

Elastic CNAF DataCenter extension via opportunistic resources

Friday, 18 March 2016 10:00 (30 minutes)

CNAF/ Bologna, the biggest WLCG Computing Center in Italy, serves all WLCG Experiments plus more than other 20 non WLCG Virtual Organizations, and currently deploys more than 180 kHS06 of Computing Power and more than 20 PB of Disk and 40 PB of tape via a GPFS SAN.

The Center has started a program to evaluate the possibility to extend its resources on external entities, either commercial or opportunistic or simply remote, in order to be prepared for future upgrades or temporary burst in the activity from experiments. The approach followed is meant to be completely transparent to users, with additional external resources directly added to the CNAF LSF batch system; several variants are possible, like the use of VPN tunnels in order to establish LSF communications between hosts, a multi-master LSF approach, or in the longer term the use of HTCondor. Concerning the storage, the simplest approach is to use Xrootd fallback to CNAF storage, unfortunately viable only for some experiments; a more transparent approach involves the use of GPFS/AFM module in order to cache files directly on the remote facilities. In this presentation we focus on the technical aspects of the integration, and assess the difficulties using different remote virtualisation technologies, as made available at different sites. A set of benchmark is provided in order to allow for an evaluation of the solution for CPU and Data intensive workflows.

Summary

CNAF/ Bologna, the biggest WLCG Computing Center in Italy, serves all WLCG Experiments plus more than other 20 non WLCG Virtual Organizations, and currently deploys more than 180 kHS06 of Computing Power and more than 20 PB of Disk and 40 PB of tape via a GPFS SAN.

The Center has started a program to evaluate the possibility to extend its resources on external entities, either commercial or opportunistic or simply remote, in order to be prepared for future upgrades or temporary burst in the activity from experiments. The approach followed is meant to be completely transparent to users, with additional external resources directly added to the CNAF LSF batch system; several variants are possible, like the use of VPN tunnels in order to establish LSF communications between hosts, a multi-master LSF approach, or in the longer term the use of HTCondor. Concerning the storage, the simplest approach is to use Xrootd fallback to CNAF storage, unfortunately viable only for some experiments; a more transparent approach involves the use of GPFS/AFM module in order to cache files directly on the remote facilities. In this presentation we focus on the technical aspects of the integration, and assess the difficulties using different remote virtualisation technologies, as made available at different sites. A set of benchmark is provided in order to allow for an evaluation of the solution for CPU and Data intensive workflows.

Primary authors: Dr DAL PRA, Stefano (INFN-CNAF); Dr CIASCHINI, Vincenzo (INFN-CNAF); Dr BOCCALI, tommaso (INFN)

Presenters: Dr DAL PRA, Stefano (INFN-CNAF); Dr CIASCHINI, Vincenzo (INFN-CNAF); Dr BOCCALI, tommaso (INFN)

Session Classification: Infrastructure Clouds and Virtualisation Session II

Track Classification: Infrastructure Clouds and Virtualisation