



CERN Compute Accounting

miguel.coelho.santos@cern.ch

IT-CM

pre-GDB April 2016

CERN Cloud accounting

- cASO is a pluggable extractor of Cloud Accounting Usage Records for OpenStack.
 - https://wiki.egi.eu/wiki/Federated_Cloud_Accounting
 - <https://github.com/IFCA/caso>
- Scale tests with Ceilometer data showed scaling issue with Ceilometer
 - many (many!) keystone calls
 - very large infrastructure requirements
 - Gnocchi is being tested, will need to be reassessed later

CERN Cloud accounting

- Adopting cASO for producing Cloud reports of CERN Openstack
 - Data collected from Nova
 - Patch to support Keystone v3 (for upstream),
 - Hope to have a productive collaboration.

External Batch

- New project for using HTCondor as the Batch System on fixed resources in commercial Cloud(s).
- Can Batch cluster use commercial clouds in a similar manner to the Wigner Data Centre extension, i.e. with flat capacity?

Accounting on commercial clouds

- Commercial cloud providers have similar usage reports that are used for billing

Amazon Web Services

Usage Report

Service	Operation	UsageType	StartTime	EndTime	UsageValue
AmazonEC2	RunInstances	EUC1-BoxUsage:t2.micro	04/01/16 00:00:00	04/02/16 00:00:00	12
AmazonEC2	RunInstances	EUC1-DataTransfer-In-Bytes	04/01/16 00:00:00	04/02/16 00:00:00	5115349
AmazonEC2	RunInstances	EUC1-DataTransfer-Out-Bytes	04/01/16 00:00:00	04/02/16 00:00:00	1029352
AmazonEC2	RunInstances	EUC1-C3DataTransfer-In-Bytes	04/01/16 00:00:00	04/02/16 00:00:00	85319
AmazonEC2	RunInstances	EUC1-C3DataTransfer-Out-Bytes	04/01/16 00:00:00	04/02/16 00:00:00	31470

Usage Report

Date	Meter Category	Meter Id	Meter Sub-category	Meter Name	Consume Quantity
5/4/2015	"Virtual Machines"	"fee1c571-5e0d-47ef-b031-3ba67f6a67c2"	"BASIC.A1 VM (Windows)"	"Compute Hours"	24
5/4/2015	"Virtual Machines"	"96cbeb60-b981-46b1-8ae6-101fd3367c34"	"Standard_D1 VM (Windows)"	"Compute Hours"	4.083355
5/5/2015	"Virtual Machines"	"fee1c571-5e0d-47ef-b031-3ba67f6a67c2"	"BASIC.A1 VM (Windows)"	"Compute Hours"	24

Billing

Billing Period	Subscription Name	Order Id	Description	Quantity	Currency	Value
201506(5/3/2015 - 6/2/2015)	"Virtual Machines"	"BASIC.A1 VM (Windows)"	"Compute Hours"	112	USD	\$5.71 USD
201506(5/3/2015 - 6/2/2015)	"Networking"	"Public IP Addresses"	"IP Address Hours"	4.1	USD	\$0.02 USD
201506(5/3/2015 - 6/2/2015)	"Virtual Machines"	"Standard_D1 VM (Windows)"	"Compute Hours"	4.083355	USD	\$0.35 USD

Google Cloud Platform

Report Date	MeasurementId	Quantity	Unit	Resource URI	Resource ID	Location
02/13/2014	<code>com.google.cloud/services/compute-engine/VmimageN1Standard_1</code>	86400	seconds	<code>https://www.googleapis.com/compute/v1/projects/myproject/zones/us-central1-a/instances/my-instance</code>	16557630484	us-central1-a

Summary

- AWS, Google, Azure, Rackspace produce usage reports, typically for one or two meters per service, for example:
 - time and bandwidth,
 - time and capacity,
 - storage and #operations,
 - just time,
 - just #operations

Summary

- These usage reports typically have:
 - one line per meter (not per VM, no complex records),
 - a daily frequency,
 - a simple format (CSV, JSON, both),
 - stored in a bucket.
- A monthly billing is produced:
 - aggregated data (per “service”, per meter),
 - to each of these, a rating is applied to calculate cost.

Final remarks:Resource Monitoring

- Our recent test with commercial clouds have shown some interesting points (see Domenico)
- It's also important to monitor your resources and compare with the usage metering reports and billing (note, the problem also exists for owned data centres).
- Idling or failed nodes have a cost (good automated node orchestration is key).

Final remarks:Resource Performance

- Typically Commercial Cloud resources have somewhat varying performance:
 - from resource to resource,
 - over time.
- Important to monitor resource performance, which is important for efficiency and OPEX.
- Hard to do accounting based on (fluctuating) performance measurements. It's like a (fluctuating) efficiency rating applied to time (wall-clock). Being that efficiency and throughput are complex measurements.
- Typical solution seems to be real-time monitoring of application performance coupled sometimes with usage of reserved instances.



www.cern.ch