

This chapter covers the following topics, which you need to master to pass the CCNP/CCDP/CCIP BSCI exam. This list of exam objectives stated for the Building Scalable Cisco Internetworks (BSCI) exam (640-901) was obtained from the following web page:

[http://www.cisco.com/warp/public/10/wwtraining/certprog/testing/current\\_exams/640-901.html](http://www.cisco.com/warp/public/10/wwtraining/certprog/testing/current_exams/640-901.html)

Cisco often updates these pages, and this URL may change. If you have problems, go to Cisco.com and follow the directions to the exam certifications.

- The issues with interconnecting multiple areas.
- The differences between the possible types of areas, routers, and LSAs.
- How OSPF operates across multiple areas using NBMA.
- How OSPF supports the use of VLSM and summarization.
- The Cisco defaults in OSPF, the Cisco commands for implementing OSPF for multiple areas, and Cisco commands for reviewing the configuration.

# Using OSPF Across Multiple Areas

---

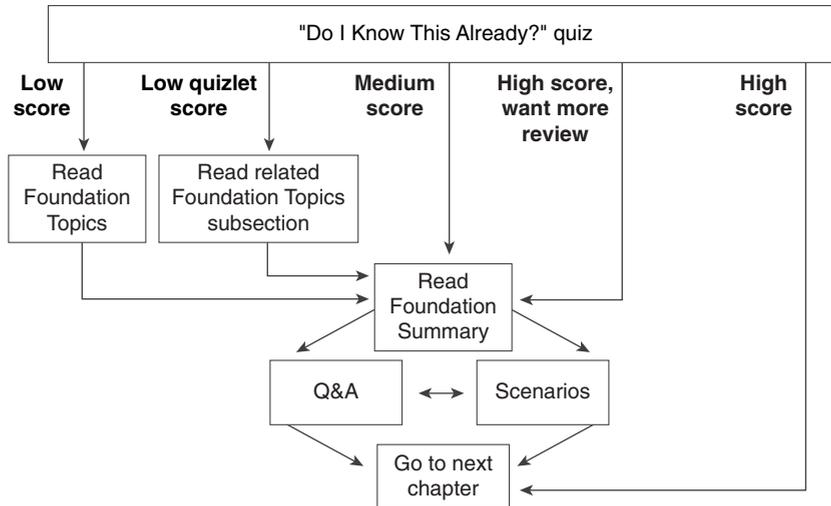
The topics in this chapter detail the routing protocol OSPF across multiple areas. This chapter assumes your knowledge of the previous chapter, which dealt with OSPF concepts and its configuration in a single area. This chapter builds on this understanding and explains how OSPF works within a large multiarea network. Each of these chapters covers two major sections. The first deals theoretically with how the protocol works. The second covers how to implement and manage an OSPF network. This chapter introduces OSPF areas and explains the operation of the protocol across those areas. Both the network communication that the protocol uses and its configuration are explained in this chapter.

The topics in this chapter will directly reflect questions on the exam. OSPF is the industry-standard interior routing protocol designed for use in large networks. Therefore, it is an obligatory subject in an exam on IP routing protocols. Nine percent of the BSCI course material is devoted to interconnecting multiple OSPF areas, and you can expect approximately four to five questions on the BSCI exam to be directly related to this subject.

## How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and the answers for all your work with this book in one place, for easy reference.
- When you take a quiz, write down your answers. Studies show that retention significantly increases by writing down facts and concepts, even if you never look at the information again.
- Use the diagram in Figure 6-1 to guide you to the next step.

**Figure 6-1** *How to Use This Chapter*

If you skip to the Foundation Summary, Q&A, and scenarios sections and have trouble with the material there, you should go back to the Foundation Topics section.

## “Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

This 16-question quiz helps you determine how to spend your limited study time. The quiz is sectioned into four smaller four-question “quizlets,” which correspond to four major topics in the chapter. Figure 6-1 outlines suggestions on how to spend your time in this chapter. Use Table 6-1 to record your scores.

**Table 6-1** *Score Sheet for Quiz and Quizlets*

Quizlet Number	Topic	Questions	Score
1	Issues with connecting multiple areas and NBMA	1 to 4	
2	Differences between routers, areas, and LSAs	5 to 8	
3	Summarization and VLSM	9 to 12	
4	Configuring and verifying OSPF across multiple areas	13 to 16	
All questions	All	1 to 16	

**1** A virtual link in OSPF is used to solve what problem?

---



---



---

**2** State one disadvantage for making an NBMA Frame Relay cloud Area 0.

---



---



---

**3** State one advantage in making the centralized routers and network resources dwell in Area 0 while the Frame Relay cloud and the stub remote LANs reside in satellite stub areas.

---



---



---

**4** How does creating a number of areas in OSPF reduce the number of SPF calculations?

---



---



---

5 How does a stub area differ from the backbone area?

---

---

---

6 How does a totally stubby area differ from a stub area?

---

---

---

7 State the different LSA types.

---

---

---

8 Where does the backbone router reside, and what is its function?

---

---

---

9 Are there any considerations for OSPF configured with VLSM sending routing updates into RIPv1?

---

---

---

10 There are two types of summarization. What are they?

---

---

---

- 11** Can the following subnets with a mask of 255.255.255.0 be summarized? If so, state the subnet and mask that can be used.

19.44.16.0	19.44.24.0
19.44.17.0	19.44.25.0
19.44.18.0	19.44.26.0
19.44.19.0	19.44.27.0
19.44.20.0	19.44.28.0
19.44.21.0	19.44.29.0
19.44.22.0	19.44.30.0
19.44.23.0	19.44.31.0

---



---



---

- 12** Why can interarea summarization be configured only on ABRs?

---



---



---

- 13** What command would be used to create a totally stubby area?

---



---



---

- 14** What is a virtual link, and what command would be used to create it?

---



---



---

- 15** Where would you issue the command to summarize IP subnets? State the command that would be used.

---



---



---

16 How would you summarize external routes before injecting them into the OSPF domain?

---

---

---

The answers to this quiz are found in Appendix A, “Answers to Quiz Questions.” The suggested choices for your next step are as follows:

- **2 or less on any quizlet**—Review the appropriate sections of the “Foundation Topics” portion of this chapter, based on Table 6-1. Then move on to the “Foundation Summary” section, the “Q&A” section, and the “Scenarios” at the end of the chapter.
- **8 or less overall score**—Read the entire chapter. This includes the “Foundation Topics” and “Foundation Summary” sections, the “Q&A” section, and the “Scenarios” at the end of the chapter.
- **9 to 12 overall score**—Begin with the “Foundation Summary” section, and then go to the “Q&A” section and the “Scenarios” at the end of the chapter. If you have trouble with these exercises, read the appropriate sections in “Foundation Topics.”
- **13 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” section, and then go to the “Q&A” section and the “Scenarios” at the end of the chapter. Otherwise, move to the next chapter.

## Foundation Topics

The following section examines the implications of using OSPF in a larger network.

### OSPF in a Multiple Area Network

The first consideration must be why multiple areas are needed. There is a lot of noise about multiple areas in OSPF, and indeed it is one of the main distinguishing features between the distance vector protocols and the link-state OSPF.

#### Case Study

Returning to the company Jackanory.com, you will recall that in the previous chapter, Jackanory.com was acquired along with some other companies. The company is still operating intact but must communicate with the other companies and with the holding company that purchased it. Most of the communication occurs with the parent company, although there is occasional work done on a project basis with the other companies. The parent company needs to be capable of communicating with all the small companies that it holds.

The infrastructure within the companies is to stay the same, but a Frame Relay network has been created to connect the sites. OSPF has been implemented internally at Jackanory.com, and it is working well. Now it is necessary to implement OSPF at the other sites and to have a complete understanding of all the networks available.

The network administrators have been charged with making this happen. There is more to be done than simply typing in the appropriate commands, of course. The administrators need to understand and decide on several things. These include whether summarization is possible and desired for the Frame Relay design, what design is appropriate for the OSPF areas, and whether any of the areas should be stub, totally stubby, or not so stubby areas. To make these decisions, the administrators must understand the OSPF operation over multiple areas. This chapter addresses these subjects and then, in the final section, shows the configuration options and requirements.

#### Why Multiple Areas?

An area is a logical grouping of routers that are running OSPF with identical topological databases. It is a subdivision of the greater OSPF domain. The creation of multiple areas solves the problem of a large network outgrowing its capacity to communicate the details of the network to the routing devices charged with maintaining control and connectivity throughout the network.

The division of the AS into areas allows routers in each area to maintain their own topological databases. This limits the size of the topological databases, and summary and external links ensure connectivity between areas and networks outside the AS.

## How to Determine Area Boundaries

Although there is an obvious need for the multiple areas, the practical question of how this is implemented arises. There are two approaches. The first approach is to grow a single area until it becomes unmanageable. The second approach is to design the network with multiple areas, which are very small, in the expectation that the networks will grow to fit comfortably into their areas.

Both approaches are valid. The first approach requires less initial work and configuration. Great care should be put into the design of the network, however, because this may cause problems in the future, particularly in addressing.

In practice, many companies convert their networks into OSPF from a distance vector routing protocol when they realize that they have outgrown the existing routing protocol. This allows the planned implementation of the second approach.

Now consider the implications of implementing the first approach to OSPF—that of configuring one area and adding others as needed. By looking at the issues, you can learn many things beyond just how to create multiple areas.

## Problems with OSPF in a Single Area

To understand the true benefits of multiple areas, consider why someone might decide to create multiple areas from one area.

The following symptoms that you will observe on the network provide a clue that a single area is becoming overpowered:

- The frequency of the SPF algorithm being run will increase. The larger the network, the greater the probability of a network change and, thus, a recalculation of the entire area. Each recalculation will also take longer.
- The routing table will become extremely large. The routing table is *not* sent out wholesale as in a distance vector routing protocol; however, the greater the size of the table, the longer each lookup becomes. The memory requirements on the router also increase.
- The topological database will increase in size and will eventually become unmanageable for the same reasons. The topology table is exchanged between adjacent routers at least every 30 minutes.

- As the various databases increase in size and the calculations become increasingly frequent, the CPU utilization will increase as the available memory decreases. This will make the network response time very sluggish (not because of congestion on the line, but because of congestion within the router itself). It can also cause congestion on the link.

---

**TIP** To check the CPU utilization on the router, use the **show processes cpu** command. To check the memory utilization, issue the **show memory free** command.

---

## OSPF Areas

Now that you understand why the size of the areas should be controlled, it is important to consider the design issues for the different areas, including the technology that underpins them and their communication (both within and between the areas).

### OSPF Within an Area

One of the main strengths of OSPF is its capability to scale and to support large networks. It achieves this by creating areas from groups of subnets. The area is seen internally almost as if it is a small organization or entity on its own. It communicates with the other areas, exchanging routing information; this exchange is kept to a minimum, however, allowing only that which is required for connectivity. All computation is kept within the area.

In this way, a router is not overwhelmed by the entirety of the organization's network. This is crucial because the nature of a link-state routing protocol is more CPU- and memory-intensive.

### Router Types

Given the hierarchical nature of the OSPF network, you will see routers operating within an area, routers connecting areas, and routers connecting the organization or autonomous system to the outside world. Each of these routers will have a different set of responsibilities, depending on their position and functionality within the OSPF hierarchical design.

The following list identifies the different OSPF routers:

- **Internal router**—Within an area, the functionality of the router is straightforward. It is responsible for maintaining a current and accurate database of every subnet within the area. It is also responsible for forwarding data to other networks by the shortest path. Flooding of routing updates is confined to the area. All interfaces on this router are within the same area.

- **Backbone router**—The design rules for OSPF require that all the areas be connected through a single area, known as the *backbone area* or *Area 0* or 0.0.0.0. A router within this area is referred to as a *backbone router*. It may also be an internal router or an Area Border Router.
- **Area Border Router (ABR)**—This router is responsible for connecting two or more areas. It holds a full topological database for each area to which it is connected and sends LSA updates between the areas. These LSA updates are summary updates of the subnets within an area. It is at the area border that summarization should be configured for OSPF because this is where the LSAs make use of the reduced routing updates to minimize the routing overhead on both the network and the routers.
- **Autonomous System Boundary Router (ASBR)**—To connect to the outside world or to any other routing protocol, you need to leave the OSPF domain. OSPF is an *interior routing protocol* or *Interior Gateway Protocol (IGP)*; *gateway* is an older term for a router. The router configured for this duty is the ASBR. If there is any redistribution between other protocols to OSPF on a router, it will become an ASBR because the other routing protocols are outside the OSPF autonomous systems. Although you can place this router anywhere in the OSPF hierarchical design, it should reside in the backbone area. Because any traffic leaving the OSPF domain will also likely leave the router's area, it makes sense to place the ASBR in a central location that all traffic leaving its area must traverse.

Figure 6-2 shows how the different router types are interrelated.

**Figure 6-2** Router Definitions for OSPF

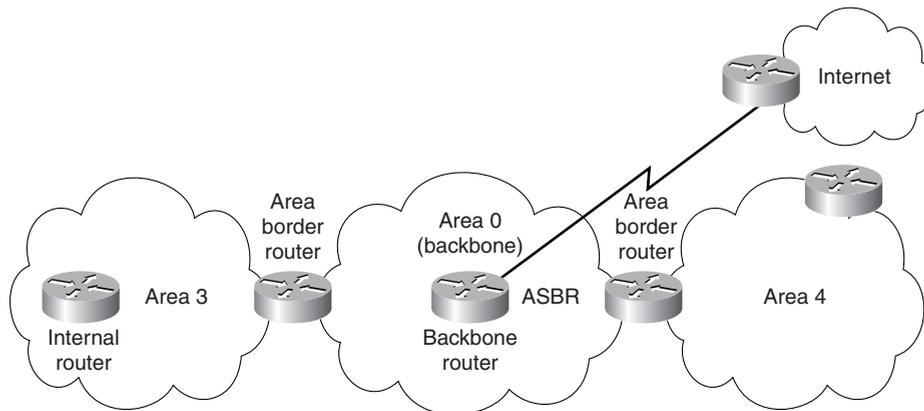
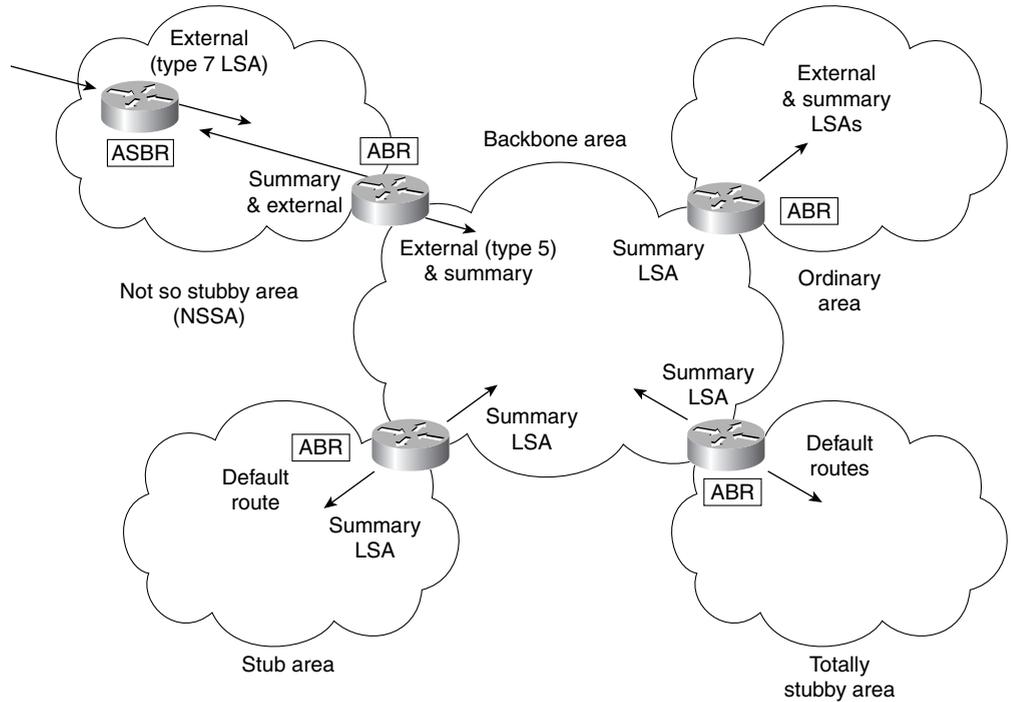


Figure 6-3 shows the connectivity and functionality of the different areas.

The routers will send out routing updates and other network information through LSAs. The function or type of router will determine the LSAs that are sent.

**Figure 6-3** *The Different Types of OSPF Areas and LSA Propagation*



## The Link-State Advertisements

Five commonly used types of link-state advertisements (LSAs) exist. The OSPF standard lists 11 types of LSA, although currently not all of these are in use. Six LSAs are listed here; this discussion focuses only on those advertisements sent between routers in the same area and on the same segment.

The six link-state advertisements are as follows:

**The router link LSA**—This LSA is generated for each area to which the router belongs. This LSA gives the link states to all other routers within an area. This LSA is flooded into an area. This is identified as a Type 1 LSA.

**The network link LSA**—This LSA is sent out by the designated router and lists all the routers on the segment for which it is the designated router and has a neighbor relationship. The LSA is flooded to the whole area. This is identified as a Type 2 LSA.

**The network summary link LSA**—This LSA is sent between areas and summarizes the IP networks from one area to another. It is generated by an ABR. This is identified as a Type 3 LSA.

**The AS external ASBR summary link LSA**—This LSA is sent to a router that connects to the outside world (ASBR). It is sent from the Area Border Router to the Autonomous System Boundary Router. The LSA contains the metric cost from the ABR to the ASBR. This is identified as a Type 4 LSA.

**The external link LSA**—This LSA is originated by AS boundary routers and is flooded throughout the AS. Each external advertisement describes a route to a destination in another autonomous system. Default routes for the AS can also be described by AS external advertisements. This is identified as a Type 5 LSA.

**The NSSA External LSA**—Identified as Type 7, these LSAs are created by the ASBR residing in a not so stubby area (NSSA). This LSA is very similar to an autonomous system external LSA, except that this LSA is contained within the NSSA area and is not propagated into other areas.

Figure 6-4 clearly shows the relationships between the different LSAs. This section discusses the router and network LSAs. The LSAs concerned with communication outside an area are considered later.

## The ABRs and ASBR Propagation of LSAs

When a router is configured as an ABR, it generates summary LSAs and floods them into the backbone area. Routes generated within an area are Type 1 or Type 2, and these are injected as Type 3 summaries into the backbone. These summaries are then injected by the other ABRs into their own areas, unless they are configured as totally stubby areas. Any Type 3 or Type 4 LSA received from the backbone will be forwarded into the area by the ABR.

The backbone will also forward external routes both ways unless the ABR is a stub router, in which case they are blocked.

If a summary is received from within the area, it cannot be forwarded, and summaries received from the backbone cannot be further summarized.

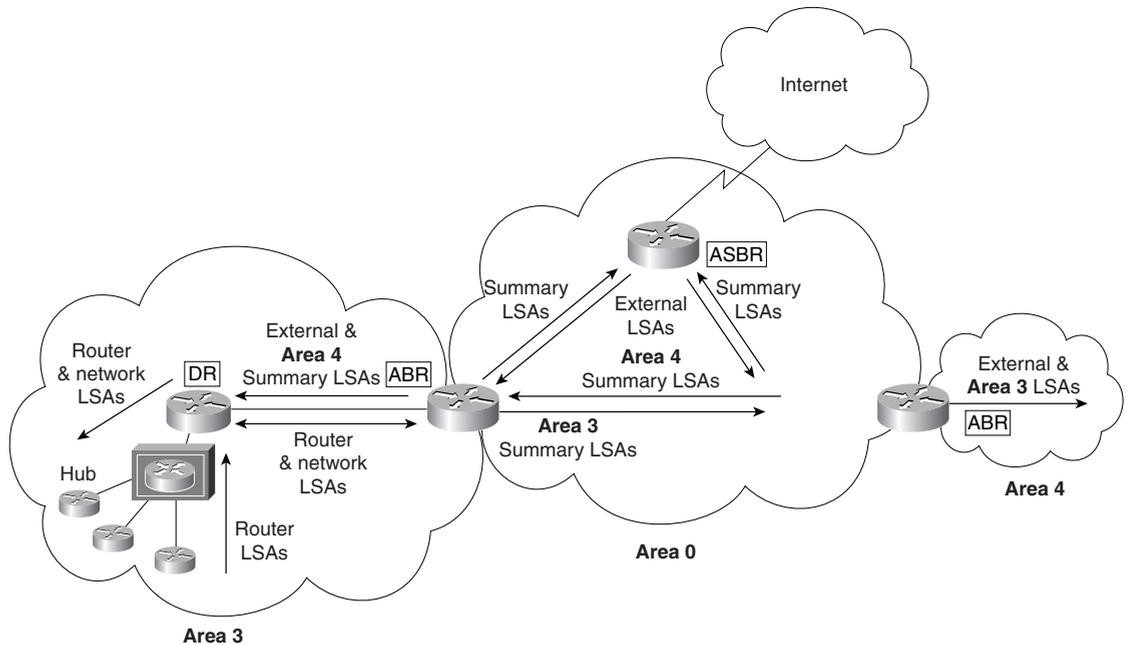
---

**NOTE**

The different types of areas mentioned in this section are described in the later section “The Different Types of Areas.”

---

**Figure 6-4** *The Propagation of LSAs*



It is necessary for certain conditions to be met before any LSAs can be flooded out of all interface:

- The LSA was not received through the interface.
- The interface is in a state of exchange or full adjacency.
- The interface is not connected to a stub area (no LSA Type 5 will be flooded).
- The interface is not connected to a totally stubby area (no Type 3,4, or 5 will be propagated).

## OSPF Path Selection Between Areas

The OSPF routing table that exists on a router depends on the following factors:

- The position that the router has in the area and the status of the network
- The type of area that the router is located in
- Whether there are multiple areas in the domain
- Whether there are communications outside the autonomous system

Remember the sequence of events. The router receives LSAs. It builds the topological database. Then it runs the Dijkstra algorithm, from which the shortest path first is chosen and entered into the routing table. The routing table is therefore the conclusion of the decision-making process. It holds information on how that decision was made by including the metric for each link. This enables the network administrator to view the operation of the network.

Different LSAs hold different weighting in the decision-making process. It is preferable to take an internal route (within the area) to a remote network rather than to traverse multiple areas just to arrive at the same place. Not only does multiple-area traveling create unnecessary traffic, but it also can create a loop within the network.

The routing table reflects the network topology information and indicates where the remote network sits in relation to the local router.

The router will process the LSAs in this order:

- Step 1** The internal LSA (Type 1 and 2).
- Step 2** The LSAs of the AS (Type 3 and 4). If there is a route to the chosen network within the area (Type 1 or 2), this path will be kept.
- Step 3** The external LSAs (Type 5).

## Calculating the Cost of a Path to Another Area

There are paths to networks in other areas, and then there are paths to networks in another AS. The costs of these paths are calculated slightly differently.

### The Path to Another Area

This is calculated as the smallest cost to the ABR, added to the smallest cost to the backbone. Thus, if there were two paths from the ABR into the backbone, the shortest (lowest-cost) path would be added to the cost of the path to the ABR.

### The Path to Another AS

External routes are routes passed between a router within the OSPF domain and a router in another autonomous system or routing domain. The routes discovered by OSPF in this way can have the cost of the path calculated in one of two ways:

- **E1**—The cost of the path to the ASBR is added to the external cost to reach the next-hop router outside the AS.
- **E2**—The external cost of the path from the ASBR is all that is considered in the calculation. This is the default configuration. This is used when there is only one router advertising the route and no path selection is required. If both an E1 and an E2 path are offered to the remote network, the E1 path will be used.

At the side of the routing table is a column indicating the source of the routing information. Typically, this is the routing protocol. In the instance of OSPF, however, it includes the LSA type that provided the path.

Table 6-2 shows the codes used in the routing table.

**Table 6-2** *OSPF Routing Table Codes and Associated LSAs*

LSA Type	Routing Table Entry	Description
1 Router Link	O	This is generated by the router, listing all the links to which it is connected, their status, and their cost. It is propagated within the area.
2 Network Link	O	This is generated by the designated router on a multiaccess LAN to the area.
3 or 4 Summary Link (between areas)	OIA	LSA Type 3 includes the networks or subnets within an area that may have been summarized and that are sent into the backbone and between ABRs. LSA Type 4 is information sent to the ASBR from the ABR. These routes are not sent into totally stubby areas.
5 Summary Link/External Link (between autonomous systems)	OE1 or OE2	The routes in this LSA are external to the autonomous system. They can be configured to have one of two values. E1 will include the internal cost to the ASBR added to the external cost reported by the ASBR. E2 does not compute the internal cost—it just reports the external cost to the remote destination.

**TIP** The exam focuses on the LSA Types 1 through 5.

## The Different Types of Areas

The only obligatory area is Area 0, also known as the backbone area or Area 0.0.0.0. In addition to the backbone area, which connects the other areas, OSPF networks use several other types of areas. The following are the different types of areas:

- **An ordinary or standard area**—This type of area connects to the backbone. The area is seen as an entity unto itself. Every router knows about every network in the area, and each router has the same topological database. However, the routing tables will be unique from the perspective of the router and its position within the area.

- **A stub area**—This is an area that will not accept external summary routes. The LSAs blocked are Types 4 (summary link LSAs that are generated by the ABRs) and 5. The consequence is that the only way that a router within the stub area can see outside the autonomous system is via the configuration of a default route. Every router within the area can see every network within the area and the networks (summarized or not) within other areas. It is typically used in a hub-and-spoke network design.
- **A totally stubby area**—This area does not accept summary LSAs from the other areas or the external summary LSAs from outside the autonomous system. The LSAs blocked are Types 3, 4, and 5. The only way out of the totally stubby area is via a configured default route. A default route is indicated as the network 0.0.0.0. This type of area is particularly useful for remote sites that have few networks and limited connectivity with the rest of the network. This is a proprietary solution offered only by Cisco. Cisco recommends this solution if you have a totally Cisco shop because it keeps the topological databases and routing tables as small as possible.
- **A not so stubby area (NSSA)**—This area is used primarily to connect to ISPs, or when redistribution is required. In most respects, it is the same as the stub area. External routes are not propagated into or out of the area. It does not allow Type 4 or Type 5 LSAs. This area was designed as a special stub area for applications like an area with a few stub networks but with a connection to a router that runs only RIP, or an area with its own connection to an Internet resource needed only by a certain division.

A NSSA is an area that is seen as a stub area but that can receive external routes that it will not propagate into the backbone area, and thus the rest of the OSPF domain. Another LSA, Type 7, is created specifically for the NSSA. This LSA may be originated and communicated throughout the area, but it will *not* be propagated into other areas, including Area 0. If the information is to be propagated throughout the AS, it is translated into an LSA Type 5 at the NSSA ABR.

It is not always possible to design the network and determine where redistribution is to occur. RFC 1587 deals with this subject.

- **The backbone area**—This area is often referred to as Area 0, and it connects all the other areas. It can propagate all the LSAs except for LSA Type 7, which would have been translated into LSA Type 5 by the ABR.

Some restrictions govern creating a stub or totally stubby area. Because no external routes are allowed in these areas, the following restrictions are in place:

- No external routes are allowed.
- No virtual links are allowed.
- No redistribution is allowed.
- No ASBR routers are allowed.
- The area is not the backbone area.
- All the routers are configured to be stub routers.

Now that you understand many components of OSPF, it is important to focus on some of the design implications of creating multiple areas. This focus will reinforce the concepts detailed in the chapter.

## Design Considerations in Multiple Area OSPF

The major design consideration in OSPF is how to divide the areas. This is of interest because it impacts the addressing scheme for IP within the network.

An OSPF network works best with a hierarchical design, in which the movement of data from one area to another comprises only a subset of the traffic within the area itself.

It is important to remember that with all the interarea traffic disseminated by the backbone, any reduction of overhead through a solid hierarchical design and summarization is beneficial. The lower the number of summary LSAs that need to be forwarded into the backbone area, the greater the benefit to the entire network. This will allow the network to grow more easily because the network overhead is at a minimum.

With this in mind, summarization is the natural consequence. As shown in Chapter 3, “IP Addressing,” summarization is not something that can be imposed on a network. It must be part of the initial network design. The addressing scheme must be devised to support the use of summarization.

---

**WARNING** Although it is possible to have more than three areas (per router) in OSPF, the Cisco Technical Assistance Center (TAC) recommends that a greater number of areas be created only after careful consideration. The results of having more areas will vary depending on the router (memory and CPU), as well as network topology and how many LSAs are generated. It is recommended that you not exceed 50 routers in an OSPF area, but again, this is a guideline and not a strict rule. Remember that OSPF is very CPU-intensive in its maintenance of the databases and in the flooding of LSAs, as well as when it calculates the routing table, a process based on LSAs. Therefore, it is not strictly the number of routers or areas that is important, but the number of routes and the stability of the network. These issues must be considered because the number of LSAs in your network is proportional to the amount of router resources required.

With this understanding, the general rules stated by Cisco for OSPF design are that the following numbers should not be exceeded:

- Routers per area: 50
  - Neighbors per router: 60
  - Areas per router: 3
  - A router may not be a DR or BDR for more than 1 LAN
-

## Summarization

Two types of summarization exist:

- **Interarea summarization**—This is performed at the ABR and creates Type 3 and 4 LSAs.
- **External summarization**—This is performed at the ASBR and creates Type 5 LSAs.

Both have the same fundamental requirement of contiguous addressing.

OSPF is stringent in its demand for a solid hierarchical design—so much so that it has devised some commands to deal with situations that break its rules of structure.

## The Virtual Link

The main dictate in OSPF is that the multiple areas must all connect directly to the backbone area. The connection to the backbone area is via an ABR, which is resident in both areas and holds a full topological database for each area.

OSPF has provided for the unhappy occasion that this rule cannot be followed. The solution is called a *virtual link*. If the new area cannot connect directly to the backbone area, a router is configured to connect to an area that does have direct connectivity.

The configuration commands create a tunnel to the ABR in the intermediary area. From the viewpoint of OSPF, it has a direct connection.

The reasons such a situation may occur are listed here:

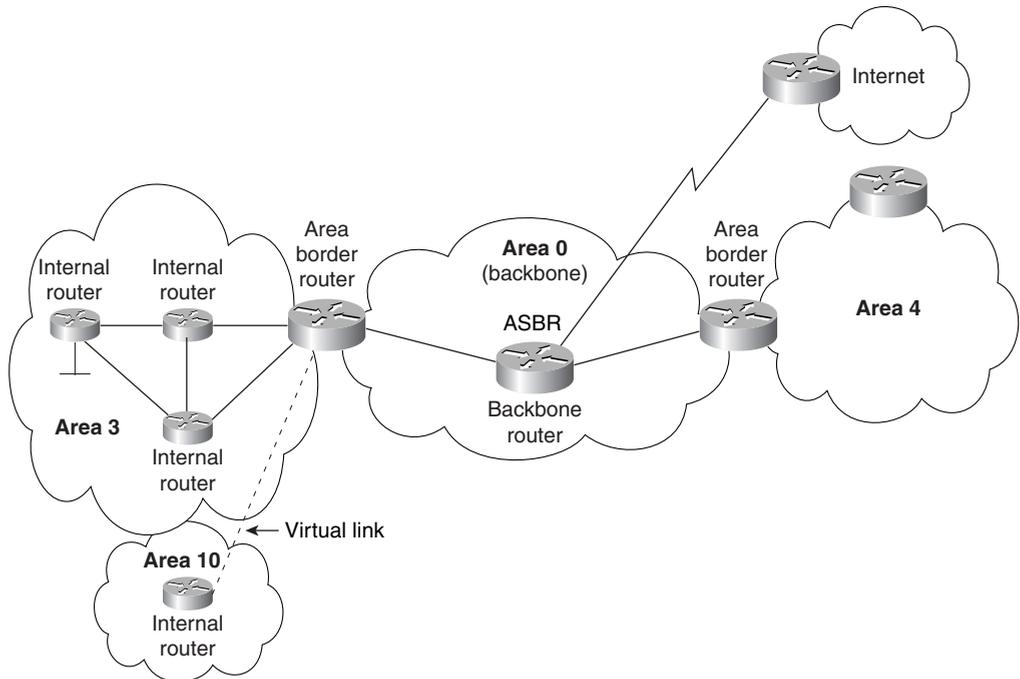
- There is no physical connection to Area 0. This may be because the organization has recently merged with another or because of a network failure.
- There are two Area 0s because of a network merger. These Area 0s are connected by another area (for example, Area 5).
- The area is critical to the company, and an extra link has been configured for redundancy.

Although this is an extremely powerful command, it is not recommended as part of the design strategy for your network; instead, it is a temporary solution to a connectivity problem. It is necessary to ensure that the following is observed in creating a virtual link:

- Both routers must share a common area.
- One of the routers must be connected to Area 0.

Figure 6-5 illustrates the use of a virtual link to provide a router in Area 10 connectivity to the backbone in Area 0.

Figure 6-5 Virtual Links in a Multiple-Area OSPF Network



## OSPF over an NBMA Network

Another design consideration is the design of the NBMA network as part of the OSPF domain. There are two main ways to approach the inclusion of an NBMA network:

- The NBMA network can be defined as Area 0. The reasoning is that if the NBMA is used to connect all remote sites, all traffic will have to traverse this network. If the remote sites are made satellite areas, then all traffic will have to traverse the NBMA, so it makes sense to make it the backbone area. This works well in a full-mesh environment, although it will result in a large number of LSAs being flooded into the WAN and puts extra demands on the routers connecting to the NBMA network.
- In a hub-and-spoke NBMA network, it makes sense to assign the hub network as Area 0 with the other remote sites and the NBMA network as other areas. This is a good design if the satellite areas are stub areas because it means that the routing information—and, thus, network overhead—is kept to a minimum over the NBMA cloud. Depending on the design, the rest of the network may constitute one other area or multiple areas. This will depend on the size and growth expectations of the OSPF domain.

After the design of the network is in place, it is time to configure the routers. The configuration of a basic OSPF network is demonstrated in Chapter 5, “Using OSPF in a Single Area”.

## Configuring OSPF on a Multiarea Network

Some of these commands were dealt with in Chapter 5, in the section “Configuring OSPF in a Single Area.” Commands that have been covered already are reviewed briefly here, and the additional parameters for configuration in a multiarea environment are explained in detail.

### Configuration Commands for a Multiarea OSPF Network

The following commands are necessary to configure a multiarea OSPF network:

- The OSPF **network** command
- The OSPF **area range** command for an ABR
- The OSPF **summary-address** command for an ASBR
- The OSPF command for a stub area
- The OSPF command for a totally stubby area
- The OSPF command for the cost of a default route propagated into the area
- The command for configuring a virtual link

#### The **network** Command

The first command to consider is one that was dealt with in Chapter 5. The **network** command was explained in terms of identifying the interfaces that participated in the OSPF routing process. The command will now be used to identify not only the interfaces that are sending and receiving OSPF updates, but also the area in which they reside. This configuration is used on an ABR.

The following is the syntax for the OSPF **network** command:

```
network network-number wildcard-mask area area-number
```

---

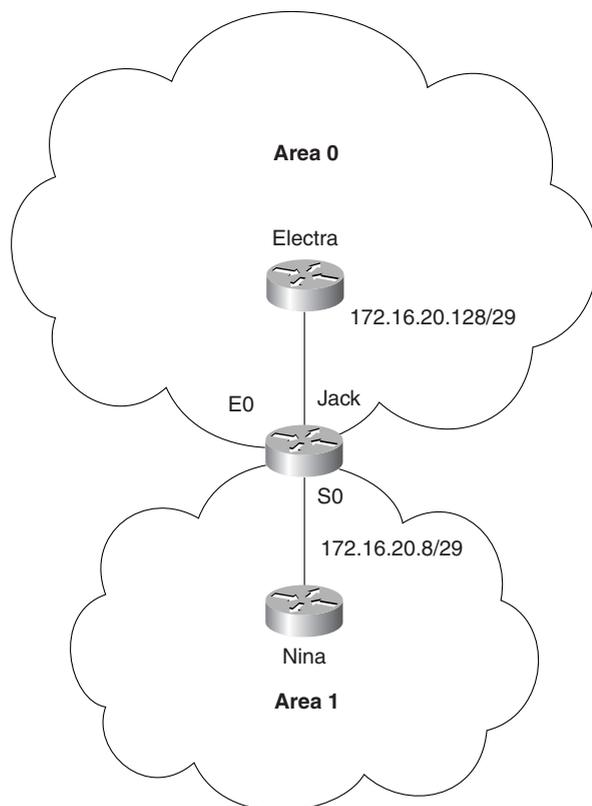
#### NOTE

The area requested in the preceding syntax is the area in which the interface or interfaces configured with the network address reside.

---

Care must be taken now in the use of the wildcard mask. In a single-area configuration, all the interfaces are in the same area. The network commands just identify the network numbers in use. Therefore, they may be configured to the Internet number, as they are in IGRP and RIP. The only reason to be more specific would be to exclude some interfaces from the OSPF domain.

Figure 6-6 illustrates the example configuration that follows.

Figure 6-6 The `network` Command

Example 6-1 shows two interfaces, each with a subnet of the same major network where the interfaces are placed into different areas. The network number has been subnetted into the last octet so that you can truly see the power and granularity of the wildcard mask at work.

Example 6-1 The `network` Command for Router Jack

```
network 172.16.20.128 0.0.0.7 area 0
network 172.16.20.8 0.0.0.7 area 1
```

The need now to identify areas on a per interface basis brings into use the other part of the command. Although the command itself is very simple, it adds complexity to the use of the mask. It is to be remembered that the **network** command follows the rule of a linked list. The order of the statements is important: The most specific should be stated first because the OSPF process will act on the first match that is found.

## The OSPF **area range** Command for an ABR

The **area range** command is configured on an ABR because it dictates the networks that will be advertised out of the area.

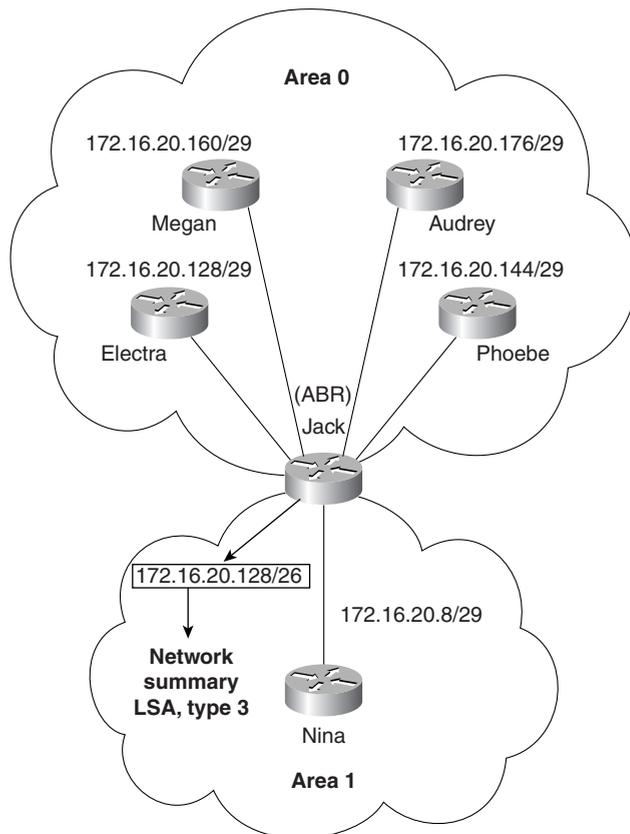
Use the **area** router configuration command with the **range** keyword to consolidate and summarize routes at an area boundary. Use the **no** form of this command to disable this function for the specified area:

```
area area-id range address mask  
no area area-id range address mask
```

In the preceding syntax, *area-id* is the identifier (ID) of the area about which routes are to be summarized. It can be specified as either a decimal value or an IP address. Here, *address* is the IP address, and *mask* is the IP mask.

Figure 6-7 illustrates the example configuration that follows.

**Figure 6-7** The OSPF **area range** Command for an ABR



Example 6-2 shows the configuration required to summarize the following five individual subnets (which can address six hosts each) into one subnet. This summarized subnet may then be propagated across the OSPF network, saving both bandwidth and CPU:

- 172.16.20.128 /29
- 172.16.20.144 /29
- 172.16.20.160 /29
- 172.16.20.176 /29

These subnets are summarized into one subnet:

- 172.16.20.128 /26

This one subnet will then be propagated into Area 1.

**Example 6-2** *The OSPF area range Command for an ABR*

```
Router Jack
router ospf 100
network 172.16.20.128 0.0.0.7 area 0
network 172.16.20.8 0.0.0.7 area 1
area 0 range 172.16.20.128 255.255.255.192
```

**NOTE**

The area ID requested is the area that the subnets originated from. It is not the destination area. The summarization update populates the topological databases of the routers in the destination area. These routers will need to know the source area for the summarized subnet to know where to send the data traffic.

## The OSPF summarization Command for an ASBR

The **summary-address** command is used on the ASBR to summarize the networks to be advertised to the outside world.

The syntax for the OSPF **summary-address** command for an ASBR is as follows:

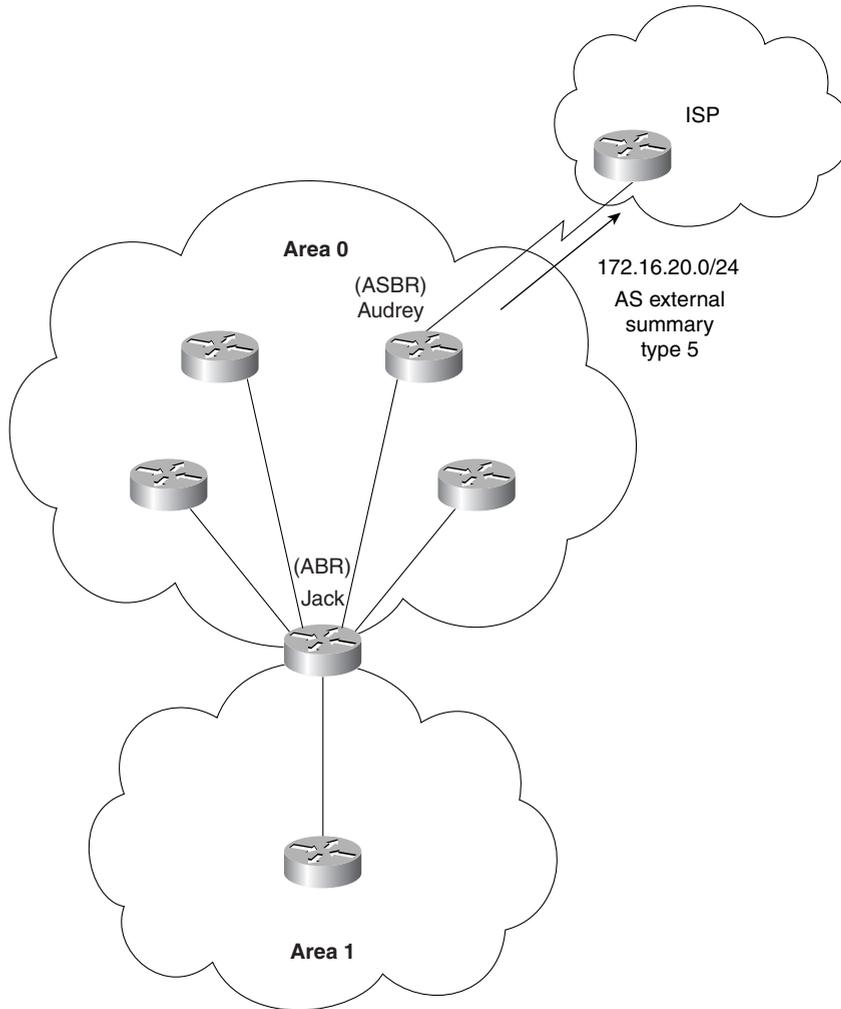
```
summary-address address mask [not advertise][tag tag]
```

In the preceding syntax, *address* is the summary address designated for a range of addresses, and *mask* is the IP subnet mask used for the summary route.

The design and implementation of the addressing scheme are crucial to the success of the OSPF network and cannot be stressed too strongly. Refer to Chapter 3 for details on IP addressing and summarization.

Figure 6-8 illustrates the example configuration that follows.

**Figure 6-8** The *OSPF summary-address* Command for an ASBR



Example 6-3 shows the summarization of the network address 172.16.20.0, which will be propagated into the ISP's autonomous system.

**Example 6-3** *The OSPF summary-address Command for an ASBR*

```
Router Audrey
router ospf 100
network 172.16.20.176 0.0.0.7 area 0
summary-address 172.16.20.0 255.255.255.0
```

## The OSPF Command for a Stub Area

After designing the addressing scheme for the network, it should be clear which areas, if any, are suitable candidates for configuration as a stub, totally stubby, or not so stubby areas.

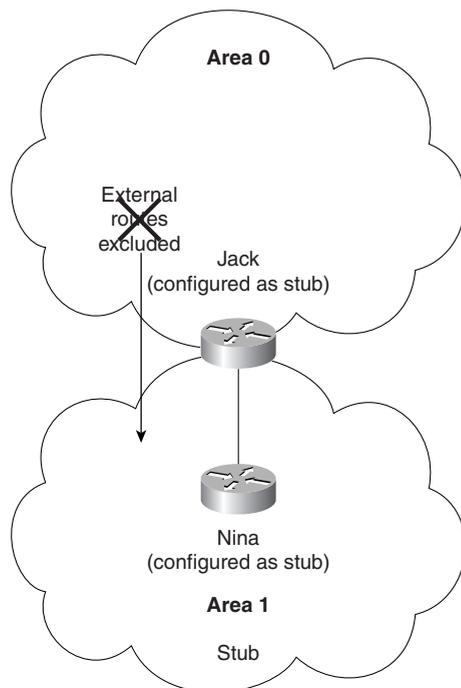
**NOTE**

In this age of jargon and complex language, it is refreshing that the industry sense of humor allows such descriptive yet slightly ludicrous terms to have official status.

The syntax for the OSPF router command for a stub area is as follows:

```
area area-id stub
```

Figure 6-9 illustrates the example configuration that follows.

**Figure 6-9** *The Configuration of a Stub Area*

Example 6-4 shows the creation of a stub area. Note that both the ABR and the internal router share the stub area configuration.

**Example 6-4** *The Configuration of a Stub Area*

```
Router Nina
ospf 100
network 0.0.0.0 255.255.255.255 area 1
area 1 stub
Router Jack
router ospf 100
network 172.16.20.128 0.0.0.7 area 0
network 172.16.20.8 0.0.0.7 area 1
area 0 range 172.16.20.128 255.255.255.192
area 1 stub
```

---

**NOTE**

All OSPF routers inside a stub area must be configured as stub routers. This is because whenever an area is configured as a stub, all interfaces that belong to that area will start exchanging hello packets with a flag that indicates that the interface is a stub. Actually, this is just a bit in the hello packet (E bit) that gets set to 0. All routers that have a common segment must agree on that flag. If they don't, they will not become neighbors, and routing will not take effect.

---

## The OSPF Command for a Totally Stubby Area

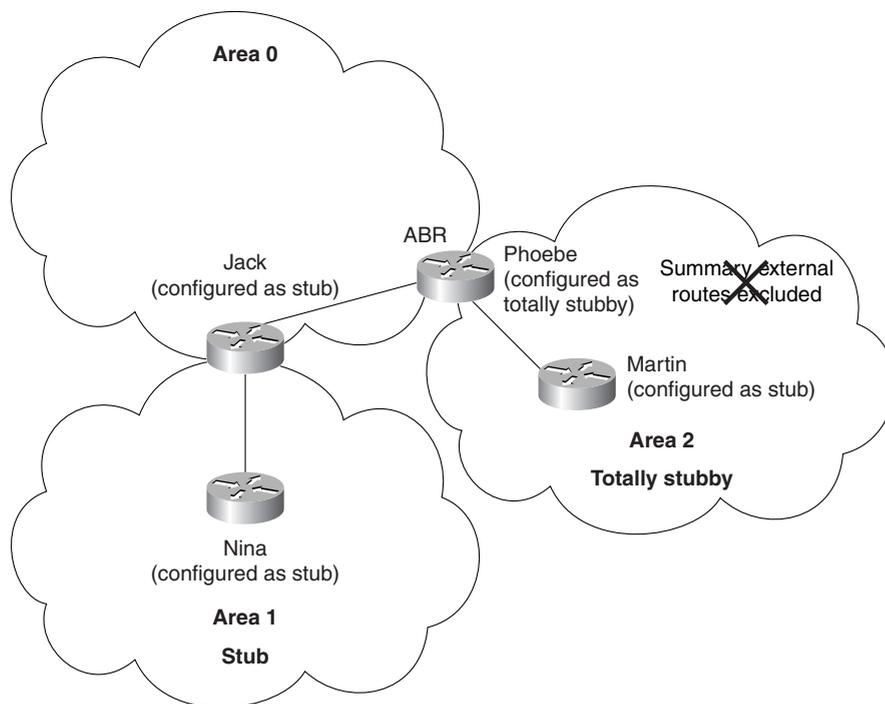
The syntax for the OSPF command for a totally stubby area is as follows:

```
area area-id stub no-summary
```

This addition of the **no-summary** parameter informs the ABR not to send summary updates from other areas into the area. This command needs to be configured only on the ABR because it is the only router with this responsibility. This command is configurable only on a Cisco router because it is a proprietary feature. All the other routers are configured as stub-area internal routers.

Figure 6-10 illustrates the example configuration of a totally stubby area that follows.

**Figure 6-10** *The Configuration of a Totally Stubby Area*



Example 6-5 shows the configuration of a totally stubby area.

**Example 6-5** *The Configuration of a Totally Stubby Area*

```
Router Phoebe
router ospf 100
network 172.16.20.144 0.0.0.7 area 0
network 172.16.20.16 0.0.0.7 area 2
area 2 stub no-summary
area 0 range 172.16.20.128 255.255.255.192
Router Martin
router ospf 100
network 0.0.0.0 255.255.255.255 area 2
area 2 stub
```

As a totally stubby area, no summary or external routes are propagated by the ABR into the area. To reach networks and hosts outside their area, a workstation must send to a default route, which the ABR advertises into the area.

## The OSPF Command for the Cost of a Default Route Propagated into the Area

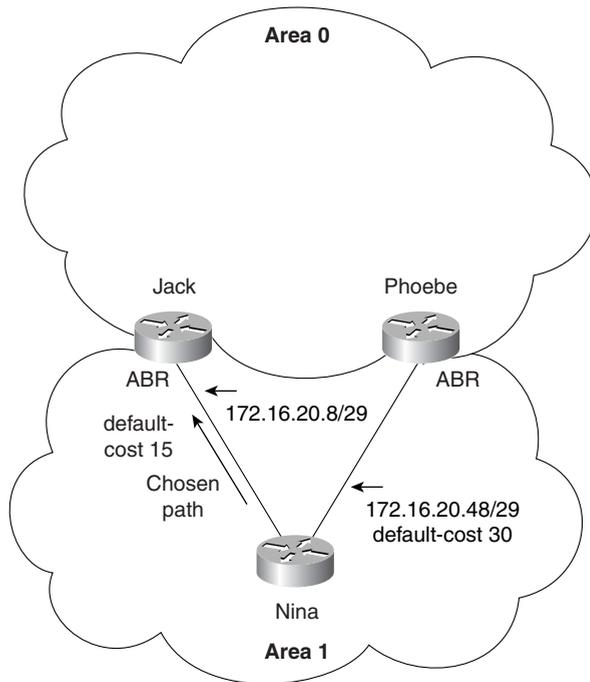
To define the cost to the default route, the following command is used. If the cost is not specified, the path will be calculated as the internal area cost plus 1.

The syntax for the OSPF command for the cost of a default route propagated into the area is as follows:

```
area area-id default-cost cost
```

The ABR attached to the stub area automatically generates and advertises a default route with a destination of 0.0.0.0 into the stub area. Figure 6-11 illustrates the example configuration that follows.

**Figure 6-11** *The OSPF Command for the Default Route Propagated into the Area*



Example 6-6 shows how the default cost can be set in the configuration. Setting a cost on the default route is useful when the stub area has more than one ABR: It allows the ABR used to exit the area to be determined.

**Example 6-6** *The OSPF Command for the Default Route Propagated into the Area*

```
Router Nina
router ospf 100
network 0.0.0.0 255.255.255.255 area 1
area 1 stub
Router Jack
router ospf 100
network 172.16.20.128 0.0.0.7 area 0
network 172.16.20.8 0.0.0.7 area 1
area 0 range 172.16.20.128 255.255.255.192
area 1 stub
area 1 default-cost 15
Router Phoebe
router ospf 100
network 172.16.20.144 0.0.0.7 area 0
network 172.16.20.48.0 0.0.0.7 area 1
area 1 stub
area 1 default-cost 30
area 0 range 172.16.20.128 255.255.255.192
```

---

**NOTE** The **area default-cost** command needs to be configured only on the ABR because it is the only router with this responsibility. However, Example 6-6 has shown the configuration on both routers to illustrate the choice. The second ABR, Phoebe, will only be used if Jack fails. If there were no configuration on Jack, it would still be used by all internal routers as the ABR because the default cost is 1.

---

## Configuring a Virtual Link

When it is not possible to connect an area to Area 0 directly, a solution is to create an IP tunnel called a virtual link. This is remarkably easy to configure. As with many things in OSPF, of course, this ease of configuration belies the complexity of the technology being used. Many things can go wrong. The most common problem is in the address of the other end of the virtual link. The command is given between ABRs, at least one of which must be in Area 0. The command, issued at both ABRs, states the transit area and the router ID of the remote destination ABR. This creates essentially a tunnel through the transit area, which, although it may involve many routers to forward the traffic, appears to the remote ABRs as next hops.

---

**NOTE** The **area virtual-link** command is potentially included in the BSCI exam and, for that reason, is worth mentioning. In practice, virtual links are a design nightmare and are best avoided. They are useful when mending a network on a temporary basis while awaiting a moment's peace to rectify the design of the network.

---

The syntax to configure a virtual link is as follows:

```
area area-id virtual-link router-id
```

Here, *area-id* is the ID assigned to the transit area for the virtual link.

In addition, *router-id* is the router ID of the virtual link neighbor.

Seeing how these commands work in context makes their use and functionality much more apparent.

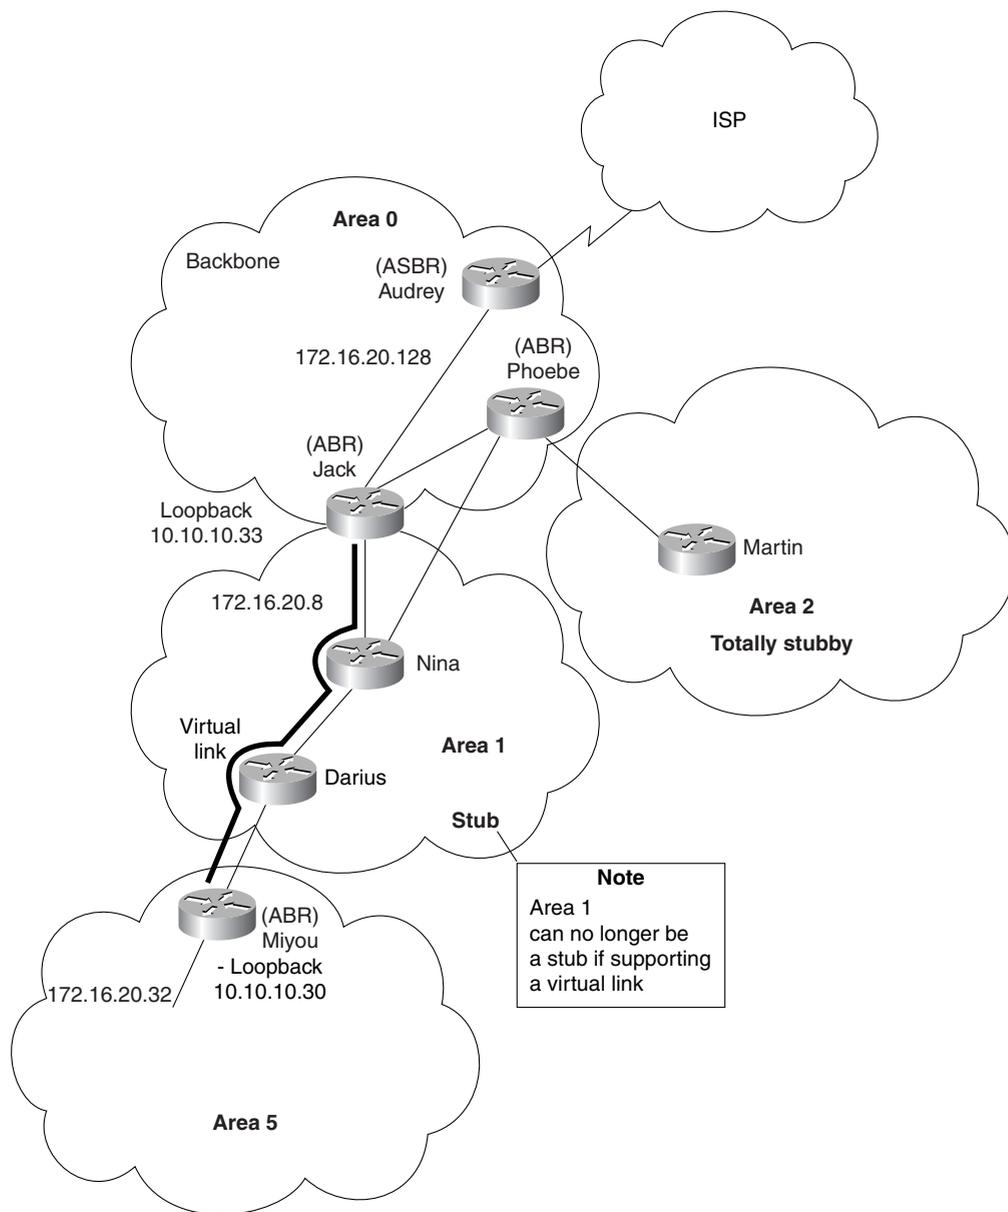
Figure 6-12 illustrates the example configuration that follows.

Example 6-7 shows the setting of the loopback interfaces that provide the router ID. It then shows the configuration of the virtual link through the network.

**Example 6-7** *Configuring a Virtual Link*

```
Router Jack
interface loopback 0
ip address 10.10.10.33 255.255.255.255
router ospf 100
network 172.16.20.128 0.0.0.7 area 0
network 172.16.20.8 0.0.0.7 area 1
area 0 range 172.16.20.128 255.255.255.192
area 1 default-cost 15
area 1 virtual-link 10.10.10.30
Router Miyou
loopback interface 0
ip address 10.10.10.30 255.255.255.255
ospf 100
network 172.16.20.32 0.0.0.7 area 5
network 172.16.20.64 0.0.0.7 area 1
area 1 virtual-link 10.10.10.33
```

Figure 6-12 Configuring a Virtual Link

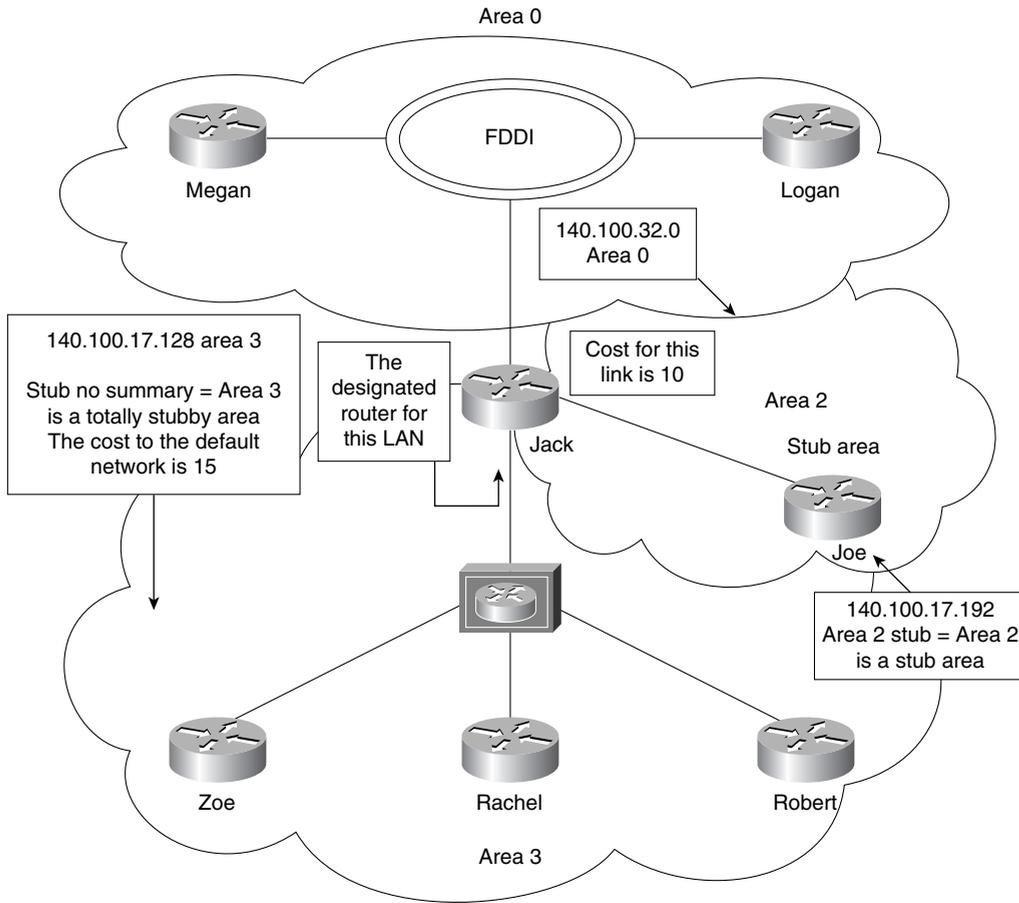


## A Working Configuration of OSPF on a Multiarea Network

Example 6-8 is a working configuration tested for verification. It includes many of the commands explained earlier in this chapter. This is so that you see an entire working configuration rather than the relevant segment for configuring a particular networking nuance.

The configuration should be used in conjunction with Figure 6-13.

Figure 6-13 Diagram of the Example 6-8 Network



**Example 6-8** *Configuring OSPF in a Multiarea Network on Router Jack*

```
Router(config)#router ospf 100
Router(config-router)#network 140.100.17.128 0.0.0.15 area 3
Router(config-router)#network 140.100.17.192 0.0.0.15 area 2
Router(config-router)#network 140.100.32.0 0.0.0.255 area 0
Router(config-router)#area 2 stub
Router(config-router)#area 3 stub no-summary
Router(config-router)#area 3 default-cost 15
!
Router(config-router)#interface FastEthernet0
Router(config-if)#ip address 140.100.17.129 255.255.255.240
Router(config-if)#no ip directed-broadcast
Router(config-if)#ip ospf priority 100
!
Router(config-if)#interface FastEthernet1
Router(config-if)#ip address 140.100.17.193 255.255.255.240
Router(config-if)#no ip directed-broadcast
Router(config-if)#ip ospf cost 10
!
Router(config-if)#interface Fddi0
Router(config-if)#ip address 140.100.32.10 255.255.255.240
Router(config-if)#no ip directed-broadcast
Router(config-if)#no keepalive

Router(config-if)#exit
```

Of course, it is imperative to check any configuration on a network device because any errors could potentially bring down the entire network. To verify the configuration, a wealth of Cisco commands are available. They are covered in the following section.

## Checking the Configuration of OSPF on a Multiarea Network

The **show** commands shown here are in addition to the commands described in Chapter 5, in the section, “Checking the Configuration of OSPF on a Single Router.” The single router commands are also extremely useful in a multiarea configuration. They are all invaluable in both the configuration and maintenance of a live network. They are particularly useful in troubleshooting the network. The following are the additional commands that you can use in conjunction with single router commands when verifying OSPF operation on a multiarea network:

- The **show ip ospf border-routers** command
- The **show ip ospf virtual-links** command

The capability to analyze the output of a **show** command demonstrates more than rote learning; it also demonstrates an understanding of the concepts that make up the foundations of OSPF design and configuration.

The commands explained in this book constitute a small subset of the commands available in OSPF. Because the OSPF command set is very comprehensive, the capability to monitor the network and thereby maintain and troubleshoot it requires advanced OSPF knowledge.

## The show ip ospf border-routers Command

```
show ip ospf border-routers
```

This command shows the OSPF ABRs and ASBRs for which the internal router has entries in its routing table. This command is excellent for troubleshooting configuration errors and understanding how the network is communicating about its routes.

Example 6-9 shows the output of this command.

### Example 6-9 show ip ospf border-routers Output

Router# show ip ospf border-routers								
OSPF Process 100 internal Routing Table								
Destination	Next Hop	Cost	Type	Rte Type	Area	SPF	No	
160.89.97.53	144.144.1.53	10	ABR	INTRA	0.0.0.3		3	
160.89.103.51	160.89.96.51	10	ABR	INTRA	0.0.0.3		3	
160.89.103.52	160.89.96.51	20	ASBR	INTER	0.0.0.3		3	
160.89.103.52	144.144.1.53	22	ASBR	INTER	0.0.0.3		3	

Table 6-3 explains the meaning of the important fields in the output of the **show ip ospf border-routers** command.

**Table 6-3** Explanation of the show ip ospf border-routers Command Output

Field	Explanation
OSPF Process 100 Internal Routing Table	This is the OSPF routing process ID for the router.
Destination	This is the router ID of the destination router, whether an ABR or an ASBR.
Next Hop	If the ABR or ASBR is not directly connected, this is the address of the next logical hop in the chosen path to the ABR or ASBR.
Cost	This is the metric or cost of taking this path to the destination.
Type	This states whether the destination router is an ABR or ASBR or both.

**Table 6-3** *Explanation of the show ip ospf border-routers Command Output (Continued)*

Field	Explanation
Rte Type	The is the type of this route; it is either an intra-area or interarea route.
Area	This is the area ID of the area that this route is learned from.
SPF No	This is the SPF calculation number that installed this route into the routing table.

The **show ip ospf border-routers** command is useful to verify that the configuration has worked and that the OSPF network is functioning correctly. In a multiarea network, **show ip ospf border-routers** command can immediately indicate why users cannot connect outside their area.

It is helpful to extract this information from what could be a long routing table, within which this information is scattered.

## The show ip ospf virtual-links Command

This command shows the configured virtual links that are in existence.

```
show ip ospf virtual-links
```

Another command to use in conjunction with this is **show ip ospf neighbors**.

Example 6-10 shows the output of the **show ip ospf virtual-links** command.

### Example 6-10 show ip ospf virtual-links Output

```
Router# show ip ospf virtual-links
Virtual Link to router 140.100.32.10 is up
Transit area 0.0.0.1, via interface Ethernet0, Cost of using 10
Transmit Delay is 1 sec, State DROTHER
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 0:00:08
Adjacency State FULL
```

Table 6-4 explains the meaning of the important fields in the output of the **show ip ospf virtual-links** command.

**Table 6-4** *Explanation of the show ip ospf virtual-links Command*

Field	Explanation
Virtual Link to router 140.100.32.10 is up	This shows the router ID of the other end of the virtual link, which is seen as a neighbor.
Transit area 0.0.0.1	This is the area through which the virtual link is tunneled: area 0.0.0.1 or area 1 in this case.
via interface Ethernet0	This is the outgoing interface on the router that connects the virtual link to Area 0.
Cost of using 10	This is the cost of reaching the OSPF neighbor through the virtual link.
Transmit Delay is 1 sec	This is the delay of the link, how long it will take to transmit an LSA. This value must be less than the retransmit timer setting.
State DROTHER	This gives the state of the OSPF neighbor.
Hello 10	This gives the timed update interval for the Hello protocol, in seconds. The default is 10 seconds.
Dead 40	This tells how long the router will wait without hearing a hello from the neighbor before it declares the neighbor dead. The default is 40 seconds.
Retransmit 5	The retransmit interval is the time in seconds that the router will wait without hearing an acknowledgment for the LSA that it has sent to its neighbor. The default is 5 seconds.
Hello due in 0:00:08	This shows the amount of time before the next Hello is expected from the neighbor.
Adjacency State FULL	This specifies the state of the neighbor adjacency. The two routers have fully synchronized their topological databases.

## Conclusion

OSPF has the advantage of being an industry standard, ensuring international interoperability as long as everyone conforms to the standard. It is a very stable protocol, having been available for many years. At the heart of OSPF is the concept of areas that allow summarization and the capability to contain the breadth of LSA flooding and thus knowledge of the network and SPF computation. This capability to impose a hierarchy upon a network was revolutionary at the time that OSPF was conceptualized. OSPF is still a hot contender to EIGRP because it is the only viable option of providing the open forum with scalable growth for the network. As can be seen by this chapter, although many solutions may be implemented, the health of the network still depends on well-conceived designs based on a careful analysis of the network.

## Foundation Summary

The “Foundation Summary” is a collection of quick reference information that provides a convenient review of many key concepts in this chapter. For those of you who already feel comfortable with the topics in this chapter, this summary will help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final preparations before the exam, these tables and figures will be a convenient way to review the day before the exam.

## OSPF Routers

The following list identifies the different OSPF routers:

- **Internal router**—Within an area, the functionality of the router is straightforward. It is responsible for maintaining a current and accurate database of every subnet within the area. It is also responsible for forwarding data to other networks by the shortest path. Flooding of routing updates is confined to the area.
- **Backbone router**—The design rules for OSPF require that all the areas be connected through a single area known as the *backbone area* or *Area 0*. A router within this area is referred to as a *backbone router*. It may also be an internal router or an Area Border Router.
- **Area Border Router (ABR)**—This router is responsible for connecting two or more areas. It holds a full topological database for each area to which it is connected and sends LSA updates between the areas. These LSA updates are summary updates of the subnets within an area. It is at the area border that summarization should be configured for OSPF because this is where the LSAs make use of the reduced routing updates to minimize the routing overhead on both the network and the routers.
- **Autonomous System Boundary Router (ASBR)**—To connect to the outside world, or to any other routing protocol, you need to leave the OSPF domain. OSPF is an *interior routing protocol* or *Interior Gateway Protocol (IGP)*; *gateway* is an older term for a router. The router configured for this duty is the ASBR. Although you can place this router anywhere in the OSPF hierarchical design, it should reside in the backbone area. Because any traffic leaving the OSPF domain will likely also leave the router’s area, it makes sense to place the ASBR in a central location that all traffic leaving its area must traverse.

## Link-State Advertisements

The five link-state advertisements are as follows:

- **The router link**—This LSA states all the links to the router sending out the LSA. The list is of all the neighbors attached to the router. The LSA is flooded to the area.
- **The network link**—This LSA is sent out by the designated router and lists all the routers on the segment for which it is the designated router and has a neighbor relationship. The LSA is flooded to the whole area.
- **The network summary link**—This LSA is sent between areas and summarizes the IP networks from one area to another. It is generated by an ABR.
- **The AS external (ASBR) summary link**—This LSA is sent to a router that connects to the outside world (ASBR). It is sent from the ABR to the ASBR. The LSA contains the metric cost from the ABR to the ASBR.
- **The external link**—This LSA is originated by AS boundary routers and flooded throughout the AS. Each external advertisement describes a route to a destination in another autonomous system. Default routes for the AS can also be described by AS external advertisements.

## Routing Table Codes

Table 6-5 shows the codes used in the routing table.

**Table 6-5** *OSPF Routing Table Codes And Associated LSAs*

LSA Type	Routing Table Entry	Description
1 Router Link	O	This is generated by the router, listing all the links to which it is connected, their status, their and cost. It is propagated within the area.
2 Network Link	O	This is generated by the designated router on a multiaccess LAN to the area.
3 or 4 Summary Link (between areas)	OIA	LSA Type 3 includes the networks or subnets within an area that may have been summarized and that are sent into the backbone and between ABRs. LSA Type 4 is information sent to the ASBR from the ABR. These routes are not sent into totally stubby areas.

**Table 6-5** *OSPF Routing Table Codes And Associated LSAs (Continued)*

<b>LSA Type</b>	<b>Routing Table Entry</b>	<b>Description</b>
5 Summary Link/External Link (between autonomous systems)	0E1 or 0E2	The routes in this LSA are external to the autonomous system. They can be configured to have one of two values. E1 will include the internal cost to the ASBR added to the external cost reported by the ASBR. E2 does not compute the internal cost—it just reports the external cost to the remote destination.

## Command Summaries

This section contains a list of the commands explained in this chapter. This list is not intended to teach the use of the commands, but to remind you of the options available.

**Table 6-6** *OSPF Command Summary*

<b>Command</b>	<b>Description</b>
<b>network</b> <i>network-number wildcard-mask area area-id</i>	This command identifies the interfaces that are running OSPF and places them in the appropriate area.
<b>summary-address</b> <i>address mask [not-advertise][tag tag]</i>	This command consolidates routes into a summary route before injecting them into the external world. Remember that the mask is the subnet, not the wildcard mask. The options to not advertise will suppress routes that match the mask. The <b>tag</b> option is for use with redistribution.
<b>area</b> <i>area-id range address mask</i>	This command is used to summarize the routes at the ABR before injecting them into another area.
<b>no area</b> <i>area-id range address mask</i>	This command disables the configured summarization at the ABR.
<b>area</b> <i>area-id stub</i>	This command turns the area into a stub area. This command must be configured on every router in the area.
<b>area</b> <i>area-id stub no-summary</i>	This command turns an area into a totally stubby area.
<b>area</b> <i>area-id default-cost cost</i>	This configures the cost for the default summary route used for a stub or totally stub area. The default is 1.
<b>area</b> <i>area-id virtual-link router-id</i>	This command creates a virtual link.
<b>show ip ospf border-routers</b>	This command lists the ABR routers in the AS.
<b>show ip ospf virtual-links</b>	This command shows the virtual links and the current parameters.

## Q&A

The following questions test your understanding of the topics covered in this chapter. The final questions in this section repeat the opening “Do I Know This Already?” questions. These are repeated to enable you to test your progress. After you have answered the questions, find the answers in Appendix A. If you get an answer wrong, review the answer and ensure that you understand the reason for your mistake. If you are confused by the answer, refer to the appropriate text in the chapter to review the concepts.

- 1 In a totally stubby area, which routes are not propagated into the area?

---

---

---

- 2 Where would you see whether a learned network was within the same area as the router you were looking at?

---

---

---

- 3 An Area Border Router must be resident in which area?

---

---

---

- 4 What does *ABR* stand for, and what LSAs will it forward?

---

---

---

- 5 State two advantages in creating areas in OSPF.

---

---

---

**6** What is an external route, and on which type of router will this route be introduced?

---

---

---

**7** Which command in OSPF shows the network LSA information?

---

---

---

**8** Why is the use of VLSM important in the design of OSPF?

---

---

---

**9** Given the networks 144.12.8.0 and 144.12.12.0, both with the mask of 255.255.255.0, is it possible to summarize these into one subnet? If so, state the new subnet and mask.

---

---

---

**10** What is a discontinuous network?

---

---

---

**11** A virtual link in OSPF is used to solve what problem?

---

---

---

**12** State one disadvantage for making an NBMA Frame Relay cloud Area 0.

---

---

---

- 13** State one advantage in making the centralized routers and network resources dwell in Area 0 while the Frame Relay cloud and the stub remote LANs reside in satellite stub areas.

---

---

---

- 14** How does creating a number of areas in OSPF reduce the number of SPF calculations?

---

---

---

- 15** How does a stub area differ from the backbone area?

---

---

---

- 16** How does a totally stubby area differ from a stub area?

---

---

---

- 17** State the different LSA types.

---

---

---

- 18** Where does the backbone router reside, and what is its function?

---

---

---

- 19 Are there any considerations for OSPF configured with VLSM sending routing updates into RIPv1?

---

---

---

- 20 There are two types of summarization. What are they?

---

---

---

- 21 Can the following subnets with a mask of 255.255.255.0 be summarized? If so, state the subnet and mask that can be used.

19.44.16.0	19.44.24.0
19.44.17.0	19.44.25.0
19.44.18.0	19.44.26.0
19.44.19.0	19.44.27.0
19.44.20.0	19.44.28.0
19.44.21.0	19.44.29.0
19.44.22.0	19.44.30.0
19.44.23.0	19.44.31.0

---

---

---

- 22 Why can interarea summarization be configured only on ABRs?

---

---

---

- 23 What command would be used to create a totally stubby area?

---

---

---

**24** What is a virtual link, and what command would be used to create it?

---

---

---

**25** Where would you issue the command to summarize IP subnets? State the command that would be used.

---

---

---

**26** How would you summarize external routes before injecting them into the OSPF domain?

---

---

---

## Scenarios

The following scenarios and questions are designed to draw together the content of the chapter and to exercise your understanding of the concepts. There is not necessarily a right answer. The thought process and practice in manipulating the concepts is the goal of this section. The answers to the scenario questions are found at the end of this chapter.

### Scenario 6-1

Refer to Figure 6-14, and design the addressing scheme for the network. After this is designed, write the configuration for the central router.

Ensure that you include the following:

- 1 Address the network using the private network 10.0.0.0. Design the addressing scheme so that it allows for the summarization of addresses between areas. Show the summarization that you allocate, and explain your reasons for your choices.

Area 0 is using a prefix of 28 bits within the area.

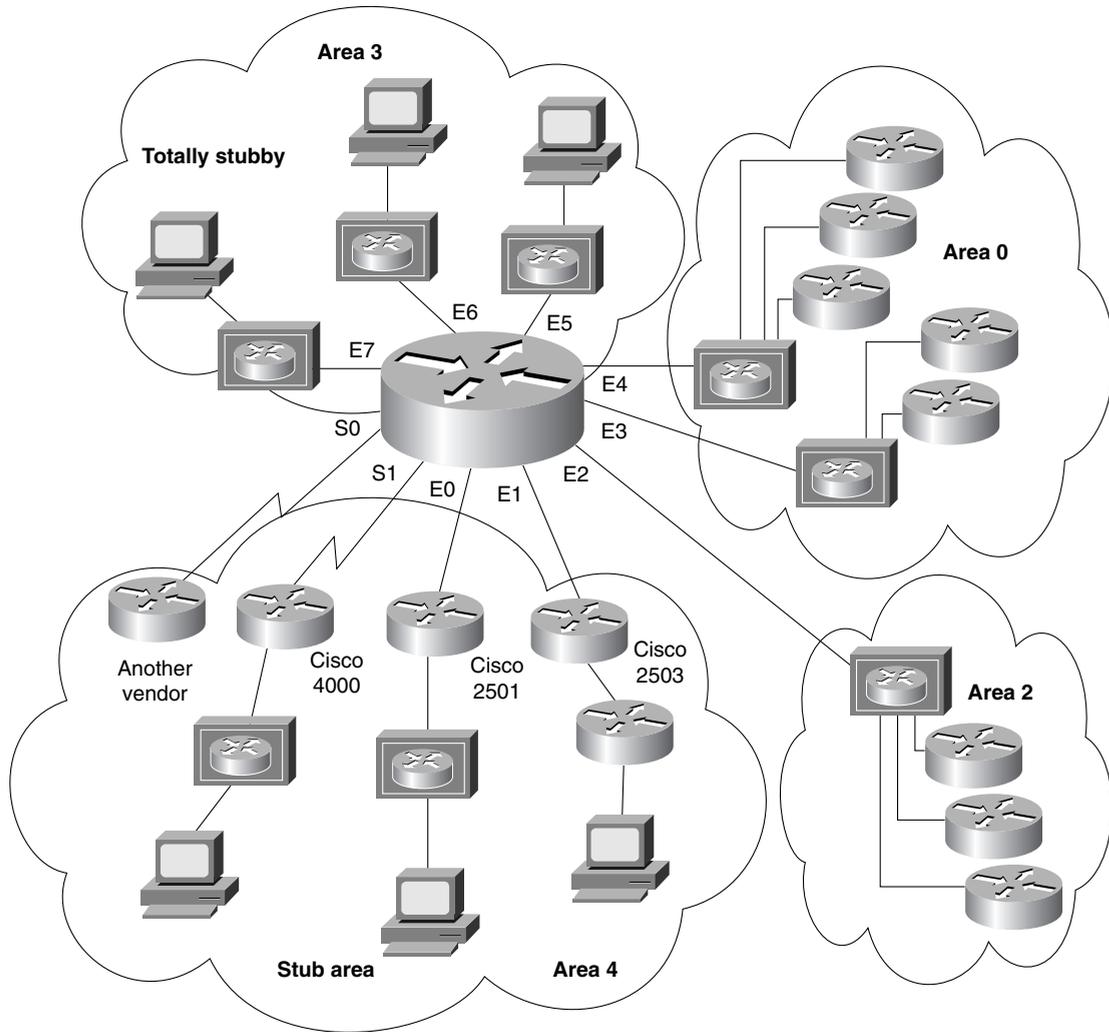
Area 2 is using a prefix of 22 bits within the area.

Area 3 is using a prefix of 24 bits within the area.

Area 4 is using a prefix of 30 bits for the serial connections. It is using a 28-bit prefix for the connections to the Ethernet routers. Do not include the subnets attached to the LANs in Area 4.

- 2 Issue the commands for the main router Figure 6-14 to configure the following:
  - The router ID
  - The network commands to place the appropriate interfaces into the correct areas
  - The configuration of the totally stubby area (Area 3)
  - The configuration of the stub (Area 4)
  - Summarization between areas
  - The election of the central router as designated router, where appropriate

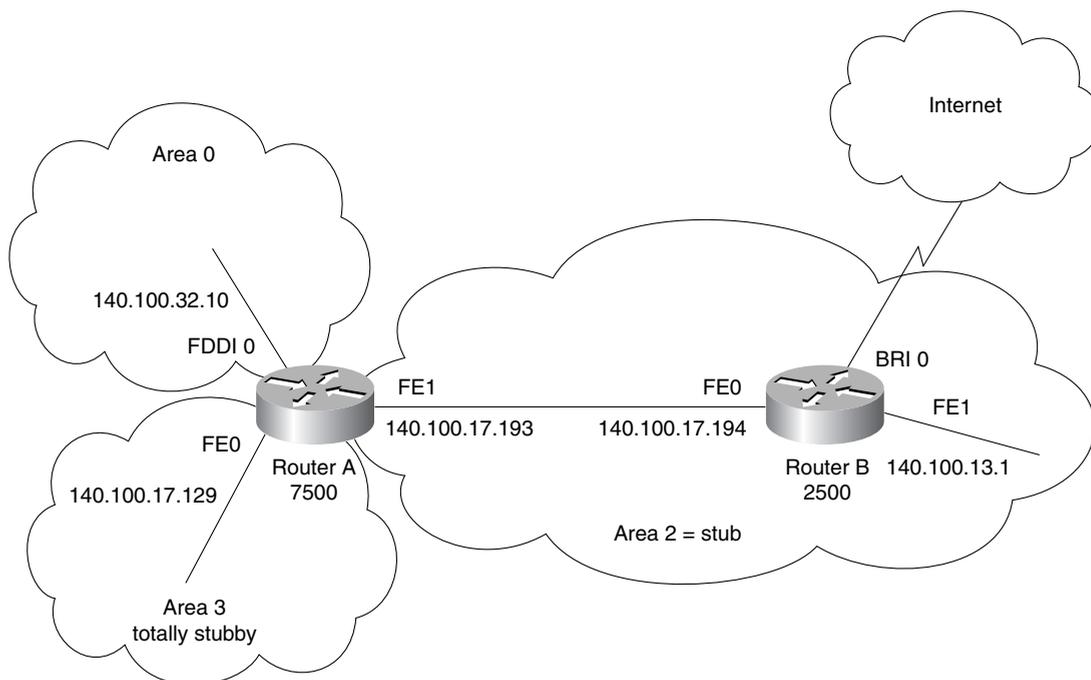
Figure 6-14 The Diagram for Configuration Scenario 6-1



## Scenario 6-2

Using Figure 6-15, answer the questions that follow.

**Figure 6-15** *The Diagram for Configuration Scenario 6-2*



The users in the network are complaining about the slowness of the network, particularly when trying to access the Internet. Examine the configuration in Example 6-11 in conjunction with Figure 6-15, and give reasons for any slowness or lack of connectivity that you can see on the network. Provide current configuration commands to correct any errors that you find.

**Example 6-11** *Configuring OSPF Example for Scenario 6-2*

```
ROUTER A
Router(config)#router ospf 100
Router(config-router)#network 140.100.17.128 0.0.0.15 area 3
Router(config-router)#network 140.100.17.192 0.0.0.15 area 2
Router(config-router)#network 140.100.32.0 0.0.0.0 area 0
Router(config-router)#area 2 stub
Router(config-router)#area 3 stub no-summary
Router(config-router)#area 3 default-cost 15
```

*continues*

**Example 6-11** *Configuring OSPF Example for Scenario 6-2 (Continued)*

```

!
Router(config-router)#interface FastEthernet0
Router(config-if)#ip address 140.100.17.129 255.255.255.240
Router(config-if)#no ip directed-broadcast
!
Router(config-if)#interface FastEthernet1
Router(config-if)#ip address 140.100.17.193 255.255.255.240
Router(config-if)#no ip directed-broadcast
!
Router(config-if)#interface Fddi0
Router(config-if)#ip address 140.100.32.10 255.255.255.240
Router(config-if)#no ip directed-broadcast
Router(config-if)#no keepalive

Router(config-if)#exit
ROUTER B
Router(config)#router ospf 100
Router(config-router)# network 140.100.0.0 0.0.0.15 area 2
!
Router(config-router)#interface FastEthernet0
Router(config-if)#ip address 140.100.17.194 255.255.255.240
Router(config-if)#no ip directed-broadcast
Router(config-if)#ip ospf priority 100
!
Router(config-if)#interface FastEthernet1
Router(config-if)#ip address 140.100.13.1 255.255.255.240
Router(config-if)#no ip directed-broadcast
!
Router(config-if)#exit

```

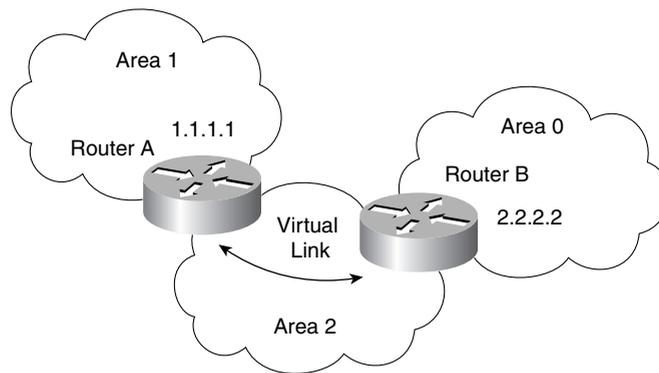
- 1 There are problems with router B. There is inconsistency in the routing table, and the system is extremely slow. What commands would be used to identify the problem? In examining the diagram and configuration, what problems can you see?
- 2 Router A is having problems connecting to Area 0, which is causing problems in other areas because this router is used to connect to Area 0. What commands would be used to identify the problem? In examining the diagram and configuration, what problems can you see?
- 3 Issue the commands that would be used to correct the configuration problems that you see in the example configuration for routers A and B.
- 4 When you issue the **show ip ospf interface** command, you notice that there is a discrepancy in the timers on the link between routers A and B. The transmit timer on Router A is set to 5, and the retransmit timer is set to 1. What problems would this cause? What command would be used to change the timers, and what are the default settings?

- 5 There is an ISDN link into the Internet from Router B. The network manager has suggested that this link is the cause of some performance problems on the router. You have noticed that the interface is included in the OSPF **network** command. What might be the cause of the problem, and how could it be fixed?

## Scenario 6-3

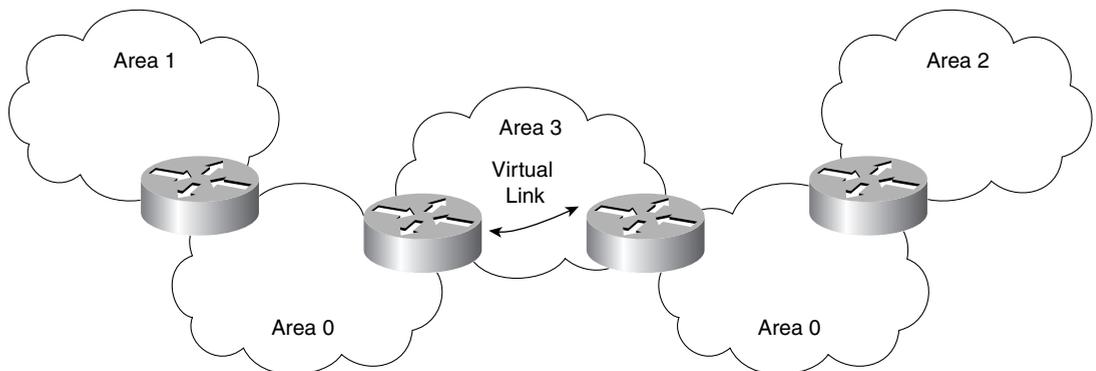
- 1 Explain the purpose of the virtual link in Figure 6-16.

Figure 6-16 Network Diagram #1 for Scenario 6-3



- 2 Is the configuration of the OSPF network shown in Figure 6-17 a valid configuration?

Figure 6-17 Network Diagram #2 for Scenario 6-17



- 3 Why would a company implement this design?

## Scenario Answers

The answers are in **bold**. The answers provided in this section are not necessarily the only possible answers to the questions. The questions are designed to test your knowledge and to give practical exercise in certain key areas. This section is intended to test and exercise skills and concepts detailed in the body of this chapter.

If your answer is different, ask yourself whether it follows the tenets explained in the answers provided. Your answer is correct not if it matches the solution provided in the book, but rather if it has included the principles of design laid out in the chapter.

In this way, the testing provided in these scenarios is deeper: It examines not only your knowledge, but also your understanding and ability to apply that knowledge to problems.

If you do not get the correct answer, refer back to the text and review the subject tested. Be certain to also review your notes on the question to ensure that you understand the principles of the subject.

### Scenario 6-1 Answers

Table 6-7 shows a possible addressing scheme using the criteria stated in Scenario 6-1.

Taking the private address 10.0.0.0, there is a great deal of flexibility in the addressing scheme that can be devised. It must be remembered, however, that careful filtering is required if the organization is to connect to the Internet.

The addressing scheme proposed here is broken out by area. It is not exhaustive in terms of designing an addressing policy down to the LAN level, as was shown in Chapter 3; instead, it deals with the principles of addressing and summarization.

**Note in the allocation of addresses that this scenario deals with the allocation of subnets. The addresses displayed in this table are the address ranges of the available subnets, given the prefix length.**

**Table 6-7** *Allocation of Addresses for Scenario 6-1*

Area	Subnet/Prefix	Subnet Range	Reasons
0	10.0.0.0/28	10.0.0.16 to 10.0.255.224	The use of the zeros in the second octet is an easy reminder that you are in Area 0. Because Area 0 is a transit area, it will be small. The addressing within the area would be allocated the prefix of 28 bits, allowing the range of SUBNETS shown. The number of hosts on each of these subnets would be 15.

**Table 6-7** Allocation of Addresses for Scenario 6-1 (Continued)

Area	Subnet/Prefix	Subnet Range	Reasons
2	10.2.0.0/22	10.2.0.0 to 10.2.252.0	Again, the private addressing of 10.0.0.0 as a Class A address is so large that full use may be made of the documentation advantages of the addressing scheme. The second octet allows Area 2 to be identified. The prefix of 22 bits is used within the area. This allows 1022 hosts on each network, which is good for further VLSM and VLANs in switched environments.
3	10.3.0.0/24	10.3.0.0 to 10.3.255.0	The second octet identifies the area. Within the area, a 24-bit prefix is used to address the LANs.
4	10.4.0.0/28 ... 10.4.100.0/30	Ethernet: 10.4.0.16 to 10.4.255.240  Serial: 10.4.100.0 to 10.4.100.250	The second octet identifies the area. Within the area, a 30-bit mask is used to identify the serial links on which only two addresses are needed. The subnet 10.4.100.0 was chosen as the subnet to sub-subnet for the serial links simply to ease troubleshooting—all serial links in the company in any area would be assigned the third octet of 100. A 28-bit mask was chosen for the Ethernet connections to allow the creation of many subnets.

**Example 6-12 demonstrates a sample configuration for Scenario 6-1. The configuration file is for the central router.**

**Example 6-12** Sample Configuration of Scenario 6-1

```
Router(config)# router ospf 100
! Network Commands
Router(config-router)# network 10.0.0.0 0.0.255.255 area 0
Router(config-router)# network 10.2.0.0 0.0.255.255 area 2
Router(config-router)# network 10.3.0.0 0.0.255.255 area 3
Router(config-router)# network 10.4.0.0 0.0.255.255 area 4
! Totally Stubby Area
Router(config-router)# area 3 stub no-summary
! Stub Area
Router(config-router)# area 4 stub
! Summarization between areas
Router(config-router)# area 0 range 10.0.0.0 255.255.255.0
Router(config-router)# area 2 range 10.2.0.0 255.255.0.0
Router(config-router)# area 3 range 10.3.0.0 255.255.0.0
Router(config-router)# area 4 range 10.4.0.0 255.255.0.0
```

*continues*

**Example 6-12** *Sample Configuration of Scenario 6-1 (Continued)*

```
Router(config)# interface e0
Router(config-if)# ip address 10.4.0.33 255.255.255.240
Router(config)# interface e1
Router(config-if)# ip address 10.4.0.17 255.255.255.240
Router(config)# interface e2
Router(config-if)# ip address 10.2.4.1 255.255.252.0
!Election of Designated Router
Router(config-if)# ip ospf priority 64
Router(config)# interface e3
Router(config-if)# ip address 10.0.0.193 255.255.255.240
!Ensures Router not elected as Designated Router
Router(config-if)# ip ospf priority 0
Router(config)# interface e4
Router(config-if)# ip address 10.0.0.129 255.255.255.240
!Ensures Router not elected as Designated Router
Router(config-if)# ip ospf priority 0
Router(config)# interface e5
Router(config-if)# ip address 10.3.3.1 255.255.255.0
Router(config)# interface e6
Router(config-if)# ip address 10.3.2.1 255.255.255.0
Router(config)# interface e7
Router(config-if)# ip address 10.3.1.1 255.255.255.0
Router(config)# interface s0
Router(config-if)# ip address 10.4.0.9 255.255.255.252
Router(config)# interface s1
Router(config-if)# ip address 10.4.0.5 255.255.255.252
!Router ID set by configuring the Router ID
Router(config)# interface loopback 0
Router(config-if)# ip address 10.100.100.101 255.255.255.255
```

## Scenario 6-2 Answers

- 1 There are problems with Router B. There is inconsistency in the routing table, and the system is extremely slow. What commands would be used to identify the problem? In examining the diagram and configuration, what problems can you see?

**Router B has been configured to be the designated router for the LAN, which means that it is dealing with all the traffic on the LAN associated with the management of OSPF. Given that the system is a 2500, it is a poor choice for a designated router. A better choice would be Router A, which is a larger system that connects directly to Area 0, making it a better choice from the standpoint of the network design. If Router B were a larger system than a 2500, there could be an argument for making it the designated router to elevate Router A, which would otherwise be functioning as the ABR as well as the designated router.**

**The router has not been configured as a stub, so the communication between Router A and Router B will be confused, preventing any communication between the two routers.**

- 2 Router A is having problems connecting to Area 0, which is causing problems in other areas because this router is used to connect to Area 0. What commands would be used to identify the problem? In examining the diagram and configuration, what problems can you see?

**Router A is configured incorrectly. The command that would show the problem would be either `show ip route`, `show ip protocols`, or `show ip ospf database`. The lack of LSA traffic would indicate a configuration problem. When examining the configuration, you would see that the mask on the configuration of the network command for Area 0 is wrong. The mask for 140.100.32.0 should be 0.0.0.15. Using the mask of 0.0.0.0 will place only interfaces with the IP address 140.100.32.0, which is a subnet address not an interface address. Therefore, there will be no communication of OSPF LSAs between the areas.**

- 3 Issue the commands that would be used to correct the configuration problems that you see in the example configuration for routers A and B.

**The commands that would solve these problems are as follows:**

**On Router A:**

```
router ospf 100
network 140.100.32.0 0.0.0.15 area 0
interface fastethernet 1
no ip ospf cost 10
ip ospf priority 100
```

**On Router B:**

```
router ospf 200
network 140.100.13.0 0.0.0.15 area 2
area 2 stub
interface FastEthernet0
no ip ospf priority 100
```

- 4 When you issue the `show ip ospf interface` command, you notice that there is a discrepancy in the timers on the link between routers A and B. The transmit timer on Router A is set to 5, and the retransmit timer is set to 1. What problems would this cause? What command would be used to change the timers, and what are the default settings?

**The default setting for the transmit timer is set to 1 second, and the retransmit timer is set to 5 seconds. The transmit timer determines the estimated number of seconds that it takes to send a LSA to a neighbor. The retransmit timer states the number of seconds to wait for an acknowledgment before retransmitting an LSA.**

**If the transmit timer is not smaller than the retransmit timer, the interface retransmits in the belief that the other side did not receive the LSA. This leads to excess traffic, confusion in the topology database, and the possibility of flapping links. To correct the settings, issue the following subinterface commands:**

```
ip ospf retransmit-interval seconds
ip ospf transmit-delay seconds
```

- 5 There is an ISDN link into the Internet from Router B. The network manager has suggested that this link is the cause of some performance problems on the router. You have noticed that the interface is included in the OSPF **network** command. What might be the cause of the problem, and how could it be fixed?

**If the ISDN interface is configured for dial-on-demand routing (DDR) and is also included in OSPF network commands, you may find that the link that the DDR process establishes will cause the routing updates to be propagated throughout the network, causing additional CPU utilization on the routers and flooding of packets throughout the network. The solution is to ensure that the interface is not included in the network command to the OSPF process. A more important problem is that Router B is in a stub area and will not track external routes. Router B cannot connect to the Internet as an ASBR because it will not propagate the Type 5 LSAs. The BRI interface cannot partake in the OSPF network. Therefore, the network will not be slow; it will be down.**

## Scenario 6-3 Answers

- 1 Explain the purpose of the virtual link in Figure 6-16.

**In this example, Area 1 does not have a direct physical connection into Area 0. A virtual link must be configured between RTA and RTB. Area 2 is to be used as a transit area, and RTB is the entry point into Area 0. This way, RTA and Area 1 will have a logical connection to the backbone.**

- 2 Is the configuration of the OSPF network shown in Figure 6-17 a valid configuration?

**Yes, the configuration is a valid one.**

- 3 Why would a company implement this design?

**OSPF allows for linking discontinuous parts of the backbone using a virtual link. In some cases, different Area 0s need to be linked together. This can occur, for example, if a company is trying to merge two separate OSPF networks into one network with a common Area 0. In other instances, virtual links are added for redundancy in case some router failure causes the backbone to be split in two. Whatever the reason may be, a virtual link can be configured between separate ABRs that touch Area 0 from each side and that have a common area between them.**

