



Contribution ID: 495

Type: poster presentation

A Grid and cloud computing farm exploiting HTCondor

This work presents the result of several tests that demonstrate the capabilities of HTCondor as batch system for a big computing farm serving both LHC use cases and others scientists.

The HTCondor testbed hosted at INFN-Bari is made of about 300 nodes and 15'000 CPU slots, and meant to sustain about 50'000 job in the queue. The computing farm is used both from Grid users of many VOs (HEP, Bioinformatics, Astrophysics, etc) and from local users, that come from complete different environments: meteorologists, HEP, bioinformatics, humanities, computational chemistry, remote sensing, statistics, etc. These very different use-cases express a wide range of requirements to the batch system.

For instance, the batch system must address both a huge number of very fast jobs and very large MPI based multi-CPU jobs.

The HTCondor configuration tuned in the test aims to allow several scheduling functionalities: Hierarchical FairShare, Quality of Service, priority on users, priority on group, limit on the number of jobs per user/group/queue, job age scheduling, job size scheduling, "consumable resources" scheduling. Also it has been tested the ability to manage different job types: serial job, MPI job, multi-thread jobs, interactive jobs, whole node jobs.

HTCondor has been tested as well on a HPC Cluster with GPU card and Low Latency Infiniband connection, connected to the standard HTC farm at the same time.

Another point of interest is the High Availability configuration for the master node. This possibility is indeed very useful in order to let the cluster working also in case of hardware failure of the server hosting the batch system.

Other interesting features involved in our tests are: capabilities of using ACLs on resources in general; capabilities to deal with the pre-execution and post-execution script, capabilities of Condor_rooster to manage the start-up the worker node as soon as there are job waiting in the queue, and to shutdown them as soon as they are not needed anymore.

A very new item is also instantiating computational resources in a cloud environment. We have developed a plug-in that is able to request the needed resources to our OpenStack local facility, choosing the right image and the flavour that fulfills the user requirements expressed with the ClassAD in the jobs. In this way there is the possibility to support the computational user requests exploiting cloud resources, and this will help the site admin to efficiently and dynamically share resources, among different use-cases (batch jobs, Cloud applications, etc), exploiting the flexibility of a Cloud Computing infrastructure.

Finally, results will be shown in terms of performances achieved in job submission/scheduled/executed, configuration available and behaviour in terms of scheduling features, best practices on how to configure the batch master for High Availability.

The work will also provide the results in terms of stress tests for the configuration of the CREAM Computing Element when exploiting the HTCondor cluster.

Hence, this work summarizes the capabilities of HTCondor when supporting a big computing farm made up of diverse computational resources (Grid, Local HTC Cluster and Cloud Computing).

Authors: Mr ITALIANO, Alessandro (INFN-Bari); Mr DIACONO, Domenico (INFN-Bari); Dr DONVITO, Giacinto (INFN-Bari); Mr VALENTINI, Roberto (INFN-Bari); SPINOSO, Vincenzo (Universita e INFN (IT))

Presenter: Mr ITALIANO, Alessandro (INFN-Bari)

Track Classification: Track4: Middleware, software development and tools, experiment frameworks, tools for distributed computing