CAOS: a tool for OpenStack accounting management

F. Chiarello, S. Traldi, P. Andreetto

INFN-Padova

What is CAOS?

The analysis and understanding of resources utilization in shared infrastructures, such as cloud environments, is crucial in order to provide better performance, administration and capacity planning.

CAOS (Control Application for OpenStack) is a tool which we have been implementing to collect, manage and present the data concerning resource usage of our OpenStack-based cloud infrastructures hosted at INFN-Padova cluster: the CloudVeneto and the INFN-PADOVA-STACK instance of the EGI Federated Cloud.

By gathering data from both the Ceilometer service and OpenStack API, CAOS enables us to track resource usage at different levels (e.g. per project), in such a way that both current and past consumption of resources can be easily determined, stored and presented.

The issue

Ceilometer, the OpenStack component responsible to collect and manage accounting information, has some limiting problems related to the way it handles information:

- the imbalance between storage and data retention requirements (due to redundant metadata accompaining each data point)
- the complexity in computing custom metrics
- metrics refer to single resources: aggregating data (e.g. per project) must be done on the fly and can be extremely slow
- it is not straightforward to get information involving different metrics (e.g. the efficiency in CPU usage)
- there exists some known bugs (fixed on some versions) that require a post-processing to possibly correct the data

Architecture

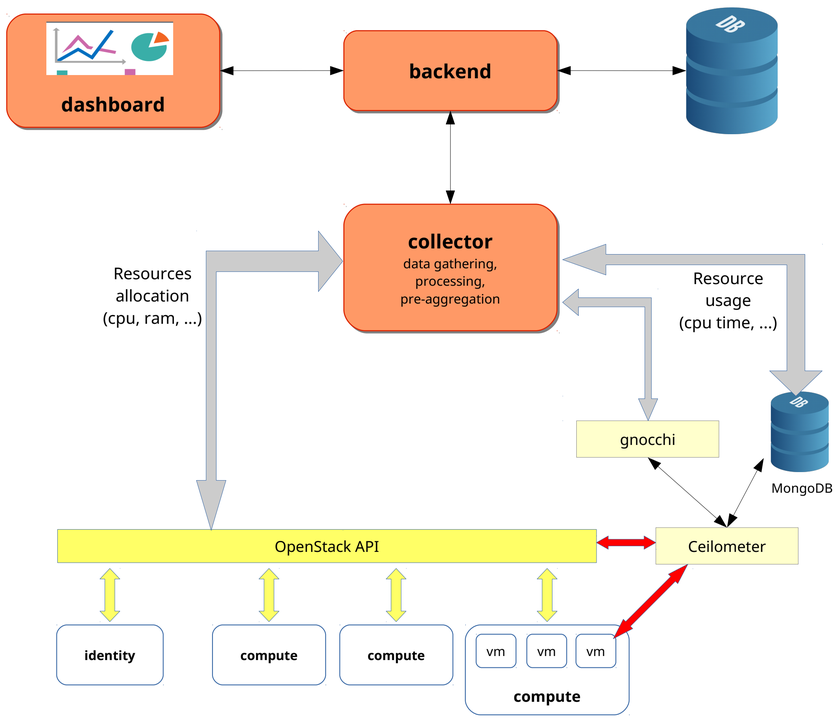

The collector gathers data at regular intervals from both the OpenStack API and the Ceilometer service. The data is then analyzed and pre-aggregated at coarser granularity (e.g. hourly), and metrics involving operations across different metrics are computed.

The backend provides a time series framework for writing and reading metrics. The data can be aggregated or downsampled at a given resolution or time range to provide resources accounting.

The dashboard allows the Cloud administrator to easily get resource usage information for a given time slot period.

Features

- Metrics aggregated per project and per compute node:
  - CPU time, Wall Clock time, CPU efficiency
  - Quotas
  - Instances
  - VCPUs
  - VRAM
- Support for different OpenStack releases:
  - Kilo
  - Mitaka
  - Newton
  - Ocata
- Support for both MongoDB and Gnocchi
- Docker enabled

Resources monitoring

The collected data can be used to monitor resources used and allocated to projects.

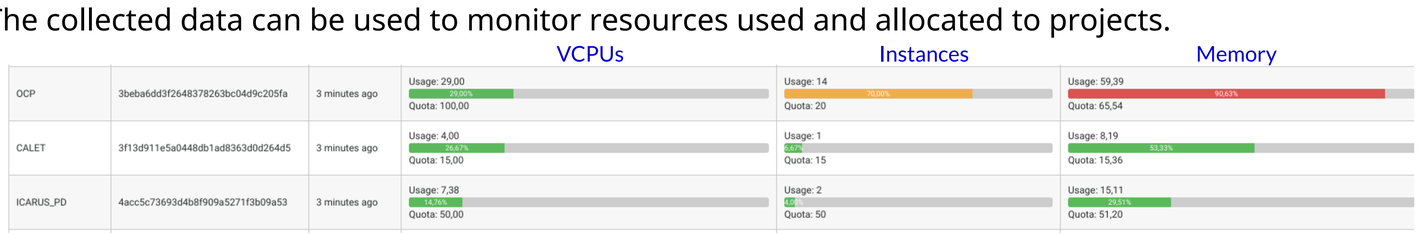

It's also possible to monitor how the compute nodes have been used.

Name	Last Updated	State / Status	Running VMs	Workload	Allocated VCPUs	Allocated VRAMs	Allocated VCPUs (w.r.t bare)	Allocated VRAMs (w.r.t bare)	Load
cld-nl-03.cloud.pd.infn.it	11 minutes ago	up / enabled	19	0	56,25% 72 of 128	31,10% 39,72GB of 127,73GB	225,00% 72 of 32	40,43% 39,72GB of 98,26GB	5m 356,00% 10m 379,00% 15m 389,00%
cld-nl-08.cloud.pd.infn.it	11 minutes ago	up / enabled	17	0	66,41% 85 of 128	20,22% 34,39GB of 170,10GB	265,63% 85 of 32	26,28% 34,39GB of 130,85GB	5m 20,00% 10m 17,00% 15m 15,00%

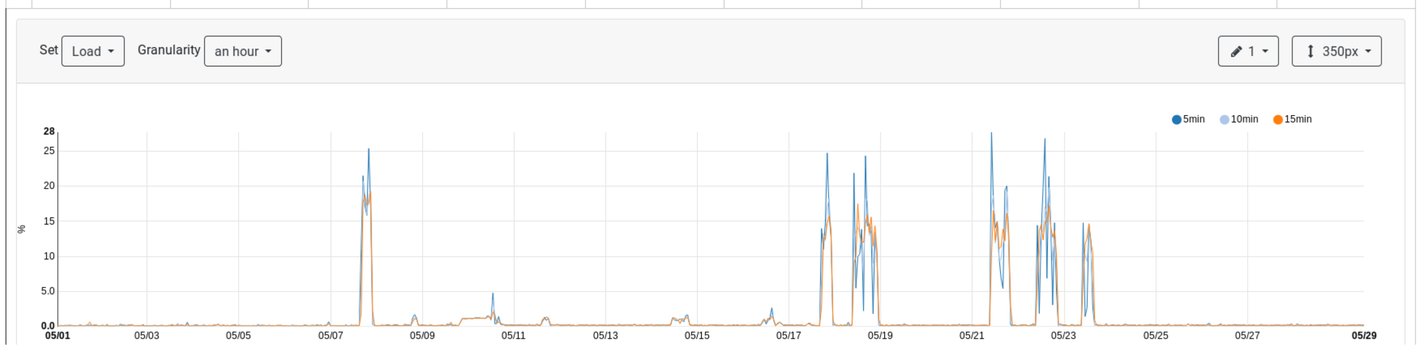

Accounting

Accounting information (e.g. wall clock time, CPU time, efficiency) can be displayed and aggregated for a selected time window for one or more projects.

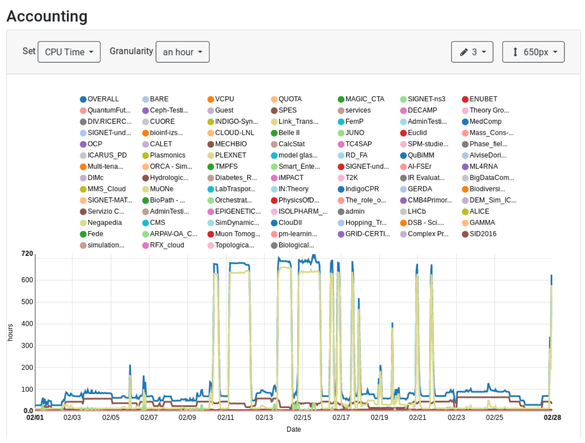

Aggregates from 2017-02-01 00:00 CET to 2017-02-28 00:00 CET

Project	Sum [hours]	Percentage
OVERALL	**106.298,32**	**100,00%**
CMS	73.667,72	69,30%
SPES	23.759,09	22,35%
AdminTesting	4.814,98	4,53%

CPU time

Project	Sum [hours]	Percentage
OVERALL	**1.079.272,28**	**100,00%**
SPES	531.792,20	49,27%
CMS	242.158,97	22,44%
LHCb	63.500,56	5,88%

wall clock time

Used resources can be analyzed for specific use cases (e.g. CPU intensive workload of clusters).