ATLAS Distributed Computing Shift Operation in the first 2 full years of LHC data taking

Jaroslava Schovancová[1], Alessandro Di Girolamo[2], Johannes Elmsheuser[3], Stephane Jézéquel[4], Guidone Negri[2], Nurcan Ozturk[5], Hiroshi Sakamoto[6], Mark Slater[7], Yuri Smirnov[8], I. Ueda[6], Daniel Colin Van Der Ster[2], Jaehoon Yu[5]
for the ATLAS Collaboration

[1] Institute of Physics, Academy of Sciences of the Czech Republic, Na Slovance 2, CZ-18221 Prague 8, Czech Republic
[2] CERN, CH - 1211 Geneva 23, Switzerland
[3] Ludwig-Maximilians-Universitat München, Fakultät für Physik, Am Coulombwall 1, DE - 85748, Garching, Germany
[4] LAPP, Universite de Savoie, CNRS/IN2P3, Annecy-le-Vieux, France
[5] The University of Texas at Arlington, Department of Physics, Box 19059, Arlington, TX 76019, United States of America
[6] The University of Tokyo, International Center for Elementary Particle Physics and Department of Physics, 7-3-1 Hongo, Bunkyo-ku, JP - Tokyo 113-0033, Japan
[7] University of Birmingham, School of Physics and Astronomy, Edgbaston, Birmingham B15 2TT, United Kingdom
[8] Brookhaven National Laboratory, Physics Department, Bldg. 510A, Upton, NY 11973, United States of America

E-mail: schovan@fzu.cz

ATL-SOFT-PROC-2012-015
20/05/2012

Abstract. ATLAS Distributed Computing organized 3 teams to support data processing at Tier-0 facility at CERN, data reprocessing, data management operations, Monte Carlo simulation production, and physics analysis at the ATLAS computing centers located worldwide. In this paper we describe how these teams ensure that the ATLAS experiment data is delivered to the ATLAS physicists in a timely manner in the glamorous era of the LHC data taking. We describe experience with ways how to improve degraded service performance, we detail on the Distributed Analysis support over the exciting period of the computing model evolution.

1. Introduction

The ATLAS Experiment [1] extensively uses grid resources provided by over a hundred WLCG sites all over the world. The ATLAS Distributed Computing is a very complex system [2]. For the workload management backend the ATLAS Experiment uses PanDA [3], as for the Distributed Data Management the DQ2 system [4] is used. Each activity within the ATLAS Distributed Computing is monitored through a system of monitoring tools. For the critical activities there are email notifications to the particular activity experts triggered when that particular activity gets bellow a certain threshold.

ATLAS runs continuously over 100k simultaneous jobs, over 20k job slots occupied by Analysis jobs, over 80k jobslots occupied by the ATLAS central production or the production carried by the ATLAS Physics groups, with average of 200k-300k jobs per day all over the world using the ATLAS Grid resources.

The DQ2 system is capable to accomodate throughput over 10 GBps, including all the activities: the data export from the Tier-0 to Tier-1s and calibration Tier-2s, Data Brokering, Data Consolidation, Production transfers, Subscriptions of Physics Groups or User data, and also Functional Tests at a marginal level of the overall throughput.

In order to provide reliable and efficient quality of service to the ATLAS Collaboration, the system has to be monitored and its functionality regularly tested. There are three areas covered by three different teams of shifters, who work from different places with different distance from an activity expert: the Distributed Analysis Support Team, the Comp@P1 Shift Team, and the ATLAS Distributed Computing Operations Shift Team.

In this paper we describe structure of the Shift Teams, their mutual communication, issue reporting and lessons learned in the past 2 years of the data taking era with the ATLAS Experiment.

2. Organization

The ATLAS Distributed Computing Operations Shift (ADCoS) Team was organized in 2008 to support 24/7 simulation production of the ATLAS experiment, data reprocessing and data management operations at more than a hundred of computing centers located world-wide. In order to support the ATLAS Physics Analysis the ATLAS Distributed Analysis Support Team (DAST) was formed in 2008. In addition, a team to support data processing at Tier-0 and data export from Tier-0 to Tier-1s, the Comp@P1, was formed in 2009.

The User analysis activity is monitored by the DAST. The DAST provides first line support to grid-related questions of the ATLAS physicists.

The data processing at Tier-0 and data export from Tier-0 to Tier-1s and calibration Tier-2s are monitored by the Comp@P1 Shift Team located at CERN in the ATLAS Control Room.

The Grid production activity, including the data transfers among Tier-1s, Tier-2s, and Tier-3 sites, data processing, data reprocessing, Physics Group production, and Monte Carlo Simulations, and Functional Testing, is monitored by the ADCoS Team.

The ATLAS Distributed Computing Operations Shift (ADCoS) and Distributed Analysis Support Team (DAST) are world-wide distributed teams organized in three different time zones, natural to the location of the team members: the Asia-Pacific time zone (00-08 hrs CE(S)T), the Europe time zone (08-16 hrs CE(S)T), and the Americas time zone (16-24 hrs CE(S)T). The Comp@P1 shift, which is located at CERN, follows similar time zone scheme.

The WLCG sites supporting the Virtual Organization atlas are organised in 11 geographical areas called clouds. Each cloud organises a team of experts, the Cloud Squad, who can provide first line support for the site administrators in ATLAS matters.

In addition to the ATLAS Distributed Computing shift teams and the Cloud Squads, there is one ATLAS Distributed Computing Manager on Duty, who can address or escalate issues with the core services of the ATLAS Distributed Computing.

3. Training

The Shifter training consists of three phases: the shift tutorial, self-study of shift instructions, and the supervized "Trainee" shift.

The shift tutorials are organised with different frequency reflecting the current needs of the corresponding shift team. The tutorial frequency varies from once per year to once per month for the ATLAS Distributed Computing Shifts.

Each ATLAS Distributed Computing Shift Team has its own TWiki page to document Shift workflow, to describe procedures, and to propose solutions or a list of steps to be followed in order to escalate the issue to activity expert attention, for every category of issues. The instructions point to every monitoring [5] or command-line tool, which can provide useful information in order to debug an issue. The instructions are available to every member of the ATLAS Experiment. The Shift instructions are regularly updated. The self-study of shift instructions gives the Trainee shifter a first hands-on glance of the Shift activity and duties.

During the supervised Trainee shift the Trainee Shifter can perfect her/his skills in issue troubleshooting, so that she/he is fairly familiar with the Shift duties for his solo shift. At any point of the training or the solo shift the Shifters are encouraged to ask questions to the Shift experts.

4. Communication

The ATLAS Distributed Computing Shifters use different technologies to communicate with others on duty. Each of the three Shift Teams use mailing lists for intra-/inter- shift communication. In addition to e-mail communication the Shift teams can instantly communicate in the ATLAS Distributed Computing Virtual Control Room based on the XMMP protocol. Every message to a mailing list or to the Virtual Control Room is archived, therefore can be used as a knowledge bank for the Shift Team.

The ATLAS Distributed Computing Shift Teams use Global Grid User Support (GGUS) [6] tool to escalate issues with services provided by the sites to the site administrators. In addition, when a site is failing to perform well in an activity, site is excluded from that activity and the exclusion and subsequent actions are recorded in a Savannah ticket. Any time an ATLAS Distributed Computing Shifter takes an action, such an action is then recorded in an electronic logbook. Every ATLAS Distributed Computing Shifter, expert, or Cloud Squad member can then access a full record of manual actions which affect his/her area of interest.

In the first two years of data taking there were created approx. 200 GGUS tickets per month to address issues at ATLAS sites, over 97 % of the GGUS tickets were created by the ATLAS Distributed Computing Shifters. The ATLAS Distributed Computing Operations are trying to increase level of automation of the repetitive tasks, in order to decrease the need for repetitive manual operations taken by site administrators, experts, or Shifters.

5. Automation

In order to provide resilient analysis infrastructure to the ATLAS physicists, various tests to regularly examine functionality of services have been put in place. Currently, there are tests probing quality of network links among sites (Sonar, perfSonar [7]), measuring availability of services (SAM/Nagios tests [8], [9]), and simulating workload (HammerCloud Production and Analysis Functional Tests [10]).

6. Conclusions

The ATLAS Experiment is collecting data for more than 2 years already. The ATLAS Distributed Computing successfully fulfills its mission to deliver data to the ATLAS physicists. The ATLAS Distributed Computing Shift Teams, the Distributed Analysis Support Team, the ATLAS Distributed Computing Operations Shift Team, and the Comp@P1 Shift Team, together with the Experts, Cloud Squads members, and the site administrators, significantly contribute to this effort, making sure that the physicists benefit from Computing efficiently and in timely manner.

7. Acknowledgements

Jaroslava Schovancová greatfully appreciates support from the Academy of Sciences of the Czech Republic. Support from the grant LA08032 of the MEYS (MŠMT), Czech Republic, and from ACEOLE: A Marie Curie ITN Project is greatly acknowledged. The work was supported by the grant SVV-2012-265309 of the Charles University in Prague.

References

[1] ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider," JINST **3** (2008) S08003.

[2] S. Jézéquel *et al.*, "ATLAS Distributed Computing Operations: Experience and improvements after 2 full years of data-taking," to appear in Proceedings of the Computing in High Energy and Nuclear Physics 2012 International Conference

[3] T. Maeno, "PanDA: Distributed production and distributed analysis system for ATLAS," J. Phys. Conf. Ser. **119** (2008) 062036.

[4] M. Branco *et al.*, "Managing ATLAS data on a petabyte-scale with DQ2," J. Phys. Conf. Ser. **119** (2008) 062017.

[5] J. Schovancová for the ATLAS Collaboration, "ATLAS Distributed Computing Monitoring tools after full 2 years of LHC data taking," to appear in Proceedings of the Computing in High Energy and Nuclear Physics 2012 International Conference

[6] T. Antoni *et al.*, "Global grid user supportbuilding a worldwide distributed user support infrastructure," J. Phys. Conf. Ser. **119** (2008) 052002.

[7] P. Laurens *et al.*, "Monitoring the US ATLAS Network Infrastructure with perfSONAR-PS," to appear in Proceedings of the Computing in High Energy and Nuclear Physics 2012 International Conference

[8] A. Di Girolamo *et al.*, "New solutions for large scale functional tests in the WLCG infrastructure with SAM/Nagios: the experiments experience," to appear in Proceedings of the Computing in High Energy and Nuclear Physics 2012 International Conference

[9] P. M. Rodrigues De Sousa Andrade *et al.*, "Service Availability Monitoring framework based on commodity software," to appear in Proceedings of the Computing in High Energy and Nuclear Physics 2012 International Conference

[10] D. C. van der Ster *et al.*, "HammerCloud: A Stress Testing System for Distributed Analysis," J. Phys. Conf. Ser. **331** (2011) 072036.