PLANNING FOR PREDICTABLE NETWORK PERFORMANCE IN THE ATLAS TDAQ

C. Meirosu[*#], B. Martin[*], A. Topurov[*], A. Al-Shabibi[†]

Abstract

The Trigger and Data Acquisition System of the ATLAS experiment is currently being installed at CERN. A significant amount of computing resources will be deployed in the Online computing system. More than 3000 high-performance computers will be supported by networks composed of about 200 Ethernet switches. The architecture of the networks was optimised for the particular traffic profile generated by data transfer protocols with real-time delivery constraints. In this paper, we summarise the operational requirements imposed on the TDAQ networks. We describe the architecture of the network management solution that fulfils the complete set of requirements. Commercial and custom-developed applications will be integrated in a solution that will provide a maximum of relevant information to the physics operator on shift and enable the networking team to analyse trends and predict the network performance.

INTRODUCTION

The infrastructure of the ATLAS experiment is currently being installed at CERN. The ATLAS detector will observe the results of proton-proton collisions and produce a raw data rate of about 60 TB/s. A Trigger and Data Acquisition System (TDAQ) [1] will analyse the data in real time and select the most interesting results of the collisions. The data rate will thus be reduced to about 300 MB/s at the input of the permanent storage system. Part of the TDAQ is implemented through large computer farms supported by an Ethernet network infrastructure [2]. More than 3000 high-performance computers will be interconnected by networks composed of about 200 Ethernet switches. Two types of networks can be distinguished in the TDAQ environment:

- Data networks: they were designed and dimensioned to support the request-response traffic that is typical to the TDAQ data path [3]. Each vertical component of the TDAQ data path comes with its own set of requirements in terms of the average transit time allowed. For example, a time budget of 15-20 ms is allocated for the Data Collection system [1], whereby typical processing takes about 1s in the Event Filter [1].
- Control networks: These are more general purpose networks, designed to support the network traffic associated with the control path of the TDAQ system.

Best practices in the industry recommend that network infrastructures as large as the one supporting the TDAQ system must be managed through dedicated software, as it is practically impossible to be monitored and diagnosed by human operators alone.

The article is organised as follows. After a brief introduction to the context of the TDAQ, we review the standards for network management. We then present the functional requirements of a network management solution adapted to the TDAQ. We continue by laying out the overall architecture of the proposed solution, followed by details regarding the actual implementation.

NETWORK MANAGEMENT OVERVIEW

A generic network management system has to integrate functionality that covers the following categories, as defined by the International Standards Institute (ISO) in the X.700 recommendation:

- The *security management* controls the access to the network resources
- The *accounting management* associates network traffic to the user or group of users that generated it, for tracking and/or billing purposes.
- The *fault management* identifies problems related to physical or logical failures in the network topology, stores the related information and notifies the network administrator about the problem. Under certain conditions, this component may implement automatic repair for basic problems.
- The *performance management* measures the current performance of the network, in terms of traffic load and response time for certain applications.
- The *configuration management* tracks and allows for controlled changes in the running configuration of the network devices.

The most popular paradigm of network management is based on the manager-agent model. The network resources are modelled as managed objects and expose

[*] CERN, Geneva, Switzerland
[#] "Politehnica" University of Bucharest, Romania
[†] EPFL, Lausanne, Swirzerland
Contact: catalin.meirosu@cern.ch

their properties across a set of standard interfaces provided by an agent. A manager application accesses the managed objects through the agent and implements the management policies. The communication between the manager and the agent takes place according to a management protocol.

The Internet and most local area networks are managed today through a manager-agent model based on the Simple Network Management Protocol (SNMP) [6] and the associated standards defined by the Internet Engineering Task Force (IETF). The data and statistics associated with the current state of operation of a network device are represented according to the Structure of Management Information (SMI) standard [4]. The network management information is organised in the form of Management Information Bases (MIBs) [5].

Even though standards exist, they are in a continuous transformation and upgrade process. New proposals are also initiated. It is customary practice in the industry to provide non-standard capabilities related to the management of specific devices, in addition to the standard-compliant ones. There are no standards to describe the configuration of a network device, even though efforts are being made in this direction by IETF through the Netconf proposal [7].

The day to day operation of IT infrastructures is guided by best practices. The Information Technology Infrastructure Library (ITIL) [8] is one of the most popular collections of best practices, drawn from the experience of both public and private sector enterprises. According to the Network Service Management (NSM) model defined by ITIL, the focus in a networked system should be on the service to be provided and not on the infrastructure that supports it [9].

The high availability of the service depends on the network architecture and on the operational model chosen for the management of the network resources.

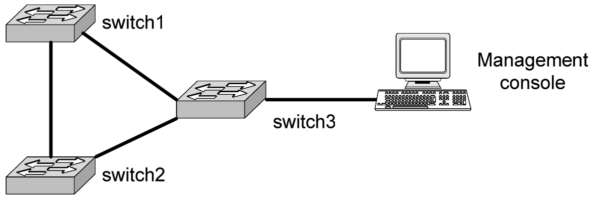

Figure 1 – Simple managed network architecture

A network is said to be operated using a *reactive* model in case the management system relies on collecting data about the network status at fixed intervals and can thus discover problems only after a certain time has passed. In the network presented in Figure 1, a failure of switch1 would only be observed when an information request issued by the management station times out.

A network is said to be operated in a *proactive* mode in case the management system relies on a combination of on demand collection and voluntary transmission of status information by the network devices. In the network shown in Figure 1, a failure of switch1 would be noticed immediately because both switch2 and switch3 would voluntarily report to the management station that they lost direct connectivity with switch1.

A network is said to be operated in a *predictive* mode in case the management system is proactive and also tries to predict the evolution of the network by analysing past performance and fault data. The predictive systems may take steps to automatically fix certain categories of problems before they arise, eliminating potential interruptions of the service.

Today’s networks have long passed the stage when they could be operated in a pure reactive way, with the network administrator solving problems as they appear. Collections of standards describe the available data and the exchange mechanisms. Best-practices suggest industry-proven ways of organising network operations. To deliver a predictable service, the planning of network operations should allow for potential problems to be anticipated centrally and addressed before they are encountered by the users.

FUNCTIONAL REQUIREMENTS

The networks that support the TDAQ are an integral part of the Online processing system attached to the ATLAS detector. The Technical Design Requirements document stipulates that the TDAQ should maximise the network uptime [1], in view of the rarity of the events that are sought by the detector (for example, a Standard Model Higgs boson may be observed at a rate below 0.001 Hz [1]).

Two strategies were employed in order to maximise the network uptime. The first step consisted in choosing an adequate topology [2] and network components with both a high level of internal redundancy and hot-swap capability for the switch line cards. The second step involves the efficient management of the network, using a predictive operational model. In the following, we will concentrate on some of the requirements which are specific to the TDAQ system. The complete set of requirements is presented in [10].

The TDAQ networks will be located at the ATLAS experimental site and are completely isolated from the CERN site standard networking infrastructure [11]. Therefore, in view of the strict access rules and the reduced number of applications that will send data over these networks, we consider that the *security* and *accounting management* components of a network management framework are of lesser importance in this context. The computer and network security-related aspects of the experimental infrastructure are covered in detail elsewhere [11], [13].

A fault in the network may generate a large number of messages that report how various devices are affected. Root cause analysis is one of the methods recommended by literature for finding the primary source of an error [21]. The *fault management* component of the network management solution must use knowledge about the topology of the network in order to perform efficient root cause analysis on the messages that are transmitted by network devices via the MIB / SNMP mechanism.

The TDAQ applications rely on the Data networks for transferring the data acquired by the detector towards the computers in the filtering and analysis system. The traffic profile that is expected to be generated by these applications during regular production runs was thoroughly analysed in [12]. The networks were designed to respond optimally to this particular traffic profile [3]. However, the actual traffic profile in the experiment will be determined by the distribution of the physics data gathered by the different components of the detector system. This distribution may, momentarily and unpredictably, differ substantially from stable-state operation assumptions. The *performance management* component of the TDAQ network management solution has to allow rapid identification of traffic hotspots in the network.

In our view, the network management solution has to enable application developers to gather information about the network traffic to help debugging problems in the data transfers. This information must be obtained at the exact locations on the network that are relevant for a particular application.

To allow the debugging of application data transfer problems, the *performance management* will have to include the following functionality:

- to acquire information about the network traffic with high granularity timewise
- to observe traffic with low granularity in order to predict long term trends
- to capture the traffic at many points on the system and extract detailed statistics that would allow an off-line reconstruction of the traffic pattern

The same application may be used in different roles. The amount of data delivered by the detectors to the TDAQ system is highly dependent on the role that is imposed on the application. It is thus possible for the Data networks to be oversubscribed during particular debugging or calibration runs, especially when different detectors are performing such tests concurrently. Different levels of network assured quality of service have to be assigned for the different roles of the applications, in order to keep potential traffic loss under control. The assignment of network attached resources to application roles is performed at runtime, depending on the particular needs identified by the physics operators on shift. The *configuration management* has to be able to impose, dynamically and on demand from the application, the configuration that would guarantee predictable results for the applications running on the data network.

The fact that the TDAQ networks are an integral part of a high-energy physics experimental infrastructure determined an additional and rather non-standard requirement. The TDAQ network management solution has to signal the state of the network to the software that controls the experimental infrastructure. It is unusual for a network management system to report network status to end-users, but this fact is highly justified in our scenario. The physics operator on shift needs timely information about the entire infrastructure of the experiment, and the TDAQ networks are just another component of this infrastructure.

OVERALL ARCHITECTURE AND IMPLEMENTATION

In view of the requirements that are specific to the operation in the TDAQ environment, we were unable to find a network management suite, commercial or open source, which would provide all the necessary functionality. These requirements lead to the architecture presented in Figure 2.

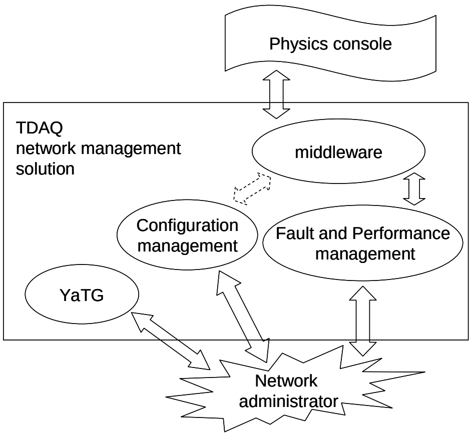

Figure 2: A generic architecture for the TDAQ network management solution.

The generic architecture of the TDAQ management solution consists of four components, each of them covering one or more of the functional requirements. The following options were considered in terms of implementation and will be discussed below for each individual component:

- use open source software
- purchase commercial software
- develop in-house, perhaps based on existing open-source libraries

Fault management

Fault management is still a topical research area. The complexity of the problems to be handled made in-house development a non-option in this case. Open source programs (like OpenNMS [18]) contend themselves with passing messages received from the network devices to the operator by means of a console, with no filtering or reasoning about it. All of the major commercial software packages that specialise in network management include a module that employs root cause analysis or related methods for filtering received messages. Therefore, commercial software was the obvious choice for covering the fault management including root cause analysis requirement. In further selecting between the options available on the market, we preferred the solution that ensured interoperability with the ATLAS technical network.

Performance management

As with other commercial packages, the network management software we chose provides support for low-frequency (at polling intervals higher than about 1 minute) monitoring of the network traffic and a plethora of reporting options. Therefore, the same component implements the *fault management* and part of the *performance management* capability in our system.

A program called YaTG [15] is being developed for the purpose of obtaining high granularity views of the network traffic. The high-frequency monitoring of the network answers the requirement for helping application developers. We believe that detailed views of the network traffic, together with the combined expertise of the application developers and network support staff, will provide a rapid way for identifying potential sources problems.

A first version of YaTG was developed [15] and used for studying the SNMP implementation on several network devices. We discovered, for example, that few manufacturers have taken the necessary steps to allow the SNMP agent to respond with updated counter values for a fast polling process such as the one induced by our tool. Other issues related to SNMP implementations are described in [15]. The existence of the preliminary version of YaTG allowed us to open discussions with network device manufacturers about improving the SNMP implementation.

Configuration management

The role of the configuration management in the context of the TDAQ networks is still under discussion. The minimum functionality would include detecting unauthorised changes in the configuration of the network devices and the possibility to restore the configuration of the entire network in a known state.

An open-source collection of scripts called Rancid [19] is capable of monitoring configuration changes in a broad category of network devices via a versioning system. Commercial software comes with highly increased functionality, which includes ITIL compliance. The market targeted by such implementations is that of Internet Service Providers and large enterprises, which is unfortunately reflected by the price of such solution.

Being dimensioned to accommodate the traffic generated by ATLAS at design luminosity, the TDAQ networks will require virtually no upgrades for many years. Once installed, the physical topology of the TDAQ networks will be static. The operational model allows for few changes in the logical topology of the network. The ITIL-compliant tools provide standardised methods for applying changes in the configuration, together with analysing how a specific change might affect the current configuration. However, in view of the rather static physical and logical configuration of the networks, we consider the use of such a commercial tool to be overkill in this scenario.

The fact that TDAQ applications in different roles may use different sets of resources attached to the network calls for traffic differentiation at the network level. Several solutions could be used for providing such differentiation. The minimum functionality of the *configuration management* will have to be augmented in order to support either of the options presented below.

One of the options calls for the network to dynamically adapt to the role of the application, for example by dynamically defining VLANs according to the roles. The existence of a network configuration service that interacts directly with the application would be required in this case. Software supporting similar functionality is available for wide-area networks [20]. However, it has yet to reach production grade and the level of interactivity required by a daily use in the TDAQ environment.

Another potential solution relies on changing traffic prioritisation profiles while maintaining the rest of the network configuration. Each role of an application would be associated with a quality of service profile, defined by the priority field in the VLAN tag. Up to 8 roles could be thus accommodated. This approach relies on the end nodes to tag packets according to the role of the application.

Integration on the physics console

The Online software [14] controls the operation of the entire TDAQ system. The physics operator on shift accesses a graphical user interface (IGUI) that displays the operational status of all the components of the system. The Online software is built in Java and uses a CORBA broker to communicate with the other components of the Online infrastructure [14]. The network management system also uses a CORBA interface to expose the state of the network. A middleware package will be developed to interface the fault management to the experimental control software.

The middleware will be considered part of the software infrastructure category of the TDAQ Online applications. A unique instance will be active at any given time in the system. The delivery of network status information will follow a subscriber-based policy. Multiple instances of the IGUI may receive the network status information on request or the information may be pushed to all the subscribers as it arrives from the network management system.

The physics operator on shift will have the possibility of opening a graphical view of the network, as generated by the network management system, through a dedicated panel on the IGUI. This will help in rapidly assessing the implications of a potential network problem.

CONCLUSION

The TDAQ networks are built using off-the-shelf components and optimised for the particular requirements of the ATLAS system. Achieving predictable network performance is a strong necessity in such an environment. The network was dimensioned and the devices acquired having this goal in mind.

We reviewed the industry best practices and considered the standards for network management in the context of the TDAQ networks. We designed a solution that integrates commercially available and open source components to respond to generic requirements in terms of fault, performance and configuration management. Specific needs, like providing help to the application developers and reporting the network status to the physics operator on shift require in-house software developments.

Future work is concentrated in carrying out the software development and further re-examining the specific configuration management requirements and solutions.

REFERENCES

[1] The ATLAS TDAQ Technical Design Report, June 2003

[2] S. Stancu, C. Meirosu, M. Ciobotaru, L. Leahu, B. Martin, "Networks for the ATLAS TDAQ", in proceedings of CHEP 2006, Bombay, India, February 2006.

[3] S. Stancu, M. Ciobotaru, K. Korcyl, B. Martin, "ATLAS TDAQ DataFlow network architecture analysis and upgrade proposal", in proceedings of the 14th IEEE - NPSS Real Time Conference 2005 Nuclear Plasma Sciences Society RT 2005 , Stockholm, Sweden , 4 - 10 Jun 2005.

[4] K. McCloghrie, D. Perkins, J. Schoenwaelder, "Structure of Management Information Version 2 (SMIv2)", RFC 2578

[5] K. McCloghrie, M. Rose, "Management Information Base for Network Management of TCP/IP-based internets:MIB-II", RFC 1213

[6] D. Harrington, R. Presuhn, B. Wijnen, "An Architecture for Describing SNMP Management Frameworks", RFC 2571

[7] R. Enns (ed.), "NETCONF Configuration Protocol", draft-ietf-netconf-prot-10, work in progress

[8] http://www.itil.co.uk/

[9] L. J. G. T. van Hemmen, "Models supporting the network management organization", Int. J. Network Mgmt 2000; 10:299 – 314

[10] ATLAS TDAQ Network Management System Requirements, work in progress

[11] M. Dobson et al., "The architecture and administration of the ATLAS online computing system",

[12] S. Stancu – PhD thesis.

[13] http://wg-cnic.web.cern.ch/wg-cnic/

[14] The ATLAS Online Software, http://atlas-onlsw.web.cern.ch/Atlas-onlsw

[15] A. Al-Shabibi, S. Stancu, B. Martin, C. Meirosu, "A high speed monitoring tool for the ATLAS TDAQ Network", submitted for publication

[16] M. Ciobotaru, S. Stancu, M.J. LeVine, B. Martin, "GETB - A Gigabit Ethernet Application Platform", in proceedings of the 14th IEEE - NPSS Real Time Conference 2005 Nuclear Plasma Sciences Society RT 2005 , Stockholm, Sweden , 4 - 10 Jun 2005.

[[17] T. Chiu, "Getting Proactive Network Management From Reactive Network Management Tools", Int. J. Network Mgmt., 8, 12–17 (1998)

[18] The OpenNMS project home page, http://www.opennms.org/wiki/

[19] RANCID - Really Awesome New Cisco confIg Differ, http://www.shrubbery.net/rancid/

[20] The Advanced UCLP Services - A Graphical Resource Management Tool for Creating & Managing Articulated Private Networks, http://www.uclp.ca/uclpv2.html

[21] A.T. Bouloutas, S. Calo, A. Finkel, Alarmcorrelation and fault identification in communication networks, IEEE Transactions on Communications 42 (2–4), 523–533 (1994).