Forensic Logging Requirements

Dirk Brouwer RE RA and Perry Mertens MSIT

Based on experience we know that fraud investigations in the financial industry are often hampered by the poor quality of logging by IT systems. Especially because fraudsters are using new techniques like APT's (Advance Persistent Threats) and "anti-forensics" tooling. In general a forensic analysis of the logging should provide insight into who did what, where, when, how and with what result. In this article we share the bad as well as best practices we encountered with respect to logging and audit trails. Finally we propose a 6 steps approach on how to realize logging and audit trails that adequately facilitate forensic investigations. It will also help companies to strengthen their internal control environment and be better able to comply with the regulatory reporting requirements.

D. Brouwer RE RA
is Senior Audit Manager at NN Group.
dirk.brouwer@nn-group.com

P. Mertens MSIT
is Audit Supervisor at NN Group.
perry.mertens@nn-group.com

Introduction

The Association of Certified Fraud Examiners (ACFE) estimates that 14.3 percent of all internal fraud cases occur at financial institutions with an average loss of $258,000 per case ([Feigo8]). Many of these frauds are committed via IT systems. For the investigation of these frauds it is important that the investigators can make use of the correct logging and audit trails. However in practice forensic investigators are often hampered by weak application logging and audit trails. Although implementation of an adequate logging solution sounds easy, it proves to be rather complex. Complicating factors are:

- *The complexity of IT.* In general, a business process does not make use of just a single IT system. In practice several user applications are part of a process chain. And in addition many IT components and thus system software and hardware are involved. To mention a few: operating systems, databases, network components such as routers and firewalls, user access control systems etc. All of them (should) provide the right audit trail.
- *The sheer amount of data.* The amount of data that is transferred and processed is still growing rapidly due to bigger data centers, faster processors, faster networks and new technologies like cloud platforms, Big Data ([Univ]) and Internet of Things ([Höll14]). On top of this, every IT device and application generates log files. However, there really are no standards for how these logs present their data. As a result, an investigator either has to learn what the log files are telling him or develop technologies to normalize these logs into some common and useable format.
- *Latest developments to wipe traces or circumvent logging and detection.* Very old techniques used by hackers to frustrate forensic investigations are hard disk scrubbing and file wiping by overwriting areas of disk over and over again. Also the use of encryption (like Pretty Good Privacy) ([Geigo6]) and physical destruction of hardware were commonly used. However, nowadays, specialized so called "anti-forensic tooling" is available to try to manipulate logging remotely. The number is steadily growing and the techniques are getting more and more sophisticated. An additional complication factor is the hacker technique of APT's (Advanced Persistent Threats). In this way hackers spread their activities over a long period of time – a couple of months is not unusual – while making as little "noise" as possible. Aim of this technique is to stay under the radar screen and plan for a huge "hit and run" attack after the system is fully compromised.

In this article we will investigate to which logging requirements the IT systems in the financial industry

must comply to produce an audit trail that is adequate to support forensic investigations.

Note: this article is based on a paper for the graduation assignment of the Executive Program "forensic accounting expert" at the Erasmus University of Rotterdam.

Definition of the Problem

Many organizations are confronted with illegitimate actions via IT systems, like data leakage or fraudulent money transfers. According to the ACFE (Association of Certified Fraud Examiners) organizations suffer an average of more than 52 incidents of insider fraud annually. In such cases it is important to have a sound audit trail. However, while many organizations maintain access logs most of these logs are insufficient due to the following 3 limitations ([Geigo6]):

- *The logs are missing record and field-level data and focus solely on a given transaction.* Most existing logs only contain information at the transaction level, such as: Which users accessed which transaction at what time? In these cases, critical information is still missing. Such as "Which specific records and fields did the user access?' and "What did the user do with the data?"
- *Main existing systems fail to log read-only actions, leaving gaps in the records.* Most existing logs only record update activities. This leaves critical information about what was viewed, queried or simply accessed out of the audit trail entirely. In these cases, there is often no record of the points in time that information was accessed without being changed. This information is extremely important for preventing and investigating information leakage and data theft. Another area where this absence of information reveals significant gaps is in demonstrating access to private or privileged information.
- *Logs represent an incomplete view of activities that is often "hidden" across multiple systems and difficult to correlate.* Many logs are maintained in separate systems or applications that don't "talk" to each other. This makes it difficult to find and correlate relevant information – or respond quickly to an audit request. This reality often aids the malicious insider in obscuring their activity.

Legacy systems that were developed a decade or two ago and even many newer systems were not designed for collecting detailed data access logs. Altering logging capabilities or introducing a logging mechanism to these applications frequently required the addition of a logging component to each online program. In a large enterprise, this can add up to tens of thousands of lines of code.

As a result most forensic investigations are hampered by the poor quality of the logging by IT systems. This is also evidenced by the input we received during our interviews with forensic investigators of three Dutch financial companies. Poor logging makes it very difficult for investigators to determine what actually happened and to trace transactions back to natural persons.

Next to a poor design of the logging, the quality of the log data can also be affected by the use of so-called "anti-forensic" ([Kedz]) tooling. Hackers make use of these tools when attempting to destroy any traces of their malicious actions.

Objective and Scope

The objective of this article is to investigate which logging requirements IT systems in the financial industry must comply with to produce and safeguard an audit trail that is adequate to support forensic investigations. This includes the legal requirements in the Netherlands.

The IT controls to prevent fraud are out of scope for this article.

Approach

For this article we have used the following approach:

- Interviewing forensic staff of three Dutch financial institutions about the problems they encountered during special investigations with respect to the implemented logging and with respect to the company policies
- Study of regulatory logging requirements for the financial industry in the Netherlands as prescribed in:
  - Wet Financieel Toezicht
  - Toetsingskader DNB 2014
- Study of best practices and industry standards regarding logging specifics:
  - ISO 27001 Information security management systems (ISO = International Standardization Organization)
  - COBIT Assurance Guide 4.1 (COBIT=Control Objectives for Information and related Technology)
  - PCI Data Security Standard (PCI = Payment Card Industry)
  - ISF Standard of Good Practices (ISF = International Security Forum)

- Analysis of logging
- Drawing a conclusion

Regulations

Financial institutions in the Netherlands have to comply with the Financial Supervision Act (Wet op het Financieel Toezicht) and/or the Pension Act (Pensioenwet). Pursuant to these acts, the Dutch Central Bank (DNB) holds that financial institutions have to implement adequate procedures and measures to control IT risks. These risks relate to, among other things, the continuity of IT and the security of information. In this context, 'adequate' means proportionate: the measures and procedures should be in line with the nature of the financial institution and the complexity of its organizational structure. The procedures must be in conformity with generally accepted standards (good practices). Examples of such standards are Cobit ([ITGI07]) and ISO27000. These standards provide for measures which are, in principle, considered adequate by DNB.

In order to test the security of information within financial institutions, DNB has developed an assessment framework consisting of a selection from Cobit. DNB developed this framework, which comprises 54 COBIT controls, in 2010. The framework was updated in 2014 and split into a "questionnaire" and a "points to consider" document.

The DNB assessment framework expects the financial companies to implement a logging and monitoring function to enable the early prevention and / or detection and subsequent timely reporting of unusual and/or abnormal activities that may need to be addressed. To this aim DNB refers to the following COBIT requirements ([DNB], "Points to Consider"):

- Enquire whether and confirm that an inventory of all network devices, services and applications exists and that each component has been assigned a security risk rating.
- Determine if security baselines exist for all IT utilized by the organization.
- Determine if all organization-critical, higher-risk network assets are routinely monitored for security events.
- Determine if the IT security management function has been integrated within the organization's project management initiatives to ensure that security is considered in development, design and testing requirements, to minimize the risk of new or existing systems introducing security vulnerabilities

Some more detailed requirements are included regarding logging. But these requirements are not focused on applications (software) but more on the IT infrastructure level and in particular on components that are aimed at protecting against cybercrime ([Glen11]) (like firewalls ([Ches03])).

Company Policies

From a regulatory perspective it is not obliged to have a specific policy on logging. However the DNB Self-Assessment Framework ([DNB]) expects financial companies to develop a log retention policy. It is considered a best practice to include guidelines on logging in a Security Monitoring Policy. In addition it is advised to establish a Code of Conduct for internal staff and have them sign that they have read this code. The code should specify among others the actions that staff should refrain from like unethical behavior such as browsing non-business related data on the Internet. Staff should be made aware that violations will be logged and the log data can be used during special investigations.

From our interviews with the staff of 3 Dutch financial companies we learned that all companies have a Code of Conduct that is to be signed by the staff. In addition, logging requirements are addressed in internal policies but only at a very high level. In practice logging requirements are specified per application for most of the business critical applications and for internet facing applications. Data integrity protection requirements (like hashing to be able to detect unauthorized changes) are not specifically set for audit trails. In addition the policies do not prescribe that logging should be stored in a centralized location. In practice most logging for laptops and workstations is stored locally and this very much hampers the forensic investigations.

Best Practices and Industry Standards

For this article we have studied the following best practices and industry standards:

- ISO 27001 Information security management systems (ISO = International Standardization Organization)
- COBIT Assurance Guide 4.1 (COBIT=Control Objectives for Information and related Technology)
- PCI Data Security Standard (PCI = Payment Card Industry)

- ISF Standard of Good Practices (ISF = International Security Forum)

ISO and COBIT are well known industry standards as also evidenced by the fact that DNB is specifically referring to these standards (see previous paragraph). ISO 27001 Information security management system provides some basic input related to logging requirements in section A10.10 Monitoring. Although the standard specifies that audit logs should be produced and kept for an agreed period to assist in future investigations not many details are provided about such requirements.

Several sections in COBIT 4 address logging and audit trail related information and details. Compared to other examined standards, COBIT gives specific focus to audit trails with respect to IT management processes such as configuration management, change management and backup processes. In addition at a general level attention is paid to logging related to hardware and software components.

PCI ([PCI15]) is a standard for credit card transactions. In our view PCI contains a best practice for logging that is to be provided by individual applications because it specifies:

- That the audit trails should link all access to system components to an individual user.
- Which events should be logged related to system components.
- Which details should be included in the event log.
- Time synchronization requirements be able to compare log files from different systems and establish an exact sequence of events.
- The need to secure audit trails so they cannot be altered.
- The need to retain audit trails so that investigators have sufficient log history to better determine the length of time of a potential breach and the potential system(s) impacted.
- Review requirements, since checking logs daily minimizes the amount of time and exposure of a potential breach.

However, the standard would further improve by explicitly requiring that:

- the logging specifies which specific records and fields the user accessed and what the user did with the data. For example in the case of adjusting data, both the content before the change and after the change should be recorded.
- the logging also includes information about read-only actions on highly confidential data.

Along with this, the ISF Standard ([ISF14]) of Good Practices provides a good addition to the PCI since it provides a more holistic view. ISF contains best practices about logging procedures as well as the kind of information systems for which security event logging should be performed. The ISF prescribes that security event logging should be performed on information systems that:

a. are critical for the organization (e.g. financial databases, servers storing medical records or key network devices)
b. have experienced a major information security incident
c. are subject to legislative or regulatory mandates

In our view it would be best to subdivide “ad a)” in:

a1. are critical for the organization (e.g. financial databases or key network devices)
a2. contain privacy related data (e.g. medical records or other confidential customer and personal data)

Bad Practices

During the interviews with financial institutions the following bad practices / common mistakes were identified;

a. *Lack of time synchronization.* It is crucial for forensic analysis in the event of a breach that the exact sequence of events can be established. Therefore time synchronization technology should be used to synchronize clocks on multiple systems. When clocks are not properly synchronized, it can be difficult, if not impossible, to compare log files from different systems and to establish the sequence of events. For a forensic investigator the accuracy and consistency of time across all systems and the time of each activity is critical in determining how the systems were compromised.
b. *Logging is not activated.* For performance reasons the logging function is mostly deactivated by default. So unless the logging is specifically activated no events are logged. In case of Internet applications it has to be noted that both the web application and the web server logging are needed and have to be activated separately.
3. *Logging is only activated in production systems.* For performance and cost cutting reasons logging is often not activated in development, test and acceptance environments. This enables hackers to attack these systems without being noticed. After successfully compromising one of these systems the hacker can launch an attack on the production system via the network. Or simply use retrieved userid/password combinations as

they are often the same on production and non-production systems.

d. *Logging is activated but log data gets overwritten.* Most of the time a maximum storage space capacity is defined for logging data. When this capacity is reached the system starts to write the next log data from the beginning of the file again. Thus the old log data is overwritten and gets lost if it has not been backed up in the meantime.
e. *Logging is activated but log data is over-detailed.* Log data needs to be tuned. If no filtering is applied the log file is overloaded with irrelevant events. This hampers efficient analysis in case of forensic investigations.
f. *Application logging requirements are not defined.* During development of applications the asset owner as well as the developers often overlook defining the log requirements.

Logging Requirements

A forensic analysis of the logging should provide insight into who did what, where, when, how and with what result. Therefore the logging should contain complete data and be kept for a sufficient period of time in a secured manner. After studying the logging requirements listed by the regulations and best practices used for this article (see appendices B to F), we have divided the logging requirements into 5 categories that are relevant from a forensic investigations point of view:

a. Retention requirements
b. Correlation requirements
c. Content requirements
d. Integrity requirements
e. Review requirements

Ad a) Retention requirements
The logging should be kept for a sufficient period of time. With the latest attack techniques (like APT = advanced persistent threat) hackers stealthily break into systems. As a result organizations are often compromised for several months without detection. The logs must therefore be kept for a longer period of time than it takes an organization to detect an attack so they can accurately determine what occurred. We recommend retaining audit logs for at least one year, with a minimum of three months immediately available online.

Ad b) Correlation requirements
A forensic analysis of the logging of IT systems must result in a chronological sequence of events related to natural persons. It is therefore important that logged activities can be correlated to individual users as well as to activities logged on other systems. These activities include business transactions as well as actions to get read access to sensitive data. As a business process chain contains several applications, the logging needs to be arranged for every application where users can modify transaction data or view sensitive data.

Because different applications running on different platforms each log a part of a transaction process, it is required to have a proper time synchronization between all computers. This is to ensure that timestamps in logs are consistent. To facilitate correlations, logging is best provided in a standardized format (like the Common Event Expression - CEE standard[1]).

Ad c) Content requirements
The main requirement content wise is that the logging provides information for all relevant events on who did what, where, when, how and with what result. Both application as well as infrastructure logging should provide this insight.

The PCI standard gives a good reference for application logging, but the precise log content should be decided on basis of individual applications and follow a risk based approach. In many cases this will result in the requirement to also log read-only actions, something that is now often forgotten.

The DNB assessment framework provides a good reference for infrastructure logging, especially regarding network components that are aimed to protect against cybercrime (such as firewalls).

Ad d) Integrity requirements
An attacker will often attempt to edit the audit logs in order to hide his activity. In the event of a successful compromise the audit logs can be rendered useless as an investigation tool. Therefore logging should be protected to guarantee completeness, accuracy and integrity.

Ad e) Review requirements
Security-related event logs should be analyzed regularly to help identify anomalies. This analysis is best done using automated security information and event management (SIEM) tools or equivalent. The analysis should include:

1 CEE™ is the Common Event Expression initiative being developed by a community representing the vendors, researchers and end users, and coordinated by MITRE. The primary goal of the effort is to standardize the representation and exchange of logs from electronic systems. See https://cee.mitre.org.

- processing of key security-related events (e.g. using techniques such as normalization, aggregation and correlation)
- interpreting key security-related events (e.g. identification of unusual activity)
- responding to key security-related events (e.g. passing the relevant event log details to an information security incident management team).

Analysis of Logging

Many breaches occur over days or months before being detected. Regularly checking logs minimizes the amount of time and exposure of potential breaches. If exceptions and anomalies identified during the log-review process are not investigated, the entity may be unaware of unauthorized and potentially malicious activities that are occurring within their own systems.

Therefore a daily review of security events is necessary to identify potential issues. This encompasses, for example, notifications or alerts that identify suspicious or anomalous activities. Especially in logs from critical system components and logs from systems that perform security functions, such as firewalls, Intrusion Detection Systems, Intrusion Prevention Systems, etc.

Logs for all other system components should also be periodically reviewed to identify indications of potential issues or attempts to gain access to sensitive systems via less-sensitive systems. The frequency of the reviews should be determined by an entity's annual risk assessment.

Note that the determination of "security event" will vary for each organization and may include consideration for the type of technology, location and function of the device. Organizations may also wish to maintain a baseline of "normal" traffic to help identify anomalous behavior.

The log review process does not have to be manual. The use of log harvesting, parsing, and alerting tools can help facilitate the process by identifying log events that need to be reviewed. Many security event manager tools are available that can be used to analyze and correlate every event for the purpose of compliance and risk management. Such a tool sorts through millions of log records and correlates them to find the critical events. This can be used for *a posteriori* analysis such as forensic investigations. But also for real time analysis to produce dashboards, notifications and reports to *a tempo* prioritize security risks and compliance violations.

DNB recommends to deploy SIM/SEM (security incident management/security event management) or log analytic tools for log aggregation and consolidation from multiple machines and for log correlation and analysis. SIEM (Security Information and Event Management) is a term for software products and services combining both SIM and SEM.

SIEM software such as Splunk and Arcsight combine traditional security event monitoring with network intelligence, context correlation, anomaly detection, historical analysis tools and automated remediation. Both are a multi-level solution that can be used by network security analysts, system administrators and business users. But also by forensic investigators to effectively reconstruct many system and user activities on a computer system.

Conclusion and Recommendations

At first sight it seems relatively easy to arrange for adequate logging. However, going into more detail there is much more to it. Logging should provide information on who did what, where, when, how and with what result. Both application as well as infrastructure logging should provide this insight. The sheer amount of data, the complex IT systems, the large varieties in hardware, software and also logging formats themselves often hamper forensic investigations within financial institutions. It is therefore very difficult and time consuming to perform a forensic investigation on an IT system and populate a chronological sequence of events related to natural persons. Especially since in practice most logging does not contain all relevant data for forensic investigations.

The root cause for this is in our opinion twofold: lack of specific regulations and company policies and lack of priority for logging requirements during the development of new systems. In our view a best practice for achieving an adequate audit trail during system development would be:

- *Risk assessment.* Perform a risk assessment as the first step of the development phase and make an inventory of the key operational risks (including the risk of fraud and data leakage).
- *Key controls.* Define the key controls to mitigate the key risks to an acceptable level.

- *Event definition*. Related to the key controls, define which events should be logged and which attributes (details) should be logged per event to prove that the controls are working effectively. Have the event definition and previous steps reviewed by staff of Internal Audit and/or Special Investigations.
- *Logging*. Design the logging based on the defined events and attributes, taking into account the general logging requirements regarding retention, correlation, content, integrity and review (see "Logging Requirements").
- *Implementation*. Implement the system and perform a User Acceptance Test. This test should include testing the effectiveness of the key controls and the completeness of the logging.
- *Monitoring*. Periodically monitor the logging to help identify suspicious or unauthorized activity. Because of the sheer amount of data this is best done with an automated tool (SIEM solution).

Such an approach would not only facilitate forensic investigations. It will also help companies to strengthen their internal control environment and be better able to comply with the regulatory reporting requirements.

References

[Ches03] W.R. Cheswick, S.M. Bellovin and A.D. Rubin, *Firewalls and Internet Security: Repelling the Wily Hacker* (2nd ed.), 2003.

[DNB] DNB Self Assessment Framework, http://www.toezicht.dnb.nl/binaries/50-230771.XLSX

[Feig08] N. Feig, *Internal Fraud Still a Danger as Banks Adjust Strategy*, 31 January 2008, http://www.banktech.com/internal-fraud-still-a-danger-as-banks-adjust-strategy/d/d-id/1291705?

[Geig06] M. Geiger, *Counter-Forensic Tools: Analysis and Data Recovery*, 2006, http://www.first.org/conference/2006/program/counter-forensic_tools__analysis_and_data_recovery.html

[Glen11] M. Glenny, *DarkMarket: Cyberthieves, Cybercops and You*. New York: Alfred A. Knopf, 2011.

[Höll14] J. Höller, V. Tsiatsis, C. Mulligan, S. Karnouskos, S. Avesand and D. Boyle, *From Machine-to-Machine to the Internet of Things: Introduction to a New Age of Intelligence*, Elsevier, 2014.

[ISF14] ISF, *The Standard of Good Practice for Information Security*, https://www.securityforum.org/tools/sogp/

[ITGI07] IT Governance Institute, *COBIT 4.1 Excerpt*, 2007, https://www.isaca.org/Knowledge-Center/cobit/Documents/COBIT4.pdf

[Kedz] M. Kedziora, *Anti-Forensics*, http://www.forensics-research.com/index.php/anti-forensics/#index

[PCI15] PCI Security Standards Council, *PCI SSC Data Security Standards Overview*, https://www.pcisecuritystandards.org/security_standards/documents.php?agreements=pcidss&association=pcidss

[Univ] University Alliance, *What is Big Data?*, Villanova University, http://www.villanovau.com/resources/bi/what-is-big-data/#.VSAwR_msXuQ

[Wiki15] Wikipedia, *Security information and event management*, 2015, http://en.wikipedia.org/wiki/Security_information_and_event_management

About the Authors

D. Brouwer RE RA is a Senior Audit Manager at NN Group. He holds degrees in science, finance and audit. In his professional life he has long-standing IT audit experience (since 1988) and is combining this with business audit responsibilities since 2013.

P. Mertens MSIT is an Audit Supervisor at NN Group. Perry has long-standing experience in technical IT audit and is an expert in the area of IT security and IT forensics. He frequently publishes on LinkedIn about recent IT security developments.