

Video Frame Info RTP Header Extension

draft-ietf-avtext-framemarking-04

Note Well: <https://datatracker.ietf.org/ipr/2876/>

Espen Berger, Suhas Nandakumar, Mo Zanaty

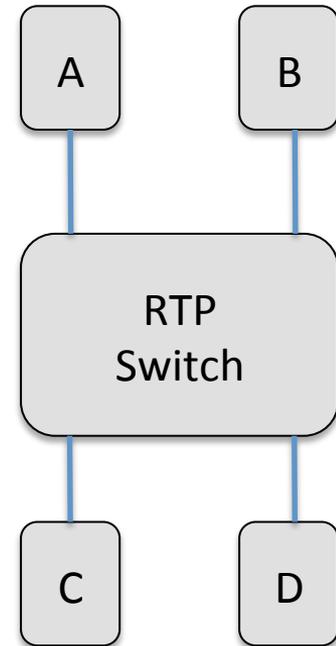
AVTEXT WG

IETF 98 – Chicago, IL – March 28, 2017

Review: Main Motivation

Payload-Agnostic RTP Switch

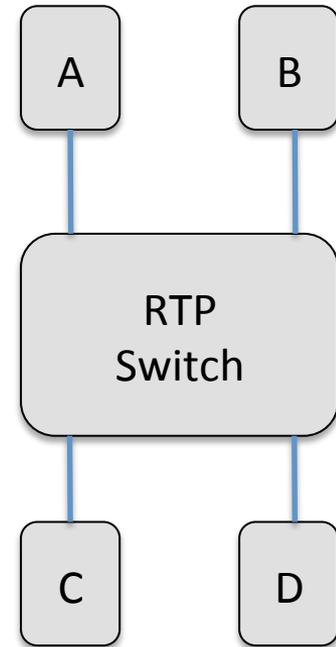
- Payload may be encrypted
 - Avoid decryption cost to improve switch scale and latency
- Payload may be encrypted end-to-end
 - Impossible to decrypt / inspect payload without end-to-end keys
- Payload may be unknown format
 - Codec-agnostic switching can support any format, old or new



Review: More Motivations

Smarter RTP Switch

- Clean video switching at intra-frames
- Better recovery during packet loss
- Drop least important packets during congestion
- Drop scalable enhancement layers for constrained endpoints



Smarter Endpoints

- Better recovery during packet loss

Changes in version -04

- Allowed the short header for scalable streams.

3.1. Extension for Non-Scalable Streams

The following RTP header extension is RECOMMENDED for non-scalable streams. It MAY also be used for scalable streams if the sender has limited or no information about stream scalability.

- Allowed the long header for non-scalable streams.

3.2. Extension for Scalable Streams

The following RTP header extension is RECOMMENDED for scalable streams. It MAY also be used for non-scalable streams, in which case TID, LID and TLOPICIDX MUST be 0.

Changes in version -04

- Added section for LRR

3.4.1. Relation to Layer Refresh Request (LRR)

Receivers can use the Layer Refresh Request (LRR) [I-D.ietf-avtext-lrr] RTCP feedback message to upgrade to a higher layer in scalable encodings. The TID/LID values and formats used in LRR messages correspond to the same values and formats specified in Section 3.2.

Changes in version -04

- Clarified this only applies to video Source RTP Streams, not audio nor RTX/FEC Redundancy RTP Streams.

This extension is only specified for Source (not Redundancy) RTP Streams [RFC7656] that carry video payloads. It is not specified for audio payloads, nor is it specified for Redundancy RTP Streams. The (separate) specifications for Redundancy RTP Streams often include provisions for recovering any header extensions that were part of the original source packet. Such provisions SHALL be followed to recover the Frame Marking RTP header extension of the original source packet.

Changes in version -04

- Imported relevant information from:

draft-aboba-avtcore-sfu-rtp

- MUST match the corresponding info in video payload structures
 - H.264-SVC NAL/PACSI, H.265 NAL/PACI, etc.

Open Issues

- Marking for Redundancy RTP Streams?
 - Not intended for RTX, FEC
 - Recommend a separate draft if RTX or FEC need some form of marking
- RTP Padding Only bit proposal for PERC use
 - Can the Discardable bit be used instead?
 - Can RTX/FEC PT be used instead of real media PT?
 - If not, do not recommend adding this oddity

Next Steps

- Ready for WGLC after resolving open issues.
- Questions?
- Thank you!