



STRENGTHENING EVIDENCE THROUGH SMARTER INFRASTRUCTURE DESIGN



Jason Ventresco
Principal Solutions Engineer
EMC
Jason.Ventresco@emc.com

EMC²

Table of Contents

Executive Summary	3
Introduction.....	5
Separation of Duties	7
The Websense Example	9
Security as Part of the Model, not as an Afterthought	12
Extending the SDLC Model	13
Trust, but Verify	16
Creating the (Nearly) Infallible Log	18
Foundations in Forensics	21
Conclusion.....	26
Appendix.....	28

Disclaimer: The views, processes, or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation's views, processes, or methodologies.

Executive Summary

A modern information technology infrastructure generates a large amount of data that could potentially be used to investigate a crime. A brief list of sources of this data includes:

- Firewall log data including but not limited to alerts generated based on predefined event thresholds and perimeter network traffic statistics.
- Router logs that contain detailed data about various types of intra and extra-organizational network traffic.
- Antimalware activity logs that contain details of attacks, what time that they occurred, and their overall severity.
- Intrusion detection or prevention logs that contain data about detected events, what time they occurred, and even packet data from the event.
- Server time-stamped access logs that show logon successes, logon failures, and other access attempts.
- Network directory services logs that contain historical data about changes, additions, and removals from the directory.
- Platform monitoring tools that monitor server, workstation, and network device activities for events of interest or concern.
- Web server logs that contain historical information about server utilization and details about what components of the website were accessed and when.
- Remote access or application logs that record usage information.
- Desktop, laptop, personal digital assistant, and smartphone logs all contain a variety of data that may or may not be of use during an investigation.

Data obtained from one of these sources may not be compelling by itself, but when placed into context with other available information an investigator may be able to start drawing conclusions as to the details of the crime. The time to question the usefulness of a particular source as evidence is after an incident has occurred, not before, which means that companies should not limit monitoring activities or archival of monitoring information based on a preconceived notion about the “value” of the information.

Once you have accepted the fact that the average information technology infrastructure produces a large amount of potentially useful data, you must then consider how you protect the integrity of that information and ensure that it can stand up to whatever scrutiny that may arise. There are numerous techniques for achieving this of course, a good deal of which revolve around technical or procedural

methods. This Knowledge Sharing article discusses a number of these technical and procedural techniques within the context of preserving the integrity of information in general, with a focus on that which is likely to be used during an investigation.

Introduction

The field of information security is anything but stagnant. The driving force behind many of the innovations within the information security arena appears to be the never ending battle between criminals and those organizations or individuals they target. Criminals recognize that technology is becoming more ingrained in our day-to-day lives, and in response they are trading their lock picks and guns for a laptop and an Internet connection. Gone are the stereotypical criminal thugs made famous in movies; they have been replaced with individuals who wield technological skills that are sufficient to earn them a job at any number of premier technology companies.

While information security is a complex topic, the underlying concepts behind it are not much different than those which define security itself. The minute one decides to engage the outside world in the pursuit of profit or even altruism you expose yourself to a certain amount of risk, and as members of a civilized society we must find civilized methods to mitigate that risk or at least make it manageable.

Businesses have a few options when it comes to guarding against criminals.

- Technological defenses – Firewalls, intrusion prevention devices, network access control platforms, Internet security and monitoring platforms, antimalware software, are just some of the tools that companies can use to protect and control their electronic assets.
- Employee policies – Be they actively enforced or simply an agreement between employer and employee, these policies restrict how technology is supposed to be used within an organization.

These are a small portion of tools and techniques that companies depend on to keep their businesses safe.

Unfortunately, no security solution is perfect. Though companies such as Perimeter eSecurity spend millions of dollars developing custom security solutions, in the end there is still a human being behind their technological curtain making sure that everything is working as it should. Vendors and solution providers hype security products with impressive client lists and high price tags; after all, if it is good enough for a Fortune 100 company it must be good enough for my own infrastructure.

Companies need to start living in the real world; the one described in the 2010/2011 CSI survey where 21.6 percent of respondents were subjected to a targeted attack, and the average loss approached six figures. In all their effort to stay on the defensive, companies rarely stop to think about going on the

offensive. In this case, offensive refers to the only action that companies can take outside of implementing more security; initiating formal investigations and prosecuting the accused in court.

Prosecuting computer crimes requires the same things that prosecuting any crime does; evidence. This article will examine what companies can do to strengthen their offensive capabilities by using the information contained within their infrastructure to identify and convict those that have attacked it. The questions this article will answer include:

- How can I configure my infrastructure so that it can provide the information that I need to prosecute a security breach?
- What can I do to make sure that the information provided by my infrastructure is reliable? This includes all data involved with or that may lead to an investigation.
- What can I do to prevent a criminal or even simple internal error from compromising the evidence that I need to perform an investigation?

Answers to these questions and others are important if an organization is to obtain reliable, dependable information from their infrastructure.

Separation of Duties

The Information Systems Audit and Control Association defines Separation of Duties as “a basic internal control that prevents or detects errors and irregularities by assigning to separate individuals responsibility for initiating and recording transactions and custody of assets to separate individuals”. Information technology departments implement separation of duties for many reasons and to prevent a number of different scenarios. One way to illustrate the concept of separation of duties is with a table of common tasks involved with building and maintaining an information technology infrastructure.

	Design changes to the infrastructure	Approve changes to the infrastructure	Make changes to the firewall	Make changes to the servers	Make changes to the telecom systems	Make changes to the user security policies
Design or suggest changes to the infrastructure		X	X	X	X	X
Approve changes to the infrastructure	X		X	X	X	X
Make changes to the firewall	OK	X		X	X	X
Make changes to the servers	OK	X	X		X	X
Make changes to the telecommunications systems	OK	X	X	X		X
Make changes to the security policies	OK	X	X	X	X	

Table 1: Common tasks performed by IT departments

Table 1 includes six different tasks common to many information technology departments. These tasks include changes to the firewall, servers, telecommunications systems, security policies, and infrastructure design, as well as the approval of those changes.

In this example, “X” designates situations where the responsibilities of one individual or group should not intersect with a related responsibility. The theory behind the concept of separation of duties is that no one individual (or group) has the authority to fully compromise the infrastructure. In the example provided we have essentially three unique functions:

- Those responsible for designing the infrastructure.
- Those responsible for approving changes to the infrastructure.
- Those responsible for implementing changes to the four key components of the infrastructure.

Ideally, these three functions will be handled by wholly distinct groups within an organization. While nothing can prevent individuals from different groups from colluding with one another to subvert the organizations’ security measures, allowing one individual or group to control multiple functions of the table will threaten the very balance that idea of separation of duties hopes to provide. In some cases a combination of duties would be acceptable, which in the example provided is marked with an “OK”. It is assumed that those that implement the infrastructure would have some hand in how it is designed, although in the interest of separation of duties they should still be excluded from the approval process.

The concept of separation of duties can be applied to the process of gathering, analyzing, and retaining data that could be used as evidence. The tools that organizations buy to secure, monitor, and maintain their information technology infrastructures produce large amounts of data, yet the responsibility of securing the information generated by these tools lies with those that configure them.

The Websense Example

Let us examine a fairly basic installation of the Websense Web Filter and Monitoring Platform, an enterprise-capable solution for monitoring Internet activity, enforcing browsing restrictions, and protecting the organization from web-based malware threats. Websense requires a number of different components in order to function, which means by extension there are a number of areas where the reliability of the platform can be compromised.

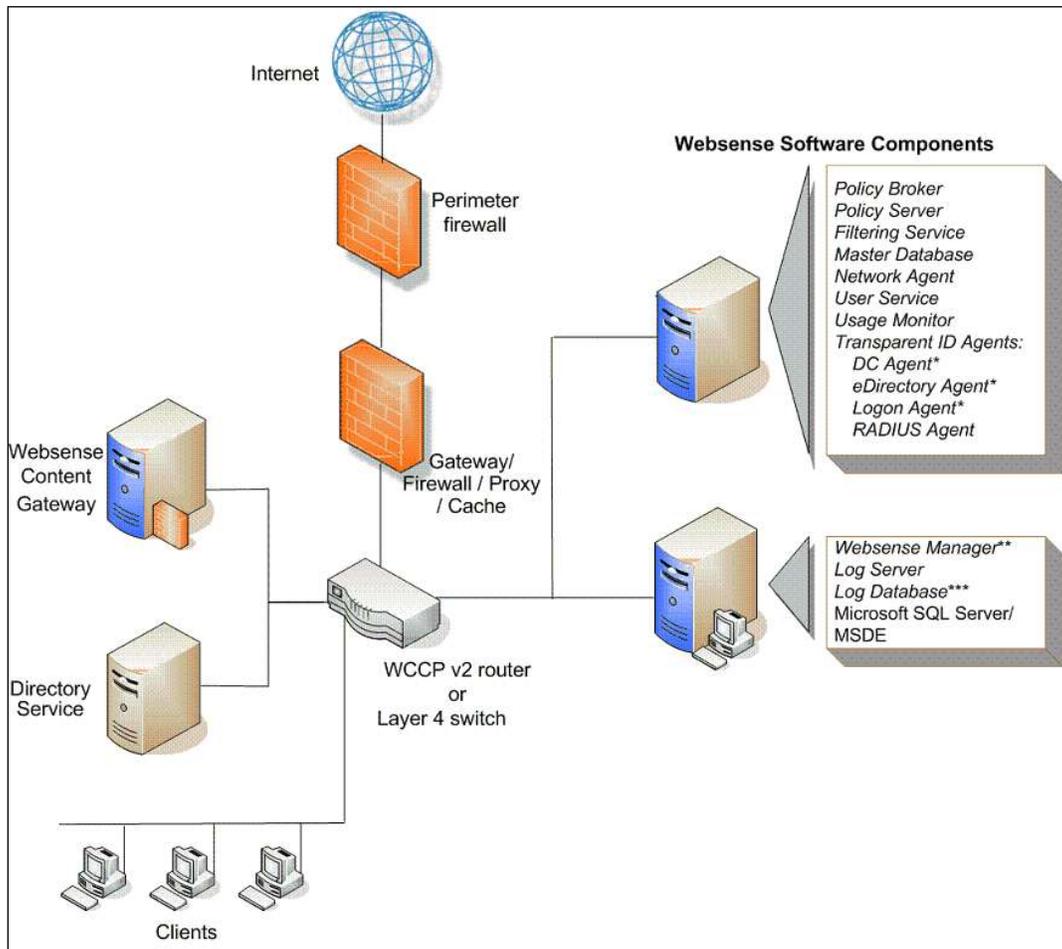


Figure 1: Websense Architecture

To function properly, the Websense platform requires the collective efforts of six functional groups similar to those described in the initial separation of duties table:

- The server team is responsible for the deployment of the Websense server-based components as needed within the infrastructure.

- The networking team is responsible for configuring the network equipment to enable the actual monitoring.
- The security policy team is responsible for drafting the policies that impact how the Websense platform will operate:
 - Workstation security policies that ensure that a workstation cannot be compromised by someone who wishes to use it to hide their own Internet usage.
 - Websense blocking policies that restrict Internet access.
 - Websense notification policies that inform the appropriate personnel when specific Internet usage patterns are observed.
- The database administration team is responsible for a number of tasks involving the Websense database:
 - Archival of Websense monitoring data as needed to ensure database performance and manageability.
 - General database maintenance as needed to maintain the reliability of the product.
- The backup and recovery team is responsible for performing ongoing backups of the current monitoring data and the Websense server and network device configurations.
- The monitoring team is responsible for responding to Websense alerts and ensuring that the platform is operating correctly.

The distribution of workload between the six functional groups is critical to maintaining the concept of separation of duties. Assuming that each group has access only to the components for which they are responsible, the most that any one group could do is cause an interruption in the general availability of the platform. There are a number of examples of how the separation of duties should allow the organization to continue in the event that the Websense platform has been compromised:

- The server, networking, or security teams make a change in either of their respective systems in order to subvert the Websense platform.
 - The monitoring team will notice that the platform is not functioning as expected.

- The backup and recovery team can utilize their backups to restore the system to a point in time where it was known to be working.
- The database administration team alters the contents of a database in order to erase their activities.
 - The monitoring team will notice that the database has been accessed by something other than the Websense application itself.
 - The backup and recovery team can utilize their backups to restore the erased data.

These are just some scenarios that outline how the separation of duties can help organizations preserve the integrity of their infrastructure and the availability of evidence needed to investigate any incidents that may occur.

Implementing the concept of separation of duties requires more than I have outlined. For instance, in the example I assumed that the archives made by the backup and recovery team are monitored in accordance with internal policies or even local, state, or federal statutes. One example of this is Sarbanes-Oxley legislation; it requires companies to institute various controls over their data backup and retention processes in the interest of preventing financial fraud. While these regulations were not created for the purpose of improving infrastructure security, adherence should assist companies in maintaining the availability and integrity of their data, most notably that which may be used as evidence.

Security as Part of the Model, not as an Afterthought

Measuring the overall security stance of an organizations' information technology infrastructure is not a trivial task. In my many years of working as an IT consultant I have seen companies use a number of different approaches to ensure that their infrastructures are secure, including:

- Use of established infrastructure deployment standards that are published by vendors or other third parties.
 - Institutions such as the National Security Agency publish guides that assist organizations in deploying secure computing systems.
- Retaining the services of qualified personnel when deploying new platforms, performing infrastructure upgrades, or when making changes that affect large portions of the infrastructure.
 - Information technology companies such as Cisco, Microsoft, and Red Hat work with vendors and solution providers so that they can develop and maintain expertise in their products.
 - Vendors and solution providers engage in partnerships with these companies in order to show a level of commitment to their products.
 - Information technology workers utilize training and certification programs to further their expertise in these products and establish themselves as an expert.
- Perform regular security audits including but not limited to penetration testing, vulnerability assessments, and random checks of infrastructure records.

These represent a small sample of the ways that organizations can ensure that their information technology infrastructures are secured.

Most information technology professionals would tend to agree that the examples I have provided represent a reasonable starting point for improving the security posture of an organization. Each of the three approaches mentioned involves seeking outside assistance, be it in the form of subject matter experts or simply as a third-party "unbiased" examination of the state of the infrastructure. Unfortunately, there is no panacea among the methods that I list, and organizations that fail to realize this are putting themselves at risk by not seeking an unbiased view of their infrastructure.

The impact of shortsightedness as it relates to the security of a modern information technology infrastructure brings to mind the cliché “a chain is only as strong as its weakest link”. There are a number of scenarios that can occur within an organization that lead to the creation of weak links; managerial shortcomings, office politics, lack of expertise or understanding of the challenges facing the infrastructure, or just simple human error. One way to deal with this problem is to standardize the approach by which an organization designs, implements, and maintains key infrastructure components. The Systems Development Life-Cycle is an example of a development methodology that, when properly applied, can ensure an organization is keeping security in the forefront throughout all phases of the infrastructure lifecycle. The Information Systems Audit and Control Association defines the Systems Development Life Cycle as “the process, involving multiple stages (from establishing the feasibility to carrying out post implementation reviews), used to convert a management need into an application system, which is custom-developed or purchased or is a combination of both”.

Table 2 shows how the Systems Development Life-Cycle could be extended to ensure that an organization is mindful of information security throughout all phases of the infrastructure design and management process.

Extending the SDLC Model

SDLC Phases	Extending that Phase to Involve Security
<p>Project Definition – The project definition phase is used by organizations to determine the feasibility of a given project, what the goals are, and an investigation of alternative solutions.</p>	<p>Define how this project will impact the security posture of the organization including but not limited to risk analysis, monitoring requirements, regulatory requirements, and disaster recovery considerations.</p>
	<p>Investigate what if any impact that the security needs of the project have on the net benefits that the project is to provide.</p>
<p>User Requirements Definition – The user requirements definition phase is used to define how users will interact with the proposed system, and what capabilities and features the users would expect.</p>	<p>Investigate the impact that the end user requirements will have on the security posture of the organization. Example: Will the projects require any alterations to the current policies and restrictions in place between the client/server, client/workstation, or workstation/server relationship?</p>

<p>System/Data Requirements Definition – The system/data requirements definition phase is where user requirements are analyzed and the high level architecture of the system is created.</p>	<p>Investigate the impact that the architecture will have across the rest of the enterprise. Example: Does the project introduce a weak link into a chain that is currently regarded as strong?</p>
	<p>Investigate the impact that the end user requirements are having on the overall system architecture, specifically if the user requirements are making the project inherently less secure or more difficult to secure.</p>
<p>Analysis and Design – The analysis and design phase is where the final architecture of the system is decided upon; it is impacted by the needs of the users, the cost of one solution versus another, purchasing restrictions, or by any restrictions that exist within the current infrastructure.</p>	<p>Evaluate if the external factors that impact the analysis and design of the project are having a negative impact on security. Example: Are cost considerations or purchasing restrictions taking precedence over functionality or security?</p>
	<p>Verify that the architecture of the project supports the concept of separation of duties.</p>
	<p>Evaluate the final design of the project to determine if any previously identified security concerns have been addressed, and if any new concerns have been identified.</p>
<p>System Build/Prototype/Pilot – The system build/prototype/pilot phase is where the project is built and deployed in a test status. This phase gives users an opportunity to validate the overall design of the solution, suggest areas for improvement, develop the details of any training that will be required, and generally verify that the project is meeting the expectations set forth in the project definition phase.</p>	<p>Evaluate the security of the pilot system and identify any deviations in expected performance.</p>
	<p>As the long term maintenance needs of the project become more apparent during the pilot phase, determine if they are any different than what was outlined during the project definition phase. Example: In the event that the project requires more maintenance or has greater disaster recovery or regulatory needs, how does that impact the ongoing maintenance requirements of the solution?</p>

	Evaluate any impact that user feature requests may have on the overall security of the project.
Implementation and Training – The implementation and training phase includes all actions that are required to move a prototype or pilot system into production, including training the end user community and those that will support the project.	Evaluate what if any security protocols will need to be relaxed to place the finished product into production.
	Evaluate the ongoing maintenance of the finished product and ensure that separation of duties is being maintained.

Table 2: Systems Development Life-Cycle model

The intent of the Systems Development Life-Cycle may have been designed as a conceptual model for project management, but as evidenced by my simple extensions to the model it can also assist an organization in becoming more secure.

Trust, but Verify

Former US president Ronald Reagan used the phrase “trust, but verify” throughout his presidential career. Those involved in an investigative process of any sort also live by those words, particularly if they want their evidence to stand up to scrutiny, be it hostile or cursory. The naïve among us believe that verification is the rule and not the exception.

- We assume that our public defenders do not indict individuals unless they know for certain that they are guilty.
- We assume that our expensive security monitoring products are incapable of making mistakes, and that we are protected against any type of attack.
- We assume that because something worked well yesterday, the day before, and the day before that will probably work well tomorrow.

Julie Amero makes no such assumptions and, after their evidence against her was so thoroughly discredited, nor does the government that prosecuted her. Julie Amero was a substitute teacher that was accused of viewing pornography on a school computer. The assumptions made during this trial were significant and numerous:

- The school district assumed that their computers were secure.
- The law enforcement agency assumed that the individual examining the evidence was qualified to do so, and would use methods that would stand up in court.
- The individual who examined the evidence thought that he knew what he was doing, and that the evidence he found indicated that Julie Amero was guilty.
- The judge that first worked on the case assumed that the individual who examined the evidence was competent by not challenging his credentials.

While Julie Amero was eventually found innocent, her case is a textbook example of what can happen when fundamentally weak evidence is used in an attempt to prosecute a case.

Let us pretend that we are the prosecuting attorney in the Amero case, but for this example let us assume that she is indeed guilty. As an attorney we are not an expert in the field of information security; the extent of our knowledge about computer security begins and ends with the computer we have at home and in our office, and last time we checked there were no middle school age children using our

computers to “explore” the Internet. As far as we are concerned we will never get spyware because we never use our computer for any questionable activities, which as we know is the only way that you can get spyware.

One day a local school calls and informs us that a teacher is viewing pornography in the classroom. Our first thought is that this sounds implausible, but then again in our twenty years deciding law we have seen and heard of much worse. We contact our “expert” on computer related matters and send him to collect the evidence. He collects all available evidence and starts his investigation.

The investigator examines all available evidence and feels that the claim is valid, and that it appears she did browse the sites in question. Based on the investigation and analysis of the investigator, the suspect is indicted and a court date set.

Now fast forward to the court case. Mrs. Amero has retained the services of an attorney who knows a little bit about information technology. This attorney obtained detailed records from the school district information technology department that outline password policies, screensaver lockout policies, log retention information, information about log storage facilities, antivirus client statistics, and information about past help desk calls concerning spyware infections.

The expert witness retained by the state has just finished giving his testimony. He feels that the evidence he has is strong and that the cross examination will be short. He is wrong of course; he just does not know it yet.

The defense attorney starts talking about the evidence that the investigator has presented. The questions asked include:

- How was the Internet browsing history verified? Is there any way to validate that the timestamps on the computer were indeed accurate?
- With a screensaver timeout being only ten minutes how do we know for certain that Mrs. Amero was the one using the computer?
- How are the Internet activity records gathered, stored, reviewed, and archived?
 - Who has access to the records at each of those stages?
 - How is the integrity of the records maintained, and can that integrity be proven?

- The antivirus client statistics show that on average at least one-third of the computers lack current antivirus definitions, including the computer Mrs. Amero was using.
 - Given that some types of malware are spread over a local area network, can the school district produce evidence that all other computers in the building were free of malware on the day that the incident took place?
- The evidence that you have presented includes data that is time stamped after the incident has occurred; what is the reason for this?
 - Does this indicate that none of the timestamps are reliable, or only that the evidence has been altered between the time the alleged incident has occurred and the trial.

One could assume that at this point the jury is starting to understand how reasonable doubt can be applied in cases that involve electronic evidence. The lack of controls and consistency throughout the school district network will likely jeopardize the only real evidence in this case, which in my example is the hard data provided by the tools that the school district uses to monitor the status of their infrastructure.

Creating the (Nearly) Infallible Log

To solve the problems that arose during my version of the Amero trial we need to take a closer look at how we generate, gather, and archive monitoring data that could be used as potential evidence. In essence, we must address the following problems:

- The method used to identify and gather the monitoring data.
- The method used to transfer the data to a remote destination, if applicable.
- The processes applied to the data once it has reached its destination.
 - These processes include any actions that alter or analyze the raw data prior to storage or presentation.
- The process by which the raw or processed data is stored, archived, and backed up.
- What if any methods exist that can associate the monitoring data with a specific individual?

- The separation of duties as it applies to the log file lifecycle; with particular attention to how current processes and security policies may compromise the data gathering, processing, and archival process.

The Amero case is an example where the monitoring data was insufficient to gain a conviction. A group of information technology security professionals headed up by Sunbelt Software CEO Alex Eckelberry was given the chance to examine the evidence in the Amero trial and their analysis raised a number of concerns about the validity of the evidence used in the trial.

- The firewall used by the school is not capable of identifying the specific user that visited the websites in question.
- The antimalware platform used by the school did not include a centralized monitoring system, only local logs.
 - The lack of a centralized logging system means that local computer logs are the only records kept, and increases the likelihood that anti-forensic techniques will be able to alter the timeline and make investigations difficult.

The Amero case is a lesson in what happens not only when you lack the proper monitoring platforms, but when the limited evidence that you do have is judged useless because of the poor condition of the platform or tool that gathered it.

A general framework for strengthening data monitoring activities is fairly straightforward if you approach the process from beginning to end.

1. Gathering and monitoring data.
 - a. Solutions such as the LANDesk Server Manager platform can monitor the majority of components within an infrastructure for data collection purposes as well as notify the appropriate personnel when events of interest occur.
 - i. The data collection process is tightly controlled through policies set by assigned LANDesk administrators.
 - b. Server and workstation event log data can be protected against unauthorized access.
2. Transfer of monitoring data for processing or storage.

- a. Platforms from vendors such as Cisco and Microsoft support the use of point-to-point encryption using the IPSec standard.
 - i. The use of encryption between data collection platforms and the processing and storage facilities reduces the likelihood that the monitoring data can be compromised or intercepted while in transit.
- 3. Processes applied to data once it has reached its destination.
 - a. Monitoring platforms should offer the ability to transmit data in raw format so that additional analysis can be performed as needed.
 - i. The syslog protocol enables a standardized approach for sending event messages across a network, and is supported by a wide variety of platforms.
- 4. Short term and long term storage of log data.
 - a. Once received, monitoring data should be secured against unauthorized access or tampering.
 - i. Backup products such as Symantec NetBackup support encryption of data at the host before it is transmitted over the network to the backup medium.
 - ii. Backups of monitoring data can be performed using unique credentials whose usage is tightly controlled and monitored.
- 5. Ensuring that monitoring data can be attributed to a specific individual.
 - a. Many monitoring platforms such as SurfControl Web Monitoring, HP Tipping Point, and Microsoft Internet Security and Acceleration Server have the ability to integrate with network directory platforms such as Microsoft Active Directory for authentication and record keeping purposes.
 - b. Tightening access control policies and procedures may be necessary to protect an employee's workstation from misuse.
 - i. Screensaver timeouts, restrictions on the time of day network accounts can be used, limiting privileges to the workstation operating system, and complex password policies are just some of the ways that organizations can secure workstations against misuse.

- c. Multi-factor authentication models can assist organizations in validating infrastructure access records.
 - i. Hardware-based access control devices such as smart cards can work in tandem with other security procedures to add an additional authentication layer.
 - ii. Companies such as CEM Systems have released solutions that enable organizations to implement up to a three-factor authentications process.
- 6. Separation of duties across the monitoring data lifecycle.
 - a. Organizations may wish to extend the concept of separation of duties to include the systems that perform infrastructure monitoring. Example scenarios include:
 - i. The group who review changes or updates to the monitoring platforms should be separate from those who make the changes or implement the updates.
 - ii. The group who reviews the monitoring data should be separate from those that archive the monitoring data.

Ensuring the reliability of monitoring data is difficult; perhaps the only way to approach it is to assume that the data will be used as evidence someday and treat it as such. This means securing it not only after it has been collected, but starting with the system that is used to collect it and working in reverse.

Foundations in Forensics

The Amero case presented in this article featured a computer forensic investigator (Mark Lounsbury) who used investigative techniques that, according to a team of security experts, were not in line with established standards. The information technology staff also failed to consider what information may be needed in the event an incident were to occur, as neither they nor the investigative team properly secured the computer in question after the incident occurred.

The first steps that an organization must take have already been discussed in this article; it involves bringing security to the forefront of all aspects of the infrastructure especially as it relates to those systems which gather and store monitoring data. Also important are the access control procedures that tie users to the monitoring data; weaknesses in this regard will jeopardize the value of the monitoring data itself. The second step is far more difficult and requires companies to become adept at understanding the strengths and weaknesses of their infrastructure as well as their end user community. Understanding the impact that the end user community has on the security of the

infrastructure is important, especially as it relates to the process of gathering or even identifying potential evidence.

Organizations such as SANS offer a number of different courses related to the field of information technology forensics. One such course is Computer Forensics and e-Discovery training, which includes lessons about:

- Ethical standards in investigative methods, performing internal investigations, and working with law enforcement agencies.
- Forensic essentials including evidence integrity, evidence volatility, proper selection of investigative tools, and data loss.
- Forensic methodologies such as investigative techniques, identifying potential evidence, and properly documenting findings.
- Investigative techniques including technical concepts that are important to understand when performing investigations.

SANS targets the course toward security personnel, particularly those involved the incident response process:

“Most incident response and security personnel will need to be familiar with core forensic techniques in order to respond to a variety of incidents for their organizations. This course teaches investigators how to follow the trail typical for intrusions and incidents that they might encounter. Incident responders should learn how intruders breached the infrastructure to identify additional systems/networks that are compromised. You will learn how to investigate traces left by complex attacks using the latest exploit methodologies”.

While targeted toward those involved in the incident response process, this SANS course and others can assist companies in understanding the strengths and weaknesses of their infrastructure, and how those findings impact the investigative process.

The SANS Forensics and e-Discovery class teaches a number of concepts that can help information technology professionals identify weaknesses within their infrastructure that could lead to the loss of evidence. Many of the concepts are technical in nature but not something that many individuals would need to know in order to design, implement, or maintain an information technology infrastructure. A closer examination of the topics covered by SANS identifies several subjects that could assist

organizations in building an infrastructure that, by design, supports more thorough and reliable investigations. Table 3 shows how the topics covered in the SANS class might influence the design and management of an IT infrastructure.

Topic from the SANS Class	Applying that Topic to Infrastructure Design and Management
Evidence integrity and volatility	Organizations that understand the challenges that evidence typically faces will better understand how to bolster the reliability of their own data.
	Understanding the impact that the infrastructure as a whole has on the reliability of evidence can assist organizations in identifying potential areas for improvement.
File system guidelines	Organizations that understand (from an evidentiary standpoint) the consequences of choosing one file system over another may make different choices when building out the infrastructure.
	Choices in file systems can impact evidence availability. Examples of actions that can differ from one file system to the next include: what happens when files are deleted and how that impacts the ability to recover data, what if any metadata is stored concerning files, and how files are written to the system and how that will impact a forensic analysis.
Chain of custody	Understanding how the concept of separation of duties, specifically the lack thereof, can impact the chain of custody of evidence.
	Organizations that understand the impact that their infrastructure management policies and data management procedures have on the chain of custody will be able to implement changes to mitigate the issue.

Investigative techniques involving Windows restore points, shadow copies, and registry data	The native features of various infrastructure components can support or extend the capabilities of an investigation.
Evidentiary value of various Windows metadata (Office files, thumbnails, recycle bins, link files, and prefetch and superfetch files)	The generation of various sources of Windows metadata can be controlled using Active Directory group policies. Organizations should consider how current or proposed policy changes will limit or extend the availability of data that may be needed for an investigation.
Password recovery techniques	Understanding the difficulty involved in various password recovery processes may lead an organization to enact guidelines or restrictions on what types of password protection should be used internally.
	Organizations that allow use of unmanaged high bit encryption internally may encounter difficulties during an internal investigation; an alternative would be to deploy an internal public key infrastructure with a key archival and recovery capability.
Legal authority needed to gather data	Organizations must be mindful of local and federal laws concerning data gathering activities, and should ensure that any activities are outlined as needed in employee manuals or terms of use agreements.
Anti-forensics	Organizations that are familiar with techniques used by individuals to conceal or otherwise disguise their activities will at the very least understand their own weaknesses, and ideally will take steps to enhance the security and monitoring techniques used within their infrastructure.

Table 3: SANS course topics

The topics covered in the SANS course will provide information technology architects and administrators with a greater understanding of the challenges that their networks will face should they become a component of or subject of an investigation. While SANS aims this class toward those that are on an incident response team, there is no doubt that if all levels of the technical staff had this knowledge they could position the infrastructure to provide better support to those that handle investigative duties.

Conclusion

One of the more difficult concepts for an organization to adopt is how to approach their data from an evidentiary standpoint. During my work as an IT consultant, I encountered clients with multiple platforms that gather data which could be of value during an investigation that involves the infrastructure. However, data from only a handful of these platforms was typically retained. There is no specific reason why the other data is not considered as important; the only correlation that could be drawn was that the systems whose data is archived are typically involved in the general auditing process of the organization. These organizations need to approach this data not with bias toward perceived value, but with the assumption that if or when the need were to arise, such data may represent another, if not the only, source of evidence.

This article outlines four key methods that organizations can employ to increase the availability and integrity and any potential sources of evidence originating from their information technology infrastructure:

- Separation of duties – While the concept of separation of duties is not anything new, it took legislation such as Sarbanes-Oxley to get organizations to think about how it impacts information technology security. Sarbanes-Oxley introduced separation of duties into the IT field as way to prevent software developers from making changes to production systems that could impact the financial statements of the organization. The same concept can be easily extended to data that may need used as evidence; a lack of control over how data is created, transferred, and even stored may cast its authenticity or reliability as evidence into doubt.
- Security as Part of the Model, Not an Afterthought – Contrary to what staffing firms, manufacturers, and solution providers will tell you an organization cannot hire, buy, or consult its way to perfection; at least not permanently. Organizations that are concerned about their security, including the integrity of the information they use when assessing it, must think of security as a never ending journey that requires teamwork, expertise, honesty with one another, and most importantly an agreement on the goals among all those involved. Ideally, security will become the standard that the organization follows, and not something they resign to resolve after the fact. The unfortunate truth is that an organization with a lax or inadequate security stance will find that the reliability of any evidence they produce will be unpredictable at best, or easily repudiated at worst.
- Trust, but Verify – Separation of duties is a valuable tool for helping organizations reduce the likelihood of evidence being compromised by individuals within the organization, but it does not

make it generally more reliable against scrutiny. Trust, but verify tells us that we must lock down and secure the entire process of identifying the individual who is generating the data, then the process that gathers, transports, and secures data for later review. The reason why it is so important to verify the individual that is generating the evidence is obvious; evidence that cannot be tied to a specific individual is not very useful during an investigation. Once that evidence is generated, the process by which it is gathered, transferred, stored, and archived must be secured and protected to prevent unauthorized review or even tampering.

- Foundations in Forensics – While most people have an understanding of what forensics is, the techniques used by those who practice it are far less known. Though organizations such as SANS offer courses targeted toward individuals who participate in internal incident response teams, the reality is that these classes can help any information technology professional understand issues their organization may face were to an investigation occur. Any organization concerned about their ability to contribute to an investigation should consider sending some of their information technology staff to training such as that offered by SANS. Classes like the SANS Computer Forensics and eDiscovery will provide staff with the skills needed to ensure that the infrastructure can provide meaningful assistance during an investigation. Particularly important is the training in anti-forensic methods; such training will help an organization understand what types of challenges they may face should an investigation take place, and such knowledge will enable them to better prepare.

This article provides an introduction into some of the difficulties that an organization will need to anticipate in the event that they are asked to participate in an investigation. A strong security posture is certainly an important first step, but if an organization is serious about protecting potential evidence they need to take a close look at their data protection and infrastructure monitoring methods, policies, and procedures.

Appendix

BBC.CO.UK, *Pop-up porn case to get new trial*, Retrieved from:

<http://news.bbc.co.uk/2/hi/technology/6729905.stm>

Byrd, Payton, *What Does the Sarbanes-Oxley Act Mean to Developers?*, Retrieved from:

<http://it.toolbox.com/blogs/paytonbyrd/what-does-the-sarbanesoxley-act-mean-to-developers-1029>

Computer Security Institute, *CSI Computer Crime and Security Survey 2010/2011*, Retrieved from:

<http://gocsi.com/survey>

Faqs.org, *RFC 3164 – The BSD Syslog Protocol*, Retrieved from: <http://www.faqs.org/rfcs/rfc3164.html>

Findlaw.com, *Sarbanes-Oxley Act of 2002*, Retrieved from:

<http://f11.findlaw.com/news.findlaw.com/hdocs/docs/gwbush/sarbanesoxley072302.pdf>

HP.com, *HP Network*, Retrieved from: [http://h17007.www1.hp.com/us/en/products/network-](http://h17007.www1.hp.com/us/en/products/network-security/index.aspx)

[security/index.aspx](http://h17007.www1.hp.com/us/en/products/network-security/index.aspx)

Information Systems Audit and Control Association, *ISACA Glossary of Terms*, Retrieved from:

<http://www.isaca.org/Pages/Glossary.aspx?tid=815&char=S>

Information Systems Audit and Control Association, *Systems Development Life Cycle – Definition*,

Retrieved from: <http://www.isaca.org/Pages/Glossary.aspx?tid=877&char=S>

Lacey, David, *The Advancing Science of Anti-Forensics*, Retrieved from:

http://www.computerweekly.com/blogs/david_lacey/2008/10/the_advancing_science_of_antif.html

LANDesk, *LANDesk Server Manager User Guide, pg. 165-168*, Retrieved from:

<http://community.landesk.com/support/docs/DOC-1391>

Microsoft.com, *AD CS and PKI Step-by-Steps, Labs, Walkthroughs, HowTo, and Examples*, Retrieved

from: <http://social.technet.microsoft.com/wiki/contents/articles/4797.ad-cs-and-pki-step-by-steps-labs-walkthroughs-howto-and-examples.aspx>

Microsoft.com, *Configure security for Office 2013*, Retrieved from: <http://technet.microsoft.com/en-us/library/ff400327.aspx>

National Security Agency, *Security Configuration Guides*, Retrieved from:

http://www.nsa.gov/ia/mitigation_guidance/security_configuration_guides/

PCWorld, *Spyware Case Finally Closed for Teacher Julie Amero*, Retrieved from:

http://www.pcworld.com/article/154366/spyware_teacher_case.html

Perimeter eSecurity, *SaaS Network Security*, Retrieved from:

<http://www.perimeterusa.com/services/network-security/>

Reagan, Ronald, *Farewell Address to the Nation*, Retrieved from:

http://www.reaganfoundation.org/pdf/Farewell_Address_011189.pdf

SANS, *Computer Forensics and e-Discovery*, Retrieved from: <http://computer-forensics.sans.org/>

Securitypark.net, *S610f fingerprinting card reader offers card, PIN, and biometric verification in one device*, Retrieved from: http://www.securitypark.co.uk/security_article262016.html

Smallvoid.com, *Restrict Guest Access to Event Logs*, Retrieved from: <http://smallvoid.com/article/winnt-eventlog-guest.html>

Smartcardalliance.org, *Smart Card Primer*, Retrieved from:

<http://www.smartcardalliance.org/pages/smart-cards-intro-primer>

Sunbelt Software, *Technical Review of the Trial Testimony State of Connecticut vs. Julie Amero*,

Retrieved from: <http://www.sunbelt-software.com/ihs/alex/julieamerosummary.pdf>

SurfControl.com, *SurfControl Web Filter System Requirements*, Retrieved from:

http://kb.websense.com/pf/12/webfiles/WBSN%20Documentation/Surf%20Documents/SurfControl_Web_Filter_for_Windows_System_Requirements.pdf

Symantec.com, *Symantec Expands Data Encryption Options for Veritas NetBackup*, Retrieved from:

http://www.symantec.com/en/ca/about/news/release/article.jsp?prid=20061212_01

Techtarget.com, *What is systems development life cycle (SDLC)?*, Retrieved from:

http://searchsoftwarequality.techtarget.com/sDefinition/0,,sid92_gci755068,00.html

U.S. House of Representatives, *Executive Summary – Systems Development Life Cycle*, Retrieved

from: <http://www.house.gov/content/cao/procurement/ref-docs/SDLCPOL.pdf>

Websense, *Websense Small Deployment*, Retrieved from:

http://kb.websense.com/pf/12/webfiles/KB%20Article%20Images/WWS-WWF/v7/Small_Deploy.gif

Websense, *Websense Web Security*, Retrieved from:

http://www.websense.com/assets/pdf/WebSecurity_Datasheet_2B.PDF

Windowsecurity.com, *Microsoft ISA Server, Part I*, Retrieved from:

http://www.windowsecurity.com/articles/Microsoft_ISA_Server_Part_I__introduction_installation_configuration_Web_caching_and_Internet_access.html

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.