SACM Notes

SACM had two sessions scheduled during IETF 96, with the first on Monday and the second on Friday.

In-room consensus calls (Friday and Monday combined):

- With respect to SWID M&A, the following consensus calls were taken in the room:
    - Can we live with both the exchange of identifers and descriptive metadata?
        - Result: YES.
    - Should we define the information model for software descriptions before we choose a data model?
        - Result: YES.
    - Should we use the ISO SWID tag elements as a starting point?
        - Result: YES.
- With respect to versioning Information Elements
    - Hum was too close to call
    - Chairs intervened to defer the decision on IE versioning until a later time
- With respect to Requirements draft
    - Should we open the requirements draft beyond the quick rewording of DM-001?
        - Result: NO

These initial consensus polls will be confirmed on the list.

While we started without open issues on our requirements draft, Danny Haynes did bring up two issues regarding DM-001 and DM-002. We determined that we would open the requirements draft for additional changes only to DM-001 (with several people volunteering to provide Nancy with candidate text), and that we will keep the text of DM-002 intact.

Resolution to open issues with respect to the Information Model were not readily had during the Monday session, but were worked on during the week.

Dave Waltermire suggested that we start using one session at face-to-face meetings, and many in the room agreed—no one seemed to object. It is likely that we will move to a single session per meeting approach, while maintaining our usual virtual interims.

On Friday, the chairs deferred the need to determine an Information Element versioning scheme, and the room hummed in favor of *not* reopening the requirements draft any further.

The next virtual interims will be sometime the week of September 12 and the week of October 10 (Doodle polls will be established for each).

ACTIONS

- Update to SWID M&A draft prior to September VI
- Update to requirements draft prior to September VI
- Update Requirement (DM-001) and submit to IESG after IETF 96 (within a week or two)
- Complete WGLC on Vulnerability Scenario after IETF 96 (within a week or two)
- Update charter for dates (extend for 6 months)
- Rough Roadmap developed by September VI
- Concise Software Identifier discussion on the list

RAW NOTES:

MONDAY
Dave Waltermire

IETF 96 - SACM WG - Monday, July 18, 2016 10:00-12:00

Note Takers:
Michael Jones
David Waltermire
Jabber Scribe: Chris Ignacio

1. Logistics, note takers - charis - 5 minutes

2. WG status - chairs - 10 minutes

Nancy - No open issues on requirements
Danny - I have two issues on the requirements. Will discuss during another presentation.

3. Software Identification Draft Open Issues - Charles/Gunnar - 60 minutes

The unique identifiers in SWID M&A relate to a SWID document
Lots of discussion on identifier formats, but what about descriptive formats for software? Which do we want?

Henk Birkholz: It doesn't matter if the data is by reference or by inclusion. By including there can be scalability issues. You should be able to choose which to do.
Charles: Should an identifer reference always be resolvable?
Henk: You shouldn't be forced to resolve, but yes.
Charles: The same unique identifier must be used consistantly for the same software
Henk: A locally generated SWID could have a different identifier.
Charles: The goal of SWID M&A is to send to the server a list of installed applications.

CONSENNSUS CALL: Can we live with both the exchange of identifers and descriptive metadata? CONSENSUS: Yes.

Charles: Does anyone have any concerns with use of the ISO SWID specification?
Dan Romascanu: The background and use cases are not clear.
David Waltermire: The NISTIR 8060 provides this info.

CONSENSUS CALL: Should we define the information model for software descriptions before we choose a data model? CONSENSUS: yes
CONSENSUS CALL: Should we use the ISO SWID tag elements as a starting point? CONSENSUS: yes

Charles will provide an updated draft before the next virtual interim.

4. Information Model Open Issues - Danny - 20 minutes

There was a long sidebar discussion on how versioning of information elements (IE) are handled in IPFIX.

David: For issue #2, use vendor id instead of vendor name to avoid naming collisions. (Danny agreed)

Chris/David/Joe Saloway: Discussion about using a pair of IEs for version and the version format (as an enumeration of format identifiers). Agreement at the MIC about this approach.

Henk/David/Danny/Nancy: We need to work out what aspects of the information model need to be addressed in a given data model to perform some SACM operation
Jim: The SACM data model must fully implement the SACM IM
Nancy: We agreed early in this WG that we would have multiple data models. We need a single information model to validate the data models.
Jim: Isn't necessary to not only state what software is present, but also what endpoint it is on?
Nancy: We don't need a full chain of context.

Chairs: We will work to clarify the DM-001 requirement only.
Henk: When we claify, the result must not be to force the SWID model to implement the full SACM information model.

DM-002: Discussed updating the text as suggested, but consensus was to keep the text as-is.

6. Terminology draft update - Henk - 20 minutes

Dave/Henk/Adam: Discused the need for discovery capabilities. Dave suggested that we need to complete other work before we do anything significant here. Agreement to write some text on the capabilities we have defined already.

5. Endpoint Information Data Format - Danny - 20 minutes

7. Work plan for the week - chairs/all - 15 minutes

Mike Jones

SACM 10am Monday 18-Jul-16, IETF 96 Berlin

There are no open issues for the Requirements draft
- It is ready for submission
- We will undergo a privacy review concurrently

We will probably put the Vulnerability draft into WGLC soon

Charles Schmidt talked about data models draft
- SWID M&A -01 - draft-coffin-sacm-nea-swid-patnc-01
- A lot of people appeared to be proposing identifier formats rather than descriptive information
- Asked the WG how want to report our software inventories
    - Descriptions or identifiers?
- Henk Birkholtz said that the outcome is immaterial to the problem
    - It should be easy to switch between both or even mix them
- Charles: Possible when unique identifiers can always be resolved to descriptive information
- Henk: For some use cases you only need identifiers
- Charles: A SWID goal is to send a list of installed applications
- Jim Schaad: Has a slight problem with the statement that globally unique identifiers always need to be resolvable
- John Straussner: You need metadata
    - He has an example that shows how to attach metadata to any object
    - You can query for characteristics you need
    - That also means you need a registry
- Dave Waltermire: Talked about use of evidence
- Charles responded to Kathleen Moriarty that he is trying to determine the working group consensus on the topic
- He needs consensus on what the metadata looks like
- Adam Montville: This needs to be able to resolved
- Kathleen Moriarty: Said that she hopes that the chairs have a plan to resolve this issue
    - It seems like it's at a higher level than just this document
- Charles: I didn't hear anyone saying that the ability to convey metadata is unnecessary
    - Some are saying that you don't need to report descriptive information every time
    - But in other circumstances, you need the metadata
    - You want to be able to control which you get when
- Chris from CMU asked for a 3-way hum
- Karen O'Donoghue asked Charles for a clear question

- Charles: When inventory is reported do we?
    - Always use unique identifiers

Always use descriptive information
Allow both

Kathleen: Does it need to be identified in advance?
I don't want the WG to arguing about SWID identifiers for another year

Henk: It matters whether the endpoint creates the identifier or whether the vendor does
Focus on what the SWID expresses

Kathleen: Do we need to decide among the three choices in advance?

Bob Moskowitz: Do we need to publish information about capabilities in advance?

Dave: Could we ask the question "Could everyone live with both?"?

Charles: I think we have consensus ...

Kathleen: No, we had a few people at the microphone. Let the chairs do their job.

Hum: Can everyone live with both?
Unanimous for both!

Charles: There should be one data model
Possibly use ISO SWID 2015
Only 4 people in the room have looked at this
Getting a copy of the document requires payment

Charles asked if any of the 4 people had concerns about this specification

Charles said that NIST published a profile on how to use it

A lot of the specification is concerned with things like document lifecycle, which we don't need
We are only interested in the data representation

Dave Waltermire: Was on the ISO committee that developed the specification
The document doesn't provide any use cases
That's why the NIST document including use cases was created
The NIST document is freely available

Charles: SACM Information Model doc is going down the road of rolling their own data format

Charles: Asked if anyone knows of another metadata format they'd like we to consider
No one responded

Charles: Asked the chairs to determine whether there's consensus for our own format or the SWID standard?

Dave: NIST created a CPE spec with 11 strings identifying attributes of software
He will send this spec to the list. The number is 8085. It's still a draft.

Jim Schaad: Are we rushing to prematurely choose a data model
Karen: "Heavily rushing" hardly seems to apply to anything we've been doing

Karen: Kathleen is interested in us making a decision

Three options:
Use ISO SWID 2015 schema
Define our own
Defer

Dave: Differentiated between data model and information model

Jim: Prefer to determine the information model and then decide the data model
There may be things we would need to add to particular data models

Hum: Is there consensus to use ISO SWID 2015
Unanimous for yes
Hum: Is there consensus that we need to determine the information model and then decide the data model?
Unanimous for yes

Charles: Will make the specification more data model independent

Danny Haynes talked about IPFIX issues
Danny asked how/whether we want to use versions and/or flags
Dave Waltermire: Asked what the version number would be used for
Chris Inacio, CMU: Revisions can be made to the IANA registry based on expert review
John Strossner: It's good to be able to track when something has been deprecated
Danny asked about IPFIX naming conventions
He will send a note to the list asking the question
Dave asked about use of the vendor ID
Asked about datatypes
Discussion ensued about version strings
Nancy Cam-Winget: There will be contexts in which software inventories are used
... There needs to be guidance in the information model on how to use that
Jim Schaad: Wants one data model - not different ones in different contexts
Nancy: SWID only describes software. SACM wants to also include hardware descriptions.
... She thought that when the working group was started that there would be more than one data model
Jim: The data models need to be complete in the information model
Jim: A SACM data model needs to deal with the complete SACM information model
Karen: Danny asked if needed revise the requirements because of the work happening with the information model
People would like to see DM-001 be clarified
Karen: I would like a concise description of how to fix this specific problem
Karen: Nancy has an action to clarify this one statement
John, Dave, and Henk are willing to contribute text on the topic
Danny presented about DM-002
Asked whether we wanted to break this into two requirements
Karen: It seems that we don't need this change
Danny: That's fine as long as we need to maintain the relationships in the information model

Henk Birkholz talked about SACM Terminology
SACM and I2NSF are collaborating on terminology
There are Capability terms
There are Component terms
There are Task terms
Hank asked how we want to proceed
Dave Waltermire: We have a lot of work to do in SACM before finalizing the terminology work

Adam Montville: As we are going along, we'll discover what else we need
Dave: It's not appropriate to define a discovery mechanism now
Karen O'Donoghue: What is the timeline for this document?
Henk: This document is intended to finish at the same time as the core documents

Karen O'Donoghue: We will move the final endpoint discussion to Friday
She encouraged people to make progress this week
Danny: We will make progress on the information model this week

Dave Waltermire: Suggested we stop having the Friday meetings
The Friday meetings encourage people to have side meetings during the week, conflicting with participating in other working groups
Karen and Adam would also prefer one session a week
No one spoke up for multiple meetings
Karen: The virtual interims are a good mechanism for helping make progress

John Strassner

SACM Monday 7/18

Chairs Slides
IM, terminology, Software ID drafts updated
Vulnerability I-D will go into last call soon
Requirements Draft will be discussed

SWID Draft
Want to pin down issues regarding data model.

Do people want the endpoint to report a unique ID, or descriptive info,
or something else?
Henk talked about the difference between payload and evidence attributes.
John talked about using metadata and software versions.
Evidence information tags SHOULD have the same tags when generated.
Therefore, locally generated tags MUST be resolved by the resolving entity
Henk talks about how the information is created.
John agreed, and said it's about context. Don't overload the function
of an ID with what the content carried by an attribute is!
Consensus: live with both IDs and metadata.
Discussion of descriptive structures.
ISO SWID 2015 vs XORCISM vs Building our own - consensus is use
ISO SWID 2015 in info model, then develop superset of all possible
data in the data model.

Info Model update
Attributes and versioning - John mentioned using metadata and semver.
Chris Inacio said that IPFIX uses a different structure; John disagreed
and John and Chris will work offline to resolve/understand.
Naming conventions - no controversy.
Datatypes - uses IPFIX datatype. Note that time zone is not specified
as part of dateTime, but the time zone is established at time of connection
Versioning is important, so attributes need to be able to say whether
this version has a vulnerability, or this vulnerability occurred before

or after this version.
DM-001: needs more specific wording; John and Henk volunteered to help write text on this.
John and Henk posited that you further need to distinguish between a DM for data plane and a DM for control or management plane. John worried about exposing management information to endpoints that do not need to know such information.

Terminology
Started collaboration with I2NSF terminology
- imported AAA definitions from SACM
- based Controller, Control Plane, I2NSF Capability definitions from SACM
- based Capability, SACM Component, SACM Interface, SACM Role based on definitions from I2NSF

Capability
- discoverability
- endpoint management capability vs vulnerability management capability

Management is much more than a data store (change to SACM)
Name of a Capability should reflect on what it does
- Capabilities can be registered and discovered in a SACM domain
- IANA registration of a list of capabilities seems viable

Component imported from I2NSF
- can have software and hardware components
- need to define chains of components, and the tasks they perform

Selecting a Data Format for an Endpoint Information Data Model

FRIDAY

Charles Schmidt note taking...

Note well, Agenda (no changes)

Information Model
ıny Haynes – A while back when we submitted SWID M&A we submitted OVAL. Hadn't done much since, but maybe want to pick that up. At Buenos Aires I proposed a list of data models we needed out of that OVAL work. I didn't get much feedback so I just started with representing endpoint information. OVAL is based on XML and some people showed support and others wanted JSON/CBOR.
– Agenda
– Considerations – have to look at the SACM requirements. OVAL is community built and serves its purpose there. What does SACM need – slightly different requirements.
– Requirements consideration – starting to go through requirements, especially data model requirements. What requirements do we care about?
– CBOR – It is compact. Pros: compact; CDDL allows structure definition. People have complained about the size of OVAL XML. Cons – newer – not as tested. Fewer

implementations. Not human readable.
– JSON – In the middle: more compact than XML, bigger than CBOR. Pros: lots of adoption. There is a JSON schema but just an expired I-D.
– XML – You know it. Robust, lots of tools (XPath, XQuery). Cons – verbose and can become complex. But easiest since OVAL already is expressed in XML.
nleen Moriarty – You have a large install base with XML. Something to understand – how would that user base be effected. Would changing increase user base and/or impact usability (so it may be a different OVAL anyway). Are current users really stuck in XML? Is it locked in?
– All current implementations use OVAL XML and have done so for a while.
– About 50?
– About. If we want XML, we can – just want to get opinions. There is a clear understanding in OVAL community that what we do will not be compatible. But if we want them to move to this, we need their input.
– CBOR seems to be taking off more than I realized. Whatever the group decides is fine. I just want to make sure we consider all points. If they are going to be incompatible, may not matter. JSON is great. CBOR may let us move into IoT.
.m Montville – I think because this is breaking backwards compatibility, there is an opening to consider changes. Not sure I have a favorite and not sure it matters. Not sure how many OVAL implementations are directly processing OVAL vs converting to an internal format. May just mean new translators. With JSON, I think that we go more into the realm of convention. The reason so many use it is that there is no schema – more loose. Might not be specific enough for our needs. I have no favorite.
– I don't know who is using OVAL natively vs. translating. I know of both practices.
– To keep going, you can convert between CBOR and JSON since CBOR builds on JSON. Currently just guidelines.
ık Berkholtz – CBOR is a superset of JSON. If you start with CBOR you might lose something if you convert to JSON, but the other way is very easy. The general superstructures are all the same. If you are using both, you can create a data definition that covers both. There is a way to retain all features of a complex type in XML. Can incorporate a namespace, etc. into CBOR. Then rebuild XML. If you are forced to do it, it can be done – retain all XML capabilities in JSON or CBOR. Better to refactor the structure and use that.
– Clarifying question: Talking about serialization. Are we talking about deciding on one? Last time we talked about IM -> DM -> serialization. I can see some working better in different places.
– Understand all three and figure out what makes the most sense. We have to pick something, right? Which one do we want to start with? But could do all.
nan D. – Could start with an information model.
– Actually one layer below. Could have multiple bindings between data model and serialization. We are not necessarily constrained to just one serialization.
– I agree, but SACM does have to select MTI.
– So no information?
– No – saying something different.
Schaad - <Jessica> - Have you determined which parts of our large information model this applies to?
– Just configuration information. Just what you use OVAL for.
– JSON Objects. One interesting thing we found: it wasn't always clear when creating objects if the order matters and how to handle two members with the same name. The IETF created

RFC 7493 – said the order doesn't matter and don't duplicate names.
– That impacts what we do – in the information model we have ordered lists. Have to deal with that if we use JSON. Could use arrays and put each information element in there. Makes the format more complex, though.
– On XML – DTDs don't support open content – hard to extend. Use schema instead.
– When designing your schema, have to decide where you want this extensibility. Have to be careful.
– Maybe investigate YANG. Maybe convert between.
– I think we investigated that. YANG is heavily used. YANG is closely intertwined in that it does both data at rest and in motion. There are some restrictions in YANG – no circular references. Can be hard to handle XML. There are some hidden restrictions. One should be very cautious.
– Yes. Something we started but not finished – looking at requirements and mapping. Eventually, want to pick something and develop a data model. Maybe do some prototyping and see what things look like and what people like.
- One thing you didn't talk about but need to include in the evaluation is the security representations for the various serialization. COSI, HOSE, XML digital encryption. Include that. The things you can express in those is not the same.

rmation Model Update
– Quick update from Monday. IPFIX syntax issue #1 – template properties. We had a proposal to add a new structure property. I didn't hear any strong objections to not including it. It is really just for certain data types – makes life easier. Jim is writing some code to process/auto-generate information elements. Previously, these constructs were jammed in the description with prose text or embedded in data type property. Breaking out makes sense. Also a suggestion to update the status field. Add more concepts. There were concerns about overloading one field. Was a suggestion to use other fields. There was a request for more detailed versioning. I think it is more of an operational issue once we have a registry, so was thinking about not dealing with this now.
– I think the new structure – don't want to invent too many semantics into one thing. I prefer to have one label mean one thing rather than many.
– Agree with regard to the second. I think by having the structure will help with that as well. Do we need to deal with issue 2 now?
– Paraphrase: is versioning scheme needed to decide now? "Do we need to describe the versioning of information elements now? <close to a tie> Hand vote: decide now <4>; defer <7>
– Not definitive – too close. Chairs?
– Chairs say defer.
– On naming conventions – Jerome didn't really like option 4 so we'll go forward and keep things as they are. Might be beneficial to use Private Enterprise Number values – more unique and consistent.
– Data types – I had created different version data types as primitives. Going to convert to information elements. There was a request for an enumeration primitive data type. Seems reasonable to express lists. Also received a proposal for a "map" datatype – need to figure that out. Regarding data type semantics – I heard "yes and no". It isn't clear to me why we definitely need it. Looking at the list with identifier quantity, etc. – not super useful.
is Nacio – Agree, but having the underlying ability to specify... it isn't for building the IM that

we agree on. It is for allowing extension of DM later without reconvening. Having some of those underlying types allows people to build extensions and have them interoperate without IANA updates etc.
– By not including the now I'm not saying never to include. I just didn't want to add if we didn't have a good idea of use.
– Fair enough. Not clear to me that having a list and pointer is not sufficient. If you wanted a list of IP and netmask, it is about how you associate those in encodings. That can be tricky.
– I would strongly vote for at least quantity. The identifier argument is valid – work around in the IM. The DM need them. Semantics like choices are interesting – for example, deciding if you want a set of attributes or identifiers. You could, for example, have an attribute set, have why. You cannot express this if you don't have quantities and choices. My argument: quantity – yes! If you leave them out now, you have to add them later.
- Adding them later says we add them to the template. That is zero work. Chris – I heard you say you were using data types to establish serializations in data models. I would like to keep the relationships in the IM itself. If we are doing this mostly for serialization, I don't see the purpose.
– We need them in the IM so we can express that in the IM. How that gets mapped to DM and serialization, that has to be determined. But need a common IM expression – needs to be thought of there.
- I don't know if I agree with you. I need examples.
– I'm happy to let this continue on the list.
– IM/DM requirements. We talked about the first DM requirement – every DM must implement whole IM. We decided no – So SWID (partial) can be used in SACM. We also brought up this point – IM is a combination of two different DM – information we get off the endpoint (SWID, OVAL, etc.) and then the control etc. (metadata, guidance on interaction, etc.). I think we need to talk about this break up between control plane and data plane a bit more. Just thinking about this I don't know if it has an impact on requirements.
– At the moment DM-014 says that every IM element needs to be in the DM, which is true for a DM which glues everything together. For other DM such as SWID, it can be included. So two levels of DM – the DM that we leverage that don't cover the whole IM, and then there has to be a more simple subset of things you want to talk about as this glue DM. This should probably, as worded, map all IM items into this glue DM. Is this the way we want to see things – 2 levels. Will only take tiny tweaks to requirements to reflect that.
– It may be useful for us to look at them with respect to that. I put DM 14 up here since we brought it up. We should discuss this more on the list.
– The requirements are fine for the most part. Currently they seem a bit conflicting without requirements.
ɐn O – Do you want to open that door?
– No. I don't think a DM needs to satisfy all of them? If there is 100% match, great, but we can live with incomplete. This isn't a problem with requirements. Charles said individually they make sense, but in combination they are potentially problematic. I'm fine moving forward – I just want to note that there might be an overload of the term DM.
– I would like to not open the requirements can of worms. There is no such thing as perfect. Done is better than perfect. If there is a problem we fix it later?
- Also if we define multiple layers of data model we can get around this. I don't want to block this, but wanted to highlight.
– Hum if you think we need to delve into this at this point of time? <silent> Against

<HUMMM!!!!!!!>
– Part of the rewrite of DM 1 – broken into 3 parts. The first part says that if you are going to create a DM for SACM it must implement something from the IM. Second part is that a DM can contain things not in the IM. Covers non-standard extensions. Last part is an example describing that. Should we bring to the list?
– How close is that discussion to being done?
– Close.
– On Monday we said we would tweak to clarify. If you are close, send to the list. No wishy-washy language. “If you object strongly, provide alternative text”.
– Jon, Danny, and me made 4 iterations. Everyone says it is good enough. Note that this conflicts DM 14 still.
– For DM 2, we decided we could live with that. To capture that I will send this out to the list so it is archived.
– DM 14 – first part talks about the need for an attribute dictionary. I wasn’t really sure how the first part of this could be done – seems subjective. Get rid of? Second part: DM must include all attributes of the IM. This conflicts with DM1. Maybe simplify just to say the DM should be extensible given the framework in the IM. Can leave if people want, but suggest fixing second part.
– Other updates – I pulled section 8 out of the IM. The usage scenario doesn’t really belong in the IM. Shows how you are going to use the data to achieve an objective. There are some information elements covered in there – can pull those into the IM as necessary, but don’t really need the usage scenario.
– Next steps – same as Monday. When we get back, we’ll fix the IPFIX syntax issues that are resolved and work the rest on the list. Once done we can take elements from Nancy/Henk draft and pull those in. Guidance is still a pretty big gap – we want to define more clearly.

ıcise Software Identifiers
– This is a draft I am writing with Charles, Jessica, and Dave W. It is a mapping of the existing ISO SWID standard.
- We are encoding SWIDs in CBOR, defining a CDDF, and signing with COSE. We have a lot of new security considerations. One highlight – we can provide security considerations in other things: vulnerability assessment, golden measurement, etc. Also have a description of attributes of software identifiers. We will elaborate on the XML description where needed. Better udnerstand the consequences of using these documents. It already maps lots of the SACM IM items – we can elaborate on that in this document if we want.
– Questions?
– Who has read? <?> We need more people to read and review and provide discussion.
– I’ll push the pointers to the list and ask for a review.

ıerability Assessment Scenario
– We said at IETF 95 we would try to get to WGLC
– I think there is only one issue open if we pull in Server Discover. There is one more on the level of detail of the information needs, but that is an IM issue – might not need to hold up the draft.
– We’ll put it through.

/ forward

– The IM merge is not done yet?
– Yes. Just a conversion thing once the syntax is more fixed.
– Timeline?
– Yes. Prior to September
– Thinking of having a September and October VIM. I would like to shoot for week of Sept 12 and Oct 10. Will do Doodle polls.
- Object! The week of Labor Day is better.
– Update to SWID M&A before Sept?
S – Yes.
– Update to requirement and send to IESG before sept?
– Maybe next week?
– Same for Vulnerability last call?
– We need to just do the small update to the charter about the dates – we are going to extend the charter 6 mo.
– It would be helpful for everyone if we had a roadmap of next steps for documents.
– We can put something together.
– Doesn't need to be formal, but there are lots of bits. Doesn't need to be set in stone but good to know where we are headed.
– To move forward, I would ask for decisions on how registration/authorization/authentication is going to work. This is a gap at the moment. I wouldn't know how to build a SACM component because I don't know how registration works. This is required to move forward. If we have this and can propagate capabilities, this would help implementers.
– Required to move forward from this point?
– No, but it has been open so long... someone needs to work on that.
– Agree we need to discuss.
– It is impossible to build a proof of concept.
– I agree it would be bad at some point in the future. Not sure it is bad now. Still trying to figure out basic things now. Finish the IM.
– My impression was the architecture was stable because we parked it. Maybe my impression was wrong.
– Duly noted. We cannot finish work without doing those, but maybe not critical for now.
– I would be curious what people think are the important next steps. What should we focus on?
- I think there are lots of things we are doing. We need to make sure we are doing them in the right order. We are making decisions we may need to backtrack later. The number one thing we should be doing is defining the information model. Once we have that, DM flow from that and serializations flow from that. Then you can figure out how to ask questions of or publish into DM. Until the IM is there, we don't have anything to talk about.
– Yes, and in the process of doing the IM, operations take... no? Ok. I'll be unconcerned.
· <Jess> We need to address the architecture. It was parked and we need to figure out how the architecture meets our needed solutions.
– For the VIM, would be nice to have a roadmap. As a group we need to figure out how to make faster, more consistent progress. I don't have the answers. I would like to see a rough roadmap and some thoughts on how to execute that roadmap. Maybe properly prep that conversation.
– I'm glad to hear you bring that up. It is a positive step. This week was a positive step – there were some long going conversations that we had decision. There were perceived

senses of consensus, but without the chairs there is no way to measure. That type of thing will help a lot.