

Role Based Access Control system in the ATLAS experiment

M L Valsan^{1,9}, M Dobson², G Lehmann Miotto², D A Scannicchio³, S Schlenker²,
V Filimonov⁴, V Khomoutnikov⁴, I Dumitru¹, A S Zaytsev⁵, A A Korol⁵, A
Bogdantchikov⁵, G Avolio², C Caramarcu⁶, S Ballestrero⁷, G L Darlea¹, M
Twomey⁸, F Bujor¹

¹ Politehnica University of Bucharest, Romania

² CERN, Geneva, Switzerland

³ University of California at Irvine, USA

⁴ Petersburg Nuclear Physics Institute, Russia

⁵ Budker Institute of Nuclear Physics, Russia

⁶ National Institute of Physics and Nuclear Engineering, Bucharest, Romania

⁷ University of Johannesburg, South Africa

⁸ University of Washington, USA

⁹ E-mail: liviu.valsan@cern.ch

Abstract. The complexity of the ATLAS experiment motivated the deployment of an integrated Access Control System in order to guarantee safe and optimal access for a large number of users to the various software and hardware resources. Such an integrated system was foreseen since the design of the infrastructure and is now central to the operations model. In order to cope with the ever growing needs of restricting access to all resources used within the experiment, the Roles Based Access Control (RBAC) previously developed has been extended and improved. The paper starts with a short presentation of the RBAC design, implementation and the changes made to the system to allow the management and usage of roles to control access to the vast and diverse set of resources. The RBAC implementation uses a directory service based on Lightweight Directory Access Protocol to store the users (~3000), roles (~320), groups (~80) and access policies. The information is kept in sync with various other databases and directory services: human resources, central CERN IT, CERN Active Directory and the Access Control Database used by DCS. The paper concludes with a detailed description of the integration across all areas of the system.

1. Introduction

The Role Based Access Control (RBAC) system used inside the ATLAS [1] experiment allows secure, reliable and optimal access to the software and hardware resources for a large number of users. The ATLAS experiment has a large user base (around 3000 users) requiring access either locally (from inside the experimental network, ATLAS Technical & Control Network - ATCN), from inside the CERN's General Purpose Network (GPN) or remotely from any of the external institutes.

The access control system regulates which operations can be executed on data and resources, preventing possible intentional or unintentional harmful actions.

The RBAC model was found to be best suited for implementing the access policy for the ATLAS experiment, as it offers the flexibility to accommodate a large number of users, roles and resources [2].



The ATLAS operational model defines activity areas with their tasks and responsibilities for the various systems and the resulting structure is naturally reflected by the roles associated to systems and sub-systems. Therefore the hierarchical RBAC was the best choice for an underlying access control schema as it is a model for controlling access to resources where permitted actions are identified with roles rather than with individual subject identities.

2. System design

The scope of the ATLAS access control system is the ATLAS online computing environment where most of the hardware and the software resources are located or are controlled from. The wide scope requires a design for the access control system that allows for a centralized management of the access policies and, at the same time, permits implementations from the lowest operating system (OS) [3] level to the highest software application level.

The NIST RBAC reference model [4] defines four model components: Core RBAC, Hierarchical RBAC, Static Separation of Duty Relations (SSD), and Dynamic Separation of Duty Relations (DSD).

Inside ATLAS a simplified RBAC system has been adopted, without separation of duty constrains, either SSD or DSD. The data elements of the ATLAS RBAC system are:

- users: computing accounts associated with human beings or software applications;
- roles: job functions or job titles which define an authority level;
- resources: object which supports a set of possible actions;
- permissions: approvals to perform an action on a given resource;
- sessions: mappings between a user and a subset of enabled roles.

The RBAC system used in the ATLAS experiment takes the access decision for an individual user based on the roles enabled for the user. A user is assigned a role when he has the expertise required for said role; the role is then enabled when the user needs to act in that role. The role, on the basis of policies defined, determines which resources can be accessed and permission is being granted only if the user has the required role enabled.

The element sets and the relations between them are shown in figure 1.

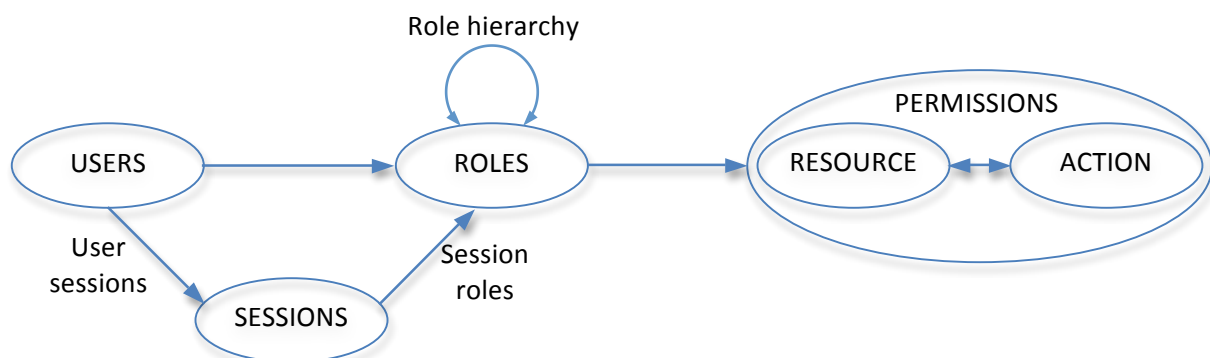


Figure 1. The RBAC element sets and relations.

All these elements and the relations among them are stored in a database-like directory: users, roles, roles hierarchies, permissions, and user assignments to roles. Implementing the data repository for all RBAC constituents as a single storage allows for better control over the coherence of the access control policies.

3. Implementation

Roles used inside ATLAS have been designed to have both an organizational (subsystem or project) and a functional component (reflecting the expertise levels within a subsystem). The components of a role are depicted in figure 2.

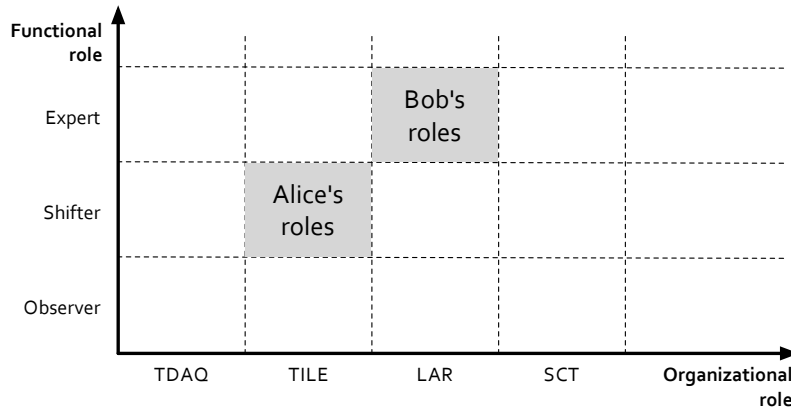


Figure 2. The functional and organizational components of ATLAS roles.

Permissions are attached to roles located at the bottom of a role hierarchy (like a task based organization). The roles assigned directly to users are at the top of hierarchy. The two classes of roles are connected via intermediary roles, allowing for a fine-grained and flexible permission allocation.

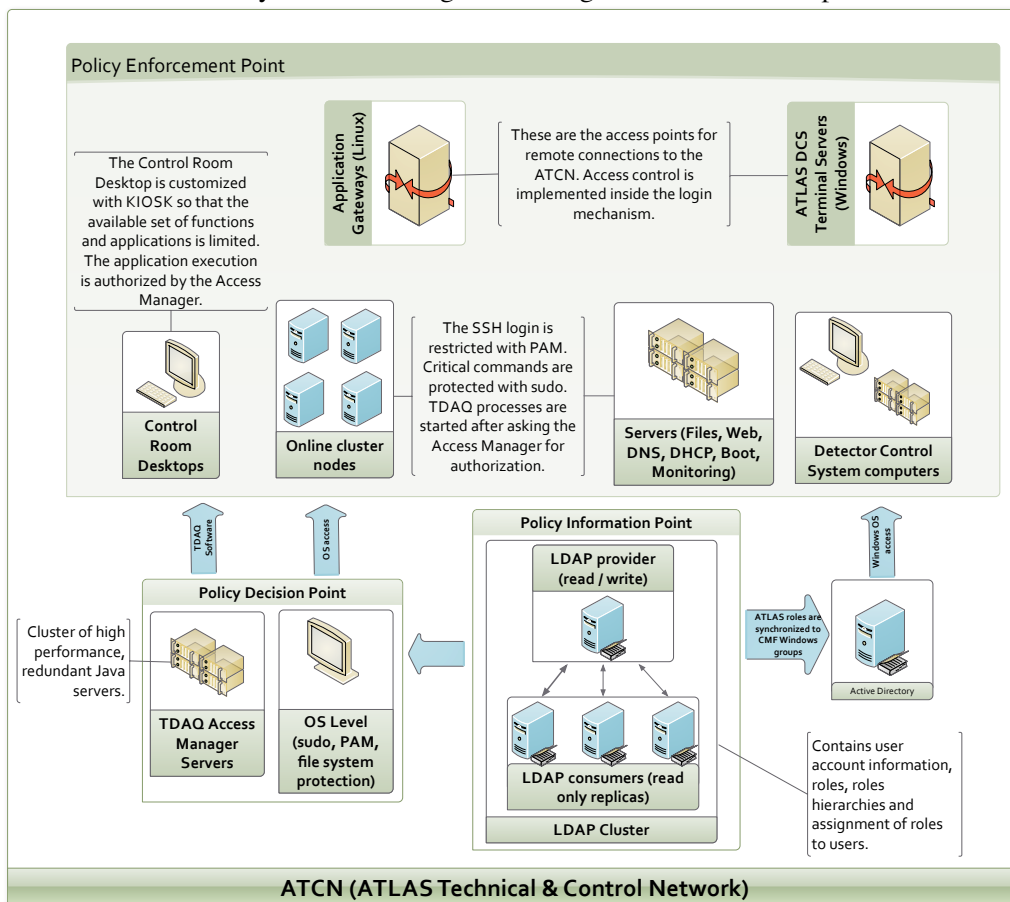


Figure 3. Access management inside the ATLAS Technical & Control Network (ATCN).

Global base roles are used as building blocks. These are not directly assignable to users and generic access policies are defined for them. More specific roles are assignable to users, inherit other roles including base roles not assignable to users and can get additional permissions.

Roles are defined in LDAP (Lightweight Directory Access Protocol) as NIS (Network Information Service) netgroups. This approach offers several advantages, including the ability to model a hierarchical structure and integration at the operating system / application level.

A global overview of the access control system deployed inside the ATLAS Technical & Control Network can be seen in figure 3.

4. The TDAQ Access Manager component

The Access Manager [5] is the component that implements the access management for Trigger and Data Acquisition (TDAQ) software. It has been designed as a highly scalable system, capable of handling hundreds of requests in parallel and has been implemented using a client-server architecture:

- The client sends authorization requests to the server
- The server processes the requests and sends back the responses to the client

Three main classes compose the Access Manager dataflow (see figure 3):

- Policy Information Point: it acts as a source of attribute values required for the policy evaluation. The required information is retrieved from LDAP and is cached for a certain amount of time.
- Policy Decision Point: the place where an authorization request is evaluated against one or more policies to produce a decision.
- Policy Enforcement Point: places where access control is enforced by making a request and applying the authorization decision.

The server component has been developed as a high performance Java server, designed around a reactor pattern. For message passing Sun's implementation of OASIS eXtensible Access Control Markup Language [6] is being used. The server listens on two ports:

- Service port: receives authorization requests from clients and sends back the authorization decisions; messages are XACML formatted, sent over TCP/IP.
- Control port: receives short informative messages or questions and sends back answers (notification messages when LDAP has been updated, monitoring information, etc); messages are string formatted, clear text, sent over UDP.

The various TDAQ software components act as Access Manager clients, using the client API, available for Java and C++.

5. Integration of the RBAC system

In order to have a consistent security policy across all ATLAS subsystems, a synchronization mechanism was implemented to integrate the RBAC policies with various other security components, depicted in figure 4:

- The Detector Control System (DCS) Access Control
- Login restrictions to Windows systems
- POSIX groups used at the filesystem level

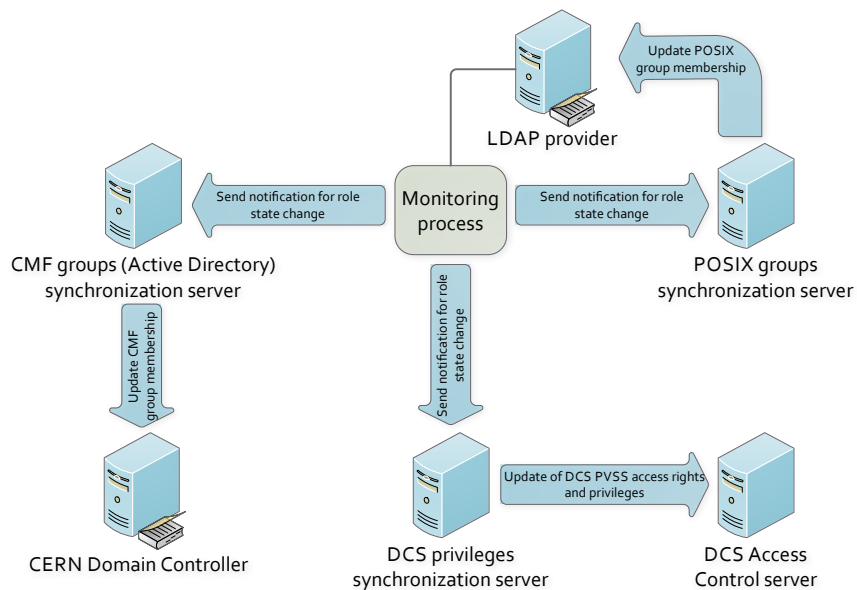


Figure 4. Synchronization of role information with other systems.

The synchronization mechanism is based on a monitoring process, running on the LDAP provider server, which sends notifications to several synchronization processes whenever a role to user relation changes. The connection between the two peers is encrypted with SSL using host certificates.

Whenever a change has been detected the monitoring process analyzes all changes from the last checkpoint and constructs two lists of users: one with users for which only an enable / disable operation has taken place and a list of users for which any kind of role operation was performed (enable / disable / assign / revoke). The two lists are needed as some synchronizations are done only after an enable / disable operation while others are done regardless of the type of operation. The trigger sent consists of a list of users for which roles associations changed. After receiving a trigger the listener process connects to one of the LDAP servers, gets the updated user to role relationships and performs the required changes.

5.1. Integration of access control policies across different operating systems

The first protection level for the ATLAS computing system is provided by the application gateways and terminal servers that separate ATCN from the CERN General Purpose network. Users outside ATCN wishing to login on a node inside ATCN have to login to the application gateways and then hop to the desired machine where the access is restricted with the operating system native access control mechanism and, in addition, taking into account the role of the user in the system.

The Pluggable Authentication Modules (PAM) are used for controlling access on the Linux systems. Having the roles defined in LDAP as netgroups means that they can be used directly by PAM. The Windows systems used inside the ATLAS experimental network are managed via the Computer Management Framework (CMF), which in turn stores entities needed for enforcing access in the CERN Domain Controller. In order to have a consistent access policy, roles are automatically synchronized to the CMF specific Active Directory groups. These groups are used for enforcing access restrictions via the Active Directory policies.

5.2. Integration with the Detector Control System (DCS)

The DCS Access Control system uses PVSS (Prozessvisualisierung und Steuerungssystem) specific groups to control access and to enforce restrictions. For each of the users triggered the full list of enabled DCS-related roles (explicit and inherited) is obtained. This list is then passed to the DCS Access Control system to avoid the need for duplicating and synchronizing the whole role hierarchy.

5.3. Integration between RBAC roles and POSIX groups

It's not always possible to directly use the RBAC roles at the OS level, due to the fact that netgroups are not supported in all instances. For example, file system restrictions can't be enforced based on netgroups. In order to solve this issue, a synchronization mechanism has been set in place to automatically create and update POSIX groups. The LDAP schema used for defining roles supports an argument, specifying whether a role should be automatically synchronized to a POSIX group or not. The synchronization process takes care of automatically creating / removing the POSIX group, as well as keeping their contents up to date with the user to role associations.

6. Performance

The Access Manager server component is distributed over a number of 7 nodes (6 primary servers and a backup one) in a redundant, high performance manner. The load gets distributed among all the nodes, access being done through a round-robin DNS alias, for the primary servers. In case the client can't contact one of the primary servers it automatically fails over to the backup server.

A single computing node (2 x Intel Xeon CPU 5130 2.00GHz, 4GB RAM) is capable of handling up to 800 requests per second. The aggregate number of requests served by the Access Manager servers used inside the ATLAS experiment is around 2500 requests per second.

Figure 5 shows the number of requests served per second by each of the primary Access Manager servers along with the combined values. The performance obtained in the current setup is adequate for the ATLAS online computing farm and the architecture is scalable.

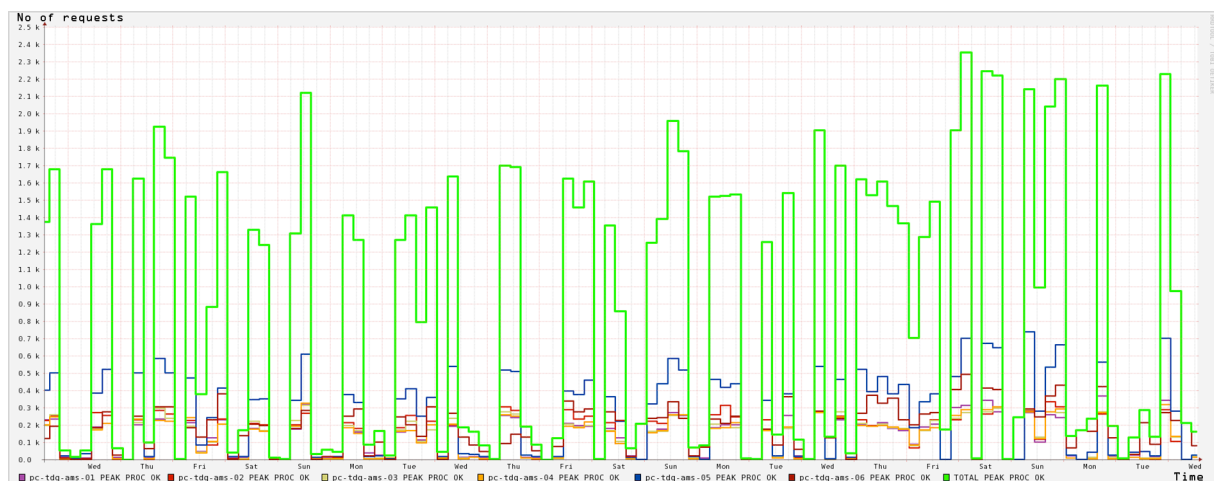


Figure 5. Number of requests served by the Access Manager servers.

References

- [1] The ATLAS Collaboration *et al* 2008 The ATLAS Experiment at the CERN Large Hadron Collider *JINST* **3** S08003
- [2] Leahu M C, Dobson M and Avolio G 2008 Access Control Design and Implementations in the ATLAS Experiment *IEEE Trans. Nucl. Sci.* **55** 386-391
- [3] Adeel-Ur-Rehman A *et al* 2010 System administration of ATLAS TDAQ computing environment *J. Phys.: Conf. Ser.* **219** 022048
- [4] Ferraiolo D *et al* 2001 Proposed NIST standard for role-based access control *ACM Trans. Inf. Syst. Security* **4** 224-274
- [5] Sloper J E, Leahu M, Dobson M and Lehmann G 2006 Access management in the ATLAS TDAQ *IEEE Trans. Nucl. Sci.* **53** 986-989
- [6] Moses T 2005 OASIS Standard for eXtensible Access Control Markup Language (XACML) Version 2.0 http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf