Jabber: xmpp:rats@jabber.ietf.org (mailto:xmpp:rats@jabber.ietf.org)?join
MeetEcho: https://www.meetecho.com/ietf110/rats (https://www.meetecho.com/ietf110/rats)
Notes <rel="self">: https://codimd.ietf.org/notes-ietf-110-rats (https://codimd.ietf.org/notes-ietf-110-rats)

Notes takers:
Wei Pan
Eric Voit
Nancy Cam-Winget
Thomas Fossati
Jessica Fitzgerald-McKay

RATS Agenda March 9th

Room 9
Time zone: UTC+1

13:00 Agenda Bash & Logistics
(5 min) Ned Smith, Nancy Cam-Winget, Kathleen Moriarty

Roman: What is the realistic deadline for covering the documents which are looking to be completed?

13:05 Status of TPM-based Network Device Remote Integrity Verification (RIV)
(5 min) Nancy Cam-Winget

Nancy: Through WGLC, but it references many documents. One of the documents is not adopted (TUDA). Should we wait and package or just send through the IESG?

Roman: I will have an answer for tomorrow's meeting.

Eric: The links to other drafts are for informational content.

Roman: That makes it easier.

13:10 Attestation Sets
(10 min) Kathleen Moriarty

Kathleen: How do we scale posture assessment better? We need to summarize, and provide something back with an EAT. It would be nice to have an easy button that auto configures based on a profile.

Jessica: How does the policy model work to show if you are conformant?

Kathleen: Still trying to figure out how to do SCAP at scale. There might be some technologies other than attestation that apply, but it makes for a useful start. The goal is to make things [processing attestation related messages] simpler and easier.

Dave: @Jessica there are Categories of things that appear in Claims (such as debug posture), and there are other things that appear in configuration.

Jessica: Need to consider how to walk back when there is a problem, and figure the source of the problem.

Guy: How does this intersect with TPR [Trusted Path Routing]? The trust vector might apply here. This could be a useful component.

Kathleen: will have to look at the draft to see if it could converge, we want to reinforce the idea that client systems and servers vs. devices can be captured better.

Eric: Trustworthiness levels were designed to be extracted from the TPR draft. They should be able to be encoded in EAT [Entity Attestation Token] defined structures, and work for both Systems and Devices. [I'm] Happy to chat offline to figure out how to progress.

Kathleen: [I] Will look at this and get back to you.

13:20 RATS Architecture update and status
(15 min) Dave Thaler

Dave: Lots and Lots of changes in response to feedback from many sources.
Key areas of addition are in Reference Values and Freshness.
3 styles of freshness: Synched clocks, nonces, handles.

Brendan: Are handles PSKs [Pre-shared Keys]?

Dave: It is a predictable value which is known to both parties, so yes - PSK.

Henk: [Handles] Can either tie to Global time, or can do locally relative time within a certain domain.

Russ: has problem with the name "handle", RFC 3650 uses it already, so need to avoid name collision.

Dave: Should we use PSK? Russ: seems fine with that (Others later think this term has issues.)

Henk: PSK is not necessarily a secret. Will work this out offline.

Someone on the chat suggested “public PSK”.

Dave: questioned whether confidentiality is needed. You don’t need it to be a secret once the epoch is known. [handles] Could be broadcast [or multicast] to many entities.

Brendan: maybe it is an unpredictable ID?

Michael: how to you validate the legitimacy of the handle? [if it is unpredictable]

Dave: Require origin validation, this ensures the source of the handle. This includes integrity and handle spoofing protections.

Michael: Looking for implementation guidance to see how this works. E.g., how to find and validate the source’s keys for the handle. Can see that timeframe synch can be non-trivial.

Eric: The handle approach will scale better than nonces on devices with lots of peers.

Dave: [yes] Nonces have state requirements whereas handles may only need to keep one or two; synchronized-clocks approach also has issues.

Brendan: trying to understand use cases for handles; handles are a form of synchronized-clocks.

Dave: Handles are less precise than synchronized-clocks.

Nancy: We need to work the handle vs. clock discussion offline.

13:35 Reference Interaction Models update and status
(5 min) Henk Birkholz

Henk: My hair looks grey on this camera :-)
There is a new implementation on a veraison Github. 20 issues from Wei are posted on draft GitHub.

13:40 Time-based Unidirectional Attestation
(10 min) Henk Birkholz

Henk: [draft authors are] Leaning towards keeping the TPM1.2 specific content vs. dropping it. [authors are] Looking to support both SNMP and YANG.

Henk: Requesting chairs consider WG adoption. Should we take the adoption question to the list?

Nancy: [I’m] Running a poll right now in the session. The vast majority of people have not read the latest version.
Need more people to read the draft before we have a poll of the WG. We’ll solicit feedback on

the mail list and work towards adoption there.

13:50 RATS Attestation Result Claims for SUIT
(10 min) Henk Birkholz

Henk: There have been various fixes. For example, SUIT claims have parallels to the TPR [Trusted Path Routing] draft notion of ‘trustworthiness vectors’.
The TPR’s trustworthiness vectors will be broken out into a new draft.
[draft authors] Might be able to differentiate attestation results claim definitions from the sets of claims that might be relevant for a system/device.

Laurence: What is the relation to EAT? How do you convey the claims? CWT claims? JWT claims?

Henk: EAT is a direct candidate for the conveyance of these claims. We need to figure out how EAT and AR [Attestation Results] relate.

Laurence: Don’t we have to nail this down to a system?

Henk: Don’t think this is necessary.

Laurence: [EAT draft authors are considereing] Putting text into EAT to state that it can be something specifying Attestation Results.

Dave: This discussion is about SUIT. Teep protocol has a normative dependency on values assigned. [I’m] Hoping to see these [TEEP and SUIT draft updates] soon.
[I] Will be talking about this more in the Teep WG. Come to that session.

14:00 CBOR Tag for Unprotected CWT Claims
(10 min) Jeremy O’Donoghue

Jeremy: Use case for UCCS [Unprotected CWT ClaimS] is to pass this [Evidence] over a channel without incremental CPU/energy requirements of more encryption and/or signing.
Comments resolved on bi-directional integrity protection, security, others.
Ask: we need a WG for this effort. Several contributors want this WG to be RATS. (other choices COSE, ACE.)

Brendan: We did this [sort of thing] in SUIT. There were people who thought this was a bad idea. Reasons: using an existing channel might not be a good match as someone who compromises the channel could also pivot to work other stacks/protocols to peer attacks. [It is good practice to limit] One key per functional purpose?

Jeremy: since this use case is about attestation evidence conveyance, perhaps there are not the same ‘pivoting security’ considerations [that the SUIT WG identified] given a reasonably designed [attestation] system?

Brendan: suggests that one key is used for just this purpose; that of signing attestation evidence.

Jeremy: this key is just for transport protocol payload protection.

Laurence: There should be less independent documents/coverage about security. This [work] should be done in RATS. I have done an implementation.

Henk: This [considerations regarding the expected use of keys] is a good reason to tie this [definition of UCCS bindings to conveyance protocols] to the RATS WG.

Roman: We need to talk about parsing the documents and the claims. If you want a narrow scope without going outside the RATS concepts/constraints.

Voting: those mostly in favor of adoption (14 rasied hand, 2 not)

Brendan: We need a threat model for this scope.

Jeremy: We could do a threat model.

14:10 YANG Data Model for Challenge-Response aka CHARRA update and status
(10 min) Eric Voit
Eric: updates on the draft.

Voting: 4 have read the latest version. Plus 8 authors.

Chair: will issue a WGLC.

Eric: Can do YANG Dr. WGLC first, then do the RATS WGLC if that is easier?

14:20 Attestation Event Stream Subscription
(5min) Eric Voit

Eric: No need to adopt draft until after Charra is complete. It is possible the scope will increase beyond network devices, which is something we are in the process of considering.

14:25 RATS Use Cases
(10 min) Meiling Chen

Meiling: Application authentication of attestation results. Also includes vTPM. Does this WG cover the contents of apps?
Also there are requirements for trusted routing?

Eric: There are drafts like TPR [Trusted Path Routing] and Michael's Use Case draft. Does this match those documents? It would be great to better understand the purpose this document is trying to achieve.

Roman: There was a decision not to adopt any use case drafts in this WG. Understanding this will help define future document scopes.

Nancy: We should take this to the mailing list to understand the purpose.

Dave: Some RATS authors believe that this is the right place to handle some of the questions?

Meiling: Will take the questions to the mailing list.

Henk: supports trying to better handle multi-TPM device attestation.

14:35 Open mic
(25 min)

Dave: Would be good to better understand the WG depedencies such as TEEP.

Michael: One thing not adopted is ADD working group [https://datatracker.ietf.org/wg/add/about/ (https://datatracker.ietf.org/wg/add/about/)]. Attestation quality and privacy of DNS server. [The RATS WG] Needs to have the ADD WG understand RATS attestation.

Nancy: Agree we might need to raise awareness. Hard part is figuring out how.

Michael: Adopt and publish is likely the best way.

RATS Agenda March 10th

Room 8
Time zone: UTC+1

15:30 Agenda Bash & Logistics
(5 min) Ned Smith, Nancy Cam-Winget, Kathleen Moriarty

Ned: If there is time, the chairs will do a milestone update at the end.

15:35 EAT Adoption status update
(5 min) Laurence Lundblade

- **Slide 2:** what goes in an EAT [Entity Attestation Token] token (colour coded)
- **Slide 3:** progress status:

    - Location nearly done.
    - HW identification nearly done.
    - Submodules and nested EATs are in pretty good shape.
    - CoSWID [Concise Software Identifier] tags supports measurement of running software, identify verifier input seems to lack sufficient definition and progress.
- **Slide 4:** re-sync EAT with RATS architecture nomenclature still needs to be done. [EAT draft authors are] Considering inluding information beyond claims, but want to do some testing/research work before committing to a solution.
- **Slide 5:** changes since IETF 109: UEID of type IMEI are 14 bytes; added a bunch of new claims (keys, HW version, intended use, boot seed), added profile section, reworked CBOR interoperability issues.
- **Slide 6:** no progress for measurements [Evidence], Attestation Results and Verifier input [Endorsements, Reference Values, Appraisal Policy for Verification].

15:40 Overview of EAT Profiles
(20 min) Laurence Lundblade

- **Slide 7-8:** Profile claim. This is key to getting to interoperable EAT. Narrow down optionality inherited from COSE and CWT. Negotiation isn't always possible, and two implementations of EAT may not interoperate because of the possiblity of variance between protection algorithms, serialisation formats, key identification conventions, claim sets, freshness and reply protection conventions.

((Dave Thaler @ Jabber)): TEEP might be the first EAT profile.

((Hannes Tschofenig @ Jabber)) points out that COSE is not meant to be used without profiling. It is just a toolbox.

- **Slide 9:** A profile narrows the definition of EAT so that interop is possible. It's a human-readable document which addresses the "checklist" in the profile section of the EAT draft. Profile docs can come from IETF, other SDOs [Standards Development Organizations], vendors, private entities, etc. Profiles will name the document (via OID or URI). A profile needs to address: serialization, protection (algorithms and parameters used for signatures and encryption/integrity protection), keys identification, required and/or prohibited claims, further constraints on claims types. It also probably needs to say something about the freshness and replay protection models that are acceptable.
- **Slide 10:** serialisation requirements regarding CBOR and JSON. [provided status update]
- **Slide 11:** token protection: Profiles should include guidance on how signing, encryption, MAC options, and what algortithms are allowed. They should be tight enough to guarantee interoperability.

- **Slide 12:** Verifier requires a verification key [credential], and usually requires an Endorsement [from an Endorser]. Profile should say how that is supposed to happen.
- **Slide 13:** claim set definition [answers the questions]: which EAT claims belong to the profiled EAT [Attester]; which additional constraints apply; which EAT claims are not allowed; and which *new* claims are needed.

Dave: do you think a protocol using a profiled EAT would specify what profile it is, or a protocol would negotiate the profile to use in an exchange?

Laurence: I don't have much to say about that.

Giri: in our case we don't negotiate: the device/environment is too constrained to allow negotiation.

Dave Thaler: Would a protocol ever need to advertise what is expected from an attester? Maybe say it is only allowed ot pass EATs using a particular profile?

Laurence: possibly.

Dave: I'll make it an issue in the TEEP spec to make the TEEP Evidence a profile of EAT according to the profiling guidelines you defined.

Laurence: maybe we can propose some text about verifier options regarding EAT profiles.

Henk: how do you want to express profiles?

Laurence: anyway you want, it's completely open-ended.

Henk: what about in the IETF context?

Laurence: I'm open to suggestions. I didn't intend to narrow it at all, but we could maybe develop some conventions around this.

Henk: seems like a subset of the profile issue you described here, at least for COSE EATs.

Thomas: you could come up with a draft template in markdown or xml2rfc with options pre-populated that one could select from.

Henk: Profile can state your options in CDDL. If you write it in English, you would be doing the same thing, but with less precision.

Thomas: some profile dimensions will not be covered by CDDL. Some English text will be necessary.

Hannes: Disgrees with Henk. Some fields, CDDL will be overkill.

((Dave Thaler @ Jabber)): +1 to Hannes.

Hannes: don’t mandate the profile to be written in CDDL. [Note: Hannes expressed a reversal of opinion later on the RATS list]

Henk: maybe not mandate, but put it side by side with the English implementation.

Hannes: COSE is full of CDDL, but not interoperable.

((Dave Thaler @ Jabber)): can’t express hash of another field in CDDL.

Ben: can’t you define a new constraint operator?

((Thomas @ Jabber)): CDDL can’t express the fact a ‘bstr’ [binary string] needs to be a pseudo-random value.

Laurence: my take away is we might need English text in a profile, but can also use CDDL to narrow [add more precision about] what [data] must be from what is described in English text.

16:00 Discussion about CoSWID in EAT
(30 min) Laurence Lundblade

- **Slide 15:** Two kinds of SW description:
  - One that is created outside the device, signed by the manufacturer, put on the device at installation time. This should be relayed to the Verifier/RP.
  - The other is akin to an on-the-fly inventory of what’s in the device, signed as part of the attestation evidence.
- **Slide 15 (cont):** Current thinking is EAT should support CoSWID. SUIT manifest and CoMiD [Concise Module Identifier tags] are other options, but having too many options will not do well for interoperability.
- **Slide 16:** Proposal: EAT must be able to carry CoSWIDs. May or may not be signed and/or encrypted. No XML SWIDs [ISO/IEC 19770-2:2015]. Sign/encrypt in COSE. Whether they provide payload evidence can be determined by examining the CoSWID. How to package CoSWID in a EAT:
  - option 1: “coswid” claim as an array.
  - option 2: leave it to the profile.
  - option 3: a single claim for CoSWID::evidence and one for CoSWID::payload.
- **Slide 16 (cont):** CoSWIDs are file system oriented, which doesn’t work for all devices. Can work around that with extensions. In phone world, lots of software is in form of applications. CoSWIDs doesn’t specifically address applications. TEEs also have notion of apps.

((Dave Thaler @ Jabber)): are all three Options allowed or is this a question to the WG to choose?

Laurence: it’s a question for the [RATS] WG.

Dave: Option 3 is the cleanest, option 2 is next, I don’t like option 1.

Henk: I can see how you want to have two different types (payload and evidence), but why don’t you do that with an array instead of submods in option 3?

Laurence: more options possible, I’m open to discussion.

Dave: if you don’t have just a list of CoSWIDs but you also want to add info about each CoSWID, option 1 is not sufficient and you need to do either 2 or 3. Three seems to be the cleanest to achieve per-CoSWID metadata embedding.

Henk: [A] small issue with [option] 3 is the possibility to make errors (*scribe note: I couldn’t follow completely what Henk said*)

[Ned’s Note: CoSWID supports using the tag for both Evidence and for Reference Values. The CosWID schema doesn’t prevent the tag from containing both or from containing Reference Values when Evidence was intended or vise versa. A profile or other specification is needed to restrict the cases that lead to Verifier undecidability]

Ned: what about CoSWID signatures?

Laurence: CoSWID::evidence is weird to sign (it’s implicitly signed by the signature on the Evidence) [Ned’s Note: there is a single signature covering everything contained in the tag. If the tag contains both Evidence and Reference Values, then it can confuse Verifiers], CoSWID::payload [aka Reference Values] is typically signed by the manufacturer [aka Reference Value Provider].

Laurence: it looks like consensus is to go with CoSWID.

Henk: *(scribe note: I missed this part)*
[Ned’s Note: I believe Henk +1-ed Laurence, but with appropriate attention being given to address the areas of ambiguity that were just discussed]

Thomas: what about the situation when you don’t have a filesystem to anchor your CoSWIDs, e.g., firmware?
[Ned’s Note: SWID/CoSWID tags use a filesystem abstraction to organize Claims - which are impilcitly defined as the digest of the files that are named by the filesystem path. The file digests are copied into the CoSWID tag then a digest of the tag is signed.]

Laurence: If that is unworkable, or there are other options beyond CoSWIDs, that is not a problem. This isn't the only possible solution. Maybe we can tweak CoSWIDs to meet your needs better?

Thomas: we have done that exercise already. Let's take conversation off-line.

16:20 Issue 98 – UEID permanence
(10 min) Giri

- **Slide 1:** FIDO developing a device onboarding spec with a dependancy on EAT.
  We are using the concept of G(lobal)UID and wanted to use UEID in lieu but we are not sure about the "permanent" semantics associated with UEIDs. This probably needs some clarification in EAT to make FIDO [Fast IDentity Online] more comfortable.

Laurence: define a different claim with "semi-permanent" semantics; there's another possibility; binding the lifetime of a UUID to an authentication session…

(*scribe note: not sure I captured what Laurence said correctly*)

[Ned's Note: Non-cryptographic identifiers such as GUIDs/UUIDs can be used as device identifiers if they're associated with an authentication credential. Some credentials are not themselves device identifiers as they could use zero-knowledge proofs or could be a group credential such as a shared key. Use of one of these alternate credential forms would allow the UUID to change to suit use case requirements such as when a device changes hands (owners) or to supply a different UUID to different service providers to prevent them from colluding on the backend, e.g. tracking transactions. For example, if a UUID context was specific to a protocol session then multiple sessions couldn't be correlated via analytics on the UUID. Some identifiers are intended to be permanent such as serial number, IMEI while others can be generated dynamically such as GUID/UUID.]

Guy: 802.1AR has manufacturer ID (IDevID) separated from ID supplied by owner (LDevID).

Giri: OK, we might revisit the decision in FIDO.

((Geoffrey Cooper @ Jabber)): Purpose of replacing the GUID is to remove ability to correlate to device in operation via factory-issued GUID.

((Ned @ Jabber)): + 1 Geoffrey. IETF will have to consider privacy implications.

Chairs wrap up

Nancy: let’s move this discussion on the mailing list.

Kathleen: milestones update discussion, there’s no time now. we’ll do that on list too.

Nancy: we are done!