Security Automation Developer Days

July 9-13, 2012

Table of Contents

Introduction

Security Automation Developer Days was held on July 9 - 13, 2012 at The MITRE Corporation in Bedford, MA. This event was the most recent chapter in an ongoing series of workshops, beginning in June of 2009 at MITRE.

One hundred eleven people registered for the event, and 80-90 people were present each day of the workshop. Over the five days, thirty-one sessions were held and this document contains a comprehensive summary of each of those sessions.

As you prepare to review these minutes, the authors would once again remind you that the standards cannot continue to advance without ongoing discussion of the issues throughout the year. This is accomplished through dialogue in the email discussion lists. A complete list of these email discussion lists can be found here: http://measurablesecurity.mitre.org/participation/index.html . Please sign up for those lists that interest you.

What follows is a detailed summary of the discussions from the event.

Attendee List

Advanced Cyber Security	-	Andrew Ferguson
Alcatel-Lucent	-	Fabio Jaramillo
Apex Assurance	-	Luis Nunez
BAE Systems.	-	Vince Ellspermann
	-	Stephen Frieda
	-	Jeremy Unruh
BCF Solutions	-	Tim Foerster
Belarc	-	Gary Newman
	-	Sumin Tchen
Booz Allen Hamilton	-	Larry Feldman
	-	Adam Halbardier
	-	Tim Harrison
	-	Timothy Nary
CENTECH Group	-	Mahadevan LakshmiNarayanan
Cisco Systems	-	Anthony Busciglio
	-	Omar Santos
CSC	-	Michael Rains
Deloitte & Touche	-	Craig Astrich
DHS	-	Rich Struse
DISA	-	Brady Alleman
	-	Jared Joels
	-	Matthew Lempka
	-	Jason Mackanick
	-	Jordan Shuhart
	-	Leland Steinke
	-	Joseph Wolfkiel
DOD	-	Edward Wienholt
DoD PEOMA	-	Pete Schmidt
DTCC	-	Aharon Chernin
EMC	-	John Field
	-	Kathleen Moriarty
EPA/CSC	-	Vincent Ross
Epstein Becker & Green	-	Robert Hudock

G2, Inc. .	- Matthew Kerr - Shane Shaffer - Faith Wingate
General Dynamics	- Alan Chen - Jeremy Wyant
gOCIL.org	- David Ries
IBM	- Louis Lam - Charles McClain
Information-technology Promotion Agency	- Masato Terada
jOVAL.org	- David Solin
JP Morgan Chase & Co	- Ryan Clough - Matthew Wong
Juniper Networks	- Steve Hanna
L3	- Michael Carter
McAfee	- Donald Campbell - Kent Landfield - Dick Whitehurst
Microsoft	- Michael Tan
MITRE	- Jonathan Baker - Sean Barnum - Steve Boczenowski - Penny Chase - Brant Cheikes - Michael Chisholm - Mark Davidson - Tom Graves - Kayla Green - Matthew Hansbury - Danny Haynes - Jasen Jacobsen - Ivan Kirillov - Gerry McGuire - Stelios Melachrinoudis - Michael Peck - Matt Richard - David Rothenberg - Charles Schmidt - Bryan Worrell - John Wunder - Margie Zuk
Modulo	- Marlon Machado
NASA/DB Consulting	- Gary Gapinski

Navy	- Chris Hairston
NetIQ	- Jose Palma - Kirk Sievers - Mark Slosberg
Nexagen Networks, Inc.	- Rupal Parikh
NIST	- John Banghart - Harold Booth - David Waltermire
Northrop Grumman	- Richard Galloway - Jason Liu
NSA	- Mike Kinney
Owl Computing Tech	- Jose Gonzalez
Polycom	- Erik Cockrell
Rapid7	- Jyoti Kedia
Raytheon	- Mike Kuhnkey
Red Hat	- William Heinbockel
Samsung	- Venkata Kiran Chegu
SecPod Technologies	- Chandrashekhar Basavanna - Greg Pottebaum
SPAWAR	- Jack Vander Pol
Symantec	- Jason Meinhart - John Richardson - John Williams
TagVault.org	- Steve Klos
Teledyne Brown Eng	- Randy Wynn
Telos	- Justin Furniss
Tenable Network Security	- Mehul Revankar
Tripwire	- Brian Cox - Rob Etzel - Robert Huffman - Tom Pearson - Todd Whitaker
Triumfant, Inc	- Bill Goodrich
US Army	- Randy Dague
Varen Technologies	- Jim Ronayne

Monday July 9th

CCE

Presenter
David Mann, The MITRE Corporation

Overview

Dave Mann started the presentation giving an overview of what CCE is and the current work being done on the project. Dave discussed the five use cases for CCE: Configuration Management Lifecycle, Guide Document Authoring and System Design, Configuration Tool Configuration, Audit Tool Result Integration, and Regulatory Compliance.

Mike Kinney noted that some of the use cases seem different from the others and Dave responded that CCE is not a complete solution for all of the use cases but can be used to support their correlation needs.

Dave continued the presentation discussing the outreach program MITRE is currently engaged in. He said that Cisco, CIS, and DISA are currently working with MITRE to potentially produce CCEs and notes that there has been a shift away from vendors and toward security guide authors. It seems that security guide authors better understand the configuration management life cycle better than vendors currently do.

Dave moved on to describe the progress that CCE has made over the past year. New platform groups were created for Polycom and Microsoft Exchange 2007 and 2010. It was noted that Polycom did a great job applying CCE content decisions and creating candidates. There is work currently being done on creating platform groups for Apache HTTP server and Apache Tomcat and a team working for the NSA is currently working on creating new CCEs and platform groups. Dave also mentioned that the Content Decisions are being review, which was the focus of the remaining part of this presentation.

It was asked if the CCE team was creating platform groups for SCADA devices. Dave responded that there are not currently any platform groups of the kind in the works but if there were someone who published a security guide and was willing to work with MITRE to create CCEs, then a platform group would be created.

The rest of the discussion was over the problems MITRE's CCE team is encountering. Dave emphasized in the lead up to these discussions that everything in CCE is a judgment call, and CCE has decided to go one way but others may go another.

Publicly Available Reference Documents

MITRE does not research and validate configuration controls so MITRE uses public security guides as validation of a configuration control. Not all configuration controls should be in the CCE corpus. Listing

every configuration control would be too much information so MITRE limits the configurations to those in a security guide.

Problems

- References disappear
- References are loosely versioned

MITRE's Policy

- MITRE will require that the reference be public at a time. No other requirements are placed on the reference.
- MITRE suggests that the community work with the guide creators to implement versioning and archives.

Discussion

- Is there a process in place to work with guide authors to update CCE references when they change?
  - There is a process to work with the authors
  - But updating references has a low priority compared to many of the other issues that need to be dealt with. MITRE would like to this happen.
- Who is considered a CNA (Candidate Naming Authority)? Are CIS and DISA CNAs?
  - Anyone who is willing to work with MITRE to create CCE candidates will be considered a CNA.
- Has MITRE looked at using something like DOIs?
  - CCE used some Dublin Core descriptions.
  - However that is a different question from archival and versioning.
- Must reference documents be a security guide? Can it be a reference manual?
  - "Security guides" is old language. CCE will accept reference manuals and technical documents like Microsoft TechNet™.

Proprietary IDs

Overview

For the CVE initiative, it has been easier to perform mapping when there are identifiers at both ends. This would likely be the case for CCE but there are very few other identifiers for configurations. Primary source vendors cannot solve this problem. Just as primary source vendors often think in terms of patches and not vulnerabilities, they talk about source code branches and features and not configuration items.

Problem

- CCE cannot cover all platform groups. There are various reasons for this: size, proprietary nature, etc.
- CCE IDs are need for the majority of software used by an organization
- CCEs are being used as primary keys and thus are considered a "bottleneck."

MITRE's Policy

- CCE team encourages others to create their own proprietary IDs.
- Creating IDs means the creators have to create "chunks" of configurations which makes comparison easier.
- CCEs should not be used as primary keys in local knowledge repositories.

Discussion

- There was a question earlier about not having CCEs for small platforms. Dave Mann pointed out that proprietary IDs solve the problem where there is not a CCE platform group for software being managed by an organization.
- What happens when those creating the proprietary IDs organize the configurations in a different way than CCE?
    - Like with CVE, this would be acceptable and expected to happen but the proprietary IDs would allow MITRE to map it similarly.
    - CCE tries to create a consistent way to count configurations. CCEs use the content decisions to try and identify configurations in a consistent manner.

Platform groups vs Codebases

Little was discussed in this section. The major points made were:

- CCE is shifting away from working with primary source vendors. They do not think the same way about configuration policies
- Defining what Linux consists of is difficult but the configuration management community thinks in terms versions of Linux like Red Hat Linux.
- This is likely due to a lack of maturation of the community. MITRE encourages the community to pressure primary source vendors to think in the ways of the configuration management community.

Content Decisions

As discussed earlier, the content decisions document is currently under review. Developer Days was not intended to make decisions on the revisions. Discussions will continue on the CCE working group mailing list.

Default Objects

- Each default object gets its own CCE identifier.

Default Sets

- There are default ways to create sets of objects such as files in etc/password.
- We split on the default sets.

All Objects

- One CCE will be issued for a configuration that must be applied to all object of a type, e.g. user accounts.

- CCE includes identifies for configurations in the system and those that are specified in a configuration guide. Some configuration guides are more like policy statements and do not necessarily have controls that can actually enforce them.

User Defined Sets

- Some sets of configuration items in a system can be created by users. An example would user groups.
- These user defined sets are not necessarily common to all organization and are not appropriate for CCE identifiers. CCEs are meant to common.
- Security audit tools provide capability to define a user created set and CCE should provide a mechanism to do this as well.

Discussion

- Question: Is the difference between the two examples (minimum password for the Accounting user group is an acceptable CCE and the startup type of nonessential service would not be an acceptable CCE) that one applies to everything in a group and the other applies to a conditional subgroup?
  - The difference is that there is no grouping mechanism in the system to define sets of service.
- Question: It is possible to parameterize a list of nonessential services and give it to a tool. Why would it not be acceptable for there to be a CCE for that configuration?
  - The CCE team recognizes that it is possible to go this way but whitelists and blacklists have been avoided historically.
- Remark: In some cases where there are defined ranges then they can be standardized into white and black lists.
  - The content decision does provide for parameters. CCE calls these TARGET parameters and a collection of these would create your lists.
- Remark: Organizations can make their own list of “nonessential” services, why can’t these lists be created for CCEs.
  - This content decision only deals with grouping mechanisms. There are other content decisions that can be used to provide the list functionality.
- Remark: It is splitting hairs to say that each individual service has to be designated as essential or nonessential.
- Remark: These are proposed content decisions. If the community does not feel that they are necessary or sufficient then the CCE team will change them, but that discussion should take place on the CCE working group mailing list.
- Remark: There may not be a CCE to relate to each Common Remediation Enumeration.
- Question: How to relate proprietary IDs and CCEs
  - Based on the experience with CVE, it will be easier if you have atomized your configuration controls and assigned ids. It is harder to complete a mapping when the boundaries of these configurations are not well defined.

- Remark: MITRE has created a bottleneck since they decide which configurations are legitimate, resulting in the situation where there are no CCEs for some configurations. If a system has to be created based on proprietary id, there is not an incentive to go back and add CCEs. The manual review process is inefficient and it would be better to implement a pass through and deprecation system instead.
    - MITRE requires a public reference for a configuration. It is extremely rare that MITRE rejects a candidate CCE because they did not meet the requirement of a content decision. The biggest problem is the public reference.
- Question: Why can't MITRE be its own reference?
    - The problem is that MITRE is not funded to create or validate configuration information. MITRE does not write the CCE content. As long as there is a public reference and someone creates the candidate CCEs using the content decisions, MITRE will publish the CCEs.
- Remark: There are several costs to creating CCEs. There is the cost of creating the CCE information and there is the cost of de-conflicting the CCEs. This is similar to the problem with CVEs. The problem is that the current governance model makes CCE creation expensive. Maybe we should look at different governance models for CCE.
    - A different governance model is to use proprietary IDs.
- Question: Is the creation cost offsetting the benefits realized by having the CCE for correlation?
    - Some configuration management tool creators have a problem in that they cannot dump their proprietary data and put it on the web. This was tried by some in the early days with CCE but the previously propriety data was then used by others to create a competing product.
    - Primary source vendors have a problem in that the way they think about configurations is not useful for configuration management. A vendor may have a set of features but naming is not consistent. Letting vendors create their own CCEs is essentially creating a federated name space but this would not support the correlation role of CCEs.
    - The way CCEs are assigned now standardizes the way configurations are abstracted, which allows for the correlation that CCE was intended to enable.
    - Primary source vendors do not see the financial incentive to create CCEs.
    - Kent Landfield of McAfee pointed out that CCE uses a different level of abstraction than CVE. CCE uses platform groups while CVE applies to affected software.
        - Higher level platform types (e.g. UNIX) were tried in earlier versions of CCE but it was found not to work, which is why it was changed in CCE 5.
        - Also, CCEs are not bounded by a temporal fact like CVE is. CCEs come out in bursts when individual platforms are released. CVE has a fairly steady stream vulnerabilities coming out.
    - Question: CCEs are more of a policy statement. CCEs are not actionable and automatable. So why are they platform dependent? Is the same configuration for Red Hat™ 4, 5, and 6 the same?

- The answer is that a new platform release may have changed their underlying security model. An example would be the change from XP and Windows 7 with file permissions.

User Defined Objects

- There will be one CCE issue for each configuration on a user defined object, with the object being defined by the TARGET parameter.
- This allows for whitelists and blacklists but you have to specify each individual object.

Platform Groups

- It is not clear what should be included in a Linux platform group.
- Red Hat does not own many of these platforms and does not want to support the CCEs for these platforms
- Apache HTTP server brings up the problem of software that applies to both to Linux and Windows.

MITRE's Suggestions

- Platform groups are to be created when it is useful, which is implied by the way security guides and configuration management tools define platform groups.
- Some rules to create platform groups might be:
  - The platform is supported by an organization
  - The platform should be versioned
  - The platform should have a reasonable number of configuration statements
  - The platform group is useful to the community
- Remark: Cisco is using CPEs to define their platform groups.

Underlying Platform Group Content Decisions

- Some content decisions were defined to help guide when platform groups run on top of other platform groups. These were not discussed in any detail.

SCAP Repository Demonstration

Presenter

Chandra Basavanna, SecPod

Introduction

Chandra from SecPod Technologies began the section with a short presentation regarding the SCAP Content Repository that is now hosted by SecPod Technologies. Currently content is hosted by a variety of different repositories. Ideally an organization could retrieve all of their content from a single location. SecPod's repository provides such a location, and has the following interfaces:

- GUI-based Search Interface

- Admin Interface (importing/exporting content)
- Web Services Interface (automated interface for downloading content)

Demo

Chandra went through a demonstration of the current SecPod Content Repository functionality. First he showed several ways to search through the repository via the browser-based search capability. This included numerous ways to search for and acquire SCAP content of all types. He also demonstrated some complex queries and advanced search features.

Next he showed off the ability to use a command line interface to administer the system, allowing a user to import and export content into the system, as well as view statistics related to the repository.

Chandra also showed the Web Service client developed by SecPod to access the Web Services interface, allowing the user to download SCAP content.

Discussion

The following items were discussed during this section:

- It was asked if the Web Services interface was intended to pull in SCAP content into an SCAP-validated tool. Chandra confirmed that that was the intent.
- Another question was asked about how SCAP bundles would work in the context of downloading them via the Web Service interface. It was explained that downloading the SCAP bundle would get the entire SCAP bundle, including the OVAL Definitions, CVEs, etc.
- Another question revolved around the subscription model, to which Chandra explained that they had several subscription models that broke content up by patch, vulnerability, etc.
- It was asked if the various Schemas were adequate to implement the search capabilities. The presenter answered that they needed to make a series of mappings to make the search work correctly.
- The question was asked about how the repository handles errors for things like content not found. The presenter answered that an error would be returned.

CPE – SWID Tagging

Presenters
Steve Klos, TagVault
Brant Cheikes, The MITRE Corporation

Focus of the session was on joint MITRE-TagVault.org efforts to promote interoperability between the CPE standard and the ISO/IEC-19770-2:2009 standard on software identification tags.

CPE is currently in version 2.3 (released August 2011) and NIST is working on operationalizing this latest version of CPE. Since July 2011, MITRE and TagVault.org have been collaborating on a technical

proposal to integrate CPE-related information into future software identification (SWID) tags. An early audience "poll" revealed that almost nobody in the audience was "very familiar" with SWID tagging and the related standard. Today's session will present a mini-tutorial on SWID tagging, then walk the audience thru the current technical proposal involving the integration of CPE information into SWID tags.

Steve Klos, Executive Director of TagVault.org, presented the mini-tutorial. (Slides 4-23 in the session briefing package.) Q/A during that portion of the meeting is noted below. Readers are referred to the accompanying briefing charts and to the TagVault.org website to learn more about SWID tags and the associated ISO/IEC standard. Key fact: the SWID standard mandates a unique ID as an element of SWID tags. This is an obvious point for interoperability with CPE.

Q: Does Microsoft plan to patch products that are already released to incorporate SWID tags?

A: Heather Young of Microsoft made a presentation at the recent Software ID Summit. See the TagVault.org website for the video. In brief, MS has committed to releasing SWID tags in new products. The Windows 8 beta 2nd release already installs with a tag. They have also committed to supporting reading tags in System Center and Map toolkit. For older products, future patches will include tags for both the product and for the patch, but no detailed roadmap has been released for when this will happen.

Q: Can tags be spoofed?

A: Yes, if a tag is not digitally signed. Tags are just XML files installed on the computing endpoint. Digital signing ensures that the values in the tag haven't been tampered with. However, the presence of a tag is not a guarantee that the product is actually installed. Tags provide a means for the "package footprint" to be explicitly listed, which if used would enable a significant degree of validation that the tagged product is actually installed. So spoofing is possible to some degree. But tags do provide mechanisms to at least partially verify.

Q: Does a tag provide just affirmation of existence, or does it document all the pieces that comprise a software product?

A: It does provide affirmation of product presence. If the tag creator takes advantage of the optional "package footprint" element and certifies the tag at level higher than basic, then the tag will include details about the files comprising the product. Note, however, that the details of how the package footprint will be specified are still being worked out as part of the next revision to the ISO/IEC 19770-2 standard.

TagVault plans to establish a repository for storage of certified tags, but access will likely be behind a paywall.

Q: Will TagVault members also have to pay to access the repository? Would non-members be enjoined from establishing repositories of their own?

A: Members would not have to pay to access a TagVault-established repository. Non-members would not be enjoined from establishing tag repositories of their own.

Q: Tags provide a capability to encode "component_of" relationships among products. Does this enable multiple distinct products to, e.g., record the fact that MS SQL Server Express is a component?

A: Yes. Though sometimes multiple distinct instances can be installed, as opposed to a single shared instance. But tags can handle either case.

After Steve Klos' mini-tutorial on SWID tagging, the group spent time reviewing a draft technical proposal, currently circulated within both the SWID and CPE communities, but targeted at the SWID community, outlining an approach to embedding CPE name-related information into SWID tags. Below we capture Q&A related to this proposal. (Slides 24-34 in the briefing package.)

Five new elements are planned to be added to the SWID tag standard, in the next revision of the standard beginning in August 2012. Those elements are currently named:

- Product_name
- Licensing_version
- Product_edition
- Target_platform
- Product_update

By combining values provided for these elements, a CPE name can be mechanically constructed. There are no current plans to embed a CPE name in its entirety in a SWID tag.

Q: How will normalization of element values, both within and across vendors and their product lines, be accomplished?

A: Normalization methods will vary for each element. Some element values will be normalized by TagVault and the software publisher as part of the tag creation process.

Q: Need to consider adding a deprecation process for legacy CPE names. That is, if the tag normalization process leads to the creation of new names which are inconsistent with legacy names, it may be necessary for dictionary deprecations to bring all names into consistency.

Q: What is the intended meaning of the target_platform element?

A: The new SWID element "target_platform" is currently intended to correspond to the CPE v2.3 notion of "target_hw". The CPE v2.3 notion of "target_sw" does not currently have a corresponding element in the proposed revision to the SWID tag standard. To date, the target_sw CPE element has not been well specified.

Q: But what if the publisher puts target software information explicitly in the product name, e.g., Microsoft Office for Macintosh?

A: At this point, that kind of target software information would still be contained in the product name, but would not be parsed out into a separate element.

Discussion of proposed SWID equivalent for CPE 'part' (slide 27):

Thus far there is no provision for embedding 'part' related information in the SWID tag. This raises the question of how one could automatically create a CPE name from a SWID tag that assigns a correct value to 'part'.

Q: We can probably do without the 'part' value; it seems to have relatively little value. (Though it's helpful to be able to sort applications separately from operating systems.) But we've heard that SWIDs can also be assigned to patches, which isn't true for CPE names. If one had a SWID tag for a patch, would there be a way to create a CPE name that properly distinguished the CPE name as the name of a patch?

A: In tagging space, tags can explicitly represent parent/child relations. Also, there's a tag type element in the tag (coming in the next rev of the standard). This is closely correlated to the CPE part, but has a richer set of possible values. One option might be to just use the SWID tag type as the CPE 'part'. Unfortunately, v2.3 of the CPE Name spec limits the value of the 'part' field to only 'a', 'o', and 'h'.

There was some general discussion about the opportunity created by SWID tags for patches to have CPE names. Today's CPE standard doesn't really support naming of patches, so this would need to be fixed in a new release of the CPE standard.

Patching is a complex area. There are "patch sets" and sequences/dependencies of patches. Naming strategies for patches would need careful thought.

Audience feedback suggested that there would be value in techniques for naming patches.

Another part-related complication: some products that in the past were considered applications are increasingly looking/acting like operating systems. The lines have blurred particularly where virtualization applications are involved.

Discussion of proposed SWID equivalent for CPE "vendor":

The SWID proposal proposes to use the fully-qualified publisher domain name as the source for the CPE "vendor" element. This would be inconsistent with current CPE convention. CPE names today use, e.g., "acme" rather than "acme.com" for the "vendor" element. If there were a name conflict, e.g., products from "acme.org" and "acme.com", then CPE would use the fully-qualified domain name for the second vendor to appear in the product dictionary. Thus far, this has never occurred, so the point may be moot.

Since we're no longer thinking about embedding a full CPE name in a SWID tag, this is probably not a significant issue. By looking into the SWID tag, one can determine the fully qualified domain name of the product vendor. This could be mechanically altered if desired to reduce, e.g., "microsoft.com" to just "microsoft".

Q: How would this work with small vendors, open source software, etc.?

A: The SWID community has an approach which will work; currently, the plan is to base the registration id on email addresses for those cases in which websites/fully qualified domain names don't meaningfully apply.

Q: Could a product have multiple software creators?

A: SWID tags make provisions for different namespaces. The regid can allow nested namespaces to be represented. Not clear how this could be handled in CPE names. Should CPE be able to represent the fact that "product A" from IBM is actually from their "database" division, while "product B" is from IBM's "analytics" division.

Q: What about internationalized domain names? (Unicode compliance.)

A: Lots of issues with internationalization. CPE names are currently based on the ASCII character set. Dictionary entries can contain product names in other languages/character sets.

Q: It appears as though this proposal creates the possibility for lots of new CPE names to be created (or creatable) that are not consistent with existing practice. What do we do with those names?

A: That's a CPE community decision. The ideal is that we find a way to accommodate those new names into the dictionary, possibly deprecating legacy names as needed to maintain consistency. Or we accept some limits to full dictionary consistency in the interest of being able to rapidly incorporate new names for newly released products that come with SWID tags.

Discussion of proposed SWID equivalent for CPE "product":

Two elements being considered on the SWID side, which could be combined to create a CPE "product" value: product_name and licensing version. Licensing version is intended to capture other version-related information that is frequently embedded in product titles (like the "2010" in "Office 2010"), and keep that separate from codebase version information. Licensing version would only be specified if it made sense for a given product.

Q: Could one generate a SWID tag from a CPE name?

A: No. CPE names do not contain the minimum information required for a valid SWID tag.

Discussion of proposed SWID equivalent for CPE "product_update":

Q: Will vendors be consistent across their product lines in terms of how updates are specified?

A: Not necessarily. Normalization across vendors will definitely not occur. For example, vendor A might refer to updates as "updates" while vendor B refers to updates as "service packs". There is a possibility that a individual vendor might be consistent in terms of how they label updates. But certain very large vendors with many products in many business units may use a variety of naming strategies.

Session wrapped up without covering a couple of other elements. Community members should expect to see an update to the integration proposal posted to the CPE discussion list. Today's session was effective in equipping the attendees with good information about SWID tags, so they can better review the integration proposal and provide actionable feedback.

MILE and Information Sharing

Presenter
Kathleen Moriarty, EMC

Note: Only the discussion points are captured here. Briefing content is not captured.

Q: Within IODEF, is there any way to correlate similar incidents?

A: There is a field to represent related incidents & related incident IDs within IODEF.

Q: There is no common format for reporting vulnerabilities. Has there been any thought given to expanding MILE to disseminate 0-Day vulnerability information?

A: There are a variety of options (RID/IODEF/CVRF). Not sure what the right option is. A full discussion on the various solutions would be beneficial.

Q: Is extending IODEF simply a choice between using a structured representation (like XML) or some text that may or may not be useful to anyone but the originator?

A: Sounds like 2 questions:

A1: You can embed a schema, extending IODEF

A2: You can include an enumerated value, we would likely use an IANA table for these values.

Dave Waltermire of NIST and Adam Montville of TripWire have had a lot of discussion on this topic on the MILE mailing list. Discussion on the mailing list will help drive the discussion.

Q: When we were talking about 'mandatory to implement', it was asked whether or not there are varying levels of 'mandatory to implement'.

A: This is an opinion, and it is important to discuss on the mailing list.

Q: [Inaudible]

A: IANA tables are needed for both instances, so you can parse them and know what you might expect.

Q: Some XML representations lack a schema. If somebody decides to use ASN1, is that a potential representation scheme, or are we stuck with XML?

A: The nice part about IODEF with the additional data class is that you can include all kinds of data. IODEF was also written as a UML description. It was written as a data model first, then it was implemented as a schema.

Q: So, I think you probably have the requirement right, in terms of what the goal is. Given a certain spec ID, I would like to process the content. We are trying to provide the flexibility here.

A: The questions here are about which enumerated values make the most sense. Anything can be included now. Your sharing partner would have to know what to expect in order to parse it.

Q: The one thing that might have helped is to use some actual examples.

A: Yes, I'm seeing that now.

Q: Must a conforming implementation handle all registered formats?

A: No, that's a reason for IETF's mandatory to implement.

Q: There are likely solicited and unsolicited communications. Is it possible for the solicitor to specify which formats are acceptable?

A: Currently, the answer is no. However, any shortcomings can be addressed when IODEF is opened back up.

Q: What is the difference between RID & GRC Exchange?

A: RID was written specific to incident response, GRC is a generalization of RID and is meant to cover any schema.

Q: We are going to be having a lot of discussion about GRC exchange on Tuesday. We are working to incorporate those concepts into the Continuous Monitoring reference model work. I think the important thing to point out relating to GRC, is that you can use it to exchange any kind of information.

A: A very different example may be a regulator that you need to deliver a regular report with. You could use GRC exchange to send that report.

Q: How do you specify what kind of response information could be accepted?

A: That's something we'll have to look at. Dave Waltermire of NIST will talk more about GRC exchange during Tuesday's sessions.

Q: From a business perspective, do you have a list of some of the major centers that are adopting this: Verizon, AT&T, ISACs, etc.? Any products?

A: ArcSight, DFLaboratories are two - There are other areas that are very active in this category.

Future SCAP Releases Discussion

Presenter
John Banghart, NIST

Quick Summary

Major points for this talk included:

- Instability, in the form of causing the specification to change faster than those changes can be supported, was acknowledged as something that needs to be avoided. Both the frequency of change and the magnitude of individual changes were noted as contributing factors to this problem.
- There were multiple discussions about the differences between SCAP and the specifications that make up SCAP as far as how they can be supported and evolved. Multiple parties noted that it could be possible to evolve individual specifications without an expectation of immediate changes to SCAP itself. It was also noted that some standards and even some use cases might never be made part of SCAP, and that this would not necessarily be a bad thing.
- It was observed that having reference implementations of standards that are intended for operational use within a functional architecture (as opposed to serving as stand-alone proof-of-concepts) would be invaluable to the community. These would reveal weaknesses in specifications before the specification was finalized, and would also help guide the validation program by identifying critical functions and interfaces.

Minutes

The discussion began with an overview of the history of SCAP policy and administration from a US government perspective. This began with the Federal Desktop Core Configuration (FDCC). The idea behind this was to have a strong set of configuration recommendations for Windows. The goal of this was to both improve security across the US government and also to simplify management by having a single baseline from which everyone was operating. NIST worked with multiple parties to come up with the appropriate settings. The challenge was then how to deploy this configuration and assess that deployment. As a result, it became clear that there was a need include automation as part of FDCC.

At that same time, there had been community efforts to put SCAP together. One of SCAP's use cases aligned perfectly with FDCC's need to characterize settings in a machine-readable format. Seeing this, the US Office of Management and Budget (OMB) directed NIST to put FDCC in the machine-readable SCAP format. In addition to pushing out the content, OMB tasked NIST with creating an SCAP validation program to test tools for SCAP compatibility. NIST acted on this direction, creating the content, the validation program, and eventually a formal SCAP specification. Today the OMB mandate (now called the US Government Configuration Baseline, or USGCB) remains in place - a US government initiative that requires NIST and other agencies to use and advance SCAP, and to maintain a validation program so US government agencies can use SCAP tools to manage their networks. It is important to remember this background and understand that there are structures within the US government that have influence over how SCAP moves forward. The individual specifications, such as CCE, XCCDF, OVAL and others, all stand alone; SCAP pulls these together into specific use cases and initiatives that have strong, policy-centric aspects within the US government.

When considering what to do next in SCAP, the major question from NIST's perspective is "what are the criteria for making another set of changes to SCAP?" If OMB told NIST to create SCAP 2.0, this would certainly drive NIST to comply with this mandate. More importantly, NIST wants to hear from the

community as to where there are gaps and deficiencies, as well as aspects that are working well and should not be changed. NIST wants the community to have input as to when is the right time to update SCAP. A major risk is that if SCAP updates too quickly, it creates instability. The validation program must update whenever the specification updates, because from NIST's perspective, the specification and the validation program are joined together. Continually changing the validation program would likely create instability in the marketplace. While reaching forward is good, it is important not to create change just for the sake of change. This is why it is important to hear from the vendors and users in order to know when the right time to consider another set of changes is.

John Banghart of NIST was asked, "Given the role the US government plays in SCAP, was NIST receiving inputs from the US government on what changes should be made to SCAP?" John responded that there has been input from the US government, but that the government is quite large with diverse needs and thus does not speak with a single voice when identifying what features are important and when they are needed. He argued that this is why there is a need for a forum for the US government to discuss and coordinate with regard to new features, not just for security automation but also for policy. For example, NIST, as part of the Department of Commerce, is very interested in teaming up with international standards bodies, because compliance with international standards brings greater access to additional markets and therefore more jobs. Other agencies are worried about defending the national infrastructure today, and might be concerned that operating through international standards groups could slow SCAP's ability to respond to immediate needs. There have been specific suggestions from the US government, for example, with regard to USGCB and its support, as well as this week's recommendations on making OCIL more useful in an enterprise setting.

Mike Kinney (NSA) observed that, while there are many things that people are talking about incorporating, the real issue is *when* they will be incorporated and when they will be mandatory to support. He noted that a few years ago, SCAP had a defined, 3-year development cycle, so asking "should we make changes" seems like a departure from previous plans. Mike also noted that internationalizing would likely both delay the adoption of changes and alter the US government's ability to influence those changes. He noted that there is an existing need to make changes because SCAP 1.2 does not meet specific, important needs, such as support for OCIL in the enterprise and support for remediation. While he noted that whether remediation becomes a part of SCAP is an open question, but he felt that the fixes to OCIL should be made, and should be made before SCAP was turned over to an international standards body so that OCIL was usable without waiting on these bodies to make changes. John Banghart noted that whether SCAP goes international is not entirely up to NIST, since SCAP is not copyrighted and anyone could bring it forward. He also emphasized Mike's point about remediation - that not everything belongs in SCAP. He argued this was a very important question that needs to be asked. He observed that there are many useful specifications that are not, and should not, be rolled into SCAP. He noted that SCAP was developed with a few, very specific use cases in mind, and that there was a need to avoid any perception that incorporation into SCAP was necessary to make the specification useful.

With regard to Mike Kinney's comments on OCIL, John Banghart observed that the evolution of any standard is a continual process and there will always be changes that people want made. However, there must be a balance so vendors are not dealing with a constantly shifting target. Mike agreed with that sentiment, but noted that there is currently no requirement in SCAP 1.2 to support OCIL, and thus the validation program does not test OCIL compliance. Given this, Mike argued that there would be no impact to evolving the OCIL specification in its own cycle, outside of the SCAP cycle. Kent Landfield (McAfee) emphatically agreed with that sentiment, arguing that there is a need to separate "security

automation" from SCAP. He observed that SCAP "has a box around it" in that it represents a specific point of integration. He argued that security automation, however, should take the approach that Mike Kinney had alluded to, where there were official certifications for individual specifications (covering the tool support for the specification and any necessary integration with other standards). Today's validation program tests for compliance with the SCAP box, not the individual standards. As a vendor, Kent wanted to see more things that were outside of the SCAP box. For example, there could be a standards-based remediation solution that had a validation program specific to supporting only that remediation use cases. He argued that if we continue bundling specification into one validation program, it will quickly become intractable to support. Instead, Kent argued for a more piecemeal approach to validation, focused on individual specifications and specific use cases, as being both easier to support and capable of supporting changes in a more rapid timeframe.

John Banghart responded that perhaps it might be beneficial to start looking at validation from a use case perspective instead of the current program which looks across the breadth of the languages. In the future, it might be useful to define specific functions in a continuous monitoring architecture and validate on an application's ability to support individual functions. However, John noted that he is somewhat restricted in what he can do by the OMB directive to support "SCAP" - anything that becomes part of SCAP (in any version) must be managed according to the OMB directives. The OMB directive may talk about a specific use case but the directive itself is a manifestation of government political policy rather than technology.

John Banghart went on to say that there is a need to evolve these standards in a way that is not US government-centric. While the US government supported the initial development of these specifications, there is a need to migrate away from that perspective so the entire community is not letting one group define what constitutes "validation" of this technology. Towards this end, there needs to be feedback from industry to drive the evolution of these standards. The OMB directive allows NIST to decide what goes into SCAP and what does not, and NIST recognizes that such decisions have a long-term downstream impact on what the US government does and must evaluate changes based on that understanding. However, the criteria by which such and evaluation is made will not be equally relevant outside a US government context. There is nothing that prevents the individual specifications from being evolved independently - NIST, in its role to support the US government, can simply take the evolved specs and compose them into a new version of SCAP as necessary.

Kent Landfield observed that the recently released SCAP 1.2 was really a major change with substantial changes that will take significant effort to integrate. In the past there was an attempt to make releases on a yearly basis, although that never quite succeeded. If the community wants the acceptance that we saw in 1.0, there needs to be feedback on new releases when they come out. We need to digest and integrate SCAP 1.2 before we have another round of churn of the specification. Making smaller, incremental changes would be fine so long as we avoid the huge changes (at least in the near term). The original yearly release cycle had a purpose in that the vendors knew when to expect changes. The problem this time was that the latest change was too much. He argued that the quantity of change was potentially more destabilizing than having smaller yearly changes. Mike Kinney agreed, arguing that there is no reason specs cannot evolve between SCAP cycles.

Mike Kinney also noted a need to prototype reference implementation, arguing that the lack of such an implementation was a contributing factor to the problems we are having with OCIL today. Kent Landfield agreed with this, noting that this is part of the reason the IETF requires "running code" with any new standards effort. He also noted that it was good for OCIL to be evolved outside of SCAP, arguing that this allowed those vendors interested in OCIL capabilities to pursue them without having a lot of

excess baggage to support SCAP as well. John Banghart agreed on the benefit of having sample implementations that are used in actual architectures in order to better target the needs of the validation program to the functions and interfaces such implementations demonstrated.

Kent Landfield argued that there is also a need to look at existing work that has been done outside of the SCAP community rather than attempting to create a home-grown standard for everything. He noted that both the Trusted Computing Group and the IETF have relevant work that could be leveraged. John Banghart agreed, noting that the validation program's derived test requirements could point to any suitably approved specification.

John Banghart also observed that people have been talking about being driven by use cases. Towards that end, he requested that people, especially from the US government, share some cohesive, formal use cases. He argued this would help define what everyone needs NIST to build. Kent Landfield noted that this would also be useful to spur adoption, since it is hard to sell a new specification within a company without any context, but including a well-defined use case that the specification supports can make the value clearer to decision makers.

Dave Waltermire noted that having a reference architecture to demonstrate how standards work together is extremely useful. He noted that the Internet was not created because of TCP or IP, but because devices were developed that allow these standards to be used together to meet a need. He argued that we have not really had this in the past a technical reference architecture would be a huge benefit to the community.

John Banghart argued that predictability was the real issue when people talked about cadence - that it wasn't so much that there was a need to have releases on a particular interval (since different groups would want different intervals) but that vendors needed to have sufficient advanced warning as to when there would be new releases. He felt that this was possible, but that democratic consensus needed to be built into the process. Getting a functional reference architecture and defined use cases will help make the process more predictable.

Tuesday July 10th

Standardizing Access to Organizational SCAP Content

Presenters
Dave Waltermire, NIST
Kent Landfield, McAfee

Introduction

Kent Landfield opened the morning session by presenting an overview of past content management, and highlighted a number of shortcomings with the current ways of accomplishing this. While the tools that run SCAP content adhere to a set of standards in order to provide a consistent and interoperable experience, the delivery of this content is left up to the individual organizations and/or each vendor. This has led to difficulty in acquiring and/or tailoring of content and also pain in managing larger security automation environments, especially heterogeneous ones.

The talks today are focused on the requirements to help solve the content management issues, not commercial repositories/aggregation, ownership, or vendor proprietary capabilities.

Current Problems

During this introduction, Kent mentioned a number of issues with the status quo:

- Dependency on single source of baseline checklists (NIST)
- Difficulty in assigning authority to content
- Multiple tools in an enterprise is difficult to manage without central content
- Content delivery is challenging and costly in dollars and staffing
- Organizations want to maintain control over content and distribution
- No ability to easily search for content
- Content is available from multiple sources, at potentially different versions

Repositories Direction

After highlighting some of the issues, Kent spent some time introducing some high level concepts that might help these issues:

- A team composed of several folks from the SACM group has begun work on a specification to standardize content repositories/delivery to allow better content management.
- Highlighting a long stated desire, the best candidate for writing and owning content for specific software products is the product's vendor.
- Organizational content serves should be defined to allow enterprise local content to be hosted internally for use by all of the standardized content for the given organization.
    - Not limited to SCAP
    - Provides authoritative (within the organization) content for the various tools
    - Allows querying for content

- Additional needs:
  - Global, automated content delivered to organizations
  - Guidance authors should be clearly understood and registered.
  - Capabilities for organizations to locate new content, provide existing locally defined content, and to ensure the correct versions of content is understood.

Discussion

During the presentation of this introduction, the following discussion points were brought up:

- It was mentioned that the National Checklist Program (NCP) isn't intended to host all of the SCAP content, but instead provide pointers to content. Additionally it was said that a specification around repositories would make the NCP more effective.
- The capability to capture metadata was brought up as an important feature of this work. Some of this is covered in a later section today.
- It was brought up that licenses and Service Level Agreements (SLAs) are an important aspect to be managed to make this effort successful. The presenters agreed with this point, but pointed out that it was more critical to deal first with the simplest content delivery case first.
- The overall process that would be put in place was discussed with some concern for how authorship would be assigned, especially in context of a modified check from some authoritative source. The legal aspects were also mentioned as an area of concern.
- The fact that not all SCAP tools support the same set of features (OVAL Tests, etc.) was mentioned as an impediment to this work. Additionally, the fact that some tools do not directly ingest SCAP content was brought up as a concern that impacts compatibility. The presenters agreed with both of these, but believed that as the work matures, much of this will be worked out.
- Lastly, it was mentioned that some of these issues had already been worked out by the MILE/RID communities and that those results may help this discussion.

Automating Content Data Exchange

Session Goals

Following Kent's introduction, Dave Waltermire spent the rest of the session discussing the technical challenges presented by the problem space defined in Kent's presentation. Specifically, he spent time discussing the various challenges presented.

The goals for the session were to start the development of consensus for the approaches laid out, to use consensus to start the writing of a specification, and to inform prototyping efforts of these consensus items.

Challenges

Dave went on to present each challenge and lead discussion. Due the volume of discussion on each challenge, the below breakdown includes the discussion of each challenge separately:

Challenge 1: Content Identification and Reference

The first challenge presented was the issue of content identification and reference. The issue here is that the numerous identification schemes used by the various security automation efforts makes it difficult to reliably request and identify specific content for retrieval.

Additionally, references to given identifiers present challenges with regards to things like versions, etc.

Three proposals were floated as possible solutions to this challenge:

- Normalize legacy identifiers – Use a standard notation to format existing identifier formats. This works very well for legacy IDs.
- Define new globally unique content IDs – Every piece of content would have its own newly generated ID. This works well for newly created content.
- Hybrid of both – By combining both normalized legacy IDs and new globally unique IDs, the best of both worlds may be achieved. This was proposed as the preferred solution.

Discussion

During the presentation, the following discussion items were brought up:

- GRC Exchange was mentioned as a possible transport protocol.
- There was some discussion regarding SCAP bundles and the ability to decompose the items in the stream. The presenters felt that was a viable way to get around having to bundle all of the content in the single package. Others in the audience disagreed that this was a good choice.
- During the normalize legacy IDs proposal, the namespace component of the legacy ID was asked about. Two questions were posed:
    - First it was asked if the namespace meant that the author must host a content repository. Dave Waltermire answered that it didn't necessarily mean that the author stood up a content repository, but rather that they simply had posted it in a content repository.
    - Additionally, it was asked if this meant that the author needed to know a lot of details about the repository. Dave answered that while an author needs to know some about the repository host that hosted the content in order to refer to it via legacy ID, it meant that the author need only know where the content is hosted, but no other details.
- It was mentioned that CIS used this type of naming scheme.
- What is meant by authoritative and publisher was asked. The presenters suggested that authoritative meant the best place to contact with questions, and that publisher meant the active supporter of the content.
- Versioning was discussed with Dave proposing that the underlying versioning scheme of the content would be used, while others suggested that it might be best to move to a more consistent versioning scheme. The presenters added that it could be suggested, but not enforced. Additionally, the idea of using dates as versions was floated. Ultimately, it was suggested that while this is important, the focus of this conversation was for the identification, not the versioning, so the topic was tabled.

- Additional conversation occurred around the usage of URLs in the IDs. It was pointed out that domain name changes could cause issues with this naming scheme. On the one hand, using the domain name kept a simple solution to provide unique names and to provide information about the source. This comes at the expense of being somewhat fragile, in the case a domain change occurs. Others suggested that you could decouple the ID from the domain using a registry, though this introduces further complexity into the mix. Lastly, it was said that one could use DNS redirects to handle domain name changes, but that this has a cost as well, since you need to continue to operate the domain name and it could lose some metadata context about the content's author. No consensus was achieved here.
- It was asked if specifying too much about where to find the content (such as the domain name, etc.) would create a situation that would make it difficult to stand up a local content repository to host an enterprise's content, specifically in the case of mixed publicly available content and locally defined non-public content. The presenters suggested that it might be problematic, and that later sections in the day would discuss content caching.
- The question about using the term URL, as opposed to URI was brought up. The intent is that the namespace be a URI, however it may also be useful for it to be usable as a URL.
- It was asked regarding the global IDs whether the GUID would be expected to change for different versions of the content. The proposal suggests that that would be the case, which was pointed out was different that many or most of the standards currently do.
- The question was asked if the legacy IDs were to be considered absolute identifiers, would there be a 'relative' option to avoid embedding location information in the ID. The presenters suggested that urlbase could help with this, though this doesn't entirely solve the issue of the author needing to know the content location. Others pointed out that the content doesn't necessarily contain the location, but rather a logical pointer to the content.
- Additionally, it was pointed out that by using the naming scheme(s) proposed, all content that uses the same namespace, must come from the same repository. The presenters confirmed that this is a base assumption.
- Some other points that were made:
    - The id type in the legacy ID provides the type of identifier that is embedded in the ID.
    - We could make the 'global' ID another id type for the legacy IDs, thereby combining the two.

Challenge 2: Supporting Varying XML Models and Model Revisions

The second challenge presented was that the set of security automation content that is to be managed by content repositories make use of varying XML Models, and that further these models are modified from version to version. Further, the retrieval and management of such content is often tied to the XML model.

The proposal that was made to solve this issue was to create an abstraction layer over the XML models for the content. This would hide the specific XML Details for each standard from the management of the content. Using a metadata-driven approach, the abstraction layer defines Entities and Relationships,

with Entities capturing important model constructs, and Relationships capturing the referencing of one Entity in another. Examples of each were given.

Discussion

During the presentation, the following discussion items were brought up:

- The question as to whether metadata aside from the raw XML was stored as part of the data model. Dave responded that in fact, a large amount of metadata can be extracted from the content.
- It was asked why we need such an abstraction layer. The presenters answered that by using this type of metadata model, fewer coding and infrastructure updates would be required for each revision of the standards.
- It was suggested that RDF (Resource Description Framework) might be able to capture the required metadata and may be able to be used in place of defining a new language. Dave mentioned that the prototyping was doing exactly that under the covers.

Challenge 3: Content Reuse

The third challenge laid out by Dave was the issue of Content Reuse, which is the concept of making use of already written, existing content in the context of new content. With no formal mechanism for searching and maintaining content repositories, content reuse is difficult to accomplish.

In addition to the difficulties in acquiring and making use of existing content, the content formats themselves can prove challenging. Often what is needed is the ability to make use of sub-components of a given document, not the entire document. However many tool are not designed to work in this way.

To solve this problem, the presenters proposed supporting tailoring of documents, much like XCCDF allows as of XCCDF 1.2. Additionally, standard interfaces for content querying would be exposed to facilitate the reuse of content by making it easier to locate.

Discussion

During the presentation, the following discussion items were brought up:

- It was asked if this could be expanded to capture additional metadata about the content available via search and how the search capability would handle searching for embedded text. Dave suggested that additional metadata could easily be captured and exposed by this method and that full text searching would accomplish some of the more challenging search criteria.
- Several questions regarding Intellectual Property (IP) in the context of reuse were brought up. The presenters felt that this would be an area that needed more conversation, but that was certainly important. The goal here was to discuss inputs and outputs not deep technical details.
- It was voiced that some of this metadata was crossing into implementation, as opposed to interface. The presenters believe that it is important to define a common language here to support the use cases.

Challenges 4 & 5: Content Integrity & Content Confidentiality

At this point in the presentation, time became short, so there was a brief overview of the 4th and 5th challenges. The 4th challenge dealt with content integrity. There are several layers of integrity in question here:

- Data Model – This is the integrity of the actual underlying content model.
- Persistence – This refers to the integrity of the content within the data store. In other words what is stored in the data store is the same as that which comes out of the data store.
- Communication – This is integrity of the data between two communication points.
- Transport – This refers to the idea that the content you sent is the same that was received.

The proposal here was to leverage the existing integrity schemes at each of the aforementioned levels, including XML Digital Signatures for the Data Model level, TLS for the Transport level, etc.

Due to time no substantial discussion occurred regarding the 5th challenge, Content Confidentiality.

Discussion

During the presentation, the following discussion items were brought up:

- It was asked how the data integrity issues apply to byte streams. The presenters mentioned that in the prototype work, they used an RDF store as well as an XML store and used placeholders to achieve decomposition/re-composition. This code for this is available on Google code.

Content Repository Interface Discussions

Presenters
Dave Waltermire, NIST
Adam Halbardier,Booz Allen Hamilton

Introduction

Adam Halbardier from Booz Allen Hamilton, continued the morning discussion on content repositories with a discussion regarding interfaces that have been proposed as part of the content repository specification that is under development. He opened by highlighting that the intent of this section was to foster discussion regarding the different proposed interfaces.

During his opening comments Adam asked the community if they felt that Content Repositories was a problem that needs to be addressed soon. The general consensus was that it should be addressed soon, though there were several questions as to whether a clear problem statement was made here. It was also pointed out that vendors have long product lifecycles, which means that getting these things defined sooner than later would be beneficial.

Additionally, it was pointed out that many existing tools can be used to accomplish a lot of the content repository capabilities, without the need for creating new specifications, etc.

The question of how much must be addressed at what stage was asked, with some feeling that we need to focus on the simple case of a client talking to a repository, while others believed that it was difficult to ignore cases like repositories acting as clients to other repositories. While no conclusion was reached here, the presenters felt that it was appropriate to begin the discussion around the interfaces that might be part of a solution.

Challenges

Adam began by highlighting some of the specific challenges that drive the requirements for the interfaces that are discussed here. He mentioned distribution, content reuse, interoperability, and revision management as driving factors in the interface design. These issues echoed some of the points brought up earlier in the day, but from a deeper technical aspect.

Additionally, a set of basic assumptions was laid out, including that REST services would be use, that XML Digital Signatures would be used to sign content when appropriate, and that clients to the services would need only process content delivered, not knowing the implementation details of the repository.

Discussed proposals

Due to time constraints not all of the proposed interfaces were discussed. The following captures the high level discussion points that were discussed:

Interface 1: Retrieve

The first discussion focused on Interface 1, for retrieval of content. The interface allows making a retrieval request for a single piece of known content by its model ID, for instance a single XCCDF Rule, or an OVAL Definition. Several parameters that might be defined for such an interface were discussed including a flag to determine whether or not to include relevant metadata or not, as well as a concept of depth, which referred to how much of the dependent content would be retrieved along with the primary piece of content, as opposed to using placeholders for non-retrieved components using x:include.

The discussion included some notional XML examples of what this type of query might look like, and included examples of retrieving by both the legacy IDs and using global IDs (as discussed earlier in the morning).

Discussion

The following discussion points were raised during this section:

- Some specific questions about the interface were asked:
    - If you request an XCCDF Rule using the highest level of depth, would you get the OVAL Definition as well? The answer was given that you would get the entire XCCDF Rule and its XCCDF dependencies, but not the OVAL Definition, though the OVAL Definition ID would be part of the metadata.
    - If one is tasked with creating a new benchmark, how would I discover existing content? This is a complex use case, partly supported by the Search interface.
    - Would one need to call out an SCAP Stream's ID to retrieve the content? The presenter responded yes to this question.

  - What would be the use case for calling the interface with metadata="false" and why is that the default? Typically, when a client does a retrieve, they will know the details about the content, and need not retrieve the metadata. It defaults to false, as that would allow a dumb client to do an HTTP GET for the content and be able to successfully retrieve and make use of the returned content.
  - If requesting all of the content for a specific item, would I expect it in document order? Adam confirmed that it would.
  - Would the only type of content supported be XML? The presenters confirmed that XML would be initially all that is supported, but that in the future other types may be considered.
  - It was mentioned that by requiring a specific URI for each piece of content retrieval, it makes it more challenging to acquire more general types of content, like "all of the USGCB content for Windows 8". It would be nice to have a way to more abstractly request content. The presenters agreed that this is a concern, and agreed that policy must be considered here. It was suggested that search could help, but not solve this issue.
- There was also some discussion around the types of clients that would be expected to use this interface. While the support of "dumb" clients (those that understood little about what they were asking for) was expected, it was also pointed out that many clients would be quite intelligent and know exactly how to handle the content retrieved. The presenters agreed that both cases need to be supported.
- It was mentioned that there was a desire for this to be more abstract, to allow a more general content request. The presenters understood this, though would need to discuss it further.

Interface 2: Search

The next discussion was around the Search interface, which is designed to allow robust querying of a content repository. It is expected that this interface would be used to discover the metadata regarding a set of content.

The proposed solutions for this use case included a short term possibility of creating a new query language that could be marshaled/unmarshaled via JSON or XML, supporting querying by attributes as well as relationships, and a longer term option to build a search engine for repositories. Examples were shown of what these queries might look like.

Initially only the case of querying a single repository is being covered.

Discussion

The following discussion points were raised during this section:

- There were several comments about the proposal to create a new query language around content repositories, with the general consensus being that if possible, it would be better to use existing technologies to accomplish this. It was pointed out that this query language is acting on the Metamodel, as opposed to the actual XML content, which makes using things like XQuery and/or XPath more difficult.

Content Caching

After skipping over some slides for time, Adam continued the discussion talking about caching of content. The section laid out a few different proposals for caching, including the pros and cons for each. The proposals included:

- Use an HTTP Proxy server to cache the content
- Use a dedicated local content repository to cache the content (recommended approach)

Discussion

The following discussion points were raised during this section:

- It was pointed out that using TLS one cannot cache the data. The presenters agreed and said that it is a challenge, and that XML Encryption could help here.
- Additionally, it was pointed out that perhaps a hybrid solution may be appropriate here to allow for different use cases. The presenters suggested that that would introduce an additional level of complexity.

Interface 3: Metamodel Exchange

The final discussion of the section revolved around the issues of how a client could understand the Metamodel used by a given content repository. The proposal made was to:

- Allow each content repository to support one or more Metamodels.
- Each of the Metamodels must be updated in a backwards compatible manner.
- The retrieve interface response must include the name of the Metamodel used to create the metadata returned.
- An interface will be designed to allow retrieve of a Metamodel.
- Each Metamodel would have a unique ID.

Discussion

The following discussion points were raised during this section:

- It was asked if this approach seems reasonable. Several folks mentioned that while this was reasonable, it was also difficult to implement.
- It was also asked if the backward compatibility requirement would inhibit evolution of the Metamodels. The presenters suggested that one could deactivate an old Metamodel in favor of a newer one as one way to deal with this issue. It was pointed out that would force a maintainer to continue to support the model. In the end it was mentioned that really the coordination of such an issue would need to be dealt with as a community.
- It was asked where the discussion for these topics would take place. The presenters suggested that the SACM list would be an appropriate venue.

gOCIL Interpreter Demonstration

Presenter
David Ries, gOCIL

Introduction

gOCIL is the OCIL interpreter that they have recently developed. The current state of OCIL is that there is very little adoption, very little content, and very little tooling of OCIL.

Two questions came up during the discussions on Monday:

1) Is the OCIL schema ready for prime time? Does it support real-world, enterprise use cases?
2) What would an enterprise OCIL deployment look like?

David stated that they had the same two questions, and this project was a practical exploration of answering them.

David's discussion of OCIL concentrated on compliance, e.g. STIGs, as this is the most obviously applicable use case for the attendees at Dev Days. Some obvious places where OCIL is useful, e.g. manual checks that can't be automated. For any comprehensive assessment of compliance status, one needs to complement OVAL checks with OCIL questionnaires.

David described the goals of the gOCIL project:

1) Support the full OCIL 2.0 schema, be true to the OCIL standard, no proprietary extensions.
2) Friendly, usable user interface.
3) Enable interface to the OCIL interpreter with an XCCDF interpreter or a host application. gOCIL is available via a full featured REST API for transaction management
4) Offer flexible deployment options, such as on-premise or accessible via the Internet with a SaaS model and can be embedded or white-labeled

gOCIL Architecture

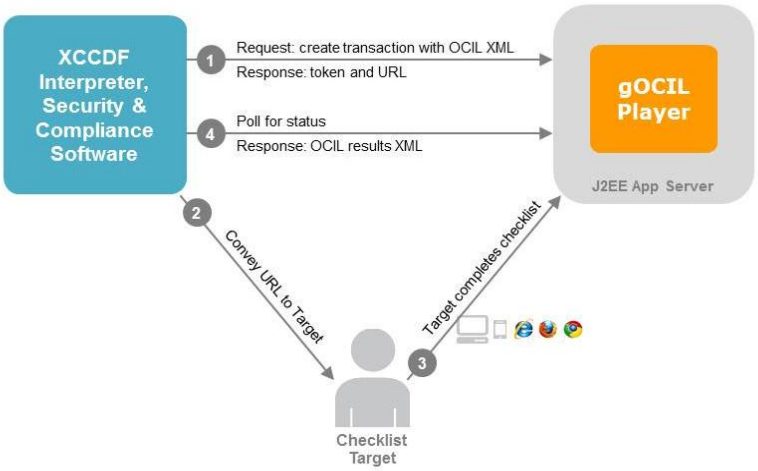

Three main players: host application (e.g. XCCDF interpreter), checklist target, and gOCIL Player (built as a WAR file). The gOCIL exposes two interfaces: an API and a web interface.

Q & A

Gary Gapinski asked whether the survey solely interactive or could it be accomplished by another means, such as form-based email. David replied that the gOCIL Player exposes the survey in a web page, so he didn't think a form-based email would support it, but it could be possible. However, there are no plans to support such a feature.

In response to another [inaudible] question, David replied that you needed a J2EE app server running the gOCIL application.

Gary asked a follow-up question about which process keeps track of outstanding evaluations, and David replied both the host application and the gOCIL Player.

Another [inaudible] question was asked. David explained that when targeting is talked about, three items are usually referred to: audit trail information which can be embedded in OVAL documents or OCIL documents – basically the log of what happened; applicability information which exists at the XCCDF level; and then information which identifies the target person and their role, which takes place at a level of a system which understands organizational assets and is out-of-scope for XCCDF.

David then demonstrated the gOCIL Player and gOCIL.org.

Mike Kinney asked a question about aggregating result and David responded that results aggregation would need to be performed by the application, not the OCIL interpreter.

Jason Mackanick mentioned that OCIL content developed by DISA make use of a Responsibility field which specifies the person or role that is responsible for answering the questions. Jim Ronayne suggested that this issue be discussed during the OCIL session on Thursday.

Justin Furniss asked how it is handled if there two checklists assigned to the same individual. David described that gOCIL assigns one transaction at a time to a single person. But they both agreed that the application could handle the workflow management in a variety of ways.

A questioner asked how the results are extracting from the interpreter, and David explained that the results are embedded in the document. He went on to explain that the API provided provides features to extract the data. The three REST API calls are: one creates a transaction which gives you the token, the second checks the status do your application can poll for status to determine when it's been updated. When status request returns the state "submitted", which means that the checklist has been completed by the participant, then you can make REST request to get the checklist which will provide the results.

Security Automation and the International Community

Presenter
Kent Landfield,McAfee

With SCAP gathering international interest, there have been several movements towards internationalization of the standards efforts. The SACM working group will be introduced, and the IETF, ISO, and ITU-T international standards bodies will be addressed as possible candidates.

Reasons for Transitioning to International Development

It would be much more desirable to expand the security automation efforts to a much broader audience. These other countries have the same problems that we are having with security automation, and it is necessary to have a global understanding of these foundations. To expand the audience, we must address the issue that the current perception of SCAP seems largely a US-only effort. With that belief, they exclude themselves from feeling responsible to participate.

Additionally, in the development of the specifications, there have been too few efforts working on prototyping along the development cycle that may lead to revisions later on. Some groups require working code as a maturity indicator. In order to get more people looking at the specs, there must be a larger number of eyes looking at the problem. The current model of development has shown that most do not participate in creating new features. Rather, people wait for someone else to introduce the innovations. Vendors should be taking a larger part in pushing forward with ideas. We need to be able to create derivative standards, and for those standards to be normative references. Existing work from other countries should also be recognized so that we do not have to reinvent the wheel every time the same problem arises. It is not the individual work of these efforts that is important, it is the end goal.

Timeline up to present

At the 6th ITSAC, several vendors expressed their feeling that nothing was happening with the security automation efforts and that it was time to start a working group to get things moving. In November of 2012, at the IETF meeting in Beijing, an educational BOF was intended to raise awareness of their intentions. From this perspective, it was successful in getting others talking about SCAP, including mention at the MILE working group. However, this meeting was interpreted as a working group forming BOF.

Following this meeting, many vendors started discussions on what would be important in organizing potential ways of getting into the international community. Research was done on several international standards bodies. At the 7th ITSAC, all the vendors voted that the place to work would be the IETF.

A letter was drafted to NIST, NSA, and DHS to move the automation efforts into the international direction, signed by 16 SCAP vendors, mostly CTOs. The letter was a recommendation to move technical specifications to the IETF. It was positively received since NIST was at the next IETF meeting as a presenter.

Security Automation & Continuous Monitoring (SACM) and SACM timeline

The proposed SACM working group is intended to continue to grow the foundation laid by SCAP. Potential work areas have been discussed but not finalized in a draft charter yet. A subset of these potential work areas includes semantic vocabularies, remediation, content distribution, and enterprise reporting. Real work is already underway, through draft documents for use cases in security automation that ties in these efforts to existing IETF working groups such as NEA and MILE. There has also been a content repository started.

At the Paris IETF meeting, a side meeting was successful in creating actions that include drafting the use case document. At the Vancouver IETF meeting, it is desired to discuss the potential charter required for the working group. It was added by Kathleen Moriarty that the SACM mailing list was a good place to voice opinions on what use cases to target, and that by doing so the IETF can see that the discussions are active. This would be desired before the Atlanta IETF, where there will be a potential SACM working group forming BOF. The mailing list currently has a sizeable amount of traffic to date.

ITU-T

The ITU-T is a UN based body and has been working to create a Cybersecurity Information Exchange (CYBEX). The X.1500 series utilizes SCAP component pieces including CVE, CVSS, and CWE. When dealing with the ITU-T, the previous work they've done with the related standards has been beneficial, but it would be better to have these standards in a technical body like the IETF.

Kathleen Moriarty talked to the work going on in the ITU-T in relation to CYBEX and MILE. In general, the CYBEX framework got visibility by tying to the existing specifications, which brought it to the attention of many countries who were previously unaware. The work from MILE was represented, but technically inaccurate. From working with the IETF and the ITU-T, the resolution was to have a pointer in the ITU-T document to the IETF standard, such that nothing was transferred to the ITU. This was done since anything given to the ITU has them assume copyright. In using the pointer, this left change control within the IETF.

ISO SC27

XCCDF 1.2 has been submitted by NIST to the ISO international standards body with a request for the IETF to be named the maintenance organization moving forward. This allows the specification to be fast-tracked due to the IETF's relation with the ISO. Currently the US CS1 national board has a ballot underway questioning whether the XCCDF specification should be submitted to the SC27 to become a standard within ISO. The CS1 recommended a YES with COMMENTS vote to the executive board, with the comments being that the standard should be freely available. It is likely the US will approve of the vote. If the vote fails, then the fast track is over. The vote will be known in August.

Some discussion followed that questioned the ability to make changes going forward. Since IETF was the maintenance organization, then SACM would be the place where XCCDF development moving forward would occur. A question was asked about the relation between the IETF and the ITU. These two bodies have a liaison agreement, which is a technical agreement to recognize standards from the other body much quicker. In this specification, the IETF would be an extension of the ISO, with any new revisions

pushed back to the ISO. There was another question about the speed at which a document could be revised. It was said that revisions get pushed every few months.

In addition, it was asked that once moved to the IETF, would the standards be for sale. The standards would be freely available in this case. Normally the ISO charges unless a request has been made for them to be freely available. Another question was raised about the intellectual property issues of IETF. The IETF considers all submissions to be an "IETF contribution." All patents must be documented in relation to the submissions.

TCG TNC and SCAP

The Trusted Computing Group is working with network access control and looking at SCAP into their environment. They wish to use SCAP to do health checking on machines.

Discussions

One question was raised about what traffic should be switched to the SACM. Since this is not defined yet, it was determined to do so as appropriate to the community.

Steve Hanna shared his experience with the TCG in moving standards over to the IETF. They had determined that some of the standards they had should be moved to the international standards area to be more broadly adopted. The ones that were well understood made up the subset that was moved. They started the NEA working group and transferred the change control of the standards to the IETF. They found that the same people active in defining the standards in TCG remained active in defining the standards within IETF, along with the global community. TCG continues to develop new standards as extensions. Out of their dozen standards, 4 were moved over with the rest to follow as needed.

CEE

Presenter

John Wunder, The MITRE Corporation

Introduction

CEE is a logging standard and the website is cee.mitre.org. The current version is 1.0 Alpha. Today we will be talking about 1.0Beta. John Wunder is now the CEE Project Leader and has recently rejoined MITRE. We will first do a brief introduction to CEE and review some terms. An EVENT is a single occurrence within an environment, usually involving an attempted state change. A user attempting to change his password is an example of an attempted state change. An EVENT RECORD is a collection of EVENT FIELDS that, together, describe a single event. A LOG is simply a collection of event records. We often say EVENT to mean EVENT RECORD.

Design Goals

Compatibility is our overarching design goal. We are trying to write into the existing practices. There are many applications running logs and so it is very important that CEE work with existing products. We are

not going to ask developers to re-write their products to adapt to CEE. Efficiency is another important CEE goal. We use open standards whenever possible.

Consuming Events

Receiving a raw event from the wire is the first step for consuming events. Next, the data is PARSED so that the individual fields can be located. You need to understand what the event means. Next, the event is normalized and analyzed. The final step is to present the results to the consumer who will do something useful with the results.

Problem

The Problem today with current event logs is effective analysis to make a decision. The difficulty stems from the vast number of different formats. All the vendors do things differently. Parsing events is difficult. Comprehending the events is even harder. There is often no reliable way to figure out the event type or even what the event means. Lastly, today's event systems lack security and resiliency of log protocols. The UNIX world is using syslog which is showing its age.

Solution

First step, for receiving events is to provide the Common Log Transport (CLT) which describes how you want to receive events from the wire. It is basically the mappings to several existing log protocols. The Common Log Syntax (CLS) describes an event at the markup level, which turns into mappings to XML and JSON. The CEE Dictionary and Event Taxonomy (CDET) is the list of fields that matter for an event and the classification of an event. For instance, logon name is a classification. Lastly, the CEE Event Log Recommendation (CELR) is the case where if you have a regulatory body you might have everyone under that body to comply with your recommendations.

The New Approach

The community has decided that the event dictionary, taxonomy and requirements languages are similar and could be put into a thing called a CEE Profile. Some fields are required and some are optional. With CEE profile we get a shared vocabulary and shared taxonomy. Now the comprehension phase is separated from the requirements phase. At a high level an event is produced. The CLS will encode the event into either XML or JSON. The CLT will map the event to an existing profile. We have a shared understanding using CEE profile of what the event means, a shared encoding using CLS of how to represent the event and a shared transport using CLT.

Event Modeling

An event is set of Fields and Values for those fields. Events are just a series of fields and tags.

The Field is a name and value(s) associated with an object or property of an event. The Tag is the event "type". Action tags might include login, remove, read, block, or search. The status tags might include success, fail, and error. Other tag examples are HIPAA, audit, critical, warning, and info. Conceptually, tags are just fields with enumerated values used to classify events. In addition we have the classification into taxonomy.

CEE Log Syntax

The CEE Log Syntax is the description of structured encoding for Event Records.

CLS Overview

The CLS Event Specification defines a generic CEE Event Record Structure. The CLS Encoding Specifications define encodings to and from this structure to two common formats: JSON and XML. Both formats represent the same data but have some differences. It is relatively easy to convert JSON to XML or to convert XML to JSON. Both work well with hierarchical structures. There are different use cases than lend JSON to be better than XML. On the other hand XML may be necessary to facilitate other use cases, such as audit compliancy. In the future, we plan to write a tool to convert from XML to JSON and the reverse.

Event Record

The Field list may be derived from the CEE Field Dictionary, referenced profile, or custom. Fields may be optional or required based on profile. Additionally, you can have the optional classification into the CEE Taxonomy. Fields are either required or optional depending upon which profile you are using.

Event Comprehension & Analysis – CEE Profile

This is where we get into the CEE profile and get the most value out of CEE. This is where we develop the shared vocabulary. CEE is a shared syntax. It's a common way to represent events. The CEE Profile Specification documents the features and usage of a CEE Profile document. The CEE profile is divided in to a couple different areas. We have the actual specification of what a CEE profile looks like. This is how do you describe events that are going to conform to a profile. This is how you define the list of field and whether or not they are required or optional.

The CEE Core Profile is the official field dictionary and taxonomy. In itself is a profile. The CEE Core Profile is the "base" profile that all others inherit from. The base profile is how CEE provides a vocabulary for very common event types.

The CEE Profile Repository is the collection of CEE Profile XML Documents. We haven't put the CEE Profile Repository on-line yet. We would like the vendors to contribute to the repository.

CEE Profile Types

The CEE Core Profile is provided by CEE and contains core field dictionary and taxonomy. This is how CEE provides event types for very common activities.

The Function Profile defines the event profiles for events associated with a specific function. Examples might be Firewall and Session Management Profile. We are encouraging functional profiles because that's how we gain a shared understanding within CEE. The Product Profile defines event profiles for events that a specific product may generate.

John Field, of EMC, asked how many function profiles exist? John Wunder responded that today, no function profiles exist. We're currently exploring how to do a function profile specification. Once we're

out of BETA we will encourage vendors to step up to the plat and provide functional profiles. We may also set up a function profile repository.

An observation was made by John Field: "It seems that some products might contain multiple profiles." Tom Graves and John Wunder concurred. John Field suggested that we keep the focus on the technical controls within the standards body. "Think in terms of the profile for the specific capability. Not so much on the regulatory controls. Regulatory controls can change over time."

Sharing CEE Events - Common Log Transport (CLT) CLT Overview

Sticking an event on the wire is the lowest level and this is where the Log Transport comes into play. This specification is divided into two areas. The first is CLT requirements. Different requirements give you different capabilities. Some use cases care about message reliability. Other use cases require message integrity or chain of custody needs. CLT has a set of mandatory and optional requirements. For instance, with UDP message reliability is not guaranteed. CLT has four conformance levels (0-3). Level 0 is the mandatory set of requirements. Level 3 gives you data integrity back to the issuing source.

The 2nd important aspect of CLT is the CLT Protocol Mappings. Today we have mappings to the syslog (RFC3164, RFC5424) specification. Syslog is in IETF protocol. This describes how to send CLS Encoded CEE Events over specific protocols.

CLT Protocol Mapping

Specification defines how to encode a CEE Event and transmit over a protocol. Currently, we have the CLT Mapping Syslog supported in CEE. CEE events are encoded using the CLS JSON Spec.

Today we are looking specifically at syslog transport. Basically this describes how to encode a CLS event into a syslog message and then use syslog to transport the message over the wire. We stick a JSON encodes CEE event represented by the @cee: flag. Then we stick this at the end of the syslog.

Full chain from producer to consumer would use Syslog to transport, CEE JSON encoding, and the CEE base profile to represent the data.

Next Steps

Alpha is on the website now. Beta will go up soon. Architecture and syntax are close to finalized. There is still work to go on core spec for 1.0-final. CEE Core Profile content (field dictionary and taxonomy) is still under development, especially, the taxonomy parts. People can participate now to help build the taxonomy. More work on profiles is needed. Profiles for popular use cases, products, and audit requirements are needed. We also plan to identify further CEE Profile requirements and iterate specification. We believe we need more CLT Mappings. Today we have only syslog. We are looking for a higher assurance protocol to support validation, high assurance, etc. We are looking for something that can transmit XML.

Ongoing Activities

An adoption program is important to the success of CEE. We need support in vendor products. Training materials and publicity are important.

Today, we have a reference implementation call Lumberjack (libumberlog). This is mainly syslog work using CEE compatible JSON.

In addition to adoption, we are looking at a validation program. This is difficult because of JSON. It is difficult to validate JSON. XML is easier to validate. What do you validate? How do you validate? Bill Heinbockel mentioned that although JSON to XML and XML to JSON conversions are possible, there is no agreed upon conversion standard.

A versioning policy is important in order to support backward compatibility. We need to address how the spec will be versioned. Should the Core Profile be versioned with the spec? There are benefits if we version the profile outside the core specification.

As stated earlier, we do plan to stand-up the profile repository. We've always had plans to go to the IETF with this standard. We are talking internally at this time.

Enterprise Asset Reporting

Presenters
Adam Halbardier,Booz Allen Hamilton
Dave Waltermire, NIST

Note: Only discussion points are captured here. Briefing content is not captured.

Q: [On Asset Identification] So the uniqueness is guaranteed by the couple of the namespace and the identifier, or is the identifier supposed to be globally unique?
A: The synthetic identifier is the tuple of the namespace and the identifier. The identifier is unique within that namespace. The idea behind this is that the namespace can be controlled by the organization that defines the namespace.

Q: Is there guidance on whether or not those identifiers must persist across time? One of the vendors we are working with is using IP address as an asset identifier. When a device changes IP addresses, they just re-associate all the old data so that you can't rely on IP address as an identifier.
A: I don't think there is guidance to that effect. That sounds important for an update to the spec. There is a fundamental assumption that identifiers should persist over time, perhaps forever. Perhaps that needs to be added to the spec.

Q: We ran into a problem where we migrated our OS, effectively making the systems completely new systems. However, the hostnames all stayed the same.
A: This is something that makes sense to talk about in the future.

Comment: I think that it's important to keep identifiers unique and around for a long period of time. There are times when you need to keep historical data for trend reporting. Relating to synthetic identifiers, I think those should be assigned by the users that are using the system. Usually the

customers tend to think about assets like “I own this server”, or “This is the server for HR”. Synthetic identifiers should be marked by the users.

Comment: We are trying to move other vendors to using synthetic identifiers that are assigned by the tool. There’s combination of hostname, DNS name, MAC address, the GUID of the OS. If you can look at those and say that there are 4 of those that persist across time, that’s the same device. If you can come up with a matching on that, then assign that a synthetic ID, there’s a good chance I’ll use the same ID every time.

Q: [Inaudible]

A: They can have their own identifiers. We talk about software identifiers in the spec. You are not restricted to one or two identifiers. You are not restricted to any number of namespaces.

Q: We are using the synthetic ID to direct OCIL questionnaires. We use primary keys when sending that. We should make it standard to do something like that.

A: Excellent, thank you for the feedback.

Q: Will there be an opportunity to open ARF back up? It’s nice that it’s there, but because we haven’t defined any other data structures, we can’t use it for interoperability. There is no object model behind ARF. It leaves me kind of confused on how to use it. What happens when I package all this “stuff” and send it to somebody?

A1: There are a couple of things that need to be done with ARF. SCAP 1.2 defined a specific way to use ARF to ensure interoperability.

A2: If you need to express something using ARF, either you should define a format or you should seek a standards development organization, or we can work together to define a format.

Q: Is that something I put out on the list?

A: I think we should probably start with the use cases you are trying to address, then talk about what solutions you could use to address those use cases.

Q: I thought we agreed that the device model would be the next thing for the asset development group.

A: Thank you for reminding us about that. That is something we can work on.

Q: I would like to point out CIM as well. At the risk of prefetching the answer without a use case, if you are describing assets CIM may be something you want to use.

A: CIM is not SCAP aware. We are going to have to effectively rewrite CIM to get CIM to do what we want.

A: AI can be the bridge between CIM and SCAP. AI is a perfect fit if you can do AI on SIEM.

A: The BMC products were one of the ones we did cross walks with. We looked at all the device models we could find. Every place we could go, we took all the data elements and lined them up and did a cross walk. Then we pared it down and added SCAP stuff.

A: That’s why the CMDBs have all adopted some kind of federation or another. CMDBs only need enough information to de-conflict data.

A: We should probably have some conversation about this on the list.

Q: [Inaudible]

A: In this case there are three attributes from each record. What these attributes mean is defined in the record set type. The one option you have here is each record can have a list of assets. Assets in the list would refer to the list of assets in the data source. The assets listed compose the count.

Q: Was GRC the only thing you looked at?

A: We've explored using WSNotification as well as WS Make Connection and a few other web service features for Continuous Monitoring work. One of the challenges we started running into is that in general SOAP web services are not very popular. They are seen to be relatively heavyweight. We were looking at other approaches, I found out about GRC exchange with some of the work we are doing with MILE, and it seemed to be a pretty good fit. GRC exchange is essentially a very extensible container that we can use to support a variety of different continuous monitoring applications.

Q: My fundamental problem is that we seem like we aren't being prescriptive enough to drive interoperability. It's like saying "For exchange, we are going to use XML" and expecting it to work. I am hoping that we can do something a little more deterministic.

A: I understand and agree. We are going to show how we are going to use specific data models to solve specific use cases.

Q: For GRC, are you required to have a request before you send a report?

A: There are three communication patterns. One of which that allows you to send an asynchronous report to an endpoint.

Q: Could you use XML/TLS? HTTP is just going to encapsulate the XML.

A: Yes, you can use XML over HTTP over TLS.

A: You could set a content type of text/xml and handle it that way.

Q: Had you thought about using HTTP posts and a RESTful web services?

A: That's another possibility. We haven't really achieved any consensus around what the transport protocol should be. REST services are still HTTP.

Q: Some of those things are redundant, then you add layers of complexity. I'm not sure that we always get something out of all the extra complexity.

A: The thing we get is the rigor for specifying the communication flows that aren't specified in the underlying data specifications.

Q: Is GRC going to request retransmits when data is not delivered?

A: At this point, the basic transport protocol is undefined, if you talk about ASR. If we are talking about interoperability, how do we define what the transport looks like?

A: It may be important to address the retry pattern.

Q: Could you have more than one piece to a GRC report, or is it always a single transmission?

A: That would be handled in the transport layer. GRC handles the policy aspects.

Q: Do we have the maturity to leverage that (The policy aspects)?

A: No one is suggesting that you have to.

Q: Suppose we build GRC into our standard, it costs overhead. Is this another thing that we have to spend resources to implement? We need to believe that there is something of value in order to build this in.

A: If we want to build interoperability, we need to specify a transport protocol.

Q: Other than GRC, we could just do an HTTP post.

A: Just pointing to a REST service doesn't define how you do enterprise reporting.

A: Some of the other advantages, you can do digital signatures/encryption. There are use cases that are further along, like the reporting to regulatory bodies. The other thing to remember is that this is new and can be totally reshaped. The reason that the transport bindings are in a separate draft is so that transport can be changed.

Q: When do you expect publishers and vendors to have this built into their process?

A: I think the context of this presentation is that we are putting these ideas out there, a reasonable way forward. We don't have a firm timeline on the rollout.

A: To a large degree, the things that go on in the working group aren't up to me. We are going to be moving forward with some proposals over the next couple of months.

A: I can't speak to what vendors are going to do. You'll have to ask them.

Q: I thought you said you were going to go final in the next few months.

A: ASR is a completely different topic than this.

Q: ASR will be specified as XML?

A: Yes. ASR is pretty much done.

Q: Where would this GRC report fit? Is it an SCAP thing? Is it a Continuous Monitoring thing? You have it listed under enterprise reporting.

A: I think it would fit in Continuous Monitoring. It fits into this notion of enterprise asset reporting, mainly as a format. Now we are talking about protocols for exchanging the information across the enterprise, which hasn't been specified anywhere in SCAP.

Q: So I look and I see 'use HTTP/TLS to send XML' and I was just looking for a definitive answer on "This is a proposal for SCAP", or "This is a proposal for Continuous Monitoring".

A: There will probably be better context for this conversation tomorrow during the CAESARS FE conversation.

Q: Will ASR be allowed to be transmitted in a separate envelope, or will it always be in ARF?

A: This proposal is for ASR within ARF. If we went the other route, we'd need to register ASR as a schema for GRC exchange with IANA.

Q: So the first one says "Send report". What is the acknowledgement?

A: The acknowledgement could be an "I have it, but I have not processed it yet". In the event that the receiving system denies the report, the acknowledgement might be a denied message.

Q: Will vendors be required to implement these workflows in their business logic?

A: I would argue that we are not defining workflow, but more of a communication flow.

Q: It implies state management on the end system.

A: GRC exchange doesn't require that you reply with any acknowledgement. You don't have to queue information, you can respond to it immediately when you receive it.

Q: This may be a question for Kathleen. Can you compare and contrast, at a high level, GRC and RID?

A: GRC generalizes RID. We are hoping that the community will start commenting on the draft, shaping it up in a way that works for various use cases. GRC generalizes it and gets it out of the incident space.

Q: There seems to be a slight asymmetry when a command generates a request, since there is already a mechanism for generate a request.

A: The command would indicate that additional data needs to be collected, rather than just requesting a report.

Wednesday July 11th

Continuous Monitoring (CM) History and Directions

Presenter

Kent Landfield,McAfee

- How do you define continuous monitoring?
  - Different groups have different views
  - Presented definition from NIST (on slides). "Risk management" point of view of continuous monitoring.
- Reduce cost of managing infrastructure
  - More information about assets
  - Fewer required resources
- Plug+play is good, it enables competition because you can easily switch products
  - Drives more innovation
  - No sunk costs
- Will need to be phased in - longer term architecture but cannot be but in place immediately
  - Most of the day will be focused on those interfaces
  - Leads to no real aggregation & correlation of data (or, aggregation only to a point)
    - Instead, today we push a lot of buttons (i.e. use a bunch of different tools)
- Went through expected capabilities of continuous monitoring:
  - Asset visibility
  - Automated data feeds
  - Quantification of risk
  - Ensures continued effectiveness of controls
  - Remediation and prioritization of remediations
  - Empowers employees at multiple levels
- Started as a federal initiative, but not a federal problem
  - Which means developing solutions makes sense for vendors
- FISMA is not effective
  - Not accurate
  - Followed the letter of the law
- CyberScope is a baby step:
  - Didn't use approved standards, everything was draft or brand new
  - Needed it fast
  - Outside of normal product roadmaps
  - Frustrating to agencies (and vendors?)
  - Double reporting was possible
  - There was a lot of handcrafting of XML
  - But the intent was right
- Point from BAH: Agreed with everything, added that even FISMA-compliant systems aren't by definition secure

- Explanation of CAESARS, CAESARS-FE, etc. history:
  - CAESARS came out, CAESARS-FE identified gap and extended
  - CAESARS-FE allows for a tiered "continuous monitoring instance"
  - Dave Waltermire: "CAESARS-FE provides a lot more rigor around the specifications than CAESARS"
  - Problem w/ things being under-defined; vendors were not clear what to implement
  - NIST held vendor working groups to help clarify this, which led to development of the reference model
- Logistical problems slide
  - Large installed base, lack of plug and play, lack of automated reporting, architecture is more notional than reality
  - Must achieve some successes first...don't do it all at once
- CyberScope vs. CAESARS-FE slide
  - CAESARS-FE is more automated, CyberScope more manual
  - CAESARS-FE is focused at all levels, CyberScope more at very high level
- Poll: How many vendors are doing CyberScope? About 10ish. How many have updated from FY10 to FY11 metrics...none?
  - Point is that no vendors that he knows of have gone back to update to new metrics
- Mike Kinney, of NSA: it's only going to get worse, referencing CMSS and new requirements.
  - Kent Landfield:
    - Agreed, but only if we proceed incorrectly. In the past we got ahead of ourselves.
    - This provides new capabilities, it's basically investing in changing a model around network administration.
    - Vendors get a chance to revamp their product lines
- No guidance on what to do about CyberScope. Lots of information was specified that would be valuable, but the means to achieving it was not obvious.
  - Adam Halbardier, of BAH: Also not obvious which questions were actually required for CyberScope
- CAESARS slides
  - Goes through capabilities (see slides)
  - Points out that there are no asset management capabilities, need to fill that in
  - Mike Kinney: There's a requirement to scan every IP. Has that assumption been built into the document? Maybe other gov't requirements are missing from the framework arch.
  - Aharon Chernin, of DTCC: Need to define Continuous Monitoring. Commercial space is extremely fragmented, defined by vendors and organizations. Don't want people to get turned off to it b/c vendors are presenting it inconsistently.
  - Aharon Chernin: Risk is a loosely used term in CVSS and throughout CM. Commercial space has defined and mature risk organizations...unlike here. To create risk numbers you need something to do w/ threat.
    - Kent Landfield: Agreed, that's a big missing piece
- iPost
  - Succeeded in changing network posture
  - But also succeeded in changing the culture by changing the incentives
  - (May have been questionable from some perspectives [Editor: I think this is a reference to focus on vulnerability and compliance, not risk posture. Also didn't use standard data.])
- Jon Baker, of MITRE: Chief engineer at MITRE is involved through incentivizing compliance

- Mike Kinney: iPost worked because they had an organizational structure to support it. But in other places you don't have that, where the lines of responsibility are not solidified. Until that's sorted out you can't really do this.
- (BAH): Commenting on NIST CAESARS-FE and CM documents...it's about the history of CM. At DHS industry days they had a pretty clear picture of what needed to happen. Success depends on Security Automation community though.
- Jason Liu, of Northrup-Grumman: Helped State Dept deploy iPost...was network performance tool w/ security components added. Did include scoring though.
- Kent Landfield: if you're a vendor who's not participating, you're losing out
- Gary Gapinski, of NASA: You mentioned that some of ref. architecture is more notional than concrete...ITIL v3 is a reasonably good ref. architecture for config and asset management. Why is there no mention of ITIL?
    - Dave Waltermire, of NIST: Good observation. We're trying to define one level lower to promote actual interoperability based on ITIL.
    - Kent Landfield: agreed, we can't keep reinventing the wheel
- Joe Wolfkiel, of DISA: Is anyone working on how to actually identify and move organizational information (info management). What about tasking? Seems to be a lack of interest.
    - Kent Landfield...we're trying to get there…Joe is ahead of his time.
    - Joe Wolfkiel is frustrated with having to rebuild stuff when standards catch up
        - Vendors all chimed in that they have the same problem
    - Dave Waltermire: standards should be following, not leading. Will require some retrofitting. But yes, we'll need that kind of organization information
    - Kent Landfield: please continue to give feedback
    - Mike Kinney: How much more metadata could you handle than what you already have?
    - Joe Wolfkiel: That's not really the question...more, given the national focus the money will keep flowing in. The question is what is the MINIMUM metadata I need to make information make sense
    - Mike Kinney/Joe Wolfkiel discussion about whether you do things at the enclave level or higher
        - Mike Kinney: shouldn't we do this at the enclave level?
        - Joe Wolfkiel: Sort of. Locally there isn't an understanding of what individual IT assets are doing. Most tools don't tell you which assets support which missions, it's an ops problem at every level. Best DONE at a local level, best SOLVED at an enterprise level
        - Mike Kinney: Hasn't been required, but you can get to it.
        - Joe Wolfkiel: But it isn't required so you can't get interfaces off the shelf, have to custom-develop it.
        - Joe Wolfkiel: Just saying that we need to start thinking about that, and about how to use ITIL.
        - Dave Waltermire: And how to communicate that to the rest of the community. Bringing Joe's requirements to the world…
        - (Discussion continued into next section)

Continuous Monitoring CAESARS-FE Overview

Presenter
Dave Waltermire, NIST

- There was some continuing discussion into this session, notes are on last session
- Session will be an overview then a deep dive into individual components and interfaces
- Slide of goals and data domains
- Comment (BAH): CM is not replacing any data domains, it supports data domains (i.e. if you don't have good license management then CM can help, but will not replace that as a practice)
- In depth slide on goals:
  - Support decision making
  - Enable effective measurement
  - Correlate across data domain
  - Provide situation awareness at all levels
- Goes through the components one by one…
  - Presentation and Reporting
  - Content
  - Collection
  - Aggregation
  - Analysis and Scoring
  - Tasking
  - Enforcement (optional and notional)
    - Remediation?
- Challenge 1: Lack of modularity
  - Inability to plug and play components, you buy one big solution
- Challenge 2: Cross-product/instance orchestration
  - We have standard content but not standard interfaces
- Challenge 3: Static data sets
  - Query and analysis is more difficult, especially at different org. levels
- Challenge 4: Lack of a Tie to Enforcement Mechanisms
  - (Remediation)
  - Mike Kinney, of NSA: You haven't addressed defense in depth issues, or other security issues that make that complicated. You'll end up with a central console
    - Dave Waltermire, of NIST: hopefully we can leverage existing controls and expose them in a way to enforce access restrictions centrally. Security needs will vary.
  - Question about whether companies will accept automated remediation
    - Kathleen Moriarty, of EMC: a lot of it is about "getting up to snuff" and not breaking things. A lot of poor products turn clients off. But no clients have the resources.
    - Dave Waltermire: Also striking the right balance between human approval and automated action
    - Gary Gapinski, of NASA: GPOs can be considered remediations, and those are applied automatically. Only reason we don't do it in the tool is b/c we can't

        - Charles McClain, of IBM: In the database arena customers have very complex environments. Always have provided guidance but never applied it in the tool. It's far too easy to do damage.
        - Dave Waltermire: Is that because of a lack of situational awareness? Is the possibility to do damage only because of a lack of information?
        - Charles McClain: To some extent, but it might be impossible to know enough to automatically remediate
        - Michael: In terms of remediation...automated vs. manual can be decided on a product by product basis.
        - Mike Kinney: We'll talk about remediation later. Want to remediate to our baselines, want to be able to mitigate certain things. Preapproved situations can be remediated automatically.
        - Kathleen Moriarty: This could be fed into ITIL processes, like change management workflows
        - Steven Piliero, of CIS: Majority of people will do the workaround approach, because a lot of times the security team doesn't have insight into application itself. Unless changes are integrated into development lifecycle, there's always a risk of making fixes directly to apps.
        - Larry Feldman, of Booz: There should be a phased approach.
        - Mike Kinney: A lot of the problem is SA.
        - Point: Product vendors need to understand how guidance affects their products
- Challenge 5: Monitoring data is not normalized
    - Specifically, CMDB type data
- Challenge 6: Many tools only collect findings (i.e. too high level)
    - Doesn't allow for analysis
    - Joe Wolfkiel, of DISA: it's a trade-off...depends on the value and size of low-level data. A lot of the time it's worth sending low-level data to compute compliance on the server, but sometimes it's too much data.
    - Gary Gapinski: It's called "continuous", but what is the line between reporting many times on the same data and collecting new data. It's confusing in the architecture, shouldn't be called continuous unless it's live data
    - Mike Kinney: You can't collect data that fast. You do it at a periodicity that is necessary and possible for your organization. Should really be called continual, not continuous. You'll never have perfect continuous SA.
    - Steve Hanna: TNC uses a lot of event-driven notifications, as opposed to polling. Sometimes a better model.
    - Dave Waltermire: Advantages to a multidisciplinary view...do what you need and can do, potentially also do both polling and event driven. Framework should support all of this.
    - Gary Newman, of Belarc: Wasn't implying that there was any one answer, but thinks the wrong approach is to talk about continual monitoring. We shouldn't talk about "scans"...take the data when you can and tag it. Talking about continuous monitoring is too vague. Tagging data is important.
    - Steven Piliero: The danger is going all one way...in the end a hybrid approach is best.
    - Dave Waltermire: Absence of information can be meaningful if you have a heartbeat that isn't there.

    - Don Campbell, of McAfee: A lot of this is based on experience with current scanning technologies. Is it more appropriate to talk about how often you need data and defer how to get that to the collection tools.
    - Joe Wolfkiel: From a bandwidth perspective, if you're covering a whole bunch of stuff roll-up reporting might be more useful. At different levels of the hierarchy you can have different reporting architectures. Different expenses to collecting different data.
    - Mike Kinney: Don't know what's on the networks, that's why we have to scan.
    - Steve Hanna: Is this architecture intended to support both very short term (second-level) responses as well as more long term (months?)
- Challenge 7: Standardized analytics
    - Makes it difficult to compare results across organizations
    - Challenging to put together guidance
- Reference Model Specification Layers
    - Describes a 5 layer model from communications (layer 1) to model (layer 5).
- Layer 5: The model - NISTIR 7756 (not much discussion)
- Layer 4: General Specifications - NISTIR 7799 (not much discussion)
- Layer 3: Data synthesis - difficult to find best practices.
    - Gary Gapinski: Is this marshaling of information or is it aggregated information?
    - Dave Waltermire: This layer deals with taking data in, performing an analysis, and coming out with an analyzed data set
    - Attendee from BAH: Just trying to provide first-phase capabilities, so there are definitely missing areas and we can talk about next steps
- Layer 2: Data Binding and Handling - NISTIR 7800

CAESARS-FE Subsystem Components

Presenter

Dave Waltermire, NIST

Review and discussion of each subsystem

- Visualization engine
    - Dashboard and reporting engine
    - Reports at multiple levels
    - Works with task manager to pass on user queries to get data
    - Question: Publishing of results isn't in this subsystem, right? Confusion about results vs. reports.
    - Mike Kinney, of NSA: How do you do dynamic querying?
    - Dave Waltermire: Layer 3 specification defines named queries, then higher levels refer to those by names. Challenge is doing this in a technology-agnostic way.
- Content Subsystem
    - Serves as the repository for structured content
    - Needs to provide content when asked
    - Can ask for content from higher levels
    - Basic create/update/read/delete functionality for content
    - Mike Kinney: Unstructured data may be more valuable than structured data in some cases.

    - John Field, of EMC: Structured can mean tabular-SQL, semi-structured is XML, plus there's unstructured. Need to support all three.
- Task Management Subsystem
    - Orchestration of data collection, analysis, and enforcement
    - Joe Wolfkiel, of DISA: is there any concept of the orchestration working with GRC access control? Also sounds a lot like an ESB with PEP and PDP.
        - Dave Waltermire: Undecided, these are fairly abstract. But in general yes, we'd include that.
        - Dave Waltermire: Yes, some of this could be implemented using WS-*, but other people don't want to do that
        - Joe Wolfkiel: It's hard and expensive, but that doesn't mean that you can change the name and protocols to make it less hard and expensive.
    - Query orchestrator component: has knowledge of assets and can exercise queries against them
    - Collection component: collects results data from collection subsystem and responds back to query controller
    - Enforcement controller: handles enforcement and responds back to query controller
    - McAfee: Does this subsystem handle recurring schedules?
        - Dave Waltermire: You could do it either here or in the collection subsystem. But if you need orchestration across several products it would need to happen here.
- Collection subsystem
    - Supports whatever data collection use case, generally through instrumenting an existing sensor to accept tasking from task management subsystem
    - Continuous vs. continual vs. on-demand is an important distinction. Need to support one or more.
- Data Aggregation Subsystem
    - Performs data aggregation and correlation of collected data
    - May be performed by more than one application
    - Stores:
        - Asset data
        - System state and findings
        - Metrics
        - Metadata
- Analysis and Scoring Subsystem
    - Provides analysis and scoring of aggregated and correlated data
    - May be federated across several instances
    - Publishes results back to data aggregation (NOT directly to presentation)
- Enforcement Subsystem
    - Remediation
    - Ticketing
    - Network Enforcement
- Use Case: Monitoring Configuration Baselines
    - Walks through the process, from creating content, dissemination, tasking, reporting, aggregation, analysis and scoring
    - Includes creating a list of applicable assets from general guidance
    - McAfee: Saying "all of the data is collected" is problematic.
        - Dave Waltermire: Yes, sometimes you only collect partial data. You would need to track that.

- John Field, of EMC: There's a lot of freedom to innovate in the Task Manager, so don't want to over-specify it.
- Gary Newman, of Belarc: It's really just a data call.
- Mike Rains, of CSC: How will this work with multiple collection points? Where does it get reconciled?
  - Dave Waltermire: Depends on how those collection points work. Can be aggregated at different levels.
- Larry Feldman, BAH: The question is who knows about the asset?
- Steve Hanna, of Juniper: Are there use cases regarding responding in short order to incoming events? Do the interfaces provide the capability for asynchronous notifications?
  - Dave Waltermire: Potentially yes, but they're not defined. Maybe IF-MAP.
- Gary Gapinski, of NASA: Do sufficient standards exist now to implement this? Can we decompose it to the level where we can implement this?
- Tom Pearson, of Tripwire: If I find something on my network, how do I know if it's CAESARS-FE?
  - Dave Waltermire: You're talking about service registration? We could use things like UDDI...good question, not something we've really tackled. Base assumption is that the organization should know whether they have this capability.
- Steve Hanna: Of course you don't just want to trust things you find on your network.

- Prototype
  - Implemented content repository
  - Implemented rudimentary collection subsystem
  - Starting work on task manager
  - Google Code project for CAESARS-FE
  - Sourceforge project for content repository

Lessons/Tips of Creating PowerShell Configurations for OVAL

Presenter

Michael Tan, Microsoft

Introduction

In 2011, Microsoft started to work on the Microsoft Exchange baseline and found that the majority of the configuration settings could only be collected using Microsoft PowerShell cmdlets. This presented a problem because every baseline that is shipped with Microsoft Security Compliance Manager (SCM) can be exported as SCAP, however with the Microsoft Exchange baseline, it could not be exported as SCAP because there was no support for checking configuration settings using Microsoft PowerShell cmdlets.
As a result, at Summer Security Automation Developer Days 2011, Microsoft presented a session introducing the win-def:cmdlet_test that leveraged cmdlets in PowerShell to check a system's configuration.

http://oval.mitre.org/community/docs/OVAL_Developer_Days_2011_Minutes.pdf

The proposal was discussed with the community and eventually incorporated into the OVAL 5.10 release and was made a part of SCAP 1.2 in September 2011. This session focuses on the lessons learned from

creating configuration baselines using SCAP 1.2 and more specifically the win-def:cmdlet_test introduced in OVAL 5.10.

PowerShell Configuration Data Model

PowerShell cmdlets are utilities that the operating system provides to perform specific tasks. They consist of a verb that describes the action to be performed (e.g. get, set, etc.), a noun that describes the object on which to perform the action (e.g. service, ACL, etc.), and other input parameters that can be used to further specify what information to collect or operate on. For example, if you specified *Get-Service <service>* you would retrieve information about the particular service, but, if you just specified *Get-Service*, you would retrieve information about every service on the system. The primary difference between PowerShell and a traditional shell is that everything is represented as a .NET object and you can specify exactly which properties, in the object, you would like to target using the *Select-Object* cmdlet.

The following is an example of a win-def:cmdlet_object from the Microsoft Exchange configuration baseline which was shipped this year as well as the corresponding command line that would be executed in the PowerShell environment.

```
<cmdlet_object id="oval:microsoft.com:obj:1">
 <module_name>Microsoft.SolutionAccelerator.Baseline.Exchange</module_name>
 <module_id>{f1486d15-04a1-4782-82e9-981c2b986f4d}</module_id>
 <module_version datatype="version">1.0</module_version>
 <verb>Get</verb>
 <noun>ExchangeConfiguration</noun>
 <parameters datatype="record" operation="equals">
  <oval-def:field name="configtype">AdministratorAuditLogging</oval-def:field>
 </parameters>
 <select datatype="record">
  <oval-def:field name="property">SettingData</oval-def:field>
 </select>
</cmdlet_object>
```

```
>Get-ExchangeConfiguration -configType AdministratorAuditLogging | Select-Object -Property SettingData
```

After a year of working with this, Microsoft thinks this model is very simple, scalable, and easy to adopt.

Problems

To start off the discussion, the following questions were asked.

Question: How many use PowerShell in their daily lives?
Response: Many of you use PowerShell.

Question: How many have SCAP content and OVAL that leverage PowerShell?
Response: We haven't done much with it. I think only tool that has support for it is the reference implementation so if anyone has tools that support it let us know we would like to start experimenting with it.

PowerShell is a language, engine, and a tool for automation. Furthermore, because it operates on objects, what used to take hundreds of lines of code can now be accomplished in just a couple of lines of code, which makes the automation of tasks very simple. Cmdlets are production oriented, meaning that they are developed for customers to use right away to perform various tasks. However, the problem is that when a product is shipped they do not necessarily consider compliance as a requirement in the design stage. That is, as a cmdlet developer, you will not necessarily take requirements from a compliance aspect. Typically, a cmdlet developer will focus on management tasks that need to be performed by a user for that product and will build the cmdlets around that. User interfaces can then be built upon the cmdlet. However, in focusing on management tasks, they do not necessarily think about compliance. As a result, we usually need to run one cmdlet to retrieve data, and then feed that data into another cmdlet, and build up that logic to expose the configuration information. An example of this is checking for expired certificates on Internet Information Services (IIS). Since a cmdlet doesn't exist, we would need to use a cmdlet to enumerate the certificates, another cmdlet to get the properties, check the date, and so on. The reality is that compliance is usually the last requirement.

Solution

We had to build our own custom cmdlets which are a wrapper around the Microsoft Exchange cmdlets to really support SCAP and our model in OVAL 5.10. In PowerShell v2.0, creating cmdlets are relatively simple in that you only need to be able to write PowerShell functions, unlike in PowerShell v1.0, where you need to know how to write C# code and then build the cmdlet into a module. The downside to having to build our own cmdlets is that it requires an additional deployment of cmdlets to the target machine.

Question: How are you going to distribute the custom cmdlets that you are going to write? With the SCAP content? Does the SCAP content have to worry about checking the right cmdlets are deployed? Right version of PowerShell is deployed? How do all these very important things fit into the work flow?
Response: A very valid question. When we ship Exchange, and Exchange comes with SCM 2.5, SCM doesn't have the capability to take the baseline content to generate the SCAP 1.2 data format. So what we do is any baseline we ship contains documentation and a user guide to document features and we have baseline attachment for SCAP 1.2 datastream as an attachment in Exchange baseline. Same thing for custom cmdlets, we ZIP an MSI in the Exchange Baseline to ship custom cmdlets to customers.

Question: If an SCAP tool wants to use those custom cmdlets, does a tool need to install them on the system?
Response: Yes, the custom cmdlets need to be installed.

Question: If it's not there beforehand. Obviously, with an auditing tool, we don't want to modify the system right? So, for any SCAP validated tool, you would need to tell it what the cmdlet prerequisites are so you can error out if perquisites are not installed.
Response: Yes, that is the prerequisite for consuming the content. That is not necessarily a tool prerequisite.

Comment: The tool needs to be aware whether or not the cmdlets are there or else it is not going to work.
Response: Yes, that is correct. The tool needs to support PowerShell in OVAL 5.10 anyway. So, the tool needs to be aware of the content it has in the cmdlet definition in OVAL. But again, not a tool prerequisite, but, a content prerequisite.

Comment: Just want to add that if a cmdlet is not found on the system when you try to evaluate that OVAL Definition, there is a defined behavior in the OVAL language, as it is now, for what to do when not found, but, obviously your results are going to end up being trickled up to an error result.

Comment: If you look at the example, I highlighted the critical data and if you look at data model there is information for identifying cmdlets. There is module ID, module name, and version. So, that means custom cmdlets live in that PowerShell module. You have the prerequisites and a tool should have the minimal information to indicate what is wrong with the system.

Question: Is there a road map item so if Microsoft is going to have cmdlets for performing various configuration checks that those cmdlets will get installed on the system alongside the installation? Otherwise, that is really bad.
Response: I totally agree. So, I will leave that as lessons and tips and I can summarize after the demo.

Demo

The first part of the demo described how in PowerShell v2.0 there is new file extension (*.psm1*) for creating your own cmdlets in a module and that it is much easier now that you do not need to know how to write C# code.

The next part of the demo showed how there are two key files to create a cmdlet: a module manifest file (SCAPDemoCmdlet.psd1) and a script or binary module file (SCAPDemoCmdlet.psm1) that is associated with the manifest file. The module manifest file describes the versioning, metadata, and contents of a module as well as other configuration information that specifies how the module should be processed. A module manifest is needed whether you are creating a cmdlet using PowerShell scripting or C# and an initial template manifest file can be generated using the *Create-Module* cmdlet. Once generated, the relevant information just needs to be filled in. It is important to remember to make sure that the module name, defined in the manifest file, aligns with the name of the module file and specifying the function that you want to export which is basically the cmdlet that you want to export and use. The module file requires that you define a function using PowerShell scripting. The function name should be *Verb-Noun,* as specified in the module manifest file, and the function can be decorated with *[CmdletBinding()]* to specify the input parameters of the function. The demo cmdlet specified a cmdlet name of *Get-MyConfiguration* and input parameters *InputParam1* and *InputParam2*. Then the process of the function can be implemented which describes the logic used to expose the information that you would like to collect. The demo also showed how you can extend the properties of an object returned by a cmdlet, at runtime, using the *NoteProperty*. The *NoteProperty* was used to create two output properties *OutputProperty1* and *OutputProperty2*. The module manifest and module files used in the demo will be available in the ZIP file containing the slides for this year's conference.

Next the demo showed the different steps required to setup the environment to run the custom cmdlet. First, the module needs to be imported into the PowerShell environment using the *Import-Module* cmdlet (*Import-Module SCAPDemoCmdlet.psm1*). You can then check if a module is available in the environment by using the *Get-Module* cmdlet. Next, the following was executed in the PowerShell environment and its output shown.

Get-MyConfiguration –InputParam1 1 –InputParam2 3

Lastly, the following was executed in the PowerShell environment and its output shown noting that only *OutputProperty2* was targeted.

Get-MyConfiguration –InputParam1 1 –InputParam2 3 | Select-Object –Property *OutputProperty2*

Tips and Lessons Learned

The first lesson learned is that we should consider compliance requirements when we ship products. At Microsoft, there are the Common Engineering Criteria (CEC) requirements that they must adhere to when building server products. PowerShell scripting and cmdlets for automating management tasks are one of those requirements. It doesn't say compliance is one of those requirements so we are working with the CEC team to push for compliance to be one of the requirements considered at the design stage instead of after the product is shipped.

We are also trying to push our custom cmdlets back into the product for Exchange, SQLServer, IIS, and SharePoint to solve the problem of having to do an additional deployment. If the cmdlets are pushed back into the product that means, when the product is installed, the cmdlets will be available for the customer to use. However, until then, having the cmdlets as an attachment is a transition solution to solve the problem because many of the products being shipped today do not have the necessary configuration cmdlets.

Lastly, we realize that deployment is the main concern, but building the necessary configuration cmdlets is not that complicated as demonstrated.

Question: I just wanted to make sure that I understood. So, basically you are on Microsoft team developing baselines for Microsoft products. From my understanding of your position, relative to the product development team, is you guys aren't directly connected and you are not the same folks developing Exchange or SQLServer. It's after fact that you have to develop baselines. So, what you are dealing with is the challenge that the initial product development team has provided necessary hooks or cmdlets. Really, so you are developing cmdlets and then going back trying to educate and encourage them to incorporate those cmldets into, or the capabilities needed for those cmdlets back into the products in their releases.
Response: Exactly.

Question: One of my questions is how big of a gap did you find in the Exchange baseline and how many cmdlets did you have to develop on your own versus what was provided?
Response: The gap is not small. It's almost like for every configuration we had to build custom cmdlets.
Response: Because that was definitely the vision that the products would have the cmdlets in there and I remember Jeffery's architecture diagram; that all the GUIs might have for controlling their product in the Common Engineering Criteria would really behind the scenes be interfacing with cmdlets and those same cmdlets would be exposed so we could call *Get-Verb* to read the different properties that the GUIs were setting that we cared about in our baseline.
Response: That is right. But, remember though, it's not very typical for Microsoft products to ship and build in some compliance solutions at, or with, the product. They rely on other products to fill that gap like configuration manager has the capability to do that. So, the compliance requirement is not actually there. That is basically the situation.

Question: Quick question here. So there is this middleware of cmdlets now. Doesn't that kind of assert the ability to have version issues with cmdlets or the cmdlets not returning back accurate data? So, if someone else wrote a cmdlet that looks and feels like what you are expecting it to do and returns back values, do you have any assurance that it is actually returning what you are asking for? Because, it

seems to me like you have some middleware where you could be returning back true regardless of what the actual setting is. Or, maybe I am missing something?
Response: Right. So, I just want to say, if you can do the threat model for any application right, I don't consider this different from any other application; just go to registry versus WMI or cmdlet there is no real difference. So, in terms of security, I think we talked about this, but, PowerShell has capability for signing; like all the custom cmdlets we built are signed. So, that is one of the ways to mitigate any random change for the script and PowerShell has whole other capabilities in terms of security and you can also leverage that.

Question: So, as part of OVAL, is there something to verify the signatures to know that it actually came from Microsoft or another trusted source? I should know this, but, I haven't read that far ahead yet in the spec. I know we are working on it and have some guys doing it, but, I'm kind of new on this.
Response: Again, in terms of the content integrity, the OVAL content, it is no different than any other type of construct, they all have the same issue if the content is not signed.

Comment: I think what Jack is asking is there a way to check if the cmdlet is signed when you use it in OVAL?
Response: Yes, just like when we tried to get the demo working yesterday, when you have PowerShell initially installed there is policy called ExecutionPolicy. In the initial ExecutionPolicy it is restricted so that means the PowerShell cmdlet has to be signed, but, you can configure that because my demo is not signed so we have to get around that. So, there is no way you can launch the cmdlet if your policy says the script has to be signed. There is no way for your tool to launch the cmdlet.

Question: Can you actually run a cmdlet to check the execution environment is set up to check for signed cmdlets?
Response: Yes.
Response: What you could do is create a check for that and then.
Response: Right. So, normally when we ship a baseline, the execution policy is one of things we check first. That is one configuration that we had no need to build a custom cmdlet for.

Question: If you changed a cmdlet. If you versioned a cmdlet to fix an error or do something different in a later version in the content, it will have a completely different ID or a completely different name so there won't be any version control problems between benchmarks?
Response: Versioning is actually already covered see the module has ID, etc.

Question: So, if you make any changes it will be a different module ID effectively and there won't be any confusion between versions of cmdlets?
Response: Right, minor revision will be same GUID as the version changes. But, if a new module, it will be a new GUID.

Question: Did you say it would be same GUID? I guess what I was asking was could the situation arise where you have a new version of a benchmark and a new version of a cmdlet and now in your environment you have two versions of each running around. Could a new version of the benchmark mistakenly run the old version of the cmdlet? Could you have bad data because they are not synchronized? Are there checks in there or do you need have the OVAL check for that?
Response: That can happen.

Comment: I think what Jack was going for was for is there are two versions to control.

Response: Let's say we ship Exchange 1.0 and it worked with cmdlet module 1.0 and later on we find there is a bug in our baseline content and need to revise the content. At the same time, we find we made a mistake in the custom cmdlet and we did something wrong. Then we can actually ship 1.1 with the same GUID and module and then the new content, the new benchmark, you can define the particular settings and actually reference 1.1 not 1.0. So, that portion of the requirement would mean you need to get a new module installed to be able to check the new benchmark. But, that is kind of the version control we have today in place to deal with that.

Question: So, it might be useful to have in your benchmark a check that the right cmdlet is installed? Because, otherwise, if you just rely on it didn't work and throw back an error or something then you're not helping the user figure out what to do.
Response: No, that's not the case. Yeah, you have the version check.

Comment: Oh, it does?
Response: Yeah, it does (pointing to the module_version in the sample content).

Question: I know it does here, but, if that cmdlet is not installed, then do you have a check in your benchmark?
Response: That would be a runtime error.

Question: Is it part of your policy to make sure that the correct cmdlets are installed so you can have a rule that you could fail? Then the user would know I failed or got errors in all my tests because I didn't install all the cmdlets versus everything explodes and you have no idea why.
Response: Normally, for any application to deal with that kind of situation, because you could have a runtime error right not only just whatever the component it depends on you could have lost network you can have many reasons to fail at the runtime and the tool has to deal with that through the reporting. But, the unmatched version could be one of these.

Question: But, it will be the tools responsibility to check that the right versions of the cmdlets are installed then report if they are not? Or, will that be part of the content?
Response: Nope, that is not what I am saying though. It is the tool can deal with the runtime error which in this case cannot find the module to match with the benchmark and then reporting back to the user. When you get the results, in the results, you should be able to know if this rule failed, but, a fail is not actually the rule validation failed because you have not got the data yet, right, because you cannot load the cmdlet. But, that information should be in the reporting to indicate this rule failed is not actually the rule validation failed it is actually the execution failed.

Question: Is it clear for implementers that it is a requirement for them to build too?
Response: What I am saying is it is not a requirement for the tool to check the cmdlet version because of the content itself. Think about my tool asks to do something, but, because the underneath component is not installed in this case, because I updated my content, it will fail because when you load let say this module version 1.0, but, I am asking 1.1 and then the module load will fail. When that fail happens then you cannot check the rule right. So, this setting will fail, but, the reporting mechanism will say setting failed because it did not actually fail on the rule validation.

Question: What reporting mechanisms are you talking about? So, I guess where my confusion is you are saying it will fail.

Response: I am not sure. Maybe some of the people, I am not familiar with the XCCDF Results and OVAL Results design. Is one of the enumerations the runtime fail?
Response: If you have this OVAL content and you try to run it on your system. If a tool tries to execute this, it is supposed to report that there was an error that something happened. You can find out that the cmdlet didn't exist and the tool can put that in the system characteristics file and say that we didn't find this cmdlet on the system. Whether or not the tool propagates that up to the XCCDF, I don't really know.

Comment: Same thing for registry and WMI. WMI is a typical example. Provider has a new version and now for some reason you do not have the new version or you made the mistake and deleted the WMI DLL, then what happens with the tool? You have to get some information back, in the report, to say sorry I didn't get the data because not actually rule validation failed, because I don't have data to validate, but, I got some runtime error because I cannot find the DLL.

Comment: So, that is left to the implementer because I am just guessing tools right now will just fail and throw an error.
Response: Yeah, that is nothing different from say PowerShell versus any other mechanism.

Question: There is a module_version in the cmdlet_object, how do you expose the module version to address the versioning issue?
Response: You can define the version (pointing to the module_version in the sample content).

Question: I see it there, but, where is it exposed in the module itself so you can validate it?
Response: Going back to manifest. The version is just specified and expressed in the manifest definition and when you ship new module make sure you update this manifest.

Comment: Right, but, I don't know that tool vendors are going to ship modules because then if I write a module then Jack's tool isn't going to be able to run it because he hasn't got it. So, for the content to be portable…
Response: Well, I agree, that is one of the sort of additional tasks to put into the customers right. I agree that's a tough decision, but, that is the only solution. When you download Exchange and the Exchange attachment actually has one of the MSIs that you can just install. So, it's not something say we ship a bunch of files and then you just manually install that is not the case okay. So, you just run the MSI then everything is beautiful. Cmdlets are available for the system.

Comment: Right, I think we are going to want to use those; the cmdlets from Microsoft is going to provide alongside these.
Response: When we ship Exchange, we ship those cmdlets with the baseline.

Question: Right, so you guys will increment the module versions?
Response: Yes, absolutely. What I am talking about is any application for any product outside Microsoft. Well, with Microsoft, you can make a decision to build another way.

Question: I guess I had something slightly related to David's question. You said a couple of times the baseline will come with the MSI that's going to package the cmdlets so there will be a nice install. So, what about when we want to deploy that to any number of systems in enterprises, is there a way to easily deploy those cmdlets so that we can monitor many systems across the enterprise?
Response: ***Unknown response because the microphone was off.***

Response: So, is that the answer basically?
Response: Yeah. SCM also supports the CM format called the DCM and the same time we actually support the software distribution package that works with Microsoft technology. That's the technology mechanism with one baseline you can target thousands of machines, but, for SCAP that's kind of challenging because the SCAP customer does not necessarily have this SMS or CM in the environment so I am kind of wondering what is the other technology products used to deploy any product?

Comment: The cmdlets could be packaged in the object definition, which I only put out there as a horrible idea, but, that's an idea. Right, that would guarantee that if you have a piece of content that we all have the definitions.
Response: Has the dependency on the deployment right?
Response: That it doesn't have external dependencies. Right, so that is to say, you could have a string entity. We could extend the cmdlet; make a cmdlet511_test that allows you to embed the module. Who else hates this idea?

Comment: Well, I don't know. This is a similar situation as you deploy the baseline for Exchange and target 100 machines which only three machines have Exchange server running and so then you will fail actually on the many servers which don't have the component.

Comment: Well if, I mean if you're lucky, there is a non-cmdlet applicability check in XCCDF and then you don't have to deploy it.

Comment: Yeah, that's CPE. CPE is different because this is dependency check versus the applicability check. I really can't think of any reason.

Question: So is your SCM tool going to still output OVAL standard? Or, is it going to be turned all into PowerShell?
Response: Yeah, it does actually still support exporting an SCAP datastream.

Response: And that's intended to go forward? So, we can always fallback on that position if we don't want to use the PowerShell cmdlet?
Response: Yes, absolutely.
Response: No, I don't think that is what he is saying. He can export OVAL and for that OVAL to run successfully he will need to have the cmdlets installed because he is using the cmdlet_test.
Response: Oh, so that isn't what I just asked him. So, are you still going to output standard OVAL without PowerShell with your SCM tool?
Response: Correct, we are going to continue the baseline with OS and major applications using group policy based like Office and we continue even though they have cmdlets probably continue to use the traditional definitions. The reason behind this is not many applications support cmdlet.
Response: But, this whole issue is because Exchange didn't have the cmdlets to check those things.
Response: That is why it's kind of like the chicken and the egg. If more content supports PowerShell then applications will follow right? So, then we can actually gradually migrate, but, that takes some transition period.

Question: So, for example, for Microsoft Exchange, the exported SCAP datastream will not be capable of all evaluations, particularly those using PowerShell ,for which a specific custom function need be declared and present on the destination system. So, will that datastream then be thus reduced and only handle the registry entries and the like?

Response: There is a very small portion of the registry for Exchange. Very small, I am talking the majority are actually PowerShell.

Comment: Thank you.

A final lesson learned is that supporting the *Where-Object* cmdlet could improve our data model because it could reduce the data sizes and improve performance through less post-processing of the data. This was discussed during the initial proposal last year, however, there was concern around the fact it would require adding support for the PowerShell language which means you would allow content authors to include arbitrary PowerShell code in the *Where-Object* input parameters. We are currently trying to work with the PowerShell team to get some type of light-weight support for the *Where-Object* cmdlet that doesn't include that requirement.

CAESARS-FE Interfaces

Presenter

Dave Waltermire, NIST

Joe Wolfkiel, of DISA, had a general question about "how much should I care?". I.e. how soon until this is finalized enough to put into acquisition guidelines? Dave Waltermire: 5-6 months at the earliest...draft versions first.

Gary Newman, of Belarc, asked about the context and where DHS is taking CAESARS-FE. Dave Waltermire said it was out of his scope, nobody could really answer it.

I1: Result reporting (SOAP vs. GRC/REST)

- Presented extension of AI to carry asset information, used synthetic IDs to carry random info
  - Intended to be a stopgap, others wonder whether the stopgap will become permanent
- Gary Gapinski, of NASA, pointed that out, Dave Waltermire admits it's a problem and is semantically meaningless
- Joe Wolfkiel points out that this really should be aggregate reporting, individual reporting is not a strong use case
- McAfee: scale is definitely an issue
- Gary Newman: What are the DHS goals for asset information
  - Dave Waltermire: I don't know...we're trying to see what the vendors want to do
  - John Banghart, of NIST: We also consider non-government uses as well
  - Gary Gapinski: are there any defined use cases?
- SCAP Results Interface
  - Joe Wolfkiel: how do you tune which types of results you want?
  - Dave Waltermire: We're limited by what's in SCAP, but yes you want to do that.
  - Joe Wolfkiel: offers up work on ARMOR project to specify result tunings
  - Dave Waltermire: recommend you bring up these capabilities during next SCAP rev discussions
- GRC
  - Joe Wolfkiel: does this require me to go through an entire GRC workflow with all the caveats? Or can I just send it?

        - Dave Waltermire: you can just send it
    - Gary Gapinski prefers GRC over SOAP. Also wonders how policy information will be exchanged? Where does policy "occur"
        - Dave Waltermire: use GRC exchange policy capabilities
    - More discussion on policy enforcement
- I2.1 - Content retrieval (SOAP vs. REST)
    - John Field, of EMC: Prefers #2 (REST). They've already done something very similar and it works well.
    - Mark Davidson, of MITRE: Stick to standard GET/POST
    - Don Campbell, of McAfee: Simple is better, #2.
    - Gary Gapinski: REST is preferable...but note that HTTPS (TLS) is necessary to avoid potential insecure proxies
    - Question about compression, Dave Waltermire pointed out HTTP supports this natively
- I2.2 - Content CRUD Operations
    - Gary Newman: Confused about CRUD...we talked about the controller previously, but now we're talking about push operations from the scanner and CRUD from client to server?
        - Dave Waltermire: These are specifically for content management, not scan or results data
    - Question about why they were SOAP
        - Consistency is good, yes
        - Joe Wolfkiel: concepts of permissions, licensing, etc. you'll need access control. SOAP provides some of that for you.
            - EMC commented on some options for REST, but in particular proxy authentication is not handled as well in REST
        - Kathleen Moriarty: Really you want to build on top of transport bindings that can be swapped in and out.
- I3.1 - Tasking
    - Question about how assets are listed. Answer is AI.
    - Concern w/ query results scripting...not all data could be put into a single result type.
        - Dave Waltermire...that could just be an error?
    - McAfee: Still concerns about how data is reported when not all of it may be current. What happens when all data isn't available.
    - Don Campbell: Assumption that if the collect bit is given but can't be collected there's an error. (EMC): Same is true of opposite, you can have cases where collect bit is false but you can't provide the data without collecting
    - Tripwire: Ability to run analysis descriptors twice makes the assumption that results get cached. So that would always have to be true.
    - McAfee: Is this data supposed to be used for trending? What sorts of historical data are kept?
    - John Field, of EMC: There's a distinction between expectations of the client (things that won't be useful otherwise) plus hints to the server (if old data is around, you can give it to me)
- I3.2 - Data Collection
    - Question: A result descriptor is really only necessary at the lowest level. Dave Waltermire: yes, but they're hierarchical
    - Question: Is there an interface used to propagate tasks down to subordinate task managers. Dave Waltermire: yes, 3.1 would do that.
- I3.3 - Analysis
    - Question: In a rollup architecture, we've seen 3 approaches: warehouse, mediation, and meta-queries. Is this agnostic to the mechanism for doing rollups?

- Don Campbell: is there any specification of a "timeout" window? Dave Waltermire: not yet, but it could be added
  - Joe Wolfkiel: this would be necessary.
- Issues with some of this is that the data queries don't scale. A very big danger in this. But if you define a data model at a high level you can let tools decide how to perform queries across it.
  - Joe Wolfkiel: we do this now with something very generic.
- Discussion of data age and a hierarchy of the different ages you want to allow (i.e. how much caching is allowed).
- Gary Gapinski: GRC does not enforce a response...the response could just be "ACK". Then the response could arrive at any time after that.
  - Kathleen Moriarty: In GRC, the only capability there is a response of pending...but that isn't a yes or a no, just a "we'll look at it"
  - Dave Waltermire: What about a way to ask for status? Also not there
- Gary Newman: Couldn't you just repeat task IDs to ensure that repeats could be ignored? Then if you don't get a request for awhile just make it again w/ the same ID.
- Don Campbell: What if there's a failure and you're stuck in a state of non-response? At some point you need to resolve that.

- I4 - Query data out of data aggregation
  - Mark Davidson: is GRC in the current proposal?
    - Dave Waltermire: only for the interfaces where the communication flows: I1 and I3
    - Mark Davidson: Can't you use callbacks then (in RID).
      - Mark Davidson: You've mentioned REST...any thoughts on URLs? Potentially having those well-defined would be a good step. Dave Waltermire: there's no convention defined yet, probably drive discussion based on prototyping efforts. Mark Davidson: need to also make sure to include version info.

CM Discussions and Next Steps

Presenter

Dave Waltermire, NIST

Dave Waltermire asked what the next steps should be and threw out a few options: query language, asset data, tasking language. Which should we prioritize?

- John Field, of EMC: The data models and concepts should come before the languages and transports
- Jim Ronayne, of Varen Technologies: We've been talking about the asset data model for a long time, we should do that. But applicability language might be related in that it needs to select things based on an asset data model.

Automated Checking of Windows User Configuration Settings

Presenter
Jack Vander Pol, SPAWAR

Introduction

There is an automation gap in SCAP related to the ability to review Windows User Config settings. With the increasing amount of products being configured on a per user basis such as Microsoft Office, Internet Explorer, and many OS setting, SCAP, especially the OVAL language, need to be updated to allow for this automation.

Think of settings at a user level GPO, not a computer level GPO, that is the scope of this conversation. There are many settings that are user based, not computer based – especially Internet Explorer. We did some research and found a method of reading the NTUser.dat file.

The Automation Gap – User Account Configuration

I didn't realize how big the gap was. 1453 of 3102 settings, about half, of the GPOs in Windows Server 2008 are per user based settings. Unless I'm missing something, there's no way of checking any of those in SCAP. In our tool, you kind of can, but it's kind of a hack and it's not consistent with anyone else. We want to share our hack with the community and see if we can incorporate it somehow.

Comment from audience: Many of the settings are only preferences, not security settings.

Jack Vander Pol: Yes, I did not filter for security settings. But I still think there is a sizable amount of items that we may be interested in.

Current Capability

Currently use HKEY_CURRENT_USER hive (HKCU), but that is only applicable for the currently logged in user. If there are three users on a system, could get three different results. Also, HKCU is not available remotely.

Our Method

We implemented this in 2005 or 2006 in a non-SCAP tool; then ported it into SCC.

HKCU is a pointer to HKUS\<SID>. So if you know who is logged on, you can search each user. So you could use HKCU remotely if you know who is logged in.

HKUS\<SID> corresponds to each user's NTUser.dat file. So we read the NTUser.dat file of any other users on the system. And then roll that in with some OVAL messages to say here's the user, here's where we found their data, and here's what the setting was. So you end up with an abnormal response back. If the SCAP test content asks for HKCU you might get one result or five, or however many users are on the machine.

Dave Waltermire, of NIST: Is this something that you could use multiple instantiation for? XCCDF has some support for something like this.

Answer from audience: The problem is that at the OVAL level you have multiple results, but at the XCCDF level you have only one pass or fail, so you have no idea what user it is.

Jack Vander Pol: With any recursive, one-to-many OVAL check you can end up with one result in XCCDF.

From audience: USGCB kind of solved this problem for patches, but we should solve it for all these common cases.

(Back to presentation)

I emailed the OVAL developer list with our approach and the response was that you're covering the requirement, but also going above and beyond. There's no problem with that.

Proposal: Simplified Overview

Figure out who's logged into the system, read their data from HKUS\<SID>. Determine the User Profiles directory – you can read that from the registry. Filter out some unneeded NTUser.dat files. Perform the checks on each user found. We logged the username, last logon, and profile directory as info messages in OVAL.

Example Results

Presented an example of checking a setting in Outlook 2007

Gary Gapinski, of NASA: Aren't there some GPOs that are not permanent? Do these go away when the user is not logged in?

Jack Vander Pol: I don't know. From our observation, everything is saved. There may be some.

Audience: Roaming users.

Jack Vander Pol: Roaming profiles is another whole question.

(Some discussion of roaming profiles and Active Directory servers.)

Jack Vander Pol: Checking roaming profiles is out of scope of this proposal.

Dave Waltermire: Is there a cmdlet that would do this?

Jack Vander Pol: We did look at this briefly, but as of now, I don't know.

Audience: Michael Tan of MS was asked about this, and he's going to check with his engineers, but there's no obvious solution.

Question: What happens with a terminal server?

Jack Vander Pol: Depends on if it's an application server.

Discussion: OVAL Updates?

We implemented this into a system that is running on about 100,000 boxes once a month.

Is per user profile settings something that should be put into the specification?

Three scenarios:

1. Overload HKEY_CURRENT_USER registry documentation to explain that it means HKUS\<SID> and all unlocked NTUser.dat files. (This is the approach we currently have implemented.)
2. Update the registry_test specification. If you see HKCU then you need to iterate. But then what if you really only want info on the current user.
   Audience: This was an issue on the Mac with the plist test. Many OVAL tests seem to be written with the assumption that you are running locally with an account. We concluded on the mailing list that the plist test just didn't make sense.

3. Create a new OVAL test for NTUser.dat
   This is probably the right way, but will take a year or so to get into the spec.

(Some discussion of other implementation approaches that were determined not to be workable. E.g. using a wildcard variable, domain GPO issues)

Pros and Cons: Reuse registry_test

Redocument. Inconsistent tools until everyone complies.

Pros and Cons: New NTUser.dat test

More granular. But not available until OVAL 5.11 (or whenever)

Discussion Questions

Jack Vander Pol asked for direction from the attendees. If we were to proceed, what would be the way to go?

Dave Waltermire: Are we approaching this the wrong way? If the goal is to evaluate domain accounts (or local), should we go to the domain controller?

Jack Vander Pol: Querying domain controllers could lead to querying a couple hundred thousand accounts.

Audience: A little beyond this, you want to be able to look into a user's profile. If we had this new capability, we could use it as a hunting tool.

Don Campbell, of McAfee: I would go with a new test. Will give the vendors a chance to come out with the same thing working the same way at the same time.

Dave Waltermire: I think that's reasonable. But we should wait till we see if there's a cmdlet.

Jack Vander Pol: If the cmdlet is not installed, that is a challenge. And doing remote scanning is another challenge. And it's not going to be available in XP or Vista.

Audience: Agree on making new test. New behavior with new results. Not sure about cmdlet approach.

Danny Haynes, of MITRE: Even if it is available in PowerShell, we could create a test around it.

Audience: We loaded the hives and read the SIDs. Non-cached roaming profiles could end up copying gigs of data to the machine being tested.

Jack Vander Pol: This proposal is only testing accounts on the system.

Audience: Are you volunteering to write the new OVAL test for the community?

Jack Vander Pol: Yes. It's not hard. We're basically recycling some of the registry key stuff and adding a few things.

Audience: Can we also put on the list for XCCDF to deal with the instance results? Not just for this, but with any iterated thing?

Audience: OVAL does not support instance results.

Audience: Why not use the OVAL results?

Audience: Frequently do not have the detailed OVAL results at the higher level.

(Much discussion of OVAL & XCCDF reporting concerns.)

Audience: XCCDF needs a way to name all the instances but OVAL doesn't have a way to tell it.

Jon Baker, of MITRE: Put contributions into the sandbox on GitHub so community can iterate over it.

Audience: Can we vote on any of this?

Audience: It needs to go to the list.

Jon Baker: For OVAL we were hoping to take back to the list more mature ideas.

Dave Waltermire: We need to draw a line in the sand to say this or that should be in 5.11

Jason Mackanick, of DISA: The issue of querying domain controller or LDAP directories is big for me and my users. OVAL is not "enterprisable".

Dave Waltermire: Agreed. We can't target in a guide that these checks apply to the domain controller and these are local.

Jack Vander Pol: In a new test we could specify whether you want local accounts, domain accounts, or both.

Audience: Right now, if we only want to target the local or domain, we can't.

Dave Waltermire: This relates to remediation. If we can't tell where the problem is we can't fix it.

Dave Waltermire: Part of the problem here is we are considering this issue piecemeal in each specification rather than the broad use case. I think we've been having this conversation the last two or three years and we haven't made any progress.

Jack Vander Pol: My focus for this solution was for things like Office 2007, IE. We have an implementation that addresses a particular need. That's why I put all this stuff at the end; it's by no means a perfect solution.

Audience: Rough consensus & running code.

Implications Continued

Do not load each NTUser.dat as a registry hive – things will break.

Audience: From a DISA perspective having users screaming about Gold Disk going away and then they do a comparison analysis of how much OVAL can do vs. Gold Disk; these checks really start to show number differences big time. It's hard to argue with people who say, "Gold Disk used to be able to do this and now I can't do this."

Dave Waltermire: It's important to be able to tailor the reporting results.

(Some more discussion about domain controllers & LDAP & how remediation fits in with this.)

Thursday July 12th

OVAL – Mobile Device Assessment

Presenters
Chandra Basavanna,SecPod
Tim Nary, Booz Allen Hamilton

Quick Summary

Major points from this talk include:

- An overview was provided of the main mobile operating systems and their capabilities. It was noted that there is a great deal of variety in how much security control is supported programmatically between the different operating systems. It was further noted that the Android operating system itself has a huge amount of variability in its security capabilities since each device vendor tends to tweak their installation.
- There was a presentation on the use of Mobile Device Management suites to support SCAP-based assessment of mobile devices.
- There was a demonstration of the use of SCAP content (both existing and experimental) to scan Android devices for indications of vulnerabilities. This involved a lightweight OVAL agent on the device itself.
- It was noted that the relative immaturity of the mobile device market means that there is both a great deal of variance between operating systems and supporting applications, and also that the software and capabilities change at a rapid rate. Both of these represent significant challenges for development of standards for assessing mobile device security.

Minutes

Mobile devices are ubiquitous, but consumers generally select devices for features and performance rather than their security characteristics. This is not to say that they are inherently insecure, but they are generally not expected to live up to the security constraints we expect for our enterprise desktops. There are Mobile Device Management (MDM) suites that are expected to manage these devices, including their security, but their capabilities are limited by what the mobile device itself can enforce.

There is some existing guidance for securing mobile devices. We looked at mobile device STIGs and NIST guidance for mobile devices. Mobile operating systems have a range of capabilities. Blackberry is very mature with lots of security controls, but MDM vendors don't provide much support, although there is some. Apple iOS has a fairly good number of security controls. There are, however, some important things that cannot be managed programmatically, such as WiFi and Bluetooth. For Android, because (unlike iOS and Blackberry) the OS author is not the same as the hardware manufacturer, you have different "flavors" of the OS as different hardware vendors seek to differentiate their products through tweaks to the OS. As such, there is the "vanilla" Android and then many vendor "flavors". Vanilla Android has only 16 security commands, mostly dealing with passwords. However, because vendors may have altered the OS, it cannot be assumed that every flavor of Android will respond to all security

commands in the same way. There are also enhanced security APIs that can be added to devices to give MDM's greater abilities to support security.

Currently there are over 100 MDM vendors in the market-space, although only a few are really differentiating their products and providing unique capabilities. Most just implement the standard features provided by the operating systems. There is an expectation that there will be consolidation in the future as companies are acquired. It was noted that a couple MDM vendors have expressed interest in SCAP. MDM suites rely on Push Notification Services (PNS) - pushing commands to the device rather than relying on the device to check in. These communications are dependent on the device being connected to the Internet. It was noted that for Android devices, because of the open-source nature of the OS, anyone can stand up their own PNS to manage Android devices. By contrast, any PNS communication with Apple devices must use the Apple Push Notification Service (APNS). Microsoft Exchange ActiveSync is also commonly used on Android and Apple devices for communication and management.

In our project, we demonstrated the use of SCAP to assess mobile devices using MDMs. We chose to use MDMs because this meant that agents did not need to be developed and deployed for all the different types of devices we wanted to assess. MDMs often support more than one platform, and many are open, making them good targets for this proof-of-concept work. We looked at several MDMs and learned that many use a backend SQL database. We thought this would let us use the independent SQL tests in OVAL, but found that the lack of support for these tests in OVAL was a barrier. We also found that, because a single MDM will manage many devices, the SQL queries needed to be tailored to target individual devices within the MDM database. Because of these challenges, we developed the MDM-SCAP Middleware (MSM) which ingests SCAP (OVAL SQL tests) and then makes the appropriate SQL queries for each device. To succeed in this required support from the chosen MDM vendors to gain insight into their database structures. The implementers felt that it would be easier if there was a mobile device test for OVAL to express assertions about mobile device policy that could then be mapped to the individual mechanisms used by the various devices by the device vendors.

Based on this work, there are two things we would like to see in the future:

1. MDMs that are able to read and interpret SCAP content
2. MDMs that report using a standardized format

A suggestion was made that, instead of developing an OVAL schema for mobile devices, that some other management standards might be looked into for their ability to serve this use case. It was agreed that this might be worth looking into. It was also suggested that, because the mobile device market is relatively immature, building market share will be more important to device vendors than standards compliance. As such, many may feel that moving to standards is a distraction. Nonetheless, there is probably value in beginning that conversation, especially with the MDM and middleware vendors.

Chandra then presented an overview of a demonstration architecture that was developed for scanning mobile devices. He then presented of a demonstration where an Android mobile device was scanned for vulnerabilities using SCAP content. In this case, the OVAL interpreter was actually on the Android device

and the OVAL definitions were pushed to the device from a laptop and OVAL system characteristics returned. Actual evaluations are performed on the evaluator (laptop) using those system characteristics files. The OVAL tests themselves were primarily *uname* tests, detecting the version of Android as an indication of the presence of vulnerabilities. Also, some new OVAL tests were created by SecPod as part of the proof-of-concept. SecPod felt that this work successfully showed that the creation of Android schemas for OVAL is plausible and useful. A draft Android schema is in the OVAL sandbox.

Chandra was asked if there was a plan to release the applications developed in support of this demo. Chandra responded that there had not yet been a decision on whether to do that. There was also a question as to whether the developed tests could only be resolved on the device or if they could be resolved on an MDM. Chandra noted that all functionality used by the tests would also be accessible to an MDM - thus it was not unreasonable for an MDM to collect all relevant information and then have the OVAL run entirely on the MDM. The challenge in that situation would be that the MDM would have records for many devices - OVAL generally assumes it is being applied against a single host so there could be complexity in having tests that differentiated between individual devices on the MDM. It was also noted that any data on the MDM reflect state when the mobile device last synced its settings with the MDM, so test results might not reflect current device state.

A suggestion was made that it might be useful to have generic tests that described behaviors, such as which applications could delete from a file system. It was noted that the number of mobile device applications is growing rapidly and having separate tests for each application is intractable. This way one could identify which application had risky levels of access to system capabilities and make a determination as to whether this represented a concern. There was a sense that this might be possible, but others cautioned against building such policies as this got into author reputation and trying to define characteristics indicative of dangerous software rather than specific system artifacts. Most applications at least have the capability to do something bad to the system, so a simple list of capabilities wouldn't be useful. Making application characterizations that differentiated between two applications with similar basic access but which represented different threats to the system could be an extremely long and difficult path.

It was observed that, while the proposed alternative to multiple, torturous SQL queries was desirable, the diversity of devices and the rate at which they were being changed made the OVAL path extremely difficult at best, since each individual test would need to go through the process of adopting it into the official schema. Some observed that if there were abstractions that could be utilized across multiple devices, such as convincing vendors to utilize NETCONF and then writing tests to that interface, that OVAL might be able to help in that way. Again, it was noted that the immaturity of the mobile device market and the churn this generates makes addressing these challenges difficult.

OVAL for Inter-networking Devices

Presenters
Luis Nunez, Apex Assuance Group
Chandra Basavanna,SecPod
David Solin, jOVAL

OVAL for Inter-networking Devices

This session will focus on OVAL for inter-networking devices (routers and switches) specifically the introduction of a new platform component schema for Juniper JunOS as well as a demonstration of running code and sample content. It will also discuss some issues related to inter-networking devices some of which are similar to those in the mobile space.

jOVAL, SecPod, and Apex Assurance

This was a collaborative effort between jOVAL, SecPod, and Apex Assurance leveraging each other's expertise (domain knowledge, content development, and tool development) related to security automation to develop a proof-of-concept demonstrating the extension of OVAL to Juniper JunOS inter-networking devices.

Project Martini Goals

Given the common objective to add support for Juniper JunOS, within the security automation space, specifically OVAL, there was a need to develop a proof-of-concept and provide "rough consensus and running code" to show that it could actually work. This came in the form of developing an OVAL component schema for JunOS, creating content (OVAL, XCCDF, CCE, and CPE), and extending a tool jOVAL(jovaldi, Xpert) to run the newly created content. All with the end goal of getting the new component schema accepted into the next official release of the OVAL Language.

Current List of Platforms Supported on OVAL

A current list of platforms supported in the OVAL Language was presented. Currently, Cisco is the only vendor with inter-networking device component schemas (CatOS, IOS, and PixOS) in the OVAL Language. The issues for Juniper JunOS and Cisco IOS are very similar being that they are both inter-networking devices and you will notice that their schemas are very similar too.

Ingredients to Making this Work

In order to make this proof-of-concept work, three things are needed: schema specification support for JunOS, content, and tools. From a tool perspective, Xpert and jovaldi will be demonstrated running the SCAP 1.2 datastream content (OVAL, CPE, CCE, XCCDF) that was created based on the DISA STIG.

Juniper JunOS OVAL Schema

The new Juniper JunOS component schema contains four tests: the global_test, the line_test, the version_test, and the xml_line_test. Note that if you compare these tests to the Cisco IOS tests, they will appear very similar.

OVAL Tests (Inter-networking Devices)

A diagram showing the various OVAL tests associated with inter-networking devices was displayed. The diagram included the JunOS (experimental), NETCONF (experimental), Cisco IOS, Cisco CatOS, and Cisco PixOS platforms.

DISA Network Infrastructure STIG

For the proof-of-concept content, the DISA Network Infrastructure STIG was selected because there is a need for the content and because there is content for Cisco IOS. There is currently XCCDF content for Cisco IOS and for Juniper JunOS STIGs, however, there is no OVAL content available. As a result, a sample of the DISA STIG was taken. Specifically, the network infrastructure router and the content was created around that. The DISA Network Infrastructure Router STIG contains several other specific STIGs including STIGs for firewalls, networking perimeter devices, and general infrastructure devices.

Juniper JunOS Content – SCAP 1.2 Datastream

The content created, based on a sample of the DISA Network Infrastructure STIG, was represented in the SCAP 1.2 datastream format. It included XCCDF that described the policies, OVAL content based on the policies in the XCCDF, CPEs, and some conceptual candidate CCEs.

DISA STIG NET0400 Test

One example policy that was used from the DISA Network Infrastructure STIG was DISA STIG ID NET0400 which checks to see that you are following best practices and using interior routing protocols that are secured by authenticating with its peers. In this case, it is looking at the Open Shortest Path First (OSPF) protocol. An example JunOS configuration setting for this policy was displayed. The key settings that you want to be looking for are protocols OSPF, area, interface, and authentication. It was also noted that there are several different ways to output the configuration of the device. The first method that was shown was the curly-brace method. The second method that was shown was the set method. The set method was used because it seemed easier to parse the configuration information.

DISA STIG NET0340 Test

Another policy that was used from the DISA Network Infrastructure STIG was DISA STIG ID NET0340 which checks to see that the login banner is non-existent or not DoD approved. The main goal was to create some content that expressed some rudimentary checks against the system and it seems that the tests we have now will allow us to check for the DISA STIGs. However, as we look into this further, we may come across more complex checks and the need to introduce additional tests.

Demo

For the demo, we will go over the JunOS schema, the content, and then run that content through the jovaldi and Xpert tools, connect to the Juniper router, retrieve the configuration of the router, analyze it, and generate the results. Basically, what is done when scanning other devices.

Question: Would you mind describing the motivation for this? In particular, did Juniper ever participate in this? Or, do you think Juniper is going to dedicate some resources for this?

Response: Juniper is here in the room with us and from the beginning we engaged Juniper on the whole concept for providing support for JunOS within OVAL and within SCAP. So, they are very supportive of it and they are here today so maybe I will let Steve answer.

Steve Hanna, of Juniper: Yeah, this is something that Juniper has wanted to have for some time and we have been looking for how we can support SCAP and the assessment of our devices in a way that is compatible with the way that the devices operate and with maintaining an appropriate security profile shall we say for the devices and this approach looks like a good one to us.

Question: Would you say it is arising from a customer need as opposed to it being an academic exercise?
Response: Certainly, this is something that our customers have asked us for.

Question: Are you configuring these devices or are you auditing? I saw you setting and you were running the set command.
Response: No, we're showing the set command.

Question: Is that what the fix would be?
Response: That is what the fix would be. Yes.

Comment: It depends right. If you are running a test because you don't want something set, but, you are comparing it to the set command, then maybe your fix would be the delete equivalent of that command, but, let me just explain how we structured the Juniper schema.
Response: Okay.

Comment: But, just to be clear, we are not making any remediation or configuration changes to the device. All we are doing is retrieving and analyzing. We are not making any changes.

Comment: They are dumping the configuration of the device which shows up as set command.
Response: Yes.

Comment: So, that is what you would want to send if you wanted to configure the device in that manner, but, you are not actually reconfiguring the device.
Response: Yeah and that will kind of lead into a remediation at a later point if we need to.

The demo began with a walk through of the JunOS component schemas which can be found in the OVAL Language Sandbox at the following links.

https://raw.github.com/OVALProject/Sandbox/master/x-junos-definitions-schema.xsd

https://raw.github.com/OVALProject/Sandbox/07b9fcddcf8b19695d3b92f97bf7effafd3e5175/x-junos-system-characteristics-schema.xsd

This walk through included a description of the global_state, line_state, version_state, and xml_line_state. It was also noted that if you wanted to write checks that utilized the set view of a JunOS configuration, you would want to use the global_test. If you wanted to write content that took advantage of the curly-brace configuration, you would want to use the line_test. If you wanted to write content to check the result of *show version*, you would use the version_test. Lastly, there is the

xml_line_test which allows you to evaluate an XPath expression against the output of an arbitrary *show command* on the router. The SCAP 1.2 datastream that was created by SecPod, using these tests, was shown. It included an XCCDF Benchmark, OVAL Definitions for the CPE check, and the OVAL Definitions for the policies described in the XCCDF Benchmark. The CPE OVAL Definitions utilized the version_test and the majority of the OVAL Definitions to check the policies utilized the xml_line_test. Sample content, for the JunOS platform, can be downloaded from SCAPrepo (http://www.scaprepo.com/) by searching for "junos definitions".

Question: David, could you maybe just explain what jOVAL is and how it relates to Xpert and gOCIL. There is all this new terminology coming up.
Response: Yes. jOVAL is an open source Java OVAL Definition interpreter library and we have a couple of command line programs that accompany it. One is called jovaldi which is ovaldi except it runs the jOVAL library to do its interpreting. jOVAL has three different plugins that you can use for various use cases. It has a local plugin, which is the default, which scans the machine on which you are running it. It has a remote plugin which allows you to scan machines over the network. It supports Windows, AIX, Linux, Cisco, Juniper, NETCONF, Solaris, and it allows you to scan any of these machines from any other machine. For example, I can scan Windows from Linux and so on and then there is an offline plugin that allows you to evaluate content that you supply. So, either a datastream that comes from a Cisco IOS device, or a Juniper JunOS device, or a Mac. We also support MacOS 10. A MacOS plist file if you wanted to run the Apple IOS XCCDF, you can use the offline plugin to do that. What we are going to show is the remote plugin which is really the use case that we designed the jOVAL library around was to scan remote devices. And then later on, we built the XCCDF Processing Engine and Reporting Tool (Xpert) because it seemed important to provide a capability to do the higher-level XCCDF scanning or evaluation.

Next, the SCAP content was run in the Xpert tool. It ran the applicability check, evaluated the checks that were selected in the XCCDF, and reported which checks passed and which checks failed. An HTML view of the XCCDF results was displayed showing that two checks failed (2 and 3) and two checks passed (1 and 4). Next, one of the failed checks was manually remediated from the command line and Xpert was run again. Then the OVAL Results from that scan were displayed.

Question: What do you think is needed next for developing this JunOS schema? What challenges have you had in drafting the schema as it was? What should we be thinking about in order to take this as it is and mature it so that it is ready to be incorporated into the language?
Response: I don't see a need to expand the schema any further. We are still evolving and figuring out what additional checks we will need, but, I think for right now we have all the right tests in there to do. I think at this point 70-80% of the STIG.

Question: Steve, has Juniper been able to look at it? And I am assuming Juniper provides some hardening guidance, does it cover your needs there for expressivity?
Steve Hanna: We've taken a look at it. It looks like a good start, but, I think there are some more things that we will probably want to suggest to be added to it.

Comment: Yeah, I think certainly from a next step it's really on the content side; further building out content based on STIGs or any other security recommendations from Juniper.

Question: One of the challenges with any new construct we introduce to the language is promoting awareness of the construct. Driving awareness at the end user side helps drives demand at both the

operating system or application level and at the product that would do the evaluation level. So, kind of a question for Steve and for you Luis, from Juniper's perspective, is this something that once it's matured, you guys would look to embrace and publish content leveraging this extension? Similarly, for Luis, what other things do you think we could do to help raise awareness and drive adoption?
Steve Hanna: I am not able to commit on behalf of Juniper, but, certainly we have to look at what the customer wants and the customers have been asking us for this so I think that it's something that we would want to look very seriously at.
Response: Just to kind of follow up on Steve's comment. Part of what we did here is to show; demonstrate. Many times, I would do presentations to executives and management. XCCDF, OVAL, CPE, CCE, it is very hard to get so something like this where you can actually run something and produce based on something like the DISA content; its powerful. So, it's really good to show Juniper and really everyone here in the community and to customers. So, maybe my question is really is to the community, is this something that you want to adopt and further evolve? Or, are we still focusing on the current platforms that we need to further evolve?
Response: I absolutely want to adopt it. I want to expand OVAL and its reach to every platform we can out there and having JunOS supported is absolutely a plus.

Comment: Alright, so you guys can download this build if you are not up for forking your own clone or cloning your own fork or whatever on GitHub of the jOVAL project. That said, https://github.com/joval or you can just go to our website (http://joval.org/).

Question: I saw references earlier to the textual and the XML based query schemes. I was wondering if either of those have proved adequate to the task of answering all likely questions that might be posted?
Response: I don't know if it will answer all. We are still kind of building out the content and pushing to see how far we can leverage these tests. I think the next talk with NETCONF would get more into the XML and XPath aspects of it and I had the slide up there where we had the various tests. In the NETCONF schema, there is only one test as opposed to either JunOS or Cisco where there are several tests so having one test is pretty nice, but, we are still looking to see what the limitations are.

Question: Could you talk a little bit about the xml_line_test that you had in the content?
Response: Sure, here is an example. It is pretty simple, you just say here is the show sub command that I want to run. I guess this one is going to be show configuration protocol router advertisement interface em0.0 etc. and you specify the XPath that you want to evaluate against it. This particular object ID is 102. So, if I go back and actually look at the OVAL output and look for 102, there's the definition, there's the state, there's the result definition. Alright, so this test refers to item number 1 in the system characteristics and item 1 in the system characteristics – well, there you go, it has a value of does not exist. That's not very fun. This one is a little more interesting item 2 so you can see this XPath for this show sub command evaluated to this password hash.

Comment: So, this is the actual configurations that we are kind of keying off to say that, yes, it is configured on security best practices. XPath was pretty elegant in being able to quickly identify that and, again, I think really from a JunOS operating system, it puts this configuration out there in an XML format that is elegant and we can use XPath.

Question: From a Juniper perspective, is there a preferred way to do this? Is it preferred to do the XML output format? Or, just out of curiosity…
Response: There all supported so whatever is most convenient for you.

Question: I know that the example here was based off the STIG content so how was it developing the XCCDF plus OVAL content here? What sort of effort was there? Is there a way to take some of the existing content sources out there and auto-generate, or in an efficient manner, get to the XCCDF plus OVAL content?
Response: Yeah, I don't know about auto-generation taken from a STIG that's probably another endeavor. So, this supplies STIGs based on the XCCDF format, we didn't actually use the DISA XCCDF. We kind of rolled our own to make a quicker, smaller XCCDF file, but, it obviously still has the same characteristics of it.
Response: This was an unfunded side effort so I just threw four tests out there and it was enough that Chandra didn't ask me to do anymore.

Comment: Is there anything that we should be thinking about in terms of reducing content development costs for those that are authoring the content?
Response: You could probably just limit this whole thing to down to just the xml_line_test that was obviously pretty powerful. That's what all these particular checks were implemented in. Personally, I think it is really nice to have the ability to run XPath against the data that you are retrieving. I think that it simplifies content development, but, maybe I am not a normal content developer.

Question: When Juniper or other security agencies issue bulletins and guidance about how to identifying and remediating configuration issues, do they typically provide that guidance in all three approaches? Or, typically, what format is the guidance provided in? Does that make sense?
Response: So you mean an advisory that DoD would put out?
Response: Well, someone I assume is issuing security bulletins about like…
Response: Like IAVAs?
Response: Yeah.
Response: I'm not with DoD, but, from what I have seen it's a real lot of emails or text based. There's really no XML or automation piece to it which, technically, we can wrap some remediation around it.
Response: I think to the extent guidance is given, it is usually in the set format not the XML format.
Response: That would make the set, or global_test that we wrote, useful if you just want to copy and paste out of there.

Comment: We have some of the Juniper content as well. It's on our road map to match up the Cisco and Juniper content. We already have export and some format capabilities so if Juniper and some of the other community members jOVAL and other folks in the room that have expertise and are willing to commit and work together with us, I am willing to commit to get this done before the end of the year.
Response: I really like the community aspect and collaboration. Leveraging expertise and doing it as a community is a whole lot better and efficient and you get quality out of it. So, yeah, I would like to do some more collaboration with other folks. That would be great.

SCAP and the Network Configuration Protocol (NETCONF)

Presenters
Luis Nunez, Apex Assuance Group
Chandra Basavanna,SecPod
David Solin, jOVAL

Introduction

Further expanding the discussion on inter-networking devices (Routers and Switches) the NETCONF protocol will be discussed. NETCONF is an open standard protocol supported by major inter-networking vendors. This session looks to leveraging the NETCONF schema to retrieve the configuration files from inter-networking devices. This session will cover issues and challenges related to:

- Security Automation and inter-networking devices.
- Access methods to retrieve and process device configuration settings.

SCAP & NETCONF

The team was initially experimenting with JUNOS OVAL tests. They finished this work quickly and had some time left over so decided to take a look implementing a NETCONF OVAL test. They wanted to demonstrate leveraging an existing protocol to do SCAP OVAL assessment. Luis considers NETCONF an elegant way to access networking devices.

NETCONF

NETCONF is an open protocol with an RFC. It is not that complex, but allows sophisticated tasks do be done with it. Those are some of the primary goals for getting JUNOS into the OVAL specification. They were able to do this in less than three months.

Why NETCONF for SCAP?

Again mentioned the desire to leverage existing protocols within SCAP.

There are other similar protocols, e.g. SNMP. But NETCONF is specific to configuration management. NETCONF uses XML as does OVAL, so there is some inherent synergy. XPath can be used with NETCONF which works well with OVAL.

NETCONF Capabilities

NETCONF is used to connect to a device and get configuration information. Luis provided an overview of some of the NETCONF request types.

XML based, so could be parsed down to specific configuration items you wanted.

YANG data models are a logical next step to explore. The team did not use them, but is considered an area that could be leveraged, and they are being discussed in the IETF. NETCONF can be transported over secure protocols such as SSH.

Major device vendors support NETCONF out of the box. The configuration information is not in a standard format. JUNOS uses XML which made it easy to work with.

OVAL Tests

Really one test. There's something nice and simple about that. But the implementation is still evolving. They are trying to see how far they can push the one test concept.

Scanning Methods

Three OVAL SCAP scanning methods:

Local – OVAL scanner running on the device. Not supported by inter-networking devices (yet).

Remote – Connect to device via SSH or other protocol and retrieve information using NETCONF.

Offline – Connect to a configuration server and get the device's configuration information from there. Similar to an NDM.

Demo Content

David Solin of jOVAL showed the NETCONF OVAL schema they developed and discussed some of its features including XPath filtering of results.

They implemented the NETCONF test on both Cisco and JUNOS. Cisco provides a text – not XML – result, which limits you to using the substring and regular expression functions rather than XPath.

Discussion

Jim Ronayne, of Varen Technologies: I was thinking that using NETCONF would be standardizing the output. We don't want to prescribe the access method to the configuration information. We should write at a level of abstraction that does not specify "how" the configuration information is retrieved.

David Solin: That's an excellent point. NETCONF is an access protocol. But what you know is that the configuration is going to come back in a predictable way vs. the native schemas. If you have a generic schema, you kind of have to specify the way in which you retrieve the data. What you are saying is let's model the configuration in a schema and then it doesn't matter how it was retrieved.

Jim Ronayne: Right, NETCONF may be the way to do that. My concern is that there is a new dependency that the content is not going to be consistent if there is not that level of standardization.

David Solin: We were hoping the configuration information would be expressed in a uniform way; unfortunately once we got into this we found that not to be the case. There is work going on with YANG that may address that, but we did not delve into that. YANG seems to be a way to standardize applying configuration information; if there were a reverse transform, then that could be quite beneficial. But NETCONF is not ready for that.

Jon Baker, of MITRE: One of the things that struck me when looking at the NETCONF schema is whether we needed it at all. Could we just use the XML file content structure? NETCONF test looks very similar to the existing test. Is there a way we could make that work?

David Solin: That kind of reminds me of the way we hacked the plist test to work with offline scan for iOS; we ignore the file path. The other thing about the XML file object is it could be conceptually confusing because here (NETCONF) we are connecting to a remote device over the network; unless we hacked the filename to specify a NETCONF connection somehow.

Jon Baker: I understand that we don't want to hack the filepath. But the similarities, ignoring the file path, are striking.

David Solin: If the existing schema elements were what we wanted, then we wouldn't need a new schema element.

David Waltermire, of NIST: But on some level, I think it's the wrong question. We have plenty of other things that have similar capabilities within OVAL, e.g. PowerShell. The reason to move toward one of these management standards is to better align with the vendor efforts to expose these capabilities. It's something we've really struggled with – getting the vendors involved in the development of the OVAL tests.

Jim Ronayne: Are you suggesting vendors would be more willing to standardize their applet in a way that's useful to OVAL if working with NETCONF? Because right now I don't see that NETCONF brings anything to OVAL. But if it could be a standard form that we are getting from them, then that would be valuable.

David Waltermire: Yes, because NETCONF has clients outside of OVAL.

David Solin: In theory you could start writing checks for some vendor that supports NETCONF but does not support OVAL.

Jim Ronayne: But how would you write a check? You just showed JUNOS and Cisco provide very different content. You're basically doing a regex.

David Solin: But you would know about it. If I am familiar with what will be returned, I can start writing content.

Jim Ronayne: Why do you need NETCONF for that?

David Solin: At least you have a known means for accessing it.

Jim Ronayne: Would you want to put that in OVAL? We don't want to prescribe how you would get that. We want the tools to figure out how to get it.

David Solin: We can't get a file from the machine. I can't run a command on the machine. There's nothing in the OVAL schema that describes what it is that I want out of that machine. A NETCONF test prescribes an access method, but it also provides a generic capability to connect to any device that supports NETCONF.

Jim Ronayne: A key piece is that you know what the content is going to look like.

David Waltermire: Another piece is how do we get vendors to expose this instrumentation natively in their software? It's been difficult for us to get vendors to expose OVAL functionality in their software.

Jim Ronayne: There are lots of tools that already get config files & manage config files. That's everything that OVAL does, but not in a standard way. Why would we want to get involved in how they get and manage those files?

David Solin: Is the objection against retrieving generic information and then using domain knowledge to create compliance checks for it?

(There was continued discussion on whether using NETCONF as a way to get "generic" content that other tools can already retrieve is valuable vs. relying on other tools to get that content and use existing OVAL tests to evaluate the content.)

Adam Halbardier, of BAH: We should not rely on any approach that requires vendors to implement instrumentation. They have no incentive to do so.

David Waltermire: YANG or some other mechanism for specifying the schemata of what is returned.

David Solin: How is text file any better than NETCONF? Same issue of not knowing what the content is.

Wrap Up

Luis Nunez: Great discussion. Both Juniper and Cisco provide ways to expose their APIs into the network, so there may be opportunities there.

One last slide, another place NETCONF can be leveraged: Software Defined Networks (SDN) uses NETCONF and may be an opportunity for a NETCONF test.

We need to further engage network management.

OVAL – Database Vulnerability Assessment

Presenters
Charlie McClain,IBM
Louis Lam, IBM

Why should you care about database vulnerability assessment?

Databases are the leading source of compromised assets and according to Verizon, in terms of the percent of compromised records, databases accounted for 75 % of all compromised records in 2009 and increased to 92% in 2010. There's no reason to think it has decreased since then or will decrease in the future. A few interesting cases of breaches, that have occurred in the past in the past, were shown and the main takeaway is that 88% of all organizations have had at least one data breach associated with a database and when they occur; they are very costly.

What's different about database vulnerability assessment?

The first difference between database vulnerability assessment and other types of assessments that may be carried out, with OVAL, is that most breaches are perpetrated by insiders and these insiders

misuse their database privileges. Secondly, privileged information is stored in the database and its catalog. Furthermore, the catalog can be either centralized (e.g. Oracle) or de-centralized (e.g. SQLServer) and the structure of the catalog and its contents can vary depending on the type of database. This can add additional complexity to queries. Another way that database vulnerability assessment is different is because when determining whether or not a test passes or fails could be based on a query that has returned thousands of rows. It may also involve combinations of database users, objects, and specific privileges against those objects and a simple pass or fail will not be enough for customers to act upon. You need to let them know exactly why they passed or failed.

Why not just use <sql57_test>?

The first reason that the sql57_test shouldn't be used is because the <connection_string> entity exposes login credentials as well as database identification credentials which opens the door for attackers to target the built-in user accounts and default passwords that may be present on the database. Another reason is that the <connection_string> is associated with each test which hurts its usability. The next reason that the sql57_test shouldn't be used is because there is no link between the asset identification for a database and the actual connection string which leads to the need for some type of datasource. The sql57_test is also missing a few important DBMS engine types and doesn't distinguish between DB2 products. The current sql57_test documentation also does not make it clear whether or not complex multi-database queries are permitted and how to handle databases with decentralized catalogs. There may also be a need to support test categories and the ability to identify custom code mechanisms that support a test. Lastly, there is a need to support failure detail and exceptions from tests.

<connection_string> exposes credentials

The deprecated <sql_test> and the newer <sql57_test> provide for a <connection_string> element, but the username and password elements are in clear text. We know from experience that this is unacceptable to our customers, who are obligated to satisfy their auditors and legislative requirements;

<connection_string> exposes database identification

The <connection_string> entity contains information sufficient to uniquely identify a database. Even without credentials, intruders can use database identification information to "shotgun" login attempts to known users, using default passwords. Again, this is unacceptable to customers.

No link between <connection_string> and asset identification

The above issues make a case for removing the <connection_string> entity from the <sql57_test>. Furthermore, in asset identification, the only entity within a database asset type is instance-name and to uniquely identify a database, in some cases, you need to know the server name, port, and instance name. For example, if you have multiple instances of MySQL on a server, the only thing that distinguishes one instance from another is the port that the database is listening on. To correlate the database instance name, server name, and port, in Asset Identification (AI), a relationship could be used to join the ai:database and ai:service constructs where the ai:database provides the instance name and the ai:service provides the hostname and port.

We also need a way to tell the tool what database it needs to connect to without using the <connection_string>. To do this, we propose the introduction of a new asset type in AI called a datasource which would consist of a datasource-id and a datasource-name. There would be nothing here that identifies the actual database or exposes credentials, but, it would provide enough information to know which database was being talked about. The datasource referenced in the

proposed ai:datasource type would be an independent object (note: not OVAL object), created somewhere in the tool, encapsulating the connection information.

Question: Where are you suggesting putting that information? Are you suggesting that it is part of the content?
Response: No.

Question: So the example XML assets that you put up there, where would that live? How does that fit into the flow of things?
Response: This asset is part of the content and it identifies the datasource, but it does not provide connection details.

Question: A question that I would ask to the community is I know this connection string was in there, but, we really don't want that in there at all do we? We should be writing content more generically and maybe using some applicability tags to say that this content applies to an Oracle 11g datasource and maybe applies to finance applications or some categories of things, but, I didn't think we would want to have authors writing content for specific instances of specific databases. I think we want it more broadly, right?
Response: I think we touched on this before but it's exactly analogous to machines. You are not necessarily running an interpreter on the machine that you are scanning. You need to be able to supply credentials and it's out of the scope of SCAP.
Response: Right, I guess what I am suggesting is not just the credentials, but, all of that seems to be like a tool problem then a content problem. A content author shouldn't be involved in that at all.
Response: Yeah, exactly.
Response: And we are agreeing.

Comment: Because encoding a connection…now, I would think any practical use of the sql57_test would have to rely on external_variables to supply the connection strings and maybe that was the original intent.

Comment: For the most part, I agree that the tool really should be the one figuring out what the databases are on there and scanning. However, I think there are certain use cases where content may want to want to run only against this one instance so we did allow for the capability of saying if an instance is passed in of a database or name of the server to just do that one. That example was I think some of the HBSS content. It may not apply, but, if there is more than one database on the system, you may have content that is only applicable to a certain instance of the database.

Comment: I actually think there are a couple of challenges here. One is the applicability challenge. You could just as easily say I want to target this specific instance as you would say I want to target instances that look like this. Or, that have these properties? Or, something like that. And you're kind of getting to that. The other problem too is that you may want to target. You might want to use different credentials essentially with different checks so you might want to say I want to do this check at this security level and I want to do this other check at this security level right and so to that degree. The content author would have some sensibility about what security access or what level of access they would need to do those different checks. So, in those cases, we almost need a place holder in content that would allow the tool to essentially inject the appropriate credentials at the appropriate time.

Question: Should that apply to a more broad case? Should that apply more broadly like with other checks where we were talking about that yesterday with the HKEY_CURRENT_USER check? One possible solution, and I have seen it done manually, is to say this check should be run as a specific user, or all user accounts, as opposed to system. Should the content author, for any test, be allowed to prescribe at what level it is run?

Response: I think to answer both of your questions later on as we move to the slides to show how Guardium handles this you will see our approach in that how we do it today. The datasource itself is an independent object. The test itself has no relation to the datasource. At some point, where we actually want to run the test that is, when we create an assessment and we link any particular datasources to a test and then the tool itself will know that oh this is for an Oracle database, this is for a db2 database so therefore on that does the linking there. In addition to that, going back to Dave's question, is that the user we use is also another thing. Typically, in our environment, all of these financial companies they don't just give you a user that can access the database and do just about anything you want so in every DBMS type we do have a script creating the requirement; the privilege that we need, at a very minimum, to be able to access a database and successfully run the test so typically the requirement is very low. We do everything we can to prevent having system admin privileges. Where there may be cases where we are looking at password hashes, trying to break the hashes, and trying to figure it out to eventually figure out that password is really a default user from the manufacturer. Some of these get very complex where we would at some point need a much higher privilege. In some cases we, rather than pass or fail a test, we would actually throw an error saying that we just can't do this unless you create another datasource and grant additional privileges.

Question: I had a couple comments. It appears that the commonality elements in the content should not include information like connection string or logons or passwords obviously. Does that mean we create content to a certain level? Can we use the same content for multiple instances? If you have to custom make your content, for every instance that you are checking, then you are going to have something you can create your content on the fly. So, the whole deal with this is to be able to share content, write content once, and share it with other people and they can use that. If you try to do specifically to a datasource then you have lost that functionality or you are going to have to go in and make modifications every time you have a different instance. So, I think that is something that should be tool specific. In other words, that data should be able to be entered into your tool that says for this instance run this content rather than go to the point where you actually embed it into the test.

Response: I don't think I have explained it well. In our view of the world, the tests are associated with a DBMS type. Oracle, for example. Or, in some cases, with a specific version of Oracle or from this version on up, but, those tests are independent of what they are being executed against.

Comment: So, that would make it an applicability statement that you would run say in XCCDF and would only run these instances against these types of databases. So, normally you run XCCDF as a pair so it's the same thing as a computer right? So, it says only run these tests if this platform is XP or only run these tests if the platform is Vista, or whatever.
Response: Correct.

Question: Same with this right? So, it's more of an applicability statement that would go in XCCDF?
Response: Yes, but, the bottom line is that when you get down to it you need to be able to connect to the database somehow some way to be able to apply the content. So, the tool has to know what you are applying it too.

Response: That is really a tool problem and not a content problem?
Response: Only partially.

Comment: To clarify: I think Dave and you have actually said one of the challenges was that for certain queries you need to have, the query expects a certain level of access to the database and as you said you have want to. Luis talked about script you have to ensure that you create the right user with the minimal access need to execute the query.
Response: Correct.

Comment: There is a need to associate some level of permissions that are required to evaluate a given SQL statement and that is one thing, but that is not the same as the connection string. That's like they are kind of different needs. What you have suggested there is removing the connection string and when you started talking about the need to identify the asset, in our previous conversations, the way I walked away from it, from an OVAL perspective, what you want to be able to do is think of a system characteristics file where you have system information that talks about the system that was assessed. That it needs to better describe the system, and in this case, the system is kind of what you have there; it's your datasource. It's your server name, plus port, plus database name that you actually accessed. I see it as two different things that we are interchanging. I don't think you are trying to make any suggestion that we want to tie content to specific instances of databases in any way. I think you've said that content maybe it's for Oracle with a version number, or MySQL with a version number, or Postgres with a version number, but there is a need to have that content associate with a level of access to the database and there is a need to better define the system that you assessed on the output side.
Response: Yes.
Comment: I believe that it is appropriate to remove information from the test or evaluation descriptions that have no particular applicability to instance, venue, and the like, but, if we are to have a tasking system, as has been described earlier in the week, and have this type of information instance, the machine on which the database is located, the credentials with which to access the database, or the system on which the database is located, it would be prudent to adopt a uniform means by which such could be described because if it is relegated to the tool vendors we would likely have that many tool vendors or more methods of describing such information. So, while the ssql57_test is probably not the place to put this, it does have to lodge itself somewhere in the constellation of the XML languages used to describe SCAP in general.
Response: Okay.

Comment: I agree with what Gary is suggesting. We are essentially parameterizing the task with the credential information. I think the challenge will be how we inject that into the appropriate place so we need some type of place holder in the content that would actually receive the parameterizations so that the tool would know how to essentially deal with the appropriate parameters.
Response: Agreed, and what we threw up here on the slide was one proposal. There are other ways to do it and possibly better ways, but at the end of the day, the issue of concealing not only the credentials and the specific database identification information is the core issue here. The solution has to encompass all that; whatever the solution is, but, we are not in a position to say that we have the solution as we stand today. We simply know what the issue is.

Comment: I wonder if we are intermingling concepts about assets and then authentications. When we specify a file test, you need to have credentials to access that system in order to test the file. We don't specify those for any of the overall tests so I am struggling against doing that for databases. Who are we to tell folks how to authenticate to those systems. For example, with Microsoft it is very typical to setup

a service with the appropriate privileges where no one actually passes directly those credentials over the wire. Likewise, when we are talking about assets, this is really an asset problem. What we do need to specify is what is appropriate for a particular technology. For instance, Oracle, you could have a single database, but, any medium to large size organizations these days it's very typical for them to have a cluster of these databases.
Response: Thousands in some cases.
Response: Yes. Now we are talking test paths and then what versions of that technology is that applicable to, but even that, we have some mechanisms to handle. For example, this check is associated with this CPE so just some feedback and thoughts.

<engine> enumeration updates

The ind-def:EntityObjectEngineType, ind-def:EntityStateEngineType, and ind-sc:EntityItemEngineType enumerations need to be updated to include two new values: *netezza* and *teradata*. The enumeration value for *db2*, in these enumerations, should also be deprecated and replaced by *db2-luw* (DB2 for Linux, UNIX, and Windows) and *db2-zos* (DB2 for z/OS) because they have different code bases and catalog structures and should be treated as two distinct database types.

Multi-database queries

Some DBMS's (e.g., Oracle) store catalog information centrally, but others (e.g., Microsoft SQL Server) store certain catalog information inside each database. For these "decentralized" DBMS's, privilege tests must query the central catalog to retrieve a list of databases, and then loop through this list, executing the same query against each database's catalog. Pass/Fail is based on the results of the entire loop, encompassing multiple queries;

To address the issue of Pass/Fail being dependent upon result of assessing all databases, the current OVAL Test check_existence and check properties will probably suffice.

Possible Solution 1 (OVAL-centric)
Add a "decentralized" behavior to the <sql57_object> that can have a value of "true" or "false" as well as a <catalog_sql> entity that will hold the SQL query to discover the databases associated with a decentralized DBMS. If the "decentralized" has a value of "true", the tool would first execute the query contained in the <catalog_sql> entity, to discover the databases, and then execute the query specified in the <sql> entity on all databases found. Alternatively, you could force the user to write a SQL statement that will list the databases. Then use variables to concatenate another SQL statement in a for-each style loop. This second approach forces the content author to write significantly more complex OVAL content and more complex SQL statements. It is also difficult to test.

Possible Solution 2 (Tool-centric)
The existing <engine> entity, which is already present and required, is enough to tell tools whether the subject DBMS has a centralized or decentralized catalog. If the <engine> value indicates a DBMS type that is decentralized, the burden is then on the tool to carry out a two-step operation: 1.) Execute an appropriate database discovery tool for that DBMS; 2.) Execute the query contained in the <sql> entity iteratively for each discovered database.

Comment: I think that is one interpretation, but, there are cases where you don't want to preclude being specific to one database. So, if I want to say, you can only have these permissions or these privileges in the master database in SQLServer, I am looking at that as an independent database versus MSDB versus all of the user databases that have very different security requirements and I am going to

secure those independently and want independent reporting. So, I think it is useful to say that if anything on any of these databases fail, but, we have to make sure that we can still treat each as an individual thing.
Response: Certainly. I agree with you.

Question: It's my understanding that the way you would enumerate the different databases on a product, that has a catalog approach, is not consistent across products in any way. The tool vendors know how to do that now just out of necessity and it is going to be tied to the engine type. Were you suggesting that the method for enumerating things in the catalog would be expressed in the catalog_sql entity? Or, what were you hoping to convey in that entity?
Response: The actual method. Yes.

Question: Is that method always a SQL statement or is that method sometimes a script?
Response: That is a SQL statement always. And we do this today, as you said out of necessity, obviously in a proprietary way. Once we know the database engine and we know the version because I said the discovery queries don't change; sometimes they do by version. So, you need to know two things: the DBMS engine and the version of that engine. Once you know that, you know the form of the discovery query you need. Okay two options, we think this (option 2) is probably the better option, but, it clearly involves telling the tools that they know they have to behave that way when they see a test of that type with a decentralized catalog DBMS.

Question: If we take that option, it is basically the engine entity drives – the engine and the version entity in the test will basically tell the vendor okay use this option or this method to catalog things and that method to catalog if it is a different version.
Response: Correct.

Question: Were either of these to be adopted would that suffice to accommodate either extent or conceivable checklist content that might be applied to such a database. In other words, is this alone sufficient to describe everything that is necessary to be able to write a checklist regarding the security posture of a database?
Response: I think this is the point where I get to a point over here and say Louis do you want to take that one.
Response: If we can loop through every database, this is one of the challenges. But, as you see when Charlie goes through more slides, the other enhancements that we are hoping that the OVAL Language could improve to fully accommodate what we are doing today in a proprietary way as a full blown database vulnerability assessment.
Response: The short answer is that based on what we know today, from the way we do it in a proprietary format, yes, this would give us enough information to do this as a tool vendor for any of the databases we know about today. Okay, can we say that for every possible conceivable database that will ever be invented including even some of the unstructured databases, it will be adequate? No, I can't say that.

Question: Have you looked at unstructured databases like Splunk, or Hadoop, or Mongo DB, or any of those? Have you?
Response: I bought a book.
Response: Oh, okay. Thank you.

Comment: I think those are more file based. I think you would do a totally different test. I don't think you would be using <sql_test> or anything like this.
Response: I know it would be a different test, I was just curious if had done any work on that. That's all.
Response: We have been looking at the Hadoop ecosystem, and the big data HBase, and stuff like that.

Question: I am a little confused about option two because if you have a query you want to execute against every database in the SQLServer DBMS. How do you distinguish that from sql57 query that you want to execute on a specific database in that DBMS?
Response: By engine and version. That is the only way you can know.

Question: Wouldn't you need some kind of parameter in there to say we want to run on sqlserver1 right? You would have to say what server instance you would want to run it on. Or, have a ".*" or something? You have to have some wildcard to say hey we meant to run this test on every database you find or this one because there may be certain instances where…

Comment: I don't like the first option because I don't think you should have to specify how to iterate through the database, but you might need something to distinguish between the two.
Response: That is possible. Yes.

Comment: Well, I would think that in the case where the database don't have the same admin privileges you kind of have to iterate to find out which ones to access which way.
Response: In today's world, the way we do it is actually handled test-by-test for a DBMS type so the test itself will know that it needs to iterate through all the databases type in a given instance. If it doesn't, that is sort of like a partial assessment right so you can't really say that that particular SQLServer instance, for example, passes the test because you know you could either have lack of privilege, or we run into situations where customer put the database offline for one reason or another, or put it in a single user mode where you can't really iterate though that database, but the test itself, this is something that should be done at the test level. The test itself needs to know whether that particular test needs to iterate through the database or not. And then, to take care of that at the beginning is, from the datasource end, in our current solution. Take for example SQLServer, at the very beginning, we create a script which creates a role and we iterate through every possible database and make sure that role is in every database and have the required to privilege. If not, you run into issue later on it could be any test.

Support for Test Categories, Identification of Executable

SQL tests require a supporting code mechanism in any tool, and it is of course impractical to have a separate code mechanism for each test. However, SQL tests can be categorized and a code mechanism developed (and identified in <sql57_object>) for test categories. Some complex tests (e.g. CVE test) will still require a custom code mechanism, and there needs to be a way to identify a specific executable or Java class that supports a test because the <engine> and <sql> entities do not always provide the information needed to do so.

Possible Solution 1

Add an integer <category> entity to <sql57_object> and leave the use of <category> to the tool.

Possible Solution 2

Add a <category> enumeration to <sql57_object>, limiting the use of categories to enumerated items.

Question: Just to make it clear, what would the values of that category enumeration be? Hypothetically, do you have an example value that you might have there?
Response: CVE. It would tell me that when I see this test that it is a CVE test and I go to my piece of code that responds to CVE tests.
Response: But, why? It doesn't matter why you are doing something it is just a technical mechanism if you are running a SQL query to get an answer it doesn't matter if you are doing a CVE or whatever. I can see these categories if you have different technical mechanisms, but, what are you doing that's different from CVE versus something else?
Response: Oh, we do in fact run a SQL query to determine the patch level of the instance, but, then there is other code that compares that to the CVE and carries out a lot of different steps. Okay, it's not a simple SQL test, but, there is a SQL query necessary to respond to it.
Response: For each CVE, we actually have a repository of metadata say you know what version of Oracle, is it a rack environment , is it a grid environment, in the Oracle world, it is quite complex. It is not like oh you don't have those patches and you failed the CVE. Oracle addressed patching in the most complex way in all DBMS type there is CPU, there's PSU, there's bundle patches, exit data environment, there is bundle patch in the Windows environment, they all are different, and the patch numbers they don't increase in a sequential manner as well. So, we do have sort of like a mechanism that can fill in each test to populate and figure it out so we can actively say that this DBMS actually do pass this particular CVE.

Question: I know that Oracle patching is very messed up, but, shouldn't all of these things you are checking be encapsulated in the OVAL Definitions and then you have a higher level definition that describes the CVE? Because, if there is secret knowledge behind evaluating for a particular CVE that's not in the OVAL? Then it's not really an OVAL test.
Response: Okay.

Question: Is the reason that you are identifying this as a shortcoming because you have a test, the nature of which is a single test that asks the system have you any outstanding CVE vulnerabilities, as opposed to having that accompanied by SQL statement that is executed in the context of a CVE test?
Response: If I understand you correctly, the only way I can respond is we are not running a single test that says do you have any CVE vulnerabilities. Each CVE, and CVE is simply one category of test, but, each CVE is its own test which is applied by our CVE test mechanism.
Response: Let me ask it a separate way then, what would the content of the SQL element be for a test whose category was CVE?
Response: Ah, it would be a query that would determine the patch level of the database and obviously each DBMS has a different way of determining that patch level.
Response: But, that query is insufficient to the task of interrogating the...providing sufficient information to interrogate the system for purposes of divining a CVE?
Response: That's correct.

Question: If your SQL statement collects some information and then, based on the category saying CVE, you have a follow-on script that you run. Is there a way to take what's in the follow-on script and represent that in OVAL so that we can open that up so it allows other vendors to be able take and reuse content?
Response: I know that you can relate tests so that the success of one test is dependent on the success of another and so forth and that's related acknowledge that it is possible.
Response: I would offer that if you have a few examples that you could share, we would be delighted to work with you to figure out how to do it if it is possible in OVAL.

Response: We would absolutely be interested in that.

Complex Queries

Some database vulnerabilities require very complex queries, including, but not limited to: A.) Multi-table queries involving complex JOIN conditions; B.) subqueries, both simple and coordinated; C.) UNION SELECT queries involving potentially a dozen tables or more; D.) Anonymous block execution. We have queries supporting some of our vulnerability assessment tests that run to 100 – 200 lines of code. I'm not sure the existing OVAL <sql57_test> type can support all the variety of queries that we need;

Possible Solution

Include clear specification documentation indicating that DBMS-specific capabilities are supported by the test. This will require vendors to implement in a way that will allow for these DBMS-specific capabilities.

Support for Failure Detail

For complex privilege tests, the test query is typically something like "select count(*) where <condition> is true". The test should fail if "count" > 0. This seems to be supported by <sql57_test>, <sql57_object>, and <sql57_state>. However, our customers demand more than simple pass/fail information. They require specifics as to what caused them to fail one of our tests. For example, if we have a test that says, "No non-DBA user should have access to database objects of type X", our customers require that our assessment results tell them which non-DBA users have such privileges, and on which objects of type X. In both cases – the users involved and the objects involved – the count could run into the 100's or even 1,000's in a large environment, and telling the customer that he has to take our results and go off and manually identify the users and objects that caused him to fail is unacceptable. The implication of that is that when such a test fails, it needs to be able to execute a *second* query that will provide the failure detail.

Possible Solution 1

Add a <detail_sql> query to <sql57_object>, to be executed only in the event of test failure, which will return detail sufficient for the user to take remedial action;

Possible Solution 2

Stipulate that the contents of <sql> is instead the detail query, and that test passing/failure is determined by <sql57_state>.

Question: So, in the second option you would be returning the actual data that you would normally be counting in the previous query that you showed?
Response: Right.
Response: I think that is almost the kind of pattern we tend to use in a lot of the security automation content and, in that case, you would essentially be allowing OVAL to determine the criteria for success or failure.

Question: What if there were a million reasons for failure?
Response: In a large installation, I absolutely agree with Gary. That's absolutely a possibility, and just a sidebar here, the first question we have to answer when we go into any customer is what kind of impact is this going to have on my database and the answer had better be very, very small. So, throughout our tool, we try to do everything with the least performance impact on the monitored database that we can

possibly have. The guy is only interested in complete details if he fails. So, that's why we chose this method. Is it the only method that will work? No.

Comment: You could put a limit on the results returned as a way to try to mitigate that.
Response: There is no language support for that at this time.
Response: Well, no. The SQL supports that though is my point and you could sort the results so that you get the most important results potentially as well there are a variety of options.
Response: What I found is that it presents a shift in how people, who are in the database vulnerability assessment field, think about things. I think that what you have presented here with the first query and then the follow on if you fail query from everything that I have seen is very much the norm and what we kind of talked through the idea of it basically having the follow-on query be the initial query and then use an OVAL state to make assertions about the return values or counts and things like that it definitely represents a change there. And so building in things like using the limit capabilities and that sort of stuff into the SQL statements might be even further rethinking how we are doing those queries. On the other hand, there is a pretty mature industry there that does database vulnerability assessments. I think it would be naïve of us to say we are going to go in and change how everybody does everything.

Comment: We do that kind of thing already, in a lot of places, and it probably wouldn't work in all cases, but we could select users that have bad permissions and then we could have that you failed based on the presence of something there and the detailed data that you need is already there so I don't think it would work in all cases, and it wouldn't work with how it is defined today, but, there is a precedent for doing it in many other places in OVAL.
Response: Okay.

Comment: I think this doesn't just pertain to databases, this problem actually is spread out throughout the security configuration space. What I see here is that sometimes there is a need to collect evidence so that you can go ahead and act upon it. The problem here is that sometimes you can have way too much information or no information. You get a fail, but, you don't know why. One of the things, and we haven't solved this problem, but, one of the things that I think will help is that at a much higher level, when you are the security officers running these queries, you should have an option to say show me what's wrong with it. There are times when I really don't even know if I pass - okay. If this is a vulnerability, I pass. I am good with it, but, there are times where certain privileges, where certain users have a privilege that I don't want to them to have, and then I get a fail. Who is failing? Who is the user that has too much privilege that I need to address?
Response: And what object is it on and which privilege is it.
Response: Exactly. So, what I think right now the OVAL and standards that we have in place are not sufficient yet. It's evolving, but, we need to get there and one more thing is that a pattern that seemed to me could work is that you return the result set and say which one passed and which one failed because even when you have one that passes, you like to know why and you can filter. I want to see the ones that failed. I want to see the ones that passed. So, I could return all the users that had the privileges, but, I just want to pass on the ones that fail and have too many privileges and the ones that pass that didn't have the privileges.
Response: So, in addition to just displaying the users that actually just failed, the privileges and the majority of all customers demanded that, while we know that user needed that privilege, and in addition to that we also provide at a very high level which isn't really related to a specific test, a facility to put in an exception.
Response: That is a later slide.

Response: So you could have a violation, but maybe they acknowledge this is what they wanted so don't fail me because of this. So, it does run into a performance issue because then you are comparing potential thousands of violations against another list.

Exceptions from Tests

In the real world, particularly for database vulnerability tool vendors whose products are used by many customers, it is impossible to define "one-size-fits-all" tests that account for every possible variation in the way databases are configured. A good example of this is the CIS-mandated "Only DBA access to SYS.USER$" (CIS v2.01, Item 9.07). Certain installable Oracle components (e.g., APEX and XDB) install administrative users with such privileges, and the components *will not function* if you revoke those privileges and this may be unacceptable to a customer. For customers using these functions, then, tests must provide a way to exclude these users from the test.

Possible Solution 1 (OVAL-centric)

Add an <exception> entity to <sql57_object> to be observed by the test. The problem with this approach is that it exposes database user information in clear text;

Possible Solution 2 (Tool-centric)

Put the exception responsibility on the tool, so that the <sql_57> test knows nothing about exceptions. The problem with this approach is that you then have tests which, on their face, should fail, but they in fact pass with no visible explanation as to why.

Question: Is that exception always a user or does it have structure? Or, what is that?
Response: No, it isn't always a user. It could be an object.
Response: So, it's an OVAL object?
Response: No, it's not an OVAL object. It's a database object.
Response: I guess I am stuck with how it's different from a conditional written in OVAL or why it needs to be something special here as opposed to any other piece of OVAL content that I choose not to enforce. I don't want to have my password 14 characters. I choose to not comply with that policy. There are exceptions all over the place. What's special about the database case? How is this implemented?
Response: You may be right. First of all, if I am understanding you correctly, you are talking about a test that is not applied at all or an exception to that test?
Response: I guess I am stuck with, if this is an exception, as in, how is it different from an OR condition essentially on the when you wrote the test or some other logical combination of special tests? What is in here and what makes this? Why is this necessary as opposed to writing other types of content? Or, just tailoring content for particular users? I guess I just don't understand what's in here, what the structure of this would be and how it would be used.
Response: It's a string or a regular expression in our tool and it can be applied to whatever is returned from the <sql_test>.
Response: What do you mean by applied?
Response: That means it's compared to.
Response: So, it's a regular expression that's compared to what?
Response: Whatever the query returns.
Response: Whatever the results are in that record that comes back in the OVAL Results, you run a regex against, and if that string is in there, you say pass?
Response: Correct.
Response: Doesn't OVAL already have facilities for doing that?

Response: No, let me clarify. We exclude that from the test results if that was the only condition that caused it to fail and it's the only row in the result set then, yes, we pass. If there are others we fail.
Response: Why would you do that using the tools that oval already gives you? Why do you need a special case added on?
Response: Like a state. You can compare items to states. But, for instance, a regular expression in a state. Or, if you didn't want it to be in the content, you could have it refer to an external variable.

Comment: Can you go back one slide maybe? So, I think what he is trying to get at, and tell me if I clarify this incorrectly, is so here we have an example of a configuration item where you could imagine checking to make sure that only the DBA user has access to the sys.user.$ table or something like that, and in this one organization, Charlie is an extra important guy and he needs access to that and…
Response: Can you tell my wife that.
Response: Well, she might like it if it is just one organization. I guess so really the way I equated this need is more along the lines of the POA&M process whereby, sure there is a finding here, you are not compliant with this policy, but, at the organizational level we've said that's fine. Charlie needs to have that access and so it wasn't necessarily something we wanted to change about the check itself. The check still stands as it is and there is nothing that needs to be changed there. It's that our organization has a policy that says Charlie needs to have access to all of these database tables.
Response: That is actually what I was trying to say earlier, but, was not doing a very good job of it.
Response: So who writes that and where does it apply?
Response: Wouldn't you just implement that as a new test for your organization? You would be substituting, you might add a filter, or a set complement operation, or a state that takes a variable, or any number of things that wouldn't require adding anything.
Response: So, I had taken a slightly different route I guess. Thinking more along the lines of the POA&M process there, where any vulnerability or configuration audit finding you might have to say we are going to accept that risk and we don't want to be yelled at over and over again for that finding. We have acknowledged the risk and we've moved on. At MITRE, we have an internal process whereby we waiver things so that rather than get cited for not having the latest patches on my system. Well, we realized it is a special system, and gets waivered, so it doesn't get that finding anymore. While the check has changed, there is no variable input to the check, the check is still the same.

Question: I thought OVAL was simply to make statements about the state of a system. It's either the sys.user.$ table is accessible only by DBA, or it's not. It's true or false. If you want to make qualitative statements, then that is up to XCCDF, and in XCCDF, you would use deviation. Wouldn't you? That seems like what this is attempting to solve?
Response: But you are not deviating from the entire test, you are still applying the test. It is only a particular user, object, privilege, whatever, a particular thing that you are excluding from that test.

Comment: First of all, for the apex and xdb scenario, it would be prudent for the benchmark author to structure the test to include the more complex scenario: if not apex nor xdb are installed then only DBA access should be found and not something other than that, but, because in the absence of that it's a bald statement. If it's not DBA, you fail the test.
Response: Right.
Response: Having failed the test then there would be some other user that would be found to have access to sys.user.$. If that is acceptable to the organization, and as Jon mentioned, you wish not want to be nagged by this, a generic benchmark would not going to have any *a priori* knowledge that the organization confidences someone else having access to that table, but, were the organization to do it rather than casting it as a benchmark item, it would likely be done outside the context of OVAL. The

only other way this could be done is either edit the benchmark, the organization would edit the benchmark using one of the techniques mentioned such as filter and so forth to apply this exception, or it would be done as Jon mentioned totally outside of OVAL and this would be something that would be applied to the results of the evaluation, post-evaluation to say that this is okay. That last part is a little bit difficult outside the context of OVAL because some of the information necessary to detect the precise nature of the violation would not necessarily be available.
Response: I don't disagree with the thing you said. The only response I can make to you is to tell you what we hear from our customers. Apex or xdb aside, things that we can disable or, somehow or another not include in the test, we do have customers who just say I want Fred to have access to this table even though I have a test in here that says nobody other than a DBA can have access to this table. That is what they tell us.
Response: If that is the case then it would be necessary to create OVAL that included such conditional evaluation.
Response: And, I understand from an OVAL standpoint as a tool vendor with predefined tests that are encapsulated in, from the user, without the user paying me extra to create a separate test or creating his own user-defined test. I don't have that option.
Response: That is a totally antithetical approach to the way that OVAL has been practiced over the years. Everything is explicated in an OVAL document. None of it is hidden behind the scenes with tests of varying natures are not triggered by information strictly supplied in the OVAL document.
Response: Understood and I don't disagree, I am just telling you how our customers respond.

Comment: I think it hints at a more complex use case that we have discussed in the past where you're configuration guidance might be: "only appropriate users should have access" and the way that you would do that, because appropriate is subjective, you would use need both OCIL and OVAL where you would use OVAL to see who does have access and then you would send that to an OVAL check which would be presented to an authorized person to say, yes, he should have access or no he shouldn't have access. So, unless you are going to predefine, in some OVAL Variables, a list of who approved, or who should have, and define who should have access, the only other way would be to combine the two and that is not a use case that we really built to yet.
Response: Okay, I want to move on because I don't want to take up all the extra time on this one issue. The issue has been explained and there's a lot of discussion, but, we are not going to resolve it here. So, let's move on.

Where We Want to Be

To demonstrate where we want to be as opposed to what is currently possible with the sql57_test, we will demonstrate the capabilities of our product using various screen captures of it.

Demo of IBM InfoSphere Guardium with Screen Shots

The first screen capture shows how a datasource can be added to Guardium. This is how Guardium currently handles the information contained within the connection string. This information is stored in a locked-down version of a MySQL database where all of the information is encrypted. In the datasource, you can also specify the datasource name, type, severity, login, and password. Note that a script is used to properly grant a role or group to that user to carry out the vulnerability assessment. Also, some fields are left blank because they don't apply to the particular database. Lastly, at the bottom, there is Change Audit System (CAS) section. CAS provides the ability to find out things that change at the operating system level or see that a database configuration file has changed even though it may not be possible to query for that information. CAS is used to look at important binary file permissions and group

ownership. For example, it is possible to determine if someone changes one of the files, it will be reported and the tool can determine if it is a violation or not.

The next screen capture was of the security assessment screen which showed that numerous assessments were created for Teradata, Oracle, and Sybase IQ databases among other things.

After that, a screen capture showing how to create a security assessment was shown. This involved specifying an assessment name, how long we want to keep the information around, and other things as well as selecting a datasource(s) to use in the assessment. This screen shows that the assessment should be performed against DB2 10.1, MSSQL 2012, Oracle 11, and Sybase 15.5 databases. During this setup it was also pointed out that there was another facility (CAS) to configure what type of operating system, file, etc. that you wanted to monitor.

The next screen capture showed a list of tests available for each DBMS which is denoted by a specific tab. For example, DB2, z/OS, Informix, SQLServer, MySQL, Netezza, Oracle, Postgres, Sybase IQ, and Teradata are all supported. There are plans to add support for DB2 for the iSeries in the next release and add support for the Hadoop and the big data ecosystem further down the road. It was also noted that the tests are broken down by type: Pre-defined, Query-based, CVE, APAR, or All.

The following screen capture displayed the results summary of running the vulnerability assessment. It gives you a score, a breakdown of the tests, findings, recommendations, and its size could vary from a few pages to hundreds of pages. It was also noted that the summary can be exported into a PDF, it has a query facility where users can query the results, and in Version 9, the ability to export the results as AXIS XML or an SCAP result set was introduced.

Next, a screen capture detailing the results of a specific test was shown. Here you can see if a test passed or failed, what the test was checking, as well as provides recommendations on how to fix the problem.

After that a screen capture of the exception test facility was shown. This allows users to add a regular expression to describe exceptions for a test in an assessment such that the next time a test was run, if no violations are found after the exceptions are applied, the test may change from a fail to a pass.

Finally, the last screen capture shows an example of a CVE test (CVE-2012-0082). Specifically, a CVE test that applies to an Oracle database. It was noted that most of the information about the database was imported from the CVE registry.

Comment: The approach that we agreed on with the OVAL people is that we have to fix the immediate problems with the sql57_test first and add support to OVALDI as a proof-of-concept so that people have a supporting tool to actually develop some tests and try out some of these scenarios to see what does work, what doesn't work, and then more issues will arise from that and it is going to be a longer discussion.

Question: What does a score mean? Is it the number of pass and fail tests? Or, the sum of the CVSS?
Response: On the CVE test, they actually have a specific score for pass or fail, but, if you look at the only very screen here, it gives an overall score. It's really not a weighted score, it is just the total number of tests that passes versus failed and just give you a score. We haven't gotten to the point where we put in a weighted average for tests so for certain tests where you fail, it may have heavier weight than others.

Response: We report the CVE weighted score, but in the overall assessment score, that test is only counted as 1 because our other tests don't have a weighted scoring system and there's been some discussion, early in the week, about scoring systems and consistency in scoring, but, we are not there yet certainly.

TAXII Adoption

Presenter
Aharon Chernin, DTCC

Note: Only discussion points are captured here. Briefing content is not captured.

Q: I am wondering if you are planning to provide a list of products that conform to these specifications.

A: I'll have to check with legal to see if it's possible, but I would like to do that. We would not recommend products; we would just list products that meet a certain criteria.

Q: Do you have plans to bring this into working groups that exist?

A: I don't care where this ends up; I just want my use case met. The venue doesn't matter as long as the capability exists.

Q: What is the volume of data that you have for threat intelligence from your members so far?

A: I won't answer that, but it's not extremely heavy. Some of it may be due to the nature in which we have to process it.

Q: Is anyone using RDF? If so, what ontologies are they using?

A: I do not know of anyone using RDF.

Q: If a vulnerability scanner was to have a TAXII interface, do you envision the ability to raise issues?

A: If a vendor wants to integrate it, it will work better.

Q: Have you addressed the transport mechanisms?

A: We see that topic as an issue, and we are not sure we want to fix that issue in the first few phases. We are trying to share the data using existing transport mechanisms.

Q: What's the model for sharing platforms?

A: Right now we are using existing distribution models.

Q: So you say you want to take baby steps, can you tell us what your ultimate view is?

A: In the ten years from now, I would hope that there would be an API, automation, and that tool vendors are automatically creating threat data.

Q: What do you mean by API?

A: To me, if a web service can duplicate what a programming API can give me, I'm OK with that as well. It's whatever allows me to get my job done.

Q: What is the ultimate goal here?

A: Threat intelligence today is seeing whether or not you've already been compromised.

A: The ideal is that one institution will share their information, and others will use that information to protect themselves if they can.

Q: You said that hopefully you can identify who the bad guy is.

A: Definitely.

Q: I'm still not understanding the request for an API. A programmatic language binding vs a web service.

A: I need a programmatic API to build the XML.

Q: Right now there are python bindings to XML, but that's very close to the XML schema. What I felt that you were describing is a more abstracted mechanism for building the XML.

Q: What do you see as a problem from a legal perspective?

A: We are already doing this, and there are already legal agreements in place. Legally, all member firms should be aware of what their sharing and that it's not tool signatures.

Q: That works for two types: Def Gov't and Vertical Industry partnerships. How does that work for the rest of the world?

A: From the DHS/US-CERT perspective, one of the roles we see is to help bridge the gap between the different sectors and the different ISACs. It is important to us that we enable sharing across sectors. The hard part is getting the lawyers to agree, the easier) part is the transporting the data.

Q: What about those who don't want to share that they are currently compromised?

A: You can either not share it, or specify that it should not be shared within the transport format. This is the same data we are sharing already, just in a different format.

A: When you start sharing context, it gets dicey, since you are sharing internal configuration information, etc.

OCIL

Presenters

Jim Ronayne, NSA

Shane Shaffer, G2, Inc

Overview

The plan for the session is to run through several proposals for enhancements for OCIL. All proposals have been described on the OCIL Discussion Forum, so none of them are new. Since the issues have been on the list for a while, the hope is that final decisions can be made at this event, the specification will be updated and NIST will be asked to put it out for the final 30 day review period. A reference implementation has been written for these proposals and tested, so the issues have been worked through.

There are two major categories of proposals: enterprise usability proposals and capability proposals. Some of these proposals will involve other standards, such as XCCDF and CPE.

Enterprise Usability Proposals

Jim Ronayne presented this first section of the OCIL session.

Proposal 1: Questionnaire routing (XCCDF)

Routing is the process of getting the questions to the right respondent. The expectation is that, within a given benchmark, one person is not going to be able to provide all the answers for all the questions, but we want to avoid the necessity of creating different benchmarks for different classes of people. The assumption is that OCIL questioning will usually be managed centrally and vendors are responsible for the method to assign questions to respondents. We should allow the content author to recommend the type of person who should provide data for each rule.

The solution is to use the existing metadata tag in XCCDF; not a requirement in the specification, rather a best practice to make it easier for the tool to direct it to the right people.

Optional feature, processing not required.

Dave Waltermire suggested using a queue name for the text content, which allows you to associate the name-space to a prefix. Jim concurred.

Dick Whitehurst asked if there was some way to wrap a context around that? E.g. the Financial DBA, as opposed to the Engineering DBA. Jim suggested this question was more in line with the Applicability proposal. There was then a discussion as to how specific question targeting can be or should be, and Jim suggested that only characteristics of the intended respondent should be provided.

A questioner asked how a vendor should interpret a list of roles provided by an “href” statement, was it a simple list or expressed in XML. Jim stated that the assumption was that it was a simple list.

Jim then mentioned that they were proposing to use the meta-date at the Rule level. But they also considered that XCCDF currently only has meta-date at the Benchmark, Group, and Rule level, but not at the Check level, and, it might be useful to have the meta-data at the Check level. But since that requires a change to XCCDF, they are recommending to specify it at the Rule level. However, they are likely to propose adding meta-data to the Check level at some point in the future.

Jim also explained that they considered whether modifying XCCDF for this feature is appropriate, or would it be better to modify OCIL. By making the change in XCCDF, the one use case that is precluded is rules that go to questionnaires which have dependent questions that go to different people, i.e. a questionnaire that goes to Alice and then, based on her answer, may or may not go to Bob. When tested, this use case was difficult to implement and had no obvious adaptations in the real world.

Justin Furniss suggested that there is a small use case which makes it appealing to make this change at the OCIL level. That is, it would allow an OCIL interpreter to have a smaller footprint, i.e. an XCCDF interpreter would not be needed, and thereby making it more practical to deploy on a small mobile

device. Jim observed that there are not many good places to put meta-data into OCIL and, regardless of this proposal, it would be a good idea to insert some meta-data into OCIL for future experimentation.

An attendee suggested that, since there were no practical use cases for specifying roles outside of OCIL, therefore the specification of a role belonged in OCIL. Charles Schmidt countered that benchmarks can have some questions for clients and some for servers, so moving the specification of the role in XCCDF was justified.

Jim provided an example for Routing:

```
<Rule id="xccdf_com.example_rule_Rule1" selected="false">
    <title>Unused DBMS components should not be installed.</title>
    <description>Are there any unused DBMS components installed?</description>
    <metadata>
        <eocil:target_respondent_role href="http://iase.disa.mil/roles">DBA</eocil:target_respondent_role>
    </metadata>
    <check system="http://scap.nist.gov/schema/ocil/2">
        <check-content-ref href="example-ocil.xml" name="ocil:com.example:questionnaire:1"/>
    </check>
</Rule>
```

Don Campbell commented that conventions for the definitions of roles should be documented somewhere. Don also made the point that the group should be careful not to confuse the definitions of Role and Applicability.

David Solin suggested that these proposed changes don't go far enough, that Roles could refer to more than just OCIL checks, but also to other checks. Dave Waltermire asked how such a change would handle the situation where more than one respondent was needed to answer various stages of a questionnaire. Solin suggested breaking it into multiple questionnaires. The Waltermire asked how to relate previous questions in those cases, and Solin suggested that import/export could be used. Jim Ronayne then cut in and said that one goal was to not create the dependency on one person's questions being dependent on another person's answers, and then opened the issue to the attendees and asked if that was a desired goal.

Jim then summarized by saying this capability was desired, and using the meta-data tag is acceptable, and that XCCDF should add a meta-data tag at the Check level. There was further discussion about other alternatives, but the group decided that a different approach would lead to more complexity which was not desirable at his time.

Proposal 2: Applicability (CPE-Language, XCCDF)

This issue has been discussed at previous Developer Days meetings, and the current proposal has grown from those discussions. There is currently a notional Applicability capability in XCCDF using the Platform

tag. A key aspect to this proposal is that it has no effect on CPE. The goal is to develop a machine-readable mechanism for determining where a benchmark should go – i.e. this involves more than just OCIL.

The proposal is to add DataFactRefType to include externally defined attributes.

Dave Waltermire suggested that the implication that a noun implies that it should be processed could be dangerous and further suggested that instead, a flag should be added to the new check that would say what to do in the case that it's unknown. Jim acknowledged the point.

Jim then added that FactRefType could be expanded to include other standard identifiers, e.g. CVEs and CCEs. Jim then showed the following example:

```
<                                                                    platform-specification>
                                <platform                                    id="windows_7_mac-1_public">
                          <title        xml:lang="en-US">Windows        7        (MAC-1_Public)</title>
                                <logical-test                operator="AND"                negate="false">
                                 <fact-ref        name="cpe:2.3:o:microsoft:windows_7:*:*:*:*:*:*:*:*"
system="http://www.mitre.org/CPE"/>
                                                <data-fact-ref                description="MAC-1"                >
                                               <asset_property     xmlns="http://mil.disa.asset_properties/1.0">
                                                                                <mac_level>1</mac_level>
                                                                                </asset_property>
                                                                                </data-fact-ref>
                                                <data-fact-ref                description="Public"                >
                                               <asset_property     xmlns="http://mil.disa.asset_properties/1.0">
                                                                <confidentiality_level>Public</confidentiality_level>
                                                                                </asset_property>
                                                                                </data-fact-ref>
                                                                                </logical-test>
                                                                                </platform>
</platform-specification>
```

Dick Whitehurst asked about the example and Jim explained that there should be a schema behind the namespace shown in the example. And this schema would constrain the values that are chosen. Jim then added that the use case assumes that there would be an asset management system that is closely tied to the evaluation system.

In response to a comment, Jim made the point that this proposal only works if the asset management system uses the same tags as the benchmark author.

Dave Waltermire suggested using the same attributes that are defined for ASR to represent meta-data concepts. Jim concurred.

Jim provided two more examples:

```
<platform-specification>
<platform id="network">
<title xml:lang="en-US">Network</title>
<logical-test operator="AND" negate="false">
<data-fact-ref description="network" >
<asset_property xmlns="http://mil.disa.asset_properties/1.0">
<asset_type>network</asset_type>
</asset_property>
</data-fact-ref>
</logical-test>
</platform>
</platform-specification>
```

```
<platform-specification>
<platform id="CVE-CCE_combo">
<title xml:lang="en-US">CVE and CCE combo</title>
<logical-test operator="AND" negate="false">
<fact-ref name="CVE-2012-1234" system="http://cve.mitre.org"/>
<fact-ref name="CCE-4321-6" system="http://cce.mitre.org"/>
</logical-test>
</platform>
</platform-specification>
```

David Solin asked where the Applicability change would be made. Jim explained that the idea is to take the "CPE Language Specification", which is one of the four CPE specs. This proposal would take that document and take "CPE" out of the title and make it a larger, more encompassing spec. David then wondered about running benchmarks, and who decides to run which benchmark on which assets, which is the "unspoken" policy that lives on top of all this SCAP stuff. Jim responded that this Applicability proposal allows an author to declare that this policy applies to assets that have the following properties, basically defining what's in and what's out. David insisted that there should be some way to make this designation outside of OCIL. Charles Schmidt explained that what matters is that if the benchmark fails the platform check, then the result becomes not applicable (not a failure).

Proposal 3: Timestamp granularity (OCIL)

This proposal adds a timestamp attribute so that the time that each respondent answers the questions in a questionnaire. One outstanding question is at what level should the time be: at the questionnaire, the test action, the question? Jim was inclined to put it at the lowest level.

A questioner asked whether all the responses from multiple respondents would be combined into a single document. Jim replied that, yes, he assumed that all responses would be combined in a single document.

Proposal 4: Inclusion of AI (OCIL)

This proposal is to make inclusion of the AI schema follow the same model used by OVAL. One clarifying point that Jim wanted to make is that it should not be assumed that when a tool is running content the system it's on is not necessarily the subject of the content. Another minor point, the term "system" is used in the schema, and Jim plans to replace that with "subject."

Charles Schmidt wanted clarification on whether "subject" referred to the party providing answers or the party that is the subject of the assessment. Jim replied that "subject" refers to the "thing it's about" and "respondent" refers to the thing providing answers.

Dave Waltermire wondered whether "respondent" is a person or an asset that provided answers. Jim replied that practically it's a person, but in theory it is an asset.

Proposal 5: Response provider on question result (OCIL)

In the case where there are multiple people responding to questions, each person needs to be identified as a "submitter."

Proposal 6: Multiple Results Sets (OCIL)

This proposal allows for more than one set of results in a single file. This will be done by adding ResultType under ResultsType

Jim provided an example:

```
<results>

        <result start_time="2012-06-01T00:00:00" end_time="2012-06-01T00:10:00"
subject_id="ocil:com.example:system:1">
                <questionnaire_results>
                ...
        </result>
        <result … subject_id="ocil:com.example:system:2">
                …
        </result>
        <targets>
                <system id="ocil:com.example:system:1">
                        <ai:computing-device>
                          <ai:hostname>myBox</ai:hostname>
        …
        </targets>
<results>
```

Capability Proposals

Shane Shaffer presented the capability proposals for OCIL.

Proposal 1: Exceptional answer handling

Exceptional answers are things such as "not applicable", "unknown", etc. The question is, when should they appear for the user to choose them? The specification is not clear on this. Since the same question can be asked multiple times, we want to avoid the situation in which "not applicable" is sometimes a legal answer for a question and at other times "not applicable" is not legal.

This proposal will allow the question to explicitly identify the allowed exceptional answers. If "not applicable" is not added as a legal response to the question, it does not mean that the result cannot come back as "not applicable" – it just means that the user cannot say "not applicable" as an answer to that question.

Shane provided the following example:

```
<numeric_question id="ocil:org.namespace:question:1">
<question_text>How many times a month do you floss your teeth?</question_text>
<allowed_exceptional_answers>
        <answer>UNKNOWN</answer>
        <answer>NOT_APPLICABLE</answer>
</allowed_exceptional_answers>

</numeric_question>
```

Dave Waltermire asked how the result could be "not applicable" when the user cannot answer as "not applicable." Shane explained that if it's a legal answer within the handler, then interpreter can apply that legal answer to be returned as a result "not applicable." Dave and Shane agreed that such a situation would be very unusual. But Shane pointed out that there's a difference between the response "not applicable" and the result "not applicable."

As a bonus enhancement, Shane pointed out that if there is no handler for the "not applicable" then by default return "not applicable."

Proposal 2: when_else handler

Currently in OCIL, unhandled answers return an error. Therefore, every response that you expect must be explicitly handled, which can be very difficult in some circumstances. The proposal is to add a when_else handler to all types that extend QuestionTestActionType and it acts like an "else" statement. If the answer is not handled, the when_else handler is invoked. If there is no when_else, then an error is still returned. This applies only to non-exceptional answers.

Shane provided this example:

```
<numeric_question_test_action question_ref="ocil:org.namespace:question:1"
id="ocil:org.namespace:testaction:1">
        <when_equals>
        <result>PASS</result>
        <value>2</value>
        <value>3</value>
        <value>5</value>
        <value>7</value>
        <value>11</value>
        <value>13</value>
        <value>17</value>
        <value>19</value>
        <value>23</value>
</when_equals>
<when_else>
        <result>FAIL</result>
</when_else>
</numeric_question_test_action>
```

Proposal 3: Answer restrictions

Today there is no way to constrain answers to open-ended questions, which can result in getting back bad data. The proposal is to add some restrictions to open ended questions:

- Numeric
  - Integer vs. Decimal
    - Decimal precision
  - Min/Max values
- String
  - Pattern
- Choice *
  - Number of choices

Shane provided these examples:

```
<numeric_question id="ocil:org.namespace:question:1">
<question_text>What percentage of people are left handed?</question_text>
<answer_restriction>
        <datatype>DECIMAL</datatype>
        <precision>2</precision>
        <range>
                <min>0</min>
                <max>100</max>
        </range>
</answer_restriction>
</numeric_question>
```

```
<string_question id="ocil:org.namespace:question:2">
<question_text>What is a 8 letter word starting with a q that is not followed by a u and does not end with s (please use lowercase)?</question_text>
<answer_restriction>
        <pattern>^q[a-tv-z]{6}[a-rt-z]$</pattern>
</answer_restriction>
</string_question>
```

Proposal 4: Multiple Response questions

Currently, choice questions only support a single answer. The spec says this is the case and suggests how you can work around, if needed. There is a desire to allow the respondent to select multiple answers. The proposed solution makes changes within ChoiceQuestionType: add an optional attribute, “multi” (true/false); chang default_answer_ref from a single choice ID to a list of Choice IDs; and add an answer restriction.

Shane provided this example:

```
<choice_question id="ocil:org.namespace:question:1" multi="true"
default_answer_ref="ocil:org.namespace:choice:1 ocil:org.namespace:choice:2
ocil:org.namespace:choice:3 ocil:org.namespace:choice:4 ocil:org.namespace:choice:5">
<question_text>Which days of the week do you work?</question_text>
<choice id="ocil:org.namespace:choice:1">Monday</choice>
<choice id="ocil:org.namespace:choice:2">Tuesday</choice>
<choice id="ocil:org.namespace:choice:3">Wednesday</choice>
<choice id="ocil:org.namespace:choice:4">Thursday</choice>
<choice id="ocil:org.namespace:choice:5">Friday</choice>
<choice id="ocil:org.namespace:choice:6">Saturday</choice>
<choice id="ocil:org.namespace:choice:7">Sunday</choice>
</choice_question>
```

Changes must also be made to ChoiceQuestionTestActionType, where the logic of processing the answer is done - this gets more complicated. There are no changes needed for when_choice. Changes for when_count is use it to define a set of choices and then declare how many from that set have to match. Finally, when_choices is a set of choice_ref, choice_count, and choices, and is used to combine when_choice, when_count, when_choices with boolean operators.

Shane provided this example:

```
<choice_question_test_action question_ref="ocil:org.namespace:question:1"
id="ocil:org.namespace:testaction:1">
<when_choice>
        <result>PASS</result>
        <choice_ref>ocil:org.namespace:choice:6</choice_ref>
</when_choice>
<when_count>
        <result>PASS</result>
        <check_count operation="greater than or equal">2</check_count>
        <choice_ref>ocil:org.namespace:choice:1</choice_ref>
        <choice_ref>ocil:org.namespace:choice:2</choice_ref>
        <choice_ref>ocil:org.namespace:choice:3</choice_ref>
        <choice_ref>ocil:org.namespace:choice:4</choice_ref>
</when_count>
```

```
        <when_choices>
        <result>FAIL</result>
        <choices                                                    operator="OR">
                <choice_ref>ocil:org.namespace:choice:7</choice_ref>
                <choice_count>
                        <check_count          operation="less          than">2</check_count>
                        <choice_ref>ocil:org.namespace:choice:1</choice_ref>
                        <choice_ref>ocil:org.namespace:choice:2</choice_ref>
                        <choice_ref>ocil:org.namespace:choice:3</choice_ref>
                </choice_count>
        </choices>
</when_choices>

</choice_question_test_action>
```

Proposal 5: Information only

Currently, OCIL evaluates everything as pass or fail. This is not necessarily appropriate for some questions, e.g. where you would like to collect information. There was some discussion over the relative merits of returning pass/fail values, including consideration as to how OVAL handles this.

The proposed solution is to add a new value to ExceptionalResultType, which is INFO_ONLY, which would be used as the result of a Questionnaire/Test Action when you just want to collect data but not judge that constituted a pass or a fail.

Don Campbell suggested that this might muddy the water because it would make OCIL less analogous to OVAL. David Solin concurred.

Proposal 6: Display hints

OCIL purposely doesn't dictate how the questions are presented to users, but being able to include something to influence the display may be useful. This is not just aesthetic concern, because there can be impediments to data entry. There will be two pre-defined (for string questions and choice questions, but beyond that It is extensible. For a string question, you can specify essay or short answer; and for a choice question, you can specify dropdown or radio.

Harold Booth suggested adding check box and list for choice questions.

Kent Landfield requested that, if DateTime is added, then it supports global DateTime format and not just the US. Don Campbell added that the same comment goes for addresses and phone numbers.

Shane provided three examples:

```
<string_question          id="ocil:org.namespace:question:2"          display_hint="short_answer">
<question_text>What          is          your          name?</question_text>
</string_question>
```

```
<string_question id="ocil:org.namespace:question:4" display_hint="essay">
<question_text>What did you do on your summer vacation?</question_text>
</string_question>
```

```
<string_question id="ocil:org.namespace:question:5"
display_hint="ocil:hint:vendor_x:mycustomformat">
<question_text>How else would you like this question displayed?</question_text>
</string_question>
```

John Wunder mentioned that this is already done by HTML5 and he wondered if it had been considered to just adopt those conventions. Shane said he would look into it.

Proposal 7: "Other" option

Frequently questions can have an "other" option, e.g. "what's your favorite flavor of ice cream: vanilla, chocolate, strawberry, or other?" This can be annoying to implement in OCIL. The proposed solution is to add an "other_option" to choice questions. This would allow inline "follow-up" as the initial response, allowing you to ask a numeric or string question.

Shane provided this example:

```
<choice_question id="ocil:org.namespace:question:1">
<question_text>What is your favorite flavor of ice cream?</question_text>
<choice id="ocil:org.namespace:choice:1">Vanilla</choice>
<choice id="ocil:org.namespace:choice:2">Chocolate</choice>
<choice id="ocil:org.namespace:choice:3">Strawberry</choice>
<other_option type="string"/>
</choice_question>
```

```
<choice_question_test_action question_ref="ocil:org.namespace:question:1"
id="ocil:org.namespace:testaction:1">
<when_choice>
        <result>PASS</result>
        <choice_ref>ocil:org.namespace:choice:1</choice_ref>
        <choice_ref>ocil:org.namespace:choice:3</choice_ref>
</when_choice>
<when_choice>
        <result>FAIL</result>
        <choice_ref>ocil:org.namespace:choice:2</choice_ref>
</when_choice>
<when_other>
        <when_pattern>
                <result>PASS</result>
                <pattern>^.*berry.*$</pattern>
        </when_pattern>
        <when_pattern>
                <result>FAIL</result>
                <pattern>^.*mint.*$</pattern>
        </when_pattern>
</when_other>
</choice_question_test_action>
```

David Ries thought this would be better handled as a display optimization by the tool, that for the "other" option the test action handler can route you to a string question and it displays it then inline.

Proposal 8: Sequence questions

A sequence question is one where user is asked to put items in order. To implement this, add a SequenceType, which is ordered set of ChoiceType instances; add a SequenceQuestionType; and a SequnceQuestionTestActionType.

Shane provided this example:

```
<sequence_question id="ocil:org.namespace:question:1">
<question_text>Put these events in chronological order, starting with the oldest</question_text>
<choice id="ocil:org.namespace:choice:1">America is "discovered"</choice>
<choice id="ocil:org.namespace:choice:2">Neil Armstrong walks on the moon</choice>
<choice id="ocil:org.namespace:choice:3">OCIL 2.0 is released</choice>
<choice id="ocil:org.namespace:choice:4">Invention of the wheel</choice>
<default_answer>
        <choice_ref>ocil:org.namespace:choice:2</choice_ref>
        <choice_ref>ocil:org.namespace:choice:3</choice_ref>
        <choice_ref>ocil:org.namespace:choice:1</choice_ref>
        <choice_ref>ocil:org.namespace:choice:4</choice_ref>
</default_answer>
</sequence_question>

<sequence_question_test_action id="ocil:org.namespace:testaction:1"
question_ref="ocil:org.namespace:question:1">
<when_sequence>
        <result>PASS</result>
        <sequence>
            <choice_ref>ocil:org.namespace:choice:4</choice_ref>
            <choice_ref>ocil:org.namespace:choice:1</choice_ref>
            <choice_ref>ocil:org.namespace:choice:2</choice_ref>
            <choice_ref>ocil:org.namespace:choice:3</choice_ref>
        </sequence>
</when_sequence>
</sequence_question_test_action>
```

Proposal 9: Matching questions

Matching questions are used when the user is asked to match items from one list with items from another list. To implement this, add a MatchingPairType; add a MatchingQuestionType; and a MatchingQuestionTestActionType.

Shane provided these examples:

```
<matching_question id="ocil:org.namespace:question:1">
<question_text>Match these baby animals to their non-baby counterparts</question_text>
<set_A>
        <choice id="ocil:org.namespace:choice:1">Kitten</choice>
        <choice id="ocil:org.namespace:choice:2">Puppy</choice>
        <choice id="ocil:org.namespace:choice:3">Calf</choice>
</set_A>
<set_B>
        <choice id="ocil:org.namespace:choice:4">Cow</choice>
        <choice id="ocil:org.namespace:choice:5">Dog</choice>
        <choice id="ocil:org.namespace:choice:6">Cat</choice>
</set_B>
</matching_question>
```

```
<matching_question_test_action id="ocil:org.namespace:testaction:1"
question_ref="ocil:org.namespace:question:1">
<when_matches>
        <result>PASS</result>
        <check>all</check>
        <match>
                <item_A>ocil:org.namespace:choice:1</item_A>
                <item_B>ocil:org.namespace:choice:6</item_B>
        </match>
        <match>
                <item_A>ocil:org.namespace:choice:2</item_A>
                <item_B>ocil:org.namespace:choice:5</item_B>
        </match>
        <match>
                <item_A>ocil:org.namespace:choice:3</item_A>
                <item_B>ocil:org.namespace:choice:4</item_B>
        </match>
        </when_matches>
</matching_question_test_action>
```

David Ries asked for an use case applicable to security automation and Mike Kinney answered that this is for more general use cases, and is just an attempt at “thinking outside the box.”

Next Steps

Jim Ronayne stated that the next steps would be to update the specifications and to start formal acceptance process. He encouraged people to watch the discussions lists for news and updates.

Jim asked whether people thought there would be a benefit to having an "adoption program" – an opportunity for vendors to state their intention of adopting the OCIL specification. Kent Landfield said he saw no benefit in an adoption program, because it did not require vendors to implement the spec properly. Kent thought an interoperability event would be much more beneficial. John Banghart agreed with Kent that an interoperability event would be a good thing. John encouraged vendors to step forward to NIST, or elsewhere, to express their interest in such an exercise.

Reinvigorating Remediation

Presenters

Jim Ronayne, NSA

Kent Landfield,McAfee

Session Objectives

Kent Landfield led off this presentation and started with some general observations. The remediation specifications have been slow to develop. The CRE specification has been released in draft form but has had little traction so far.

This session will first briefly discuss a specific proposal for including CRE content in XCCDF benchmarks. A significant number of new benchmarks will be created this year and will include basic remediation content. A method for declaring remediation policy is required to make the content usable.

The rest of the session will focus on the strategy for making remediation standards viable within the security automation community. Fundamental assumptions about future development and use of remediation standards will be questioned and a plan of action for specification, content, and tool development will be determined.

Using CRE

Jim Ronayne then briefly discussed a specific proposal for including CRE content in XCCDF benchmarks. A significant number of new benchmarks will be created this year and will include basic remediation content. A method for declaring remediation policy is required to make the content usable.

The CRE spec is draft but stable; it needs implementation experience. USG plans to create significant amount of SCAP content this year and desires the inclusion of CREs. These CREs will be created in the DoD namespace. There is a requirement to express remediation policy.

CRE is used in XXCDF today and it validates against the 1.1.4 schema. It used the fixtext tags within XCCDF with a human-oriented description of how to apply the fix. It uses XCCDF variables (values) to set CRE parameters. Potentially re-using check variable when possible, otherwise create new variables. Use

fix tag to indicate CRE and parameter. The initial use case is for human readable fix output lists. A Secondary use case is for tools to consume XCCDF results and XCCDF fix policy.

Jim provided an example:

```
<fixtext fixref=" cre_com.example_31-5_fix"> Set Domain Group Policy Object (Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options) using the IGroupPolicyObject interface. Network security: LAN Manager authentication level should be set to <sub idref="lan_manager_authentication_level_var" />.</fixtext>
<fix id="cre_com.example_31-5_fix" system="http://cre.mite.org/cre.xsd">
    cre:com.example:31-5:
lan_manager_authentication_level:<sub idref="lan_manager_authentication_level_var" />
</fix>
```

CRE in XCCDF in the future might look like this:

```
<fixtext fixref="maximum_password_age_fix">Set the maximum password age to <sub idref="password_maximum_age_var_cre" /> days directly in the local SAM database (e.g., via NetUserModalsSet()).</fixtext>
<fix system="http://cre.mitre.org" id="maximum_password_age_fix">
    <fix-export export-name="cre:org.mitre.cre.draft:var:117"
value-id="password_maximum_age_var_cre" />
    <fix-content-ref name="cre:org.mitre.cre.draft:117" />
</fix>
```

There are two outstanding questions that need to be addressed:

- Should fix be defined in the XCCDF spec or separately (for use in XCCDF)?
- Should we start working on changes to XCCDF to support remediation now or wait until we have some experience with the current method?

Steve Piliero stated that he was supportive of the approach described by Jim. Joe Wolfkiel stated that he would be in favor of keeping CRE definitions separate from XCCDF. David Solin said that if CRE was to stay in XCCDF, there would be a lot to do to de-couple compliance from remediation to insert a change control process.

Remediation Topics

Kent then focused the discussion on the strategy for making remediation standards viable within the security automation community. Fundamental assumptions about future development and use of remediation standards were questioned and he wanted to develop a plan of action for specification, content, and tool development.

Derived Requirements for Remediation (Draft NIST IR 7670)

- DR1. Method for uniquely identifying a remediation (CRE)
- DR2. Definition of an exchange format for basic remediation information (exCRE)
- DR3. Definition of additional data about a remediation, including mappings to applicable platforms, related vulnerabilities, or configuration issues (ERI)
- DR4. Definition of a language for the exchange of the additional remediation data identified in DR3
- DR5. Method for specifying remediations for classes of assets
- DR6. Method for applying remediations to specific assets in an enterprise environment
- DR7. Method for reporting the results of an attempted remediation
- DR8. Method for expressing how to perform a remediation in a precise, machine-readable fashion

Common Remediation Enumeration (DR1 - Draft NIST IR 7670)

- Similar to a CVE
- The scope of a CRE entry is the set of actions that must be taken to accomplish a distinct remediation objective (e.g., installing a software patch or changing the system configuration). As such, a single CRE could require that multiple atomic actions, such as changing a configuration value and installing a patch, be performed to achieve the desired end state.
- A CRE entry consists of only the minimum amount of data required to differentiate one remediation from another:
- **Unique Identifier** - textual ID for the specific remediation being referred to. Because there is a need to enumerate organization-specific remediations in addition to those universally recognized, CRE will accommodate local identifiers. For example, an organization may choose to issue local CRE identifiers for internal, custom applications or for remediation actions that are specific to their operational environment. The CRE ID will contain a namespace component that identifies the organization that issued and controls the CRE entry. The remainder of a CRE ID is a non-semantic unique ID; it does not convey or encode any information about the remediation or impart any meaning.

- **Description** - brief paragraph intended for a human audience. The description, in conjunction with the supporting references, must provide sufficient information to allow a person to differentiate one remediation from another. The description is not intended to convey the details of the remediation actions, but only a concise description.
- **Supporting References** - links to authoritative sources where the remediation has been described (e.g., configuration guides, vendor security bulletins, patches). The references may provide additional supporting information about the CRE, including why it was created, how it is distinct from other similar CREs or additional technical discussions regarding the remediation.
- **Metadata** - Information about the CRE entries themselves will also be maintained, such as creation and modification dates, deprecation status, version information, and provenance.

CRE Data Exchange Format (exCRE) (DR2 - Draft NIST IR 7670)

- An exchange format for CRE entries and related metadata is required to enable the transfer of CREs between parties and tools.
- This transport format allows the exchange of either the standard CRE list or organization-specific CREs.
- The CRE data exchange format is envisioned as a lightweight, XML-based schema that serves as the standard import, export, and exchange format for basic remediation information as provided by CRE.
- The CRE data exchange format will be described in a forthcoming specification.

Extended Remediation Information (ERI) (DR3 - Draft NIST IR 7670)

- As CRE is analogous to CVE, so is Extended Remediation Information (ERI) analogous to the additional CVE-related information available in the National Vulnerability Database (NVD).
- Extended Remediation Information defines additional information about CRE entries necessary to fully support enterprise remediation workflows. While a sizeable collection of remediation information exists today, it lacks structural consistency, varies in completeness from vendor to vendor, and often must be retrieved from multiple sources. By specifying desired ERI, providers of remediation information have a template that describes the desired content.
- ERI may describe:
  - Applicable platforms (i.e., CPEs) for the remediation
  - Vulnerabilities (i.e., CVEs) that a remediation is intended to resolve
  - Misconfigurations (i.e., CCEs) that a remediation is intended to resolve
  - Human- or machine-readable prerequisites for remediation (e.g., other remediations)
  - Descriptions of remediation actions (human- or machine-readable)
  - Required actions on success or failure of an attempt to apply the remediation (human- or machine-readable)

- ERI does not prescribe a database format or schema or any other presentation model. It simply identifies the additional data that may be required to support the identified technical use cases, beyond the base CRE entries.
- ERI as described provides the information necessary to decide which remediations to include in an enterprise remediation policy, or to facilitate the selection of appropriate remediations to apply based on assessment results.
- The ability to fully support the breadth of identified use cases, enabling maximum automation and tool integration, requires that ERI for all critical remediations be managed and maintained by some centralized authority or authorities.
- ERI will be fully described in a forthcoming specification.

Extended Remediation Information Data Exchange Format (exERI) (DR4 - Draft NIST IR 7670

- A common representation of ERI is required to facilitate data exchange and to foster tool interoperability. The Extended Remediation Information data exchange format is proposed as a means of enabling efficient interchange of ERI data.
- While ERI defines the remediation data necessary to support the described use cases, the data exchange format specifies a standardized format for the automated exchange of ERI between remediation information sources and remediation tools. ERI may also appear in machine-readable remediation policy documents.
- The ERI data exchange format is envisioned as an XML-based schema that extends the CRE schema, allowing ERI documents to refer to the CRE entries they extend by CRE ID alone, or to contain the full contents of the CRE entry.
- The ERI data exchange format will be fully described in a forthcoming specification document.

Remediation Policy Specification (RP) (DR5 - Draft NIST IR 7670)

- The Remediation Policy Specification defines how to associate particular remediations with various classes or types of IT assets. Such a capability allows organizations to specify allowed, preferred, or required remediations for specified collections of IT assets.
- Those asset types may be defined by:
  - Platform type (e.g., desktop, notebook, server)
  - Software inventory (i.e., presence of a particular product)
  - Presence of specific vulnerabilities
  - Current configuration of the IT asset
  - Functional categories (e.g., web server, database server)
  - Organizational boundaries
  - Combinations of the above
- The Remediation Policy Specification provides a standard format that enables an organization to constrain the full set of *possible* remediation options for a given circumstance to a smaller

allowed subset. For example, suppose there are two known CRE entries for a particular vulnerability, one identifying a patch and the other a mitigating workaround. An organization's remediation policy might indicate that in most cases, the patch should be installed, but in cases where a third-party application with known conflicts with the patch is also present, the workaround should be applied instead.

- A remediation policy in effect conveys remediation decisions that have been made in advance, simplifying the decisions that must be made synchronously in a remediation workflow. In cases where the remediation policy specifies a single remediation for a given situation, full automation of remediation action may be possible. The Remediation Policy Specification defines how remediation policies may be expressed and exchanged in an open, unambiguous, and machine-readable format.
- Initial discussion of the requirements for the Remediation Policy Specification suggests XCCDF could potentially be used for this purpose, either in its current form or with some modifications. The use of XCCDF as potentially be used for this purpose, either in its current form or with some modifications. The use of XCCDF as this expression will be investigated, as will other viable alternatives.
- The Remediation Policy Specification will be fully described in a forthcoming specification document.

Remediation Tasking Language (RTL) (DR6 - Draft NIST IR 7670)

- In contrast to the Remediation Policy Specification, which assigns remediations to classes of assets, the proposed Remediation Tasking Language (RTL) provides a standardized format to direct compliant tools to enact specific remediations on specific assets. RTL documents represent the output of the remediation decision process, and function as a standardized input format for remediation tools.
- Remediation Tasking Language documents specify:
- Which assets to remediate
- Which remediation actions to perform
- What values are to be used in performing each remediation (e.g., number of characters to set as 335 the minimum password length)
- Other operational parameters, such as deferral options, may also be included.
- Development of the Remediation Tasking Language will take into consideration other emerging reporting and control specifications being considered in the overall security automation architecture. This evaluation will include assessing conceptual alignment and the potential for schema reuse.
- The Remediation Tasking Language will be fully described in a forthcoming specification document.

Remediation Results (RR) (DR7 - Draft NIST IR 7670)

- In order to determine what follow-up steps, if any, are necessary, the results of a remediation attempt must be communicated back to the tool or process that requested the remediation. These Remediation Results convey the outcome (e.g., success/failure/error) of attempted remediation actions as reported by the remediation tool. Remediation Results also enable roll-up reporting and provide enhanced situational awareness.
- These results include, by asset:
  - Outcome of the attempted remediation
  - Explanatory information, when the remediation attempt was unsuccessful
  - Date and time the remediation was performed
  - Date and time the remediation is scheduled to be performed, if deferred
  - Initiator of the deferral action
- Remediation Results are not intended to serve as an authoritative assertion of whether an asset is still subject to a vulnerability or misconfiguration that a remediation was intended to address. Initiating a reassessment of the affected asset using the appropriate assessment tool is the preferred method for making such a determination. Remediation Results are most ideally suited for supporting follow-on decisions in the remediation workflow, such as whether to attempt a failed remediation again, whether to override the deferral of a remediation by a user, or as decision support material in determining the need for further assessment.
- Development of the Remediation Results will take into consideration other emerging reporting formats being considered in the overall security automation architecture. This evaluation will include assessing conceptual alignment and the potential for schema reuse.
- Remediation Results will be fully described in a forthcoming specification document.
- OVRL will be fully described in a forthcoming specification document.

Kent summarized by stating that there are currently two NIST documents which address the subject of remediation:

- CRE format and usage described in NIST IR 7831.
- The CRE data exchange format is described Appendix B in NIST IR 7831.

And there are additional future plans to address this topic:

- ERI will be fully described in a forthcoming specification.
- The ERI data exchange format will be fully described in a forthcoming specification document.
- The Remediation Policy Specification will be fully described in a forthcoming specification document.
- The Remediation Tasking Language will be fully described in a forthcoming specification document.

- Remediation Results will be fully described in a forthcoming specification document.
- OVRL will be fully described in a forthcoming specification document.

General Discussion

Kent then proposed several questions to drive the conversation.

Which use cases do we want remediation to support? Which is most important? Which should we focus on first (i.e., which is most achievable in a short time frame)?

- Only one person wanted to just address configuration changes.
- Mike Kinney stated a fundamental concept of software distribution is that patches do not get pushed, they get pulled by the user.
- Joe Wolfkiel agreed with Mike and then added that his office does not auto-deploy GPOs to Active Directory either.
- Another attendee countered that rather than discussing the types of remediations to be performed, the discussion should center on where remediation sits. Will there really be a policy that says there's remediation content that the government will require to be executed.
- Kent says we're already there. There is policy being developed that will require remediation.
- Mike Kinney stated that the first question to be answered is whether we want to pursue OVRL?
- Jim Ronayne boiled down this issue a bit – do we want to agree to develop languages that fix things on the fly, or do we want to restrict remediation content to be developed by the vendor?
- JoeWolfkiel described an approach his office is taking to remediation: a list of patches and system upgrades are made available and a user can hit a button to automate these to a bunch of systems. He felt this approach was very workable.

Kent queried vendors in the audience if they would feel comfortable in getting instructions on Patch Tuesday from outside their organization to distribute in their packages to their customers. The overwhelming consensus was that they would not be comfortable, and Kent agreed with that assessment, saying that he is very dependent on his QA organization to test fixes before they are distributed. Jim Ronayne added that the deliberate decision had already been made in the past not to pursue OVRL.

An attendee made the point that the act of changing configuration settings is vastly different from pushing patches.

Aharon Chernin spoke as a consumer and said he would like to own his OVRL content. He'd like to buy a detection feed from one vendor and an OVRL feed from another vendor. He wondered why he could own detection content, but could not own remediation content. Kent suggested that he could own remediation content by contracting with any of many vendors who offer such a solution. Aharon countered that that would require him to be locked into that product. Kent expressed concern on the part on vendors because an OVRL solution would require them to a) re-architect their products and b) not be able to test fixes when they came out. Kent cited an example of a problem with FDCC and he had difficulty proving to his customer, a government agency, that the problem was caused by NIST and not

his company. Aharon again countered saying that Kent's point limited the use case of his product to the ones he thought of, assuming that the product was just doing vulnerability management and compliance, but a customer may be using the remediation for threat mitigation. Dave Waltermire agreed with Aharon.

Mike Kinney said his office has been looking at this problem for a while and they've taken the identifier approach, as opposed to the OVRL approach, because he feels the customer shouldn't specificy the approach that vendors take to remediate. He feels the OVRL approach will lead to that. Dave Waltermire disagreed and stated that the identity approach would require customers to enumerate every possibility that needs to be done and he wondered if that would lead to a scalability problem. Dave also stated that comparing products is likely to be important to organizations. Mike countered by suggesting that the OVRL approach will take a very long time to implement and that the identity approach will allow the community to make progress right away. Joe Wolfkiel expressed agreement with Mike.

Jim Ronayne suggested that an earlier decision to pursue CRE and delay OCRL is how we got into this current situation – i.e. no progress has been made on OVRL because no one's been working on it. Dave Waltermire suggested that one of the reasons no progress has been made is that vendors are not participating, but rather waiting for government. Dave suggested that vendors need to show more initiative in OVRL.

Steve Piliero stated his organization is pursuing the CRE approach for the reasons identified by Mike.

Kent then summarized the discussion by saying there is disagreement on whether to pursue OVRL; it makes sense to take slow and measured steps toward the identity solution, but there should also be an eye to the horizon. Kent suggested that the community should work toward the identity solution and gauge how much progress can be made before investing in the OVRL solution.

Dave Waltermire stated that the CRE spec has been released and wondered what else needs to be done. Mike Kinney responded that the ERI spec also needs to be completed. Jack Vander Pol said his group did some prototyping with CRE, and wondered about the problem of synchronizing remediation IDs with policy. Dave Waltermire suggested that adding parameters to CREs could help this problem. But Jack countered that his experience is that on simple items that works, but on more complicated things like file permissions, auditing, and user rights are tougher to implement and then you end up with OVRL. Dave said they're trying to strike a happy medium. Jim Ronayne chimed in that while you want to add structure so that it's predictable, but at what point of adding structure have you effectively implemented OVRL? Jim, commenting on Jack's work, said adding structure seemed like it was a slippery slope. Jim went on to say that he thought they ended up at a reasonable compromise between the two approaches and he thinks the community should pursue this course, but more collaboration is needed.

Jim Ronayne described that there's an ambiguous relationship between individual CCEs and individual CREs. Mike Kinney said that he wanted an ID for everything they wanted to fix, but that CCE doesn't

provide that. Harold Booth asked Mike if it would helpful if CCE could be used to identify things that needed to be fixed. Mike thought that it would help. Jim Ronayne asked whether that wouldn't effectively be making CCE and OVAL redundant. Harold concurred. Joe Wolfkiel chimed in by saying he thought it would be a mistake to try to make CRE fix CCE. Jon Baker pointed out that it's possible for multiple CREs to point to one remediation.

Gary Gapinski stated that the current CRE spec does not clearly the specify the method to be used to remediate systems. Jim Ronayne suggested that one of the difficulties in producing CREs is that they're supposed to be precise in their descriptions, but in reality one can only be so precise without using an OVRL approach. So different people will use different descriptions for CRE and in the end we'll be dependent upon the way vendors interpret these descriptions, with the goal of different methods yielding the same result. Gary added that he detects some discomfort within the community on this situation. Jim added that OVRL is the best way to force the intended results, but that it's a lot more work to develop such a spec, and that is our dilemma.

Mike Kinney stated he would like to take a list of the approved DoD CREs and map them to the fixes in vendor's remediation database, plug it into my network, run a STIG, and have it fix anything that wasn't compliant. Mike said he wants that functionality today, but believes we can get there faster doing that mapping than any other way we're considering. Dave Waltermire warned that you have to make sure the mappings are correct because pretty chaotic things can happen if they're not. Mike countered that there are always other optional approaches. Dave said that any system that is dependent on human input can be flawed because humans make mistakes.

Kent wanted to turn the discussion toward ownership and responsibility of CREs. He asked whether anyone thought there would be a central repository of CREs, and nobody thought there would be. Mike said he foresaw vendors adding a column in their database of fixes to list their own CREs, and the government informing each vendor what they wanted fixed and what that means. The vendors would understand how to map from their database of fixes into the things the government wants fixed. Kent described that there's a timing issues that makes him dislike this approach, e.g. on Patch Tuesday, his organization would be under a time crunch to get all of the fixes out for a large amount of organizations. Mike said he only wants this functionality for configuration settings, but Kent insisted that it's a very big problem. Kent summarized by saying that this is an outstanding issue to be resolved and his organization needs to look into mapping CREs into his database, in a way that scales across all of the products they support.

Kent asked, where is the boundary between existing remediation tools and the standards? Kent then answered his own question by saying that CRE might be that initial boundary.

Kent suggested that, based upon the discussion, it seems like the only short term approach that makes sense is to pursue CRE. Mike Kinney proposed that the community also continue down a parallel path of OVRL, as work continued on CRE, i.e. finding out how much progress could be made with CRE. Jim Ronayne suggested that CRE be declared "final" and then build some content around it because the thing the community needs most is some experience using it, and we're going to need some vendor

participation using it. Another audience member opined that he thought OVRL could be valuable, but was skeptical about CRE because he thought there would be difficulty in mapping them. Rather than such a "kluge" he preferred a free-form remediation text in XCCDF, making it optional and not automated. This idea was met with criticism.

Steve Hanna expressed concern that CRE could represent a huge barrier to entry. Kent tried to re-focus the discussion by saying the goal of this session was to agree on a short term approach that can make progress toward the longer term goal. With that in mind, a short term approach which includes some proprietary solutions is acceptable, as long as it can lead to longer term progress. Dave Waltermire cautioned that sometimes short term approaches end up as the final solution.

Kent pointed out that yet to be discussed were exchange formats, remediation policy, and remediation tasking language. The relative merits of these topics were discussed, with no general consensus about any of them.

Mike Kinney asked for a show of hands of who believed that OVRL was the proper long term goal – but Gary Gapinski asked to change the question from OVRL to "a uniform means by which to express remediation." With a good showing of hands, Mike offered that he is no longer in a position to fund OVRL and asked for someone else to step forward and lead the effort for the development of a proof-of-concept. Dave Waltermire also made an appeal for volunteers to work on this effort. Mike Kinney, Kent Landfield, Gary Gapinski, and another person volunteered to work on this effort on their own time.

Kent suggested using the existing remediation mailing list to carry on discussions on this topic - /remediation/dev/subscribe@nist.gov

Friday July 13th

TAXII / STIX

Presenters
Rich Struse,DHSA
Sean Barnum, The MITRE Corporation

Note: Only discussion points are captured here. Briefing content is not captured.

Q: Are you aware of any people who have gone from XML schema to a semantic representation?

A: Yes, domains that are well understood do this. When you don't know what you don't know, it's harder to do that.

Q: Is it essential for the three stages of maturation/standardization to go in the order on the slides?

A: No, we just felt that it made the most sense for this effort.

A: We picked standards as the end of the roadmap because standards tend to move slower.

So there are some standards bodies where you have to come in with something more formal, but within bodies like the IETF you can iterate quickly. I just want to make sure that everything is understood.

Q: Where does something like IODEF fit into the STIX architecture?

A: One of the things that we're going to do is provide a tool that goes from STIX to IODEF and IODEF to STIX to provide interoperability where appropriate. STIX embeds IODEF currently.

Q: Where do we find the schemas and some of this content?

A: We are going to build a website & github site and post it there.

Q: Can you make sure that announcement gets posted to this community?

A: Yes.

Q: How do object schemas get added?

A: The core schema is decoupled from CybOX objects. We are still working on the community process.

Q: Is INDeX used as a common language for sensors?

A: INDeX wouldn't be the language for sensors. CybOX could be.

Q: How would you indicate how accurate the sensor is?

A: What we've built in so far is a fairly rich representation of that kind of information.

Q: Is there any way to express how current the data is?

A: At the CybOX level, you capture exactly when it happened. For temporal patterns, that is captured in INDeX.

Q: Is anybody already at the edge of information overload?

A: Part of the valid time window is to help with that.

Q: Your use cases indicate that CybOX may be used as a bridge for converting tool data for different sensors. Is there any study that has been done to show that there is not any data lost in conversion?

A: No. We cannot cover all tools, but we capture the core of what those tools deal with.

High-Level CybOX

Presenter

Sean Barnum, The MITRE Corporation

Don Campbell (McAfee): How do object schemas get added?

Sean Barnum (MITRE): The way the core schema actually works is that it has a Defined Object element, which is an abstract type. All of the objects have separate schemas and namespaces, which are extensions of that abstract type. As such, the core schema is decoupled from any and all objects, which means that one can use any objects that they want, and also allows us to create objects independently and without needing to modify the core schema. This also permits users to create any new object as an extension of that abstract type, which can then be used in their content. Of course, this means that all consumers would need that object in order to understand the content. We're working out exactly the process for how new objects will be added into CybOX through community input.

Don Campbell: Thank you - I was mostly interested in how objects get added through the community.

Sean Barnum: We're still working on that and getting feedback. We've yet to have a completely independent party submit a new object. Right now it's more of a discussion of what objects we need and the fields that they should have, but we recognize that we'll hit these sorts of issues going forward and are working out the process for how this will be accomplished.

Dave Waltermire (NIST): What do you mean by excavating incident context, and how does this support it?

Sean Barnum: Basically, in an incident response investigation, something gets fired off and the investigators will go look and ask "what actually happened?" As they're going through to see what occurred, what the systems look like, what the traffic might have looked like within logs, the incident response team is trying to get a handle on what happened and what it might mean. Really what they're doing throughout the course of that investigation is observing what the system looked like, the logs that are relevant to this picture, etc. There's the potential to use CybOX to capture that sort of context,

whereas today it's typically an incident response analyst keeping a set of notes. Perhaps at the end they capture these notes in an email or pdf, but this gives them the ability to potentially to capture these details in a structured fashion as they're going through the investigation, so that at the end you now have a much better bundled set of information regarding the context of what happened than you did before. This is especially useful in cases where you want to generate new indicators out of that data, since you have the structure right there at hand. To use the example that Rich Struse used with the tool developed by the DTCC, you now have that structure to say "now I want to generate an indicator for not just what came in, but also what we discovered during the investigation."

Mike Kinney (NSA): Can IndEX be used as a common language for sensors?

Sean Barnum: IndEX wouldn't be the language for sensors. IndEX is a layer of context around an observable that characterizes it as an indicator, e.g. what does it mean, what you can do with it, etc. The CybOX underlying it can be used as a common language for sensors, since the sensors describe what they saw, or the facts of something that was seen.

Mike Kinney: I need a language that's common to all of our different kinds of sensors so that I can build an ID once instead of 27 different ones. So CybOX can do that?

Sean Barnum: CybOX is probably the closest thing to that. Now that doesn't mean that each of the sensors has to natively represent it, but it would permit the interoperability between them. So some of them may be purpose built and you could use this, other ones are going to use proprietary standards, and still others may be commercial solutions. CybOX can be the language in which you pull from all of those different things and do correlation and aggregation across them.

Dave Waltermire: How do you provide context around the observations? How would you indicate whether the sensor had the correct picture of what it's looking at or whether it was obscured in some way, how do you get that kind of context? How do you determine how accurate the sensor is? That affects how you interpret the observables, and if that gets disconnected from the observables you lose a lot of information, essentially, about what they mean.

Sean Barnum: So what we've built in so far is that in any of these observables there's a fairly rich description of the information source, where it came from, when it came, tooling, the context of the environment that it was run in, etc. This tries to enable some of that determination, but again this is a language rather than an action. CybOX, IndEX, and STIX are all languages in the vein of OVAL and MAEC, but they're not the tooling. When I say that we provide utilities, they're not in terms of analysis or heuristics or any of that stuff. Rather, we want to try and provide richly expressed enough information so that you can do that sort of analysis; accordingly, we've built in such information in some places which we've recognized and understood. But again, we want all of this to be used and prototyped in the real world, as this will help us find things that "we didn't know we didn't know", to use a Donald Rumsfeld quote.

Unknown Speaker: Is currency, or the lack thereof of information explicit, implicit, or both within STIX? So how useful, how timely, how pertinent, or the negation of any of those?

Sean Barnum: So at the CybOX level you capture where it went; if it's an instance of an observable, you can capture exactly when it was done. If it's a pattern you're not necessarily describing when it's going to be valid because you're not trying to attach meaning. At the IndEX level for indicators, there is the ability to define valid time windows, so you can actually say that "this set of observables is relevant within this context during this period of time and then it will expire and be less relevant." There's also the ability inside IndEX and STIX to capture information regarding confidence of something, and it can be attached to time windows as well. So there is a temporal dimension to some of these things, but it's not at the CybOX level for patterns, only for instances. When you start to abstract to patterns, that temporal context is captured within IndEX for indicators.

Unknown Speaker: Now that people have a tidy way to represent this information, voluminous information can be obtained and stashed away without needing to purge it. I was wondering if anyone is at the edge of where they have more information than they can profitably or efficiently exploit?

Sean Barnum: There are absolutely such organizations; there are some who want as much data as they can get, but there are others as well. So part of that valid time window is the ability to actually manage the lifecycle for some of these things. Again, these languages are not about the process of managing this information, but rather they try to give you the information to help you manage that process, such that you could say "this thing has moved outside of its valid time window and we're going to change its status to something like deprecated." Now, the actually persistent store of this data could go through any transition at that level, e.g. does it get archived, does it get thrown away, etc. Especially with this kind of information within a public private space, you're dealing with privacy concerns as well. We're trying to provide the right level of information to help track these things and enable the processes to do so, but we're not about the processes themselves.

Unknown Speaker: Your use cases indicate that CybOX may be used as a bridge for converting between different sensor outputs. Has any study been done to prove that no characteristics are lost during such a conversion?

Sean Barnum: No, but this is a targeted use case. Depending on which tools you're going to and from, and the expressivity of what they capture, there's no guarantee. We can't possibly cover all tools, but what we're trying to do is to provide an information structure that captures the core of what some of those tools deal with. It's very possible that certain proprietary tools may have a little extra flavor that we don't have, so if you went from those to someplace else, you're going to lose that. Over time, if we look at where those gaps exist from transforming those things, and from communities coming to us and saying "oh by the way, your model is missing this", if the broader community feels that this is valuable then we would add it. But there's always the possibility for losing some of the granularity in those transitions. We're trying to aim at the core so you can do transitions between some of the obvious ones like Snort, OVAL, etc. Today we make no guarantees, and in fact know that it will not cover everything, and make no guarantees in the future that we will have full coverage of everything, as that's an impossible task to undertake to some degree.

High Level MAEC

Presenter
Penny Chase, The MITRE Corporation

When discussing TTPs, malware can provide insight into additional TTPs or attribution. MAEC represents the malware focus of Structured Threat Information eXpression (STIX). This talk gives the overview of the MAEC language.

Malware Attribute Enumeration and Characterization (MAEC)

MAEC is a language with three main components that can share structured threat information about malware. The first component is the grammar, which is the schema that defines the language. Secondly, there is the vocabulary that has been created for the common terms. Lastly, there is all the relevant information about a piece of malware that gets stored in a bundled format. MAEC focuses more on attributes and behaviors rather than signatures. In this sense you can come up with families of malware. It is often the case that these attributes and behaviors become indicators for the malware.

Some operational use cases for MAEC include being able to model information for threat analysis, intrusion detection, and incident management. Another large use case is the analysis of malware, where the use of MAEC can provide a standardized tool output that could allow for information sharing between repositories. It can also help guide the analysis process.

MAEC structure overview

The core has been developed at three levels of abstraction. At the lowest level represented you will find actions. These actions may include data at the implementation level such as specific API calls used, but in general will span multiple architectures. The next level up represents the behaviors found that explain the actions. A created file might represent the behavior for a malicious binary instantiation. At the highest level you will find mechanisms, which is similar to behaviors, but attempts to note the intent of such abstractions.

In MAEC, the language aims to characterize malware through the actions and behaviors. MAEC makes use of the CybOX specification at its core, through actions and objects imported from CybOX. In this capacity, the observations are made once then can be characterized later.

MAEC's bundle.

The MAEC bundle is composed of several sections. Some of these contain metadata about the bundle itself for validation purposes and include the globally unique ID and schema version. The relevant information about malware is also collected in the bundle under the analyses, behaviors, actions, objects, indicators, or pools components. The analyses component contains information about how the analysis was performed, in what environment this information was obtained, and with what tools. Indicators are the attributes that a tool may identify as being a feature of the malware. These are usually candidates for inclusion as an indicator of the malware. All of these additional components are optional, and may be added in at a later time. The different components may be attained through various forms of malware analysis.

Analysis process

The malware binary undergoes a pre-screen/static triage possibly involving running it through various antivirus engines to extract information about the structure. Possible information obtained from this could be strings, hashes, or function import/exports. The MAEC bundle is created with this new object information. In some cases this is enough for representing interesting malware. If additional information is desired, then one could perform some dynamic analysis, where the malware is executed on a sandbox and have actions/additional objects reported as well as their relationship. Furthermore, an analyst could perform manual analysis to figure out additional actions and behaviors, all of which get added to the bundle. All of these can be done incrementally with a valid bundle being represented at any step.

MAEC development

MAEC had looked into leveraging different resources to prevent having to reinvent the interesting collection data through their requirements and what they wished to express. The team also worked with Mandiant's openIOC for having their objects represented in MAEC and CybOX. Other related efforts include the IEEE ISCG Malware Metadata Exchange Format working group wishing to incorporate behavioral metadata, and which will likely import MAEC/CybOX. The MILE working group is also looking into MAEC for external sources of information.

A question was raised about the IEEE effort and whether it would be incorporated into an IEEE standard. The answer is currently unknown. The first version of the schema was not a standard, and the group will have to wait until the changes are made to the schema to decide whether it should become one.

MAEC roadmap

Version 1.1 of MAEC was released in January 2011 and was focused on dynamic and static analysis elements. The following year for version 2.0, CybOX 0.7 was incorporated, as well as the ability to add relationships and indicators to the bundle. The current version of MAEC is 2.1, which updated the CybOX objects to the 1.0 schema.

Moving forward, some desired goals for MAEC version 2.x include the ability to apply views to a MAEC bundle. Bundle management is another goal, which would allow the ability to merge bundles. If two tools were to create a bundle for the same malware, then it would be desirable to be able to compare the common and unique portions. Bundle views relate to the trimming of information within a bundle to only report back data relevant to a particular role. For example, an incident responder might only be interested in indicators of the malware.

An extended goal intended for MAEC version 3.x includes the creation of taxonomies for behaviors and mechanisms. Better bundle abstraction is desired such that related bundles could be considered a family of malware. Incorporation of mobile malware into MAEC/CybOX is a more recent goal now that it has been on the rise.

The goal to create common vocabulary for the different indicators of malware has had some work done on it already. The team will be closely looking at this to make sure they align with the previous efforts. It is expected to expand the behaviors and classifications for different classes of malware to be more accurate.

Discussion

The only question asked was how CybOX fits into the schema for MAEC. MAEC imports CybOX for objects and action level implementations.

MAEC Utilities

Presenter

Ivan Kirillov, The MITRE Corporation

This presentation will delve into the available utilities for assisting with MAEC content. This will cover the existing tools as well as some that are currently in development.

MAEC tools overview

Within MAEC there are three general classes of tools. The first are those that help with content creation and manipulation. With MAEC being rooted in analysis of dynamic analysis, some transformations have been created to take the output from several tools and create MAEC content. Next, some tools exist that allow for analysis of the MAEC content to delve into the components rather than trying to parse the files directly. Lastly, there are some utilities for viewing MAEC content. These are a simple MAEC to HTML transformation. This is more human readable than direct XML inspection.

Many vendors helped out with the creation of the translators by providing schema and support for their respective outputs. Some of these vendors are in talks to have MAEC become natively supported. In addition, some open source projects took the initiative to implement MAEC. Someone else wrote a tool to compare the common output of MAEC bundles while researching how different dynamic analysis tools represented objects.

From the tools available, some support MAEC version 2.1 while others support MAEC version 1.1. The tools are either written in Python or XSL.

MAEC schema bindings

The MAEC bindings take the MAEC and CybOX schemas and create a class for each datatype. This allows for easier creation and editing of MAEC content. The bindings support the full 1.0 CybOX schema, and were generated from a tool called GenerateDS with some fine tuning. Several examples of created content and its related output were given to show how one may use these bindings. It is still a requirement of the content author to understand the schema in order to be able to accurately use the bindings.

MAEC translators

Many translators have been created for the results of multiple dynamic analysis tools. Currently only Anubis is available at present for conversions using the MAEC 2.1 schema, while the other translators only support version 1.1. Some of these tools include Threat Expert, GFI Sandbox, FireEye, PE File, and Norman Sandbox. Furthermore there is a MAEC to OVAL, which is in essence a CybOX to OVAL

transformation. The CybOX objects have been mapped to OVAL objects and can then be used to take a MAEC bundle and create the system checks for that malware. It was noted that due to OVAL being incomplete in terms of collecting malware objects representable by CybOX, that there has been an OVAL for artifact hunting use case effort that recently spawned. Finally, the MAEC to HTML XSL transform will eventually take into account considerations such as bundle views, since a MAEC bundle can be rather large in size.

Making MAEC operational

The operational use case for MAEC involves performing the malware analysis and creating the MAEC bundle. In one example, there is a beaconing behavior being observed and a registry key being created on a system. A Snort signature could be generated and deployed to your IDS, and furthermore one could use the available Snort to CybOX translator. Conversely, if you have experienced that malware before one could use the soon-to-be-released CybOX to Snort translator. A CybOX to OVAL translator would give you the host-based check. All this could be created in a single automated fashion. In the future, it would be advantageous to have all the analysis tools output MAEC natively instead of having to rely on the translators.

Future MAEC tools

Since the current editing of MAEC content, even with the bindings, requires knowledge of the schema, it is a future goal to have a set of APIs to handle the creation of objects without explicitly knowing all the elements you may require. In working towards this goal, it was noticed that there were several common components to generating MAEC content. This resulted in the “MAEC Helper” tool which adds a layer between the bindings and the schema to handle the commonalities by providing a file name and file path. This tool is likely to be the basis for future API work.

Another desired tool is the ability to easily manage MAEC bundles. This would be able to split a MAEC bundle depending on parameters, or merge multiple bundles into a single file. Similarly, being able to construct MAEC views through a tool that lets you indicate what aspects you may be interested.

MAEC community

A Github repository has been established be the public location for MAEC scripts and tools. It contains the python bindings, the translators, and the MAEC comparator tool. These tools are in the process of being updated to MAEC 2.1 schema.

Lastly, a MITRE “Handshake” group has been created for collaboration on schema development and networking with others. This stemmed from the need to remove possible malware binaries from the public mailing lists. The group is invite-only so it would be required to send the MAEC team an email.

Discussions

One question was asked about the results of looking for false positives and false negatives with the Snort signatures, but the team has not performed that type of in depth testing yet.

OVAL Artifact Hunting

Presenter
Charles Schmidt, The MITRE Corporation

Quick Summary

Major points for this talk include:

- The use of OVAL for Artifact Hunting aligns with OVAL's capabilities and purpose. Thus it makes sense to enhance OVAL to support this use case rather than create a new language with redundant capabilities.
- Any incorporation of new artifact-hunting features into the official OVAL language will be undertaken only after deliberation and consensus within the existing OVAL community. It was noted that creating features focused on this use case could impact the perception of how well existing tools support the OVAL language as a whole. As such, new ways of handling this integration may need to be developed.

Minutes

The talk began by addressing why OVAL was being applied to the "new" use case of hunting for malware artifacts. There is a community-expressed need for the ability to hunt for artifacts using an open, actionable, objective, multi-platform, and vendor neutral mechanism. It was noted that hunting for these artifacts involves the same types of assertions about system state that are needed for configuration testing and vulnerability analysis and, as such, are the same class of activity OVAL already supports. In fact, artifact hunting has been listed among the uses of OVAL for a number of years, although little actual work has been done towards this end. The proposed work is to develop new and extended tests to better support the capabilities required by this use case.

It was emphasized that plans to support artifact hunting are being scoped with some specific limits. Most importantly, there is an absolute need to ensure that support for the artifact hunting use case does not disrupt support for the other OVAL use cases, such as policy compliance and vulnerability scanning. Also, there is no plan to use OVAL for network scanning, since OVAL focuses on host-based scanning. It was also noted that it there is an expectation that any content developed in support of artifact hunting would require a different management lifecycle than that currently used for vulnerability and configuration definitions and, as such, there is no assumption that artifact hunting definitions would be consolidated in the OVAL Repository at this time. Finally, there is no plan or expectation that OVAL for artifact hunting will replace signature-based scanning technologies.

It was noted that the ability to generate artifact hunting definitions automatically based on information expressed using malware standards is an important objective of this work. This would allow a malware analysis to become immediately actionable in an OVAL-supporting enterprise. It was noted that manual creation is also to be supported and, as such, there is no attempt to bind OVAL to any particular malware standard.

There was a quick overview of the history of this work, starting with the presentation last year by CyberESI. It was noted that there have been some tweaks to the OVAL language driven by this work, but that there has not been much development of these capabilities in the last 12 months. It was noted that there has been an analysis of the gaps between existing malware standards, such as CybOX for MAEC and OpenIOC. This is available on the github site: https://github.com/OVALProject/Sandbox/blob/master/resources/x-win-artifact-hunting/CybOX_OpenIOC_to_OVAL.xlsx?raw=true. It was emphasized that we have been looking to these standards for direction primarily because they represent a compilation of community experience with regard to malware artifacts rather than a desire to exclude other sources. Any input as to useful targets for this work are welcome.

The plan moving forward is to develop high-priority capabilities based on the needs of the malware community. Development would include schemas, OVAL DI modules and, if the capability has an analog in some malware standards, scripts to auto-generate the requisite OVAL from those standards. Once this is completed, the OVAL community will decide what to do with the new capabilities - whether these should go into the official OVAL language or if they should remain separate.

Kent Landfield noted that he had previously voiced concern about pursuing this capability, primarily because some of the desired tests involved undocumented capabilities in software that vendors would be reluctant to rely upon, and also because it is moving OVAL into a very different usage model. He noted that, while officially one can say they support OVAL without implementing all OVAL capabilities, failure to support everything can lead to a negative impression of a particular product. On that note, he observed that a vendor's artifact hunting capabilities might be in a completely different product from their compliance or vulnerability scanning capabilities. He emphasized that if these new capabilities were added, that there needed to be boxes built around them so expectations could be managed. He also emphasized that he was not opposed to this work, just that he had significant concerns that it could be disruptive to customers' understanding of OVAL capabilities. Charles responded that the avoidance of such disruption is a top priority. He noted that the OVAL community will have the final say as to what is included in the language, and that keeping the work that is only of benefit to the artifact hunting community in a separate box is one of the possible options. It was suggested that, towards that end, OVAL "profiles" could be established and tools could identify the profiles that they support. This might mitigate the confusion if those profiles came with a clear statement of what it meant to be compliant with them. Jon Baker also noted that artifact hunting was one of the driving forces behind the github sandbox - to provide a place for development and experimentation outside the main OVAL development environment.

The conclusion was that this work deserves to go forward, but that new capabilities will only be made "official" after community review and approval, where the nature of the "official" incorporation might differ from current practices of extending the OVAL language. Further discussion on how to incorporate new capabilities will be had when there are some sample capabilities that can guide that discussion.