

Commissioning of a CERN Production and Analysis Facility Based on xrootd

Simone Campana, Daniel C. van der Ster, Alessandro Di Girolamo, Andreas J. Peters, Dirk Duellmann, Miguel Coelho Dos Santos, Jan Iven, Tim Bell

CERN, Geneva, Switzerland

E-mail: simone.campana@cern.ch

Abstract. The CERN facility hosts the Tier-0 of the four LHC experiments, but as part of WLCG it also offers a platform for production activities and user analysis. The CERN CASTOR storage technology has been extensively tested and utilized for LHC data recording and exporting to external sites according to experiments computing model. On the other hand, to accommodate Grid data processing activities and, more importantly, chaotic user analysis, it was realized that additional functionality was needed including a different throttling mechanism for file access. This paper will describe the xroot-based CERN production and analysis facility for the ATLAS experiment and in particular the experiment use case and data access scenario, the xrootd redirector setup on top of the CASTOR storage system, the commissioning of the system and real life experience for data processing and data analysis.

The ATLAS experiment [1] relies on the WLCG [2] infrastructure for the largest part of its offline activities. In fact, ATLAS benefits from resources in three different Grid flavors: EGI [3], OSG [4] and NorduGrid [5]. In order to guarantee transparent access to resources accessible via different protocols and services, ATLAS built a service layer on top of the grid services, accessible by the high level applications. ATLAS Distributed Data Management system (DDM) [6] guarantees data placement, organization and bookkeeping, while the Production and Data Analysis service (PanDA) [7] has been developed to cover the workload management use cases and currently manages the simulated data production the data reprocessing and most of the physics group and physics user analysis. In the analysis scenario, users submit workloads (analysis code and input parameters) using the native PanDA client (pathena) or the more generic job submission framework Ganga. In both cases, the user task definition is translated into PanDA jobs. Before the task can start, PanDA executes a prolog job, compiling the user code, storing the binary objects in a DDM dataset at the local storage of the site. The bulk of the workload will then be executed at the given site. All the data handling is local to the site (no cross site data transfers): the precompiled binary together with input datasets are pulled from the local storage, the files with the results of the analysis are stored in transient space in the local site as well. Probably the most challenging aspect of user analysis consists of the data access to the input datasets, since the data volume can be considerably large (up to 15GB for a 1 hour WallClockTime job) compared to the output files sizes (a few MB). Concurrent accesses from many jobs to the same disk servers often combine to overload the storage element and slow down the daemons serving the files. Historically, more than 60% of data analysis jobs failures have been due to data access. The ATLAS framework allows to dynamically choose






Figure 2. The CASTOR setup for ATLAS at CERN.

high level services like PanDA, but local (non-Grid) access to data would still remain an issue. Also the usage of SRM rather than the native RFIO protocol would partially solve the data access issues in terms of storage complexity, but would introduce an unacceptable overhead in terms of performance. On this last topic, even the more performing access via RFIO suffers from delays due to CASTOR internal scheduling, probably rather appropriate for organized activities like intra-sites bulk data movement, but similarly inefficient for the chaotic access use case.

For the reasons mentioned above, the CERN CASTOR setup has been extended with an extra layer, integrating the XROOT technology into the existing storage solution. The new X2CASTOR [10] setup (shown in Figure 2) consists of an XROOT redirector on top of the CASTOR namespace and XROOT daemons running on each disk server. In order to provide transparent data access, X2CASTOR maps directories in the namespaces with pools in CASTOR (and space tokens in SRM) so that specifying the filename as in the namespace is sufficient for the system to quickly locate a file in the pool (compare 100ms of lookup time in X2CASTOR with few seconds for via SRM). In addition, the XROOT protocol allows throttling of file access across various pools and bypasses the CASTOR batch scheduling. Finally, X2CASTOR allows authentication both via X505 and Kerberos, which partially resolves the dichotomy of Grid versus local authorization.

The X2CASTOR setup at CERN has been tested and finally commissioned with real ATLAS analysis workflows, using HammerCloud [11]. HammerCloud is a framework designed and implemented in collaboration between ATLAS and CERN IT to (stress) test sites and




Figure 3. Results of the final HammerCloud test of the X2CASTOR service.

particularly the storage elements with analysis jobs. Tests can be defined and scheduled by central ATLAS computing operations, site contacts and cloud contacts, so that a given number of jobs is delivered to a given site or cloud, running an analysis pattern defined by the submitters for a configurable number of hours or days. The job monitoring (successes versus failures) is also provided by the framework, together with the mining of the most important measurable metrics such as the processed events per second, CPU utilization, time to retrieve input and time to store the outputs. The framework has not just been used for site commissioning (including CERN analysis) but also to evaluate changes to site software and configurations, evaluate changes to experiment software and benchmark different sites to provide a realistic comparison.

The HammerCloud tests on the CERN ATLAS analysis infrastructure, based on X2CASTOR started in March 2010, running the MuonTriggerAnalysis with Athena 15.5.1. As input datasets we choose the merged AOD files (2GB average size) from the mc08 simulation campaign. Initially, the maximum number of running jobs in parallel was capped to 50 and the test lasted 40 hours. The copy-to-scratch file access approach was used; the FileStager method has also been tested but could not be fully exploited due to the current limitations of Athena in supporting X509 authentication via ROOT. The tests concluded at the end of April 2010 when the cap mentioned above was first raised to 500 jobs and finally removed (up to 20 3K jobs did run concurrently on the farm for few days). The final results of the commissioning activity can be seen in Figure 3: the overall job efficiency reached 99%, with more than 15Hz average event rate and an approximate 50% Wall/CPU Time Efficiency. Such performance ranks CERN among the best sites for ATLAS analysis.

Given the positive outcome of the commissioning phase, the CERN site was enabled for ATLAS analysis, production and reprocessing activities on the 19th of May 2010. In the first 24 hours of real activity, 8020 jobs finished and only 61 failed, confirming and exceeding the expectations from the HammerCloud test (and confirming the effectiveness of the testing framework itself in simulating load from real activity). Figure 4 shows the ramp-up of ATLAS




Figure 4. PanDA plots showing the number of jobs running at CERN (in pink) on the first day (left) and for the past year (right).




Figure 5. Example saturation of a X2CASTOR disk server due to hot files.

running analysis jobs at CERN together with the daily number of running jobs over an extended period of time; for the comparison, one has to consider that the CERN cloud consists of only one site (CERN itself), while all other clouds count many different sites. Because of the large data sample hosted by the site, together with the performance of the infrastructure, CERN has become one of the most utilized sites for user analysis. At the same time, CERN played a fundamental role in the ATLAS Autumn reprocessing campaign, running a large part of the reprocessing of the express stream and serving as fallback solution for sites temporarily unavailable. In total, CERN reprocessed around 10% of the data in the mentioned campaign, while it ran little if no reprocessing at all in campaigns prior to the X2CASTOR system deployment and commissioning.

Despite all the achievements just described, the X2CASTOR setup still presents limitations hardly addressable with the current core storage system. One of the major issues of the current system appears in trying to handle hot files, i.e. files access by many batch nodes concurrently (the typical case is the access to files containing information concerning the detector conditions in the given period of the run). The absence (or limited) replication of such files across multiple servers (with is rather difficult and inflexible in CASTOR and therefore X2CASTOR) in several cases created hotspots in the storage, exhausting the max I/O deliverable by a single disk server (approx 120MB/s for a standard CERN disk), resulting in massive file access failures. See the saturation effect for a hot disk server in Figure 5. To overcome this issue, the number of job slots for analysis at CERN has been throttled on many occasions down to less than 1000.

To provide a better solution to the mentioned issues and together provide a more maintainable storage solution in the long term, CERN IT has been working for many months on a new XROOT-based disk pool solution (EOS) [12], which in the medium term should replace X2CASTOR for the disk resident data, while data access to tapes would still remain within the CASTOR system. ATLAS and CERN IT are collaborating to evaluate this new technology, in the so-called Large Scale Test initiative, born after the WLCG Storage Workshop in Amsterdam




Figure 6. Early HammerCloud results of the new EOS service with 8 disk servers configured with JBOD disks.

in June 2010. Early tests have proven the functionality and shown good performance. In Figure 6 one can see the outcome of the test (efficiency, event rate, CPU/WallClock time efficiency) for the EOS system, with disk servers first organized according to the RAID0 standard and then configured as an 8 disk server JBOD where disks are independently addressed and the redundancy (and therefore data reliability and I/O performance) is achieved through the storage implementation itself, rather than the disk arrays driven by a controller. Those tests are preliminary and the outcome should be interpreted more qualitatively rather than quantitatively, at least at the time of this contribution in October 2010.

Acknowledgements

Partial support for this work was provided by EGI-InSPIRE INFSO-RI-261323.

References

- [1] The ATLAS Collaboration Website: atlas.cern.ch
- [2] The Worldwide LHC Compute Grid Website: lcg.web.cern.ch
- [3] The European Grid Infrastructure Website: www.egi.eu
- [4] The Open Science Grid Website: www.opensciencegrid.org
- [5] The NorduGrid Website: www.nordugrid.org
- [6] V. Garonne et al. Status, News and Update of the ATLAS Distributed Data Management Software Project: DQ2. In these proceedings.
- [7] T. Maeno et al. Overview of ATLAS PanDA Workload Management. In these proceedings.
- [8] The CASTOR Website: castor.web.cern.ch
- [9] A. Shim and A. Shosani, editors. The Storage Resource Manager Interface Specification. Version 2.2. Available online: <https://sdm.lbl.gov/srm-wg/doc/SRM.v2.2.html>
- [10] X2CASTOR service website: <https://twiki.cern.ch/twiki/bin/view/DataManagement/X2CASTOR>
- [11] D. van der Ster et al. HammerCloud: A Stress Testing System for Distributed Analysis. In these proceedings.
- [12] A. Peters et al. Exabyte Scale Storage at CERN. In these proceedings.