Proceedings of the International Conference on Security and Management (SAM'04) - ISBN: 1-932415-37-8
Computer Science Research, Education, and Applications Press, June 21-24, 2004

A Security Architecture for a Web Portal of Sensitive Archival Records

Binh Q. Nguyen
Army Research Laboratory

Abstract

Web portals are ubiquitous, but an operational portal containing sensitive electronic archival records has not existed in a public network due to security concerns; therefore, setting one up requires serious consideration of defensive security measures. The key topic of this paper is a security architecture for the protection of an experimental portal that provides its stakeholders with convenient and low-cost means for sharing and processing sensitive electronic archival records. The portal also serves as a test bed for conducting empirical research activities in support of the future building of a secure operational portal of sensitive electronic archival records. The proposed architectural framework applies the technical facet of the defense-in-depth strategy that was developed and widely implemented by the Department of Defense.

Keywords: security architecture, sensitive electronic archives, defense in depth strategy.

1. Introduction

A computer programmer views electronic archives as collections of computer files that need to be saved for future uses. An archivist, on the other hand, views them as computer-generated ordered collections of related files that provide the tangible proof of past activities, whose contents, structure, and context must be authentically preserved [1]. The geographically dispersed partners, researchers, and administrators of the electronic records all desire to have a portal that provides convenient and low-cost means for (1) uploading and downloading sensitive archival data files and processing tools, (2) monitoring and sharing research results, and (3) simultaneously performing empirical experiments with defensive security strategies, tactics, and technologies that are potentially capable of meeting the security requirements of a fully operational portal in the near future.

The term *portal* used in this paper refers to an experimental centralized server located within an institution having a physical connection to the Internet. The word *convenient* relates to the process of transferring archival data files via electronic means rather than physical means. The phrase *low-cost* is a relative term that pertains to the cost of running the portal over a public network instead of a private network. The expression *sensitive* used in this paper concerns the secret or confidential information in some of the electronic archival records.

Connecting a portal of such sensitive archival records to the Internet requires extreme care to reasonably assure the confidentiality and integrity of the records

and to provide authentication, availability, and non-repudiation services for their stakeholders. The reason is simply that the Internet is a public network to which every Internet user is physically connected to the portal and that the threats to the portal and its electronic archival records are "real and present." A successful attack against the electronic archival records would potentially damage the reputation of their stakeholders, complicate the operation of the responsible organizations, and deprive future references to historical events.

Amid the phenomenal growth of electronic-commerce sites, a portal containing sensitive electronic archival records operating over the Internet could be constructed by implementing the defense-in-depth strategy [2]. This strategy requires the building of multi-layered defenses. The layers consist of (1) the network and infrastructure, (2) the enclave boundary, (3) the computing environment, and (4) the supporting infrastructure. Although the strategy relies on people, operation, and technology, this paper focuses on the technology feature of the strategy by presenting an architectural framework and describing some proactive measures for safeguarding an experimental portal containing sensitive electronic archival records.

The next section describes the environment in which the portal operates and its protection layers. Section 3 discusses security issues related to the use of the experimental portal of sensitive electronic records and describes some on-going empirical research activities that support the distributed processing of sensitive electronic archival records. Section 4 summarizes key ideas presented in the paper and concludes the paper.

2. Method

Applying the defense-in-depth strategy, the networked computing environment in which the experimental portal runs is placed within an internal subnet of the information infrastructure of the associated research institution as illustrated in Figure 1.

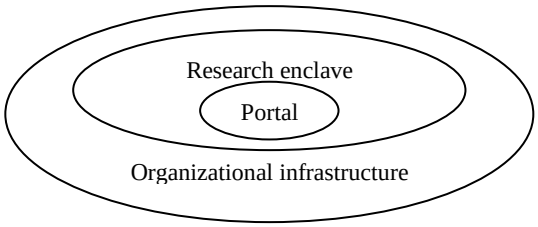

Figure 1. Portal environment

The security measures of the institution are presumed to provide some protection for the portal environment; however, further proactive approaches to safeguarding the portal and its contents still need to be implemented for added layers of defenses to counter all types of attacks against the portal. The potential attacks consist of active, passive, inside, close in, and distribution types [2]. As shown in Figure 1, the portal computing environment is inside the research enclave but isolated from other computing systems within the research enclave. The portal environment directly connects to the Internet and has different security policies than other computing systems that are outside the portal environment. If the portal is compromised by a successful attack, it is prevented to be used by the attacker as a stepping stone for launching attacks against other systems within the research enclave and the information infrastructure of the host organization.

Safeguarding the portal computing environment includes implementation of

technologies that can successfully secure and protect all five basic information assurance services offered by the portal. These services consist of confidentiality, integrity, authentication, non-repudiation, and availability services. Confidentiality services provide a reasonable assurance that the contents of the archives are only disclosed to authorized users while the archives are stored in the portal and the user's computer and when they are uploaded to and downloaded from the portal. Integrity services ensure the wholeness of the electronic records archives and provide a way for detecting a change to the archives. Authentication services provide mutual assurance of identities: the identity of the portal itself and the identity of its users. Non-repudiation services provide the stakeholders of the portal a way for obtaining credible evidence of the use of electronic archival records stored at the portal. Availability services ensure that the portal services are readily accessible to authorized users whenever the electronic archives are needed.

Incoming and outgoing packets traveling between the Internet and the portal must pass through several layers of defense. Each layer has its own set of security policy and defensive security measures. Although describing detailed defense mechanisms of the outer layers of the portal environments is outside the scope of this paper, their functionality can be summarized as follows. The outermost layer usually has an overarching security policy that affects all users and the computing systems that are connected to the information infrastructure of the organization. The security policy applied at the enclave (e.g., electronic records research enclave) further limits access to research data, tools, and documents. The portal itself also has a security policy for accessing the portal and its sensitive electronic archives. Figure 2 illustrates this multi-layered defense mechanism.

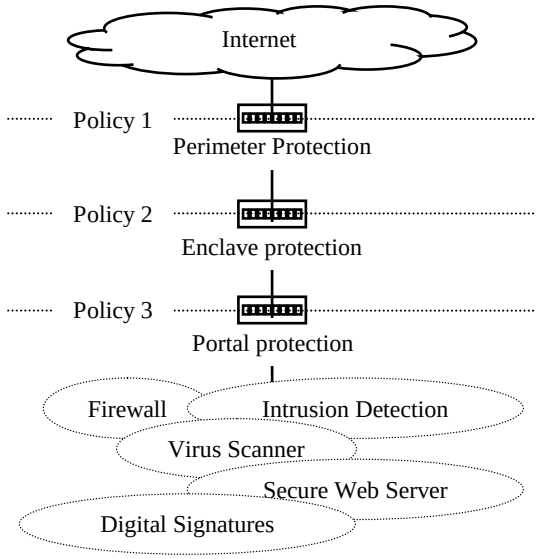

Figure 2. Layers of defense

Using an appropriate security policy at the very first line of defense, the outermost checkpoint inspects every incoming packet and decides whether the inspected packet should be dropped or sent to a destined subnet within the organizational infrastructure. Once an incoming packet has passed this inspection, it is subject to another similar inspection at the perimeter of the research enclave. If it passes the inspection, it is then routed to an appropriate destination within the research enclave. If the destination is the portal of electronic records, then the gatekeeper of the research enclave once again scrutinizes the incoming packet before sending it to the portal for services using the local security policy of the portal.

The security posture at the portal includes the employment of several defensive

mechanisms for the protection of sensitive electronic archives and the physical assets deployed at the portal. All electronic accesses to the portal are considered to originate from untrusted networks; therefore, they must first go through the firewall and the intrusion detection system deployed at the portal computing enclave. The firewall inspects the header of each incoming packet and decides the fate of the packet. The intrusion detection system inspects the contents of incoming packets for possible malicious payload. These defensive mechanisms are also a way to prevent unauthorized insiders from accessing the portal. Figure 3 shows the security process at the portal environment.

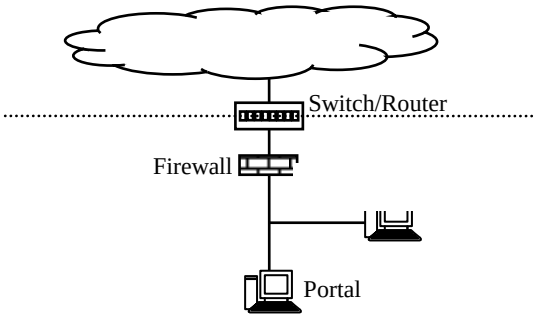

Figure 3. Security posture at the portal

Controlling access to the portal services consists of user authentication and authorization. The user authentication process verifies the identity of the user, which consists of the verification of the legitimacy and correctness of user-presented credentials, including digital certificates in combination of user name, password, or biometric information, and possibly the unique address of the computer from which the incoming packets originated. Once the user authentication is successful, the portal in turn sends its digital certificate to the user to provide the user with a reasonable assurance that the portal is the bona fide server of the sensitive electronic records. The digital certificates used by the user and the portal can be issued by the owner, the administrator of the electronic records, the administrator of the portal, or a third party that is mutually trusted by all the stakeholders of the sensitive electronic records. Authorization mainly concerns restricting the privileges of an authenticated authorized user to access a certain computing resource available at the portal. Authorization depends on the local security policy at the portal.

Monitoring access to the portal services includes logging all the activities that occur in the portal such that the logs can be used for creating credible proof of an activity performed by a user (non-repudiation). A provable activity could be sending a request for a particular electronic record or receiving a requested record.

Protecting the integrity and availability of the portal consists of several proactive measures. Some of these preventive measures include:

- Avoiding potential distribution attacks (e.g., spy ware and Trojan horse) by deploying validated software and hardware products in the portal [3],
- Excluding unauthorized executable contents from being installed at the portal and scanning every newly stored electronic file for possible malicious code (e.g, viruses and worms),
- Rejecting incoming requests for unauthorized network services and monitoring potential abuses of authorized services (e.g., denial-of-service attacks),
- Detecting and preventing unauthorized changes to electronic records and system configuration by implementing access

control mechanisms and by monitoring changes to system and data files as well as their attributes (e.g, ownership, type, and permission),

- Preventing temporary loss of electrical power that the portal needs by connecting the portal to an uninterruptible power system having sufficient power to keep the portal up and running while the main source of energy is being restored,
- Backing up portal resources regularly for the replacement of lost or damaged electronic records and restoring system services, and
- Preventing close-in attacks by restricting access to the physical area of the portal, the computing and network infrastructure, and the source of electrical energy.

3. Discussion

Accessing the portal using a web browser (e.g., Microsoft Internet Explorer) is envisioned to be primary means for accessing sensitive electronic archives as it is commonly used to conduct electronic commerce over the Internet. To buy a product over the Web, a potential customer is not usually required to have a digital certificate, but a legitimate vendor often relies on a digital certificate issued by a third party (e.g., Verisign) to gain the trust from a potential buyer. A communication scenario at the experimental portal differs from a commercial web site in that mutual authentication is required. A user of the portal must present a digital certificate to the portal for verification and validation of the presented identity. Commercially available cryptographic products and services that are enabling this type of secure electronic commerce activities [4] can meet some of the requirements for mutual authentication and for other information-assurance services, including integrity, confidentiality, and non-repudiation services.

Given the current (128-bit) cipher strength of a typical web browser operating in a secure mode, the portal possesses only electronic records having a relatively low level of sensitivity. Using Type 1 cryptographic products to protect classified and highly sensitive electronic records stored at the portal is a possibility, but doing so will require additional expenses in terms of product acquisition costs and procedural and administrative overhead. Therefore, as the main purpose of the experimental portal at the present time is to support geographically dispersed researchers and administrators, the use of commercial cryptographic products is sufficient.

These information-assurance services must be provided to protect sensitive electronic records in all information states. According to Maconachy et al [5], information is found in one or more of the three states: stored, processed, or transmitted. The proposed security implementation at the portal can provide reasonable information-assurance services for the electronic archival records just while they are stored at the portal and in transit. Once an electronic archival record has been transferred from the portal to an external host, it also needs information-assurance services at the remote host to preserve its contents, structure, and context, especially while it is in processing.

Other on-going information-assurance research activities that support the building of an experimental portal and the distributed processing of electronic archival records include several tasks: (1) performing empirical experiments to assess government-

validated security products that are potentially appropriate for the protection of sensitive electronic archives and the portal computing environment, (2) developing a set of performance metrics that are potentially suitable for evaluating the effectiveness of the deployed defensive security products in the experimental portal environment, (3) measuring the performance of the portal operating in unsecured and secured modes, and (4) investigating technical, financial, and operational issues related to the implementation of a particular security product capable of providing information-assurance services at the portal.

4. Conclusion

This paper presented a security architecture for an experimental portal of sensitive electronic archival records operating in a public network. The portal provides indispensable information-assurance services for the electronic archival records to provide their stakeholders with a reasonable assurance that the contents, structure, and context of the records are authentically preserved. The main themes of the architecture incorporate the use of government-validated defensive information assurance products, secure electronic commerce technology and services, and the defense-in-depth strategy.

Acknowledgements

The author sincerely appreciates the support for this work from the U.S. National Archives and Records Administration and the constructive comments of anonymous reviewers.

Disclaimer

The findings in this paper are not to be construed as an official Department of the Army position unless so designated by other authorized documents. Citation of manufacturer's trade names does not constitute an official endorsement or approval of the use thereof.

References

[1] Thibodeau, K, "Building the Archives of the Future", *D-Lib Magazine*, Vol. 7, No. 2, February 2001. URL: http://www.dlib.org/dlib/february01/thibodeau/02thibodeau.html (accessed 13 Jan 04).

[2] National Security Agency, *Information Assurance Technical Framework*, Release 3.1, September 2002, National Security Agency, Fort Meade, Maryland, 20755-6730, URL: https://www.iatf.net/ (accessed 15 January 2004).

[3] National Information Assurance Partnership, Common Criteria, *Validated Products*, URL: http://niap.nist.gov (accessed 15 January 2004).

[4] Ford, W. and Baum, M., *Secure Electronic Commerce*, 2nd Ed, Prentice Hall PTR, Upper Saddle River, NJ, 2001.

[5] Maconachy V., Schou C., Welch D., and Ragsdale D. J., "A Model for Information Assurance: An Integrated Approach," *Proceedings of the 2nd Annual IEEE Systems, Man, and Cybernetics Information Assurance Workshop*, West Point, NY, June 5-6, 2001, pp. 306-310.