        Secure Automation and Continuous Monitoring (SACM) Requirements
                      draft-ietf-sacm-requirements-10

Abstract

   This document defines the scope and set of requirements for the
   Secure Automation and Continuous Monitoring (SACM) architecture, data
   model and transport protocols. The requirements and scope are based
   on the agreed upon use cases.

Status of This Memo

   This Internet-Draft is submitted in full conformance with the
   provisions of BCP 78 and BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF). Note that other groups may also distribute
   working documents as Internet-Drafts. The list of current Internet-
   Drafts is at http://datatracker.ietf.org/drafts/current/.

   Internet-Drafts are draft documents valid for a maximum of six months
   and may be updated, replaced, or obsoleted by other documents at any
   time. It is inappropriate to use Internet-Drafts as reference
   material or to cite them other than as "work in progress."

   This Internet-Draft will expire on April 3, 2016.

Table of Contents

1. Introduction

   Today's environment of rapidly-evolving security threats highlights
   the need to automate the sharing of such information while protecting
   user information as well as the systems that store, process, and
   transmit this information. Security threats can be detected in a
   number of ways. SACM's 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
   [I-D.ietf-sacm-use-cases]. Also, this document uses terminology
   defined in [I-D.ietf-sacm-terminology].

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].

   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
   defined in [I-D.ietf-sacm-use-cases]. This section describes the
   requirements used by SACM to assess and compare candidate data
   models, interfaces, and protocols, to suit the SACM architecture.
   These requirements express characteristics or features that a
   candidate protocol or data model must be capable of offering to
   ensure security and interoperability.

   Multiple data models, protocols, and transports may be employed in a
   SACM environment. A SACM transport protocol is one that runs on top
   of L3 protocols such as TCP/IP or L4 protocols such as HTTP, carries
   operations (requests / responses), and moves data.

   SACM defines an architecture and information model focused on
   addressing the needs for determining, sharing, and using posture
   information via Posture Information Providers and Posture Information
   Consumers mediated by a Controller. With the information model
   defining assets and attributes to facilitate the guidance,
   collection, and assessment of posture, these are some of the tasks
   that should be considered:

   1. Asset Classification: Map 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. This is what we want to
       know about a target endpoint. For instance, organizations will
       want to know what software is installed and its many 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 those values which are good as well as those

values which 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 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 use cases defined in [I-D.ietf-sacm-use-cases]. 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 transports defined by SACM MUST be designed to allow
 support for future extensions, including both standard and
 proprietary transport protocols and data models.

1. The information model and interfaces MUST support the ability to
    add new operations while maintaining backwards compatibility.
    SACM-defined transport protocols MUST have extensibility to
    allow them to transport 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 to follow; 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 MUST be
 scalable unless they are specifically defined to apply to a special-
 purpose scenario, such as constrained devices. A SACM transport
 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 datagrams: 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 datagram or a very large set of
    assessments (up to multiple gigabytes in size).

 * 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 Agility: 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.

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. The set of attributes defined by the information
 model MUST be well defined.

G-006 Data Integrity: To protect the information being shared, SACM
 components MUST protect the integrity and confidentiality of data in
 transit (while information is being transferred between providers
 and consumers, and through proxies and/or repositories) and data at
 rest (for information stored on repositories and on providers /
 consumers). Mechanisms for this protection are unspecified but
 should include industry best practices such as encrypted storage,
 encrypted transports, data checksums, etc. 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 transport protocols, data models, specific attributes within data models, standard attribute expression sets, etc.) MUST be supported. Negotiation for both versioning and capability 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, providing a method 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. Use Case 2.1.2 describes the need to discover endpoints and their composition.

G-011 Push and Pull Access: Three methods of data access MUST be supported: the standard Pull model as well as solicited and unsolicited Push models. 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 use cases 2.1.3, 2.1.4 and 2.1.5.

G-012 Device 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 transport protocols, and network transport protocols to enable SACM components to communicate. [Revised per our 6/29 conversation - does this work? -LL]

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 or 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

At the simplest abstraction, the SACM architecture represents the
core components and interfaces needed to perform the production and
consumption of posture assessment information. Requirements relating
to SACM's architecture include:

ARCH-001 Scalability: The architectural components MUST account for
  a range of deployments, from very small sets of endpoints to very
  large deployments.

ARCH-002 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-003 Separation of Data and Management Functions: SACM MUST
  define both the configuration and management of the SACM data models
  and protocols used to transport and share posture assessment
  information.

ARCH-004 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.

ARCH-005 Capability Negotiation: Announcement and negotiation of
 functional capabilities (such as authentication protocols,
 authorization schemes, data models, transport protocols, etc.) must
 be supported, enabling a SACM component to make inquiries about the
 capabilities of other components in the SACM ecosystem.

ARCH-006 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-007 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-008 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 (unsolicited) or pull (solicited) 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 (and network layer transport).
    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 uniquely. 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 should the
    data 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 sub-components) for ease of
    maintenance and extension. For example, endpoint identification
    could be an independent sub-component 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 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: The data model MUST contain a data model
 element for each information model element (e.g. endpoint, IP
 address, asset). In other words, for every item in the information
 model, there must be an item in the data model. The data model can
 also include elements that do not exist in the information model.

DM-002 Data Model Structure: The data model can be structured either
 as one single module or separated into modules and sub-modules 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. 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 Provider Identification: The interfaces and actions in the
 data model MUST include the ability to identify data from a specific
 provider. For example, a SACM consumer should be able to request
 all data to come from a specific provider (e.g. Provider A) as
 there can be a larger set of providers.

DM-007 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 as well as whether the attributes are
 conditional, optional, or mandatory.

DM-008 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-009 Data Origin: The data model MUST include the ability for
 consumers to identify the data origin (provider that collected the
 data).

DM-010 Origination Time: The data model SHOULD allow the provider to
 include the information's origination time.

DM-011 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-012 Data Source: The data model MUST allow the provider to
 include attributes defining the data source (target endpoint from
 which the data was collected) - e.g. hostname, domain (DNS) name or
 application name.

DM-013 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-014 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-015 Attribute Extensibility: Use Cases in the whole of Section 2
 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.

DM-016 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 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-017 Transport Agnostic: The data model MUST be transport
agnostic, to allow for the data operations to leverage the most
appropriate SACM transport 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 transport
requirements in that operations requirements are requirements on the
application performing requests and responses, whereas transport
requirements are requirements on the transport protocol carrying the
requests / 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 directly be 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 which 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. Use Case
2.1.2 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. Use Case 2.1.3 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 Data Model Scalability: The operation resulting from a query
operation MUST be able to handle the return and receipt of large
amounts of data. Use Cases 2.1.4 and 2.1.5 describes 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 other SACM components can determine whether they are
constructing an equivalent target endpoint (and their 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.

[2.6](#). Requirements for SACM Transport Protocols

The term transport protocol is frequently overloaded. The term SACM
transport protocol is intended to be distinguished from underlying
layer 3 and 4 protocols such as TCP/IP and TLS. However, it is
possible that a layer 3 or 4 protocol may be used as a SACM transport
protocol, either alone or as part of a SACM transport protocol (i.e.
using HTTP over TLS with XML as the content). The SACM transport
protocol is focused on moving data and performing necessary access
control operations; it is agnostic to the data model operations.

The requirements for SACM transport protocols include:

T-001 Multiple Transport Protocol Support: SACM transport protocols
MUST support different network transport protocols in a deployment
to support different transport layer requirements, different device
capabilities, and system configurations dealing with connectivity.

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

   T-003 Data Confidentiality: SACM transport protocols MUST be able to
    support data confidentiality. SACM transport 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 transport MUST be able to support data
    confidentiality, implementations MAY choose to make confidentiality
    optional. Protection for data at rest is not in scope for transport
    protocols. Data protection MAY be used for both privacy and non-
    privacy scenarios.

   T-004 Transport Protection: SACM transport protocols MUST be capable
    of supporting mutual authentication and replay protection.

   T-005 Transport Reliability: SACM transport protocols MUST provide
    reliable delivery of data. This includes the ability to perform
    fragmentation and reassembly, and to detect replays.

   T-006 Transport Layer Requirements: Each SACM transport 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. For example, a proxy server between UDP and TCP can
    provide a connection that correctly fulfills the connectivity and
    replay requirements as well as data link requirements (through the
    use of TLS and DTLS) but would be unable to provide a channel
    binding requirement, as that implies there is no MITM to look at the
    data.

[3](). **Acknowledgements**

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

[4](). **IANA Considerations**

   This memo includes no request to IANA.

5. Security Considerations

   This document defines the requirements for SACM. As such, it is
   expected that several data models, protocols, and transports may be
   defined or reused from already existing standards. This section will
   highlight security considerations that may apply to SACM based on the
   architecture and standards applied in SACM. In particular,
   highlights to security considerations that may apply to the SACM
   reference architecture and standard data models and transports will
   be discussed.

   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 / 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 sharing of posture information for
   improving 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.

5.1. 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 either at the provider, the
   requesting consumer, or devices that are involved in the transport
   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 transport 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.

It is common to, in general, 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 transport 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.

A requestor beyond distrusting the network service provider, must also 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 transport 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 is 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 transport
   layer.

o Either or both confidentiality and integrity at the data layer.

5.2. 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 advertised capabilities (and roles, e.g.
SACM components enabled) by the endpoint may be construed as private
information. There may be applications as well as business and
regulatory practicess 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 unecessary 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 for which SACM protocols and information and data
models should follow.

6. Change Log

6.1. -05 to -06

Updated G-005 to clarify the MUST to allow non-standard extensions,
SHOULD avoid collisions and ensure interoperability.

Cleaned up and clarified IM-003, DM-001.

Cleaned up some of the OP-XXX and ARCH-XXX per Jim Schaad's comments.

Updated some of the text around Editor notes and removed all 'Editor
Note' comments

7. References

7.1. Normative References

[I-D.ietf-sacm-terminology]
          Birkholz, H., "Secure Automation and Continuous Monitoring
          (SACM) Terminology", draft-ietf-sacm-terminology-07 (work
          in progress), July 2015.

   [I-D.ietf-sacm-use-cases]
              Waltermire, D. and D. Harrington, "Endpoint Security
              Posture Assessment - Enterprise Use Cases", draft-ietf-
              sacm-use-cases-10 (work in progress), July 2015.

   [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
              Requirement Levels", BCP 14, RFC 2119,
              DOI 10.17487/RFC2119, March 1997,
              <http://www.rfc-editor.org/info/rfc2119>.

   [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,
              <http://www.rfc-editor.org/info/rfc5209>.

7.2. Informative References

   [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,
              <http://www.rfc-editor.org/info/rfc6973>.

Authors' Addresses

   Nancy Cam-Winget
   Cisco Systems
   3550 Cisco Way
   San Jose, CA 95134
   US

   Email: ncamwing@cisco.com

   Lisa Lorenzin
   Pulse Secure
   2700 Zanker Rd., Suite 200
   San Jose, CA 95134
   US

   Email: llorenzin@pulsesecure.net