
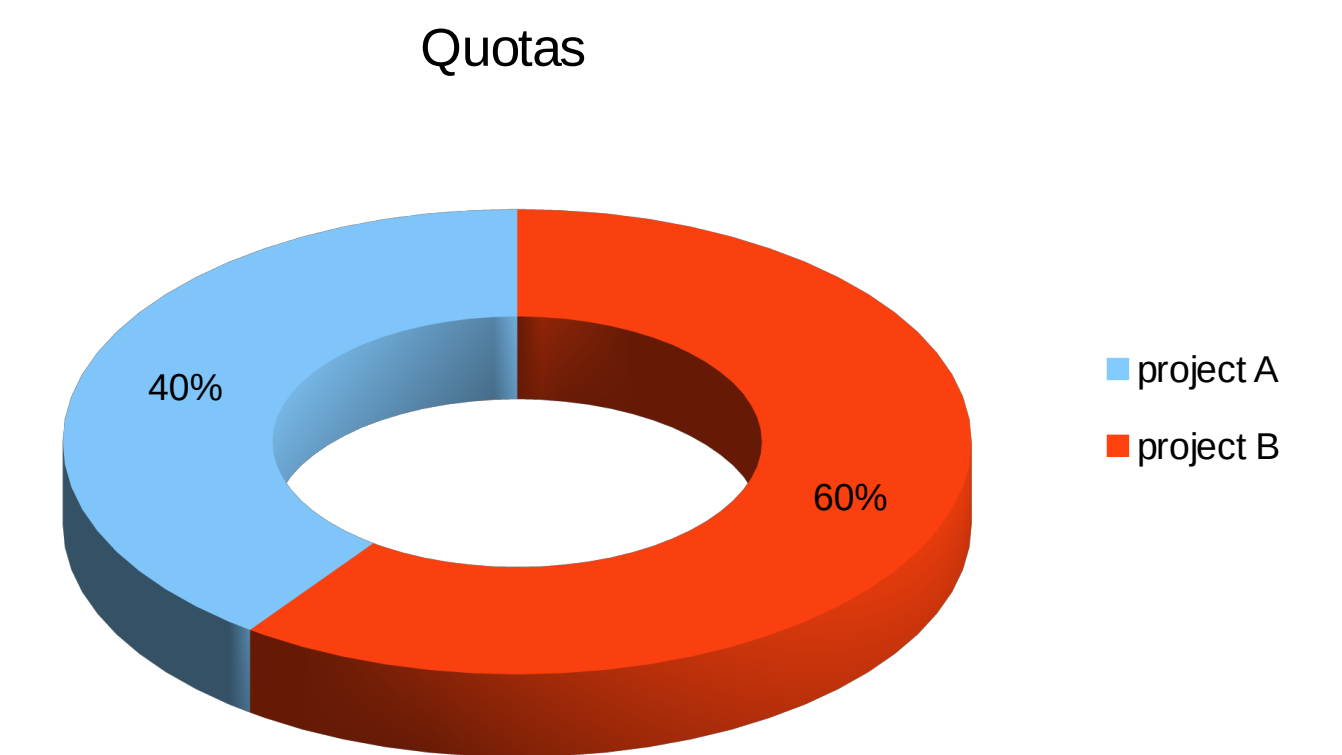


The issue

OpenStack allows only a static partitioning of resources among the user communities through fixed quotas. The quotas cannot be neither shared nor exceeded. In case of quota saturation new user requests are simply rejected even if there are several unused resources but allocated to different projects.



Synergy

This issue has been addressed by Synergy, implemented in the context of the INDIGO-DataCloud project. It provides a new advanced allocation model based on a new kind of quota: "shared". The shared quota is a subset of the total resources, can be shared among projects and allows to en-queue requests that cannot be immediately fulfilled.

Its resources are fairly distributed according to a set of fair-share policies defined by the Cloud administrator as:

- the shares for the different projects (e.g. project A=70%, project B=30%)
- the shares among different users of the same project (optional)
- the maximum lifetime (e.g. 48 hours) of the relevant resources (virtual machines, containers)

Projects can access to two distinct quota kinds:

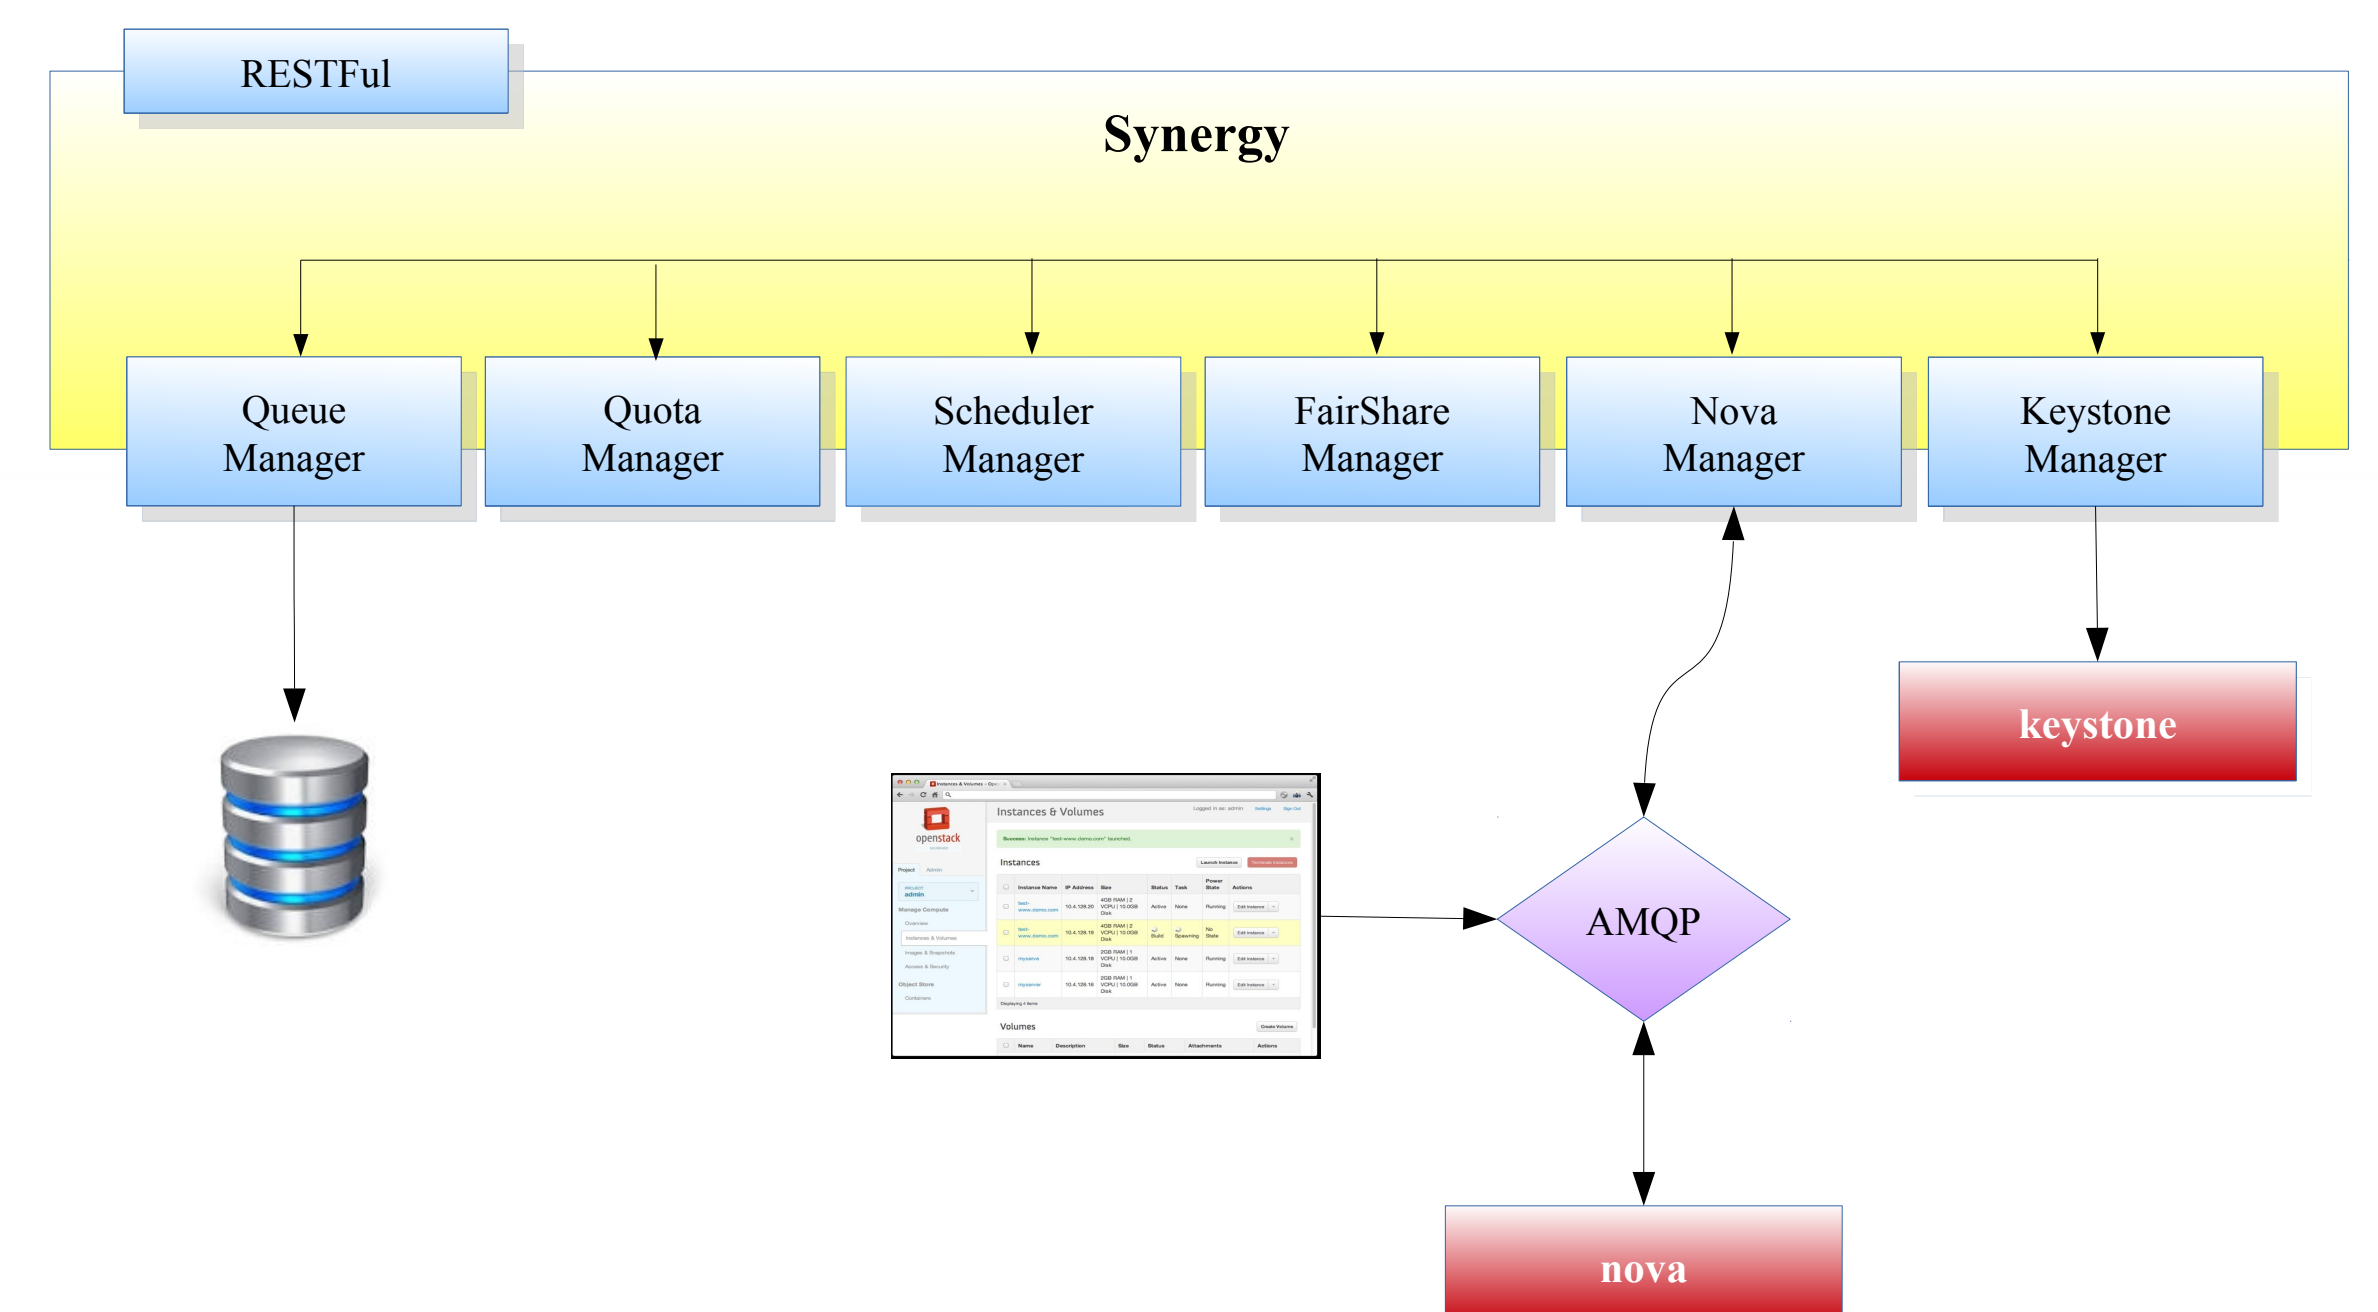
- private quota: the standard OpenStack quota, fixed and statically allocated
- shared quota: access to extra resources shared among different projects

The Synergy architecture

The advanced resource allocation model has been implemented with a set of specific Synergy managers:

- **FairShare-Manager:** implements the main fair-share scheduling logic adopting the SLURM multifactor priority algorithm
- **Scheduler-Manager:** processes the requests from Nova-Manager and schedules them according with the private or shared quota policy
- **Queue-Manager:** provides a persistent priority queue service where the en-queued requests are scheduled with a backfill strategy
- **Quota-Manager:** it is in charge of handling the quota of all projects
- **Nova-Manager:** it interacts with the OpenStack Nova components
- **Keystone-Manager:** it interacts with the OpenStack Keystone service

Synergy does not replace any existing OpenStack service



Testing results in production environment

Tests were carried out at INFN-Padova OpenStack production site of the EGI Federated Cloud. Shared resources amounted the 20% of the total and more than 20.000 VMs (with different flavors) were executed over two days. The project A and project B resource usage accounted at the end of each period was the expected one (70% and 30%) within 1%. Tests coexisted and did not interfere/degrade normal operations of the other production projects (not involved in the fair-share computation).