ACM Computing Surveys, 50th anniversary commemorative issue, Volume 28, Number 1, March 1996, pages 241-243.

AUTHENTICATION, ACCESS CONTROL AND AUDIT

Ravi Sandhu, George Mason University, Fairfax, Virginia, USA
Pierangela Samarati, University of Milan, Milan, Italy

Authentication, access control and audit together provide the foundation for information and system security.

> Authentication establishes the identity of one party to another. Most commonly authentication establishes the identity of a user to some part of the system typically by means of a password. More generally, authentication can be computer-to-computer or process-to-process and mutual in both directions.
>
> Access control determines what one party will allow another to do with respect to resources and objects mediated by the former. Access control usually requires authentication as a prerequisite.
>
> The audit process gathers data about activity in the system and analyzes it to discover security violations or diagnose their cause. Analysis can occur off-line after the fact or it can occur on-line in real time. In the latter case the process is usually called intrusion detection

We now discuss these three mutually supportive technologies in turn.

AUTHENTICATION

User-to-computer authentication can be based on one or more of the following:

> something the user knows, such as a password,
>
> something the user possesses, such as a credit-card sized cryptographic token or smart card, or
>
> something the user is, exhibited in a biometric signature such as a fingerprint or voice-print.

Password-based authentication is the most common technique but it has significant problems. Passwords can be surreptitiously observed or guessed. Password management is required to prod

users to regularly change their passwords, to select good ones and to protect them with care. Excessive password management makes adversaries of users and security administrators which can be counter-productive. An intrinsic flaw of passwords is that users can share them with other users which breaks down accountability. However, passwords can be effective and are cheap, so they are likely to remain in use.

The second technique authenticates the token rather than the user. Each token has a unique secret cryptographic key stored within it, used to establish the token's identity via a challenge-response handshake. The party establishing the authentication issues a challenge to which a response is computed using the secret key. Sometimes the challenge is implicitly taken to be the current time. The secret key should never leave the token. Attempts to break the token open to recover the key should cause the key to be destroyed. User-to-token authentication can be based on passwords in the form of a PIN (personal identification number).

Biometric authentication has been used for some time for high-end applications. The biometric signature should be different every time, for example, voice-print check of a different challenge phrase on each occasion, or require an active input, for example, dynamics of handwritten signatures.

Technically the best combination would be user-to-token biometric authentication, followed by mutual cryptographic authentication between the token and system services. This combination may emerge sooner than we imagine, although there are social issues in addition to technical ones.

Token-based authentication is a technical reality today, but it still lacks major market penetration. Many existing systems use the desktop workstation as a "token" for authentication with the rest of the network. A cryptographic key is computed from the user's password by the workstation on basis of which workstation authenticates to the network. Kaufman et al [KPC95] describe some of the techniques in current use.

ACCESS CONTROL

Access controls apply usually after authentication has been established. Access control can take on several forms [SS94].

> Discretionary access control (DAC) is based on the idea that the owner of data should determine who has access to it. DAC allows data to be freely copied from object to object, so even if access to the original data is denied access to a copy can be obtained.
>
> Lattice-based access controls [SAN93], also known as mandatory access controls (MAC), confine the transfer of information to occur only in one direction in a lattice of security labels, e.g., low to high but not high to low. MAC emerged from confidentiality requirements of the military but has broad applications for integrity and separation objectives.

Role-based access control (RBAC) requires that access rights are assigned to roles rather than to individual users (as in DAC) [SCFY96]. Users obtain these rights by virtue of being assigned membership in appropriate roles. This simple idea greatly eases the administration of authorizations.

Other forms of access control also exist and this remains a fertile area for further research and development.

Existing systems often take a feature-based approach to access control in which multiple interacting access control facilities are provided to be configured by security administrators to meet their policy objectives. Unfortunately, these access-control features are all too often poorly documented and their interactions poorly understood.

AUDIT

Audit has two components,

collection and organization of audit data [JGB95], and

its analysis to discover or diagnose security violations [LUN93, MHL94].

Audit data needs protection from modification by an intruder. Vast amounts of audit data can be recorded. Audit data tends to be captured at a low level of abstraction. Analysis of audit data is often performed only when violations are suspected. Even so, only audit data connected with the suspected violation are examined.

Intrusion detection systems seek to help carry out audit controls.

Passive intrusion detection systems analyze the audit data, usually off-line, and bring possible intrusions or violations to the attention of the auditor.

Active systems analyze audit data in real time and may take immediate protective response, such as killing the suspected process and disabling the account.

The problem is what to look for in audit data and how to automatically determine whether a violation has occurred or is being perpetrated. The following approaches have been tried.

Anomaly detection is based on the assumption that the exploitation of the vulnerabilities of the system involves abnormal use of the system.

Misuse detection is based on rules specifying events, sequences of events, or observable properties of the system, symptomatic of violations.

Finally, we note that audit analysis is an empirical discipline in which we currently have little historical data.

REFERENCES

[JGB95] Jajodia, S., Gadia, S. and Bhargava, G. ‘‘Logical Design of Audit Information in Relational Databases.” In Abrams, Jajodia and Podell (editors), Information Security: An Integrated Collection of Essays, IEEE Computer Society Press, 1995, pages 585-595.

[KPC95] Charles Kaufman, Radia Perlman and Mike Speciner, ‘‘Network Security,” Prentice Hall, 1995.

[LUN93] T.F. Lunt, ‘‘A Survey of Intrusion Detection Techniques,” Computer & Security, vol. 12, no. 4, pages 405-418.

[MHL94] B. Mukherjee, L.T. Heberlein, and K.N. Levitt, ‘‘Network Intrusion Detection,” IEEE Network, May/June 1994, pages 26-41.

[SAN94] Sandhu, R.S. ‘‘Lattice-Based Access Control Models.” IEEE Computer, Volume 26, Number 11, November 1993, pages 9-19.

[SS94] Sandhu, R.S. and Samarati, P. ‘‘Access Control: Principles and Practice.” IEEE Communications, Volume 32, Number 9, September 1994, pages 40-48.

[SCFY96] Sandhu, R.S., Coyne, E.J., Feinstein, H.L. and Youman, C.E. ‘‘Role-Based Access Control Models.” IEEE Computer, Volume 29, Number 2, February 1996.