< draft-ietf-lisp-gpe-14.txt	draft-ietf-lisp-gpe-15.txt >
Internet Engineering Task Force F. Maino, Ed. Internet-Draft Cisco Intended status: Standards Track J. Lemon Expires: July 11, 2020 Broadcom P. Agarwal Innovium D. Lewis M. Smith Cisco January 8, 2020	Internet Engineering Task Force F. Maino, Ed. Internet-Draft Cisco Intended status: Standards Track J. Lemon Expires: December 2, 2020 Broadcom P. Agarwal Innovium D. Lewis M. Smith Cisco May 31, 2020
LISP Generic Protocol Extension draft-ietf-lisp-gpe-14	LISP Generic Protocol Extension draft-ietf-lisp-gpe-15
Abstract This document describes extentions to the Locator/ID Separation Protocol (LISP) Data-Plane, via changes to the LISP header, to support multi-protocol encapsulation.	Abstract This document describes extentions to the Locator/ID Separation Protocol (LISP) Data-Plane, via changes to the LISP header, to support multi-protocol encapsulation.
Status of This Memo This Internet-Draft is submitted in full conformance with the	Status of This Memo This Internet-Draft is submitted in full conformance with the
skipping to change at *page 1, line 38*	**skipping to change at *page 1, line 38***
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 July 11, 2020.	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 December 2, 2020.
Copyright Notice Copyright (c) 2020 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 (https://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents	Copyright Notice Copyright (c) 2020 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 (https://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents
skipping to change at *page 2, line 21*	**skipping to change at *page 2, line 21***
1. Introduction . . . 2 1.1. Conventions . . . 3 1.2. Definition of Terms . . . 3 2. LISP Header Without Protocol Extensions . . . 3 3. Generic Protocol Extension for LISP (LISP-GPE) . . . 4 4. Implementation and Deployment Considerations . . . 6 4.1. Applicability Statement . . . 6 4.2. Congestion Control Functionality . . . 7 4.3. UDP Checksum . . . 7 4.3.1. UDP Zero Checksum Handling with IPv6 . . . 8 4.4. Ethernet Encapsulated Payloads . . . 9 5. Backward Compatibility . . . 10 5.1. Detection of ETR Capabilities . . . 10 6. IANA Considerations . . . 10 6.1. LISP-GPE Next Protocol Registry . . . 10 7. Security Considerations . . . 10 8. Acknowledgements and Contributors . . . 11 9. References . . . 11 9.1. Normative References . . . 11 9.2. Informative References . . . 12 Authors' Addresses . . . 13	1. Introduction . . . 2 1.1. Conventions . . . 3 1.2. Definition of Terms . . . 3 2. LISP Header Without Protocol Extensions . . . 3 3. Generic Protocol Extension for LISP (LISP-GPE) . . . 4 4. Implementation and Deployment Considerations . . . 6 4.1. Applicability Statement . . . 6 4.2. Congestion Control Functionality . . . 7 4.3. UDP Checksum . . . 7 4.3.1. UDP Zero Checksum Handling with IPv6 . . . 8 4.4. DSCP, ECN and TTL . . . 9 5. Backward Compatibility . . . 10 5.1. Detection of ETR Capabilities . . . 10 6. IANA Considerations . . . 11 6.1. LISP-GPE Next Protocol Registry . . . 11 7. Security Considerations . . . 11 8. Acknowledgements and Contributors . . . 11 9. References . . . 12 9.1. Normative References . . . 12 9.2. Informative References . . . 13 Authors' Addresses . . . 14
1. Introduction The LISP Data-Plane is defined in [I-D.ietf-lisp-rfc6830bis]. It specifies an encapsulation format that carries IPv4 or IPv6 packets (henceforth jointly referred to as IP) in a LISP header and outer UDP/IP transport. The LISP Data-Plane header does not specify the protocol being encapsulated and therefore is currently limited to encapsulating only	1. Introduction The LISP Data-Plane is defined in [I-D.ietf-lisp-rfc6830bis]. It specifies an encapsulation format that carries IPv4 or IPv6 packets (henceforth jointly referred to as IP) in a LISP header and outer UDP/IP transport. The LISP Data-Plane header does not specify the protocol being encapsulated and therefore is currently limited to encapsulating only
skipping to change at *page 5, line 29*	**skipping to change at *page 5, line 29***
This document defines the following Next Protocol values: 0x01 : IPv4 0x02 : IPv6 0x03 : Ethernet 0x04 : Network Service Header (NSH) [RFC8300] 0x05 to 0x7F: Unassigned 0x80 to 0xFF: Unassigned (shim headers)	This document defines the following Next Protocol values: 0x01 : IPv4 0x02 : IPv6 0x03 : Ethernet 0x04 : Network Service Header (NSH) [RFC8300] 0x05 to 0x7D Unassigned 0x7E to 0x7F: Experimentation and testing 0x80 to 0xFD: Unassigned (shim headers) 0xFE to 0xFF: Experimentation and testing

The values are tracked in the IANA LISP-GPE Next Protocol Registry as described in Section 6.1.

Next protocol values from 0x80 to 0xFF are assigned to protocols encoded as generic "shim" headers. All shim protocols MUST use the header structure in Figure 4, which includes a Next Protocol field. When a shim header is used with other protocols identified by next protocol values from 0x0 to 0x7F, the shim header MUST come before the further protocol, and the next header of the shim will indicate which protocol follows the shim header.

Shim headers can be used to incrementally deploy new GPE features, keeping the processing of shim headers known to a given xTR implementation in the 'fast' path (typically an ASIC), while punting the processing of the remaining new GPE features to the 'slow' path.

Shim protocols MUST have the first 32 bits defined as:

The values are tracked in the IANA LISP-GPE Next Protocol Registry as described in Section 6.1.

Next protocol values 0x7E, 0x7F and 0xFE, 0xFF are assigned for experimentation and testing as per [RFC3692].

Next protocol values from 0x80 to 0xFD are assigned to protocols encoded as generic "shim" headers. All shim protocols MUST use the header structure in Figure 4, which includes a Next Protocol field. When a shim header is used with other protocols identified by next protocol values from 0x0 to 0x7D, the shim header MUST come before the further protocol, and the next header of the shim will indicate which protocol follows the shim header.

Shim headers can be used to incrementally deploy new GPE features, keeping the processing of shim headers known to a given xTR implementation in the 'fast' path (typically an ASIC), while punting the processing of the remaining new GPE features to the 'slow' path.

Shim protocols MUST have the first 32 bits defined as:

skipping to change at ***page 9, line 40*** **skipping to change at** ***page 9, line 46***

specified in [RFC6936].

The requirement to check the source IPv6 address in addition to the destination IPv6 address, plus the recommendation against reuse of source IPv6 addresses among LISP-GPE tunnels collectively provide some mitigation for the absence of UDP checksum coverage of the IPv6 header. A traffic-managed controlled environment that satisfies at least one of three conditions listed at the beginning of this section provides additional assurance.

4.4. Ethernet Encapsulated Payloads

When a LISP-GPE router performs Ethernet encapsulation, the inner 802.1Q [IEEE.802.1Q_2014] 3-bit priority code point (PCP) field MAY be mapped from the encapsulated frame to the 3-bit Type of Service field in the outer IPv4 header, or in the case of IPv6 the 'Traffic Class' field.

When a LISP-GPE router performs Ethernet encapsulation, the inner header 802.1Q [IEEE.802.1Q_2014] VLAN Identifier (VID) MAY be mapped to, or used to determine the LISP Instance IDentifier (IID) field.

5. Backward Compatibility

LISP-GPE uses the same UDP destination port (4341) allocated to LISP.

When encapsulating IP packets to a non LISP-GPE capable router the

specified in [RFC6936].

The requirement to check the source IPv6 address in addition to the destination IPv6 address, plus the recommendation against reuse of source IPv6 addresses among LISP-GPE tunnels collectively provide some mitigation for the absence of UDP checksum coverage of the IPv6 header. A traffic-managed controlled environment that satisfies at least one of three conditions listed at the beginning of this section provides additional assurance.

4.4. DSCP, ECN and TTL

When encapsulating IP (including over Ethernet) packets [RFC2983] provides guidance for mapping DSCP between inner and outer IP headers. The Pipe model typically fits better Network virtualization. The DSCP value on the tunnel header is set based on a policy (which may be a fixed value, one based on the inner traffic class, or some other mechanism for grouping traffic). Some aspects of the Uniform model (which treats the inner and outer DSCP value as a single field by copying on ingress and egress) may also apply, such as the ability to remark the inner header on tunnel egress based on transit marking. However, the Uniform model is not conceptually consistent with network virtualization, which seeks to provide strong isolation between encapsulated traffic and the physical network.

[RFC6040] describes the mechanism for exposing ECN capabilities on IP tunnels and propagating congestion markers to the inner packets. This behavior MUST be followed for IP packets encapsulated in LISP-GPE.

Though Uniform or Pipe models could be used for TTL (or Hop Limit in case of IPv6) handling when tunneling IP packets, Pipe model is more aligned with network virtualization. [RFC2003] provides guidance on handling TTL between inner IP header and outer IP tunnels; this model is more aligned with the Pipe model and is recommended for use with LISP-GPE for network virtualization applications.

When a LISP-GPE router performs Ethernet encapsulation, the inner 802.1Q [IEEE.802.1Q_2014] 3-bit priority code point (PCP) field MAY be mapped from the encapsulated frame to the DSCP codepoint of the DS field defined in [RFC2474].

When a LISP-GPE router performs Ethernet encapsulation, the inner header 802.1Q [IEEE.802.1Q_2014] VLAN Identifier (VID) MAY be mapped to, or used to determine the LISP Instance IDentifier (IID) field.

5. Backward Compatibility

LISP-GPE uses the same UDP destination port (4341) allocated to LISP.

When encapsulating IP packets to a non LISP-GPE capable router the

skipping to change at ***page 10, line 35*** **skipping to change at** ***page 11, line 16***

6.1. LISP-GPE Next Protocol Registry

IANA is requested to set up a registry of LISP-GPE "Next Protocol". These are 8-bit values. Next Protocol values in the table below are defined in this document. New values are assigned under the Specification Required policy [RFC8126]. The protocols that are being assigned values do not themselves need to be IETF standards track protocols.

Next Protocol	Description	Reference
0x00	Reserved	This Document
0x01	IPv4	This Document
0x02	IPv6	This Document
0x03	Ethernet	This Document
0x04	NSH	This Document
0x05..0x7F	Unassigned	
0x82..0xFF	Unassigned	

6.1. LISP-GPE Next Protocol Registry

IANA is requested to set up a registry of LISP-GPE "Next Protocol". These are 8-bit values. Next Protocol values in the table below are defined in this document. New values are assigned under the Specification Required policy [RFC8126]. The protocols that are being assigned values do not themselves need to be IETF standards track protocols.

Next Protocol	Description	Reference
0x0	Reserved	This Document
0x1	IPv4	This Document
0x2	IPv6	This Document
0x3	Ethernet	This Document
0x4	NSH	This Document
0x05..0x7D	Unassigned	
0x7E..0x7F	Experimentation and testing	This Document
0x80..0xFD	Unassigned (shim headers)	
0x8E..0x8F	Experimentation and testing	This Document

	+---------------+----------------------------+---------------+
7. Security Considerations LISP-GPE security considerations are similar to the LISP security considerations and mitigation techniques documented in [RFC7835]. LISP-GPE, as many encapsulations that use optional extensions, is subject to on-path adversaries that by manipulating the P-Bit and the packet itself can remove part of the payload or claim to encapsulate any protocol payload type. Typical integrity protection mechanisms	7. Security Considerations LISP-GPE security considerations are similar to the LISP security considerations and mitigation techniques documented in [RFC7835]. LISP-GPE, as many encapsulations that use optional extensions, is subject to on-path adversaries that by manipulating the P-Bit and the packet itself can remove part of the payload or claim to encapsulate any protocol payload type. Typical integrity protection mechanisms
skipping to change at *page 11, line 51*	**skipping to change at *page 12, line 32***
o Puneet Agarwal, Innovium 9. References 9.1. Normative References [I-D.ietf-lisp-rfc6830bis] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos-Aparicio, "The Locator/ID Separation Protocol (LISP)", draft-ietf-lisp-rfc6830bis-28 (work in progress), November 2019.	o Puneet Agarwal, Innovium 9. References 9.1. Normative References [I-D.ietf-lisp-rfc6830bis] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos-Aparicio, "The Locator/ID Separation Protocol (LISP)", draft-ietf-lisp-rfc6830bis-32 (work in progress), March 2020.
[IEEE.802.1Q_2014] IEEE, "IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks", IEEE 802.1Q-2014, DOI 10.1109/ieeestd.2014.6991462, December 2014, <http://ieeexplore.ieee.org/servlet/opac?punumber=6991460>. [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119,	[IEEE.802.1Q_2014] IEEE, "IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks", IEEE 802.1Q-2014, DOI 10.1109/ieeestd.2014.6991462, December 2014, <http://ieeexplore.ieee.org/servlet/opac?punumber=6991460>. [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119,
skipping to change at *page 12, line 41*	**skipping to change at *page 13, line 25***
Briscoe, B., Kaippallimalil, J., and P. Thaler, "Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP", draft-ietf-tsvwg-ecn-encap-guidelines-13 (work in progress), May 2019. [I-D.lemon-vxlan-lisp-gpe-gbp] Lemon, J., Maino, F., Smith, M., and A. Isaac, "Group Policy Encoding with VXLAN-GPE and LISP-GPE", draft-lemon-vxlan-lisp-gpe-gbp-02 (work in progress), April 2019.	Briscoe, B., Kaippallimalil, J., and P. Thaler, "Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP", draft-ietf-tsvwg-ecn-encap-guidelines-13 (work in progress), May 2019. [I-D.lemon-vxlan-lisp-gpe-gbp] Lemon, J., Maino, F., Smith, M., and A. Isaac, "Group Policy Encoding with VXLAN-GPE and LISP-GPE", draft-lemon-vxlan-lisp-gpe-gbp-02 (work in progress), April 2019.
	[RFC2003] Perkins, C., "IP Encapsulation within IP", RFC 2003, DOI 10.17487/RFC2003, October 1996, <https://www.rfc-editor.org/info/rfc2003>.
[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <https://www.rfc-editor.org/info/rfc2460>.	[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <https://www.rfc-editor.org/info/rfc2460>.
	[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, <https://www.rfc-editor.org/info/rfc2474>. [RFC2983] Black, D., "Differentiated Services and Tunnels", RFC 2983, DOI 10.17487/RFC2983, October 2000, <https://www.rfc-editor.org/info/rfc2983>. [RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", BCP 82, RFC 3692, DOI 10.17487/RFC3692, January 2004, <https://www.rfc-editor.org/info/rfc3692>.
[RFC6935] Eubanks, M., Chimento, P., and M. Westerlund, "IPv6 and UDP Checksums for Tunneled Packets", RFC 6935, DOI 10.17487/RFC6935, April 2013, <https://www.rfc-editor.org/info/rfc6935>. [RFC6936] Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", RFC 6936, DOI 10.17487/RFC6936, April 2013, <https://www.rfc-editor.org/info/rfc6936>.	[RFC6935] Eubanks, M., Chimento, P., and M. Westerlund, "IPv6 and UDP Checksums for Tunneled Packets", RFC 6935, DOI 10.17487/RFC6935, April 2013, <https://www.rfc-editor.org/info/rfc6935>. [RFC6936] Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", RFC 6936, DOI 10.17487/RFC6936, April 2013, <https://www.rfc-editor.org/info/rfc6936>.
End of changes. 17 change blocks.	
33 lines changed or deleted	*85 lines changed or added*