SACM Virtual Interim Notes – 2015-06-29

Note takers:

- Gunnar Engelbach
- David Waltermire
- Lisa Lorenzin

Agenda

1. Logistics, note takers - chairs - 5 min
2. Requirements draft - Nancy/Lisa/chairs - 60 minutes
3. Endpoint ID design update and Information Model discussion - Dave/Cliff/Danny - 20 minutes
4. (NOT COVERED) draft-fitzgeraldmckay-sacm-endpointcompliance-00 - Jess - 15 minutes
5. (NOT COVERED) draft-hansbury-sacm-oval-info-model-mapping-00 - Matt/Danny - 15 minutes
6. Next steps / way forward - chairs - 5 minutes

A Note on Notes

This collection of notes is an adaptation from several sources, using the form of one of those sources as a primary guide. Full notes provided by Lisa Lorenzin (a near transcription really) are provided further below. Browsing the actions, and the summary notes followed by looking for bold-faced items in the more complete notes could be a good strategy to get caught up quickly.

Action Items

Dan Romascanu – Ask Linda Dunbar about status of her review
Dan Romascanu – Provide examples of various uses of the term transport
WG – Review specific T-XX requirements to recommend if any should be moved to OP-XX
Requirements Editors – Clarify difference between collection and query operations
Dave Waltermire – Dig up endpoint ID outline and send it to the list
Henk Birkholz – Send something to the list pertaining to the relationship between IP addresses and MAC addresses

Summary Notes

Requirements

- WGLC on Requirements is designed to drive comments on the draft. The draft will not progress unless there is consensus on the draft from both the authors and the working group.

Open Issues – 1 (slide)

- With regards to non-reputation, there is a need to describe requirements for data provenance/origination (e.g., source authentication, message integrity). Need to point to authoritative sources to avoid the need to create new (requirement/terminology) text. Use definitions in RFC4949. Kathleen will assist if needed.
- Need to drive a conversation about what we mean with regards to "transport requirements". Where are these transport requirements addressed? Should the requirements apply to SACM sessions/communications? Define transport as "protocols used to move data between SACM components." Discuss the semantics of various layers in a protocol stack that need to be addressed.
- Dan will provide examples of various uses of the term transport.
- WG will review specific T-XXX requirements to recommend if any should be moved to OP-XXX.
- Editors - With regards to Issue #42, clarify the difference between a collection and a query operation. What components are involved in each operation? What is the origin of information (e.g., provider entity performing collection)? What is the object (e.g., target endpoint)? Use target endpoint instead of object (consensus: no objections).
- Should all attribute queries be defined as a finite list? A wildcard should be allowed, which is resolved into a finite list of attribute values. Both querying for a specific list and all attributes must be supported.
- Issue #37: T-001 – Agreement on providing clarification.

Open Issues – 2

- In section 2.6 should be in sync with agreement on terminology for transport above.
- Issue #34: Section 2.2 – Consensus on one information model. The single information model may point to other information models, but no alternatives are to be provided. Consensus that the information model may be modular to allow for easier updates.
- Issue #32: DM-010 – Consensus that attributes must be extensible. The IM provides a starting list of attributes that may be extended. Extensibility must be handled in a way that each data model can pickup extensions to the IM. Data models must also be able to introduce new attributes. Define a type model that all data models can utilize. Address potential collisions in attributes. Allow clients to pass-through attributes that they don't understand. This is BOTH a requirements and an IM issue. IM defines the types, data models define the encoding.

Open Issues – 3

- DM-006 – See previous discussion around origin vs. endpoint for solution.
- Issue #16: G-009 is discovery of schema. G-010 is discovery of target endpoints. Also need to allow for discovery of what components have what information, a third type.
- Issue #13: Skip this until we address the data origination issues from above.

Open Issues – 4

- Issue #9: G-003 – Transports need to be scalable above and beyond the lower level transport (e.g., TCP/IP). There is a need to define what needs to be scalable and how. A transport needs to be scalable enough for the intended use of the transport. A transport model must have a section defining scalability constraints. Must require scalability considerations and allow for exemptions to be made.

Architecture

There was a brief discussion about when the architecture draft might be updated next, as comments for it from January have not yet been addressed. An attempt to address and update before IETF 93 submission cutoff was pledged.

Endpoint ID Design Update

Dave Waltermire walked through the endpoint ID design team update. The team has been:

- Defining a list of common identifying attributes
- Developing a catalog of these various ID types
- Next steps is to get this work into the information model

Lisa Lorenzin pointed out that most recently it seems that the endpoint ID effort was working to identify a requirement to gather relationships between these attributes. Some discussion followed her observation, and the notion of using triples to represent such relationships came up. There is still an open question as to whether we move forward with the triple (i.e. graph) way of thinking or if we approach the problem using another way (i.e. hierarchical). Dave Waltermire asked a good question: To what point do we need to define this representation in our information model vs. the data model?

Also discussed:

- We may need a new design team (or a "recharter" of this endpoint ID design team) to work on characterizing endpoint posture or classification of endpoints (i.e. in support of targeting endpoints for assessment)
- Scope of identifier collection. How many attributes should be collected in support of identification? As many as possible? What about privacy considerations? Are there performance concerns?
- Provenance. Everyone on the call seemed to agree that some form of "provenance" should be included; likely in the form of providing identifying information for the target as well as the SACM component doing the collection.

Way Forward

We noted that our use cases draft was "stuck" but this was likely due to Stephen being out on holiday for a while. The goal was to get this updated before the IETF 93 cutoff.

- Update the information model
- Update the requirements draft; when ready, we'll issue second WGLC
- Work on critical path drafts (terminology -> requirements -> architecture -> information model)
- IETF 93 Focus
    - Requirements
    - Architecture
    - Information model
    - Terminology
    - OVAL Assessment
    - NEA Assessment

More Complete Notes

150629 SACM interim F2F

Notes taken by Lisa Lorenzin

Agenda:

7. Logistics, note takers - chairs - 5 min
8. Requirements draft - Nancy/Lisa/chairs - 60 minutes
9. Endpoint ID design update and Information Model discussion - Dave/Cliff/Danny - 20 minutes
10. draft-fitzgeraldmckay-sacm-endpointcompliance-00 - Jess - 15 minutes
11. draft-hansbury-sacm-oval-info-model-mapping-00 - Matt/Danny - 15 minutes
12. Next steps / way forward - chairs - 5 minutes

Dan takes AI to engage Linda Dunbar about her NEA review

Requirements

- LL - this draft was moved into WGLC while several participants were not comfortable, including me
    - Thought that required consensus - what is the process there?
    - DR - not sure it needed consensus from WG, needed from authors
    - Sent a query to authors
    - Purpose of WGLC is not to send the document, but to generate activity
    - Certainly would not send it forward until we close the consensus from open issues
    - LL - didn't receive that query - please check to make sure you're using my Pulse Secure address and not my old Juniper address
    - Thanks for the explanation - I understand the process now
- #46 - Section 5 - non-repudiation
    - We do desire to find the origin
    - Ensuring non-repudiation is harder than we want to tackle
    - Do we want to explicitly state that it's out of scope?
    - If so, need to define what we're placing out of scope
    - JS - don't know that you need to say anything about non-repudiation
    - Term I don't care for, would like to see the term deleted
    - KM - data provenance and origin
    - DW - issues of source authentication of messages
    - KM - that's more origin
    - **LL - could delete that mention to address this issue**
    - **[No objection]**
- #13 - Terms like integrity, origin of data, confidentiality, authentication
    - DR - integrity in SACM is meant as defined here by reference
    - LL - add to terminology draft, refer out to authoritative source
    - DR - don't want to spend the time to redefine when already defined
    - HB - more easy targets, but for example origin of data and data provenance can have subtle differences
    - Can be resolved in terminology draft
    - Highlight here - data provenance might include chain of custody of data

    - Origin might be the origin of data where created, or where observed for the first time - ambiguity there
    - Handle in terminology draft, provide a few examples for the next update
    - **KM - I'd start with RFC 4949, see if that's good enough**
    - If not, being discussed in many WGs, silly to re-do this work and have different definitions from rest of IETF
    - Let me know if you need help tracking this down
    - **HB - will review and contact you if needed**
    - CK - question about integrity and confidentiality - data in motion, at rest, or both
    - That would tell us what kind of definitions we're shopping for
    - For data in motion, pretty easy to come by, sorts of things you'll find in IETF sources
    - Integrity and confidentiality at rest, not so much
    - LL - data in motion - defining a transport protocol for data, not specifying how to store data
    - **JS - probably need to talk about both of them - data in motion from provider creating data to consumer, need to be integrity protected at rest on repository in the middle**
    - **LL - good point - maybe address them as sub-definitions of same term**
    - DW - worth considering what we address in MUST/SHOULD vs. security considerations?
    - May not mandate any implementation around storage security, make it a security consideration
    - **LL - we could say "you MUST store it securely", don't have to say what that means, in the security considerations**
    - **[No objection]**
- #44 - T-XXX vs. OP-XXX
    - LL - agree with Nancy that they're distinct
    - JS - what's transport? Is that TCP/IP, or some protocol running on top of TCP/IP?
    - LL - need to agree on that, think that's orthogonal
    - JS - some things you want to do, can't do on TCP/IP
    - Suggest they be operational - the protocol you use to move data from point A to point B was an OP, not a transport
    - Separate or moved depends on how you define transport
    - CK - if so, SACM not having to do with transport - shouldn't be calling it transport, call it SACM sessions or communication
    - Then transport means L1-3, all we've been talking about is above that - vocabulary problem
    - KM - could talk about data in transit, HTTP with data in it
    - Lots of ways to describe protocols that ride on top of TCP/IP
    - JS - don't have a problem with using the word transport as long as we're clear about what we're defining transport as
    - My understanding of transport is transport is the protocols that move data from point A to point B, SACM defined
    - **LL - three levels - TCP/IP underlayer, transport which is the protocols we specify, operations which those protocols must implement**
    - JS - need to look at how those are defined
    - LL - operations are like publish, subscribe, authorize - any protocol can implement in a different way
    - HB - strongly advise not to use the term transport in any other way than L4, especially in the IETF

    - IM - TCP and TLS are L4 protocols, other arbitrary protocols like HTTP are not L4 protocols, have added another concept of resource
    - DR - and yet, widely used
    - Agree that transport is used in a broad term, needs clarification
    - In 2015, insisting on 2015 meaning L4 is very limiting
    - This usage of transport, HTTP whether we like it or not is already well ahead in other WGs
    - HB - understand that there's a new way to apply this term
    - It is about transporting data, very intuitive, comes easily to mind
    - With regard to existing definition, I'm not sure it's wise to do it
    - If it's been done, new trend to do it, not sure it's a good idea
    - DR - not new - SNMP defining new ways of transport, NETCONF talks about making transport like SOAP which are even on top of HTTP
    - Not new in the IETF
    - HB - maybe I'm too much like a hard-liner in this
    - Think re-using the same term is a source for confusion
    - We could do another definition of transport and hope that everyone will get the meaning
    - DR - explain that we're using broader concept of transport, use that association with other examples in the IETF
    - HB - I was always a fan of using the word acquisition, components acquire information
    - Looking at it from another perspective, might resolve ambiguity
    - Okay with it if the group decides this is a good idea
    - **LL - sounds like consensus to use transport to refer to protocols that run on top of TCP/IP or even HTTP**
    - **DR - and move data**
    - **IM - distinguished by the fact that they're not really the application, still some application on top that has some operations**
    - **HB - Dan, if you can provide a small list of examples to the list, that would be good**
    - DR - sure, NETCONF would be the top, but there are others
    - DW - is the consensus that we leave the requirements under transport?
    - **LL - operations requirements are requirements on the application, transport requirements are requirements on the protocol**
    - **JS - need to walk through the transport requirements to decide whether they need to be moved**
    - Separate issue
    - **[General agreement, Henk is okay with consensus of group]**
- #42 - OP-002 vs. OP-004
    - LL - what would you like clarified?
    - JS - when I read that, collection abstraction, the way you're implementing collection abstraction is doing attribute-based queries
    - Or will you support queries which are not attribute-based
    - LL - I'd like to be able to make a query for everything about an endpoint - is that attribute-based?
    - JS - an endpoint is identified by attributes
    - LL - semantic rabbit-hole - if I want to query by publisher or by timeframe, those could be considered attributes - at that point, everything is an attribute
    - Text change to propose?

- JS - GitHub issue #42
- Issue I'm interested in is different from what's on the list
- Difference between what a collection operation is and what a query operation is
- Sometimes we talk about querying endpoints for their attributes, sometimes talk about collecting things from endpoints
- If you're something that's going to do an evaluation against policy, that is always a query
- **LL - one option would be to tightly define the act of gathering data from an endpoint as collection, act of requesting data from a provider as query**
- JS - that would clear things up
- DW - would that mean a query could lead to a collection?
- LL - absolutely
- HB - don't know how the terms data plane and management plane come into that
- Both somehow involved, good example to specify that
- Would like to ask if the collection is from the origin to the first SACM component, query is from SACM component to component
- Are communications between SACM components always only queries and publish/subscribe mechanisms, or is there something else?
- Are those the two acquisition methods?
- LL - collection is done by SACM component against target endpoint
- Need to be clear whether term origin means target endpoint or provider
- Personally, would prefer to have the origin be the provider, object be target endpoint
- Some attributes won't be collected directly from endpoint
- **JS - would expect origin to be entity to be performing the collection, not thing it was collected against**
- **LL - would we say object is thing it's collected against?**
- **JS - isn't that the endpoint?**
- LL - of course, no need for a separate word for that, sorry
- JFM - what if endpoint is self-reporting?
- LL - then it's both
- HB - target endpoint that is self-reporting contains SACM component that's a collector
- GE - you're differentiating between collecting for specific attributes vs. all attributes
- I'm not seeing the distinction - when we do all, it's still a finite set of attributes, still the same thing
- All known attributes, still a finite list
- LL - was trying to make sure that we weren't limiting it so we could not collect all attributes
- IM - desirable that collection operation have some wildcard - not list of attributes names, ALL
- GE - result of that is still known
- IM - result comes back same way, list of attributes with values
- HB - sometimes the entity querying, SACM component in this context knows what it wants or can state specific topics or types of data
- Sometimes there is no context information available
- The only thing you can do is get ALL without any further qualification
- LL - think we all agree that we don't want to preclude that
- JS - repeat what we're agreeing to?
- **LL - that both querying for specific attributes and querying for ALL attributes must be supported**

    - GE - what about something intermediary? All of a certain type?
    - LL - covered in the filter language
    - IM - if you have filters, ALL is one of the special keywords, can have other named filters
    - **[Agreement]**
- #37 - T-001 - just do it
- #36 - Clarification of intent of transport
    - LL - think we covered this in previous conversation
    - **DR - say see 2.1 / 2.2, get clarification**
    - **We probably need to be specific when we say TCP/UDP, say these are transports in IETF sense, understand transport in broader sense**
    - **[Agreement]**
- #34 - 1 information model - consensus?
    - DR - discussed at previous interims
    - Consensus in call, strong dissenting voice
    - AM - which GitHub issue?
    - LL - couldn't figure that out today
    - JS - one overall information model with sub-information models, or single
    - Identification won't be a sub-component
    - DR - could be modular
    - CK - the point is that there aren't alternatives
    - Might be a conjunction, several pieces, but that's different
    - Not you can do this OR that, whichever you like
    - LL - is there a specific demand to allow the information model to be modular?
    - JS - if you make it modular, easier to do things like update just endpoint identification
    - As opposed to having to update the entire thing
    - **LL - sounds like we have consensus - one information model, can be a modular information model**
    - IM - sounds like rough consensus that it should be modular
    - CK - if it adds work, I disagree, or don't agree that we know enough to know that it needs to be split into multiple RFCs
    - Don't know if we know
    - **LL - leave as may be modular**
    - IM - could be modular and still in one RFC, hierarchical
    - CK - then I don't at all know what you mean
    - LL - want to stick with what we have
    - [**Agreement]**
- #32 - DM-010 attribute dictionary
    - LL - allow data model to define attributes that are not in information model, give guidance on how that can be done
    - HB - naively suggest that if the IM encompasses the ability to be modular, would define at the information model level how a data model would be enabled to define its additional sets of attributes
    - For example, standard way of doing this - highlighted via the IM, the DM would inherit this capability
    - CK - like to suggest a different way of - agree with the conclusion, think the word modular is richly ambiguous
    - **Principle of extensibility - anyone should be able to add attributes is implied by that principle**

- **If anyone can, would be strange to say that RFCs may not**
- DR - agree with this point
- One more argument - purpose of separation between information model and data model, in some cases inevitable that you need to add attributes because of the very nature of the specific nature of the data model
- For example, some data models are not disjoint in their architecture from X protocols
- Need to add protocol elements that are transparent at level of information model because information model is protocol independent
- If we cannot expand data models, not achievable
- **LL- so we need to revise this requirement to focus on attribute extensibility, remove phrasing of attribute dictionary because limits how it can be extended?**
- DR - didn't take attribute dictionary in a narrow sense
- Maybe we should take it like specific matters, drop that term
- DW - are we saying that the information model effectively provides a starting list of attributes that may be extended?
- **DR - information model provides list of information elements, may be incorporated into different data models**
- **Data models not limited only to attributes it gets from information model for transmission of information elements**
- CK - expressing concern, or just clarifying?
- DW - clarifying
- IM - attributes that are artifacts of the data model, its relationship to specific transports, are harmless for interoperability
- The attributes that are entire new groups of identity attributes not in the information model are invisible at some level in SACM because they're data model specific
- The operations, I hope, expose the information model
- CK - so you would say yes, extensible, but in such a way that every data model would embrace the extension
- **IM - data model would define in the information model language the extension so other data models could pick it up and interoperate**
- **DW - does that mean the data model must have type and meta information to support that extensibility?**
- LL - one way to approach this is how we did with IF-MAP
- We allowed prefixing of metadata, requirement that map server accept metadata it didn't understand and pass through
- Requirement that clients ignore metadata they didn't understand
- CK - so in effect data is labeled to avoid collisions
- Question of whether we want to constrain data models in that way
- CI - question becomes do you define a type model that all of the data models must understand
- If you can ignore the metadata or typing information, you might as well not even have it
- Doesn't add to understandability
- **Otherwise, must have exchangeable typing so data models can understand extensions**
- **IM - do want that in the information model at some level; otherwise data models don't give interoperability**
- CK - I felt that we needed that too
- Think the information model doc, latest draft we have, kind of says that
- Not saying it's authoritative, says something to that effect

- May not be detailed enough - I'm sure it's not - but is that not a good way to go?
- Reservations?
- HB - not sure I understand this
- The information model would define that an IP address is a type of 4 bytes only
- IM - abstract type, IPv4 address
- HB - talked about AVP, attribute-value pair, not a type
- IM - suggesting it's not necessarily 4 bytes in the information model
- Go see various SNMP, NETCONF, other encodings
- HB - talking about IP address, string data model is using as an IP address
- If talking about AVP at information model level, I'm okay with that
- DW - would one way of handling this be to develop an IANA table at information model level that would enumerate attributes, require that the data models provide a way of binding to that IANA table to express a specific attribute and value?
- HB - would take this further
- CI - don't know if that's sufficient
- Unless you have basic data elements defined, bootstrapping is hard
- CK - agree we need some basic ones
- **Coming back to requirement question that launched this discussion, seems like we're saying that the vocabulary of attribute types is basically orthogonal to data model**
- **Can add new attribute types, it's extensible**
- **Need a controlled vocabulary to get interoperability**
- Different question from which data model you're using
- Think Ira was suggesting that these are orthogonal
- IM - data model might send string 101.49.1.57, another data model might send 4 binary octets with same meaning
- Other IETF protocols do that - okay as long as information model has authoritative definitions of what an IPv4 address is
- At one level, at the information model level, it's binary octets
- Not sure you can always squish it that hard in information model
- **Want flexibility of wire encoding and particular data models**
- Hard to achieve, I know
- CK - hard to achieve, but if the information model is describing byte encoding, it's violating the rules for what an information model is allowed to do
- IM - think so too - got to have some fluidity
- If have to have an abstract operation encoding exposed by an operation, stuck with some representation, don't want to go there if we can avoid it in the information model
- **A MAC address is a MAC address, IPv6 is IPv6 address, may have various wire representations in the data models**
- JS - I read this requirement as saying a data model could define a new attribute which is not in the information model, but which is evaluatable against
- Could define as part of an endpoint identifier, a $foo address
- $foo addrss are not part of the standard information model - not metadata, a new attribute
- LL - that is the intent of this requirement
- IM - since other vendors can do that, rough consensus that data model can do that too
- Does do some harm to interoperability, because if not done in all data models, that new data item is less accessible

    - **GE - that ability to add new data items is a necessity to keep up with rapid changes in the area**
    - **Supporting explicitly extensibility is important**
    - Next step is how are those brought back into the standard to create extensibility
    - **LL - sounds like this requirement needs to be rewritten to express those goals - requirement for extensibility and for framework to impose interoperability on that extensibility**
    - **[Agreement]**
- [time check - agreement to keep going, drop -00 draft reviews]
- #31 - DM-006
    - LL - goes back to previous discussion
    - Apply language about target endpoint vs. origin, will fall out of that
    - **[Agreement]**
- #16 - G-009
    - HB - are target endpoint discovery and endpoint discovery different things?
    - SACM component can be on an endpoint
    - Want to discover producer of information, another kind of discovery, more SACM-internal
    - LL - great point
    - **Three types of information - discovery of target endpoints, discover of which producer has what information, discovery of what schemas are in use**
    - **[Agreement]**
- #13 - Data integrity
    - JS - skip until we get data origination and transport issues resolved
- #12 - Data attribute - editorial fix
- #9 - Scalability
    - LL - want to make sure we have consensus that scalability applies to transport, not to operations or underlying TCP/IP protocol
    - JS - are you saying a collection protocol needs scalability? Not sure I agree
    - LL - very much does need it - if collecting SWID messages, may get 4GB data
    - JS - if I'm querying thermostat in my building, can be specialized transport that does not need to be scalable
    - HB - one million thermostats might need some kind of scalability
    - IM - what is scalable is very badly undefined in this conversation
    - CK - the other question is, do we need to preclude a specialized transport that will work for Jim's case, not for all other cases
    - Do we need all transports to be as scalable as all others?
    - GE - in that case, you will have decided that the scalability of a particular transport isn't inherent in the transport, it was in the decision that that usage of the transport is not subject to restrictions on scalability
    - JS - if I create a specialized transport for that particular issue, my transport doesn't meet the scalability requirements
    - CK - unless scalability means scalability enough for its use, which is a pretty sensible understanding of scalability, but gets to Ira's point
    - Wonder if it's impossible to define scalability in the abstract
    - IM - without more precision on what we mean by transport being scalability
    - LL - I put this requirement here, want to avoid limitations we ran into with IF-MAP

    - **Would like to say transports should be scalability unless they are spec purpose which are exempt**
    - **IM - include language such as "(especially for constrained devices)"**
    - **JS - do for data models, too?**
    - **LL - makes sense**
    - CK - thought constrained devices were low priority
    - LL - not specifying them here, just making sure we can spec in future
    - IM - someone can write a small data model and exclude SWID tags
    - JS - if you read G-003 in document, does define what scalability is
    - Not talking about in the abstract
    - **CK - one possibility is to say a transport must have a section on scalability considerations that talks about how many endpoints, how much information it could scale to**
    - LL - don't want that to be the only requirement
    - Would prefer to put requirements by default, have exemptions for special purposes, and then require scalability considerations in addition to that
    - **[No objection]**
- #1 - Period vs. colon - looks like we agreed on colon, editorial scrub
- LL - got through all open issues in this list
    - I have opened a few more today, would like to deal with them later
- Resolved issues still open - propose to close them all
    - JS - want to wait to close until text is written
    - LL - not proposing to close all the ones we just discussed - only to close the ones listed on this slide which are already addressed
    - **[#41, #40, #38, #3, #2]**
    - **Unless someone says they are explicitly still open, I will close them on Friday**
    - JS - okay with that
    - **[Agreement]**

Architecture

- JS - new architecture draft in the next week?
    - Comments from January not addressed yet
    - DR - submission cutoff July 6th
    - LL - don't know, have the impression Nancy is traveling, will sync up with her
    - Will try to, cannot promise

<u>Endpoint ID design update</u>

- DW - primary focus of our design team work has been to define a list of what we're calling common identifying attributes
    - These are attributes that might not be able to be provided for all device types, but tend to represent the common set of attributes that most devices will be able to provide
    - Currently seven attributes listed that we're focusing on
    - These represent to some degree types of identifiers for which there might be specific instances of those types
    - For IP addresses, could be v4/v6
    - For public keys, different types of keys
    - For tool-specific, BIOS identifier, motherboard ID, asset tags, things like that
    - Meant hardware-specific identifier

- Been working on developing a catalog of these various identifier types, define what we mean by those different kinds of identifiers
    - What potential sources of those identifiers are, whether they can be provided by the endpoint itself, or by an external authority, or through observing that identifier through some other means
    - Been talking about issues around multiplicity for a given type of attribute
    - Can that attribute appear multiple times on a given endpoint? Is there potential for the same attribute value to appear on multiple endpoints?
    - The current state of that work is available on the link provided on this slide
    - Cliff has been maintaining this list
- Next step is to get this work into the information model
    - Working towards doing that ahead of the IETF draft submission deadline
- LL - what is the specific plan for merging this into the information model?
    - DW - something we're still trying to formalize
    - Working on creating a new section based on the outline that we've been discussing during the design team
    - Don't have that information available right now to share
    - Can probably circulate that
    - Adam, do you have the outline on your Dropbox account?
    - Outline we talked about 4-5 weeks ago on the EIDT
    - AM - should be at that link that I send out with all the notes
    - **DW - could dig that up and send that around**
    - That's the current proposed outline that we're working towards addressing
    - Need to drive more discussion around the organization of the information model, how we'll address the specifics of endpoint identity relative to that outline
    - See that as one of the focuses of the EIDT
- LL - in my intermittent attendance recently, sounded like we were identifying a requirement to gather relationships between these attributes
    - IP address associated with MAC address, etc.
    - **HB - promised to send something to the list, didn't do yet**
    - **How would this look if not talking about relationship between attributes, but relationship between attributes**
    - **Triple, not a single data point**
    - Still composing that, was not able to send before this interim meeting
    - DW - based on this conversations, current notion of endpoint identity assertion is that flat listing
    - Work that needs to be done in information model to capture this grouping notion of attributes
    - CK - think there's a fork in our road
    - Number of us agree that we need structure, flat list won't capture what we need to capture
    - Nice try to keep it that simple, no such luck
    - **Question of whether to approach in the way Henk describes - relationships, triples, graph way of thinking about it**
    - **Or whether to approach with some other structure, groups of attributes, possibility hierarchical**
    - How to approach it seems kind of critical, don't think we have really a plan for how to get that resolved

    - Except maybe to sketch out a couple possibilities and then debate
    - **DW - to that point, need to be thinking about, to what point do we need to define this representation in information model vs. data model**
- DW - that's our current work effort
    - Have identified open questions in current scope of work
- One issue is need to define whether we're characterizing not just endpoint but posture
    - Not something we've tackled in the EIDT
    - Working towards information model around asserting information about endpoint
    - Uncertain whether should be focusing efforts on identifying information model for endpoint identity as well as asserting posture
    - DR - don't understand the two formalities here
    - Isn't this already part of the information model work?
    - **Not really identity - maybe we need a design team, maybe design teams are the same**
    - Not identity work
    - LL - what do you mean by characterization of endpoint posture?
    - **CK - wouldn't have said characterization of endpoint posture, but of kind of endpoint**
    - **Identifying whether running Oracle, critical server, what department it belongs to**
    - **Identifying in those ways so we know whether to monitor it properly**
    - **Came to rough consensus within EIDT that you could call that a kind of identity**
    - **Didn't want to call it that, wanted to call it characterization**
    - Different section of the information model
    - Pre-screening to define what further diagnostic tests to run
    - **LL - a more common term for that might be classification? Is that what you mean?**
    - CK - think it is
    - DW - part of the challenge we've run into with identity work is that there's a rather blurry line between collection of attributes that can be used to identify an endpoint and other types of classification
    - Because effectively you can use aspects of classification to create a rough concept of the identity of an endpoint
    - It's an endpoint running Oracle, has this collection of software
    - We've really struggled in the EIDT around creating a really clear delineation between attributes that have a high degree of likelihood of uniquely identifying endpoint vs. some of these other types of categorization
    - Caused us to struggle as to what level of effort we should focus the EIDT, focused on identity, around some of these other things
    - LL - classification is a giant rabbit-hole - unless a compelling argument for it to be in scope, nice not to have to tackle it
    - DW - certainly in SACM's charter
    - LL - we don't specify how to classify, we provide the data consumers can use to classify
    - GE - classification isn't usually something that can be found on the endpoint itself
    - If you can positively identify an endpoint, you can use that to look up what it's classification is
    - Things like is this an accounting machine
    - But you can't query the machine and say do you belong to the accounting department, are you intended to run at secret level or higher, etc.
    - HB - on the other hand, if not enough identifying attributes available to uniquely identify an endpoint, might be enough to classify it
    - In some scenarios, might be the primary use case

    - **Certainly use for that, but maybe we only need to enable the information model to provide measurement for classification**
    - Leave up to data model for how to be done
    - LL - that would certainly be my preference
    - GE - classification is necessary for knowing which policies apply to an endpoint
    - That's still an external lookup, will change from organization to organization how things are classified
    - Not something you can make as part of the endpoint identifier
    - **HB - correct, more like a framework that can be domain specific**
    - **DW - I'm hearing that we probably shouldn't focus our efforts right now on this issue**
    - CK - I think you're hearing that right
    - **We're hearing something stronger than that, which is that issue should be out of scope for the EIDT, if it's going to be tackled at all**
    - My personal view, the EIDT should be finishing up its work and dissolve itself very soon
- DW - another question, scoping issue
    - We've recognized that for a given set of consumers, different identifying attributes might be more desirable over others
    - Often impossible to know a priori at the point of collection what the set of identifying attributes are that downstream consumers might need
    - One thing we talked about in the last couple of weeks has been maybe using guidance as a way to control what endpoint identifying attributes a given collector will collect
    - Do folks see this as something that should be in scope
    - **LL - why wouldn't you just collect as many attributes as possible?**
    - **Let the downstream consumers use whatever they find useful**
    - **HB - generally speaking, issue of privacy in some domains**
    - Sparse collection is a plus in this use case
    - **In general, I would agree that it could be a downstream-weighted problem**
    - CI - always a downstream problem - making a value proposition for folks that may not share their same values
    - Consumer may have different issues
    - CK - assumed guidance would have to be a union
    - **Find out what consumers were interested, put together guidance, collector would have to provide anything that any consumer wanted to know**
    - True for all attributes, not just identifying attributes
    - If we need to engineer something where we send some subset of all possible attributes, okay
    - If we don't, we shouldn't
    - Identifying attributes may be part of mechanism for getting guidance to endpoint
    - Circular problem
    - DW - one thing we talked about was maybe having minimal set of identity attributes that would bootstrap that process, help to understand what guidance would apply
    - CK - this creates a pile of complexity, don't' think this question needs an answer
    - May have run away with ourselves
    - If we can get away with something as simple as send it all, then why don't we
    - DW - maybe that's a question we should grapple with before take on more work
    - LL - haven't heard an answer other than the privacy issue
    - **CK - were imagining there might be a performance concern**

    - Device being monitored has a number of certs that it uses, no real need to send all of them if it's enough to send one of them
    - Imagining something like that, didn't spell out a performance requirement and figure out whether we had a problem
    - May have put the cart before the horse
    - **Think we're agreeing that we don't pursue it unless we have a good reason, and at the moment we don't have a good reason**
    - **DW - agree, put some thought into coming up with answers to Lisa's question**
    - **Put some brainstorming into that, revisit this question at Prague**
- DW - third questions, more conceptual, applies to broader set of information that SACM may want to include
    - One thing we talked about historically in EIDT is notion of provenance
    - Based on previous calls within EIDT, larger working group
    - Tackling aspect of provenance is probably too much to do right now
    - Not necessarily at core of what we're trying to address, would be a rat-hole
    - That being said, left with a few aspects of provenance-oriented information that we think is actually important to still consider relative to our work in the EIDT
    - One question that came up was, should we provide - as part of the information model - the ability to provide identification attributes for both subject, target endpoint, as well as the SACM component doing the collection?
    - Came up during the requirements discussion, specifically in the case where those things differ
    - If the same, avoid the case where we're providing a duplicate set of attributes
    - Does anyone think it would be a good idea to provide identification information for the target endpoint as well as the SACM component doing the collection?
    - LL - think it's an absolute requirement
    - ?? - but not for EIDT
    - **DR - question could be asked the other way - does anyone think we should not do it?**
    - **HB - all agree that it should be in there**
    - Lisa elaborated on that in the EIDT
    - Four scenarios covered here, might be a little more complex than only target endpoint and origin
    - **Also about - is this expected state or observed state?**
    - Those two can be in practice difficult to differentiate, may need to be highlighted specifically
    - Attribute you're checking against or actual value you managed to collect through observation or self-reporting
    - **DR - let's continue this discussion currently in the design team**
    - **If we get to the conclusion that we need a new design team or an expansion of mandate or something like this, we can decide on the list or at IETF 93**
- DW - all we had to talk about
    - As far as ongoing work, making a primary focus of our activities around updating the information model
    - Starting to focus more heavily to driving changes to text in information model
    - Plan to continue to make updates, would like to submit our current state on July 6th
    - Will continue to work on the information model up to IETF 93
    - **Will be in a position where we'll have to drive some conversation at IETF 93 around changes not included in a draft submission on July 6th**

- DR - are you using GitHub?
- DW - yes, discussed during last meeting, planned to make use of GitHub for that

Way Forward

- DR - what happened with the use cases? Is this back with AD or IESG?
  - IESG provided feedback, right?
  - DW - I addressed a couple questions that Stephen had on the use cases before it would get finalized for publication
  - Sent out an update on the use cases to Kathleen, waiting to hear back on those issues
  - Think Stephen was out on vacation, waiting for him to come back
  - DR - he's back
  - **Maybe Kathleen could help to ping Stephen and if possible submit this update before the deadline**
  - If it's actually approved, announcement can be sent, revised I-D agreed by Stephen is submitted, goes to the ?? and we can declare victory
- AM - avoid serialization where possible
  - Seems like we don't have enough actively working people to parallelize
  - **Get information model draft update**
  - **Get requirements draft updated for IETF 93, like to see it done before that**
  - Have a suspicion we'll still be discussing it in the next month
  - DR - at some point in time, will issue a second WGLC
  - Making progress
  - AM - important part is that we are meeting next month in Prague
  - Focus areas - these three things
  - **LL - would like to add a discussion of the I-D Jess and others submitted**
  - **DR - protocol solution discussions, need to focus on this**
  - Don't need to wait for everything to be finalized on the information model or requirements to start working on solutions
  - DW - I think if we can get to a point where our terminology is fairly stable, think that will drive a lot of changes to the architecture, information model, requirements
  - Would be good to get some degree of consensus around the terminology so we can move forward with that
  - DR - don't forget that terminology needs to be aligned with the rest of the work
  - Something we keep open as long as we are… [?]
  - Need consistency but terminology is not by itself
  - DW - if we want to make consistent use of terminology, need terminology in a stable state where we can begin to make those updates
  - **Good to have WGLC on current state of terminology so we can get to that point**
  - AM - know there are some updates coming on terminology draft
  - It sounds like we just said everything is a high priority, and we don't have enough people to do all that work
  - Can say we want to focus on 6 different things, don't think we have enough people to do all that work well
  - DW - which are critical path issues?
  - **LL - think terminology is critical path**
  - HB - I tried to raise most glaring problems with terminology in our draft in issue tracker on GitHub
  - **AM - what's after terminology?**

- **DW - requirements are also critical path, setting the tone**
- **AM - follow that with architecture or information model?**
- **LL - think architecture**
- DW - agree
- DR - got mail from Kathleen, cannot make voice heard
- She plans tomorrow to talk with Stephen about the use case draft, hopefully will have response on this by end of tomorrow

Adjourned at 1500 Eastern