



Security and Privacy Auditing in Health Care Information Technology

This Paper was developed by the
Joint NEMA/COCIR/JIRA Security and Privacy Committee (SPC)

The Paper has been approved by:
NEMA (National Electrical Manufacturers Association)
COCIR (European Coordination Committee of the Radiological and
Electromedical Industry)
JIRA (Japan Industries Association of Radiological Systems)

November 2001

© JOINT NEMA/COCIR/JIRA SECURITY AND PRIVACY COMMITTEE (SPC)

www.nema.org/medical

Secretariat: NEMA (National Electrical Manufacturers Association) www.nema.org
1300 North 17th Street, Suite 1847, Rosslyn, VA 22209 USA tel:703-841-3200 fax:703-841-5900
Secretary: Stephen Vastagh t: 703-841-3281 fax 703-841-3381 E-mailto:ste_vastagh@nema.org

May be quoted if reference and credit to SPC is properly indicated.

SECURITY AND PRIVACY AUDITING IN HEALTH CARE INFORMATION TECHNOLOGY

November 2001

The white paper 'Security and Privacy: An Introduction to HIPAA' by the Joint NEMA/COCIR/JIRA Security and Privacy Committee (SPC) (discusses security and privacy concepts in health care, that require both administrative and technical measures (www.nema.org/medical/spc - see white papers.) Currently data security and data privacy legislations are evolving in e.g. the US (HIPAA), the European Community (country specific implementation of the directive EC 95/46 is required), and Japan (HPB 517). Because of differing cultures, legislation in some countries simply focuses on principles of protection whereas in other countries requirements are defined in order to guarantee such principles to the individual. But eventually a certain set of common technological requirements may be derived comprising, among other things, risk assessment, risk mitigation, and internal auditing. This paper describes how auditing in a medical IT environment, including electronic medical devices, can effectively meet legal mandates and provide the individual accountability and anomaly detection called for in privacy and security regulations. Follow-up SPC efforts in early 2002 will deal with guidelines for real-world security auditing implementations based upon these requirements that can be achieved with contemporary technology.

1 Audit Controls – a Basic Security Requirement

Administrative and technical procedures and tools need to be used by Covered Entities (CEs), who are the health care providers and others that are entrusted with Protected Health Information (PHI), to guarantee that PHI is accessed only as authorized by and consented to by their patients. One suitable administrative procedure is to conduct internal audits, which are in-house reviews of the records of system activity. The mechanisms employed to record and examine system activity are called audit controls, and the data collected and potentially used to facilitate a security audit is called the audit trail that in turn may consist of several audit files located on different entities of a network. Both audit controls and the audit trail are needed to monitor the security and privacy behavior of a system.

The local implementation of audit controls, including audits and audit trails, not only depend on the technology currently available, but also on the level of risk a CE uncovers and accepts in accordance with its risk analysis and risk management processes. For this reason no single recommendation can provide exact specifications for every situation. This white paper recognizes that auditing can also be useful and may be required for several purposes including, for example, patient accounting and equipment usage. However, the present paper is intentionally focused on discussing suggestions for auditing in health care IT that is specifically tailored to individual accountability as it relates to accessing PHI, and is based upon best practices and reasonableness.

2 Audit Trail Events

The overall goals when constructing an audit trail are to record who did what to which object, when and on which system. In principle it is possible to detect and record every keystroke of any user at any modality or any workstation. However, this would lead to a vast amount of collected data that in and by itself would not guarantee effective auditing. In fact it is necessary and possible to control the amount of data collected while allowing the above goals to be realized at the same time.

Detecting abnormalities often depends on first being able to define normalcy. We know that even normal events can be interpreted with both positive and negative connotation. For example, a user logging in at an unusual time can be interpreted in the negative way as a potential unauthor-

ized access or in the positive way as a mere shift change. The correlation of different log files located on different machines can help sufficiently answer such questions.

Well-conceived and properly enforced audit controls define specific events that should be logged as audit entries. Together with clear guidelines for evaluation of the data collected they allow the audit trail to become a valuable investigative tool. In the section to follow we suggest numerous specific events that should be logged in an audit trail. Note, however, that ultimate determination of what events need to be entered depends upon the local situation, policies and procedures, and the results of a careful risk analysis.

The following events should be entered in the audit trail at the discretion of the Security Administrator.

2.1 PHI-Related Events

Privacy rules strictly restrict PHI use and disclosure; therefore many events dealing with PHI should be able to be logged. PHI-related events are those that specifically deal with the use and potential disclosure of protected patient data. This section will identify audit entries that specifically deal with access to PHI.

2.1.1 Create Events

PHI Create events cover system or user actions resulting in the creation of PHI, such as:

- Creation of records that contain PHI (e.g., images, input of data records, patient histories, billing and insurance data).
- Import of records that contain PHI.

2.1.2 Modify Events

Since PHI contains the history of a patient during its medical treatment it is rarely static. New data are added or specific records have to be updated by physicians, nurses and authorized others. Thus the original data are modified during routine work. In case they are modified in an unauthorized or accidental manner audit files could be the only tools available to reconstruct the original data. Modify events can include:

- Editing of data (e.g. appending, merging, modifying).
- Re-association of data.
- De-identifying of PHI.

2.1.3 View Events

The simple act of viewing PHI can lead to a compromise of its confidentiality if viewed by any inappropriate person. To be able to reconstruct a record of which individuals accessed what PHI if an investigation becomes necessary, many events should be logged, falling into the following three categories:

- Access to PHI by any user.
- Export of PHI to digital media or network.
- Print or FAX of PHI.

While this may produce big audit logs, the ability to audit suspicious activity for a period of time can be of great use during monitoring or investigations. Analysis of view events can be used by the Security Administrator to search for:

- Access to PHI by anyone not directly related to the patients treatment, payment or health care operation.
- Access to information not corresponding to the role of the user.
- Access to PHI of VIPs or community figures.
- Access to records that have not been accessed in a long time
- Access to the PHI of an employee.
- Access to the PHI of a terminated employee.
- Access to sensitive records such as psychiatric records.

- Access to PHI of minors.
- Data recorded without a corresponding order.

2.1.4 Delete Events

Once created PHI has to be maintained according to legal requirements or local policies for many years. Deletion of PHI may harm the patient and hinder a successful medical treatment or even make it impossible. To avoid negative consequences for the patient any deletion of PHI should be logged carefully, including:

- User command to delete PHI (e.g., data records at business associates after performing their duty, image data stored locally on a modality after transfer to a central archive).
- User command to delete PHI before the correct transmission of PHI was confirmed by the receiver.
- Automated command to delete PHI.

2.2 Non-PHI Events

2.2.1 General

There are many events that occur during routine operation of an IT system that are not PHI- or security-relevant of and by themselves. However, as shown above even routine events can become abnormal if they happen under specific circumstances, perhaps depending on the local situation. Therefore the following events should be captured:

- Machine startup and shutdown.
- Successful login and logout of users.
- Changes to user accounts (creation, modification, deletion).
- Automatic logout of a user after exceeding a locally-defined time of inactivity.
- Switching to another user's access or privileges after logging in with one's own identification.
- Software or hardware modification.
- Update of virus signatures.

2.2.2 Operational Events

This class covers auditable events related to general operation of a system, including those related to user activity, and should include:

- Login attempts with failed identification or authentication, also known as failed login attempts.
- Changes of the time or date of the system.
- Emergency mode operation.
- Detection of a virus.
- Detectable hardware errors.
- Changes to log files (creation, deletion, configuration).

2.2.3 Communication Events

Networked IT systems present special auditing needs, and should include the following events:

- Network link failures.
- Device connection failure due to device identification or authentication failure (also known as a failed connection attempt).
- Network and device connections dropped.
- Data integrity verification failure for information transmitted over a network.
- Message authentication failure for information transmitted over a network.
- Evocation of a network abnormality alarm.
- IP addresses of successful and unsuccessful connections.
- Changes to network security configuration (e.g., firewalls if part of a medical IT system)

3 Content of Audit Trails

Audit trails containing only the event itself without any additional information would be nearly worthless. Useful audit trails need to include specifics in terms of time and actor in order to fulfill

the above-mentioned overall goal of investigative usefulness. The following should be contained with each audit entry when appropriate or available:

- Date and time of the event (Note: In case of a network audit trail time on the various distributed systems or workstations has to be standardized. This can be a standard time for the local network or UTC as the global standard. The precision should be a millisecond as standard practice in network security today.)
- The ID of the user who caused the event.
- The application that created the audit event.
- The application(s) responsible for executing the event.
- The component or workstation where the event happened.
- Description of the event (e.g. patient record identifier).

3.1 Location of Audit Trails

Just as important as the content is the location of the audit trail. A comprehensive internal audit must regularly evaluate any local audit files. There are some security-related events that can only be detected via a synoptic view of more than one audit file, for example, if one user is working on two different workstations at the same time. It may be difficult to bring together many files from many systems maintained in many different vendor specific formats. However, the different audit files need to contain some well-defined and accurate information.

For those devices that are always connected to another device or to the local IT-network the audit trail may be maintained locally or in some central part of a system. In case a system of devices is used, combined audit trails can show interactions between these devices. This would result in a more complex audit trail that can be useful in providing a more complete accountability picture. However, there are also many devices and systems that work on a stand-alone basis. These systems may never be connected to other equipment or systems with data interfaces, or may be connected only some of the time. Hence, the only way to ensure integrity of the audit trail is to maintain audit files locally. The audit control policy will need to specify rules for the frequency of gathering these distributed files in order to allow the evaluation of the data they contain.

3.2 Level of Logging

Each device, system, or application will likely come with its own vendor-dependent set of events that can be tracked. The local policy at each facility will likely define different levels of auditing from each device or system. In a networked system the software layer on which logging should take place has to be carefully defined. For example, a large number of authorized requests in proper format for the transmission of very large amounts of data could appear normal at the network layer, e.g., at the firewall, if the syntax of every data packet is correct. However, if every request is addressed to the same application on the same workstation, then a denial-of-service attack might be in progress.

In general audit trails should contain events at a functional level. For example, it might be important to record changed parameters of a medical image, but it might not be necessary to store the whole before and after images themselves. Finally periodic review of the Security Policy and the actual setting will also assure the integrity of the data.

4 Maintenance and Storage Term of Audit Trails

Wherever the audit trail is created and maintained, access to it needs to be controlled to ensure its integrity. Audit trails must be accessible only to authorized staff and must not be accessible to others. This audit trail access control should be the same as the normal access control structure and should not allow modification of the audit trails.

Audit trail retention will vary depending upon legal requirements and operational usefulness. Some data will only be retained for a few days. Other data may be retained on-line for rapid access for a short period (e.g., 14 days), and then archived off-line for an extended period (e.g. 2

years). Laws and regulations frequently require longer retention such as the 6-year requirement in HIPAA for disclosure accounting. Facility operational plans should clearly identify retention and restoral intervals for different audit trail information.

5 Conclusion

This document provides a comprehensive discussion of audit trails and an introduction to audit controls. Obviously there can be no one-size-fits-all audit control. However, audit controls are a valuable tool to support deterrence, local risk management, and the effort of a covered entity to comply with privacy legislation. It's important to keep in mind that even the most sophisticated audit controls detect only past violations of data security and privacy.

6 Glossary

| | |
|-------------------------------|---|
| audit | In-house review of the records of system activity. |
| audit controls | The mechanisms employed to record and examine system activity. |
| audit file | The smallest part of an audit trail in an IT system, usually a single data file containing the audit data of only one entity. |
| audit trail | Data collected and potentially used to facilitate a security audit |
| compromise of confidentiality | A violation of the privacy of any secret information, e.g. PHI, caused by unauthorized release or disclosure of that information. |
| data integrity | The property that data has not been altered or destroyed in an unauthorised manner. |
| denial-of-service attack | Generally, the intentional and malicious act of flooding available resources with unnecessary or unwanted data or commands for the purpose of disrupting service. |
| de-identification | A process whereby the individual identity of a specific patient is removed from data. |
| functional level | A level of abstraction whereby the intent of the action is communicated instead of the specifics. |
| message authentication | A process in communication whereby the identity of the sender is verified by the receiver to a level of assurance acceptable to the receiver. |
| risk management | A process of assessing risk, taking steps to reduce risk to an acceptable level, and maintaining that level of risk. |
| security administrator | A user of a system with privilege to operate and manage specific functions, usually related to security tasks such as user identification, authentication, authorization, and accountability. |

-End-