

SACM Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 22, 2016

H. Birkholz
Fraunhofer SIT
N. Cam-Winget
Cisco Systems
March 21, 2016

SACM Information Model
draft-cam-winget-sacm-information-model-00

Abstract

This document defines the data types and data relations and operations that comprise the information model for Security Automation and Continuous Monitoring (SACM) of posture information. This information model is maintained as the IANA "SACM Information Elements" registry. This document defines the initial set and contents to address SACM's use cases (RFC7632).

Please help this paragraph becoming an abstract.

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 September 22, 2016.

Copyright Notice

Copyright (c) 2016 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 (<http://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
2. Requirements notation	3
3. Information Elements (IE)	3
4. Structure of Information Elements	3
4.1. Atomic Information Elements (AIE)	4
4.2. Composite Information Elements (CIE)	4
4.3. SACM Statements	4
4.4. SACM Content Elements	5
4.5. Relationship Types	5
4.6. Events	6
5. Information Element Vocabulary	6
5.1. Vocabulary of Categories	7
5.2. Vocabulary of Atomic Information Elements	7
5.3. Vocabulary of Composite Information Elements	20
6. Example composition of SACM statements	29
7. IANA considerations	31
8. Security Considerations	31
9. Acknowledgements	31
10. Change Log	31
11. Contributors	31
12. References	31
12.1. Normative References	31
12.2. Informative References	32
Authors' Addresses	32

1. Introduction

The purpose of the SACM Information Model (IM) is to ensure interoperability between SACM data models that are used as transport encoding and to provide a base set of information elements and operations that may be exposed or shared between SACM components. A complete set of requirements imposed on the IM can be found in [I-D.ietf-sacm-requirements]. The SACM IM leverages existing definitions of information elements and references the sources in the corresponding descriptions so as to minimize re-invention and duplication. The SACM IM itself is intended to be used for data exchange between SACM components (data in motion). Nevertheless, the Information Elements (IEs) defined in this document can be leveraged to create and align corresponding data models for data at rest.

2. Requirements notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119, BCP 14 [RFC2119].

3. Information Elements (IE)

Every type or group of information, e.g. the information elements, defined in this document represent content transported by a SACM component and are associated with a unique label: their name. This document defines that set of IEs standardized by SACM. A SACM data model MAY include additional IEs that are not defined in this document. The labels of additional IEs included in different SACM data models MUST NOT conflict with the labels of the IEs defined by this information model, and the names of additional IEs MUST NOT conflict with each other or across multiple data models. In order to avoid naming conflicts, the labels of additional IEs SHOULD be prefixed to avoid collision across extensions. The prefix MUST include an organizational identifier and therefore, for example, MAY be an IANA enterprise number, a (partial) name space URI or an organization name abbreviation.

4. Structure of Information Elements

The IEs defined in this document are differentiated into two basic types of Information Elements:

- o Atomic Information Elements: an atomic IE is the simplest IE structure comprised of a single attribute value pairing (atomic IEs are listed in Section 5.2).
- o Composite Information Elements: a composite IE is a richer structure that can be comprised of one or more attribute value pairings (composite IE are listed in Section 5.3).

To associate metadata (e.g. an observation time stamp) with an atomic information element is the equivalent of creating a composite information element that includes the initial atomic information element and an additional information element that represents the time stamp. The resulting composite information element is associated with its own unique name.

Four general structures are expressed via the the two basic types of IE and are used throughout the information model:

- o SACM statements

- o SACM content elements
- o Relationship Types
- o Events

4.1. Atomic Information Elements (AIE)

Atomic IEs represent the smallest building blocks for SACM content, including, for example, a SACM endpoint attribute, a policy entry, a configuration item, an expected states, or a threshold value. AIE can be bundled into composite IE. The set of AIEs defined by the SACM IM is described in section Section 5.2.

In essence, AIEs are attribute value pairs that constitute the "leaves" in a SACM semantic structure. While the SACM IM sometimes does elaborate on the structure of values (e.g. an IPv6 address is an octet string with a maximum length of 16 that may be collapsed in certain conditions), it does not prescribe specific types used in the data model representation (e.g. an unbounded character string).

Every AIE is registered as an corresponding entry at the IANA registry. The Integer Index of the IANA SMI number tables can be used by SACM data models.

4.2. Composite Information Elements (CIE)

Composite IEs constitute bundles of atomic AIEs and/or composite IEs. A CIE represents a specific set of related information that share a semantic relationship, e.g. a SACM statement metadata or state information about a network interface. The set of CIEs defined by the SACM IM is described in section Section 5.3. In essence, CIEs are a "named container" construct that can be used to compose additional CIEs that go beyond the ones standardized by the SACM information model.

The SACM IM allows for recursive or circular nesting of composite IEs. A SACM data Model (DM) MUST include the "default-depth" base AIE that is part of the SACM content metadata.

4.3. SACM Statements

The data exchanged between SACM components is always embedded in a SACM statement. SACM Statements contain one or more CIEs and/or AIEs. A SACM statement functions as an "envelope" type that is associated with metadata about the providing SACM component. The SACM statement metadata can be used to resolve conflicting

information, retrace the provenance of information or to locate archived information in data repositories.

Examples of SACM statement metadata information elements:

- o SACM Domain Identifier: a globally unique identifier that enables the differentiation of SACM statements across SACM domains.
- o Data Origin: the SACM domain unique identifier associated with a SACM component.
- o Statement Identifier: an identifier that enables to uniquely reference this specific statement.

SACM statements are comprised of one or more CIEs; Section 6 provides examples for constructing SACM statements.

4.4. SACM Content Elements

SACM Content Elements are categorized CIEs. The content elements can be composed of one or more AIEs and/or CIEs or it can be another representation that is embedded in the statement, for example, an IPFIX Template Record. Each SACM content element has its own Content Metadata associated with it (analogously to the way that each SACM statement has metadata associated to it). Content element metadata include information about its type, data source (the result produced by a collector) or data origin (the result produced by most other SACM components).

Examples of SACM content element metadata information elements:

- o Target Endpoint Label: an identifier that enables to distinctly identify a target endpoint as a SACM content element.
- o Relationship Identifier(s): a set of semantic relationships that associate this SACM content element with other SACM content elements via their content element identifier.
- o Content Element Identifier: an identifier that enables to uniquely reference this specific content element.

SACM content elements are described in section FIXME.

4.5. Relationship Types

Relationships are expressed via AIE contained within a CIE. There are two ways SACM content elements are associated with each other. "A Flow" associated with "A User", for example, would be a typical

case, in which two separate SACM content elements could be associated with each other.

One way is to include the Relationships AIE in the content element metadata that precludes the actual content (in this example, the content element metadata of the flow record). Relationship Types are uni-directional. For example, the "is-associated-with-user" Relationship AIE included in the content element metadata points to a specific user via a corresponding content element identifier.

The alternative way is to include the reference of associated information directly into the content of the content element. A session CIE, for instance, could refer to a specific user by including identifying attributes about that user. While this is a valid way of creating a relationship between different kinds of content, it requires careful matching or the introduction of another appropriate identifier mechanism (that does not conflict with other SACM statements and SACM content element identifiers). If a SACM data model allows for transport of other representations as payload of a content element (e.g. a pcap fragment containing suspicious packets, for example), there might be no alternative as to use the content element metadata to include relationships to other content elements.

4.6. Events

Events are a specific type of CIE that are always associated with a time stamp and represent a change of state or configuration that can be expressed as a SACM content. The time an event was published by a SACM component is recorded in its corresponding SACM statement metadata, the time it was created (or initially observed) is recorded in its content element metadata. It is also recorded in the CIE itself, which is somewhat redundant but can improve performance in some scenarios. Event CIE can also include the past state or configuration before the change occurred, or - if applicable - a threshold or trigger condition that lead to the creation of the event.

5. Information Element Vocabulary

The vocabulary of Information Element names standardized by the SACM IM does not prescribe the use of these exact same names in every SACM data model. If terms diverge, a mapping has to be provided in the corresponding SACM data model document.

A subset of the names of the information elements defined in this document are appended with "-type". This indicates that the IM defines a set of values for these information elements (e.g. the

interface types defined by the IANA registry or the relationship types).

5.1. Vocabulary of Categories

Categories are special Information Elements that enable to refer to multiple types of IEs via just one name. Therefore, they are similar to a type-choice. A prominent example of a category is network-address. Network-address is a category that every kind of network address is associated with, e.g. mac-address, ipv4-address, ipv6-address, or typed-network-address. If a CIE includes network-address as one of its components, any of that categories members is valid to be used in its stead.

Another prominent example is EndpointIdentifier. Some IEs can be used to identify (and over time re-recognize) target endpoints - those are associated with the category endpoint-identifier.

content: this is a very broad category. Content is the payload of a content element in a SACM statement. Formally, metadata is the complement to content and everything that is not part of SACM statement metadata or content element metadata is therefore considered to be content. Every IE can be content (although the same type of IE can be used in the metadata at the same time - and those would not be content as described before). Annotating every IE with this category would be highly redundant and is therefore omitted for brevity.

network-address: (work-in-progress)

 ipv4-address

 ipv6-address

 mac-address

endpoint-identifier: (work-in-progress)

software-component: (work-in-progress)

software-label: (work-in-progress)

5.2. Vocabulary of Atomic Information Elements

The content of every Atomic Information Element is expressed in a single value. Note that while this section lists AIEs, some of them may also be represented as a CIE (especially if metadata is used).

access-privilege-type: a set of types that represents access privileges (e.g. read, write, none)

References: none

account-name: a label that uniquely identifies an account that can require some form of (user) authentication to access

References: none

administrative-domain: a label the is supposed to uniquely identify an administrative domain

References [IFMAP]

address-association-type: a set of types that defines the type of address associations (e.g. broadcast-domain-member-list, ip-subnet-member-list, ip-mac, shared-backhaul-interface, etc.)

References: none

address-mask-value: a value that expresses a generic address subnetting bitmask

address-type: a set of types that specifies the type of address that is expressed in an address CIE (e.g. ethernet, modbus, zigbee)

References: none

address-value: a value that expresses a generic network address

References: none

Category: network-address

application-component: a label that references a "sub"-application that is part of the application (e.g. an add-on, a chiper-suite, a library)

References: [SWID]

Category: software-component

application-label: a label that is supposed to uniquely reference an application

References: [SWID]

Category: software-label

application-type: a set of types (FIXME maybe a finite set is not realistic here - value not enumerator?) that identifies the type of (user-space) application (e.g. text-editor, policy-editor, service-client, service-server, calender, rouge-like RPG)

References: [SWID]

Category: software-type

application-manufacturer: the name of the vendor that created the application

References: [SWID]

Category: software-manufacturer

application-name: a value that represents the name of an application given by the manufacturer

References: [SWID]

application-version: a version string that identifies a specific version of an application

References: [SWID]

Category: software-version

authenticator: a label that references a SACM component that can authenticate target endpoints (can be used in a target-endpoint CIE to express that the target endpoint was authenticated by that SACM component)

References: none

attribute-name: a value that can express the attribute name of generic Attribute-Value-Pair CIE

References: none

attribute-value: a value that can express the attribute value of generic Attribute-Value-Pair CIE

References: none

authentication-type: a set of types that expresses which type of authentication was used to enable a network interaction/connection

References: [PXGRID]

birthdate: a label for the registered day of birth of a natural person (e.g. the date of birth of a person as an ISO date string <http://rs.tdwg.org/ontology/voc/Person#birthdate>)

References: [SCAP-AI]

bytes-received: a value that represents a number of octets received on a network interface

Reference : [PXGRID]

bytes-sent: a value that represents a number of octets sent on a network interface

Reference : [PXGRID]

certificate: a value that expresses a certificate that can be collected from a target endpoint

References: none

Category: endpoint-identifier

collection-task-type: a set of types that defines how collected SACM content was acquired (e.g. network-observation, remote-acquisition, self-reported)

Reference: none

confidence: a representation of the subjective probability that the assessed value is correct. If no confidence value is given it is assumed that the confidence is 1 (limits confidence values to the range between zero and one)

References: [ARF]

content-action: a set of types that expresses a type of action (e.g. add, delete, update). Can be associated, for instance, with an event CIE or with an network observation

References: [ARF]

content-elements: a value that represents the number of content-elements included in a SACM statement

References: none

content-topic: a set of types that defines what kind of concept the information is included in a content element (e.g. Session, User, Interface, PostureProfile, Flow, PostureAssessment, TargetEndpoint)

References: none

content-type: a set of types that defines what kind of information is included in a content element (e.g. EndpointConfiguration, EndpointState, DirectoryEntry, Event, Incident)

References: none

country-code: a set of types according to ISO 3166-1 trigraphic codes of countries

References: FIXME

data-origin: a label that uniquely identifies a SACM component in and across SACM domains

References: none

Aliases: sacm-component-id

data-source: a label that is supposed to uniquely identify the data source (e.g. a target endpoint or sensor) that provided an initial endpoint attribute record

References: [ARF]

Aliases: te-id (work-in-progress)

decimal-fraction-denominator: a denominator value to express a decimal fraction time stamp (e.g. in timestamp)

References: none

decimal-fraction-numerator: a numerator value to express a decimal fraction time stamp (e.g. in timestamp)

default-depth: a value that expresses how often a circular reference of CIE is allowed to repeat, or how deep a recursive nesting may occur, respectively.

References: none

discoverer: a label that refers to the SACM component that discovered a target endpoint (can be used in a target-endpoint CIE to express, for example, that the target endpoint was authenticated by that SACM component)

References: none

email-address: a value that expresses an email-address

References: none

event-type: a set of types that define the categories of an event (e.g. access-level-change, change-of-privilege, change-of-authorization, environmental-event, or provisioning-event)

Reference: none

event-threshold: if applicable, a value that can be included in an event CIE to indicate what numeric threshold value was crossed to trigger that event

Reference: none

event-threshold-name: if an event is created due to a crossed threshold, the threshold might have a name associated with it that can be expressed via this value

References: none

event-trigger: this value is used to express more complex trigger conditions that may cause the creation of an event.

firmware-id: a label that represents the BIOS or firmware ID of a specific target endpoint

Reference: none

Category: endpoint-identifier

hardware-serial-number: a value that identifies a piece of hardware that is a component of a composite target endpoint (in essence,

every target endpoint is a composite) and can be acquired from a target endpoint by a collection task

Reference: none

Category: endpoint-identifier

host-name: a label typically associated with an endpoint but not always intended to be unique in a given scope

References [ARF], [SCAP-AI]

Category: endpoint-identifier

interface-label: a unique label a network interface can be referenced with

Reference: none

ipv6-address-subnet-mask-cidrnot: an IPv6 subnet bit mask in CIDR notation

References: TBD

ipv6-address-value: an IPv4 address value

References: TBD

Category: endpoint-identifier, network-address

ipv4-address-subnet-mask-cidrnot: an IPv4 subnet bit mask in CIDR notation

References: TBD

ipv4-address-subnet-mask: an IPv4 subnet mask

References: TBD

ipv4-address-value: an IPv4 address value

References: TBD

Category: endpoint-identifier, network-address

layer2-interface-type: a set of types referenced by IANA ifType

References: [RFC3635], [RFC2863]

layer4-port-address: a layer 4 port address (typically used, for example, with TCP and UDP)

References: none

Category: network-address

layer4-protocol: a set of types that express a layer 4 protocol (e.g. UDP or TCP)

location-name: a value that represents a named region of space FIXME

References: [IFMAP], [ARF], [SCAP-AI]

mac-address: a value that expresses an Ethernet address

References: [IFMAP], [ARF], [SCAP-AI]

Category: endpoint-identifier, network-address

method-label: a label that references a specific method registered and used in a SACM domain (e.g. method to match and re-identify target endpoints via identifying attributes)

References: none

method-repository: a label that references a SACM component methods can be registered at and that can provide guidance in the form of registered methods to other SACM components

References: none

network-access-level-type: a set of types that expresses categories of network access-levels (e.g. block, quarantine, etc.)

References: [IFMAP]

network-id: most networks, such as AS, an OSBF domains, or vlans, can have an ID that is represented via this AIE

References: none

network-interface-name: a label that uniquely identifies an interface associated with a distinguishable endpoint

References: FIXME

network-layer: a set of layers that express the specific network layer an interface operate on (typically layer 2-4)

References: FIXME

network-name: a label that is associated with a network. Some networks, for example effective layer2-broadcast-domains, are difficult to "grasp" and therefore quite complicated to name

References: none

organization-id: a label that is supposed to uniquely identify an organization

References: [ARF]

organization-name: a value that represents the name of an organization

References: [ARF]

os-component: a label that references a "sub-component" that is part of the operating system (e.g. a kernel module, microcode, or ACPI table)

References: [SWID]

Category: software-component

os-label: a label that references a specific version of an operating system, including patches and hotfixes

References: [SWID]

Category: software-label

os-manufacturer: the name of the manufacturer of an operating system

References: [IFMAP]

Category: software-manufacturer

os-name: the name of an operating system

References: [IFMAP]

Category: software-name

os-type: a set of types that identifies the type of an operating system (e.g. real-time, security-enhanced, consumer, server)

References: none

Category: software-type

os-version: a value that represents the version of an operating-system

Category: software-version

patch-id: a label the uniquely identifies a specific software patch

References: [ARF]

patch-name: the vendor's name of a software patch

References: [ARF], [SWID]

person-first-name: the first name of a natural person

References: [ARF], [SCAP-AI]

person-last-name: the last name of a natural person

References: [ARF], [SCAP-AI]

person-middle-name: the first name of a natural person

References: [ARF], [SCAP-AI]

phone-number: a label that expresses the u.s. national phone number (e.g. pattern value="((\d{3}))?\d{3}-\d{4}")

References: [ARF], [SCAP-AI]

phone-number-type: a set of types that express the type of a phone number (e.g. DSN, Fax, Home, Mobile, Pager, Secure, Unsecure, Work, Other)

References: [ARF]

privilege-name: the attribute-name of the privilege represented as an AVP

References: none

privilege-value: the value-content of the privilege represented as an AVP

References: none

protocol: a set of types that defines specific protocols above layer 4 (e.g. http, https, dns, ipp, or unknown)

References: none

public-key: the value of a public key (regardless of its method of creation, crypto-system, or signature scheme) that can be collected from a target endpoint

Reference: none

Category: endpoint-identifier

relationship-content-element-guid: a reference to a specific content element used in a relationship CIE

References: none

relationship-statement-guid: a reference to a specific SACM statement used in a relationship CIE

References: none

relationship-object-label: a reference to a specific label used in content (e.g. a te-label or a user-id). This reference is typically used if matching content AIE can be done efficiently and can also be included in addition to a relationship-content-element-guid reference.

References: none

relationship-type: a set of types that is in every instance of a relationship CIE to highlight what kind of relationship exists between the CIE the relationship is included in (e.g. associated_with_user, applies_to_session, seen_on_interface, associated_with_flow, contains_virtual_device)

References: none

role-name: a label that references a collection of privileges assigned to a specific entity (identity? FIXME)

References: FIXME

session-state-type: a set of types a discernible session (an ongoing network interaction) can be in (e.g. Authenticating, Authenticated, Postured, Started, Disconnected)

References: [PXGRID]

statement-guid: a label that expresses a global unique ID referencing a specific SACM statement that was produced by a SACM component

References: none

statement-type: a set of types that define the type of content that is included in a SACM statement (e.g. Observation, DirectoryContent, Correlation, Assessment, Guidance)

References: none

status: a set of types that defines possible result values for a finding in general (e.g. true, false, error, unknown, not applicable, not evaluated)

References: [ARF]

sub-administrative-domain: a label for related child domains an administrative domain can be composed of (used in the CIE administrative-domain)

References: none

sub-interface-label: a unique label a sub network interface (e.g. a tagged vlan on a trunk) can be referenced with

References: none

super-administrative-domain: a label for related parent domains an administrative domain is part of (used in the CIE administrative-domain)

References: none

super-interface-label: a unique label a super network interface (e.g. a physical interface a tunnel interface terminates on) can be referenced with

References: none

te-assessment-state: a set of types that defines the state of assessment of a target-endpoint (e.g. in-discovery, discovered, in-classification, classified, in-assessment, assessed)

References: [ARF]

te-label: an identifying label created from a set of identifying attributes used to reference a specific target endpoint

References: none

te-id: an identifying label that is created randomly, is supposed to be unique, and used to reference a specific target endpoint

References: [ARF], [SWID]

Aliases: data-source

timestamp: a timestamp that expresses a specific point in time

References: [IFMAP], [ARF]

timestamp-type: a set of types that express what type of action or event happened at that point of time (e.g. discovered, classified, collected, published). Can be included in a generic timestamp CIE

References: none

units-received: a value that represents a number of units (e.g. frames, packets, cells or segments) received on a network interface

Reference : [PXGRID]

units-sent: a value that represents a number of units (e.g. frames, packets, cells or segments) sent on a network interface

Reference : [PXGRID]

username: a part of the credentials required to access an account that can be collected from a target endpoint

References: none

Category: endpoint-identifier

user-directory: a label that identifies a specific type of user-directory (e.g. ldap, active-directory, local-user)

Reference: [PXGRID]

user-id: a label that references a specific user known in a SACM domain

References: [PXGRID]

web-site: a URI that references a web-site

References: [ARF]

WGS84-longitude: a label that represents WGS 84 rev 2004 longitude

References: [SCAP-AI]

WGS84-latitude: a label that represents WGS 84 rev 2004 latitude

References: [SCAP-AI]

WGS84-altitude: a label that represents WGS 84 rev 2004 altitude

References: [SCAP-AI]

5.3. Vocabulary of Composite Information Elements

The content of every Composite Information Element is expressed by the mandatory and optional IE it can be composed of. The components of an CIE can have a cardinality associated with them:

- o (*): zero to unbounded occurrences
- o (+): one to unbounded occurrences
- o (?): zero or one occurrence
- o (n*m): between n and m occurrences
- o no cardinality: one occurrence

If there is no cardinality highlighted or the cardinality (+) or (n*m) is used, including this IE in the CIE is mandatory. In contrast, optional IE are expressed via the cardinality (?) or (*). An CIE can prescribe a strict sequence to the component IE it contains. This is indicated by an (s).

address-association (s): some addresses are associated with each other, e.g. a mac-address can be associated with a number of IP addresses or a sensor address can be associated with the external

address of its two redundant IP gateways. The first address is the address a number of addresses with the same type is associated with. An address type SHOULD be included and the addresses associated with the first address entry MUST be of the same type.
NANCY FIXME

address

address-type (?)

address (+)

address-type (?)

administrative-domain: this CIE is intended to express more complex setups of interconnected administrative domains

administrative-domain

sub-administrative-domain (*)

super-administrative-domain (?)

location (?)

application: an application is software that is not part of the kernel space (therefore typically runs in the user space. An application can depend on specific running party of an operating system.

application-label (?)

application-name

application-type (*)

application-component (*)

application-manufacturer (?)

application-version (?)

application-instance: a specific instance of an application that is installed on an endpoint. The application-label is used to refer to corresponding information stored in an application CIE

application-label

target-endpoint

attribute-value-pair: a generic CIE that is used to express various AVP (e.g. Radius Attributes)

attribute-name

attribute-value

content-creation-timestamp: a decimal fraction timestamp that specifies the point in time the content element was created by a SACM component

decimal-fraction-denominator

decimal-fraction-numerator

content-element: content produced by a SACM component is encapsulated in content-elements that also include content-metadata regarding that content

content-metadata (+)

content (+)

content-metadata: metadata regarding the content included in a specific content-element. The content the metadata annotates can be initially collected content - in this case a data-source has to be included in the metadata. Content can also be the product of a SACM component (e.g. an evaluator), which requires a data-origin IE instead that references the producer of information.

content-element-guid

content-creation-timestamp

content-topic

content-type

data-source (?)

data-origin (?)

relationship (*)

data-source: a CIE that refers to a target endpoint that is the source of SACM content - either via a label (data-source, which

could also be used without this CIE), or via a list of endpoint-identifiers (category). Both can be included at the same time but MUST NOT conflict.

data-source (?)

endpoint-identifier (*)

dst-flow-element: identifies the destination of a flow. The port number SHOULD be included if the network-address is an IP-address.

network-address

layer4-port-address (?)

ethernet-interface: the only two mandatory component of this CIE is the mac-address and the generated label (to distinguish non-unique addresses). This acknowledges the fact that in many cases this is the only information available about an Ethernet interface. If there is more detail information available it MUST be included to avoid ambiguity and to increase the usefulness for consumer of information. The exception are sub-interface-labels and super-interface-labels, which SHOULD be included.

interface-label

network-interface-name (?)

mac-address

network-name (?)

network-id (?)

layer2-interface-type (?)

sub-interface-label (*)

super-interface-label (*)

event (s): this a special purpose CIE that represents the change of content. As with content-elements basically every content can be included in the two content entries. The mandatory content entry represents the "after" state of the content and the optional content entry can represent the "before" state if available or required.

event-type (?)

event-threshold (?)

event-threshold-name (?)

event-trigger (?)

typed-timestamp

content

content (?)

flow-record: a composite that expresses a single flow and its statistics. If applicable, protocol and layer4-protocol SHOULD be included

src-flow-element

dst-flow-element

protocol (?)

layer4-protocol (?)

flow-statistics

flow-statistics: this CIE aggregates bytes and units send and received

bytes-received

bytes-sent

units-received

units-sent

group: insert text here (work in progress)

ipv4-address: an IPv4 address is always associated with a subnet. This CIE combines these both tightly nit values. Either a subnet mask or a CIDR notation bitmask SHOULD be included.

ipv4-address-value

ipv4-address-subnet-mask-cidrnot (?)

ipv4-address-subnet-mask (?)

ipv6-address: an IPv6 address is always associated with a subnet. This CIE combines these both tightly nit values. A CIDR notation bitmask SHOULD be included.

ipv6-address-value

ipv6-address-subnet-mask-cidrnot (?)

location: a CIE that aggregates potential details about a location

location-name

WGS84-longitude

WGS84-latitude

WGS84-altitude

operation-system: an operation-system is software that is directly interacting with the hardware, provides the runtime environment for the user-space and corresponding interfaces to hardware functions.

os-label (?)

os-name

os-type (*)

os-component (*)

os-manufacturer (?)

os-version (?)

organization: this CIE aggregates information about an organization and can be references via its id

organization-id

organization-name

location (?)

person: a CIE that aggregates the details about a person and combines it with a identifier unique to SACM domains

person-first-name

person-last-name

person-middle-name (*)

phone-contact (*)

email-address (*)

phone-contact: this CIE can be used to reference a phone number and how it functions as a contact

phone-number

phone-number-type (?)

privilege: a CIE to express privileges via a specific name/value pair

privilege-name

privilege-value

relationship: the relationship CIE enables to associate the CIE it is included in with other CIE if they contain a unique identifier or label - providing an alternative to including attributes of other content CIE as a means to map them (which remains a valid alternative, though). The relationship CIE MUST at least reference one relationship object (either a SACM statement identifier)

relationship-type

relationship-content-element-guid (*)

relationship-statement-guid (*)

relationship-object-label (*)

sacm-statement: every SACM component produces information in this format. This CIE can be considered the root IE for every SACM message generated. There MUST be at least one content element included in a SACM statement and if there are more than one, they are ordered in a sequence.

statement-metadata

content-element (+)(s)

session: represents an ongoing network interaction that can be in various states of authentication or assesement

session-state-type

(work-in-progress)

src-flow-element: identifies the source of a flow. The port number SHOULD be included if the network-address is an IP-address.

network-address

layer4-port-address (?)

statement-creation-timestamp: a decimal fraction timestamp that specifies the point in time the SACM statement was created by a SACM component

decimal-fraction-denominator

decimal-fraction-numerator

statement-publish-timestamp: a decimal fraction timestamp that specifies the point in time the SACM component attempted to publish the SACM statement (if successful, this will result in the publish-timestamp send with the SACM statement).

decimal-fraction-denominator

decimal-fraction-numerator

statement-metadata: every SACM statement includes statement metadata about the SACM component it was produced by and a general category that indicates what this statement is about

statement-guid

data-origin

statement-creation-timestamp (?)

statement-publish-timestamp

statement-type

content-elements

target-endpoint: this is a central CIE used in the process chains a SACM domain can compose. Theoretically every kind of information can be associated with a target endpoint CIE via its corresponding content element. A few select IE can be stored in the CIE itself to reduce the overhead of following references that would occur in most scenarios. If the hostname is unknown the value has to be set as an equivalent to "not available" (e.g. NULL). Comment from the authors: This is "work in progress" an a good basis for discussion

host-name

te-label

administrative-domain (?)

application-instance (*)

ethernet-interface (*)

address-association (*)

data-source (?)

operation-system (?)

te-profile: a set of expected states, policies and pieces of guidance that can be matched to a target endpoint (or a class of target endpoints "work in progress")

typed-timestamp: a flexible timestamp CIE that can express the specific type of timestamp via its content. This is an alternative to the "named" timestamps that do not include a timestamp-type

decimal-fraction-denominator

decimal-fraction-numerator

timestamp-type

user: a CIE that references details of a specific user known in a SACM domain active on a specific target endpoint

user-id

username (?)

data-source (?)

user-directory (?)

6. Example composition of SACM statements

This section illustrates how SACM statements can be composed of content information elements, how relationship CIEs can be used in content metadata, and how the categories statement-type, content-topic and content-type are intended to be used.

The SACM statements instances are written in pseudo code. AIE end with a colon. Some AIE include exemplary values to, for example, present how references to guid and labels can be used. For the sake of brevity, not all mandatory IE that are part of a CIE are always included (e.g. as it is the case with target-endpoint).

The example shows three SACM statements that were produced by three different SACM components that overall include four related content elements.

This is (work in progress).

```
sacm statement
  statement-metadata
    statement-guid: example-sguid-one
    data-origin: SACM-component-label-one
    statement-publish-timestamp: exmample-TS-one
    statement-type: Observation
  content-element
    content-metadata
      content-element-guid: example-cguid-one
      content-creation-timestamp:
      content-topic: Flow
      content-type: EndpointState
    relationship
      relationship-type: is-associated-with-user
      relationship-content-object: example-cguid-three
    relationship
      relationship-type: is-associated-with-te
      relationship-content-object: example-cguid-two
    relationship
      relationship-type: is-associated-with-te
      relationship-content-object: example-te-label
  flow-record
    src-flow-element
      network-address (ipv4-address)
      ipv4-address-value:
```

```
        ipv4-address-subnet-mask-cidrnot:
        layer4-port-address: 23111
dst-flow-element
  network-address (IPv4-address)
    ipv4-address-value:
    ipv4-address-subnet-mask-cidrnot:
    layer4-port-address: 22
  protocol: ssh
  layer4-protocol: tcp
  flow-statistics
    bytes-received:
    bytes-sent:
    units-received:
    units-sent:
content-element
  content-metadata
    content-element-guid: example-cguid-two
    content-creation-timestamp:
    content-topic: TargetEndpoint
    content-type: EndpointConfiguration
  target-endpoint
    te-label: example-te-label
    host-name: example-host-name
    ethernet-interface: example-interface

sacm statement
  statement-metadata
    statement-guid: example-sguid-two
    data-origin: SACM-component-label-two
    statement-publish-timestamp: exmample-TS-two
    statement-type: DirectoryContent
  content-element
    content-metadata
      content-element-guid: example-cguid-three
      content-creation-timestamp:
      content-topic: User
      content-type: DirectoryEntry
  user
    user-name: example-username
    user-directory: component-id

sacm statement
  statement-metadata
    statement-guid: example-sguid-three
    data-origin: SACM-component-label-three
    statement-publish-timestamp: exmample-TS-three
    statement-type: Observation
  content-element
```

```
content-metadata
  content-element-guid: example-cguid-four
  content-creation-timestamp:
  content-topic: Priviledges
  content-type: Event
  relationship
    relationship-type: is-associated-with-user
    relationship-content-object: example-cguid-three
event
  event-type: change-of-priviledge
  typed-timestamp
    decimal-fraction-denominator:
    decimal-fraction-numerator:
    timestamp-type: time-of-observation
priviledge
  privilege-name: super-user-escalation
  privilege-value: true
priviledge
  privilege-name: super-user-escalation
  privilege-value: false
```

7. IANA considerations

This document includes requests to IANA.

8. Security Considerations

9. Acknowledgements

10. Change Log

First version -00

11. Contributors

12. References

12.1. Normative References

- [ARF] Corporation., T., "Assessment Results Format", 2010.
- [IFMAP] "TCG Trusted Network Communications - TNC IF-MAP Metadata for Network Security Specification Version 1.1r9", May 2012.

- [PXGRID] Appala, S., Cam-Winget, N., McGrew, D., and J. Verma, "An Actionable Threat Intelligence system using a Publish-Subscribe communications model", ACM Proceedings of the 2nd ACM Workshop on Information Sharing and Collaborative Security, page 61-70, DOI 10.1145/2808128.2808131, ISBN 978-1-4503-3822-6.
- [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>>.
- [RFC2863] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", RFC 2863, DOI 10.17487/RFC2863, June 2000, <<http://www.rfc-editor.org/info/rfc2863>>.
- [RFC3635] Flick, J., "Definitions of Managed Objects for the Ethernet-like Interface Types", RFC 3635, DOI 10.17487/RFC3635, September 2003, <<http://www.rfc-editor.org/info/rfc3635>>.
- [SCAP-AI] Wunder, J., Halbardier, A., and D. Waltermire, "Specification for Asset Identification 1.1", NIST Interagency Report 7693 , 2011.
- [SWID] "Information technology - Software asset management - Part 2: Software identification tag'", ISO/IEC 19770-2:2015, October 2015.

12.2. Informative References

- [I-D.ietf-sacm-requirements]
Cam-Winget, N. and L. Lorenzin, "Security Automation and Continuous Monitoring (SACM) Requirements", draft-ietf-sacm-requirements-13 (work in progress), March 2016.

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
Darmstadt 64295
Germany

Email: henk.birkholz@sit.fraunhofer.de

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

Email: ncamwing@cisco.com

SACM
Internet-Draft
Intended status: Standards Track
Expires: May 4, 2017

C. Coffin
D. Haynes
C. Schmidt
The MITRE Corporation
J. Fitzgerald-McKay
Department of Defense
October 31, 2016

Software Inventory Message and Attributes (SWIMA) for PA-TNC
draft-coffin-sacm-nea-swid-patnc-03

Abstract

This document specifies the Software Inventory Message and Attributes for PA-TNC. It extends the PA-TNC specification [RFC5792] by providing specific attributes and message exchanges to allow endpoints to report their installed software inventory information to a NEA server (as described in [RFC5209]).

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 May 4, 2017.

Copyright Notice

Copyright (c) 2016 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 (<http://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	4
1.1. Network Endpoint Assessment (NEA)	5
1.2. Keywords	7
1.3. Definitions	7
2. Background	8
2.1. Supported Use Cases	8
2.1.1. Use Software Inventory as a Factor in Determining Endpoint Access	8
2.1.2. Maintain a Central Repository Reflecting an Endpoint's Software Inventory	9
2.1.3. PA-TNC Use Cases	10
2.2. Non-supported Use Cases	10
2.3. Specification Requirements	11
2.4. Non-Requirements	12
2.5. Assumptions	12
2.6. Non-Assumptions	13
2.7. Software Inventory Message and Attributes for PA-TNC Diagram Conventions	13
3. Software Inventory Message and Attributes for PA-TNC System Requirements	14
3.1. Basic Attribute Exchange	14
3.2. Core Software Reporting Information	15
3.2.1. Software Identifiers	15
3.2.1.1. Record Identifiers	16
3.2.1.2. Software Locators	17
3.2.1.3. Using Software and Record Identifiers in SW Attributes	18
3.3. Targeted Requests	18
3.4. Monitoring Changes in an Endpoint's Software Inventory Evidence Collection	19
3.5. Reporting Change Events	22
3.5.1. Event Identifiers	22
3.5.2. Core Event Tracking Information	23
3.5.3. Updating Inventory Knowledge Based on Events	24
3.5.4. Using Event Records in SW Attributes	24
3.5.5. Partial and Complete Lists of Event Records in SW Attributes	25
3.5.6. Synchronizing Event Identifiers and Epochs	27
3.6. Subscriptions	28
3.6.1. Establishing Subscriptions	29
3.6.2. Managing Subscriptions	29
3.6.3. Terminating Subscriptions	30

3.6.4.	Subscription Status	30
3.6.5.	Fulfilling Subscriptions	31
3.6.5.1.	Subscriptions Reporting Inventories	32
3.6.5.2.	Subscriptions Reporting Events	32
3.6.5.3.	Targeted Subscriptions	34
3.6.5.4.	No Subscription Consolidation	34
3.6.5.5.	Delayed Subscription Fulfillment	35
3.7.	Multiple Sources of Software Inventory Evidence Records .	35
3.8.	Error Handling	37
4.	Software Inventory Message and Attributes for PA-TNC Protocol	38
4.1.	PA Subtype (AKA PA-TNC Component Type)	38
4.2.	SW Attribute Overview	39
4.3.	SW Attribute Exchanges	41
4.4.	Software Inventory Message and Attributes for PA-TNC Attribute Enumeration	43
4.5.	Normalization of Text Encoding	44
4.6.	Request IDs	44
4.7.	SW Request	45
4.8.	Software Identifier Inventory	49
4.9.	Software Identifier Events	52
4.10.	Software Inventory	57
4.11.	Software Events	60
4.12.	Subscription Status Request	64
4.13.	Subscription Status Response	65
4.14.	PA-TNC Error as Used by Software Inventory Message and Attributes for PA-TNC	67
4.14.1.	SW_ERROR, SW_SUBSCRIPTION_DENIED_ERROR and SW_SUBSCRIPTION_ID_REUSE_ERROR Information	69
4.14.2.	SW_RESPONSE_TOO_LARGE_ERROR Information	70
4.14.3.	SW_SUBSCRIPTION_FULFILLMENT_ERROR Information	72
5.	Supported Data Models	74
5.1.	ISO 2015 SWID Tags using XML	74
5.1.1.	Guidance on Normalizing Source Data to ISO 2015 SWID Tags using XML	74
5.1.2.	Guidance on Creation of Software Identifiers from ISO 2015 SWID Tags	75
5.2.	ISO 2009 SWID Tags using XML	75
5.2.1.	Guidance on Normalizing Source Data to ISO 2015 SWID Tags using XML	75
5.2.2.	Guidance on Creation of Software Identifiers from ISO 2015 SWID Tags	75
6.	Security Considerations	76
6.1.	Evidentiary Value of Software Inventory Evidence Records	76
6.2.	Sensitivity of Collected Records	76
6.3.	Integrity of Endpoint Records	78
6.4.	SW-PC Access Permissions	78
6.5.	Sanitization of Record Fields	79
6.6.	PA-TNC Security Threats	79

7. Privacy Considerations	79
8. Relationship to Other Specifications	79
9. IANA Considerations	80
9.1. Registry for PA-TNC Attribute Types	80
9.2. Registry for PA-TNC Error Codes	81
9.3. Registry for PA Subtypes	82
9.4. Registry for Software Data Models	83
10. References	83
10.1. Normative References	83
10.2. Informative References	84
Authors' Addresses	84

1. Introduction

Possession of a list of an endpoint's installed software is very useful in understanding and maintaining the security state of an enterprise. For example, if an enterprise policy requires the presence of certain pieces of software and/or prohibits the presence of other software, reported software installation inventory lists can be used to indicate compliance or non-compliance with these requirements. Software installation inventories and the patch level of the identified software can be compared to vulnerability or threat alerts to determine an endpoint's exposure to attack. All of these uses make an understanding of an endpoint's installed software inventory highly useful to NEA Servers and other enterprise security applications.

There is a need for a standardized method for exchanging software inventory that carries a software identifier (a unique identifier associated with a specific software product and version thereof). In some cases, it may also be necessary to convey information that characterizes this product (i.e., provides metadata about the product beyond its identifier) as well as observable installation information. These "Messages and Attributes" would enable software identification, installation, and characterization information to be provided for any software installed on any endpoint that supports this specification.

To that end, this specification defines a new set of PA-TNC attributes, carried over PA-TNC messages, which are used to communicate requests for software information and software events, and for conveying that information back to a NEA Server.

This specification is designed only to report software that is installed on an endpoint. In particular, it does not monitor or report information about what software is running on the endpoint. Likewise, it is not intended to report individual files, libraries, installation packages, or similar artifacts. While all of this

information has its uses, this information requires different metadata and different methods of monitoring the endpoint. As a result, this specification focuses solely on installed software, leaving reporting of other classes of endpoint information to other specifications.

1.1. Network Endpoint Assessment (NEA)

The Network Endpoint Assessment (NEA) Working Group defines an open solution architecture that enables network operators to collect and utilize information about endpoint configuration and state. This information can be used to enforce policies, monitor endpoint health, and for many other activities. Information about the software present on an endpoint is an important consideration for such activities. Such information can come from multiple sources, including tag files (such as ISO SWID tags [SWID], reports from third party inventory tools, output from package managers, and other sources. The attributes defined in this document are used to communicate software inventory evidence, collected from a range of possible sources, from the posture collector on the endpoint to the posture validator on a NEA Server using the PA-TNC interface, as shown in Figure 1 below.




Figure 1: NEA Reference Model

Before reading this specification any further, the reader should review and understand the NEA reference architecture as described in the Network Endpoint Assessment (NEA): Overview and Requirements [RFC5209]. The reader should also understand the capabilities and requirements common to PA-TNC interfaces as defined in RFC 5792 [RFC5792].

This document is based on standards published by the Trusted Computing Group's Trusted Network Communications (TNC) workgroup. The TNC and NEA architectures are interoperable and the following components are equivalent:

TNC Component	NEA Component
Integrity Measurement Verifier (IMV)	Posture Validator
Integrity Measurement Collector (IMC)	Posture Collector
TNC Server (TNCS)	Posture Broker Server
TNC Client (TNCC)	Posture Broker Client

Table 1: Equivalent components in TNC and NEA architectures

1.2. Keywords

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this specification are to be interpreted as described in RFC 2119 [RFC2119]. This specification does not distinguish blocks of informative comments and normative requirements. Therefore, for the sake of clarity, note that lower case instances of must, should, etc. do not indicate normative requirements.

1.3. Definitions

This section defines terms with special meaning within this document.

SW-PC - A Posture Collector (PC) that conforms to this specification. Note that such a posture collector might also support other PA-TNC exchanges beyond Software Inventory Message and Attributes for PA-TNC.

SW-PV - A Posture Validator (PV) that conforms to this specification. Note that such a posture verifier might also support other PA-TNC exchanges beyond Software Inventory Message and Attributes for PA-TNC.

SW Attribute - This is a PA-TNC attribute (as defined in RFC 5792 [RFC5792] extension for conveying software inventory information. This specification defines several new attribute types.

Endpoint's Software Inventory Evidence Collection - The set of information regarding the set of software installed on an endpoint. An endpoint's software inventory evidence collection might include information created by or derived from multiple sources, including but not limited to SWID tag files deposited on the file system during software installation, information generated to report output from

software discovery tools, and information dynamically generated by a software or package management system on an endpoint.

Software Inventory Evidence Record - The endpoint's Software Inventory Evidence Collection is composed of "records". Each record corresponds to one installed instance of a particular software product as reported by some data source. It is possible for a single installed instance to have multiple software inventory evidence records in an endpoint's Software Inventory Evidence Collection - this can happen if multiple sources all report the same software installation instance.

Software Identifier - A string associated with a specific version of a specific software product. Each supported Software Inventory Message and Attributes for PA-TNC data model has its own rules for how a Software Identifier (see Section 3.2.1) is derived from the representation of the given software product using that data model, and different sources for this information might populate relevant information differently. As such, while each Software Identifier uniquely identifies a specific software product, the same software product might be associated with multiple Software Identifiers reflecting differences between different information sources and supported data models.

2. Background

2.1. Supported Use Cases

This section describes the Software Inventory Message and Attributes for PA-TNC use cases supported by this specification. The primary use of exchanging software inventory information over the PA-TNC interface is to enable a challenger (e.g. NEA Server) to obtain inventory evidence about some system in a way that conforms to NEA procedures and expressed using a standard format. Collected software information can support a range of security activities including determining whether an endpoint is permitted to connect to the enterprise, determining which endpoints contain software that requires patching, and similar activities.

2.1.1. Use Software Inventory as a Factor in Determining Endpoint Access

Some enterprises might define security policies that require connected endpoints to have certain pieces of security software installed. By contrast, some security policies might prevent access by endpoints that have certain prohibited pieces of software installed, such as applications that pose known security risks. To support such policies, the NEA Server needs to collect evidence

indicating the software inventory of an endpoint that is seeking to initiate or continue connectivity to the enterprise.

Software Inventory Message and Attributes for PA-TNC facilitates policy decisions that consider an endpoint's software inventory by providing the NEA Server with software inventory information from the endpoint. The SW-PC can provide a complete or partial inventory to the SW-PV as required to determine policy compliance. The SW-PV can then use this as evidence of compliance or non-compliance with enterprise policy.

2.1.2. Maintain a Central Repository Reflecting an Endpoint's Software Inventory

Many tools can use information about an endpoint's software inventory to monitor and enforce the security of an enterprise. For example, a software patching service can use an endpoint's software inventory to determine whether certain endpoints have software that requires patching. A vulnerability management tool might identify endpoints with known vulnerabilities (patched or otherwise) and use this to gauge enterprise exposure to attack. A license management tool might verify that all copies of a particular piece of software are accounted for within the enterprise. The presence of a central repository representing a real-time understanding of each endpoint's software inventory facilitates all such activities. Using a central repository that can ensure the freshness of its collected information is generally more efficient than having each tool collect the same inventory information from each endpoint individually and leads to a more consistent understanding of enterprise state.

Software Inventory Message and Attributes for PA-TNC supports this activity through a number of mechanisms. As noted above, it allows a SW-PC to provide a complete list of software present in an endpoint's Software Inventory Evidence Collection to the SW-PV, which can then pass this information on to a central repository such as a Configuration Management Database (CMDB) or similar application. In addition, SW-PCs are required to be able to monitor for changes to an endpoint's Software Inventory Evidence Collection in near real-time and push reports of changes to the SW-PV as soon as those changes are detected. Thus any central repository fed by a SW-PV receiving such information can be updated soon after the change occurs. Keeping such a central repository synchronized with the state of each endpoint's Software Inventory Evidence Collection allows tools that use this information for their own security activities to make decisions in a consistent, efficient manner.

2.1.3. PA-TNC Use Cases

Software Inventory Message and Attributes for PA-TNC are intended to operate over the PA-TNC interface and, as such, are intended to meet the use cases set out in the PA-TNC specification.

2.2. Non-supported Use Cases

Some use cases not covered by this version of Software Inventory Message and Attributes for PA-TNC include:

- o This specification does not address how the endpoint's Software Inventory Evidence Collection is populated. In particular, NEA components are not expected to perform software discovery activities beyond compiling information in an endpoint's Software Inventory Evidence Collection. This collection might potentially come from multiple sources on the endpoint (e.g., information generated dynamically by package management tools or discovery tools, as well as SWID tag files discovered on the file system). While an enterprise might make use of software discovery procedures to identify installed software such procedures are outside the scope of this specification.
- o This specification does not address converting inventory information expressed in a proprietary format into formats used in the attributes described in this specification. Instead, it focuses exclusively on defining interfaces for the transportation of software information in the expectation that this is the format around which reporting tools will converge.
- o This specification provides no mechanisms for a posture validator to request a specific list of software information based on arbitrary software properties. For example, requesting only information about software from a particular vendor is not supported. After the endpoint's software inventory evidence collection has been copied to some central location, such as the CMDB, processes there can perform queries based on any criteria present in the collected information, but this specification does not address using such queries to constrain the initial collection of this information from the endpoint.
- o This specification does not address utilization of properties of certain sources of software information that might facilitate local tests (i.e., on the endpoint) of endpoint state. For example, the optional package_footprint field of an ISO SWID tag can contain a list of files and hash values associated with the software indicated by the tag. Tools on the endpoint can use the values in this field to test for the presence of the indicated

files. Successful evaluation of such tests leads to greater assurance that the indicated software is present on the endpoint. Currently, most SWID tag creators do not provide values for tag fields that support local testing. For this reason, the added complexity of supporting endpoint testing using these fields is out of scope for this specification. Future versions of this specification might add support for such testing.

2.3. Specification Requirements

Below are the requirements that the Software Inventory Message and Attributes for PA-TNC specification is required to meet in order to successfully play its role in the NEA architecture.

o Efficient

The NEA architecture enables delay of network access until the endpoint is determined not to pose a security threat to the network based on its asserted integrity information. To minimize user frustration, the Software Inventory Message and Attributes for PA-TNC ought to minimize overhead delays and make PA-TNC communications as rapid and efficient as possible.

Efficiency is also important when one considers that some network endpoints are small and low powered, some networks are low bandwidth and/or high latency, and some transport protocols (such as PT-EAP, Posture Transport (PT) Protocol for Extensible Authentication Protocol (EAP) Tunnel Methods [RFC7171]) or their underlying carrier protocol might allow only one packet in flight at a time or only one roundtrip. However, when the underlying PT protocol imposes fewer constraints on communications, this protocol ought to be capable of taking advantage of more robust communication channels (e.g. using larger messages or multiple roundtrips).

o Scalable

Software Inventory Message and Attributes for PA-TNC needs to be usable in enterprises that contain tens of thousands of endpoints or more. As such, it needs to allow a security tools to make decisions based on up-to-date information about an endpoint's software inventory without creating an excessive burden on the enterprise's network.

o Interoperable

This specification defines the protocol for how PCs and PVs can exchange and use software information to provide a NEA Server with information about an endpoint's software inventory. Therefore a key

goal for this specification is ensuring that all SW PCs and PVs, regardless of the vendor who created them, are able to interoperate in their performance of these duties.

- o Support precise and complete historical reporting

This specification outlines capabilities that support real-time understanding of the state of endpoint in a network in a way that can be used by other tools. One means of facilitating such an outcome is for a Configuration Management Database (CMDB) to be able to contain information about all endpoints connected to the enterprise for all points in time between the endpoint's first connection and the present. In such a scenario, it is necessary that any PC be able to report any changes to its software inventory evidence collection in near real-time while connected and, upon reconnection to the enterprise, be able to update the NEA Server (and through it the CMDB) with regard to the state of its software inventory evidence collection throughout the entire interval when it was not connected.

2.4. Non-Requirements

There are certain requirements that the Software Inventory Message and Attributes for PA-TNC specification explicitly is not required to meet. This list is not exhaustive.

- o End to End Confidentiality

This specification does not define mechanism for confidentiality, nor is this property automatically provided by PA-TNC interface use. Confidentiality is generally provided by the underlying transport protocols, such as the PT Binding to TLS [RFC6876] or PT-EAP Posture Transport for Tunneled EAP Methods [RFC7171] - see Section 8 for more information on related standards. Should users wish confidentiality protection of assessment instructions or results, this needs to be provided by parts of the NEA architecture other than this specification.

2.5. Assumptions

Here are the assumptions that Software Inventory Message and Attributes for PA-TNC makes about other components in the NEA architecture.

- o Reliable Message Delivery

The Posture Broker Client and Posture Broker Server are assumed to provide reliable delivery for PA-TNC messages and therefore the Attributes sent between the SW PCs and the PVs. In the event that

reliable delivery cannot be provided, the Posture Collector or Posture Validator is expected to terminate the connection.

2.6. Non-Assumptions

The Software Inventory Message and Attributes for PA-TNC specification explicitly does not assume:

- o Authenticity and Accuracy of the Software Inventory Evidence Collection with Regard to Endpoint Inventory

This specification makes no assumption as to whether the software information that it reports correctly reflect the software state on the endpoint. This specification does not attempt to detect when the endpoint is providing false information, either through malice or error, but instead focuses on correctly and reliably providing the reported Software Inventory Evidence Collection to the NEA Server. Similarly, this specification makes no assumption with regard to the completeness of the Software Inventory Evidence Collection's coverage of the total set of software installed on the endpoint. It is possible, and even likely, that some installed software is not represented by a record in an endpoints Software Inventory Evidence Collection. See Section 6 for more on this security consideration.

2.7. Software Inventory Message and Attributes for PA-TNC Diagram Conventions

This specification defines the syntax of the Software Inventory Message and Attributes for PA-TNC using diagrams. Each diagram depicts the format and size of each field in bits. Implementations MUST send the bits in each diagram as they are shown from left to right for each 32-bit quantity traversing the diagram from top to bottom. Multi-octet fields representing numeric values MUST be sent in network (big endian) byte order.

Descriptions of bit fields (e.g. flags) values refer to the position of the bit within the field. These bit positions are numbered from the most significant bit through the least significant bit. As such, an octet with only bit 0 set would have a value of 0x80 (1000 0000), an octet with only bit 1 set would have a value of 0x40 (0100 0000), and an octet with only bit 7 set would have a value of 0x01 (0000 0001).

3. Software Inventory Message and Attributes for PA-TNC System Requirements

The Software Inventory Message and Attributes for PA-TNC specification facilitates the exchange of software inventory and event information. Specifically, each application supporting Software Inventory Message and Attributes for PA-TNC includes a component known as the SW-PC that receives messages sent with the SW Attributes component type. The SW-PC is also responsible for sending appropriate SW Attributes back to the SW-PV in response. This section outlines what software inventories and events are and the requirements on SW-PCs and SW-PVs in order to support the stated use cases of this specification.

3.1. Basic Attribute Exchange

In the most basic exchange supported by this specification, a SW-PV sends a request to the SW-PC requesting some type of information about the endpoint's software inventory. This simple exchange is shown in Figure 2.




Figure 2: Basic SW Attribute Exchange

Upon receiving such a SW Request from the SW-PV, the SW-PC is expected to collect all the relevant software inventory information from the endpoint's software evidence collection and place it within its response attribute.

SW-PVs MUST discard without error any SW Response attributes that they receive for which they do not know the SW Request parameters that led to this SW Response. This is due to the fact that the SW Request includes parameters that control the nature of the response (as will be described in the following sections) and without knowing those parameters the SW Response cannot be reliably interpreted. Most often receiving an unsolicited SW Response attribute happens when a NEA Server has multiple SW-PVs; one SW-PV sends a SW Request but, unless exclusive delivery is used by the SW-PC in sending the

response, both SW-PVs receive copies of the resulting SW Response. In this case, the SW-PV that didn't send the SW Request would lack the context necessary to correctly interpret the SW Response it received and would simply discard it. Note, however, that proprietary measures might allow a SW-PV to discover the SW Request parameters for a SW Response even if that SW-PV did not send the given SW Request. As such, there is no blanket requirement for a SW-PV to discard all SW Responses to SW Request the SW-PV did not generate itself, only that SW-PVs are required to discard SW Responses for which they cannot get the necessary context to interpret.

In the case that it is possible to do so, the SW-PC SHOULD send its SW Response attribute to the SW-PV that requested it using exclusive delivery as described in section 4.5 of RFC 5793 (PB-TNC) [RFC5793]. Exclusive delivery ensures that only the sender of the SW Request receives the resulting SW Response.

3.2. Core Software Reporting Information

Different parameters in the SW Request can influence what information is returned in the SW Response. However, while each SW Response provides different additional information about this installed software, they all share a common set of fields that support reliable software identification on an endpoint. These fields include: Software Identifiers, Record Identifiers, and Software Locators. These three fields are present for each reported piece of software in each type of SW Response. The following sections examine these three types of information in more detail.

3.2.1. Software Identifiers

A Software Identifier uniquely identifies a specific version of a specific software product. The Software Inventory Message and Attributes for PA-TNC specification does not dictate the structure of a Software Identifier (beyond stating that it is a string) or define how it is created. Instead, each data model described in the Software Data Model IANA table (Section 9.4) includes its own rules for how a Software Identifier is created based on a record in the Endpoint's Software Inventory Evidence Collection expressed in that data model. Other data models will have their own procedures for the creation of associated Software Identifiers. Within the Software Inventory Message and Attributes for PA-TNC specification, the Software Identifier is simply an opaque string and there is never any need to unpack any information that might be part of that identifier.

A Software Identifier is a fraction of the size of the inventory record from which it is derived. For some combinations of data

models and sources, the full record might never be necessary as the identifier can be directly correlated to the contents of the full record. This is possible with authoritative SWID tags, since these tags always have the same contents and thus a Software Identifier derived from these tags can be used as a lookup to a local copy of the full tag. For other combinations of source and data model, a server might not be able to determine the specific software product and version associated with the identifier without requesting delivery of the full record. However, even in those cases, downstream consumers of this information might never need the full record as long as the Software Identifiers they receive can be tracked reliably. A SW-PV can use Software Identifiers to track the presence of specific software products on an endpoint over time in a bandwidth-efficient manner.

There are two important limitations of Software Identifiers to keep in mind:

1. The identifiers do not necessarily change when the associated record changes. In some situations, a record in the endpoint's Software Inventory Evidence Collection will change due to new information becoming available or in order to correct prior errors in that information. Such changes might or might not result in changes to the Software Identifier, depending on the nature of the changes and the rules governing how Software Identifiers are derived from records of the appropriate data model.
2. It is possible that a single software product is installed on a single endpoint multiple times. If both of these installation instances are reported by the same source using the same data format, then this can result in identical Software Identifiers for each installation instances. In other words, Software Identifiers might not uniquely identify installation instances; they just are intended to uniquely identify software products (which might have more than one installation instance). Instead, to reliably distinguish between multiple instances of a single software product, one needs to make use of Record Identifiers, described in the following section.

3.2.1.1. Record Identifiers

A Record Identifier is a 4-byte integer generated by the SW-PC that is uniquely associated with a specific record within the Endpoint's Software Inventory Evidence Collection. The SW-PC MUST assign a unique identifier to each record when it is added to the Endpoint's Software Inventory Evidence Collection. The Record Identifier SHOULD remain unchanged if that record is modified. The SWID-PC might wish

to assign Record Identifiers sequentially, but any scheme is acceptable provided that no two records receive the same identifier.

Servers can use Record Identifiers to distinguish between multiple instances of a single software product installed on an endpoint. Since each installation instance of a software product is associated with a separate record, servers can use the record identifier to distinguish between instances. For example, if an event is reported (as described in Section 3.5), the record identifier will allow the server to discern which instance of a software product is involved.

3.2.1.2. Software Locators

In addition to the need to identify software products, many use cases of inventory information need to know where software is located on the endpoint. This information might be needed to direct remediation actions or similar processes. For this reason, every reported software product also includes a Software Locator to identify where the software is installed on the endpoint.

If the location is not provided directly by the record source the SW-PC is responsible for attempting to determine the location of the software product. The "location" of a product SHOULD be the directory in which the software products executables are kept. However, if that directory is shared by other software products, the "location" SHOULD be the location of the primary executable associated with the software product. The source and/or SW-PC MUST be consistent in reporting the location of a software product (i.e., it cannot use the executable location in one report and the directory location in another).

The location is expressed as a URI string consisting of a scheme and path. [RFC3986] The location URI does not include an authority part. The URI schema describes the context of the described location. For example, in most cases the location of the installed software product will be expressed in terms of its path in the filesystem. For such locations, the location URI scheme MUST be "file". It is possible that other schemes could be used to represent other location contexts. Apart from reserving the "file" and "unknown" (described below) scheme to indicate an installation location expressed using a path in the endpoint's file system, this specification does not reserve schemes. When representing software products in other location contexts, tools MUST be consistent in their use of schemes, but the exact string used in those schemes is not normatively defined here.

It is possible, that a SW-PC is unable to determine the location of a reported software product. In this case, the SW-PC MUST assign that

software product a location of "unknown:". (I.e., the "unknown" scheme and an empty path.) However, SW-PCs SHOULD only make such an assignment as a last resort. Even a probable location for a software product is preferable to using the unknown indicator.

3.2.1.3. Using Software and Record Identifiers in SW Attributes

A SW Attribute reporting an endpoint's Software Inventory Evidence Collection always contains the Software Identifiers associated with the identified software products. A SW Attribute might or might not also contain copies of software inventory evidence records. The attribute exchange is identical to the diagram shown in Figure 2 regardless of whether software inventory evidence records are included. The SW Request attribute indicates whether the response is required to include software inventory evidence records. Excluding software inventory evidence records can reduce the attribute size of the response by multiple orders of magnitude when compared to sending the same inventory with full records.

3.3. Targeted Requests

Sometimes a SW-PV does not require information about every piece of software on an endpoint but only needs to receive updates about certain software instances. For example, enterprise endpoints might be required to have certain software products installed and to keep these updated. Instead of requesting a complete inventory just to see if these products are present, the SW-PV can make a "targeted request" for the software in question.

Targeted requests follow the same attribute exchange described in Figure 2. The SW-PV targets its request by providing one or more Software Identifiers in its SW Request attribute. The SW-PC MUST then limit its response to contain only records that match the indicated Software Identifier(s). This allows the network exchange to exclude information that is not relevant to a given policy question, thus reducing unnecessary bandwidth consumption. The SW-PC's response might or might not include software inventory evidence records, depending on the parameters of the SW Request.

Note that targeted requests identify the software relevant to the request only through Software Identifiers. This specification does not support arbitrary, parameterized querying of records. For example, one cannot request all records from a certain software publisher, or all records created by a particular record source. Targeted requests only allow a requestor to request specific software (as identified by their Software Identifiers) and receive a response that is limited to the named software.

There is no assumption that a SW-PC will recognize "synonymous records" - that is, records from different sources for the same software. Recall that different sources and data models may use different Software Identifier strings for the same software product. The SW-PC returns only records that match the Software Identifiers in the SW Request, even if there might be other records in the endpoint's Software Inventory Evidence Collection for the same software product. This is necessary because SW-PCs might not have the ability to determine that two Software Identifiers refer to the same product.

Targeted requests do not include Record Identifiers or Software Locators. The response to a targeted request MUST include all records associated with the named Software Identifiers, including the case where there are multiple records associated with a single Software Identifier.

SW-PCs MUST accept targeted requests and process them correctly as described above. SW-PVs MUST be capable of making targeted requests and processing the responses thereto.

3.4. Monitoring Changes in an Endpoint's Software Inventory Evidence Collection

The software collection on an endpoint is not static. As software is installed, uninstalled, patched, or updated, the Software Inventory Evidence Collection is expected to change to reflect the new software state on the endpoint. Different record sources might update the evidence collection at different rates. For example, a package manager might update its records in the Software Inventory Evidence Collection immediately whenever it is used to add or remove a software product. By contrast, sources that perform periodic examination of the endpoint would likely only update their records in the Software Inventory Evidence Collection after each examination.

All SW-PCs MUST be able to detect changes to the Software Inventory Evidence Collection. Specifically, SW-PCs MUST be able to detect:

- o The creation of records
- o The deletion of records
- o The alteration of records

An "alteration" is anything that modifies the contents of a record (or would modify it, if the record is dynamically generated on

demand) in any way, regardless of whether the change is functionally meaningful.

SW-PCs MUST detect such changes to the endpoint's Software Inventory Evidence Collection in close to real-time (i.e., within seconds) when the Posture Collector is operating. In addition, in the case where there is a period during which the SW-PC is not operating, the SW-PC MUST be able to determine the net change to the endpoint's Software Inventory Evidence Collection over the period it was not operational. Specifically, the "net change" represents the difference between the state of the endpoint's Software Inventory Evidence Collection when the SW-PC was last operational and monitoring its state, and the state of the endpoint's software inventory evidence collection when the SW-PC resumed operation. Note that a net change might not reflect the total number of change events over this interval. For example, if a record was altered three times during a period when the SW-PC was unable to monitor for changes, the net change of this interval might only note that there was an alteration to the record, but not how many individual alteration events occurred. It is sufficient for a SW-PC's determination of a net change to note that there was a difference between the earlier and current state rather than enumerating all the individual events that allowed the current state to be reached.

The SW-PC MUST assign a time to each detected change in the endpoint's Software Inventory Evidence Collection. These timestamps correspond to the SW-PC's best understanding as to when the detected change occurred. These timestamps MUST be as accurate as possible. For changes to the endpoint's Software Inventory Evidence Collection that occur while the SW-PC is operating, the SW-PC ought to be able to assign a time to the event that is accurate to within a few seconds. For changes to the endpoint's Software Inventory Evidence Collection that occur while the SW-PC is not operational, upon becoming operational the SW-PC needs to make a best guess as to the time of the relevant events (possibly by looking at timestamps on files), but these values might be off. In the case of dynamically generated records, the time of change is the time at which the data from which the records are generated changes, not the time at which a changed record is generated. For example, if records are dynamically generated based on data in an RPM database, the time of change would be when the RPM database changed.

With regard to deletions of records, the SW-PC needs to detect the deletion and MUST retain a copy of the full deleted record along with the associated Record Identifier and Software Locator so that the record and associated information can be provided to the SW-PV upon request. This copy of the record MUST be retained for a reasonable amount of time. Vendors and administrators determine what

"reasonable" means, but a copy of the record SHOULD be retained for as long as the event recording the deletion of the record remains in the SW-PC's event log (as described in Section 3.5). This is recommended because, as long as the event is in the SW-PC's change logs, the SW-PC might send an event attribute (described in Section 3.5) that references this record, and a copy of the record is needed if the SW-PV wanted a full copy of the relevant records.

With regard to alterations to a record, SW-PCs MUST detect any alterations to the contents of a record. Alterations need to be detected even if they have no functional impact on the record. A good guideline is that any alteration to a record that might change the value of a hash taken on the record's contents needs to be detected by the SW-PC. A SW-PC might be unable to distinguish modifications to the content of a record from modifications to the metadata the file system associates with the record. For example, a SW-PC might use the "last modification" timestamp as an indication of alteration to a given record, but a record's last modification time can change for reasons other than modifications to the record contents. A SW-PC is still considered compliant with this specification if it also reports metadata change events that do not change the record itself as alterations to the record. In other words, while SW-PC authors are encouraged to exclude modifications that do not affect the bytes within the record, discriminating between modifications to file contents and changes to file metadata can be difficult and time consuming on some systems. As such, as long as the alterations detected by a SW-PC always cover all modifications to the contents of record, the SW-PC is considered compliant even if it also registers alterations that do not modify the contents of a record as well. When recording an alteration to a record, the SW-PC is only required to note that an alteration occurred. The SW-PC is not required to note or record how the record altered, nor is it possible to include such details in SW Attributes reporting the change to a SW-PV. There is no need to retain a copy of the original record.

When a record changes it SHOULD retain the same Record Identifier. The Software Locator might or might not change, depending on whether the software changed locations during the changes that led to the record change. A record change MUST retain the same Software Identifier. This means that any action that changes a software product (e.g., application of a patch that results in a change to the product's version) MUST NOT be reflected by a record change but instead MUST result in the deletion of the old record and the creation of a new record. This reflects the requirement that a record in the endpoint's Software Inventory Evidence Collection correspond directly with an instance of a specific software product.

3.5. Reporting Change Events

As noted in the preceding section, SW-PCs MUST be able to detect changes to the endpoints Software Inventory Evidence Collection (creation, deletion, and alteration) in near real-time while the SW-PC is operational, and MUST be able to account for any net change to the endpoint's Software Inventory Evidence Collection that occurs when the SW-PC is not operational. However, to be of use to the enterprise, the NEA Server needs to be able to receive these events and be able to understand how new changes relate to earlier changes. In Software Inventory Message and Attributes for PA-TNC, this is facilitated by reporting change events. All SW-PCs MUST be capable of receiving requests for change events and sending change event attributes. All SW-PVs MUST be capable of requesting and receiving change event attributes.

3.5.1. Event Identifiers

To be useful, change events need to be correctly ordered. Ordering of events is facilitated by two pieces of information: an Event Identifier (EID) value and an EID Epoch value.

An EID is a 4-byte unsigned integer that the SW-PC assigns sequentially to each observed event (whether detected in real-time or deduced by looking for net changes over a period of SW-PC inactivity). All EIDs exist within the context of some "EID Epoch", which is also represented as a 4-byte unsigned integer. EID Epochs are used to ensure synchronization between the SW-PC and any SW-PVs with which it communicates. EID Epoch values SHOULD be generated randomly and in such a way that it is unlikely that the same EID Epoch is generated twice, even if the SW-PC reverted to an earlier state (e.g., resetting it to factory defaults). In the case where a SW-PC needs to reset its EID counter, either because it has exhausted all available EID values or because the SW-PC's event log becomes corrupted, then a new EID Epoch MUST be selected.

Within an Epoch, EIDs MUST be assigned sequentially, so that if a particular event is assigned an EID of N, the next observed event is given an EID of N+1. In some cases, events might occur simultaneously, or the SW-PC might not otherwise be able to determine an ordering for events. In these cases, the SW-PC creates an arbitrary ordering of the events and assigns EIDs according to this ordering. Two change events MUST NOT ever be assigned the same EID within the same EID Epoch. No meaningful comparison can be made between EID values of different Epochs.

The EID value of 0 is reserved and MUST NOT be associated with any event. Specifically, an EID of 0 in a SW Request attribute indicates

that a SW-PV wants an inventory response rather than an event response, while an EID of 0 in a SW Response is used to indicate the initial state of the endpoint's Software Inventory Evidence Collection prior to the observation of any events. Thus the very first recorded event in a SW-PC's records within an EID Epoch MUST be assigned a value of 1 or greater. Note that EID and EID Epoch values are assigned by the SW-PC without regard to whether events are being reported to one or more SW-PVs. The SW-PC records events and assigns EIDs during its operation. Any and all SW-PVs that request event information from the SW-PC will have those requests served from the same event records and thus will see the same EIDs and EID Epochs for the same events.

The SW-PC MUST ensure there is no coverage gap (i.e., change events that are not recorded in the SW-PC's records) in its change event records. This is necessary because a coverage gap might give a SW-PV a false impression of the endpoint's state. For example, if a SW-PV saw an event indicating that a particular record had been added to the endpoint's software inventory evidence collection, and saw no subsequent events indicating that record had been deleted, it might reasonably assume that this record was still present and thus that the indicated software was still installed (assuming the Epoch has not changed). If there is a coverage gap in the SW-PC's event records, however, this assumption could be false. For this reason, the SW-PC's event records MUST NOT contain gaps. In the case where there are periods where it is possible that changes occurred without the SW-PC detecting or recording them, the SW-PC MUST either compute a net change and update its event records appropriately, or pick a new EID Epoch to indicate a discontinuity with previous event records.

Within a given Epoch, once a particular event has been assigned an EID, this association MUST NOT be changed. That is, within an Epoch, once an EID is assigned to an event, that EID cannot be reassigned to a different event, and the event cannot be assigned a different EID. When the SW-PC's Epoch changes, all of these associations between EIDs and events are cancelled, and EID values once again become free for assignment.

3.5.2. Core Event Tracking Information

Whether reporting events or full inventories it is important to know how the reported information fits into the overall timeline of change events. This is why all SW Response attributes include fields to place that response within the sequence of detected events. Specifically, all SW Responses include a Last EID and EID Epoch field. The EID Epoch field identifies the EID Epoch in which the SW Response was sent. If the SW Response is reporting events, all

reported events occurred within the named EID Epoch. The Last EID (which is also always from the named EID Epoch) indicates the EID of the last recorded change event at the time that the SW Response was sent. These two fields allow any response to be placed in the context of the complete set of detected change events within a given EID Epoch.

3.5.3. Updating Inventory Knowledge Based on Events

Modern endpoints can have hundreds of software products installed, most of which are unlikely to change from one day to the next. As such, instead of exchanging a complete list of an endpoint's inventory on a regular basis, one might wish to only identify changes since some earlier known state of this inventory. This is readily facilitated by the use of EIDs to place change events in a context relative to earlier state.

As noted above, every SW Response sent by a SW-PC to a SW-PV (as described in Section 3.1 through Section 3.3) includes the EID Epoch and EID of the last event recorded prior to that response being compiled. This allows the SW-PV to place all subsequently received event records in context relative to this SW Response attribute (since the EIDs represent a total ordering of all changes to the endpoint's software inventory evidence collection). Specifically, a SW-PV (or, more likely, a database that collects and records its findings) can record an endpoint's full inventory and also the EID and Epoch of the most recent event reflected at the time of that inventory. From that point on, if change events are observed, the attribute describing these events indicates the nature of the change, the affected records, and the order in which these events occurred (as indicated by the sequential EIDs). Using this information, any remote record of the endpoint's Software Inventory Evidence Collection can be updated appropriately.

3.5.4. Using Event Records in SW Attributes

A SW-PV MUST be able to request a list of event records instead of an inventory. The attribute flow in such an exchange looks the same as the basic flow shown in Figure 2. The only difference is that, in the SW Request attribute, the SW-PV provides an EID other than 0. (A value of 0 in these fields represents a request for an inventory.) When the SW-PC receives such a request, instead of identifying records from the endpoint's Software Inventory Evidence Collection, it consults its list of detected changes. The SW-PC MUST add an event record to the SW Response attribute for each recorded change event with an EID greater than or equal to the EID in the SW Request attribute (although targeting of requests, as described in the next

paragraph, might limit this list). A list of event records MUST only contain events with EIDs that all come from the current Epoch.

SW-PVs can target requests for event records by including one or more Software Identifiers, as described in Section 3.3, in the SW Request that requests an event record list. A targeted request for event records is used to indicate that only events affecting software that matches one of the provided Software Identifiers are to be returned. Specifically, in response to a targeted request for event records, the SW-PC MUST exclude any event records that are less than the indicated EID (within the current EID Epoch) and exclude any event records where the affected software does not match one of the provided Software Identifiers. This might mean that the resulting list of event records sent in the response attribute does not provide a continuous sequence of EIDs. Both SW-PCs and SWIC-PVs MUST support targeted request for event records.

3.5.5. Partial and Complete Lists of Event Records in SW Attributes

Over time, a SW-PC might record a large number of change events. If a SW-PV requests all change events covering a large period of time, the resulting SW Response attribute might be extremely large, especially if the SW-PV requests inclusion of software inventory evidence records in the response. In the case that the resulting attribute is too large to send (either because it exceeds the 4GB attribute size limit imposed by the PA-TNC specification, or because it exceeds some smaller size limit imposed on the SW-PC) the SW-PC MAY send a partial list of event records back to the SW-PV.

Generation of a partial list of events in a SW Response attribute requires the SW-PC to identify a "consulted range" of EIDs. A consulted range is the set of event records that are examined for inclusion in the SW Response attribute and that are included in that attribute if applicable. Recall that, if a SW Request is targeted, only event records that involve the indicated software would be applicable. (See Section 3.3 for more on Targeted Request.) If a request is not targeted, all event records in the considered range are applicable and included in the SW Response attribute.

The lower bound of the consulted range MUST be the EID provided in the SW Request. (Recall that a SW Request indicates a request for event records by providing a non-0 EID value in the SW Request. See Section 3.5.4.) The upper bound of the consulted range is the EID of the latest event record (as ordered by EID values) that is included in the SW Response attribute if it is applicable to the request. The EID of this last event record is called the "Last Consulted EID". The SW-PC chooses this Last Consulted EID based on the size of the event record list it is willing to provide to the SW-PV.

A partial result list MUST include all applicable event records within the consulted range. This means that for any applicable event record (i.e., any event record in an un-targeted request, or any event record associated with software matching a requested Software Identifier in a targeted request) whose EID is greater than or equal to the EID provided in the SW Request and whose EID is less than or equal to the Last Consulted EID, that event record MUST be included in the SW Response conveying this partial list of event records. This ensures that every partial list of event records is always complete within its indicated range.

All SW Response attributes that convey event records include a Last Consulted EID field. This is in addition to the EID Epoch and Last EID fields that are present in all SW Responses. Note that, if responding to a targeted SW Request, the SW Response attribute might not contain the event record whose EID matches the Last Consulted EID value. For example, the last consulted EID record might have been deemed inapplicable because it did not match the specified list of Software Identifiers in the SW Request.

If a SW-PV receives a SW Response attribute where the Last EID and Last Consulted EID fields are identical, the SW-PV knows that it has received a result list that is complete, given the parameters of the request, up to the present time.

On the other hand, if the Last EID is greater than the Last Consulted EID, the SW-PV has received a partial result list. (The Last Consulted EID MUST NOT exceed the Last EID.) In this case, if the SW-PV wishes to try to collect the rest of the partially delivered result list it then sends a new SW Request whose EID is one greater than the Last Consulted EID in the preceding response. Doing this causes the SW-PC to generate another SW Response attribute containing event records where the earliest reported event record is the one immediately after the event record with the Last Consulted EID (since EIDs are assigned sequentially). By repeating this process until it receives a SW Response where the Last EID and Last Consulted EID are equal, the SW-PV is able to collect all event records over a given range, even if the complete set of event records would be too large to deliver via a single attribute.

Implementers need to be aware that a SW Request might specify an EID that is greater than the EID of the last event recorded by a SW-PC. In accordance with the behaviors described in Section 3.5.4, a SW-PC MUST respond to such a request with a SW Response attribute that contains zero event records. This is because the SW-PC has recorded no event records with EIDs greater than or equal to the EID in the SW Request. In such a case, the Last Consulted EID field MUST be set to the same value as the Last EID field in this SW Response attribute.

This case is called out because the consulted range on a SW-PC in such a situation is a negative range, where the "first" EID in the range (provided in the SW Request) is greater than the "last" EID in the range (this being the EID of the last recorded event on the SW-PC). Implementers need to ensure that SW-PCs do not experience problems in such a circumstance.

Note that this specification only supports the returning of partial results when returning event records. There is no way to return a partial inventory list under this specification.

3.5.6. Synchronizing Event Identifiers and Epochs

Since EIDs are sequential within an Epoch, if a SW-PV's list of event records contains gaps in the EID values within a single Epoch, the SW-PV knows that there are events that have not been accounted for. The SW-PV can either request a new event list to collect the missing events or request a full inventory to re-sync its understanding of the state of the endpoint's Software Inventory Evidence Collection. In either case, after the SW-PV's record of the endpoint's Software Inventory Evidence Collection has been updated, the SW-PV can record the new latest EID value and track events normally from that point on.

If the SW-PV receives any attribute from a SW-PC where the EID Epoch differs from the EID Epoch that was used previously, then SW-PV or any entity using this information to track the endpoint's Software Inventory Evidence Collection knows that there is a discontinuity in their understanding of the endpoint's state. To move past this discontinuity and reestablish a current understanding of the state of the endpoint's Software Inventory Evidence Collection, the SW-PV needs to receive a full inventory from the endpoint. The SW-PV cannot be brought in sync with the endpoint's state through the collection of any set of event records in this situation. This is because it is not possible to account for all events on the SW-PC since the previous Epoch was used, because there is no way to query for EIDs from a previous Epoch. Once the SW-PV has received a full inventory for the new Epoch, the SW-PV records the latest EID reported in this new Epoch and can track further events normally.

A SW-PC MUST NOT report events with EIDs from any Epoch other than the current EID Epoch. The SW-PC MAY choose to purge all event records from a previous Epoch from memory after an Epoch change. Alternately, the SW-PC MAY choose to retain some event records from a previous EID Epoch and assign them new EIDs in the current Epoch. However, in the case where a SW-PC chooses the latter option it MUST ensure that the order of events according to their EIDs is unchanged and that there is no coverage gap between the first retained event

recorded during the previous Epoch (now reassigned with an EID in the current Epoch) and the first event recorded during the current Epoch. In particular, the SW-PC MUST ensure that all change events that occurred after the last recorded event from the previous Epoch are known and recorded. (This might not be possible if the Epoch change is due to state corruption on the SW-PC.) A SW-PC might choose to reassign EIDs to records from a preceding Epoch to create a "sliding window" of events, where each Epoch change represents a shift in the window of available events.

In the case where a SW-PC suffers a crash and loses track of its current EID Epoch or current EID, then it MUST generate a new EID Epoch value and begin assigning EIDs within that Epoch. In this case, the SW-PC MUST purge all event records from before the crash as it cannot ensure that there is not a gap between the last of those records and the next detected event. The process for generating a new EID Epoch MUST minimize the possibility that the newly generated EID Epoch is the same as a previously used EID Epoch.

The SW-PV will normally never receive an attribute indicating that the latest EID is less than the latest EID reported in a previous attribute within the same EID Epoch. If this occurs, the SW-PC has suffered an error of some kind, possibly indicative of at least partial corruption of its event log. In this case, the SW-PV SHOULD treat the situation as if there was a change in Epoch and treat any local copy of the endpoint's Software Inventory Evidence Collection as out-of-sync until a full inventory can be reported by the SW-PC. In this case, the SW-PV SHOULD flag the occurrence so the SW-PC can be examined to ensure it is now operating properly.

3.6. Subscriptions

Thus far, all attribute exchanges discussed assume that a SW-PV sent an SW Request attribute and the SW-PC is providing a direct response to that request. The Software Inventory Message and Attributes for PA-TNC specification also supports the ability for a SW-PC to send a SW Response to the SW-PV in response to observed changes in the endpoint's software inventory evidence collection, instead of in direct response to a SW Request. An agreement by a SW-PC to send content when certain changes are detected to the endpoint's Software Inventory Evidence Collection is referred to in this specification as a "subscription", and the SW-PV that receives this content is said to be "subscribed to" the given SW-PC. All SW-PCs and SW-PVs MUST support the use of subscriptions.

3.6.1. Establishing Subscriptions

A SW-PV establishes a subscription on a particular SW-PC by sending a SW Request attribute with the Subscription flag set. The SW Request attribute is otherwise identical to the SW Requests discussed in previous sections. Specifically, such a SW Request might or might not request inclusion of software inventory evidence records, might or might not be targeted, and might request change event records or endpoint inventory. Assuming no error is encountered, a SW-PC MUST send a SW Response attribute in direct response to this SW Request attribute, just as if the Subscription flag was not set. As such, the attribute exchange that establishes a new subscription in a SW-PC has the same flow seen in the previous attribute exchanges, as depicted in Figure 2. If the SW-PV does not receive a PA-TNC Error attribute (as described in Section 3.8 and Section 4.14) in response to their subscription request, the subscription has been successfully established on the SW-PC. The SW Request attribute that establishes a new subscription is referred to as the "establishing request" for that subscription.

When a subscription is established it is assigned a Subscription ID value. The Subscription ID is equal to the value of the Request ID of the establishing request. (For more about Request IDs, see Section 4.6.)

A SW-PC MUST have the ability to record and support multiple simultaneous subscriptions from a single party and from multiple parties. A SW-PV MUST have the ability to record and support multiple simultaneous subscriptions to a single party and subscriptions to multiple parties.

3.6.2. Managing Subscriptions

The SW-PC MUST record each accepted subscription along with the identity of the party to whom attributes are to be pushed in compliance with the subscription. This identity includes both the NEA Server's connection ID and the Posture Validator Identifier from the PB-PA message that delivered the request.

Likewise, SW-PVs MUST record each accepted subscription for which they are the subscribing party along with the associated Subscription ID and the identity of the SW-PC that will be fulfilling the subscription. The SW-PV needs to retain this information in order to correctly interpret pushed SW Response attributes sent in fulfillment of the subscription. The identity of the SW-PC is given in the Posture Collector Identifier of the PB-PA message header in all messages from that SW-PC.

3.6.3. Terminating Subscriptions

Subscriptions MAY be terminated at any time by the subscribing SW-PV by setting the Clear Subscriptions flag in a SW Request. (See Section 4.7 for more on using this flag.) In the case that a SW Request with the Clear Subscriptions flag set is received the SW-PC MUST only clear subscriptions that match both the NEA server connection ID and the SW-PV ID for this SW Request, and MUST clear all such subscriptions.

This specification does not give the SW-PV the ability to terminate subscriptions individually - all subscriptions to the SW-PV are cleared when the Clear Subscriptions flag is set.

This specification does not give the SW-PC the ability to unilaterally terminate a subscription. However, if the SW-PC experiences a fatal error fulfilling a subscription, resulting in sending a PA-TNC Error attribute of type `SW_SUBSCRIPTION_FULFILLMENT_ERROR`, then the subscription whose fulfillment led to the error MUST be treated as terminated by both the SW-PC and the SW-PV. Only the subscription experiencing the error is cancelled and other subscriptions are unaffected. See Section 3.8 for more on this error condition.

Finally, a subscription is terminated if the connection between the SW-PC and SW-PV is deleted. This occurs when the connection ID used in the messages between the SW-PC and the SW-PV becomes unbound. Loss of this connection ID would prevent the SW-PC from sending messages in fulfillment of this subscription. As such, loss of the connection ID necessarily forces subscription termination between the affected parties.

3.6.4. Subscription Status

A SW-PV can request that a SW-PC report the list of active subscriptions for which the SW-PV is the subscriber. A SW-PV can use this to recover lost information about active subscriptions. A SW-PV can also use this capability to verify that a SW-PC has not forgotten any of its subscriptions. The latter is especially useful where a SW-PC does not send any attributes in fulfillment of a given subscription for a long period of time. The SW-PV can check the list of active subscriptions on the SW-PC and verify whether the inactivity is due to a lack of reportable events or due to the SW-PC forgetting its obligations to fulfill a given subscription.

A SW-PV requests a list of its subscriptions on a given SW-PC by sending that SW-PC a Subscription Status Request. The SW-PC MUST then respond with a Subscription Status Response (or a PA-TNC Error

if an error condition is experienced). The Subscription Status Response MUST contain one subscription record for each of the active subscriptions for which the SW-PV is the subscribing party.

3.6.5. Fulfilling Subscriptions

As noted in Section 3.4 SW-PCs MUST have the ability to automatically detect changes to an endpoint's Software Inventory Evidence Collection in near real-time. For every active subscription, the SW-PC MUST send an attribute to the subscribed SW-PV whenever a change is detected to relevant records within the endpoint's Software Inventory Evidence Collection. Such an attribute is said to be sent "in fulfillment of" the given subscription and any such attribute MUST include that subscription's Subscription ID. If the establishing request for that subscription was a targeted request, then only records that match the Software Identifiers provided in that establishing request are considered relevant. Otherwise, (i.e., for non-targeted requests) any record is considered relevant for this purpose. Figure 3 shows a sample attribute exchange where a subscription is established and then later attributes are sent from the SW-PC in fulfillment of the established subscription.




Figure 3: Subscription Establishment and Fulfillment

The contents of an attribute sent in fulfillment of a subscription depend on the parameters provided in the establishing request for that subscription. Specifically, the contents of an attribute sent in fulfillment of a subscription have the same format as would a direct response to the establishing request. For example, if the establishing request stipulated a response that contained an event record list that included software inventory evidence records, all attributes sent in fulfillment of this subscription will also consist of event record lists with software inventory evidence records. As such, all SW Responses displayed in the exchange depicted in Figure 3 have the same format. A SW Response generated in fulfillment of an active subscription MUST be a valid SW Response attribute according to all the rules outlined in the preceding sections. In other words, an attribute constructed in fulfillment of a subscription will look the same as an attribute sent in direct response to an explicit request from a SW-PV that had the same request parameters and which arrived immediately after the given change event. There are a few special rules that expand on this guideline:

3.6.5.1. Subscriptions Reporting Inventories

In the case that a SW-PV subscribes to a SW-PC requesting an inventory attribute whenever changes are detected (i.e. the EID in the establishing request is 0), then the SW-PC MUST send the requested inventory whenever a relevant change is detected. (A "relevant change" is any change for untargeted requests, or a change to an indicated record in a targeted request.) Upon detection of a relevant change for an active subscription, the SW-PC sends the appropriate inventory information as if it had just received the establishing request. Attributes sent in fulfillment of this subscription will probably have a large amount of redundancy, as the same records are likely to be present in each of these SW Attributes. The role of an inventory subscription is not to report records just for the items that changed - that is the role of a subscription that reports events (see Section 3.6.5.2). A SW-PC MUST NOT exclude a record from an attribute sent in fulfillment of an inventory subscription simply because that record was not involved in the triggering event (although a record might be excluded for other reasons, such as if the subscription is targeted - see Section 3.6.5.3).

3.6.5.2. Subscriptions Reporting Events

The way in which a SW-PV indicates it wishes to establish a subscription requesting event records is by providing a non-zero EID in the SW Request establishing the subscription (see Section 3.5.1). However, when the SW-PC constructs an attribute in fulfillment of the subscription (other than the direct response to the establishing

request), it MUST only include event records for the detected change(s) that precipitated this response attribute. In other words, it MUST NOT send a complete list of all changes starting with the indicated EID, up through the latest change, every time a new event is detected. In effect, the EID in the establishing request is treated as being updated every time an attribute is sent in fulfillment of this subscription, such that a single event is not reported twice in fulfillment of a single subscription. As such, every SW-PC MUST track the EID of the last event that triggered an attribute for the given subscription. When the next event (or set of events) is detected, the SW-PC MUST only report events with EIDs after the last reported event. In the case that the EID Epoch of the SW-PC changes, the SW-PC MUST treat EID values in the new Epoch as being after all EIDs assigned in the previous Epoch regardless of the relative numeric values of these EIDs.

Note that while a subscription is active, the subscribing SW-PV MAY make other requests for event records that overlap with events that are reported in fulfillment of a subscription. Such requests are unaffected by the presence of the subscription, nor is the subscription affected by such requests. In other words, a given request will get the same results back whether or not there was a subscription. Likewise, an attribute sent in fulfillment of a subscription will contain the same information whether or not other requests had been received from the SW-PV.

A SW-PV needs to pay attention to the EID Epoch in these attributes, as changes in the Epoch might create discontinuities in the SW-PV's understanding of the endpoint's Software Inventory Evidence Collection state, as discussed in Section 3.5.6. In particular, once the EID Epoch changes, a SW-PV is unable have confidence that it has a correct understanding of the state of an endpoint's Software Inventory Evidence Collection until after the SW-PV collects a complete inventory.

SW-PCs MAY send partial lists of event records in fulfillment of a subscription. (See Section 3.5.5 for more on partial list of event records.) In the case that a SW-PC sends a partial list of event records in fulfillment of a subscription, it MUST immediately send the next consecutive partial list, and continue doing so until it has sent the equivalent of the complete list of event records. In other words, if the SW-PC sends a partial list it does not wait for another change event to send another SW Response, but continues sending SW Responses until it has sent all event records that would have been included in a complete fulfillment of the subscription. Note that the direct response to the establishing request is not considered to be sent in fulfillment of a subscription. However, in this case the SW-PC MUST treat the presence of unreported events as a triggering

event for pushing additional messages in fulfillment of the newly established subscription. As such, the net effect is that, if the direct response to the establishing request (i.e., the Subscription Fulfillment flag is unset) is partial, the SW-PC will immediately follow this with additional attributes (with the Subscription Fulfillment flag set) until the complete set of events has been sent to the SW-PV.

3.6.5.3. Targeted Subscriptions

Subscriptions MAY be targeted to only apply to records that match a given set of Software Identifiers. In the case where changes are detected that affect multiple records, some matching the establishing request's Software Identifiers and some not, the attribute sent in fulfillment of the subscription MUST only include inventory or events (as appropriate) for records that match the establishing request's Software Identifiers. The SW-PC MUST NOT include non-matching records in the attribute, even if those non-matching records experienced change events that were co-temporal with change events on the matching records.

In addition, a SW-PC MUST send an attribute in fulfillment of a targeted subscription only when changes to the endpoint's Software Inventory Evidence Collection impact one or more records matching the subscription's establishing request's Software Identifiers. A SW-PC MUST NOT send any attribute in fulfillment of a targeted subscription based on detected change to the endpoint's Software Inventory Evidence Collection that did not involve any of the records targeted by that subscription.

3.6.5.4. No Subscription Consolidation

A SW-PV MAY establish multiple subscriptions to a given SW-PC. If this is the case, it is possible that a single change event on the endpoint might require fulfillment by multiple subscriptions, and that the information included in attributes that fulfill each of these subscriptions might overlap. The SW-PC MUST send separate attributes for each established subscription that requires a response due to the given event. Each of these attributes MUST contain all information required to fulfill that individual subscription, even if that information is also sent in other attributes sent in fulfillment of other subscriptions at the same time. In other words, SW-PCs MUST NOT attempt to combine information when fulfilling multiple subscriptions simultaneously, even if this results in some redundancy in the attributes sent to the SW-PV.

3.6.5.5. Delayed Subscription Fulfillment

A SW-PC MAY delay the fulfillment of a subscription following a change event in the interest of waiting to see if additional change events are forthcoming and, if so, conveying the relevant records back to the SW-PV in a single SW Response attribute. This can help reduce network bandwidth consumption between the SW-PC and the SW-PV. For example, consider a situation where 10 changes occur a tenth of a second apart. If the SW-PC does not delay in assembling and sending SW Response attributes, the SW-PV will receive 10 separate SW Response attributes over a period of 1 second. However, if the SW-PC waits half a second after the initial event before assembling a SW Response, the SW-PV only receives two SW Response attributes over the same period of time.

Note that the ability to consolidate events for a single subscription over a given period of time does not contradict the rules in Section 3.6.5.4 prohibiting consolidation across multiple subscriptions. When delaying fulfillment of subscriptions, SW-PCs are still required to fulfill each individual subscription separately. Moreover, in the case that change events within the delay window cancel each other out (e.g., a record is deleted and then re-added), the SW-PC MUST still report each change event rather than just reporting the net effect of changes over the delay period. In other words, delayed fulfillment can decrease the number of attributes sent by the SW-PC, but it does not reduce the total number of change events reported.

SW-PCs are not required to support delayed fulfillment of subscriptions. However, in the case that the SW-PC does support delayed subscription fulfillment, it MUST be possible to configure the SW-PC to disable delayed fulfillment. In other words, parties deploying SW-PCs need to be allowed to disable delayed subscription fulfillment in their SW-PCs. The manner in which such configuration occurs is left to the discretion of implementers, although implementers MUST protect the configuration procedure from unauthorized tampering. In other words, there needs to be some assurance that unauthorized individuals are not able to introduce long delays in subscription fulfillment.

3.7. Multiple Sources of Software Inventory Evidence Records

The records in an endpoint's software inventory evidence collection might potentially come from multiple sources. For example, records might be derived from ISO SWID tags deposited on the file system and collected therefrom. Records might also be generated by tools such as software and package managers (e.g., RPM or YUM) or might be translated from software discovery reports.

A SW-PC is not required to identify every possible source of software information on its endpoint. Some SW-PCs might be explicitly tied only to one or a handful of software inventory sources. For all software inventory evidence sources that a particular SW-PC supports, it MUST completely support all requirements of this specification with regard to those sources. In other words, for supported sources, the SW-PC is required to be capable of providing a complete set of the provided software inventory evidence records; monitoring for changes in the records reported by those sources, correctly providing responses for both full and targeted requests for records from those sources, and delivering complete software inventory evidence records as appropriate. In all cases, the SW-PC MUST also be capable of deriving a Software Identifier from the resulting record and also assigning that record a unique Record Identifier. The SW-PC MUST NOT provide any inventory or event information from software inventory sources for which it cannot provide this full support. Note that the SW-PC SHOULD be able to provide a Software Locator for each software product reported by a given source, but it is recognized that this might not be possible in all circumstances and the inability to do so does not preclude use of the given source.

When providing a SW Response (either in direct response to a SW Request or in fulfillment of a subscription) the SW-PC MUST include the complete set of relevant data from all supported sources of software inventory evidence. In other words, a full inventory is required to contain all records from all supported sources, a targeted inventory is required to contain all relevant records from all sources, and event tracking is required to cover all events from all sources. With regard to events, a SW-PC's assignment of EIDs MUST reflect the presence and order of all events on the endpoint (at least for supported sources) regardless of the source. This means that if source A experiences an event, and then source B experiences two events, and then source A experiences another two events, the SW-PC is required to capture five events with consecutive EID values reflecting the order in which the events occur.

Note that, if a SW-PC collects data from multiple sources, it is possible that some software products might be "double counted". This can happen if both sources of inventory evidence provide a record for a single installation of a software product. Moreover, each of these provided records might have different Software Identifier and Software Locator values due to the different ways a source might report its information. When a SW-PC reports information or records events from multiple inventory evidence sources, it MUST use the information those sources provide, rather than attempting to perform some form of reduction. In other words, if multiple sources report records corresponding to a single installation of a software product, all such records from each source are required to be part of the SW-

PC's processing even if this might lead to multiple reporting, and the SW-PC is not to ignore some records to avoid such multiple reporting. Similarly, in the case that multiple sources report an event, the SW-PC MUST create separate event records with separate EIDs for each of these, even if there is the chance that they represent multiple sources reporting the same action on the endpoint. Entities tracking software inventory information collected via SW-PCs and SW-PVs need to be aware that such double-reporting might occur. How (or if) such occurrences are detected and resolved is up to the implementers of those entities.

3.8. Error Handling

In the case where the SW-PC detects an error in a SW Request attribute that it receives it MUST respond with a PA-TNC Error attribute with an error code appropriate to the nature of the error. (See Section 4.2.8 of PA-TNC [RFC5792] for more details about PA-TNC Error attributes and error codes as well as Section 4.14 in this specification for error codes specific to SW Attributes.) In the case that an error is detected in a SW Request the SW-PC MUST NOT take any action requested by this SW Request, even if partial completion of the SW is possible. In other words, a SW Request that contains an error is completely ignored by the SW-PC (beyond sending a PA-TNC Error attribute, and possibly logging the error locally) rather than being partially executed.

In the case where the SW-PC receives a valid SW Request attribute but experiences an error during the process of responding to that attribute's instructions where that error prevents the SW-PC from properly or completely fulfilling that request, the SW-PC MUST send a PA-TNC Error attribute with an error code appropriate to the nature of the error. In the case where a PA-TNC Error attribute is sent, the SW-PC MUST NOT take any of the actions requested by the SW Request attribute which led to the detected error. This is the case even if some actions could have been completed successfully, and might even require the SW-PC to reverse some successful actions already taken before the error condition was detected. In other words, either all aspects of a SW Request complete fully and successfully (in which case the SW-PC sends a SW Response attribute), or no aspects of the SW Request occur (in which case the SW-PC sends a PA-TNC Error attribute). In the case that a SW-PC sends a PA-TNC Error attribute in response to a SW Request then the SW-PC MUST NOT also send any SW Response attribute in response to the same SW Request. For this reason, the sending of a SW Response attribute MUST be the last action taken by a SW-PC in response to a SW Request to avoid the possibility of a processing error occurring after that SW Response attribute is sent.

In the case that the SW-PC detects an error that prevents it from properly or completely fulfilling its obligations under an active subscription, the SW-PC MUST send a PA-TNC Error attribute of type `SW_SUBSCRIPTION_FULFILLMENT_ERROR` to the SW-PV that established this subscription. This type of PA-TNC Error attribute identifies the specific subscription that cannot be adequately honored due to the error condition as well as an error "sub-type". The error sub-type is used to indicate the type of error condition the SW-PC experienced that prevented it from honoring the given subscription. In the case that the error condition cannot be identified or does not align with any of the defined error codes, the `SW_ERROR` error code SHOULD be used in the sub-type. In the case that a `SW_SUBSCRIPTION_FULFILLMENT_ERROR` is sent, the associated subscription MUST be treated as cancelled by both the SW-PC and SW-PV.

The SW-PV MUST NOT send any PA-TNC Error attributes to SW-PCs. In the case that a SW-PV detects an error condition, it SHOULD log this error but does not inform any SW-PC's of this event. Errors might include, but are not limited to, detection of malformed SW Response attributes sent from a given SW-PC, as well as detection of error conditions when the SW-PV processes SW Responses.

Both SW-PCs and SW-PVs SHOULD log errors so that administrators can trace the causes of errors. Log entries SHOULD include the type of the error, the time it was detected, and additional descriptive information to aid in understanding the nature and cause of the error.

4. Software Inventory Message and Attributes for PA-TNC Protocol

This section describes the format and semantics of the Software Inventory Message and Attributes for PA-TNC protocol. Software Inventory Message and Attributes for PA-TNC uses the standard PA-TNC message header format. See the PA-TNC specification [RFC5792] for information on this header format.

4.1. PA Subtype (AKA PA-TNC Component Type)

The NEA PB-TNC interface provides a general message-batching protocol capable of carrying one or more PA-TNC messages between the Posture Broker Client and Posture Broker Server. When PB-TNC is carrying a PA-TNC message, the PB-TNC message headers contain a 32 bit identifier called the PA Subtype. The PA Subtype field indicates the type of component associated with all of the PA-TNC attributes carried by the PB-TNC message. The core set of PA Subtypes is defined in the PA-TNC specification. This specification adds the following enumeration element to the IANA registry defined in section

7.2 of the PA-TNC specification [RFC5792] using the IETF Standard name space (SMI Private Enterprise Number 0x000000):

PEN	Integer	Name	Defining Specification
0	9	SW Attributes	Software Inventory Message and Attributes for PA-TNC

Table 2: PA Subtype

Each PA-TNC attribute described in this specification is intended to be sent between the SW-PC and SW-PV, so will be carried in a PB-TNC message indicating a PA Subtype of SW Attributes. Note that although the PA-TNC Error attribute is defined in the PA-TNC specification, when it is used in a SW Attribute exchange, it uses the SW Attributes Component Definition Value, as described in Section 4.2.8 of the PA-TNC specification [RFC5792]. PB-TNC messages MUST always include the SW Attributes Subtype defined in this section when carrying SW Attributes over PA-TNC.

For more information on PB-TNC and PA-TNC messages and message headers, see the PB-TNC [RFC5793] and PA-TNC [RFC5792] specifications, respectively.

4.2. SW Attribute Overview

The attributes defined in this specification appear below with a short summary of their purposes. Each attribute is described in greater detail in subsequent sections.

- o SW Request - This attribute is used to request a software inventory or software event list from an endpoint. This attribute might also establish a subscription on the recipient SW-PC. A SW-PC MUST NOT send this attribute.
- o Software Identifier Inventory - This attribute is used to convey an inventory without the inclusion of software inventory evidence records. When a SW-PC receives a SW Request attribute requesting an inventory without software inventory evidence records, the SW-PC MUST send a Software Identifier Inventory attribute (or a PA-TNC Error) in response. This attribute also MAY be sent by the SW-PC in fulfillment of an active subscription. A SW-PV MUST NOT send this attribute.
- o Software Identifier Events - This attribute is used to convey a list of events concerning changes to an endpoint's Software

Inventory Evidence Collection. Reported events do not include software inventory evidence records. When a SW-PC receives a SW Request attribute requesting an event collection without software inventory evidence records, the SW-PC MUST send a Software Identifier Events attribute (or a PA-TNC Error) in response. This attribute also MAY be sent by the SW-PC in fulfillment of an active subscription. A SW-PV MUST NOT send this attribute.

- o Software Inventory - This attribute is used to convey an inventory expressed including software inventory evidence records. When a SW-PC receives a SW Request attribute requesting an inventory including software inventory evidence records, the SW-PC MUST send a Software Inventory attribute (or a PA-TNC Error) in response. This attribute also MAY be sent by the SW-PC in fulfillment of an active subscription. A SW-PV MUST NOT send this attribute.
- o Software Events - This attribute is used to convey a list of events concerning changes to an endpoint's inventory evidence collection. Reported events include software inventory evidence records. When a SW-PC receives a SW Request attribute requesting an event collection including software inventory evidence records, the SW-PC MUST send a Software Events attribute (or a PA-TNC Error) in response. This attribute also MAY be sent by the SW-PC in fulfillment of an active subscription. A SW-PV MUST NOT send this attribute.
- o Subscription Status Request - This attribute is used to request a SW-PC send a summary of all the active subscriptions it has where the requesting party is the subscriber. The SW-PC MUST respond with a Subscription Status Response (or a PA-TNC Error). A SW-PC MUST NOT send this attribute.
- o Subscription Status Response - This attribute is used to convey information about the active subscriptions that a SW-PC has for a given subscriber. A SW-PV MUST NOT send this attribute.
- o PA-TNC Error - This is the standard PA-TNC Error attribute as defined in PA-TNC [RFC5792] and is used to indicate that an error was encountered during a SW Attribute exchange. It MUST be sent by a SW-PC in response to a SW Request in the case where the SW-PC encounters a fatal error (i.e., an error that prevents further processing of an exchange) relating to the attribute exchange. A SW-PV MUST NOT send this attribute. The SW-PC MUST then ignore the erroneous attribute after a PA-TNC Error attribute is sent (i.e., do not attempt to act on an attribute that generated a PA-TNC Error beyond sending the PA-TNC Error). In the case where the SW-PV experiences a fatal error, it MUST ignore the erroneous attribute without sending a PA-TNC Error attribute. It MAY take


other actions in response to the error, such as logging the cause of the error, or even taking actions to isolate the endpoint

Because one of the Software Identifier Inventory, Software Identifier Events, Software Inventory, or Software Events attributes is expected to be sent to a SW-PV in direct response to a SW Request attribute or in fulfillment of an active subscription, those four attribute types are referred to collectively in this document as "SW Response" attributes.

All SW-PVs MUST be capable of sending SW Requests and be capable of receiving and processing all SW Response attributes as well as PA-TNC Error attributes. All SW-PCs MUST be capable of receiving and processing SW Requests and be capable of sending all types of SW Response attributes as well as PA-TNC Error attributes. In other words, both SW-PVs and SW-PCs are required to support their role in exchanges using any of the attribute types defined in this section. SW-PVs MUST ignore any SW Request attributes that they receive. SW-PCs MUST ignore any SW Response attributes or PA-TNC Error attributes that they receive.

4.3. SW Attribute Exchanges

A SW Attribute Exchange is used to provide the SW-PV with a software inventory or event collection from the queried endpoint.



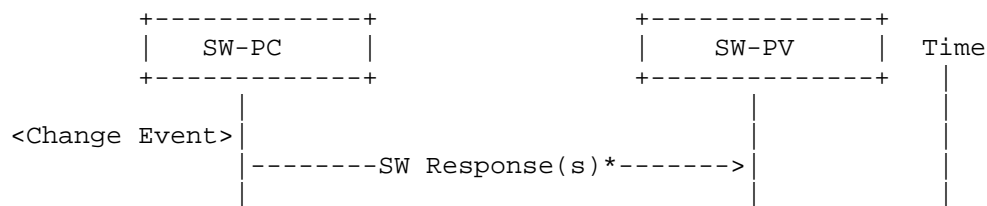
*SW Response is one of the following: Software Identifier Inventory, Software Identifier Events, Software Inventory, or Software Events.

Figure 4: SW Attribute Exchange (Direct Response to SW Request)

In this exchange, the SW-PV indicates to the SW-PC, via a SW Request, the nature of the information it wishes to receive (inventory vs. events, full or targeted) and how it wishes the returned inventory to

be expressed (with or without software inventory evidence records). The SW-PC responds with the requested information using the appropriate attribute type. A single SW Request MUST only lead to a single SW Response or PA-TNC Error that is in direct response to that request.

In addition, if there is an active subscription on the endpoint, the SW-PC sends a SW Response to the SW-PV following a change event on the endpoint in fulfillment of that subscription. Such an exchange is shown in Figure 5.



*SW Response is one of the following: Software Identifier Inventory, Software Identifier Events, Software Inventory, or Software Events.

Figure 5: SW Attribute Exchange (In Fulfillment of an Active Subscription)

Note that, unlike direct responses to a SW Request, a single change event can precipitate multiple SW Responses for a single subscription, but only if all but the last of those SW Responses convey partial lists of event records, and the last of those SW Responses conveys a complete list of event records. (That is, the initial responses are partial lists and the last response is the remainder of the relevant event records, completing the delivery of all relevant events at the time of the change event.) A single Change Event MUST NOT otherwise be followed by multiple SW Response or PA-TNC Error attributes in any combination.

All SW-PVs and SW-PCs MUST support both types of exchanges. In particular, SW-PCs MUST be capable of pushing a SW Response to a SW-PV immediately upon detection of a change to the endpoint's Software Inventory Evidence Collection in fulfillment of established SW-PV subscriptions, as described in Section 3.6.

4.4. Software Inventory Message and Attributes for PA-TNC Attribute Enumeration

PA-TNC attribute types are identified in the PA-TNC Attribute Header via the Attribute Type Vendor ID and Attribute Type fields. Table 3 identifies the appropriate values for these fields for each attribute type used within the Software Inventory Message and Attributes for PA-TNC protocol.

Attribute Name	PEN	Integer	Description
SW Request	0x000000	0x00000011	Request from a SW-PV to a SW-PC for the SW-PC to provide a software inventory or event list
Software Identifier Inventory	0x000000	0x00000012	An inventory sent without software inventory evidence records sent from a SW-PC.
Software Identifier Events	0x000000	0x00000013	A collection of events impacting the endpoint's Software Inventory Evidence Collection, where events do not include software inventory evidence records.
Software Inventory	0x000000	0x00000014	An inventory including software inventory evidence records sent from a SW-PC.
Software Events	0x000000	0x00000015	A collection of events impacting the endpoint's Software Inventory Evidence Collection, where events include software inventory evidence records.
Subscription Status Request	0x000000	0x00000016	A request for a list of a SW-PV's active subscription.
Subscription	0x000000	0x00000017	A list of a SW-PV's active

Status Response			subscriptions.
Reserved	0x000000	0x00000018 - 0x0000001F	These attribute types are reserved for future use in revisions to Software Inventory Message and Attributes for PA-TNC.
PA-TNC Error	0x000000	0x00000008	An error attribute as defined in the PA-TNC specification [RFC5792].

Table 3: SW Attribute Enumeration

4.5. Normalization of Text Encoding

In order to ensure consistency of transmitted attributes some fields require normalization of their format. When this is necessary, this is indicated in the field's description. In such cases, the field contents MUST be normalized to Network Unicode format as defined in RFC 5198 [RFC5198]. Network Unicode format defines a refinement of UTF-8 that ensures a normalized expression of characters. SW-PCs and SW-PVs MUST NOT perform conversion and normalization on any field values except those specifically identified as requiring normalization in the following sections. Note, however, that some data models require additional normalization before source information is added to an Endpoint's Inventory Evidence Collection as a record. The references from the Software Data Model IANA table (see Section 9.4) will note where this is necessary.

4.6. Request IDs

All SW Request attributes MUST include a Request ID value. The Request ID field provides a value that identifies a given request relative to other requests between a SW-PV and the receiving SW-PC. Specifically, the SW-PV assigns each SW Request attribute a Request ID value that is intended to be unique within the lifetime of a given network connection ID as assigned by the SW-PV's Posture Broker Server. In the case where all possible Request ID values have been exhausted within the lifetime of a single network connection ID, the sender MAY reuse previously used Request IDs within the same network connection that are not being used as Subscription IDs. (See below in this section for an explanation of Subscription ID assignment.) In this case of Request ID reuse, Request IDs SHOULD be reused in the order of their original use. In other words, a SW-PC SHOULD NOT use a given Request ID value more than once within a persistent

connection between a given Posture Broker Client-Posture Broker Server pair, but, in the case where reuse is necessary due to exhaustion of possible ID values, the SW-PC SHOULD structure the reuse to maximize the time between original and subsequent use. The Request ID value is included in a SW Response attribute directly responding to this SW Request to indicate which SW Request was received and caused the response. Request IDs can be randomly generated or sequential, as long as values are not repeated per the rules in this paragraph. SW-PCs are not required to check for duplicate Request IDs.

In the case that a SW Request requests the establishment of a subscription and the receiving SW-PC agrees to that subscription, the Request ID of that SW Request (i.e., the establishing request of the subscription) becomes that subscription's Subscription ID. All attributes sent in fulfillment of this subscription include a flag indicating that the attribute fulfills a subscription and the subscription's Subscription ID. A SW-PV MUST NOT reuse a Request ID value in communicating to a given SW-PC while that Request ID is also serving as a Subscription ID for an active subscription with that SW-PC. In the case where a SW-PC receives a SW Request from a given SW-PV where that Request ID is also the Subscription ID of an active subscription with that SW-PV, the SW-PC MUST respond with a PA-TNC Error attribute with an error code of SW_SUBSCRIPTION_ID_REUSE_ERROR. Note that this error does not cancel the indicated subscription.

Subscription Status Requests and Subscription Status Responses do not include Request IDs.

4.7. SW Request

A SW-PV sends this attribute to a SW-PC to request that the SW-PC send software inventory information to the SW-PV. A SW-PC MUST NOT send this attribute.

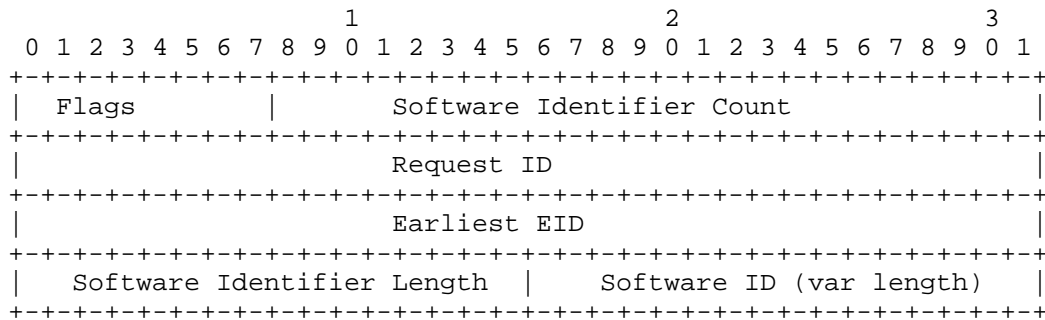


Figure 6: SW Request Attribute

Field	Description
Flags: Bit 0 - Clear Subscriptions	If set (1), the SW-PC MUST delete all subscriptions established by the requesting SW-PV (barring any errors).
Flags: Bit 1 - Subscribe	If set (1), in addition to responding to the request as described, the SW-PC MUST establish a subscription with parameters matching those in the request attribute (barring any errors).
Flags: Bit 2 - Result Type	If unset (0), the SW-PC's response MUST include software inventory evidence records and thus the response MUST be a Software Inventory, a Software Events, or a PA-TNC Error attribute. If set (1), the response MUST NOT include software inventory evidence records and thus the response MUST be a Software Identifier Inventory, a Software Identifier Events, or a PA-TNC Error attribute.
Flags: Bit 3-7 - Reserved	Reserved for future use. This field MUST be set to zero on transmission and ignored upon reception.
Software Identifier Count	A 3-byte unsigned integer indicating the number of Software Identifiers that follow. If this value is non-zero, this is a targeted request, as described in Section 3.3. The Software Identifier Length and Software ID fields are repeated, in order, the number of times indicated in this field. In the case where Software

	Identifiers are present, the SW-PC MUST only report software that corresponds to the identifiers the SW-PV provided in this attribute (or with a PA-TNC Error attribute). This field value MAY be 0, in which case there are no instances of the Software Identifier Length and Software ID fields. In this case, the SW-PV is indicating an interest in all software inventory evidence records on the endpoint (i.e., this is not a targeted request).
Request ID	A value that uniquely identifies this SW Request from a particular SW-PV.
Earliest EID	In the case where the SW-PV is requesting software events, this field contains the EID value of the earliest event the SW-PV wishes to have reported. (Note - the report will be inclusive of the event with this EID value.) In the case where the SW-PV is requesting an inventory, then this field MUST be 0. (0x00000000) In the case where this field is non-zero, the SW-PV is requesting events and the SW-PC MUST respond using a Software Events, Software Identifier Events, or a PA-TNC Error attribute. In the case where this field is zero, the SW-PV is requesting an inventory and the SW-PC MUST respond using a Software Inventory, a Software Identifier Inventory, or a PA-TNC Error attribute.
Software Identifier Length	A 2-byte unsigned integer indicating the length in bytes of the Software ID field.
Software ID	A string containing the Software Identifier value from a software inventory evidence record. This field value MUST be normalized to Network Unicode format, as described in Section 4.5. This string MUST NOT be NULL terminated.

Table 4: SW Request Attribute Fields

The SW-PV sends the SW Request attribute to a SW-PC to request the indicated information. Note that between the Result Type flag and the Earliest EID field, the SW-PC is constrained to a single possible

SW Response attribute type (or a PA-TNC Error attribute) in its response to the request.

The Subscribe and Clear Subscription flags are used to manage subscriptions for the requesting SW-PV on the receiving SW-PC. Specifically, an attribute with the Subscribe flag set seeks to establish a new subscription by the requesting SW-PV to the given SW-PC, while an attribute with the Clear Subscription flag seeks to delete all existing subscriptions by the requesting SW-PV on the given SW-PC. Note that, in the latter case, only the subscriptions associated with the Connection ID and the Posture Validator ID of the requester are deleted as described in Section 3.6.3. A newly established subscription has the parameters outlined in the Request attribute. Specifically, the Result Type flag indicates the type of result to send in fulfillment of the subscription, the value of the Earliest EID field indicates whether the fulfillment attributes list inventories or events, and the fields describing Software Identifiers (if present) indicate if and how a subscription is targeted. In the case that the SW-PC is unable or unwilling to comply with the SW-PV's request to establish or clear subscriptions, the SW-PC MUST respond with a PA-TNC Error attribute with the `SW_SUBSCRIPTION_DENIED_ERROR` error code. (Note that if the SW-PV requests that subscriptions be cleared but has no existing subscriptions, this is not an error.)

An attribute requesting the establishment of a subscription is effectively doing double-duty, as it is a request for an immediate response from the SW-PC in addition to setting up the subscription. Assuming the SW-PC is willing to comply with the subscription it MUST send an appropriate response attribute to a request with the Subscribe flag set containing all requested information. The same is true of the Clear Subscription flag - assuming there is no error the SW-PC MUST generate a response attribute without regard to the presence of this flag in addition to clearing its subscription list.

Both the Subscribe and Clear Subscription flags MAY be set in a single SW Request attribute. In the case where this request is successful, the end result MUST be equivalent to the SW-PC clearing its subscription list for the given SW-PV first and then creating a new subscription in accordance with the request parameters. (In other words, do not first create the new subscription and then clear all the subscriptions including the one that was just created.) In the case that the requested actions are successfully completed, the SW-PC MUST respond with a SW Response attribute. (The specific type of SW Response attribute depends on the Result Type and Earliest EID fields, as described above.) In the case where there is a failure that prevents some part this request from completing, the SW-PC MUST NOT add a new subscription, MUST NOT clear the old subscriptions, and the SW-PC MUST respond with a PA-TNC Error attribute. In other

words, the SW-PC MUST NOT partially succeed at implementing such a request; either all actions succeed, or none succeed.

The Earliest EID field is used to indicate whether the SW-PV is requesting an inventory or event list from the SW-PC. A value of 0 (0x00000000) represents a request for inventory information. Otherwise, the SW-PV is requesting event information. For Earliest EID values other than 0, the SW-PC's response MUST respond with event records, as described in Section 3.5. Note that the request does not identify a particular EID Epoch, since responses can only include events in the SW-PC's current EID Epoch.

The Software Identifier Count indicates the number of Software Identifiers in the attribute. This number might be any value between 0 and 16,777,216, inclusive. A single Software Identifier is represented by the following fields: Software Identifier Length and Software ID. These fields are repeated a number of times equal to the Software Identifier Count. Note that this could be 0 times. The Software Identifier Length field indicates the number of bytes allocated to the Software ID field. The Software Identifier field contains a Software Identifier as describe in Section 3.2.1. The presence of one or more Software Identifiers is used by the SW-PV to indicate a targeted request, which seeks only inventories of or events affecting software corresponding to the given identifiers. The SW-PC MUST only report software that matched the Software Identifiers provided in the SW-PVs SW Request attribute.

4.8. Software Identifier Inventory

A SW-PC sends this attribute to a SW-PV to convey the inventory of the endpoint's Software Inventory Evidence Collection without the inclusion of software inventory evidence records. This list might represent a complete inventory or a targeted list of records, depending on the parameters in the SW-PV's request. A SW-PV MUST NOT send this attribute. The SW-PC either sends this attribute in fulfillment of an existing subscription where the establishing request has a Result Type of 1 and the Earliest EID is zero, or in direct response to a SW Request attribute where the Result Type is 1 and the Earliest EID is zero.

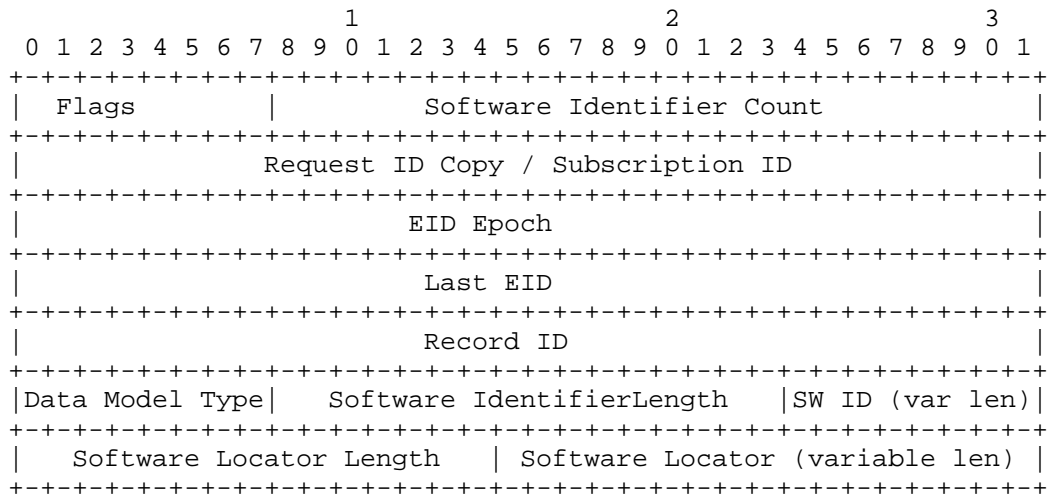


Figure 7: Software Identifier Inventory Attribute

Field	Description
Flags: Bit 0 - Subscription Fulfillment	In the case that this attribute is sent in fulfillment of a subscription this bit MUST be set (1). In the case that this attribute is a direct response to a SW Request, this bit MUST be unset (0).
Flags: Bit 1-7 - Reserved	Reserved for future use. This field MUST be set to zero on transmission and ignored upon reception.
Software Identifier Count	The number of Software Identifiers that follow. This field is an unsigned integer. The Record ID, Data Model Type, Software Identifier Length, SW ID, Software Locator Length, and Software Locator fields are repeated, in order, the number of times indicated in this field. This field value MAY be 0, in which case there are no instances of these fields.
Request ID Copy / Subscription ID	In the case where this attribute is in direct response to a SW Request attribute from a SW-PV, this field MUST contain an exact copy of the Request ID field from that SW Request. In the

	case where this attribute is sent in fulfillment of an active subscription, this field MUST contain the Subscription ID of the subscription being fulfilled by this attribute.
EID Epoch	The EID Epoch of the Last EID value. This field is an unsigned 4-byte integer.
Last EID	The EID of the last event recorded by the SW-PC, or 0 if the SW-PC has no recorded events. This field is an unsigned 4-byte integer.
Record ID	A 4-byte, unsigned integer containing the Record Identifier value from a software inventory evidence record.
Data Model Type	A 1-byte unsigned integer containing an identifier number from the Software Data Model IANA table that identifies the data model of the reported record.
Software Identifier Length	A 2-byte unsigned integer indicating the length in bytes of the SW ID field.
SW ID	A string containing the Software Identifier value from a software inventory evidence record. This field value MUST be normalized to Network Unicode format, as described in Section 4.5. This string MUST NOT be NULL terminated.
Software Locator Length	A 2-byte unsigned integer indicating the length in bytes of the Software Locator field.
Software Locator	A string containing the Software Locator value. This is expressed as a URI. This field value MUST be normalized to Network Unicode format as described in Section 3.2.1.2. This string MUST NOT be NULL terminated.

Table 5: Software Identifier Inventory Attribute Fields

In the case that this attribute is sent in fulfillment of a subscription, the Subscription Fulfillment bit MUST be set (1). In the case that this attribute is sent in direct response to a SW Request, the Subscription Fulfillment bit MUST be unset (0). Note

that the SW Response attribute sent in direct response to a SW Request that establishes a subscription (i.e., a subscription's establishing request) MUST be treated as a direct response to that SW Request (and thus the Subscription Fulfillment bit is unset). SW Response attributes are only treated as being in fulfillment of a subscription (i.e., Subscription Fulfillment bit set) if they are sent following a change event, as shown in Figure 3.

The Software Identifier Count field indicates the number of Software Identifiers present in this inventory. Each Software Identifier is represented by the following set of fields: Record ID, Data Model Type, Software Identifier Length, SW ID, Software Locator Length, and Software Locator. These fields will appear once for each reported record.

When responding directly to a SW Request attribute, the Request ID Copy / Subscription ID field MUST contain an exact copy of the Request ID field from that SW Request. When this attribute is sent in fulfillment of an existing subscription on this Posture Collector, then this field MUST contain the Subscription ID of the fulfilled subscription.

The EID Epoch field indicates the EID Epoch of the Last EID value. The Last EID field MUST contain the EID of the last recorded change event (see Section 3.5 for more about EIDs and recorded events) at the time this inventory was collected. In the case where there are no recorded change events at the time that this inventory was collected, the Last EID field MUST contain 0. These fields can be interpreted to indicate that the provided inventory reflects the state of the endpoint after all changes up to and including this last event have been accounted for.

4.9. Software Identifier Events

A SW-PC sends this attribute to a SW-PV to convey events where the affected records are reported without software inventory evidence records. A SW-PV MUST NOT send this attribute. The SW-PC either sends this attribute in fulfillment of an existing subscription where the establishing request has a Result Type is 1 and the Earliest EID is non-zero, or in direct response to a SW Request attribute where the Result Type is 1 and the Earliest EID is non-zero.

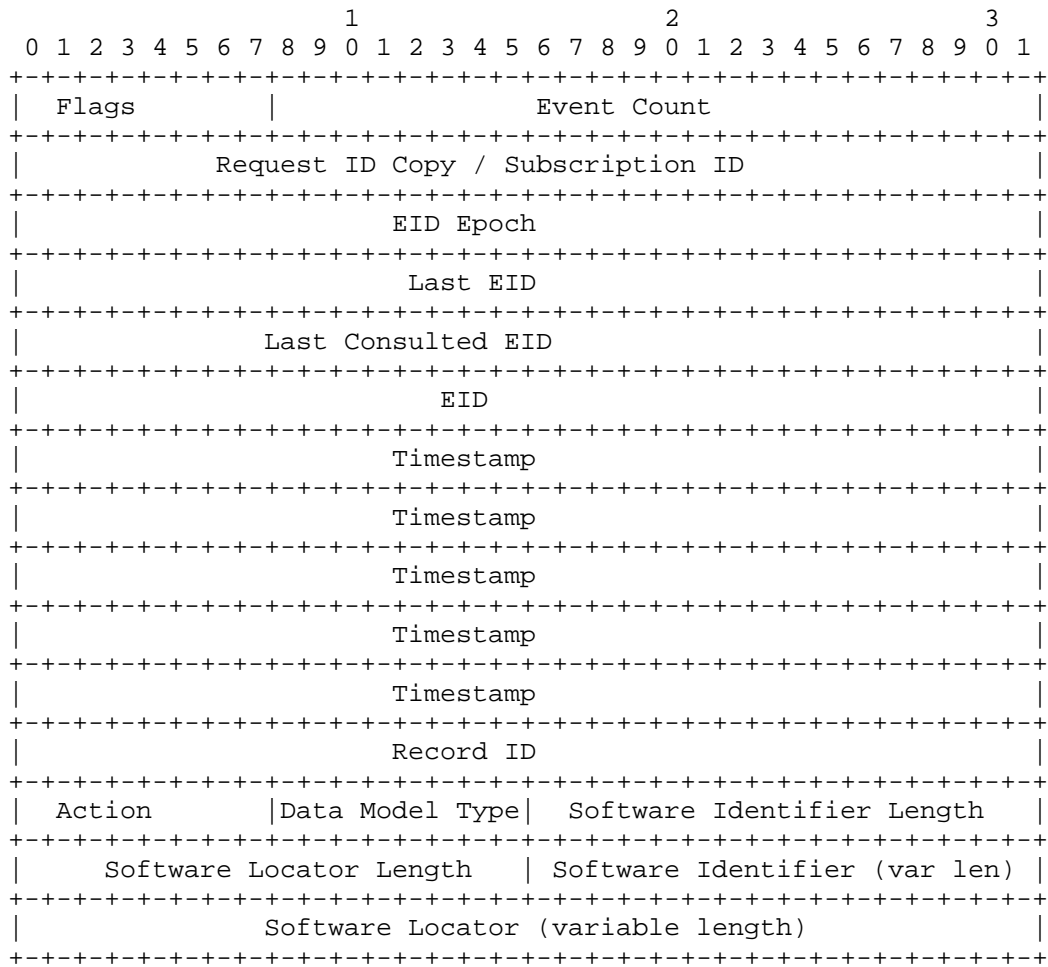


Figure 8: Software Identifier Events Attribute

Field	Description
Flags: Bit 0 - Subscription Fulfillment	In the case that this attribute is sent in fulfillment of a subscription this bit MUST be set (1). In the case that this attribute is a direct response to a SW Request, this bit MUST be unset (0).
Flags: Bit 1-7 -	Reserved for future use. This field MUST be set to zero on transmission and ignored upon reception.

Reserved	
Event Count	The number of events that are reported in this attribute. This field is a 3-byte unsigned integer. The EID, Timestamp, Record ID, Action, Data Model Type, Software Identifier Length, Software Locator Length, Software Identifier, and Software Locator fields are repeated, in order, the number of times indicated in this field. This field value MAY be 0, in which case there are no instances of these fields.
Request ID Copy / Subscription ID	In the case where this attribute is in direct response to a SW Request attribute from a SW-PV, this field MUST contain an exact copy of the Request ID field from that SW Request. In the case where this attribute is sent in fulfillment of an active subscription, this field MUST contain the Subscription ID of the subscription being fulfilled by this attribute.
EID Epoch	The EID Epoch of the Last EID value. This field is an unsigned 4-byte integer.
Last EID	The EID of the last event recorded by the SW-PC, or 0 if the SW-PC has no recorded events. This field contains the EID of the SW-PC's last recorded change event (which might or might not be included as an event record in this attribute).
Last Consulted EID	The EID of the last event record that was consulted when generating the event record list included in this attribute. This is different from the Last EID field value if and only if this attribute is conveying a partial list of event records. See Section 3.5.5 for more on partial list of event records.
EID	The EID of the event in this event record.
Timestamp	The timestamp associated with the event in this event record. This timestamp is the SW-PC's best understanding of when the given event occurred. Note that this timestamp might be an estimate. The Timestamp date and time MUST be represented as an RFC 3339 [5] compliant ASCII string in Coordinated Universal Time (UTC) time with the additional restrictions that the 'T' delimiter and

	the 'Z' suffix MUST be capitalized and fractional seconds (time-secfrac) MUST NOT be included. This field conforms to the date-time ABNF production from section 5.6 of RFC 3339 [RFC3339] with the above restrictions. Leap seconds are permitted and SW-PVs MUST support them. The Timestamp string MUST NOT be NULL terminated or padded in any way. The length of this field is always 20 octets.
Record ID	A 4-byte, unsigned integer containing the Record Identifier value from a software inventory evidence record.
Action	The type of event that is recorded in this event record. Possible values are: 1 = CREATION - the addition of a record to the endpoint's Software Inventory Evidence Collection; 2 = DELETION - the removal of a record from the endpoint's Software Inventory Evidence Collection; 3 = ALTERATION - There was an alteration to a record within the endpoint's Software Inventory Evidence Collection. All other values are reserved for future use and MUST NOT be used when sending attributes. In the case where a SW-PV receives an event record that uses an action value other than the ones defined here, it MUST ignore that event record but SHOULD process other event records in this attribute as normal.
Data Model Type	A 1-byte unsigned integer containing an identifier number from the Software Data Model IANA table that identifies the data model of the reported record.
Software Identifier Length	A 2-byte unsigned integer indicating the length in bytes of the Software Identifier field.
Software Locator Length	A 2-byte unsigned integer indicating the length in bytes of the Software Locator field.
Software Identifier	A string containing the Software Identifier value from a software inventory evidence record. This field value MUST be normalized to Network Unicode format, as described in Section 4.5. This string MUST NOT be NULL terminated.

Software Locator	A string containing the Software Locator value. This is expressed as a URI. This field value MUST be normalized to Network Unicode format as described in Section 3.2.1.2. This string MUST NOT be NULL terminated.
------------------	---

Table 6: Software Identifier Events Attribute Fields

The first few fields in the Software Identifier Events attribute mirror those in the Software Identifier Inventory attribute. The primary difference is that, instead of conveying an inventory, the attribute conveys zero or more event records, consisting of the EID, timestamp, Record ID, action type, data model type, Software Identifier, and Software Locator of the affected software inventory evidence record.

With regard to the Timestamp field, it is important to note that clock skew between the SW-PC and SW-PV as well as between different SW-PCs within an enterprise might make correlation of timestamp values difficult. This specification does not attempt to resolve clock skew issues, although other mechanisms outside of this specification do exist to reduce the impact of clock skew and make the timestamp more useful for such correlation. Instead, Software Inventory Message and Attributes for PA-TNC uses Timestamp value primarily as a means to indicate the amount of time between two events on a single endpoint. For example, by taking the difference of the times for when a record was removed and then subsequently re-added, one can get an indication as to how long the system was without the given record (and, thus without the associated software). Since this will involve comparison of timestamp values all originating on the same system, clock skew between the SW-PC and SW-PV is not an issue. However, if the SW-PC's clock was adjusted between two recorded events, it is possible for such a calculation to lead to incorrect understandings of the temporal distance between events. Users of this field need to be aware of the possibility for such occurrences. In the case where the Timestamp values of two events appear to contradict the EID ordering of those events (i.e., the later EID has an earlier timestamp) the recipient MUST treat the EID ordering as correct.

All events recorded in a Software Identifier Events Attribute are required to be part of the same EID Epoch. Specifically, all reported events MUST have an EID from the same EID Epoch as each other, and which is the same as the EID Epoch of the Last EID and Last Consulted EID values. The SW-PC MUST NOT report events with EIDs from different EID Epochs.

The Last Consulted EID field contains the EID of the last event record considered for inclusion in this attribute. If this attribute contains a partial event set (as described in Section 3.5.5) this field value will be less than the Last EID value; if this attribute contains a complete event set, the Last EID and Last Consulted EID values are identical.

If multiple events are sent in a Software Identifier Events attribute, the order in which they appear within the attribute is not significant. The EIDs associated with them are used for ordering the indicated events appropriately. Also note that a single software record might be reported multiple times in an attribute, such as if multiple events involving the associated record were being reported.

4.10. Software Inventory

A SW-PC sends this attribute to a SW-PV to convey a list of inventory records. A SW-PV MUST NOT send this attribute. The SW-PC either sends this attribute in fulfillment of an existing subscription where the establishing request had a Result Type of 0 and the Earliest EID is zero, or in direct response to a SW Request attribute where the Result Type is 0 and the Earliest EID is zero.




Figure 9: Software Inventory Attribute

Field	Description
Flags: Bit 0 - Subscription Fulfillment	In the case that this attribute is sent in fulfillment of a subscription this bit MUST be set (1). In the case that this attribute is a direct response to a SW Request, this bit MUST be unset (0).
Flags: Bit 1-7 - Reserved	Reserved for future use. This field MUST be set to zero on transmission and ignored upon reception.
Record Count	The number of records that follow. This field is a 3-byte unsigned integer. The Record ID, Software Identifier Length, Software Locator Length, Record Length, Data Model Type, Software Identifier, Software Locator, and Record fields are repeated, in order, the number of times indicated in this field. This field value MAY be 0 in which case there are no instances of these fields.

Request ID Copy / Subscription ID	In the case where this attribute is in direct response to a SW Request attribute from a SW-PV, this field MUST contain an exact copy of the Request ID field from that SW Request. In the case where this attribute is sent in fulfillment of an active subscription, this field MUST contain the Subscription ID of the subscription being fulfilled by this attribute.
EID Epoch	The EID Epoch of the Last EID value. This field is an unsigned 4-byte integer.
Last EID	The EID of the last event recorded by the SW-PC, or 0 if the SW-PC has no recorded events. This field is an unsigned 4-byte integer.
Record ID	A 4-byte, unsigned integer containing the Record Identifier value from a software inventory evidence record.
Software Identifier Length	A 2-byte unsigned integer indicating the length in bytes of the Software Identifier field.
Software Locator Length	A 2-byte unsigned integer indicating the length in bytes of the Software Locator field.
Record Len	This is a 4-byte unsigned integer indicating the length of the Record field in bytes.
Data Model Type	A 1-byte unsigned integer containing an identifier number from the Software Data Model IANA table that identifies the data model of the reported record.
Software Identifier	A string containing the Software Identifier value from a software inventory evidence record. This field value MUST be normalized to Network Unicode format, as described in Section 4.5. This string MUST NOT be NULL terminated.
Software Locator	A string containing the Software Locator value. This is expressed as a URI. This field value MUST be normalized to Network Unicode format as described in Section 3.2.1.2. This string MUST NOT be NULL terminated.

Record	A software inventory evidence record as a string. The record MUST be converted and normalized to Network Unicode format, as described in Section 4.5. This string MUST NOT be NULL terminated.
--------	--

Table 7: Software Inventory Attribute Fields

The Software Inventory attribute contains some number of software inventory evidence records along with the core response attribute fields. Given that the size of records can vary considerably, the length of this attribute is highly variable and, if transmitting a complete inventory, can be extremely large. Enterprises might wish to constrain the use of Software Inventory attributes to targeted requests to avoid over-burdening the network unnecessarily.

When copying a software inventory evidence record into the Record field, the record MUST be converted and normalized to use Network Unicode format prior to its inclusion in the record field.

4.11. Software Events

A SW-PC sends this attribute to a SW-PV to convey a list of events that include software inventory evidence records. A SW-PV MUST NOT send this attribute. The SW-PC either sends this attribute in fulfillment of an existing subscription where the establishing request has a Result Type of 0 and the Earliest EID is non-zero, or in direct response to a SW Request attribute where the Result Type is 0 and the Earliest EID is non-zero.

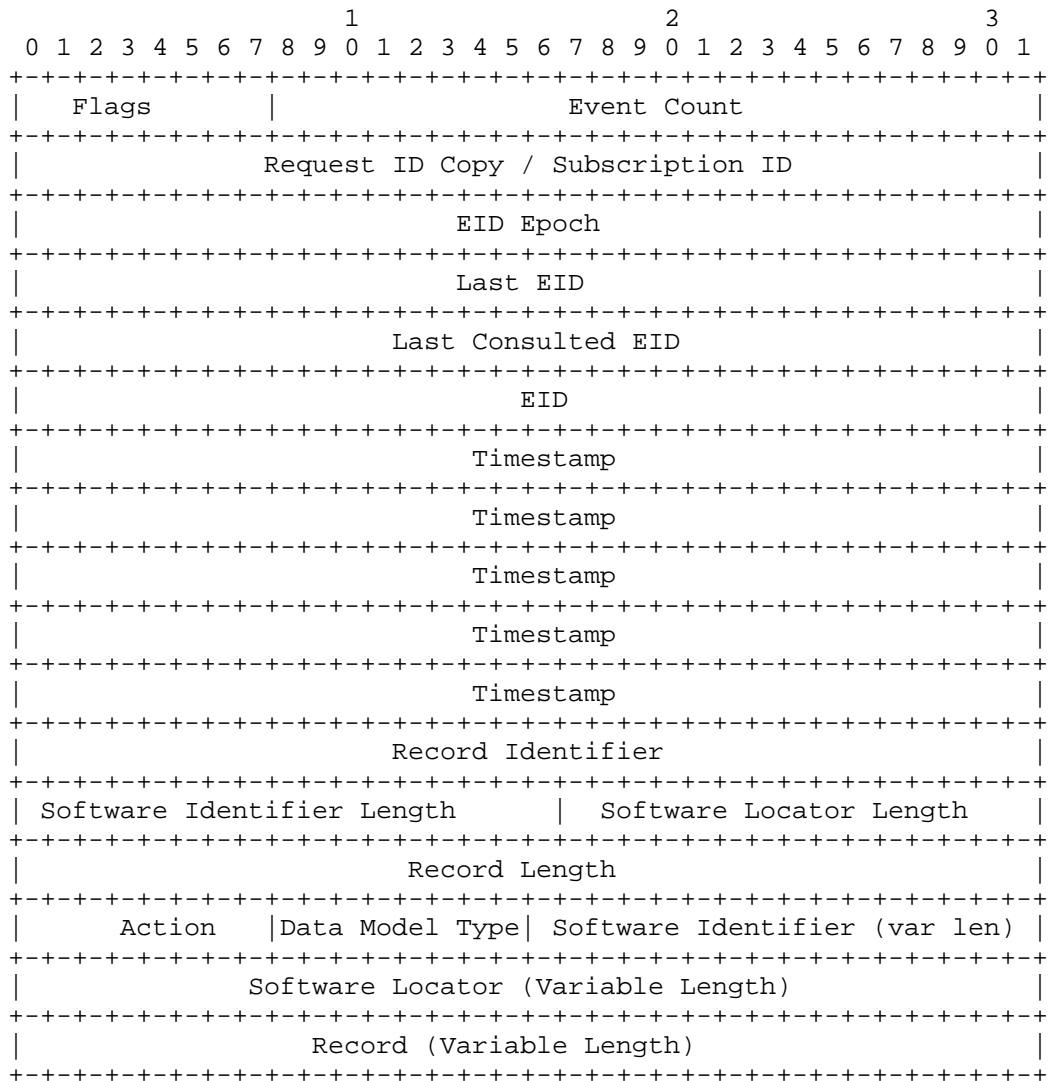


Figure 10: Software Events Attribute

Field	Description
Flags: Bit 0	In the case that this attribute is sent in fulfillment of a subscription this bit MUST be set
- Subscription	(1). In the case that this attribute is a direct response to a SW Request, this bit MUST be unset
Fulfillment	

	(0).
Flags: Bit 1-7 - Reserved	Reserved for future use. This field MUST be set to zero on transmission and ignored upon reception.
Event Count	The number of events being reported in this attribute. This field is a 3-byte unsigned integer. The EID, Timestamp, Record Identifier, Software Identifier Length, Software Locator Length, Record Length, Action, Data Model Type, Software Identifier, Software Locator, and Record fields are repeated, in order, the number of times indicated in this field. This field value MAY be 0, in which case there are no instances of these fields.
Request ID Copy / Subscription ID	In the case where this attribute is in direct response to a SW Request attribute from a SW-PV, this field MUST contain an exact copy of the Request ID field from that SW Request. In the case where this attribute is sent in fulfillment of an active subscription, this field MUST contain the Subscription ID of the subscription being fulfilled by this attribute.
EID Epoch	The EID Epoch of the Last EID value. This field is an unsigned 4-byte integer.
Last EID	The EID of the last event recorded by the SW-PC, or 0 if the SW-PC has no recorded events. This field contains the EID of the SW-PC's last recorded change event (which might or might not be included as an event record in this attribute).
Last Consulted EID	The EID of the last event record that was consulted when generating the event record list included in this attribute. This is different from the Last EID field value if and only if this attribute is conveying a partial list of event records. See Section 3.5.5 for more on partial list of event records.
EID	The EID of the event in this event record.
Timestamp	The timestamp associated with the event in this event record. This timestamp is the SW-PC's best understanding of when the given event occurred.

	Note that this timestamp might be an estimate. The Timestamp date and time MUST be represented as an RFC 3339 [5] compliant ASCII string in Coordinated Universal Time (UTC) time with the additional restrictions that the 'T' delimiter and the 'Z' suffix MUST be capitalized and fractional seconds (time-secfrac) MUST NOT be included. This field conforms to the date-time ABNF production from section 5.6 of RFC 3339 [RFC3339] with the above restrictions. Leap seconds are permitted and SW-PVs MUST support them. The Timestamp string MUST NOT be NULL terminated or padded in any way. The length of this field is always 20 octets.
Record Identifier	A 4-byte, unsigned integer containing the Record Identifier value from a software inventory evidence record.
Software Identifier Length	A 2-byte unsigned integer indicating the length in bytes of the Software Identifier field.
Software Locator Length	A 2-byte unsigned integer indicating the length in bytes of the Software Locator field.
Record Len	This is a 4-byte unsigned integer indicating the length of the Record field in bytes.
Action	The type of event that is recorded in this event record. Possible values are: 1 = CREATION - the addition of a record to the endpoint's Software Inventory Evidence Collection; 2 = DELETION - the removal of a record from the endpoint's Software Inventory Evidence Collection; 3 = ALTERATION - There was an alteration to a record within the endpoint's Software Inventory Evidence Collection. All other values are reserved for future use and MUST NOT be used when sending attributes. In the case where a SW-PV receives an event record that uses an action value other than the ones defined here, it MUST ignore that event record but SHOULD process other event records in this attribute as normal.
Data Model Type	A 1-byte unsigned integer containing an identifier number from the Software Data Model IANA table

	that identifies the data model of the reported record.
Software Identifier	A string containing the Software Identifier value from a software inventory evidence record. This field value MUST be normalized to Network Unicode format, as described in Section 4.5. This string MUST NOT be NULL terminated.
Software Locator	A string containing the Software Locator value. This is expressed as a URI. This field value MUST be normalized to Network Unicode format as described in Section 3.2.1.2. This string MUST NOT be NULL terminated.
Record	A software inventory evidence record as a string. The record MUST be converted and normalized to Network Unicode format, as described in Section 4.5. This string MUST NOT be NULL terminated.

Table 8: Software Events Attribute Fields

The fields of this attribute are used in the same way as the corresponding fields of the previous attributes. As with the Software Inventory attribute, a Software Events attribute can be quite large if many events have occurred following the event indicated by a request's Earliest EID. As such, it is recommended that the SW Request attributes only request full records be sent (Result Type set to 0) in a targeted request, thus constraining the response just to records that match a given set of Software Identifiers.

As with the Software Identifier Events Attribute, this attribute MUST only contain event records with EIDs coming from the current EID Epoch of the SW-PC.

As with the Software Inventory Attribute, the SW-PC MUST perform conversion and normalization of the record.

4.12. Subscription Status Request

A SW-PV sends this attribute to a SW-PC to request a list of active subscriptions for which the requesting SW-PV is the subscriber. A SW-PC MUST NOT send this attribute.

This attribute has no fields.

A SW-PC MUST respond to this attribute by sending a Subscription Status Response attribute (or a PA-TNC Error attribute if it is unable to correctly provide a response).

4.13. Subscription Status Response

A SW-PC sends this attribute to a SW-PV to report the list of active subscriptions for which the receiving SW-PV is the subscriber. A SW-PV MUST NOT send this attribute.

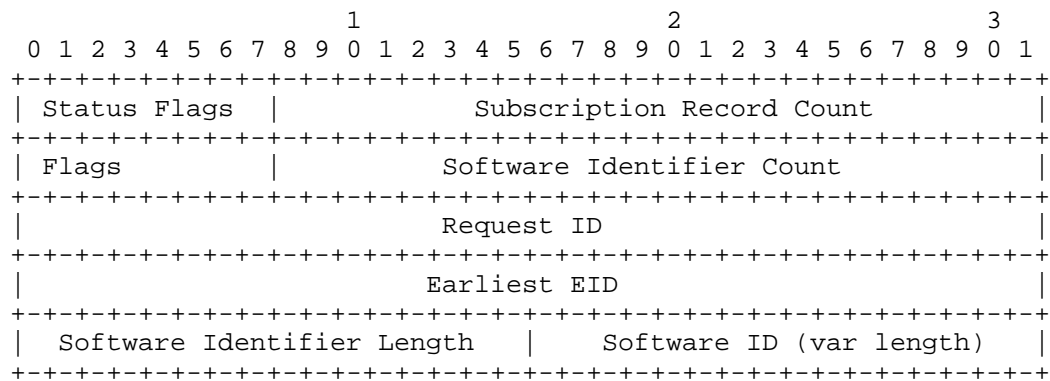


Figure 11: Subscription Status Response Attribute

Field	Description
Flags: Bit 0-7 - Reserved	Reserved for future use. This field MUST be set to zero on transmission and ignored upon reception.
Subscription Record Count	The number of subscription records that follow. This field is a 3-byte unsigned integer. The Flags, Software Identifier Count, Request ID, Earliest EID, Software Identifier Length, and Software ID fields are repeated, in order, the number of times indicated in this field. This field value MAY be 0 in which case there are no instances of these fields.
Flags, Software Identifier Count, Request ID, Earliest EID, Software Identifier Length, and Software ID	For each active subscription, these fields contain an exact copy of the fields with the same name as provided in the subscription's establishing request.

Table 9: Subscription Status Response Fields

A Subscription Status Response contains zero or more subscription records. Specifically, it MUST contain one subscription record for each active subscription associated with the party that sent the Subscription Status Request to which this attribute is a response. As described in Section 3.6.2, the SW-PC MUST use the requester's Connection ID and its Posture Validator ID to determine which subscriptions are associated with the requester.

A SW-PC MUST send a Subscription Status Response attribute in response to a Subscription Status Request attribute. The only exception to this is if the SW-PC experiences an error condition that prevents it from correctly populating the Subscription Status Response attribute, in which case it MUST respond with a PA-TNC Error attribute appropriate to the type of error experienced. If there are no active subscriptions associated with the requesting party, the Subscription Status Response attribute will consist of its Status Flags field, a Subscription Record Count field with a value of 0, and no additional fields.

Each subscription record included in a Subscription Status Response attribute duplicates the fields of a SW Request attribute that was the establishing request of a subscription. Note that the Request ID field in the record captures the Subscription ID associated with the given subscription record (since the Subscription ID is the same as the Request ID of the establishing request). Note also that if the establishing request is targeted, then its Record Count field will be non-zero and, within that subscription record, the Software Identifier Length and Software Identifier fields are repeated, in order, the number of times indicated in the Record Count field. As such, each subscription record can be different sizes. If the establishing request is not targeted (Record Count field is 0), the subscription record has no Software Identifier Length or Software Identifier fields.

When a SW-PV compares the information received in a Subscription Status Response to its own records of active subscriptions it should be aware that the SW-PC might be unable to distinguish this SW-PV from other SW-PVs on the same NEA Server. As a result, it is possible that the SW-PC will report more subscription records than the SW-PV recognizes. For this reason, SW-PVs SHOULD NOT automatically assume that extra subscriptions reported in a Subscription Status Response indicate a problem.

4.14. PA-TNC Error as Used by Software Inventory Message and Attributes for PA-TNC

The PA-TNC Error attribute is defined in the PA-TNC specification [RFC5792], and its use here conforms to that specification. A PA-TNC Error can be sent due to any error in the PA-TNC exchange and might also be sent in response to error conditions specific to the Software Inventory Message and Attributes for PA-TNC exchange. The latter case utilizes error codes defined below.

A PA-TNC Error attribute is sent instead of a SW Response attribute due to factors that prevent the reliable creation of a SW Response. As such, a SW-PC MUST NOT send both a PA-TNC Error attribute and a SW Response attribute in response to a single SW Request attribute.

Table 10 lists the Error Code values for the PA-TNC Error attribute specific to the Software Inventory Message and Attributes for PA-TNC exchange. In all of these cases, the Error Code Vendor ID field MUST be set to 0x000000, corresponding to the IETF SMI Private Enterprise Number. The Error Information structures for each error type are described in the following subsections.

Note that a message with a Software Inventory Message and Attributes for PA-TNC attribute might also result in an error condition covered

by the Standard PA-TNC Error Codes defined in PA-TNC. For example, a SW Attribute might have an invalid parameter, leading to an error code of "Invalid Parameter". In this case, the SW-PC MUST use the appropriate PA-TNC Error Code value as defined in Section 4.2.8 of PA-TNC specification.

Error Code Value	Description
0x00000020	SW_ERROR. This indicates a fatal error (i.e., an error that precludes the creation of a suitable response attribute) other than the errors described below but still specific to the processing of SW Attributes. The Description field SHOULD contain additional diagnostic information.
0x00000021	SW_SUBSCRIPTION_DENIED_ERROR. This indicates that the SW-PC denied the SW-PV's request to establish a subscription. The Description field SHOULD contain additional diagnostic information.
0x00000022	SW_RESPONSE_TOO_LARGE_ERROR. This indicates that the SW-PC's response to the SW-PV's request was too large to be serviced. The error information structure indicates the largest possible size of a response supported by the SW-PC (see Section 4.14.2). The Description field SHOULD contain additional diagnostic information.
0x00000023	SW_SUBSCRIPTION_FULFILLMENT_ERROR. This indicates that the SW-PC experienced an error fulfilling a given subscription. The error information includes the Subscription ID of the relevant subscription, as well as a sub-error that describes the nature of the error the SW-PC experienced. The SW-PC and SW-PV MUST treat the identified subscription as cancelled.
0x00000024	SW_SUBSCRIPTION_ID_REUSE_ERROR. This indicates that the SW-PC received a SW Request from a given SW-PV where the

	Request ID of that SW Request is currently used as the Subscription ID of an active subscription with that SW-PV. This error does not cancel the identified subscription.
0x00000025-0x0000002F	RESERVED. These Error Codes are reserved for use by future revisions of the Software Inventory Message and Attributes for PA-TNC specification. Any PA-TNC Error attribute using one of these Error Codes MUST be treated as indicating a fatal error on the sender without further interpretation.

Table 10: PA-TNC Error Codes for Software Inventory Message and Attributes for PA-TNC

The following subsections describe the structures present in the Error Information fields.

4.14.1. SW_ERROR, SW_SUBSCRIPTION_DENIED_ERROR and SW_SUBSCRIPTION_ID_REUSE_ERROR Information

The SW_ERROR error code indicates that the sender (the SW-PC) has encountered an error related to the processing of a SW Request attribute but which is not covered by more specific SW error codes. The SW_SUBSCRIPTION_DENIED_ERROR is used when the SW-PV requests to establish a subscription or clear all subscriptions from the given SW-PV, but the SW-PC is unable or unwilling to comply with this request. The SW_SUBSCRIPTION_ID_REUSE_ERROR is used when the SW-PC receives a SW Request whose Request ID duplicates a Subscription ID of an active subscription with the request's sender. All of these error codes use the following error information structure.




Figure 12: SW Error, Subscription Error, and Subscription ID Reuse Information

Field	Description
Copy of Request ID / Subscription ID	In the case that this error condition is generated in direct response to a SW Request attribute, this field MUST contain an exact copy of the Request ID field in the SW Request attribute that caused this error. In the case that the attribute in question is generated in fulfillment of an active subscription, this field MUST contain the Subscription ID of the subscription for which the attribute was generated. (This is only possible if the error code is SW_ERROR as the other errors are not generated by subscription fulfillment.) Note that, in this case, the indicated error appears as a sub-error for a SW_SUBSCRIPTION_FULFILLMENT_ERROR, as described in Section 4.14.3.
Description	A UTF-8 string describing the condition that caused this error. This field MAY be 0-length. However, senders SHOULD include some description in all PA-TNC Error attributes of these types. This field MUST NOT be NULL terminated.

Table 11: SW Error, Subscription Error, and Subscription ID Reuse Information Fields

This error information structure is used with SW_ERROR, SW_SUBSCRIPTION_DENIED_ERROR, and SW_SUBSCRIPTION_ID_REUSE_ERROR status codes to identify the SW Request attribute that precipitated the error condition and to describe the error. The Description field contains text describing the error. The SW-PC MAY encode machine-interpretable information in this field, but SHOULD also include a human-readable description of the error, since the receiving SW-PV might not recognize the SW-PC's encoded information.

4.14.2. SW_RESPONSE_TOO_LARGE_ERROR Information

The SW_RESPONSE_TOO_LARGE_ERROR error code indicates that a response generated by a SW-PC in response to a SW-PV's SW Request attribute was too large to send.




Figure 13: SW Response Too Large Error Information

Field	Description
Copy of Request ID / Subscription ID	In the case that the attribute in question is generated in direct response to a SW Request, this field MUST contain an exact copy of the Request ID field in the SW Request attribute that caused this error. In the case that the attribute in question is generated in fulfillment of an active subscription, this field MUST contain the Subscription ID of the subscription for which the attribute was generated. Note that, in the latter case, the <code>SW_RESPONSE_TOO_LARGE_ERROR</code> appears as a sub-error for a <code>SW_SUBSCRIPTION_FULFILLMENT_ERROR</code> , as described in Section 4.14.3.
Maximum Allowed Size	This field MUST contain an unsigned integer indicating the largest permissible size, in bytes, of SW Attribute that the SW-PC is currently willing to send in response to a SW Request attribute.
Description	A UTF-8 string describing the condition that caused this error. This field MAY be 0-length. However, senders SHOULD include some description in all PA-TNC Error attributes of these types. This field MUST NOT be NULL terminated.

Table 12: SW Response Too Large Error Information Fields

This error structure is used with the `SW_RESPONSE_TOO_LARGE_ERROR` status code to identify the SW Request attribute that precipitated the error condition and to describe the error. The Maximum Allowed Size field indicates the largest attribute the SW-PC is willing to

send in response to a SW Request under the current circumstances. Note that under other circumstances, the SW-PC might be willing to return larger or smaller responses than indicated (such as if the endpoint connects to the NEA Server using a different network protocol). The other fields in this error information structure have the same meanings as corresponding fields in the SW_ERROR and SW_SUBSCRIPTION_DENIED_ERROR information structure.

4.14.3. SW_SUBSCRIPTION_FULFILLMENT_ERROR Information

The SW_SUBSCRIPTION_FULFILLMENT_ERROR error code indicates that the SW-PC encountered an error while fulfilling a subscription. The bytes after the first 4 octets duplicate a PA-TNC Error attribute (as described in Section 4.2.8 of PA-TNC) that is used to identify the nature of the encountered error.




Figure 14: SW Subscription Fulfillment Error Information

Field	Description
Subscription ID	This field MUST contain the Subscription ID of the subscription whose fulfillment caused this error.
Reserved	This field MUST contain the value of the Reserved field of a PA-TNC Error attribute that describes the error condition encountered during subscription processing.
Sub Error Code Vendor ID	This field MUST contain the value of the Error Code Vendor ID field of a PA-TNC Error attribute that describes the error condition encountered during subscription processing.
Sub Error Code	This field MUST contain the value of the Error Code field of a PA-TNC Error attribute that describes the error condition encountered during subscription processing.
Sub Error Information	This field MUST contain the value of the Error Information field of a PA-TNC Error attribute that describes the error condition encountered during subscription processing.

Table 13: SW Subscription Fulfillment Error Information Fields

This error structure is used with the `SW_SUBSCRIPTION_FULFILLMENT_ERROR` status code. The first 4 octets of this error structure contain the Subscription ID of the subscription that was being fulfilled when the error occurred. The remaining fields of this error structure duplicate the fields of a PA-TNC Error attribute, referred to as the "sub-error". The error code of the sub-error corresponds to the type of error that the SW-PC encountered while fulfilling the given subscription. The sub-error MUST NOT have an error code of `SW_SUBSCRIPTION_FULFILLMENT_ERROR`.

The SW-PC sending a PA-TNC Error attribute with this error code, and the SW-PV receiving it, MUST treat the subscription identified by the Subscription ID field as cancelled. All other subscriptions are unaffected.

5. Supported Data Models

Software Inventory Message and Attributes for PA-TNC supports an extensible list of data models for representing and transmitting software inventory information. This list of data models appears in the Software Data Model IANA table (see Section 9.4). This document provides guidance for an initial set of data models. Other documents might provide guidance on the use of new data models by Software Inventory Message and Attributes for PA-TNC, and will be referenced by extensions to the Software Data Model IANA table.

5.1. ISO 2015 SWID Tags using XML

The International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC) published the specification governing SWID tag construction and use in 2009 with a revised version published in 2015. [SWID] Since that time, a growing number of vendors have integrated SWID tags into their software products. Doing so significantly simplifies the task of identifying these pieces of software: instead of relying on discovery processes that look for clues as to software presence, such as the presence of particular files or registry keys, a readily available list of SWID tags provides simple and immediate evidence as to the presence of the given piece of software.

SWID Message and Attributes for PA-TNC has no reliance on the presence or management of SWID tags on an endpoint as described in the ISO specification. However, the data model for describing software that is presented in the ISO specification provides a robust and comprehensive way of describing software and is adopted here as a means of representing and transmitting software information. It should be emphasized, the use of the ISO SWID tag data model makes no assumption as to whether the source of the recorded information was, in fact, an ISO SWID tag harvested from the endpoint or whether the information was created using some other source and normalized to the SWID format.

5.1.1. Guidance on Normalizing Source Data to ISO 2015 SWID Tags using XML

TBD

Don't violate the specification

Use your own Tag Creator RegID or the Unknown Tag Creator RegID. Do not use some other party's RegID, especially not the RegID of the software author if you are not the author.

5.1.2. Guidance on Creation of Software Identifiers from ISO 2015 SWID Tags

TBD

Use combination of Tag Creator RegID and Unique ID fields. Specifically, format should be `NUMBER::TAG_CREATOR_REGID UNIQUE_ID`, where `NUMBER` is the length of `TAG_CREATOR_REGID` in bytes. The rest of the Software Identifier MUST be the concatenation of the Tag Creator RegID and the Unique ID from the tag, without any connecting character or whitespace.

5.2. ISO 2009 SWID Tags using XML

As noted above, ISO's SWID tag specification provides a useful data model for representation of software information. As of the writing of this specification, while the 2015 specification is considered more comprehensive and addresses some issues with the 2009 specification, 2009-format SWID tags remain far more common in deployments. For this reason, ISO 2009 SWID tags are included in the Software Data Model IANA table.

5.2.1. Guidance on Normalizing Source Data to ISO 2015 SWID Tags using XML

TBD

Don't violate the specification

Use your own Tag Creator RegID or the Unknown Tag Creator RegID. Do not use some other party's RegID, especially not the RegID of the software author if you are not the author.

5.2.2. Guidance on Creation of Software Identifiers from ISO 2015 SWID Tags

TBD

Use combination of Tag Creator RegID and Unique ID fields. Specifically, format should be `NUMBER::TAG_CREATOR_REGID UNIQUE_ID`, where `NUMBER` is the length of `TAG_CREATOR_REGID` in bytes. The rest of the Software Identifier MUST be the concatenation of the Tag Creator RegID and the Unique ID from the tag, without any connecting character or whitespace.

6. Security Considerations

This section discusses some of the security threats facing Posture Collectors and Posture Validators that implement the Software Inventory Message and Attributes for PA-TNC protocol. This section primarily notes potential issues for implementers to consider, although it does contain a handful of normative requirements to address certain security issues. Implementers need to make the final decision as to how their implementations address the given issues. The issues identified below focus on capabilities specific to this document. Implementers are advised to consult other relevant NEA specifications for security issues that are applicable to such components in general.

Reading the Security Considerations section of any well-written specification can be discouraging, as a long list of possible threats is catalogued. Keep in mind that no security measure is absolute, but each one can be beneficial. By understanding the weaknesses of each security measure, we can put in place countermeasures to protect against exploitation of these weaknesses.

6.1. Evidentiary Value of Software Inventory Evidence Records

It must be remembered that the accuracy of an endpoints Software Inventory Evidence Collection as an indicator of the endpoints software load and changes thereon is only as accurate as the tools that populate and manage the software inventory evidence records in this collection. Some tools might not be designed to update records in the Software Inventory Evidence Collection in real time resulting in a collection that is out-of-step with actual system state. Moreover, tools might inaccurately characterize software, or fail to properly record its removal. Finally, it is likely that there will be software on the endpoint that is not tracked by any source and thus is not reflected in the Software Inventory Evidence Collection. Users of collected software inventory evidence records need to understand that the information provided by the Software Inventory Message and Attributes for PA-TNC capability cannot be treated as completely accurate. Nonetheless, having endpoints report this information can still provide useful insights into the state of the endpoint's software load, and can alert administrators and policy tools of situations that require remediation.

6.2. Sensitivity of Collected Records

Software records on an endpoint are generally not considered to be sensitive, although there can be exceptions to this generalization as noted in the section on Privacy Considerations. In general, an endpoint's Software Inventory Evidence Collection can be browsed and

individual records read by any party with access to the endpoint. This is generally not considered to be problematic, as those with access to the endpoint can usually learn of everything disclosed by that endpoint's records simply by inspecting other parts of the endpoint.

The situation changes when an endpoint's inventory records are collected and stored off of the endpoint itself, such as on a NEA Server or CMDB. Inventory records represent a wealth of information about the endpoint in question and, for an adversary who does not already have access to the endpoint, a collection of the endpoint's inventory records might provide many details that are useful for mounting an attack. A list of the inventory records associated with an endpoint reveals a list of software installed on the endpoint. This list can be very detailed, noting specific versions and even patch levels, which an adversary can use to identify vulnerable software and design efficacious attacks.

In addition, other information might also be gleaned from a collection of software inventory records:

- o An inventory record might include information about where the product was installed on a given endpoint. This can reveal details about the file organization of that endpoint that an attacker can utilize.
- o An inventory record might include information about how the software was provided to the endpoint, who in an organization signs off on the package release, and who packaged the product for installation. This information might be used as a starting point for the development of supply chain attacks.
- o Events affecting inventory records are reported with timestamps indicating when each given event occurred. This can give the attacker an indication of how quickly an organization distributes patches and updates, helping the attacker determine how long an attack window might remain open.

Any consolidated software inventory is a potential risk, because such an inventory can provide an adversary an insight into the enterprise's configuration and management process. It is recommended that a centralized software inventory record collection be protected against unauthorized access. Mechanisms to accomplish this can include encrypting the data at rest, ensuring that access to the data is limited only to necessary individuals and processes, and other basic security precautions.

6.3. Integrity of Endpoint Records

SW-PCs maintain records of detected changes to the endpoint's Software Inventory Evidence Collection. These records are used to respond to a SW-PV's request for change events. The SW-PV might use a list of reported events to update its understanding of the endpoint's Software Inventory Evidence Collection without needing to receive a full inventory report from the SW-PC. For this reason, preserving the integrity of the SW-PC's record of events is extremely important. If an attacker modifies the SW-PC's record of changes to the endpoint's Software Inventory Evidence Collection, this might cause the SW-PV's understanding of the endpoint's Software Inventory Evidence Collection to differ from its actual state. Results might include leading the SW-PV to believe that absent software was present, that present software was absent, that patches have been installed even if this is not the case, or to be unaware of other changes to Software Inventory Evidence Records. As such, the SW-PC MUST take steps to protect the integrity of its event records.

In addition, records of established SW-PV subscriptions also require protection against manipulation or corruption. If an attacker is able to modify or delete records of an established subscription by a SW-PV, the SW-PC might fail to correctly fulfill this subscription. The SW-PV would not be aware that its subscription was not being correctly fulfilled unless it received additional information that indicated a discrepancy. For example, the SW-PV might collect a full inventory and realize from this that certain events had not been correctly reported in accordance with an established subscription. For this reason, the SW-PC MUST protect the integrity of subscription records.

6.4. SW-PC Access Permissions

A SW-PC requires sufficient permissions to collect Software Inventory Evidence Records from all of its supported sources, as well as sufficient permissions to interact with the endpoint's Posture Broker Client. With regard to the former, this might require permissions to read the contents of directories throughout the file system. Depending on the operating environment and other activities undertaken by a SW-PC (or software that incorporates a SW-PC as one of its capabilities) additional permissions might be required by the SW-PC software. The SW-PC SHOULD NOT be granted permissions beyond what it needs in order to fulfill its duties.

6.5. Sanitization of Record Fields

Not all sources of software inventory evidence are necessarily tightly controlled. For example, consider a source that gathers .swid files from the endpoint's file system. Any party could create a new .swid file that could be collected and turned into a Software Inventory Evidence Record. As a result, it is important that the contents of source information not be automatically trusted. In particular, tools that read source information and the Software Inventory Evidence Records derived therefrom, including SW-PCs, need to be careful to sanitize input to prevent buffer overflow attacks, encoding attacks, and other weaknesses that might be exploited by an adversary who can control the contents of a record.

6.6. PA-TNC Security Threats

In addition to the aforementioned considerations the Software Inventory Message and Attributes for PA-TNC protocol is subject to the same security threats as other PA-TNC transactions, as noted in Section 5.2 of PA-TNC [RFC5792]. These include, but are not limited to, attribute theft, message fabrication, attribute modification, attribute replay, attribute insertion, and denial of service. Implementers are advised to consult the PA-TNC specification to better understand these security issues.

7. Privacy Considerations

As noted in Section 6.2, if an adversary can gain an understanding of the software installed on an endpoint, they can utilize this to launch attacks and maintain footholds on this endpoint. For this reason, the NEA Server needs to ensure adequate safeguards are in place to prevent exposure of collected inventory records. For similar reasons, it is advisable that an endpoint only send records to a NEA Server that is authorized to receive this information and that can be trusted to safeguard this information after collection.

8. Relationship to Other Specifications

This specification makes frequent reference to and use of other specifications. This section describes these relationships.

This specification is expected to participate in a standard NEA architecture. As such, it is expected to be used in conjunction with the other protocols used in a NEA exchange. In particular, SW Attributes are conveyed over PB-TNC [RFC5793], which is in turn conveyed over some variant of PT (either PT-TLS [RFC6876] or PT-EAP [RFC7171]). These protocols have an especially important role, as they are responsible for ensuring that attributes defined under this

specification are delivered reliably, securely, and to the appropriate party.

It is important to note that the Product Information, Numeric Version, and String Version attributes defined in the PA-TNC specification [RFC5792] are also meant to convey information about installed applications and the versions thereof. As such, there is some conceptual overlap between those attributes and the intent of this specification. However, PA-TNC was designed to respond to very specific queries about specific classes of products, while the Software Inventory Message and Attributes for PA-TNC specification is able to convey a broader query, resulting in a more comprehensive set of evidence regarding an endpoint's installed software. As such, this specification provides important capabilities not present in the PA-TNC specification.

9. IANA Considerations

This section extends multiple existing IANA registries. Specifically, it extends the PA-TNC Attribute Types and PA-TNC Error Codes defined in the PA-TNC specification [RFC5792] and the PA-Subtypes registry defined in the PB-TNC specification [RFC5793] and extended in PA-TNC. This specification only adds values to these registries and does not alter how these registries work or are maintained. Consult the appropriate specifications for details on the operations and maintenance of these registries.

9.1. Registry for PA-TNC Attribute Types

Section 4.4 of this specification defines several new PA-TNC attributes. The following values are added to the registry for PA-TNC Attribute Types defined in the PA-TNC specification. Note that Table 3 in Section 4.4 lists these attributes but uses a hexadecimal value to identify their associated integer. The integer values given in that table are identical to those provided here. Note also that Table 3 includes an entry for PA-TNC Error attributes, but the IANA information associated with that attribute is already defined in the PA-TNC specification and is not reproduced here.

PEN	Integer	Name	Defining Specification
0	17	SW Request	Software Inventory Message and Attributes for PA-TNC
0	18	Software Identifier Inventory	Software Inventory Message and Attributes for PA-TNC
0	19	Software Identifier Events	Software Inventory Message and Attributes for PA-TNC
0	20	Software Inventory	Software Inventory Message and Attributes for PA-TNC
0	21	Software Events	Software Inventory Message and Attributes for PA-TNC
0	22	Subscription Status Request	Software Inventory Message and Attributes for PA-TNC
0	23	Subscription Status Response	Software Inventory Message and Attributes for PA-TNC
0	24	Subscription Status Response	Software Inventory Message and Attributes for PA-TNC
0	25 - 31	Reserved for future use	Software Inventory Message and Attributes for PA-TNC

9.2. Registry for PA-TNC Error Codes

Section 4.14 of this specification defines several new PA-TNC Error Codes. The following values are added to the registry for PA-TNC Error Codes defined in the PA-TNC specification. Note that Table 10 in Section 4.14 lists these codes but uses a hexadecimal value to identify their associated integer. The integer values given in that table are identical to those provided here.

PEN	Integer	Name	Defining Specification
0	32	SW_ERROR	Software Inventory Message and Attributes for PA-TNC
0	33	SW_SUBSCRIPTION_DENIED_ERROR	Software Inventory Message and Attributes for PA-TNC
0	34	SW_RESPONSE_TOO_LARGE_ERROR	Software Inventory Message and Attributes for PA-TNC
0	35	SW_SUBSCRIPTION_FULFILLMENT_ERROR	Software Inventory Message and Attributes for PA-TNC
0	36	SW_SUBSCRIPTION_ID_REUSE_ERROR	Software Inventory Message and Attributes for PA-TNC
0	37-47	Reserved for future use	Software Inventory Message and Attributes for PA-TNC

9.3. Registry for PA Subtypes

Section 4.1 of this specification defines one new PA Subtype. The following value is added to the registry for PA Subtypes defined in the PB-TNC specification.

PEN	Integer	Name	Defining Specification
0	9	SW Attributes	Software Inventory Message and Attributes for PA-TNC

9.4. Registry for Software Data Models

TBD

10. References

10.1. Normative References

- [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>>.
- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, DOI 10.17487/RFC3339, July 2002, <<http://www.rfc-editor.org/info/rfc3339>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC5198] Klensin, J. and M. Padlipsky, "Unicode Format for Network Interchange", RFC 5198, DOI 10.17487/RFC5198, March 2008, <<http://www.rfc-editor.org/info/rfc5198>>.
- [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>>.
- [RFC5792] Sangster, P. and K. Narayan, "PA-TNC: A Posture Attribute (PA) Protocol Compatible with Trusted Network Connect (TNC)", RFC 5792, DOI 10.17487/RFC5792, March 2010, <<http://www.rfc-editor.org/info/rfc5792>>.
- [SWID] The International Organization for Standardization/International Electrotechnical Commission, "Information Technology - Software Asset Management - Part 2: Software Identification Tag, ISO/IEC 19770-2", November 2009.

10.2. Informative References

[NIST-SWID]

The National Institute of Standards and Technology and The MITRE Corporation, "Guidelines for the Creation of Interoperable Software Identification (SWID) Tags", August 2013.

[RFC5793] Sahita, R., Hanna, S., Hurst, R., and K. Narayan, "PB-TNC: A Posture Broker (PB) Protocol Compatible with Trusted Network Connect (TNC)", RFC 5793, DOI 10.17487/RFC5793, March 2010, <<http://www.rfc-editor.org/info/rfc5793>>.

[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, <<http://www.rfc-editor.org/info/rfc6876>>.

[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, <<http://www.rfc-editor.org/info/rfc7171>>.

Authors' Addresses

Chris Coffin
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
USA

Email: ccoffin@mitre.org

Daniel Haynes
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
USA

Email: dhaynes@mitre.org

Charles Schmidt
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
USA

Email: cmschmidt@mitre.org

Jessica Fitzgerald-McKay
Department of Defense
9800 Savage Road
Ft. Meade, Maryland
USA

Email: jmfitz2@nsa.gov

SACM
Internet-Draft
Intended status: Informational
Expires: July 25, 2016

C. Coffin
B. Cheikes
C. Schmidt
D. Haynes
The MITRE Corporation
J. Fitzgerald-McKay
Department of Defense
D. Waltermire
National Institute of Standards and Technology
January 22, 2016

SACM Vulnerability Assessment Scenario
draft-coffin-sacm-vuln-scenario-01

Abstract

This document provides a core narrative that walks through an automated enterprise vulnerability assessment scenario. It is aligned with the SACM use cases and begins with an enterprise ingesting vulnerability description data, followed by identifying endpoints on the network and collecting and storing information about them to enable posture assessment, and finally ends with assessing these endpoints against the vulnerability description data to determine which ones are affected. Processes that specifically overlap between this scenario and SACM use cases will be noted where applicable. Specifically, the relationship between this document and the SACM use case building block capabilities and the usage scenarios will be covered.

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 July 25, 2016.

Copyright Notice

Copyright (c) 2016 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 (<http://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. Scope	3
2. Assumptions	4
3. Endpoint Identification and Initial (Pre-Assessment) Data Collection	5
3.1. Identification	6
3.1.1. SACM Use Case Alignment	6
3.2. Processing Artifacts	6
3.3. Endpoint Data Collection	7
3.3.1. SACM Use Case Alignment	8
3.4. Implementation Examples	9
4. Vulnerability Description Data	9
4.1. SACM Use Case Alignment	10
4.2. Implementation Examples	10
5. Endpoint Applicability and Assessment	10
5.1. Applicability	11
5.1.1. SACM Use Case Alignment	11
5.2. Secondary Assessment	11
5.2.1. SACM Use Case Alignment	12
5.3. Implementation Examples	13
6. Assessment Results	13
6.1. SACM Use Case Alignment	14
6.2. Implementation Examples	15
7. IANA Considerations	15
8. Security Considerations	15
9. Informative References	15
Appendix A. Change Log	16
A.1. Changes in Revision 01	16
Appendix B. Continuous Vulnerability Assessment	17
Appendix C. Priority	18
Appendix D. Data Attribute Table and Definitions	19
D.1. Table	19

D.2. Definitions	22
Appendix E. Alignment with Other Existing Works	24
E.1. Critical Security Controls	24
E.1.1. Continuous Vulnerability Assessment	24
E.1.2. Hardware and Software Inventories	26
Appendix F. SACM Usage Scenarios	26
Appendix G. SACM Requirements and Charter - Future Work	28
Authors' Addresses	28

1. Scope

The purpose of this document is to describe a detailed scenario for vulnerability assessment, and identify aspects of this scenario that could be used in the development of an information model. This includes classes of data, major roles, and a high-level description of role interactions. Additionally, this scenario intends to inform engineering work on protocol and data model development. The focus of the document is entirely intra-organizational and covers enterprise handling of vulnerability description data. The document does not attempt to cover the security disclosure itself and any prior activities of the security researcher or discloser, nor does it attempt to cover the specific activities of the vendor whose software is the focus of the vulnerability description data (i.e., the vulnerable software).

For the purposes of this document, the term "vulnerability description data" is intended to mean: "Data intended to alert enterprise IT resources to the existence of a flaw or flaws in software, hardware, and/or firmware, which could potentially have an impact on enterprise functionality and/or security." For the purpose of this scenario, such data also includes information that can be used to determine (to some level of accuracy, although possibly not conclusively) whether or not the flaw is present within an enterprise, when compared to information about the state of the enterprise's endpoints. For those who are familiar with current security practices and terminology, the use of vulnerability description data is also synonymous with security bulletin or advisory.

This document makes no attempt to provide a definition of a normalized data format (e.g. industry standard) for vulnerability description data although there is nothing precluding the development of such a normalized data format. Also, it does not attempt to define procedures by which a vulnerability discoverer coordinates the release of vulnerability description data to other parties.

2. Assumptions

A number of assumptions must be stated in order to further clarify the position and scope of this document.

- o The document begins with the assumption that the enterprise has received vulnerability description data, and that the data has already been processed into a format that the enterprise's security software tools can understand and use. In particular, this document:
 - * Does not discuss how the enterprise identifies potentially relevant vulnerability description data.
 - * Does not discuss how the enterprise collects the vulnerability description data.
 - * Does not discuss how the enterprise assesses the authenticity of the vulnerability description data.
 - * Does not discuss parsing of the vulnerability description data into a usable format.
- o The document assumes that the enterprise has a means of identifying enterprise endpoints. This could mean identifying endpoints as they join the network, actively scanning for connected endpoints, passive scanning of network traffic to identify connected endpoints, or some other method of accounting for the presence of all endpoints in the enterprise. The document also does not distinguish between physical endpoints and virtualized endpoints.
- o The document assumes that the enterprise has a means of extracting relevant information about enterprise endpoints. Moreover, this extracted information is expressed in a format that is compatible with the information extracted from the vulnerability description data. The document:
 - * Does not specify how relevant information is identified.
 - * Does not specify the mechanics of how relevant information is extracted from the data sources (such as the endpoint itself).
 - * Does not specify how extracted endpoint information and vulnerability description data is normalized to be compatible.

Note that having a means of extracting relevant information about enterprise endpoints is within the scope of the SACM Endpoint

Security Posture Assessment process. In the case of this document, this sub-process is assumed to be existent.

- o The document assumes that all information described in the steps below is available in the vulnerability description data and serves as the basis of this assessment. Likewise, the document assumes that the enterprise can provide all relevant information about any endpoint needed to perform the described analysis. The authors recognize that this will not always be the case, but these assumptions are taken in order to show the breadth of data utilization in this scenario. Less complete information may require variations to the described steps.
- o The document assumes that the enterprise has a policy by which assessment of endpoints based on vulnerability description data is prioritized. The document:
 - * Does not specify how prioritization occurs.
 - * Does not specify how prioritization impacts assessment behaviors.
- o The document assumes that the enterprise has a mechanism for long-term storage of vulnerability description data and endpoint assessment results (e.g., a data repository).
- o This document assumes that the enterprise has a procedure for reassessment of endpoints at some point after initial assessment. The document:
 - * Does not specify how a reassessment would impact individual assessment behaviors. (i.e., it is agnostic as to whether the assessment procedure is the same regardless of whether this is the first or a subsequent assessment for some set of vulnerability description data.)
 - * Does not provide recommendations or specifics on reassessment intervals.

3. Endpoint Identification and Initial (Pre-Assessment) Data Collection

The first step in this scenario involves identifying endpoints and collecting the basic or minimum set of system information attributes from them such as operating system type and version. Further examples of system information and attributes can be found below in the section titled Endpoint Data Collection. This identification occurs prior to the receipt of any specific vulnerability description data and is part of the regular, ongoing monitoring of endpoints

within an enterprise. This process is not meant to report on, or gather data for any specific vulnerabilities. The information gathered during this step could be applied in many enterprise automation efforts. Specifically, in addition to vulnerability management, it could be used by configuration and license management tasks. All of the information collected during this step is stored in a central location such as a Repository.

This activity involves the following sub-steps:

3.1. Identification

Prior to any other steps, the identification of endpoints must occur. This involves locating (at least virtually) and distinguishing between endpoints on the network in a way that allows each endpoint to be recognized in future interactions and selected for specific treatment. This not only allows later steps to determine the scope of what endpoints need to be assessed, but also allows for the unique identification of each endpoint. Unique and persistent endpoint IDs are used to allow for endpoints to be tracked over time and between sensors as well as allow for proper counts of assets during inventories and other similar collections. Endpoint identity can be established by collecting certain attributes that allow for unique and persistent tracking of endpoints on the enterprise network. Examples include, but are not limited to, IP address, MAC address, FQDNs, pre-provisioned identifiers such as GUIDs or copies of serial numbers, certificates, hardware identity values, or similar attributes. It is important to note that the persistency of these attributes will likely vary depending on the enterprise. For example, a statically assigned IP address is much more persistent than an IP address assigned via DHCP.

3.1.1. SACM Use Case Alignment

This sub-step aligns with the Endpoint Discovery, Endpoint Characterization, and Endpoint Target Identification building block capabilities. The alignment is due to the fact that the purpose of this sub-step is to discover, identify, and characterize all endpoints on an enterprise network.

3.2. Processing Artifacts

Processing artifacts, such as the date and time the collection was performed, should be collected and stored. This timestamp is extremely important when performing later assessments, as it is needed for data freshness computations. The organization may develop rules for stale data and when a new data collection is required. This metadata is also helpful in correlating information across

multiple data collections. This includes correlating both pre-assessment data and secondary assessment data (sections 4.3 Endpoint Data Collection and 6.2 Secondary Assessment).

3.3. Endpoint Data Collection

The enterprise should perform ongoing collection of basic endpoint information such as operating system and version information, and an installed software inventory. This information is collected for general system monitoring as well as its potential use in activities such as vulnerability assessment.

Some examples of basic information to collect about endpoints in this pre-assessment process could include:

- o Endpoint type - traditional (e.g., workstation, server, etc.) network infrastructure (e.g., switches, routers, etc.), mobile (e.g., cell phones, tablets, laptops, etc.), and constrained (e.g., industrial control systems, Internet of Things, etc.)
- o Hardware version/firmware - e.g., BIOS version, firmware revision, etc.
- o Operating system - e.g., Windows, Linux, Mac OS, Android
- o Operating system attributes - e.g., version, patch level, service pack level, internationalized or localized version, etc.
- o Installed software inventory - Would include the software names and versions and possibly other high-level attributes. Could be used to quickly determine endpoint applicability when new vulnerability description data arrives.

Some additional and more advanced information to collect from endpoints in this pre-assessment process could include:

- o Open ports and enabled services - This would include applications listening for incoming connections on open ports as well as services that are starting, running, suspended, or enabled to run pending some event.
- o Operating system optional component inventory - some OS' have optional components that can be installed which may not show up as separate pieces of software (e.g., web and ftp servers, demo web pages, shared libraries, etc.). Note that this could also occur within third-party applications as well.

- o Endpoint location - physical location (e.g., department, room, Global Positioning System (GPS), etc.), logical location (e.g., what network infrastructure endpoints (e.g. switches, wireless access point, etc.) an endpoint is connected to, etc.
- o Purpose - describes how the endpoint is used within the enterprise (e.g., end-user system, database server, public web server, etc.)
- o Criticality - enterprise defined rating (possibly a score) that helps determine the criticality of the endpoint. If this endpoint is attacked or lost, what is the impact to the overall enterprise?

It is important to note that some of these attributes may exist natively on the endpoint whereas other attributes may be assigned by a human, computed, or derived from other data and may or may not be available for collection on the endpoint.

Furthermore, the possibility should be left open for enterprises to define their own custom queries and algorithms to gather and derive enterprise-specific attributes that are deemed of interest to regular enterprise operations.

In addition to collecting these attributes, metadata about the attributes should also be collected which could include:

Data origin - where the data originated from

Data source - what provided the data

Date and time of collection - when the data was collected

3.3.1. SACM Use Case Alignment

This sub-step aligns with the Data Publication building block capability because this section involves storage of endpoint attributes within an enterprise Repository. This sub-step also aligns with the Endpoint Characterization and Endpoint Target Identification building block capabilities because it further characterizes the endpoint through automated and possibly manual means. There is direct alignment with the Endpoint Component Inventory, Posture Attribute Identification, and Posture Attribute Value Collection building block capabilities since the purpose of this sub-step is to perform an initial inventory of the endpoint and collect basic attributes and their values. Last, there is alignment with the Collection Guidance Acquisition building block capabilities as the inventory and collection of endpoint attributes would be directed by some type of enterprise or third-party guidance.

3.4. Implementation Examples

Within the SACM Architecture, the Internal and External Collector components could be used to allow enterprises to collect posture attributes that demonstrate compliance with enterprise policy. Endpoints can be required to provide posture attributes, which may include identification attributes to enable persistent communications.

The SWID Message and Attributes for IF-M standard defines collection and validation of software identities using the ISO Software Identification Tag Standard. Using this standard, the identity of all installed software including the endpoint operating system, could be collected and used for later assessment.

The OVAL Definitions Model provides a data model that can be used to specify what posture attributes to collect as well as their expected values which can be used to drive an assessment.

The OVAL System Characteristics Model can be used to capture information about an endpoint. The model is specifically suited to expressing OS information, endpoint identification information (such as IP and MAC addresses), and other endpoint metadata.

4. Vulnerability Description Data

The next step in the Vulnerability Assessment scenario begins after vulnerability description data has been received and processed into a form that can be used in the assessment of the enterprise. As a part of the enterprise process for managing vulnerability description data, the enterprise should store all received and processed vulnerability description data in a Repository. The stored vulnerability description data can be used and compared with later vulnerability description data for the purpose of duplicate detection and in some cases, guidance on how to handle similar issues.

All vulnerability description data should be assigned an internal tracking ID by the enterprise as a first step as this helps compensate for the fact that incoming vulnerability description data might not have a global identifier when it is received, and might never be assigned one.

High-level vulnerability description data metadata to store would include:

- o Ingest date and time - the date and time that the vulnerability description data was received by the enterprise.

- o Date and time of vulnerability description data release (i.e., publication or disclosure date and time) - Some older vulnerability description data may be ingested long after publication. This can be useful when reviewing historical enterprise information to (potentially) identify the period when a particular endpoint was first assessed as vulnerable. Sometimes this information will help to differentiate between similar vulnerability description data.
- o Version - the version or iteration of the vulnerability description data according to the author, if applicable.
- o External Vulnerability Description Data ID(s) (if applicable) - any external or third-party IDs assigned to the vulnerability description data should be tracked. There could be multiple IDs in some cases (e.g., vendor bug id, global ID, discoverer's local ID, third-party vulnerability database ID, etc.).
- o Severity Score (if available) - these may be useful for later mitigation prioritization.

In addition to the described metadata, the raw or original vulnerability description data would be stored along with the specific information extracted from it that is to be used in the applicability and assessment process.

4.1. SACM Use Case Alignment

This step aligns with the Data Publication and Data Retrieval building block capabilities because this section details storage of vulnerability description data within an enterprise Repository and later retrieval of the same.

4.2. Implementation Examples

The Common Vulnerability Reporting Framework (CVRF) is an XML-based language that attempts to standardize the creation of vulnerability report documentation. Using CVRF, the enterprise could create automated tools based on the standardized schema which would obtain the needed and relevant information useful for later assessments and assessment results.

5. Endpoint Applicability and Assessment

When new vulnerability description data is received by the enterprise, applicable enterprise endpoints must be identified and assessed. Endpoints are first examined using the already obtained pre-assessment data. If this is not sufficient to determine endpoint

applicability, a secondary data collection for additional data and attributes may be performed to determine status with regard to the vulnerability description data.

5.1. Applicability

The applicability of an endpoint and its vulnerability status can, in many cases, be determined entirely by the existence of a particular version of installed software on the endpoint. This data may have been collected in the pre-assessment data collection. If the applicability and vulnerability status of an endpoint can be determined entirely by the pre-collected data attribute set, no further data collection is required.

Other cases may require specific data (i.e., file system attributes, specific configuration parameters, etc.) to be collected for the assessment of a particular vulnerability description data. In these cases, a secondary, targeted vulnerability assessment is required. Administrators may want to evaluate applicability to the vulnerability description data iteratively. Specifically, the process would compare against pre-collected data first (easy to do and the data is stored in a Repository), and then if needed, query endpoints that are not already excluded from applicability for additional required data. (I.e., A "fast-fail" model). To do this, the criteria for determining applicability must be separable, so that some conclusions can be drawn based on the possession of partial data.

5.1.1. SACM Use Case Alignment

This sub-step aligns with the Data Retrieval, Data Query, and Posture Attribute Value Query building block capabilities because, in this sub-step, the process is attempting to determine the vulnerability status of the endpoint using the data that has previously been collected.

5.2. Secondary Assessment

If the applicability and vulnerability status of an endpoint cannot be determined by the pre-assessment data collection, a secondary and targeted assessment of the endpoint will be required. A secondary assessment may also be required in the case that data on-hand (either from pre-assessment or from prior secondary assessments) is stale or out-of-date.

The following data types and attributes are examples of what might be required in the case of a secondary and targeted assessment:

- o Specific files and attributes - i.e., file name, versions, size, write date, modified date, checksum, etc. Some vulnerabilities may only be distinguishable through the presence or absence of specific files or their attributes.
- o Shared libraries - Some vulnerabilities will affect many products across multiple vendors. In these cases the vulnerability may apply to a shared library. Under these circumstances, product versions may be less helpful than looking for the presence of one or more specific files and their attributes.
- o Other software configuration information (if applicable) - e.g., Microsoft Windows registry queries, Apple configuration profiles, GConf, Proc filesystem, text configuration files and their parameters, and the installation paths. Sometimes vulnerabilities only affect certain software configurations and in some cases these are not the default configurations. Certain configuration attributes can be used to determine the current configuration state.

Note that the secondary assessment described here does not need to be a pull assessment that is initiated by the server. The secondary assessment could also be part of a push to the server when the endpoint detects a change to a vulnerability assessment baseline.

5.2.1. SACM Use Case Alignment

This sub-step aligns with the Data Publication building block capability because this section details storage of endpoint attributes within an enterprise Repository. The sub-step also aligns with the Collection Guidance Acquisition building block capability since the vulnerability description data (guidance) drives the collection of additional endpoint attributes.

This sub-step aligns with the Endpoint Characterization (both manual and automated) and Endpoint Target Identification building block capabilities because it could further characterize the endpoint through automated and possibly manual means. There is direct alignment with the Endpoint Component Inventory, Posture Attribute Identification, and Posture Attribute Value Collection building block capabilities since the purpose of this sub-step is to perform additional and more specific component inventories and collections of endpoint attributes and their values.

5.3. Implementation Examples

Within the SACM Architecture, the assessment task would be handled by the Evaluator component. If pre-assessment data is used, this would be stored on and obtained from a Data Store component.

Within the SACM Architecture, the Internal and External Collector components could be used to allow enterprises to collect posture attributes that demonstrate compliance with enterprise policy. Endpoints can be required to provide posture attributes, which may include identification attributes to enable persistent communications.

The SWID Message and Attributes for IF-M standard defines collection and validation of software identities using the ISO Software Identification Tag Standard. Using this standard, all installed software including the endpoint operating system could be collected and stored for later assessment.

The OVAL Definitions Model provides a data model that can be used to specify what posture attributes to collect as well as their expected values which can be used to drive an assessment.

The OVAL System Characteristics Model can be used to capture information about an endpoint. The model is specifically suited to expressing OS information, endpoint identification information (such as IP and MAC addresses), and other endpoint metadata.

The SACM Internal and External Attribute Collector components can be used to allow enterprises to collect posture attributes that demonstrate compliance with enterprise policy. Endpoints can be required to provide posture attributes, which may include identification attributes to enable persistent communications.

6. Assessment Results

Assessment results present the results of an assessment, along with sufficient context so a human or machine can make the appropriate response. This context might include a description of the issue provided by the vulnerability description data, the endpoint attributes that indicate applicability, or other information needed to respond to the results of the assessment. Data in this step is stored for auditing and forensic purposes.

The following details are important to track in assessment results. Note that information may be "included" by providing pointers to other records stored in a Repository (e.g., vulnerability description data, endpoint data, etc.).

- o Date and time of assessment - The date and time that the assessment was performed. To understand when the data was compared against the vulnerability description data and what conclusions were drawn.
- o Data collection/attribute age - The age of the data used in the assessment to make the endpoint status determination.
- o Endpoint ID - The endpoint itself must be identified for tracking results over time.
- o Vulnerability description data ID(s) - May include both the internally defined ID as well as one or more externally defined IDs if they exist. The internally assigned ID allows linkage to the correct vulnerability description data. If available, external IDs provide a "pivot point" to additional external information.
- o Vulnerable software product(s) - Identifies the software products on the endpoint that resulted in the endpoint being declared applicable. Since some vulnerability description data identify vulnerabilities in multiple products, this will help identify the specific product (or products) found to be vulnerable in the endpoint assessment.
- o Endpoint vulnerability status - The endpoint status based on the vulnerability description data. Does the vulnerability exist on the endpoint?
- o Vulnerability description - Not needed for automated assessment but probably should be included for human review. The reason for inclusion is to support the human user understanding of the vulnerability assessment results within the application front-end or interface.
- o Vulnerability remediation - Similar to the above, remediation or vendor patch information would be useful for a human response. In many cases, this information may be a part of the description information described above. Note that patch information may change over time due to supersession of the vendor patches.

6.1. SACM Use Case Alignment

This step aligns with the Data Publication and Data Retrieval building block capabilities because this section details storage of vulnerability assessment results within an enterprise Repository and later retrieval of the same.

6.2. Implementation Examples

The OVAL Results Model provides a data model to encode the results of the assessment, which could then be stored in a Repository and later accessed. The assessment results described in this scenario could be stored and later accessed using the OVAL Results Model. Note that the use of the OVAL Results Model for sharing results is not recommended per section 7.3 of the OVAL and the SACM Information Model [draft-hansbury-sacm-oval-info-model-mapping-01].

Within the SACM Architecture, the generation of the assessment results would occur in the Report Generator component. Those results might then be moved to a Data Store component for later sharing and retrieval as defined by SACM.

7. IANA Considerations

This memo includes no request to IANA.

8. Security Considerations

This document provides a core narrative that walks through an automated enterprise vulnerability assessment scenario and is aligned with SACM "Endpoint Security Posture Assessment: Enterprise Use Cases" [RFC7632]. As a result, the security considerations for [RFC7632] apply to this document. Furthermore, the vulnerability description data may provide attackers with useful information such as what software an enterprise is running on their endpoints. As a result, organizations should properly protect the vulnerability description data it ingests.***TODO IS THIS COVERED BY RFC7632???***

9. Informative References

[charter-ietf-sacm-01]

Security Automation and Continuous Monitoring, "Charter, Version 1.0", July 2013.

[critical-controls]

Council on CyberSecurity, "Critical Security Controls, Version 5.1".

[draft-hansbury-sacm-oval-info-model-mapping-01]

Security Automation and Continuous Monitoring, "OVAL and the SACM Information Model", November 2015.

[I-D.ietf-sacm-requirements]

Cam-Winget, N. and L. Lorenzin, "Secure Automation and Continuous Monitoring (SACM) Requirements", draft-ietf-sacm-requirements-11 (work in progress), November 2015.

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

Appendix A. Change Log

A.1. Changes in Revision 01

Clarification of the vulnerability description data IDs in sections 4 and 6.

Added "vulnerability remediation" to the Assessment Results and Data Attribute Table and Definitions sections.

Added Implementation Examples to Endpoint Identification and Initial (Pre-Assessment) Data Collection, Vulnerability Description Data, Endpoint Applicability and Assessment, and Assessment Results sections.

Added an example to vulnerability description data in the scope section.

Added a sentence to clarify vulnerability description data definition in the scope section.

Added data repository example for long-term storage scope item.

Added sentence to direct reader to examples of basic system information in endpoint identification section.

Split the examples of information to collect in the pre-assessment collection section into a basic and advanced list.

Added examples of data stored in the repository in the Assessment Results section.

Added sentence for human-assigned attributes in the Future Work section.

Replaced "vulnerability report" to "vulnerability description data" because the term report was causing confusion. Similarly, replaced "assessment report" with "assessment results".

Replaced "Configuration Management Database (CMDB)" with "Repository" which is SACM's term for a data store.

Replaced endpoint "Role" with "Purpose" because "Role" is already defined in SACM. Also, removed "Function" because it too is already defined in SACM.

Clarified that the document does not try to define a normalized data format for vulnerability description data although it does not preclude the creation of such a format.

Included additional examples of software configuration information.

Clarified the section around endpoint identification to make it clear designation attributes used to correlate and identify endpoints are both persistent and unique. Furthermore, text was added to explain how the persistency of attributes may vary. This was based on knowledge gained from the Endpoint ID Design Team.

Updated the Security Considerations section to mention those described in [RFC7632].

Removed text around Bring Your Own Device (BYOD). While important, BYOD just adds complexity to this initial draft. BYOD should be addressed in a later revision.

Merged the list of "basic endpoint information" and the list of "human-assigned endpoint attributes" as both represent data we want to collect about an endpoint. Whether or not that data is natively available on the endpoint for collection or assigned by a human, computed, or derived from other data which may or may not be available on the endpoint for collection seems arbitrary. With this scenario, we primarily care about expressing information needs rather than how the information is collected or from where.

Appendix B. Continuous Vulnerability Assessment

It is not sufficient to perform a single assessment when vulnerability description data is published without any further checking. Doing so does not address the possibility that the reported vulnerability might be introduced to the enterprise environment after the initial assessment completes. For example, new endpoints can be introduced to the environment which have old software or are not up-to-date with patches. Another example is where unauthorized or obsolete software is installed on an existing endpoint by enterprise users after vulnerability description data and initial assessment has taken place. Moreover, enterprises might not wish to, or be able to, assess all vulnerability description data

immediately when they come in. Conflicts with other critical activities or limited resources might mean that some alerts, especially those that the enterprise deems as "low priority", are not used to guide enterprise assessments until sometime after the initial receipt.

The scenario above describes a single assessment of endpoints. However, it does not make any assumptions as to when this assessment occurs relative to the original receipt of the vulnerability description data that led to this assessment. The assessment could immediately follow ingest of the vulnerability description data, could be delayed, or the assessment might represent a reassessment of some vulnerability description data against which endpoints had previously been assessed. Moreover, the scenario incorporates long-term storage of collected data, vulnerability description data, and assessment results in order to facilitate meaningful and ongoing reassessment.

Appendix C. Priority

Priorities associated with the vulnerability description data, assessment results, and any remedy is important, but is treated as a separate challenge and, as such, has not been integrated into the description of this scenario. Nevertheless, it is important to point out and describe the use of priorities in the overall vulnerability description data scenario as they separable issues with their own sets of requirements.

Priority in regard to vulnerability description data, can be viewed in a couple of different ways within an enterprise. The assessment prioritization involves prioritization of the vulnerability description data assessment process. This determines what vulnerability description data is assessed, and in what order it is assessed in. For instance, a vulnerability affecting an operating system or application used throughout the enterprise would likely be prioritized higher than a vulnerability in an application which is used only on a few, low-criticality endpoints.

The prioritization of remedies relates to the enterprise remediation and mitigation process based on the discovered vulnerabilities. Once an assessment has been performed and applicable endpoints identified, enterprise vulnerability managers must determine where to focus their efforts to apply appropriate remedies. For example, a vulnerability that is easily exploitable and which can allow arbitrary code execution might be remedied before a vulnerability that is more difficult to exploit or which just degrades performance.

Some vulnerability description data include severities and/or other information that places the vulnerability in context. This information can be used in both of the priority types discussed above. In other cases, enterprise administrators may need to prioritize based only on what they know about their enterprise and the description provided in the vulnerability description data.

Examples of data attributes specific to priority of assessments and/or remedies include (but not limited to) the following:

- o Enterprise - defined purpose of the device, criticality of the device, exposure of the device, etc.
- o Severity attributes - A rating or score that attempts to provide the level of severity or criticality associated with a given vulnerability.
- o Cyber threat intelligence - information such as tactics, techniques, and procedures of threat actors, indicators of compromise, incidents, courses of action, etc. that help the enterprise understand relevant threats and how to detect, mitigate, or respond to them.

Appendix D. Data Attribute Table and Definitions

D.1. Table

The following table maps all major data attributes against each major process where they are used.

	vulnerability description data	Endpoint Identification and Initial (Pre-Assessment) Data Collection	Endpoint Applicability and Assessment	Assessment Results
Endpoint				
Collection date/time		X	X	
Endpoint type		X	X	
Hardware version	X	X	X	

sion/firmware				
Operating system	X	X	X	
Operating system attributes (e.g., version, service pack level, edition, etc.)	X	X	X	
Installed software name	X	X	X	X
Installed software attributes (e.g., version, patch level, install path, etc.)	X	X	X	X
Open ports/services	X	X	X	
Operating system optional component inventory	X	X	X	
Location		X		X
Purpose		X		X
Criticality		X		X
File system attributes (e.g., versions,	X		X	

size, write date, modified date, checksum, etc.)				
Shared libraries	X		X	
Other software configuration information	X		X	
External vulnerability description data				
Ingest Date	X		X	
Date of Release	X		X	
Version	X		X	
External vuln ID	X		X	X
Severity Score				X
Assessment Results				
Date of assessment			X	X
Date of data collection		X	X	X
Endpoint identification and/or locally assigned ID		X	X	X

Vulnerable software product(s)	X	X	X	X
Endpoint vulnerability status			X	X
Vulnerability description	X			X
Vulnerability remediation	X			X

Table 1: Vulnerability Assessment Attributes

D.2. Definitions

Endpoint

- o Collection date/time - the date and time of data collection
- o Endpoint type - the device type of the endpoint (e.g., standard computer, printer, router, mobile device, tablet, etc.)
- o Hardware version/firmware - the hardware or firmware version if applicable (e.g., BIOS version, firmware revision, etc.)
- o Operating system - Operating system name
- o Operating system attributes - Operating system high-level attributes (e.g., version, service pack level, edition, etc.). Would not include configuration details.
- o Installed software name - List of all installed software packages (i.e., software inventory). May or may not include software installed by the operating system.
- o Installed software attributes - Software high-level attributes (e.g., version, patch level, install path, etc.). Would not include configuration details.
- o Open ports/enabled services - Listening network ports (e.g., TCP, UDP, etc.) as well as services that are starting, running, suspended, or enabled to run pending some event.

- o Operating system optional component inventory - Operating system specific components and software (when NOT already included in the general software inventory)
- o Location - The physical location of an enterprise endpoint (e.g., department, room, etc.)
- o Purpose - describes how the endpoint is used within the enterprise (e.g., end user system, database server, public web server, etc.)
- o Criticality - An enterprise-defined rating (possibly a score) that helps determine the criticality of the endpoint. If this endpoint is attacked or lost, what is the impact to the overall enterprise?
- o File system attributes - Attributes that describe the file or directory (e.g., versions, size, write date, modified date, checksum, etc.)
- o Shared libraries - libraries that can be used by and installed with many different software applications. A shared library vulnerability could affect multiple software applications in the same way.
- o Other software configuration information - operating system or software application configuration attributes that go beyond that basic information already captured (e.g., Microsoft Windows registry, Apple configuration profiles, GConf, Proc filesystem, text configuration files and their parameters, and the installation paths.)

External vulnerability description data

- o Ingest Date - the date that the vulnerability description data was received by the enterprise.
- o Date of Release - publication or disclosure date of the vulnerability description data
- o Version - the version or iteration of the vulnerability description data according to the author, if applicable.
- o External vuln ID - external or third-party IDs assigned to the vulnerability description data. Could be multiple IDs in some cases (e.g., vendor bug id, global ID, discoverer's local ID, third-party vulnerability database ID, etc.).

- o Severity Score - the severity of the vulnerability description data according to the vulnerability description data author, if applicable.

Assessment Results

- o Date of assessment - The date that the assessment was performed against an endpoint.
- o Date of data collection - The age of the data used in the assessment to make the endpoint status determination.
- o Endpoint identification and/or locally assigned ID - The ID assigned to the enterprise endpoint. Must be assigned for tracking results over time.
- o Vulnerable software product(s) - The vulnerable software products identified as being installed on the endpoint.
- o Endpoint vulnerability status - Overall vulnerability status of the enterprise endpoint (i.e., Pass or Fail)
- o Vulnerability description - A human-consumable description of a vulnerability. Supports the human user understanding of the vulnerability assessment results within an application front-end or user interface.
- o Vulnerability remediation - The fix, workaround, or patch information for a vulnerability. This information may be a part of the vulnerability description described previously. Note that this information can change over time due to vendor patch supersession.

Appendix E. Alignment with Other Existing Works

E.1. Critical Security Controls

The Council on CyberSecurity's Critical Security Controls [critical-controls] includes security controls for a number of use scenarios, some of which are covered in this document. This section documents the alignment between the Council's controls and the relevant elements of the scenario.

E.1.1. Continuous Vulnerability Assessment

"CSC 4: Continuous Vulnerability Assessment and Remediation," which is described by the Council on CyberSecurity as "Continuously acquire, assess, and take action on new information in order to

identify vulnerabilities, remediate, and minimize the window of opportunity for attackers." The scenario described in this document is aligned with CSC 4 in multiple ways:

CSC 4-1 applies to this scenario in that it calls for running regular, automated scanning to deliver prioritized lists of vulnerabilities with which to respond. The scenario described in this document is intended to be executed on a continuous basis, and the priorities of both vulnerability description data and the remedy of vulnerabilities are discussed in the Priority section earlier in this document.

This scenario assumes that the enterprise already has a source for vulnerability description data as described in CSC 4-4.

Both CSC 4-2 and 4-7 are made possible by writing information to a Repository since this makes previously collected data available for later analysis.

While this scenario does not go into the details of how prioritization would be calculated or applied, it does touch on some of the important ways in which prioritization would impact the endpoint assessment process in the Priority section. As such, the Priority section aligns with CSC 4-10, which deals with vulnerability priority. Vulnerability priority in this scenario is discussed in terms of the vulnerability description data priority during receipt, as well as the vulnerability priority with regards to remedies.

The described scenario does not address the details of applying a remedy based on assessment results. As such, CSC 4-5, 4-8, and 4-9, which all deal with mitigations and patching, are out of scope for this work. Similarly, CSC 4-3 prescribes performing scans in authenticated mode and CSC 4-6 prescribes monitoring logs. This scenario does not get into the means by which data is collected, focusing on "what" to collect rather than "how", and as such does not have corresponding sections, although the procedures described are not incompatible with either of these controls.

The CSC 4 System Entity Relationship diagram and numbered steps directly align with the scenario described in this document with the exception of step 7 (patch response). Steps 1 -6 in CSC 4 describe the overall process for vulnerability management starting with obtaining the vulnerability description data from the source in Step 1, to producing assessment results in step 6.

E.1.2. Hardware and Software Inventories

This scenario is also aligned with, and describes a process for, collecting and maintaining hardware and software inventories, which are covered by the Council on CyberSecurity CSC 1 "Inventory of Authorized and Unauthorized Devices" and CSC 2 "Inventory of Authorized and Unauthorized Software." This scenario documents a process that is specific to collecting and maintaining hardware and software data attributes for vulnerability assessment purposes, but the collection of the hardware attributes and software inventory documented in the Endpoint Data Collection section that follows can also be used for the purpose of implementing authorized and unauthorized hardware and software management processes (e.g., scanning tools looking for unauthorized software). Moreover, the ability to accurately identify endpoints and, to a lesser degree, applications is integral to effective endpoint data collection and vulnerability management.

The Endpoint Data Collection section does not have coverage for the specific details described in CSC 1 and 2 as they are different processes and would be out-of-scope of this scenario, but the section does provide the data necessary to support the controls.

The Endpoint Identification and Endpoint Data Collection sections within this scenario align with CSC 1-1 and 1-4 by identifying enterprise endpoints and collecting their hardware and network attributes. The Endpoint Data Collection section aligns with and supports CSC 2-3 and 2-4 by defining a software inventory process and a method of obtaining operating system and file system attributes. The rest of the items from CSC 1 and 2 deal with implementation details and would be out-of-scope for this document.

CSC 2-9 describes the use of a software ID tag in XML format. SWID tags (https://en.wikipedia.org/wiki/ISO/IEC_19770) would also be a possible implementation for the Endpoint Data Collection section described in this scenario.

Appendix F. SACM Usage Scenarios

The SACM "Endpoint Security Posture Assessment: Enterprise Use Cases" document ([RFC7632]) defines multiple usage scenarios that are meant to provide examples of implementing the use cases and building block capabilities. Below is a brief summary of some of these usage scenarios and how this document aligns and/or adds additional value to the identified usage scenarios.

- o Automated Checklist Verification (2.2.2) - "An enterprise operates a heterogeneous IT environment. 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 ensure adequate coverage of all IT assets." The usage scenario, as defined in the RFC, is targeted at the checklist level and can be interpreted as being specific to endpoint configuration. There is mention of patch assessment and vulnerability mitigation, but the usage scenario could be expanded upon by including vulnerability verification. Replacing the idea of a checklist in the SACM usage scenario with vulnerability would allow the usage scenario to align almost exactly with the scenario described in this document. Instead of collecting automatable security configuration checklists, the enterprise would collect automatable vulnerability description data available from the vendor as described or possibly from other interested third-parties.

- o Detection of Posture Deviations (2.2.3) - "An enterprise has established secure configuration baselines for each different type of endpoint within their IT environment. When an endpoint connects to the network, the appropriate baseline configuration is communicated to the endpoint. 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. When the endpoint detects a posture change, an alert is generated identifying the specific changes in posture." This usage scenario would support the concept of endpoints signaling or alerting the enterprise to changes in the posture relates to endpoint vulnerabilities in the same way that it would for configurations. Replacing the idea of a checklist with vulnerability description data allows the SACM usage scenario and the scenario described in this document to align in their objectives.
- o Asynchronous Compliance/Vulnerability Assessment at Ice Station Zebra (2.2.5) - "An isolated arctic IT environment that is separated from the main university network. The only network communications are via an intermittent, low-speed, high-latency, high-cost satellite link. Remote network admins will need to show continued compliance with the security policies of the university, the government, and the provider of the satellite network, as well as keep current on vulnerability testing." This SACM usage scenario describes vulnerability assessment and aligns well with the vulnerability scenario described in this document. The endpoint assets are identified and associated data is published in a Repository. Vulnerability description data is collected and saved in a Repository as it is released. The vulnerability description data is queued for later assessment, then the

assessment results and vulnerability description data are stored after assessment. The only real difference in this SACM usage scenario is the timing of the assessments. The scenario described within this document would have no problems adjusting to the timing of this SACM usage scenario or anything similar.

Appendix G. SACM Requirements and Charter - Future Work

In the course authoring this document, some additional considerations for possible future work were noted. The following points were taken from the SACM Requirements [I-D.ietf-sacm-requirements], SACM Charter [charter-ietf-sacm-01], and SACM Use Cases ([RFC7632]) documents and represent work that may be necessary to support the tasks or goals of SACM going forward.

- o The SACM requirements mentions "Result Reporting" with applications but no detail around what the assessment results data set should include. In the case of vulnerability assessment results, context is important and details beyond just a Pass or Fail result are needed in order to take action. A good example of this might be the Priority of the vulnerability itself and how many systems it affects within the enterprise. With this in mind, it might be worthwhile to investigate a minimum data set or schema for assessment results. The concern here is with vulnerability description data, but this could apply to other enterprise processes as well.
- o The "Human-assigned endpoint attributes" mentioned previously in this scenario are touched on in the SACM use cases, but the topic could probably be explored in much more depth. Enterprise policy and behaviors could be greatly influenced by endpoint attributes such as locations, how the endpoint is used, and criticality. When and how these data attributes are collected, as well as what the minimum or common set might look like, would be good topics for future related SACM work. In addition, the storage of these attributes could be central (stored in a data repository) or they could be assigned and stored on the endpoints themselves.

Authors' Addresses

Christopher Coffin
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
USA

Email: ccoffin@mitre.org

Brant Cheikes
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
USA

Email: bcheikes@mitre.org

Charles Schmidt
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
USA

Email: cmschmidt@mitre.org

Daniel Haynes
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
USA

Email: dhaynes@mitre.org

Jessica Fitzgerald-McKay
Department of Defense
9800 Savage Road
Ft. Meade, Maryland
USA

Email: jmfitz2@nsa.gov

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

Email: david.waltermire@nist.gov

SACM
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2017

D. Haynes
The MITRE Corporation
J. Fitzgerald-McKay
Department of Defense
L. Lorenzin
Pulse Secure
March 9, 2017

Endpoint Compliance Profile
draft-haynes-sacm-ecp-02

Abstract

This document specifies the Endpoint Compliance Profile, a high-level specification that describes a specific combination and application of NEA and TNC protocols and interfaces specifically designed to support ongoing assessment of endpoint posture and the controlled exposure of collected posture information to appropriate security applications. This document is a subset of the Trusted Computing Group's Endpoint Compliance Profile Version 1.0 specification.

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 September 10, 2017.

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 (<http://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	3
1.1. Preventative Posture Assessments	4
1.2. Standardized Schema	5
1.3. Secure Standardized Protocols	5
1.4. Keywords	6
2. Terminology	6
3. Endpoint Compliance Profile	7
3.1. Posture Assessments	7
3.2. Data Storage	7
3.3. Follow-up Actions	8
4. Background	8
4.1. Purpose of the Endpoint Compliance Profile	8
4.2. Supported Use Cases	8
4.2.1. Connected and Compliant	8
4.2.2. Exposing Data to the Network	10
4.2.2.1. Asset Management	12
4.2.2.2. Vulnerability Searches	12
4.2.2.3. Threat Detection and Analysis	12
4.2.3. Non-supported Use Cases	12
4.2.4. Profile Requirements	13
4.2.5. Assumptions	14
5. Endpoint Compliance Requirements	16
5.1. Endpoint Pre-Provisioning	17
5.1.1. SWID Tags	17
5.1.2. Endpoint Identity and Machine Certificate	17
5.2. Posture Validators and Posture Collectors	17
5.2.1. SWID Posture Collectors and Posture Validators	18
5.2.1.1. The SWID Posture Collector	18
5.2.1.2. The SWID Posture Validator	18
5.3. NEA Client (NEAC) and NEA Server (NEAS)	19
5.3.1. NEAC	19
5.3.2. NEAS	19
5.4. Repository	19
6. Posture Transport Client (PTC) and Posture Transport Server (PTS)	20
7. Administrative Interface and API	20
8. Endpoint Compliance Profile Examples	21
8.1. Continuous Posture Assessment of an Endpoint	21
8.1.1. Change on Endpoint Triggers Posture Assessment	22
8.2. Administrator Searches for Vulnerable Endpoints	24

9. Acknowledgements	25
10. IANA Considerations	27
11. Security Considerations	27
11.1. Security Benefits of Endpoint Compliance Profile	27
11.2. Threat Model	29
11.2.1. Endpoint Attacks	30
11.2.2. Network Attacks	30
11.2.3. Server Attacks	30
11.2.4. Repository Attacks	31
11.3. Countermeasures	31
11.3.1. Countermeasures for Endpoint Attacks	31
11.3.2. Countermeasures for Network Attacks	32
11.3.3. Countermeasures for Server Attacks	32
11.3.4. Countermeasures for Repository Attacks	33
12. Privacy-Considerations	34
13. Change Log	34
13.1. -00 to -01	34
13.2. -01 to -02	34
14. References	34
14.1. Informative References	34
14.2. Normative References	35
Authors' Addresses	36

1. Introduction

The IETF NEA WG has defined an open architecture for network security, including standard protocols for endpoint posture assessment. The Endpoint Compliance Profile (ECP) builds on the NEA protocols, along with complementary interfaces from the Trusted Network Communications (TNC) WG of the Trusted Computing Group [TNC], to determine the posture of any type of endpoint on a network including user endpoints, servers, and infrastructure. The first generation of this specification focuses on reducing the security exposure of a network by confirming that all network-connected endpoints are:

- o known and authorized
- o running applications that are known and authorized
- o running applications that are patched and up-to-date; and,
- o applications with known vulnerabilities can be located and patched

When ECP is used, posture information is gathered by the NEA Client (NEAC) running on the endpoint and is forwarded to the NEA Server (NEAS), which stores it in a repository. This information is gathered while the endpoint is already connected to the network.

Administrators will query the repository to determine the compliance status of an endpoint. For example, if a vulnerability is discovered in a product, an administrator may query the repository to determine which endpoints have the vulnerable software installed and thus require some follow-up action.

Future versions of the ECP may want to address how to expose information--such as endpoint purpose, the software that is supposed to be running on an endpoint, and the activities an endpoint is supposed to be performing--to sensors that are looking for indicators of attacks and malicious activity on the network.

1.1. Preventative Posture Assessments

The value of continuous endpoint posture assessment is well established. Security experts have for years identified software updating and patching as a critical step for preventing intrusions. Application white listing, patching applications and operating systems, and using the latest versions of applications top the Defense Signals Directorate's "Top 4 Mitigations to Protect Your ICT System". [DSD] "Inventory of Authorized and Unauthorized Endpoints", "Inventory of Authorized and Unauthorized Software", and "Continuous Vulnerability Assessment and Remediation" are Critical Controls 1, 2, and 4, respectively, of the SANS "20 Critical Security Controls". [SANS] While there are commercially available solutions that attempt to address these security controls, these solutions do not run on all types of endpoints; consistently interoperate with other tools that could make use of the data collected; collect posture information from all types of endpoints in a consistent, standardized schema; or require vetted, standardized protocols that have been evaluated by the international community for cryptographic soundness.

As is true of most solutions offered today, the solution found in the ECP does not attempt to solve the lying endpoint problem. An endpoint that has already been infected with malicious software can provide false information about its identity and the software it is running. The primary purpose of the ECP is not to detect infected endpoints; rather, it focuses on ensuring that healthy endpoints remain healthy by keeping software up-to-date and patched. The first goal of the ECP is to help an administrator be able to readily determine which endpoints require some follow-up action. Future versions of the ECP may want to address how to expose posture information to sensors to aid the detection of attacks on endpoints and drive follow-up actions.

1.2. Standardized Schema

The ECP requires the use of standardized schema for the exchange of posture information. This helps to ensure that the posture information sent from endpoints to the repository can be easily stored, due to their known format, and shared with authorized endpoints and users. Standardized schema also enable collection from myriad types of endpoints. Such standardization saves implementers time and money--time that does not have to be spent integrating new schema into the enterprise's reporting mechanisms, and money that does not have to be spent on developing tools to parse information from each type of endpoint connected to the network. Standardized schema also enable the development of standardized client software. This allows endpoint vendors to include their own client software that can interoperate with posture assessment infrastructure and thus not have to introduce third party code in their products.

1.3. Secure Standardized Protocols

Posture information must be sent over mature, standardized protocols to ensure the confidentiality and authenticity of this data while in transit. The ECP requires use of the NEA PT-TLS protocol [RFC6876] for communication between the endpoint and the server. This protocol allows networks that implement this solution to collect large amounts of posture information from an endpoint in order to make decisions about that endpoint's compliance to some policy. This Profile offers a solution for all endpoints already connected to the network. Periodic assessments and automated reporting of changes to installed software allow for instantaneous identification of connected endpoints that are no longer compliant to some policy.

The IETF NEA WG has designed an architecture to support endpoint posture assessment. Figure 1 illustrates the architectural components used in the Endpoint Compliance Profile:




Figure 1: The Endpoint Compliance Architecture

Note that the SWID Posture Collector and SWID Posture Validator are implementations of NEA's Posture Collector (PC) and Posture Validator (PV) architectural components, respectively. Requirements for each of the components in the diagram above are contained in this profile. The reader should consult [RFC5209] for additional information on these components. All current repository requirements are contained within the Endpoint Compliance Profile.

1.4. Keywords

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 [RFC2119]. This specification does not distinguish blocks of informative comments and normative requirements. Therefore, for the sake of clarity, note that lower case instances of must, should, etc. do not indicate normative requirements.

2. Terminology

This document uses terms as defined in [I-D.ietf-sacm-terminology] unless otherwise specified.

3. Endpoint Compliance Profile

The Endpoint Compliance Profile describes how NEA and TNC specifications can be used to support the posture assessment of endpoints on a network. This profile does not generate new schema or protocols; rather, it offers a full end-to-end solution for posture assessment, as well as a fresh perspective on how existing standards can be leveraged against vulnerabilities.

3.1. Posture Assessments

The Endpoint Compliance Profile 1.0 describes how NEA and TNC specifications make it possible to perform posture assessments against all network-connected endpoints by:

1. uniquely identifying the endpoint;
2. collecting and assessing posture based on data from the endpoint;
3. creating a secure, authenticated, confidential channel between the endpoint and the server;
4. enabling the endpoint to notify the server about changes to its configuration;
5. enabling the server to request information about the configuration of the endpoint; and
6. storing the posture information in a repository linked to the identifier for the endpoint.

3.2. Data Storage

The ISO/IEC Software Identification Tag standard [SWID] has defined a schema for identifying applications installed on endpoints and their patch status. The Endpoint Compliance Profile 1.0 focuses on being able to collect this information from an endpoint and store it in a repository. This makes posture information from a network's endpoints available to authorized parties. Uses of this data are innumerable--vulnerability management, asset management, software asset management, and configuration management solutions, analytics tools, endpoints that need to make connectivity decisions, and metrics reporting scripts, among others, are all able to reference the data stored in the repository to achieve their purposes.

3.3. Follow-up Actions

The ability of the endpoint to notify the server whenever a modification is made to the endpoint enables immediate identification of endpoints that fall out of compliance. The Endpoint Compliance Profile 1.0 does not specify requirements for how these endpoints should be addressed. However, the TNC specifications do support the ability to send instructions that drive access control enforcement decisions for a non-compliant endpoint. Additional information about the types of follow-up actions an enterprise may want to support can be found in [RFC7632].

There is a clear need for nuanced, automated instructions sent from the server to the endpoint (for example, to update an endpoint's software, or remove a piece of non-compliant software). Those messages are complicated to define and may have to be tailored to a particular operating system. Future versions of this specification may want to address which instructions can be defined based on the configuration content that is collected from endpoints.

4. Background

4.1. Purpose of the Endpoint Compliance Profile

The Endpoint Compliance Profile describes a standard way to communicate endpoint posture information such as software identity and software version and to make it available to other authorized parties. The Endpoint Compliance Profile 1.0 focuses on collecting the application information available in SWID tags, as specified in [SWID]. Future versions of the Endpoint Compliance Profile could describe how additional types of posture information can be collected and communicated in a standardized way.

4.2. Supported Use Cases

The Endpoint Compliance Profile focuses on the posture assessment of enterprise endpoints on enterprise networks. Use cases supported by the Endpoint Compliance Profile 1.0 are as follows:

4.2.1. Connected and Compliant

A network-connected endpoint sends posture information using standard schemas such as SWID over NEA protocols.




Figure 2: Connected and Compliant Use Case

1. If necessary, the endpoint finds and validates the server in compliance with [Server-Discovery].
2. The Posture Transport Client (PTC) on the endpoint and Posture Transport Server (PTS) on the server complete a TLS handshake, during which endpoint identity information is exchanged.
3. Either the NEA Server (NEAS) on the server or the NEA Client (NEAC) on the endpoint initiates a posture assessment. Checks may be triggered for multiple reasons, including:
 - (a) policy states that a previous assessment has aged out and become invalid;

- (b) the NEAC notices that the relevant posture information on the endpoint has changed, (for example, due to application updates, deletions or additions); or
- (c) the NEAS is alerted by a sensor or an administrator (via the server's user interface) that an assessment must be completed.

All information exchanges between the PCs and PVs are subject to the enterprise's policy, which may limit the content or size of information sent between the endpoint and the server.

4. The SWID Posture Collector on the endpoint collects from the SWID tag directory on the endpoint. This data is sent via the NEAC and PTC to the server.
5. Once the posture information is received by the PTS, it is forwarded to the SWID Posture Validator via the NEAS. The SWID Posture Validator also forwards the posture information to the repository. The posture information is stored along with past posture information collected about the endpoint.

4.2.2. Exposing Data to the Network

Because the endpoint posture information was sent in a standards-based schema (ISO/IEC 19770-2:2009) over secure, standardized protocols, and the SWID tags are stored in a centralized repository linked to unique endpoint identifiers, authorized parties are able to access the posture information. Such authorized parties may include, but are not limited to, administrators or endpoint owners (via the server's administrative interface), and other pieces of infrastructure that can make use of this data (via the server's API). The server will provide:

- o a standard administrative interface that allows data sharing with authorized parties;
- o a standard API that allows data sharing with authorized infrastructure and software;
- o a persistent account of endpoints that have connected to the network over a period of time set by the administrator;
- o the identities provided by those endpoints; and
- o what SWIDs were reported by the endpoint.

The endpoint will publish updates as its local SWID directory changes, as well as each time it disconnects and reconnects to the network.

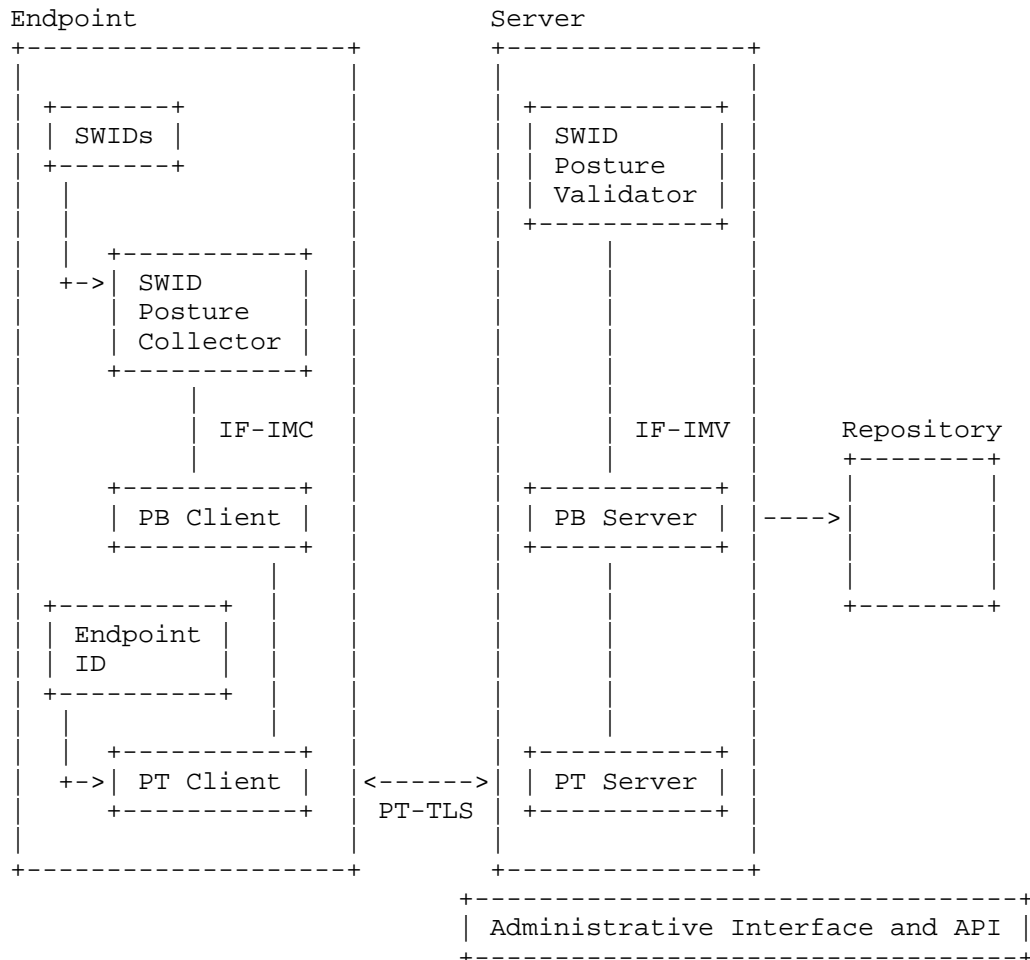


Figure 3: Exposing Data to the Network

It should be noted that the neither the Endpoint Compliance Profile nor the protocols, interfaces, and data models that it references provide solutions to the server capabilities listed above. However, these capabilities are useful and solutions for them should be pursued in the future.

4.2.2.1. Asset Management

Using the administrative interface on the server, an authorized user can learn:

- o what endpoints are connected to the network at any given time; and
- o what SWID tags were reported for the endpoints.

The ability to answer these questions offers a standards-based approach to asset management, which is a vital part of enterprise processes such as compliance report generation for the Federal Information Security Modernization Act (FISMA), Payment Card Industry Data Security Standard (PCI DSS), Health Insurance Portability and Accountability Act (HIPAA), etc.

4.2.2.2. Vulnerability Searches

The administrative interface also provides the ability for authorized users or infrastructure to locate endpoints running software for which vulnerabilities have been announced. Because of

1. the unique IDs assigned to each endpoint; and
2. the rich application data provided in the endpoints' posture information,

the repository can be queried to find all endpoints running a vulnerable application. Endpoints suspected of being vulnerable can be addressed by the administrator or flagged for further scrutiny.

4.2.2.3. Threat Detection and Analysis

The repository's standardized API allows authorized infrastructure endpoints and software to search endpoint posture assessment information for evidence that an endpoint's software inventory has changed, and can make endpoint software inventory data available to other endpoints. This automates security data sharing in a way that expedites the correlation of relevant network data, allowing administrators and infrastructure endpoints to identify odd endpoint behavior and configuration using secure, standards-based schema and protocols.

4.2.3. Non-supported Use Cases

Several use cases, including but not limited to these, are not covered by the Endpoint Compliance Profile 1.0:

- o Gathering other types of posture information: The Endpoint Compliance Profile does not prevent administrators from collecting other types of posture information other than SWIDs from the endpoint; however it does not set requirements for doing so.
- o Solving the lying endpoint problem: The Endpoint Compliance Profile does not address the lying endpoint problem; the Profile makes no assertions that it can catch an endpoint that is, either maliciously or accidentally, reporting false posture information to the server. However, other solutions may be able to use the posture information collected using the capabilities described in this profile to catch an endpoint in a lie. For example, a sensor may be able to compare the posture information it has collected on an endpoint's activity on the network to what the endpoint reported to the server and flag discrepancies. However, these particular capabilities are not described in this profile.
- o Publish/subscribe repository interface: Future versions of the Endpoint Compliance Profile may specify a publish/subscribe interface for the repository, so infrastructure endpoint can subscribe to and receive published posture assessment results from the repository regarding endpoint configuration changes. However, the Endpoint Compliance Profile 1.0 includes no such requirements.

4.2.4. Profile Requirements

Here are the requirements that the Endpoint Compliance Profile protocol must meet in order to successfully fit in the SACM architecture.

- o Meets the needs of the SACM architecture: The Endpoint Compliance Profile must support the use cases described in [RFC7632] as they apply to endpoint self-reporting and endpoint posture assessment.
- o Efficient: To minimize user frustration, it is essential to minimize delays by making endpoint posture information collection, transmission, and assessment as brief and efficient as possible.
- o Extensible: The Endpoint Compliance Profile needs to expand over time as new features are added to the SACM architecture. The solution must allow new features to be added easily, providing for a smooth transition and allowing newer and older architectural components to continue to work together. Further, the Endpoint Compliance Profile and the specifications referenced here must define safe extensibility mechanisms that enable innovation without breaking interoperability.

- o Easy to implement: The Endpoint Compliance Profile should be easy for vendors to implement in their products, and should result in products that are easy for administrators to implement on their networks. Products conformant to the Endpoint Compliance Profile should interoperate seamlessly, and be simple to integrate into existing network infrastructure.
- o Easy to use: The Endpoint Compliance Profile should describe a simple, integrated user interface that administrators can use to perform the activities listed in the profile's use cases. The Endpoint Compliance Profile should not constrain innovation by specifying details of the user interface but rather functional requirements.
- o Platform-independent: Since network environments may contain many different types of endpoints, the solution should operate independently of the endpoint platform.
- o Scalable: The Endpoint Compliance Profile must be designed to scale to very large numbers of endpoints.

4.2.5. Assumptions

Here are the assumptions that the Endpoint Compliance Profile makes about other components in the SACM architecture.

- o Existence of a server and repository: The Endpoint Compliance Profile assumes that a server and repository exist.
- o Endpoint SWID installation: The Endpoint Compliance Profile assumes that an endpoint has been pre-provisioned with Software Identification Tags for its applications, and that these SWID tags are formatted and stored in conformance with [SWID].
- o Certificate provisioning: In order to implement the most secure endpoint identification option, the Endpoint Compliance Profile assumes that the enterprise has set up a certificate root authority, and has provisioned each endpoint with an endpoint identification certificate. This is not required if an enterprise chooses to use other endpoint authentication methods.

In addition, the Endpoint Compliance Profile makes the following assumptions about the SACM ecosystem:

- o All network-connected endpoints are endpoints: As defined by [I-D.ietf-sacm-terminology], an endpoint is any physical or virtual computing endpoint that can be connected to a network. Posture assessment against policy is equally, if not more,

important for continuously connected endpoints, such as enterprise workstations and infrastructure endpoints, as it is for sporadically connected endpoints. Continuously connected endpoints are just as likely to fall out of compliance with policy, and a standardized posture assessment method is necessary to ensure they can be properly handled.

- o All endpoints on the network must be uniquely identified: Many administrators struggle to identify what endpoints are connected at any given time. By requiring a standardized method of endpoint identity, the Endpoint Compliance Profile will enable administrators to answer the basic question, "What is on my network?" Unique endpoint identification also enables the comparison of current and past endpoint posture assessments, by allowing administrators to correlate assessments from the same endpoint. This makes it easier to flag suspicious changes in endpoint posture for manual or automatic review, and helps to swiftly identify malicious changes to endpoint applications.
- o Posture assessments must occur over secure, standardized protocols: Endpoint identity and application information is very valuable, both to administrators and to attackers. Therefore, it must be kept confidential, using secure protocols to transport it from the endpoint to network infrastructure endpoints. Additionally, it is critical that only authorized parties be capable of requesting information, receiving information, or taking action to change an endpoint's connectivity status. Relying on standardized protocols to provide this security enables greater interoperability and compatibility between endpoints, and allows for the development of compliance testing to ensure that each endpoint operates securely and in conformance with appropriate specifications. A standards body provides a process for experts in protocols and cryptography to evaluate the soundness of protocols and security management procedures; a set of security standards allows an enterprise to make the most effective use of their investment in a security management infrastructure.
- o Posture assessment results must be formatted using standardized schema: Well-known, standard schema allow for a universal language for generating compliance reports. With each endpoint speaking the same language, the Endpoint Compliance Profile enables information sharing between user endpoints and infrastructure endpoints, and between infrastructure endpoints that perform different security tasks.
- o Posture information must be stored by the repository and must be exposed to an interface at the server: A standard schema enables

standard queries from an interface exposed to an administrator at the server console. A repository must retain any current posture information retrieved from the endpoint and store it indexed by the unique identifier for the endpoint. Any PV specified by this profile must be able to ascertain from its corresponding PC whether the posture information is up to date. An interface on the server must support a request to the PV to obtain up-to-date information when an endpoint is connected. This interface must also support the ability to make a standard set of queries about the posture information stored by the repository. In the future, some forms of posture information might be retained at the endpoint. The interface on the server must accommodate the ability to make a request through the PV to the corresponding PC about the posture of the endpoint. Standard schema and protocols also enable the security of posture assessment results. By storing these results indexed under the endpoint's unique identification, secure storage itself enables endpoint posture information correlation, and ensures that the enterprise's Repositories always offer the freshest, most up-to-date view of the enterprise's endpoint posture information possible.

- o Posture information can be shared: By exposing posture information using a standard interface and API, other security and operational components have a high level of insight into the enterprise's endpoints and the software installed on them. This will support innovation in the areas of asset management, vulnerability scanning, and administrative interfaces, as any authorized infrastructure endpoint can interact with the posture information.
- o Owners and administrators must have complete control of posture information, policy, and endpoint mitigation: Enterprise asset posture information belongs to the enterprise. Standardized schema, protocols and interfaces help to ensure that this posture information is not locked in proprietary databases, but is made available to its owners. This enables administrators to develop as nuanced a policy as necessary to keep their networks secure.

5. Endpoint Compliance Requirements

These requirements are written with a view to performing a posture assessment on an endpoint; as the Endpoint Compliance Profile grows and evolves, these requirements will be expanded to address issues that arise. Note that these requirements refer to defined components of the NEA architecture. As with the NEA architecture, implementers have discretion as to how these NEA components map to separate pieces of software or endpoints.

5.1. Endpoint Pre-Provisioning

The following requirements assume that the platform or OS vendor supports the use of SWID tags and has identified a standard directory location for the SWID tags to be located as specified by [SWID].

5.1.1. SWID Tags

The primary content for the Endpoint Compliance Profile 1.0 is the information conveyed in the elements of a SWID tag.

The endpoint MUST have SWID tags stored in a directory specified in [SWID]. The tags SHOULD be provided by the software vendor; they MAY also be generated by:

- o the software installer; or
- o third-party software that creates tags based on the applications it sees installed on the endpoint.

The elements in the SWID tag MUST be populated as specified in [SWID]. These tags, and the directory in which they are stored, MUST be updated as software is added, removed, or updated.

5.1.2. Endpoint Identity and Machine Certificate

The endpoint SHOULD authenticate to the server using a machine certificate during the establishment of the outer tunnel achieved with PT. [IF-IMV] specifies how to pull an endpoint ID out of a machine certificate. An endpoint ID SHOULD be created in conformance with [IF-IMV] from a machine certificate sent via [RFC6876].

In the future, the identity could be a hardware certificate compliant with [IEEE-802-1ar]; ideally, this ID SHOULD be associated with the identity of a hardware cryptographic module, in accordance with [IEEE-802-1ar], if present on the endpoint. The enterprise SHOULD stand up a certificate root authority; install its root certificate on endpoints and on the server; and provision the endpoints and the server with machine certificates. The endpoint MAY authenticate to the server using a combination of the machine account and password; however, this is less secure and not recommended.

5.2. Posture Validators and Posture Collectors

Any PC used in an Endpoint Compliance Profile solution MUST be conformant with [IF-IMC]; an Internet-Draft, under development, that is a subset of the TCG TNC Integrity Measurement Collector interface [IF-IMC] and will be submitted in the near future. Any Posture

Validator used in an Endpoint Compliance Profile solution MUST be conformant with [IF-IMV].

5.2.1. SWID Posture Collectors and Posture Validators

5.2.1.1. The SWID Posture Collector

For the Endpoint Compliance Profile, the SWID Posture Collector MUST be conformant with [I-D.ietf-sacm-nea-swid-patnc], which includes requirements for:

1. Collecting SWID tags from the SWID directory
2. Monitoring the SWID directory for changes
3. Initiating a session with the server to report changes to the directory
4. Maintaining a list of changes to the SWID directory when updates take place and no PT-TLS connection can be created with the server
5. Responding to a request for SWID tags from the SWID Posture Validator on the server
6. Responding to a query from the SWID Posture Validator as to whether all updates have been sent

The SWID Posture Collector is not responsible for detecting that the SWID directory was not updated when an application was either installed or uninstalled.

5.2.1.2. The SWID Posture Validator

Conformance to [I-D.ietf-sacm-nea-swid-patnc] enables the SWID Posture Validator to:

1. Send messages to the SWID Posture Collector (at the behest of the administrator at the server console) requesting updates for SWID tags located on endpoint
2. Ask the SWID Posture Collector whether all updates to the SWID directory located at the server have been sent
3. Compare an endpoint's SWID posture information to policy, and make a recommendation to the NEAS about the endpoint

In addition to these requirements, a SWID Posture Validator used in conformance with this profile MUST be capable of passing information from the posture assessment results and the endpoint identity associated with those results to the repository for storage.

5.3. NEA Client (NEAC) and NEA Server (NEAS)

[RFC5793] describes a standard way for the NEAC and the NEAS to exchange messages.

5.3.1. NEAC

The NEAC MUST conform to [RFC5793], which levies a number of requirements against the NEAC. A NEAC that complies with these requirements will be able to:

1. attempt to initiate a session with the NEAS if the SWID Posture Collector makes a request to send an update to the SWID directory to the server;
2. notify the SWID Posture Collector if no PT-TLS session with the server can be created;
3. notify the SWID Posture Collector when a PT-TLS session is established; and
4. receive information from the PCs, forward this information to the server via the PTC.

The NEAC MUST also conform to [IF-IMC] to enable communications with the SWID Posture Collector.

5.3.2. NEAS

The NEAS MUST conform to all requirements in the [RFC5793] and [IF-IMV] specifications. Conformance to [IF-IMV] enables the NEAS to obtain endpoint identity information from the PTS, and pass this information to any IMVs on the server.

5.4. Repository

ECP 1.0 requires a simple administrative interface for the repository. PVs on the server receive the endpoint data via PA-TNC [RFC5792] messages sent from corresponding PCs on an endpoint and store this information in the repository linked to the identity of the endpoint where the PCs are located.

The administrative interface SHOULD enable an administrator to:

1. Query which endpoints have reported SWID tags for a particular application
2. Query which SWID tags are installed on a particular endpoint
3. Query tags based on characteristics, such as vendor, publisher, etc.

In the future, if SACM decides to develop an interface to the repository server, it should consider requirements for:

1. Creating a secure channel between a publisher and the repository
2. Creating a secure channel between a subscriber and the repository
3. The types of interactions that must be supported between publishers and subscribers to a repository

6. Posture Transport Client (PTC) and Posture Transport Server (PTS)

The PT-TLS protocol provides a transport service for carrying the PB-TNC protocol messages between the endpoint and the server.

The PTC and PTS MUST implement PT-TLS, since a connection is needed that:

- o Can handle large volumes of data, which might require multiple roundtrips, to be sent while the endpoint is connected
- o Allows either the NEAC or NEAS to initiate a connection
- o Supports secure transport based on machine certificates at both ends of the connection

The PTC and PTS MUST support the use of machine certificates for TLS at each endpoint consistent with the requirements stipulated in [RFC6876] and [Server-Discovery].

The PTC MUST be able to locate an authorized server, and switch to a new server when required by the network, in conformance with [Server-Discovery].

7. Administrative Interface and API

An interface is necessary to allow administrators to manage the endpoints and software used in the Endpoint Compliance Profile. This interface SHOULD be accessible either on or through (as in the case

of a remotely hosted interface) the server. Using this interface, an authorized user or administrator SHOULD be able to:

- o Query the repository
- o Send commands to the PVs, requesting information from the associated PCs residing on network endpoints
- o Update the policy that resides on the server

An API is necessary to allow infrastructure endpoints and software access to the information stored in the repository. Using this API, an authorized endpoint SHOULD be able to:

- o Query the repository

8. Endpoint Compliance Profile Examples

8.1. Continuous Posture Assessment of an Endpoint




Figure 4: Continuous Posture Assessment of an Endpoint

8.1.1.1. Change on Endpoint Triggers Posture Assessment

A new application is installed on the endpoint, and the SWID directory is updated. This triggers an update from the SWID Posture Collector to the SWID Posture Validator. The message is sent down the NEA stack, encapsulated by NEA protocols until it is sent by the PTC to the PTS. The PTS then forwards it up through the stack, where the layers of encapsulation are removed until the SWID Message arrives at the SWID Posture Validator.

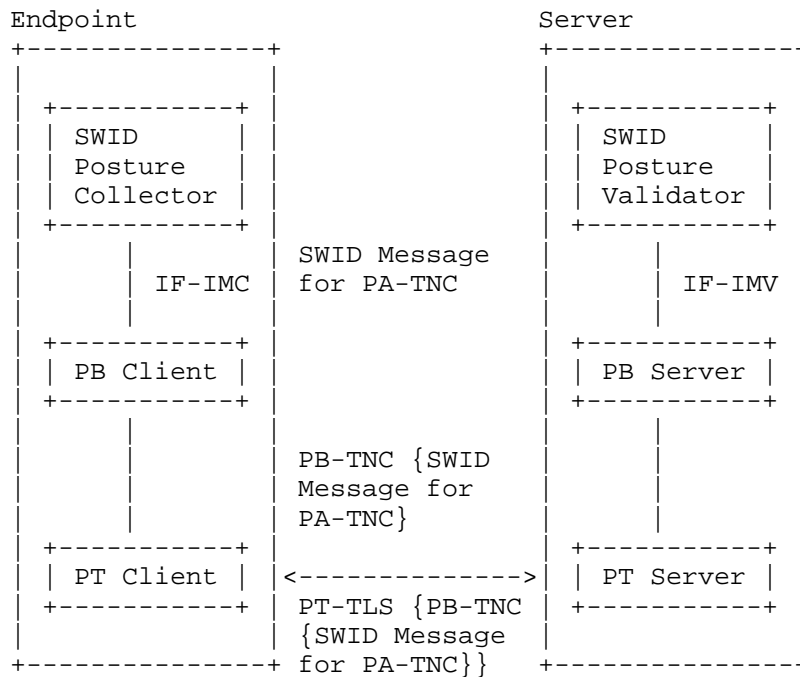


Figure 5: Compliance Protocol Encapsulation

The SWID Posture Validator stores the new tag information in the repository. If the tag indicates that the endpoint is compliant to the policy, then the process is complete until the next time an update is needed (either because policy states that the endpoint must submit posture assessment results periodically or because an install/uninstall/update on the endpoint triggers a posture assessment).

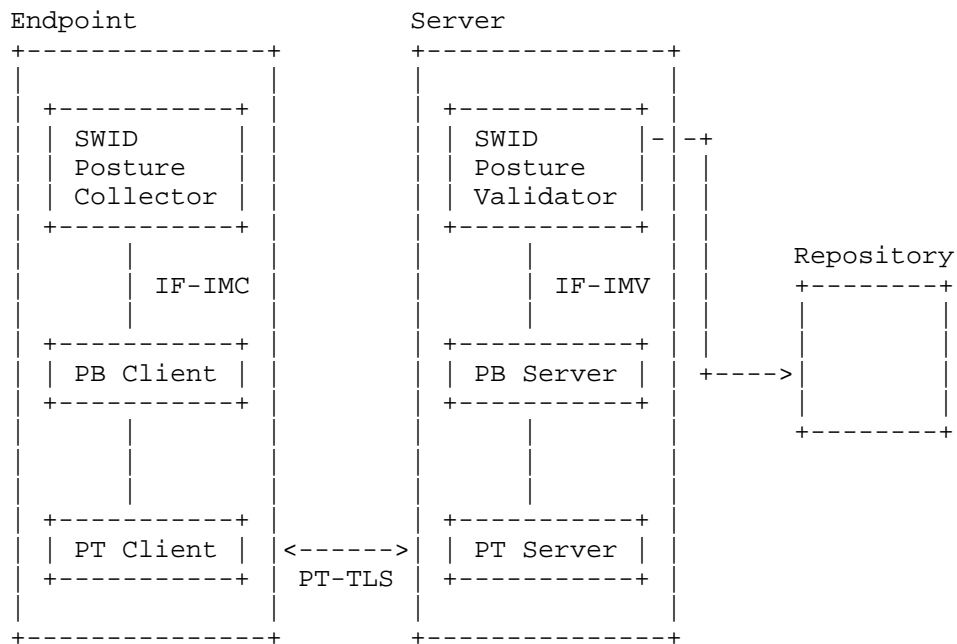


Figure 6: Storing SWIDs in the Repository

If the endpoint has fallen out of compliance with a policy, the server can alert the administrator via the server's administrative interface. The administrator can then take steps to address the problem. If the administrator has already established a policy for automatically addressing this problem, that policy will be followed.




Figure 7: Server Alerts Network Admin

8.2. Administrator Searches for Vulnerable Endpoints

An announcement is made that a particular version of a piece of software has a vulnerability. The administrator uses the Administrative Interface on the server to search the repository for endpoints that reported the SWID tag for the vulnerable software.

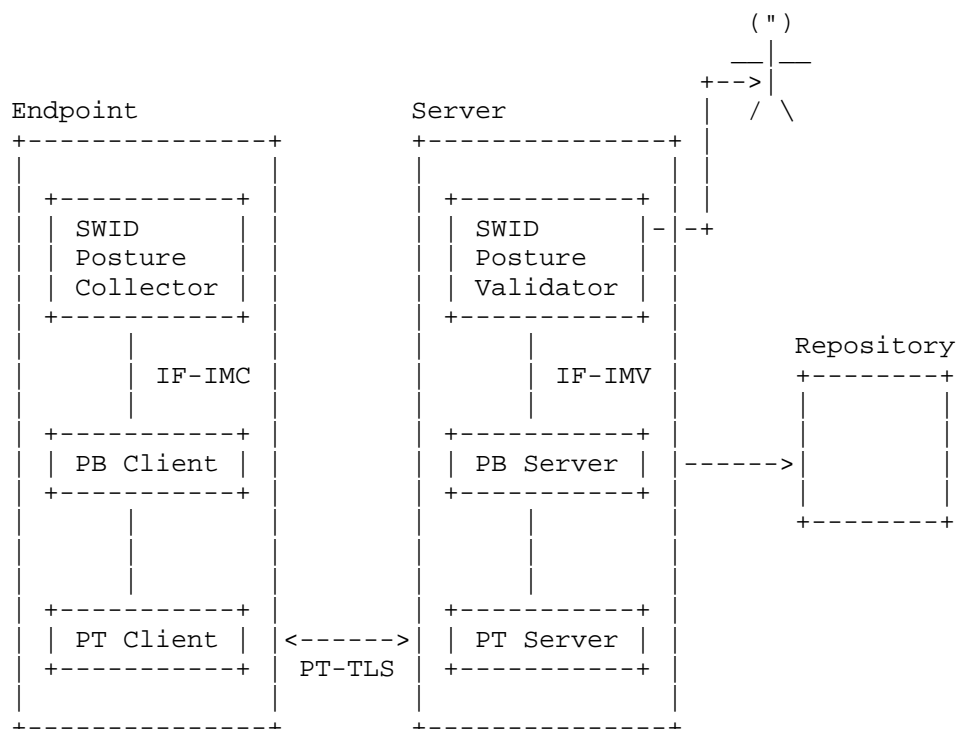


Figure 8: Admin Searches for Vulnerable Endpoints

The repository returns a list of entries in the matching the administrator's search. The administrator can then address the vulnerable endpoints by taking some follow-up action such as removing it from the network, quarantining it, or updating the vulnerable software.

9. Acknowledgements

The authors wish to thank all of those in the TCG TNC work group who contributed to development of the TNC ECP specification upon which this document is based.

Member	Organization
Padma Krishnaswamy	Battelle Memorial Institute
Eric Fleischman	Boeing
Richard Hill	Boeing

Steven Venema	Boeing
Nancy Cam-Winget	Cisco Systems
Scott Pope	Cisco Systems
Max Pritikin	Cisco Systems
Allan Thompson	Cisco Systems
Nicolai Kuntze	Fraunhofer Institute for Secure Information Technology (SIT)
Ira McDonald	High North
Dr. Andreas Steffen	HSR University of Applied Sciences Rapperswil
Josef von Helden	Hochschule Hannover
James Tan	Infoblox
Steve Hanna (TNC-WG Co-Chair)	Juniper Networks
Cliff Kahn	Juniper Networks
Lisa Lorenzin	Juniper Networks
Atul Shah (TNC-WG Co- Chair)	Microsoft
Jon Baker	MITRE
Charles Schmidt	MITRE
Rainer Enders	NCP Engineering
Dick Wilkins	Phoenix Technologies
David Waltermire	NIST
Mike Boyle	U.S. Government
Emily Doll	U.S. Government
Jessica Fitzgerald-	U.S. Government

McKay	
Mary Lessels	U.S. Government
Chris Salter	U.S. Government

Table 1: Members of the TNC Work Group that Contributed to the Document

10. IANA Considerations

This document does not define any new IANA registries. However, this document does reference other documents that do define IANA registries. As a result, the IANA Considerations section of the referenced documents should be consulted.

11. Security Considerations

The Endpoint Compliance Profile offers substantial improvements in endpoint security, as evidenced by the Australian Defense Signals Directorate's analysis that 85% of targeted cyber intrusions can be prevented through application white listing, patching applications and operating systems, and using the latest versions of applications. [DSD] Despite these gains, some security risks continue to exist and must be considered.

To ensure that these benefits and risks are properly understood, this Security Considerations section includes an analysis of the benefits provided by the Endpoint Compliance Profile (Section 11.1), the attacks that may be mounted against systems that implement the Endpoint Compliance Profile (Section 11.2), and the countermeasures that may be used to prevent or mitigate these attacks (Section 11.3). Overall, a substantial reduction in cyber risk can be achieved.

11.1. Security Benefits of Endpoint Compliance Profile

Security weaknesses of the components for this profile should be considered in light of the practical considerations that must be addressed to have a viable solution.

Posture assessment has two parts: assessment and follow-up actions. The point of posture assessment is to ensure that authorized users are using authorized software configured to be as resilient as possible against an attack.

Posture assessment answers the question whether the endpoint is healthy. Our goal for posture assessment is to make it harder for an

adversary to execute code on one of our endpoints. This profile represents an important first step in reaching that goal. If we keep our endpoints healthier, we are able to prevent more attacks on our endpoints and thus on our information systems.

The goal of ECP is to address posture assessment in stages. Stage 1 is the ability to ascertain whether all endpoints are authorized and whether all applications are authorized and up to date. Stage 2 will attempt to address the harder problem of whether all software is configured safely. Eventually, the goal is to also address remediation which is currently out-of-scope for the SACM WG; that presents a far greater security challenge than reporting, since remediation implies the ability of a remote party to modify software or its settings on endpoints.

A second security consideration is how to gain visibility over every type of endpoint and every piece of software installed on the endpoint. This is a problem of scaling and observation. A solution is needed that can report from every type of endpoint. All software on the endpoint has to be discovered. Information about the software has to be up to date and accurate. The information that is discovered has to be reported in a consistent format, so administrators do not have to squander time deciphering proprietary systems and the information can be made readily useful for other security automation purposes.

ECP is based on a model of a standards-based schema, a standards-based set of protocols and interfaces, and the existence of an oversight group, the IETF, that can update the schemas and protocols to meet new use cases and security issues that may be discovered.

The data elements in the schema determine what work can be done consistently for every endpoint and every piece of software. How the data gets populated is an important consideration. ECP leverages the SWID tags from ISO 19770-2 because the tag originates with a single authoritative source, the application vendor itself. Moreover, there is a natural incentive for the vendor to create this content, since it makes it easier for enterprises and vendors to track whether software is licensed. Practical considerations are security considerations. A sustainable business model for obtaining all the necessary content is a fundamental requirement.

The NEA model is based on having a NEAC run on an endpoint that publishes posture information to a server. The advantages are easy to list. A platform vendor can implement its own NEAC and have it be compatible with the NEAS from a different vendor. The interfaces are layered on top of mature protocols such as TLS. TLS is the protocol of choice for ECP, since:

- o it has proven secure properties,
- o it can be implemented on most types of endpoints,
- o it allows the gathering of large amounts of information when a endpoint is connected, and
- o it enables use of a mechanism to ensure that the client is authenticated (authorized) - a client certificate - which also provides a consistent identifier.

Mature protocols that can be implemented on most types of endpoints and a standards-based schema with a sustainable business model are both critical security considerations for compliance.

Additionally, it is important to consider the future stages for ECP such as a posture assessment being followed up by some action (e.g. remediation, alert, etc.). Ensuring that clients are taking instructions only from authorized parties will be critical. Inasmuch as it is practical, enterprises will want to use the same infrastructure and investment in PKI to send those instructions to a client.

Likewise, as more information with more value is gathered from endpoints, we will also want to ensure that this information is only released to authorized applications and parties. For the next stage of ECP, SACM may want to define an interface on the repository that can be queried by other security automation applications to make it easier to detect attacks and for other security automation applications. This interface has to be standards-based for enterprises to reap the benefits of innovation that can be achieved by making the enterprise's data available to other tools and services.

11.2. Threat Model

This section lists the attacks that can be mounted on an Endpoint Compliance Profile environment. The following section (Section 11.3) describes countermeasures.

Because the Endpoint Compliance Profile describes a specific use case for NEA components, many security considerations for these components are addressed in more detail in the technical specifications: [I-D.ietf-sacm-nea-swid-patnc], [IF-IMC], [RFC5793], [Server-Discovery], [RFC6876], [IF-IMV].

11.2.1. Endpoint Attacks

While the Endpoint Compliance Profile provides substantial improvements in endpoint security as described in Section 11.1, a certain percentage of endpoints will always get compromised. For this reason, all parties must regard data coming from endpoints as potentially unreliable or even malicious. An analogy can be drawn with human testimony in an investigation or trial. Human testimony is essential but must be regarded with suspicion.

- o Compromise of endpoint: A compromised endpoint may report false information to confuse or even provide maliciously crafted information with a goal of infecting others.
- o Putting bad information in SWID directory: Even if an endpoint is not completely compromised, some of the software running on it may be unreliable or even malicious. This software, potentially including the SWID generation or discovery tool, or malicious software pretending to be a SWID generation or discovery tool, can place incorrect or maliciously crafted information into the SWID directory. Endpoint users may even place such information in the directory, whether motivated by curiosity or confusion or a desire to bypass restrictions on their use of the endpoint.
- o Identity spoofing (impersonation): A compromised endpoint may attempt to impersonate another endpoint to gain its privileges or to besmirch the reputation of that other endpoint.

11.2.2. Network Attacks

A variety of attacks can be mounted using the network. Generally, the network cannot be trusted.

- o Eavesdropping, modification, injection, replay, deletion
- o Traffic analysis
- o Denial of service and blocking traffic

11.2.3. Server Attacks

The server is a critical security element and therefore merits considerable scrutiny.

- o Compromised trusted server: A compromised server or a malicious party that is able to impersonate a server can incorrectly grant or deny access to endpoints, place incorrect information into the repository, or send malicious messages to endpoints

- o Misconfiguration of trusted server: Accidental or purposeful misconfiguration of a trusted server can cause effects that are similar to those listed for compromised trusted server.
- o Malicious untrusted server: An untrusted server cannot mount any significant attacks because all properly implemented endpoints will refuse to engage in any meaningful dialog with such a server.

11.2.4. Repository Attacks

The repository is also an important security element and therefore merits careful scrutiny.

- o Putting bad information into trusted repository: An authorized repository client such as a server may be able to put incorrect information into a trusted repository or delete or modify historical information, causing incorrect decisions about endpoint security. Placing maliciously crafted data in the repository could even lead to compromise of repository clients, if they fail to carefully check such data.
- o Compromised trusted repository: A compromised trusted repository or a malicious untrusted repository that is able to impersonate a trusted repository can lead to effects similar to those listed for "Putting bad information into trusted repository". Further, a compromised trusted repository can report different results to different repository clients or deny access to the repository for selected repository clients.
- o Misconfiguration of trusted repository: Accidental or purposeful misconfiguration of a trusted repository can deny access to the repository or result in loss of historical data.
- o Malicious untrusted repository: An untrusted repository cannot mount any significant attacks because all properly implemented repository clients will refuse to engage in any meaningful dialog with such a repository.

11.3. Countermeasures

This section lists the countermeasures that can be used in an Endpoint Compliance Profile environment.

11.3.1. Countermeasures for Endpoint Attacks

This profile is in and of itself a countermeasure for a compromised endpoint. A primary defense for an endpoint is to run up to date software configured to be run as safely as possible.

Ensuring that anti-virus signatures are up to date and that a firewall is configured are also protections for an endpoint that are supported by the current NEA specifications.

Endpoints that have hardware cryptographic modules that are provisioned by the enterprise, in accordance with [IEEE-802-1ar], can protect the private keys used for authentication and help prevent adversaries from stealing credentials that can be used for impersonation. Future versions of the Endpoint Compliance Profile may want to discuss in greater detail how to use a hardware cryptographic module, in accordance with [IEEE-802-1ar], to protect credentials and to protect the integrity of the code that executes during the bootstrap process.

11.3.2. Countermeasures for Network Attacks

To address network attacks, [RFC6876] includes required encryption, authentication, integrity protection, and replay protection. [Server-Discovery] also includes authorization checks to ensure that only authorized servers are trusted by endpoints. Any unspecified or not yet specified network protocols employed in the Endpoint Compliance Profile (e.g. the protocol used to interface with the repository) should include similar protections.

These protections reduce the scope of the network threat to traffic analysis and denial of service. Countermeasures for traffic analysis (e.g. masking) are usually impractical but may be employed. Countermeasures for denial of service (e.g. detecting and blocking particular sources) SHOULD be used when appropriate to detect and block denial of service attacks. These are routine practices in network security.

11.3.3. Countermeasures for Server Attacks

Because of the serious consequences of server compromise, servers SHOULD be especially well hardened against attack and minimized to reduce their attack surface. They SHOULD be monitored using the NEA protocols to ensure the integrity of the behavior and analysis data stored on the server and SHOULD utilize a [IEEE-802-1ar] compliant hardware cryptographic module for identity and/or integrity measurements of the server. They should be well managed to minimize vulnerabilities in the underlying platform and in systems upon which the server depends. Network security measures such as firewalls or intrusion detection systems may be used to monitor and limit traffic to and from the server. Personnel with administrative access to the server should be carefully screened and monitored to detect problems as soon as possible. Server administrators should not use password-based authentication but should instead use non-reusable credentials

and multi-factor authentication (where available). Physical security measures should be employed to prevent physical attacks on servers.

To ease detection of server compromise should it occur, server behavior should be monitored to detect unusual behavior (such as a server reboot, unusual traffic patterns, or other odd behavior). Endpoints should log and/or notify users and/or administrators when peculiar server behavior is detected. To aid forensic investigation, permanent read-only audit logs of security-relevant information pertaining to servers (especially administrative actions) should be maintained. If server compromise is detected, the server's certificate should be revoked and careful analysis should be performed of the source and impact of this compromise. Any reusable credentials that may have been compromised should be reissued.

Endpoints can reduce the threat of server compromise by minimizing the number of trusted servers, using the mechanisms described in [Server-Discovery].

11.3.4. Countermeasures for Repository Attacks

If the host for the repository is located on its own endpoint, it should be protected with the same measures taken to protect the server. In this circumstance, all messages between the server and repository should be protected with a mature security protocol such as TLS or IPsec.

The repository can aid in the detection of compromised endpoints if an adversary cannot tamper with its contents. For instance, if an endpoint reports that it does not have an application with a known vulnerability installed, an administrator can check whether the endpoint might be lying by querying the repository for the history of what applications were installed on the endpoint.

To help prevent tampering with the information in the repository:

1. Only authorized parties should have privilege to run code on the endpoint and to change the repository.
2. If a separate endpoint hosts the repository, then the functionality of that endpoint should be limited to hosting the repository. The firewall on the repository should only allow access to the server and to any endpoint authorized for administration.
3. The repository should ideally use "write once" media to archive the history of what was placed in the repository, to include a snapshot of the current status of applications on endpoints.

12. Privacy-Considerations

The Endpoint Compliance Profile specifically addresses the collection of posture data from enterprise endpoints by an enterprise network. As such, privacy is not going to often arise as a concern for those deploying this solution.

A possible exception may be the concerns a user may have when attempting to connect a personal endpoint (such as a phone or mobile endpoint) to an enterprise network. The user may not want to share certain details, such as an endpoint identifier or SWID tags, with the enterprise. The user can configure their NEAC to reject requests for this information; however, it is possible that the enterprise policy will not allow the user's endpoint to connect to the network without providing the requested data.

13. Change Log

13.1. -00 to -01

There are no textual changes associated with this revision. This revision simply reflects a resubmission of the document so that it remains in active status.

13.2. -01 to -02

Added references to the Software Inventory Message and Attributes (SWIMA) for PA-TNC I-D.

Replaced references to PC-TNC with IF-IMC.

Removed erroneous hyphens from a couple of section titles.

Made a few minor editorial changes.

14. References

14.1. Informative References

[DSD] http://www.dsd.gov.au/publications/csocprotect/top_4_mitigations.htm, "Top 4 Mitigation Strategies to Protect Your ICT System", November 2012.

[IEEE-802-1ar]
Institute of Electrical and Electronics Engineers, "IEEE 802.1ar", December 2009.

- [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>>.
- [SANS] <http://www.sans.org/critical-security-controls/>, "CIS Critical Security Controls".
- [TNC] Trusted Computing Group, "TCG Trusted Network Connect TNC Architecture for Interoperability, Version 1.5", February 2012.

14.2. Normative References

- [I-D.ietf-sacm-nea-swid-patnc] Schmidt, C., Haynes, D., Coffin, C., and J. Fitzgerald-McKay, "Software Inventory Message and Attributes (SWIMA) for PA-TNC", draft-ietf-sacm-nea-swid-patnc-00 (work in progress), January 2017.
- [I-D.ietf-sacm-terminology] Waltermire, D., Montville, A., Harrington, D., and N. Cam-Winget, "Terminology for Security Assessment", draft-ietf-sacm-terminology-05 (work in progress), August 2014.
- [IF-IMC] Trusted Computing Group, "TCG Trusted Network Connect TNC IF-IMC, Version 1.3", February 2013.
- [IF-IMV] Trusted Computing Group, "TCG Trusted Network Connect TNC IF-IMV, Version 1.4", December 2014.
- [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>>.
- [RFC5792] Sangster, P. and K. Narayan, "PA-TNC: A Posture Attribute (PA) Protocol Compatible with Trusted Network Connect (TNC)", RFC 5792, DOI 10.17487/RFC5792, March 2010, <<http://www.rfc-editor.org/info/rfc5792>>.
- [RFC5793] Sahita, R., Hanna, S., Hurst, R., and K. Narayan, "PB-TNC: A Posture Broker (PB) Protocol Compatible with Trusted Network Connect (TNC)", RFC 5793, DOI 10.17487/RFC5793, March 2010, <<http://www.rfc-editor.org/info/rfc5793>>.

- [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, <<http://www.rfc-editor.org/info/rfc6876>>.
- [RFC7632] Waltermire, D. and D. Harrington, "Endpoint Security Posture Assessment: Enterprise Use Cases", RFC 7632, DOI 10.17487/RFC7632, September 2015, <<http://www.rfc-editor.org/info/rfc7632>>.
- [Server-Discovery] Trusted Computing Group, "DRAFT: TCG Trusted Network Connect PDP Discovery and Validation, Version 1.0", October 2015.
- [SWID] "Information technology--Software asset management--Part 2: Software identification tag", ISO/IEC 9899:1999, 2009.

Authors' Addresses

Danny Haynes
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
USA

Email: dhaynes@mitre.org

Jessica Fitzgerald-McKay
Department of Defense
9800 Savage Road
Ft. Meade, Maryland
USA

Email: jmfitz2@nsa.gov

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

Email: llorenzin@pulsesecure.net

SACM
Internet-Draft
Intended status: Standards Track
Expires: October 29, 2017

D. Waltermire, Ed.
NIST
K. Watson
DHS
C. Kahn
L. Lorenzin
Pulse Secure, LLC
M. Cokus
D. Haynes
The MITRE Corporation
H. Birkholz
Fraunhofer SIT
April 27, 2017

SACM Information Model
draft-ietf-sacm-information-model-10

Abstract

This document defines the Information Elements that are transported between SACM components and their interconnected relationships. The primary purpose of the Secure Automation and Continuous Monitoring (SACM) Information Model is to ensure the interoperability of corresponding SACM data models and addresses the use cases defined by SACM. The Information Elements and corresponding types are maintained as the IANA "SACM Information Elements" registry.

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 October 29, 2017.

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 (<http://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	12
2. Conventions used in this document	13
2.1. Requirements Language	13
2.2. Information Element Examples	13
3. Information Elements	13
3.1. Context of Information Elements	14
3.2. Extensibility of Information Elements	14
4. Structure of Information Elements	14
4.1. Information Element Naming Convention	17
4.2. SACM Content Elements	18
4.3. SACM Statements	18
4.4. Relationships	20
4.5. Event	22
4.6. Categories	23
5. Abstract Data Types	23
5.1. Simple Datatypes	23
5.1.1. IPFIX Datatypes	23
5.2. Structured Datatypes	24
5.2.1. List Datatypes	24
5.2.2. Enumeration Datatype	25
5.2.3. Category Datatype	26
6. Information Model Assets	26
6.1. Asset	27
6.2. Endpoint	28
6.3. Hardware Component	28
6.4. Software Component	29
6.4.1. Software Instance	29
6.5. Identity	29
6.6. Guidance	29
6.6.1. Collection Guidance	30
6.6.2. Evaluation Guidance	30

6.6.3.	Classification Guidance	31
6.6.4.	Storage Guidance	31
6.6.5.	Evaluation Results	31
7.	Information Model Elements	32
7.1.	sacmStatement	32
7.2.	sacmStatementMetadata	32
7.3.	sacmContentElement	32
7.4.	sacmContentElementMetadata	33
7.5.	targetEndpoint	33
7.6.	targetEndpointIdentifier	33
7.7.	targetEndpointLabel	33
7.8.	anyIE	34
7.9.	accessPrivilegeType	34
7.10.	accountName	34
7.11.	administrativeDomainType	34
7.12.	addressAssociationType	34
7.13.	addressMaskValue	35
7.14.	addressType	35
7.15.	addressValue	35
7.16.	applicationComponent	35
7.17.	applicationLabel	36
7.18.	applicationType	36
7.19.	applicationManufacturer	36
7.20.	authenticator	36
7.21.	authenticationType	36
7.22.	birthdate	37
7.23.	bytesReceived	37
7.24.	bytesReceived	37
7.25.	bytesSent	37
7.26.	certificate	38
7.27.	collectionTaskType	38
7.28.	confidence	38
7.29.	contentAction	38
7.30.	countryCode	38
7.31.	dataOrigin	39
7.32.	dataSource	39
7.33.	default-depth	39
7.34.	discoverer	39
7.35.	emailAddress	40
7.36.	eventType	40
7.37.	eventThreshold	40
7.38.	eventThresholdName	40
7.39.	eventTrigger	40
7.40.	firmwareId	41
7.41.	hostName	41
7.42.	interfaceLabel	41
7.43.	ipv6AddressSubnetMask	41
7.44.	ipv6AddressSubnetMaskCidrNotation	41

7.45. ipv6AddressValue	42
7.46. ipv4AddressSubnetMask	42
7.47. ipv4AddressSubnetMaskCidrNotation	42
7.48. ipv4AddressValue	42
7.49. layer2InterfaceType	42
7.50. layer4PortAddress	42
7.51. layer4Protocol	43
7.52. locationName	43
7.53. networkZoneLocation	43
7.54. layer2NetworkLocation	43
7.55. layer3NetworkLocation	44
7.56. macAddressValue	44
7.57. methodLabel	44
7.58. methodRepository	44
7.59. networkAccessLevelType	44
7.60. networkId	45
7.61. networkInterfaceName	45
7.62. networkLayer	45
7.63. networkName	45
7.64. organizationId	45
7.65. patchId	46
7.66. patchName	46
7.67. personFirstName	46
7.68. personLastName	46
7.69. personMiddleName	46
7.70. phoneNumber	46
7.71. phoneNumberType	47
7.72. privilegeName	47
7.73. privilegeValue	47
7.74. protocol	47
7.75. publicKey	48
7.76. relationshipContentElementGuid	48
7.77. relationshipStatementElementGuid	48
7.78. relationshipObjectLabel	48
7.79. relationshipType	48
7.80. roleName	49
7.81. sessionStateType	49
7.82. statementGuid	49
7.83. statementType	49
7.84. status	50
7.85. subAdministrativeDomain	50
7.86. subInterfaceLabel	50
7.87. superAdministrativeDomain	50
7.88. superInterfaceLabel	51
7.89. teAssessmentState	51
7.90. teLabel	51
7.91. teId	51
7.92. timestampType	51

7.93. unitsReceived	52
7.94. unitsSent	52
7.95. userDirectory	52
7.96. sacmUserId	52
7.97. webSite	53
7.98. WGS84Longitude	53
7.99. WGS84Latitude	53
7.100. WGS84Altitude	53
7.101. hardwareSerialNumber	53
7.102. interfaceName	54
7.103. interfaceIndex	54
7.104. interfaceMacAddress	54
7.105. interfaceType	54
7.106. interfaceFlags	54
7.107. networkInterface	55
7.108. softwareIdentifier	55
7.109. softwareTitle	55
7.110. softwareCreator	56
7.111. simpleSoftwareVersion	56
7.112. rpmSoftwareVersion	56
7.113. ciscoTrainSoftwareVersion	56
7.114. softwareVersion	56
7.115. softwareLastUpdated	57
7.116. softwareClass	57
7.117. softwareInstance	58
7.118. globallyUniqueIdentifier	59
7.119. creationTimestamp	59
7.120. collectionTimestamp	59
7.121. publicationTimestamp	59
7.122. relayTimestamp	59
7.123. storageTimestamp	60
7.124. type	60
7.125. protocolIdentifier	60
7.126. sourceTransportPort	60
7.127. sourceIPv4PrefixLength	61
7.128. ingressInterface	61
7.129. destinationTransportPort	61
7.130. sourceIPv6PrefixLength	61
7.131. sourceIPv4Prefix	62
7.132. destinationIPv4Prefix	62
7.133. sourceMacAddress	62
7.134. ipVersion	62
7.135. interfaceDescription	62
7.136. applicationDescription	62
7.137. applicationId	63
7.138. applicationName	63
7.139. exporterIPv4Address	63
7.140. exporterIPv6Address	63

7.141. portId	63
7.142. templateId	64
7.143. collectorIPv4Address	64
7.144. collectorIPv6Address	64
7.145. informationElementIndex	65
7.146. informationElementId	65
7.147. informationElementDataType	65
7.148. informationElementDescription	65
7.149. informationElementName	66
7.150. informationElementRangeBegin	66
7.151. informationElementRangeEnd	66
7.152. informationElementSemantics	67
7.153. informationElementUnits	67
7.154. applicationCategoryName	68
7.155. mibObjectValueInteger	68
7.156. mibObjectValueOctetString	69
7.157. mibObjectValueOID	69
7.158. mibObjectValueBits	69
7.159. mibObjectValueIPAddress	70
7.160. mibObjectValueCounter	70
7.161. mibObjectValueGauge	71
7.162. mibObjectValueTimeTicks	71
7.163. mibObjectValueUnsigned	72
7.164. mibObjectValueTable	72
7.165. mibObjectValueRow	72
7.166. mibObjectIdentifier	73
7.167. mibSubIdentifier	73
7.168. mibIndexIndicator	73
7.169. mibCaptureTimeSemantics	74
7.170. mibContextEngineID	75
7.171. mibContextName	76
7.172. mibObjectName	76
7.173. mibObjectDescription	76
7.174. mibObjectSyntax	76
7.175. mibModuleName	76
7.176. interface	77
7.177. iflisteners	77
7.178. physicalProtocol	77
7.179. hwAddress	78
7.180. programName	79
7.181. userId	79
7.182. inetlisteningserver	79
7.183. transportProtocol	79
7.184. localAddress	79
7.185. localPort	80
7.186. localFullAddress	80
7.187. foreignAddress	80
7.188. foreignFullAddress	80

7.189.	selinuxboolean	80
7.190.	selinuxName	81
7.191.	currentStatus	81
7.192.	pendingStatus	81
7.193.	selinuxsecuritycontext	81
7.194.	filepath	82
7.195.	path	82
7.196.	filename	82
7.197.	pid	82
7.198.	role	82
7.199.	domainType	83
7.200.	lowSensitivity	83
7.201.	lowCategory	83
7.202.	highSensitivity	83
7.203.	highCategory	83
7.204.	rawlowSensitivity	84
7.205.	rawlowCategory	84
7.206.	rawhighSensitivity	84
7.207.	rawhighCategory	84
7.208.	systemdunitdependency	84
7.209.	unit	85
7.210.	dependency	85
7.211.	systemdunitproperty	85
7.212.	property	85
7.213.	systemdunitValue	85
7.214.	file	86
7.215.	fileType	86
7.216.	groupId	86
7.217.	aTime	86
7.218.	cTime	86
7.219.	mTime	87
7.220.	size	87
7.221.	suid	87
7.222.	sgid	87
7.223.	sticky	87
7.224.	hasExtendedAcl	88
7.225.	inetd	88
7.226.	serverProgram	88
7.227.	inetdEndpointType	88
7.228.	execAsUser	89
7.229.	waitStatus	89
7.230.	inetAddr	90
7.231.	netmask	90
7.232.	passwordInfo	90
7.233.	username	91
7.234.	password	91
7.235.	gcos	91
7.236.	homeDir	91

7.237. loginShell	91
7.238. lastLogin	92
7.239. process	92
7.240. commandLine	92
7.241. ppid	92
7.242. priority	93
7.243. startTime	93
7.244. routingtable	93
7.245. destination	93
7.246. gateway	93
7.247. runlevelInfo	94
7.248. runlevel	94
7.249. start	94
7.250. kill	94
7.251. shadowItem	94
7.252. chgLst	95
7.253. chgAllow	95
7.254. chgReq	95
7.255. expWarn	95
7.256. expInact	95
7.257. expDate	96
7.258. encryptMethod	96
7.259. symlink	96
7.260. symlinkFilepath	96
7.261. canonicalPath	97
7.262. sysctl	97
7.263. kernelParameterName	97
7.264. kernelParameterValue	97
7.265. uname	98
7.266. machineClass	98
7.267. nodeName	98
7.268. osName	98
7.269. osRelease	98
7.270. processorType	99
7.271. internetService	99
7.272. serviceProtocol	99
7.273. serviceName	99
7.274. flags	99
7.275. noAccess	100
7.276. onlyFrom	100
7.277. port	100
7.278. server	100
7.279. serverArguments	100
7.280. socketType	101
7.281. registeredServiceType	101
7.282. wait	101
7.283. disabled	102
7.284. windowsView	102

7.285.	fileauditedpermissions	102
7.286.	trusteeName	103
7.287.	auditStandardDelete	103
7.288.	auditStandardReadControl	103
7.289.	auditStandardWriteDac	104
7.290.	auditStandardWriteOwner	104
7.291.	auditStandardSynchronize	105
7.292.	auditAccessSystemSecurity	105
7.293.	auditGenericRead	106
7.294.	auditGenericWrite	106
7.295.	auditGenericExecute	107
7.296.	auditGenericAll	107
7.297.	auditFileReadData	108
7.298.	auditFileWriteData	108
7.299.	auditFileAppendData	109
7.300.	auditFileReadEa	109
7.301.	auditFileWriteEa	110
7.302.	auditFileExecute	110
7.303.	auditFileDeleteChild	111
7.304.	auditFileReadAttributes	111
7.305.	auditFileWriteAttributes	112
7.306.	fileeffectiverights	112
7.307.	standardDelete	113
7.308.	standardReadControl	113
7.309.	standardWriteDac	113
7.310.	standardWriteOwner	114
7.311.	standardSynchronize	114
7.312.	accessSystemSecurity	114
7.313.	genericRead	114
7.314.	genericWrite	114
7.315.	genericExecute	115
7.316.	genericAll	115
7.317.	fileReadData	115
7.318.	fileWriteData	115
7.319.	fileAppendData	115
7.320.	fileReadEa	116
7.321.	fileWriteEa	116
7.322.	fileExecute	116
7.323.	fileDeleteChild	116
7.324.	fileReadAttributes	116
7.325.	fileWriteAttributes	117
7.326.	groupInfo	117
7.327.	group	117
7.328.	subgroup	117
7.329.	groupSidInfo	117
7.330.	userSidInfo	118
7.331.	userSid	118
7.332.	subgroupSid	118

7.333. lockoutpolicy	118
7.334. forceLogoff	118
7.335. lockoutDuration	119
7.336. lockoutObservationWindow	119
7.337. lockoutThreshold	119
7.338. passwordpolicy	119
7.339. maxPasswdAge	120
7.340. minPasswdAge	120
7.341. minPasswdLen	120
7.342. passwordHistLen	121
7.343. passwordComplexity	121
7.344. reversibleEncryption	121
7.345. portInfo	121
7.346. foreignPort	121
7.347. printereffectiverights	122
7.348. printerName	122
7.349. printerAccessAdminister	122
7.350. printerAccessUse	122
7.351. jobAccessAdminister	122
7.352. jobAccessRead	123
7.353. registry	123
7.354. registryHive	123
7.355. registryKey	124
7.356. registryKeyName	124
7.357. lastWriteTime	124
7.358. registryKeyType	125
7.359. registryKeyValue	126
7.360. regkeyauditedpermissions	127
7.361. auditKeyQueryValue	128
7.362. auditKeySetValue	128
7.363. auditKeyCreateSubKey	129
7.364. auditKeyEnumerateSubKeys	129
7.365. auditKeyNotify	130
7.366. auditKeyCreateLink	130
7.367. auditKeyWow6464Key	131
7.368. auditKeyWow6432Key	131
7.369. auditKeyWow64Res	132
7.370. regkeyeffectiverights	132
7.371. keyQueryValue	133
7.372. keySetValue	133
7.373. keyCreateSubKey	133
7.374. keyEnumerateSubKeys	134
7.375. keyNotify	134
7.376. keyCreateLink	134
7.377. keyWow6464Key	134
7.378. keyWow6432Key	134
7.379. keyWow64Res	134
7.380. service	135

7.381. displayName	135
7.382. description	135
7.383. serviceType	135
7.384. startType	136
7.385. currentState	137
7.386. controlsAccepted	138
7.387. startName	140
7.388. serviceFlag	140
7.389. dependencies	140
7.390. serviceeffectiverights	140
7.391. trusteeSid	141
7.392. serviceQueryConf	141
7.393. serviceChangeConf	141
7.394. serviceQueryStat	141
7.395. serviceEnumDependents	141
7.396. serviceStart	142
7.397. serviceStop	142
7.398. servicePause	142
7.399. serviceInterrogate	142
7.400. serviceUserDefined	142
7.401. sharedresourceauditedpermissions	143
7.402. netname	143
7.403. sharedresourceeffectiverights	143
7.404. user	144
7.405. enabled	144
7.406. lastLogon	144
7.407. groupSid	144
7.408. endpointType	144
7.409. endpointPurpose	145
7.410. endpointCriticality	145
7.411. ingestTimestamp	145
7.412. vulnerabilityVersion	146
7.413. vulnerabilityExternalId	146
7.414. vulnerabilitySeverity	146
7.415. assessmentTimestamp	146
7.416. vulnerableSoftware	146
7.417. endpointVulnerabilityStatus	147
7.418. vulnerabilityDescription	147
8. Acknowledgements	147
9. IANA Considerations	148
10. Security Considerations	148
11. Operational Considerations	149
11.1. Endpoint Designation	149
11.2. Timestamp Accuracy	150
12. Privacy Considerations	151
13. References	151
13.1. Normative References	151
13.2. Informative References	151

Appendix A. Change Log	152
A.1. Changes in Revision 01	152
A.2. Changes in Revision 02	154
A.3. Changes in Revision 03	154
A.4. Changes in Revision 04	154
A.5. Changes in Revision 05	155
A.6. Changes in Revision 06	155
A.7. Changes in Revision 07	155
A.8. Changes in Revision 08	156
A.9. Changes in Revision 09	156
A.10. Changes in Revision 10	157
Authors' Addresses	157

1. Introduction

The SACM Information Model (IM) serves multiple purposes:

- o to ensure interoperability between SACM data models that are used as transport encodings,
- o to provide a standardized set of Information Elements - the SACM Vocabulary - to enable the exchange of content vital to automated security posture assessment, and
- o to enable secure information sharing in a scalable and extensible fashion in order to support the tasks conducted by SACM components.

A complete set of requirements imposed on the IM can be found in [I-D.ietf-sacm-requirements]. The SACM IM is intended to be used for standardized data exchange between SACM components (data in motion). Nevertheless, the Information Elements (IE) and their relationships defined in this document can be leveraged to create and align corresponding data models for data at rest.

The information model expresses, for example, target endpoint (TE) attributes, guidance, and evaluation results. The corresponding Information Elements are consumed and produced by SACM components as they carry out tasks.

The primary tasks that this information model supports (on data, control, and management plane) are:

- o TE Discovery
- o TE Characterization
- o TE Classification

- o Collection
- o Evaluation
- o Information Sharing
- o SACM Component Discovery
- o SACM Component Authentication
- o SACM Component Authorization
- o SACM Component Registration

These tasks are defined in [I-D.ietf-sacm-terminology].

2. Conventions used in this document

2.1. Requirements Language

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

2.2. Information Element Examples

The notation used to define the SACM Information Elements (IEs) is based on a customized version of the IPFIX information model syntax [RFC7012] which is described in Figure 2. However, there are several examples presented throughout the document that use a simplified pseudo-code to illustrate the basic structure. It should be noted that while they include actual names of subjects and attributes as well as values, they are not intended to influence how corresponding SACM IEs should be defined in Section 7. The examples are provided for demonstration purposes only.

3. Information Elements

The IEs defined in this document comprise the building blocks by which all SACM content is composed. They are consumed and provided by SACM components on the data plane. Every Information Element has a unique label: its name. Every type of IE defined by the SACM IM is registered as a type at the IANA registry. The Integer Index of the IANA SMI number tables can be used by SACM data models.

3.1. Context of Information Elements

The IEs in this information model represent information related to assets in the following areas (based on the use cases described in [RFC7632]):

- o Endpoint Management
- o Software Inventory Management
- o Hardware Inventory Management
- o Configuration Management
- o Vulnerability Management

3.2. Extensibility of Information Elements

A SACM data model based on this information model MAY include additional information elements that are not defined here. The labels of additional Information Elements included in different SACM data models MUST NOT conflict with the labels of the Information Elements defined by this information model, and the names of additional Information Elements MUST NOT conflict with each other or across multiple data models. In order to avoid naming conflicts, the labels of additional IEs SHOULD be prefixed to avoid collisions across extensions. The prefix MUST include an organizational identifier and therefore, for example, MAY be an IANA enterprise number, a (partial) name space URI, or an organization name abbreviation.

4. Structure of Information Elements

There are two basic types of IEs:

- o Attributes: Atomic information elements that are equivalent to name-value-pairs and can be components of Subjects.
- o Subjects: Composite information elements that have a name and are made up of Attributes and/or other Subjects. Every IE that is part of a Subject can have a quantity associated with it (e.g. zero-one, none-unbounded). The content IEs of a Subject can be ordered or unordered.

```
Example Instance of an Attribute:  
hostname = "arbutus"
```

```
Example Instance of a Subject:  
coordinates = (  
  latitude = N27.99619,  
  longitude = E86.92761  
)
```

Figure 1: Example instance of an attribute and subject.

In general, every piece of information that enables security posture assessment or further enriches the quality of the assessment process can be associated with metadata. In the SACM IM, metadata is represented by specific subjects and is bundled with other attributes or subjects to provide additional information about them. The IM explicitly defines two kinds of metadata:

- o Metadata focusing on the data origin (the SACM component that provides the information to the SACM domain)
- o Metadata focusing on the data source (the target endpoint that is assessed)

Metadata can also include relationships that refer to other associated IEs (or SACM content in general) by using referencing labels that have to be included in the metadata of the associated IE.

Subjects can be nested and the SACM IM allows for circular or recursive nesting. The association of IEs via nesting results in a tree-like structure wherein subjects compose the root and intermediary nodes and attributes the leaves of the tree. This semantic structure does not impose a specific structure on SACM data models regarding data in motion or data repository schemata for data at rest.

The SACM IM provides two conceptual top-level subjects that are used to ensure a homogeneous structure for SACM content and its associated metadata: SACM statements and SACM content-elements. Every set of IEs that is provided by a SACM component must provide the information contained in these two subjects although it is up to the implementer whether or not the subjects are explicitly defined in a data model.

The notation the SACM IM is defined in is based on a modified version of the IP Information Flow Export (IPFIX) Information Model syntax described in Section 2.1 of [RFC7012]. The customized syntax used by the SACM IM is defined below in Figure 2.

- `elementId` (required): The numeric identifier of the Information Element. It is used for the compact identification of an Information Element. If this identifier is used without an `enterpriseId`, then the `elementId` must be unique, and the description of allowed values is administrated by IANA. The value "TBD" may be used during development of the information model until an `elementId` is assigned by IANA and filled in at publication time.
- `enterpriseId` (optional): Enterprises may wish to define Information Elements without registering them with IANA, for example, for enterprise-internal purposes. For such Information Elements, the `elementId` is not sufficient when used outside the enterprise. If specifications of enterprise-specific Information Elements are made public and/or if enterprise-specific identifiers are used by SACM components outside the enterprise, then the enterprise-specific identifier MUST be made globally unique by combining it with an enterprise identifier. Valid values for the `enterpriseId` are defined by IANA as Structure of Management Information (SMI) network management private enterprise numbers.
- `name` (required): A unique and meaningful name for the Information Element.
- `dataType` (required): There are two kinds of datatypes: simple and structured. Attributes are defined using simple datatypes and subjects are defined using structured datatypes. The contents of the datatype field will be either a reference to one of the simple

	datatypes listed in Section 5.1, or the specification of structured datatype as defined in Section 5.2.
status (required):	The status of the specification of the Information Element. Allowed values are "current" and "deprecated". All newly defined Information Elements have "current" status. The process for moving Information Elements to the "deprecated" status is TBD.
description (required):	Describes the meaning of the Information Element, how it is derived, conditions for its use, etc.
structure (optional):	A parsable property that provides details about the definition of structured Information Elements as described in Section 5.2.
references (optional):	Identifies other RFCs or documents outside the IETF which provide additional information or context about the Information Element.

Figure 2: Information Element Specification Template

4.1. Information Element Naming Convention

SACM Information Elements must adhere to the following naming conventions.

- o Names SHOULD be descriptive
- o Names MUST be unique within the SACM registry. Enterprise-specific names SHOULD be prefixed with a Private Enterprise Number [PEN].
- o Names MUST start with lowercase letters unless it begins with a Private Enterprise Number
- o Composed names MUST use capital letters for the first letter of each part

4.2. SACM Content Elements

Every piece of information that is provided by a SACM Component is always associated with a set of data source metadata (e.g. the timestamp when the information was collected, the target endpoint from which the this set of information is about, etc.) which is provided in the SACM Content Element Metadata. The SACM Content Element is the subject information element that associates the information with the SACM Content Element Metadata. The SACM Content Element Metadata may also include relationships that express associations with other SACM Content Elements.

```
content-element = (  
  content-metadata = (  
    collection-timestamp = 146193322,  
    data-source = fb02e551-7101-4e68-8dec-1fde6bd10981  
  ),  
  hostname = "arbutus",  
  coordinates = (  
    latitude = N27.99619,  
    longitude = E86.92761  
  )  
)
```

Figure 3: Example set of IEs associated with a timestamp and a target endpoint label.

4.3. SACM Statements

One or more SACM Content Elements are bundled in a SACM Statement. In contrast to SACM Content Element Metadata, SACM Statement Metadata focuses on the providing information about the SACM Component that provided it rather than the target endpoint that the content is about. The only content-specific metadata included in the SACM Statement is the statement-type IE. Therefore, multiple SACM Content Elements that share the same SACM Statement Metadata and are of the same statement-type can be included in a single SACM Statement. A SACM Statement functions similar to an envelope or a header and is the subject information element that associates SACM Statement Metadata with security automation information provided in its SACM Content Element(s). Its purpose is to enable the tracking of the origin of data inside a SACM domain and more importantly to enable the mitigation of conflicting information that may originate from different SACM Components. How a consuming SACM Component actually deals with conflicting information is out-of-scope of the SACM IM. Semantically, the term statement implies that the SACM content provided by a SACM Component might not be correct in every context,

but, rather is the result of a best-effort to produce correct information.

```
sacm-statement = (  
  statement-metadata = (  
    publish-timestamp = 1461934031,  
    data-origin = 24e67957-3d31-4878-8892-da2b35e121c2,  
    statement-type = observation  
  ),  
  content-element = (  
    content-metadata = (  
      collection-timestamp = 146193322,  
      data-source = fb02e551-7101-4e68-8dec-1fde6bd10981  
    ),  
    hostname = "arbutus"  
  )  
)
```

Figure 4: Example of a simple SACM statement including a single content-element.

```
sacm-statement = (  
  statement-metadata = (  
    publish-timestamp = 1461934031,  
    data-origin = 24e67957-3d31-4878-8892-da2b35e121c2  
    statement-type = observation  
  ),  
  content-element = (  
    content-metadata = (  
      collection-timestamp = 146193322,  
      data-source = fb02e551-7101-4e68-8dec-1fde6bd10981  
    ),  
    coordinates = (  
      latitude = N27.99619,  
      longitude = E86.92761  
    )  
  )  
)  
  
sacm-statement = (  
  statement-metadata = (  
    publish-timestamp = 1461934744,  
    data-origin = e42885a1-0270-44e9-bb5c-865cf6bd4800,  
    statement-type = observation  
  ),  
  content-element = (  
    content-metadata = (  
      collection-timestamp = 146193821,  
      te-label = fb02e551-7101-4e68-8dec-1fde6bd10981  
    ),  
    coordinates = (  
      latitude = N16.67622,  
      longitude = E141.55321  
    )  
  )  
)
```

Figure 5: Example of conflicting information originating from different SACM components.

4.4. Relationships

An IE can be associated with another IE, e.g. a user-name attribute can be associated with a content-authorization subject. These references are expressed via the relationships subject, which can be included in a corresponding content-metadata subject. The relationships subject includes a list of one or more references. The SACM IM does not enforce a SACM domain to use unique identifiers as

references. Therefore, there are at least two ways to reference another

- o The value of a reference represents a specific content-label that is unique in a SACM domain (and has to be included in the corresponding content-element metadata in order to be referenced), or
- o The reference is a subject that includes an appropriate number of IEs in order to identify the referenced content-element by its actual content.

It is recommended to provide unique identifiers in a SACM domain and the SACM IM provides a corresponding naming-convention as a reference in Section 4.1. The alternative highlighted above summarizes a valid approach that does not require unique identifiers and is similar to the approach of referencing target endpoints via identifying attributes included in a characterization record.

```
content-element = (  
  content-metadata = (  
    collection-timestamp = 1461934031,  
    te-label =  
    fb02e551-7101-4e68-8dec-1fde6bd10981  
    relationships = (  
      associated-with-user-account =  
      f3d70ef4-7e18-42af-a894-8955ba87c95d  
    )  
  ),  
  hostname = "arbutus"  
)  
  
content-element = (  
  content-metadata = (  
    content-label = f3d70ef4-7e18-42af-a894-8955ba87c95d  
  ),  
  user-account = (  
    username = romeo  
    authentication = local  
  )  
)
```

Figure 6: Example instance of a content-element subject associated with another subject via its content metadata.

4.5. Event

Event subjects provide a structure to represent the change of IE values that was detected by a collection task at a specific point of time. It is mandatory to include the new values and the collection timestamp in an event subject and it is recommended to include the past values and a collection timestamp that were replaced by the new IE values. Every event can also be associated with a subject-specific event-timestamp and a lastseen-timestamp that might differ from the corresponding collection-timestamps. If these are omitted the collection-timestamp that is included in the content-metadata subject is used instead.

```
sacm-statement = (  
  statement-metadata = (  
    publish-timestamp = 1461934031,  
    data-origin = 24e67957-3d31-4878-8892-da2b35e121c2,  
    statement-type = event  
  ),  
  event = (  
    event-attributes = (  
      event-name = "host-name change",  
      content-element = (  
        content-metadata = (  
          collection-timestamp = 146193322,  
          data-source =  
            fb02e551-7101-4e68-8dec-1fde6bd10981,  
          event-component = past-state  
        ),  
        hostname = "arbutus"  
      ),  
      content-element = (  
        content-metadata = (  
          collection-timestamp = 146195723,  
          data-source =  
            fb02e551-7101-4e68-8dec-1fde6bd10981,  
          event-component = current-state  
        ),  
        hostname = "lilac"  
      )  
    )  
  )  
)
```

Figure 7: Example of a SACM statement containing an event.

4.6. Categories

Categories are special IEs that refer to multiple types of IEs via just one name. Therefore, they are similar to a type-choice. A prominent example of a category is when identifying a target endpoint. In some cases, a target endpoint will be identified by a set of identifying attributes and in other cases a target endpoint will be identified by a target endpoint label which is unique within a SACM domain. If a subject includes the targetEndpoint information element as one of its components, any of the category members (targetEndpointIdentifier or targetEndpointLabel) are valid to be used in its place.

5. Abstract Data Types

This section describes the set of valid abstract data types that can be used for the specification of the SACM Information Elements in Section 7. SACM currently supports two classes of datatypes that can be used to define Information Elements.

- o Simple: Datatypes that are atomic and are used to define the type of data represented by an attribute Information Element.
- o Structured: Datatypes that can be used to define the type of data represented by a subject Information Element.

Note that further abstract data types may be specified by future extensions of the SACM information model.

5.1. Simple Datatypes

5.1.1. IPFIX Datatypes

To facilitate the use of existing work, SACM supports the following abstract data types defined in Section 3 of [RFC7012].

- o unsigned8, unsigned16, unsigned32, unsigned64
- o signed8, signed16, signed32, signed64
- o float32, float64
- o boolean
- o macAddress
- o octetArray

- o string
- o dateTimeSeconds, dateTimeMilliseconds, dateTimeMicroseconds, dateTimeNanoSeconds
- o ipv4Address, ipv6Address

5.2. Structured Datatypes

5.2.1. List Datatypes

SACM defines the following abstract list data types that are used to represent the structured data associated with subjects.

- o list: indicates that the Information Element order is not significant but MAY be preserved.
- o orderedList: indicates that Information Element order is significant and MUST be preserved.

The notation for defining a SACM structured datatype is based on regular expressions, which are composed of the keywords "list" or "orderedList" and an Information Element expression. IE expressions use some of the regular expression syntax and operators, but the terms in the expression are the names of defined Information Elements instead of character classes. The syntax for defining list and orderedList datatypes is described below, using BNF:

```

<list-def> -> ("list"|"orderedList") "(" <ie-expression> ")"

<ie-expression> -> <ie-name> <cardinality>?
                  ( ("," | "|") <ie-name> <cardinality>?)*

<cardinality> -> "*" | "+" | "?" |
                  ( "(" <non-neg-int> ("," <non-neg-int>)? ")" )

```

Figure 8: Syntax for Defining List Datatypes

As seen above, multiple occurrences of an Information Element may be present in a structured datatype. The cardinality of an Information Element within a structured Information Element definition is defined by the following operators:

- * - zero or more occurrences
- + - one or more occurrences
- ? - zero or one occurrence
- (m,n) - between m and n occurrences

Figure 9: Specifying Cardinality for Structured Datatypes

The absence of a cardinality operator implies one mandatory occurrence of the Information Element.

Below is an example of a structured Information Element definition.

```
personInfo = list(firstName, middleNames?, lastName)
firstName = string
middleNames = orderedList(middleName+)
middleName = string
lastName = string
```

As an example, consider the name "John Ronald Reuel Tolkien". Below are instances of this name, structured according to the personInfo definition.

```
personInfo = (firstName="John", middleNames(middleName="Ronald",
middleName="Reuel"), lastName="Tolkien")

personInfo = (middleNames(middleName="Ronald", middleName=" Reuel"),
lastName="Tolkien", firstName="John")
```

The instance below is not legal with respect to the definition of personInfo because the order in middleNames is not preserved.

```
personInfo = (firstName="John", middleNames(middleName=" Reuel",
middleName="Ronald"), lastName="Tolkien")
```

Figure 10: Example of Defining a Structured List Datatype

5.2.2. Enumeration Datatype

SACM defines the following abstract enumeration datatype that is used to represent the restriction of an attribute value to a set of values.


```

name, hex-value, description
<enumeration-def> -> -> <name> ";" <hex-value> ";" <description>
<name> -> [0-9a-zA-Z]+
<hex-value> -> 0x[0-9a-fA-F]+
<description> -> [0-9a-zA-Z\.\,]+

```

Figure 11: Syntax for Defining an Enumeration Datatype

Below is an example of a structured Information Element definition for an enumeration.

```

Red      ; 0x1  ; The color is red.
Orange   ; 0x2  ; The color is orange.
Yellow   ; 0x3  ; The color is yellow.
Green    ; 0x4  ; The color is green.
...

```

Figure 12: Example of Defining a Structured Enumeration Datatype

5.2.3. Category Datatype

SACM defines the following abstract category datatype that is used to represent a type-choice between a set of information elements.

```

<category-def> -> "category(" <ie-expression> ")"
<ie-expression> -> <ie-name> ("|" <ie-name>)*
<name> -> [0-9a-zA-Z]+

```

Figure 13: Syntax for Defining an Category Datatype

Below is an example of a structured Information Element definition for a category.

```

targetEndpoint = category(targetEndpointIdentifier |
                           targetEndpointLabel)

```

Figure 14: Example of Defining a Structured Category Datatype

6. Information Model Assets

In order to represent the Information Elements related to the areas listed in Section 3.1, the information model defines the information needs (or metadata about those information needs) related to following types of assets which are defined in [I-D.ietf-sacm-terminology] (and included below for convenience) which are of interest to SACM. Specifically:

- o Endpoint

- o Software Component
- o Hardware Component
- o Identity
- o Guidance
- o Evaluation Results

The following figure shows the make up of an Endpoint asset which contains zero or more hardware components and zero or more software components each of which may have zero or more instances running an endpoint at any given time as well as zero or more identities that act on behalf of the endpoint when interfacing with other endpoints, tools, or services. An endpoint may also contain other endpoints in the case of a virtualized environment.




Figure 15: Model of an Endpoint

6.1. Asset

As defined in [RFC4949], an asset is a system resource that is (a) required to be protected by an information system's security policy,

(b) intended to be protected by a countermeasure, or (c) required for a system's mission.

In the scope of SACM, an asset can be composed of other assets. Examples of Assets include: Endpoints, Software, Guidance, or Identity. Furthermore, an asset is not necessarily owned by an organization.

6.2. Endpoint

From [RFC5209], an endpoint is any computing device that can be connected to a network. Such devices normally are associated with a particular link layer address before joining the network and potentially an IP address once on the network. This includes: laptops, desktops, servers, cell phones, or any device that may have an IP address.

To further clarify, an endpoint is any physical or virtual device that may have a network address. Note that, network infrastructure devices (e.g. switches, routers, firewalls), which fit the definition, are also considered to be endpoints within this document.

Physical endpoints are always composites that are composed of hardware components and software components. Virtual endpoints are composed entirely of software components and rely on software components that provide functions equivalent to hardware components.

The SACM architecture differentiates two essential categories of endpoints: Endpoints whose security posture is intended to be assessed (target endpoints) and endpoints that are specifically excluded from endpoint posture assessment (excluded endpoints).

6.3. Hardware Component

Hardware components are the distinguishable physical components that compose an endpoint. The composition of an endpoint can be changed over time by adding or removing hardware components. In essence, every physical endpoint is potentially a composite of multiple hardware components, typically resulting in a hierarchical composition of hardware components. The composition of hardware components is based on interconnects provided by specific hardware types (e.g. mainboard is a hardware type that provides local busses as an interconnect). In general, a hardware component can be distinguished by its serial number.

Examples of a hardware components include: motherboards, network interfaces, graphics cards, hard drives, etc.

6.4. Software Component

A software package installed on an endpoint (including the operating system) as well as a unique serial number if present (e.g. a text editor associated with a unique license key).

It should be noted that this includes both benign and harmful software packages. Examples of benign software components include: applications, patches, operating system kernel, boot loader, firmware, code embedded on a webpage, etc. Examples of malicious software components include: malware, trojans, viruses, etc.

6.4.1. Software Instance

A running instance of the software component (e.g. on a multi-user system, one logged-in user has one instance of a text editor running and another logged-in user has another instance of the same text editor running, or on a single-user system, a user could have multiple independent instances of the same text editor running).

6.5. Identity

Any mechanism that can be used to identify an asset during an authentication process. Examples include usernames, user and device certificates, etc. Note, that this is different than the identity of assets in the context of designation as described in Section 11.1.

6.6. Guidance

Guidance is input instructions to processes and tasks, such as collection or evaluation. Guidance influences the behavior of a SACM component and is considered content of the management plane. Guidance can be manually or automatically generated or provided. Typically, the tasks that provide guidance to SACM components have a low-frequency and tend to be sporadic. A prominent example of guidance are target endpoint profiles, but guidance can have many forms, including:

Configuration, e.g. a SACM component's name, or a CMDB's IPv6 address.

Profiles, e.g. a set of expected states for network behavior associated with target endpoints employed by specific users.

Policies, e.g. an interval to refresh the registration of a SACM component, or a list of required capabilities for SACM components in a specific location.

6.6.1. Collection Guidance

A collector may need guidance to govern what it collects and when. Collection Guidance provides instructions for a Collector that specifies which endpoint attributes to collect, when to collect them, and how to collect them. Collection Guidance is composed of Target Endpoint Attribute Guidance, Frequency Guidance, and Method Guidance.

- o Target Endpoint Attribute Guidance: Set of endpoint attributes that are supposed to be collected from a target endpoint. The definition of the set of endpoint attributes is typically based on an endpoint characterization record.
- o Frequency Guidance: Specifies when endpoint attributes are to be collected.
- o Method Guidance: Indicates how endpoint attributes are to be collected.

6.6.2. Evaluation Guidance

An evaluator typically needs guidance to govern what it considers to be a good or bad security posture. Evaluation Guidance provides instructions for an Evaluator that specifies which endpoint attributes to evaluate, the desired state of those endpoint attributes, and any special requirements that enable an Evaluator to determine if the endpoint attributes can be used in the evaluation (e.g. freshness of data, how it was collected, etc.). Evaluation Guidance is composed of Target Endpoint Attribute Guidance, Expected Endpoint Attribute Value Guidance, and Frequency Guidance.

- o Target Endpoint Attribute Guidance: Set of target endpoint attributes that are supposed to be used in an evaluation as well as any requirements on the endpoint attributes. The definition of the set of endpoint attributes is typically based on an endpoint characterization record.
- o Expected Endpoint Attribute Value Guidance: The expected values of the endpoint attributes described in the Target Endpoint Attribute Guidance.
- o Frequency Guidance: Specifies when endpoint attributes are to be evaluated.
- o Method Guidance: Indicates how endpoint attributes are to be collected.

6.6.3. Classification Guidance

A SACM Component carrying out the Target Endpoint Classification Task may need guidance on how to classify an endpoint. Specifically, how to associate endpoint classes with a specific target endpoint characterization record. Target Endpoint Classes function as guidance for collection, evaluation, remediation and security posture assessment in general. Classification Guidance is composed of Target Endpoint Attribute Guidance and Class Guidance.

- o Target Endpoint Attribute Guidance: Set of target endpoint attributes that are supposed to be used to identify the endpoint characterization record.
- o Class Guidance: A list of target endpoint classes that are to be associated with the identified target endpoint characterization record.

6.6.4. Storage Guidance

An SACM Component typically needs guidance to govern what information it should store and where. Storage Guidance provides instructions for a SACM Component that specifies which security automation information should be stored, for how long, and on which endpoint. Storage Guidance is composed of Target Endpoint Attribute Guidance, Expected Security Automation Information Guidance, and Retention Guidance.

- o Target Endpoint Attribute Guidance: Set of target endpoint attributes that are supposed to be used to identify the endpoint where the security automation information is to be stored.
- o Expected Security Automation Information Guidance: The security automation information that is expected to be stored (guidance, collected posture attributes, results, etc.).
- o Retention Guidance: Specifies how long the security automation information should be stored.

6.6.5. Evaluation Results

Evaluation Results are the output of comparing the actual state of an endpoint against the expected state of an endpoint. In addition to the actual results of the comparison, Evaluation Results should include the Evaluation Guidance and actual target endpoint attributes values used to perform the evaluation.

7. Information Model Elements

This section defines the specific Information Elements and relationships that will be implemented by data models and transported between SACM Components.

7.1. sacmStatement

```
elementId: TBD
name: sacmStatement
dataType: orderedList
status: current
description: Associates SACM Statement Metadata
             which provides data origin information about
             the providing SACM Component with one or more
             SACM Content Elements that contain security
             automation information.
structure: orderedList(sacmStatementMetadata,
                       sacmContentElement+)
```

7.2. sacmStatementMetadata

```
elementId: TBD
name: sacmStatementMetadata
dataType: orderedList
status: current
description: Contains IEs that provide
             information about the data origin of the
             providing SACM Component as well as the
             information necessary for other SACM
             Components to understand the type of
             security automation information in the
             SACM Statement's SACM Content Element(s).
structure: orderedList(publicationTimestamp,
                       dataOrigin, anyIE*)
```

7.3. sacmContentElement

```
elementId: TBD
name: sacmContentElement
dataType: list
status: current
description: Associates SACM Content Element
             Metadata which provides information about the
             data source and type of security automation
             information with the actual security automation
             information.
structure: TODO
```

7.4. sacmContentElementMetadata

```
elementId: TBD
name: sacmContentElementMetadata
dataType: orderedList
status: current
description: Contains IEs that provide
information about the data source and type of
security automation information such that other
SACM Components are able to parse and understand
the security automation information contained
within the SACM Statement's SACM Content Element(s).
structure: orderedList(collectionTimestamp,
                        targetEndpoint, anyIE*)
```

7.5. targetEndpoint

```
elementId: TBD
name: targetEndpoint
dataType: category
status: current
description: Information that identifies a target
endpoint on the network. This may be a set of
attributes that can be used to identify an endpoint
on the network or a label that is unique to a SACM
domain.
structure: category(targetEndpointIdentifier |
                    targetEndpointLabel)
```

7.6. targetEndpointIdentifier

```
elementId: TBD
name: targetEndpointIdentifier
dataType: list
status: current
description: A set of attributes that uniquely
identify a target endpoint on the network.
structure: list(anyIE+)
```

7.7. targetEndpointLabel

```
elementId: TBD
name: targetEndpointLabel
dataType: string
status: current
description: A label that uniquely identifies
a target endpoint on SACM domain.
```


7.8. anyIE

elementId: TBD
name: anyIE
dataType: category
status: current
description: This category is a placeholder
for any information element defined within
the SACM Information Model. Its purpose is
to provide an extension point in other
information elements that enable them to
support the specific needs of an enterprise,
user, product, or service.

7.9. accessPrivilegeType

elementId: TBD
name: accessPrivilegeType
dataType: string
status: current
description: A set of types that represent access
privileges (read, write, none, etc.).

7.10. accountName

elementId: TBD
name: accountName
dataType: string
status: current
description: A label that uniquely identifies an account
that can require some form of (user) authentication to
access.

7.11. administrativeDomainType

elementId: TBD
name: administrativeDomainType
dataType: string
status: current
description: A label the is supposed to uniquely
identify an administrative domain.

7.12. addressAssociationType

elementId: TBD
name: addressAssociationType
dataType: string
status: current
description: A label the is supposed to uniquely
identify an administrative domain.

7.13. addressMaskValue

elementId: TBD
name: addressMaskValue
dataType: string
status: current
description: A value that expresses a generic address
subnetting bitmask.

7.14. addressType

elementId: TBD
name: addressType
dataType: string
status: current
description: A set of types that specifies the type
of address that is expressed in an address subject
(e.g. ethernet, modbus, zigbee).

7.15. addressValue

elementId: TBD
name: addressValue
dataType: string
status: current
description: A value that expresses a generic network
address.

7.16. applicationComponent

elementId: TBD
name: applicationComponent
dataType: string
status: current
description: A label that references a "sub"-application
that is part of the application (e.g. an add-on, a
cipher-suite, a library).

7.17. applicationLabel

elementId: TBD
name: applicationLabel
dataType: string
status: current
description: A label that is supposed to uniquely reference an application.

7.18. applicationType

elementId: TBD
name: applicationType
dataType: string
status: current
description: A set of types (FIXME maybe a finite set is not realistic here - value not enumerator?) that identifies the type of (user-space) application (e.g. text-editor, policy-editor, service-client, service-server, calendar, rouge-like RPG).

7.19. applicationManufacturer

elementId: TBD
name: applicationManufacturer
dataType: string
status: current
description: The name of the vendor that created the application.

7.20. authenticator

elementId: TBD
name: authenticator
dataType: string
status: current
description: A label that references a SACM component that can authenticate target endpoints (can be used in a target-endpoint subject to express that the target endpoint was authenticated by that SACM component).

7.21. authenticationType

elementId: TBD
name: authenticationType
dataType: string
status: current
description: A set of types that express which type of authentication was used to enable a network interaction/connection.

7.22. birthdate

elementId: TBD
name: birthdate
dataType: string
status: current
description: A label for the registered day of birth of a natural person (e.g. the date of birth of a person as an ISO date string).
references: <http://rs.tdwg.org/ontology/voc/Person#birthdate>

7.23. bytesReceived

elementId: TBD
name: bytesReceived
dataType: string
status: current
description: A value that represents a number of octets received on a network interface.

7.24. bytesReceived

elementId: TBD
name: bytesReceived
dataType: string
status: current
description: A value that represents the number of octets received on a network interface.

7.25. bytesSent

elementId: TBD
name: bytesSent
dataType: string
status: current
description: A value that represents the number of octets sent on a network interface.

7.26. certificate

elementId: TBD
name: certificate
dataType: string
status: current
description: A value that expresses a certificate that can be collected from a target endpoint.

7.27. collectionTaskType

elementId: TBD
name: collectionTaskType
dataType: string
status: current
description: A set of types that defines how collected SACM content was acquired (e.g. network-observation, remote-acquisition, self-reported, derived, authority, verified).

7.28. confidence

elementId: TBD
name: confidence
dataType: string
status: current
description: A representation of the subjective probability that the assessed value is correct. If no confidence value is given, it is assumed that the confidence is 1. Acceptable values are between 0 and 1.

7.29. contentAction

elementId: TBD
name: contentAction
dataType: string
status: current
description: A set of types that express a type of action (e.g. add, delete, update). It can be associated, for instance, with an event subject or with a network observation.

7.30. countryCode

elementId: TBD
name: countryCode
dataType: string
status: current
description: A set of types according to ISO 3166-1.

7.31. dataOrigin

elementId: TBD
name: dataOrigin
dataType: string
status: current
description: A label that uniquely identifies a SACM component in and across SACM domains.

7.32. dataSource

elementId: TBD
name: dataSource
dataType: string
status: current
description: A label that is supposed to uniquely identify the data source (e.g. a target endpoint or sensor) that provided an initial endpoint attribute record.

7.33. default-depth

elementId: TBD
name: default-depth
dataType: string
status: current
description: A value that expresses how often a circular reference of subject is allowed to repeat, or how deep a recursive nesting may occur, respectively.

7.34. discoverer

elementId: TBD
name: discoverer
dataType: string
status: current
description: A label that refers to the SACM component that discovered a target endpoint (can be used in a target-endpoint subject to express, for example, that the target endpoint was authenticated by that SACM component).

7.35. emailAddress

elementId: TBD
name: emailAddress
dataType: string
status: current
description: A value that expresses an email-address.

7.36. eventType

elementId: TBD
name: eventType
dataType: string
status: current
description: a set of types that define the categories of an event (e.g. access-level-change, change-of-privilege, change-of-authorization, environmental-event, or provisioning-event).

7.37. eventThreshold

elementId: TBD
name: eventThreshold
dataType: string
status: current
description: If applicable, a value that can be included in an event subject to indicate what numeric threshold value was crossed to trigger that event.

7.38. eventThresholdName

elementId: TBD
name: eventThresholdName
dataType: string
status: current
description: If an event is created due to a crossed threshold, the threshold might have a name associated with it that can be expressed via this value.

7.39. eventTrigger

elementId: TBD
name: eventTrigger
dataType: string
status: current
description: This value is used to express more complex trigger conditions that may cause the creation of an event.

7.40. firmwareId

elementId: TBD
name: firmwareId
dataType: string
status: current
description: A label that represents the BIOS or
firmware ID of a specific target endpoint.

7.41. hostName

elementId: TBD
name: hostName
dataType: string
status: current
description: A label typically associated with an
endpoint, but, not always intended to be unique given
scope.

7.42. interfaceLabel

elementId: TBD
name: interfaceLabel
dataType: string
status: current
description: A unique label that can be used to
reference a network interface.

7.43. ipv6AddressSubnetMask

elementId: TBD
name: ipv6AddressSubnetMask
dataType: string
status: current
description: An IPv6 subnet bitmask.

7.44. ipv6AddressSubnetMaskCidrNotation

elementId: TBD
name: ipv6AddressSubnetMaskCidrNotation
dataType: string
status: current
description: An IPv6 subnet bitmask in CIDR notation.

7.45. ipv6AddressValue

elementId: TBD
name: ipv6AddressValue
dataType: ipv6Address
status: current
description: An IPv6 subnet bitmask in CIDR notation.

7.46. ipv4AddressSubnetMask

elementId: TBD
name: ipv4AddressSubnetMask
dataType: string
status: current
description: An IPv4 subnet bitmask.

7.47. ipv4AddressSubnetMaskCidrNotation

elementId: TBD
name: ipv4AddressSubnetMaskCidrNotation
dataType: string
status: current
description: An IPv4 subnet bitmask in CIDR notation.

7.48. ipv4AddressValue

elementId: TBD
name: ipv4AddressValue
dataType: ipv4Address
status: current
description: An IPv4 address value.

7.49. layer2InterfaceType

elementId: TBD
name: layer2InterfaceType
dataType: string
status: current
description: A set of types referenced by IANA ifType.

7.50. layer4PortAddress

elementId: TBD
name: layer4PortAddress
dataType: unsigned32
status: current
description: A layer 4 port address
typically associated with TCP and UDP
protocols.

7.51. layer4Protocol

elementId: TBD
name: layer4Protocol
dataType: string
status: current
description: A set of types that express a layer 4
protocol (e.g. UDP or TCP).

7.52. locationName

elementId: TBD
name: locationName
dataType: string
status: current
description: A value that represents a named region of
physical space.

7.53. networkZoneLocation

elementId: TBD
name: networkZoneLocation
dataType: string
status: current
description: The zone location of an endpoint on the
network (e.g. internet, enterprise DMZ,
enterprise WAN, enclave DMZ, enclave).

7.54. layer2NetworkLocation

elementId: TBD
name: layer2NetworkLocation
dataType: string
status: current
description: The location of a layer-2 interface on
the network (e.g. link-layer neighborhood,
shared broadcast domain).

7.55. layer3NetworkLocation

elementId: TBD
name: layer3NetworkLocation
dataType: string
status: current
description: The location of a layer-3 interface on the network (e.g. next-hop routing neighbor).

7.56. macAddressValue

elementId: TBD
name: macAddressValue
dataType: string
status: current
description: A value that expresses an Ethernet address.

7.57. methodLabel

elementId: TBD
name: methodLabel
dataType: string
status: current
description: A label that references a specific method registered and used in a SACM domain (e.g. method to match and re-identify target endpoints via identifying attributes).

7.58. methodRepository

elementId: TBD
name: methodRepository
dataType: string
status: current
description: A label that references a SACM component methods can be registered at and that can provide guidance in the form of registered methods to other SACM components.

7.59. networkAccessLevelType

elementId: TBD
name: networkAccessLevelType
dataType: string
status: current
description: A set of types that express categories of network access-levels (e.g. block, quarantine, etc.).

7.60. networkId

elementId: TBD
name: networkId
dataType: string
status: current
description: Most networks such as AS, OSBF domains,
or VLANs can have an ID.

7.61. networkInterfaceName

elementId: TBD
name: networkInterfaceName
dataType: string
status: current
description: A label that uniquely identifies an
interface associated with a distinguishable endpoint.

7.62. networkLayer

elementId: TBD
name: networkLayer
dataType: string
status: current
description: A set of layers that expresses the specific
network layer an interface operates on.

7.63. networkName

elementId: TBD
name: networkName
dataType: string
status: current
description: A label that is associated with a network.
Some networks, for example, effective
layer2-broadcast-domains are difficult to "grasp" and
therefore quite difficult to name.

7.64. organizationId

elementId: TBD
name: organizationId
dataType: string
status: current
description: A label that uniquely identifies an
organization via a PEN.

7.65. patchId

elementId: TBD
name: patchId
dataType: string
status: current
description: A label that uniquely identifies a specific software patch.

7.66. patchName

elementId: TBD
name: patchName
dataType: string
status: current
description: The vendor's name of a software patch.

7.67. personFirstName

elementId: TBD
name: personFirstName
dataType: string
status: current
description: The first name of a natural person.

7.68. personLastName

elementId: TBD
name: personLastName
dataType: string
status: current
description: The last name of a natural person.

7.69. personMiddleName

elementId: TBD
name: personMiddleName
dataType: string
status: current
description: The middle name of a natural person.

7.70. phoneNumber

elementId: TBD
name: phoneNumber
dataType: string
status: current
description: A label that expresses the U.S. national
phone number (e.g. pattern value="((\d{3}))?\d{3}-\d{4}").

7.71. phoneNumberType

elementId: TBD
name: phoneNumberType
dataType: string
status: current
description: A set of types that express the type of
a phone number (e.g. DSN, Fax, Home, Mobile, Pager,
Secure, Unsecure, Work, Other).

7.72. privilegeName

elementId: TBD
name: privilegeName
dataType: string
status: current
description: The attribute name of the privilege
represented as an AVP.

7.73. privilegeValue

elementId: TBD
name: privilegeValue
dataType: string
status: current
description: The value content of the privilege
represented as an AVP.

7.74. protocol

elementId: TBD
name: protocol
dataType: string
status: current
description: A set of types that defines specific
protocols above layer 4 (e.g. http, https, dns, ipp,
or unknown).

7.75. publicKey

elementId: TBD
name: publicKey
dataType: string
status: current
description: The value of a public key (regardless of its method of creation, crypto-system, or signature scheme) that can be collected from a target endpoint.

7.76. relationshipContentElementGuid

elementId: TBD
name: relationshipContentElementGuid
dataType: string
status: current
description: A reference to a specific content element used in a relationship subject.

7.77. relationshipStatementElementGuid

elementId: TBD
name: relationshipStatementElementGuid
dataType: string
status: current
description: A reference to a specific SACM statement used in a relationship subject.

7.78. relationshipObjectLabel

elementId: TBD
name: relationshipObjectLabel
dataType: string
status: current
description: A reference to a specific label used in content (e.g. a te-label or a user-id). This reference is typically used if matching content attribute can be done efficiently and can also be included in addition to a relationship-content-element-guid reference.

7.79. relationshipType

elementId: TBD
name: relationshipType
dataType: string
status: current
description: A set of types that is in every instance of a relationship subject to highlight what kind of relationship exists between the subject the relationship is included in (e.g. associated_with_user, applies_to_session, seen_on_interface, associated_with_flow, contains_virtual_device).

7.80. roleName

elementId: TBD
name: roleName
dataType: string
status: current
description: A label that references a collection of privileges assigned to a specific entity.

7.81. sessionStateType

elementId: TBD
name: sessionStateType
dataType: string
status: current
description: A set of types a discernible session (an ongoing network interaction) can be in (e.g. Authenticating, Authenticated, Postured, Started, Disconnected).

7.82. statementGuid

elementId: TBD
name: statementGuid
dataType: string
status: current
description: A label that expresses a global unique ID referencing a specific SACM statement that was produced by a SACM component.

7.83. statementType

elementId: TBD
name: statementType
dataType: string
status: current
description: A set of types that define the type of content that is included in a SACM statement (e.g. Observation, DirectoryContent, Correlation, Assessment, Guidance, Event).

7.84. status

elementId: TBD
name: status
dataType: string
status: current
description: A set of types that defines possible result values for a finding in general (e.g. true, false, error, unknown, not applicable, not evaluated).

7.85. subAdministrativeDomain

elementId: TBD
name: subAdministrativeDomain
dataType: string
status: current
description: A label for related child domains an administrative domain can be composed of (used in the subject administrativeDomain).

7.86. subInterfaceLabel

elementId: TBD
name: subInterfaceLabel
dataType: string
status: current
description: A unique label a sub network interface (e.g. a tagged vlan on a trunk) can be referenced with.

7.87. superAdministrativeDomain

elementId: TBD
name: superAdministrativeDomain
dataType: string
status: current
description: a label for related parent domains an administrative domain is part of (used in the subject administrativeDomain).

7.88. superInterfaceLabel

elementId: TBD
name: superInterfaceLabel
dataType: string
status: current
description: a unique label a super network interface
(e.g. a physical interface a tunnel
interface terminates on) can be referenced
with.

7.89. teAssessmentState

elementId: TBD
name: teAssessmentState
dataType: string
status: current
description: a set of types that defines the state of
assessment of a target-endpoint (e.g.
in-discovery, discovered, in-classification,
classified, in-assessment, assessed).

7.90. teLabel

elementId: TBD
name: teLabel
dataType: string
status: current
description: an identifying label created from a set
of identifying attributes used to reference
a specific target endpoint.

7.91. teId

elementId: TBD
name: teId
dataType: string
status: current
description: an identifying label that is created
randomly, is supposed to be unique, and
used to reference a specific target
endpoint.

7.92. timestampType

elementId: TBD
name: timestampType
dataType: string
status: current
description: a set of types that express what type of action or event happened at that point of time (e.g. discovered, classified, collected, published). Can be included in a generic timestamp subject.

7.93. unitsReceived

elementId: TBD
name: unitsReceived
dataType: string
status: current
description: a value that represents a number of units (e.g. frames, packets, cells or segments) received on a network interface.

7.94. unitsSent

elementId: TBD
name: unitsSent
dataType: string
status: current
description: a value that represents a number of units (e.g. frames, packets, cells or segments) sent on a network interface.

7.95. userDirectory

elementId: TBD
name: userDirectory
dataType: string
status: current
description: a label that identifies a specific type of user-directory (e.g. ldap, active-directory, local-user).

7.96. sacmUserId

elementId: TBD
name: sacmUserId
dataType: string
status: current
description: a label that references a specific user known in a SACM domain.

7.97. webSite

elementId: TBD
name: webSite
dataType: string
status: current
description: a URI that references a web-site.

7.98. WGS84Longitude

elementId: TBD
name: WGS84Longitude
dataType: float64
status: current
description: a label that represents WGS 84 rev 2004 longitude.

7.99. WGS84Latitude

elementId: TBD
name: WGS84Latitude
dataType: float64
status: current
description: a label that represents WGS 84 rev 2004 latitude.

7.100. WGS84Altitude

elementId: TBD
name: WGS84Altitude
dataType: float64
status: current
description: a label that represents WGS 84 rev 2004 altitude.

7.101. hardwareSerialNumber

elementId: TBD
name: hardwareSerialNumber
dataType: string
status: current
description: A globally unique identifier for a particular piece of hardware assigned by the vendor.

7.102. interfaceName

elementId: TBD
name: interfaceName
dataType: string
status: current
description: A short name uniquely describing an interface, e.g. "Eth1/0". See [RFC2863] for the definition of the ifName object.

7.103. interfaceIndex

elementId: TBD
name: interfaceIndex
dataType: unsigned32
status: current
description: The index of an interface installed on an endpoint. The value matches the value of managed object 'ifIndex' as defined in [RFC2863]. Note that ifIndex values are not assigned statically to an interface and that the interfaces may be renumbered every time the device's management system is re-initialized, as specified in [RFC2863].

7.104. interfaceMacAddress

elementId: TBD
name: interfaceMacAddress
dataType: macAddress
status: current
description: The IEEE 802 MAC address associated with a network interface on an endpoint.

7.105. interfaceType

elementId: TBD
name: interfaceType
dataType: unsigned32
status: current
description: The type of a network interface. The value matches the value of managed object 'ifType' as defined in [IANA registry ianaiftype-mib].

7.106. interfaceFlags

elementId: TBD
name: interfaceFlags
dataType: unsigned16
status: current
description: This information element specifies the flags
 associated with a network interface. Possible
 values include:
structure:
 Up ; 0x1 ; Interface is up.
 Broadcast ; 0x2 ; Broadcast address valid.
 Debug ; 0x4 ; Turn on debugging.
 Loopback ; 0x8 ; Is a loopback net.
 Point-to-point ; 0x10 ; Interface is point-to-point
 link.
 No trailers ; 0x20 ; Avoid use of trailers.
 Resources allocated ; 0x40 ; Resources allocated.
 No ARP ; 0x80 ; No address resolution protocol.
 Receive all ; 0x100 ; Receive all packets.

7.107. networkInterface

elementId: TBD
name: networkInterface
dataType: orderedList
status: current
description: Information about a network interface
 installed on an endpoint. The
 following high-level digram
 describes the structure of
 networkInterface information
 element.
structure: orderedList(interfaceName, interfaceIndex, macAddress,
 interfaceType, flags)

7.108. softwareIdentifier

elementId: TBD
name: softwareIdentifier
dataType: string
status: current
description: A globally unique identifier for a particular
 software application.

7.109. softwareTitle

elementId: TBD
name: softwareTitle
dataType: string
status: current
description: The title of the software application.

7.110. softwareCreator

elementId: TBD
name: softwareCreator
dataType: string
status: current
description: The software developer (e.g., vendor or author).

7.111. simpleSoftwareVersion

elementId: TBD
name: simpleSoftwareVersion
dataType: string
status: current
description: The version string for a software application that conforms to the format of a list of hierarchical non-negative integers separated by a single character delimiter format.

7.112. rpmSoftwareVersion

elementId: TBD
name: rpmSoftwareVersion
dataType: string
status: current
description: The version string for a software application that conforms to the EPOCH:VERSION-RELEASE format.

7.113. ciscoTrainSoftwareVersion

elementId: TBD
name: ciscoTrainSoftwareVersion
dataType: string
status: current
description: The version string for a software application that conforms to the Cisco IOS Train string format.

7.114. softwareVersion

```
elementId: TBD
name: softwareVerison
dataType: category
status: current
description: The version of the software application. Software
             applications may be versioned using a number of
             schemas. The following high-level digram describes
             the structure of the softwareVersion information
             element.
structure: category(simpleSoftwareVersion | rpmSoftwareVersion |
                   ciscoTrainSoftwareVersion)
```

7.115. softwareLastUpdated

```
elementId: TBD
name: softwareLastUpdated
dataType: dateTimeSeconds
status: current
description: The date and time when the software instance
             was last updated on the system (e.g., new
             version installed or patch applied)
```

7.116. softwareClass


```
    elementId: TBD
    name: softwareClass
    dataType: enumeration
    status: current
    description: The class of the software instance.
    structure:
      Unknown          ; 0x1 ; The class is not known.
      Other            ; 0x2 ; The class is known, but,
                        ;     ; something other than a value
                        ;     ; listed in the enumeration.
      Driver           ; 0x3 ; The class is a device driver.
      Configuration Software ; 0x4 ; The class is configuration
                        ;     ; software.
      Application Software ; 0x5 ; The class is application
                        ;     ; software.
      Instrumentation    ; 0x6 ; The class is instrumentation.
      Diagnostic Software ; 0x8 ; The class is diagnostic
                        ;     ; software.
      Operating System   ; 0x9 ; The class is operating
                        ;     ; system.
      Middleware         ; 0xA ; The class is middleware.
      Firmware           ; 0xB ; The class is firmware.
      BIOS/FCode         ; 0xC ; The class is BIOS or FCode.
      Support/Service Pack ; 0xD ; The class is a support or
                        ;     ; service pack.
      Software Bundle    ; 0xE ; The class is a software
                        ;     ; bundle.
    References: See Classifications of the DMTF
                CIM_SoftwareIdentity schema.
```

7.117. softwareInstance

```
    elementId: TBD
    name: softwareInstance
    dataType: orderedList
    status: current
    description: Information about an instance of software
                installed on an endpoint. The following
                high-level digram describes the structure of
                the softwareInstance information element.
    structure: orderedList(softwareIdentifier, softwareTitle,
                          softwareCreator, softwareVersion,
                          softwareLastUpdated, softwareClass)
```

7.118. globallyUniqueIdentifier

elementId: TBD
name: globallyUniqueIdentifier
dataType: unsigned8
status: current
description: TODO.

7.119. creationTimestamp

elementId: TBD
name: creationTimestamp
dataType: dateTimeSeconds
status: current
description: The date and time when the posture
information was created by a SACM Component.

7.120. collectionTimestamp

elementId: TBD
name: collectionTimestamp
dataType: dateTimeSeconds
status: current
description: The date and time when the posture
information was collected or observed by a SACM
Component.

7.121. publicationTimestamp

elementId: TBD
name: publicationTimestamp
dataType: dateTimeSeconds
status: current
description: The date and time when the posture
information was published.

7.122. relayTimestamp

elementId: TBD
name: relayTimestamp
dataType: dateTimeSeconds
status: current
description: The date and time when the posture
information was relayed to another SACM Component.

7.123. storageTimestamp

elementId: TBD
name: storageTimestamp
dataType: dateTimeSeconds
status: current
description: The date and time when the posture
information was stored in a Repository.

7.124. type

elementId: TBD
name: type
dataType: enumeration
status: current
description: The type of data model use to represent
some set of endpoint information. The following
table lists the set of data models supported by SACM.
structure: TBD

7.125. protocolIdentifier

elementId: TBD
name: protocolIdentifier
dataType: unsigned8
status: current
description: The value of the protocol number in the IP packet
header. The protocol number identifies the IP packet
payload type. Protocol numbers are defined in the
IANA Protocol Numbers registry.

In Internet Protocol version 4 (IPv4), this is
carried in the Protocol field. In Internet Protocol
version 6 (IPv6), this is carried in the Next Header
field in the last extension header of the packet.

7.126. sourceTransportPort

elementId: TBD
name: sourceTransportPort
dataType: unsigned16
status: current
description: The source port identifier in the transport header.
For the transport protocols UDP, TCP, and SCTP, this
is the source port number given in the respective
header. This field MAY also be used for future
transport protocols that have 16-bit source port
identifiers.

7.127. sourceIPv4PrefixLength

elementId: TBD
name: sourceIPv4PrefixLength
dataType: unsigned8
status: current
description: The number of contiguous bits that are relevant in the sourceIPv4Prefix Information Element.

7.128. ingressInterface

elementId: TBD
name: ingressInterface
dataType: unsigned32
status: current
description: The index of the IP interface where packets of this Flow are being received. The value matches the value of managed object 'ifIndex' as defined in [RFC2863]. Note that ifIndex values are not assigned statically to an interface and that the interfaces may be renumbered every time the device's management system is re-initialized, as specified in [RFC2863].

7.129. destinationTransportPort

elementId: TBD
name: destinationTransportPort
dataType: unsigned16
status: current
description: The destination port identifier in the transport header. For the transport protocols UDP, TCP, and SCTP, this is the destination port number given in the respective header. This field MAY also be used for future transport protocols that have 16-bit destination port identifiers.

7.130. sourceIPv6PrefixLength

elementId: TBD
name: sourceIPv6PrefixLength
dataType: unsigned8
status: current
description: The number of contiguous bits that are relevant in the sourceIPv6Prefix Information Element.

7.131. sourceIPv4Prefix

elementId: TBD
name: sourceIPv4Prefix
dataType: ipv4Address
status: current
description: IPv4 source address prefix.

7.132. destinationIPv4Prefix

elementId: TBD
name: destinationIPv4Prefix
dataType: ipv4Address
status: current
description: IPv4 destination address prefix.

7.133. sourceMacAddress

elementId: TBD
name: sourceMacAddress
dataType: macAddress
status: current
description: The IEEE 802 source MAC address field.

7.134. ipVersion

elementId: TBD
name: ipVersion
dataType: unsigned8
status: current
description: The IP version field in the IP packet header.

7.135. interfaceDescription

elementId: TBD
name: interfaceDescription
dataType: string
status: current
description: The description of an interface, e.g.
"FastEthernet 1/0" or "ISP connection".

7.136. applicationDescription

elementId: TBD
name: applicationDescription
dataType: string
status: current
description: Specifies the description of an application.

7.137. applicationId

elementId: TBD
name: applicationId
dataType: octetArray
status: current
description: Specifies an Application ID per [RFC6759].

7.138. applicationName

elementId: TBD
name: applicationName
dataType: string
status: current
description: Specifies the name of an application.

7.139. exporterIPv4Address

elementId: TBD
name: exporterIPv4Address
dataType: ipv4Address
status: current
description: The IPv4 address used by the Exporting Process.
This is used by the Collector to identify the
Exporter in cases where the identity of the Exporter
may have been obscured by the use of a proxy.

7.140. exporterIPv6Address

elementId: TBD
name: exporterIPv6Address
dataType: ipv6Address
status: current
description: The IPv6 address used by the Exporting Process.
This is used by the Collector to identify the
Exporter in cases where the identity of the
Exporter may have been obscured by the use of a
proxy.

7.141. portId

elementId: TBD
name: portId
dataType: unsigned32
status: current
description: An identifier of a line port that is unique per
IPFIX Device hosting an Observation Point.
Typically, this Information Element is used for
limiting the scope of other Information Elements.

7.142. templateId

elementId: TBD
name: templateId
dataType: unsigned16
status: current
description: An identifier of a Template that is locally unique
within a combination of a Transport session and an
Observation Domain.

Template IDs 0-255 are reserved for Template Sets,
Options Template Sets, and other reserved Sets yet
to be created. Template IDs of Data Sets are
numbered from 256 to 65535.

Typically, this Information Element is used for
limiting the scope of other Information Elements.
Note that after a re-start of the Exporting Process
Template identifiers may be re-assigned.

7.143. collectorIPv4Address

elementId: TBD
name: collectorIPv4Address
dataType: ipv4Address
status: current
description: An IPv4 address to which the Exporting Process sends
Flow information.

7.144. collectorIPv6Address

elementId: TBD
name: collectorIPv6Address
dataType: ipv6Address
status: current
description: An IPv6 address to which the Exporting Process sends
Flow information.

7.145. informationElementIndex

elementId: TBD
name: informationElementIndex
dataType: unsigned16
status: current
description: A zero-based index of an Information Element
referenced by informationElementId within a Template
referenced by templateId; used to disambiguate
scope for templates containing multiple identical
Information Elements.

7.146. informationElementId

elementId: TBD
name: informationElementId
dataType: unsigned16
status: current
description: This Information Element contains the ID of another
Information Element.

7.147. informationElementDataType

elementId: TBD
name: informationElementDataType
dataType: unsigned8
status: current
description: A description of the abstract data type of an IPFIX
information element. These are taken from the
abstract data types defined in section 3.1 of the
IPFIX Information Model [RFC5102]; see that section
for more information on the types described in the
informationElementDataType sub-registry.

These types are registered in the IANA IPFIX
Information Element Data Type subregistry. This
subregistry is intended to assign numbers for type
names, not to provide a mechanism for adding data
types to the IPFIX Protocol, and as such requires a
Standards Action [RFC5226] to modify.

7.148. informationElementDescription

elementId: TBD
name: informationElementDescription
dataType: string
status: current
description: A UTF-8 [RFC3629] encoded Unicode string containing a human-readable description of an Information Element. The content of the informationElementDescription MAY be annotated with one or more language tags [RFC4646], encoded in-line [RFC2482] within the UTF-8 string, in order to specify the language in which the description is written. Description text in multiple languages MAY tag each section with its own language tag; in this case, the description information in each language SHOULD have equivalent meaning. In the absence of any language tag, the "i-default" [RFC2277] language SHOULD be assumed. See the Security Considerations section for notes on string handling for Information Element type records.

7.149. informationElementName

elementId: TBD
name: informationElementName
dataType: string
status: current
description: A UTF-8 [RFC3629] encoded Unicode string containing the name of an Information Element, intended as a simple identifier. See the Security Considerations section for notes on string handling for Information Element type records.

7.150. informationElementRangeBegin

elementId: TBD
name: informationElementRangeBegin
dataType: unsigned64
status: current
description: Contains the inclusive low end of the range of acceptable values for an Information Element.

7.151. informationElementRangeEnd

elementId: TBD
name: informationElementRangeEnd
dataType: unsigned64
status: current
description: Contains the inclusive high end of the range of
 acceptable values for an Information Element.

7.152. informationElementSemantics

elementId: TBD
name: informationElementSemantics
dataType: unsigned8
status: current
description: A description of the semantics of an IPFIX
 Information Element. These are taken from the data
 type semantics defined in section 3.2 of the IPFIX
 Information Model [RFC5102]; see that section for
 more information on the types defined in the
 informationElementSemantics sub-registry. This
 field may take the values in Table ; the special
 value 0x00 (default) is used to note that no
 semantics apply to the field; it cannot be
 manipulated by a Collecting Process or File Reader
 that does not understand it a priori.

These semantics are registered in the IANA IPFIX
Information Element Semantics subregistry. This
subregistry is intended to assign numbers for
semantics names, not to provide a mechanism for
adding semantics to the IPFIX Protocol, and as such
requires a Standards Action [RFC5226] to modify.

7.153. informationElementUnits

elementId: TBD
name: informationElementUnits
dataType: unsigned16
status: current
description: A description of the units of an IPFIX Information Element. These correspond to the units implicitly defined in the Information Element definitions in section 5 of the IPFIX Information Model [RFC5102]; see that section for more information on the types described in the informationElementsUnits sub-registry. This field may take the values in Table 3 below; the special value 0x00 (none) is used to note that the field is unitless.

These types are registered in the IANA IPFIX Information Element Units subregistry; new types may be added on a First Come First Served [RFC5226] basis.

7.154. applicationCategoryName

elementId: TBD
name: applicationCategoryName
dataType: string
status: current
description: An attribute that provides a first level categorization for each Application ID.

7.155. mibObjectValueInteger

elementId: TBD
name: mibObjectValueInteger
dataType: signed64
status: current
description: An IPFIX Information Element which denotes that the integer value of a MIB object will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with the Base Syntax of Integer32 and INTEGER with IPFIX Reduced Size Encoding used as required. The value is encoded as per the standard IPFIX Abstract Data Type of signed64.

7.156. mibObjectValueOctetString

elementId: TBD
name: mibObjectValueOctetString
dataType: octetArray
status: current
description: An IPFIX Information Element which denotes that an Octet String or Opaque value of a MIB object will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with the Base Syntax of OCTET STRING and Opaque. The value is encoded as per the standard IPFIX Abstract Data Type of octetArray.

7.157. mibObjectValueOID

elementId: TBD
name: mibObjectValueOID
dataType: octetArray
status: current
description: An IPFIX Information Element which denotes that an Object Identifier or OID value of a MIB object will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with the Base Syntax of OBJECT IDENTIFIER. Note - In this case the "mibObjectIdentifier" will define which MIB object is being exported while the value contained in this Information Element will be an OID as a value. The mibObjectValueOID Information Element is encoded as ASN.1/BER [BER] in an octetArray.

7.158. mibObjectValueBits

elementId: TBD
name: mibObjectValueBits
dataType: octetArray
status: current
description: An IPFIX Information Element which denotes that a set of Enumerated flags or bits from a MIB object will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with the Base Syntax of BITS. The flags or bits are encoded as per the standard IPFIX Abstract Data Type of octetArray, with sufficient length to accommodate the required number of bits. If the number of bits is not an integer multiple of octets then the most significant bits at end of the octetArray MUST be set to zero.

7.159. mibObjectValueIPAddress

elementId: TBD
name: mibObjectValueIPAddress
dataType: ipv4Address
status: current
description: An IPFIX Information Element which denotes that the IPv4 Address of a MIB object will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with the Base Syntax of IPAddress. The value is encoded as per the standard IPFIX Abstract Data Type of ipv4Address.

7.160. mibObjectValueCounter

elementId: TBD
name: mibObjectValueCounter
dataType: unsigned64
status: current
description: An IPFIX Information Element which denotes that the counter value of a MIB object will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with the Base Syntax of Counter32 or Counter64 with IPFIX Reduced Size Encoding used as required. The value is encoded as per the standard IPFIX Abstract Data Type of unsigned64.

7.161. mibObjectValueGauge

elementId: TBD
name: mibObjectValueGauge
dataType: unsigned32
status: current
description: An IPFIX Information Element which denotes that the Gauge value of a MIB object will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with the Base Syntax of Gauge32. The value is encoded as per the standard IPFIX Abstract Data Type of unsigned64. This value will represent a non-negative integer, which may increase or decrease, but shall never exceed a maximum value, nor fall below a minimum value.

7.162. mibObjectValueTimeTicks

elementId: TBD
name: mibObjectValueTimeTicks
dataType: unsigned32
status: current
description: An IPFIX Information Element which denotes that the TimeTicks value of a MIB object will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with the Base Syntax of TimeTicks. The value is encoded as per the standard IPFIX Abstract Data Type of unsigned32.

7.163. mibObjectValueUnsigned

elementId: TBD
name: mibObjectValueUnsigned
dataType: unsigned64
status: current
description: An IPFIX Information Element which denotes that an unsigned integer value of a MIB object will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with the Base Syntax of unsigned64 with IPFIX Reduced Size Encoding used as required. The value is encoded as per the standard IPFIX Abstract Data Type of unsigned64.

7.164. mibObjectValueTable

elementId: TBD
name: mibObjectValueTable
dataType: orderedList
status: current
description: An IPFIX Information Element which denotes that a complete or partial conceptual table will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with a SYNTAX of SEQUENCE. This is encoded as a subTemplateList of mibObjectValue Information Elements. The template specified in the subTemplateList MUST be an Options Template and MUST include all the Objects listed in the INDEX clause as Scope Fields.
structure: orderedList(mibObjectValueRow+)

7.165. mibObjectValueRow

elementId: TBD
name: mibObjectValueRow
dataType: orderedList
status: current
description: An IPFIX Information Element which denotes that a single row of a conceptual table will be exported. The MIB Object Identifier ("mibObjectIdentifier") for this field MUST be exported in a MIB Field Option or via another means. This Information Element is used for MIB objects with a SYNTAX of SEQUENCE. This is encoded as a subTemplateList of mibObjectValue Information Elements. The subTemplateList exported MUST contain exactly one row (i.e., one instance of the subtemplate). The template specified in the subTemplateList MUST be an Options Template and MUST include all the Objects listed in the INDEX clause as Scope Fields.
structure: orderedList(mibObjectValue+)

7.166. mibObjectIdentifier

elementId: TBD
name: mibObjectIdentifier
dataType: octetArray
status: current
description: An IPFIX Information Element which denotes that a MIB Object Identifier (MIB OID) is exported in the (Options) Template Record. The mibObjectIdentifier Information Element contains the OID assigned to the MIB Object Type Definition encoded as ASN.1/BER [BER].

7.167. mibSubIdentifier

elementId: TBD
name: mibSubIdentifier
dataType: unsigned32
status: current
description: A non-negative sub-identifier of an Object Identifier (OID).

7.168. mibIndexIndicator

elementId: TBD
name: mibIndexIndicator
dataType: unsigned64
status: current
description: This set of bit fields is used for marking the Information Elements of a Data Record that serve as INDEX MIB objects for an indexed Columnar MIB object. Each bit represents an Information Element in the Data Record with the n-th bit representing the n-th Information Element. A bit set to value 1 indicates that the corresponding Information Element is an index of the Columnar Object represented by the mibFieldValue. A bit set to value 0 indicates that this is not the case.

If the Data Record contains more than 64 Information Elements, the corresponding Template SHOULD be designed such that all INDEX Fields are among the first 64 Information Elements, because the mibIndexIndicator only contains 64 bits. If the Data Record contains less than 64 Information Elements, then the extra bits in the mibIndexIndicator for which no corresponding Information Element exists MUST have the value 0, and must be disregarded by the Collector. This Information Element may be exported with IPFIX Reduced Size Encoding.

7.169. mibCaptureTimeSemantics

elementId: TBD
name: mibCaptureTimeSemantics
dataType: unsigned8
status: current
description: Indicates when in the lifetime of the flow the MIB value was retrieved from the MIB for a mibObjectIdentifier. This is used to indicate if the value exported was collected from the MIB closer to flow creation or flow export time and will refer to the Timestamp fields included in the same record. This field SHOULD be used when exporting a mibObjectValue that specifies counters or statistics.

If the MIB value was sampled by SNMP prior to the IPFIX Metering Process or Exporting Process retrieving the value (i.e., the data is already stale) and it's important to know the exact sampling time, then an additional observationTime* element should be paired with the OID using structured data. Similarly, if different mibCaptureTimeSemantics apply to different mibObject elements within the Data Record, then individual mibCaptureTimeSemantics should be paired with each OID using structured data.

Values:

- 0. undefined
- 1. begin - The value for the MIB object is captured from the MIB when the Flow is first observed
- 2. end - The value for the MIB object is captured from the MIB when the Flow ends
- 3. export - The value for the MIB object is captured from the MIB at export time
- 4. average - The value for the MIB object is an average of multiple captures from the MIB over the observed life of the Flow

7.170. mibContextEngineID

elementId: TBD
name: mibContextEngineID
dataType: octetArray
status: current
description: A mibContextEngineID that specifies the SNMP engine ID for a MIB field being exported over IPFIX. Definition as per [RFC3411] section 3.3.

7.171. mibContextName

elementId: TBD
name: mibContextName
dataType: string
status: current
description: This Information Element denotes that a MIB Context Name is specified for a MIB field being exported over IPFIX. Reference [RFC3411] section 3.3.

7.172. mibObjectName

elementId: TBD
name: mibObjectName
dataType: string
status: current
description: The name (called a descriptor in [RFC2578]) of an object type definition.

7.173. mibObjectDescription

elementId: TBD
name: mibObjectDescription
dataType: string
status: current
description: The value of the DESCRIPTION clause of an MIB object type definition.

7.174. mibObjectSyntax

elementId: TBD
name: mibObjectSyntax
dataType: string
status: current
description: The value of the SYNTAX clause of an MIB object type definition, which may include a Textual Convention or Subtyping. See [RFC2578].

7.175. mibModuleName

elementId: TBD
name: mibModuleName
dataType: string
status: current
description: The textual name of the MIB module that defines a MIB Object.

7.176. interface

elementId: TBD
name: interface
dataType: list
structure: list (interfaceName, hwAddress, inetAddr, netmask)
status: current
description: Represents an interface and its configuration options.

7.177. iflisteners

elementId: TBD
name: iflisteners
dataType: list
structure: list (interfaceName, physicalProtocol, hwAddress, programName, pid, userId)
status: current
description: Stores the results of checking for applications that are bound to an ethernet interface on the system.

7.178. physicalProtocol

elementId: TBD
name: physicalProtocol
dataType: enumeration
structure:
ETH_P_LOOP ; 0x1 ; Ethernet loopback packet.
ETH_P_PUP ; 0x2 ; Xerox PUP packet.
ETH_P_PUPAT ; 0x3 ; Xerox PUP Address Transport packet.
ETH_P_IP ; 0x4 ; Internet protocol packet.
ETH_P_X25 ; 0x5 ; CCITT X.25 packet.
ETH_P_ARP ; 0x6 ; Address resolution packet.
ETH_P_BPQ ; 0x7 ; G8BPQ AX.25 ethernet packet.
ETH_P_IEEE802.3_PUP ; 0x8 ; Xerox IEEE802.3 PUP packet.
ETH_P_IEEE802.3_PUPAT ; 0x9 ; Xerox IEEE802.3 PUP address transport packet.
ETH_P_DEC ; 0xA ; DEC assigned protocol.
ETH_P_DNA_DL ; 0xB ; DEC DNA Dump/Load.
ETH_P_DNA_RC ; 0xC ; DEC DNA Remote Console.
ETH_P_DNA_RT ; 0xD ; DEC DNA Routing.
ETH_P_LAT ; 0xE ; DEC LAT.
ETH_P_DIAG ; 0xF ; DEC Diagnostics.
ETH_P_CUST ; 0x10 ; DEC Customer use.
ETH_P_SCA ; 0x11 ; DEC Systems Comms Arch.
ETH_P_RARP ; 0x12 ; Reverse address resolution packet.
ETH_P_ATALK ; 0x13 ; Appletalk DDP.
ETH_P_AARP ; 0x14 ; Appletalk AARP.

ETH_P_8021Q ; 0x15 ; 802.1Q VLAN Extended Header.
ETH_P_IPX ; 0x16 ; IPX over DIX.
ETH_P_IPV6 ; 0x17 ; IPv6 over bluebook.
ETH_P_SLOW ; 0x18 ; Slow Protocol. See 802.3ad 43B.
ETH_P_WCCP ; 0x19 ; Web-cache coordination protocol.
ETH_P_PPP_DISC ; 0x1A ; PPPoE discovery messages.
ETH_P_PPP_SES ; 0x1B ; PPPoE session messages.
ETH_P_MPLS_UC ; 0x1C ; MPLS Unicast traffic.
ETH_P_MPLS_MC ; 0x1D ; MPLS Multicast traffic.
ETH_P_ATMMPOA ; 0x1E ; MultiProtocol Over ATM.
ETH_P_ATMFATE ; 0x1F ; Frame-based ATM Transport over Ethernet.
ETH_P_AOE ; 0x20 ; ATA over Ethernet.
ETH_P_TIPC ; 0x21 ; TIPC.
ETH_P_802_3 ; 0x22 ; Dummy type for 802.3 frames.
ETH_P_AX25 ; 0x23 ; Dummy protocol id for AX.25.
ETH_P_ALL ; 0x24 ; Every packet.
ETH_P_802_2 ; 0x25 ; 802.2 frames.
ETH_P_SNAP ; 0x26 ; Internal only.
ETH_P_DDCMP ; 0x27 ; DEC DDCMP: Internal only
ETH_P_WAN_PPP ; 0x28 ; Dummy type for WAN PPP frames.
ETH_P_PPP_MP ; 0x29 ; Dummy type for PPP MP frames.
ETH_P_PPPTALK ; 0x2A ; Dummy type for Atalk over PPP.
ETH_P_LOCALTALK ; 0x2B ; Localtalk pseudo type.
ETH_P_TR_802_2 ; 0x2C ; 802.2 frames.
ETH_P_MOBITEX ; 0x2D ; Mobitex.
ETH_P_CONTROL ; 0x2E ; Card specific control frames.
ETH_P_IRDA ; 0x2F ; Linux-IrDA.
ETH_P_ECONET ; 0x30 ; Acorn Econet.
ETH_P_HDLC ; 0x31 ; HDLC frames.
ETH_P_ARCNET ; 0x32 ; 1A for ArcNet.
 ; 0x33 ; The empty string value is permitted here
 to allow for detailed error reporting.
status: current
description: The physical layer protocol used by the AF_PACKET
socket.

7.179. hwAddress

elementId: TBD
name: hwAddress
dataType: string
status: current
description: The hardware address associated
 with the interface.

7.180. programName

elementId: TBD
name: programName
dataType: string
status: current
description: The name of the communicating program.

7.181. userId

elementId: TBD
name: userId
dataType: unsigned32
status: current
description: The numeric user id.

7.182. inetlisteningserver

elementId: TBD
name: inetlisteningserver
dataType: list
structure: list (transportProtocol, localAddress, localPort, localFullAddress, programName, foreignAddress, foreignPort, foreignFullAddress, pid, userId)
status: current
description: Stores the results of checking for network servers currently active on a system. It holds information pertaining to a specific protocol-address-port combination.

7.183. transportProtocol

elementId: TBD
name: transportProtocol
dataType: string
status: current
description: The transport-layer protocol (tcp or udp).

7.184. localAddress

elementId: TBD
name: localAddress
dataType: ipAddress
status: current
description: This is the IP address being listened to. Note that the IP address can be IPv4 or IPv6.

7.185. localPort

elementId: TBD
name: localPort
dataType: unsigned32
status: current
description: This is the TCP or UDP port
being listened to.

7.186. localFullAddress

elementId: TBD
name: localFullAddress
dataType: string
status: current
description: The IP address and network port on which the program
listens, including the local address and the local port. Note
that the IP address can be IPv4 or IPv6.

7.187. foreignAddress

elementId: TBD
name: foreignAddress
dataType: ipAddress
status: current
description: The IP address with which the program is
communicating, or with which it will communicate. Note that the
IP address can be IPv4 or IPv6.

7.188. foreignFullAddress

elementId: TBD
name: foreignFullAddress
dataType: ipAddress
status: current
description: The IP address and network port to which the program
is communicating or will accept communications from, including
the foreign address and foreign port. Note that the IP address
can be IPv4 or IPv6.

7.189. selinuxboolean

elementId: TBD
name: selinuxboolean
dataType: list
structure: list (selinuxName, currentStatus,
pendingStatus)
status: current
description: Describes the current and pending status of a
SELinux boolean.

7.190. selinuxName

elementId: TBD
name: selinuxName
dataType: string
status: current
description: The name of the SELinux
boolean.

7.191. currentStatus

elementId: TBD
name: currentStatus
dataType: boolean
status: current
description: Indicates current state of
the specified SELinux boolean.

7.192. pendingStatus

elementId: TBD
name: pendingStatus
dataType: boolean
status: current
description: Indicates the pending
state of the specified SELinux boolean.

7.193. selinuxsecuritycontext

elementId: TBD
name: selinuxsecuritycontext
dataType: list
structure: list (filepath, path, filename, pid,
username, role, domainType, lowSensitivity, lowCategory,
highSensitivity, highCategory, rawlowSensitivity,
rawlowCategory, rawhighSensitivity, rawhighCategory)
status: current
description: Describes the SELinux security
context of a file or process on the local system.

7.194. filepath

elementId: TBD
name: filepath
dataType: string
status: current
description: Specifies the absolute path for a file on the machine. A directory cannot be specified as a filepath.

7.195. path

elementId: TBD
name: path
dataType: string
status: current
description: Specifies the directory component of the absolute path to a file on the machine.

7.196. filename

elementId: TBD
name: filename
dataType: string
status: current
description: The name of the file.

7.197. pid

elementId: TBD
name: pid
dataType: unsigned32
status: current
description: The process ID of the process.

7.198. role

elementId: TBD
name: role
dataType: string
status: current
description: Specifies the types that a process may transition to (domain transitions).

7.199. domainType

elementId: TBD
name: domainType
dataType: string
status: current
description: Specifies the domain in which the file is accessible
or the domain in which a process executes.

7.200. lowSensitivity

elementId: TBD
name: lowSensitivity
dataType: string
status: current
description: Specifies the current sensitivity of a file or
process.

7.201. lowCategory

elementId: TBD
name: lowCategory
dataType: string
status: current
description: Specifies the set of
categories associated with the low sensitivity.

7.202. highSensitivity

elementId: TBD
name: highSensitivity
dataType: string
status: current
description: Specifies the maximum
range for a file or the clearance for a process.

7.203. highCategory

elementId: TBD
name: highCategory
dataType: string
status: current
description: Specifies the set of
categories associated with the high sensitivity.

7.204. rawlowSensitivity

elementId: TBD
name: rawlowSensitivity
dataType: string
status: current
description: Specifies the current sensitivity of a file or process but in its raw context.

7.205. rawlowCategory

elementId: TBD
name: rawlowCategory
dataType: string
status: current
description: Specifies the set of categories associated with the low sensitivity but in its raw context.

7.206. rawhighSensitivity

elementId: TBD
name: rawhighSensitivity
dataType: string
status: current
description: Specifies the maximum range for a file or the clearance for a process but in its raw context.

7.207. rawhighCategory

elementId: TBD
name: rawhighCategory
dataType: string
status: current
description: Specifies the set of categories associated with the high sensitivity but in its raw context.

7.208. systemdunitdependency

elementId: TBD
name: systemdunitdependency
dataType: list
structure: list (unit, dependency)
status: current

description: Stores the dependencies of the systemd unit.

7.209. unit

elementId: TBD
name: unit
dataType: string
status: current
description: Refers to the full systemd unit name, which has a form of "\$name.\$type". For example "cupsd.service". This name is usually also the filename of the unit configuration file.

7.210. dependency

elementId: TBD
name: dependency
dataType: string
status: current
description: Refers to the name of a unit that was confirmed to be a dependency of the given unit.

7.211. systemdunitproperty

elementId: TBD
name: systemdunitproperty
dataType: list
structure: list (unit, property, systemdunitValue)

status: current
description: Stores the properties and values of a systemd unit.

7.212. property

elementId: TBD
name: property
dataType: string
status: current
description: The property associated with a systemd unit.

7.213. systemdunitValue

elementId: TBD
name: systemdunitValue
dataType: string
status: current
description: The value of the property associated with a systemd unit. Exactly one value shall be used for all property types except dbus arrays - each array element shall be represented by one value.

7.214. file

elementId: TBD
name: file
dataType: list
structure: list (filepath, path, filename, fileType, userId, aTime, cTime, mTime, size)
status: current
description: The metadata associated with a file on the endpoint.

7.215. fileType

elementId: TBD
name: fileType
dataType: string
status: current
description: The file's type (e.g., regular file (regular), directory, named pipe (fifo), symbolic link, socket or block special.)

7.216. groupId

elementId: TBD
name: groupId
dataType: unsigned32
status: current
description: The group owner of the file, by group number.

7.217. aTime

elementId: TBD
name: aTime
dataType: dateTimeSeconds
status: current
description: The time that the file was last accessed.

7.218. cTime

elementId: TBD
name: cTime
dataType: dateTimeSeconds
status: current
description: The time of the last change to the file's inode.

7.219. mTime

elementId: TBD
name: mTime
dataType: dateTimeSeconds
status: current
description: The time of the last change to
the file's contents.

7.220. size

elementId: TBD
name: size
dataType: unsigned32
status: current
description: This is the size of the file in
bytes.

7.221. suid

elementId: TBD
name: suid
dataType: boolean
status: current
description: Indicates whether the program runs with the uid
(thus privileges) of the file's owner, rather than the calling
user.

7.222. sgid

elementId: TBD
name: sgid
dataType: boolean
status: current
description: Indicates whether the program runs with the gid
(thus privileges) of the file's group owner, rather than the
calling user's group.

7.223. sticky

elementId: TBD
name: sticky
dataType: boolean
status: current
description: Indicates whether users can delete each other's
files in this directory, when said directory is writable by
those users.

7.224. hasExtendedAcl

elementId: TBD
name: hasExtendedAcl
dataType: boolean
status: current
description: Indicates whether the file or directory has ACL permissions applied to it. If a system supports ACLs and the file or directory doesn't have an ACL, or it matches the standard UNIX permissions, the entity will have a status of 'exists' and a value of 'false'. If the system supports ACLs and the file or directory has an ACL, the entity will have a status of 'exists' and a value of 'true'. Lastly, if a system doesn't support ACLs, the entity will have a status of 'does not exist'.

7.225. inetd

elementId: TBD
name: inetd
dataType: list
structure: list (serviceProtocol, serviceName, serverProgram, serverArguments, inetdEndpointType, execAsUser, waitStatus)
status: current
description: Holds information associated with different Internet services.

7.226. serverProgram

elementId: TBD
name: serverProgram
dataType: string
status: current
description: Either the pathname of a server program to be invoked by inetd to perform the requested service, or the value internal if inetd itself provides the service.

7.227. inetdEndpointType

elementId: TBD
name: inetdEndpointType
dataType: enumeration
structure:
 stream ; 0x1 ; The stream value is used to describe a stream socket.
 dgram ; 0x2 ; The dgram value is used to describe a datagram socket.
 raw ; 0x3 ; The raw value is used to describe a raw socket.
 seqpacket ; 0x4 ; The seqpacket value is used to describe a sequenced packet socket.
 tli ; 0x5 ; The tli value is used to describe all TLI endpoints.
 sunrpc_tcp ; 0x6 ; The sunrpc_tcp value is used to describe all SUNRPC TCP endpoints.
 sunrpc_udp ; 0x7 ; The sunrpc_udp value is used to describe all SUNRPC UDP endpoints.
 ; 0x8 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: The endpoint type (aka, socket type) associated with the service.

7.228. execAsUser

elementId: TBD
name: execAsUser
dataType: string
status: current
description: The user id of the user the server program should run under.

7.229. waitStatus

elementId: TBD
name: waitStatus
dataType: enumeration
structure: wait ; 0x1 ; The value of 'wait' specifies that the server that is invoked by inetd will take over the listening socket associated with the service, and once launched, inetd will wait for that server to exit, if ever, before it resumes listening for new service requests.

nowait ; 0x2 ; The value of 'nowait' specifies that the server that is invoked by inetd will not wait for any existing server to finish before taking over the listening socket associated with the service.

; 0x3 ; The empty string value is permitted here to allow for detailed error reporting.

status: current

description: Specifies whether the server that is invoked by inetd will take over the listening socket associated with the service, and whether once launched, inetd will wait for that server to exit, if ever, before it resumes listening for new service requests. The legal values are "wait" or "nowait".

7.230. inetAddr

elementId: TBD
name: inetAddr
dataType: ipAddress
status: current
description: The IP address of the specific interface. Note that the IP address can be IPv4 or IPv6.

7.231. netmask

elementId: TBD
name: netmask
dataType: ipAddress
status: current
description: The bitmask used to calculate the interface's IP network.

7.232. passwordInfo

elementId: TBD
name: passwordInfo
dataType: list
structure: list (username, password, userId, groupId, gcos,
homeDir, loginShell, lastLogin)
status: current
description: Describes user account information for a
system.

7.233. username

elementId: TBD
name: username
dataType: string
status: current
description: The name of the user.

7.234. password

elementId: TBD
name: password
dataType: string
status: current
description: The encrypted version of the
user's password.

7.235. gcos

elementId: TBD
name: gcos
dataType: string
status: current
description:

7.236. homeDir

elementId: TBD
name: homeDir
dataType: string
status: current
description: The user's home
directory.

7.237. loginShell

elementId: TBD
name: loginShell
dataType: string
status: current
description: The user's shell
program.

7.238. lastLogin

elementId: TBD
name: lastLogin
dataType: unsigned32
status: current
description: The date and time when the
last login occurred.

7.239. process

elementId: TBD
name: process
dataType: list
structure: list (commandLine, pid, ppid, priority, startTime)

status: current
description: Information about a process running on an endpoint.

7.240. commandLine

elementId: TBD
name: commandLine
dataType: string
status: current
description: The string used to start the
process. This includes any parameters that are part of the
command line.

7.241. ppid

elementId: TBD
name: ppid
dataType: unsigned32
status: current
description: The process ID of the process's
parent process.

7.242. priority

elementId: TBD
name: priority
dataType: unsigned32
status: current
description: The scheduling priority with
which the process runs.

7.243. startTime

elementId: TBD
name: startTime
dataType: string
status: current
description: The time of day the process
started.

7.244. routingtable

elementId: TBD
name: routingtable
dataType: list
structure: list (destination, gateway, flags,
interfaceName)
status: current
description: Holds information about an individual routing table
entry found in a system's primary routing table.

7.245. destination

elementId: TBD
name: destination
dataType: ipAddress
status: current
description: The destination IP address
prefix of the routing table entry.

7.246. gateway

elementId: TBD
name: gateway
dataType: ipAddress
status: current
description: The gateway of the specified
routing table entry.

7.247. runlevelInfo

elementId: TBD
name: runlevelInfo
dataType: list
structure: list (serviceName, runlevel, start, kill)

status: current
description: Information about the start or kill state of a specified service at a given runlevel.

7.248. runlevel

elementId: TBD
name: runlevel
dataType: string
status: current
description: Specifies the system runlevel associated with a service.

7.249. start

elementId: TBD
name: start
dataType: boolean
status: current
description: Specifies whether the service is scheduled to start at the runlevel.

7.250. kill

elementId: TBD
name: kill
dataType: boolean
status: current
description: Specifies whether the service is scheduled to be killed at the runlevel.

7.251. shadowItem

elementId: TBD
name: shadowItem
dataType: list
structure: list (username, password, chgLst, chgAllow, chgReq, expWarn, expInact, expDate, flags, encryptMethod)
status: current
description:

7.252. chgLst

elementId: TBD
name: chgLst
dataType: dateTimeSeconds
status: current
description: The date of the last password change.

7.253. chgAllow

elementId: TBD
name: chgAllow
dataType: unsigned32
status: current
description: Specifies how often in days a user may change their password. It can also be thought of as the minimum age of a password.

7.254. chgReq

elementId: TBD
name: chgReq
dataType: unsigned32
status: current
description: Describes how long a user can keep a password before the system forces her to change it.

7.255. expWarn

elementId: TBD
name: expWarn
dataType: unsigned32
status: current
description: Describes how long before password expiration the system begins warning the user.

7.256. expInact

elementId: TBD
name: expInact
dataType: unsigned32
status: current
description: Describes how many days of account inactivity the system will wait after a password expires before locking the account.

7.257. expDate

elementId: TBD
name: expDate
dataType: dateTimeSeconds
status: current
description: Specifies when will the
account's password expire.

7.258. encryptMethod

elementId: TBD
name: encryptMethod
dataType: enumeration
structure: DES ; 0x1 ; The DES method corresponds to the (none)
prefix.
 BSDi ; 0x2 ; The BSDi method corresponds to BSDi modified
 DES or the '_' prefix.
 MD5 ; 0x3 ; The MD5 method corresponds to MD5 for Linux/BSD
 or the \$1\$ prefix.
 Blowfish ; 0x4 ; The Blowfish method corresponds to Blowfish
 (OpenBSD) or the \$2\$ or \$2a\$ prefixes.
 Sun MD5 ; 0x5 ; The Sun MD5 method corresponds to the \$md5\$
 prefix.
 SHA-256 ; 0x6 ; The SHA-256 method corresponds to the \$5\$
 prefix.
 SHA-512 ; 0x7 ; The SHA-512 method corresponds to the \$6\$
 prefix. ; 0x8 ; The empty string value is permitted here to
 allow for empty elements associated with variable references.
status: current
description: Describes method that is used for hashing
passwords.

7.259. symlink

elementId: TBD
name: symlink
dataType: list
structure: list (symlinkFilepath, canonicalPath)
status: current

description: Identifies the result generated for a symlink.

7.260. symlinkFilepath

elementId: TBD
name: symlinkFilepath
dataType: string
status: current
description: Specifies the filepath to
the subject symbolic link file.

7.261. canonicalPath

elementId: TBD
name: canonicalPath
dataType: string
status: current
description: Specifies the canonical
path for the target of the symbolic link file specified by
the filepath.

7.262. sysctl

elementId: TBD
name: sysctl
dataType: list
structure: list (kernelParameterName, kernelParameterValue+,
uname, machineClass, nodeName, osName, osRelease,
osVersion, processorType)
status: current
description: Stores
information retrieved from the local system about a kernel
parameter and its respective value(s).

7.263. kernelParameterName

elementId: TBD
name: kernelParameterName
dataType: string
status: current
description: The name of a kernel
parameter that was collected from the local system.

7.264. kernelParameterValue

elementId: TBD
name: kernelParameterValue
dataType: string
status: current
description: The current value(s)
for the specified kernel parameter on the local system.

7.265. uname

elementId: TBD
name: uname
dataType: list
structure: list (machineClass, nodeName, osName, osRelease,
osVersion, processorType)
status: current
description: Information about the hardware the machine is running
on.

7.266. machineClass

elementId: TBD
name: machineClass
dataType: string
status: current
description: Specifies the machine
hardware name.

7.267. nodeName

elementId: TBD
name: nodeName
dataType: string
status: current
description: Specifies the host
name.

7.268. osName

elementId: TBD
name: osName
dataType: string
status: current
description: Specifies the operating system
name.

7.269. osRelease

elementId: TBD
name: osRelease
dataType: string
status: current
description: Specifies the build
version.

7.270. processorType

elementId: TBD
name: processorType
dataType: string
status: current
description: Specifies the processor
 type.

7.271. internetService

elementId: TBD
name: internetService
dataType: list
structure: list (serviceProtocol, serviceName, flags,
 noAccess, onlyFrom, port, server, serverArguments,
 socketType, registeredServiceType, user, wait, disabled)

status: current
description: Holds information associated with Internet services.

7.272. serviceProtocol

elementId: TBD
name: serviceProtocol
dataType: string
status: current
description: Specifies the protocol
 that is used by the service.

7.273. serviceName

elementId: TBD
name: serviceName
dataType: string
status: current
description: Specifies the name of the
 service.

7.274. flags

elementId: TBD
name: flags
dataType: string
status: current
description: Specifies miscellaneous settings
 associated with the service with executing a program.

7.275. noAccess

elementId: TBD
name: noAccess
dataType: string
status: current
description: Specifies the remote hosts to
which the service is unavailable.

7.276. onlyFrom

elementId: TBD
name: onlyFrom
dataType: ipAddress
status: current
description: Specifies the remote hosts to
which the service is available.

7.277. port

elementId: TBD
name: port
dataType: unsigned32
status: current
description: The port entity specifies the port
used by the service.

7.278. server

elementId: TBD
name: server
dataType: string
status: current
description: Specifies the executable that is
used to launch the service.

7.279. serverArguments

elementId: TBD
name: serverArguments
dataType: string
status: current
description: Specifies the arguments
that are passed to the executable when launching the service.

7.280. socketType

elementId: TBD
name: socketType
dataType: string
status: current
description: Specifies the type of socket
that is used by the service. Possible values include: stream,
dgram, raw, or seqpacket.

7.281. registeredServiceType

elementId: TBD
name: registeredServiceType
dataType: enumeration
structure: INTERNAL ; 0x1 ; The INTERNAL type is used to describe
services like echo, chargen, and others whose functionality is
supplied by xinetd itself.
RPC ; 0x2 ; The RPC type is used to describe services that
use remote procedure call ala NFS.
UNLISTED ; 0x3 ; The UNLISTED type is used to describe
services that aren't listed in /etc/protocols or /etc/rpc.
TCPMUX ; 0x4 ; The TCPMUX type is used to describe services
that conform to RFC 1078. This type indicates that the service
is responsible for handling the protocol handshake.
TCPMUXPLUS ; 0x5 ; The TCPMUXPLUS type is used to describe
services that conform to RFC 1078. This type indicates that
xinetd is responsible for handling the protocol
handshake.
; 0x6 ; The empty string value is permitted here to allow
for detailed error reporting.
status: current

description: Specifies the type of internet service.

7.282. wait

elementId: TBD
name: wait
dataType: boolean
status: current
description: Specifies whether or not the service is single-threaded
or multi-threaded and whether or not xinetd accepts the connection
or the service accepts the connection. A value of 'true' indicates
that the service is single-threaded and the service will accept the
connection. A value of 'false' indicates that the service is multi-
threaded and xinetd will accept the connection.

7.283. disabled

elementId: TBD
name: disabled
dataType: boolean
status: current
description: Specifies whether or not the service is disabled. A value of 'true' indicates that the service is disabled and will not start. A value of 'false' indicates that the service is not disabled.

7.284. windowsView

elementId: TBD
name: windowsView
dataType: enumeration
structure: 32_bit ; 0x1 ; Indicates the 32_bit windows view.
64_bit ; 0x2 ; Indicates the 64_bit windows view.
; 0x3 ; The empty string value is permitted here to allow for empty elements associated with error conditions.
status: current
description: Indicates from which view (32-bit or 64-bit), the information was collected. A value of '32_bit' indicates the Item was collected from the 32-bit view. A value of '64-bit' indicates the Item was collected from the 64-bit view.

7.285. fileauditedpermissions

elementId: TBD
name: fileauditedpermissions
dataType: list
structure: list (filepath, path, filename, trusteeSid, trusteeName, auditStandardDelete, auditStandardReadControl, auditStandardWriteDac, auditStandardWriteOwner, auditStandardSynchronize, auditAccessSystemSecurity, auditGenericRead, auditGenericWrite, auditGenericExecute, auditGenericAll, auditFileReadData, auditFileWriteData, auditFileAppendData, auditFileReadEa, auditFileWriteEa, auditFileExecute, auditFileDeleteChild, auditFileReadAttributes, auditFileWriteAttributes, windowsView)
status: current
description: Stores the audited access rights of a file that a system access control list (SACL) structure grants to a specified trustee. The trustee's audited access rights are determined checking all access control entries (ACEs) in the SACL.

7.286. trusteeName

elementId: TBD
name: trusteeName
dataType: string
status: current
description: Specifies the trustee name. A trustee can be a user, group, or program (such as a Windows service).

7.287. auditStandardDelete

elementId: TBD
name: auditStandardDelete
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: The right to delete the object.

7.288. auditStandardReadControl

elementId: TBD
name: auditStandardReadControl
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: The right to read the information in the object's security descriptor, not including the information in the SACL.

7.289. auditStandardWriteDac

elementId: TBD
name: auditStandardWriteDac
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: The right to modify the DACL in the object's security descriptor.

7.290. auditStandardWriteOwner

elementId: TBD
name: auditStandardWriteOwner
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: The right to change the owner in the object's security descriptor.

7.291. auditStandardSynchronize

elementId: TBD
name: auditStandardSynchronize
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: The right to use the object for synchronization. This enables a thread to wait until the object is in the signaled state. Some object types do not support this access right.

7.292. auditAccessSystemSecurity

elementId: TBD
name: auditAccessSystemSecurity
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Indicates access to a system access control list (SACL).

7.293. auditGenericRead

elementId: TBD
name: auditGenericRead
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Read access.

7.294. auditGenericWrite

elementId: TBD
name: auditGenericWrite
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Write access.

7.295. auditGenericExecute

elementId: TBD
name: auditGenericExecute
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Execute access.

7.296. auditGenericAll

elementId: TBD
name: auditGenericAll
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Read, write, and execute access.

7.297. auditFileReadData

elementId: TBD
name: auditFileReadData
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Grants the right to read data from the file.

7.298. auditFileWriteData

elementId: TBD
name: auditFileWriteData
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Grants the right to write data to the file.

7.299. auditFileAppendData

elementId: TBD
name: auditFileAppendData
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Grants the right to append data to the file.

7.300. auditFileReadEa

elementId: TBD
name: auditFileReadEa
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Grants the right to read extended attributes.

7.301. auditFileWriteEa

elementId: TBD
name: auditFileWriteEa
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Grants the right to write extended attributes.

7.302. auditFileExecute

elementId: TBD
name: auditFileExecute
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Grants the right to execute a file.

7.303. auditFileDeleteChild

elementId: TBD
name: auditFileDeleteChild
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Right to delete a directory and all the files it contains (its children), even if the files are read-only.

7.304. auditFileReadAttributes

elementId: TBD
name: auditFileReadAttributes
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Grants the right to read file attributes.

7.305. auditFileWriteAttributes

elementId: TBD
name: auditFileWriteAttributes
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: Grants the right to change file attributes.

7.306. fileeffectiverights

elementId: TBD
name: fileeffectiverights
dataType: list
structure: list (filepath, path, filename,
trusteeSid, trusteeName, standardDelete, standardReadControl,
standardWriteDac, standardWriteOwner,
standardSynchronize, accessSystemSecurity, genericRead,
genericWrite, genericExecute, genericAll, fileReadData,
fileWriteData, fileAppendData, fileReadEa, fileWriteEa,
fileExecute, fileDeleteChild, fileReadAttributes,
fileWriteAttributes, windowsView)
status: current
description: Stores the effective rights of a file that a
discretionary access control list (DACL) structure grants
to a specified trustee. The trustee's effective rights
are determined checking all access-allowed and access-denied
access control entries (ACEs) in the DACL.

7.307. standardDelete

elementId: TBD
name: standardDelete
dataType: boolean
status: current
description: The right to delete the
object.

7.308. standardReadControl

elementId: TBD
name: standardReadControl
dataType: boolean
status: current
description: The right to read
the information in the object's security descriptor, not
including the information in the SACL.

7.309. standardWriteDac

elementId: TBD
name: standardWriteDac
dataType: boolean
status: current
description: The right to modify the
DACL in the object's security descriptor.

7.310. standardWriteOwner

elementId: TBD
name: standardWriteOwner
dataType: boolean
status: current
description: The right to change
the owner in the object's security descriptor.

7.311. standardSynchronize

elementId: TBD
name: standardSynchronize
dataType: boolean
status: current
description: The right to use the
object for synchronization. This enables a thread to wait
until the object is in the signaled state. Some object
types do not support this access right.

7.312. accessSystemSecurity

elementId: TBD
name: accessSystemSecurity
dataType: boolean
status: current
description: Indicates access to
a system access control list (SACL).

7.313. genericRead

elementId: TBD
name: genericRead
dataType: boolean
status: current
description: Read access.

7.314. genericWrite

elementId: TBD
name: genericWrite
dataType: boolean
status: current
description: Write access.

7.315. genericExecute

elementId: TBD
name: genericExecute
dataType: boolean
status: current
description: Execute access.

7.316. genericAll

elementId: TBD
name: genericAll
dataType: boolean
status: current
description: Read, write, and execute
access.

7.317. fileReadData

elementId: TBD
name: fileReadData
dataType: boolean
status: current
description: Grants the right to read
data from the file

7.318. fileWriteData

elementId: TBD
name: fileWriteData
dataType: boolean
status: current
description: Grants the right to write
data to the file.

7.319. fileAppendData

elementId: TBD
name: fileAppendData
dataType: boolean
status: current
description: Grants the right to
append data to the file.

7.320. fileReadEa

elementId: TBD
name: fileReadEa
dataType: boolean
status: current
description: Grants the right to read
extended attributes.

7.321. fileWriteEa

elementId: TBD
name: fileWriteEa
dataType: boolean
status: current
description: Grants the right to write
extended attributes.

7.322. fileExecute

elementId: TBD
name: fileExecute
dataType: boolean
status: current
description: Grants the right to execute
a file.

7.323. fileDeleteChild

elementId: TBD
name: fileDeleteChild
dataType: boolean
status: current
description: Right to delete a
directory and all the files it contains (its children),
even if the files are read-only.

7.324. fileReadAttributes

elementId: TBD
name: fileReadAttributes
dataType: boolean
status: current
description: Grants the right to
read file attributes.

7.325. fileWriteAttributes

elementId: TBD
name: fileWriteAttributes
dataType: boolean
status: current
description: Grants the right to
change file attributes.

7.326. groupInfo

elementId: TBD
name: groupInfo
dataType: list
structure: list (group, username, subgroup)
status: current
description: Specifies the different users and subgroups, that
directly belong to specific groups.

7.327. group

elementId: TBD
name: group
dataType: string
status: current
description: Represents the name of a particular
group.

7.328. subgroup

elementId: TBD
name: subgroup
dataType: string
status: current
description: Represents the name of a
particular subgroup in the specified group.

7.329. groupSidInfo

elementId: TBD
name: groupSidInfo
dataType: list
structure: list (groupSid, userSid, subgroupSid)
status: current
description: Specifies the different users and subgroups, that
directly belong to specific groups
(identified by SID).

7.330. userSidInfo

elementId: TBD
name: userSidInfo
dataType: list
structure: list (userSid, enabled, groupSid, lastLogon)

status: current
description: Specifies the different groups (identified by SID)
that a user belongs to.

7.331. userSid

elementId: TBD
name: userSid
dataType: string
status: current
description: Represents the SID of a
particular user.

7.332. subgroupSid

elementId: TBD
name: subgroupSid
dataType: string
status: current
description: Represents the SID of a
particular subgroup.

7.333. lockoutpolicy

elementId: TBD
name: lockoutpolicy
dataType: list
structure: list (forceLogoff, lockoutDuration,
lockoutObservationWindow, lockoutThreshold)
status: current
description: Specifies various attributes associated
with lockout information for users and global groups in the
security database.

7.334. forceLogoff

elementId: TBD
name: forceLogoff
dataType: unsigned32
status: current
description: Specifies, in seconds, the
amount of time between the end of the valid logon time and
the time when the user is forced to log off the
network.

7.335. lockoutDuration

elementId: TBD
name: lockoutDuration
dataType: unsigned32
status: current
description: Specifies, in seconds,
how long a locked account remains locked before it is
automatically unlocked.

7.336. lockoutObservationWindow

elementId: TBD
name: lockoutObservationWindow
dataType: unsigned32
status: current
description: Specifies the
maximum time, in seconds, that can elapse between any two
failed logon attempts before lockout occurs.

7.337. lockoutThreshold

elementId: TBD
name: lockoutThreshold
dataType: unsigned32
status: current
description: Specifies the number of
invalid password authentications that can occur before an
account is marked "locked out."

7.338. passwordpolicy

elementId: TBD
name: passwordpolicy
dataType: list
structure: list (maxPasswdAge, minPasswdAge,
minPasswdLen, passwordHistLen, passwordComplexity,
reversibleEncryption)
status: current
description: Specifies
policy information associated with passwords.

7.339. maxPasswdAge

elementId: TBD
name: maxPasswdAge
dataType: unsigned32
status: current
description: Specifies, in seconds (from
a DWORD), the maximum allowable password age. A value of
TIMEQ_FOREVER (max DWORD value, 4294967295) indicates
that the password never expires. The minimum valid value
for this element is ONE_DAY (86400). See the
USER_MODAL_INFO_0 structure returned by a call to
NetUserModalsGet().

7.340. minPasswdAge

elementId: TBD
name: minPasswdAge
dataType: unsigned32
status: current
description: Specifies the minimum
number of seconds that can elapse between the time a password
changes and when it can be changed again. A value of
zero indicates that no delay is required between password
updates.

7.341. minPasswdLen

elementId: TBD
name: minPasswdLen
dataType: unsigned32
status: current
description: Specifies the minimum
allowable password length. Valid values for this element are
zero through PWLEN.

7.342. passwordHistLen

elementId: TBD
name: passwordHistLen
dataType: unsigned32
status: current
description: Specifies the length of
password history maintained. A new password cannot match any
of the previous usrmod0_password_hist_len passwords.
Valid values for this element are zero through DEF_MAX_PWHIST.

7.343. passwordComplexity

elementId: TBD
name: passwordComplexity
dataType: boolean
status: current
description: Indicates whether
passwords must meet the complexity requirements put forth
by the operating system.

7.344. reversibleEncryption

elementId: TBD
name: reversibleEncryption
dataType: boolean
status: current
description: Indicates whether
or not passwords are stored using reversible encryption.

7.345. portInfo

elementId: TBD
name: portInfo
dataType: list
structure: list (localAddress, localPort, transportProtocol,
pid, foreignAddress, foreignPort)
status: current
description: Information about open listening ports.

7.346. foreignPort

elementId: TBD
name: foreignPort
dataType: string
status: current
description: The TCP or UDP port to which
the program communicates.

7.347. printereffectiverights

elementId: TBD
name: printereffectiverights
dataType: list
structure: list (printerName, trusteeSid,
standardDelete, standardReadControl, standardWriteDac,
standardWriteOwner, standardSynchronize,
accessSystemSecurity, genericRead, genericWrite,
genericExecute, genericAll, printerAccessAdminister,
printerAccessUse, jobAccessAdminister, jobAccessRead)
status: current
description: Stores the effective rights of a printer that a
discretionary access control list (DACL) structure grants to a
specified trustee. The trustee's effective rights are determined
checking all access-allowed and access-denied access control
entries (ACEs) in the DACL.

7.348. printerName

elementId: TBD
name: printerName
dataType: string
status: current
description: Specifies the name of the
printer.

7.349. printerAccessAdminister

elementId: TBD
name: printerAccessAdminister
dataType: boolean
status: current
description:

7.350. printerAccessUse

elementId: TBD
name: printerAccessUse
dataType: boolean
status: current
description:

7.351. jobAccessAdminister

elementId: TBD
name: jobAccessAdminister
dataType: boolean
status: current
description:

7.352. jobAccessRead

elementId: TBD
name: jobAccessRead
dataType: boolean
status: current
description:

7.353. registry

elementId: TBD
name: registry
dataType: list
structure: list (registryHive, registryKey, registryKeyName,
lastWriteTime, registryKeyType, registryKeyValue,
windowsView)
status: current
description: Specifies information that can be
collected about a particular registry key.

7.354. registryHive

elementId: TBD
name: registryHive
dataType: enumeration
structure: HKEY_CLASSES_ROOT ; 0x1 ; This registry subtree contains information that associates file types with programs and configuration data for automation (e.g. COM objects and Visual Basic Programs).
HKEY_CURRENT_CONFIG ; 0x2 ; This registry subtree contains configuration data for the current hardware profile.
HKEY_CURRENT_USER ; 0x3 ; This registry subtree contains the user profile of the user that is currently logged into the system.
HKEY_LOCAL_MACHINE ; 0x4 ; This registry subtree contains information about the local system.
HKEY_USERS ; 0x5 ; This registry subtree contains user-specific data.
; 0x6 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description: The
hive that the registry key belongs to.

7.355. registryKey

elementId: TBD
name: registryKey
dataType: string
status: current
description: Describes the registry key.
Note that the hive portion of the string should not be included, as this data can be found under the hive element.

7.356. registryKeyName

elementId: TBD
name: registryKeyName
dataType: string
status: current
description: Describes the name of a
registry key.

7.357. lastWriteTime

elementId: TBD
name: lastWriteTime
dataType: unsigned64
status: current
description: The last time that the key or any of its value entries were modified. The value of this entity represents the FILETIME structure which is a 64-bit value representing the number of 100-nanosecond intervals since January 1, 1601 (UTC). Last write time can be queried on any key, with hives being classified as a type of key. When collecting only information about a registry hive or key the last write time will be the time the key or any of its entries were modified. When collecting only information about a registry name the last write time will be the time the containing key was modified. Thus when collecting information about a registry name, the last write time does not correlate directly to the specified name. See the RegQueryInfoKey function lpftLastWriteTime.

7.358. registryKeyType

elementId: TBD
name: registryKeyType
dataType: enumeration
structure: reg_binary ; 0x1 ; The reg_binary type is used by registry keys that specify binary data in any form.
reg_dword ; 0x2 ; The reg_dword type is used by registry keys that specify an unsigned 32-bit integer.
reg_dword_little_endian ; 0x3 ; The reg_dword_little_endian type is used by registry keys that specify an unsigned 32-bit little-endian integer. It is designed to run on little-endian computer architectures.
reg_dword_big_endian ; 0x4 ; The reg_dword_big_endian type is used by registry keys that specify an unsigned 32-bit big-endian integer. It is designed to run on big-endian computer architectures.
reg_expand_sz ; 0x5 ; The reg_expand_sz type is used by registry keys to specify a null-terminated string that contains unexpanded references to environment variables (for example, "%PATH%").
reg_link ; 0x6 ; The reg_link type is used by the registry keys for null-terminated unicode strings. It is related to target path of a symbolic link created by the RegCreateKeyEx function.
reg_multi_sz ; 0x7 ; The reg_multi_sz type is used by registry keys that specify an array of null-terminated strings, terminated by two null characters.

reg_none; 0x8 ;
The reg_none type is used by registry keys that have no defined value type.
reg_qword; 0x9 ; The reg_qword type is used by registry keys that specify an unsigned 64-bit integer.
reg_qword_little_endian; 0xA ; The reg_qword_little_endian type is used by registry keys that specify an unsigned 64-bit integer in little-endian computer architectures.
reg_sz; 0xB ; The reg_sz type is used by registry keys that specify a single null-terminated string.
reg_resource_list; 0xC ; The reg_resource_list type is used by registry keys that specify a resource list.
reg_full_resource_descriptor; 0xD ; The reg_full_resource_descriptor type is used by registry keys that specify a full resource descriptor.
reg_resource_requirements_list; 0xE ; The reg_resource_requirements_list type is used by registry keys that specify a resource requirements list.
; 0xF ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description:
Specifies the type of data stored by the registry key.

7.359. registryKeyValue

elementId: TBD
name: registryKeyValue
dataType: string
status: current
description: Holds the actual value of the specified registry key. The representation of the value as well as the associated datatype attribute depends on type of data stored in the registry key. If the value being tested is of type REG_BINARY, then the datatype attribute should be set to 'binary' and the data represented by the value entity should follow the xsd:hexBinary form. (each binary octet is encoded as two hex digits) If the value being tested is of type REG_DWORD, REG_QWORD, REG_DWORD_LITTLE_ENDIAN, REG_DWORD_BIG_ENDIAN, or REG_QWORD_LITTLE_ENDIAN then the datatype attribute should be set to 'int' and the value entity should represent the data as an unsigned integer. DWORD and QWORD values represent unsigned 32-bit and 64-bit integers, respectively. If the value being tested is of type REG_EXPAND_SZ, then the datatype attribute should be set to 'string' and the pre-expanded string should be represented by the value entity. If the value being tested is of type REG_MULTI_SZ, then only a single string (one of the multiple strings) should be tested using the value entity with the datatype attribute set to 'string'. In order to test multiple values, multiple OVAL registry tests should be used. If the specified registry key is of type REG_SZ, then the datatype should be 'string' and the value entity should be a copy of the string. If the value being tested is of type REG_LINK, then the datatype attribute should be set to 'string' and the null-terminated Unicode string should be represented by the value entity.

7.360. regkeyauditedpermissions

elementId: TBD
name: regkeyauditedpermissions
dataType: list
structure: list (registryKey, trusteeSid, trusteeName,
standardDelete, standardReadControl, standardWriteDac,
standardWriteOwner, standardSynchronize,
accessSystemSecurity, genericRead, genericWrite,
genericExecute, genericAll, keyQueryValue, keySetValue,
keyCreateSubKey, keyEnumerateSubKeys, keyNotify,
keyCreateLink, keyWow6464Key, keyWow6432Key, keyWow64Res,
windowsView)
status: current
description: Stores the audited access rights of a registry key
that a system access control list (SACL) structure grants to a
specified trustee. The trustee's audited access rights are
determined checking all access control entries (ACEs) in the SACL.

7.361. auditKeyQueryValue

elementId: TBD
name: auditKeyQueryValue
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is
used to perform audits on all unsuccessful occurrences of
specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel
all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to
perform audits on all successful occurrences of the specified
events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE
is used to perform audits on all successful and unsuccessful
occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for
detailed error reporting.
status: current
description:

7.362. auditKeySetValue

elementId: TBD
name: auditKeySetValue
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description:

7.363. auditKeyCreateSubKey

elementId: TBD
name: auditKeyCreateSubKey
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description:

7.364. auditKeyEnumerateSubKeys

elementId: TBD
name: auditKeyEnumerateSubKeys
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description:

7.365. auditKeyNotify

elementId: TBD
name: auditKeyNotify
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description:

7.366. auditKeyCreateLink

elementId: TBD
name: auditKeyCreateLink
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description:

7.367. auditKeyWow6464Key

elementId: TBD
name: auditKeyWow6464Key
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description:

7.368. auditKeyWow6432Key

elementId: TBD
name: auditKeyWow6432Key
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description:

7.369. auditKeyWow64Res

elementId: TBD
name: auditKeyWow64Res
dataType: enumeration
structure: AUDIT_FAILURE ; 0x1 ; The audit type AUDIT_FAILURE is used to perform audits on all unsuccessful occurrences of specified events when auditing is enabled.
AUDIT_NONE ; 0x2 ; The audit type AUDIT_NONE is used to cancel all auditing options for the specified events.
AUDIT_SUCCESS ; 0x3 ; The audit type AUDIT_SUCCESS is used to perform audits on all successful occurrences of the specified events when auditing is enabled.
AUDIT_SUCCESS_FAILURE ; 0x4 ; The audit type AUDIT_SUCCESS_FAILURE is used to perform audits on all successful and unsuccessful occurrences of the specified events when auditing is enabled.
; 0x5 ; The empty string value is permitted here to allow for detailed error reporting.
status: current
description:

7.370. regkeyeffectiverights

elementId: TBD
name: regkeyeffectiverights
dataType: list
structure: list (registryHive, registryKey, trusteeSid, trusteeName, standardDelete, standardReadControl, standardWriteDac, standardWriteOwner, standardSynchronize, accessSystemSecurity, genericRead, genericWrite, genericExecute, genericAll, keyQueryValue, keySetValue, keyCreateSubKey, keyEnumerateSubKeys, keyNotify, keyCreateLink, keyWow6464Key, keyWow6432Key, keyWow64Res, windowsView)
status: current
description: Stores the effective rights of a registry key that a discretionary access control list (DACL) structure grants to a specified trustee. The trustee's effective rights are determined checking all access-allowed and access-denied access control entries (ACEs) in the DACL.

7.371. keyQueryValue

elementId: TBD
name: keyQueryValue
dataType: boolean
status: current
description: Specifies whether or not
permission is granted to query the key's value.

7.372. keySetValue

elementId: TBD
name: keySetValue
dataType: boolean
status: current
description: Specifies whether or not
permission is granted to set the key's value.

7.373. keyCreateSubKey

elementId: TBD
name: keyCreateSubKey
dataType: boolean
status: current
description: Specifies whether or not
permission is granted to create a subkey.

7.374. keyEnumerateSubKeys

elementId: TBD
name: keyEnumerateSubKeys
dataType: boolean
status: current
description: Specifies whether or
not permission is granted to list the subkeys associated
with key.

7.375. keyNotify

elementId: TBD
name: keyNotify
dataType: boolean
status: current
description:

7.376. keyCreateLink

elementId: TBD
name: keyCreateLink
dataType: boolean
status: current
description:

7.377. keyWow6464Key

elementId: TBD
name: keyWow6464Key
dataType: boolean
status: current
description:

7.378. keyWow6432Key

elementId: TBD
name: keyWow6432Key
dataType: boolean
status: current
description:

7.379. keyWow64Res

elementId: TBD
name: keyWow64Res
dataType: boolean
status: current
description:

7.380. service

elementId: TBD
name: service
dataType: list
structure: list (serviceName, displayName, description,
 serviceType, startType, currentState, controlsAccepted,
 startName, path, pid, serviceFlag, dependencies)
status: current
description: Stores information about Windows services that are
 present on the system.

7.381. displayName

elementId: TBD
name: displayName
dataType: string
status: current
description: Specifies the name of the
 service as specified in administrative tools.

7.382. description

elementId: TBD
name: description
dataType: string
status: current
description: Specifies the description of
 the service.

7.383. serviceType

elementId: TBD
name: serviceType
dataType: enumeration
structure: SERVICE_FILE_SYSTEM_DRIVER ; 0x1 ; The
SERVICE_FILE_SYSTEM_DRIVER type means that the service is
a file system driver. The DWORD value that this
corresponds to is 0x00000002.
SERVICE_KERNEL_DRIVER ; 0x2 ; The SERVICE_KERNEL_DRIVER type
means that the service is a driver. The DWORD value that
this corresponds to is 0x00000001.
SERVICE_WIN32_OWN_PROCESS ; 0x3 ; The SERVICE_WIN32_OWN_PROCESS
type means that the service runs in its own process. The DWORD
value that this corresponds to is 0x00000010.
SERVICE_WIN32_SHARE_PROCESS ; 0x4 ; The
SERVICE_WIN32_SHARE_PROCESS type means that the service runs
in a process with other services. The DWORD value that this
corresponds to is 0x00000020.
SERVICE_INTERACTIVE_PROCESS ; 0x5 ; The
SERVICE_WIN32_SHARE_PROCESS type means that the service runs
in a process with other services. The DWORD value that this
corresponds to is 0x00000100.
; 0x6 ; The empty string value is permitted here to allow for
empty elements associated with error conditions.
status: current
description:
Specifies the type of the service.

7.384. startType

elementId: TBD
name: startType
dataType: enumeration
structure: SERVICE_AUTO_START ; 0x1 ; The SERVICE_AUTO_START type means that the service is started automatically by the Service Control Manager (SCM) during startup. The DWORD value that this corresponds to is 0x00000002.
SERVICE_BOOT_START ; 0x2 ; The SERVICE_BOOT_START type means that the driver service is started by the system loader. The DWORD value that this corresponds to is 0x00000000.
SERVICE_DEMAND_START ; 0x3 ; The SERVICE_DEMAND_START type means that the service is started by the Service Control Manager (SCM) when StartService() is called. The DWORD value that this corresponds to is 0x00000003.
SERVICE_DISABLED ; 0x4 ; The SERVICE_DISABLED type means that the service cannot be started. The DWORD value that this corresponds to is 0x00000004.
SERVICE_SYSTEM_START ; 0x5 ; The SERVICE_SYSTEM_START type means that the service is a device driver started by IoInitSystem(). The DWORD value that this corresponds to is 0x00000001.
; 0x6 ; The empty string value is permitted here to allow for empty elements associated with error conditions.
status: current
description: Specifies when the service should be started.

7.385. currentState

elementId: TBD
name: currentState
dataType: enumeration
structure: SERVICE_CONTINUE_PENDING ; 0x1 ; The
SERVICE_CONTINUE_PENDING type means that the service has been
sent a command to continue, however, the command has
not yet been executed. The DWORD value that this corresponds
to is 0x00000005. SERVICE_PAUSE_PENDING ; 0x2 ; The
SERVICE_PAUSE_PENDING type means that the service has been
sent a command to pause, however, the command has not
yet been executed. The DWORD value that this corresponds to
is 0x00000006.
SERVICE_PAUSED ; 0x3 ; The SERVICE_PAUSED type means that
the service is paused. The DWORD value that this corresponds
to is 0x00000007.
SERVICE_RUNNING ; 0x4 ; The SERVICE_RUNNING type means that
the service is running. The DWORD value that this
corresponds to is 0x00000004.
SERVICE_START_PENDING ; 0x5 ; The SERVICE_START_PENDING type
means that the service has been sent a command to start,
however, the command has not yet been executed. The DWORD
value that this corresponds to is 0x00000002.
SERVICE_STOP_PENDING ; 0x6 ; The SERVICE_STOP_PENDING type
means that the service
has been sent a command to stop, however, the command has
not yet been executed. The DWORD value that this
corresponds to is 0x00000003.
SERVICE_STOPPED ; 0x7 ; The SERVICE_STOPPED type means that
the service is stopped. The DWORD value that this corresponds
to is 0x00000001.
; 0x8 ; The empty string value is permitted here to allow
for empty elements associated with error conditions.
status: current
description: Specifies the current state of
the service.

7.386. controlsAccepted

elementId: TBD
name: controlsAccepted
dataType: enumeration
structure:
SERVICE_ACCEPT_NETBINDCHANGE ; 0x1 ;
The SERVICE_ACCEPT_NETBINDCHANGE type means that the
service is a network component and can accept changes in its
binding without being stopped or restarted. The DWORD value
that this corresponds to is 0x00000010.
SERVICE_ACCEPT_PARAMCHANGE ; 0x2 ; The SERVICE_ACCEPT_PARAMCHANGE

type means that the service can re-read its startup parameters without being stopped or restarted. The DWORD value that this corresponds to is 0x00000008.

SERVICE_ACCEPT_PAUSE_CONTINUE ; 0x3 ; The SERVICE_ACCEPT_PAUSE_CONTINUE type means that the service can be paused or continued. The DWORD value that this corresponds to is 0x00000002.

SERVICE_ACCEPT_PRESHUTDOWN ; 0x4 ; The SERVICE_ACCEPT_PRESHUTDOWN type means that the service can receive pre-shutdown notifications. The DWORD value that this corresponds to is 0x00000100.

SERVICE_ACCEPT_SHUTDOWN ; 0x5 ; The SERVICE_ACCEPT_SHUTDOWN type means that the service can receive shutdown notifications. The DWORD value that this corresponds to is 0x00000004.

SERVICE_ACCEPT_STOP ; 0x6 ; The SERVICE_ACCEPT_STOP type means that the service can be stopped. The DWORD value that this corresponds to is 0x00000001.

SERVICE_ACCEPT_HARDWAREPROFILECHANGE ; 0x7 ; The SERVICE_ACCEPT_HARDWAREPROFILECHANGE type means that the service can receive notifications when the system's hardware profile changes. The DWORD value that this corresponds to is 0x00000020.

SERVICE_ACCEPT_POWEREVENT ; 0x8 ; The SERVICE_ACCEPT_POWEREVENT type means that the service can receive notifications when the system's power status has changed. The DWORD value that this corresponds to is 0x00000040.

SERVICE_ACCEPT_SESSIONCHANGE ; 0x9 ; The SERVICE_ACCEPT_SESSIONCHANGE type means that the service can receive notifications when the system's session status has changed. The DWORD value that this corresponds to is 0x00000080.

SERVICE_ACCEPT_TIMECHANGE ; 0xA ; The SERVICE_ACCEPT_TIMECHANGE type means that the service can receive notifications when the system time changes. The DWORD value that this corresponds to is 0x00000200.

SERVICE_ACCEPT_TRIGGEREVENT ; 0xB ; The SERVICE_ACCEPT_TRIGGEREVENT type means that the service can receive notifications when an event that the service has registered for occurs on the system. The DWORD value that this corresponds to is 0x00000400.

; 0xC ; The empty string value is permitted here to allow for empty elements associated with error conditions.

status: current

description: Specifies the control codes that a service will accept and process.

7.387. startName

elementId: TBD
name: startName
dataType: string
status: current
description: Specifies the account under
which the process should run.

7.388. serviceFlag

elementId: TBD
name: serviceFlag
dataType: boolean
status: current
description: Specifies whether the
service is in a system process that must always run (true)
or if the service is in a non-system process or is not
running (false).

7.389. dependencies

elementId: TBD
name: dependencies
dataType: string
status: current
description: Specifies the dependencies
of this service on other services.

7.390. serviceeffectiverights

elementId: TBD
name: serviceeffectiverights
dataType: list
structure: list (serviceName, trusteeSid,
standardDelete, standardReadControl, standardWriteDac,
standardWriteOwner, genericRead, genericWrite,
genericExecute, serviceQueryConf, serviceChangeConf,
serviceQueryStat, serviceEnumDependents, serviceStart,
serviceStop, servicePause, serviceInterrogate,
serviceUserDefined)
status: current
description: Stores the
effective rights of a service that a discretionary access
control list (DACL) structure grants to a specified
trustee. The trustee's effective rights are determined by
checking all access-allowed and access-denied access
control entries (ACEs) in the DACL.

7.391. trusteeSid

elementId: TBD
name: trusteeSid
dataType: string
status: current
description: Specifies the SID that is
associated with a user, group, system, or program (such as a
Windows service).

7.392. serviceQueryConf

elementId: TBD
name: serviceQueryConf
dataType: boolean
status: current
description: Specifies whether or
not permission is granted to query the service configuration.

7.393. serviceChangeConf

elementId: TBD
name: serviceChangeConf
dataType: boolean
status: current
description: Specifies whether or
not permission is granted to change service configuration.

7.394. serviceQueryStat

elementId: TBD
name: serviceQueryStat
dataType: boolean
status: current
description: Specifies whether or
not permission is granted to query the service control
manager about the status of the service.

7.395. serviceEnumDependents

elementId: TBD
name: serviceEnumDependents
dataType: boolean
status: current
description: Specifies whether
or not permission is granted to query for an enumeration of
all the services dependent on the service.

7.396. serviceStart

elementId: TBD
name: serviceStart
dataType: boolean
status: current
description: Specifies whether or not
 permission is granted to start the service.

7.397. serviceStop

elementId: TBD
name: serviceStop
dataType: boolean
status: current
description: Specifies whether or not
 permission is granted to stop the service.

7.398. servicePause

elementId: TBD
name: servicePause
dataType: boolean
status: current
description: Specifies whether or not
 permission is granted to pause or continue the service.

7.399. serviceInterrogate

elementId: TBD
name: serviceInterrogate
dataType: boolean
status: current
description: Specifies whether or not permission is granted to
 request the service to report its status immediately.

7.400. serviceUserDefined

elementId: TBD
name: serviceUserDefined
dataType: boolean
status: current
description: Specifies whether or
 not permission is granted to specify a user-defined
 control code.

7.401. sharedresourceauditedpermissions

elementId: TBD
name: sharedresourceauditedpermissions
dataType: list
structure: list (netname, trusteeSid,
 standardDelete, standardReadControl, standardWriteDac,
 standardWriteOwner, standardSynchronize,
 accessSystemSecurity, genericRead, genericWrite,
 genericExecute, genericAll)
status: current
description: Stores
 the audited access rights of a shared resource that a system
 access control list (SACL) structure grants to a
 specified trustee. The trustee's audited access rights are
 determined checking all access control entries (ACEs)
 in the SACL.

7.402. netname

elementId: TBD
name: netname
dataType: string
status: current
description: Specifies the name associated
 with a particular shared resource.

7.403. sharedresourceeffectiverights

elementId: TBD
name: sharedresourceeffectiverights
dataType: list
structure: list (netname, trusteeSid,
 standardDelete, standardReadControl, standardWriteDac,
 standardWriteOwner, standardSynchronize,
 accessSystemSecurity, genericRead, genericWrite,
 genericExecute, genericAll)
status: current
description: Stores
 the effective rights of a shared resource that a
 discretionary access control list (DACL) structure grants
 to a specified trustee. The trustee's effective rights are
 determined checking all access-allowed and access-denied
 access control entries (ACEs) in the DACL.

7.404. user

elementId: TBD
name: user
dataType: list
structure: list (username, enabled, group, lastLogon)
status: current
description: Specifies the groups to which a user belongs.

7.405. enabled

elementId: TBD
name: enabled
dataType: boolean
status: current
description: Represents whether the
particular user is enabled or not.

7.406. lastLogon

elementId: TBD
name: lastLogon
dataType: unsigned32
status: current
description: The date and time when the
last logon occurred.

7.407. groupSid

elementId: TBD
name: groupSid
dataType: string
status: current
description: Represents the SID of a
particular group. If the specified user belongs to more than
one group, then multiple groupSid elements are
applicable. If the specified user is not a member of a single
group, then a single groupSid element should be
included with a status of 'does not exist'. If there is an
error determining the groups that the user belongs to,
then a single groupSid element should be included with a
status of 'error'.

7.408. endpointType

elementId: TBD
name: endpointType
dataType: enumeration
status: current
description: The possible types of endpoint in the enterprise.

structure:
workstation; 0x1; Workstation Endpoint
printer; 0x2; Printer Endpoint
router; 0x3; Router Endpoint
tablet; 0x4; Tablet Endpoint

7.409. endpointPurpose

elementId: TBD
name: endpointPurpose
dataType: string
status: current
description: A description of how the endpoint is used within the enterprise.
Examples include end user system,
and public web server.

7.410. endpointCriticality

elementId: TBD
name: endpointCriticality
dataType: string
status: current
description: An enterprise-defined rating which indicates the criticality of the endpoint. The rating should be specific enough to assess the impact to the overall enterprise if the endpoint is attacked or lost.

7.411. ingestTimestamp

elementId: TBD
name: ingestTimestamp
dataType: dateTimeSeconds
status: current
description: The point in time that the description of a vulnerability was received by the enterprise.

7.412. vulnerabilityVersion

elementId: TBD
name: vulnerabilityVersion
dataType: string
status: current
description: The version or iteration of the
vulnerability description information
(reported by the author, if applicable).

7.413. vulnerabilityExternalId

elementId: TBD
name: vulnerabilityExternalId
dataType: string
status: current
description: An external or third-party ID
assigned to the vulnerability
description. This could be multiple
IDs in some cases (e.g., vendor bug
ID, global ID, discoverer's local ID,
third-party vulnerability database
ID, etc.).

7.414. vulnerabilitySeverity

elementId: TBD
name: vulnerabilitySeverity
dataType: string
status: current
description: The severity of the vulnerability
(reported by the author, if applicable).

7.415. assessmentTimestamp

elementId: TBD
name: assessmentTimestamp
dataType: dateTimeSeconds
status: current
description: The point in time that the assessment
was performed against an endpoint.

7.416. vulnerableSoftware

```
elementId: TBD
name: vulnerableSoftware
dataType: list
status: current
description: A listing of software products
             installed on the endpoint which are
             known to have vulnerabilities.
structure: list(softwareInstance*)
```

7.417. endpointVulnerabilityStatus

```
elementId: TBD
name: endpointVulnerabilityStatus
dataType: enumeration
status: current
description: Overall vulnerability status of an
             enterprise endpoint.
structure: Pass; 0x1; Endpoint passed the
             vulnerability test(s).
             Fail; 0x2; Endpoint failed the
             vulnerability test(s).
```

7.418. vulnerabilityDescription

```
elementId: TBD
name: vulnerabilityDescription
dataType: string
status: current
description: A human-readable description of the
             vulnerability.
```

8. Acknowledgements

Many of the specifications in this document have been developed in a public-private partnership with vendors and end-users. The hard work of the SCAP community is appreciated in advancing these efforts to their current level of adoption.

Over the course of developing the initial draft, Brant Cheikes, Matt Hansbury, Daniel Haynes, Scott Pope, Charles Schmidt, and Steve Venema have contributed text to many sections of this document.

9. IANA Considerations

This document specifies an initial set of Information Elements for SACM in Section 7. An Internet Assigned Numbers Authority (IANA) registry will be created and populated with the Information Elements in Section 7. New assignments for SACM Information Elements will be administered by IANA through Expert Review [RFC2434]. The designated experts MUST check the requested Information Elements for completeness and accuracy of the submission with respect to the template and requirements expressed in Section 4 and Section 4.1. Requests for Information Elements that duplicate the functionality of existing Information Elements SHOULD be declined. The smallest available Information Element identifier SHOULD be assigned to a new Information Element. The definition of new Information Elements MUST be published using a well-established and persistent publication medium.

10. Security Considerations

Posture Assessments need to be performed in a safe and secure manner. In that regard, there are multiple aspects of security that apply to the communications between components as well as the capabilities themselves. This information model only contains an initial listing of items that need to be considered with respect to security and will need to be augmented as the model continues to be developed.

Security considerations include:

Authentication: Every SACM Component and asset needs to be able to identify itself and verify the identity of other SACM Components and assets.

Confidentiality: Communications between SACM Components need to be protected from eavesdropping or unauthorized collection. Some communications between SACM Components and assets may need to be protected as well.

Integrity: The information exchanged between SACM Components needs to be protected from modification. Some exchanges between assets and SACM Components will also have this requirement.

Restricted Access: Access to the information collected, evaluated, reported, and stored should only be viewable and consumable to authenticated and authorized entities.

Considerations with respect to the operational aspects of collection, evaluation, and storage security automation information can be found in Section 11.

Considerations concerning the privacy of security automation information can be found in Section 12.

11. Operational Considerations

The following sections outline a series of operational considerations for SACM deployments within an organization. This section may be expanded to include other considerations as the WG gains additional operational experience with SACM deployments and extending the information model.

11.1. Endpoint Designation

In order to successfully carry out endpoint posture assessment, it is necessary to be able to identify the endpoints on a network and track the changes to them over time. Specifically, enabling SACM Components to:

- o Tell whether two endpoint attribute assertions concern the same endpoint
- o Respond to compliance measurements, for example by reporting, remediating, and quarantining (SACM does not specify these responses, but SACM exists to enable them).

Ideally, every endpoint would be identified by a unique identifier present on the endpoint, but, this is complicated due to different factors such as the variety of endpoints on a network, the ability of tools to reliably access such an identifier, and the ability of tools to correlate disparate identifiers. As a result, it is necessary for an endpoint to be identified by a set of attributes that uniquely identify it on a network. The set of attributes that uniquely identify an endpoint on a network will likely vary by organization; however, there are a number of properties to consider when selecting identifying attributes as some are better suited for identification purposes than others.

Multiplicity: Is the attribute typically associated with a single endpoint or with multiple endpoints? If the attribute is associated with a single endpoint, it is better for identifying an endpoint on a network.

Persistence: How likely is the attribute to change? Does it never change? Does it only change when the endpoint is reprovisioned? Does it only change due to an event? Does it change on an ad-hoc and often unpredictable basis? Does it constantly change? The less likely it is for an attribute to

change over time, the better it is for identifying an endpoint on a network.

Immutability: How difficult is it to change the attribute? Is the attribute hardware rooted and never changes? Can the attribute be changed by a user/process with the appropriate access? Can the attribute be changed without controlled access. The less likely an attribute is to change over time, the better chance it will be usable to identify an endpoint over time.

Verifiable: Can the attribute be corroborated? Can the attribute be externally verified with source authentication? Can the attribute be externally verified without source authentication? Is it impossible to externally verify the attribute. Attributes that can be externally verified are more likely to be accurate and are better for identifying endpoints on a network.

With that said, requiring SACM Components and end users to constantly refer to a set of attributes to identify an endpoint, is particularly burdensome. As a result, SACM supports the concept of a target endpoint label which associates an identifier (unique to a SACM domain) with the set of attributes used by an organization to identify endpoints on a network. Once defined for an endpoint, the target endpoint label can be used in place of the set of identifying attributes.

11.2. Timestamp Accuracy

An organization will likely have different collectors deployed across the network that will be configured to collect posture attributes on varying frequencies (periodic, ad-hoc, event-driven, on endpoint, off endpoint, etc.). Some collectors will detect changes as soon as they occur whereas others will detect them at a later point during a periodic scan or when an event has triggered the collection of posture attributes. Furthermore, some changes will be detected on the endpoint and others will be observed off of the endpoint. As a result of these differences, the accuracy of the timestamp associated with the collected information will vary. For example, if a collector is only running once every 12 hours, the change probably happened at some point in time prior to the scan and the timestamp is likely not accurate. Due to this, it is important for system administrators to determine if the accuracy of a timestamp is good enough for their intended purposes.

12. Privacy Considerations

In the IETF, there are privacy concerns with respect to endpoint identity and monitoring. This is especially true when the activity on an endpoint can be linked to a particular person. For example, by correlating endpoint attributes such as usernames, certificates, etc. with browser activity, it may be possible to gain insight in to user behavior and trends beyond what is required to carry out endpoint posture assessments. In the hands of the wrong person, this information could be used to negatively influence a user's behavior or to plan attacks against the organization's infrastructure.

As a result, SACM data models should incorporate a mechanism by which an organization can designate which endpoint attributes are considered sensitive with respect to privacy. This will allow SACM Components to handle endpoint attributes in a manner consistent with the organization's privacy policies. Furthermore, organization's should put the proper mechanism in place to ensure endpoint attributes are protected when transmitted, stored, and accessed to ensure only authorized parties are granted access.

It should also be noted that some of this is often mitigated by organizational policies that require a user of an organization's network to consent to some level of monitoring in return for access to the network and other resources. The information that is monitored and collected will vary by organization and further highlights the need for a mechanism by which an organization can specify what constitutes privacy sensitive information for them.

13. References

13.1. Normative References

- [PEN] Internet Assigned Numbers Authority, "Private Enterprise Numbers", July 2016, <<https://www.iana.org/assignments/enterprise-numbers/enterprise-numbers>>.
- [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>>.

13.2. Informative References

- [I-D.ietf-sacm-requirements] Cam-Winget, N. and L. Lorenzin, "Secure Automation and Continuous Monitoring (SACM) Requirements", draft-ietf-sacm-requirements-01 (work in progress), October 2014.

[I-D.ietf-sacm-terminology]

Waltermire, D., Montville, A., Harrington, D., and N. Cam-Winget, "Terminology for Security Assessment", draft-ietf-sacm-terminology-05 (work in progress), August 2014.

[RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 2434, DOI 10.17487/RFC2434, October 1998, <<http://www.rfc-editor.org/info/rfc2434>>.

[RFC3580] Congdon, P., Aboba, B., Smith, A., Zorn, G., and J. Roese, "IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines", RFC 3580, DOI 10.17487/RFC3580, September 2003, <<http://www.rfc-editor.org/info/rfc3580>>.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<http://www.rfc-editor.org/info/rfc4949>>.

[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>>.

[RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, DOI 10.17487/RFC7012, September 2013, <<http://www.rfc-editor.org/info/rfc7012>>.

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

Appendix A. Change Log

A.1. Changes in Revision 01

Added some proposed normative text.

For provenance:

Added a class "Method"

Added the produced-using relationship between an AVP and a method

Added the produced-by relationship between a Guidance and a SACM Component

Added the hosted-by relationship between a SACM Component and an Endpoint

asserted-by and summarized-by have been renamed to produced-by.

"User" is now "Account". If a user has different credentials, SACM cannot know that they belong to the same user. But, per Kim W, many organizations do have accounts that associate credentials.

The multiplicity of the based-on relationships has been corrected.

More relationships now have labels, per UML convention.

The diagram no longer has causal arrow. They had become redundant and were nonstandard and clutter.

Renamed "credential" to "identity", following industry usage. A credential includes proof, such as a key or password. A username or a distinguished name is called an "identity".

Removed Session, because an endpoint's network activity is not SACM's initial focus

Removed Authorization, for the same reason

Added many-to-many relationship between Hardware Component and Endpoint, for clarity

Added many-to-many relationship between Software Component and Endpoint, for clarity

Added "contains" relationship between Network Interface and Network Interface

Removed relationship between Network Interface and Account. The endpoint knows the identity it used to gain network access. The PDP also knows that. But they probably do not know the account.

Added relationship between Network Interface and Identity. The endpoint and the PDP will typically know the identity.

Made identity-to-account a many-to-one relationship.

A.2. Changes in Revision 02

Added Section Identifying Attributes.

Split the figure into Figure Model of Endpoint and Figure Information Elements.

Added Figure Information Elements Take 2, proposing a triple-store model.

Some editorial cleanup

A.3. Changes in Revision 03

Moved Appendix A.1, Appendix A.2, and Mapping to SACM Use Cases into the Appendix. Added a reference to it in Section 1

Added the Section 4 section. Provided notes for the type of information we need to add in this section.

Added the Section 6 section. Moved sections on Endpoint, Hardware Component, Software Component, Hardware Instance, and Software Instance there. Provided notes for the type of information we need to add in this section.

Removed the Provenance of Information Section. SACM is not going to solve provenance rather give organizations enough information to figure it out.

Updated references to the Endpoint Security Posture Assessment: Enterprise Use Cases document to reflect that it was published as an RFC.

Fixed the formatting of a few figures.

Included references to [RFC3580] where RADIUS is mentioned.

A.4. Changes in Revision 04

Integrated the IPFIX [RFC7012] syntax into Section 4.

Converted many of the existing SACM Information Elements to the IPFIX syntax.

Included existing IPFIX Information Elements and datatypes that could likely be reused for SACM in Section 7 and Section 4 respectively.

Removed the sections related to reports as described in <https://github.com/sacmwg/draft-ietf-sacm-information-model/issues/30>.

Cleaned up other text throughout the document.

A.5. Changes in Revision 05

Merged proposed changes from the I-D IM into the WG IM (<https://github.com/sacmwg/draft-ietf-sacm-information-model/issues/41>).

Fixed some formatting warnings.

Removed a duplicate IE and added a few IE datatypes that were missing.

A.6. Changes in Revision 06

Clarified that the SACM statement and content-element subjects are conceptual and that they do not need to be explicitly defined in a data model as long as the necessary information is provided.

Updated the IPFIX syntax used to define Information Elements. There are still a couple of open issues that need to be resolved.

Updated some of the Information Elements contained in Section 7 to use the revised IPFIX syntax. The rest of the Information Elements will be converted in a later revision.

Performed various clean-up and refactoring in Sections 6 and 7. Still need to go through Section 8.

Removed appendices that were not referenced in the body of the draft. The text from them is still available in previous revisions of this document if needed.

A.7. Changes in Revision 07

Made various changes to the IPFIX syntax based on discussions at the IETF 96 Meeting. Changes included the addition of a structure property to the IE specification template, the creation of an enumeration datatype, and the specification of an IE naming convention.

Provided text to define Collection Guidance, Evaluation Guidance, Classification Guidance, Storage Guidance, and Evaluation Results.

Included additional IEs related to software, configuration, and the vulnerability assessment scenario.

Added text for the IANA considerations, security considerations, operational considerations, and privacy considerations sections.

Performed various other editorial changes and clean-up.

A.8. Changes in Revision 08

Clarified text that describes subjects and attributes.

Clarified text that describes SACM Statements and Content Elements.

Removed stray metadata property fields from the definitions of several IEs.

Specified a syntax for defining category IEs.

Added an anyCategory IE that represents any IE in the IM.

Fixed several errors reported by the Travis-CI continuous integration service.

Performed various other editorial changes and clean-up.

A.9. Changes in Revision 09

Added "derived", "authority", and "verified" to the collectionTaskType IE (<https://github.com/sacmwg/draft-ietf-sacm-information-model/issues/18>).

Updated IE examples that use content-type to use statement-type (<https://github.com/sacmwg/draft-ietf-sacm-information-model/issues/56>).

Added "networkZoneLocation", "layer2NetworkLocation", and "layer3NetworkLocation" IEs (<https://github.com/sacmwg/draft-ietf-sacm-information-model/issues/9>).

Created a softwareClass attribute IE and added it to the softwareInstance subject IE. Also, removed the os* attribute IEs (<https://github.com/sacmwg/draft-ietf-sacm-information-model/issues/10>).

A.10. Changes in Revision 10

Added several IEs necessary for the SACM Vulnerability Assessment Scenario (<https://github.com/sacmwg/draft-ietf-sacm-information-model/issues/43>).

Fixed various typos and formatting issues.

Authors' Addresses

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

Email: david.waltermire@nist.gov

Kim Watson
United States Department of Homeland Security
DHS/CS&C/FNR
245 Murray Ln. SW, Bldg 410
MS0613
Washington, DC 20528
USA

Email: kimberly.watson@hq.dhs.gov

Clifford Kahn
Pulse Secure, LLC
2700 Zanker Road, Suite 200
San Jose, CA 95134
USA

Email: cliffordk@pulsesecure.net

Lisa Lorenzin
Pulse Secure, LLC
2700 Zanker Road, Suite 200
San Jose, CA 95134
USA

Email: llorenzin@pulsesecure.net

Michael Cokus
The MITRE Corporation
903 Enterprise Parkway, Suite 200
Hampton, VA 23666
USA

Email: msc@mitre.org

Daniel Haynes
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
USA

Email: dhaynes@mitre.org

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
Darmstadt 64295
Germany

Email: henk.birkholz@sit.fraunhofer.de

SACM Working Group
Internet-Draft
Intended status: Informational
Expires: June 17, 2019

H. Birkholz
Fraunhofer SIT
J. Lu
Oracle Corporation
J. Strassner
Huawei Technologies
N. Cam-Winget
Cisco Systems
A. Montville
CIS
December 14, 2018

Security Automation and Continuous Monitoring (SACM) Terminology
draft-ietf-sacm-terminology-16

Abstract

This memo documents terminology used in the documents produced by SACM (Security Automation and Continuous Monitoring).

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 <https://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 June 17, 2019.

Copyright Notice

Copyright (c) 2018 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
2. Terms and Definitions	2
3. IANA Considerations	21
4. Security Considerations	21
5. Acknowledgements	22
6. Change Log	22
7. Contributors	26
8. References	27
8.1. Normative References	28
8.2. Informative References	28
Appendix A. The Attic	29
Authors' Addresses	29

1. Introduction

Our goal with this document is to improve our agreement on the terminology used in documents produced by the IETF Working Group for Security Automation and Continuous Monitoring. Agreeing on terminology should help reach consensus on which problems we're trying to solve, and propose solutions and decide which ones to use.

2. Terms and Definitions

This section describes terms that have been defined by other RFC's and defines new ones. The predefined terms will reference the RFC and where appropriate will be annotated with the specific context by which the term is used in SACM. Note that explanatory or informational augmentation to definitions are segregated from the definitions themselves. The definition for the term immediately follows the term on the same line, whereas expository text is contained in subsequent paragraphs immediately following the definition.

Assertion: Defined by the ITU in [X.1252] as "a statement made by an entity without accompanying evidence of its validity".

In the context of SACM, an assertion is the output of a SACM Component in the form of a SACM Statement (including metadata about the data source and data origin, e.g. timestamps). While the validity of an assertion about Content and Content Metadata cannot be verified without, for example, Integrity Proofing of the

Data Source, an assertion (and therefore a SACM statement, respectively) of the validity of Statement Metadata can be enabled by including corresponding Integrity Evidence created by the Data Origin.

Assessment: Defined in [RFC5209] as "the process of collecting posture for a set of capabilities on the endpoint (e.g., host-based firewall) such that the appropriate validators may evaluate the posture against compliance policy."

Attribute: Is a data element, as defined in [RFC5209], that is atomic.

In the context of SACM, attributes are "atomic" information elements and an equivalent to attribute-value-pairs. Attributes can be components of Subjects, the basic composite definitions that are defined in the SACM Information Model.

Capability: A set of features that are available from a SACM Component.

See also "capability" in [I-D.ietf-i2nsf-terminology].

In the context of SACM, the extent of a SACM component's ability is enabled by the functions it is composed of. Capabilities are registered at a SACM broker (potentially also at a proxy or a repository component if it includes broker functions) by a SACM component via the SACM component registration task and can be discovered by or negotiated with other SACM components via the corresponding tasks. For example, the capability of a SACM provider may be to provide target endpoint records (declarative guidance about well-known or potential target endpoints), or only a subset of that data.

A capability's description is in itself imperative guidance on what functions are exposed to other SACM components in a SACM domain and how to use them in workflows.

The SACM Vulnerability Assessment Scenario [I-D.ietf-sacm-vuln-scenario] defines the terms Endpoint Management Capabilities, Vulnerability Management Capabilities, and Vulnerability Assessment Capabilities, which illustrate specific sets of SACM capabilities on an enterprise IT department's point of view and therefore compose sets of declarative guidance.

Collection Result: Is a composition of one or more content elements carrying information about a target endpoint, that is produced by a collector when conducting a collection task.

Collection Task: A targeted task that collects attributes and/or corresponding attribute values from target endpoint.

There are four types of frequency collection tasks can be conducted with:

ad-hoc, e.g. triggered by a unsolicited query

conditional, e.g. triggered in accordance with policies included in the compositions of workflows

scheduled, e.g. in regular intervals, such as every minute or weekly

continuously, e.g. a network behavior observation

There are three types of collection methods, each requiring an appropriate set of functions to be included in the SACM component conducting the collection task:

Self-Reporting: A SACM component located on the target endpoint itself conducts the collection task.

Remote-Acquisition: A SACM component located on an Endpoint different from the target endpoint conducts the collection task via interfaces available on the target endpoint, e.g. SNMP/NETCONF or WMI.

Behavior-Observation: A SACM component located on an Endpoint different from the target endpoint observes network traffic related to the target endpoint and conducts the collection task via interpretation of that network traffic.

Collector: A piece of software that acquires information about one or more target endpoints by conducting collection tasks.

A collector can be distributed across multiple endpoints, e.g. across a target endpoint and a SACM component. The separate parts of the collector can communicate with a specialized protocol, such as PA-TNC [RFC5792]. At least one part of a distributed collector has to take on the role of a provider of information by providing SACM interfaces to propagate capabilities and to provide SACM content in the form of collection results.

Configuration: A non-volatile subset of the endpoint attributes of a endpoint that is intended to be unaffected by a normal reboot-cycle.

Configuration is a type of imperative guidance that is stored in files (files dedicated to contain configuration and/ or files that are software components), directly on block devices, or on specific hardware components that can be accessed via corresponding software components. Modification of configuration can be conducted manually or automatically via management (plane) interfaces that support management protocols, such as SNMP or WMI. A change of configuration can occur during both run-time and down-time of an endpoint. It is common practice to schedule a change of configuration during or directly after the completion of a boot-cycle via corresponding software components located on the target endpoint itself.

Examples: The static association of an IP address and a MAC address in a DHCP server configuration, a directory-path that identifies a log-file directory, a registry entry.

Configuration Drift: The disposition of endpoint characteristics to change over time.

Configuration drift exists for both hardware components and software components. Typically, the frequency and scale of configuration drift of software components is significantly higher than the configuration drift of hardware components.

Consumer: A SACM Role that requires a SACM Component to include SACM Functions enabling it to receive information from other SACM Components.

Content Element: Content elements constitute the payload data (SACM content) transferred via statement Subjects emitted by providers of information. Every content element Subject includes a specific content Subject and a corresponding content metadata Subject.

Content Metadata: Data about content Subjects. Every content-element includes a content metadata Subject. The Subject can include any information element that can annotate the content transferred. Examples include time stamps or data provenance Subjects.

Control Plane: An architectural component that provides common control functions to all SACM components.

Typically used as a term in the context of routing, e.g. [RFC6192]. SACM components may include authentication, authorization, (capability) discovery or negotiation, registration and subscription. The control plane orchestrates the flow on the data plane according to imperative guidance (i.e. configuration) received via the management plane. SACM components with interfaces to the control plane have knowledge of the capabilities of other SACM components within a SACM domain.

Controller: A controller is a SACM Role that is assigned to a SACM component containing control plane functions managing and facilitating information sharing or execute on security functions.

There are three types of SACM controllers: Broker, Proxy, and Repository. Depending on its type, a controller can also contain functions that have interfaces on the data plane.

Data Confidentiality: Defined in [RFC4949] as "the property that data is not disclosed to system entities unless they have been authorized to know the data."

Data In Motion: Data that is being transported via a network; also referred to as "Data in Transit" or "Data in Flight".

Data in motion requires a data model to transfer the data using a specific encoding. Typically, data in motion is serialized (marshalling) into a transport encoding by a provider of information and deserialized (unmarshalling) by a consumer of information. The termination points of provider of information and consumer of information data is transferred between are interfaces. In regard to data in motion, the interpretation of the roles consumer of information and provider of information depends on the corresponding OSI layer (e.g. on layer2: between interfaces connected to a broadcast domain, on layer4: between interfaces that maintain a TCP connection). In the context of SACM, consumer of information and provider of information are SACM components.

Data At Rest: Data that is stored.

Data at rest requires a data model to encode the data to be stored. In the context of SACM, data at rest located on a SACM component can be provided to other SACM components via discoverable capabilities.

Data Integrity: Defined in [RFC4949] as "the property that data has not been changed, destroyed, or lost in an unauthorized or accidental manner."

Data Origin: The SACM Component that initially acquired or produced data about an endpoint.

Data Origin enables a SACM component to identify the SACM component that initially acquired or produced data about a (target) endpoint (e.g. via collection from a data source) and made it available to a SACM domain via a SACM statement. Data Origin can be expressed by an endpoint label information element (e.g. to be used as metadata in statement).

Data Plane: Is an architectural component providing operational functions enabling information exchange that is not command and control or management related.

Typically used as a term in the context of routing (and used as a synonym for forwarding plane, e.g. [RFC6192]). In the context of SACM, the data plane is an architectural component providing operational functions to enable a SACM component to provide and consume SACM statements and therefore SACM content, which composes the actual SACM content. The data plane in a SACM domain is used to conduct distributed SACM tasks by transporting SACM content via specific transport encodings and corresponding operations defined by SACM data models.

Data Provenance: An historical record of the sources, origins and evolution, as it pertains to data, that is influenced by inputs, entities, functions and processes.

Additional Information - In the context of SACM, data provenance is expressed as metadata that identifies SACM statements and corresponding content elements a new statement is created from. In a downstream process, this references can cascade, creating a data provenance tree that enables SACM components to trace back the original data sources involved in the creation of SACM statements and take into account their characteristics and trustworthiness.

Data Source: Is an endpoint from which a particular set of attributes and/or attribute values have been collected.

Data Source enables a SACM component to identify - and potentially characterize - a (target) endpoint that is claimed to be the original source of endpoint attributes in a SACM statement. Data Source can be expressed as metadata by an endpoint label information element or a corresponding subject of identifying endpoint attributes.

Endpoint: Defined in [RFC5209] as "any computing device that can be connected to a network."

Additional Information - The [RFC5209] definition continues, "Such devices normally are associated with a particular link layer address before joining the network and potentially an IP address once on the network. This includes: laptops, desktops, servers, cell phones, or any device that may have an IP address."

To further clarify the [RFC5209] definition, an endpoint is any physical or virtual device that may have a network address. Note that, network infrastructure devices (e.g. switches, routers, firewalls), which fit the definition, are also considered to be endpoints within this document.

Physical endpoints are always composites that are composed of hardware components and software components. Virtual endpoints are composed entirely of software components and rely on software components that provide functions equivalent to hardware components.

The SACM architecture differentiates two essential categories of endpoints: Endpoints whose security posture is intended to be assessed (target endpoints) and endpoints that are specifically excluded from endpoint posture assessment (excluded endpoints).

Based on the definition of an asset, an endpoint is a type of asset.

Endpoint Attribute: Is a discreet endpoint characteristic that is computably observable.

Endpoint Attributes typically constitute Attributes that can be bundled into Subject (e.g. information about a specific network interface can be represented via a set of multiple AVP).

Endpoint Characteristics: The state, configuration and composition of the software components and (virtual) hardware components a target endpoint is composed of, including observable behavior, e.g. sys-calls, log-files, or PDU emission on a network.

In SACM work-flows, (Target) Endpoint Characteristics are represented via Information Elements.

Endpoint Characterization Task: The task of endpoint characterization that uses endpoint attributes that represent distinct endpoint characteristics.

Endpoint Classification: The categorization of of the endpoint into one or more taxonomic structures.

Endpoint classification requires declarative guidance in the form of an endpoint profile, discovery results and potentially collection results. Types, classes or the characteristics of an individual target endpoint are defined via endpoint profiles.

Endpoint Classification Task: The task of endpoint classification that uses an endpoint's characteristics to determine how to categorize the given endpoint into one or more taxonomic structures.

Endpoint Label: A unique label associated with a unique endpoint.

Endpoint specializations have corresponding endpoint label specializations. For example, an endpoint label used on a SACM Component is a SACM Component Label.

Endpoint Management Capabilities: Enterprise IT management capabilities that are tailored to manage endpoint identity, endpoint information, and associated metadata.

Evaluation Task: A task by which an endpoint's asserted attribute value is evaluated against a policy-compliant attribute value.

Evaluation Result: The resulting value from having evaluated a set of posture attributes.

Expected Endpoint Attribute State: The policy-compliant state of an endpoint attribute that is to be compared against.

Sets of expected endpoint attribute states are transported as declarative guidance in target endpoint profiles via the management plane. This, for example, can be a policy, but also a recorded past state. An expected state is represented by an Attribute or a Subject that represents a set of multiple attribute value pairs.

Guidance: Machine-processable input directing SACM processes or tasks.

Examples of such processes/tasks include automated device management, remediation, collection, evaluation. Guidance influences the behavior of a SACM Component and is considered content of the management plane. In the context of SACM, guidance is machine-readable and can be manually or automatically generated

or provided. Typically, the tasks that provide guidance to SACM components have a low-frequency and tend to be sporadic.

There are two types of guidance:

Declarative Guidance: Guidance that defines the configuration or state an endpoint is supposed to be in, without providing specific actions or methods to produce that desired state. Examples include Target Endpoint Profiles or network topology based requirements.

Imperative Guidance: Guidance that prescribes specific actions to be conducted or methods to be used in order to achieve an outcome. Examples include a targeted Collection Task or the IP-Address of a SACM Component that provides a registration function.

Prominent examples include: modification of the configuration of a SACM component or updating a target endpoint profile that resides on an evaluator. In essence, guidance is transported via the management plane.

Endpoint Hardware Inventory: The set of hardware components that compose a specific endpoint representing its hardware configuration.

Hardware Component: A distinguishable physical component used to compose an endpoint.

The composition of an endpoint can be changed over time by adding or removing hardware components. In essence, every physical endpoint is potentially a composite of multiple hardware components, typically resulting in a hierarchical composition of hardware components. The composition of hardware components is based on interconnects provided by specific hardware types (e.g. FRU in a chassis are connected via redundant busses). In general, a hardware component can be distinguished by its serial number. Occasionally, hardware components are referred to as power sucking aliens.

Information Element: A representation of information about physical and virtual "objects of interest".

Information elements are the building blocks that constitute the SACM information model. In the context of SACM, an information element that expresses a single value with a specific name is referred to as an Attribute (analogous to an attribute-value-pair). A set of attributes that is bundled into a more complex composite information element is referred to as a Subject. Every

information element in the SACM information model has a unique name. Endpoint attributes or time stamps, for example, are represented as information elements in the SACM information model.

Information Model: An abstract representation of data, their properties, relationships between data and the operations that can be performed on the data.

While there is some overlap with a data model, [RFC3444] distinguishes an information model as being protocol and implementation neutral whereas a data model would provide such details. The purpose of the SACM information model is to ensure interoperability between SACM data models (that are used as transport encoding) and to provide a standardized set of information elements for communication between SACM components.

Interaction Model: The definition of specific sequences regarding the exchange of messages (data in motion), including, for example, conditional branching, thresholds and timers.

An interaction model, for example, can be used to define operations, such as registration or discovery, on the control plane. A composition of data models for data in motion and a corresponding interaction model is a protocol.

Internal Collector: A collector that runs on a target endpoint to acquire information from that target endpoint.

Management Plane: An architectural component providing common functions to steer the behavior of SACM components, e.g. their behavior on the control plane.

Typically, a SACM component can fulfill its purpose without continuous input from the management plane. In contrast, without continuous availability of control plane functions a typical SACM component could not function properly. In general, interaction on the management plane is less frequent and less regular than on the control plane. Input via the management plane can be manual (e.g. via a CLI), or can be automated via management plane functions that are part of other SACM components.

Network Address: A layer-specific address that follows a layer-specific address scheme.

The following characteristics are a summary derived from the Common Information Model and ITU-T X.213. Each Network Interface of a specific layer can be associated with one or more addresses appropriate for that layer. There is no guarantee that a network

address is globally unique. A dedicated authority entity can provide a level of assurance that a network address is unique in its given scope. In essence, there is always a scope to a network address, in which it is intended to be unique.

Examples include: physical Ethernet port with a MAC address, layer 2 VLAN interface with a MAC address, layer 3 interface with multiple IPv6 addresses, layer 3 tunnel ingress or egress with an IPv4 address.

Network Interface: An Endpoint is connected to a network via one or more Network Interfaces. Network Interfaces can be physical (Hardware Component) or logical (virtual Hardware component, i.e. a dedicated Software Component). Network Interfaces of an Endpoint can operate on different layers, most prominently what is now commonly called layer 2 and 3. Within a layer, interfaces can be nested.

In SACM, the association of Endpoints and Network Addresses via Network Interfaces is vital to maintain interdependent autonomous processes that can be targeted at Target Endpoints, unambiguously.

Examples include: physical Ethernet port, layer 2 VLAN interface, a MC-LAG setup, layer 3 Point-to-Point tunnel ingress or egress.

Metadata: Data about data.

In the SACM information model, data is referred to as Content. Metadata about the content is referred to as Content-Metadata, respectively. Content and Content-Metadata are combined into Subjects called Content-Elements in the SACM information model. Some information elements defined by the SACM information model can be part of the Content or the Content-Metadata. Therefore, if an information element is considered data or data about data depends on which kind of Subject it is associated with. The SACM information model also defines metadata about the data origin via the Subject Statement-Metadata. Typical examples of metadata are time stamps, data origin or data source.

Posture: Defined in [RFC5209] as "configuration and/or status of hardware or software on an endpoint as it pertains to an organization's security policy."

This term is used within the scope of SACM to represent the configuration and state information that is collected from a target endpoint in the form of endpoint attributes (e.g. software/hardware inventory, configuration settings, dynamically assigned

addresses). This information may constitute one or more posture attributes.

Posture Attributes: Defined in [RFC5209] as "attributes describing the configuration or status (posture) of a feature of the endpoint. A Posture Attribute represents a single property of an observed state. For example, a Posture Attribute might describe the version of the operating system installed on the system."

Within this document this term represents a specific assertion about endpoint configuration or state (e.g. configuration setting, installed software, hardware) represented via endpoint attributes. The phrase "features of the endpoint" highlighted above refers to installed software or software components.

Provider: A provider is a SACM role assigned to a SACM component that provides role-specific functions to provide information to other SACM components.

Repository: A repository is a controller that contains functions to consume, store and provide information of a particular kind.

Such information is typically data transported on the data plane, but potentially also data and metadata from the control and management plane. A single repository may provide the functions of more than one specific repository type (i.e. configuration baseline repository, assessment results repository, etc.)

SACM Broker Controller: A SACM Broker Controller is a controller that contains control plane functions to provide and/or connect services on behalf of other SACM components via interfaces on the control plane.

A broker may provide, for example, authorization services and find, upon request, SACM components providing requested services.

SACM Component: Is a component, as defined in [I-D.ietf-i2nsf-terminology], that is composed of SACM capabilities.

In the context of SACM, a set of SACM functions composes a SACM component. A SACM component conducts SACM tasks, acting on control plane, data plane and/or management plane via corresponding SACM interfaces. SACM defines a set of standard components (e.g. a collector, a broker, or a data store). A SACM component contains at least a basic set of control plane functions and can contain data plane and management plane functions. A SACM component residing on an endpoint assigns one or more SACM roles

to the corresponding endpoint due to the SACM functions it is composed of. A SACM component "resides on" an endpoint and an endpoint "contains" a SACM component, correspondingly. For example, a SACM component that is composed solely of functions that provide information would only take on the role of a provider.

SACM Component Discovery: The task of discovering the capabilities provided by SACM components within a SACM domain.

This is likely to be performed via an appropriate set of control plane functions.

SACM Component Label: A specific endpoint label that is used to identify a SACM component.

In content-metadata, this label is called data origin.

SACM Content: The payload provided by SACM components to the SACM domain on the data plane.

SACM content includes the SACM data models.

SACM Domain: Endpoints that include a SACM component compose a SACM domain.

(To be revised, additional definition content TBD, possible dependencies to SACM architecture)

SACM Function: A behavioral aspect of a SACM component that provides external SACM Interfaces or internal interfaces to other SACM Functionse.

For example, a SACM Function with SACM Interfaces on the Control Plane can provide a brokering function to other SACM Components. Via Data Plane interfaces, a SACM Function can act as a provider and/or as a consumer of information. SACM Functions can be propagated as the Capabilities of a SACM Component and can be discovered by or negotiated with other SACM Components.

SACM Interface: An interface, as defined in [I-D.ietf-i2nsf-terminology], that provides SACM-specific operations.

[I-D.ietf-i2nsf-terminology] defines interface as a "set of operations one object knows it can invoke on, and expose to, another object," and further defines interface by stating that an interface "decouples the implementation of the operation from its

specification. An interface is a subset of all operations that a given object implements. The same object may have multiple types of interfaces to serve different purposes."

In the context of SACM, SACM Functions provide SACM Interfaces on the management, control, or data plane. Operations a SACM Interface provides are based on corresponding data model defined by SACM. SACM Interfaces are used for communication between SACM components.

SACM Proxy Controller: A SACM Proxy Controller is a controller that provides data plane and control plane functions, information, or services on behalf of another component, which is not directly participating in the SACM architecture.

SACM Role: Is a role, as defined in [I-D.ietf-i2nsf-terminology], that requires the SACM Component assuming the role to bear a set of SACM functions or interfaces.

SACM Roles provide three important benefits. First, it enables different behavior to be supported by the same Component for different contexts. Second, it enables the behavior of a Component to be adjusted dynamically (i.e., at runtime, in response) to changes in context, by using one or more Roles to define the behavior desired for each context. Third, it decouples the Roles of a Component from the Applications that use that Component."

In the context of SACM, SACM roles are associated with SACM components and are defined by the set of functions and interfaces a SACM component includes. There are three SACM roles: provider, consumer, and controller. The roles associated with a SACM component are determined by the purpose of the SACM functions and corresponding SACM interfaces the SACM component is composed of.

SACM Statement: Is an assertion that is made by a SACM Component.

Security Automation: The process of which security alerts can be automated through the use of different components to monitor, analyze and assess endpoints and network traffic for the purposes of detecting misconfigurations, misbehaviors or threats.

Security Automation is intended to identify target endpoints that cannot be trusted (see "trusted" in [RFC4949]). This goal is achieved by creating and processing evidence (assessment statements) that a target endpoint is not a trusted system [RFC4949].

Software Package: A generic software package (e.g. a text editor).

Software Component: A software package installed on an endpoint.

The software component may include a unique serial number (e.g. a text editor associated with a unique license key).

Software Instance: A running instance of a software component.

For example, on a multi-user system, one logged-in user has one instance of a text editor running and another logged-in user has another instance of the same text editor running, or on a single-user system, a user could have multiple independent instances of the same text editor running.

State: A volatile set of endpoint attributes of a (target) endpoint that is affected by a reboot-cycle.

Local state is created by the interaction of components with other components via the control plane, via processing data plane payload, or via the functional properties of local hardware and software components. Dynamic configuration (e.g. IP address distributed dynamically via an address distribution and management services, such as DHCP) is considered state that is the result of the interaction with another component (e.g. provided by a DHCP server with a specific configuration).

Examples: The static association of an IP address and a MAC address in a DHCP server configuration, a directory-path that identifies a log-file directory, a registry entry.

Statement: A statement is the root/top-level subject defined in the SACM information model.

A statement is used to bundle Content Elements into one subject and includes metadata about the data origin.

Subject: A semantic composite information element pertaining to a system entity that is a target endpoint.

Like Attributes, subjects have a name and are composed of attributes and/or other subjects. Every IE that is part of a subject can have a quantity associated with it (e.g. zero-one, none-unbounded). The content IE of a subject can be an unordered or an ordered list.

In contrast to the definitions of subject provided by [RFC4949], a subject in the scope of SACM is neither "a system entity that

causes information to flow among objects or changes the system state" nor "a name of a system entity that is bound to the data items in a digital certificate".

In the context of SACM, a subject is a semantic composite of information elements about a system entity that is a target endpoint. Every acquirable subject-as defined in the scope of SACM-about a target endpoint represents and therefore identifies every subject-as defined by [RFC4949]-that is a component of that target endpoint. The semantic difference between both definitions can be subtle in practice and is in consequence important to highlight.

Supplicant: A component seeking to be authenticated via the control plane for the purpose of participating in a SACM domain.

System Resource: Defined in [RFC4949] as "data contained in an information system; or a service provided by a system; or a system capacity, such as processing power or communication bandwidth; or an item of system equipment (i.e., hardware, firmware, software, or documentation); or a facility that houses system operations and equipment."

Target Endpoint: Is an endpoint that is under assessment at some point in, or region of, time.

Every endpoint that is not specifically designated as an excluded endpoint is a target endpoint. A target endpoint is not part of a SACM domain unless it contains a SACM component (e.g. a SACM component that publishes collection results coming from an internal collector).

A target endpoint is similar to a device that is a Target of Evaluation (TOE) as defined in Common Criteria and as referenced by {{RFC4949}}.

Target Endpoint Address: An address that is layer specific and which follows layer specific address schemes.

Each interface of a specific layer can be associated with one or more addresses appropriate for that layer. There is no guarantee that an address is globally unique. In general, there is a scope to an address in which it is intended to be unique.

Examples include: physical Ethernet port with a MAC address, layer 2 VLAN interface with a MAC address, layer 3 interface with multiple IPv6 addresses, layer 3 tunnel ingress or egress with an IPv4 address.

Target Endpoint Characterization: The description of the distinctive nature of a target endpoint, that is based on its characteristics.

Target Endpoint Characterization Record: A set of endpoint attributes about a target endpoint that was encountered in a SACM domain, which are associated with that target endpoint as a result of a Target Endpoint Characterization Task.

A characterization record is intended to be a representation of an endpoint. It cannot be assured that a record distinctly represents a single target endpoint unless a set of one or more endpoint attributes that compose a unique set of identifying endpoint attributes are included in the record. Otherwise, the set of identifying attributes included in a record can match more than one target endpoints, which are - in consequence - indistinguishable to a SACM domain until more qualifying endpoint attributes can be acquired and added to the record. A characterization record is maintained over time in order to assert that acquired endpoint attributes are either about an endpoint that was encountered before or an endpoint that has not been encountered before in a SACM domain. A characterization record can include, for example, acquired configuration, state or observed behavior of a specific target endpoint. Multiple and even conflicting instances of this information can be included in a characterization record by using timestamps and/or data origins to differentiate them. The endpoint attributes included in a characterization record can be used to re-identify a distinct target endpoint over time. Classes or profiles can be associated with a characterization record via the Classification Task in order to guide collection, evaluation or remediation tasks.

Target Endpoint Characterization Task: An ongoing task of continuously adding acquired endpoint attributes to a corresponding record. The TE characterization task manages the representation of encountered target endpoints in the SACM domain in the form of characterization records. For example, the output of a target endpoint discovery task or a collection task can be processed by the characterization task and added to the record. The TE characterization Task also manages these representations of target endpoints encountered in the SACM domain by splitting or merging the corresponding records as new or more refined endpoint attributes become available.

Target Endpoint Classification Task: The task of associating a class from an extensible list of classes with an endpoint characterization record. TE classes function as imperative and declarative guidance for collection, evaluation, remediation and security posture assessment in general.

Target Endpoint Discovery Task: The ongoing task of detecting previously unknown interaction of a potential target endpoint in the SACM domain. TE Discovery is not directly targeted at a specific target endpoint and therefore an un-targeted task. SACM Components conducting the discovery task as a part of their function are typically distributed and located, for example, on infrastructure components or collect from those remotely via appropriate interfaces. Examples of infrastructure components that are of interest to the discovery task include routers, switches, VM hosting or VM managing components, AAA servers, or servers handling dynamic address distribution.

Target Endpoint Identifier: The target endpoint discovery task and the collection tasks can result in a set of identifying endpoint attributes added to a corresponding Characterization Record. This subset of the endpoint attributes included in the record is used as a target endpoint identifier, by which a specific target endpoint can be referenced. Depending on the available identifying attributes, this reference can be ambiguous and is a "best-effort" mechanism. Every distinct set of identifying endpoint attributes can be associated with a target endpoint label that is unique in a SACM domain.

Target Endpoint Label: An endpoint label that identifies a specific target endpoint.

Target Endpoint Profile: A bundle of expected or desired component composition, configurations and states that is associated with a target endpoint.

The corresponding task by which the association with a target endpoint takes places is the endpoint classification task. The task by which an endpoint profile is created is the endpoint characterization task. A type or class of target endpoints can be defined via a target endpoint profile. Examples include: printers, smartphones, or an office PC.

In respect to [RFC4949], a target endpoint profile is a protection profile as defined by Common Criteria (analogous to the target endpoint being the target of evaluation).

SACM Task: Is a task conducted within the scope of a SACM domain by one or more SACM functions that achieves a SACM-defined outcome.

A SACM task can be triggered by other operations or functions (e.g. a query from another SACM component or an unsolicited push on the data plane due to an ongoing subscription). A task is part of a SACM process chain. A task starts at a given point in time

and ends in a deterministic state. With the exception of a collection task, a SACM task consumes SACM statements provided by other SACM components. The output of a task is a result that can be provided (e.g. published) on the data plane.

The following tasks are defined by SACM:

Target Endpoint Discovery

Target Endpoint Characterization

Target Endpoint Classification

Collection

Evaluation [TBD]

Information Sharing [TBD]

SACM Component Discovery

SACM Component Authentication [TBD]

SACM Component Authorization [TBD]

SACM Component Registration [TBD]

Timestamps : Defined in [RFC4949] as "with respect to a data object, a label or marking in which is recorded the time (time of day or other instant of elapsed time) at which the label or marking was affixed to the data object".

A timestamp always requires context, i.e. additional information elements that are associated with it. Therefore, all timestamps wrt information elements are always metadata. Timestamps in SACM Content Elements may be generated outside a SACM Domain and may be encoded in an unknown representation. Inside a SACM domain the representation of timestamps is well-defined and unambiguous.

Virtual Endpoint: An endpoint composed entirely of logical system components (see [RFC4949]).

The most common example is a virtual machine/host running on a target endpoint. Effectively, target endpoints can be nested and at the time of this writing the most common example of target endpoint characteristics about virtual components is the EntLogicalEntry in [RFC6933].

Vulnerability Assessment: An assessment specifically tailored to determining whether a set of endpoints is vulnerable according to the information contained in the vulnerability description information.

Vulnerability Description Information: Information pertaining to the existence of a flaw or flaws in software, hardware, and/or firmware, which could potentially have an adverse impact on enterprise IT functionality and/or security.

Vulnerability description information should contain enough information to support vulnerability detection.

Vulnerability Detection Data: A type of imperative guidance extracted or derived from vulnerability description information that describes the specific mechanisms of vulnerability detection that is used by an enterprise's vulnerability management capabilities to determine if a vulnerability is present on an endpoint.

Vulnerability Management Capabilities: An IT management capability tailored toward managing endpoint vulnerabilities and associated metadata on an ongoing basis by ingesting vulnerability description information and vulnerability detection data, and performing vulnerability assessments.

Vulnerability assessment capabilities: An assessment capability that is tailored toward determining whether a set of endpoints is vulnerable according to vulnerability description information.

Workflow: A workflow is a modular composition of tasks that can contain loops, conditionals, multiple starting points and multiple endpoints.

The most prominent workflow in SACM is the assessment workflow.

3. IANA Considerations

This memo includes no request to IANA.

4. Security Considerations

This memo documents terminology for security automation. While it is about security, it does not affect security.

5. Acknowledgements

6. Change Log

Changes from version 00 to version 01:

- o Added simple list of terms extracted from UC draft -05. It is expected that comments will be received on this list of terms as to whether they should be kept in this document. Those that are kept will be appropriately defined or cited.

Changes from version 01 to version 02:

- o Added Vulnerability, Vulnerability Management, xposure, Misconfiguration, and Software flaw.

Changes from version 02 to version 03:

- o Removed Section 2.1. Cleaned up some editing nits; broke terms into 2 sections (predefined and newly defined terms). Added some of the relevant terms per the proposed list discussed in the IETF 89 meeting.

Changes from version 03 to version 04:

- o TODO

Changes from version 04 to version 05:

- o TODO

Changes from version 05 to version 06:

- o Updated author information.
- o Combined "Pre-defined Terms" with "New Terms and Definitions".
- o Removed "Requirements language".
- o Removed unused reference to use case draft; resulted in removal of normative references.
- o Removed introductory text from Section 1 indicating that this document is intended to be temporary.
- o Added placeholders for missing change log entries.

Changes from version 06 to version 07:

- o Added Contributors section.
- o Updated author list.
- o Changed title from "Terminology for Security Assessment" to "Secure Automation and Continuous Monitoring (SACM) Terminology".
- o Changed abbrev from "SACM-Terms" to "SACM Terminology".
- o Added appendix The Attic to stash terms for future updates.
- o Added Authentication, Authorization, Data Confidentiality, Data Integrity, Data Origin, Data Provenance, SACM Component, SACM Component Discovery, Target Endpoint Discovery.
- o Major updates to Building Block, Function, SACM Role, Target Endpoint.
- o Minor updates to Broker, Capability, Collection Task, Evaluation Task, Posture.
- o Relabeled Role to SACM Role, Endpoint Target to Target Endpoint, Endpoint Discovery to Endpoint Identification.
- o Moved Asset Targeting, Client, Endpoint Identification to The Attic.
- o Endpoint Attributes added as a TODO.
- o Changed the structure of the Change Log.

Changes from version 07 to version 08:

- o Added Assertion, Collection Result, Collector, Excluded Endpoint, Internal Collector, Network Address, Network Interface, SACM Domain, Statement, Target Endpoint Identifier, Target Endpoint Label, Timestamp.
- o Major updates to Attributes, Broker, Collection Task, Consumer, Controller, Control Plane, Endpoint Attributes, Expected Endpoint State, SACM Function, Provider, Proxy, Repository, SACM Role, Target Endpoint.
- o Minor updates to Asset, Building Block, Data Origin, Data Source, Data Provenance, Endpoint, Management Plane, Posture, Posture Attribute, SACM Component, SACM Component Discovery, Target Endpoint Discovery.

- o Relabeled Function to SACM Function.

Changes from version 08 to version 09:

- o Updated author list.
- o Added Data Plane, Endpoint Characterization, Endpoint Classification, Guidance, Interaction Model, Software Component, Software Instance, Software Package, Statement, Target Endpoint Profile, SACM Task.
- o Removed Building Block.
- o Major updates to Control Plane, Endpoint Attribute, Expected Endpoint State, Information Model, Management Plane.
- o Minor updates to Attribute, Capabilities, SACM Function, SACM Component, Collection Task.
- o Moved Asset Characterization to The Attic.

Changes from version 09 to version 10:

- o Added Configuration Drift, Data in Motion, Data at Rest, Endpoint Management Capability, Hardware Component, Hardware Inventory, Hardware Type, SACM Interface, Target Endpoint Characterization Record, Target Endpoint Characterization Task, Target Endpoint Classification Task, Target Endpoint Discovery Task, Vulnerability Description Information, Vulnerability Detection Data, Vulnerability Management Capability, Vulnerability Assessment
- o Added references to i2nsf definitions in Capability, SACM Component, SACM Interface, SACM Role.
- o Added i2nsf Terminology I-D Reference.
- o Major Updates to Endpoint, SACM Task, Target Endpoint Identifier.
- o Minor Updates to Guidance, SACM Component Discovery, Target Endpoint Label, Target Endpoint Profile.
- o Relabeled SACM Task
- o Removed Target Endpoint Discovery

Changes from version 10 to version 11:

- o Added Content Element, Content Metadata, Endpoint Label, Information Element, Metadata, SACM Component Label, Workflow.
- o Major Updates to Assessment, Capability, Collector, Endpoint Management Capabilities, Guidance, Vulnerability Assessment Capabilities, Vulnerability Detection Data, Vulnerability Assessment Capabilities.
- o Minor updates to Collection Result, Control Plane, Data in Motion, Data at Rest, Data Origin, Network Interface, Statement, Target Endpoint Label.
- o Relabeled Endpoint Management Capability, Vulnerability Management Capability, Vulnerability Assessment.

Changes from version 11 to version 12:

- o Added Configuration, Endpoint Characteristic, Event, SACM Content, State, Subject.
- o Major Updates to Assertion, Data in Motion, Data Provenance, Data Source, Interaction Model.
- o Minor Updates to Attribute, Control Plane, Data Origin, Data Provenance, Expected Endpoint State, Guidance, Target Endpoint Classification Task, Vulnerability Detection Data.

Changes from version 12 to version 13:

- o Added Virtual Component.
- o Major Updates to Capability, Collection Task, Hardware Component, Hardware Type, Security Automation, Subject, Target Endpoint, Target Endpoint Profile.
- o Minor Updates to Assertion, Data Plane, Endpoint Characteristics.

Changes from version 13 to version 14:

- o Handled a plethora of issues listed in GitHub.
- o Pruned some commonly understood terms.
- o Narrowing term labels per their definitions.
- o In some cases, excised expositional text.

- o Where expository text was left intact, it has been separated from the actual definition of a term.

Changes from version 14 to version 16:

- o moved obsolete definitions into the Appendix (attic).

7. Contributors

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

Email: david.waltermire@nist.gov

Adam W. Montville
Center for Internet Security
31 Tech Valley Drive
East Greenbush, NY 12061
USA

Email: adam.w.montville@gmail.com

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

Email: ietfdbh@comcast.net

Brian Ford
Lancope
3650 Brookside Parkway, Suite 500
Alpharetta, GA 30022
USA

Email: bford@lancope.com

Merike Kaeo
Double Shot Security
3518 Fremont Avenue North, Suite 363
Seattle, WA 98103
USA

Email: merike@doubleshotsecurity.com

8. References

8.1. Normative References

- [RFC5792] Sangster, P. and K. Narayan, "PA-TNC: A Posture Attribute (PA) Protocol Compatible with Trusted Network Connect (TNC)", RFC 5792, DOI 10.17487/RFC5792, March 2010, <<https://www.rfc-editor.org/info/rfc5792>>.
- [RFC6933] Bierman, A., Romascanu, D., Quittek, J., and M. Chandramouli, "Entity MIB (Version 4)", RFC 6933, DOI 10.17487/RFC6933, May 2013, <<https://www.rfc-editor.org/info/rfc6933>>.

8.2. Informative References

- [I-D.ietf-i2nsf-terminology]
Hares, S., Strassner, J., Lopez, D., Xia, L., and H. Birkholz, "Interface to Network Security Functions (I2NSF) Terminology", draft-ietf-i2nsf-terminology-06 (work in progress), July 2018.
- [I-D.ietf-netmod-entity]
Bierman, A., Bjorklund, M., Dong, J., and D. Romascanu, "A YANG Data Model for Hardware Management", draft-ietf-netmod-entity-08 (work in progress), January 2018.
- [I-D.ietf-sacm-vuln-scenario]
Coffin, C., Cheikes, B., Schmidt, C., Haynes, D., Fitzgerald-McKay, J., and D. Waltermire, "SACM Vulnerability Assessment Scenario", draft-ietf-sacm-vuln-scenario-02 (work in progress), September 2016.
- [RFC3444] Pras, A. and J. Schoenwaelder, "On the Difference between Information Models and Data Models", RFC 3444, DOI 10.17487/RFC3444, January 2003, <<https://www.rfc-editor.org/info/rfc3444>>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<https://www.rfc-editor.org/info/rfc4949>>.
- [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>>.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, DOI 10.17487/RFC6192, March 2011, <<https://www.rfc-editor.org/info/rfc6192>>.

[X.1252] "ITU-T X.1252 (04/2010)", n.d..

Appendix A. The Attic

The following terms are stashed for now and will be updated later:

Asset: Is a system resource, as defined in [RFC4949], that may be composed of other assets.

Examples of Assets include: Endpoints, Software, Guidance, or X.509 public key certificates. An asset is not necessarily owned by an organization.

Asset Management: The IT process by which assets are provisioned, updated, maintained and deprecated.

Asset Characterization: Asset characterization is the process of defining attributes that describe properties of an identified asset.

Asset Targeting: Asset targeting is the use of asset identification and categorization information to drive human-directed, automated decision making for data collection and analysis in support of endpoint posture assessment.

Client: An architectural component receiving services from another architectural component.

Endpoint Identification (TBD per list; was "Endpoint Discovery"):
The process by which an endpoint can be identified.

Authors' Addresses

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
Darmstadt 64295
Germany

Email: henk.birkholz@sit.fraunhofer.de

Jarrett Lu
Oracle Corporation
4180 Network Circle
Santa Clara, CA 95054
USA

Email: jarrett.lu@oracle.com

John Strassner
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95138
USA

Email: john.sc.strassner@huawei.com

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

Email: ncamwing@cisco.com

Adam Montville
Center for Internet Security
31 Tech Valley Drive
East Greenbush, NY 12061
USA

Email: adam.w.montville@gmail.com