

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 17, 2016

J. Dong
Z. Li
Huawei Technologies
J. Tantsura
Ericsson
H. Gredler
Private Contributor
March 16, 2016

BGP Link-State Extension for Distribution of IP Tunnel Information
draft-dong-idr-ls-ip-tunnel-00

Abstract

This document specifies extensions to BGP-LS for the collection and distribution of IP tunnel information. Such information can be distributed to external components for service mapping and tunnel selection.

Requirements Language

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 2119](#) [[RFC2119](#)].

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 17, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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
2.	Carrying IP Tunnel Information in BGP-LS	3
2.1.	IP Tunnel Identifier Information	3
2.2.	IP Tunnel Parameters TLV	6
3.	Operational Considerations	8
4.	IANA Considerations	8
4.1.	BGP-LS NLRI-Types	8
4.2.	BGP-LS Protocol-IDs	9
4.3.	BGP-LS Attribute TLVs	9
5.	Security Considerations	9
6.	Acknowledgements	10
7.	References	10
7.1.	Normative References	10
7.2.	Informative References	11
	Authors' Addresses	11

[1.](#) Introduction

BGP has been extended to distribute the link-state [[I-D.ietf-idr-ls-distribution](#)] and TE-LSP information [[I-D.ietf-idr-te-lsp-distribution](#)] to external components. When IP tunnel technologies, such as Generic Routing Encapsulation (GRE), Layer Two Tunneling Protocol - Version 3 (L2TPv3), VxLAN, NVGRE, etc., are used in the network, it is necessary to collect the information of IP tunnels in the network and share with the external components. Such information can be distributed to external components for service mapping and tunnel selection. One typical use case of IP tunnel information is described in [[I-D.hao-idr-flowspec-redirect-tunnel](#)]. This document specifies extensions to BGP-LS for the collection and distribution of IP tunnel information.

2. Carrying IP Tunnel Information in BGP-LS

2.1. IP Tunnel Identifier Information

The IP tunnel Identifier information is advertised in BGP UPDATE messages using the MP_REACH_NLRI and MP_UNREACH_NLRI attributes [RFC4760]. The "Link-State NLRI" defined in [I-D.ietf-idr-ls-distribution] is extended to carry the IP Tunnel information. BGP speakers that wish to exchange IP Tunnel information MUST use the BGP Multiprotocol Extensions Capability Code (1) to advertise the corresponding (AFI, SAFI) pair, as specified in [RFC4760].

The format of "Link-State NLRI" is defined in [I-D.ietf-idr-ls-distribution]. A new "NLRI Type" is defined for IP Tunnel Identifier Information as following:

- o NLRI Type = TBA: IPv4/IPv6 Tunnel NLRI

The IPv4/IPv6 Tunnel NLRI is shown in the following figure:

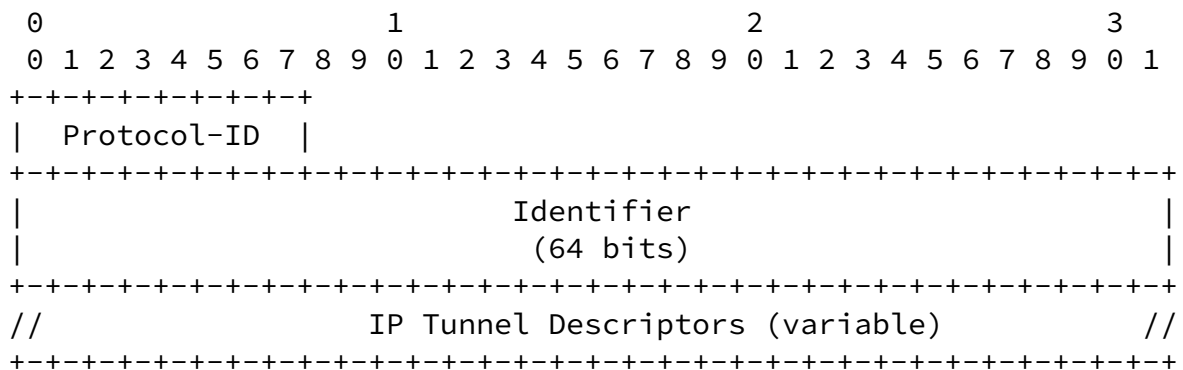


Figure 1. IPv4/IPv6 Tunnel NLRI

Where

The 'Protocol-ID' field is used to identify the source of the

advertised NLRI. For IPv4/IPv6 Tunnel NLRI, according to the method of tunnel establishment, the Protocol-ID field can be set to either "Static configuration" or the specific signaling protocol of the IP tunnel. Several new Protocol-IDs are defined as below:


Protocol-ID	NLRI information source protocol
TBD	L2TPv3
TBD	GTPv2-C

As defined in [[I-D.ietf-idr-ls-distribution](#)], the 64-Bit 'Identifier' field is used to identify the "routing universe" where the NLRI belongs.

The "IP Tunnel Descriptors" field consists of a set of Descriptor TLVs which together identifies the IP tunnel. The following Descriptor TLVs as defined in [[I-D.ietf-idr-te-lsp-distribution](#)] are reused for IPv4/IPv6 Tunnel NLRI:

o Tunnel ID

The Tunnel Identifier TLV contains the Tunnel ID defined in [[RFC3209](#)] and has the following format:



where:

- + Type: To be assigned by IANA (suggested value: 267)

- + Length: 2 octets.
- + Tunnel ID: 2 octets as defined in [[RFC3209](#)].
- o IPv4/6 Tunnel Head-end address

The IPv4/IPv6 Tunnel Head-End Address TLV contains the Tunnel Head- End Address defined in [[RFC3209](#)] and has following format:

```

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               |           Length               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//           IPv4/IPv6 Tunnel Head-End Address (variable)           //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

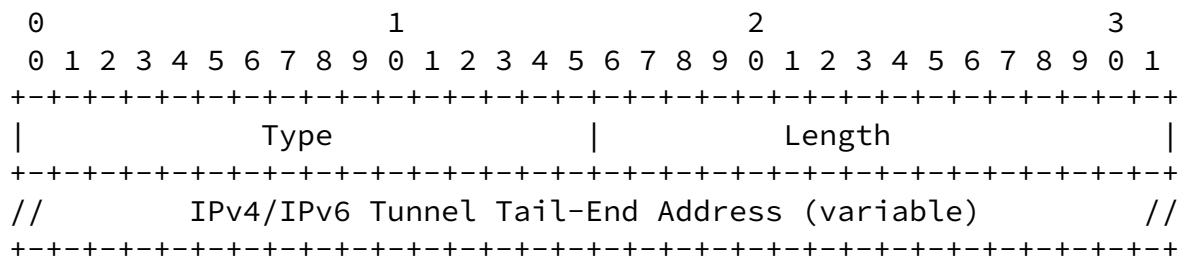
- + Type: To be assigned by IANA (suggested value: 269)
- + Length: 4 or 16 octets.

When the IPv4/IPv6 Tunnel Head-end Address TLV contains an IPv4 address, its length is 4 (octets).

When the IPv4/IPv6 Tunnel Head-end Address TLV contains an IPv6 address, its length is 16 (octets).

o IPv4/6 Tunnel Tail-end address

The IPv4/IPv6 Tunnel Tail-End Address TLV contains the Tunnel Tail- End Address defined in [[RFC3209](#)] and has following format:




where:

- + Type: To be assigned by IANA (suggested value: 270)
- + Length: 4 or 16 octets.

When the IPv4/IPv6 Tunnel Tail-end Address TLV contains an IPv4 address, its length is 4 (octets).

When the IPv4/IPv6 Tunnel Tail-end Address TLV contains an IPv6 address, its length is 16 (octets).

In addition, a new descriptor TLV called "Tunnel Type TLV" is defined for IP Tunnel as below:



```

+---+---+---+---+---+---+---+---+---+

```

Figure 2. Tunnel Type TLV

- o Type: TBA.
- o Length: 2 octets.
- o Value: The 2-octet Tunnel Type identifies the type of tunneling technology as defined in the "BGP Tunnel Encapsulation Attribute Tunnel Types" registry [[RFC5512](#)].

The IPv4/6 Tunnel Head-end address TLV, IPv4/6 Tunnel Tail-end address, Tunnel-type TLV and the Tunnel ID TLV together uniquely identify the IP tunnel.

[2.2.](#) IP Tunnel Parameters TLV

A new TLV called "IP Tunnel Parameters TLV" is defined to describe the detailed information of the IP tunnels, which is carried in the optional non-transitive BGP Attribute "LINK_STATE Attribute" defined in [[I-D.ietf-idr-ls-distribution](#)]. The IP Tunnel Parameters TLV SHOULD only be used with IPv4/IPv6 Tunnel NLRI.





Figure 3. IP Tunnel Parameters TLV

The "Value" field of the IP Tunnel Parameters TLV is composed of a set of sub-TLVs. The sub-TLV is structured as below:



~ Sub-TLV Value (Variable) ~
+-----+

The following sub-TLVs defined in this document can be carried in the Value field of the IP Tunnel Parameters TLV:

o Tunnel Name sub-TLV:

Type: TBA

Length: Variable

Value: A string identifies the name of the IP tunnel.

o Description sub-TLV

Type: TBA

Length: Variable

Value: A string which contains the textual description of the IP tunnel.

o Status sub-TLV:

Type: TBA

Length: 1 octet

Value: 8-bit flags which indicate the status of the IP tunnel. Bit 0 is defined as the Up/Down bit, which SHOULD be set to 1 if there is no available route for the tunnel destination. The other bits are reserved which MUST be set to 0 on transmission and ignored on receipt.

o Encapsulation sub-TLV:

Type: TBA

Length: Variable

Value: The encapsulation information of the IP tunnel, syntax

and semantics of which are determined by the Tunnel Type. The format of Encapsulation sub-TLVs are defined in [[RFC5512](#)] and [[I-D.ietf-idr-tunnel-encaps](#)].

- o CoS Sub-TLV:

Type: TBA

Length: 1 octet

Value: the class of differentiated services that can be provided by the tunnel. The format is same as the DS field as defined in [[RFC2474](#)].

- o MTU sub-TLV:

Type: TBA

Length: 2 octets

Value: the Maximum Transmission Unit (MTU) of the IP tunnel.

[3.](#) Operational Considerations

The Existing BGP operational procedures apply to this document. No new operation procedures are defined in this document. The operational considerations as specified in [[I-D.ietf-idr-ls-distribution](#)] apply to this document.

In general the ingress nodes of the IP Tunnels are responsible for the distribution of the IP tunnel information, while the egress nodes of the IP tunnels MAY report the IP tunnel information if needed.

[4.](#) IANA Considerations

IANA is requested to administer the assignment of new values defined in this document and summarized in this section.

[4.1.](#) BGP-LS NLRI-Types

IANA maintains a registry called "Border Gateway Protocol - Link State (BGP-LS) Parameters" with a sub-registry called "BGP-LS NLRI-Types".

IANA is requested to assign two new NLRI-Types:

Type	NLRI Type	Reference
TBD	IPv4/v6 Tunnel NLRI	this document

4.2. BGP-LS Protocol-IDs

IANA maintains a registry called "Border Gateway Protocol - Link State (BGP-LS) Parameters" with a sub-registry called "BGP-LS Protocol-IDs".

IANA is requested to assign two new Protocol-IDs:

Protocol-ID	NLRI information source protocol	Reference
TBD	L2TPv3	this document
TBD	GTPv2-C	this document

4.3. BGP-LS Attribute TLVs

IANA maintains a registry called "Border Gateway Protocol - Link State (BGP-LS) Parameters" with a sub-registry called "Node Anchor, Link Descriptor and Link Attribute TLVs".

IANA is requested to assign one new TLV code point:

TLV Code Point	Description	IS-IS TLV/ Sub-TLV	Value defined in:
TBD	IP Tunnel Parameters TLV	---	this document

5. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the 'Security Considerations' section of [\[RFC4271\]](#) for a discussion of BGP security. Also refer to [\[RFC4272\]](#) and [\[RFC6952\]](#) for analysis of security issues for BGP.

6. Acknowledgements

The authors would like to thank Nan Wu, Shunwan Zhuang and Xia Chen for their review and valuable comments.

7. References

7.1. Normative References

[I-D.ietf-idr-ls-distribution]

Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", [draft-ietf-idr-ls-distribution-13](#) (work in progress), October 2015.

[I-D.ietf-idr-te-lsp-distribution]

Dong, J., Chen, M., Gredler, H., Previdi, S., and J. Tantsura, "Distribution of MPLS Traffic Engineering (TE) LSP State using BGP", [draft-ietf-idr-te-lsp-distribution-04](#) (work in progress), December 2015.

[I-D.ietf-idr-tunnel-encaps]

Rosen, E., Patel, K., and G. Velde, "The BGP Tunnel Encapsulation Attribute", [draft-ietf-idr-tunnel-encaps-01](#) (work in progress), December 2015.

[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>>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<http://www.rfc-editor.org/info/rfc3209>>.

[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>>.

Dong, et al.

Expires September 17, 2016

[Page 10]

Internet-Draft

BGP-LS for IP Tunnel Distribution

March 2016

[7.2](#). Informative References

- [I-D.hao-idr-flowspec-redirect-tunnel]
Weiguo, H., Li, Z., and L. Yong, "BGP Flow-Spec Redirect to Tunnel action", [draft-hao-idr-flowspec-redirect-tunnel-00](#) (work in progress), October 2015.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", [RFC 2474](#), DOI 10.17487/RFC2474, December 1998, <<http://www.rfc-editor.org/info/rfc2474>>.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", [RFC 4272](#), DOI 10.17487/RFC4272, January 2006, <<http://www.rfc-editor.org/info/rfc4272>>.
- [RFC5342] Eastlake 3rd, D., "IANA Considerations and IETF Protocol Usage for IEEE 802 Parameters", [RFC 5342](#), DOI 10.17487/RFC5342, September 2008, <<http://www.rfc-editor.org/info/rfc5342>>.
- [RFC5512] Mohapatra, P. and E. Rosen, "The BGP Encapsulation Subsequent Address Family Identifier (SAFI) and the BGP Tunnel Encapsulation Attribute", [RFC 5512](#), DOI 10.17487/RFC5512, April 2009, <<http://www.rfc-editor.org/info/rfc5512>>.
- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", [RFC 6952](#), DOI 10.17487/RFC6952, May 2013,

<<http://www.rfc-editor.org/info/rfc6952>>.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Campus, No.156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

Dong, et al. Expires September 17, 2016 [Page 11]

Internet-Draft BGP-LS for IP Tunnel Distribution March 2016

Zhenbin Li
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, CA 95134
US

Email: jeff.tantsura@ericsson.com

Hannes Gredler
Private Contributor

Email: hannes@gredler.at

