rack1SW4#show ip ospf database network 155.1.58.8
```

```
OSPF Router with ID (223.255.255.255) (Process ID 1)
```

```
Net Link States (Area 3)
```

```
Routing Bit Set on this LSA
LS age: 111
Options: (No TOS-capability, DC)
LS Type: Network Links
Link State ID: 155.1.58.8 (address of Designated Router)
Advertising Router: 150.1.8.8
LS Seq Number: 80000001
Checksum: 0xEA7F
Length: 32
Network Mask: /24
Attached Router: 150.1.8.8
Attached Router: 150.1.5.5
```

Since the DR is directly adjacent with 150.1.5.5, SW4 infers the total cost to 150.1.5.5 as 150 + 30. Since R5 is advertising the metric 19731 in its LSA 3 for 150.1.6.0, and SW4's metric to R5 is 180, we now know why the total metric in the routing table of SW4 is 19911. However we do not yet know how R5's calculation to the final destination occurs. Since area 3 has no visibility on complete Inter-Area calculations, this can only be verified starting on R5.

```

Rack1R5#show ip route 150.1.6.6
Routing entry for 150.1.6.0/24
  Known via "ospf 1", distance 110, metric 19731, type inter area
  Last update from 155.1.0.1 on Serial0/0, 00:06:48 ago
  Routing Descriptor Blocks:
    155.1.45.4, from 223.255.255.255, 00:06:48 ago, via Serial0/1
      Route metric is 19731, traffic share count is 1
    * 155.1.0.4, from 223.255.255.255, 00:06:48 ago, via Serial0/0
      Route metric is 19731, traffic share count is 1
    155.1.0.1, from 150.1.1.1, 00:06:48 ago, via Serial0/0
      Route metric is 19731, traffic share count is 1

```

R5 says that it has three possible paths to the destination 150.1.6.6 via the longest match 150.1.6.0/24. One of these comes from the area 0 router 150.1.1.1 via the next-hop 150.1.0.1 (R1 via the Frame Relay segment), and two of these come from the area 0 router 223.255.255.255 via the next-hops 155.1.45.4 and 155.1.0.4 (R4 via the point-to-point Serial and Frame Relay segments respectively). Since R5 sees this prefix as Inter-Area, it must ask the area 0 ABRs how they are calculating the path, and then run SPF on the ABRs themselves to find the ultimate shortest path. This can be verified on R5 by first viewing the possible paths to the Network Summary LSAs (LSA 3) that the ABRs are advertising, and then to the ABRs that are originating it via their Router LSAs (LSA 1).

```
Rack1R5#show ip ospf database summary 150.1.6.0
```

```
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: 524
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.6.0 (summary Network Number)
Advertising Router: 150.1.1.1
LS Seq Number: 80000001
Checksum: 0x2D5
Length: 28
Network Mask: /24
TOS: 0 Metric: 301

```

```

Routing Bit Set on this LSA
LS age: 524
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.6.0 (summary Network Number)
Advertising Router: 223.255.255.255
LS Seq Number: 80000001
Checksum: 0x830E
Length: 28
Network Mask: /24
TOS: 0 Metric: 301

```

## Summary Net Link States (Area 3)

```

LS age: 523
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.6.0 (summary Network Number)
Advertising Router: 150.1.5.5
LS Seq Number: 80000004
Checksum: 0x7129
Length: 28
Network Mask: /24
TOS: 0 Metric: 19731

```

R5 sees three possible paths to the destination in the OSPF database. The first is in area 0 via the ABR 150.1.1.1 (R1), the second is in area 0 via the ABR 223.255.255.255 (R4), and the third in area 3 via the ABR 150.1.5.5 (R5). Since the area 3 LSA 3 is originated by R5 itself, this one is immediately eliminated. Now R5 must find the shortest path to the two ABRs 150.1.1.1 and 223.255.255.255 in area 0. This is accomplished the same way it was in area 3, by finding out the cost to R5's adjacent neighbors, and then recursing over and over to find their cost and their adjacent neighbors' cost towards the ABRs. In this case the topology is fairly simple because R5 is directly adjacent with both of these area 0 routers.

```
Rack1R5#show ip ospf database router 150.1.5.5 self-originate
```

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

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

```

LS age: 769
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.5.5
Advertising Router: 150.1.5.5
LS Seq Number: 80000052
Checksum: 0xDCE9
Length: 108
Area Border Router
Number of Links: 7

```

```

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.1.1
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 19430

```

```

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.3.3
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 19430

```

```

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.2.2
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 19430

```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 223.255.255.255
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 19430
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 155.1.0.5
(Link Data) Network Mask: 255.255.255.255
Number of TOS metrics: 0
TOS 0 Metrics: 0
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 223.255.255.255
(Link Data) Router Interface address: 155.1.45.5
Number of TOS metrics: 0
TOS 0 Metrics: 19430
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 155.1.45.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 19430
```

<output omitted>

Two points should be noted about the above output. First, the area 3 output that R5 would normally have originated is omitted, because this does not affect its path selection out towards area 1. Secondly, the output of the three area 0 learned routes indicates that the information is from a *Link connected to: another Router (point-to-point)*. This is due to the fact that the OSPF network type point-to-multipoint non-broadcast, which is running over the Frame Relay link, and the network type point-to-point, which is running over the direct link to R4, do not support the DR/BDR election. Recall that this output in area 3 was different because SW4 and SW2 could only learn LSA 1 information via the DR's pseudonode from LSA 2. Network types point-to-multipoint, point-to-multipoint non-broadcast, and point-to-point do not support the DR/BDR election, and therefore do not require the Network LSA (LSA 2) to recurse towards the adjacent neighbor.

At this point we can now see why R5 is installing three paths to the same destination. R1 is originating the LSA 3 with a metric of 301, as is R4. R5 has a metric of 19430 to R1 and R4 via the Frame Relay, and to R4 via the point-to-point link. Since the end-to-end metric is the same all three routes are installed.

Since the goal of the section is to route only via the point-to-point link to R4, we can now see where changes can and cannot be made. R1's metric to R6 could be raised, which would raise its LSA 3 advertisement to R5. This in turn would cause R5 to choose only R4 and not R1. However this still results in R5 routing towards both links to R4. A lowering of R4's metric to R6 will make it lower the LSA 3 advertisement to R5, and in turn make it more preferred over R1. Again however this still results in R5 routing towards both links to R4, because it affects the LSA 3 sent out both of R4's links. Therefore the only way to force the path to only the point-to-point link is either to lower the cost of the point-to-point link on R5, or to raise the cost of the Frame Relay. In this solution the point-to-point link's cost is lowered from 19430 to 10000. The result of this can be seen in R5's LSA 1 advertisement.

```
Rack1R5#show ip ospf database router 150.1.5.5 self-originate
```

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

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

```
LS age: 75
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.5.5
Advertising Router: 150.1.5.5
LS Seq Number: 8000000D
Checksum: 0xCA37
Length: 108
Area Border Router
Number of Links: 7
```

```
<output omitted>
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 223.255.255.255
(Link Data) Router Interface address: 155.1.45.5
Number of TOS metrics: 0
TOS 0 Metrics: 10000
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 155.1.45.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 10000
```

Final verification of the change in traffic flow can be seen in R5's routing table and in a traceroute from SW4.

```
Rack1R5#show ip route 150.1.6.6
Routing entry for 150.1.6.0/24
  Known via "ospf 1", distance 110, metric 10301, type inter area
  Last update from 155.1.45.4 on Serial0/1, 00:03:22 ago
  Routing Descriptor Blocks:
  * 155.1.45.4, from 223.255.255.255, 00:03:22 ago, via Serial0/1
    Route metric is 10301, traffic share count is 1
```

```
Rack1SW4#traceroute 150.1.6.6
```

```
Type escape sequence to abort.
Tracing the route to 150.1.6.6
```

```
 1 155.1.108.8 0 msec 4 msec 0 msec
 2 155.1.58.5 0 msec 4 msec 0 msec
 3 155.1.45.4 16 msec 16 msec 12 msec
 4 155.1.146.6 24 msec * 24 msec
```



## 6.10 OSPF Path Selection with Bandwidth

- Using the interface level `bandwidth` keyword modify the OSPF domain so that traffic from R6 to the Loopback0 network of SW2 is first sent to R1.

### Configuration

```
R1:
interface Serial0/0
 bandwidth 10000
```

### Verification

#### Note

Like the previous section it is first important to understand why the current path selection occurs before modifying it. In this case R6 has two routes to the destination 150.1.8.8 via the longest match 150.1.8.0/24 from R1 and R4.

```
Rack1R6#show ip route 150.1.8.8
Routing entry for 150.1.8.0/24
  Known via "ospf 1", distance 110, metric 20031, type inter area
  Last update from 155.1.146.4 on FastEthernet0/0.146, 00:27:54 ago
  Routing Descriptor Blocks:
    155.1.146.4, from 223.255.255.255, 00:27:54 ago, via FastEthernet0/0.146
      Route metric is 20031, traffic share count is 1
  * 155.1.146.1, from 150.1.1.1, 00:27:54 ago, via FastEthernet0/0.146
      Route metric is 20031, traffic share count is 1
```

The metric 20031 through both R1 and R4 implies that the addition of R1 and R4's LSA 3 advertisement for 150.1.8.0 plus R6's metric to the DR for VLAN 146 is the same. This can be verified by viewing the ABRs' LSA 3 advertisements, and R6's LSA 1 advertisement.

```
Rack1R6#show ip ospf database summary 150.1.8.0
```

```
OSPF Router with ID (150.1.6.6) (Process ID 1)
```

```
Summary Net Link States (Area 1)
```

```
Routing Bit Set on this LSA
LS age: 395
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.8.0 (summary Network Number)
Advertising Router: 150.1.1.1
LS Seq Number: 80000002
Checksum: 0x930F
Length: 28
Network Mask: /24
TOS: 0 Metric: 19731
```

```
Routing Bit Set on this LSA
LS age: 379
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.8.0 (summary Network Number)
Advertising Router: 223.255.255.255
LS Seq Number: 80000002
Checksum: 0x1547
Length: 28
Network Mask: /24
TOS: 0 Metric: 19731
```

```
Summary Net Link States (Area 2)
```

```
Routing Bit Set on this LSA
LS age: 380
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.8.0 (summary Network Number)
Advertising Router: 150.1.3.3
LS Seq Number: 80000002
Checksum: 0xCD03
Length: 28
Network Mask: /24
TOS: 0 Metric: 65836
```

R6 sees the Inter-Area LSA 3 route via three ABRs, to R1 with a metric of 19731, to R4 with a metric of 19731, and to R3 with the metric 65836. R6 must then recurse to LSA 1 to find the shortest Intra-Area paths to these ABRs.

```
Rack1R6#show ip ospf database router 150.1.6.6 self-originate
```

```
OSPF Router with ID (150.1.6.6) (Process ID 1)
```

```
Router Link States (Area 1)
```

```
LS age: 627
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.6.6
Advertising Router: 150.1.6.6
LS Seq Number: 80000008
Checksum: 0x415E
Length: 48
Number of Links: 2
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 150.1.6.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 1
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.146.6
(Link Data) Router Interface address: 155.1.146.6
Number of TOS metrics: 0
TOS 0 Metrics: 300
```

```
Router Link States (Area 2)
```

```
LS age: 1147
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.6.6
Advertising Router: 150.1.6.6
LS Seq Number: 80000005
Checksum: 0x935B
Length: 36
Number of Links: 1
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.67.7
(Link Data) Router Interface address: 155.1.67.6
Number of TOS metrics: 0
TOS 0 Metrics: 300
```

R6 says that the metric to reach the DR 155.1.146.6 for VLAN 146 is 300. Since we know from the topology diagram that the only neighbors on VLAN 146 are R1, R4, and R6, this implies the metric to reach the ABRs R1 and R4 is 300. Since R1 and R4 are advertising the LSA 3 with a metric of 19731, and R3 is advertising it with a metric of 65836, we can assume at this point that traffic from R6 to the destination will transit area 1 and not area 2, and we can exclude R3 from the problem at hand. Next we come to the decision of where cost manipulations can and cannot occur.

By changing R6's metric to the DR of VLAN 146, it would affect the metric to both ABRs R1 and R4. While this will change the overall end-to-end cost, it will not change R6's path selection. By changing R5's area 3 path to the destination this would affect the LSA 3 advertisement R5 sends into area 0, but would still result in the same LSA 3 being received by R6 via R1 and R4, just with a different metric than the default. By raising R4's metric of the LSA 3 sent into area 1, or by lowering R1's metric of the LSA 3 sent into area 1, this would cause R6 to route to R1. In this particular solution R1's metric is lowered in the LSA 3.

Since the OSPF cost value is derived from the bandwidth, changing the **bandwidth** keyword at the interface level changes the cost.

```
Rack1R1#show interface Serial0/0 | include BW
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
```

```
Rack1R1#show ip ospf interface Serial0/0 | include Cost
  Process ID 1, Router ID 150.1.1.1, Network Type POINT_TO_MULTIPOINT, Cost:
  19430
```

```
Rack1R1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R1(config)#interface Serial0/0
Rack1R1(config-if)# bandwidth 10000
Rack1R1(config-if)#end
```

```
Rack1R1#show interface Serial0/0 | include BW
  MTU 1500 bytes, BW 10000 Kbit, DLY 20000 usec,
```

```
Rack1R1#show ip ospf interface Serial0/0 | include Cost
  Process ID 1, Router ID 150.1.1.1, Network Type POINT_TO_MULTIPOINT, Cost:
  3000
```

Based on this change R6 now sees different values in the LSA 3 summary reported by R1 and R4.

```
Rack1R6#show ip ospf database summary 150.1.8.0
```

```
OSPF Router with ID (150.1.6.6) (Process ID 1)
```

```
Summary Net Link States (Area 1)
```

```
Routing Bit Set on this LSA
```

```
LS age: 95
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 150.1.8.0 (summary Network Number)
```

```
Advertising Router: 150.1.1.1
```

```
LS Seq Number: 80000005
```

```
Checksum: 0x7E8F
```

```
Length: 28
```

```
Network Mask: /24
```

```
TOS: 0 Metric: 3301
```

```
Routing Bit Set on this LSA
```

```
LS age: 772
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 150.1.8.0 (summary Network Number)
```

```
Advertising Router: 223.255.255.255
```

```
LS Seq Number: 80000004
```

```
Checksum: 0x1149
```

```
Length: 28
```

```
Network Mask: /24
```

```
TOS: 0 Metric: 19731
```

```
Summary Net Link States (Area 2)
```

```
LS age: 762
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 150.1.8.0 (summary Network Number)
```

```
Advertising Router: 150.1.3.3
```

```
LS Seq Number: 80000004
```

```
Checksum: 0xC905
```

```
Length: 28
```

```
Network Mask: /24
```

```
TOS: 0 Metric: 65836
```

Since R6's cost of 300 to the VLAN 146 DR did not change R6 now has a metric of 300 + 3301 if it routes to R1, or 300 + 19731 if it routes to R4. The end result is that the path selection is changed to route only to R1 with a metric of 3601.

```
Rack1R6#show ip route 150.1.8.8
Routing entry for 150.1.8.0/24
  Known via "ospf 1", distance 110, metric 3601, type inter area
  Last update from 155.1.146.1 on FastEthernet0/0.146, 00:05:18 ago
  Routing Descriptor Blocks:
  * 155.1.146.1, from 150.1.1.1, 00:05:18 ago, via FastEthernet0/0.146
    Route metric is 3601, traffic share count is 1
```

Final verification of this can be seen via a traceroute from R6.

```
Rack1R6#traceroute 150.1.8.8

Type escape sequence to abort.
Tracing the route to 150.1.8.8

 0 155.1.146.1 4 msec 4 msec 0 msec
 1 155.1.0.5 24 msec 20 msec 20 msec
 2 155.1.58.8 20 msec * 20 msec
```

## 6.11 OSPF Path Selection with Per-Neighbor Cost

- The Frame Relay PVC between R1 & R5 has been provisioned with 512Kbps of bandwidth by the service provider, while the PVC between R4 & R5 has been provisioned at 768Kbps. Configure per-neighbor cost values on R5 to reflect the underlying bandwidth of the Frame Relay circuits connecting it to R1 and R4.

### Configuration

```
router ospf 1
 neighbor 155.1.0.1 cost 58593
 neighbor 155.1.0.4 cost 39062
```

### Verification

#### Note

OSPF cost calculation for a segment is based on the bandwidth value of the outgoing interface. In certain network topologies this can be an issue if the underlying layer 2 network does not map directly to the bandwidth value of the interface connecting to it.

For example in this topology R5's single Serial interface connecting to the Frame Relay cloud has multiple underlying layer 2 PVCs. Although the physical Serial interface has a bandwidth value of T1, which based on the WIC card being used, the PVCs themselves are not provisioned at this rate. Normally the easiest way to fix this is to just change the bandwidth value to match the underlying circuit speed. However the interface does not support setting the bandwidth on a per circuit basis, only on a per interface basis. Therefore even though the circuit to R1 is slower than the circuit to R4, there is no way for OSPF to tell the difference based on the interface itself. In order to resolve this the network types point-to-multipoint and point-to-multipoint non-broadcast support the setting of the OSPF cost value on a per-neighbor basis.

To solve this task we must first determine what the cost values associated with bandwidths of 512Kbps and 768Kbps are. One way is to manually compute this based on the auto-cost formula. The other way is to use the command line output to determine it.

```
Rack1R5#show interface Serial0/0 | include BW
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
Rack1R5#show ip ospf interface Serial0/0 | include Cost
  Process ID 1, Router ID 150.1.5.5, Network Type POINT_TO_MULTIPOINT, Cost:
19430
```

From the above output we can see that R5's T1 Serial interface has a cost of 19430.

```
Rack1R5#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R5(config)#interface Serial0/0
Rack1R5(config-if)#bandwidth 512
Rack1R5(config-if)#do show ip ospf interface Serial0/0 | include Cost
  Process ID 1, Router ID 150.1.5.5, Network Type POINT_TO_MULTIPOINT, Cost:
58593
```

When the bandwidth is changed to 512Kbps, the cost is calculated as 58593.

```
Rack1R5(config-if)#bandwidth 768
Rack1R5(config-if)#do show ip ospf interface Serial0/0 | include Cost
  Process ID 1, Router ID 150.1.5.5, Network Type POINT_TO_MULTIPOINT, Cost:
39062
```

When the bandwidth is changed to 768Kbps, the cost is calculated as 39062. These values can now be configured under the OSPF process as the per-neighbor costs.

```
Rack1R5(config-if)#no bandwidth
Rack1R5(config-if)#router ospf 1
Rack1R5(config-router)#neighbor 155.1.0.1 cost 58593
Rack1R5(config-router)#neighbor 155.1.0.4 cost 39062
Rack1R5(config-router)#end
Rack1R5#
```



Once complete we can see that the interface cost value stays the same as the original default, however R5's router LSA (LSA 1) is changed to reflect the cost values defined on a per-neighbor basis.

```
Rack1R5#show ip ospf interface Serial0/0 | include Cost
  Process ID 1, Router ID 150.1.5.5, Network Type POINT_TO_MULTIPOINT, Cost:
19430
```

```
Rack1R5#show ip ospf database router 150.1.5.5 self-originate
```

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

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

```
LS age: 19
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.5.5
Advertising Router: 150.1.5.5
LS Seq Number: 80000011
Checksum: 0xA2C9
Length: 108
Area Border Router
Number of Links: 7
```

```
  Link connected to: another Router (point-to-point)
    (Link ID) Neighboring Router ID: 150.1.3.3
    (Link Data) Router Interface address: 155.1.0.5
    Number of TOS metrics: 0
    TOS 0 Metrics: 19430
```

```
  Link connected to: another Router (point-to-point)
    (Link ID) Neighboring Router ID: 150.1.1.1
    (Link Data) Router Interface address: 155.1.0.5
    Number of TOS metrics: 0
    TOS 0 Metrics: 58593
```

```
  Link connected to: another Router (point-to-point)
    (Link ID) Neighboring Router ID: 150.1.2.2
    (Link Data) Router Interface address: 155.1.0.5
    Number of TOS metrics: 0
    TOS 0 Metrics: 19430
```

```
  Link connected to: another Router (point-to-point)
    (Link ID) Neighboring Router ID: 223.255.255.255
    (Link Data) Router Interface address: 155.1.0.5
    Number of TOS metrics: 0
    TOS 0 Metrics: 39062
```

```
<output omitted>
```

## 6.12 Repairing Discontiguous OSPF Areas with Virtual-Links

- Configure the network so that full reachability is maintained if R3's connection to SW1 goes down.
- Do not apply any configuration onto R4 to solve this task.

### Configuration

```
R1:
router ospf 1
 area 1 virtual-link 150.1.6.6
```

```
R6:
router ospf 1
 area 1 virtual-link 150.1.1.1
```

### Verification

#### Note

Devices in OSPF area 2 have two exit points out the rest of the network, R3 via area 0, and R6 via area 1. However, although R6 has interfaces in multiple areas, it is technically not an ABR. This can be verified in the OSPF database of devices area 2, as only R3 can originate Summary Network LSAs (LSA 3) to describe Inter-Area destinations.

```
Rack1SW3#show ip ospf database
```

```
OSPF Router with ID (150.1.9.9) (Process ID 1)
```

```
Router Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.3.3	150.1.3.3	1894	0x80000010	0x00FD34	1
150.1.6.6	150.1.6.6	1960	0x8000000F	0x007F65	1
150.1.7.7	150.1.7.7	1038	0x80000012	0x000C7F	5
150.1.9.9	150.1.9.9	1878	0x8000000F	0x0046FE	3

```
Net Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.37.3	150.1.3.3	109	0x8000000E	0x000384
155.1.67.7	150.1.7.7	29	0x8000000E	0x008CCA
155.1.79.9	150.1.9.9	136	0x8000000E	0x001B24

```
Summary Net Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum
150.1.1.0	150.1.3.3	1118	0x8000000D	0x002112
150.1.2.0	150.1.3.3	357	0x80000011	0x001E24
150.1.3.0	150.1.3.3	1378	0x8000000D	0x002EC9
150.1.4.0	150.1.3.3	1898	0x8000000C	0x00249C
150.1.5.0	150.1.3.3	1121	0x8000000C	0x0018DE
150.1.6.0	150.1.3.3	1898	0x8000000C	0x00D0C0

150.1.8.0	150.1.3.3	361	0x80000013	0x001BCC
150.1.10.0	150.1.3.3	361	0x80000013	0x00E668
155.1.0.1	150.1.3.3	1121	0x8000000D	0x00D658
155.1.0.2	150.1.3.3	361	0x80000011	0x00D469
155.1.0.3	150.1.3.3	1382	0x8000000D	0x00E50E
155.1.0.4	150.1.3.3	1898	0x8000000C	0x00DCDF
155.1.0.5	150.1.3.3	1122	0x8000000C	0x00D31F
155.1.5.0	150.1.3.3	361	0x80000011	0x00B803
155.1.8.0	150.1.3.3	361	0x80000013	0x00FCC8
155.1.10.0	150.1.3.3	361	0x80000013	0x00C864
155.1.13.0	150.1.3.3	1643	0x8000000E	0x007276
155.1.23.0	150.1.3.3	1643	0x8000000E	0x0004DA
155.1.45.0	150.1.3.3	1899	0x8000000C	0x00147F
155.1.58.0	150.1.3.3	362	0x80000013	0x00A70A
155.1.108.0	150.1.3.3	362	0x80000013	0x006187
155.1.146.0	150.1.3.3	1899	0x8000000C	0x007B85
192.10.1.0	150.1.3.3	362	0x80000011	0x005192

## Summary ASB Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum
192.10.1.254	150.1.3.3	362	0x80000011	0x004D96

## Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
51.51.51.51	192.10.1.254	1302	0x80000039	0x002CA4	0

Based on this output we can infer that if area 2 loses its connection to R3, all Inter-Area connectivity to area 2 will be lost. This can be demonstrated by shutting down R3's connection to SW1.

```
Rack1R3#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R3(config)#interface FastEthernet0/0
Rack1R3(config-if)#shutdown
Rack1R3(config-if)#end
Rack1R3#
```

Once the link to R3 is down SW1 issues a withdraw message, but the devices in area 2 cache the LSA information from R3 in the OSPF database until the LSA age expires.

```
Rack1SW3#show ip ospf database
```

```
OSPF Router with ID (150.1.9.9) (Process ID 1)
```

```
Router Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.3.3	150.1.3.3	351	0x80000011	0x00FB35	1
150.1.6.6	150.1.6.6	75	0x80000011	0x007172	1
150.1.7.7	150.1.7.7	54	0x80000016	0x00FC57	4
150.1.9.9	150.1.9.9	54	0x80000002	0x003420	3

```
Net Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.37.3	150.1.3.3	592	0x8000000E	0x000384
155.1.67.6	150.1.6.6	76	0x80000001	0x00C5A1
155.1.79.7	150.1.7.7	56	0x80000001	0x0073DE

```
Summary Net Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum
150.1.1.0	150.1.3.3	1600	0x8000000D	0x002112
150.1.2.0	150.1.3.3	840	0x80000011	0x001E24
150.1.3.0	150.1.3.3	1861	0x8000000D	0x002EC9
150.1.4.0	150.1.3.3	352	0x8000000D	0x00229D
150.1.5.0	150.1.3.3	1601	0x8000000C	0x0018DE
150.1.6.0	150.1.3.3	352	0x8000000D	0x00CEC1
150.1.8.0	150.1.3.3	841	0x80000013	0x001BCC
150.1.10.0	150.1.3.3	841	0x80000013	0x00E668
155.1.0.1	150.1.3.3	1602	0x8000000D	0x00D658
155.1.0.2	150.1.3.3	841	0x80000011	0x00D469
155.1.0.3	150.1.3.3	1862	0x8000000D	0x00E50E
155.1.0.4	150.1.3.3	353	0x8000000D	0x00DAE0
155.1.0.5	150.1.3.3	1603	0x8000000C	0x00D31F
155.1.5.0	150.1.3.3	842	0x80000011	0x00B803
155.1.8.0	150.1.3.3	842	0x80000013	0x00FCC8
155.1.10.0	150.1.3.3	842	0x80000013	0x00C864
155.1.13.0	150.1.3.3	2124	0x8000000E	0x007276
155.1.23.0	150.1.3.3	2124	0x8000000E	0x0004DA
155.1.45.0	150.1.3.3	353	0x8000000D	0x001280
155.1.58.0	150.1.3.3	842	0x80000013	0x00A70A
155.1.108.0	150.1.3.3	842	0x80000013	0x006187
155.1.146.0	150.1.3.3	353	0x8000000D	0x007986
192.10.1.0	150.1.3.3	842	0x80000011	0x005192

```
Summary ASB Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.10.1.254	150.1.3.3	842	0x80000011	0x004D96

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

Link ID	ADV Router	Age	Seq#	Checksum	Tag
51.51.51.51	192.10.1.254	1781	0x80000039	0x002CA4	0

Although the LSA 3's still show up as reachable via R3, the recursion process towards R3 fails because the LSA 1 describing R3 is unreachable.

```
Rack1SW3#show ip ospf database router 150.1.3.3

      OSPF Router with ID (150.1.9.9) (Process ID 1)

      Router Link States (Area 2)

    Adv Router is not-reachable
    LS age: 363
    Options: (No TOS-capability, DC)
    LS Type: Router Links
    Link State ID: 150.1.3.3
    Advertising Router: 150.1.3.3
    LS Seq Number: 80000011
    Checksum: 0xFB35
    Length: 36
    Area Border Router
    Number of Links: 1

    Link connected to: a Transit Network
    (Link ID) Designated Router address: 155.1.37.3
    (Link Data) Router Interface address: 155.1.37.3
    Number of TOS metrics: 0
    TOS 0 Metrics: 300
```

The end result of this is that devices in area 2 cannot reach devices in other areas.

```
Rack1SW3#show ip route ospf
 155.1.0.0/24 is subnetted, 4 subnets
O    155.1.7.0 [110/60] via 155.1.79.7, 00:00:55, Vlan79
O    155.1.67.0 [110/60] via 155.1.79.7, 00:00:55, Vlan79
 150.1.0.0/24 is subnetted, 2 subnets
O    150.1.7.0 [110/31] via 155.1.79.7, 00:00:55, Vlan79
```

To resolve this problem R6 must offer devices in area 2 an alternate path to area 0. This can be done with a new link being added to R6 that is in area 0, such as another Ethernet interface or a Tunnel interface, or by configuring an OSPF virtual-link.

An OSPF virtual-link allows the forming of an indirect area 0 adjacency, which can be used to repair breaks in the OSPF domain or solve traffic engineering requirements. In this particular case R6 can virtual-link to the neighbors R1 or R4 who are in area 1, since they both have connections to area 0. R1 is chosen in this solution since the question states not to configure R4.

The first important point to note about the virtual-link is that the neighbor value specified in the virtual-link syntax is the router-id of the neighbor in the transit area. This means that if for some reason the router-id changes, i.e. a new higher Loopback interface is added, or the `router-id` command is changed, the virtual-link will fail.

Additionally the neighbors forming adjacency over the virtual-link do not have to be directly connected; they simply need to know how to recurse towards each others' LSA 1 advertisements. This means that the traffic flow via the virtual-link should naturally follow the Intra-Area SPF calculation between the routers' LSA 1 advertisements. Verification of this can be seen as follows.

```
Rack1R1#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R1(config)#router ospf 1
Rack1R1(config-router)#area 1 virtual-link 150.1.6.6
Rack1R1(config-router)#end
Rack1R1#

Rack1R6#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R6(config)#router ospf 1
Rack1R6(config-router)#area 1 virtual-link 150.1.1.1
Rack1R6(config-router)#end
Rack1R6#
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.1.1 on OSPF_VL0 from LOADING to FULL,
Loading Done

Rack1R6#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.1.1	0	FULL/ -	-	155.1.146.1	OSPF_VL0
150.1.1.1	1	FULL/BDR	00:00:39	155.1.146.1	FastEthernet0/0.146
223.255.255.255	0	FULL/DROTHER	00:00:37	155.1.146.4	FastEthernet0/0.146
150.1.7.7	1	FULL/BDR	00:00:39	155.1.67.7	FastEthernet0/0.67

R1 and R6 form adjacency over the virtual-link. `show ip ospf interface` indicates that the virtual-link is an area 0 interface.

```
Rack1R6#show ip ospf interface
OSPF_VL0 is up, line protocol is up
  Internet Address 155.1.146.6/24, Area 0
  Process ID 1, Router ID 150.1.6.6, Network Type VIRTUAL_LINK, Cost: 300
  Configured as demand circuit.
  Run as demand circuit.
  DoNotAge LSA allowed.
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  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 1/4, 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 150.1.1.1 (Hello suppressed)
  Suppress hello for 1 neighbor(s)
```

The cost value of the virtual-link is based on R6's cost to reach the router 150.1.1.1 in area 1. This can be calculated via LSA 1 recursion.

```
Rack1R6#show ip ospf database router 150.1.6.6 self-originate
```

```
      OSPF Router with ID (150.1.6.6) (Process ID 1)
```

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

```
LS age: 3
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.6.6
Advertising Router: 150.1.6.6
LS Seq Number: 80000002
Checksum: 0xEDE
Length: 36
Area Border Router
Number of Links: 1

  Link connected to: a Virtual Link
    (Link ID) Neighboring Router ID: 150.1.1.1
    (Link Data) Router Interface address: 155.1.146.6
  Number of TOS metrics: 0
    TOS 0 Metrics: 300
```

```
          Router Link States (Area 1)
```

```
LS age: 3
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.6.6
Advertising Router: 150.1.6.6
LS Seq Number: 80000017
Checksum: 0x3259
Length: 48
Area Border Router
Virtual Link Endpoint
```

Number of Links: 2

Link connected to: a Stub Network  
 (Link ID) Network/subnet number: 150.1.6.0  
 (Link Data) Network Mask: 255.255.255.0  
 Number of TOS metrics: 0  
 TOS 0 Metrics: 1

Link connected to: a Transit Network  
 (Link ID) Designated Router address: 155.1.146.6  
 (Link Data) Router Interface address: 155.1.146.6  
 Number of TOS metrics: 0  
 TOS 0 Metrics: 300

#### Router Link States (Area 2)

LS age: 13  
 Options: (No TOS-capability, DC)  
 LS Type: Router Links  
 Link State ID: 150.1.6.6  
 Advertising Router: 150.1.6.6  
 LS Seq Number: 80000012  
 Checksum: 0x726F  
 Length: 36  
 Area Border Router  
 Number of Links: 1

Link connected to: a Transit Network  
 (Link ID) Designated Router address: 155.1.67.6  
 (Link Data) Router Interface address: 155.1.67.6  
 Number of TOS metrics: 0  
 TOS 0 Metrics: 300

R6 now generates an LSA 1 for area 0, along with its previous LSA 1's in areas 1 and 2. Inside the area 0 LSA 1 it says that it is adjacent to R1 over a virtual-link. Inside the area 1 LSA 1 it says that it is the endpoint for the virtual-link, and is adjacent with the DR 155.1.146.6 (itself) with a cost of 300. Asking the DR who it is adjacent with via LSA 2 it implies that R6 is on the same segment as R1, with a cost of 300. This metric of 300 is what the virtual-link inherits. Now since R6 is a true ABR, it should be generating LSA 3 into area 2 about area 0 (and the other areas), and should be generating LSA 3 into area 0 about area 2.

Rack1SW3#show ip ospf database

OSPF Router with ID (150.1.9.9) (Process ID 1)

#### Router Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.3.3	150.1.3.3	2257	0x80000011	0x00FB35	1
150.1.6.6	150.1.6.6	481	0x80000012	0x00726F	1
150.1.7.7	150.1.7.7	131	0x80000017	0x00FA58	4
150.1.9.9	150.1.9.9	26	0x80000003	0x003221	3



## Net Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum
155.1.37.3	150.1.3.3	2498	0x8000000E	0x000384
155.1.67.6	150.1.6.6	34	0x80000002	0x00C3A2
155.1.79.7	150.1.7.7	132	0x80000002	0x0071DF

## Summary Net Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum
150.1.1.0	150.1.3.3	3507	0x8000000D	0x002112
150.1.1.0	150.1.6.6	467	0x80000001	0x00F7DA
150.1.2.0	150.1.3.3	2747	0x80000011	0x001E24
150.1.2.0	150.1.6.6	468	0x80000001	0x0075E5
150.1.3.0	150.1.6.6	468	0x80000001	0x006AEF
150.1.4.0	150.1.3.3	2258	0x8000000D	0x00229D
150.1.4.0	150.1.6.6	463	0x80000002	0x00D4F9
150.1.5.0	150.1.3.3	3507	0x8000000C	0x0018DE
150.1.5.0	150.1.6.6	468	0x80000001	0x0066A4
150.1.6.0	150.1.3.3	2258	0x8000000D	0x00CEC1
150.1.6.0	150.1.6.6	483	0x80000001	0x00FDFC
150.1.8.0	150.1.3.3	2748	0x80000013	0x001BCC
150.1.8.0	150.1.6.6	468	0x80000001	0x00778B
150.1.10.0	150.1.3.3	2748	0x80000013	0x00E668
150.1.10.0	150.1.6.6	468	0x80000001	0x004228
155.1.0.1	150.1.3.3	3509	0x8000000D	0x00D658
155.1.0.1	150.1.6.6	468	0x80000001	0x00AD21
155.1.0.2	150.1.3.3	2748	0x80000011	0x00D469
155.1.0.2	150.1.6.6	468	0x80000001	0x002C2B
155.1.0.3	150.1.6.6	469	0x80000001	0x002234
155.1.0.4	150.1.3.3	2259	0x8000000D	0x00DAE0
155.1.0.4	150.1.6.6	464	0x80000002	0x008D3D
155.1.0.5	150.1.3.3	3509	0x8000000C	0x00D31F
155.1.0.5	150.1.6.6	469	0x80000001	0x0020E6
155.1.5.0	150.1.3.3	2748	0x80000011	0x00B803
155.1.5.0	150.1.6.6	469	0x80000001	0x0010C4
155.1.8.0	150.1.3.3	2748	0x80000013	0x00FCC8
155.1.8.0	150.1.6.6	470	0x80000001	0x005888
155.1.10.0	150.1.3.3	2749	0x80000013	0x00C864
155.1.10.0	150.1.6.6	470	0x80000001	0x002424
155.1.13.0	150.1.6.6	470	0x80000001	0x00D1BE
155.1.23.0	150.1.6.6	470	0x80000001	0x00EB24
155.1.45.0	150.1.3.3	2260	0x8000000D	0x001280
155.1.45.0	150.1.6.6	470	0x80000001	0x006245
155.1.58.0	150.1.3.3	2749	0x80000013	0x00A70A
155.1.58.0	150.1.6.6	470	0x80000001	0x0004C8
155.1.108.0	150.1.3.3	2750	0x80000013	0x006187
155.1.108.0	150.1.6.6	470	0x80000001	0x00BC47
155.1.146.0	150.1.3.3	2261	0x8000000D	0x007986
155.1.146.0	150.1.6.6	485	0x80000001	0x006BD1
192.10.1.0	150.1.3.3	2750	0x80000011	0x005192
192.10.1.0	150.1.6.6	470	0x80000001	0x00A854

## Summary ASB Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum
192.10.1.254	150.1.3.3	2750	0x80000011	0x004D96
192.10.1.254	150.1.6.6	471	0x80000001	0x00A458

## Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
51.51.51.51	192.10.1.254	1685	0x8000003A	0x002AA5	0

Note that in the above output LSA 3's appear in area 2 from both R3 and R6, even though R3 is still unreachable. This is due to the fact that R3's LSAs are still aging out, even though the router R3 itself is still unreachable. The end result of the virtual-link is that traffic from area 2 to area 0, along with the other areas, now transits area 1 via the virtual-link.

```
Rack1SW3#show ip route ospf
 51.0.0.0/32 is subnetted, 1 subnets
O E2   51.51.51.51 [110/20] via 155.1.79.7, 00:07:51, Vlan79
155.1.0.0/16 is variably subnetted, 18 subnets, 2 masks
O IA   155.1.146.0/24 [110/360] via 155.1.79.7, 00:08:01, Vlan79
O IA   155.1.23.0/24 [110/26028] via 155.1.79.7, 00:07:52, Vlan79
O IA   155.1.10.0/24 [110/3545] via 155.1.79.7, 00:07:52, Vlan79
O IA   155.1.8.0/24 [110/3395] via 155.1.79.7, 00:07:52, Vlan79
O IA   155.1.13.0/24 [110/19790] via 155.1.79.7, 00:07:52, Vlan79
O IA   155.1.0.2/32 [110/6598] via 155.1.79.7, 00:07:52, Vlan79
O IA   155.1.0.3/32 [110/6598] via 155.1.79.7, 00:07:52, Vlan79
O IA   155.1.0.1/32 [110/360] via 155.1.79.7, 00:07:52, Vlan79
O     155.1.7.0/24 [110/60] via 155.1.79.7, 00:08:02, Vlan79
O IA   155.1.0.4/32 [110/360] via 155.1.79.7, 00:07:47, Vlan79
O IA   155.1.5.0/24 [110/3410] via 155.1.79.7, 00:07:52, Vlan79
O IA   155.1.0.5/32 [110/3360] via 155.1.79.7, 00:07:52, Vlan79
O IA   155.1.58.0/24 [110/3365] via 155.1.79.7, 00:07:52, Vlan79
O IA   155.1.45.0/24 [110/13360] via 155.1.79.7, 00:07:53, Vlan79
O     155.1.67.0/24 [110/60] via 155.1.79.7, 00:08:02, Vlan79
O IA   155.1.108.0/24 [110/3515] via 155.1.79.7, 00:07:53, Vlan79
O IA 192.10.1.0/24 [110/6898] via 155.1.79.7, 00:07:53, Vlan79
150.1.0.0/24 is subnetted, 10 subnets
O     150.1.7.0 [110/31] via 155.1.79.7, 00:08:02, Vlan79
O IA 150.1.6.0 [110/61] via 155.1.79.7, 00:08:02, Vlan79
O IA 150.1.5.0 [110/3361] via 155.1.79.7, 00:07:53, Vlan79
O IA 150.1.4.0 [110/361] via 155.1.79.7, 00:07:48, Vlan79
O IA 150.1.3.0 [110/6599] via 155.1.79.7, 00:07:53, Vlan79
O IA 150.1.2.0 [110/6599] via 155.1.79.7, 00:07:53, Vlan79
O IA 150.1.1.0 [110/361] via 155.1.79.7, 00:07:53, Vlan79
O IA 150.1.10.0 [110/3516] via 155.1.79.7, 00:07:53, Vlan79
O IA 150.1.8.0 [110/3366] via 155.1.79.7, 00:07:53, Vlan79
```

```
Rack1SW3#traceroute 155.1.10.10
```

```
Type escape sequence to abort.
Tracing the route to 155.1.10.10
```

```
 1 155.1.79.7 0 msec 0 msec 0 msec
 2 155.1.67.6 4 msec 0 msec 4 msec
 3 155.1.146.1 0 msec 4 msec 0 msec
 4 155.1.0.5 32 msec 28 msec 28 msec
 5 155.1.58.8 32 msec 28 msec 28 msec
 6 155.1.108.10 32 msec * 28 msec
```

Once R3's connection to SW1 comes back up, traffic from area 2 still uses R6 as the exit point due to the shorter SPF. For example take SW3's path to the VLAN 10 network 155.1.10.0/24.

```
Rack1R3#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R3(config)#
Rack1R3(config)#interface FastEthernet0/0
Rack1R3(config-if)#no shutdown
Rack1R3(config-if)#end
Rack1R3#
<output omitted>
Rack1R3#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.5.5	0	FULL/ -	00:01:33	155.1.0.5	Serial1/0
150.1.7.7	1	FULL/DR	00:00:37	155.1.37.7	FastEthernet0/0
150.1.1.1	0	FULL/ -	00:00:30	155.1.13.1	Serial1/2
150.1.2.2	0	FULL/ -	00:00:35	155.1.23.2	Serial1/3

```
Rack1SW3#show ip route 150.1.10.0
Routing entry for 150.1.10.0/24
  Known via "ospf 1", distance 110, metric 3516, type inter area
  Last update from 155.1.79.7 on Vlan79, 00:02:30 ago
  Routing Descriptor Blocks:
  * 155.1.79.7, from 150.1.6.6, 00:02:30 ago, via Vlan79
    Route metric is 3516, traffic share count is 1
```

SW3 says the metric to reach 150.1.10.0/24 is 3516 via the ABR 150.1.6.6. Why this path is being chosen can be verified by checking the metric that R3 and R6 are originating as LSA 3, and then find how the LSA 1 Intra-Area SPF calculation occurs to the ABR.

```
Rack1SW3#show ip ospf database summary 155.1.10.0

          OSPF Router with ID (150.1.9.9) (Process ID 1)

          Summary Net Link States (Area 2)

LS age: 1805
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 155.1.10.0 (summary Network Number)
Advertising Router: 150.1.3.3
LS Seq Number: 80000014
Checksum: 0xC665
Length: 28
Network Mask: /24
          TOS: 0 Metric: 65720
```

```

Routing Bit Set on this LSA
LS age: 1504
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 155.1.10.0 (summary Network Number)
Advertising Router: 150.1.6.6
LS Seq Number: 80000001
Checksum: 0x2424
Length: 28
Network Mask: /24
      TOS: 0 Metric: 3485

```

R3 says that the prefix 155.1.10.0/24 is reachable with the metric 65720, while R6 is advertising it with the metric 3485. Now we must determine the shortest path to the ABRs.

```
Rack1SW3#show ip ospf database router 150.1.9.9 self-originate
```

```
      OSPF Router with ID (150.1.9.9) (Process ID 1)
```

```
          Router Link States (Area 2)
```

```

LS age: 700
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.9.9
Advertising Router: 150.1.9.9
LS Seq Number: 80000011
Checksum: 0x162F
Length: 60
Number of Links: 3

```

```

Link connected to: a Stub Network
(Link ID) Network/subnet number: 150.1.9.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
      TOS 0 Metrics: 1

```

```

Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.79.7
(Link Data) Router Interface address: 155.1.79.9
Number of TOS metrics: 0
      TOS 0 Metrics: 30

```

```
<output omitted>
```

SW3's metric to the VLAN 79 DR (and hence SW1) is 30. The total metric to R3's advertisement is now 65750, while R6's advertisement is 3515.

```
Rack1SW1#show ip ospf database router 150.1.7.7 self-originate
```

```
OSPF Router with ID (150.1.7.7) (Process ID 1)
```

```
Router Link States (Area 2)
```

```
LS age: 767
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.7.7
Advertising Router: 150.1.7.7
LS Seq Number: 80000019
Checksum: 0xE39F
Length: 84
Number of Links: 5
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 150.1.7.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 1
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.37.7
(Link Data) Router Interface address: 155.1.37.7
Number of TOS metrics: 0
TOS 0 Metrics: 300
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.79.7
(Link Data) Router Interface address: 155.1.79.7
Number of TOS metrics: 0
TOS 0 Metrics: 30
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.67.6
(Link Data) Router Interface address: 155.1.67.7
Number of TOS metrics: 0
TOS 0 Metrics: 30
```

```
<output omitted>
```

SW1's metric to DR on the link to R3 (and hence R3) is 300, while the metric to the VLAN 67 DR (and hence R6) is 30. The total metric through R3 is now 66050, while the total metric to R6 is 3545. This results in the path through R6 being installed in the routing table.

```
Rack1SW3#show ip route 155.1.10.0
Routing entry for 155.1.10.0/24
Known via "ospf 1", distance 110, metric 3545, type inter area
Last update from 155.1.79.7 on Vlan79, 01:00:04 ago
Routing Descriptor Blocks:
* 155.1.79.7, from 150.1.6.6, 01:00:04 ago, via Vlan79
Route metric is 3545, traffic share count is 1
```

### 6.13 OSPF Path Selection with Non-Backbone Transit Areas

- Configure R4's Frame Relay and point-to-point Serial links connecting to R5 to have an OSPF cost of 1.
- Modify the SPF calculation in the OSPF domain so that R6 cannot use R4 to reach area 0 by transiting area 1.
- Verify this by ensuring that traffic from R6 going to the Loopback0 network of SW2 is still first sent to R1, and then directly to R5.

#### Configuration

```
R1:
router ospf 1
  no capability transit

R4:
interface Serial0/0
  ip ospf cost 1
!
interface Serial0/1
  ip ospf cost 1

R6:
router ospf 1
  no capability transit
```

#### Verification

##### Note

Per RFC 2328 (OSPF Version 2) section 16.3, *Examining transit areas' summary-LSAs*, non-backbone (not area 0) areas can be used for inter-area transit if a shorter path can be found through them, and if the “*TransitCapability parameter has been set to TRUE in Step 2 of the Dijkstra algorithm*”. In Cisco's IOS implementation this flag is controlled with the `capability transit` routing process level command, and is on by default.

The design case when the feature is used is very specific, and has to do with a shorter Inter-Area path being found via a non-area 0 router as compared to the target router of a virtual-link. To understand this in detail let's first see R6's path selection to 150.1.8.0/24 when R4's links to R5 both have a cost of 1.

```
Rack1R6#show ip route 150.1.8.8
Routing entry for 150.1.8.0/24
  Known via "ospf 1", distance 110, metric 307, type inter area
  Last update from 155.1.146.4 on FastEthernet0/0.146, 00:12:31 ago
  Routing Descriptor Blocks:
  * 155.1.146.4, from 150.1.5.5, 00:12:31 ago, via FastEthernet0/0.146
    Route metric is 307, traffic share count is 1
```

R6's virtual-link to area 0 is via R1. Normally we would assume that R6 must route via the virtual-link path to R1 to reach Inter-Area destinations advertised by area 0. However in this case R6 says that the Inter-Area route 150.1.8.0/24 is via the next-hop 155.1.146.4 (R4), but originated by the router 150.1.5.5 (R5). The below traceroute output indicates that R1 is not used in the transit path for this traffic.

```
Rack1R6#traceroute 150.1.8.8

Type escape sequence to abort.
Tracing the route to 150.1.8.8

 0 155.1.146.4 4 msec 0 msec 0 msec
 1 155.1.0.5 20 msec 20 msec 20 msec
 2 155.1.58.8 20 msec * 20 msec
```

To determine why R6 chooses this path we must first find out how R6 is learning the Inter-Area LSA 3 route.

```
Rack1R6#show ip ospf database summary 150.1.8.0

          OSPF Router with ID (150.1.6.6) (Process ID 1)

          Summary Net Link States (Area 0)

Routing Bit Set on this LSA
LS age: 342 (DoNotAge)
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.8.0 (summary Network Number)
Advertising Router: 150.1.5.5
LS Seq Number: 80000014
Checksum: 0x1E1
Length: 28
Network Mask: /24
          TOS: 0 Metric: 6
<output omitted>
```

This output indicates that 150.1.5.5 (R5) is the originating ABR, and has a metric of 6. We must now figure out how to route towards the ABR R5.

```
Rack1R6#show ip ospf database router 150.1.5.5

      OSPF Router with ID (150.1.6.6) (Process ID 1)

      Router Link States (Area 0)

Routing Bit Set on this LSA
LS age: 342 (DoNotAge)
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.5.5
Advertising Router: 150.1.5.5
LS Seq Number: 8000001E
Checksum: 0x2539
Length: 108
Area Border Router
Number of Links: 7

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.3.3
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 3238

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.1.1
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 58593

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.2.2
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 3238

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 223.255.255.255
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 39062

Link connected to: a Stub Network
(Link ID) Network/subnet number: 155.1.0.5
(Link Data) Network Mask: 255.255.255.255
Number of TOS metrics: 0
TOS 0 Metrics: 0

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 223.255.255.255
(Link Data) Router Interface address: 155.1.45.5
Number of TOS metrics: 0
TOS 0 Metrics: 10000

<output omitted>
```



R5's area 0 Router LSA (LSA 1) states that it is adjacent with R1 over the Frame Relay network, and R4 over the Frame Relay and point-to-point Serial links. Now we must determine what R1 and R4's costs to R5 are.

```
Rack1R6#show ip ospf database router 150.1.1.1

      OSPF Router with ID (150.1.6.6) (Process ID 1)

      Router Link States (Area 0)

Routing Bit Set on this LSA
LS age: 1 (DoNotAge)
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.1.1
Advertising Router: 150.1.1.1
LS Seq Number: 8000001B
Checksum: 0x4E2C
Length: 72
Area Border Router
Number of Links: 4

  Link connected to: a Virtual Link
    (Link ID) Neighboring Router ID: 150.1.6.6
    (Link Data) Router Interface address: 155.1.146.1
    Number of TOS metrics: 0
      TOS 0 Metrics: 300

  Link connected to: a Stub Network
    (Link ID) Network/subnet number: 150.1.1.0
    (Link Data) Network Mask: 255.255.255.0
    Number of TOS metrics: 0
      TOS 0 Metrics: 1

  Link connected to: another Router (point-to-point)
    (Link ID) Neighboring Router ID: 150.1.5.5
    (Link Data) Router Interface address: 155.1.0.1
    Number of TOS metrics: 0
      TOS 0 Metrics: 3000
```

<output omitted>

R1 says it's metric to R5 is 3000 via the Frame Relay link.

```
Rack1R6#show ip ospf database router 223.255.255.255
```

```
      OSPF Router with ID (150.1.6.6) (Process ID 1)
```

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

```
Routing Bit Set on this LSA
```

```
LS age: 3 (DoNotAge)
```

```
Options: (No TOS-capability, DC)
```

```
LS Type: Router Links
```

```
Link State ID: 223.255.255.255
```

```
Advertising Router: 223.255.255.255
```

```
LS Seq Number: 8000001D
```

```
Checksum: 0x9EBF
```

```
Length: 84
```

```
Area Border Router
```

```
Number of Links: 5
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 150.1.4.0
```

```
(Link Data) Network Mask: 255.255.255.0
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 1
```

```
Link connected to: another Router (point-to-point)
```

```
(Link ID) Neighboring Router ID: 150.1.5.5
```

```
(Link Data) Router Interface address: 155.1.0.4
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 1
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 155.1.0.4
```

```
(Link Data) Network Mask: 255.255.255.255
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 0
```

```
Link connected to: another Router (point-to-point)
```

```
(Link ID) Neighboring Router ID: 150.1.5.5
```

```
(Link Data) Router Interface address: 155.1.45.4
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 1
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 155.1.45.0
```

```
(Link Data) Network Mask: 255.255.255.0
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 1
```

R4 says its metrics are both 1 to R5 over the Frame Relay and point-to-point links. Now we must find what R6's metrics to R1 and R4 are.

```
Rack1R6#show ip ospf database router 150.1.6.6 self-originate
```

```
OSPF Router with ID (150.1.6.6) (Process ID 1)
```

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

```
LS age: 50
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.6.6
Advertising Router: 150.1.6.6
LS Seq Number: 80000023
Checksum: 0xCBFF
Length: 36
Area Border Router
Number of Links: 1
```

```
Link connected to: a Virtual Link
(Link ID) Neighboring Router ID: 150.1.1.1
(Link Data) Router Interface address: 155.1.146.6
Number of TOS metrics: 0
TOS 0 Metrics: 300
```

```
Router Link States (Area 1)
```

```
LS age: 50
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.6.6
Advertising Router: 150.1.6.6
LS Seq Number: 80000038
Checksum: 0xEF7A
Length: 48
Area Border Router
Virtual Link Endpoint
Number of Links: 2
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 150.1.6.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 1
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.146.6
(Link Data) Router Interface address: 155.1.146.6
Number of TOS metrics: 0
TOS 0 Metrics: 300
```

```
<output omitted>
```

R6 has reachability to R1 via the virtual-link with a metric of 300. R6 has reachability to R4 via the DR of VLAN 146 in area 1 with a metric of 300. Now R6 must make its final determination how to route.

If R6 routes through R1 via the virtual-link the metric is 300 to R1, 3000 from R1 to R5, and R5's metric of 6 to the final destination, for a total of 3306. If R6 routes through R4 via area 1 the metric is 300 to R4, 1 from R4 to R5, and R5's metric of 6 to the final destination, for a total of 307.

Under normal conditions R6 should choose to route through R1 since this is the route through area 0. However if the **capability transit** feature is enabled, which it is by default in IOS 12.4, R6 can choose the shorter path through area 1 to R4. This results in the final metric of 307 via R4 being installed in the routing table.

```
Rack1R6#show ip route 150.1.8.8
Routing entry for 150.1.8.0/24
  Known via "ospf 1", distance 110, metric 307, type inter area
  Last update from 155.1.146.4 on FastEthernet0/0.146, 17:01:50 ago
  Routing Descriptor Blocks:
  * 155.1.146.4, from 150.1.5.5, 17:01:50 ago, via FastEthernet0/0.146
    Route metric is 307, traffic share count is 1
```

If area 1 is not allowed to be used as transit, by issuing the **no capability transit**, R6 cannot choose this route.

```
Rack1R6#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R6(config)#router ospf 1
Rack1R6(config-router)#no capability transit
Rack1R6(config-router)#end
Rack1R6#

Rack1R6#show ip route 150.1.8.8
Routing entry for 150.1.8.0/24
  Known via "ospf 1", distance 110, metric 3306, type inter area
  Last update from 155.1.146.1 on FastEthernet0/0.146, 00:00:06 ago
  Routing Descriptor Blocks:
  * 155.1.146.1, from 150.1.5.5, 00:00:06 ago, via FastEthernet0/0.146
    Route metric is 3306, traffic share count is 1
```

R6 now says the route is through R1 with a metric of 3306, which matches what we calculated from the OSPF database. The traceroute output however is still not exactly the desired result.

```
Rack1R6#traceroute 150.1.8.8
```

```
Type escape sequence to abort.
Tracing the route to 150.1.8.8
```

```
 1 155.1.146.1 0 msec 4 msec 4 msec
 2 155.1.146.4 4 msec 0 msec 4 msec
 3 155.1.0.5 20 msec 20 msec 20 msec
 4 155.1.58.8 20 msec * 20 msec
```

R6 sends its packet to R1, but R1 takes it and redirects it out the same interface back to R4. This is due to the fact that R1 now says its shorter path to the router R5 is through area 1 via R4.

```
Rack1R1#show ip ospf database router 150.1.1.1 self-originate
```

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

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

```
LS age: 740
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.1.1
Advertising Router: 150.1.1.1
LS Seq Number: 8000003C
Checksum: 0xC4D
Length: 72
Area Border Router
Number of Links: 4
```

```
Link connected to: a Virtual Link
(Link ID) Neighboring Router ID: 150.1.6.6
(Link Data) Router Interface address: 155.1.146.1
Number of TOS metrics: 0
TOS 0 Metrics: 300
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 150.1.1.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 1
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.5.5
(Link Data) Router Interface address: 155.1.0.1
Number of TOS metrics: 0
TOS 0 Metrics: 3000
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 155.1.0.1
(Link Data) Network Mask: 255.255.255.255
Number of TOS metrics: 0
TOS 0 Metrics: 0
```

## Router Link States (Area 1)

```
LS age: 744
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.1.1
Advertising Router: 150.1.1.1
LS Seq Number: 80000035
Checksum: 0x6DC7
Length: 36
Area Border Router
Virtual Link Endpoint
Number of Links: 1
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.146.6
(Link Data) Router Interface address: 155.1.146.1
Number of TOS metrics: 0
TOS 0 Metrics: 300
```

R1 says that its direct adjacency with R5 over the Frame Relay link is via a metric of 3000 in area 0. R1's adjacency with the VLAN 146 DR (and hence R4) is 300 via area 1. Since R4's metric to R5 is 1, R1 says that the metric of 301 via area 1 is closer than the metric of 3000 via area 0. Since capability transit is on R1, it can use R4 to reach R5 instead of using area 0. With R5's metric of 6 to the destination R1 then inserts the total metric of 307 via R4, but from R5.

```
Rack1R1#show ip route 150.1.8.8
Routing entry for 150.1.8.0/24
  Known via "ospf 1", distance 110, metric 307, type inter area
  Last update from 155.1.146.4 on FastEthernet0/0, 00:11:39 ago
  Routing Descriptor Blocks:
    * 155.1.146.4, from 150.1.5.5, 00:11:39 ago, via FastEthernet0/0
      Route metric is 307, traffic share count is 1
```

Disabling R1's ability to use area 1 as transit to reach area 0 forces R1 to route directly to R5, instead of to R4. The final result is that the traffic flow from R6 to SW2 is now as desired.

```
Rack1R1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R1(config)#router ospf 1
Rack1R1(config-router)#no capability transit
Rack1R1(config-router)#end
Rack1R1#

Rack1R1#show ip route 150.1.8.8
Routing entry for 150.1.8.0/24
  Known via "ospf 1", distance 110, metric 3006, type inter area
  Last update from 155.1.0.5 on Serial0/0, 00:00:06 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.5.5, 00:00:06 ago, via Serial0/0
      Route metric is 3006, traffic share count is 1

Rack1R6#traceroute 150.1.8.8

Type escape sequence to abort.
Tracing the route to 150.1.8.8

 1 155.1.146.1 4 msec 4 msec 0 msec
 2 155.1.0.5 24 msec 20 msec 24 msec
 3 155.1.58.8 20 msec * 20 msec
```

## 6.14 OSPF Path Selection with Virtual-Links

- Configure the OSPF domain so that traffic from SW3 going to OSPF area 51 transits the point-to-point Serial link between R3 and R2.

### Configuration

---

```
R2:
router ospf 1
 area 5 virtual-link 150.1.3.3
```

```
R3:
interface Serial1/3
 ip ospf cost 1
!
router ospf 1
 area 5 virtual-link 150.1.2.2
```

### Verification

---

#### Note

As we saw in the previous section where the virtual-link was added between R1 and R6, traffic from area 2 into area 0 transited R6. This was due to the fact that R3's slow serial links have a much higher cost value than R6 has via R1. This can be verified from the below traceroute from SW3 to area 51.

```
Rack1SW3#traceroute 51.51.51.51
```

```
Type escape sequence to abort.
Tracing the route to 51.51.51.51
```

```
 1 155.1.79.7 0 msec 0 msec 4 msec
 2 155.1.67.6 0 msec 4 msec 0 msec
 3 155.1.146.1 4 msec 0 msec 4 msec
 4 155.1.0.5 28 msec 32 msec 28 msec
 5 155.1.0.2 28 msec 32 msec 28 msec
 6 192.10.1.254 32 msec * 32 msec
```



The goal of this section is to modify this path selection so that area 2 exits via R3, and uses the link through area 5. Recall though that inter-area routing can only occur through area 0, or in the previous specific design, via a shorter path to area 0 via the non-area 0 transit area. Based on this design however the area 5 link can never be used to reach area 51 regardless of the cost. This can be illustrated as follows.

```
Rack1R3#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R3(config)#interface Serial1/3
Rack1R3(config-if)#ip ospf cost 1
Rack1R3(config-if)#end
Rack1R3#

Rack1R3#show ip route 51.51.51.51
Routing entry for 51.51.51.51/32
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric
  69073
  Last update from 155.1.0.5 on Serial1/0, 00:00:14 ago
  Routing Descriptor Blocks:
  * 155.1.0.5, from 192.10.1.254, 00:00:14 ago, via Serial1/0
    Route metric is 20, traffic share count is 1

Rack1R3#traceroute 51.51.51.51

Type escape sequence to abort.
Tracing the route to 51.51.51.51

 1 155.1.0.5 28 msec 32 msec 28 msec
 2 155.1.0.2 64 msec 56 msec 56 msec
 3 192.10.1.254 60 msec * 60 msec
```

R3 sets the cost of the link to R2 via area 5 to 1. Even though the cost is lower through this link than through area 0 to R5, the direct route to R2 cannot be used because R3 does not know how to reach the ASBR for area 51 via area 5. This can be seen by recursing towards the external route 51.51.51.51/32 that comes from area 51.

```
Rack1R3#show ip ospf database external 51.51.51.51
```

```
OSPF Router with ID (150.1.3.3) (Process ID 1)
```

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

```
Routing Bit Set on this LSA
LS age: 342
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 51.51.51.51 (External Network Number )
Advertising Router: 192.10.1.254
LS Seq Number: 8000005D
Checksum: 0xE3C8
Length: 36
Network Mask: /32
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

R3 says the prefix 51.51.51.51 has a forward address of 0.0.0.0. This means that R3 needs to lookup the advertising routing for the external LSA, who in this case is 192.10.1.254.

```
Rack1R3#show ip ospf database asbr-summary 192.10.1.254

      OSPF Router with ID (150.1.3.3) (Process ID 1)

          Summary ASB Link States (Area 0)

Routing Bit Set on this LSA
LS age: 519
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(AS Boundary Router)
Link State ID: 192.10.1.254 (AS Boundary Router address)
Advertising Router: 150.1.2.2
LS Seq Number: 80000032
Checksum: 0x2B4C
Length: 28
Network Mask: /0
      TOS: 0  Metric: 300
```

<output omitted>

```
          Summary ASB Link States (Area 5)

Adv Router is not-reachable
LS age: 283
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(AS Boundary Router)
Link State ID: 192.10.1.254 (AS Boundary Router address)
Advertising Router: 150.1.2.2
LS Seq Number: 80000032
Checksum: 0x2B4C
Length: 28
Network Mask: /0
      TOS: 0  Metric: 300
```

<output omitted>

The above output indicates that R3 has multiple paths to the ASBR of area 51, however most are invalid. R3 can use the area 0 route towards R2 to reach 192.10.1.254, but it cannot use the area 5 route towards R2, as the *Adv Router is not-reachable* through this path. This output essentially means that area 5 cannot be used as transit to reach area 51, only area 0 can. Therefore the only way to force the traffic to use the point-to-point Serial link between R2 and R3 is to give R3 an area 0 route towards R2 via this link. This can be accomplished by configuring a tunnel interface over the link in area 0, or more simply, through a virtual-link.

```
Rack1R3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R3(config)#router ospf 1
Rack1R3(config-router)#area 5 virtual-link 150.1.2.2
Rack1R3(config-router)#end
Rack1R3#
```

```
Rack1R2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R2(config)#router ospf 1
Rack1R2(config-router)#area 5 virtual-link 150.1.3.3
Rack1R2(config-router)#end
Rack1R2#
```

Now that R3 has an area 0 route to R2 via the virtual-link through area 5, and the cost of the area 5 link is 1, R3 chooses this as the shortest path towards the ASBR. The final result is that R3 advertises a better cost towards area 51 into area 2 than R6 does.

```
Rack1SW3#show ip ospf database external 51.51.51.51
```

```
OSPF Router with ID (150.1.9.9) (Process ID 1)
```

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

```
Routing Bit Set on this LSA
LS age: 1592
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 51.51.51.51 (External Network Number )
Advertising Router: 192.10.1.254
LS Seq Number: 8000005D
Checksum: 0xE3C8
Length: 36
Network Mask: /32
Metric Type: 2 (Larger than any link state path)
TOS: 0
Metric: 20
Forward Address: 0.0.0.0
External Route Tag: 0
```

SW3 now says that the external prefix 51.51.51.51/32 is through the ASBR 192.10.1.254, which is the same output it would have shown before the changes on R3. What is different now, however, is how area 2 routes towards the ASBR.

```

Rack1SW3#show ip ospf database asbr-summary 192.10.1.254

      OSPF Router with ID (150.1.9.9) (Process ID 1)

          Summary ASB Link States (Area 2)

Routing Bit Set on this LSA
LS age: 140
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(AS Boundary Router)
Link State ID: 192.10.1.254 (AS Boundary Router address)
Advertising Router: 150.1.3.3
LS Seq Number: 80000039
Checksum: 0x1A53
Length: 28
Network Mask: /0
      TOS: 0  Metric: 301

LS age: 57
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(AS Boundary Router)
Link State ID: 192.10.1.254 (AS Boundary Router address)
Advertising Router: 150.1.6.6
LS Seq Number: 80000025
Checksum: 0x5C7C
Length: 28
Network Mask: /0
      TOS: 0  Metric: 6838

```

SW3 says the ASBR is reachable via the area 2 ABRs 150.1.3.3 (R3) and 150.1.6.6 (R6). R3's reported metric is 301, while R6's metric is 6838. Based on this SW3 chooses to use R3 as the exit point, and the traffic is routed as desired.

```

Rack1SW3#traceroute 51.51.51.51

Type escape sequence to abort.
Tracing the route to 51.51.51.51

 0 155.1.79.7 0 msec 0 msec 0 msec
 1 155.1.37.3 0 msec 0 msec 4 msec
 2 155.1.23.2 16 msec 12 msec 16 msec
 3 192.10.1.254 20 msec * 16 msec

```

## 6.15 OSPF Demand Circuit

- The point-to-point between R4 and R5 is a low-speed Serial link. Configure the OSPF demand circuit feature on this link to reduce periodic OSPF hello transmission and paranoid update flooding.

### Configuration

---

```
R4:
interface Serial0/1
 ip ospf demand-circuit
```

### Verification

---

#### Note

Per RFC 1793, *Extending OSPF to Support Demand Circuits*, “OSPF Hellos and the refresh of OSPF routing information are suppressed on demand circuits, allowing the underlying data-link connections to be closed when not carrying application traffic.” This feature allows low-speed and pay-per-use links, such as analog dial, ISDN, etc. to run OSPF without the need for periodic hellos and LSA flooding.

This feature is enabled with the interface level command **ip ospf demand-circuit**, and is negotiated as part of the neighbor adjacency establishment. The below output dictates the change in the adjacency between R4 and R5 once the demand circuit feature is enabled.

```
Rack1R4#show ip ospf interface Serial0/1
Serial0/1 is up, line protocol is up
  Internet Address 155.1.45.4/24, Area 0
  Process ID 1, Router ID 223.255.255.255, Network Type POINT_TO_POINT, Cost: 1
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  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 2/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 5 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.5.5
  Suppress hello for 0 neighbor(s)
```

```

Rack1R4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R4(config)#interface Serial0/1
Rack1R4(config-if)#ip ospf demand-circuit
Rack1R4(config-if)#end
Rack1R4#
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.5.5 on Serial0/1 from FULL to DOWN,
Neighbor Down: Interface down or detached
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.5.5 on Serial0/1 from LOADING to FULL,
Loading Done
%SYS-5-CONFIG_I: Configured from console by console

Rack1R4#show ip ospf interface Serial0/1
Serial0/1 is up, line protocol is up
  Internet Address 155.1.45.4/24, Area 0
  Process ID 1, Router ID 223.255.255.255, Network Type POINT_TO_POINT, Cost: 1
  Enabled by interface config, including secondary ip addresses
  Configured as demand circuit.
  Run as demand circuit.
  DoNotAge LSA allowed.
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  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 2/3, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 5 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.5.5  (Hello suppressed)
  Suppress hello for 1 neighbor(s)

Rack1R4#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address        Interface
150.1.5.5        0     FULL/ -         -           155.1.45.5    Serial0/1
150.1.5.5        0     FULL/ -         00:01:53   155.1.0.5     Serial0/0
150.1.1.1        1     FULL/BDR        00:00:35   155.1.146.1   FastEthernet0/1
150.1.6.6        255   FULL/DR         00:00:39   155.1.146.6   FastEthernet0/1

```

The above output indicates that the link between R4 and R5 is set to *Run as demand circuit*, that *DoNotAge LSA* is allowed, and that R4 is going to *Suppress hello for 1 neighbor*. This indicates that periodic hellos are not sent, and the “paranoid” flooding of LSAs is disabled. Normally once an LSA reaches an age of 30 minutes it must be re-flooded, regardless if the network is stable and has not changed. The lack of periodic hello exchange can also be seen in the `show ip ospf neighbor` output, as the Dead Time field for the adjacency to R5 is null.

## 6.16 OSPF Flooding Reduction

- Configure R5, SW2, and SW4 so that links in OSPF area 3 do not participate in periodic paranoid update LSA flooding.

### Configuration

---

```
R5:
interface FastEthernet0/0
  ip ospf flood-reduction
!
interface FastEthernet0/1
  ip ospf flood-reduction
```

```
SW2:
interface Vlan8
  ip ospf flood-reduction
!
interface Vlan58
  ip ospf flood-reduction
!
interface Port-channel1
  ip ospf flood-reduction
```

```
SW4:
interface Vlan10
  ip ospf flood-reduction
!
interface Port-channel1
  ip ospf flood-reduction
```

### Verification

---

#### Note

Per RFC 2328, *OSPF Version 2*, "an LSA's LS age is never incremented past the value MaxAge." When the Link State Age reaches MaxAge "the router must attempt to flush the LSA... by reflooding the MaxAge LSA just as if it was a newly originated LSA".

In Cisco's IOS implementation of OSPF the MaxAge value is 3600 seconds, or 30 minutes. This means that each LSA is reflooded after 30 minutes, regardless if the topology is stable. This periodic flooding behavior is commonly referred to as the "paranoid update". The `ip ospf flood-reduction` feature stops unnecessary LSA flooding by setting the DoNotAge (DNA) bit in the LSA, removing the requirement for the periodic refresh.



In the below output SW4 has the (DNA) field next to all LSAs learned from area 3, which indicates that they do not need to be periodically flooded.

```
Rack1SW4#show ip ospf database
```

```
OSPF Router with ID (223.255.255.255) (Process ID 1)
```

```
Router Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.5.5	150.1.5.5	2	(DNA) 0x80000048	0x0071E8	3
150.1.8.8	150.1.8.8	1	(DNA) 0x80000042	0x0088F6	4
223.255.255.255	223.255.255.255	32	0x8000003C	0x00BF37	3

```
Net Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.58.5	150.1.5.5	2	(DNA) 0x80000001	0x00482B
155.1.108.8	150.1.8.8	1	(DNA) 0x80000001	0x0055A4

```
Summary Net Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
150.1.1.0	150.1.5.5	13	(DNA) 0x80000037	0x00B450
150.1.2.0	150.1.5.5	1546	(DNA) 0x8000003A	0x00B361
150.1.3.0	150.1.5.5	1546	(DNA) 0x8000003A	0x00A86B
150.1.4.0	150.1.5.5	767	(DNA) 0x80000039	0x00B1DD
150.1.6.0	150.1.5.5	767	(DNA) 0x8000003A	0x005C03
150.1.7.0	150.1.5.5	13	(DNA) 0x8000002A	0x005F93
150.1.9.0	150.1.5.5	13	(DNA) 0x8000002A	0x00765C
155.1.0.1	150.1.5.5	13	(DNA) 0x80000037	0x006A96
155.1.0.2	150.1.5.5	1546	(DNA) 0x8000003A	0x006AA6
155.1.0.3	150.1.5.5	1546	(DNA) 0x8000003A	0x0060AF
155.1.0.4	150.1.5.5	767	(DNA) 0x80000039	0x006A21
155.1.0.5	150.1.5.5	13	(DNA) 0x80000038	0x006161
155.1.7.0	150.1.5.5	13	(DNA) 0x8000002A	0x00418F
155.1.9.0	150.1.5.5	13	(DNA) 0x8000002A	0x005759
155.1.13.0	150.1.5.5	1546	(DNA) 0x8000003B	0x00EA1A
155.1.23.0	150.1.5.5	1029	(DNA) 0x8000003B	0x008871
155.1.37.0	150.1.5.5	13	(DNA) 0x80000029	0x00CA07
155.1.45.0	150.1.5.5	767	(DNA) 0x8000003A	0x009FC1
155.1.67.0	150.1.5.5	13	(DNA) 0x8000002A	0x00AAE9
155.1.79.0	150.1.5.5	13	(DNA) 0x8000002A	0x002662
155.1.146.0	150.1.5.5	767	(DNA) 0x8000003A	0x0007C7
192.10.1.0	150.1.5.5	1546	(DNA) 0x8000003A	0x00E6CF

```
Summary ASB Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.10.1.254	150.1.5.5	1546	(DNA) 0x8000003A	0x00E2D3

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

Link ID	ADV Router	Age	Seq#	Checksum	Tag
51.51.51.51	192.10.1.254	587	(DNA) 0x80000063	0x00D7CE	0

## 6.17 OSPF Clear Text Authentication

- Configure clear-text OSPF authentication for all adjacencies in area 2 using the password CLEARKEY.
- SW1 should enable authentication on all interfaces in area 2 with one single command.
- R3, R6, and SW3 should only enable authentication on their links connecting to SW1.

### **Configuration**

---

```
R3:
interface FastEthernet0/0
 ip ospf authentication
 ip ospf authentication-key CLEARKEY
```

```
R6:
interface FastEthernet0/0.67
 ip ospf authentication
 ip ospf authentication-key CLEARKEY
```

```
SW1:
interface Vlan67
 ip ospf authentication-key CLEARKEY
!
interface Vlan79
 ip ospf authentication-key CLEARKEY
!
interface FastEthernet0/3
 ip ospf authentication-key CLEARKEY
!
router ospf 1
 area 2 authentication
```

```
SW3:
interface Vlan79
 ip ospf authentication
 ip ospf authentication-key CLEARKEY
```

## Verification

### Note

OSPF supports three types of authentication, type 0, or null authentication (no authentication), type 1, or clear text authentication, and type 2, or MD5 authentication. The type of authentication can be configured at the interface level with the `ip ospf authentication` command, or at the process level with the `area [id] authentication` command. The only difference between these commands is whether authentication is enabled on all interfaces in the area at the same time, or on a per link basis. In either case the password must still be configured at the interface level with the `ip ospf authentication-key` or `ip ospf message-digest-key` commands.

A successfully authenticated adjacency can be verified as follows.

```
Rack1SW3#show ip ospf interface Vlan79 | include authentication
Simple password authentication enabled
```

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.7.7	1	FULL/DR	00:00:30	155.1.79.7	Vlan79

### Pitfall

A failure in authentication can occur for two reasons, a mismatch in authentication type, or a mismatch in authentication key. An authentication type mismatch occurs when one neighbor is configured with clear text, while the other is running MD5, or one is running MD5, and the other is running null, etc. A key mismatch is simply when the routers are using different passwords. Failure in authentication type can be verified as follows.

```
Rack1SW3#debug ip ospf adj
OSPF adjacency events debugging is on

Rack1SW3#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1SW3(config)#interface Vlan79
Rack1SW3(config-if)#ip ospf authentication message-digest
Rack1SW3(config-if)#
OSPF: Send with youngest Key 0
OSPF: Rcv pkt from 155.1.79.7, Vlan79 : Mismatch Authentication type. Input
packet specified type 1, we use type 2
```

The above output from `debug ip ospf adj` indicates a *Mismatch Authentication type*, where the inbound OSPF packet is using type 1 and the local router is using type 2. This indicates that the local router is trying to do MD5 authentication and the remote router is using clear text.

```
Rack1SW3#debug ip ospf adj
OSPF adjacency events debugging is on

Rack1SW3#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1SW3(config)#interface Vlan79
Rack1SW3(config-if)#ip ospf authentication null
OSPF: Rcv pkt from 155.1.79.7, Vlan79 : Mismatch Authentication type. Input
packet specified type 1, we use type 0
```

The above output indicates that SW3 is not doing authentication, i.e. type 0, and the remote end is doing clear text authentication, i.e. type 1. A mismatch in the password itself can be seen as follows.

```
Rack1SW3#debug ip ospf adj
OSPF adjacency events debugging is on

Rack1SW3#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1SW3(config)#interface Vlan79
Rack1SW3(config-if)#ip ospf authentication-key WRONGKEY
Rack1SW3(config-if)#
OSPF: Rcv pkt from 155.1.79.7, Vlan79 : Mismatch Authentication Key - Clear
Text
```

## 6.18 OSPF MD5 Authentication

- Configure MD5 based OSPF authentication for all adjacencies in area 0, including the Virtual-Links, using the password MD5KEY.
- R1 should enable MD5 authentication on all interfaces in area 0 with a single command.
- All other devices in area 0 should enable MD5 authentication on a per interface basis.

### Configuration

---

```
R1:
interface Serial0/0
 ip ospf message-digest-key 1 md5 MD5KEY
!
router ospf 1
 area 0 authentication message-digest
 area 1 virtual-link 150.1.6.6 message-digest-key 1 md5 MD5KEY

R2:
interface Serial0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 MD5KEY
!
router ospf 1
 area 5 virtual-link 150.1.3.3 authentication message-digest
 area 5 virtual-link 150.1.3.3 message-digest-key 1 md5 MD5KEY

R3:
interface Serial1/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 MD5KEY
!
router ospf 1
 area 5 virtual-link 150.1.2.2 authentication message-digest
 area 5 virtual-link 150.1.2.2 message-digest-key 1 md5 MD5KEY

R4:
interface Serial0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 MD5KEY
!
interface Serial0/1
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 MD5KEY

R5:
interface Serial0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 MD5KEY
!
interface Serial0/1
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 MD5KEY
```

```
R6:
router ospf 1
 area 1 virtual-link 150.1.1.1 authentication message-digest
 area 1 virtual-link 150.1.1.1 message-digest-key 1 md5 MD5KEY
```

### Verification

#### Note

Like clear text authentication, MD5 authentication can be enabled on a per link basis with the interface level command **ip ospf authentication message-digest**, or for all links in the area with the process level command **area [id] authentication message-digest**. For successful MD5 authentication the authentication type, the password, and the key ID must match. Correctly configured authentication can be verified as follows.

```
Rack1R5#show ip ospf interface | include protocol|authentication|key
Serial0/0 is up, line protocol is up
  Message digest authentication enabled
  Youngest key id is 1
Serial0/1 is up, line protocol is up
  Message digest authentication enabled
  Youngest key id is 1
FastEthernet0/1 is up, line protocol is up
FastEthernet0/0 is up, line protocol is up
Loopback0 is up, line protocol is up
```

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.3.3	0	FULL/ -	00:01:31	155.1.0.3	Serial0/0
150.1.1.1	0	FULL/ -	00:01:31	155.1.0.1	Serial0/0
150.1.2.2	0	FULL/ -	00:01:31	155.1.0.2	Serial0/0
223.255.255.255	0	FULL/ -	00:01:31	155.1.0.4	Serial0/0
223.255.255.255	0	FULL/ -	-	155.1.45.4	Serial0/1
150.1.8.8	1	FULL/BDR	00:00:33	155.1.58.8	FastEthernet0/0

**☠ Pitfall**

Remember that a virtual-link is an interface in area 0. This means that if the `area 0 authentication [message-digest]` command is enabled, authentication is also enabled on the virtual-link. In this example MD5 authentication is enabled on the virtual-link on R1 with the `area 0 authentication message-digest` command, but the password must still be assigned on the virtual-link interface itself with the `area 1 virtual-link 150.1.6.6 message-digest-key 1 md5 MD5KEY` command. On R6 authentication is enabled at the virtual-link interface level with the `area 1 virtual-link 150.1.1.1 authentication message-digest` command, and the key is applied with the `area 1 virtual-link 150.1.1.1 message-digest-key 1 md5 MD5KEY` command. In some versions these two commands are combined automatically in the running config to the single statement `area 1 virtual-link 150.1.1.1 authentication message-digest message-digest-key 1 md5 MD5KEY`, however the result of either syntax is the same.

Once authentication is enabled on the virtual-link ensure to issue the `clear ip ospf process` command, because the virtual-link does not support periodic hellos. This means that if the authentication is wrong, the virtual-link interface will not immediately go down, but if there is a change in the topology it won't actually be propagated across the virtual-link.

Failures in MD5 authentication can occur because of a type mismatch, a password mismatch, or a key ID mismatch, as verified below.

```
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 ip ospf message-digest-key 1
Rack1R5(config-if)#ip ospf message-digest-key 1 md5 WRONGKEY
Rack1R5(config-if)#
OSPF: Send with youngest Key 1
OSPF: Rcv pkt from 155.1.0.3, Serial0/0 : Mismatch Authentication Key - Message
Digest Key 1
```

This output indicates a password mismatch.

```
Rack1R5(config-if)#no ip ospf message-digest-key 1
Rack1R5(config-if)#ip ospf message-digest-key 2 md5 MD5KEY
Rack1R5(config-if)#
OSPF: Rcv pkt from 155.1.0.3, Serial0/0 : Mismatch Authentication Key - No
message digest key 1 on interface
Rack1R5(config-if)#
OSPF: Send with youngest Key 2
```

This output indicates a mismatch in key ID. R5 is sending key 2, but R3 is sending key 1.

```
Rack1R5(config-if)#ip ospf authentication null
Rack1R5(config-if)#
OSPF: Rcv pkt from 155.1.0.3, Serial0/0 : Mismatch Authentication type. Input
packet specified type 2, we use type 0
```

This output indicates a mismatch in authentication type. R5 is using null authentication (type 0), and R3 is using MD5 authentication (type 2).



## 6.19 OSPF Null Authentication

- Modify SW1 so that if a new router is added to area 2 on VLAN 7 it does not require authentication.

### Configuration

```
SW1:
interface Vlan7
 ip ospf authentication null
```

### Verification

#### Note

The authentication type configured at the interface level overrides the authentication type configured at the process level. In this example SW1 disables authentication on Vlan7 by configuring type 0 (null) authentication at the interface level. Note that the other links in area 2 are not affected by this change.

```
Rack1SW1#show ip ospf interface | include authentication|protocol
Loopback0 is up, line protocol is up
  Simple password authentication enabled
FastEthernet0/3 is up, line protocol is up (connected)
  Simple password authentication enabled
Vlan79 is up, line protocol is up
  Simple password authentication enabled
Vlan67 is up, line protocol is up
  Simple password authentication enabled
Vlan7 is up, line protocol is up
  Simple password authentication enabled
```

```
Rack1SW1#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1SW1(config)#interface Vlan7
Rack1SW1(config-if)#ip ospf authentication null
Rack1SW1(config-if)#end
Rack1SW1#show ip ospf interface | include authentication|protocol
Loopback0 is up, line protocol is up
  Simple password authentication enabled
FastEthernet0/3 is up, line protocol is up (connected)
  Simple password authentication enabled
Vlan79 is up, line protocol is up
  Simple password authentication enabled
Vlan67 is up, line protocol is up
  Simple password authentication enabled
Vlan7 is up, line protocol is up
```

## 6.20 OSPF MD5 Authentication with Multiple Keys

- Enable MD5 authentication in OSPF area 1.
- R1 should authenticate R6 using the password R1R6KEY.
- R4 should authenticate R6 using the password R4R6KEY.

### Configuration

```
R1:
interface FastEthernet0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 16 md5 R1R6KEY

R4:
interface FastEthernet0/1
 ip ospf authentication message-digest
 ip ospf message-digest-key 46 md5 R4R6KEY

R6:
interface FastEthernet0/0.146
 ip ospf authentication message-digest
 ip ospf message-digest-key 16 md5 R1R6KEY
 ip ospf message-digest-key 46 md5 R4R6KEY
```

### Verification

#### Note

OSPF supports multiple MD5 keys applied to a single interface in order to allow for key rotation. The normal application of this would be that all neighbors are configured with key 1. Key 2 is then added on all neighbors, and key 1 is removed. Since both keys are temporarily sent, there is no loss of adjacency while the old key is being removed. In this design multiple keys are used to authenticate different neighbors on the same interface.

```
Rack1R6#debug ip ospf adj
OSPF adjacency events debugging is on
Rack1R6#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R6(config)#interface FastEthernet0/0.146
Rack1R6(config-subif)#ip ospf authentication message-digest
OSPF: Send with youngest Key 0
OSPF: Rcv pkt from 155.1.146.4, FastEthernet0/0.146 : Mismatch Authentication
Key - No message digest key 46 on interface
OSPF: Rcv pkt from 155.1.146.1, FastEthernet0/0.146 : Mismatch Authentication
Key - No message digest key 16 on interface
```

R6 has MD5 authentication enabled, but no keys configured. We can see that it receives key 46 from R4 and key 16 from R1.

```
Rack1R6(config-subif)#do undebug all
All possible debugging has been turned off
Rack1R6(config-subif)# ip ospf message-digest-key 16 md5 R1R6KEY
Rack1R6(config-subif)# ip ospf message-digest-key 46 md5 R4R6KEY
Rack1R6(config-subif)#end
Rack1R6#
%OSPF-5-ADJCHG: Process 1, Nbr 223.255.255.255 on FastEthernet0/0.146 from
LOADING to FULL, Loading Done
%OSPF-5-ADJCHG: Process 1, Nbr 150.1.1.1 on FastEthernet0/0.146 from LOADING to
FULL, Loading Done
```

Once both keys are configured adjacency is established.

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.1.1	0	FULL/ -	-	155.1.146.1	OSPF_VL0
150.1.1.1	1	FULL/BDR	00:00:34	155.1.146.1	FastEthernet0/0.146
223.255.255.255	0	FULL/DROTHER	00:00:34	155.1.146.4	FastEthernet0/0.146
150.1.7.7	1	FULL/BDR	00:00:32	155.1.67.7	FastEthernet0/0.67

R6 sends both key 16 and 46.

```
Rack1R6#debug ip ospf adj
OSPF adjacency events debugging is on
OSPF: Send with key 16
OSPF: Send with key 46
```

## 6.21 OSPF Internal Summarization

- Configure R5 to advertise a summary route for the VLAN 8 and 10 prefixes as they are sent into area 0.
- This summary should be as specific as possible while still encompassing all addresses in both subnets.

### Configuration

---

```
R5:
router ospf 1
 area 3 range 155.1.8.0 255.255.252.0
```

### Verification

---

#### Note

Since devices in an OSPF area require the same copy of the database in order to compute correct SPF, filtering or summarization of routes in the database can only occur between areas or domains, not within an area. The below configuration illustrates how an Intra-Area Summary Network LSA (LSA 3) can be summarized as it is originated by an ABR.

Without any modification R1 sees two separate routes and LSAs to reach the networks 155.1.8.0/24 and 155.1.10.0/24. Since they are both originated by R5 (150.1.5.5), it implies that only R5 can modify R1's view of these. By configuring the `area 3 range 155.1.8.0 255.255.252.0` command R5 stops sending the specific LSAs from area 3 into area 0, and groups them into the single route 155.1.8.0/22.

```
Rack1R1#show ip route 155.1.8.8
Routing entry for 155.1.8.0/24
  Known via "ospf 1", distance 110, metric 3035, type inter area
  Last update from 155.1.0.5 on Serial0/0, 21:42:03 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.5.5, 21:42:03 ago, via Serial0/0
      Route metric is 3035, traffic share count is 1
```

```
Rack1R1#show ip route 155.1.10.10
Routing entry for 155.1.10.0/24
  Known via "ospf 1", distance 110, metric 3185, type inter area
  Last update from 155.1.0.5 on Serial0/0, 21:42:06 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.5.5, 21:42:06 ago, via Serial0/0
      Route metric is 3185, traffic share count is 1
```

```
Rack1R5#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R5(config)#router ospf 1
Rack1R5(config-router)#area 3 range 155.1.8.0 255.255.252.0
Rack1R5(config-router)#

Rack1R1#show ip route 155.1.8.8
Routing entry for 155.1.8.0/22
  Known via "ospf 1", distance 110, metric 3035, type inter area
  Last update from 155.1.0.5 on Serial0/0, 00:00:05 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.5.5, 00:00:05 ago, via Serial0/0
      Route metric is 3035, traffic share count is 1

Rack1R1#show ip route 155.1.10.10
Routing entry for 155.1.8.0/22
  Known via "ospf 1", distance 110, metric 3035, type inter area
  Last update from 155.1.0.5 on Serial0/0, 00:00:07 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.5.5, 00:00:07 ago, via Serial0/0
      Route metric is 3035, traffic share count is 1
```

R5 now originates a single LSA 3 for 155.1.8.0/22.

```
Rack1R1#show ip ospf database summary 155.1.8.0

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

          Summary Net Link States (Area 0)

Routing Bit Set on this LSA
LS age: 1636
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 155.1.8.0 (summary Network Number)
Advertising Router: 150.1.5.5
LS Seq Number: 80000029
Checksum: 0xA905
Length: 28
Network Mask: /22
          TOS: 0  Metric: 35

<output omitted>
```

## 6.22 OSPF Path Selection with Summarization

- Configure R3 to originate the summary route 155.X.6.0/23 in order to force traffic destined to VLAN 7 to transit R6.
- If R6's connection to VLAN 146 is down, traffic for VLAN 7 should transit R3.

### Configuration

---

```
R3:
router ospf 1
 area 2 range 155.1.6.0 255.255.254.0
```

### Verification

---

#### Note

When a router does a routing lookup on a destination it always chooses the longest match route for the path. This means that if the routes is trying to reach the destination 1.2.3.4 it will choose the route 1.2.3.0/24 over 1.2.0.0/18, or 1.2.0.0/18 over 0.0.0.0/0. This principle can be used for traffic engineering purposes by selectively summarizing prefixes into the IGP domain as seen below.

```
Rack1R5#show ip ospf database summary 155.1.7.0
```

```
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: 141
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 155.1.7.0 (summary Network Number)
Advertising Router: 150.1.3.3
LS Seq Number: 80000001
Checksum: 0xBEF1
Length: 28
Network Mask: /24
TOS: 0 Metric: 330
```

```
Routing Bit Set on this LSA
LS age: 12 (DoNotAge)
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 155.1.7.0 (summary Network Number)
Advertising Router: 150.1.6.6
LS Seq Number: 80000001
Checksum: 0x9713
Length: 28
Network Mask: /24
    TOS: 0 Metric: 330
```

<output omitted>

R5 learns the Summary Network LSA (LSA 3) route for 155.1.7.0/24 both from R3 and R6 with a metric of 330. Since R5's cost to R6 is higher than the cost to R3, R5 uses R3 as the exit point to this area 2 route.

```
Rack1R5#show ip route 155.1.7.7
Routing entry for 155.1.7.0/24
  Known via "ospf 1", distance 110, metric 3568, type inter area
  Last update from 155.1.0.3 on Serial0/0, 00:02:09 ago
  Routing Descriptor Blocks:
  * 155.1.0.3, from 150.1.3.3, 00:02:09 ago, via Serial0/0
    Route metric is 3568, traffic share count is 1
```

The path selection can be verified with a traceroute from area 3.

```
Rack1SW4#traceroute 155.1.7.7

Type escape sequence to abort.
Tracing the route to 155.1.7.7

 1 155.1.108.8 4 msec 4 msec 4 msec
 2 155.1.58.5 0 msec 4 msec 0 msec
 3 155.1.0.3 36 msec 36 msec 36 msec
 4 155.1.37.7 28 msec * 28 msec
```

By changing R3's origination of this route from 155.1.7.0/24 to 155.1.6.0/23, the longest match for all devices outside of area 2 is now 155.1.7.0/24 through R6.

```
Rack1R5#show ip ospf database summary 155.1.7.0
```

```
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: 12 (DoNotAge)
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 155.1.7.0 (summary Network Number)
Advertising Router: 150.1.6.6
LS Seq Number: 80000001
Checksum: 0x9713
Length: 28
Network Mask: /24
TOS: 0 Metric: 330
```

```
<output omitted>
```

```
Rack1R5#show ip route 155.1.7.7
```

```
Routing entry for 155.1.7.0/24
```

```
Known via "ospf 1", distance 110, metric 59223, type inter area
```

```
Last update from 155.1.0.1 on Serial0/0, 00:00:17 ago
```

```
Routing Descriptor Blocks:
```

```
* 155.1.0.1, from 150.1.6.6, 00:00:17 ago, via Serial0/0
```

```
Route metric is 59223, traffic share count is 1
```

```
Rack1SW4#traceroute 155.1.7.7
```

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

```
Tracing the route to 155.1.7.7
```

```
 1 155.1.108.8 4 msec 0 msec 4 msec
 2 155.1.58.5 0 msec 0 msec 4 msec
 3 155.1.0.1 28 msec 28 msec 28 msec
 4 155.1.146.6 28 msec 32 msec 28 msec
 5 155.1.67.7 32 msec * 28 msec
```



Due to the longest match routing principle, 155.1.6.0/23 can only be used to reach 155.1.7.7 if R6 withdraws 155.1.7.0/24.

```
Rack1R5#show ip ospf database summary 155.1.6.0
```

```
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: 81
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 155.1.6.0 (summary Network Number)
Advertising Router: 150.1.3.3
LS Seq Number: 80000001
Checksum: 0xC4ED
Length: 28
Network Mask: /23
TOS: 0 Metric: 330
```

```
Rack1R5#show ip route 155.1.7.7
```

```
Routing entry for 155.1.7.0/24
```

```
Known via "ospf 1", distance 110, metric 59223, type inter area
```

```
Last update from 155.1.0.1 on Serial0/0, 00:00:17 ago
```

```
Routing Descriptor Blocks:
```

```
* 155.1.0.1, from 150.1.6.6, 00:00:17 ago, via Serial0/0
```

```
Route metric is 59223, traffic share count is 1
```

```
Rack1R6#config t
```

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

```
Rack1R6(config)#interface FastEthernet0/0
```

```
Rack1R6(config-if)#shutdown
```

```
Rack1R6(config-if)#
```

```
Rack1R5#show ip route 155.1.7.7
```

```
Routing entry for 155.1.6.0/23
```

```
Known via "ospf 1", distance 110, metric 3568, type inter area
```

```
Last update from 155.1.0.3 on Serial0/0, 00:00:00 ago
```

```
Routing Descriptor Blocks:
```

```
* 155.1.0.3, from 150.1.3.3, 00:00:00 ago, via Serial0/0
```

```
Route metric is 3568, traffic share count is 1
```

```
Rack1SW4#traceroute 155.1.7.7
```

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

```
Tracing the route to 155.1.7.7
```

```
 1 155.1.108.8 0 msec 4 msec 4 msec
 2 155.1.58.5 4 msec 4 msec 4 msec
 3 155.1.0.3 36 msec 40 msec 36 msec
 4 155.1.37.7 40 msec * 36 msec
```

## 6.23 OSPF External Summarization

- Redistribute between RIPv2 and OSPF on R4.
- Redistribute between EIGRP and OSPF on R6.
- R4 should advertise two summary routes that have a cost of 50 into OSPF for the prefixes learned from BB3.
- R6 should advertise a single summary route into OSPF for the prefixes learned from BB1 that has a cost of 100, and includes the cost needed to reach R6.

### **Configuration**

---

R4:

```
router ospf 1
  router-id 150.1.4.4
  summary-address 30.0.0.0 255.252.0.0
  summary-address 31.0.0.0 255.252.0.0
  redistribute rip metric 50 subnets
!
router rip
  redistribute ospf 1 metric 1
```

R6:

```
router eigrp 10
  redistribute ospf 1 metric 10000 1000 255 1 1500
!
router ospf 1
  summary-address 200.0.0.0 255.255.252.0
  redistribute eigrp 10 metric 100 metric-type 1 subnets
```

SW4:

```
router ospf 1
  router-id 150.1.10.10
```

## Verification

### ☠ Pitfall

The first issue that should be evident with this section is that R4 and SW4 have the same OSPF router-id, due to a Loopback interface preconfigured with the IP address 223.255.255.255 from the initial configurations. Previously we saw a problem where two routers in the same area had the same router-id, and the result was that SPF throughout that area failed. In this case the routers are in different areas, but one of them is originating external LSA type 5 advertisements through redistribution.

Since SW4 sees 223.255.255.255 as the router-id of the originator, which matches its own router-id, it thinks it is receiving its own advertisement back. Since SW4 is not originating these routes it sends a withdraw message back. The result can be seen from the following log messages on R4. The specifics behind this problem can be seen in section 13.4, *Receiving self-originated LSAs, of RFC 2328, OSPF Version 2.*

```
R4#
%OSPF-4-FLOOD_WAR: Process 1 re-originates LSA ID 204.12.1.0 type-5
adv-rtr 223.255.255.255 in area 0
%OSPF-4-FLOOD_WAR: Process 1 re-originates LSA ID 31.2.0.0 type-5 adv-
rtr 223.255.255.255 in area 0
%OSPF-4-FLOOD_WAR: Process 1 re-originates LSA ID 31.3.0.0 type-5 adv-
rtr 223.255.255.255 in area 0
%OSPF-4-FLOOD_WAR: Process 1 re-originates LSA ID 204.12.1.0 type-5
adv-rtr 223.255.255.255 in area 0
%OSPF-4-FLOOD_WAR: Process 1 re-originates LSA ID 31.2.0.0 type-5 adv-
rtr 223.255.255.255 in area 0
%OSPF-4-FLOOD_WAR: Process 1 re-originates LSA ID 31.3.0.0 type-5 adv-
rtr 223.255.255.255 in area 0
%OSPF-4-FLOOD_WAR: Process 1 re-originates LSA ID 31.1.0.0 type-5 adv-
rtr
<output omitted>
```

Additionally SW4 cannot install the external routes from R4 in the routing table or OSPF database.

```
Rack1SW4#show ip route ospf | include E1|E2
O E2 51.51.51.51 [110/20] via 155.1.108.8, 01:41:14, Port-channel1
O E2 200.0.0.0/24 [110/20] via 155.1.108.8, 01:41:14, Port-channel1
O E2 54.1.1.0 [110/20] via 155.1.108.8, 01:41:14, Port-channel1
O E2 200.0.1.0/24 [110/20] via 155.1.108.8, 01:41:14, Port-channel1
O E2 200.0.2.0/24 [110/20] via 155.1.108.8, 01:41:14, Port-channel1
O E2 200.0.3.0/24 [110/20] via 155.1.108.8, 01:41:14, Port-channel1
```

```
Rack1SW4#show ip ospf database | begin External
      Type-5 AS External Link States
```

Link ID	ADV Router	Age	Seq#	Checksum	Tag
51.51.51.51	192.10.1.254	26	(DNA) 0x80000001	0x009C6C	0
54.1.1.0	150.1.6.6	4	(DNA) 0x80000001	0x006E56	0
200.0.0.0	150.1.6.6	4	(DNA) 0x80000001	0x001420	0
200.0.1.0	150.1.6.6	4	(DNA) 0x80000001	0x00092A	0
200.0.2.0	150.1.6.6	4	(DNA) 0x80000001	0x00FD34	0
200.0.3.0	150.1.6.6	4	(DNA) 0x80000001	0x00F23E	0

To solve this problem R4 and SW4's router-ids should be set to be something unique.

```
Rack1R4#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R4(config)#router ospf 1
Rack1R4(config-router)#router-id 150.1.4.4
Reload or use "clear ip ospf process" command, for this to take effect
Rack1R4(config-router)#end
Rack1R4#clear ip ospf 1 process
Reset OSPF process? [no]: yes
Rack1R4#
```

```
Rack1SW4#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1SW4(config)#router ospf 1
Rack1SW4(config-router)#router-id 150.1.10.10
Reload or use "clear ip ospf process" command, for this to take effect
Rack1SW4(config-router)#end
Rack1SW4#clear ip ospf 1 process
Reset OSPF process? [no]: yes
Rack1SW4#
```

Now SW4 can install all redistributed routes.

```
Rack1SW4#show ip route ospf | include E1|E2
O E2 51.51.51.51 [110/20] via 155.1.108.8, 00:31:05, Port-channel1
O E2 204.12.1.0/24 [110/50] via 155.1.108.8, 00:31:05, Port-channel1
O E1 200.0.0.0/24 [110/59173] via 155.1.108.8, 00:00:14, Port-channel1
O E1 54.1.1.0 [110/59173] via 155.1.108.8, 00:17:37, Port-channel1
O E1 200.0.1.0/24 [110/59173] via 155.1.108.8, 00:00:19, Port-channel1
O E1 200.0.2.0/24 [110/59173] via 155.1.108.8, 00:00:19, Port-channel1
O E1 200.0.3.0/24 [110/59173] via 155.1.108.8, 00:00:19, Port-channel1
O E2 31.3.0.0 [110/50] via 155.1.108.8, 00:00:00, Port-channel1
O E2 31.2.0.0 [110/50] via 155.1.108.8, 00:00:00, Port-channel1
O E2 31.1.0.0 [110/50] via 155.1.108.8, 00:00:00, Port-channel1
O E2 30.2.0.0 [110/50] via 155.1.108.8, 00:00:10, Port-channel1
O E2 30.3.0.0 [110/50] via 155.1.108.8, 00:00:10, Port-channel1
O E2 30.0.0.0 [110/50] via 155.1.108.8, 00:00:05, Port-channel1
O E2 30.1.0.0 [110/50] via 155.1.108.8, 00:00:10, Port-channel1
```

 **Note**

External OSPF summarization is configured at the redistribution point between routing domains with the **summary-address** command. These summaries inherit their attributes from the subnets that make them up. For example a summary comprised of External Type-1 routes will result in an External Type-1 summary. This means that on R6 in this configuration the **metric-type 1** command is set at the time of redistribution instead of on the summary itself.

External Type-2 OSPF routes, which are the default, do not install the end-to-end metric in the routing table. Instead only the metric that was reported via the ASBR is installed. The actual routing path is determined by the addition of the reported metric and the metric towards the ASBR, which is called the forward metric.

```
Rack1SW4#show ip ospf database external 30.0.0.0
```

```
OSPF Router with ID (150.1.10.10) (Process ID 1)
```

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

```
Routing Bit Set on this LSA
LS age: 3 (DoNotAge)
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 30.0.0.0 (External Network Number )
Advertising Router: 150.1.4.4
LS Seq Number: 80000005
Checksum: 0xF1D1
Length: 36
Network Mask: /14
Metric Type: 2 (Larger than any link state path)
TOS: 0
Metric: 50
Forward Address: 0.0.0.0
External Route Tag: 0
```

In the above output SW4 sees the summary 30.0.0.0/14 as an External Type-2 route originated by the ASBR 150.1.4.4. The *Forward Address: 0.0.0.0* field means that SW4 must now compute the metric towards the advertising router, R4, and install this metric in the routing table as the forward metric. Specifically this is calculated as follows.

```
Rack1SW4#show ip ospf database router 150.1.10.10 self-originate
```

```
OSPF Router with ID (150.1.10.10) (Process ID 1)
```

```
Router Link States (Area 3)
```

```
LS age: 1928
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.10.10
Advertising Router: 150.1.10.10
LS Seq Number: 80000004
Checksum: 0x1384
Length: 60
Number of Links: 3
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.108.8
(Link Data) Router Interface address: 155.1.108.10
Number of TOS metrics: 0
TOS 0 Metrics: 150
```

```
<output omitted>
```

**SW4's metric to the DR on the link to SW2 is 150.**

```
Rack1SW4#show ip ospf database router 150.1.8.8
```

```
OSPF Router with ID (150.1.10.10) (Process ID 1)
```

```
Router Link States (Area 3)
```

```
LS age: 1 (DoNotAge)
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.8.8
Advertising Router: 150.1.8.8
LS Seq Number: 8000002F
Checksum: 0xF09E
Length: 72
Number of Links: 4
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.108.8
(Link Data) Router Interface address: 155.1.108.8
Number of TOS metrics: 0
TOS 0 Metrics: 150
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.58.8
(Link Data) Router Interface address: 155.1.58.8
Number of TOS metrics: 0
TOS 0 Metrics: 30
```

```
<output omitted>
```

**SW2's metric to the DR on the link to R5 is 30.**

```
Rack1R5#show ip ospf database router 150.1.5.5 self-originate
```

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

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

```
LS age: 1743
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.5.5
Advertising Router: 150.1.5.5
LS Seq Number: 80000041
Checksum: 0x2596
Length: 108
Area Border Router
Number of Links: 7
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.4.4
(Link Data) Router Interface address: 155.1.45.5
Number of TOS metrics: 0
TOS 0 Metrics: 10000
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 155.1.45.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 10000
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.3.3
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 3238
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.1.1
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 58593
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.2.2
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 3238
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.4.4
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 39062
```

```
<output omitted>
```

R5's metric to R4 via the point-to-point Serial link is 10000. Therefore SW4's total metric towards R4, the forward metric, is 10180.

The routing table output for the specific routing lookup should now show a metric of 50 to the destination, but a forward metric of 10180 towards the ASBR. This forward metric is what is used to determine the path towards the exit point.

```
Rack1SW4#show ip route | include E1|E2
      E1 - OSPF external type 1, E2 - OSPF external type 2
O E2   51.51.51.51 [110/20] via 155.1.108.8, 00:27:44, Port-channel1
O E2 204.12.1.0/24 [110/50] via 155.1.108.8, 00:27:44, Port-channel1
O E1   54.1.1.0 [110/59173] via 155.1.108.8, 00:14:17, Port-channel1
O E2   31.0.0.0 [110/50] via 155.1.108.8, 00:27:44, Port-channel1
O E2   30.0.0.0 [110/50] via 155.1.108.8, 00:27:44, Port-channel1
O E1 200.0.0.0/22 [110/59173] via 155.1.108.8, 00:14:17, Port-channel1

Rack1SW4#show ip route 30.0.0.1
Routing entry for 30.0.0.0/14
  Known via "ospf 1", distance 110, metric 50, type extern 2, forward
metric 10180
  Last update from 155.1.108.8 on Port-channel1, 00:27:49 ago
  Routing Descriptor Blocks:
  * 155.1.108.8, from 150.1.4.4, 00:27:49 ago, via Port-channel1
    Route metric is 50, traffic share count is 1
```

From the calculation of the forward metric we can predict that the traffic flow from SW4 to the final destination, 30.0.0.1, will be first sent to SW2, then to R5, over the point-to-point Serial link to R4, and finally out to the RIP domain via BB3.

```
Rack1SW4#traceroute 30.0.0.1

Type escape sequence to abort.
Tracing the route to 30.0.0.1

 1 155.1.108.8 4 msec 4 msec 4 msec
 2 155.1.58.5 0 msec 4 msec 0 msec
 3 155.1.45.4 16 msec 16 msec 12 msec
 4 204.12.1.254 28 msec * 24 msec
```



The calculation for External Type-1 OSPF routes does not distinguish in the routing table between the metric reported by the ASBR, and the metric to the ASBR via the forward metric. Instead External Type-1 routes represent the metric as one cumulative value of the reported metric and the metric to the ASBR.

```
Rack1SW4#show ip route ospf | include E1
O E1 54.1.1.0 [110/59173] via 155.1.108.8, 00:56:00, Port-channell
O E1 200.0.0.0/22 [110/59173] via 155.1.108.8, 00:02:03, Port-channell
```

```
Rack1SW4#show ip route 200.0.0.1
Routing entry for 200.0.0.0/22, supernet
  Known via "ospf 1", distance 110, metric 59173, type extern 1
  Last update from 155.1.108.8 on Port-channell, 00:03:53 ago
  Routing Descriptor Blocks:
  * 155.1.108.8, from 150.1.6.6, 00:03:53 ago, via Port-channell
    Route metric is 59173, traffic share count is 1
```

The metric for this route is calculated as follows.

```
Rack1SW4#show ip ospf database external 200.0.0.0

      OSPF Router with ID (150.1.10.10) (Process ID 1)

      Type-5 AS External Link States

Routing Bit Set on this LSA
LS age: 4 (DoNotAge)
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 200.0.0.0 (External Network Number )
Advertising Router: 150.1.6.6
LS Seq Number: 8000000E
Checksum: 0x8ACF
Length: 36
Network Mask: /22
  Metric Type: 1 (Comparable directly to link state metric)
  TOS: 0
  Metric: 100
  Forward Address: 0.0.0.0
  External Route Tag: 0
```

SW4 sees that the ASBR 150.1.6.6 (R6) originated the route with a metric of 100. Now SW4 must find the closest path to forward towards the ASBR.

```
Rack1SW4#show ip ospf database router 150.1.6.6
```

```
OSPF Router with ID (150.1.10.10) (Process ID 1)
```

SW4 tries to find an Intra-Area path to the ASBR 150.1.6.6. Since R6 is not in the same area as SW4, SW4 must find which ABR it should use to exit towards R6.

```
Rack1SW4#show ip ospf database asbr-summary 150.1.6.6
```

```
OSPF Router with ID (150.1.10.10) (Process ID 1)
```

```
Summary ASB Link States (Area 3)
```

```
Routing Bit Set on this LSA
LS age: 2 (DoNotAge)
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(AS Boundary Router)
Link State ID: 150.1.6.6 (AS Boundary Router address)
Advertising Router: 150.1.5.5
LS Seq Number: 80000003
Checksum: 0x53AD
Length: 28
Network Mask: /0
TOS: 0 Metric: 58893
```

SW4's Intra-Area exit point towards R6 is R5. R5 has a metric of 58893 to reach R6. SW4 now runs SPF to reach R5.

```
Rack1SW4#show ip ospf database router 150.1.10.10 self-originate
```

```
OSPF Router with ID (150.1.10.10) (Process ID 1)
```

```
Router Link States (Area 3)
```

```
LS age: 1111
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.10.10
Advertising Router: 150.1.10.10
LS Seq Number: 80000006
Checksum: 0xF86
Length: 60
Number of Links: 3

Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.108.8
(Link Data) Router Interface address: 155.1.108.10
Number of TOS metrics: 0
TOS 0 Metrics: 150
```

```
<output omitted>
```

SW4's metric to the DR for the link to SW2 is 150.

```
Rack1SW2#show ip ospf database router 150.1.8.8 self-originate
```

```
OSPF Router with ID (150.1.8.8) (Process ID 1)
```

```
Router Link States (Area 3)
```

```
LS age: 1170
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.8.8
Advertising Router: 150.1.8.8
LS Seq Number: 80000035
Checksum: 0xE4A4
Length: 72
Number of Links: 4
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.108.8
(Link Data) Router Interface address: 155.1.108.8
Number of TOS metrics: 0
TOS 0 Metrics: 150
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.58.8
(Link Data) Router Interface address: 155.1.58.8
Number of TOS metrics: 0
TOS 0 Metrics: 30
```

```
<output omitted>
```

SW2's metric to the DR on the link to R5 is 30. This means that SW4's metric to R5 is 150 + 30, R5's metric to R6 is 58893, and R6's metric to the external domain is 100, for a total of 59173. This is the value seen in the routing table as the final metric.

```
Rack1SW4#show ip route 200.0.0.1
Routing entry for 200.0.0.0/22, supernet
  Known via "ospf 1", distance 110, metric 59173, type extern 1
  Last update from 155.1.108.8 on Port-channel1, 00:15:10 ago
  Routing Descriptor Blocks:
  * 155.1.108.8, from 150.1.6.6, 00:15:10 ago, via Port-channel1
    Route metric is 59173, traffic share count is 1
```

We can predict the traffic flow from SW4 to this destination by examining the Intra-Area routing of area 0. Since R6 is the ASBR originating the route, and R6 is in area 0, R5 will compute Intra-Area SPF to reach R6. Since R6 connects to area 0 via the virtual-link to R1, and R1 connects to R5 via the Frame Relay link, traffic must flow from SW4 to SW2, to R5, to R1 over the Frame Relay link, to R6 over VLAN 146, and finally out to BB1. Even if R5 has a lower cost to R6 via the point-to-point link to R4, and from R4 to VLAN 146, this path cannot be considered because it is Inter-Area, versus the Intra-Area path through the virtual-link. The result is as follows.

```
Rack1SW4#traceroute 200.0.0.1
```

```
Type escape sequence to abort.  
Tracing the route to 200.0.0.1
```

```
 1 155.1.108.8 4 msec 4 msec 0 msec  
 2 155.1.58.5 0 msec 4 msec 0 msec  
 3 155.1.0.1 28 msec 32 msec 28 msec  
 4 155.1.146.6 28 msec 28 msec 28 msec  
 5 54.1.1.254 52 msec * 48 msec
```

## 6.24 OSPF Stub Areas

- Configure OSPF area 3 so that R5 filters External routes out as they are sent from area 0 to area 3.
- Devices in area 3 should still have reachability to routes external to the OSPF domain.

### Configuration

---

```
R5:  
router ospf 1  
  area 3 stub
```

```
SW2:  
router ospf 1  
  area 3 stub
```

```
SW4:  
router ospf 1  
  area 3 stub
```

### Verification

---

#### Note

OSPF stub areas types are used to filter information out of the OSPF database based on the role of the LSA. The stub flag is part of the OSPF adjacency formation, which implies that all devices in an area must agree on the parameter in order for adjacencies to establish. The four stub types that IOS supports are stub areas, totally stubby areas, not-so-stubby areas (NSSA), and not-so-totally-stubby areas. The first option, the stub area, is used to remove Type-5 External link states from the database, and replace them with a default route. The logic behind this feature stems from how external lookups between areas occur in OSPF.

When an OSPF router redistributes a route into the domain, it originates a Type-5 External LSA representing the route and its attributes. Inside this LSA the originating router sets the *advertising router* field to its local router-id, and generally, the *forward address* field to 0.0.0.0.

When an OSPF router in the same area as the originator looks up the Type-5 LSA, it looks at the forward address. If the forward address is set to 0.0.0.0, it means that the traffic should be sent towards the advertising router to reach the destination. To find out how to reach the advertising router, the advertising router's Type-1 Router LSA is consulted, and intra-area SPF is performed. This is similar to inter-area routing logic, because the router doing the lookup does not compute SPF to the final destination, only the intermediary advertising router. For external routing between areas, the logic is modified slightly.

When an Area Border Router receives a Type-5 External LSA from a device in its own area, and passes it into a different area, a Type-4 ASBR Summary LSA is generated. The Type-4 LSA tells devices in the new area how to forward towards the ASBR, which in turn tells them how to forward towards the external route. For example take the following situation in this topology. R4 redistributes the RIP route 30.1.0.0/24 into OSPF, originating a Type-5 External LSA.

```
Rack1R4#show ip ospf database external 30.1.0.0

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

      Type-5 AS External Link States

LS age: 1651
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 30.1.0.0 (External Network Number )
Advertising Router: 150.1.4.4
LS Seq Number: 80000005
Checksum: 0xC419
Length: 36
Network Mask: /16
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

When R5 wants to reach the destination 30.1.0.0, it sees that the forward address is 0.0.0.0, and the advertising router is 150.1.4.4. R5 now does a Type-1 Router LSA lookup on 150.1.4.4 (R4).

```
Rack1R5#show ip ospf database router 150.1.4.4

<snip>

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.5.5
(Link Data) Router Interface address: 155.1.45.4
Number of TOS metrics: 0
    TOS 0 Metrics: 64

<snip>

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.5.5
(Link Data) Router Interface address: 155.1.0.4
Number of TOS metrics: 0
    TOS 0 Metrics: 64
```

From this R5 knows that R4 is directly adjacent via two paths, both with a metric of 64.

```
Rack1R5#show ip route 30.1.0.0
Routing entry for 30.1.0.0/16
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 64
  Last update from 155.1.0.4 on Serial0/0, 00:24:16 ago
  Routing Descriptor Blocks:
    * 155.1.45.4, from 150.1.4.4, 00:24:16 ago, via Serial0/1
      Route metric is 20, traffic share count is 1
    155.1.0.4, from 150.1.4.4, 00:24:16 ago, via Serial0/0
      Route metric is 20, traffic share count is 1
```

R5 installs both paths in the routing table with a metric of 20 from the Type-5 External LSA, and a forward metric of 64 to reach R4. If this route were redistributed as Type-1 External, as opposed to Type-2 External, the total metric would be 84 (the advertised metric plus the forward metric). Now R5 sends the Type-5 LSA from area 0 into area 3 to SW2.

```
Rack1SW2#show ip ospf database external 30.1.0.0

      OSPF Router with ID (150.1.8.8) (Process ID 1)

          Type-5 AS External Link States

Routing Bit Set on this LSA
LS age: 902
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 30.1.0.0 (External Network Number )
Advertising Router: 150.1.4.4
LS Seq Number: 80000006
Checksum: 0xC21A
Length: 36
Network Mask: /16
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

SW2, like R5, sees the Forward Address for the route as 0.0.0.0, meaning a lookup on 150.1.4.4 must be performed. The difference here however, is that when SW2 looks for a Type-1 Router LSA, none is found.

```
Rack1SW2#show ip ospf database router 150.1.4.4

      OSPF Router with ID (150.1.8.8) (Process ID 1)
```

This essentially means that the advertising router (R4) for the external route is not in the same area as SW2. SW2 now checks to see which ABRs are advertising reachability information about R4.

```
Rack1SW2#show ip ospf database asbr-summary 150.1.4.4
```

```
OSPF Router with ID (150.1.8.8) (Process ID 1)
```

```
Summary ASB Link States (Area 3)
```

```
Routing Bit Set on this LSA
LS age: 1830
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(AS Boundary Router)
Link State ID: 150.1.4.4 (AS Boundary Router address)
Advertising Router: 150.1.5.5
LS Seq Number: 80000001
Checksum: 0x6357
Length: 28
Network Mask: /0
TOS: 0 Metric: 64
```

SW2 sees that 150.1.4.4 (R4) is known via the advertising router 150.1.5.5 (R5) with a metric of 64. An intra-area lookup is now performed on 150.1.5.5.

```
Rack1SW2#show ip ospf database router 150.1.5.5
```

```
OSPF Router with ID (150.1.8.8) (Process ID 1)
```

```
Router Link States (Area 3)
```

```
Routing Bit Set on this LSA
LS age: 12
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.5.5
Advertising Router: 150.1.5.5
LS Seq Number: 8000000C
Checksum: 0xE8DA
Length: 60
Area Border Router
Number of Links: 3

Link connected to: a Stub Network
(Link ID) Network/subnet number: 155.1.5.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 1

Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.58.8
(Link Data) Router Interface address: 155.1.58.5
Number of TOS metrics: 0
TOS 0 Metrics: 1
```

```
<snip>
```



The recursion process continues until SW2 sees that it is adjacent with the DR 155.1.58.8 as well, meaning that R5 is reachable out Vlan58. Based on this SW2 knows that packets for 30.1.0.0 should go towards R5, who sends them towards R4, who send them to the final destination.

The key point about this external lookup between areas, however, is that for all external destinations outside of area 3, SW2 will see that R5 is the ABR that it must transit through. This redundant information can be seen in the database view of area 3 as follows.

```
Rack1SW4#show ip ospf database | begin Type-5
```

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

Link ID	ADV Router	Age	Seq#	Checksum	Tag
30.0.0.0	150.1.4.4	1864	0x80000001	0x00D80A	0
30.1.0.0	150.1.4.4	1864	0x80000001	0x00CC15	0
30.2.0.0	150.1.4.4	1864	0x80000001	0x00C020	0
30.3.0.0	150.1.4.4	1864	0x80000001	0x00B42B	0
31.0.0.0	150.1.4.4	1864	0x80000001	0x00CB16	0
31.1.0.0	150.1.4.4	1864	0x80000001	0x00BF21	0
31.2.0.0	150.1.4.4	1864	0x80000001	0x00B32C	0
31.3.0.0	150.1.4.4	1864	0x80000001	0x00A737	0
51.51.51.51	192.10.1.254	1226	0x80000002	0x009A6D	0
54.1.1.0	150.1.6.6	1878	0x80000001	0x006E56	0
200.0.0.0	150.1.6.6	1885	0x80000001	0x001420	0
200.0.1.0	150.1.6.6	1885	0x80000001	0x00092A	0
200.0.2.0	150.1.6.6	1885	0x80000001	0x00FD34	0
200.0.3.0	150.1.6.6	1885	0x80000001	0x00F23E	0
204.12.1.0	150.1.4.4	1870	0x80000001	0x005EC8	0

SW4 sees multiple Type-5 LSAs, most reachable via R4 and R6. Recursion for both R4 and R6 points to R5, as R5 is the only ABR servicing area 3. In a design such as this, stub areas can be used to optimize the OSPF database by replacing the redundant Type-5 External and Type-4 ASBR routing information with default information. In this particular case the changes seen are as follows.

```
Rack1R5#conf t
```

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

```
Rack1R5(config)#router ospf 1
```

```
Rack1R5(config-router)#area 3 stub
```

```
Rack1R5(config-router)#end
```

```
Rack1R5#
```

```
Rack1SW2#conf t
```

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

```
Rack1SW2(config)#router ospf 1
```

```
Rack1SW2(config-router)#area 3 stub
```

```
Rack1SW2(config-router)#end
```

```
Rack1SW2#
```

```
Rack1SW4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1SW4(config)#router ospf 1
Rack1SW4(config-router)#area 3 stub
Rack1SW4(config-router)#end
Rack1SW4#
```

With area 3 configured as a stub area, Type-5 External LSAs and Type-4 ASBR Summary LSAs no longer exist in the area 3 database.

```
Rack1SW4#show ip ospf database external 30.1.0.0

      OSPF Router with ID (150.1.10.10) (Process ID 1)

Rack1SW4#show ip ospf database asbr-summary 150.1.4.4

      OSPF Router with ID (150.1.10.10) (Process ID 1)
```

This implies that SW4 no longer has a specific route to 30.1.0.0, due to the deletion of the Type-4 and Type-5 LSAs.

```
Rack1SW4#show ip route 30.1.0.0
% Network not in table
```

What has been added to the database, however, is a default route as a Type-3 Summary LSA via the ABR.

```
Rack1SW4#show ip ospf database summary 0.0.0.0

      OSPF Router with ID (150.1.10.10) (Process ID 1)

      Summary Net Link States (Area 3)

Routing Bit Set on this LSA
LS age: 580
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 0.0.0.0 (summary Network Number)
Advertising Router: 150.1.5.5
LS Seq Number: 80000002
Checksum: 0x1B80
Length: 28
Network Mask: /0
      TOS: 0  Metric: 1
```

The result of this new default route is that area 3 maintains connectivity to the external routes by using the default route, but the size of the OSPF database and the routing table is much smaller.

```
Rack1SW4#show ip route | include _0.0.0.0
```

```
Gateway of last resort is 155.1.108.8 to network 0.0.0.0
```

```
O*IA 0.0.0.0/0 [110/3] via 155.1.108.8, 00:32:01, Port-channel1
```

```
Rack1SW4#traceroute 30.1.0.1
```

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

```
Tracing the route to 30.1.0.1
```

```
 1 155.1.108.8 4 msec 0 msec 4 msec
 2 155.1.58.5 0 msec 4 msec 0 msec
 3 155.1.45.4 16 msec
   155.1.0.4 24 msec
   155.1.45.4 20 msec
 4 204.12.1.254 24 msec * 24 msec
```

## 6.25 OSPF Totally Stubby Areas

- Configure OSPF area 3 so that R5 filters Inter-Area and External routes as they are sent from area 0 to area 3.
- Devices in area 3 should still have reachability to routes external to the OSPF domain.

### Configuration

```
R5:
router ospf 1
 area 3 stub no-summary
```

### Verification

#### Note

In the previous case with we saw that with area 3 converted to a stub area, the size of the routing table and OSPF database was reduced with no negative impact on connectivity. Specifically SW4's view of the topology was as follows.

```
Rack1SW4#show ip ospf database
```

```
OSPF Router with ID (150.1.10.10) (Process ID 1)
```

```
Router Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.5.5	150.1.5.5	807	0x8000000F	0x0001C1	3
150.1.8.8	150.1.8.8	464	0x80000014	0x00165C	4
150.1.10.10	150.1.10.10	819	0x8000000F	0x00969E	3

```
Net Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.58.8	150.1.8.8	715	0x8000000D	0x00F06F
155.1.108.10	150.1.10.10	819	0x80000009	0x001A08

```
Summary Net Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	150.1.5.5	41	0x80000003	0x001981
150.1.1.1	150.1.5.5	12	0x80000001	0x00D8E9
150.1.2.2	150.1.5.5	12	0x80000001	0x00C3FC
150.1.3.3	150.1.5.5	12	0x80000001	0x00AE10
150.1.4.4	150.1.5.5	12	0x80000001	0x009923
150.1.6.6	150.1.5.5	12	0x80000001	0x00793E
150.1.7.7	150.1.5.5	12	0x80000001	0x006451
150.1.9.9	150.1.5.5	12	0x80000001	0x00446C
155.1.0.1	150.1.5.5	13	0x80000001	0x009827
155.1.0.2	150.1.5.5	13	0x80000001	0x008E30
155.1.0.3	150.1.5.5	13	0x80000001	0x008439
155.1.0.4	150.1.5.5	13	0x80000001	0x007A42
155.1.0.5	150.1.5.5	13	0x80000001	0x00ED0E
155.1.7.0	150.1.5.5	13	0x80000001	0x00694E

155.1.9.0	150.1.5.5	13	0x80000001	0x005D57
155.1.13.0	150.1.5.5	13	0x80000001	0x0095DD
155.1.23.0	150.1.5.5	13	0x80000001	0x002742
155.1.37.0	150.1.5.5	13	0x80000001	0x001486
155.1.45.0	150.1.5.5	13	0x80000001	0x00B1E1
155.1.67.0	150.1.5.5	13	0x80000001	0x00D2A8
155.1.79.0	150.1.5.5	13	0x80000001	0x004E21
155.1.146.0	150.1.5.5	13	0x80000001	0x0060CC
192.10.1.0	150.1.5.5	13	0x80000001	0x00523E

The above output indicates that the intra-area information from Type-1 Router LSAs and Type-2 Network LSAs still exist, along with the inter-area Type-3 Summary LSA information, but Type-4 ASBR Summary LSAs and Type-5 External LSAs have been removed.

Recall that in the previous case for external routes we saw that every Type-4 ASBR Summary LSA inside of area 3 always recursed back to R5, since R5 was the only ABR connecting area 3 to area 0. By configuring area 3 as stub, this information was replaced with a default route which recursed to R5. The same connectivity resulted, but space was saved in the routing table and OSPF database. The next logical step in further optimizing the database then is to summarize the redundant Type-3 Summary LSAs that represent the inter-area routes. This is where configuring the area as totally stubby can be advantageous.

Where a stub area optimizes the database by removing external routes, and replacing it with a default route, a totally stubby area optimized the database further by removing the inter-area and external routes, replacing them both with a default route. This is accomplished by telling the area border router not to inject Type-3 Summary Network LSAs from area 0, hence the **no-summary** argument of the **stub** command.

Note that only the ABR(s) connecting the stub area to area 0 need the **no-summary** argument on the stub command, since they are the only devices that are allowed to originate Type-3 LSAs. While it won't "break" the OSPF design to add the command to other routers inside the totally stubby area, it is technically incorrect to configure the option this way. As long as all devices in the area agree on the stub flag in the first place, it is the ABR(s) duty to determine if the area is totally stubby or not.

To verify the operation of the totally stubby area, view the changes to the database on the area 3 routers.

```
Rack1SW4#show ip ospf database
```

```
OSPF Router with ID (150.1.10.10) (Process ID 1)
```

```
Router Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.5.5	150.1.5.5	168	0x80000010	0x00FEC2	3
150.1.8.8	150.1.8.8	1854	0x80000014	0x00165C	4
150.1.10.10	150.1.10.10	184	0x80000010	0x00949F	3

```
Net Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.58.8	150.1.8.8	86	0x8000000E	0x00EE70
155.1.108.10	150.1.10.10	185	0x8000000A	0x001809

```
Summary Net Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	150.1.5.5	170	0x80000004	0x001782

Note the major change that has occurred is that SW4 no longer has specific inter-area routing information listed under the *Summary Net Link States* field, which represents the Type-3 Summary LSAs. Since SW4 still has a default route via R5, connectivity is not affected by this change.

```
Rack1SW4#show ip route 155.1.9.9
```

```
% Subnet not in table
```

```
Rack1SW4#show ip route | include _0.0.0.0
```

```
Gateway of last resort is 155.1.108.8 to network 0.0.0.0
```

```
O*IA 0.0.0.0/0 [110/3] via 155.1.108.8, 00:05:36, Port-channel1
```

```
Rack1SW4#traceroute 155.1.9.9
```

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

```
Tracing the route to 155.1.9.9
```

```

 1 155.1.108.8 4 msec 0 msec 0 msec
 2 155.1.58.5 0 msec 4 msec 0 msec
 3 155.1.0.3 28 msec 32 msec 32 msec
 4 155.1.37.7 24 msec 24 msec 20 msec
 5 155.1.79.9 32 msec * 28 msec
```

## 6.26 OSPF Not-So-Stubby Areas

- Shutdown R6's Loopback0 interface and connection to VLAN 146.
- Configure OSPF area 2 so that R3 filters External routes out as they are sent from area 0 to area 2.
- R6 should still be allowed to redistribute EIGRP into OSPF.
- Note any reachability problems throughout the domain.

### Configuration

---

```
R3:
router ospf 1
 area 2 nssa

R6:
interface Loopback0
 shutdown
!
interface FastEthernet0/0.146
 shutdown
!
router ospf 1
 area 2 nssa

SW1:
router ospf 1
 area 2 nssa

SW3:
router ospf 1
 area 2 nssa
```

### Verification

---

#### Note

The *OSPF Not-So-Stubby Area (NSSA) Option*, as defined in RFC 3101, extends the functionality of a stub area to allow the importing of a subset of external routes into the area. Recall that with the stub area, Type-5 External LSA information is suppressed from entering the database, and is replaced with a default route originated by the ABR(s). Since all Type-5 LSAs are suppressed, this also implies that redistribution cannot occur within the area as well. This problem can be seen from the parser error generated when redistribution and stub areas are configured together, as follows.

```
Rack1SW4#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1SW4(config)#router ospf 1
Rack1SW4(config-router)#redistribute connected subnets
OSPF-4-ASBR_WITHOUT_VALID_AREA: Router is currently an ASBR while having only
one area which is a stub area
```

The OSPF NSSA option changes this behavior by allowing redistribution to occur *within* the stub area, while still blocking external routes from entering the area through the ABR(s). Specifically this is implemented through the introduction of a new link-state advertisement type, the Type-7 NSSA External LSA.

Routes that are redistributed directly into the NSSA are generated as Type-7 NSSA External LSAs. Like Type-5 External LSAs, two subtypes of Type-7 NSSA External LSAs exist, type 1 (N1) and type 2 (N2). N1, similar to E1, considers the metric that the ASBR reports into the OSPF domain along with the metric needed to reach the ASBR. N2, similar to E2, separates the metric into the flat value that the ASBR reports into the OSPF domain, which is installed in the routing table, and the value needed to reach the ASBR, known as the forward metric.

From the below output we can see that with the default redistribution values, R6 originates the Type-7 NSSA External LSAs as metric-type 2, with a metric value of 20. The detailed output from SW1's routing table indicates a metric of 20 reported in by R6, and a forward metric of 1000 – SW1's metric to reach R6.

```
Rack1R6#show run | section router ospf
router ospf 1
  router-id 150.1.6.6
  area 1 virtual-link 150.1.1.1
  area 2 nssa
  redistribute eigrp 10 subnets
```

```
Rack1SW1#show ip route | include N1|N2
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
O N2 200.0.0.0/24 [110/20] via 155.1.67.6, 07:10:50, Vlan67
O N2   54.1.1.0 [110/20] via 155.1.67.6, 07:10:50, Vlan67
O N2 200.0.1.0/24 [110/20] via 155.1.67.6, 07:10:50, Vlan67
O N2 200.0.2.0/24 [110/20] via 155.1.67.6, 07:10:50, Vlan67
O N2 200.0.3.0/24 [110/20] via 155.1.67.6, 07:10:50, Vlan67
```

```
Rack1SW1#show ip route 200.0.0.0
Routing entry for 200.0.0.0/24
  Known via "ospf 1", distance 110, metric 20, type NSSA extern 2, forward
metric 1000
  Last update from 155.1.67.6 on Vlan67, 07:10:58 ago
  Routing Descriptor Blocks:
  * 155.1.67.6, from 150.1.6.6, 07:10:58 ago, via Vlan67
    Route metric is 20, traffic share count is 1
```



When the Type-7 NSSA External LSA is received by the ABR, and is moved into area 0, the information contained in the Type-7 LSA is translated to a normal Type-5 External LSA. If multiple ABRs exist, only one of them performs the translation through an election process, which is discussed in depth in a later task. In this fashion OSPF devices outside of the NSSA do not know that the NSSA exists, which is analogous to how a Confederation works in BGP.

In this particular case R3 learns the Type-7 NSSA External LSAs that are originated by R6.

```
Rack1R3#show ip ospf database | begin Type-7
      Type-7 AS External Link States (Area 2)

Link ID      ADV Router   Age          Seq#         Checksum Tag
54.1.1.0     150.1.6.6   129         0x80000001  0x005480 0
200.0.0.0    150.1.6.6   129         0x80000001  0x00F94A 0
200.0.1.0    150.1.6.6   129         0x80000001  0x00EE54 0
200.0.2.0    150.1.6.6   129         0x80000001  0x00E35E 0
200.0.3.0    150.1.6.6   129         0x80000001  0x00D868 0
<snip>
```

The detailed information contained in the LSA can be viewed with the `show ip ospf database nssa-external [Link ID]` command.

```
Rack1R3#show ip ospf database nssa-external 200.0.0.0
      OSPF Router with ID (150.1.3.3) (Process ID 1)
      Type-7 AS External Link States (Area 2)

Routing Bit Set on this LSA
LS age: 312
Options: (No TOS-capability, Type 7/5 translation, DC)
LS Type: AS External Link
Link State ID: 200.0.0.0 (External Network Number )
Advertising Router: 150.1.6.6
LS Seq Number: 80000001
Checksum: 0xF94A
Length: 36
Network Mask: /24
  Metric Type: 2 (Larger than any link state path)
  TOS: 0
  Metric: 20
  Forward Address: 155.1.67.6
  External Route Tag: 0
```

Note that R3 receives the Type-7 NSSA External LSA with the forward address set to 155.1.67.6, which is R6's interface pointing into area 2. With the previous Type-5 external lookups we saw the forward address set to 0.0.0.0, which meant to route towards the advertising router to reach the final destination. In this case the forward address is non-zero, which causes the lookup to be performed towards 155.1.67.6. This is a subtle difference in the lookup, and in this case results in the same path selection even if the lookup occurred on the advertising router (150.1.6.6) instead of the forward address (155.1.67.6). There can however be certain designs where there is a shorter path to the forward address than the advertising router's address, which is explored in the task related to multiple exit points out of the NSSA.

The result of the translation on R3 is that devices in area 0 see the routes as Type-5 External LSAs, not Type-7.

```
Rack1R5#show ip ospf database | begin Type-5
```

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

Link ID	ADV Router	Age	Seq#	Checksum	Tag
30.0.0.0	150.1.4.4	982	0x80000001	0x00D80A	0
30.1.0.0	150.1.4.4	982	0x80000001	0x00CC15	0
30.2.0.0	150.1.4.4	982	0x80000001	0x00C020	0
30.3.0.0	150.1.4.4	982	0x80000001	0x00B42B	0
31.0.0.0	150.1.4.4	982	0x80000001	0x00CB16	0
31.1.0.0	150.1.4.4	982	0x80000001	0x00BF21	0
31.2.0.0	150.1.4.4	982	0x80000001	0x00B32C	0
31.3.0.0	150.1.4.4	982	0x80000001	0x00A737	0
51.51.51.51	192.10.1.254	1334	0x8000012D	0x00419A	0
54.1.1.0	150.1.3.3	781	0x80000001	0x0010D4	0
200.0.0.0	150.1.3.3	779	0x80000001	0x00B59E	0
200.0.1.0	150.1.3.3	779	0x80000001	0x00AAA8	0
200.0.2.0	150.1.3.3	780	0x80000001	0x009FB2	0
200.0.3.0	150.1.3.3	780	0x80000001	0x0094BC	0
204.12.1.0	150.1.4.4	984	0x80000001	0x005EC8	0

```
Rack1R5#show ip ospf database external 200.0.0.0
```

```
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: 173
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 200.0.0.0 (External Network Number )
Advertising Router: 150.1.3.3
LS Seq Number: 80000002
Checksum: 0xB39F
Length: 36
Network Mask: /24
Metric Type: 2 (Larger than any link state path)
TOS: 0
Metric: 20
Forward Address: 155.1.67.6
External Route Tag: 0
```

R5 performs a lookup on the now Type-5 External LSA, and like R3, sees the forward address set to 155.1.67.6. Again note that the lookup process for this translated Type-7 LSA is performed differently than a normal inter-area Type-5 External LSA lookup, because R5 computes its metric towards 155.1.67.6, and not a Type-4 LSA describing the ASBR. Furthermore, note that R3 does not generate a Type-4 ASBR Summary LSA describing R6.

```
Rack1R5#show ip ospf database asbr-summary 150.1.6.6
```

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

```
Rack1R5#show ip route 155.1.67.6
```

```
Routing entry for 155.1.67.0/24
Known via "ospf 1", distance 110, metric 1065, type inter area
Last update from 155.1.0.3 on Serial0/0, 00:35:16 ago
Routing Descriptor Blocks:
* 155.1.0.3, from 150.1.3.3, 00:35:16 ago, via Serial0/0
Route metric is 1065, traffic share count is 1
```

R5's metric to the forwarding address 155.1.67.6 is 1065 via R3. This is the value installed as the forward metric for the translated Type-7 LSA, with a metric of 20 from the Type-5 LSA itself.

```
Rack1R5#show ip route 200.0.0.0
```

```
Routing entry for 200.0.0.0/24
Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric
1065
Last update from 155.1.0.3 on Serial0/0, 00:35:51 ago
Routing Descriptor Blocks:
* 155.1.0.3, from 150.1.3.3, 00:35:51 ago, via Serial0/0
Route metric is 20, traffic share count is 1
```

Similar to the stub area, the NSSA flag must be agreed upon by all devices in the area, or adjacency cannot occur. This implies that the area is a normal area, a stub area, or an NSSA, but no combination of the three.

Furthermore, like the stub area, Type-5 External LSAs are blocked from entering the NSSA area on the ABR(s).

```
Rack1SW1#show ip route | include E1|E2
      E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
Rack1SW1#show ip ospf database | begin Type-5
```

### ☠ Pitfall

The other key difference between stub and NSSA areas is how default routing works. The stub area removes external LSAs and replaces them with a default route. The totally stubby area extends this by replacing external LSAs and inter-area LSAs with a default route. However *with the NSSA a default route is not automatically originated by the ABR*. This means that devices within the NSSA will have reachability to their own area and to other areas, but not to destinations outside of the OSPF domain.

```
Rack1R3#show ip route 30.0.0.1
Routing entry for 30.0.0.0/16
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 845
  Last update from 155.1.0.5 on Serial1/0, 00:02:09 ago
  Routing Descriptor Blocks:
    * 155.1.0.5, from 150.1.4.4, 00:02:09 ago, via Serial1/0
      Route metric is 20, traffic share count is 1
```

```
Rack1R3#ping 30.0.0.1
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 30.0.0.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 104/113/128 ms
```

```
Rack1SW1#show ip route 30.0.0.1
% Network not in table
```

```
Rack1SW1#show ip route | include Gateway
Gateway of last resort is not set
```

```
Rack1SW1#ping 30.0.0.1
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 30.0.0.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

## 6.27 OSPF Not-So-Stubby Areas and Default Routing

- Configure R3 to advertise a default route into area 2.
- This default route should have a cost of 500.

### Configuration

```
R3:
router ospf 1
 area 2 nssa default-information-originate
 area 2 default-cost 500
```

### Verification

#### Note

Unlike the stub area, totally-stubby area, or not-so-totally-stubby area, the ABR(s) of an NSSA do not automatically originate a default route. A default route can be originated as a Type-7 NSSA External LSA into the NSSA by adding the **default-information-originate** option onto the **area [id] nssa** statement. The cost that the ABR advertises for the default can be modified with the **area [id] default-cost** command.

```
Rack1SW1#show ip route | include _0.0.0.0
Gateway of last resort is 155.1.37.3 to network 0.0.0.0
O*N2 0.0.0.0/0 [110/500] via 155.1.37.3, 00:00:17, FastEthernet0/3
```

```
Rack1SW1#show ip ospf database | begin Type-7
```

Type-7 AS External Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	150.1.3.3	31	0x80000001	0x00EB6E	0
54.1.1.0	150.1.6.6	56	0x80000010	0x00368F	0
200.0.0.0	150.1.6.6	57	0x80000010	0x00DB59	0
200.0.1.0	150.1.6.6	57	0x80000010	0x00D063	0
200.0.2.0	150.1.6.6	57	0x80000010	0x00C56D	0
200.0.3.0	150.1.6.6	57	0x80000010	0x00BA77	0

## 6.28 OSPF Not-So-Totally-Stubby Areas

- Remove the previous default routing configuration on R3.
- Modify the area 2 configuration so that R3 filters Inter-Area and External routes as they are sent from area 0 to area 2.
- R6 should still be allowed to redistribute EIGRP into OSPF.

### Configuration

```
R3:
router ospf 1
 no area 2 nssa
 area 2 nssa no-summary
```

### Verification

#### Note

The not-so-totally-stubby area is the combination of the totally-stubby area and the NSSA. Like the totally-stubby area, Type-3 Summary LSAs, Type-4 ASBR Summary LSAs, and Type-5 External LSAs are removed and replaced with a Type-3 Summary LSA default route. Like the NSSA, Type-7 NSSA External LSAs are allowed to be originated inside the area.

The combination of these two result in the blocking of all inter-area OSPF routes & routes external to the OSPF domain, replacing them with a default route, and allowing redistribution to occur. SW3's routing table output indicates this because only intra-area, NSSA external, and an inter-area default route are installed. Note that the cost of the default route is derived from the intra-area cost to R3, plus the `area 2 default-cost 500` previously configured on R3.

```
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, 00:00:53, Vlan79
O       155.1.37.0 [110/2] via 155.1.79.7, 00:00:53, Vlan79
O       155.1.67.0 [110/1001] via 155.1.79.7, 00:00:53, Vlan79
O N2   200.0.0.0/24 [110/20] via 155.1.79.7, 00:00:53, Vlan79
    54.0.0.0/24 is subnetted, 1 subnets
O N2   54.1.1.0 [110/20] via 155.1.79.7, 00:00:53, Vlan79
O N2   200.0.1.0/24 [110/20] via 155.1.79.7, 00:00:53, Vlan79
O N2   200.0.2.0/24 [110/20] via 155.1.79.7, 00:00:53, Vlan79
O N2   200.0.3.0/24 [110/20] via 155.1.79.7, 00:00:54, 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, 00:00:54, Vlan79
O*IA   0.0.0.0/0 [110/502] via 155.1.79.7, 00:00:54, Vlan79
```

The database output indicates that the only Type-3 Summary LSA is the default route originated by the ABR.

**Rack1SW3#show ip ospf database**

OSPF Router with ID (150.1.9.9) (Process ID 1)

Router Link States (Area 2)

Link ID count	ADV Router	Age	Seq#	Checksum	Link
150.1.3.3	150.1.3.3	512	0x80000015	0x00AD9F	1
150.1.6.6	150.1.6.6	1213	0x80000014	0x0007FC	1
150.1.7.7	150.1.7.7	60	0x8000001A	0x007D8D	5
150.1.9.9	150.1.9.9	59	0x80000015	0x004327	3

Net Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum
155.1.37.7	150.1.7.7	1235	0x80000012	0x00244D
155.1.67.7	150.1.7.7	1236	0x80000012	0x002A23
155.1.79.9	150.1.9.9	60	0x80000001	0x00DA6B

Summary Net Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	150.1.3.3	566	0x80000001	0x004F54

Type-7 AS External Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum	Tag
54.1.1.0	150.1.6.6	1214	0x80000010	0x00368F	0
200.0.0.0	150.1.6.6	1215	0x80000010	0x00DB59	0
200.0.1.0	150.1.6.6	1215	0x80000010	0x00D063	0
200.0.2.0	150.1.6.6	1215	0x80000010	0x00C56D	0
200.0.3.0	150.1.6.6	1215	0x80000010	0x00BA77	0

SW3 does not have a longer match to 30.0.0.1, but can use its default information to reach it.

**Rack1SW3#show ip route 30.0.0.1**

% Network not in table

**Rack1SW3#ping 30.0.0.1**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 30.0.0.1, timeout is 2 seconds:

!!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 92/104/120 ms

## 6.29 OSPF Stub Areas with Multiple Exit Points

- Re-enable R6's Loopback0 interface and connection to VLAN 146.
- Modify area 2 in such a way that traffic from devices in area 2 going to area 3 uses R6 as the exit point, while traffic from area 2 going to the RIP routes learned from BB3 uses R3.
- If R3 loses its connection to area 2, traffic for External destinations should be rerouted to R6.
- If R6 loses its connection to area 2, traffic for Inter-Area destinations should be rerouted to R3.
- Do not modify the cost of any links in area 2 to accomplish this.

### Configuration

---

```
R3:
router ospf 1
 area 2 nssa no-summary
```

```
R6:
router ospf 1
 area 2 nssa default-information-originate
```

```
SW1:
router ospf 1
 area 2 nssa
```

```
SW3:
router ospf 1
 area 2 nssa
```

### Verification

---

#### Note

In addition to database filtering, stub areas can be used for inter-area traffic engineering. In this particular case there are multiple exit points out of area 2, R3 and R6. R3 advertises only a default route as a Type-3 Summary LSA, due to its not-so-totally-stubby configuration. R6 advertises all Type-3 Summary LSAs, plus a Type-7 NSSA External default route. For inter-area routing from devices in area 2 this means that the longest match learned from R6 will always be used, and for default routing the Type-3 default will be used from R3. The default preference through R3 occurs since OSPF always prefers routes in the sequence intra-area > inter-area > external > nssa-external.



The result of this can be seen on SW1 in the database and the routing table.

```
Rack1SW1#show ip ospf database | begin Summary
```

```
Summary Net Link States (Area 2)

Link ID          ADV Router      Age             Seq#            Checksum
0.0.0.0          150.1.3.3      1457           0x80000002     0x004D55
150.1.1.1        150.1.6.6      323            0x80000002     0x00D81E
150.1.2.2        150.1.6.6      323            0x80000002     0x00C8AB
150.1.3.3        150.1.6.6      324            0x80000002     0x00B3BE
150.1.4.4        150.1.6.6      324            0x80000003     0x009758
150.1.5.5        150.1.6.6      324            0x80000002     0x0007A7
150.1.6.6        150.1.6.6      324            0x80000002     0x006588
150.1.8.8        150.1.6.6      324            0x80000002     0x00D1D5
150.1.10.10     150.1.6.6      324            0x80000002     0x00B1F0
<snip>
```

The only inter-area route learned from R3 is a default route. This means traffic to the inter-area destination 150.1.10.10 must use R6 as the exit point.

```
Rack1SW1#show ip route 150.1.10.10
```

```
Routing entry for 150.1.10.10/32
  Known via "ospf 1", distance 110, metric 1068, type inter area
  Last update from 155.1.67.6 on Vlan67, 00:39:46 ago
  Routing Descriptor Blocks:
  * 155.1.67.6, from 150.1.6.6, 00:39:46 ago, via Vlan67
    Route metric is 1068, traffic share count is 1
```

```
Rack1SW1#traceroute 150.1.10.10
```

```
Type escape sequence to abort.
Tracing the route to 150.1.10.10
```

```
 1 155.1.67.6 0 msec 0 msec 8 msec
 2 155.1.146.4 0 msec 8 msec 0 msec
 3 155.1.0.5 34 msec 34 msec 25 msec
 4 155.1.58.8 33 msec 34 msec 33 msec
 5 155.1.108.10 34 msec * 25 msec
```

SW1 does not have a longer match for the destination 30.0.0.1, but it does have default information through R3.

```
Rack1SW1#show ip route 30.0.0.1
```

```
% Network not in table
```

```
Rack1SW1#show ip route | include _0.0.0.0
```

```
Gateway of last resort is 155.1.37.3 to network 0.0.0.0
O*IA 0.0.0.0/0 [110/501] via 155.1.37.3, 00:40:44, FastEthernet0/3
```

The result is that for routes outside of the OSPF domain, such as the RIP routes redistributed into OSPF on R4, R3 is used as the exit point.

```
Rack1SW1#traceroute 30.0.0.1
```

```
Type escape sequence to abort.
Tracing the route to 30.0.0.1
```

```
 1 155.1.37.3 0 msec 0 msec 0 msec
 2 155.1.0.5 26 msec 33 msec 34 msec
 3 155.1.45.4 17 msec
   155.1.0.4 25 msec
   155.1.45.4 16 msec
 4 204.12.1.254 26 msec * 92 msec
```

Note that in the database SW1 has two default routes, one from R3 as a Type-3 Summary LSA, and one from R6 as a Type-7 NSSA External route. Even though the metric of the NSSA route is lower than the Type-3 route, the Type-3 route via R3 is always preferred.

```
Rack1SW1#show ip ospf database summary 0.0.0.0
```

```
OSPF Router with ID (150.1.7.7) (Process ID 1)
```

```
Summary Net Link States (Area 2)
```

```
Routing Bit Set on this LSA
LS age: 1619
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 0.0.0.0 (summary Network Number)
Advertising Router: 150.1.3.3
LS Seq Number: 80000002
Checksum: 0x4D55
Length: 28
Network Mask: /0
TOS: 0 Metric: 500
```

```
Rack1SW1#show ip ospf database nssa-external 0.0.0.0
```

```
OSPF Router with ID (150.1.7.7) (Process ID 1)
```

```
Type-7 AS External Link States (Area 2)
```

```
LS age: 493
Options: (No TOS-capability, No Type 7/5 translation, DC)
LS Type: AS External Link
Link State ID: 0.0.0.0 (External Network Number )
Advertising Router: 150.1.6.6
LS Seq Number: 80000002
Checksum: 0xE343
Length: 36
Network Mask: /0
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 1
    Forward Address: 155.1.67.6
    External Route Tag: 0
```

Only once R3 withdraws its inter-area default route can R6 be used as the default exit point.

```
Rack1R3#config t
```

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

```
Rack1R3(config)#interface Fa0/0
```

```
Rack1R3(config-if)#shutdown
```

```
Rack1R3(config-if)#
```

```
Rack1SW1#show ip route | include _0.0.0.0
```

```
Gateway of last resort is 155.1.67.6 to network 0.0.0.0
```

```
O*N2 0.0.0.0/0 [110/1] via 155.1.67.6, 00:00:00, Vlan67
```

```
Rack1SW1#traceroute 30.0.0.1
```

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

```
Tracing the route to 30.0.0.1
```

```
 1 155.1.67.6 9 msec 0 msec 0 msec
 2 155.1.146.4 8 msec 0 msec 0 msec
 3 204.12.1.254 8 msec * 0 msec
```

Likewise if R6 loses connectivity to area 0, inter-area destinations are rerouted to R3 since there is no longer match than the default route.

```
Rack1R3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R3(config)#interface Fa0/0
Rack1R3(config-if)#no shutdown
Rack1R3(config-if)#

Rack1R6#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R6(config)#interface Fa0/0.146
Rack1R6(config-subif)#shutdown
Rack1R6(config-subif)#

Rack1SW1#show ip route 150.1.10.10
% Subnet not in table

Rack1SW1#show ip route | include _0.0.0.0
Gateway of last resort is 155.1.37.3 to network 0.0.0.0
O*IA 0.0.0.0/0 [110/501] via 155.1.37.3, 00:00:05, FastEthernet0/3

Rack1SW1#traceroute 150.1.10.10

Type escape sequence to abort.
Tracing the route to 150.1.10.10

 1 155.1.37.3 0 msec 0 msec 9 msec
 2 155.1.0.5 25 msec 34 msec 33 msec
 3 155.1.58.8 34 msec 33 msec 34 msec
 4 155.1.108.10 33 msec * 26 msec
```

### 6.30 OSPF NSSA Type-7 to Type-5 Translator Election

- Create a new Loopback9 interface on SW3 with the IP address 9.9.9.9/32, and redistribute it into OSPF.
- Modify the OSPF domain so that only R3 advertises this route into area 0, but traffic from R5 going to this destination transits R6.
- If R6 is down this traffic should be rerouted to R3.

#### Configuration

```
R3:
interface FastEthernet0/0
 ip ospf cost 1000
!
router ospf 1
 router-id 150.1.30.30

SW3:
interface Loopback9
 ip address 9.9.9.9 255.255.255.255
!
router ospf 1
 redistribute connected subnets
```

#### Verification

##### Note

As previously seen with OSPF Not-So-Stubby Areas, Type-7 NSSA External LSAs are translated to Type-5 External LSAs by the ABR connecting the NSSA to area 0. When multiple ABRs connect the NSSA to area 0, the ABR with the highest router-id is elected as the Type-7 to 5 translator, and is responsible for re-originating the Type-5 LSA into area 0. This election process is an optimization of the OSPF database, and relates to how the Type-7 NSSA External route uses the forward address field to ensure optimal routing.

Recall that with normal external routes, only one Type-5 LSA is originated by the router performing the redistribution. When the route moves between areas, each ABR originates a Type-4 ASBR Summary LSA advertising their reachability to the ASBR. This means that for all Type-5 External LSA inter-area lookups OSPF would require  $Ext\_Routes + Num\_ABRs + Num\_Routers$  LSAs, where *Ext\_Routes* is the number of Type-5 LSAs, *Num\_ABRs* is the number of ABRs generating Type-4 ASBR summaries, and *Num\_Routers* is the number of Type-1 LSAs from the routers in the local area.

Now with Type-7 LSAs the situation becomes more complicated, because this information needs to be re-originated at the ABR level as the route moves into area 0. Let's suppose for the sake of argument that each ABR connecting the NSSA to area 0 *did* do a translation of Type-7 to 5. This would mean for all inter-area lookups on a Type-5 External LSAs that were translated from Type-7, there would be  $(NSSA\_Routes * Num\_ABRs) + Num\_ABRs + Num\_Routers$  LSAs, where  $NSSA\_Routes$  is the number of Type-7 LSAs to start.

This operation would be highly redundant and inefficient, because each ABR would re-originate the same Type-5 LSA, each with the same forwarding address. To avoid this only one ABR performs the Type-7 to 5 translation, but maintains the forward address field, essentially separating the relationship between the routing advertisement and the traffic flow. This principle can be illustrated as follows.

Before any modification in the OSPF domain, R5 performs a lookup on the Type-5 LSA for 9.9.9.9 that was translated from a Type-7 LSA. At this point R3 has an OSPF Router-ID of 150.1.3.3, and R6 has 150.1.6.6. The advertising router R5 sees is 150.1.6.6 (R6), since R6 won the translator election due to the higher RID. Note however, the forward address is set to 150.1.9.9 (SW3). This means that R5 next needs to figure out how to route towards 150.1.9.9.

```
Rack1R5#show ip ospf database external 9.9.9.9
```

```
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: 224
```

```
Options: (No TOS-capability, DC)
```

```
LS Type: AS External Link
```

```
Link State ID: 9.9.9.9 (External Network Number )
```

```
Advertising Router: 150.1.6.6
```

```
LS Seq Number: 80000001
```

```
Checksum: 0x1F10
```

```
Length: 36
```

```
Network Mask: /32
```

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

```
  TOS: 0
```

```
  Metric: 20
```

```
  Forward Address: 150.1.9.9
```

```
  External Route Tag: 0
```

Since 150.1.9.9 does not belong to a device in the same area as R5, an inter-area lookup is performed on the Type-3 LSA. R5 finds that two ABRs are advertising the route to 150.1.9.9, 150.1.3.3 (R3) and 150.1.6.6 (R6), both with a metric of 3.

```
Rack1R5#show ip ospf database summary 150.1.9.9
```

```
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: 603
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 150.1.9.9 (summary Network Number)
```

```
Advertising Router: 150.1.3.3
```

```
LS Seq Number: 80000001
```

```
Checksum: 0xBD35
```

```
Length: 28
```

```
Network Mask: /32
```

```
TOS: 0 Metric: 3
```

```
LS age: 4 (DoNotAge)
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 150.1.9.9 (summary Network Number)
```

```
Advertising Router: 150.1.6.6
```

```
LS Seq Number: 80000001
```

```
Checksum: 0x9656
```

```
Length: 28
```

```
Network Mask: /32
```

```
TOS: 0 Metric: 3
```

R5 now needs to find what the metric is to reach these ABRs. R5 checks its locally originated Type-1 Router LSA and finds that 150.1.1.1 (R1) and 150.1.3.3 (R3) are directly attached, both with a metric of 64.

```
Rack1R5#show ip ospf database router 150.1.5.5 self-originate
```

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

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

```
LS age: 1620
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.5.5
Advertising Router: 150.1.5.5
LS Seq Number: 80000019
Checksum: 0x7DAE
Length: 108
Area Border Router
Number of Links: 7
```

```
<snip>
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.1.1
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 64
```

```
<snip>
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.3.3
(Link Data) Router Interface address: 155.1.0.5
Number of TOS metrics: 0
TOS 0 Metrics: 64
```

```
<snip>
```



R5 asks R1 who it is adjacent with, and finds 150.1.6.6 (R6) is connected over the a Virtual-Link with a metric of 1.

```
Rack1R5#show ip ospf database router 150.1.1.1
```

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

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

```
Routing Bit Set on this LSA
LS age: 773
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.1.1
Advertising Router: 150.1.1.1
LS Seq Number: 8000001B
Checksum: 0xFA2E
Length: 72
Area Border Router
Number of Links: 4
```

```
Link connected to: a Virtual Link
(Link ID) Neighboring Router ID: 150.1.6.6
(Link Data) Router Interface address: 155.1.146.1
Number of TOS metrics: 0
TOS 0 Metrics: 1
```

This means that R5's intra-area cost to R3 is 64, and to R6 is 65. Since both R3 and R6 reported a cost of 3 to the forward address 150.1.9.9, the total forward metric through R6 is 65+3 = 68, but is only 64+3 = 67 through R3. Therefore although the route is originated into area 0 by R6, the Type-7 to 5 translator, the traffic does not actually flow through R6. Instead the path through R3 installed with the default redistribution metric of 20 for the E2 route, and a forward metric of 67 through R3.

```
Rack1R5#show ip route 9.9.9.9
```

```
Routing entry for 9.9.9.9/32
Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 67
Last update from 155.1.0.3 on Serial0/0, 00:17:27 ago
Routing Descriptor Blocks:
* 155.1.0.3, from 150.1.6.6, 00:17:27 ago, via Serial0/0
Route metric is 20, traffic share count is 1
```

```
Rack1R5#traceroute 9.9.9.9
```

```
Type escape sequence to abort.
Tracing the route to 9.9.9.9
```

```
 1 155.1.0.3 28 msec 32 msec 32 msec
 2 155.1.37.7 32 msec 32 msec 32 msec
 3 155.1.79.9 28 msec * 28 msec
```

This illustrates why a Type-5 External route that was translated from a Type-7 NSSA External route does not use a Type-4 ASBR Summary LSA, because the forward address lookup replaces the need for the ASBR Summary lookup. Since the forward address is preserved only one router needs to do the translation, while the calculation of the final forwarding path stays independent

The Type-7 to 5 translator election can be modified by increasing R3's router-id to be higher than R6's.

```
Rack1R3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R3(config)#router ospf 1
Rack1R3(config-router)#router-id 150.1.30.30
Reload or use "clear ip ospf process" command, for this to take effect
Rack1R3(config-router)#end
Rack1R3#clear ip ospf 1 process
Reset OSPF process? [no]: yes
Rack1R3#
```

```
Rack1R5#show ip ospf database external 9.9.9.9
```

```
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: 42
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 9.9.9.9 (External Network Number )
Advertising Router: 150.1.30.30
LS Seq Number: 80000001
Checksum: 0xE519
Length: 36
Network Mask: /32
Metric Type: 2 (Larger than any link state path)
TOS: 0
Metric: 20
Forward Address: 150.1.9.9
External Route Tag: 0
```

R5 now sees the advertising router as 150.1.30.30 (R3), because this is the highest router-id of the ABRs connecting the NSSA to area 0.

Although the advertising router has changed, the forward address is still 150.1.9.9, which means the traffic flow has not changed.

```
Rack1R5#traceroute 9.9.9.9
```

```
Type escape sequence to abort.
Tracing the route to 9.9.9.9
```

```
 1 155.1.0.3 28 msec 32 msec 28 msec
 2 155.1.37.7 32 msec 32 msec 32 msec
 3 155.1.79.9 28 msec * 28 msec
```

Only once the forward metric via R3 is higher than the forward metric via R6 will the path selection change. This can be accomplished by changing the OSPF cost on R3's link to SW1.

```
Rack1R3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R3(config)#interface Fa0/0
Rack1R3(config-if)#ip ospf cost 1000
Rack1R3(config-if)#end
```

Now when R5 computes the forward metric through R3 it sees the same intra-area cost of 64 to R3, but R3 increased the advertised metric to 150.1.9.9 to 1002.

```
Rack1R5#show ip ospf database summary 150.1.9.9 adv-router 150.1.30.30
```

```
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: 70
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.9.9 (summary Network Number)
Advertising Router: 150.1.30.30
LS Seq Number: 80000002
Checksum: 0x854B
Length: 28
Network Mask: /32
TOS: 0 Metric: 1002
```

The final result is that the route is installed via R6 with a metric of 20 reported by the ASBR, plus the forward metric of 68 through R6.

```
Rack1R5#show ip route 9.9.9.9
```

```
Routing entry for 9.9.9.9/32
```

```
Known via "ospf 1", distance 110, metric 20, type extern 2, forward  
metric 68
```

```
Last update from 155.1.0.1 on Serial0/0, 00:02:06 ago
```

```
Routing Descriptor Blocks:
```

```
* 155.1.0.1, from 150.1.30.30, 00:02:06 ago, via Serial0/0  
Route metric is 20, traffic share count is 1
```

```
Rack1R5#traceroute 9.9.9.9
```

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

```
Tracing the route to 9.9.9.9
```

```
 1 155.1.0.1 29 msec 28 msec 28 msec  
 2 155.1.146.6 28 msec 32 msec 28 msec  
 3 155.1.67.7 32 msec 32 msec 28 msec  
 4 155.1.79.9 32 msec * 28 msec
```

If R6 loses connectivity to area 0 its route to the forward address is lost, and traffic is rerouted to R3.

```
Rack1R6#config t
```

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

```
Rack1R6(config)#interface Fa0/0.146
```

```
Rack1R6(config-subif)#shutdown
```

```
Rack1R6(config-subif)#
```

```
Rack1R5#traceroute 9.9.9.9
```

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

```
Tracing the route to 9.9.9.9
```

```
 1 155.1.0.3 28 msec 32 msec 28 msec  
 2 155.1.37.7 28 msec 28 msec 28 msec  
 3 155.1.79.9 32 msec * 28 msec
```

### 6.31 OSPF NSSA Redistribution Filtering

- Create new Loopback interfaces on R5 and SW4 with IP addresses 5.5.5.5/32 and 10.10.10.10/32 respectively, and redistribute them into the OSPF domain.
- Configure area 3 as an NSSA so R5 blocks all LSA types 3, 4, and 5 and replaces them with a default route.
- Modify area 3 so that R5's redistributed Loopback is injected into area 0 as LSA Type-5, but is not injected into area 3 as LSA Type-7.

#### Configuration

---

```
R5:
interface Loopback5
 ip address 5.5.5.5 255.255.255.255
!
router ospf 1
 redistribute connected subnets
 area 3 nssa no-redistribution no-summary
```

```
SW2:
router ospf 1
 area 3 nssa
```

```
SW4:
interface Loopback10
 ip address 10.10.10.10 255.255.255.255
!
router ospf 1
 redistribute connected subnets
 area 3 nssa
```

#### Verification

---

##### Note

In certain NSSA designs the ABR can be an ASBR at the same time. When routes are redistributed directly on the ABR they are originated into area 0 as Type-5 External LSAs, and into the NSSA as Type-7 NSSA External LSAs. The origination as Type-7 into the NSSA may be unneeded overhead if the ABR performing redistribution is the only exit point out of the area. In this particular case R5 is both an ABR and ASBR, and is the only exit point for SW2 and SW4 to route packets into area 0.

By configuring the **area 3 nssa no-summary** option on R5, along with the **area 3 nssa** option on SW2 and SW4, the amount of routes contained in the area 3 database are minimized, while still allowing redistribution on SW4. Before R5 performs redistribution the database in area 3 looks as follows.

```
Rack1SW2#show ip ospf database
```

```
OSPF Router with ID (150.1.8.8) (Process ID 1)
```

```
Router Link States (Area 3)
```

Link ID count	ADV Router	Age	Seq#	Checksum	Link
150.1.5.5	150.1.5.5	142	0x80000035	0x004250	3
150.1.8.8	150.1.8.8	138	0x80000033	0x005FEB	4
150.1.10.10	150.1.10.10	638	0x80000030	0x00E128	3

```
Net Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.58.8	150.1.8.8	138	0x80000030	0x003203
155.1.108.10	150.1.10.10	638	0x8000002E	0x00579D

```
Summary Net Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	150.1.5.5	148	0x80000003	0x00A0F1

```
Type-7 AS External Link States (Area 3)
```

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.10.10.10	150.1.10.10	643	0x80000001	0x0043D3	0

SW2 knows about the three routers in the area via Type-1 Router LSAs, the two DRs in the area via Type-2 Network LSAs, an inter-area default route originated by R5 as a Type-3 Summary LSA, and SW4's redistributed Loopback via a Type-7 NSSA External LSA. This is essentially the minimal information needed in the database to perform intra-area SPF, use default routing to leave the area, and still allow redistribution. Next, R5 performs redistribution into OSPF.

```
Rack1R5#config t
```

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

```
Rack1R5(config)#router ospf 1
```

```
Rack1R5(config-router)#redistribute connected subnets
```

R5 originates the link 5.5.5.5/32 into area 3 as a Type-7 NSSA External LSA, as well as into area 0 as a Type-5 External LSA.

```
Rack1R5#show ip ospf database | begin Type-7
Type-7 AS External Link States (Area 3)

Link ID          ADV Router      Age             Seq#            Checksum Tag
5.5.5.5          150.1.5.5      156            0x80000001     0x005098 0
10.10.10.10     150.1.10.10   1219          0x80000001     0x0043D3 0

Type-5 AS External Link States

Link ID          ADV Router      Age             Seq#            Checksum Tag
5.5.5.5          150.1.5.5      156            0x80000001     0x006C7E 0
9.9.9.9          150.1.30.30   1484          0x8000001A     0x00B332 0
10.10.10.10     150.1.5.5      701           0x80000001     0x001912 0
<snip>
```

The problem with this design is that unnecessary information is now in the database of area 3. Since the area 3 routers already had a default route via R5, having specific reachability information about the network 5.5.5.5/32 is redundant.

```
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, 00:14:41, Port-channell
O    155.1.5.0 [110/2] via 155.1.58.5, 00:14:41, Vlan58
5.0.0.0/32 is subnetted, 1 subnets
O N2 5.5.5.5 [110/20] via 155.1.58.5, 00:05:32, Vlan58
10.0.0.0/32 is subnetted, 1 subnets
O N2 10.10.10.10 [110/20] via 155.1.108.10, 00:14:41, Port-channell
150.1.0.0/16 is variably subnetted, 3 subnets, 2 masks
O    150.1.10.10/32 [110/2] via 155.1.108.10, 00:14:41, Port-channell
O    150.1.5.5/32 [110/2] via 155.1.58.5, 00:14:41, Vlan58
O*IA 0.0.0.0/0 [110/2] via 155.1.58.5, 00:14:41, Vlan58
```

Therefore this design is a good candidate for Type-7 LSA suppression on the ABR itself. By adding the **no-redistribution** keyword onto the **area 3 nssa** statement of R5, Type-7 LSAs are not generated for locally redistributed routes. This does not, however, prevent other devices inside the NSSA from performing redistribution, such as SW4.

```
Rack1R5#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R5(config)#router ospf 1
Rack1R5(config-router)#area 3 nssa no-redistribution no-summary
```

Devices in area 3 no longer have a specific route to 5.5.5.5/32, but can reach it using the default route. Also a Type-7 NSSA External LSA still exists for 10.10.10.10/32.

**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, 00:21:20, Port-channell
O    155.1.5.0 [110/2] via 155.1.58.5, 00:21:20, Vlan58
10.0.0.0/32 is subnetted, 1 subnets
O N2 10.10.10.10 [110/20] via 155.1.108.10, 00:04:49, Port-channell
150.1.0.0/16 is variably subnetted, 3 subnets, 2 masks
O    150.1.10.10/32 [110/2] via 155.1.108.10, 00:21:20, Port-channell
O    150.1.5.5/32 [110/2] via 155.1.58.5, 00:21:20, Vlan58
O*IA 0.0.0.0/0 [110/2] via 155.1.58.5, 00:04:54, Vlan58
```

**Rack1SW2#ping 5.5.5.5**

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:

!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms

Devices in area 0 and beyond have a specific route for 5.5.5.5/32 as a Type-5 External LSA.

**Rack1R3#show ip route 5.5.5.5**

Routing entry for 5.5.5.5/32

Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 781

Last update from 155.1.0.5 on Serial1/0, 00:12:20 ago

Routing Descriptor Blocks:

\* 155.1.0.5, from 150.1.5.5, 00:12:20 ago, via Serial1/0

Route metric is 20, traffic share count is 1



## 6.32 OSPF LSA Type-3 Filtering

- Configure LSA Type-3 Filtering on R5 so that devices in area 0 do not have reachability information about the layer 3 EtherChannel link between SW2 and SW4 or SW4's Loopback0 interface.
- Configure LSA Type-3 Filtering on R2 so that devices in area 51 do not have reachability information about the Loopback0 interface of R5; this should not affect devices in area 5.

### Configuration

---

```
R2:
ip prefix-list R5_LOOPBACK deny 150.1.5.5/32
ip prefix-list R5_LOOPBACK permit 0.0.0.0/0 le 32
!
router ospf 1
 area 51 filter-list prefix R5_LOOPBACK in

R5:
ip prefix-list AREA_3_ROUTES deny 150.1.10.10/32
ip prefix-list AREA_3_ROUTES deny 155.1.108.0/24
ip prefix-list AREA_3_ROUTES permit 0.0.0.0/0 le 32
!
router ospf 1
 area 3 filter-list prefix AREA_3_ROUTES out
```

### Verification

---

#### Note

LSA Type-3 Filtering, like stub areas, is used to remove LSAs from the database as advertisements move between areas. Unlike stub areas however, Type-3 LSA Filtering can be used to permit or deny any arbitrary inter-area routes based on a prefix-list.

The filter-list syntax supports the arguments *in* and *out*, which are used to allow more control on ABRs that terminate multiple areas. In the case of R5, who terminates only area 0 and area 3, the syntax **area 3 filter-list prefix AREA\_3\_ROUTES out** has the same result as **area 0 filter-list prefix AREA\_3\_ROUTES in**. However on R2, who terminates areas 0, 5, and 51, the syntax **area 51 filter-list prefix R5\_LOOPBACK in** applies only to prefixes leaving areas 0 and 5 going into area 51, while the syntax **area 0 filter-list prefix R5\_LOOPBACK out** would apply to prefixes leaving area 0 going into area 5 and 51.

In other words if **area 0 filter-list prefix R5\_LOOPBACK out** was applied, R5's loopback would not enter area 5 or 51 on R2, but with **area 51 filter-list prefix R5\_LOOPBACK in** applied, R5's loopback enters area 5 but not area 51. This configuration can be verified by viewing the database and the routing table. Prior to filtering R2 originated 150.1.5.5 into area 51 as a Type-3 Summary LSA.

```
Rack1R2#show ip ospf database | begin Area 51
```

```

Router Link States (Area 51)

Link ID        ADV Router    Age           Seq#           Checksum Link count
150.1.2.2      150.1.2.2    1062         0x8000002F   0x00C76C 1
192.10.1.254   192.10.1.254 707          0x8000015E   0x00CFCE 1

Net Link States (Area 51)

Link ID        ADV Router    Age           Seq#           Checksum
192.10.1.254   192.10.1.254 707          0x8000002E   0x0039D7

Summary Net Link States (Area 51)

Link ID        ADV Router    Age           Seq#           Checksum
150.1.1.1      150.1.2.2    310          0x8000002E   0x000A4F
150.1.2.2      150.1.2.2    1062         0x8000002E   0x00EFE7
150.1.3.3      150.1.2.2    1827         0x8000001B   0x000662
150.1.4.4      150.1.2.2    310          0x8000002E   0x00CA88
150.1.5.5      150.1.2.2    310          0x8000002E   0x00335E
150.1.7.7      150.1.2.2    1827         0x80000023   0x00D497
150.1.8.8      150.1.2.2    1323         0x80000002   0x005660
150.1.9.9      150.1.2.2    1827         0x80000023   0x00B4B2
<snip>

```

After the filter is applied this route is withdrawn from area 51, but not from area 5.

```

Rack1R2#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R2(config)#ip prefix-list R5_LOOPBACK deny 150.1.5.5/32
Rack1R2(config)#ip prefix-list R5_LOOPBACK permit 0.0.0.0/0 le 32
Rack1R2(config)#router ospf 1
Rack1R2(config-router)#area 51 filter-list prefix R5_LOOPBACK in
Rack1R2(config-router)#end
Rack1R2#

```

```
Rack1R2#show ip ospf database | begin Area 51
```

```
Router Link States (Area 51)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link
count					
150.1.2.2	150.1.2.2	1080	0x8000002F	0x00C76C	1
192.10.1.254	192.10.1.254	725	0x8000015E	0x00CFCE	1

```
Net Link States (Area 51)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.10.1.254	192.10.1.254	725	0x8000002E	0x0039D7

```
Summary Net Link States (Area 51)
```

Link ID	ADV Router	Age	Seq#	Checksum
150.1.1.1	150.1.2.2	329	0x8000002E	0x000A4F
150.1.2.2	150.1.2.2	1080	0x8000002E	0x00EFE7
150.1.3.3	150.1.2.2	1845	0x8000001B	0x000662
150.1.4.4	150.1.2.2	329	0x8000002E	0x00CA88
150.1.7.7	150.1.2.2	1845	0x80000023	0x00D497
150.1.8.8	150.1.2.2	1341	0x80000002	0x005660
150.1.9.9	150.1.2.2	1845	0x80000023	0x00B4B2

<snip>

```
Rack1R2#show ip ospf database | begin Area 5
```

```
Router Link States (Area 5)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link
count					
150.1.2.2	150.1.2.2	376	0x80000031	0x008B73	2
150.1.30.30	150.1.30.30	427	0x8000001D	0x00E651	2

```
Summary Net Link States (Area 5)
```

Link ID	ADV Router	Age	Seq#	Checksum
150.1.1.1	150.1.2.2	870	0x8000002E	0x000A4F
150.1.1.1	150.1.30.30	427	0x8000001C	0x00DB87
150.1.2.2	150.1.2.2	1622	0x8000002E	0x00EFE7
150.1.2.2	150.1.30.30	427	0x8000001C	0x00C69A
150.1.3.3	150.1.2.2	376	0x8000001C	0x000463
150.1.3.3	150.1.30.30	427	0x8000001C	0x00911E
150.1.4.4	150.1.2.2	870	0x8000002E	0x00CA88
150.1.4.4	150.1.30.30	427	0x8000001C	0x009CC0
150.1.5.5	150.1.2.2	870	0x8000002E	0x00335E
150.1.5.5	150.1.30.30	427	0x8000001C	0x000596
150.1.7.7	150.1.2.2	376	0x80000024	0x00D298
150.1.7.7	150.1.30.30	428	0x8000001E	0x006D4C
150.1.8.8	150.1.2.2	1884	0x80000002	0x005660
150.1.8.8	150.1.30.30	1698	0x80000002	0x0004AA

<snip>

R5 knows how to reach SW4's Loopback0, along with the link between SW2 and SW4, but devices in area 0 do not.

```
Rack1R5#show ip route | include 150.1.10.10|155.1.108.0
O          155.1.108.0/24 [110/2] via 155.1.58.8, 00:33:37, FastEthernet0/0
O          150.1.10.10/32 [110/3] via 155.1.58.8, 00:33:37, FastEthernet0/0
```

```
Rack1R3#show ip route | include 150.1.10.10|155.1.108.0
<null>
```

### Pitfall

Note that devices in area 0 have the Type-5 External LSA 10.10.10.10/32 installed in the database, but not in the routing table. This problem is related to a lookup failure on the forward address 150.1.10.10, which is explored in detail in the next section.

```
Rack1R3#show ip ospf database external 10.10.10.10
```

```
OSPF Router with ID (150.1.30.30) (Process ID 1)
```

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

```
Routing Bit Set on this LSA
LS age: 1605
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 10.10.10.10 (External Network Number )
Advertising Router: 150.1.5.5
LS Seq Number: 80000001
Checksum: 0x1912
Length: 36
Network Mask: /32
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 150.1.10.10
    External Route Tag: 0
```

```
Rack1R3#show ip route 10.10.10.10
% Network not in table
```

### 6.33 OSPF Forwarding Address Suppression

- Modify R5's NSSA configuration so that devices outside of area 3 maintain connectivity to SW4's redistributed Loopback 10.10.10.10/32.

#### Configuration

```
R5:
router ospf 1
 area 3 nssa no-redistribution no-summary translate type7 suppress-fa
```

#### Verification

##### Note

Recall that with OSPF database lookups on external routes, the *Forward Address* field determines who the next recursive lookup should be performed towards. With typical Type-5 External LSAs, such as the RIP routes R4 redistributes into OSPF, the forward address is normally set to 0.0.0.0. This means that the next lookup should be performed towards the *Advertising Router*.

For example, when R4 redistributes the route 30.0.0.0/16 from RIP into OSPF, R1 performs a lookup on the Type-5 External LSA as follows.

```
Rack1R1#show ip ospf database external 30.0.0.0
```

```

      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: 1595
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 30.0.0.0 (External Network Number )
Advertising Router: 150.1.4.4
LS Seq Number: 8000002E
Checksum: 0x7E37
Length: 36
Network Mask: /16
  Metric Type: 2 (Larger than any link state path)
  TOS: 0
  Metric: 20
  Forward Address: 0.0.0.0
  External Route Tag: 0
```

R1 sees the forward address field set to 0.0.0.0, which means a lookup should be performed on the advertising router 150.1.4.4.

R1 sees that on VLAN 146 it is adjacent with the DR 155.1.146.1, with a cost of 1.

```
Rack1R1#show ip ospf database router 150.1.1.1 self-originate | begin Area 1
```

```
Router Link States (Area 1)
```

```
LS age: 725
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.1.1
Advertising Router: 150.1.1.1
LS Seq Number: 80000036
Checksum: 0x1F46
Length: 36
Area Border Router
Virtual Link Endpoint
Number of Links: 1

Link connected to: a Transit Network
(Link ID) Designated Router address: 155.1.146.1
(Link Data) Router Interface address: 155.1.146.1
Number of TOS metrics: 0
TOS 0 Metrics: 1
```

R1 asks the DR who it is adjacent with, and finds that 150.1.4.4 is on the local segment.

```
Rack1R1#show ip ospf database network 155.1.146.1
```

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

```
Net Link States (Area 1)
```

```
Routing Bit Set on this LSA
LS age: 749
Options: (No TOS-capability, DC)
LS Type: Network Links
Link State ID: 155.1.146.1 (address of Designated Router)
Advertising Router: 150.1.1.1
LS Seq Number: 80000036
Checksum: 0x73E6
Length: 36
Network Mask: /24
Attached Router: 150.1.1.1
Attached Router: 150.1.4.4
Attached Router: 150.1.6.6
```

This means that when R1 installs 30.0.0.0/16 in the routing table the metric will be 20, from R4's default metric during redistribution, and a forward metric of 1 to reach R4.

```
Rack1R1#show ip route 30.0.0.1
```

```
Routing entry for 30.0.0.0/16
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Last update from 155.1.146.4 on FastEthernet0/0, 00:14:03 ago
  Routing Descriptor Blocks:
  * 155.1.146.4, from 150.1.4.4, 00:14:03 ago, via FastEthernet0/0
    Route metric is 20, traffic share count is 1
```

Now let's compare this normal Type-5 External LSA lookup to a Type-5 External LSA that was translated from a Type-7 NSSA External LSA. In this design SW4 redistributes the route 10.10.10.10/32 into area 3 as a Type-7 NSSA External LSA, and R5 translates it into a Type-5 External LSA as it moves to area 0.

R3 performs a lookup on the external route, and sees the advertising router set to 150.1.5.5 (R5), and the forward address set to 150.1.10.10.

```
Rack1R3#show ip ospf database external 10.10.10.10
```

```
OSPF Router with ID (150.1.30.30) (Process ID 1)
```

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

```
LS age: 463
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 10.10.10.10 (External Network Number )
Advertising Router: 150.1.5.5
LS Seq Number: 80000003
Checksum: 0x1514
Length: 36
Network Mask: /32
  Metric Type: 2 (Larger than any link state path)
  TOS: 0
  Metric: 20
  Forward Address: 150.1.10.10
  External Route Tag: 0
```

Since the forward address is non-zero, the next recursive lookup is performed towards 150.1.10.10, instead of the advertising router 150.1.5.5.

```
Rack1R3#show ip route 150.1.10.10
```

```
% Subnet not in table
```

The problem with this design however, is that the prefix 150.1.10.10 was filtered out from the previous Type-3 LSA Filter. The result is that recursion towards the external route fails, and it cannot be installed in the routing table.

```
Rack1R3#show ip route 10.10.10.10
% Subnet not in table
```

To resolve this problem a very specific feature can be configured on R5, which is known as OSPF Forwarding Address Suppression in Translated Type-5 LSAs. This feature, configured by adding the **translate type7 suppress-fa** argument onto the **area 3 nssa** statement, instructs the ABR to not preserve the value in the forward address field as a Type-7 NSSA External LSA is translated into a Type-5 External LSA.

```
Rack1R5#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R5(config)#router ospf 1
Rack1R5(config-router)#area 3 nssa no-redistribution no-summary
translate type7 suppress-fa
Rack1R5(config-router)#end
Rack1R5#
```

```
Rack1R3#show ip ospf database external 10.10.10.10
```

```
OSPF Router with ID (150.1.30.30) (Process ID 1)
```

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

```
Routing Bit Set on this LSA
LS age: 36
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 10.10.10.10 (External Network Number )
Advertising Router: 150.1.5.5
LS Seq Number: 80000004
Checksum: 0x7F54
Length: 36
Network Mask: /32
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

The result of this configuration is that R3 now sees the forward address for the Type-5 External LSA 10.10.10.10/32 set to 0.0.0.0, which means a lookup must now be performed on the advertising router 150.1.5.5.



In this particular case R3 finds that 150.1.5.5 is directly connected with a metric of 781 via a point-to-point adjacency.

```
Rack1R3#show ip ospf database router 150.1.30.30 self-originate
```

```
OSPF Router with ID (150.1.30.30) (Process ID 1)
```

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

```
LS age: 1098
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.30.30
Advertising Router: 150.1.30.30
LS Seq Number: 8000001D
Checksum: 0x5668
Length: 60
Area Border Router
AS Boundary Router
Number of Links: 3
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 150.1.5.5
(Link Data) Router Interface address: 155.1.0.3
Number of TOS metrics: 0
TOS 0 Metrics: 781
```

The final result is that the external route is installed with a metric of 20, which derives from SW4's default redistribution metric, plus a forward metric of 781 to reach R5.

```
Rack1R3#show ip route 10.10.10.10
```

```
Routing entry for 10.10.10.10/32
```

```
Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 781
```

```
Last update from 155.1.0.5 on Serial1/0, 00:00:57 ago
```

```
Routing Descriptor Blocks:
```

```
* 155.1.0.5, from 150.1.5.5, 00:00:57 ago, via Serial1/0
```

```
Route metric is 20, traffic share count is 1
```

Although this feature fixes the problem introduced by the Type-3 LSA Filter, suboptimal routing may be introduced when there are multiple exit points out of the NSSA. As previously discussed, normally the Type-7 to 5 translator election and the forward address calculation are kept separate, which means the control plane advertisement of the route does not need to follow the traffic forwarding plane, but with forwarding address suppression the traffic will always flow through the Type-7 to 5 translator.

This can be seen with the previously configured case regarding SW3's Loopback interface 9.9.9.9/32 that was redistributed into area 2. R5 sees the external LSA 9.9.9.9/32 with an advertising router of 150.1.30.30 – R3, the Type-7 to 5 translator – and a forward address of 150.1.9.9.

```
Rack1R5#show ip ospf database external 9.9.9.9
```

```

      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: 436
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 9.9.9.9 (External Network Number )
Advertising Router: 150.1.30.30
LS Seq Number: 8000001D
Checksum: 0xAD35
Length: 36
Network Mask: /32
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 20
    Forward Address: 150.1.9.9
    External Route Tag: 0

```

The lookup towards the forward address 150.1.9.9 results in using R6 as an exit point. This means that although R3 is advertising the route into area 0, the traffic flows through R6.

```
Rack1R5#show ip route 150.1.9.9
```

```

Routing entry for 150.1.9.9/32
  Known via "ospf 1", distance 110, metric 68, type inter area
  Last update from 155.1.0.1 on Serial0/0, 00:23:35 ago
  Routing Descriptor Blocks:
  * 155.1.0.1, from 150.1.6.6, 00:23:35 ago, via Serial0/0
    Route metric is 68, traffic share count is 1

```

```
Rack1R5#traceroute 9.9.9.9
```

```
Type escape sequence to abort.
Tracing the route to 9.9.9.9
```

```
 1 155.1.0.1 28 msec 28 msec 28 msec
 2 155.1.146.6 28 msec 32 msec 28 msec
 3 155.1.67.7 32 msec 32 msec 28 msec
 4 155.1.79.9 32 msec * 28 msec
```

Now let's see what happens when R3, the Type-7 to 5 translator, suppresses the forward address field.

```
Rack1R3#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R3(config)#router ospf 1
Rack1R3(config-router)#area 2 nssa no-summary translate type7 suppress-fa
```

```
Rack1R5#show ip ospf database external 9.9.9.9
```

```
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: 10
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 9.9.9.9 (External Network Number )
Advertising Router: 150.1.30.30
LS Seq Number: 8000001E
Checksum: 0x3358
Length: 36
Network Mask: /32
Metric Type: 2 (Larger than any link state path)
TOS: 0
Metric: 20
Forward Address: 0.0.0.0
External Route Tag: 0
```

R5 now sees the forward address as 0.0.0.0, which means the recursive lookup should be performed on the advertising router 150.1.30.30. The result is that the traffic flows through R3, the Type-7 to 5 translator, which technically is sub-optimal in this design.

```
Rack1R5#traceroute 9.9.9.9
```

```
Type escape sequence to abort.
Tracing the route to 9.9.9.9
```

```
 1 155.1.0.3 32 msec 28 msec 28 msec
 2 155.1.37.7 32 msec 28 msec 32 msec
 3 155.1.79.9 32 msec * 28 msec
```

## 6.34 OSPF Default Routing

- Configure R6 with a static default route pointing to BB1.
- Configure R6 to originate an external type-2 default route with a metric of 60 into OSPF as long as the static default route is installed in the routing table.
- Do not use a route-map to accomplish this.
- Configure R4 to originate an external type-1 default route with a metric of 40 regardless of whether it has a default route already installed in the routing table.

### Configuration

---

```
R4:
router ospf 1
  default-information originate always metric 40 metric-type 1
```

```
R6:
ip route 0.0.0.0 0.0.0.0 54.1.1.254
!
router ospf 1
  default-information originate metric 60
```

### Verification

---

#### Note

Default routing for non-stub areas in OSPF is accomplished through the origination of Type-5 External LSAs via the **default-information originate** command. Without any additional arguments, the OSPF process first checks to see if a default route is installed in the routing table. If a default route is already installed, such as via a static route or learned via BGP, the OSPF default route is originated. If the default route is not found, no origination occurs. This behavior is typically desirable in designs with multiple exit points out of the OSPF domain to upstream networks.

For example, imagine an OSPF network with exit points A and B out to the Internet. Both router A and B are running BGP with upstream peers, and learning a default route via BGP. As long as both device maintain their upstream peerings, a default route can be advertised into OSPF. However if A's link to the upstream neighbor is lost, and hence its default route via BGP is lost, its OSPF default route advertisement is withdrawn. The result of this design is that an individual exit point will only collect default traffic if they themselves have a default exit point to upstream networks.

This behavior can be modified by adding the **always** argument to the **default-information originate** statement, which essentially skips over the checking for a default route already being installed in the table.

The below view of the OSPF database on R1 indicates that both R4 and R6 are originating a default route. Without additional arguments on the command the default route would have been advertised as a Type-2 External route with a metric of 20. The same route lookup logic is applied to these default routes as normal Type-5 External LSAs, where E1 is preferred over E2, and if multiple E2 routes exist with the same metric, the forward metrics are compared.

```
Rack1R1#show ip ospf database external 0.0.0.0
```

```
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: 297
```

```
Options: (No TOS-capability, DC)
```

```
LS Type: AS External Link
```

```
Link State ID: 0.0.0.0 (External Network Number )
```

```
Advertising Router: 150.1.4.4
```

```
LS Seq Number: 80000001
```

```
Checksum: 0xB7B4
```

```
Length: 36
```

```
Network Mask: /0
```

```
  Metric Type: 1 (Comparable directly to link state metric)
```

```
  TOS: 0
```

```
  Metric: 40
```

```
  Forward Address: 0.0.0.0
```

```
  External Route Tag: 1
```

```
LS age: 136
```

```
Options: (No TOS-capability, DC)
```

```
LS Type: AS External Link
```

```
Link State ID: 0.0.0.0 (External Network Number )
```

```
Advertising Router: 150.1.6.6
```

```
LS Seq Number: 80000002
```

```
Checksum: 0xE7EA
```

```
Length: 36
```

```
Network Mask: /0
```

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

```
  TOS: 0
```

```
  Metric: 60
```

```
  Forward Address: 0.0.0.0
```

```
  External Route Tag: 1
```

The conditional checking for the already installed default route on R6 can be verified as seen below. As long as R6 has the static default route installed in the routing table, a default route is originated into OSPF. Once the route is removed, only R4 continues to originate default information.

```
Rack1R6#show ip route static
```

```
S* 0.0.0.0/0 [1/0] via 54.1.1.254
```

```
Rack1R6#show ip ospf database | include _0.0.0.0
```

```
0.0.0.0      150.1.4.4      608      0x80000001 0x00B7B4 1
```

```
0.0.0.0      150.1.6.6      11       0x80000001 0x00E9E9 1
```

```
Rack1R6#config t
```

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

```
Rack1R6(config)#no ip route 0.0.0.0 0.0.0.0 54.1.1.254
```

```
Rack1R6(config)#do show ip ospf database | include _0.0.0.0
```

```
0.0.0.0      150.1.4.4      621      0x80000001 0x00B7B4 1
```

## 6.35 OSPF Conditional Default Routing

- Remove the previous default route advertisement on R6.
- Configure R6 to originate a default route into OSPF, but only if the local Serial interface connecting to BB1 is up.

### Configuration

```
R6:
router ospf 1
 default-information originate always route-map TRACK_LINK_TO_BB1
 !
ip prefix-list LINK_TO_BB1 seq 5 permit 54.1.1.0/24
 !
route-map TRACK_LINK_TO_BB1 permit 10
 match ip address prefix-list LINK_TO_BB1
```

### Verification

#### Note

Conditional default-information origination in OSPF uses a route-map to check for the existence of a specific prefix in the IP routing table before the default route is originated. Recall from the previous section that the default condition is to check for a default route already installed in the IP routing table. In this example the check for the existing default route is circumvented with the **always** keyword added to the **default-information originate** statement.

In this design R6 is configured to check for the prefix 54.1.1.0/24 in the routing table. If the route exists, the default route is originated. As seen below once this prefix is no longer in the routing table, the default advertisement is withdrawn.

```
Rack1R6#show ip ospf database | include _0.0.0.0
0.0.0.0          150.1.4.4          856          0x80000001 0x00B7B4 1
0.0.0.0          150.1.6.6          47           0x80000001 0x009975 1
```

```
Rack1R6#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R6(config)#interface Serial0/0
Rack1R6(config-if)#frame-relay lmi-type q933a
Rack1R6(config-if)#end
Rack1R6#
%SYS-5-CONFIG_I: Configured from console by console
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0, changed
state to down
```

```
Rack1R6#show ip ospf database | include _0.0.0.0
0.0.0.0          150.1.4.4          916          0x80000001 0x00B7B4 1
```

## 6.36 OSPF Reliable Conditional Default Routing

- Remove the previous default route advertisement on R4.
- Configure an IP SLA instance on R4 to check ICMP reachability to BB3 every five seconds.
- R4 should advertise a default route into OSPF, but only if the SLA monitoring reports its status as OK.

### Configuration

---

```
ip sla monitor 1
  type echo protocol ipIcmpEcho 204.12.1.254
  timeout 2000
  frequency 5
ip sla monitor schedule 1 start now life forever
!
track 1 rtr 1
!
ip route 169.254.0.1 255.255.255.255 Null0 track 1
!
ip prefix-list PLACEHOLDER seq 5 permit 169.254.0.1/32
!
route-map TRACK_PLACEHOLDER permit 10
  match ip address prefix-list PLACEHOLDER
!
router ospf 1
  default-information originate always route-map TRACK_PLACEHOLDER
```

### Verification

---

#### Note

In the previous example R6 was configured to track the status of its Frame Relay interface connecting to BB1. When the link was down, and its IP prefix withdrawn from the routing table, the default route origination was withdrawn based on the conditional checking. For this redundancy design to work properly, we would have to assume that the end-to-end reachability status on the Frame Relay circuit to BB1 can be accurately determined by the Layer 2 interface status. However, as we have seen in previous designs before, tracking an interface status directly is *not* a good indication of end-to-end reachability, as the interface could be UP/UP locally, but the circuit itself could be down.



Another example that is common in today's networks is with Metro Ethernet. Since the router's local Ethernet interface only tracks link status to its attached switch, end-to-end reachability can not be inferred by checking this link status. In the output below we see a case where R4 wants to originate a default route only when the link to BB3 is viable. As the prefix 204.12.1.0/24 is currently in the routing table, the default route is originated.

```
R4:
router ospf 1
  default-information originate always route-map TRACK_LINK_TO_BB3
!
ip prefix-list LINK_TO_BB3 seq 5 permit 204.12.1.0/24
!
route-map TRACK_LINK_TO_BB3 permit 10
  match ip address prefix-list LINK_TO_BB3
```

```
Rack1R4#show ip ospf database | include _0.0.0.0
0.0.0.0          150.1.4.4          72          0x80000001 0x00B35F 1
0.0.0.0          150.1.6.6          392         0x80000001 0x009975 1
```

Now we simulate a case where a failure occurs in the network, but not on the segment attached directly to R4. In this example the switchport between SW3 and BB3 is shutdown. This effectively breaks the end-to-end connectivity between R4 and BB3, but leaves the local link status of R4 as UP/UP, which meets the default route condition.

```
Rack1SW3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1SW3(config)#interface Fa0/24
Rack1SW3(config-if)#shutdown
Rack1SW3(config-if)#
```

```
Rack1R4#ping 204.12.1.254
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 204.12.1.254, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
Rack1R4#show ip ospf database | include _0.0.0.0
0.0.0.0          150.1.4.4          96          0x80000001 0x00B35F 1
0.0.0.0          150.1.6.6          416         0x80000001 0x009975 1
```

**Rack1R4#show ip route 204.12.1.0**

```

Routing entry for 204.12.1.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via ospf 1, rip
  Advertised by ospf 1
    rip
  Routing Descriptor Blocks:
  * directly connected, via FastEthernet0/0
    Route metric is 0, traffic share count is 1

```

As we can see, even though R4's uplink is no longer viable to end-to-end connectivity, the default route is still advertised since 204.12.1.0/24 is still in the routing table. To add more intelligence to this design, an IP SLA tracking instance can be used for more reliable information than the layer 2 interface status.

```

Rack1R4(config)#ip sla monitor 1
Rack1R4(config-sla-monitor)#type echo protocol ipIcmpEcho 204.12.1.254
Rack1R4(config-sla-monitor-echo)#timeout 2000
Rack1R4(config-sla-monitor-echo)#frequency 5
Rack1R4(config-sla-monitor-echo)#exit
Rack1R4(config)#ip sla monitor schedule 1 start now life forever
Rack1R4(config)#track 1 rtr 1
Rack1R4(config-track)#exit
Rack1R4(config)#ip route 169.254.0.1 255.255.255.255 null0 track 1
Rack1R4(config)#ip prefix-list PLACEHOLDER permit 169.254.0.1/32
Rack1R4(config)#route-map TRACK_PLACEHOLDER
Rack1R4(config-route-map)#match ip address prefix-list PLACEHOLDER
Rack1R4(config-route-map)#end
Rack1R4#

```

The IP SLA instance is instructed to ping 204.12.1.254 every five seconds, and wait two seconds for an echo-reply. The SLA instance is called from a tracked object, which in turn is called from an arbitrary placeholder static route. As seen below, with the tracked object reporting the status code OK, the default route is originated.

**Rack1R4#show track**

```

Track 1
  Response Time Reporter 1 state
  State is Up
    3 changes, last change 00:00:04
  Latest operation return code: OK
  Latest RTT (milliseconds) 3
  Tracked by:
    STATIC-IP-ROUTING 0

```

**Rack1R4#show ip ospf database | include \_0.0.0.0**

```

0.0.0.0          150.1.4.4          4          0x80000001 0x00B35F 1
0.0.0.0          150.1.6.6          806        0x80000001 0x009975 1

```

Now the same indirect failure scenario is simulated as previously seen. Specifically SW3's link to BB3 is disabled, while R4's local link stays up.

```
Rack1SW3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1SW3(config)#interface Fa0/24
Rack1SW3(config-if)#shutdown
Rack1SW3(config-if)#
```

R4 detects the failure of the SLA instance, and the tracked object transitions to down.

```
Rack1R4#
%TRACKING-5-STATE: 1 rtr 1 state Up->Down
Rack1R4#show track
Track 1
  Response Time Reporter 1 state
  State is Down
    4 changes, last change 00:00:04
  Latest operation return code: Timeout
  Tracked by:
    STATIC-IP-ROUTING 0
```

The failure of the tracked object causes the static route to be withdrawn. Since the route-map condition for the default-information originate statement is looking for this prefix to be installed in the routing table, the default route is withdrawn when 169.254.0.1 is withdrawn.

```
Rack1R4#show ip route 169.254.0.1
% Network not in table

Rack1R4#show ip ospf database | include _0.0.0.0
0.0.0.0          150.1.6.6      785           0x80000001 0x009975 1
```

## 6.37 OSPF Filtering with Distribute-Lists

- Configure distribute-list filtering on R5, SW2, and SW4 so that these devices do not install routes to the Loopback0 networks of R1 and R2.

### Configuration

---

```
R5:
router ospf 1
 distribute-list 1 in
!
access-list 1 deny 150.1.1.1
access-list 1 deny 150.1.2.2
access-list 1 permit any
```

```
SW2:
router ospf 1
 distribute-list 1 in
!
access-list 1 deny 150.1.1.1
access-list 1 deny 150.1.2.2
access-list 1 permit any
```

```
SW4:
router ospf 1
 distribute-list 1 in
!
access-list 1 deny 150.1.1.1
access-list 1 deny 150.1.2.2
access-list 1 permit any
```

### Verification

---

#### Note

Recall that to properly compute SPF, all routers with an OSPF area must agree on their view of the database. This implies that OSPF filtering in the database can be accomplished between areas, but not within an area. Inter-area filtering has been previously demonstrated with stub areas, and the Type-3 LSA Filter. Intra-area filtering *can* be accomplished in OSPF with an inbound distribute-list, however this filtering only affects the local routing table, not the OSPF database.

Once the distribute-list has been applied on R5, the routes 150.1.2.2/32 and 150.1.1.1/32 no longer appear in the routing table.

```
Rack1R5#show ip route | include 150.
    150.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
C       150.1.5.0/24 is directly connected, Loopback0
O       150.1.10.10/32 [110/3] via 155.1.58.8, 00:01:42, FastEthernet0/0
O IA    150.1.9.9/32 [110/67] via 155.1.0.3, 00:01:42, Serial0/0
O       150.1.8.8/32 [110/2] via 155.1.58.8, 00:01:42, FastEthernet0/0
O IA    150.1.7.7/32 [110/66] via 155.1.0.3, 00:01:42, Serial0/0
O IA    150.1.6.6/32 [110/66] via 155.1.45.4, 00:01:42, Serial0/1
O       150.1.4.4/32 [110/65] via 155.1.45.4, 00:01:42, Serial0/1
O       150.1.3.3/32 [110/65] via 155.1.0.3, 00:01:42, Serial0/0
```

Viewing the OSPF database on SW4, however, indicates that the Type-3 (Inter-Area) LSA's 150.1.1.1 and 150.1.2.2 still exist.

```
Rack1SW4#show ip ospf database summary 150.1.1.1

          OSPF Router with ID (150.1.10.10) (Process ID 1)

          Summary Net Link States (Area 3)

Routing Bit Set on this LSA
LS age: 181
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 150.1.1.1 (summary Network Number)
Advertising Router: 150.1.5.5
LS Seq Number: 80000003
Checksum: 0xB608
Length: 28
Network Mask: /32
          TOS: 0   Metric: 65
```

### Pitfall

This type of design can result in traffic black holes unless implemented carefully. If R5 was configured with the distribute-list, but SW2 and SW4 were not, traffic from SW4 towards 150.1.1.1 would be sent to SW2, from SW2 to R5, and then black holed on R5. When implementing inbound distribute-list filtering ensure that all routers still agree on the forwarding paths in the network.

### 6.38 OSPF Summarization and Discard Routes

- Configure R5 to advertise the summary 150.X.0.0/22 into area 3.
- Ensure that R5, SW2, and SW4 can still reach the Loopback0 networks of R1 and R2, even with the previously configured distribute-list filter applied.

#### Configuration

```
R5:
router ospf 1
 no discard-route internal
 area 0 range 150.1.0.0 255.255.252.0
```

#### Verification

##### Note

When summarization is configured in OSPF, similar to EIGRP and BGP, a matching route to Null0 for the summary is installed locally in the routing table. This “discard” route is used to prevent the forwarding of traffic towards a shorter match, such as a default route, if no specific route towards the actual destination exists in the network.

The automatic origination of the discard route can be disabled with the **no discard-route [internal | external]**, where *internal* refers to inter-area summarization performed with the **area range** command, while *external* refers to redistributed summarization performed with the **summary-address** command. The operation of the discard route can be illustrated as follows.

Per the previous task, R5, SW2, and SW4 have the prefixes 150.1.1.1/32 and 150.1.2.2/32 filtered out of the routing table with a distribute-list. Additionally, R5 is originating the summary 150.1.0.0/22 into area 3, which encompasses addresses 150.1.0.0 through 150.1.3.255.

```
Rack1SW2#show ip route | include 150.
      150.1.0.0/16 is variably subnetted, 8 subnets, 3 masks
O       150.1.10.10/32 [110/2] via 155.1.108.10, 00:06:54, Port-channel1
O IA    150.1.9.9/32 [110/68] via 155.1.58.5, 00:06:54, Vlan58
O IA    150.1.7.7/32 [110/67] via 155.1.58.5, 00:06:54, Vlan58
O IA    150.1.6.6/32 [110/67] via 155.1.58.5, 00:06:54, Vlan58
O       150.1.5.5/32 [110/2] via 155.1.58.5, 00:06:54, Vlan58
O IA    150.1.4.4/32 [110/66] via 155.1.58.5, 00:06:54, Vlan58
O IA    150.1.0.0/22 [110/66] via 155.1.58.5, 00:00:41, Vlan58
C       150.1.8.0/24 is directly connected, Loopback0
```

Reachability to the network 150.1.3.3 is obtained from SW2, but reachability to 150.1.1.1 is not.

```
Rack1SW2#ping 150.1.3.3
```

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

```
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 58/60/67 ms
```

```
Rack1SW2#ping 150.1.1.1
```

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

```
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:
```

```
U.U.U
```

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

This is due to the fact that R5's longest match to 150.1.1.1 is the discard route via Null0. Although R5 does have a valid default path to 150.1.1.1 via R4, this cannot be used because /22 is a longer match than /0.

```
Rack1R5#show ip route 150.1.1.1
```

```
Routing entry for 150.1.0.0/22
```

```
Known via "ospf 1", distance 110, metric 0, type intra area
```

```
Routing Descriptor Blocks:
```

```
* directly connected, via Null0
```

```
Route metric is 0, traffic share count is 1
```

```
Rack1R5#show ip route 0.0.0.0
```

```
Routing entry for 0.0.0.0/0, supernet
```

```
Known via "ospf 1", distance 110, metric 1, candidate default path
```

```
Tag 1, type extern 2, forward metric 64
```

```
Last update from 155.1.0.4 on Serial0/0, 00:14:54 ago
```

```
Routing Descriptor Blocks:
```

```
* 155.1.45.4, from 150.1.4.4, 00:14:54 ago, via Serial0/1
```

```
Route metric is 1, traffic share count is 1
```

```
Route tag 1
```

```
155.1.0.4, from 150.1.4.4, 00:14:54 ago, via Serial0/0
```

```
Route metric is 1, traffic share count is 1
```

```
Route tag 1
```

With the discard route removed on R5, the longest match to 150.1.1.1 is now 0.0.0.0/0.

```
Rack1R5(config)#router ospf 1
Rack1R5(config-router)#no discard-route internal
Rack1R5(config-router)#end
Rack1R5#
%SYS-5-CONFIG_I: Configured from console by console

Rack1R5#show ip route 150.1.1.1
% Subnet not in table

Rack1R5#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "ospf 1", distance 110, metric 1, candidate default path
  Tag 1, type extern 2, forward metric 64
  Last update from 155.1.0.4 on Serial0/0, 00:00:09 ago
  Routing Descriptor Blocks:
  * 155.1.45.4, from 150.1.4.4, 00:00:09 ago, via Serial0/1
    Route metric is 1, traffic share count is 1
    Route tag 1
  155.1.0.4, from 150.1.4.4, 00:00:09 ago, via Serial0/0
    Route metric is 1, traffic share count is 1
    Route tag 1
```

The final result is that SW2 uses the 150.1.0.0/22 prefix to route traffic for 150.1.1.1 towards R5, while R5 uses the 0.0.0.0/0 prefix to route the traffic towards R4.

```
Rack1SW2#ping 150.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:
!!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 58/60/67 ms

Rack1SW2#traceroute 150.1.1.1

Type escape sequence to abort.
Tracing the route to 150.1.1.1

 1 155.1.58.5 0 msec 9 msec 0 msec
 2 155.1.0.4 33 msec 34 msec 25 msec
 3 155.1.146.1 34 msec * 25 msec
```



## 6.39 OSPF Filtering with Administrative Distance

- Configure administrative distance filtering on R5 so that traffic going towards the VLAN 67 network is sent towards R3.

### Configuration

---

```
R5:
router ospf 1
  distance 255 150.1.6.6 0.0.0.0 67
!
access-list 67 permit 155.1.67.0
```

### Verification

---

#### Note

Like the other routing protocols, administrative distance can be changed on a per-prefix and per-neighbor basis in OSPF. One key difference however, is the address field in the **distance** command refers to the originator of the prefix into the area, not necessarily the neighbor you are learning the route from.

In the below output we can see that R5 has two routes to the prefix 155.1.67.0/24. Both are reachable out Serial0/0, one via the next hop 155.1.0.1 (R1), and one via the next-hop 155.1.0.1 (R3). Note however, the difference between the next-hop values and the “from” field.

The prefix reachable via 155.1.0.3 is from 150.1.3.3, which is the router-id of R3. The prefix reachable via 155.1.0.1 is from 150.1.6.6, which is the router-id of R6. This means that R3 is using itself as the ABR exit point to the inter-area destination, while R1 is using R6 as the ABR. Therefore to change the distance of one of these prefixes over the other, the from 150.1.3.3 or 150.1.6.6 fields should be matched in the **distance** command.

```
Rack1R5#show ip route 155.1.67.0
Routing entry for 155.1.67.0/24
  Known via "ospf 1", distance 110, metric 66, type inter area
  Last update from 155.1.0.1 on Serial0/0, 00:06:18 ago
  Routing Descriptor Blocks:
  * 155.1.0.3, from 150.1.3.3, 00:06:18 ago, via Serial0/0
    Route metric is 66, traffic share count is 1
    155.1.0.1, from 150.1.6.6, 00:06:18 ago, via Serial0/0
      Route metric is 66, traffic share count is 1
```

```
Rack1R5#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R5(config)#access-list 67 permit 155.1.67.0
Rack1R5(config)#router ospf 1
Rack1R5(config-router)#distance 255 150.1.6.6 0.0.0.0 67
Rack1R5(config-router)#end

Rack1R5#show ip route 155.1.67.0
Routing entry for 155.1.67.0/24
  Known via "ospf 1", distance 110, metric 66, type inter area
  Last update from 155.1.0.3 on Serial0/0, 00:00:03 ago
  Routing Descriptor Blocks:
    * 155.1.0.3, from 150.1.3.3, 00:00:03 ago, via Serial0/0
      Route metric is 66, traffic share count is 1
```

Once the distance has be modified the only route left is the one via R3, hence the modified traffic flow as seen below.

```
Rack1R5#traceroute 155.1.67.6
```

```
Type escape sequence to abort.
Tracing the route to 155.1.67.6
```

```
 1 155.1.0.3 32 msec 28 msec 28 msec
 2 155.1.37.7 33 msec 32 msec 32 msec
 3 155.1.67.6 28 msec * 28 msec
```

## 6.40 OSPF Filtering with Route-Maps

- Configure route-map filtering on R6 so that traffic going towards the Loopback0 network of R3 is sent towards R1.

### Configuration

```
R6:
router ospf 1
  distribute-list route-map DENY_R3_LOOPBACK_FROM_R4 in
!
access-list 3 permit 150.1.3.3
access-list 4 permit 155.1.146.4
!
route-map DENY_R3_LOOPBACK_FROM_R4 deny 10
  match ip address 3
  match ip next-hop 4
!
route-map DENY_R3_LOOPBACK_FROM_R4 permit 20
```

### Verification

#### Note

Referencing a route-map with a distribute-list in OSPF extends the filtering capability with additional match criteria. Specifically the matches of interface (outgoing interface in the routing table), ip address, ip next-hop, ip route-source (router-id of the originating router), metric, route-type (intra-area, inter-area, etc.), and tag are supported inside the route-map.

Note the limitation of this feature is that only inbound filtering is supported, and the filter still is only local to the router's routing table, i.e. the filter does not affect the OSPF database advertisements.

In the below output we can see that R6 has two routes to the prefix 150.1.3.3/32. One via 155.1.146.4 via the ABR 150.1.3.3, and one via 155.1.146.1 via the ABR 150.1.3.3. Both are reachable out the same interface.

```
Rack1R6#show ip route 150.1.3.3
Routing entry for 150.1.3.3/32
  Known via "ospf 1", distance 110, metric 130, type intra area
  Redistributing via eigrp 10
  Advertised by eigrp 10 metric 100000 1000 255 1 1500
  Last update from 155.1.146.1 on FastEthernet0/0.146, 00:00:02 ago
  Routing Descriptor Blocks:
    155.1.146.4, from 150.1.3.3, 00:00:02 ago, via FastEthernet0/0.146
      Route metric is 130, traffic share count is 1
    * 155.1.146.1, from 150.1.3.3, 00:00:02 ago, via FastEthernet0/0.146
      Route metric is 130, traffic share count is 1
```

With a normal distribute-list filter referencing an access-list, there would be no way to distinguish between these two prefixes since they have the same address. Likewise with an administrative distance filter, there is no way to distinguish between them because they both are via the same ABR. By matching both the address and the next-hop, one prefix can be discarded in favor of the other.

```
Rack1R6#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R6(config)#router ospf 1
Rack1R6(config-router)# distribute-list route-map
DENY_R3_LOOPBACK_FROM_R4 in
Rack1R6(config-router)#end

Rack1R6#clear ip route *
Rack1R6#show ip route 150.1.3.3
Routing entry for 150.1.3.3/32
  Known via "ospf 1", distance 110, metric 130, type intra area
  Redistributing via eigrp 10
  Advertised by eigrp 10 metric 100000 1000 255 1 1500
  Last update from 155.1.146.1 on FastEthernet0/0.146, 00:00:02 ago
  Routing Descriptor Blocks:
  * 155.1.146.1, from 150.1.3.3, 00:00:02 ago, via FastEthernet0/0.146
    Route metric is 130, traffic share count is 1

Rack1R6#ping 150.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/101/105 ms
```

The result in this case is that traffic no longer follows the path via R4, and is sent only to R1.

```
Rack1R6#traceroute 150.1.3.3

Type escape sequence to abort.
Tracing the route to 150.1.3.3

 0 155.1.146.1 4 msec 0 msec 0 msec
 1 155.1.0.5 32 msec 28 msec 32 msec
 2 155.1.0.3 61 msec * 57 msec
```

## 6.41 OSPF NSSA ABR External Prefix Filtering

- Disable R6's link to VLAN 146 and Loopback0.
- Configure area 2 as an NSSA.
- Configure summarization on R3 so that devices outside of area 2 do not have a route to the network 200.0.0.0/24.
- This filter should not affect any other prefixes.

### Configuration

---

```
R3:
router ospf 1
 area 2 nssa
 summary-address 200.0.0.0 255.255.255.0 not-advertise
```

```
R6:
interface FastEthernet0/0.146
 shutdown
!
interface Loopback0
 shutdown
!
router ospf 1
 area 2 nssa
```

```
SW1:
router ospf 1
 area 2 nssa
```

```
SW3:
router ospf 1
 area 2 nssa
```

### Verification

---

#### Note

As previously discussed, only one ABR exit point out of an NSSA converts a Type-7 NSSA External LSA to a Type-5 External LSA, based on the translator election process. The NSSA translator can be configured to suppress the origination of the Type-5 External LSA into area 0 via the **summary-address** command.

As seen below, R5 currently learns the prefix 200.0.0.0/24 as a Type-5 External route, originated by the NSSA translator R3.

```
Rack1R5#show ip route | include E2(.* )via 155.1.0.3
O E2 200.0.0.0/24 [110/20] via 155.1.0.3, 00:00:37, Serial0/0
O E2 54.1.1.0 [110/20] via 155.1.0.3, 00:06:56, Serial0/0
O E2 200.0.1.0/24 [110/20] via 155.1.0.3, 00:06:56, Serial0/0
O E2 200.0.2.0/24 [110/20] via 155.1.0.3, 00:06:56, Serial0/0
O E2 200.0.3.0/24 [110/20] via 155.1.0.3, 00:06:56, Serial0/0
```

To suppress the advertisement of this prefix into area 0, R3 configures a **summary-address** with an identical mask of the original NSSA external route, but adds the **not-advertise** argument to the summary.

```
Rack1R3#config t
Enter configuration commands, one per line. End with CNTL/Z.
Rack1R3(config)#router ospf 1
Rack1R3(config-router)#summary-address 200.0.0.0 255.255.255.0 not-
advertise
Rack1R3(config-router)#end
Rack1R3#
```

```
Rack1R5#show ip route | include E2(.* )via 155.1.0.3
O E2 54.1.1.0 [110/20] via 155.1.0.3, 00:07:13, Serial0/0
O E2 200.0.1.0/24 [110/20] via 155.1.0.3, 00:07:13, Serial0/0
O E2 200.0.2.0/24 [110/20] via 155.1.0.3, 00:07:13, Serial0/0
O E2 200.0.3.0/24 [110/20] via 155.1.0.3, 00:07:13, Serial0/0
```

The key difference with this filtering technique, versus the previously seen distribute-lists and administrative distance filters, is that the prefix is filtered out of the *database*, not just the routing table. This can be verified by R5's lack of LSA information as seen below.

```
Rack1R5#show ip ospf database external 200.0.0.0

                OSPF Router with ID (150.1.5.5) (Process ID 1)
Rack1R5#
```

## 6.42 OSPF Database Filtering

- Configure SW1 so that SW3 cannot learn any OSPF routes from SW1, but SW1 can still learn OSPF routes from SW3.
- Configure R5 so that R2 cannot learn from R5, but R5 can still learn OSPF routes from R2.

### Configuration

```
R5:
router ospf 1
 neighbor 155.1.0.2 database-filter all out
```

```
SW1:
interface Vlan79
 ip ospf database-filter all out
```

### Verification

#### Note

The OSPF command **database-filter all out** is similar in operation to the **passive-interface** command in RIPv2. This feature allows the formation of OSPF neighbors, since hello packets are not filtered out, but stops the advertisements of all LSAs out the interface or to the neighbor in question.

As seen below, SW1 learns prefixes from SW3, however SW3 has no LSAs in the database, with the exception of locally originated ones.

```
Rack1SW1#show ip route ospf | include Vlan79
```

```
O      155.1.9.0/24 [110/2] via 155.1.79.9, 00:00:54, Vlan79
O      150.1.9.9/32 [110/2] via 155.1.79.9, 00:00:54, Vlan79
```

```
Rack1SW3#show ip route ospf
```

```
Rack1SW3#show ip ospf database
```

```
OSPF Router with ID (150.1.9.9) (Process ID 1)
```

```
Router Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.9.9	150.1.9.9	140	0x80000002	0x006914	3

```
Net Link States (Area 2)
```

Link ID	ADV Router	Age	Seq#	Checksum
155.1.79.9	150.1.9.9	140	0x80000001	0x00DA6B

Likewise R5 learns LSAs from R2, but R2 cannot learn any LSAs from R5.

**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, 00:00:30, FastEthernet0/0
```

**Rack1R5#show ip route ospf | include via 155.1.0.2**

```
O E2    51.51.51.51 [110/20] via 155.1.0.2, 00:00:41, Serial0/0
O IA    155.1.23.0/24 [110/128] via 155.1.0.2, 00:00:51, Serial0/0
O       155.1.0.2/32 [110/64] via 155.1.0.2, 00:00:51, Serial0/0
O IA    192.10.1.0/24 [110/65] via 155.1.0.2, 00:00:51, Serial0/0
O       150.1.2.2/32 [110/65] via 155.1.0.2, 00:00:51, Serial0/0
```

**Rack1R2#show ip ospf database**

OSPF Router with ID (150.1.2.2) (Process ID 1)

Router Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
150.1.2.2	150.1.2.2	31	0x80000002	0x0041DE	3

Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
155.1.23.0	150.1.2.2	32	0x80000001	0x00ADFF
192.10.1.0	150.1.2.2	32	0x80000001	0x00D8FB

Summary ASB Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
150.1.3.3	150.1.2.2	32	0x80000001	0x009F23
192.10.1.254	150.1.2.2	22	0x80000001	0x00D4FF

<output omitted>



## 6.43 OSPF Stub Router Advertisement

- Configure R4 to advertise the maximum metric value to all neighbors inside its Type-1 Router LSA.

### Configuration

```
R4:
router ospf 1
 max-metric router-lsa
```

### Verification

#### Note

The OSPF Stub Router Advertisement feature, not to be confused with OSPF stub *areas*, is used to prevent traffic black holes due to device adds or removes from the OSPF topology. Essentially this feature causes the router to advertise a maximum metric for non-stub destinations, making it the worst cost path to all destinations. The result is that upon initializing the OSPF process, transit traffic will not flow through the stub router unless it is the only possible path. Once the routing domain is fully converged, the max metric value can be withdrawn, allowing normal forwarding to occur through the device.

The `max-metric router-lsa` syntax unconditionally advertises the maximum metric until the command is removed, while the `max-metric router-lsa on-startup wait-for-bgp` option causes the router to advertise the maximum metric until BGP keepalives are received from all neighbors (keepalives in BGP indicate convergence is complete), and the `max-metric router-lsa on-startup announce-time` controls how long the router should advertise the maximum metric after a reload.

In the below output we can see that R5 learns the prefix 155.1.146.0/24 from R4 via multiple links, all with a metric of 65.

```
Rack1R5#show ip route 155.1.146.0
Routing entry for 155.1.146.0/24
  Known via "ospf 1", distance 110, metric 65, type intra area
  Last update from 155.1.0.4 on Serial0/0, 00:00:08 ago
  Routing Descriptor Blocks:
    155.1.45.4, from 150.1.4.4, 00:00:08 ago, via Serial0/1
      Route metric is 65, traffic share count is 1
    155.1.0.4, from 150.1.4.4, 00:00:08 ago, via Serial0/0
      Route metric is 65, traffic share count is 1
  * 155.1.0.1, from 150.1.4.4, 00:00:08 ago, via Serial0/0
      Route metric is 65, traffic share count is 1
```

```
Rack1R5#show ip ospf database router 150.1.4.4

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

          Router Link States (Area 0)

Routing Bit Set on this LSA
LS age: 19
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.4.4
Advertising Router: 150.1.4.4
LS Seq Number: 80000014
Checksum: 0x238B
Length: 96
AS Boundary Router
Number of Links: 6

  Link connected to: a Transit Network
    (Link ID) Designated Router address: 155.1.146.4
    (Link Data) Router Interface address: 155.1.146.4
      Number of TOS metrics: 0
        TOS 0 Metrics: 1
<output omitted>
```

R4 currently advertises a metric of 1 to its transit link of VLAN 146.

After the max-metric router-lsa command is configured, R4 advertises the maximum cost of 65535 for VLAN 146, essentially making it the worst cost path.

```
Rack1R4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rack1R4(config)#router ospf 1
Rack1R4(config-router)#max-metric router-lsa
Rack1R4(config-router)#end
Rack1R4#
```

```
Rack1R5#show ip ospf database router 150.1.4.4
```

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

                Router Link States (Area 0)

Routing Bit Set on this LSA
LS age: 10
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 150.1.4.4
Advertising Router: 150.1.4.4
LS Seq Number: 80000015
Checksum: 0xF935
Length: 96
AS Boundary Router
Number of Links: 6

    Link connected to: a Transit Network
      (Link ID) Designated Router address: 155.1.146.4
      (Link Data) Router Interface address: 155.1.146.4
      Number of TOS metrics: 0
      TOS 0 Metrics: 65535
```

The final result is that R5 no longer uses R4 to reach this destination.

```
Rack1R5#show ip route 155.1.146.0
Routing entry for 155.1.146.0/24
  Known via "ospf 1", distance 110, metric 65, type intra area
  Last update from 155.1.0.1 on Serial0/0, 00:00:04 ago
  Routing Descriptor Blocks:
    * 155.1.0.1, from 150.1.4.4, 00:00:04 ago, via Serial0/0
      Route metric is 65, traffic share count is 1
```

## 6.44 OSPF Interface Timers

- Configure R5 to send OSPF hello packets out its link to the Frame Relay network every five seconds, and wait for seven seconds before declaring a neighbor down.
- Configure R4 to send OSPF hello packets every 250ms on the point-to-point link to R5.
- Configure R4 to send OSPF hello packets every 250ms on the point-to-point link to R5.

### **Configuration**

---

```
R1:
interface Serial0/0
 ip ospf hello-interval 5
 ip ospf dead-interval 7
```

```
R2:
interface Serial0/0
 ip ospf hello-interval 5
 ip ospf dead-interval 7
```

```
R3:
interface Serial1/0
 ip ospf hello-interval 5
 ip ospf dead-interval 7
```

```
R4:
interface Serial0/0
 ip ospf hello-interval 5
 ip ospf dead-interval 7
!
interface Serial0/1
 ip ospf dead-interval minimal hello-multiplier 4
```

```
R5:
interface Serial0/0
 ip ospf hello-interval 5
 ip ospf dead-interval 7
!
interface Serial0/1
 ip ospf dead-interval minimal hello-multiplier 4
```

**Verification** **Note**

OSPF hello and dead timers must match for adjacency to occur. When the `ip ospf hello-interval` command is modified without the `ip ospf dead-interval` command, the dead timer is automatically set to be four times the configured hello.

With the OSPF sub-second hello feature, the dead timer is set to one second, and the hello interval is set to 1000ms/hello-multiplier. Per the above command `ip ospf dead-interval minimal hello-multiplier 4`, the dead interval is one second, and the hello interval is 250ms (four times per second).

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
150.1.4.4	0	FULL/ -	944 msec	155.1.45.4	Serial0/1
150.1.4.4	0	FULL/ -	00:00:02	155.1.0.4	Serial0/0
150.1.1.1	0	FULL/ -	00:00:02	155.1.0.1	Serial0/0
150.1.2.2	0	FULL/ -	00:00:05	155.1.0.2	Serial0/0
150.1.3.3	0	FULL/ -	00:00:05	155.1.0.3	Serial0/0
150.1.8.8	1	FULL/DR	00:00:38	155.1.58.8	FastEthernet0/0

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

```
Serial0/1 is up, line protocol is up
  Internet Address 155.1.45.5/24, Area 0
  Process ID 1, Router ID 150.1.5.5, Network Type POINT_TO_POINT, Cost: 64
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 250 msec, Dead 1, Wait 1, Retransmit 5
    oob-resync timeout 40
    Hello due in 22 msec
<output omitted>
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 POINT_TO_MULTIPOINT, Cost: 64
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT,
  Timer intervals configured, Hello 5, Dead 7, Wait 7, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:03
<output omitted>
```

## 6.45 OSPF Global Timers

- Modify R4 and R5's OSPF timers as follows:
  - Configure SPF throttling to start new re-calculation at least 100ms after a new LSA arrives. The second SPF calculation should occur in no less than 1s after the first one, and the maximum wait time should be no more than 10 seconds.
  - Configure LSA pacing to wait at least 50ms between consecutive link-state updates. LSA retransmissions should be paced at least 75ms apart.
  - Configure LSA throttling to generate subsequent LSAs after 10ms, to wait at least 4 seconds to generate the next LSA, and no more than 6 seconds between generation of the same LSA.
  - Configure LSA arrival throttling to wait 2 seconds between reception of the same LSA from a neighbor.
  - of the re-calculation at least 100ms after a new LSA arrives. The second SPF calculation should occur in no less than 1s after the first one, and the maximum wait time should be no more than 10 seconds.
- Configure R4 and R5 to assume that LSA transmission takes 2 seconds on the point-to-point link between them. LSA retransmission should occur if an acknowledgement is not received within 10 seconds over this link.

### Configuration

---

```
R4:
router ospf 1
  timers throttle spf 100 1000 10000
  timers pacing flood 50
  timers pacing retransmission 75
  timers throttle lsa all 10 4000 6000
  timers lsa arrival 2000
!
interface Serial0/1
  ip ospf transmit-delay 2
  ip ospf retransmit-interval 10
```

```
R5:
router ospf 1
  timers throttle spf 100 1000 10000
  timers pacing flood 50
  timers pacing retransmission 75
  timers throttle lsa all 10 4000 6000
  timers lsa arrival 2000
!
interface Serial0/1
  ip ospf transmit-delay 2
  ip ospf retransmit-interval 10
```

## Verification

### Note

OSPF packet and SPF pacing/throttling timers control how fast OSPF responds to convergence events. In the majority of deployments the default values should not need modification. Within the scope of the CCIE Lab Exam, determining which timers control which events should be self explanatory based on the usage guidelines of the commands in the OSPF command reference section of the documentation. These timers can be verified as follows.

#### **Rack1R5#show ip ospf**

```
Routing Process "ospf 1" with ID 150.1.5.5
Start time: 00:00:31.038, Time elapsed: 01:30:15.971
Supports only single TOS(TOS0) routes
Supports opaque LSA
Supports Link-local Signaling (LLS)
Supports area transit capability
It is an area border router
Router is not originating router-LSAs with maximum metric
Initial SPF schedule delay 100 msec
Minimum hold time between two consecutive SPF's 1000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Incremental-SPF disabled
Initial LSA throttle delay 10 msec
Minimum hold time for LSA throttle 4000 msec
Maximum wait time for LSA throttle 6000 msec
Minimum LSA arrival 2000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 50 msec
Retransmission pacing timer 75 msec
Number of external LSA 15. Checksum S
<output omitted>
```

## 6.46 OSPF Resource Limiting

- Configure R4 so that no more than 5000 LSAs can exist in the database.
- No more than 500 of these routes should be originated through redistribution.
- R4's OSPF process should be allowed to use up to 20% CPU utilization before allowing interrupts from higher priority processes.

### Configuration

---

```
R4:
router ospf 1
  max-lsa 5000
  redistribute maximum-prefix 500
  process-min-time percent 20
```

### Verification

---

#### Note

LSA and redistributed prefix limiting in OSPF is used to prevent against attacks or misconfigurations in the OSPF domain which could interrupt normal forwarding, such as if the full Internet BGP table is accidentally redistributed into IGP. Likewise the `process-min-time` command rate limits OSPF's CPU usage so that other processes, such as IP Input, can be serviced.

#### Rack1R4#show ip ospf

```
Routing Process "ospf 1" with ID 150.1.4.4
Start time: 00:00:31.423, Time elapsed: 01:39:24.674
Supports only single TOS(TOS0) routes
Supports opaque LSA
Supports Link-local Signaling (LLS)
Supports area transit capability
Maximum number of non self-generated LSA allowed 5000
  Threshold for warning message 75%
  Ignore-time 5 minutes, reset-time 10 minutes
  Ignore-count allowed 5, current ignore-count 0
It is an area border and autonomous system boundary router
Redistributing External Routes from,
  rip, includes subnets in redistribution
  Maximum limit of redistributed prefixes 500
  Threshold for warning message 75%
<output omitted>
```



## 6.47 Miscellaneous OSPF Features

- Configure SW4 so that it goes not generate a log message upon receipt of a Type-6 LSA advertisement.
- Configure SW4 so that it does not take the MTU value into account when establishing adjacencies on its Port-Channel1 interface.
- Configure SW4 to reflect the following output:

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
Rack1SW2	1	FULL/BDR	00:00:38	155.1.108.8	Port-channel1

### Configuration

```
SW4:
ip host Rack1SW2 150.1.8.8
!
ip ospf name-lookup
!
interface Port-channel1
 ip ospf mtu-ignore
!
router ospf 1
 ignore lsa mospf
```

### Verification

#### Note

Cisco's implementation of OSPFv2 does not support Multicast OSPF, which is advertised through LSA Type-6. Upon receipt of this LSA type from a non-Cisco OSPF router, a log message is generated. To disable this issue the **ignore lsa mospf** command under the OSPF process.

A neighbor relationship cannot occur if two OSPF neighbors have different MTU values on their interfaces. In the case that the MTU difference is by design, the interface level command **ip ospf mtu-ignore** removes this requirement from the adjacency establishment.

The **ip ospf name-lookup** command performs DNS resolution on the OSPF router-id value in show commands to simplify the identification of neighbors.

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
Rack1SW2	1	FULL/BDR	00:00:38	155.1.108.8	Port-channel1

