

INTAREA/Network Working Group
Internet Draft
Intended status: Informational
Expires: October 1, 2021

J. Zhu
Intel
S. Kanugovi
Nokia
April 1, 2021

Generic Multi-Access (GMA) Encapsulation Protocol
draft-zhu-intarea-gma-09

Abstract

Today, a device can be simultaneously connected to multiple networks, e.g. Wi-Fi, LTE, 5G, and DSL. It is desirable to combine them seamlessly below transport layer (L4) to improve quality of experience for applications that do not have such (multi-path) capability built in. Such optimization requires additional control information, e.g. Sequence Number, in each (IP) data packet. This document presents a new light weight and flexible encapsulation protocol for this need. The solution has been developed by the authors based on their experiences in multiple standards bodies including the IETF and 3GPP, is not an Internet Standard and does not represent the consensus opinion of the IETF. This document will enable other developers to build interoperable implementations.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on October 1, 2021.

Copyright Notice

Copyright (c) 2021 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Conventions used in this document	4
3.	Use Case	4
4.	GMA Encapsulation Methods	5
4.1.	Trailer-based IP Encapsulation	6
4.2.	Header-based IP Encapsulation	9
4.3.	(Header-based) non-IP Encapsulation	10
5.	Fragmentation	10
6.	Concatenation	11
7.	Security Considerations	12
8.	IANA Considerations	13
9.	References	13
9.1.	Normative References	13
9.2.	Informative References	13

[1.](#) Introduction

Today, a device can be simultaneously connected to multiple networks, e.g. Wi-Fi, LTE, 5G, and DSL. It is desirable to combine them seamlessly below transport layer (L4) to improve quality of experience for applications that do not have such (multi-path) capability built in.

Figure 1 shows the Multi-Access Management Service (MAMS) user-plane protocol stack [[MAMS](#)], which has been used in today's multi-access solutions [[ATSSS](#)] [[LWIPEP](#)] [[GRE](#)]. It consists of two layers: convergence and adaptation.

The convergence layer is responsible for multi-access operations, including multi-link (path) aggregation, splitting/reordering, lossless switching/retransmission, fragmentation, concatenation, etc. It operates on top of the adaptation layer in the protocol stack. From the Transmitter perspective, a User Payload (e.g. IP packet) is processed by the convergence layer first, and then by the adaptation layer before being transported over a delivery connection; from the Receiver perspective, an IP packet received over a delivery connection is processed by the adaptation layer first, and then by the convergence layer.

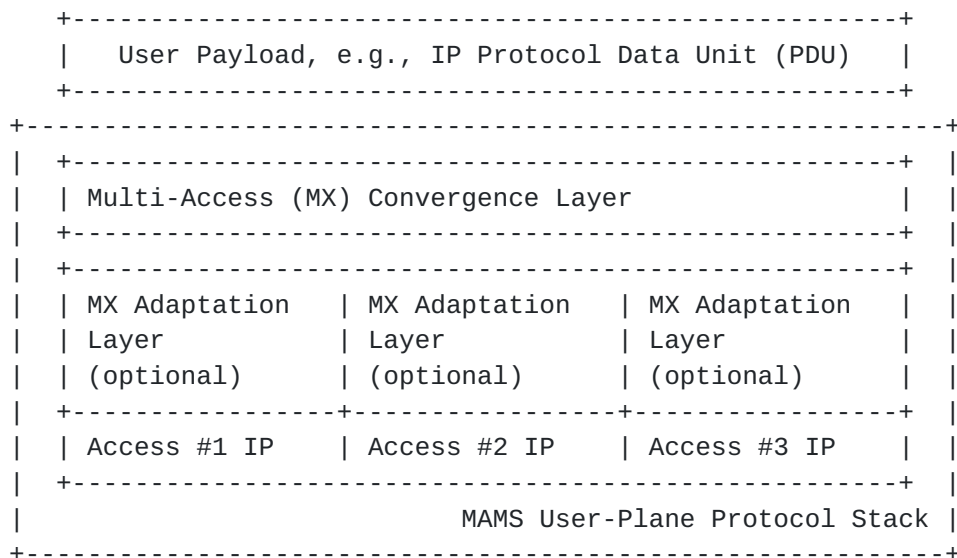


Figure 1: MAMS User-Plane Protocol Stack [[MAMS](#)]

Today, GRE is used [[LWIPPEP](#)] [[GRE1](#)] [[GRE2](#)] as the encapsulation protocol at the convergence layer to encode additional control information, e.g. Key, Sequence Number. However, there are two main drawbacks with this approach:

- o IP-over-IP tunnelling (required for GRE) leads to higher overhead especially for small packet;
- o It is difficult to introduce new control fields.

For example, the overhead of IP-over-IP/GRE tunnelling with both Key and Sequence Number is 32 Bytes (20 Bytes IP header + 12 Bytes GRE header), which is 80% of a 40 Bytes TCP ACK packet;

This document presents a light-weight GMA encapsulation protocol for the convergence layer. It supports three encapsulation

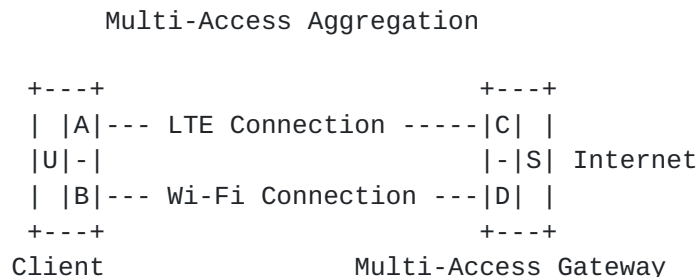
methods: trailer-based IP encapsulation, header-based IP encapsulation, and non-IP encapsulation. Particularly, the IP encapsulation methods avoid IP-over-IP tunneling overhead (20 Bytes), which is 50% of a 40 Bytes TCP ACK packet. Moreover, it introduces new control fields to support fragmentation and concatenation, which are not available in today's GRE-based solutions [[LWIPEP](#)] [GRE].

GMA protocol SHALL only operate between endpoints that have been configured to operate with GMA through additional control messages and procedures, for example [[MAMS](#)]. Moreover, UDP or IPSec tunneling MAY be used at the adaptation sublayer to protect GMA operation from intermediate nodes.

2. Conventions used in this document

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.

3. Use Case



A: The adaptation layer endpoint of the LTE connection resides in the client

B: The adaptation layer endpoint of the Wi-Fi connection resides in the client

C: The adaptation layer endpoint of the LTE connection resides in the Multi-Access Gateway, aka N-MADP (Network Multi-Access Data Proxy) in [[MAMS](#)]

D: The adaptation layer endpoint of the Wi-Fi connection resides in the Multi-Access Gateway

U: The convergence layer endpoint resides in the client

S: The convergence layer endpoint resides in the Multi-Access Gateway

Figure 2: GMA-based Multi-Access Aggregation

As shown in Figure 2, a client device (e.g. Smartphone, Laptop, etc.) may connect to Internet via both Wi-Fi and LTE connections, one of which (e.g. LTE) may operate as the anchor connection, and the other (e.g. Wi-Fi) may operate as the delivery connection. The anchor connection provides the IP address and connectivity for end-to-end Internet access, and the delivery connection provides additional path between client and Multi-Access Gateway for multi-access optimizations.

For example, per-packet aggregation allows a single IP flow to use the combined bandwidth of the two connections. In another example, packets lost due to temporarily link outage may be retransmitted. Moreover, packets may be duplicated over multiple connections to achieve high reliability and low latency, and duplicated packets should be eliminated by the receiving side. Such multi-access optimization requires additional control information, e.g. Sequence Number, in each IP data packet, which can be supported by the GMA encapsulation protocol described in this document.

The GMA protocol in this document is designed for multiple connections, but it may also be used when there is only one connection between two end-points. For example, it may be used for loss detection and recovery. In another example, it may be used to concatenate multiple small packets and reduce per packet overhead.

4. GMA Encapsulation Methods

The GMA encapsulation protocol supports the following three methods:

- o Trailer-based IP Encapsulation (4.1)
- o Header-based IP Encapsulation (4.2)
- o (Header-based) non-IP Encapsulation (4.3)

Trailer-based IP encapsulation SHOULD be used as long as implementation allows.

Header-based encapsulation SHOULD be used if trailer-based encapsulation is not feasible due to any reason, e.g. implementation constraints. In this case, if the adaptation layer, e.g. UDP tunnelling, supports non-IP packet format, header-based non-IP encapsulation SHOULD be used; otherwise, header-based IP encapsulation SHOULD be used.

If non-IP encapsulation is configured, GMA header SHOULD always be present in every packet. In comparison, if IP encapsulation is configured, GMA header or trailer may be added dynamically on per-packet basis, and it indicates the presence of GMA header (or trailer) to set the protocol type of the GMA PDU to "114".

The GMA endpoints MAY configure the GMA encapsulation method through control signalling or pre-configuration. For example, the "MX UP Setup Configuration Request" message as specified in Multi-Access Management Service [[MAMS](#)] includes "MX Convergence Method Parameters", which provides the list of parameters to configure the convergence layer, and can be extended to indicate the GMA encapsulation method.

4.1. Trailer-based IP Encapsulation




Figure 3: GMA PDU Format with Trailer-based IP Encapsulation

Figure 3 shows the trailer-based IP encapsulation GMA PDU (protocol data unit) format. A (GMA) PDU may carry one or multiple IP packets, aka (GMA) SDUs (service data unit), in the payload, or a fragment of the SDU.

The Protocol Type field in the IP header of the GMA PDU MUST be changed to 114 (Any 0-Hop Protocol) [[IANA](#)] to indicate the presence of the GMA trailer.

If the original IP packet is IPv4, the following three IP header fields SHOULD be changed:

- o IP Length: add the length of "GMA Trailer" to the length of the original IP packet
- o Time To Live (TTL): set to "1"
- o IP checksum: recalculate after changing "Protocol Type", "TTL" and "IP Length"

If the original IP packet is IPv6, the following two IP header fields SHOULD be changed:

- o IP Length: add the length of "GMA Trailer" to the length of the original IP packet
- o Hop-Limit (HL): set the HL field to "0"

The GMA (Generic Multi-Access) trailer MUST consist of two mandatory fields (the last 3 bytes): Next Header and Flags, defined as follows:

- o Next Header (1 Byte): the IP protocol type of the (first) SDU in a PDU, and it stores the value before it was overwritten to 114.
- o Flags (2 Bytes): Bit 0 is the most significant bit (MSB), and Bit 15 is the least significant bit (LSB)
 - + Checksum Present (bit 0): If the Checksum Present bit is set to 1, then the Checksum field is present;
 - + Concatenation Present (bit 1): If the Concatenation Present bit is set to 1, then the PDU carries multiple SDUs, and the First SDU Length field is present;
 - + Connection ID Present (bit 2): If the Connection ID Present bit is set to 1, then the Connection ID field is present;
 - + Flow ID Present (bit 3): If the Flow ID Present bit is set to 1, then the Flow ID field is present;
 - + Fragmentation Present (bit 4): If the Fragmentation Present bit is set to 1, then the PDU carry a fragment of the SDU and the Fragmentation Control field is present;
 - + Delivery SN Present (bit 5): If the Delivery SN (Sequence Number) Present bit is set to 1, then the Delivery SN field is present and contains the valid information;
 - + Flow SN Present (bit 6): If the Flow SN Present bit is set to 1, then the Sequence Number field is present;
 - + Timestamp Present (bit 7): If the Timestamp Present bit is set to 1, then the Timestamp field is present;
 - + TTL Present (bit 8): If the TTL Present bit is set to 1, then the TTL field is present;
 - + Reserved (bit 9-12): set to "0" and ignored on receipt;
 - + Version (bit 13-15): GMA version number, set to 0 for the GMA encapsulation protocol specified in this document.

The Flags field is at the end of the PDU, and the Next Header field is the second last. The Receiver SHOULD first decode the Flags field to determine the length of the GMA trailer, and then decode each optional field accordingly. The GMA (Generic Multi-Access) trailer MAY consist of the following optional fields:

- o Checksum (1 Byte): to contain the (one's complement) checksum sum of all the 8 bits in the trailer. For purposes of computing the checksum, the value of the checksum field is

- zero. This field is present only if the Checksum Present bit is set to one.
- o First SDU Length (2 Bytes): the length of the first IP packet in the PDU, only included if a PDU contains multiple IP packets. This field is present only if the Concatenation Present bit is set to one.
 - o Connection ID (1 Byte): an unsigned integer to identify the anchor and delivery connection of the GMA PDU. This field is present only if the Connection ID Present bit is set to one.
 - + Anchor Connection ID (MSB 4 Bits): an unsigned integer to identify the anchor connection
 - + Delivery Connection ID (LSB 4 Bits): an unsigned integer to identify the delivery connection
 - o Flow ID (1 Byte): an unsigned integer to identify the IP flow that a PDU belongs to, for example Data Radio Bearer (DRB) ID [[LWIPPEP](#)] for a cellular (e.g. LTE) connection. This field is present only if the Flow ID Present bit is set to one.
 - o Fragmentation Control (FC) (e.g. 1 Byte): to provide necessary information for re-assembly, only needed if a PDU carries fragments. This field is present only if the Fragmentation Present bit is set to one. Please refer to [section 5](#) for its detailed format and usage.
 - o Delivery SN (1 Byte): an auto-incremented integer to indicate the GMA PDU transmission order on a delivery connection. Delivery SN is needed to measure packet loss of each delivery connection and therefore generated per delivery connection per flow. This field is present only if the Delivery SN Present bit is set to one.
 - o Flow SN (3 Bytes): an auto-incremented integer to indicate the GMA SDU (IP packet) order of a flow. Flow SN is needed for retransmission, reordering, and fragmentation. It SHALL be generated per flow. This field is present only if the Flow SN Present bit is set to one.
 - o Timestamp (4 Bytes): to contain the current value of the timestamp clock of the transmitter in the unit of 1 millisecond. This field is present only if the Timestamp Present bit is set to one.
 - o TTL (1 Byte): to contain the TTL value of the original IP header if the GMA SDU is IPv4, or the Hop-Limit value of the IP header if the GMA SDU is IPv6. This field is present only if the TTL Present bit is set to one.

Figure 4 shows the GMA trailer format with all the fields present, and the order of the GMA control fields SHALL follow the bit order in the Flags field. For example, Bit 0 (the MSB) of the Flags field is the Checksum Present bit, and the Checksum field is the last in the trailer except of the two mandatory fields. Bit 1 is

the Concatenation present bit, and the FSL field is the second last.

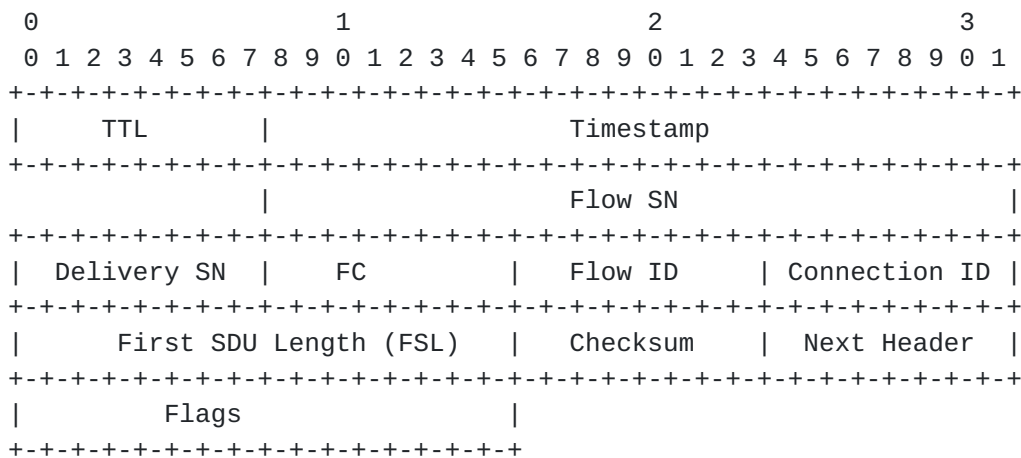


Figure 4: GMA Trailer Format

4.2. Header-based IP Encapsulation

Figure 5 shows the header-based IP encapsulation format. Here, the GMA header is inserted right after the IP header of the GMA SDU, and the IP header fields of the GMA PDU SHOULD be changed in the same way as trailer-based IP encapsulation.




Figure 5: GMA PDU Format with Header-based IP Encapsulation

Figure 6 shows the GMA header format. In comparison to GMA trailer, the only difference is that the Flags field is now in the front so that the Receiver can first decode the Flags field to determine the GMA header length. Moreover, the IP header fields of the GMA SDU SHOULD be changed the same way as in trailer-based IP encapsulation.




Figure 6: GMA Header Format

4.3. (Header-based) non-IP Encapsulation

Figure 7 shows the header-based non-IP encapsulation format. Here, "UDP Tunnelling" is configured at the MX adaptation layer. "TTL", "FSL", and "Next Header" are no longer needed. Moreover, the IP header fields of the GMA SDU remain unchanged.




Figure 7: GMA PDU Format with Non-IP Encapsulation

5. Fragmentation

The convergence layer MAY support fragmentation if a delivery connection has a smaller maximum transmission unit (MTU) than the original IP packet (SDU). The fragmentation procedure at the convergence sublayer is similar to IP fragmentation [[RFC791](#)] in principle, but with the following two differences for less overhead:

- o The fragment offset field is expressed in number of fragments
- o The maximum number of fragments per SDU is 2^7 (=128)

The Fragmentation Control (FC) field in the GMA trailer (or header) contains the following bits:

- o Bit #7: a More Fragment (MF) flag to indicate if the fragment is the last one (0) or not (1)
- o Bit #0~#6: Fragment Offset (in units of fragments) to specify the offset of a particular fragment relative to the beginning of the SDU

A PDU carries a whole SDU without fragmentation if the FC field is set to all "0"s or the FC field is not present in the trailer. Otherwise, the PDU contains a fragment of the SDU.

The Flow SN field in the trailer is used to distinguish the fragments of one SDU from those of another. The Fragment Offset (FO) field tells the receiver the position of a fragment in the original SDU. The More Fragment (MF) flag indicates the last fragment.

To fragment a long SDU, the transmitter creates n PDUs and copies the content of the IP header fields from the long PDU into the IP header of all the PDUs. The length field in the IP header of PDU SHOULD be changed to the length of the PDU, and the protocol type SHOULD be changed to 114.

The data of the long SDU is divided into n portions based on the MTU size of the delivery connection. The first portion of the data is placed in the first PDU. The MF flag is set to "1", and the FO field is set to "0". The i -th portion of the data is placed in the i -th PDU. The MF flag is set to "0" if it is the last fragment and set to "1" otherwise. The FO field is set to $i-1$.

To assemble the fragments of a SDU, the receiver combines PDUs that all have the same Flow SN. The combination is done by placing the data portion of each fragment in the relative order indicated by the Fragment Offset in that fragment's GMA trailer (or header). The first fragment will have the Fragment Offset set to "0", and the last fragment will have the More-Fragments flag set to "0".

GMA fragmentation operates above the IP layer of individual access connection (Wi-Fi, LTE) and between the two end points of convergence layer. The convergence layer end points (client, multi-access gateway) SHOULD obtain the MTU of individual connection through either manual configuration or implementing PMTUD as suggested in [[RFC8900](#)].

6. Concatenation

The convergence sublayer MAY support concatenation if a delivery connection has a larger maximum transmission unit (MTU) than the original IP packet (SDU). Only the SDUs with the same client IP address, and the same Flow ID MAY be concatenated.

If the (trailer or header based) IP encapsulation method is used, the First SDU Length (FSL) field SHOULD be included in the GMA trailer (or header) to indicate the length of the first SDU. Otherwise, the FSL field SHOULD not be included.

```
+-----+
|IP hdr| IP payload   |IP hdr|   IP payload | GMA Trailer |
+-----+
```

Figure 8: Example of GMA PDU Format with Concatenation and Trailer-based IP Encapsulation

To concatenate two or more SDUs, the transmitter creates one PDU and copies the content of the IP header field from the first SDU

into the IP header of the PDU. The data of the first SDU is placed in the first portion of the data of the PDU. The whole second SDU is then placed in the second portion of the data of the PDU (Figure 8). The procedure continues till the PDU size reaches the MTU of the delivery connection. If the FSL field is present, the IP length field of the PDU SHOULD be updated to include all concatenated SDUs and the trailer (or header), and the IP checksum field SHOULD be recalculated if the packet is IPv4.

To disaggregate a PDU, if the (header or trailer based) IP encapsulation method is used, the receiver first obtains the length of the first SDU from the FSL field and decodes the first SDU. The receiver then obtains the length of the second SDU based on the length field in the second SDU IP header and decodes the second SDU. The procedure continues till no byte is left in the PDU. If the non-IP encapsulation method (Figure 7) is used, the IP header of the first SDU will not change during the encapsulation process, and the receiver SHOULD obtain the length of the first SDU directly from its IP header (Figure 9).




Figure 9: Example of GMA PDU Format with Concatenation and Header-based Non-IP (UDP) Encapsulation

If a PDU contains multiple SDUs, the Flow SN field is for the last SDU, and the Flow SN of other SDU carried by the same PDU can be obtained according to its order in the PDU. For example, if the SN field is 6 and a PDU contains 3 SDUs (IP packets), the SN is 4, 5, and 6 for the first, second, and last SDU respectively.

GMA concatenation can be used for packing small packets of a single application, e.g. TCP ACKs, or from multiple applications. Notice that a single GMA flow may carry multiple application flows (TCP, UDP, etc.).

7. Security Considerations

Security in a network using GMA should be relatively similar to security in a normal IP network. The GMA protocol at the convergence sublayer is a 0-hop protocol and relies on the

security of the underlying network transport paths. When this cannot be assumed, appropriate security protocols (IPsec, DTLS, etc.) SHOULD be configured at the adaptation sublayer. On the other hand, packet filtering requires either that a firewall looks inside the GMA packet or that the filtering is done on the GMA endpoints. In those environments in which this is considered to be a security issue it may be desirable to terminate the GMA operation at the firewall.

Local-only packet leak prevention (HL=0, TTL=1) SHOULD be on preventing the leak of the local-only GMA PDUs outside the "local domain" to the Internet or to another domain which could use the same IP protocol type, i.e. 114.

8. IANA Considerations

This draft makes no requests of IANA

9. References

9.1. Normative References

- [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>>.
- [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>>.
- [GRE1] Dommety, G., "Key and Sequence Number Extensions to GRE", <<https://www.rfc-editor.org/info/rfc2890>> .

9.2. Informative References

- [MAMS] S. Kanugovi, F. Baboescu, J. Zhu, and S. Seo "Multi-Access Management Services
(MAMS)<https://tools.ietf.org/rfc/rfc8743.txt>
- [IANA] <https://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml>
- [LWIPEP] 3GPP TS 36.361, "Evolved Universal Terrestrial Radio Access (E-UTRA); LTE-WLAN Radio Level Integration Using Ipsec Tunnel (LWIP) encapsulation; Protocol specification"

[RFC791] Internet Protocol, September 1981

[ATSSS] 3GPP TR 23.793, Study on access traffic steering, switch and splitting support in the 5G system architecture.

[GRE2] [RFC 8157](#), Huawei's GRE Tunnel Bonding Protocol, May 2017

[RFC8900] Bonica, R., Baker, F., Huston, G., Hinden, R., Troan, O., and F. Gont, "IP Fragmentation Considered Fragile", [BCP 230](#), [RFC 8900](#), DOI 10.17487/RFC8900, September 2020, <<https://www.rfc-editor.org/info/rfc8900>>.

Authors' Addresses

Jing Zhu

Intel

Email: jing.z.zhu@intel.com

Satish Kanugovi

Nokia

Email: satish.k@nokia.com