# Information system activity review in an academic medical center

David McKelvey
dpm0217@ecu.edu

***Abstract* — Information System Activity Review (ISAR) is intended to detect and limit damage to the confidentiality, integrity, and availability of a system. ISAR is required of covered entities, like Academic Medical Centers, for compliance with the Health Insurance Portability and Accountability Act (HIPAA) Final Security Regulation. The Audit and Accountability family of controls in the National Institute of Standards and Technology (NIST) Special Publication 800-53 is used as a guide. Security Event Management (SEM) is a more comprehensive approach to ISAR.**

***Index Terms* — Information System Activity Review (ISAR), security information and event management (SIEM), security event management (SEM), HIPAA, information security, cybersecurity.**

## I. Introduction

Information system activity review (ISAR) is required of covered entities by the Health Insurance Portability and Accountability Act (HIPAA) final security rule [1]. ISAR is a required control and states "Implement procedures to regularly review records of information system activity, such as audit logs, access reports, and security incident tracking reports." A covered entity is defined at § 160.103 as one of the following: (1) A health plan; (2) a health care clearinghouse; (3) a health care provider who transmits any health information in electronic form in connection with a transaction covered by part 162 of title 45 of the Code of Federal Regulations (CFR).

ISAR is classified as a passive control, that is, it primarily is used to detect and limit damage to the Confidentiality, Integrity and Availability (CIA) of a system instead of directly blocking compromises to CIA. Confidentiality is the assurance that only authorized individuals have access to the services provided by the system and to the data contained therein, typically controlled through an ID and password. Integrity is the assurance that the information held by the system is correct and has not been tampered with. Availability is the assurance that the services provided by the system respond to authorized users as intended.

The purpose of this paper is to provide some guidance on information system activity review from the viewpoint of an Academic Medical Center (AMC). Applicable best practices, the HIPAA Privacy and Security rules, NIST guidance and the experience of the author gained at a large AMC will be used as a guide. As with other controls required by HIPAA, interpretation of the requirement is left to the covered entity.

An AMC has two mutually exclusive sets of principles. The academic principles value openness among faculty, staff, and students both within the institution and between institutions. Collaboration, learning, and research require an open network with few restrictions. Academic Medical Center principles must include strong protections for the privacy and security of protected health information (PHI). A network with many restrictions, the opposite of an open network, is most conducive for compliance with the HIPAA Privacy and Security rules. The challenge of security in this environment is to comply with both the letter and spirit of the rules while minimizing the impact of privacy and security measures on academic principles.

**Figure 1 – Audit and Accountability Outline of controls from NIST Special Publication 800-53**


Manuscript received November 20, 2005.
David McKelvey is with the Information Security Office of a large academic medical center (e-mail: dpm0217@ecu.edu).

A good starting point to interpret ISAR is with the Audit and Accountability family of 11 controls contained within NIST Special Publication 800-53 Recommended Security Controls for Federal Information Systems [2], see Figure 1.

I chose to use the Audit and Accountability family in NIST Special Publication 800-53 because in my opinion it: 1) covers all the elements one needs to consider for ISAR, and 2) will be required for federal information systems in 2006. Note however that the publication is being revised at this time [3].

## II. AU-1 Audit and Accountability Policy and Procedures

ISAR begins with and is defined by policy. Policy is needed to document the why, what, who, and how. The "why" should include: compliance, information confidentiality, system integrity, and system availability. The "what" should describe the information to be logged and the source. The "who" should include: administrators that have to implement the controls, management that are responsible for system CIA, and users who generate auditable events. The "how" spells out expectations for implementation of the technical controls.

The policy should have an owner and be reviewed and updated on a regular basis e.g. yearly. The policy needs to be communicated to everyone with roles that touch or are touched by this policy. All applicable security policies should be communicated to employees initially in an orientation session and annually thereafter. A short test should be administered to prove competence and a record maintained of completion. Automatic reminders make it easier to maintain compliance.

Ensure that users who access systems that are audited are warned that their actions are being logged. Disciplinary actions based on system logs may not be taken against users who are unaware that their actions may be logged. User notification of logging also has the effect of reducing inappropriate access when users know they can be monitored. The user notification requirement may be met by a banner or system usage memo made available to users when they logon. Lastly the policy needs to have sanctions which are the documented expected consequences of non-compliance.

## III. AU-2 auditable events

Which events should be audited? The decision of which events to log and audit should be risk based. That is, log and audit events that the failure of which has negative consequences for system CIA. There are two classes of auditable events; system, and application.

System auditable events at a minimum should include: system resources, attempted and successful user logon, and logoff. Application auditable events is a little more tricky to determine and can vary between systems.

The impact to the resources of the system needs to be considered when determining what to log. System resources most likely to be affected include: CPU, memory, and disk I/O. The logging of system auditable events usually does not cause a resource problem because they usually constitute a manageable volume. However, application logging can be a resource problem because if file accesses are logged, the volume can be significant, especially if read access is logged. Read access may be required because of a HIPAA privacy rule that requires a covered entity to be able to provide a record of accountable disclosures upon patient demand that are not for treatment, payment, or operations (TPO). Examples of accountable disclosures include state public health regulations e.g. in North Carolina reportable neoplasms are required to be registered with the NC tumor registry.

## IV. AU-3 content of audit records

The two classes of auditable events; system, and application can require very different contents but do have some fields in common. The fields in common that are required in every audit record include: timestamp of the event, the system component, the type of event, an identifier, and the result of the event.

System audit records are well defined because they are based on standard operating systems and each performs similar functions. I am grouping audit records from network security devices like firewalls and VPN concentrators with system audit records for purposes of this discussion. Network security devices have attributes similar to operating systems.

Application audit records are not well defined because they can vary widely by application. The content of patient records in a pharmacy database is very different from the records in a laboratory results database. The HIPAA privacy rule requires that only those with a valid reason to access a given patients records is permitted to do so. However, the way that an AMC operates does not lend itself to being able to implement technical controls on who can access individual patient records. A large number of clinicians routinely need access to individual patient records because of the division of labor prevalent in an AMC. Consequently, compliance relies primarily on education, training, and awareness. Log analysis and the threat of log analysis also play a compliance role. In order to do log analysis for cause, or randomly, the application must have logged sufficient details to be able to determine who looked at what for whom. This level of detail is application specific and typically requires that a massive amount of information be kept.

Because of the massive amount of information one might be tempted to do sampling by time interval, application, patient, or user. However, sampling is not an equal replacement for full logging. I would not recommend sampling because disk space is relatively cheap, there is too much opportunity to miss intrusions, and sampling would not demonstrate due diligence in my opinion.

## V. AU-4 AUDIT STORAGE CAPACITY

The amount and type of storage needs to be considered. Storage capacity needs to be sufficient to meet the logging requirements identified by risk analysis and should be documented in policy.

Consider a tiered storage approach. For example, one could have three tiers: raw, normalized, and archived. Raw log records could be kept for only a relatively short time, just long enough to ensure that the raw data was successfully normalized and entered into regular storage. Normalized log records are in a form that is standard across systems to allow for easy reporting using standard tools. Log records in a normalized state are where routine analysis is performed. Archived log records are normalized records that are moved to lower cost storage when sufficient time has passed to make it unlikely that they will be needed for analysis.

## VI. AU-5 AUDIT PROCESSING

The audit processing control is intended to monitor audit logging operations. The two conditions to monitor for include: out of space and any other failure that could result in a loss of log records. An organization that goes to the trouble and expense of deploying an ISAR system should ensure that their investment is not wasted.

For an AMC, compliance with HIPAA rules requires an AMC to maintain an uninterrupted log record. The monitoring of the ISAR system demonstrates due diligence or reasonable care on the part of the AMC. Even if something bad was to happen and log records were lost due diligence should protect an AMC from sanctions. To prove due diligence the monitoring program should itself produce a log.

## VII. AU-6 AUDIT MONITORING, ANALYSIS, AND REPORTING

Audit monitoring, analysis, and reporting are the main functions of an ISAR system. The goals of monitoring and analysis include: detection of inappropriate and unusual activity, performance issues, and application problems.

Inappropriate and unusual activity may be caused by authorized users, unauthorized users and system failures. The detection means for the different sources of potential problems are different. Inappropriate and unusual activity may adversely affect the confidentiality, integrity, and availability of a system.

Performance issues are detected through the monitoring of system resources like CPU, memory, disk space, temperature, etc. If there is a problem or shortage of system resources availability of the system is compromised.

Application problems may be detected through monitoring of the application log. Application problems can result in a compromise of the availability of the system.

Automated tools are essential to be able to effectively perform monitoring and analysis in medium to large organizations. There are a number of vendors that offer tools. The tools may be marketed as Security Information and Event Management (SIEM) or more simply as Security Event Management (SEM), interchangeable terms in recent use. A security event is something that was observed in a security system that may or may not be a risk to the CIA of the system, sensitive data on the system, or the network the system is connected to. A SEM system is a system that can collect events, monitor them, issue alerts based on defined rules, and produce reports [4].

The idea of SEM is the ability to pull together security and resource information from a wide variety of security and system devices and make sense of it [5]. Analysis of security information in more traditional log analysis systems was only after the fact, retrospective. The new SEM systems add real-time analysis to existing retrospective abilities. NetForensics and ArcSight are some of the better known vendors offering SEM products [6]. In theory, integrating information from all security devices deployed by an organization should allow better and more complete event correlation of inappropriate and unusual activity. However, Jim Hurley, an Aberdeen Group analyst, warns that “vendors have a long way to go when it comes to linking security event data to business impact.”

An alternative to the centralized SEM model described in the preceding paragraph is what’s called the intelligent agent model [7]. This model deploys an agent on each security device to be monitored. The agent filters on events, interacts with other agents, and deliberates on what it detects acting on its deliberation when it determines that it is being attacked. I would describe the intelligent agent model as decentralized.

A different approach to a decentralized model was described by J. Reynolds and L. Clough [8]. They point out that network-based intrusion detection systems (IDS) have been found to have detection accuracy of between 30% and 90% for known attacks and 0% for attacks never seen before. In response they developed what they call Windows Intrusion Monitor and Response (WIMR) host-based IDS. Their process continually monitors the system log for running processes, system resources like CPU and memory, and changes to critical files. Within a second, unapproved running processes are terminated, files created by unauthorized processes are deleted and critical files changed by unapproved processes are replaced from read-only medium. WIMR supplements information not found in the system log with data from other sources including internal resources and SNORT, a popular open source IDS.

Analysis and reporting, even with proper tools, can be time consuming. The more security devices and tools that send log records to a central repository the more material tools have to work with to detect inappropriate and unusual activity. The downside to all this data is the need to achieve a delicate balance between false negatives and false positives, and the need for more storage space. Balance is only achieved

through constant vigilance and fine tuning of analysis and reporting parameters.

The HIPAA Security Educational Paper Series "Security Standards – Administrative Safeguards" [9] offers some guidance that I found helpful:

- ISAR should promote continual awareness of any information system activity that could suggest a security incident.
- ISAR should be customized to meet the covered entity's risk management.
- ISAR may be different for each covered entity.

## VIII. AU-7 AUDIT REDUCTION AND REPORT GENERATION

An important component of an ISAR system is the ability to reduce the raw logs and report on the result. Reduction of logs is essential to be able to analyze the data. Reduction is primarily a function for after-the-fact analysis because real-time analysis is more concerned with reacting to individual raw log records. A requirement of log reduction is to not alter the original log records.

Vendors have not yet agreed on standards to apply to the problem. Some potential standards are included in Figure 2. However, none of the standards listed are comprehensive enough to address all the needs of a SEM system. CVE is concerned with viruses and worms, IDXP and IDMEF are concerned with signature-based IDS, and SESA is a vendor developed solution.

**Figure 2 - Log Data Normalization Potential Standards**

| Who | What |
|---|---|
| Common Vulnerabilities and Exposures (CVE) | Security event descriptions |
| Computer Emergency Response Team (CERT) | Security event descriptions |
| Internet Emergency Task Force (IETF) | Intrusion Detection Exchange Protocol (IDXP) |
| Internet Emergency Task Force (IETF) | Intrusion Detection Message Exchange Format (IDMEF) |
| Symantec | Symantec Enterprise Security Architecture (SESA) |
| Distributed Management Task Force (DMTF) | Logging, alerting, and reporting |

## IX. AU-8 TIME STAMPS

This control is a requirement to synchronize the timestamps on all audit log records. Network Time Protocol (NTP) support is included with all off-the-shelf operating systems and many embedded operating systems. However, one needs to ensure that the client is started on each device, that the clients point to a common NTP server, and that they update themselves at about the same time each day. Particular attention needs to be paid to the times twice a year when daylight savings time begins and ends. Some operating systems require a reboot or reset after the clock is changed.

It is important to synchronize time on all systems that audit logs are pulled from to enable the correlation of data. For intrusions especially, it would not be possible to get a clear picture of what happened or even know that an intrusion was under way without knowing that the timestamps from different security devices were synchronized.

## X. AU-9 PROTECTION OF AUDIT INFORMATION

The audit log and audit log tools themselves need protection. Steps taken by attackers of hosts include making changes to the operating system to make their processes invisible and changing or deleting log records. Their goal is to eliminate or reduce evidence of their presence [10].

Techniques that may be used to protect the audit log include the following: encryption of the file, writing audit log records to write-only mediums e.g. CD-R, and not storing the audit log records locally by writing to a separate log server. A separate log server is probably most common and confers the added benefit of the ability to correlate events across security devices and hosts.

## XI. AU-10 NON-REPUDIATION

The system must ensure that individuals can not later falsely claim that actions or records attributed to them were not theirs. Some of the non-repudiation elements include: unique IDs, strong passwords, secure audit logs, digital signatures, and timestamps.

## XII. AU-11 AUDIT RETENTION

The factors that go into determining how long to keep audit log records include: the data retention policy of the institution, the amount of data, and a risk analysis. NIST Special Publication 800-16 Computer Security Incident Handling Guide [11] recommends a minimum of a few weeks and preferably for a minimum of a few months. Consider a tiered retention strategy that archives audit log records to CD-R, DVD, disk, or tape after a minimum of a few months. Archived audit records should never be changed, thus there is some attraction to using a write-once DVD or CD-R to guarantee the integrity of the archived data.

The reason to retain audit log records for long periods is for data analysis of intrusions which may be needed up to many months later. In the news recently was the case of an institution that discovered that a root kit on one of their servers had been in place for approximately 2 years.

There is no consensus among covered entities as to the retention period for HIPAA security rule compliance. There are references within the HIPAA Final Security Rule for a retention period of 6 years for other controls but no explicit retention period is given for audit log retention. The guidance from the Centers for Medicare & Medicaid Services (CMS) on ISAR is that it should be based on risk analysis.

## REFERENCES


[1] "Health Insurance Reform: Security Standards; Final Rule," in Federal Register, February 20, 2003, 34 Information system activity review ((§ 164.308(a)(1))

[2] "NIST Special Publication 800-53 Recommended Security Controls for Federal Information Systems," February 2005, AC-13, AU-1-11

[3] F. Olsen, (2005, November). "NIST to tweak mandatory security controls," Federal Computer Week [Online]. Available: http://www.fcw.com/article91381-11-11-05-Web&RSS=yes (November 2005)

[4] C. King, "Security Management: Making Sense of Events," Business Communications Review, Sep. 2001, VOl. 31 Issue 9, p32

[5] D. Margulius, "Managing it all," InfoWorld, Jan. 2003, Vol. 25 Issue 2, p42

[6] J. Vijayan, "New Tools Help Users Manage Security Events," Computerworld, Feb. 2004, Vol. 38 Issue 7, p16

[7] K. Boudaoud, N. Foukia, Z. Guessoum, "An Intelligent Agent Approach for Security Management," in Proceedings of the HP OpenView University Association HP-OVUA 7th Plenary Workshop, Santorini, Greece, 2000, session 6-1. [Online] Available: http://www.hpovua.org/PUBLICATIONS/PROCEEDINGS/7_HPOVUAWS/ (November 2005).

[8] J. Reynolds, L. Clough, "Continual repair for windows using the event log," in Proceedings of the 2003 ACM workshop on Survivable and self-regenerative systems: in association with 10th ACM Conference on Computer and Communications Security

[9] "HIPAA Security Series - Administrative Safeguards," June 2005, p6-7. [Online]. Available: http://www.cms.hhs.gov/hipaa/hipaa2/education/HIPAA%20Security%20Series%20Administrative%20Safeguards.pdf

[10] F. Buchholz, C. Shields, "Providing process origin information to aid in computer forensic investigations," in Journal of Computer Security 12 (2004) 753-776 IOS Press

[11] "NIST Special Publication 800-16 ",Computer Security Incident Handling Guide," January 2004, p. 3-10