


Chapter
5

CEN445
Network Protocols & Algorithms

OSPF

Prepared by
Dr. Mohammed Amer Arafah
Summer 2008




OSPF – The Interior Gateway Routing Protocol


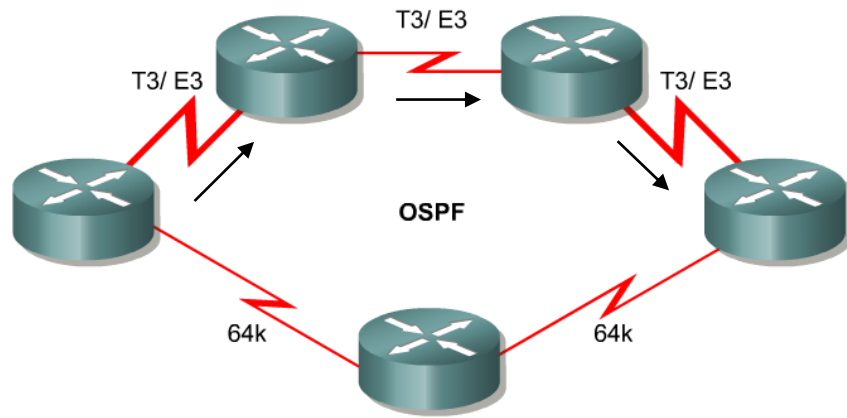
- ✦ A routing algorithm within an AS is called an **Interior Gateway Protocol**, and an algorithm for routing between ASes is called an **Exterior Gateway Protocol**.
- ✦ The original Internet interior gateway protocol was a distance vector algorithm (**DVA**) based on Bellman-Ford algorithm. However, it was replaced by a link state protocol in May 1979. In 1990, the successor, called **OSPF (Open Shortest Path First)** became a standard.


Characteristics of OSPF:

1. *Open*.
2. A *variety of metrics* is supported, including physical distance, delay, and so on.
3. *Dynamic*.
4. Routing based on *type of service* (Differentiating between real-time traffic and others).
5. *Load balancing* (Splitting the load over multiple lines).
6. *Hierarchical systems* are supported
7. *Security* is provided (Preventing from spoofing router by sending them false routing information).



OSPF versus RIP Route Selection






OSPF – The Interior Gateway Routing Protocol

OSPF Network Types:


Point-to-Point



Broadcast Multiaccess



Nonbroadcast Multiaccess (NBMA)





OSPF – The Interior Gateway Routing Protocol

- ✦ OSPF supports **three** types of connections or networks:
 1. *Point-to-Point lines* between exactly two routers.
 2. *Multiaccess networks with broadcasting* (e.g., most LANs).
 3. *Multiaccess networks without broadcasting* (e.g., most packet-switched WAN).

- ✦ A **multiaccess** network is one that can *have multiple routers on it, each pair of routers on such a network is able to communicate directly.*




OSPF – The Interior Gateway Routing Protocol


First: A Simple Autonomous System:

- ❖ OSPF works by abstracting a collection of actual networks, routers, and lines into a *directed graph* in which:
 - ❖ Each arc is assigned a cost (distance, delay, reliability, etc.). Shortest path is computed based on these weights on the arcs.
 - ❖ A connection between two nodes is represented by a pair of arcs, one in each direction.
 - ❖ Each network is represented by a node.
 - ❖ Each router is also represented by a node.


- ❖ OSPF represents actual networks as a directed graph and computes the shortest path from every router to every other router.



OSPF – The Interior Gateway Routing Protocol



(a)



(b)

(a) An autonomous system

(b) A graph representation of (a)



OSPF – The Interior Gateway Routing Protocol


Second: A Large Autonomous System:

- ⊕ Many of the ASes in the Internet are themselves large and nontrivial to manage.
- ⊕ OSPF allows them to be divided up into numbered **areas**, where an area is a network or a set of contiguous networks. Areas do not overlap but need not to be exhaustive, that is, some routers may belong to no area.
- ⊕ Each AS has a **backbone area**, called **area 0**. All areas are connected to the backbone.
- ⊕ *Each router that is connected to two or more areas is part of the backbone.*
- ⊕ Within an area, each router has the *same link state database* and runs the *shortest path algorithm*. Its main job to calculate the shortest path from itself to every other router in the area, including the router that is connected to the backbone, of which there must be at least one.




OSPF – The Interior Gateway Routing Protocol

- ⊕ The way OSPF handles **type of service** routing is to have multiple graphs. One uses **delay** as the metric, one uses **throughput** as the metric, and the other uses the **reliability** as the metric. Although this triples the computation needed, it allows separate routes for optimizing delay, throughput, and reliability.
- ⊕ Three kinds of routes may be needed:
 1. **Intra-area**
 2. **Interarea**
 3. **InterAS.**
- ⊕ **Intra-area** routes are the easiest, since the source router knows the shortest path to the destination router.
- ⊕ **Interarea routing** always proceeds in **three steps**: *go from the source to the backbone; go across the backbone to the destination area; go to the destination.*



OSPF




The relation between ASes, backbones, and areas in OSPF.



OSPF – The Interior Gateway Routing Protocol

- ❖ OSPF distinguishes four classes of routers:
 - ❑ **Internal routers** are wholly within one area.
 - ❑ **Area border routers** connect one or more OSPF areas to the backbone.
 - ❑ **Backbone routers** are on the backbone.
 - ❑ **AS boundary routers** talk to routers in other ASes.


- ❖ Note that all the area border routers are part of the *backbone*. In addition, a router that is in the backbone but not part of any other area is also an *internal router*.



OSPF – The Interior Gateway Routing Protocol

Procedure:

- ⊕ When a router boots, it sends **HELLO** packets on all of its *point-to-point* lines and multicasts them on LANs. On WANs, it needs some configuration information to know *who to contact*. From the responses, each router learns who its neighbors are.
- ⊕ The **Hello** carries information about which *all neighbors must agree to form an adjacency and exchange link-state information*.
- ⊕ On **multi-access networks** the **Designated Router (DR)** and **Backup Designated Router (BDR)** maintain adjacencies with all other OSPF routers on the network.
- ⊕ OSPF works by exchanging information between **adjacent routers**, which is not the same as between neighboring routers. For example, *it is not efficient to have every router on the LAN talk to every other router on the LAN*. To avoid this situation, on router is elected as the **DR**. It is adjacent to all the other routers on it LAN, and exchanges information with them.
- ⊕ A **BDR** is always is *kept up* and *up-to-date* to ease the transition if the **DR** crashes.



OSPF – The Interior Gateway Routing Protocol

Procedure (continued):


- ⊕ **Adjacent routers** go through a *sequence of states*. *Adjacent routers must be in the full state before routing tables are created and traffic routed.*
- ⊕ Each router sends *link-state advertisements (LSA) in link-state update (LSU)* packets. These LSAs *describe all of the routers links*. Each router that receives an LSA from its neighbor records the LSA in the link-state database. This process is repeated for all routers in the OSPF network.
- ⊕ When a router *detects up or down link or its cost changes*, the router sends **LINK STATE UPDATE (LSU)** packets to each of its adjacent routers. Each **LSU** carries a collection of **link state advertisements (LSAs)** *one hop further from its origin*. Several link state advertisements may be included in a single packet. The **LSU** provides the *costs of the links used in the topological database*.
- ⊕ The **flooding** packets are acknowledged, to make them reliable. Each **LSU** has a *sequence number*, so a router can see whether an incoming **LSU** is older or newer than what currently has.
- ⊕ A **LINK STATE ACK (LSAck)** packet is used to acknowledge each **LSU**, to make them reliable.



OSPF – The Interior Gateway Routing Protocol

Procedure (continued):

- ✦ **Database Description packets** announce the link state entries currently held by the sender. By comparing its own values with those of the sender, the receiver can determine who has the most recent values. *This packet is used when a line is brought up.*
- ✦ Either partnet can request link state information from the other one by using **Link State Request (LSR)** packets. The result is that each pair of adjacent routers check to see who has the most recent data, and new information is spread throughout the area.
- ✦ Finally, we can put all the pieces together. Using *flooding*, each router informs all the other routers in its area of its neighbors and costs. This information allows each router to construct a *graph* for its area and compute the shortest path (Dijkstra).
- ✦ Routing information is now maintained. When there is a change in a link-state, routers use a flooding process to notify other routers on the network about the change. The Hello protocol *dead interval* provides a simple mechanism for determining that an adjacent neighbor is down.



OSPF – The Interior Gateway Routing Protocol

Procedure (continued):


- ⊕ The backbone area does this too. In addition, the backbone routers accept information from the *area border routers* in order to *compute the best route from each backbone router to every other router*.
- ⊕ This information is propagated back to the *area border routers*, which advertise it within their areas. Using this information, a route about to *send an interarea packet can select the best exit router to the backbone*.



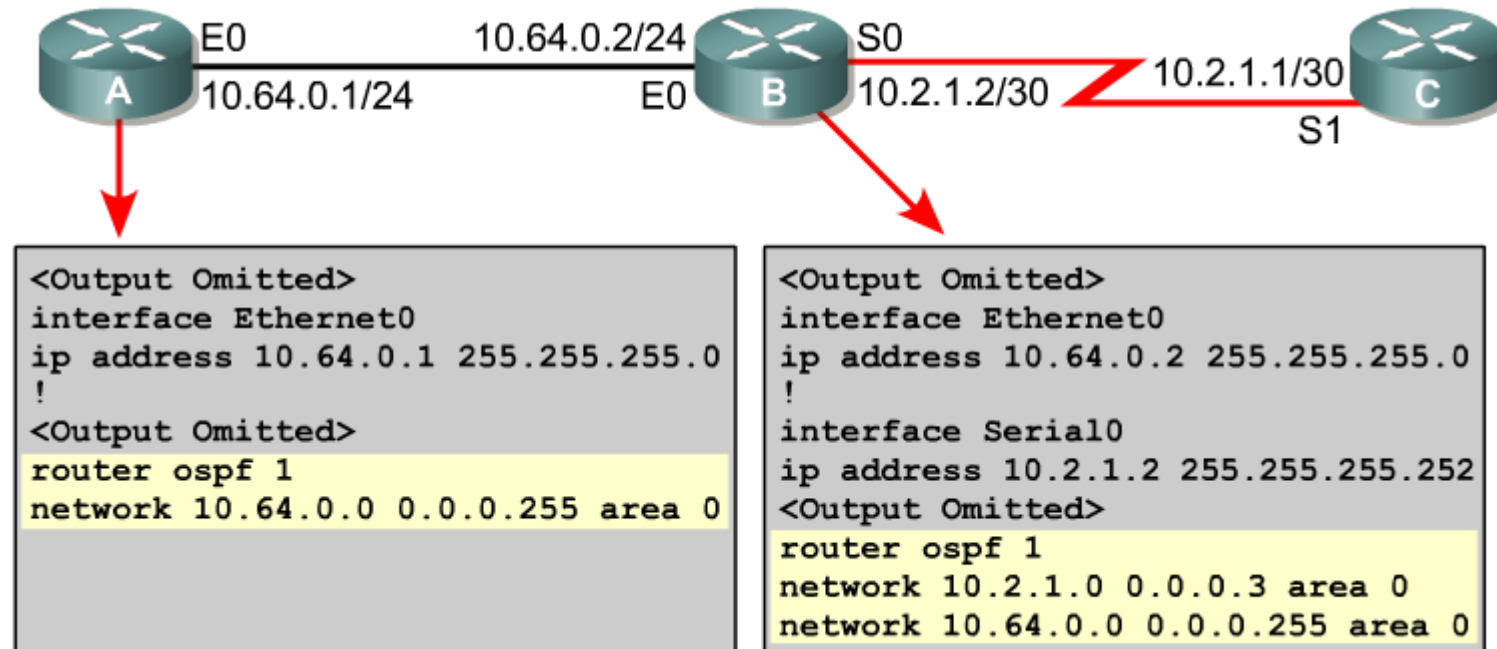
OSPF


Message type	Description
Hello	Used to discover who the neighbors are
Link state update	Provides the sender's costs to its neighbors
Link state ack	Acknowledges link state update
Database description	Announces which updates the sender has
Link state request	Requests information from the partner

The five types of OSPF messages



Basic OSPF Configuration






OSPF – The Interior Gateway Routing Protocol

OSPF Databases:

Database	Description
Adjacencies Database	List all of the neighbor routers to which a router has established bidirectional communication. This is <i>unique</i> for each router.
Link-State Database (Topological Database)	List of information about all other routers in the network . This database shows the network topology. All routers within an area have <i>identical</i> link state databases.
Forwarding Database (Routing Table)	List of routes generated when an algorithm (Dijkstra) is run on the link-state database. Each router's routing table is <i>unique</i> and contains information on how and where to send packets to other routers.




OSPF – The Interior Gateway Routing

Protocol

OSPF Packet Types:

Type	OSPF Packet Type	Description
1	Hello	Establishes and maintains <i>adjacency information with neighbors</i>
2	Database Description Packet	Describes the contents of an OSPF router's <i>link-state database</i>
3	Link State Request	Requests <i>specific pieces</i> of router's <i>link-state database</i>
4	Link State Update (LSU)	Transports <i>link-state</i> advertisements (LSAs) to neighbors routers
5	Link State Acknowledgement (LSACK)	Acknowledgement receipt of a neighbor's <i>LSA</i> .



OSPF – The Interior Gateway Routing Protocol

OSPF States:

- ✦ OSPF neighbor relationships progress through the seven states, one at a time in the following order:
 1. **Down**
 2. **Init**
 3. **Two-way**
 4. **ExStart**
 5. **Exchange**
 6. **Loading**
 7. **Full Adjacency**



OSPF – The Interior Gateway Routing Protocol

1. Down State:

- ⊕ In the Down state, the OSPF process *has not yet exchange information with any neighbor*. OSPF is waiting to enter the Init State.

2. Init State:


- ⊕ OSPF routers send Type 1 (Hello) packets at regular interval (usually 10 seconds) to establish relationship with neighbor routers. *When an interface receives its first Hello packet, the router enters the Init state*, which means the router knows a neighbor is out there and is waiting to take the relationship to the next step.




OSPF – The Interior Gateway Routing Protocol

3. Two-Way State:


- ☉ Hello packets include a list of the sender's known OSPF neighbors. A router enters the two-way state when it sees itself in a neighbor's Hello.



- ☉ The routing information is not shared between routers in two-way state. To learn about other routers' link state, every OSPF router must form at least one *adjacency*. The first step toward full adjacency is the **ExStart State**.



OSPF – The Interior Gateway Routing Protocol

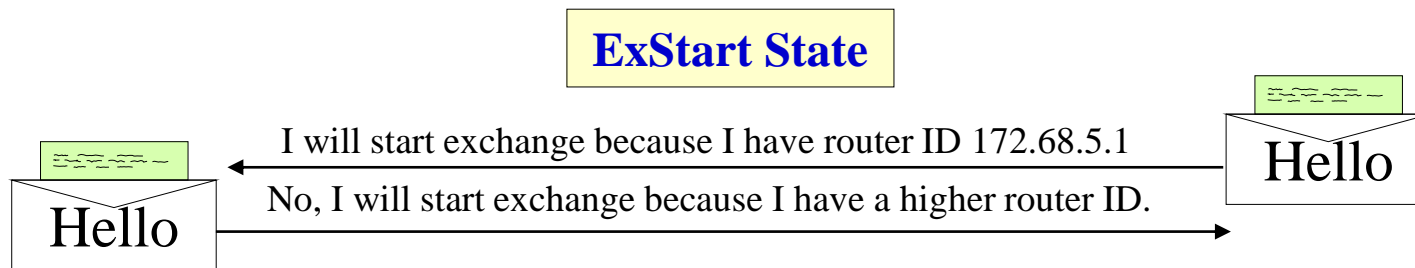




OSPF – The Interior Gateway Routing Protocol

4. ExStart State:

- ✦ The two neighbor routers use Hello packets to negotiate who is the *master* and who is the *slave* in their relationship.
- ✦ The router with the highest OSPF router ID becomes the master.
- ✦ When the neighbors establish their roles as master and slave, they enter the *Exchange state* and begin sending routing information.

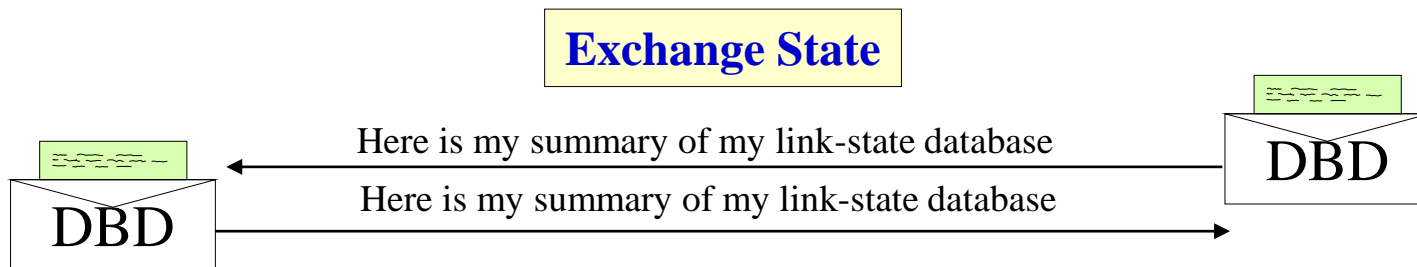





OSPF – The Interior Gateway Routing Protocol

5. Exchange State:

- Neighbor routers use Type 2 *database description (DBD) packets* to send each other their *link-state information*.
- The routers compare what they learn with their existing link-state databases.
- If either router of the router receives information about a link that is not already in its database, the router requests a complete update from its neighbor.

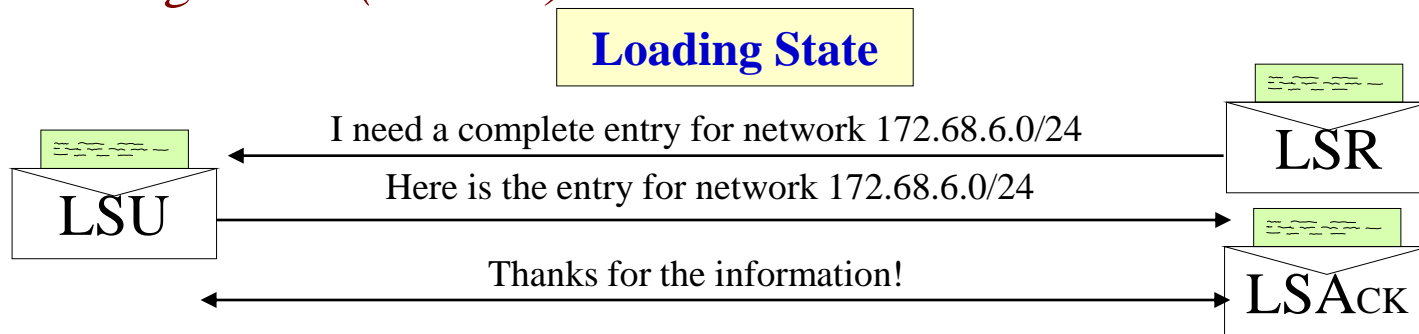




OSPF – The Interior Gateway Routing Protocol

6. Loading State:

- After each router describes its database to the others, the router can request information that is more complete by using Type 3 packet, *Link State Requests (LSRs)*.
- When a router receives an LSR, it responds with an update using Type 4 *Link State Update (LSU) packet*. These Type 4 LSU contain the actual *link-state advertisements (LSAs)*, which is the heart of the link-state routing protocols.
- Type 4 LSUs are acknowledged using Type 5 packets, called *Link-State Acknowledgements (LSAcks)*.






OSPF – The Interior Gateway Routing Protocol


7. Full Adjacency:

- ⊕ With Loading state complete, the routers are fully adjacent.
- ⊕ Each router keeps a list of adjacent neighbors, called the *adjacency database*.

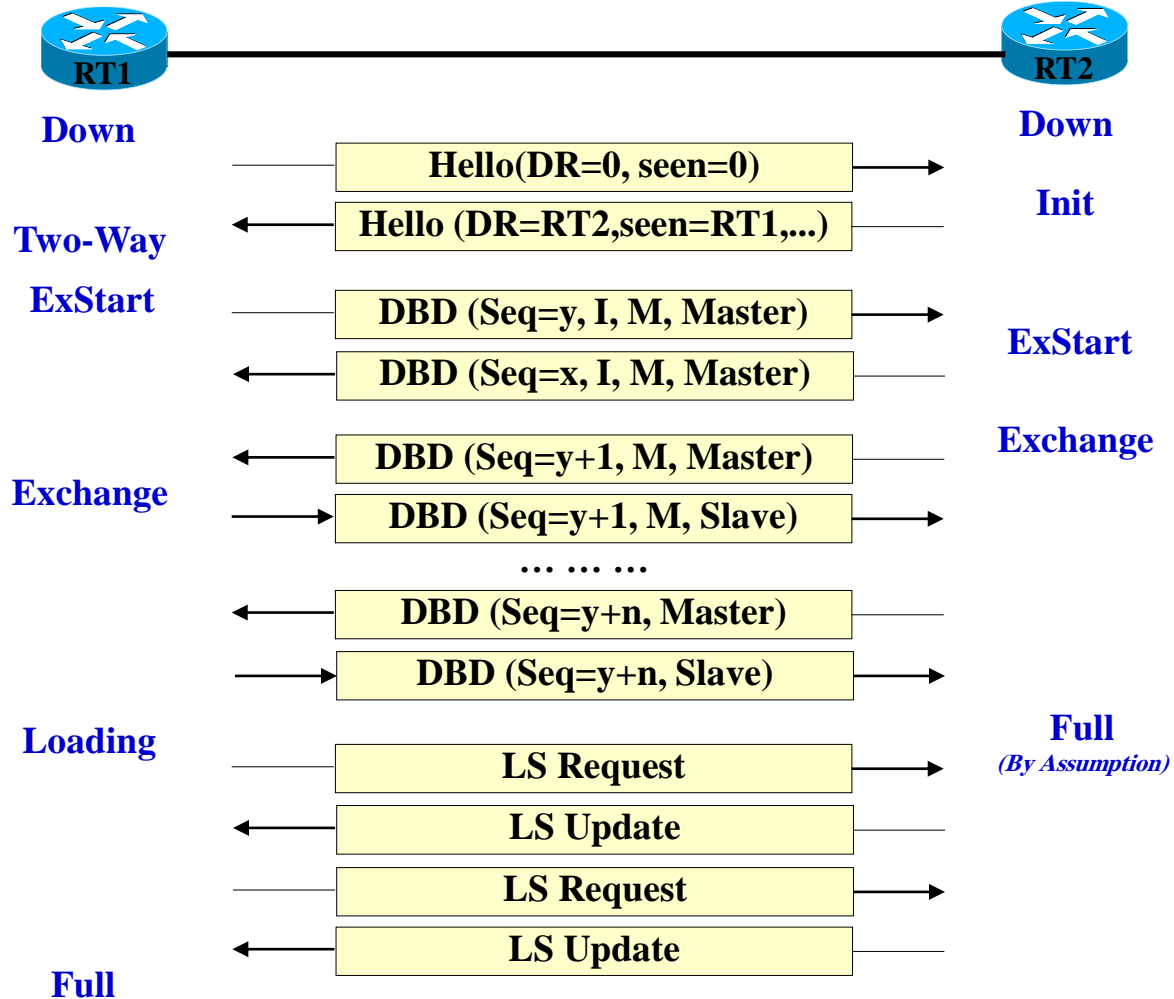


OSPF - Example

- ✦ Routers RT1 and RT2 are both connected to a broadcast network. It is assumed that *RT2 is the Designated Router for the network*, and that *RT2 has a higher Router ID than Router RT1*.
- ✦ Router RT1's interface to the network becomes operational. It begins sending Hello Packets, although it doesn't know the identity of the Designated Router or of any other neighboring routers. Router RT2 hears this hello (moving the neighbor to Init state), and in its *next Hello Packet indicates that it is itself the Designated Router* and that it has heard Hello Packets from RT1. This in turn causes RT1 to go to state ExStart, as it starts to bring up the adjacency.
- ✦ RT1 begins by asserting itself as the master. When it sees that RT2 is indeed the master (because of RT2's higher Router ID), RT1 transitions to slave state and adopts its neighbor's DD sequence number. *Database Description packets are then exchanged, with polls coming from the master (RT2) and responses from the slave (RT1). This sequence of Database Description Packets ends when both the poll and associated response has the M-bit off.*
- ✦ In this example, it is assumed that RT2 has a completely up to date database. In that case, RT2 goes immediately into Full state. RT1 will go into Full state after updating the necessary parts of its database. This is done by sending **Link State Request Packets**, and receiving **Link State Update Packets** in response. Note that, while RT1 has waited until a complete set of Database Description Packets has been received (from RT2) before sending any Link State Request Packets, this need not be the case. RT1 could have interleaved the sending of Link State Request Packets with the reception of Database Description Packets.



OSPF - Example




An adjacency bring-up example



OSPF – The Interior Gateway Routing Protocol

OSPF Hello Protocol:

- ⊕ When a router starts an OSPF routing process on an interface, it sends a Hello packet, and continues to send Hellos at regular intervals.
- ⊕ The Hello packets are addressed to the multicast address 224.0.0.5.
- ⊕ This address is “*all OSPF routers.*”
- ⊕ OSPF routers *use Hello packets to initiate new adjacencies and to ensure that neighbor routers are still functioning.*
- ⊕ Hellos are sent every 10 seconds, by default, on *broadcast multicast* and *point-to-point* networks.
- ⊕ On interfaces that connect to *NBMA networks*, such as Frame Relay, Hellos are sent every 30 seconds.
- ⊕ On multiaccess networks, the Hello protocol elects a *designated router (DR)* and *a backup designated router (BDR).*



OSPF – The Interior Gateway Routing Protocol

OSPF Hello Header:

+	Bits 0–7	8–15	16–18	19–31
0	Version	Type	Packet Length	
32	Router ID			
64	Area ID			
96	Checksum		Authentication Type	
128	Authentication			
160	Authentication			
192	Network Mask			
224	Hello Interval		Options	Router Priority
256	Router Dead Interval			
288	Designated Router			
320	Backup Designated Router			
352	Neighbor ID			
384	...			



OSPF – The Interior Gateway Routing Protocol

Steps in OSPF Operations:

- ⊕ OSPF routers progress through five distinct steps operation:
 1. Establish router adjacencies.
 2. Elect a DR and BDR (if necessary).
 3. Discover routes.
 4. Select the appropriate routes to use.
 5. Maintain routing information.



OSPF – The Interior Gateway Routing Protocol

Step 1: Establishing a Router Adjacency


- ⊕ A router's first step in OSPF operation is to establish router adjacencies.
- ⊕ To become adjacent with another router, RTB sends **Hello** packets, advertising its own router ID.
- ⊕ RTB multicasts Hello packets on both S0 and E0.
- ⊕ RTA and RTC should both receive the **Hello** packets.
- ⊕ RTA and RTC add RTB to the *Neighbor ID field* of their **Hello** packets and enter the **Init state** with RTB.
- ⊕ RTB receives **Hello** packets from both of its neighbors and see its own ID number (10.6.0.1) in the *Neighbor ID field*.
- ⊕ RTB declares a **Two-Way** state between itself and RTA, and a **Two-Way** state between itself and RTC.



OSPF – The Interior Gateway Routing Protocol

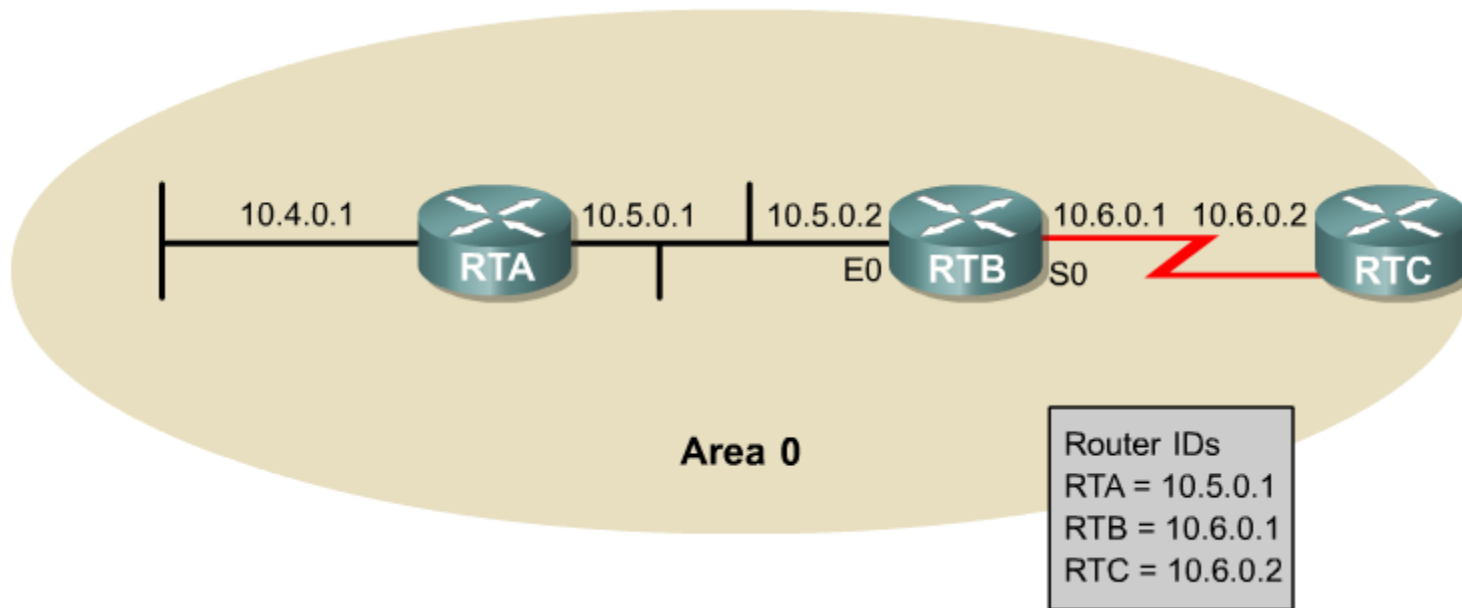
Step 1: Establishing a Router Adjacency

- ⊕ If the network type is *point-to-point*, the router becomes **adjacent** with its link partner.
- ⊕ If the network type is *multiaccess*, RTB enters a election process to become a **DR** or **BDR**.
- ⊕ If an election is necessary, OSPF routers proceeds to *elect a DR or BDR*. If an election is not necessary, OSPF routers enter the **ExStart** state to perform *discovering routes*.




OSPF – The Interior Gateway Routing Protocol

Step 1: Establishing a Router Adjacency



An OSPF router tries to form an adjacency with at least one neighbor for each IP network to which it is connected.



OSPF – The Interior Gateway Routing

Protocol

Step 2: Electing a DR and BDR


- ⊕ The purpose of **DR** is to reduce network traffic by *providing a source for routing updates*, the **DR** maintains a complete topology table of the network and sends the updates to the other routers via multicast.
- ⊕ All routers in an area will *form a slave/master relationship with the DR*. They will form adjacencies with the **DR** and **BDR** only.
- ⊕ Every time a router sends an update, it sends it to the **DR** and **BDR** on the multicast address **224.0.0.6**.
- ⊕ The DR will then send the update out to all other routers in the area, to the multicast address **224.0.0.5**. This way all the routers do not have to constantly update each other, and can rather get all their updates from a single source. The use of multicasting further reduces the network load.
- ⊕ **DRs** and **BDRs** are always setup/elected on *Broadcast networks* (Ethernet). **DR's** can also be elected on *NBMA (Non-Broadcast Multi-Access) networks* such as Frame Relay or ATM. **DRs** or **BDRs** are not elected on *point-to-point* links (such as a point-to-point WAN connection) because the two routers on either sides of the link must become fully adjacent and the bandwidth between them cannot be further optimized.



OSPF – The Interior Gateway Routing Protocol

Step 2: Electing a DR and BDR

- OSPF routers *record information about their neighbors* in the adjacency database. To reduce the number of routing information exchanges among several neighbors on the same network, OSPF routers elects a *designated router (DR)* and a *backup designated router (BDR)* that serve as central points for routing information exchange.
- For example, 5 routers would require 10 adjacency relationships and 10 link states to be sent. In general, for n routers, $n \times (n-1) / 2$ adjacencies would need to be formed.




5 routers would require 10 adjacency relationships and 10 link states



OSPF – The Interior Gateway Routing Protocol


Step 2: Electing a DR and BDR

- ✦ The solution of this overhead is to elect a *designated router (DR)*. This router becomes adjacent to all routers in the broadcast segment. All other routers on the segment send their link state information to the DR. For preceding example, only 5 sets of link state need to be sent.
- ✦ The DR sends link-state information to all other routers on the segment using multicast address for all OSPF routers; **224.0.0.5**.
- ✦ Despite the gain in *efficiency*, the DR represents a *single-point of failure*.
- ✦ A second router is elected as a *backup designated router (BDR)* to take over the DR's duties if it fails.
- ✦ All the routers on the segment send their LSAs to the designated routers (DR and BDR) using the multicast address, **224.0.0.6**.




OSPF – The Interior Gateway Routing Protocol

Step 2: Electing a DR and BDR



DR and BDR receive LSAs



OSPF – The Interior Gateway Routing Protocol

Step 2: Electing a DR and BDR


Network Type	Examples	DR Election
Broadcast Multiaccess	Ethernet, Token Ring, FDDI	YES
Nonbroadcast Multiaccess	Frame Relay, X.25, SMDS	YES
Point-to-Point	PPP, HDLC	NO
Point-to-Multipoint	Configured by the administrator	NO



OSPF – The Interior Gateway Routing Protocol


Step 2: Electing a DR and BDR

- ⊕ OSPF elects a DR to be the central point of all link-state updates and LSAs.
- ⊕ The DR's role is critical, therefore, a BDR is elected to take over smoothly if the DR fails.
- ⊕ The election ballots are Hello packets, which contain a router's *ID* and *priority* fields.
- ⊕ The router with the highest priority value among adjacent neighbors win the election and becomes the DR. The router with the second-highest priority is elected the BDR.
- ⊕ OSPF routers all have the same priority value by default 1.
- ⊕ You can assign a priority from 0 (lowest) to 255 (highest).



OSPF – The Interior Gateway Routing Protocol

Step 2: Electing a DR and BDR




OSPF routers perform DR and BDR elections only on multiaccess IP networks.



OSPF – The Interior Gateway Routing Protocol

Step 3: Discovering Routes

- ⊕ On a *multiaccess* network, the exchange of routing information occurs between the **DR** or **BDR** and every other router on the network.
- ⊕ A *point-to-point* network also engage in the exchange of routing information.
- ⊕ However, *who goes first for the exchange process?*
- ⊕ The purpose of **ExStart** state to establish a *master/slave* relationship between to routers.
- ⊕ The router that announces the highest *router ID* in the Hello packet acts as *master*, while the slave router responds to the master.
- ⊕ After the routers define their roles as master and slave, they enter the Exchange state. The *master leads the slave through the swap of DBDs (Type 2)* that describe *each router's link state data base in limited detail*.



OSPF – The Interior Gateway Routing Protocol

Step 3: Discovering Routes


- ⊕ These descriptions include:
 - ⊞ advertising router's address
 - ⊞ cost of the link
 - ⊞ sequence number
- ⊕ The routers acknowledge the receipt of a DBD by sending an *LSAck (Type 5) packet*.
- ⊕ Each router compares the information that it receives in the DBD with the information it already has.
- ⊕ If the DBD advertises a *new* or *more up-to-date link state*, the router will enter the **Loading state** by sending *LSR (Type 3) packet* about that entry.
- ⊕ In response to the LSR, a router sends the complete link-state information, using *LSU (Type 4) packet*, LSUs carry LSAs.
- ⊕ With the **Loading state** complete, the routers achieve full adjacency. At this point, the neighbor routers should all have *identical* link-state databases.



OSPF – The Interior Gateway Routing Protocol

Step 4: Selecting the Best Routes

- After a router has a *complete link-state database*, it is ready to create its routing table.
- The default cost value is based on the media *bandwidth*. Cost decreases as the speed of the link increases.
- To calculate the lowest cost to destination, a router uses the SPF algorithm.




The link-state database is processed by the shortest path first algorithm and the best routes are selected.



OSPF – The Interior Gateway Routing Protocol

Step 5: Maintaining Route Information

- ⊕ When a change occurs in a link state, OSPF routers use a flooding process to notify other routers on the network about the change.
- ⊕ The Hello protocol's *dead interval* provides a simple mechanism for declaring a link partner down.
- ⊕ If a router does not hear from his neighbor for a time period exceeding the dead interval (usually 40 seconds), the router declares that its link to his neighbor is down.



OSPF – The Interior Gateway Routing Protocol

Step 5: Maintaining Route Information


- ⊕ The router then sends an *LSU packet* containing a new link-state information as follows:
 - ⊞ On a *point-to-point* network, no DR or BDR exists. New link state information is sent to the 224.0.0.5 multicast address.
 - ⊞ On a *multiaccess network*, a DR and BDR exist and maintain adjacencies with all other routers on the network. If a DR or BDR needs to send a link-state update, it sends the update to all OSPF routers at 224.0.0.5. However, non-DR/BDR routers send their LSUs to 224.0.0.6 multicast address to all DR/BDR routers.
- ⊕ When the DR receives and acknowledges the LSU destined to 224.0.0.6, it floods the LSU to all OSPF routers at 224.0.0.5. Each router acknowledges receipt of the LSU with an LSAck.



OSPF – The Interior Gateway Routing Protocol

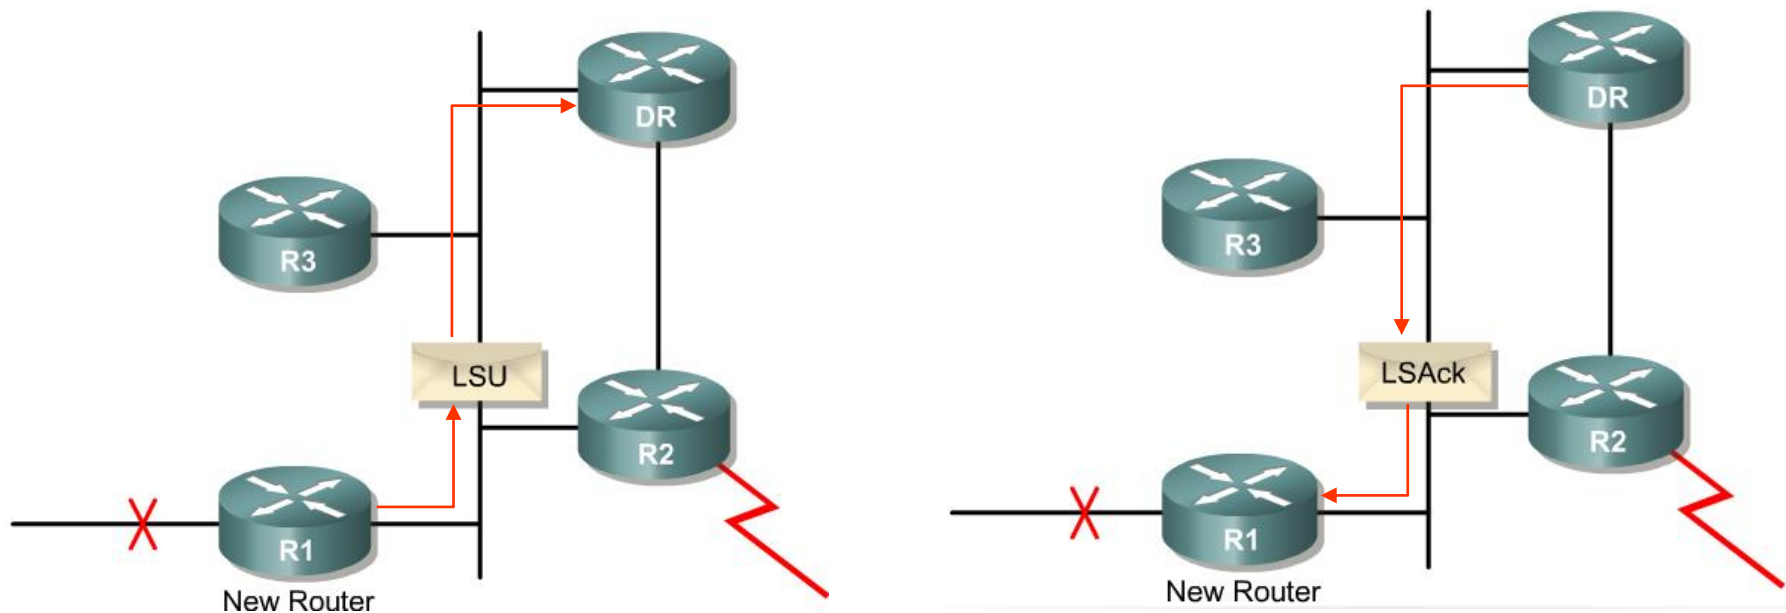
Step 5: Maintaining Route Information

- ✦ If an OSPF router is connected to another network, it floods the LSU to other networks by forwarding the LSU to the *multiaccess network*'s DR, or to the adjacent router if in *point-to-point* network. The DR, in turn, multicasts the LSU to other OSPF router in that network.
- ✦ After receiving an LSU, an OSPF router updates its link-state database. It then runs the SPF algorithm to recalculate the routing table.
- ✦ After the *SPF hold time* expires, the router switches over to the new routing table.
- ✦ If the SPF is calculating a new route, the router will not use that route until after the SPF calculation is complete.
- ✦ Even if a change in link state does not occur, OSPF routing information is periodically refreshed. Each LSA entry has its own *age timer*. The default timer value is 30 minutes. After an LSA entry ages out, the router that originate that entry sends an LSU to the network to verify that the link is still active.




OSPF – The Interior Gateway Routing Protocol

Step 5: Maintaining Route Information

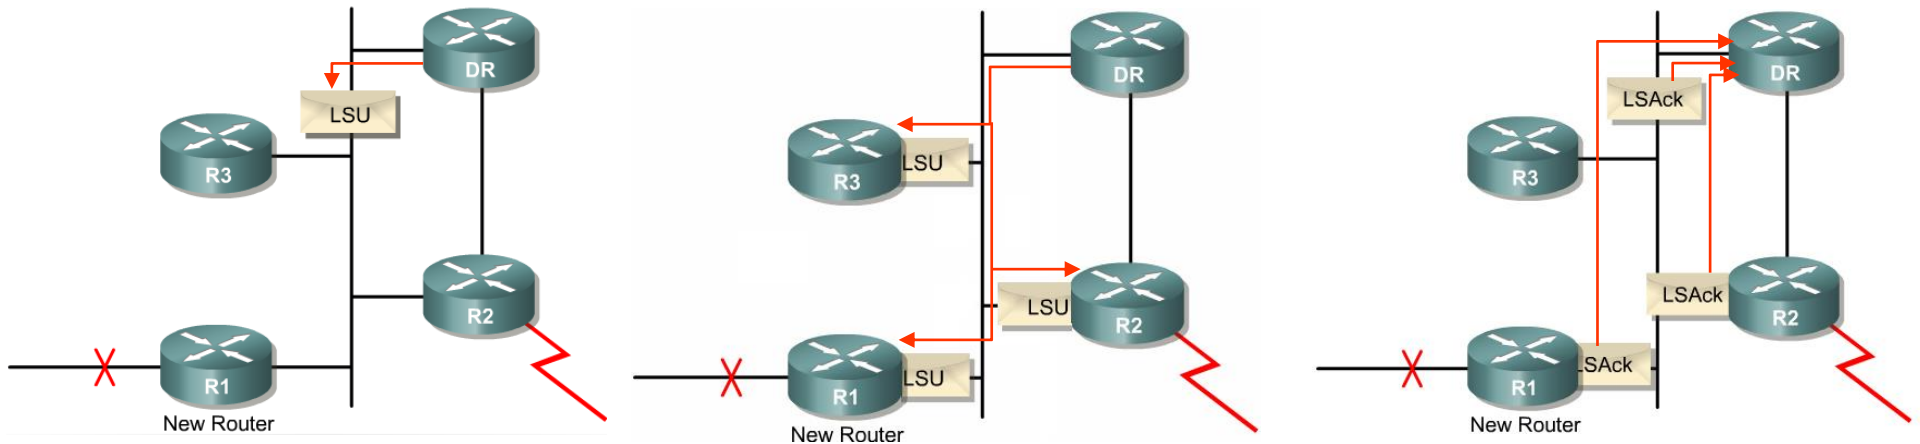


R1 detects a downed link and sends an LSU to the DR.
The DR acknowledges receipt of the LSU.



OSPF – The Interior Gateway Routing Protocol

Step 5: Maintaining Route Information




After the DR acknowledges receipt of the LSU, it floods the LSU to all OSPF routers on the network using the multicast address 224.0.0.5. Each router acknowledges receipt of the LSU with an LSAck.




OSPF – The Interior Gateway Routing Protocol

Step 5: Maintaining Route Information




If an OSPF router is attached to another network, it floods the LSU to other networks.



OSPF – The Interior Gateway Routing Protocol

Step 5: Maintaining Route Information




After an OSPF router receives an LSU that includes new information, it updates its link-state database. It then runs the SPF algorithm with the new information to recalculate the routing table.



OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/100.htm>)

Encapsulation of OSPF Packets:

- ✦ OSPF runs directly over the Internet Protocol's network layer. *OSPF packets are therefore encapsulated solely by IP and local data-link headers.*
- ✦ OSPF does not define a way to fragment its protocol packets, and depends on IP fragmentation when transmitting packets larger than the network MTU. The OSPF packet types that are likely to be large (Database Description Packets, Link State Request, Link State Update, and Link State Acknowledgment packets) can usually be split into several separate protocol packets, without loss of functionality. This is recommended; *IP fragmentation should be avoided whenever possible.* Using this reasoning, an attempt should be made to limit the sizes of packets sent over virtual links to 576 bytes. However, if necessary, the length of OSPF packets can be up to 65,535 bytes (including the IP header).
- ✦ The other important features of OSPF's IP encapsulation are:
 - ❑ Some OSPF messages are **multicast**, when sent over multi-access networks. Two distinct IP multicast addresses are used. Packets sent to these multicast addresses should never be forwarded; they are meant to *travel a single hop only. To ensure that these packets will not travel multiple hops, their IP TTL must be set to 1.*
 - This multicast address has been assigned the value **224.0.0.5**. *All routers running OSPF should be prepared to receive packets sent to this address.* Hello packets are always sent to this destination. Also, certain OSPF protocol packets are sent to this address during the flooding procedure.
 - This multicast address has been assigned the value **224.0.0.6**. *Both the Designated Router and Backup Designated Router must be prepared to receive packets destined to this address.* Certain OSPF protocol packets are sent to this address during the flooding procedure.
 - ❑ OSPF is IP protocol number 89.
 - ❑ Routing protocol packets are sent with IP TOS of 0.
 - ❑ OSPF protocol packets should be given precedence over regular IP data traffic, in both sending and receiving.



OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/101.htm>)


The OSPF Options Field:

- ✦ The OSPF Options field is present in OSPF *Hello packets, Database Description packets and all link state advertisements*.
- ✦ When used in **Hello packets**, the Options field allows a router to reject a neighbor because of a capability mismatch. Alternatively, when capabilities are exchanged in **Database Description packets** a router can choose not to forward certain link state advertisements to a neighbor because of its reduced functionality. Lastly, listing capabilities in **link state advertisements** allows routers to route traffic around reduced functionality routers, by excluding them from parts of the routing table calculation.
- ✦ Two capabilities are currently defined:



Option Field


- ❖ **T-bit:** This describes the router's TOS capability. If the T-bit is reset, then the router supports only a single TOS (TOS 0). Such a router is also said to be incapable of TOS-routing.
- ❖ **E-bit:** When the E-bit is reset in the Hello packet sent out a particular interface, it means that the router will neither send nor receive AS external link state advertisements on that interface. Two routers will not become neighbors unless they agree on the state of the E-bit



OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/102.htm>)

OSPF Packet Formats:

- ✦ There are five distinct OSPF packet types. *All OSPF packet types begin with a standard 24 byte header (OSPF packet Header).*
- ✦ All OSPF packet types are:
 - ❑ The Hello packet
 - ❑ The Database Description packet
 - ❑ The Link State Request packet
 - ❑ The Link State Update packet
 - ❑ The Link State Acknowledgment packet

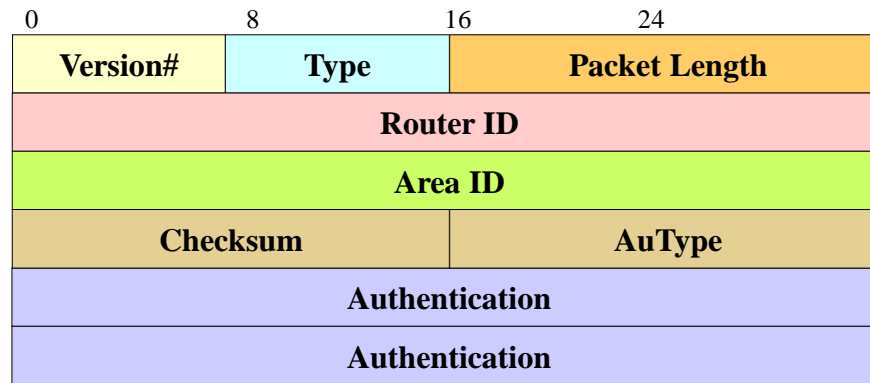


OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/103.htm>)

The OSPF Packet Header:

- ✦ Every OSPF packet starts with a *common 24 byte header*.



The OSPF Packet Header

Version #: The OSPF version number.

Type: The OSPF packet types are as follows:

Type	Description
1	Hello
2	Database Description
3	Link State Request
4	Link State Update
5	Link State Acknowledgment



OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/103.htm>)

The OSPF Packet Header:

Packet length: The length of the protocol packet in bytes. This length includes the standard OSPF header.


Router ID: *The Router ID of the packet's source.* In OSPF, the source and destination of a routing protocol packet are the two ends of an (potential) adjacency.

Area ID: A 32 bit number *identifying the area* that this packet belongs to.

Checksum: The standard IP checksum of the entire contents of the packet, starting with the OSPF packet header but excluding the 64-bit authentication field. This checksum is calculated as the 16-bit one's complement of the one's complement sum of all the 16-bit words in the packet, excepting the authentication field.

AuType: Identifies the authentication scheme to be used for the packet.

Authentication: A 64-bit field for use by the authentication scheme.



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/104.htm>)

The Hello Packet:

0	8	16	24
Version#	Type	Packet Length	
Router ID			
Area ID			
Checksum		AuType	
Authentication			
Authentication			
Network Mask			
Hello Interval		Options	Rtr Pri
Dead Interval			
Designated Router			
Backup Designated Router			
Neighbor Router ID			
Neighbor Router ID			
Additional Neighbor Router ID Fields			

The Hello Packet



OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/104.htm>)

The Hello Packet:

Network mask: The network mask associated with this interface. For example, if the interface is to a class B network, the network mask is 0xfffff00.

Options: The optional capabilities supported by the router.

HelloInterval: *The number of seconds between this router's Hello packets.*


Rtr Pri: This router's Router Priority. Used in *Designated Router and Backup Designated Router election*. If set to 0, the router will be ineligible to become Designated Router or Backup Designated Router.

RouterDeadInterval: The number of seconds before *declaring a silent router down*.

Designated Router: The *identity of the Designated Router for this network*, in the view of the advertising router. The Designated Router is identified here by its IP interface address on the network. Set to 0.0.0.0 if there is no Designated Router.

Backup Designated Router: The identity of the Backup Designated Router for this network, in the view of the advertising router. The Backup Designated Router is identified here by its IP interface address on the network. Set to 0.0.0.0 if there is no Backup Designated Router.

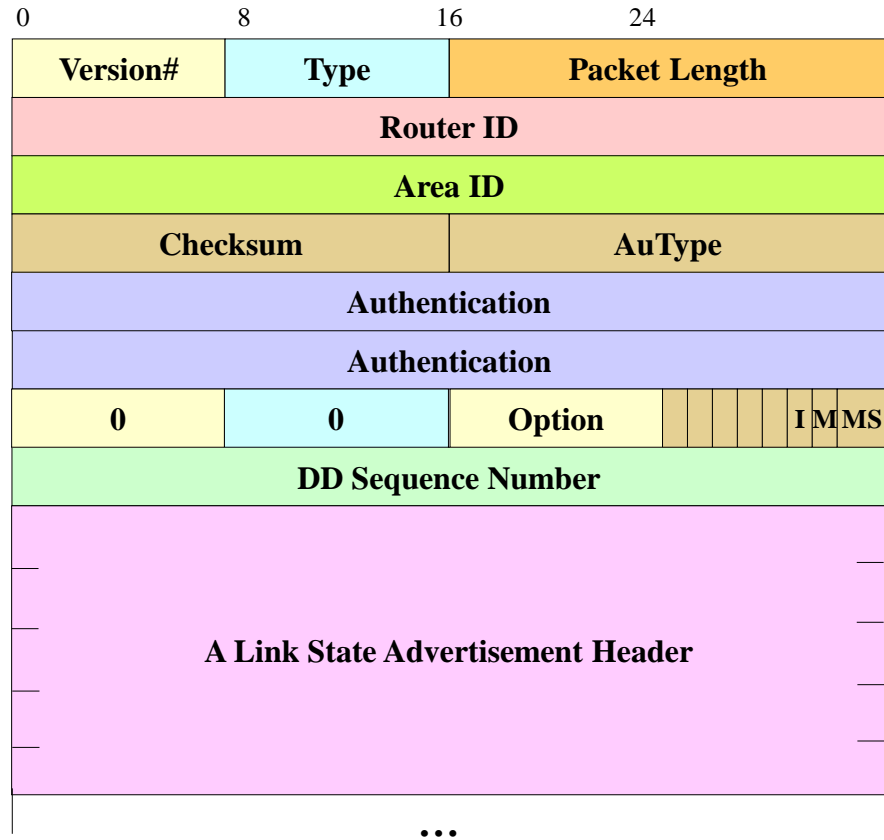
Neighbor: The Router IDs of each router from whom valid Hello packets have been seen recently on the network. Recently means in the last *RouterDeadInterval* seconds.



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/105.htm>)

The Database Description Packet:



The Database Description Packet



OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/105.htm>)

The Database Description Packet:

- ⊕ Database Description packets are OSPF packet *type 2*. These packets are exchanged when an adjacency is being initialized. They describe the contents of the topological database. Multiple packets may be used to describe the database. For this purpose a poll-response procedure is used. One of the routers is designated to be master, the other a slave. The master sends Database Description packets (polls) which are acknowledged by Database Description packets sent by the slave (responses). The responses are linked to the polls via the packets' DD sequence numbers.

Options: The optional capabilities supported by the router.


I-bit: The Init bit. When set to 1, this packet is *the first in the sequence* of Database Description Packets.

M-bit: The More bit. When set to 1, it indicates that *more Database Description Packets are to follow*.

MS-bit: The *Master/Slave bit*. When set to 1, it indicates that the router is the master during the Database Exchange process. Otherwise, the router is the slave.

DD sequence number: Used to sequence the collection of Database Description Packets. The initial value (indicated by the Init bit being set) should be unique. The DD sequence number then increments until the complete database description has been sent.

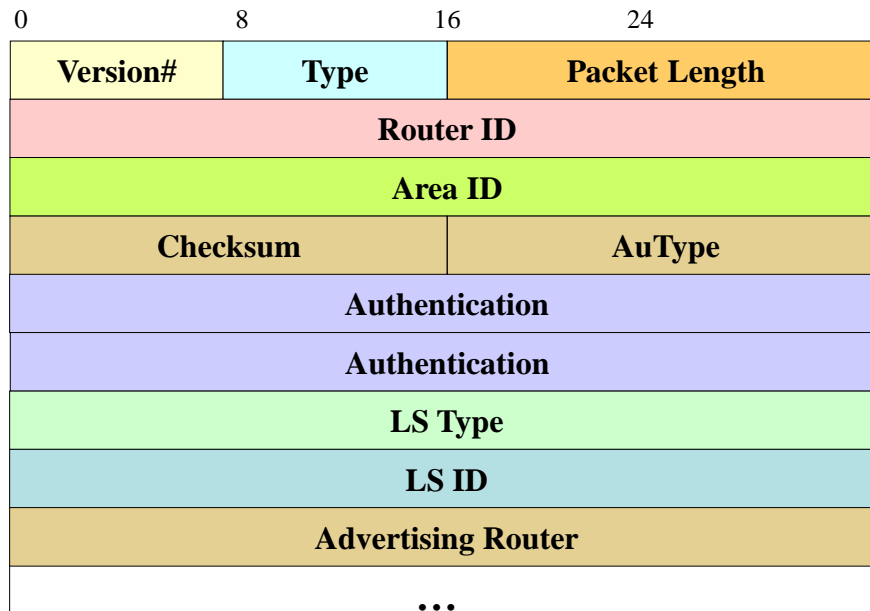
- ⊕ The rest of the packet consists of a (possibly partial) list of the topological database's pieces. Each *link state advertisement (LSA)* in the database is described by its link state advertisement header.




OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/106.htm>)

The Link State Request Packet:




The Link State Request Packet



OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/106.htm>)

The Link State Request Packet:

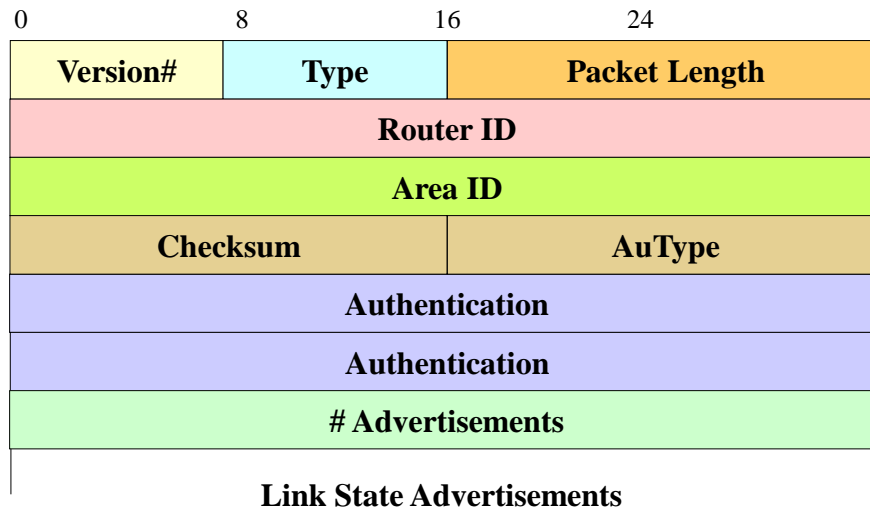
- ⊕ Link State Request packets are OSPF packet *type 3*. After exchanging Database Description packets with a neighboring router, a router may find that parts of its topological database are out of date. The Link State Request packet is used to *request the pieces of the neighbor's database that are more up to date*. Multiple Link State Request packets may need to be used. The sending of Link State Request packets is the last step in bringing up an *adjacency*.
- ⊕ A router that sends a Link State Request packet has in mind the precise instance of the database pieces it is requesting, *defined by LS sequence number, LS checksum, and LS age*, although these fields are not specified in the Link State Request Packet itself. *The router may receive even more recent instances in response*.
- ⊕ Each advertisement requested is specified by its *LS type, Link State ID, and Advertising Router*. This uniquely identifies the advertisement. Link State Request packets are understood to be requests for the most recent instance.



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/107.htm>)

The Link State Update Packet:



...

The Link State Update Packet



OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/107.htm>)

The Link State Update Packet:

- ⊕ Link State Update packets are OSPF packet *type 4*. These packets implement the flooding of link state advertisements. Each Link State Update packet carries a collection of link state advertisements one hop further from its origin. Several link state advertisements may be included in a single packet.
- ⊕ Link State Update packets are *multicast* on those physical networks that support multicast/broadcast. In order to make the flooding procedure reliable, flooded advertisements are *acknowledged* in Link State Acknowledgment packets. If retransmission of certain advertisements is necessary, the *retransmitted advertisements* are always carried by *unicast* Link State Update packets.

advertisements: The number of link state advertisements included in this update.

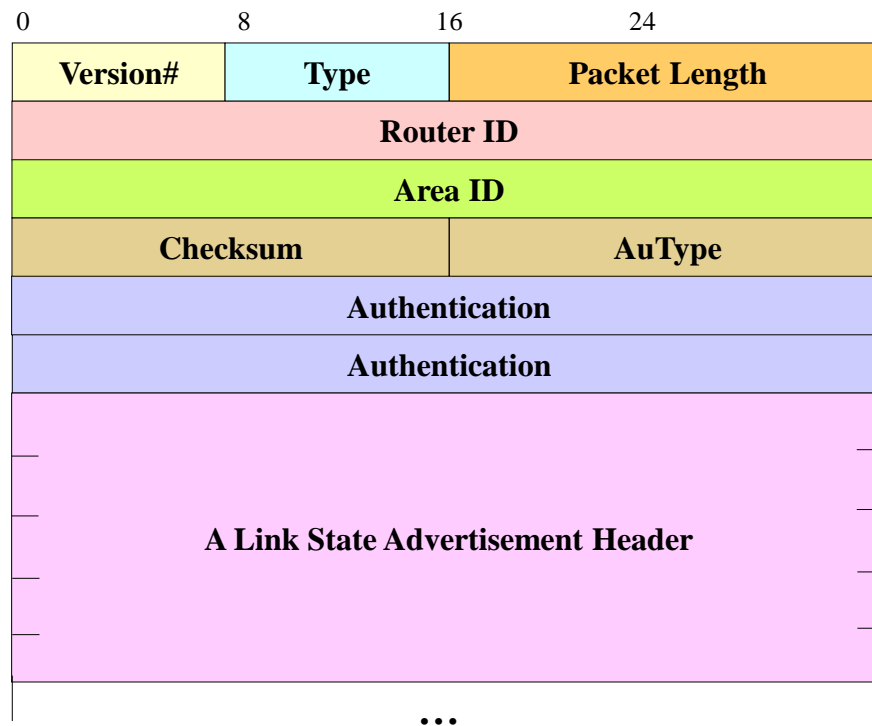
- ⊕ The body of the Link State Update packet consists of a list of link state advertisements. Each advertisement begins with a *common 20 byte header, the link state advertisement header*.



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/108.htm>)

The Link State Acknowledgment Packet:




The Link State Acknowledgment Packet



OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/108.htm>)

The Link State Acknowledgment Packet:

- ⊕ Link State Acknowledgment Packets are OSPF packet *type 5*. To make the flooding of link state advertisements reliable, flooded advertisements are explicitly acknowledged. This acknowledgment is accomplished through the sending and receiving of Link State Acknowledgment packets. Multiple link state advertisements can be acknowledged in a single Link State Acknowledgment packet.
- ⊕ Depending on the state of the sending interface and the source of the advertisements being acknowledged, a Link State Acknowledgment packet is sent either to the *multicast address AllSPFRouters*, to the *multicast address AllDRouters*, or as a *unicast*.
- ⊕ The format of this packet is similar to that of the Data Description packet. The body of both packets is simply a list of link state advertisement headers.




OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/109.htm>)

Link state advertisement formats:

- ⊕ There are five distinct types of link state advertisements. Each link state advertisement *begins with a standard 20-byte link state advertisement header.*
- ⊕ Each link state advertisement describes a piece of the OSPF routing domain. Every router originates *a router links advertisement*. In addition, whenever the router is elected Designated Router, it originates *a network links advertisement*. All link state advertisements are then flooded throughout the OSPF routing domain. The flooding algorithm is reliable, ensuring that all routers have the same collection of link state advertisements. This collection of advertisements is called the link state (or topological) database.
- ⊕ From the link state database, each router constructs a shortest path tree with itself as root. This yields a routing table.
- ⊕ The Link State Advertisements are:
 - ❏ Router links advertisements
 - ❏ Network links advertisements
 - ❏ Summary link advertisements
 - ❏ AS external link advertisements

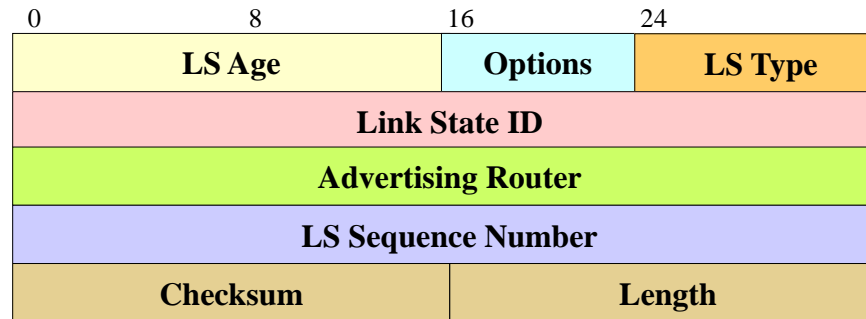
OSPF – (Optional -



<http://www.freesoft.org/CIE/RFC/1583/110.htm>)

The Link State Advertisement Header:

- All link state advertisements begin with a common 20 byte header. This header contains enough information to uniquely identify the advertisement (LS type, Link State ID, and Advertising Router). Multiple instances of the link state advertisement may exist in the routing domain at the same time. It is then necessary to determine which instance is more recent. This is accomplished by examining the LS age, LS sequence number and LS checksum fields that are also contained in the link state advertisement header.



The Link State Advertisement Header

LS age: The time in seconds since the link state advertisement was originated.

LS type: The type of the link state advertisement.

Each link state type has a separate advertisement format:

LS Type	Description
1	Router Links
2	Network Links
3	Summary link (IP network)
4	Summary link (ASBR)
5	AS external link



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/110.htm>)

The Link State Advertisement Header:


Link State ID: This field identifies the *portion of the internet environment that is being described by the advertisement*. The contents of this field depend on the advertisement's LS type. For example, in network links advertisements the Link State ID is set to the *IP interface address of the network's Designated Router*.

Advertising Router: The Router ID of the router that originated the link state advertisement. For example, in network links advertisements this field is set to the Router ID of the network's Designated Router.

LS sequence number: Detects old or duplicate link state advertisements. Successive instances of a link state advertisement are given successive LS sequence numbers.

LS checksum: The Fletcher checksum of the complete contents of the link state advertisement, including the link state advertisement header but excepting the LS age field.

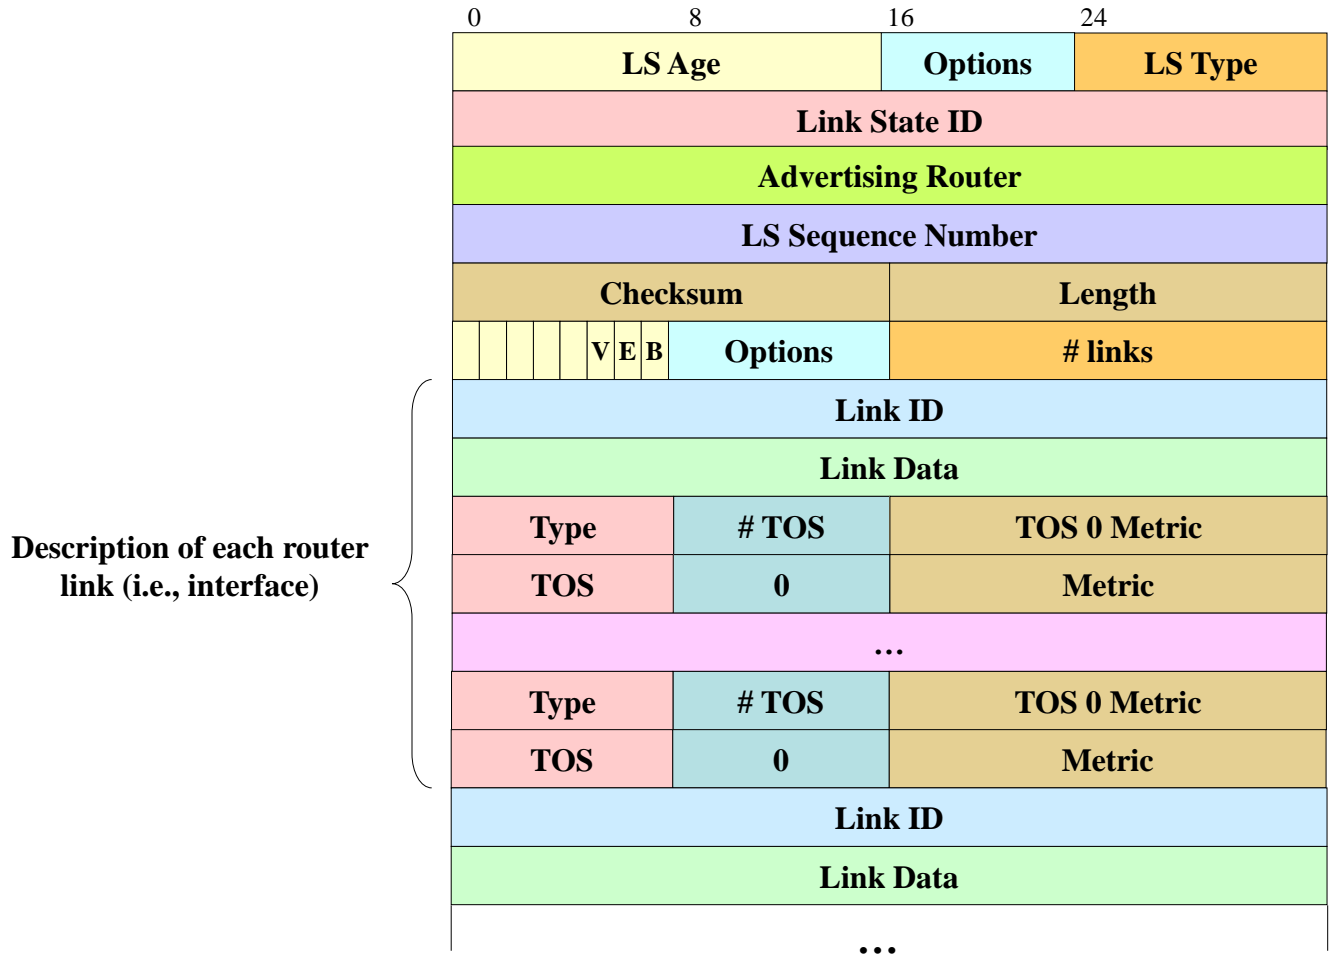
Length: The length in bytes of the link state advertisement. This includes the 20 byte link state advertisement header.



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/111.htm>)

Router Links Advertisements:





OSPF – (Optional - <http://www.freesoft.org/CIE/RFC/1583/111.htm>)

Router Links Advertisements:

- ⊕ Router links advertisements are the *Type 1 link state advertisements*. Each router in an area originates a router links advertisement. The advertisement describes *the state and cost of the router's links (i.e., interfaces)* to the area. *All of the router's links to the area must be described in a single router links advertisement.*
- ⊕ In router links advertisements, the Link State ID field is set to *the router's OSPF Router ID*. The T-bit is set in the advertisement's Option field if and only if the router is able to calculate a separate set of routes for each IP TOS. *Router links advertisements are flooded throughout a single area only.*

bit V: When set, the router is an endpoint of an active virtual link that is using the described area as a Transit area (V is for virtual link endpoint).

bit E: When set, the router is an AS boundary router (E is for external)

bit B: When set, the router is an area border router (B is for border)

links: The number of router links described by this advertisement. This must be the total collection of router links (i.e., interfaces) to the area.



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/111.htm>)

Router Links Advertisements:

- ⊕ The following fields are used to *describe each router link (i.e., interface)*. Each router link is typed. The *Type* field indicates the kind of link being described. It may be a link to a *transit network*, to *another router* or to *a stub network*. The values of all the other fields describing a router link depend on the link's Type. For example, each link has an associated 32-bit data field. *For links to stub networks this field specifies the network's IP address mask*. For other link types the Link Data specifies the router's associated IP interface address.

Type: A quick description of the router link. One of the following. Note that host routes are classified as links to stub networks whose network mask is 0xffffffff.

Type	Description
1	Point-to-point connection to another router
2	Connection to a transit network
3	Connection to a stub network
4	Virtual link



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/111.htm>)

Router Links Advertisements:

Link ID: Identifies the object that this router link connects to. Value depends on the link's Type. When connecting to an object that also originates a link state advertisement (i.e., another router or a transit network) the Link ID is equal to the neighboring advertisement's Link State ID. This provides the key for looking up said advertisement in the link state database.

Type	Description
1	Neighboring router's Router ID
2	IP address of Designated Router
3	IP network/subnet number
4	Neighboring router's Router ID

Link Data: Contents again depend on the link's Type field.

TOS: The number of different TOS metrics given for this link, not counting the required metric for TOS 0. For example, if no additional TOS metrics are given, this field should be set to 0.

TOS 0 metric: The cost of using this router link for TOS 0.



OSPF – (Optional -


<http://www.freesoft.org/CIE/RFC/1583/111.htm>)

Router Links Advertisements:

- ✦ For each link, separate metrics may be specified for each Type of Service (TOS). The metric for TOS 0 must always be included, and was discussed above. Metrics for non-zero TOS are described below. The encoding of TOS in OSPF link state advertisements is described in Section 12.3. Note that the cost for non-zero TOS values that are not specified defaults to the TOS 0 cost. Metrics must be listed in order of increasing TOS encoding. For example, the metric for TOS 16 must always follow the metric for TOS 8 when both are specified.

TOS: IP Type of Service that this metric refers to.

Metric: The cost of using this outbound router link, for traffic of the specified TOS.

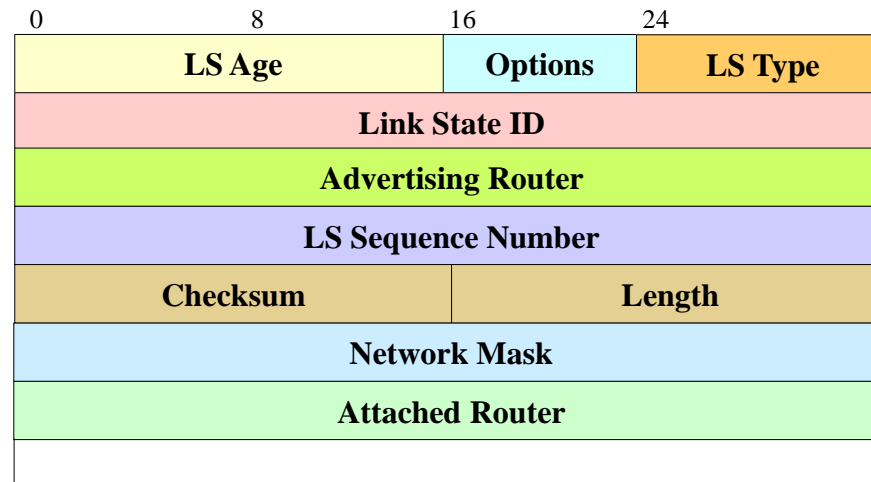


OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/112.htm>)

Network Links Advertisements:

- Network links advertisements are the *Type 2 link state advertisements*. A network links advertisement is originated for each transit network in the area. A transit network is a multi- access network that has more than one attached router. *The network links advertisement is originated by the network's Designated Router. The advertisement describes all routers attached to the network, including the Designated Router itself. The advertisement's Link State ID field lists the IP interface address of the Designated Router.*
- The distance from the network to all attached routers is zero, for all Types of Service. This is why the TOS and metric fields need not be specified in the network links advertisement.



Network Links Advertisements




OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/112.htm>)

Network Links Advertisements:

Network Mask: The IP address mask for the network.

Attached Router: The Router IDs of each of the routers attached to the network. Actually, *only those routers that are fully adjacent to the Designated Router are listed*. The Designated Router includes itself in this list. The number of routers included can be deduced from the link state advertisement header's length field.

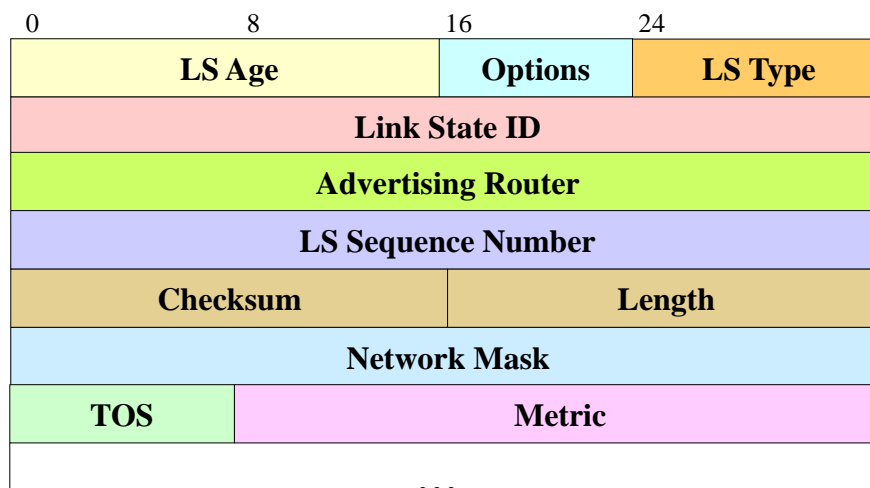


OSPF – (Optional -


<http://www.freesoft.org/CIE/RFC/1583/113.htm>)

Summary Link Advertisements:

- Summary link advertisements are the *Type 3 and 4 link state advertisements*. These advertisements are originated by *area border routers*. A separate summary link advertisement is made for each destination (known to the router) which belongs to the AS, yet is outside the area.
- Type 3 link state advertisements are used when the destination is an IP network. In this case the advertisement's Link State ID field is an IP network number. When the destination is an AS boundary router, a Type 4 advertisement is used, and the Link State ID field is the AS boundary router's OSPF Router ID.* Other than the difference in the Link State ID field, the format of Type 3 and 4 link state advertisements is identical.



Summary Link Advertisements



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/113.htm>)

Summary Link Advertisements:


- ✦ For stub areas, Type 3 summary link advertisements can also be used to describe a (per-area) default route. Default summary routes are used in stub areas instead of flooding a complete set of external routes. When describing a default summary route, the advertisement's Link State ID is always set to DefaultDestination (0.0.0.0) and the Network Mask is set to 0.0.0.0.
- ✦ Separate costs may be advertised for each IP Type of Service. The encoding of TOS in OSPF link state advertisements is described in Section 12.3. Note that the cost for TOS 0 must be included, and is always listed first. If the T-bit is reset in the advertisement's Option field, only a route for TOS 0 is described by the advertisement. Otherwise, routes for the other TOS values are also described; if a cost for a certain TOS is not included, its cost defaults to that specified for TOS 0.

Network Mask: For Type 3 link state advertisements, this indicates the destination network's IP address mask. For example, when advertising the location of a class A network the value 0xff000000 would be used. This field is not meaningful and must be zero for Type 4 link state advertisements.

- ✦ For each specified Type of Service, the following fields are defined. The number of TOS routes included can be calculated from the link state advertisement header's length field. Values for TOS 0 must be specified; they are listed first. Other values must be listed in order of increasing TOS encoding. For example, the cost for TOS 16 must always follow the cost for TOS 8 when both are specified.

TOS: The Type of Service that the following cost concerns.

Metric: The cost of this route. Expressed in the same units as the interface costs in the router links advertisements.

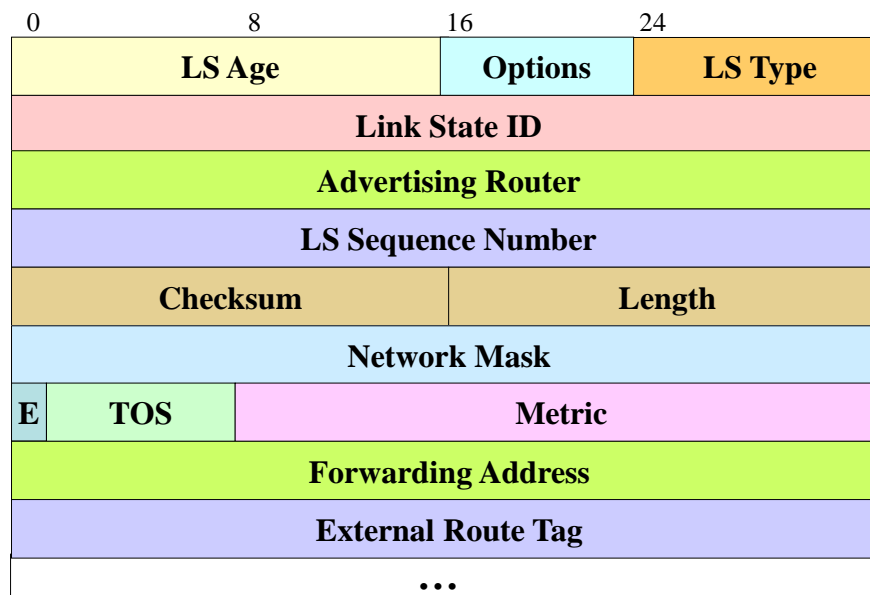



OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/114.htm>)

AS External Link Advertisements:

- AS external link advertisements are the Type 5 link state advertisements. These advertisements are originated by AS boundary routers. A separate advertisement is made for each destination (known to the router) which is external to the AS.
- AS external link advertisements usually describe a particular external destination. For these advertisements the Link State ID field specifies an IP network number. AS external link advertisements are also used to describe a default route. Default routes are used when no specific route exists to the destination. When describing a default route, the Link State ID is always set to DefaultDestination (0.0.0.0) and the Network Mask is set to 0.0.0.0.





OSPF – (Optional -

<http://www.freesoft.org/CIE/RFC/1583/114.htm>)

AS External Link Advertisements:

- ✦ Separate costs may be advertised for each IP Type of Service. The encoding of TOS in OSPF link state advertisements is described in Section 12.3. Note that the cost for TOS 0 must be included, and is always listed first. If the T-bit is reset in the advertisement's Option field, only a route for TOS 0 is described by the advertisement. Otherwise, routes for the other TOS values are also described; if a cost for a certain TOS is not included, its cost defaults to that specified for TOS 0.

Network Mask: The IP address mask for the advertised destination.

- ✦ For each specified Type of Service, the following fields are defined. The number of TOS routes included can be calculated from the link state advertisement header's length field. Values for TOS 0 must be specified; they are listed first. Other values must be listed in order of increasing TOS encoding. For example, the cost for TOS 16 must always follow the cost for TOS 8 when both are specified.

bit E: The type of external metric. If bit E is set, the metric specified is a Type 2 external metric. This means the metric is considered larger than any link state path. If bit E is zero, the specified metric is a Type 1 external metric. This means that it is comparable directly (without translation) to the link state metric.

Forwarding address: Data traffic for the advertised destination will be forwarded to this address. If the Forwarding address is set to 0.0.0.0, data traffic will be forwarded instead to the advertisement's originator (i.e., the responsible AS boundary router).

TOS: The Type of Service that the following cost concerns.

Metric: The cost of this route. Interpretation depends on the external type indication (bit E above).

External Route Tag: A 32-bit field attached to each external route. This is not used by the OSPF protocol itself. It may be used to communicate information between AS boundary routers; the precise nature of such information is outside the scope of this specification.