

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 4, 2018

A. Lindem
Cisco Systems
K. Patel
Arrcus, Inc
S. Zandi
LinkedIn
J. Haas
Juniper Networks, Inc
X. Xu
Huawei
July 3, 2017

**BGP Logical Link Discovery Protocol (LLDP) Peer Discovery
draft-acee-idr-lldp-peer-discovery-01.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 January 4, 2018.

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 [3](#)
- [2.](#) LLDP Extensions [3](#)
- [2.1.](#) LLDP Organizationally Specific TLV Format [3](#)
- [2.2.](#) BGP Config OS-TLV Format [4](#)
- [2.2.1.](#) BGP Config OS-TLV - Peering Address Sub-TLV [5](#)
- [2.2.2.](#) BGP Config OS-TLV - BGP Local AS Sub-TLV [6](#)
- [2.2.3.](#) BGP Config OS-TLV - BGP Identifier Sub-TLV [7](#)
- [2.2.4.](#) BGP Config OS-TLV - Session Group-ID Sub-TLV [8](#)
- 2.2.5. BGP Config OS-TLV - BGP Session Capabilities Sub-TLV [9](#)
- [2.2.6.](#) BGP Config OS-TLV - Key Chain Sub-TLV [10](#)
- [3.](#) BGP LLDP Peer Discovery Operations [11](#)
- [3.1.](#) Advertising BGP Speaker [11](#)
- [3.2.](#) Receiving BGP Speaker [11](#)
- [4.](#) Security Considerations [12](#)
- [5.](#) IANA Considerations [13](#)
- [5.1.](#) IANA Assigned LLDP Subtype [13](#)
- [5.2.](#) BGP Config LLDP OS-TLV Sub-TLVs [13](#)
- [6.](#) Contributors [14](#)
- [7.](#) References [14](#)
- [7.1.](#) Normative References [14](#)
- [7.2.](#) Informative References [15](#)
- [Appendix A.](#) Acknowledgments [15](#)
- Authors' Addresses [16](#)

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.

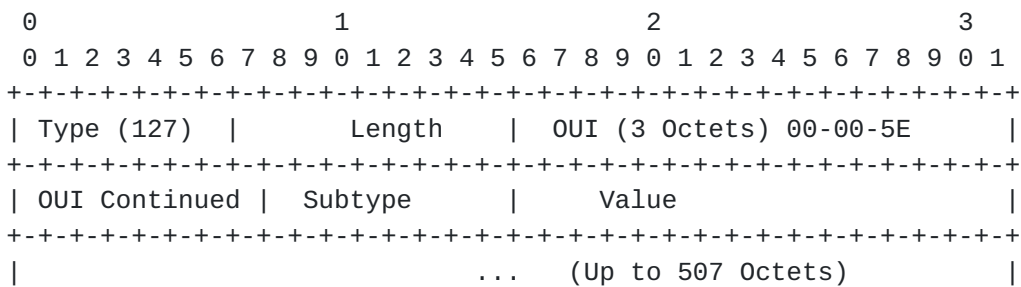
1.1. Requirements Notation

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.



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 IANA, this is value is 00-00-5E as specified in [IEEE-802-IANA].

Subtype IETF specific subtype

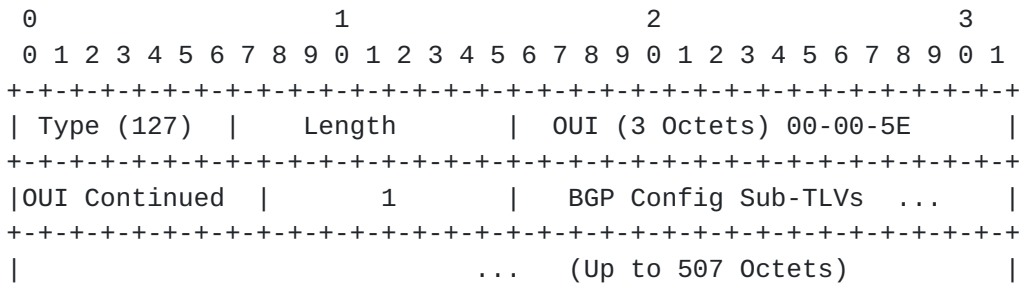
Value Value for organizationally specific TLV. The Length of the value is 4 octets less than the TLV length.

LLDP Organizationally Specific TLV

The OUI for 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.

2.2. 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.

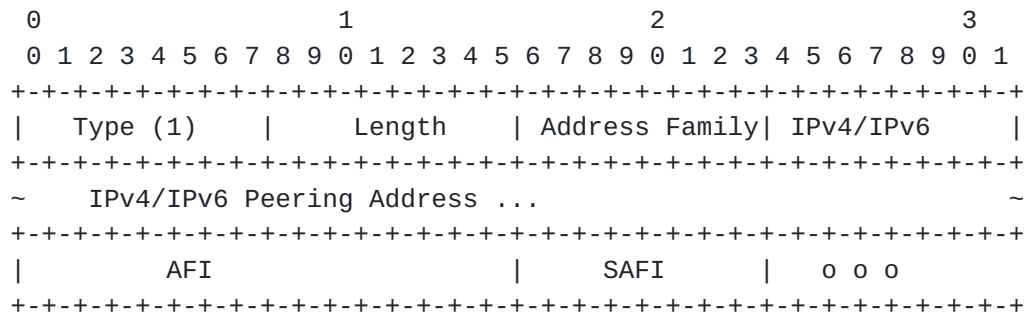


- Length The length of the BGP TLV.
- Subtype IETF specific subtype for BGP Config OS-TLV. The value shall be 1.
- Value BGP Config Sub-TLVs each with a 1 byte Type and Length. The Length will include solely the value portion of the TLV and not the Type and Length fields themselves.

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

The BGP OS-TLV Peering Address Sub-TLV will be used to advertise the local IP addresses used for BGP sessions and the associated address families specified by AFI/SAFI tuples. The AFI/SAFI tuple, 0/0, indicates to use the associated peering address for all locally configured address families without an explicit peering address specification. As always, the address families supported for a given BGP session will be determined during capabilities negotiation [MP-BGP]. It is RECOMMENDED that the wildcard AFI/SAFI be used in deployments with fairly homogenous address family usage.

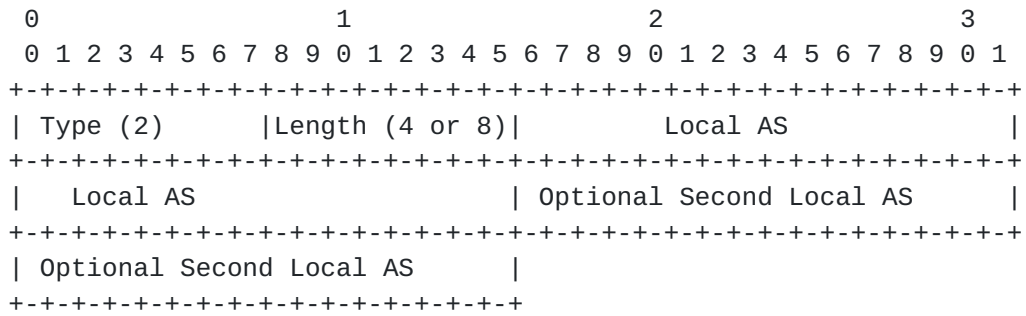
The format of the BGP Peering Address Sub-TLV is shown below.



- 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 tuples.
- Address Family IANA Address family (1 for IPv4 or 2 for IPv6)
- Peering Address An IPv4 address (4 octets) or an IPv6 address (16 octets)
- AFI/SAFI Pairs One or more AFI/SAFI tuples for BGP session using this peering address. The AFI/SAFI tuple, 0/0, is a wildcard indicating to attempt negotiation for all AFI/SAFIs.

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

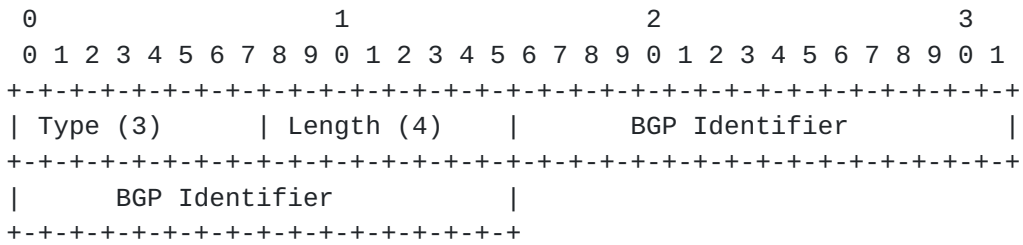
The BGP Config OS-TLV Local AS Sub-TLV will be used to advertise the 4-octet local Autonomous System (AS) number(s). For AS transitions, a second local AS number may be specified. The format of the BGP Local AS Sub-TLV is shown below.



- Type The Sub-TLV Type value shall be 2.
- Length The Sub-TLV Length will be 4 or 8 octets.
- Local AS Local Autonomous System (AS)
- Second Local AS Local Autonomous System (AS)

2.2.3. BGP Config OS-TLV - BGP Identifier Sub-TLV

The BGP Config OS-TLV BGP Identifier Sub-TLV will be used to advertise the 4-octet local BGP Identifier. The BGP Identifier is used for debugging purposes and possibly to reduce the likelihood of BGP connection collisions. The format of the BGP Identifier Sub-TLV is shown below.



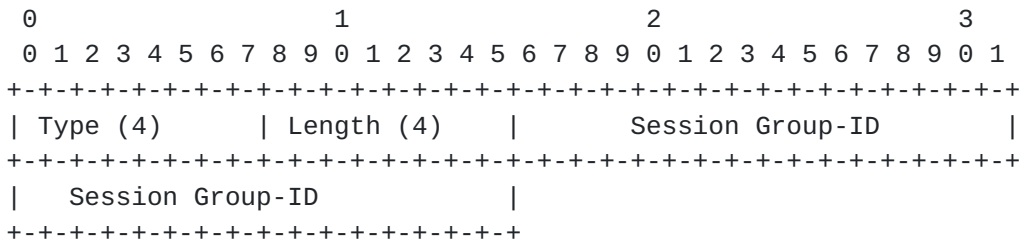
Type The Sub-TLV Type value shall be f32.

Length The Sub-TLV Length will be 4 octets.

BGP Identifier Local BGP Identifier (aka, BGP Router ID)

2.2.4. BGP Config OS-TLV - Session Group-ID Sub-TLV

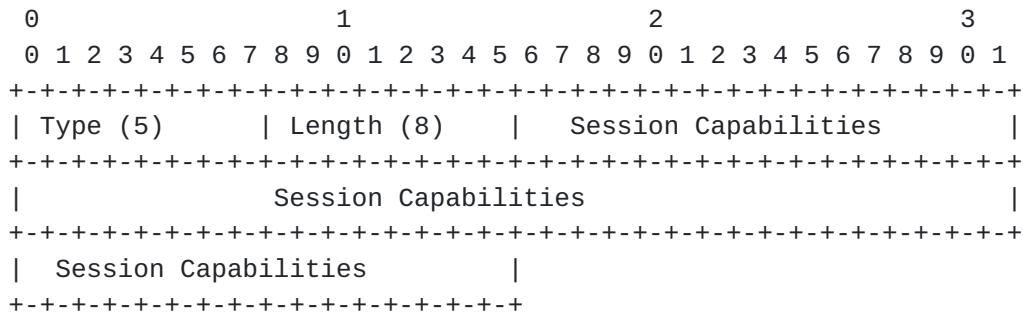
The BGP Config OS-TLV Session Group-ID Sub-TLV is an opaque 4-octet value that is used to represent a category of BGP session that is supported on the interface. The format of the Session Group-ID Sub-TLV is shown below.



- Type The Sub-TLV Type value shall be 4.
- Length The Sub-TLV Length will be 4 octets.
- Session Group-ID The session group-id used to indicate a class or category of BGP session supported on the interface.

2.2.5. BGP Config OS-TLV - BGP Session Capabilities Sub-TLV

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



- Type The Sub-TLV Type value shall be 5.
- Length The Sub-TLV Length will be 8 octets.
- Session Capabilities Bit fields identify BGP session capabilities

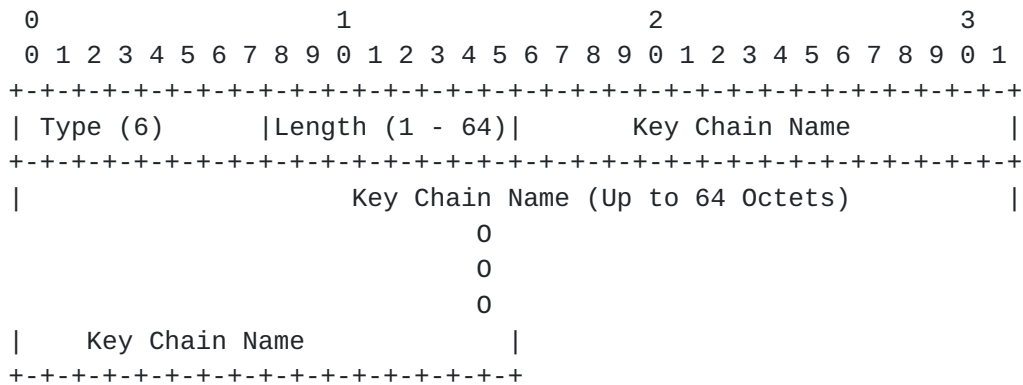
The BGP Session Capabilities is an 8-octet bit field. The most significant bit is the first bit (Bit 1) of the Session Capabilities. The following bits are defined:

- Bit 1: This bit indicates 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.

2.2.6. BGP Config OS-TLV - Key Chain Sub-TLV

The BGP Config OS-TLV Key Chain Sub-TLV is a string specifying the name for the key chain used for session authentication. Key chains [[YANG-KEY-CHAIN](#)] are a commonly used for protocol authentication and encryption key specification. Given the limited length of all BGP configuration information, the key chain name will be limited to 64 characters and will not include a trailing string delimiter. The format of the Session Group-ID Sub-TLV is shown below.



- Type The Sub-TLV Type value shall be 6.
- Length The Sub-TLV Length will be 1 - 64 octets.
- Key Chain Name The name of a key chain to be used for MD5 or TCP-A0 authentication.

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 use case where a loopback address or other local address is advertised as the peering address is also supported. However, reachability to a peering address other than the interface address is beyond the scope of this document.

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 the BGP-OS TLV. This can be an IPv4 or IPv6 local address associated with the LLDP link for 1-hop peering. For 2-hop peering, it could be a loopback address or any other address that is local to the node but not the LLDP link. As noted above, reachability to the loopback address is beyond the scope of this document.

A BGP speaker MAY advertise its local AS number using the BGP Local AS Sub-TLV of the BGP-OS TLV. During AS transitions, a second local AS number may be included in the Local AS Sub-TLV. The local BGP identifier may also be advertised using the BGP Identifier Sub-TLV of the BGP-OS TLV. While not specifically required for session establishment, the values may be used for validation, troubleshooting, and connection collision avoidance. A BGP speaker may also announce a Session Group-ID indicating the class or category of session(s) supported and/or mapping to a set of session parameters. Additionally, a BGP speaker MAY also announce relevant capabilities using BGP Session Capabilities Sub-TLV of the BGP-OS TLV.

If TCP MD5 authentication [[TCP-MD5](#)] or TCP Authentication Option (TCP-AO) [[TCP-AO](#)] is to be used on the session, the Key Chain Sub-TLV of the BGP-OS TLV MAY be used to specify the key chain name.

3.2. Receiving BGP Speaker

A BGP speaker configured for LLDP peer discovery WILL attempt to establish BGP sessions 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 is 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 peer must be reachable for the session to be established and the mechanisms for establishing reachability are beyond the scope of this specification. If the BGP speaker receives

the same BGP peering address in LLDP PDUs received on multiple links, it will not establish multiple sessions. Rather, a single 2-hop session will be established.

When the deployment of address families is fairly homogenous across the deployment, the wildcard AFI/SAFI can be utilized to simplify LLDP advertisement. When there is variance in the address families supported, usage of the wildcard could result in session establishment delay due to capabilities negotiation [[BGP-CAP](#)].

A BGP speaker MAY receive a remote neighbor's local AS number(s) in an LLDP PDU in the BGP Local AS Sub-TLV of the BGP-OS TLV. A BGP speaker MAY use the received local AS number(s) to perform validation checking of the AS received in the OPEN message. A BGP speaker MAY receive a remote neighbor's BGP Identifier in the BGP Identifier Sub-TLV of the BGP-OS TLV. This can be used to avoid connection collisions by delaying session establishment if the remote BGP Identifier is greater than the receiving speaker's BGP Identifier.

A BGP speaker MAY receive a Session Group-ID Sub-TLV in the LLDP BGP-OS TLV. This Session Group-ID may be used for validation and/or mapping the session to a particular set of session parameters. For example, the Session Group-ID could be mapped to a spine, leaf, or Top-of-Rack (ToR) session in a data center deployment and can be used to detect cabling problems when an unexpected Session Group-ID is received.

Additionally, A BGP speaker MAY receive a remote neighbor's capabilities in LLDP in the BGP Session Capabilities Sub-TLV of the BGP-OS TLV. A BGP speaker MAY use the received capabilities to ensure appropriate local neighbor configuration in order to facilitate session establishment.

If TCP MD5 authentication [[TCP-MD5](#)]. or TCP Authentication Option (TCP-AO) [[TCP-AO](#)] is to be used on the session as determined either via the Session Capabilities Sub-TLV, Session Group-ID, or local policy, the key chain name in the Key Chain Sub-TLV of the BGP-OS TLV MAY be used to identify the correct key chain [[YANG-KEY-CHAIN](#)].

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 BGP Config OS-TLV.

Range	Assignment Policy
0	Reserved (not to be assigned)
1	BGP Configuration
2-127	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 2-127 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 BGP Identifier Sub-TLV, 4 for the BGP Session Group-ID, 5 for the Session Capabilities Sub-TLV, and 6 for the Key Chain Name.

Range	Assignment Policy
0	Reserved (not to be assigned)
1	Peering Address
2	Local AS
3	BGP Identifier
4	Session Group-ID
5	Session Capabilities
6	Key Chain Name
7-127	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 7-127 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. Contributors

Contributors' Addresses

7. References

7.1. Normative References

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

[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.

7.2. Informative References

[BGP-CAP] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), February 2009.

[BGP-DC] Lapukhov, P., Premji, A., and J. Mitchell, "BGP Routing in Data Centers", [RFC 7938](#), August 2016.

[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.

[MP-BGP] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), January 2007.

[TCP-A0] 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.

[YANG-KEY-CHAIN]

Lindem, A., Qu, Y., Yeung, D., Chen, I., and J. Zhang, "YANG Data Model for Key Chains", [RFC 8177](#), June 2017.

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

Jeff Haas
Juniper Networks, Inc
1133 Innovation, Inc.
Sunnyvale, CA 94089
USA

Email: jhaas@juniper.net

Xiaohu Xu
Huawei

Email: xuxiaohu@huawei.com

