

Distributed Access Control

- BIBSYS and the FEIDE-solution

Sigbjørn Holmslet, BIBSYS

Ingrid Melve, UNINET

June 2004

Introduction

This paper will discuss the challenges BIBSYS meets regarding access control in a distributed environment, and explain why and how BIBSYS intend to make use of the FEIDE (Federated Electronic Identity for Education) solution [FEIDE], and give a brief overview of what FEIDE is and how it works.

Distributed access control in general

This chapter will briefly discuss general issues concerning distributed access control. First we will define some important terms used throughout this document.

Authentication - Process of providing the identity of a previously registered user (Who are you?)

Authorization - Process of granting or denying access rights for a resource to an authenticated user (What are you allowed to do?)

Credentials - Information that includes identification and proof of identification that is used to gain access to resources. Examples of credentials are user names and passwords, smart cards, and certificates.

SSO (Single Sign On) - A system that enables a user to access multiple computer platforms or application systems after being authenticated only once.

The typical situation for a user accessing distributed services today is shown in Figure 1 [AAI].

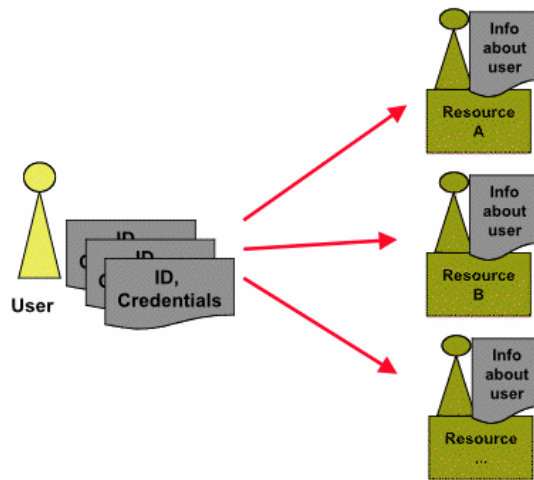


Figure 1 - Access to networked services without a distributed access control system.

The user has several username and password pairs (credentials) for each of the resources he wants to access. The problem is that the user has to deal with lots of credentials and lots of different registration procedures. A way to solve the problems is shown in Figure 2.

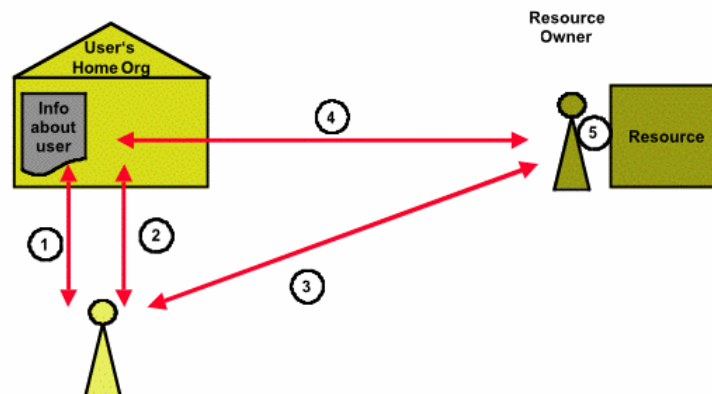


Figure 2 - Access to networked services with a distributed infrastructure for authentication and authorization.

1. Registration at the user's home organization; information about a user is stored and updated by the home organization. This is usually done once per user, e.g. at matriculation time in the case of a student
2. Authentication by the user's home organization
3. Access request
4. Delivery of additional information about the user to the resource upon resource's request and user's consent
5. Authorization based on the retrieved information about the user

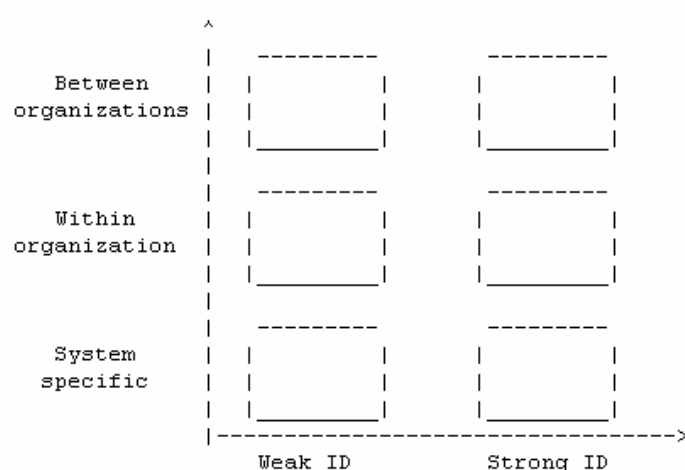


Figure 3 - The trend is going towards a higher security level and more cross-organizational access control

Figure 3 illustrates two important factors in distributed access control. The horizontal axis shows the level of security and the vertical axis the extent of interoperability.

Security levels may vary, depending on the needs of the service. Read- and write access are examples of different needs. For some applications qualified digital signatures are required, others may never need any high level of security. Anyway, the common distributed access control system must always satisfy those with the highest demands.

When it comes to interoperability, Single Sign On (SSO) is an important subject. But SSO also raises new difficult technical and security issues. How do we solve it on the practical level? A proxy with one entry to several target recourses is one possibility but have many drawbacks concerning performance and flexibility. You have to consider whether to use cookies or not, decide timeout values, how to handle sessions and so on. And at last but not least - SSO also means that both users and services may have to share and reuse credentials. How can you handle that in a secure way?

Facts about BIBSYS and the users

BIBSYS is an integrated library system used by all Norwegian University Libraries, the National Library, all college libraries, and a number of research libraries We define the users of the BIBSYS-system in to categories.

The primary users

Consists of ca. 2500 librarians

The end-users

- Ca. 600.000 - patrons (not all of them are active users)
- Ca. 4000 - academic users for our research database

- 1000+ - users of different smaller systems

History of access control in BIBSYS

Until the early nineties BIBSYS had only one main centralized application (the legacy system). The application was based on telnet and used the underlying UNIX access control. The users were librarians, and unidentified patrons who used a public terminal in the library to search the base.

In the middle nineties BIBSYS developed a web-based interface for search and for online loan and copy ordering. To make it easy for the patrons the authentication was very weak. Patrons just had to state patron ID and last name in order to get the request fulfilled. But the weak authentication made it impossible to provide services showing personal and sensitive information, e.g. a list of the patron's standing orders.

We also offered a web-application for search in the ISI database, which required authentication. To solve the problem we used IP source address filtering [Lynch98]. We maintained a list of IP addresses with authorized access, and used standard mechanisms in the Apache web-server to implement the filtering.

In addition we started to release new smaller systems with their own access control.

At this point we had two major problems:

- 1) We were offering various systems and services with their own independent access control solutions.
- 2) Our end-users needed personalized services. We only offered personal authentication (username/password) for the librarians. A similar solution for our end users would raise significant new demands regarding user administration, both technical and organisational.

We started an access control project and tried to solve the problems listed above in the following way:

- 1) We gathered *all* access related information for *all* of our users (librarians, patrons, institutions, employees, etc.) in to one centralized role based access control system. The new system is accessible from all BIBSYS applications, not only the legacy system.
- 2) To avoid user administration, we use the patrons email address, which we already have, and on demand send a machine generated password to the patron.

In 2001 BIBSYS released a new web-based application with personalized services for both librarians and patrons. The authentication was based on the HTTP Basic protocol. Still some of the old systems use their old access control, so the migration is not 100% completed.

In the process we also evaluated two commercial access control systems - ISOS/Athens [Athens] and Candle/Cactus [Candle2000]. Economy was the main reason we didn't choose either of them. They were too expensive and we felt they were targeted against bigger cross-organizational solutions, maybe at a national or international level.

The BIBSYS challenges of today

BIBSYS still have unsolved challenges. In the daily life most of our users use online services provided by other actors than BIBSYS. This may be services closely related to BIBSYS, e.g. scientific reference databases, electronic journals, educational portals/systems, etc. The need of interoperability is obvious. In order to achieve interoperability between BIBSYS and the other systems access control is important. There we have two major goals:

- 1) *One username/password.* When a user logs on a BIBSYS service, he can use the same username and password as elsewhere.
- 2) *Single sign on.* If the user is logged on another system, he is automatically authenticated and authorized when entering a BIBSYS service.

Since we have no intention of developing or hosting a cross-organizational access control system our selves, we had to look elsewhere. Fortunately there were an ongoing project called FEIDE (Federated Electronic Identity for Education), which addressed these problems. Since most of the BIBSYS end-users are in the educational sector, which FEIDE covers, the solution was very interesting for BIBSYS.

Some facts about FEIDE

FEIDE (Federated Electronic Identity for Education) is a project with the goal of establishing a common, secure electronic identity for Norwegian academic users. FEIDE is implementing the academic sector's system for reliable user data handling, secure identification of internet-service users and assignment of user access-rights. This poses requirements on the administration procedures of students and employees.

Common data model

FEIDE has defined a common data model for information about persons, including the user information. Based on this, universities grant a common integrated well-defined information structure about all their users, accessed through FEIDE.

User management systems

The user management system is in most cases a meta directory solution, three software solutions (Cerebrum, Dir-XML+eDirectory, MIIS+AD) are currently used, they supply the authentication service with the information required to check the identity of the users.

Central log-in server

FEIDE operates a central log-in server, running Moria. A central service is provided for the convenience of service providers, assuring high availability for the login part of the service. FEIDE was initially built for Initial Sign-On, not for Single Sign On, but this is now being changed, as the demand for an integrated cross-institutional SSO-solution is rising. SSO brings in a wide variety of security considerations, because SSO requires sharing a (authentication) secret.

For more information about FEIDE we highly recommend the following URLs:

<http://www.feide.no/dokumenter/arkitektur.html> [FEIDEarc] and

http://www.feide.no/programvare/authentication_model.html [MORIA].

The practical experience with FEIDE

In April 2003 BIBSYS tested a pilot that made use of the FEIDE system. We used our new system that offers personalized services for patrons and librarians, and made the pilot available for a limited group of users. The objective was to try out the technical issues. From FEIDE we got a Java-library, a certificate and a Java servlet filter. All we had to do was to configure our Tomcat-server so that all requests went through the filter, and configure the filter. Everything worked out fine; neither BIBSYS nor the users experienced any problems using the FEIDE system. I should be mentioned we did not test performance or scalability.

The biggest obstacle although, was not on the technical side, but on the modelling and organizational level. The problem is: how can we map a FEIDE-user to a BIBSYS-user? The FEIDE username is not in our database. Some months ago BIBSYS landed on the following solution: FEIDE has an internal key to uniquely identify a person. This key is not visible for the user and is identical with the National Identity Number, a number used by all public services and many private corporations to identify persons. In these days BIBSYS is extending the user database to include National Identity Number.

When a BIBSYS service use the FEIDE system the logon process looks briefly like the illustration in figure 4.




Figure 4 - A simplified overview of the logon process, before it reaches the BIBSYS service.

1. The user makes a request to a BIBSYS service.
2. The filter requests an authentication session from Moria
3. The filter redirect the browser to Moria
4. The user state username and password
5. Moria checks username and password against the AT
6. If ok, Moria redirect the browser to the filter.
7. The filter retrieves user attributes, like the National Identity Number.
8. The National Identity Number is passed on to the BIBSYS-service.
9. The BIBSYS-service uses the National Identity Number look up user data in our base.

- All communication with MORIA, both from the filter and the user, are encrypted with SSL (Secure Socket Layer).
- FEIDE implements the authentication, and BIBSYS is responsible for authorization (step 9)
- In the authorization-process BIBSYS also can make use of user attributes retrieved from Moria containing role information.

The plans are that we within the next 3-4 months will release our first application using the FEIDE solution. The next step is to try out the SSO features FEIDE offers. We also intend to make use of user attributes from FEIDE in the authorization process.

References

[AAI] Switch “Authentication and Authorization Infrastructure” <http://www.switch.ch/aai/>

[Athens] “Eduserv Athens for education” <http://www.athens.ac.uk/>

[Candle2000] C. Farmakis, D. Martakos, Andrew Cox “Managing Networked Information Services: the CANDLE Case.” Russian Digital Libraries Journal - 2000 - Vol 3 - Issue 3, <http://www.iis.ru/el-bib/2000/200003/FCM/fcm.en.html>

[FEIDE] The FEIDE project, <http://www.feide.no/>

[FEIDEarc] FEIDE System Architecture, <http://www.feide.no/dokumenter/arkitektur.html>

[Lynch98] Clifford Lynch “A White Paper on Authentication and Access Management Issues in Cross-organizational Use of Networked Information Resources” – 1998, <http://www.cni.org/projects/authentication/authentication-wp.html>

[MORIA] “Moria - FEIDE's HTTP Authentication Service” http://www.feide.no/programvare/authentication_model.html