ERICSSON

RTP Field Considerations

draft-westerlund-perc-rtp-field-considerations-00

Magnus Westerlund

Outline

- Methodology
- Usage Scenario
- Attackers
- RTP Fields
  - Field
  - Attacks
  - Recommendations
- Summary of Fields

Methodology

- Have analyzed each RTP packet field
- Considered need for end-to-end RTCP
- For each RTP packet field
  - Can the MDD modify it?
  - Does the receiving endpoint need the original value?
  - Does the field need end-to-end authentication?
  - Does the field need end-to-end confidentiality?
  - Motivation for the above
    - Including explaining attacks
  - Hop-by-hop protection will be noted separately at the end
    - Not focus in this presentation

Usage Scenario

- Consider one or more source RTP stream sent from one endpoint (Sender1)
- Though cascaded MDDs
- Arriving at receiving endpoint (Receiver 1)
- Acknowledge that there will be multiple sending and receiving endpoints

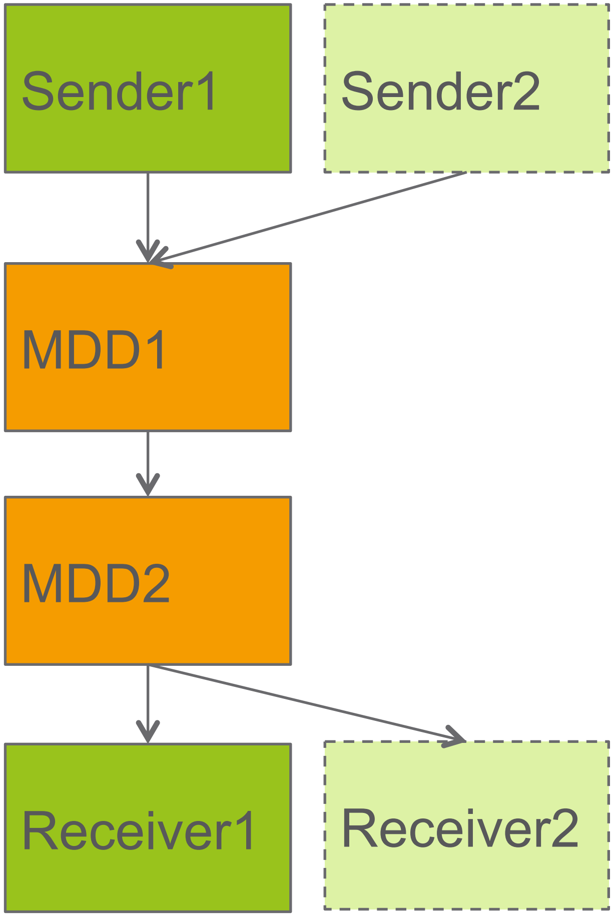

Attackers

- Third Parties
  - Modify, block or inject traffic between nodes
- Malicious MDDs
  - Semi-trusted
  - Have active role
  - Prevent abuse of role
  - Ensure confidentiality of media and sensitive meta data
- Malicious Endpoints
  - Trusted Entity

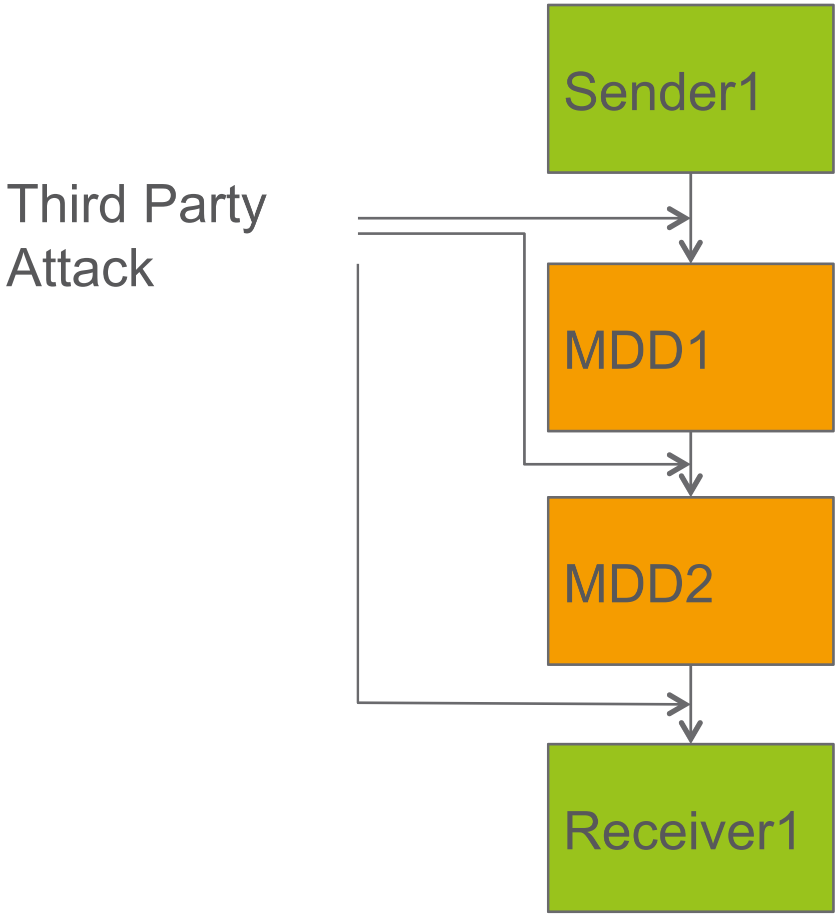

RTP Fields

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|X|  CC   |M|     PT      |       sequence number         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           timestamp                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           synchronization source (SSRC) identifier            |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|            contributing source (CSRC) identifiers             |
|                             ....                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   RTP extension (OPTIONAL)                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          payload  ...                         |
|                               +-------------------------------+
|                               | RTP padding   | RTP pad count |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Version (V)

- Current RTP has value 2
- Will only change if new RTP version is defined
  - Processing dependent of version
- Can the MDD modify it?
  - No
- Does the receiving endpoint need the original value?
  - No, needs to be supported RTP version
- Does the field need end-to-end authentication?
  - No, implicitly protected, but could be included
- Does the field need end-to-end confidentiality?
  - No

Padding (P)

- Indicates the presence of padding in the end of the RTP payload field
- Assumption that padding may be added by originating endpoint
  - To improve privacy by hiding actual payload length end-to-end
- Can the MDD modify it?
  - No

- Does the receiving endpoint need the original value?
  - Yes
- Does the field need end-to-end authentication?
  - Yes
  - Prevent padding processed by Payload format depacketizer
- Does the field need end-to-end confidentiality?
  - Desirable, but not necessary
  - Leaks info that padding is present

Extension Indicator bit (X)

- Indicates presence of header extensions
- Can the MDD modify it?
  - Needs to able
  - Adding or removing header extensions can result in value change
- Does the receiving endpoint need the original value?
  - No
- Does the field need end-to-end authentication?
  - No
- Does the field need end-to-end confidentiality?
  - No

CSRC Count (CC)

- Indicates the number of Contributing Sources (CSRC) that are present
- See CSRC List for discussion of how and why the CSRC count may change
- Media Switching Mixer is one reason to add CSRC list

- Can the MDD modify it?
  - Media Switching Mixer needs to
- Does the receiving endpoint need the original value?
  - Maybe?
- Does the field need end-to-end authentication?
  - Depends
- Does the field need end-to-end confidentiality?
  - No

Marker Bit (M)

- Semantics Payload Format Dependent
  - Video: End of Frame marker
  - Audio: Start of talkspurt
  - May be other semantics
- Leaking media related information to MDD
  - Audio: Talkspurt indication reveals media content
  - Useful for switch start
  - Should be confidentiality protected?

- Video:
- End of Frame not particular sensitive
- Frame marking draft also reveals end of frame
- Necessary for efficient switching on frame boundary
- To indicate to receiver a switch
- Audio's talkspurt indication could be beneficial for this
- Propose using other methods

Marker Bit (M)

- Can the MDD modify it?
  - No
- Does the receiving endpoint need the original value?
  - Yes
- Does the field need end-to-end authentication?
  - Yes
- Does the field need end-to-end confidentiality?
  - Desirable?

Payload Type (PT)

- Indicates the format of the RTP Payload
- Values mapped to formats and parameters using signalling
  - Dependent on direction and pair of nodes
  - Example: H.264 can be:
    - PT=97 on Sender to MDD1 leg
    - PT=101 on MDD1 to MDD2 leg
    - PT=98 on MDD2 to receiver leg

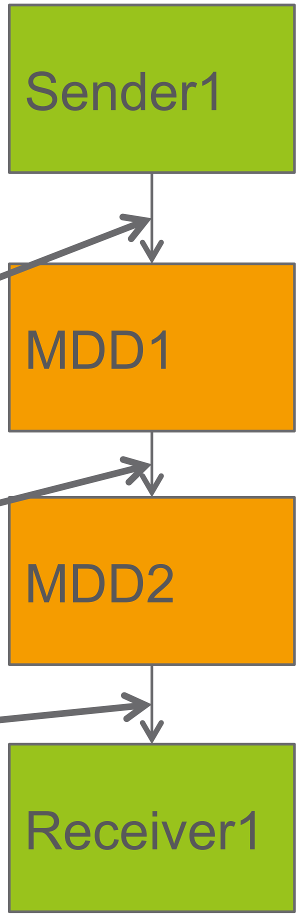

PT Modification Attack

- An attacker modifies the PT value
  - Points to different format than originating sender used
  - Decoded by wrong Payload Depacketizer and media decoder
- Issues:
  - Not sufficiently robust decoders can crash or enable buffer overrun exploits
- Issues:
  - Robust decoders can still produce garbage:
    - Encoded video as PCM
  - Can poison codec state and may trigger concealment actions
- Difficult to exploit buffer overruns in PERC setting
  - Difficult to control input
  - PCM into codec X most likely to succeed

Payload Type (PT)

- Can the MDD modify it?
  - Needs to cope with different assignment
- Does the receiving endpoint need the original value?
  - Yes
  - Original PT to media type mapping also needed
  - Alt. Control signalling so common PT space across all legs
- Does the field need end-to-end authentication?
  - Yes, original value
- Does the field need end-to-end confidentiality?
  - No, difference between media types will commonly be detectable even if E2E protected
  - Protecting it would create difficult signaling requirements

Sequence Number

- Originating Sequence number provides sending order and payload sequence
  - E2E sequence needed for decoding in correct order
  - Expected IV basis
- MDD will need to be able to rewrite the RTP sequence number
  - Stream on/off behavior

- Otherwise switching causes:
  - Loss of transport functionality
    - Loss Detection
    - Inconsistent RTCP reporting
- Packet Sequence Attacks
  - Replay Attack
  - Delay Attack

Replay Attack

- The attacker saves packet sequences sent by source.
- At suitable time attacker replaces source's current packets with some sequence of old packets.
  - Can turn a spoken Yes into a No!

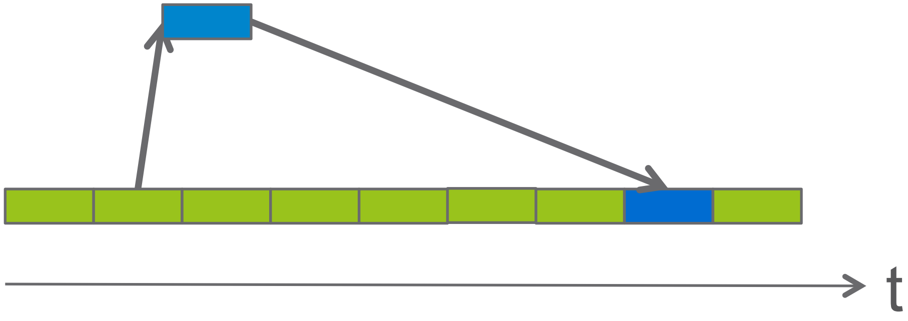

- Replay Protection needed!
- Authenticated original sequence number or equivalent needed
- Only accept newer packets or very near newest received to cope with re-ordering
- Handle sequence number wraps and rekeying events

Delay Attack

- Even with Replay Protection, the MDD can hold packets (Stream Switched off).

- When turning on, use any packet between latest sent to receiver and newest received by MDD.
  - Can be minutes of content

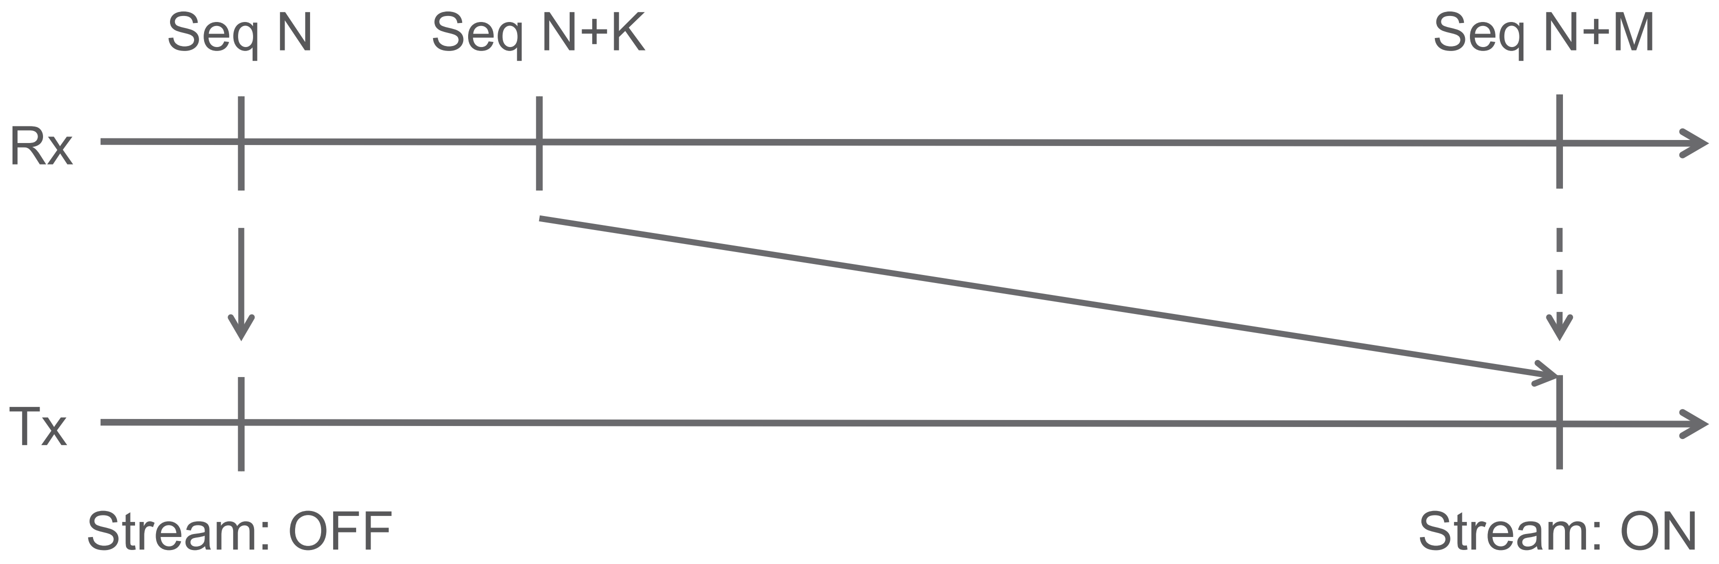

Delay Attack

- End-to-End Sequence numbers don’t solve Delay Attack
- Receiver don’t know:
  - How many packet source sent
  - May have paused at source
- Other Solution needed:
  - Time based
    - RTP Timestamp?
  - End-to-End Reporting

Sequence Number

- Can the MDD modify it?
  - Needs to
- Does the receiving endpoint need the original value?
  - Yes
- Does the field need end-to-end authentication?
  - Yes
- Does the field need end-to-end confidentiality?
  - No

Timestamp

- Expresses Media Timeline
- Switching Media Mixer
  - Need to rewrite Timestamp as outgoing streams SSRC has its own timeline
    - Created by concatenating the different contributing stream's time lines
- Delay Attack Protection
  - Possible use Timestamp
    - Wall clock and Timestamp needs to progress consistently
    - Deal with Clock Skew

- Can the MDD modify it?
  - Needs to given Switching Mixer
- Does the receiving endpoint need the original value?
  - Yes
- Does the field need end-to-end authentication?
  - Yes, if end-to-end
- Does the field need end-to-end confidentiality?
  - No
  - Leaks media time line, but linked to packet sequence for interactive

SSRC

- Sender Source
- Identifies stream context
  - Sequence number space
  - Timestamp space
  - Likely Identifying crypto context
- Media Switching Mixer
  - Has it's own SSRCs
    - Can use CSRC to indicate original SSRC

- Proposed to be THE Source ID in the solution
- Splicing Attack

Splicing Attack

- A Malicious MDD replaces part of sender 1’s stream with sender 2’s stream
- Would be simpler if SSRC Collision can occur
  - MDD can generate collisions and force sources to switch
- Protection
  - Authenticate original source
  - Ensure unique source IDs
  - Prevent media protection rekeying until source ID verified

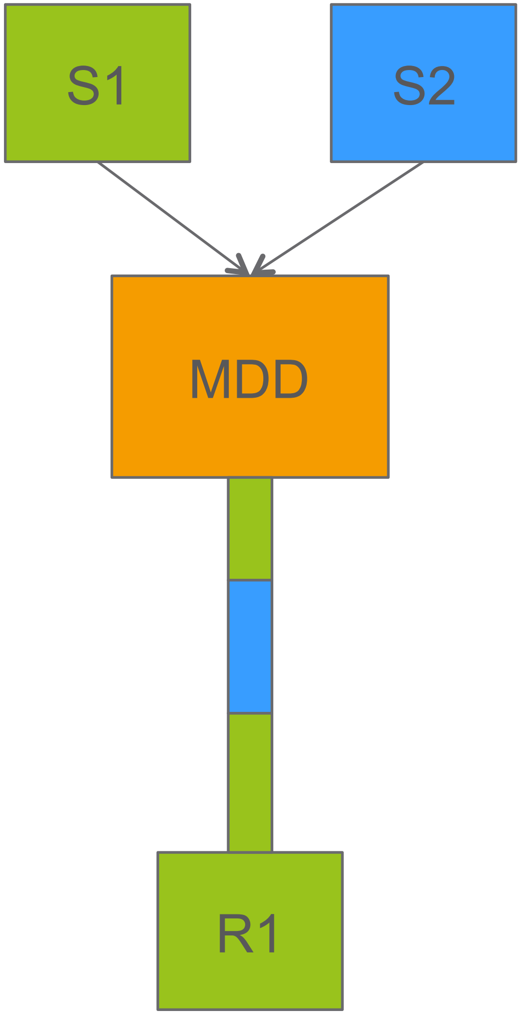

SSRC

- Can the MDD modify it?
  - No
  - May be copied into CSRC by switching mixer
- Does the receiving endpoint need the original value?
  - Yes
- Does the field need end-to-end authentication?
  - Yes
- Does the field need end-to-end confidentiality?
  - No

CSRC List

- Switching Media Mixer
  - Can use CSRC field to indicate original SSRC value
  - Possible solution for knowing originating SSRC for the payload

- Payload Originating Source
  - Indicate that it produces a mix of sources as indicated by CSRC list
  - Not compatible with Switching Media Mixer
  - Mixing PERC endpoints
    - Are they needed?

CSRC

- Can the MDD modify it?
  - Yes, if switching mixer
  - Copy SSRC without modification
- Does the receiving endpoint need the original value?
  - Yes
- Does the field need end-to-end authentication?
  - Yes
- Does the field need end-to-end confidentiality?
  - No

Header Extensions

- Assumes RFC 5285
- Header Extension Id values have the same properties as PTs:
  - Dynamically assigned
  - Depending on signalling
  - Can vary between conference legs
  - Malicious change of IDs could have substantial impact on application

- Need for privacy and confidentiality depends on individual header extensions
- MDD can consume and generate some header extensions
  - Which can be authenticated end-to-end
  - Which needs confidentiality end-to-end

Header Extensions

- Transmission Time offsets
- Gives Transmission time
  - Used by for example congestion control
  - When using hop-by-hop adaptation
    - Rewrite when sending from MDD
    - Measure individual leg
- MDD Modify: Yes
- Original value: No
- End-to-End Auth: No
- End-to-End Conf: No

- SMPTE time-code mapping
  - Unlikely to use by interactive media source
  - Would reveal source information if not continuously increasing
  - However, should come from source if used
- MDD Modify: No
- Original value: Yes
- End-to-End Auth: Yes
- End-to-End Conf: Probably

Header Extensions

- Synchronisation metadata
  - Provides the equivalent of RTCP SR NTP to TS mapping
  - Needed by MDD, especially if Switching Media Mixer
- MDD Modify: No
- Original value: Yes
- End-to-End Auth: Yes
- End-to-End Conf: No

- Client to Mixer Audio Level
  - May be used by MDD to make stream forwarding decision
  - At the same time privacy sensitive, may leak media content [RFC6562]
- MDD Modify: Yes, remove
- Original value: Yes
- End-to-End Auth: Yes, but conditionally
- End-to-End Conf: Desirable, but prevents its use

Header Extensions

- Mixer-to-client audio level
  - Provided for streams with mixed media
  - Does not appear likely in PERC context
  - Not Relevant

- Coordination of video orientation (CVO)
  - Provides video streams orientation (Rotation)
  - Reveals end user actions
    - How they rotate device
    - Privacy sensitive
- MDD Modify: No
- Original value: Yes
- End-to-End Auth: Yes
- End-to-End Conf: Yes

Header Extensions

- Region-of-interest (ROI)
  - Identifies the sub-selection of the video picture provided
  - Controlled by receiver
    - Privacy sensitive
- MDD Modify: No
- Original value: Yes
- End-to-End Auth: Yes
- End-to-End Conf: Yes

- SDES Information
  - Provides SDES items like CNAME, MID and RID
  - CNAME can be sensitive
    - Can be made safe
- MDD Modify: No
- Original value: Yes
- End-to-End Auth: Yes
- End-to-End Conf: No (Maybe)

Header Extension

- Treatment depends on header extensions:
  - MDD changeable
  - End-to-End Authenticated
  - End-to-End Confidentiality
- The whole header extension framework can be added and removed

- Notes that end-to-end authenticated header extension has an issue with ID of extensions

Payload

- Contains the media content that PERC shall confidentiality protect end-to-end.
- Can the MDD modify it?
  - No
- Does the receiving endpoint need the original value?
  - Yes
- Does the field need end-to-end authentication?
  - Yes
- Does the field need end-to-end confidentiality?
  - Yes

Padding

- The Padding consists of a Padding counter and up to 255 bytes of Null Padding
- Can be used to conceal the size of the encoded payload
- Can the MDD modify it?
  - No
- Does the receiving endpoint need the original value?
  - Yes
- Does the field need end-to-end authentication?
  - Yes
- Does the field need end-to-end confidentiality?
  - Yes

RTCP

- A lot of the RTCP information will be leg specific
  - RTCP SR/RR
  - RTCP FB messages related to transport
- Some information is end-to-end
- RTCP SDES items
  - Some are privacy sensitive
    - Name, Location,…
  - Some needed by MDD
    - CNAME, MID, RID
  - SDES: CNAME, MID
    - If changeable by MDD
      - Miss-associate streams
      - Miss-sync with wrong streams
    - Needs End-to-End authentication to prevent attacks

RTCP

- RTCP FB
  - ROI requests
    - E2E
    - Privacy sensitive
  - AFB – Application Layer Feedback
    - Unknown
- RTCP APP
  - Unknown content

- To me it appear that we will have to define both:
  - End-to-End authenticated
  - End-to-End confidential
- Issue with End-to-End is that any source IDs (SSRC) needs to be common space
  - No SSRC translation in MDD

Hop-by-hop protection

- No reason to not authenticate all data sent hop-by-hop
- Confidentiality can be discussed on field per field basis
  - See draft
- SRTP is not the state of the art in preserving privacy