SACM Virtual Interim Notes

2017-09-26

Agenda (Bashed)

1. WG Status
2. Re-charter Discussion
3. (Added) Review of in-progress working-group-adopted drafts
4. Hackathon update
5. Way forward

Summary

The main accomplishment of this meeting was getting through items 2 and 3 on the agenda. We spent a substantial amount of time going through each of the re-charter issues raised on the list, and a revised charter has since been published to the list for final review [1]. We then went through each of the in-progress work-group items, and intend to have most of them updated prior to the next virtual interim in late October. Hackathon efforts are continuing down the path of pulling the two related hackathon efforts from IETF 99 together – hackathon documentation will be updated at [2] and the project planning and progress can be followed at [3].

[1] https://mailarchive.ietf.org/arch/msg/sacm/IB_kCFizupDkiNY9lxPuxLrc0PA
[2] https://github.com/sacmwg/vulnerability-scenario/tree/master/ietf_hackathon
[3] https://github.com/sacmwg/vulnerability-scenario/projects/2

Raw Notes

Jess

Status

- Requirements doc published (good job, Nancy!)
- Hackathon planning ongoing
- Charter discussion ongoing
    - Do we include work items in charter or elsewhere?

ECP

- Do we need a line between the posture manager and the orchestrator?
    - Henk thinks we should swap out the repository with the orchestrator in the diagram
    - Kathleen thinks swapping them meets certain use cases better
    - Dave says we shouldn't prohibit the posture manager to store data in the repository
        - Can put repository and orchestrator immediately to the left of the manager
    - Charles says the orchestrator may just be a control place, data would not pass through it
    - Henk likes idea of putting orchestrator and repository on same column
    - Charles envisioning orchestrator as a band across the top of the repository, manager and evaluator
    - Jess will draw up both ways

Charter Proposal

- Updated "describe" to "identify and characterize"
- Frank's comment
    - Is SACM going to describe guidance for when to collect, and how?
        - Eric's edits addressed them
    - Why do we need to standardize evaluation?
        - Henk wants to standardize evaluation format to enable post processing
        - Dave says paragraph B implies we need a query language, or OVAL evaluation language
        - Henk asked if anyone is planning to create a language. Dave says he is.
        - Frank says his comment has been adequately addressed
    - Does scope include network infrastructure?
        - Dave says yes, as per our definition of endpoint.
        - Nancy wants to point to draft terminology for endpoint definition. Dave thinks it is best to just include the definition here. Henk agrees.
- Specific work items section
    - Karen tried to uplevel Stephen's additions to the text
        - Group edited some of Karen's proposed text regarding the NEA extension
        - Group discussed wording of the COSWID, Dave provided some edits
        - Frank doesn't see his YANG work reflected in the SACM charter draft text
        - Group discussed the ROLIE software descriptor work, Stephen says the description of this work does not preclude other work in this area. Karen made some clarifying edits.
    - Kathleen is going to take a close read and provide comments
    - Karen will post the draft to list
- Status Updates

    - ROLIE Software Descriptor
        - Stephen editing draft for clarity, update will be out in advance of October Interim meeting. Dave hopes to get some review on the next update
        - CoSWIDs- consolidating attributes descriptions in the specification, update should be available prior to next interim
        - ECP- updates based on this conversation will be done, but probably not prior to next interim
        - Information Model- no updates currently being worked. We need additional work group resources to carry it forward.
        - SWIMA- Latest draft addresses all comments. There were no normative changes Adam and Karen will decide on next steps
        - Terminology- working on a new update. May have dependencies on ECP, COSWIDS, and coming work.
- Hackathon
    - Last hackathon had two SACM related efforts
    - This year will try to combine those efforts
    - Two champions- Adam and Henk
        - Other participants may include CIS folks, Andreas, Henk and co., Dave and Stephen, put them at one table instead of two
        - Henk hopeful Juniper will join, too
- Chairs Way Forward
    - Next Virtual Interim- October 23
        - Folks with existing drafts may want to request time for that meeting
        - Same for folks with new work that has not yet been adopted by the WG

Danny

Charter

KO – DW, do you think there are more edits based on what EV provided?
DW – No.
HB – FX asked if evaluation rules need to be communicated between SACM components.
DW – Seems like evaluation rules are another form of imperative guidance. If that is the case, I am not sure what other text we would need to include above and beyond what we already have.
HB – Results can be post–processed. The transport might be agnostic to the content. We just need some way to identify it such as by content type. Actual guidance is not so important to specify, but, we need to be able to type it, identify it, and get the correct results.
DW – Different types of collection could be OVAL, query, etc.
KO – By your assessment, does (b) describe what you are going to do?
DW – Yes.
KO – FX, does that clarify things for you?
FX – Yes.

DW – I would have concern with pointing to a draft in the charter because they can change. We should just include descriptions of what would be done.
DW – Are the terms in RFC 5209 sufficient?
HB – The only term that is confusing is endpoint.
DW – Should we say it can be any type (router, laptop, server, etc.)?
HB – We need to distinguish between user devices and SACM components. I think highlighting that is enough?
DW – Different technologies are going to be needed for different devices. A network device is an endpoint.
KO – Does this address this set of comments?
FX – Yes.
KO – Discussed some changes around how SACM wants to support collection of posture information from traditional computing devices (servers, laptops, etc.) and move it up above the more detailed work descriptions.
JMF – This describes what NEA did. We are not trying to standardize the communication to collect data from an endpoint. Rather, we are standardizing a NEA extension to collect and deliver information about installed software on an endpoint.
AM – Do we just want installed software or should we collect other things?
JMF – Well, we have SWIMA for installed software. I am not sure what other people are doing or planning to work on.
HB – Huawei is working on network devices.
NCW – I am not sure that is the same thing so we should be explicit.
JMF – Yes, let's say firmware, operating systems, software, etc.
HB – Yes.
KO – Adam, is your more general question answered above?
HB – Yes.
AM – Yes, I just didn't want to be bound to only using NEA to collect installed software information and the following text calls out other things so we should be good.
SB – The new text looks good. So, I am happy if everyone else is.
DW – I wrote this original text and it was written to represent a work item for the COSWID work which is an adopted draft in the WG. This specificity is defining the work being done in COSWID. It might be useful to include some of that specificity in the item if that is the work we are doing.
AM – To play devil's advocate, what if someone decided instead of CBOR they would rather use binary JSON? If we make it too specific, will it preclude other work for being done?
JMF – I think this would be covered by the above statement.
AM – But, this is saying the work is being done to be suitable for constrained devices.
DW – I think we are aiming to identify the work we are going to do now. We can always re–charter to do additional work.
KO – What we are trying to do is create a more robust charter so that we don't have to re–charter when we want to take on new work.
JMF – I think it's good to say here is what the WG is currently working on – collection, evaluation, orchestration, etc. I don't think people would say we cannot do this in SACM because we have similar work items below.

FX – We are already doing this for network devices (collection, evaluation, etc.). Right now, I don't see anything in the charter to describe this work. Should we do it now or later?
KO – Maybe we should say we are doing the following work, but, we don't have actual work going on for these things so the language did not quite work.
JMF – I see where you are going and the challenges with going to the higher–level descriptions. I think part of the benefit, in the more specific charter, is that it's a commitment to doing certain work and making progress.
DW – How about we update the text to say "a CBOR format based on ISO/IEC 19770–2 software identification SWID tag standard".
KO – <KO made some edits>. How is it now?
DW – Good.
HB – Well, there are drafts introduced by FX to the SACM WG that indicate interest and doing "that" work apparently. Or, am I mistaken?
HB – So, that answers KO's question. There are people interested in doing the work.
KO – Now to ROLIE.
SB – This doesn't preclude anyone using YANG or something else. I think this is good because it represents work that has interest in the WG at the moment.
KO – <KO updated text to indicate ROLIE could be used to share software, configuration, etc. information>,
AM – We may want to mention that there can be other types of guidance that are not tied to ROLIE.
KO – How is the top sentence.
DW – It looks fine.
Stephen – It looks fine to me too.
***<missed some discussion about how the IETF works between KO and IM>.
KO – KM, are you still on the call?
KM – Yes.
KO – Is there anything that you see as a problem?
KM – I will read it carefully and try to make some edits and try to channel other ADs to help it there.
KO – Was there anything that you saw as red flags in the discussion?
KM – no.
DW – There are lots of work items in the charter that your peers will comment on. I think it is important to note that many of the items are drafts already being worked on and will be completed shortly.
KM – Just make sure the milestones appropriately reflect that.
KO – I think we have a solid draft. Need to check what HB was saying.
HB – We output an experimental data model as part of the SACM Hackathon. I think it is implicitly discussed in the charter. Should we make it explicit?
KO – I think if it's there, even implicitly, we should just be done with this.
HB – Ok.

Hackathon Discussion

AM – This is just a quick update. Now that we have a closer charter, we can probably just leave this for one of the upcoming hackathon sessions.
DW – Can you reiterate what is being planned for the hackathon?
AM – We are going to focus on trying to bring together the two pieces – the ECP approach and the XMPP approach. Like what the ECP draft is currently working to describe. How do we get the orchestration piece involved? How do we get them to play together, etc.?
JMF – Who is participating in the hackathon?
AM – Right now, there are two champions – HB and AM. We are expecting that we will get folks from CIS, AS, HB and his people, DW and SB. Maybe others that are involved in SACM as well.
HB – In October, I will have discussions with Juniper on how they might be able to join the effort.

Status of Work Items
KO – Now, I just wanted to step through existing work and see where the status is.

ROLIE
SB – Currently, we are working on a sizeable update to that draft to make it organizationally easier to read. The content will largely be the same. We will also have an update for software descriptor extension.
KO – Will it be ready for the October Virtual Interim Meeting?
SB – Yes.
DW – We are trying to make progress on that.

COSWID
HB – We are consolidating the attributes involved in the draft and creating a YANG module that describes what will be used.
KO – Do you have a timeline for an update?
HB – I am the bottleneck. I think I should have an update for next virtual interim?

ECP
JMF – It was updated. Now that have feedback from the WG, I should be able to make some good progress.

IM
DH – I am not currently working on this because I am not funded to do so. As a result, someone else will have to take the lead on this work. With that said, I would be happy to help answer any questions and get someone up to speed, but, I can't really drive it forward.

SWIMA
CM – I updated the draft based on conversations with NCW and AM. It covers all the latest comments. No normative changes. It mostly consists of clarifications and using different words, etc. Thanks for their feedback.
KO – What are the next steps for this?

CS – It was in WGLC. The comments that were addressed were non–normative. So, I am not sure if that means anything.
KO – I will check with AM, but, do you think all comments are resolved?
CS – Yes. I didn't agree with all comments. Like NCW's to tie back to SACM Requirements. Since SWIMA has no dependency on SACM, I don't see any reason to tie it to it. That is one example of a comment.
DW – Ultimately, when SACM completes its work a thousand years from now, this document will stand on its own. I see that as less critical to include.
KO – But, none of the things that you are talking about are normative changes. It is additional information or things like that.
CS – There were no requests for those changes so I didn't make any. Sometimes things were clarified, but, it didn't change the meaning.
KO – Does anyone think this document shouldn't proceed to the IESG.

Terminology
KO – HB, I know we talked about an update at some point. Any idea when you might be able to get to that?
HB – We did a recent update of non–critical items. Now, we will do an update with more significant changes. With the rekindling of the architecture that might influence this work, John Strassner and the ECA model, and I2NSF work. We cannot do anything with that until we see a first draft of the evaluation language. So, we will just try to bring it in alignment with ECP and COSWID.
KO – Do you have a timeline?
HB – In i2nsf, everyone wants to rely on them. Since the SACM Charter changed along the way, JL, NCW, and I wanted to wait and realign them. When the prime documents are done, then we will also have a stable terminology. But, until then, it will be difficult to get the document stable.
KO – I think that is fabulous.

Way Forward
KO – We have another Virtual Interim Meeting scheduled for October 23rd. When AM sent out the invite for this only JMF requested time. If anyone has updates, they should request time. AM, do you have anything else?
AM – I don't have anything else. But for agenda items, we wanted to focus on open issues rather than status updates. However, we could also use this time to introduce new work to the WG or work that is being done but not adopted by the WG.
KO – Any other business from the floor? <No>. Thanks to our notetaker and for the productive charter discussion.
AM – Thanks KO for walking through that. It was not easy.

Charles (covering ECP for Jess)

ECP
JFM – As I was updating draft, I hit some problems. Most comments were requesting “upleveling” to describe the process of SACM collection. Seems reasonable. Challenge: we have a document describing

using NEA for collection. Recently, there has also been interest in YANG push. So – how would we use these collections? Currently, the ECP shows an endpoint with a Posture Collector/Posture Client/Comm Client and matching Server roles.
HB – Always implied that the thing that collects does the validation?
JFM – Not in the ECP. In ECP, the thing that collects (the server) just stores it in a repository for later evaluation.
IM – So diagram is not right because the validator doesn't Validate?
JFM – Correct. Was based on TNC. Our point is to separate collection from validation.
IM – Just syntax validation and storage.
JFM – Right. Seems reasonable.
JFM – That looks nothing at all like the NETCONF/YANG architecture. In NETCONF the device you are collecting from (Server) has data that gathers data and sends to a client. There is less going on in the architecture and no details about what the client does. Also some nomenclature issues (Server-endpoint).
HB – That is why we initially created "consumer-provider" to avoid this issue.
JFM – But the way we use "consumer-provider"... where does the endpoint fit. We have never pinned down. I would argue the endpoint is the provider. But not clear in architecture.
HB – It does depend. You could say that the role of a YANG provider is different.
JFM – Proposed ECP diagram: Not married to anything here. Starting from the right hand side, the endpoint has something doing collection. (I call them Posture Collectors.) There is some engine that is communicating this data to someone who cares. I call this the Posture Manager. The Manager manages collection and might do some level of validation. (Sanity check). The posture manager than dumps to a repository (like the hackathon). Maybe the repository is co-located, but could be separate. From the repository, it can be made available to the rest of the SACM domain. Evaluator or orchestrator (XMPP-Grid?).
JFM – So those are my initial thoughts. What are yours?
AM – Is there a communication between the orchestrator and the posture manager? They are using the rep as a communication mechanism? If you want to instruct the endpoint what to collect or how frequently. Who does that. Orchestrator or Posture Manager?
HB – Maybe by changing the titles Orchestrator and Repository. The Posture manager pushes to an orchestrator which can dump to a repository?
JFM – Endpoint wouldn't point to a repository. Posture manager has to be next. So you are saying direct communication between orchestrator and Posture Manager.
HB – Yes. Exchange names and then the posture manager is orchestrated as to who it pushes data to.
JFM – I could see a use case for posture managers and orchestrator to talk, but I'm not sure all communications go through. By having repo in between, you can avoid an orchestrator. Not all sites have an orchestrator. But I see the value in a direct communication. I need to hear from someone know knows CMPP-Grid.
KM – In SDN, your orchestrator would be your security controller. They would be swapped in that scenario.
IM – Swap. The orchestrator can still be optional. (Implicit routing to the repository if absent.)
DW – What I am hearding Jess ask, if the posture manager speaks directly with orchestrator (and there are use cases for that) we want to make sure we are not prohibiting the posture manager from storing directly to a repository. So have the orchestrator and respository to the immediate left.
HB – Reasonable.
CMS – What about orchestrator in control plane but not data.
HB – Tell a stream to go to various points or route itself. Orchestrator is not required to multiplex all data. It just directs the data stream to the repository. If you are omitting the orchestrator, then

everything is statically configured. (Possible, but doesn't scale.) Yes, could put repo and orchestrator in same column. Would indicate
CMS – Orchestrator above.
HB – Agree. Could be the data router (not ideal) or could facilitate direct links between. Having both the repository and orchestrator as options for direct communications makes sense.
JFM – I'll try drawing both ways and will share with the group. Devil will be in the details.