





OSPF – LSAs

Why there are a dirty dozen of them
Part 3



*“An algorithm
must be seen
to be believed”*




Donald E. Knuth

LSA Sequence Number



- In order to stop flooding, each LSA carries a sequence number
- Only increased if LSA has changed
 - ◆ So each router can check if a particular LSA had already been forwarded
 - ◆ To avoid LSA storms
- 32 bit number



(C) Herbert Haas 2005/03/11

3

When reaching the end of the 32 bit sequence number the associated router will wait for an hour so that this LSA ages out in each link state database. Then the router resets the sequence number (lowest negative number i. e. MSB=1, 80000001) and continues to flood this LSA.

Each LSA carries also a 16 bit age value, which is set to zero when originated and increased by every router during flooding. LSAs are also aged as they are held in each router's database. If sequence numbers are the same, the router compares the ages the younger the better but only if the age difference between the recently received LSA is greater than MaxAgeDiff; otherwise both LSAs are considered to be identical.

Note:

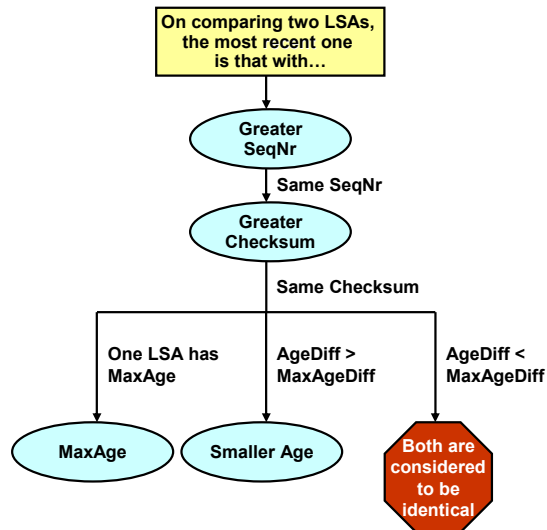
Radia Perlman proposed a "Lollipop" sequence number space but today a linear space is used as described above.

Since signed integers are used to describe sequence numbers, 8000001 represents the most-negative number in a hexadecimal format. To verify this, the 2-complement of this number must be calculated. This can be done in two steps. First calculate the 1-complement by simply inverting the binary number, that is the most significant byte "0x80" which is "1000 000" is transformed to "0111 111", the least significant byte "0x01" which is "0000 0001" is transformed to "1111 1110" and all other bytes inbetween are now "1111 1111". Secondly, in order to receive the 2-complement, "1" must be added. Then the final result is "0111 1111 1111 1111 1111 1111 1111 1111", which is the absolute number (without sign).

Detailed Flooding Decisions



- LSA is identified by its
 - ♦ LS type
 - ♦ Link State ID
 - ♦ Advertising Router
- The most recent one of two instances of the same LSA is determined by:
 - ♦ LS sequence number
 - ♦ LS checksum
 - ♦ LS age
- MaxAgeDiff (15 min) as tolerance value



Each LSA carries also a 16 bit age value, which is set to zero when originated and increased by every router during flooding. LSAs are also aged as they are held in each router's database. If sequence numbers are the same, the router compares the ages the younger the better but only if the age difference between the recently received LSA is greater than MaxAgeDiff; otherwise both LSAs are considered to be identical.

LS Age



- **Originating router sets LS age = 0 seconds**
- **Increased during flooding by InfTransDelay by every router**
- **Also increased while stored in database**
- **Age is never incremented past MaxAge (60 min)**
- **LSAs having MaxAge:**
 - ◆ Are not used in routing table calculation anymore
 - ◆ Are reflooded immediately
 - ◆ Are always considered as most recent
 - ◆ Thus quickly flushed from routing domain
- **Responsible router maintains LSRefreshTime (30 min) to refresh LSAs periodically**

Router LSA – Type 1



- **Router ID (Highest IP address)**
- **Number of Links**
- **Link Descriptions**
 - ◆ **Link type (P2P, Stub, ...)**
 - ◆ **Neighboring router ID**
 - ◆ **Router interface address**
 - ◆ **ToS (typically not supported today)**
 - ◆ **Metrics**

Network LSA – Type 2



- **DR's IP address**
- **One Subnet mask for this broadcast segment**
- **List of Router-IDs of all routers in the broadcast segment**

Network Summary LSA – Type 3



- Originated by **ABRs** only
- Each LSA Type 3 contains a number of
 - ◆ Destination networks + Subnet masks
 - ◆ Metric for each destination network
- This is basically a distance-vector routing information (!)

ASBR Summary LSA – Type 4



- Originated by **ABRs**
- Advertise routes to ASBRs
- Nearly identical to Type 3
 - ◆ Except destination is ASBR not a network
- Each LSA Type 4 contains
 - ◆ Router IDs of ASBRs
 - ◆ Mask 0.0.0.0 (host route)
 - ◆ Metric

AS External LSA – Type 5



- **Originated by ASBRs**
 - ◆ External type 1
 - ◆ External type 2 (default)
- **Advertises**
 - ◆ External routes
 - ◆ Default route
- **Contains**
 - ◆ External Net-ID + Mask
 - ◆ Metric
 - ◆ Next hop (external, not ASBR)

NSSA External LSA – Type 7



- **Originated by ASBRs within NSSAs**
- **Almost identical to Type 5**
 - ◆ **But only flooded within NSSA**
- **RFC 1587**

Other LSAs



- **Group Membership LSA (6)**
 - ♦ For MOSPF
- **External Attribute LSA (8)**
 - ♦ Alternative to IBGP
 - ♦ Should transport BGP information within an OSPF domain
 - ♦ Not yet implemented, no RFC yet (?)
- **Opaque LSA (9)**
 - ♦ Application specific information
 - ♦ Link local scope

Opaque LSAs are e. g. used as load indication messages with MPLS.

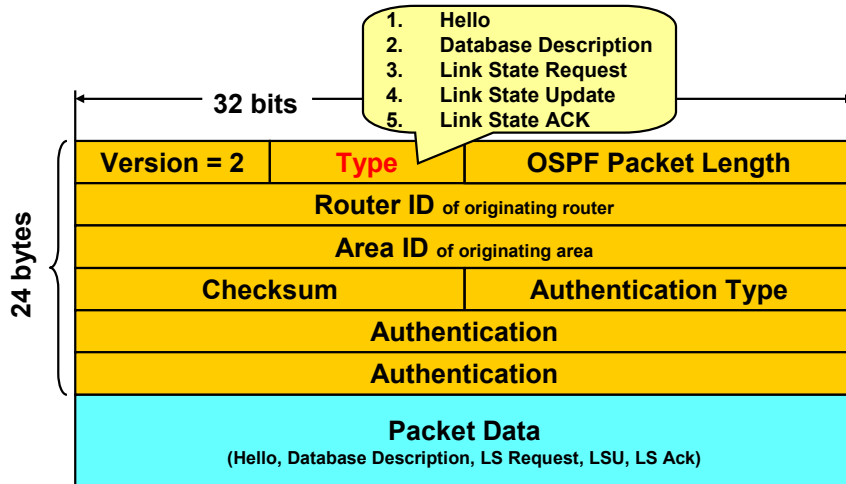
Other LSAs



- **Opaque LSA (10)**
 - ◆ Application specific information
 - ◆ Area-local scope
- **Opaque LSA (11)**
 - ◆ Application specific information
 - ◆ AS scope

Opaque LSAs are e. g. used as load indication messages with MPLS.

General OSPF Packet Structure



- Carried directly in IP (protocol number 89)
- **All OSPF packets begin with a 24-byte OSPF packet header**

The OSPF version we use today is version 2. The packet type identifies the actual OSPF message type that is carried in the packet data area at the bottom. The OSPF packet length describes the number of bytes of the OSPF packet including the OSPF header. Router and Area IDs identify the originator of this packet. If a packet is sent over a virtual link, the Area ID will be 0.0.0.0, because virtual links are considered part of the backbone area. The checksum is calculated over the entire packet including the header.

Three authentication types had been defined:

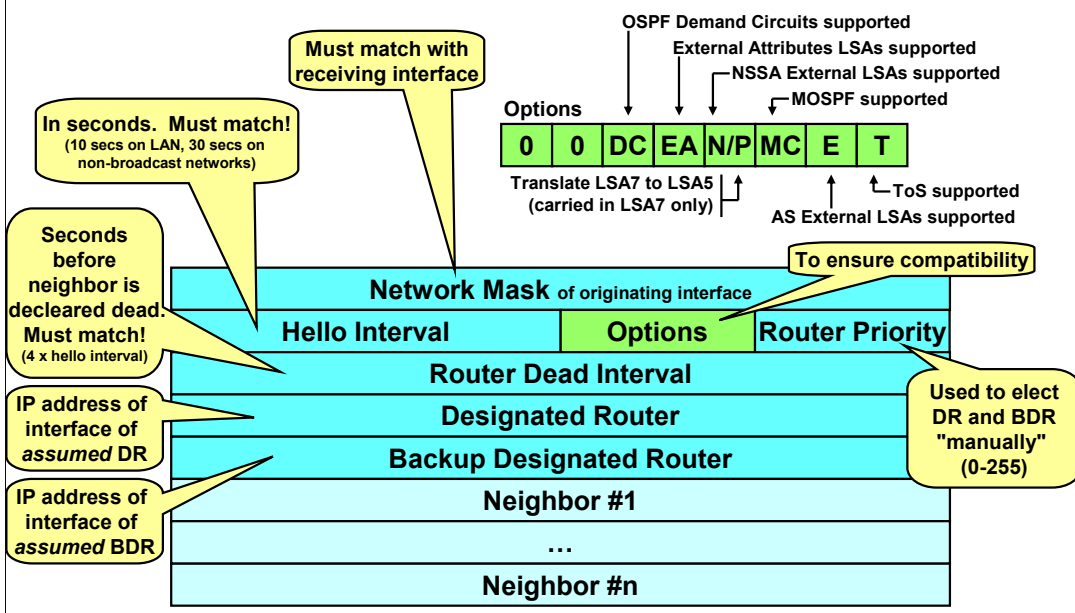
- 0 No authentication
- 1 Simple clear text password authentication
- 2 MD5 Checksum

If the Authentication Type = 1, then a 64 bit clear text password is carried in the authentication fields. If the Authentication Type = 2, then the authentication fields contain a key-ID, the length of the message digest, and a nondecreasing cryptographic sequence number to prevent replay attacks. The actual message digest would be appended at the end of the packet.

The efficiency of routing updates also depends on the maximum transfer unit (MTU) defined. Cisco defined a MTU of 1500 bytes for OSPF.

Hello Packet

Type 1



(C) Herbert Haas 2005/03/11

15

The network mask must match the mask on the receiving interface, ensuring that they share a segment and network.

The Options field is also used by other message types. If the Router Priority is set to zero this router cannot become DR or BDR.

Note that the fields "Designated Router" and "Backup Designated Router" only contain the interface IP address of the DR or BDR on that network, not the router ID !!

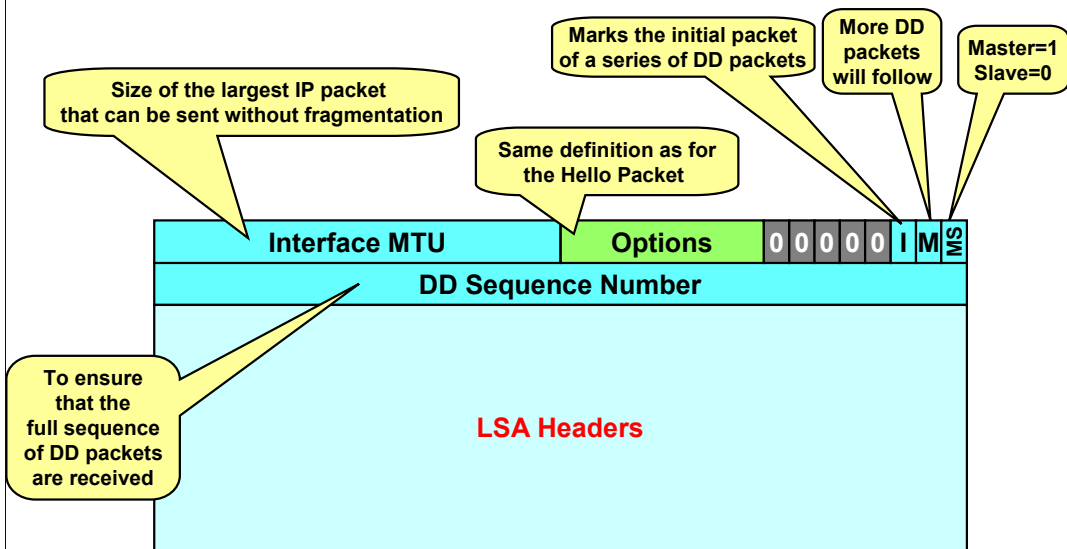
If these numbers are unknown or not necessary (other network type) then these fields are set to 0.0.0.0.

It is important to know that neighbors must have configured identical Hello and Dead Intervals.

Database Description Packet

Also called "DDP"

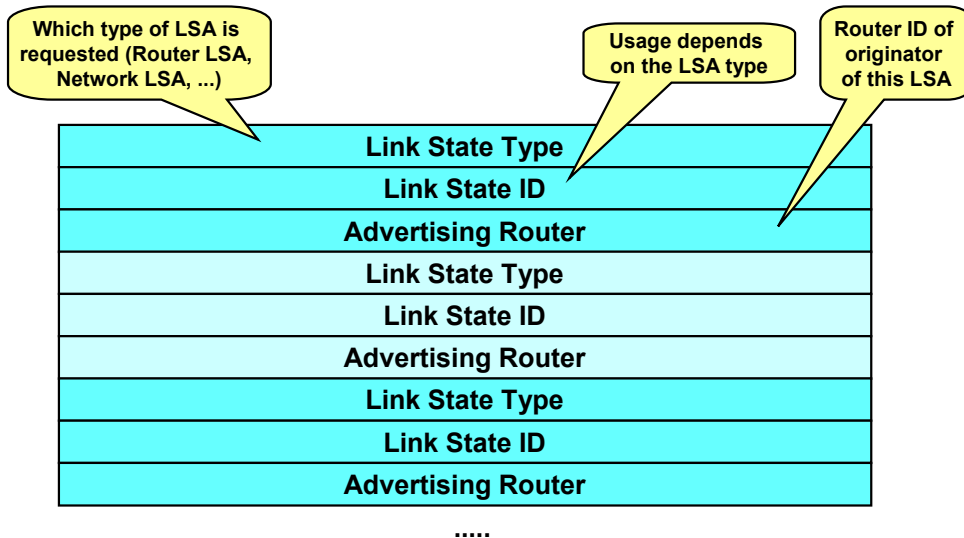
Type 2



The DD sequence number is set by the master to some unique value in the first DD packet. This number will be incremented in subsequent packets.

Link State Request Packet


Type 3



Note that the Link State Request Packet uniquely identifies the LSA by Type, ID, and advertising router fields of its header. It does not include the sequence number, checksum, and age, because the requestor is not interested in a specific instance of the LSA but in the most recent instance.

Link State Update Packet

Type 4



- LSUs contain one or more LSAs (limited by MTU)
- Used for flooding and response to LS requests
- LSUs are carried hop-by-hop

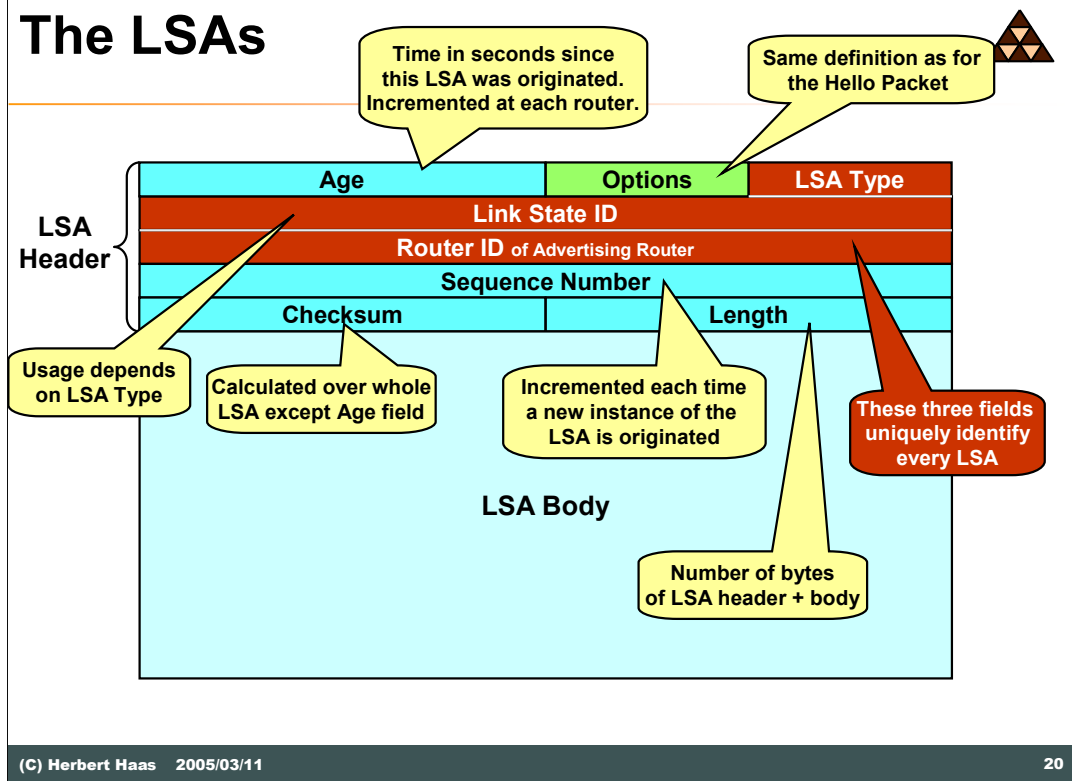
Link State ACK Packet

Type 5



- Each LSA received must be **explicitly** acknowledged → reliable flooding!
- Acknowledged LSA is identified by **LSA header**
- Single Link State ACK packet can acknowledge multiple LSAs

The LSAs



(C) Herbert Haas 2005/03/11

20

All LSAs have the LSA header at the beginning. This LSA header is also used in Database Description and Link State Acknowledgement packets.

The **Age** is incremented by **InfTransDelay** seconds at each router interface this LSA exits. The Age is also incremented in seconds as it resides in a link state database.

The **Options** field describes optional capabilities supported at that topological portion described by this LSA.

The **LSA Type** describes which information is carried in the LSA Body. Here the structural differences between Router LSAs, Network LSAs, etc. are identified.

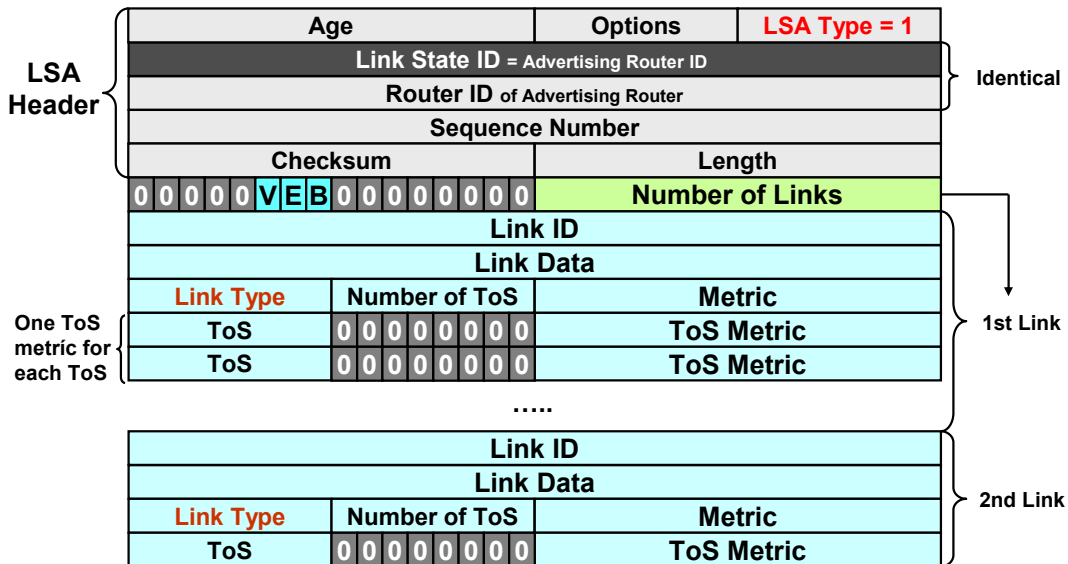
The **Link State ID** is used differently by the LSA types. Basically this field contains some information identifying the topological portion described by this LSA. For example a Router ID or an interface address is used here. The following slides will explain this field for each LSA type.

The **Router ID** identifies the originating router of this LSA.

The **Sequence Number** helps routers to identify the most recent instance of this LSA.

The **Checksum** is a so-called 8 bit Fletcher checksum, providing more protection than traditional checksum methods such as used for TCP. The first eight bits contain the 1's complement sum of all octets, while the second eight bits contain a high-order sum of the running sums. See RFC 1146 for more details.

Router LSA



Router LSAs are generated by all OSPF routers and must describe all links of the originating router!

The **V-bit** (Virtual Link Endpoint) is set to one if the originating router is a virtual link endpoint and this area is a transit area. The **E-bit** (External) is set if the originating router is an ASBR. The **B-bit** (Border) is set if the originating router is an ABR.

The **Link ID** and **Link Data** depend on the **Link Type** field which describes the general type of connection the link provides.

Link Type 1 is a point-to-point link, the Link ID describes the Neighbor Router ID and the Link Data field contains the IP address of the originating router's interface to the network.

Link Type 2 is a link to a transit network, the Link ID describes the interface address of the Designated Router and the Link Data field contains the IP address of the originating router's interface to the network.

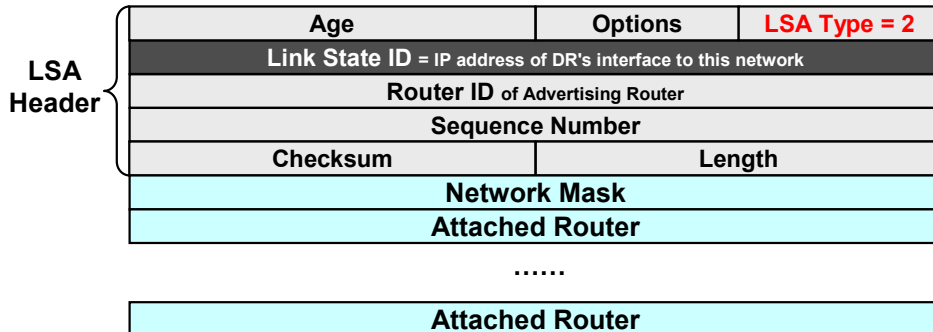
Link Type 3 is a link to stub network, the Link ID describes the IP network number or subnet address and the Link Data field contains the network's IP address or subnet mask.

Link Type 4 is a virtual link, the Link ID describes the neighboring router's Router ID and the Link Data contains the MIB-II ifIndex value for the originating router's interface.

Number of ToS specifies the number of **ToS Metrics** listed for this link. For each ToS an additional line is appended to this link state section. Generally, ToS is not used today anymore and the Number of ToS field is set to all-zero.

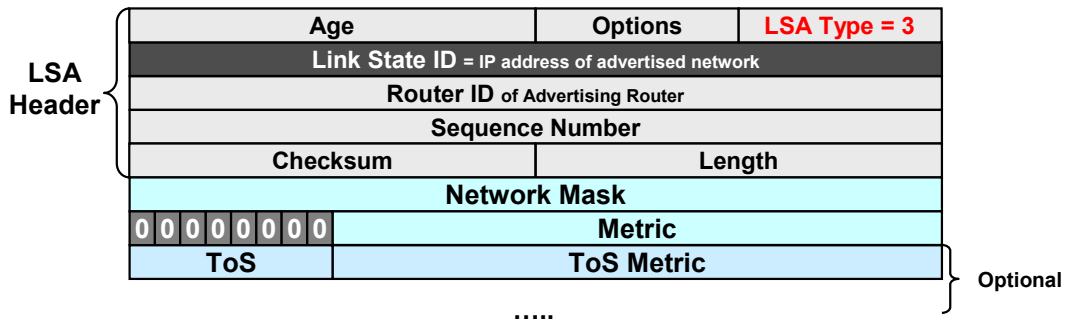
Metric is the cost of the interface that established this link.

Network LSA



Network LSAs are originated by DRs and describe the multi-access network and all routers attached to it, including the DR.

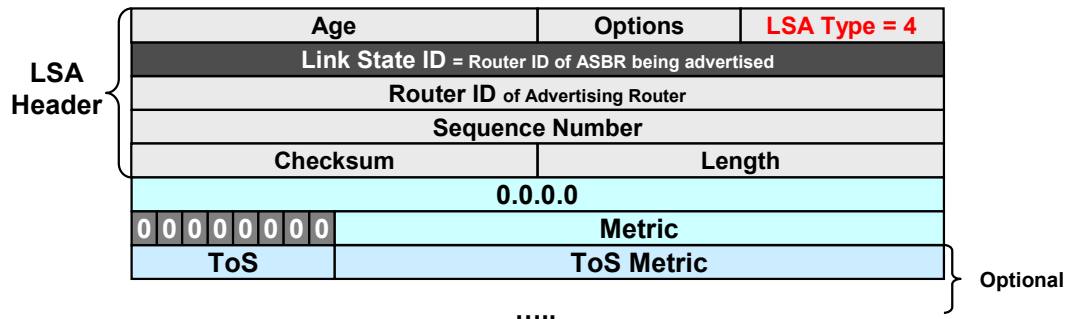
Network Summary LSA



- If a **default route** is advertised, both the Link State ID and the Network Mask fields will be 0.0.0.0
- Also used for route summarization
- Note: Cisco only supports ToS=0

A Network Summary LSA is originated by an ABR and advertises networks external to an area.

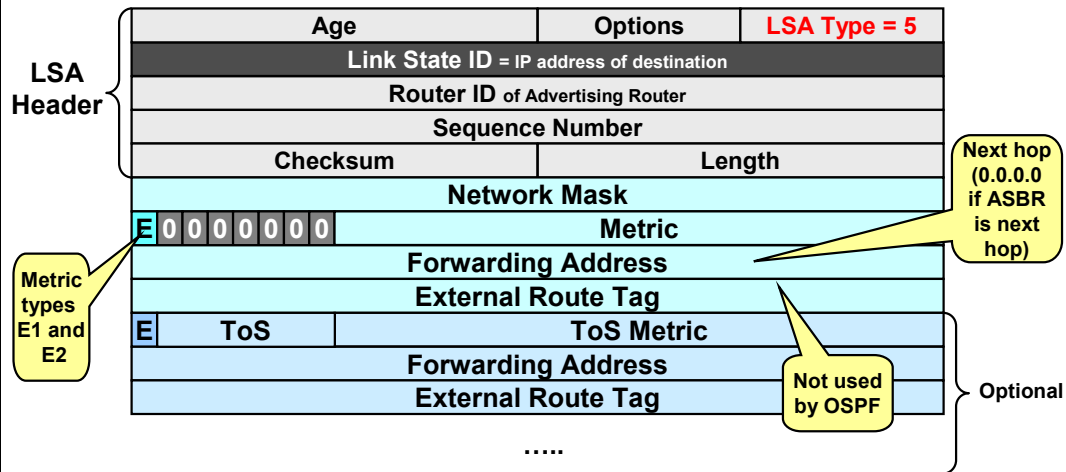
ASBR Summary LSA



- Note: Cisco only supports ToS=0

A ASBR Summary LSA is originated by an ABR and advertises ASBRs external to an area.

Autonomous System External LSA



- When describing a default route, both the Link State ID and the Network Mask are set to 0.0.0.0.

NSSA External LSA



- **Same structure as AS External LSA**
- **Forwarding address is**
 - ◆ **Next hop address for the network between NSSA and adjacent AS, if this network is advertised as internal route**
 - ◆ **Router ID of NSSA-ASBR otherwise**