Security Automation and Continuous Monitoring WG D. Waltermire
Internet-Draft NIST
Intended status: Informational D. Harrington
Expires: April 17, 2014 Effective Software
October 14, 2013

Endpoint Security Posture Assessment - Enterprise Use Cases
draft-ietf-sacm-use-cases-02

Abstract

This memo documents a sampling of use cases for securely aggregating configuration and operational data and assessing that data to determine an organization's security posture. From these operational use cases, we can derive common functional capabilities and requirements to guide development of vendor-neutral, interoperable standards for aggregating and assessing data relevant to security posture.

Table of Contents

1. Introduction

Our goal with this document is to improve our agreement on which problems we're trying to solve. We need to start with short, simple problem statements and discuss those by email and in person. Once we agree on which problems we're trying to solve, we can move on to propose various solutions and decide which ones to use.

This document describes example use cases for endpoint posture assessment for enterprises. It provides a sampling of use cases for securely aggregating configuration and operational data and assessing that data to determine the security posture of individual endpoints, and, in the aggregate, the security posture of an enterprise.

These use cases cross many IT security information domains. From these operational use cases, we can derive common concepts, common information expressions, functional capabilities and requirements to

guide development of vendor-neutral, interoperable standards for aggregating and assessing data relevant to security posture.

Using this standard data, tools can analyze the state of endpoints, user activities and behaviour, and assess the security posture of an organization. Common expression of information should enable interoperability between tools (whether customized, commercial, or freely available), and 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.

The goal is to enable organizations to make informed decisions that support organizational objectives, to enforce policies for hardening systems, to prevent network misuse, to quantify business risk, and to collaborate with partners to identify and mitigate threats.

It is expected that use cases for enterprises and for service providers will largely overlap, but there are additional complications for service providers, especially in handling information that crosses administrative domains.

The output of endpoint posture assessment is expected to feed into additional processes, such as policy-based enforcement of acceptable state, verification and monitoring of security controls, and compliance to regulatory requirements.

2. Endpoint Posture Assessment

Endpoint posture assessment involves orchestrating and performing data collection and analysis pertaining to the posture of a given endpoint. Typically, endpoint posture information is gathered and then published to appropriate data repositories to make collected information available for further analysis supporting organizational security processes.

Endpoint posture assessment typically includes:

- Collecting the posture of a given endpoint;
- Making that posture available to the enterprise for further analysis and action; and
- Performing analysis to assess that the endpoint's posture is in compliance with enterprise standards and policy.

As part of these activities it is often necessary to identify and acquire any supporting content that is needed to drive data collection and analysis.

The following is a typical workflow scenario for assessing endpoint posture:

1. Define a target endpoint to be assessed

2. Select policies applicable to the target, and identify what posture attributes need to be collected for assessment.

3. Verify the identity of the target being assessed

4. Collect posture attributes from the target

5. Communicate target identity and collected attributes to external system for evaluation

6. Evaluator compares collected posture attributes from the target with expected values as expressed in policies

The following subsections detail specific use cases for data collection, analysis, and related operations pertaining to the publication and use of supporting content.

2.1. Definition and Publication of Automatable Configuration Guides

A network device vendor manufactures a number of enterprise grade routers and other network devices. They also develop and maintain an operating system for these devices that enables end-user organizations to configure a number of security and operational settings for these devices. As part of their customer support activities, they publish a number of secure configuration guides that provide minimum security guidelines for configuring their devices.

Each guide they produce applies to a specific model of device and version of the operating system and provides a number of specialized configurations depending on the devices intended function and what add-on hardware modules and software licenses are installed on the device. To enable their customers to assess the security posture of their devices to ensure that all appropriate minimal security settings are enabled, they publish an automatable configuration checklist using a popular data format that defines what settings to check using a network management protocol and appropriate values for each setting. They publish these guides to a public content repository that customers can query to retrieve applicable guides for their deployed enterprise network infrastructure endpoints.

Guides could also come from sources other than a device vendor, such as industry groups or regulatory authorities, or enterprises could develop their own checklists.

QUESTION: This use case applies equally to vendors representing other endpoint types. Should this be generalized to capture this notion?

QUESTION: Is providing traceability to functional capabilities useful? If so, we need to replicate this for the other use cases.

2.2. Automated Checklist Verification

A financial services company operates a heterogeneous IT environment. In support of their risk management program, they utilize vendor provided automatable security configuration checklists for each operating system and application used within their IT environment. Multiple checklists are used from different vendors to insure adequate coverage of all IT assets.

To identify what checklists are needed, they use automation to gather an inventory of the software versions utilized by all IT assets in the enterprise. This data gathering will involve querying existing data stores of previously collected endpoint software inventory posture data and actively collecting data from reachable endpoints as needed utilizing network and systems management protocols. Previously collected data may be provided by periodic data collection, network connection-driven data collection, or ongoing event-driven monitoring of endpoint posture changes.

Using the gathered software inventory data and associated asset management data indicating the organizational defined functions of each endpoint, they locate and query each vendors content repository for the appropriate checklists. These checklists are cached locally to reduce the need to download the checklist multiple times.

Driven by the setting data provided in the checklist, a combination of existing configuration data stores and data collection methods are used to gather the appropriate posture information from each endpoint. Specific data is gathered based on the defined enterprise function and software inventory of each endpoint. The data collection paths used to collect software inventory posture will be used again for this purpose. Once the data is gathered, the actual state is evaluated against the expected state criteria in each applicable checklist. Deficiencies are identified and reported to the appropriate endpoint operators for remedy.

Checklists could also come from sources other than the application or OS vendor, such as industry groups or regulatory authorities, or enterprises could develop their own checklists.

2.3. Organizational Software Policy Compliance

Example Corporation, in support of compliance requirements, has identified a number of secure baselines for different endpoint types that exist across their enterprise IT environment. Determining which baseline applies to a given endpoint is based on the organizationally defined function of the device.

Each baseline, defined using an automatable standardized data format, identifies the expected hardware, software and patch inventory, and software configuration item values for each endpoint type. As part of their compliance activities, they require that all endpoints connecting to their network meet the appropriate baselines. Each endpoint is checked to make sure it complies with the appropriate baseline whenever it connects to the network and at least once a day thereafter. These daily compliance checks assess the posture of each endpoint and report on its compliance with the appropriate baseline.

[TODO: Need to speak to how the baselines are identified for a given endpoint connecting to the network.]

2.4. Detection of Posture Deviations

Example corporation has established secure configuration baselines for each different type of endpoint within their enterprise including: network infrastructure, mobile, client, and server computing platforms. These baselines define an approved list of hardware, software (i.e., operating system, applications, and patches), and associated required configurations. When an endpoint connects to the network, the appropriate baseline configuration is communicated to the endpoint based on its location in the network, the expected function of the device, and other asset management data. It is checked for compliance with the baseline indicating any deviations to the device's operators. Once the baseline has been established, the endpoint is monitored for any change events pertaining to the baseline on an ongoing basis. When a change occurs to posture defined in the baseline, updated posture information is exchanged allowing operators to be notified and/or automated action to be taken.

2.5. Search for Signs of Infection

TODO

2.6. Remediation and Mitigation

TODO

2.7. Endpoint Information Analysis and Reporting

TODO

2.8. Others...

Additional use cases will be identified as we work through other domains.

3. IANA Considerations

This memo includes no request to IANA.

4. Security Considerations

This memo documents, for Informational purposes, use cases for security automation. While it is about security, it does not affect security.

5. Acknowledgements

The National Institute of Standards and Technology (NIST) and/or the MITRE Corporation have developed specifications under the general term "Security Automation" including languages, protocols, enumerations, and metrics.

The authors would like to recognize and thank Adam Montville for his work on early edits of this draft. Additionally, the authors would like to thank Kathleen Moriarty and Stephen Hanna for contributing text to this document. The authors 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.

6. Change Log

6.1. -00- to -01-

Changed title

removed section 4, expecting it will be moved into the requirements document.

removed the list of proposed caabilities from section 3.1

Added empty sections for Search for Signs of Infection, Remediation and Mitigation, and Endpoint Information Analysis and Reporting.

Removed Requirements Language section and rfc2119 reference.

Removed unused references (which ended up being all references).

6.2. -00- to -01-

- Work on this revision has been focused on document content relating primarily to use of asset management data and functions.
- Made significant updates to section 3 including:
  - Reworked introductory text.
  - Replaced the single example with multiple use cases that focus on more discrete uses of asset management data to support hardware and software inventory, and configuration management use cases.
  - For one of the use cases, added mapping to functional capabilities used. If popular, this will be added to the other use cases as well.
  - Additional use cases will be added in the next revision capturing additional discussion from the list.
- Made significant updates to section 4 including:
  - Renamed the section heading from "Use Cases" to "Functional Capabilities" since use cases are covered in section 3. This section now extrapolates specific functions that are needed to support the use cases.
  - Started work to flatten the section, moving select subsections up from under asset management.
  - Removed the subsections for: Asset Discovery, Endpoint Components and Asset Composition, Asset Resources, and Asset Life Cycle.
  - Renamed the subsection "Asset Representation Reconciliation" to "Deconfliction of Asset Identities".
  - Expanded the subsections for: Asset Identification, Asset Characterization, and Deconfliction of Asset Identities.

    * Added a new subsection for Asset Targeting.

    * Moved remaining sections to "Other Unedited Content" for future updating.

6.3. draft-waltermire-sacm-use-cases-05 to draft-ietf-sacm-use-cases-00

- Transitioned from individual I/D to WG I/D based on WG consensus call.
- Fixed a number of spelling errors. Thank you Erik!
- Added keywords to the front matter.
- Removed the terminology section from the draft. Terms have been moved to: draft-dbh-sacm-terminology-00
- Removed requirements to be moved into a new I/D.
- Extracted the functionality from the examples and made the examples less prominent.
- Renamed "Functional Capabilities and Requirements" section to "Use Cases".
    * Reorganized the "Asset Management" sub-section. Added new text throughout.
        + Renamed a few sub-section headings.
        + Added text to the "Asset Characterization" sub-section.
- Renamed "Security Configuration Management" to "Endpoint Configuration Management". Not sure if the "security" distinction is important.
    * Added new sections, partially integrated existing content.
    * Additional text is needed in all of the sub-sections.
- Changed "Security Change Management" to "Endpoint Posture Change Management". Added new skeletal outline sections for future updates.

6.4. -04- to -05-

- Are we including user activities and behavior in the scope of this work? That seems to be layer 8 stuff, appropriate to an IDS/IPS application, not Internet stuff.

- I removed the references to what the WG will do because this belongs in the charter, not the (potentially long-lived) use cases document. I removed mention of charter objectives because the charter may go through multiple iterations over time; there is a website for hosting the charter; this document is not the correct place for that discussion.

- I moved the discussion of NIST specifications to the acknowledgements section.

- Removed the portion of the introduction that describes the chapters; we have a table of concepts, and the existing text seemed redundant.

- Removed marketing claims, to focus on technical concepts and technical analysis, that would enable subsequent engineering effort.

- Removed (commented out in XML) UC2 and UC3, and eliminated some text that referred to these use cases.

- Modified IANA and Security Consideration sections.

- Moved Terms to the front, so we can use them in the subsequent text.

- Removed the "Key Concepts" section, since the concepts of ORM and IRM were not otherwise mentioned in the document. This would seem more appropriate to the arch doc rather than use cases.

- Removed role=editor from David Waltermire's info, since there are three editors on the document. The editor is most important when one person writes the document that represents the work of multiple people. When there are three editors, this role marking isn't necessary.

- Modified text to describe that this was specific to enterprises, and that it was expected to overlap with service provider use cases, and described the context of this scoped work within a larger context of policy enforcement, and verification.

- The document had asset management, but the charter mentioned asset, change, configuration, and vulnerability management, so I added sections for each of those categories.

- Added text to Introduction explaining goal of the document.

- Added sections on various example use cases for asset management, config management, change management, and vulnerability management.

7. References

7.1. Normative References

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

7.2. Informative References

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

Authors' Addresses

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

Email: david.waltermire@nist.gov

David Harrington
Effective Software
50 Harding Rd
Portsmouth, NH 03801
USA

Email: ietfdbh@comcast.net