Fall 10-2012

Implementing OSPF for ISP

Milot Cakaj

Faculty of Computer Sciences and Engineering

Implementing OSPF for ISP

Bachelor Thesis

Milot Cakaj

October / 2012
Prishtina

Faculty of Computer Sciences and Engineering

Bachelor Thesis
Academic Year 2011 – 2012

Student: Milot Cakaj

Implementing OSPF for ISP

Supervisor: Petrit Shala

Date: October/2012

This thesis is submitted in partial fulfillment of the requirements for a Bachelor Degree

ABSTRACT

Routing protocol has a vital role in the modern communication networks. As there are various routing protocols available, it is very important to select a right protocol among them. This selection depends on several parameters such as network convergence time, network scalability and bandwidth requirements.

Each of the different routing protocols, in the context of routing protocol performance, has different architecture, adaptability, route processing delays and convergence capabilities.

Among them, Open Shortest Path First (OSPF) is considered as one of the pre-eminent routing protocols for large networks. OSPF is dynamic routing protocols used in large networks to disseminate network topology towards the adjacent routers. OSPF is considered an open-standards protocol, with many vendors offering its implementation based on RFC 2328.

This thesis tries to argue the suitability of implementing OSPF for ISPs, namely the advantages of transition from static routing to deployment of dynamic routing protocol. As a dynamic protocol, OSPF performs real-time calculations and promptly reacts to topological changes within a network, thus offering fault-tolerance if a link goes down. With the possibility of configuring multiple areas, OSPF offers scalability of the network by minimizing the routes and reducing the size of routing table.

ACKNOWLEDGEMENT

Initially, I would like to thank my supervisor, Professor Petrit Shala for his continuous guidance and advice in order to successfully finish the thesis. He has been a patient and very considerate mentor from the very beginning and managed to do all of this with a lot of humor. I would also like to thank Professor Selman Haxhijaha for all his technical expertise and for keeping me focused on the thesis. He helped in improving the quality of the thesis by giving invaluable feedback.

Very special thanks, goes to Professor Fisnik Prekazi for all the encouragement, for being a friend and for believing when I didn't.

Lastly, I would like to thank Rina and my family for their love, understanding and support.

CONTENT

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

1.1. Background

Open Shortest Path First (OSPF) is a TCP/IP internet routing protocol that was developed by Internet Engineering Task Force (IETF). OSPF is classified as an Interior Gateway Protocol (IGP), meaning that it distributes routing information between routers belonging to a single Autonomous System [1]. In 1989, the first version of OSPF was defined as OSPFv1, which was published in RFC 1131. The second version of OSPFv2 was introduced in 1998, which was defined in RFC 2328. In 1999, the third version of OSPFv3 for IPv6 was released in RFC 2740 (superseded by RFC 5340 in 2008). The entire list of RFC related to OSPF can be found in IETF OSPF Working Group [2].

OSPF is a link-state routing protocol. Each OSPF router maintains an identical database describing the Autonomous System's topology. From this database, a routing table is calculated by constructing a shortest-path tree. OSPF recalculates routes quickly in the face of topological changes, utilizing a minimum of routing protocol traffic. OSPF provides support for equal-cost multipath. An area routing capability is provided, enabling an additional level of routing protection and a reduction in routing protocol traffic [1].

1.2. Problem statement

A Kosovar Internet Service Provider (ISP) company is expanding its business rapidly. The company is based in Kosova's capitol Prishtina, where its main operation lies. Besides Prishtina, the company has also started offering its services in other major cities of Kosova, thus becoming a national provider. The company uses static routing throughout their network infrastructure.

In static routing, network administrator must manually create, maintain and update the routing tables on each router. With the network increasing, the job of maintaining and updating the routing tables becomes overwhelming, meaning it becomes very difficult to scale up, while the network becomes unmanageable. The impossibility to quickly scale up the network is impeding the business growth of the ISP.

1.3. Aims and Objectives

The main goal of this thesis is to introduce the advantages of dynamic routing protocols, specifically OSPF, for large networks, such as the case of ISP. Due to very complex internal nature of OSPF, the thesis aims to give a detailed account of OSPF as interior routing protocol, explaining its operation in single and multiple areas and stating the benefits that OSPF implementation brings for large networks. The concepts, which this thesis tries to describe, are:

- OSPF fundamentals,
- OSPF features,
- OSPF operation in a single area,
- OSPF network topologies,
- The purpose of using OSPF in a multiple area network,
- The features of multiple area OSPF,
- The operation of OSPF across multiple areas,
- Design considerations for multiple area OSPF.

Lastly, the thesis proposes the topology for solving ISP's problems.

1.4. Thesis Outline

This thesis consists of 8 chapters. Chapter 1 describes the background of the thesis. Chapter 2 explains the basic logic behind link-state protocols and gives a general overview for OSPF. Chapter 3 describes in detail the features and operation of OSPF in single area. Chapter 4 explains in detail the features and operation of OSPF in multiple areas. Chapter 5 gives the proposed topology for implementing the OSPF based solution, stating the advantages of the solution. Chapter 6 concludes the entire thesis with the summary. References are given in Chapter 7, while OSPF terminology is given as annex in Chapter 8.

2. LINK-STATE ROUTING PROTOCOL FEATURES

Link-State Routing protocols, such as OSPF, were created to overcome the limitations of early distance vector protocols. Instead of sending the entire routing table periodically via broadcasts, as it was done by the original distance vector protocols, for example RIPv1, the link-state routing protocols use multicast addressing and incremental updates. The full routing update is sent only every 30 minutes, compared to every 30 seconds for the early distance vector protocols.

A link refers to the connection between routers, that is, the physical connection or medium between the routers, over which a logical link is formed. A link-state routing protocol is therefore a protocol that sends information about the links between routers, when there is a change in the state of one of those links [3]. In case that there is a failure of connection between two routers, those routers propagate an update in order to inform the entire network that the link between them is in the down.

While the distance vector protocols propagate the routes, the link-state routing protocols propagate the information concerning only the local links connected to the router. These links are propagated to every other router in the network. In this way, every router will have the same image of the network, which was created from the original updates from every other router in the network.

The fact that the link-state routing protocols send updates only about links is what makes them efficient protocols. This is because one link might affect many routes. Off course, this means that the router itself must calculate the routes and this computation uses o a lot of router CPU, but the on the other hand the network bandwidth is preserved.

The link-state routing protocols develop and maintain the neighbor relationship with routers on the same link. This is done by sending a simple hello message. The exchange is connection-oriented. After the routers have synchronized their routing tables by exchanging routing updates, they are adjacent neighbors. For OSPF, this process will be explained in detail in the next section.

As long as the Hello protocol is received, the neighbor relationship is maintained. Routers must have the same subnet mask and hello timers, for this to function. The continuous neighbor relationship enables quick and efficient exchange of information. Link changes in the network are detected very quickly, for example, whether the neighbor is dead, because Hello protocol messages are no longer received from the neighbor. When a problem is identified, the routing protocol immediately sends out a message, without waiting for the update timer to expire. This is called a triggered update and it is an incremental update because it contains only the network change. The incremental update improves convergence time and also reduces the amount of information that needs to be sent across the network. The network overhead on the physical media is eased, allowing more bandwidth for data [3]. Requiring fewer network resources makes link-state routing protocols very convenient to be used in larger networks.

Link-state routing protocols attempt to reduce network overhead by:

- Using multicast addressing

- Sending triggered updates
- Sending network summaries infrequently, if at all
- Using small packets from every router to describe their local connectivity, instead of the entire routing table [3]

The network map of every link encountered by the link-state routing protocol is called the topology database and it is identical for every router in the network. From the topology database the routing table is calculated. The link-state routing protocols use the Dijkstra algorithm, or Shortest Path First algorithm, to select the best path to a destination, based on the metric called cost. Cost is protocol specific.

2.1. OSPF overview

OSPF was created to overcome the limitations of early distance vector protocols, such as RIPv1. The improvements introduced by OSPF, as a link-state routing protocol, compared to distance vector routing protocols, such as RIPv1, are:

- Using bandwidth far more efficiently by sending only incremental updates, while requiring more resources from the router (greater memory and CPU to calculate the Dijkstra algorithm).
- The updates are not broadcast as in RIPv1, but directed to multicast addresses 224.0.0.5 and 224.0.0.6.
- Propagating changes in the network more quickly through neighbor relationships and incremental updates.
- Not being limited in size by a maximum hop count of 15.
- By using VLSM, it allows for variation in network design throughout the organization.
- It has security options.
- The metric can be defined manually, offering greater sophistication in the path determination.
- Better responsiveness to network changes.
- Flexibility in network addressing and design, resulting in allowing the network to scale [5].

Especially important is the last point that OSPF is designed to offer the greatest flexibility in network design. On the other hand, as an open standard, OSPF is required to offer interoperability while allowing the network to grow. These requirements make OSPF a highly complex routing protocol.

The main characteristics of OSPF include the following:

- Maintaining adjacent neighbors.
- Using hello timers to maintain adjacencies. The neighbor is declared dead, if nothing has been heard from him within four times the hello timer. In such cases the generation of an LSA is required.
- Sending the minimum amount of information with incremental updates, when there is a change in the network. This results in fast network convergence. If have been no updates within 30 minutes, a compressed update is sent.
- It divides the network into areas, adding another level of hierarchy to the IP address.
- It’s a classless routing protocol.
- It uses VLSM and both manual and automatic summarization at the IANA class boundary.

- It uses cost as the metric. The default defined by Cisco is the inverse bandwidth, calculated with the formula 10^8/bandwidth (in bps).
- It assigns specific functionality to different routers to streamline the process of communication change in the network.
- It operates as an interior routing protocol within an organization [3].

2.1.1. OSPF Convergence

Different routing protocols use different methods of updating the routing table. These methods affect the time of convergence of the routing tables. For OSPF convergence, the following steps occur:

1. When a link failure is detected by a router, it sends an LSA to its neighbors. If the router is on a multiaccess link, the update will be sent only to the designated router (DR) and the backup designated router (BDR).
2. The path is removed from the originating router's tables.
3. On receipt of the LSA, all routers update the topology table and flood the LSA out their interfaces.
4. The Dijkstra algorithm will be run to rebuild the routing table.

For OSPF, convergence is detection time, plus LSA flooding, plus 5 seconds before computing the topology table. This amounts to a few seconds [3].

3. OSPF IN A SINGLE AREA

As mentioned in the introduction, OSPF is an open standard built by a standards committee. This means that it is not vendor specific and it makes OSPF an obvious option for connecting various technologies and vendor solutions. As a routing protocol, OSPF's function is to convey routing information to every router within the organizational network. The link-state technology, on which OSPF is based, was designed to be very efficient in the way it propagates updates. This allows the organizational network to grow or scale.

3.1. OSPF Features

The most important features of OSPF are briefly explained in the following section in the context of the simplest OSPF network design, that of a single area. The terminology related to OSPF is given as Annex A, at the end of this document.

3.1.1. OSPF Neighbors

The router that shares the same network link or the same physical segment is called a neighbor in OSPF. A router running OSPF discovers its neighbors by sending and receiving a simple protocol called the Hello protocol. The small hello packet is sent periodically, for example on broadcast multiaccess media the default is every 10 seconds. The hello packet has a source address of the router and a multicast destination address set to AllSPFRouters (224.0.0.5). The other routers, which are also running OSPF, listen to the protocol and send their own hello packets periodically.

3.1.2. Adjacent OSPF Neighbors

After neighbor relationship has been set up by the Hello protocol, the neighbors exchange routing updates. Via updates the information about the network is gathered and it is entered into a database, called the topology table. From the topology table, the best paths to destinations are calculated and entered into the routing table.

When the topology databases of the neighbors are identical, it is said that they are synchronized and that the neighbors are fully adjacent. The Hello protocol will continue to transmit, in order to ensure that the link is maintained and that the topology databases are updated and accurate. A router and its networks will exist in the topology database for as long as the other routers receive the Hello protocol from him.

Hence, having a neighbor relationship is a mechanism for determining that a router has gone down, because its neighbor no longer sends hello packets. The other advantage is the resulting streamlined communication, because after the topological databases are synchronized, only incremental updates will be sent to the neighbors as soon as a change is perceived. Consequently, a much faster convergence of the network can be achieved.

3.1.3. The Designated Router (DR)

When routers are connected to a broadcast segment, one router on the segment can be assigned the duty of maintaining adjacencies with all the routers on that segment. If this is the case, this router will be the designated router (DR). DR is elected by the use of the Hello protocol. The election is determined by either the priority of the router or the

highest IP address. The priority of the router can be set between 0–255, where the higher priority increases the likelihood that the router will be selected as the DR. After the election of DR, all other routers will peer with the DR.

Designated routers are created on multiaccess links, because if there are many routers on the same segment, the intermesh of neighbor relationships becomes complex. Mathematically speaking, the number of adjacencies required for a full mesh is n(n-1)/2 and for a DR/BDR situation is 2n-2 [3].

Even though the Hello protocol is not networking intensive, maintaining the relationships requires additional CPU cycles. Also, there is a sharp increase in the number of LSAs (Link State Advertisements – explained in next section of the document) generated. Therefore, if one router is DR, responsible for maintaining adjacencies and forwarding updates, this dramatically reduces the overhead on the network.

3.1.3.1. Backup Designated Router (BDR)

The centralized responsibility of the DR creates the situation of a single point of failure. Therefore redundancy is created into the network with the backup designated router (BDR). The BDR knows all the links for the segment. All routers have an adjacency not only with the DR, but also with the BDR, which in turn has an adjacency with the DR. If the DR fails, the BDR immediately becomes the new DR.

3.1.3.2. Electing the DR and BDR

The DR and BDR can be manually elected or the Hello protocol can select them dynamically.

3.1.3.2.1. Dynamic Election of the DR

The dynamic selection is made on the basis of the highest router ID or IP address present on the network segment. After the DR and BDR have been elected, all routers on the broadcast medium will communicate directly with the DR. They will use the multicast address to all DRs. The BDR will listen but will not respond. The DR will send out multicast messages if it receives any information pertinent to the connected routers for which it is responsible.

3.1.3.2.2. Manual Configuration of the DR

The priority of the router should be set in order to determine manually which router will be the DR. The router interface can have a priority of 0 to 255, where the value of 0 means that the router cannot be a DR or BDR. On the other hand, the higher the priority, the chances are more favorable for winning the election. When there are few routers on the segment with the same priority level, the election process picks the router with the highest router ID. The default priority on a Cisco router is 1 [5].

3.1.3.2.3. Steps for electing the DR

The process for electing the DR and BDR is as follows:

All the neighbors, which have a priority greater than 0 are listed.

1. The neighbor with the highest priority is elected as the BDR.
2. If there is no DR, the BDR is promoted as DR.
3. From the remaining routers, the router with the highest priority is elected as the BDR.

4. If the priority has not been configured, there will be a tie, because the default is to set the priority to 1.
5. If there is a tie because the priority has not been configured, the highest router ID is used [3].

3.1.4. The Hello Packet

The hello packet is transmitted by routers running OSPF to establish the neighbor relations, but it also serves other functions. The various fields in the hello packet have specific responsibilities. Figure 1 shows the format of the hello packet for OSPF version 2, while figure 2 shows the format of the hello packet for OSPF version 3 (OSPF for IPv6). Table 1 describes each field.

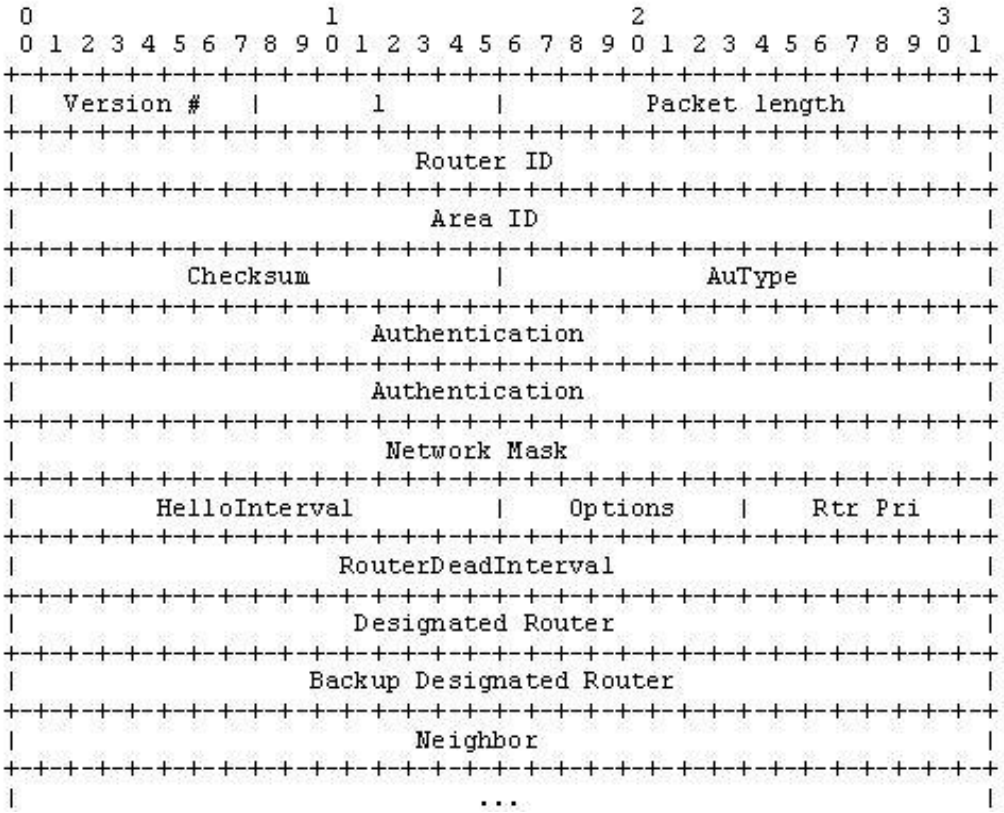

Figure 1: Hello packet for OSPF v2 [1]

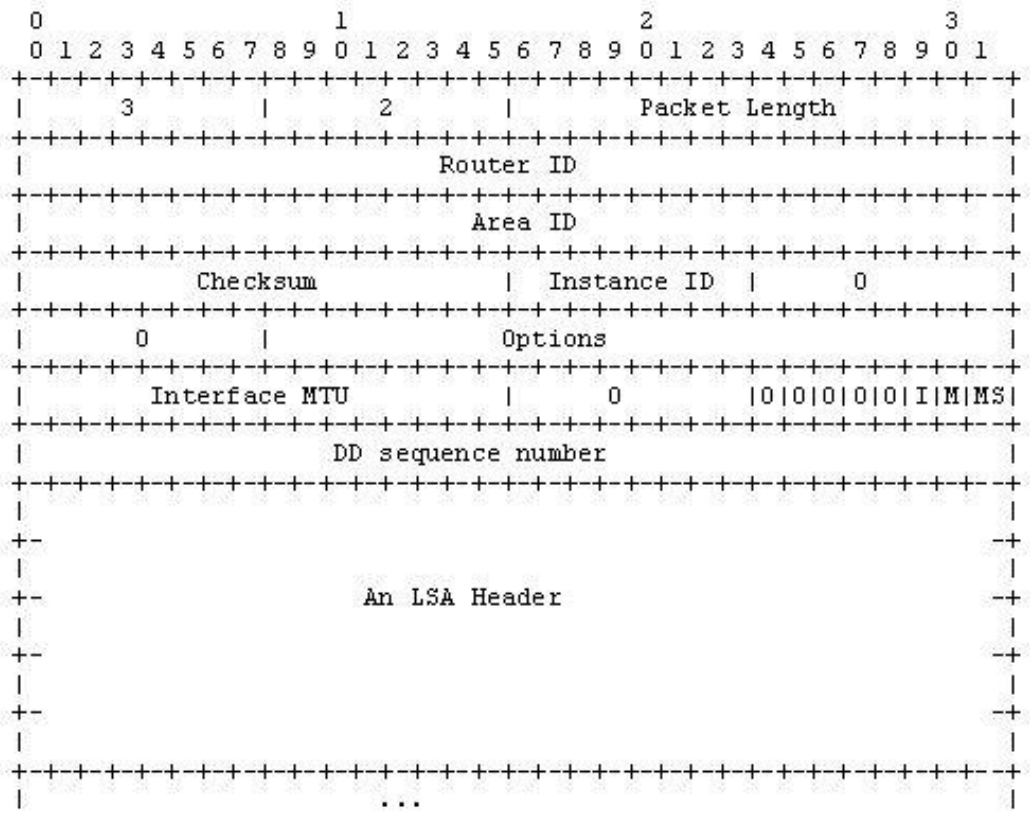

Figure 2: Hello packet for OSPF v3 [4]

Field	Characteristics	Function
Common OSPF Header		
Version #	The version of OSPF, for example version 2	To ensure the versions of OSPF are compatible
Packet Type=1	This states the type of OSPF packet after the header.	The Type 1 header is the hello packet.
Packet Length	This is the length of the packet including the OSPF header.	This field is used to identify the packet length.
Router ID	This is a 32-bit number. The highest IP address on the router is used as the ID. If a loopback address is configured, this will be used, even if it is not the highest address. If there are multiple loopbacks, the highest address is chosen.	This field identifies the router within the autonomous system. It is the ID of the originating router.
Area ID	This is the area ID of the originating router's interface.	The hello packet must come from a router within the same area to be valid.
Checksum	A checksum on the entire OSPF packet excluding the authentication field.	This is used to ensure the integrity of the packet.
AU Type	States the type of authentication used	Ensures the same authentication is used between systems
Authentication	64 bit authentication	Used for security between systems
Hello Packet Format		
Network Mask	The network mask for the transmitting interface	The mask must match the mask on the receiving interface,

		ensuring that they share a segment and network.
Hello Interval Options RouterDeadInterval	Used on broadcast, multiaccess networks: Dead Interval=40 Hello=10 sec Used on nonbroadcast networks: Dead Interval=120 Hello = 30	Hello maintains the presence of the router in its neighbor's databases. It works like a keepalive. The dead interval is how long the router waits before it determines that a neighbor is unavailable because it has not heard a hello packet within the prescribed time, that is, four times the hello timer.
Neighbor	The router ID of a neighbor is entered in the neighbor table when a two-way (bidirectional) communication is established within the RouterDeadInterval. The communication is established when the router sees itself listed as a neighbor in the hello packet generated by another router on the same physical segment.	A neighbor is another router with which updates will be exchanged to synchronize databases.
Rtr Pri	This is the router priority of the source router interface. The higher the priority, the higher the likelihood of the router being selected as a DR or BDR.	This field is used to select the DR and BDR manually.
Designated Router	This is the address of the existing DR.	This field is used to allow the router to create unicast traffic to the DR router.
Backup Designated Router	This is the address of the existing BDR.	This field is used to allow the router to create unicast traffic to the BDR router.
Authentication	This specifies the authentication type and information. If set, the password must match the password stated on the router.	This field is used as security.

Table 1: The Hello Packet fields for OSPF v2 [3].

3.2. OSPF Operation in a Single Area

OSPF operates as a typical link-state routing protocol, using the topology database and the SPF tree as the basis of the SPF algorithm (Shortest Path First algorithm). The SPF algorithm, created by Edsger Wybe Dijkstra [3], creates an SPF tree from the topology database. After calculating the algorithm on the SPF tree, the forwarding table is created. The forwarding table is the routing table.

3.2.1. OSPF Routing Table – Creation and Maintenance

After a neighbor is discovered in OSPF, an adjacency is formed. The process of creating the neighbor adjacency should be explained in the context of the messages that the routers receive.

There are two situations on how the routing table is built. Either an established database has to adjust to a change in the network, or a new router has to create the topology and forwarding databases when it enters the network. The difference between these two situations is in fact that when a new router connects to a network, it will find a neighbor using the Hello protocol and will exchange the routing information. While when a change occurs in an existing network, the router that sees the change will flood the area with the new routing information.

However these events must occur in the stated order, because the new router must learn the network topology, but its addition is a change for the rest of the network.

3.2.2. Building the Routing Table on a New Router

When a new router is added to the network, it builds a routing table by listening to the established routers with complete routing tables. Since every router within an area will have the same database and it will know every network within the area. The routing table built from database is unique for each router because the decisions depend on the individual router's position within the area, relative to the remote destination network.

When the routing table is built for the first time, five packet types are used:

- Hello protocol – is used to find neighbors and to determine the DR and BDR. The propagation of the Hello protocol continuous and it maintains the transmitting router in the topology database of those that hear the message.
- Database descriptor (DDP) – is used to send summary information to neighbors to synchronize topology databases.
- LSR (Link-State Request) - is a request for more detailed information, which is sent when the router receives a database descriptor that contains new information.
- LSU (Link-State Update) - works as the LSA (Link-State Advertisement) packet issued in response to the request for database information in the LSR packet. (note: the different types of LSA (Link-State Advertisements) will be described when explaining the Multiple Area OSPF).
- LSACK (Link-State Acknowledgement) - acknowledges the LSU [7].

Figure 3 shows the process of new router joining the OSPF network.

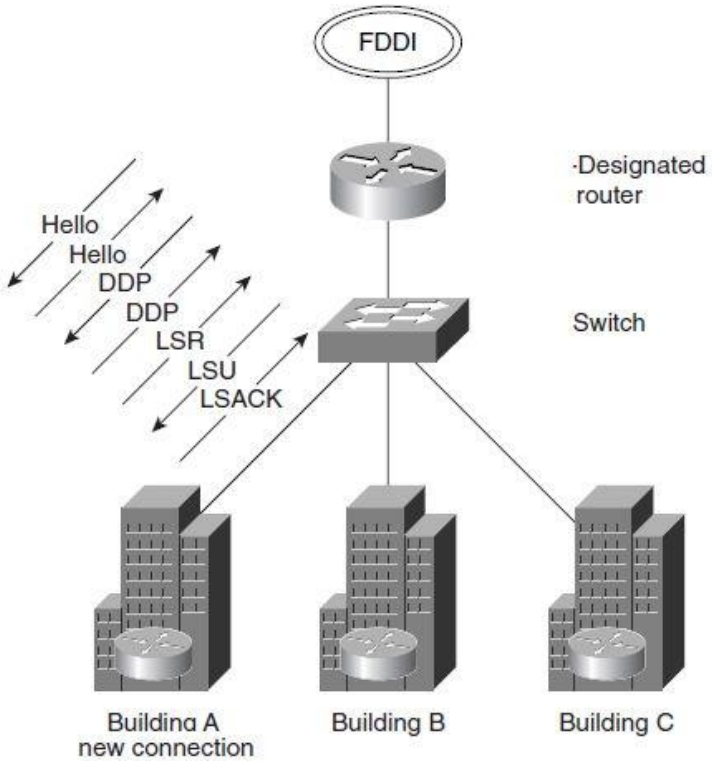

Figure 3: Joining the OSPF network [3]

When it is connected to the network, a router that has been configured to run OSPF will start the process of finding its neighbors and the process of exchange. The new router must learn the network from the systems that are up and running. The entire process goes through seven states. First three states are the creation of the neighbor relationship, while the last four states regard the exchange process. They are the following:

State 1 - The down state - The new router starts in a down state. The new router 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). The DR and BDR are set in the hello to 0.0.0.0.

State 2 - The Init State - The new router waits for a reply. Typically this is four times the length of the hello timer. When the new router receives the first Hello packet, the router enters the Init state. From the hello packet from another router, the new router learns the DR and the BDR. If there is no DR or BDR stated in the incoming hello, an election takes place.

State 3 - The Two-Way State - The received Hello packet includes a list of the sender's known OSPF neighbors. The new router enters the two-way state when it sees itself in a neighbor's Hello. The neighbor relationship is established.

The two-way state is the most basic relationship that OSPF neighbors can have. The routing information is not yet shared between the routers in this relationship. To learn about the link states of other routers and eventually build a routing table, every OSPF router must form at least one adjacency. An adjacency is an advanced relationship between OSPF routers that involves a series of progressive states that rely not just on Hellos, but also on the other four types of OSPF packets. Routers attempting to become

adjacent to one another exchange routing information even before the adjacency is fully established. The first step toward full adjacency is the ExStart state.

State 4 - The ExStart State - Technically, when a router and its neighbor enter the ExStart state, their conversation is characterized as an adjacency, but they have not become fully adjacent. ExStart is established using Type 2 database description packets (DDP). The two neighbor routers use Hello packets to negotiate who is the "master" and who is the "slave" in their relationship and use DBD packets to exchange databases. The router with the highest OSPF router ID "wins" and becomes the master. This designation is not significant - it just determines which router starts the communication. When the neighbors establish their roles as master and slave, they enter the Exchange state and begin sending routing information.

State 5 - The Exchange State - In this state, neighbor routers use Type 2 DDP packets to send each other their link-state information. The routers describe their link-state databases to each other. The routers compare what they learn with their existing link-state databases. 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 DDP will not contain all the necessary information, but just a summary. This summary is enough for the receiving router to determine whether more information is required or whether it already contains that entry in its database. If either of the routers receives information about a link that is not already in its database, the router requests a complete update from its neighbor. Complete routing information is exchanged in the Loading state.

Sate 6 - the Loading State - After the databases have been described to each router, they may request information that is more complete by using Type 3 packets, link-state requests (LSR). When a router receives an LSR, it responds with an update by using a Type 4 link-state update (LSU) packet. These Type 4 LSU packets contain the actual link-state advertisements (LSA). Type 4 LSUs are acknowledged using Type 5 packets, called link-state acknowledgments (LSAck).While the router is awaiting the LSUs from its neighbor it is in the loading state.

State 7 - Full Adjacency - With the Loading state complete, the routers are fully adjacent. Their databases are updated and synchronized. Each router keeps a list of adjacent neighbors, called the adjacency database. The adjacency database should not be confused with the link-state database or the forwarding database [7].

The following diagram illustrates this process (Figure 4).

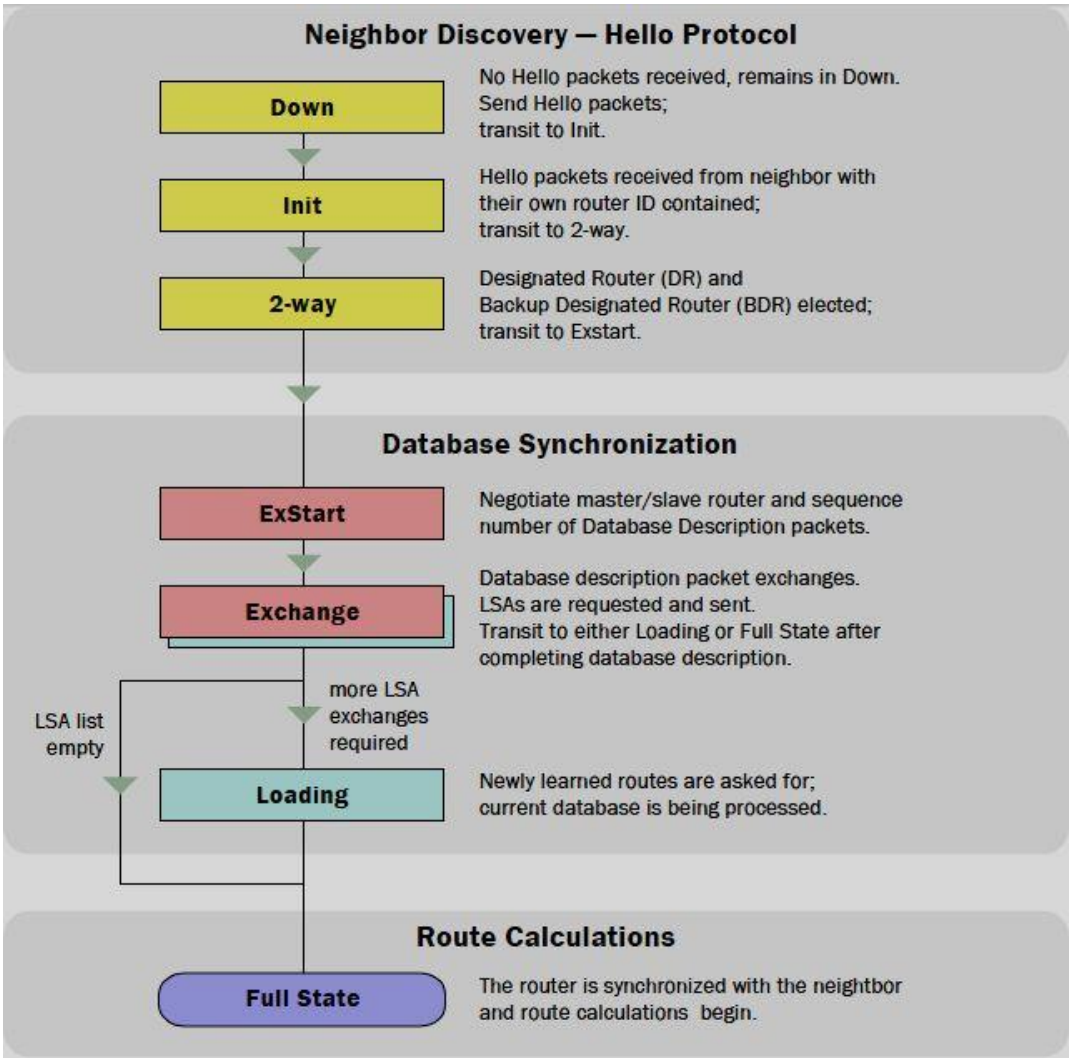

Figure 4: OSPF adjacency process [8]

3.2.3. The Topology Database (the Link-State Database)

The topology database is the router's view of the network within the area. The topology database includes every OSPF router within the area and all the connected networks. The topology database is updated by the LSAs. Within the area, each router has exactly the same topology database. This is a must, otherwise, confusion, routing loops, and loss of connectivity will result. The synchronization of the topology maps is ensured by the complex use of sequence numbers in the LSA headers.

The routing database is constructed from the topology map. The routing database is unique for each router. The router creates the routing database by running the Shortest Path First SPF algorithm, called the Dijkstra algorithm. Each router uses this algorithm to determine the best path to each network by creating an SPF tree. The router will place itself at the top, or in the root of the SPF tree. If there are equal metrics for a remote network, OSPF will include all the paths and will load balance the routed data traffic among them. RFC 2328 does not state the number of multiple equal-cost paths that can be used at the same time. Cisco has defined a maximum of six paths that can be used simultaneously for load balancing [6].

3.2.3.1. Maintaining the Topological Database and the Routing Table

When the router realizes that a change has happened in the network topology, he is responsible for informing the rest of the routers in the area. The reasons for identifying the change in the state of one of its links are the following:

- The router loses the physical or data link layer connectivity on a connected network. On a multiaccess network, the router would propagate an LSU and send it to the DR. On the point-to-point network, the router would propagate an LSU and send it to the adjacent router. From there, in both cases, it is flooded to the network.
- The router fails to hear either an OSPF Hello protocol or a data link Hello protocol. The router propagates an LSU and sends it to the DR on a multiaccess network or the adjacent router in a point-to-point network. From there, it is flooded to the network.
- The router receives an LSA update from an adjacent neighbor, informing it about a change in the network topology. The LSU is acknowledged and flooded out the other OSPF interfaces.

In any of these cases, the router will generate an LSA and flood it to all its neighbors.

3.2.3.2. Addition of new routes

The internal complexity of OSPF can be seen by the process that the router goes through when he receives network LSA update from the DR. The router goes through the following steps:

1. The router takes the first entry from the update - the first network with information about the state of its link.
2. The router verifies that the type of LSA is one that can be accepted by this router.
3. Having determined that it is a valid LSA which it can receive, the router issues a lookup to its topological database.
4. If the LSA entry is not in the topological database, it is instantly flooded out all the OSPF interfaces, except for the receiving interface.
5. If the LSA entry is in the topological database, further questions are required.
6. The router determines whether the new LSA has a more recent (higher) sequence number.
7. If the sequence numbers are the same, the router calculates the checksum for the LSAs and uses the LSA with the higher checksum.
8. If the checksum numbers are the same, the router checks the MaxAge field to decide which update is the most recent one.
9. Having found that the latest LSU is the one that was received, the router determines whether it has arrived outside the wait period, before another computation is allowed (minsLSarrival).
10. If the new LSA entry passes these tests, it is flooded out all the OSPF interfaces, except for the receiving interface.
11. The current copy replaces the old LSA entry. If there was no entry, the current copy is just placed in the database.
12. The received LSA is acknowledged.
13. If the LSA entry was in the database, but the LSA that has just been received has an older sequence number, the router asks whether the information in the database is the same.

14. If the information is the same and the new LSA has an older sequence number, the process discards the packet. There is no inconsistency in the database.
15. If the information is different and the newly received LSA has an older sequence number, however, the receiving router discards the LSA update. It issues a copy of the LSA it has in its database, sending it out of the receiving interface to the source address of the out-of-date LSA. The reason is that the sending router has bad or old information and must be updated because its topological database is obviously not synchronized with the rest of the area. This ensures that any packets that get out of sequence will be verified before action is taken. It also attempts to rectify a problem that it sees - that of multiple routers offering different paths because their topological databases are completely confused.
16. After the initial flood, things calm down, and updates are sent only when there are changes in the area or when the 30-minute timer goes off. This timer ensures that the databases stay synchronized [3].

3.2.3.3. Choosing the Shortest Path First and Building the Routing Table

OSPF selects the shortest, most direct path to the destination. To be able to do this, OSPF needs to examine all the available paths to every network that it has information about. Also, this decision is based on the metric OSPF uses to select the shortest path. After the shortest path is determined, the routing table should be build with the information needed for the forwarding process.

3.2.3.3.1. The Metric

The decision for which path to choose from the routing table is based on the metric used by the routing protocol. For example, RIP uses hop count, which shows how many routers must be passed through to get to the destination. This is for historical reasons, because when CPU and memory speeds were slower, the latency of traveling through the router had much higher implications on network performance. Now days, such constraints do not exist. OSPF chooses the metric of cost. However, cost is not defined within the OSPF specification. It depends on the implementation of the protocol. Cisco's implementation of a dynamic and default cost uses a predefined value based on the bandwidth of the router interface. This default can be manually overridden by the network administrator.

Sometimes, the metric will determine more than one path to the destination. These are known as multiple equal-cost paths.

The cost is applied to the outgoing interface. The routing process will select the lowest accumulated cost of the interfaces to the remote network [3].

3.2.3.3.2. Information Needed in the Routing Table

When the shortest path has been determined, the routing process needs to supply additional information. This information enables the forwarding of the data down the chosen path and includes the next logical hop, link and outgoing interface. This information is required by the routing table. The routing table is also known as the forwarding database.

3.3. OSPF Network Topologies

The way that the OSPF protocol will communicate via the Hello protocol to its neighbors depends on the physical medium being used. Five distinct network types or technologies are identified for OSPF:

- Broadcast multi-access
- Point-to-point
- Point-to-multipoint
- Nonbroadcast multiaccess (NBMA)
- Virtual links [10]

3.3.1. Broadcast Multiaccess Network

Broadcast multiaccess is any LAN network, such as Ethernet, Token Ring, or FDDI. OSPF will send out multicast traffic in this environment. A DR and a BDR will be elected. Figure 5 shows a broadcast multiaccess network, DR and BDR.

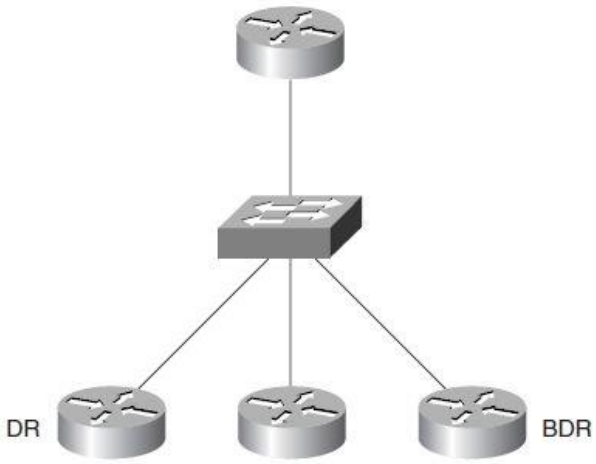

Figure 5: Broadcast Multiaccess Network [3]

3.3.2. Point-to-Point Network

Point-to-point technology is typically used where there is only one other router directly connected to the transmitting or receiving router. The serial line is an example for point-to-point network. There is no need for a DR or BDR in this scenario. OSPF messaging is sent using the multicast address for AllSPFRouters, 224.0.0.5. Figure 6 shows a point-to-point network.

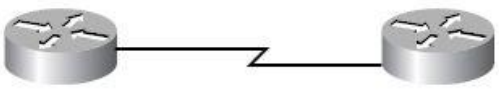

Figure 6: Point-to-Point Network [3]

The following three topologies (point-to-multipoint, nonbroadcast multiaccess and virtual links) go beyond the scope of this thesis. They are mainly designed to support legacy solutions. A brief explanation is given for correctness.

3.3.3. Point-to-Multipoint Network

In the point-to-multipoint network, a single interface connects to multiple destinations. The underlying network treats the network as a series of point-to-point circuits. The

LSA packets are replicated for each circuit and the OSPF traffic is sent as multicast. There is no DR or BDR election. This technology uses one IP subnet for all endpoints on the network [10].

3.3.4. Nonbroadcast Multiaccess (NMBA) Network

As certain point-to-multipoint networks cannot support multicast or broadcast traffic from the physical point of view, NBMA topology, is required. NBMA physically resembles a point-to-point line, but in fact, many destinations are possible. NBMA uses a fully meshed or partially meshed network, which sees as a broadcast network, which will be represented by one IP subnet. NMBA requires manual configuration of the neighbors and the DR and BDR selection, which must be directly connected to their neighbors. All network traffic sent between neighbors will be replicated for each physical circuit using unicast addresses, because multicast and broadcast addresses are not understood [10].

3.3.5. Virtual Links

When the network becomes segmented the virtual link is used. The virtual link is a virtual connection to a remote area that does not have any connections to the backbone (Area 0). OSPF treats this link as a direct, single-hop connection to the backbone area. In fact, it is a virtual connection that tunnels through the network. The OSPF network traffic is sent in unicast datagrams across these links [10].

4. USING OSPF ACROSS MULTIPLE AREAS

One of the main distinctive features between distance vector protocols and the link-state protocols, such as OSPF, are the multiple area networks. While an OSPF area is a logical grouping of routers that are running OSPF with identical topological databases, it is through multiple areas that OSPF prevents a large network from outgrowing its capacity. In this sense, an OSPF area is a subdivision of the greater OSPF domain, called the autonomous system. Partition of the autonomous system into areas allows routers in each area to maintain their own topological databases. This decreases the size of the topological databases, while connectivity between areas and networks outside the autonomous system is achieved with summaries and external links.

4.1. Motivation for OSPF in Multiple Areas

Indications that the single area OSPF is becoming overcrowded are for instance the frequent running of the SPF algorithm. This is because the larger the network, the greater the probability of a network change and consequently a recalculation of the entire area. Also each recalculation will take a longer time to compute. Moreover, the larger the area, the greater the size of the routing table. Although, the entire routing table is not sent out, as with distance vector routing protocols, however, the greater the size of the table, the longer each lookup becomes. This increases the memory requirements on the router.

For the same reasons as above, with the area becoming larger the topological database increases in size and eventually becomes unmanageable. The topology table is exchanged between adjacent routers at least every 30 minutes.

With the various databases increasing in size, the calculations become more frequent, the CPU utilization increases while the available memory decreases. All of this will make the network response time sluggish. This is not because of congestion on the line, but because of congestion within the router itself. It can also cause congestion on the link. These can result in various additional problems, such as loss of connectivity, loss of packets, and system hangs [3].

In all these cases the system would benefit if split into multiple areas.

4.2. The Features of Multiple Areas OSPF

The design issues for the different areas, the underlying technology and the communication, both within and between the areas, are considered in the following section.

4.2.1. Importance of Area Boundaries

One of the main strong points of OSPF is its capability to scale and by scaling to support large networks. This is done by creating areas from groups of subnets. The area is seen internally as if it was an entity of its own. An area communicates with the other areas and it exchanges routing information. However, this exchange is kept to a minimum, allowing only the minimum that is required for connectivity. All computation is kept within the area and therefore a router is not overwhelmed by the entire autonomous system. This is vital, because the nature of a link-state routing protocol is more CPU-intensive and more memory-intensive [3].

4.2.2. Router Types

Depending on its position and function within the OSPF hierarchical design, routers have different responsibilities. There are routers operating within an area, routers connecting areas, and routers connecting the organization or autonomous system to the outside world.

Different OSPF routers are:

- Internal router - All interfaces on the internal router are within the same area. As it operates 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. The internal router and the Autonomous System Boundary Router (ASBR) are the only routers that can operate in a single area OSPF network.
- Backbone router - One of the design rules for OSPF requires that all the areas should be connected through a single area, called the backbone area, Area 0, or 0.0.0.0. A router within this area is referred to as a backbone router. The backbone router can also be an internal router, an ASBR, or an Area Border Router (ABR).
- Area Border Router (ABR) - is responsible for connecting two or more areas. ABR needs to have a full topological database for each area to which it is connected. It sends LSA updates between the areas. These LSA updates are summary updates of the subnets within an area. In order to minimize the routing overhead on both the network and the routers, summarization should be configured for OSPF at the area border because this is where the LSAs make use of the reduced routing updates.
- Autonomous System Boundary Router (ASBR) - Since OSPF is a interior routing protocol, you need to leave the OSPF domain to be able to connect to the outside world or to any other routing protocol. This responsibility belongs to ASBR. If any other routing protocols are being redistributed to OSPF on a router, the router will become an ASBR. This is because the other routing protocols are outside the OSPF autonomous systems. ASBR should reside in the backbone area, although it can be placed anywhere in the OSPF hierarchical design. The reason for ASBR to reside in the backbone area is because any traffic leaving the OSPF domain is also likely to leave the router's area, so it makes sense to place the ASBR in a central location that all traffic leaving its area must traverse. This router can be configured within a single OSPF area as a connection to the outside world [9].

Figure 7 shows different router types for OSPF. It should be noted that all the routers in the backbone area, or in area 0, are not only performing the function of ABR, or ASBR as labeled, but are also backbone routers.

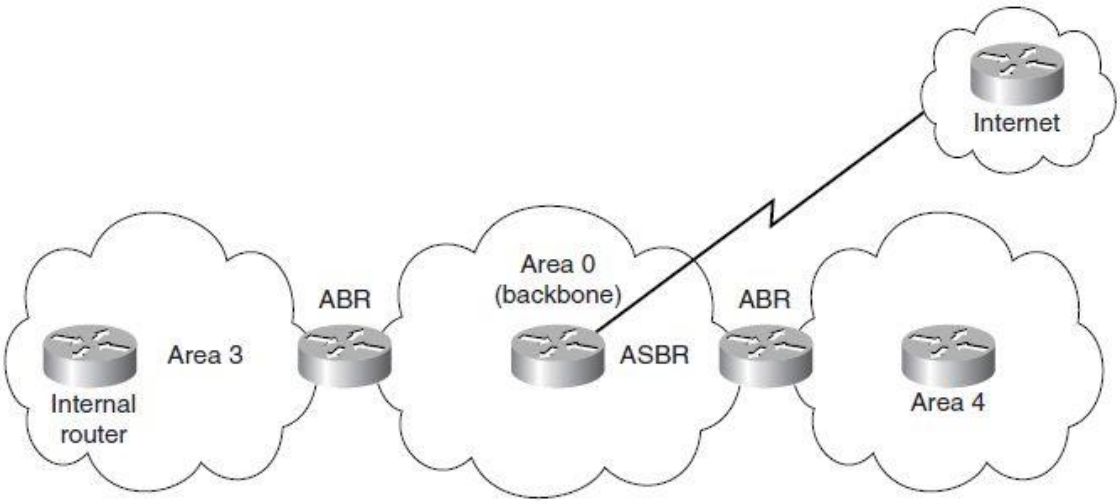

Figure 7: Different router types for OSPF [3]

4.2.3. Link-State Advertisements

There are five types of link-state advertisements (LSAs) defined in the OSPF specification [1]. Cisco has added the sixth type, which is specific only for Cisco equipment. These LSAs are:

- Type 1 - The router link LSA - is generated for each area to which the router belongs. This
- LSA gives the link states to all other routers within an area and it is this LSA that is flooded into an area.
- Type 2 - The network link LSA - is sent by the DR and lists all the routers on the segment for which it is the DR and has a neighbor relationship. The LSA is flooded to the whole area.
- Type 3 - The network summary LSA - is sent between areas and summarizes the IP networks from one area to another. It is generated by an ABR.
- Type 4 - The AS external ASBR summary link LSA is sent to a router that connects to the outside world (ASBR). This LSA is sent from the ABR to the ASBR and it contains the metric cost from the ABR to the ASBR.
- Type 5 - The external link 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.
- Type 7 - The NSSA external LSA - is created by the ASBR residing in a not so stubby area (NSSA). This LSA is similar to an autonomous system external LSA, except that this LSA is contained within the NSSA area and is not propagated into other areas, but it is converted into a Type 5 LSA by the ABR [7].

4.2.4. The Different Types of Areas

OSPF network can be constructed with only one area, a single area. This area is known as the backbone area or Area 0. If more than one area is used, OSPF networks can have several other types of areas. The backbone area must be connected to all the other areas. The different types of OSPF areas are:

- An ordinary or standard area - In this are every router knows about every network in the area, and each router has the same topological database. The

routing tables are unique and depend from the position of the router within the area. The area connects to the backbone and is an entity in itself.

- A stub area - The Type 5 LSAs are blocked for this kind of area, so it will not accept external summary routes. This means that the only way that a router within the stub area can see outside the autonomous system is by the use of a default route. Routers within the stub area can see every network within the area and within other areas.
- A totally stubby area - Type 3, 4 and 5 LSAs are blocked for this area, meaning it will not accept summary LSAs from the other areas or the external summary LSAs from outside the AS. The default route (which is indicated as the network 0.0.0.0.), is the only way out of the totally stubby area. This is a proprietary solution offered by Cisco and it is recommended for remote sites that have few networks and limited connectivity with the rest of the network.
- A not so stubby area (NSSA) - Type 4 and 5 LSAs are blocked for this area. It is used primarily to connect to ISPs, or when redistribution is required. External routes are not propagated into or out of the area. In many ways, NSSA it is the same as the stub area, but which was designed as a special stub area for applications such as an area with its own connection to an Internet resource needed only by a certain division. NSSA is like a stub area but which can receive external routes, but which it will not propagate into the backbone area and thus the rest of the OSPF domain. The Type 7 LSA is created specifically for the NSSA. This LSA can 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.
- The backbone area - called also the Area 0. As mentioned above, it connects all the other areas and it can propagate all the LSAs, except for Type 7 LSA, which is translated into LSA Type 5 by the ABR [3].

The fact that no external routes are allowed, imposes constraints when creating a stub area or a totally stubby area. These constraints are:

- 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 [3].

Figure 8 shows the connectivity and functionality of the different areas. The routing updates and other network information are sent out via LSAs. The function or type of router will determine the LSAs that are sent.

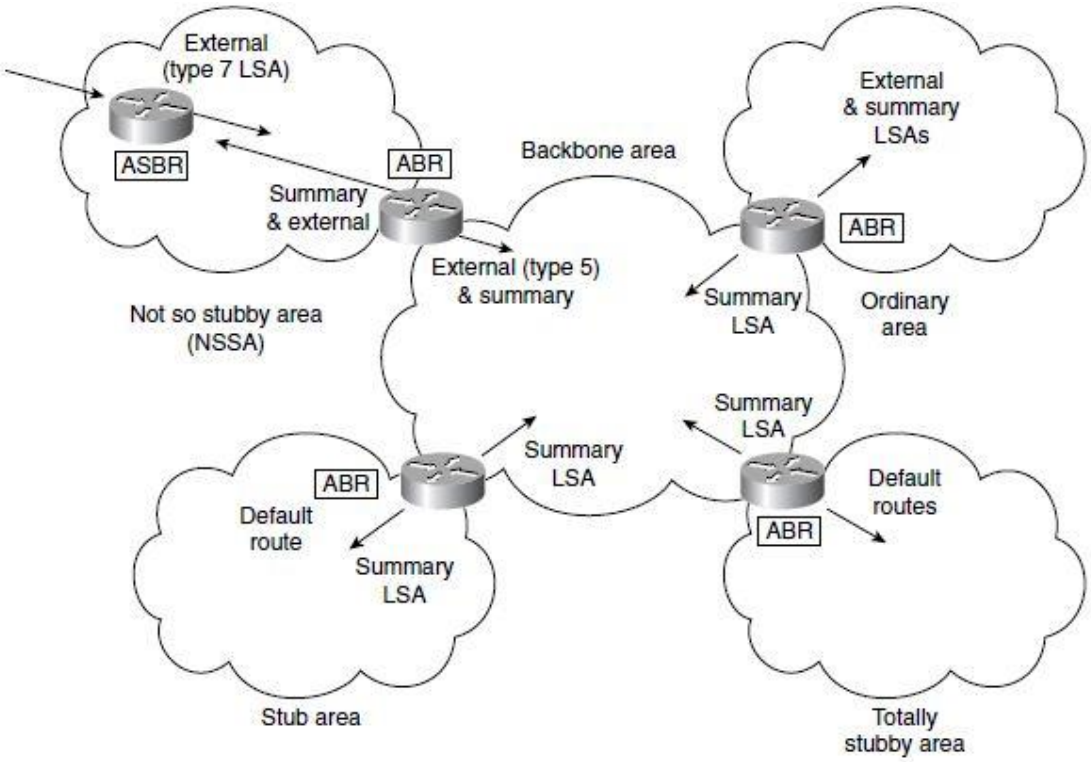

Figure 8: Different types of OSPF Areas and LSA Propagation [3]

4.3. The Operation of OSPF across Multiple Areas

This section deals with the operation of OSPF across the various areas, in order to maintain coherence throughout the autonomous system.

4.3.1. The ABRs and ASBR Propagation of LSAs

LSAs generated within an area are Type 1 or Type 2, and these LSAs are injected as Type 3 summaries by ABR into the backbone. These summaries are then further injected by the other ABRs into their own areas. This is unless those areas are configured as totally stubby areas. This also functions the other way around, where any Type 3 or Type 4 LSA received from the backbone are forwarded into the area by the ABR. The backbone also forwards external routes both ways unless the ABR is a stub router, in which case they are blocked. Summaries received from within the area, cannot be forwarded and summaries received from the backbone cannot be further summarized. The propagation of LSAs within and between areas is shown in Figure 9.

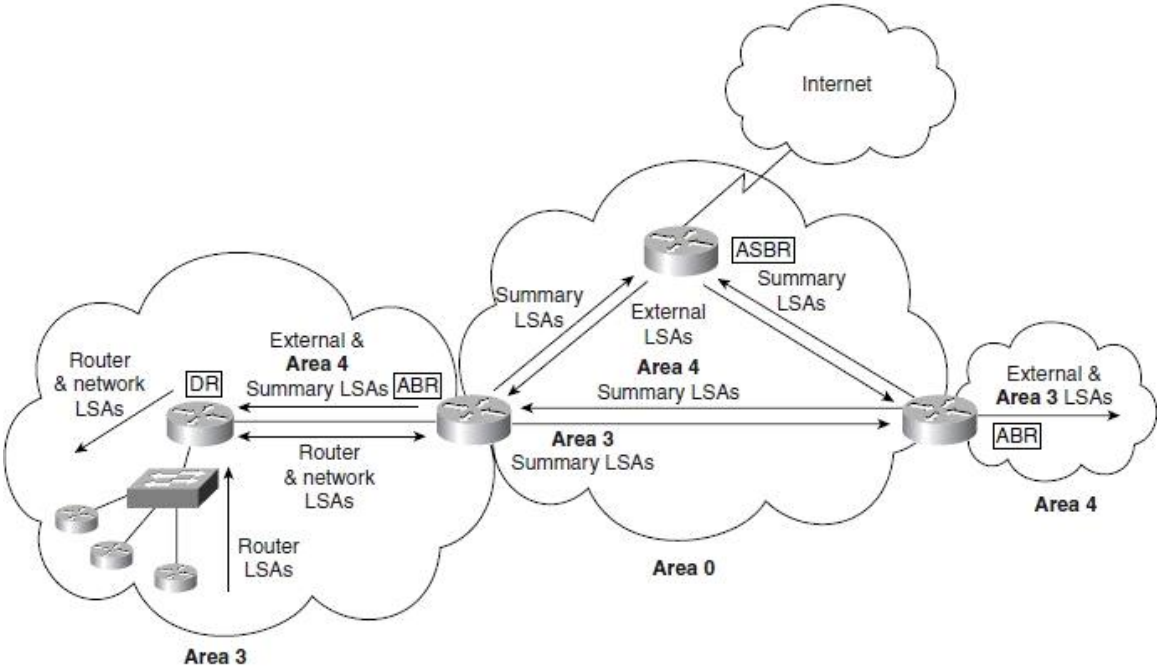

Figure 9: The propagation of LSAs [3]

Couple of conditions must be fulfilled before any LSAs can be flooded out of all interfaces. These conditions are:

- 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) [3].

4.3.2. OSPF path selection between Areas

The routing table depends on the position that the router within in the area and the status of the network. It will also depend from the type of area in which the router is located and whether there are multiple areas in the domain. Lastly, it will depend whether there are communications outside the autonomous system.

When the LSA is received, the router builds the topological database. Then it runs the Dijkstra algorithm, from which the shortest path is chosen and entered into the routing table. Thus, the routing table is the conclusion of the decision-making process. The metric for each link is included. Since 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, LSAs are weighted differently in the decision-making process. The routing table reflects the network topology and indicates where the remote network stands in relation to the local router.

The router will process the LSAs in this order:

1. The internal LSA (Type 1 and 2).
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), that path will be chosen.
3. The external LSAs (Type 5) [3].

4.3.3. Calculating the Cost of a Path to another Area

The path to another area is calculated as the smallest cost to the ABR, added to the smallest cost to the backbone

External routes are routes between a router within the OSPF domain and a router in another autonomous system. The external routes can have two ways of calculating the cost of the path:

- 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 - Only the external cost of the path from the ASBR is taken in the calculation [5].

If both an E1 and an E2 path are offered to the remote network, the E1 path will be used. In the routing table there is a column indicating the source of the routing information, where OSPF includes the LSA type that provided the path. Table 2 shows the codes used in the routing table.

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)	O IA	LSA Type 3 includes the networks or subnets within an area that might 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	O E1 or O E2	The routes in this LSA are external to the autonomous system. They can be

(between autonomous systems)		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.

Table 2: OSPF Routing Table Codes and Associated LSAs [3]

4.4. Design Considerations in Multiple Area OSPF

Dividing the areas is one of the major design considerations in OSPF, because it impacts the addressing scheme for IP within the network. OSPF network requires 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. Since, all the interarea traffic is distributed by the backbone, any reduction of overhead through a solid hierarchical design and summarization is beneficial. The entire network benefits when fewer summary LSAs need to be forwarded into the backbone area [3].

With minimizing the network overhead, the network can grow easily. Therefore, summarization is the natural outcome. Summarization can not be imposed on a network, but it must be part of the initial network design. The addressing scheme must be planed to support the use of summarization.

As with designing any network, the resources available should be considered to make sure that none of these resources are overwhelmed, either initially or in the future. The means, by which the network can grow, are provided in OSPF with the possibility of creating the multiple areas. The network should be designed to run efficiently within the limits of the resources available.

4.4.1. Capacity Planning in OSPF

The first recommendation is not to have more than three areas per router, though it is possible. Having more than three areas will depend on the router's memory and CPU, on network topology and the number of LSAs generated. The second recommendation is not to exceed 50 routers in an OSPF area, since 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. These are only guidelines and not strict rules, because it is not only the number of routers or areas that is important, but the number of routes and the stability of the network.

The general rules 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 [3]

Again it should be stated that these are not strict rules. Elements of the network that use resources, CPU, memory, and bandwidth must be evaluated and provided for when the network is designed. Among factors that influence the number of routers per area is for example the type of area (stub, totally stub, or backbone), because this determines the

number of LSAs and how often and how much CPU and memory each SPF computation requires.

The level of computing power of the routers within the area is also important, as smaller routers are not designed to manage large databases and to run the SPF algorithm continually. The media on which the network is build determines the bandwidth, consequently the higher the bandwidth on the link, the less congestion within the router as it queues the packets for transmission.

Stability of the network is another factor that determines how often LSAs will be propagated because of topology changes and consequently the need for bandwidth, CPU and memory resources. For areas that have external connections, the number of external LSAs is very important. Servicing the external connections with a default link, requires far less memory and CPU than if hundreds of external Internet links are propagated into the network.

Lastly, it is very important to have a hierarchical design with summarization, because the greater the summarization, the smaller and fewer the LSA packets that need to be propagated [5].

4.4.2. Number of Neighbors per Router

If there are many routers on the link, the router that performs the DR function could become overloaded. Thus it is advisable to select the DR manually. The DR should be the router with the most available CPU and memory on the segment. This router should not be selected to be the DR for more than one link.

4.4.3. Number of Areas per ABR

The ABR will have a full topology table for each area that is connected, which could result in overloading the router before it has attempted to compute the best path. The number of areas that a router can support depends on the type of the router and the size of the area. However, when the maintenance of the areas is spread over a few routers it is considered a good hierarchical design. This not only shares the resources, but also builds in a level of redundancy [5].

4.4.4. Summarization

The ability to scale the network is one of the major strengths of OSPF. This can be achieved not only through the creation of multiple areas, which limit the computation and propagation of routing updates, but also through the use of summarization. Two types of summarization are possible in OSPF:

- Interarea summarization - performed at the ABR by creating Type 3 and 4 LSAs.
- External summarization - performed at the ASBR by creating Type 5 LSAs.

The fundamental requirement for both is contiguous addressing [3].

5. IMPLEMENTING OSPF FOR ISP

As mentioned in the problem statement, the Kosovar ISP is using static routing in their network. A static routing table is created, maintained, and updated by a network administrator, manually. A static route to every network must be configured on every router for full connectivity. This provides a granular level of control over routing, but quickly becomes impractical on large networks.

Since routers do not share updates with each other, besides reducing CPU/RAM overhead on the router, more importantly it saves bandwidth by using zero bandwidth. However, static routing is not fault-tolerant, as any change to the routing infrastructure, such as a link going down or a new network added, requires manual intervention. This also means that the administrator must understand the complete internetwork and how each router is connected to configure properly. With the network increasing, manually calculating all the routes and keeping them up-to-date becomes an overwhelming task and very prone to error.

For these reasons, the ISP would greatly benefit from implementing a dynamic routing protocol, such as OSPF, for which a dynamic routing table is created, maintained, and updated by a routing protocol running on the router. In this case, routers do share dynamic routing updates with each other and it increases CPU, RAM, and bandwidth usage. However, OSPF is capable of dynamically choosing a different or a better path when there is a change to the routing infrastructure. More importantly, OSPF allows the network to scale.

Off course these features can be achieved with other dynamic routing protocols as well, but the main reason for choosing OSPF it is because it is an open standard. OSPF offers to ISP the ability to purchase equipment from different vendors, because it supports interoperability, as OSPF is not a proprietary protocol.

5.1. Proposed topology for ISP

The proposed topology starts with dividing the areas. Initially six areas are created for every major city in Kosova. The number of areas can be increased in the future, if needed and when needed. The backbone, or area 0, is situated in Prishtina, where the biggest numbers of clients are, with the intention of minimizing the interarea traffic. The other areas, for each major city in Kosova, are created as standard areas. For all the areas the broadcast multiaccess topology is chosen. The DR and BDR should be elected manually.

The proposed solution is shown in Figure 10.

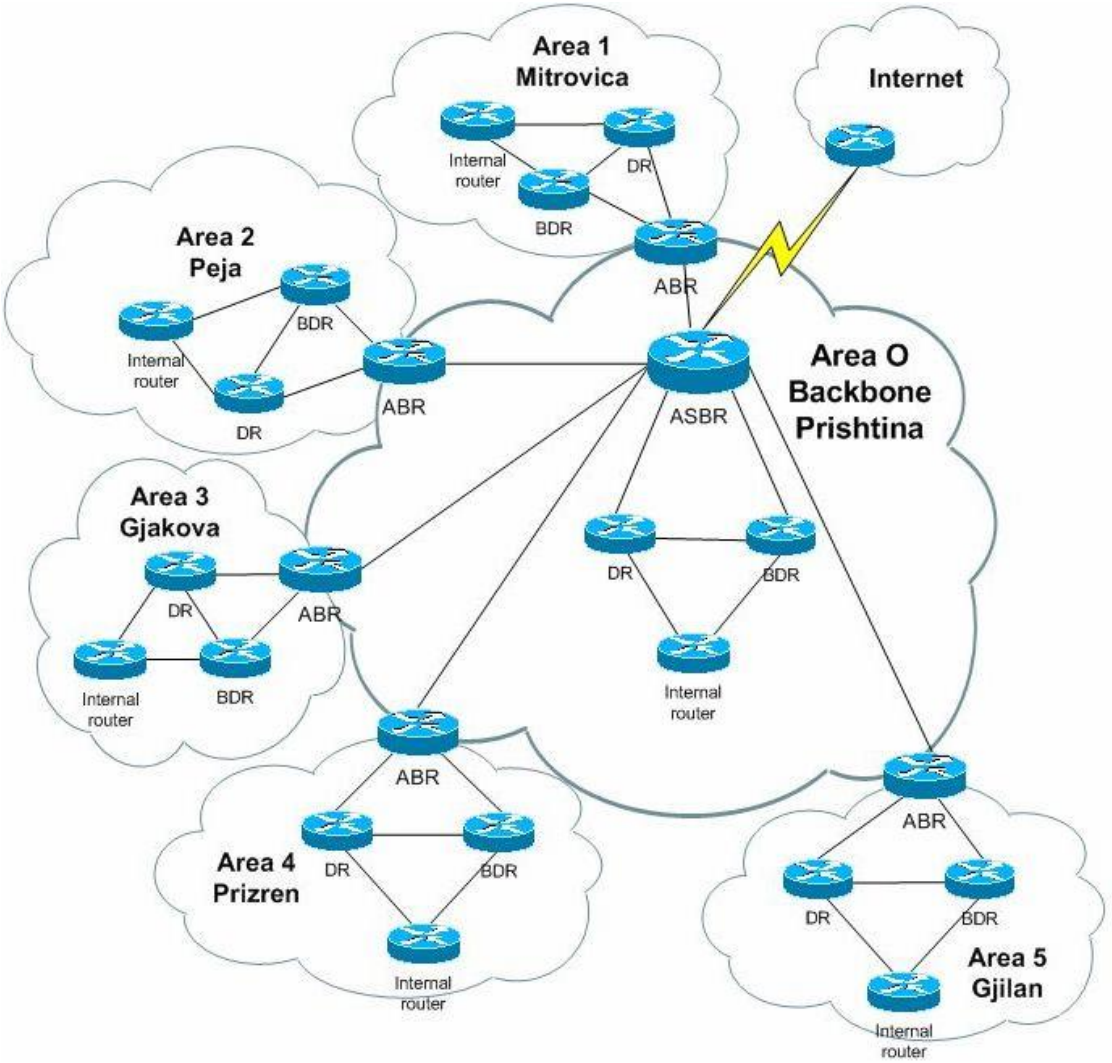

Figure 10: Proposed topology for implementing OSPF for ISP

5.2. Benefits of implementing OSPF

The ISP's benefits from implementing OSPF and particularly the proposed solution include:

- Dynamic determination of loop free routes.
- Fast convergence, if any changes occur in the network it updates fast.
- Minimizing the routes and reducing the size of routing table by configuring multiple areas.
- Multiple routes are supported, thus enabling load balancing.
- Support for variable length subnet masking.
- Ability for the network to scale.
- Not being vendor specific.

6. SUMMARY

OSPF is a scalable routing solution for the ever growing IP networks of today. Its complex and descriptive topology advertisement and the concept of hierarchical routing areas satisfy the demands of the most diversified network designs. Quick convergence and the robustness of link-state database exchanges are the key features of OSPF networks. Also important is OSPF's improved design for network security.

With enhanced features and capabilities of OSPF networks, OSPF offers a number of key benefits:

- It supports cost-based routing metrics to indicate the bandwidth and capability of each individual network link. This gives administrators great flexibility in engineering networks.
- It is able to support multiple paths to a given destination, giving routers a choice in deciding how to balance data forwarding over multiple paths. This increases the capability to implement load balancing and redundancy.
- It supports the hierarchical routing through the concept of distinct areas and multilevel routing, which provides the scalability to accommodate systems with large number of routers.
- It uses Shortest Path First algorithm. This fundamental capability of OSPF allows it to calculate a cost to every destination in the network.
- It maintains a consistent topology table among all routers within an area, it can converge very quickly and is immune to routing loops.

The specification for OSPF is an ongoing work. Couples of drafts, that enhance various aspects of OSPF, are in final stages before being approved as RFC, then to be implemented as practical solutions by different vendors. This leaves a lot work for future designs, implementations, testing and discussions regarding OSPF.

7. REFERENCES

[1] J. Moy, “OSPF Version 2”, Internet Engineering Task Force, RFC 2328, April 1998
(online documentation, last accessed on 09.09.2012)
http://tools.ietf.org/html/rfc2328

[2] Internet Engineering Task Force OSPF Working Group,
(online documentation, last accessed on 09.09.2012)
http://datatracker.ietf.org/wg/ospf/

[3] C. Gough, CCNP Scalable Cisco Internetworks (BSCI) Exam Certification Guide, Third Edition, Cisco Press, 2004, ISBN: 1-58720-085-6.

[4] R. Coltun, et al, “OSPF for IPv6”, Internet Engineering Task Force, RFC 5340, July 2008
(online documentation, last accessed on 09.09.2012)
http://tools.ietf.org/html/rfc5340

[5] Cisco Systems, OSPF Design Guide, August 2005, Cisco Systems
(online documentation, last accessed on 09.09.2012)
http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml

[6] H. Benjamin, CCNP Practical Studies: Routing, Cisco Press, 2002, ISBN: 1-58720-054-6.

[7] Cisco Systems, CCNP 1: Advanced Routing Companion Guide (Cisco Networking Academy Program), 2005, published by Cisco Press.

[8] D. Lee, Open Shortest Path First (OSPF) Conformance and Performance Testing, White Paper, 2004, Ixia.

[9] Juniper Networks, Junos OS OSPF Configuration Guide, November 2011, Juniper Networks, Inc.
(online documentation, last accessed on 09.09.2012)
http://www.juniper.net/techpubs/en_US/junos11.4/information-products/topic-collections/config-guide-routing/gen/config-guide-routing-ospf.pdf

[10] Cisco Systems, IP Routing: OSPF Configuration Guide, Cisco IOS Release 12.4T, 2011, Cisco Systems
(online documentation, last accessed on 09.09.2012)
http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_ospf/configuration/12-4t/iro-12-4t-book.pdf

8. APPENDIX A – OSPF TERMINOLOGY

Term	Description
Adjacency	Formed when two neighboring routers have exchanged information and have the same topology table. The databases are synchronized, and they both see the same networks.
Area	A group of routers that share the same area ID. Each router in the area has the same topology table. Each router in the area is an internal router. The area is defined on an interface basis in the configuration of OSPF.
Autonomous system	Routers that share the same routing protocol within the organization.
Backup designated router (BDR)	The backup to the designated router (DR), in case the DR fails. The BDR performs none of the DR functions while the DR is operating correctly.
Cost	The metric for OSPF. It is not defined in the standard with a value. Cisco use the default of the inverse of bandwidth so that the higher the speed of the link, the lower the cost, and, therefore, the more attractive the path. This default can be overridden by a manual configuration. This should be done only if you have a full knowledge of the network.
Database descriptor	Referred to as DBDs or database descriptor packets (DDPs). These are packets exchanged between neighbors during the exchange state. The DDPs contain partial LSAs, which summarize the links of every router in the neighbor's topology table.
Designated router (DR)	Router responsible for making adjacencies with all neighbors on a multiaccess network, such as Ethernet or FDDI. The DR represents the multiaccess network, in that it ensures that every router on the link has the same topology database.
Dijkstra algorithm	A complex algorithm used by routers running link-state routing protocols to find the shortest path to the destination.
Exchange state	State in which two neighboring routers discover the map of the network. When these routers become adjacent, they must first exchange DDPs to ensure that they have the same topology table.
Exstart state	State in which the neighboring routers determine the sequence number of the DDPs and establish the master/slave relationship.
Flood	A term that refers to network information. When network information is flooded, it is sent to every network device in the domain.
Fully adjacent	When the routing tables of the two neighbors are fully synchronized, with exactly the same view of the network.
Init state	State in which a hello packet has been sent from the router, which is waiting for a reply to establish two-way communication.
Internal router	A router that has all its interfaces in the same area.
Link-state advertisement (LSA)	A packet describing a router's links and the state of those links. There are different types of LSAs to describe the different types of links.
Link-state database	Otherwise known as the topology map, the link-state database has a map of every router, its links, and the state of the links. It also has a map of every network and every path to each network.
Link-state request (LSR)	When the router receives a DDP complete with a partial LSA, it compares the summarized information against the topological database. If either the LSA entry is not present or the entry is older than the DDP, it will request further information.
Link-state update (LSU)	Update sent in response to the LSR. It is the LSA that was requested.

Loading state	State in which, if the receiving router requires more information during the process in which two routers are creating an adjacency, it will request that particular link in more detail using the LSR packet. The LSR will prompt the master router to send the LSU packet. This is the same as an LSA used to flood the network with routing information. While the receiving router is awaiting the LSUs from its neighbor, it is in the loading state.
Neighbor	A router on the same link with whom routing information is exchanged.
Neighbor table	A table built from the hello messages received from the neighbors. The hello message also carries a list of the neighbors.
Priority	A Cisco tool by which the DR can be manually elected or, conversely, prevented from taking part in the DR/BDR election.
Shortest Path First (SPF)	The same as the Dijkstra algorithm, which is the algorithm used to find the shortest path.
SPF tree	A tree of the topological network. It can be drawn after the SPF algorithm has been run. The algorithm prunes the database of alternative paths and creates a loop-free shortest path to all networks. The router is at the root of the network, which is perceived from its perspective.
Topology table	The same as a link-state database. The table contains every link in the wider network.
Two-way state	State during the process in which two routers are creating an adjacency. The new router sees its own router ID in the list of neighbors, and a neighbor relationship is established. This is the stage before routing information is exchanged.

Table 3: OSPF Terminology [3]