Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 5, 2015

D. Lewis
Cisco Systems, Inc.
P. Agarwal
Broadcom
L. Kreeger
F. Maino
P. Quinn
M. Smith
N. Yadav
Cisco Systems, Inc.
July 4, 2014

LISP Generic Protocol Extension
draft-lewis-lisp-gpe-02.txt

Abstract

This draft describes extending the Locator/ID Separation Protocol (LISP) [RFC6830], via changes to the LISP header, with three new capabilities: support for multi-protocol encapsulation, operations, administration and management (OAM) signaling, and explicit versioning.

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 5, 2015.

Table of Contents

1. Introduction

LISP [RFC6830] defines an encapsulation format that carries IPv4 or IPv6 (henceforth referred to as IP) packets in a LISP header and outer UDP/IP transport.

The LISP header does not specify the protocol being encapsulated and therefore is currently limited to encapsulating only IP packet payloads. Other protocols, most notably VXLAN [VXLAN] (which defines a similar header format to LISP), are used to encapsulate L2 protocols such as Ethernet. LISP [RFC6830] can be extended to indicate the inner protocol, enabling the encapsulation of Ethernet, IP or any other desired protocol all the while ensuring compatibility with existing LISP [RFC6830] deployments.

As LISP is deployed, there's also the need to provide increased visibility and diagnostic capabilities within the overlay.

This document describes extending LISP ([RFC6830]) via the following changes:

Next Protocol Bit (P bit): A reserved flag bit is allocated, and set in the LISP-gpe header to indicate that a next protocol field is present.

OAM Flag Bit (O bit): A reserved flag bit is allocated, and set in the LISP-gpe header, to indicate that the packet is an OAM packet.

Version: Two reserved bits are allocated, and set in the LISP-gpe header, to indicate LISP-gpe protocol version.

Next protocol: An 8 bit next protocol field is present in the LISP-gpe header.

2. LISP Header Without Protocol Extensions

As described in the introduction, the LISP header has no protocol identifier that indicates the type of payload being carried by LISP. Because of this, LISP is limited to an IP payload. Furthermore, the LISP header has no mechanism to signal OAM packets.

The LISP header contains flags (some defined, some reserved), a Nonce/Map-version field and an instance ID/Locator-status-bit field. The flags provide flexibility to define how the reserved bits can be used to change the definition of the LISP header.

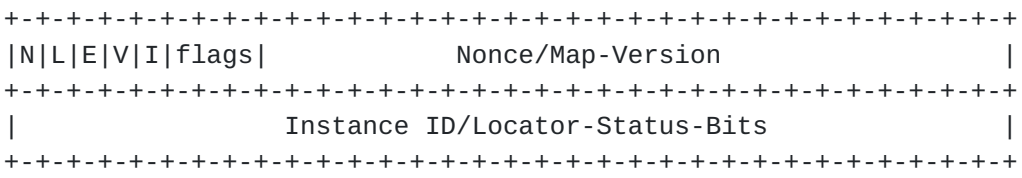

Figure 1: LISP Header

3. Generic Protocol Extension for LISP (LISP-gpe)

3.1. Multi Protocol Support

This draft defines the following changes to the LISP header in order to support multi-protocol encapsulation.

P Bit: Flag bit 5 is defined as the Next Protocol bit. The P bit MUST be set to 1 to indicate the presence of the 8 bit next protocol field.

P = 0 indicates that the payload MUST conform to LISP as defined in [RFC6830].

Flag bit 5 was chosen as the P bit because this flag bit is currently unallocated in LISP [RFC6830].

Next Protocol Field: The lower 8 bits of the first word are used to carry a next protocol. This next protocol field contains the protocol of the encapsulated payload packet.

LISP [RFC6830] uses the lower 16 bits of the first word for either a nonce, an echo-nonce ([RFC6830]) or to support map-versioning ([RFC6834]). These are all optional capabilities that are indicated by setting the N, E, and the V bit respectively.

To maintain the desired data plane compatibility, when the P bit is set, the N, E, and V bits MUST be set to zero.

A new protocol registry will be requested from IANA for the Next Protocol field. This draft defines the following Next Protocol values:

0x1 : IPv4

0x2 : IPv6

0x3 : Ethernet

0x4: Network Service Header

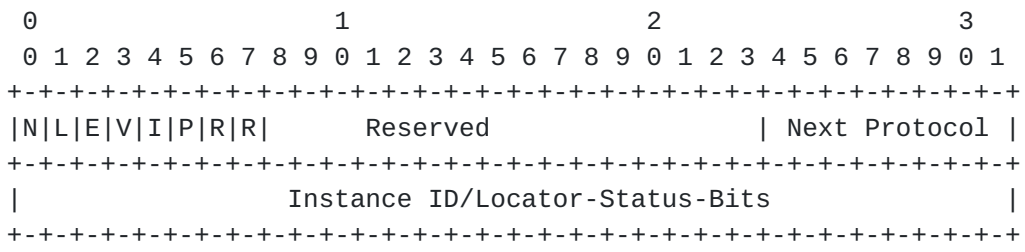

Figure 2: LISP-gpe Next Protocol (P=1)

3.2. OAM Support

Flag bit 7 is defined as the O bit. When the O bit is set to 1, the packet is an OAM packet and OAM processing MUST occur. The OAM protocol details are out of scope for this document. As with the P-bit, bit 7 is currently a reserved flag in [RFC6830].

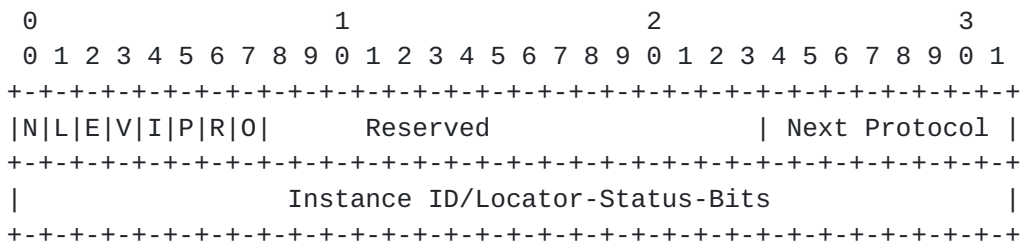

Figure 3: LISP-gpe OAM bit (P=1)

3.3. Version Bits

LISP-gpe bits8 and 9 are defined as version bits. The version field is used to ensure backward compatibility going forward with future LISP-gpe updates.

The initial version for LISP-gpe is 0.

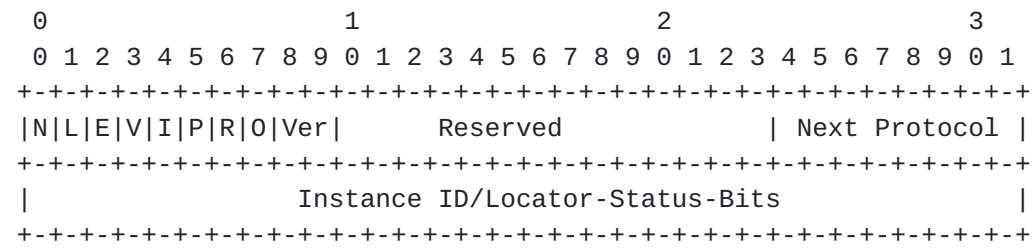

Figure 4: LISP-gpe Version bits (P=1)

4. Backward Compatibility

Undefined (in RFC6830) flag bits 5 and 7, LISP-gpe P and O bits, were selected to ensure compatibility with existing LISP [RFC6830] deployments.

Similarly, using P = 0 to indicate that the format of the header and payload conforms to [RFC6830] ensures compatibility with existing LISP hardware forwarding platforms.

4.1. LISP-gpe Routers to (legacy) LISP Routers

A LISP-gpe router MUST not encapsulate non-IP packet nor OAM packets to a LISP router. A method for determining the capabilities of a LISP router (gpe or "legacy") is out of the scope of this draft.

When encapsulating IP packets to a LISP router the P bit SHOULD be set to 1 and the UDP port MUST be set to 4341. OAM bit MUST be set to 0. The Next Protocol field SHOULD be 0x1 (IPv4) or 0x2 (IPv6). The (legacy) LISP router will ignore the P bit and the protocol type field. The (legacy) LISP router will treat the packet as a LISP packet and inspect the first nibble of the payload to determine the IP version.

When the P bit is set, the N, E, and V bits MUST be set to zero. The receiving (legacy) LISP router will ignore N, E and V bits, when the P bit is set.

4.2. (legacy) LISP Routers to LISP-gpe Routers

When a LISP-gpe router receives a packet from a (legacy) LISP router, the P bit MUST not be set and the UDP port MUST be 4341. The payload MUST be IP, and the LISP-gpe router will inspect the first nibble of the payload to determine IP version.

4.3. Type of Service

When a LISP-gpe router performs Ethernet encapsulation, the inner 802.1Q [IEEE8021Q] priority code point (PCP) field MAY be mapped from the encapsulated frame to the Type of Service field in the outer IPv4 header, or in the case of IPv6 the 'Traffic Class' field.

4.4. VLAN Identifier (VID)

When a LISP-gpe router performs Ethernet encapsulation, the inner header 802.1Q [IEEE8021Q] VLAN Identifier (VID) MAY be mapped to, or used to determine the LISP Instance ID field.

5. LISP-gpe Examples

This section provides two examples of IP protocols, and one example of Ethernet encapsulated LISP-gpe using the generic extension described in this document.

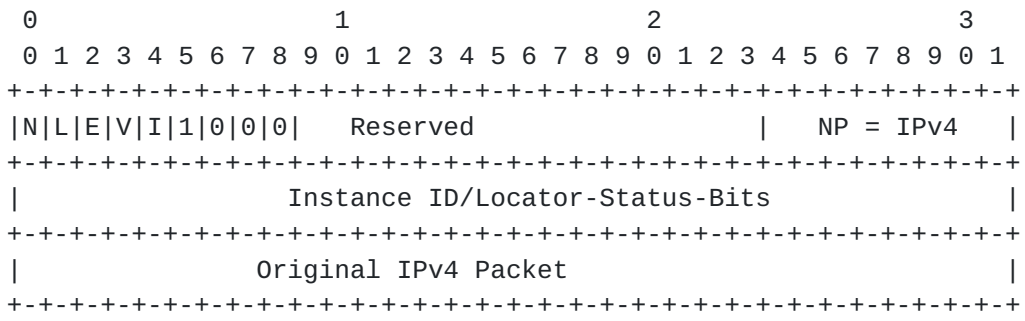

Figure 5: IPv4 and LISP-gpe

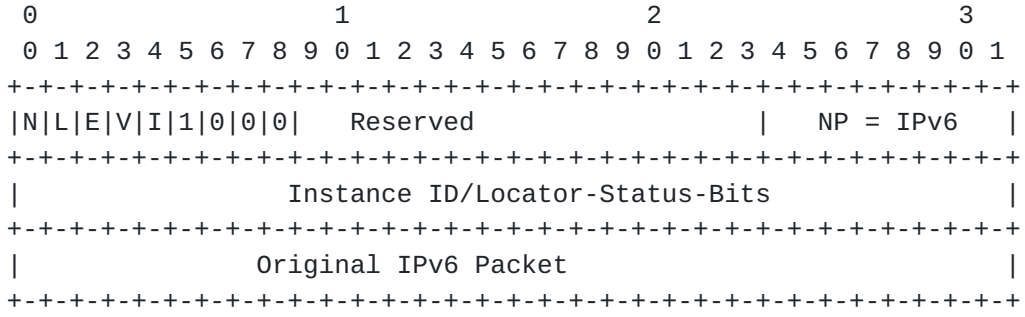

Figure 6: IPv6 and LISP-gpe

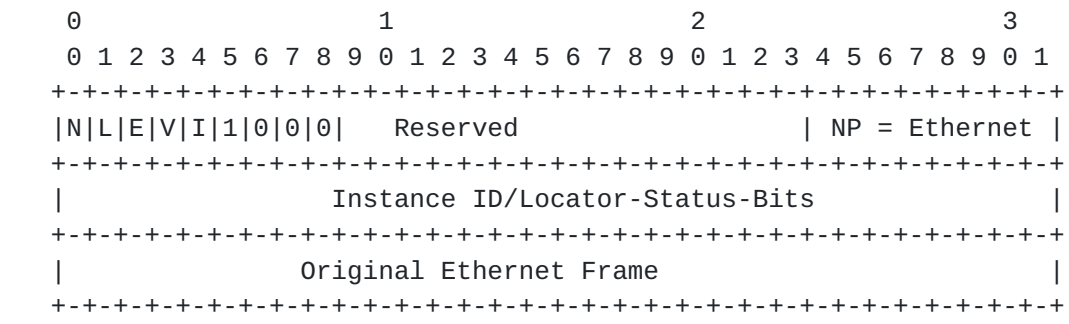

Figure 7: Ethernet and LISP-gpe

6. Security Considerations

LISP-gpe security considerations are similar to the LISP security considerations documented at length in LISP [RFC6830]. With LISP-gpe, issues such as dataplane spoofing, flooding, and traffic redirection are dependent on the particular protocol payload encapsulated.

7. Acknowledgments

A special thank you goes to Dino Farinacci for his guidance and detailed review.

8. IANA Considerations

IANA is requested to set up a registry of "Next Protocol". These are 8-bit values. Next Protocol values 0, 1, 2, 3 and 4 are defined in this draft. New values are assigned via Standards Action [RFC5226].

Next Protocol	Description	Reference
0	Reserved	This document
1	IPv4	This document
2	IPv6	This document
3	Ethernet	This document
4	NSH	This document
5..253	Unassigned	

Table 1

There are ten bits at the beginning of the LISP-gpe header. New bits are assigned via Standards Action [RFC5226].

Bits 0-3 - Assigned by LISP [RFC6830]
Bit 4 - Instance ID (I bit)
Bit 5 - Next Protocol (P bit)
Bit 6 - Reserved
Bit 7 - OAM (O bit)
Bits 8-9 - Version

9. References

9.1. Normative References

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

9.2. Informative References

[ETYPES] The IEEE Registration Authority, "IEEE 802 Numbers", 2012, <http://www.iana.org/assignments/ieee-802-numbers/ieee-802-numbers.xml>.

[IEEE8021Q] The IEEE Computer Society, "Media Access Control (MAC) Bridges and Virtual Bridge Local Area Networks", August 2012, <http://standards.ieee.org/getieee802/download/802.1Q-2011.pdf>.

[RFC1700] Reynolds, J. and J. Postel, "Assigned Numbers", RFC 1700, October 1994.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, January 2013.

[RFC6834] Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", RFC 6834, January 2013.

[VXLAN] Dutt, D., Mahalingam, M., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", 2013.

Authors' Addresses

Darrel Lewis
Cisco Systems, Inc.

Email: darlewis@cisco.com

Puneet Agarwal
Broadcom

Email: pagarwal@broadcom.com

Larry Kreeger
Cisco Systems, Inc.

Email: kreeger@cisco.com

Fabio Maino
Cisco Systems, Inc.

Email: fmaino@cisco.com

Paul Quinn
Cisco Systems, Inc.

Email: paulq@cisco.com

Michael Smith
Cisco Systems, Inc.

Email: michsmit@cisco.com

Navindra Yadav
Cisco Systems, Inc.

Email: nyadav@cisco.com