Adoption of Cloud Auditing Data Federation (CADF) standard by IBM Spectrum Virtualize

Version 1.0

Overview

Challenge:

- *Developer* – How to make use of new infrastructure in IBM Spectrum Virtualize to relay event as per CADF specification?
- *Cloud administrator* – How to understand events of Spectrum Virtualize which are relayed as per CADF specification?
- *Traditional storage administrator* – How to correlate between traditional Spectrum Virtualize format of an event and a new format?
- *Common* – How to perceive the intended meaning of taxonomy value in CADF specification and understand CADF specification better in general?

Solution:

- Describing the CADF event and data model using examples and good explanation of events relayed by Transparent Cloud Tiering feature of Spectrum Virtualize
- A perfect mix of theoretical and practical knowledge to equip the reader about CADF and applying the same

Table of contents

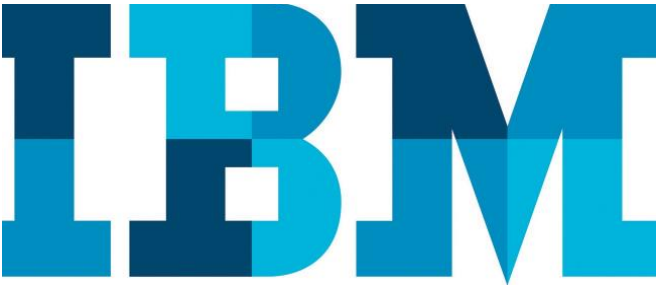

Abstract

This white paper introduces the IBM® Spectrum Virtualize™ Transparent Cloud Tiering feature in the 7.8.0.0 release. This white paper includes the adoption of the Cloud Auditing Data Federation (CADF) specification to log an event for this feature, which can help IBM Spectrum Virtualize to interact better with other entities in the cloud world.

Intended audience and scope

This technical report is intended to present information in a way that it could benefit first time readers, Spectrum Virtualize users, and administrators, though it also equally benefit developers and testers who work on Spectrum Virtualize to get on board to understand the CADF specification and Spectrum Virtualizes adoption of it. The beneficiary includes:

- Traditional storage administrator
- Cloud administrator
- Developer
- Tester

Prerequisites

This technical paper assumes familiarity with the following prerequisites:

1. Basic knowledge of Spectrum Virtualize range of products. IBM System Storage® SAN Volume Controller (SVC) or IBM Storwize® systems installed with the Spectrum Virtualize software 7.8.0 code level or later (that supports Transparent Cloud Tiering feature) and should have basic configuration [such as Ethernet/Fibre Channel (FC)/Fibre Channel over Ethernet (FCoE) connections, Domain Name System (DNS), switches, and so on].
2. Storage system should have valid licenses for IBM FlashCopy®, encryption and Transparent Cloud Tiering (for Storwize only).
3. User should have account with supported cloud service provider (CSP). Also, CSP account should be accessible through management network of IBM SVC/ Storwize cluster
4. Users requiring data encryption on cloud should have access to IBM Security Key Lifecycle Manager or physical access of USBs on IBM SVC/Spectrum Virtualize.

CADF requirement

Customers often hesitate to adopt the cloud deployment model due to security and integration issues.

The main ask is that security policy which binds to their application are enforced in the cloud in the same way as off cloud on-premises environment.

The other important point is having capability to analyze events and notification from all the heterogeneous entities in a timeline and promote better application integration.

This can be achieved if all the entities in the cloud world that interact with each other are able to provide specific audit event, log, and report information. This entails all the hardware, software, and network infrastructure involved in cloud deployment. This is not limited to just CSP elements, but also other external vendors who take part in cloud deployment in any way.

Spectrum Virtualize making way into the cloud deployment as an entity that helps customers to back up their data into the cloud, also faces the same requirement like other entities in the cloud world.

A proven method to address such needs is to develop open standards to enable information sharing. CADF is one such standard for describing events (the 'audit' part) observed by different computer systems that can be stored in a single place to give an administrator a single view of what has happened (the 'federation' part CADF data Model

The CADF data model is designed to provide the information that the auditors are looking forward to track activities in cloud environments. The data in an event can record the **WHO**, **WHAT, WHEN, WHERE, FROM WHERE** and **TOWHERE** of an activity. This is also referred to as the 7 Ws of audit and compliance.

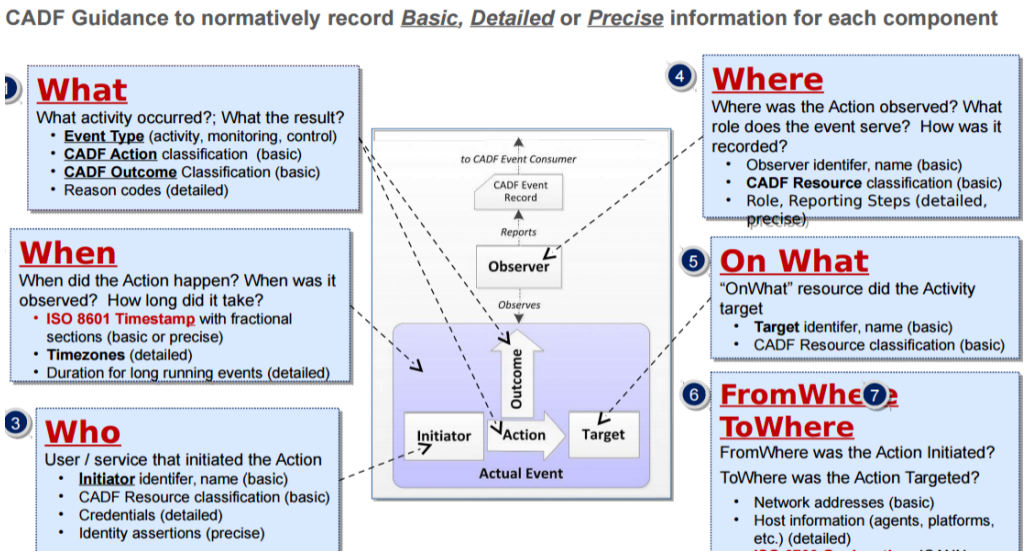

Figure 1: CADF – the 7 'W's of audit

Event model

The event model uses the concept of a resource that is used within multiple defined event components. A resource is an entity that can provide or consume services or information within the context of a cloud Infrastructure. Examples of resources include traditional IT infrastructure components such as servers and network devices, software components such as databases and applications, operation and business entities, such as accounts, users, and roles (used for security).
CADF allows the event model to be extended to include new event types that can be used for other domains.

Included in the event model are taxonomies for specific field values. The taxonomies ensure that event field values are consistent when the events come from different sources (that is, different cloud providers).
The taxonomies include:

- Resource taxonomy – used to classify the event by the logical IT or cloud resources that are related to the event's action. For example, values of this taxonomy could be used to classify the resource that observed the action or the resource that was the (intended) target of the action.
- Action taxonomy – used to classify the event by the activity that caused it to be generated.
- Outcome taxonomy – used to describe the outcome of the attempted action of the event.

Every CADF event has the following five most important resources to be identified:

- Observer
- Initiator
- Target
- Action
- Outcome

A careful procedure to choose what will be observer, initiator, action, and target resources for the CADF event to be logged is having utmost importance. It gives a clear idea to users about what is going on (back and forth) from that entity (for example, storage array) perspective. The outcome gives a clear idea about the result of that operation.

The following figure and the table explain the data model and its element.

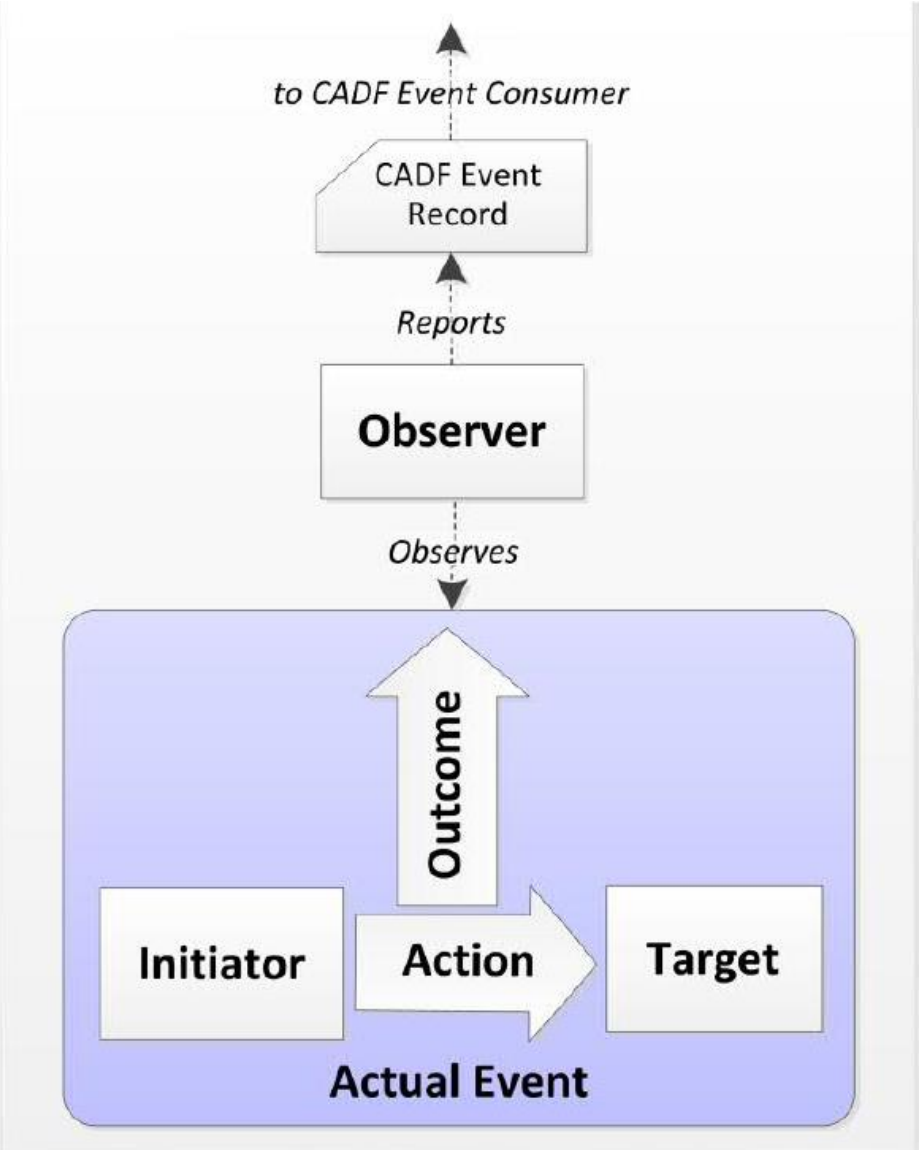

Figure 2: CADF event model components

Model component	CADF definition
Observer	The resource that generates the CADF event record based on its observation (directly or indirectly) of the actual event
Initiator	The resource that initiated, originated, or instigated the event's action, according to the observer
Action	The operation or activity the initiator has performed or attempted to perform or has pending against the event's target, according to the observer
Target	The resource against which the action of a CADF event record was performed, was attempted, or is pending according to the observer **Note**: A target (in the CADF event model) can represent a plurality of target resources
Outcome	The result or status of the action against the target, according to the observer

Table 1: CADF definition of model components

The following table lists the most important event components that make one CADF event record and also echo Spectrum Virtualize view sometime.

Common attribute/component	Comment
Event unique identifier	Every CADF data must have a unique identifier. Spectrum Virtualize uses system ID and the time of the event to generate this unique identifier.
Timestamp of the event	This refers to the time when the event happened. CADF *event records* seek to represent time so that consumers can make intelligent decisions about how each event (within the same activity domain) relates to other events temporally. The UTC offset is always required (not optional) to remove ambiguity. Timestamps in CADF event records is recorded in local time, (that is, the 24-hour clock time for the local time zone) with explicit reference to the Coordinated Universal Time (UTC) time zone offset. For example: "New York City, United States during Eastern Standard Time (EST) or UTC-05:00" During the period when Eastern Standard Time (EST) is in effect, the UTC offset for New York City would be UTC minus five hours or UTC-05:00. An example of a valid timestamp typed value for New York City during EST would be: 2012-02-25T09:00:00-05:00 This timestamp

	represents the date February 25th, 2012 at 9:00 AM (EST) local time in New York City.
Event type	Every CADF data most have an event type. Event type can be one of the following: • Activity • Control • Monitor
Outcome	It refers to the result of the event. It can be one of the following: • Success • Failure • Pending
Event action	It refers to the action that was performed in the cloud. For example, taking backup/restore.
Event type URI	It can be used to declare versioning/schema of events.
Tags	It is a label that can be added to a CADF event record to qualify or categorize an event. Tags provide a powerful mechanism for adding domain-specific identifiers and classifications to CADF event records that can be referenced by the CADF query interface. Customers can construct custom reports or views on the event data held by a provider for a specific domain of interest. **For example, all backup-related events can carry a tag, *backup*.** Note: A CADF event record can have multiple tags that enable cross-domain analysis.
Attachments	It is a container for data or *content* that might follow any structure – from an atomic type in a complex hierarchy used to describe the other CADF entity, such as observer/initiator/target or the event record itself. As the name suggests, it is used to give more information to know better about the said resource.
Initiator resource	It is the resource that initiates the request.
Target resource	It is the target service where the request will be performed.
Observer resource	It is the service that logs or monitors the audit event (for example, logger module in a Spectrum Virtualize cluster).
Measurements	They are used to report various stat attributes related to the event. Measurements are mandatory when the event type is *monitor*, else

	optional. For example, the data which is uploaded as part of backup operation can be reported and the unit could be GB.
Reason	It contains a means to provide additional details and further classify the top-level *outcome* of the *action* included in a CADF event record. This is mandatory when the event type is *control*, else optional. Reason code is taken from canonical sources so that each entity can understand what does this mean. One excellent canonical source for event reason codes is the HTTP Status Codes, which are defined by the URI: http://www.iana.org/assignments/http-status-codes/http-status-codes.xml. For example, any request to access a resource for which proper authorization has not been provided can result in a *401 reasonCode* property value, which corresponds to *Unauthorized*.

Table 2: Event components

Benefits of following the CADF specification

Customers might not trust clouds to host their workloads and data without the ability to self-audit and monitor the provider's platform and infrastructure. The CADF specification provides a comprehensive event model that can be used to solve this problem and prove compliance against corporate, industry, or regional policies in any cloud deployment, CADF can also provide accurate metering data that can be used for service level agreement (SLA) monitoring, real-time analytics and problem diagnosis in cloud infrastructures.
The CADF specification maps out functionality that makes it indispensable for a cloud platform. Having an audit trail of the events that occur on cloud platforms is critical for enterprise applications, and the auditing functions of CADF can also be used for security and troubleshooting. The data captured is ideal for data mining and is easily analyzed.

The CADF will develop specifications for federating audit event data, including interface definitions and a compatible interaction model that describes interactions between IT resources for cloud deployment models. The CADF is also working closely with the Distributed Management Task Force (DMTF) Cloud Management Working Group (CMWG) to reference their resource model and interface protocol work.

Benefits to storage arrays:

- Ability to self-manage auditing of their data
- Similar reports from different cloud providers
- Aggregate audit data from different clouds / partners
- Auditing processes and tools unchanged

IBM Spectrum Virtualize adoption – Identification of event nodes and enablement

In IBM Spectrum Virtualize Transparent Cloud Tiering, for each CADF event record, there is an event type associated with it and it can take three values.

- Activity – Characterizes events that provide information about actions having occurred or intended to occur. Such events typically report on regular operations of a cloud infrastructure or services. An action may also require multiple activity events to describe it completely if it is asynchronous. For instance, taking a snapshot is an asynchronous activity if accepted. So, the snapshot started event will be generated at time T1 and snapshot completion event will be generated at time T2 where T2 > T1.
  However, to avoid duplication with audit log, Spectrum Virtualize relays events as per CADF specification only when an operation is complete and the result is available. So, initiation part of any operation would not result into an event right now.
- Control – Characterizes events that reflect on or provide information about the application of a policy or business rule, or more generally, express the outcome of a decision-making process. The property *reason* has to be populated mandatorily in this kind of event.
- Monitor – Characterizes events that provide information about the status of a resource or its attributes or properties. Such events typically report on measurements or periodic probes on cloud resources, and might produce aggregate data such as statistical or summary metrics. Measurement is a mandatory component of this kind of event.

Treatment (identification) of observer/initiator/target from Spectrum Virtualize perspective

You can choose the right taxonomy values for the actors using the following two rules:

- Find the highest-level taxonomy value by answering the following question yourself.

Where do the actors sit in the platform? – Compute, network, storage, or service layer.

- The subsequent level values for taxonomy should be chosen to equip query platform to find all the requests initiated by the user easily.

Suggestion for observer/initiator/target taxonomy value

- **Observer taxonomy suggestion**

  The observer should be under the service branch of the CADF resource taxonomy. For all the events that take place, this role is played by cluster software of Spectrum Virtualize, and specifically by the logging service of this software.

  Use service/network/cluster/logger to identify the observer. The ID will be the Spectrum Virtualize system ID *(which is unique)* and name would be Spectrum Virtualize system name.

- **Initiator taxonomy suggestion**
  The initiator should be under the following branches of the CADF resource taxonomy:

  - Data/Security
  - Network/Node
  - Service

  For all the events that take place in the Spectrum Virtualize case, a service that exists on a configuration (master) node plays this role in the cluster, So, use **service/network/node** as the taxonomy value.

- **Target taxonomy suggestion**
  The target either denotes volume that exist in our storage controller or the cloud account on which an operation is being targeted.

You can use the following taxonomy values:

storage/volume: A case when restoring of a backed-up volume from a cloud service provider to a block storage appliance happens.

service/storage/object: A case when backing up a volume from a block storage appliance to an object storage maintained by a cloud service provider (CSP).

Reason code:

Spectrum Virtualize follows the reason code of the following two canonical sources:

- **HTTP status code registry:** http://www.iana.org/assignments/http-status-codes/http-status-codes.xml
- **Distributed audit service code registry:** http://www.opengroup.org/bookstore/catalog/p441.htm

Enabling syslog server in Spectrum virtualize

Prerequisite

Configuring Spectrum Virtualize system for logging CADF events requires a remote syslog server that can receive syslog data from syslog clients (Spectrum Virtualize system). Usually, Linux hosts have default syslog servers installed. Few syslog servers are: syslog, syslog-n:g, rsyslog.

Note: The test team used rsyslogd 5.8.10 on a RHEL Linux 6.4 based host for sample data collection.

To configure a remote syslog server for sending CADF data on Spectrum Virtualize, issue the following command:

```
svctask mksyslogserver -ip <Syslog server IP
address> -error <on/off> -warning <on/off> -info
<on/off> -cadf on
```

Note: Events in Spectrum Virtualize can be of type `error`/`warning`/`info`. Using flags, `-error`/`-warning`/`-info`, an administrator can configure the type of events that can reach the remote syslog server. For general purposes, keep all the types '**on**' so that your configured syslog server can receive all of them. As of now, you might not be able to view an event of type, `warning`. A successful Transparent Cloud Tiering operation results in an `info` event and an unsuccessful operation in an `error` event.

```
IBM_2145:CAYMAN:admin>mksyslogserver -ip  9.193.231.103 -error on -warning on -info on -cadf  on
Syslog Server id [0] successfully created
IBM_2145:CAYMAN:admin>lssyslogserver
id name    IP_address    facility error warning info cadf
0  syslog0 9.193.231.103          on    on      on   on
```

Figure 3: ***How to enable CADF in Spectrum Virtualize***

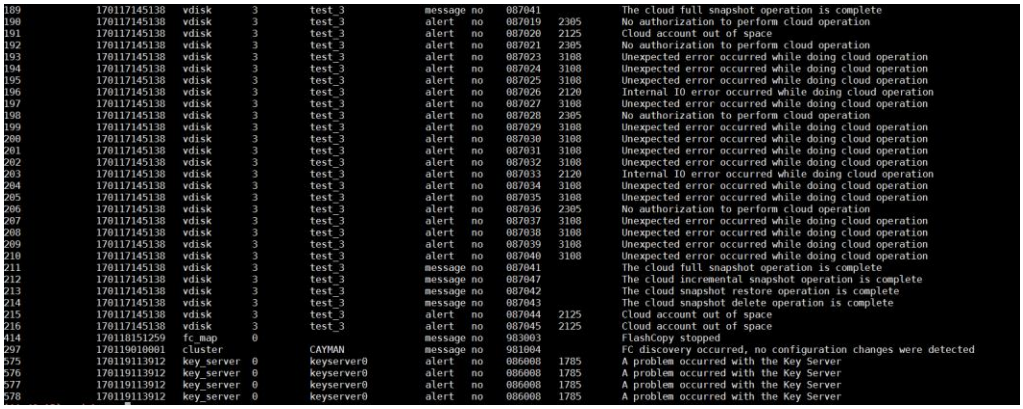

189	170117145138	vdisk	3	test_3	message	no	087041		The cloud full snapshot operation is complete
190	170117145138	vdisk	3	test_3	alert	no	087019	2305	No authorization to perform cloud operation
191	170117145138	vdisk	3	test_3	alert	no	087020	2125	Cloud account out of space
192	170117145138	vdisk	3	test_3	alert	no	087021	2305	No authorization to perform cloud operation
193	170117145138	vdisk	3	test_3	alert	no	087023	3108	Unexpected error occurred while doing cloud operation
194	170117145138	vdisk	3	test_3	alert	no	087024	3108	Unexpected error occurred while doing cloud operation
195	170117145138	vdisk	3	test_3	alert	no	087025	3108	Unexpected error occurred while doing cloud operation
196	170117145138	vdisk	3	test_3	alert	no	087026	2120	Internal IO error occurred while doing cloud operation
197	170117145138	vdisk	3	test_3	alert	no	087027	3108	Unexpected error occurred while doing cloud operation
198	170117145138	vdisk	3	test_3	alert	no	087028	2305	No authorization to perform cloud operation
199	170117145138	vdisk	3	test_3	alert	no	087029	3108	Unexpected error occurred while doing cloud operation
200	170117145138	vdisk	3	test_3	alert	no	087030	3108	Unexpected error occurred while doing cloud operation
201	170117145138	vdisk	3	test_3	alert	no	087031	3108	Unexpected error occurred while doing cloud operation
202	170117145138	vdisk	3	test_3	alert	no	087032	3108	Unexpected error occurred while doing cloud operation
203	170117145138	vdisk	3	test_3	alert	no	087033	2120	Internal IO error occurred while doing cloud operation
204	170117145138	vdisk	3	test_3	alert	no	087034	3108	Unexpected error occurred while doing cloud operation
205	170117145138	vdisk	3	test_3	alert	no	087035	3108	Unexpected error occurred while doing cloud operation
206	170117145138	vdisk	3	test_3	alert	no	087036	2305	No authorization to perform cloud operation
207	170117145138	vdisk	3	test_3	alert	no	087037	3108	Unexpected error occurred while doing cloud operation
208	170117145138	vdisk	3	test_3	alert	no	087038	3108	Unexpected error occurred while doing cloud operation
209	170117145138	vdisk	3	test_3	alert	no	087039	3108	Unexpected error occurred while doing cloud operation
210	170117145138	vdisk	3	test_3	alert	no	087040	3108	Unexpected error occurred while doing cloud operation
211	170117145138	vdisk	3	test_3	message	no	087041		The cloud full snapshot operation is complete
212	170117145138	vdisk	3	test_3	message	no	087047		The cloud incremental snapshot operation is complete
213	170117145138	vdisk	3	test_3	message	no	087042		The cloud snapshot restore operation is complete
214	170117145138	vdisk	3	test_3	message	no	087043		The cloud snapshot delete operation is complete
215	170117145138	vdisk	3	test_3	alert	no	087044	2125	Cloud account out of space
216	170117145138	vdisk	3	test_3	alert	no	087045	2125	Cloud account out of space
414	170110151259	fc_map	0		message	no	983003		FlashCopy stopped
297	170119010801	cluster		CAYMAN	message	no	981004		FC discovery occurred, no configuration changes were detected
575	170119113912	key_server	0	keyserver0	alert	no	086008	1785	A problem occurred with the Key Server
576	170119113912	key_server	0	keyserver0	alert	no	086008	1785	A problem occurred with the Key Server
577	170119113912	key_server	0	keyserver0	alert	no	086008	1785	A problem occurred with the Key Server
578	170119113912	key_server	0	keyserver0	alert	no	086008	1785	A problem occurred with the Key Server

Figure 4: ***Transparent Cloud Tiering Events in Spectrum Virtualize***

Sample Spectrum Virtualize events as per CADF specification

As Transparent Cloud Tiering is the first service out of Spectrum Virtualize which offers event per CADF specification, here are a few sample events in the JSON format related to Transparent Cloud Tiering.

View CADF events in the syslog server by looking into /var/log/message:


```
{"typeURI":
"http://schemas.dmtf.org/cloud/audit/1.0/event",
"eventTime": "2016-10-26T16:07:07.000000+0000", "target":
{"typeURI": "storage/volume", "attachments": [{"content":
"0", "typeURI": "text/plain", "name":
"cloud_account_id"}], "name": "test_1", "id":
"600507680C8C0002400000000000001"}, "observer":
{"typeURI": "service/network/cluster/logger", "id":
"20323000090", "name": "SVC_NODE"}, "tags": ["Restore"],
"eventType": "activity", "measurements": [{"metric":
{"metricId":
"www.ibm.com/svc/Cloud/Backup_Time/1477498027/197/1",
"name": "Time of backup being copied or restored",
"unit": "YYMMDDHHMMSS"}, "result":
"1970/01/01/00/00/00"}, {"metric": {"metricId":
"www.ibm.com/svc/Cloud/Backup_Generation_Number/147749802
7/197/2", "name": "Volume backup generation number",
"unit": "Natural Number"}, "result": "All"}],
"initiator": {"typeURI": "service/network/node", "host":
{"address": "IP_ADDRESS"}, "id": "2", "name": "node2"},
"reason": {"reasonCode": "200", "reasonType":
"http://www.iana.org/assignments/http-status-codes/http-
status-codes.xml"}, "action": "restore", "outcome":
"success", "id": "20323000090-1477498027-197"}
Oct 26 16:07:42 IP_ADDRESS IBM2145: {"typeURI":
"http://schemas.dmtf.org/cloud/audit/1.0/event",
```

"eventTime": "2016-10-26T16:07:07.000000+0000", "target": {"typeURI": "storage/volume", "attachments": [{"content": "0", "typeURI": "text/plain", "name": "cloud_account_id"}], "name": "test_1", "id": "600507680C8C00024000000000000001"}, "observer": {"typeURI": "service/network/cluster/logger", "id": "20323000090", "name": "SVC_NODE"}, "tags": ["Restore"], "eventType": "activity", "measurements": [{"metric": {"metricId": "www.ibm.com/svc/Cloud/Backup_Time/1477498027/199/1", "name": "Time of backup being copied or restored", "unit": "YYMMDDHHMMSS"}, "result": "1970/01/01/00/00/00"}, {"metric": {"metricId": "www.ibm.com/svc/Cloud/Backup_Generation_Number/1477498027/199/2", "name": "Volume backup generation number", "unit": "Natural Number"}, "result": "All"}], "initiator": {"typeURI": "service/network/node", "host": {"address": "IP_ADDRESS"}, "id": "2", "name": "node2"}, "reason": {"reasonCode": "507", "reasonType": "http://www.iana.org/assignments/http-status-codes/http-status-codes.xml"}, "action": "restore", "outcome": "failure", "id": "20323000090-1477498027-199"}

Sample CADF events (a user friendly presentation prepared by JSON formatter)

- A successful full backup (cloud full snapshot operation) event


```
{
   "typeURI":
"http://schemas.dmtf.org/cloud/audit/1.0/even
t",
   "eventTime": "2016-10-
12T20:02:30.000000+0000",
   "target": {
      "typeURI": "service/storage/object",
      "id": "0",
      "name": "cloudaccount0"
   },
   "observer": {
      "typeURI":
"service/network/cluster/logger",
      "id": "10032004394",
      "name": "Cluster_9.193.231.50"
   },
   "tags": [
      "Backup"
   ],
   "eventType": "activity",
   "measurements": [
      {
        "metric": {
```

```
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Time/1476302550
/110/1",
            "name": "Time of backup being
copied or restored",
            "unit": "YYMMDDHHMMSS"
         },
         "result": "2016/10/12/20/02/30"
      },
      {
         "metric": {
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Generation_Numb
er/1476302550/110/2",
            "name": "Volume backup generation
number",
            "unit": "Natural Number"
         },
         "result": "1"
      }
   ],
   "initiator": {
      "typeURI": "service/network/node",
      "host": {
         "address": "9.193.231.50"
      },
      "attachments": [
         {
            "content":
"6005076400C8010E5000000000000000",
            "typeURI": "text/plain",
            "name": "volume_uuid"
        }
      ],
      "name": "node1",
      "id": "1"
   },
   "reason": {
      "reasonCode": "200",
      "reasonType":
"http://www.iana.org/assignments/http-status-
codes/http-status-codes.xml"
   },
   "action": "backup",
   "outcome": "success",
   "id": "10032004394-1476302550-110"
}
```


- A successful incremental backup (cloud incremental snapshot operation) event


```
{
   "typeURI":
"http://schemas.dmtf.org/cloud/audit/1.0/event",
   "eventTime": "2016-10-12T20:02:55.000000+0000",
   "target": {
```

```
      "typeURI": "service/storage/object",
      "id": "0",
     "name": "cloudaccount0"
   },
   "observer": {
      "typeURI": "service/network/cluster/logger",
     "id": "10032004394",
      "name": "Cluster_9.193.231.50"
   },
   "tags": [
      "Backup"
   ],
   "eventType": "activity",
   "measurements": [
      {
         "metric": {
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Time/1476302575/111/1
",
            "name": "Time of backup being copied or
restored",
            "unit": "YYMMDDHHMMSS"
         },
         "result": "2016/10/12/20/02/55"
      },
      {
         "metric": {
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Generation_Number/147
6302575/111/2",
            "name": "Volume backup generation
number",
            "unit": "Natural Number"
         },
         "result": "2"
      }
   ],
   "initiator": {
      "typeURI": "service/network/node",
      "host": {
         "address": "9.193.231.50"
      },
      "attachments": [
         {
            "content":
"6005076400C8010E5000000000000000",
            "typeURI": "text/plain",
            "name": "volume_uuid"
         }
      ],
      "name": "node1",
      "id": "1"
   },
   "reason": {
      "reasonCode": "200",
```

```
      "reasonType":
"http://www.iana.org/assignments/http-status-
codes/http-status-codes.xml"
   },
   "action": "backup/incremental",
   "outcome": "success",
   "id": "10032004394-1476302575-111"
}
```


- A successful restore (cloud snapshot restore operation) event


```
 {
   "typeURI":
"http://schemas.dmtf.org/cloud/audit/1.0/event",
   "eventTime": "2016-10-12T20:56:29.000000+0000",
   "target": {
      "typeURI": "storage/volume",
      "attachments": [
         {
            "content": "0",
            "typeURI": "text/plain",
            "name": "cloud_account_id"
         }
      ],

      "name": "vdisk0",
      "id": "6005076400C8810E5000000000000000"
   },
   "observer": {
      "typeURI": "service/network/cluster/logger",
      "id": "10032204394",
      "name": "Cluster_9.193.231.50"
   },

   "tags": [
      "Restore"
   ],
   "eventType": "activity",
   "measurements": [
      {
         "metric": {
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Time/1476305789/111/1
",
            "name": "Time of backup being copied or
restored",
            "unit": "YYMMDDHHMMSS"
         },
         "result": "2016/10/12/20/56/14"
      },
      {
        "metric": {
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Generation_Number/147
6305789/111/2",
```

```
            "name": "Volume backup generation
number",
            "unit": "Natural Number"
         },
         "result": "2"
      }
   ],
   "initiator": {
      "typeURI": "service/network/node",
      "host": {
         "address": "9.193.231.50"
      },
      "id": "1",
      "name": "node1"
   },
   "reason": {
      "reasonCode": "200",
      "reasonType":
"http://www.iana.org/assignments/http-status-
codes/http-status-codes.xml"
   },
   "action": "restore",
   "outcome": "success",
   "id": "10032204394-1476305789-111"
}
```


- A successful delete (cloud snapshot delete operation) event


```
{
   "typeURI":
"http://schemas.dmtf.org/cloud/audit/1.0/event",
   "eventTime": "2016-10-12T11:23:13.000000+0000",
   "target": {
      "typeURI": "service/storage/object",
      "id": "0",
      "name": "deepak"
   },
   "observer": {
      "typeURI": "service/network/cluster/logger",
      "id": "20322800090",
      "name": "SVC_NODE"
   },
   "tags": [
      "Delete"
   ],
   "eventType": "activity",
   "measurements": [
      {
         "metric": {
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Generation_Number/147
6271393/350/1",
            "name": "Volume backup generation
number",
            "unit": "Natural Number"
```

```
            },
            "result": "All"
         }
      ],
      "initiator": {
         "typeURI": "service/network/node",
         "host": {
            "address": "IP_ADDRESS"
         },
         "attachments": [
            {
               "content":
"600507680C8900024000000000000012",
               "typeURI": "text/plain",
               "name": "volume_uuid"
            }
         ],
         "name": "node1",
         "id": "3"
      },
      "reason": {
         "reasonCode": "200",
         "reasonType":
"http://www.iana.org/assignments/http-status-
codes/http-status-codes.xml"
      },
      "action": "delete",
      "outcome": "success",
      "id": "20322800090-1476271393-350"
}
```


- A failed backup operation due to space crunch event


```
{
   "typeURI":
"http://schemas.dmtf.org/cloud/audit/1.0/event",
   "eventTime": "2016-10-12T11:23:13.000000+0000",
   "target": {
      "typeURI": "service/storage/object",
      "id": "0",
      "name": "hrms_account"
   },
   "observer": {
      "typeURI": "service/network/cluster/logger",
      "id": "20322800090",
      "name": "SVC_NODE"
   },
   "tags": [
      "Backup"
   ],
   "eventType": "activity",
   "measurements": [
      {
         "metric": {
```

```
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Time/1476271393/327/1
",
            "name": "Time of backup being copied or
restored",
            "unit": "YYMMDDHHMMSS"
         },
         "result": "1970/01/01/00/00/00"
      },
      {
         "metric": {
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Generation_Number/147
6271393/327/2",
            "name": "Volume backup generation
number",
            "unit": "Natural Number"
         },
         "result": "68"
      }
   ],
   "initiator": {
      "typeURI": "service/network/node",
      "host": {
         "address": "IP_ADDRESS"
      },

      "attachments": [
         {
            "content":
"600507680C8900024000000000000012",
           "typeURI": "text/plain",
            "name": "volume_uuid"
         }
      ],
      "name": "node1",
      "id": "3"
   },
   "reason": {
      "reasonCode": "507",
      "reasonType":
"http://www.iana.org/assignments/http-status-
codes/http-status-codes.xml"
   },
   "action": "backup",
  "outcome": "failure",
   "id": "20322800090-1476271393-327"
}
```


- A failed restore operation due to Spectrum Virtualize internal issue eventA


```
{
   "typeURI":
"http://schemas.dmtf.org/cloud/audit/1.0/event",
```

```
   "eventTime": "2016-10-12T11:23:13.000000+0000",
   "target": {
      "typeURI": "storage/volume",
      "id": "0",
      "name": "hrms_account"
   },
   "observer": {
      "typeURI": "service/network/cluster/logger",
      "id": "20322800090",
      "name": "SVC_NODE"
   },
   "tags": [
      "Restore"
   ],
   "eventType": "activity",
   "measurements": [
      {
         "metric": {
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Time/1476271393/340/1
",
            "name": "Time of backup being copied or
restored",
            "unit": "YYMMDDHHMMSS"
         },
         "result": "1970/01/01/00/00/00"
      },
      {
         "metric": {
            "metricId":
"www.ibm.com/svc/Cloud/Backup_Generation_Number/147
6271393/340/2",
            "name": "Volume backup generation
number",
            "unit": "Natural Number"
         },
         "result": "All"
      }

   ],

   "initiator": {
      "typeURI": "service/network/node",
      "host": {
         "address": "IP_ADDRESS"
      },
      "attachments": [
         {
            "content":
"600507680C8900024000000000000012",
            "typeURI": "text/plain",
            "name": "volume_uuid"
         }
      ],
      "name": "node1",
```

```
      "id": "3"
   },
   "reason": {
      "reasonCode": "500",
     "reasonType":
"http://www.iana.org/assignments/http-status-
codes/http-status-codes.xml"
   },
   "action": "restore",
   "outcome": "failure",
   "id": "20322800090-1476271393-340"
}
```


Future – Wider adoption by Spectrum Virtualize

Spectrum Virtualize might use this specification better in future to become a first-class citizen in cloud deployment. Here are a few enhancements that can widen the ambit of CADF adoption in IBM Spectrum Virtualize.

- In future, events generated due to user initiated action might be recorded in the CADF format.
- Spectrum Virtualize might log an initiation part (when the event/action starts). Right now, it is result focused (care to log only when an event/action completes with result).
- Spectrum Virtualize is publishing events related to Transparent Cloud Tiering as a first step and later might include other events related to different services provided by it.
- Spectrum Virtualize might publish them as part of *informational* document attached with each OpenStack release.
- You can see the CADF specification as an evolving entity. One such evolution is adding new key-value pairs versus using *attachments* using Internet Assigned Numbers Authority (IANA) mime types. This might erase the requirement of having attachment altogether from the Spectrum Virtualize perspective.

The CADF working group might potentially consider the following items in future versions of this specification:

- Support for summarization of sets of like events into a single CADF event record
- Support for aggregation of sets of like events into a single CADF event record
- Support for secure signing of CADF events, logs, and reports.
- Support for indicating precision (granularity) of a CADF timestamp
- Provide guidance on use of metric standards for use in the CADF metric data type

Resources

The following websites provide useful references to supplement the information contained in this paper:

- What is CADF
  http://www.dmtf.org/standards/cadf
- CADF working group link
  https://www.dmtf.org/standards/cadfwg
- CADF latest specification
  http://www.dmtf.org/sites/default/files/standards/documents/DSP0262_1.0.0.pdf
- Who manage the CADF
  https://en.wikipedia.org/wiki/Distributed_Management_Task_Force
- A developer works article from a DMTF founding member
  https://developer.ibm.com/open/2015/12/24/cadf-on-go-is-a-go

About the authors

Sandeep Bangur is a senior staff software engineer in the IBM Systems SVC organization. You can reach Sandeep at sandeep.bangur@in.ibm.com.

Dhiraj Verma is a senior storage test engineer in the IBM Systems SVC organization. You can reach Dhiraj at dhiverma@in.ibm.com.

IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America

Please recycle