RTCWEB

Administrivia

Note taker: Bo Burman
Jabber scribe: Olle E. Johansson
The chairs pointed out that the WG progress depends on a few drafts in other WG (ICE bis, Trickle ICE and STUN bis) that urgently need review and solicited volunteers to review them. Eric Rescorla and Magnus Westerlund accepted to review.

DSCP Black-holing Issue

David Black (TSVWG co-chair) presented the DSCP black-hole issue with -rtcweb-transports draft that was recently discussed on the list. This issue needs to be solved and described, even though both -rtcweb-transports and the referenced draft-ietf-tsvwg-rtcweb-qos has gone through IESG review. Magnus Westerlund has suggested a solution to the list, but what should the -rtcweb-transports draft say about DSCP black-holing and the possibility to use ICE to avoid it? The WG discussed this and concluded that the issue should be described by the -rtcweb-transports draft. Ted Hardie summarized the discussion by suggesting a text formulation for a resolution that seemed acceptable to the WG: “We will treat DSCP-induced path failure parallel with other types of path failures and resolve it by using ICE restart. Note: There is a problem with multiple DSCP codepoints on a single transport, where one might be blocked and other might get through. In this case, the ICE probes, using one DSCP codepoint, may succeed while others fail. This is complex and should be warned about. A likely viable solution is ICE restart with DSCP markings turned off, but detection requires watching the multiple-DCSP-codepoint-using channels for differential failures”. If there are other proposals for resolution, please contact Harald. Cullen Jennings asked David if this solution was acceptable, but David wanted to see the text proposal. The -rtcweb-transports author Harald Alvestrand took on the action item and will work with Justin Uberti to send a text proposal to the list.

Draft status update

JSEP

A number of tracker issues were discussed and concluded on.
Codecs in re-offer [PR #269]
The proposed solution is that codecs in an SDP re-offer can only be a subset of what was negotiated, not what the remote sent (which will be strange if the re-offering part is the originally

answering one). Paul Kyzivat (remote) commented that this is an old story in SIP Offer/Answer and turned out to be a bad answer. The topic is discussed in RFC 6337. In an SDP re-offer, you should offer everything you are capable of and think is interesting. A brief discussion concluded that this would cause change of codec depending on who sends the re-offer, for example if A and B offer (foo, bar) and (bar, foo), respectively (with order reflecting preferences). So, the SDP re-offer should, as proposed in the slide, include a subset of what was negotiated before. Paul was not happy with this, but did not think it worth pursuing.

Enforce max-bundle on offer processing [PR #282]

There was a long discussion on how to enforce max-bundle; the conclusion is basically to go with the proposal on the slide (only the first m= line is accepted), per design and for symmetry reasons. A clarification to that rule was found when an offer is received with some m= lines bundled and some not; you accept all m= lines that are possible to bundle with the first m= line.

Rewrite LS handling text to indicate edge cases and that we're living with them [PR #276 & #263]

The conclusion was to go with the proposal, even if this is a simple LS group strategy. Justin Uberti commented (via Jabber) that it is always possible to let the JavaScript (app) to create new m= lines that have a new LS group if it wants to.

addIceCandidate() and ICE restart [#250]

The conclusion was that we need an unambiguous indication in ICE: “this is the last candidate” for a certain ufrag, which should (probably) be specified in JSEP and not in any ICE WG specification.

Roll back ICE restart [#250]

Decision to go with what is described on the slide.

SDP o= line increment [#239]

While the proposed solution on the slide was accepted, it was noted that the proposal on the slide goes against what was agreed at last meeting. Emil Ivov also commented that Trickle ICE for SIP does exactly what is described on the slide. Jonathan Lennox said that we should make sure this is compatible with whatever is done in trickle ICE SIP stacks.

addTrack assignment [#288]

Peter Thatcher commented that it is necessary to scope this down even more compared to what is proposed on the slide, to in some cases not allow re-using m-sections after an ICE restart. There’s still some risk that another edge case will turn up. Adam Roach suggested that a better approach might be to instead detail what “compatible” means. Justin Uberti wondered how to normatively define that m-sections are “compatible”, but supports adding stricter text that allows for better determinism. Peter said that it is always possible to use addTransceiver to avoid any ambiguity with addTrack. Ted commented that this is a substantial algorithm change. He then gave Peter and Justin the task to generate fresh text that captures Peter’s point for scoping down even more and send it to the mailing list.

draft-ietf-rtcweb-sdp

Open issue: a=rtcp usage

There are three apparently conflicting statements around use of “a=rtcp” in JSEP, BUNDLE-31 and ICE-SIP-SDP-08. This seems more appropriate to discuss in MMUSIC WG. JSEP should be consistent with whatever solution is specified by BUNDLE and ICE-SIP-SDP, but should be allowed to make it clear what the implementation should do. Bug #296 was opened in the JSEP tracker. Ted suggested to send the entire slide deck to MMUSIC WG and ask for time to present it during their session on Friday July 22nd.

Open issue: a=fingerprint

BUNDLE-31 was thought to be correct and JSEP should align with that.

Open issue: a=rtcp-mux-only

The conclusion is to update JSEP to align with BUNDLE-31. JSEP bug #297 opened.

JSEP SDP Usage vs Generic SDP Usage

Christer commented that it would be good to note that provisional answer is used in JSEP, but is not possible in generic (RFC 3264) SDP usage without specific signaling.

ICEbis vs RFC 5245 reference

Cullen suggested that this draft should reference ICE bis, trickle ICE and STUN bis rather than the predecessors, which was not challenged by the WG.

Next steps

This draft is really hard to review and the chairs are really grateful for Paul’s ongoing review of it, but it would be great if someone else was able to review it too.

Jonathan commented that it would be good if the SDP examples in this draft could be automatically extracted and used for testing. Cullen explained that this is the intent. He also clarified that the examples are currently not machine generated, mainly because none of the current implementations support all of what is in this draft and it would take much work to support all of it. If someone is able to provide machine-generated examples that could replace the hand-generated ones currently in the draft, they would happily be accepted. Justin Uberti offered via Jabber to write some JavaScript to machine-generate SDP examples. Tim Panton proposed to have a look at the possibility having ORTC generate SDP.

Summing Up

Ted stressed once more that review of RTCWEB documents, mainly the JSEP draft, is really needed if we are going to be able to take all needed documents to WGLC before IETF 97 in Seoul in November.

Alissa Cooper (ART co-AD) asked if the previously discussed, organized draft review activities are still planned. The RTCWEB chairs explained that the situation has proven to be less bad than was originally thought, but some reviews will likely will needed. We will likely need some review activity during September, maybe as virtual interim meetings or in some type of review teams, which will hopefully take us to WGLC for all the documents before IETF 97 in Seoul.