CMSGu2012-05

National Computer Board

Mauritian Computer Emergency Response Team

Enhancing Cyber Security in Mauritius

Guideline on Auditing and Log Management

CERT-MU

National Computer Board
Mauritius

Version 1.1

July 2012

Issue No. 5

Table of Contents

DISCLAIMER: *This guideline is provided "as is" for informational purposes only. Information in this guideline, including references, is subject to change without notice. The products mentioned herein are the trademarks of their respective owners.*

1.0 Introduction

1.1 Purpose and Scope

This guideline aims at assisting organisations in understanding the need for sound computer auditing and log management and the best practices that need to be followed to meet existing challenges.

1.2 Audience

The target audience for this document includes computer security staff and program managers; system, network, and application administrators; computer security incident response teams; and others who are responsible for performing duties related to computer security audit and log management.

1.3 Document Structure

This document is organised into the following sections:

Section 1 contains the document's content, the targeted audience and the document's structure.

Section 2 presents a background on auditing and log management.

Section 3 gives details on computer security logs.

Section 4 details auditing and log management.

Sec*tion 5* explains how to overcome the existing challenges.

Section 6 concludes the document.

Section 7 contains a list of references that have been used in this document.

Appendix A defines a set of acronyms used in this document.

2.0 Background

A log is a record of the events that occur within an organisation's systems and networks. Logs are made up of log entries, each entry containing information related to a specific event that has occurred within a system or network. Initially, logs were used mainly for troubleshooting problems, but logs now provide many functions within most organisations, such as optimizing system and network performance, recording users' actions, and providing data useful for investigation in case of malicious activities.

Logs have advanced and now contain information related to many different types of events occurring within networks and systems. Within an organisation, many logs contain records related to computer security; common examples of these computer security logs are audit logs that track user authentication attempts and security device logs that record possible attacks.

Cybersecurity teams always drive the need for strong event auditing and log management in order to support their incident response process. This is particularly true at network level events with the deployment of network intrusion detection systems (NIDS). As the type, origin, and sophistication of attacks against computer networks has changed significantly, changes in techniques for auditing and logging for host, application, data store, user access control requires significant upgrades as well to improve network data monitoring to be able to detect advanced intrusion attempts.

3.0 Computer Security Logs

Logs can contain a wide variety of information on the events occurring within systems and networks. The following sections describe the different categories of logs.

3.1 Security Software Logs

Most organisations use several types of network-based and host-based security software to detect malicious activity, protect systems and data, and support incident response efforts. Accordingly, security software is a major source of computer security log data. Common types of network-based and host-based security software include the following:

- **Antimalware Software**

  The most common form of antimalware software is antivirus software, which typically records all instances of detected malware, file and system disinfection attempts, and file quarantines. Additionally, antivirus software might also record when malware scans were performed and when antivirus signature or software updates occurred. Antispyware software and other types of antimalware software (e.g., rootkit detectors) are also common sources of security information.

- **Intrusion Detection and Intrusion Prevention Systems**

  Intrusion detection and intrusion prevention systems record detailed information on suspicious behaviour and detected attacks, as well as any actions intrusion prevention systems performed to stop malicious activity in progress. Some intrusion detection systems, such as file integrity checking software, run periodically in order to generate log entries in batches instead of on an ongoing basis.

- **Remote Access Software**

  Remote access is often granted and secured through virtual private networking (VPN). VPN systems typically log successful and failed login attempts, as well as the dates and times each user connected and disconnected, and the amount of data sent and received in each user session. VPN systems that support granular access control, such as many Secure Sockets Layer (SSL) VPNs, may record detailed information about the use of resources.

- **Web Proxies**

  Web proxies are intermediate hosts through which Web sites are accessed. Web proxies make Web page requests on behalf of users, and they cache copies of retrieved Web pages to make additional accesses to those pages more efficient. Web proxies can also be used to restrict Web access and to add a layer of protection between Web clients and Web servers. Web proxies often keep a record of all URLs accessed through them.

- **Vulnerability Management Software**

  Vulnerability management software, which includes patch management software and vulnerability assessment software, typically logs the patch installation history and vulnerability status of each host, which includes known vulnerabilities and missing software updates. Vulnerability management software may also record additional information about hosts' configurations. Vulnerability management software typically runs occasionally, not continuously, and is likely to generate large batches of log entries.

- **Authentication Servers**

  Authentication servers, including directory servers and single sign-on servers, typically log each authentication attempt, including its origin, username, success or failure, and date and time.

- **Routers**

  Routers may be configured to allow or deny certain types of network traffic based on a policy. Routers that block traffic are usually configured to log only the most basic characteristics of blocked activity.

- **Firewalls**

  Like routers, firewalls allow or deny activity based on a policy; however, firewalls use much more sophisticated methods to examine network traffic. Firewalls can also track the state of network traffic and perform content inspection. Firewalls tend to have more complex policies and generate more detailed logs of activity than routers.

- **Network Quarantine Servers**

  Some organisations check each remote host's security level before allowing it to join the network. This is often done through a network quarantine server and agents placed on each host. Hosts that do not respond to the server's checks or that fail the checks are quarantined on a separate virtual local area network (VLAN) segment. Network quarantine servers log information about the status of checks, including which hosts were quarantined and for what reasons.

3.2 Operating Systems

Operating systems (OS) for servers, workstations, and networking devices (e.g., routers, switches) usually log a variety of information related to security. The most common types of security-related OS logs are as follows:

- **System Events**

  System events are operational actions performed by OS components, such as shutting down the system or starting a service. Typically, failed events and the most significant successful events are logged, but many OSs permit administrators to specify which types of events will be logged. The details logged for each event also vary widely; each event is usually timestamped, and other supporting information could include event, status, and error codes; service name; and user or system account associated with an event.

- **Audit Records**

  Audit records contain security event information such as successful and failed authentication attempts, file accesses, security policy changes, account changes (e.g., account creation and deletion, account privilege assignment), and use of privileges. OSs typically allow system administrators to specify which types of events should be audited and whether successful and/or failed attempts to perform certain actions should be logged.

- **Information from security software and other applications running on the system**

  They are most beneficial for identifying or investigating suspicious activity involving a particular host. After suspicious activity is identified by security

software, OS logs are often consulted to get more information on the activity. For example, a network security device might detect an attack against a particular host; that host's OS logs might indicate if a user was logged into the host at the time of the attack and if the attack was successful. Many OS logs are created in syslog format. Other OS logs, such as those on Windows systems, are stored in proprietary formats.

3.3 Applications

Operating systems and security software provide the foundation and protection for applications, which are used to store, access, and manipulate the data used for the organisation's business processes. Most organisations rely on a variety of commercial off-the-shelf (COTS) applications, such as e-mail servers and clients, Web servers and browsers, file servers and file sharing clients, and database servers and clients.

Some applications generate their own log files, while others use the logging capabilities of the OS on which they are installed. Applications vary significantly in the types of information that they log. The following lists some of the most commonly logged types of information and the potential benefits of each:

- **Client requests and server responses**

  These can be very helpful in reconstructing sequences of events and determining their apparent outcome. If the application logs successful user authentications, it is usually possible to determine which user made each request. Some applications can perform highly detailed logging, such as e-mail servers recording the sender, recipients, subject name, and attachment names for each e-mail; Web servers recording each URL requested and the type of response provided by the server; and business applications recording which financial records were accessed by each user. This information can be used to identify or investigate incidents and to monitor application usage for compliance and auditing purposes.

- **Account information**

  This is primarily successful and failed authentication attempts, account changes (e.g., account creation and deletion, account privilege assignment), and use of privileges. In addition to identifying security events such as brute force password

guessing and escalation of privileges, it can be used to identify who has used the application and when each person has used it.

- **Usage information**

  This is mainly the number of transactions occurring in a certain period (e.g., minute, hour) and the size of transactions (e.g., e-mail message size, file transfer size). This can be useful for certain types of security monitoring (e.g., a ten-fold increase in e-mail activity might indicate a new e-mail-borne malware threat; an unusually large outbound e-mail message might indicate inappropriate release of information).

- **Significant operational actions**

  This is typically application startup and shutdown, application failures, and major application configuration changes. This can be used to identify security compromises and operational failures.

4.0 Auditing and log management

Log management can benefit an organisation in many ways. It helps to ensure that computer security records are stored in sufficient detail for an appropriate period of time. Routine log reviews and analysis are beneficial for identifying security incidents, policy violations, fraudulent activity, and operational problems shortly after they have occurred, and for providing information useful for resolving such problems. Logs can also be useful for performing auditing and forensic analysis, supporting the organisation's internal investigations, establishing baselines, and identifying operational trends and long-term problems.

4.1 The Challenges of auditing and log management

Most organisations face similar log management-related challenges, which have the same underlying problem: effectively balancing a limited amount of log management resources with an ever-increasing supply of log data. This section discusses the most common types of challenges, divided into three groups. First, there are several potential problems with the initial generation of logs because of their variety and prevalence. Second, the confidentiality, integrity, and availability of generated logs could be breached inadvertently or intentionally. Finally, the people responsible for performing log analysis are often inadequately prepared and supported. The sections below discuss each of these three categories of log challenges.

4.1.1 Log Generation and Storage

In a typical organisation, many hosts' OSs, security software, and other applications generate and store logs. This complicates log management in the following ways:

- **Many Log Sources**

  Logs are located on many hosts throughout the organisation, necessitating log management to be performed throughout the organisation. Moreover, a single log source can generate multiple logs, for example, an application storing authentication attempts in one log and network activity in another log.

- **Inconsistent Log Content**

  Each log source records certain pieces of information in its log entries, such as host IP addresses and usernames. For efficiency, log sources often record only the pieces of information that they consider most important. This can make it difficult

to link events recorded by different log sources because they may not have any common values recorded (e.g., source 1 records the source IP address but not the username, and source 2 records the username but not the source IP address). Each type of log source may also represent values differently; these differences may be slight, such as one date being in "MMDDYYYY" format and another being in "MM-DD-YYYY" format, or they may be much more complex, such as use of the File Transfer Protocol (FTP) being identified by name, e.g. "FTP" in one log and by port number, e.g. "21" in another log. This further complicates the process of linking events recorded by different log sources.

- **Inconsistent Timestamps**

  Each host that generates logs typically references its internal clock when setting a timestamp for each log entry. If a host's clock is inaccurate, the timestamps in its logs will also be inaccurate. This can make analysis of logs more difficult, particularly when logs from multiple hosts are being analyzed. For example, timestamps might indicate that event A happened 45 seconds before event B, when event A actually happened two minutes after event B.

- **Inconsistent Log Formats**

  Many of the log source types use different formats for their logs, such as comma-separated or tab-separated text files, databases, syslog, Simple Network Management Protocol (SNMP), Extensible Markup Language (XML), and binary files. Some logs are designed for humans to read, while others are not; some logs use standard formats, while others use proprietary formats. Some logs are created not for local storage in a file, but for transmission to another system for processing; a common example of this is SNMP traps. For some output formats, particularly text files, there are many possibilities for the sequence of the values in each log entry and the delimiters between the values (e.g., comma-separated values, tab-delimited values, XML).

To facilitate analysis of logs, organisations often need to implement automated methods of converting logs with different content and formats to a single standard format with consistent data field representations. Inconsistent log formats and data field representations also present

challenges to people reviewing logs, who need to understand the meaning of various data fields in each log to perform a thorough review.

Because most hosts within an organisation typically log some computer security-related information, often with multiple logs per host, the number of logs within an organisation can be quite high. Many logs record large volumes of data on a daily basis, so the total daily volume of log data within an organisation is often overwhelming. This impacts the resources needed to store the data for the appropriate length of time and to perform reviews of the data. The distributed nature of logs, inconsistent log formats, and volume of logs all make the management of log generation and storage challenging.

4.1.2 Log Protection

Because logs contain records of system and network security, they need to be protected from breaches of their confidentiality and integrity. For example, logs might intentionally or inadvertently capture sensitive information such as users' passwords and the content of e-mails. This raises security and privacy concerns involving both the individuals that review the logs and others that might be able to access the logs through authorized or unauthorized means. Logs that are secured improperly in storage or in transit might also be susceptible to intentional and unintentional alteration and destruction. This could cause a variety of impacts, including allowing malicious activities to go unnoticed and manipulating evidence to conceal the identity of a malicious party. For example, many rootkits are specifically designed to alter logs to remove any evidence of the rootkits' installation or execution.

Organisations also need to protect the availability of their logs. Many logs have a maximum size, such as storing the 10,000 most recent events, or keeping 100 megabytes of log data. When the size limit is reached, the log might overwrite old data with new data or stop logging altogether, both of which would cause a loss of log data availability. To meet data retention requirements, organisations might need to keep copies of log files for a longer period of time than the original log sources can support, which necessitates establishing log archival processes. Because of the volume of logs, it might be appropriate in some cases to reduce the logs by filtering out log entries that do not need to be archived. The confidentiality and integrity of the archived logs also need to be protected.

4.1.3 Log Analysis

Within most organisations, network and system administrators have traditionally been responsible for performing log analysis, i.e. studying log entries to identify events of interest. It has often been treated as a low-priority task by administrators and management because other duties of administrators, such as handling operational problems and resolving security vulnerabilities, necessitate rapid responses.

Administrators who are responsible for performing log analysis often receive no training on doing it efficiently and effectively, particularly on prioritization. Also, administrators often do not receive tools that are effective at automating much of the analysis process, such as scripts and security software tools (e.g., host-based intrusion detection products, security information and event management software).

Many of these tools are particularly helpful in finding patterns that humans cannot easily see, such as correlating entries from multiple logs that relate to the same event. Another problem is that many administrators consider log analysis to be boring and to provide little benefit for the amount of time required.

Log analysis is often treated as reactive, something to be done after a problem has been identified through other means rather than proactive, to identify ongoing activity and look for signs of impending problems. Traditionally, most logs have not been analyzed in a real-time or near-real-time manner. Without sound processes for analyzing logs, the value of the logs is significantly reduced.

5.0 Overcoming the challenges

Despite the many challenges an organisation faces in log management, there are a few key practices an organisation can follow to avoid and even solve many of these obstacles it confronts. The following four measures give a brief explanation of these solutions:

- **Prioritize log management appropriately throughout the organisation**
  An organisation should define its requirements and goals for performing logging and monitoring logs to include applicable laws, regulations, and existing organisational policies. The organisation can then prioritize its goals based on balancing the organisation's reduction of risk with the time and resources needed to perform log management functions.

- **Establish policies and procedures for log management**
  Policies and procedures are beneficial because they ensure a consistent approach throughout the organisation as well as ensuring that laws and regulatory requirements are being met. Periodic audits are one way to confirm that logging standards and guidelines are being followed throughout the organisation. Testing and validation can further ensure that the policies and procedures in the log management process are being performed properly.

- **Create and maintain a secure log management infrastructure**
  It is very helpful for an organisation to create components of a log management infrastructure and determine how these components interact. This aids in preserving the integrity of log data from accidental or intentional modification or deletion, and also in maintaining the confidentiality of log data. It is also critical to create an infrastructure robust enough to handle not only expected volumes of log data, but also peak volumes during extreme situations (e.g., widespread malware incident, penetration testing, vulnerability scans).

- **Provide adequate support for all staff with log management responsibilities**
  While defining the log management scheme, organisations should ensure that they provide the necessary training to relevant staff regarding their log management responsibilities as well as skill instruction for the needed resources to support log management. Support also includes providing log management tools and tool

documentation, providing technical guidance on log management activities, and disseminating information to log management staff.

6.0 Conclusion

Many logs within an organisation contain records related to computer security events occurring within systems and networks. Auditing and log management help to ensure that computer security records are stored in adequate detail for an appropriate period of time. Log management encompasses a few challenges, but if key practices are correctly followed, these challenges can be met.

7.0 References

- Infosecisland: **http://www.infosecisland.com**
- Guide to Computer Security Log Management, NIST
- Guidelines for Auditing and Logging, CERT-In