Document Number: DSP2028

Date: 2012-06-26

Version: 1.0.0a

Cloud Auditing Data Federation (CADF) Use Case White Paper

Information for Work-in-Progress version:

IMPORTANT: This document is not a standard. It does not necessarily reflect the views of the DMTF or all of its members. Because this document is a Work in Progress, it may still change, perhaps profoundly. This document is available for public review and comment until the stated expiration date.

It expires on: 2012-10-30

Provide any comments through the DMTF Feedback Portal:
http://www.dmtf.org/standards/feedback

Document Type: DMTF Informational

Document Status: Work In Progress

Document Language: en-US

Copyright Notice

Copyright © 2012 Distributed Management Task Force, Inc. (DMTF). All rights reserved.

DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems management and interoperability. Members and non-members may reproduce DMTF specifications and documents, provided that correct attribution is given. As DMTF specifications may be revised from time to time, the particular version and release date should always be noted.

Implementation of certain elements of this standard or proposed standard may be subject to third party patent rights, including provisional patent rights (herein "patent rights"). DMTF makes no representations to users of the standard as to the existence of such rights, and is not responsible to recognize, disclose, or identify any or all such third party patent right, owners or claimants, nor for any incomplete or inaccurate identification or disclosure of such rights, owners or claimants. DMTF shall have no liability to any party, in any manner or circumstance, under any legal theory whatsoever, for failure to recognize, disclose, or identify any such third party patent rights, or for such party's reliance on the standard or incorporation thereof in its product, protocols or testing procedures. DMTF shall have no liability to any party implementing such standard, whether such implementation is foreseeable or not, nor to any patent owner or claimant, and shall have no liability or responsibility for costs or losses incurred if a standard is withdrawn or modified after publication, and shall be indemnified and held harmless by any party implementing the standard from any and all claims of infringement by a patent owner for such implementations.

For information about patents held by third-parties which have notified the DMTF that, in their opinion, such patent may relate to or impact implementations of DMTF standards, visit http://www.dmtf.org/about/policies/disclosures.php.

Contents

Abstract

The Cloud Auditing Data Federation (CADF) Working Group determined to develop and publish granular use cases around cloud auditing and data federation that will be used as input for development of their data format and interface specification. The use cases contained within are not normative or comprehensive but represent submissions by working group members for specific consideration.

The use cases included in this whitepaper (or portions of) are intended to contribute to the development of DSP0262 "CADF *Data Format and Interface Definitions Specification*" by providing input material that may be considered.

The creation of the use cases listed in this white paper is permitted by the CADF charter as "In Scope" under the "WG Deliverables" section.

CADF WG Charter excerpt:

WG deliverables

a) Cloud Audit Event Data Model Specification
    a) Including Resource, Action and Outcome Taxonomies
    b) Including Guidance and Best Practices for Use of the Data Model.
b) Cloud Audit Event API Specification
    a) Including an exemplary Component Model
    b) Including Use Cases

Foreword

The *Cloud Auditing Data Federation (CADF) Use Case White Paper* (DSP2028) was prepared by the Cloud Auditing Data Federation (CADF) Working Group.

DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems management and interoperability. For information about the DMTF, see http://www.dmtf.org.

Acknowledgments

The DMTF acknowledges the following individuals for their contribution to this document:

- Alvin Black – CA Technologies (Editor)
- Winston Bumpus – VMWare
- Rick Cohen – IBM
- David Corlette – NetIQ (Co-Chair)
- Jacques Durand – Fujitsu
- Il-Sung Lee – Microsoft
- Monica Martin – Microsoft
- Steve Neely – Cisco
- Davi Ottenheimer – VMWare
- John Parchem – Microsoft
- Martin Pohlman – EMC
- Matt Rutkowski – IBM (Co-Chair, Editor)
- Hemal Shah – Broadcom
- Song, Zhexuan – Huawei

Introduction

Target audience

The target audience for this white paper is those developing standards for cloud auditing including the members of the Cloud Auditing Data Federation (CADF) Working Group.

Cloud Auditing Data Federation (CADF) - Use Case White Paper

1 Executive summary

This document is intended to provide a set of real-world use cases representing certain auditing considerations of cloud based resources. These considerations include the types of data, resources and interactions expected by entities responsible for auditing the compliance of systems, applications, and data hosted in cloud deployments. These entities include data and application administrators, corporate security and compliance officers and corporate auditors, and service and tool vendors in the cloud auditing ecosystem.

The use cases in this document represent the use cases proposed by the companies or individuals who submitted them. They may use terminology or semantics which is not consistent with the specification being developed.

The use cases in the document will guide the development of a CADF specification and is intended to help ensure the specification meets real-world cloud auditing needs. However, during the development of the specification, the CADF WG reserves the right to choose to modify, extend, deliberately ignore, or add to the use cases contained in this document.

2 Terms and definitions

In this document, some terms have a specific meaning beyond the normal English meaning. Those terms are defined in this clause.

The terms "shall" ("required"), "shall not,", "should" ("recommended"), "should not" ("not recommended"), "may," "need not" ("not required"), "can" and "cannot" in this document are to be interpreted as described in ISO/IEC Directives, Part 2, Annex H. The terms in parenthesis are alternatives for the preceding term, for use in exceptional cases when the preceding term cannot be used for linguistic reasons. Note that ISO/IEC Directives, Part 2, Annex H specifies additional alternatives. Occurrences of such additional alternatives shall be interpreted in their normal English meaning.

The terms "clause,", "subclause,", "paragraph,", and "annex" in this document are to be interpreted as described in ISO/IEC Directives, Part 2, Clause 5.

The terms "normative" and "informative" in this document are to be interpreted as described in ISO/IEC Directives, Part 2, Clause 3. In this document, clauses, subclauses, or annexes labeled "(informative)" do not contain normative content. Notes and examples are always informative elements.

The terms defined in DSP0004, DSP0223, and DSP1001 apply to this document. The following additional terms are used in this document:

2.1

Aggregation

Aggregation refers to the combination within a single event of two or more other events (or references to those events). Aggregation is typically a bundling of separate events that preserves and keeps the original events accessible.

2.2

Control Objective

A control objective refers to a security compliance related requirement or practice. Control objectives are often abstracted statements of requirements from specific security regulations or frameworks. For example, “Separation of Duties (SoD)” is a common security control objective that focuses on the best

practice of requiring different people to perform different duties in order to provide a level of checks and balances in a system.

2.3
Event Consumer | Consumer

A consumer of events is an entity that needs to process, report on, or otherwise use CADF events.

2.4
Event Provider

An event provider is an entity that can produce events in a CADF event format.

2.5
Filtering

Filtering refers to the process of reducing the events that are returned in a query. This filtering is based on the filter parameters within the query.

2.6
Geolocation | Geo-location

Geolocation refers to the identification of the geographical location of a resource or entity related to an event. The identification of the physical location of a resource or player is important from a legal compliance perspective to ensure or audit compliance with the laws of various countries, regions, or logical boundaries that dictate where information must be stored.

2.7
Geo-routing

Geo-routing refers to the geographical tracking of an event from its origin through the various resources that participated in the event or the handling an event.

2.8
Summarization

Summarization refers to the consolidation of multiple similar or identical events in to a single event, typically for storage, bandwidth, or other optimization purposes. Summarization is typically destructive of the original events, as opposed to aggregation, which preserves the original events.

2.9
Suppression

Suppression refers to the dropping/elimination of events from an event stream or event store. From an auditing perspective, the entity that drops the events will typically create a “meta” event indicating the count and type of event being dropped. From a semantic perspective, suppression refers to events that have been removed from an event store, and not from a query result set. This differs from the concept of filtering, which refers to removing events from a result set returned from a query.

3 Symbols and abbreviated terms

The abbreviations defined in DSP0004, DSP0223, and DSP1001 apply to this document. The following additional abbreviations are used in this document.

3.1
Access Control List
ACL
A security object that lists entities that have various access rights to a given resource.

3.2
Cloud Management Working Group
CMWG
The CMWG is a DMTF working group.

3.3
Identity and Access Management
IAM
<abbrev. term definition>

3.4
Service License Agreement
SLA
<abbrev. term definition>

3.5
Virtual Machine
VM
<abbrev. term definition>

4 References

The following non-normative references are used by this white paper:

Tag	Reference
[IDCloud-CN]	OASIS Committee Note (CN), Identity in the Cloud TC, Rutkowski, et al. "Identity in the Cloud Use Cases Version 1.0", May 08, 2012, http://docs.oasis-open.org/id-cloud/IDCloud-usecases/v1.0/cn01/IDCloud-usecases-v1.0-cn01.pdf
[DSP-IS0301]	DMTF White Paper, Software License Management (SLM) Incubator, Version 0.9.2, Work In Progress Draft, December 9, 2011, http://dmtf.org/sites/default/files/standards/documents/DSP-IS0301_1.0.0a_0.pdf
[CiscoVPN]	Cisco Systems, Cisco Security Appliance Command Line Configuration Guide, Version 8.0, 2009, http://www.cisco.com/en/US/docs/security/asa/asa80/configuration/guide/asa80cfg.pdf
[CiscoLogs]	Cisco Systems, Cisco Security Appliance System Log Messages, Version 7.0, 2005, http://www.cisco.com/en/US/docs/security/asa/asa70/system/message/asa_msgs.pdf
[IPSecSimpl]	Peter J. Welcher, web article, "IPsec Simplified", January 8, 2001, http://www.netcraftsmen.net/resources/archived-articles/446.html

5 Use cases by category

This white paper contains use cases developed to exhibit data, resource, and interaction requirements for the following audit categories:

Table 1 - Granular audit event use case categories

Category	Demonstrates Audit Requirements for:
Binary and Metadata	Consider the need for inclusion of binary and meta data within the event (format)
Compliance Control Based	Explore representation of common Control Based auditing frameworks (e.g., HIPAA, PCI DSS, COBIT, etc.)
Correlation	Correlating related events that span service, infrastructure and deployment boundaries (e.g., monetary transactions, network routes, etc.)
Data Tagging	Tagging events with domain and non-domain based classification values to achieve custom reports/views
Informational Events	Multiple language considerations, character encoding needs
Location Based	Representing physical location of event resources (e.g., geo or regional) and representation of location data
Network	Representation of network events and their characteristics, such as location and protocol representations
Operational	Treatment, demarcation, and representation of audit report (event) data filtered by query parameters
Obfuscation	Treatment, demarcation, and representation of Personally Identifiable Information (PII)
Report Chaining	Tracking and identifying the resources that create, modify, or surface auditable events
Security	Exhibit the needs of security related events; e.g., normalized representation of identity, tokens, policy, etc.
Service Level Agreement (SLA)	Representation of SLA monitoring events, their metrics, and rule representations
Software License Management (SLM)	Representation of SLM monitoring events, their metrics, and rule representations
Signature	The need to sign audit information at various granularities (e.g., event, report, and log level)
Summarizing	Treatment, demarcation, and representation of repeated events that are collapsed into a single event for reporting (for compactness)
Temporal	Attributing event actions with time-based information (e.g., granularity of measurement observed time, best time, modification time, etc.)

5.1 Binary data

The following use cases consider the need for inclusion of binary data within the event (format).

5.1.1 Binary data as an element or property of an event

5.1.1.1 Description

A consumer of certain events needs to be able to obtain certain properties or elements of the event, which properties or elements are inherently binary in nature.

For example:

1) An anti-virus product emits an event that includes a virus signature (and/or an identifier for the virus signature).
2) An IDS/IPS system emits an event that includes an attack signature (and/or an identifier for the attach signature).
3) An “IT Screen Recorder” emits a package of data that contains the data necessary to “replay” a recorded session.

5.1.1.2 Requirements and considerations

- The CADF event format should support optional binary properties or attributes for an event.
- An entity that queries for events should be able to receive and process events that contain binary data.
- Inclusion of binary data within an event can significantly increase the average size of an event. Hence, considerations related to storage, bandwidth, processing performance, etc., and/or recommendations related to when binary data should/should not be included in an event may need to be addressed.
- Entities other than the original event reporter can add binary data to an event.
- Binary data should be able to be included in an event via reference.

5.1.1.3 Assumptions

- Binary data will typically not be involved as a field that can be queried.

5.1.1.4 Event classification data

Reporter	Initiator	Action	Target	Outcome
Anti-virus	AV Software	Virus Detected	Platform, files, or other resources affected	Repaired Not Repaired Unknown
IDS/IPS system	IDS/IPS System	Attack Detected	The resources under attack	Blocked Not Blocked Unknown
IT Recorder	IT Recorder	Session Record	The system and session recorded	Success/Failure Partial Success

5.2 Compliance control based

Explore representation of common control based auditing frameworks (e.g., HIPAA, PCI DSS, COBIT, etc.).

5.2.1 Supporting security control requirements (PCI DSS and COBIT)

5.2.1.1 Description

A certain consumer of events is primarily interested in using CADF event data for insuring or auditing compliance with certain security control objectives, such as those found within the PCI DSS or COBIT security frameworks. This consumer desires events to have been identified as to their relevance to the standard/framework and/or specific sections or subsections of the standard or framework and/or identification of a control objective of a meta security framework that abstracts the control objectives of multiple security standards or frameworks.

5.2.1.2 Requirements and considerations

- There should be a mechanism to associate audit events to specific compliance frameworks (domains) and controls for which they may be an indicator of compliance.
- Provide a means to associate events with compliance controls that are part of compliance standard frameworks without using a "tag".
- The approach to security compliance differs greatly from industry to industry and even from enterprise to enterprise. Consideration must be given to the fact that a single event may not be able to be definitely mapped to a fixed set of security control objectives in all cases.
- Use of data "tagging" may not be the best means to convey adherence to a compliance standard if mixed with other tags that are free form or proprietary unless tags that referenced such agreed upon standards were given a special classification to differentiate them from non-standard tags.
- There is an assumption that there may be more prescriptive and structured data elements or attributes that should be developed to attach to events that go beyond tagging.
- We will need to determine the following questions about tags:
  - What does a tag look like?
  - At what levels can tags be applied? Event level? Can they be applied to hosts, etc.?
  - How do queries work with tags or other compliance objective mechanisms?
- Are there other mechanisms for achieving this requirement other than tagging? Could this be determined by the intersection of taxonomies, privileged users, and critical system lists?

5.2.1.3 Assumptions

None

5.2.1.4 Event classification data

None

5.3 Correlation

The following use cases demonstrate the need to correlate related events that span service, infrastructure, and deployment boundaries (e.g., resource management interfaces monetary transactions, network routes, etc.).

5.3.1 Correlating similarities

5.3.1.1 Description

An event consumer with events collected from a wide variety of different event providers wishes to analyze that data and track the activity by/on specific resources – hosts, users, files, etc.

5.3.1.2 Requirements and considerations

- Consumer should be able to query for “activity caused by resource X” or “activity affecting resource Y”. Reports should be similarly capable.
- The key factor for this use case is that resource identifiers must be presented consistently, and this must be true regardless of where the event record is generated, how it is delivered, and so forth. The following notes list some common issues that should be considered:

- ID vs. name: roughly speaking, many IT resources have an internal, machine-readable identifier and a more human-consumable name (IP/hostname, user id/username, etc.). If some systems present the ID and some present the name, correlation across events becomes impossible without external referential information.
- Data presentation: in many cases, data can be presented in several different forms (IP addresses can be in dotted-quad/hex/binary, network/host order, etc.). For each common data type, we must define the standard form (for example, we can always use the standard IPv6 format, even for IPv4 addresses).
- Name-spacing: In many cases, resource IDs and names are not unique, sometimes not even across a single system. Event should always include namespace information to ensure unique identification of a particular resource is possible.
- Examples: usernames on a Linux box should indicate the local host as the namespace; LDAP directories should indicate the container, database table names should indicate the database (and host) as namespace.
- Considerations: Many namespaces are hierarchical – will we need to “unwrap” the entire path, or just treat the full path as a single namespace identifier?
- Completeness: In many cases the observer does not have all the relevant information about a particular resource, but every effort should be made to include enough data to uniquely identify the resource.

5.3.1.3 Assumptions

None

5.3.1.4 Event classification data

None

5.3.1.5 Classification notes

None

5.3.2 Grouping

5.3.2.1 Description

A consumer is interested in tracking activity on a busy database server. Because there are several users on the system simultaneously, and in many cases they are modifying the same tables, the consumer needs ways to distinguish one user's activity from another's. Because the database is front-ended by a website and a proxy account is used, the username is not sufficient.

5.3.2.2 Requirements and considerations

- There are a number of ways to indicate that a set of events is related as part of a single transaction:
  - Provide a transaction ID that is referenced in all related events
  - If the events are identical except for a small set of attributes, collapse them into a single event with an array for the varying attributes.
- I like to distinguish between transactions that take place at a single level, e.g., within a single process, and transactions that take place across different processes, like a client-server app. This use case is focused on the former.
- Here is an example: User X writes a complex set of data to the database that affects multiple tables. We can either:

399 – { event: {action: { transactionID: “123” }, { target: { database1, table1 }}}
400 – { event:{ transactionID: “123” }, { target: { database1, table2 }}}
401 – { event: { transactionID: “123” }, { target: { database1, table3 }}}
402 • Or, we could:
403 – { event: { target: { database1, [table1, table2, table3] }}}
404 • The second option is obviously more compact and does not require a new transactionID field, but
405 obviously the action has to be identical in all respects except for the target table name (must all be
406 writes, must all be successful, etc.). If that is not the case, we will be forced into the first option.
407 • This transaction identifier or event grouping is something that should be generated by a single
408 observer, as part of a single process. In other words, the interaction scenario is something like this:
409 “Hey, I just received a request to perform (some complicated action). This will require several small
410 sub-actions, so I will generate a transaction ID, perform each action, and then inject that ID into each
411 event.” The idea is that this transaction ID is not globally unique but is tied to the observer and
412 possibly even to the exact process ID from which the event was generated.

413 #### 5.3.2.3 Assumptions

414 • The consumer would like to see that several related events are correlated in some way as part of the
415 same transaction.
416 • The relevant correlation here has to do with associating multiple events from the database audit trail
417 together to show that they are part of the same transaction. Not covered is how to correlate that
418 transaction with the request from the web front-end.

419 #### 5.3.2.4 Event classification data

420 #### 5.3.2.5 Classification notes

Notes:

- Other data needed: Transaction ID (May be related to the action component)

421 ### 5.3.3 Correlation of a Cloud management API request from authorization to resource
422 modification

423 #### 5.3.3.1 Description

424 A certain consumer of events wishes to be able to correlate events through the vertical layers of a cloud
425 infrastructure that are a result of a single external action (transaction) while maintaining the reporting
426 chain and unique information items related to each step (i.e., events generated at each step).

427 #### 5.3.3.2 Requirements and considerations

428 A very generic pattern for "correlation" is when a remote request goes through several processing layers:

429 1. Authentication/authorization
430 2. Cloud management API operation
431 3. Cloud resource state change/modification

432 In this use case, the event is a remote access to Cloud management API for a management operation
433 (e.g., start/stop virtual servers). The initiator is an end-user (or a client application) that generates the
434 management request. In terms of the pattern described above:

1. The request is first authorized/authenticated by security/ACL module. The access event is reported by the request handling layer as the authorized request moves along to the Cloud management API.
2. The Cloud management API (as defined by CMWG) reports that the management request went through.
3. The actual resource targeted gets operated or modified by the request when completed. This gets reported as a resource modification.

So this sequence could be seen as separate events logged in different logs - but there is such a strong correlation between them that they could be viewed as a consolidated event with some means to associate them.

5.3.3.3 Assumptions

- The various processing layers are able to propagate a request to the next layer of infrastructure, log an event at each layer, and correlate them.

5.3.3.4 Event classification data

Reporter	Initiator	Action	Target	Outcome
Protocol / API Request handler (e.g., an HTTP request handler)	Client-side software or end-user	Any (e.g., Start VM)	Virtual Server	Success/ Failure
Cloud management module (e.g., a cloud based service / web service endpoint)	Authorized Account/ User (ID)			
Target cloud resource (e.g., a virtual server)	Authorization Token/ Identity Token			

448 #### 5.3.3.5 Classification notes

Reporter Notes: • Each reporter is processing/adding some information to the event (User ID/credentials, resource management operation, resource status). • The client may include some identifier of the initiator (e.g., a client, user or account ID (along with some information about the authorization outcome).
Initiator Notes: • Either an end-user or a client application.
Action Notes: • Any operation available for cloud management.
Target Notes: • A cloud resource as defined by CMWG.
Timestamp Notes: • The various reporters can timestamp their reporting. The actual request processing may last some time especially for the last reporter (cloud resource).
Compliance Area: • Security - Administration or management of cloud resources.
Tags / Tag Description: • "Access Management" since correlated event represents access to a cloud resource.

449 ## 5.4 Data tagging

450 ### 5.4.1 Supporting security control requirements (PCI DSS and COBIT)

451 #### 5.4.1.1 Description

452 A cloud provider needs to show that they enforce PCI DSS v2.0 Control Requirement #4 'Encrypt
453 Transmission of Cardholder Data across open public networks'.

454 The provider determines that this PCI DSS control requirement is supported by several COBIT Control
455 Objectives:

456 • DS5.8 Cryptographic key management

457 • DS5.10 Network security

458 • DS11.6 Security requirements for data management

459 • DS5.9 Malicious software prevention, detection and correction

460 • PO8.3 Development and acquisition standards

461 #### 5.4.1.2 Requirements and considerations

462 • A provider needs to be able to tag all applicable events in their infrastructure to show governance of
463 both these compliance standards in order to report these events to their tenant customers.

464 • A provider needs to be able to tag compliance events by control standard (i.e., PCI DSS or COBIT).

- A provider needs to be able to obfuscate any cardholder data that is considered Personal Privacy Information.
- Acquisitions of cardholder data need to be able to be tracked from entry (e.g., from an application or endpoint) to when it is securely stored.
- Cardholder data is securely stored and securely managed.
- Cardholder data can be tracked or correlated for network encryption and transmission.
- Use of data "tagging" may not be the best means to convey adherence to a compliance standard if mixed with other tags that are free form or proprietary unless tags that referenced such agreed upon standards were given a special classification to differentiate them from non-standard tags.

5.4.1.3 Assumptions

None

5.4.1.4 Event classification data

Reporter	Initiator	Action	Target	Outcome
Various	Various	Various	"Cardholder Data"	Any

5.4.1.5 Classification notes

Action Notes:

- All actions (including data reads) that target resources that manage cardholder data must raise events and be logged.

Target Notes:

- The target would be any logical resource that manages/handles "Cardholder Data".

Outcome Notes:

- All interactions with cardholder data are reported regardless of success or failure

Compliance Area:

- Compliance (Security, Industry)

Tags / Tag Description:

- PCI DSS, COBIT
- Identify domain of owning control standard and version
- Identity perhaps could be established by using a URI that identifies the control standard (namespace), version
- The URI may also include the control objective name/value <or> this may be represented as separate attributed value.
  - e.g., <tag type="control objective" domain="//pcidss.org/v2.0/" control="4.0">, or <tag control="//pcidss.org/v2.0/control/4.0">,
- The tag may need a "type" such as "control objective" to differentiate this tag type from others for filtering and parsing purposes.

Additional Data:

- Some means to correlate or identify cardholder data without disclosing Personally Identifiable Information (PII).

Notes:

- The Cloud Security Alliance (CSA) namespaces may be used to identify the compliance "domain" and "control objective".
- It would be possible to "tag" any PII data as such.
- The cloud provider may choose to "tag" all applicable controls using one or more standards (e.g., COBIT) and then later map these to other compliance standards (e.g., such as PCI DSS). This "cross mapping" (between compliance control frameworks) could be part of (and performed during) the "query" of audit events (as described in the CADF spec.)
- *** The cloud provider COULD publish its own security and compliance policies (that conform to SLAs) and provide events that are tagged with an identifier they publish so that their customers can create reports to verify SLA compliance.

5.4.2 Consumer cloud application tags business events (Process based)

5.4.2.1 Description

A company that hosts an application on a public cloud uses the cloud provider's platform services to generate audit events from their application with "tags" that prove compliance to the company's business and operational policies.

5.4.2.2 Requirements and considerations

None

5.4.2.3 Assumptions

- The cloud provider makes available an interface (method) and service that enables the cloud consumer to generate auditable events from their cloud based applications.
- Tagging of this nature would be done by the consumer application at event generation time.

5.4.2.4 Event classification data

None

5.4.2.5 Classification notes

Reporter Notes: • Reporter would be the cloud consumer's application or service.
Initiator Notes: • The initiator would be a human user or service entity that interacts with the cloud consumer's application or service.
Action Notes: • Any action the cloud consumer needs to audit to reflect their compliance policies.
Target Notes: • Any resource object which is meaningful to the cloud consumer's application or service.
Outcome Notes: • Any outcome deemed interesting to the cloud consumer's compliance policies.
Compliance Area: • Compliance (Security, Industry, Regulatory, etc.)
Tags/Tag Description: • Identify domain of "consumer" organization and the compliance policy (and perhaps versions) that is unique and meaningful to that consumer. • These policies may be application or service specific and this may need to be reflected in the identifier. • Identity could perhaps be accomplished by using a URI that identifies the control standard (namespace), version. • e.g., <tag type="consumer" domain="//mycompany.com/business/policy/A99.10"> • The tag may need a "type" such as "consumer" or to differentiate this tag type from others for filtering and parsing purposes.

5.5 International data in events

5.5.1 Configuring audit reports for different consumer locales (Globalization)

5.5.1.1 Description

A consumer company has locations in multiple countries (e.g., US and France). The auditor in one location wants to run a report on all company-specific user activity in the cloud environment. Different cloud systems and services are configured by using different locales and the event data being recorded has locale specific data. The auditor wants to see the report output that uses his preferred locale.

5.5.1.2 Requirements and considerations

- The output report should use the preferred locale to display information about the events. This is really a function of the service or application displaying the report.
- All event fields will need to be able to support international content.

5.5.1.3 Assumptions

- Any locale specific data in the events will be displayed in the locale used to record the data. This data can include resource names and event description data that is mapped from raw event data.
- All metadata having to do with classification taxonomies will not be translated.
- The taxonomies we define will likely either be numeric or else English “codes”. The display of a translated display string for a taxonomy value will outside the scope of the standard.
- Data in an event (such as the raw event data) will not generally be translated. The idea behind the taxonomies, etc., is so that the raw event really would not need to be referenced.

5.5.1.4 Event classification data

None

5.5.1.5 Classification notes

None

5.6 Location based

5.6.1 Control of data geolocation

5.6.1.1 Description

A consumer wishes to audit:

- location of data in transit
- location where data is exposed and executed
- locations through which data is routed
- location where data can be stored

5.6.1.2 Requirements and considerations

- The data must be able to be classified based on the jurisdictional constraints.

- Each geolocation must be evaluated for the data-specific constraints and entitlements that apply to the region or domain. (Possibly outside the realm of the standard, because the standard does not deal with policies).
- Data, including virtual machines, must be classified in such a way that regulatory constraints may be applied.
- Telecommunications and networking infrastructure must be capable of routing and constraining the transport of data based on categorization and policy. (Probably outside the realm of the standard).
- Based on predefined user criteria, regulatory and routing constraints may be overridden under the consent of the data owner. (Probably outside the realm of the standard. The standard can only provide information to report or audit compliance. It cannot enforce policy).
- Based on the jurisdictional routing and permissions enabled by the data owner and individual lawfully empowered to enforce the laws of a jurisdiction may lawfully intercept a data element within residence of their jurisdiction or in transit through their jurisdiction. It is understood that a data owner by permitting transport or instantiation had agreed to be legally subject to the laws of a specific jurisdiction
- Description of regulatory mandates in a machine-readable format (Probably outside the scope of the standard).
- Need to have meta-tags on an event that describe data and resources for policy enforcement.
- The standard needs to be able to support identity continuity within cloud infrastructure and across cloud deployment models for the purpose of non-repudiation of identity associated with an action permitted against security policy.

5.6.2 Assumptions

- We can only monitor (and not control data) as described by this use case.
- Legislative jurisdictions have continually changing legislative mandates that require regular policy revision.
- Data classification is necessary to ensure proper routing and handling.
- Entities need to be able to manage the location and routing path of data in transit.
- Entities need to be able to manage the creation, modification, or deletion of policies that govern access to data based on geo-location.
- Entities need to be able to manage the routing path of data based on geo-location policy.
- Entities need to be able to manage the computation and execution of code based on geo-location policy.

5.6.3 Classification notes

- The Data Owner, as described in these use cases, is assumed to be legally obligated and entitled to control the data based on legislative jurisdiction.
- Geo-location of the initiator must link to policy domain.

5.6.4 Administrator: Geo-location of events and resources

5.6.4.1 Description

A consumer of events wishes to be able to report on the geographical location of certain resources (including data) related to the event. This requirement includes data that is in transit, in storage, or being processed.

5.6.4.2 Requirements and assumptions

This use case drives a requirement to allow the association of geographical location data with an event resource or the event itself.

5.6.4.3 Assumptions

Geographical information is optional.

5.6.4.4 Event classification data

Data Location	Reporter	Initiator	Target	Outcome
In Storage	Location Service	Data Steward	Partner	GPS Coordinate
In Transit	Routing Service	Network Provider	Partner	GPS Coordinate
In Process	Hypervisor	Data Steward	Partner	GPS Coordinate

5.6.4.5 Classification notes

None

Classification notes:

- A "Data Owner" represents a logical data record that contains information about a partner (external to the cloud provider) that provides services to the provider and/or its consumers (customers). Partner information may include security information, such as its identity, Endpoints/URLs, Physical Address, Location, Certificates, (Web) Services, Security Policies, Protocols, etc.
- A "Network" represents a logical data path through which the data transits and? is subject to a set of legal constraints

Tags / Tag Description:

- Category Tag: "Geolocation"

5.7 Network

5.7.1 Description

An event consumer wishes to be able to query for events and/or run reports that distinguish between “inbound” and “outbound” connections for network devices.

In addition, a consumer wishes to be able to identify the network protocol implementation related to an event.

A consumer of events wishes to be able to target events in a query related to a given type of network resource.

5.7.2 Requirements and considerations

- The standard needs to allow classification of events (perhaps through tagging, or through an event class hierarchy) as inbound or outbound, and support the query of events based on this classification.
- There may be a need to have "tags" to reflect specific network protocol implementations at various levels of the IP stack.

589 • We will need to make sure the various network entities and abstractions are represented in the
590 resource taxonomy.

591 ### 5.7.3 Local terms

592 The following terminology is provided as background information only for Network use cases [CiscoVPN]
593 [CiscoLogs] [IPSecSimpl]:

594 **Network security appliances**

595 Rely upon named profiles to manage VPNs or "tunnel connections". These profiles contain connection
596 policies that determine which security protocols are used for a connection and which servers they should
597 use to authenticate and account for users.

598 **[Network] Connection profiles**

599 A connection profile consists of a set of records that determines tunnel connection policies. These
600 records identify the servers to which the tunnel user is authenticated, as well as the accounting servers, if
601 any, to which connection information is sent. They also identify a default group policy for the connection,
602 and they contain protocol-specific connection parameters. Connection profiles include a small number of
603 attributes that pertain to creating the tunnel itself. Connection profiles include a pointer to a group policy
604 that defines user-oriented attributes.

605 **[Network] Connection profile name**

606 You specify a connection profile name when you add or edit a connection profile. The following
607 considerations apply:

608 For clients that use preshared keys to authenticate, the connection profile name is the same as the group
609 name that an IPSec client passes to the security appliance.

610 Clients that use certificates to authenticate pass this name as part of the certificate, and the security
611 appliance extracts the name from the certificate.

612 **[Network] Connection type**

613 Connection types include IPSec remote access, IPSec LAN-to-LAN, and clientless SSL VPN. A
614 connection profile can have only one connection type. These connections

615 **Network authentication, Authorization, and Accounting servers**

616 These parameters identify the server groups or lists that the security appliance uses for the following
617 purposes:

618 • Authenticating users

619 • Obtaining information about services users are authorized to access

620 • Storing accounting records

621 • Network Security Appliance Events

622 #### 5.7.3.1 Description

623 Network security appliances typically emit in the following situations:

624 • **Use Case A**: A user authentication to the network security appliance fails.

625 • **Use Case B**: A failure occurs when an administrator is removing a peer connection (an IP address
626 entry plus other data) from peer table during management of an IPSec VPN configuration.

627 • **Use Case C**: A network connection policy group for a network user is retrieved.

- **Use Case D**: When a VPN loses connection to a remote peer during an Internet Key Exchange (IKE), this typically results in a deletion of a peer connection entry from the VPN's peer table.
- **Use Case E**: A responder (request from a user or origin IP address) attempts to (request) or force a change of IPSec key (in a running network device).
- **Use Case F**: An ICMP message from an external interface is denied. Internet Control Message Protocols (ICMP) are designed to announce network errors and problems and support troubleshooting (and perhaps impact SLA compliance) on IP-based networks.
- **Use Case G**: An IPSec receives an ESP (Encapsulating Security Payload) request that an anti-replay (attack) check failed.
- **Use Case H**: An ESP message packet fails authentication.
- **Use Case I**: When Network Access Control (NAC) for a host is disabled.
- **Use Case J**: An outbound TCP connection is built.
- **Use Case K**: The security negotiation is complete for an inbound connection.
- **Use Case L**: An inbound connection, remote security access (RSA) is created.
- **Use Case M**: Remote user assigned private address (VPN).
- **Use Case N**: Network connection fingerprint created for a user based upon network factors, such as public/private IP addresses, identity group, client (host), etc. For example: User = joshia2@skynet.com has a fingerprint based upon IP Address= 128.124.58.50, Client Type: WinNT Client Application Version: 4.8.0.
- **Use Case O**: A user (e.g., john.arroyo@skynet.com) at IP = 192.143.245.178 received an unsupported transaction message.
- **Use Case P**: The network device has accepted an authentication request from a user (from an IP address) and indicated that the request has been committed.
- **Use Case Q**: A consumer of events wishes to audit Automatic NAT Detection Status. For example, a remote (external) end point is detected to be behind a NAT device; however, the provider's (internal) endpoint is NOT behind a NAT.
- **Use Case R**: A user requests disconnection from the network. For example "User claudia@skynet.com (at IP = 128.231.155.95) disconnected her session (of type: IPSecOverNat) and was connected for Duration: 0h:31m:41s"
- **Use Case S**: The network device denies or blocks a connection. For example: "Deny TCP (no connection) from 10.16.252.100/1943 to 10.18.8.49/445 flags RST on interface outside"
- **Use Case T**: The network device does a build or teardown of an ICMP/UDP/TCP connection. Examples:
  - "Teardown ICMP connection for address 199.11.1.248/79"
  - Built inbound UDP connection 43326033 for outside:10.16.252.158/1026 (10.16.252.158/1026) to inside:10.18.8.20/53 (10.18.8.20/53), which maps to user joshia@skynet.com."
  - Built outbound TCP connection 43326039 for outside:10.16.252.163/139 (10.16.252.163/139) to inside:10.18.8.20/4908 (10.18.8.20/4908)
- **Use Case U**: The network device detects a "spoof" attack. For example: "Deny IP spoof from an IP Address (e.g., 10.16.69.254) to another IP address (e.g., 10.18.8.18) from an interface internal or external to the provider."

A consumer of events wishes to be able to run reports on these events to accomplish goals such as the following:

- **Use Case A**: Track failed login attempts to a specific device, a class of network security appliances, or across all systems (not just network security appliances).
- **Use Case C**: Track the retrieval of an associated network connection policy group for a network user in order to report on policy management actions for network devices.
- **Use Case E**: Track IPSec key changes.
- **Use Case F**: Track failures at the Internet Control Message Protocol (ICMP) level.
- **Use Case K**: Audit network (peer based) connection messages.
- **Use Case M**: Track usage of private addresses via a VPN.
- **Use Case N**: Track the creation of dynamic identities (or fingerprints as they are known in networks) for compliance and security purposes.

5.7.3.2 Requirements and considerations

- **All:**
    - The CADF action taxonomy needs to be able to support the events above in its taxonomy (some of the actions may not be unique to a network security appliance)
    - The CADF event schema needs to be able to support all fields that are likely to be queried related to the events above.
    - The CADF query needs to be able to support queries based on target type.
- **Use Case E**: Key information likely needs to be obfuscated/encrypted. The CADF event format may need to support obfuscated data.
- **Use Case F**: ICMP messages will correspond to some auditable event (log) that can be surfaced through the CADF standard.

5.7.3.3 Assumptions

- **Use Case A**: The appliance may have its own IAM system with its own set of usernames and passwords. This may be true for many "appliances" used in the cloud.
- **Use Case C**: Tracking any action related to management of VPN (or tunnel) connections.

696 **5.7.3.4 Event classification data**

Use Case	Reporter	Initiator	Action	Target	Outcome
A	Network Security Appliance (NSA)	User	Authenticate	Network Security Appliance	Failure
B		[Privileged] User	Remove	Connection, IP Address	
C		User	Retrieve [Policy]	Policy	Success
D		NSA	Disconnect	[Peer] Connection	
E		User/IP Address	Modify	IPSec [Config]	Success, Failure, Unavailable
F			Receive	[Protocol] Message	Failure
G			Verify	[Protocol Message] Packet/Payload	
H			Authenticate		
I		User	Disable	Network Access Control (NAC) (on a Host)	Success, Failure
J		NSA	Complete	[TCP] Connection [Outbound]	Success
K		IP Address	Complete	Connection [Inbound]	
L			Create		
M		User IP Address (Public)	Set	User IP Address	Success, Failure
N		NSA	Fingerprint	User	Success, Failure
O		IP Address	Receive	User (record), IP Address (record)	Success
P		User IP Address	Commit	User (record)	
Q		[Connection] IP Address	[Resource] Alert	Connection, IP Address	Warning
R		User IP Address	Disconnect	Network Device, User (record)	Success/Failure
S		Network Device	Connect	Host / Endpoint	Failure (Denied)
T		Network Device	Build (Connect), Teardown (Disconnect)	ICMP/UDP/TCP Connection,IP Address	Success
U		IP Address (From)	Attack	IP Address (To)	Failure (Denied)

697 **5.7.3.5 Classification notes**

Category:

- Network – Configuration

Reporter Notes:

- Can be any type of NSA
- Use Case A: Granular to the NSA itself or some specific component of the NSA such as its Identity and Access Management.system

Initiator Notes:

- Use cases assume some identity object and/or credentials are passed to the NSA from a user
- **Use Case M**: User is identifiable (e.g., joshua@skynet.com) and has a public IP address.
- **Use Case N**: The network security appliance would initiate this fingerprinting on its own if it was able to establish a set of identifying information based upon what it had available from network messages.

Target Notes:

- **Use Case B**: There is some "Peer Table" that tracks remote IP addresses. The Peer Address is associated with a Tunnel group (name)
- **Use Case M**: Either the user record is the target of the "set" action or some IP address table,

Action Notes:

- **Use Case A**: Authenticate is a granular, (network) message level request.
- **Use Case B**: Configuration failures of network devices are highly interesting. VPN (peer connection) removal is an important action to track.
- **Use Case J, K**: The term "Complete" seems to be used consistently. This seems to be another type of "Monitoring" event.
- **Use Case M**: The term "Set" is used in networking for setting a user's (private) IP address.

Outcome Notes:

- Use Case A: Only Auth failures are interesting.
- Use Case E: Unavailable is an outcome (response) for a request to change the network appliance at a time when it cannot fulfill; retry is implied. In this situation(?), a security (encryption) change is being requested.
- Use Case O: Apparently, it is common to indicate that such network messages are successfully received and "handled" as a "success".
- Use Case R: A "Reason Code" accompanies the normal outcome (e.g., a reason code of "User Requested" would be associated to the disconnect). These reason codes are present on success as well as failure outcomes.

Tags/Tag Description:

- **All**: Network Security, Network Configuration,
- **Use Case E, F**: Network Protocol
- **Use Case R**: “User Access”

Additional Data: • **Use Case A**: User / Identity, Credentials, IP Address • **Use Case B**: Peer Table Information, Peer Table Entry Information • **Use Case D, E, F**: Origin IP Address • **Use Case E**: Metric: Duration in Milliseconds (Sometimes information on retry may be provided (e.g., from 28800 to 3600 seconds)) • **Use Case E**: IPSec (Config) reference, Peer Table Entry(s) affected. • **Use Case R**: Connection Type, Connection Duration
Notes: • **Use Case F**: RFC 792 was referenced in conjunction with this use case. • **Use Case G**: Auditing/tracking of specific port usage is an important compliance objective.

5.8 Operational

5.8.1.1 4.8.1 Data set integrity when filtering results

5.8.1.2 Description

A consumer of cloud services wishes to query for all events from a particular service, but in the interest of resource conservation, wants to filter out low-level records like data reads. At the same time, the consumer wants to ensure that there are no gaps in the data; e.g., there were no events lost during transmission through the reporter chain.

5.8.1.3 Requirements and considerations

For any stream of events from a single source, guaranteeing that all generated events arrive at the consumer will probably mean that the source will need to inject a monotonically increasing sequence ID into the event data.

If this is done on a per-reporter level, one could in fact determine whether any given reporter filtered out some set of the event data, and how many records were filtered.

Example:

- Event1: (auth event) <rep type=obs seqid=001><rep type=relay seqid=101>
- Event2: (other event) <rep type=obs seqid=002><rep type=relay seqid=102>
- Event3: (auth event) <rep type=obs seqid=003><rep type=relay seqid=103>
- Event4: (auth event) <rep type=obs seqid=005><rep type=relay seqid=104>

If we are getting events directly from the relay, we can examine this event sequence and determine that the relay sent us all the events it intended to because it incremented the sequence ID each time, with no gaps. Comparing that result with the observer, however, we can see that the relay dropped the event with seqid 004, perhaps due to some filter.

Let's say we query an aggregator asking for just auth events; if the aggregator stamps each event it delivers with a sequence ID:

- Event1: (auth event) <rep type=obs seqid=001><rep type=relay seqid=101><rep type=aggregator seqid=201>

- Event2: (auth event) <rep type=obs seqid=003><rep type=relay seqid=103><rep type=aggregator seqid=202>
- Event3: (auth event) <rep type=obs seqid=005><rep type=relay seqid=104><rep type=aggregator seqid=203>

Again, we can determine that the aggregator has filtered out an event from the relay, and some upstream component has filtered out two events from the observer.

We need to determine if this sort of analysis is useful and justifies the cost.

- Q: Is calculating and injecting a sequence ID possible for all reporters?
- Q: This is also useful as an anchor point for batch queries.
- N: This is not simple
- Q: What about rollover of sequence ID?
- Q: Also, what if the query engine provided a summary ahead of the result set, so that you can at least tell if you got all events?
- N: This would be an optional checkbox feature for only environments that require it.

Other requirements considerations:

- The preservation of completeness needs to be done within the context of the query, which will have a more narrow result set than the entire set of events.
- We also need to consider the fact that we are dealing with aggregators and need to be able to show the integrity of the query or sequence back to the source.
- The sequence ID might need to be on a per-query basis, because the result set might be different for each query. This, of course, complicates the providers need to maintain cursors.
- May need an anchor point concept instead to guarantee results are not overlapping when retrieving large result sets in chunks or for ongoing situations where all events are being retrieved (potentially in near-real-time.)

5.8.1.4 Assumptions

5.8.1.5 Event classification data

5.8.1.6 Classification notes

Reporter Notes: • Reporter may need to include a unique sequence ID.
Compliance Area: • Goal would be to preserve data completeness for compliance related queries.
Additional Data: • Sequence ID

5.8.2 Event driven collection – No event repository at service provider

5.8.2.1 Description

A service provider offers a service, but auditing of that service is only desired by certain customers. The service provider is willing to emit compliance-related (or SLA/SLM) events for those customers who

require SIEM compliance and are willing to do it in a format that is compatible with CADF events. However, they are not willing to maintain an archive of the events and/or provide a query service against that archive. They wish to be able to send the events, in near real-time, to a registered event sink for a given customer in an event-driven, non-query model.

5.8.2.2 Requirements and considerations

This use case forces us to think outside the “query-centric” model to make sure the standard works when no queries are involved.

5.8.2.3 Assumptions

None

5.8.2.4 Event classification data

None

5.8.2.5 Classification notes

None

5.9 Data privacy

5.9.1 Obfuscation for data privacy

5.9.1.1 Description

In Germany and other countries, strict privacy laws require that all displayed event data be protected to ensure that personally identifiable information (e.g., usernames, IP addresses) is not visible. In theory consumers are supposed to be able to resolve such data at some later point, if, for example, litigation becomes necessary.

A user working at a large corporation accesses a cloud-hosted database that contains health information about the symptoms and treatment of HIV. The corporation wants to monitor access to the database to ensure that it is providing the right resources, but does not want to expose itself to litigation from the user if some network administrator sees the user's name in event data and spreads rumors.

5.9.1.2 Requirements and considerations

There are three possibilities as to what we can require:

1) The reporter just deletes the sensitive information
    a) Pro: Data is obfuscated
    b) Con: Harder to interpret the event; no easy way to resolve the data if needed; cannot correlate or summarize
2) The reporter replaces the sensitive data with junk, like '*****'
    a) Pro: Data is obfuscated; can still properly interpret the event
    b) Con: No easy way to resolve the data if needed; cannot correlate or summarize
3) The reporter replaces the data with a unique token, and provides a “resolver” service that, with proper approvals, can be used to restore the original event data.
    a) Pro: Data is obfuscated; event can still be interpreted/correlated/summarized; data can be resolved later as needed with proper authorization

792 b) Con: Increased complexity as unique tokens must be generated/stored/made available

793 • Of course, we could provide multiple options that providers can select based on the expected use of
794 the data.

795 • Other considerations include the “raw” event data – how will that be obfuscated? Encrypted only to
796 be unlocked with approval?

797 • Q: What about stored data? Can anyone else store the un-obfuscated data?

798 • Q: Who defines the policy? Consumer or provider?

799 • Q: Is there a special “obfuscator” reporter class? Or just one option for “modifier”?

800 • Q: Does this affect our modeling at all? Or is this just a prescriptive best practice that we document?

801 – Need to have a way for the event itself to show that the data was obfuscated

802 – Need to have this resolution mechanism supported

803 – Need link back to raw data record

804 • Q: Scope: what about correlation across many different event streams? Do they have to obfuscate in
805 the same way (e.g., provide the same hash) to support correlation?

806 #### 5.9.1.3 Assumptions

807 • We will assume that depending on the software consuming the events to obfuscate the data is not
808 sufficient.

809 #### 5.9.1.4 Event classification data

Reporter	Initiator	Action	Target	Outcome
(cloud-hosted database)	User info (obfuscated) Client IP (obfuscated)	Any	Patient health documents	Any

5.9.1.5 Classification notes

Reporter Notes:

- A typical reporter chain for this activity might include:
  - type=observer class=database vendor=PostgreSQL product=database process=psql
  - type=relay class=service vendor=Novell product=SLES11 process=evtsvc
  - type=processor class=aggregator vendor=Amazon product=CloudAudit

Initiator Notes:

- The initiator for this activity is the user working on his corporate desktop. The initiator object would then include an “account” sub-object and also a “host” sub-object:
- class=account/security/data name=(obfuscated) domain=dc=data\O=company\OU=users
- class=host/endpoint/network name=(obfuscated) IP=(obfuscated)
- rlate: account “using” host

Target Notes:

- The target here is sensitive information, although many other targets might be similarly sensitive. This example might say:
  - target: name=HIV_SYMPTOMS class=table/relational/database/storage namespace=HEALTHDB

Compliance Area:

- Privacy Laws

Tags/Tag Description:

- NA, unless we want to tag the event in some way to indicate that it was obfuscated.

Additional Data:

- Resolver URI (perhaps?)

5.9.2 Protection of proprietary data

5.9.2.1 Description

A cloud provider must share relevant security and compliance information with its consumers, but does not wish to reveal proprietary information about the cloud infrastructure. For example, they do not want consumers to know what sort of VM technology their systems are hosted by, but at the same time they wish to share information about which VMs were started by the customer and when they were started.

In addition, events being fed through a non-production (i.e.: test system) may need to be security sanitized to allow the test system? to be as close to real-world as possible without adding security risk by exposing real-world systems information to a broader group of individuals.

5.9.2.2 Requirements and considerations

The requirements and considerations for obfuscating of proprietary data are similar to the requirements and considerations of obfuscating data for privacy purposes.

5.9.2.3 Assumptions

Data is coming from a wide variety of sources, and must be obfuscated by the cloud provider before delivery to the consumer.

5.9.2.4 Event classification data

None

5.9.2.5 Classification notes

Reporter Notes: • Aspects of the reporter(s) are likely to need to be obfuscated, for example vendor/product information about the systems hosting the cloud.
Initiator Notes: • Aspects of the initiator(s) are likely to need to be obfuscated, for example details of the account management infrastructure.
Action Notes: • Aspects of the action(s) are likely to need to be obfuscated, such as specific vendor event codes.
Target Notes: • Aspects of the target(s) are likely to need to be obfuscated, for example details of the hosting environment.
Outcome Notes: • Aspects of the outcome(s) are likely to need to be obfuscated, for example additional vendor error codes.
Metric Notes: • Aspects of the metric(s) may need to be obfuscated, for example licensing restrictions.
Compliance Area: • Security, Compliance, Privacy
Additional Data: • In this scenario, it may be sufficient for the provider to simply delete or overwrite the fields desired to be obfuscated. • It may be desirable to provide an indication of which fields the provider modified, to support override requests for the full dataset.

5.10 Query driven

5.10.1 Selecting data sets for compactness

5.10.1.1 Description

Consumer wants to fetch a set of events from a cloud provider to serve some reporting needs, but due to the requirements of the reports and/or summaries the consumer is creating, does not need to get all the detailed event data – only the most important fields. In particular, the consumer may want to:

- Ask “simple” questions about the set of events based on high-level classifications of data that cut across any product, not on vendor-specific data (example: “show me all logins”).
- Just get summary counts based on a small tuple of common data (“how many times has each user logged into each system?”).
- Satisfy common regulatory requirements such as PCI: “Record at least the following audit trail entries for all system components for each event:

User identification, Type of event, Date and time, Success or failure indication, Origination of event, Identity or name of affected data, system component, or resource”

In particular, the consumer wants to conserve bandwidth and processing load, perhaps because:

- Connection to cloud provider is a slow link
- Set of events is large (logins for a global enterprise, for example)
- Report generation must be quick

5.10.1.2 Requirements and considerations

- Consumer wants to be able to issue a simple query using a standard interface, such as REST; no particular client implementation should be assumed.
- Consumer does not want to have to construct a complex query to specify exactly which data structures to include in the result set; consumer wants to get the “standard” data that is usually required by auditors, management, etc.
  - This use case requires methods for the query API to “select” a set of output fields. This could be implemented in any number of ways, from SQL-like 'SELECT' statements to simple query flags.
  - Unless we propose to support arbitrary SQL-like syntax (and even if we do), this use case would seem to indicate that some form of best practice or recommendation of which fields should be included at different query “levels” is necessary. This best practice could take the form of simple documentation, or could be implemented as an explicit “field profile” that the customer selects as part of the query.
  - The concept of a “field profile” is particularly attractive, because such profiles could then be tagged with some sense of the use cases they implement – such as 'PCI' or 'NIST SP 800-53'. CADF could define some basic profiles, but downstream reporting systems could then define custom profiles that specify the data they need in the form of a profile, which could be passed upstream for greater efficiency.

5.10.1.3 Assumptions

None

5.10.1.4 Event classification data

This use case does not necessarily require the definition of any new fields within the CADF data format

871 #### 5.10.1.5 Classification notes

Reporter Notes: • **Critical**: The critical reporter data would include basic identity information about the Observer. • **Important**: Important data would include the rest of the reporter chain. • **All**: Additional data would include details about any event data modifications made by mid-stream reporters, plus any vendor extensions.
Initiator Notes: • **Critical**: Critical initiator data would include the identity of the most proximate initiator, including name/ID and namespace information. • **Important**: Important information about the initiator would include contextual information about the initiator – host environment, group/role/access information, etc. • **All**: Additional data could include more detailed relationship data, plus any vendor extensions.
Action Notes: • **Critical**: Critical information about the action could include all the high-level classification, plus a vendor-supplied display message. Also timestamps. • **Important**: NA • **All**: Additional data could include any vendor extensions
Target Notes: • **Critical**: Critical data about the target could include the identity of the primary target, including name/ID and namespace information. • **Important:** Important information about the target would include contextual information about the target – host environment, group/role/access information, etc. • **All:** Additional data could include more detailed relationship data, plus any vendor extensions.
Outcome Notes: • **Critical**: Critical information about the outcome could include the high-level classification of outcome and result. • **Important**: Important information about the outcome could include vendor code and message. • **All**: Additional data could include any further vendor extensions.
Tags / Tag Description: • Tags and their values may be subject to selection depending on the level of information requested on a query. Certain tags could be classified as being of higher order consideration for selection purposes.

872 ## 5.11 Reporter chain auditing

873 #### 5.11.1.1 Description

874 A company wishes to offer event collection/aggregation services from multiple other service providers.
875 This company queries and collects events, perhaps performing some value-added processing on them,
876 and then makes them available for query to another aggregator and/or to the ultimate service, which is
877 used by a customer to perform compliance reporting.

From a compliance auditing perspective, a compliance auditor who is reviewing compliance reports needs to be able to have an audit trail that shows who handled (and/or processed any given event that appears in a report). It is important that there is enough meta-data maintained in the event that the auditor is satisfied that they can forensically determine the complete path travelled by an event before it arrived in a report.

Key information that the auditor may wish to review includes:

- The entities that handled, processed, or otherwise had access to an event
- The date and time range of that access
- Information about which physical systems collected the event (host, IP, etc.)
- Details related to what information was changed, or supplemented by each reporter who touches the event, including, where possible, the original copy of the event before it was processed by a reporter

5.11.1.2 Requirements and considerations

None

5.11.1.3 Assumptions

None

5.11.1.4 Event classification data

None

5.11.1.5 Event classification notes

None

5.12 Related event correlation

5.12.1 Related event correlation

5.12.1.1 Description

A consumer of events wishes to be able to report on various low-level events that are closely related to each other as if they were a single event.

For example, a remote request could go through several process layers such as:

- Authentication/Authorization
- Cloud Management API Operation
- Cloud Resource State Change/Modification (i.e.: the event is related to a remote access for a cloud management API for an operation, such as start/stop a virtual server)

In the example, the initiator is an end user (or client application) that generates a management request.

- The request is first authorized/authenticated by a security/ACL module. The access event is reported by the request handling layer as the authorized request moves along to the Cloud Management API.
- The Cloud Management API (as defined by CMWG) reports the management request went through.

- The actual Resource Targeted gets operated/modified by the request when requested. This <result | process?> gets reported as a resource modification.

These could be seen as separate events logged in different logs, but the strong correlation between the events creates a need at the event consumer level to report on these as a single action.

5.12.1.2 Requirements and considerations

None

5.12.1.3 Assumptions

None

5.12.1.4 Event classification data

Reporter	Initiator	Action	Target
1. Request handler is the initial "observer" (reporting: initiator ID, authorization outcome) 2. Cloud management module (reporting: API operation, initial outcome) 3. The target cloud resource (reporting: detailed outcome of operation – e.g., in case it lasts for some time)	Client-side software or end-user	Start	Virtual Server

5.13 Security

5.13.1 Categorizations

Security use cases, in this section of the white paper, are categorized based upon a taxonomy described within the OASIS "Identity in the Cloud Use Cases Version 1.0" document.

5.13.2 Challenges

- How do we proscribe auditors query these types of security events as described by these use cases
- It seems that most security events of interest are assumed to be tied to a particular account or a particular security object linked to the account. Perhaps there is a pattern where some of these data elements would be required.
- It also seems that when these security events are queried using an interface into the cloud provider that there would be some indicator on the query or the protocol that provides identification of the account and the required security credentials to perform the query.

5.13.3 General notes

5.13.3.1 General identity and access manager functions

- Manage the creation, modification, and termination of user privileges, user groups, and roles throughout the entire user or entity lifecycle.
- Manage the creation, modification, or deletion of policies that govern access to users, groups and roles, and resources.

- Manage roles, accounts, group membership, and passwords.
- Role and group management provides the ability to add, remove, or change attributes.
- Help govern user access to services and resources against user rights, privileges, and credentials (often provided by group assignment or role attribution).

5.13.3.2 Data format field notes

The CADF will not address the defining of roles or ACLs in its specification work.

5.13.3.3 Reporter notes

- The Identity and Access Manager (or provider), as described in these use cases, are assumed to be part of the cloud provider's management platform; however, if the provider uses Federated IdM, these functions could be external to the cloud (perhaps third-party provider(s) that both the provider and consumer recognize and have a trust relationship established with).
- The consumer (i.e., the tenant business, customer, etc.) could act as its own identity and access management provider or reference a third-party provider that can be trusted by the cloud provider to establish an identity federation.
- Access Managers often are used to manage policies and their rules and provide evaluation of those rules against the identities and attributes (e.g., roles) presented to determine access (grant or deny).
- The provider's Identity and Access Manager may be distinct (separate) from the one that is used to manage consumer identities.

5.13.3.4 Initiator notes

- Provider Administrator is a type of "privileged user" within the provider infrastructure and may affect multiple consumer accounts.
- Consumer Administrator is a type of "privileged user" for a particular account.
- Self-service administration in clouds is common. For example, a consumer administrator may use a web portal to add, remove, and modify account services themselves.

5.13.4 Infrastructure trust establishment

This use case features establishment of trust between cloud providers their partners and customers and includes consideration of topics such as certificate services (e.g., x.509), signature validation, transaction validation, non-repudiation, etc.

5.13.4.1 Description

A consumer of events wishes to be able to report on actions taken by cross-vendor Identity Management systems. These actions include the following:

- Consumer Administrator Create Partner
- Consumer Administrator Create [Delete, Update] Federation
- Consumer Administrator Certificate Management
    - A cloud consumer wants to create and send an X.509 certificate for use by one its suppliers. The certificate's keys will be used to sign messages that arrive into one of its cloud-hosted applications. A secure trust relationship between a cloud consumer and one its business partners is rooted by the creation and exchange of unique certificates that contain identifiers and cryptographic keys that can be used to establish further credentials.
- Provider Administrator Certificate Management

- A provider administrator receives a request from a partner service provider for an x.509 certificate that can be used to exchange security keys that they can use to secure messages (and process, workflows, etc.) between them. A secure trust relationship between the provider and a consumer or partner is rooted by the creation and exchange of unique certificates that contain identifiers and cryptographic keys that can be used to establish further credentials.

5.13.4.2 Requirements and considerations

CADF action taxonomies need to consider these events.

5.13.4.3 Assumptions

None

5.13.4.4 Event classification data

Reporter	Initiator	Action	Target	Outcome
Identity Manager	Consumer Administrator	Create	Partner	Success/Failure
		Create [Delete, Update]	Federation	
		Add, Create, Remove, Request, Send, Receive	Certificate	
	Provider Administrator	Add, Create, Remove, Request, Send, Receive		

5.13.4.5 Classification notes

A "Partner" represents a logical data record that contains information about a partner (external to the cloud provider) that provides services to the provider and or its consumers (customers). Partner information may include security information, such as its identity, endpoints/URLs, physical address, location, certificates, (Web) services, security policies, protocols, etc.

A "Federation" represents a logical data record that contains information about an identity representation that spans deployment boundaries. For example, it could represent a person's electronic representation of identity and attributes, and how that identity may be stored or represented across multiple distinct identity management systems. See http://en.wikipedia.org/wiki/Federated_identity - cite_note-0.

5.13.5 Infrastructure identity management

This use case features virtualization, separation of identities across different IT infrastructural layers (e.g., server platform, operating system (OS), middleware, virtual machine (VM), application, etc.).

No use cases currently submitted for this category.

5.13.6 Authentication

This use case features general authentication use cases (non-SSO), as well as ones that reference Single Sign-On (SSO) patterns across cloud deployment models.

5.13.6.1 Description

A consumer of events wishes to audit the authentication activities of users. These actions include:

- A user authentication or re-authentication
- A user logon/logoff

- A privileged user (sudo) user logon
  This action includes the audit of “superuser” logon to system which has the power/privilege to perform some actions on behalf of another user and assume their identity (and hence their rights and permissions).

5.13.6.2 Requirements and considerations

5.13.6.3 Assumptions

Authentication is a distinct function from Authorization and it may be integrated into an identity and access manager service or its own standalone service or set of services.

5.13.6.4 Event classification data

Reporter	Initiator	Action	Target	Outcome
Authentication Service	User	Authenticate, Re-authenticate	User [Account]	Success/Failure
	User	Logon, Logoff	User	
	Privileged [Sudo] User	Logon	User	

1017 **5.13.6.5 Classification notes**

Reporter Notes: • A cloud provider may support multiple, protocol-specific authorization services.
Initiator Notes: • "User" is a "Person" that has presented itself as having an established identity in the system (i.e., within or recognized by the cloud provider).
Action Notes: • The notion of a "re-authentication" due to some policy (e.g., time or access based)
Target Notes: • Tracking of last authentication (time)
Compliance Area: • Security
Tags/Tag Description: • "Access Management". "Authentication"
Additional Data: • User Identity (token, identifier, etc.) • User Credentials (any presented at authentication time) • Policy references
Notes: • Authentication failure reporting is a significant aspect of this use case. • Correlation to auth. policies (or reauthentication policies)
Reporter Notes: • A cloud provider may support multiple, protocol-specific authorization services.
Initiator Notes: • "**User**" is a "Person" that has presented itself as having an established identity in the system (i.e., within or recognized by the cloud provider). • "**Superuser**" is a privileged user that has presented itself as having an established identity in the system (i.e., within or recognized by the cloud provider).
Action Notes: • Logon and Logoff both may need to be audited as a pair where possible. • There is also the notion of a system "logoff" due to timeout or some other error.
Target Notes: • User status/state may change when logged in and the user usage or connection time may be tracked (audited).
Compliance Area: • Security

Tags/Tag Description: • "Access Management". "Authentication", "Logon", "Logoff" (especially logon failures) • (Privileged user only): "Privileged Action"
Additional Data: • User Identity (token, identifier, etc.) • User Credentials (any presented at logon)
Notes: • Logon failure reporting is a significant aspect of this use case. • Tracking and correlation of SUDO (privileged user) events is highly significant for auditing security in any compliance framework.

5.13.7 Authorization

This use case features general authorization.

5.13.7.1 Description

A consumer of events wishes to audit authorization to resources that have been granted to a user.

5.13.7.2 Requirements and considerations

None

5.13.7.3 Assumptions

None

5.13.7.4 Event classification data

Reporter	Initiator	Action	Target	Outcome
Access Manager or Policy Enforcement Point (PEP)	User	Authorize	Resource [File, DB, etc.]	Success / Failure

5.13.7.5 Classification notes

Reporter Notes: • See information on "Policy Enforcement Points" and "Policy Decision Points" in the Notes section below.
Initiator Notes: • A user or entity with an identity is authorized (or not) access to a specified resource. • If the reporter is a PDP, there may be information to correlate the auth. request back to the PEP, which is a resource in the system.
Target Notes: • Resource for which the initiator is requesting authorization
Compliance Area: • Security
Tags/Tag Description: • Category Tag: "Access Management"
Notes: • "Policy Enforcement Points" (PEPs) usually sit in front of resources and present user identity representations (e.g., ID, token, etc.) to an Access Manager that acts as a "Policy Decision Point" (or PDP) that evaluates the ID against its stored attributes/privileges/roles to determine whether access is permitted to the resource in question based upon policies that govern access permissions (e.g., via rules). • Both PEPs and PDPs may act as the event reporter.

5.13.8 Account and attribute management

This category includes use cases that feature account establishment or provisioning and security policy at tributes and their management. Use cases may include descriptions of established provisioning techniques, as well as developing examples of Just-In-Time (JIT) Account Provisioning.

5.13.8.1 Provider/Consumer administrator management

5.13.8.2 Description

Some provider administrators have special privileges (perhaps via roles) to do the following:

- **Use Case A**: Manage cloud consumer accounts. Some privileged functions (actions) include create, delete, update, enable, and disable of accounts.
- **Use Case B**: Manage account level resources to his consumer users (customers). A company's consumer administrator creates and configures a compute resource that will be used to run applications by a department within their company. A provider administrator adds access to storage, network, compute and composition services to a consumer account in accordance with the service license agreement (SLA)
- **Use Case C**: A consumer account administrator is able to manage account level resources to his consumer users (customers). A company's consumer administrator creates and configures a compute resource that will be used to run applications by a department within their company.

A consumer of events wishes to track the actions of the users with special privileges as part of a security auditing function

5.13.8.3 Requirements and considerations

None

5.13.8.4 Assumptions

- **Use Case A, B**: Security policies and roles exist within the provider to distinguish this logical class of user.
- **Use Case C:** Consumer administration of resources is an account level role/function.

5.13.8.5 Event classification data

Use Case	Reporter	Initiator	Action	Target	Outcome
A	Identity or Access Manager	Provider Administrator	Create, Read, Update, Delete, Enable, Disable	Consumer Account	Success/ Failure
B			Add, Remove, Update	Service [Workflow]	
C		Consumer Administrator	Create, Read, Update, Delete	Cloud Resource [Compute]	

5.13.8.6 Classification notes

Reporter Notes:

- The provider Identity and Access Manager may be distinct (separate) from the one that is used to manage consumer identities.

Initiator Notes:

- Self-service administration in clouds is common. For example, a consumer administrator may use a web portal to add, remove, and modify account services themselves.

Action Notes:

- May be subclasses of some parent action classification.
- Consumer accounts can be enabled or disabled for various reasons (e.g., disabled due to non-payment or violation of terms).
- Add is not a “create.” Add makes the service workflow accessible (available) to an account or account group.
- Remove is not a delete. Remove implies the service is no longer accessible (available) to an account or account group.
- Cloud resource management may resolve to standard CRUD operations.

Target Notes:

- Consumer Account Directory represents some logical "store" for all consumer account information at a provider that can be implemented in many ways.
- Target can be any secure resource that has an associated security policy. In most cloud provider architectures the security policies are managed at an account level for consumers.
- Accounts can exist at various levels within a cloud IT infrastructure; here, we are focusing on cloud consumer accounts.
- Cloud consumers may have their own class of privileged users (e.g., administrators) that manage access to account resources (e.g., account licensed resources and services or hosted applications and data).
- In most cloud architectures security policies are managed at an account level for consumers. Target can be any secure resource that has an associated security policy.

Compliance Area:

- Security

Tags/Tag Description:

- **All**: "Account Management"
- **Use Case C**: “Access Management”

Additional Data:

- Consumer Account Directory (logical) (Use Case A)
- Consumer Account [or a reference to one]
- Other Contextual Data (data store, etc.) where available (Use Case A)
- Service, Service Workflow (Operational or Business) (Use Case B)

5.13.9 Identity and access management - auditing privileged user accesses to cloud hosted resources

5.13.9.1 Description

A consumer of events wishes to audit key identity management actions, including the following:

- **Use Case A**: An administrator managing user identities
- **Use Case B**: A consumer administrator managing account users
- **Use Case C**: A consumer administrator granting or revoking access to cloud hosted resources (such as access to its customer database to an authenticated cloud-based Customer Relationship Mgmt. (CRM) service)
- **Use Case D:** A consumer administrator grants or revokes an access privilege to all users belonging to a logical group of users (i.e. a policy group) such as managers, developers, testers, etc.
- **Use Case E**: A provider administrator managing consumer account privileged users
- **Use Case F**: An administrator locks a master server configuration so that other privileged users may not alter that resource until they have completed an update.
- **Use Case G**: A consumer account administrator grants access to its customer database to an authenticated cloud-based Customer Relationship Mgmt. (CRM) service which may be external to the cloud provider.
- **Use Case H**: A consumer administrator manages consumer user credentials
- **Use Case I**: A consumer administrator creates a consumer user credential group. A credential group is an administrator-defined set of domains that share the same set of access credentials. (You can think of a credential group as all the login services that use the same username and password.)

5.13.9.2 Requirements and considerations

- We need to make sure we consider third-party Identity Providers (IDPs), because identities may not be owned by the identity manager and/or cloud provider. This can include very complex federation scenarios with identity chaining from one IDP to another IDP. We must also consider identity claim tokens from different sources.

5.13.9.3 Assumptions

- **Use Case G**: The Identity Manager (or provider) here is described to part of the cloud provider's management platform; however, if the provider uses Federated IdM, the Identity Manager or Identity Provider could be external to the cloud (perhaps a third-party provider that both provider and consumer recognize and have a trust relationship with or the consumer has its own identity provider).
- **Use Case G**: The Customer Relationship Management service is a recognized (authorized) service hosted by the same cloud provider.
- **Use Case I**: A credential group may have any number of authentication mechanisms (also known as "credential group elements"). The security manager supports any number of credential groups.

1093 **5.13.9.4 Event classification data**

Use Case	Reporter	Initiator	Action	Target	Outcome
A	Identity or Access Manager	Provider or Consumer Administrator	Create, Delete, Modify, Move, Enable, Disable	User (Person)	Success / Failure
B		Consumer Administrator	Add, Remove	Account / Credential Group	
C			Grant/Revoke Access	Cloud Resource [e.g., Customer Database, CRM Service, etc.]	
D				Account/Policy Group	
E		Provider Administrator	Create, Delete, Enable, Disable, Modify	[Privileged] User	
F		Provider or Consumer Administrator	Lock, Unlock, Refresh	Configuration Repository	
G		Consumer Administrator	Grant/Revoke Access	Consumer Account	
H		Consumer Administrator	Create, Validate	Credential	
I		Consumer Administrator	Create, Modify, Refresh, Copy	Credential Group	

5.13.9.5 Classification notes

Compliance Area:

- **All**: Security
- **Use Case C, D, G**: Security-Access Management

Tags/Tag Description:

- "Access Management"

Additional Data:

- Consumer Account
- **Use Case C, G**: Access/Permission Rule, Customer Database, CRM Service
- **Use Case D**: Policy Group (e.g., an account), Name of policy group (e.g., a distinguished name), Location of policy group record, container, etc., Representation of Policy (expression standards?)
- **Use Case G**: Customer Database, CRM Service

Initiator Notes:

- **Use Case C, G**: Consumer account administrator has privileges to Grant or Revoke Access to a consumer resource.
- **Use Case D**: Administrator is a type of "privileged user".

Action Notes:

- **Use Case C, D, G**: "Grant" and "Revoke are typical verbs used to describe access control functions, but these verbs are typically accompanied by a logical object that describes "what" is being granted or revoked, in this case "Access" is the "what". Another term could be "Permission", etc.

Target Notes:

- Use Case C, G: Access could be granted to an entity (e.g., a web service) or a person (e.g., a user the consumer account acknowledges).
- Use Case D: Account is a type of policy group for cloud consumers.

Other Notes (Use Case C, G)

- Notes: There is potentially "Other" Information, such as identifying the ("On what") target resource to which access was granted.
- Access can be granted to a logical "Group" that has already been defined within the IdM component that implicitly grants access to a group of entities or users (or both).

5.13.10 Identity and access management - Auditing consumer users accesses to cloud hosted resources

5.13.10.1 Description

A consumer of events wishes to audit all levels of user accesses to all resources. This may include different actions that represent access and the actions could be dependent on the type of resource being access. This access can include:

Use Case A: A non-foreign user attempting any type of access to a resource.

Use Case B: A foreign user attempting any type of access to a resource.

- For example, a consumer account administrator is running a virtual server within a cloud provider. The administrator wishes to use SSH to connect to the server to configure it. This use case is for IaaS cloud providers (e.g., AWS, Rightscale, etc.) that permit connecting/attaching to a hosted/running application (image) server.

Use Case C: Any consumer user executing an application.

- This use case is for IaaS cloud providers (e.g., AWS, Rightscale, etc.) that permit connecting/attaching to a hosted/running application (image) server.

5.13.10.2 Assumptions

- **Use Case B, C**: Cloud provider allows access to users via "foreign" workstations (clients or applications) using some credential that can be used to authorize access to some cloud hosted resource. That is access is not from cloud provider hosted interfaces or portals and performed through some other protocol using an identity or credential that may not be coupled to a fully defined user identity (e.g., a shared admin identity to view and monitor an application hosted on the cloud).
- **Use Case B, C**: Identity Provider/Manager is able to account for "foreign" user connections.

5.13.10.3 Event classification data

Use Case	Reporter	Initiator	Action	Target	Outcome
A	Identity or Access Manager	Consumer Account User	Resource Dependent	Cloud Resource	Success/Failure
B		Foreign User	Attach, Detach, Enable, Disable	Application (Image) Server	
C	Access Manager	User	Execute	Application	

5.13.10.4 Classification notes

Additional Data (e.g., Other Stuff):

- Use Case A: User ID, Credentials (e.g., tokens, etc.)
- Use Case B, C: Need to identify (virtual) server or virtual machine

Tags/Tag Description)

- "Access Management". "Account Management"

Reporter Notes:

- Use Case B, C: The Identity Manager (or provider) here is described to part of the cloud provider's management platform; however, if the provider uses Federated IdM, the Identity Manager could be external to the cloud (perhaps a third-party provider that both provider and consumer recognize and have a trust relationship with or the consumer has its own identity provider).
- Use Case B, C: Perhaps needs correlation to some network connection where foreign user accessed cloud.

Initiator Notes:

- **Use Case B, C**: Foreign user is a concept used when handling users who use workstations that are NOT part of the local domain. Traditional examples are SSH, Telnet access, etc. Often, Access Management Systems will permit the allocation of UIDs (from some assigned pool) that keep the identity of the foreign user separate from users that are members of the domain (in this case users that are known to the cloud provider).
- **Use Case B, C**: Foreign User is a special type of user on a associated with a particular cloud consumer account

Action Notes:

- **Use Case B, C**: Need better examples or another use case to show enable/disable actions

Target Notes:

- **Use Case B, C**: Need to identify (virtual) server or virtual machine, perhaps with ID of owning account. Need to convey credentials (e.g., SSH key or token)

Outcome Notes:

- **Use Case B**: Success, Failure (with additional provider specific Information if failure)

Compliance Area:

- Security

5.13.11 Identity and attribute provisioning

On-boarding of consumer accounts, identities, roles, attributes, policies, etc.

No use cases currently submitted for this category.

5.13.12 Security tokens

This category includes use cases that feature Security Token Formats and Token Services including Token Transformation and Token Proofing.

No use cases currently submitted for this category.

5.13.13 Audit and compliance

This category includes use cases that feature Identity continuity within cloud infrastructure and across cloud deployment models for the purpose of nonrepudiation of identity associated with an action permitted against security policy.

5.13.13.1 Auditing of audit-related configurations and actions

5.13.13.2 Description

For auditing purposes, a consumer of events needs to be able to audit the configuration of audit-related changed and actions. This includes auditing of the following:

Use Case A: When a consumer administrator configures per-account and per-application audit logging. This may include things like:

- Configuring location for logs and reports
- Configuring customizable report filters
- Configuring alerts/emails

Use Case B: The start and/or stop of the service that is actually gathering or providing audit data to verify it has not been tampered with during a particular time period.

5.13.13.3 Requirements and considerations

Use Case B: The start/stop times must be normative to the event times reported in the audit reports or logs.

5.13.13.4 Assumptions

None

5.13.13.5 Event classification data

Use Case	Reporter	Initiator	Action	Target	Outcome
A	Access Manager	Consumer Admin	Set/Change	Audit Configuration	Success/Failure
B	[Cloud Management] Platform	User/Entity	Start/Stop/Pause	[Audit] Service	

5.13.13.6 Classification notes

Compliance Area: • Security Auditing
Tags/Tag Description: • **Use Case A**: Access Management • **Use Case B**: "Security Compliance", "Priority Alerts"

5.13.14 Password management

5.13.14.1 Description

5.13.14.2 Requirements and considerations

A consumer of events wishes to audit user password changes on consumer accounts.

Use Case A: Provider administrators may have the authority to change passwords for users on consumer accounts (including "privileged users" types such as a consumer account administrator). In turn, consumer account administrators may have the authority (privilege) to manage passwords for users on the same consumer account that belong to different access control groups.

Use Case B: A provider administrators may have the authority to change passwords for "Foreign User" accounts

5.13.14.3 Assumptions

None

5.13.14.4 Event classification data

Use Case	Reporter	Initiator	Action	Target	Outcome
A	Identity or Access Manager	Provider or Consumer Administrator	Change/Reset	Password	Success/Failure
B		Administrator or Entity			

1163 **5.13.14.5 Classification notes**

Reporter Notes: • **Use Case B**: Perhaps needs correlation to some network connection where foreign user accessed cloud
Initiator Notes: • Describing this auditable activity for both Provider and Consumer Administrators • **Use Case B**: Foreign user is a concept used when handling users who use workstations that are NOT part of the local domain. Traditional examples are SSH, Telnet access, etc. Often, Access Management Systems will permit the allocation of UIDs (from some assigned pool) that keep the identity of the foreign user separate from users that are members of the domain (in this case users that are known to the cloud provider). • **Use Case B**: Foreign User is a special type of user on a associated with a particular cloud consumer account
Action Notes: • None
Target Notes: • Password is a security object (resource).
Outcome Notes: • Success, Failure (with additional provider specific Information if failure) • Failure perhaps information on password policy rule not met
Tags/Tag Description: • "Access Management". "Account Management", "Password Change"
Additional Data: • **All**: “User” or “Account” (associated with password), Credentials • **All**: Password Policy • **Use Case B**: Foreign User, Credentials, Network Connection, Protocol

1164 **5.13.14.6 Assumptions**

• Both use cases apply to IaaS cloud providers (e.g., AWS, Rightscale, etc.) that permit connecting/attaching to a hosted/running application (image) server. • Cloud provider supports password management functions and password policies for either or both their provider and consumer admins.
• Cloud Provider allows access to users via "foreign" workstations (clients or applications) using some credential that can be used to authorize access to some cloud hosted resource. • Identity Provider/Manager is able to account for "foreign" user connections.

5.13.15 Policy management

5.13.15.1 Policy management use cases

5.13.15.2 Description

A consumer of events wishes to audit policy management activities, including changes to audit policies (a security object). These include auditing the following:

Use Case A: When a consumer administrator performs policy management activities.

Use Case B: When a consumer administrator performs policy rule management activities.

Use Case C: When a consumer administrator manages account roles.

Use Case D: When a provider administrator manages platform (system) services.

Use Case E: When a user or resource reads a policy file.

5.13.15.3 Requirements and considerations

Policies can be applied and managed at various levels within a consumer account (e.g., at the account level itself, application or service level, resource level, etc.)

5.13.15.4 Assumptions

None

5.13.15.5 Event classification data

Use Case	Reporter	Initiator	Action	Target	Outcome
A	Policy Manager	Consumer Administrator	Create, Modify, Delete, Activate, Deactivate	Policy	Success, Failure
B			Create, Modify, Refresh	Rule	
C			Create, Delete, Modify, Grant, Revoke	Role	
D		Provider Administrator	Add, Delete, Activate, Deactivate, Remove	Platform [System] Service	
E		User or Entity [Resource]	Read	Policy	

5.13.15.6 Classification notes

Action Notes:

- There are also requests to alter the state of a policy represented by the "Activate" and "Deactivate" actions.

Target Notes:

- The name "Policy" can represent any type of compliance policy including security policies.
- **Use Case D**: Some "Platform Services" may be managed as "groups" for example all "Storage Services" may be managed as a group by a cloud provider's administrator.

Compliance Area:

- Security

Tags/Tag Description:

- **Use Case A, C, D, E**: Policy Management
- **Use Case B**: Access Management
- **Use Case D**: "System Processes", "Platform Services"

Additional Data:

- **Use Case A, B**: Policy:
  - Policy Setting, Attribute, Rule, etc.
  - Resource Policy Applies to
- **Use Case C, D, E**: Roles can be associated with:
  - Users, Policy Groups (Accounts) and perhaps other Roles
  - These objects need to be reference-able from the event

Assumptions:

- Consumer Administrator management of policy is supported.

5.13.16 Profile Management

5.13.16.1 Consumer administrator profile management

5.13.16.2 Description

A consumer administrator is able to create, delete, or modify security profiles that are used to govern the types of security the provider.

Typically security profiles describe the security and governance required within a domain for exchange of policies, consent directives, and authorizations between entities (e.g., between provider and a partner, or between cloud hosted applications and services).

These profiles may include descriptions of acceptable security methods for confirming auditable identities, authorization status, and role attributes for entities/actors/users that interact with a cloud hosted account, application, service or workflow (as defined by the consumer).

A consumer of events wishes to audit the management of profiles that define security information/settings on resources and services.

5.13.16.3 Requirements and considerations

- These "profiles" are managed as separate objects that are deployed with cloud services or applications that define security parameters, policy references, permissions, etc.
- These profiles can be managed from cloud consumer accounts via provider supplied interfaces.
- Profiles can be treated as secured, controlled structured documents.
  - Security profiles may be embodied as standardized documents such as those defined by OASIS Cross-Enterprise Security and Privacy Authorization (XSPA) TC.
- Profiles can be identified as a cloud resource within an auditable event.

5.13.16.4 Assumptions

None

5.13.16.5 Event classification data

The following event classification data provided as examples for this use case:

Reporter	Initiator	Action	Target	Outcome
Identity or Access Manager	Consumer Administrator	Create, Modify, Delete	Profile	Success, Failure

5.13.16.6 Classification notes

Target Notes:

- Profile includes such things as "Service Profiles" and "Resource Profiles" that are used to define and govern access to cloud based applications and services.

Compliance Area:

- Security

Tags/Tag Description:

- "Profile Management", "Access Control Information Management"

5.14 Service Level Agreement (SLA)

Please see section titled "Undeveloped summarizing SLA use case idea" for an SLA related scenario.

No use cases currently submitted for this category.

5.15 Software License Management (SLM)

Please see scenarios as provided within the DMTF SLM Incubator's "Software Identification and Entitlement Metrics" white paper.

5.16 Signature

5.16.1 General notes

In general, we are "tamper proofing" of event documents (e.g., reports and logs) to a level acceptable to establish trust in the data received from the provider.

Established signing techniques for documents and messages (i.e., transmitted on interfaces) may involve the consideration of the following topics:

- Use of Ciphers/Keys and Key Lengths
- Timestamps and Nonces (variated data)
- Message Authentication Codes (MACs)
- Hashing algorithms including seeds

NOTE: The intent of these use cases is strictly to assure audit document formats are suitable for signing and not to suggest that the work group prescribe how documents be secured outside the boundary of the provider.

5.16.2 Use case 1: Cloud provider signing reports or logs for a cloud consumer

5.16.2.1 Description

A cloud consumer auditor wishes to examine and obtain a report or log and have the entirety of the document signed by the cloud provider by using an agreed upon (e.g., shared) key

Can be signed at consumer account or application/service level (i.e., using standardized signing techniques that may be tied to a specific consumer account or application)

5.16.2.2 Requirements and considerations

The primary concern is one of guaranteeing signed delivery at hand-off. We do not want to get involved in key management “six years later.”

5.16.2.3 Assumptions

- Provider Granular. That is the signing of the audit report or log is done by the cloud provider (as an entity) and not individual IT component resources or services.

5.16.2.4 Event classification data

Not applicable

5.16.2.5 Classification notes

Not applicable

5.16.3 Use Case 2: Cloud provider signing one or more events within a report or log for a cloud consumer

5.16.3.1 Description

Reports or logs may contain a mix of informational events that may not have compliance or auditing implications (not tied to any legal framework) along with those that may ties to compliance controls and auditing frameworks. This means that it may be desirable to sign individual events that

A provider of events wishes to be able to make sure individual events contained within a log or report can be independently signed.

5.16.3.2 Assumptions

- Provider Granular. That is the signing of certain events within a larger audit report or log may done by the cloud provider and appear along with other events that the provider does not wish to sign (including partner services).

5.16.3.3 Event classification data

Not applicable

5.16.3.4 Classification notes

Not applicable

5.16.4 Use Case 3: Cloud provider signing a group of events within a report or log for a cloud consumer

5.16.4.1 Description

A cloud provider may have a set of related events (perhaps from the same resource or events that reflect a correlation or transaction).

Instead of signing each individually, they may be signed as a group by using some correlating identifier along.

A provider of events wants to assure a method exists to sign groups of "like" events (perhaps from the same secure database) instead of having to sign them individually.

5.16.4.2 Requirements and considerations

To assure a method exists to sign groups of "like" events (perhaps from the same secure database) instead of having to sign them individually.

- That components (such as network appliances) may be able to sign their own events.
- That components would only sign events that they generated.
- That components that modify events may need a means to sign (or resign) an already signed event (perhaps look into use the "report chain" to capture signing info).
- In order to make signing groups of events efficient, that temporal order of events in reports may need to be non-linear.

5.16.4.3 Assumptions

- Provider, Consumer or Component Granular. That is the signing of groups of like events within an audit report or log may done by the cloud provider (as an entity), cloud consumer (perhaps at an account or application level) or by and not individual IT component resources or services (including partner services).

5.16.4.4 Event classification data

Not applicable

5.16.4.5 Classification notes

Not applicable

5.16.5 Use Case 4: Cloud partners or customers signing a one or more events for submission to cloud provider

5.16.5.1 Description

Partners and other federated (distributed) services that contribute to a cloud application/service/workflow may need a means to sign their event submissions to the cloud provider that will end up on cloud consumer/customer logs and reports.

Cloud consumers may be permitted to submit events from their hosted applications/services via some interface supported by the cloud provider.

Cloud consumer applications or cloud partners (third-party service providers) wish to be able to submit audit events that to the cloud provider in a format that can contribute to the entirety of the providers audit stream.

5.16.5.2 Requirements and considerations

To assure that cloud consumer applications or cloud partners (third-party service providers) are able to submit audit events to the cloud provider in a format that can contribute to the entirety of the providers audit stream.

- This use case may include "message-level" signing of one or more events being submitted over an interface to the cloud provider from a cloud consumer.

5.16.5.3 Assumptions

- Partner Service, or Consumer granular.

5.16.5.4 Event classification data

Not applicable

5.16.5.5 Classification notes

Not applicable

5.16.6 Use Case 5: Cloud infrastructure components signing events

5.16.6.1 Description

Some components in a cloud infrastructure may have the ability to identify themselves and sign events they generate with their own keys that have been established with the cloud provider or perhaps even a cloud consumer who has dedicated resources within the provider's infrastructure (e.g., a database appliance, or a web server appliance).

Component resources (e.g., appliances such as a database, web server, or network appliance) and hosted cloud services (including partner services hosted within the cloud provider) wish to be able to submit signed audit events that to the cloud provider in a format that can contribute to the entirety of the providers audit stream.

5.16.6.2 Requirements and considerations

To assure that component resources (e.g., appliances, such as a database, web server, or network appliance) and hosted cloud services (including partner services hosted within the cloud provider) are able to submit signed audit events that to the cloud provider in a format that can contribute to the entirety of the providers audit stream.

5.16.6.3 Assumptions

- Component Resource or Service Granular

5.16.6.4 Event classification data

Not applicable

5.16.6.5 Classification notes

Not applicable

5.17 Summarization and suppression

5.17.1 Summarization

5.17.1.1 Description

Certain raw event sources are very noisy and may create a large number of identical or significantly similar events. For storage reduction, bandwidth reduction, and processing reduction, there is a need to be able to summarize these events as close to the log source as possible, while still preserving the essence of the nature of these events.

Service Level Monitor Examples:

1. A hosting provider emits a “status okay” event for a given hosted application every 30 seconds. Over the course of minutes, hours, or even days, these emitted events may be nearly identical, differing only in the time stamps of the events. A SLM compliance service needs to query these status okay events, but desires to have all identical events within a time range collapsed in to a summarized event so that the query result set it obtains is smaller, requires less bandwidth to transfer, less space to store, and less computing resources to process.
2. A hosting provider emits a “resource usage” event for a given resource every 5 seconds. Large groups of similar events will exist, differing only in the time stamps of the events. A capacity forecasting component needs the data contained in these events, but does not need the individual events

Security Compliance Examples:

1. A reseller of banking web services wishes to provide to its customers suspicious event information related to attempted accesses to its hosted banking services so that its customers can be in compliance with defined control objectives. The reseller is subjected to a massive access-attempt DDoS attack, which generates several billion access logs from a million node bot net. In delivering these event records to its customers, the reseller desires to summarize events based on time and time range, but for practical purposes cannot preserve the originating event sources in the summarized events.

In all of the above use examples, there is a need for the summarized event to have the following:

- An indicator (either implicit or explicit) that an event in a query result set is a summarization of other events.
- A time range which indicates the earliest and latest event times being summarized
- A count indicating the number of events that have been collapsed or summarized
- A preservation of all properties which were identical across the summarized events.

5.17.1.2 Requirements and considerations

None

5.17.1.3 Assumptions

None

5.17.1.4 Event classification data

None

5.17.1.5 Classification notes

Additional Data:

- Time Range Represented by the event
- Count of events summarized
- Indicator (implicit or explicit) that event is a summarized event
- Information about fields which were not identical (other than event time) where information was dropped in the event

5.17.2 Event suppression

5.17.2.1 Description

A cloud provider generates a large number of events. For practical purposes, events that are deemed irrelevant are often dropped/suppressed/filtered at various points in an event ecosystem. For example, a security device within the cloud provider may generate large numbers of events that a reporter does not deem necessary (according to some compliance policy). However, a consumer of the events wished to use them for compliance auditing. In such usage, it is often important to have meta-events in the system that record the fact that events of a certain type were dropped, together with a count of the events that were dropped.

5.17.2.2 Requirements and considerations

Suppression meta-events are similar to summarization events, but differ in their lack of a need to preserve key values from the original events.

5.17.2.3 Assumptions

None

5.17.2.4 Event classification data

None

5.17.2.5 Classification notes

None

5.17.3 Undeveloped summarizing SLA use case idea

A couple of cases that may be seen as "summarization":

1. A typical case of aggregation: events logged for a complex resource (e.g., a virtual system in a cloud), can be an aggregation of events from the components of this resource. For example:

"Starting" a virtual system in a cloud, will require starting every one of its components. A "successful" system start event can be logged only when all components have been started successfully. If only a subset of the components have started and nothing happens for the remaining components over some time, a "start failure" could be logged for the system.

- In many cases you could argue that there is a system entity doing this aggregation for you. But in other cases, e.g., a distributed system, the aggregation/summarization has to be done from a log.

2. A "metrics" event that keeps track of a response time average, for SLA / SLO tracking purpose. The event may aggregate all response times over a day (or from beginning of an SLA measurement period), and can be used as alarm in case of failure to satisfy SLO.

5.18 Temporal

5.18.1.1 Description

A consumer of CADF events is concerned with several issues related to time and the events in a result set, including the following:

- Accuracy of a time stamp (i.e., is there a way to understand the accuracy of the time on the host which recorded the event? Is there any protection against post-action event injection due to server time stamp adjustment?) This is particularly important in trying to correlate events that occur worldwide on different hosts against each other.
- Time zone of the Initiator, and time zone of each reporter – In particular, the time zone of the Initiator is important because an event consumer may wish to detect actions which occur at some unusually local time. A normalized GMT time is not sufficient for all scenarios. The time zone of the reporters is import for similar auditing reasons.
- Precision of a time stamp – This is of particular concern when a consumer wishes to make repeated queries to collect ALL events and does not want to run in to situations where it misses events or gets duplicate events at the “overlap” of two queries.
- Latency of processing of events – Especially for processing/querying that occurs in near-real time, it may be important to understand some aspect of the latency of event collection throughout the entire system, to ensure, where possible, event sequencing and event correlation integrity.
- Time stamps for each reporter, or each key operation on an event, for audit trail purposes. For example, if a group of events are aggregated, the time the aggregation occurs is important.

5.18.1.2 Requirements and considerations

None

5.18.1.3 Assumptions

None

5.18.1.4 Event classification data

None

5.18.1.5 Classification notes

None

1430

Change log

Version	Date	Description
1.0.0a	2012-06-07	Alvin Black (CA), Matt Rutkowski (IBM) Final editor draft candidate. for WIP public review

Bibliography

DMTF DSP4004, *DMTF Release Process 2.4*,
http://www.dmtf.org/sites/default/files/standards/documents/DSP4004_2.4.pdf