

draft-ietf-lisp-gpe-11

F. Maino, J. Lemon, P. Agarwal, D. Lewis, M. Smith

IETF 106 – Singapore

Nov. 2019

Changes since -06

- Removed support for Nonce, Map-Version, and LSB (section 3)
- “Shim” Headers (section 3)
- Implementation and deployment considerations (new section 4)
- Multiple Data-Planes Encapsulation Bitmap Registry (section 6)

Removed support for Nonce, Map-Version, and LSB




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-nouncing, 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.

Shim Headers

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 value 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.

Transit nodes that are not aware of a given shim header **type** **MUST** ignore the shim header and proceed to parse the next protocol.

Shim headers can be used to incrementally deploy new GPE features without updating the implementation of each transit node between two tunnel endpoints, and without punting the packet with shim headers of unknown type to the 'slow' path.

Shim Header Format

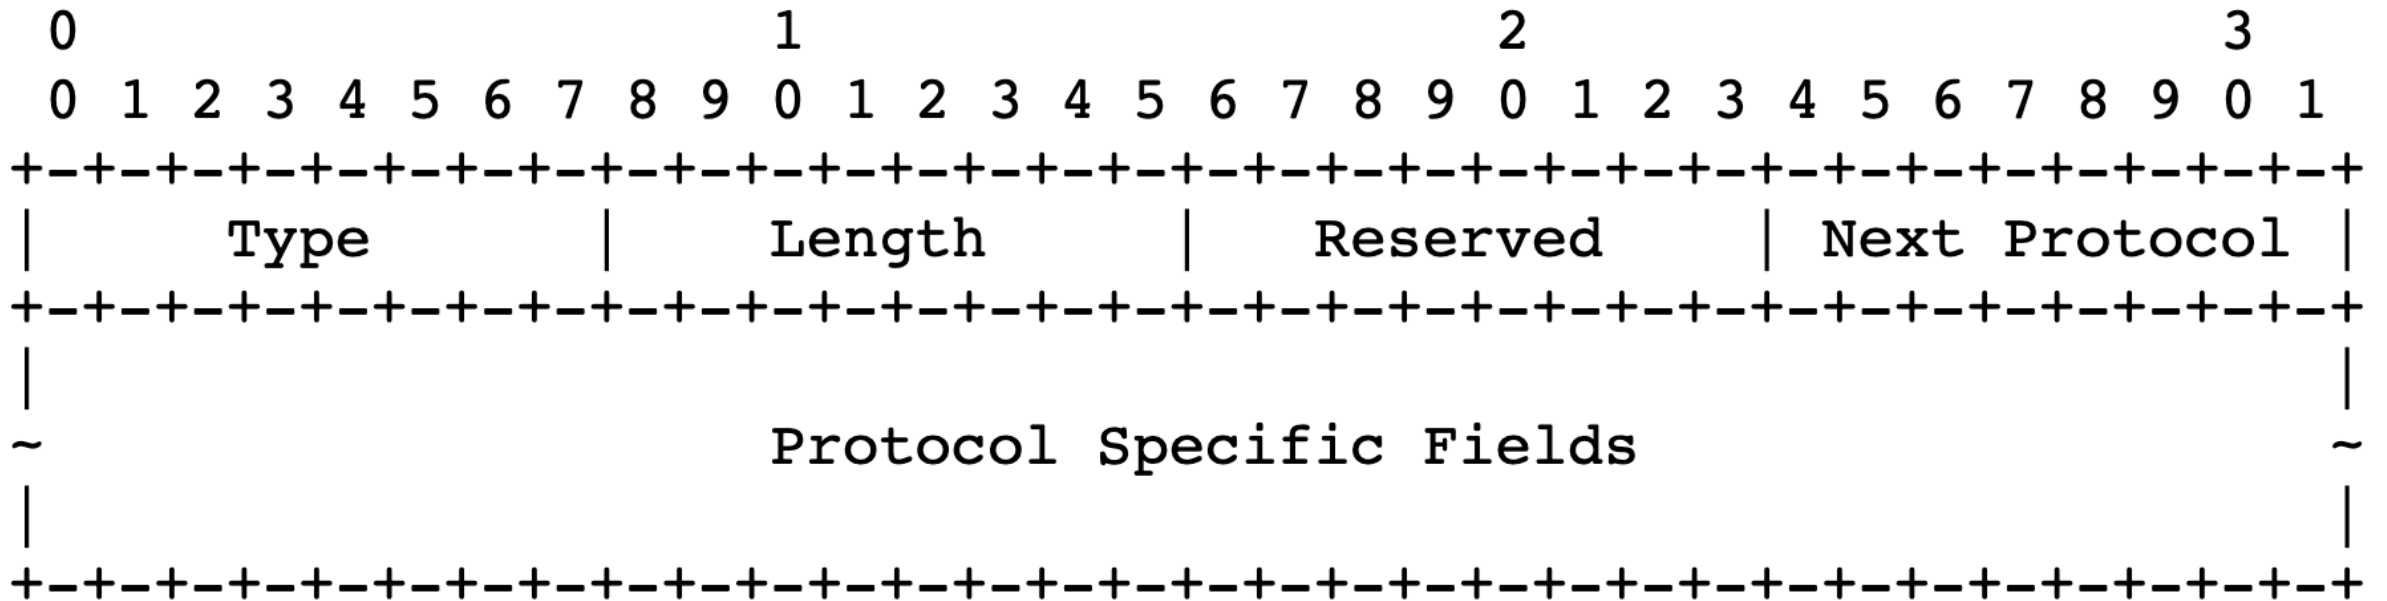


Figure 3: Shim Header

Shim Headers Examples

- GBP - draft-lemon-vxlan-lisp-gpe-gbp
- iOAM - draft-brockners-ippm-ioam-vxlan-gpe-03
- Echo noncing, map-versioning, LSB

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].

Implementation and deployment considerations (new section 4)


- Considerations about Congestion Control, UDP Checksum, PCP/ToS bits in ethernet frames are now in section 4
- We followed the guidelines of RFC8085 (UDP Usage Guidelines) and RFC8086 (GRE-in-UDP Encapsulation)

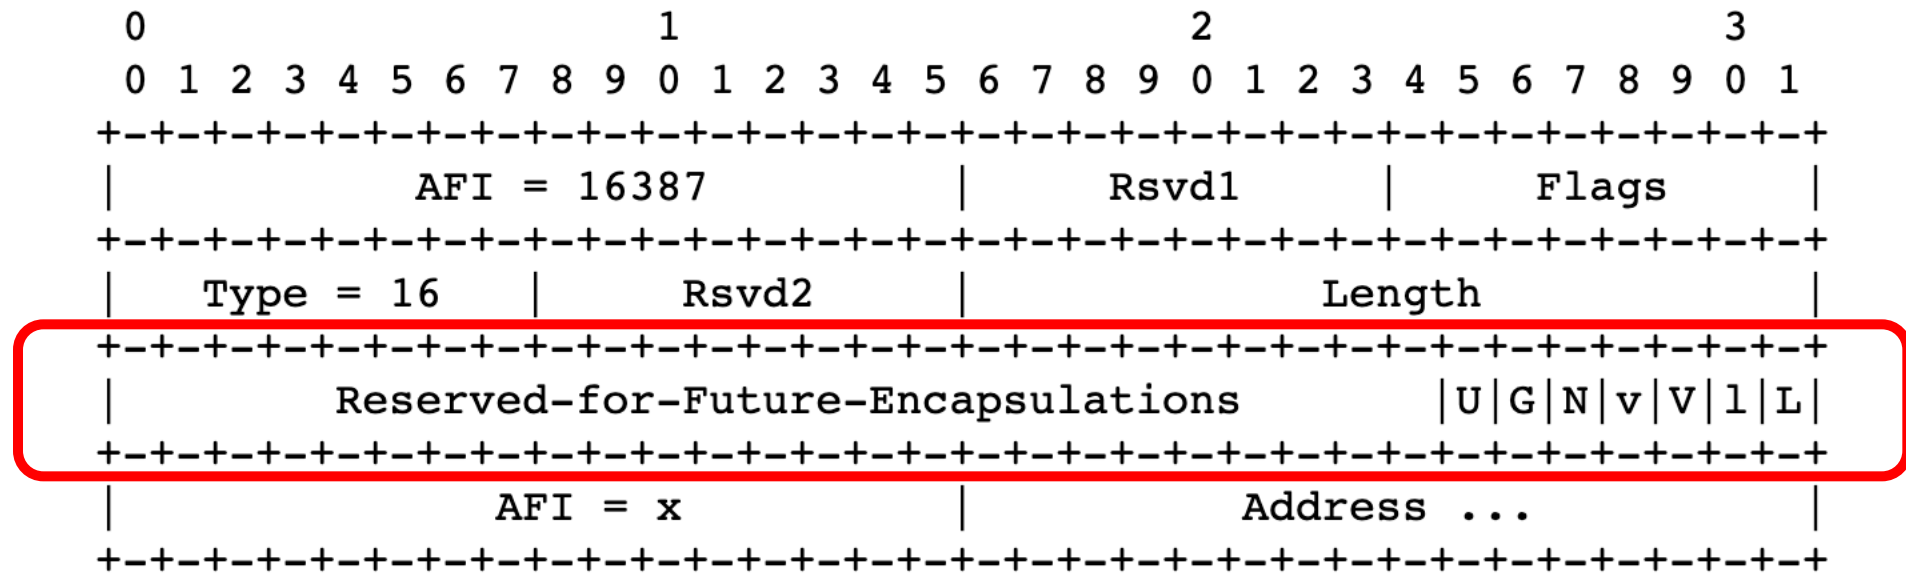
4.	Implementation and Deployment Considerations	7
4.1.	Applicability Statement	7
4.2.	Congestion Control Functionality	7
4.3.	UDP Checksum	8
4.3.1.	UDP Zero Checksum Handling with IPv6	8
4.4.	Ethernet Encapsulated Payloads	10

Applicability Statement: Traffic-Managed Controlled Environment (TMCE)

- RFC8085 outlines two applicability scenarios for UDP applications,
 - 1) general Internet and 2) controlled environment.
- 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.

Multiple Data-Planes Encapsulation Bitmap Registry (section 6)

- RFC8060 defines a “Multiple Data-planes” LCAF type (16)
 - Used to represent the encapsulation formats supported by an RLOC
- No IANA registry was defined to support future encapsualtions
 - Allocating a registry in LISP-GPE is a “Down Ref” to an Experimental RFC 



Multiple Data-Planes Encapsulation Bitmap Registry (section 6): Addressing the “Down Ref”

- Section 5.1 will exclude “Detection of ETR Capabilities” from the scope of LISP-GPE

Next Steps

- Publish rev -12 ASAP to address the comments received so far
- Document should go back to Last Call