Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 29, 2019

R. Bush
Arrcus & IIJ
K. Patel
Arrcus
April 27, 2019

L3DL Upper Layer Protocol Configuration
draft-ymbk-lsvr-l3dl-ulpc-00

Abstract

This document users the Layer 3 Liveness and Discovery protocol to communicate the parameters needed to exchange inter-device Upper Layer Protocol Configuration for upper layer protocols such as the BGP family.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 https://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 October 29, 2019.

Copyright Notice

Table of Contents

1. Introduction

Massive Data Centers (MDCs) which use upper layer protocols such as BGP4, BGP-LS, BGP-SPF, etc. may use the Layer 3 Liveness and Discovery Protocol, L3DP, [I-D.ietf-lsvr-l3dl] to reveal the inter-device links of the topology. It is desirable for devices to facilitate the configuration parameters of those upper layer protocols to enable more hands-free configuration. This document defines a new L3DP PDU to communicate these Upper Layer Protocol Configuration parameters.

2. Reading and Terminology

The reader is assumed to have read Layer 3 Discovery and Liveness [I-D.ietf-lsvr-l3dl]. The terminology and PDUs there are assumed here.

Familiarity with the BGP4 Protocol [RFC4271] is assumed. Familiarity with BGP-SPF, [I-D.ietf-lsvr-bgp-spf], might be useful.

3. Upper Layer Protocol Configuration PDU

To communicate parameters required to configure peering and operation of Upper Layer Protocols at IP layer 3 and above, e.g., BGP sessions on a link, a neutral sub-TLV based Upper Layer Protocol PDU is defined as follows:

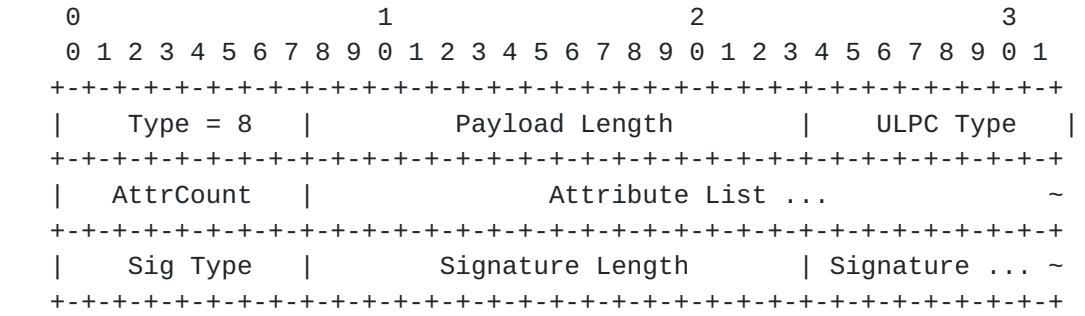

The Type and Payload Length are defined in [I-D.ietf-lsvr-l3dl].

ULPC Type:

Bit 0 : BGP
Bit 1-15: Must be 0

The AttrCount is the number of attribute sub-TLVs in the Attribute List.

The Attribute List is a, possibly null, set of sub-TLVs describing the configuration attributes of the specific upper layer protocol.

3.1. BGP ULPC Attribute sub-TLVs

The parameters needed for BGP peering on a link are exchanged in sub-TLVs within an Upper Layer Protocol PDU. The following describe the various sub-TLVs for BGP.

The goal is to provide the minimal set of configuration parameters needed by BGP OPEN to successfully start a BGP peering. The goal is specifically not to replace or conflict with data exchanged during BGP OPEN. Multiple sources of truth are a recipe for complexity and hence pain.

If there are multiple BGP sessions on a link, e.g., IPv4 and IPv6, then multiple sets of BGP sub-TLVs are exchanged within the BGP ULPC PDU.

A peer receiving BGP ULPC PDUs has only one active BGP ULPC PDU at any point in time; receipt of a new BGP ULPC PDU replaces any

previous one. If there are one or more open BGP sessions, receipt of a new BGP ULPC PDU does not affect these sessions.

As a link may have multiple encapsulations and multiple addresses for an IP encapsulation, which address of which encapsulation are to be used for the BGP session MUST be specified.

For each BGP peering on a link here MUST be one agreed encapsulation, and the addresses used MUST be in the corresponding L3DP IPv4/IPv6 Announcement PDUs. If a peering address has been announced as a loopback, a two or three (one or both ends could be loopbacks) hop BGP session will be established. Otherwise a direct one hop session is used.

3.1.1. BGP ASN

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

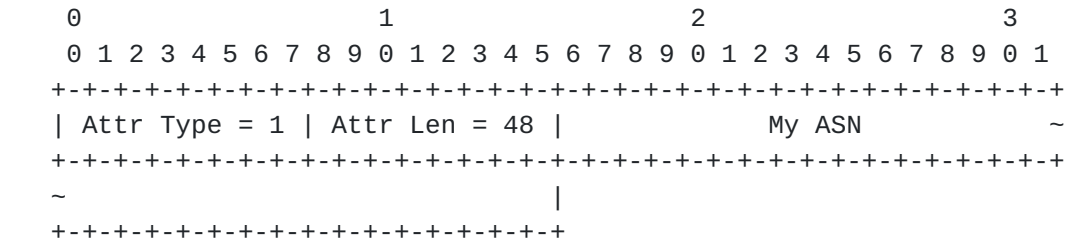

3.1.2. BGP IPv4 Address

The BGP IPv4 Address sub-TLV announces the sender's IPv4 BGP peering source address to be used by the receiver. At least one of IPv4 or IPv6 BGP source addresses MUST be announced.

As usual, the BGP OPEN capability negotiation will determine the AFI/SAFIs to be transported over the peering, see [RFC4760] .

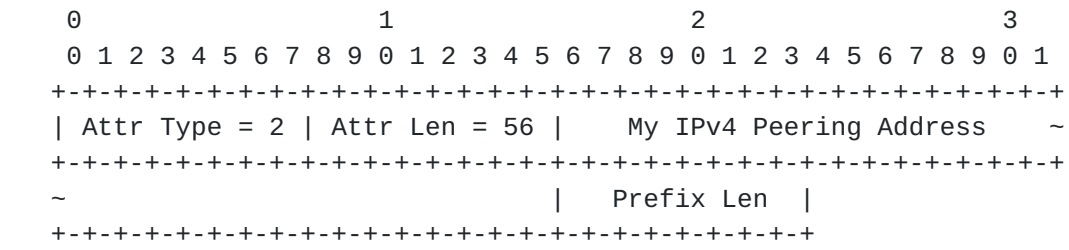

3.1.3. BGP IPv6 Address

The BGP IPv6 Address sub-TLV announces the sender's IPv6 BGP peering source address to be used by the receiver. At least one of IPv4 or IPv6 BGP source addresses MUST be announced.

As usual, the BGP OPEN capability negotiation will determine the AFI/SAFIs to be transported over the peering, see [RFC4760] .

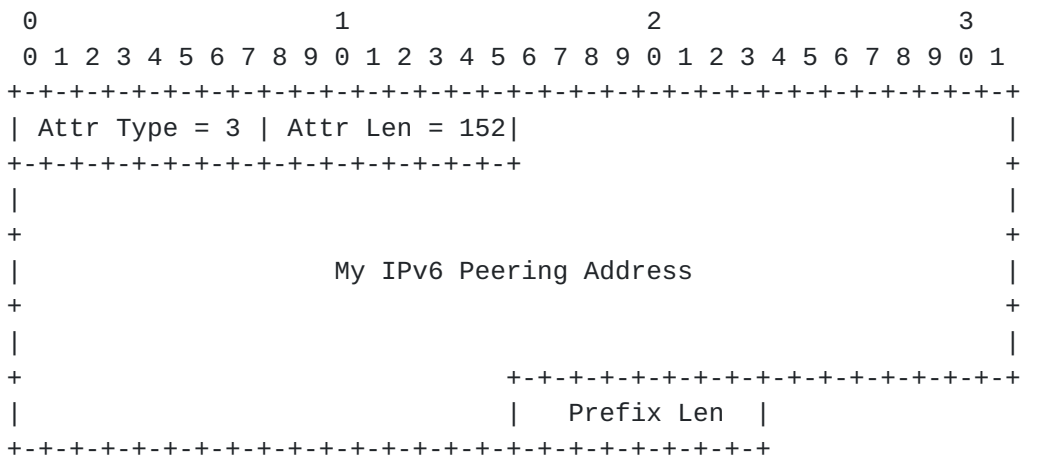

3.1.4. BGP Authentication sub-TLV

The BGP Authentication sub-TLV provides any authentication data needed to OPEN the BGP session. Depending on operator configuration of the environment, it might be a simple MD5 key (see [RFC2385]), the name of a key chain a KARP database (see [RFC7210]), or one of multiple Authentication sub-TLVs to support hop[RFC4808].

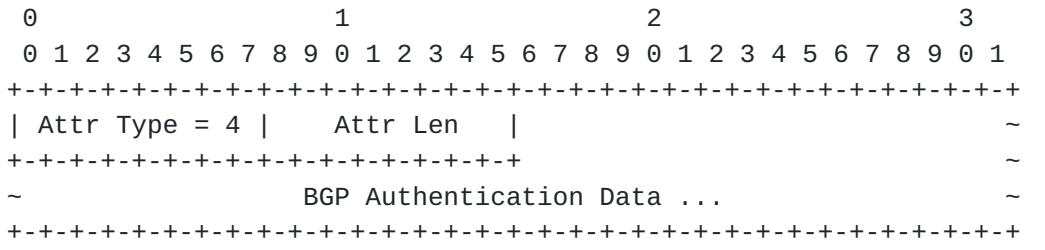

3.1.5. BGP Miscellaneous Flags

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

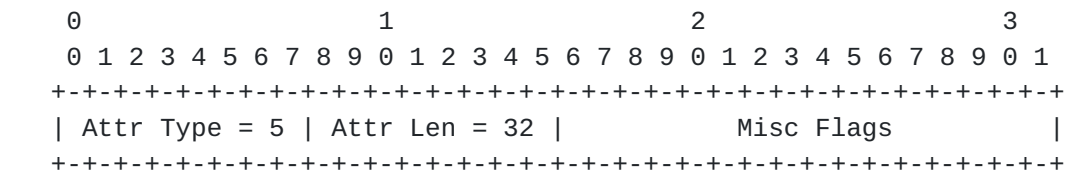

Misc Attrs:

Bit 0: Ghu knows what
Bit 1-15: Must be zero

4. Security Considerations

All the Security considerations of [I-D.ietf-lsvr-l3dl] apply to this PDU.

As the ULPC PDU may contain keying material, see Section 3.1.4, it SHOULD BE signed.

Any keying material in the PDU SHOULD BE salted ad hashed.

The BGP Authentication sub-TLV provides for provisioning MD5, which is a quite weak hash, horribly out of fashion, and kills puppies. But, like it or not, it is what BGP deployments use.

5. IANA Considerations

This document requests the IANA create a new entry in the L3DL PDU Type registry as follows:

PDU Code	PDU Name
9	ULPC

This document requests the IANA create a registry for L3DL ULPC Type, which may range from 0 to 255. The name of the registry should be L3DL-ULPC-Type. The policy for adding to the registry is RFC Required per [RFC5226], either standards track or experimental. The initial entries should be the following:

Bit	Bit Name
0	Reserved
1	BGP
2-255	Reserved

6. References

6.1. Normative References

[I-D.ietf-lsvr-l3dl]
Bush, R., Austein, R., and K. Patel, "Layer 3 Discovery and Liveness", draft-ietf-lsvr-l3dl-00 (work in progress), April 2019.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <http://www.rfc-editor.org/info/rfc4271>.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <http://www.rfc-editor.org/info/rfc4760>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <http://www.rfc-editor.org/info/rfc8174>.

6.2. Informative References

[I-D.ietf-lsvr-bgp-spf]
Patel, K., Lindem, A., Zandi, S., and W. Henderickx, "Shortest Path Routing Extensions for BGP Protocol", draft-ietf-lsvr-bgp-spf-04 (work in progress), December 2018.

[RFC2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, DOI 10.17487/RFC2385, August 1998, <http://www.rfc-editor.org/info/rfc2385>.

[RFC4808] Bellovin, S., "Key Change Strategies for TCP-MD5", RFC 4808, DOI 10.17487/RFC4808, March 2007, <http://www.rfc-editor.org/info/rfc4808>.

[RFC7210] Housley, R., Polk, T., Hartman, S., and D. Zhang, "Database of Long-Lived Symmetric Cryptographic Keys", RFC 7210, DOI 10.17487/RFC7210, April 2014, <http://www.rfc-editor.org/info/rfc7210>.

Authors' Addresses

Randy Bush
Arrcus & IIJ
5147 Crystal Springs
Bainbridge Island, WA 98110
United States of America

Email: randy@psg.com

Keyur Patel
Arrcus
2077 Gateway Place, Suite #400
San Jose, CA 95119
United States of America

Email: keyur@arrcus.com