CAOS: a tool for accounting management in OpenStack

P. Andreetto[1], F. Chiarello[1], F. Costa[1], A. Crescente[1], A. Dorigo[1], S. Fantinel[2], F. Fanzago[1], E. Konomi[1], M. Segatta[1], M. Sgaravatto[1], S. Traldi[1], N. Tritto[1], M. Verlato[1], L. Zangrando[1]

[1] INFN-PD [2] INFN-LNL

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: the Cloud Area Padovana 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

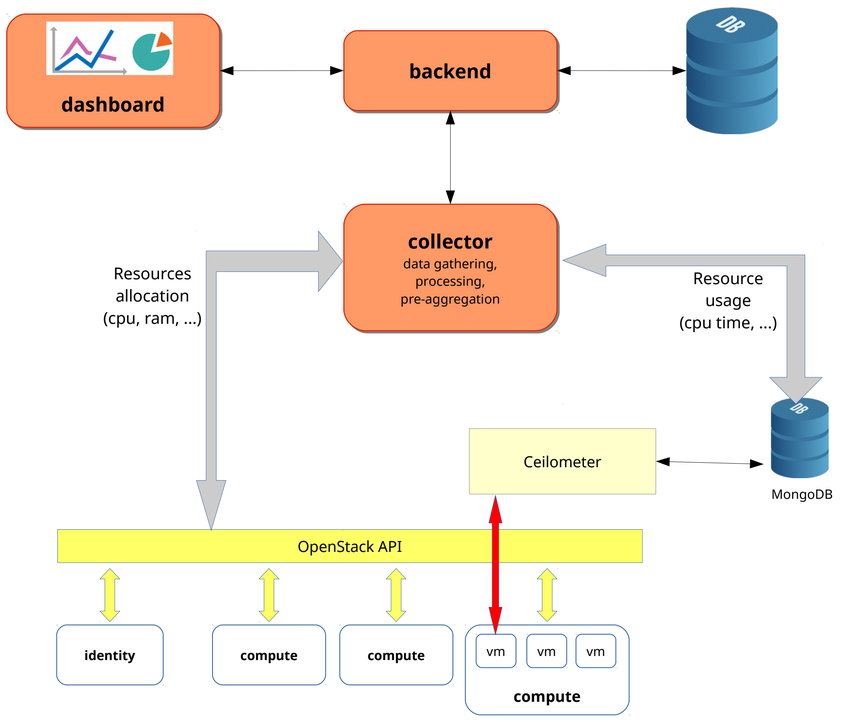

The collector gathers data at regular intervals from both the OpenStack API and the Ceilometer service. The data is then analyzed and pre-aggregated at coarser granularity (e.g. hourly), and metrics involving operations across different metrics are computed.

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.

Resources monitoring

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

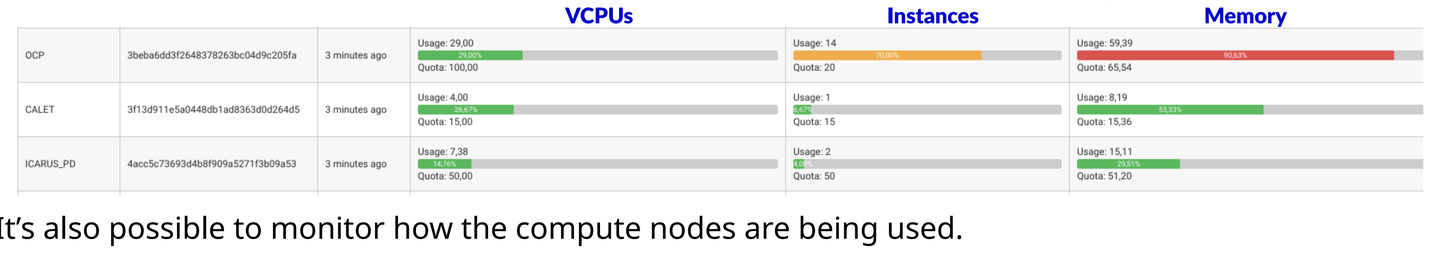

It's also possible to monitor how the compute nodes are being used.

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-07.cloud.pd.inf	11 minutes ago	up / enabled	14	0	60,94% 78 of 128	48,33% 93,18GB of 192,80GB	243,75% 78 of 32	72,50% 93,18GB of 128,53GB	5m 0,37% 10m 0,40% 15m 0,41%
cld-nl-06.cloud.pd.inf	11 minutes ago	up / enabled	13	0	51,56% 66 of 128	64,40% 93,18GB of 144,70GB	206,25% 66 of 32	96,59% 93,18GB of 96,47GB	5m 1,07% 10m 1,29% 15m 1,15%
cld-nl-08.cloud.pd.inf	11 minutes ago	up / enabled	12	0	64,84% 83 of 128	54,44% 104,96GB of 192,80GB	259,38% 83 of 32	81,66% 104,96GB of 128,53GB	5m 1,40% 10m 1,32% 15m 1,26%

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.

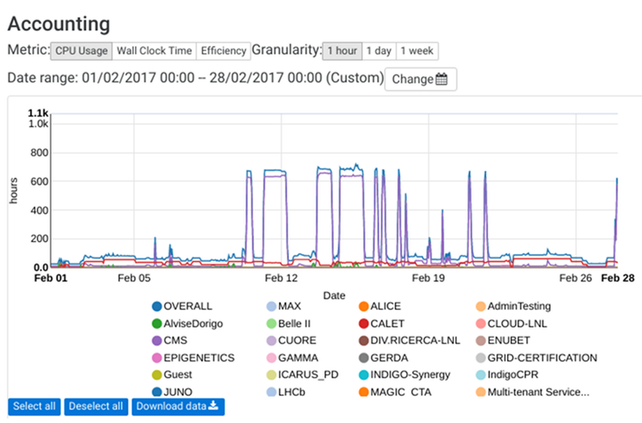

Aggregates from 2017-02-01 00:00 CET to 2017-02-28 00:00 CET

Project	Sum [hours]	Percentage
OVERALL	106.298,32	100,00%
CMS	73.667,72	69,30%
SPES	23.759,09	22,35%

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%

wall clock time

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