

Grid Site Testing for ATLAS with HammerCloud

Johannes Elmsheuser¹, Federica Legger¹, Friedrich Hönig¹

Ramon Medrano Llamas², Gianfranco Sciacca³,

Daniel van der Ster²

for the ATLAS collaboration

¹Ludwig-Maximilians-Universität München, Germany, ²CERN, Switzerland, ³Universität Bern, Switzerland

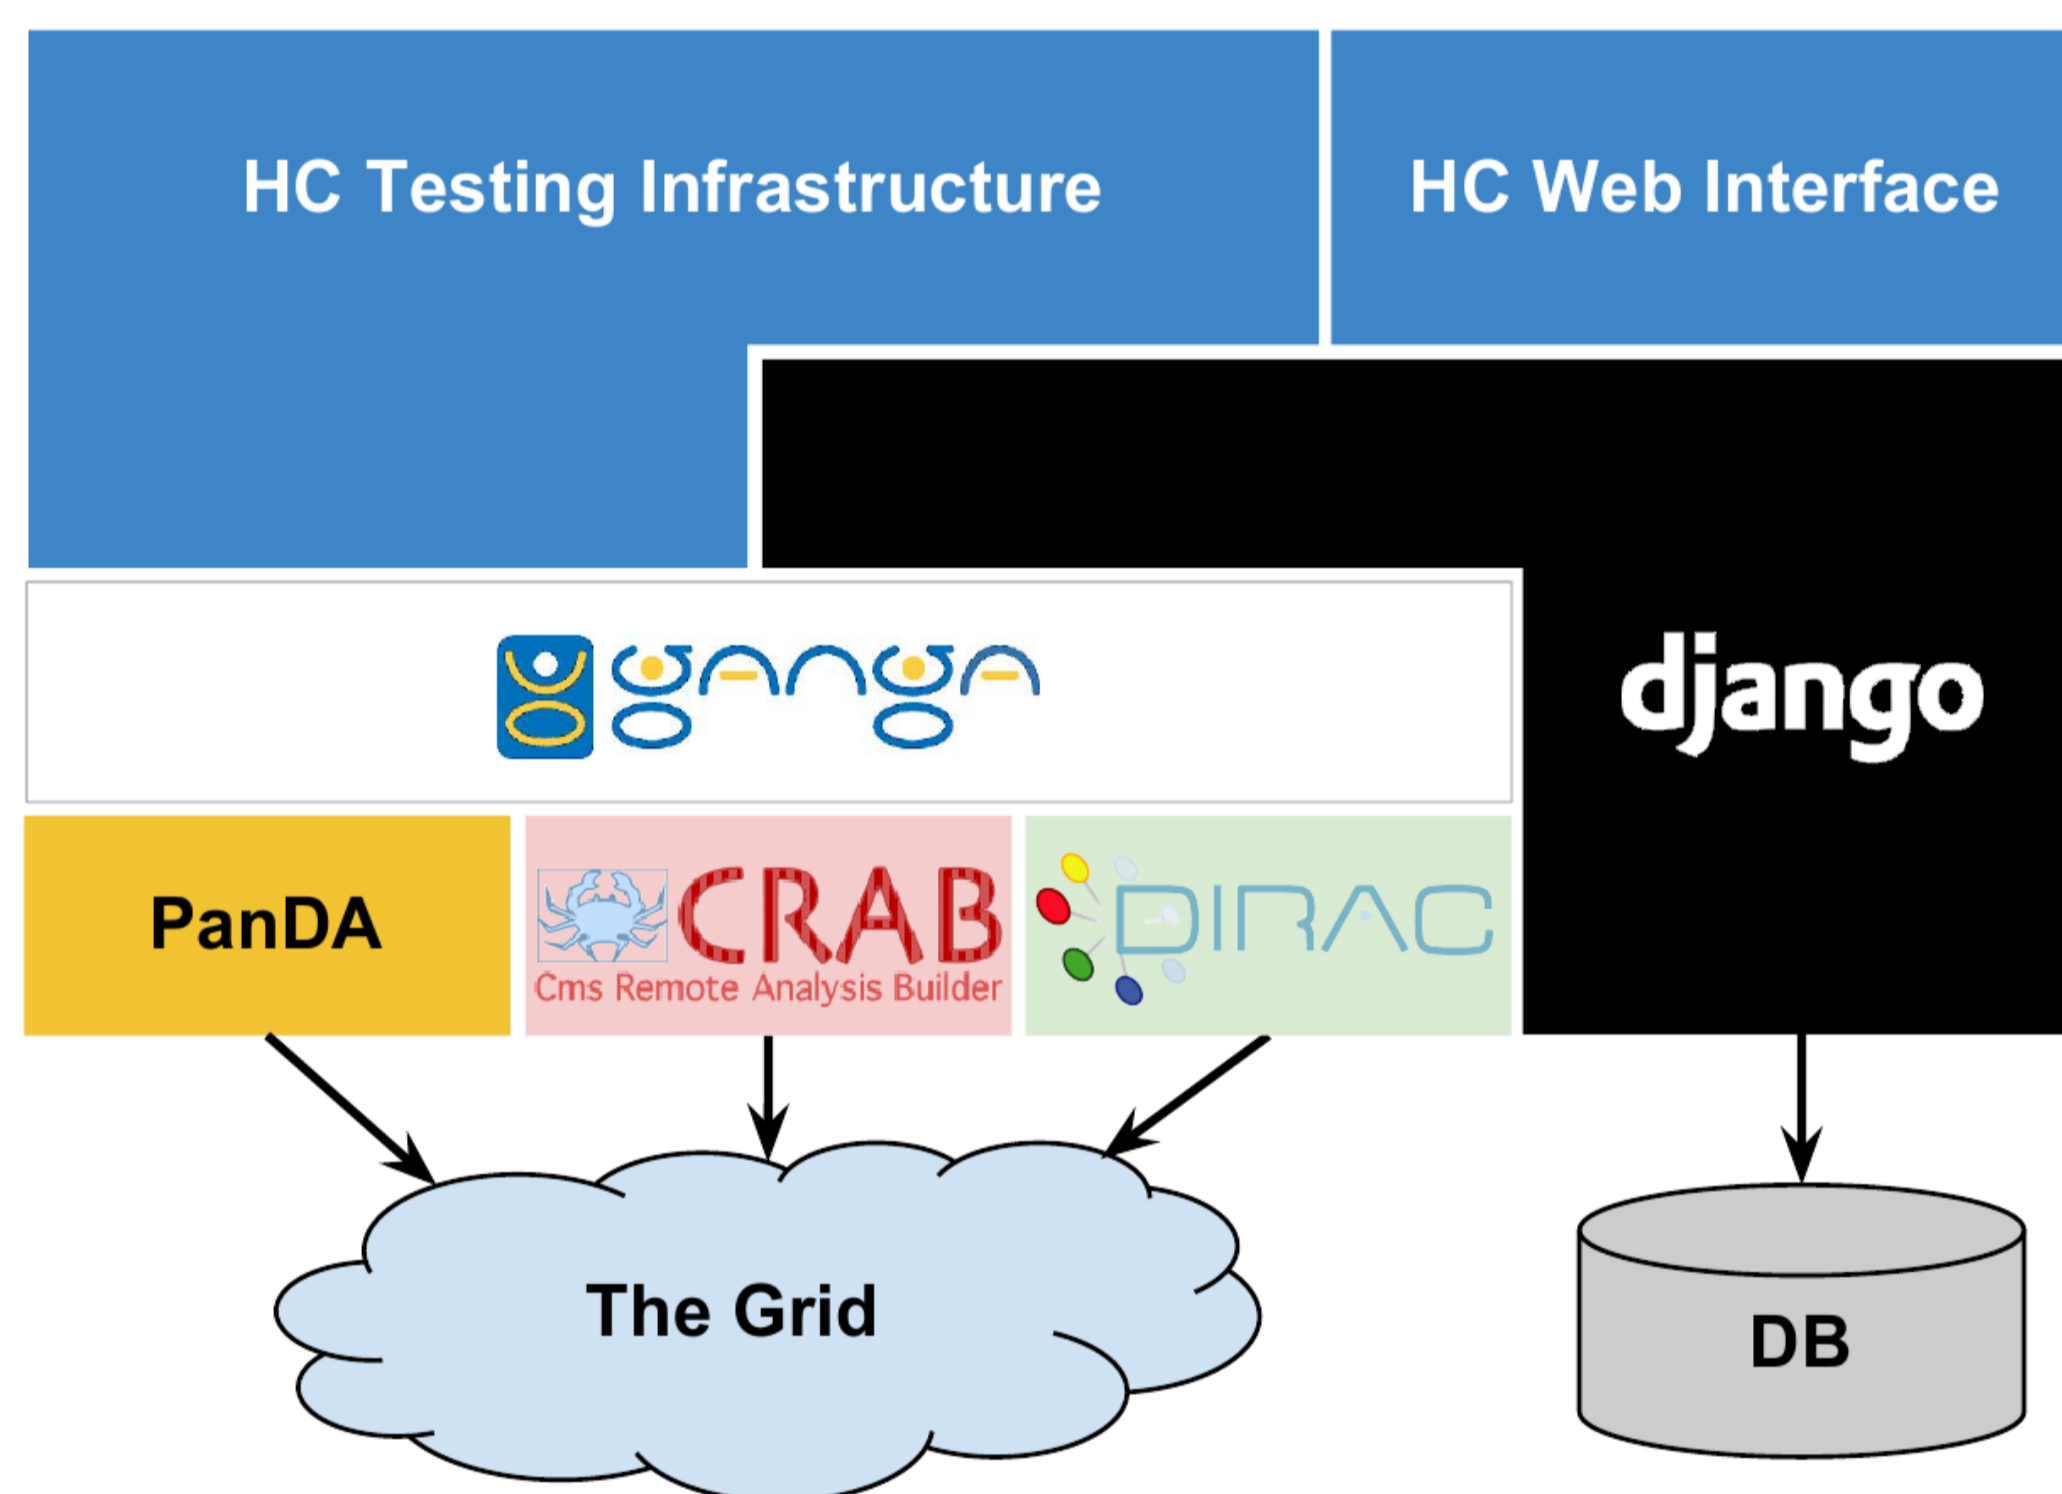


1. Introduction

With the exponential growth of LHC (Large Hadron Collider) data in 2012, distributed computing has become the established way to analyze collider data. The ATLAS grid infrastructure includes more than 130 sites worldwide, ranging from large national computing centers to smaller university clusters. HammerCloud was previously introduced with the goals of enabling VO- and site-administrators to run validation tests of the site and software infrastructure in an automated or on-demand manner. The HammerCloud infrastructure has been constantly improved to support the addition of new test workflows. These new workflows comprise e.g. tests of the ATLAS nightly build system, ATLAS MC production system, XRootD federation FAX and new site stress test workflows. We report on the development, optimization and results of the various components in the HammerCloud framework.



2. HammerCloud overview



HammerCloud (HC) in a nutshell:

- HammerCloud is a PYTHON application using GANGA.
- GANGA is a job management system that has direct access to the interfaces of ATLAS Grid middleware stack: Panda and DQ2/Rucio.
- Grid jobs and test results are stored in the HammerCloud DB and displayed on the webpage: <http://hammercloud.cern.ch>

3. HammerCloud tests

ATLAS functional tests (FT):

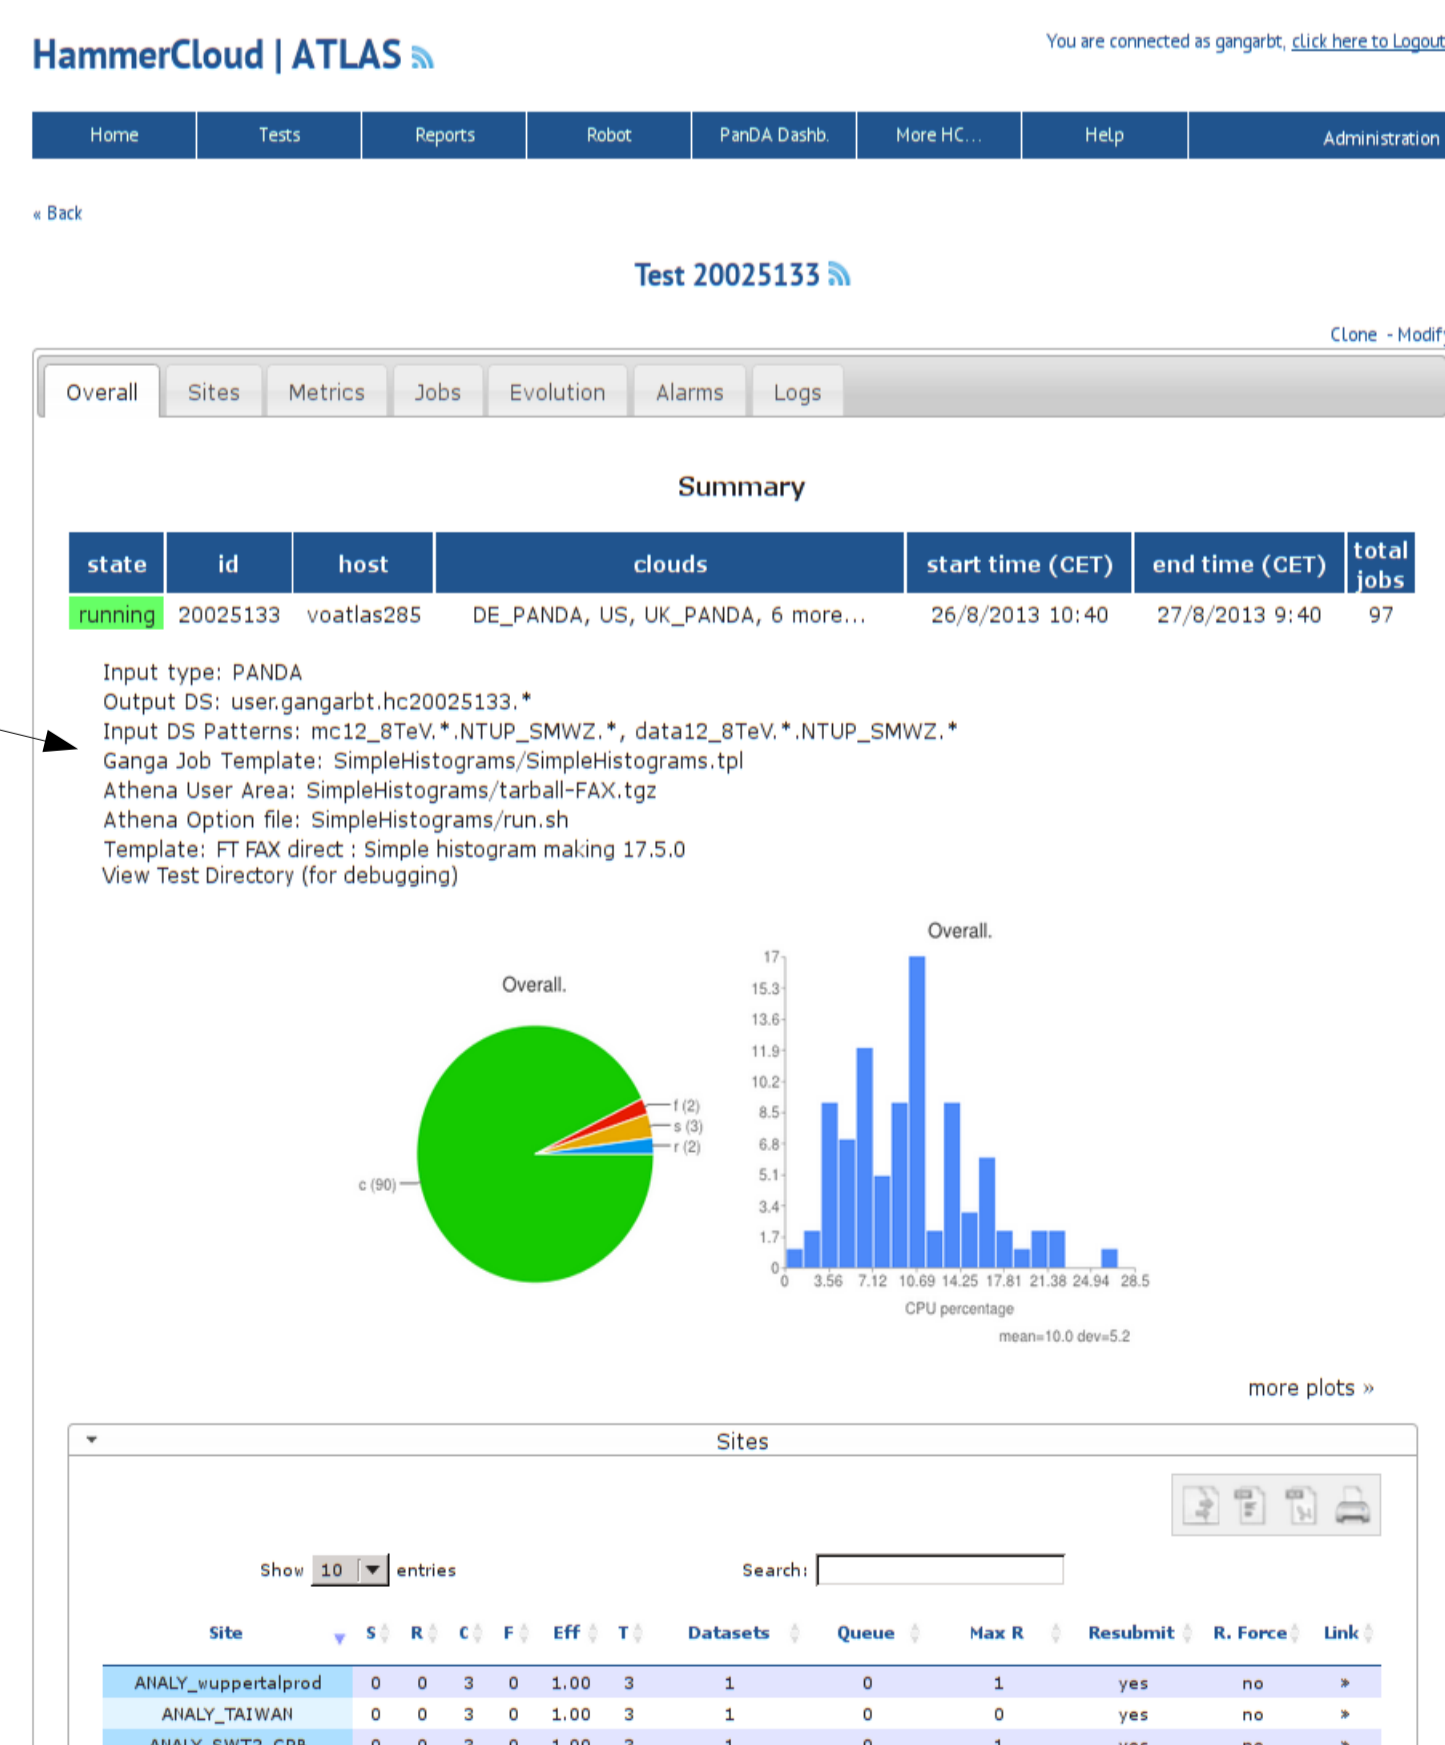
- Analysis queues (AFT)
- Production queues (PFT)
- both used for queue **black-listing**
- Panda Pilot development
- Multiple Processor Core
- Athena Nightly Build System
- Federated XRootD (FAX)
- Cloud

ATLAS stress tests:

- Standard analysis or production workflows used for site or new component commissioning with large number of jobs

Functional Test operations:

- Pre-defined in templates
- Tests executed in 24h rhythm
- Single jobs
- Results displayed on webpages
- Email notifications for shifters and experts



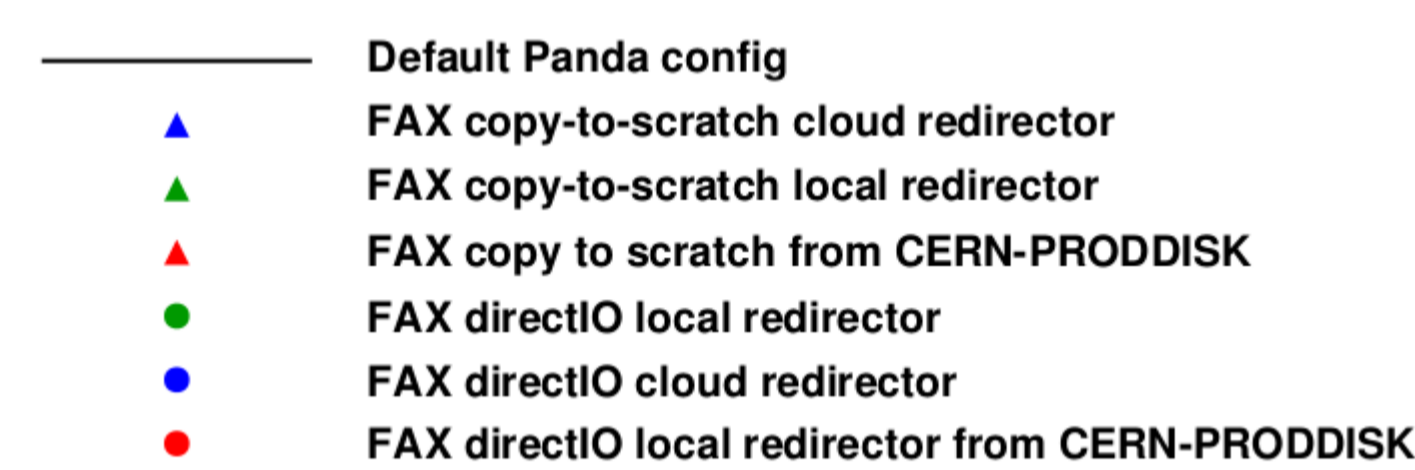
4. HammerCloud FAX results

FAX (Federating ATLAS storage systems using XRootD) brings Tier1, Tier2 and Tier3 storage resources together into a common namespace, accessible from anywhere, thus relaxing the traditional requirement of data-CPU locality. Client software tools like ROOT or xrdcp interact with FAX behind the scenes to reach data regardless of its location. Improved network bandwidth, reduced latency and data structure aware caching mechanism like TTreeCache make this possible.


HammerCloud is used to setup and test the FAX system:

- Define and setup Panda FAX configuration
- Functional tests of:
 - Local, cloud and global redirector
 - Direct I/O and copy mode
- Determine best work mode and heavy load in stress tests

Executed **Stress Test** with typical ATLAS user analysis at various sites for 48-72h with different **input access modes**:



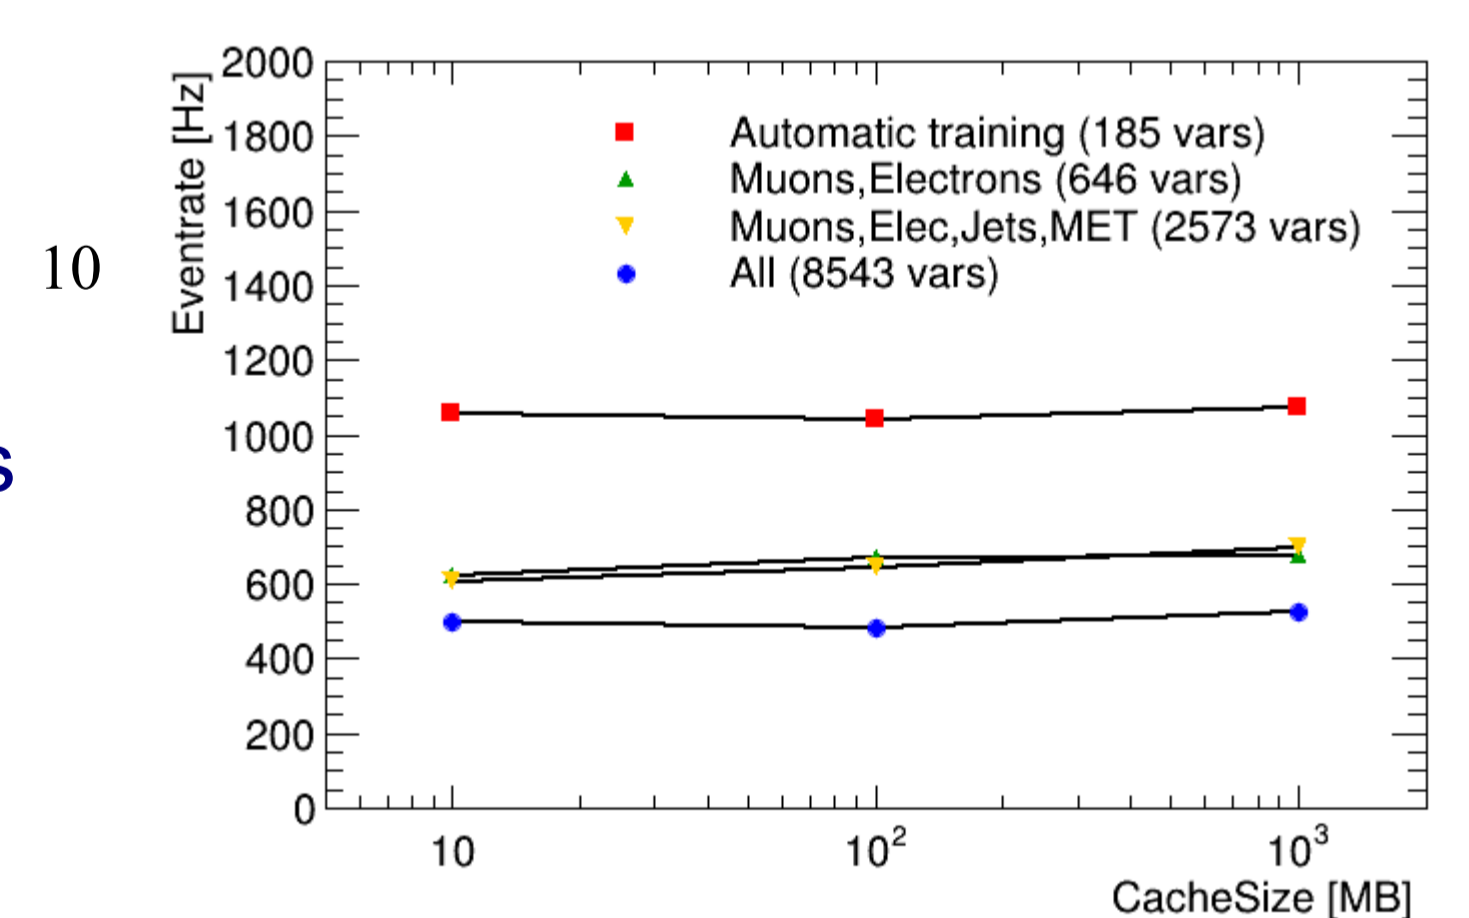
Result for example sites of different storage systems:



Default direct I/O at dCache, Storm, EOS fine – large improvements for DPM with FAX/XRootD direct I/O.

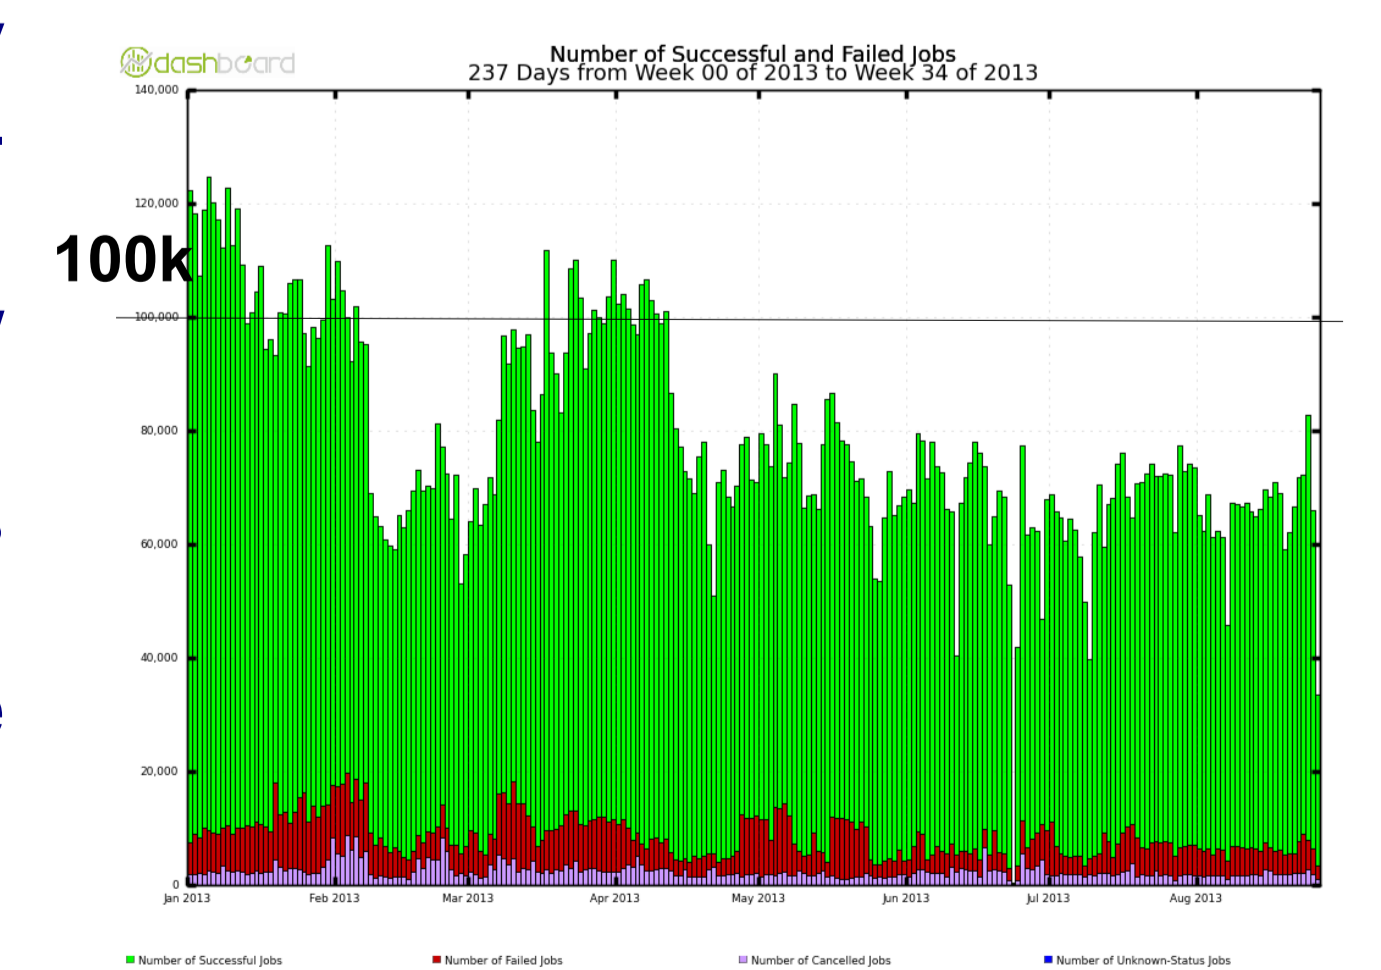
Usage of TTreeCache esp. for remote reading essential.

Best: automatic TTreeCache Variable training independent of Cache size used for this analysis



5. Current usage and plans

HammerCloud is very successfully used in many aspects of ATLAS day-to-day Grid computing: functional tests with site black-listing, new component testing like e.g. FAX and many other things. HammerCloud is executing on average 50k functional test jobs per day. Periods with more jobs per day are due to additional stress tests running.



ATLAS is moving to new Grid components before the 2015 LHC data taking: Panda will be upgraded to the job task system JEDI and the data management system DQ2 will be upgraded to RUCIO. These two core components have to be integrated into the HammerCloud workflows.