

CCNA 3 v7.0 Curriculum: Module 1 – Single-Area OSPFv2 Concepts

itexamanswers.net/ccna-3-v7-0-curriculum-module-1-single-area-ospfv2-concepts.html

April 6, 2020

1.0. Introduction

1.0.1. Why should I take this module?

Welcome to Single-Area OSPFv2 Concepts!

Welcome to the first module in CCNA Enterprise Networking, Security, and Automation v7.0 (ENSA)!

Imagine that it is time for your family to visit your grandparents. You pack your bags and load them into the car. But this takes a bit longer than you planned for and now you are running late. You pull out your map. There are three different routes. One route is no good because there is a lot of construction on the main road and it is temporarily closed. Another route is very scenic, but it takes an additional hour to get to your destination. The third route is not as pretty but it includes a highway, which is much faster. In fact, it is so much faster that you might actually be on time if you take it.

In networking, packets do not need to take the scenic route. The fastest available route is always the best. Open Shortest Path First (OSPF) is designed to find the fastest available path for a packet from source to destination. This module covers the basic concepts of single-area OSPFv2. Let's get started!

1.0.2. What will I learn to do in this module?

Module Title: Single-Area OSPF Concepts

Module Objective: Explain how single-area OSPF operates in both point-to-point and broadcast multiaccess networks.

Topic Title	Topic Objective
OSPF Features and Characteristics	Describe basic OSPF features and characteristics.
OSPF Packets	Describe the OSPF packet types used in single-area OSPF.
OSPF Operation	Explain how single-area OSPF operates.

1.1. OSPF Features and Characteristics

1.1.1. Introduction to OSPF

This topic is a brief overview of Open Shortest Path First (OSPF), which includes single-area and multiarea. OSPFv2 is used for IPv4 networks. OSPFv3 is used for IPv6 networks. The primary focus of this entire module is single-area OSPFv2.

OSPF is a link-state routing protocol that was developed as an alternative for the distance vector Routing Information Protocol (RIP). RIP was an acceptable routing protocol in the early days of networking and the internet. However, the RIP reliance on hop count as the only metric for determining best route quickly became problematic. Using hop count does not scale well in larger networks with multiple paths of varying speeds. OSPF has significant advantages over RIP in that it offers faster convergence and scales to much larger network implementations.

OSPF is a link-state routing protocol that uses the concept of areas. A network administrator can divide the routing domain into distinct areas that help control routing update traffic. A link is an interface on a router. A link is also a network segment that connects two routers, or a stub network such as an Ethernet LAN that is connected to a single router. Information about the state of a link is known as a link-state. All link-state information includes the network prefix, prefix length, and cost.

This module covers basic, single-area OSPF implementations and configurations.

1.0.5 Packet Tracer – Logical and Physical Mode Exploration

The network model in this Packet Tracer Physical Mode (PTPM) activity incorporates many of the technologies that you can master in Cisco Networking Academy courses. It represents a simplified version of how a small to medium-sized business network might look.

Most of the devices in the Seward branch office and Warrenton data center are already deployed and configured. You have just been hired to review the devices and networks deployed. It is not important that you understand everything you see and do in this activity. Feel free to explore the network on your own. If you wish to proceed more systematically, follow the steps below. Answer the questions to the best of your ability.

1.0.5 Packet Tracer – Logical and Physical Mode Exploration

1.1.2. Components of OSPF

All routing protocols share similar components. They all use routing protocol messages to exchange route information. The messages help build data structures, which are then processed using a routing algorithm.

Click each OSPF component below for more information.

Routers running OSPF exchange messages to convey routing information using five types of packets. These packets, as shown in the figure, are as follows:

- Hello packet
- Database description packet
- Link-state request packet
- Link-state update packet
- Link-state acknowledgment packet

These packets are used to discover neighboring routers and also to exchange routing information to maintain accurate information about the network.



1.1.3. Link-State Operation

To maintain routing information, OSPF routers complete a generic link-state routing process to reach a state of convergence. The figure shows a five router topology. Each link between routers is labeled with a cost value. In OSPF, cost is used to determine the best path to the destination. The following are the link-state routing steps that are completed by a router:

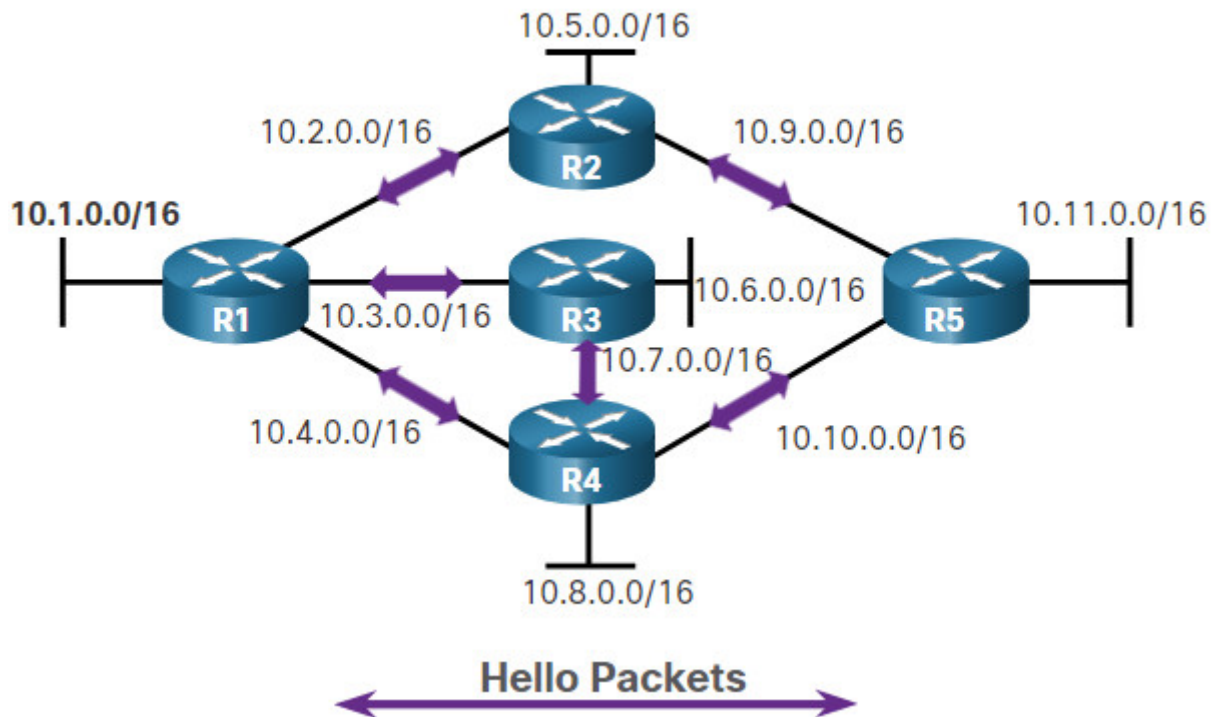
1. Establish Neighbor Adjacencies
2. Exchange Link-State Advertisements
3. Build the Link State Database
4. Execute the SPF Algorithm
5. Choose the Best Route

Click each button for an illustration of the steps in the link-state routing process that R1 uses to reach convergence.

1. Establish Neighbor Adjacencies

OSPF-enabled routers must recognize each other on the network before they can share information. An OSPF-enabled router sends Hello packets out all OSPF-enabled interfaces to determine if neighbors are present on those links. If a neighbor is present, the OSPF-enabled router attempts to establish a neighbor adjacency with that neighbor.

Routers Exchange Hello Packets



1.1.4. Single-Area and Multiarea OSPF

To make OSPF more efficient and scalable, OSPF supports hierarchical routing using areas. An OSPF area is a group of routers that share the same link-state information in their LSDBs. OSPF can be implemented in one of two ways, as follows:

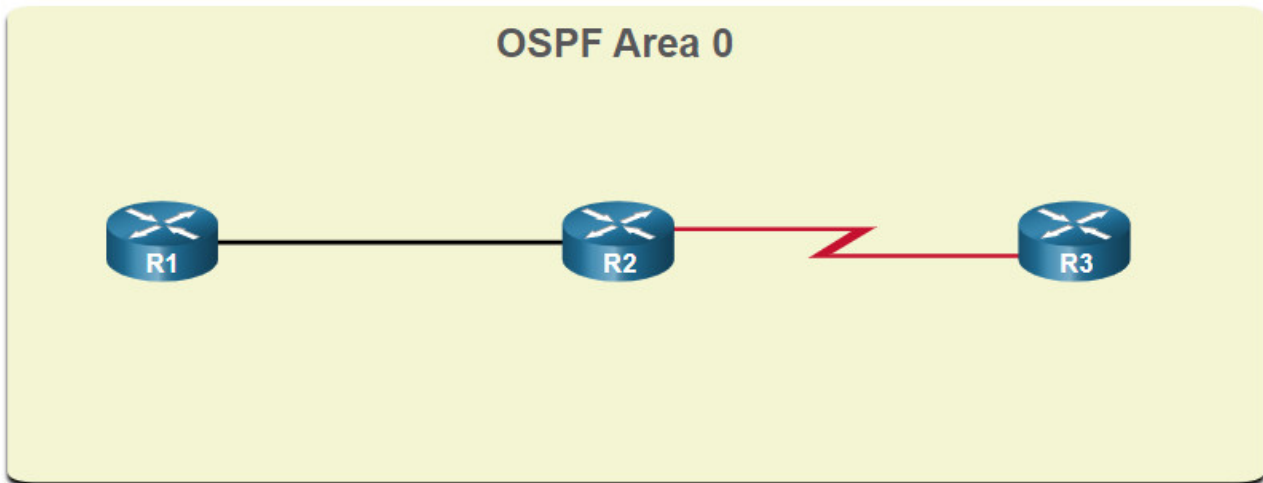
- **Single-Area OSPF** – All routers are in one area. Best practice is to use area 0.
- **Multiarea OSPF** – OSPF is implemented using multiple areas, in a hierarchical fashion. All areas must connect to the backbone area (area 0). Routers interconnecting the areas are referred to as Area Border Routers (ABRs).

The focus of this module is on single-area OSPFv2.

Click each button to compare single-area and multiarea OSPF.

- [Single-Area OSPF](#)
- [Multiarea OSPF](#)

Single-Area OSPF



1.1.5. Multiarea OSPF

With multiarea OSPF, one large routing domain can be divided into smaller areas, to support hierarchical routing. Routing still occurs between the areas (interarea routing), while many of the processor intensive routing operations, such as recalculating the database, are kept within an area.

For instance, any time a router receives new information about a topology change within the area (including the addition, deletion, or modification of a link) the router must rerun the SPF algorithm, create a new SPF tree, and update the routing table. The SPF algorithm is CPU-intensive and the time it takes for calculation depends on the size of the area.

Note: Routers in other areas receive updates regarding topology changes, but these routers only update the routing table, not rerun the SPF algorithm.

Too many routers in one area would make the LSDBs very large and increase the load on the CPU. Therefore, arranging routers into areas effectively partitions a potentially large database into smaller and more manageable databases.


The hierarchical-topology design options with multiarea OSPF can offer the following advantages.

- **Smaller routing tables** – Tables are smaller because there are fewer routing table entries. This is because network addresses can be summarized between areas. Route summarization is not enabled by default.
- **Reduced link-state update overhead** – Designing multiarea OSPF with smaller areas minimizes processing and memory requirements.
- **Reduced frequency of SPF calculations** – Multiarea OSPF localize the impact of a topology change within an area. For instance, it minimizes routing update impact because LSA flooding stops at the area boundary.

For example, in the figure R2 is an ABR for area 51. A topology change in area 51 would cause all area 51 routers to rerun the SPF algorithm, create a new SPF tree and update their IP routing tables. The ABR, R2, would send an LSA to routers in the area 0, which

would eventually be flooded to all routers in the OSPF routing domain. This type of LSA does not cause routers in other areas to rerun the SPF algorithm. They only have to update their LSDB and routing table.

Link Change Impacts Local Area Only



- Link failure affects the local area only (area 51).
- The ABR (R2) isolates the flooding of a specific LSA to area 51.
- Routers in areas 0 and 1 do not need to run the SPF algorithm.

1.1.6. OSPFv3

OSPFv3 is the OSPFv2 equivalent for exchanging IPv6 prefixes. Recall that in IPv6, the network address is referred to as the prefix and the subnet mask is called the prefix-length.

Similar to its IPv4 counterpart, OSPFv3 exchanges routing information to populate the IPv6 routing table with remote prefixes.

Note: With the OSPFv3 Address Families feature, OSPFv3 includes support for both IPv4 and IPv6. OSPF Address Families is beyond the scope of this curriculum.

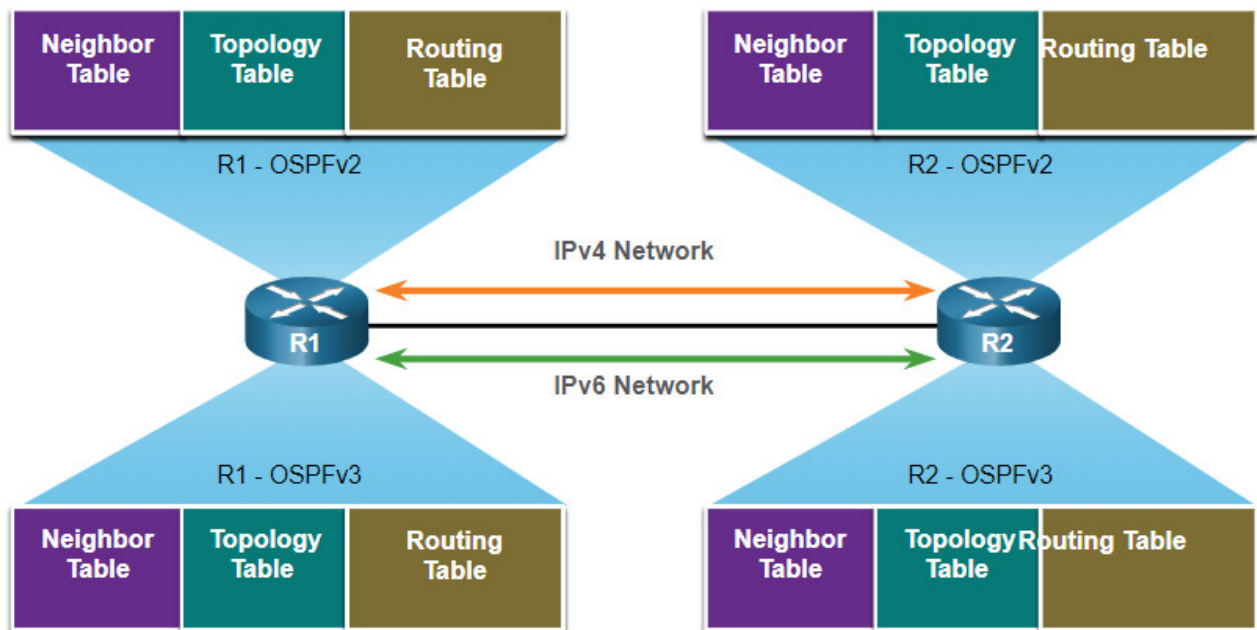
OSPFv2 runs over the IPv4 network layer, communicating with other OSPF IPv4 peers, and advertising only IPv4 routes.

OSPFv3 has the same functionality as OSPFv2, but uses IPv6 as the network layer transport, communicating with OSPFv3 peers and advertising IPv6 routes. OSPFv3 also uses the SPF algorithm as the computation engine to determine the best paths throughout the routing domain.

OSPFv3 has separate processes from its IPv4 counterpart. The processes and operations are basically the same as in the IPv4 routing protocol, but run independently. OSPFv2 and OSPFv3 each have separate adjacency tables, OSPF topology tables, and IP routing tables, as shown in the figure.

The OSPFv3 configuration and verification commands are similar to those used in OSPFv2.

OSPFv2 and OSPFv3 Data Structures



1.2. OSPF Packets

1.2.1 Video – OSPF Packets

Click Play in the figure to view a video about OSPF packets.

1.2.2. Types of OSPF Packets

Link-state packets are the tools used by OSPF to help determine the fastest available route for a packet. OSPF uses the following link-state packets (LSPs) to establish and maintain neighbor adjacencies and exchange routing updates. Each packet serves a specific purpose in the OSPF routing process, as follows:

- **Type 1: Hello packet** – This is used to establish and maintain adjacency with other OSPF routers.
- **Type 2: Database Description (DBD) packet** – This contains an abbreviated list of the LSDB of the sending router and is used by receiving routers to check against the local LSDB. The LSDB must be identical on all link-state routers within an area to construct an accurate SPF tree.
- **Type 3: Link-State Request (LSR) packet** – Receiving routers can then request more information about any entry in the DBD by sending an LSR.
- **Type 4: Link-State Update (LSU) packet** – This is used to reply to LSRs and to announce new information. LSUs contain several different types of LSAs.

- **Type 5: Link-State Acknowledgment (LSAck) packet** – When an LSU is received, the router sends an LSAck to confirm receipt of the LSU. The LSAck data field is empty.

The table summarizes the five different types of LSPs used by OSPFv2. OSPFv3 has similar packet types.

Type	Packet Name	Description
1	Hello	Discovers neighbors and builds adjacencies between them
2	Database Description (DBD)	Checks for database synchronization between routers
3	Link-State Request (LSR)	Requests specific link-state records from router to router
4	Link-State Update (LSU)	Sends specifically requested link-state records
5	Link-State Acknowledgment (LSAck)	Acknowledges the other packet types

1.2.3. Link-State Updates

Routers initially exchange Type 2 DBD packets, which is an abbreviated list of the LSDB of the sending router. It is used by receiving routers to check against the local LSDB.

A Type 3 LSR packet is used by the receiving routers to request more information about an entry in the DBD.

The Type 4 LSU packet is used to reply to an LSR packet.

A Type 5 packet is used to acknowledge the receipt of a Type 4 LSU.

LSUs are also used to forward OSPF routing updates, such as link changes. Specifically, an LSU packet can contain 11 different types of OSPFv2 LSAs, with some of the more common ones shown in the figure. OSPFv3 renamed several of these LSAs and also contains two additional LSAs.

Note: The difference between the LSU and LSA terms can sometimes be confusing because these terms are often used interchangeably. However, an LSU contains one or more LSAs.

LSUs Contain LSAs

LSUs		
Type	Packet Name	Description
1	Hello	Discovers neighbors and builds adjacencies between them
2	DBD	Checks for database synchronization between routers
3	LSR	Requests specific link-state records from router to router
4	LSU	Sends specifically requested link-state records
5	LSAck	Acknowledges the other packet types



LSAs	
LSA Type	Description
1	Router LSAs
2	Checks for database synchronization between routers
3 or 4	Summary LSAs
5	Autonomous System External LSAs
6	Multicast OSPF LSAs
7	Defined for Not-So-Stubby Areas
8	External Attributes LSA for Border Gateway Patrol (BGPs)

- An LSU contains one or more LSAs.
- LSAs contain route information for destination networks.

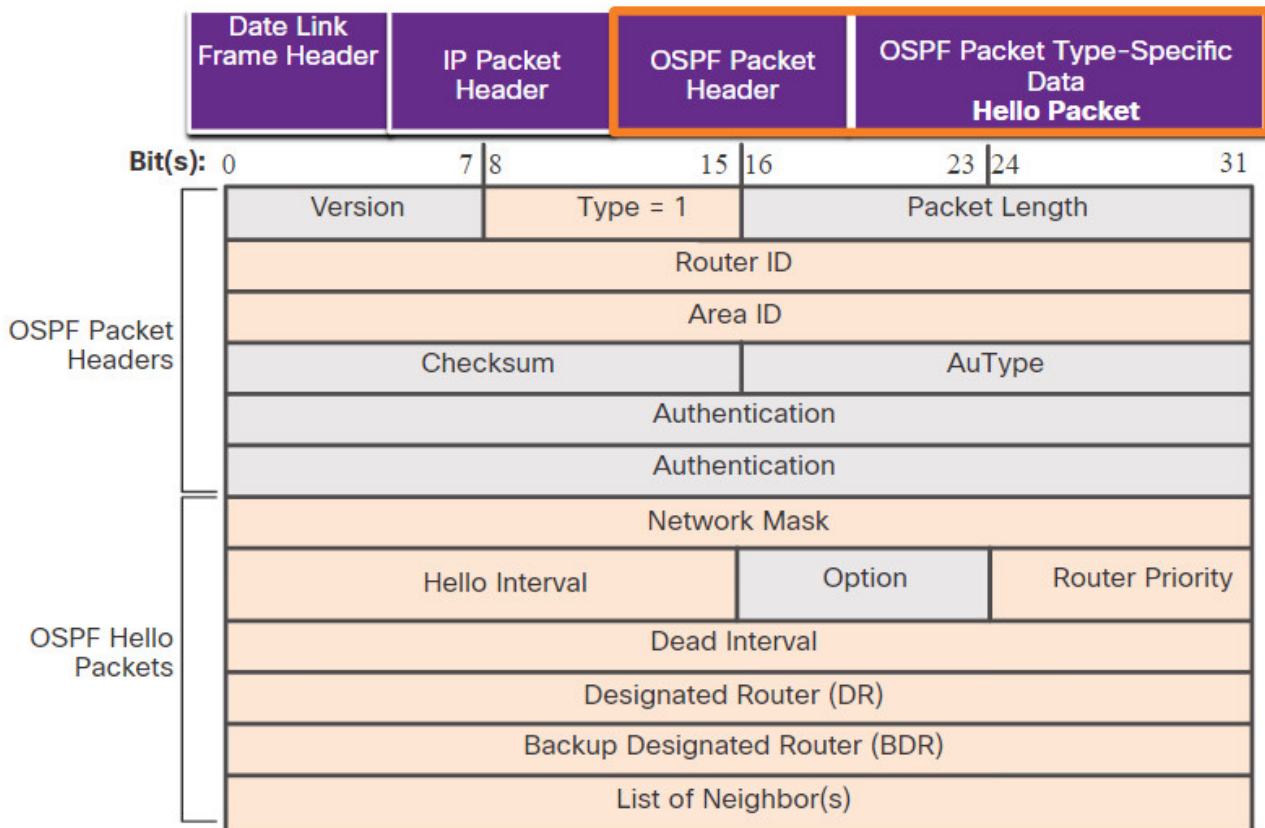
1.2.4. Hello Packet

The OSPF Type 1 packet is the Hello packet. Hello packets are used to do the following:

- Discover OSPF neighbors and establish neighbor adjacencies.
- Advertise parameters on which two routers must agree to become neighbors.
- Elect the Designated Router (DR) and Backup Designated Router (BDR) on multiaccess networks like Ethernet. Point-to-point links do not require DR or BDR.

The figure displays the fields contained in the OSPFv2 Type 1 Hello packet.

OSPF Hello Packet Content



Important fields shown in the figure include the following:

- **Type** – This identifies the type of packet. A value 1 indicates a Hello packet. A value 2 identifies a DBD packet, 3 an LSR packet, 4 an LSU packet, and 5 an LSAck packet.
- **Router ID** – A 32-bit value expressed in dotted decimal notation (like an IPv4 address) is used to uniquely identify the originating router.
- **Area ID** – This is the number of the area from which the packet originated.
- **Network Mask** – This is the subnet mask associated with the sending interface.
- **Hello Interval** – This specifies the frequency, in seconds, at which a router sends Hello packets. The default Hello interval on multiaccess networks is 10 seconds. This timer must be the same on neighboring routers; otherwise, an adjacency is not established.
- **Router Priority** – This is used in a DR/BDR election. The default priority for all OSPF routers is 1, but can be manually altered from 0 to 255. The higher the value, the more likely the router becomes the DR on the link.
- **Dead Interval** – This is the time, in seconds, that a router waits to hear from a neighbor before declaring the neighboring router out of service. By default, the router Dead Interval is four times the Hello interval. This timer must be the same on neighboring routers; otherwise, an adjacency is not established.
- **Designated Router (DR)** – This is the router ID of the DR.
- **Backup Designated Router (BDR)** – This is the router ID of the BDR.
- **List of Neighbors** – This list identifies the router IDs of all adjacent routers.

1.3. OSPF Operation

1.3.1 Video – OSPF Operation

Click Play in the figure to view a video about OSPF operation.

1.3.2. OSPF Operational States

Now that you know about the OSPF link-state packets, this topic explains how they work with OSPF-enabled routers. When an OSPF router is initially connected to a network, it attempts to:

- Create adjacencies with neighbors
- Exchange routing information
- Calculate the best routes
- Reach convergence

The table details the states OSPF progresses through while attempting to reach convergence:

State	Description
Down State	<ul style="list-style-type: none">• No Hello packets received = Down.• Router sends Hello packets.• Transition to Init state.
Init State	<ul style="list-style-type: none">• Hello packets are received from the neighbor.• They contain the Router ID of the sending router.• Transition to Two-Way state.
Two-Way State	<ul style="list-style-type: none">• In this state, communication between the two routers is bidirectional.• On multiaccess links, the routers elect a DR and a BDR.• Transition to ExStart state.
ExStart State	On point-to-point networks, the two routers decide which router will initiate the DBD packet exchange and decide upon the initial DBD packet sequence number.
Exchange State	<ul style="list-style-type: none">• Routers exchange DBD packets.• If additional router information is required then transition to Loading; otherwise, transition to the Full state.
Loading State	<ul style="list-style-type: none">• LSRs and LSUs are used to gain additional route information.• Routes are processed using the SPF algorithm.• Transition to the Full state.
Full State	The link-state database of the router is fully synchronized.

1.3.3. Establish Neighbor Adjacencies

When OSPF is enabled on an interface, the router must determine if there is another OSPF neighbor on the link. To accomplish this, the router sends a Hello packet that contains its router ID out all OSPF-enabled interfaces. The Hello packet is sent to the reserved All OSPF Routers IPv4 multicast address 224.0.0.5. Only OSPFv2 routers will process these packets. The OSPF router ID is used by the OSPF process to uniquely identify each router in the OSPF area. A router ID is a 32-bit number formatted like an IPv4 address and assigned to uniquely identify a router among OSPF peers.

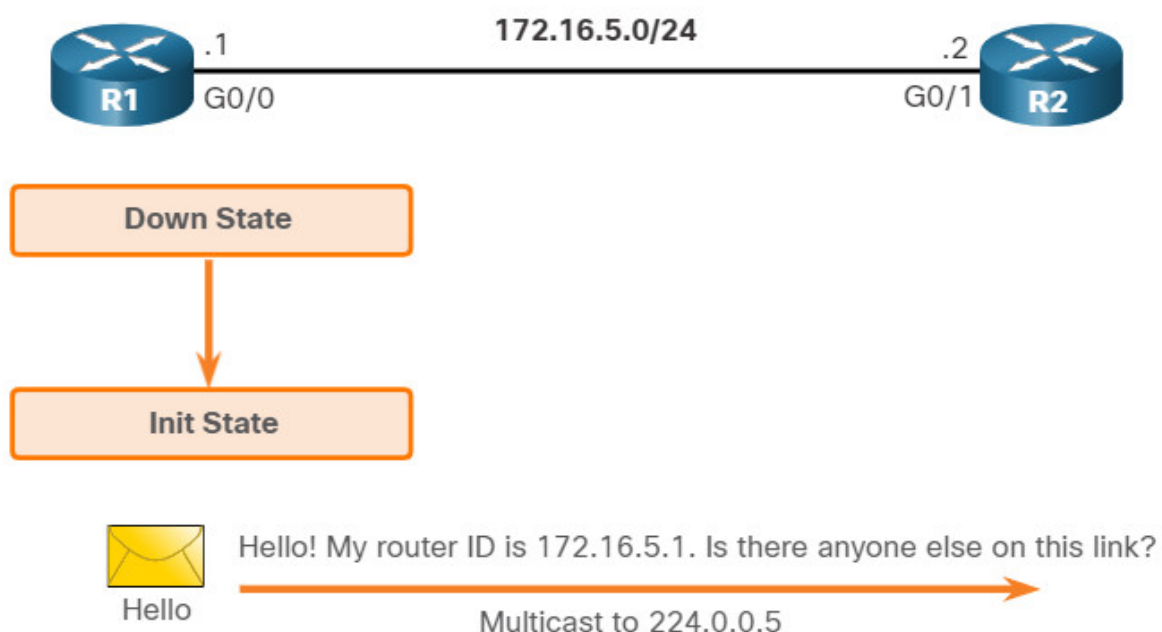
When a neighboring OSPF-enabled router receives a Hello packet with a router ID that is not within its neighbor list, the receiving router attempts to establish an adjacency with the initiating router.

Click each button below to step through the process routers use to establish adjacency on a multiaccess network.

- [1. Down State to Init State](#)
- [2. The Init State](#)
- [3. Two-Way State](#)
- [4. Elect the DR and BDR](#)

1. Down State to Init State

When OSPFv2 is enabled, the enabled Gigabit Ethernet 0/0 interface transitions from the Down state to the Init state. R1 starts sending Hello packets out all OSPF-enabled interfaces to discover OSPF neighbors to develop adjacencies with.



1.3.4. Synchronizing OSPF Databases

After the Two-Way state, routers transition to database synchronization states. While the Hello packet was used to establish neighbor adjacencies, the other four types of OSPF packets are used during the process of exchanging and synchronizing LSDBs. This is a three step process, as follows:

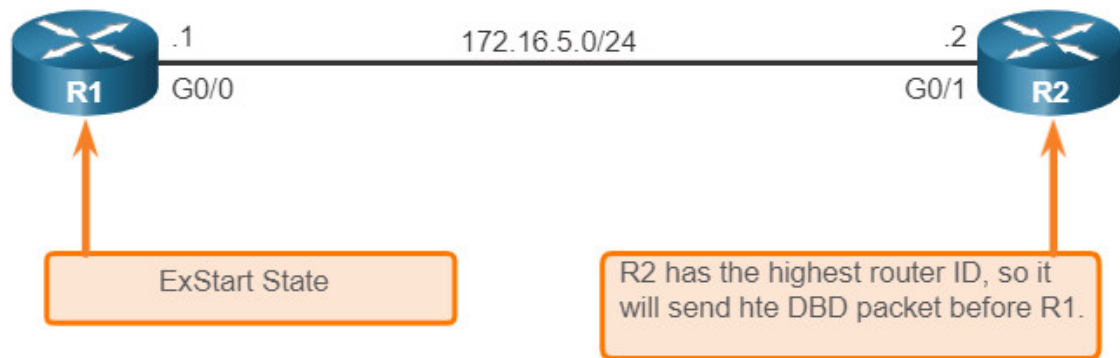
1. Decide first router
2. Exchange DBDs
3. Send an LSR

Click each button below to step through the process routers use to synchronize their LSDBs.

- [1. Decide First Router](#)
- [2. Exchange DBDs](#)
- [3. Send an LSR](#)

1. Decide First Router

In the ExStart state, the two routers decide which router will send the DBD packets first. The router with the higher router ID will be the first router to send DBD packets during the Exchange state. In the figure, R2 has the higher router ID and sends its DBD packets first.



Decide Which Router Sends the First DBD

1.3.5. The Need for a DR

Why is a DR and BDR election necessary?

Multiaccess networks can create two challenges for OSPF regarding the flooding of LSAs, as follows:

- **Creation of multiple adjacencies** – Ethernet networks could potentially interconnect many OSPF routers over a common link. Creating adjacencies with every router is unnecessary and undesirable. It would lead to an excessive number of LSAs exchanged between routers on the same network.
- **Extensive flooding of LSAs** – Link-state routers flood their LSAs any time OSPF is initialized, or when there is a change in the topology. This flooding can become excessive.

To understand the problem with multiple adjacencies, we must study a formula:


For any number of routers (designated as n) on a multiaccess network, there are $n(n - 1) / 2$ adjacencies.

For example, the figure shows a simple topology of five routers, all of which are attached to the same multiaccess Ethernet network. Without some type of mechanism to reduce the number of adjacencies, collectively these routers would form 10 adjacencies:

$$5 (5 - 1) / 2 = 10$$

This may not seem like much, but as routers are added to the network, the number of adjacencies increases dramatically. For example, a multiaccess network with 20 routers would create 190 adjacencies.

Creating Adjacencies With Every Neighbor



- Number of Adjacencies = $n (n - 1) / 2$
- n = number of routers
- Example: $5 (5 - 1) / 2 = 10$ adjacencies

1.3.6. LSA Flooding With a DR

A dramatic increase in the number of routers also dramatically increases the number of LSAs exchanged between the routers. This flooding of LSAs significantly impacts the operation of OSPF.


Click each button to compare the flooding of LSAs without and with a DR.

- [Flooding LSAs](#)
- [LSAs and DR](#)

Flooding LSAs

To understand the problem of extensive flooding of LSAs, play the animation in the figure. In the animation, R2 sends out an LSA. This event triggers every other router to also send out an LSA. Not shown in the animation are the required acknowledgments sent for

every LSA received. If every router in a multiaccess network had to flood and acknowledge all received LSAs to all other routers on that same multiaccess network, the network traffic would become quite chaotic.



1.4. Module Practice and Quiz

1.4.1. What did I learn in this module?

OSPF Features and Characteristics

Open Shortest Path First (OSPF) is a link-state routing protocol that was developed as an alternative for the distance vector Routing Information Protocol (RIP). OSPF has significant advantages over RIP in that it offers faster convergence and scales to much larger network implementations. OSPF is a link-state routing protocol that uses the concept of areas for scalability. A link is an interface on a router. A link is also a network segment that connects two routers, or a stub network such as an Ethernet LAN that is connected to a single router. All link-state information includes the network prefix, prefix length, and cost. All routing protocols use routing protocol messages to exchange route information. The messages help build data structures, which are then processed using a routing algorithm. Routers running OSPF exchange messages to convey routing information using five types of packets: the Hello packet, the database description packet, the link-state request packet, the link-state update packet, and the link-state acknowledgment packet. OSPF messages are used to create and maintain three OSPF databases: the adjacency database creates the neighbor table, the link-state database (LSDB) creates the topology table, and the forwarding database creates the routing table. The router builds the topology table using results of calculations based on the Dijkstra SPF (shortest-path first) algorithm. The SPF algorithm is based on the cumulative cost to reach a destination. In OSPF, cost is used to determine the best path to the destination. To maintain routing information, OSPF routers complete a generic link-state routing process to reach a state of convergence:

1. Establish Neighbor Adjacencies
2. Exchange Link-State Advertisements
3. Build the Link State Database
4. Execute the SPF Algorithm
5. Choose the Best Route

With single-area OSPF any number can be used for the area, best practice is to use area 0. Single-area OSPF is useful in smaller networks with few routers. With multiarea OSPF, one large routing domain can be divided into smaller areas, to support hierarchical routing. Routing still occurs between the areas (interarea routing), while many of the processor intensive routing operations, such as recalculating the database, are kept within an area. OSPFv3 is the OSPFv2 equivalent for exchanging IPv6 prefixes. Recall that in IPv6, the network address is referred to as the prefix and the subnet mask is called the prefix-length.

OSPF Packets

OSPF uses the following link-state packets (LSPs) to establish and maintain neighbor adjacencies and exchange routing updates: 1 Hello, 2 DBD, 3 LSR, 4 LSU, and 5 LSAck. LSUs are also used to forward OSPF routing updates, such as link changes. Hello packets are used to:

- Discover OSPF neighbors and establish neighbor adjacencies.
- Advertise parameters on which two routers must agree to become neighbors.
- Elect the Designated Router (DR) and Backup Designated Router (BDR) on multiaccess networks like Ethernet. Point-to-point links do not require DR or BDR.

Some important fields in the Hello packet are type, router ID, area ID, network mask, hello interval, router priority, dead interval, DR, BDR and list of neighbors.

OSPF Operation

When an OSPF router is initially connected to a network, it attempts to:

- Create adjacencies with neighbors
- Exchange routing information
- Calculate the best routes
- Reach convergence

The states that OSPF progresses through to do this are down state, init state, two-way state, ExStart state, Exchange state, loading state, and full state. When OSPF is enabled on an interface, the router must determine if there is another OSPF neighbor on the link by sending a Hello packet that contains its router ID out all OSPF-enabled interfaces. The Hello packet is sent to the reserved All OSPF Routers IPv4 multicast address 224.0.0.5. Only OSPFv2 routers will process these packets. When a neighboring OSPF-enabled router receives a Hello packet with a router ID that is not within its neighbor list, the

receiving router attempts to establish an adjacency with the initiating router. After the Two-Way state, routers transition to database synchronization states, which is a three step process:

1. Decide First Router
2. Exchange DBDs
3. Send an LSR

Multiaccess networks can create two challenges for OSPF regarding the flooding of LSAs: the creation of multiple adjacencies and extensive flooding of LSAs. A dramatic increase in the number of routers also dramatically increases the number of LSAs exchanged between the routers. This flooding of LSAs significantly impact the operation of OSPF. If every router in a multiaccess network had to flood and acknowledge all received LSAs to all other routers on that same multiaccess network, the network traffic would become quite chaotic. This is why DR and BDR election is necessary. On multiaccess networks, OSPF elects a DR to be the collection and distribution point for LSAs sent and received. A BDR is also elected in case the DR fails.

1.4.2 Module Quiz – Single-Area OSPFv2 Concepts

Post navigation

[← Previous Article](#)

[CCNA 1 v7.0 Curriculum: Module 17 – Build a Small Network](#)

[Next Article →](#)

[CCNA 3 v7.0 Curriculum: Module 2 – Single-Area OSPFv2 Configuration](#)