

INTAREA
Internet Draft
Intended status: Standards Track
Expires: April 19, 2019

J. Zhu
Intel
S. Seo
Korea Telecom
S. Kanugovi
Nokia
S. Peng
Huawei
October 19, 2018

User-Plane Protocols for Multiple Access Management Service
draft-zhu-intarea-mams-user-protocol-06

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/lid-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 April 19, 2009.

Copyright Notice

Copyright (c) 2018 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.

Abstract

Today, a device can be simultaneously connected to multiple communication networks based on different technology implementations and network architectures like WiFi, LTE, and DSL. In such multi-connectivity scenario, it is desirable to combine multiple access networks or select the best one to improve quality of experience for a user and improve overall network utilization and efficiency. This document presents the u-plane protocols for a multi access management services (MAMS) framework that can be used to flexibly select the combination of uplink and downlink access and core network paths having the optimal performance, and user plane treatment for improving network utilization and efficiency and enhanced quality of experience for user applications.

Table of Contents

1. Introduction.....	3
2. Terminologies.....	3
3. Conventions used in this document.....	3
4 MAMS User-Plane Protocols.....	4
4.1 MX Adaptation Sublayer.....	4
4.2 Trailer-based MX Convergence Sublayer.....	5
4.2.1 Trailer-based MX PDU Format.....	5
4.2.2 MX Fragmentation.....	8
4.2.3 MX Concatenation.....	9
4.3 MPTCP-based MX Convergence Sublayer.....	10
4.4 GRE as MX Convergence Sublayer.....	11
4.4.1 Transmitter Procedures.....	11
4.4.2 Receiver Procedures.....	12
4.5 Co-existence of MX Adaptation and MX Convergence Sublayers	12
5. MX Convergence Control Message.....	12
5.1 Keep-Alive Message.....	13
5.2 Probe REQ/ACK Message.....	14
5.3 Packet Loss Report (PLR) Message.....	15
5.4 First Sequence Number (FSN) Message.....	15
5.5 Coded MX SDU (CMS) Message.....	16
5.6 Traffic Splitting Update (TSU) Message.....	17
5.7 Traffic Splitting Acknowledgement (TSA) Message.....	18
6 Security Considerations.....	18
7 IANA Considerations.....	19
8 Contributing Authors.....	19
9 References.....	19

9.1	Normative References.....	19
9.2	Informative References.....	19

1. Introduction

Multi Access Management Service (MAMS) [MAMS] is a programmable framework to select and configure network paths, as well as adapt to dynamic network conditions, when multiple network connections can serve a client device. It is based on principles of user plane interworking that enables the solution to be deployed as an overlay without impacting the underlying networks.

This document presents the u-plane protocols for enabling the MAMS framework. It co-exists and complements the existing protocols by providing a way to negotiate and configure the protocols based on client and network capabilities. Further it allows exchange of network state information and leveraging network intelligence to optimize the performance of such protocols. An important goal for MAMS is to ensure that there is minimal or no dependency on the actual access technology of the participating links. This allows the scheme to be scalable for addition of newer access technologies and for independent evolution of the existing access technologies.

2. Terminologies

Anchor Connection: refers to the network path from the N-MADP to the Application Server that corresponds to a specific IP anchor that has assigned an IP address to the client.

Delivery Connection: refers to the network path from the N-MADP to the C-MADP.

"Network Connection Manager" (NCM), "Client Connection Manager" (CCM), "Network Multi Access Data Proxy" (N-MADP), and "Client Multi Access Data Proxy" (C-MADP) in this document are to be interpreted as described in [MAMS].

3. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The terminologies "Network Connection Manager" (NCM), "Client Connection Manager" (CCM), "Network Multi Access Data Proxy" (N-MADP), and "Client Multi Access Data Proxy" (C-MADP) in this document are to be interpreted as described in [MAMS].

4 MAMS User-Plane Protocols

Figure 1 shows the MAMS u-plane protocol stack as specified in [MAMS_CP].




Figure 1: MAMS U-plane Protocol Stack

It consists of the following two Sublayers:

- o Multi-Access (MX) Convergence Sublayer: This layer performs multi-access specific tasks, e.g., access (path) selection, multi-link (path) aggregation, splitting/reordering, lossless switching, fragmentation, concatenation, keep-alive, and probing etc.
- o Multi-Access (MX) Adaptation Sublayer: This layer performs functions to handle tunneling, network layer security, and NAT.

The MX convergence sublayer operates on top of the MX adaptation sublayer in the protocol stack. From the Transmitter perspective, a User Payload (e.g. IP PDU) is processed by the convergence sublayer first, and then by the adaptation sublayer before being transported over a delivery access connection; from the Receiver perspective, an IP packet received over a delivery connection is processed by the MX adaptation sublayer first, and then by the MX convergence sublayer.

4.1 MX Adaptation Sublayer

The MX adaptation sublayer supports the following mechanisms and protocols while transmitting user plane packets on the network path:

- o UDP Tunneling: The user plane packets of the anchor connection can be encapsulated in a UDP tunnel of a delivery connection between the N-MADP and C-MADP.
- o IPsec Tunneling: The user plane packets of the anchor connection are sent through an IPsec tunnel of a delivery connection.
- o Client Net Address Translation (NAT): The Client IP address of user plane packet of the anchor connection is changed, and sent over a delivery connection.
- o Pass Through: The user plane packets are passing through without any change over the anchor connection.

The MX adaptation sublayer also supports the following mechanisms and protocols to ensure security of user plane packets over the network path.

- o IPsec Tunneling: An IPsec [RFC7296] tunnel is established between the N-MADP and C-MADP on the network path that is considered untrusted.
- o DTLS: If UDP tunneling is used on the network path that is considered "untrusted", DTLS (Datagram Transport Layer Security) [RFC6347] can be used.

The Client NAT method is the most efficient due to no tunneling overhead. It SHOULD be used if a delivery connection is "trusted" and without NAT function on the path.

The UDP or IPsec Tunnelling method SHOULD be used if a delivery connection has a NAT function placed on the path.

4.2 Trailer-based MX Convergence Sublayer

4.2.1 Trailer-based MX PDU Format

Trailer-based MX convergence integrates multiple connections into a single e2e IP connection. It operates between Layer 2 (L2) and Layer 3 (network/IP).




Figure 2 shows the trailer-based Multi-Access (MX) PDU (Protocol Data Unit) format. A MX PDU MAY carry multiple IP PDUs in the payload if concatenation is supported, and MAY carry a fragment of the IP PDU if fragmentation is supported.

The MX trailer may consist of the following fields:

- o MX flags (e.g. 1 Byte): Bit 0 is the most significant bit, bit 7 is the least significant bit. Bit 6 and 7 are reserved for future.
 - + Next Header Present (bit 0): If the Next Header Present bit is set to 1, then the Next Header field is present and contains valid information.
 - + Connection ID Present (bit 1): If the Connection ID Present bit is set to 1, then the Connection ID field is present and contains valid information.
 - + Traffic Class Present (bit 2): If the Traffic Class Present bit is set to 1, then the Traffic Class field is present and contains valid information.
 - + Sequence Number Present (bit 3): If the Sequence Number Present bit is set to 1, then the Sequence Number field is present and contains valid information.
 - + Packet Length Present (bit 4): If the Packet Length Present bit is set to 1, then the First SDU (Service Data Unit) Length field is present and contains valid information.
 - + Fragmentation Control Present (bit 5): If the Fragmentation Control Present bit is set to 1, then the Fragmentation Control field is present and contains valid information.
 - + Traffic Splitting Flag (bit 6): The bit will be flipped (0 or 1) when the traffic splitting configuration changes
 - + Bit 7: reserved
- o Next Header (e.g. 1 Byte): the IP protocol type of the (first) IP packet in a MX PDU
- o Connection ID (e.g.1 Byte): an unsigned integer to identify the anchor connection of the IP packets in a MX PDU
- o Traffic Class (TC) ID (e.g. 1 Byte): an unsigned integer to identify the traffic class of the IP packets in a MX PDU, for example Data Radio Bearer (DRB) ID [LWIPPEP] for a cellular (e.g. LTE) connection
- o Sequence Number (e.g. 2 Bytes): an auto-incremented integer to indicate order of transmission of the MX SDU (e.g. IP packet), needed for lossless switching or multi-link (path) aggregation or fragmentation. Sequence Number SHALL be generated on a per Connection and per Traffic Class (TC) basis.
- o First SDU Length (e.g. 2 Bytes): the length of the first IP packet, only included if a MX PDU contains multiple IP packets, i.e. concatenation.
- o Fragmentation Control (FC) (e.g. 1 Byte): to provide necessary information for re-assembly, only needed if a MX PDU carries fragments, i.e. fragmentation.

Figure 3 shows the MX trailer format with all the fields present. The MX flags are always encoded in the last octet of the MX Trailer at the end of a MX PDU. Hence, the Receiver SHOULD first decode the MX flags

field to determine the length of the MX trailer, and then decode each MX field accordingly.




Figure 3: MX Trailer Format

Moreover, the following field of the IP header of the MX PDU SHOULD be changed:

- o Protocol Type: "114" to indicate that the presence of MX trailer (i.e. the trailer based MAMS u-plane protocol is a "0-hop" protocol, not subject to IP routing)

If the MX PDU is transported with the MX adaptation method of IPSec tunnelling, Client NAT, or Pass Through, the following fields of the IP header of the MX PDU SHOULD also be changed:

- o IP length: add the length of "MX Trailer" to the length of the original IP packet
- o IP checksum: recalculate after changing "Protocol Type" and "IP Length"

If the MX adaptation method is UDP tunnelling and "MX header optimization" in the "MX_UP_Setup_Configuration_Request" message [MAMS] is true, the "IP length" and "IP checksum" header fields of the MX PDU SHOULD remain unchanged.

The MX u-plane protocol can support multiple Anchor connections simultaneously, each of which is uniquely identified by Connection ID. It can also support multiple traffic classes per connection, each of which is identified by Traffic Class ID.

Moreover, the MX trailer format MAY be negotiated dynamically between NCM and CCM. For example, NCM can send a control message to indicate which of the above fields SHOULD be included for individual delivery connection, on downlink and uplink, respectively.

4.2.2 MX Fragmentation

The Trailer-based MX Convergence Layer SHOULD support MX fragmentation if a delivery connection has a smaller maximum transmission unit (MTU) than the original IP packet (MX SDU), and IP fragmentation is not supported or enabled on the connection. The MX fragmentation procedure 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 not 8-bytes blocks
- o The maximum number of fragments per MX SDU is 2^7 (=128)

The Fragmentation Control (FC) field in the MX Trailer 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 MX SDU

A MX PDU carries a whole MX 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 MX PDU contains a fragment of the MX SDU.

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

To fragment a long MX SDU, the MADP transmitter creates two MX PDUs and copies the content of the IP header fields from the long MX PDU into the IP header of both MX PDUs. The length field in the IP header of MX PDU SHOULD be changed to the length of the MX PDU, and the protocol type SHOULD be changed to "114", indicating the presence of the MX trailer.

The data of the long MX SDU is divided into two portions based on the MTU size of the delivery connection. The first portion of the data is placed in the first MX PDU. The MF flag is set to "1", and the FO field is set to "0". The second portion of the data is placed in the second

MX PDU. The MF flag is set to "0", and the FO field is set to "1". This procedure can be generalized for an n-way split, rather than the two-way split described the above.

To assemble the fragments of a MX SDU, the MADP receiver combines MX PDUs that all have the same MX Sequence Number (in the trailer). 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 MX trailer. The first fragment will have the Fragment Offset set to "0", and the last fragment will have the More-Fragments flag reset to "0".

4.2.3 MX Concatenation

The Trailer-based MX Convergence Layer MAY support MX concatenation if a delivery connection has a larger maximum transmission unit (MTU) than the original IP packet (MX SDU). Only the MX SDUs with the same client address, the same anchor connection and the same Traffic Class MAY be concatenated.

If the MX adaptation method is IPSec tunnelling, Client NAT, or Pass Through, The First SDU Length (FSL) field SHOULD be included in the MX Trailer to indicate the length of the first MX SDU.

If the MX adaptation method is UDP tunneling and "MX header optimization" in the "MX_UP_Setup_Configuration_Request" message [MAMS] is true, the FSL field SHOULD not be present, or the entire MX trailer MAY not be present. The MADP receiver compares the IP length field of the MX PDU and the actual length of the MX PDU to determine if the MX PDU contains multiple MX SDUs. If the MX PDU is larger than what the IP length field indicates, the MX PDU contains multiple MX SDUs; otherwise, the MX PDU contains only one MX SDU. To concatenate two or more MX SDUs, the MADP transmitter creates one MX PDU and copies the content of the IP header field from the first MX SDU into the IP header of the MX PDU. The data of the first MX SDU is placed in the first portion of the data of the MX PDU. The whole second MX SDU is then placed in the second portion of the data of the MX PDU (Figure 4). The procedure continues till the MX PDU size reaches the MTU of the delivery connection. If the FSL field is present in the MX Trailer, the IP length field of the MX PDU SHOULD be updated to include all concatenated SDUs and the trailer, and the IP checksum field SHOULD be recalculated.

To disaggregate a MX PDU, the MADP receiver first obtains the length of the first MX SDU from the FSL field in the trailer, and decodes the first MX SDU. If the FSL field or the MX Trailer is not present, the MADP receiver obtains the length of the first MX SDU directly from the IP length field of the MX PDU. The MADP receiver then obtains the length of the second MX SDU based on the length field in the second MX

SDU IP header, and decodes the second MX SDU. The procedure continues till no byte is left in the MX PDU.

If a MX PDU contains multiple SDUs, the SN field in the MX trailer is for the last MX SDU, and the 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 MX PDU contains 3 SDUs (IP packets), then the SN is 4, 5, and 6 for the first, second, and the last IP packet in the PDU, respectively.




Figure 4: MX PDU Format with Concatenation

4.3 MPTCP-based MX Convergence Sublayer

Figure 5 shows the MAMS u-plane protocol stack based on MPTCP. Here, MPTCP is reused as the "MX Convergence Sublayer" protocol. Multiple access networks are combined into a single MPTCP connection. Hence, no new u-plane protocol or PDU format is needed in this case.




Figure 5: MAMS U-plane Protocol Stack with MPTCP as MX Convergence Layer

If NCM determines that N-MADP is to be instantiated with MPTCP as the MX Convergence Protocol, it exchanges the support of MPTCP capability in the discovery and capability exchange procedures [MAMS_CP]. MPTCP proxy protocols [MPPProxy][MPPlain] SHOULD be used to manage traffic steering and aggregation over multiple delivery connections.

4.4 GRE as MX Convergence Sublayer

Figure 6 shows the MAMS u-plane protocol stack based on GRE (Generic Routing Encapsulation) [GRE2784]. Here, GRE is reused as the "MX Convergence sub-layer" protocol. Multiple access networks are combined into a single GRE connection. Hence, no new u-plane protocol or PDU format is needed in this case.




Figure 6: MAMS U-plane Protocol Stack with GRE as MX Convergence Layer

If NCM determines that N-MADP is to be instantiated with GRE as the MX Convergence Protocol, it exchanges the support of GRE capability in the discovery and capability exchange procedures [MAMS_CP].

4.4.1 Transmitter Procedures

Transmitter is the N-MADP or C-MADP instance, instantiated with GRE as the convergence protocol that transmits the GRE packets. The Transmitter receives the User Payload (e.g. IP PDU), encapsulates it with a GRE header and Delivery Protocol (e.g. IP) header to generate the GRE Convergence PDU.

When IP is used as the GRE delivery protocol, the IP header information (e.g. IP address) can be created using the IP header of the user payload or a virtual IP address. The "Protocol Type" field of the delivery header is set to 47 (or 0X2F, i.e. GRE) [IANA].

The GRE header fields are set as specified below,

- If the transmitter is a C-MADP instance, then sets the LSB 16 bits to the value of Connection ID for the Anchor Connection associated with the user payload or sets to 0xFFFF if no Anchor Connection ID needs to be specified.

- All other fields in the GRE header including the remaining bits in the key fields are set as per [GRE_2784][GRE_2890].

4.4.2 Receiver Procedures

Receiver is the N-MADP or C-MADP instance, instantiated with GRE as the convergence protocol that receives the GRE packets. The receiver processes the received packets per the GRE procedures [GRE_2784, GRE_2890] and retrieves the GRE header.

- If the Receiver is an N-MADP instance,
 - o Unless the LSB 16 Bits of the Key field are 0xFFFF, they are interpreted as the Connection ID of Anchor Connection for the user payload. This is used to identify the network path over which the User Payload (GRE Payload) is to be transmitted.
- All other fields in the GRE header, including the remaining bits in the Key fields, are processed as per [GRE_2784][GRE_2890].

The GRE Convergence PDU is passed onto the MX Adaptation Layer (if present) before delivery over one of the network paths.

4.5 Co-existence of MX Adaptation and MX Convergence Sublayers

MAMS u-plane protocols support multiple combinations and instances of user plane protocols to be used in the MX Adaptation and the Convergence sublayers.

For example, one instance of the MX Convergence Layer can be MPTCP Proxy [MPProxy][MPPlain] and another instance can be Trailer-based. The MX Adaptation for each can be either UDP tunnel or IPsec. IPsec may be set up for network paths considered as untrusted by the operator, to protect the TCP subflow between client and MPTCP proxy traversing that network path.

Each of the instances of MAMS user plane, i.e. combination of MX Convergence and MX Adaptation layer protocols, can coexist simultaneously and independently handle different traffic types.

5. MX Convergence Control Message

A UDP connection may be configured between C-MADP and N-MADP to exchange control messages for keep-alive or path quality estimation. The N-MADP end-point IP address and UDP port number of the UDP connection is used to identify MX control PDU. Figure 7 shows the MX control PDU format with the following fields:

- o Type (1 Byte): the type of the MX control message
 - + 0: Keep-Alive

- + 1: Probe REQ/ACK
- + 2: Packet Loss Report (PLR)
- + 3: First Sequence Number (FSN)
- + 4: Coded MX SDU (CMS)
- + 5: Traffic Splitting Update (TSU)
- + 6: Traffic Splitting Acknowledgement (TSA)
- + Others: reserved
- o CID (1 Byte): the connection ID of the delivery connection for sending out the MX control message
- o MX Control Message (variable): the payload of the MX control message

Figure 8 shows the MX convergence control protocol stack, and MX control PDU goes through the MX adaptation sublayer the same way as MX data PDU.




Figure 7: MX Control PDU Format




Figure 8: MX Convergence Control Protocol Stack

5.1 Keep-Alive Message

The "Type" field is set to "0" for Keep-Alive messages. C-MADP may send out Keep-Alive message periodically over one or multiple delivery connections, especially if UDP tunneling is used as the adaptation method for the delivery connection with a NAT function on the path.

A Keep-Alive message is 2 Bytes long, and consists of the following fields:

- o Keep-Alive Sequence Number (2 Bytes): the sequence number of the keep-alive message

5.2 Probe REQ/ACK Message

The "Type" field is set to "1" for Probe REQ/ACK messages.

N-MADP may send out the Probe REQ message for path quality estimation. In response, C-MADP may send back the Probe ACK message.

A Probe REQ message consists of the following fields:

- o Probing Sequence Number (2 Bytes): the sequence number of the Probe REQ message
- o Probing Flag (1 Byte):
 - + Bit #0: a Probe ACK flag to indicate if the Probe ACK message is expected (1) or not (0);
 - + Bit #1: a Probe Type flag to indicate if the Probe REQ/ACK message is sent during the initialization phase (0) when the network path is not included for transmission of user data or the active phase (1) when the network path is included for transmission of user data;
 - + Bit #2: a bit flag to indicate the presence of the Reverse Connection ID (R-CID) field.
 - + Bit #3~7: reserved
- o Reverse Connection ID (1 Byte): the connection ID of the delivery connection for sending out the Probe ACK message on the reverse path
- o Padding (variable)

The "R-CID" field is only present if both Bit #0 and Bit #2 of the "Probing Flag" field are set to "1". Moreover, Bit #2 of the "Probing Flag" field SHOULD be set to "0" if the Bit #0 is "0", indicating the Probe ACK message is not expected.

If the "R-CID" field is not present but the Bit #0 of the "Probing Flag" field is set to "1", the Probe ACK message SHOULD be sent over the same delivery connection as the Probe REQ message.

The "Padding" field is used to control the length of Probe REQ message.

C-MADP SHOULD send out the Probe ACK message in response to a Probe REQ message with the Probe ACK flag set to "1".

A Probe ACK message is 2 Bytes long, and consists of the following fields:

- o Probing Acknowledgement Number (2 Bytes): the sequence number of the corresponding Probe REQ message

5.3 Packet Loss Report (PLR) Message

The "Type" field is set to "2" for PLR messages.

C-MADP may send out the PLR messages to report lost MX SDU for example during handover. In response, C-MADP may retransmit the lost MX SDU accordingly.

A PLR message consists of the following fields:

- o Connection ID (1 Byte): an unsigned integer to identify the anchor connection which the ACK message is for;
- o Traffic Class ID (1 Byte): an unsigned integer to identify the traffic class of the anchor connection which the ACK message is for;
- o ACK number (2 Bytes): the next (in-order) sequence number (SN) that the sender of the PLR message is expecting
- o Number of Loss Bursts (1 Byte)
 - For each loss burst, include the following
 - + Sequence Number of the first lost MX SDU in a burst (2 Bytes)
 - + Number of consecutive lost MX SDUs in the burst (1 Byte)




Figure 9: MAMS Retransmission Procedure

Figure 9 shows the MAMS retransmission procedure in an example where the lost packet is found and retransmitted.

5.4 First Sequence Number (FSN) Message

The "Type" field is set to "3" for FSN messages.

N-MADP may send out the FSN messages to indicate the oldest MX SDU in its buffer if a lost MX SDU is not found in the buffer after receiving

the PLR message from C-MADP. In response, C-MADP SHALL only report packet loss with SN not smaller than FSN.

A FSN message consists of the following fields:

- o Connection ID (1 Byte): an unsigned integer to identify the anchor connection which the FSN message is for;
- o Traffic Class ID (1 Byte): an unsigned integer to identify the traffic class of the anchor connection which the FSN message is for;
- o First Sequence Number (2 Bytes): the sequence number (SN) of the oldest MX SDU in the (retransmission) buffer of the sender of the FSN message.

Figure 10 shows the MAMS retransmission procedure in an example where the lost packet is not found.




Figure 10: MAMS Retransmission Procedure with FSN

5.5 Coded MX SDU (CMS) Message

The "Type" field is set to "4" for CMS messages.

N-MADP (or C-MADP) may send out the CMS message to support downlink (or uplink) packet loss recovery through coding, e.g. [CRLNC], [CTCP], [RLNC]. A coded MX SDU is generated by applying a network coding algorithm to multiple consecutive (uncoded) MX SDUs, and it is used for fast recovery without retransmission if any of the MX SDUs is lost.

A Coded MX SDU message consists of the following fields:

- o Connection ID (1 Byte): an unsigned integer to identify the anchor connection of the coded MX SDU;
- o Traffic Class ID (1 Byte): an unsigned integer to identify the traffic class of the coded MX;
- o Sequence Number (2 Bytes): the sequence number of the first (uncoded) MX SDU used to generate the coded MX SDU.
- o Fragmentation Control (FC) (1 Byte): to provide necessary information for re-assembly, only needed if the coded MX SDU is too long to transport in a single MX control PDU.
- o N (1 Byte): the number of consecutive MX SDUs used to generate the coded MX SDU
- o K (1 Byte): the length (in terms of bits) of the coding coefficient field
- o Coding Coefficient (N x K / 8 Bytes)
 - + a(i): the coding coefficient of the i-th (uncoded) MX SDU
 - + padding
- o Coded MX SDU (variable): the coded MX SDU

If N = 2 and K = 0, the simple XOR method is used to generate the Coded MX SDU from two consecutive uncoded MX SDUs, and the a(i) fields are not included in the message.

If the coded MX SDU is too long, it can be fragmented, and transported by multiple MX control PDUs. The N, K, and a(i) fields are only included in the MX PDU carrying the first fragment of the coded MX SDU.




Figure 11: MAMS Packet Recovery Procedure with XOR Coding

5.6 Traffic Splitting Update (TSU) Message

The "Type" field is set to "5" for TSU messages.

N-MADP (or C-MADP) may send out a TSU message if downlink (or uplink) traffic splitting configuration has changed.

A TSU message consists of the following fields:

- o Connection ID (1 Byte): an unsigned integer to identify the anchor connection;
- o Traffic Class ID (1 Byte): an unsigned integer to identify the traffic class;
- o Sequence Number (2 Bytes): an unsigned integer to identify the TSU message.
- o Traffic Splitting Configuration Parameters (3 + (N -1) Bytes):
 - + StartSN (2 Bytes): the sequence number of the first MX SDU using the traffic splitting configuration provided by the TSU message
 - + L (1 Byte): the traffic splitting burst size
 - + K(i): the traffic splitting threshold of the i-th delivery connection, where connections are ordered according to their Connection ID.

Let's use $f(x)$ to denote the traffic splitting function, which maps a MX SDU Sequence Number "x" to the i-th delivery connection.

$$f(x)=i, \quad \text{if } K[i-1] < \text{or } = \text{mod}(x - \text{StartSN}, L) < K[i]$$

Wherein, $1 < \text{or } = i < N$, $K[0]=0$, and $K[N]=L$.

N is the total number of connections for delivering a data flow, identified by (anchor) Connection ID and Traffic Class ID.

5.7 Traffic Splitting Acknowledgement (TSA) Message

The "Type" field is set to "6" for TSA messages.

C-MADP (or N-MADP) SHOULD send out the TSA message in response to the successful reception of a TSU message.

A TSU message consists of the following fields:

- o Connection ID (1 Byte): an unsigned integer to identify the anchor connection;
- o Traffic Class ID (1 Byte): an unsigned integer to identify the traffic class;
- o Acknowledgment Number (2 Bytes): the sequence number of the received TSU message.

6 Security Considerations

User data in MAMS framework rely on the security of the underlying network transport paths. When this cannot be assumed, NCM configures use of appropriate protocols for security, e.g. IPsec [RFC4301] [RFC3948], DTLS [RFC6347].

7 IANA Considerations

TBD

8 Contributing Authors

The editors gratefully acknowledge the following additional contributors in alphabetical order: Salil Agarwal/Nokia, Hema Pentakota/Nokia.

9 References

9.1 Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI10.17487/RFC4301, December 2005, <<http://www.rfc-editor.org/info/rfc4301>>.

9.2 Informative References

- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012, <<http://www.rfc-editor.org/info/rfc6347>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<http://www.rfc-editor.org/info/rfc7296>>.
- [RFC3948] Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and M. Stenberg, "UDP Encapsulation of IPsec ESP Packets", RFC 3948, DOI 10.17487/RFC3948, January 2005, <<http://www.rfc-editor.org/info/rfc3948>>.
- [MPProxy] X. Wei, C. Xiong, and E. Lopez, "MPTCP proxy mechanisms", <https://tools.ietf.org/html/draft-wei-mptcp-proxy-mechanism-02>
- [MPPlain] M. Boucadair et al, "An MPTCP Option for Network-Assisted MPTCP", <https://www.ietf.org/id/draft-boucadair-mptcp-plain-mode-09.txt>

- [MAMS] S. Kanugovi, S. Vasudevan, F. Baboescu, and J. Zhu, "Multiple Access Management Protocol",
<https://tools.ietf.org/html/draft-kanugovi-intarea-mams-protocol-03>
- [MAMS_CP] S. Kanugovi, et al., "Control Plane Protocols and Procedures for Multiple Access Management Services"
- [GRE2784] D. Farinacci, et al., "Generic Routing Encapsulation (GRE)", RFC 2784 March 2000, <<http://www.rfc-editor.org/info/rfc2784>>.
- [GRE2890] G. Dommety, "Key and Sequence Number Extensions to GRE", RFC 2890 September 2000, <<http://www.rfc-editor.org/info/rfc2890>>.
- [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 Ipv6 Tunnel (LWIP) encapsulation; Protocol specification"
- [RFC791] Internet Protocol, September 1981
- [CRLNC] S Wunderlich, F Gabriel, S Pandi, et al. Caterpillar RLNC (CRLNC): A Practical Finite Sliding Window RLNC Approach, IEEE Access, 2017
- [CTCP] M. Kim, et al. Network Coded TCP (CTCP), eprint arXiv:1212.2291, 2012
- [RLNC] J. Heide, et al. Random Linear Network Coding (RLNC)-Based Symbol Representation, <https://www.ietf.org/id/draft-heide-nwcrp-rlnc-00.txt>

Authors' Addresses

Jing Zhu

Intel

Email: jing.z.zhu@intel.com

SungHoon Seo

Korea Telecom

Email: sh.seo@kt.com

Satish Kanugovi

Nokia

Email: satish.k@nokia.com

Shuping Peng

Huawei

Email: pengshuping@huawei.com

