



Legal and Technical Issues Concerning Evidence in Data Breach Cases

Published in 2012 PLI Privacy and Data Security Law Institute (Thirteenth Annual) Course Handbook

By Winston Krone, Esq.

March 2012

KIVU CONSULTING, Inc.
44 Montgomery Street
Suite 700
San Francisco, CA 94104
Tel: (415) 524-7320
Fax: (415) 524-7325
www.kivuconsulting.com
PI License #26798

Legal and Technical Issues Concerning Evidence in Data Breach Cases

by Winston Krone, Esq., Kivu Consulting, Inc.

March 2012

Introduction

Prior to California SB 1386¹, California's breach notification law, and the resulting flood of states' laws and regulatory interest, the investigation of data breaches was mainly an exercise in identifying and patching a vulnerability, plus trying (often in vain) to identify the perpetrator. When reviewing accounts of network breaches in the 1980s and 90s, the impression is often of a highbrow game of chess².

California SB 1386, and subsequent breach notification legislation, has radically changed the nature and tenor of responding to data breaches by introducing mandatory notification requirements and opening the door to significant regulatory fines and civil damages. This has affected the way evidence in data breaches must be collected and treated.

The gathering of evidence has always been a significant step in analyzing the cause and extent of data breaches. However, the use of forensic evidence and methodologies (i.e. preserving data so that findings can be verified and authenticated in litigation³) has

¹ SB-1386, The California Security Breach Information Act, California Civil Code Sec. 1798.80-1798.82.

² See for example the classic account of an early network breach investigation in Eric Stoll's *The Cuckoo's Egg: Tracking a Spy through the Maze of Computer Espionage* (1989)

³ Digital forensics (also referred to as computer or network forensics) is considered the application of science to the identification, collection, and

grown in importance in the last couple of years. Factors driving this trend have included:

1. A realization that many network breaches do not automatically lead to unauthorized access of PII or PHI, and that forensic analysis can obviate the need for expensive notification and detrimental publicity
2. A growing trend for regulators to question the procedures used by organizations in determining the scope of a breach and the numbers of persons to be notified
3. A surge in civil claims following data breaches, with plaintiffs' attorneys being ready to argue that inadvertent loss of data during the initial breach response is spoliation of relevant evidence leading to sanctions and negative presumptions against the breached organization.

The role of evidence in data breach cases poses unique challenges, both technically and legally. Technically, the gathering of evidence is a crucial first step in analyzing the cause and extent of data breaches. Such evidence often consists of automatically generated logs and audits and is easily lost or overwritten if not collected immediately. In addition, the very configuration and file structure of infected machines can be crucial in determining the source and extent of an intrusion. This is exactly the type of evidence that is often lost as in-house IT teams struggle to get an organization safely back online. Finally, attackers have been known to specifically target and erase the evidence of their intrusion that was left on the compromised networks. This has all lead to IT teams adopting more advanced tools and methodologies to preserve evidence and analyze the data immediately following a data breach in a way that findings can be later verified and authenticated.

analysis of data while preserving the integrity of the information and maintaining a strict chain of custody for the evidence.

The Rise of Forensic Collection and Analysis of Evidence in Data Breaches

The use of forensic methodologies to gather evidence following a data breach has grown in importance in the last couple of years. Three factors driving this are set out below.

1. Data Breach –v- Network Breach

Senior executives and outside counsel have finally understood a fact that has been appreciated by IT personnel for many years – that there is a conceptual and quantifiable difference between a “network breach” and a “data breach”. Simply put, IT networks are commonly designed with the concept of “defense in depth” – network perimeters may fail without confidential data being put at risk. For example, parts of an organization’s network may be placed in a DMZ (demilitarized zone) where even a successful attack will not compromise the rest of the network. Honey pots and Glue Traps may be set up on the periphery or even within a network to deliberately route intruders into digital blind alleys to exhaust their efforts and allow defenders to observe the attack protocols. And important data may be sequestered within a network with its own firewalls, access controls, and encryption. However, current IT security does a poor job in addressing the critical question of proving the negative – i.e. that confidential data has not been accessed.

In the absence of a comprehensive federal data breach notification law, the majority of states have passed bills or introduced legislation to require businesses and/or government agencies to notify persons affected by breaches involving their sensitive personal information. Most states model their definitions off California’s SB1386 (e.g. a security breach is the unauthorized acquisition of computerized data that compromises the security, confidentiality, or integrity of personal information). In some states, even if a breach occurs, notice is not required where there is no reasonable likelihood of harm to the persons whose information

is compromised⁴, or if the breach is not material⁵. However, given the vagueness of the term “unauthorized acquisition”, organizations have been reluctant to argue that a compromised network has not lead to some form of compromise of their data’s security, confidentiality or integrity, unless there has been extremely strong evidence that the data could not possibly have been accessed by intruders or otherwise exposed. Current IT infrastructure (and the post-breach response plans, including evidence collection) is often not designed to give this level of certainty to an organization following a data breach.

Examples of such shortcomings include:

- The failure to audit last access dates and times of critical folders and files
- Being able to identify the (hopefully limited) individuals accessing confidential databases during set time periods
- Maintaining access or security event logs for too short a period⁶ – so that IT can only confirm the absence of compromise during a prior relatively short period (when the intrusion might be found to have taken place outside of that time period)
- Reliance on monitoring systems that only can show exfiltration of data from a network, when the applicable breach notification law focuses on “acquisition” and “compromise”, not actual physical theft or removal

While IT departments might be confident that no data has been taken or actually compromised, it is an awkward moment when the

⁴ States with such provisions include Arkansas, Florida, Louisiana, Montana, New Jersey, North Carolina, and Rhode Island.

⁵ Montana and Nevada.

⁶ See Appendix A for a list of suggested log file retention times.

organization's general counsel explains that the absence of evidence is not necessarily evidence of absence.

2. The Skeptical Regulator

There is a growing trend for regulators to question the procedures used by organizations in determining the scope of a breach and the numbers of persons to be notified, and a willingness for regulators to second guess the correct level of information security in a particular case. For example, in May 2011, the Federal Trade Commission ("FTC") settled with Ceridian Corporation following a data breach that had exposed over 28,000 customer employee records⁷. It is noticeable that the FTC went so far as to claim that Ceridian failed to take "readily available, free or low-cost defenses" against reasonably foreseeable SQL injection attacks, and also criticized Ceridian's storage of PII in (unencrypted) clear, readable text. Under the terms of the settlement, Ceridian has implemented a comprehensive information security program and is subject to independent audits every other year for 20 years. In a more recent case, the Attorney General of Minnesota brought a case against Accretive Health, Inc. in what is believed to be the first state case alleging violations of the Federal HIPAA/ HITECH Act⁸. Accretive, acting as a business associate to two Minnesota-based hospitals, was accused of losing an unencrypted laptop containing the patient data of over 23,000 Minnesota residents. In particular, the lawsuit, filed in the United States District Court in Minnesota, alleges that Accretive violated HIPAA by, among other things, failing "to identify and respond to suspected or known security incidents and to mitigate, to the extent practicable, harmful effects of security incidents that were known to them in violation of 45 C.F.R. § 164.308(a)(6). Tellingly, according to the complaint in the Accretive case, Accretive carried out its own

⁷ The FTC's complain is at
<http://www.ftc.gov/os/caselist/1023160/110503ceridiancmpt.pdf>

⁸ A copy of the complaint is at
<http://www.ag.state.mn.us/PDF/Consumer/AccretiveHealth20120119.pdf>

analysis of the number of patients whose data might have been lost and supplied this number to the two hospitals. However, at least one of the hospitals later retained its own computer expert “to undertake an independent forensic investigation to corroborate the representations of Accretive about the nature, scope, and extent of the lost data as it relates to [the hospital’s] patients.” According to the complaint, this expert “discovered an additional 6,690 patients whose names and data were believed to be on the laptop but who were not revealed to be on the laptop by Accretive.”

Clearly, corporations and organizations will be under growing scrutiny as regards the quality and extent of their technical response to data breaches. In fact, the ability of organizations to present their findings to government agencies in a verifiable format and without destroying valuable clues was specifically identified by the National Institute of Standards and Technology⁹ as one of several reasons advocating the forensic collection of digital evidence in data breach investigations. As demonstrated by the recent settlement between the Department of Health and Human Services, Office for Civil Rights, and BlueCross BlueShield of Tennessee for \$1.5m¹⁰, a thorough technical investigation of the cause and consequences of a data breach will be seen as a minimum requirement by regulators.

3. The threat of civil litigation

So far, consumer lawsuits resulting from data breaches have generally been dismissed due to the plaintiffs' inability to allege cognizable harm/damages. However, this is probably a false comfort for organizations and their attorneys dealing with post-breach situations. Even if consumer/ plaintiff cases fail, the next wave of post-breach litigation is likely to involve contractual and

⁹ Guide to Integrating Forensic Techniques into Incident Response - Recommendations of the National Institute of Standards and Technology, NIST Special Publication 800-86 (2006)

¹⁰ See http://www.hhs.gov/ocr/privacy/hipaa/enforcement/examples/resolution_agreement_and_cap.pdf

tort claims from damaged business partners, and insurance companies vigorously resisting claims brought under general and cyber-risk insurance policies. In both types of cases, key evidence will relate to whether a breached organization was negligent or had otherwise fallen below minimum standards or its contractual obligations as it related to information security. In these cases, the effected organization will need to collect and preserve a full picture of its security infrastructure at the time of the breach, even as its IT personnel dismantle, take off-line, and reinstall large parts of its network.

Evidence and Forensic Analysis

What are the goals of a forensic analysis in a data breach?

Data breach investigations are initiated with the identification, collection and assessment of potentially relevant digital evidence. However, different types of evidence will be relevant for each of the goals of an investigation (e.g. determining if a breach has occurred; determining the extent of the loss; preserving evidence/ preparing findings for regulators/ law enforcement/ civil litigation). Significantly, first responders to a potential data breach (often the organization's own IT personnel) are frequently only focused on the initial stages – determining if a data breach has occurred and remediating the cause of the breach. However, from the beginning of an incident, an organization must keep in mind potential subsequent stages – e.g. quantifying the amount of data that has been lost (and number of persons to be notified) and being able to present a cogent explanation to regulators and other interested parties of the organization's response to the incident. Evidence relevant to these later stages is frequently lost in the first 48 hours of an incident because the incident response team is unaware of the significance of such evidence (or is overwhelmed by simply responding to or technically remediating the breach).

Determining if a data breach took place

To establish whether a loss occurred, forensic experts rely on sources of digital evidence that address four key questions: who, what, where and when.

Determining whether a breach has occurred is an analysis function that is singularly dependent on the amount of available relevant evidence. Network breaches are usually identified by one of three events - the tripping of a specific monitoring or network intrusion device; the detection of an unexplained anomaly within the network (e.g. unusual traffic flow); or a full or partial failure of the network that was a direct (if unintended) consequence of the intrusion. In any of these events, relevant evidence might be unavailable for a successful analysis due to:

- There is insufficient logging or monitoring of the network. Much auditing of internal access (i.e. being able to tell who has accessed data and when) is not turned on – IT departments may not have the budget or resources to monitor or store such data. However, without such records, hopefully going back weeks before an incident is identified, organizations are not able to form baselines for “normal” network activity or be able to “prove the negative” of conclusively determining that certain data has not been accessed without authorization. In addition, internal network traffic logs can be critical in establishing that there was no access from a compromised server to a separate server containing PII/ PHI.
- Much of the relevant evidence will be automatically generated logs and system audits. Such data is often quickly overwritten (it is not unusual for even important access logs to be overwritten after 24 hours). If only minimal logs are being maintained, these may be overwritten in the time it takes for a recognized anomaly to be identified as an actual network breach.

- The totality of evidence can come from different types of monitoring and logging. For example, relevant evidence may include logs from firewalls, detection intrusion systems, servers and applications related to critical business systems (e.g. access to payroll, finance or customer databases) such as payroll systems; logs from third-party integrated systems such as payroll, banking or uploads for third-party data processing; and logs or digital records of physical access. Such cumulative evidence is most powerful when the extent and time period for which such different logs are kept is synchronized.
- Cash-strapped IT departments, tasked with maintaining business continuity, may react to a network breach by wiping and reformatting potentially compromised servers and computers. Even if they are aware of the need to preserve evidence (such as the configuration and file structure of infected machines), they may rely on backups should that evidence be required. However, backups and disaster recover/archiving systems usually do not capture critical system logs (e.g. records of security events or attempts to access the network itself). In addition, a forensic analysis will include reviewing deleted data¹¹ that is inevitably lost during the reformatting process and not captured by backups.

Quantifying the loss of PII/ PHI

The main steps to quantifying the loss include:

1. An analysis of the data that has been compromised (assuming it still exists) including a complete collection of

¹¹ Examples of deleted data that are frequently relevant in forensic analysis of data breaches include fragments of malware that auto-delete once installed; deleted compressed files used to transmit data from systems compromised by key loggers; monitoring data deliberately deleted by hackers or network intruders to cover their tracks.

the data believed to be compromised that preserves its original metadata and location, together with all logs relating to such data.

2. A verifiable explanation of how the analysis was carried out, examples of sampling or other analytics used, and a critique of the reliability of the data (e.g. could hackers, in addition to stealing PII/ PHI, have also deleted data in order to give the impression that the theft was smaller or to delay a full detection? In such cases, should the organization compare its current data with backups to determine if data is not only accessed but missing?). Evidence will include the work product and the original data used for the analysis. In addition, internal emails or memos may be highly relevant in supporting (or contradicting) the findings.
3. An analysis of backups or similar computers (e.g. if the data is on a missing laptop), with a similar verifiable explanation and critique of the analysis. In cases of missing laptops, evidence could include not only copies of backups but also forensically preserved copies of similar media (for comparison purposes) and any logs or evidence showing downloads, access or usage of the missing media. Even where the missing laptop or storage device is believed to have been encrypted, evidence of installation and usage of the encryption should be identified and preserved.¹²

If a breach is determined to have taken place, the most pressing question becomes the number of individuals or businesses affected. Correctly and accurately quantifying the loss is critical to

¹² Although laptops can be installed with full disk encryption, frequently this will not be set to apply in hibernation mode. Thus an encrypted laptop, stolen while in hibernation mode, may arguably not have been encrypted in a way that would be a defense under many states' breach notification laws.

determining the response to the breach. Examples include referring to backups/ archives to determine how much data was stored on a compromised system; reviewing the metadata of files on the compromised system to determine how many files may have been accessed during the breach; and forensically determining if data containing PII/ PHI was deleted or modified during the breach. Also, the organization must determine a protocol for calculating the amount of PII/ PHI that may have been compromised. Unless an organization has already carried out (and maintained) an accurate review and audit of the sensitive data stored in its network, this will require a combined automatic and manual review of individual files, and sampling/ searching (using a protocol that must stand up to regulatory scrutiny) – all at a time when the clock is ticking down for the deadline to notify those effected by the breach.

Some Legal Implications of Evidence in Data Breaches

What evidence could be discoverable?

The general rule is that anything relevant and non-privileged can be discoverable (Fed. R. Civ. P. Rule 26(b)(1)). Litigation subsequent to a data breach is likely to focus on whether:

- An organization was negligent in planning or implementing its IT security
- An organization had failed to purge redundant data or had failed to use available data storage security such as encryption
- An organization's response to the actual breach did not properly follow its own response plan or failed industry best practices
- The decision to notify (or not notify) specific persons that their PII or PHI might have been compromised was based on faulty or incomplete analysis

Spoliation of evidence

Clearly the types of evidence that are relevant in most civil cases (e.g. email, employee created documents, third-party reports) are also likely to be relevant in such litigation. However, unique types of evidence might include automatically created log files, network configurations, and metadata. As discussed above, such evidence is easily lost during the first 24 hours of a data breach.

Whether or not the loss of such data amounts to spoliation depends on whether the duty to preserve has been triggered by the actual or reasonable anticipation of litigation. Regulatory action and/ or civil litigation are almost inevitable once an organization has made a public notification of a breach. It is therefore arguable that an organization has a duty to preserve as soon as it realizes it may have lost PII/ PHI data in a breach. In the context of preservation by the incident response team during the initial reaction to a data breach, it will be interesting to see how the courts judge preservation decisions. In *Pension Committee*¹³, Judge Scheindlin described this hindsight review as “a judgment call” where the court will employ “ ‘a gut reaction’ based on years of experience as to whether a litigant has complied with its discovery obligations and how hard it worked to comply.” The court in *Pension Committee* found that a failure to issue a written legal hold notice per se constituted gross negligence. Could it also be argued that an incident response plan (often drafted with the assistance of counsel) that does not contain similar preservation obligations on the IT/ incident response team is insufficient given the foreseeable potential for evidence loss/ destruction?

Authentication of Evidence in post-breach litigation

As referenced above, data breach gives rise to unique forms of potentially relevant technical evidence. Such evidence may also

¹³ *Pension Committee of the University of Montreal Pension Plan v. Banc of America Securities, LLC*, 2010 WL 184312 (S.D.N.Y. Jan. 15, 2010)

provide significant authentication issues. In the case of computer generated evidence (such as network logs), this may require someone who has personally knowledge of the origins of such logs and can identify them. The problem would be that outside forensic experts can't necessarily authenticate such logs if they didn't personally collect them from the network (which is often the case, given the collaborative nature of breach response). In practice, this may mean that the only person who can authenticate critical digital evidence is an organization's IT personnel. It would therefore be prudent for counsel to take near contemporary declarations from the relevant IT personnel, especially given the possibility of disciplinary action or employee departure following a serious network breach.

Evidence held by third parties

A significant evidentiary problem arises where an organization has outsourced some or all of its network infrastructure (or network logging) to a third-party. This is a particular concern in data breaches, given the use of cloud and off-shore storage of data. Critical evidence relating to access, security settings, and the amount of data in existence may only be held in the possession of such third parties. It is absolutely crucial that an organization not only confirms the contractual obligations on third-party vendors to preserve/ produce but also put in place the logistics of a notification process. While a legal grey area (with some judges being more flexible to the conundrum faced by organizations with outsourced data), the courts are not likely to accept that a U.S. based corporation can outsource its obligations to preserve discoverable data by a creative relationship with a third-party provider.

Appendix A

Log Files and Suggested Retention Times

Forensics is the application of science, not magic. If the evidence is gone, forensic analysis will be limited. Following an incident, it is highly advisable that the logs set out below (as a minimum) will be available, having been retained for the suggested times. Also remember: logs are favorite targets of hackers. Logs need to be securely stored, preferably outside the possible damage zone.

1) Logs from the border devices: firewall, VPN, Exchange/OWA, and any router that is the next hop out from the firewall (if controlled by the client and not the same appliance/device as the firewall appliance). These logs should be kept at least 90 days if possible, 6 months ideal.

2) Active Directory related: Object access (success/failure), Login/Logout (success/failure). These logs fill quickly, 90 days would be a good start.

3) Public facing servers: Web server logs, FTP logs (if running a FTP server), OWA logs, SSH logs, collective Linux syslogs or entire contents of /var/log (messages, wtmp, sshd, cron, etc.). These logs should be kept for 6 - 12 months.

4) Internal netflow logs: logs from managed switches, wireless access point logs, DHCP logs. These logs should be kept 6 months.

5) In addition, for any of these logs, forensic experts would like +72 hours of regular "non-compromised" activity to create a baseline and compare to the incident time-line. This is mainly an issue where logs are only being kept for a few days – in those cases, it may not be possible to determine what "normal" activity is since the entire log will reflect compromised activity.