AVTCORE WG Minutes for IETF 85:

Session One: 5 Nov 2012

Chairs: Magnus Westerlund Roni Even
Note Takers: Stephen Botzko and Kevin Gross

WG Chairs opened the meeting presenting the Note Well and the WG status.
RFC 6679 published and draft-ietf-avtcore-monarch-22 in editor queue.
draft-ietf-avt-srtp-not-mandatory-10-09 is in AD follow up. Robert Sparks (WG's AD) commented that he wants to bring srtp-not-mandatory back into working group when path forward has been agreed on. There was a general call for reviewers on the other working group documents. Eric Rescorla volunteered to review the SRTP crypto suit Aria. The status was finished up with a review of the WG's Milestones.

Multiple Media type in an RTP session

Authors: Colin Perkins, Magnus Westerlund, Jonathan Lennox
draft-ietf-avtcore-multi-media-rtp-session-01
Colin Perkins presenting

This drafts propose updating RFFC 3550, and 3551 to allow multiplexing and to identify cases where it is safe. The updated draft (-01) includes some clarifications and notes some open issues. Including a overrule section for RFC 3551 separating audio and video. Updates the draft's section 3.3 to clarify what is meant by architectural equality.

Single RTCP reporting interval constraint (Slide 4)
-determines participant timeout period
-does not depend on sending rate for SSRC
-all participants report at roughly the same rate
-reporting interval constraint is inherent to RTP design

Issue: RTCP packet rate can match the audio RTP packet rate when multiplexing audio and video. Note, this is not unique to audio/video multiplexing; it is just the most common example. The open Question: Is this really a problem, if so what kind of fix is needed. Harald Alvestrand doesn't see it as a problem. The real question is how much feedback is needed. Roni Even (as individual) commented that one could report video based on frames per second, not packets. Jonathan Lennox though maybe there should be guidance on feedback rate independent of multiplexing, as the current formulas sometimes create too much feedback to be useful. Colin responded since you can tune the bandwidth as you like, the guidance is when not to use the defaults and what alternative tuning should be used to get sensible behavior.

Roni Even (as individual) asked why should RTCP feedback be based on data rate and not packet rate anyway? Jonathan responded that just changes the problem to a mismatched packet rate. The serious problems happen when the RTCP packet rate exceeds the RTP rate for one of the streams.

Kevin Gross asked if this is a problem also happening with the other multiplex proposals. Colin Perkins commented that the shim approach doesn't need this.
Jonathan Lennox commented that even if the shim approach doesn't need it, we still might need a fix to ensure sensible behavior. Harald had the opinion that we don't need to change it, we've lived with it for 20 years. Colin Perkins disagreed, as we are starting to use the protocols in different ways.

Issue: SSRC Timeout (Slide 6)
-Need to ensure RTCP timeout.
-Currently are two options (reduced minimum vs 5 sec)
-Proposals to go with 5 second

Next Steps, is this a separate clarification for RFC 3550 or not? No feedback.

RTP Considerations for Endpoints sending multiple media streams

Authors: Jonathan Lennox, Magnus Westerlund
draft-lennox-avtcore-rtp-multi-stream-01
Jonathan Lennox presented.

This draft is to clarify usage of RTP/RTCP with multiple sources per session. Usages that are being intended in CLUE WG, with the Bundle/MMT proposals in the MMUSIC WG, and being used by multi-source RTP mixers. The draft has been updated by adding an explicit RTCP SDES item to describe reporting groups, added calculations motivating use of reporting groups, several additional open issues.

Reporting groups is a Proposal for how to indicate which SSRCs that are common to an end-point and where only a subset of SSRC will send RTCP Receiver reports. It include a number of rules that applies to reporting groups and their RTCP usage.

An Open issue: how to signal/negotiate in SDP as legacy devices won't interoperate smoothly with these reporting groups.

Another OPEN ISSUE: The calculation of the RTCP variable avg_rtcp_size. The transmission interval calculation doesn't work properly if several sources are aggregated into a single compound RTCP. Do we need to change how transmission interval (and/or avg_rtcp_size) is calculated?

Cullen Jennings asked; what if SSRCs in a reporting groups use different DSCPs when sending media? Jonathan Lennox responded that a receiving endpoint still report on individual sources independently. Colin Perkins commented that we may need to clarify that one can only use reporting grouping if not experiencing different network behavior for different stream. RFC 3550 requires cross reporting and thus the different combination will be exercised.

Colin Perkins wondered if an endpoint has 5 sources, does the endpoint get 5 report slots to send whenever it needs to, or can it simply pick the source to use? Johathan responded that he haden't thought about that yet. Magnus Westerlund (As Individual) commented that in early mode the endpoint's reports could use the bandwidth however they desire.

Roni Even asked ??? Jonathan responded that one can send one compound packet with SR/RR for two streams rather than one per stream. Even if there are no report blocks in the SR/RR included. Colin Perkins commented that we need to define correct behavior when this is used (for instance, stripping SDES can't be allowed).
Jonathan commented that if you're using Report Groups intermediate devices must forward SDES.

Colin Perkins commented that he had not been convinced that this makes sense. However, he has been convinced that it is useful. The proposal does appear to be a reasonable approach.

Magnus Westerlund (As Chair): There is another presentation on similar topic. Let's see what we have there before making any decisions

Multi-source endpoint RTCP Advertisement

Authors: Qin Wu, Roni Even
draft-wu-avtcore-multisrc-endpoint-adver-02

This draft is a proposal for how to achieve suppression of redundant RTCP reception reports in two scenarios. (a) multiple send sources in the same session. (b) RTCP self-reporting.

Colin Perkins commented that the overhead on slide 4 is higher than 16 bits per sdes, though that does not affect the conclusion that the savings are significant.

The proposed solution is to define a single media source as the "primary active" report source. This proposal have some open issues; requires negotiation in SDP; how does this changes RTCP timing and timeout rules, and also the rules for compound packets; the current draft omits source projection mixer case; current advertisement only applies to multiplexed media in the same session (does not apply to multiplexed media from different sessions).

Jonathan Lennox commented that source projection case is important, as the savings can be very significant there. Jonathan also had a question what the report source list is on slide 5. Answer: it describes the sources on whose behalf you are reporting from the same endpoint. Jonathan commented that this could be large. Qin answered that it doesn’t need to be sent every time. Jonathan commented that this requires one device to report for all remote sources.

Discussion of both solutions:

Roni, Jonathan and Colin had comments along the line that both solutions would work simply different trade-offs. Colin thought there is actually a third way we could do this. Sources could delegate reporting to another named source (saying xxxx is reporting for me). Jonathan remarked that one drawback of his own current draft is that it doesn't let you identify if an RTCP report block was lost from a source timeout. Harald Alvestrand commented that Jonathan's scheme has all sources pointing to the reporter. This scheme (Qin Wu’s) has the reporting describe all the delegated sources. This is potentially a security issue if a source generates spoofed reports. Though this might not be a real issue, since if you care about security you need to use SRTP.

Magnus (as chair) stated that since all proposals work, perhaps the proponents can get together and describe the tradeoffs and perhaps generate a single proposal.

Circuit Breakers for Unicast Sessions

Authors: Colin Perkins, Varun Singh
draft-ietf-avtcore-rtp-circuit-breakers-01

The new version of the draft has added RTCP timout, several clarifications, and expanded security considerations. An open issue is to clarify that RTCP timeout doesn't occur if RTP data packets have been received, this since some devices don't do RTCP and sometimes RTCP is blocked.

Roni Even (as individual) asked how we can support endpoints that never send RTCP. Collin responded that this needs to be clarified in a revision. One possibility may be that endpoints that don't do RTCP at all simply don't work with the circuit breakers.

Bernard Aboba asked what the success criteria for a circuit breaker are. Proposes that congestive collapse should be the goal. What do the simulations have to show? Colin responded that he doesn't think preventing collapse can be proven (Bernard agrees). Colin thinks it is better that it is better than what we have today. Bernard agrees, but still has the question. Bernard followed up with an example. For instance, if I send 3 mbits on a 1 mbit link, and it stops, does that mean I am successful? Colin responded that given the time it takes to get RTCP feedback, it likely won't prevent/eliminate collapse. Cullen Jennings commented that this is just a basic circuit breaker, just intended to make things better. No need for simulations. RMCAT is working on the more complete solution. Kevin Gross commented that we need simulation to make sure there will not be false triggers. Collin stated that this especially true for circuit breaker #3.

Jonathan Lennox commented that maybe some other things should provoke the trigger (ICE packets for instance)

Presentation moved on to the Meaning of "Cease Transmission". Where it is clarified to mean - don't start unless User requests it, or you have some information that congestion has ceased. This still has OPEN ISSUEs; is explicit user action clear in the draft; do we need more guidance on automatic restart; do we need to specify the sending rate?

Kevin asked why current draft thinks 1/10th the bandwidth is equivalent to stopping.
Colin: circuit breaker is comparing to TCP at 1/10th the bandwidth. Colin agrees that clarification is needed (that particular clause is in the wrong place). Jonathan wondered if "Cease transmission to that destination" or alternatively "cease transmission to all destinations"? Colin responded that this should be to the same 5-tuple and it needs to be clarified in the document.

Eric Rescorla wanted clarification of whether javascript can restart. (in particular wants it to be clear that the browser can't be responsible for what the javascript does). Colin's view is that the user must explicitly press a button.

Comment from Christian Hones? was missed.

Moving on to behaviour with RTP/AVP Early Feedback. Where non-compound RTCP shouldn't trigger the breaker. Jonathan commented that the important thing is ignore reception reports that don't apply to your source. Colin thought that makes sense.

Behavior with RTCP XR, for consistency and ease of implementation they are ignored. Colin asked for objections? None heard.

Behavior with ECN. If ECN is negotiated and in use, treat ECN-CE as lost when calculating. But ignore ECN-CE marks for media timeout or RTCP time circuit breakers. People seem comfortable with the principle

Security Considerations:
Attackers can use fake RTCP reports to trigger the breaker. Use SRTP if that is a concern.

Next Steps: Simulations, and address open issues.

Guidelines for Choosing RTP CNAMEs

Authors: Eric Rescorla, Ali Begen
[draft-rescorla-avtcore-6222bis-00]

RFC 6222 has security defect. Solution is simply generate a suitable random number instead of generating the digest, etc in the current section 5.

The WG was asked if this was ready for adoption? Colin commented that this conceptually makes sense, we should do it. There are some textual and editorial issues to be addressed. Also need to check the CNAME values length to avoid padding (Is really 16 + n*32 are the right lengths?). Magnus takes hum. Strong support for adopting. ADs will be informed, and a milestone added.

End of First Session

Second Session Thursday afternoon Nov 8 2012

Scribe: Bill VerSteeg

Chairs do Note Well

Media Multiplexing with RTP subsessions

draft-ejzak-avtcore-rtp-subsessions-02
Richard presents the slides

Slide 4 claims that there is no provision for dealing with SSRC collisions with MSID.
Jonathan Lennox - clarifies how SSRC collisions work in RFC5576
Cullen - this is underspecified, and there probably work to be done in this context
Colin Perkins --- RFC5576 deals with this, browsers may not do this, but that is a browser problem, not a specification problem.
Richard – it is awkward, but you have to deal with it
Jonathan - RFC5576 also specifies not to use SSRCs that you have ever seen.
Richard – use cases where you would not have that information
Jonathan – yes, there are use cases that are hard
Richard – a few words of clarification would be helpful in the draft

Slide 5 – options & Alternatives
Jonathan – proposal 1 is better than proposal 2, no need to use DSCP marking – how does filter set DSCP at all?
Magnus as Individual – This is out of scope for this WG
Richard – assume the middlebox can construct the filters

Colin – main reasons for different 5-tuples is to provide QOS. Bundling takes this away. After bundling, we need different QOS, why bundle
Richard – bundling eliminates ICE for multiple ports, DTLS key exchange, etc
Colin – need to justify having one more solution within bundling
Jonathan – distinguish between flow based and packet based QOS, there are use cases for DSCP differences within a flow already
Colin – if suggestion is different DSCP within a flow, that would be OK
Magnus – lets limit that discussion here, maybe a small note on QOS
Jonathan – at rmcat, there was alarm with a single flow having multiple DSCP markings

Richard – LTE policy control typically uses 5-tuples, using SSRC is conceptually simple, but requires DPI
Hadriel Kaplan – this is DPI – what is the goal of the draft in this WG
Richard – if we can segregate the flows, this draft is less important.

Update to Recommended Codecs for the AVP RTP Profile

draft-terriberry-avp-codecs-00

Tim Presented the slides

Colin – need to look at capitalization ("may" in section 3.1)

Chairs – hum to adopt. Very strong consensus for adoption

Multiple RTP Session on a Single Lower-Layer Transport

draft-westerlund-avtcore-transport-multiplexing-04
Magnus presents the slides

Font bashing
Mathew – can you define full RTP session?
Magnus – A single complete SSRC space. Multiple full RTP sessions allows each session to re-using the same SSRC value.
Colin – we have a multiplexing architecture draft – some payloads won't work –
Cullen – many drafts that nobody implements – do we care about payloads that won't work in this context
Mathew- any time there are two that need to interoperate – back to question, is there anything else that full means?
Jonathan – can also send separate BWs in the sessions
Hadriel – are you asking for additional advantages, or something else
Hadriel – why 14 bits?
Magnus – 2 bytes are currently used
Hadriel – why not just use 32 bits? And reserve 2 bits
Magnus – 1 byte was too little, how important is space vs alignment
Hadriel – we could have used the M= line port values for the SHIM header value
Magnus – currently using an attribute line
xxx- sent a note to the list
Magnus – is there sufficient interest on the list to adopt?
Roni – would like to ask 2 things

1- Should AVTCORE work on Multiplexing multiple RTP sessions inside an RTP session

Cullen – we already decided that we will.
Roni – 2nd question – should we adopt this document to support this WG item
Mathew- there are at least two option – not multiplex, or use multiple sessions in another manner
Colin – the group has previously agreed that putting audio and video on single 5-tuple. Also agreed that we need multiple sessions on a 5-tuple.

Colin – saying that we do not need the function is not a solution to the problem
Richard – not sure that we should not ask the question again. He does not see the need to progress the draft
Jonathan – once you have the draft, you will need SDP to signal it. Bundle is easy. Signaling with MMT is impossible.
Roni – do we need a document to signal this? Normative reference?
Jonathan – does not know
Colin – everything this group does has signaling documents – we should do this as well
Cullen – we did peel the onion, we got agreement to do this as part of the agreement to do it the other way. In favor of taking the hum
Magnus – need a solution to signal this, as well as whatever the MMUSIC solution is
Harald – will hum against – wrong solution to wrong problem, will never address again
Hadriel – you do not have to implement it – far cleaner than the other bundle proposals – if MMT does not work with this, MMT has a problem
Roni – calls for a hum – first hum was not conclusive
Hum again – some consensus, but no consensus – take it to the list.
Roni – will discuss with AD.

Note from Roni: There was an agreement to have a milestone for multiple RTP sessions on a single lower layer transport. There is only this document to address the milestone and if not adopted cannot address the milestone

RTP Clock Source Signalling

[draft-ietf-avtcore-clksrc-01]
Aiden from the slides
Slide 4 –
Magnus (as individual) – traceable has no link to clock source
Aiden – traceable means that the source master is connected to a global source of time
Magnus – Swedish time exists, but is not the same as the TAI.
Aiden – the way traceable means traceable back to a high stratum server. In this draft, any source that is labeled as traceable means "traceable to something", but do not need to identify the source ip address.
Kevin - are we refereeing to TAI, do we need to add that to the draft
Aiden – OK
Slide 6 -
Colin – Is this SSRC a separate RTP stream just conveying time?
Aiden – not necessarily can be a RTP stream using the RTP time stamp or header extension.
Colin - would it make sense to define payload format for conveying time stream
Aiden – May make sense, we can talk.
Jonathan – master as a=ssrc and slave as mediaclock is it always like that the master is per source and the slave per m-line
Aiden – SSRC of master is required.
Magnus – a=ssrc is actually not required for some cases, example all streams from the mixer in slide 5 to C and D are master.
Aiden – agreed

Aiden – how many people read the draft – 4 or 5 raised hands. Feels that it is close to WGLC. Please review after the changes.
Magnus – typing up comments, needs offer/answer section, may need more than one revision, will send comments to the list

RTP and Leap Second

[draft-ietf-avtcore-leap-second-01]
Kevin from the slides
Slide 4
Magnus (as Individual) – on first bullet depends on reporting interval if you can skip reporting interval – if you expect to send one every 100ms or so, you will time out
Kevin – maybe put a zero in the NTP timestamp, which would be legal, but not well tested for receiver behavior.
Colin - sending a zero would be legal but may confuse receiver. Could also just send a receiver report rather than a sender report. That avoids a timeout

Colin – time stamps not synched during the last two seconds of the affected day – the timestamp being reported are not synchronized to anything. the clock may not be accurate, so need to say the reporting clock is the reference rather than the wall clock
Kevin – yes thank you.
Colin – do leap seconds always happen at the same time?
Kevin – scheduled potentially 4 times per year
Colin – should we say never send sender reports at these times

Kevin – if one knows about the leap second and one does not, you may have problems
Jonathan – if two ends disagree about whether there was a leap second or not, you still have problems. This will be a problem with support of UTC. No problem if people use TAI.
Kevin this is what the draft recommends.

Kevin – trying to provide recommendations, one is you must have a communications channel for when leap second happens.
Magnus – if one receiver takes the leap second and the other does not, one receiver would see a time shift.
Kevin – playout time will be off, and to late, will adjust its buffers
Jonathan – if trying to synchronize and the devices disagree on leap seconds, they will eventually resynch
Jonathan – clock discontinuities for other reasons, and devices need to be resilient
Kevin – these are all should

End of Second Session