

IT INFRASTRUCTURE DESIGN AND IMPLEMENTATION CONSIDERATIONS FOR THE ATLAS TDAQ SYSTEM

M. Dobson, G. Unel

University of California at Irvine, USA

C. Caramarcu

National Institute of Physics and Nuclear Engineering, Romania

I. Dumitru, L. Valsan, G.L. Darlea, F. Bujor

Politehnica University of Bucharest, Romania

A. Bogdanchikov, A. Korol, A. Zaytsev

Budker Institute of Nuclear Physics, Novosibirsk, Russia

S. Ballestrero

University of Johannesburg, South Africa

sergio.ballestrero@cern.ch

Keywords: HEP, LHC, ATLAS, DAQ, online computing, computing farms, parallel processing, access management, IT security, accounting systems, system administration, Linux.

Abstract: This paper gives a thorough overview of the ATLAS TDAQ SysAdmin group activities which deals with administration of the TDAQ computing environment supporting Front End detector hardware, Data Flow, Event Filter and other subsystems of the ATLAS detector operating on the LHC accelerator at CERN. The current installation consists of approximately 1500 netbooted nodes managed by more than 60 dedicated servers, a high performance centralized storage system, about 50 multi-screen user interface systems installed in the control rooms and various hardware and critical service monitoring machines. In the final configuration, the online computer farm will be capable of hosting tens of thousands applications running simultaneously. The ATLAS TDAQ computing environment is now serving more than 3000 users subdivided into approximately 300 categories in correspondence with their roles in the system. The access and role management system is custom built on top of an LDAP schema. The engineering infrastructure of the ATLAS experiment provides 340 racks for hardware components and 4 MW of cooling capacity. The estimated data flow rate exported by the ATLAS TDAQ system for future long term analysis is about 2.5 PB/year. The number of CPU cores installed in the system will exceed 10000 during 2010.

1 INTRODUCTION

The Trigger and Data Acquisition (TDAQ) System (Padilla, 2009; Zhang, 2010) of the ATLAS experiment (Aad, 2008) exploits a large online computing farm for the readout of the detector front-end data, the second level trigger and Event Filter farms, plus all the ancillary functions (monitoring, control, etc.). These systems are deployed underground (USA15 service cavern) and on the

surface (SDX1 hall, ATLAS main and satellite control rooms, etc.) at the experimental site. Two of these areas hold the majority of TDAQ equipment:

- USA15, provided with 220 racks (deployed on 3 floors) which are 70% occupied by TDAQ and ATLAS sub-detectors equipment. The equipment in this area uses 1 MW of power at the present moment while 2.5 MW of cooling capacity is available for future upgrades.



- SDX1, provided with 120 racks (deployed on 2 floors) which are currently 50% occupied by TDAQ equipment. The power consumption of this area is 0.5 MW at the moment and up to 1.5 MW of cooling capacity is available for the final design system.

At the moment, the ATLAS TDAQ system exploits more than 1500 computers and tens of thousands instances of various applications. These machines need to be administrated in a coherent and optimal way in order to maintain the computing farms at the highest level of availability in order for ATLAS to make the best use of available luminosity provided by the LHC collider.

A dedicated group of system administrators (the ATLAS TDAQ SysAdmin Group) is dealing with these tasks and, in addition, with the support of shifters and users of the ATLAS online computing system, on a 24/7 basis.

2 ATLAS TDAQ SYSADMIN RESPONSIBILITIES

The group maintains multiple ATLAS TDAQ related computing areas across the CERN sites:

- ATLAS Point1: SDX1, USA15, ACR & SCR (Main & Satellite Control Room),
- 3 laboratories across the CERN Meyrin site.

Everyday activities include:

- dealing with ATLAS Point 1 user requests,
- handling IT security issues,
- 24/7 service (shift or on-call, from mid-2008),
- hardware and software monitoring and maintenance of the computing infrastructure,
- commissioning of new hardware items.

All these tasks are handled in close cooperation with other relevant groups (ATLAS TDAQ Networking Team, ATLAS Technical Coordination, CERN IT Department) dealing with other aspects of maintenance and operation of ATLAS experiment. In addition the group carries out the development and validation of tools and solutions for automated user, software and hardware management, monitoring and control.

A variety of tools were developed within the group in order to automate the most frequently executed operations, for instance:

- ATLAS Point 1 user and role management scripts and web UI,
- Boot With Me (BWM) project components (control of netbooted nodes),
- storage area synchronization scripts,

- tools for registering new entities in the CERN network database (LanDB),
- configuration and command execution tool for clients and servers (ConfdbUI),
- tools for bulk firmware upgrade for the High Level Trigger Processor Unit (HLT PU) and Local File Server (LFS) nodes.

These tools are being intensively used and continuously improved.

3 SYSTEM DESIGN AND IMPLEMENTATION

3.1 Generic Design Overview

A major concern in any high availability computing farm is the mean time between failures. This is highly correlated to the mean time between failures of only a few of the components in the computers, such as disks. For this reason, TDAQ decided minimize the use of disks in the data acquisition system, by using diskless nodes, which are booted into Linux over the network. This approach has other advantages, such as ease of maintenance, reproducibility on a large scale, and the like. The BWM project was developed in order to respond to the need for a flexible system to build boot images and configure the booting of the diskless nodes.

Another major cornerstone of the system is therefore the boot servers. These are designed to serve the DHCP requests and boot images, and to provide network mounted disks for the main part of the OS and for the applications. In this kind of system, the servers are single points of failure. To overcome this limitation, the responsibility of booting and providing network drives is shared across two or more servers. This redundancy ensures the high availability of the clients independently of that of individual servers.

Another requirement for the system is the ability to run the experiment and take data for up to 24 hours while having lost the connection to the IT department and Tier0 centre (responsible for long term storage of the data, distribution of it to Tier1 centres, and also analysis of some of the data). The implications are that the system should replicate any services in IT which are vital such as DNS, NTP, user authentication; and should be able to buffer the event data. The latter is done using a few servers with large disk caches (12 TB each), able to handle the incoming rate of selected events to the disks, 300 MB/s, and to simultaneously sustain twice the output

rate from the disks to the permanent storage in IT (in order to catch up any connectivity loss).

3.2 Current Configuration

Currently the SDX1 and USA15 areas together host more than 200 fully occupied racks of equipment from ATLAS TDAQ and detector sub-systems, which can be subdivided into the following groups:

- core services, comprising a NAS, 2 gateways, 2 DNS, 2 CFS (Central File Server), web servers, LDAP and MySQL clusters,
- more than 1500 netbooted nodes served by 60 LFS (Local File Server), including 154 ROS (Read-Out System), 841 HLT PU (6728 CPU cores combined), 77 EB (Event Building nodes), 6 SFO (Sub-Farm Output buffer node to permanent storage), 64 online and monitoring nodes, 161 SBC (Single Board Computer),
- many special purpose, locally installed systems: 24 ACR (standard Control Room machines, 4-screen each), 45 SCR machines, Detector Control System (DCS) nodes, sub-detector PCs, public nodes.

The number of HLT Processor Units (PU) deployed in SDX1 up to now represent only 50% of total SDX1 capacity. This is assumed to be sufficient for the initial phase of the accelerator programme and will be increased to meet demand.

3.3 Increasing System Availability

Given the complexity and interdependency of the IT infrastructure components, it takes a significant amount of time to bring it to the production state after a complete shutdown. Measures were therefore taken to protect the TDAQ infrastructure from the power cuts of various origins. The whole facility is provided with two centralized UPS lines with diesel generators backup, plus the independent UPS lines that are available in SDX1 for mission critical equipment. Nowadays about 5% of equipment deployed in SDX1 is on UPS lines, some on dual power (both UPS and mains power), and just a few machines like CFS nodes are dual powered from both a dedicated locally installed UPS and general UPS power.

3.4 Hardware and Service Monitoring

Host monitoring is currently being done using the NAGIOS platform. NAGIOS allows the monitoring of various services for the machines. For all

machines, basic OS, network, and hardware state is monitored. On certain hosts, such as LFS or Application Gateways, also specific services are monitored, such as NTP, NFS, DHCP, etc.

The NAGIOS graphs collected by the monitoring system are stored on disks in the form of more than 29400 RRD files of total size approximately 5.5 GB. Status information of all the nodes and NAGIOS graphs for various parameters are published automatically in the monitoring section of the private ATLAS Operations web server. For selected indicators which are of crucial importance to the proper functioning of the ATLAS TDAQ infrastructure E-mail/SMS alerts are provided.

3.5 Node Management Tools

Due to the large number of netbooted nodes being managed, nontrivial automation is required to reduce the amount of time consumed by performing various routine operations, like rebooting a group of nodes, assigning clients to the LFS servers, etc. A dedicated set of tools based on a MySQL database and the GUI front end to them called ConfdbUI, were developed to cover these requirements. With few exceptions, all the nodes are currently running Scientific Linux CERN (SLC5).

The standard solution for the remote hardware management in ATLAS Point 1 is based exclusively on IPMI. Locally installed nodes, such as LFS, Control Room, central servers are installed and configured centrally, using a Quattor (Leiva, 2005) based configuration management system.

3.6 Remote Access Subsystems

The access to the ATLAS Technical and Control Network (ATCN) from outside of Point 1 is highly restricted and only allowed via one of the following gateway systems:

- Point 1 Linux gateways, allowing expert users to access a particular set of machines within Point 1 via the SSH protocol,
- remote Monitoring System, providing the graphical terminal services for remote monitoring of ATLAS sub-detector systems,
- Windows Terminal Servers, allowing experts to access the DCS SCADA system.

All the gateways and remote monitoring nodes are provided with both host based and network based accounting, security monitoring and intrusion prevention systems. The gateways are implemented on a XEN based virtualization solution which

ensures high availability and manageability for this vital subsystem.

3.7 Centralized Storage System

The storage solution used in the ATLAS TDAQ environment is based on a high performance storage solution fulfilling the following requirements:

- 3.2 TB of online high speed storage capacity, with scalability up to 10 TB,
- NFSv3, NFSv4, CIFS, and iSCSI support,
- serving up to 2500 clients simultaneously without degradation of performance,
- multiple levels of hardware and software redundancy (2N redundancy schemas).

The storage solution ensures the scalability of the centralized storage system with the increase of computing power in the TDAQ environment foreseen over the coming years.

3.8 User and Role Management

Due to the requirement to be operationally independent from the services provided by CERN IT, and to allow more flexibility, the experimental area has its own user directory stored on an LDAP cluster based on OpenLDAP software. The system is standalone but for consistency it is synchronized with IT for the usernames.

For authentication, the consistency is maintained by having a local slave replica of the central CERN Windows Domain Controller against which user credentials are validated. The only exceptions are the locally defined service accounts, for which the authentication information is stored in the LDAP.

Contrary to the way some of the previous experiments have been run, it has been decided to have user based authentication and not group based authentication, in order to have accountability and traceability of actions, as well as increased security.

The reasons for wanting to use group accounts in past experiments comes from the natural splitting of users into categories of people and tasks which they are allowed to do, for example some users are shifters, detector experts, TDAQ experts.

In order to address this categorization and still have the user authentication for accountability and traceability, it has been decided in TDAQ to implement a Role Based Access Control (RBAC) authorization system. Currently ATLAS Point1 is provided with the dedicated role based access control and authorization system currently holding more than 300 unique roles in a hierarchical structure. The total number of users registered is

more than 3000, but only a small fraction of them would be allowed to access Point1 remotely during the data taking period (experts which are on shift or on call).

3.9 Future Activities

The list of major milestones which are to be encountered in 2010 contains the following items:

- installing 80 new HLT PU machines (64 logical CPU cores/2U) thus increasing the amount of computing power installed in SDX1 beyond 10000 CPU cores,
- deployment of the new web servers provided with sophisticated load balancing solution,
- upgrade and increase of redundancy for the Active Directory service,
- expansion of the storage capacity on the central NAS,
- test and deployment of NFSv4 + Kerberos5 to overcome user/group limits of NFSv3.

4 CONCLUSIONS

The ATLAS TDAQ system is fully operational and now is re-entering an extended period of data taking for the ATLAS detector. The design of the system and the supporting computer infrastructure has been validated and the IT environment of ATLAS online systems has been running steadily since 2009.

However the activities devoted to increasing the usability of the system and the amount of computing power available for the experiment are still ongoing. This would allow the ATLAS collaboration to take full advantage of the increased luminosity of the accelerator over the next 20 months of data taking.

REFERENCES

- Padilla, C., *et al.*, 2009. The ATLAS Trigger System. *16th IEEE NPSS Real Time Conference 2009, Beijing, China, 10-15 May 2009.*
- Zhang, J., *et al.*, 2010. ATLAS Data Acquisition. *16th IEEE NPSS Real Time Conference 2009, Beijing, China, 10-15 May 2009.*
- Aad, G., *et al.*, 2008. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST 3 (2008) S08003.*
- Leiva, R., *et al.*, 2005. Quattor: Tools and Techniques for the Configuration, Installation and Management of Large-Scale Grid Computing Fabrics. *Journal of Grid Computing (2004) 2: 313-322. Springer.*