

Synergy: a service for optimising resource allocation in cloud based environments (remote presentation)

Wednesday, 16 March 2016 17:00 (20 minutes)

Computing activities performed by user groups in Public Research and Administration are usually not uniformly distributed over periods in the order of a year. The amount of computing resources effectively used by such groups may vary significantly. In those well defined environments customers generally stipulate contracts with their Data Centers to guarantee the provisioning of an average computing level during a long period of time rather than of an agreed amount of computing resources that should be available at any given point in time.

In the current OpenStack scheduling model only fixed quotas can be granted to user groups. Those resources cannot be exceeded by one group even if there are unused resources allocated to other groups. Therefore, in a scenario of full resource usage for a specific project, new requests are simply rejected.

When resources are statically allocated amongst user groups, the global efficiency in the Data Center's resource usage may become quite low. The recently started European INDIGO-DataCloud project will address this issue through 'Synergy', a new advanced scheduling service to be integrated in the OpenStack Cloud Management Framework. 'Synergy' will adopt a resource provisioning model based on a fair-share algorithm to maximize resources usage. The INDIGO team is considering to use the SLURM multifactor fair-share algorithm for its first release. Beside the improved usage of resources, the algorithm guarantees that those resources are equally distributed among users and groups by accounting for the portion of the resources allocated to them and the resources already consumed in previous usage periods. Moreover, the mechanism will provide a persistent priority-queuing for handling user requests that can't be immediately fulfilled. Those requests will be processed as resources will become available again.

As our improvements are well structured and only have to be applied once to the upstream code repository, we don't expect significant maintenance efforts.

Starting from a list of selected requirements, 'Synergy' has to satisfy, this paper will provide a high level architecture of services, focusing on integration and interoperability aspects of existing OpenStack components, preferably those, which don't need to be extended. Along with preliminary results, we'll elaborate on the status of the current 'Synergy' implementation.

Summary

The presentation will describe 'Synergy', a new advanced scheduling service developed by the INDIGO-DataCloud European project, to be integrated in the OpenStack Cloud Management Framework.

Primary author: Dr ZANGRANDO, L. (INFN - Sez. Padova)

Co-authors: Dr FANZAGO, Federica (INFN); Dr VERLATO, Marco (INFN); Dr SGARAVATTO, Massimo (INFN)

Presenters: Dr ZANGRANDO, L. (INFN - Sez. Padova); Dr VERLATO, Marco (INFN)

Session Classification: Infrastructure Clouds and Virtualisation Session I

Track Classification: Infrastructure Clouds and Virtualisation