WEBSEC WG IETF 84 Alexey, Yoav, Tobias

Thanks to Ted Hardie for taking these notes.

Agenda

The chairs asked if there was any agenda bashing. There was initially none, so the agenda as given was used until after the frame-options drafts were discussed, when a new presentation from Brad Hill and Jeff Hodges on CSP was uploaded and discussed.

The chairs reviewed the status of Working Group drafts. During the discussion of the mime sniffing draft, Larry Masinter suggested that we could get rid of sniffing when there is a new version of HTTP, as browsers sniff only because of bad content types delivered by broken servers. We could eliminate this for HTTP 2.0, by saying that browsers should not sniff there.

Note that he is not argueing that a specification is not needed; it is still needed, but this may change. A speaker then noted that the W3C still needs this as a reference, so they will need a version. The chairs and speaker will look for authors.

HSTS (draft-ietf-websec-strict-transport-sec-11)

Jeff Hodges came up to talk about the gen-art review message by Ben Campbell (see http:/ /www.ietf.org/mail-archive/web/gen-art/current/msg07648.html). Minor issue: should this document update RFC 2616 or RFC 2818? Jeff believes that it does not update 2616, but RFC 2818 is a trickier question. Jeff and ekr agreed that they do not need to do an update for RFC 2818 for this. A larger effort for a RFC 2818bis would be useful, but is not part of this charter item.

After discussion, Jeff took an action item to clarify the text related to Ben's question on UA conflicts between the policy records using includeSubdomain and any records for subdomains.

The group then discussed the registry management, particularly around "mandatory to understand" extensions. The sense of the room was that the group did not want to create mandatory to understand extensions ever. If such extensions are needed, they will need a new header. Where the draft says other specifications can update this one, say that any necessary registries may be created when such an update comes around.

For the comments in section 7.2, Jeff believes that the text as-is is clear. No argument from the room. For Ben's comments on Section 8.4, Jeff agrees, but this is derived from the cited specifications, and is not needed here. The chairs agreed that this did not require new text. The review then moved to nits, for

which some text changes have already been incorporated. The group is now waiting for Ben's comments. After that, Barry will await a new version, to be posted for ballot.

draft-ietf-websec-key-pinning-02 (Ryan Sleevi)

Ryan talked about issues seen by the Chrome team. First issue was non-determinism in path building because of multiple paths to a trust anchor, which can cause false negatives. It is possible to DoS yourself by defining a policy that cannot be validated. Question about system roots (e.g. enterprise roots); is this a specification issue or an implementation issue?

Ryan wanted to compare with TACK and HPKP. Eric Rescorla (as TLS chair), said TACK has been proposed before, but it is not adopted. Paul H pointed out that TACK and HPKP were not really on topic, even though they were impacted by non-determinism. The chairs asked about rollover--the backup key, though not yet well specified, is meant to address that. In any case, pinning the current subscriber key and a backup key eliminates the problems because then it matches regardless of how the path is constructed.

Pin Time-out and Revocation were then discussed. This is very challenging. Ryan reviewed aging and activation in TACK. Would the number of backup keys be a subject of specification? If so, is one backup key sufficient? Or is this an operational decision? Eric Rescorla came up to talk about what benefits the aging and activation has in normal operation. Essentially, it means that the more visited sites are validated for longer. He's more worried about malicious pins than incompetent pinning errors. There was a suggestions that pinning will be deployed in "learn" or "inactive" mode, and then turned on when enough confidence exists. Tobias from the jabber noted that "deployment activation" sounds like the "report-only" feature you see at CSP (and which we also discussed for HSTS previously) where you can test first. Probably makes sense in the context of key pinning that you can shut yourself out. ekr believes that there needs to be something that allows operators to test pinning, but he's not sure that this is the best approach since people can be shut out. PHB notes that almost all of the complexity is based on the idea that the validation is being done in a client, which has no discretion. Other validation methods may be better places to experiment. Chairs noted that a tracker issue should be created here and further discussion on this issue continued on the list.

Ease of use is clearly an issue here. Eric Rescorla goes through a hierarchy of deployment pain, and he notes as TLS chair that he would be happy to see an extension to do this in TLS, but as an individual notes we shouldn't wait on this, but should continue to work on pinning here.

frame-options and x-frame-options (Tobias via Skype)

Question to the room about how many people are aware of CSP? 3 or 4. Jeff

notes that they have slides on CSP (Sent to the list) on the question of whether the frame-options work should roll into CSP. It's not clear to Tobias what the criteria should be for rolling into CSP (for example, would key pinning headers benefit from be rolled in to CSP?). He notes personally that he doesn't care whether he's an author on this draft--if this gets moved to resolve header bloat, he's fine. He does note that the experience of the current system seems to be going well; not sure of the impact.

Jeff then presented his slides on "UI Safety" in CSP work. Note that Transclusion is out of scope, but framing and embedding is in scope. Basic issue is due to incomplete exclusion at the client execution context. Jeff believes that the "frame options" could be handled there, because it is essentially the same context, and that there are advantages to a single conveyance. He suggests that we drop frame options and roll it into CSP in W3C.

Discussion between chairs, ADs and W3C counterparts will follow.