

Sending Multiple Media Streams in a Single RTP Session

draft-ietf-avtcore-rtp-multi-stream-05

draft-ietf-avtcore-rtp-multi-stream-optimisation-03

Jonathan Lennox – Vidyo

Magnus Westerlund – Ericsson

Qin Wu – Huawei

Colin Perkins – University of Glasgow

Sending Multiple Media Streams in an RTP Session

- Clarifications to RFC 3550 regarding behaviour of end points that send and receive multiple streams simultaneously
- Changes in draft-ietf-avtcore-rtp-multi-stream-05:
 - Update Section 5.4.1 (“Choice of SSRC for feedback packets”) to clarify which SSRC is used when sending RTCP feedback
 - “For example, if one sets a TMMBR limitation, the same SSRC needs to be used to remove the limitation...” and similar types of clarification
 - Rewrite the first half of the draft (up to end of Section 5.3.1) for clarity – further clarifications needed on later parts of the draft

Grouping RTCP Reception Statistics and Other Feedback

- Optimisations to reduce cross-reporting for end points that have multiple SSRCs
 - e.g., bundle end points that send and receive audio and video streams simultaneously
 - draft-ietf-avtcore-rtp-multi-stream-optimisation-03

Changes in -multi-stream-optimisation-03 (1/3)

- Clarify that a new SDES item type number is to be used for the RGRP item
 - It has the same syntax as the CNAME item, but uses an item type number to be assigned by IANA

Changes in -multi-stream-optimisation-03 (2/3)

- Add examples of cases where the SSRC used in feedback has semantic meaning, to motivate why each SSRC sends separate feedback
 - “For example, an RTP video sender might choose to treat packet loss feedback received from SSRCs known to be audio receivers with less urgency than feedback that it receives from video receivers when deciding what packets to retransmit, and a multimedia receiver using reporting groups might want to chose the outgoing SSRC for feedback packets to reflect this.”
 - “For example, it makes no sense for an SSRC that is receiving video to send a VoIP metric RTCP XR block...”
 - RTP/AVPF duplicate report suppression ensures that only a single report is sent, and grouping suppresses RR blocks not entire reports, so sending reports from a semantically meaningful SSRC has no extra cost

Changes in -multi-stream-optimisation-03 (3/3)

- **Adds Middlebox Considerations (Section 3.5)**
 - Middleboxes that have no SSRC of their own cannot use reporting groups
 - They have no reports to group...
 - Middleboxes that have >1 SSRC of their own can use reporting groups to group the RTCP reports they generate
 - E.g., can occur in a bundled session if a middlebox is acting as both an RTP audio mixer and an RTP video mixer for flows that are multiplexed in a single RTP session, and hence has an SSRCs for each the two mixes, and that wants to avoid cross reporting
 - Middleboxes cannot use reporting groups to group the RTCP packets that they are forwarding
 - But can aggregate forwarded RTCP packets into compound RTCP packets following Section 6.1 of [RFC3550] and Section 5.3 of [draft-ietf-avtcore-rtp-multi-stream]
 - Middleboxes that forward packets sent by members of a reporting group **MUST** forward the RTCP SDES RGRP items and RTCP RGRS packets
 - Middleboxes that rewrite SSRCs **MUST** make corresponding changes to RTCP SDES RGRP items and RTCP RGRS packets

Status and Next Steps

- draft-ietf-avtcore-rtp-multi-stream-05 needs further editorial work; algorithms need simulation and validation
- draft-ietf-avtcore-rtp-multi-stream-optimisation-03 is essentially done – please review