



Contribution ID: 389

Type: oral presentation

Integrated Monitoring-as-a-service for Scientific Computing Cloud applications using the Elasticsearch ecosystem

Monday 13 April 2015 17:45 (15 minutes)

The INFN computing centre in Torino hosts a private Cloud, which is managed with the OpenNebula cloud controller. The infrastructure offers IaaS services to different scientific computing applications. The main stakeholders of the facility are a grid Tier-2 site for the ALICE collaboration at LHC, an interactive analysis facility for the same experiment and a separate grid Tier-2 site for the BES-III collaboration, plus an increasing number of other smaller tenants. The dynamic allocation of resources to tenants is partially automated. This feature requires detailed monitoring and accounting of the resource usage.

We set up a monitoring framework to inspect the site activities both in terms of IaaS and applications running on the hosted virtual instances. For this purpose we used the Elasticsearch, Logstash and Kibana stack. The infrastructure relies on an SQL database back-end for data preservation and to ensure flexibility to switch to a different monitoring solution if needed. The heterogeneous accounting information is transferred from the database to the Elasticsearch engine via a custom Logstash plugin. Each use-case is indexed separately in Elasticsearch and we setup a set of Kibana dashboards with pre-defined queries in order to monitor the relevant information in each case.

For the IaaS metering, we developed sensors for the OpenNebula API. The IaaS level information gathered through the API is sent to the MySQL database through an ad-hoc developed RESTful web service. Moreover, we have developed a billing system for our private Cloud, which relies on the RabbitMQ message queue for asynchronous communication to the database and on the ELK stack for its graphical interface.

Concerning the application level, we used the Root plugin TProofMonSenderSQL to collect accounting data from the interactive analysis facility. The BES-III virtual instances used to be monitored with Zabbix, as a proof of concept we also retrieve the information contained in the Zabbix database.

Finally, we have defined a model for monitoring-as-a-service, based on the tools described above, which the Cloud tenants can easily configure to suit their needs.

In this way we have achieved a uniform monitoring interface for both the IaaS and the scientific applications, mostly leveraging off-the-shelf tools.

Primary authors: GUARISE, Andrea (Istituto Nazionale Fisica Nucleare (IT)); BERZANO, Dario (CERN); MASERA, Massimo (Universita e INFN (IT)); VALLERO, Sara (Universita e INFN (IT)); Dr BAGNASCO, Stefano (I.N.F.N. TORINO); LUSSO, Stefano (INFN-TO)

Presenter: VALLERO, Sara (Universita e INFN (IT))

Session Classification: Track 7 Session

Track Classification: Track7: Clouds and virtualization