Chapter 7: Advanced OSPF

Instructor Materials

CCNP Enterprise: Advanced Routing

Chapter 7 Content

This chapter covers the following content:

- **Link-State Advertisements -** This section explains how Open Shortest Path First (OSPF) stores, communicates, and builds the topology from link-state advertisements (LSAs).
- **OSPF Stubby Areas -** This section explains the method OSPF provides for filtering external routes while still providing connectivity to them.
- **OSPF Path Selection -** This section explains how OSPF makes path selection choices for routes learned within the OSPF routing domain.
- **Summarization of Routes -** This section explains how network summarization works with OSPF.
- **Discontiguous Network -** This section demonstrates a discontiguous network and explains why it cannot distribute routes to all areas properly.
- **Virtual Links -** This section explains how OSPF repairs a discontiguous network.

Link-State Advertisements

- An OSPF link-state advertisement (LSA) contains the link state and link metric to a neighboring router.
- Received LSAs are stored in a local database called the link-state database (LSDB); the LSDB advertises the link-state information to neighboring routers exactly as the original advertising router advertised it.
- All OSPF routers in the same area maintain a synchronized identical copy of the LSDB for that area.
- The LSDB provides the topology of the network, providing the router a complete map of the network.
- When OSPF neighbors become adjacent, the LSDBs synchronize between the OSPF routers. As an OSPF router adds or removes a directly connected network link to or from its database, the router floods the LSA out all active OSPF interfaces.

LSA Types

The OSPF LSA contains a complete list of networks advertised from that router. OSPF uses six LSA types for IPv4 routing:

- **Type 1, router -** LSAs that advertise network prefixes within an area
- **Type 2, network -** LSAs that indicate the routers attached to broadcast segment within an area
- **Type 3, summary -** LSAs that advertise network prefixes that originate from a different area
- **Type 4, ASBR summary -** LSA used to locate the ASBR from a different area
- **Type 5, AS external -** LSA that advertises network prefixes that were redistributed into OSPF
- **Type 7, NSSA external -** LSA for external network prefixes that were redistributed in a local NSSA area

LSA Types 1, 2, and 3 are used for building the SPF tree for intra-area and interarea route routes. LSA Types 4, 5, and 7 are related to external OSPF routes (that is, routes that were redistributed into the OSPF routing domain).

LSA Types (Cont.)

Figure 7-1 shows a packet capture of an OSPF update LSA and outlines the important components of the LSA: the LSA type, LSA age, the sequence, and the advertising router. Because this is a Type 1 LSA, the link IDs add relevance as they list the attached networks and the associated OSPF cost for each interface.

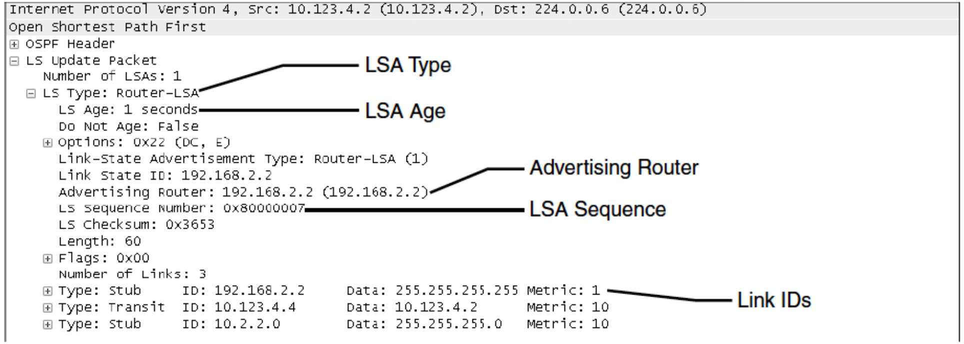

Figure 7-1 *Packet Capture of an LSA Update for the Second Interface*

Figure 7-2 shows a sample topology to demonstrate the different LSA types. In this topology:

- R1, R2, and R3 are member (internal) routers.
- R4 and R5 are area border routers (ABRs).
- R6 is the ASBR, which is redistributing the 172.16.6.0/24 network into OSPF.

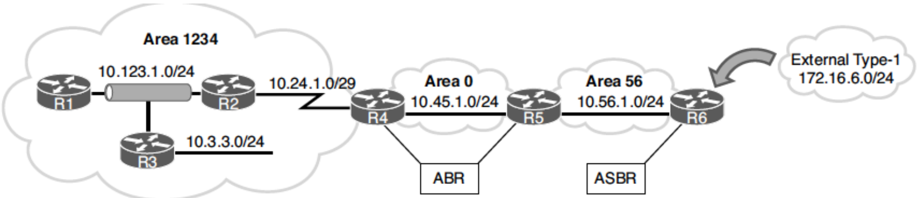

Figure 7-2 *LSA Reference Topology for LSAs*

LSA Sequences

- OSPF uses the sequence number to overcome problems caused by delays in LSA propagation in a network.
- The LSA sequence number is a 32-bit number used to control versioning. When the originating router sends out LSAs, the LSA sequence number is incremented.
- If a router receives an LSA sequence that is greater than the one in the LSDB, it processes the LSA.
- If the LSA sequence number is lower than the one in the LSDB, the router deems the LSA old and discards it.

LSA Age and Flooding

- Every OSPF LSA includes an age that is entered into the local LSDB that increments by 1 every second.
- When a router's OSPF LSA age exceeds 1800 seconds (that is, 30 minutes) for its networks, the originating router advertises a new LSA with the LSA age set to 0.
- As each router forwards the LSA, the LSA age is incremented with a calculated delay that reflects the link (which is minimal).
- If the LSA age reaches 3600, the LSA is deemed invalid and is purged from the LSDB.
- The repetitive flooding of LSAs is a secondary safety mechanism to ensure that all routers maintain a consistent LSDB within an area.

LSA Type 1

Every OSPF router advertises a Type 1 LSA. Type 1 LSAs are the essential building blocks in the LSDB. A Type 1 LSA entry exists for each OSPF-enabled link (that is, an interface and its attached networks). Figure 7-3 shows that the Type 1 LSAs are not advertised outside Area 1234, thus making the underlying topology in an area invisible to other areas.

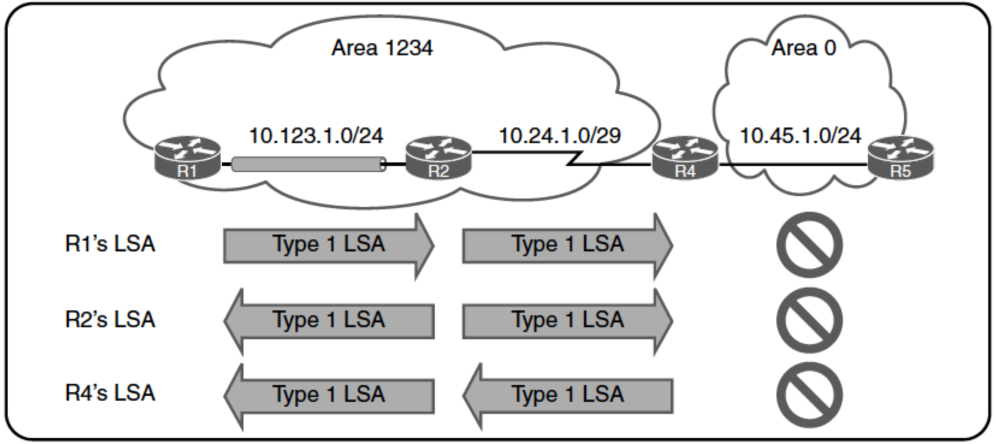

Figure 7-3 *Type 1 LSA Flooding in an Area*

LSA Type 1 (Cont.)

For a brief summary view of the Type 1 LSAs for an area, look under the Router Link States column within the LSDB, as shown in Example 7-1.

Table 7-2 provides an overview of the fields used within the LSDB output.

Example 7-1 *Generic OSPF LSA Output for Type 1 LSAs*

```
R1# show ip ospf database
            OSPF Router with ID (192.168.1.1) (Process ID 1)

                Router Link States (Area 1234)

Link ID         ADV Router      Age         Seq#       Checksum Link count
192.168.1.1     192.168.1.1     14          0x80000006 0x009EA7 1
192.168.2.2     192.168.2.2     2020        0x80000006 0x00AD43 3
192.168.3.3     192.168.3.3     6           0x80000006 0x0056C4 2
192.168.4.4     192.168.4.4     61          0x80000005 0x007F8C 2
```

Table 7-2 OSPF LSDB Fields

Field	Description
Link ID	Identifies the object that the link connects to. It can refer to the neighboring router's RID, the IP address of the DR's interface, or the IP network address.
ADV Router	The OSPF router ID for this LSA.
AGE	The age of the LSA on the router on which the command is being run. Values over 1800 are expected to refresh soon.
Seq #	The sequence number for the LSA to protect out-of-order LSAs.
Checksum	The checksum of the LSA to verify integrity during flooding.
Link Count	The number of links on this router in the Type 1 LSA.

Examining Type 1 LSAs

You can examine the Type 1 OSPF LSAs by using the command **show ip ospf database router**, as demonstrated in Example 7-2.

Example 7-2 *OSPF Type 1 LSAs for Area 1234*

```
R1# show ip ospf database router
! Output omitted for brevity
            OSPF Router with ID (192.168.1.1) (Process ID 1)

                Router Link States (Area 1234)

  LS age: 352
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 192.168.1.1
  Advertising Router: 192.168.1.1
  LS Seq Number: 80000014
```

```
Length: 36
Number of Links: 1

 Link connected to: a Transit Network
   (Link ID) Designated Router address: 10.123.1.3
   (Link Data) Router Interface address: 10.123.1.1
     TOS 0 Metrics: 1


LS age: 381
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 192.168.2.2
Advertising Router: 192.168.2.2
LS Seq Number: 80000015
Length: 60
Number of Links: 3

  Link connected to: another Router (point-to-point)
   (Link ID) Neighboring Router ID: 192.168.4.4
   (Link Data) Router Interface address: 10.24.1.1
     TOS 0 Metrics: 64

  Link connected to: a Stub Network
   (Link ID) Network/subnet number: 10.24.1.0
   (Link Data) Network Mask: 255.255.255.248
     TOS 0 Metrics: 64

  Link connected to: a Transit Network
   (Link ID) Designated Router address: 10.123.1.3
   (Link Data) Router Interface address: 10.123.1.2
     TOS 0 Metrics: 1
```

LSA Type 1 Neighbor States

Each OSPF-enabled interface is listed under the number of links for each router. Each network link on a router contains the following information, in this order:

- Link type, as shown in Table 7-3, after the Link connected to
- Link ID, using the values based on the link type listed in Table 7-3
- Link data (when applicable)
- Metric for the interface

Table 7-3 OSPF Neighbor States for Type 1 LSAs

Description	Link Type	Link ID Value	Link Data
Point-to-point link (IP address assigned)	1	Neighbor RID	Interface IP address
Point-to-point link (using IP unnumbered)	1	Neighbor RID	MIB II IfIndex value
Link to transit network	2	Interface address of DR	Interface IP address
Link to stub network	3	Network address	Subnet mask
Virtual link	4	Neighbor RID	Interface IP address

cisco

Link Types

During the SPF tree calculation, network link types are one of the following:

- Transit - A transit network indicates that an adjacency was formed and that a DR was elected on that link.
- Point-to-point - Point-to-point links indicate that an adjacency was formed on a network type that does not use a DR. Interfaces using the OSPF point-to-point network type advertise two links. One link is the point-to-point link type that identifies the OSPF neighbor RID for that segment, and the other link is a stub network link that provides the subnet mask for that network.
- Stub - A stub network indicates that no neighbor adjacencies were established on that link. Point-to-point and transit link types that did not become adjacent with another OSPF router are classified as a stub network link type. When an OSPF adjacency forms, the link type changes to the appropriate type: point-to-point or transit.

Visualization of Type 1 LSAs

Figure 7-5 demonstrates the topology built by all routers in Area 1234, using the LSA attributes for Area 1234 from all four routers. Using only Type 1 LSAs, a connection is made between R2 and R4 because they point to each other's RID in the point-to-point LSA.

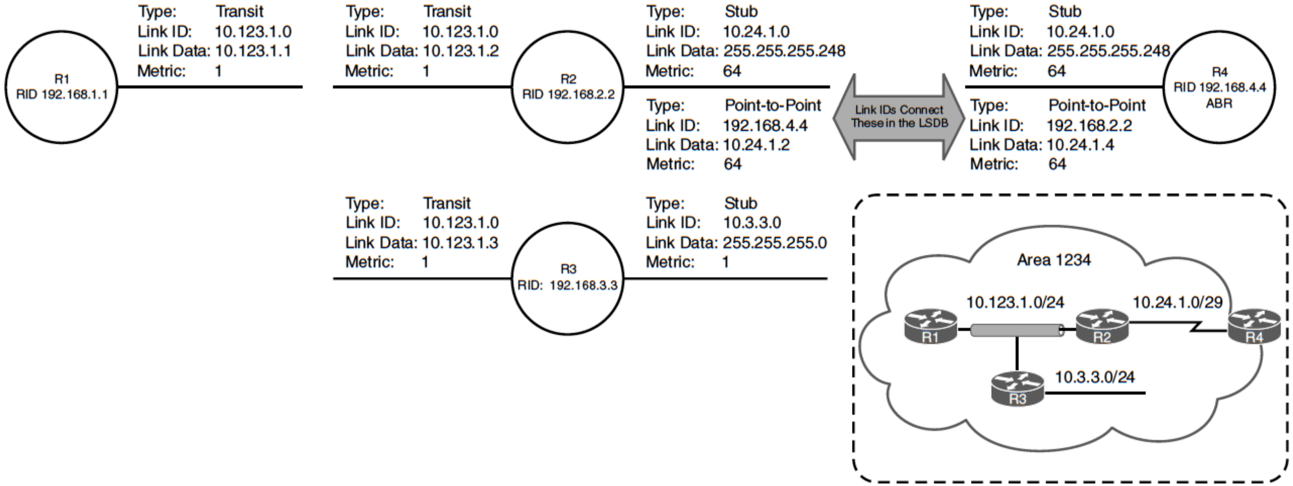

Figure 7-5 *Visualization of Type 1 LSAs*

LSA Type 2

A Type 2 LSA represents a multi-access network segment that uses a DR. The DR always advertises the Type 2 LSA and identifies all the routers attached to that network segment.

If a DR has not been elected, a Type 2 LSA is not present in the LSDB because the corresponding Type 1 transit link type LSA is a stub. Type 2 LSAs are not flooded outside the originating OSPF area in an identical fashion to Type 1 LSAs.

Example 7-3 provides the output for Type 2 LSAs in Area 1234 from the reference topology.

Example 7-3 *Generic OSPF LSA Output for Type 2 LSAs*

```
R1# show ip ospf database
! Output omitted for brevity
            OSPF Router with ID (192.168.1.1) (Process ID 1)
..
                Net Link States (Area 1234)

Link ID         ADV Router         Age          Seq#        Checksum
10.123.1.3      10.192.168.3.3     1752         0x80000012 0x00ADC5
```

LSA Type 2 (Cont.)

To see detailed Type 2 LSA information, you use the command **show ip ospf database network**. Example 7-4 shows the Type 2 LSA that is advertised by R3 and shows that the link-state ID 10.123.1.3 attaches to R1, R2, and R3 (by listing their RIDs at the bottom). The network mask for the subnet is included in the Type 2 LSA.

Example 7-4 *Detailed Output for OSPF Type 2 LSAs*

```
R1# show ip ospf database network
            OSPF Router with ID (192.168.1.1) (Process ID 1)

                Net Link States (Area 1234)

  LS age: 356
  Options: (No TOS-capability, DC)
  LS Type: Network Links
  Link State ID: 10.123.1.3 (address of Designated Router)
  Advertising Router: 192.168.3.3
  LS Seq Number: 80000014
  Checksum: 0x4DD
  Length: 36
  Network Mask: /24
        Attached Router: 192.168.3.3
        Attached Router: 192.168.1.1
        Attached Router: 192.168.2.2
```

Visualization of Type 1 and Type 2 LSAs

Figure 7-6 provides a visualization of the Type 1 and Type 2 LSAs; it corresponds with Area 1234 perfectly.

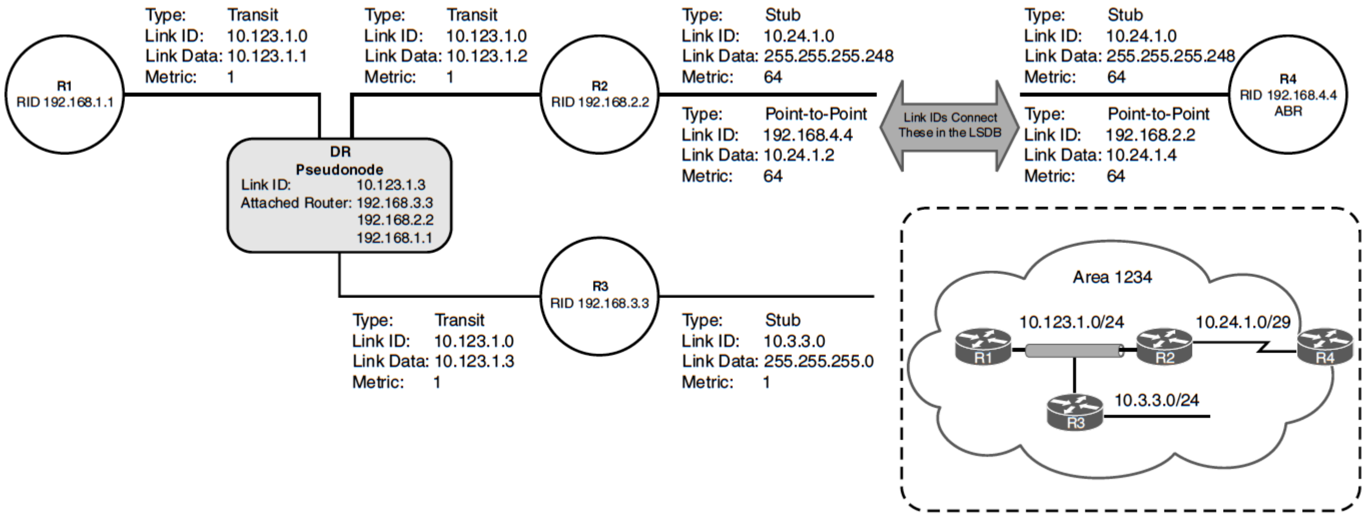

Figure 7-6 *Visualization of Area 1234 with Type 1 and Type 2 LSAs*

LSA Type 3

Type 3 LSAs represent networks from other areas. ABRs do not forward Type 1 or Type 2 LSAs into other areas. When an ABR receives a Type 1 LSA, it creates a Type 3 LSA referencing the network in the original Type 1 LSA.

The ABR then advertises the Type 3 LSA into other areas. If an ABR receives a Type 3 LSA from Area 0 (backbone), it regenerates a new Type 3 LSA for the nonbackbone area and lists itself as the advertising router with the additional cost metric.

Figure 7-7 demonstrates the concept of a Type 3 LSA interaction with Type 1 LSAs. Notice that the Type 1 LSAs exist only in the area of origination and convert to Type 3 when they cross the ABRs (R4 and R5).

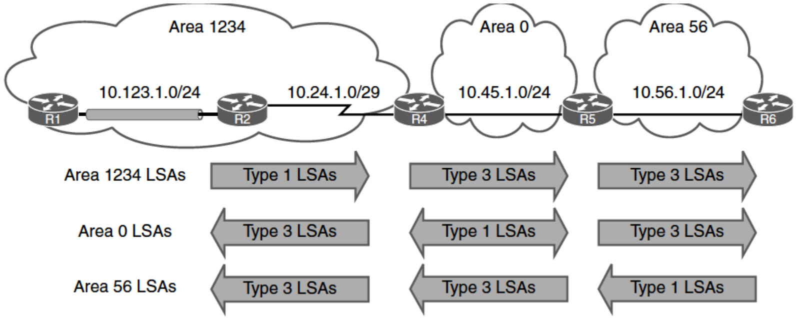

Figure 7-7 *Type 3 LSA Conceptual Diagram*

LSA Type 3 (Cont.)

For a summary view of the Type 3 LSAs, look under Summary Net Link States, as shown in Example 7-5.

The Type 3 LSAs show up under the appropriate area where they exist in the OSPF domain.

For example, the 10.56.1.0 Type 3 LSA exists only in Area 0 and Area 1234 on R4. R5 contains the 10.56.1.0 Type 3 LSA only for Area 0, but not for Area 56 because Area 56 has a Type 1 LSA.

Example 7-5 *Generic OSPF LSA Output for Type 3 LSAs*

```
R4# show ip ospf database
! Output omitted for brevity
            OSPF Router with ID (192.168.4.4) (Process ID 1)
..
                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.3.3.0        192.168.4.4     813         0x80000013 0x00F373
10.24.1.0       192.168.4.4     813         0x80000013 0x00CE8E
10.56.1.0       192.168.5.5     591         0x80000013 0x00F181
10.123.1.0      192.168.4.4     813         0x80000013 0x005A97
..
                Summary Net Link States (Area 1234)

Link ID         ADV Router      Age         Seq#       Checksum
10.45.1.0       192.168.4.4     813         0x80000013 0x0083FC
10.56.1.0       192.168.4.4     813         0x80000013 0x00096B
```

```
R5# show ip ospf database
! Output omitted for brevity
            OSPF Router with ID (192.168.5.5) (Process ID 1)
..
                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.3.3.0        192.168.4.4     893         0x80000013 0x00F373
10.24.1.0       192.168.4.4     893         0x80000013 0x00CE8E
10.56.1.0       192.168.5.5     668         0x80000013 0x00F181
10.123.1.0      192.168.4.4     893         0x80000013 0x005A97
..
                Summary Net Link States (Area 56)

Link ID         ADV Router      Age         Seq#       Checksum
10.3.3.0        192.168.5.5     668         0x80000013 0x00F073
10.24.1.0       192.168.5.5     668         0x80000013 0x00CB8E
10.45.1.0       192.168.5.5     668         0x80000013 0x007608
10.123.1.0      192.168.5.5     668         0x80000013 0x005797
```

LSA Type 3 (Cont.)

To see detailed Type 3 LSA information, you use the **command show ip ospf database summary**. You can restrict the output to a specific LSA by adding the prefix to the end of the command. The advertising router for Type 3 LSAs is the last ABR that advertises the prefix.

The metric in the Type 3 LSA uses the following logic:

- If the Type 3 LSA is created from a Type 1 LSA, it is the total path metric to reach the originating router in the Type 1 LSA.

- If the Type 3 LSA is created from a Type 3 LSA from Area 0, it is the total path metric to the ABR plus the metric in the original Type 3 LSA.

LSA Type 3 (Cont.)

Example 7-6 shows the Type 3 LSA for the Area 56 prefix (10.56.1.0/24) from R4's LSDB. Notice that the metric increases in Area 1234's LSA compared to Area 0's LSA.

Table 7-4 provides an explanation of the fields in a Type 3 LSA.

Table 7-4 Type 3 LSA Fields

Field	Description
Link ID	Network number
Advertising Router	RID of the router advertising the route (ABR)
Network Mask	Prefix length for the advertised network
Metric	Metric for the LSA

Example 7-6 *Detailed Output for OSPF Type 3 LSAs*

```
R4# show ip ospf database summary 10.56.1.0
            OSPF Router with ID (192.168.4.4) (Process ID 1)

                Summary Net Link States (Area 0)

  LS age: 754
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 10.56.1.0 (summary Network Number)
  Advertising Router: 192.168.5.5
  LS Seq Number: 80000013
  Checksum: 0xF181
  Length: 28
  Network Mask: /24
        MTID: 0         Metric: 1


                Summary Net Link States (Area 1234)

  LS age: 977
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 10.56.1.0 (summary Network Number)
  Advertising Router: 192.168.4.4
  LS Seq Number: 80000013
  Checksum: 0x96B
  Length: 28
  Network Mask: /24
        MTID: 0         Metric: 2
```

Visualization of Type 3 LSAs

Understanding the metric in Type 3 LSAs is an important concept. Figure 7-8 provides R4's perspective of the Type 3 LSA created by ABR (R5) for the 10.56.1.0/24 network.

Figure 7-9 provides R3's perspective of the Type 3 LSA created by the ABR (R4) for the 10.56.1.0/24.

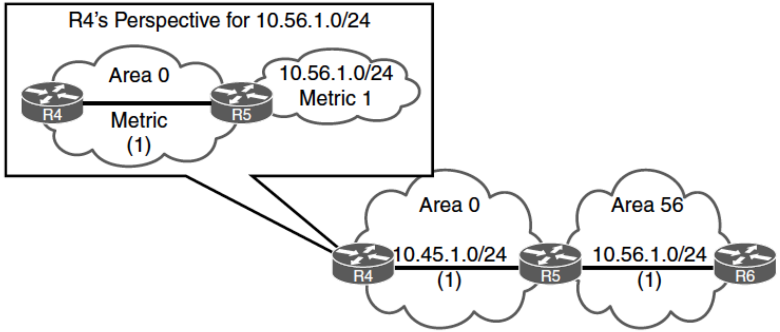

Figure 7-8 *Visualization of the 10.56.1.0/24 Type 3 LSA from Area 0*

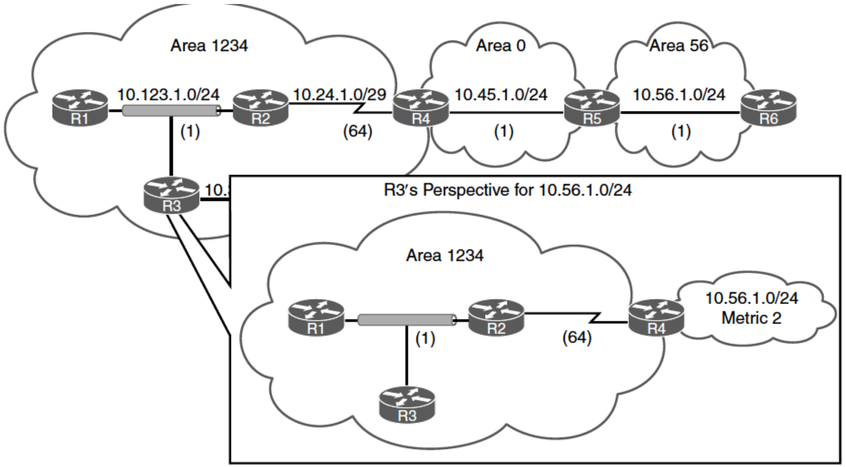

Figure 7-9 *Visualization of the 10.56.1.0/24 Type 3 LSA from Area 1234*

LSA Type 5

When a route is redistributed into OSPF, the router is known as an autonomous system boundary router (ASBR). The external route is flooded throughout the entire OSPF domain as a Type 5 LSA. Type 5 LSAs are not associated with a specific area and are flooded throughout the OSPF domain.

Figure 7-10 shows R6 redistributing the 172.16.6.0/24 static route into the OSPF domain.

Example 7-7 provides a brief summary view of the Type 5 LSAs under Type 5 AS External Link States.

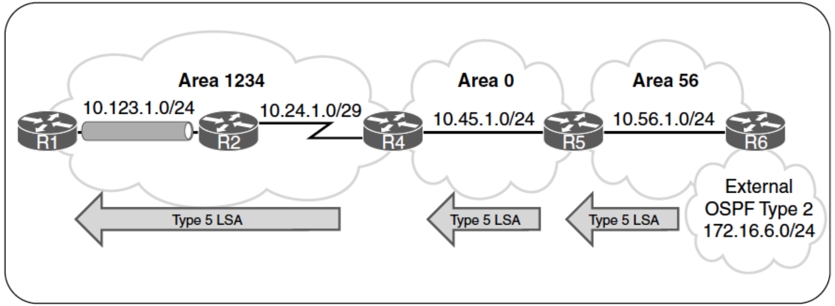

Figure 7-10 *OSPF Type 5 LSA Flooding*

Example 7-7 *Generic OSPF LSA Output for Type 5 LSAs*

```
R6# show ip ospf database
! Output omitted for brevity
                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
172.16.6.0      192.168.6.6     11          0x80000001 0x000866 0
```

LSA Type 5 (Cont.)

You can see Type 5 LSAs in detail by using the command **show ip ospf database external**. ABRs only modify the LSA age as the Type 5 LSA propagates through the OSPF domain.

Example 7-8 provides detailed output for the external OSPF LSAs in the OSPF domain.

Table 7-5 provides an explanation of the fields in a Type 5 LSA.

Table 7-5 Type 5 LSA Fields

Field	Description
Link ID	External network number
Network Mask	Subnet mask for the external network
Advertising Router	RID of the router advertising the route (ASBR)
Metric Type	OSPF external metric type (Type 1 O E1 or Type 2 O E2)
Metric	Metric upon redistribution
External Route Tag	32-bit field included with an external route, which is not used by OSPF under normal operations but can be used to communicate AS boundaries or other relevant information to prevent routing loops

Example 7-8 *Detailed Output for OSPF Type 5 LSAs*

```
R6# show ip ospf database external
            OSPF Router with ID (192.168.6.6) (Process ID 1)

                Type-5 AS External Link States

  LS age: 720
  Options: (No TOS-capability, DC, Upward)
  LS Type: AS External Link
  Link State ID: 172.16.6.0 (External Network Number )
  Advertising Router: 192.168.6.6
  LS Seq Number: 8000000F
  Checksum: 0xA9B0
  Length: 36
  Network Mask: /24
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 0
```


LSA Type 4

A Type 4 LSA locates the ASBR for a Type 5 LSA. Type 4 LSAs provide a way for routers to locate the ASBR when the router is in a different area from the ASBR. A Type 4 LSA is created by the first ABR, and it provides a summary route strictly for the ASBR of a Type 5 LSA.

The metric for a Type 4 LSA uses the following logic:

- When the Type 5 LSA crosses the first ABR, the ABR creates a Type 4 LSA with a metric set to the total path metric to the ASBR.
- When an ABR receives a Type 4 LSA from Area 0, the ABR creates a new Type 4 LSA with a metric set to the total path metric of the first ABR plus the metric in the original Type 4 LSA.

Figure 7-11 shows how the ABRs (R4 and R5) create Type 4 LSAs for the ASBR (R6).

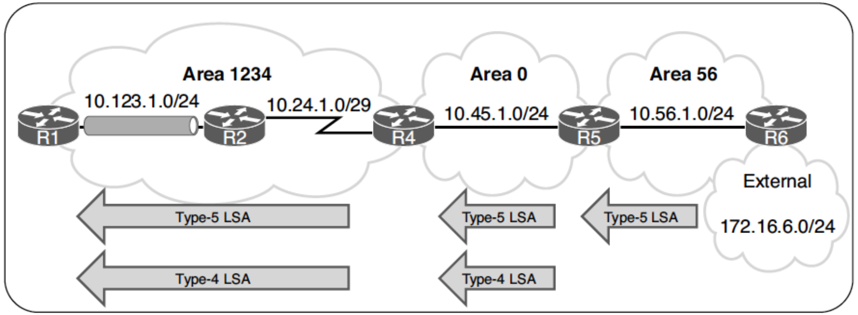

Figure 7-11 *OSPF Type 4 and Type 5 LSA Flooding Within an OSPF Domain*

LSA Type 4 (Cont.)

Example 7-9 provides a brief summary view of the Type 4 LSAs for R4 in the LSDB under Summary ASB Link States.

To view the details of the Type 4 LSAs, you use the command **show ip ospf database asbr-summary**. Example 7-10 provides detailed output of the Type 4 LSA on R4.

Example 7-9 *Generic OSPF LSA Output for Type 4 LSAs*

```
R4# show ip ospf database
! Output omitted for brevity
            OSPF Router with ID (192.168.4.4) (Process ID 1)
..
                Summary ASB Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
192.168.6.6     192.168.5.5     930         0x8000000F 0x00EB58
..
                Summary ASB Link States (Area 1234)

Link ID         ADV Router      Age         Seq#       Checksum
192.168.6.6     192.168.4.4     1153        0x8000000F 0x000342
```

Example 7-10 *Detailed Output for Type 4 LSAs*

```
R4# show ip ospf database asbr-summary
! Output omitted for brevity
            OSPF Router with ID (192.168.4.4) (Process ID 1)

                Summary ASB Link States (Area 0)
  LS age: 1039
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(AS Boundary Router)
  Link State ID: 192.168.6.6 (AS Boundary Router address)
  Advertising Router: 192.168.5.5
  Length: 28
  Network Mask: /0
        MTID: 0         Metric: 1


                Summary ASB Link States (Area 1234)

  LS age: 1262
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(AS Boundary Router)
  Link State ID: 192.168.6.6 (AS Boundary Router address)
  Advertising Router: 192.168.4.4
  Length: 28
  Network Mask: /0
        MTID: 0         Metric: 2
```

LSA Type 7

A Type 7 LSA exists only in NSSAs where route redistribution is occurring. An ASBR injects external routes as Type 7 LSAs in an NSSA. The ABR does not advertise Type 7 LSAs outside the originating NSSA, but it converts the Type 7 LSA into a Type 5 LSA for the other OSPF areas. If the Type 5 LSA crosses Area 0, the second ABR creates a Type 4 LSA for the Type 5 LSA.

Figure 7-12 shows Area 55 as an NSSA and R6 redistributing the 172.16.6.0/24 prefix.

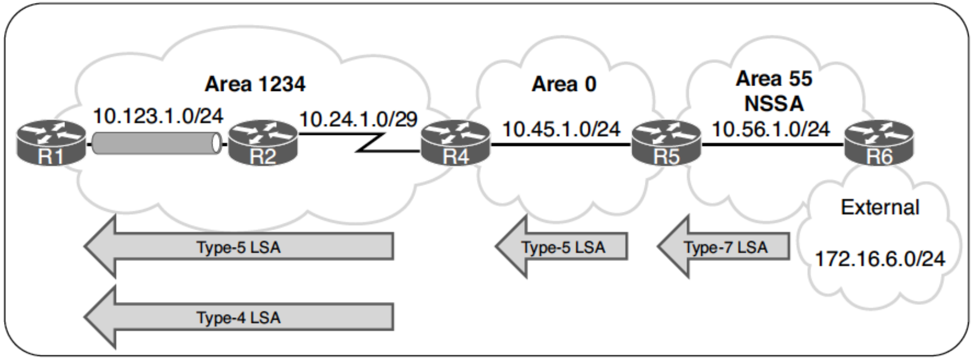

Figure 7-12 *OSPF Type 7 LSAs*

LSA Type 7 (Cont.)

Example 7-11 provides a brief summary view of the Type 7 LSAs, under Type 7 AS External Link States.

Type 7 LSAs are present only in the OSPF NSSA where redistribution is occurring. Notice that the Type 7 LSA is not present on R4; R4 contains a Type 5 LSA that was created by R5 and the Type 4 LSA (created by R4 for Area 1234).

Example 7-11 *Generic OSPF LSA Output for Type 7 LSAs*

```
R5# show ip ospf database
! Output omitted for brevity
            OSPF Router with ID (192.168.5.5) (Process ID 1)

..
Type-7 AS External Link States (Area 56)

Link ID         ADV Router      Age         Seq#       Checksum Tag
172.16.6.0      192.168.6.6     46          0x80000001 0x00A371 0

!   Notice that no Type-4 LSA has been generated. Only the Type-7 LSA for Area 56
!   and the Type-5 LSA for the other areas. R5 advertises the Type-5 LSA
                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
172.16.6.0      192.168.5.5     38          0x80000001 0x0045DB
```

```
R4# show ip ospf database
! Output omitted for brevity
            OSPF Router with ID (192.168.4.4) (Process ID 1)
..
                Summary ASB Link States (Area 1234)

Link ID         ADV Router      Age         Seq#       Checksum
192.168.5.5     192.168.4.4     193         0x80000001 0x002A2C

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
172.16.6.0      192.168.5.5     176         0x80000001 0x0045DB 0
```

LSA Type 7 (Cont.)

To see the specific Type 7 LSA details, you use the command **show ip ospf database nssa-external**.

Example 7-12 shows this command executed on R5.

Table 7-6 provides an explanation of the fields in a Type 7 LSA.

Example 7-12 *Detailed Output for OSPF Type 7 LSAs*

```
R5# show ip ospf database nssa-external
            OSPF Router with ID (192.168.5.5) (Process ID 1)

                Type-7 AS External Link States (Area 56)
  LS age: 122
  Options: (No TOS-capability, Type 7/5 translation, DC, Upward)
  LS Type: AS External Link
  Link State ID: 172.16.6.0 (External Network Number )
  Advertising Router: 192.168.6.6
  LS Seq Number: 80000001
  Checksum: 0xA371
  Length: 36
  Network Mask: /24
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 20
        Forward Address: 10.56.1.6
        External Route Tag: 0
```

Table 7-6 Type 7 LSA Fields

Field	Description
Link ID	External network number
Network Mask	Subnet mask for the external network
Advertising Router	RID of the router advertising the route (ASBR)
Metric Type	OSPF external metric type (Type 1 O N1 or Type 2 O N2)
Metric	Metric upon redistribution
External Route Tag	32-bit field that is included with an external route, which is not used by OSPF itself and can be used to communicate AS boundaries or other relevant information to prevent routing loops

LSA Types Summary

Table 7-7 provides a summary of the OSPF LSAs discussed.

Figure 7-13 shows the network prefixes from the sample topology and the relevant LSAs that are present.

Table 7-7 OSPF LSA Types

LSA Type	Description
1	Router link
2	Network link
3	Summary link
4	ASBR summary
5	AS external
7	NSSA external

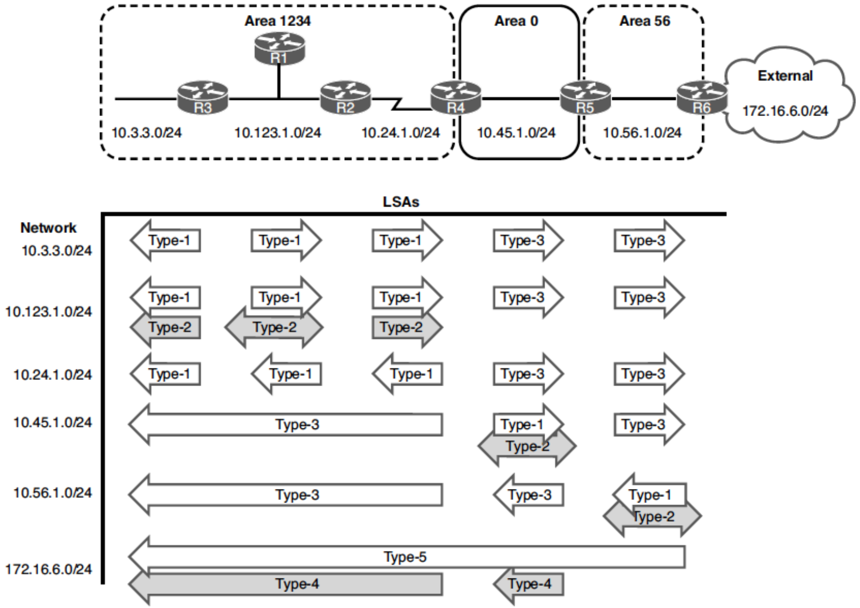

Figure 7-13 *Overview of LSA Types from the Sample Topology*

OSPF Stubby Areas

- OSPF stubby areas provide a method to filter out external routes and the option to block interarea routes.

- The following sections explain the four types of OSPF stubby areas in more detail:
    - Stub areas
    - Totally stubby areas
    - Not-so-stubby areas (NSSAs)
    - Totally NSSAs

Stub Areas

OSPF stub areas prohibit Type 5 LSAs (external routes) and Type 4 LSAs (ASBR summary LSAs) from entering the area at the ABR. RFC 2328 states that when a Type 5 LSA reaches the ABR of a stub area, the ABR generates a default route for the stub via a Type 3 LSA. A Cisco ABR generates a default route when the area is configured as a stub and has an OSPF-enabled interface configured for Area 0. Figure 7-14 demonstrates the concept.

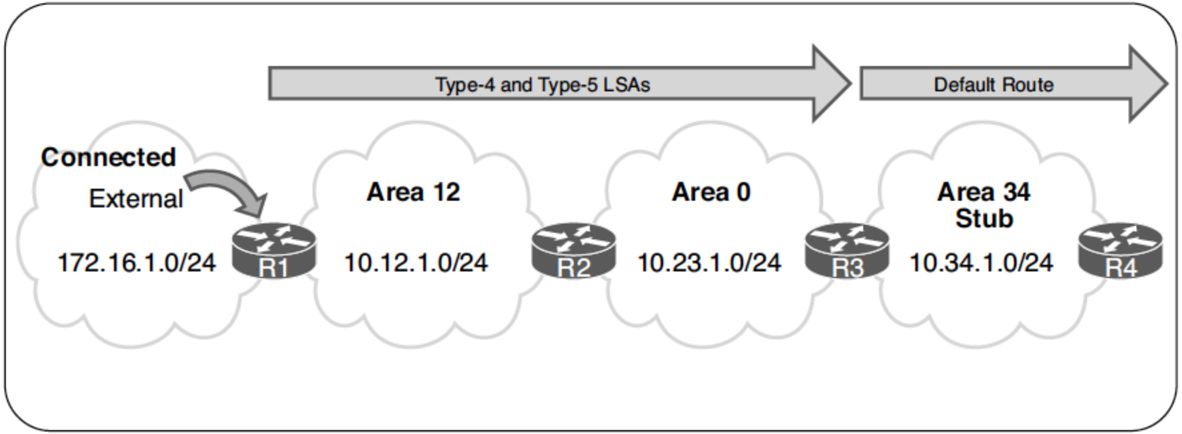

Figure 7-14 *OSPF Stub Area Concept*

OSPF Stub Configuration

Example 7-13 shows the routing tables for R3 and R4 before Area 34 is configured as a stub area.

All routers in the stub area must be configured as stubs, or an adjacency cannot form because the area type flags in the hello packets do not match. An area is configured as a stub with the OSPF process command **area** *area-id* **stub**.

Example 7-14 demonstrates the configuration for R3 and R4 making Area 34 an OSPF stub area.

Example 7-13 *Routing Table in Area 1 and Area 2 Without Stub*

```
R3# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     10.12.1.0/24 [110/2] via 10.23.1.2, 00:01:36, GigabitEthernet0/1
    172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/22] via 10.23.1.2, 00:01:36, GigabitEthernet0/1
O IA     192.168.1.1 [110/3] via 10.23.1.2, 00:01:36, GigabitEthernet0/1
O        192.168.2.2 [110/2] via 10.23.1.2, 00:01:46, GigabitEthernet0/1
O        192.168.4.4 [110/2] via 10.34.1.4, 00:01:46, GigabitEthernet0/0
```

```
R4# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O IA     10.12.1.0/24 [110/3] via 10.34.1.3, 00:00:51, GigabitEthernet0/0
O IA     10.23.1.0/24 [110/2] via 10.34.1.3, 00:00:58, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/23] via 10.34.1.3, 00:00:46, GigabitEthernet0/0
O IA     192.168.1.1 [110/4] via 10.34.1.3, 00:00:51, GigabitEthernet0/0
O IA     192.168.2.2 [110/3] via 10.34.1.3, 00:00:58, GigabitEthernet0/0
O IA     192.168.3.3 [110/2] via 10.34.1.3, 00:00:58, GigabitEthernet0/0
```

Example 7-14 *OSPF Stub Configuration for Area 34*

```
R3# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# router ospf 1
R3(config-router)# area 34 stub
```

```
R4# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)# router ospf 1
R4(config-router)# area 34 stub
```

OSPF Stubby Areas

Routing Table After Stub Area Configuration

Example 7-15 shows the routing table for R3 and R4 after Area 34 is made an OSPF stub area.

The routing table from R3's perspective is not modified as it receives the Type 4 and Type 5 LSAs from Area 0.

When the Type 5 LSA (172.16.1.0/24) reaches the ABR (R3), the ABR generates a default route by using a Type 3 LSA.

While R4 does not see the route to the 172.16.1.0/24 network in its routing table, it has connectivity to that network through the default route.

Example 7-15 *Routing Table After Stub Area Configuration*

```
R3# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     10.12.1.0/24 [110/2] via 10.23.1.2, 00:03:10, GigabitEthernet0/1
     172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/22] via 10.23.1.2, 00:03:10, GigabitEthernet0/1
O IA     192.168.1.1 [110/3] via 10.23.1.2, 00:03:10, GigabitEthernet0/1
O        192.168.2.2 [110/2] via 10.23.1.2, 00:03:10, GigabitEthernet0/1
O        192.168.4.4 [110/2] via 10.34.1.4, 00:01:57, GigabitEthernet0/0
```

```
R4# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is 10.34.1.3 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/2] via 10.34.1.3, 00:02:45, GigabitEthernet0/0
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O IA     10.12.1.0/24 [110/3] via 10.34.1.3, 00:02:45, GigabitEthernet0/0
O IA     10.23.1.0/24 [110/2] via 10.34.1.3, 00:02:45, GigabitEthernet0/0
O IA     192.168.1.1 [110/4] via 10.34.1.3, 00:02:45, GigabitEthernet0/0
O IA     192.168.2.2 [110/3] via 10.34.1.3, 00:02:45, GigabitEthernet0/0
O IA     192.168.3.3 [110/2] via 10.34.1.3, 00:02:45, GigabitEthernet0/0
```

Totally Stubby Areas

Totally stubby areas prohibit Type 3 LSAs (interarea), Type 4 LSAs (ASBR summary LSAs), and Type 5 LSAs (external routes) from entering the area at the ABR. When an ABR of a totally stubby area receives a Type 3 or Type 5 LSA, the ABR generates a default route for the totally stubby area.

In fact, ABRs for totally stubby areas advertise the default route into the totally stubby areas the instant an interface is assigned to Area 0. Assigning the interface acts as the trigger for the Type 3 LSA that leads to the generation of the default route. Only intra-area and default routes should exist within a totally stubby area. Figure 7-15 illustrates the totally stubby area concept.

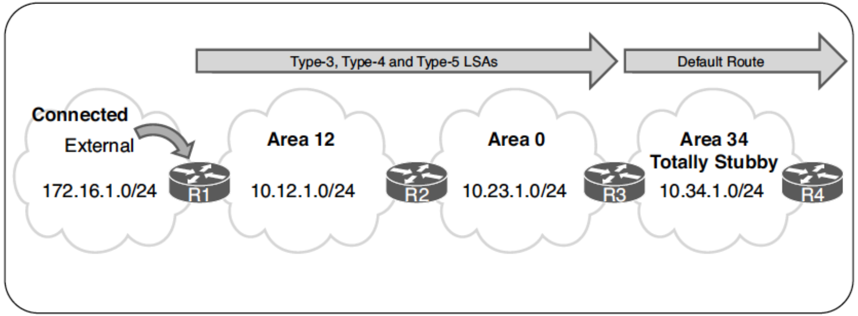

Figure 7-15 *Totally Stubby Area Concept*

Totally Stubby Area Configuration

ABRs of a totally stubby area have **no-summary** appended to the configuration. The command **area area-id stub no-summary** is configured under the OSPF process.

The keyword **no-summary** does exactly what it states: It blocks all Type 3 (summary) LSAs going into the stub area, making it a totally stubby area. Example 7-17 demonstrates the configuration of R3 and R4 for making Area 34 a totally stubby area for both routers.

Example 7-17 *Totally Stubby Area Configurations*

```
R3# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# router ospf 1
R3(config-router)# area 34 stub no-summary
```

```
R4# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)# router ospf 1
R4(config-router)# area 34 stub
```

Routing Table After Totally Stubby Area

Example 7-16 displays R3 and R4’s routing tables before Area 34 is converted to an OSPF totally stubby area.

Example 7-16 *Routing Tables of R3 and R4 Before Totally Stubby Areas*

```
R3# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     10.12.1.0/24 [110/2] via 10.23.1.2, 00:01:36, GigabitEthernet0/1
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/22] via 10.23.1.2, 00:01:36, GigabitEthernet0/1
O IA     192.168.1.1 [110/3] via 10.23.1.2, 00:01:36, GigabitEthernet0/1
O        192.168.2.2 [110/2] via 10.23.1.2, 00:01:46, GigabitEthernet0/1
O        192.168.4.4 [110/2] via 10.34.1.4, 00:01:46, GigabitEthernet0/0
```

```
R4# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O IA     10.12.1.0/24 [110/3] via 10.34.1.3, 00:00:51, GigabitEthernet0/0
O IA     10.23.1.0/24 [110/2] via 10.34.1.3, 00:00:58, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/23] via 10.34.1.3, 00:00:46, GigabitEthernet0/0
O IA     192.168.1.1 [110/4] via 10.34.1.3, 00:00:51, GigabitEthernet0/0
O IA     192.168.2.2 [110/3] via 10.34.1.3, 00:00:58, GigabitEthernet0/0
O IA     192.168.3.3 [110/2] via 10.34.1.3, 00:00:58, GigabitEthernet0/0
```

Example 7-18 shows the routing tables for R3 and R4 after Area 34 is converted to a totally stubby area.

Example 7-18 *Routing Tables After Area 34 Is Converted to a Totally Stubby Area*

```
R3# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     10.12.1.0/24 [110/2] via 10.23.1.2, 00:02:34, GigabitEthernet0/1
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/22] via 10.23.1.2, 00:02:34, GigabitEthernet0/1
O IA     192.168.1.1 [110/3] via 10.23.1.2, 00:02:34, GigabitEthernet0/1
O        192.168.2.2 [110/2] via 10.23.1.2, 00:02:34, GigabitEthernet0/1
O        192.168.4.4 [110/2] via 10.34.1.4, 00:03:23, GigabitEthernet0/0
```

```
R4# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is 10.34.1.3 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/2] via 10.34.1.3, 00:02:24, GigabitEthernet0/0
```

Not-So-Stubby Areas

The not-so-stubby-area (NSSA) prohibits Type 5 LSAs from entering at the ABR but allows for redistribution of external routes into the NSSA. As the ASBR redistributes the network into OSPF in the NSSA, the ASBR advertises the network with a Type 7 LSA instead of a Type 5 LSA. When the Type 7 LSA reaches the ABR, the ABR converts the Type 7 LSA to a Type 5 LSA. The ABR does not automatically advertise a default route when a Type 5 or Type 7 LSA is blocked. Figure 7-16 demonstrates the concept of LSAs being processed on the ABR for the NSSA. Notice that the default route is optional and depends on the configuration.

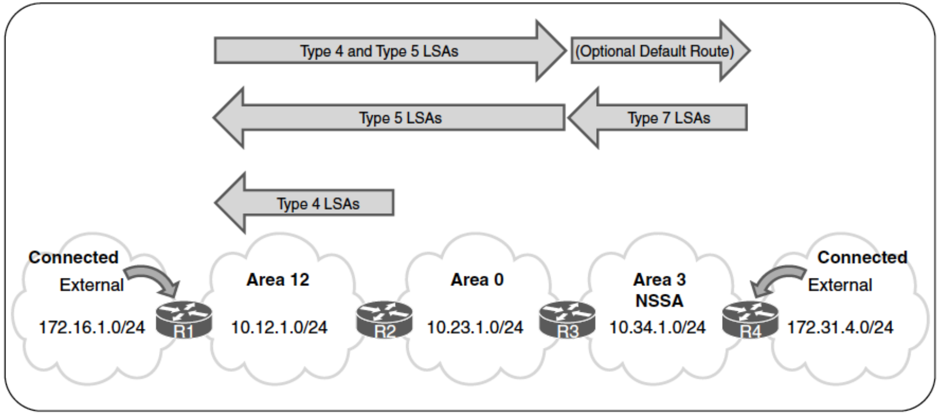

Figure 7-16 *NSSA Concept*

NSSA Configuration

The command **area** *area-id* **nssa** [**default-information-originate**] is placed under the OSPF process on the ABR. All routers in an NSSA must be configured with the **nssa** option, or they do not become adjacent because the area type flags must match in the OSPF hello protocol in order to become adjacent.

A default route is not injected on the ABRs automatically for NSSAs, but the optional command **default-information-originate** can be appended to the configuration if a default route is needed in the NSSA. Example 7-20 shows the OSPF configuration of R3 and R4 after making Area 34 an NSSA.

Example 7-20 *NSSA Configuration for Area 34 Routers*

```
R3# show run | section router ospf
router ospf 1
 router-id 192.168.3.3
 area 34 nssa default-information-originate
 network 10.23.1.0 0.0.0.255 area 0
 network 10.34.1.0 0.0.0.255 area 34
 network 192.168.3.3 0.0.0.0 area 0
```

```
R4# show run | section router ospf
router ospf 1
 router-id 192.168.4.4
 area 34 nssa
 redistribute connected metric-type 1 subnets
 network 10.34.1.0 0.0.0.255 area 34
 network 192.168.4.4 0.0.0.0 area 34
```

Routing Tables Before and After NSSA

Example 7-19 shows the routing tables of R1, R3, and R4 before Area 34 is converted to an NSSA.

Example 7-19 *R1, R3, and R4's Routing Tables Before Area 34 Is Converted to an NSSA*

```
R1# show ip route ospf | section 172.31
      172.31.0.0/24 is subnetted, 1 subnets
O E1     172.31.4.0 [110/23] via 10.12.1.2, 00:00:38, GigabitEthernet0/0
```

```
R3# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     10.12.1.0/24 [110/2] via 10.23.1.2, 00:01:34, GigabitEthernet0/1
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/22] via 10.23.1.2, 00:01:34, GigabitEthernet0/1
      172.31.0.0/24 is subnetted, 1 subnets
O E1     172.31.4.0 [110/21] via 10.34.1.4, 00:01:12, GigabitEthernet0/0
O IA     192.168.1.1 [110/3] via 10.23.1.2, 00:01:34, GigabitEthernet0/1
O        192.168.2.2 [110/2] via 10.23.1.2, 00:01:34, GigabitEthernet0/1
O        192.168.4.4 [110/2] via 10.34.1.4, 00:01:12, GigabitEthernet0/0
```

```
R4# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O IA     10.12.1.0/24 [110/3] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
O IA     10.23.1.0/24 [110/2] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/23] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
O IA     192.168.1.1 [110/4] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
O IA     192.168.2.2 [110/3] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
O IA     192.168.3.3 [110/2] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
```

Example 7-21 shows the routing tables of R3 and R4 after converting Area 34 to an NSSA.

Example 7-21 *R3 and R4 OSPF NSSA Routing Tables*

```
R3# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     10.12.1.0/24 [110/2] via 10.23.1.2, 00:04:13, GigabitEthernet0/1
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/22] via 10.23.1.2, 00:04:13, GigabitEthernet0/1
      172.31.0.0/24 is subnetted, 1 subnets
O N1     172.31.4.0 [110/22] via 10.34.1.4, 00:03:53, GigabitEthernet0/0
O IA     192.168.1.1 [110/3] via 10.23.1.2, 00:04:13, GigabitEthernet0/1
O        192.168.2.2 [110/2] via 10.23.1.2, 00:04:13, GigabitEthernet0/1
O        192.168.4.4 [110/2] via 10.34.1.4, 00:03:53, GigabitEthernet0/0
```

```
R4# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is 10.34.1.3 to network 0.0.0.0

O*N2  0.0.0.0/0 [110/1] via 10.34.1.3, 00:03:13, GigabitEthernet0/0
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O IA     10.12.1.0/24 [110/3] via 10.34.1.3, 00:03:23, GigabitEthernet0/0
O IA     10.23.1.0/24 [110/2] via 10.34.1.3, 00:03:23, GigabitEthernet0/0
      192.168.1.0/32 is subnetted, 1 subnets
O IA     192.168.1.1 [110/4] via 10.34.1.3, 00:03:23, GigabitEthernet0/0
      192.168.2.0/32 is subnetted, 1 subnets
O IA     192.168.2.2 [110/3] via 10.34.1.3, 00:03:23, GigabitEthernet0/0
      192.168.3.0/32 is subnetted, 1 subnets
O IA     192.168.3.3 [110/2] via 10.34.1.3, 00:03:23, GigabitEthernet0/0
```

Totally NSSAs

Totally stubby areas prohibit Type 3 LSAs (interarea), Type 4 LSAs (ASBR summary LSAs), and Type 5 LSAs (external routes) from entering the area at the ABR, and they prohibit routes from being redistributed within that area. OSPF areas that need to block Type 3 and Type 5 LSAs and still provide the capability of redistributing external networks into OSPF should use the totally NSSA.

Figure 7-17 demonstrates how the LSAs are processed on the ABR for a totally NSSA.

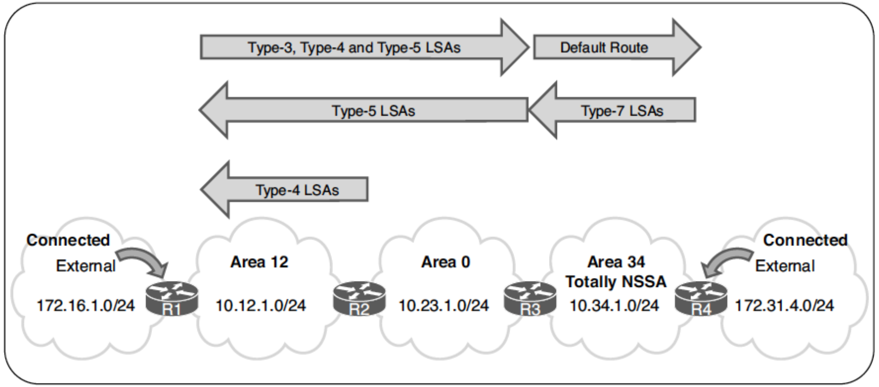

Figure 7-17 *Totally NSSA Concept*

Totally NSSA Configuration

Member routers of a totally NSSA use the same configuration as members of an NSSA. ABRs of a totally stubby area have no summary appended to the configuration. The command **area** *area-id* **nssa no-summary** is configured under the OSPF process.

Example 7-23 shows R3's and R4's OSPF configuration to convert Area 34 into a totally NSSA. Notice the **no-summary** keyword appended to R3's **nssa** command.

Example 7-23 *Totally NSSA Configuration*

```
R3# show run | section router ospf 1
router ospf 1
 router-id 192.168.3.3
 area 34 nssa no-summary
 network 10.23.1.0 0.0.0.255 area 0
 network 10.34.1.0 0.0.0.255 area 34
 network 192.168.3.3 0.0.0.0 area 0
```

```
R4# show run | section router ospf 1
router ospf 1
 router-id 192.168.4.4
 area 34 nssa
 redistribute connected metric-type 1 subnets
 network 10.34.1.0 0.0.0.255 area 34
 network 192.168.4.4 0.0.0.0 area 34
```

Routing Tables Before and After Totally NSSA

Example 7-22 displays the routing tables of R1, R3, and R4 before Area 34 is converted into an OSPF totally NSSA.

Example 7-22 *R1, R3, and R4's Routing Tables Before Area 34 Is a Totally NSSA*

```
R1# show ip route ospf | section 172.31
      172.31.0.0/24 is subnetted, 1 subnets
O E1     172.31.4.0 [110/23] via 10.12.1.2, 00:00:38, GigabitEthernet0/0
```

```
R3# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     10.12.1.0/24 [110/2] via 10.23.1.2, 00:01:34, GigabitEthernet0/1
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/22] via 10.23.1.2, 00:01:34, GigabitEthernet0/1
      172.31.0.0/24 is subnetted, 1 subnets
O E1     172.31.4.0 [110/21] via 10.34.1.4, 00:01:12, GigabitEthernet0/0
O IA     192.168.1.1 [110/3] via 10.23.1.2, 00:01:34, GigabitEthernet0/1
O        192.168.2.2 [110/2] via 10.23.1.2, 00:01:34, GigabitEthernet0/1
O        192.168.4.4 [110/2] via 10.34.1.4, 00:01:12, GigabitEthernet0/0
```

```
R4# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O IA     10.12.1.0/24 [110/3] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
O IA     10.23.1.0/24 [110/2] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/23] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
O IA     192.168.1.1 [110/4] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
O IA     192.168.2.2 [110/3] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
O IA     192.168.3.3 [110/2] via 10.34.1.3, 00:02:28, GigabitEthernet0/0
```

Example 7-24 shows the routing tables of R3 and R4 after Area 34 is converted into a totally NSSA.

Example 7-24 *R3's and R4's Routing Tables After Area 34 Is Made a Totally NSSA*

```
R3# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     10.12.1.0/24 [110/2] via 10.23.1.2, 00:02:14, GigabitEthernet0/1
      172.16.0.0/24 is subnetted, 1 subnets
O E1     172.16.1.0 [110/22] via 10.23.1.2, 00:02:14, GigabitEthernet0/1
      172.31.0.0/24 is subnetted, 1 subnets
O N1     172.31.4.0 [110/22] via 10.34.1.4, 00:02:04, GigabitEthernet0/0
O IA     192.168.1.1 [110/3] via 10.23.1.2, 00:02:14, GigabitEthernet0/1
O        192.168.2.2 [110/2] via 10.23.1.2, 00:02:14, GigabitEthernet0/1
O        192.168.4.4 [110/2] via 10.34.1.4, 00:02:04, GigabitEthernet0/0
```

```
R4# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is 10.34.1.3 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/2] via 10.34.1.3, 00:04:21, GigabitEthernet0/0
```

OSPF Path Selection

- OSPF executes Dijkstra’s shortest path first (SPF) algorithm to create a loop-free topology of shortest paths.
- All routers use the same logic to calculate the shortest path for each network.
- Path selection prioritizes paths by using the following logic:
  1. Intra-area
  2. Interarea
  3. External Type 1
  4. External Type 2

Link Costs

Interface cost is an essential component for Dijkstra's SPF calculation because the shortest path metric is based on the cumulative interface cost (that is, metric) from the router to the destination. OSPF assigns the OSPF link cost (that is, metric) for an interface using the formula in Figure 7-18.

The default reference bandwidth is 100 Mbps. There is no differentiation in the link cost associated to a Fast Ethernet interface and a 10-Gigabit Ethernet interface. Under the OSPF process, the command **auto-cost reference-bandwidth** *bandwidth-in-mbps* changes the reference bandwidth for all OSPF interfaces associated with that process. If the reference bandwidth is changed on one router, it is a best practice to set the same reference bandwidth for all OSPF routers. The OSPF cost can be set manually with the command **ip ospf cost** *1-65535* under the interface.

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

Figure 7-18 *OSPF Interface Cost Formula*

Intra-Area Routes

If multiple intra-area routes exist, the path with the lowest total path metric is installed in the OSPF Routing Information Base (RIB) and is then presented to the router's global RIB. If there is a tie in metric, both routes are installed into the OSPF RIB.

In Figure 7-19, R1 is computing the route to 10.4.4.0/24. Instead of taking the faster Ethernet connection (R1→R2→R4), R1 takes the path across the slower serial link (R1→R3→R4) to R4 because that is the intra-area path.

Example 7-25 shows R1's routing table entry for the 10.4.4.0/24 network.

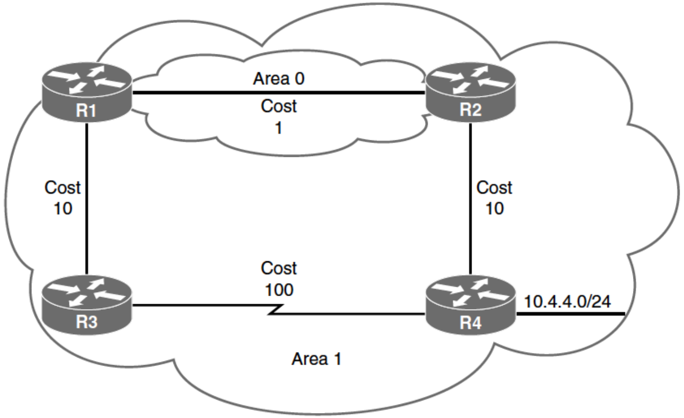

Figure 7-19 *Intra-Area Routes over Interarea Routes*

Example 7-25 *R1's Routing Table for the 10.4.4.0/24 Network*

```
R1# show ip route 10.4.4.0
Routing entry for 10.4.4.0/24
  Known via "ospf 1", distance 110, metric 111, type intra area
  Last update from 10.13.1.3 on GigabitEthernet0/1, 00:00:42 ago
  Routing Descriptor Blocks:
  * 10.13.1.3, from 10.34.1.4, 00:00:42 ago, via GigabitEthernet0/1
      Route metric is 111, traffic share count is 1
```

Interarea Route Selection

The next priority for selecting a path to a network is selection of the path with the lowest total path metric to the destination. If there is a tie in metric, both routes are installed in the OSPF RIB. All interarea paths for a route must go through Area 0 to be considered. In Figure 7-20, R1 is computing the path to R6. R1 uses the path R1→R3→R5→R6 because its total path metric is 35 rather than the R1→R2→R4→R6 path with a metric of 40.

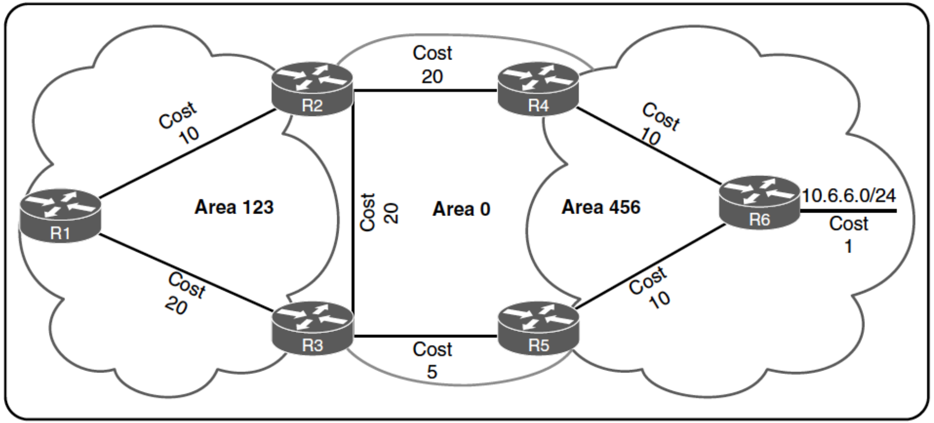

Figure 7-20 *Interarea Route Selection*

External Route Selection

External routes are classified as Type 1 or Type 2. The main differences between Type 1 and Type 2 external OSPF routes are as follows:

- Type 1 routes are preferred over Type 2 routes.
- The Type 1 metric equals the redistribution metric plus the total path metric to the ASBR. In other words, as the LSA propagates away from the originating ASBR, the metric increases.
- The Type 2 metric equals only the redistribution metric. The metric is the same for the router next to the ASBR as it is for the router 30 hops away from the originating ASBR. This is the default external metric type that OSPF uses.

E1 and N1 External Routes

External OSPF Type 1 route calculation uses the redistribution metric plus the lowest path metric to reach the ASBR that advertised the network.

Type 1 path metrics are lower for routers closer to the originating ASBR, whereas the path metric is higher for a router 10 hops away from the ASBR. If there is a tie in the path metric, both routes are installed into the RIB.

If the ASBR is in a different area, the path of the traffic must go through Area 0.

An ABR does not install O E1 and O N1 routes into the RIB at the same time. O N1 is always given preference for a typical NSSA, and its presence prevents the O E1 from being installed on the ABR.

OSPF Path Selection

E2 and N2 External Routes

External OSPF Type 2 routes do not increment in metric, regardless of the path metric to the ASBR. If there is a tie in the redistribution metric, the router compares the forwarding cost.

The forwarding cost is the metric to the ASBR that advertised the network, and the lower forwarding cost is preferred. If there is a tie in forwarding cost, both routes are installed into the routing table.

An ABR does not install O E2 and O N2 routes into the RIB at the same time. O N2 is always given preference for a typical NSSA, and its presence prevents the O E2 from being installed on the ABR.

E2 and N2 External Routes (Cont.)

Figure 7-21 shows the topology for R1 computing a path to the external network (172.16.0.0/24) that is being redistributed.

Example 7-26 shows R1's metric and forwarding metric to the 172.16.0.0/24 network.

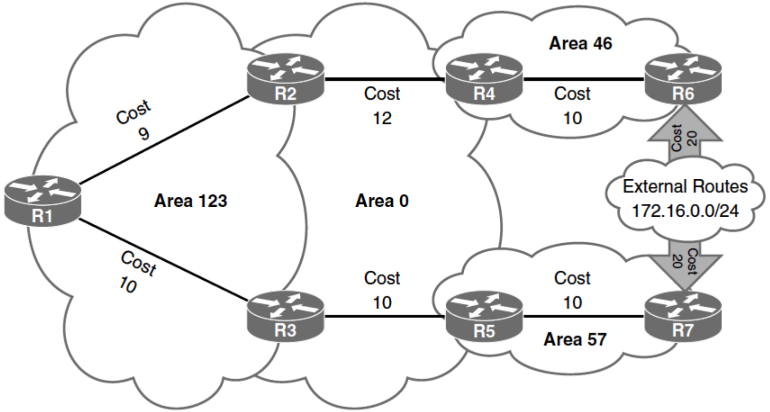

Figure 7-21 *External Type 2 Route Selection Topology*

Example 7-26 *OSPF Forwarding Metric*

```
R1# show ip route 172.16.0.0
Routing entry for 172.16.0.0/24
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 30
  Last update from 10.13.1.3 on GigabitEthernet0/1, 00:12:40 ago
  Routing Descriptor Blocks:
  * 10.13.1.3, from 192.168.7.7, 00:12:40 ago, via GigabitEthernet0/1
      Route metric is 20, traffic share count is 1
```

Equal-Cost Multipathing

If OSPF identifies multiple paths in the algorithms discussed so far in this chapter, those routes are installed into the routing table using equal-cost multipathing (ECMP).

The default maximum number of ECMP paths is four.

The default ECMP setting can be overwritten with the command **maximum-paths** *maximum-paths* under the OSPF process to modify the default setting.

Summarization of Routes

- Splitting up an OSPF routing domain into multiple areas reduces the size of each area’s LSDB.
- While the number of routers and networks remains the same within the OSPF routing domain, the detailed Type 1 and Type 2 LSAs are exchanged for simpler Type 3 LSAs.

Summarization Fundamentals

Summarization of routes helps SPF calculations run faster. Because all routers within an area must maintain an identical copy of the LSDB, summarization occurs between areas on the ABRs. Summarization can eliminate the SPF calculation outside the area for the summarized prefixes because the smaller prefixes are hidden.

Figure 7-23 provides a simple network topology in which the serial link significantly adds to the path metric, and all traffic uses the path through R2 to reach the 172.16.46.0/24 network.

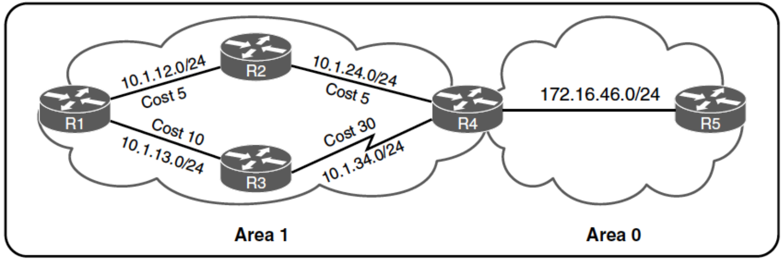

Figure 7-23 *The Impact of Summarization on SPF Topology Calculation*

Figure 7-24 shows the networks in Area 1 being summarized at the ABR into the aggregate 10.1.0.0/18 prefix.

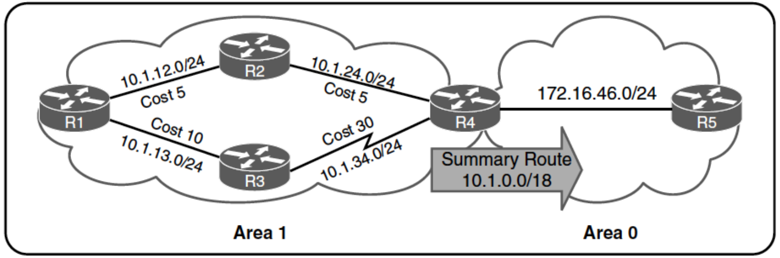

Figure 7-24 *Topology Example with Summarization*

Interarea Summarization

Interarea summarization reduces the number of Type 3 LSAs that an ABR advertises into an area when it receives Type 1 LSAs. The network summarization range is associated with a specific source area for Type 1 LSAs.

When a Type 1 LSA in the summarization range reaches the ABR from the source area, the ABR creates a Type 3 LSA for the summarized network range. The ABR suppresses the more specific Type 1 LSAs, thereby reducing the number of Type 1 LSAs being generated. Interarea summarization does not impact the Type 1 LSAs within the source area.

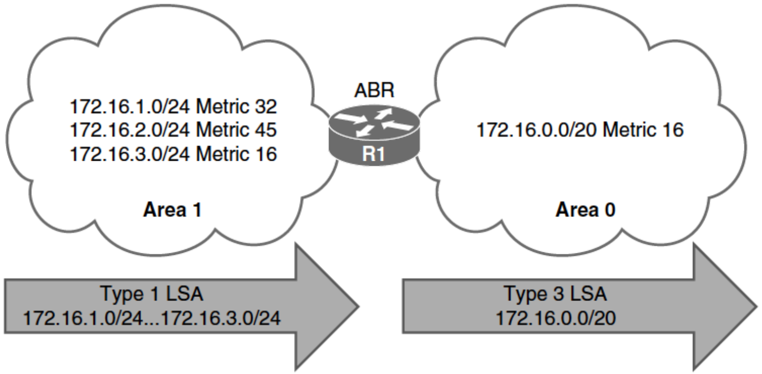

Figure 7-25 *OSPF Interarea Summarization Concept*

Figure 7-25 illustrates this concept, with the three Type 1 LSAs (172.16.1.0/24, 172.16.2.0/24, and 172.16.3.0/24) being summarized into one Type 3 LSA, as the 172.16.0.0/20 network.

Configuration of Interarea Summarization

You define the summarization range and associated area by using the command **area** *area-id* **range** *network subnet-mask* [**advertise** | **not-advertise**] [**cost** *metric*] under the OSPF process. The default behavior is to advertise the summary prefix, so the keyword advertise is not necessary. Appending cost metric to the command statically sets the metric on the summary route.

Figure 7-26 provides a topology example in which R1 is advertising the 172.16.1.0/24, 172.16.2.0/24, and 172.16.3.0/24 networks.

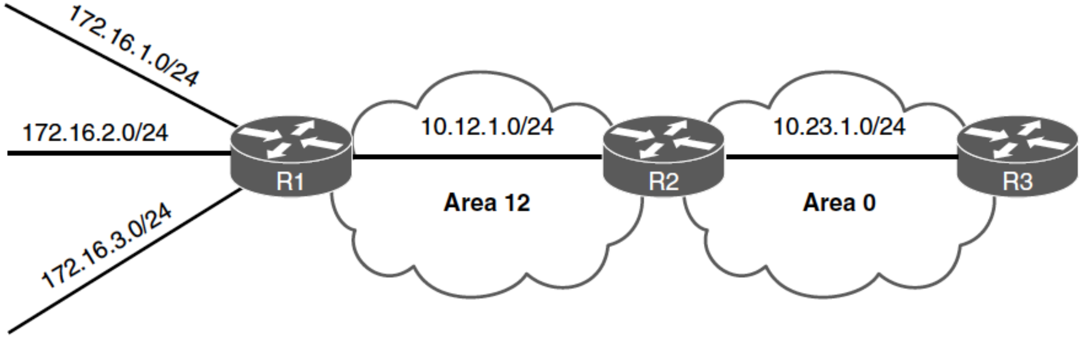

Figure 7-26 *OSPF Interarea Summarization Example*

Interarea Route Summarization Configuration

Example 7-27 shows the routing table on R3 before summarization. Notice that the 172.16.1.0/24, 172.16.2.0/24, and 172.16.3.0/24 networks are all present.

Example 7-27 *Routing Table Before OSPF Interarea Route Summarization*

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

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     10.12.1.0/24 [110/20] via 10.23.1.2, 00:02:22, GigabitEthernet0/1
      172.16.0.0/24 is subnetted, 3 subnets
O IA     172.16.1.0 [110/3] via 10.23.1.2, 00:02:12, GigabitEthernet0/1
O IA     172.16.2.0 [110/3] via 10.23.1.2, 00:02:12, GigabitEthernet0/1
O IA     172.16.3.0 [110/3] via 10.23.1.2, 00:02:12, GigabitEthernet0/1
```

Example 7-28 provides R2's configuration for interarea summarization into an aggregate route 172.16.0.0/16. A static cost of 45 is added to the summary route to reduce CPU load if any of the three networks flap.

Example 7-28 *R2's Interarea Route Summarization Configuration*

```
router ospf 1
 router-id 192.168.2.2
 area 12 range 172.16.0.0 255.255.0.0 cost 45
 network 10.12.0.0 0.0.255.255 area 12
 network 10.23.0.0 0.0.255.255 area 0
```

Interarea Summarization Verification

Example 7-29 shows R3's routing table for verification that the smaller routes were suppressed while the summary route was aggregated.

The ABR performing interarea summarization installs discard routes, which are routes to the Null0 interface that match the summarized network.

Example 7-30 shows the discard route to Null0 on R2.

Example 7-29 *Routing Table After OSPF Interarea Route Summarization*

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

      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O IA    10.12.1.0/24 [110/2] via 10.23.1.2, 00:02:04, GigabitEthernet0/1
O IA  172.16.0.0/16 [110/46] via 10.23.1.2, 00:00:22, GigabitEthernet0/1
```

Example 7-30 *Discard Route for Loop Prevention*

```
R2# show ip route ospf | begin Gateway
Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
O        172.16.0.0/16 is a summary, 00:03:11, Null0
O        172.16.1.0/24 [110/2] via 10.12.1.1, 00:01:26, GigabitEthernet0/0
O        172.16.2.0/24 [110/2] via 10.12.1.1, 00:01:26, GigabitEthernet0/0
O        172.16.3.0/24 [110/2] via 10.12.1.1, 00:01:26, GigabitEthernet0/0
```

External Summarization

External summarization reduces the number of external LSAs in an OSPF domain. An external network summarization range is configured on the ASBR router, and network prefixes that match the network range do not generate a Type 5/Type 7 LSA for the specific prefix. Instead, a Type 5/Type 7 LSA with the external network summarization range is created, and the smaller routes in the summary range are suppressed.

Figure 7-27 demonstrates the concept with the external network summarization range 17.16.0.0/20 configured on the ASBR (R6). The ASBR creates only one Type 5/Type 7 LSA in Area 56 when EIGRP redistributes routes into OSPF.

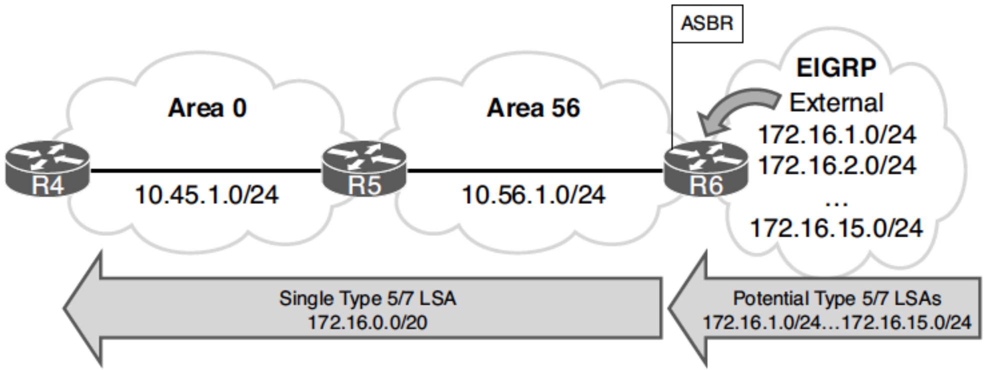

Figure 7-27 *External Summarization Concept*

External Summarization Configuration

Example 7-31 provides the routing table on R5 before external route summarization.

To configure external summarization, you use the command **summary-address** *network subnet-mask* under the OSPF process.

Example 7-32 demonstrates the configuration for external route summarization on R6 (the ASBR).

Example 7-31 *Routing Table Before External Summarization*

```
R5# show ip route ospf | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
O IA     10.3.3.0/24 [110/67] via 10.45.1.4, 00:01:58, GigabitEthernet0/0
O IA     10.24.1.0/29 [110/65] via 10.45.1.4, 00:01:58, GigabitEthernet0/0
O IA     10.123.1.0/24 [110/66] via 10.45.1.4, 00:01:58, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 15 subnets
O E2     172.16.1.0 [110/20] via 10.56.1.6, 00:01:00, GigabitEthernet0/1
O E2     172.16.2.0 [110/20] via 10.56.1.6, 00:00:43, GigabitEthernet0/1
..
O E2     172.16.14.0 [110/20] via 10.56.1.6, 00:00:19, GigabitEthernet0/1
O E2     172.16.15.0 [110/20] via 10.56.1.6, 00:00:15, GigabitEthernet0/1
```

Example 7-32 *OSPF External Summarization Configuration*

```
R6
router ospf 1
 router-id 192.168.6.6
 summary-address 172.16.0.0 255.255.240.0
 redistribute eigrp 1 subnets
 network 10.56.1.0 0.0.0.255 area 56
```

Routing Table After External Summarization

Example 7-33 shows R5's routing table, which verifies that the component routes were summarized into the 172.16.0.0/20 summary network.

Example 7-33 *Routing Table After External Summarization*

```
R5# show ip route ospf | begin Gateway
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
O IA     10.3.3.0/24 [110/67] via 10.45.1.4, 00:04:55, GigabitEthernet0/0
O IA     10.24.1.0/29 [110/65] via 10.45.1.4, 00:04:55, GigabitEthernet0/0
O IA     10.123.1.0/24 [110/66] via 10.45.1.4, 00:04:55, GigabitEthernet0/0
      172.16.0.0/20 is subnetted, 1 subnets
O E2     172.16.0.0 [110/20] via 10.56.1.6, 00:00:02, GigabitEthernet0/1
```

```
R5# show ip route 172.16.0.0 255.255.240.0
Routing entry for 172.16.0.0/20
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Last update from 10.56.1.6 on GigabitEthernet0/1, 00:02:14 ago
  Routing Descriptor Blocks:
  * 10.56.1.6, from 192.168.6.6, 00:02:14 ago, via GigabitEthernet0/1
      Route metric is 20, traffic share count is 1
```

The summarizing ASBR installs a discard route to Null0 that matches the summarized network range as part of a loop-prevention mechanism. Example 7-34 shows the routing table of R6 with the external summary discard route.

Example 7-34 *R6 Discard Route Verification*

```
R6# show ip route ospf | begin Gateway
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
O IA     10.3.3.0/24 [110/68] via 10.56.1.5, 00:08:36, GigabitEthernet0/1
O IA     10.24.1.0/29 [110/66] via 10.56.1.5, 00:08:36, GigabitEthernet0/1
O IA     10.45.1.0/24 [110/2] via 10.56.1.5, 00:08:36, GigabitEthernet0/1
O IA     10.123.1.0/24 [110/67] via 10.56.1.5, 00:08:36, GigabitEthernet0/1
      172.16.0.0/16 is variably subnetted, 15 subnets, 3 masks
O        172.16.0.0/20 is a summary, 00:03:52, Null0
```

Discontiguous Network

- A network engineer who does not fully understand OSPF design might create a topology where interarea traffic crosses through a non-backbone area.

Discontiguous Network

While R2 and R3 have OSPF interfaces in Area 0, traffic from Area 12 must cross Area 234 to reach Area 45. An OSPF network with this design is discontiguous because interarea traffic is trying to cross a non-backbone area.

The routing tables on R2 and R3 in Figure 7-29, looks like routes are being advertised across area 234.

The 10.45.1.0/24 network was received by R4, injected into R4's Area 0 and then re-advertised to Area 234, where R2 installs the route.

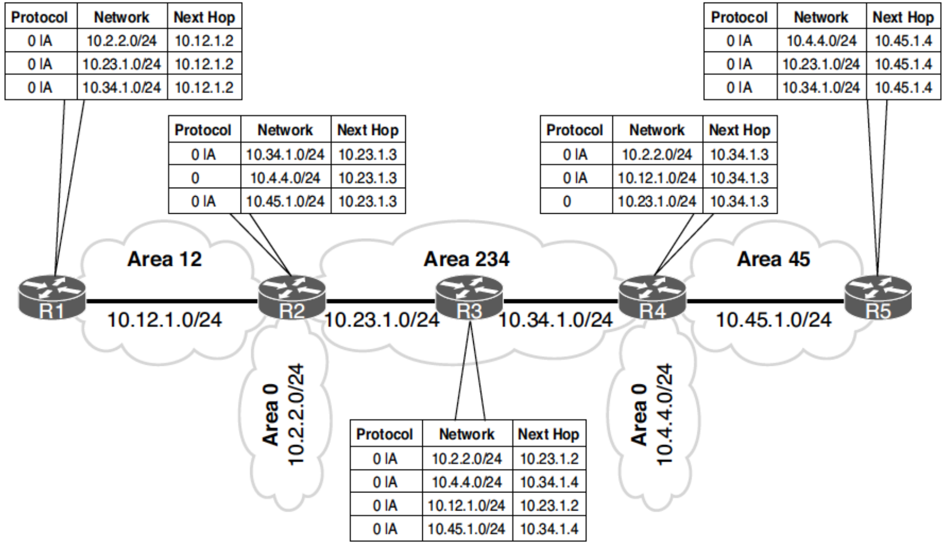

Figure 7-29 *OSPF Routes for a Discontiguous Network*

Discontiguous Network (Cont.)

Most people would assume that R1 would learn about the route learned by Area 45 because R4 is an ABR. However, they would be wrong. ABRs follow three fundamental rules for creating Type 3 LSAs:

- Type 1 LSAs received from an area create Type 3 LSAs into backbone area and non-backbone areas.
- Type 3 LSAs received from Area 0 are created for the non-backbone area.
- Type 3 LSAs received from a non-backbone area are only inserted into the LSDB for the source area. ABRs do not create a Type 3 LSA for the other areas (including a segmented Area 0).

The simplest fix for a discontiguous network is to ensure that Area 0 is contiguous and convert the interfaces on R2, R3, and R4 for the 10.23.1.0/24 and 10.34.1.0/24 networks to be members of Area 0.

Virtual Links

- OSPF virtual links provide a method to overcome discontiguous networks.

Virtual Links

Using a virtual link is similar to running a virtual tunnel within OSPF between an ABR and another multi-area OSPF router. The tunnel belongs to the backbone (Area 0), and therefore the router terminating the virtual link becomes an ABR if it does not have an interface already associated to Area 0. Figure 7-30 shows a contiguous backbone between R2 and R4, with a virtual link across Area 234. With the virtual link established, the routes from Area 12 are advertised into Area 45 and vice versa.

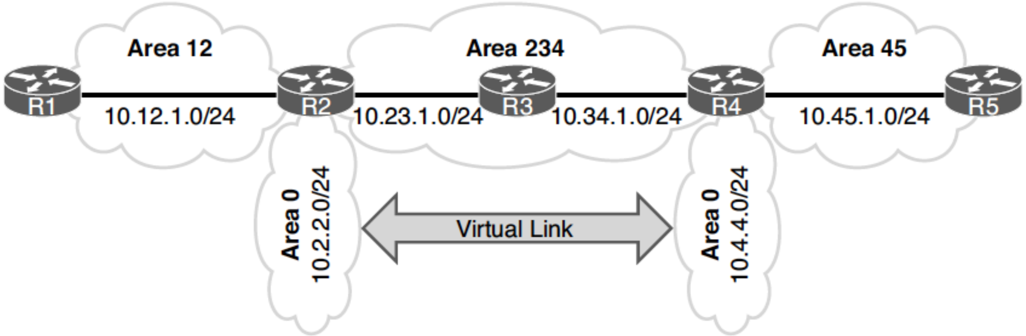

Figure 7-30 *Discontiguous Network*

Virtual Links (Cont.)

Virtual links are built between routers in the same area. The area in which the virtual link endpoints are established is known as the transit area. Each router identifies the remote router by its RID. The virtual link can be one hop or multiple hops away from the remote device. The virtual link is built using Type 1 LSAs, where the neighbor state is Type 4.

You configure the virtual link by using the command **area** *area-id* **virtual-link** *endpoint-RID*. The configuration is performed on both endpoints of the virtual link. At least one endpoint virtual link router has to be a member of Area 0, and virtual links cannot be formed on any OSPF stubby areas.

Virtual Link Configuration

Virtual links are built between routers in the same area. The area in which the virtual link endpoints are established is known as the transit area. Each router identifies the remote router by its RID. The virtual link can be one hop or multiple hops away from the remote device. The virtual link is built using Type 1 LSAs, where the neighbor state is Type 4.

You configure the virtual link by using the command **area** *area-id* **virtual-link** *endpoint-RID*. The configuration is performed on both endpoints of the virtual link. At least one endpoint virtual link router has to be a member of Area 0, and virtual links cannot be formed on any OSPF stubby areas.

Example 7-35 *OSPF Virtual Link Configuration*

```
R2
router ospf 1
 router-id 192.168.2.2
 area 234 virtual-link 192.168.4.4
 network 10.2.2.2 0.0.0.0 area 0
 network 10.12.1.2 0.0.0.0 area 12
 network 10.23.1.2 0.0.0.0 area 234
```

```
R4
router ospf 1
 router-id 192.168.4.4
 area 234 virtual-link 192.168.2.2
 network 10.4.4.4 0.0.0.0 area 0
 network 10.34.1.4 0.0.0.0 area 234
 network 10.45.1.4 0.0.0.0 area 45
```

Virtual Links Verification

To verify the virtual link status, you use the command **show ip ospf virtual-links**. Example 7-36 shows the output.

A virtual link appears as a specific interface, as demonstrated in Example 7-37. Notice that the cost here is 2, which accounts for the metrics between R2 and R4.

Example 7-37 *OSPF Virtual Link as an OSPF Interface*

```
R4# show ip ospf interface brief
Interface    PID   Area         IP Address/Mask    Cost  State Nbrs F/C
Gi0/2        1     0            10.4.4.4/24        1     DR    0/0
VL0          1     0            10.34.1.4/24       2     P2P   1/1
Lo0          1     34           192.168.4.4/32     1     DOWN  0/0
Gi0/1        1     45           10.45.1.4/24       1     BDR   1/1
Gi0/0        1     234          10.34.1.4/24       1     BDR   1/1
```

Example 7-36 *OSPF Virtual Link Verification*

```
R2# show ip ospf virtual-links
Virtual Link OSPF_VL0 to router 192.168.4.4 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 234, via interface GigabitEthernet0/1
 Topology-MTID    Cost    Disabled     Shutdown      Topology Name
        0           2         no          no            Base
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
    Adjacency State FULL (Hello suppressed)
    Index 1/1/3, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

```
R4# show ip ospf virtual-links
! Output omitted for brevity
Virtual Link OSPF_VL0 to router 192.168.2.2 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 234, via interface GigabitEthernet0/0
 Topology-MTID    Cost    Disabled     Shutdown      Topology Name
        0           2         no          no            Base
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
    Adjacency State FULL (Hello suppressed)
```

Routing Table After Virtual Link

The virtual link routers form a point-to-point adjacency, as demonstrated in Example 7-38. Notice that R4 views R2 as a neighbor even though it is not directly connected.

Example 7-39 shows the routing tables of R1 and R5 after the virtual link is established.

Example 7-39 *R1's and R5's Routing Tables After the Virtual Link is Created*

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

      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O IA     10.2.2.0/24 [110/2] via 10.12.1.2, 00:00:10, GigabitEthernet0/0
O IA     10.4.4.0/24 [110/4] via 10.12.1.2, 00:00:05, GigabitEthernet0/0
O IA     10.23.1.0/24 [110/2] via 10.12.1.2, 00:00:10, GigabitEthernet0/0
O IA     10.34.1.0/24 [110/3] via 10.12.1.2, 00:00:10, GigabitEthernet0/0
O IA     10.45.1.0/24 [110/4] via 10.12.1.2, 00:00:05, GigabitEthernet0/0
```

```
R5# show ip route ospf | begin Gateway
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O IA     10.2.2.0/24 [110/4] via 10.45.1.4, 00:00:43, GigabitEthernet0/1
O IA     10.4.4.0/24 [110/2] via 10.45.1.4, 00:01:48, GigabitEthernet0/1
O IA     10.12.1.0/24 [110/4] via 10.45.1.4, 00:00:43, GigabitEthernet0/1
O IA     10.23.1.0/24 [110/3] via 10.45.1.4, 00:01:48, GigabitEthernet0/1
O IA     10.34.1.0/24 [110/2] via 10.45.1.4, 00:01:48, GigabitEthernet0/1
```

Example 7-38 *A Virtual Link Displayed as an OSPF Neighbor*

```
R4# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
192.168.2.2       0   FULL/  -            -       10.23.1.2       OSPF_VL0
192.168.5.5       1   FULL/DR         00:00:34    10.45.1.5       GigabitEthernet0/1
192.168.3.3       1   FULL/DR         00:00:38    10.34.1.3       GigabitEthernet0/0
```

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 7

Description	
Link-state advertisements	OSPF Type 5 LSA flooding
LSA Type 1: router links	LSA Type 4: ASBR summary
LSA Type 1 link types	OSPF Type 4 and Type 5 LSA flooding
LSA Type 2: network link	LSA Type 7: NSSA external summary
Visualization of Area 1234 with Type 1 and Type 2 LSAs	OSPF Type 7 LSA advertisement
LSA Type 3: summary link	Overview of LSA types from a sample topology
LSA Type 3 metric calculations	Stub area
LSA Type 3 topology	OSPF Stub area concept
LSA Type 5: external routes	Totally stubby areas

cisco

Prepare for the Exam

Key Topics for Chapter 7 (Cont.)

Description	
Totally stubby area concept	E2 and N2 external routes
Not-so-stubby area	LSA reduction through area segmentation
Not-so-stubby area concept	Interarea summarization
Totally not-so-stubby area	Summarization of Type 1 LSAs
Totally not-so-stubby area concept	Summarization of discard routes
OSPF path selection	External summarization
Intra-area routes	A discontiguous network
Interarea routes	Virtual links
E1 and N1 external routes	Locating virtual link endpoints, by RID

Prepare for the Exam

Key Terms for Chapter 7

Key Terms	
area border router (ABR)	NSSA external LSA
autonomous system border router (ASBR)	OSPF stub area
backbone area	OSPF totally stubby area
Router LSAs	OSPF not-so-stubby area (NSSA)
network LSA	OSPF totally NSSA
summary LSA	discontiguous network
ASBR summary LSA	virtual link
external LSA	Locating virtual link endpoints

Command Reference for Chapter 7

Task	Command Syntax
Initialize the OSPF process	**router ospf** *process-id*
Display the generic LSA listings in a router's LSDB and the type and count for each type from an advertising router	**show ip ospf database**
Display the specific information for a Type 1, 2, 3, 4, 5, or 7 LSA	**show ip ospf database {router | network | summary | asbr-summary | external | nssa-external}**
Configure all routers in an OSPF area as an OSPF stub area	**area** *area-id* **stub**
Configure an ABR as a totally stubby area router ABR	**area** *area-id* **stub no-summary**
Configure all routers in an OSPF area as an OSPF NSSA area	**area** *area-id* **nssa**

Command Reference for Chapter 7 (Cont.)

Task	Command Syntax
Configure an ABR for an NSSA with optional default route injection	**area** *area-id* **nssa** [**default-information-originate**]
Configure an ABR as a totally NSSA router ABR	**area** *area-id* **nssa no-summary**
Modify the OSPF reference bandwidth for dynamic interface metric costing	**auto-cost reference-bandwidth** *bandwidth-in-mbps*
Statically set the OSPF metric for an interface	**ip ospf cost** *1-65535*
Configure internal route summarization on the first ABR attached to the source network	**area** *area-id* range *network subnet-mask* [*advertise* | *not-advertise*] [*cost metric*]
Configure external route summarization on the ASBR	**summary-address** *network subnet-mask*
Configure an OSPF virtual link to extend Area 0	**area** *area-id* **virtual-link** *endpoint-RID*