

(IN) SECURE

OPEN. INFORMATIVE. TO THE POINT. Issue 16 - April 2008



THE EFFECTIVENESS OF SECURITY CERTIFICATIONS
SECURITY POLICY CONSIDERATIONS FOR VIRTUAL WORLDS
ELECTIONS AND CYBERCRIME

SECURITY AS A SERVICE

NOW AVAILABLE AT A BROWSER NEAR YOU

Software-as-a-Service (SaaS) has been described as the most disruptive delivery model to ever face the enterprise software market for one simple reason: *it works*

Qualys is the first company to deliver an on demand solution for security risk and compliance management. QualysGuard® is the widest deployed security on demand platform in the world, performing over 150 million IP audits per year – with no software to install and maintain.

For a free trial, go to a browser near you.

www.qualys.com/SaaS_Trial



TABLE OF CONTENTS

- Page 05 - **Corporate security news**
- Page 08 - Security policy considerations for virtual worlds
- Page 11 - US political elections and cybercrime
- Page 14 - Using packet analysis for network troubleshooting
- Page 23 - **Latest additions to our bookshelf**
- Page 27 - The effectiveness of industry certifications
- Page 31 - Building a secure future: lessons learned from 2007's highest-profile security events
- Page 34 - Advanced social engineering and human exploitation, part 2
- Page 39 - **Events around the world**
- Page 41 - Interview with Nitesh Dhanjani, Senior Manager at Ernst & Young
- Page 45 - Is your data safe? Secure your web apps
- Page 47 - RSA Conference 2008
- Page 57 - Producing secure software with security enhanced software development processes
- Page 68 - Network event analysis with Net/FSE
- Page 74 - Security risks for mobile computing on public WLANs: hotspot registration
- Page 78 - **Security software spotlight**
- Page 79 - Black Hat Europe 2008 Briefings & Training
- Page 83 - A Japanese perspective on Software Configuration Management
- Page 86 - Windows log forensics: did you cover your tracks?
- Page 97 - Traditional vs. non-traditional database auditing
- Page 101 - Payment card data: know your defense options

Welcome to (IN)SECURE 16 the digital security magazine



It's been a busy few months since the last issue of the magazine. We've attended InfoSec World in Orlando, the Black Hat Briefings in Amsterdam as well as the immense RSA Conference in San Francisco. It was a pleasure meeting our contributors and readers and we hope to catch many more of you at upcoming events. This issue contains news from the conferences along with a variety of photos.

This time around we bring forward very interesting topics such as social engineering, digital forensics, security in the software development lifecycle and wireless security problems. Topics we haven't featured before include elections and cybercrime, security policy considerations for virtual worlds and a deep look into the effectiveness of security certifications.

Keep reading and sending feedback, it's always a pleasure to hear from you.

Mirko Zorz
Chief Editor

Visit the magazine website at www.insecuremag.com

(IN)SECURE Magazine contacts

Feedback and contributions: Mirko Zorz, Chief Editor - editor@insecuremag.com

Marketing: Berislav Kucan, Director of Marketing - marketing@insecuremag.com

Distribution

(IN)SECURE Magazine can be freely distributed in the form of the original, non modified PDF document. Distribution of modified versions of (IN)SECURE Magazine content is prohibited without the explicit permission from the editor.

Copyright HNS Consulting Ltd. 2008.

Corporate security news



Payment Application Data Security Standard released



The PCI Security Standards Council, a global, open industry standards body providing management of the Payment Card Industry Data Security Standard (DSS), PCI PIN Entry Device (PED) Security Requirements and the Payment Application Data Security Standard (PA-DSS), announced the release of version 1.1 of the Payment Application Data Security Standard (PA-DSS). Following release of the PA-DSS, this fall the Council will also roll out a program to include maintenance of a list of validated payment applications. This list will enable buyers to identify the payment applications that have been recognized by the PCI SSC and meet the new standard. (www.pcisecuritystandards.org)

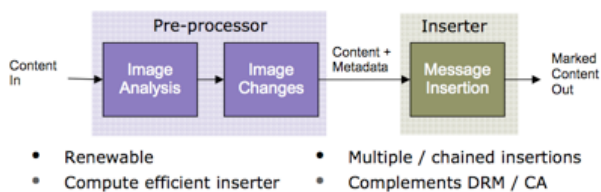
Sony enhances lineup of security cameras

Sony is entering the megapixel marketplace with several new additions to its lineup of IP-based security network cameras. Sony's first megapixel cameras include the SNC-DM110 Megapixel Normal mini-dome, SNC-CM120 Megapixel CS mount Day/Night, and the SNC-DM160 Megapixel Rugged Day/Night mini-dome models.

The units are the first to feature Light Funnel technology, which combines image data gathered from multiple horizontally and vertically aligned pixels to provide extremely bright image output even when monitoring moving objects. This function can be activated automatically in response to surrounding light conditions or on a pre-specified time schedule. (www.sony.com)



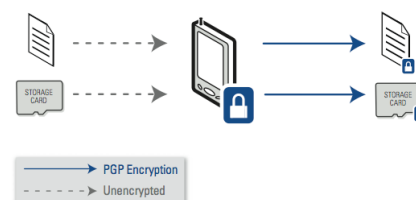
Forensic watermarking of encrypted content



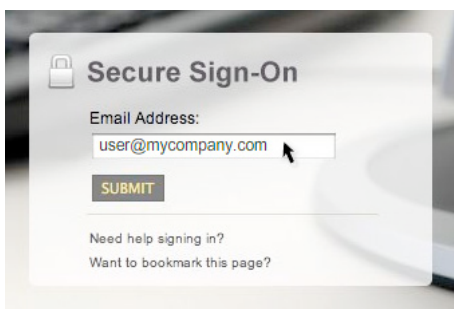
Cinea announced that its Running Marks technology now enables traceable, stream-specific forensic watermarking of encrypted content. The new approach allows Running Marks to embed a unique serial number into individual customer video streams, without requiring providers or system operators to expose the high-value content in an unencrypted format. The new functionality enables distributors of on-demand content in the cable, telecommunications, satellite, and video-download markets to provide a secure distribution chain that minimizes cost, time, and exposure to theft and piracy. (www.cinea.com)

PGP brings enterprise data protection to smartphone users

PGP released PGP Mobile, an encryption application that allows enterprise users to easily protect data on smartphones. PGP Mobile joins an award-winning family of applications that are part of the PGP Encryption Platform, enabling organizations to protect data while reducing the operational costs associated with managing encryption keys, users, policy, and reporting for multiple point encryption products. (www.pgp.com)



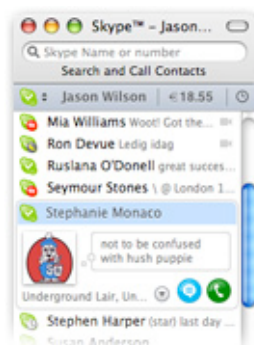
TriCipher secures access to social networks



TriCipher announced the integration of popular social network sites MySpace, LinkedIn, Plaxo, YouTube, Classmates.com, Friendster and others with myOneLogin, the first Web-based service to combine strong authentication and single sign-on. myOneLogin helps businesses implement controls that protect employees' access to Web-based applications. It streamlines password policies and compliance reporting, mitigates the risk of phishing and eliminates the need for expensive authentication hardware and software. (www.tricipher.com)

Malware prevention for Skype

FaceTime Communications announced enhancements to its Greynet Enterprise Manager including detection of malicious URLs entering the enterprise network via Skype instant messaging conversations. Skype is encrypted using a proprietary method, making it impossible for traditional security products to view the content of a Skype text conversation. Working in partnership with Skype over the last year, FaceTime is the only security vendor with the ability to examine the content of a Skype instant message as it enters the network. Using its leading malware signature database maintained by FaceTime Security Labs, FaceTime's products verify that content is safe and free of malicious URL links before entering the network. (www.facetime.com)



Red Hat Certificate System source code released

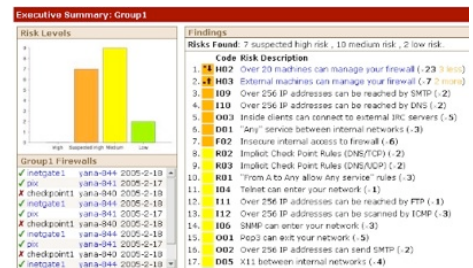


Much of the technology in Red Hat Certificate System was already open source, including the Apache web server, Red Hat Directory Server and the FIPS140-2 level 2 validated NSS cryptographic libraries, but this move further demonstrates Red Hat's belief that the open source development model creates more secure software. With the Certificate System

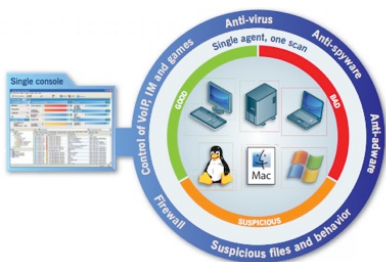
code now available under an open source license, it will be much easier to integrate these proven technologies with other open source projects. One specific example of this is the Red Hat-sponsored freeIPA project. freeIPA provides central management of (I)dentify, (P)olicy and (A)udit for the Unix and Linux world through the use of open source and open standards. (pki-svn.fedora.redhat.com)

New Firewall Analyzer product suite

AlgoSec announced the availability of the AlgoSec Firewall Analyzer (AFA) product suite. The new suite, which improves the overall security and efficiency of enterprise firewalls, is built on three distinct software modules, each focused on addressing a specific set of technical requirements within the enterprise: Firewall Operations Management, Policy Optimization and Risk Management. AlgoSec's firewall analysis solutions work across all available platforms from the big 3 enterprise firewall vendors: Cisco, Check Point and Juniper/NetScreen. (www.algosec.com)



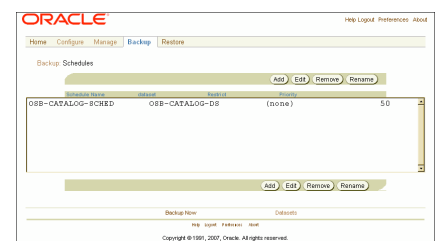
New Sophos Endpoint Security and Control 8.0



Sophos announced Sophos Endpoint Security and Control 8.0, incorporating Network Access Control technology to go beyond reactive and proactive anti-virus – offering businesses complete preventive computer protection. Since last year's acquisition of Endforce, Sophos has seamlessly integrated the core policy control and assessment capabilities into its flagship endpoint solution without requiring a separate agent deployment or additional software license costs. (www.sophos.com)

Oracle releases Secure Backup 10.2

Oracle announced the general availability of Oracle Secure Backup 10.2, Oracle's high-performance tape backup solution for Oracle Databases and NAS storage devices. It provides policy-based encryption at the domain, host, backup or tape level using AES128, AES192 or AES256 encryption algorithms. Oracle Secure Backup provides automated management of all encryption keys associated with tape backups, simplifying the job of the database administrator. Encryption keys are centrally stored on the Oracle Secure Backup Administrative Server, which seamlessly manages decryption during restoration. (www.oracle.com)



Security policy considerations for virtual worlds

By Jeff Surratt



Increasingly virtual worlds, sometimes called 3D Internet environments, offer significant outreach and business development opportunities to companies, governments, and the world at large.

These Virtual worlds – such as Second Life, World of Warcraft, Entropia Universe, EVE On-Line and others – allow people to interact through digital personas or avatars. As these worlds evolve and grow in popularity and acceptance, and become more integrated into many aspects of business and society, they offer new and uncharted terrain for security practitioners to embrace, explore and apply corporate governance and information security policy.

As security practitioners, there are many things to consider before advising our business leaders on how to make the leap into virtual realms. In many cases, the old rules apply, but some things we simply have not had to think about before. Walk into the uncharted terrain, but go with trepidation.

Terms of Service - Many third-party companies now provide virtual services to individuals

and businesses. Most require a monthly fee but some are free of charge. All of them require participants to agree to the company's terms of service. These terms of service are an attempt to protect the virtual world service provider's control over all aspects of the service, content and data generated in that virtual world. Thus, ability to remain a member of a virtual world or to have or use a space within that world is behavior dependent and not always guaranteed.

Each virtual world provider has its own unique characteristics and terms of service. In signing up to participate, businesses should fully understand the terms and conditions to which they are agreeing as members of that community. Users who sign up for service should also recognize that, unless specifically directed by their management, they are signing those terms and conditions as an individual.

Therefore, the individual and not the business are responsible for all aspects of participation in a virtual world. If an account is for business purposes, ensure that it is not paid for with personal funds and review the terms of service with appropriate legal counsel.

Public Forum - It is also important to remember that virtual worlds are public, software-based, open societies in which having a dialogue is similar to having a discussion or meeting in a public place such as a hotel lobby or an airport. Individuals acting for themselves or as part of a directed business venture should operate on the assumption that all actions, communications and data can be seen, heard and recorded by anyone, including the service provider – which may not,

and often does not, have any obligation to protect your communications or information.

In the conduct of personal or employer business, many rules of the physical world apply to the virtual world(s):

1. Do not run the client on an account with Administrative privilege.
2. Do not disclose proprietary information or talk about company business in a non-company forum.
3. Do not use the same password to access the virtual world as you do for internal company or personal business.
4. Never give out your password.
5. Follow the company dress code.

Be mindful that all actions will be public and may be visible for a long time.

Now you may be asking yourself “Follow the dress code?” Yes! If your avatar is conducting company business then follow the company dress code. In many virtual worlds, there are many clothing types. Do you want users to conduct business as (or with) a ring-tailed lemur in bunny slippers or would you rather work with an avatar in business casual attire? If your business does not enforce dress codes, then insist people use good judgment.

Be mindful that all actions will be public and may be visible for a long time. Many software tools exist to screen capture. If you or your users can see other avatar, they can see you. Every user within sight of an avatar has the ability to immortalize a single poor choice for the entire Web to see.

Understand the software client and connection – Each virtual world uses a client to connect to the server cluster. Every virtual world mentioned above requires that numerous ports and protocols are opened through the corporate firewall. Contact each virtual world service provider and research how each is open port is used. Just like their web counterparts, virtual world clients are subject to attacks.

The years 2006 and 2007 for example, saw an increase in the number of malware and Trojan programs written with the primary purpose of stealing passwords from virtual world users. These malicious programs utilized exploitable vulnerabilities (for example, the Web browser) to install password-stealing software or account "harvesting" programs.

The virtual world developed by Linden Lab, Second Life, has a client with its own XML HTTP Request that uses asynchronous callbacks; it gives the platform the capability to communicate with the Web on demand.

Every object is scriptable and can be aware and active. In Second Life, a QuickTime file can automatically play on the machine of a user who enters another user's virtual land or accepts a scripted object. Once played, the malicious QuickTime file can cause the user's machine to do anything that file tells it to do (especially if the user is running Second Life under an account with administrative access (remember rule number 1!).

If you build it, accept that good and bad will come – Aside from password stealing on the client side, there is fraud that can be committed within the virtual world itself.

There is also the potential for exploitation of scenarios that the developers (or users) never envisioned when designing the world or user created content.

For example: The attacks against Second Life such as the "Grey Goo" infestation where replicating objects brought about a shut down of the Second Life Grid for all but Linden Lab staff or the exploit in Blizzard Entertainment's World of Warcraft that allowed "item duping" before the developers implemented a patch.

Ok, so what is the big deal with duplication? In the case of the "Grey Goo" item replication resulted in a denial of service (DoS) due to database loads. Outside of that, in virtual worlds, items hold a value. Just as it is in the

real economy value is determined by demand or rarity. If I can flood the market with perfect copies of your virtual product then I can drive down demand by making it less rare or I can sell my copies for much less and undersell you. In Second Life, this could lead to legal consequences as the users hold the intellectual property rights to items they create. If a user is in the virtual world on a directed business initiative, the employer may hold the rights depending on the employment contract or agreement.


There is much to gain by embracing a growing population of virtual world users, but there is nothing virtual about the business consequences of a lapse in judgment.

Jeff Surratt (CISSP MSIA) has seventeen years of experience in IT and information security. Jeff holds a degree in Internetworking Technology from Strayer University and a Masters degree in Information Assurance from Norwich University.

Want to reach a large audience of security professionals by writing for (IN)SECURE?



Send your idea to editor@insecuremag.com



US political elections and cybercrime

By Oliver Friedrichs

Phishing, adware, and crimeware. These issues, along with a variety of others, have long affected consumers and enterprises alike; but can these also infringe upon the forum of political campaigns? Recent research indicates that potential IT security risks can and do pose a threat to the nature of the electoral process.

Over time, the Internet has become an increasingly popular and effective mode of communication between political candidates and constituents. As candidates continue to leverage the Internet to connect with people - whether it be to communicate their opinions, rally supporters, or to sway critics--it is vital to understand the potential IT security risks associated with the increased use of technology in the election process. Some of these risks include, among others, the diversion of online campaign donations, dissemination of misinformation, fraud, phishing, and privacy invasion. Some of these attacks have the potential to hamper voters' faith in our electoral system.

Research has revealed that many of the Internet-based attacks which affect enterprises and consumers can materialize within the election process as well. In hopes to increase awareness of these and other plausible risks in the 2008 or future elections, Symantec has assessed Internet-based attack agents that are most likely to immediately affect the 2008 presidential election (or any fu-

ture election), and how these vectors may affect the election process up until voting day.

Symantec concludes that perpetrators would likely fall into two possible categories: those with political motives and those seeking to profit from these attacks. However, identifying an actual motive based on a given attack may be difficult. Symantec discovered several potential threats, including abuse of candidate's Web domains and malicious attacks.

Abuse of candidate's Internet domain names and typo squatting

It would be quite easy to use a domain name of this type for phishing or crimeware installation, since voters may not know the URL associated with their political party or their chosen candidate's web site. Also, legitimate-sounding domain names may not be as they appear.

In order to determine the current level of domain name speculation and typo squatting

in the 2008 federal U.S. election, Symantec performed an analysis of 17 well-known candidate domain names in order to seek out domain speculators and typo squatters. From this analysis, the company drew two clear conclusions: firstly, a high number of typo and cousin (correctly spelled with additional wording) domain names have been registered by parties other than the candidate's own campaign. Several of the registered web sites are registered in order to boost traffic to advertising web sites. Many typo and cousin domain names were registered; some of the typo domain names were being used in bad faith, while the cousin domain names were being used in support of a candidate, as well as to detract from a candidate.

Secondly, most candidates have not done a good job at protecting themselves by proactively registering typo domains to eliminate possible abuse. Symantec was only able to find one typo web site that had been registered by a candidate's campaign (www.mittromney.com).

Phishing

During the 2004 presidential election, phishing was still in its infancy and therefore posed minimal risk. Now, however, the situation is altogether different. Today, candidates extensively use the Internet to communicate with supporters and raise campaign contributions electronically.

A phisher may pose as a candidate and ask the recipients for money, with the typical goal being to steal the credentials of his victims. Attackers may also attempt to disrupt a candidate's fundraising efforts by introducing illegitimate payments from stolen credit cards, which causes confusion. This can be an issue regardless of whether the candidate uses email for contribution requests or not.

Symantec performed an analysis of campaign web sites to find out the degree to which they allow contributions to be made online. Because typical Internet users would not be readily familiar with political candidate domains, a risk exists that phishers would use a similar web site in order to gather credentials from victims.

Phishers can copy legitimate fundraising emails easily to make people submit their sensitive information or download crimeware. This can be done with the direct goal of disseminating malware, or also to deliver a negative blow to political candidates who tend to rely largely on the Internet for interacting with their supporters.

Another attack may involve the diversion of donations intended for one candidate to the web site of a different candidate. This is a more worrisome attack, since the potential impact is financial and psychological, weakening a contributor's faith in electronic donations.

Campaigns can take immediate and clear steps to invest in typo domains before they fall into the wrong hands. As of now, however, many have failed to do so.

Adware and spyware

Adware may be used in several different ways to manipulate or influence users during the election process. First, adware may simply present the user with ads promoting a candidate, directing the user to the candidate's web site when clicked. Adware may also be surreptitiously used to replace advertisements for one candidate with that of a different candidate. This may help sway undecided or uncommitted voters.

Spyware has the capability of capturing and recording user behavior patterns (including Web browsing, online campaign contributions, party affiliation, and email traffic) without users' knowledge or consent. This dramatically alters the landscape when it comes to election-related data collection, and poses a new risk to the accumulation of election-related statistics used to track trends in elections.

Keyloggers and crimeware

Crimeware can collect person information about people which malicious actors can leverage to intimate voters or hold for ransom to sway votes. A keylogger has the ability to cause material damage to a candidate. Such code may be also targeted towards campaign staff or others who may be deemed material to the particular candidate's efforts.

An infection like this can result in the monitoring of all communications, including email messages and web site access initiated on the infected computer. This monitoring would provide the would-be-attacker with unsurpassed insight into the plans, progress, and disposition of the campaign, including speeches and otherwise sensitive information vital to the candidate's campaign.

Campaign web site security

The breaching of a candidate's web site would enable an attacker to have direct control over all content observed by visitors, thereby potentially allowing for the posting of misinformation or the deployment of malicious code to unsecured visitors. Examples of misinformation include a candidate's decision to drop out of the race, legal trouble, health issues, or a fake scandal. It may also include subtle information that could be portrayed as legitimate, such as a shift in position on a particular issue, thus possibly leading to a loss of supporters who may feel strongly about that particular subject.

Denial of Service attacks

Several high profile and wide scale attacks have demonstrated the effect that a distributed denial of service (DDoS) attack can have. The attack launched against the country Estonia in May 2007 is one of the best known largest attacks, and serves as an example of one that was politically motivated.

In 2006, Joe Lieberman's web site also faced a DDoS attack. The attack forced the site offline, paralyzing the joe2006.com domain, thus preventing campaign officials from using their official campaign email accounts.

Therefore, the implications of these attacks are severe and clear: they prevent voters from reaching campaign web sites, and also pre-

vent campaign officials from communicating with these very voters.

Public voter information sources

The Federal Election Commission maintains a record of all campaign contributions, which is available to the public and contains contributor's personal information. This information can allow a person to build a history of political contributions for any U.S. citizen contained in the record. Thus, appearing in this database may expose high-net worth contributors to targeted phishing or malicious code attacks if that person's name can be connected to their email address; this can make those individuals listed in the record increasingly susceptible to some of the attacks discussed earlier.

Intercepting voice communications

The infection of a candidate, campaign staff, or candidate's family's cell phone could have grave consequences. Private and personal conversations could be monitored. Moreover, opinions and facts not shared with the public have the potential to be heard and recorded, making room for widespread negative exposure or damage, thus affecting the campaign itself as well the candidate's private life and personal affairs.


As campaigns increasingly rely on the Internet to bolster their campaigns, it is very important to realize the possible risks that come along with it. Although this article discusses some of these key risks, it is important to consider that there are several other remaining ones, as well as more complex combination threats. Research has shown that some campaigns are steadily beginning to understand the risks of online advocacy, but have yet to take the vital steps and precautions in order to protect themselves.

Oliver Friedrichs is the Director of Emerging Technologies in Symantec Security Response, the organization which is responsible for the delivery of antivirus definitions, intrusion detection updates, and early warning technologies.

Prior to his role at Symantec, Oliver served as co-founder and director of engineering at SecurityFocus until the company's acquisition by Symantec in 2002, as well as the co-founder and vice president of engineering at Secure Networks, Inc. He has over 15 years of experience in security technologies and has shared his expertise with many of the world's most powerful organizations, including the Department of Homeland Security, the IRS, and NASA.

Using packet analysis for network troubleshooting

By Chris Sanders



A million different things can go wrong with a computer network on any given day – from a simple spyware infection to a complex router configuration error – and it is impossible to solve every problem immediately. The best we can hope to do is be fully prepared with the knowledge and the tools it takes to respond to these types of issues.

All network problems stem from the packet level, where even the prettiest-looking applications can reveal their horrible implementations and seemingly trustworthy protocols can prove malicious.

To better understand and solve network problems, we go to the packet level where nothing is hidden from us, where nothing is obscured by misleading menu structures, eye catching graphics, or untrustworthy employees. Here there are no secrets, and the more we can do at the packet level, the more we can control our network and solve problems. This is the world of packet analysis. In this article I will show you how you can effectively use packet analysis to isolate and troubleshoot issues on your network.

How packet sniffers work

Simply put, a packet sniffer is a program used to capture and analyze packets on a network. These packets are the fundamental building blocks of network communications.

The packet sniffing process involved three basic steps: collection, conversion, and analysis. In the first step, the packet sniffer switches the network interface card of the computer it is running on into promiscuous mode. In this mode an NIC can listen for all network traffic on its particular network segment. The sniffer uses this feature to capture the raw binary data flowing across the network hardware. Once this is completed, the conversion process begins and the capture binary data is converted into a readable form.

All that is left at this point is the analysis of these capture packets. This is where a sniffer takes the captured and converted data and verifies what protocol it is. The various features of these protocols are then evaluated by the sniffer and displayed to the user in an easily readable format.

There are a variety of enterprise level and free packet sniffers available. Some of the more popular ones include Omnipcap, TCPDump, and Wireshark. Which one you use is based upon your personal preference. I prefer Wireshark due to its large community support and easy to use interface, so that's what we will be using for the rest of this article.

Capturing packets with Wireshark

The Wireshark program is distributed freely and can be downloaded at

www.wireshark.org. Installing this software is simple enough that I am not going to go through it here, although it is important to note that this software relies on the WinPcap driver which is included in the installation package.

Once you have installed Wireshark along with the WinPcap driver you should be ready to dive in head first. The first thing we are going to do is a simple packet capture. To do this you will first need to select your capture interface by clicking the "List available capture interfaces" button to the far left hand side of the main toolbar. Once you have done this you will be presented with a window listing all available capture interfaces. Clicking the "Capture" button next to the interface you wish to use will begin capturing packets from this connection. Wait for a few minutes until a significant amount of packets have been collected and then click the "Stop" button.



You should then be returned back to the main screen with a whole bunch of new data. Congratulations! You have just completed your first successful packet capture! You may now be asking yourself how you interpret this capture data, but we aren't quite ready for that yet. We need to go over sniffer placement first.

Tapping into the wire

Now that you know how to do a basic packet capture in Wireshark it is important to learn how to capture the right traffic. Assuming you are on a switched Ethernet network (which most everybody is these days), all of the traffic you just captured was your own. That is, all traffic was either coming to or going from the computer from which you initiated the packet capture. This is basically how a switched network functions. The switch only sends data to

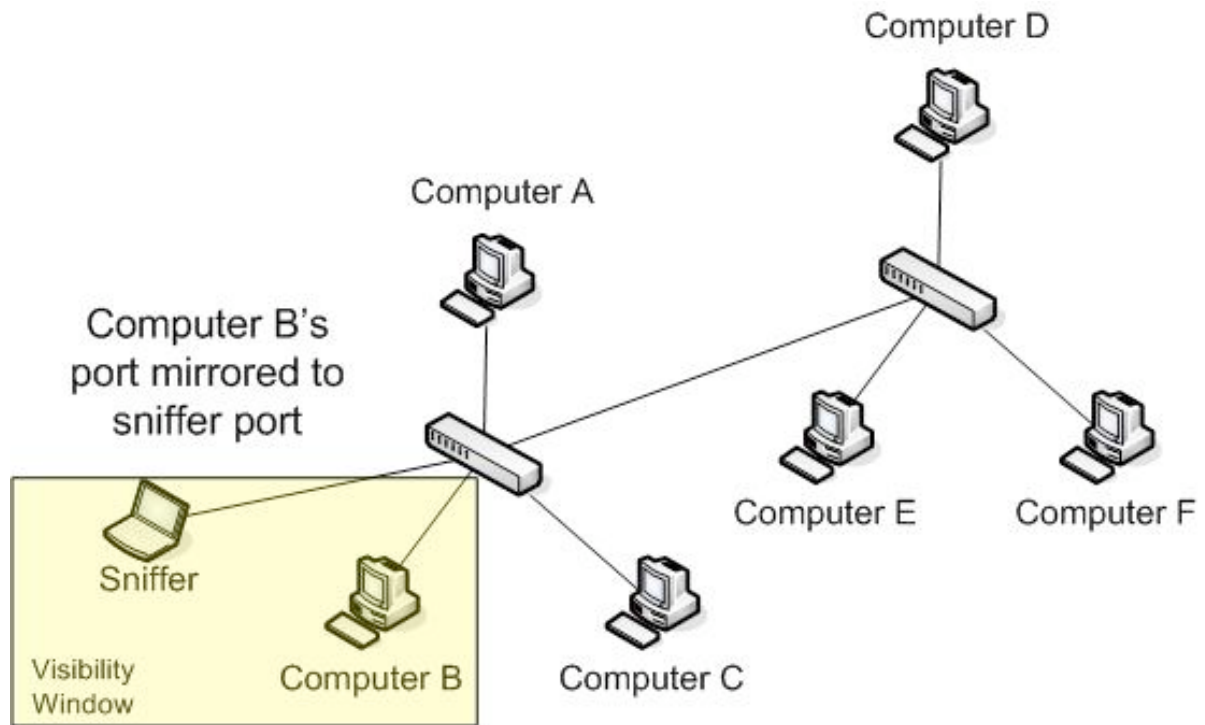
those ports in which it is destined. This brings up the question how do you capture traffic from a computer that packet sniffing software is not installed on?

There are a couple of different methods that can be used in a situation like this. Three of the most common techniques are port mirroring, hubbing out, and ARP cache poisoning.

Port Mirroring is probably one of the easiest ways to capture the traffic you are looking for. Also called port spanning, this is a feature available on most managed network switches. This is configurable by accessing the command line or GUI management for the switch the target and sniffer systems are plugged in to and entering commands which mirror the traffic of one port to another. For instance, to capture the traffic of a device plugged in

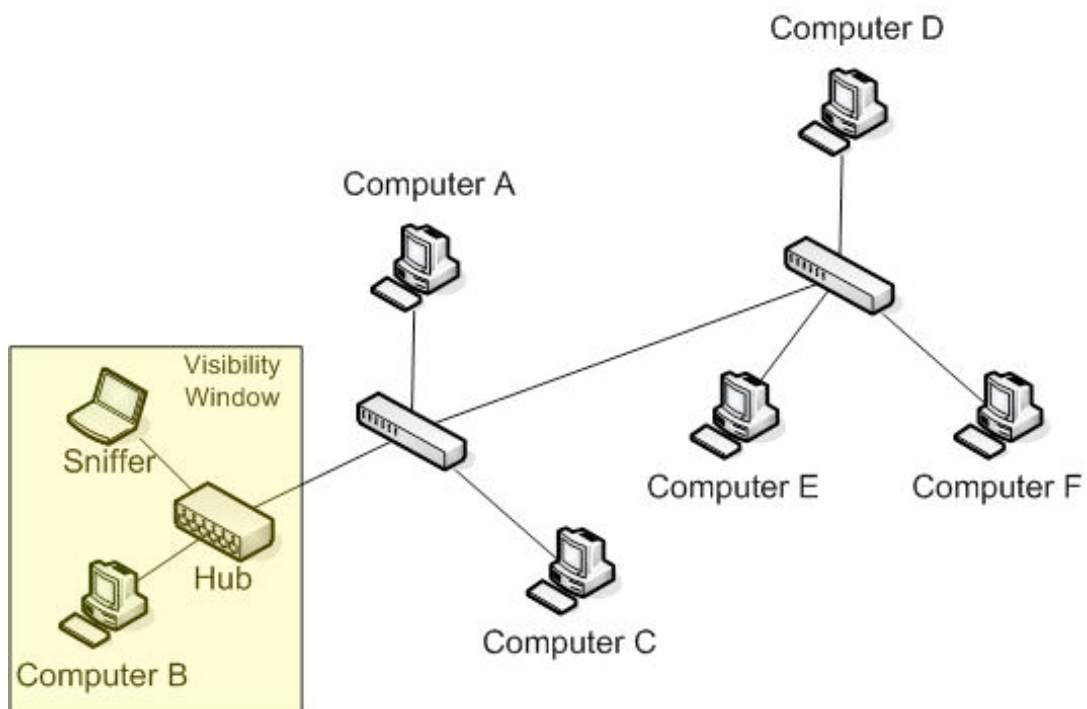
to port 3 on a switch, you could plug your sniffer into port 6 and enter a vendor specific mir-

roring command that mirrors port 3 to port 6.



Hubbing out is a technique in which you localize the target device and your analyzer system on the same network segment by plugging them directly in to a hub. In order to do this, all you need is an old hub and a few network cables. Simply go to the switch that the target computer resides on and unplug it from the network. Plug the targets network cable, along with the cable for your sniffer, into the hub, and then plug the hub into the network switch.

This will put your sniffer and the target machine on the same broadcast domain and allow you to see all of the packets going to and from the target machine, as well as yours. Since this does involve a brief moment of connectivity loss, I do highly recommend letting the user of the target system know that you will be briefly disrupting their connectivity, especially if it is someone in management!



The last and most advanced technique when placing your network sniffing is ARP cache poisoning. This requires the use of some third party tools and a deeper understanding of the ARP protocol, which is a little beyond this article. You can see an explanation and tutorial of this technique here: www.chrissanders.org/?p=113.

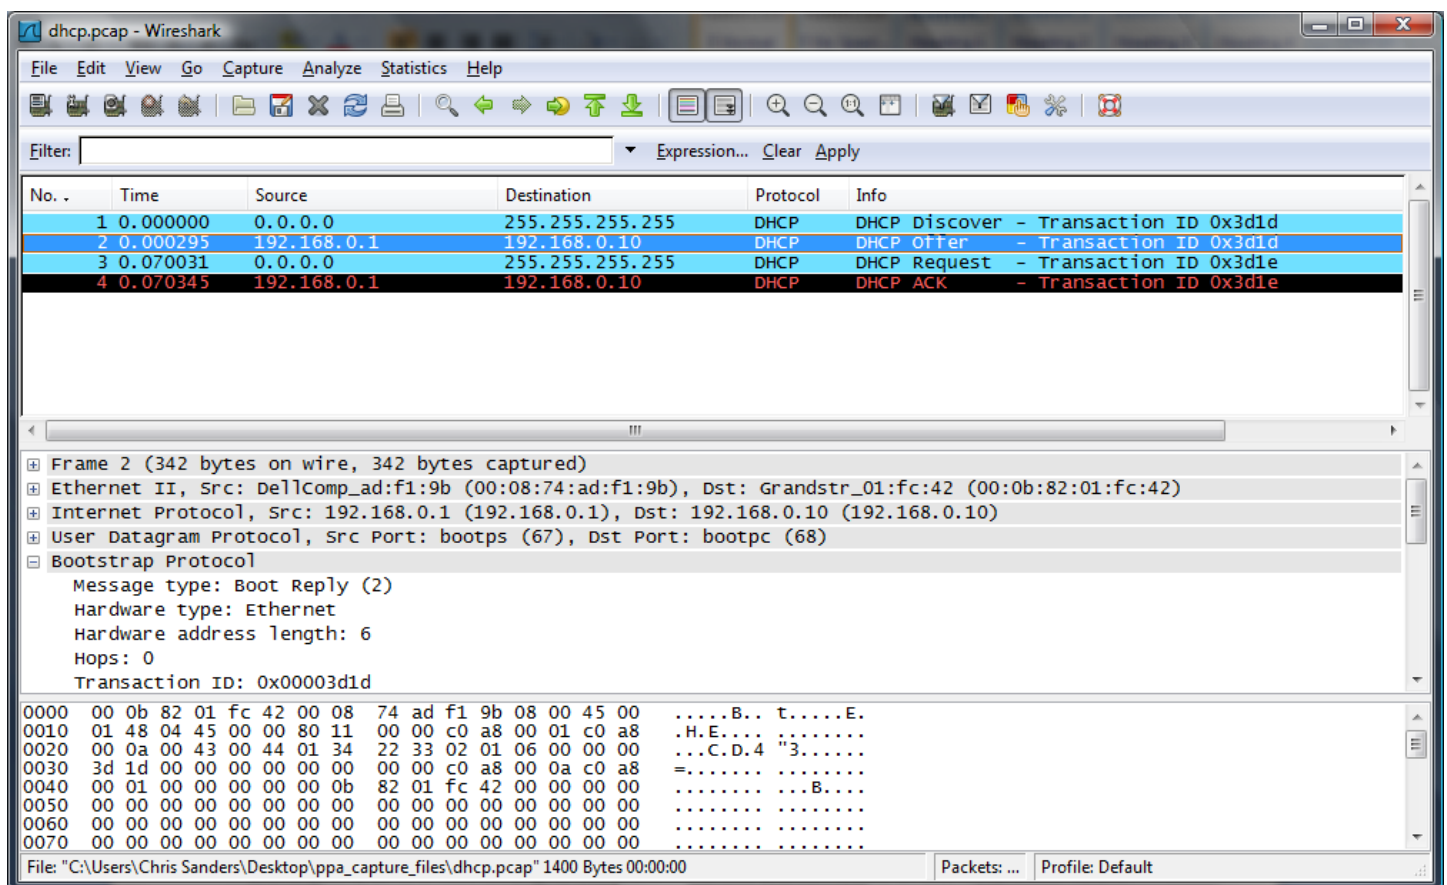
Breaking down Wireshark

Now that you have read the crash course in packet sniffer placement, we can jump back into Wireshark and how to use it to analyze packets. If you still have your first packet capture from earlier open, then great. If not, go ahead and step through that process again so that you have some data to look at.

There are three main sections of the Wireshark program. These are the packet list pane, the packet details pane, and the packet bytes pane. The packet list pane is the top most

window and displays a table containing all packets in the current capture file. This section is divided into various columns including the packet number, the relative time that the packet was captured, the source and destination of the packet, the packets protocol, and some general information about the packet based on its protocol. The packet details pane is just below the packet list pane, and contains a hierarchical display of all information capture about a single packet. The items in this section can be expanded or collapsed for ease of viewing.

The last section is the packet bytes pane. This pane, at the bottom of the window, displays information about the selected individual packet in its raw unprocessed form. It's basically the same data that is displayed in the packet details pane, but without all of the warm and fuzzy stuff to make it easier to interpret. It is very important to understand how these different panes relate to each other.



If you are anything like me, you may have an aversion to shiny objects and pretty colors. If that is the case, the first thing you probably noticed when captured your first packets were the different colors of the displayed packets in

the packet list pane. Each packet is displayed as a certain color for a reason. For example, you may notice that all DNS traffic is blue and all HTTP traffic is green. These colors reflect the packets protocol.

The color coding allows you to quickly differentiate among various protocols so that you don't have to read the protocol field in the packet list pane for each individual packet. You will find that this greatly speeds up the time it takes to browse through large capture files.

Looking at some typical network traffic

Before you can identify problematic traffic you must first understand what normal traffic looks like. Let's quickly take a look at some ARP traffic to get a little more comfortable with looking at traffic on the packet level.

An overview of ARP

The basic idea behind ARP is for a machine to broadcast its IP address and MAC address to all of the clients in its broadcast domain in order to find out the IP address associated with a particular MAC address it wishes to transmit data to. Basically put, it looks like this:

Computer A – “Hey everybody, my IP address is XX.XX.XX.XX, and my MAC address is XX:XX:XX:XX:XX:XX. I need to send something to whoever has the IP address XX.XX.XX.XX, but I don't know what their hardware address is. Will whoever has this IP address please respond back with their MAC address?”

All of the other computers that receive the broadcast will simply ignore it, however, the

one who does have the requested IP address will send its MAC address to Computer A. With this information in hand, the exchange of data can begin.

Computer B – “Hey Computer A. I am who you are looking for with the IP address of XX.XX.XX.XX. My MAC address is XX:XX:XX:XX:XX:XX.

One of the best ways I've seen this concept described is through the limousine driver analogy. If you have ever flown, then chances are when you get off of a plane, you have seen a limo driver standing with a sign bearing someone's last name. Here, the driver knows the name of the person he is picking up, but doesn't know what they look like. The driver holds up the sign so that everyone can see it. All of the people getting off of the plane see the sign, and if it isn't them, they simply ignore it. The person whose name is on the card however, sees it, approaches the driver, and identifies himself.

ARP at the packet level

Understanding the basic concept of ARP, we can take a look at some packets to see how it actually functions. Here we will step through the entire ARP process, start to finish. In this scenario Computer A needs to communicate with Computer B. You can take a look at this capture for yourself by downloading the capture file from www.chrissanders.org/resource/arp.pcap

```
Frame 1 (42 bytes on wire, 42 bytes captured)
Ethernet II, Src: HonHaiPr_6e:8b:24 (00:16:ce:6e:8b:24), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
Address Resolution Protocol (request)
  Hardware type: Ethernet (0x0001)
  Protocol type: IP (0x0800)
  Hardware size: 6
  Protocol size: 4
  Opcode: request (0x0001)
  Sender MAC address: HonHaiPr_6e:8b:24 (00:16:ce:6e:8b:24)
  Sender IP address: 192.168.0.114 (192.168.0.114)
  Target MAC address: 00:00:00_00:00:00 (00:00:00:00:00:00)
  Target IP address: 192.168.0.1 (192.168.0.1)
```

The packet details window of the first packet in the capture file is very straightforward. The computer at 192.168.0.114 needs to communicate with the computer at 192.168.0.1, but doesn't know its MAC address. Notice that the target MAC address here is 00:00:00:00:00:00.

This being the case, it sends a packet with the destination address ff:ff:ff:ff:ff:ff, in turn broadcasting that packet to everything on the current network segment. This is the basic ARP Request packet, as stated in the Opcode field.

```
Frame 2 (46 bytes on wire, 46 bytes captured)
Ethernet II, Src: D-Link_0b:22:ba (00:13:46:0b:22:ba), Dst: HonHaiPr_6e:8b:24 (00:16:ce:6e:8b:24)
Address Resolution Protocol (reply)
  Hardware type: Ethernet (0x0001)
  Protocol type: IP (0x0800)
  Hardware size: 6
  Protocol size: 4
  Opcode: reply (0x0002)
  Sender MAC address: D-Link_0b:22:ba (00:13:46:0b:22:ba)
  Sender IP address: 192.168.0.1 (192.168.0.1)
  Target MAC address: HonHaiPr_6e:8b:24 (00:16:ce:6e:8b:24)
  Target IP address: 192.168.0.114 (192.168.0.114)
```

The second packet is our reply from 192.168.0.1. This device received the ARP Request in step one, and generated this reply addressed to 192.168.0.114. Notice that this reply contains the information that 192.168.0.114 needs to communicate properly. This is the sender MAC address in this second packet. You can tell immediately that this is an ARP reply by looking at the Opcode field.

Once the device at 192.168.0.114 receives the ARP Reply it can then take the MAC address of 192.168.0.1 and put it in its ARP table for future use. With this new information, ARP can successfully translate between layer two and layer three so that communication can move on to the physical medium.

Troubleshooting a real network Issue

Now it's time to get even deeper into the nuts and bolts of network communication. We are going to look at a real network troubleshooting scenario and how to figure it out at the packet level.

In this scenario, an organization has an FTP server that it uses to maintain all of its pre-release software. Lately, the technician in charge of maintaining this server has received several reports from the company's programmers that when they are working late at night the upload/download performance of the server is severely degraded. This FTP server is running a very simple FTP application so the logging features of it don't really give us any usable information. This is the perfect scenario for using a packet sniffer.

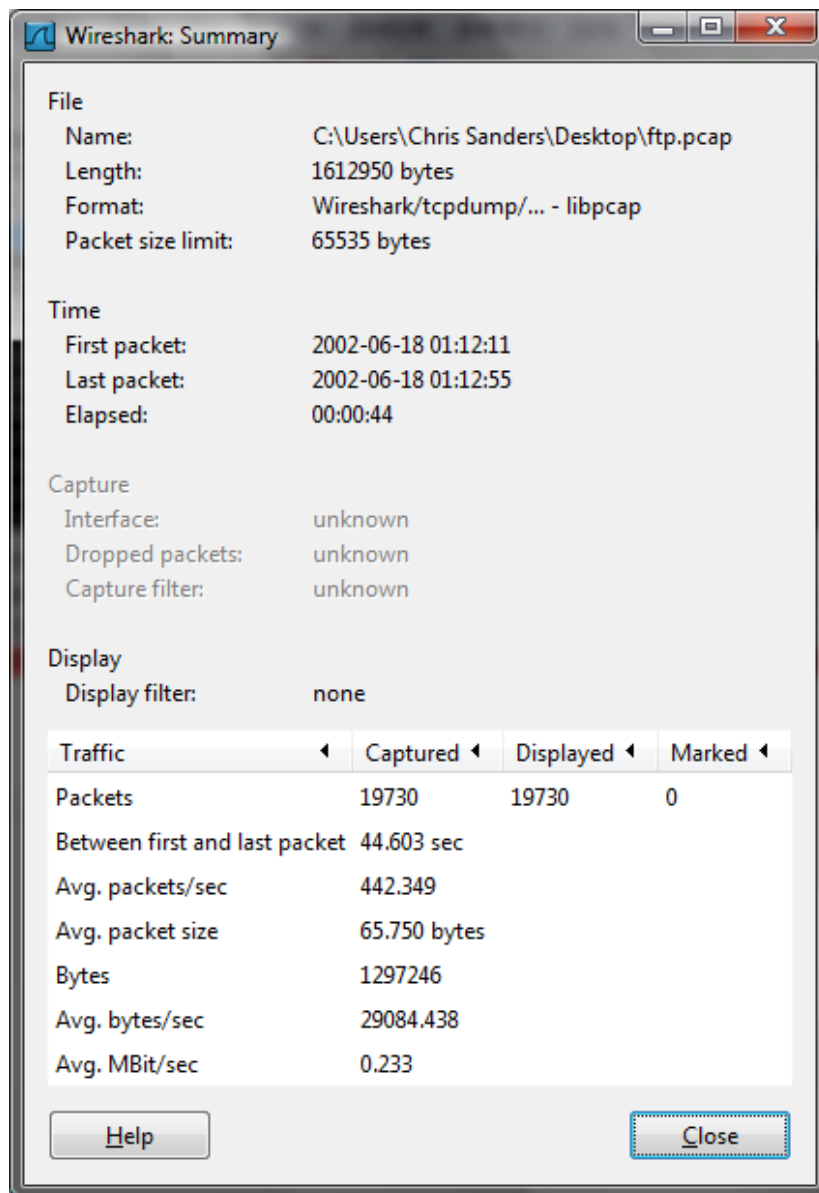
This capture file can be best obtained by using the port mirroring technique to get on the wire and capture the appropriate data. Some might think that you are best off installing Wireshark

directly to the machine in question, but if performance is an issue then this is risky as severe performance issues may cause packets to be dropped and reduce the validity of our capture. If you want to see this sample capture file, you can download it from www.chrissanders.org/resource/ftp.pcap.

When you open this capture file, you will see a whole lot happening in a very short amount of time. Notice the time column in the packet list window shows the first two hundred packets coming across the wire in less than a second. This number doesn't really mean a whole lot since we can't use that too effectively to see the rate of data flowing across the wire, but there is another way to do this. Select Statistics from the drop down menu at the top of the screen and then choose Summary. This will display a summary screen which will give an overview of some statistics for this entire capture process.

If you look at the last piece of data displayed on the screen, you will see that the average MBit/sec rate is .233. This isn't a LOT of data, but it is significant enough to be a concern.

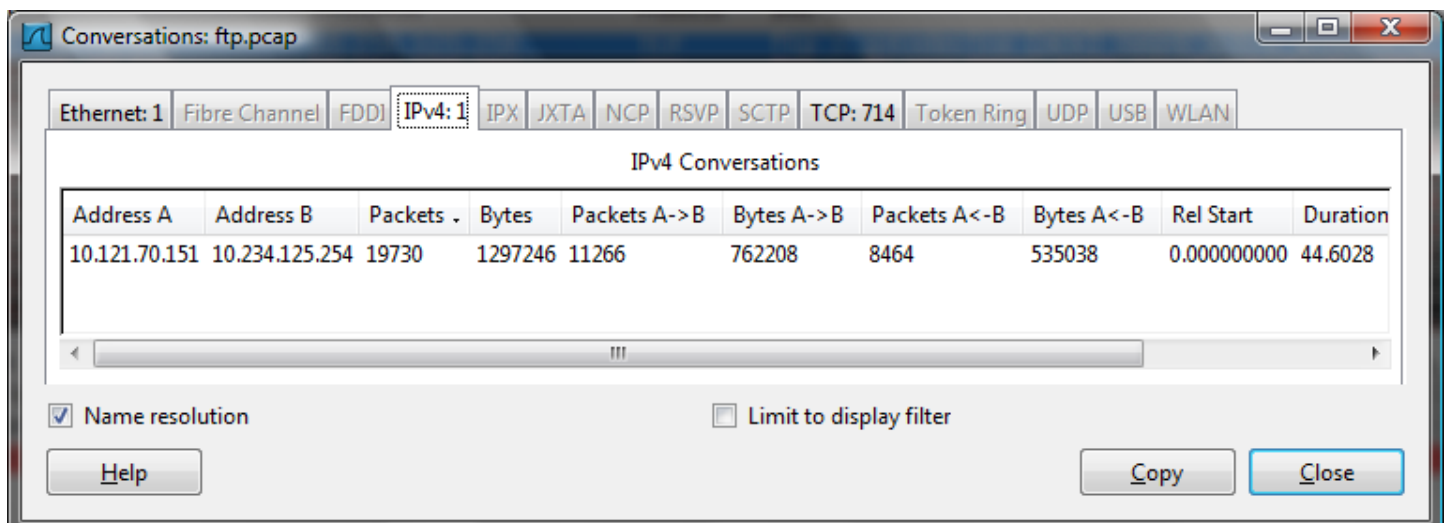
Looking back towards the packet list window, there is a lot of data to digest; around 20,000 packets in all. When looking at this many packets it's often a good idea to start by seeing how we can narrow down the field of examination. One good way to do this is to look at the conversations window. A conversation on a network, just like a conversation between two people, describes the communication that takes place between two hosts (also known as endpoints). You can access the conversations dialog by selecting Statistics at the top of the screen and choosing Conversations. Doing this will display a list of all conversations in a capture along with some general information about each conversation.



In this capture however, only one conversation is listed. This means that all 20,000 packets shown are being transmitted between the FTP server and the same host.

Unable to narrow down our search using the conversations window, our next step is to lev-

erage a filter to accomplish this task. We know that this is an FTP server, and looking at the packet list pane we have a mix of TCP and FTP traffic. A good strategy may be just isolating the FTP traffic so we can look at it. This can be done by creating a display filter.



A display filter is a filter that tells Wireshark to only display packets that match certain criteria. We can create one of these by typing the criteria we want to filter by into the Filter dialog docked directly above the packet list pane. If you type “ftp” without the quote into this box and hit enter, only the FTP traffic will be displayed in the packet list pane.

The packet details pane can be looked at to get an idea of what each FTP packet is doing. To do this, select a packet in the packet list pane and then expand the FTP section in the packet details pane. If you do this for the first FTP packet, you will see a response being sent from the server (10.121.70.151) to the client (10.234.125.254) stating that an attempted login had failed.

```

⊕ Frame 4 (76 bytes on wire, 76 bytes captured)
⊕ Ethernet II, Src: Cisco_3c:3f:a8 (00:01:96:3c:3f:a8), Dst: AmbitMic_aa:af:80 (00:d0:59:aa:af:80)
⊕ Internet Protocol, Src: 10.121.70.151 (10.121.70.151), Dst: 10.234.125.254 (10.234.125.254)
⊕ Transmission Control Protocol, Src Port: ftp (21), Dst Port: gotodevice (2217), Seq: 1, Ack: 1, Len: 22
⊖ File Transfer Protocol (FTP)
  ⊖ 530 Login incorrect.\r\n
    Response code: Not logged in (530)
    Response arg: Login incorrect.
  
```

If you continue along this course of action, looking at the details for each FTP packet, you will start seeing a trend. This entire capture consists of a remote user attempting and failing to log into this FTP server. Not only can you see these login attempts, but since the traffic is unencrypted, you can also see the FTP passwords being attempted. Looking at packets 11, 17, 21, and 47 you can see that the remote client has tried to login with the passwords merlin, mercury, mets, and mgr re-

spectively. Not only that, but if you look at the time column, these login attempts all happen within two tenths of a second.

Just to verify this a bit further, we can use a more advanced display filter to show all of the login attempts to this FTP server. Using the display filter “ftp.request.command == “PASS”” you will be shown all of the passwords that the remote client used to attempt to login to the server with.

No. .	Time	Source	Destination	Protocol	Info
11	0.032425	10.234.125.254	10.121.70.151	FTP	Request: PASS merlin
17	0.051363	10.234.125.254	10.121.70.151	FTP	Request: PASS mercury
21	0.059593	10.234.125.254	10.121.70.151	FTP	Request: PASS mets
47	0.195865	10.234.125.254	10.121.70.151	FTP	Request: PASS mgr
76	0.284369	10.234.125.254	10.121.70.151	FTP	Request: PASS mickey
80	0.333312	10.234.125.254	10.121.70.151	FTP	Request: PASS michael
91	0.419464	10.234.125.254	10.121.70.151	FTP	Request: PASS michelle
99	0.468459	10.234.125.254	10.121.70.151	FTP	Request: PASS michelle
116	0.506413	10.234.125.254	10.121.70.151	FTP	Request: PASS minimum
139	0.572819	10.234.125.254	10.121.70.151	FTP	Request: PASS mit
144	0.587005	10.234.125.254	10.121.70.151	FTP	Request: PASS minsky

Let’s see here, multiple rapid login attempts using alphabetically sequentially passwords?

This is a sure sign that someone is attempting to break into your FTP server by using a brute force attack.

We have successfully used Wireshark to not only track down what could be the cause of degraded FTP server performance during non peak hours, but we have also identified a potential intruder.

Conclusion

I often like to compare a network analyst working on a computer network to a doctor working on a human body. Regardless of whether you are seeing a cardiac, neurological, or orthopedic specialist, all of these doctors start with basic measurements of your overall well being. Where as a doctor might complete a blood culture, a network analyst will view a protocol hierarchy; where a doctor would complete a full medical history to get baseline of the patients overall health, a network

analyst will perform a few packet captures to get a baseline of the networks overall health. The idea here is that you have to know what makes something tick before you can focus in on a specific problem.

Visualizing a problem on a network isn't as easy as capturing a couple of packets and looking for the word "ERROR" in big bold print. You have to know what things look like when they are working properly to find the small subtleties that make the difference between a network in optimal health and one that creeps along at an alarming pace. The ONLY way to do this effectively is to be able to interpret the packets that are flowing across the wire.

This article is designed to give you a taste of what you can do with a packet sniffer and how essential of a skill packet analysis is. There are a few more resources I would recommend

if you want to learn more about packet sniffing and analysis:

- tinyurl.com/4va5gu - Practical Packet Analysis, published by No Starch Press last year is my full length book on using packet sniffing as an effective network troubleshooting tool. I'm a bit partial, but I think it's THE book to have if you are serious about learning packet analysis techniques.
- www.wireshark.org - The official Wireshark website contains downloads for Wireshark as well as community support links.
- www.openpacket.org - The mission of OpenPacket.org is to provide quality network traffic traces to researchers, analysts, and other members of the digital security community. This is a great learning resource.
- www.wiresharku.com - Wireshark University is an excellent self-paced DVD course on packet analysis, founded by Laura Chappell, Wireshark trainer extraordinaire.

Chris Sanders is a Senior Support Engineer for KeeFORCE, a technology consulting firm. Chris writes and speaks on various topics including packet analysis, network security, Microsoft technologies, and general network administration. His personal blog at www.chrissanders.org contains a great deal of articles and resources on all of these topics.



MailScanner www.mailscanner.info

The world's most widely-used e-mail security and anti-spam system that protects over 1 billion e-mails every day.

Over 1 million downloads!
Get your **FREE** copy today:
www.mailscanner.info






Latest additions to our bookshelf

Mac OS X Leopard Phrasebook

By Brian Tieman

Addison-Wesley Professional, ISBN: 0672329549



Mac OS X Leopard Phrasebook gives you the complete command phrases you need to take full advantage of the Leopard's hidden and undocumented power underneath the graphical user interface.

It contains time-saving solutions for effectively working with files, folders, the Finder, Spotlight, text files, servers, disks, CDs/DVDs, permissions, printing, applications, Expose, networking, security, and much more.

Router Security Strategies: Securing IP Network Traffic Planes

By Gregg Schudel, David J. Smith

Cisco Press, ISBN: 1587053365




Router Security Strategies: Securing IP Network Traffic Planes provides a comprehensive approach to understand and implement IP traffic plane separation and protection on IP routers.

This book details the distinct traffic planes of IP networks and the advanced techniques necessary to operationally secure them. This includes the data, control, management, and services planes that provide the infrastructure for IP networking.

Security Power Tools

By multiple authors

O'Reilly Media, ISBN: 0596009631




Designed to be browsed, Security Power Tools offers you multiple approaches to network security via 23 cross-referenced chapters that review the best security tools on the planet for both black hat techniques and white hat defense tactics. It's a valuable reference for network administrators, engineers and consultants with tips, tricks, and how-to advice for an assortment of freeware and commercial tools, ranging from intermediate level command-line operations to advanced programming of self-hiding exploits.

The Book of PF: A No-Nonsense Guide to the OpenBSD Firewall

By Peter Hansteen

No Starch Press, ISBN: 1593271654



The Book of PF is a current, no-nonsense guidebook to harnessing the power of PF. Its contents include coverage of network address translation, wireless networking, spam fighting, traffic shaping, failover provisioning, and logging.

Written for anyone who has felt lost in PF's manual pages or baffled by its massive feature set, author Peter Hansteen helps readers confidently build the high-performance, low maintenance network they need.

Network Warrior

By Gary A. Donahue

O'Reilly Media, ISBN: 0596101511



Network Warrior provides a thorough and practical introduction to the entire network infrastructure, from cabling to the routers. What you need to learn to pass a Cisco certification exam such as CCNA and what you need to know to survive in the real world are two very different things.

The strategies that this book offers weren't on the exam, but they 're exactly what you need to do your job well.

Geekonomics: The Real Cost of Insecure Software

By David Rice

Addison-Wesley Professional, ISBN: 0321477898




In this book author David Rice reveals why the software industry is rewarded for carelessness, and how we can revamp the industry's incentives to get the reliability and security we desperately need and deserve.

You'll discover why the software industry still has shockingly little accountability—and what we must do to fix that.

The Book of Wireless: A Painless Guide to Wi-Fi and Broadband Wireless

By John Ross

No Starch Press, ISBN: 1593271697



In the Book of Wireless, 2nd Edition, you'll learn how to set up your own home wireless network and how to use public wireless networks, safely and securely. This plain-English guide demystifies configuring and using wireless networks—everything from shopping for parts to securing your network.

You'll also learn about new and forthcoming broadband wireless standards and how to choose the right service provider and equipment. With up-to-date information on wireless routers, network interface cards, antennas, security and software.

Cross-Platform Development in C++: Building Mac OS X, Linux, and Windows Applications

By Syd Logan

Addison-Wesley Professional, ISBN: 032124642X



Cross-Platform Development in C++ is the definitive guide to developing portable C/C++ application code that will run natively on Windows, Macintosh, and Linux/Unix platforms without compromising functionality, usability, or quality.

This book will be an indispensable resource for every software professional and technical manager who is building new cross-platform software, porting existing C/C++ software, or planning software that may someday require cross-platform support.

Backup & Recovery

By W. Curtis Preston

O'Reilly Media, ISBN: 0596102461




Packed with practical, freely available backup and recovery solutions for Unix, Linux, Windows, and Mac OS X systems - as well as various databases - this new guide is a complete overhaul of Unix Backup & Recovery by the same author, now revised and expanded with over 75% new material.

It starts with a complete overview of backup philosophy and design, including the basic backup utilities of tar, dump, cpio, nbackup, ditto, and rsync. It then explains several open source backup products that automate backups using those utilities, including AMANDA, Bacula, BackupPC, rdiff-backup, etc.

READ BOOK AND SOFTWARE REVIEWS
www.net-security.org/reviews.php

ONE PRODUCT. FIVE DEFENDERS.
FIVE ANTI-VIRUS ENGINES. ONE CHOICE.



ENHANCE YOUR EMAIL DEFENSES
TODAY WITH GFI MAILSECURITY

GFI MailSecurity

Complete email security with up to five anti-virus engines for Exchange/SMTP/Lotus

No single anti-virus vendor scanner is the BEST and can stop ALL viruses. To obtain maximum security, you need GFI MailSecurity which uses not one, but up to five virus scanners to check all company email, with limited or no effect on network and server performance.

GFI MailSecurity is better priced than most single anti-virus engine solutions on the market. With multiple anti-virus engines you:

- React fastest to the latest virus threats by receiving the quickest virus signature updates
- Take advantage of all their strengths because no single anti-virus scanner is the BEST
- Virtually eliminate the chances of an infection.

Download your **FREE** trial version from www.gfi.com/ehns/



GFI NETWORK SECURITY
CONTENT SECURITY
MESSAGING



McAfee
NORMAN

bitdefender
secure your every bit

AVG Anti-Virus

The effectiveness of industry certifications

By Anthony Fama

This certifies that

Date _____

Over the course of time, each professional in the information technology field flips through countless magazines and comes across innumerable ads that start off with eye-catching phrases similar to the following:

"Become a ____, " or "Earn your ____ in four weeks!"

Invariably, the task of whether or not to pursue one (or more) certifications becomes challenging. This becomes a very personal decision that has no right or wrong answer. Additionally, the topic itself can become very contentious amongst industry professionals who work in circles where some individuals have certifications and some do not. The hope is that this article will help some of you with the anxiety inherent in making that very tough judgment call.

When one uses a cliché, it is often because she cannot creatively get a point across and must defer to a popular phrase most people are familiar with. Since the "proof is in the pudding" is a cliché most people understand, several integration firms in New York City were used as a baseline for study (to prove a

point). Essentially, in all cases, integration firms who attempted to "align" themselves with hardware manufacturers by sending technicians to training experienced the same results. In speaking with clients of these companies it was found that the most common feedback was that expertise had dropped over time (for whatever reason, especially when the technology being used was voice).

One firm in particular got great references from all clients several years ago. Since they aligned themselves with a hardware manufacturer by placing incompetent technicians in role that once had very competent people, took some simple steps and become certified and then, proceeded to perform such great feats as cut off key infrastructure for one of the biggest electronic trading firms in the

world. They continue to lose untold business to their former employees while their existing team of certified people lower the bar on their effectiveness. We will leave the names out to protect the unfortunate.

The time frame for this private study by an independent New York City firm started in 2001 and ended in 2007. The primary question leading to the unofficial study was whether or not to send an already elite team of engineers for training leading to certification. A direct correlation was found between time spent pursuing certification and "decrease" in effectiveness on site, in "read-life" applications.

Why would this truth unfold? One issue with engineers (or manufacturers) focused on "marketing" themselves is that they are directed to curriculum-style material that spoon-feeds the information necessary to pass tests

or to take labs. In contract, by the time a particular technology (or product) has generated enough interest to be included into a curriculum, and then added to test material, some time has gone by; in some cases, a significant amount of time.

As we all know, technology is time-sensitive. For example, knowing how to install outdated versions of Novell, Microsoft XP or Cisco 3550 switches (today) is not something engineers are scrambling to learn to do to give them an edge over their competition. Yet, these things are still required to pass some exams. Of course, these products are still out there and should be known by a field technician. In contrast however, larger firms that make a habit of installing the latest and greatest products want to know that the engineering resource responsible has worked on the technology already.

Manufacturers force certifications to be more prestigious than they really should be.

The more successful firms and engineers are proficient in these areas (in some cases) before the product is available, while their certification-bound counterparts continue to take practice exams! What's more is that the most successful engineers are the creative ones. Notice I say "engineer" here and not technician. The reason certifications may decrease effectiveness of an engineer and turn her into a technician is that it (in many cases) dictates how to approach a design, troubleshoot a problem, or recommend a product.

This is fine if the individual becoming familiar with the material is of moderate intelligence. Be that as it may, the most intelligent individuals going into the IT field are creative enough and resourceful enough to provide solutions outside the confines of any curriculum – essentially written by another individual in the field who is not in demand in the "street." It should be stated here that before the letters start pouring in, that the sources for this article are holders of the highest levels of certifications in several areas.

There is another impetus for the proliferation of certifications. Manufacturers invest tons of

money into promoting them. Why? Because if you become a CCIE, you will most likely recommend Cisco products. Manufacturers know this and do all they can to market these certifications as if they were college degrees.

Manufacturers force certifications to be more prestigious than they really should be. For example, an "engineer" is a person who has typically gone to an engineering "school." Every one of the industry certifications that has the word engineer in it should replace the "engineer" with technician.

In any case, that would seem less prestigious to those who pursue it (and then recommend those products). Luckily for Cisco, they used the word Expert. A Cisco Certified Internetworking Engineer is a very far cry from the sacrifices made when attending a four-year computer science or electrical engineering school. Finally, now that CCIEs believe they hold some pinnacle of technician classification, Cisco is releasing the CCDE, which according to sources, holds a higher level of prestige, isn't that throwing CCIEs under the bus. Evidently, the loyalty does not go in both directions.

Additionally, many IT professionals are unfamiliar with the business phrase "opportunity cost." As such, the up and coming starry-eyed engineer may not realize that yes, if she studies hard, she will pass tests, but there are other things that she could be doing besides studying. For example, visiting clients and promoting. It is a very tough call to quantify what the return will be on 'marketing' oneself. Economists spend a significant amount of time trying to figure this out. Yet, there are no advertisements in trade journals that read "Go out and tell everyone about you." Instead, they read, "Three weeks to a ___ certification." Thus, one gravitates towards certifications in lieu of simply telling the world what services can be rendered.

This 'long-winded' explanation is referred to as opportunity cost. In different terms, what one pays in order to pursue an opportunity. The payment may not be money; most often it is time, or the absence of pursuit of something else. This is a very tough concept to grasp, especially for an IT professional who does not have MBA experience. Individuals choosing between two potential mates go through this!

If I choose option A, what returns did I not realize had I chosen option B? That is opportunity cost. As it applies to IT, what else could one do instead of pursue a certification based on what one vendor believes I should know about a product.

We will now address this conundrum with an example: Alex wished to become a CCIE. It was his dream. The pursuit of a dream that requires self-sacrifice deserves praise and respect. Brian, was Alex' best friend. He wished to build his career by simply hitting the streets and seeing what potential clients had as objectives. In six weeks Brian had two clients who wanted nothing more than for Brian to orchestrate the installation of a local carrier's Internet circuit. In contrast, Alex was scheduling his first written exam. Within the next four weeks, the two clients Brian found paid Brian a total of \$2,000 for 20 hours of work to get the circuits in place and ensure that end-users were protected. Alex, had failed his first attempt at becoming a CCIE. Total cost, \$2,500 including study material. We'll leave out the remaining details of the story as you can see where this is going.

The moment a person reveals she is a CISSP, one immediately knows specified areas within which to plant (seed) information to 'steer' a person towards what we wish them to believe is the truth.

Please be advised that we are not saying certification is a bad thing because you would not want to be operated on by a doctor who was not board-certified! Certifications do serve a purpose: they convey to a potential client, who may not know your field, that you do know something about what you say you do and are not simply blowing smoke. Certifications can also get reseller better discounts on hardware to pass along to clients. Nonetheless, with the hardware market being affected by more services like eBay, this argument loses some strength. Another area that has been recently flooded with certification options has been security. It is in this area that industry certification is the least useful and we will now examine why.

Security, by definition, is safety. As it applies to information technology, safety refers to the protection of data from modification and/or

theft. Information, including knowledge of somebody's background can be used as a springboard into an attempt at finding out more information. The pursuit of which could lead to the discovery of a password and/or information that could facilitate a further breach of security. Why do we say this? Simple - the moment a person reveals she is a CISSP, one immediately knows specified areas within which to plant (seed) information to 'steer' a person towards what we wish them to believe is the truth.

A top-notch security expert never reveals her certifications and/or expertise and programming ability. The moment this rule is violated, the certification becomes less valuable. As such, the value of the CISSP or any other security-related certification drifts downward as we tell the whole world what it is all about. Unfortunately, in the area of security, this

makes the quest to be the best of the best a thankless pursuit. Marketing and security negatively affect each other in many ways. The number one tenet in security is that "everyone is on a need to know basis". If you don't need to know I am the best infiltration expert in the world, then you should not know it. Period. Thus, in security, the best of the best maintain small circles and do not advertise. Their clients know this and pay premiums since they know they are being serviced by professionals that are not going to the world marketing themselves (and therefore the procedures used to secure their clients – to some degree directly or indirectly).

There is one other aspect pertaining to industry certifications that we should mention. Recently, we had the opportunity to work indirectly with a technician (we'll simply call him Joe) who claimed he was a Master CNE. By conveying to his peers that he maintains an area of expertise indicates that he must know something about the topic. The fact of the matter was that some of the technicians working alongside Joe actually knew quite a bit more than Joe did but were not certified in anything. They were simply very knowledgeable.

To the non-technical managers who worked in the environment, it appeared somewhat illogical that a certified technician (in some cases) was following the lead of other individuals, even when it pertained to the area Joe was certified in. Thus, at this Japanese Bank, Joe is not held in high regard by the senior staff because he was expected to know more than the others yet fell short. Hence, by waving the certification flag in front of your peers and superiors, you immediately put yourself on defense to know certain things.

A good quote offered by Mayor Giuliani of New York City was "promise small and deliver big." Had Joe simply did his job and came up with the answer once in awhile, he may have surprised those around him and looked good. Instead, the need to tell the world he knew


something backfired on him. This centers on an insecurity issue better suited for an article on psychology. Still, we state it here simply to convey to the reader that once you tell the world you are certified, you are expected to lead and this adds stress to your career. Certifications leading to both money (potentially) and stress is something you need to consider before investing such a significant amount of time and money.

Certifications are an area that can help a newcomer to the IT industry. If a computer technician used to work at a local car wash, it helps a great deal to present an A+ certification when applying for a job. The morale of the story seems to be that the value of a certification is directly proportional to you and how many years you have in the industry. Thus, the more experience you accumulate, the less valuable your existing certifications are as a percentage of your overall worth. When viewed this way, given the examples we provide able, the decision starts to make more sense

If you have 10+ years as a switching and routing expert, a CCIE may not help you as much as it will help an engineer who is moving from C+ programming into networking with zero years of networking experience. As Einstein said, it's all relative; thus, a colleague who was performing security audits twenty years prior to the individuals who came up with the curriculum for the CISSP, mentioned "If I get a chance to get away from all my security clients, I will check out the curriculum and let them know what they missed."

Again, there are many arguments for and against certifications and the procurement of some certifications is admirable from the point of view of self-sacrifice. Ultimately, it may have a large self-esteem component for some people, quite honestly, no value can be placed on that. At the end of the day, you will make your own choices based on your own situation. As an old wise man once said - choose wisely.

Anthony Fama, Executive Vice President and CIO of Executive Technology Group, currently speaks to audiences and provides solutions for many high-profile local and global firms based in New York City. He has held numerous industry certifications, including some from Novell, Cisco, and Sun. He holds a Bachelor of Science Degree in Electrical Engineering and Physics and is an MBA as well. He also has a Masters Degree in Astrophysics and has been a member of Mensa for over twenty years.



Building a secure future: lessons learned from 2007's highest-profile security events

By Vijay Basani

Between the proliferation of worms, Trojans and botnets, and the outbreaks of cyber attacks, 2007 was a daunting year for even the most hardened IT security professional.

Clearly the most worrisome event, however, was the surge of targeted attacks and identity theft, which was best exemplified by the high-profile TJX Companies' security breach. While many details about the TJX breach have not been released, what is known is that hackers penetrated a wireless connection, gained access to a server holding sensitive data, installed rogue applications and stole over 90 million customer records from a central database. And, like many large-scale breaches, the operation was carried out over a period of several months.

Given that the financial gain possible with such identity and personal data theft is substantial, we should expect attacks of this sort to be one of the most troublesome areas for enterprise IT departments moving forward. So, what lessons can be learned from the attacks that occurred over the last year? And how can organizations more effectively mitigate these potentially catastrophic crimes moving forward?

First, recent high-profile security breaches suggest that companies assume their log-based security information management (SIM) solutions will help detect and identify all breaches. Unfortunately, this is often not the case and results in a false sense of security, especially when facing TJX-like "low and slow" targeted attacks that touch different parts of the network over an extended period of time.

Close examination of these breaches suggests companies should apply three basic guidelines for preventing catastrophic breaches, which can lead to lawsuits, financial loss and damage to brand reputation.

Guideline one: collect and analyze more than just log data

Analyzing log data is certainly a good place to start, but the problem is that most log management solutions focus solely on event logs, which do not contain all relevant security-related data.

For example, log files typically do not contain information about new accounts created, new users added, configurations changed, sensitive data accessed, new applications installed or new processes started. Nor does log data give you the ability to detect unauthorized access to application or system resources. After analyzing the breaches of 2007, it is clear that examining log data is not enough. Organizations also need to collect other key, more comprehensive information and analyze it for possible correlation with log data to more effectively predict threat patterns. Examples of such key information are vulnerability, configuration, asset, performance and network behavioral anomaly (NBA) data.

Guideline two: collect and correlate data over months, not days

Another lesson learned from recent breaches arises from the fact that targeted attacks often occur over weeks or even months. Many of the current SIM products that only analyze recent data are stymied by slow-evolving breaches. Short collection and analysis periods make it impossible to establish normal behavior, and thus, proactively detect anomalous behavior.

ous behavior. Data collection needs to span several months, not days, so that data correlation—across all data types and over time—can reveal targeted, slowly evolving identity attacks. Thus, to be effective, data collection and correlation must be “broad and deep”: broad in type and deep duration.

Guideline three: automatically correlate multi-source data

Now that it has been established that organizations must collect and correlate all security data—log, vulnerability, configuration, asset, performance and NBA—over several months, the third guideline pertains to data processing. Analysts working in a Security Operations Center (SOC) that uses multiple point solutions have to manually analyze data from numerous, separate IT silos—a time-consuming, inefficient and ineffective method that results in most important incidents going undetected. IT departments need to take a different approach to data analysis, with systems that enable automatic correlation and analysis of all relevant security data in real time across the entire enterprise.

Short collection and analysis periods make it impossible to establish normal behavior, and thus, proactively detect anomalous behavior.

A new approach to security, risk & audit management

In the combination of these three guidelines lies the essence of a new approach to security, risk and audit management. Revolutionary? No. Evolutionary? Yes, and it is similar to how network systems management frameworks, such as HP OpenView, IBM Tivoli, and CA Unicenter evolved to help large network operations teams reduce management complexity and improve operational efficiency. Now, in the world of security, enterprises must consider platforms that integrate security, risk and audit management capabilities, and provide a single unified data management and correlation engine for processing information from multiple data silos.

While this all sounds very futuristic, this solution category has already emerged and is being

leveraged by Global 2000 enterprises. Designed to complement existing point technologies, these integrated platforms provide broad, deep and automated data correlation and the ability to immediately detect suspicious activity, rapidly analyze root causes and proactively remediate problems.

Once SOC and networks operation center (NOC) teams identify a breach using an integrated platform like this, they can use built-in tools to collaborate with each other and quickly understand the “what”, “when”, “where”, “why” and “how” of any breach. With the consolidation of multiple data silos, users are provided with the complete context of any event in a single pane, thus avoiding the “swivel chair” approach. This reduces the time to detect, understand and mitigate a breach to seconds or minutes instead of days and months, which is often already too late.

In addition to security breaches, the demands of evolving compliance regulations continue to challenge organizations. These new integrated platforms also address the rigid IT requirements of regulatory mandates and best practice policies. Organizations can identify all necessary information around compliance violations, with continuous self-assessment that eliminates the worry over audit reviews mandated by regulations such as PCI DSS, SOX, FISMA, HIPAA, GLBA and more.

Many IT security experts predict that the record-setting security challenges of 2007 will only get worse in 2008 and beyond. However, there are few better teachers than past security breaches. With the emergence of integrated security, risk and audit management platforms, organizations will be better equipped to rapidly detect and respond to security incidents, while supporting regulations and best practice implementations—laying the foundation to help ensure they avoid becoming part of future targeted attacks.

While disparate log, vulnerability, configuration, asset, performance and network behavioral anomaly solutions can be used, the volumes of data collected creates individual silos.

The point solution approach

It is worth taking a moment to discuss a position that some would argue, and that is that the collection of all relevant security data can be accomplished with multiple point solutions.

While disparate log, vulnerability, configuration, asset, performance and network behavioral anomaly solutions can be used, the volumes of data collected creates individual silos.

This multiple-silo approach introduces a number of business challenges including the following:

- Cost: increases both deployment expenses because each point solution requires its own expert.

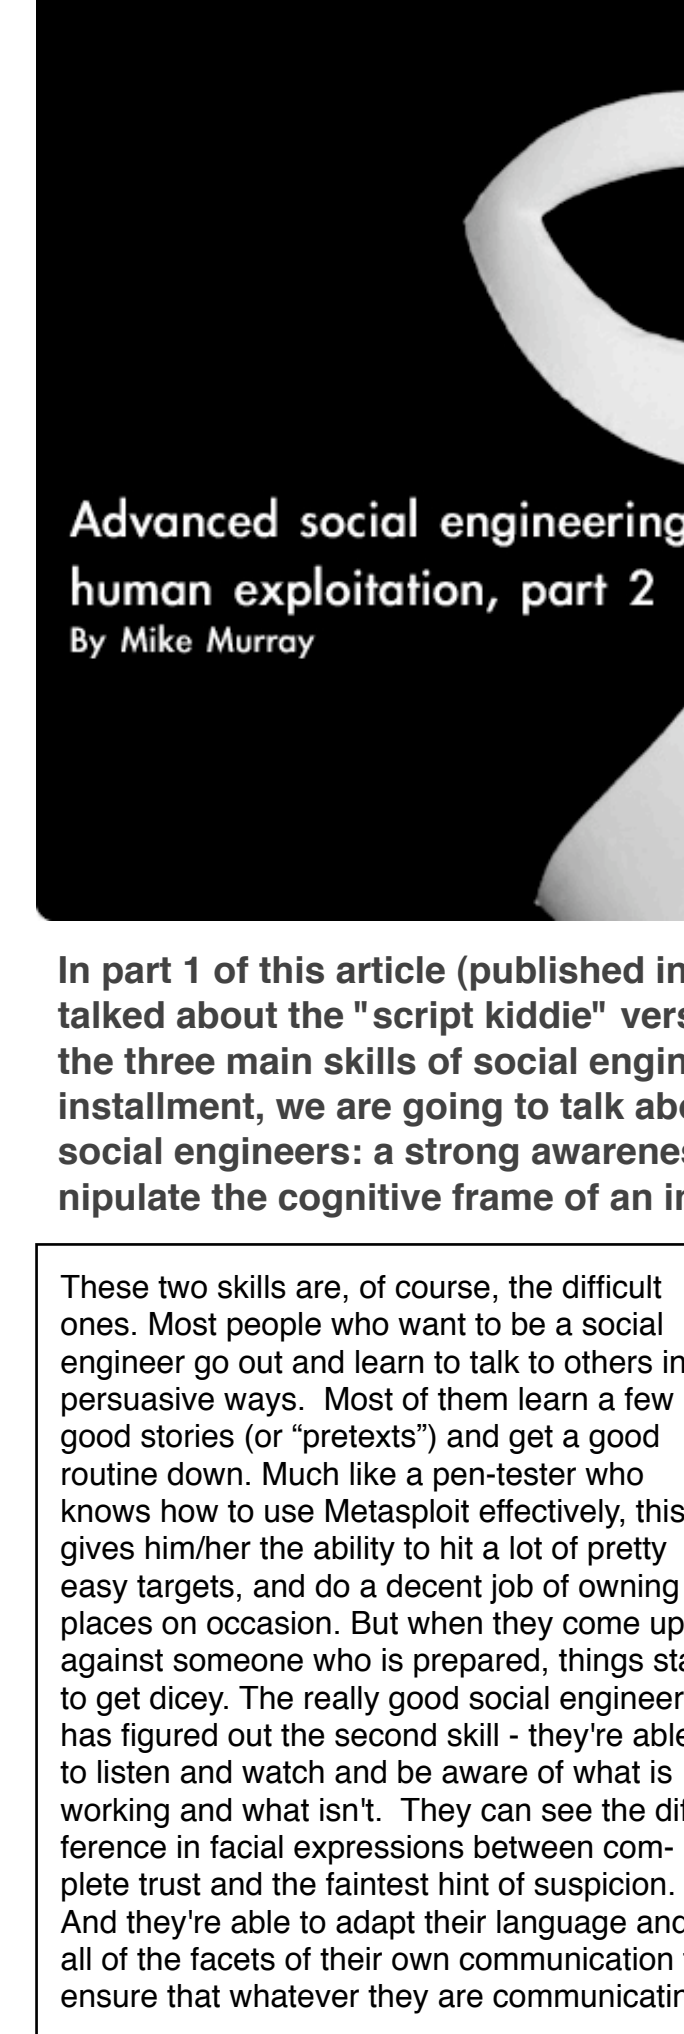
- Management: results in multiple data repositories that need to be continually correlated and managed.
- Analysis: requires the manual connection of dots to identify patterns and anomalies - this can add weeks, months or years to the time it takes to determine incident root cause.

The end result of such a cobbled together approach presents a tactical solution that increases management complexity and cost while failing to proactively and promptly identify security breaches.

Fortunately and as indicated by research from various industry analysts, IT departments are moving towards a more cost-effective, single solution approach that integrates all relevant security data to facilitate early breach detection and mitigation.

Vijay Basani is CEO & Co-Founder of eIQnetworks. Prior to starting eIQnetworks, he founded both ApplQ, Inc., an application storage resource management provider acquired by Hewlett Packard in October 2005, and WebManage Technologies, Inc., a policy driven content delivery solution provider acquired by Network Appliance, Inc. in August 2000. His experience also includes numerous senior executive positions in the financial industry at Spencer Trask Securities and Josephthal Lyon & Ross, Inc. Vijay, the co-owner of five patents for the architecture and design of the WebManage Content Delivery system, Adaptive Policy Engine and SLA Management, holds a Bachelor of Engineering in electronics and instrumentation as well as MBA and Post MBA degrees from Baruch College in New York.

www.net-security.org
Get up-to-date security information now.



Advanced social engineering and human exploitation, part 2

By Mike Murray

In part 1 of this article (published in issue 15 of (IN)SECURE Magazine), we talked about the "script kiddie" version of social engineering, and the first of the three main skills of social engineering – the artful use of language. In this installment, we are going to talk about the other two skills of really advanced social engineers: a strong awareness of other people, and the ability to manipulate the cognitive frame of an interaction.

These two skills are, of course, the difficult ones. Most people who want to be a social engineer go out and learn to talk to others in persuasive ways. Most of them learn a few good stories (or "pretexts") and get a good routine down. Much like a pen-tester who knows how to use Metasploit effectively, this gives him/her the ability to hit a lot of pretty easy targets, and do a decent job of owning places on occasion. But when they come up against someone who is prepared, things start to get dicey. The really good social engineer has figured out the second skill - they're able to listen and watch and be aware of what is working and what isn't. They can see the difference in facial expressions between complete trust and the faintest hint of suspicion. And they're able to adapt their language and all of the facets of their own communication to ensure that whatever they are communicating

keeps working. It's that awareness of others that we're going to start off with this time.

Awareness – your internal compass

Your awareness of others is fundamentally like having a compass. If your goal is to go to a destination that lies due north, it is incredibly helpful to know that you happen, at this moment, to be walking east. The compass will tell you that a direction change is required. That is exactly what your ability to read others will enable you to do. It gives you the ability to change, adapt and restructure communications on the fly relies entirely upon the ability to see the effects that your communications are having. The awareness of others within the communication is the true skill that conveys the ability to know whether your language is having the intended effect.

This compass is the basis of all good communicators, and can be best seen in the actions of the average six-year old.

"Can I have ice cream?"

"No", the parent says.

"Please?"

"No", the parent says with a little bit of irritation.

"I want it!!!", the child yells with a bit of foot stomping.

We've all seen this interaction and know where it ends up most of the time - eventually, the child either gets the ice cream, or gives up. While perseverance plays an important part in the conversation, the child's ability to accurately understand what their parent is going to do is of the utmost importance to whether they end up with chocolate dripping

down their chin or end up banished to their room.

In NLP, this skill is referred to as calibration which is the ability to notice the emotional, physical and mental state of another and to accurately apply that state to oneself. If the parent in the above example happens to be in a loving mood, the choice of temper tantrum may not be the most effective. Whereas, if the parent is already stressed, it may be. The point is not what the calibration is - the skill is to notice facial expressions and body language at an extremely precise level, and accurately represent them within oneself in order to structure communication more effectively.

If this sounds hard, it's because it should be. But, luckily for us, our brains are designed to overcome the challenge.

IF THIS SOUNDS HARD, IT'S BECAUSE IT SHOULD BE.

Mirror neurons - our brain's inner reflector

In the mid-90s, a bunch of Italian scientists were studying a particular type of monkey and attempting to map the monkey's brain. They would have the monkey pick up a banana and take note of which parts of its brain were active as it did so - they did this over and over again so that they could figure out the brain regions that were involved in that particular task.

They decided to take a break from the experiment, but left one of the monkeys wired up to the machine. One of the experimenters decided to have a banana himself, and reached over to pick up one. At that exact moment, the monkey's brain lit up in all of the same regions that it had when the monkey picked up a banana itself.

They could have dismissed that as a fluke, but, luckily for us, they didn't. That study led to a whole pile of other studies that found that our brain (i.e. that of all primates) is actually wired to respond in reaction to other people. If I'm looking at you, every thing that you do, from picking up a banana to smiling to laughing is represented in my brain at the same

time that you're doing it. Our brain has structures that science has deemed as "mirror neurons" that are designed to represent the experience of everyone around us in our brains at the same time.

And, what's really cool is that you can improve that response through work on calibration. We can learn to sensitize our mirror neurons to have more sensitivity to those around us. For example, one of the "calibration games" that is often played by those practicing NLP skills is known as "red rose / white rose". It involves sitting across from a partner and having them think of a red rose, and then, after a quick break, a white rose. Then, the partner is to pick either a red rose or a white rose to think of, and it's the goal of the practitioner to guess what they're thinking of.

The game works because there are subtle differences in muscle tone, facial expression, breathing, etc. as the person thinks of two different things. Because of those differences, someone who has practiced their ability to be sensitive to such differences is able to perform far better than random chance on this type of game.

Noticing others - body language and the human face

Calibration of others can be broken down into two main skill sets - the ability to observe body language and facial expression. While there are hundreds of books on body language, the majority of them actually lead to wrong conclusions for most people who want to be effective in a social engineering setting.

The goal of most books on body language is to attempt to apply meaning to some individual piece of body language in order to give a "reading" of what a person is thinking or feeling. A frequently used example is the belief that if you stand with your arms crossed it indicates that you are emotionally closed. This is one of the most common theories of "body language", and many people expound it as though it is a universal truth of human behavior.

Unfortunately, it isn't ever that simple. While some patterns exist in human behavior, the types of absolute rules that most proponents of "body language analysis" give us are simply not very good with people. There are very, very few rules that are true of every person, all the time, in every situation. The problem with this type of "meaning-based" analysis of body language is that alternate meanings are

equally possible given certain contexts and certain other cues in almost every situation. For example, in the situation above, the person could simply be cold. Or they could have a stain on their shirt. Or any number of other things that have nothing to do with the interpretation of their body language. This same problem exists with most of the work on "facial expressions" that exist. NLP is actually the worst of all worlds for this - anyone who has studied even the most rudimentary NLP-type language is likely familiar with what are commonly known as "eye accessing cues" with its familiar picture (an example of which is at tinyurl.com/3tfasb).

Unfortunately, this theory has the same limitations as most of the "body language" theories - none of the rules work for all people all of the time. What is an aspiring social engineer to do? Simply put, practice calibration. Work on becoming more and more precise in the ability to notice and represent the smallest facial expressions and physiology changes within their own neural network. This is not as complex a task as it sounds - a large number of exercises exist that are designed to help develop calibration skills. For a place to start, I encourage looking at the work of Paul Ekman and his METT/SETT instructional tools (tinyurl.com/3hsc22l).

WHILE THERE ARE HUNDREDS OF BOOKS ON BODY LANGUAGE, THE MAJORITY OF THEM ACTUALLY LEAD TO WRONG CONCLUSIONS FOR MOST PEOPLE WHO WANT TO BE EFFECTIVE IN A SOCIAL ENGINEERING SETTING.

Rapport - magical influence

The purpose of obtaining a really deep and nuanced representation of those we are influencing is ultimately to more effectively enable us to influence those people. Wouldn't it be amazing if there was a state we could enter with another person that would make influence even easier? Luckily for us, there is: it is known in most reading on the subject as "rapport".

Rapport is described by Wikipedia as: "It is commonality of perspective, being in "sync", being on the same "wavelength" as the person with whom you are talking."

(en.wikipedia.org/wiki/Rapport). There have been hundreds of books from sales to psychology that offer theories of how to enter rapport, and almost all of them have focused on doing rapport backwards. Most experts on rapport will tell you that the best way to enter rapport with someone is to "match and mirror" their posture and physiology as you are communicating with them - breathe when they breathe, blink when they blink, and sit or stand how they're sitting or standing.

While it is true that people in rapport actually do end up matching and mirroring each other, this is backwards. NLP trainer Tom Vizzini says it best when he says that the way

that most people approach obtaining rapport is like "trying to stuff exhaust gasses up the tail pipe of a car and expecting the engine to start". The real skill of rapport comes back to calibration - the goal is to replicate the other person's state in oneself. If you're talking to someone who is happy, being like them is not likely to create the feeling of "being in sync" in that person. The goal is to, rather than matching particular parts of their physiology, to learn to match the state that causes the physiology.

Many of us have had that experience of rapport where we were so in sync with a person that we could feel what the next sentence out of their mouth would be (this happens most often when we are newly in love, but that's another topic entirely). Regardless of whether,

in that moment, you were matching posture and blinking rates, you were matching the person's state in that situation. This is how effective rapport happens in the real world, and it is how we can train ourselves to make it happen in our social engineering engagements.

That ability to match state will create endless situations of rapport in our social engineering, and is the real defining characteristic of a strong social engineer. However, the best social engineers take it to another level. That level involves an understanding not only of the interaction between you and the person (or people) you are working with, but of the context or "frame" around the situation.

WHAT STRONG SOCIAL ENGINEERS ARE ABLE TO DO IS MANIPULATE THE PERCEPTIONS OF THOSE AROUND THEM IN ORDER TO ENSURE THAT THE CONTENT THAT THEY ARE USING IS INTERPRETED IN THE WAY THAT MAKES THEM MOST ABLE TO GET WHAT THEY WANT.

Cognitive frames

Frames are defined by Wikipedia as: "the inevitable process of selective influence over the individual's perception of the meanings attributed to words or phrases. Framing defines the packaging of an element of rhetoric in such a way as to encourage certain interpretations and to discourage others".

In effect, the frame of a conversation is the context that allows us to shape the meaning that is given to the content. Most of the time, we spend our lives absorbed in content. When having a conversation, we are mostly engrossed in what the person is saying and what we are saying in response. However, the frame around the communication has a great deal more impact on its effectiveness than does the content of the communication itself. What strong social engineers are able to do is manipulate the perceptions of those around them in order to ensure that the content that they are using is interpreted in the way that makes them most able to get what they want. Let's look at an example. Imagine the sentence: "Give me your password." Now, let's try placing a few different frames around that.

- Someone on the street saying it just after you have introduced yourself.
- Your boss saying it after he has told you that you're terminated.
- A police officer saying it in an interrogation room while holding your laptop.
- Your lover saying it sweetly while you're laying in bed.

Each of these contexts radically changes the meaning of the words in the sentence. Though the context of a social engineering engagement will not shift nearly as wildly, a truly advanced social engineer is able to shift, move and alter the frame of a conversation subtly to evoke the conditions that are most likely to allow him or her to succeed in a given engagement.

Frame control – the art of manipulating context

When thinking about changing frames, it makes the most sense to think of a frame as a physical thing - imagine a real picture frame. If we have a picture frame, there are a few things that we can do in order to manipulate the way that the frame holds the picture.

- Transformation: Changing the shape of the frame so that it focuses differently.
- Extension/Contraction: Making the frame larger or smaller to incorporate more or less of the situation.
- Combination: Joining two frames or bringing a second frame into play within the first.
- Amplification/Compression: Changing the intensity of the content within the frame.

While this may seem confusing, a common example of extension and contraction should suffice to help you understand. Imagine sitting in a movie theatre in the middle of a movie. The entire context of a movie theatre is designed to create a situation where the picture on the screen is framed by darkness - in fact, if you're like most people, when sitting in a movie theatre, the picture on the screen is the limit of your visual perception. Your context for most of the movie is limited to the boundaries of the screen - this is what allows you to suspend disbelief and get engrossed in the story. Imagine, though, extending the frame and watching the movie while simultaneously maintaining an awareness of all of the people in front of you in the theatre as well. Do you notice (even though this is simply a thought experiment) how that changes your perception of the movie on the screen? This happens in human interaction as well. We have all been in a conversation where the entire room appears to disappear - where we're so intent on the conversation that is occurring that we stop noticing anything outside of the conversation. Conversely, a time when someone was talking to us and we couldn't concentrate because we were unable to limit the frame to just what was being said.

Common frame elements

There are some well known elements that a great social engineer can bring in to a frame (through choice of words, actions or props) that can assist the social engineer in setting the appropriate frame. As has been discovered through years of marketing and sales, creating frames with certain meanings can make people more likely to be influenced. Some of those elements are:

- Reciprocation: we want to do things for those who do things for us. As well, we want to do things for those we have done things for in the past.
- Commitment/Consistency: if we take a position publicly, we are much more likely to act in a way that makes us consistent with the position in the future.
- Social Proof: if other people do something, it makes us likely to do it as well.
- Authority: we are likely to follow along with authority.
- Scarcity: if we believe that a resource is scarce, we are likely to want it.
- Amygdala Hijack: the use of emotional response to lessen the critical response.
- Confusion: in confusion, our mind searches for an escape to confusion with less critical awareness of the options.

Note that these frame elements are not the same as the "pretexts" that are often talked about in describing social engineering, though most of the common pretexts that are chosen by social engineers incorporate one or more of these frame elements naturally. However, almost any context can be manipulated by a skilled communicator to integrate some of these elements.

Seeing the threat

As I'm sure you can see, the truly advanced social engineer knows a great deal more than how to "just ask" for the information that he/she wants. A real social engineer is part artist and part magician, and has the skills to make subtle communication into a weapon in ways that few in our industry have yet to become aware of. There is no technology that will protect enterprises from this type of exploitation - only the ability to prepare themselves through their control environments and their people. Unfortunately for our industry, we have yet to really understand the severity of the threat. The attacker who has these skills is like someone who understood IDS evasion in 1997 - they simply aren't getting caught. They're in and out like Derren Brown paying in a store with blank paper - they're gone before anybody ever realizes what happened.

Mike Murray is an experienced social engineer, trained hypnotherapist, and long-time information security professional. He currently is the Director of Neohapsis product testing lab, and is the author of the upcoming book "Social Engineering: Advanced Human Exploitation". Read his blog at www.episteme.ca.

Events around the world



LayerOne 2008

17 May-18 May 2008 - Pasadena Hilton, Pasadena, CA
www.layerone.info

OWASP AppSec Europe 2008

20 May-23 May 2008 - Ghent, Belgium
www.owasp.org

EUSecWest 2008

21 May-22 May 2008 - Victoria Park Plaza, London, UK
www.eusecwest.com/

Hacker Halted USA 2008

28 May-4 June 2008 - Myrtle Beach, SC, USA
www.eccouncil.org/hhusa

Shakacon 2008

9 June-13 June 2008 - Dole Cannery Ballroom, Honolulu
www.shakacon.org

Recon 2008

13 June-15 June 2008 - DoubleTree Plaza Montreal, Canada
www.recon.cx/2008

SyScan 2008

3 July-4 July 2008 - Novotel Clarke Quay, Singapore
www.syscan.org

Black Hat **USA** 2008

**“DEFENSE IS THE STRONGER FORM
OF WAGING WAR”**

-Karl von Clausewitz

The war for your data rages on.
Be certain your defenses are up to the job.

Black Hat USA convenes the best infosec minds on the planet for six days of intense, hands-on security education and peer-to-peer networking. Our speakers and trainers are the world's leading voices from academia, research and the underground. The breadth and depth of topics is unmatched. You will gain actionable knowledge, discover new tools, and learn expert techniques for digital self defense.

12 tracks 80 presentations 40 training sessions

August 2-7 2008
Caesars Palace



Las Vegas
Nevada, USA

Diamond Sponsor

Microsoft

Platinum Sponsors



Gold Sponsors



Google

Silver Sponsors



Interview with Nitesh Dhanjani, Senior Manager at Ernst & Young

By Mirko Zorz



Nitesh Dhanjani is a well known security researcher, author, and speaker. Dhanjani is currently Senior Manager at Ernst & Young LLP where he leads their Application Security efforts. Dhanjani is responsible for evangelizing new application security service lines, ensuring current service lines stay bleeding edge, and helping enterprises develop world-class application security strategies. Dhanjani is the author of several books and has been invited to talk at various information security events such as the Black Hat Briefings, RSA, Hack in the Box, and OSCON.

Enterprises need to formulate high-level goals for application security efforts before implementing specific service lines. What are the key areas they have to cover in order to make their endeavor successful?

I agree. You've got to strategize high-level goals before deciding on specifics. Most businesses are not in the business of being secure. They are in the business of generating revenue, protecting their brand, and their intellectual property. Application security goals must derive from and support these business goals to promise risk reduction across the enterprise cheaply and effectively. Such promises in turn require specific implementations

and processes such as hiring the right talent, laying the right framework, hooking security into the development lifecycles, training, metrics, and executive support.

Despite owning a plethora of software and hardware solutions, the critical asset to an organization is still the security professional who works with those acquisitions. How exactly important is the security team?

Talent is key. What good is an application scanner or code analyzer if you don't have professionals in your team who actually understand the results? The fastest way to lose credibility with the business is to employ

individuals who cannot go beyond running assessment tools and exporting reports. The job of a security team in any organization is not to hire people who can point and click their way into running assessment tools, but to establish a world-class effort that serves the needs of the business. You do that by hiring subject matter experts. You do that by hiring talent that can impress the business and demonstrate tangible value and progress.

With the threat landscape constantly evolving and old issues still not resolved, the organization has to battle problems such as a lack of security awareness that bring in a myriad of complications. What is the right approach to take in order to battle difficulties one can't completely protect against?

That's a two-part question: how to deal with known issues, and how to keep up with the latest attack vectors. First, you've got to establish a process that aims to remove security gaps at the root. Training and awareness offers the best ROI in this regard: bugs that don't get created in the first place - imagine that! It is also vital to embed security into the development lifecycles of applications. However many organizations have trouble deciding where to begin.

The solution is to assign efforts based on risk. Start by understanding what applications you

own and what their business impact is. What type of data do these applications read and write? What is the business criticality of these applications? Once you have a good understanding of your application portfolio, it will be much easier to assign effort so you can focus clearly.

As for the second part of the question, the solution is to invest effort into research and development so you continue to understand how the latest attack vectors may target your software. Yet again, training and awareness wins in this regard. Set aside a budget to send your team to information security conferences and training programs so they can soak up new knowledge. Allow analysts to take some time out to investigate the latest attack techniques. Most hands on security professionals are scientists at heart - understand what makes them thrive and support their talent. Support their desire to learn new ways to break security controls.

Finally, capture and communicate this knowledge to the business. For example, ensure your threat modeling attack libraries are up to date and reflect the latest attack vectors, that your code review and assessment methodologies are bleeding edge, and that you take time out to brief the architects and developers on what they need to know to keep up.

Allow analysts to take some time out to investigate the latest attack techniques.

Although applications have the largest attack vector today, CSOs don't take this into account when strategizing security spending. What kind of issues can this bring in the near future?

I overwhelmingly agree - the security spend of many organizations is out of whack with the real threat landscape. I feel there are multiple reasons behind this situation, the most common reason being, in my personal opinion, that many individuals who have been hands on in the past remember and hold on to the notion of the network layer being the only big thing to worry about. I can sympathize with that view - if you rewind a couple of years,

majority of the high impact attacks could be identified and blocked via network controls because the attack surface of applications was low compared to today's scenario where a typical enterprise level web application is comprised of millions of lines of custom code.

Perhaps another reasoning for this situation is that solutions around application security do not provide the instant gratification of throwing in a few appliances to solve problems. Well, perhaps, I should take that back, there are a few web application firewall solutions in the form of appliances that are starting to be marketed that way, but that's the nature of marketing and I'll save my rant for another forum.

Also, quite unfortunately, there are going to be more situations in the coming future where many security efforts will align with reality after learning the hard way - i.e. after an application related exposure has already taken place. That said, the onus is still on the security pro-

fessionals and researchers - we need to do an even better job of demonstrating impact and educating decision makers on why a solid application security strategy is vital to any organization's overall security effort.

Black-box assessments, coupled with the right strategy, do have their place.

While having consultants come in and perform black box penetration audits of applications every year is more costly than investing in a solid SDLC process, many organizations still believe it to be the proper strategy. What should they take into consideration when making this decision?

Black-box penetration tests are useful, yet they are extremely expensive and ineffective when relied upon as an exclusive solution. Paraphrasing a colleague of mine, "Companies that solely on black-box assessments to guide their security efforts do something similar to having consultants come in, throw a grenade at them, and have the consultants close the door on their way out". Gary McGraw likens this situation to what he calls a "Badnessometer". Black-box assessments correlate to symptoms that reflect the level of trouble you may be in. The response shouldn't be to just fix the black-box assessment results, but to respond to the situation strategically, and ensure you are responding to and eliminating the root cause of your problems to make sure they do not re-occur.

The solution is to "push left". The ingredients of a typical application development cycle, from left to right, includes the requirements phase, followed by design, then implementation, test, and production. The more effort you put into implementing the right security controls at an earlier tollgate, the lesser it will cost you. For example, assume that a review of security controls during the design phase results in an architect having to re-engineer the authentication mechanism. Now, imagine if this issue was not caught during the design phase, but uncovered during an attack & penetration assessment after the application is in production. It is not trivial to re-engineer a product that is already in production. And it is extremely costly - multiple times costlier.

Black-box assessments, coupled with the right strategy, do have their place. Going back to Gary McGraw's point, they help uncover symptoms. These assessments can be used to further augment and enhance an existing security SDLC process. For example, if you are finding too many issues via black-box assessments during pre-production that you missed to uncover during design and test, then it is time to re-evaluate your SDLC process and approach.

Besides having an excellent technical background, the CSO has to be good at demonstrating a tangible impact of his actions to the management in order to justify security spending. This ability is becoming increasingly important and can take the focus off the main areas of security he should be working on. How can the CSO lighten this load?

Justifying security spend for application security is not difficult. In addition, application security efforts can have positive political side effects for the CSO and the security team. I'll tackle these two points separately. First, application security must tie into the overall IT risk strategy. Start with asking what the company's business goals are, and how you want to demonstrate value. Map these goals to efforts based on risk that will flow into specific tactics that can demonstrate ROI. For example, if your organization currently relies on yearly black-box assessments, calculate the cost of performing the assessments in addition to the cost of remediation. Compare the cost with progress you've made by evaluating the last two assessment results for the same application. Most likely, you will find that you have made little tangible progress in the form of risk reduction and that the remediation cost has been high. Now calculate the cost of investing in a "push left" scenario.

Put these two scenarios side by side, and you'll have a solid case for ROI. The returns for a good application strategy are tremendous. The important thing is to continue to measure returns by formulating a good metrics program. Keep track of how security is helping business and technology improve, measure the drop of defects per lines of code, measure the amount of risk reduction. Show value. Demonstrate how your program is embedding security into the organization's DNA.

To my second point, a well thought out application security strategy can help the CSO politically. If you hire the right talent, and approach business and technology with the right attitude, i.e. to enable and not disable, you'll make a lot of new friends. Security departments often complain that the revenue generating business units view them with disapproval, yet the quest for application security is a fantastic opportunity for the security team to work closely with business. Do it right, and the business will love you for it. You will win their credibility and gratitude.

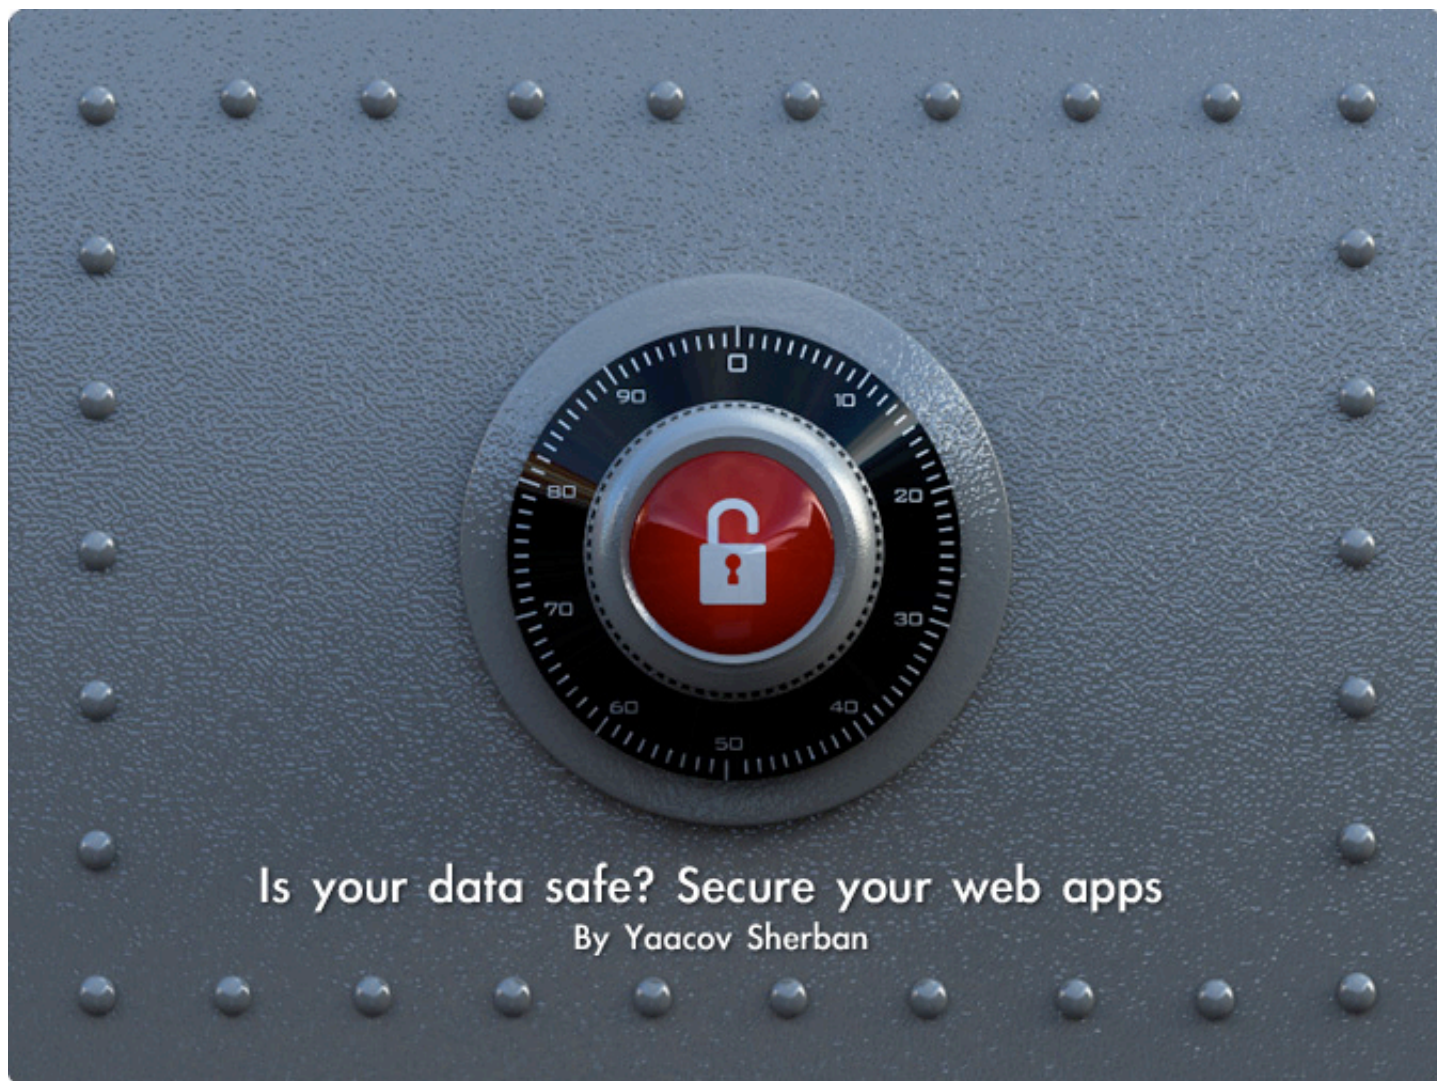
How important is threat modeling?

If you want to do application security right, you've got to invest in threat modeling. The goal of a threat model is to enumerate the malicious entity's goals even if the threats being enumerated have been mitigated. This helps the business, developers, architects, and security analysts understand the real world threats to their applications. Threat modeling should be initiated during the design phase of the application and it should be treated as a living document. As the application development process progresses, the threat model can be further enhanced so it is increasingly valuable. For example, a threat model created during the design phase can be further augmented to map to actual code review results to help developers and architects understand areas they need to improve on and areas where they are doing a good job. Threat modeling is a core component of the push-left strategy, so you eliminate defects as early as possible.

What recommendations would you give to a new organization that is just starting to build an application security strategy?

Once you've derived from your overall business and security goals, it's time to list specifics.

- 1. Talent and Framework.** Hire the right talent and lay the framework: policies, requirements, best practices, and methodologies.
- 2. Kick start efforts on critical applications.** Kick start your efforts on your critical applications: work with business and technology to help them understand the risks to their applications and what they can do to eliminate them as early as possible. Help them invest their by offering advise on their architecture level security controls and threat modeling. Give the development teams guidance on secure coding policies. Assess the code for security defects, followed by a penetration test before the application is turned over to production. Ensure proper application logging mechanisms are built in and monitored.
- 3. Application portfolio.** Come up with a formula to calculate business impact of an application based on key questions. Rank the applications by impact, and assign effort. At this point, you may want to take regulatory requirements into account, most often based on the type of data handled.
- 4. Invest in training and awareness.** The security team, business, and technology must have access to continuous security training. Calculate metrics from code review results to target security training to certain business units. After a code review assessment, get a few of the developers into a room and show them the impact of the vulnerabilities found. Work together to enhance the threat model and possibly fixing the defects. The goal is training, awareness, and knowledge transfer.
- 5. Metrics.** Demonstrate value, for example, a graph showing defects per lines of code decrease within the span of the last few months. Demonstrate risk reduction per business unit which often leads to some healthy competition, and that's a good thing. Overall, you've got to show application risk reduction across the enterprise.
- 6. Stay cutting edge. Retain talent.** Treat your team well. Understand the latest attack vectors. Invest in research. Communicate and support the business - they are your clients and they need your help.



Application security is the greatest online threat of the moment, and is at the top of the influential SANS Institute's Top 20 Vulnerabilities list. It is widely held by experts that a huge percentage of corporate website compromises are due to attacks on the application layer, but many companies are yet to take any action. Many are not aware of the problem at all, and those that are may not know where to start.

Additionally, businesses that handle credit cards have until June 30, 2008 to comply with the Payment Card Industry Data Security Standard (PCI DSS) which requires specific action to be taken on application security. The standard was designed by credit card issuers including American Express, Discover Financial Services, JCB, MasterCard Worldwide and Visa International to curb the number of data breaches occurring.

Failure to comply with PCI can have serious consequences in the event of a data breach, with fines of hundreds of thousands of dollars and trading restrictions among them. Fines are thought to be in the region of \$8 per compromised account, a breach fee in excess of \$158,000 per incident, as well as possible re-

strictions on the merchant and permanent prohibition of the merchant's credit card facility.

Many of the PCI core requirements are related to basic good practices in IT security. Aspects such as deploying a perimeter firewall and not using default passwords are certainly mandatory for any online business, while protecting stored data and using HTTPS/SSL encryption as a minimum is also everyday practice. Additionally, maintaining up-to-date anti-virus protection and restricting physical access to customer data is all part of a basic information security policy.

However, the PCI standard really gets tough around application security, and this is where

many companies slip up. Applications must be patched within one month of a new vendor release, and new vulnerabilities must be tracked by a specific internal process. Any custom applications must be built according to best practices, and secure coding guidelines must be used. Common coding vulnerabilities such as cross-site scripting (XSS) attacks, SQL injections, buffer overflows, improper error handling and denial of service must be prevented.

The June 30 deadline relates to the additional PCI requirement that all web-facing applications must be protected against known attacks by regular code reviews, or a Web Application Firewall (WAF).

Having code reviewed annually by specialists may at first sight be an attractive one, as it might be initially cheaper and the result will be customized to your application configuration. However, a full code review takes time, and a specialist application security company will require full and unfettered access to your servers. The duration will vary according to the size of your company, but anything from a few days to weeks should be expected. A full report of known vulnerabilities will be provided at the end of the review, and these must be corrected before compliance can be claimed.

The alternative, a hardware or software WAF is designed to inspect the contents of the application layer of an IP packet - in other words any data that is processed by an application. Many solutions claim to contain WAF technology, some incorporating other functions such as packet filtering, proxying, SSL termination and load balancing, but not all will ensure compliance with PCI DSS 6.6.

A genuine web application firewall will inspect all application input and respond based on rules or policy, prevent data leakage by inspecting output based on rules or policy, and support both positive (white-list) and negative (blacklist) security models. It will also inspect both web page content, such as HTML, DHTML and CSS as well as the underlying

protocols that deliver content, such as Hyper Text Transport Protocol (HTTP) and Hyper Text Transport Protocol over SSL (HTTPS). A WAF should also be able to inspect web services messages such as Simple Object Access Protocol (SOAP) and eXtensible Markup Language (XML), both document- and RPC-oriented models, in addition to HTTP, and defend against threats that target the WAF itself.

There are both hardware and software WAFs on the market, with associated pros and cons. As a hardware appliance a WAF will have a high initial cost of ownership, and ongoing specialist maintenance costs. Hardware WAFs use learning or statistical algorithms to determine if a specific input is malicious, while software WAFs will generally use pre-determined rules to make the same decision. The learning algorithms can consume large amounts of processing power, CPU and memory resources, and also require dedicated security expertise to deal with any false positives. Software logs on the other hand can be monitored on a more periodic basis.

Hardware application WAFs are more suited for static applications. Software WAFs work well with dynamic applications and are particularly appealing for smaller businesses that are unwilling to invest the tens of thousands of dollars on an appliance, and commit to the regular maintenance required. This is especially relevant today as a wider range of companies are required to become PCI DSS compliant. Additionally, software WAFs will not only be pre-configured, but should receive automatic updates that ensure you are not only secure, but compliant with PCI DSS on an ongoing basis.

It's worth bearing in mind that media coverage of breaches is at an all time high, and is unlikely to decrease as more public organizations and private institutions suffer serious, multi-million-individual breaches. Failing to secure applications could cost a company's its brand, not just a hefty fine.

Yaacov Sherban is CEO of Applicure Technologies (www.applicure.com). The company develops the leading multi-platform web application security software products that protect web servers and internal applications from external and internal attacks. Built upon years of research into hacker behavior, Applicure solutions feature a comprehensive knowledge base to identify attacks accurately, and stop them before they reach the website or application. Applicure's flagship products dotDefender and dotDefender Monitor are deployed internationally and serviced by offices and distributors in the US, UK and Israel.



The largest and most important computer security event of the year - RSA Conference 2008 - took place in San Francisco in April. 17,000 computer security professionals were able to choose from 19 class tracks and more than 220 sessions delivered by the brightest the industry has to offer.

World-class keynote addresses have been delivered by companies such as Microsoft, RSA,

Oracle, IBM, The Security Division of EMC, Computer Associates, VeriSign, Symantec and TippingPoint. Attendees were able to discover and evaluate products and services offered by more than 350 exhibitors.

What follows is a roundup of the most interesting product releases and photos from the show floor.



First integrated SaaS solution for security and compliance



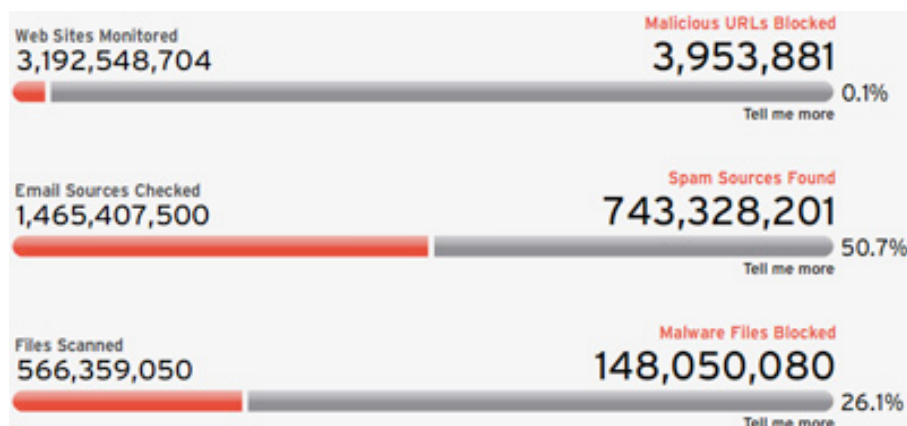
Qualys (www.qualys.com) introduced the QualysGuard Security and Compliance Suite, a suite of SaaS products aimed at helping global organizations to better manage the operational challenges and costs associated with securing their IT infrastructure, and complying with the ever increasing set of regulations. The suite comes in two editions:

1. Enterprise Edition - ideal for large, distributed organizations. Annual subscriptions start at \$25,000, which includes unlimited vulnerability and compliance scans in multiple locations, unlimited number of users, enterprise and scorecard reports and 24x7 customer support.
2. Express Edition - ideal for small to medium-sized organizations. Annual subscriptions start at \$2,500, which includes unlimited vulnerability and compliance scans and 24x7 customer support.



Count web attacks and preventions with TrendTracker

Trend Micro (itw.trendmicro.com) announced that visitors to the company's new online tracker will be able to see, in action, the number of email, URLs and files scanned for malicious content, and subsequently blocked by Trend Micro's in-the-cloud Web threat protection technology.



By scanning "in-the-cloud," Trend Micro blocks access to unsafe Web sites that can automatically download malware. From the spam standpoint, the product shows that of the billions of emails scanned by Trend Micro technology, about 95 percent of them are spam messages; and approximately 32 percent of over 12 billion files scanned contained malware.

Two-factor authentication on SanDisk Secure USB flash drives



SanDisk Corporation (www.sandisk.com) announced the ability to deploy, store and use RSA SecurID software tokens from RSA on SanDisk Cruzer Enterprise USB flash drives. Now available for order, this "two-for-one" solution gives users a single device for secure data storage and two-factor authentication, an alternative to carrying both a flash drive and a separate hardware authenticator.


Cruzer Enterprise is managed by SanDisk's CMC server software, making it easy for IT managers to provision and monitor flash drives throughout their lifecycle. At the same time, Cruzer Enterprise flash drives that contain RSA SecurID software tokens provide two-factor authentication capabilities for remote and mobile network access – requiring something users have and something users know.

SafeNet releases encryption solution for IBM z/OS mainframes

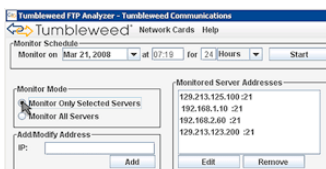
SafeNet (www.safenet-inc.com) introduced a mainframe solution which leverages Ingrian Networks' advanced encryption capabilities to allow customers to protect sensitive information in z/OS environments. Ingrian Networks, a provider of data privacy solutions, was acquired by SafeNet on April 3, 2008.



SafeNet's mainframe solution helps retail, banking, and other financial institution customers achieve PCI compliance by allowing them to quickly encrypt and decrypt critical data in z/OS environments as well as manage bulk encryptions and decryptions within flat files with the SafeNet Transform Utility.



Analyzer for File Transfer Protocol traffic



Tumbleweed (www.tumbleweed.com) announced the availability of FTP Analyzer, a tool providing organizations visibility into their FTP traffic. FTP has emerged as a new security attack vector, resulting in a wave of high-profile data breaches, as well as organized criminal activity, stemming from vulnerabilities in susceptible FTP servers.

The tool provides IT departments with the ability to monitor FTP activity on their networks, in order to assess and address the exposure of sensitive data via FTP. Additionally, this tool generates a two-page executive-level report. (www.tumbleweed.com)

New enterprise biometric and smart card authentication solution

IdentiPHI (www.identiphi.net) launched its new flagship security software, IdentiPHI SAFsolution Enterprise Edition 5. It allows organizations to improve security by replacing cumbersome passwords with biometric or smart card authentication methods. It offers broad support for enterprise platforms and security standards, enabling tight systems integration and simplified deployment.

SAFsolution 5 is currently in pilot deployments with key customers across the U.S. with deployment numbers expected to climb over 100-thousand users by the end of 2008.





IronKey delivers improved mobile data security

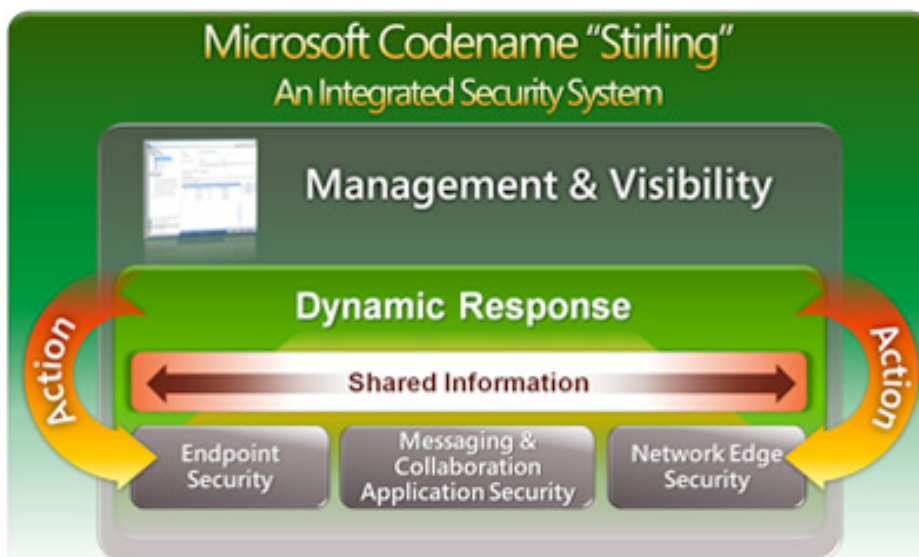
IronKey (www.ironkey.com) announced that RSA has certified interoperability between the IronKey Enterprise line of secure USB flash drives and RSA SecurID one-time password technology through the RSA Secured partner program. RSA SecurID software tokens will now be stored on IronKey Enterprise devices, available with standard features for software token seed provisioning.



The interoperability partnership enables strong authentication and secure access to information through RSA SecurID technology embedded on an IronKey Enterprise flash drive. The combination of two-factor authentication with the IronKey Enterprise's hardware-based encryption, centralized management capabilities, and policy enforcement provides improved security in a convenient format that is easy for IT administrators to deploy and manage.

Beta of Microsoft Forefront security solution

Microsoft (www.microsoft.com) announced the availability of the public beta release of its next-generation Microsoft Forefront security solution, currently code-named "Stirling." It's an integrated security system that is designed to deliver comprehensive, coordinated protection, making it easy to control, access and manage security capabilities across an organization's IT infrastructure.



Forefront "Stirling" includes a central management console for security configuration and enterprise-wide visibility, combined with the next-generation Forefront products that span the client, server and network edge. They include Forefront Client Security, Forefront Security for Exchange Server, Forefront Security for SharePoint and the next generation of Microsoft Internet Security and Acceleration Server (ISA Server), Forefront Threat Management Gateway.

Barracuda Networks launches Barracuda Web Site Firewall

Barracuda Networks (www.barracudanetworks.com) launched the Barracuda Web Site Firewall product line which leverages the capabilities of the award-winning Web Application Controller product line acquired from NetContinuum. Targeted at businesses of all sizes requiring Web application security and PCI compliance, the Barracuda Web Site Firewall starts at \$4,999.



By harnessing the same powerful protection offered by the Barracuda Web Application Controllers, the Barracuda Web Site Firewall secures Web sites against data theft, denial of service or defacement. As a full proxy, the Barracuda Web Site Firewall blocks or cloaks attacks, such as SQL injections, cross-site scripting attacks or buffer overflows, while preventing outbound sensitive data leakage.

The Barracuda Web Site Firewall product line integrates varied degrees of traffic management capabilities, including SSL offloading, hardware-based SSL acceleration and load balancing, which increases both performance and availability of the applications.



Next-generation Secure Access SSL VPN appliances

Juniper Networks (www.juniper.net) announced the next generation of its industry leading Secure Access (SA) SSL VPN platforms – the SA 2500, 4500 and 6500 appliances. The new Secure Access appliances provide enterprises and service providers with best-in-class performance, scalability and redundancy to ensure fast, reliable and secure remote access to applications and services for even the most complex and demanding secure environments.



Juniper's Secure Access SSL VPN appliances enable high-performance businesses to deliver anytime, anywhere access of corporate resources and applications to their remote and mobile employees, customers, and partners. They offer investment protection by providing a single platform to handle remote access to web applications, terminal services, client/server applications, and for the rising use of mobile devices.



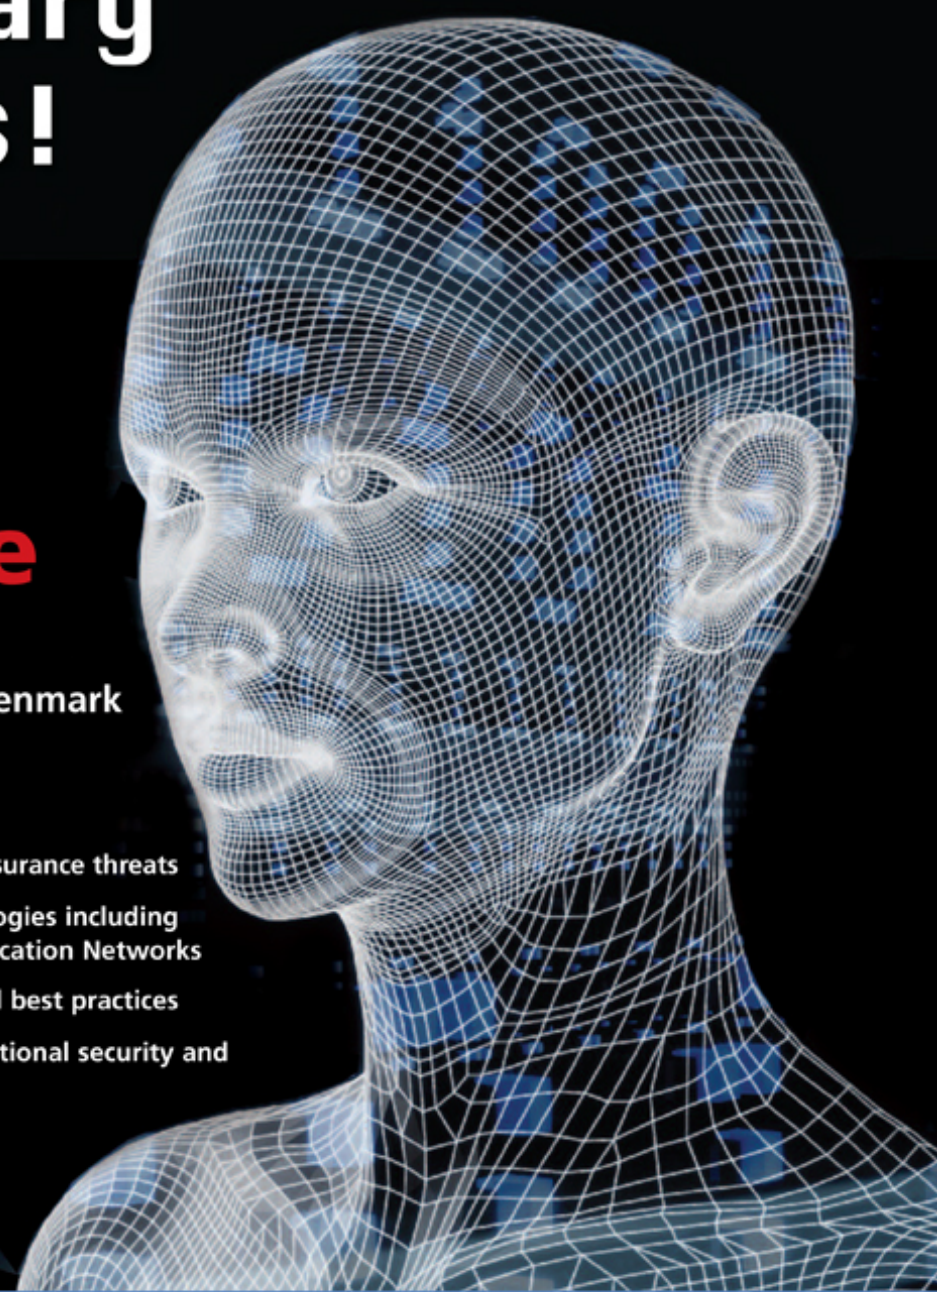
Protect Critical Infrastructure and Military Networks!

Cyber Defence Conference

14th & 15th May 2008,
Marriott Hotel, Copenhagen, Denmark

Gain exclusive insights into:

- ✓ Evolving Cyber Security & Information Assurance threats
- ✓ Vulnerabilities of state-of-the-art technologies including Voice over IP, IPV6 and Wireless Communication Networks
- ✓ Estonian Cyber attack - lessons learnt and best practices
- ✓ Military-civil cooperation in protecting national security and critical infrastructure
- ✓ Latest developments in SCADA Security, intrusion detections and resilient systems




www.smi-online.co.uk/cyberdefence.asp

Alternatively contact Yasir Ansari on Tel +44 (0) 20 7827 6164

Sponsored by





Producing secure software with security enhanced software development processes

By Marco M. Morana

For most organizations, the first approach to application security consists on tactical security assessments such as secure code reviews and web application penetration tests. The scope of these activities is to address any potential security issues introduced during coding and implementation.

With this approach security is considered as a feature to be added to the application rather than the outcome of a security software strategy to identify risks that can be mitigated as the application is being engineered. A security software strategy consists on proactively considering threats and countermeasures throughout the life cycle of the application development from cradle to grave. In a nutshell adopting a security software strategy means embracing security software engineering practices to produce secure software. This includes the designing of the artifacts that the application is engineered with: use cases, requirement documents, high level and detail level design documents, source code, unit test cases, system test plans and the production build.

A pre-condition before rolling out such initiative is to understand the organization context in which software is developed that is the organization's software engineering practices

and processes and the roles played by stakeholders at different level including, Information Security Officers, Project Managers, Architects, Developer Leads and least but not last Software Developers.

Another pre-condition is to deliver software security training and awareness training that in essence means communicating to project stakeholders what software security entitles to and when and how security needs to be applied to the application: in other terms explaining the difference between application security and software security.

The difference between applications and software security is best described with the Chinese philosophical concept of "yin" and "yang": opposing and, at the same time, complementary (completing) aspects of one software development process that can be seen from different perspectives.

For example, from the “time” perspective, application vs. software security consists on applying security after application is already build vs. while the application is being build, from the problem solution perspective is looking at the symptoms instead of the root causes, from the solution approach perspective is catch and patch vs. fixing security bugs and flaws throughout the software life cycle, from the risk perspective is reactive response to security incidents vs. proactive threat analysis, risk assessment and implementation of countermeasures as necessary.

Besides requiring different perspectives, software security also requires different people skills such as programming (e.g. coding) and information security skills. From the programming/coding skills perspective, this represents a major challenge for organizations when these people skills cannot be found in-house. In fact is very likely that such skills cannot be found among both application

security and information security people within the organization. If this is the case, for your software security initiative to be successful is a prerequisite to build a team of secure software technologists that can understand both coding and security.

In practical terms, it means to look for software security skills that usually are not found among typical information security professionals such as:

- Ethical hackers that know how to break into the applications but cannot tell you how to build a secure one.
- Security engineers professionals that know how to run security assessment tools but do not have grass roots (aka experience) in software engineering, designing application and coding.
- Information security professionals with little or no experience with checking security compliance in web applications.

Ethical hackers that know how to break into the applications but cannot tell you how to build a secure one

Assuming that you have finally developed a team of secure software technologists the next step in adopting the secure software initiative is to make the case for it to your organization. If your organization is mature on information security processes and very little on software and application security, you might face challenges in deliver a software security initiative within your organization. In your day to day activities, you might need to fight a typical mentality that is detrimental for software security: looking for the magic solution tool or process, in essence the silver bullet mentality.

For any security initiative (not just secure software) a tool (e.g. source code analysis, web scanning tools) or a technology (e.g. application firewall, multi-factor authentication control) alone does not buy your organization a solution of all the software security problems, it is at the best only part of the problem you are trying to solve. If you would like that a software security initiative becomes a solution

of the software security problems within your organization you need to approach it from the software engineering and risk management perspective. To approach security from the software engineering perspective means taking a step back (from the strategy) and trying to understand better how software is developed within your organization, what where and how secure software can be built and how security risks can be effectively managed.

If you are the promoter of the software security initiative within your organization you need to be realistic on what goals can be achieved in the short and in the long term and on which strengths, people, process, tools the software security initiative can rely upon. The first step is to ask basic questions, a sort of organization level software security assessment. From the business perspective for example, does my company develops shrink wrapped software or delivers services via web applications developed in-house?

From the process maturity perspective, do I develop software with life-cycle methodologies, unified processes and tools or with ad-hoc methodologies and scope-limited quality processes and tools? Do different engineering departments within my organization use different software development methodologies or there is a common process to develop software across the organization? Do I know which software development processes are used within my organization? Are there any standard life-cycles used across the organization? To build my security in to the SDLC (Software Development Life Cycle) can I adopt a security enhanced SDLC such as MS SDL, CLASP, "Touchpoints" or I should build a security software framework and adapt my software security best practices to my SDLCs being waterfall, spiral, Agile, and RUP? Are tools being used sparsely and different from department to department or there is a common tool and process used across the organization? Does my organization use a consistent metrics for dealing with security issues across departments or it is just left to ad-hoc initiatives to know what to track and what to measure? How are my security risks analyzed and managed? I act upon incidents or I try to proactively identify potential risks and proactively trying to address them? There is any risk management process?

The gold rule for any software security initiative is undertaking the essential main steps: (1) assessment (2) implementation and (3) measurement.

Step 1 - Assessment: Answering specific questions on your organization software engineering and risk management processes is one way to perform the assessment. Formal methods can help assess how the organization processes are currently managed and how can be implemented and improved.

One of the objectives of every process is to improve over time and become standardized and adopted across different department of the organization. The Capability Maturity Model (CMM) provides a model to assess the level of maturity in processes and what activities need to be performed to reach a higher maturity level. This can be applied to the organization software security processes, the people and the tools. By applying software

security through the CMM you can tie the adoption of software security activities to certain level of software security assurance. For example, for most organizations moving from traditional network centric and application centric approach to software security approach means adopting a software security framework. This adoption cannot happen overnight, by stating a famous phrase: "Rome cannot be build in one day", secure software will not be built in one day either. The main reason is because you need time to mature you processes, training and tools. From the tools perspective for example, it means starting from a proof of concept and then seek adoption throughout your organization, from the process perspective means standardize organization wide software development and risk management processes. From the people perspective it means to provide training for the organization as a whole instead ad-hoc training for some departments as required.

An example of reaching maturity security levels depending on the software security activities and the time frames, from low maturity to medium maturity in the short term and from low maturity to high maturity in the long tem is included herein:

1. *Low Maturity (CMM 0-1)*

- a. No formal security requirements
- b. Issues identified through penetration testing and security incidents
- c. Issues are identified and corrected late in the lifecycle
- d. Penetrate and patch and reactive approach

2. *Medium Maturity (CMM 2-3)*

- a. Applications are reviewed for security design flaws with threat modeling
- b. Source code is reviewed for security with source code analysis
- c. Penetration tests validate issues dealt with earlier in the SDLC

3. *Advanced Maturity (CCM 4-5)*

- a. Threat analysis in each phase of the SDLC
- b. Risk metrics and measurements are used for improving security engineering and risk management processes.

Step 2 - Implementation: This means implementing the software security initiative and adopting software security best practices throughout the SDLC such as secure requirements during definition, secure

architecture during design, secure coding during implementation and secure tests during validation.

When such best practices are documented as security requirements, you can start implementing checkpoints as part of information security review process: the essential ones are the security reviews during design, secure coding reviews during development and vulnerability assessments (aka penetration tests) during testing.

To be effective in mitigating security risks, software security activities needs to be threat driven: threats identified during threat modeling can be used to derive test cases for security tests for example. A complete software security process also takes into account of the security configuration of the application when is deployed and is tied back to code changes

that are necessary to support secure operations.

Ultimately a mature software security process blends both information risk management and software engineering processes in a software security framework. For example, threat modeling will identify threats and technical impacts during design that are used as a factor along with business impact in the calculation of the overall risk. Ideally, such mature software security process should integrate software security activities in each phase of the SDLC. This also included activities that are performed as part of the operation life-cycle such as patch and incident management as well foundational activities such as metrics and measurements and security training and awareness. An example of such software security framework for a waterfall SDLC is included in Figure 1:

SDLC Phases	Requirements	Design	Development	Testing	Deployment and Operations		
Secure Software Best Practices	Preliminary Software Risk Analysis	Security Requirements Engineering	Security Risk-Driven Design	Secure Code Implementation	Security Tests	Security Configuration & Deployment	Secure Operations
Ongoing S-SDLC Activities Metrics and Measurements, Training, and Awareness							
S-SDLC Activities	Define Use & Misuse Cases	Define Security Requirements	Secure Architecture & Design Patterns Threat Modeling Security Test Planning Security Architecture Review	Peer Code Review Automated Static and Dynamic Code Review Security Unit Tests	Functional Test Risk Driven Tests Systems Tests White Box Testing Black Box Testing	Secure Configuration Secure Deployment	
Other Disciplines	High-Level Risk Assessments		Technical Risk Assessment				Incident Management Patch Management

Figure 1: Software Security Framework

Ideally, for a software security framework to be useful in practice, it needs to be applicable to the different software development methodologies used by your organization. This might include evolutionary-interactive software life-cycles such as spiral, RUP, XP and Agile besides the traditional waterfall model.

Currently, there are several security vetted security-enhanced lifecycle process (S-SDLC) models that can be adopted by organizations while designing building and testing secure software:

1. Microsoft Secure Development Lifecycle (SDL)
2. Common Lightweight Application Security Process (CLASP)
3. Touch Point Model (TP).

From the perspective of software engineering and vulnerability analysis, all these models have a common set of criteria for software security:

1. The definition of software security best practices as security activities that can be performed at different phases of the SDLC
2. The definition of checkpoints (e.g. software security tollgates)
3. How such activities should be conducted (e.g. guidelines) and by who
4. A set of tactical resources (e.g. checklists, common vulnerability lists)
5. A provision for the use of automation tools (e.g. static code parsers, threat modeling tool, risk management tools).

Security-enhanced lifecycle process models differ in some of the aspects of the implementation, the scope of the analysis and the different emphasis to people, process and tools/technology factors.

Instead of trying to exhaustively describe each of the models (not feasible in the context of this article) I would like to describe what the

essential elements of each model are and emphasize the pros and the cons of each of these models. Specific information about these security-enhanced lifecycle process models can be found through public available documentation on-line and the essential bibliography that is provided in the resources section of the article.

Microsoft Secure Development Lifecycle (SDL)

Essential elements: Verification check points, developer's training and tools.

The aim of the Software Development Lifecycle SDL is to integrate tasks and checkpoints in the software development organization's processes so that software security can be improved with well defined security reviews and security deliverables.

Critical to the effectiveness of the SDL implementation is the ongoing education and training for software developers and the techniques and technologies required for secure development. Figure 2 below depicts at high level the phases of the SDL and the software security activities. The SDL software security activities (the yellow blocks) map to the SDLC (the big arrows) such as:

1. Definition of security feature requirements during the requirements phase
2. Security architecture design, best practices and threat modeling during design
3. Static analysis via automation tools during implementation
4. Application of coding and testing standards during implementation
5. Code reviews during implementation
6. Security focused testing, pen and fuzzing testing, during the verification phase
7. Security servicing and response during operations (post release).

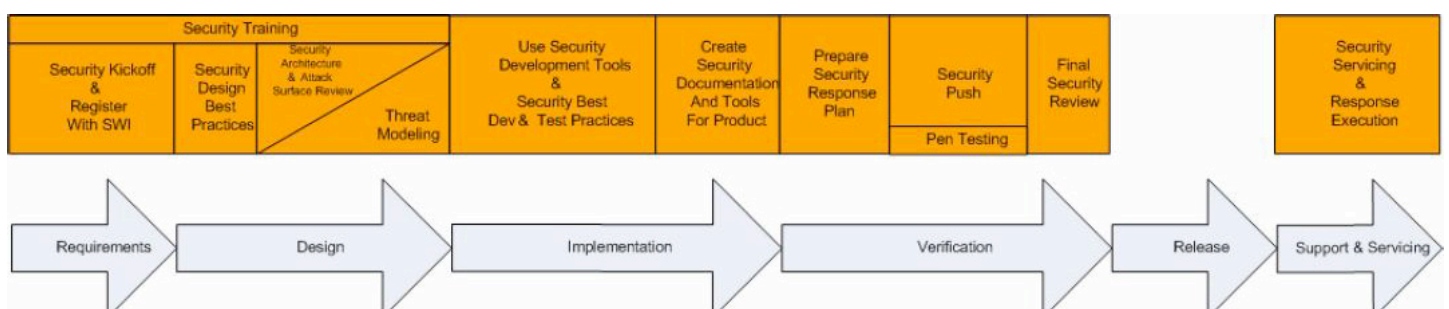


Figure 2: MS SDL

Pros and cons

Pro: Microsoft products that have adopted the SDL such as SQL Server 2003, Vista and IE6 had significant less security issues than products that did not use it. This is data validated by comparing with the number of vulnerabilities in software versions developed prior to the introduction of the SDL: software developed under the SDL have exhibited over a 50% reduction in vulnerabilities.

Pro: The SDL methodology is adequately supported by tools such as threat modeling and source code analysis tools. Some of these tools (e.g. Prefix and Prefast) are part of Microsoft Integrated Development Environment for which a license is needed while others are free (e.g. ACE Torpedo TM). So far SDL is the only security-enhanced lifecycle process model that support threat analysis with a tool.

Pro: MS SDL is very well documented, easily accessible via MSDN web site and very extensive including secure software handbooks for developers, threat and countermeasures guidelines and checklists etc.

Con: Threat analysis as is done in the MS SDL is very product oriented and static because is build around the Microsoft risk model STRIDE/DREAD. This risk model might work well for a software development shop and for shrink wrapped software development but it will not be suitable for applications developed for organizations that use proprietary risk management models such as financial and banking institutions for example.

OWASP Common Lightweight Application Security Process (CLASP)

Essential elements: Emphasis on roles and responsibilities, lightweight process phases and interactions.

CLASP (Comprehensive, Lightweight Application Security Process) helps organizations in producing secure applications with a not too cumbersome and still very systematic process that can be adapted to different software development methodologies. In CLASP security activities are role driven, that is these activities are assigned to the roles with designating owners and participants. The roles include the project manager, the requirement specifier, the architect, the implementor, the test analyst and the security analyst. For each role-based activity CLASP describes:

1. When and how the activity should be performed
2. The level of risk associated with omitting the activity
3. The estimated cost for implementing the activity.

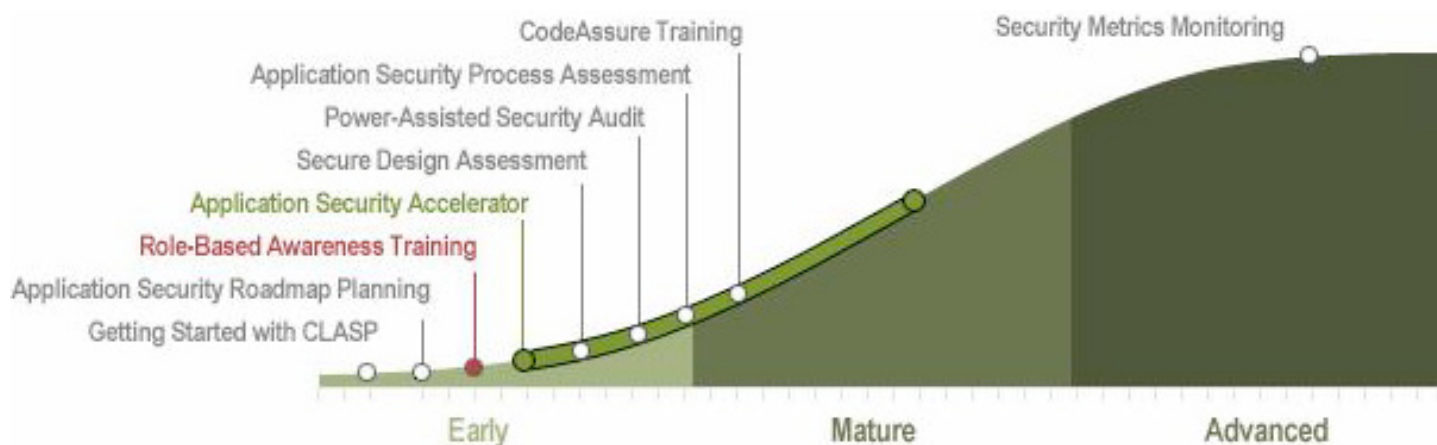
As a lightweight process CLASP is designed for adaptation to any software development life-cycle: it can be either integrated with interactive software development life-cycles such as Rational Unified Processes (RUP) or explicitly mapped to life-cycle phases such as requirements, architecture & design, development, testing and deployment. The process includes several security activities that are build around security best practices depicted in Figure 3.



Figure 3: CLASP best practices

Another aspect of CLASP being lightweight is in the guidance on when to perform the activities such as “when it is necessary”, “as required” and “as needed”. Every activity is characterized by a frequency assuming an interactive process such as RUP for example. Most of the activities are once per life-cycle interaction such as “Document security-relevant requirements” and “Apply security principles to design” while other activities can be performed multiple times per interaction such as: “Detail misuse cases” and “Identify, implement, and perform security tests”. Other activities are phase specified. “Threat modeling” for example should be performed for each lifecycle once initial requirements are specified and once near feature complete

(end of construction) while “Define attack surface” should be performed once during design and ongoing during elaboration. “Source code analysis” should be performed at the end of each implementation (construction) interaction. All the other activities can be performed “as needed”, “as required” and “as necessary”. A peculiar aspect of CLASP is the choice of roadmaps to both legacy application and new applications (Greenfield roadmap). Another unique aspect of CLASP is the emphasis on process improvement and maturity as depicted in Figure 4. Security activities are introduced only after some maturity level is reached such as “Role-Based Awareness Training” for example.



The area under the curve represents the cost for achieving maturity in terms of training, tools and activity implementation. The steeper the curve the highest the cost

Figure 4: CLASP Maturity Curve

Pros and cons

Pro: Allows integrating security related activities into existing application development processes.

Pro: Includes instructions, guidance and checklists for security activities and a collection of methods and best practices.

Pro: Each activity description includes when and how it should be performed, the level of risk of omitting the activity and the estimated effort.

Pro: Contains a vulnerability taxonomy that can be enforced through use of automated tools for static analysis of source code.

Pro: Provide two roadmaps, one that support focus on development of new systems using

an “interactive” process and one that focuses on maintenance of deployed systems.

Pro: Activities implicitly map to general SDLC phases and interactions to explicitly map to RUP.

Cons: Provide guidance to specific roles that are not always representative of project stakeholder roles in some organizations.

Cons: There is no specific mapping to software development methodologies such as Agile and XP.

Touch Point (TP) Model

Essential elements: Emphasis touchpoints, risk management, security knowledge bases.

The touchpoint software security methodology is authored by Dr. Gary McGraw and is dealt

extensively in his book: *Software Security: Building Security In* (Addison-Wesley, 2006).

“Touchpoints” are “lightweight” best practices that can be applied to any software development lifecycle. There are 7+1 bonus “Touchpoints” in this enhanced software security development lifecycle the + 1 being the external review, in which analysts from outside the organization perform un-biased security reviews also referred as vendor security assessments.

“Touchpoints” activities can be performed in different phases to the SDLC and are tied to “artifacts” rather than SDLC phases so that can be interactively applied to both linear (e.g. waterfall) and evolutionary (e.g. spiral, RUP) software development methodologies.

Figure 5 shows how “Touchpoint” activities are applied to software development artifacts and the bonus external review activity (the arc arrows).

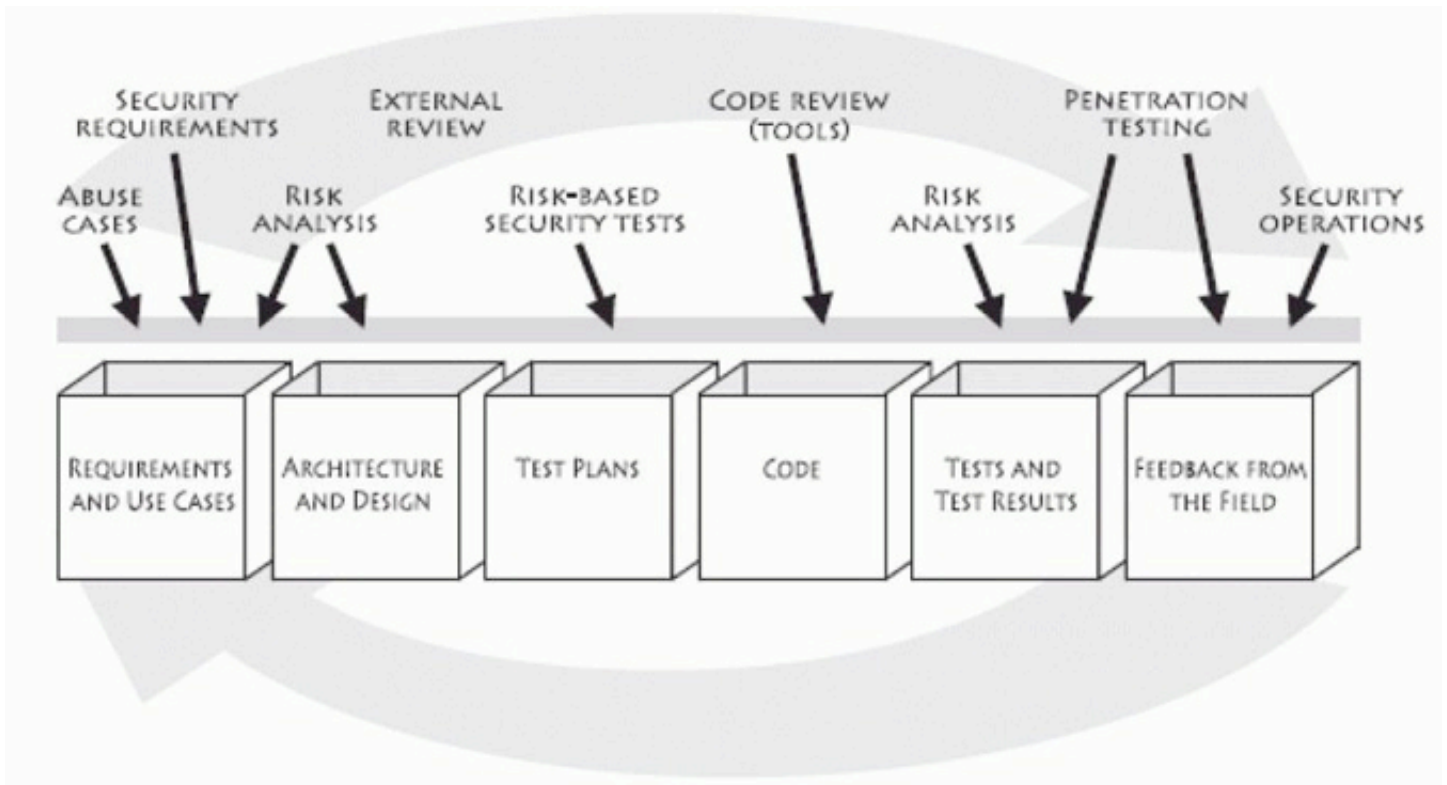


Figure 5: TouchPoints SDLC

In summary this is the description of the 7 necessary “Touchpoints” to deliver secure software:

Abuse cases: methodology to derive abusive cases from use cases so that countermeasures for such abusive scenarios could be documented in requirements and use cases

Security requirements: encompasses both functional requirements for security controls and risk mitigation driven requirements from the abuse case scenarios

Risk analysis: is the methodology to identify risks that can be documented in software requirements and assessed both during design and validation tests. In the case of design, risk

analysis is equivalent to threat modeling that is categorization of design flaws, vulnerabilities and the associated risks (the impact resulting from the exploitation of the vulnerability). In case of security tests, risk analysis consists of testing the application based upon threat scenarios that are identified via abuse cases and attack patterns. Part of risk analysis is also risk monitoring that is collection of risk metrics throughout the SDLC

Security risk based tests: Include both functional tests and risk driven tests and includes how to derive test cases from abuse cases and attack patterns.

Code reviews: Consists in the review of source code using source code analysis tools

(e.g. static parsers) to identify vulnerabilities in source code.

Penetration tests: Is the use of automation tests to perform ethical hacking to verify the resilience of the application to common application vulnerabilities as well as specific vulnerabilities identified during the architectural risk analysis. As a touch point, penetration tests can be performed both against new applications to produce test reports and existing applications in production to produce a “feedback from the field”.

Security operations: this is an activity that is not part of the SDLC and is meant to feed secure software engineering (the other 6 activities) with knowledge base risk data gained when the application is exposed to security incidents, attacks and exploits during operations.

Pros and cons

Pro: Most of touch point activities are driven by risk analysis that is identifying threat scenarios as early as during the inception of the software with abuse cases and architectural risk design as well as for testing during the verification of countermeasures.

Pro: There is emphasis on building a risk knowledge base where identified software risks are measured and reported for further analysis. This support risk data driven initiatives and help to build a business case for software security.

Pro: Software security activities target artifacts rather than SDLC phases so can be used for different software development methodologies.

Cons: There is not explicit emphasis of software security training as being a pre-requisite to perform the security activities. As on-going activity training and awareness are necessary to provide project stakeholders the software security training and education in software engineering and risk analysis.

Cons: There is no guidance on how to adopt touch points as effective software development practice within an organization based upon on maturity of software development process and risk management disciplines.

Step 3 - Measurements: By assessing the maturity of secure software engineering prac-

tices within your organization it is possible to set the achievable goals such as which activities can be performed in the short term, mid term and long term. It is important to measure how the goals are achieved and which ones.

Adopting a software security metrics is a critical factor in measuring the effectiveness of secure software activities within the organization. For example, setting a baseline for the security posture of existing applications allows comparing results after the security activities have been implemented. Security metrics can be tied to specific objectives and goals when are SMART that is Specific, Measurable, Attainable, Realistic/Relevant and Time-bounded. An example of a SMART metrics goal can be reducing the overall number of vulnerabilities by 30% by fixing all low hanging fruits with source code analysis during construction.

Having a software security metrics in hand is the best way to make the case for software security to the software security initiative stakeholders. If the data shows what you are doing best and when as well as what your gaps are it is much easier to make the case for software security budget and resources since you can prove how effective these can be for your organization.

One of the most critical aspects of enacting a software security program is gain support from different levels of management within your organization: this might require fighting some misconceptions such as security impacts performance, security impacts costs and security impacts development. This “fighting” might involve different battles depending on your role within your organization: as developer lead you need to make the case to developers that are tired to rebuild their application because a security auditor changed his mind on requirements. As engineering director you have to make the case to project managers that worry about missing deadlines and how software security activities could affect the budget, the costs and the performance. As information security officer you need to make the case to the CIOs that worry about putting money in a process a tool or a technology and not being able to show the return on investment to the company executives.

In all role cases this means communicating effectively the case for software security, showing where the problem is (for example source code!) by supporting it with data but also by providing potential solutions and trade-off options to project stakeholders involved in the software security initiative.

To software developers software security metrics can show that software can be build more securely and efficiently using source code analysis tools when secure coding standards are followed and software security training is available.

To project managers, software security metrics shows that projects are on schedule and moving on target for the delivery dates and are getting better during tests. Such metrics also builds the business case if the initiative comes from information security officers. For example, it needs to provide evidence that security checkpoints and assessments that are part of the SDLC do not impact the project delivery but rather reduce the amount of resources and the workload needed to address vulnerabilities later in production and allow to deliver the project on time. It shows that projects do not get stack during the pre-production phase fixing high risk vulnerabilities for example.

To compliance auditors a software security metrics provides a level of assurance that security standard compliance is addressed through the security review processes within the organization.

Ultimately the software security metrics needs to be tied to the business case and support the business decision making. For a CIO and CISOs that need to make decisions on budgeting for right resources and technologies this means showing the Return Of Security Investment (ROSI) metrics for example. Return on Investment (ROI) over-simplified means that if I spend \$100K on something, I want to know that in a certain period of time the money I spent is going to return something back to the organization.

A good metrics for software security should support both cost and risk management decision making and answer some critical questions such as:

1. Is the investment in the security solution (e.g. SSDLC activity, tool etc) financially justified?
2. Is the cost of the security solution less than the cost of loosing my assets?
3. What are the most cost effective security solutions?
4. Which investment in security activities is most effective in reducing costs in fixing security defects?

Making the business case for software security needs to take into consideration the culture of the organization, the responsibility of each role and the commitment to the software security initiative. For large organizations a CIO/CEOs commitment is necessary to sign-off on software security standards and policies and define the strategic risk posture and influence the culture of the organization. A CISO/ISOs commitment is required to document new software security standards and build them into the risk management, compliance and audit processes.

Your organization commitment on the adoption of the software security initiative is really the essence of what you can achieve: you might have the best processes, people and tools but without commitment there is not security as outcome, in mathematical terms it is like a value multiplier (e.g. multiply by zero produce zero result).

Commitment in the software security initiative really depends on how successful your management is in pushing it throughout the organization. It might come from the top CIO/CEOs and pushed down to different levels and become a reality in very short time. A top down software security initiative is what Bill Gates did at Microsoft in 2002 with the security push of "The Trustworthy Security Initiative": two months of security freeze for security training, funding the "Find the Bug" contexts and delivery of the Security Development Lifecycle SDL across Microsoft projects.

From the bottom up perspective software security could be pushed from information security officers that can show control gaps in compliance with software security standards, from project managers, architects and engineering leads that can show that can address

vulnerabilities in the SDLC more efficiently with less impact on the life-cycle delivery of projects.

Pushing a secure software initiative will require day to day partnering with different en-

gineering departments within your organization, build the business case with project stakeholders and showing with data (e.g. metrics) how such software security initiative is successful so that the case for it is re-enforced and the plan can be executed.

Marco serves as a leader of the OWASP (Open Web Application Security Project - owasp.org) where he contributed to write the OWASP Security Testing Guide. Marco and also works as Technology Information Security Officer for a large financial organization with key roles in defining the web application security roadmap and activities, document security standards and guidelines, perform security assessments for software security as well as training software developers and project managers on the software security and information security processes. In the past, Marco served as senior security consultant within a major security consulting company in USA where his responsibilities included providing software security services for several clients in the banking, telecommunication, computers and financial business sectors. Marco also had a career in the software industry in diverse professional roles such as contractor, senior software engineer and project manager with responsibility to design and to develop business critical security software products for several FORTUNE 500 companies as well for the US Government (i.e. NASA).

Marco is active in publishing on the topic of software security for several professional organizations (ISSA, OWASP) as well as on his blog at securesoftware.blogspot.com. Marco can be contacted at marco.morana@gmail.com



OWASP

The Open Web Application Security Project

JOIN US! OWASP is a free and open community dedicated to improving application security for everyone.

You'll find free tools, books, articles, best practices, mailing lists, conferences, and local chapters around the world to help you build secure code.

www.owasp.org



Network event analysis with Net/FSE

By Ben Uphoff

Incident response and network operations teams rely on IP-based network data, or network events, for a variety of tasks in their daily operations: alert analysis, network monitoring, trend analysis and network forensics among others. Alert analysis tasks involve daily investigation of alerts from intrusion detection and prevention systems (IDS and IPS), network behavior analysis (NBA) tools, security information management (SIM) systems and various other alerting tools.

Network monitoring and trend analysis is an ongoing task that is necessary to understand the daily rhythm of the network. Network forensics comes into play in the event of an investigation and involves analyzing weeks or months of network data to determine the scope and extent of a network breach.

Ultimately, these four tasks are data analysis problems and the predominant data source in network security is network event data. For the purposes of this article we consider network event data to be log records containing IP address information generated by a device on the network or any alerts, in the form of logs as well, generated by systems processing packets (such as intrusion detection systems) or log files (i.e. network behavior analysis).

Network event data is crucial to network security analysts for two primary reasons: it is ubiquitous and it is passively collected. Enterprise routers, and some switches, can generate high volumes of information-rich flow data (NetFlow, sFlow, J-Flow or IPFIX depending on the vendor), that can be leveraged for network forensics, insider threat detection and anomaly detection (including network behavior analysis).

Additionally, firewalls, web servers and security appliances generate logs that are highly relevant to network security operations. Unlike network-based logging, host-based logging can be easily deactivated or erased once an attacker has compromised the host.

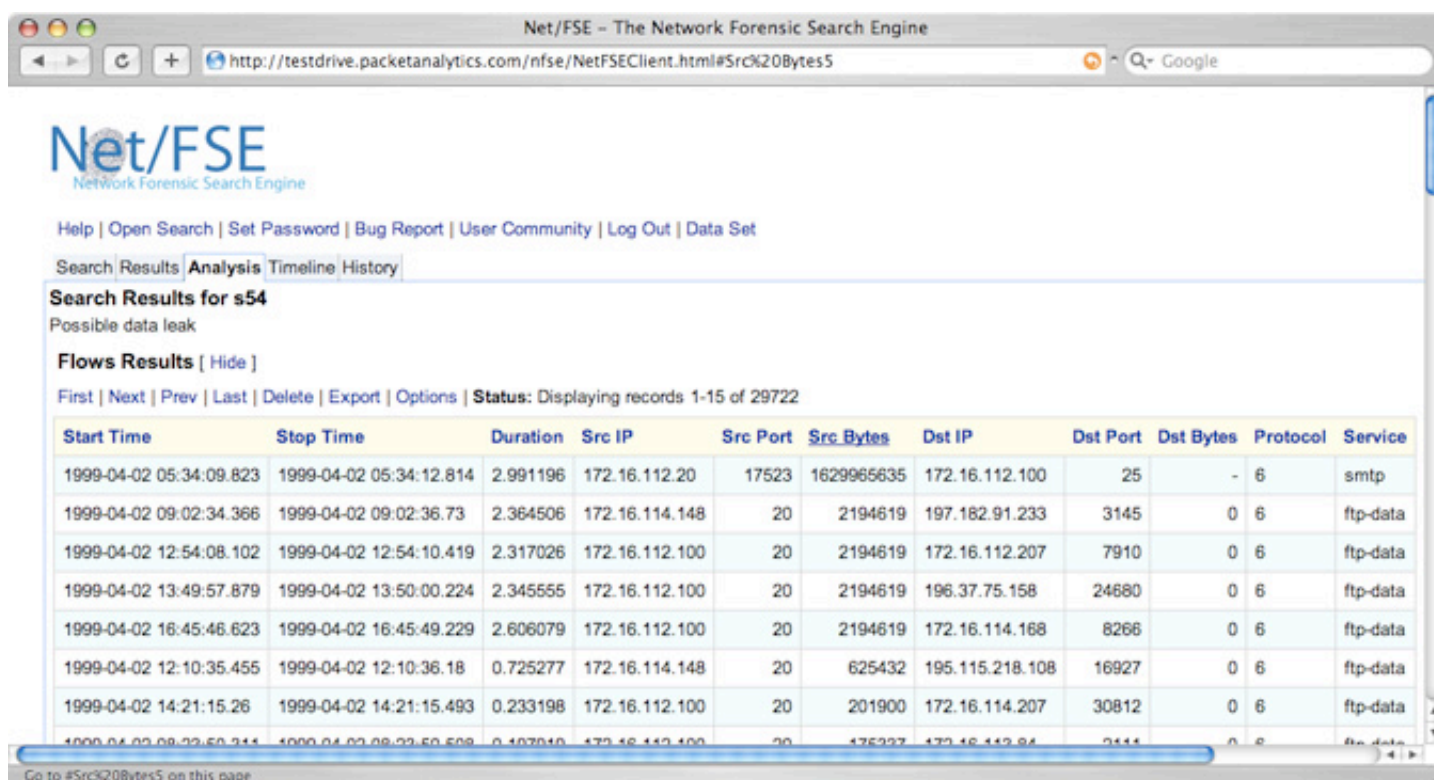
Network event data is generated automatically and passively, so an intruder on the network cannot avoid generating network events that can later be used by analysts to track back the activities of the intruder.

Incidents are inevitable

A core belief at Packet Analytics is that despite the best efforts of security vendors and practitioners, incidents are inevitable. There are simply too many threats and too many angles of attack. Technology on enterprise networks evolves so quickly that it is nearly impossible to keep up with the ever-changing threat landscape. For this reason, network breaches and security incidents must be seen as part of doing business in a connected world. Enterprises can mitigate the risk of a breach by following best practices and preparing a comprehensive incident response and recovery plan.

Best practices in network security involve pairing diverse, layered defenses with the ability to rapidly react, respond and recover when the inevitable security incident occurs. Defenses include the usual: firewalls, anti-virus, patch management, intrusion detection, logging for forensics and policy enforcement. These defenses need to be deployed throughout the enterprise for defense in depth, not just the perimeter as is still common in many networks.

Since even the best defenses fail at times, reacting, responding and recovering from security incidents are critical aspects of incident response that cannot be overlooked in the enterprise security plan. Preparation saves enterprises money by minimizing the impact of a network breach. A formal incident response procedure coupled with effective analysis tools reduces the turnaround time on investigations and leads to definitive answers.



Users can quickly drill down to analyze security incidents and perform network forensics.

Net/FSE

Packet Analytics' Net/FSE, the network forensic search engine, is the first commercial solution available to network security analysts that is built from the ground up to make network

event analysis operations cost effective, faster and more efficient. Net/FSE, available as a free download, brings together event data from network devices and gives security analysts the ability to correlate and analyze billions of events in real time.

Net/FSE gives the security team the ability to collect any type of network event data, including flow data (unlike many SIM and log management solutions) that can be generated by almost every enterprise network router and is essentially a free resource of forensic information. Other valuable information sources for Net/FSE include alerts from IDS, IPS, SIM and NBA, firewall logs, web server logs, authentication logs and database server access logs.

Net/FSE is the only product on the market designed specifically to address the daily data analysis challenges faced by security analysts. Log management, security information management and network behavior analysis systems all fall short in this area due to the fact that they are solutions designed to solve other problems but are often applied in practice for data analysis. Net/FSE is not a replacement for any of these solutions but is instead a complimentary tool that works as a value-add, providing security analysts with the tools they need to quickly and efficiently solve analysis problems.

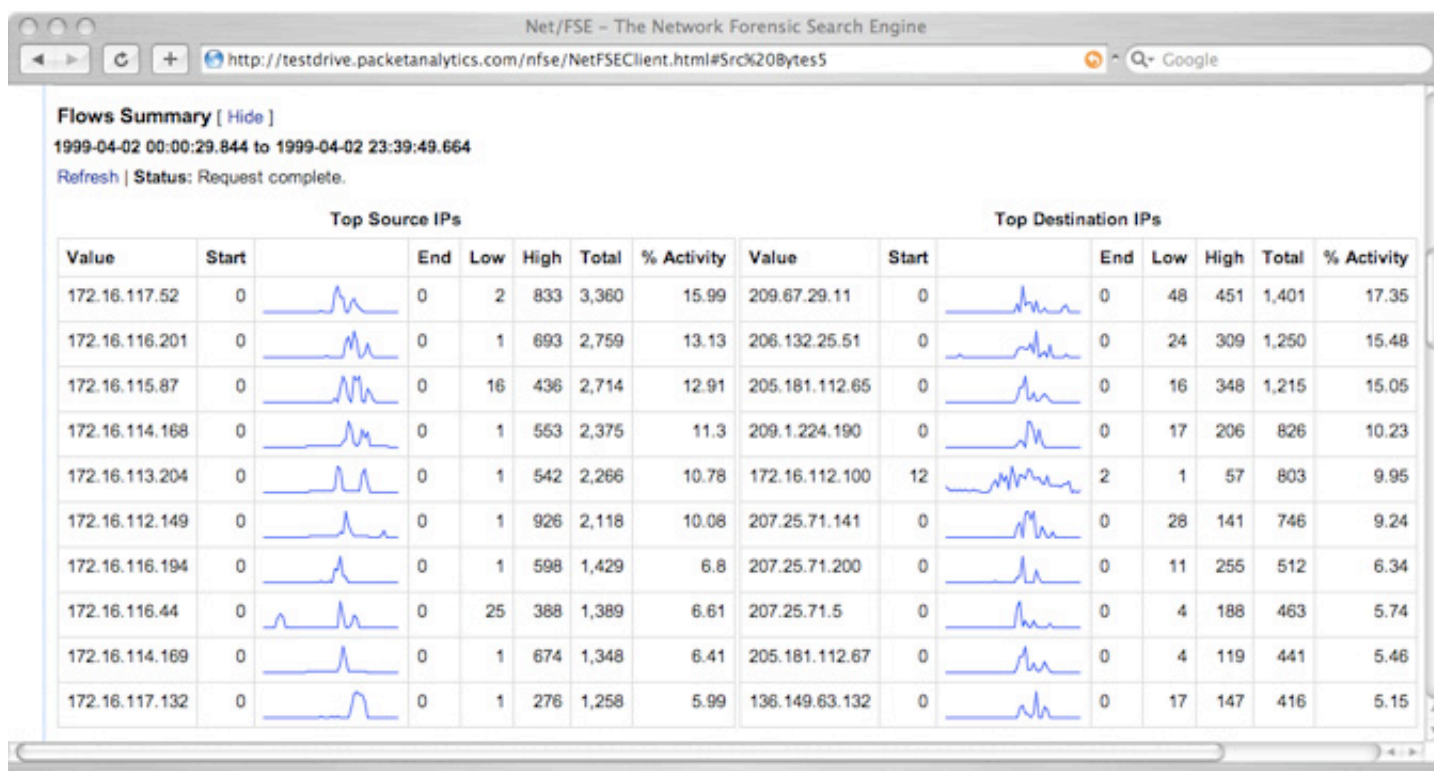
Keep it all

Every day, enterprise networks generate millions of network events that are relevant to the network event analysis tasks outlined above.

Firewalls, web servers, routers, switches, authentication servers - practically any network device - are constantly creating log records that can be collected, searched and analyzed. More and more enterprises are beginning to see the value in their event streams but few are leveraging them for effective incident response and network data analysis. Even fewer are keeping a large and diverse event repository for long-term storage.

One challenge with working with network event data is that you can never be sure what event information is relevant until after the fact. For example, enterprises did not see value in storing DNS logs until DNS exfiltration attacks started appearing. With no historical log of DNS activity, those that fell victim to such attacks had no way of definitively knowing the extent of the data leakage resulting from the breach.

What seemed like noise yesterday can quickly become highly relevant in the event of a security incident, as was the case for DNS events for most enterprises. This leaves one option for security operations: keep it all. Disk is cheap and nothing compresses better than log files, so long as the tools are available to provide access to large, potentially compressed, log repositories.



Net/FSE's incident visualization capabilities provide a visual and statistical overview of a security incident.

Contrary to the “keep it all” approach, SIMs try to reduce data volume at the collection points by aggregating similar events into statistical summaries that are then fed into the correlation engine, losing potentially valuable information in the process. Summaries are useful for the correlation engine but not for deep analysis of network events.

Another approach SIM vendors use to reduce data volume is to allow for enterprises to filter out “unnecessary” events from a data stream. This of course requires a prior knowledge of what information is important, which, as we have already discussed in the example of DNS events, is not possible.

SIM solutions perform data reduction because adding all that perceived noise into the system will quickly impair performance on the backend correlation engine and lead to inaccurate event correlation – the primary mission of any SIM. This differs drastically from the design of Net/FSE, which is designed from the ground up to store any and all network event data with minimal storage overhead.

As SIMs do not meet the data analysis needs of enterprises, there is a place for both a SIM and Net/FSE in any security-conscious network. By combining these two technologies, enterprises get the benefits of powerful alerting capabilities coupled with deep forensic and analysis.

Search first, then analyze

Finding interesting information quickly is a key to any search engine application. Simply displaying search results is not enough for network security analysts. There is simply too much information, most of which is noise. Unfortunately, today’s log management solutions focus on presenting information in the form of charts, graphs and reports. While useful for IT functions like compliance and quality of service, these features do little for helping network security practitioners analyze network event data.

Search is the first step in launching an investigation or analyzing event data. Generally some basic information is known about an incident: time ranges, hosts involved in the event (both internal and external) and a gen-

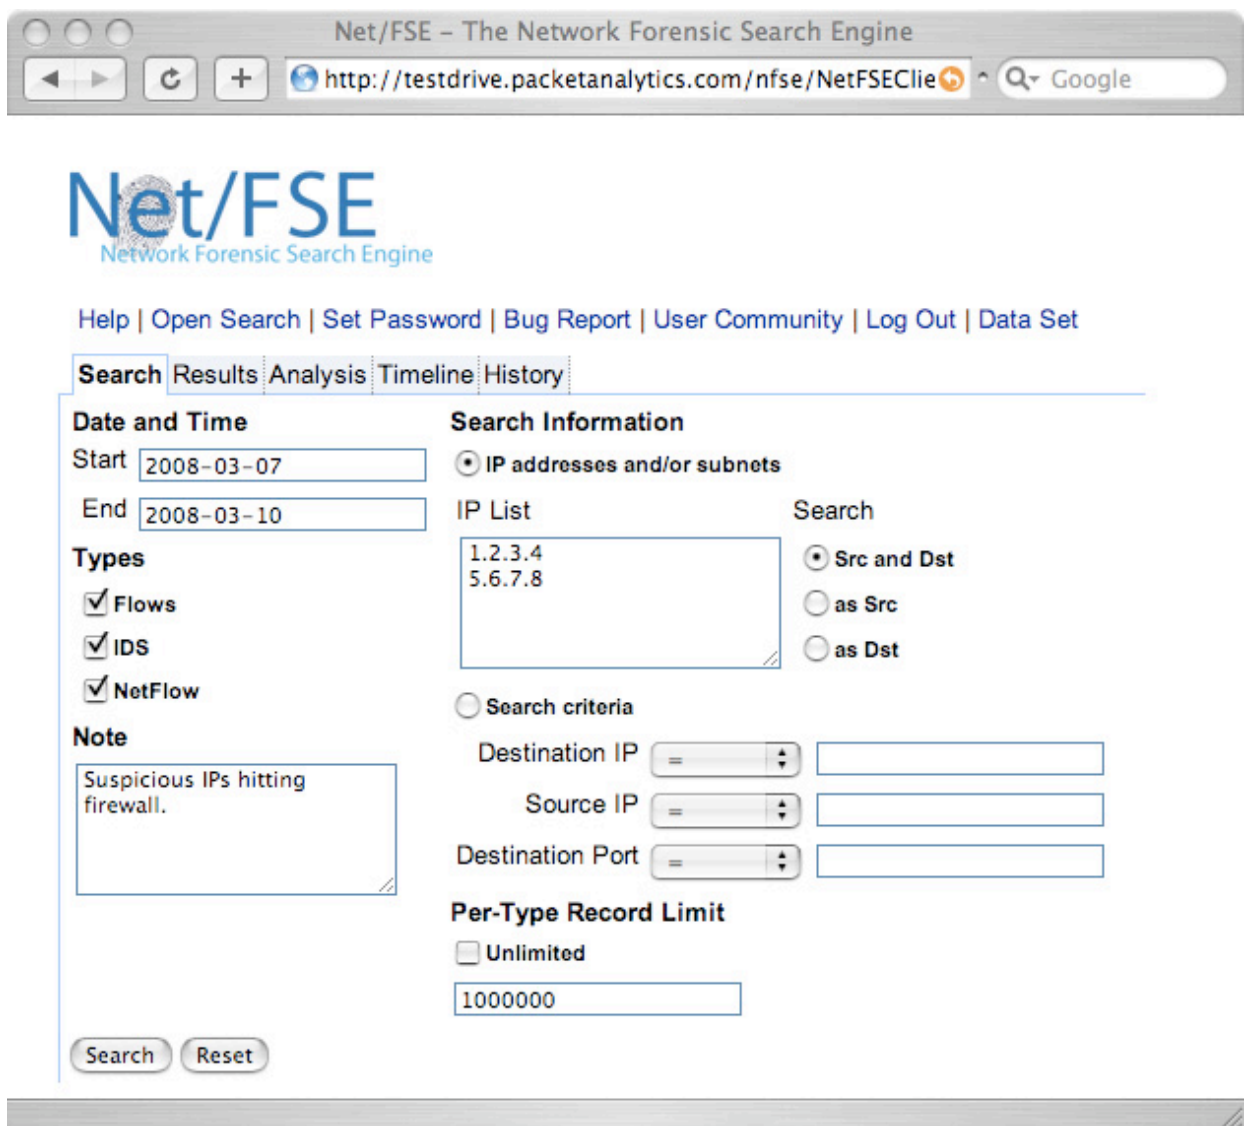
eral idea of what network event information will be relevant to the investigation. Net/FSE’s search interface is engineered to take this information and quickly identify events that meet the security analyst’s criteria. Because Net/FSE’s user interface is designed for security analysts by security analysts, new users can be trained to start searching and analyzing data in as little as one hour.

Search results in Net/FSE are stored in temporary relational database tables for responsive and powerful analytic capabilities. Users can work concurrently on the same search results for collaborative incident response and data analysis. An analyst can have any number of search results active at once and search history is automatically recorded on the server. These features are very useful when returning from a long weekend and trying to remember what tasks were in progress the week before.

Dig deep

Intrusion detection systems and network behavior analysis systems have never been known for their accuracy. Adding a SIM to provide event correlation capabilities on top of an IDS or NBA helps in identifying interesting alerts but no correlation algorithm can make up for bad alerts or incomplete data. However, no matter where an alert is coming from, additional context is always required. There is never sufficient information presented by an IDS, NBA or SIM to make a definitive conclusion about the importance of a given alert. Further analysis by a trained security expert is always needed to identify false positives and avoid unnecessary investigations into dead ends. Net/FSE is designed to be the ultimate tool for digging deep into network alerts for network forensic operations.

Until the recent past, "forensics" has been primarily used in the physical world and defined as "scientific analysis of physical evidence (as from a crime scene)." Forensics as it relates to an enterprise network can be applied as it does in the physical world in that it is the search and analysis of network data. The activity of performing forensics is based in the reality that you never know what you don't know.



Net/FSE's security-purposed search interface is designed for incident response and network forensics.

The only way to perform a definitive forensics investigation is to be sure you have collected all of the evidence. It is no different when it comes to a network forensics investigation. In most cases information is provided in the form of a lead or a tip (in network parlance - a network alert or tip from a user) and it is the job of the security analyst to turn find the "evidence" which in this case is the network event.

In network forensics a security analyst must:

1. Have a complete collection of all of the network data for an extended period of time.
2. Have real-time access to that data.
3. Be able to perform rapid, deep search and analysis on that data.

Without satisfying the above criteria it is impossible to definitively conclude that an exhaustive network forensics investigation has been performed. Although some commercial solutions begin to address these points, Net/

FSE is the only commercially supported product on the market that is purposely built to address these needs.

Response and recovery

In any enterprise where data on the network is mission critical, effective incident response and recovery is crucial to an organization's ability to avoid what could be potentially disastrous results from the impact of a network breach. Enterprise security solutions continue to improve their alerting capabilities through the use of SIMs, NBA systems and next generation IDS/IPS. Security analysts must be able to quickly assess, react and put definitive context around network alerts to determine the level of response resources necessary. Net/FSE is designed to streamline response and recovery operations by enabling security analysts to quickly and efficiently answer these key questions:

Net/FSE - The Network Forensic Search Engine

http://testdrive.packetanalytics.com/nfse/NetFSEClient.html

Net/FSE
Network Forensic Search Engine

Help | Open Search | Set Password | Bug Report | User Community | Log Out | Data Set

Search Results Analysis Timeline History

Actions: Refresh | Delete | Merge | Make Permanent

	Search ID	Note	Expires	Status	Flows	IDS	NetFlow	Total Records
<input type="checkbox"/>	s54	Possible data leak	2008-03-17 12:43:40.388	Complete	29722	110		29832
<input type="checkbox"/>	s6	IDS alerts 2000	2008-03-17 12:36:58.96	Complete		225		225
<input type="checkbox"/>	s435	None	2008-03-17 12:36:38.484	Complete	0	0	0	0

The Results tab gives an overview of the analyst's in-process investigations.

1. What other hosts have been involved with the incident?
2. How long has the event been going on? (typically the alert is not the first indication of the incident)
3. Is the activity that generated the alert still going on?
4. What ingress and egress points were used for the suspicious activity?

With the current suite of tools in place within most enterprise networks, it is virtually impossible to perform the above tasks in a reasonable amount of time. Network event data is scattered around different servers in the form of log files or only accessible via application specific ("point") consoles. This approach makes data analysis for response and recovery slow, cumbersome and error prone. Users must be trained on each point console, each of which has differing capabilities.

Net/FSE gives security analysts the ability to collect all of the enterprise's network event data and analyze it from a flexible and easily deployed web interface. If an enterprise is already collecting event data, Net/FSE can retrospectively index the data to minimize duplication of data. Net/FSE gives the security in-

cident response team access to network data in near real-time (as real-time as the devices that are sending data to Net/FSE). Most importantly Net/FSE gives security analysts the ability to perform deep search and analysis on all of the network data to arrive at definitive answers for response and recovery.

Conclusion

Enterprise networks have committed large portions of their IT budgets to a wide variety of network and security devices but have yet to fully leverage the valuable network event data streams generated by these systems. SIMs and log management solutions have partially addressed the needs of network security analysts but such systems are not built to provide analysis capabilities for alert analysis, in-depth network forensics or incident response. Net/FSE by Packet Analytics fills the gap in the network security market by bringing cost-effective, easy to use network event analysis capabilities to enterprise networks. Net/FSE adds value to an enterprise's existing tool suite and maximizes the value of these tools by making the organization's security practitioners more effective in their daily tasks.

Ben Uphoff is the Vice President of Research for Packet Analytics.



Security risks for mobile computing on public WLANs: hotspot registration

By Simon Ford

Wireless broadband internet access via hotspots is convenient for both the casual surfer and the internet-dependent teleworker. Unfortunately, current security technologies integrated into wireless LAN (WLAN) products offer insufficient protection here, and mobile users must be wary when accessing the central company network via a hotspot.

What is necessary is a security solution that protects the teleworkers place in all phases of connection construction on hotspots – without risky, foreboding configurations and without the help of users or administrators.

This article illuminates the effectiveness of VPN security mechanisms, data encryption, strong authentication and personal firewalls and shows how optimal protection can be achieved by dynamically integrating each of these technologies.

Risks in the WLAN

Each user can access public WLANs with correspondingly equipped terminals. They automatically obtain an IP address in the sense that they recognize the SSID (service set identifier) of the WLAN, thus putting themselves within range of the access points, and able to access permission onto the WLAN. Data security or protection of participating de-

vices from attacks is not guaranteed by the WLAN operator. Security is limited to monitoring authorized network access in order to eliminate misuse of the server administration. User identification serves solely for the acquisition and the accounting of time online.

However, how does it look regarding the protection of sensitive information during data transmission? How can the PC optimally seal itself off from attacks from the WLAN and the Internet?

Because the actual security risk on the hotspot originates from having to register with the operator outside the protected area of a VPN, as a rule it has to take place by means of the browser. During this timeframe, the terminal device is unprotected. This stands in opposition to the company's security policy that prohibits direct surfing on the Internet and that only permits certain protocols.

Basically, VPN mechanisms and data encryption serve to protect confidentiality. The corresponding security standards are IPSec tunneling and AES encryption for data, and X.509 v3 for access protection. Additional security mechanisms, such as certificates in a PKI (public key infrastructure) or onetime password tokens complement/replace the usual user-ID and password. A personal firewall offers the required protective mechanisms against attacks from the Internet and from the public WLAN. Here, stateful packet inspection is critical. If this is not provided, it is not advised to use a hotspot for mobile computing.

VPN client and external personal firewall

For a VPN solution with a separately installed firewall, the ports for http/https data traffic to the personal firewall must be activated during hotspot registration. This can take place in three different ways:

1. The firewall rules for http/https are firmly pre-configured in order to guarantee the functionality with the desired hotspots.

2. The configuration allows that the ports are opened for http/https as needed for a certain time window (e.g. two minutes).
3. The user has administration rights and independently changes the firewall rules.

In all three cases there exists the risk that the user may surf outside of the secure VPN tunnel on the Internet and encounter destructive software such as viruses, worms or Trojans. Temporarily opening the firewall creates the danger of deliberate misuse by the user on the basis of multiple actuations of the time window. If the personal firewall fundamentally permits no communication outside of the configuration, then the user has to activate the corresponding firewall rules for the duration of registration on the hotspot. This requirements-based opening of the personal firewall involves the greatest risk of misconfigurations. The user must have a firm grasp of the exact changes being made and the exact environment in which they are made. Employee security awareness and technical know-how determine the security level quality.

EMPLOYEE SECURITY AWARENESS AND TECHNICAL KNOW-HOW DETERMINE THE SECURITY LEVEL QUALITY

A large security risk also exists when user data (user ID/password) is spied out externally on the hotspot during the registration process. With the aid of his notebook a hacker can simulate both the hotspot and the WLAN SSIDs. If a user then registers on a hotspot, he does not land at the access point of the provider, but rather on the notebook of the hacker. By means of the previously mirrored access point web pages, the user still assumes that he is authenticated on the hotspot, when in reality he is on the notebook of the hacker and his personal registration data is now exposed.

Providers always attempt to protect the hotspot registration pages through SSL processing (https), but that does not always succeed. For example, a user who arrives at a manipulated hotspot obtains the following report from the browser: A problem exists with the security certificate on the web site. In the background of this report, the attacker has only recreated the hotspot registration page and does not

use the original certificate. For the lay person, this may not be recognizable at first glance, and it is incumbent to him to decide whether or not he should trust the certificate. In order not to place a user in the position of making this decision, the hotspot registration should flow transparently before construction of the VPN. A solution that has proven itself in practice is the so-called registration script that takes over the transmission of registration and the inspection of the certificate at the hotspot.

The requirements for the functionality of a personal firewall with mobile computing on WLANs are multilayered. They also apply to the critical phases during the registration and sign-off process on the hotspot. Requirements must be known at the earliest possible time and should be in place from system start. They also must remain when no VPN connection exists or has been deactivated. Furthermore, the user should be safeguarded against arbitrarily reconfiguring or completely shutting off the personal firewall.

VPN client with integrated personal firewall

The dilemma of system requirements may be resolved by a VPN solution with a client-integrated personal firewall. The advantage of the integrated variant is that a personal firewall and VPN client are functionally linked to one another. In a quasi-teamwork fashion, the existing firewall rule statements are dynamically activated in dependence on the network environment. Fundamentally, three situations may be differentiated:

1. Known networks
2. Unknown networks
3. VPN networks

Automatic recognition of the network takes place by validating different network factors. In friendly networks, permissive firewall rules apply as they do in public environments like the hotspot. The personal firewall must work with intelligent mechanisms that guarantee a

secure activation of network access via the browser, as well as a secure registration on the hotspot. The user chooses the menu point "hotspot registration" in the welcome area of a public WLAN. Subsequently, the VPN client automatically searches the hotspot and opens the web site for registration in a standard browser. For example, after successful entry of access data and activation by the operator, the VPN connection can connect to the company headquarters and communicate as securely as it would in an office.

In this manner, the PC is accessible in the WLAN in no time, and there are ports dynamically assigned for http/https for registration and logging off the hotspot. During this time, only data traffic is possible with the operator's hotspot server. Unnecessary data packets are refused. In this way, it is guaranteed that a public WLAN can use the VPN connection at the central data network and no direct internet access can take place.

A PREREQUISITE FOR SECURE REMOTE ACCESS IN WLANS IS END-TO-END SECURITY, WITH DYNAMIC INTERLOCKING SECURITY TECHNOLOGY

Inspection of security-relevant parameters

An additional important component of the implementation of company wide security directives for mobile computing on hotspots is central management of client software. With central security management, the administrator also fundamentally determines the client's firewall rules. It can enforce adherence in which the user allows no on-site possibility of an intended or unintended change.

Additionally, further security-relevant parameters such as the status of virus protection programs, operating system patch status, and software release of the VPN client must be inspected upon connection to the company network. Access to the productive network is only authorized after the clearance of all security risks.

Bottom line

A prerequisite for secure remote access in WLANs is end-to-end security, with dynamic interlocking security technology. The use of a VPN client with an integrated, intelligent personal firewall and strong user authentication is state of the art in this scenario. The firewall rules must automatically adapt to registering onto and logging off of the hotspot, and they must be inspected within the framework of an integrated endpoint security system with each connection.

Only in this way can administrators and users be consistently sure that they are securely sealing off terminal devices and data, and signing off the company network.

Simon Ford currently serves as International Director for NCP Engineering. He has been working with security technologies for more than 20 years. Simon lives in Nuremberg, Germany.



secureworld expo

2008

is your world secure?

Register Today
for the Security Conference
Built by Security Leaders
Like You.



BOSTON
MARCH 26-27



HOUSTON
APRIL 23-24



ATLANTA
APRIL 29-30



PHILADELPHIA
MAY 7-8



CHICAGO
MAY 21-22



BAY AREA
SEPTEMBER 10-11



CLEVELAND
SEPTEMBER 24-25



DETROIT
November 4-5




SEATTLE
OCTOBER 29-30



DALLAS
NOVEMBER 12-13

(IN)SECURE Magazine Special Discount -
Register With Code **NSECMW8** and Save!

REGISTER ONLINE: www.secureworldexpo.com



Software spotlight

Kiwi CatTools

<http://www.net-security.org/software.php?id=317>

CatTools is a utility for managing routers switches and other network devices. Provides automated configuration backups, password changes and scripted configuration commands. Multiple device types supported.

NuFW

<http://www.net-security.org/software.php?id=526>

NuFW is an "authenticating gateway". This means it requires authentication for any connections to be forwarded through the gateway.

NetShred X

<http://www.net-security.org/software.php?id=621>

Easy to use because it runs automatically, permanent because it shreds - not just deletes browser cache, history files, email trash and more so the data can't be recovered.

Shimo

<http://www.net-security.org/software.php?id=671>

The initial reason for the development of Shimo was the lousy software implementation of the Cisco VPNClient for Mac OS. Shimo wants to do it the MacOS way: That means to reduce the interface to the important features and integrate it right into the user interface of the OS. Thus the interaction with this piece of software is absolutely intuitive and self explanatory.



Black Hat Europe 2008 Briefings & Training

In April security researchers gathered in Amsterdam for one of the most important security gatherings on the planet - Black Hat.

Researchers from all over the world shared their findings with a knowledgeable crowd.

Here's a list of some of the very interesting talks. If you find them interesting you should consider attending the event next year.

More information on the event can be found at www.blackhat.com



Bad Sushi - Beating Phishers at Their Own Game

- Nitesh Dhanjani, Senior Manager and Leader of Application Security Services, Ernst & Young LLP
- Billy Rios, Microsoft

This talk exposed tactics and tools used by phishers, show how easy it is to hack into servers that are used to perform phishing, demonstrate how easy it is so follow a phisher's trail to find out how they share information on US Citizens including SSNs, bank account numbers, credit card numbers, ATM PINs, you name it.

Phishers usually setup their sites on servers they have compromised. In other words, the phishers have already done the hard work and it is easy to gain access to these servers. Due to the sheer volume of sites that need to be setup to perform a successful phish, phishers tend to be sloppy and leave traces everywhere.

We've interviewed the authors for Help Net Security, read the interview here: www.net-security.org/article.php?id=1110



Billy Rios

New Viral Threats of PDF Language

- Eric Filiol, Head Scientist Officer, Virology and Cryptology Lab, French Signals Academy

Adobe Portable Document Format has become the most widespread and used document description format throughout the world. It is also a true programming language of its own, strongly dedicated to document creation and manipulation which has accumulated a lot

of powerful programming features from version to version. Until now, no real, exploratory security analysis of the PDF and of its programming power with respect to malware attacks has been conducted.

This presentation presented an in-depth security analysis of the PDF programming features and capabilities, independently from any vulnerability.

Intercepting Mobile Phone/GSM Traffic

- David Hulton, Security Researcher, Pico Computing, Inc.
- Steve, Security Researcher, Pico Computing, Inc.

This talk was about GSM security. The authors explained the security, technology and

protocols of a GSM network and presented a solution to build a GSM scanner for 900 USD.

The second part of the talk unraveled a practical solution to crack the GSM encryption A5/1.



Steve and David Hulton

Developments in Cisco IOS Forensics

- Felix "FX" Lindner, Reurity Labs GmbH

Attacks on network infrastructure are not a new field. However, the increasing default protections in common operating systems, platforms and development environments in-

crease interest in the less protected infrastructure sector. Today, performing in-depth crash analysis or digital forensics is almost impossible on the most widely used routing platform. This talk showed new developments in this sector and queried the audience for their experience, input and wishes.

The Fundamentals of Physical Security

- Deviant Ollam, The Open Organization of Lockpickers

As a veteran trainer in the field of locks and physical security, Deviant has given innumerable presentations geared towards security professionals who are in charge of overseeing facilities. Those who attended this session left with a full awareness of how to best protect buildings and grounds from unauthorized ac-

cess. Discussion as well as direct example have been used to demonstrate the grave failings of low-grade hardware... much of which was opened by audience members with no prior training.

Ollam also covered what features to look for in locks and safes as well as how to invest in systems that are easiest to manage in large environments.

Hacking Second Life

- Michael Thumann, CSO, ERNW GmbH

Beyond being an online game SecondLife is a growing marketplace for big companies where lot of money is made. Living and acting in a virtual world gives the people the opportunity to do things they would never do in real life. Therefore it is not surprising that SecondLife has increasingly attracted real world hackers.

The talk covered the basic architecture of SecondLife and pointed out the possible attack vectors against SecondLife itself. It also demonstrated hacks from the inside of SecondLife against real-life systems in the internet.

Michael did a video for Help Net Security on this topic, you can view it here:

www.net-security.org/article.php?id=1125



The John Henry challenge

Iron Chef Black Hat: John Henry Challenge

- Jacob West, Fortify Software
- Pravir Chandra, Principal Consultant, Cigital
- Brian Chess, Chief Scientist, Fortify Software
- Sean Fay, Lead Engineer, Fortify Software

In the spirit of the Food Network's cult favorite show, Iron Chef, the Chairman revealed the surprise ingredient (the code), and then let the challenger and the 'Iron Hacker' face off in a frenetic security battle.

The guest panel judged the tools created and techniques used to determine which who will be victorious and who will be vanquished.

The testers had only one hour to complete their challenge and they were restricted to their respective choice of bug-finding techniques: One team used automated tools they themselves have built, while the other flexed their security muscles through manual code review.



After serious troubles with software systems in Japanese corporations like last year's temporary transaction stop at the Tokyo Stock Exchange, preventive methods like Software Configuration Management (SCM) have caught a lot of attention in Japan. SCM is a method to track, control and document processes in software development.

There are several dimensions and management objectives to SCM. This makes SCM an appealing choice, especially for managers who control budgets and who are accountable for problems in the software. At first glance, one might not consider SCM and security to be directly connected. Yet, SCM has an extremely close connection to security, particularly as a solution to prevent the inclusion of malicious code during the development process.

In this article, I will analyze which systems and tools Japanese software houses employ for SCM from a security angle.

The software development process

Before explaining how SCM can be effective for security, let's have a look at what SCM

controls specifically. In a simple definition, SCM is just considered to be the same as Version Control for source code. This opinion is widespread in Japan. Generally, this is not wrong, if you consider Version Control as the SCM of the implementation phase. The more modern definition of SCM is, however, that it does not just record the update history of the source code, but it also works as a knowledge base for other documentation that is created during the whole development process.

I have explained SCM as "a method to track, control and document the processes in software development," but what documentation occurs in each process?

In the basic model of software development, there are five phases: Requirements -> Design -> Implementation -> Test -> Deployment.

In each step, a document on the work results is created. For example, in the requirements stage, this is a requirements or bug report.

In big projects on open systems, third party source code and binaries must be coordinated with the developer's own code. It is necessary to ensure compatibility between all parts of the system to make sure it runs without troubles. On a rollback, the SCM system must be able to return all components to a certain point in time to avoid duplications and leaks. SCM does not apply only to the source code itself, but also to configuration files, for example for servers and network devices. Specification sheets and manuals should also not be forgotten.

SCM and security

Software is developed on different scales, from small projects run by just one person to huge projects with different teams and supervisors on each phase. Regardless of size, the process itself does not change much. The SCM way of thinking can be applied regardless of scale. Through SCM, the whole software development process is governed. Thus, if for example a supervisor in the implementation phase plants a backdoor, it can be found in the verification phase. This makes it difficult to slip malicious code fragments into software. Overly paranoid? Maybe, maybe not, as even the US Navy had to learn when a subcontractor inserted malicious code into a collision prevention system for submarines (tinyurl.com/5db27q). This is where SCM bolsters security significantly.

"Whole system" configuration management

In case of a system rollback, source code and setting files as well as documentation can be simultaneously restored. This prevents overlaps and leaks in the development process after the rollback. Especially leaks are a potential security hole, which makes SCM an important security tool in the development process.

The dawn of SCM in Japan

Recently, the true meaning of SCM has gotten more attention in Japan for two main reasons.

One is that software problems are caused by the so-called the "2007 problem" in Japan. The other is the requirement for stricter internal controls following the enactment of the Japanese version of the Sarbanes-Oxley Law.

The 2007 problem

In Japan, there have been a number of high-profile cases in the past year that gained widespread publicity and caused significant economic and public relations damage. Below is a quick list of some of the bigger ones:

- April 2007: Risona Bank, Saitama: ATM network failure (program failure).
- June 2007: Michinoku Bank: 341 credit record registration mistakes (program failure).
- October 2007: JR East: 4378 ticket gates out of order (insufficient testing).
- February 2008: Tokyo Stock Exchange: Temporary transaction stop at the TOPIX futures market (insufficient testing).
- March 2008: Tokyo Stock Exchange: Trading stop for two stocks (program failure).

Wondering what the unifying theme is between all of these high profile software problems? In all cases, there were indications that the failures could be attributed to the retirement of skilled key employees. The baby boomer generation that built the Japanese post-war economy is currently retiring en masse, which means that Japanese corporations are losing their technical skills. IT departments of large corporations are no exception, and especially sectors where small errors can cause large damage, such as in banking and transportation, the effect has been noticeable.

The problem is also deeply connected to the Japanese system of longtime employment. Contrary to the American system of hire-and-fire, Japanese companies have traditionally hired their employees until retirement, under the assumption of steady growth. Since specialists often have their position in a department for many years, they can amass a lot of experience, without much need for documentation or training of new employees or contractors. The negative side is that this leads to problems if the employee retires. Proper use of documentation and SCM can avoid these problems. Another reason why the situation was disregarded until these problems

occurred was that previously Japanese companies did not view software resources and developers as important assets, and did not make adequate investments.

The financial instruments and exchange law, also known as “J-SOX”

The Japanese version of the Sarbanes-Oxley Law (Financial Instruments and Exchange Act, or J-SOX for short) has also had a strong influence. With the enactment of J-SOX, internal control has risen to prominence in Japan.

There is a wide range of issues summarized under internal control, and many IT systems play a crucial role:

- Segregation of duties and administrative control without unlawful conduct or fraud.
- Maintenance of business policies.
- Administration of the development and operation of IT systems.
- Regular system monitoring.
- Strengthening security.
- Building / System / Information Access controls.
- Retrieval and archiving of audit logs.
- Establishment of a system to vindicate management.
- Maintaining a risk management framework.
- Education and oversight of employees and contractors.

In many of these requirements, new software has to be introduced or developed to meet the demands of internal control. SCM can be part of internal controls where software development is involved.

The spread of SCM in Japan

In the Japanese software development scene, there are few companies which established SCM systems. Version Control, on the other hand, is widespread, mainly because of its obvious benefits for developers. Version Control also has a whole universe of tools and utilities to help developers code effectively. Conversely, one reason why SCM has not reached maturity in the Japanese market is that there is a clear paucity of tools, and those that exist don't match the development proc-

ess well. Furthermore, from the business side, it is also difficult to measure the return of investment in SCM, which makes companies reluctant to pour money and manpower into it.

Popular SCM and version control software in Japan

As for Version Control, there are a number of popular programs, especially CVS and Subversion. As open source software, they are cost-effective to set up, and there is plenty of documentation available in Japanese. Similar to other non-English-speaking countries, software has a hard time to become popular if it is not available in a localized version. A full-feature SCM-software is not yet available in a Japanese localization. The software that can be considered to come closest are bug tracking systems with Version Control integration. Popular are the Open Source solutions Trac (trac.edgewall.org) and Mantis (www.mantisbt.org). On the side of proprietary software, there is the Microsoft Visual Studio 2005 Team Foundation Server (tinyurl.com/5k59fa) and Borland StarTeam (tinyurl.com/mp3ff).

Conclusion

The newspapers in Japan are filled with security stories about hackers and other problems. But in Japan, having the Tokyo Stock Exchange come to a screeching halt due solely to a software technical problem caught a lot of people's attention. There is clearly a heightened awareness around solving problems in the code before it becomes a problem in the real world. Japan will increasingly turn to SCM as these types of well-known software problems caused by poor documentation and the lifetime employment retirees phenomenon continue.

If companies wanted to avoid responsibility before, legally, J-SOX is now in effect for roughly 3,800 companies listed in Japan, along with their foreign subsidiaries. It is clear that IT controls are a central focus for J-SOX, so regulations and, most importantly, punishments now force the issue.

Kohei Matsushita is General Manager of the Online Business of Plat'Home Co., Ltd. in Akihabara, Tokyo, Japan. Building on his experiences building a development system that takes security into account from the system design levels, he is now responsible for a safe around-the clock user experience for the Plat'Home e-commerce website.

Windows log forensics: did you cover your tracks?

By Rob Faber



The information that is stored in logs is very useful, particularly when it comes to gathering forensic evidence related to intrusive actions, fraudulent behavior or malicious attacks. Accordingly, log-files are important sources of forensic information because they usually connect a certain event to a particular point in time. This time we will delve a bit more deeply, and investigate Microsoft Windows Event logs.

Logs can be generated and found within fire-wall solutions, web servers, Windows servers and clients, and many other devices and services. On its own, a Windows computer can generate an overwhelming number of logs and security events, and the information stored in them can be of great interest, particularly when it comes to gathering forensic evidence related to intrusive actions, fraudulent behavior or malicious attacks. That makes log-files important sources of forensic information because they usually connect a certain event to a particular point in time.

A thorough analysis of this information is necessary to present a case and, in a forensic scenario where there is a need to investigate

all of this data, you will probably need tools with which to conduct this analysis in a convenient way.

Yet, logs are not only important in the case of a computer forensic investigation. They are also a useful source of information about platform behavior and security for administrators.

In this article, I will concentrate on Windows Forensics and event logs, and suggest some methods and tools you can use to help you carry out your own investigations. I will demonstrate how events can be abstracted from a number of system logs on a Windows system and can then be used to get a better picture of what has taken place.

Logging within Windows

The Windows operating system comes with a complex architecture with which to handle events such as logging on. It is also possible to trigger the writing of an action to the log when a specific type of event occurs.

Firstly, the System and Application logs can be used in a number of ways by both applications and the Windows operating system. However, there are conventions about how to write specific events to the log.

Secondly, there is a special type of logging: security logging, during which security related events are written to the separate security log. This log can only be directly written to by the Local Security Authority Subsystem Service, or LSASS.

The LSASS is a process that runs as part of the Windows operating system, and is re-

sponsible for enforcing its security policy. Amongst the LSASS's tasks are, for example, the verification of users logging on to the Windows system, handling password changes, creating access tokens and consequently writing entries to the Windows Security Log.

The local system's audit policy controls the auditing of a particular type of security event. So, if there isn't a policy to audit a specific type of event, this won't be written to the security log! This policy is - by default - not activated until Windows Server 2003. Accordingly, you have to configure some policies yourself to turn on security auditing and logging, and to ensure that it is working in the way you want it to. The audit policy is part of the security policy. The LSASS is part of the local system, and so is configured with the Local Security Policy editor (or - of course - policies from the Windows domain).

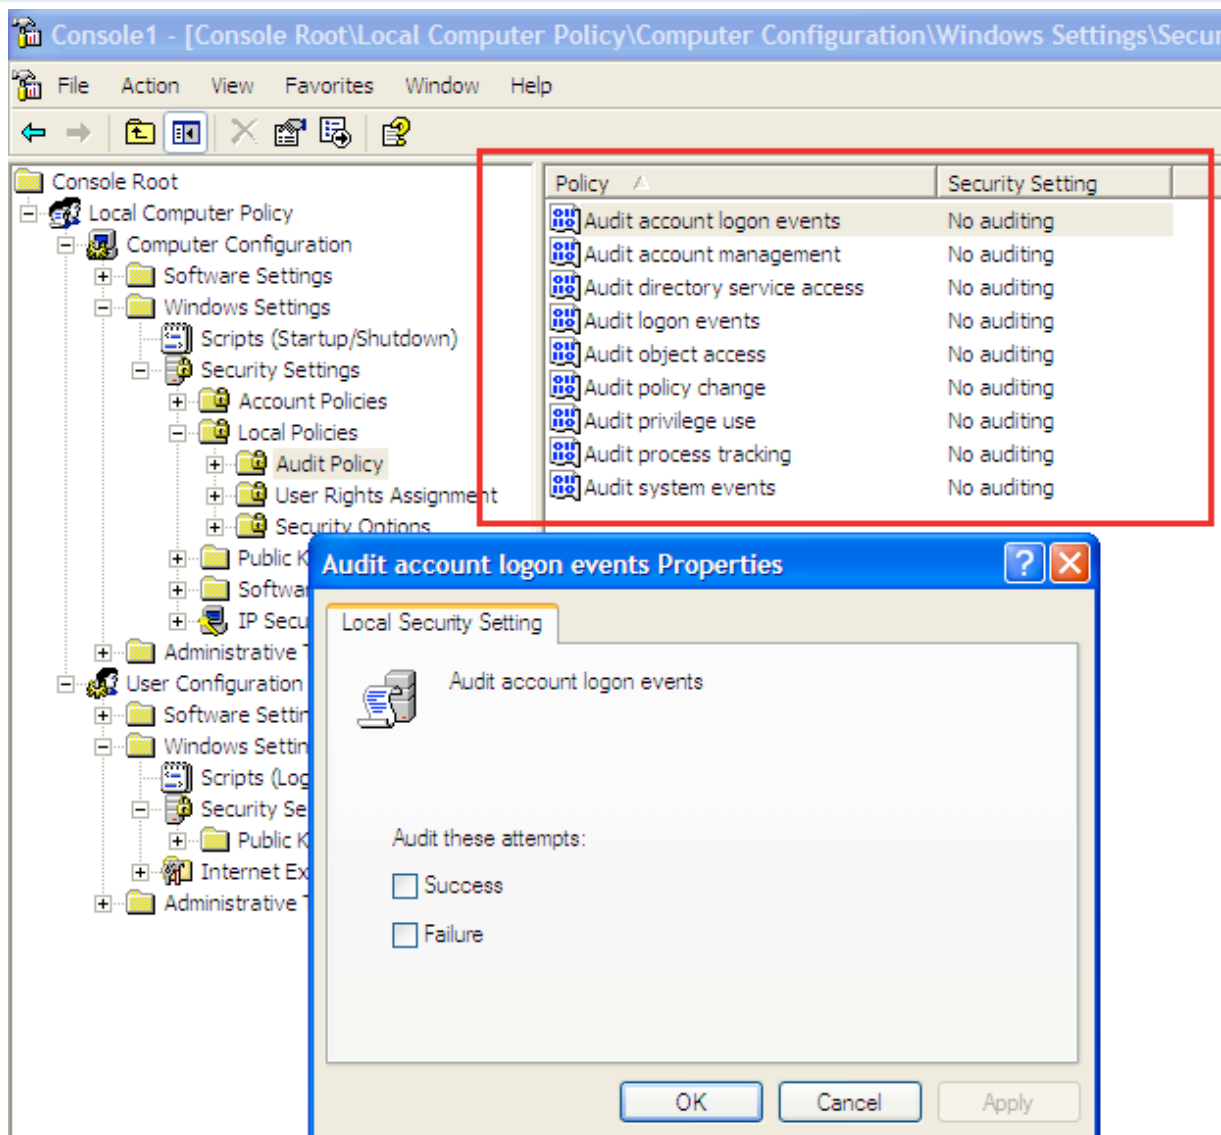


Figure 1. Local Audit Policy

Security Reference Monitor (SRM) and triggering events

In operating systems' architectures, a so called reference monitor can be found. The reference monitor's task is to control all access that subjects (for example, a user or service) have to objects (for example a file or folder). The purpose of this is both to ensure that the subjects have the necessary access rights, and to protect the objects from unauthorized access and unauthorized modification. Windows' 3.x and 9.x operating systems didn't contain such a reference monitor, whereas the (latest) operating systems such as Windows XP, Server 2003 and Vista do.

The reference monitor is not a product of Microsoft or Windows itself, but is a commonly recognized function in all security systems. Linux and other operating systems have a similar concept implemented within them (see Common Criteria and CISSP All-in-One Exam Guide, by Shon Harris for more information). To conclude with the theory, systems evaluated at EAL 5 and above (TCSEC B2) by the Common Criteria must enforce the reference monitor concept.

The LSASS communicates with the Security Reference Monitor (SRM) in Windows, sending it messages to inform it of the auditing policy at system initialization time, and when that policy changes. The LSASS is also responsible for receiving the audit records that are generated on the basis of the audit events from the SRM, editing the records, and - here it is - sending them to the Windows Event Logger. It is beyond the scope of this article to explain this in great depth, but for those of you who are interested in getting more detailed background information, a useful resource about this subject can be found in Microsoft Windows Internals, fourth edition by Mark Russinovich, chapter 8.

Types of logs and log formats on a Windows System

Logging on to a Windows System can occur in different ways as stated before. Log files concentrated on three of these (before Windows Vista): Application logging, System logging and Security events. The application logging stores information on individual applications

that are running on the system. The system logs contain operating system event details. Security logs store data from logged on users, together with other security related details about the system. Vista changed all this, but I will return to that later on in this article. So far as forensic evidence is concerned, security logging is really the most interesting of the presented log types. Accordingly, I will concentrate further on this type of event log, but have to stress that the other logs can also contain useful information and should therefore not be overlooked.

Log files will often be flat text files, which contain and present each log entry that is written to it as a single line or row. Depending on the situation, a log file can contain hundreds of these lines and the file can be up to several megabytes in size. Of course, all of these files can be a serious concern when it comes to the overall performance of the system, particularly if you don't get the balance right between collecting enough information to use in an investigation when necessary, but not so much that it puts too much pressure on the system, which then becomes unresponsive or simply crashes.

The standard way of viewing event logs is to use the Microsoft Event Viewer. Event Viewer can be invoked by typing [eventvwr] from the command prompt on Windows systems. Event Viewer uses the Management console (MMC) interface to display the information that is stored in the different log-files, and you can use it to view log files from a local or remote system. The files can be found with the .evt extension on the system. This way, it is possible to connect to a remote machine and import the files for further investigation.

The default location of the security event log within Windows XP is:

```
%SystemRoot%\System32\config\SecEvent.Evt
```

This entry can be changed if you want to relocate the files to a larger disk or partition. To modify the location of the Event Viewer log-files, use the Registry Editor. For the Security log the location is:

```
System\CurrentControlSet\Services\EventLog\Security then double-click the FILE value.
```

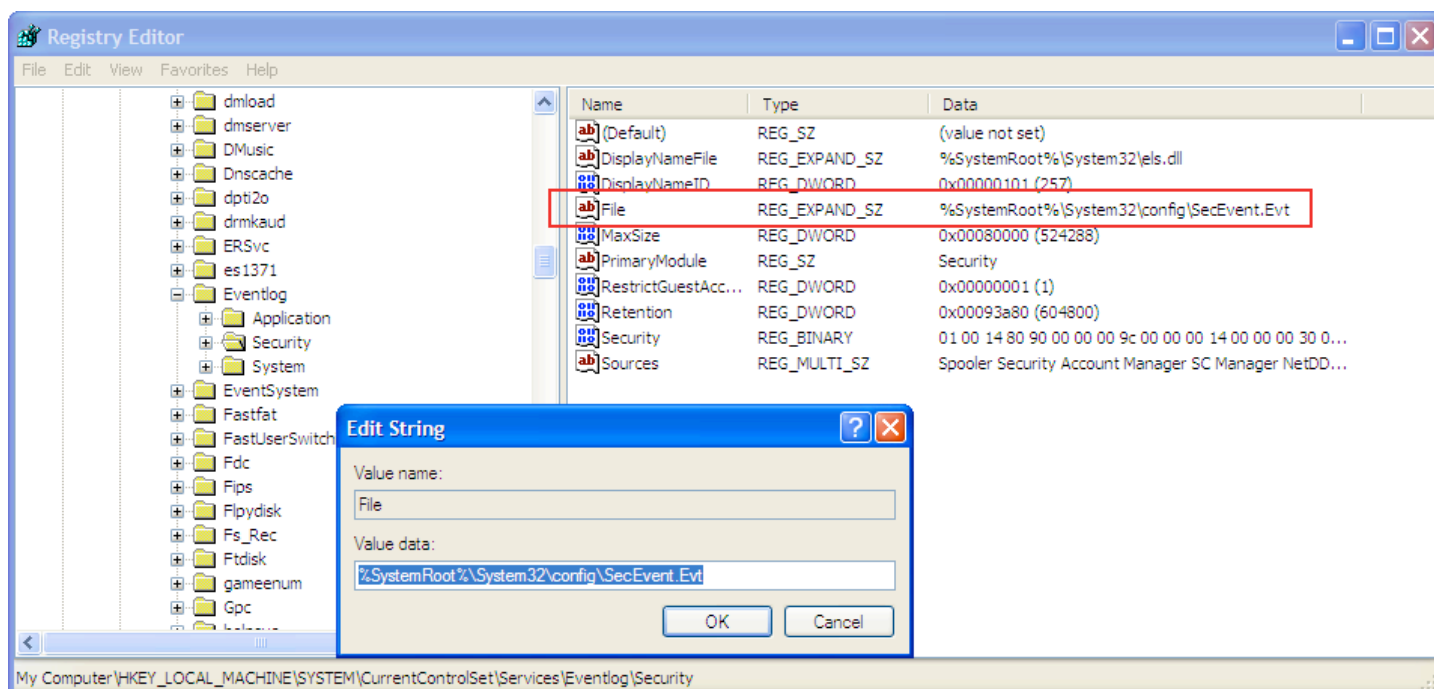


Figure 2. Default eventlog location

Interpreting Event IDs

The types of event that can be found in a column within the event viewer are probably the most important in the event log. The log contains specific numbers, and these correspond to a specific type of security event. Microsoft has assigned a unique number to every event that can be recorded in a Security log. When viewing the Security log, simply looking at the Event ID number can tell you exactly what type of event has occurred.

The column Computer, lists the computer that recorded the event. This can, however, be misleading because this doesn't always represent the computer that caused the event. Indeed, it isn't even the computer that a possible intruder was using when the incident took place because, in most cases, it is simply the system that generated the log entry.

There are a few other matters that you should also be aware of. The times presented in the log and viewed with the event viewer MMC, are time-zone related. In other words, time values are interpreted by relating this to the time-zone currently set on the system. When you have a timeline on your system, you must be aware of this behavior when investigating the data.

You should also bear in mind the fact that newly introduced versions of the Windows op-

erating system can sometimes have amended Event IDs, and other IDs are often added as time goes on. As a result, you may have some difficulty in interpreting the logs if you are using an older OS to view or interpret the newer version's log files. By way of an example, modifications were made between Windows XP and Windows Server 2003, and within the latter a new field has been added to the log: the IP address of the system that made a remote connection. Windows XP will not, however, recognize this new entry! Consequently, important information presented in the log can be overlooked if you are not completely up to date with the changes and differences between the OS versions.

The lesson is: if you don't want to miss important evidence of incidents, always view logs on a system that has an equivalent or newer OS version installed. This way you won't get misleading information or no information at all!

If you want more information about this issue, and the differences between versions and how to deal with the information presented, some great resources are:

- www.eventid.net
- tinyurl.com/2zyvun

Here you will find a comprehensive overview of the different OSs and the Event IDs utilized.

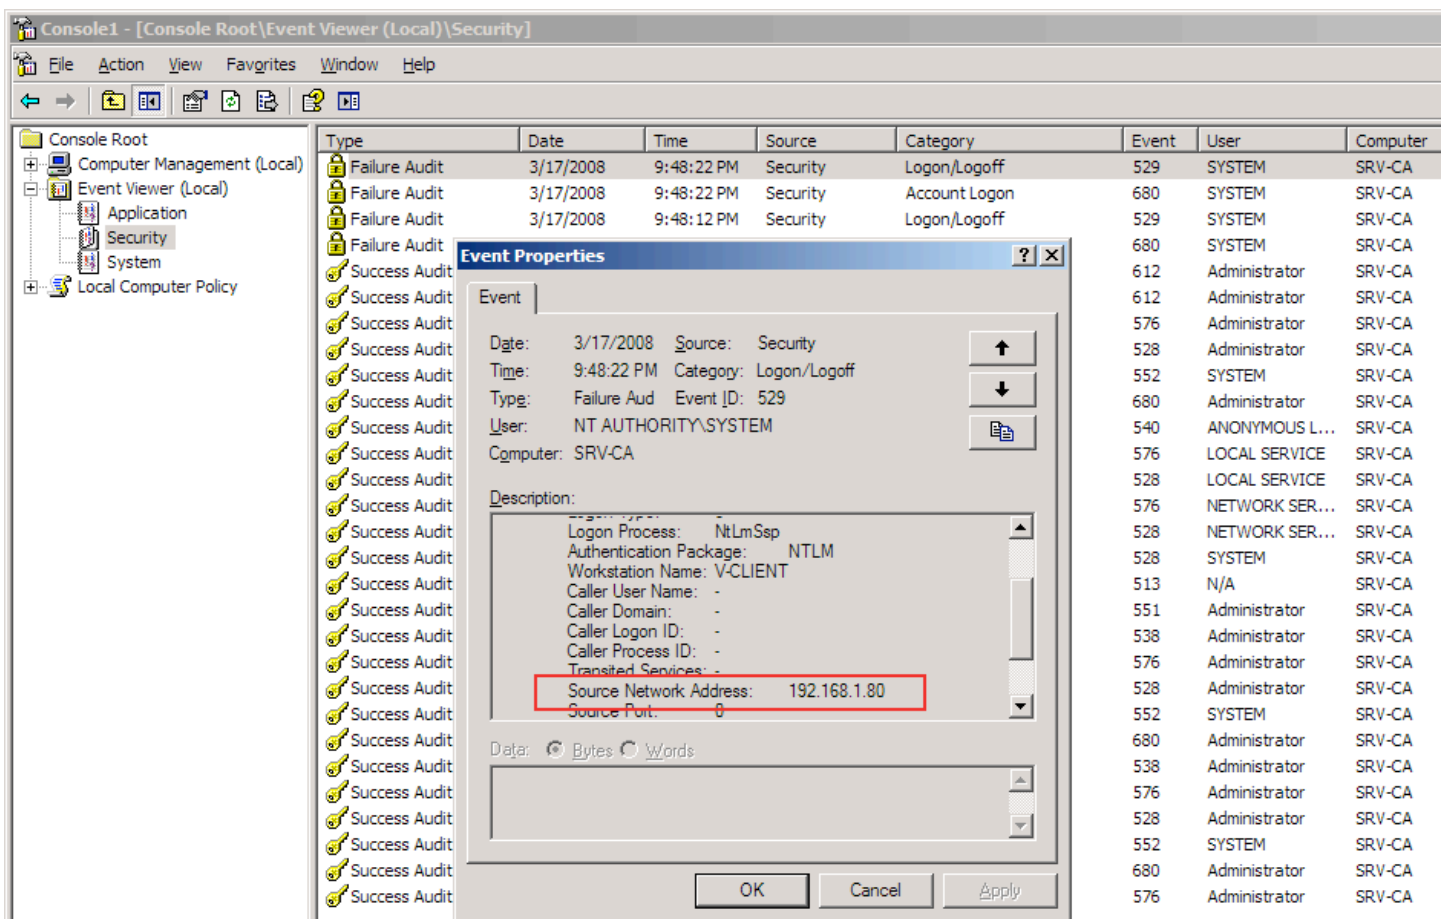


Figure 3. The IP-address presented in the Server 2003 eventlog

Log codes within Event IDs

The Event IDs 528 and 540 represent a successful log-on and event ID 538 a log-off and are therefore interesting from a forensic point of view. Other events in this category identify different reasons for a log-on failure.

However, just knowing about a successful or failed log-on attempt doesn't tell the whole story. Windows has many different ways in which you can log on to a system. All of these variations are logged separately. Logging on can occur interactively, when sitting behind the system locally, or over a network by performing a drive mapping, and through terminal services or connecting IIS.

Within all of these events, the Log-on Type code can be found, which reveals the type of log-on or log-off that triggered the event. This basically tells us how the authorization happened (success or failure) and the path that was used (interactive, remote and so on). Briefly, these types of event are:

- Type 2: Interactive
- Type 3: Network

- Type 4: Batch
- Type 5: Service
- Type 7: Unlock
- Type 8: Network Cleartext
- Type 9: New Credentials
- Type 10: Remote Interactive
- Type 11: Cached Interactive

The log-on/log-off category of the Windows security log gives you the ability to monitor all attempts made to access the local computer. This article will not examine every log type in detail, although there some resources available that deal with this topic.

Using Microsoft Log Parser on logs

Microsoft Log Parser is a free command line tool, and version 2.2 is the most recent. It is extremely useful in situations when (large) log files have to be processed, particularly in cases of security incidents. The Log Parser uses SQL queries to process, search and sort log files. It has a simple, yet effective, engine and can be used to read any text based file, file system objects and the Windows registry. It can also handle database formats.

Event Viewer (Local) Security 23 event(s)

Type	Date	Time	Source	Category	Event	User	Computer
Success Audit	3/17/2008	8:56:41 PM	Security	Logon/Logoff	538	User001	XP001
Success Audit	3/17/2008	8:56:32 PM	Security	Logon/Logoff	528	User001	XP001
Success Audit	3/17/2008	8:56:32 PM	Security	Account Logon	680	SYSTEM	XP001
Failure Audit	3/17/2008	8:56:30 PM	Security	Logon/Logoff	529	SYSTEM	XP001
Failure Audit	3/17/2008	8:56:30 PM	Security	Account Logon	680	SYSTEM	XP001
Failure Audit	3/17/2008	8:56:30 PM	Security	Logon/Logoff	529	SYSTEM	XP001
Failure Audit	3/17/2008	8:56:30 PM	Security	Account Logon	680	SYSTEM	XP001
Failure Audit	3/17/2008	8:56:26 PM	Security	Logon/Logoff	529	SYSTEM	XP001
Failure Audit	3/17/2008	8:56:26 PM	Security	Account Logon	680	SYSTEM	XP001
Failure Audit	3/17/2008	8:56:26 PM	Security	Logon/Logoff	529	SYSTEM	XP001
Failure Audit	3/17/2008	8:56:26 PM	Security	Account Logon	680	SYSTEM	XP001
Success Audit	3/17/2008	8:56:19 PM	Security	Logon/Logoff	551	User001	XP001
Success Audit	3/17/2008	8:44:58 PM	Security	Logon/Logoff	528	User001	XP001
Success Audit	3/17/2008	8:44:58 PM	Security	Account Logon	680	SYSTEM	XP001
Success Audit	3/17/2008	8:44:37 PM	Security	Logon/Logoff	551	User001	XP001
Success Audit	3/17/2008	8:43:52 PM	Security	Object Access	562	User001	XP001
Success Audit	3/17/2008	8:43:52 PM	Security	Object Access	562	User001	XP001
Success Audit	3/17/2008	8:43:52 PM	Security	Object Access	560	User001	XP001
Success Audit	3/17/2008	8:43:52 PM	Security	Object Access	560	User001	XP001
Success Audit	3/17/2008	8:43:46 PM	Security	Object Access	562	User001	XP001
Success Audit	3/17/2008	8:43:46 PM	Security	Object Access	562	User001	XP001
Success Audit	3/17/2008	8:43:46 PM	Security	Object Access	560	User001	XP001
Success Audit	3/17/2008	8:43:46 PM	Security	Object Access	560	User001	XP001

Figure 4a. Logged events in securitylog

Event Properties

Event

Date: 3/17/2008 Source: Security
 Time: 8:56:30 PM Category: Logon/Logoff
 Type: Failure Aud **Event ID: 529**
 User: NT AUTHORITY\SYSTEM
 Computer: XP001

Description:

Logon Failure:
 Reason: Unknown user name or bad password
 User Name: User001
 Domain: XP001
Logon Type: 2
 Logon Process: User32
 Authentication Package: Negotiate
 Workstation Name: XP001

For more information, see Help and Support Center at

Data: Bytes Words

OK Cancel Apply

Figure 4b. Logon failure and the logon type details

The Log Parser can format the output into a variety of formats, and can also send output directly to a log server, a database, and the DATAGRID. This final option presents you with an MS Excel like worksheet, containing the information sought or entered at the command line.

You can download Log Parser from the Microsoft website, and once this has been done double-click the setup file, LogParser.msi. The Setup Wizard will then start and Log Parser will be installed on your system.

A bit SQL

Log Parser queries are based on standard SQL queries. For those of you who are not familiar with these, they are a very powerful resource! They enable us to quickly select the fields that we need to display for further investigation, and just as useful, ignore the entries that are less important. These queries also allow the aggregation of log information, such

as by counting the number of specific events that have occurred in the log, or by searching for a special type of event. This way, a vast amount of data can be processed and the information being sought can be quickly located, and events and data from different sources can be correlated.

The basic parts of a SQL query are the SELECT, FROM, and the optional WHERE statement. SELECT indicates the fields to return, FROM indicates the data source to use in the query, and WHERE stands for any conditions, such as the rows to be included in the result. This way, a very simple query will look like this:

```
logparser -o:DATAGRID "SELECT * FROM security"
```

The "*" in this case means that the query should return every field in the row. In the following picture you can see the result of that query, presenting all the rows and fields from the security log.

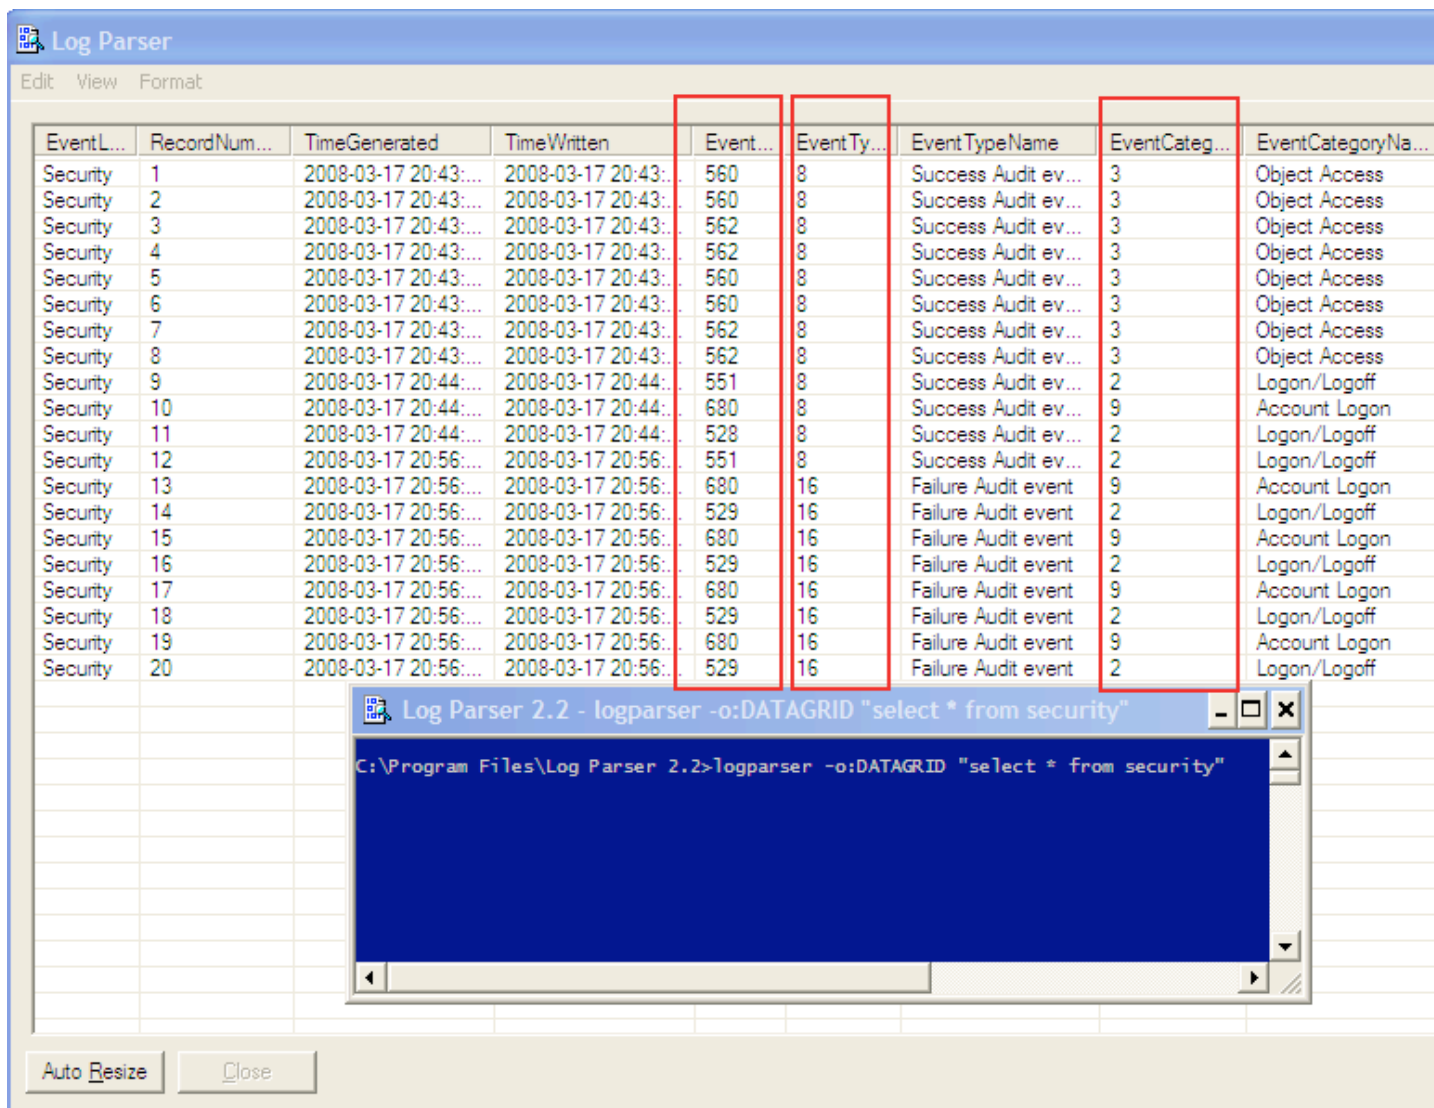


Figure 5. Simple Logparser query output

In fact, a Log Parser query can become quite detailed and complex. The following query represents only all of the event types that equal 16, sorted by the corresponding Event ID.

```
logparser -i:EVT -o:DATAGRID "SELECT TimeGenerated, TimeWritten, EventID, EventType, EventTypeName, EventCategory, Message from security WHERE EventType=16 ORDER BY EventID"
```

As you can see, queries can become more complicated. There are a few differences between standard SQL and the Log Parser SQL. With Log Parser, the input is from a file, instead of a table. The input is limited to just one file or resource. Joins are not possible, although sub-queries can be performed.

If you want to test it yourself, there is much information to be found in the html help file. Also, you can find useful information on the commandline. You can do this by typing: **Log Parser -h**.

In the online help file you can find numerous example concerning the output format, input options. There are different sections, each covering a specific topic such as input formats, output, filtering and sorting options. The Log Parser tool has a lot to offer investigations and can replace a whole bunch of other tools. I would recommend The Microsoft Logparser Toolkit by Syngress for further reading.

Windows Vista changes

Windows Vista comes with a much changed architecture with which to log all of the platform's activities. Although we can view logs with the Event Viewer as before, the Event viewer tool has changed dramatically. More items can be logged, there are better filter capabilities, and XML is introduced. Not only are the three different kinds of logging possible, but Vista also supports:

- Administrative tasks
- Applications
- Security
- Setup
- System
- Applications and Services Logs
- Forwarded Events.

When working with Windows Vista, you will find that the log file format has changed from .evt to .evtx and, as already mentioned, the new log-file format is XML based.

Unfortunately, when working with Log Parser on Vista, this causes some problems in parsing event logs from down-level systems like Windows XP. The easiest way to deal with event files in Windows Vista is to have them saved as .evtx. Vista provides the opportunity to convert exported Event Log (.evt) files from Windows XP and Windows Server 2003 to the .evtx format. This conversion can be done via the Event Viewer MMC, by choosing Save As.

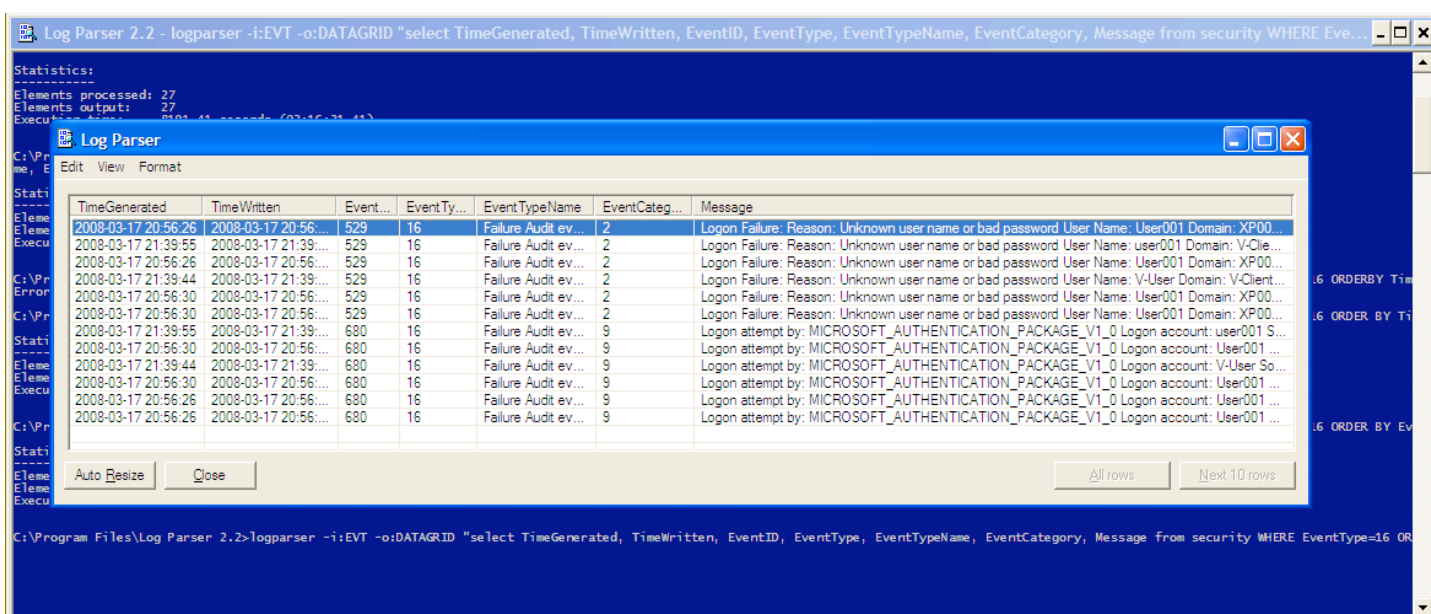


Figure 6. More complex Logparser query output

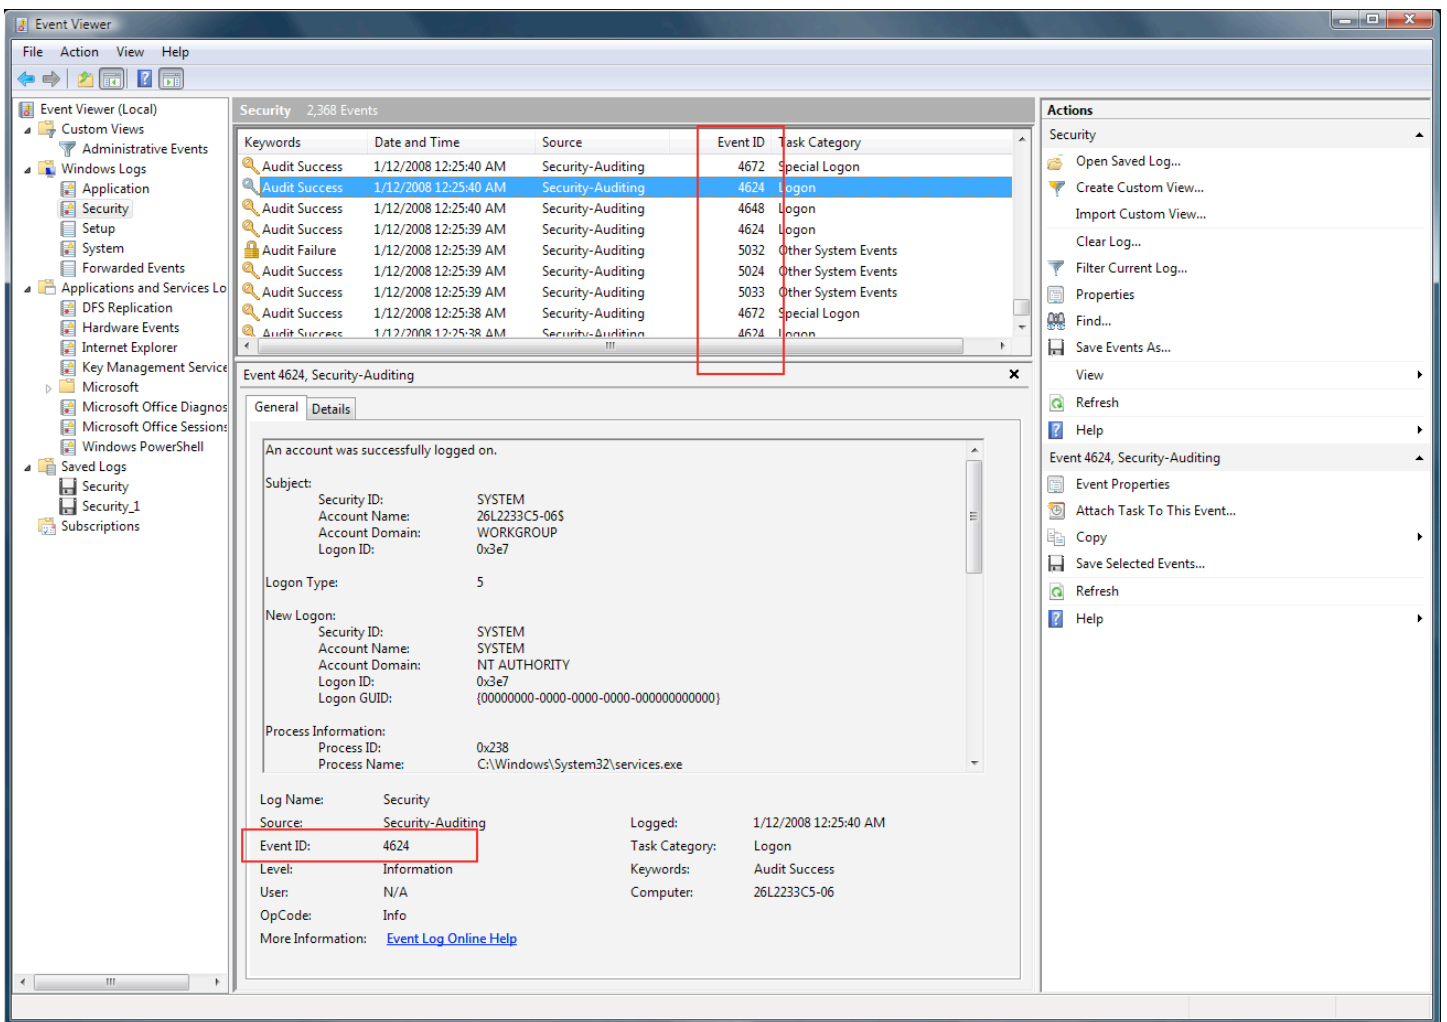


Figure 7. The new Event Viewer MMC in Vista

The Windows Events Command Line Utility (wevtutil.exe) can also be used to perform this conversion. We can make this happen by typing:

```
wevtutil epl application.evtx /lf:true
```

The 'epl' command exports the log file and the /lf:true specifies that this is an event log file.

In Vista and beyond, the wevtutil command can provide you with a wealth of information about the (security) event log used by Windows Vista / Server 2008. See for yourself by typing at the command prompt:

```
wevtutil gp Microsoft-Windows-Security-Auditing /ge /gm:true
>d:\VistaEventID.txt
```

You may have noticed that the (security) Event IDs have changed a lot in Vista, and much higher numbers are used. Some research on my part (on forums) revealed the following story.

There is a relationship between the pre-Vista security Event IDs and the new Vista security event IDs. The following rule is put in place: for the majority of security events: Vista EventId = Pre-Vista EventId + 4096.

This is because Microsoft needed to differentiate between the Vista events and the pre-Vista events. Indeed, as there were significant changes in Vista, Microsoft wanted to ensure that there was a relationship between the old and the new event IDs but, at the same time, they wanted to make the differences between them clear so that we could recognize them.

They initially considered offsetting the old IDs by adding a decimal number to get the new IDs (e.g. 6000, so the ID 528 would become 6528 and so on). However, Event IDs are declared in hexadecimals and are all 3 digits long (528 = 0x210), and so, ultimately, they decided to add 0x1000 (4096) to the existing event IDs.

The "add 4096" rule is not meant to imply that the events are all the same in the new system. Some events have changed codes, and some are combined. Simple re-numbering will therefore not work.

Security concerns

As with all solutions there are, inevitably, some concerns. Personnel with a higher level of privileges are often allowed to view and clear the log. Yet, even such a step is not always necessary, since it is possible to use a tool like Winzapper to delete specific events from the log.

Winzapper is a freeware utility / hacking tool that is used to delete events from the Microsoft Windows Security Log. It was developed to demonstrate that once the administrator account has been compromised, event logs are no longer a reliable source of information for investigating intrusions. Winzapper will allow a hacker to hide an intrusion by deleting only those log events that are relevant to the attack.

A way of defending against this, is to set up a remote log server with only necessary services present (hardened host), and allowing only console access, meaning that access at

the physical facility must be permitted. This is a step in a layered defense or "defense in depth" concept.

Given the ability of administrators to manipulate the Security Log to cover unauthorized activity, the separation of duties between operations and security-monitoring IT staff is essential, as are frequent backups of the log to a server that is accessible only to authorized personnel from the IT security or audit departments.

Naturally, there is also the risk of an incorrect configuration. A log can reach its maximum size at any time. At that point, it can start to either overwrite old events, or stop logging new events. This makes it possible for an attacker to attack the system by flooding the log and generating large numbers of new events.

However, instead of local logging and auditing, the Security Information Event Management (SIEM) can be put in place as a way of collecting and relating events from different sources and platforms, enabling them to be stored in a tamperproof way. It is also possible set up an alarm that would be triggered if certain strange and unexpected behavior is detected.

DIFFERENT TYPES OF EVENTS CAN AFFECT HOW YOU INTERPRET THEM FROM A SECURITY AND FORENSICS PERSPECTIVE.

A final word

I hope that this article about Windows Event logging, auditing and forensics will help you, as a security professional, to have a better understanding about the best ways of observing your Windows systems and network.

Like in this article, you can bring the pieces of the jigsaw together and get an idea of the different ways in which users are accessing your

computers, and how different types of Events are triggered.

Paying attention to this is crucial because different types of events can affect how you interpret them from a security and forensics perspective. More often than not, it is a couple of events from different sources and within a certain timeframe that provides us with an overview of what has really happened on our systems.

Rob P. Faber, CISSP, CEH, MCTS, MCSE, is an information security consultant. He currently works for Atos Origin, a global company and international IT services provider based in The Netherlands. His specialization and main areas of interest are Windows Platform Security, Ethical Hacking, Active Directory and Identity Management. He maintains a weblog on www.icranium.com and you can find him on the LinkedIn network.

H@cker | Halted™

USA
2008



Attend Hacker Halted 2008 Security Event in USA

May 28th - June 4th 2008, Myrtle Beach

Hacker Halted has been successfully held in Mexico City, Dubai, Singapore, Kuala Lumpur, China and now in the USA.

**Special Keynote Session
Featuring Howard Schmidt**

For more information, please visit
<http://www.hackerhalted.com>

Email: info@hackerhalted.com

<http://www.hackerhalted.com>

EC-Council

Traditional vs. non-traditional database auditing

By Michael Semaniuk



Traditional native audit tools and methods are useful for diagnosing problems at a given point in time, but they typically do not scale across the enterprise. The auditing holes that are left in their wake leave us blind to critical activities being performed within the systems that contain our most coveted trade secrets, customer lists, intellectual property, and more.

Would we be happy if our bank allowed people into the vault that contains our money without a camera monitoring their activity? Would we want to share our most personal data with a company that isn't a good fiduciary of our information?

The odds are we wouldn't want to participate in either scenario, but the reality is that this is what happens to our most private data all the time. We simply aren't aware of it because in the world of electronic data, we don't "see" what is going on. Employees and partners of companies have the ability to access our personal information in databases all over the world. And, although many of those companies have traditional security in place, most don't know what is actually happening with our

data—and the data of millions of other individuals.

In the recent past, native audit tools, such as SQL Profiler, trace functions, and triggers were all that we had. But they are no longer the only game in town. A new category of technology has emerged that empowers enterprises to "see" and immediately analyze what is going on with sensitive data.

This new technology, called Data Activity Monitoring (DAM), has the ability to monitor sensitive data as it is being accessed from data servers and analyze the activity to determine if the user, or the particular access, has the potential to endanger data or create a non-compliant situation.

We have historically shied away from performing extensive monitoring and auditing within our database environments because of performance and manageability issues and something that I call “information glut.” We can gather all sorts of interesting data with native auditing tools, but the result has always been slower systems, more management overhead and so much raw data that making sense of it is nearly impossible.

Performance and native auditing have been diametrically opposed. The more knobs and switches we enabled within a database tool

like SQL Profiler, the more overhead we introduced. This is an inherent problem because native tools leverage the same CPU and disk I/O as our production systems. While performance degradation is the downside, the upside is the plethora of data that we can extract. With auditing we can get information such as success/unsuccessful access, stored procedure activity, duration for a transaction and almost anything we can think of in relation to the activities that are taking place within our target environment. However, auditing issues have typically outweighed the benefits.

PERFORMANCE AND NATIVE AUDITING HAVE BEEN DIAMETRICALLY OPPOSED.

Another problem with native auditing is manageability. The manageability problem is in direct relation to the number of systems that we have in place. If we only have a single system to worry about, it is not a concern. Having a bit more table space to manage or a single data source to analyze means we can probably find the time to deal with it; however, multiply that by two, ten, fifty, or a hundred systems and now we have a nightmare.

Monitoring systems to make sure that they are functioning optimally for our users, system design, capacity planning and other duties tend to be the priorities in most environments and will suffer if there are a large number of audit tables to maintain. Add to that the consolidation and analyzing of the data found and it gets to be a very messy problem. Plus, native auditing tools aren't designed to do high-end analysis for trends, anomalous activity or the infamous needle in the haystack type of data we are interested in finding. They are not designed to be intelligent analytical tools. They are designed to collect data. So in order to make sense of the data, we need to build or purchase tools that can handle the analysis of the piles of data that we've collected.

As I mentioned earlier, an often-overlooked problem with native auditing is information glut. Where do we put the data that we've gathered? The last thing we need is another database so that we can dig through mounds of data to find what we're looking for. While disk space is fairly inexpensive these days, it

isn't easy to add additional storage to a system or to introduce a new server into the environment to simply act as the repository for audit data.

The DAM cure for performance and manageability issues

DAM technologies address this problem. The DAM solutions market is comprised of a number of vendors using a variety of methods to enable database auditing. Some concentrate on databases while others go beyond structured data to audit file share activity. For now, let's focus on monitoring activity in databases. As with native audit tools, there are pluses and minuses for each method. Let's review how a single or hybrid approach to monitoring activity can be accomplished without the performance issues, management overhead or information glut associated with native auditing.

The many flavors of DAM deployment

Generally, there are three DAM deployment methods that can be leveraged to monitor activity within a database environment. They are: network, agent and agent-less auditing.

The network-based approach typically utilizes a network-based appliance to monitor the activity occurring via the SQL protocol within the target environment. It will also store that information on the appliance for reporting and alerting purposes.

Proper deployment does require a bit of forethought, as there are two deployment models for network-based monitoring: inline and passive.

Inline is exactly what it sounds like: the appliance resides between the target database(s) and the network infrastructure. All SQL activity passes through the appliance before it reaches the database server. This can allow for a high level of action to take place on the appliance. For instance, if a query is extracting 1,000 rows of data, but the client is only allowed to retrieve 100 rows of data, an inline device can typically alter the response to 100 rows. Many inline devices bring additional security-like functions to the table. Protocol validation and blocking of activity are a couple of these additions. The goal of protocol validation is to help prevent worm or zero-day attacks from taking place against the database server, and blocking is the capability to shut the user down.

There are advantages to inline solutions, such as the ability to take action, alter a response and validate the protocol. The downside is latency, service interruption and scalability. Any device that acts as a forwarding element (any piece of network equipment, a patch panel, an inline device such as an Intrusion Prevention System (IPS) or database firewall) introduces latency to the environment. If the performance is a concern, then carefully weigh the cost of running inline. It also requires a service interruption to install, remove or upgrade. Your mission critical applications may have more downtime than what is acceptable for the business.

Finally, inline devices are limited in the total number of connections that can pass through them, causing the total number of devices protected to be rather low. This can be good for a point solution but not necessarily for an enterprise deployment.

THE NETWORK-BASED APPROACH WILL DO ALL OF THE HEAVY LIFTING IN TYPICAL DATABASE ENVIRONMENTS.

The second type of network-based solution is passive. The network-appliance monitors activity by capturing and evaluating copies of the data stream between clients and the database servers as presented by the network infrastructure of the target environment--similar to the way a network engineer uses a network sniffer to monitor traffic. This is similar to the inline approach, in that it monitors via the network, but its deployment model is fundamentally different. Both analyze the SQL protocol to determine what is relevant and what is not. Passive deployment allows a single appliance to scale to a large number of devices because it is not in the traffic path. Passive deployment eliminates the latency that could be introduced with an inline solution and can be installed without any service interruption.

There are also tradeoffs with a passive solution. There is no ability to alter a response or block activity. Inline solutions and passive solutions handle threats somewhat differently. If an inline solution sees a username combined with an application that it has been told to intercept, it will prevent the network packets

from reaching the target. Passive solutions typically reset a session at this point by sending a reset packet to both the client and the server, accomplishing the same goal in a different way.

The network-based approach will do all of the heavy lifting in typical database environments. They monitor and audit activity as it happens without impacting the server itself. But what about encrypted network communications or local console activity, such as database management on the local system via secure shell (SSH), remote desktop protocol (RDP) or console access via the local keyboard and mouse? Network-based solutions generally only understand the native database protocols themselves, leaving a hole in the audit trail. But there are other methods that capture the activity that network-based monitoring misses.

Agent-based and agent-less auditing are used to fill in the gaps left by network-based auditing. Like network-based auditing, agent-based comes in a couple of flavors, depending on the goal of the deployment.

Agents can be used to parse database logs (i.e. redo logs), act as a proxy for console or encrypted activity, act as a sniffer to capture console or network activity or as a hybrid. They may utilize some combination of these methods to get the data. Agents, especially hybrids, can potentially do the same things that the network-based approach does by monitoring clients as they access the environment over the network. They can also capture the local and encrypted activity. This gives an agent-based approach a wide set of capabilities. Additionally, some agent technology can present before and after values for fields and, in addition to monitoring SQL activity, may also be able to monitor database configuration files.

On the downside, agents can be very heavy on a system. If they do everything to capture the activity and aren't supplemented by a network-based solution, they will negatively impact the target system, consuming significant CPU, disk and network I/O in order to monitor the activity of interest, and on occasion may require the use of native auditing. Not to mention the need to deploy another database to store audit data. Agent-less auditing is similar to the agent-based approach, but it has a significant advantage in that it does not always require the deployment of software. Agent-less auditing is another word for native auditing, but the difference is in the details. Agent-less auditing DAM solutions focus the native audit tools to capture only the data of interest and manage that data once collected. SQL Profiler can be a powerful tool and when it's tuned to capture just console activity versus all network activity, overhead is kept to a minimum. Combine that with some form of automated data management, and we can capture the encrypted or console activity. This approach allows for visibility into the DML activity that may be occurring as we call stored procedures.

As with all of the approaches, there are downsides to agent-less auditing. Separation of duties is challenging when we use native tools, and many industry regulations require moni-

toring of DBA activity. Tuning of the native tools can also be a challenge. Databases from different vendors use different audit methods. This could lead to capturing more than just the encrypted or console activity and overtax servers.

Finding the right solution for your database auditing situation

How do we sort through the alternatives to arrive at the best solution—one that balances our need for intelligence about data activity with our need to keep business systems humming along? The key is to enable the visibility needed for data protection and compliance initiatives while protecting the performance of systems (and our sanity) by leveraging the strengths of multiple approaches.

The network-based approach alone can handle the network activity but not the console and encrypted activity. The 100% agent approach introduces data management issues and may cause significant pain on the database server itself. A 100% agent-less, like the agent-based approach, may introduce the same issues.

The key is a flexible, multi-pronged approach to database auditing. Network-based solutions are critical to the overall auditing effort but encrypted and console activity call for additional functionality. Most network-based appliances are complimented by an agent-based solution, or in some cases both an agent-based and agent-less solution. This combination of approaches allows the appliance to do the heavy lifting by monitoring the client/server traffic that originates from network-based clients. The agents and/or agent-less functionality covers encrypted and console activity. Combining the strengths of these different or non-traditional approaches allows for the deployment of a comprehensive audit and monitoring solution with all of the upside and none of the down. Now we can get the visibility into data activity that we need without the issues associated with traditional/native auditing functionality.

Michael Semaniuk (CISSP, SSCP, MCSE, CCSE+) is the Compliance and Security Solutions Architect at Tizor Systems (www.tizor.com). He has spent the last 12 years working in the Information Security field. Over this time, he has been involved with hundreds of installation and consulting engagements, including vulnerability and penetration testing, gap analysis and security product deployments. His professional experience includes: Triumph Technologies, Interliant, Akibia, Top Layer Networks and Tizor Systems.

o3: magazine

Open Source / Enterprise
Free DIGITAL magazine

<http://www.o3magazine.com>

INSIDE: Open Source Web Acceleration with Varnish Cache

o3: The Open Source Enterprise Magazine

Issue 6
August 2007

<http://www.o3magazine.com>

Deploying **Globally** Distributed Web Applications with Ruby on Rails

Production Rails Apps with Mongrel

Deploying PostgreSQL

Simple Appliance
Stacks with LFS

Secure Global
Networks with OpenVPN

Enterprise WiFi --
Thin Access Points



This issue is sponsored by:



<http://www.othello.tech.net>

INSIDE: Building Secure Postfix SMTP Appliances with LFS

o3: The Open Source Enterprise Magazine

Issue 8
September 2007

<http://www.o3magazine.com>

Designing **Scalable** Enterprise SMTP Networks for Email

Using **Dovecot** for imapd / pop3d

Encrypting Mail Protocols

Using **DSPAM** to reduce
storage requirements

Web based email with
Roundcube



This issue is sponsored by:



<http://www.arubanetworks.com>



Payment card data: know your defense options

By Ulf Mattsson

With the advent of the Payment Card Industry Data Security Standard (PCI DSS), protecting stored credit card numbers is no longer optional. Any company that stores, processes, or transmits credit card information—regardless of size or volume of transactions—must secure stored credit card data or face serious consequences for non-compliance, including fines, higher transaction fees, the loss of brand integrity, and erosion of market value.

While the PCI standard offers broad guidance - featuring rules on the proper use of firewalls, web application firewalls, computer access controls, antivirus software, and more - encryption requirements are proving to be among the most difficult for organizations to address. To complicate the situation even further, the compensating controls defined in PCI DSS 1.1 are not fully addressing the growing threat from data level attacks.

This article will review different approaches to protect credit card data that can be combined to significantly strengthen an organization's security posture, while minimizing the cost and effort of PCI compliance.

Evolving data threats

A growing number of applications perform electronic commerce, selling products, infor-

mation or services in an Internet based environment. Another category of applications that also attracts attackers attention are those that deliver services on behalf of financial firms such as rewards redemption, report delivery for banks, and merchants and banks information exchanges. Unlike traditional static internet applications, many of these applications store and process information that is strictly regulated (e.g. GLB-A, SEC) and most must satisfy SOX compliance requirements.

Typically, these applications compile databases containing hundreds, thousands, or even millions of credit card accounts and personal identifiable information. For attackers, these databases represent an excellent opportunity for theft and fraud.

One major database attack vector is via the application layer.

Back in the 90's, system configuration, buffer overflow, and other platform level type flaws were all the rage, but these have become increasingly easy to manage. Economies of scale have given ubiquity and commodity status to packet-filtering firewalls, multi-platform patch management systems, vulnerability scanners, and intrusion prevention systems. But any security system is only as strong as its weakest link, so that's what attackers look for - up to the application and even client level, and down to the system internals and driver level. At the top level you have the world of web application attacks, where web applications are used as proxies to attack the underlying databases.

Custom web application security is different than platform security. There are no vendor advisories or patches and the burden of patching code is on the company that created the application. The attackers out there are very determined - the sophistication of botnets and worms constantly attacking the internet demonstrate that organized crime is hiring very talented people to attack systems - and there are fortunes being made from this kind of theft. Criminals aren't above extortion and blackmail of highly placed insider employees, who might have access to the very routines and data you are trying to protect. All in all, it's the stuff of system administrators' nightmares.

CUSTOM WEB APPLICATION SECURITY IS DIFFERENT THAN PLATFORM SECURITY

Web applications attract hack attacks

Attackers like web applications because these apps have built in, exposed mechanisms that must have connectivity to the data the attacker is after. The attacker thinks, why compromise an entire system when you can manipulate the application into releasing the data that you're looking for? Most protection is at the network, not the application layer, so the chances of getting caught are much lower. Application attacks are much harder to catch and prevent at the network layer, because the network components don't understand the application, its logic, or which resources should be accessed and by which user roles. It is incumbent upon organizations to understand the attackable surface area represented by web applications, particularly those that store and process confidential personal or payment card data.

Common vulnerabilities and exposures across the Web include application-level attacks such as cross-site scripting, SQL injection and buffer overflow as the favorite vectors for Web attacks. SQL injection attacks are caused primarily by applications that lack input validation checks. These attacks are relatively easy to carry out. Thankfully, they are almost as easy to defend against - that is, good solutions are available that can help plug the gaping holes that allow such attacks - but most enterprises still do not own a Web application firewall - which would control the execution of files or

the handling of data by specific applications - and many don't yet do any application scanning. Furthermore, many enterprises have never had a third party audit their apps for vulnerabilities - in fact, many large enterprises don't even know how many websites they operate.

Typical threats in a retail scenario

Any substantive analysis and discussion about an application's risks must include a discussion of the likely threats the application will face during its anticipated deployed life. In the analysis of the POS, the most likely threats that the system may face, for reasons we will describe below, are 1) malicious insiders and 2) technically knowledgeable outsiders motivated by profit. A brief discussion of each, along with the respective rationale, follows below.

There is likely to be one primary category of insider threat to the POS: retail employees. They may either be enlisted by an outsider or may enlist the help of an outsider in order to attack the POS. Their motivations are likely to be either profit or to cause harm to the retailer by way of a direct denial of revenue and/or tarnishing the retailer's brand with bad publicity that would almost inevitably be the result of a successful compromise. In any of the above scenarios, the insider has learned how the POS system functions to a level significant

enough to attempt an attack.

Traditionally, insider threats are the most difficult to prevent and detect. Further, it is likely that no technology solution will be adequate to safeguard against every possible attack scenario. However, other industries (notably the financial services industry) have handled insider threats for centuries. A second category of threat that must not be neglected is outsiders, their motivation is far more likely to be profit rather than to harm the retailer's reputation.

To combat both internal fraud and attacks emanating from outside of the organization, organizations need to develop applications that prevent against all forms of attack and regularly assess their security situation for potential vulnerabilities. Conduct regular security self-assessments, checking for new hacker tactics and vulnerabilities on your infrastructure and applications, and remedying any

problems that you identify. Also conduct periodic automated network scans. Just be aware that there is no single tool that can find and fix every vulnerability. The same goes for security solutions, Web application firewalls protect against major threats, but there's simply no such thing as completely bulletproof protection. That's why layers of protection work better than over-reliance on one or two methods or solutions.

Ultimately, you want to build the vulnerability scanning and testing phase into your development process. Realistically, however, enterprises should be more concerned about the applications they've already deployed than about revamping their QA process. Enterprises should attack the problem first by identifying all their sites and the applications running on them. An audit by a third-party expert and a scan by a vulnerability scanning tool can give the enterprise a starting point for remediation.

THOUGH SOURCE CODE ANALYSIS AND SIMPLE, CLEAN APPLICATION DESIGN CAN HELP TO ACHIEVE HIGHER SECURITY ASSURANCE IN CRITICAL APPLICATIONS

Achieving higher security in critical applications

Though source code analysis and simple, clean application design can help to achieve higher security assurance in critical applications. We know that for many organizations adopting this mindset will require an overhaul for their existing software lifecycles. Few people are truly analyzing their design and their code. If the assurance of externally facing applications is of utmost importance, then design/implementation time controls (e.g. static analysis) is very important.

Methodologies must maximize the efficiency and effectiveness of the time scoped for a particular assessment activity. One way to do this is by identifying the most comprehensive automated tools. Computers are good at automating things in a repeatable, measurable way. But it's important to remember that there are many classes of vulnerabilities which automated tools have serious trouble finding.

This is partly a function of the perspective from which the tool operates, such as Source Code Analysis vs. Fault Injection:

- Fault Injection is interactive testing of a web applications including spidering, querying for known vulnerable scripts or components, testing for conditions like forceful browsing, directory traversal, and using the results of spidering to identify all points of user input to test for flaws like SQL injection, XSS, CSRF, command execution, etc.
- Source Code Analysis is often using a combination of techniques such as searching for strings, identifying user input vectors, tracing the flow of data through the application, and mapping execution paths.

It is critical that we understand what automated tools can and can't find, and develop other methods for identifying the "false negatives" – vulnerabilities that exist, but were missed. The two main approaches that exist at present for web application testing are "Fault Injection" and "Source Code Analysis."

There are also two more philosophical approaches, "White Box," and "Black Box." Awareness is growing for a different technique often referred to as "Gray Box assessment," which integrates the approaches described above. This approach combines static and fault injection testing techniques, in order to compensate for their different detection capabilities, and also integrates elements of white and black box methodologies.

Software security, in the form of code review, penetration testing and the remediation of issues uncovered by these methods, are a sound practice. But as mentioned above, the chances are good that not all flaws will be caught. For the foreseeable future the problem of bad software will remain and we will need external security products like Web application firewalls. That said, you should fix the code/patch applications whenever you can and not view Web application firewalls as a magic silver security bullet.

Options for credit card data protection

A strategic approach is to implement solutions that are automated, integrated, scalable, and secure in an enterprise environment. A mature solution should provide a choice to balance and optimize the mix of different approaches to credit card protection across different systems in the enterprise, including tokenization, encryption and hashing.

How to make hashing more secure

Hash algorithms are one-way functions that turn a message into a fingerprint, usually several dozen bytes long binary string to avoid collisions.

PCI provided standards for strong encryption keys and key management but is vague in different points regarding hashing.

Hashing can be used to secure data fields in situations where you do not need the data to do business and you never need the original data back again. Unfortunately a hash will be non-transparent to applications and database schemas since it will require long binary data type string (longer than the 20 bytes for the broken SHA-1 or two-way symmetric encryption).

An attacker can easily build a (rainbow) table to expose the relation between hash values and real credit card numbers if the solution is not based on HMAC and a rigorous key management system. Salting of the hash can also be used if data is not needed for analytics.

Some people think that it is sufficient to encrypt or hash, regardless of algorithm, key lengths, key management and other controls. Some hashing implementations are ineffective. Using a plain unsalted hash to protect a fixed length 60 bit number, to produce a 160 bit hash (SHA-1) is a bad practice. It is susceptible to rainbow attacks by hashing up a mapping table between all relevant input values and the corresponding output hash values. A weak hash scheme is easy to break.

One approach is to store the salt with the hash. Salting adds uniqueness to the output values making them harder to break by brute force. The salt should be much longer than 4-8 bytes to provide proper protection. If you know the salt, you may apply a brute force attack. The idea with the salt is that you need to attack each value individually (if you have a unique salt for each value). This will take much longer time compared to not using salt. Salting using a single large salt value for all hashed PANs - could be broken in less than hours if the salt is compromised. The first six digits of the PAN are known and greatly shrinks down the amount of storage required to generate the rainbow attack table for the remaining digits.

A more secure approach is to use a secret salt and don't store it with the hash. The salt can be as long as you like and the resulting hash is stronger. It is only better if the secret salt remains a secret. A "secret salt" is a key. An HMAC, is a type of message authentication code (MAC) calculated using a specific algorithm involving a cryptographic hash function in combination with a secret key. As with any MAC, it may be used to simultaneously verify both the data integrity and the authenticity of a message. Any iterative cryptographic hash function, such as MD5 or SHA-1, may be used in the calculation of an HMAC. The resulting MAC algorithm is termed HMAC-MD5 or HMAC-SHA-1 accordingly.

The cryptographic strength of the HMAC depends upon the cryptographic strength of the underlying hash function, on the size and quality of the key and the size of the hash output length in bits.

If you need to replace the hash values with a future more secure option you would need to have the reversible data to replace it with the new values. If so, this might be a good time to build an internal protected rainbow table for the hash values. Some applications encrypt the PAN in several ways, AES with key rotation, public key encryption for total loss recovery, and a salted SHA-1 for data integrity and pattern recognition. Some organizations need

to be able to prove transactional data in the case of a dispute for up to two years. To get the PAN info from the processor requires additional efforts, so the AES value is still needed on the system. Loss prevention applications may have a greater need for pattern recognition, but maybe not access to the actual PAN. In some cases they need to go back as far as four years. Pattern recognition is possible on the AES if the keys are rotated and kept synchronized with all data. A secret key hash (HMAC) can be another option for this business need. It would not be reasonable to decrypt the entire databases. Public key operations are not either reasonable due to the length of the result, processing cycles and other reasons.

AN OUTSOURCING ENVIRONMENT MUST BE CAREFULLY REVIEWED FROM A SECURITY POINT

How to implement secure tokenization

The basic idea behind tokens is that each credit card number that previously resided on an application or database is replaced with a token that references the credit card number. A token can be thought of as a claim check that an authorized user or system can use to obtain the associated credit card number.

Rule 3.1 of the PCI standard advises that organizations “Keep cardholder data storage to a minimum.” To do so, organizations must first identify precisely where all payment data is stored. While this may seem simple, for many large enterprises it is a complicated task for a large enterprise the data discovery process can take months of staff time to complete. And then security administrators must determine where to keep payment data and where it shouldn't be kept. It's pretty obvious that the fewer repositories housing credit card information, the fewer points of exposure and the lower the cost of encryption and PCI initiatives.

In the event of a breach of one of the business applications or databases only the tokens could be accessed, which would be of no value to a would-be attacker. All credit card numbers stored in disparate business applications and databases are removed from those systems and placed in a highly secure, cen-

tralized tokenization server that can be protected and monitored utilizing robust encryption technology. Tokenization is like network segmentation, as a way to move card data to another internal or external server. The access to the card data is still in your POS or other systems, so be careful of how a tokenized system can be attacked. Most merchants implement tokenization motivated by cost savings.

Whether a security administrator is tasked with complying with the PCI standard, adhering to another standard such as HIPAA or GLBA, or is simply looking to optimize the security of sensitive information, encryption is essential. By combining standard encryption with an approach to encryption in which sensitive data is centralized and tokenized, organizations can benefit in terms of security, efficiency, and cost savings for some application areas within an enterprise.

An option is to outsource your credit card information and the tokenizing solution. One question is how that will affect the risk and liability. A larger organization can potentially provide more secure environment in-house. An outsourcing environment must be carefully reviewed from a security point and also providing a reliable service to each globally connected endpoint.

Tokenization can minimize exposure of data and keys

As mentioned above, PCI requirement 3.1 requires that organizations keep payment data in the minimum number of locations, and this approach addresses this requirement fully. Security is immediately strengthened by minimizing the number of potential targets for would-be attackers. PCI requirements 3.5.1 and 3.5.2 mandate that access to keys is restricted to the fewest number of custodians and that keys are stored securely in the fewest possible locations. By centralizing keys on a secure server, a tokenization server optimally addresses these requirements.

Unencrypted data never resides in databases or in application memory. Application users never see payment data in clear text unless they have specific, valid permission. With this approach, if an attacker somehow bypasses both the token and encryption, they will have

access to only one card number. In contrast, with many encryption solutions, if an attacker gains access to one cryptographic key, they can potentially decrypt thousands or even hundreds of thousands of records. All keys and policies can be managed centrally, as opposed to having keys in multiple, distributed locations. This makes PCI-required tasks, such as key revocation and rotation, much faster and easier.

Also information like Social Security numbers are, in many systems, less protected than credit card data. Tokenization combined with an enterprise class encryption solution can be the fastest way to substantially reduce their risk to this data. Protection the complete 'flow of data' and the supporting the interdependencies among and referential integrity across systems can be supported by a tokenization solution combined with an enterprise class encryption solution.

TOKENIZATION CAN CREATE A SMALL NUMBER OF EXTREMELY ATTRACTIVE TARGETS FOR DATA THIEVES

Algorithmic tokens vs. lookup/index tokens

It may appear that a credit card number is 16 digits, and 10^{16} would be a very large number of tests, more than half of a card number is easily learned or is guessable. I also have the Luhn checksum as the last digit to verify a valid account number. Thus I start out knowing "ten" digits worth of a sixteen digit card number. I now have to test only 10^6 , or only one million possible tokenizations to compare to a given token. On a modern desktop that calculation would take only a few seconds.

An attacker, however, would take every advantage available to him. If the tokenization algorithm is present and can be stolen or reverse engineered an attacker could run the million tests and have a good chance of discovering the full account number associated with a token in a minute. The goal of an attacker is to acquire as many account numbers as easily as possible. With a stolen database of sales information, cracking even a fraction of the tokens into valid credit card numbers would yield great profits. A well protected

token-lookup table with completely randomized token values can still be stolen but cannot be reverse engineered by an attacker. How to choose a secure algorithm is discussed further below.

Other issues with tokenization

Tokenization can create a small number of extremely attractive targets for data thieves. Your tokenizing routine has to be guarded as securely as you would the decryption key to your account numbers. If an attacker can access the tokenizer from an unknown, unauthenticated, or unaudited source, he can use it to perform his own "testing." An added detection of frequency and volume of requests to the tokenizer could detect abnormal pattern of requests from a certain user/client but building a secure tokenizing solution that meet both the business requirements and the security requirements can be a complex job. This also applies to the lookup table. Using a key in the algorithm as for HMAC means you can't build the table yourself even if the algorithm is known. You need the key to pass the tokenizer. With an algorithm you may build your

own lookup table or use brute force attack locally. Using a random lookup table means you need to pass the tokenizer. Using a key in the algorithm as for HMAC means you can't build the table yourself even if the algorithm is known. You need the key.

For small data sets, and if the algorithm is known or at least accessible, you can easily build a lookup table or apply brute force attack. Using hashing ("SHA-1 based tokenization routine") is an obvious case – you'll have everything available and may pre-compute all values. But even if you're using some hidden tokenizer routine (hidden based on non-public algorithm or some key as for HMAC) you may have the same problem. Since the same input must produce the same output (to support equality search), if having access to the tokenizer, you may send all possible input values into the tokenizer and see what is returned. This will give you the lookup table. You don't need to know the key and/or the algorithm, it's enough having access to the tokenizer. This is the same as when using encryption with no IV. Even if you can't decrypt the data, if having encrypt access, you may encrypt all possible input values and see what is returned. You don't know the key, but there is no need to.

Secure access and key management

To implement this approach, an organization must deploy a single, secure server that will house all payment data and that will act as the central repository for managing keys, tokens, and security policies. Ensure that the server can automate routine key management tasks and intelligently handle key rotation or that it can be integrated with third-party solutions that deliver these capabilities. Capabilities need to be in place to ensure that only authorized staff can access administrative functions.

The server should track all decryption activity to provide an audit trail specifying who has decrypted sensitive payment data. The server should support high volumes of encryption routines and token requests without impeding the performance of associated applications and workflow. In addition, the server should be enabled for continuous processing, even in the event of a server outage.

Development time and expertise

Developing all the capabilities outlined in this article can present significant challenges if a security team seeks to build a solution in-house. Following are a few of the biggest hurdles an internal team could face in this endeavor.

To be implemented effectively, all applications that currently house payment data must be integrated with the centralized tokenization server. Developing either of these interfaces would require a great deal of expertise in order to ensure performance and availability. Writing an application that is capable of issuing and managing tokens in heterogeneous environments and that can support multiple field-length requirements can be complex and challenging. Furthermore, ongoing support of this application could be time consuming and difficult. Allocating dedicated resources to this large undertaking and covering for responsibilities this staff would otherwise be fulfilling could present logistical, tactical, and budgetary challenges.

For many organizations, locating the in-house expertise to develop such complex capabilities as key management, token management, policy controls, and heterogeneous application integration can be very difficult. Writing code that interfaces with multiple applications, while minimizing the performance impact on those applications, presents an array of challenges. The overhead of maintaining and enhancing a security product of this complexity can ultimately represent a huge resource investment and a distraction from an organization's core focus and expertise. Over time, an organization's security needs change and so the cryptographic algorithms or encryption mechanisms in use may also change. Once initial development has been done, the development team may need to add capabilities for integrating with a new protocol or encryption solution, which may entail a substantial rewrite of the application in use. These challenges are significant and can severely undermine the value of a tokenization server. Security administrators looking to gain the benefits of centralization and tokenization, without having to develop and support their own tokenization server, should look at vendors that offer off-the-shelf solutions.

The challenges of encryption and key management

While most security professionals recognize the merits of encrypting credit card data, they often struggle with a few common challenges. A problem for some organizations, especially those with credit card data stored in multiple systems across an enterprise, is the cost of basic native encryption, both in terms of upfront implementation costs, performance and in terms of ongoing maintenance. Integrating encryption in enterprise applications poses significant challenges. Legacy applications are equally challenging as most will have no encryption or key management functionality, necessitating significant modifications to comply with PCI. Finally, the requirement for periodic key rotation typically entails application downtime and consumes extensive system resources, stretching the performance limits of many application infrastructures.

Third party encryption solutions can help with these issues. Often, it is only after an initial encryption deployment that administrators realize how much ongoing effort is required for key management. PCI rules include the need to establish dual control of keys, prevent un-

authorized substitution of keys, revoke old or invalid keys, and rotate keys on a routine basis. Establishing and maintaining all these processes can require a tremendous amount of time and resources on an ongoing basis.

The principle behind dual control and split knowledge dual control and split knowledge is required to access the clear text key. Only a single master key will be needed under this control. The determination of any part of the key must require the collusion between at least two trusted individuals. Any feasible method to violate this axiom means that the principles of dual control and split knowledge are not being upheld. This principle is enforced by requiring both dual control, and split knowledge. That is, at least two people are required to 'reconstruct' the key, and they each must have a physical thing and they each must have some information that is required. The use of a key in memory to encipher or decipher data, or access to a key that is enciphered under another key does not require such control by PCI DSS 1.1. Keys appearing in the clear in memory, the principles of dual control and split knowledge are difficult but not impossible to enforce.

INTEGRATING ENCRYPTION IN ENTERPRISE APPLICATIONS POSES SIGNIFICANT CHALLENGES

How to manage the keys at the retail store

It is essential in many retail environments that each local POS encryption service retain a copy of the encryption key while in an operational and unlocked state. At the same time this will present a security exposure of the POS system. An attacker with access to a local encryption service could potentially peruse the system's processes and memory to acquire the key and decrypt data. The likelihood of this sort of attack succeeding is quite low, the result of a successful attack could be very high. The data encryption key should not be stored, once decrypted by the key encryption key, in a file or database. Key can be in clear form in memory but best practice is to add additional protection including compartmentalization, fragmenting, moving, zeroizing and providing decoy structures(. It should not be feasible for a DBA or Administrator to obtain or

substitute keys. The exposure of keys in a memory dump can be addressed by removing the crypto operation to a separate box or by applying the methods discussed earlier. If crypto operations are moved to a separate box, the logon context to that box will be a potentially weak point in the encryption solution. A solution that is deploying keys to local encryption environments must use an approach where a key encryption key does not reside in its entirety on such a server. Using write-only public keys can provide dual control so that no one person has access to any key that can ultimately decrypt cardholder information.

A single key custodian should not have access to decrypt cardholder data. Dual control should be implemented also at the retail POS level. During normal operation, at startup of the application, the POS should authenticate to a central security server.

The credentials for decrypting the data encryption key can then be passed to the local encryption device over an authenticated and encrypted session. Dual control of the credentials in this instance would be lost if a single individual would obtain both parts of the credentials for input. It is possible to maintain split knowledge whereby the Store Manager obtains half of the key from source A, and the IT Support person obtains half from source B. As the integrity of the credentials may be at risk following such a process, it should be noted that a key rotation of the data encryption key must then occur, as well as a change of the credentials.

Think holistically

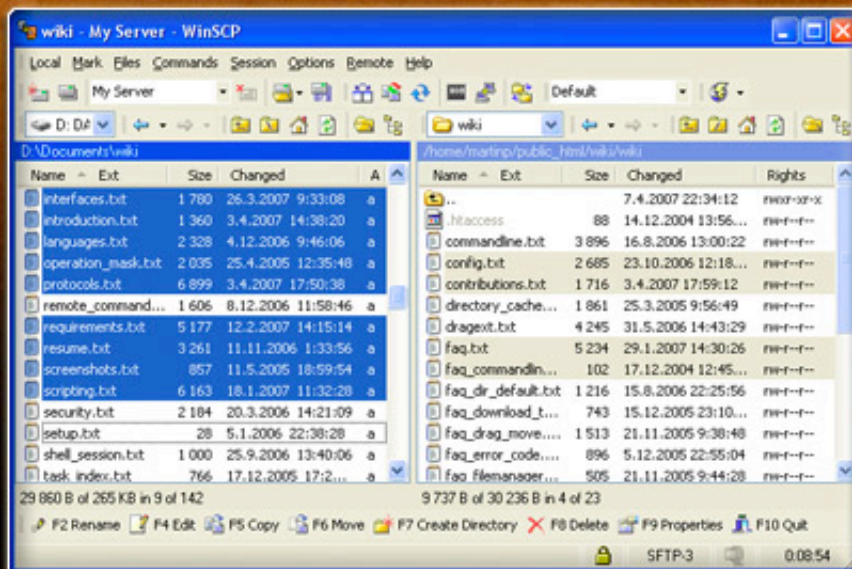
A security system is only as strong as its weakest link. We can't rely on applications to do all the work for us and we can't just throw money at the data security problem and hope it will go away. A holistic layered approach to

security is far more powerful than the fragmented practices present at too many companies. Think of your network as a municipal transit system – the system is not just about the station platforms. The tracks, trains, switches and passengers are equally critical components. Many companies approach security as if they are trying to protect the station platforms, and by focusing on this single detail they lose sight of the importance of securing the flow of information.

It's critical to take time-out from managing the crisis of the moment to look at the bigger picture. One size doesn't fit all in security so assess the data flow and risk environment within your company and devise a comprehensive plan to manage information security that dovetails with business needs. A data protection-driven holistic plan is the only way to truly secure data – it allows you to think strategically, act deliberately and get the absolute best return on your data security investment.

Ulf T. Mattsson is the CTO of Protegrity. Ulf created the initial architecture of Protegrity's database security technology, for which the company owns several key patents. His extensive IT and security industry experience includes 20 years with IBM as a manager of software development and a consulting resource to IBM's Research and Development organization, in the areas of IT Architecture and IT Security. Ulf holds a degree in electrical engineering from Polhem University, a degree in Finance from University of Stockholm and a master's degree in physics from Chalmers University of Technology.

WinSCP is freeware SFTP, FTP client for Windows using SSH. Its main function is safe copying of files between a local and a remote computer.



Download it for free at winscp.net