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

Table of Contents

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:

- 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).

- 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:

- 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 my 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:

- SACM Domain Identifier: a globally unique identifier that enables the differentiation of SACM statements across SACM domains.
- Data Origin: the SACM domain unique identifier associated with a SACM component.
- 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 mnore 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:

- Target Endpoint Label: an identifier that enables to distinctly identify a target endpoint as a SACM content element.
- Relationship Identifier(s): a set of semantic relationships that associate this SACM content element with other SACM content elements via their content element identifier.
- 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 preludes 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 occour, 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-priviledge, 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 efficiantly 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 the 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:

- (*): zero to unbounded occurrences
- (+): one to unbounded occurrences
- (?): zero or one occurrence
- (n*m): between n and m occurrences
- 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 in 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 specfific 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 fucntions as a contact

phone-number

phone-number-type (?)

priviledge: a CIE to express priviledges 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 iden

relationship-type

relationship-content-element-guid (*)

relationship-statement-guid (*)

relationship-object-label (*)

sacm-statement: every SACM components 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 assessement

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