INTEGRATED ACCESS CONTROL FOR PVSS-BASED SCADA SYSTEMS AT CERN

P. Golonka, M. Gonzalez-Berges, CERN, Geneva, Switzerland

Abstract

The protection of the PVSS-based (Process Control and Visualization [1]) Human-Machine-Interface parts of the Control Systems for the LHC accelerator and the experiments at CERN is implemented using the JCOP (The Join COntrols Project) Framework Access Control component. It ensures the protection from non-malicious actions (such as misuse due to an operator's mistake) by enabling/disabling parts of the User Interface. It extends the native PVSS mechanisms for user-authentication and makes the management of the role-based authorizations easy to configure and maintain. Ultimately, it enables the synchronization of the access-control related data across distributed systems, and allows the synchronization of this data with central user-management resources at CERN (such as Active Directory). It also automates the creation of user accounts and roles (provisioning).

INTRODUCTION

Control systems of the experiments as well as of some subsystems of the accelerators are implemented using a commercial SCADA product, PVSS[1], extended with the components of the JCOP Framework [2][3]. They are typically composed as a distributed set of collaborating and autonomous PVSS applications, with a count as high as 150. Each of this system requires to be protected from non-malicious attacks, or mistaken errors of operators through a protection of the operators' screens (HMI). The integration of the protection mechanisms as well as centralized management, account provisioning and authentication mechanisms are the task of the JCOP Framework Access Control component [3].

The Access Control Component is used in the control systems for the large physics experiments which have been built by large international collaborations of physicists; this puts additional, demanding requirements on the homogenity of the access control mechanisms. During the commissioning phase, the Detector Control Systems are operated by a number of shifters present in the control room: each of them must be able to access only parts of the system. Also later, during the LHC runs, an important number of experts would need occasional access to certain parts of the systems. Their qualifications will differ and evolve dynamically; furthermore, some of them will only stay at CERN for a limited period of time. Unlike typical industrial systems, where the trained staff operates the system throughout its lifetime with not much rotation, the CERN applications require frequent changes to user's permission and a flexible authorization model allowing for effective administration of the users and their rights.

An additional level of complication is brought by the fact that the access control configuration data needs to be uniformly replicated into many elements of large distributed PVSS systems.

The main requirement for the JCOP Framework Access Control Component [3] was the central management and deployment of the access control data into a large number of PVSS systems, as well as the ways to synchronize the data with external user directories. The other, not less important, was the availability of the proper user/group (role) management tools, giving the powers of the Role-Based Access Control (RBAC) [4] model to the administrators.

The PVSS product already provided an implementation of access control, which would fulfil the requirements of standard mid-scale industrial applications. However, in our environment with a huge variety of domains that needed protection, this mechanism needed to be extended in a scalable way, putting also the ease of use and integration with JCOP Framework tools as a focus. The Access Control components builds on top of native access control mechanisms of PVSS and maintains compatibility to the maximum possible level, to profit from PVSS system-integrity protection and activity logging.

Like for the other JCOP Framework components, portability across Linux and Windows was implemented.

In the coming chapters we will focus on the four typical areas for access control implementations, namely the authentication, authorization, subjects being protected and the access-right enforcement mechanisms. We will also present our achievements in the area of deployment in large system and provisioning mechanisms.

SUBJECTS, DOMAINS AND ENFORCEMENT MECHANISMS

A notable feature of the large control systems at CERN is their complexity not only in terms of the number of elements being operated, but also the variety of subsystems. This implies the large number of subjects that need to be targeted by access control, independent from the others. As a consequence the number of permissions that need to be defined in the system, even when rationalized, become large. This exceeds the requirements on typical industrial control systems and thus the capabilities of the so called *system authorizations* mechanism implemented natively in PVSS.

To address this issue we came up with the concept of *domains* and *privileges*. We defined the *domain* to be an entity, which can be real, conceptual or organizational, that needs to be protected. It may be understood as the subject being protected. For example, a domain can be the cryogenic system for a sector of an accelerator, or a specific hardware item such as power supply. Each domain is uniquely identified by its name.

Within a domain, we define a set of *privileges* corresponding to permissions that are specific for the domain. The privilege levels may be defined separately for each domain. The standardized levels recommended by JCOP are Monitor, *Control*, *Debug*, and *Modify* which correspond to permissions for displaying the information about the subject, operating (controlling) it, performing expert operations and modifying its structure.

The most elementary item for which the authorization could be defined is therefore a domain with a privilege inside it. We call the pair domain-privilege *access rights*, and note them as colon-separated domain name and privilege name, for instance *Calo/HighVoltage:Control*. By collapsing the two-dimensional domain-privilege model, one obtains the classical flat list of all permissions,

An additional advantage of our model is the abstract, composite and generic approach to organizing the permissions, rather than strictly referring to the particular features of, for example, a piece of equipment. For instance, through the abstraction of domain and privileges one would rather define *HighVoltage:Control* and *HighVoltage:Expert* than a set of fine-grained access rights such as *MyPwSupl:On*, *MyPwSupl:Off*, *MyPwSupl:SetVoltage*, *MyPwSupl:SetCurrent*, etc...

To enforce the permissions defined by domains and privileges, we put the protection at the level of the user-interface. This is implemented by enabling and disabling the elements of operational panels (such as command buttons), depending on the current authorizations (i.e. access rights) available to the user who logged in to the system. The recommendation was given that the UI elements should not be hidden, but rather set to a disabled state, where they could not be operated, yet showing their potential availability.

The implementation of the enforcing mechanism is the responsibility of the application developer. The easy-to-use API (Application Programming Interface) of the Access Control component provides all the necessary functionality.

The native PVSS system authorizations are mapped within our model by means of a special domain, called *System*. The permissions in this domain are linked to PVSS authorization levels and allow restricting certain actions such as modification of the structure, use of the engineering tools, or acknowledging of alerts. These restrictions are enforced by PVSS itself.

The access control component has been integrated with other components of the JCOP Framework, notably with the FSM (Finite State Machine) tool used to operate the detectors. The Access Control component is used to control the ownership of the FSM tree, and limit the access to expert commands.

ROLE-BASED MODEL

Role-based access control (RBAC)[1] [4] is a generally adopted best practice for implementation of access control systems. This model is used in the JCOP Framework Access Control component.

In the RBAC approach, permissions (authorizations) required to fulfil the activities related to a certain role are coalesced and granted to a *role*; the users, gain the authorizations by taking (or being assigned) certain roles.

The model described above, in its most simple way, corresponds to the one already present in PVSS, where permissions are assigned to groups, and the users gain these permissions through the membership to certain groups, i.e. the groups are interpreted as and correspond to roles. This model, referred to in [4] as *Flat-RBAC* needed to be extended in order to implement role-hierarchies and inheritance, i.e. to implement the second level of the specification: *Restricted Hierarchical RBAC*. This has been achieved by implementing the groups-in-groups model. In our approach, the group (role) A that contains groups B and C, has the permissions inherited from B and C, in addition to the ones explicitly granted.

The Access Control component implements also a simplified version of the *Dynamic Separation of Duties* (DSD) through the so called *strict role checking mode*. By default, this mode is inactive, and the user has at disposal all the privileges obtained through the assigned roles (or group membership). Once the *strict role checking mode* is active, the user needs to explicitly select the active role, out of all his available roles, using the access-control UI element. Only the permissions coming from the active role are available and the user needs to make a conscious choice of his current role.

AUTHENTICATION

The PVSS native authentication mechanism is based on local credentials checking: the user name and password are checked against the ones stored in encrypted form locally, similarly to traditional POSIX password-checking. For large, distributed system, a local copy of credentials data needs to be stored, and maintained.

One of the important requirements for the Access Control Component was to allow to log into a PVSS application using the same credentials as used for CERN central services (i.e. common credentials used by central Windows, email, Linux services, and many web applications at CERN). Since the credentials verification could be done against CERN Active Directory server, using the LDAP protocol [6], the LDAP extension for PVSS and a corresponding authentication method were implemented.

The LDAP authentication is not limited to the use of CERN central Active Directory servers. It has been used successfully with other identity management products, allowing for LDAP credentials verification, such as Oracle Internet Directory or Open LDAP.

[1] RBAC term used as a general concept, not to be mistaken with the RBAC project for the LHC, described in [5]

To better satisfy our users' requirements, the authentication in the Access Control component was designed and implemented as an extendible, pluggable and customizable mechanism. The flexibility allowed the implementation of numerous customized authentication methods using the standard PVSS scripting mechanisms. The following interesting concepts were applied

- using more than one authentication server and a fall-back mechanism to increase the availability, or host special accounts
- integrate RFID card readers for authentication using standard CERN access cards; LDAP query is used to associate the retrieved data with person's identity
- checking the IP address of the user interface machine, to distinguish between logins from the control room and other places (such as offices)
- marking certain accounts as local: the authentication for these account is always performed locally
- credentials caching as a fall-back mechanism for cases of authentication server not being available

LARGE-SYSTEM DEPLOYMENT

The *Access Control (AC) Server* was developed to address the issue of replication of the access control data into all subsystems of a large distributed PVSS system,.

It was designed to be an additional PVSS service that could be easily activated on any machine. In a set up with AC Server, the server becomes the master for the access control data and configuration; all the changes in the user/group/domain configuration or privileges mapping are immediately propagated to the slave systems. To prevent modifications performed by local administrators, the AC Server, on its start-up, locks the access control data of all the slave systems it manages.

Even though the AC Server may be run on any of already existing PVSS systems, it is a common practice to configure it on a dedicated system as a part of the supporting infrastructure. As such, it would act as the access control management console, often coupled with user and role provisioning system described below.

The AC Server may also act as a proxy server for authentication. In such a setup, the client machines send the credentials-verification request to the server, which in turn may perform e.g. the LDAP based credentials checking, and return the result to the clients. This allows, for instance, for an easier configuration of firewalls, as only the AC Server needs to contact the authentication server and not all the clients.

The management of hundreds of user accounts and roles for systems as large as those of the LHC experiments' control systems is undoubtedly a challenge. Even though the administration and management tools provided by the Access Control Component made this task more feasible, the use of dedicated identity management tools (like Oracle Internet Directory, Microsoft Active Directory, or OpenLDAP) for such purpose seemed to be more adequate. Dedicated user account and role provisioning mechanisms have been developed for all the requested cases; this allows for automated creation of PVSS accounts and synchronization of role-privilege mapping. In each case, the provisioning mechanism was implemented on the PVSS side: the data was queried from the identity management server using the LDAP protocol, then used to (re-)create and (re-)configure PVSS access control on the AC Server, which in turn replicated the complete data into all managed systems.

SUMMARY

The presented Access Control component has been in production for a few years now, notably prior to the LHC start-up. Deployed in over 600 controls PCs, including the ones of the 4 large LHC experiments' (Detector Control Systems, Magnet Control Systems, Gas Control Sytems, Detector Safety Systems) and many accelerator control systems (LHC cryogenics, Quench Protection System, Power Converters), it fulfils the requirements for performance and scalability, and has shown very good flexibility and extendibility. The standardized RBAC model, and the abstraction of domains and privileges make the implementation and the management of access rights straightforward. The use of the LDAP standard provided a platform-independent implementation (OS, being Linux or Windows, and also for Identity Management products being the source for provisioning). This allowed for very good integration with CERN central services (i.e. common credentials, user account and role provisioning).

REFERENCES

[1] PVSS: Process Control and Visualization SCADA, http://www.pvss.com.

[2] The JCOP Framework, http://cern.ch/en-dep-ice-scd/scd/Projects/Framework.

[3] "The JCOP Framework", O.Holme, M. Gonzalez-Berges, P. Golonka, S, Schmeling, ICALEPCS, Geneva, October 2005.

[4] "The NIST Model for Role-Based Access Control: Toward. a Unified Standard", Sandhu, R, Ferraiolo, D.F. and Kuhn, D.R., 5th ACM Workshop Role-Based Access Control.

[5] "Role-Based Access Control for the Accelerator Control Systems at CERN", S Gysin et al, ICALEPCS07, Knoxville, USA, 2007.

[6] The Lightweight Directory Access Protocol, RFC 4510.