SACM VIM Minutes (January 11, 2017)

Agenda Bashing

WG Status

- Confirm people are happy moving the Vulnerability Assessment Scenario to the wiki instead of RFC <confirmed>
- Requirements
    - Sounds like there was agreement on resolution. Need to review new text.
    - Making progress. Don't see any blocking issues.
- Enumerate interfaces/functions/requirements
    - This is related to the Vulnerability Assessment Scenario follow-up work and will be discussed later in the meeting
- Are people happy with Slack?
    - Crickets
    - Will see how it goes. It does not hurt anything having it around.

Draft Disposition

- Overcome by events

Proposed work (Vulnerability Assessment Scenario)

- Background – try to deal with architecture and IM challenges. (Grid vs. NEA; scope of IM)
- Objectives: tasks, interactions, information elements, then exercise using CVEs
- Approach (1)
    - DW – validate using CVEs: agree that it would help with IE and ensuring they are complete, but how do we figure out if our identification of tasks is appropriate?
    - DH – Harder. There may be a gap.
    - AM – Doesn't VAS lay out tasks?
    - DH – Yes, at a high level. We got some feedback that it aligned.
    - DW – Adam – Agree the VAS is a good baseline.
    - DH – I'll update so that we get feedback from vendors too.
- Approach (2)
    - DW – At what point can we publish something? Milestones that were a completed deliverable by the WG? Each phase? Other?
    - DH – Through each phase, the architecture/IM will continually be revised. Once interactions are mapped out – interfaces and date, you could probably start taking those solutions and work towards that. The phase 3 optional interactions shouldn't represent a huge change over phase 2. So as soon as phase 2 is done we could start picking things out.
    - DW – One thing we have discussed: operations piece might go in an architecture. Could publish that quick is separate from the IM. If we approached what we published incrementally, then there could be a deliverable at the end of each phase. Let us work on solutions aligned with each phase sooner rather than a long: term informational document.

    - JFM – Validating using CVEs. Understand why good conceptually? Are we using CVE or any vulnerability announcements?
    - DH – More a desire for real world examples. E.g., some security advisories that are focused on installed software and versions, others have configuration aspects. Just want some examples to make sure it works.
    - HB – We are creating more optional operations later (like discovery). That is ok. We should see that the solutions created in phase 1 are not contradicting or biasing phase 4. I'm not sure how we anticipate. Maybe we need multiple solution drafts. This approach does not inhibit. Don't want to commit too early.
    - DW – To that point, we could also focus on ensuring we have necessary modularity and extensibility. Doing this iteratively. But VAS itself is also an iterative approach to the SACM space. Always a risk that an early decision will be challenging later, but can mitigate by being modular/providing for extensions. Just work smart; don't be paralyzed.
- Deliverables (Architecture, IM, examples using CVEs)
- Next steps – we need more feedback. Needs more than just a couple people.
    - DH – Anyone interested and willing to help drive forward?
    - AM – Would be very helpful. I think it is worthwhile. I'm willing to contribute somehow.
    - DW – Phase 1 work was one of the things I was hoping to tackle with the architecture work in the new year anyway. I would be happy to contribute.
    - HB – The output of this work will be some concrete operations. Those will need to be added to documents. I can contribute to identifying.
    - DW – If we start writing, would others review?
    - SB – I would be happy to review.
    - JL – I will review too
    - DH – Will follow up on list.

SWIMA

- Status
    - CS – Latest draft covers all of the issues in GitHub, but, will focus on session #4 and #1. If there is time, we will discuss the other issues which are simpler.
- Support user/vendor-defined data models (verify PEN sizes)
    - CS – Previously, there was a one-byte field to express data models from an IANA registry. There was a concern that you would be restricted to data models in the IANA registry. The updated version introduces a 3-byte Private Enterprise Number (PEN), in addition to the 1-byte data model type field, that allows vendors to specify their PEN to control their proprietary data models or data models that have not yet been standardized. Furthermore, end-user organizations can use the PEN values between 192 – 255 to specify their own data models. A graphical representation of the new fields was shown (there is now a Data Model Type PEN and a Data Model Type field). We now have the ability to have core data models, non-core data models, and enterprise-specific data models.
    - DW – Are we guaranteed that the assignment of PENs won't exceed 3-bytes?
    - CS – The PEN is a 3-byte standard as far as I know.
    - DW – I was just looking at the PENs and couldn't find where it was specified.
    - CS – I don't know for certain, but, I am confident it should be 3-bytes. Others will have the same problem since I copied it from them. I will take this as an action item to verify.

- Data Source Identification
    - CS – Original specification did not support this capability so I will go through it here. The proposed solution currently in the draft uses a 7-bit integer and a 1-bit source first use flag. When a source is first used, the SW-PC assigns a source identification number to it and when information from a source is first reported, the source use flag is set in all attribute records from that source. This allows a SW-PV to determine whether or not two records came from the same source. However, the source identification number does not tell the SW-PV what that source is. A graphical representation of the new fields was shown.
    - CS – I also wanted to go through the source first use flag. The SW-PV could exhaust the maximum number of sources (256), but, it is unlikely. We could also get rid of the flag and the SW-PV could maintain the mapping between the two and could still correlate between the source.
    - CS – A second approach is there could be a Source Information exchange (request/response). The SW-PV would send a Source Information Request and a SW-PC could send a corresponding Source Information Response. This would use a non-standardized string that would allow for descriptive text about each Source Identification Number. It should be that while the string isn't standardized, it should keep the door open for future extension.
    - CS – Do people have a preference on which approach to take?
    - DW – What is the preference that you are asking for?
    - CS – The first is to not use the source use flag. The second is to use the approach currently defined in the specification.
    - DW – I am in favor of dropping the source use flag and binding to a session rather than the lifetime of the SW-PC. That way, if the SW-PV needs more information on the source it can get it.
    - DW – Taking a smaller bite, you could also have the sources bound to a given session as well as give index numbering, but, it would not provide a way to request more information about a source.
    - CS – Can you clarify what you mean by lifetime of a source and by session?
    - DW – When the SW-PV first becomes aware of the SW-PC and exchanges with the SW-PC over its lifetime versus a given set of exchanges with the SW-PC that may be part of the same client-server session. That is what I would be referring to that session scope.
    - CS – When I was thinking about a session, I was thinking about its operational lifetime. When the SW-PC boots until it ceases operation for whatever reason.
    - DW – You may also want to consider the epoch in which the posture is being reported. We already have a type of a session with the epoch and set a collection task and allow the baseline of knowledge to be reset.
    - CS – What do you mean by epoch? It's not defined in this specification. It is defined in IF-M Segmentation.
    - DW – If you subscribe to information in events, there is an epoch identifier.
    - CS – In the event enumeration, there are epochs of event identifiers. However, those epochs can last longer than the lifetime of a posture collector. A single epoch can cover multiple lifetimes where a SW-PC starts up, shuts down, starts up, shuts down.
    - DW – Depending on how the SW-PC is working and determines the source depending on the SW-PC. Say there is a software update to the software providing the source. Some may want to call it a different source whereas others may call it the same source. There

is no benefit to re-optimizing sources. I think it is unlikely that it will go to 256, however over time if there is a drift of sources, that may exceed 256. We also need to preserve the minimal bandwidth. That's why I thought this would be a happy medium because you wouldn't have the penalty. May want to reserve the 0 – 255 threshold.

- CS – With respect to that, it has to be a detectable event that the SW-PC can determine. Also, the fact that sessions are not long lived (not very long lived). The chances of exhausting the sessions is minimal if binding on a session basis. The downside is that if we are doing source information exchanges, every session requires the SW-PV to regather the meanings of those source identification numbers if it cares about the sources of things. In many cases, it won't. Otherwise, it should be a short lifetime.
- DW – If you could get metadata about source it would help de-conflict and be a way of getting to that type of metadata.
- CS – Here's an additional challenge. Events are recorded and persist for an entire epoch. A single event, the binding between source identifier and source, only lasts a session, there would have to be some trickery for the SW-PV to make sure that it persists across multiple sessions.
- DW – That's not so much a spec problem rather an implementation detail. The tool will have to manage that mapping and assign an appropriate index number for the source and pass that number on.
- CS – It's not so much a problem but needs to be pointed out or people will trip on it.
- DW – That's my frame of mind on this.
- CS – To summarize, the group is still in favor of 256 identification numbers. For SW-PV, we want sessions and it is up to the SW-PC to make sure the source identifiers are correlated across an epoch.
- DW – Does anyone have concerns with this recommended approach?
- CS – What are people's use cases for knowing these sources and does this meet them?
- KO – Maybe this is a question for the mailing list since we are quiet?
- NCW – I am quiet because I came in late.
- HB – We have this idea of data provenance. This could be the first point in the chain of evidence. That is one use case.
- DW – We also have the management use case. If you are collecting this inventory information, you will want to say that I want to make a change to a software load on a device. Knowing the source would help you reach out to the device to make the change.
- CS – Maybe.
- DW – In the absence of having any identifiable information about where software installed, this is two bits of information you would need to know to along with location.
- CS – I disagree. This has nothing to do with management of software. Sometimes it may be part of it, but, not really.
- DW – Location is an absolute. We need some way to determine relevancy for the observer.
- JFM – I am trying to understand how that will help me take an action as you described. I might want to update, patch, install, etc. Most of these things are controlled by the operating system. I just need to know where to send the command.
- DW – Only in some cases is it covered by the operating system. In edge cases, we have software installed on app servers, installed to data sources, etc. There is software that manages each of those installation contexts etc. If we don't have the ability to discriminate between sources and installation contexts then we don't have the ability to use this information to manage those installation contexts. It's the root of this

conversation. It's one thing to know the software installed. It's another to manage it. I would rather not create another standard for this.
    - AM – How much of this is going to be determined when we go through the Vulnerability Assessment Scenario? Will we learn more about when these contexts really matter? Do we really want to go into the weeds here? Or, can we wait until we learn more based on that exercise?
    - DW – How might it change how we understand the problem?
    - AM – We might go through the exercise, the report might have information on that context and may help us better understand what we need to address here. Did that make sense?
    - CS – Makes sense, but, I don't have an answer.
    - DW – Would we find out we need less?
    - AM – It's not really about more or less. We have a goal with the group. We should focus on what we need to do to get there.
    - DW – We need this as a hook for provenance information. How far do you want to take this capability? I don't need to build a prototype in order to address this problem. It's how far do we want to take the capability? I think that's a fact.
    - AM – I don't disagree. I just wonder if we can wait to go through the scenario.
    - CS – I don't disagree. The relationship between the entity that reports the software will be true in some cases and knowing which cases is useful sometimes. With regard to provenance, how much can we trust this information?
    - DW – In the solution, we are saying we need provenance, but, we are not saying how to represent the provenance. And, the trust issue is orthogonal to this. Same with software inventory, it is orthogonal.
    - KM – I think DW is right we can do this another way. With JOSE, encryption can be done to apply it to a data model.
    - DW – Right, that is a problem we can solve down the road. We do not have to solve it today.
    - KM – That's why I brought up JOSE.
    - AM – Agree with that.
    - KO – Looking at the time, we still have to go through the Architecture and IM.
    - CS – Yes, this is a good stopping point. Things should go to the mailing list. The other slides are more for back-up. In short, read the spec and provide feedback.

Architecture

- Status
    - DW – JS and I do not have an update.
    - KO – NCW can you provide an update from the meeting at IETF 97?
    - NCW – We had a few side meetings. Jim took the action to do a next cut based on some clarifications of what we were trying to achieve in the original architecture as well as how to reconcile those with the JS/DW architecture draft. Jim took all of the notes. I don't have any notes.
    - KO – Adam and I will reach out to JS to get a summary.
    - DH – I don't remember all the details, but, it was a starting point for the operations and interfaces that would be used in the architecture.
    - NCW – There were no solid conclusions. We were just going to start with the mapping to continue the discussion.

Information Model

- Status
- #8 Define a provenance information model
    - KM – Figuring out how to do this JOSE for all data models is a good solution.
    - HB – Will be lots of metadata available to infer provenance. Don't have to tie all together; just move on. It is there.
    - DH – The way the IM is written now; people can add new IE as they wish. If they really care about it, they can add.
- #9 Network topology and location information as identifying attributes
    - HB – Most important thing: needs to be collectable automatically.
    - DH – I think we can get rid of this issue provided we document.
- #11 IP spoofing
    - DW – This goes back to same question about software inventory. We'll never have confidence in information. Just have some knowledge about where it is coming from. If we include IPs, we cannot claim they will always be correct, but we can provide hook about where we get the information, and then protect integrity of info like Kathleen mentioned.
    - DH – So we don't need, can close out.
    - KO – Need to be mentioned in security and privacy considerations. From a consideration of security issues, they need to be called out.
    - DH – I don't think we need to mention IP specifically. Malware could be faking just about anything. (Could spoof OS information.) Should talk about spoofing in more general terms.
    - KO – There is generality and then there are exemplars. IP is the latter.
- #15 – Identification and definition of attributes
    - No objections to closing out
- #17 Identifying attributes
    - No objection to closing out
- #18 How known
    - No objection to closing out
- #25 SACM components MUST have time synchronization
    - No objection to moving to architecture
    - DW – Sounds like a best practice, not a requirement. The thing that is an IM problem is if there is a drift in time then correlation isn't possible
    - DH – Main thing: want to make sure the IM has components to track time. We have.
    - KO – It feels like best practices to me. I do lots of time synchronization work in my other job. There are rat holes here.
- KO – Is there a way to drive non-contentious issues that have been around for a while to the mailing list? Problem is that mailing list is pretty quiet. We could close some of these without taking VIM time. Maybe try going to the mailing list first.
- DH – Agree.
- KO – Need to balance between utterly non-contentious vs. rat-hole fodder

Wrap up

- Virtual interim sometime in Feb? Week of the 13th is hard for some people. Any major show-stoppers.
    - NCW – Week of Cisco Live.
    - Week of 6th is TCG meeting
    - Week of 13th is RSA
    - Conclusion: week of the 13th
- Expected topics
    - AM - Vulnerability Scenario session exercise – I think a lot hinges on that.
    - DW – Interest in the phase 1-4 approach. Maybe do some phase 1 engineering.
    - AM – Would like to get a lot of work done in between.
    - DW – Would like more time about work and less time about slides.
    - KO – Maybe more dynamic discussion of architecture issues. What about 2 shorter meetings focused on specific work items. Design team kind of work.
    - AM – That could work. Is anyone else interested.
    - JFM – Good idea.
    - DW – Good idea. Do we want to plan that now?
    - KO – I think we can come away thinking we will meet and send that out to the list. AM and I could come up with those ideas. Does that work?
    - DW: Yeah.
    - AM – Works for me.