Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 23, 2013

D. Waltermire, Ed.
NIST
A. Montville
TW

January 19, 2013

Analysis of Security Automation and Continuous Monitoring (SACM) Use Cases
draft-waltermire-sacm-use-cases-03

Abstract

This document identifies foundational use cases, derived functional capabilities, and requirements needed to provide a foundation for creating interoperable automation tools and continuous monitoring solutions that provide visibility into the state of assets, user activities, and network behavior. Stakeholders will be able to use these tools to aggregate and analyze relevant security and operational data to understand the organizations security posture, quantify business risk, and make informed decisions that support organizational objectives while protecting critical information. Organizations will be able to use these tools to augment and automate information sharing activities to collaborate with partners to identify and mitigate threats. Other automation tools will be able to integrate with these capabilities to enforce policies based on human decisions to harden systems, prevent misuse and reduce the overall attack surface.

Table of Contents

1. Introduction

This document addresses foundational use cases in security automation. These use cases may be considered when establishing a charter for the Security Automation and Continuous Monitoring (SACM) working group within the IETF. This working group will address a many of the standards needed to define an interoperable, automation infrastructure required to support timely, accurate and actionable situational awareness over an organization's IT systems. This document enumerates use cases and breaks down related concepts that cross many IT security information domains.

Sections Section 2, Section 3, and Section 4 of this document respectively focus on:

- Defining the key concepts and terminology used within the document providing a common frame of reference;
- Identifying foundational use cases that represent classes of stakeholders, goals, and usage scenarios;
- A set of derived functional capabilities and associated requirements that are needed to support the use cases;

The concepts identified in this document provide a foundation for creating interoperable automation tools and continuous monitoring solutions that provide visibility into the state of assets, user activities, and network behavior. Stakeholders will be able to use these tools to aggregate and analyze relevant security and operational data to understand the organizations security posture, quantify business risk, and make informed decisions that support organizational objectives while protecting critical information. Organizations will be able to use these tools to augment and automate information sharing activities to collaborate with partners to identify and mitigate threats. Other automation tools will be able to integrate with these capabilities to enforce policies based on human decisions to harden systems, prevent misuse and reduce the overall attack surface.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Key Concepts

The operational methods we use within the bounds of our present realities are failing us - we are falling behind. We have begun to recognize that the evolution of threat agents, increasing system complexity, rapid situational security change, and scarce resources are detrimental to our success. There have been efforts to remedy our circumstance, and these efforts are generally known as "Security Automation."

Security Automation is a general term used to reference standards and specifications originally created by the National Institute of Standards and Technology (NIST) and/or the MITRE Corporation. Security Automation generally includes languages, protocols (prescribed ways by which specification collections are used), enumerations, and metrics.

These specifications have provided an opportunity for tool vendors and enterprises building customized solutions to take the appropriate steps toward enabling Security Automation by defining common information expressions. In effect, common expression of information enables interoperability between tools (whether customized, commercial, or freely available). Another important capability common expression provides is the ability to automate portions of security processes to gain efficiency, react to new threats in a timely manner, and free up security personnel to work on more advanced problems within the processes in which they participate.

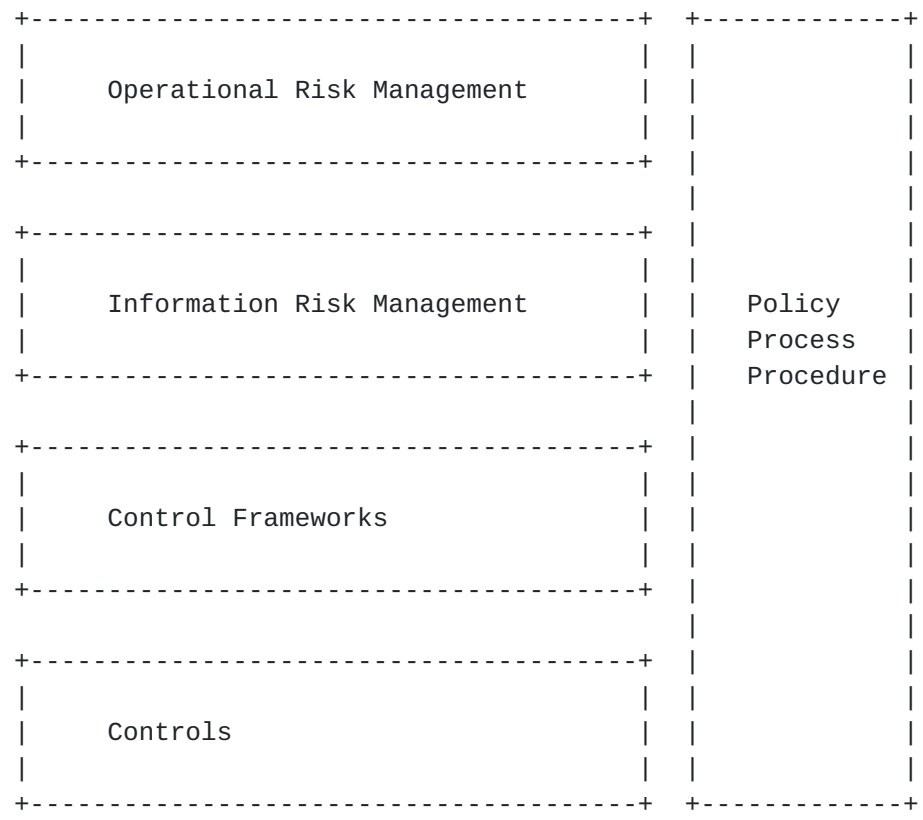

Figure 1

The figure above provides some context for our focus area. Organizations of all sizes will have a more or less formal risk management program, depending upon their maturity and organization-specific needs. A small business with only a few employees may not have a formally recognized risk management program, but they still lock the doors at night. Typically, financial entities and governments sit at the other end of the spectrum with often large, laborious risk frameworks. The point is that all organizations practice, to some degree, Operational Risk Management. An Information Risk Management program is most likely a constituent of Operational Risk Management (another constituent might be Financial Risk Management). In the Information Risk Management domain, we often use Control Frameworks to provide guidance for organizations practicing ORM in an information context, and these Control Frameworks define a variety of Controls.

From ORM, IRM, Control Frameworks, and the Controls themselves, organizations derive a set of organization-specific policies, processes, and procedures. Such policies, processes, and procedures make use of a library of supporting information commonly stipulated by the organization (i.e. enterprise acceptable use policies), but

often prescribed by external entities (i.e. Payment Card Industry Data Security Standards, Sarbanes-Oxley, or EU Data Privacy Directive). The focus of this document spans Controls, certain aspects of policy, process, and procedure, and Control Frameworks.

3. Use Cases

This document addresses three use cases: System State Assessment, Enforcement of Acceptable State, Security Control Verification and Monitoring. Currently, the first use case, System State Assessment, is being pursued under the SACM charter. The additional use cases are included to provide broader context to this work and represents additional work that may be considered by SACM or another IETF working group in the future.

3.1. UC1: System State Assessment

3.1.1. Goal

To assess the security state of a given system to be in compliance with enterprise standards and, therefore, ensure alignment with enterprise policy.

3.1.2. Main Success Scenario

1. Define target system to be assessed
2. Select acceptable state policies to apply to the defined target
3. Identify the target being assessed
4. Collect actual state values from target
5. Communicate target identity and collected state values to external system for evaluation
6. Compare actual state values collected from target with expected state values as expressed in acceptable state policies

3.2. UC2: Enforcement of Acceptable State

3.2.1. Goal

Allow or deny access to a desired resource based on system characteristics compliance with enterprise policy.

3.2.2. Main Success Scenario

1. An entity (user on a system or the system itself) requests access to a given resource (i.e. network connection)

2. Assessment of system state is achieved using Section 3.1

3. Based on assessment results (i.e. compliance level with enterprise policy)

   A. System is allowed access to requested resource, or

   B. System is denied access to requested resource

3.3. UC3: Security Control Verification and Monitoring

3.3.1. Goal

Continuous assessment of the implementation and effectiveness of security controls based on machine processable content.

3.3.2. Main Success Scenario

1. Define set of targets to be assessed.

2. Select acceptable state policies to apply to set of targets

3. Define assessment trigger based on either a

   A. Time period, or

   B. System/enterprise event.

4. Define result reporting/alerting criteria

5. Enable continuous assessment

4. Functional Capabilities and Requirements

In general, the activities of managing assets, configurations, and vulnerabilities are common between UC1, UC2, and UC3. UC2 uses these activities to either grant or deny an entity access to a requested resource. UC3 uses these activities in support of compliance measurement on a periodic basis.

At the most basic level, an enterprise needing to satisfy these use cases will need certain capabilities to be met. Specifically, we are

talking about risk management capabilities. This is the central problem domain, so it makes sense to be able to convey information about technical and non-technical controls, benchmarks, control requirements, control frameworks and other concepts in a common way.

4.1. Capabilities Supporting UC1

The capabilities in this section support assessing host and/or network state in an automated manner as described in Section Section 3.1.

4.1.1. Asset Management

Effective Asset Management is a critical foundation upon which all else in risk management is based. There are two important facets to asset management: 1) understanding coverage (how many assets are under control) and, 2) understanding specific asset details. Coverage is fairly straightforward - assessing 80% of the enterprise is better than assessing 50% of the enterprise. Getting asset details is comparatively subtle - if an enterprise does not have a precise understanding of its assets, then all acquired data and consequent actions are considered suspect. Assessing assets (managed and unmanaged) requires that we see and properly characterize our assets at the outset and over time.

4.1.1.1. Concepts

What we need to do initially is discover and characterize our assets, and then identify them in a common way. Characterization may take the form of logical characterization or security characterization, where logical characterization may include business context not otherwise related to security, but which may be used as information in support of decision making later in risk management workflows.

The following list details the requisite Asset Management capabilities:

- Discover assets in the enterprise
- Characterize assets according to security and non-security asset properties
- Identify and describe assets using a common vocabulary between implementations
- Reconcile asset representations originating from disparate tools

- Manage asset information throughout the asset's life cycle

4.1.1.2. Requirements

A method MUST be provided for identifying a target system (asset identification) as a unique entity within the enterprise.

A method MUST be provided for defining a target system (asset classification) based on a set of organizationally relevant properties (e.g. organizational affiliation, criticality, function).

4.1.2. Data Collection

Related to managing assets, and central to any automated assessment solution, is the ability to collect data from (or related to) a target device (some might call this "harvesting"). Of particular interest is data representing the security state of a target, be it a computing device, network hardware, operating system, or application. The primary interest of the activities demanding data collection is centered on object state collection, where attributes may be installed software, file properties, operating system and/or application configuration items, and network device configuration items among others.

4.1.2.1. Concepts

There are many valid perspectives to take when considering required data collection capabilities. The nature of data collected relating to assets supports a variety of information domains including: security configuration management (SCM) and vulnerability management. SCM deals with the configuration of computing and infrastructure devices, including the software installed and in use on the device. Vulnerability management involves identifying the patch level of software installed on the device and the identification of insecure custom code (e.g. web vulnerabilities). All vulnerabilities need to be addressed as part of a comprehensive risk management program, which is a superset of software vulnerabilities. Thus, the capability of assessing non-software vulnerabilities applicable to the in-scope system is required. Additionally, it may be necessary to support non-technical assessment of data relating to assets such as aspects related to operational and management controls.

The following assessment capabilities support SCM relative to a target asset:

- Collect the state of technical controls including, but not necessarily limited to:

    * Software inventory (e.g. operating system, applications, patches)

    * Configuration settings

- Collect the state of non-technical controls commonly called administrative controls (i.e. policy, process, procedure)

4.1.2.2. Requirements

One or more data formats MUST be provided to describe instructions, data collection methods, to drive data collection (e.g. technical, interrogative).

A method MUST be provided for retrieving data collection instructions from a remote host (see Section Section 4.1.4).

A data format MUST be provided to capture the results of data collection.

A mechanism MUST be provided to identify the device the results pertain to (see Section Section 4.1.1.

A mechanism MUST be provided to identify the software inventory of a device.

A mechanism MUST be provided to associate configuration settings values to the associated software.

A mechanism MUST be provided to identify additional collected attribute/value pairs related to the device, installed software, or other controls.

A mechanism MUST be provided to associate the data collection method with the collected value.

Collected data

A method of communicating data collection results to another system for further analysis MUST be identified.

4.1.3. Assessment Result Analysis

At the most basic level, the data collected needs to be analyzed for compliance to a standard stipulated by the enterprise. Analysis methods may vary between enterprises, but commonly take a similar form.

4.1.3.1. Concepts

The following capabilities support the analysis of assessment results:

- Comparing actual state to expected state
- Scoring/weighting individual comparison results
- Relating specific comparisons to benchmark-level requirements
- Relating benchmark-level requirements to one or more control frameworks

4.1.3.2. Requirements

A method MUST be provided for selecting acceptable state policy (test expression).

A method MUST be provided for comparing collected data to expected state values (test evaluation).

4.1.4. Content Management

It should be clear by now that the capabilities required to support risk management state measurement will yield volumes of content. The efficacy of risk management state measurement depends directly on the stability of the driving content, and, subsequently, the ability to change content according to enterprise needs.

4.1.4.1. Concepts

Capabilities supporting Content Management should provide the ability to create/define or modify content, as well as store and retrieve said content of at least the following types:

- Configuration Standards
- Scoring Models
- Vulnerability Information
- Patch Information
- Asset Characterization

Note that the ability to modify content is in direct support of tailoring content for enterprise-specific needs.

4.1.4.2. Requirements

A protocol MUST be identified for retrieving SACM content from a content repository

A protocol MUST be identified for querying SACM content held in a content repository. The protocol MUST support querying content by applicability to asset characteristics.

A protocol MUST be identified for curating SACM content in a content repository. Note: This might be an area where we can limit the scope of work relative to the initial SACM charter.

4.2. Capabilities Supporting UC2

UC2 is dependent upon UC1 and, therefore, includes all of the capabilities described in Section Section 4.1. UC2 describes the ability to make a resource access decision based on an assessment of the requesting system (either by the system itself or on behalf of a user operating that system). There are two chief capabilities required to meet the needs expressed in Section Section 3.2: Assessment Query and Transport, and Acceptable State Enforcement.

4.2.1. Assessment Query and Transport

Under certain circumstances, the system requesting access may be unknown, which can make querying the system problematic (consider a case where a system is connecting to the network and has no assessment software installed). Note that The Network Endpoint Assessment (NEA) protocols (PA-TNC [RFC5792], PB-TNC [RFC5793], PT-TLS [I-D.ietf-nea-pt-tls], and PT-EAP [I-D.ietf-nea-pt-eap]) may be used to query and transport the things to be measured.

4.2.2. Acceptable State Enforcement

Once the assessment has been performed a decision to allow or deny access to the requested resource can be made. Making this decision is a necessary but insufficient condition for enforcement of acceptable state, and an implementation must have the ability to actively allow or deny access to the requested resource. For example, network enforcement may be implemented with RADIUS [RFC2865] or DIAMETER [RFC6733].

4.3. Capabilities Supporting UC3

Recall that UC3 is dependent upon UC1 and therefore includes all of the capabilities described in Section 4.1. The difference in UC3 is the notion of when to assess rather than what to assess. Therefore,

the capabilities described in this section are relevant only to the "when" and not to the "what."

4.3.1. Tasking and Scheduling

The ability to task and schedule assessments is requisite for any effective risk management program. Tasking refers to the ability to create a set of instructions to be conveyed at a later time via scheduling. Tasking, therefore, involves selecting a set of assessment criteria, assigning that set to a group of assets, and expressing that information in a manner that can be consumed by a collection tool. Scheduling comes into play when the enterprise determines when to perform a specific assessment task (or set of tasks). Scheduling may be expressed in a way that constrains tasks to execute only during defined periods, can be ad hoc, or may be triggered by the analysis of previous assessment results or events detected in the enterprise.

The following capabilities support Tasking and Scheduling:

- Selection of assessment criteria
- Defining in-scope assets (i.e. targeting)
- Defining periodic assessments for a given set of tasks
- Defining assessment triggers for a given set of tasks

4.3.2. Data Aggregation and Reporting

Assessment results are produced for every asset assessed, and these results must be reported not only individually, but in the aggregate, and in accordance with enterprise needs. Enterprises should be able to aggregate and report on the data their assessments produce in a number of different ways in order to support different levels of decision making. At times, security operations personnel may be interested in understanding where the most critical risks exist in their enterprise so as to focus their remediation efforts in the most effective way (in terms of cost and return). At other times, only aggregated scores will matter, as might be the case when reporting to an information security manager or other executive-level role.

It is not the position of these capabilities to provide explicit details about how reports should be formatted for presentation, but only what information they should contain for a particular purpose. Furthermore, it is quite easy to imagine the need for a capability providing extensibility to aggregation and reporting.

Aggregating assessment results by the following capabilities supports Data Aggregation and Reporting

- By asset characterization
- By assessment criteria
- By control framework
- By benchmark
- By other attributes/properties of assessment characteristics
- Extensible aggregation and reporting

5. IANA Considerations

This memo includes no request to IANA.

All drafts are required to have an IANA considerations section (see RFC 5226 [RFC5226] for a guide). If the draft does not require IANA to do anything, the section contains an explicit statement that this is the case (as above). If there are no requirements for IANA, the section will be removed during conversion into an RFC by the RFC Editor.

6. Security Considerations

All drafts are required to have a security considerations section. See RFC 3552 [RFC3552] for a guide.

7. Acknowledgements

The author would like to thank Kathleen Moriarty and Stephen Hanna for contributing text to this document. The author would also like to acknowledge the members of the SACM mailing list for their keen and insightful feedback on the concepts and text within this document.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

[I-D.ietf-nea-pt-eap]
Cam-Winget, N. and P. Sangster, "PT-EAP: Posture Transport (PT) Protocol For EAP Tunnel Methods", draft-ietf-nea-pt-eap-06 (work in progress), December 2012.

[I-D.ietf-nea-pt-tls]
Sangster, P., Cam-Winget, N., and J. Salowey, "PT-TLS: A TLS-based Posture Transport (PT) Protocol", draft-ietf-nea-pt-tls-08 (work in progress), October 2012.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC5792] Sangster, P. and K. Narayan, "PA-TNC: A Posture Attribute (PA) Protocol Compatible with Trusted Network Connect (TNC)", RFC 5792, March 2010.

[RFC5793] Sahita, R., Hanna, S., Hurst, R., and K. Narayan, "PB-TNC: A Posture Broker (PB) Protocol Compatible with Trusted Network Connect (TNC)", RFC 5793, March 2010.

[RFC6733] Fajardo, V., Arkko, J., Loughney, J., and G. Zorn, "Diameter Base Protocol", RFC 6733, October 2012.

Authors' Addresses

David Waltermire (editor)
National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, Maryland 20877
USA

Phone:
Email: david.waltermire@nist.gov

Adam W. Montville
Tripwire, Inc.
101 SW Main Street, Suite 1500
Portland, Oregon 97204
USA

Phone:
Email: amontville@tripwire.com