

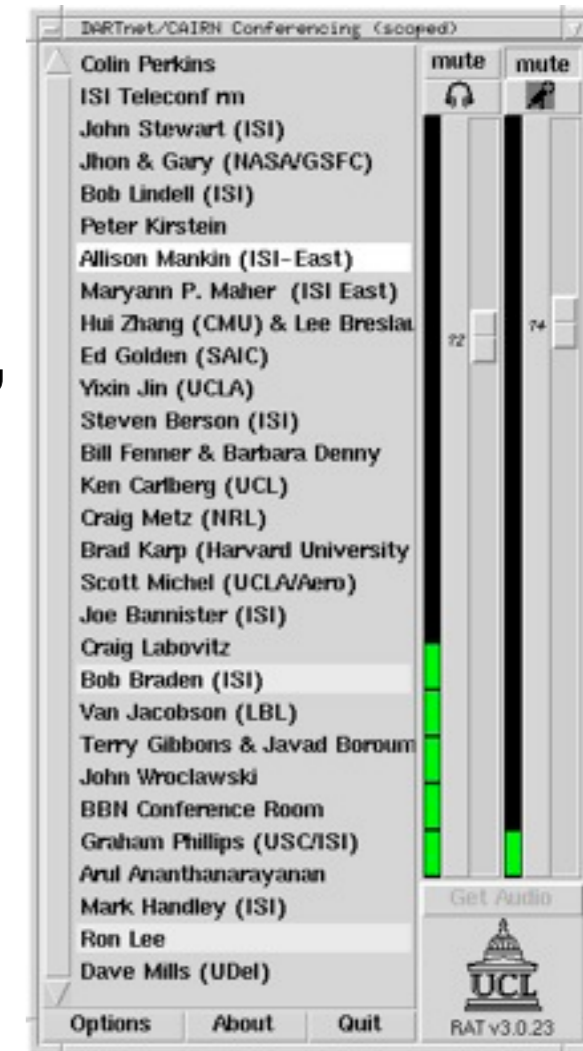
RTP Multiplexing Architecture

draft-westerlund-avtcore-multiplex-architecture-00

Magnus Westerlund
Bo Burman
Colin Perkins

Rationale

- RTP is a group communication protocol
- Early implementations made extensive use of the group communication features, and distributed multi-stream end-points
 - Multicast *and* unicast using RTP mixers and translators
- Most VoIP deployments, and many more recent conferencing systems, did not
- The RTCWeb and CLUE working groups are again considering the use of RTP for group communication and multi-stream scenarios
- This draft provides guidance on how to do this



(1997)

Overview of the Draft

- Discussion of issues
- Guidelines
- Clarifications and protocol extensions

Discussion of Issues

- RTP multiplexing points for multiple streams
 - RTP sessions and SSRCs
 - Why the payload type should not be used to demultiplex different streams
- RTP topologies and issues arising
 - Point-to-point; multipoint using multicast, RTP translators, or RTP mixers; multipoint using full unicast mesh; and distributed end-points
- When to use multiple SSRCs in an RTP session vs. when to use multiple RTP sessions
 - RTP/RTCP considerations: legacy end-points, varying sets of senders, cross-session RTCP, binding related sources, FEC, translators, multiple media types within an RTP session
 - Signalling: bandwidth, managing SSRCs, limitations of SDP m= lines
 - QoS, NAT/firewall traversal, multicast, security

Guidelines

- The draft proposes the following general guidelines:
 - Don't require the same SSRC across sessions, use RTCP to associate streams
 - Use SSRC multiplexing for additional media streams from an end-point with similar purpose
 - Use additional RTP sessions for streams with different purposes
 - When using session multiplexing, use SDP grouping
 - RTP/RTCP extensions need to support SSRC and session multiplexing
 - Extensions to improve transport reliability (e.g., FEC and retransmission) need to support SSRC and session multiplexing
- We believe these are appropriate, but discussion is welcome (...offline this week, or on the list)

Clarifications and Extensions

- We note a number of areas where RFC 3550 is not clear regarding multiple sources:
 - If an end-point has multiple SSRCs (e.g., because it's a single participant sending multiple video streams), should it send/receive RTCP for each of those SSRCs?
 - If an end-point has multiple SSRCs, should it send RTCP report relating to its own traffic?
 - Can a compound RTCP packet contain RTCP packets from multiple SSRCs?
- We note some extensions that might be useful
 - Transport multiplexing → draft-westerlund-avtcore-transport-multiplexing-01
 - Media source identification → draft-westerlund-avtext-rtcp-sdes-srcname-00
 - SSRC limitations with RTP sessions → draft-westerlund-avtcore-max-ssrc-00
 - SDP bandwidth attribute → draft-westerlund-mmusic-sdp-bw-attribute-00

Next Steps

- Does the group agree that we should provide some guidance on multiplexing, and when to use multiple RTP sessions?
- If so, is this draft a good starting point, and should it be adopted as a working group draft?
 - Looking for agreement, or otherwise, on the general approach; we will discuss details later