ZESZYTY NAUKOWE WYDZIAŁU ETI POLITECHNIKI GDAŃSKIEJ
Nr 18 Seria: Technologie Informacyjne 2010

Henryk Krawczyk[1], Paweł Lubomski[2]

[1]Gdansk University of Technology, Department of Computer Architecture
[2]Gdansk University of Technology, Information Technology Centre

GENERALIZED ACCESS CONTROL IN HIERARCHICAL COMPUTER NETWORK

Abstract

The paper presents the design of the security layer for a distributed system located in the multizone hierarchical computer network. Depending on the zone from which a client's request comes to the system and the type of the request, it will be either authorized or rejected. There is one common layer for the access to all the business services and interactions between them. Unlike the commonly used RBAC model, this system enforces a multilayer authentication and authorization. Actor's privileges are the result of the user's and the system's roles conjunction with the network zone. Unlike common systems, the privileges are given to a digital identity, not to particular accounts, so that it does not matter which account was used by the user – he will get the same privileges. Such a combination of many smaller ideas and methods results in a new and modern approach to the security aspects of the distributed service oriented systems.

1. INTRODUCTION

Modern distributed systems and platforms need to take exceptional care of the security and safety aspects. Current solutions are not sufficient. There is a great number of verified approaches but they work in some consolidated centralized systems. Information technology evolves to interoperability of e-services' based systems [1]. They are distributed through diverse extensive and publicly available networks. Securing them is a very ambitious issue that we have to face [2].

In the paper the authors present a solution developed, tested and introduced in a new distributed platform of e-services in Gdansk University of Technology. It is not the only correct approach to this issue – it should be treated as one of many possible solutions.

2. RBAC

The Role Based Access Control (RBAC) model is commonly used in many applications [3]. It is also a part of the Java Enterprise Edition (JEE 5) standard [4]. The standard introduces two approaches to this issue: a declarative and a dynamic implementation model. Regardless of the implementation model, it is typical of RBAC that the permissions to perform certain operations are assigned to some specific roles. Arranging the roles into groups of roles is a very common approach. Every user is assigned to some of these groups of roles or the roles directly. As a result, the users are not arranged by their positions but by their competences and responsibilities. Such approach is more flexible in administration, especially when one person replaces another one, e.g. during one's leave.

RBAC works in the consolidated applications but it has also one weakness – in the distributed systems it needs to be extended to face some very complex authorization cases. The following chapters describe some ways of overcoming this problem.

3. HIERACHICAL MULTIZONE COMPUTER NETWORK

First of all, we need to divide the computer network, in which the system works, into some zones. It is possible to do that on two levels: Virtual LAN (VLAN) or physical network. The first solution requires an appropriate configuration of the network routers. The second one divides network into some separate subnetworks which cannot communicate with one another. We will not concentrate on this solution, even though it is more secure, but useless in our situation.

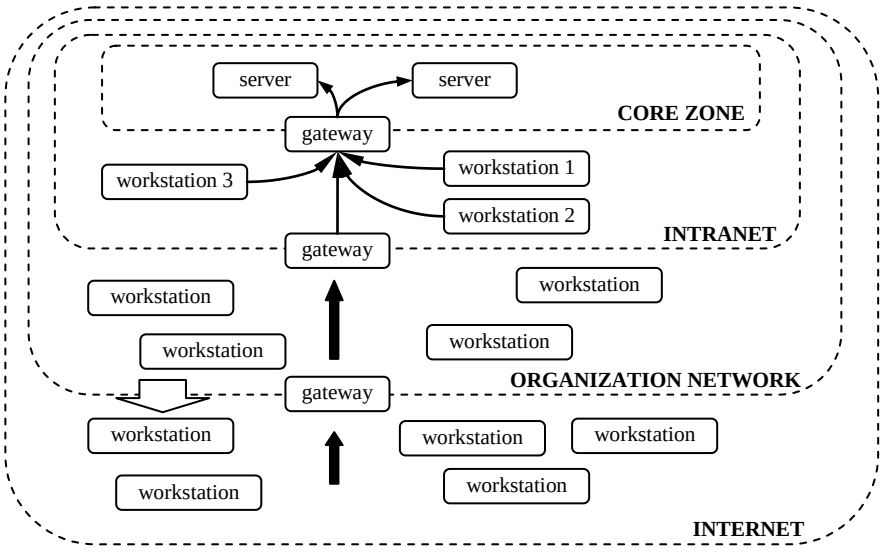

Fig. 1. Hierarchical multizone computer network

We arrange zones in the hierarchical order [5] as presented in Fig. 1. On the top there is the "demilitarized zone". We put there only our core servers. It also has only one gateway. In the next-level-zone there are some critical, well-known workstations placed. They are centrally administrated. We can extract more lower zones. They could include a higher zone, but there is one important rule: workstations from the higher zones are not visible to computers in any of the lower zones. Servers from the higher zones are only accessible via the gateways.

In the case of our university network we have distinguished four zones: the core zone, where the servers are put, the intranet, the organization network and the Internet. The intranet is a special logical zone where workstations are centrally managed and they do not have any access to the other zones – only to the servers via the gateway. These are mainly the university administration computers. We have also distinguished the organization network. For security reasons some actions can be done only from the workstations at the university, e.g. typing the students' grades.

4. AUTHENTICATION AND AUTHORIZATION

To face very complex authorization cases we developed a multilayer authentication and authorization. In the first layer the user's authentication is performed using the Single Sign On mechanism. In the second one the system (it may be a web interface or any other type of interface) authenticates itself to the business layer using the protocol based on the digital certificates. The consequential user permissions are a result of the user's roles and the system's roles.

4.1. User authentication

There is a growing number of systems that are accessible via a web browser. A person can perform any everyday operation in these systems. It is directly linked with a great number of accounts and passwords to remember. An obvious question arises: why do you have to log in to each of these systems when you go from one to another? To deal with these two issues we introduced Central Authentication Point (CAP) which implements the Single Sign On (SSO) mechanism [6]. Such an approach results in only one login and password to remember. It grants access to all the accessible to the user web systems and other services, e.g. a mail account. Another advantage is that once the user is logged in, he does not have to enter his login and password again while switching from one system to another.

4.2. System authentication

The second-level authentication is a system authentication. A client system– it may be a web system or any other kind of system – has to authenticate itself before performing any operation in the centralized business logic system located in a core of the earlier mentioned network. We developed this solution to improve the security of the whole distributed platform because of the possibility of spoofing or a man-in-the-middle attack. To improve the reliability of this authentication we use digital certificates assigned to the systems. This type of authentication has a high level of security. But the standard implementation introduces a high overhead of communication caused by the handshake protocol while making every connection. To deal with this issue we needed to introduce a solution based on the lower OSI network protocol model (transport layer) [7]. We launched an OpenVPN

(Open Virtual Private Network) server [8] on the gateway. The client systems, before accessing business logic systems, have to create a connection through this VPN. To connect to the VPN server they have to authenticate by using the protocol based on the digital certificate and then they are assigned a unique IP address (configured one per certificate), on the basis of which the system is recognised in the business logic security layer. The connection is established only once so the handshake is executed only once.

This solution has one more advantage. All the data exchange is being made through the security tunnel so the whole communication between the client and the business logic systems is encrypted. This solves another very important security problem in the distributed platform – the confidentiality of the transmitted data [9].

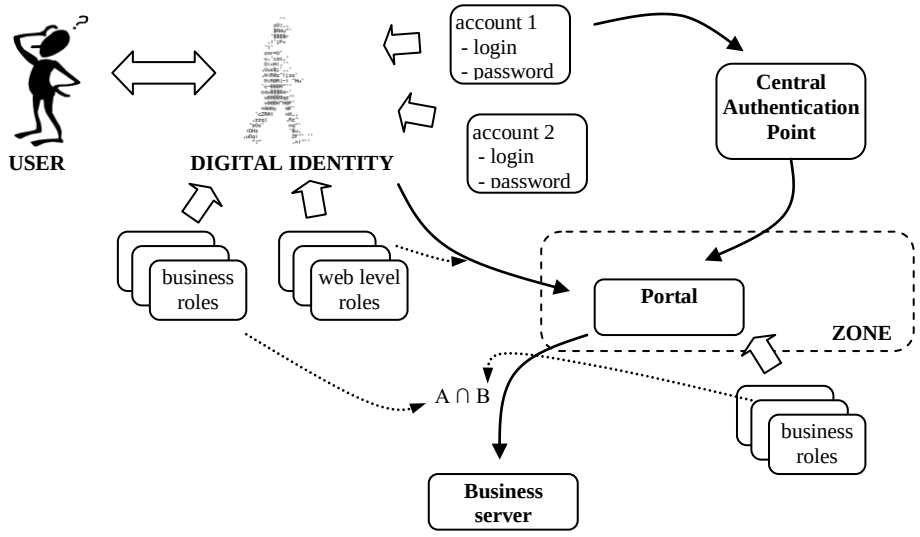

Fig.2. Digital identity and extended RBAC authorization

4.3. Authorization

The introduced authorization is an extension of the RBAC model. We use it in a more complex way – in conjunction with the network zones, authorizing the user and the client system separately, using another set of roles to organize the user's interface.

First of all, we divide the roles into business logic roles and web level roles. The first ones correspond to some small pieces of the business systems functionality. The other ones are closer to the user's position in the organisation hierarchy and his competences – on the basis of them a user can get an access to some larger parts of the interface or the whole client system.

The business roles are assigned to both: the users and the client systems. They are also arranged into some groups of roles for a more comfortable permission management. The users are assigned to the same groups of roles independently of the system they are using. The client systems– just the opposite. Each system has its own set of roles. The set varies

depending on the network zone which the system is located in. Setting the same system in another network zone will give it a different set of roles.

The resulting user permissions depend on his set of role and the system which he uses. They are determined as a common part of the set: the user's roles and the system's roles (see Fig. 2). Concluding: the same user can have different permissions in different systems or in the same systems located in different network zones. This mechanism is fully configurable. It is obvious that the earlier described distributed platform should be configured so as to assure that the higher network zone the client system and the user workstation are in, the more permissions they are assigned.

5. DIGITAL IDENTITY AND ACCOUNTS

There is one more issue to face. In some cases it is better to arrange different sets of accounts for different groups of users, e.g. mail accounts for students and employees. It is caused by assigning them different mail systems on different servers and different Internet domains. In most cases each user has only one account but in some cases, e.g. when a student works at the university during his studies, he has two different accounts – as a student and as an employee. How should the Central Authentication Point (CAP) behave? The answer is: no matter which account the user used to authenticate himself, he would be assigned the same set of roles. The system should treat him as one digital identity [10] with two accounts but only one set of roles assigned to it.

As presented in Fig. 2 we assign accounts, web level roles and business ones to the digital identity which corresponds to one real-world person. It may be a good idea that this identity has one unique ID which is used in every system to determine the author of the actions while assigning him the attributes or logging his activities. In our systems every entry in the database has some information about the date, the ID of the author and the client system he used while adding or modifying the entry. There is also a system log file where every activity is written down with the information about the time, the ID of the author and the client system he used.

Having one central repository of the digital identities, the user can activate one of the accounts that are available for him. It is more friendly for users and there is less work for the system administrators.

6. CONCLUSIONS

The integration of more and more interoperated systems leads to the need of a huge distributed platform and so far there is no solution available that secures it comprehensively.

The above described solution was thoroughly tested and introduced in practice at Gdansk University of Technology. We find it as a good approach to design a secure modern distributed platform of e-services [11, 12]. We have checked its performance, network isolation, and ease of use for user's and administrator's. The high configurability caused some performance problems at the beginning. But adopting some earlier mentioned solutions in connection with the appropriate caching has solved these problems – the authorization of any request takes no more than a few milliseconds (in most commonly used hardware), so it is a negligible cost. We also did not find such a complex authorization case which we cannot model with the presented approach.

The new strength of this solution lies in the combination of several known mechanisms and ideas. We obtain a synergies effect. But in security and safety issues work never ends – it should be developed consistently. Also, this solution needs further development, e.g. using of Zero-Knowledge Protocols. It is obvious that an appropriate information security policy should be designed, documented and introduced to keep this solution secure and safe.

BIBLIOGRAPHY

[1] Erl T.: *SOA Principles and Service Design*, SOA Systems Inc., 2008.
[2] Mather T., Kumaraswamy S., Latif S.: *Cloud Security and Privacy. An Enterprise Perspective on Risks and Compliance*, O'Reilly, 2009.
[3] Sandhu, R., Ferraiolo, D.F. and Kuhn, D.R.: *The NIST Model for Role-Based Access Control: Toward a Unified Standard*, 5th ACM Workshop Role-Based Access Control, July 2000, pp. 47–63.
[4] http://java.sun.com/javaee/technologies/javaee5.jsp
[5] Krawczyk H., Lubomski P.: *Architektury systemów informatycznych wspomagających rozwój e-uczelni*, Zeszyty naukowe WETI PG, seria Technologie Informacyjne, 2009, pp.143-150.
[6] http://en.wikipedia.org/wiki/Single_sign-on
[7] Zimmermann H.: *OSI Reference Model—The ISO Model of Architecture for Open Systems Interconnection*, IEEE Transactions on Communications, vol. 28, no. 4, April 1980, pp. 425-432
[8] http://openvpn.net
[9] Lehtinen R., Gangemi Sr. G.T.: *Computer Security Basics, Second Edition*, O'Reilly, 2006.
[10] Windley P.J.: *Digital Identity*, O'Reilly, 2005.
[11] Maamar Z., Benslimane D., Narendra N.C.: *What can Context do for Web Services?* Communications of the ACM, December 2006, pp. 98-103.
[12] Maglio P.P., Srinivasan S., Keulen J.T., Spohrer J.: *Service Systems, Service Scientists, SSME and Innovation*, Communications of the ACM, July 2006, pp. 81-85.

UOGÓLNIONA KONTROLA DOSTĘPU W HIERARCHICZNYCH SIECIACH KOMPUTEROWYCH

Streszczenie

W artykule opisano architekturę warstwy bezpieczeństwa zaprojektowaną dla rozproszonego systemu zlokalizowanego w wielostrefowej, hierarchicznej sieci komputerowej. Zależnie od lokalizacji użytkownika i systemu klienckiego, jedna wspólna warstwa bezpieczeństwa zaakceptuje żądanie lub nie. Opisane rozwiązanie jest rozwinięciem modelu RBAC. Organizacja systemu zakłada przypisywanie uprawnień nie poszczególnym kontom użytkowników, a cyfrowym tożsamościom, które odpowiadają użytkownikom. Takie połączenie wielu mniejszych pomysłów i metod tworzy z systemu całkiem nowe, nowoczesne podejście do zagadnień bezpieczeństwa rozproszonych systemów zorientowanych na usługi. Przedstawione rozwiązanie zostało gruntownie przetestowane i wdrożone na Politechnice Gdańskiej.