

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2017

A. Lindem
Cisco Systems
K. Patel
Arrcus, Inc
S. Zandi
LinkedIn
March 13, 2017

BGP Logical Link Discovery Protocol (LLDP) Peer Discovery
draft-acee-idr-lldp-peer-discovery-00.txt

Abstract

Link Layer Discovery Protocol (LLDP) or IEEE 802.1AB is implemented in networking equipment from many vendors. It is natural for IETF protocols to avail this protocol for simple discovery tasks. This document describes how BGP would use LLDP to discover directly connected and 2-hop peers when peering is based on loopback addresses.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 14, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Requirements Notation	2
2.	LLDP Extensions	3
2.1.	LLDP Organizationally Specific TLV Format	3
2.2.	Local IP Address OS-TLV Format	4
2.3.	BGP Config OS-TLV Format	5
2.3.1.	BGP Config OS-TLV - Peering Address Sub-TLV	6
2.3.2.	BGP Config OS-TLV - BGP Local AS Sub-TLV	7
2.3.3.	BGP Config OS-TLV - BGP Capabilities Sub-TLV	8
3.	BGP LLDP Peer Discovery Operations	9
3.1.	Advertising BGP Speaker	9
3.2.	Receiving BGP Speaker	9
4.	Security Considerations	10
5.	IANA Considerations	10
5.1.	IANA Assigned LLDP Subtype	10
5.2.	BGP Config LLDP OS-TLV Sub-TLVs	11
6.	References	11
6.1.	Normative References	11
6.2.	Informative References	12
Appendix A.	Acknowledgments	12
	Authors' Addresses	12

[1.](#) Introduction

Link Layer Discovery Protocol (LLDP) [[LLDP](#)] or IEEE 802.1AB is implemented in networking equipment from many vendors. It is natural for IETF protocols to avail this protocol for simple discovery tasks. This document describes how BGP [[BGP](#)] would use LLDP to discover directly connected and 2-hop peers when peering is based on loopback addresses.

[1.1.](#) Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC-KEYWORDS](#)].

2. LLDP Extensions

2.1. LLDP Organizationally Specific TLV Format

The format of the LLDP Basic Organizationally Specific TLV (OS-TLV) is defined in [LLDP]. It is shown below for completeness.

0																1																2																3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																																
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+															
Type (127)																Length																OUI (3 Octets) 00-00-5E																															
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+															
OUI Continued																Subtype																Value																															
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+															
																...																(Up to 507 Octets)																															

Type	Organizationally Specific TLV type value, 127.
------	--

Length The length of the remainder of the TLV.

OUI Organizationally unique identifier for the organization's OUI. For the IANA, this is value is 00-00-5E as specified in [[RFC7042](#)].

Subtype IETF specific subtype

Value	Value for organizationally specific TLV. The Length of the value is 4 octets less than the TLV length.
-------	--

LLDP Organizationally Specific TLV

The OUI for the IANA was allocated in section 1.4 of [\[IEEE-802-IANA\]](#). This document requests creation of a registry for IETF specific subtypes for LLDP Organizationally Specific TLVs.

2.3. BGP Config OS-TLV Format

The BGP Config Organizationally Specific TLV (OS-TLV) will be used to advertise BGP configuration information. The configuration information will be composed of Sub-TLVs. Since the length is limited to 507 octets, multiple BGP Config OS-TLVs could be included in a single LLDP advertisement.

0																1																2																3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																						
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+															
Type (127)																Length																OUI (3 Octets) 00-00-5E																															
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+															
OUI Continued																2																BGP Config Sub-TLVs ...																															
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+															
																...																(Up to 507 Octets)																															

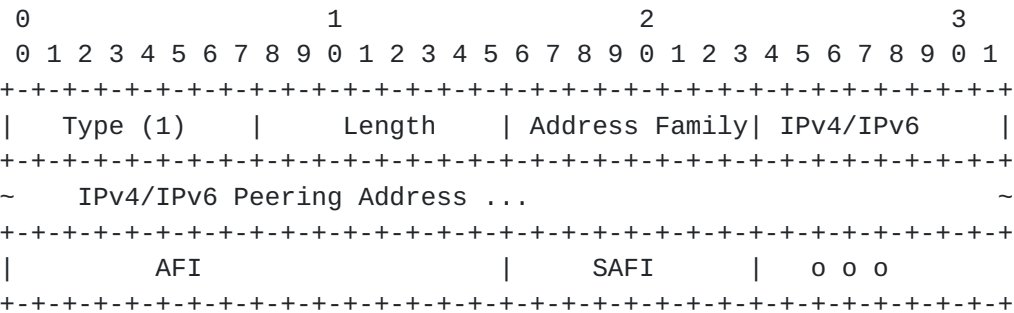
Length	The length of the BGP TLV.
--------	----------------------------

Subtype	IETF specific subtype for BGP Config OS-TLV. The value shall be 2.
---------	--

Value	BGP Config Sub-TLVs each with a 1 byte Type and Length. The Length will include solely the value portion of the TLV and not the Type and Length fields themselves.
-------	--

2.3.1. BGP Config OS-TLV - Peering Address Sub-TLV

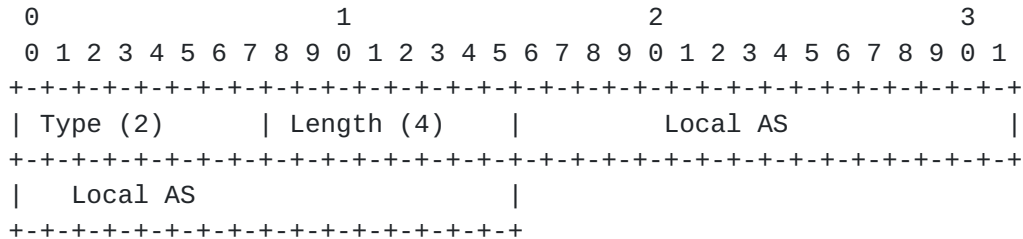
The BGP OS-TLV Peering Address Sub-TLV will be used to advertise the local IP address used for BGP sessions and the associated address families. The format of the BGP Peering Address Sub-TLV is shown below.



- Type The Sub-TLV Type value shall be 1.
- Length The Sub-TLV length in octets will be 4 for IPv4 or 16 for IPv6 plus 3 times the number of AFI/SAFI pairs.
- Address Family IANA Address family (1 for IPv4 or 2 for IPv6)
- Peering Address An IPv4 address (4 octets) or an IPv6 address (16 octet) peering address.
- AFI/SAFI Pairs One or more AFI/SAFI pairs for BGP session using this peering address.

2.3.2. BGP Config OS-TLV - BGP Local AS Sub-TLV

The BGP Config OS-TLV Local AS Sub-TLV will be used to advertise the 4-octet local Autonomous System (AS) number. The format of the BGP Local AS Sub-TLV is shown below.



Type The Sub-TLV Type value shall be 2.

Length The Sub-TLV Length will be 4 octets.

Local AS Local Autonomous System (AS)

2.3.3. BGP Config OS-TLV - BGP Capabilities Sub-TLV

The BGP Config OS-TLV Capabilities Sub-TLV will be used to advertise an 8-octet Capabilities field. The capabilities are represented as bit flags identifying the supported BGP capabilities. The format of the BGP Capabilities Sub-TLV is shown below.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Type (3)      | Length (8)    |      Capabilities      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Capabilities      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Capabilities      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type The Sub-TLV Type value shall be 3.

Length The Sub-TLV Length will be 8 octets.

Capabilities Bit fields identify BGP capabilities

The BGP Capabilities is an 8-octet bit field. The most significant bit is the first bit (Bit 1) of the Capabilities. The following bits are defined:

Bit 1: This bit indicates that support for TCP MD5 authentication [[TCP-MD5](#)].

Bit 2: This bit indicates support for TCP-AO authentication [[TCP-AO](#)].

Bit 3: This bit indicates support for Generalized TTL Security Mechanism (GTSM) [[GTSM](#)] with a configured TTL range of 254-255.

TCP MD5 authentication is described in [[TCP-MD5](#)]. The TCP Authentication Option (TCP-AO) is described in [[TCP-AO](#)]. The Generalized TTL Security Mechanism (GTSM) is described in [[GTSM](#)]. If both TCP MD5 authentication and TCP-AO authentication are specified and TCP-AO is supported, it will take precedence.

3. BGP LLDP Peer Discovery Operations

The simple use case is to just use the peer address advertised in the LLDP Packet Data Unit (PDU) to establish a 1-hop BGP peer session. This can be used in data centers using BGP as described in [[BGP-DC](#)]. The more complex use case is when a loopback address is advertised as the peering address in the LLDP PDU.

3.1. Advertising BGP Speaker

A BGP speaker MAY advertise its BGP peering address in an LLDP PDU for a link using the BGP Local Address Sub-TLV of BGP-OS TLV Format. This can be an IPv4 or IPv6 address local to the link for 1-hop peering or a loopback address for 2-hop peering.

Additionally, a BGP speaker MAY advertise one or more local addresses IPv4 and IPv6 addresses. In the case of 2-hop peering, a local address on the link can be used as a next-hop for the peering address. In this manner both the peering address and reachability can be discovered.

A BGP speaker MAY advertise its local AS number using BGP AS Sub-TLV of BGP-OS TLV format. It may also announce relevant capabilities using BGP Capabilities Sub-TLV of BGP-OS TLV format.

3.2. Receiving BGP Speaker

A BGP speaker configured for LLDP peer discovery will attempt to establish peers using the address in the BGP Local Address Sub-TLV of BGP-OS TLV format. If the peering address is directly accessible over the link on which the LLDP PDU was received, the BGP speaker will attempt to establish a 1-hop BGP session with the peer.

If the received BGP Peering Address is not directly accessible over the link, the BGP speaker may add a route to access the BGP peer. The next-hop for the route MAY be one of the addresses the BGP speaker has advertised in the Local IP Address OS-TLV. If the BGP speaker receives the same BGP peering address in LLDP PDU on multiple links, it will not establish multiple sessions. Rather a single 2-hop session will be established. Optionally, ECMP routes are added to the BGP peering session over each link on which an LLDP PDU containing the same Peering Address is received.

A BGP speaker MAY receive remote neighbor's local AS number in LLDP in BGP AS Sub-TLV of the BGP-OS TLV. A BGP speaker MAY use the received local AS number to perform validation check of AS received in the OPEN message. Furthermore, A BGP speaker MAY receive remote neighbor's capabilities in LLDP in BGP Capabilities Sub-TLV of the

BGP-OS TLV. A BGP speaker MAY use the received capabilities to ensure appropriate neighbor based configuration is done locally so as to facilitate the session establishment.

4. Security Considerations

This security considerations for BGP [BGP] apply equally to this extension.

Additionally, BGP peering address discovery should only be done on trusted links (e.g., in a data center network) since LLDP packets are not authenticated or encrypted [LLDP].

5. IANA Considerations

5.1. IANA Assigned LLDP Subtype

IANA is requested to create a registry for IANA assigned subtypes in the Organizationally Specific TLV assigned to IANA (OUI of 000-00-53 [IEEE-802-IANA]. Assignment is requested for 1 for the Local IP Address OS-TLV. Assignment is also requested for 2 for the BGP Config OS-TLV.

Range	Assignment Policy
0	Reserved (not to be assigned)
1	Local IP Address
2	BGP Configuration
3-127	Unassigned (IETF Review)
128-254	Reserved (Not to be assigned now)
255	Reserved (not to be assigned)

IANA LLDP Organizationally Specific TLV Sub-Types

- o Types in the range 3-127 are to be assigned subject to IETF Review. New values are assigned only through RFCs that have been shepherded through the IESG as AD-Sponsored or IETF WG Documents [IANA-GUIDE].
- o Types in the range 128-254 are reserved and not to be assigned at this time. Before any assignments can be made in this range,

there MUST be a Standards Track RFC that specifies IANA Considerations that covers the range being assigned.

5.2. BGP Config LLDP OS-TLV Sub-TLVs

IANA is requested to create a registry for Sub-TLVs of the BGP Config LLDP OS-TLV. Assignment is requested for 1 for the BGP Peering Address Sub-TLV. Assignment is also requested for 2 for the Local AS Sub-TLV. Additionally, assignment is requested for 3 for the Capabilities Sub-TLV.

Range	Assignment Policy
0	Reserved (not to be assigned)
1	Peering Address
2	Local AS
3	Capabilities
4-127	Unassigned (IETF Review)
128-254	Reserved (Not to be assigned now)
255	Reserved (not to be assigned)

LLDP BGP Config OS-TLV Types

- o Types in the range 4-127 are to be assigned subject to IETF Review. New values are assigned only through RFCs that have been shepherded through the IESG as AD-Sponsored or IETF WG Documents [[IANA-GUIDE](#)].
- o Types in the range 128-254 are reserved and not to be assigned at this time. Before any assignments can be made in this range, there MUST be a Standards Track RFC that specifies IANA Considerations that covers the range being assigned.

6. References

6.1. Normative References

- [BGP] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

- [BGP-DC] Lapukhov, P., Premji, A., and J. Mitchell, "BGP Routing in Data Centers", [RFC 7938](#), August 2016.
- [LLDP] IEEE, "IEEE Standard for Local and metropolitan area networks-- Station and Media Access Control Connectivity Discovery Corrigendum 2: Technical and Editorial Corrections", IEEE 802.1AB-2009/Cor 2-2015, DOI 10.1109/ieeestd.2015.7056401, March 2015.
- [RFC-KEYWORDS] Bradner, S., "Key words for use in RFC's to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[6.2. Informative References](#)

- [GTSM] Gill, V., Heasley, J., Savola, P., and C. Pignataro, "The Generalized TTL Security Mechanism", [RFC 5082](#), October 2007.
- [IANA-GUIDE] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [RFC 5226](#), May 2008.
- [IEEE-802-IANA] Eastlake, D. and J. Abley, "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", [RFC 7042](#), October 2013.
- [TCP-AO] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", [RFC 5925](#), June 2010.
- [TCP-MD5] Heffernan, A., "TCP MD5 Signature Option", [RFC 2385](#), August 1998.

[Appendix A. Acknowledgments](#)

The RFC text was produced using Marshall Rose's xml2rfc tool.

Authors' Addresses

Acee Lindem
Cisco Systems
301 Midenhall Way
Cary, NC 27513
USA

Email: acee@cisco.com

Keyur Patel
Arrcus, Inc

Email: keyur@arrcus.com

Shawn Zandi
LinkedIn
222 2nd Street
San Francisco, CA 94105
USA

Email: szandi@linkedin.com