RTCWEB March 12 & 14, 2013
Chairs: Magnus Westerlund, Ted Hardie, Cullen Jennings
AD: Gonzalo Camarillo
Jabber logs: http://www.ietf.org/jabber/logs/rtcweb/2013-03-14.html
Audio day 1:http://www.ietf.org/audio/ietf86/ietf86-caribbean3-20130312-0900-am1.mp3
Audio day 2: http://www.ietf.org/audio/ietf86/ietf86-caribbean4-20130314-0900-am1.mp3
MeetEcho: http://ietf86.conf.meetecho.com/index.php/Recorded_Sessions#RTCWEB
NB: There were initial issues with the audio levels on day 2, but were resolved within the first 10 minutes.

Agenda: March 12, 2013 9:00 to 11:30
Note takers: Stephan Wenger, Keith Drage Chairs: Magnus Westerlund, Ted Hardie

Administrivia (5 min)

Slides: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-0.pdf

The chairs reminded people of the "Note well". The agenda was presented (for Tuesday and Thursday). Chairs made it clear that we will have the codec discussion on Thursday.

Reminder to people to review the minutes from the interim meeting.

Presentation of one bottle Tequila and chocolate to Robert Sparks.

AD Message (Gonzalo Camarillo) (5 min)

Need to grow the RTCWEB pie rather than split. Need interoperability. In practical terms need to make compromises. Cannot just add meeting time or people to the problem. Need to accept solutions you can live with. Identify the essential requirements.

Bernard Aboba remarked that he does not think that problem is necessarily in the RTCWEB working group. Problem is all the other working groups. Mary Barnes commented that there are MMUSIC and AVT overlap. Gonzalo concluded that we need to make decisions, and move forward.

Data Channel

- draft-jesup-rtcweb-data-protocol-04 (20 min)

Presenter: Randell Jesup
Slides: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-5.pdf

Randell presented through his slide set and then discussion began.

Richard Ejzak asked what is our bottom line? He sees the updates in the proposal since the Interim meeting as going backwards. SDP negotiation for content/protocol in a data channel stream has been removed. Think this is a mistake. Does not suggest dynamically negotiation should not be possible like inband. Need mechanism of negotiating individual protocols on data channels, not just declaring. Randell responded that one can do it anyway one wants. Richard wants to reuse SDP for this, especially to support legacy interop.

Randell commented that nothing says SDP from JSEP is limiting. SDP does not need to be on the wire anyway, not even for non-data channel stuff. Application can do anything it wants. Nothing stops you from adding SDP at this point. Richard sees a need to standardise this. Leaving it open is bad for legacy. Chair (Ted) injected: Consider Scope. SDP is control surface between JS and browser and is subject to standardisation. But there could be signalling that is independent of that. Not a requirement that should use SDP signalling to accomplish this.

Peter Thatcher commented that he thinks the presenter's proposal (draft-jesup-rtcweb-data-protocol-04) is great. (After the meeting Peter has clarified that is was the content of the slides and to be precise “I support the symmetric SID, the option for JS to choose the SID, the optional in-band OPEN message, and the fact that SDP only negotiates the existence and transport of data channels, not individual data channels”.)

Jerome Marcon commented that he is not satisfied on how this resolves negotiation. For example T.140 Characters per second. Properties can be asymmetric. Randell responded that he sent a response on the mailing list around 3/4 am last night. Simplest method to open a data channel for negotiation - transmit offer over that. Jerome replied, but then how do we know that unspecified negotiation channel is supported by remote end. Better if things sent by SDP. Chair (Magnus) asked why this is relevant? A T.140 implementation will be in JavaScript. Therefore, negotiation protocol is with the JS. Randell added that draft-jesup-rtcweb-data-protocol-04 do have a protocol parameter with this still. It is not a replacement for negotiation, but it does indicate what a stream carries.

Harald Alvestrand stated that he Agree with this (Afterwards clarified as: “I supported what draft-jesup was suggesting - negotiating usage of SCTP in SDP, and not mentioning individual data channels in the SDP.”). Glad got data channel negotiation out of the document. There are no existing implementation this has to interwork with - this is all new. Any protocol has to define how to work on top of this and map. Not RTCWEB business to define this.

Salvatore Loreto stated that he like this proposal - should negotiate only one way and should not use SDP. Regarding the question of the even/odd issue - what is the proposal? Michael Tuexen commented that DTLS has to be the active part. Jerome Marcon added that the draft does not really contain all this material now, therefore text needs to be added. Randell agreed that there need to be added.

Chairs asked how many had read the draft(draft-jesup-rtcweb-data-protocol-04), which was couple of dozen. Chairs urge people to read and report new issues to the mailing list. Request that Randell takes the open issues in separate emails to the list.

Salvatore commented that one open issue is if we are going to define prioritization between channels. Randell responded that prioritization is tricky as it interacts with congestion control. At this point we will try to define it. Martin Thomson disagreed that this should be done. It interacts but not in particular mysterious way. Michael Tuexen commented that SCTP provides strict priorities between data channels. This doesn't really interact with congestion control. The data channels takes the bandwidth they are allowed. The open question is if this is what we want, or if we also want to make clear the prioritization with the media streams. Randell responded that he has no issue of adding this to the list of open issues for the draft-jesup-rtcweb-data-protocol-04.

Richard Ejzak asked for clarification from Ted regarding his suggestion for signalling of extensions. Ted answered that if one want to run something on top of the data channel one can make standards or common specification for that particular usage in a separate document.

Salvatore asked for a poll of the room about not using SDP to negotiate channels. Chairs responded that this is premature at the moment. Will do calls at the end of the presentations.

 - draft-thomson-rtcweb-data (15 min)

Slides: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-7.pdf

Presenter Martin Thomson went through the slide set and opened up for discussion.

Jerome Marcon commented that it would be good to know what the resulting API would look like. Martin responded that he did sent something to the list. Nonbinding but fairly straight forward.

Jerome follow up that defaulting reliability and priority appears fine, but subprotocol is more problematic. For example need to know if a channel I chat or file transfer. Martin responded that those cases are interesting ones. Need to start negotiating them to make things clear. Lot of protocols do not worry about this because the subprotocol is known, either being the same or any addition will be of a particular sub protocol. If have an chat stream then substreams can be assumed to be file transfer scheme. Jerome responded that he suggested a parameter in SDP. This is a bit non-deterministic. Sometimes you don't need negotiation, sometimes you need it. Martin clarified that the negotiation is defined when you need it.

Richard Ejzak stated that he like out of band negotiation concept. He don't see a lot of applications not wanting to know in advance what a data channel will be used to. Concern as described then seems declarative without providing a means to negotiate, to accept or reject a specific usage. He would prefer if the proposal to be able to negotiate individual streams.

Paul Kyzivat commented that the proposal uses PPIDs to separate between text and binary. What happens when want to use PPID for something else. How is it known whether it is text or binary. Martin responded this is an API problem as much as anything else. This as a PPID is on a per packet basis, and that isn't exposed. But with a limitation on a data channel having a PPID. Paul responded that there is still a problem here. If want to switch on the fly. Michael Tuexen followed up that the JS commonly needs to ID if the data is binary or text, if it is anything else, register an PPID for it. Martin tried to summarize Pauls issue being that one can't learn if it is binary or text given that it is some protocol with a PPID. A possible solution can be fixed with an extension of the websockets API.

- draft-marcon-rtcweb-data-channel-management (15 min)

Slides: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-4.pdf

Presenter Jerome Marcon went through his slides and then opened up for questions and discussion.

Randell Jesup: Note about the solution proposed for the framing issue - while there is not large overhead per stream, there is an overhead of multiple bytes per stream. Overhead that is not really necessary. Jerome requested some feedback.

Michael Tuexen: Encoding something in the bits of a stream is a bad idea. Looks like a bandwidth optimisation for setting up a lot of data channels having the same characteristic. Not sure how much is really saved if number of streams is not in the order of 64K, e.g. 100 instead. Jerome stated that the key point is always use SDP to negotiate on data channels. Whether you optimise or not is secondary.

Michael Tuexen further commented that binary is a message property, not a data channel property. Jerome responded that the agent should choose whether text or binary. Michael Tuexen responded that it the JS can call send with either a binary object or a string. Jerome, this is for the moment the same as in the WebSocket API and like Randell's proposal. The config is both a channel and a message property config.

Ted Hardie (from the floor) asked a question regarding the slides listing as benefit interoperability due to SDP O/A. Who do you believe that you get interoperability among? If this is between JS instances, then we are not heading in the right direction. If it is with something else, then can you give an example of something that already has this data channel implemented? Jerome, responded that when the end-point do not know each other's capabilities then SDP is there. Ted followed up saying that SDP is available but not required. For two download JS they can do whatever they want, including GIF images of semaphore flag to negotiate. For uses outside of two downloaded JS, what interoperability are you getting? No answer from Jerome.

Harald Alverstrand stated that this proposal seems to be uniquely qualified in managing to subsume every single disadvantage of other proposals and add its own. Extensible language embedded in SDP. Inband on every message. This is a new protocol - nothing to be interoperable with at the moment. The proposal is a vastly complex overkill that gets rid of no disadvantages.

Jerome responded that draft presented by Randell still mentions use of SDP and extended syntax. Harald clarified that when new property is added, the proposal requires one to extend SDP. Confused discussion about who adds new properties and how this works. Unclear if it requires browser modification or JS modification changing SDP to add new parameters. Harald believed the confusion illustrates the issue.

Richard Ejzak respond to Ted as individual. Notion of being able to have a standalone negotiation by passing appropriate parameter to the browser allows use of SDP offer answer and this is attractive.

Bernard Aboba commented on the discussion of out of band of versus RTT for T.140. Version of RFC 4103 that would not interoperate with RFC 4103. Therefore no interoperability issue with this. No need for additional functionality to support this. You would have a gateway to translate between the RTP/UDP T.140 and the T.140 over data channel.

- Discussion (30 min)

Adam Roach stated that we seem to be putting a lot of energy into how an application talks to itself. This functionally all one application, on the server, on the client. Don't need to standardise.

Hadriel Kaplan stated that the in band stuff makes sense. We cannot prevent two JS talking to each other anyway. The reason these two guys push the SDP offer/answer, is because they would like to use this in the SIP side eventually. But, there are no solution to interwork with for the moment. Need some way to set up separate SCTP connections. It is application decision, not a browser decision.

Richard Ejzak wanted to comment on both Hadriel and Bernards statements. With respect to a single application, no need to fix these, with single server and a single application that is true. However, if we want interoperability between applications, we need some means to control what they are speaking to one another. That way is currently SDP. In regards to T.140 (Bernard's comment) even if we need a gateway for the data, we would be forced to interwork in-band with out-of-band negotiation on the two sides, this is a pain.

Jabber: Bunch of +1s to Adam's comment coming out.

Martin Thomson: Reiterated Adam's comment, however we in this WG and W3C do a lot in the name of usability to help applications. My proposal doesn't require SDP negotiation, but it allows it for the case which desire or need to use it. If you are an application that won't needed it, then don't use it.

Robert Sparks commented that getting out of the way, is a good way of characterizing things, but not fully true. We were never in the way. Richard's comment that we have inter-application negotiation in SDP offer answer all along is just not correct. We do that in signalling layer. There is another channel, such as SIP, XCON. For example SIP Require header: ice happens in the SIP, not in the SDP. Going to support offer/answer but things don't have to use it.

Chair (Ted) asks Martin if he sees us writing a separate document on using SDP offer answer that is not part of the data signalling draft? Martin responded that the main concern is things out of the way by design by default. Randell's proposal still has an open message. That he can accept if consensus is such, but his preference is for design that has "out of the way" as the default.

Adam Roach wanted to clarify his earlier statements. He is very sensitive to idea of federating applications, but in order to do that, we are defining new protocols. The current is neither necessary nor sufficient. To do inter-domain application interoperability we need to define new protocols to meet those goals.

Justin Uberti had a question to Randell. If have external input of the stream ID, do you still need to send an open message. Randell responded, No do not need an open message in that case. There to support external notification. Justin stated that this made him happy. Martin Thomson commented that this is the reason he was happy with Randell's proposal.

- Consensus call(s) (5 min)

Chairs propose to call for adoption of drafts - indicate support of one draft which will be baseline for future work.

jesup: thomson: marcon:

Strong hum for the first. Limited support for the others.

Randell's document will be the basis for further discussion. Will confirm on list.

WGLC Issue resolution (30 min)

- draft-ietf-rtcweb-security
- draft-ietf-rtcweb-security-arch

Slides: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-8.pdf

Presenter Eric Rescorla (EKR) started going through the slides with open issues raised in the last call.

Slide 6 - Perfect Forward Secrecy (Aboba)

Bernard Aboba proposed that PFS is MUST support, not MUST use. Harald Alvestrand asked EKR to state the disadvantages of stating "MUST use"? EKR answered Performance. It is a bit slower at key-exchange time. It can also be an interoperability issue. Harald thus suggested "MUST offer"? Justin Umberti asked for clarification regarding Chrome and Firefox behavior. EKR responded that both Chrome and Firefox Offer both PFS and non-PFS mode and Server decides. They do however prefer PFS modes today, so any performance issues would already be noticeable. Jonathan Lennox asked if this is a SHOULD soley on the browser, there is no plans to expose this to the JS and let it choose. EKR responded no. Bernard Aboba commented that it is unlikely to see any performance issue in a browser to browser context. It is a large scale mixer where it might be noticeable. EKR stated that the ECDH is really fast. Justin, commented that this was the case he is worried about, but appears to be measurable but not a real issue. EKR will propose some text that ensure that it doesn't become a interop issue.

Hadriel Haplan asked regarding SDES. is that going to be decided at any point? EKR responded that expected a longer discussion on this. Not in charge of this. Magnus stated that we have promised to take it up and a virtual interim just for this question appears to be suitable way forward. Hadriel and EKR indicate that is fine. Hadriel primarily wanted the issue resolved one way or another.

Slide 7 - Name of the system (Westerlund)

Confirmed that will be "WebRTC" from now on.

Slide 8 - draft-ietf-rtcweb-security-arch-06 Slide 9 - Mixed content (Westerlund, Johston)

EKR explained the issue and the evolution of mixed content in general. Within 6 month 3 out of 4 browsers will not allow mixed content in any meaningful way. Martin Thomson commented becomes moot if browser blocking source. Martin's understanding regarding WebRTC JS over HTTP is that is allowed byt no persistent permission. Giving this don't see any reason to ban mixed content, as long mixed don't get better permissions that regular HTTP.

Phil Hallam Baker: Bad idea. No security problem from sending a channel encrypted. Concerned that rather than driving up and getting whole of page encrypted, causes a drive down and encrypt none. Example if YouTube encrypts then sites linking in YouTube requires certificate. EKR responded that what Browsers do with mixed content is out of our hand. From our perspective we are not discourage anything. Hadrial clarified that if we treat mixed page as straight HTTP then there is no discouraging.

Decision: refuse persistent permissions in mixed content settings. Mixed content is entirely unsecure. Confirm on list. Copy to W3C.

Slide 10 - Linkage Issues (Westerlund)

Martin Thomson commented that there are two types of entities that can build linkage, the web server and other clients. EKR stated that he is planning to rewrite this text to make things clear.

Slide 11 - Guy on the other end (Thomson) Slide 12 - Screen Sharing (Uberti)

Martin Thomson commented that EKR promised would not talk about the long issues. EKR responded that you can treat this as an announcement. Chair (Magnus) commented that there are two sides of this. There is the security threat that we can document. The second is to provide solutions on this which we may not have agreement on. EKR wants to check if people are interested in screen sharing before spending time on this. Martin Thomson commented that what was sent to the list was an analysis of the input and that was good. Harald Alvestrand commented that he thinks this is important to document. Probably the mitigation do not belong in IETF but in W3C. EKR responded that unfortunately some things do belong in the IETF context, like SOP.

Slide 13 - How to talk about site authentication (Johnston)

Chairs asked for volunteers to send text. Martin commented that the people you really would like to have is the people that done the different trust domains and he is not that person. Ted, plead to everyone one to reach out to their contacts.

Chairs will get a date for the next version of the drafts and announce on Thursdays meeting. Chairs expect another round of last calls on this document.

- draft-ietf-rtcweb-overview

No comments in meeting.

Harald summarised email comment: No substantive contents received.

- draft-ietf-rtcweb-use-cases-and-requirements

No comments in meeting.

RTCP-XR (15 min)

- draft-ietf-rtcweb-rtp-usage-06

Presenter Colin Perkins (CSP) Slides:
http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-2.pdf

Slide 4 - Proposals

Charles Eckel wonder about what is available from the Stats API between browser and application. This includes things that learned through RTCP and RTCP-XR. Don't see any conflict between stats API and adding RTCP-XR blocks. CSP responded that if we mandate RTCP XR we also exposes information on media path. This may be a reason for mandating its use.

Roni Even wondered why the proposal only mandates around reception. Need to also send XR report from application. This if the peer expects to receive RTCP XR information. CSP responded that the current proposal is to not mandate generate of any RTCP XR blocks. Roni asked why not. CSP, which ones (RTCP XR reports). Possibly talking at cross purposes. Browser or application. Question whether browser should be mandated or just have the flexibility. So far no compelling arguments have been provided why this is mandated.

Wolfgang Beck provided a use case. Today in SIP based communication we have monitors in the network. Would like to see similar things in RTCWEB. Jitter is one type of measurement that is part of Stats API. CSP responded that W3C is the ones deciding on whats the Stats API exposes, but jitter is available in regular RTCP SR/RR and could be exposed. The question is do we need any of the extended details that RTCP XR provides?

Martin Thomson stated that he thinks the proposal slide is exactly the right answer. There is the RTCP mux issue with respect to cannot use that particular packet slot.

Justin Uberti commented that given that Stats API exist we don't really need to have it from XR. Colin stated that the only need for RTCP XR, are cases where there things in the media plane receiving the reports or where they are better than sending them over the signalling path, for example due to federation. For example if you are gate waying a WebRTC into a SIP call.

Hadriel Kaplan understands this proposal, but for clarity. When you mean non signalled packets showing up, what is meant? Colin responded that XR packet should not cause the receiver to fall over and crash. Also don't use the RTP payload when RTP/RTCP mux. Although, this is in the middle of the reserved type for RTCP packet types. The main reason for bringing this up is that some people have been requesting inclusion.

Randell Jesup are in agreement with this proposal. Randell also remarked that due to the security the RTCP XR receiving monitor would need to be a middlebox with the security context or the remote end-point.

Roni asked how do you enable browser to send monitoring data. Colin answered that a browser is welcome to implement the reporting and should signal the usage. Roni, but that does not answer how an application will build a monitoring architecture. Roni desires that JS to get the statistics and then being able to send this out as RTCP XR. Colin, but that is W3C issue and would require an API for doing that. Martin Thomson added that you proponents for RTCP XR who really want these statistics are making it hard on yourself. This is the first step. The next step is a lot of work, but don't try to burden this to get what you want. Ron Even answered that he is not sure what Martin really meant. There are use cases for being able to send the XR report to the network for monitoring purposes. Don't need to do now. Magnus Westerlund comment that so far the interest been is limited in extended reporting. What is proposed here is allowing use cases and actual metrics to be proposed in the future.

Hadriel Kaplan commented that on the WebRTC side the information is encrypted, so compared to days unencrypted SIP deployments you anyway will need to gateway the information into the monitoring system.

Qin Wu commented that the use case he see needs for support are when gatewaying a WebRTC end-point into another system that has QoS monitoring that information needs to be provided. Therefore the web browser should support these metrics, so that the gateway do not need to synthesise the information. Colin interpreted this differently and thought what was proposed was sufficient. Qin thinks the signalling was important.

Colin, added that a thing not present in the draft is on supporting transmission time extension header. This should at least be mentioned as an important robustness issue. Justin Uberti commented that he would like to see the header extension itself as a SHOULD level item. Chair (Ted) please bring the issue to the list and discuss further there.

FEC in RTCWEB (Call for interest)

- draft-mandyam-rtcweb-fecframe-00 (5 min)

Presenter: Giri Mandyam Slides:
http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-1.pdf

Slide 1 Slide 2 - Introduction Slide 3 - FEC streaming Slide 4 - FEC streaming trade-offs Slide 5 - FEC streaming example Slide 6 - What is Requested

Call for interest and please start discussions on list. Not call for adoption today.

Mobile issues for RTCWEB (Call for review)

- http://tools.ietf.org/html/draft-reddy-rtcweb-mobile-00 (5 min)

Slides: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-3.pdf
Presenter: Richard Ejzak

Slide 6 - Possible Solutions

Hannes Tschofenig commented that bearer establishment triggered by network therefore makes difficult for browser to do anything. Richard responded that he is assuming the browser has some sort of API. Hannes remarked that this API does not exist today.

Slide 7 - Proxy Mobile IPv6 - traffic offload Slide 8 - SIPTO - Selected IP Traffic Offload Slide 9 - Next Steps

Chairs: Please review and comment on list.

Second session on Thursday. 1st session. as well as related meetings.

Meeting closed 11:29.

Agenda for March 14, 2013, 9:00 a.m. to 11:30 a.m:

JSEP Update
Presentation: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-13.pdf
Justin Uberti presenting

During the discussion, Eric Rescorla raised a question on the ICE candidate gathering, suggesting that the proposal was over-specified. Martin Thomson pointed out that it had been discussed in the W3C; he argued that it did not further discussion here. Cullen Jennings asked whether the allocation of resources was occurring during the ICE candidate gathering. Justin replied that there had been no consensus on this point.

On the topic of rollback, Eric asked whether it was possible to rollback once in the stable state; Justin confirmed, and Eric asked that this be better documented. Justin agreed. Martin asked about PRANSWER; Justin noted that it had been agreed that PRANSWER was not out of scope, but the initial focus was on offer rollback.

The group then moved to the discussion of rehydration. Justin noted that one way this could be modeled was as if it were a re-invite, and that this limited the “magic” needed. Cullen said that it might not be strictly a re-invite but “invite with replaces” or similar--but he was basically positive toward the essential approach. Question on whether MSID had already been approved; Cullen noted that it had not seen a consensus call, but that it was a working group draft and it was expected that progress would be made. Eric Rescorla noted that there were other approaches to rehydration, and though the current version was not objectionable, it was not settled. Justin asked if there were specific concerns; Eric replied that he was concerned about the consequences of having one side experience a call restart and another side not, in the case of permissions grants. Martin Thomson suggested that an updated draft with new state diagram would be appropriate. He also asked how long the property of stable MSIDs would be required. Justin noted that it could arbitrarily long. Martin noted that the media stream is a transient object for grouping and that to create that stability was going to require work. Eric and Justin then discussed what the “plumbing” object is for browser-level constructs to the underlying media stream. Eric asked where the agreement was that MSIDs should be this object? Justin asked what handles are available. Martin suggested that this topic be moved to the list. Paul Kyzivat then asked what happens if signalling occurs during the middle of this rehydration? Reply is that this fails. Justin believes that this is two RTT amount of time and that if a re-invite occurs, this just failed. Giri Mandyam commented that some of these discussions belong in the W3C, perhaps in the media stream capture task force. He also asked whether a new tab was considered a new context and thus out of scope for the rehydration. Justin said that it was possible to persist across this, but that it was much more work for the application. Cullen asked whether it would be possible to move the state from one device to another. Justin it’s a little more dicey, but it could work, and it would be an interesting call handoff choices. Martin believes that it won’t work when these are being generated on the fly. Justin notes that this presumes that this is being set, not generated on the fly. Eric notes that he’d like to see more detail as the use case Cullen’s raised would not be possible in the other approaches, so he’d like to see more detail on this. Ted Hardie asked from the floor as an individual contributor how separable is this problem? Can it move out of JSEP or is it better to keep it in a single document? Justin answered that the whole MSID stuff needs to be worked out and other W3C work needs to be coordinated, but that it still belongs in JSEP. Eric said he didn't care much about what document it is-he wants to keep this idea in mind as we do the rest of the protocol. Keeping that page of spec in JSEP does that. Andy Hutton said that he did not want to park the decision.

Mandatory to Implement Video Codec discussion
VP8 Technical: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-9.pdf
Harald Alvestrand presenting

During the discussion of testing, Bo Burman asked whether the tests used x264, and Harald confirmed. He noted that the x264 tool has a tune psnr parameter and that when they ran similar tests with the parameter and the latest x264, they did not get similar results but saw result a 1% difference between VP8 and H.264. Gaelle Martin-Cocher disagreed with the averaging of chroma and luma and said that MPEG would not do so; she also said that use of rate control was not appropriate. Harald noted that for MPEG, they had agreed to run a set of tests without rate control. During the discussion of performance testing, Cullen asked from the floor as an individual contributor whether the reflected numbers were from the same parameters as those on the previous test slide. Harald replied that there was a larger set reflected here, which included the set from the previous slide. There was a comparison slide; Xavier Marjou noted the slide indicated lots of hardware support and that these were important for battery life. Harald agreed and noted thats why they were there, but that it works well without it, in part because of better uniformity on the VP8 side--variation is not an advantage here. Justin pointed out that all the realtime apps for mobile except for Facetime, which is part of a vertically integrated stack, use the software methods--the cpu draw is not significant in the face of the other power draws. Xavier noted that the advantage VP8 has from uniformity may fade as it becomes more popular. Cullen, from the floor as an individual contributor, challenged the availability, saying that there were terms he could not agree to (sharing Cisco's marketing plans and publishing them ahead of time); Harald said he had not seen that. Matt Frost then spoke to say that it was just dropped, not that it was not being available to Cisco. He then said that the hardware designs for H.264 were a lot harder to get than VP8. Bo Burman noted that it was being maintained, but also that it was frozen. Harald clarified that there have been no bug fixes that required changes to the bitstream. If there become some, they can be evaluated then. Gaelle then spoke, saying that the venue is wrong, that the MPEG is the right place to do further testing. She wants to be sure that the test are accurate, and that the test set is problematic. Mo Zanaty noted that most applications don't use the hardware implementations, but disagreed with the reasoning. The APIs are wrong--the APIs can't control the right parameters. Daryl Malas noted that it was not objective. The reply was that it was not intended to be. He asked about deployment of WebRTC vs. H.264 devices. Kalyani Bogineni commented on the implementation of VP8 in hardware for mobile--she says that it may be in the chipsets, but it is not yet available to applications. Justin noted that there are 2 billion VP8-capable implementations shipping Chrome and Firefox.

VP8 IPR: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-10.pdf
Serge Lachapelle presenting
Gaelle reiterated that more details are needed on the recent announcement and noted that there are many more patent holders for H.264 that are not in MPEG. She asked if the IETF could send letters to solicit information from them. Ted Hardie referred her to the IETF process documents on IPR and asked her to contact the area directors if she had questions on how it works. Gonzalo noted that the area directors had already spoken to her and pointed her to the same BCPs. Cullen Jennings from the floor as an individual contributor then disagreed with Serge's characterization of the difficulty of getting an H.264 license. Serge said that it matched his experience. Cullen noted that there were only four browsers and that all but Mozilla had a fully paid license for H.264 already, so the administrative overhead was zero. Serge asked if he did not want to see a fifth browser. Cullen reiterated his disagreement with the characterization. Hadriel Kaplan then asked why this discussion was happening at this meeting, since the list had been told that there would be only technical discussions. The chairs cut the line at that point. Eric first thanked Google for arranging the MPEG-LA agreement, and asked if the upcoming statement would list from whom the sublicenses derived. Serge replied that it would. Markus Isomaki noted that since there were some IPR discussion, he wanted to repeat his statement from the list that Nokia believes it has IPR on VP8 and it is preparing an IPR disclosure. Stephan Wenger noted that it was mentioned that the IPR situation is different for different profiles; yes it is, there are non-pool patents needed for certain profiles, but that the vast majority of IPR is licensable in the pool. He also disagreed with the presentation's characterization of the difficulty. Randell Jessup noted that it was not just browser, but anything that needs to work with browsers and they would not be covered by browser licenses. Cullen Jennings, from the floor as an individual contributor, said that if the working group goes with H.264, Cisco will open source an H.264 video codec that was not GPL, since they believe it is a relevant issue. Jonathan Lennox relayed from the jabber room that servers don't have H.264, asterisk can't do anything with H.264; it's not that asterisk users don't want to pay, there's simply no agreement. Tim Terriberry then spoke in support of Randell's point that it is not just browser that are at issue, but the whole ecosystem.

H.264 MTI presentation: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-12.pdf
Bo Burnam presenting
Ted Hardie from the floor, as an individual contributor, thanked Bo for the clarity for his proposal on what profile of H.264 was proposed as MTI. He then said he didn't think there was an agreement to additionally have a Recommended profile; that is, the working group would have both MUSTs and SHOULDs. It's a base part of the system that negotiation is possible, and that negotiation might well be to the profile he indicated, and that could be the case no matter what the MTI. Tim Terriberry wanted to be clear that the profile proposed is the one that is one that is worse than VP8. Bo agreed. He went on to explain that this a trade-off between implementation efficiency and compression efficiency; the MTI correct working point is designed to allow for very constrained devices. Randell Jessup commented on the interoperablity issue; most of the uses of video conferencing (not systems) may use H.264, but are not interoperable with anything.

Facetime is an example--it uses H.264, but it is not interoperable. Harald asked if the selection of test clips would be made public. Bo said that they would; Harald thanked him, saying that the selection of test clips was important. He then asked whether he understood correctly that the tests were run outside of the profile. Bo clarified that the tests were run against the recommended profile, not the mandatory to implement. Peter St. Andre indicated that he was uncomfortable with the argument from uncertainty presented. How can you ever know you have found all the patent holders after you have licensed from the known pool? Eric asked whether there were corresponding CPU numbers for the measurements given. Bo said no, but they could get them. Eric noted that it is important to know whether constrained baseline simply allows for a lower CPU utilization or it is actually worse than the other profiles. Gaelle responded to Randell saying that the Interop issue is not an issue anymore if it is made MTI. She also clarified the IPR slide by saying that the ITU database is a consistent list of known patent holder. The MPEG-LA pool is a subset. She re-iterated her view that the IETF is not the place for clarifying this. Stephan Wenger gave his belief that 5 people from the room could detect the difference between 30% difference given on the measurements; even expert viewers find it hard. He suggests that performance numbers are not the main decision criteria, at this order of magnitude. If we do decide to do a shoot-out; the meaningful way to do it is a subjective test.

H.264 comparison: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-11.pdf
Cullen Jennings presenting (as an individual)

Commenting on the comparison between the two videos, Justin Uberti noted that it was not even the same scene. Adam Roach noted that VP8 took twenty pounds off him, given that they weren't the same aspect ratio or resolution. Cullen agreed that there were very many different ways to cook videos. He went through a series of objections. Tim Terriberry if you think we should be using something other than baseline, why haven't you proposed that? Cullen what do you think the minimum resolution should be? Tim answered 360p. Cullen said, okay, but that's not what we're going to test to--we're going to test against the common user experience, not the minimum. Tim that doesn't make sense to me, but maybe have others have other impressions. Ted Hardie from the floor, as an individual contributor, then said that the point of an MTI in this working group was to avoid negotiation failure and that you would always be able to negotiation away to a different one you share with the peer. That means that if you are testing the proposed codec you have to test it for the case of negotiation failure; that means you are testing the MTI itself. Cullen disagreed and said that he thought we were picking which codec would be broadly implemented and so we had to think about what's the user experience between a or b. Ted said that this is fundamentally different from the guidelines for the group from the beginning. Henning's guideline on this was the guideline from the beginning. Tim said that Ted's point was the one he had expressed as well. Randell then said that the question before the group is whether the MTI is sufficiently good to cover the needs of negotiation failure; the other questions are the side effects of that choice to browsers, non-browsers members of the ecosystem. Keith Moore if the MTI is not good that you want to use it, its worse than negotiation failure. Justin Uberti we don't want to limit ourselves to QVGA. Do we really think people will ship high profile?

Cullen said yes, it's the same license, and there will be open source available. Justin disagreed. Eric Rescorla first said he like the shirt in the example. Then he asked Cullen to restate Cisco's commitment to provide an open source implementation. Cullen said they would open source an H.264 AVC implementation on terms that work for firefox. Eric: High profile? Cullen: Yes. Eric: I am concerned about this testing methodology because this looks so bad, where Harald's stuff looked identical. Why? Cullen went into some reasons, including different cameras. Gaelle said that we will not have an MTI. Can we agree on a Recommended at this meeting? Ted answered that it is not a chartered work item, so it will not be asked at this meeting. Hadriel Kaplan it is about failure to negotiate; it could be ascii art animation. If they just want to see something, that's what they want. Our role is to make the SDP negotiation to not fail. I suggest that if it weren't for the IPR issue, we'd pick both. That would be the best chance of not failing negotiation. Cullen noted that it is not about whether ascii art is okay; for WebRTC to succeed, I have to build stuff that is commercially viable (Brief ascii art digression). With that goal, if WebRTC has video quality and audio quality that is grossly inferior, it won't succeed. It has to be a commercially viable codec. Stephan Wenger comments that the presentation methodology would not support scientific methodology; Cullen said plus one. Cullen said that very simple testing showed him a huge difference; he encouraged others to go get more data. He, as an individual, asked if people were interested in getting more data (no response). Justin suggested that we get subjective testing for the existing clips as an next store.

In the general discussion, Kalyani Bogineni noted the support of the 3gpp and GSMA community and asked that H.264 be chosen as MTI. Justin Uberti noted that Youtube is also a user of VP8 and they found the quality of it, bit for bit, is better for a large scale video application. Eric Rescorla liked Justin's suggestion of subjective testing on existing clips was good, with the addition that re-encoding by folks who have other H.264 settings in mind. Jeremy Fuller said that there is an entrenched beauty contest here, and that there will be groups unhappy with either choice. That tends to mean either you move on to having none or two; when do we get to that? Xavier commented that the use case of Youtube was not the same as the WebRTC use case. Gaelle asked what the next step was. Next step: listen to Martin, then listen to the AD's question. Then we'll work out the next step based on what data is needed (either on an IPR basis or technical basis). Martin, I don't believe anything we do about a beauty contest will have any influence on what will get implemented; we should let the market decide.

Sense of the room query: http://www.ietf.org/proceedings/86/slides/slides-86-rtcweb-14.pdf
Robert Sparks (as a favor to the chairs)

The questions to be asked were put on the slides. If you can't answer these question, think about what you really need. Come articulate that at the mic after the questions. If you want H.264 or can live with it had on the order of 70 hands raised; if you want vp8 or can live with had on the order of 50 hands raised. New items raised at the mic lines: Dan didn't see any discussion on performance in the presence of packet loss or bandwidth constraints. Why is that not part of the selection criteria. Keith Moore, if you don't test this over cheap hotel internet, you haven't done justice. Peter St. Andre noted that we have a clear winner: it's the ascii one. Adam

Roach: the elephant in the room is the licensing. Darryl: should we ask about no MTI? Robert: the group spent a great deal of time developing that consensus. Jonathan Lennox: there may be people who have “no mti” as second best, after their preference.

Ted Hardie thanked folks for coming and said the chairs would be working with ADs on next steps. After some of the expected statements are in, it is likely we’ll come back to the working group to see what that has changed.