< draft-ietf-lisp-gpe-09.txt

Internet Engineering Task Force F. Maino, Ed.
Internet-Draft Cisco
Intended status: Standards Track J. Lemon
Expires: April 27, 2020 Broadcom
P. Agarwal
Innovium
D. Lewis
M. Smith
Cisco
October 25, 2019

LISP Generic Protocol Extension
draft-ietf-lisp-gpe-09

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

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 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 April 27, 2020.

skipping to change at *page 2, line 16*

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

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 3, line 9*

format to LISP), are used to encapsulate Layer-2 (L2) protocols such as Ethernet.

This document defines an extension for the LISP header, as defined in [I-D.ietf-lisp-rfc6830bis], to indicate the inner protocol, enabling the encapsulation of Ethernet, IP or any other desired protocol all the while ensuring compatibility with existing LISP deployments.

draft-ietf-lisp-gpe-11.txt >

Internet Engineering Task Force F. Maino, Ed.
Internet-Draft Cisco
Intended status: Standards Track J. Lemon
Expires: May 7, 2020 Broadcom
P. Agarwal
Innovium
D. Lewis
M. Smith
Cisco
November 4, 2019

LISP Generic Protocol Extension
draft-ietf-lisp-gpe-11

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

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 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 May 7, 2020.

skipping to change at *page 2, line 16*

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

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 3, line 9*

format to LISP), are used to encapsulate Layer-2 (L2) protocols such as Ethernet.

This document defines an extension for the LISP header, as defined in [I-D.ietf-lisp-rfc6830bis], to indicate the inner protocol, enabling the encapsulation of Ethernet, IP or any other desired protocol all the while ensuring compatibility with existing LISP deployments.

A flag in the LISP header, called the P-bit, is used to signal the presence of the 8-bit Next Protocol field. The Next Protocol field, when present, uses 8 bits of the field allocated to the echo-noncing and map-versioning features. The two features are still available, albeit with a reduced length of Nonce and Map-Version.

Since all of the reserved bits of the LISP Data-Plane header have been allocated, LISP-GPE can also be used to extend the LISP Data-Plane header by defining Next Protocol "shim" headers that implements new data plane functions not supported in the LISP header. For example, the use of the Group-Based Policy (GBP) header [I-D.lemon-vxlan-lisp-gpe-gbp] or of the In-situ Operations, Administration, and Maintenance (IOAM) header [I-D.brockners-ippm-ioam-vxlan-gpe] with LISP-GPE, can be considered an extension to add support in the Data-Plane for Group-Based Policy functionalities or IOAM metadata.

1.1. Conventions

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.

1.2. Definition of Terms

skipping to change at *page 4, line 25*

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

This document defines two changes to the LISP header in order to support multi-protocol encapsulation: the introduction of the P-bit and the definition of a Next Protocol field. This is shown in Figure 2 and described below.

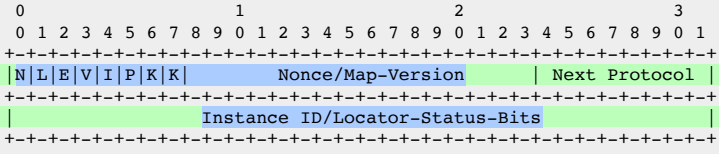

Figure 2: LISP-GPE Header

P-Bit: Flag bit 5 is defined as the Next Protocol bit.

If the P-bit is clear (0) the LISP header is bit-by-bit equivalent to the definition in [I-D.ietf-lisp-rfc6830bis].

The P-bit is set to 1 to indicate the presence of the 8 bit Next Protocol field. The combinations of bits that are allowed when the P-bit is set are the same allowed by [I-D.ietf-lisp-rfc6830bis].

Nonce/Map-Version: In [I-D.ietf-lisp-6834bis], LISP uses the lower 24 bits of the first word for a nonce, an echo-nonce, or to support map- versioning. These are all optional capabilities that are indicated in the LISP header by setting the N, E, and V bits respectively.

When the P-bit and the N-bit are set to 1, the Nonce field is the middle 16 bits (i.e., encoded in 16 bits, not 24 bits). Note that the E-bit only has meaning when the N-bit is set.

When the P-bit and the V-bit are set to 1, the Version fields use the middle 16 bits: the Source Map-Version uses the high-order 8 bits, and the Dest Map-Version uses the low-order 8 bits.

When the P-bit is set to 1 and the N-bit and the V-bit are both 0, the middle 16-bits MUST be set to 0 on transmission and ignored on receipt.

The encoding of the Nonce field in LISP-GPE, compared with the one used in [I-D.ietf-lisp-rfc6830bis] for the LISP data plane encapsulation, reduces the length of the nonce from 24 to 16 bits. As per [I-D.ietf-lisp-rfc6830bis], Ingress Tunnel Routers (ITRs)

A flag in the LISP header, called the P-bit, is used to signal the presence of the 8-bit Next Protocol field. The Next Protocol field, when present, uses 8 bits of the field that was allocated to the echo-noncing and map-versioning features in [I-D.ietf-lisp-rfc6830bis].

Since all of the reserved bits of the LISP Data-Plane header have been allocated, LISP-GPE can also be used to extend the LISP Data-Plane header by defining Next Protocol "shim" headers that implements new data plane functions not supported in the LISP header. For example, the use of the Group-Based Policy (GBP) header [I-D.lemon-vxlan-lisp-gpe-gbp] or of the In-situ Operations, Administration, and Maintenance (IOAM) header [I-D.brockners-ippm-ioam-vxlan-gpe] with LISP-GPE, can be considered an extension to add support in the Data-Plane for Group-Based Policy functionalities or IOAM metadata.

Nonce, Map-Versioning and Locator Status Bit fields are not part of the LISP-GPE header. Shim headers can be used to specify features such as echo-noncing, map-versioning or reachability by defining fields of the same size, or larger, of those specified in [I-D.ietf-lisp-rfc6830bis].

1.1. Conventions

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.

1.2. Definition of Terms

skipping to change at *page 4, line 27*

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

This document defines two changes to the LISP header in order to support multi-protocol encapsulation: the introduction of the P-bit and the definition of a Next Protocol field. This is shown in Figure 2 and described below.

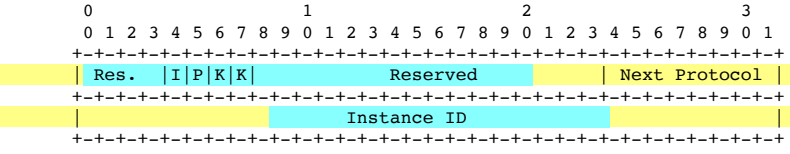

Figure 2: LISP-GPE Header

Bits 0-3 and 8-23: Bits 0-3 and 8-23 of the LISP-GPE header are Reserved. They MUST be set to zero on transmission and ignored on receipt.

Features that were implemented with bits 0-3 in [I-D.ietf-lisp-rfc6830bis], such as echo-noncing, map-versioning and reachability, can be implemented by defining the appropriate shim headers.

Instance ID When the I-Bit is set to 1 the high-order 24 bits of the Instance ID field are used as an Instance ID, as specified in [I-D.ietf-lisp-rfc6830bis]. The low-order 8 bits are set to zero, as the Locator-Status-Bits feature is not supported in LISP-GPE.

P-Bit: Flag bit 5 is defined as the Next Protocol bit.

If the P-bit is clear (0) the LISP header is bit-by-bit equivalent to the definition in [I-D.ietf-lisp-rfc6830bis] with bits N, L, E and V set to 0.

The P-bit is set to 1 to indicate the presence of the 8 bit Next Protocol field. The combinations of bits that are allowed when the P-bit is set are the same allowed by [I-D.ietf-lisp-rfc6830bis] when bits N, L, E and V are set to 0.

are required to generate different nonces when sending to different Routing Locators (RLOCs), but the same nonce can be used for a period of time when encapsulating to the same Egress Tunnel Router (ETR). The use of 16 bits nonces still allows an ITR to determine to and from reachability for up to 64k RLOCs at the same time, but reduces the overall robustness of the nonce mechanism to off-path attackers. Please refer to Section Section 7 for security considerations that apply to the use of the Nonce field.

Similarly, the encoding of the Source and Dest Map-Version fields, compared with [I-D.ietf-lisp-rfc6830bis], is reduced from 12 to 8 bits. This allows to associate only 256 different versions to each Endpoint Identifier to Routing Locator (EID-to-RLOC) mapping to inform commmunicating ITRs and ETRs about modifications of the mapping, reducing the Map-versioning wrap-around time. Please refer to Section Section 7 for security considerations that apply to the use of the Map-Versioning field.

Next Protocol: The lower 8 bits of the first 32-bit word are used to carry a Next Protocol. This Next Protocol field contains the protocol of the encapsulated payload packet.

This document defines the following Next Protocol values:

0x01 : IPv4

0x02 : IPv6

skipping to change at *page 12, line 46*

skipping to change at *page 12, line 43*

"Multiple Data-Planes Encapsulation Bitmap" registry assigning a value to bit 24 for the LISP-GPE encapsulation, assigning bits 25-31 values that are conformant with RFC8060. This will allow future allocation of values 0-23.

7. Security Considerations

LISP-GPE security considerations are similar to the LISP security considerations and mitigation techniques documented in [RFC7835].

The Echo Nonce Algorithm described in [I-D.ietf-lisp-rfc6830bis] relies on the nonce to detect reachability from ITR to ETR. In LISP-GPE the use of a 16-bit nonce, compared with the 24-bit nonce used in LISP, increases the probability of an off-path attacker to correctly guess the nonce and force the ITR to believe that a non-reachable RLOC is reachable. However, the use of common anti-spoofing mechanisms such as uRPF partially mitigates this form of attack.

The considerations made in [I-D.ietf-lisp-rfc6830bis] that Echo Nonce, Map-Versioning, and Locator-Status-Bits SHOULD NOT be used over the public Internet and SHOULD only be used in trusted and closed deployments apply to LISP-GPE as well. These considerations are even more important for LISP-GPE, considering the reduced size of the Nonce/Map-versioning field.

LISP-GPE, as many encapsulations that use optional extensions, is subject to on-path adversaries that by manipulating the g-Bit and the packet itself can remove part of the payload. Typical integrity protection mechanisms (such as IPsec) SHOULD be used in combination with LISP-GPE by those protocol extensions that want to protect from on-path attackers.

With LISP-GPE, issues such as data-plane spoofing, flooding, and traffic redirection may depend on the particular protocol payload encapsulated.

skipping to change at *page 14, line 9*

skipping to change at *page 13, line 37*

- Larry Kreeger
- John Lemon, Broadcom
- Puneet Agarwal, Innovium

9. References

9.1. Normative References

[I-D.ietf-lisp-6834bis]
Iannone, L., Saucez, D., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Map-Versioning", draft-ietf-lisp-6834bis-04 (work in progress), August 2019.

[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-27 (work in progress), June 2019.

[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,

End of changes. 19 change blocks.

80 lines changed or deleted ***41 lines changed or added***