

# ATLAS Distributed Computing Monitoring tools during the LHC Run I

J Schovancová<sup>1</sup>, S Campana<sup>2</sup>, A Di Girolamo<sup>2</sup>, S Jézéquel<sup>3</sup>, I Ueda<sup>4</sup>,  
T Wenaus<sup>1</sup> for the ATLAS Collaboration

<sup>1</sup> Brookhaven National Laboratory, Physics Department, Bldg. 510A, Upton, NY 11973, United States of America

<sup>2</sup> CERN, CH - 1211 Geneva 23, Switzerland

<sup>3</sup> LAPP, Université de Savoie, CNRS/IN2P3, Annecy-le-Vieux, France

<sup>4</sup> 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

E-mail: jschovan@bnl.gov

**Abstract.** This contribution summarizes evolution of the ATLAS Distributed Computing (ADC) Monitoring project during the LHC Run I. The ADC Monitoring targets at the three groups of customers: ADC Operations team to early identify malfunctions and escalate issues to an activity or a service expert, ATLAS national contacts and sites for the real-time monitoring and long-term measurement of the performance of the provided computing resources, and the ATLAS Management for long-term trends and accounting information about the ATLAS Distributed Computing resources. During the LHC Run I a significant development effort has been invested in standardization of the monitoring and accounting applications in order to provide extensive monitoring and accounting suite. ADC Monitoring applications separate the data layer and the visualization layer. The data layer exposes data in a predefined format. The visualization layer is designed bearing in mind visual identity of the provided graphical elements, and re-usability of the visualization bits across the different tools. A rich family of various filtering and searching options enhancing available user interfaces comes naturally with the data and visualization layer separation. With a variety of reliable monitoring data accessible through standardized interfaces, the possibility of automating actions under well defined conditions correlating multiple data sources has become feasible. In this contribution we discuss also about the automated exclusion of degraded resources and their automated recovery in various activities.

## 1. Introduction

The Large Hadron Collider (LHC) at CERN (European Organization for Nuclear Research) has been delivering stable beams colliding at the center-of-mass energy of 7 TeV since the end of March 2010, and at the center-of-mass energy of 8 TeV since mid April 2012. The ATLAS Experiment [1], one of the general purpose detectors at the LHC, has been efficiently taking data, resulting in 8 PB of RAW data accumulated over the LHC Run I. ATLAS uses the Worldwide LHC Computing Grid (WLCG) and several compute cloud technologies to process the data and simulations. The ATLAS Distributed Computing (ADC) team [2] is responsible for optimizing usage of the ATLAS grid resources.



The ADC infrastructure is a very complex system: The ATLAS grid and cloud resources (computing, storage and network resources) are spread over more than 130 computing sites distributed worldwide. ATLAS grid computing sites host over 140 PB of storage either on disk or tape systems with different storage system flavors, in addition to heterogeneous CPU resources ready to accommodate over 100k job slots. ATLAS grid sites are organized within three different flavors of grid: EGI (European Grid Infrastructure), OSG (Open Science Grid), and NeIC (Nordic e-Infrastructure Collaboration). To provide a high quality of service to the ATLAS Collaboration, it is essential for the operations of the ATLAS computing resources to easily identify issues with the infrastructure, and to address these issues. Such a challenging task is addressed by the ADC Monitoring team.

## 2. ATLAS Distributed Computing Monitoring Suite

The ADC Monitoring Suite consists of real-time monitoring tools, and of long-term accounting tools. The real-time monitoring tools [3] expose status of activities in a broad range of areas, such as CERN Analysis Facility, distributed data management, data processing system, status of database system essential for ATLAS Computing, and status of sites and services. The real-time monitoring tools address needs of the day-to-day work of the ADC Operations team – the activity experts, the distributed computing shifters, and the site administrators. The accounting tools [4] aggregate summaries of the real-time information and store it for long-term series analysis and visualization. The accounting tools provide visualization of resource utilization for the ATLAS Management and ATLAS National contacts. Architecture of the ADC Monitoring tools has been standardized to multi-layer design with clear separation of data access and visualization. Such approach favors modular design of software products, and allows for long-term maintainability.

Data access layer exposes data from possible various sources in a predefined structure in different formats, e.g. JSON (JavaScript Object Notation), XML (Extensible Markup Language), CSV (comma-separated values), utilizing variety of open-source or in-house made frameworks, e.g. Dashboard framework [5], PanDA (Production and Distributed Analysis System) web platform [6]. Monitoring data is stored primarily in Oracle database at CERN, however, in addition NoSQL data storage solutions with Apache Hadoop, and Apache Cassandra are used.

Visualization layer exposes monitoring data in form of graphics, or tables using various JavaScript libraries, such as jQuery and its plug-ins, **xbrowse** [7] or **hBrowse** [8] JavaScript libraries. The multi-layer design not only enables unified visual identity across all ADC Monitoring, but opens a variety of ways to share content.

## 3. Automation

Data available to ADC Monitoring tools help to improve automation [9] of repetitive tasks, and direct the expert manpower to address more urgent issues.

The ATLAS Grid Information System (AGIS) [10] contains detailed description of ATLAS computing resources, together with the information about resource downtimes. This kind of information enabled creation of automatic probes which not only take storage and computing elements out of production for the duration of downtime, but had drained them in advance of a downtime as well, resulting in less manual interventions from the shift team and site administrators.

Aggregation of resource status information from AGIS, automatic probes acting on resource downtime, and results of functional tests carried out by HammerCloud [11], gives rise to a centralized exclusion tool.

Evaluation of SAM (Site Availability Monitor) [12], [13] functional tests of storage elements leads to the tests enhancements, resulting in high-granularity tests. Automatic actions [14] are taken based on the high-granularity tests.

Monitoring information about network link status became part of automatic decision to determine path for data transfers in the ATLAS distributed data management system [15], [16], and for job brokerage in PanDA [17].

#### 4. Conclusions

The ATLAS Experiment has been successfully collecting data for more than 3 years during the LHC Run I. The ATLAS data is processed and analyzed at more than 130 grid and cloud sites distributed worldwide. Such a complex system is continuously monitored. The ADC has been fulfilling successfully its mission to deliver data to the ATLAS physicists. The ADC Monitoring tools have been allowing comprehensive monitoring of the distributed infrastructure, to promptly identify and address issues, and to improve automation of repetitive tasks. The future challenges to face come not only with technology upgrades of the distributed infrastructure, but with preserving and increasing the usability level of the tools themselves.

#### Acknowledgments

This manuscript has been authored by employees of Brookhaven Science Associates, LLC under Contract No. DE-AC02-98CH10886 with the U.S. Department of Energy. The publisher by accepting the manuscript for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

#### References

- [1] ATLAS Collaboration 2008 *JINST* **3** S08003
- [2] Jézéquel S *et al.* for the ATLAS Collaboration 2012 *J. Phys.: Conf. Series* **396** 032058
- [3] Schovancová J *et al.* for the ATLAS Collaboration 2012 *J. Phys.: Conf. Series* **396** 032095
- [4] Karavakis E *et al.* for the ATLAS Collaboration, “Common accounting system for monitoring the ATLAS Distributed Computing resources”, *Int. Conf. on Computing in High Energy and Nuclear Physics 2013* Amsterdam
- [5] Andreeva J *et al.* 2010 *J. Phys.: Conf. Series* **219** 062003
- [6] Maeno T *et al.* for the ATLAS Collaboration 2012 *J. Phys.: Conf. Series* **396** 032071
- [7] Tuckett D *et al.* 2012 *J. Phys.: Conf. Series* **396** 052069
- [8] Kokoszkieicz L *et al.*, “hBrowse - Generic framework for hierarchical data visualization”, Poster presented at the Computing in High Energy and Nuclear Physics 2012 International Conference
- [9] Schovancová J *et al.* for the ATLAS Collaboration, “ATLAS Distributed Computing Automation”, to appear in Proceedings of the 5th International Conference “Distributed Computing and Grid-technologies in Science and Education”, Dubna, Russian Federation, 16 – 20 Jul 2012
- [10] Anisenkov A *et al.* for the ATLAS Collaboration, “AGIS: The ATLAS Grid Information System”, *Int. Conf. on Computing in High Energy and Nuclear Physics 2013* Amsterdam
- [11] Elmsheuser J *et al.* for the ATLAS Collaboration, “Grid Site Testing for ATLAS with HammerCloud”, *Int. Conf. on Computing in High Energy and Nuclear Physics 2013* Amsterdam
- [12] Rodrigues De Sousa Andrade P M *et al.* 2012 *J. Phys.: Conf. Series* **396** 032008
- [13] Di Girolamo A *et al.* 2012 *J. Phys.: Conf. Series* **396** 032100
- [14] Tupputi S *et al.* for the ATLAS Collaboration, “Automating usability of ATLAS Distributed Computing resources”, *Int. Conf. on Computing in High Energy and Nuclear Physics 2013* Amsterdam

- [15] Branco M *et al.* 2008 *J. Phys.: Conf. Series* **119** 062017
- [16] Garonne V *et al.* for the ATLAS Collaboration, “Rucio – The next generation of large scale distributed system for ATLAS Data Management”, *Int. Conf. on Computing in High Energy and Nuclear Physics 2013* Amsterdam
- [17] Maeno T *et al.* for the ATLAS Collaboration, “Evolution of the ATLAS PanDA Workload Management System for Exascale Computational Science”, *Int. Conf. on Computing in High Energy and Nuclear Physics 2013* Amsterdam