

P. Bottorff, HPE Aruba Networking, [paul.bottorff@hpe.com](mailto:paul.bottorff@hpe.com)  
P. Congdon, Congdon Consulting, LLC, [paul.congdon@outlook.com](mailto:paul.congdon@outlook.com)  
April 22, 2025

**IEEE Std 802.1AB LLDP for IETF LSVR Neighbor Discovery and Configuration**

Abstract

IEEE Std 802.1AB, known as the Link Layer Discovery Protocol (LLDP), can be applied to LSVR neighbor discovery and configuration. This can be achieved by using a “nearest router group address” as the LLDP Scope MAC Address to target the LSVR interface and to advertise a set of LSVR-specific LLDP TLVs. These LSVR-specific TLVs are defined using LLDP Organizationally Specific TLVs, which are specified by LLDP for use by individual organizations allowing them to define their own Type-Length-Value (TLV) objects for exchange over the LLDP protocol. The IETF Organizationally Specific TLVs for LSVR can be encoded using the IETF IANA OUI (RFC 7042). This document provides an overview of applying LLDP to LSVR neighbor discovery and configuration and specifies the IETF Organizationally Specific TLVs that support link discovery for the LSVR routing protocol.

Table of Contents

1. Introduction.....2  
2. Terminology.....3  
3. Abbreviations.....4  
4. LLDP Operation.....4  
    4.1. Nearest Router Scope MAC Address for LLDP.....5  
    4.2. Logical Link Endpoint Identification.....6  
    4.3. Link Liveliness.....6  
    4.4. LLDP Organizationally Specific TLVs.....6  
        4.4.1. Type.....7  
        4.4.2. Length.....7  
        4.4.3. Organizationally Unique Identifier.....7  
        4.4.4. Subtype.....7  
        4.4.5. Information.....7  
    4.5. IETF Organizationally Specific TLVs for LSVR.....7  
    4.6. LSVR LLDP L3 data link application.....8  
5. Encapsulation TLVs for LSVR.....8  
    5.1. Encaps Flags.....9  
        5.1.1. Announce.....9  
        5.1.2. Primary.....9

5.1.3.	Under/Over.....	9
5.1.4.	Loopback.....	10
5.1.5.	Reserved.....	10
5.2.	IPv4 Announcement TLV.....	10
5.2.1.	Encaps Flags.....	10
5.2.2.	Prefix Length.....	10
5.2.3.	IPv4 Address.....	10
5.2.4.	IPv4 Announcement TLV usage rules.....	11
5.3.	IPv6 Announcement TLV.....	11
5.3.1.	Encaps Flags.....	11
5.3.2.	Prefix Length.....	11
5.3.3.	IPv6 Address.....	11
5.3.4.	IPv6 Announcement TLV usage rules.....	12
5.4.	MPLS IPv4 Announcement TLV.....	12
5.4.1.	MPLS Label List.....	12
5.4.1.1.	Label Count.....	12
5.4.1.2.	Label.....	13
5.4.1.3.	Exp.....	13
5.4.1.4.	S.....	13
5.4.2.	Encaps Flags.....	13
5.4.3.	Prefix Length.....	13
5.4.4.	IPv4 Address.....	13
5.4.5.	IPv4 Announcement TLV usage rules.....	13
5.5.	MPLS IPv6 Announcement TLV.....	13
5.5.1.	MPLS Label List.....	14
5.5.2.	Encaps Flags.....	14
5.5.3.	IPv6 Address.....	14
5.5.4.	Prefix Length.....	14
5.5.5.	MPLS IPv6 Announcement TLV usage rules.....	14
6.	Upper-Layer Protocol Configuration TLVs.....	14
6.1.1.	AttrTypes.....	15
6.1.2.	BGP ASN.....	15
6.1.3.	BGP IPv4 Peering Address.....	15
6.1.4.	BGP IPv6 Peering Address.....	15
6.1.5.	BGP Authentication Data.....	15
6.1.6.	BGP Flags.....	15
7.	References.....	16

## 1. Introduction

The IEEE Std 802.1AB [802.1AB][802.1ABcu][802.1ABdh], commonly known as the Link Layer Discovery Protocol (LLDP), is an extremely simple protocol used to advertise the identity and capabilities of router, switch, and end-station ports to network neighbors. It is widely deployed in data centers and campus networks where it provides information used to identify neighbors, capabilities, and current state. The LLDP information is used in turn by management systems to map the network topology and capabilities and by LLDP applications to perform protocol configuration.

The LLDP protocol operates over both point-to-point and multipoint networks. It can be used to advertise neighbors and capabilities both at the physical port level and at virtual interfaces within

systems. LLDP provides three standardized advertisement scopes which identify the type of port being advertised. These standard types are identified by a multicast Scope MAC Addresses which is used as an LLDP destination address. The standard port types are nearest bridge, nearest non-TPMR bridge, and nearest customer bridge. In addition, the standard allows the use of an alternate group MAC address or individual MAC address for the Scope MAC Address. For LSVR an alternate Scope MAC Address can be used to identify router ports which are internal to a switch or to identify router ports that are behind a bridged network.

LLDP provides a means for individual organizations to define their own Type-Length-Value (TLV) objects for exchange over the protocol. TLVs that belong to an organization are identified by the inclusion of the organization's OUI and an organizationally defined subtype in the initial octets of the information field.

The IETF is a standards development organization with an IANA OUI. The usage and considerations for this OUI are discussed in [RFC7042]. The IANA OUI can be used to identify IETF specific LLDP Organizationally Specific TLVs.

The LSVR working group is specifying protocols that need to discover IP Layer 3 attributes of links, encapsulations and neighbors. This document specifies a set of IETF LLDP Organizationally Specific TLVs that carry the necessary discovery attributes for LSVR working group protocols.

## 2. Terminology

**Extended Link Layer Discovery Protocol:** An LLDP extension used to retrieve the frames beyond the Normal LLDPDU of a multi-frame LLDP database.

**Extension LLDPDU:** An LLDPDU containing an Extension Identifier TLV and not containing a Time to Live TLV or Extension Request TLV.

**Extension Request LLDPDU:** An LLDPDU containing an Extension Request TLV and not containing a Time to Live TLV or Extension Identifier TLV.

**Link or Logical Link:** A logical connection between two logical ports on two devices. E.g. two VLANs between the same two ports are two links.

**Logical Link Endpoint:** One end of a logical link.

**Logical Link Endpoint Identifier:** The unique identifier for a Logical Link Endpoint.

**LLDP Protocol Data Unit:** A single L2 frame carrying LLDP TLVs.

**Manifest LLDPDU:** A Normal LLDPDU containing a Manifest TLV.

Normal LLDPDU: A LLDPDU containing a Time To Live TLV and not containing an Extension Request TLV or an Extension Identifier TLV.

Scope MAC Address: The MAC address used as the destination address in Normal LLDPDUs, which determines the limits of propagation in the network.

3. Abbreviations

- LLDP - Link Layer Discovery Protocol
- LLDPDU: LLDP Protocol Data Unit
- LLEI: Logical Link Identifier
- TLV - Type-Length-Value
- XLLDP - Extended Link Layer Discovery Protocol
- XPDU - Extension LLDPDU
- XREQ - Extension Request LLDPDU

4. LLDP Operation

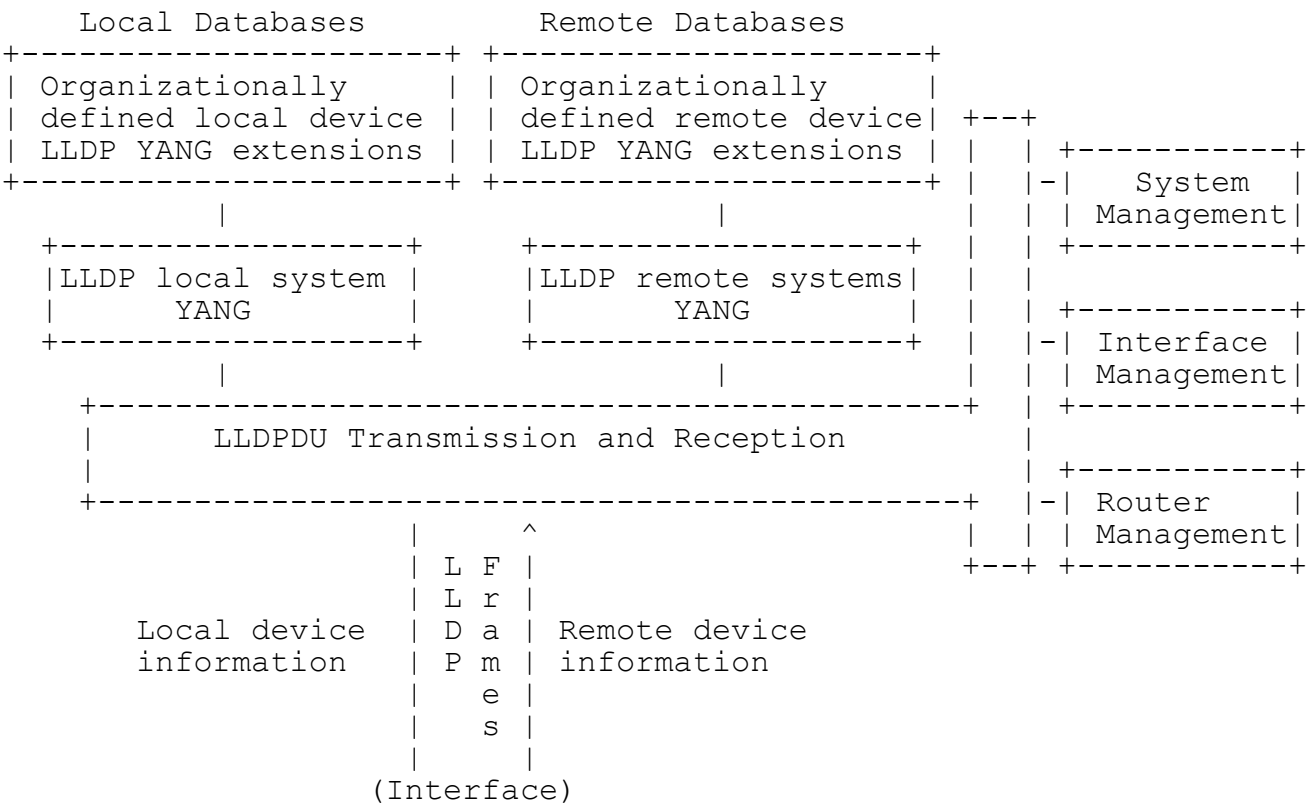


Figure 1 LLDP Agent and databases

Figure 1 is a simplified illustration of figure 6-1 from [802.1AB]. It illustrates an LLDP Agent and its relationship to its databases. For a more detailed description refer to clause 6 of [802.1AB].

LLDP is a link layer protocol that allows an Ethernet device to advertise the capabilities and status of a system interface. The local database (see figure 1) is loaded with the information to be advertised by the interface. The LLDP Agent at the interface periodically encodes the local database into LLDPDUs and transmits a Normal LLDPDU to the Scope MAC Address destination. The interfaces subscribed to the Scope MAC Address receiving incoming LLDPDUs direct them to the LLDP Agent based on the LLDPDU destination address and the Ethertype 0x88CC. The information received in LLDPDUs is stored by the LLDP Agent in the remote database. The local database contains the information advertised by LLDP for the interface, while the remote database contains the information received from each remote interface.

In the event the TLVs that encode the local database do not fit in a single LLDPDU, the LLDP Agent will encode additional TLVs in extension LLDPDUs (XPDUs) and will add a Manifest TLV to the Normal LLDPDU. The Manifest TLV contains a digest of each XPDUs. The receiver of a Normal TLV containing a Manifest TLV will decide if any of the XPDUs contains new information by comparing the digest for its existing XPDUs with the digest contained in the Manifest. If new XPDUs are found the receiver will make explicit requests for the new XPDUs using the Extension Request LLDPDU (XREQ). The transmitting LLDP Agent responds to XREQs by sending the requested XPDUs. Since the receiver makes explicit requests for XPDUs it can determine the pacing rate based on its processing overhead and resources. The receiver is responsible for timing out and retrying failed XREQs.

#### 4.1. Nearest Router Scope MAC Address for LLDP

There are two choices for how to address LLDPDUs for application to LSVR. These are to use the Nearest Bridge Address (01-80-C2-00-00-0E) [802.1AB] for the LLDP Scope MAC Address or to select an application specific Nearest Router Multicast Address for use as the LLDP Scope MAC Address.

If the IEEE standard nearest bridge address is used, the LLDP Agent will be at the physical switch port. In this case, routers attached to subinterfaces or Switched Virtual Interfaces (SVIs) will need to support information for multiple Logical Link Endpoints within the local LLDP databases on each physical port. The LLDP local databases will require the Logical Link Endpoint Identifier (LLEI), VLAN IDs, and MAC addresses associated with each Logical Link Endpoint. In addition, if the nearest bridge address is used, router interfaces with bridges between them will not be discoverable because the LLDPDUs will not travel through the bridges.

If a LSVR specific Nearest Router Multicast Address which is not a member of the reserved bridge addresses is selected for the LLDP Scope MAC Address, then the LLDP Agents can be associated with the router interfaces regardless of whether they are physical, subinterfaces, or switched virtual interfaces (SVIs). In addition, such an address will allow LLDP to operate when a bridge intervenes between the router interfaces.

Within this document a Nearest Router Multicast will be used as the only LLDP Scope MAC Address. We believe this approach is consistent with the current intent of draft-ietf-lsvr-l3dl. The Nearest Router Multicast Address for the LLDP Scope MAC Address should be reserved for router interface discovery applications by IANA and IEEE.

#### 4.2. Logical Link Endpoint Identification

Every LLDPDU begins with a ChassisID TLV and a PortID TLV [802.1AB]. These two TLVs identify the sending LLDP location which is the equivalent of the draft-ietf-lsvr-l2dl Logical Link Endpoint Identifier. The ChassisID provides a unique identifier for the system, in this case an LSVR router (i.e. a MAC address, any IANA address family identifier). The PortID provides a unique identifier for the interface within the system (i.e. port number, IfIndex).

Within this draft we assume the use of a Nearest Router Multicast Scope MAC Address for LLDP which always places LLDP Agents at the Logical Link Endpoint. Provided the LLDP Agent is advertising information from each router interface (physical or virtual) the ChassisID and PortID can be used as the Logical Link Endpoint Identifier. This draft assumes the LLDP ChassisID and PortID are used to identify Logical Link Endpoints.

#### 4.3. Link Liveliness

Every Normal LLDPDU contains Time-To-Live TLV which specifies how long the information advertised for each Logical Link Endpoint is valid. LLDP retransmits the Normal LLDPDU periodically to maintain the liveliness of the information in the remote LLDP databases. The default retransmit time for LLDP is once every 30 seconds and is settable in 1 second steps. The Time-To-Live is calculated based on the Normal LLDPDU retransmit window to allow multiple LLDPDU retransmits before expiring. A time-to-live timeout will be indicated to the coupled application layer, in this case the router, which can take action to disable the associated Logical Link(s).

#### 4.4. LLDP Organizationally Specific TLVs

IEEE Std 802.1AB defines the format of the Organizationally Specific TLVs. The format is redrawn here for convenience.

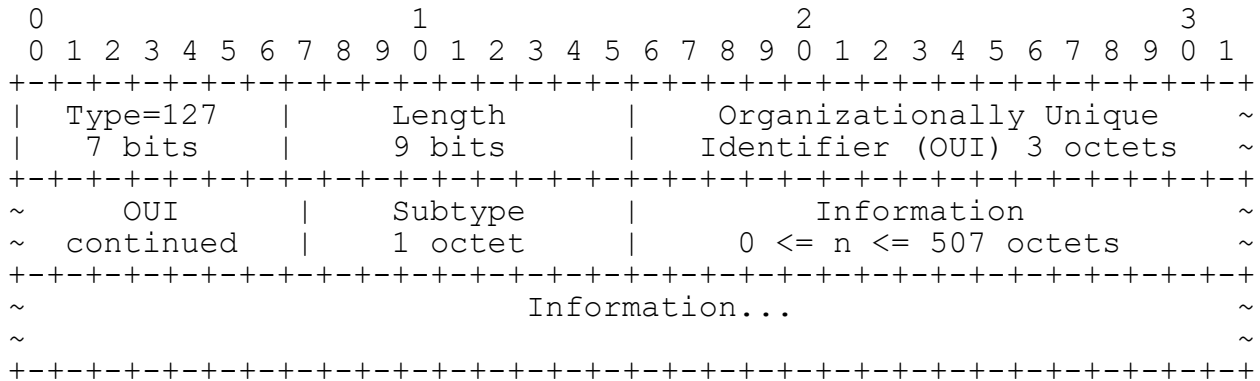


Figure 2 LLDP Organizationally Specific TLV Format

#### 4.4.1. Type

The Type field is 7 bits in length and is set to the value of 127 indicating the TLV type is an Organizationally Specific TLV.

#### 4.4.2. Length

The Length field contains the length of the TLV specific information, the OUI and subtype. The Length is from 4 to 511 octets.

#### 4.4.3. Organizationally Unique Identifier

The Organizationally Unique Identifier indicates the organization specifying the TLV. For IETF use, this field is the IANA OUI as discussed in RFC 7042 and has the hex value of 00-00-5E.

#### 4.4.4. Subtype

The Subtype field indicates the specific IETF Organizationally Specific TLV. The format of the information string is dependent upon the Subtype field.

#### 4.4.5. Information

The Information field contains the octets that are specific to the Organizationally Specific TLV. The field length is between 0 and 507 octets.

### 4.5. IETF Organizationally Specific TLVs for LSVR

This section defines the format for the IETF Organizationally Specific TLVs. The contents of the TLVs are modeled after the TLV PDUs of the L3DL protocol specified by the LSVR working group in [I-D.ietf-lsvr-l3dl].

Each TLV definition starts with the same first four fields that are described once here. The Type field SHALL be 127; the Length field SHALL be the length in octets of the LLDP TLV information string which starts immediately after the Length field; the OUI field SHALL be 00-00-0E and the Subtype field SHALL be one of the named values shown in the following table:

Subtype -----	Name -----
0	Reserved
1	IPv4 Announcement
2	IPv6 Announcement
3	MPLS IPv4 Announcement
4	MPLS IPv6 Announcement
5	Upper-Layer Protocol Configuration
6-255	Reserved

The specific TLV format for each subtype and remaining fields of each IETF Organizationally Specific TLV for LSVR are specified in the following sections.

#### 4.6. LSVR LLDP L3 data link application

The LSVR LLDP Agent advertises the local and updates the remote databases which describe the local and remote information advertised for router's logical link endpoints. An L3 data link application is used to access the LLDP databases and to determine the network configuration and operation. The application using the LLDP databases signals to the LLDP Agent whenever it changes something in the local database by using the `SomethingChangedLocal()` routine. The LLDP Agent signals the application whenever a change is detected in the remote database using the `SomethingChangedRemote()` routine.

The L3 data link application uses the information in the LLDP databases to determine active logical links and operating configurations. The application informs the LSVR router of operational links and parameters.

#### 5. Encapsulation TLVs for LSVR

LLDP advertises the encapsulations available in the local database associated with its interface. The interface is uniquely identified by the ChassisID and PortID TLVs encoded as the first two TLVs within each LLDPDU. The ChassisID and PortID TLVs together encode the Logical Link Endpoint Identifier for the router interface within the sending system associated with the LLDP Agent.

The encapsulation types the peers can exchange are IPv4, IPv6, MPLS IPv4, MPLS IPv6, and/or possibly others not defined here.

The sender of an encapsulation type MUST NOT assume the peer is capable of the same encapsulation type. Only if an interface has the encapsulation type in both the local and remote LLDP databases is it



safe for layer-3 protocols to assume they are compatible for that type.

The LLDP L3 data link application might recognize an addressing conflict, such as both ends of the link trying to use the same address. As there can be other usable addresses or encapsulations, this error might be logged, letting an upper layer topology builder deal with what works.

Further, to consider a L3 data link of an encapsulation type to formally be established, so it can be pushed up to upper layer protocols, the addressing for the type must be compatible, e.g. on the same IP subnet.

### 5.1. Encaps Flags

Each Encapsulation TLV specifies a set of flags for each address listed for the encapsulation. These flags are specified here and used in all the encapsulations.

The Encaps Flags specify attributes about the address being announced and are compatible with the Encapsulation Flags defined [I-D.ietf-lsvr-l3dl]. The Encaps Flags have the following format:

0	1	2	3	4 ... 7
-----+	-----+	-----+	-----+	-----+
Announce	Primary	Under/Over	Loopback	Reserved
-----+	-----+	-----+	-----+	-----+

Figure 3 Encapsulation Flags

#### 5.1.1. Announce

The Announce flag bit SHALL always be set to 1. [I-D.ietf-lsvr-l3dl] defines this bit as an Announce/Withdraw bit to allow encapsulations to be both announced and withdrawn by the L3DL protocol. Old encapsulations are withdrawn in LLDP by simply excluding them in the next transmission of an LLDPDU.

#### 5.1.2. Primary

The Primary flag bit indicates that the encapsulation interface is a primary interface. The value of 1 indicates the interface is primary and the value of 0 indicates it is not. If the LLEI has multiple addresses for an encapsulation type, one and only one address SHOULD be marked as primary.

#### 5.1.3. Under/Over

The Under/Over flag bit indicates if the interface address is for an underlay interface or an overlay interface. The value of 1 indicates the interface is an underlay and the value of 0 indicates it is an overlay.

#### 5.1.4. Loopback

The Loopback flag bit indicates if the interface address is a loopback address. Loopback addresses are generally not seen directly on an external interface. One or more loopback addresses MAY be announced. The value of 1 indicates that the interface address is a loopback address and the value of 0 indicates it is not.

#### 5.1.5. Reserved

MUST be transmitted as 0 and ignored on receive.

#### 5.2. IPv4 Announcement TLV

The IPv4 Announcement TLV describes a device's ability to exchange IPv4 packets on one or more subnets. It does so by announcing the interface's addresses and the corresponding prefix lengths. Multiple tuples of IPv4 address information can be announced in a single TLV, each tuple with its own set of encapsulation flags, address and prefix length. The format of the IPv4 Announcement TLV is as follows:

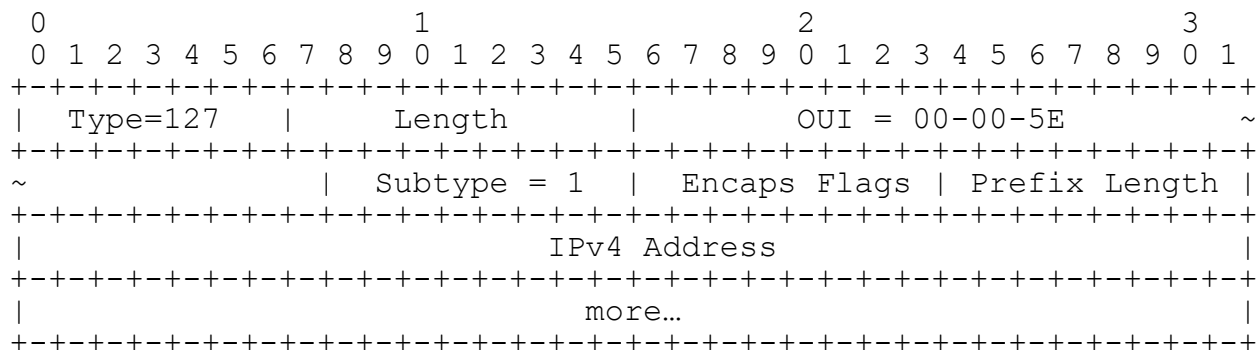


Figure 4 IPv4 Announcement TLV

##### 5.2.1. Encaps Flags

The Encaps Flags are the same flags specified in section 5.1.

##### 5.2.2. Prefix Length

The Prefix Length field indicates the number of bits in the IPv4 Address that represent the IPv4 prefix for the subnet (network portion of the address).

##### 5.2.3. IPv4 Address

The IPv4 Address field holds the IPv4 address to be announced.



#### 5.3.4. IPv6 Announcement TLV usage rules

An LLDPDU MAY contain multiple IPv6 Announcement TLVs, however, the encapsulation flags, IPv6 address and prefix length tuple for a specific IPv6 address MUST only appear once in the LLDPDU.

#### 5.4. MPLS IPv4 Announcement TLV

The MPLS IPv4 Encapsulation describes a logical link's ability to exchange labeled IPv4 packets on one or more subnets. It does so by stating the interface's addresses the corresponding prefix lengths, and the corresponding labels which will be accepted for each address. The format of the MPLS IPv4 Announcement TLV is as follows:

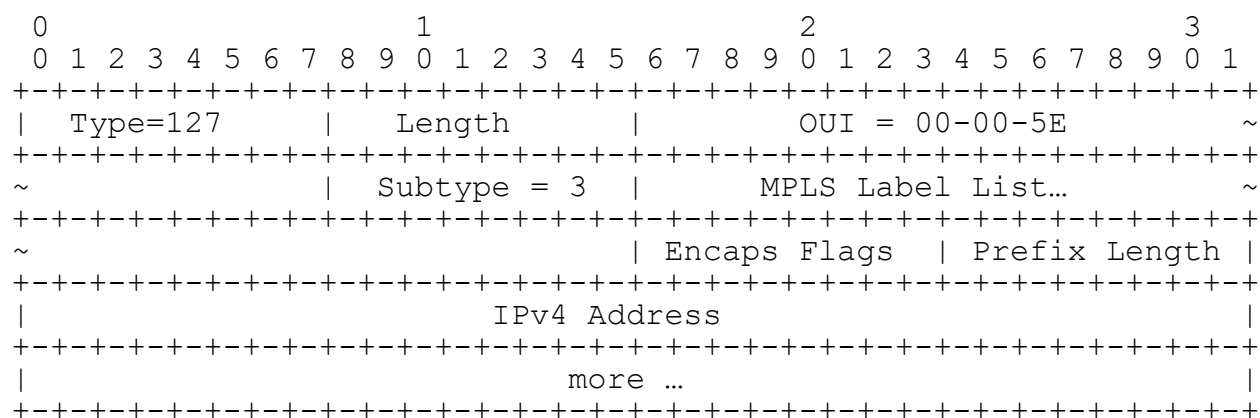


Figure 6 MPLS IPv4 Announcement TLV

##### 5.4.1. MPLS Label List

The MPLS Label List is a variable length field that contains the label stack, see [RFC3032], that the sender will accept for the prefix to which the list is attached. The format of the MPLS Label List is as follows:




Figure 7 MPLS Label List

##### 5.4.1.1. Label Count

The number of label stack entries that are included in the list. Each label stack entry has a label, experimental use bits and the bottom of stack indicator. The Label Count MUST NOT be 0.

<<Editor's note: We really don't need this count because the S bit indicates when the list ends>>

#### 5.4.1.2. Label

The Label field is the 20-bit value of the label.

#### 5.4.1.3. Exp

The Exp field is a 3-bit value reserved for experimental use.

#### 5.4.1.4. S

The S field is a single bit that indicates the last entry in the label stack. The bit is set to 1 for the last entry and 0 for all other entries in the list.

#### 5.4.2. Encaps Flags

The Encaps Flags are the same flags specified in section 5.1.

#### 5.4.3. Prefix Length

The Prefix Length field indicates the number of bits in the IPv4 Address that represent the IPv4 prefix for the subnet (network portion of the address).

#### 5.4.4. IPv4 Address

The IPv4 Address field holds the IPv4 address to be announced.

#### 5.4.5. IPv4 Announcement TLV usage rules

An LLDPDU MAY contain multiple MPLS IPv4 Announcement TLVs, however, the encapsulation flags, label list, IPv4 address and prefix length tuple for a specific IPv4 address MUST only appear once in the LLDPDU.

#### 5.5. MPLS IPv6 Announcement TLV

The MPLS IPv6 Encapsulation describes a logical link's ability to exchange labeled IPv6 packets on one or more subnets. It does so by stating the interface's addresses the corresponding prefix lengths, and the corresponding labels which will be accepted for each address. The format of the MPLS IPv6 Announcement TLV is as follows:

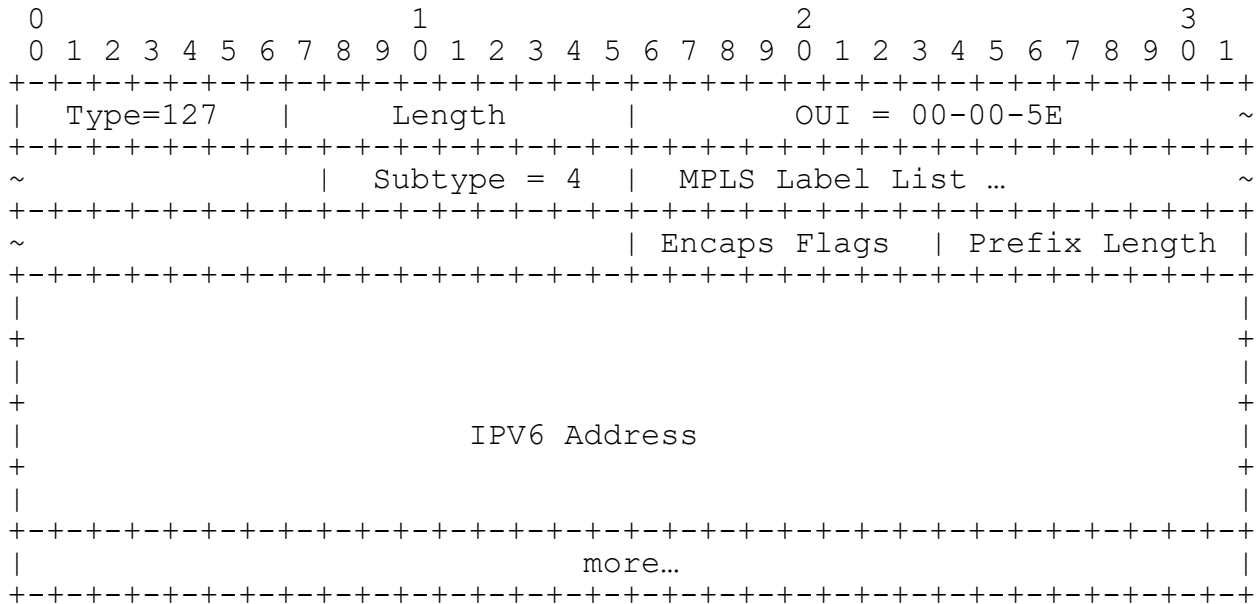


Figure 8 MPLS IPv6 Announcement TLV

#### 5.5.1. MPLS Label List

The MPLS Label List is the same as specified in section 5.4.1.

#### 5.5.2. Encaps Flags

The Encaps Flags are the same flags specified in section 5.1.

#### 5.5.3. IPv6 Address

The IPv6 Address field holds the IPv6 address to be announced.

#### 5.5.4. Prefix Length

The Prefix Length field indicates the number of bits in the IPv6 Address that represent the IPv6 prefix for the subnet (network portion of the address).

#### 5.5.5. MPLS IPv6 Announcement TLV usage rules

An LLDPDU MAY contain multiple MPLS IPv6 Announcement TLVs, however, the encapsulation flags, label list, IPv6 address and prefix length tuple for a specific IPv6 address MUST only appear once in the LLDPDU.

### 6. Upper-Layer Protocol Configuration TLVs

A router or switch supporting LSVR protocols MAY include the BGP Attribute TLVs. The format of the BGP Protocol Configuration TLVs is shown in Figure 2.

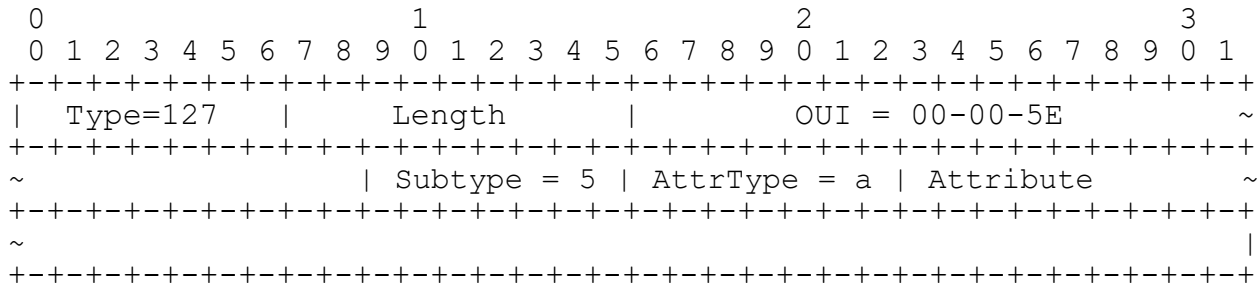


Figure 9 BGP Protocol Configuration TLV Format

### 6.1.1. AttrTypes

The subtype identifies the BGP attribute type. The types are:

- AttrType = 1: BGP ASN
- AttrType = 2: BGP IPv4 Peering Address
- AttrType = 3: BGP IPv6 Peering Address
- AttrType = 4: BGP Authentication Data
- AttrType = 5: BGP Flags, Bit 0: GTSM, Bit 1: BFD,  
Bit 2-15 reserved = 0

### 6.1.2. BGP ASN

The four octet Autonomous System number MUST be specified. If the AS Number is less than 32 bits, it is padded with high order zeros.

### 6.1.3. BGP IPv4 Peering Address

The five octet IPv4 Peering Address along with the Prefix Length.

### 6.1.4. BGP IPv6 Peering Address

The 17 octet IPv6 Peering Address along with the Prefix Length.

### 6.1.5. BGP Authentication Data

The BGP Authentication TLV provides any authentication data needed to OPEN the BGP session. Depending on operator configuration of the environment, it might be a simple MOS key (see [RFC2385]), the name of a key chain in a KARP database (see [RFC7210]), or one of multiple Authentication TLVs to support[RFC4808].

### 6.1.6. BGP Flags

The BGP session OPEN has extensive, and a bit complex, capability negotiation facilities. In case one or more extra attributes might be needed, the two octet BGP Flags TLV may be used. No flags are currently defined.

BGP Flags:

- Bit 0: GTSM

Bit 1: BFD  
Bit 2-15: Must be zero

The GTSM flag, when 1, indicates that the sender wishes to enable the [RFC5082] Generalized TTL Security Mechanism for the session.

The BFD flag, when 1, indicates that the sender wishes to enable the [RFC5880] Bidirectional Forwarding Detection for the session.

## 7. References

- [802.1AB] "IEEE Standard for Local and metropolitan area networks—Station and Media Access Control Connectivity Discovery", IEEE Std 802.1AB-2016, 29 January 2016.
- [802.1ABcu] "IEEE Standards for Local and metropolitan area networks— Station and Media Access Control Connectivity Discovery Amendment 1: YANG Data Model" IEEE Std 802.1ABcu-2021, December 2021.
- [802.1ABdh] "IEEE Standards for Local and metropolitan area networks— Station and Media Access Control Connectivity Discovery Amendment 2: Support for Multiframe Protocol Data Units", IEEE Std 802.1ABdh-2021, December 2021.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T. and Conta, A. "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [RFC7042] Eastlake, D. and Abley, J. "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", RFC 7042, October 2013.
- [I-D.ietf-lsvr-l3dl] Bush, R., Austein, R. and Patel, K., "Layer 3 Discovery and Liveness", draft-ietf-lsvr-l3dl-14.txt (work in progress), July 2019.