# Troubleshoot Packet Drops Related to SCTP Max MTU Size

## Contents



## Introduction

This document describes SCTP fragmentation and chunk bundling mechanisms in the Cisco MME and how fragmentation and bundling affect packet drops.

## Prerequisites

### Requirements

There are no specific requirements for this document.

### Components Used

The information in this document is based on these software and hardware versions:

- QVPC-SI software version 21.28.m18

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

# Background Information

Stream Control Transmission Protocol (SCTP) is a reliable transport protocol operating on top of a connectionless packet network such as IP.  It can fragment data to conform to discovered path MTU size and bundle multiple user messages into a single SCTP packet by using chunk bundling. This document provides two methods to remedy packet drops that occur due to path MTU limitations.

A case study displays fragmentation and bundling mechanisms and a means to demonstrate that packet drops, including packet drops that occur outside of the Cisco MME, rather than due to the configuration of SCTP Max MTU size.

Per RFC 4960 (Stream Control Transmission Protocol) 6.9 (Fragmentation and Reassembly):

*"If an endpoint supports fragmentation, it MUST fragment a user message if the size of the user message to be sent causes the outbound SCTP packet size to exceed the current MTU."*

In addition, multiple SCTP data chunks can be bundled into one outbound SCTP packet. The total size of the resultant IP datagram, including the SCTP packet and IP headers, MUST be less that or equal to the current Path MTU. On the MME the Path MTU is configured on the IP interface that is processing these packets.

## Fragmentation

There can be multiple streams in a SCTP association, each identified by a Stream Identifier (SI). Stream Sequence Number (SSN) identifies the order for each fragment within a specific SCTP stream. Multiple sessions can use the same SCTP stream. The Transmission Sequence Number (TSN) identifies the order for each fragment within the entire SCTP association.

When fragmented, the SCTP Payload data chunk indicates in the B and E fields whether the fragment is the beginning fragment, a middle fragment, or the end fragment:

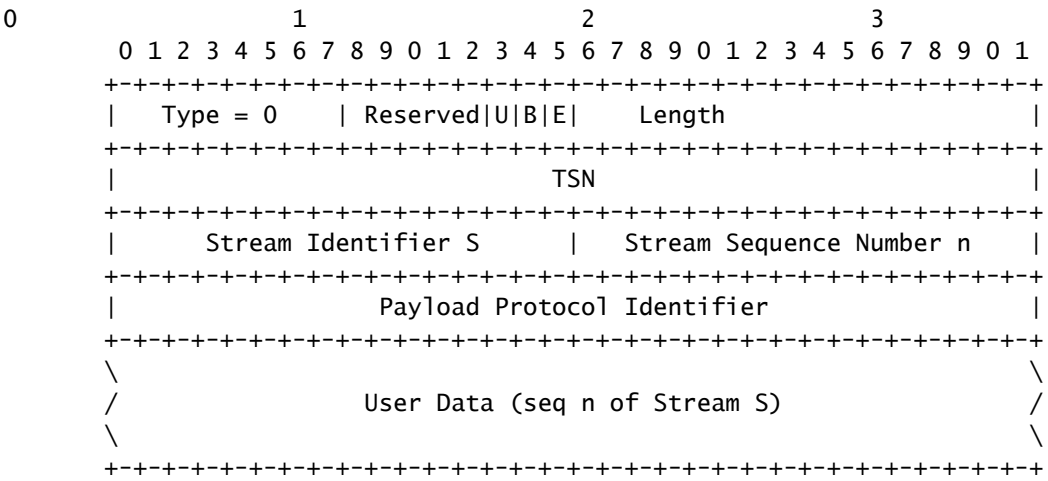


| B | E | Description |
|---|---|---|
| 1 | 0 | First piece of a fragmented user message |

```
|  0 0 | Middle piece of a fragmented user message         |
+------------------------------------------------------------+
|  0 1 | Last piece of a fragmented user message           |
+------------------------------------------------------------+
|  1 1 | Unfragmented message                              |
==============================================================
|             Table 1: Fragment Description Flags           |
==============================================================
```

Gaps in Stream Sequence Numbers can indicate that fragments within the steam or association are being dropped. Wirshark filters can help to identify gaps in SCTP streams:

```
sctp.sack_gap_block_start or sctp.sack_gap_block_end
```

## SCTP Chunk Bundling

Multiple chunks can be bundled into one SCTP packet up to the MTU size, except for the INIT, INIT ACK, and SHUTDOWN COMPLETE chunks:

```
   An SCTP packet is composed of a common header and chunks.  A chunk
   contains either control information or user data.

   The SCTP packet format is shown below:

    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                        Common Header                          |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                          Chunk #1                             |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                           ...                                 |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                          Chunk #n                             |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The chunk length for each SCTP Data Chunk is listed within each data chunk:

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |   Chunk Type  | Chunk  Flags  |        Chunk Length           |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   \                                                               \
   /                          Chunk Value                          /
   \                                                               \
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

# Configurations Impacting Sizing

## SCTP-Param-Template SCTP MTU Sizes

In the SCTP-Param-Template the SCTP-Max-MTU-Size is configured by use of the **sctp-max-mtu-size** option and/or the **sctp-start-mtu-size** option, configured in bytes. 1500 bytes is the default SCTP Max MTU Size and default SCTP Start MTU Size:

```
Exec > Global Configuration > SCTP Parameter Template Configuration
[Local]ASR5500-2# sctp-param-template template_name
[Local]ASR5500-2(sctp-param-template)# sctp-max-mtu-size 1500
[Local]ASR5500-2(sctp-param-template)# sctp-start-mtu-size 1500
```

In the MME Service the SCTP-Param-Template is associated to the MME Context:

```
Exec > Global Configuration > Context Configuration > MME Service
[mme]ASR5500-2(config-mme-service)# associate sctp-param-template S1_MME_SCTP
```

## MME Context Ethernet Interface Configuration

The interface to be used by the MME Service is configured along with the desired IP MTU size in bytes (1500 is default):

```
Exec > Global Configuration > MME Context Configuration
[mme]ASR5500-2(config-ctx)# interface int1/10_s1mme
[mme]ASR5500-2(config-if-eth)# ip address <IP ADDRESS> <SUBNET MASK>
[mme]ASR5500-2(config-if-eth)# ip mtu 2000
[mme]ASR5500-2(config-if-eth)# exit
```

## S1AP UE Radio Capability IE Size

In the MME Serice configuration, S1AP UE Radio Capability IE message sizes are configured. The default value is 9000 bytes:

```
[context_name]host_name(config-mme-service)# s1-mme ue-radio-cap size 9000
```

# Case Study

Here is an example of a fragmented INITIAL CONTEXT SETUP REQUEST/UE CAPABILITY INFORMATION message, which SCTP fragments to meet the configured SCTP Max MTU Size.

In subscriber traces the ICMP message "Fragmentation needed" is returned from the transport router toward the MME after the transmission of the SCTP fragmented INTIAL CONTEXT SETUP REQUEST message.

| No. | Time | Source | Destination | Info | Protocol | Stream identifier | Stream sequence number | Frame length on the wire |
|---|---|---|---|---|---|---|---|---|
| 1 | 2024-10-18 06:45:46.481369410 | MME | eNodeB | DATA (TSN=0) (Message Fragment) | SCTP | 0x0001 | 11339 | 1522 |
| 2 | 2024-10-18 06:45:46.488853860 | MME | eNodeB | DATA (TSN=1) (Message Fragment) | SCTP | 0x0001 | 11339 | 1522 |
| 3 | 2024-10-18 06:45:46.488855090 | MME | eNodeB | DATA (TSN=2) (Message Fragment) | SCTP | 0x0001 | 11339 | 1522 |
| 4 | 2024-10-18 06:45:46.488856320 | MME | eNodeB | DATA (TSN=3) (Message Fragment) | SCTP | 0x0001 | 11339 | 1522 |
| 5 | 2024-10-18 06:45:46.488857560 | MME | eNodeB | InitialContextSetupRequest, UECapabilityInformation, Paging | S1AP | 0x0001,0x0000 | 11339,18839 | 1530 |
| 6 | 2024-10-18 06:45:46.489096060 | Transport Router | MME | Destination unreachable (Fragmentation needed) | ICMP | | | 82 |

In Frame 5, multiple SCTP packets (2) are clubbed together under one IP packet:

```
> Frame 5: 1530 bytes on wire (12240 bits), 1530 bytes captured (12240 bits)
> Ethernet II, Src: Cisco_                              , Dst: Cisco_
> MultiProtocol Label Switching Header, Label:
> Internet Protocol Version 4, Src: MME                   , Dst: eNodeB
> Stream Control Transmission Protocol, Src Port: 36412 (36412), Dst Port: 36412 (36412)
> S1 Application Protocol
> Stream Control Transmission Protocol
> S1 Application Protocol
```

The first data chunk is the Last Segment of that fragmented message, as signified by the **1** mark in the E-Bit information element.

```
> Frame 5: 1530 bytes on wire (12240 bits), 1530 bytes captured (12240 bits)
> Ethernet II, Src: Cisco_                              , Dst: Cisco_
> MultiProtocol Label Switching Header, Label:
> Internet Protocol Version 4, Src: MME                   , Dst: eNodeB
v Stream Control Transmission Protocol, Src Port: 36412 (36412), Dst Port: 36412 (36412)
    Source port: 36412
    Destination port: 36412
    Verification tag: 0xbe183285
    [Association index: 0]
    Checksum: 0xfb290f84 [unverified]
    [Checksum Status: Unverified]
  v DATA chunk (ordered, last segment, TSN: 4, SID: 1, SSN: 11339, PPID: 18, payload length: 1367 bytes)
    > Chunk type: DATA (0)
    v Chunk flags: 0x01
        .... 0... = I-Bit: Possibly delay SACK
        .... .0.. = U-Bit: Ordered delivery
        .... ..0. = B-Bit: Subsequent segment
        .... ...1 = E-Bit: Last segment
      Chunk length: 1383
      Transmission sequence number (relative): 4
      Transmission sequence number (absolute): 3957018401
      Stream identifier: 0x0001
      Stream sequence number: 11339
      Payload protocol identifier: S1 Application Protocol (S1AP) (18)
    > Reassembled SCTP Fragments (7175 bytes, 5 fragments):
      Chunk padding: 00
> S1 Application Protocol
> Stream Control Transmission Protocol
> S1 Application Protocol
```

The second data chunk is not part of a fragmented message, as by the B-Bit and E-Bit information elements are marked with a number **1**:

```
> Frame 5: 1530 bytes on wire (12240 bits), 1530 bytes captured (12240 bits)
> Ethernet II, Src: Cisco_66:8c:90                    , Dst: Cisco_a1:d0:e3
> MultiProtocol Label Switching Header, Label:
> Internet Protocol Version 4, Src: MME                  , Dst: eNodeB
> Stream Control Transmission Protocol, Src Port: 36412 (36412), Dst Port: 36412 (36412)
> S1 Application Protocol
v Stream Control Transmission Protocol
  v DATA chunk (ordered, complete segment, TSN: 5, SID: 0, SSN: 18839, PPID: 18, payload length: 73 bytes)
    > Chunk type: DATA (0)
    v Chunk flags: 0x03
        .... 0... = I-Bit: Possibly delay SACK
        .... .0.. = U-Bit: Ordered delivery
        .... ..1. = B-Bit: First segment
        .... ...1 = E-Bit: Last segment
      Chunk length: 89
      Transmission sequence number (relative): 5
      Transmission sequence number (absolute): 3957018402
      Stream identifier: 0x0000
      Stream sequence number: 18839
      Payload protocol identifier: S1 Application Protocol (S1AP) (18)
      Chunk padding: 000000
> S1 Application Protocol
```

The first SCTP data chunk has 1383 bytes in length. The second SCTP data chunk is 89 bytes in length, so each individual SCTP packet does not exceed the configured SCTP-Max-MTU-Size of 1500 bytes:

```
v Stream Control Transmission Protocol, Src Port: 36412 (36412), Dst Port: 36412 (36412)
    Source port: 36412
    Destination port: 36412
    Verification tag: 0xbe183285
    [Association index: 0]
    Checksum: 0xfb290f84 [unverified]
    [Checksum Status: Unverified]
  v DATA chunk (ordered, last segment, TSN: 4, SID: 1, SSN: 11339, PPID: 18, payload length: 1367 bytes)
    > Chunk type: DATA (0)
    > Chunk flags: 0x01
      Chunk length: 1383
      Transmission sequence number (relative): 4
      Transmission sequence number (absolute): 3957018401
      Stream identifier: 0x0001
      Stream sequence number: 11339
      Payload protocol identifier: S1 Application Protocol (S1AP) (18)
    > Reassembled SCTP Fragments (7175 bytes, 5 fragments):
      Chunk padding: 00
> S1 Application Protocol
v Stream Control Transmission Protocol
  v DATA chunk (ordered, complete segment, TSN: 5, SID: 0, SSN: 18839, PPID: 18, payload length: 73 bytes)
    > Chunk type: DATA (0)
    v Chunk flags: 0x03
        .... 0... = I-Bit: Possibly delay SACK
        .... .0.. = U-Bit: Ordered delivery
        .... ..1. = B-Bit: First segment
        .... ...1 = E-Bit: Last segment
      Chunk length: 89
      Transmission sequence number (relative): 5
      Transmission sequence number (absolute): 3957018402
      Stream identifier: 0x0000
      Stream sequence number: 18839
      Payload protocol identifier: S1 Application Protocol (S1AP) (18)
      Chunk padding: 000000
> S1 Application Protocol
```

Because the total size of the resultant IP datagram, including the SCTP packet and IP headers is smaller than the configured IP MTU size of 2000, these SCTP chunks grouped together within one IP packet at the IP stack level on the MME.

### Solution

S1AP UE Radio Capability IE message sizes are allowed in the configuration at the default value of 9000 bytes. Fragmentation must occur at the SCTP layer to carry these messages at a level smaller than the configured SCTP Max MTU size of 1500.

**sctp-max-mtu-size** being 1500 means any SCTP packet transmitted does not exceed 1500 bytes.

The IP interface MTU size being 2000 means that any IP packet cannot exceed 2000 bytes.

One IP packet can have multiple SCTP packets and each SCTP packet can have multiple chunks. So long as the individual SCTP packet size adheres to SCTP MTU size and total size of all SCTP packets combined together adheres to IP MTU size, the node is behaving as expected.

Packet drops are not occurring within the node, if:

- The total size of SCTP chunks within packet trace frames are less than the configured value of SCTP-Max-MTU-Size in the SCTP-Param-Template and
- The total value of the SCTP PDU + the IP header are less than the configured interface MTU size on the node.

In our case study example, investigation into the transport router revealed that the transport router had an IP MTU size of 1500. Since transport node in between has an MTU of 1500, the packets were dropped by the IP interface on the transport node, promptimg the transport router to send a message to the MME indicating "Destination Unreachable".

# Workarounds to Resolve Drops

### Workaround #1: Reduce the SCTP Max MTU Size

Reduce the SCTP MTU size by configuring **sctp-max-mtu-size** to lower values until you do not see the drops.

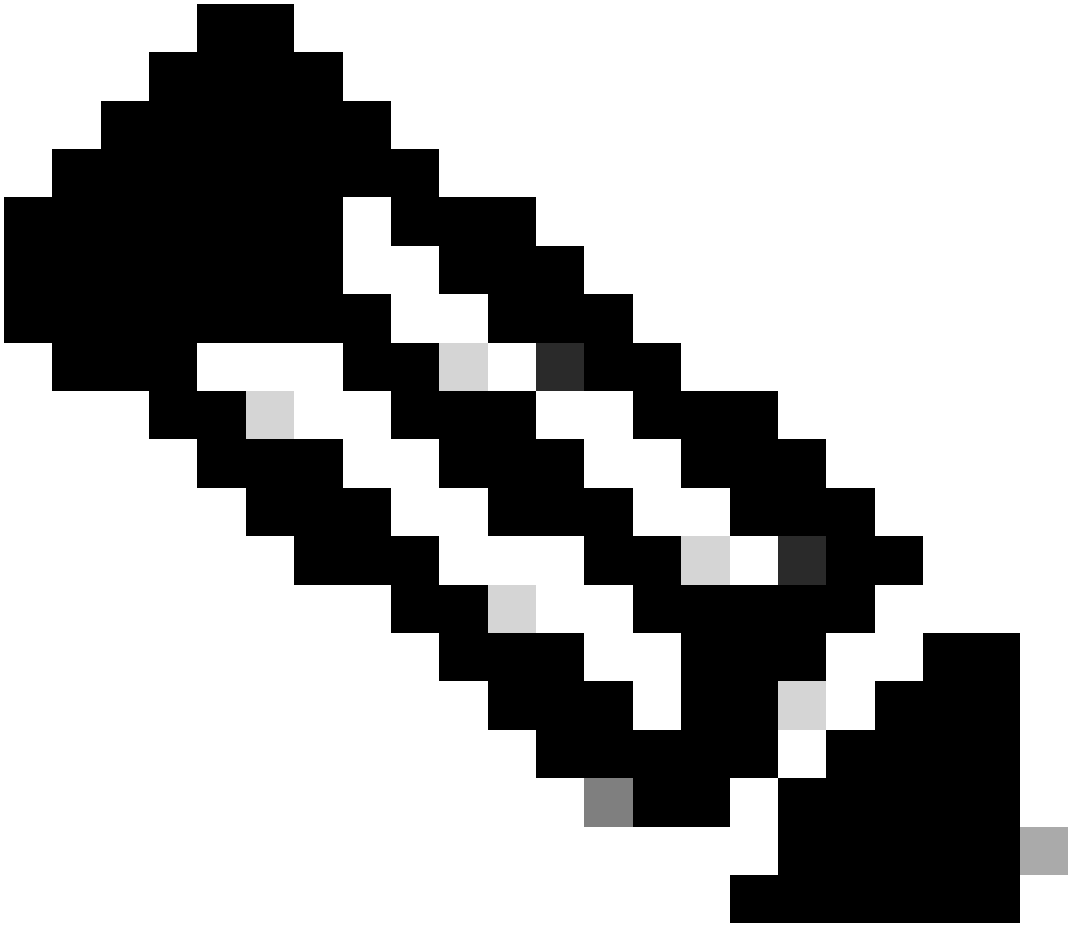

**Note**: When lowering the **sctp-max-mtu-size**, also ensure **sctp-start-mtu-size** is configured to a value less than or equal to the **sctp-max-mtu-size**.

Example: Modify the **sctp-param-template S1_MME_SCTP** (**sctp-start-mtu-size** from 1500 to 1460),

### Workaround #2: Increase the Transport Node MTU Size to Greater Than 1500

## Related Information

- RFC 4960 Stream Control Transmission Protocol

- MME Administration Guide, StarOS Release 21.28 Chapter: UE Radio Capability IE Size
- Command Line Interface Reference, Modes E - F, StarOS Release 21.28 Chapter: Ethernet Interface Configuration Mode Commands IP MTU Size
- Command Line Interface Reference, Modes E - F, StarOS, Chapter: Ethernet Interface Configuration Mode Commands: ip mtu
- SCTP Parameter Template Configuration Mode Commands: sctp-max-mtu-size

- Cisco TechNote - Configure StarOS MTU Interface, APN, and Local Subscribers