



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Hacker Tools, Techniques, and Incident Handling (Security 504)"
at <http://www.giac.org/registration/gcih>

Breach Notification in Incident Handling

GCIH Gold Certification

Author: Jeffery Buffington, jeffery.buffington@gmail.com

Adviser: John C. A. Bambenek, bambenek@gmail.com

OUTLINE

1. Introduction.....	2
2. The Rise of Breach Notification Laws.....	3
3. Breach Notification as Part of Incident Handling..	4
4. Finding Breach Notification Regulations.....	5
5. The Breach Notification "Trigger".....	7
6. When Breach Notification Is Not Required.....	8
7. Identifying Sensitive Data.....	9
8. How to Proceed with Breach Notification.....	17
9. Conclusion.....	18
10. Resources.....	19

INTRODUCTION

A computer system breach can be one of the most difficult situations for IT staff to handle. Without a response plan in place, the IT organization can become overwhelmed by internal accusations and unfounded new demands. The IT group that does have an incident response plan in place can react quickly and positively, mitigating rumors and negative reactionary behaviors. When the system breach involves the exposure of personal information such as Social Security Numbers and credit card data, the well-prepared organization will also be able to handle the public's response with a clear understanding of Breach Notification Laws.

Historically, organizations were very reluctant to reveal the fact that they had experienced a breach. Doing so would reveal some weakness that ultimately would result in lost customer confidence. But with the computerized storage of

Personally Identifiable Information, consumers realized that such breaches could have a personally detrimental impact.

This document will provide the IT professional with a general understanding of what "breach notification" is, and demonstrate some of the variety found among the legal requirements for actually conducting notification. In addition, this document will identify some of the tools currently available that may assist an incident handler with determining what data may have actually been exposed, and offer suggested means of conducting the actual notification.

THE RISE OF BREACH NOTIFICATION LAWS

The fact that companies and government agencies have been the victims of computer attacks is nothing new. However, in the past the revelations about these attacks came to light months or even years after the fact.

The birth of the Mosaic HTML browser in the early 1990s gave the public easy access to information across the Internet. It did not take long for corporations to see the commercial value in attracting customers to their own web sites, and eventually e-commerce was born. The exponential increase in personal data collected by these corporations prompted legislators to mandate that this data be safeguarded. However, hackers have likewise been increasing in numbers and skills.

Breaches into the systems holding people's personal information became almost a weekly occurrence in the early 2000s. With the data exposed, states began passing legislation

known as "Breach Notification Laws." These laws are designed to protect consumers by mandating prompt notification of a computer breach that exposed personal information, in an effort to help the consumer avoid financial losses or all-out identity theft. Such laws are also appearing on the international scene.

BREACH NOTIFICATION AS PART OF INCIDENT HANDLING

Incident handlers are typically engaged when a system breach has occurred as the result of an external hacker intrusion. The handler has to identify the type or method of intrusion, stop further access, and mitigate the damage done to the server by the installation of root kit software and other malware.

Another type of breach that might not immediately appear to be under the scope of an incident handler is the physical theft of a system. There have been several high-profile laptop thefts, one of the most damaging being the theft of a laptop from an employee of the US Department of Veterans Affairs. It was determined that the laptop contained the personal information of over 26 million military veterans.

Assisting with breach notification is quickly becoming another important step in incident handling. In general, breach notification is the process of communicating the fact that a computer system that contained personal information was compromised. The goal is to alert those whose personal information was stored on the system so that they might take special measures to protect themselves from identity theft, financial fraud or other personal injury.

FINDING BREACH NOTIFICATION REGULATIONS

To quickly determine the legal requirements for breach notification in a particular state, contacting the state's office of the Attorney General will often be an excellent starting place. Many AG offices now have special departments or task forces dealing with technology issues such as identity theft and Internet fraud. Breach notification would fall under their purview.

A state's web site is of course another place to find the specific laws of that state. If the state's AG office is not available as a direct link, a search of the site for the following terms may return results leading toward their breach notification laws: Consumer Protection, Identity Theft, Privacy, Personal Information, Business/Customer Data Security, and of course Breach Notification.

Many state sites offer searches against their specific statutes and codes. Therefore, searching for "breach notification" in such search engines will likely return many hits related to breach of contracts! The results will usually return the actual legal text of the laws, so finding the content that relates specifically to computer system compromises can be more challenging. To illustrate this concept, the California state web site (<http://california.gov>) offers a link to "Legislature" (<http://www.legislature.ca.gov>) and there is a link under "Research, Laws and Publications" where you can search the state statutes. California's legislative council also hosts a search engine against state laws (<http://www.leginfo.ca.gov/calaw.html> -- California's law is

"California Codes, Civil Code, Section 1798.80-1798.84" linked here: <http://www.leginfo.ca.gov/cgi-bin/waisgate?WAISdocID=03458915103+16+0+0&WAISaction=retrieve>).

The layman's version of California's law is more easily found in the Attorney's General web site (<http://ag.ca.gov>). There is a main section for Consumer Alerts & Information, and then a link to Identity Theft. This provides a clear, concise synopsis of the state's breach notification requirements for businesses, with a reference to the actual statute (<http://ag.ca.gov/idtheft/index.php>).

Since data breaches have been most common in the United States, other countries have been slower to adopt consumer notification rules. That is not to say that there are not privacy protection laws in place around the world. Canada's Office of the Privacy Commissioner offers an excellent list of "Key Steps for Organizations in Responding to Privacy Breaches" (http://www.privcom.gc.ca/information/guide/2007/gl_070801_02_e.asp). Based on compliance with Canada's Personal Information Protection and Electronic Documents Act (PIPEDA), these guidelines are designed to help Canadian businesses respond to a system breach in total. Although providing notification of a breach to customers is not specifically required under PIPEDA, it is addressed as an important consideration in the incident response plan. PIPEDA sets the expectation that personal information will be secured and protected appropriately, and does not address the case where such safeguards might fail and result in a system breach.

As recently as August of 2007, New Zealand and Ireland were two countries considering adding breach notification

requirements to their existing data privacy and protection laws. Driving factors for these countries included the global scope of the Payment Card Industry Data Security Standard (PCIDSS), which requires notification to Visa and MasterCard when a merchant's credit card data are compromised (PCIDSS does not have a requirement to notify consumers), and the global impact of electronic commerce on the Internet.

THE BREACH NOTIFICATION "TRIGGER"

Under most laws, it is the "reasonable belief" that sensitive data were acquired by an unauthorized person that triggers the requirement of notification. And once a system has been breached, the legal assumption is that everything on the system was accessible to the intruder. Whether incident management is handled by the in-house IT department or an outside specialist, part of the recovery process will need to include determining what data may have been exposed to an intruder. While log analysis may reveal whether or not files were accessed, skilled hackers can modify such logs to mask their activities. Therefore, the investigation should go deeper into determining whether any of the files or databases on the breached server contained sensitive data. The mere presence of files or databases containing sensitive data on a breached system is often enough to trigger breach notification requirements.

A number of states have a "risk-based" trigger. This allows for some discretion in deciding whether or not to notify based on the disposition of the data after the breach. In Arizona's law, for example, a "reasonable likelihood of substantial

economic loss" to an individual whose data are exposed would necessitate notification (<http://www.azleg.gov/ars/44/07501.htm>). Until such risk-based triggers are tested in court, companies will most likely set the threshold for the individual consumer losses significantly high, or conclude that the risk to the exposed data is sufficiently low, effectively allowing the companies to decide themselves whether or not to provide notification.

WHEN BREACH NOTIFICATION IS NOT REQUIRED

One important exemption in almost all breach notification laws is the case where the stolen information or storage device (laptop, disk, backup tape, etc.) was **encrypted**. Proof that a stolen storage device was full-disk encrypted can alleviate the need for notification, regardless of what data were on the system. The use of an at-rest full disk encryption product such as PGP or SecureDoc from WinMagic (<http://winmagic.com>) would make a stolen laptop's data inaccessible - indeed, the system usually will not start without first unlocking the disk encryption. Many full-disk encryption programs are so thorough that a thief cannot even reformat the drive. The encryption locks even the Master Boot Record (MBR), which prevents many operating systems from loading. Only a low-level reformatting solution can make the disk usable again.

Encryption on servers can be achieved in myriad ways, but may not protect data in every breach situation. Many full-disk encryption solutions for servers only encrypt the data "at rest," which typically means when the server is completely shut down. The physical theft of a server would usually mean the

system was in fact at rest at some point, and the disks should be encrypted just like the stolen laptop. However, the disks are usually NOT encrypted while the server is online. A breach of an online system is therefore NOT exempted from breach notification rules even if full-disk at rest encryption solutions were in place. Thus, only a file-level or full-time encryption solution on servers would properly secure customers' sensitive data.

Data that are copied onto a laptop or thumb drive typically ARE NOT encrypted on the destination media. It is critical that IT departments test and train their users when sensitive data are moved or copied onto portable media. File level encryption solutions like TrueCrypt (<http://www.truecrypt.org>) require the user to decrypt and mount a virtual volume on the target device, into which sensitive data could be saved. The data are then encrypted when the virtual volume is dismounted. Since this is not a full-disk solution, there is always the possibility that the sensitive files may not be saved to the proper location on the removable media, and therefore could be exposed.

Another exception specifically related to the "timeliness" of notification requirements is in the situation where law enforcement officials advise against publicizing the breach. If notification would hinder an ongoing investigation, the process can be postponed until the conclusion of the investigation.

IDENTIFYING SENSITIVE DATA

There are some variations on the definition of "sensitive" or personal data among statutes. In general, definitions will

include a person's first and last name together with one or more pieces of additional information such as Social Security Number, birth date, driver license number, or some financial account number (bank account or credit card number). Public or directory information such as name, address and telephone number by themselves are not considered sensitive. Internal account numbers or other company-specific identifiers are typically not considered sensitive either. Other laws protect the data handled by financial institutions (GLBA) and health information entities (HIPAA), and include specific remedies for breaches in these realms.

If the IT department has a current asset inventory, determining what data files are on what systems should be straightforward. In larger organizations where separate servers can be dedicated for specific single applications or asset groups, such inventories are typically well maintained. Of course, many organizations must leverage their limited server resources and host a variety of systems and files on the same physical server. Database server, file server, web server, and even email server could all be hosted on a single machine. So identifying sensitive data across these platforms can be much more difficult.

FILE SEARCHING TOOLS

There are now many tools available that can automatically scan files and databases and perform content analysis, matching patterns that could be credit card numbers, Social Security Numbers, etc. One such product is "Spider," from Cornell University (<http://www.cit.cornell.edu/security/tools/>). The free tool is available for various OS platforms, and comes with

pre-configured search options for finding credit card numbers, US Social Security numbers, and even UK and Canadian ID numbers. Another tool for Windows is "PowerGREP" by JGSoftware (<http://www.just-great-software.com>). PowerGREP is not freeware, but offers more functionality including the ability to scan across networks and perform search/replace.

Both Spider and PowerGREP have options to configure the process intensity of the scan, but even at maximum CPU priority the search can be very time consuming. For a RAID volume with 10 GB of documents and spreadsheets, Spider can take up to 4 hours to scan all the data.

USING REGULAR EXPRESSIONS

One key to improving the efficiency of the search is building the proper regular expressions. Crafting the "regex" helps the scan engine run more quickly and can significantly reduce the number of false positive matches. Of course, regular expressions can seem cryptic to an IT administrator who may not be involved with code development. The following examples will walk through the syntax of some basic regexes that could be used to search for US Social Security Numbers.

This regex "\b[0-9]{9}\b" translates as "between words (the \b means character string break), match the 'pattern,' which is the information in the square brackets (in this case is any digit zero through nine inclusive), when the pattern is repeated nine times (the bracketed nine {9} means repeat the contents of the pattern)." This search will find any 9-digit number, including nine zeros, eight zeros and the number one, eight nines, etc. This regex isn't very efficient for finding US Social Security Numbers, especially since it does not match

numbers that include spaces or dashes, either.

A far better regex for finding SSNs would be:

```
\b[0-7]\d{2}[- ]\d{2}[- ]\d{4}\b
```

Translation: find numbers that start with zero through 7 (SSNs do not start with 8 or 9), followed by any two digits (the `\d` is shorthand for digits zero through nine), followed by a dash or a space, followed by two more digits, another dash or space, and any four digits.

This is a better regex for finding SSNs, but there could still be some improvements. Consider these additional regexes:

```
^(?=((0[1-9]0)|([1-7][1-7]\d)|(00[1-9])|(0[1-9][1-9]))-(?=((1-9)0)|(0[1-9])|([1-9][1-9]))-(?=((\d{3}[1-9])$|([1-9]\d{3})$|(\d[1-9]\d{2})$|(\d{2}[1-9]\d)$))))
```

"The first three digits cannot be greater than 779, nor can they be 000. The second two digits cannot be 00. The last four digits cannot be 0000. This regex also checks formatting for numbering and dashes (###-##-####)."

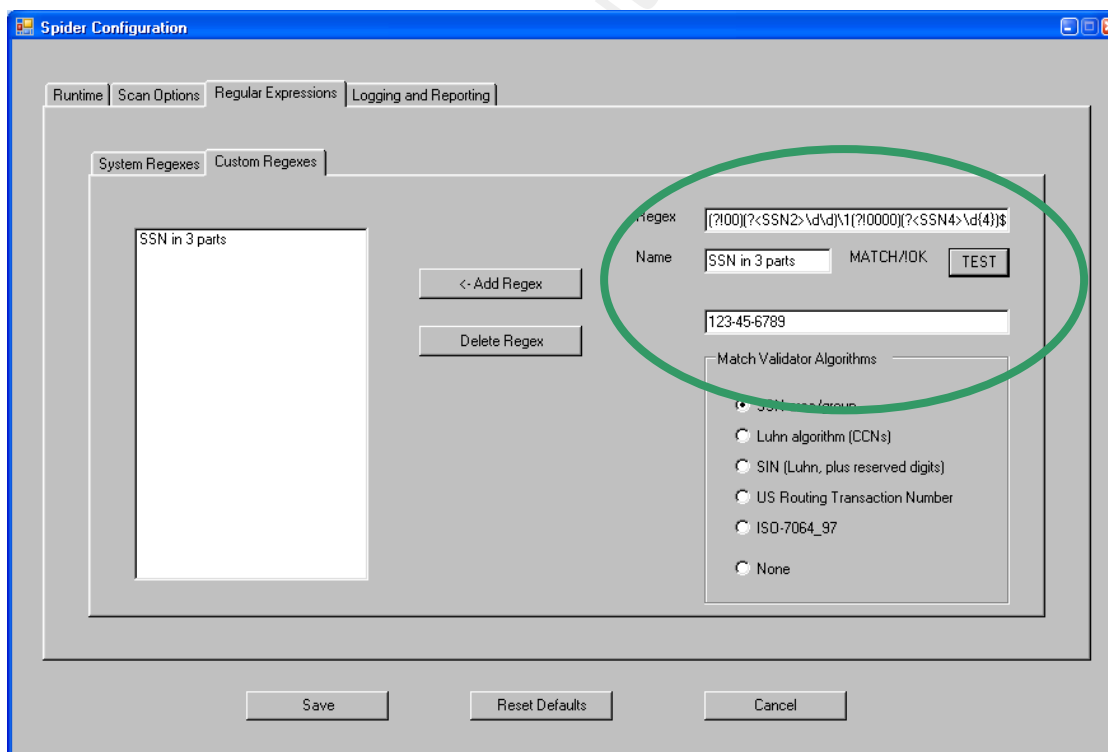
(http://regexlib.com/REDetails.aspx?regex_id=418)

```
^(?!000)(?!666)(?<SSN3>[0-6]\d{2}|7(?:[0-6]\d|7[012]))([- ]?)(?!00)(?<SSN2>\d\d)\1(?!0000)(?<SSN4>\d{4})$
```

"This regex excludes SSNs that begin with 000 or 666, then matches 0-6 plus any two digits, or (the "|") 7 and any two

digits up to 72. It accepts optional hyphens or spaces as formatting characters. It parses the three subfields of the SSN into three named sub-strings (SSN1, SSN2, and SSN3) to facilitate program use. It rejects matches on all zeros for any individual subfield of the Social Security Number." (http://regexlib.com/REDetails.aspx?regex_id=539)

These regexes and others (to find Visa/MasterCard and other credit card numbers, for example), can be found online at the "Regular Expression Library" (<http://RegExLib.com>) (The above examples are credited to their respective authors in the links provided.) RegExLib.com also offers "cheat sheet" references and other resources for creating or finding customized regexes.



Spider includes a regex test feature right in the application. In this screen shot, the last regex example above was added, then a random SSN was entered. Clicking the "Test" button reveals that the random SSN selected would be a match using the regex. Be sure to test that any formatting checks

return the desired matches, and that the exceptions will in fact not match, either. When testing regexes, realize that the "standard" format for numbers like SSNs can change over time. Notice in the examples above, one author uses "779" as the upper limit for the first part of the SSN, while the second uses "772." For additional information on US Social Security Number formats and currently valid ranges, see the Social Security web site (<http://www.socialsecurity.gov/employer/stateweb.htm>).

The top screenshot shows a regex testing interface. The 'Regex' field contains the pattern `(?000)(?<SSN2>\d\d)\1(?0000)(?<SSN4>\d(4))$`. The 'Name' field is 'SSN in 3 parts'. The 'TEST' button is highlighted with a green circle. The input field contains '123456789'. Below the input field, the 'Match Validator Algorithms' section is visible, with 'SSN area/group' selected.

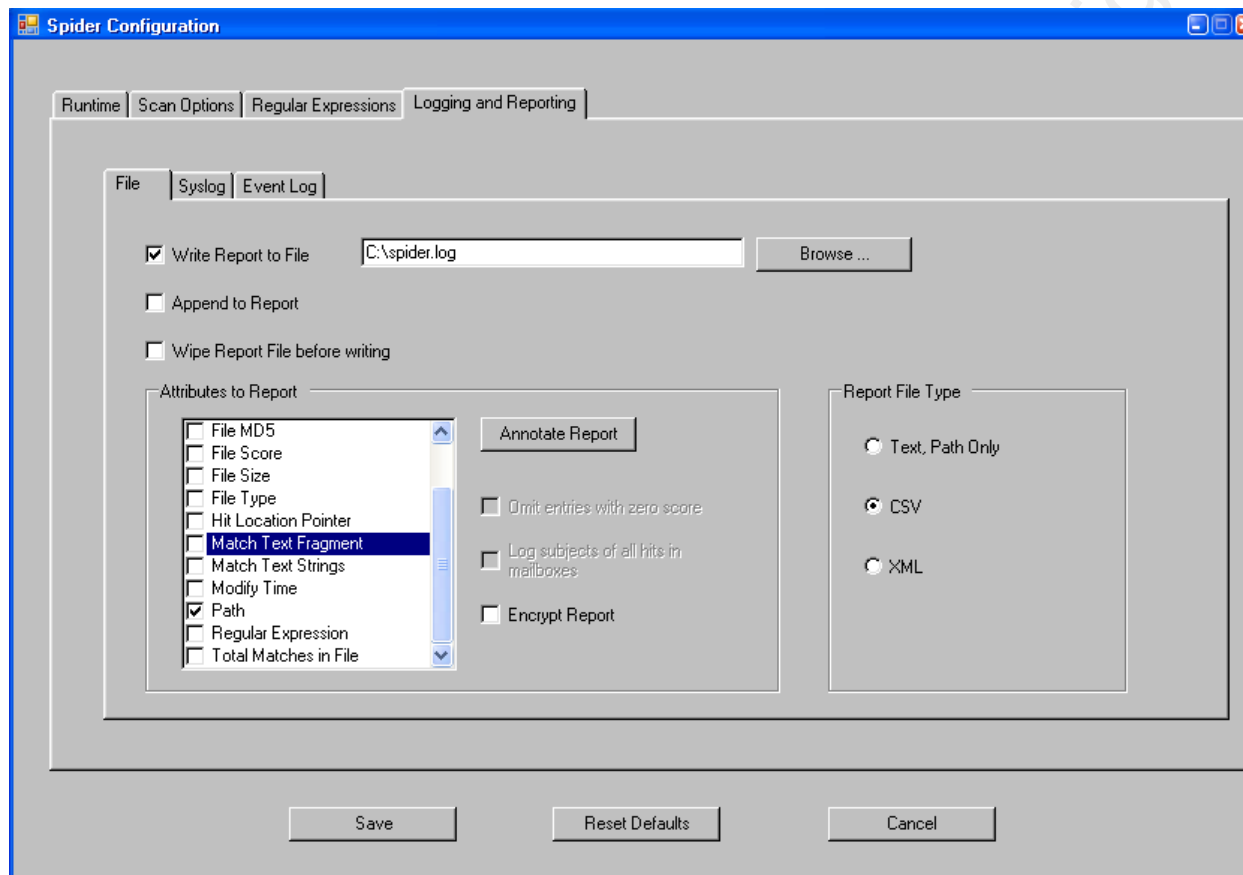
The bottom screenshot shows the same interface. The 'TEST' button is highlighted with a red circle. The input field contains '666-45-6789'. The 'Match Validator Algorithms' section is visible, with 'SSN area/group' selected.

The latest version of Spider includes "Match Validator Algorithms" which can further reduce matching false positives. One of these is the Luhn algorithm, which validates the last digit of a credit card number against the whole number. Others can validate bank account and check routing numbers. The Google Wiki "Check Digit Systems" covers these and other validation algorithms (<http://code.google.com/p/checkdigits/wiki/CheckDigitSystems>).

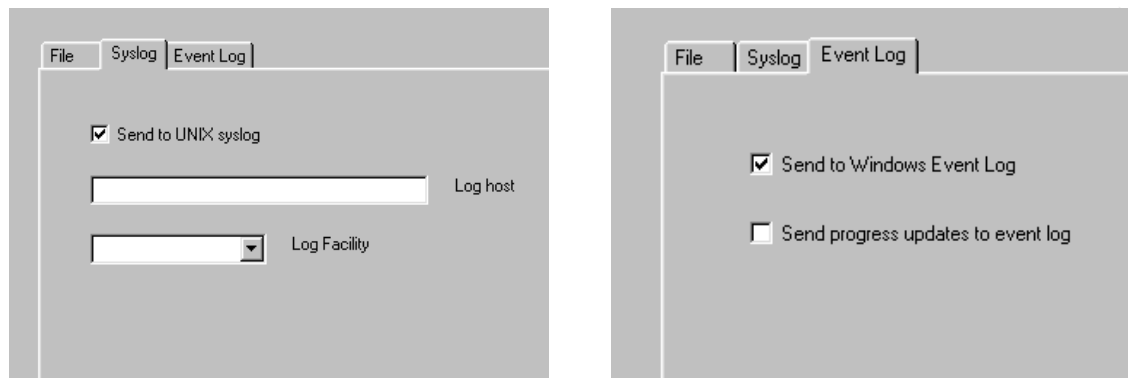
CONFIGURING OTHER SCANNING TOOL OPTIONS

Care should be taken in configuring the output from any file scanning tool. Log files may duplicate the sensitive data, and if the product can be set to save to a network share or email scan results, doing so can potentially expose sensitive data yet

again across the network.



This screen shot above from Spider shows the options for configuring the scan report file that will be generated after a search. Note in the "Attributes to Report" window that the actual Match Text can be selected. One safety feature in Spider is that the "Encrypt Report" option will be automatically selected if the Match Text options are checked. Additionally, encryption can only be performed by Spider for Text and CSV output, not XML, so that option also becomes unavailable.



The options to send the scan reports to a central UNIX syslog server or to the Windows Event Log should be used with caution. Assuming the syslog server is secure and communicating over TCP with SSL/TLS encryption, the sensitive data may still be replicated in these logs. With syslog or the Event Log, any administrator would have access to data, and it may be against corporate security policies to purge or delete such logs.

While PowerGREP is another powerful file searching tool, it was not built exclusively for finding sensitive data as was Spider. There are no corresponding options to encrypt or secure data results in PowerGREP, nor to save results to a log server.

Most file scanning products also typically modify the access time stamps on files. If using these tools after a breach, they should only be run on a backup copy of the data, or on the system itself after any required forensic analysis of the system has been completed.

Determining what was on a laptop at the time of its loss can indeed become the task of the incident handler and IT department. While the laptop owner may know and reveal that there were sensitive files on the system, retrieving copies of those files may or may not be feasible.

HOW TO PROCEED WITH BREACH NOTIFICATION

Once sensitive data are discovered on a compromised system, breach notification laws generally require that the people connected to the data be advised of the breach "immediately," which is to say as soon as possible or without undo delay. To be sure, this work is not expected to fall to the incident handler! But the handler should be knowledgeable of the requirements and acceptable methods for providing notification, in order to properly advise management.

The process of matching sensitive data to individuals, and then finding contact information for these people, can be much more difficult than it seems. In the "easy" situation, consider a university grad student who copied a spreadsheet onto her laptop from the departmental server that contained the name, address, SSN and birth date of every student who had attended the institution for the past 8 years, then the laptop was stolen from her apartment. There is little more to do than create a mail merge using the spreadsheet on the server and print letters! Of course, the breach of a database server might require more work to match up transactional data that contains names and credit card numbers for the past six months with the main customer information list in a different database (or spreadsheet!) that goes back to 2001.

Breach notification laws generally provide for several different means of delivering the actual notification to individuals whose personal information may have been exposed. The methods are typically listed in order of "most preferred" to "last resort," and will have some thresholds in terms of cost, workload and/or timeliness.

The most preferred method is written notification, but (referring to the thresholds in the California law) when the cost of producing written notification would exceed \$250,000, or the exposures involve over 500,000 individuals, notification can be made via email. Posting notifications on the corporate web site or providing press releases to public media such as newspapers and television are typically used as additional methods of notification, but these are the least preferred means by themselves. However, for exposures where individual notification may not be possible - because contact information for the individuals is not available - public notification through the web and other media may be the only options.

CONCLUSION

The IT professional with a general understanding of the breach notification laws will be better suited to respond to the exposure of sensitive data. Understanding that there are significant variations among the legal requirements for actually conducting notification, building specific breach notification procedures into a corporate incident response plan does take research and localized customization. Maintaining an accurate systems and data inventory and becoming familiar with some of the tools that help find sensitive data can significantly reduce the time needed to determine if breach notification is even a requirement in a given situation. When it is required, knowing the required and suggested means of conducting the actual notification allows management to respond as quickly as possible.

RESOURCES

Alexander, P (2007, 04, 09). Data breach notification laws: a state-by-state perspective. IntellegentEnterprise, Retrieved December, 2007, from <http://www.intelligententerprise.com/showArticle.jhtml?articleID=198800638>

(2002). Overview of American data breach notification laws. Retrieved September, 2007, from Office of the Privacy Commissioner of Canada Web site: http://www.privcom.gc.ca/parl/2007/sub_070222_06_e.asp

(2006, 03, 02). Metro state: stolen laptop had 93,000 social security numbers. Retrieved November, 2007, from Denver, Colorado Channel 7 News Web site: <http://www.thedenverchannel.com/news/7621150/detail.html>

(2006, 05, 22). US says personal data on millions of veterans stolen. Retrieved November, 2007, from WashingtonPost.com Web site: <http://www.washingtonpost.com/wp-dyn/content/article/2006/05/22/AR2006052200690.html>

(2007). State security breach notification laws. Retrieved September, 2007, from National Conference of State Legislatures Web site: <http://www.ncsl.org/programs/lis/cip/priv/breachlaws.htm>

(2007, 08, 26). New Zealand gets draft breach notification guidelines. Retrieved December, 2007, from LawFuel.com Web site: <http://lawfuel.com/show-release.asp?ID=14532>

(2007, 08, 31). Securing information about data theft. Irish Times, Ltd. / TMCNet.com, Retrieved November, 2007, from

Breach Notification in Incident Handling

<http://www.tmcnet.com/usubmit/2007/08/31/2901198.htm>

(2008, 01, 08). Open-source forensics tools for network and system administrators. Retrieved January, 2008, from Cornell University IT Security Office Web site:

<http://www.cit.cornell.edu/security/tools/>

Arizona State Legislature:

<http://www.azleg.gov/Search.asp>

Office of the Attorney General of the State of California:

<http://www.ag.ca.gov/>

Office of the Attorney General of the State of Colorado:

<http://www.ago.state.co.us/index.cfm>