

Design, Results, Evolution and Status of the ATLAS Simulation at Point1 Project

S Ballestrero¹, S M Batraneanu², F Brasolin³, C Contescu^{2,4}, D Fazio²,
A Di Girolamo², C J Lee^{1,2}, M E Pozo Astigarraga², D A Scannicchio⁵,
A Sedov^{6,7}, M S Twomey⁸, F Wang⁹ and A Zaytsev¹⁰

¹ University of Johannesburg, South Africa

² CERN, Switzerland

³ Istituto Nazionale di Fisica Nucleare Sezione di Bologna, Italy

⁴ Polytechnic University of Bucharest, Romania

⁵ University of California, Irvine, USA

⁶ Port d'Informació Científica (PIC), Barcelona, Spain

⁷ Institut de Física d'Altes Energies (IFAE), Barcelona, Spain

⁸ University of Washington Department of Physics, USA

⁹ University of Wisconsin-Madison, USA

¹⁰ Brookhaven National Laboratory (BNL), USA

E-mail: atlas-tdaq-sysadmins@cern.ch, alezayt@bnl.gov

Abstract. During the LHC Long Shutdown 1 (LS1) period, that started in 2013, the Simulation at Point1 (Sim@P1) project takes advantage, in an opportunistic way, of the TDAQ (Trigger and Data Acquisition) HLT (High-Level Trigger) farm of the ATLAS experiment. This farm provides more than 1300 compute nodes, which are particularly suited for running event generation and Monte Carlo production jobs that are mostly CPU and not I/O bound. It is capable of running up to 2700 Virtual Machines (VMs) each with 8 CPU cores, for a total of up to 22000 parallel jobs. This contribution gives a review of the design, the results, and the evolution of the Sim@P1 project, operating a large scale OpenStack based virtualized platform deployed on top of the ATLAS TDAQ HLT farm computing resources. During LS1, Sim@P1 was one of the most productive ATLAS sites: it delivered more than 33 million CPU-hours and it generated more than 1.1 billion Monte Carlo events. The design aspects are presented: the virtualization platform exploited by Sim@P1 avoids interferences with TDAQ operations and it guarantees the security and the usability of the ATLAS private network. The cloud mechanism allows the separation of the needed support on both infrastructural (hardware, virtualization layer) and logical (Grid site support) levels. This paper focuses on the operational aspects of such a large system during the upcoming LHC Run 2 period: simple, reliable, and efficient tools are needed to quickly switch from Sim@P1 to TDAQ mode and back, to exploit the resources when they are not used for the data acquisition, even for short periods. The evolution of the central OpenStack infrastructure is described, as it was upgraded from Folsom to the Icehouse release, including the scalability issues addressed.

Introduction and system overview

During the LHC Long Shutdown 1 (LS1), from Feb 2013 to Mar 2015 at CERN, the Simulation at Point1 project (Sim@P1) [1] was a great opportunity to exploit the ATLAS experiment's [2] Trigger and Data Acquisition (TDAQ) High-Level Trigger (HLT) [3] farm. More than 1300 Compute Nodes (CNs) are installed in the SDX1 area of the ATLAS experiment site, LHC Point 1 (P1) [4]. They are particularly appropriate for running Monte Carlo production jobs, which are mostly CPU bound and require limited disk and network I/O.

By design the Sim@P1 project was developed with two main constraints: the security and reliability of the whole ATLAS Detector Control System must be preserved at all times; the ATLAS production jobs need to run unmodified and without additional software requirements for the HLT nodes. For these reasons the Sim@P1 system is completely isolated in a virtualized platform.

Each physical node of the TDAQ HLT farm is connected to the TDAQ Control and Data networks as well as to a dedicated Sim@P1 network: a tagged Virtual LAN (VLAN) sharing the same top-of-the-rack switch as the TDAQ data network, with a dedicated uplink to the router which connects the farm to the Castor/EOS storage infrastructure of CERN IT, via the General Purpose Network (GPN). At a routing level, the communication between the data network and the Sim@P1 network is not permitted. Only a restricted set of servers in the CERN GPN are allowed to communicate with the Sim@P1 virtual machines (VMs), under the control of an Access Control List (ACL) implemented on the Castor router (Fig. 1). The Castor/EOS router also blocks direct connectivity to outside CERN, and if needed can be configured to limit the bandwidth to avoid interferences with ATLAS data storage requirements.

The physical nodes run Scientific Linux CERN 6 (SLC6) [5] and use the KVM hypervisor [6] to run the VMs, which are only allowed to access the dedicated VLAN interface. The use of VMs allows a fast switch between the execution of TDAQ HLT and of Sim@P1, without requiring a reboot of the physical node.

Considering the scale of the farm (over 22k HyperThread in the current configuration), a decision was made to use a cloud management layer. OpenStack [7] was chosen for this purpose as it is widely adopted, including at CERN [8].

During the first six months of production operations (2013Q2-2013Q4) the CernVM project [9] was used as a primary source of the Scientific Linux CERN (SLC5) [10] base images for the Sim@P1 project virtual machines. Starting from the beginning of 2014 the SLC6 base VM images are provided by the ATLAS Cloud Operations group [11] so as improve the management of image contents.

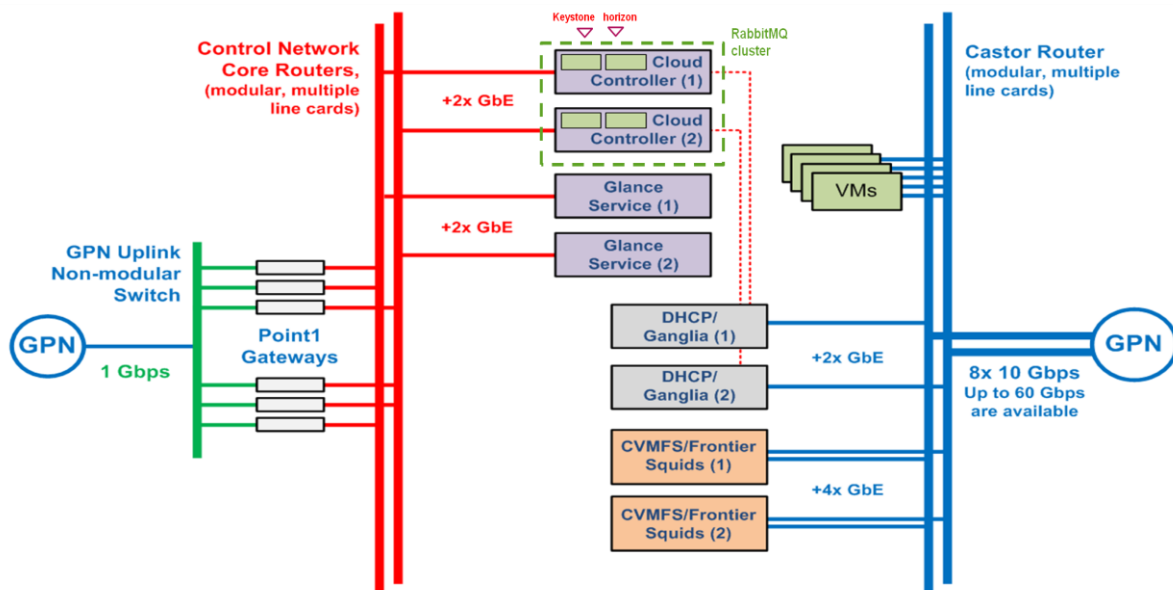


Figure 1. OpenStack services and network layout: storage and external facilities for Sim@P1 are available through the P1 Castor router.

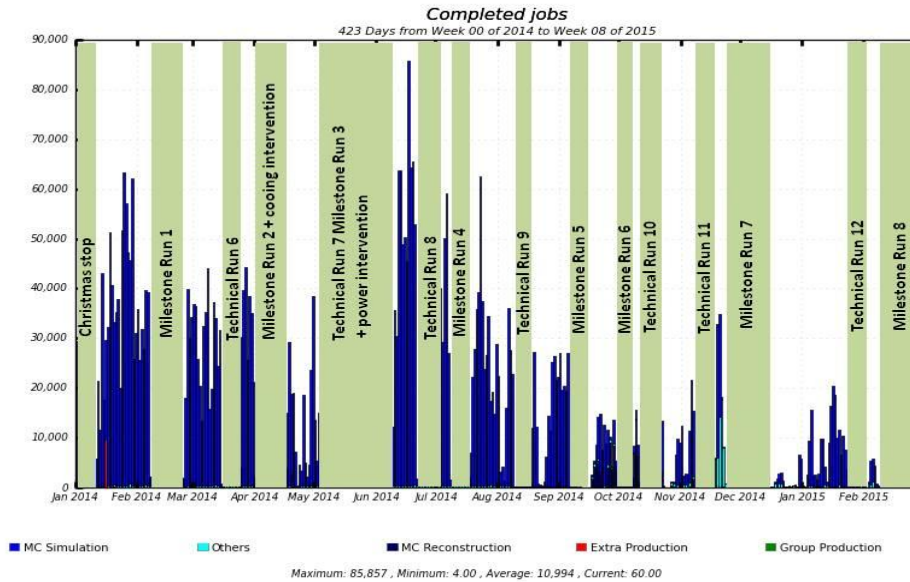


Figure 2. ATLAS HLT TDAQ farm availability for the Sim@P1 Project during the period of Jan 1st, 2014 – Feb 28th, 2015.

Availability, operation and productivity during the LS1 period

After the test and setup phase that started in 2013Q1, the Sim@P1 project entered in production in 2013Q3 and quickly became one of the most productive ATLAS sites.

By design, Sim@P1 does not run user analysis jobs, in order to minimize the risks of arbitrary workload injection in the P1 networking environment. Because of the limited disk space available on the physical hosts, and of the limited bandwidth available to each rack, it was decided to run exclusively Monte Carlo simulation and event generation jobs. The project is currently represented by the following three resources in the ATLAS workload management, Panda [12]: CERN-P1 (single core jobs), CERN-P1_MCORE (AthenaMP multicore jobs [13]) and CERN_P1_MCORE_HI (multicore jobs with high memory requirements).

The availability of the HLT farm for the Sim@P1 project was also reduced due to various other activities needed by ATLAS in preparation of the LHC Run 2, such as multiple ATLAS wide re-commissioning, TDAQ Technical Runs and P1 infrastructure cooling and power maintenance and upgrade interventions. Accounting for these periods of unavailability, in 2014 the HLT farm was used for Sim@P1 about 40% of the time, as shown in Fig. 2.

Despite of these constraints, the Sim@P1 project delivered more than 33 million CPU-hours and generated more than 1.1 billion Monte Carlo events just during the period January 1st, 2014 – February 28th, 2015 (Fig.3), and more than 50 million CPU-hours and 1.7 billion Monte Carlo events in total since the beginning of production in 2013Q2.




Figure 3. Comparison of the CPU time consumption by the successful jobs for the ATLAS Tier-1 sites and the Sim@P1 during the LS1 period (Jan 1st, 2014–Feb 28th, 2015). The plot takes into account only the SIM and EVGEN jobs types.

Improvements of the OpenStack-based Virtualized infrastructure

The choice of OpenStack leverages know-how at CERN and BNL. Only the following components are needed to support Sim@P1: Nova, Keystone, Glance and Nova-Network.

The initial implementation, at the beginning of LS1, was based on the “Folsom” 2012.2.4 version of OpenStack. In preparation for the LHC Run 2, the system was upgraded in 2014Q3 to “Icehouse” 2014.1.3.

With this release, the security has been improved by additional RPC layers on the Nova database, provided by the Nova Conductor. Also, with the improved Nova component, fewer failures have been encountered when initiating the VMs.

New wrapper scripts have been developed for handling bulk operations, such as instantiations and deletions of large groups of VMs.

Finally, to quickly move the HLT resources from TDAQ to Sim@P1 and vice-versa, automation scripts have been prepared to handle the failures during VM initiation or deletion, with automatic retrials. This allows Sim@P1 to saturate the available resources or to free some computing power, without the need of any manual intervention.

Dynamic partitioning mechanism

The Sim@P1 project uses HTCondor [14] batch system to manage the distribution of jobs to the VMs. One of the advantages of HTCondor is the easy implementation of dynamic partitioning, that is the way of allocating resources with respect to the different requirements of specific jobs.

The automatic, dynamic partitioning leads to a more effective use of available computing resources when the expected ratio of different job types is not known in advance or varying with time. In our case it means that the same VM, a Worker Node (WN) in HTCondor parlance, can be used to run single-core and multi-core jobs at different periods in time, as the priorities for ATLAS Offline simulation production change.

At any particular time each WN can either run one multi-core job or up to eight single-core jobs. Due to this limitation the whole WN needs to be drained of running jobs, before changing their type. In HTCondor this type of management is done through the group fair share mechanism and the use of the defrag tool. The first allows to set up the expected fractions of the available resources that can be used by each type of job, while the second checks the current requirements and allows draining and migration of WNs from one group to another.



Figure 4. Measurement of VM startup and graceful shutdown times for Sim@P1.

Availability and operation during Run 2

During the LHC Run 2 period, which started in April 2015, the availability of the ATLAS TDAQ HLT farm to the Sim@P1 project will be lower compared to its availability in the LS1 period. Not only the integral running time will be smaller, but, more importantly, the typical availability intervals are expected to be reduced from few weeks down to few days, and, in most cases, with no or little advance warning.

In 2014Q3-2015Q1, many improvements were made to the whole system of operations for Sim@P1, in order to maximize the use of the farm resources on short notice and for shorter time periods. In particular, a fast and fully automated procedure was established to quickly switch the resources from TDAQ to Sim@P1 mode and back. All the CNs reconfiguration steps are handled via Puppet [15], the base configuration management system of the ATLAS TDAQ computing environment.

The running mode of TDAQ HLT farm resources will change when an LHC stop longer than 24 hours is foreseen [16] and the resources are not booked for other needs of the ATLAS detector.

The shifter in the ATLAS Control Room needs to be able to perform this operation in a straightforward and easy way, at any time of day, without the intervention of the Sim@P1 experts: they simply “push the button” on a web interface and quickly the TDAQ HLT farm resources are switched from TDAQ to Sim@P1 mode or vice-versa. In this way the resources will be exploited in the best way: by TDAQ during data taking, and by Sim@P1 during the LHC stop periods. The typical switch operation will be using the pre-deployed VM images; the images are only expected to change occasionally, so the deploy overhead is not taken into account in the following. The typical switching times between the states of the TDAQ HLT farm are (Fig. 4):

- from Sim@P1 to TDAQ mode: ~12 minutes. This is the time needed by the Nova OpenStack controllers for the graceful shutdown of all VMs, and to update the Nova database.
- from TDAQ to Sim@P1 mode: ~1 hour. This is the time needed by the Nova controllers to instantiate all VMs, check the Nova database, restart bad VMs, deal with problematic CNs. This operation is indeed faster than the flow of incoming jobs.
- in addition, in case the resources are suddenly needed for ATLAS data taking, it is possible to quickly switch the whole farm in about 100 seconds. In this case the VMs are abruptly

terminated and the Nova database is not updated and will require a manual clean up, which takes approximately two hours.

Summary

During the LS1 period the Sim@P1 project was one of the most productive ATLAS sites: it delivered 33 million CPU-hours and generated more than 1.1 billion Monte Carlo events. This utilization allowed the farm to be a productive asset even beyond its design goals.

The success of the Sim@P1 project is possible due to the continuous, combined effort of all the subjects involved: the ATLAS TDAQ SysAdmin and NetAdmins teams, CERN IT SDC, the ATLAS Cloud Computing Operation team, the University of Wisconsin-Madison and the RHIC & ATLAS Computing Facility (RACF) at BNL.

The adoption of a virtualization layer managed by OpenStack allowed the separation of system administration and maintenance duties between the groups in charge of the TDAQ HLT farm infrastructure, the VM group management and the Sim@P1 WLCG [17] site operations. It also allows for the fast switching of operation mode for the resources, and provides the needed isolation for the cyber security of the ATLAS experiment.

The upgrade of OpenStack and the use of HTCondor dynamic partitioning, now provides a more robust platform and a better allocation of the available resources respect to the requirements of the different jobs.

During the LHC Run 2 the TDAQ HLT farm will mostly used for the data taking and the time slots available for Sim@P1 will be reduced: the scenario is changing and this represents a new challenge. We must take advantage of the resources available even for short period, and on short notice. For this reason it is crucial that all the steps of the switching procedure are fast, automated and reliable.

New tools have been implemented to let the ATLAS control room shifters switch the run mode of the TDAQ HLT farm. From TDAQ to Sim@P1 mode during the periods of LHC stop, even on short notice and for short period; and from Sim@P1 to TDAQ mode when LHC is in operation and the resources are needed for the data acquisition. Thus the Sim@P1 project is expected to continue its mission in the foreseeable future.

References

- [1] ATLAS TDAQ SysAdmins, Design and performance of the virtualization platform for the offline computing on the ATLAS TDAQ Farm (2014) Ser. 513 032011 (Journal of Physics: Conference Series)
- [2] ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, JINST 3 (2008) S08003
- [3] ATLAS Collaboration, ATLAS High Level Trigger, Data Acquisition and Controls: Technical Design Report, CERN/LHCC/2003-022 (2003), ISBN 92-9083-205-3
- [4] C. Lee *et al.*, ATLAS TDAQ System Administration: an overview and evolution. Submitted to Proceedings of Science The International Symposium on Grids and Clouds (ISGC) 2013, March 17-22, 2013 Academia Sinica, Taipei, Taiwan
- [5] <http://linux.web.cern.ch/linux/scientific6/>
- [6] http://www.linux-kvm.org/page/Main_Page
- [7] OpenStack Open Source Cloud Computing Software: <http://www.openstack.org>
- [8] P. Andrade *et al.*, Review of CERN Data Centre Infrastructure, J. Phys.: Conf. Ser. 396 (2012) 042002
- [9] CernVM Project: <http://cernvm.cern.ch>
- [10] <http://linux.web.cern.ch/linux/scientific5/>
- [11] <https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/AtlasCloudOperations>
- [12] <https://twiki.cern.ch/twiki/bin/view/PanDA/PanDA>
- [13] <https://inspirehep.net/record/1194221>
- [14] <https://twiki.cern.ch/twiki/bin/view/Atlas/TriggerOperationsSimAtP1>
- [15] Puppet Configuration Management System: <http://puppetlabs.com>
- [16] <http://research.cs.wisc.edu/htcondor/>
- [17] Worldwide LHC Computing GRID: <http://wlcg.web.cern.ch>