



Contribution ID: 64

Type: **Oral Presentation**

## **HammerCloud: An Automated Service for Stress and Functional Testing of Grid Sites**

*Tuesday, 12 April 2011 11:00 (30 minutes)*

### **Overview**

Many research communities rely on EGI resources to process vast quantities of data rapidly. Performance-critical activities such as this give motivation for tools which aid in the design and configuration of a grid site to ensure that its capabilities meet or exceed the requirements of the foreseen user applications.

HammerCloud (HC) is an automated testing service with two main goals. First, it is a stress-testing tool which can deliver large numbers of real jobs to objectively aid in the commissioning of new sites and to evaluate changes to site configurations. Second, it is an end-to-end functional testing tool which periodically sends short jobs to continually validate the site and related grid services.

In addition to delivering the test jobs to grid sites, HC provides a user-friendly web interface to test results. The service presents plots of job efficiency and many performance metrics to enable quick comparisons to past configurations and between other sites.

### **Impact**

The HammerCloud service empowers site administrators to undertake detailed studies of their site's capabilities without requiring any VO-specific knowledge or permissions. With only three clicks, HC users can schedule a test and shortly thereafter performance metrics are made available.

The experiences of the ATLAS Experiment in using HC demonstrate the potential of such a tool. Since late 2008, ATLAS has invested more than 200,000 CPU-days processing HC test jobs globally. The primary focus of the ATLAS stress-testing has been on optimizing the data access method at the sites. In particular, HC was used to compare strategies such as copying input files to a local disk (e.g. using `lcg-cp`, `dccp`, or `xrdcp`) against reading files directly from the site storage element using the local access protocol (e.g. `rrio`, `dcap`, `xrootd`). During large-scale global stress tests such as STEP09, HC was used to simulate the resource requirements of hundreds of real users by delivering a constant stream of up to 15,000 concurrent jobs throughout a two week period. Tests like this have led to I/O optimizations in the ATLAS software resulting in improvements to the overall job throughput on the grid.

For functional testing, HC delivers around ten types of test jobs for ATLAS. These tests validate not only the basic functionality of the grid sites, but are also used to test remote database access, to validate release-candidates of the grid middleware, and to compare data-access methods. Further, a subset of the ATLAS HC functional tests are deemed as critical – consecutive failures of these jobs result in HC taking action to blacklist the site from receiving user jobs. While the site is blacklisted, HC continues to send tests; when the jobs succeed again the site is informed and they can reset their site online at their convenience.

### **Description of the work**

HammerCloud was first developed for the ATLAS experiment at CERN. Prior to the turn on of the Large Hadron Collider (LHC), members of this research community were often participating in coordinated efforts to stress test the grid sites by simultaneously submitting their physics analysis jobs. Collating the results of these multi-user tests was time-consuming, yet the site performance evaluations these tests enabled were critical during the commissioning phase of many grid sites. In the end, these manual tests motivated the development of an automated service to carry out the stress tests and present the results via a web interface.

The HC service is composed of a backend which submits and monitors the test jobs, and a user frontend which allows users to schedule on-demand tests and to watch the progress of running tests or review completed tests. Jobs are submitted and monitored using GANGA; this tool's Grid Programming Interface provides an efficient framework to develop a grid service which needs the flexibility to submit arbitrary applications to any grid backend.

The web interface is developed using Django and is designed to provide common web views in a core HC library while allowing virtual organizations to customize their web views in VO-specific plugins. Example metrics provided by HC include job success rates, timings of the various steps of a grid job (e.g. preparing input files, execution, and storing output files), and I/O metrics including storage latency and throughput values.

The current users of HC include three LHC experiments. ATLAS is the heaviest user, having used the service since its inception in late 2008, while HC plugins for CMS and LHCb were developed in 2010. The transition of HC from a single- to multiple-VO service required the generalization of core components and the development of a plugin architecture. As a result of this work, HC can flexibly accept further additions of plugins for new communities.

## URL

<http://hammercloud.cern.ch>

## Conclusions

In summary, HammerCloud is an automated testing service which can be used for on-demand stress testing and continuous functional testing. Such a tool has been demonstrated to be useful during grid site commission and to evaluate changes to site configurations.

The LHC VOs have a strong interest in the service. ATLAS continues to use the service for ongoing stress tests and site validation tests, while CMS and LHCb are currently incorporating HC into their grid operations procedures. CMS has thirteen HC users who can schedule on-demand stress tests, and a migration to HC from a CMS-specific tool for functional testing is planned. The LHCb instance of HC has been used to validate the rollout of new storage software at RAL, and further integration to the grid operations is in progress.

The flexible architecture of HC will allow future development of new plugins for research communities who have interest in the service.

**Primary author:** Dr VAN DER STER, Daniel (CERN)

**Co-authors:** Dr SCIABA, Andrea (CERN); Dr ELMSHEUSER, Johannes (Ludwig-Maximilians-Universität München); Mr UBEDA GARCIA, Mario (CERN)

**Presenter:** Dr VAN DER STER, Daniel (CERN)

**Session Classification:** User Support Services

**Track Classification:** Producing & Deploying Technology