



Lisa Zangrando
INFN Padova

Synergy

Advanced scheduling in OpenStack

Synergy

cloud service developed in the context of the INDIGO-DataCloud European project which aims to develop a new cloud software platform for the scientific community

- <https://www.indigo-datacloud.eu/>

Main objective

enable a more effective and flexible resource allocation and utilization in open Clouds such as OpenStack


- **In the current OpenStack model:**

- resource allocation model: static partitioning
 - based on granted and fixed quotas (one per project)
 - the quotas cannot be exceeded
 - the quotas cannot be shared among projects
- scheduler too simple
 - based on the immediate First Come First Served (FCFS)
 - user requests are rejected if not immediately satisfied

- data center: very low global efficiency and increased cost
- 20 years old problem we solved by adopting batch systems
 - enhancement of our data center resources utilization from <50 to 100%
- **INDIGO addresses this issue through Synergy**

Synergy

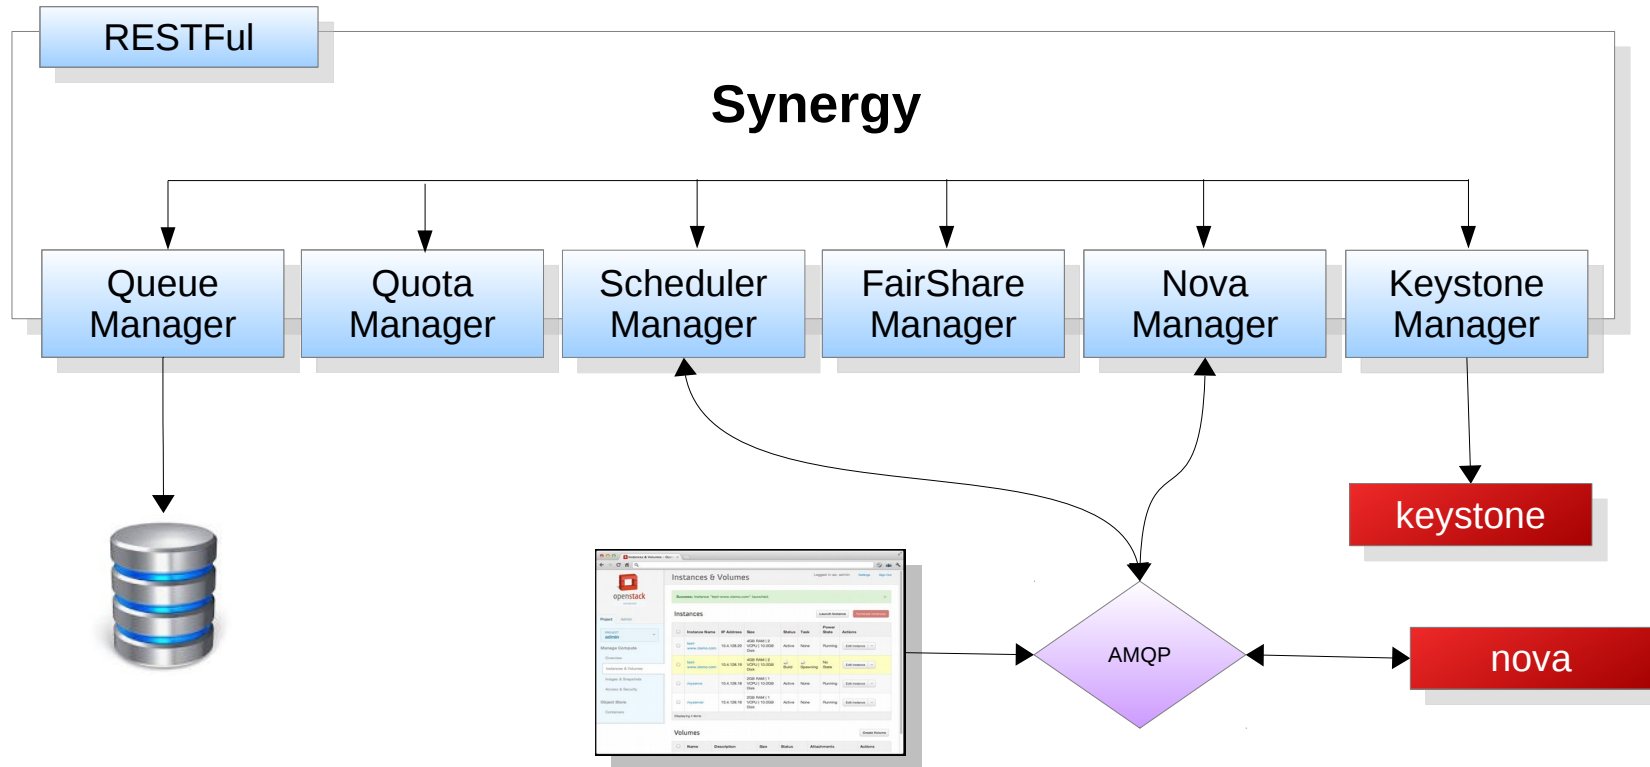
- It is a cloud service designed for executing tasks in OpenStack
- It is composed by a collection of specific and independent pluggable functionality (managers) executed periodically or interactively through a RESTful API



- **By implementing the same logic of batch systems**
- Synergy with six specific managers provides an advanced resource allocation and scheduling model

- **cloud resources can now be shared among different OpenStack projects**
 - overcomes the static partitioning limits
 - maximizes the resource utilization
- **shared resources are fairly distributed among users and projects**
 - user priority
 - project share
- **requests that can't be immediately fulfilled are enqueued (not rejected!)**

Synergy scheduler managers

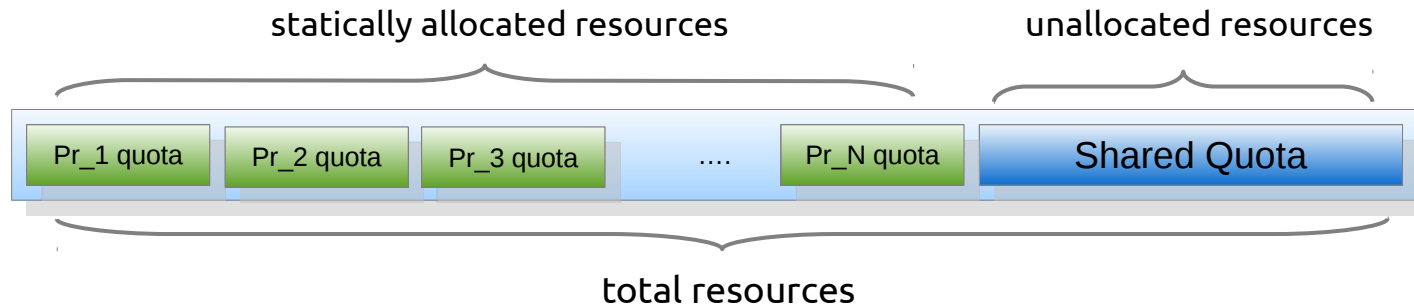


- With Synergy the OpenStack projects can now consume extra shared resources in addition to those statically assigned
- Projects can access to two distinct quota kinds:

- **private quota:**
 - the standard (i.e. fixed and statically allocated) OpenStack quota
- **shared quota:**
 - extra resources shared among projects
 - handled by Synergy
 - its size can change dynamically: amount of resources not statically allocated
 - the user requests that cannot be immediately satisfied are inserted in a **persistent priority** queue

The Shared Quota

- amount of resources not statically allocated
- its size is calculated as the difference between the total amount of cloud resources and the total resources allocated to the private quotas



- Only the projects selected by the administrator can access to the shared quota beside to their own private quota

The scheduling model

- The shared resources are fairly distributed among users according to specific fair-share policies defined by the administrator:

- list of projects allowed to access to the shared quota
- definition of shares (%) on resource usages for the selected projects
- max resource lifetime
 - VMs and Containers (instantiated via nova-docker)
 - this is needed to enforce the fair-sharing

The status

- Synergy released by INDIGO
 - support for Liberty, Mitaka and Newton
 - next release: March 2017
- Code in launchpad
 - <https://launchpad.net/synergy-service>
 - <https://launchpad.net/synergy-scheduler-manager>
- Documentation
 - <https://indigo-dc.gitbooks.io/synergy/content>
- the ultimate goal is to have it integrated in the Official OpenStack distribution

Questions?

