

Identity Management

Alberto Pace

CERN, Geneva, Switzerland

Abstract. This paper introduces identity management concepts and discusses various issues associated with its implementation. It will try to highlight technical, legal, and social aspects that must be foreseen when defining the numerous processes that an identity management infrastructure must support. Grid interoperability as well as cross platform interoperability is addressed on the technical aspect, followed by a short discussion on social consequences on user's privacy when completed traceability is enforced and some examples on how identity management has been implemented at CERN

1. Introduction

In recent years, the amount of bugs and vulnerabilities found in a large fraction of computing systems has created economic opportunities for a new branch of computer science and industry: computer security.

This has grasped the focus of the entire computing industry. Software vendors have developed new technologies to maintain and patch remotely across the internet their own software. Hardware manufacturers have multiplied the offers of boxes to strengthen the security of existing networked infrastructures (firewalls, enterprise routers, proxy servers, demilitarized zones, ...) with new classes of secure products. New opportunities were created to produce antivirus solutions for the desktops or anti spam products for the electronic mail services. Finally, IT professionals have been buying massively these new products and new job positions to handle "computer security" were opened in many companies.

This tendency has generated the incorrect belief that buying and deploying adequate security products, systematically applying the necessary patches, installing an antivirus solution, monitoring network traffic and scanning computers for vulnerabilities would be enough to secure an existing computing infrastructure to an acceptable level.

Although this strategy remains necessary, it is unfortunately no longer sufficient to improve significantly the computing security. Two additional aspects must be integrated to achieve an efficient overall secure strategy: The first is the human factor called "Social Engineering", the second is the management of the identities of all the people who have access to the corporate computing services.

The human factor is at the origin of the simplest intrusions and can be resumed to "ask the password or valid access credentials to the user himself". All these methods tries to fool the end-user either using phishing emails, malicious web sites or traditional communication techniques such as the telephone, a normal letter or even a personal visit to the victim's office.

To limit the risk factors connected to social engineering attacks, technical and behavioral training must be given to every person that needs to manipulate confidential information. This should be analyzed in a frame where also the motivation, the efficiency, and the productivity are assessed for employees that work in a highly secured environment. Unfortunately the solutions and the strategies to tackle social engineering issues go beyond the scope of this paper and will not be discussed further.

The management of the identities and of access rights is the second essential component of the enterprise strategy for a secure computing and will be discussed throughout this document.

2. Identity and Access Management (IAM)

Identity Management (IM) is the information and the set of flows which are sufficient, in legal terms, to identify the persons who access an information system. This includes all data on the persons, the workflows to create, read and modify this information, all internal process and procedures and the tools used for this purpose.

Identity Management must be associated with the concept of “Access Management (AM)” to become “Identity and Access Management (IAM)”. Access Management is the information describing what end-user can do on the corporate computing resources. It is the association of a *right* (use, read, modify, delete, open, execute, ...), a *subject* (person, account, computer, group, ...) and a *resource* (file, computer, printer, room, information system, ...).

Table 1. Example of Access management entries.

Resource	Right	Subject
cn=31-R072, ou=printers	Print	cn=Michel Blanc, o=cern, c=ch
cn=documents, ou=directories	Read	cn=IT department, o=cern, c=ch
cn=documents, ou=directories	Modify	cn=Mario Rossi, o=cern, c=ch
cn=pcit34, ou=computers	Logon	cn=John Smith, o=cern, c=ch
cn=D2342, ou=rooms	Enter	cn=Sys Admin team, o=cern, c=ch

Note that the association can be time-dependent, or location-dependent and resources can be computing resource (an application, a table in a database, a file, ...) but also physical entities (a room, a door, a terminal, ...). For example, a set of files in a folder can be readable during normal working hours and read protected during nights and weekends. Similarly, the right to open a door may depend from which side of the door the request came.

Table 1 shows few examples of entries in the access management database. These are simple logical examples: the practical implementation is typically platform specific as well as the vocabulary that changes between various implementations.

Note that an authorization may be required even for an action that does not alter the content of a resource, like a “read” permission.

3. The AAA rule: Authentication, Authorization and Accounting

From the definition of Identity and Access management, the implementation requires an architecture based on three independent components.

The first is the “*Authentication*” service which ensures the unequivocal identification of the person who is connected to the information system. The identification must be pointing to the “Identity

management” database where all potential users of the information system have been identified. This means that every process, every modification, every action done in the information system has a “label” with the name of the person who is behind the action requested.

There are several technologies to implement an authentication service. The most traditional are based on “username / password” pair where the password has been secretly communicated to the user during the initial identity validation process. But there are additional methods to obtain authentication services using electronic certificates [1], smartcards, hardware tokens or biometric recognition. In addition to the way the authentication secret is shared between the authentication authority and the agent there are multiple cryptosystems available to implement a secure authentication service. The main ones used in High Energy Physics are “Public Key Infrastructure (PKI)” [2] used in GRID infrastructures and “Kerberos” [3] used in the UNIX/Linux based AFS global file system and in Microsoft Windows based domains. In all cases there are several techniques to interoperate authentication services based on either system: For example, in a Microsoft Windows domain supports natively logon using PKI certificates, PKI based smartcards, Tokens and biometric recognition despite all the inner layer of Windows authentication are based on Kerberos [4].

The second component is the “*Authorization*” which ensures that the authenticated person has the permission to carry out a particular action on a resource. Authorization can be enforced by the operating system as a part of its native functionality like protecting a file in AFS or DSF using Access Control Lists (ACL) [5] or can be implemented by the application by a lookup of the identity validated by the authentication service into the databases of the authorizations for the resource being accessed.

Clearly the authorization component is the most complex and may require multiple implementations for multiple platforms. For each platform you can rely on the native authorization mechanism or develop your own implementation for your corporate applications. If you chose the first approach ensure you have accounted for the effort to integrate the different authentication technologies that becomes necessary when you use native or commercial implementations across multiple platforms. If you chose the second approach, remember that building a custom secure authorization mechanism may be extremely expensive.

The third and last component of an IAM architecture is the “*Accounting*” component which ensures the traceability of all actions made on the information system. It is the journal of all operations (who, when, where, what) that have been made and it gives all details on what happened in the past. In many cases, the accounting journal allows to *rollback* many transactions and therefore it can be used to recover voluntary or involuntary acts of data corruption or sabotage.

The role of the accounting component is often underestimated, despite it is a tool that allows empowering employees and boosting their productivity by granting them more authorizations than necessary, knowing that actions are accounted. This is a different approach from minimizing the “attack surface” or implementing the principle of “least privilege” where each subject is granted only the most restrictive authorizations for the performance of tasks. Therefore there is a trade that can be made between authorizations and accounting which may increase the employees’ motivation.

This applies particularly to High Energy Physics laboratories like CERN which has a long tradition of openness. For example, at CERN all users are empowered to publish on the web without requiring administrative authorization: this is a simplification for end-users possible because there is an efficient authentication/accounting mechanism behind the scenes.

Of course be aware that accounting cannot rollback all your transactions: you cannot reverse the inappropriate disclosure of confidential information and there are plenty of transactions in the

information system that have an impact in the real world (payments, for example) and cannot be rolled back. Therefore the trade between accounting and authorization must be carefully designed in parallel with an adequate risk analysis.

4. More on Identity and Access Management Components

After discussing the AAA rule, it is important to stress the fact that the Authorization component is platform and application specific. This generates difficulties in centralizing the management of Authorization.

The workaround to the dispersion of the authorization database is the implementation of “Role Based Access Control (RBAC)” which grants permissions (authorizations) to groups of subjects instead of individual subjects (person).

With this approach authorizations are managed by defining membership to groups and therefore it separates the authorization into two separate functions:

- The action of granting the authorization to a group (Role creation)
- The action of managing the group membership (Role assignment)

The “Role Creation” remains platform and application specific, but is done only once, by the expert on the resource protection. Once this has been done, the “Role assignment” i.e. managing the membership to the authorized group is a generic task which is no longer platform specific and can be delegated to administrative personnel or even automated when the group membership can be calculated from the Identity Management database.

Be aware that RBAC should remain a simplification and therefore it is important to limit the number of roles to a minimum

5. Legal Consideration

When implementing IAM, it is essential to be aware of legal constraints. The general legal frames for the implementation of IAM are the “Sarbanes Oxley Act (SOX)” [6] in the US and the “8th EU Privacy Directive” [7] in Europe which is completed by additional national or local laws.

The legal obligations may be very complex to fulfill. Laws are different in each country and it may be difficult to have a unique infrastructure for a large corporation with worldwide presence. Laws depend on the type of institute and can impose different obligation among public funded, government, privately owned, or international organizations.

Laws also depend on the sector of activity. In several areas of the economy (banking or telecom operators, for example) there is a legal obligation infrastructure in order to protect customers to have a solid IAM which imposes archiving, traceability, retention of log files and evidences. Often, the obligation of traceability may appear in contradiction with similar obligations to respect the employees or customer’s privacy. This makes difficult to find the good compromise between security / accounting / traceability and respect of privacy / personal life.

6. Why IAM ?

The motivations to deploy IAM can be very different. Cost reduction could be imagined when multiple accounts and multiple incompatible authentication mechanism could be replaced by a single one.

Another motivation is to offload IT experts from administrative tasks that have little added value, like user registration, password changes, granting permissions. This can be coupled to a technical opportunity to simplify procedures

Finally, and probably the most important reason, is the expected increased security where a global overview of the authorizations/authentication/accounting/roles is available. The centralized security policy can make possible to provide status boards of the corporate system security.

7. Architecture components of IAM

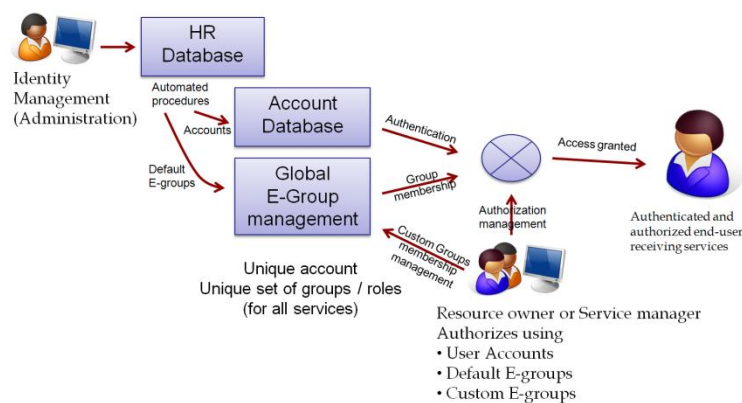


Figure 1. Components of an identity and Access Management architecture

7.1. The Identity management database

The starting point of any IAM architecture is the identity management database. This contains the information on the all persons who have access to the information system and it goes beyond the list of employees. For CERN, that has less than 3000 employees, the identity management database contains more than one order of magnitude more records, which represent the complete list of person who potentially need to access the CERN IT services and includes persons working for contractors, visitors, students, former member of personnel (including retired members), family members, ...

A database must be *accessible*. This means that an authoring portal is necessary to maintain the information in the IM database and it is used by the administration to create identities. This includes registering new persons but also modifying existing content. The typical implementation makes the authoring portal available through a web page but nothing prevents to have feature-rich, platform-specific, clients that offer an improved productivity to the staff in charge of maintaining the information in the central database.

The authoring portal has clearly the interfaces to allow administrative personnel to register and maintain identities. However there are several personal attributes for which the permission to modify could be granted to the person himself in view of reducing the administrative overhead. For example, every member of the IM database could be granted the permission to modify his/her own phone number or his/her preferred language.

In this case, the authoring portal should foresee interfaces to both end-users and administrative personnel, with approval, workflow and information validation depending on the type of data. This is a good example of trade between authorization and accounting: the principle of “least privilege” would suggest that only administrative personnel is allowed to modify IM data, but given the excellent auditing functionalities of most common databases, end-users can be allowed to modify themselves non-critical information of their own data which can be easily rolled back in case of abuse.

The public content of the database must be also accessible. Therefore the directory services for all platforms should be fed with the information from the IM database. This also provides phone book and ldap services.

7.2. *The account database*

Once the IM database is in place and you know all the persons, you need to create the computer accounts in the information system. At this point it is essential to understand what are the “administrative” requirements to “be known” by the information system. “Administrative” means that you have all information in the IAM database and remember that having an account should not grant the user any implicit permission (it is not an authorization process). For this reason, the easiest approach is to create accounts for the *entire* population you have in the IM database.

Therefore you can create a process with a well defined workflow to create automatically all accounts necessary to the person in the IM database to access the service. Obtaining the initial password / certificate / smartcard is done with the help of the administrative staff in charge of the identity registration or via the corporate helpdesk which has clear procedures to re-validate user’s identities.

Note that even in this case you need to answer some critical questions. The most difficult one is probably to decide whether you allow one person to have multiple accounts. Allowing users to have multiple accounts has the advantage of allowing users to play “multiple roles” and therefore using privileged identities only when necessary. This has the disadvantage that users could share credentials of unused additional accounts with other persons and therefore break the whole Authentication / identity management goal.

In general, it is believed that the unique account per person approach is more secure. In the implementation at CERN the identity management software is planned to create only one account per person (called the “primary” account) but will allow the user to create additional accounts himself (“Secondary” accounts) [8]. However, despite the user may have multiple accounts at CERN, several central IT services accept credentials coming only from primary accounts (for example the electronic mail service or a login to the administrative services) which is considered more secure.

The need for having multiple accounts may be reduced by granting “privileges on demand” to accounts. This means that the account is “potentially” authorized to carry out a sensitive operation but must explicitly request this right in order to execute the operation. The example of this functionality is the latest implementation of “User Account Control” available in Mac OS and Windows Vista which do no longer require to have multiple accounts (one without privileges and another with root/admin privileges) on the same computer as one single account can impersonate, on demand, the two roles.

7.3. *Authentication services*

Once the IAM database is in place together with a process to create and manage the lifetime of accounts, the authentication service can be offered.

The ideal solution is to rely on a “Single Sign on (SSO)” service used by the entire information system. The user is authenticated at the beginning of his session and then he can seamlessly access all authorized resources. Note that having a single sign on service does *not* prevent an application to request a new authentication when accessing a particularly confidential document, when signing a document or when elevating the privilege in the case of an “on demand privilege”.

For practical reasons, in some cases, it may be easier for the service managers, to have multiple independent authentication services (MAS) for the different platforms, all connected to the same IAM and account database.

Although SSO or MAS are identical in terms of security, the multiple independent authentication services has the disadvantage of forcing the user to re-authenticate whenever he/she cross the border between system managed by different authentication systems: this mean that the multiple authentication system is simpler for the service manager to implement but more complex for the personnel to use.

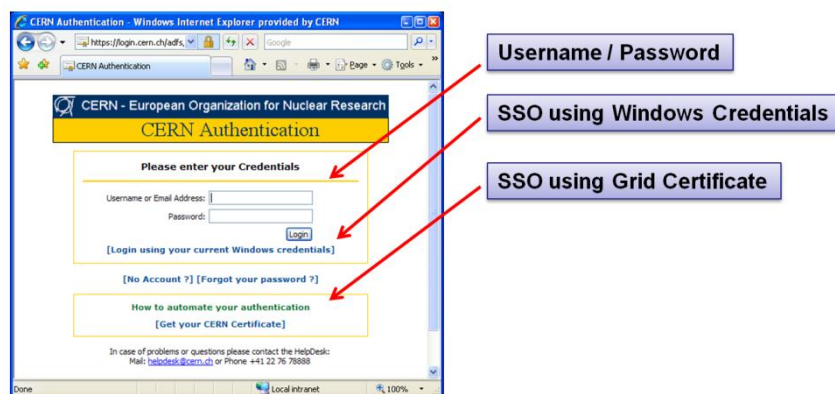


Figure 2. The Single Sign On logon dialog at CERN [9]

7.4. Managing authorizations

As already mentioned, interfaces to manage authorization are service or platform specific: Granting read permission to a file stored in an NTFS file system is done using a Windows-based user interface that is radically different from the equivalent interface on Linux to manage permission in an AFS-based directory.

Although some development effort can be invested to overcome these incompatibilities by writing web portals to manage authorizations, it can be simpler to use RBAC and recommend assignment of permissions to groups and then manage the authorization indirectly by controlling the group membership.

For this reason, an application to manage group memberships is required as an indirect way to manage authorizations. The group membership information should be integrated in the IAM database.

Two types of groups must be foreseen: groups where the membership is manually managed and groups where the membership is generated from arbitrary SQL queries in the IAM database (like “Members of the IT department”). An additional functionality that must be foreseen is the possibility to nest groups and therefore obtain mix behaviors of groups where some members are calculated from the IAM database and other have been added manually.

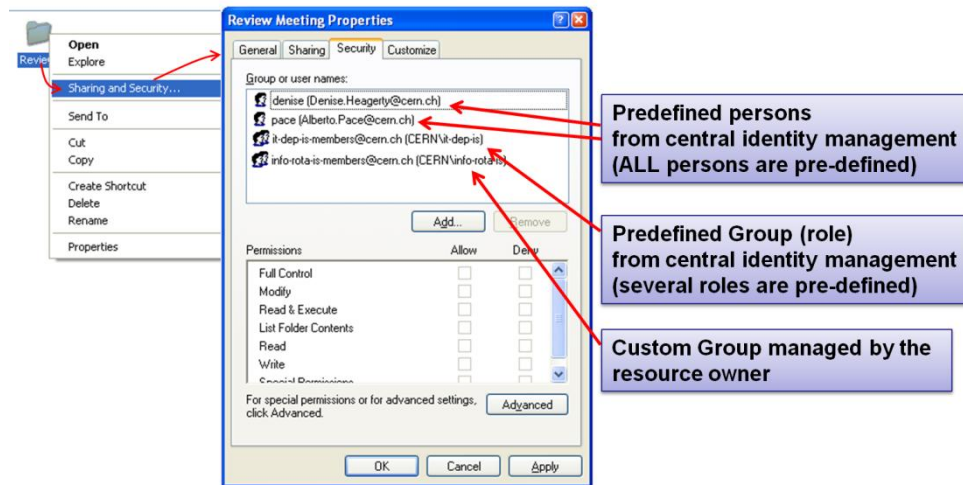


Figure 3. Example of authorizations: an “Access Control List” set on an NFS folder

8. Plans at CERN

CERN has a Human Resources (HR) database with many records (persons) who can have 23 possible status (Staff member of personnel, fellow, student, scientific associate, enterprise, external, ...) and there are heavy rules and procedures to obtain accounts to be used in the central IT infrastructure. These accounts are manually created by so called “Group Administrators” using a desktop application that is linked to the HR database.

Traditionally CERN has allowed to have multiple independent accounts across multiple services (Administrative services, Document Server, EDMS, Indico, LanDB, Mail, Oracle, Remedy, Unix/AFS, VPN, Web, Windows) as well as multiple accounts per person. This strategy is being changed towards a unique identity management system with one unique account for all services. At the time of writing 10 of these services (Administrative services, Document Server, EDMS, Indico, LanDB, Mail, Remedy, VPN, Web, Windows) have already converged to use the “CERN account”. Multiple accounts per persons are still possible but several service are only accepting credentials coming from the “Primary” account, de facto implementing the unique account per person.

The use of roles and groups to indirectly manage authorizations has been achieved by extending the existing portal to create and edit “mailing list” into the portal for to manage group membership. Basically, all the existing mailing lists have been replicated as security groups that can be used to grant/revoke permissions and authorizations. The current mailing list editor allows having mailing list with membership manually managed or dynamically calculated from the HR database and it allows nesting of mailing lists [10].

Given the important role the mailing list editor used to control authorization, it is being rewritten as an E-groups portal and a “Mailing List” is planned to become a simple E-group where the “Mail enabled” attribute has been checked.

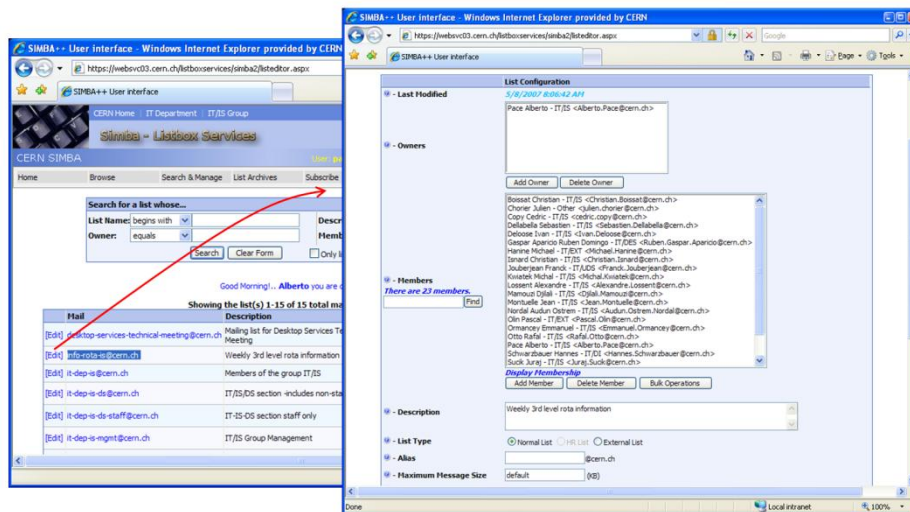


Figure 4. The portal to manage membership of mailing lists at CERN used to grant authorizations indirectly [10]

9. Integrating the big picture

A global identity management service is a strong requirement for High Energy Physics computing and Grid activities and this requirement integrates smoothly with existing IAM services available in various laboratories.

The Grid global Identity management initiative coordinated by the “International Grid trust federation (IGTF)” [11] is based on Public Key Infrastructure authentication services which allow a truly distributed set of certification authorities across the world. The coordination is done through the regional Policy Management Authorities: The Asia Pacific Grid PMA (APGRIDPMA), European Grid PMA (EUGRIDPMA) and The Americas Grid PMA (TAGPMA) [12]. The membership is composed by more than 50 accredited Certification Authorities (CA) which provides a global identity management infrastructure which covers all needs for High Energy physics, worldwide.

The CERN “local” identity management infrastructure service is directly reused within the “global” one: As the CERN IAM offer an authentication service for end-user, this service has been used by the CERN Certification Authority (CA) to issue valid Grid certificates to CERN grid users. The advantage of this approach has been that the certificate authority is online and issuing / revoking / renewing a certificate can be done online from anywhere on the internet in a matter of minutes. This was considered a strong advantage compared to an offline CA where the user had to physically move to an office at CERN to have his identity validated before obtaining a valid grid certificate. In addition, substantial manpower savings have been possible when moving from an infrastructure requiring more than one person for identity validation to a solution that is completely automated [13].



Figure 5. The online portal to request Grid certificates for CERN users

The opposite integration has also been straightforward: CERN recognizes the certificates issued by accredited by IGTF and allows an automatic mapping of these certificates to local CERN accounts for the people whose identities have been validated at CERN and are member of the IAM database.

10. Conclusion

When the number of subjects allowed to access an information system exceeds few hundreds, Identity and Access Management becomes an essential component of a secure computing infrastructure. It may take lot of effort to convert a large multitude of small home-grown authentication clusters, but when it has been done, it represents an important simplification for both end-users and system administrators who can have a global overview of the authorizations to access their systems.

References

- [1] See http://en.wikipedia.org/wiki/Digital_certificate
- [2] See www.pkiforum.org
- [3] See <http://web.mit.edu/kerberos> and <http://www.kerberos.org>
- [4] For the Microsoft Windows implementation of PKI see <http://www.microsoft.com/pki/> , for the implementation of Kerberos see: <http://www.microsoft.com/kerberos/>
- [5] For the native implementation of ACLs on AFS, see <http://www.openafs.org/pages/doc/UserGuide/auusg007.htm>. For the Windows implementation see <http://www.windowsitpro.com/Articles/ArticleID/8452/8452.html>
- [6] Available at http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=107_cong_bills&docid=f:h3763enr.tst.pdf
- [7] Directive 2002/58/EC, OJ L 201 of 31 July 2002, p. 37 (Available at http://europa.eu.int/information_society/topics/telecoms/regulatory/new_rf/index_en.htm)
- [8] The portal for account management at CERN is at <http://cra.cern.ch>
- [9] E. Ormancey, "CERN Single Sign On solution", these proceedings.
- [10] See www.cern.ch/simba for an overview of the mailing list service at CERN
- [11] See www.itgf.org
- [12] The IGTF PMAs are online at www.apgridpma.org, www.eugridpma.org, and www.tagpma.org
- [13] The CERN online certification authority is available at www.cern.ch/ca