

Open Shortest Path First (OSPF)

Open Shortest Path First (OSPF) is a routing protocol developed for Internet Protocol (IP) networks by the Interior Gateway Protocol (IGP) working group of the Internet Engineering Task Force (IETF). The working group was formed in 1988 to design an IGP based on the Shortest Path First (SPF) algorithm for use in the Internet. Similar to the Interior Gateway Routing Protocol (IGRP), OSPF was created because in the mid-1980s, the Routing Information Protocol (RIP) was increasingly incapable of serving large, heterogeneous internetworks. This chapter examines the OSPF routing environment, underlying routing algorithm, and general protocol components. **Building Routing Table on New OSPF-Configured Routers**

Five packets are used to build the routing table on a new OSPF-configured router. These are the Hello protocol; the database descriptor, which is used to send summary information to neighbors to synchronize topology databases; the link-state request, which works as a request for more detailed information that is sent when the router receives a database descriptor that contains new information; the link-state update, which works as the link-state advertisement (LSA) packet issued in response to the request for database information in the link-state request packet; and the link-state acknowledgement, which acknowledges the link-state update. When the new OSPF-configured router is connected to the network, it must learn the network from the routers that are up and running. The router goes through three stages while exchanging information: the down state, the init stage, and the two-way state. You can check what stage an interface running OSPF is in by using the `show ip ospf neighbor` command or the `debug ip ospf adjacency` command.

. The new router starts in a down state. It transmits its own Hello packets to introduce itself to the segment and to find any other OSPF-configured routers. This is sent out as a Hello to the multicast address 224.0.0.5 (AllSPFRouters). It sets the designated router (DR) and the backup designated router (BDR) in the Hello to 0.0.0.0. While the new router waits for a reply, which usually is four times the length of the Hello timer, the router is in the init state. Within the wait time, the new router hears a Hello from another router and learns the DR and the BDR. If there is no DR or BDR stated in the incoming Hello, an election takes place.

. Once the new router sees its own router ID in the list of neighbors, and a neighbor relationship is established, it changes its status to the two-way state. The new router and the DR have now established a neighbor relationship and need to ensure that the new router has all the relevant information about the network. The DR must update and synchronize the topology database of the new router. This is achieved by using the exchange protocol with the database description packets (DDPs). There are four different stages that the router goes through while exchanging routing information with a neighbor: the exstart state, the exchange state, the loading state, and the full state.

. During the exstart state, one of the routers will take seniority and become the master router, based on highest IP interface address.

. Both routers will send out database description packets, changing the state to the exchange state. At this stage, the new router has no knowledge and can inform the DR only of the networks or links to which it is directly connected. The DR sends out a series of DDPs containing the networks, referred to as links that are held in the topology database. Most of these links have been received from other routers via linkstate advertisements (LSAs). The source of the link information is referred to by the router ID. Each link will have an interface ID for the outgoing interface, a link ID, and a metric to state the value of the path. The DDPs will contain a summary rather than all the necessary information. When the router has received the DDPs from the neighboring router, it compares the received network information with that in its topology table. In the case of a new router, all the DDPs are new.

. If the new router requires more information, it will request that particular link in more detail using the link-state request packet (LSR). The LSR will prompt the master router to send the link-state update packet (LSU). This is the same as a LSA used to flood the network with routing information. While the new router is awaiting the LSUs from its neighbor, it is in the loading state.

. When these LSRs are received and the databases are updated and synchronized, the neighbors are fully adjacent. This is the full state.

OSPF Areas

There are a number of problems associated with using OSPF. These problems are related to the network size. The larger the

network, the greater the probability of a network change, which would require a recalculation of the whole area. This increases the frequency with which the SPF algorithm is being run. In addition, each recalculation will take longer. As the network grows, the size of the routing table will increase. Although the complete routing table is not sent out as in a distance vector routing protocol, the greater the size of the table, the longer each lookup becomes. The memory requirements on the router will also increase. Furthermore, the topological database will increase in size and will eventually become unmanageable. As the various databases increase in size and the calculations become increasingly frequent, the CPU utilization will increase as more of the available memory is consumed. This will have a negative impact on network response time, not because of congestion on the line but because of congestion within the router itself.

Using multiple OSPF areas solves most of the common problems with running OSPF in larger networks. The division of a large single area network into multiple areas allows routers in each area to maintain their own topological databases. This limits the size of the topological databases within an area, which results in routers requiring less memory and processing time to run SPF, and a decrease in convergence time. Summary and external links ensure connectivity between areas and networks outside the autonomous area (AS). This is achieved by creating areas from groups of subnets. Each area is treated internally as a small entity on its own. It communicates with the other areas, exchanging routing information which is kept to a minimum by allowing only that information that is required for connectivity.

There are two approaches to implementing multiple area networks. The first approach is to grow a single area until it becomes unmanageable. This approach requires less initial work and configuration but care should be put into the design of the network because this may cause problems in the future, particularly in addressing. The second approach is to design the network with multiple areas, which are very small, in anticipation that the networks will grow to fit comfortably into their areas. 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.

OSPF Area Types

Regardless of which approach is used, a multiple area OSP network has a hierarchical structure and consists a number of distinct areas.

These areas are:

- . The backbone area, which is also referred to as Area 0. All other areas must connect to the backbone area. Hence, this area is obligatory.
- . An ordinary or standard area, which is an area that connects to the backbone (Area 0) and is treated as a separate entity. All routers in a standard area have the same topological database, but their routing tables will be based on the routers position in the standard area and will thus be unique to the router.
- . A stub area, which is an area that does not accept external summary routes. A router within a stub area can only see outside the autonomous system if a default route has been configuration for it.
- . A totally stubby area, which is similar to a stub area. In this area, the default route must be configured as 0.0.0.0. This type of area is useful for remote sites that have few networks and limited connectivity with the rest of the network and is a Cisco proprietary solution.
- . A not so stubby area (NSSA), which is a stub area that can receive external routes but will not propagate those external routes into the backbone area.

Router Responsibilities

Because of the hierarchical nature of a multiple area OSPF network, routers have different responsibilities, depending on their position and functionality within the hierarchical design. These routers have different designations such as internal routers, backbone routers, area border routers (ABR), and autonomous system boundary routers (ASBR).

. The Internal Router exists within an area. 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.

. The Backbone Router exists within the backbone area, which is also called Area 0. The design rules for OSPF require that all the areas be connected through a single area, known as Area 0. Area 0 is also known as Area 0.0.0.0 on other routers. A router within this area is referred to as a backbone router. It may also be an internal router or an Area Border Router.

. The Area Border Router (ABR) 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.

. The Autonomous System Boundary Router (ASBR) is used to connect to a network or routing protocol outside the OSPF domain. OSPF is an interior routing protocol or Interior Gateway Protocol (IGP); gateway is an older term for a router. If there is any redistribution between other protocols to OSPF on a router, it will be an ASBR. This router should reside in the backbone area but you can place it anywhere in the OSPF hierarchical design.