

# Testing SLURM batch system for a grid farm: functionalities, scalability, performance and how it works with Cream-CE

*Thursday, 20 September 2012 11:00 (30 minutes)*

## Description of the work

We will show all the work done in order to install and configure the batch system itself together with the security configuration needed.

In this presentation we will show the results of the deep testing that we have done on SLURM, in order to be sure that it will cover all the needed functionalities like: priorities, fairshare, limits, QoS, failover capabilities and others.

We will report also on the possibility of exploiting this batch system within a complex mixed farm environment where grid job, local job and interactive activities are managed exploiting the same batch system.

From a point of view of the scalability we will show how the SLURM batch system is able to deal with the increasing number of node, CPU and jobs served.

We will also show the performance achieved with several client accessing the same batch server.

We also will make some comparison with other available open source batch system both in terms of performance and functionalities.

We will also provide feedback on mixed configuration with SLURM and MAUI as job scheduling.

We will also describe the work done in order to support SLURM on Cream-CE. Indeed SLURM users community has also expressed a lot of interest in having SLURM supported by the CREAM CE. The integration effort already started at the BLAHP (Batch Local ASCII Helper Protocol) layer, and a basic, prototype-level support for standard jobs is ready. Still, due to the wide range of customisation and different deployment models allowed by SLURM, the task of offering an homogeneous interface for all the deployment scenarios will not be a simple remapping of what has been done for other batch systems.

## Wider impact of this work

This activity will give to the computing farm administrator some information required to choose a new open source batch system in order to provide better scalability, without the need of buying costly proprietary batch system. This work will provide information that could help to understand the capabilities and performance of SLURM batch system and provide few feedback on the possibility to use it in a large and complex computing farm infrastructure.

Moreover in this work it will be evident the work needed to support a new batch system in the Cream-CE.

Finally, this work will give the opportunity to use the SLURM as batch system in a computing farm belonging to the EGI/IGI grid infrastructure.

## Printable Summary

As the grid computing farm are increasing in size in terms of nodes but even more in terms of CPU slots available, it become of great interest to have a scheduler solution that could scale up to tens of thousands of CPU slots and hundreds of nodes.

In order to try to keep the Total Cost of Ownership as low as possible it will be preferred to have an easy to use and open source solution.

SLURM is able to fulfil all those requirements and it looks promising also in terms of community that is supporting it, as it is used in several of the TOP500 supercomputing.

For this reason we deeply tested the SLURM batch system in order to prove if it could be a suitable solution.

In the work we will present the result of all the test executed on SLURM batch system and the results of the

development activity carried on in order to provide the possibility to use SLURM as back-end of a Cream-CE.

**Primary author:** Dr DONVITO, Giacinto (INFN)

**Co-authors:** DORIGO, Alvisè (INFN); REBATTO, David (INFN); FRIZZIERO, Eric (INFN); PRELZ, Francesco (INFN); ZANGRANDO, Luigi (INFN); MEZZADRI, Massimo (INFN); SGARAVATTO, Massimo (INFN); ANDREETTO, Paolo (INFN); BERTOCCO, Sara (INFN)

**Presenter:** Dr DONVITO, Giacinto (INFN)

**Session Classification:** Resource Infrastructure Services

**Track Classification:** Resource Infrastructure services (Peter Solagna: track leader)