RTP Considerations for Endpoints sending Multiple Media Streams

`draft-lennox-avtcore-rtp-multi-stream-00`

AVTCore, IETF 84, 2 August 2012

Jonathan Lennox

jonathan@vidyo.com

Magnus Westerlund

magnus.westerlund@ericsson.com

RTP Multi-source: Motivation

- Historically, endpoints usually only sent one source per RTP session
- A number of use cases emerging where this is changing
  - BUNDLE
  - CLUE
  - Multi-source Mixers

What this draft does

- Re-visits RFC 3550 to clarify behavior for multi-source endpoints.
- May need to update RFC 3550 to change some RTCP timing rules (to be determined).
- Gives recommendations on optimizations for reception reports.

RTP

- Stay within your share of session bandwidth (as determined by signaling and congestion control), but not necessarily independently or uniformly for all your streams.
- You can re-allocate bandwidth among your streams, depending on what you think is most useful
  - Variable-rate codecs
  - Change codec
  - Enable or disable streams

RTCP: Initial RTCP

- RFC 3550: in unicast sessions, a participant MAY send initial compound RTCP immediately.
- Recommendation: this applies to each new SSRC of a multi-stream endpoint, as well.

RTCP: combine multiple sources' RTCP packets

- RFC 3550: mixers and translators SHOULD combine RTCP packets from multiple sources into a single compound RTCP packet, up to MTU.
- Recommendation: this applies to multi-stream endpoints, as well.
- Open issue: how to calculate RTCP timing in this case?
  - Aggregate packets that are "close" in time?
  - Calculate one interval based on your share of the RTCP bandwidth?
  - Should this be different for AVP and AVPF?

Reception Reports

- An endpoint MUST send reception reports (in SR or RR packets) for every active media stream it's receiving.
- However, if you send reports from every *source* on behalf of every other source, you end up spending most of your RTCP bandwidth on redundant reports.
  - N media sources (active and inactive) in a session, and S active senders per reporting interval: N*S reception reports, or unnecessary round-robinning.
  - In the general case, quadratic.
- Proposal: change to E*S, where E is number of endpoints (often 2).

Recommendations for reports

- Endpoint SHOULD NOT send reception reports from one of its source about another of its own (“self-reports”).
- Endpoint SHOULD NOT send reception reports about remote sources from multiple local sources; instead, pick one “reporting” source per remote source.
- You still need to send SR or RR packets for every source, but for everything other than the reporting source they don’t contain any report blocks.

Consequences of report limitations

- This RTCP traffic might look like it's generated by receivers experiencing a network disconnection.
  - In the worst case, congestion control might think it's seeing a complete congestion collapse.
- But requires fairly sophisticated RTCP analysis; in most cases, senders just care about reports about themselves.
- Senders already have to be prepared that any given SR/RR doesn't describe them, due to round-robining.

Numeric estimate

- Two source-projecting mixers, 100 sources each, 8 active sources each, 16-byte CNAMEs.
- RTCP excluding reception reports: 200 SDES, 184 RR, 16 SR: approx. 6.5 kB / report interval
- Naïve reception reports: 16 * 184 + 15 * 16 report blocks: approx. additional 76 kB / report interval (vs. 2 * 8 report blocks = approx. 0.4 kB with new rules)
- To first approximation, report interval = bytes_per_interval / rtcp_bw, so the interval will be about 11 times longer.
- Gets worse as source switching happens inside report intervals, so number of active sources grows.

Alternative solution

- Explicitly signal and negotiate that you'll be doing this.
- Indicate (in RTCP) which sources originate from a single "reporting group", which won't do self-reports, and in which only one source will be doing remote reports.

Next steps

- Will address open issues, and determine whether a 3550 update would be needed.
- Does the WG want the multi-source clarifications for a WG item?
- If so, should we also do
  - Timing rule changes?
  - Reporting rule optimizations?