< draft-ietf-lisp-gpe-06.txt

Internet Engineering Task Force F. Maino, Ed.
Internet-Draft Cisco
Intended status: Standards Track J. Lemon
Expires: March 24, 2019 Broadcom
P. Agarwal
Innovium
D. Lewis
M. Smith
Cisco
September 20, 2018

LISP Generic Protocol Extension
draft-ietf-lisp-gpe-06

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 March 24, 2019.

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 11*

[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.

LISP-GPE MAY also be used to extend the LISP Data-Plane header, that has allocated all by defining a Next Protocol "shim" header that implements new data plane functions not supported in the LISP header. As an example, the use of the Network Service Header (NSH) with LISP-GPE, can be considered an extension to add support in the Data-Plane for Network Service Chaining functionalities.

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.

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 13*

[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

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 34*

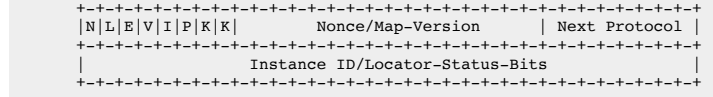

Figure 2: LISP-GPE Header

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

If the P-bit is clear (0) the LISP header conforms 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.

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.

skipping to change at *page 5, line 26*

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) 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.

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 still allows to associate 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.

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:

0x1 : IPv4

0x2 : IPv6

0x3 : Ethernet

0x4 : Network Service Header (NSH) [RFC8300]

The values are tracked in an IANA registry as described in Section 5.1.

3.1. Payload Specific Transport Interactions

To ensure that protocols that are encapsulated in LISP-GPE will work well from a transport interaction perspective, the specification of a new encapsulated payload MUST contain an analysis of how LISP-GPE SHOULD deal with outer UDP Checksum, DSCP mapping, and Explicit Congestion Notification (ECN) bits whenever they apply to the new encapsulated payload.

For IP payloads, section 5.3 of [I-D.ietf-lisp-rfc6830bis] specifies how to handle UDP Checksums encouraging implementors to consider UDP checksum usage guidelines in section 3.4 of [RFC8085] when it is desirable to protect UDP and LISP headers against corruption. Each new encapsulated payloads, when registered with LISP-GPE, MUST be accompanied by a similar analysis.

skipping to change at *page 4, line 34*

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.

skipping to change at *page 5, line 26*

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) 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

0x03 : Ethernet
0x04 : Network Service Header (NSH) [RFC8300]

0x05 to 0x7F: Unassigned

0x80 to 0xFF: Unassigned (shim headers)

The values are tracked in an IANA registry as described in Section 6.1.

Next protocol values from 0x80 to 0xFF are assigned to protocols encoded as generic "shim" headers. Shim protocols all use a common header structure, which includes a next header field using the same values as described above. When a shim header protocol is used with other data described by 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 what follows the shim protocol.

Implementations that are not aware of a given shim header MUST ignore the header and proceed to parse the next protocol. Shim protocols MUST have the first 32 bits defined as:

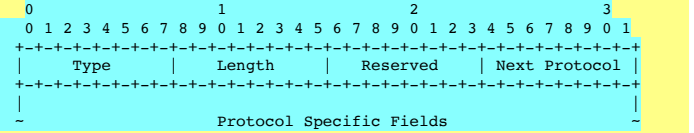

```
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: Shim Header

Where:

Type: This field identifies the different messages of this protocol.

Length: The length, in 4-octect units, of this protocol message not including the first 4 octects.

Reserved: The use of this field is reserved to the protocol defined in this message.

Next Protocol Field: This next protocol field contains the protocol of the encapsulated payload. The protocol registry will be requested from IANA as per section 10.2.

4. Implementation and Deployment Considerations

4.1. Applicability Statement

LISP-GPE conforms, as an UDP-based encapsulation protocol, to the UDP usage guidelines as specified in [RFC8085]. The applicability of these guidelines are dependent on the underlay IP network and the nature of the encapsulated payload.

[RFC8085] outlines two applicability scenarios for UDP applications, 1) general Internet and 2) controlled environment. The controlled environment means a single administrative domain or adjacent set of cooperating domains. A network in a controlled environment can be managed to operate under certain conditions whereas in general Internet this cannot be done. Hence requirements for a tunnel protocol operating under a controlled environment can be less restrictive than the requirements of general internet.

LISP-GPE scope of applicability is the same set of use cases covered by[I-D.ietf-lisp-rfc6830bis] for the LISP dataplane protocol. The common property of these use cases is a large set of cooperating entities seeking to communicate over the public Internet or other large underlay IP infrastructures, while keeping the addressing and topology of the cooperating entities separate from the underlay and Internet topology, routing, and addressing.

LISP-GPE is meant to be deployed in network environments operated by a single operator or adjacent set of cooperating network operators that fits with the definition of controlled environments in [RFC8085].

For the purpose of this document, a traffic-managed controlled environment (TMCE), outlined in [RFC8086], is defined as an IP network that is traffic-engineered and/or otherwise managed (e.g., via use of traffic rate limiters) to avoid congestion. Significant portions of text in this Section are based on [RFC8086].

It is the responsibility of the network operators to ensure that the guidelines/requirements in this section are followed as applicable to their LISP-GPE deployments

4.2. Congestion Control Functionality

LISP-GPE does not natively provide congestion control functionality and relies on the payload protocol traffic for congestion control. As such LISP-GPE MUST be used with congestion controlled traffic or within a network that is traffic managed to avoid congestion (TMCE). An operator of a traffic managed network (TMCE) may avoid congestion by careful provisioning of their networks, rate-limiting of user data traffic and traffic engineering according to path capacity.

Encapsulated payloads may have a priority field that may or may not be mapped to the DSCP field of the outer IP header (part of Type of Service in IPv4 or Traffic Class in IPv6). Such new encapsulated payloads, when registered with LISP-GPE, MUST be accompanied by an analysis similar to the one performed in Section 3.1.1 of this document for Ethernet payloads.

Encapsulated payloads may have Explicit Congestion Notification mechanisms that may or may not be mapped to the outer IP header ECN field. Such new encapsulated payoads, when registered with LISP-GPE, MUST be accompanied by a set of guidelines derived from [RFC6040].

Encapsulated payloads may have Explicit Congestion Notification mechanisms that may or may not be mapped to the outer IP header ECN field. Such new encapsulated payoads, when registered with LISP-GPE, MUST be accompanied by a set of guidelines derived from [I-D.ietf-tsvwg-ecn-encap-guidelines] and [RFC6040].

The rest of this section specifies payload specific transport interactions considerations for the two new LISP-GPE encapsulated payloads specified in this document: Ethernet and NSH.

4.3. UDP Checksum

3.1.1. Payload Specific Transport Interactions for Ethernet Encapsulated Payloads

For IP payloads, section 5.3 of [I-D.ietf-lisp-rfc6830bis] specifies how to handle UDP Checksums encouraging implementors to consider UDP checksum usage guidelines in section 3.4 of [RFC8085] when it is desirable to protect UDP and LISP headers against corruption.

The UDP Checksum considerations specified in section 5.3 of [I-D.ietf-lisp-rfc6830bis] apply to Ethernet Encapsulated Payloads. Implementors are encouraged to consider the UDP checksum usage guidelines in section 3.4 of [RFC8085] when it is desirable to protect UDP, LISP and Ethernet headers against corruption.

In order to provide integrity of LISP-GPE headers, options and payload, for example to avoid mis-delivery of payload to different tenant systems in case of data corruption, outer UDP checksum SHOULD be used with LISP-GPE when transported over IPv4. The UDP checksum provides a statistical guarantee that a payload was not corrupted in transit. These integrity checks are not strong from a coding or cryptographic perspective and are not designed to detect physical-layer errors or malicious modification of the datagram (see Section 3.4 of [RFC8085]). In deployments where such a risk exists, an operator SHOULD use additional data integrity mechanisms such as offered by IPSec.

An operator MAY choose to disable UDP checksum and use zero checksum if LISP-GPE packet integrity is provided by other data integrity mechanisms such as IPsec or additional checksums or if one of the conditions in Section 4.3.1 a, b, c are met.

By default, UDP checksum MUST be used when LISP-GPE is transported over IPv6. A tunnel endpoint MAY be configured for use with zero UDP checksum if additional requirements in Section 4.3.1 are met.

4.3.1. UDP Zero Checksum Handling with IPv6

When LISP-GPE is used over IPv6, UDP checksum is used to protect IPv6

headers, UDP headers and LISP-GPE headers and payload from potential data corruption. As such by default LISP-GPE MUST use UDP checksum when transported over IPv6. An operator MAY choose to configure to operate with zero UDP checksum if operating in a traffic managed controlled environment as stated in Section 4.1 if one of the following conditions are met:

a. It is known that the packet corruption is exceptionally unlikely (perhaps based on knowledge of equipment types in their underlay network) and the operator is willing to take a risk of undetected packet corruption

b. It is judged through observational measurements (perhaps through historic or current traffic flows that use non zero checksum) that the level of packet corruption is tolerably low and where the operator is willing to take the risk of undetected corruption

c. LISP-GPE payload is carrying applications that are tolerant of misdelivered or corrupted packets (perhaps through higher layer checksum validation and/or reliability through retransmission)

In addition LISP-GPE tunnel implementations using Zero UDP checksum MUST meet the following requirements:

1. Use of UDP checksum over IPv6 MUST be the default configuration for all LISP-GPE tunnels

2. If LISP-GPE is used with zero UDP checksum over IPv6 then such xTR implementation MUST meet all the requirements specified in section 4 of [RFC6936] and requirements 1 as specified in section 5 of [RFC6936]

3. The ETR that decapsulates the packet SHOULD check the source and destination IPv6 addresses are valid for the LISP-GPE tunnel that is configured to receive Zero UDP checksum and discard other packets for which such check fails

4. The ITR that encapsulates the packet MAY use different IPv6 source addresses for each LISP-GPE tunnel that uses Zero UDP checksum mode in order to strengthen the decapsulator's check of the IPv6 source address (i.e the same IPv6 source address is not to be used with more than one IPv6 destination address, irrespective of whether that destination address is a unicast or multicast address). When this is not possible, it is RECOMMENDED to use each source address for as few LISP-GPE tunnels that use zero UDP checksum as is feasible

5. Measures SHOULD be taken to prevent LISP-GPE traffic over IPv6 with zero UDP checksum from escaping into the general Internet. Examples of such measures include employing packet filters at the PETR and/or keeping logical or physical separation of LISP network from networks carrying General Internet

The above requirements do not change either the requirements specified in [RFC2460] as modified by [RFC6935] or the requirements 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] 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.

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.

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.

3.1.2. Payload Specific Transport Interactions for NSH Encapsulated Payloads

The UDP Checksum considerations specified in section 5.3 of [I-D.ietf-lisp-rfc6830bis] apply to NSH Encapsulated Payloads. Implementors are encouraged to consider the UDP checksum usage guidelines in section 3.4 of [RFC8085] when it is desirable to protect UDP, LISP, and NSH headers against corruption.

When a LISP-GPE router performs an NSH encapsulation, DSCP and ECN values MAY be mapped as specified for the Next Protocol encapsulated by NSH (namely IPv4, IPv6 and Ethernet).

4. Backward Compatibility

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

The next Section describes a method to determine the Data-Plane capabilities of a LISP ETR, based on the use of the "Multiple Data-Planes" LISP Canonical Address Format (LCAF) type defined in [RFC8060]. Other mechanisms can be used, including static ETR/ITR (xTR) configuration, but are out of the scope of this document.

When encapsulating IP packets to a non LISP-GPE capable router the P-bit MUST be set to 0. That is, the encapsulation format defined in this document MUST NOT be sent to a router that has not indicated that it supports this specification because such a router would ignore the P-bit (as described in [I-D.ietf-lisp-rfc6830bis]) and so would misinterpret the other LISP header fields possibly causing significant errors.

A LISP-GPE router MUST NOT encapsulate non-IP packets (that have the

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

The next Section describes a method to determine the Data-Plane capabilities of a LISP ETR, based on the use of the "Multiple Data-Planes" LISP Canonical Address Format (LCAF) type defined in [RFC8060]. Other mechanisms can be used, including static ETR/ITR (xTR) configuration, but are out of the scope of this document.

When encapsulating IP packets to a non LISP-GPE capable router the P-bit MUST be set to 0. That is, the encapsulation format defined in this document MUST NOT be sent to a router that has not indicated that it supports this specification because such a router would ignore the P-bit (as described in [I-D.ietf-lisp-rfc6830bis]) and so would misinterpret the other LISP header fields possibly causing significant errors.

5. Backward Compatibility

5.1. Use of "Multiple Data-Planes" LCAF to Determine ETR Capabilities

P-bit set to 1) to a non-LISP-GPE capable router.

4.1. Use of "Multiple Data-Planes" LCAF to Determine ETR Capabilities

LISP Canonical Address Format (LCAF) [RFC8060] defines the "Multiple Data-Planes" LCAF type, that can be included by an ETR in a Map-Reply to encode the encapsulation formats supported by a given RLOC. In this way an ITR can be made aware of the capability to support LISP-GPE, as well as other encapsulations, on a given RLOC of that ETR.

The 3rd 32-bit word of the "Multiple Data-Planes" LCAF type, as defined in [RFC8060], is a bitmap whose bits are set to one (1) to represent support for each Data-Plane encapsulation. The values are tracked in an IANA registry as described in Section 5.2.

This document defines bit 24 in the third 32-bit word of the "Multiple Data-Planes" LCAF as:

g-Bit: The RLOCs listed in the Address Family Identifier (AFI) encoded addresses in the next longword can accept LISP-GPE (Generic Protocol Extension) encapsulation using destination UDP port 4341

5. IANA Considerations

5.1. 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 via Standards Action [RFC8126]. The protocols that are being assigned values do not themselves need to be IETF standards track protocols.

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..255	Unassigned	

5.2. Multiple Data-Planes Encapsulation Bitmap Registry

IANA is requested to set up a registry of "Multiple Data-Planes Encapsulation Bitmap" to identify the encapsulations supported by an ETR in the Multiple Data-Planes LCAF Type defined in [RFC8060]. The bitmap is the 3rd 32-bit word of the Multiple Data-Planes LCAF type. Each bit of the bitmap represents a Data-Plane Encapsulation. New values are assigned via Standards Action [RFC8126].

Bits 0-23 are unassigned. This document assigns bit 24 (g-bit) to LISP-GPE. Bits 25-31 are assigned in [RFC8060]).

Bit Position	Bit Name	Assigned to	Reference
0-23		Unassigned	
24	g	LISP Generic Protocol Extension (LISP-GPE)	This Document
25	U	Generic UDP Encapsulation (GUE)	[RFC8060]
26	G	Generic Network Virtualization Encapsulation (GENEVE)	[RFC8060]
27	N	Network Virtualization - Generic Routing Encapsulation (NV-GRE)	[RFC8060]
28	v	VXLAN Generic Protocol Extension (VXLAN-GPE)	[RFC8060]
29	V	Virtual eXtensible Local Area Network (VXLAN)	[RFC8060]
30	l	Layer 2 LISP (LISP-L2)	[RFC8060]
31	L	Locator/ID Separation Protocol (LISP)	[RFC8060]

6. 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 prevents this form of attack.

LISP Canonical Address Format (LCAF) [RFC8060] defines the "Multiple Data-Planes" LCAF type, that can be included by an ETR in a Map-Reply to encode the encapsulation formats supported by a given RLOC. In this way an ITR can be made aware of the capability to support LISP-GPE, as well as other encapsulations, on a given RLOC of that ETR.

The 3rd 32-bit word of the "Multiple Data-Planes" LCAF type, as defined in [RFC8060], is a bitmap whose bits are set to one (1) to represent support for each Data-Plane encapsulation. The values are tracked in an IANA registry as described in Section 6.2.

This document defines bit 24 in the third 32-bit word of the "Multiple Data-Planes" LCAF as:

g-Bit: The RLOCs listed in the Address Family Identifier (AFI) encoded addresses in the next longword can accept LISP-GPE (Generic Protocol Extension) encapsulation using destination UDP port 4341

6. IANA Considerations

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	

6.2. Multiple Data-Planes Encapsulation Bitmap Registry

IANA is requested to set up a registry of "Multiple Data-Planes Encapsulation Bitmap" to identify the encapsulations supported by an ETR in the Multiple Data-Planes LCAF Type defined in [RFC8060]. The bitmap is the 3rd 32-bit word of the Multiple Data-Planes LCAF type. Each bit of the bitmap represents a Data-Plane Encapsulation. New values are assigned under the Specification Required policy [RFC8126].

Bits 0-23 are unassigned. This document assigns bits 24-31. Bit 24 (bit 'g') is assigned to LISP-GPE.

Bit Position	Bit Name	Assigned to	Reference
0-23		Unassigned	
24	g	LISP Generic Protocol Extension (LISP-GPE)	This Document
25	U	Generic UDP Encapsulation (GUE)	This Document
26	G	Generic Network Virtualization Encapsulation (GENEVE)	This Document
27	N	Network Virtualization - Generic Routing Encapsulation (NV-GRE)	This Document
28	v	VXLAN Generic Protocol Extension (VXLAN-GPE)	This Document
29	V	Virtual eXtensible Local Area Network (VXLAN)	This Document
30	l	Layer 2 LISP (LISP-L2)	This Document
31	L	Locator/ID Separation Protocol (LISP)	This Document

Editorial Note (The following paragraph to be removed by the RFC Editor before publication)

The "Multiple Data-Planes Encapsulation Bitmap" was "hardcoded" in RFC8060, assigning values to bits 25-31. This draft allocates the "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.

8. Acknowledgements and Contributors

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

This Working Group (WG) document originated as draft-lewis-lisp-gpe; the following are its coauthors and contributors along with their respective affiliations at the time of WG adoption. The editor of this document would like to thank and recognize them and their contributions. These coauthors and contributors provided invaluable concepts and content for this document's creation.

- Darrel Lewis, Cisco Systems, Inc.
- Fabio Maino, Cisco Systems, Inc.
- Michael Smith, Cisco Systems, Inc.
- Navindra Yadav, Cisco Systems, Inc.
- 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, <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, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

9.2. Informative References

[I-D.brockners-ippm-ioam-vxlan-gpe]
Brockners, F., Bhandari, S., Govindan, V., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Kfir, A., Gafni, B., Lapukhov, P., and M. Spiegel, "VXLAN-GPE Encapsulation for In-situ OAM Data", draft-brockners-ippm-ioam-vxlan-gpe-02 (work in progress), July 2019.

[I-D.ietf-tsvwg-ecn-encap-guidelines]
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.

[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>.

[RFC6040] Briscoe, B., "Tunnelling of Explicit Congestion Notification", RFC 6040, DOI 10.17487/RFC6040, November 2010, <https://www.rfc-editor.org/info/rfc6040>.

[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>.

[RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <https://www.rfc-editor.org/info/rfc7348>.

[RFC7835] Saucez, D., Iannone, L., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Threat Analysis", RFC 7835, DOI 10.17487/RFC7835, April 2016, <https://www.rfc-editor.org/info/rfc7835>.

[RFC8060] Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format (LCAF)", RFC 8060, DOI 10.17487/RFC8060, February 2017, <https://www.rfc-editor.org/info/rfc8060>.

[RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <https://www.rfc-editor.org/info/rfc8085>.

[RFC8086] Yong, L., Ed., Crabbe, E., Xu, X., and T. Herbert, "GRE-in-UDP Encapsulation", RFC 8086, DOI 10.17487/RFC8086, March 2017, <https://www.rfc-editor.org/info/rfc8086>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

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

[RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed.,

skipping to change at *page 12, line 19*

[RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <https://www.rfc-editor.org/info/rfc7348>.

[RFC7835] Saucez, D., Iannone, L., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Threat Analysis", RFC 7835, DOI 10.17487/RFC7835, April 2016, <https://www.rfc-editor.org/info/rfc7835>.

[RFC8060] Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format (LCAF)", RFC 8060, DOI 10.17487/RFC8060, February 2017, <https://www.rfc-editor.org/info/rfc8060>.

[RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <https://www.rfc-editor.org/info/rfc8085>.

[RFC8086] Yong, L., Ed., Crabbe, E., Xu, X., and T. Herbert, "GRE-in-UDP Encapsulation", RFC 8086, DOI 10.17487/RFC8086, March 2017, <https://www.rfc-editor.org/info/rfc8086>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

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

[RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed.,

skipping to change at *page 16, line 23*

Authors' Addresses

Fabio Maino (editor)
Cisco Systems
San Jose, CA 95134
USA

Email: fmaino@cisco.com

John Lemon
Broadcom
270 Innovation Drive
San Jose, CA 95134
USA

Email: john.lemon@broadcom.com

Puneet Agarwal
Innovium
USA

Email: puneet@acm.org

Darrel Lewis
Cisco Systems

Authors' Addresses

Fabio Maino (editor)
Cisco Systems
San Jose, CA 95134
USA

Email: fmaino@cisco.com

Jennifer Lemon
Broadcom
270 Innovation Drive
San Jose, CA 95134
USA

Email: jennifer.lemon@broadcom.com

Puneet Agarwal
Innovium
USA

Email: puneet@acm.org

Darrel Lewis
Cisco Systems

End of changes. 50 change blocks.

132 lines changed or deleted ***340 lines changed or added***