Security Technical Work Group

Audit Logging for Storage

A SNIA Security White Paper

September 21, 2005

SNIA
STORAGE NETWORKING INDUSTRY ASSOCIATION

Authors: Eric A. Hibbard, CISSP, ISSAP, ISSMP, ISSEP
Hitachi Data Systems

Richard Austin
Hewlett-Packard

Larry Hofer
McData

Many thanks to the following individuals for their contributions to this paper.

Nobuyuki Osaki
Mark Carlson
Tad Taylor
Naoki Iwami

Roger Cummings
LeRoy Budnik
Dr. David Black
Andrew Nielsen, CISSP

Table of Contents

Introduction

Experts agree that audit log management is a critical element of any organizations's risk management strategy. Audit log data (or just log data) can provide a complete record of access, activity, and configuration changes for applications, servers, and network devices. It can be used to alert management and administrators to unusual or suspicious network and system behavior. Additionally, log data can provide auditors with information required to validate security policy enforcement and proper segregation of duties. Lastly, IT staff can mine log data during root-cause analysis following a security incident; this is particularly important for the recovery and/or damage cleanup as well as the remediation activities.

Considering all of these potential uses, audit log management not only assists in achieving corporate compliance, but also reduces the risk of legal exposure from security breaches and costly network downtime. This whitepaper discusses log management from a storage security perspective and provides specific information as it relates to storage resources and networks.

Regulatory Compliance and Audit Log Management

Increasingly, data protection and privacy regulations are holding firms accountable for safeguarding their data. However, the requirements within these regulations often lack specificity and instead describe the protection using "technology neutral" language like "reasonable measures" or "best practices" that refer to mitigating risk, retaining evidentiary matter, and monitoring. The penalties, on the other hand, tend to be well defined and include remediation, fines, and even imprisonment.

Ultimately, the determination of which and how many controls constitute an effective internal control environment is made by management in consultation with an external auditor. Since there is little regulatory guidance on the specific controls, the standard of ***due care*** means doing at least what one's peers are doing as well as following security best practices frameworks (e.g., COBIT, BITS, COSO) and standards (e.g., ISO 17799 and NIST SP 800-53).

The regulations themselves do provide a little guidance on log management, especially for retention periods. The following regulatory samples show the details that are specified:

- *Healh Insurance Portability and Accountability Act (HIPAA)* – US law that requires hospitals, physicians, and managed care companies to adopt security, privacy, and data standards for medical information. It requires organizations to "audit and monitor system and user activity across the entire network, identify and investigate security breaches and suspicious behavior, and maintain an audit trail of user and network activity." HIPAA also specifies that companies should "Retain and protect log data as evidence…up to 6 years."
- *Basel II Accord* – Requires all internationally active banks to adopt similar or consistent risk management practices; banks are required to implement a comprehensive program of

risk prevention, detection, analysis and management, and mitigate operational risks associated with IT systems by 2006. The accord recommends "retaining activity logs for 3 to 7 years."

- *Sarbanes Oxley (SOX) Act* – US law that establishes corporate accountability for all public companies, requiring strict IT controls and processes. SOX requires companies to "Audit unauthorized access, misuse and fraud, in order to ensure the accuracy of corporate financial and business information" and "maintain financial records for seven years."

Elements of Audit Log Management

It is important to recognize that there are two distinct camps when it comes to audit log management. The first (and more traditional) views logging as a useful debugging and health monitoring tool and, consequently, employs ad hoc solutions and less rigor in the management process beyond such details as assuring a consistent time base. The second views logging as a source of evidentiary information, and as such, there is greater focus on the type of information logged, on maintaining a chain of custody, and on long-term protection and retention of the log data. This latter approach is much closer to what is required to address accountability aspects of regulatory compliance.

From an abstract perspective, audit log management consists of the following:

- *Aggregation* – Combination of configuring devices for external logging, transferring event data (including reliability and security measures when needed), integrity verification, filtering, relaying/forwarding, and storing the data in the necessary event log repositories (files, databases, etc.). With the central aggregation of log data, it becomes possible to implement retention strategies as well as to correlate log data across multiple sources.
- *Analysis* – Combination of real-time and historical log data mining as well as fast text-based searches; at a bare minimum, this analysis must include provisions for performing access and change control audits.
- *Alerting* – Recognizing unusual or suspicious activity and converting that recognition into an action; alerting strategies that trigger automatic notifications often use rule-based alerting, regular expression-based matching, and/or anomaly detection and intelligent threshold alerting.
- *Archiving* – Securely and reliably transferring log data in its various forms (raw, binary, database, etc.) to a centralized archive, using secure hashes to verify the integrity of the data, establishing retention and disposal parameters, and protecting the availability and confidentiality of the log data.

Effective audit log management is also dependent on the following:

- Policies on audit logging and the handling of log data
- Devices that support external logging
- Adequate storage facilities to retain the log data for the required retention periods

- Use of a reliable and trusted time source like a Stratum 1 or 2 Network Time Protocol (NTP) service

Logging and log management, using Syslog, have been around since the 1980s. What's new is the need for storage technology to natively participate in this aspect of an enterprise's security and compliance infrastructure. Of equal importance, enterprises need to include the storage layer in their monitoring and accountability strategies and approaches.

Anatomy of Audit Logging

The Information Security Forum's (ISF) *The Standard of Good Practice for Information Security* (Version 4.1, January 2005) states that the objective for logging is "To ensure individual accountability and to enable incidents, such as access violations, to be investigated and resolved." This is easy to state, but a major challenge to implement in heterogeneous environments that involve hundreds or thousands of hosts and devices. For example, simply detecting an incident can require the correlation of data from multiple systems and devices (hosts, firewalls, IDS, switches, etc.) that are then compared to a profile of past behavior to identify anomalous behavior.

So, how does one deal with this complicated and less-than-glamorous activity? A clear understanding of the audit logging requirements, which can be derived from the organization's policies and compliance audit criteria (i.e., the legal requirements and/or auditor's checklist), is a critical first step. Almost everything flows from these requirements and there are no one-size-fits-all solutions. Thus, the remainder of this section is based on dealing with a fictitious set of requirements (worse case scenario) that can be summarized as:

- Multiple compliance requirements (e.g., SOX, SEC, and HIPAA) exist and involve different types of audits
- Security incidents perpetrated by both internal and external sources are a real possibility
- An extremely heterogeneous environment is used by the organization
- Multiple levels of data sensitivity and criticality exist
- Some audit log data must be preserved as evidentiary information[1]
- Lengthy retention periods exist for some, but not all, of the log data

To address this hypothetical set of requirements, a moderately complex audit logging implementation would be necessary. As a first pass, the heterogeneous environment predisposes the solution to a cross-platform logging approach such as the de facto standard, Syslog[2], to transport log event data. Further, the approach needs flexibility in handling logging information in a variety of shapes and forms from these disparate sources. Again, the Syslog architecture

[1] To be court-admissible, log data must be: 1) complete, 2) accurate, and 3) verifiable.

[2] The Internet Engineering Task Force's (IETF) Security Issues in Network Event Logging (syslog) Work Group is actively involved in efforts to document the Syslog protocol, to identify the security and integrity problems, and to develop a standard to address these problems.

accommodates these kinds of requirements.[3] Thus, some background information on Syslog is in order before proceeding with the other requirements.

Syslog is an extremely simple utility and protocol to exchange log messages, the latter being defined in IETF RFC 3164, which is an informational RFC. The term "Syslog" is often used for the protocol, the tools that send the logs (*syslogd*), as well as the individual logs and the log files themselves.

By default, Syslog messages are sent over UDP port 514 between:

- *devices* – machines that can generate messages
- *relays* – machines that can receive messages and forward them to other machines
- *collectors* – machines that receive messages and do not relay them to any other machines (a.k.a. Syslog server)

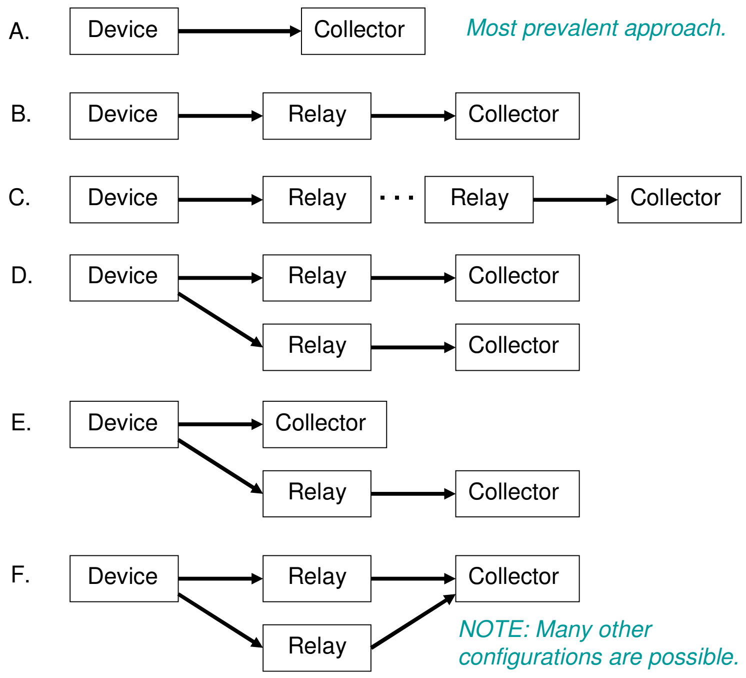

Most prevalent approach.

NOTE: Many other configurations are possible.

Figure 1. Syslog Architectures

Per RFC 3164, the Syslog architecture can be summarized as follows:

- Senders (devices and relays) send messages to relays or collectors with no knowledge of whether it is a collector or relay.
- Senders may be configured to send the same message to multiple receivers.

[3] Proprietary solutions are also available that can be used as a Syslog replacement and/or augmentation.

- Relays may send all or some of the messages that they receive to a subsequent relay or collector. In the case where they do not forward all of their messages, they are acting as both a collector and a relay. In Figure 1, these devices are designated as relays.
- Relays may also generate their own messages and send them on to subsequent relays or collectors. In that case it is acting as a device. These devices are also designated as a relay in Figure 1.

Given this simplified description of Syslog, it is now possible to resume the discussion on the hypothetical requirements listed earlier. A second pass through these requirements is needed to determine an appropriate aggregation approach; however, the following factors must also be considered during this process:

- Number of unique audit log repositories required to support the compliance activities; for example, SOX and HIPAA related events might need to be handled separately.
- Assuring the integrity of the log information as it is a common practice for an attacker to either modify or destroy log information in an attempt to conceal their activities.
- Tiering of event data, based on criticality; for example, critical systems may require aggressive real-time analysis, so care has to be exercised in controlling the amount of log data the analysis engine has to grind through
- Special handling requirements for sensitive data. For example, logs that contain regulated information may require special transfer and storage capabilities like encryption. Another example is digitally signed log events for proof of integrity or non-repudiation.
- Mandated retention periods may force separate collectors to simplify the archival process; for example, all event logs that must be retained for 3 years could be sent to a single collector, which then uses an archive solution that is configure for the appropriate retention period.
- Establishing a balance between devices and collectors; device are not created equal, so it is important to understand the volume of log events that each device will produce.
- The need for multiple copies

At the end of this process, the total number of relays and collectors as well as the special handling needs can be determined. From a storage perspective, knowing where the log events need to be shipped, how the events need to be handled, and then implementing it appropriately represents the end point; the storage layer need only concern itself with plugging into this infrastructure and then transmitting the log events.

Continuing with the hypothetical requirements and assuming that the aggregation is basically out of the way, the next step is to determine the analysis and alerting approach to be implemented. The specifics may drive the use of relays and impact the sizing of the collectors. In addition, specialized tools and databases may be required to implement the necessary analysis.

Finally, the archival needs have to be addressed. As a precursor, the appropriate enterprise policy and legal authorities should be consulted to help determine specific requirements in this area. The guidance can vary significantly, depending on the market sector, criticality and sensitive of the IT and data, and other factors. The following sample requirements are provided for illustration purposes:

- All of the raw log data must be retained for extended periods of time
- Only the analysis data (e.g., snapshots of a database) need to be backed up
- Integrity verification and protection mechanisms (e.g., WORM technology, secure hashes, etc.) must be employed
- Log data must be encrypted to ensure confidentiality of sensitive information
- All log events must be on-line and accessible

As a final note, automation is crucial to a successful audit logging strategy, so automate everything to the maximum degree possible. If it requires a human in the loop, it won't get done on a consistent basis.

SNIA Event Logging Recommendations for Storage

Event logging is an important area in which storage networks and storage resources can directly participate and contribute to an organization's efforts to achieve regulatory compliance. Further, the storage technology has an intimate relationship with the organization's data (the "crown jewels"), so implementing adequate event logging is critical. Therefore, SNIA has developed the following recommendations, which are applicable for all aspects of storage:

Recommendation #1: Include Storage in Logging Policy

In many organizations, audit logging is explicitly addressed in security policies. These policies may include the types of systems and device that are required to participate in the organization's audit logging implementation, the mandatory types of data to be collected, protection measures, retention periods, reviews, and archival. It is important to ensure that these policies extend to the storage layer.

With regard to storage systems and devices, the following elements of policy should be addressed:

- Storage systems and devices are required to participate in audit logging
- All ***significant*** storage management events will be collected
- Log data will be preserved
- Log data will be archived and retained
- The device time will be synchronized on a reliable, external source

Recommendation #2: External Event Logging

For anything other than system health monitoring and debugging, device resident logs are not recommended because they can be subjected to tampering or destruction, they have serious size restrictions because of the limited storage space available for logs, and they preclude the use of centralized automated analysis, alerting, and archiving. Therefore, storage devices must be capable of natively logging events to one, and preferably multiple, external log servers. Further, these devices should be capable of using Syslog to transfer log event; optionally, other protocols and mechanisms can be provided to address situations where Syslog is not an appropriate solution.

In addition, audit logging for which compliance, accountability, and/or security serve as the primary drivers must have devices configured to log events on a transactional basis. This means that devices are not permitted to buffer up multiple events before they are transmitted. The reason is that a buffered implementation runs the risk of losing event data when an unflushed buffer is lost or corrupted due to a system error or a malicious act.

Recommendation #3: Complete Event Logging

Not all events are created equal and it is likely that they need to be processed, based on their relative importance. Generally speaking the log event can be categorized as one of the following:

- *control* – messages that document events that provide status, warn of pending failures, identify network connections, etc.
- *management* – messages that represent explicit actions to change the configuration, move data, invoke security measures (e.g., authentication, account management, etc.).
- *data access* – messages that document transactions associated with data; for example, when a particular host accesses a LUN or a NAS users attempts to access a specific file or directory.

From an accountability perspective the management events are always of interest, the data access events are usually of limited interest (except in situations where critical files and directories need to be tightly monitored), and control events are typically of the least interest (they can provide useful information during root-cause analysis after an incident). Once the types of events to be logged have been determined, then ALL occurrences of these events must be logged. For storage resources, this means that in-band and out-of-band as well as remote (e.g., management application running on a server) and local (e.g., the management console of a storage array) events must be handled correctly and consistently.

Auditors typically check to see if the following kinds of events are being logged:

- Failed logon attempts
- Failed file and object access attempts
- Account and group profile additions, changes, and deletions
- Changes to system security configurations (e.g., zoning changes)
- System shutdown and restarts
- Privileged operations
- Use of sensitive utilities
- Access to critical data files

Ideally, each log entry will include a timestamp[4] (date and time), a severity level, the source of the log entry (distinguishing name, IP address, etc.), and a description of the event. Log entries that are cryptic are of limited value for compliance monitoring activities unless a secret decoder ring (i.e., documentation explaining the log entries) is made available.

Finally, log entries are often designated with a severity code that can be used as a basis for determining which events are logged. For illustration purposes, Table 1 shows the severity levels defined within Syslog. Assuming a Syslog based implementation for IT security infrastructure, it is a common practice to capture log entries that have been categorized as "Informational" and above (see Table 1) to help detect malicious activities; in other words, only "Debug" messages are filtered out of the audit logs. For storage, it is unlikely that such a detail level needs to be captured in the audit logs. However, it is important to capture all of the key information (especially when accountability needs to be monitored and enforced), so the enterprise logging policy should serve as the guide for determining what kind of filtering is appropriate and what level of information requires long term storage.

Numerical Code	Severity
0	*Emergency* – system is unusable
1	*Alert* – action must be taken immediately
2	*Critical* – critical conditions
3	*Error* – error conditions
4	*Warning* – warning conditions
5	*Notice* – normal but significant condition
6	*Informational* – informational messages
7	*Debug* – debug-level messages

Table 1. Syslog Message Severities

Recommendation #4: Protection

Depending on the importance of the data and/or the type of log entries, protective measures may be required to ensure the confidentiality and integrity of the log data. This protection can include securing the transfer of data from the device to the collector[5], additional hardening of the log servers, using secure hashing algorithms for integrity checks, employing digital signatures for non-repudiation, and even encryption of the log data.

In some environments, special purpose log servers are often deployed to handle these unique requirements. This approach allows the enterprise to extract the maximum benefit from its investment in the underlying technology and to concentrate its protective measures where they are needed. For example, an enterprise may require certain security events to be written

[4] External time synchronization is frequently a requirement, but there may be very specific requirements for certain market sectors (e.g., SEC regulated, NYSE regulated, etc.)

[5] A variety of mechanisms (e.g., SSH, SSL/TLS, IPsec) exist and are in use to secure the log events while they are in-flight. It is also worth noting that the IETF has defined two mappings of the Syslog protocol to TCP connections (in RFC 3195), both useful for reliable delivery of event messages.

to WORM technology; by culling out just these events on relays and then forwarding them to a special log server that uses WORM disk, the enterprise meets its requirements and minimize the use of the WORM technology to just those events that need this protection.

Recommendation #5: Retention

Typical system administration procedures call for retaining logs for up to 90 days, with weekly rotations, at the device level. For compliance purposes, this is a recipe for problems when the external auditors arrive on the scene. Use policy to establish the retention requirements (don't rely on system administrators to read the law), inventory the data and systems (know what you have and where it is), and make sure the event log data from the affected systems are handled and retained correctly.

How the event log data is retained is also important. For example, having 1300 tapes stored in a closet (and requiring one hour to read each one) is not going to be an acceptable approach when the organization is confronted with a 48-hour deadline to produce information. When dealing with log events that are associated with accountability, there may be a need to access and analyze this information very quickly, so an appropriate storage technology should be selected.

Disclaimers

This document is provided for informational and planning purposes only. The information used in compiling this document was obtained from publicly available sources and no representation is made as to the accuracy of the information, or as to the accuracy of any reading or interpretation thereof. No warranty is made or implied regarding the usefulness or suitability of this information for a particular purpose. Further, the Storage Networking Industry Association (SNIA), Hitachi Data Systems (HDS), and/or any contributors are not liable for any damages, real or consequential, arising from use of this information.

The information contained in this whitepaper is for general guidance on matters of interest only. The application and impact of laws can vary widely based on the specific facts involved. Given the changing nature of laws, rules and regulations, and the inherent hazards of electronic communication, there may be omissions or inaccuracies in information contained in this whitepaper. Accordingly, the information in this whitepaper is provided with the understanding that the authors, publishers, and contributors are not herein engaged in rendering legal or other professional advice and services. As such, it should not be used as a substitute for consultation with professional legal or other competent advisers. Before making any decision or taking any action, we recommend you consult a lawyer if you want professional assurance that our information, and your interpretation of it, is appropriate to your particular situation.

Appendix

Additional Sources of Information

- Information Security Forum (ISF), *The Standard of Good Practice for Information Security*, Version 4.1, January 2005, http://www.thestandard.org
- Internet Engineering Task Force (IETF), Security Issues in Network Event Logging (syslog) Work Group, http://www.ietf.org/html.charters/syslog-charter.html
- Internet Engineering Task Force (IETF), RFC 3164, The BSD syslog Protocol
- Internet Engineering Task Force (IETF), RFC 3195, Reliable Delivery for syslog
- Internet Engineering Task Force (IETF), RFC 1305 - Network Time Protocol (Version 3) Specification, Implementation and Analysis
- U.S. Code of Federal Regulations (CFR) Title 45 Parts 160, 162, and 164; *Health Insurance Reform: Security Standards*
- Sarbanes-Oxley Act of 2002, Section 802
- Federal Financial Institutions Examination Council (FFIEC) Information Security IT Examination Handbook
- National Institute of Standards and Technology (NIST), Special Publication 800-53, *Recommended Security Controls for Federal Information Systems*, February 2005
- ISO/IEC 17799:2005, *Information Technology–Security Techniques–Code of Practice for Information Security Management*, June 2005
- Information Systems Audit and Control Foundation, IT Governance Institute, *COBIT: Control Objectives for Information and related Technology*, 3rd Edition, July 2000
- BITS, *BITS Framework for Managing Technology Risk for IT Service Provider Relationships*, November 2003

About the Security Technical Work Group

The Storage Networking Industry Association's Security Technical Work Group (TWG) provides a focus for defining methods of increasing the security of both information residing within (or transmitting through) storage networks and information related to the management of those networks. The Security TWG:

- Defines requirements for storage network security in collaboration with other groups
- Creates architectures, interfaces, and practices that make optimal use of existing security technologies in a storage network
- Creates, or simulates creation of, new information security technologies where nothing exists that meets the requirements for use in a storage network
- Develops educational material appropriate to security in storage networks

The above may include cryptographic protection of information itself, independent of the storage or storage network system.

The scope of work encompasses storage and storage networking security for installations from the departmental level to the multi-enterprise. Technologies of interest include Fibre Channel and IP storage networks, NAS, CAS, objected-based storage, storage devices, hosts, etc. The target audiences for the TWG's deliverables are the vendors and customers of advanced storage infrastructures.

The Security TWG contributes to the body of knowledge for storage security, provides materials for inclusion in storage certification exams, participates/collaborates with formal standards bodies, compiles storage threat models, and advocates best practices for interoperable security of stored customer information. Reference implementations and/or test suites may be produced to help advance the industry. Also, the TWG tracks the activities in, and provides advice to other SNIA teams in discussing, analyzing and improving the security aspects of their work.

About the SNIA

The Storage Networking Industry Association is a not-for-profit organization made up of more than 300 companies and individuals worldwide spanning virtually the entire storage industry. SNIA members share a common goal: to set the pace of the industry by ensuring that storage networks become efficient, complete and trusted solutions across the IT community. To this end, the SNIA is uniquely committed to delivering standards, education and services that will propel open storage networking solutions into the broader market.