

MULTIPLEXING RTP SESSIONS ON A SINGLE TRANSPORT

[draft-westerlund-avtcore-transport-multiplexing-02](#)

Colin Perkins / University of Glasgow
Magnus Westerlund / Ericsson

Outline



- › Goal
- › Proposal
- › Options
- › Does Consensus Exist?
- › If Discussion is Required:
 - Background
 - Considerations
 - Determine Consensus
- › Next Steps

Goal



- › The goal is to create an AVTCORE WG consensus on how to proceed with work regarding supporting multiple media flows over one or more RTP sessions over a single transport flow

- › Definitions
 - Transport Flow = Series of packets matching the 5-tuple for the used transport protocol, such as UDP.
 - RTP session = Single SSRC space

Options




› We believe that the WG has the following options:

1. Define a mechanism that enables multiple RTP sessions over a single lower layer transport



2. Clarify the RTP specification and operation of a single RTP session with multiple media types in it



3. Do Both 1 and 2
4. Do nothing

Proposal



- › We do propose that we do **Option 3**, i.e work on both.
- › For “*Multiple RTP sessions over Single Transport*” do:
 - Support MMUSIC in their development of the signaling support to ensure that Multiple RTP sessions over a single transport functions
 - Develop a mechanism that enables multiple RTP sessions over a single transport flow
 - › [draft-westerlund-avtcore-transport-multiplexing-02](#) contains a proposal to the WG
- › For “*Multiple Media Types in a single RTP session*”:
 - Which appears to be what Harald suggests in [draft-alvestrand-rtp-sess-neutral-00](#)
 - Support MMUSIC in their development of the signaling support for multiple m= lines pointing to a single RTP session
 - Consider clarifications and optimizations for better handling of multiple media streams within any single RTP session
 - › This is not limited to be usable by multiple media types
- › Document the usage considerations of having multiple media types:
 - A potential place to do this would be a WG version of [draft-westerlund-avtcore-multiplex-architecture-01](#)

Does Consensus Exist?



- › Lets see if we can directly determine WG consensus on what work the WG should do:
 1. Work on both “Multiple RTP sessions over single transport” and “Multiple Media Types in a single session”
 2. Do only work on “Multiple RTP sessions over single transport”
 3. Do only work on “Multiple Media Types in a single session”
 4. Do Neither!

- › If not very clear consensus lets discuss things!

Background



- › In RTCWEB WG a desire to enable using multiple media types like audio and video over a single lower layer transport flow (IP/UDP) was expressed to minimize NAT traversal costs
- › In Quebec (IETF81) there was off-line discussion among a number of people:
 - In RTCWEB WG a proposal to use a single RTP session for multiple media type was presented without disagreement
 - In AVTCORE the RTCWEB intentions where presented
 - In AVTCORE the question was asked if the WG was interested in a solution addressing the presented short comings of a single RTP session with multiple media types:
 - › A Strong Consensus was for working on such a solution.

More Background



- › In Taipei (IETF 82) we (Westerlund & Perkins) come with a clear proposal for how to solve it
- › Lennox proposed that multiple RTP sessions was not needed
- › We also asked a question if the WG thought it necessary to have both:
 - Single RTP session with multiple media types
 - A solution supporting multiple RTP sessions
- › This caused a large discussion in the WG and a consensus call after the meeting.
 - The consensus call had no clear outcome
 - › Quite many was asking for more information
- › In MMUSIC WG the BUNDLE proposal has been adopted as WG item enabling multiple media types in an RTP session in SDP based signaling

Considerations




- › RFC 3550 does not forbid using multiple media types in an RTP session (Option 1)
 - It RECOMMENDS against it with less than clear motivations
- › The analysis in Section 6.7 of [draft-westerlund-avtcore-multiplex-architecture-01](#) has a number of considerations for its usage:
 - RTP Mixers and Translators can't use session context only as indication of how to process content
 - Legacy implementations can cause issues if not whole chain is verified in signaling to support multiple media types
 - › Including Deep Packet Inspection
 - Limitations in number of RTP payload types worse than for single media type
 - RTCP bandwidth can become difficult to configure to efficiently used when media is significantly different in RTCP needs.
 - Continued on next slides

Limitations



› Interworking in Multi-party session


- Clients A and B are using multiple RTP sessions in a audio and video conference
- Client C connects to RTP mixer using single session
- To avoid SSRC translation A1, A2, V1 and V2 must be unique
 - › Likely to fail as A and V are different SSRC spaces for A and B and not checked.
- SSRC translation has issues:
 - › Using SRTP incur complexities for decrypt encrypt and mixer must be in security context
 - › Potential barrier for deploying new RTCP extensions



Limitations



- › Poor support for de-composite end-points



- › If Audio and Video is in same RTP session the audio handling component must receive all video also.

Considerations



- › Despite the limitations with multiple media types in a single RTP session, there are use cases.
- › The biggest obstacle against using it in systems for which the limitations imposed are acceptable has been the lack of support in signaling
 - The Bundle WG item in MMUSIC appear to remove that obstacle
- › Unguided usage can result in applications not be aware of limitations
 - Causing significant issues, especially when interworking
- › This WG isn't selecting if it being used by application developers
 - However, we can document the considerations around its usage
 - Most optimization and fixes discussed in [draft-lennox-rtcweb-rtp-media-type-mux-00](#) and Section 10 of [draft-westerlund-avtcore-multiplex-architecture-01](#) should be considered independently as they applies to multiple media streams in a RTP session, independent if one uses one or more media types.

Considerations



- › The WG appears to have consensus that using a single transport flow is desirable in a number of use cases
- › RTP sessions are a useful tool:
 - Clearly used in a number of applications
 - We have some RTP/RTCP Extensions that rely on multiple RTP sessions, like
 - › XOR based FEC (RFC5109 and RFC2733)
 - › Retransmission (RFC4588)
 - › Scalable Video Coding (RFC6190)
- › Single RTP session has limitations
- › WG should consider RTP applications that want to use multiple RTP sessions and prefer a single transport flow
 - We need a solution also for these applications

Options




› We believe that the WG has the following options:

1. Define a mechanism that enables multiple RTP sessions over a single lower layer transport



2. Clarify the RTP specification and operation of a single RTP session with multiple media types in it



3. Do Both 1 and 2
4. Do nothing

Proposal



- › We do propose that we do **Option 3**, i.e work on both.
- › For “*Multiple RTP sessions over Single Transport*” do:
 - Support MMUSIC in their development of the signaling support to ensure that Multiple RTP sessions over a single transport functions
 - Develop a mechanism that enables multiple RTP sessions over a single transport flow
 - › [draft-westerlund-avtcore-transport-multiplexing-02](#) contains a proposal to the WG
- › For “*Multiple Media Types in a single RTP session*”:
 - Which appears to be what Harald suggests in [draft-alvestrand-rtp-sess-neutral-00](#)
 - Support MMUSIC in their development of the signaling support for multiple m= lines pointing to a single RTP session
 - Consider clarifications and optimizations for better handling of multiple media streams within any single RTP session
 - › This is not limited to be usable by multiple media types
- › Document the usage considerations of having multiple media types:
 - A potential place to do this would be a WG version of [draft-westerlund-avtcore-multiplex-architecture-01](#)

Consensus Call



- › Lets see if we can determine a WG consensus on what work the WG should do:
 1. Do only work on “Multiple RTP sessions over single transport”
 2. Do only work on “Multiple Media Types in a single session”
 3. Work on both “Multiple RTP sessions over single transport” and “Multiple Media Types in a single session”
 4. Do Neither!

Next Steps



- › If we are to do “Multiple RTP sessions over single transport” then we should aim at picking a solution by the next meeting
 - Do anyone else have proposal than [draft-westerlund-avtcore-transport-multiplexing-02?](#)
 - › Please submit a draft soon if you do.
- › If we are to do “Multiple Media Types in a single session” do we have authors for:
 - Writing up clarifications and possible optimizations for multiple media streams in an RTP session?
 - Writing up the considerations around using multiple media types?