Minutes RTCWeb WG Interim Meeting Tuesday 12 June 2012

Minute Takers: Martin Thomson, Mary Barnes

The WG chairs opened the meeting by presenting the Note Well and circulated a "Blue Sheet". The agenda was presented and discussed.

Use Cases

Draft: draft-ietf-rtcweb-use-cases-and-requirements-08
Slides:
http://www.ietf.org/proceedings/interim/2012/06/12/rtcweb/slides/slides-interim-2012-rtcweb-2-2.pdf

Stefan Håkansson led the WG through a presentation about the changes and open issues for the Use Cases draft. The recent changes were presented as listed on slide 3. A number of new requirements has been added to reflect missed requirements or agreed on use cases from the call for any changes to the use cases.

A number of use cases was proposed and not added as there was not clear consensus. These use cases was presented to the meeting for discussion and determining consensus on the interest in them.

Call Center (slide 6)

No comments and the slides conclusion that no additional requirements are raised by the use case stand.

Enterprise 1 (Slide 8)

After some discussion the conclusion was that this use case does not add any new requirements.

Enterprise 2 (Slide 9)

As the proposed solution was the same as for the previous, also this would not require any new requirements and there were no comments from the meeting participants.

Enterprise 3 (Slide 10)

Martin Thomson raised that one need to specify which port ranges to use for audio or video. Cullen Jennings commented that this is common for VoIP clients. Eric Rescorla concluded that this is a requirement on the Browser implementation. Justin raised the question how the browser becomes configured with such information. It was commented that this can be done manually or automatically like proxy configurations, such as wpad or pac.

Richard Ejzak stated that one might need to specify APN or bearers for mobile networks, not only for QoS but also to get access to the Enterprise network.

The conclusion is that also this use case doesn't require any new requirements compared to already included use cases.

Enterprise 4 (Slide 11)

No one supported adding this use case to the one supported by WebRTC.

Enterprise 5 (slide 12)

No discussion in the meeting, it was proposed that it not be included.

Low complexity central node (Slide 13-16)

There was significant amount of discussion around these slides. In the end there was not support for including any new use case from these slides. It should be noted that the use case Cullen tried to express (Slide 15) is one where key-management is done end-to-end, i.e. one that motivates the already selected DTLS-SRTP.

Priority (Slide 19)

Discussion around diffserv and if application/browser is allowed to mark its traffic sent onto the network or not. Clearly there does not yet exist a clear solution for how the browser knows that it can and how to mark the traffic. Some concerns for the size of the problem. The conclusion is to change the requirement F24 to say “The Browser SHOULD be able to use available network QoS prioritization.” There was no clear consensus on which type(s) of QoS mechanisms to be supported. There was no change to requirement F34 which was considered important by the WG.

Mobility (Slide 20)

Long discussion around a number of issues or problems for this area took place. No firm conclusions for what to do, other than it is clear this needs work and is important. No need to change the requirement or any use case. Discussion included when moving/switching interfaces occur, either as a result of a communication failure or when additional interface becomes available and could be better. How to determine if one interface is better than the other is difficult without trying it fully with intended traffic. To handle a failed interface is something that needs to be supported, any other is for discussion.

Coffee break

UDP blocking NATs (Slide 21)

The discussion yielded that people consider this important to have a competitive solution for the market. It was discussed if a solution that goes beyond simply using TCP, like TURN TCP, something looking more HTTP like. HTTPS or Websocket tunneling was of interest. It was proposed that a new requirement to handle HTTP(S) only firewalls to be added to the requirements. An individual draft for a solution would be appreciated. A question to the WG was if a conference server is considered a sufficient solution.

Other Topics

It was asked if the WG would split out the identity use cases and requirements, which was requested to be discussed with the security documents.

The other use-case related question is that there is no use-case for where the application gets statistics for the communication session. The use case document doesn’t need to be exhaustive but clearly motivating the work we are doing.

WebRTC Usage of RTP

Draft: draft-ietf-rtcweb-rtp-usage-03

Part 1: Topologies

Slides: RTP for WebRTC: Part 1 Topologies

Presenter: Magnus Westerlund

API to RTP mapping

The discussion started when Magnus reached his interpretation of how RTP maps to the API constructs with media streams and their relations (See slide 5-7). There was clear disagreement about the mapping of sources to media streams (and tracks) and SSRC and RTP. It was questioned why CNAME isn't sufficient, but also raised that the MSID (draft-alvestrand-rtcweb-msid-02) is what have so far been proposed and it is so far seen as necessary to meet the current API proposal from W3C. It was pointed out that similar binding needs exist in CLUE and those should be considered to be aligned. A lot of discussion time was spent on if creation of two MediaStreamTracks from the same source shall result in two RTP level media streams or actually optimized to a single instance if the tracks have compatible parameters.

Topologies discussion

There was no disagreement that a WebRTC end-point needs to support Point to Point, Mesh and Mixers (of different types). Colin Perkins commented that we are not defining which middlebox implementations that are to be supported, we are defining what a WebRTC compliant end-point shall support.

The Relay (Slide 16) was discussed but no support beyond the proponents for the use cases. These use cases are the ones concluded to be added. Future extensions of WebRTC can of course define what an end-point shall support.

End-Point Forwarding

This slide caused significant amount of discussion. It was concluded that there was no interest in treating A-B and B-C as part of the same RTP session. Instead this can be handled as two different RTP sessions, with possibility for B to act as any type of Mixer as previously described to provide A's media to C.

Simulcast (multiple encodings of the same source)

Slide 18 included two alternatives, there was question why two PeerConnections would be required and strong interest for accomplish this within a single PeerConnection and adding multiple instances with different encodings to that PC. Different views what making multiple adding of a media source results in, additional media streams corresponding to the same source or be optimized into a reduced set of media streams. A question is how one enables having different properties for different MediaStreamTracks based on the same media source. This is also an W3C question on how the API shall function. No conclusion beyond that more work is needed.

ACTION: Justin Uberti to provide a written IPR disclosure related to above discussion as he provided a verbal one.

RTP Mechanisms

Slides: RTP for WebRTC: Part 2 Mechanisms

Presenter Colin Perkins

Core RTP Features (Slide 6)

There was consensus that an end-point MUST support multiple SSRCs in the same RTP session and that an end-point MUST support that the same peer has multiple SSRC and uses them simultaneously.

It was consensus after discussion that an end-point MUST pick its SSRCs randomly and be capable of SSRC collision detection and resolution. The end-point MUST support signaling of SSRC using RFC 5576 but can't expect it in all cases.

When generating random SSRCs use a crypto-random one as that ensures the uniqueness properties desired. Although most SSRCs will be signaled to bind them to MediaStreamTracks there are some cases which results in non-signalled SSRCs. First case is legacy cases where all tracks goes to the default MediaStream. The second case is the SSRCs used by RTP mechanisms, like RTP retransmission that are bound to the main media carrying one by the mechanism itself. A question is if the application can influence the SSRC used. The WG expects implementers to implement RFC 3550 this only emphasis some parts which has been poorly implemented in some cases.

Consensus: Receiver MUST support RTP packets containing CSRC lists and MUST handle RTCP packets relating to CSRCs.

It is an open question how CSRC information is exposed to the application through the API. Must also the API support setting CSRC, as an application can perform a media mixing and then send that mix out.

RTCP (Slide 7)

There was consensus on Slide 7. A clarification on the requirement: MUST generate good SR packets such that a receiver could perform lip synch, NOT that the browser MUST implement lip synch on playback.

A question raised is how an WebRTC end-point interoperates with a legacy not sending RTCP. It was suggested that handling of non-conformant end-points should be noted in the specification.

Choice of RTP Profile (Slide 8-11)

There was consensus on that RTP/SAVPF MUST be used by WebRTC end-points.

Choice of RTP Payload Formats (Slide 12)

Harald Alvestrand requested that we make clear that one is not allowed to change clock rate of the payload types without changing SSRC. Colin agreed that changing clock rate is problematic and saw no problem. Cullen like to confirm that this don't cause problems with DTMF. Colin responded that this should not be an issue.

Choice of RTP Payload Formats (Slide 13)

AVP (RFC3551) has implementation of PCMU and DVI4 at SHOULD strength. Cullen Jennings proposed that DVI4 should be profiled away. Colin think we should consider updating RFC 3551. Keith Drage commented that reducing a SHOULD strength is not a profile, suggest we update RFC 3551.

RTP Session Multiplex (Slide 14)

Justin asked why a third method for MUX needs to be supported when we already have separate transport and SSRC mux. This position was supported by Cullen Jennings and Martin Thomson. Colin responded that the reason is due to the limitations of multiple media types in a single RTP session.

A consensus call between specifying implementation requirement for SHIM at SHOULD or MAY strength was tied with approx. 10 persons on each side, thus No Consensus.

RTP and RTCP Multiplexing (Slide 15)

The question was if an end-point is required to support RTP and RTCP on different ports. There was some stating that yes it should. Jonathan Lennox stated that we already lost legacy by requiring ICE. Justin Uberti asked if RTCP mux cause problems for existing end-points? Harald noted that if you are doing RTCP mux and don't like to support fallback you need new signaling. Cullen commented that a relay can trivially split RTP and RTCP for a legacy end-point. Randall supported this position and added that there is really no slowdown in doing that. Justin Uberti commented that ICE candidate gathering will take longer time and we are already failing SDP based on crypto negotiation. Lennox had arguments both ways. Synthesis of RTCP is hard, but RTCP demultiplexing is easy. The ICE cost can be avoided by simply not allocating relay resources and let the ICE checks fail.

A Consensus call was held: There was no clear consensus to only supporting RTCP multiplexing, thus the consensus is that an end-point MUST support RTP and RTCP separately.

Symmetric RTP/RTCP (Slide 16)

Clear consensus for this as proposed.

Reduced Size RTCP (Slide 17)

Consensus on requiring end-point to support RFC 5506.

Cullen commented that required to support is OK, but no reason to make a statement about supporting endpoints that don't support it. There was consensus to drop second box on Slide 17.

Generation of RTCP CNAME (Slide 18)

Eric Rescorla asked why not simply requiring pseudorandom. Justin commented that 6222 output is indistinguishable from randomness. Keith commented that there are no interoperability problems from using your own method of generating the random stuff. Cullen proposed that an end-point need to receive CNAMEs in any format, and need to send a privacy-preserving format; later we discuss whether 4.2b is the right solution. EKR commented that RFC 6222 is not preserving privacy as all the sources are guessable and thus one can track a particular end-point in different usages. Roni Even suggested that if there is a problem with RFC 6222 then the RFC should be updated to fix the issue.

Consensus that we need to come back to this issue. Eric Rescorla to express issues with RFC 6222 to AVTCORE WG.

Temporary Maximum Media Stream Bit Rate Request (Slide 21)

Consensus for the slide, i.e. MUST respect received request, MAY send request

Full Intra Request (FIR) (Slide 22)

Ted Hardi (as individual) commented that he don’t think this can be required, since this can be implemented by things other than the browser, like cameras with their own encoder. Randell Jesup commented that we can require and end-point to support it by saying that you send a full intra when you can, and there might be reasons that you can't send it. Colin added that one reason you might not send is that you don't have the bandwidth for it. It was commented that sufficient weasel words already exist in the RFC. Stephan Wenger added that this is crucial functionality for conference applications that breaks a number of use cases. Burn noted that there is a testing problem with this, how do you create a test that validates that the intra arrived as a direct result of the FIR?

Consensus was that this is required to be supported.

Semantic Loss Tolerance (1) (Slide 23)

Keith Drage suggested that PLI could be a MAY and that we simply describe the implications. Stephan Wenger thinks this is important and should be REQUIRED to minimize interop failures. Colin commented that support of this is rather trivial if you support FIR. Stephan Wenger and Randal Jesup made it clear that how you react is up to the implementation, it doesn’t have the same semantics as FIR.

Consensus to Require Support.

Semantic Loss Tolerance (2) (Slide 24)

Consensus for slides proposal of Optional support for both SLI and RPS

Controlling Codec Operation (Slide 25)

Discussion was deferred.

Other Codec Control Messages (Slide 26)

Cullen and Randal Jessup indicated interest in Temporal-Spatial trade-off request. Agreed to continue discussion of that message. No interest in H.271 back channel messages.

Header extensions (Slide 27-31)

Discussion about RFC 5285's requirement that any header extension shall be possible to ignore.

No objections to making Rapid Synchronization RECOMMNEDED (Slide 28)

Cullen Jennings commented that client-to-mixer levels RFC 6464 have some security concerns, but he believes there is no problem in some usage. Jonathan Lennox responded that the RFC requires one to use header extension encryption.

The consensus is that RFC 6464 is RECOMMNED to implement and it MUST use crypto.

Consensus that RFC 6465 is optional to support.

RFC 5450 Transmission time offset was discussed as potentially useful for congestion control. But it was agreed to leave any potential inclusion of this header extension to the decision on what should be included for that.

Consensus that RFC 5484 – SMPTE time codes are not useful in WebRTC context.

Retransmission (Slide 34)

Justin Uberti requested that RTP retransmission is made REQUIRED to support. Cullen Jennings thinks it is difficult to get the packet in interactive applications and is not particular useful. Colin commented that you can increase the jitterbuffer to make time for retransmissions. Cullen responded that he thinks concealment is better. Martin Thomson commented this is a user experience optimization prefer to see it as RECOMMENDED. Roni Even prefers FEC over retransmission. Magnus Westerlund (as individual) supports making it required and stated that he uses retransmission and it works very well, especially when you let the video fall temporary behind and then catch up, jitter in frame rate is less disturbing than packet loss artifacts and much more efficient. Randall Jessup commented that NACK has other usages than RTP retransmission, for example to determine which data is missing so that the encoder can do low cost fixes to any problems. Justin Uberti stated there is many cases where retransmission is useful and that it is more efficient and require much less data than a full intra when used with video. Google Hangouts uses retransmission for video, but not for audio. Cullen Jennings stated that Cisco has IPR on last mile NACK. Stephan Wenger has written academic paper on repair vs. restart on bandwidth limited link. For delays less than 700 ms retransmission is better restart unless you have lots of really smart stuff.

A consensus call was made between requiring retransmission and leaving it RECOMMENDED. There was no consensus so RTP retransmission continues to be RECOMMENDED to support.

Redundancy (Slide 35-36)

Some discussion on RTP payload format based redundancy. It was concluded that if supported by the selected codec(s) one can simply use it in the end-point implementation. There was also discussion around RFC 2198 and timed text. It can be noted that other options exist than Real-time text over RTP for an RTCWEB end-point, there even exist emergency service specifications (XEP-301) that does text chat in browser in javascript.

Block based FEC was discussed but it was clear consensus that this will not be included in the first version of WebRTC.

Congestion Control (Slide 38)

Martin Thomson commented that he in general is fine with hard limits, but is concerned about signaling hard limits due to application opportunity to lie. Cullen asked if this hard limit signaling is something that already exist or something that already exist. We need something more concreate on what this is, there are options. Concerning circuit-breaker it isn't done and we probably shouldn't write a blank cheque on this right now, maybe the timing isn't right to close these. Lars Eggert commented similar to pseudo wires; in their framework we simply said that you have to have limits, but you need to ensure that you reserve capacity AND have a circuit breaker; there is a precedent for doing something like this. Eric Rescorla aksed what does "signaled" mean in this context - is this the javascript indicating this? Colin responded that maybe there is some knowledge in the network, but would like to defer this discussion. Eric Rescorla stated that he is not ready to hum without more clarity. Randell Jesup commented this is setting an upper limit that may be lower than a limit that might already exist. Users might like to use something like this to cap usage for something like reserving bandwidth for other uses. Harald Alevstrand commented that circuit breakers should be included, if it is finished. Is TMMBR in this category of hard limit signaling. Does that include b=? also? Colin responded that he envision something like b=, And we've already decided to include TMMBR. Harald Alvestrand commented that the ECN should not be included if not approved. Magnus Westerlund commented that the ECN information in RTP is approved but not yet published. Harald responded but not the response to ECN indications.

No conclusions

Performance Monitoring (Slide 40)

Roni Even commented that Skype always asks about quality. Colin Perkins proposed not to say anything about this and then requests that people raise the issue separately.
Roni commented that he and some co-authors submitted a draft, but don't think it needs to be done immediately, extend later

Conclusion that the WG continue to consider if there are any XR metrics that should be supported in the end-points.