

Internet Engineering Task Force (IETF)
Request for Comments: 8248
Category: Informational
ISSN: 2070-1721

N. Cam-Winget
Cisco Systems
L. Lorenzin
Pulse Secure
September 2017

Security Automation and Continuous Monitoring (SACM) Requirements

Abstract

This document defines the scope and set of requirements for the Security Automation and Continuous Monitoring (SACM) architecture, data model, and transfer protocols. The requirements and scope are based on the agreed-upon use cases described in RFC 7632.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8248>.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction 2
 - 1.1. Requirements Language 3
- 2. Requirements 3
 - 2.1. Requirements for SACM 4
 - 2.2. Requirements for the Architecture 7
 - 2.3. Requirements for the Information Model 9
 - 2.4. Requirements for the Data Model 10
 - 2.5. Requirements for Data Model Operations 12
 - 2.6. Requirements for SACM Transfer Protocols 14
- 3. IANA Considerations 15
- 4. Security Considerations 15
 - 4.1. Trust between Provider and Requestor 16
 - 4.2. Privacy Considerations 17
- 5. References 18
 - 5.1. Normative References 18
 - 5.2. Informative References 18
- Acknowledgments 18
- Authors' Addresses 19

1. Introduction

Today's environment of rapidly evolving security threats highlights the need to automate the sharing of security information (such as posture information) while protecting user information and the systems that store, process, and transmit this information. Security threats can be detected in a number of ways. The Security Automation and Continuous Monitoring (SACM) charter focuses on how to collect and share this information based on use cases that involve posture assessment of endpoints.

Scalable and sustainable collection, expression, and evaluation of endpoint information is foundational to SACM's objectives. To secure and defend a network, one must reliably determine what devices are on the network, how those devices are configured from a hardware perspective, what software products are installed on those devices, and how those products are configured. We need to be able to determine, share, and use this information in a secure, timely, consistent, and automated manner to perform endpoint posture assessments.

This document focuses on describing the requirements for facilitating the exchange of posture assessment information in the enterprise, in particular, for the use cases as exemplified in [RFC7632].

As proposals are evaluated for SACM standardization, the documents describing each proposal are expected to include a section that describes how the enumerated requirements are addressed.

This document uses terminology defined in [TERMS].

1.1. Requirements Language

Use of each capitalized word within a sentence or phrase carries the following meaning during the SACM WG's protocol selection process:

MUST - indicates an absolute requirement

MUST NOT - indicates something absolutely prohibited

SHOULD - indicates a strong recommendation of a desired result

SHOULD NOT - indicates a strong recommendation against a result

MAY - indicates a willingness to allow an optional outcome

When the words appear in lower case, their natural language meaning is used.

2. Requirements

This document defines requirements based on the SACM use cases described in [RFC7632]. This section describes the requirements used by SACM to assess and compare candidate data models, interfaces, and protocols. These requirements express characteristics or features that a candidate protocol, information model, or data model must be capable of offering to ensure security and interoperability.

Multiple data models, protocols, and transfers may be employed in a SACM environment. A SACM transfer protocol is one that runs on top of transport-layer protocols such as TCP/IP or internet-layer protocols such as HTTP, carries operations (requests/responses), and moves data.

SACM will define an architecture and information model focused on addressing the needs for determining, sharing, and using posture information securely via posture information providers and posture information consumers. With the information model defining assets and attributes to facilitate the guidance, collection, and assessment of posture, tasks that should be considered include:

1. Asset Classification: Map the target endpoint and/or the assets on the target endpoints to asset classes. This enables

identification of the attributes needed to exchange information pertaining to the target endpoint.

2. Attribute Definition: Define the attributes desired to be collected from each target endpoint. For instance, organizations will want to know what software is installed and its critical security attributes such as patch level.
3. Policy Definition: This is where an organization can express its policy for acceptable or problematic values of an endpoint attribute. The expected values of an endpoint attribute are determined for later comparison against the actual endpoint attribute values during the evaluation process. Expected values may include both values that are good as well as values that represent problems, such as vulnerabilities. The organization can also specify the endpoint attributes that are to be present for a given target endpoint.
4. Information Collection: Collect information (attribute values) from the target endpoint to populate the endpoint data.
5. Endpoint Assessment: Evaluate the actual values of the endpoint attributes against those expressed in the policy. (An evaluation result may become additional endpoint data.)
6. Result Reporting: Report the results of the evaluation for use by other components. Examples of the use of a report would be additional evaluation, network enforcement, vulnerability detection, and license management.

2.1. Requirements for SACM

Many deployment scenarios can be instantiated to address the above tasks and the use cases defined in [RFC7632]. To ensure interoperability, scalability, and flexibility in any of these deployments, the following requirements are defined for proposed SACM standards:

G-001 (Solution Extensibility): The information model, data models, protocols, and transfers defined by SACM MUST be designed to allow support for future extensions. SACM MUST allow for both standardized and proprietary extensions.

1. The information model and programmatic interfaces (see G-012 for one example) MUST support the ability to add new operations while maintaining backwards compatibility. SACM-defined transfer protocols MUST have extensibility to allow them to transfer operations that are defined in the future.

2. The query language MUST allow for general inquiries as well as expression of specific attributes or relationships between attributes; the retrieval of specific information based on an event or on a continuous basis; and the ability to retrieve specific pieces of information, specific types or classes of information, or the entirety of available information.
3. The information model MUST accommodate the interoperable addition of new data types and/or schemas.

G-002 (Interoperability): The data models, protocols, and transports MUST be specified with enough details to ensure interoperability.

G-003 (Scalability): SACM needs to support a broad set of deployment scenarios. The data models, protocols, and transports have to be scalable unless they are specifically defined to apply to a special-purpose scenario, such as constrained devices. A SACM transfer protocol standard SHOULD include a section on scalability considerations that addresses the number of endpoints and amount of information to which it can reasonably be expected to scale. Scalability must be addressed to support:

- * Large messages: It is possible that the size of posture assessment information can vary from a single assessment that is small in size to a very large message or a very large set of assessments (up to multiple gigabytes in size).
- * Large number of messages per second: A deployment may involve many rapid or simultaneous events that require processing, generating many messages per second.
- * Large number of providers and consumers: A deployment may consist of a very large number of endpoints requesting and/or producing posture assessment information.
- * Large number of target endpoints: A deployment may be managing information of a very large number of target endpoints.

G-004 (Versatility): The data model, protocols, and transports must be suitably specified to enable implementations to fit into different deployment models and scenarios, including considerations for implementations of data models and transports operating in constrained environments. Separate solutions may be necessary to meet the needs of specific deployment models and scenarios.

G-005 (Information Extensibility): Non-standard (implementation-specific) attributes MUST be supported. A method SHOULD be defined for preventing collisions from occurring in the naming of all

attributes independent of their source. For interoperability and scope boundary, the information model MUST define the mandatory set of attributes.

G-006 (Data Protection): To protect the information being shared, SACM components MUST protect the integrity and confidentiality of data in transit (end to end) and data at rest (as information is stored in repositories). Mechanisms for this protection are unspecified but should include industry best practices. These mechanisms are required to be available (i.e., all data-handling components must support them) but are not required to be used in all cases.

G-007 (Data Partitioning): A method for partitioning data MUST be supported to accommodate considerations such as geographic, regulatory, operational requirements, overlay boundaries, and federation (where the data may be collected in multiple locations and either centralized or kept in the local region). Where replication of data is supported, it is required that methods exist to prevent update loops.

G-008 (Versioning and Backward Compatibility): Announcement and negotiation of versions, inclusive of existing capabilities (such as transfer protocols, data models, specific attributes within data models, standard attribute expression sets, etc.) MUST be supported. Negotiation for both versioning and capabilities is needed to accommodate future growth and ecosystems with mixed capabilities.

G-009 (Information Discovery): There MUST be mechanisms for components to discover what information is available across the ecosystem (i.e., a method for cataloging data available in the ecosystem and advertising it to consumers), where to go to get a specific piece of that information (i.e., which provider has the information), and what schemas are in use for organizing the information. For example, a method can be provided by which a node can locate the advertised information so that consumers are not required to have a priori knowledge to find available information.

G-010 (Target Endpoint Discovery): SACM MUST define the means by which target endpoints may be discovered. The use case in Section 2.1.2 of [RFC7632] describes the need to discover endpoints and their composition.

G-011 (Push and Pull Access): Three methods of data access MUST be supported: a Pull model, a solicited Push model, and an unsolicited Push model. All of the methods of data access MUST support the ability for the initiator to filter the set of posture assessment information to be delivered. Additionally, the provider of the

information MUST be able to filter the set of posture assessment information based on the permissions of the recipient. This requirement is driven by the use cases in Sections 2.1.3 and 2.1.4 of [RFC7632].

G-012 (SACM Component Interface): The interfaces by which SACM components communicate to share endpoint posture information MUST be well defined. That is, the interface defines the data model, SACM transfer protocols, and network transfer protocols to enable SACM components to communicate.

G-013 (Endpoint Location and Network Topology): The SACM architecture and interfaces MUST allow for the target endpoint (network) location and network topology to be modeled and understood. Where appropriate, the data model and the interfaces SHOULD allow for discovery of the target endpoint location, network topology, or both.

G-014 (Target Endpoint Identity): The SACM architecture and interfaces MUST support the ability of components to provide attributes that can be used to compose an identity for a target endpoint. These identities MAY be composed of attributes from one or more SACM components.

G-015 (Data Access Control): Methods of access control must be supported to accommodate considerations such as geographic, regulatory, operational, and federations. Entities accessing or publishing data MUST identify themselves and pass access policy.

2.2. Requirements for the Architecture

Following are the requirements for the SACM architecture:

ARCH-001 (Component Functions): At the simplest abstraction, the SACM architecture MUST represent the core components and interfaces needed to perform the production and consumption of posture assessment information.

ARCH-002 (Scalability): The architectural components MUST account for a range of deployments, from very small sets of endpoints to very large deployments.

ARCH-003 (Flexibility): The architectural components MUST account for different deployment scenarios where the architectural components may be implemented, deployed, or used within a single application, service, or network, or may comprise a federated system.

ARCH-004 (Separation of Data and Management Functions): SACM MUST define both the configuration and management of the SACM data models and protocols used to transfer and share posture assessment information.

ARCH-005 (Topology Flexibility): Both centralized and decentralized (peer-to-peer) information exchange MUST be supported. Centralized data exchange enables use of a common data format to bridge together data exchange between diverse systems and can leverage a virtual data store that centralizes and offloads all data access, storage, and maintenance to a dedicated resource. Decentralized data exchange enables simplicity of sharing data between relatively uniform systems and between small numbers of systems, especially within a single enterprise domain. The fact that a centralized or decentralized deployment is used SHOULD be invisible to a consumer. However, there may be cases where the producer chooses to include that information due to consumer preference.

ARCH-006 (Capability Negotiation): Announcement and negotiation of functional capabilities (such as authentication protocols, authorization schemes, data models, transfer protocols, etc.) MUST be supported, enabling a SACM component to make inquiries about the capabilities of other components in the SACM ecosystem.

ARCH-007 (Role-Based Authorization): The SACM architecture MUST be capable of effecting role-based authorization. Distinction of endpoints capable of and authorized to provide or consume information is required to address appropriate access controls.

ARCH-008 (Context-Based Authorization): The SACM architecture MUST be capable of effecting context-based authorization. Different policies (e.g., business, regulatory, etc.) might specify what data may be exposed to, or shared by, consumers based on one or more attributes of the consumer. The policy might specify that consumers are required to share specific information either back to the system or to administrators.

ARCH-009 (Time Synchronization): Actions or decisions based on time-sensitive data (such as user logon/logoff, endpoint connection/disconnection, endpoint behavior events, etc.) are all predicated on a synchronized understanding of time. The SACM architecture MUST provide a mechanism for all components to synchronize time. A mechanism for detecting and reporting time discrepancies SHOULD be provided by the architecture and reflected in the information model.

2.3. Requirements for the Information Model

The SACM information model represents the abstracted representation for posture assessment information to be communicated. SACM data models must adhere to and comply with the SACM information model. The requirements for the SACM information model include:

IM-001 (Extensible Attribute Vocabulary): The information model **MUST** define a minimum set of attributes for communicating posture information, to ensure interoperability between data models. (Individual data models may define attributes beyond the mandatory-to-implement minimum set.) The attributes should be defined with a clear mechanism for extensibility to enable data models to adhere to SACM's required attributes as well as allow for their own extensions. The attribute vocabulary should be defined with a clear mechanism for extensibility to enable future versions of the information model to be interoperably expanded with new attributes.

IM-002 (Posture Data Publication): The information model **MUST** allow for the data to be provided by a SACM component either solicited or unsolicited. No aspect of the information model should be dependent upon or assume a Push or Pull model of publication.

IM-003 (Data Model Negotiation): SACM's information model **MUST** allow support for different data models, data model versions, and different versions of the operations on the data models and transfer protocols. The SACM information model **MUST** include the ability to discover and negotiate the use of a particular data model or any data model.

IM-004 (Data Model Identification): The information model **MUST** provide a means to uniquely identify each data model. The identifier **MUST** contain both an identifier of the data model and a version indicator for the data model. The identifiers **SHOULD** be decomposable so that a customer can query for any version of a specific data model and compare returned values for older or newer than a desired version.

IM-005 (Data Lifetime Management): The information model **MUST** provide a means to allow data models to include data lifetime management. The information model must identify attributes that can allow data models to, at minimum, identify the data's origination time and expected time of next update or data longevity (how long the data should be assumed to still be valid).

IM-006 (Singularity and Modularity): The SACM information model **MUST** be singular (i.e., there is only one information model, not multiple alternative information models from which to choose) and **MAY** be

modular (a conjunction of several subcomponents) for ease of maintenance and extension. For example, endpoint identification could be an independent subcomponent of the information model, to simplify updating of endpoint identification attributes.

2.4. Requirements for the Data Model

The SACM information model represents an abstraction for "what" information can be communicated and "how" it is to be represented and shared. It is expected that as applications may produce posture assessment information, they may share it using a specific data model. Similarly, applications consuming or requesting posture assessment information may require that it be based on a specific data model. Thus, while there may exist different data models and schemas, they should adhere to the SACM information model and meet the requirements defined in this section.

The specific requirements for candidate data models include:

DM-001 (Element Association): A SACM information model consists of a set of SACM information model elements. A SACM data model **MUST** be derived from the SACM information model. A SACM data model consists of a set of SACM data model elements. In this derivation, a SACM data model element **MAY** map to one or more SACM information model elements. In addition, a SACM data model **MAY** include additional data model elements that are not associated with any SACM information model elements.

DM-002 (Data Model Structure): The data model can be structured either as one single module or separated into modules and submodules that allow for references between them. The data model structure **MAY** reflect structure in the information model but does not need to. For example, the data model might use one module to define endpoints, and that module might reference other modules that describe the various assets associated with the endpoint. Constraints and interfaces might further be defined to resolve or tolerate ambiguity in the references (e.g., the same IP address used in two separate networks).

DM-003 (Search Flexibility): The search interfaces and actions **MUST** include the ability to start a search anywhere within a data model structure and the ability to search based on patterns ("wildcard searches") as well as specific data elements.

DM-004 (Full vs. Partial Updates): The data model **SHOULD** include the ability to allow providers of data to provide the data as a whole or when updates occur. For example, a consumer can request a full update on initial engagement, then request to receive deltas

(updates containing only the changes since the last update) on an ongoing basis as new data is generated.

DM-005 (Loose Coupling): The data model SHOULD allow for a loose coupling between the provider and the consumer, such that the consumer can request information without being required to request it from a specific provider, and a provider can publish information without having a specific consumer targeted to receive it.

DM-006 (Data Cardinality): The data model MUST describe their constraints (e.g., cardinality). As posture information and the tasks for collection, aggregation, or evaluation could comprise one or more attributes, interfaces and actions MUST allow and account for such cardinality and for conditional, optional, or mandatory attributes.

DM-007 (Data Model Negotiation): The interfaces and actions in the data model MUST include capability negotiation to enable discovery of supported and available data types and schemas.

DM-008 (Data Origin): The data model MUST include the ability for consumers to identify the data origin (provider that collected the data).

DM-009 (Origination Time): The data model SHOULD allow the provider to include the information's origination time.

DM-010 (Data Generation): The data model MUST allow the provider to include attributes defining how the data was generated (e.g., self-reported, reported by aggregator, scan result, etc.).

DM-011 (Data Source): The data model MUST allow the provider to include attributes identifying the data source (target endpoint from which the data was collected), e.g., hostname, domain (DNS) name, or application name.

DM-012 (Data Updates): The data model SHOULD allow the provider to include attributes defining whether the information provided is a delta, partial, or full set of information.

DM-013 (Multiple Collectors): The data model MUST support the collection of attributes by a variety of collectors, including internal collectors, external collectors with an authenticated relationship with the endpoint, and external collectors based on network and other observers.

DM-014 (Attribute Extensibility): All of the use cases in Section 2 of [RFC7632] describe the need for an attribute dictionary. With

SACM's scope focused on posture assessment, the data model attribute collection and aggregation MUST have a well-understood set of attributes inclusive of their meaning or usage intent. The data model MUST include all attributes defined in the information model and MAY include additional attributes beyond those found in the information model. Additional attributes MUST be defined in accordance with the extensibility framework provided in the information model (see IM-001).

DM-015 (Solicited vs. Unsolicited Updates): The data model MUST enable a provider to publish data either solicited (in response to a request from a consumer) or unsolicited (as new data is generated, without a request required). For example, an external collector can publish data in response to a request by a consumer for information about an endpoint, or it can publish data as it observes new information about an endpoint, without any specific consumer request triggering the publication; a compliance-server provider may publish endpoint posture information in response to a request from a consumer (solicited), or it may publish posture information driven by a change in the posture of the endpoint (unsolicited).

DM-016 (Transfer Agnostic): The data model MUST be transfer agnostic, to allow for the data operations to leverage the most appropriate SACM transfer protocol.

2.5. Requirements for Data Model Operations

Posture information data adhering to a data model must also provide interfaces that include operations for access and production of the data. Operations requirements are distinct from transfer requirements in that operations requirements are requirements on the application performing requests and responses, whereas transfer requirements are requirements on the transfer protocol carrying the requests and responses. The specific requirements for such operations include:

OP-001 (Time Synchronization): Request and response operations MUST be timestamped, and published information SHOULD capture time of publication. Actions or decisions based on time-sensitive data (such as user logon/logoff, endpoint connection/disconnection, endpoint behavior events, etc.) are all predicated on a synchronized understanding of time. A method for detecting and reporting time discrepancies SHOULD be provided.

OP-002 (Collection Abstraction): Collection is the act of a SACM component gathering data from a target endpoint. The request for a data item MUST include enough information to properly identify the item to collect, but the request shall not be a command to directly

execute nor be directly applied as arguments to a command. The purpose of this requirement is primarily to reduce the potential attack vectors but has the additional benefit of abstracting the request for collection from the collection method, thereby allowing more flexibility in how collection is implemented.

OP-003 (Collection Composition): A collection request MAY be composed of multiple collection requests (which yield collected values). The desire for multiple values MUST be expressed as part of the collection request, so that the aggregation can be resolved at the point of collection without having to interact with the requestor. This requirement should not be interpreted as preventing a collector from providing attributes that were not part of the original request.

OP-004 (Attribute-Based Query): A query operation is the act of requesting data from a provider. Query operations SHOULD be based on a set of attributes. Query operations MUST support both a query for specific attributes and a query for all attributes. The use case in Section 2.1.2 of [RFC7632] describes the need for the data model to support a query operation based on a set of attributes to facilitate collection of information such as posture assessment, inventory (of endpoints or endpoint components), and configuration checklist.

OP-005 (Information-Based Query with Filtering): The query operation MUST support filtering. The use case in Section 2.1.3 of [RFC7632] describes the need for the data model to support the means for the information to be collected through a query mechanism. Furthermore, the query operation requires filtering capabilities to allow for only a subset of information to be retrieved. The query operation MAY be a synchronous request or asynchronous request.

OP-006 (Operation Scalability): The operation resulting from a query operation MUST be able to handle the return and receipt of large amounts of data. The use case in Section 2.1.4 of [RFC7632] describe the need for the data model to support scalability. For example, the query operation may result in a very large set of attributes as well as a large set of targets.

OP-007 (Data Abstraction): The data model MUST allow a SACM component to communicate what data was used to construct the target endpoint's identity, so that other SACM components can determine whether they are constructing an equivalent target endpoint (and its identity) and whether they have confidence in that identity. SACM components SHOULD have interfaces defined to transmit this data directly or to refer to where the information can be retrieved.

OP-008 (Provider Restriction): Request operations MUST include the ability to restrict the data to be provided by a specific provider or a provider with specific characteristics. Response operations MUST include the ability to identify the provider that supplied the response. For example, a SACM consumer should be able to request that all of the data come from a specific provider by identity (e.g., Provider A) or from a provider that is in a specific location (e.g., in the Boston office).

2.6. Requirements for SACM Transfer Protocols

The term "SACM transfer protocol" is intended to be distinguished from underlying transport- and internet-layer protocols such as TCP/IP or operating at an application-layer protocol such as HTTP. The SACM transfer protocol is focused on moving data and performing necessary access control operations; it is agnostic to the data model operations.

The requirements for SACM transfer protocols include:

T-001 (Multiple Transfer Protocol Support): SACM transfer protocols will vary depending on the deployment model that relies on different transfer-layer requirements, different device capabilities, and system configurations dealing with connectivity. For example, where posture attributes may be collected directly from an endpoint using the Network Endpoint Assessment (NEA) model [RFC5209], different transports may be defined to collect them using Posture Transport Protocol for Extensible Authentication Protocol Tunnel Methods (PT-EAP) [RFC7171] or Posture Transport Protocol over TLS (PT-TLS) [RFC6876], depending on the deployment scenario.

T-002 (Data Integrity): SACM transfer protocols MUST be able to ensure data integrity for data in transit.

T-003 (Data Confidentiality): SACM transfer protocols MUST be able to support data confidentiality. SACM transfer protocols MUST ensure data protection for data in transit (e.g., by encryption) to provide confidentiality, integrity, and robustness against protocol-based attacks. Note that while the transfer MUST be able to support data confidentiality, implementations MAY provide a configuration option that enables and disables confidentiality in deployments. Protection for data at rest is not in scope for transfer protocols. Data protection MAY be used for both privacy and non-privacy scenarios.

T-004 (Transfer Protection): SACM transfer protocols MUST be capable of supporting mutual authentication and replay protection.

T-005 (Transfer Reliability): SACM transfer protocols MUST provide reliable delivery of data. This includes the ability to perform fragmentation and reassembly and to detect replays. The SACM transfer may take advantage of reliability features in the network transport; however, the network transport may be unreliable (e.g., UDP), in which case the SACM transfer running over the unreliable network transport is responsible for ensuring reliability (i.e., by provisions such as confirmations and retransmits).

T-006 (Transfer-Layer Requirements): Each SACM transfer protocol MUST clearly specify the transport-layer requirements it needs to operate correctly. Examples of items that may need to be specified include connectivity requirements, replay requirements, data link encryption requirements, and/or channel-binding requirements. These requirements are needed in order for deployments to be done correctly.

T-007 (Transfer Protocol Adoption): SACM SHOULD, where reasonably possible, leverage and use existing IETF transfer protocols versus defining new ones.

3. IANA Considerations

This document does not require any IANA actions.

4. Security Considerations

This document defines the requirements for SACM. As such, it is expected that several data models, protocols, and transfer protocols may be defined or reused from already-existing standards.

To address security and privacy considerations, the data model, protocols, and transports must consider authorization based on consumer function and privileges, to only allow authorized consumers and providers to access specific information being requested or published.

To enable federation across multiple entities (such as across organizational or geographic boundaries), authorization must also extend to infrastructure elements themselves, such as central controllers, brokers, and data repositories.

In addition, authorization needs to extend to specific information or resources available in the environment. In other words, authorization is based on the subject (the information requestor), the provider (the information responder), the object (the endpoint the information is being requested on), and the attribute (what piece

of data is being requested). The method by which this authorization is applied is unspecified.

SACM's charter focuses on the workflow orchestration and the sharing of posture information for improving the efficacy of security applications such as compliance, configuration, assurance, and other threat and vulnerability reporting and remediation systems. While the goal is to facilitate the flow of information securely, it is important to note that participating endpoints may not be cooperative or trustworthy.

4.1. Trust between Provider and Requestor

The information given from the provider to a requestor may come with different levels of trustworthiness given the different potential deployment scenarios and compromise at the provider, the requesting consumer, or devices that are involved in the transfer between the provider and requestor. This section will describe the different considerations that may reduce the level of trustworthiness of the information provided.

In the information transfer flow, it is possible that some of the devices may serve as proxies or brokers and, as such, may be able to observe the communications flowing between an information provider and requestor. Without appropriate protections, it is possible for these proxies and brokers to inject and affect man-in-the-middle attacks.

In general, it is common to distrust the network service provider, unless the full hop-by-hop communications process flow is well understood. As such, the posture information provider should protect the posture information data it provides as well as the transfer it uses. Similarly, while there may be providers whose goal is to openly share its information, there may also be providers whose policy is to grant access to certain posture information based on its business or regulatory policy. In those situations, a provider may require full authentication and authorization of the requestor (or set of requestors) and share only the authorized information to the authenticated and authorized requestors.

Beyond distrusting the network service provider, a requestor must also take into account that the information received from the provider may have been communicated through an undetermined network communications system. That is, the posture information may have traversed through many devices before reaching the requestor. SACM specifications should provide the means for verifying data origin and data integrity and, at minimum, provide endpoint authentication and transfer integrity.

A requestor may require data freshness indications, both knowledge of data origination as well as time of publication, so that it can make more informed decisions about the relevance of the data based on its currency and/or age.

It is also important to note that endpoint assessment reports, especially as they may be provided by the target endpoint, may pose untrustworthy information. The considerations for this are described in Section 8 of [RFC5209].

The trustworthiness of the posture information given by the provider to one or many requestors is dependent on several considerations. Some of these include the requestor requiring:

- o Full disclosure of the network topology path to the provider(s).
- o Direct (peer-to-peer) communication with the provider.
- o Authentication and authorization of the provider.
- o Either or both confidentiality and integrity at the transfer layer.
- o Either or both confidentiality and integrity at the data layer.

4.2. Privacy Considerations

SACM information may contain sensitive information about the target endpoint as well as revealing identity information of the producer or consumer of such information. Similarly, as part of the SACM discovery mechanism, the capabilities and roles (e.g., SACM components enabled) advertised by the endpoint may be construed as private information.

In addition to identity and SACM capabilities information disclosure, the use of timestamps (or other attributes that can be used as identifiers) could be further used to determine a target endpoint or user's behavioral patterns. Such attributes may also be deemed sensitive and may require further protection or obfuscation to meet privacy concerns. That is, there may be applications as well as business and regulatory practices that require that aspects of such information be hidden from any parties that do not need to know it.

Data confidentiality can provide some level of privacy but may fall short where unnecessary data is still transmitted. In those cases, filtering requirements at the data model such as OP-005 must be applied to ensure that such data is not disclosed. [RFC6973]

provides guidelines that SACM protocols, information models, and data models should follow.

5. References

5.1. Normative References

[RFC7632] Waltermire, D. and D. Harrington, "Endpoint Security Posture Assessment: Enterprise Use Cases", RFC 7632, DOI 10.17487/RFC7632, September 2015, <<https://www.rfc-editor.org/info/rfc7632>>.

5.2. Informative References

[RFC5209] Sangster, P., Khosravi, H., Mani, M., Narayan, K., and J. Tardo, "Network Endpoint Assessment (NEA): Overview and Requirements", RFC 5209, DOI 10.17487/RFC5209, June 2008, <<https://www.rfc-editor.org/info/rfc5209>>.

[RFC6876] Sangster, P., Cam-Winget, N., and J. Salowey, "A Posture Transport Protocol over TLS (PT-TLS)", RFC 6876, DOI 10.17487/RFC6876, February 2013, <<https://www.rfc-editor.org/info/rfc6876>>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.

[RFC7171] Cam-Winget, N. and P. Sangster, "PT-EAP: Posture Transport (PT) Protocol for Extensible Authentication Protocol (EAP) Tunnel Methods", RFC 7171, DOI 10.17487/RFC7171, May 2014, <<https://www.rfc-editor.org/info/rfc7171>>.

[TERMS] Birkholz, H., Lu, J., Strassner, J., and N. Cam-Winget, "Security Automation and Continuous Monitoring (SACM) Terminology", Work in Progress, draft-ietf-sacm-terminology-13, July 2017.

Acknowledgments

The authors would like to thank Barbara Fraser, Jim Bieda, and Adam Montville for reviewing and contributing to this document. In addition, we recognize valuable comments and suggestions made by Jim Schaad and Chris Inacio.

Authors' Addresses

Nancy Cam-Winget
Cisco Systems
3550 Cisco Way
San Jose, CA 95134
United States of America

Email: ncamwing@cisco.com

Lisa Lorenzin
Pulse Secure
2700 Zanker Rd., Suite 200
San Jose, CA 95134
United States of America

Email: llorenzin-ietf@1000plus.com