

(IN) SECURE

OPEN. INFORMATIVE. TO THE POINT. Issue 9 - December 2006



INTERVIEW WITH KURT SAUER, CSO AT SKYPE
WEB 2.0 DEFENSE WITH AJAX FINGERPRINTING & FILTERING
CREATING BUSINESS THROUGH VIRTUAL TRUST

TABLE OF CONTENTS

Page 04 - [Corporate security news](#)

Page 09 - Effectiveness of security by admonition: a case study of security warnings in a web browser setting

Page 17 - Interview with Kurt Sauer, CSO at Skype

Page 20 - [Latest additions to our bookshelf](#)

Page 22 - Web 2.0 defense with AJAX fingerprinting and filtering

Page 32 - Hack In The Box Security Conference 2006

Page 36 - Where iSCSI fits in enterprise storage networking

Page 38 - [Software spotlight](#)


Page 39 - Recovering user passwords from cached domain records

Page 47 - Do portable storage solutions compromise business security?

Page 54 - [Events around the world](#)

Page 55 - Enterprise data security - a case study

Page 64 - Creating business through virtual trust: how to gain and sustain a competitive advantage using information security



Welcome to (IN)SECURE 1.9 the digital security magazine

Another year is almost over and a plethora of information security problems are behind us. To let 2006 go out in style, we bring you a feature packed issue of (IN)SECURE. As the feature interview for this issue we had the pleasure of talking with Kurt Sauer, the CSO at Skype, one of the most well-known companies in the digital world.

We'll be back next year with many new ideas in the pipeline. Stay tuned for coverage from a few conferences including the RSA Conference in San Francisco and the Black Hat Briefings & Training in Amsterdam. If you're attending, be sure to drop me an e-mail and we'll grab a drink.

We wish you a safe 2007!

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. For reprinting information please send an email to reprint@insecuremag.com or send a fax to 1-866-420-2598.



New enterprise single sign-on authentication software



digitalPersona.

DigitalPersona announced the latest version of its award-winning enterprise product, DigitalPersona Pro 4.0. The new and improved software delivers a complete, accurate and trusted fingerprint Enterprise Single Sign-On (ESSO) solution with more secure authentication, improved manageability and the broadest support available for the world's leading biometrically-enabled notebooks including models from Lenovo, HP, Dell, and Toshiba.

DigitalPersona Pro 4.0 can be deployed on all of the biometrically-enabled notebooks for the enterprise in addition to supporting DigitalPersona award-winning U.are.U readers. For more information visit www.digitalpersona.com

BT acquires Counterpane Internet Security

BT announced that it has acquired Counterpane Internet Security Inc., a provider of managed networked security services, as part of its strategy to expand and develop its global professional services capabilities.



Counterpane currently monitors 550 networks worldwide for multinational and Fortune 100 customers. The company is based in Mountain View, California. Post-acquisition, Bruce Schneier, the founder of the company, will continue in his role as CTO and Paul Stich will remain its CEO.

Andy Green, CEO BT Global Services, said: "Counterpane is a welcome addition to BT's global professional services community. As more and more of our customers seek to exploit the opportunities of globalisation, we are finding that increasingly business critical applications are dependent upon the resilience and security of their infrastructure." For more information visit www.counterpane.com

WhiteHat Security Debuts Sentinel 3.0



WhiteHat Security announced WhiteHat Sentinel 3.0, the industry's only continuous vulnerability assessment and management service for Web sites. Sentinel 3.0 reduces the burden of securing Web applications with an on-going service that provides up-to-date and comprehensive identification of the vulnerabilities that are putting online customer and corporate data at risk. It is the only solution that can assess for all 24

classes of vulnerabilities identified by the Web Application Security Consortium's (WASC) threat classification.

WhiteHat Sentinel 3.0 enables assessment each time a Web site is changed or updated, and ensures the identification of existing and new vulnerabilities. This is accomplished through a three-step process—scanning, verification and custom testing. As part of this process, WhiteHat integrates expert analysis with proprietary scanning technology which delivers more in-depth results than scanning alone, since many of the most dangerous vulnerabilities can only be detected by this combined process. For more information visit www.whitehatsec.com

SourceGuardian extends power of PHP protection with version 7.0

SourceGuardian announced the launch of Version 7.0, a powerful software protection product securing the Internet-focused PHP programming language. The upgrade of this popular software, originally launched in 2002 and already in use by thousands of customers in more than 53 countries, shows the company's commitment to remaining at the forefront of intellectual property security. Version 7.0 introduces a host of new and enhanced features.

SourceGuardian has become an essential tool for many web developers and programmers, offering the ability to protect their PHP code and therefore their intellectual property. The new version contains all of the same useful features customers have grown to expect, including bytecode encoding, time limiting scripts and locking to specific domain names or machines. For more information visit www.sourceguardian.com



New ultra secure biometric USB 2.0 flash drive



Kanguru Solutions announced the release of the Kanguru Bio Slider II, their new and improved USB 2.0 secure flash drive made complete with the most up-to-date biometric fingerprint technology.

The Kanguru Bio Slider II has taken the hassle out of remembering passwords in order to keep your information secure by using a biometric sensor that will recognize your fingerprint. The drive offers a low maintenance, effortless approach to protecting and storing your data. The built in fingerprint reader allows authorized users access to the encrypted data on the drive with the swipe of a finger. Once the print is confirmed, the user is allowed access to their confidential information. Since an individual fingerprint or thumbprint is the password, the password to your secure drive can never be stolen, forged, or forgotten. For more information visit www.kanguru.com

DevInspect 3.0 with support for Microsoft ASP.NET AJAX extensions

S.P.I. Dynamics announced that the company, in close collaboration with Microsoft, is the first Web application security vendor to provide support for Microsoft ASP.NET 2.0 AJAX Extensions (formerly code-named "Atlas") in its latest release of the company's integrated developer product, DevInspect 3.0. DevInspect is the first security product to analyze and remediate security vulnerabilities in Web applications built using ASP.NET AJAX.



"As technology such as AJAX aggressively evolves to increase the positive experience of users on the Web, Microsoft maintains a focused commitment to improving application security," said Brian Goldfarb, group product manager of the Web Platform and Tools Group at Microsoft Corp. "SPI Dynamics has worked with Microsoft and the ASP.NET AJAX team to raise awareness of application security issues and deliver developer security solutions that assist in the development of more secure software through the Microsoft Visual Studio platform." For more information visit www.spidynamics.com

McAfee SiteAdvisor Plus released



McAfee announced the launch of McAfee SiteAdvisorPlus, the first Web safety tool to actively shield consumers' computers from dangerous Web sites encountered when browsing, searching, and instant messaging or e-mailing. McAfee SiteAdvisor Plus is a premium product that extends the value of McAfee SiteAdvisor, the world's first safe search and browse technology.

McAfee SiteAdvisor Plus takes McAfee SiteAdvisor from a helpful guide to actively enforcing SiteAdvisor's safety ratings and actively shielding computers from interaction with risky sites by checking links in e-mail and instant messages, preventing users from navigating to risky sites and by adding advanced phishing site detection. Consumers may use McAfee SiteAdvisor Plus to complement their existing McAfee products, with non-McAfee security products, or as a stand-alone solution. For more information visit www.mcafee.com


Secure Computing announces Webwasher 6.0

Secure Computing Corporation announced Webwasher 6.0, a new and enhanced version of its award-winning Web Security Gateway, protecting enterprises from inbound and outbound security threats. Webwasher 6.0 marks the initial integration of CipherTrust's TrustedSource™ into the Secure Computing suite of products within 75 days of completing the merger with CipherTrust. In addition, Webwasher 6.0 adds a sophisticated anti-malware engine to proactively protect enterprises from targeted attacks.



Webwasher 6.0 is a web security gateway product now available in three newly released appliances that guards enterprises from a deluge of threats delivered via web traffic, including malware, trojans and phishing attacks, and also ensures that outbound traffic meets corporate compliance requirements. This newest version includes more than a dozen enhancements that strengthen its functionality and improve usability. For more information visit www.securecomputing.com

Check Point VPN-1 UTM Edge 7.0 now available



Check Point Software Technologies announced Check Point VPN-1 UTM Edge 7.0, an upgrade to its branch office offering, delivering powerful new protection with USB Modem Support, Wireless Roaming, and Bridge Mode capabilities. Version 7.0 provides comprehensive and easy-to-use options that help organizations become more secure and ensure constant connectivity to critical information. With USB Modem Support, customers can now have an affordable high-availability option for times when a main Internet link is down. For more information visit www.checkpoint.com

PGP is 15 years old

PGP Corporation salutes the 15th anniversary of PGP encryption technology. Developed and released in 1991 by Phil Zimmermann, Pretty Good Privacy 1.0 set the standard for safe, accessible technology to protect and share online information. Used by millions of users and tens of thousands of companies around the world, PGP technology continues to be recognized for its contributions to the software industry, Internet commerce, and the protection of privacy. Recently, PGP encryption technology was named one of the top 25 most influential products of the first 25 years of enterprise personal computing. For more information visit www.pgp.com



Breach Security releases first appliance with ModSecurity v2.0



Breach Security released the ModSecurity version 2.0 open source web application firewall on an appliance delivering the lowest cost commercial web application firewall available. The ModSecurity Pro M1000 appliance is easy to deploy and manage with rules sets for compliance with Payment Card Initiative v1.1, as well as protection for Microsoft Outlook Web Access.

“We have listened to the community and taken the ModSecurity open source project to an entirely new level—with an appliance that delivers web application security immediately. It is ideal for small-to-medium businesses or large organizations needing just-in-time virtual patching,” said Ivan Ristic, chief evangelist, Breach Security. For more information visit www.breach.com

Visual Studio to be enhanced with Dotfuscator Community Edition

Microsoft Corp. announced that an enhanced version of Dotfuscator Community Edition will be included in the next major release of Microsoft Visual Studio, code-named “Orcas.”



“Protecting intellectual property and preventing application vulnerability probing are two important issues for software developers today,” said Prashant Sridharan, group product manager in the Developer Division at Microsoft. For more information visit www.microsoft.com

Adobe delivers new hosted service for document protection



Adobe introduced Adobe Document Center, a new hosted service that enables knowledge workers to better protect, share and track the usage of Adobe PDF, Microsoft Word, and Microsoft Excel documents as part of day-to-day communications and collaboration. This new, easy-to-use, web-based service gives business professionals the power to grant and dynamically revoke access to documents distributed inside or outside the firewall, as well as audit actions such as opening, adding comments to, or printing those documents.

Adobe Document Center is designed for the professional who shares or publishes business-, time- or version-sensitive documents. Whether it's an independent graphics designer submitting designs for client review, or a legal practice exchanging sensitive files with clients, users can customize access settings, closely audit usage of their documents, and retain control over the files regardless of where they travel. Users also have the ability to set expiration dates on documents, supersede an older version once a new version is distributed, and revoke access after distribution. They even have the ability to track who has received the documents and what recipients have done, or attempted to do, with the files.

eEye Digital Security introduces Blink Personal freeware

eEye Digital Security announced the release of Blink Personal, a free version of its award-winning Blink endpoint security technology, developed for non-commercial users. Blink Personal is the first free security product available to consumers to combine multiple layers of technology that protect against identity theft, worms, trojan horses and other attack methods hackers use, into a single agent that is unobtrusive, integrated and deeply-layered with security functionality.



Deployed as a software agent on a Windows-based desktop PC or laptop, Blink Personal leverages multiple layers of protection—more so than any other endpoint security product—to shield individual digital assets from attacks and keep systems up and running. For more information visit www.eeye.com

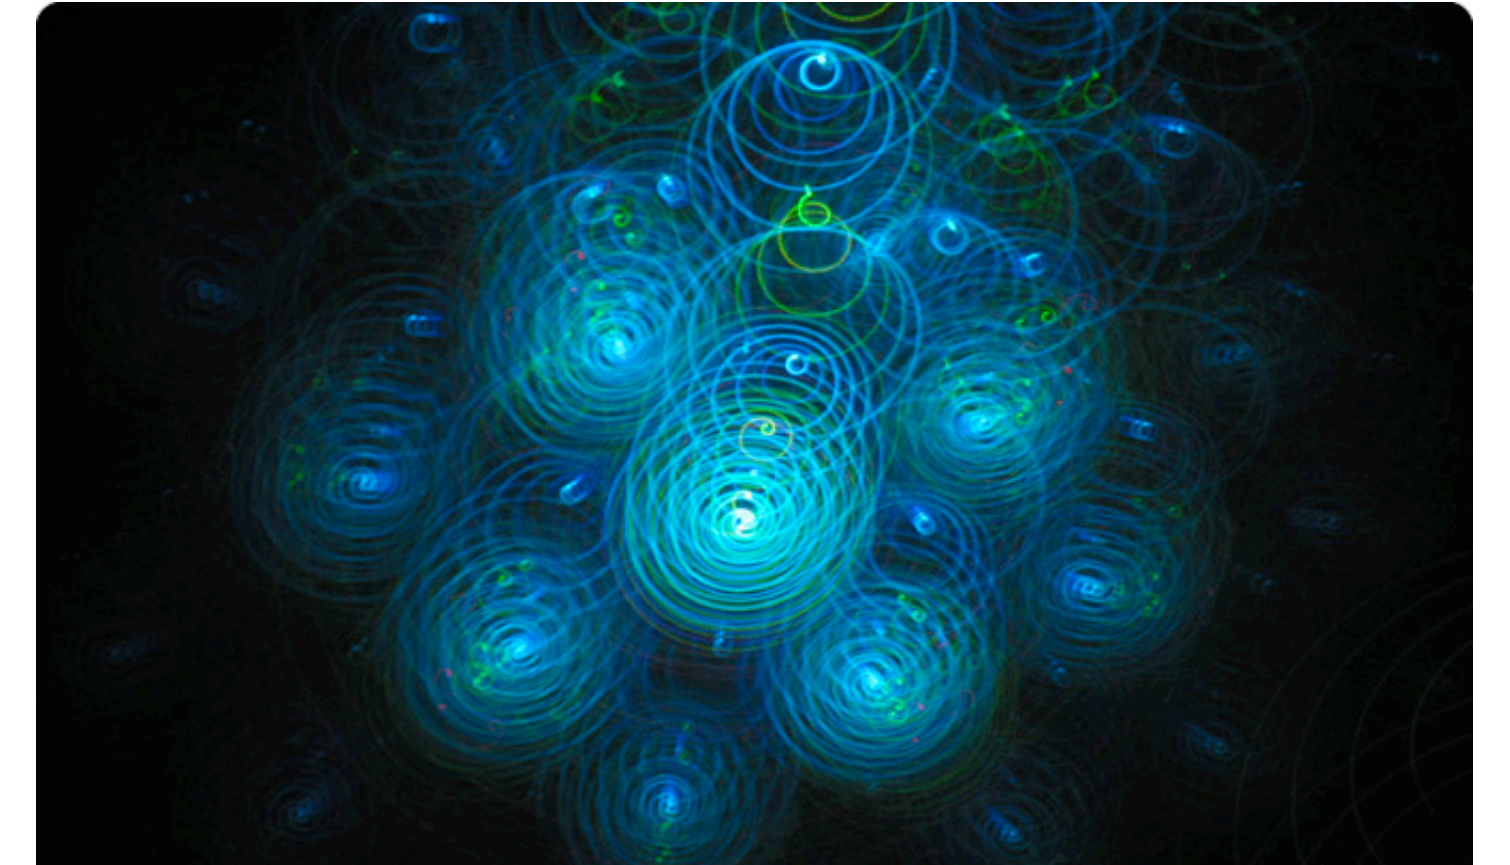
Anti-Phishing Working Group announces the Internet Crimeware Report



The Anti-Phishing Working Group has issued a joint report with the Department of Homeland Security and SRI International on the role of crimeware in enabling new forms of financial crime on the public Internet. The report is titled “The Crimeware Landscape: Malware, Phishing, Identity Theft and Beyond”.

APWG data from the 12 months between May 2005 and May 2006 tells the story of runaway proliferation of crimeware. In that time frame, the number of unique applications for password stealing that were detected in a single month grew from 79 to 215, almost tripling in detected frequency. The number of URLs employed by criminals to spread crimeware expanded at around twice the rate of crimeware code development, however, rising from 495 detected URLs in May 2005 to 2100 in May 2006 after peaking at 2683 in April, 2006.

Read the report at www.antiphishing.org/reports/APWG_CrimewareReport.pdf



Effectiveness of security by admonition: a case study of security warnings in a web browser setting

By Christian Seifert, Ian Welch and Peter Komisarczuk

Security warnings seem to be a predominant way to bridge the gap of providing rich, but potentially insecure, functionality and providing security. In this study, we investigate the effectiveness of so-called security by admonition. We present users with a web-based survey that requests the installation of a potentially insecure ActiveX component. We show that the security warning deters users from fulfilling the insecure installation request, but is ineffective in preventing it.

1. Introduction

Many vendors consider providing security as part of their products and services an important element of their business. With Bill Gates's Trustworthy Computing directive in 2002, Microsoft has been putting itself at the forefront of computer security. Their Windows XP Service Pack 2 and recently released Internet Explorer 7.0 RC1 focus on enhancing security with more secure default settings, security software patches, and new security features, such as the phishing filter.

Many of these new security features, and this is not endemic to Microsoft products, leave

the final security-relevant decision to the end user. This stems from a conflict between providing security and usability, in which security usually hinders and usability usually assists the user in achieving a task. Figure 1 illustrates one approach to overcome this conflict. The security policy represented by the box covers a wide spectrum of actions that the user may find either acceptable or unacceptable. As soon as the user performs an action that might be unacceptable to the user, but is permitted by the security policy, user confirmation of the action is required.

This principle, called security by admonition, leaves the final decision to the user.

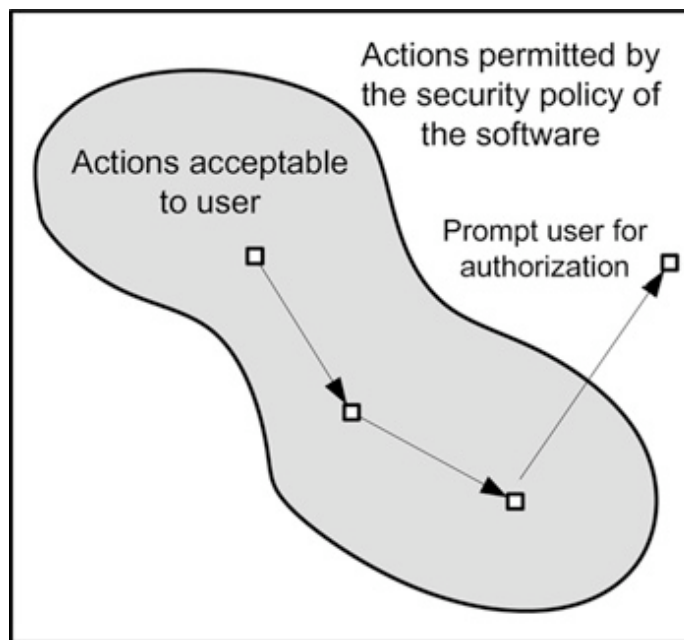


Figure 1: Security By Admonition.

Some examples of these warning messages are a firewall popping up a dialog about whether process xyz to access server aka on port abc should be allowed or denied, the Internet Explorer phishing filter warning the user that the page is potentially hazardous, and the Firefox browser asking for escalated privileges for signed JavaScript code. Security by admonition relies on the user's general knowledge about security and computing to make good decisions.

In this article, we investigate the effectiveness of such security warnings and dialog boxes via a case study on Microsoft Internet Explorer, which seems to be a prime example of the security by admonition approach. We present users with a web browsing situation that is potentially hazardous and that causes a security warning prompting the user to decide how to proceed. Specifically, we invited users to access a web page that contained a signed ActiveX control. Once the user accessed this web page, the browser displayed a security warning asking for user confirmation prior to installing and running the component. We tracked execution of the ActiveX component to determine the decision made by the user. This work allows us to answer the question of whether security warnings protect users from potential security threats.

The remainder of the article is structured the following way. Section 2 provides background

information on ActiveX controls and a description of the relevant browser behavior regarding pages that contain ActiveX controls. Section 3 describes our experiment and survey setup. In section 4 we present the data analysis and results and conclude in section 5.

2. Background

In 1996, Microsoft introduced ActiveX controls. They are lightweight programs that can be placed inside and distributed as part of a document. ActiveX controls build on top of Object Linking and Embedding technology (OLE) that allows users to place documents created in one application within documents of other applications. OLE, for example, allows placing a Microsoft Excel spreadsheet inside a Microsoft Word document. ActiveX controls, however, are not documents. They are programs that expose a defined interface that can be interacted with. The underlying technology is called Component Object Model (COM), which is a foundation of Microsoft technology. Many applications adhere to COM and therefore expose an interface to the outside world. For example, Microsoft Word's interface allows for customizations and extensions via COM.

ActiveX controls can be distributed as part of a web page with default support by Microsoft Internet Explorer.

Once a web page with its ActiveX control is retrieved, the ActiveX control is able to execute with the same permissions as the browser, which equates to the permissions of the user. As such, the control, among other things, has read/write access to files the user has access to, and can establish network connections. If the user has administrator privileges, a setup commonly encountered with home users, the ActiveX component has unrestricted access rights and can go as far as modifying the underlying operating system.

Authenticode, a Microsoft technology for digitally signing code, is the primary security mechanism for ActiveX controls. Digitally signing a program is a matter of obtaining a code signing certificate from a recognized certificate authority and using this certificate to sign the component. Code signing ensures the control's authenticity and integrity. Authenticity specifies where the code came from whereas integrity verifies that the code has not been altered since its publication. It does not, however, indicate whether the control is safe, so any signed ActiveX component could potentially be a security hazard. (Additional informa-

tion on the ActiveX security model can be found in Robert Stroud's technical report.)

The existence of an ActiveX component's signature influences the browser's behavior. Unsigned ActiveX components are disallowed to be downloaded on Microsoft Internet Explorer version 6.0 and higher, whereas a signed ActiveX component with a certificate from a recognized root certificate authority requires user confirmation prior to being downloaded and executed. Microsoft Internet Explorer 6.0 displays a dialog box as shown in figure 2. Please note that the default selection proposed by the browser is not to install and run the ActiveX component. With Service Pack 2 version of this browser, the security warning moves from a dialog box to a security bar as shown in figure 3. The security bar is less intrusive and allows the user to continue to interact without reacting to the warning. Once the user attends to the warning and chooses to install the ActiveX component, a dialog box is displayed to ask for confirmation to install the software as shown in figure 4. This behavior of Microsoft Internet Explorer 6.0 SP2 remains identical with Microsoft Internet Explorer 7.0 RC1.

Please install the activeX component before proceeding with the survey. Your above.

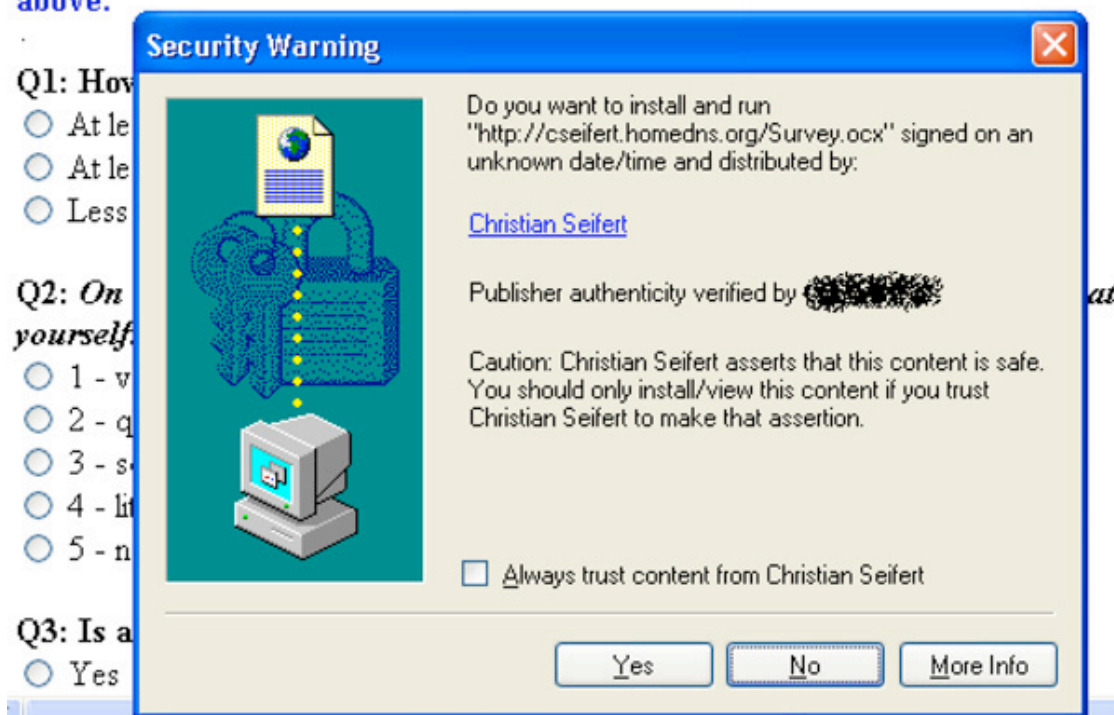


Figure 2: Internet Explorer 6.0 - security warning.

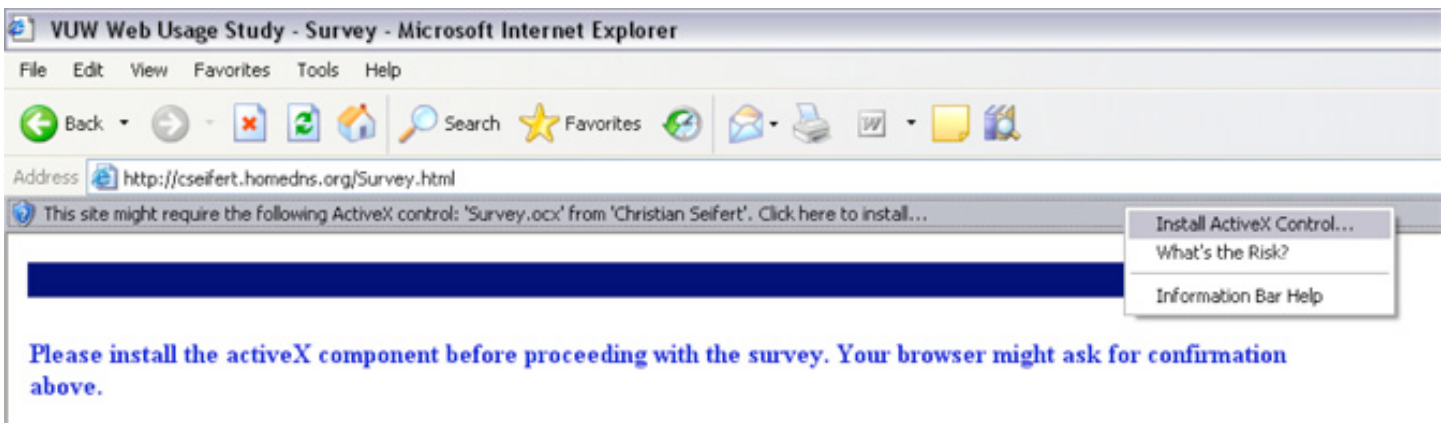


Figure 3: Internet Explorer 6.0 SP2 - security warning bar.

3. Experimental Setup

In order to learn about the decisions users make in response to these security warnings, we needed to present the users with a situation in which a warning is displayed and ignoring it could have security-threatening implications. We chose to present users with a web page in which a signed ActiveX component was embedded and track its execution, indicating the user ignored the resulting security warning and proceeded with the installation of the ActiveX component.

The Authenticode security mechanism did not hinder creation of the signed ActiveX component. The ActiveX component was digitally signed with a certificate obtained from a recognized root certificate authority. The certificate authority issued the code signing certificate free of charge within two working days. For identity verification, the certificate authority accepted a faxed New Zealand driver's li-

cense. The issued certificate was only valid for 90 days, but more than sufficient for the purposes of this experiment. We mention this to demonstrate that a signed ActiveX control could easily be created by a user with malicious intentions.

We embedded the ActiveX component in a simple web survey claiming to obtain information about web browsing behavior. The survey was entirely used to divert attention from the security relevant decision, as security is usually a secondary concern to the user. This setup was supposed to simulate a real world setting in which the user is attempting to complete a non-security-related primary task. The survey consisted of three web pages. The first page presented the information sheet in compliance with the requirements set forth by Victoria University's Human Ethics Committee. After the participant read the information sheet, they could proceed to the actual survey through a web link.

Please install the activeX component before proceeding with the survey. Your browser n above.

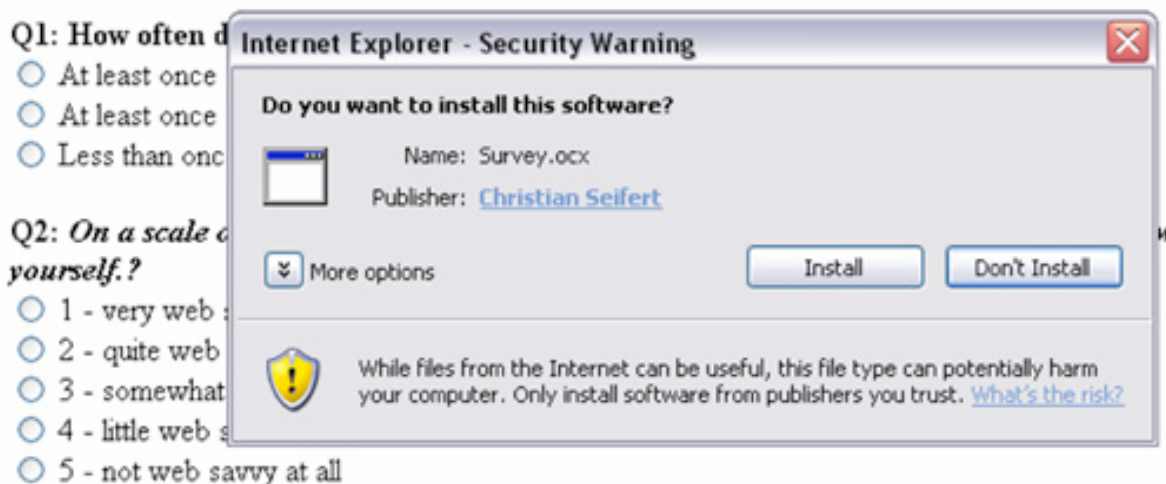


Figure 4: Internet Explorer 6.0 SP2 - security warning.

The survey consisted of seven simple questions related to web browsing behavior. Once the survey page was opened, the behavior slightly differed depending on the browser the participant used. In case of browsers other than Microsoft Internet Explorer, the user was simply presented with the survey. In cases where Microsoft Internet Explorer was used, the participant was instructed to install the ActiveX component before proceeding with the survey. Depending on the browser's setting to deal with ActiveX components (see section 2, the participant was prompted to confirm installation. Once the ActiveX component was installed and run, a new window popped up informing the participant that the ActiveX component was run and that they could now close the window. This was our way of tracking whether the component was executed. Independent of the participant's decision on whether to install the ActiveX component, they were able to complete the survey. Upon completion of the survey, the participant was presented with debriefing page explaining the true nature of the study.

We sent invitations to participate in this study to various non-security related web forums and news groups. Invitations were sent to English-speaking groups only. The groups were selected based on high frequency and membership numbers. We took care not to post to special interest groups, such as soc.retirement. We sent invitations to the following newsgroups: alt.society.zeitgeist, alt.friends, alt.philosophy, misc.legal, nz.comp, nz.general, alt.internet, uk.misc, aus.general, aus.computers, misc.consumers, soc.misc,

misc.education, alt.education as well as a YAHOO! forum on internet, psychology and education and an MSN forum on computers, technology and internet. The link to the web survey that was included in the invitation contained a tracking parameter, which allowed us to link the responses back to the invitation. The tracking parameter was included to disqualify responses in case the true nature of the study was revealed or warnings about the ActiveX component embedded in the survey were communicated in the relevant forum. This happened three times, and the corresponding results were discarded from the study.

We did not take any steps to make the invitation or the survey itself seem to come from an authorized or legitimate source that would influence the trust relationship of potential participants. While we do state that this is a study performed by a PhD student at Victoria University of Wellington, New Zealand, neither the email address used to invite participants nor the web site hosting the survey is sourced by the University. As such, there was no way to discern whether the invitation and survey did in fact come from a PhD student or from a potential imposter. The actual web survey was not created in the look and feel of the University. However, due to regulations of the Human Ethics Committee, the site did have to bear a logo of the University as well as a reference number of the Human Ethics Committee application. It seems that no participants contacted the Human Ethics Committee or any of the researchers at the University to verify the legitimacy of the study.




Figure 5: Internet Explorer breakdown.

4. Data Analysis and Results

As users participated in the study and accessed our survey pages, our web server tracked a unique identifier of the participant (IP Address), the pages accessed, as well as the browser and operating system used to access the web survey. Prior to analysis, we disregarded any data from participants that did not proceed to the survey from the initial information page or participants that withdrew from the study after completion. We disregarded any data that originated from invitation posts in which the true nature of the study was discussed or warnings issued about the ActiveX control that was embedded in the survey.

A total of 114 users participated in the study. 65 participants used a version of Microsoft Internet Explorer to access the survey. Figure 5 shows the breakdown of Microsoft Internet Explorer versions used. 49 participants used a browser other than Microsoft Internet Explorer, primarily Firefox.

The security warning about the ActiveX component is displayed as soon as a participant using Microsoft Internet Explorer accesses the survey page. Figure 6 shows that Microsoft Internet Explorer users seem to be more likely to leave the survey altogether (13 out of 65) than users with other browsers (4 out of 49). According to the chi-square test, there exists a statistical significance between the two groups (1-DOF, chi-square = 11.45, $p < 0.001$) indicating that the security warning displayed for

Microsoft Internet Explorer users deterred participants from completing the survey. They simply left the web site or closed the browser.

Of these 65 participants that used Microsoft Internet Explorer, 11 ignored the security warning and installed the ActiveX component. 3 of these 11 participants were using Microsoft Internet Explorer 6.0. Recall that the behavior of Microsoft Internet Explorer 6.0 causes a simple popup dialog box to appear in which the user has to explicitly select the installation of the ActiveX component, as shown in figure 2. For the remaining 8 participants using Microsoft Internet Explorer 6.0 SP2 or higher, the user had to click on the security bar, select "install component" and confirm installation via a popup dialog box, as shown in figure 3.

First, we calculated the confidence interval for the proportion of respondents to the security warning. We compare this confidence interval for the behavior of the 65 participants against a group without any security warning for whom visitation to the survey page would lead to automatic installation of the ActiveX component.

This, for example, is the case with earlier versions of Microsoft Internet Explorer. For this group, all or 100% would install the component. This test was designed to determine whether the presence of the security warning has a statistically significant effect on the proportion of users who installed an ActiveX component.




Figure 6: Survey completion.

The confidence interval shows that the security warning does deter users from installing the ActiveX component (p-value;0.05) since the 95% confidence interval does not contain 100% (lower limit = 70.84%, upper limit = 90.28%).

Second, we calculated the confidence interval against a group for which ActiveX components are disallowed in general leading to no installation of the component, such as users that use a non-Microsoft browser. For this group, none or 0% would install the component. This test was designed to determine whether the warning leads to secure decisions and the prevention of installations. Again, we are 95% confident that the security warning compared to a default deny decision does encourage insecure actions to take place (p-value;0.05) since the 95% confidence interval does not contain 0% (lower limit = 7.81%, upper limit = 27.82%). These statistical tests assume that the groups of users are similar and that their decisions on whether to install or not install the ActiveX component are driven by the display of the security warning.

5. Conclusion

In this article, we investigated whether security warnings inform users about security threats and successfully deter users from encountering these threats. We have demonstrated that security warnings seem to deter users from installing malicious code when browsing a web site, such as that which might be contained in an ActiveX component. However, the fact that 16.92% of participants in our study did install the ActiveX component shows that this is not enough.

By inviting study participants via newsgroups and forums, we likely attracted a pool that is overall more technology savvy than average. This may have skewed the installation percentage to a lower value. Alternately, users choosing to participate in a study solicited via newsgroup postings might have skewed the installation percentage to a higher value, if they are overall more susceptible to solicitation and more trusting to ignore security warnings. The fact that 57.01% of the participants used Microsoft Internet Explorer indicates that a representative browser distribution was present with the study as it falls within the

browser statistic published by W3 Schools for the month in which the study was conducted.

There may be several reasons why the ActiveX component was installed in 16.92% of the cases in our study. The first explanation, which is in line with a study performed by Wu on the effectiveness of security toolbars to prevent phishing attacks, is the fact that security is not the primary concern of the user. Security is important, but secondary to the actual goals of the user performing a task. If security blocks the goals, it is likely to be ignored. In our case, we asked a user to install an ActiveX component in order to complete a survey, and users might have assumed it was essential part of the survey.

The second reason why users might have ignored the security warning was a lack of knowledge regarding its possible implications. The security warning displayed does not contain enough information about the implications of the user's action. Microsoft Internet Explorer 6.0 simply states that the authenticity has been verified and that the author of the component asserts that it is safe. In Microsoft Internet Explorer 6.0SP2 and higher, the initial warning does not even contain any warning signs, but rather just states that the site might require the ActiveX component. Upon the user choosing to install this component, a security warning asks for confirmation to install the software. The implications are unknown and users are not likely to know that ActiveX component have unrestricted access rights. They might assume that the browser restricts the component from performing unacceptable behavior, which goes back to figure 1.

Solutions to preventing insecure actions are multi-fold. One could reduce the impact of a user's insecure decisions, for example through sandboxing. Explaining the implications of the user's actions in terms that are understandable would assist users in making good decisions. In our case, a warning that states that ActiveX has access to the user's personal files would be one example. Another solution would be to adjust the security policy based on the user's action, called security by designation. With such an approach, ActiveX components could be disallowed by default. However, if the browser detects a relationship to a site (e.g. through an existing bookmark),

the browser could adjust its security policy and prompt with a security warning. A simple default deny policy would be another option. While this might not be feasible in certain settings, we believe it would be appropriate in our case of an ActiveX component. We do not perceive the need to allow applications to be distributed via a browser. If rich applications are required, the user can be asked to download a program and install it, which would assist in setting the expectations of the user with respect to the program. Finally, providing a secure alternative to achieving a primary goal, could lead to users paying attention to the matter of security. In their study on phishing attacks, they suggest that the browser detects the phishing site, determines the real site and then forwards the user to the real site instead of blocking or displaying a warning about the phishing site.

Some of the solutions are already supported by existing products. Sandboxing is rather an old technology for browsers and is supported by Java Applets as well as Microsoft Internet Explorer 7.0 in the new Microsoft Vista operat-

ing system. Several add-on products exist for browsers, such as GreenBorder, which enforces a stricter security policy than the one provided by default. Security by designation products are appearing with CapDesk and Polaris, which start applications with the principal of least authority and dish out additional permissions that are inferred by the user's actions. Group policies, which are provided with many applications and operating systems, allow administrators/users to overwrite insecure default settings. For example, the default settings of the web browsers at Victoria University of Wellington do not allow ActiveX components to be downloaded.

Products and solutions do exist for certain circumstances. However, they do not seem to be widely adopted or delivered as part of a standard computing environment. Home users, the ones that are probably most vulnerable, need to be protected by standard restricted policies. We appeal to vendors to consider these points and deliver security functionality as part of their products to end consumers in the future.

Christian Seifert is a PhD student at Victoria University of Wellington and an active member of the New Zealand Honeypot Alliance. Christian's research interest lies in means of identifying malicious servers with client honeypots. Christian is the main developer of the low interaction client honeypot "HoneyC" and the sponsor of the high interaction client honeypot "Capture - HPC". Christian can be reached at cseifert@mcs.vuw.ac.nz.

Dr. Peter Komisarczuk teaches and researches at Victoria University of Wellington. Dr Komisarczuk has a technical background in network engineering and prior to Victoria University he worked for many years in the telecoms industry.

Dr. Ian Welch teaches and does research at Victoria University of Wellington. Dr Welch has a technical background in dependability and software engineering.





Interview with Kurt Sauer, CSO at
By Mirko Zorz



As Chief Security Officer at Skype Technologies, Kurt Sauer is focused on delivering trusted communications services via Skype's platforms. Both the software delivered to customers and the internal infrastructure needed to provide Skype's services are developed with an eye toward design, implementation auditing, and software life-cycle management.

Before joining Skype in 2004, Mr. Sauer was a Principal Network Security Architect for Sun Microsystems at its European research laboratory. Sauer is a member of the ACM, IEEE, Mensa and the Forum of Incident Response and Security Teams (FIRST). He holds a bachelor's degree in Computer Engineering from Texas A&M University and is fluent in English and French.

What has been your biggest challenge as the CSO of Skype?

The most difficult challenge has been keeping up with the diversity and speed of the development initiatives going on in the company. Skype is growing by leaps and bounds – it still takes a finite amount of time to investigate the nuances of the interaction among new innovations.

I remember the story told by Frederic Brooks about the development of early operating systems, which basically distills the idea that

"adding people to a problem does not necessarily solve it faster." And this is equally true at Skype – it's not having a lot of people that counts, it's having bright and adaptable people that's important.

How many active users does Skype have?

Skype currently has 136 million registered users today – in Q3 WE added nearly 23 million users, or about 250,000 new users per day – spanning more than 200 countries.

With the constant evolution of threats, what kind of technology challenges does Skype face?

One of the biggest potential threats to Skype is from attempts to conduct identity theft. Criminals and hackers are using increasingly sophisticated and targeted attacks against computer users worldwide to gain access to end-users' service and banking accounts. Internet users worldwide continue to fall prey to fake e-mail or so-called "phishing" attacks, supplying thieves with opportunities to install keystroke loggers and other malware on their computers. Skype works closely with eBay

and PayPal, as well as with other industry partners, to identify and counter these and any other kinds of attacks.

How does Skype's security compare with that of other VoIP systems?

Skype uses a sophisticated system of standards-compliant cipher and digital signature systems to preserve the security and to ensure the integrity and authenticity of the call from end-to-end. Most other VoIP systems provide no encryption or authenticity controls over the call, which puts Skype in a security leadership class of its own.

Most other VoIP systems provide no encryption or authenticity controls over the call, which puts Skype in a security leadership class of its own.

Many argue that the adoption of VoIP brings together a whole new set of security risks and problems. Do you agree? If you do, what can be done to mitigate those risks?

Most of the problems identified in the area of VoIP have to do with the complexities of inter-connecting VoIP switches and other hardware components in an enterprise configuration. In addition to this, there have been persistent arguments that VoIP is insecure because the vast majority of VoIP systems do not provide any level of encryption by default for their users.

Efforts in the VoIP industry to use encryption more pervasively, to reduce the risk of equipment configuration errors, and to reduce the amount of infrastructure components needed to deploy the service will help. Skype has a distinct advantage in this area because its peer-to-peer design eschews hardware switches, thereby eliminating the risk of mis-configuration, and uses only encrypted communications links.

What is your general strategy for making Skype more secure?

Keeping Skype simple to use and retaining a public key infrastructure-based (PKI) authentication system are the keys to ensuring continued security for Skype.

In the old days it was all about phreaking, nowadays the term of VoIPhreaking is making its way into the news. Have you had any experience with it or is it just media hype?

The term "phone phreaking" predates "malicious hacking" and the myriad of Internet-age terms that have come to represent the analogue of phone phreaking in the modern age. By their very nature, all security systems pose a challenge to those who perceive themselves as being on the outside of the barrier.

What I think is the biggest sea change in telecommunications security is in the area of motivation. Phone phreakers were, by and large, interested in the security of telecommunications systems per se; it was viewed by the phreakers as a mostly intellectual pursuit.

Today, however, we see a bifurcation of objectives: while some continue their pursuit – rightly or wrongly – for purely intellectual challenge, the commercial benefits in the areas of unsolicited commercial calling (spam messaging) and in industrial espionage are perceived to be so great that very well-financed and sophisticated attacks are appearing at an alarming rate on the Internet. This is not just a risk for VoIP, but for the general computing milieu of which VoIP is merely one part.

What challenges do you face in the marketplace? What do you see as your advantages?

While Skype is a leader in the area of peer-to-peer communications and in converged messaging, there is always the possibility of becoming obsolete due to competition. The challenge we face is partly organizational – making sure we use our resources effectively and remain lean – and partly technological, ensur-

ing that our developments are relevant, innovative and easy-to-use. I suppose that the challenges we face in the marketplace are the same as any other new company: gaining customer acceptance and focusing on delighting our users every single day.

What are your future plans? Any exciting new projects?

That would be telling...

HNS SECURITY SOFTWARE DATABASE

Get the largest selection of the best security software for Windows, Linux, Mac OS X and Windows Mobile platforms.

20 CATEGORIES
2.4 MILLION DOWNLOADS SO FAR

net-security.org




Latest additions to our bookshelf

Microsoft Windows Vista Unveiled

by Paul McFedries

Sams, ISBN: 0672328933



Microsoft Windows Vista Unveiled is an in-depth exploration of the public release Beta 2 version of Windows Vista. Whether you're just planning ahead for a future upgrade or running Beta 2 already, Microsoft Windows Vista Unveiled takes you on a detailed tour of the new and improved technologies, features, tools, and programs that were added to Vista. Some of the interesting sections include a review of Vista's new performance and stability features and a critical look at the beefed-up security features, including Windows Defender, the new Firewall, and User Account Control.

Configuring IPCop Firewalls: Closing Borders with Open Source

by Barrie Dempster, James Eaton-Lee

Packt Publishing, ISBN: 1904811361



IPCop is a powerful, open source, Linux based firewall distribution for primarily SOHHO etworks. This book is an easy-to-read guide to using IPCop in a variety of different roles within the network. It first covers basic IPCop concepts, then moves to introduce basic IPCop configurations, before covering advanced uses of IPCop. To book's target market is basically anyone interested in securing their networks with IPCop - from those new to networking and firewalls, to networking and IT Professionals with previous experience of IPCop.

Hacking the Cable Modem: What Cable Companies Don't Want You to Know

by DerEngel

No Starch Press, ISBN: 1593271018




In the beginning there was dial-up, and it was slow; then came broadband in the form of cable, which redefined how we access the internet, share information, and communicate with each other online. Hacking the Cable Modem goes inside the device that makes Internet via cable possible and, along the way, reveals secrets of many popular cable modems, including products from Motorola, RCA, WebSTAR, D-Link and more. The book features step-by-step tutorials with easy to follow diagrams, source code examples, hardware schematics, links to software, and previously unreleased cable modem hacks.

Mac OS X Internals: A Systems Approach

by Amit Singh

Addison-Wesley Professional, ISBN: 0321278542



Mac OS X Internals: A Systems Approach focuses on dissecting the internals of the system. It provides valuable information on learning the roles of the firmware, the bootloader, the Mach and BSD kernel components (including the process, virtual memory, IPC, and file system layers), the object-oriented I/O Kit driver framework, user libraries, and other core pieces of software. The book is all about technical aspects of OS X and is full of useful information and programming examples. It also covers several key areas of the Intel-based Macintosh computers.

Own Your Space: Keep Yourself and Your Stuff Safe Online

by Linda McCarthy

Addison-Wesley Professional, ISBN: 0321426428




McCarthy's book is mainly targeted towards teenagers and average Internet users. You have been immersed in the technology since preschool. You download music, Google your homework, and constantly IM your friends. You check your email before dinner, tweak your MySpace page, and bypass that hardcopy diary for your own 31337 space in the blogosphere. While you're doing that, you also need to protect yourself. This book is about keeping safe - protecting your data, your identity, and yourself without giving up all the great stuff the Net puts at your fingertips.

Practical Guide to Computer Forensics

by David Benton and Frank Grindstaff


BookSurge Publishing, ISBN: 1419623877



Practical Guide to Computer Forensics discusses the history of computer forensics, along with policies, standard operating procedures and legal considerations. The authors, both experts in the field, delve into what makes a qualified computer forensics specialist and what is the best way to find one. An ideal read for anyone needing an extensive overview of computer forensics, Practical Guide to Computer Forensics is also an intriguing look into our increasingly technical world.

Web 2.0 defense with AJAX fingerprinting and filtering

By Shreeraj Shah



Ajax is fast becoming an integral part of new generation Web applications known as Web 2.0 applications. This evolution has led to new attack vectors coming into existence around these new technologies.

To combat these new threats one needs to look at different strategies as well. In this paper we shall look at different approaches and tools to improve security posture at both, the server as well as browser ends. Listed below are the key learning objectives:

- The need for Ajax fingerprinting and content filtering.
- The concept of Ajax fingerprinting and its implementation in the browser using XHR.
- Processing Ajax fingerprints on the Web server.
- Implementation using `mod_security` for Apache and `HTTPModule` for IIS (.NET)
- Strengthening browser security using HTTP response content filtering of untrusted information directed at the browser in the form of RSS feeds or blogs.
- Web application firewall (WAF) for content filtering and defense against Cross-Site Scripting (XSS)

Requirement for Ajax fingerprints and filtering

Ajax is being used very liberally in next generation Web applications, forming an invisible layer in the browser's transport stack and bringing to the fore numerous browser-related attacks, all centered around Ajax. Although Ajax applications hold a lot of promise, there are loopholes being exploited by viruses, worms and malicious attackers in Web 2.0 applications that need to be looked at a little more closely. Ajax hides a lot of server-side critical resources due to its calling mechanism, bringing in sloppiness in coding patterns and fueling vulnerabilities in the server-side application layer as well. Untrusted resource processing from blogs, feeds and mash-ups are making Ajax vulnerabilities relatively easy to exploit. In such situations Ajax request and response fingerprinting and filtering mechanisms can enhance the security posture of

Web applications.

Web 2.0 applications have a special set of resources that are accessed by the web browsers over Ajax calls using the `XMLHttpRequest` (XHR) object. Resources can be grouped into two broad spaces – one with “Ajax-only” ac-

cess and other non-Ajax (traditional) resources. In the application architecture, one can wrap security around Ajax resources by creating a separate virtual sandbox for all incoming and outgoing Ajax calls as shown in Figure 1.0.

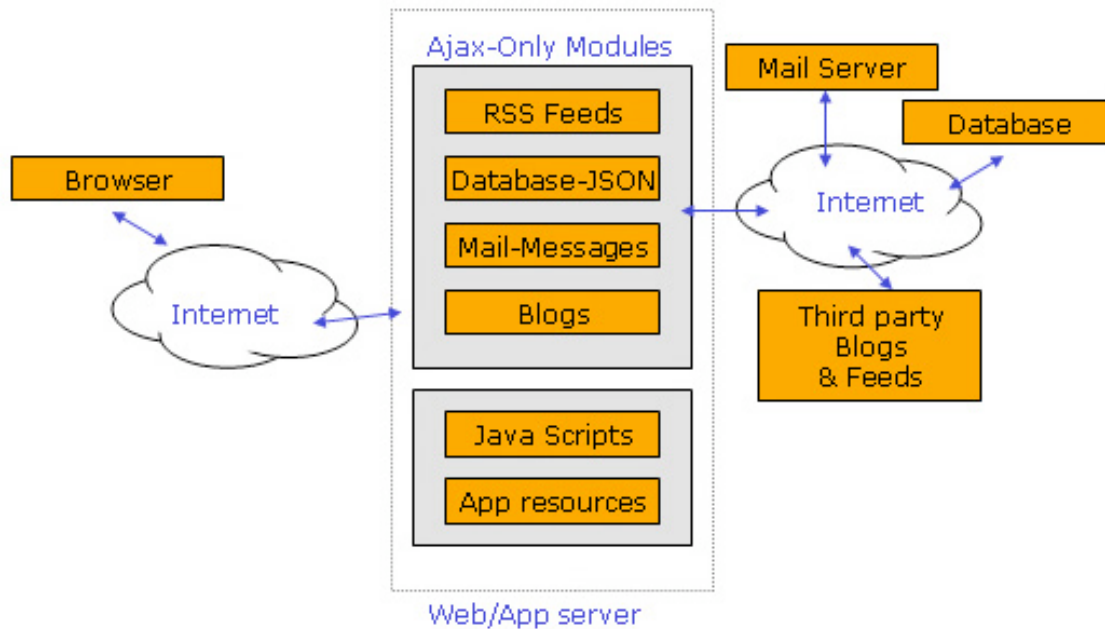


Figure 1.0 – Ajax sandbox on the server-side.

“Ajax-Only” modules access third-party resources such as blogs and feeds using their own proxy code. These proxies are essential since direct cross-domain access with Ajax is not possible. However, JavaScript scripts residing in the browser can access database streams directly over JSON or a JavaScript array as shown in Figure 1.0. Ajax resources serve a lot of untrusted and unfiltered information to the browser, in the process leaving an end-users’ browser vulnerable to several client side attacks such as XSS and XSRF.

To provide better security framework to both applications and browsers, Ajax resources on the server-side can be defended by applying Ajax fingerprinting methods. The key question, however, that we need to ask is, “is there a way to identify an HTTP Ajax call?”. It would be easy to build several security controls for both application and browser security provided an Ajax call can be fingerprinted. This is the topic of discussion in this article. Applying firewall rules for incoming traffic is always important, but in an Ajax-Only frame-

work, filtering outgoing traffic is of greater importance given the fact that the application serves untrusted information to the browser in the current application DOM context. Put simply, if a DOM-based XSS attack is successful, the client session can be hijacked with ease. This application may be running a banking system, financial transactions, mailing system or blogs. Losing session information can result in financial or non-financial losses.

Implementing Ajax fingerprinting – Adding extra HTTP headers

To implement Ajax fingerprinting, we need to first identify the `HTTP GET` and `POST` request structure for Ajax calls. Figure 2.0 illustrates a simple example of an Ajax call. The browser loads the “news.html” page. Clicking the link “Get today’s headline”, will make a backend Ajax call to the server requesting for the “/ajax-only/headline” resource. The code snippet in Listing 1.0 gets executed by the browser when a “click” action occurs. i.e. the `getHeadline()` function is invoked.

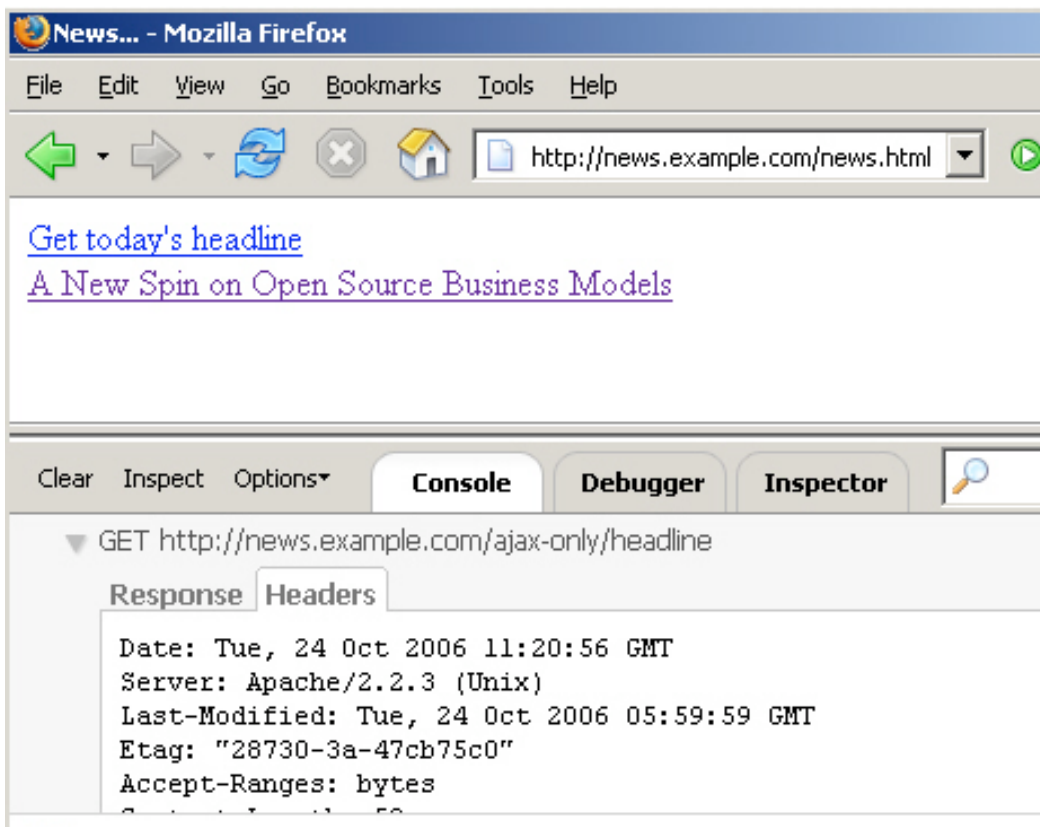


Figure 2.0 – Sample Ajax call.

```
function getHeadline()
{
    // Intializing the XHR Object
    var http;
    if(window.XMLHttpRequest) {
        http = new XMLHttpRequest();
    } else if (window.ActiveXObject) {
        http=new ActiveXObject("Msxml2.XMLHTTP");
    }
    if (! http) {
        http=new ActiveXObject("Microsoft.XMLHTTP");
    }
    }

    // Building a request
    http.open("GET", "/ajax-only/headline", true);

    // Getting ready for response processing
    http.onreadystatechange = function()
    {
        if (http.readyState == 4) {
            var response = http.responseText;
            document.getElementById('result').innerHTML = response;
        }
    }

    //Sending Async request on the wire
    http.send(null);
}
}
```

Listing 1.0 - The getHeadline() function.

The following GET request will be sent to the Web server:

```
GET /ajax-only/headline HTTP/1.1
Host: news.example.com
User-Agent: Mozilla/5.0 (Windows; U; Windows NT 5.0; en-US; rv:1.8.0.6) Gecko/20060728
Firefox/1.5.0.6
Accept:
text/xml,application/xml,application/xhtml+xml,text/html;q=0.9,text/plain;q=0.8,image/png,*/*;q=0.5
Accept-Language: en,en-us;q=0.5
Accept-Encoding: gzip,deflate
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
Keep-Alive: 300
Connection: keep-alive
```

A cursory glance at the request gives no indication that the request is made by the XHR object from within the browser. It is possible to

add an extra header to the HTTP request as per XHR's methods that would aid in identifying and fingerprinting the Ajax call.

```
// Building request
http.open("GET", "/ajax-only/headline", true);
http.setRequestHeader("Ajax-Timestamp", Date());
```

Modify the code snippet in Listing 1.0 to attach an "Ajax-Timestamp" header to the outgoing HTTP requests. By using the output of the `Date()` function and a browser fingerprinting technique, we can identify browsers as well.

Now, click the same link again.

This is the GET request that will be generated on the wire:

```
GET /ajax-only/headline HTTP/1.1
Host: news.example.com
User-Agent: Mozilla/5.0 (Windows; U; Windows NT 5.0; en-US; rv:1.8.0.6) Gecko/20060728
Firefox/1.5.0.6
Accept:
text/xml,application/xml,application/xhtml+xml,text/html;q=0.9,text/plain;q=0.8,image/png,*/*;q=0.5
Accept-Language: en,en-us;q=0.5
Accept-Encoding: gzip,deflate
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
Keep-Alive: 300
Connection: keep-alive
Ajax-Timestamp: Tue Oct 24 2006 17:37:46 GMT+0530 (India Standard Time)
```

Look closely at the GET request. From this GET request we can determine the fingerprint

of the Ajax call. On the server we receive the following timestamp header:

```
Ajax-Timestamp: Tue Oct 24 2006 17:37:46 GMT+0530 (India Standard Time)
```

This fingerprinting technique helps in determining the type of client code that has sent this request. It is possible to lockdown resources for just the right client on the server-side as well. This type of header is harder to

add by automated crawlers and bots since the logic and calls need to be understood first.

Consequently, automated attacks on your Ajax resources can be avoided.

Fingerprinting is just a starting point for securing Ajax resources. It is possible to build a security control around this extra header mechanism. You can add JavaScript libraries in your client-side code and use MD5 hashing and other encryption methods. The XHR object controls the POST method along with buffer that the client sends to the server. A secure tunnel can be built over HTTP using Ajax calls by encrypting data as well along with the extra header – another option that needs to be explored.

Detecting Ajax fingerprints on the Web server

We have Ajax fingerprints on an outgoing request from the browser. The Web application passes JavaScript to the browser in such a way that each legitimate request made by the browser has correct fingerprints. All that re-

mains to be done is to process the request on the Web server prior to serving the resource to the browser. This will be our first line of defense for Ajax-locked resources. We can build a defense bundled into the Web application firewall. We shall take two approaches here: one, for the Apache Web server and the other, for IIS with the .NET platform. Let us see each approach in a little detail.

Leveraging mod_security application firewall

`mod_security` is an application-level firewall that fits into the Apache Web server as a module. After the firewall is loaded into Apache (by modifying `httpd.conf`), start adding filtering rules. We shall add a specific ruleset for Ajax fingerprinting. Here is a sample rule.

```
<IfModule mod_security.c>
  SecFilterEngine On
  SecFilterScanPOST On
  SecFilterDebugLog logs/modsec_debug_log
  SecFilterDebugLevel 3
  SecAuditEngine On
  SecAuditLog logs/mod_audit_log

  <Location /ajax-only/>
    # Filtering incoming content
    SecFilterInheritance On
    SecFilterSelective "HTTP_Ajax-Timestamp" "^$" "deny,log,status:500"
  </Location>
</IfModule>
```

In above code snippet, the first few lines will set up the engine with logging enabled. The most critical ruleset that we want to set up is for the “Ajax-Only” section. All Ajax-serving resources reside in the `/ajax-only/` folder. Hence, we define our Ajax sandbox on the server by adding the “Location” tag with the correct folder.

All incoming requests to “Ajax-Only” must have a proper Ajax-Timestamp.

Apache will not serve any request that does not include this timestamp. This is the key filter ruleset at the application firewall.

```
SecFilterSelective "HTTP_Ajax-Timestamp" "^$" "deny,log,status:500"
```

We chop off the “Ajax-Timestamp” header; if it is empty or not present, a “500” error is

thrown back, as shown in the screenshot on the following page.

```
root@wsrd:/home/shreeraj# nc news.example.com 80
GET /ajax-only/header HTTP/1.0
```

```
HTTP/1.1 500 Internal Server Error
Date: Wed, 25 Oct 2006 15:17:21 GMT
Server: Apache/2.2.3 (Unix)
Content-Length: 607
Connection: close
Content-Type: text/html; charset=iso-8859-1
```

```
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>500 Internal Server Error</title>
</head><body>
```

Now, if we send a “proper” header to the server, we receive this response:

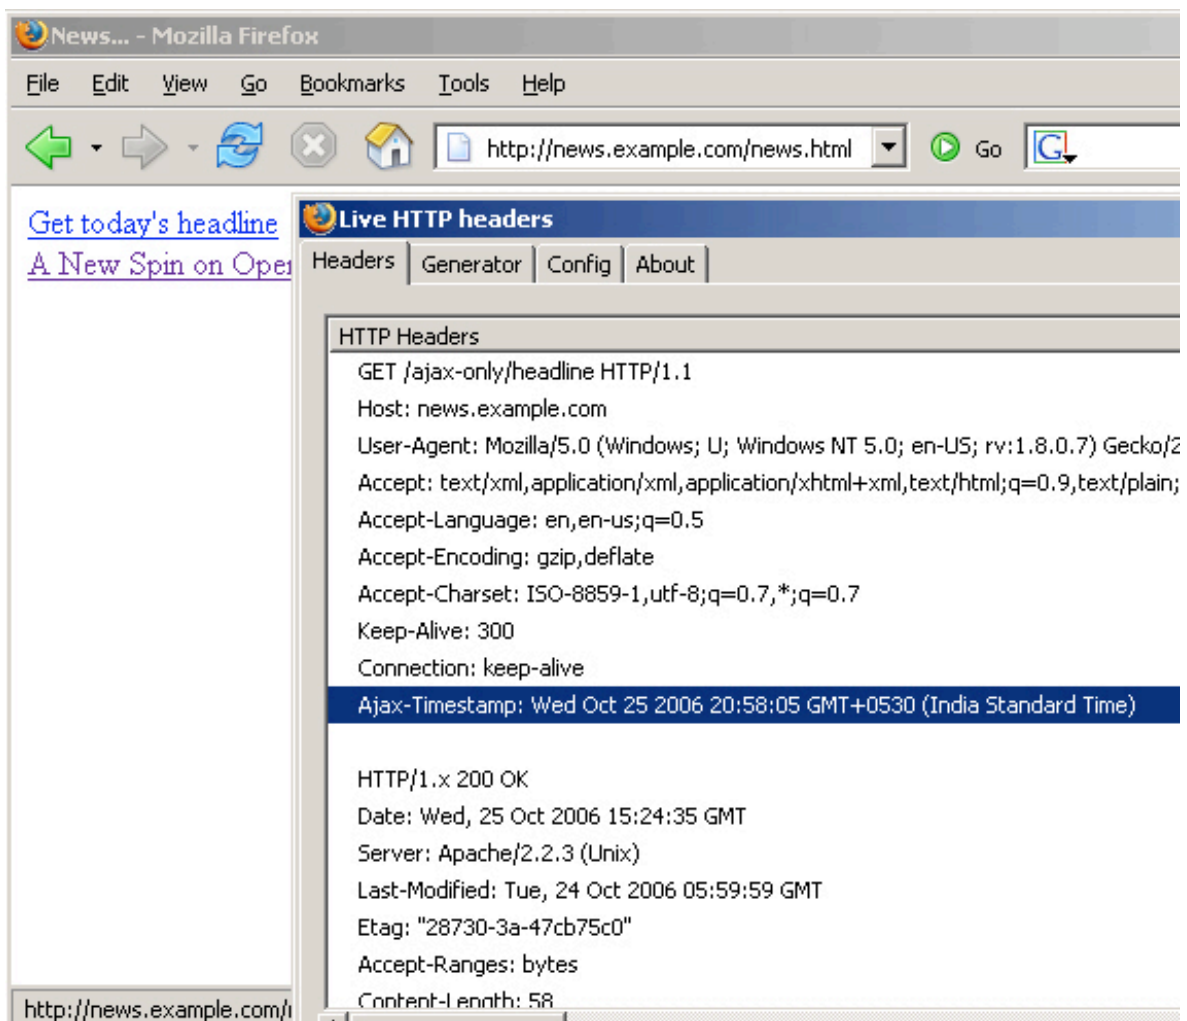


Figure 3.0 – Ajax request with the correct Timestamp.

The correct Ajax fingerprint in the HTTP request provides an entry to resources.

This example demonstrates that a web application firewall can be utilized in the right context for Ajax resources.

Leveraging HTTPModule for .NET applications

On IIS running with .NET, implement `IHttpModule` and access HTTP pipe for incoming HTTP requests. By setting up this hook it is possible to filter incoming traffic.

One can set up “Ajax-Only” virtual folder on IIS and put library code in the form of DLL in the “/bin” sub-folder inside “Ajax-Only”.

Listing 2.0 shows the function that can be overridden. The HTTP request processing hook can be set up.

```
public void ProcessRequest(object o, EventArgs ea)
{
    HttpApplication app = (HttpApplication)o;
    string ajax = app.Request.Headers["Ajax-Timestamp"];
    if (ajax == null)
    {
        app.Response.Write("Error!");
        app.Response.End();
    }
}
```

Listing 2.0 - The ProcessRequest () function.

The preceding code will throw an error if the HTTP header doesn't have a proper “Ajax-Timestamp” in the block. The entire `HTTPModule` code, compilation and implementation in-

structions are listed in Exhibit 1. Now, let's try to access a resource without an Ajax fingerprint. We get the following result.

```
D:\csharp\Ajaxwall\csc> nc example.com 80
GET /ajax-only/hi.aspx HTTP/1.0
```

```
HTTP/1.1 200 OK
Server: Microsoft-IIS/5.0
Date: Sun, 29 Oct 2006 04:21:55 GMT
X-Powered-By: ASP.NET
X-AspNet-Version: 2.0.50727
Cache-Control: private
Content-Type: text/html; charset=utf-8
Content-Length: 6
```

Error!

Ajaxwall has blocked the request by throwing an “Error!” and has protected “hi.aspx” re-

source. Now, let's send the same request with the correct fingerprint.

```
D:\csharp\Ajaxwall\csc> nc example.com 80
GET /ajax-only/hi.aspx HTTP/1.0
Ajax-Timestamp: Tue Oct 24 2006 17:37:46 GMT+0530 (India Standard Time)
```

```
HTTP/1.1 200 OK
Server: Microsoft-IIS/5.0
Date: Sun, 29 Oct 2006 04:22:22 GMT
X-Powered-By: ASP.NET
X-AspNet-Version: 2.0.50727
Cache-Control: private
Content-Type: text/html; charset=utf-8
Content-Length: 2
```

Hi

The resource is served by Ajaxwall. This demonstration of how Ajax fingerprinting on both IIS and Apache can be implemented serves as a starting point for building firewall and filtering mechanisms using this Ajax fingerprinting technique.

Implementing content filtering to defend against XSS 2.0

XSS attacks are steadily mounting in Ajax frameworks. Ajax makes a backend call to various third-party resources such as RSS

feeds or blogs. Since Ajax can not directly make these calls to the target site, calls are routed through server-side proxy code. It is important to filter out bad content originating from third-party untrusted sources and directed to the end user's browser. One of the approaches that can be adopted is by adding rulesets into the Web application firewall (WAF) for all third-party information modules. Here is an example that demonstrates this approach. Once again we can use mod_security to enable response filtering on HTTP/HTTPS content. We add certain rules.

```
<Location /ajax-only/>
  # Filtering incoming content
  SecFilterInheritance On
  SecFilterSelective "HTTP_Ajax-Timestamp" "^$"
  "deny,log,status:500"
  # Filtering outgoing content
  SecFilterScanOutput On
  SecFilterSelective OUTPUT "javascript:" "deny,status:500"
  SecFilterSelective OUTPUT "<\s*script.*?\s*>"
  "deny,status:500"
</Location>
```

The following line enables scanning for outgoing content: `SecFilterScanOutput On`

The next two lines ensure that HREFs are not injected with "javascript". Any attempt to inject the `<script>` tag in the HTTP response will also be blocked.

```
SecFilterSelective OUTPUT "javascript:" "deny,status:500"
SecFilterSelective OUTPUT "<\s*script.*?\s*>" "deny,status:500"
```

Any malicious content present in third-party information will cause a "500" error to be thrown. The user's browser stays secure. We have the following resource that fetches RSS feeds' XML file from the target server.

```
/ajax-only/rss?feed=http://sample.org/daily.xml
```

/rss is proxy code that will fetch the RSS feed from `http://sample.org/daily.xml` and send it back to the browser. `daily.xml` has the pattern "javascript" in one of the links. If the link is clicked, malicious code will get executed and the browser session may get compromised. With response filtering on HTTP/HTTPS content enabled, the same request responds with a "500" error.

```
C:\Documents and Settings\Administrator> nc news.example.com 80
GET /ajax-only/rss?feed=http://sample.org/daily.xml HTTP/1.0
Ajax-Timestamp: Tue Oct 24 2006 17:37:46 GMT+0530 (India Standard Time)
```

```
HTTP/1.1 500 Internal Server Error
Date: Sun, 29 Oct 2006 06:45:56 GMT
Server: Apache/2.2.3 (Unix)
Connection: close
Content-Type: text/html; charset=iso-8859-1
```

```
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>500 Internal Server Error</title>
```

Similarly, the <script> tag will be filtered out too. This filtering approach will help in securing a Web client.

Conclusion

Ajax security is a major issue for next generation Web applications. The techniques discussed in this article can give a headstart to security professionals to improve the security posture of Web applications. Web 2.0 applications try to integrate various sources, including untrusted information sources, at one place. This trait of Web 2.0 applications adds

new attack vectors to the landscape. The advantage of Ajax fingerprinting with XHR is twofold: one, it gives a clear idea about the origin of a request and, two, it makes it harder for automated attacks and crawler modules to launch discovery techniques. With Web application firewalls becoming an important part of Web application defense, one can leverage this mechanism to defend the web browser as well. Tools such as mod_security and HTTPModule can help in building better and secure deployment.

Exhibit 1 - AjaxWall Source Code and Implementation

```
--- Ajaxwall.cs [Source code file] ---

using System;
using System.Collections.Generic;
using System.Text;
using System.Web;
using System.Text.RegularExpressions;

namespace Ajaxwall
{
    public class Ajaxshield : IHttpModule
    {
        public void Init(HttpApplication App)
        {
            App.BeginRequest += new EventHandler(this.ProcessRequest);
        }

        public void Dispose()
        {
        }

        public void ProcessRequest(object o, EventArgs ea)
        {
            HttpApplication app = (HttpApplication)o;
            string ajax = app.Request.Headers["Ajax-Timestamp"];
            if (ajax == null)
            {
                app.Response.Write("Error!");
                app.Response.End();
            }
        }
    }
}

--- Ajaxwall.cs ends ---
```

Compiling the code

.Net 2.0 command prompt:

```
D:\csharp\Ajaxwall\csc> dir
Volume in drive D is Local Disk
Volume Serial Number is 9033-9D55

Directory of D:\csharp\Ajaxwall\csc

10/29/2006 09:34a    <DIR>        .
10/29/2006 09:34a    <DIR>        ..
10/22/2006 11:41p             768 AjaxWall.cs
                1 File(s)        768 bytes
                2 Dir(s)      987,975,680 bytes free
```

Compiling with csc:

```
D:\csharp\Ajaxwall\csc> csc /t:library AjaxWall.cs
Microsoft (R) Visual C# 2005 Compiler version 8.00.50727.42
for Microsoft (R) Windows (R) 2005 Framework version 2.0.50727
Copyright (C) Microsoft Corporation 2001-2005. All rights reserved.
```

Generation of Ajaxwall DLL:

```
D:\csharp\Ajaxwall\csc> dir
Volume in drive D is Local Disk
Volume Serial Number is 9033-9D55

Directory of D:\csharp\Ajaxwall\csc

10/29/2006 09:34a    <DIR>        .
10/29/2006 09:34a    <DIR>        ..
10/22/2006 11:41p             768 AjaxWall.cs
10/29/2006 09:34a             3,584 AjaxWall.dll
                2 File(s)        4,352 bytes
                2 Dir(s)      987,971,584 bytes free
```

Implementing on IIS:

1. Put AjaxWall.dll into the /bin/ folder of "Ajax-Only".
2. Add the following lines to the web.config file (this is required to load this module)

```
<httpModules>
  <add type="Ajaxwall.Ajaxshield, Ajaxwall" name="Ajaxshield" />
</httpModules>
```

Now that this hook has been turned on, every request hitting ASP.NET resources such as .aspx or .asmx will be processed. Ajaxwall is up and running.

Shreeraj Shah, BE, MSCS, MBA, is the founder of Net-Square and leads Net-Square's consulting, training and R&D activities. He previously worked with Foundstone, Chase Manhattan Bank and IBM. He is also the author of Hacking Web Services (Thomson) and co-author of Web Hacking: Attacks and Defense (Addison-Wesley). In addition, he has published several advisories, tools, and whitepapers, and has presented at numerous conferences including RSA, AusCERT, InfosecWorld (Misti), HackInTheBox, Blackhat, OSCON, Bellua, Syscan, etc. His articles are published on Securityfocus, O'Reilly, InformIT and HNS. You can read his blog at shreeraj.blogspot.com.



Hack In The Box Security Conference 2006
September 18-21 - Kuala Lumpur, Malaysia

This year's HITBSecConf was yet again another rockin' event which saw over 600+ attendees from around the world in Kuala Lumpur for 4 days of deep-knowledge network security training and talks.

The buzz began during the first 2 days of training on the 18th and 19th of September. This year's classes featured 6 tracks conducted by 10 trainers.

One of the highlight classes was a 'Government and Law Enforcement Only' wireless security and war driving exercise conducted by Anthony Zboralski and Jim Geovedi of Bellua Asia Pacific. The training saw participants taken in a specially equipped 44-seater bus on a wireless network hunt and security survey.

The conference proper kicked off on the 20th with the first participant signing in at the crack of dawn of 6:45am! The crowd slowly grew and by the time 9am rolled around, well over 500 participants from 25 countries were gathered and ready to go.

In all, HITBSecConf2006 – Malaysia saw 37 world renowned speakers down to share their latest research and findings. Day 1 keynote

speaker Bruce Schneier however could not make his flight and instead delivered his keynote via a live webcast. He did promise however to make it over for HITBSecConf2007 – Malaysia.

On Day 2, Mark Curphey and John Viega took the stage with an entertaining and informative keynote highlighting the limitation of automated protection and assessment tools.

Another presentation that was eagerly anticipated was Joanna Rutkowska's Blue Pill paper which dealt with the issue of stealth malware utilizing the latest virtualization technology from AMD.

The Microsoft Windows Vista team was also around this year to present for the first time in Asia, an inside look at the security workings in Windows Vista and the BitLocker drive encryption technology. Microsoft also sponsored the post conference party for conference speakers and invited guests.



At the end of every HITBSecConf, HITB organizes a charity auction in aid of a Malaysian beneficiary. In 2004/2005 all donations went to the National Kidney Foundation. For 2006 and 2007, the charity beneficiary is the Malaysian

National Cancer Council. Up for grabs were autographed HITB CREW t-shirts, Firefox laptop bags, 3 copies of autographed Bruce Schneier's "Beyond Fear", and more.





The Capture The Flag competition this year featured teams from Malaysia, Singapore and for the first time a team from Europe as well as Korea! In all 9 teams competed in the 2-day competition which was said to be one of

the most difficult and challenging CTFs ever run by HITB. For more coverage and photos of HITBSecConf2006 go to www.tinyurl.com/yx34r5



Stay Ahead



Technology moves at the speed of light in the world of cyber security attacks and defenses. Black Hat DC will once again gather the world's information and computer security elite to share their knowledge and experience with you.

Two days. Ten Classes. Thirty presentations.



Black Hat[®]

Briefings & Training DC 2007

February 27-March 1 • Sheraton Crystal City

www.blackhat.com

sponsors

diamond

Microsoft

platinum

BAE SYSTEMS

gold

ArcSight

BlackBerry

Configuresoft

CORE
SECURITY TECHNOLOGIES

IOActive

NORMAN



Where iSCSI fits in enterprise storage networking

By Todd Bundy

With growing business and regulatory pressures, enterprises are facing increased demands to enhance their off-site storage strategies.

Backing up data to disk technologies in the room next door is not an effective solution for ensuring disaster recovery or business continuity in the event of unavailability at a site of operations. Backup data centers are being located further and further off site - sometimes as far as 200 km away from the primary data centre.

A variety of transport options are available for transmitting data among distributed facilities. Depending on factors such as cost requirements and the distances to be covered, an enterprise might consider deploying distributed-storage solutions over IP, Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) and/or Wavelength Division Multiplexing (WDM) transport networks. Many small and medium-sized businesses – already accustomed to building, managing and maintaining Ethernet networks within a given corporate location – are finding a cost-effective, simple solution for distributed storage in an emerging IP-based protocol, Internet Small Computer Systems Interface

(iSCSI). In many cases, enterprises are successfully turning to iSCSI to back up traffic that is not mission critical, as a complement to Fibre Channel applications running across SONET/SDH and WDM infrastructures.

The Move to Distributed Storage

Perhaps there have been no more dominant trends in enterprise networking over the last five years than the growing alarm over the cost of network downtime and the surge in government regulations stipulating how companies back up and secure their data. The loss of mission-critical data can prove fatal for an enterprise, as the ramifications of idle staff, reputation damage and revenue loss build upon one another. Understanding this, enterprises have sought to minimize network unavailability with state-of-the-art Storage Area Network (SAN) connectivity solutions, either deployed and managed by the enterprise itself or offered as fully managed storage services by a carrier. Furthermore, enterprise IT managers and CIOs have been deluged on the

regulatory front. The Sarbanes-Oxley Act, Graham-Leach-Bliley Act, Basel 2 and the HIAA, in-addition to other new legislation has put unprecedented pressure on enterprises to plan for disaster recovery and business continuity, enhance network security and disclose compromises.

The result of the business and regulatory pressures: Never has so much or so great a range of confidential data – financial, personal, competitive, medical, etc. – been networked across metro, regional and wide area networks. Distributed data centre and storage connectivity solutions have been adopted by perhaps as many as 70 percent of the world's Fortune 1000 companies. Whether deploying and managing SAN themselves or contracting for managed storage service offerings from carriers, enterprises of every size have sought to put in place more powerful business continuity and disaster recovery capabilities.

Transport Options

IP, SONET/SDH or WDM-enhanced optical networks can transport an enterprise's storage traffic among distributed data vaults. Choosing which storage traffic to entrust to which transport option is dependent upon a variety of factors including cost requirements, the speeds of the storage applications to be networked, the mission-criticality of the traffic and the distances to be covered.

Packet-based IP is a cost-effective, simple option. Built on this packet-based foundation, Ethernet is ubiquitous across enterprise networking – widely deployed and well understood by enterprise IT managers. Enterprises can immediately leverage their existing IP networks to support iSCSI storage applications for backing up low-priority data over long distances.

Ratified by the Internet Engineering Task Force (IETF) in 2003, iSCSI has its limitations. iSCSI header and commands must be added to the Ethernet packets being transported across the SAN, and this introduces a protocol

overhead that renders iSCSI's performance insufficient for runtime-sensitive, synchronous storage applications such as Geographically Dispersed Parallel Sysplex, Fiber Connection (FICON) and Enterprise Systems Connectivity (ESCON). These applications have terrific bandwidth requirements (several terabits per second) and low tolerance for latency and only WDM-enhanced optical networks deliver the reliable performance they demand. But, particularly among cost-conscious enterprises that seek to rely on existing IP networks and interfaces that are widely familiar among IT staff, iSCSI is emerging as an increasingly prevalent solution for affordably supporting remote backup and linking storage facilities. Furthermore, performance is sure to improve as a gradually wider range of SAN systems and subsystems supports iSCSI, a relatively new protocol.

A Key Thread in the Storage Fabric

IP-based iSCSI storage networking should not be regarded as a replacement for Fibre Channel applications running through SONET/SDH gateways or across WDM-enhanced optical networks. The most sophisticated enterprise storage strategies will leverage all three transport options. For example, WDM can be relied upon to provide reliable, transparent, protocol-agnostic connectivity between iSCSI and Fibre Channel SAN islands, delivering optimal disaster recovery and business continuity capabilities as well as greater, cost-efficient storage and server utilization.

As more and more enterprises seek to realize the unprecedented capabilities of sophisticated storage networking, transport and service options are expanding to meet varying performance, distance and cost requirements. iSCSI fills an important role within that realm. With a flexible, WDM-enhanced optical network foundation underlying their storage strategies, enterprises are able to cost-effectively implement the capabilities they require today and position for tomorrow's new requirements.

Todd Bundy (tbundy@advaoptical.com) is the Director of Business Development and Alliances at ADVA Optical Networking. He has 24 years experience in the storage networking industry and is a recognized expert in SAN and optical networking; specializing in storage applications to meet corporate contingency plans.

Software spotlight

WINDOWS - SSL-Explorer

<http://www.net-security.org/software.php?id=579>

This entry-level VPN solution provides users and businesses alike with a means of securely managing their IT infrastructure from outside the network perimeter, armed with just a standard web browser.

LINUX - Sussen

<http://www.net-security.org/software.php?id=497>

Sussen is a client for the Nessus Security Scanner. It is easy to use; you can perform a vulnerability assessment with just a few mouse clicks. It has a Glade-based user interface, Druids for common tasks, GConf support, and Anjuta project support.

MAC OS X - Data Guardian

<http://www.net-security.org/software.php?id=662>

Data Guardian is a secure database application for storing passwords, credit card numbers, addresses, notes, customer databases, and more.

POCKET PC - Sentry 2020

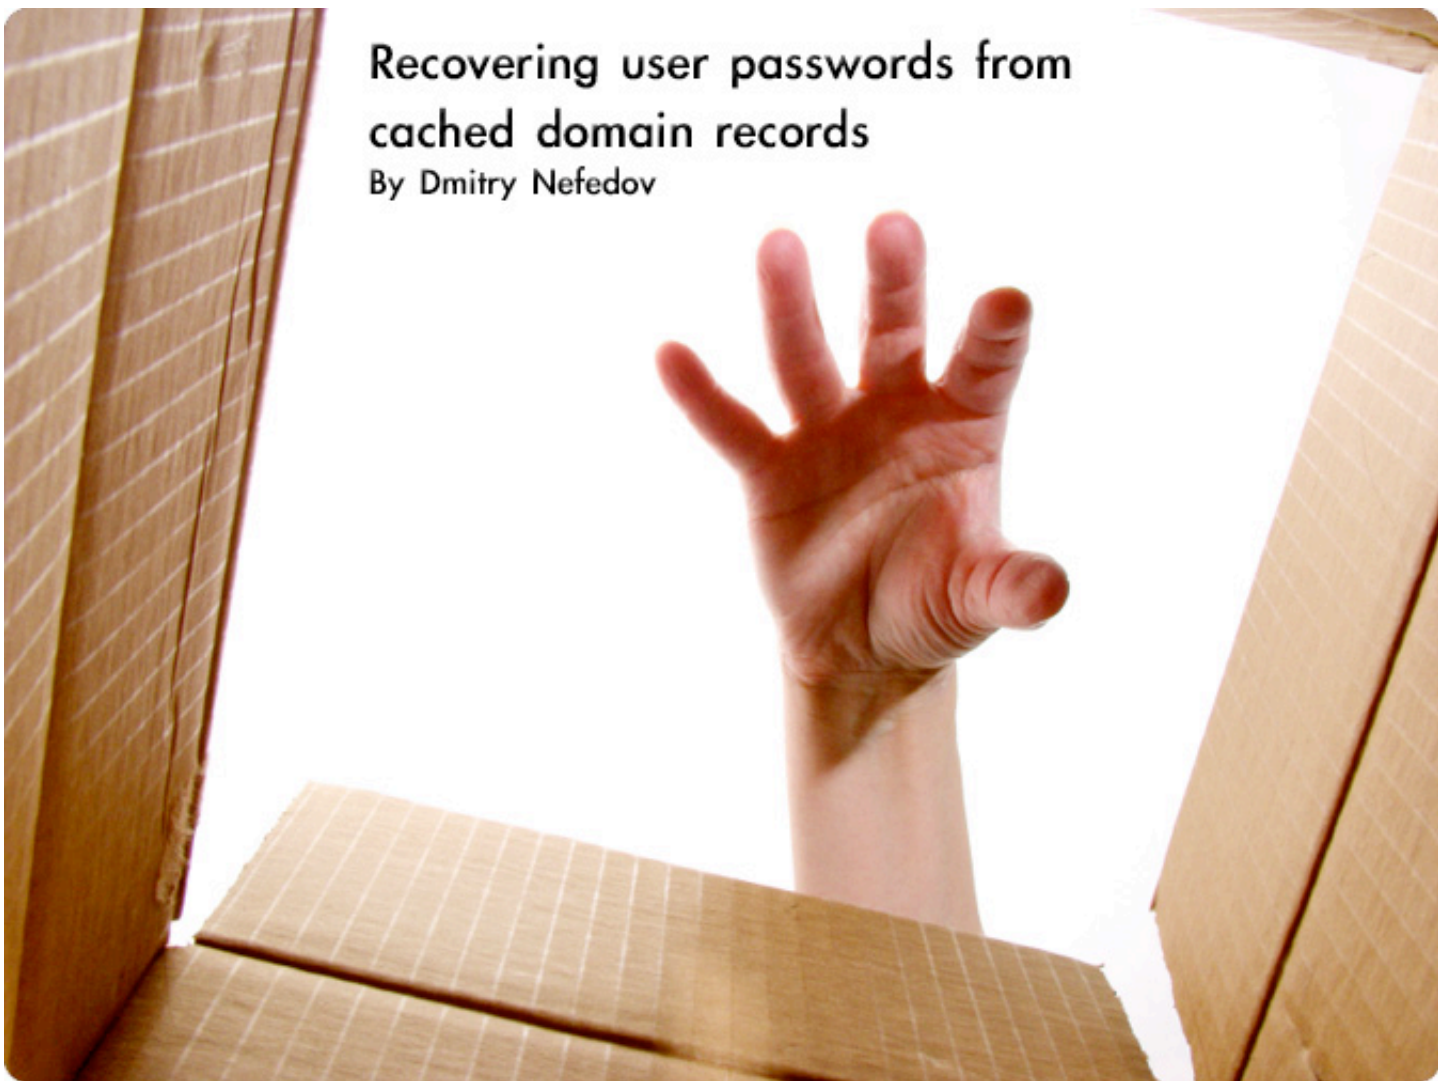
<http://www.net-security.org/software.php?id=541>

Sentry 2020 acts like a vault and provides extremely strong security so no one but you can access your confidential information.

If you want your software title included in the HNS Software Database e-mail us at software@net-security.org

Recovering user passwords from cached domain records

By Dmitry Nefedov



Operating systems based on the Windows NT series can cache (store) user logon information on users that enter the domain. This feature is designed to bypass the authorization procedure after the server has been unavailable for one reason or the other.

What is Domain Cached Credentials (DCC), and what are they for? Microsoft's website provides the details at the following addresses:

- 1) Cached Logon Information (tinyurl.com/ybsk39)
- 2) Microsoft Windows XP - Logging On Using Domain Credentials (tinyurl.com/87b7a)

Along with the general information on a domain user, which includes actual user information, domain information, and general information (the DCC common record structure will be covered below), DCC contains the user's password hash. This article's objective is to help you understand and figure out how cached user records are stored, whether passwords can be recovered from the structure, and if it creates a security threat.

Where are password hashes stored?

Let's start with the point that cached records (which store password hashes) are controlled with the `HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\Current Version\Winlogon` registry key. This branch must contain the `CachedLogonsCount` string parameter, which holds information on the number of records being cached. It may carry values ranged between 0 and 50. 0 means caching is disabled. If the value specified is greater than 50, only 50 records will be cached anyway.

By default, the `CachedLogonsCount` value is set to 10. That means caching is enabled by default. This value can be changed manually in the registry or through the domain security policy.

The actual cached records are also stored in the registry, in the encrypted form. The registry branch of `HKEY_LOCAL_MACHINE\SECURITY\Cache` contains two types of values:

1) The `NL$Control` value, which structure carries the cached records version and the number of records in the specified branch. The cached record version is necessary for controlling the records during updates, making changes in the record structure, changing encryption algorithm, etc.

Here are cached records versions for the entire NT series:

```
- NT 3.0      0x00010000
- NT 3.5      0x00010002
- NT 4.0 SP4  0x00010003
- NT 5.0      0x00010004
```

2) Values with names in the format of `NL$x`, where `x` stands for the cached record number. Each record contains information on user (name, profile, home folder, RID, group RID, etc.), domain (name, SID, last accessed time, etc.) Besides that, a cached record contains user's hashed password. It will be covered in detail further on.

Cached records format and encryption algorithms

In the previous section, we have learned that cached record is stored in the registry, and it contains hashed passwords. The access rights of `HKEY_LOCAL_MACHINE\SECURITY\Cache` disallow opening it for reading even for a domain administrator by default.

However, having the administrator access permissions and a bit of imagination, you can bypass that restriction. Anyway, if the cached records could not be decrypted automatically, you can do all that manually or modify the access permissions with `regedit` (if you are sure enough that you know what you are doing).

Additionally, it must be said that the old Windows NT cached records were stored in a different format and used a different (weaker) encryption algorithm. However, let's skip the review of this atavism in the article due to - let's put it this way - infrequent use of Windows NT these days.

Now, let's get started. Here is a more detailed description of a cached record (the structure description is provided in C++ in order to make it clearer).

```
typedef struct _tagDOMAIN_CACHE_ENTRY
{
    LPWSTR wszUserName;           //user name
    LPWSTR wszDomainName;        //fomain name user log on
    LPWSTR wszFullName;          //full user name
    LPWSTR wszLogonScript;       //user logon script name
    LPWSTR wszProfilePath;       //and profile path
    LPWSTR wszHomeDirectory;     //home directory
    DWORD dwUserId;              //RID
    DWORD dwPrimaryGroupId;      //primary group RID
    //... cut something

    LPWSTR wszLogonDomainName;    //domain name
    int64 LastAccessTime;         //last time user access the domain
    DWORD dwRevision;            //current structure version
    BOOL Valid;                  //is current entry active or not
                                //End of fixed part

    CHAR CypherKey[16];           //this path of structure is not encrypted
    CHAR Sign[16];               //not encrypted
    //start of entry encryption
    CACHE_PASSWORDS CachePasswords;
} DOMAIN_CACHE_ENTRY, *PDOMAIN_CACHE_ENTRY;
```


For the sake of simplicity, some fields were modified or cut out of this structure. Indeed, the record structure is somewhat more complicated, and contains the following fields (those who don't want to bother to read the technical details may skip this part):

- `cOffsets` - encrypted structure fields offset related to the beginning of data. Its length is 64 bytes for the version 0x00010004.
- `CypherKey` - encryption key. Its length is 16 bytes.
- `sign` - signature for verifying data integrity - 16 bytes.
- `CACHE_PASSWORDS` - password hashes. From this moment on, the cached record is encrypted.
- `cReserved` - purpose unknown. Probably, this is a reserved field. 32 bytes.
- `cData` - data of variable length.

Very ingenious, as always. However, if they considered it, they could have arranged storing the variable-length data more universally. Nevertheless, the source code roots and thoughts go far to the past, to the good old 80s. Don't be too harsh on the software developers and hold your criticism and judgments. Actually, we need to focus on the last three fields.

- `CypherKey` - 128-bit password encryption key. Every time a new cached record is created, it is filled up with random data.
- `sign` - md5 hmac signature. Digital signature for certifying integrity of data.
- `CachePasswords` - hash structure for the encrypted password. It is arranged, approximately, as follows:

```
typedef struct _tagCACHE_PASSWORDS
{
    CHAR          NtHash[16];
    CHAR          LmHash[16];
    BOOL          NtHashPresent;
    BOOL          LmHashPresent;
    BOOL          HashesNotActive;
} CACHE_PASSWORDS, *PCACHE_PASSWORDS;
```

Let's review the encryption algorithms. Unlike in Windows NT, elaborate enough and interesting password encryption algorithms are used here. By saying 'passwords', I certainly mean password hashes, for passwords are not meant to be stored in the unencrypted form in the Windows NT operating system. To be more accurate, they are not meant to be stored. The reality, however, is far not that fantastic. Let's get off the good old WinNT; yet under Win2K, for instance, you can retrieve the logged on user's original password by pulling it right out of the memory (it sits there until the user logs off the system). In WinXP, the user's original password (and the previous one, if necessary!) can also be recovered by retrieving it from the secrets. You don't even have to log on to the system as an administrator - just boot the computer under another operating system installed on your hard disk or boot up from the recovery disk and then copy the registry files of SAM, SYSTEM, and SECURITY. From the SOFTWARE file you can

retrieve the password hint (WinXP or higher only).

This is why, should you have ever lost your password, we highly DISCOURAGE you to use password-resetting software like this one (home.eunet.no/~pnordahl/ntpasswd). Before you delete your password hash, make sure the original password cannot be recovered.

Besides that, having once used the Offline NT Password & Registry Editor (or a similar program), although you can reset the SAM account's password and then successfully logon to the system, ALL private data, desktop configuration, personal certificates, and saved web site and network resources passwords will be lost irretrievably. Not having your personal certificates handy, you will not be able to access your encrypted files (if any) and e-mail messages encrypted with your private key. Use such programs at your own risk or as an extreme measure for gaining access to your system.

Until this point, I have only covered the security weak spots. Yet we should give honor to Microsoft's programmers - the area of security in WindowsXP has been greatly improved, and it has got not less of good things than *NIXes have. However, please pardon me for the off-topic lyrics; recovering the SAM or AD user password is another major topic for a separate article.

Now let's turn to encryption algorithms. Those are standard for Microsoft and include RC4, MD4 and HMACMD5. The last one is a pretty interesting interpretation of MD5, which eliminates just about all shortages of its ancestor.

In order to decrypt a domain user's password, you must complete two steps (of a large path). The first one is decrypting the actual cached record from `HKEY_LOCAL_MACHINE\SECURITY\Cache`, since it is stored, as it has been said, in the encrypted form. Once the record has been decrypted, we will have the modified password hashes at our disposal. The "modified" means the hashes will be in the form of `HASH(HASH(password)+UserName)`.

The second step is to "guess" the password for the received hashes using usual for such cases methods - the dictionary attack or the brute force attack. Recently, the new wave of attack algorithms has come up - the rainbow

attacks. In our case, this type of attacks is useless. Why? The description of hashes stored in the records makes it clear.

Recovering user passwords

Let's have a closer look at the record encryption algorithm. It was mentioned above that the cached record decryption algorithm uses the 16-byte `CypherKey` field. Actually, besides the `CypherKey` field, the encryption algorithm uses the LSA secret named `NL$KM`, which (when being created) is also initialized with random data. Thus, the record can be decrypted only on the local computer, since the `NL$KM` value is different for each computer.

Physically, `NL$KM`, along with other secrets, is stored in the `%WINDIR%\SYSTEM32\CONFIG\SECURITY` registry file, and, in its turn, is also encrypted (just like all the secrets are). However, it can be retrieved directly with the `WINAPI LsaRetrievePrivateData`. Or - if you do not have sufficient rights to read this secret (and the default setting provides the rights that are exactly insufficient) - do that via the registry. Suppose we do have a 64-byte `NL$KM` secret. What is next? Next, we are going to deal with a pretty interesting decryption algorithm. Here is a fragment from the source code that decrypts the cached record.

```
//Decrypt inplace
BOOL CDomainCache::DecryptCachedEntry(PDOMAIN_CACHE_ENTRY pCacheEntry, DWORD dwEntrySize)
{
    assert(pCacheEntry && dwEntrySize);
    CRC4 rc4;
    CMD5 md5;
    BYTE pDerivedKey[MD5_DIGESTSIZE], pMAC[MD5_DIGESTSIZE];
    PDOMAIN_CACHE_ENTRY pEntry;
    PBYTE pbData;
    ULONG cbData;

    if ( pCacheEntry->Revision!=2K_CACHE_REVISION ) //old NT version
    {
        m_dwLastError=E_OLDREVISION;
        return FALSE;
    }

    pEntry=(PDOMAIN_CACHE_ENTRY) new BYTE[dwEntrySize];
    if ( !pEntry ) //not enough memory
    {
        m_dwLastError=E_NOMEMORY;
        return FALSE;
    }

    memcpy(pEntry,pCacheEntry,dwEntrySize);
```

```

//Get key
    if ( m_CurrentSecret.Buffer && m_CurrentSecret.Length==64 )//try current NL$KM
first
    md5.HMACInit((LPBYTE)m_CurrentSecret.Buffer,64);
    else
    md5.HMACInit((LPBYTE)m_OldSecret.Buffer,64);
    md5.HMACUpdate((LPBYTE)pEntry->CypherKey,sizeof(pEntry->CypherKey)); //than random
entry key
    md5.HMACFinal(pDerivedKey);

    //Compute offset and length
    pbData=(LPBYTE)&(pEntry->CachePasswords);
cbData=dwEntrySize-(DWORD)(pbData-(LPBYTE)pEntry);

    //Try to decrypt
    rc4.SetKey(pDerivedKey,MD5_DIGESTSIZE);
    rc4.Decrypt(pbData,cbData);

    //Check integrity
    md5.HMACInit(pDerivedKey,MD5_DIGESTSIZE);
    md5.HMACUpdate(pbData,cbData);
    md5.HMACFinal(pMAC);

    if ( memcmp(pEntry->Sign,pMAC,MD5_DIGESTSIZE)==0 )
    {
        memcpy(&pCacheEntry->CachePasswords,pbData,cbData);
        deletenull(pEntry);
        return TRUE;
    }

    //////////////////////////////////////
    //Try old secret
    //Just almost the same.... Skipped and cut

    m_dwLastError=E_DECRYPTION;
    return FALSE;
}

```

What do we have here? At the point of entry, we have `pCacheEntry` - the original encrypted cached record - and `dwEntrySize` - its length.

First, we are going to check the version number. If it is obsolete, we are not going to mess with it, and simply exit. Then, we allocate local memory for duplicating the record and then copy the record to that memory. Next, we check whether the current `NL$KM` secret is available. If it is available (we should have read it before this point), we begin initialization of the encryption key; otherwise we initialize it with the old one (where do we get the old one – that's another subject for discussion). Next, we take `CypherKey` from our record and continue initializing the encryption key with the MD5 algorithm.

Once the HMACMD5 algorithm has completed its job, we will have the `pDerivedKey` encryption key. With the RC4 algorithm, we decrypt our record using the obtained `pDerivedKey` encryption key. At the same time, we must not forget to verify whether the decryption was

done right (do you still remember the `sign` field from the `DOMAIN_CACHE_ENTRY`?) If the decryption was done wrong, get the old `NL$KM` encryption key and start all over.

As you may see here, everything is both pretty quick very competent. All honor to programmers from Microsoft!

So far, we have decrypted the cached record and may be excited that we already have all data from `DOMAIN_CACHE_ENTRY` (user name, RID, home folder, etc.) at our disposal. But that is just the beginning. Now we are going to try decrypting the actual user password

For that purpose we need to know the user name; without it, we will be unable to get the password. The user name can be easily taken from the record we have just decrypted. In our hands we now have these hashes, `HASH(NTLM+UserName)`. The `UserName` stands for user name in the lower case, and `NTLM` is `HASH(password)` – the user password hash.

So, we shall check, for instance, the password "123" for user Dima. For that purpose, we obtain MD4 hash of our password. Now, convert the user name to the lower case; we will have "dima" at the output. Then we convert "dima" to UNICODE and add it to the password hash. In this algorithm, "dima" will be the salt. The obtained 24 bytes (16-byte hash + 8-byte user name) again make up the MD4 hash.

That's it. We can now compare the obtained hash with what we have got after decryption. If they match – bingo! The user's password is "123".

I will to skip writing the source code here, since everything is pretty trivial and, I hope, clear enough. You can try to recover the password using the dictionary or brute force attack.

What do we benefit from this?

Some reader may ask, "So, is there any use of that?" Yes there is. If you have a user password, you can retrieve and decrypt all other system passwords on this user's computer, and you don't even have to logon to it. Those are the Internet account passwords, network passwords, IE and OE passwords, access to files encrypted with EFS, etc.

This is an even greater treasure for those who have lost their passwords (my tongue automatically want to say, "a real boon for a spy"), which happens pretty often. We can now try to recover the lost or forgotten user password if the SAM database is unavailable for one reason or the other.

Vulnerability or risk?

Is it a risk, vulnerability – or – there is nothing to worry about? Let's consider this question from all sides.

By the highest standards, there is, certainly, nothing to worry about. All data and passwords are encrypted correctly, and calling this a vulnerability would not really be right.

First, competent usage of encryption algorithms provides no opportunity to "knock on the back door" (like it was with Win95) and find vulnerability in encryption algorithms.

Second, access to `SECURITY\cache` is closed even for the domain administrator. Third, in order to decrypt a cached record, one needs to know the LSA secret, which is also unavailable by default. Fourth, even with a decrypted record available, there is literally no way to find out the user password if the password meets some security policy.

Let's count out roughly, how much time would it take one to uncover, for instance, this password: "ABcd12@#". One will need the Charset of A-Z, a-z, 0-9, !-~ (26+26+10+34) 96 characters. Thus, one will have $96^8=7213895789838336$ password options available total. Divide this number by the number of passwords one can try per day $4600000*60*60*24=397440000000$ passwords per day. Now, divide 7213895789838336 by 397440000000 and receive 18151 days. So, in order to find a matching password, one will need approximately 100 years. Not bad, really.

Now, let's look at this from the other side.

First, encryption algorithms used in Windows are poorly optimized for the new processor types. On the one hand, this is, certainly, a benefit – such algorithms (written, as usually, in C) are universal and platform-independent. On the other hand, however, such algorithms can be re-written in the assembler language, and their performance will improve times and times. For instance, on my AMD Athlon 2500+ computer, the speed of the incremental search attack in the PSPR program is approximately 600,000 passwords per second with the standard MD4 algorithm used and around 4.6 million passwords per seconds when the algorithm is optimized for the MMX CPU. Again, if we optimize the encryption algorithm for the SSE processor, the search speed will increase another 1.5 times or so.

Second, although the `SECURITY\Cache` registry is unavailable for reading for administrators, this does not mean it cannot be read. This is another topic for a separate discussion.

Third, the same also applies to LSA secrets

Fourth, there is the law about a weak link in a chain. If we consider an average domain,

hardly over 80 percent of its users use reliable passwords. And if at least one password within the network becomes known, most of the remaining passwords will open up without much of effort.

In our first example, we considered quite a complicated password. But if user's password is, for instance, "abc123", it will take only about 15 minutes to recover it. However, I must notice, it is impossible to define a certain password's resistance level in advance.

You should not forget about the dictionary analysis either.

The truth, as people say, will be revealed in a comparison. Let's compare the password search speed for the Win98 OS and our domain passwords. In the case with Win98, the password search speed will be almost twice as slow! Here goes your "old and unreliable" Win98. What do we read in this? Exactly – that programmers from M\$ do not really learn much from their mistakes. Or they lack to cooperate with one another. Again, I must say, that is not always the case. If I ever take my head to write another article on how network passwords and SSO credentials are encrypted, you surely will see that this encryption algorithm is a role model.

Resume

An eagle-eyed reader, may now ask, why password hash analysis in the `CACHE_PASSWORDS` structure is only performed for the NTLM hashes (the `NtHash[16]` field)? Recovering a password by incremental search of the LM hash (the `LmHash[16]` field) is more efficient.

Indeed, the LM hash is weaker for cryptanalysis compared to NTLM, and it is formed out of two halves by 7 characters each by converting the password to the ANSI format and then converting it to uppercase. Let's count roughly, how much time would we need to recover LM hash for the password "ABcd12@#". Here we will get two halves - ABcd12@ and #. The second part's hash will be found almost immediately, but doing the same for the first part will generate $(26+10+34)^7$ choices, which, at the same speed of 4.6 million passwords per

second, will be searched through within approximately 21 days.

It is significantly sooner, compared to 36301 days in the case with the NTLM. Actually, even the 41 days value is overstated. We supposed that the search speed will be the same, although in the reality we should take into account that LM hashes are encrypted with the DES algorithm, which is faster. Accordingly, the password search speed should be greater as well. Besides that, in practice, one usually has to search through the entire range of passwords. Though, after we have flown so high, we will have to come down to earth, because LM hashes are not stored in cached domain records. At least, starting with Windows 2000 and on. Most likely, the developers have decided to make Windows more secure, since LM hashes is the weakest link in this chain. Not so bad! It is better to do nothing than do something well... What a unique way of thinking! They say the password-based authentication system will soon sink into oblivion. Sure, the computer market offers a great choice of HW-based authentication devices that read biometric data, and they become cheaper day after day. The future is theirs, no doubt. But what shall the millions that still use passwords do? Why Windows security policy forces us to use and remember passwords like H&2a)9%m1LD*#M\$?

I had a talk with another programmer. In the dispute on how the advanced programmer differs from the regular user, we have come to the same conclusion: the programmer always uses passwords like '123' or 'qwerty'.

Hmm, but indeed, what can be done to make even such programmers' passwords more secure? Forget both LM and NTLM hashes, you are saying? Well, this may be a solution. Caching within domain can also be disallowed, but that is not always applicable. We can configure password security policy in a way that doesn't allow users use simple passwords. Lately, some advanced users use all kinds of tricks generating and remembering such passwords. For instance, if the password contains the carriage return character (it can be entered by pressing the Alt+0+1+3 keys on the keyboard numeric pad), such passwords will not be properly processed with the popular Lophcrack program.

Some use a passage from a poem or a collection of capital letters from it as their password. For instance, the "tbontbtitq" password derives from "To Be Or Not To Be That Is The Question".

We are going to go our own way and try to give the developers a little bit of hard time. Our "salted" hashes are not afraid of the rainbow attack, and that is quite pleasant. But the HASH (HASH(password)+username) operation is processed too fast. The first thing that gets in our mind is to give up the MD4 algorithm and use the more reliable SHA algorithm instead. Although this should improve the security - still not to the extent we really need. Enclosing the hash calculation operation in a loop wouldn't be a bad idea either. However, if we set up a constant to define the number of iterations in the loop, we are risking to "get all wet" again in about 5-10 years (let's think of

our future as well), when the computing power will make another step forward. Therefore, the constant may be stored in the `DOMAIN_CACHE_ENTRY` structure. It would be also good to bind this constant to the computer's power, so that the hash computing operation would take the same time regardless of the fact whether it is running on Pentium II or Pentium V. For instance, you could set that value to the direct dependence on the CPU clock. For example, `1:1000000`. Thus, for a processor with 2.5 GHz clock, the constant value would be 2500. And one operation of HASH (HASH(password)+username) will be run at the speed of approximately not more than 920 times per second (on any computer). Certainly, you should also take into account performance downfall when the value is too great; therefore, you need to find the optimum between performance and reliability. We hope it will not take too much time.

Mr. Dmitry Nefedov is the Chief Software Developer at ElcomSoft (www.elcomsoft.com). ElcomSoft's award-winning password file protection retrieval software uses powerful algorithms, which are constantly under development.

LAVASOFT
protect your privacy

***The leading antispyware developer
now delivers the best personal firewall protection***



LAVASOFT PERSONAL FIREWALL
Superior security shield against hackers, worms and Trojans

www.lavasoft.com



Do portable storage solutions compromise business security?

By Rob Faber

Many companies surprisingly aren't worried about employees using a private USB stick to get some data from the company network and do some work at home, but analyst Gartner has warned repeatedly that portable storage devices pose a serious security threat. It can be lost or intentionally be used to leak sensitive data, and introduce serious trouble into a network. So how to cope with this and find a solution that is supporting business needs and still underline the corporate security policy?

After years of learning from attacks at the workstation level, we were able to understand that there was need for a proper anti-virus solution and, if possible, to further lock-down the client and introduce a corporate policy that prescribes how to log on as an end-user, with preferably a strong password.

But did we really ultimately secured the workstation and can we now take a deep breath? Those of us who read the news know better.

This article discusses and provides information in general about the risk that is to be expected from portable storage devices like USB memory sticks, iPods and portable hard drives. We also take a look at the possible solutions for your company to deal with this threat.

Gartner warns business

Storage devices became cheaper and cheaper and now can hold a huge amount of data. So what? Well, they can pose a serious threat.

Analyst Gartner has warned repeatedly that portable storage devices pose a security threat to companies and that there is an urgent need to do something about this situation. During 2006 there has been an impressive number of (major) incidents involving portable storage devices. The consequence: loss of unprotected and really sensitive military, commercial or police investigation data, a painful loss of a careful built-up image, serious damage from a financial perspective or even problems with the law.

The ingredients for disaster are there: high data capacity, very high transfer rates and widely platform support mean that an attached USB device has the capacity to quickly download lots of high value company information, which can then be easily leaked. A big problem is also the fact that these devices can be used to intentionally or unintentionally introduce malicious software into a network.

The first reaction would possibly be: simply ban all portable storage devices that can ever connect on a workstation. If you're in a small office this is easy to achieve but in a large organization that has many employees with different roles and needs? Not that simple.

In many cases there will be a business need to transfer or transport data, even if it is sensitive and you don't like it.

You have to give the possibility to safely store, secure and transport that company data. At the same time, you need to block every other device that doesn't meet your policy. In other words: why tolerate an iPod connecting to a workstation if that is not directly related to work and holds a serious security threat?

In order to solve this problem you have to start with a good plan for a security policy. After that you need full management support.

Corporate policy

We have to start with a corporate policy that says that no device can connect directly to a workstation or server, unless this is explicitly allowed by company rules. Also very important is to strictly forbid the taking of company data out of its context and to store that data on non-company devices. To make this work, you have to give your employees the devices they are allowed to store data on. The management is still responsible and accountable for this matter.

Part 1 - Working with secure USB devices

So now there is a policy. To facilitate the end-user there must be a solution for the secure transportation of data. Really stress the fact of safe transport, because most USB keys (not all!) only give a solution for the secure transportation of the company data. You can copy

the data on a memory stick, travel home and connect the memory stick at your private laptop at home and copy that same data to your local hard disk. This means that the data is still used in an environment where company rules don't count at all. So that is why - no data is allowed to be stored on non-company devices.

The most popular solution at the moment are biometric protected USB memory sticks. There are a lot of considerations to keep in mind here. Lets look at some points of interest concerning these devices.

Background fingerprint devices

Although some fingerprint recognition systems do the comparison on the basis of actual recognition of the pattern, most systems use only specific characteristics in the pattern of ridges. These special points are called *minutiae* and, although in general a fingerprint can contain lots of minutiae, the fingerprint area that is scanned by most sensors usually contains about 30 *minutiae*. For a positive identification at least 12 *minutiae* have to be identified in the fingerprint.

Types of sensors

Sensors that measure the temperature of a fingerprint can be smaller than the size of a finger. In general there are plate sensors and swipe sensors. Although either width or height should exceed the size of the finger, the other dimension can be fairly small since a temperature scan can be obtained by sweeping the finger over the sensor. The sensor contains an array of temperature measurement pixels which make a distinction between the temperature of the skin, the ridges and the temperature of the space between the ridges.

Accuracy

Is the fingerprint reader accurate enough concerning the point mentioned before? A fingerprint sensor can have an array of fingerprint sensor that consist of a 256 column x 300 row array of tiny sensors, giving a huge amount of sensitive pixels. Most of the times we talk about fingerprint sensor resolution (DPI) and a scan area in *mm*. Most common and adequate to capture a frame of the central portion

of a fingerprint is a 500 DPI scanner.

FRR and FAR


When there is biometric verification, a scan of a person is made and compared with the characteristics that are stored in a profile database. In general, a certain margin of error is allowed between the observed and stored characteristics. If this margin is too small, the system will refuse a legitimate person more often while if this margin is too large, malicious persons will be accepted by the system. These can be measured and are called False Reject Rate (FRR) and False Accept Rate (FAR). When using a biometric system, one would naturally want to minimize both rates, but unfortunately these are not independent. An optimum between FRR and FAR has to be

found related to the risk and protection. If you want to be certain no false acceptance can occur then the FAR will be very low or zero with the inevitable consequence of more rejections and stricter policies.

Standards and cost

It is a good idea to choose a product that meets certain regulations or government standards like the FIPS140-2, Security Requirements for Cryptographic Modules. For more information visit csrc.nist.gov/cryptval.

Biometric USB sticks are much more expensive than regular ones. Also, the costs of the implementation and administration should play an important part in the decision to choose one of these devices.



Installation and general functioning

To install specific software or drivers on an OS like Windows XP, you'll need administrator rights. In a bit larger infrastructure like the one I'm working in, software solutions like Microsoft Systems Management Server (SMS) in combination with MSI, install software under specific local system rights. In many cases a normal user will not be granted that right. It is best to choose a device that is driver-less which will enable the user to use the device without having to install additional software.

Encryption protection

What kind of encryption and protection is being used? Some USB devices do have biometric protection but store the data in a non-encrypted standard. So if you can temper with the chip or break the seal it is possible to extract the unprotected data. Other devices do have strong encryption and a complete solution (like the one Utimaco offers), makes data worthless if it is taken out of the company in-

frastructure. You will need the necessary key to decrypt the information stored.

Some manufacturers don't give insight into the encryption method or algorithm they use and I'm not sure you should trust them. There are minimal directives that are formulated by certain public authorities such as the NSA and the FBI. The new encryption standard is AES and the algorithm is Rijndael. A variety of USB sticks use this protocol with a key length of 256 bit and I think that is sufficient.

Administration of devices in organizations

There has to be control in case a company uses portable storage devices. In a large organization you must streamline the process of handing out such devices and give some support in case a stick is lost, blocked or a help-desk call is being made about it. It's a good idea to allow the administrator access to the stick, to pre-define some settings like the times you can have a false try after which the device is blocked.

This is just the top of the iceberg as large organizations will have more issues. You have to decide what kind of solution is best for you, but do it based on valid arguments: what skills do employees have and how much money and support can be spend for such a solution.

For small organizations you can also think about encryption solutions like PGP or TrueCrypt. Both are excellent tools and defacto standards on the market and certainly capable to encrypt data in a secure way.

For serious biometric protected USB sticks you can have a look at the following manufacturers and types where you must keep in mind that this is not meant to be a complete list:

The RiTech BioSlimDisk iCool, Ritech BioSlimDisk 2.0, Kobil mIDentity, MXI MXP Stealth, SafeBoot Phantom, Kingston Data-Traveler Elite, Transcend Jetflash 210, SanDisk Cruzer Profile.

Part 2 - Device control

Now that we have a security policy in place, we need to take additional measures to protect your infrastructure against unwanted devices. The need to block every other device that no longer meet your policy.

Here is Device Control coming in the picture. Device control makes it technically possible to block all uncertified devices. When you think about Device Control you can look (not to be complete) at the following solutions:

Safeboot | PortControl,
Sanctuary | SecureWave,
Smartline | DeviceLock,
Utimaco | Safeguard Advanced Security,
UBM | Drivelock,
GFI | Endpoint Security.

There are several suppliers / manufacturers that are bringing solutions on the market. Because this kind of solutions are relatively new, it is hard to find reports and tests about them. Keep in mind that there has to be a distinction between client protection (like Safeboot and Utimaco offers) and Device control solutions. Keep in mind that I'm limiting this article to the solutions based on the Windows platform.

Working with Active Directory and Group Policy

You can achieve your goal when working with group policies. If you are in a Windows domain environment, then you can apply the security setting to the registry and block access to `HKLM\System\CurrentControlSet\Enum\USBSTOR`

Since the system must write data here for the USB key to work, you as an administrator can block it from reading the data and this would stop the USB memory key from working. You could control this with Group Policies in a Windows 2000 or 2003 Active Directory as well as a Windows XP workstation environment.

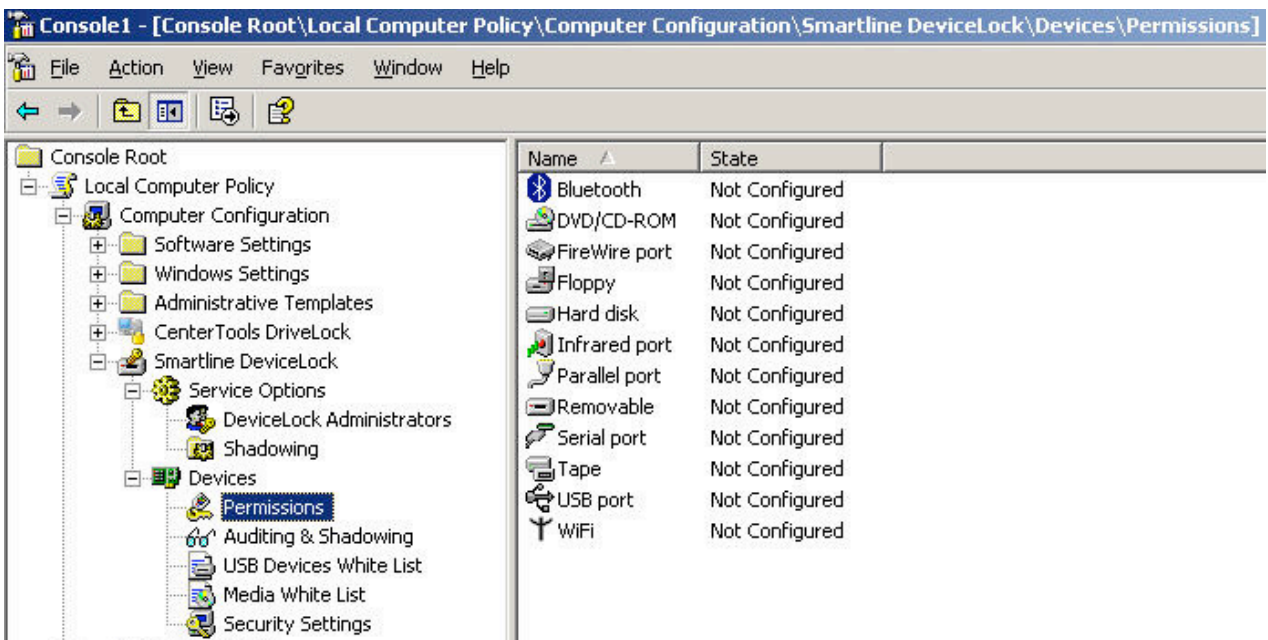
This way you can enable or disable all the USB interfaces or just USB storage devices. Because this registry setting falls under the concept of tattooing, even if you remove the policy, the setting remains there. The biggest disadvantage is that this is not really fine-grained, it has limited possibilities and is hard to achieve. Why not buy a mature product to help you?

The underlying concept is relatively simple: create a white-list of devices that you accept or are part of your standard workstation suite and all other devices that are not meeting your policy and need to be blocked (by default).

In every case, the software will contain a server and client portion. The server in that case will establish the centralized control of connected devices. It will also store policy data in the Active Directory or a separate database. The client part has to be installed on the workstation that controls the blocking or acceptance of devices that users connect to the workstation.

Invoke at kernel level

The client software must protect the workstation completely and therefore it has to be incorporated on the kernel level of the operating system. You need a certified solution and you most certainly don't want to be confronted with 10,000 workstations acting buggy after the client installation.



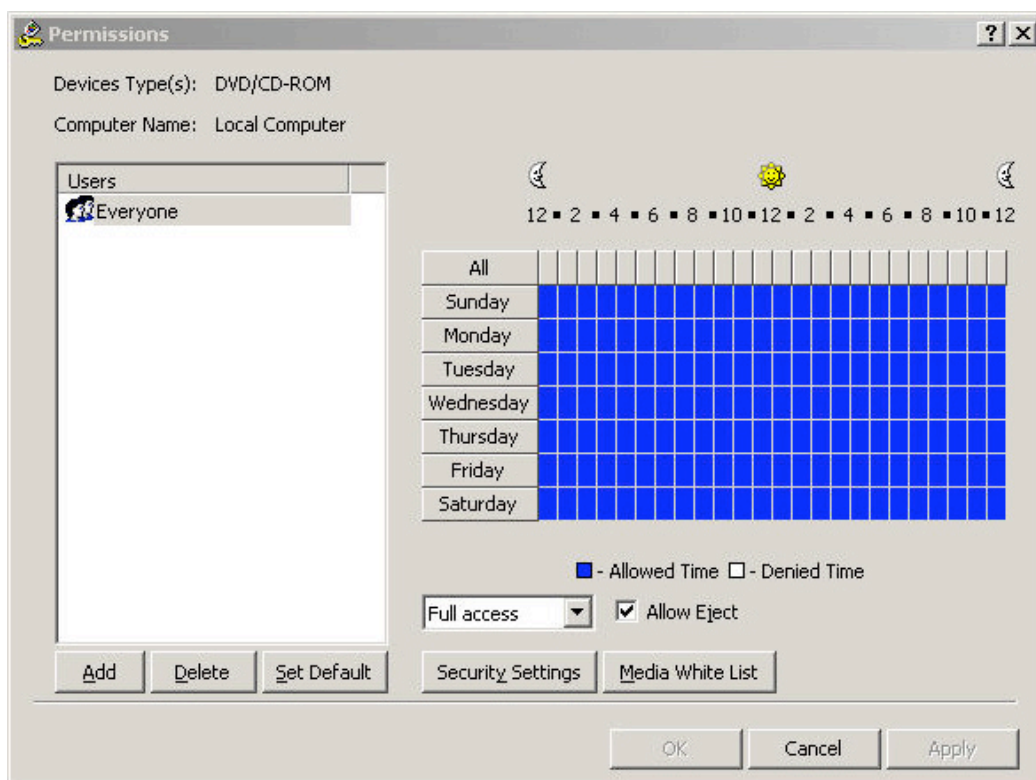
Creating policies for devices with DeviceLock.

Some other practical things the administrator has to keep in mind are the requirements of the software that comes with the device (like a version of .NET Framework) and the method used to install the client. If the installation can be distributed with a tool over the network it makes the setup much quicker.

When the software comes with an interface that is easy to understand and use, the learn-

ing curve will be low and the implementation quicker.

To be compliant and to meet certain rules and laws you have to implement exceptional logging and auditing features. Is the product supporting this kind of logging and reporting functions? Is logging stored on a central location automatically or on the workstation itself? All of these are valid questions you need to answer before implementation.



Tuning access control to devices with DeviceLock.

Working with policies

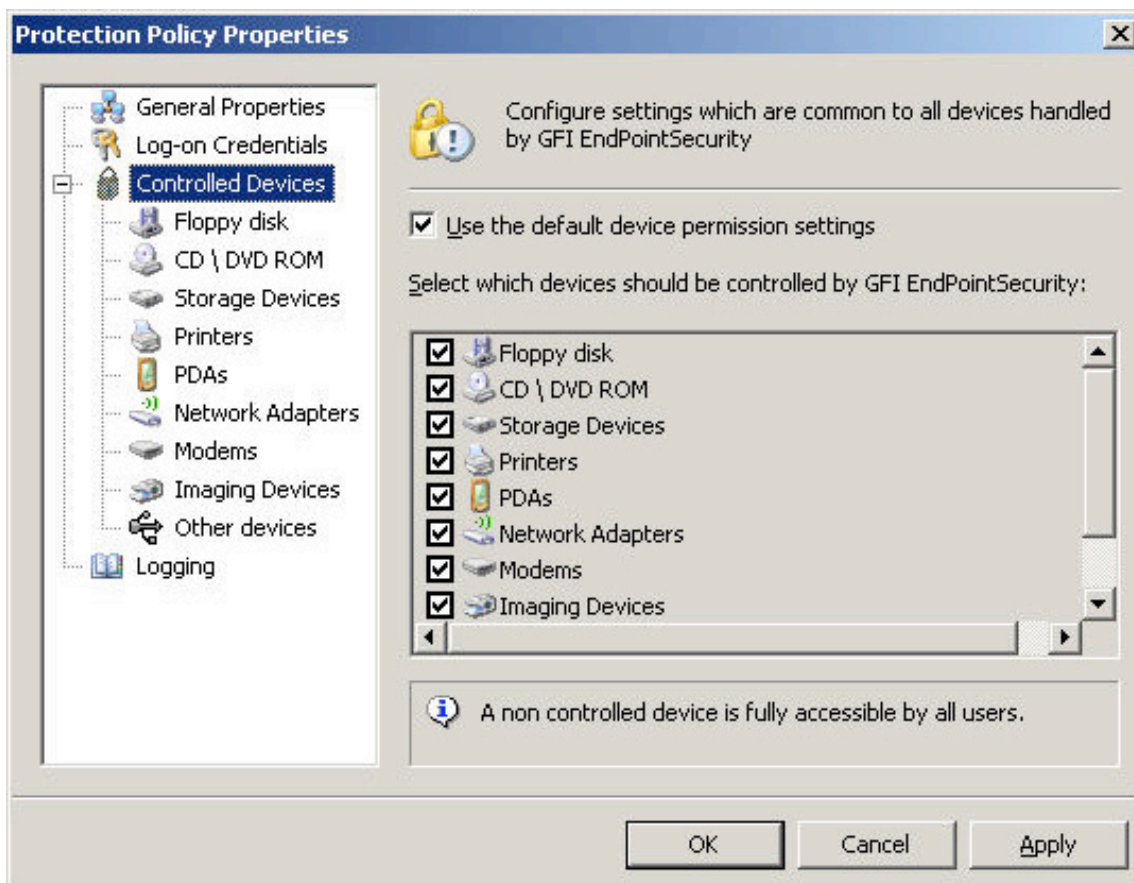
It is a good idea to make an overall policy that blocks devices. This way you can make decisions about exceptions for a specific group of users. A bonus is the opportunity to establish a connection to the Active Directory so you can use the already defined users and groups. What streamlines the integration is also the possibility to easily incorporate a new device in to the policy and also to block unknown devices by default.

White-listing

An interesting feature would be to make a difference between a whitelist of USB devices

such as specific biometric USB sticks that are allowed, handed out and registered by your organization, and the same type that is not handed out by your organization and bought privately by an employee. Some biometric USB sticks have the possibility to store a unique ID or have unique media descriptors so you can distinguish them from others.

Can you prevent specific file types to be transferred? It would be good if you can allow the transfer of .jpg files from a memory card that is used in a digital camera but you may want to prevent the transfer of a text document to that same camera because it is not something that should normally occur. Some of the solutions on the market do have this possibility.



Controlling your environment with GFI Endpoint Security.

I hope the information discussed in this article will help you keep yourself on track and make the proper decision for your organization. Some products are more mature than others

and so be aware that there is a big difference in understanding the demand of large organizations and the suitability of the products.

Rob Faber is an infrastructure architect and senior engineer. He works for an insurance company with 22.000 clients in The Netherlands. His main working area is (Windows Platform) Security, Active Directory and Identity Management. You can reach him at rob.faber@icranium.com

Confidential Notes is a practical and easy to use solution that instantly provides you with a high level of security for your mobile data.

For more information on Confidential Notes visit www.pocketpcsecurity.com



Confidential Notes 13:39

Enter password 1:

Enter password 2:

Forgot password? Enter

123 1 2 3 4 5 6 7 8 9 0 - = <

Tab q w e r t y u i o p []

CAP a s d f g h j k l ; ' <

Shift z x c v b n m , . / <

Ctl á ü ` \ <

Confidential Notes 13:17

Main Folder	Date	
ipaq software	13:08	4k
inet banking info	13:06	151k
shopping weekend	13:04	149b
target market	13:04	2k
city center plan	13:03	1k
dan's cellular	13:02	29b
early sketches	13:01	1024b
audio Q&A in NY	13:01	245k
wilderness sounds	13:00	225k
anna's NYSE column	12:59	892b
stock portfolio	12:58	1k
apple store london	12:57	3k
VC capital thoughts	12:57	145k

New Options

Confidential Notes 12:26

interview with the marketing manager

ARTICLE

Besides the overview on the success of the past year's event and a very positive forecast for this April's conference, journalists were presented with a rather new concept in the field of IT events - assistance for overseas visitors. I should note that he term "overseas" in this case is obviously connected to visitors outside the United Kingdom. As the Infosecurity conference is UK's top information security conference, UK Trade & Investment, the British Government agency that supports overseas enterprises

New Edit Options



Events around the world

LISA '06: 20th Large Installation Systems Administration Conference

3 December-8 December 2006 – Washington, D.C.

<http://www.usenix.org/lisa06/>

Black Hat DC Briefings & Trainings 2007

26 February-1 March 2007 – Sheraton Crystal City

<http://www.blackhat.com>

InfoSecurity India 2006

5 December-7 December 2006 – KTPO, Bangalore, India

<http://www.infosecurityindia.com>

2nd European Conference on Computer Network Defense

14 December-15 December 2006 – School of Computing, University of Glamorgan, UK

<http://www.comp.glam.ac.uk>

23rd Chaos Communication Congress

27 December-30 December 2006 – Berliner Congress Center, Berlin, Germany

<http://events.ccc.de/congress/2006>

RSA Conference 2007

5 February-9 February 2007 – Moscone Centre, San Francisco, USA

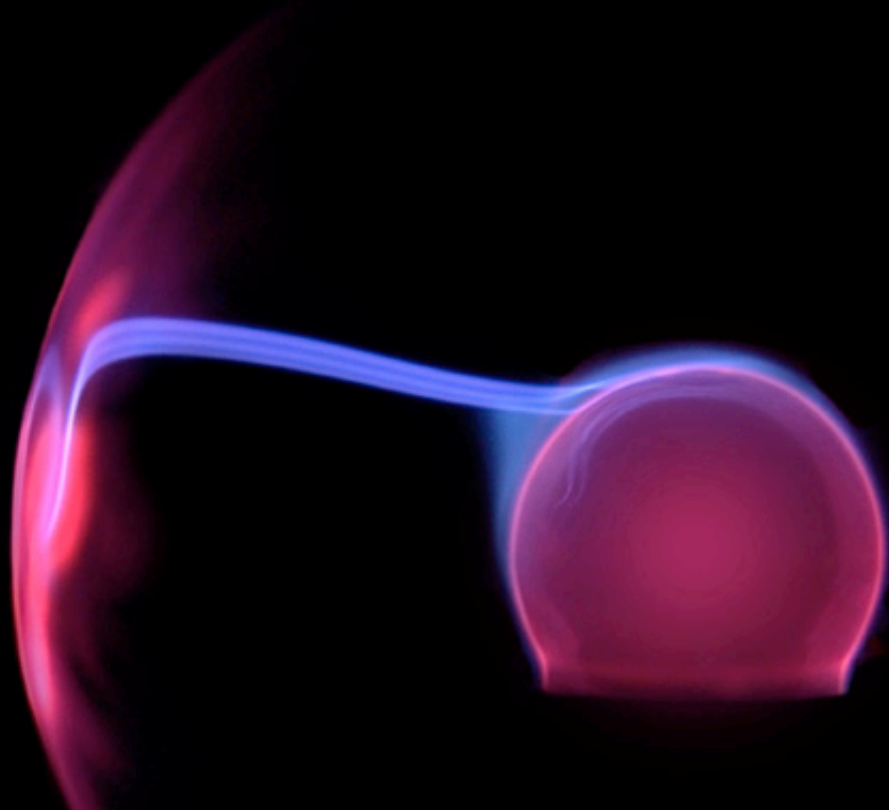
<https://2007.rsaconference.com/US/>

InfoSecurity Italia 2007

6 February-8 February 2007 – Fiera Milano, Italy

<http://www.infosecurity.it>

If you want your event included in the HNS calendar e-mail us at press@net-security.org



Enterprise data security - a case study

By Ulf T. Mattsson

This article is a case study about an Enterprise Data Security project including the strategy that addresses key areas of focus for database security encompassing all major RDBMS platforms. It presents the current state of database security tools and processes, the current needs of a typical enterprise, and a plan for evolving the data security.

This strategy will help set direction for the blueprint of data security and provide a composite high level view of data security policies and procedures for the purpose of satisfying growing regulatory and compliance requirements and develop high level timeline and for all steps of development. This article presents a three steps strategy to address current outstanding audit concerns and positioning to more readily address the evolving regulatory landscape.

1. Overview

As security, regulatory, and compliance pressures continue to be a key driver for XYZ Company, the technical environment supporting our business will need to be continually reviewed and enhanced to ensure all requirements are met. The database environment is

extremely sensitive based on the fact that a large percentage of data at XYZ Company resides in our RDBMS platforms. These environments have been audited and scrutinized on a regular basis and will continue to be as we move forward. Although the database environments at XYZ Company are protected by tightened perimeter security measures, advanced authentication, authorization and access control security measures, and are considered to be a secure environment which effectively protect XYZ Company data from external intrusions, we must continue to look for opportunities to increase the overall security and compliance of these environment based on evolving needs, as well as, new technologies that can enhance the environment. Through compliance activities such as internal audits, SOX, GLBA, PCI, and others, other opportunities have been identified to better

secure this environment.

2. The primary problem

The primary problem with many compliance initiatives is a focus on existing security infrastructure that addresses only the network and server software threats. But the data security capabilities required to be compliant goes far beyond these technologies. Network and server software protections (network firewalls, Intrusion Prevention Systems), while important, provide no insight into data-level attacks targeted directly against a database or indirectly via a web application. Regulatory compliance requires an understanding of who is allowed to access sensitive information. Regulatory compliance requires an understanding of who is allowed to access sensitive information? From where did they access information? When was data accessed? How was data used? The bottom line is that data security requires a new approach that extends the breadth and depth of IT's ability to secure information.

2.1 Stronger database security is needed to accommodate new requirements

Another driver is our extended partnership with non-XYZ Company parties, more and more tasks will be performed outside the physical boundaries of our facilities which will add another level of due diligence we must take into account. Stronger database security policies and procedures must be in place to accommodate the new environment. Centralized database management security must be considered to reduce cost. As we have been presented with opportunities to solidify the environment, we have continued to evolve the existing environment. This, at times, has led to implementing "point" or manual solutions which become harder to manage as the environment continues to grow and become more complex.

Centralized database management environment must be considered as a solution to increase efficiency, reduce implementation complexity, and in turn to reduce cost.

DATABASE SECURITY IS AN ONGOING PROCESS

3. A solution that is addressing external and internal attacks

Understand that database security is an ongoing process. More and more enterprises make database security a top priority to meet growing compliance requirements and to protect themselves from increased intrusions – both external and internal attacks.

3.1 Define strong policies and procedures

Work with auditor, security group, and IT department to outline strong policies and procedures for databases. Information security policies and procedures should dictate databases' security policies and not vice versa. Revisit security policies and procedures every quarter to ensure that they continue to meet business requirements, and strive to adapt to newer technologies. Each compliance requirement is different; therefore, make sure to understand each compliance implication for the enterprise databases. For example, SOX mainly requires that production financial databases be protected and no inappropriate changes be

made, while HIPAA requires that all personnel information be protected from unprivileged users in all environments, including test and development.

3.2 Focus on an overall, unified security strategy

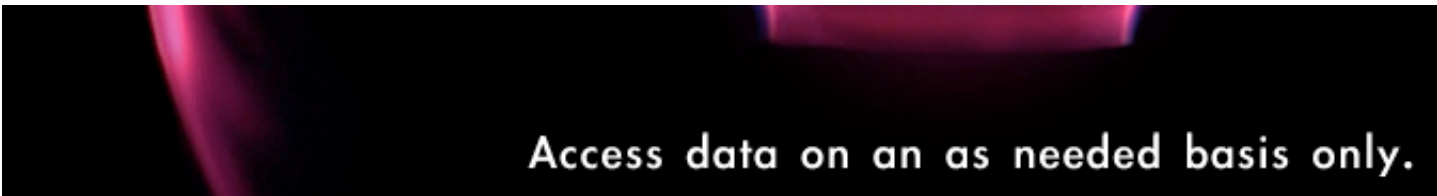
To have a robust security implementation, database security must be integrated with application-, IT-, network-, and infrastructure-level security. End-to-end security implementation should be the goal for enterprises. DBMS vendors are churning out security patches faster than ever before as new vulnerabilities are discovered. Although security patches are critical, not all databases need them, so check to ensure that they are applicable. While DBMS vendors will continue to work on simplifying security patch deployment, enterprises are seeking security patch management solutions to ensure critical patches are applied in a timely manner. Documentation remains important, not only for formalizing data security practices, but also in a court of law, should the situation arise.

4. The enterprise database environment

The typical enterprise database environment today, at a high level, consists of a selection of RDBMS platforms including Oracle, DB2, Sybase, Teradata, Informix, and MySQL. The server platforms typically includes UNIX, Mainframe and Windows. Databases reside on multiple network segments and protected by network security measures. Those network segments also include Development and Contingency Production. Overall, the database and application tiers are separated onto individual platforms. Security Facilities typically includes DB2 security via RACF, TopSecret or ACF2, and Oracle Security and Sybase Security.

4.1 Growing percentage of internal intrusion incidents

While XYZ company databases are protected by perimeter security measures and built in RDBMS security functionality, they are exposed to legitimate internal users at some degree. Due to the fragmented distribution of database environments, real time patch management, granular auditing, vulnerability assessment, and intrusion detection become hard to achieve. With the growing percentage of internal intrusion incidents in the industry and tougher regulatory and compliance requirements, XYZ company is facing tough challenges to protect XYZ company sensitive data against internal threats and meet regulatory and compliance requirements.



Access data on an as needed basis only.

5. Enforceable database security policy at an enterprise level

Define enterprise level, enforceable database security policy and procedure. This must include separation of duties. Some encryption is today performed based on individual application needs but not across the board. Many user administration tasks are split between a Central Security Team and Corporate IT with a medium degree of manual processes.

Some corporate IT database security policies are in place and enforced. Examples including Oracle Patch Management Policy, Oracle Security Standards, Oracle Userid Management Policy, Oracle Schema Owner Database, Password Management Policy, Oracle Sys and System Database Password Management Policy and 'Corporate Data Architecture' Standards.

5.1 Control data access and DBA accountability

Access data on an as needed basis only. Refine production application data access role, to allow access to only the necessary data and privilege based on different business function. Read only access should be limited, devel-

oper access to application accounts (backdoor data access) should be prohibited. Refine DBA role into different categories. For example, production support level I, production support level II and application support role.

Only necessary role should be granted to DBA based on their individual responsibility for various applications. DBA should not access application data on regular basis, but a close review of database administration for different RDBMS platform must be performed to ensure that the ability to respond and resolve is not hampered dramatically that eventually affect the business.

5.2 Establish database security officer and user administration

Establish database security officer (group) to define and enforce database security policies and procedures, to close monitor industry trends and adopt new technology. For example, administer "database firewall" and generate audit report to comply with different regulatory requirement. Database user account should be managed (creation/modification/removal) by security group and centralized user management should be adopted to simplify the process.

5.3 Protect production data in non-production environments

Test databases are often replicas of production data. All non-production databases that hold personal identifiable data should enforce strong security measures. Production data must not exist in development environment. Due to the nature of test environment, production data is needed at times functionality and performance testing, the same production security measures must be applied in the test environment. We have found this to be among the most common security vulnerabilities. It is also specifically mentioned in the PCI DSS requirements.

Production data that is used in a development environment should be replaced with either scrambled or surrogate data. While application development groups have often designed their test procedures to use production data, it is far less risky, if somewhat more costly, to use encrypted or surrogate data for testing purposes. It is generally less costly to encrypt confidential data (without changing its data type) so that it can be used in development and test environments. Because this can affect development productivity, we consider this to be a relatively long term project.

5.4 Use of granular auditing, vulnerability assessment and intrusion detection

As our user base becomes more varied and wide, the ability to monitor and detect inappropriate behavior becomes even more critical to ensuring that our information is protected. Taking into account the aforementioned requirements, auditing of activity adds another level of detection that can be utilized to enhance overall security and meet regulatory and compliance requirements. This must be done as efficiently as possible; the following functions must be considered.

5.5 Use of database activity monitors, audit and database vulnerability scanners

Network appliances or servers that monitor database and log activity that is external to the database server, and can generate real time alerts based on unusual behavior or policy violations. Be able to collect and store a rich set of audit data and provide built-in reporting

capabilities flexible enough to meet all internal or external compliance requirements. For example, PCI requires one year audit data that include all accesses to card holder private data. Software tools for scanning databases for known vulnerabilities. Those tools are similar to other vulnerability scanners, but can perform more-advanced database configuration and structural scans.

6. The need for heterogeneous database platform support

All XYZ Company database platforms should be supported with a minimum Impact on Database Performance, Stability, or Administration. The solution should have minimum or zero impact on database performance and stability, and should be administrated by security officer with minimum database expertise requirement. RDBMS security patches must be reviewed regularly; Centralized patch management solution must be in place to ensure critical security patch are applied in a timely manner.

6.1 Database encryption

Encryption is a critical component of data protection. If an intruder circumvents other network security controls and gains access to encrypted data, without the proper cryptographic keys, the data is unreadable and unusable to that person.

Other effective methods of protecting stored data should be considered as potential risk mitigation opportunities. For example, methods for minimizing risk include not storing the sensitive data unless absolutely necessary, truncating the sensitive data elements if the full content is not needed, and not sending sensitive data elements in unencrypted e-mails.

There are various levels of encryption that must be taken into account: at rest, in motion, based on environment, etc. While database-level encryption does not protect data from all kinds of attacks, it offers some level of data protection by ensuring that only authorized users can see the data and it protects database backups in case of loss, theft, or other compromise of backup media.

6.1.1 Implement column-level database encryption

We have worked with hundreds of financial services industry firms and merchants to implement column-level database encryption over the past several years. Most of these implementations have been driven directly or indirectly by the need to comply with the PCI DSS requirements for protecting confidential data at rest. One of the main (and comparatively difficult) requirements of PCI DSS is the requirement to encrypt cardholder data in databases.

While the release of the PCI 1.1 specifications does allow the use of compensating controls to mitigate the database encryption requirement, it does explicitly exclude any controls mentioned in the PCI DSS requirements from being considered a compensating control. As a result, controls such as firewalls, network access controls, passwords, anti-virus, and most other network-based controls are not acceptable as compensating controls according

to the Appendix to the PCI DSS 1.1 requirements. There is some variability among assessors on the issue of certifying compensating controls, but because we know that most major financial institutions and many of the largest merchants have implemented column-level database encryption, we believe this is a reasonable project for XYZ. This recommended action is also associated with database security, but it is included here because we have found that mainframe databases of financial institutions contain such high volumes of confidential data that they present a very attractive target.

Most of the risk is from internal threats – persons who are able to view confidential data and, in some cases, may be able to copy confidential data to less-secure environments. While all financial institutions in our experience, has extensive mainframe access controls (e.g., ACF2), the vulnerability to internal threats is one of the main reasons why the PCI DSS requirements include the encryption of cardholder data in databases.

ONE OF THE MAIN (AND COMPARATIVELY DIFFICULT) REQUIREMENTS OF PCI DSS IS THE REQUIREMENT TO ENCRYPT CARDHOLDER DATA IN DATABASES.

6.1.2 Software to encrypt data is more scalable and performs better

Network-based hardware appliances for encrypting data with relational database management systems at the table or column level, through traffic interception or API calls are unfortunately not providing the off-loading encryption load that was expected, but can be effectively be used for key management. Network-based hardware appliances for key management. Software to encrypt data at the table or column levels within relational database management systems (RDBMSs) is far more scalable and performs better on most of the platforms in an enterprise, when executing locally on the database server box.

6.1.3 Implement centralized key management

As XYZ adds several more data encryption projects over the next year, the number of key management systems and APIs will increasingly make it difficult to manage all the differ-

ent key management systems without some type of centralized architecture and management too.

6.2 Use of database only network segment(s)

Currently there are multiple different database segments based in different geographic locations that require administration. This environment is difficult to administer due to "over segregation", sometimes people have to break "other" rules to administer effectively. Many security/auditing tasks have to be duplicated for every environment where database resides. We need to explore the viability of this approach. The database only network segment(s) have advantages including centralized entry point to manage and monitor all activity, administrative tools can effectively manage security/auditing tasks, database environments are brought together, in a reduced number of environments, in "back office", and separate database from application further reducing access to environments.

Adequate network segmentation, which isolates systems that store, process, or transmit sensitive data from those that do not, may reduce the vulnerability of the data environment. Network components include firewalls, switches, routers, wireless access points, network appliances, and other security appliances. Server types include but web, database, authentication, mail, proxy, and DNS. Applications include all purchased and custom applications, including internal and external (Internet) applications.

7. General positioning of compensating controls

Compensating controls may be considered when an organization cannot meet a technical specification of a requirement, but has sufficiently mitigated the associated risk. The effectiveness of a compensating control is dependent on the specifics of the environment in which the control is implemented, the surrounding security controls, and the configuration of the control. Organizations should be aware that a particular compensating control will not be effective in all environments. Each compensating control must be thoroughly evaluated after implementation to ensure effectiveness. For organizations unable to render sensitive data unreadable (for example, by encryption) due to technical constraints or business limitations, compensating controls may be considered.

7.1 Perform a risk analysis before using compensating controls

Only organizations that have undertaken a risk analysis and have legitimate technological or documented business constraints should consider the use of compensating controls to achieve protection. Organizations that consider compensating controls for rendering sensitive data unreadable must understand the risk to the data posed by maintaining readable data. Generally, the controls must provide additional protection to mitigate any additional risk posed by maintaining readable data. Compensating controls should consist of a comprehensive set of controls covering additional segmentation/abstraction (for example, at the network-layer), providing ability to restrict access to data or databases based on IP address/Mac address, application/service,

user accounts/groups, Data type (packet filtering), restrict logical access to the database, control logical access to the database (providing separation of duties) and prevent/detect common application or database attacks (for example, SQL injection).

7.2 Conclusion regarding compensating controls

The basic conclusion from this analysis is that a combination of application firewalls, plus the use of data access monitoring and logging may, if effectively applied, provide reasonable equivalency for the use of data encryption across the enterprise. Such a combination of controls does have some weak spots, mainly when it comes to preventing damage from careless behavior of employees or weak procedures in development and separation of duties. Also, since some of XYZ's controls are specified in the PCI DSS requirements, a separate PCI assessment of whether these controls meet the PCI DSS definition of compensating controls is recommended.

8. File encryption

This is a top priority due to the fact that there are large numbers of flat files and other data files on a large number of Unix servers throughout XYZ which contain credit card PANs and other confidential data, but these files are not currently encrypted. At the same time, our experience with comparable situations in many other enterprises has shown us that many such files can be encrypted with minimal performance impact. The cost of this type of encryption (whether software or HSM) is comparatively less than other types of encryption, yet the benefit in terms of complying with the PCI DSS requirements for data at rest encryption (Requirement #3) offer a significant upside. This type of encryption project can be deployed more gradually than other types of encryption (e.g., Microsoft's EFS), which is another reason to make this a high priority project in implementing the XYZ Encryption Strategy. Encrypting bulk file transfers is another candidate for gradual deployment, and at a relatively lower cost. The primary drivers which helped make SAN/NAS encryption a top priority are the large volume of unencrypted confidential data stored in XYZ's SAN/NAS environment and the value to XYZ of

complying with the PCI requirement for encrypting sensitive data at rest (Requirement #3). The ability to roll out SAN/NAS encryption gradually is another argument for making it a relatively high priority.

9. Project implementation steps

In order to implement solutions mentioned above, we have divided efforts into Steps. Initiatives to prevent immediate threats and resolve open audit concerns are addressed in Step 1, Step 2 will continue efforts to enhance and refine our environment to meet regulatory and compliance requirements. Step 3 will include efforts to further reduce database security risks efficiently and effectively, and to address new challenges as environments continue to evolve.

All efforts within this strategy will be coordinated, where appropriate, with other projects ongoing at XYZ Company including (but not limited to) Data Protection Strategy, GLBA Remediation and PCI Compliance.

9.1 Project step 1

9.1.1 Define and implement DBA roles

Define and implement DBA roles and privileges on the Oracle environment only in Step 1. Due to less sensitive data, Sybase and other platforms will be Step 2. The following key drivers are being used to build the necessary roles- DBAs should not perform database user administration tasks, these tasks should be managed by Central Security Team, DBAs should not be allowed to view sensitive production application data on a normal basis.

As production issue arises, DBAs should be able to view production data to debug and solve the production issue but necessary audit trail must be provided. Key milestones include activities to develop and test solutions to meet above requirements for each DBMS platform, and have the DBA group define role template for different level of DBA support. The Central Security Team will also take over user administration responsibility, and define solutions to audit DBA user activities. The solutions are implemented into development, test and production environment in this Step.

9.1.2 Review production data access privilege for non-DBA accounts

Developer should not have access to sensitive production data on a normal basis. Necessary audit trail must be provided when production issue arises and examining production data is part of the solution. Also, backdoor access (using generic application account) to production database should be prohibited. Database roles for application generic account and developer account need to be carefully reviewed and refined. Unnecessary access to sensitive data should be minimized. Key milestones for this Step include that the Primary DBA for each application review application role privilege, identify roles/accounts which have privileges to access production data, and to communicate with development team for each application, and document the usage of those roles/accounts identified in Activity 1.

9.1.3 Review database security technology

Perform analysis, via RFI Process, to review database security technology. Review new technologies that provide additional database security levels including database security firewalls, intrusion detection software, vulnerability assessment software, etc. These products will be based on the requirements mentioned in the previous section. Key milestones and costs include NDA preparation, vendor product review, analysis and document (including scorecards).

9.2 Project Step 2

All Activities included in Step 2 will be based on analysis, where needed. The subsequent funding of these Activities will be based on normal project process. Impacted groups are Corporate IT, Central Security Team, Network Support, and Corporate System Delivery, Legal. This step will encrypt files on Unix file servers, encrypt bulk file transfers and encrypt SAN/NAS storage. Project Step 2-A implement Sybase and DB2 DBA security roles. Project Step 2-B is to proceed with RFP process for any identified database security technology acquisitions based on analysis performed in Step1. This will include to complete RFPs, secure funding for effort, install technology for knowledge building, build processes and procedures and roll out product(s).

Project Step 2-C will update XYZ company database security policy for best practice and publication. This includes updating current policy and the addition of any new technology/processes being introduced into the environment. High level overview is included in previous section “Define enterprise level, enforceable database security policy and procedure”.

Project Step 2-D will explore database only network segment(s) for different geographic locations to further secure database environment for all platforms. Project Step 2-E will refine and deploy Security Patch management process to ensure RDBMS vendor security patch are reviewed and critical patch are applied in a timely manner. Current process is adequate but the application of patches needs to be better managed for all database platforms. Project Step 2-F implements production data in test security measures and ensure no production data in development databases.

9.3 Project step 3 - Execution and enforcement

All Activities included in Step 3 will be based on analysis, where needed. The subsequent funding of these Activities will be based on normal project process. Impacted groups: Corporate IT, Central Security Team, Network Support. Project Step 3-A is to execute on

decisions, if any, based on database only network segment for different geographic locations and migrate database servers to identified segment. Project Step 3-B is to enforce the XYZ company database security policy at corporate level. Ensure the policy is adopted in the entire application development cycle. This step will encrypt static application passwords used to access databases, encrypt user-specific device passwords and add encryption of databases at column-level, encrypt interim/backup files and encrypt production data in development and test.

10. Conclusion

Database security is becoming top priority of enterprises to meet growing compliance requirements and protect sensitive data from increased intrusions. By implementing solutions documented above, we should be in a better position to face growing database security challenges, to proactively meet regulatory and compliance requirements and to better control our sensitive data. Database security is an ongoing process, we must revisit and refine our strategy regularly to adopt new technologies and address new challenges as environment continue to evolve. Due to the complexity of XYZ company’s current network layout, in-line intrusion prevention is currently not included in the project at this point.

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.

The art of information security awareness



A smart way to reach total security.



InfoSecurityLab

www.infosecuritylab.com

Sponsored by

USENIX & [sage]
THE USENIX SIG FOR
SYSADMINS

6 DAYS OF TRAINING
BY INDUSTRY EXPERTS, INCLUDING:

- Gerald Carter on Ethereal and the Art of Debugging Networks
- Richard Bejtlich on TCP/IP Weapons
- Aleen Frisch on Administering Linux in Production Environments
- Chip Salzenberg on Higher-Order Perl
- And 55 other tutorials

LISA'06

20TH LARGE INSTALLATION
SYSTEM ADMINISTRATION CONFERENCE

A **Blueprint** for Real World
System Administration

DECEMBER 3-8, 2006 | WASHINGTON, D.C.

3-DAY TECHNICAL PROGRAM

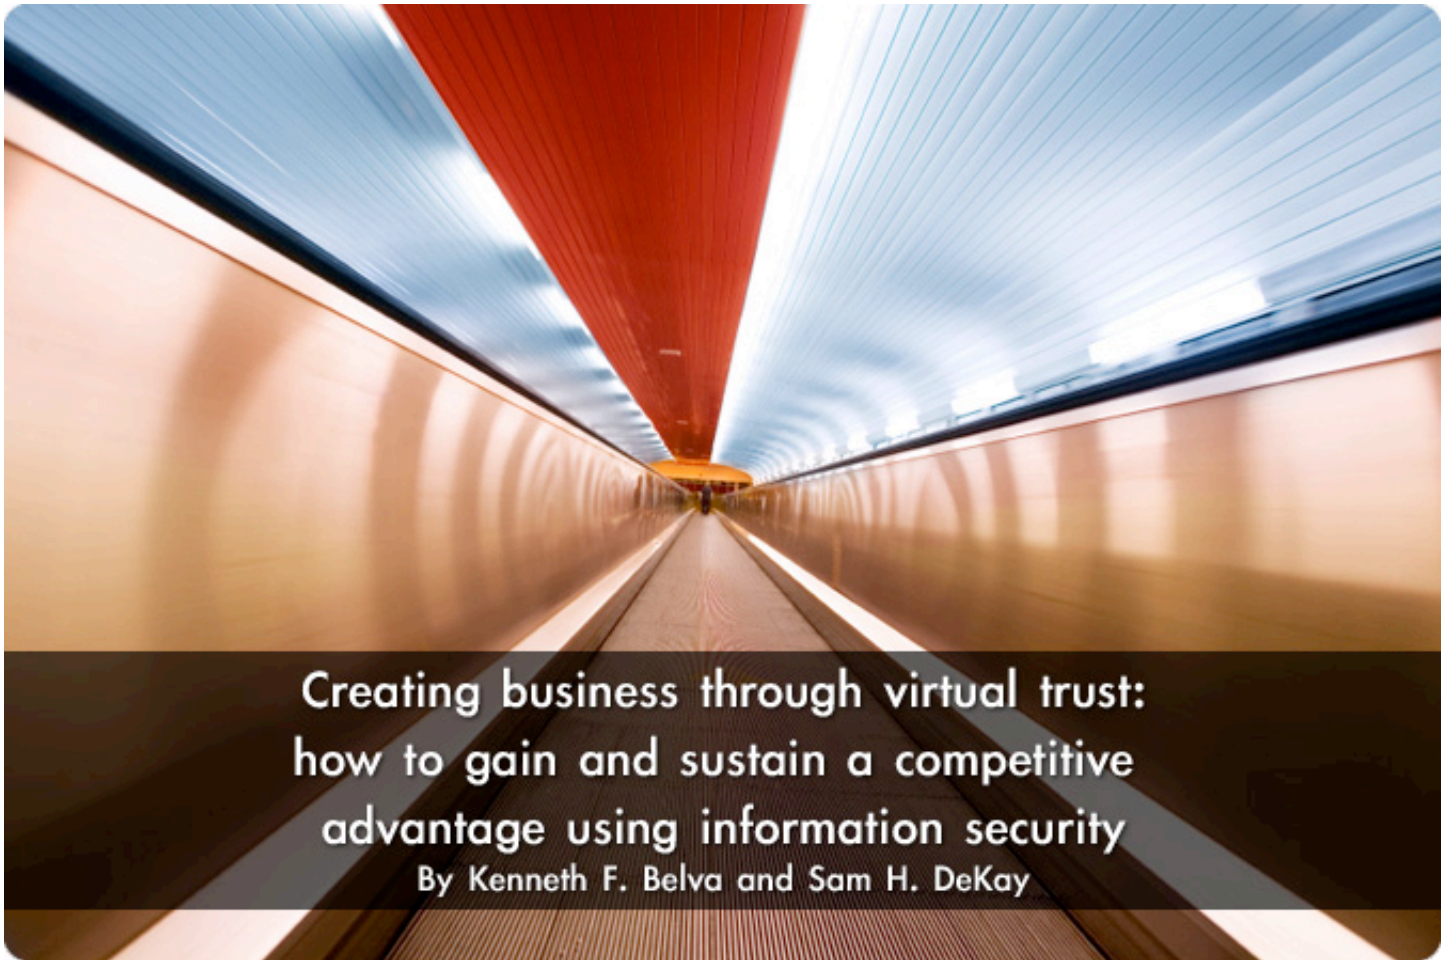
Keynote: Cory Doctorow, science fiction writer, co-editor of Boing Boing, and former Director of European Affairs for the EFF, on Hollywood's Secret War on Your NOC

20+ Invited Talks, including:

- Simple Nomad, Vernier Networks, Inc., "Corporate Security: A Hacker Perspective"
- DJ Byrne, Jet Propulsion Laboratory, "Open Source Software and Its Role in Space Exploration"
- Mazda Marvasti, Integrien, "Everything You Know About Monitoring Is Wrong"

Refereed Papers, Hit the Ground Running Track, Guru Is In Sessions, Vendor Exhibition, Workshops, BoFs, WIPs, and more!

LISA '06 offers the most in-depth, real-world system administration training available!



Creating business through virtual trust: how to gain and sustain a competitive advantage using information security

By Kenneth F. Belva and Sam H. DeKay

For at least 15 years, information security professionals have debated whether their work consists essentially of protecting the confidentiality, integrity, and availability of data, or whether security is a critical enabler of business.

Although the debate has frequently been joined by articulate and passionate voices, the argument often assumes the form of a semantic game: Is the purpose of information security merely to prevent loss, or is the preventing of loss a means of enabling business? This debate, framed within the dualism of loss prevention v. enablement, remains a favorite topic in books, journal articles, conference presentations, websites, blogs, and wherever two or more IS professionals are gathered.

This article is yet another participant in the ongoing discussion. However, the purpose of the paper is to introduce new rules to the debate, to propose new ways of framing the oft-repeated arguments.

The article maintains that the “either/or” approach to information security—either the primary function of security is to protect information assets or to enable business—is not a helpful dichotomy. Rather, a new hypothesis is

proposed: certain mechanisms and protocols that comprise security technology—including, but not limited to, authentication and encryption (as used in DRM and SSL)—actually create (not merely enable) business and makes possible the generation of revenue. Other functions associated with information security, such as intrusion detection and prevention, are more closely associated with the role of asset protection.

Security mechanisms that create business establish “virtual trust” between consumers, businesses, and government entities. This trust is “virtual” because e-business transactions involve consumers and commercial entities that are not physically present; without “virtual trust”, the transactions will not occur and business relationships cannot be built.

The article compares the concept of virtual trust to the traditional notion of information security as providing protection against loss

I. A Tale of Two Paradigms: The “Insurance Model” and “Virtual Trust”

The CIO of a major bank in Australia, speaking before a gathering of his peers from other financial institutions, recently announced that they “should use the tactics of Fear, Uncertainty, and Doubt to convince senior management to invest in security”. “While senior management may be aware of the risks to their information infrastructures,” he advised, “they often do not fully understand the damage that a breach in security can cause a business. Fear, Uncertainty, and Doubt can also motivate board members to take direct action to mitigate risks”.

The Australian CIO did not originate the term “Fear, Uncertainty, and Doubt.” In fact, these three words are so familiar to many information security professionals that they have become abbreviated as an unappealing acronym, FUD. Originally invented to describe the sales tactics of major hard- and software manufacturers, FUD has now become associated with a persuasive means of convincing corporate managers that human and financial resources should be allocated to the information security function. A prominent American technology professional, quoted in CIO Magazine, claims simply: “Fear, uncertainty, and doubt—FUD--has been used to sell security. If you scare them, they will spend”. Information security is the antidote to FUD; its purpose is to introduce controls to dispel fears of losing data, funds, and privacy. As succinctly stated by Shelton Waggener, another American CIO: “Security is really an insurance policy”. Eric Goldman, Director of the High Technology Law Institute at the Santa Clara University School of Law, offers a rephrase of Waggener’s sentiment: “There is no real wealth created by the investments in security, it is just a cost of everything we do in our lives”.

This article offers an altogether different view of the function of information security. We propose that information security is not just a kind of insurance, but a means of actually creating business and generating profit. We suggest that the typically double-negative rationale that justifies the existence of security—preventing loss—no longer accurately describes the full role of information security in

banking, commerce, education, health, and law.

Since the 1960s, the dominant paradigm, or understanding, of information security has been the prevention or mitigation of loss. To paraphrase Waggener, the “insurance model” has been accepted as the governing concept by which information security justifies its existence. And this model has, indeed, been successful: in fiscal 2006, the U.S. federal government alone spent \$5.1 billion on products, personnel, and services that prevent, or reduce the likelihood of, incidents that may adversely affect the confidentiality, integrity, and availability of data.

Within the last ten years, however, information security has commenced to serve a new purpose: establishing trust between people and between businesses and their customers. This new purpose has implications, not only for our understanding of the functions of information security, but also for future legislation, technology, and business opportunities. The function of establishing trust may also transform the traditional approach toward information security—that of a cost center—to a new view of security as a critical enabler of business.

We are not claiming that the “insurance model” must be abandoned and replaced with the new concept of “virtual trust.” However, we hope to present evidence demonstrating the emergence of a new role for security as a driver of business. Information security can no longer be characterized as a field and practice that merely prevents or mitigates loss.

For nearly forty years, the “insurance model” has provided a clear and compelling means by which to perceive the essential rationale and function of information security. In October, 1967, the Defense Science Board Task Force on Computer Security was commissioned by the Department of Defense to develop and document effective measures to secure data processed by resource-sharing systems. The Board’s published report, released in 1970, was entitled simply: Security Controls for Computer Systems. This document describes security as a “problem” that involves preserving the integrity of data and programs.

In 1973, Harry Katzan, Jr. authored *Computer Data Security*, one of the first publications concerning information security in non-government environments. Katzan, quoting an earlier IBM document, succinctly describes the essential function of security: “Data security can be defined as the protection of data from accidental or intentional disclosure by unauthorized persons and from unauthorized modification”. Three years later, Donn B. Parker’s *Crime by Computer* offered harrowing anecdotes concerning real-life crimes perpetrated via computer. Parker concluded that “‘Computer abuse’ is broadly defined to be any incident associated with computer technology in which a victim suffered or could have suffered loss”. Parker’s book was still available in bookstores when John McNeil published *The Consultant*, the first crime novel featuring computer fraud and a

technically-savvy detective; the novel offered a popularized notion of information security as an antidote to criminal activity.

The dual themes of loss and prevention, prevalent in the information security literature of the 1970s, continue to retain their potency today. CobiT, the organized set of IT controls recommended by the Information Systems Audit and Control Association, maintains that the function of systems security is “maintaining the integrity of information and processing infrastructure and minimizing the impact of security vulnerabilities and incidents”. The FFIEC IT Examination Handbook, familiar to most security professionals in the U.S. financial services industry, states that “An information security strategy is a plan to mitigate risks while complying with legal, statutory, contractual, and internally developed requirements”.

PREVENTING LOSS, THE MAJOR FOCUS OF THE “INSURANCE MODEL” OF INFORMATION SECURITY, IS SO PERVASIVE THAT THE PARADIGM HAS SPAWNED A RELATED EFFORT TO JUSTIFY THE EXPENDITURES ASSOCIATED WITH SECURITY.

Preventing loss, the major focus of the “insurance model” of information security, is so pervasive that the paradigm has spawned a related effort to justify the expenditures associated with security. This effort is prominently highlighted in the annual CSI/FBI Computer Crime and Security Survey, which describes the methods used by participating organizations to quantify the cost/benefit aspects of computer security. The Survey cites three metrics traditionally used to provide a rationale for allocating dollars to the information security function: Return On Investment (ROI), Net Present Value (NPV), and Internal Rate of Return (IRR). Each of these metrics involves the calculation of actual or potential economic loss due to lack of adequate security controls. The “insurance model” has become a means of justifying the increasingly costly resources required to support information security.

Recent federal and state legislation is also driven by the desire to prevent loss—of confidentiality, privacy, or money. The Gramm-Leach-Bliley Act of 1999 (GLBA), for example, seeks to ensure that sensitive customer information, and especially financial data, will

be secure both from identity theft and from old-fashioned monetary thievery. HIPAA, the Health Insurance Portability and Accountability Act, focuses upon the confidentiality of health-related information. A proliferation of recent state legislative initiatives, beginning with California’s SB1386, is intended to thwart the loss of personal data that could, in unauthorized hands, result in identity theft. Similar laws are currently under consideration on Capitol Hill.

Indeed, from the late 1960s until today, the “insurance model” remains a powerful engine, capable of generating new jobs in the field of information security, new hardware and software intended to secure data, new regulations, and even new kinds of crime.

II. Understanding Traditional Trust and “Virtual Trust” Concepts

The “insurance model” uses a set of concepts which we call Traditional Trust. We contrast this with a new conceptual framework, Virtual Trust. Both models rely on the core concept of trust, a necessary component of business.

Definitions of trust vary. For our purposes it is enough to mention some of the qualities and effects of trust. When we trust another person or entity, we have confidence that the outcome we expect to happen will happen because the person or entity recognizes, pursues, and completes the ends they are bound to by their word or deed. In short, if I tell you I will do something, it will be done.

Unfortunately, trusting another entity is more difficult than one might expect. Entities are often motivated by self-interest; they behave in ways that benefit them. Sometimes acting in one's self-interest dictates acting in ways that include breaking existing trust relationships, deceiving those who have trusted us. (The question of ethics is not dealt with in this paper. We are simply describing the real-world mechanics of trust, not the way they should work.)

From an historical perspective, how is such trust between entities established? Entities have used seals, signatures, contracts, deeds, treaties, notarized documents, handshakes and code words, among other methods, to create trust. These non-electronic (and/or physical) tokens of trust may be categorized as Traditional Guarantors of Trust, mainly for historical reasons. Often there exists a system -- such as a government or coalition of governments -- to settle disputes should they arise between the parties that made the agreement. In a less abstract context, the U.S. state and federal court systems are examples of independent entities enforcing Traditional Guarantors of Trust.

In contrast to Traditional Guarantors of Trust, we categorize electronic, non-physical tokens of trust as Virtual Guarantors of Trust. Examples of Virtual Guarantors of Trust include digital certificates, digital signatures, user names and passwords, public and private keys, the digital numeric sequence in two-factor authentication tokens, the electronic representation of a biological identifier, checksums and hashes.

Therefore, we make a distinction between Traditional Trust (which uses Traditional Guarantors of Trust) and Virtual Trust (which relies upon Virtual Guarantors of Trust).

Virtual trust is a reality. We encounter examples of virtual trust on a daily basis, but usually do not recognize or name it as such; it is merely taken for granted. The purpose of this paper is to make clearly visible the existence and value of virtual trust, an electronic means of establishing trust relationships that has largely remained invisible, even to information security professionals. We will provide examples of virtual trust and will demonstrate how virtual trust was used in the past and present and indicate possible ways in which it may be used in the future.

As noted earlier, the field of information security traditionally looks at guaranteeing that the internal controls that support virtual trust are not destroyed or weakened (by failed internal processes). This guaranteeing is the protection function described in the introduction of this paper. In direct contrast, we indicate how to create virtual trust. In effect, we describe how to create business using the mechanisms of information security. Additionally, we discuss how to create and maintain a competitive advantage using virtual trust.

III. Creating Business Through Virtual Trust: A Macro Perspective

Professionals in the field of information security are well aware that the "insurance model" is quite capable of creating business; each day, their email in-boxes are laden with offers from vendors or consulting services seeking to promote new and better means of preventing potential loss. It is not likely that the volume of these messages will decrease, because the current regulatory climate, coupled with an increasing reliance upon globally interconnected systems, has also generated new vulnerabilities and threats. Fear, uncertainty, and doubt is not merely a cynical sales tactic; FUD is often a legitimate response to real problems. Senior managers are anxious to protect critical information assets from potentially destructive or damaging forces; information security, like an insurance policy, is the price paid for ensuring that the destruction or damage is reduced to a minimum.

However, this is not the kind of business created by the "virtual trust" paradigm. Virtual trust is not intended merely to protect information, but to create or enable an asset for the

purpose of generating profit. The concept of trust, as mentioned previously, is not focused upon preventing loss. Rather, establishing trust is a precondition for conducting commerce. If trust does not exist between businesses and customers, then commercial transactions will not occur. Profit-making enterprises can thrive only when an assumption of trust is reasonably justified. Without a high degree of virtual trust, certain kinds of business would not be possible: automated teller machines, for example, would not function if their users are not securely identified and authenticated. Similarly, most e-commerce transactions could not occur if remote, unseen vendors cannot be trusted to identify themselves accurately. Virtual trust permits transactions between two parties who are at a distance and yet who can trust one another because they have been mutually authenticated. Because of this trust, business can occur; a flow of cash is made possible.

Robert Metcalfe, creator of Ethernet, is credited with developing a mathematical formula that attributes significance to the growth of communication networks. According to Metcalfe's Law, the "value" or "power" of a network increases in proportion to the square of the number of nodes on the network. The vir-

tual trust model of information security adds a new dimension to this law: the business/commercial "value" or "power" of a network increases in proportion to the square of the number of SECURE notes on the network. This suggests that, as the number of securely authenticated businesses and customers increases, the volume of commercial transactions and of cash flow also increases exponentially.

A physical metaphor that aptly illustrates the conceptual framework of virtual trust is a bridge. Consider, for example, the great Brooklyn Bridge that first opened for traffic on May 24, 1883. Today, this architectural masterpiece is perhaps best known as an element of a joke: "If you believe that, I have a bridge to sell you." But, when this structure was completed more than a century ago, the bridge itself was conducting the selling: it enabled new commerce, and had a profound impact upon the economies of both cities—New York (now Manhattan) and Brooklyn—connected by the bridge. Persons crossing the bridge never doubted its structural soundness. Its gargantuan stone towers, firmly planted on enormous caissons, and its strong 15-inch suspension cables became symbols of a bold, robust city and nation.

VIRTUAL TRUST PERMITS TRANSACTIONS BETWEEN TWO PARTIES WHO ARE AT A DISTANCE AND YET WHO CAN TRUST ONE ANOTHER BECAUSE THEY HAVE BEEN MUTUALLY AUTHENTICATED.

Virtual trust also functions as a bridge and establishes trust in much the same manner. Two entities must be connected for a commercial transaction to occur: the buyer and seller. But, as in the historical example of the Brooklyn Bridge, the electronic connection must also be secure; participating entities are loathe to lose money or critical information. The structural elements of a bridge, its stone superstructure and its steel cables, provide assurance that the edifice will withstand pressure and weather. Similarly, the components of virtual trust—digital certificates, signatures, and other forms of authentication—offer this assurance for buyers and sellers. And, like the Brooklyn Bridge, the "bridge" of virtual trust

creates new possibilities for commercial activity and economic growth.

IV. Creating and Maintaining a Competitive Advantage Using Virtual Trust: A Micro Perspective

Now that we understand how business is created through virtual trust via the bridge-building example, we will explore methods of creating and maintaining a competitive business advantage using the mechanisms of information security.

Virtual trust is created mainly by two mechanisms: authentication and non-repudiation. Authentication occurs when one can establish

who one is. Entering a user name/password and/or using a biometric device allows a system to identify you to establish your credentials and rights on that system.

Non-repudiation occurs when an individual or entity cannot deny that specific actions have been taken. A digital signature, for example, is intended to establish that it was you, and no one else, who sent a specific message.

For the purposes of this article, we will assume that the concept of authentication includes non-repudiation, although they are two logically distinct concepts. From this point forward we will mention Authentication only, unless we specifically mention otherwise.

Different methods of authentication yield varying degrees of trust because some mechanisms are stronger than others. Information security professionals, for example, have attained consensus that user names and passwords are not as strong as other forms of authentication, such as biometrics. Different authentication mechanisms can provide different kinds of information. For instance, digital certificates tell us about the individual presenting the certificate.

But exactly what can be done with authentication? The answer to this question, in turn, answers how to create and maintain a competitive advantage. We are seeking to perform two activities. First, we want to create or change a cash flow into the business. Second, we are seeking to decrease operating expenses though increased productivity.

Let us turn to the creation and change of cash flow. Digital certificates facilitate SSL encryption which, in combination with user names and passwords, enables business to conduct ecommerce via the Internet. This commerce represents a cash flow. For some businesses, like Amazon.com, the Internet is their only channel. For others, such as retail clothing stores, it is an additional means of communicating with customers. The security mechanisms of digital certificates and user names/passwords create a cash flow and generate opportunities for business.

A primary example of creating business through security is the SWIFT network and

application. SWIFT enables a secure messaging system for financial institutions. For the purposes of this paper, it is enough to know that entities that use SWIFT may send and receive payments as well as conduct other forms of business. The encryption used by SWIFT is primarily for purposes of non-repudiation. However, SWIFT will become a PKI over the course of the next two years (2008). In 2005, SWIFT message traffic generated revenues of 346,410,000 Euros, yielding a net profit of 7,790,000 Euros. Clearly, secure messaging generates revenue and is profitable. Not only may we say that SWIFT was enabled though secure messaging; it is more accurate to state that the business is secure messaging. Without the security, the message would be worthless and the model would fail.

The creation of a cash flow by incorporating security controls was stated by Microsoft when it suggested that people pay for sending email. People would be issued "Caller IDs" that would identify them as the legitimate sender of the email. The rationale was that by charging for securely authenticated email, the amount of SPAM would be decreased because the SPAMMers' cost would increase. The effect of such a policy would be that providers who charge for a secure email service would generate significant revenue.

Apple's iTunes employed Digital Rights Management (DRM) technologies to create a new product and, hence, a new revenue stream. Over 1 billion songs have been downloaded from iTunes. In the case of iTunes, DRM works by restricting the number of CPUs on which the .mp3 will play. The songs are also stored in a proprietary, encrypted format. These two factors, at minimum, erect a prohibitive barrier and thereby reduce the likelihood that an end user will trade songs. The various security mechanisms used by Apple's iTunes DRM created the Virtual Trust necessary to persuade the music industry that their rights will be protected digitally and be profitable.

ExxonMobil SpeedPass offers another example of cash flow made possible by using information security mechanisms. Before SpeedPass, a customer was presented with two options when paying for gas: cash or

credit. SpeedPass is an RFID token that the customer may link with a credit card. When the customer stops at a gas pump which accepts SpeedPass, they are immediately authenticated via RFID and the charge is billed. This linking of an arbitrary credit card with the SpeedPass RFID token allows the customer to alter the flow of cash at their discretion.

Citibank is using RFID to create PayPass. In a similar way that ExxonMobil uses Speed-

Pass, PayPass allows the consumer to swipe an RFID tag at certain locations to immediately debit their account. There are transactional security mechanisms—such as not being able to charge more than a certain dollar amount—built into the process to protect the consumer. This offers an example of the “insurance” and “virtual trust” concepts used in combination. Virtual trust is employed to create a new channel for cash flow and the insurance model ensures that security controls are placed upon this revenue stream.



Encryption is strongly used in governmental communication in the field of battle.

In the Northeastern United States, EZPass is an RFID system that allows customers to pass through highway tolls and be automatically billed per month. In addition to collecting dollars and cents, RFID / EZPass created a new revenue stream for the states that use it.

Security professionals are aware that RFID is not an overly secure protocol. Research has revealed that RFID transmissions may be captured and burned onto other RFID chips (cloned) or replayed. But we should not dismiss the concepts of virtual trust so quickly. Rightly or wrongly, consumers and corporations are relying on RFID tokens to conduct commercial activity. Research may indicate that the level of trust individuals are placing in this technology is unfounded and that the security protocols employed in RFID must be strengthened. Nonetheless, it is clear that people are trusting the RFID tokens to conduct transactions. RFID is an example of virtual trust created by the authentication security mechanisms used in the protocol.

Let us turn to the reduction of operating expenses. There is historical precedent for the use of virtual trust to reduce operating expenses; however, the reduction is usually attributed to technology rather than to security. VPN connectivity is a classic study. Businesses reduce their communication expenses by a secure connection via the Internet. Before VPN technology, business needed to

purchase lines for each remote location. However, with VPN, these lines were not necessary and publicly available channels requiring less overhead could be used.

Encryption is strongly used in governmental communication in the field of battle. Without such encryption, the enemy could intercept a message and disrupt military strategy. Current technology encrypts and decrypts traffic in real time, enabling secure communications between base and infield units. Traditionally, using the insurance model of information security, we would say that encryption protects the content of the messages. However, from the perspective of virtual trust, encryption enables communication to occur in the manner intended. Without encryption, the value of these messages—due, for example, to interception—would be worthless.

While it is beyond the scope of this article, it should be noted that Voice-over-IP (VoIP) appears to be the next technology that will demonstrate cost-efficiency in business environments. As often the case with technology that is implemented for internal use only, VoIP is frequently not deployed in a secure manner; security will be assumed because the internal network traffic is presumed to be trusted. However, as in the case of VPN technology, when VoIP is deployed across the Internet, encryption can be used to create virtual trust.

The most potent counter-argument against the above cases is that the technology, not the security, is the essential business driver. However, this is an inaccurate perception. We feel that, if people are using a specific technology, a certain degree of trust can be assumed when using it. This trust is linked with our belief that the technology is secure enough for whatever use we ascribe to it. Email is a prime example. Most of the time we do not consider encrypting our email; we send it plain text across the Internet. Plain text is secure enough for most email transmissions between entities since we are not really sharing anything of value. However, plain text email is not secure enough to send Social Security numbers, bank account numbers and the password to our online banking account. If I wanted to transmit these pieces of information in a secure manner, I would need to enable my application with the proper control mechanisms. Then the trust would be sufficient to allow me to conduct the necessary transactions. It is precisely this type of assumption that permits our internal network traffic to remain unencrypted, yet also allows us to feel secure about the corporate network.

V. Further Virtual Trust Concepts Defined

We have examined virtual trust from both a macro and micro perspective. At this point, it is worthwhile to further distinguish the virtual trust model from the "insurance" model as defined earlier in this paper.

The model of virtual trust incorporates a series of concepts intended to achieve enablement instead of protection. We should note here again that enablement does not displace protection as a valid model. Indeed, the protection mechanisms play an integral role within the enablement process. But the enablement model represents a significant change in mindset.

The concepts that comprise enablement, or virtual trust, are derived from both business and from information security. As has been mentioned previously, the security component is often mistakenly identified with the technology within which security controls are embedded. That is, we merely presume that a specific technology—such as a communication network—is trusted in two ways: first, we as-

sume, rightly or wrongly, that security is inherent in the technology (internal networks are a good example of this); second, the things we do with technology are assumed to be secure to the extent that we can use them and they are useful. In reality, however, security controls—such as encryption—are separable from the technology over which controls have been established. Virtual trust is not merely an accidental byproduct of implementing technology; it is the result of a convergence between specific technologies that perform security functions (e.g., encryption, digital certificates) and technologies that transmit or store data (e.g., VPN, VoIP).

The insurance paradigm has adopted a triad of concepts as its central mission-- securing the confidentiality, integrity and availability (CIA) of data. At the root of these concepts is the notion of protection: we must protect the data from being viewed by an individual who is not authorized; we must protect the data from being changed when it should not; and we must protect our resource so that when an authorized individual needs access to the data they may be able to do so.

Our core virtual trust concept is cash flow. Essentially, cash flow is a revenue stream for a business; it's how a business makes money through invoicing. We will not examine this concept in detail, as we have discussed it earlier.

The next concept, reducing operating expense, will also be treated lightly because it, too, has been considered previously. Basically, the reduction of operating expense—viewed from the perspective of information technology—usually consists of automating a process to reduce overhead.

Productivity is our next term. We may define productivity as "increasing cash flow while reducing operating expense." Traditionally, productivity is associated with technology; however, with the proper security, we can also increase productivity. BlackBerry devices are an example of enhancing productivity via secure transmission of information to end points.

One may object and say that we have insecure BlackBerry configurations and this is just as productive, if not greater, than secured

BlackBerries. This response, however, is not satisfactory because it does not account for the risks being taken in terms of trust and asset value. As mentioned earlier, the proper response to this objection is that we run at a trust level—rightly or wrongly—associated with the perceived risks associated with using a specific technology. We don't use the devices without security; rather, we use them at a security level consistent with our perception of the benefits of operating the device by comparison with our perception of a particular level of risk.

Transparency is another relevant concept. To end users, the enabling security must not be overly visible because security mechanisms are often viewed as obstacles to efficiency.

The end user must be able to take for granted that the system they are using is secure. This transparency is dependent upon trust. Transparency allows for seamless transactions between various parties using, for example, digital certificates as the authentication mechanism. This transparency allows for an increase in productivity from an internal business perspective. It also enables a greater generation of cash flow because the end user, with, for instance, an RFID token, does not need to think before scanning the token (impulse buying). Also, the transparency of security decreases the effort to purchase items (e.g., coins are not put in soda machine, credit cards receipts do not need to be signed, traffic passes smoothly through toll booths).

TO END USERS, THE ENABLING SECURITY MUST NOT BE OVERLY VISIBLE BECAUSE SECURITY MECHANISMS ARE OFTEN VIEWED AS OBSTACLES TO EFFICIENCY.

Data flow is our final term. Data flow exists at two levels: on a communication network and within an application. At the network level, it represents the passing of data from one end point to another. At the application level, data flow is the moving of information (such as records) from one entity's processing unit to another. Data flow can be considered the equivalent of cash flow. That is, data flow is the business transaction itself and, therefore, is how cash flows are created through technology. Traditionally, we thought that we must protect data as it flows from one end point to another. In our new virtual trust model, we say that we need to create a secure data flow from one end point to another. We assume that within this enablement we will protect the data. The purpose of rephrasing is not just rhetorical; the purpose is to show that our goals are, in fact, different.

Cash flow is the goal of commercial enterprise, and as we enable and build trust through authentication, we will not forget that the CIA triad helps to create a reliable environment within which commercial transactions can occur. But it is the secure enablement of virtual trust between parties that will allow for the creation of business through security, rather than using security only as a mecha-

nism to protect our assets. We will, in fact, create assets rather than just build walls around them.

VI. The Future of Virtual Trust

Every quarter, the Census Bureau publishes a statistical analysis of the growth, or decrease, of retail sales in the United States. Since at least the first quarter of 2005, the analysis has identified trends related to e-commerce, defined as the placement of orders and the negotiation of terms of sale over an online system. The most recent data indicate that electronic commerce has experienced significant growth during the past five quarters—from \$20 billion in 1Q2005 to more than \$25 billion in 1Q2006. The report also states that sales attributed to e-commerce transactions have increased an average of 5.66% during each of the measured quarters. By comparison, retail sales not conducted online experienced an average growth of only 1.84% in this time frame. These data suggest that electronic commerce will represent an increasing source of sales activity in the future, despite the warnings of some commentators that consumers' fears of identity theft may stifle online business activity.

“Virtual trust” has made possible a considerable proportion of this sales volume, at least regarding Internet transactions. (At this time, we cannot speculate concerning the extent to which virtual trust enables business occurring via extranets, Electronic Data Interchange networks, or email.) As the cash flow attributed to e-commerce continues to expand and to represent an ever-greater portion of the total value of retail sales, and as this growth is increasingly dependent upon the technical mechanisms that provide trust between businesses and their partners and customers, it is possible to make several predictions concerning the future of virtual trust. This future has implications for developments in electronic commerce and other online transactions, legislation, the evolution of information security, the metrics by which the benefits of information security are measured, and the security software industry.

Developments in electronic commerce

In the not-so-distant past, you used a personal computer to browse the Web and a cell phone to make telephone calls. Now, of course, these distinctions are passé: cells are equipped with browsers, and voice-over-IP enables telephone communications via the PC. This “device convergence”—the tendency to enable many functions from a single computing device—is now a reality, as attested by even a cursory visit to your local electronics store. And, also verifiable at your store, the bundled functions are increasingly supported by portable devices, such as PDAs, instant messaging equipment, and mobile phones.

These powerful devices are capable of supporting the software required for virtual trust. PDAs and cell phones, for example, can receive and store cookies, exchange digital certificates, and conduct transactions in an encrypted session. The portability of these devices, and the accompanying capability to engage in virtually trusted communication and commerce, have eliminated many of the traditional barriers to the conduct of business. Websites can communicate with customers on a global basis, and at any time. eBay, for example, will accept bids from cell phones on a 24x7 basis. Music can be downloaded to a cell phone whenever desired by the consumer. Retail banks can notify customers, via

mobile phone text messages, concerning authorized—and potentially unauthorized—transactions. Radio Frequency Identification, RFID, permits orders from vending machines—even at 3:00am. Time and geography are not necessarily obstacles within this worldwide marketplace. And the “virtual trust” paradigm of information security has, to a considerable extent, helped create the marketplace.

The ultimate objective of electronic commerce is to enable *any transaction, anywhere, and at anytime*. Obviously, there are limits to the feasibility of this objective: you may have your RFID device, but the vending machine is nowhere in sight. However, the bundling of functions within portable devices, the expansion of communications networks, and the increased reliance upon virtual trust are making e-commerce a ubiquitous reality. The major remaining obstacles are, it seems, cultural rather than technical. Language barriers, for example, may inhibit some web-based transactions. Mechanisms to bridge these cultural divides must be devised before customers are able to conduct any transaction, anywhere, and at any time.

Legislation

The Federal Deposit Insurance Corporation has required national banks within the United States to implement “strong authentication” for electronic banking transactions. Use of an ID and a password or PIN is no longer sufficient to access an account from the Internet. The FDIC mandate is, it seems, based upon several assumptions: (1) electronic banking will represent an increasing volume of retail activity; (2) a high degree of virtual trust between banks and their customers is required in order to transact business remotely; and (3) so-called “dual authentication”—consisting traditionally of a user ID and password—does not provide sufficient trust.

If these assumptions are valid, it is likely that the concept of “strong authentication” will be applied to ever-expanding array of online transactions, especially those involving sensitive information. It seems probable, for example, that the transmission of medical data may require a greater degree of virtual trust than is currently required. Similarly, email messages

that represent the conduct of financial or legal transactions will require strong authentication between sending and receiving parties. It is anticipated that legislation requiring the implementation of this authentication will occur, possibly at both the state and federal levels.

Revisiting the “Insurance Model”

The virtual trust paradigm of information security is essentially concerned with the issue of

authentication. A digital certificate, for example, is issued in order to assure a customer that he or she is, in fact, conducting business with a specific organization. However, virtual trust also usually involves the encryption of data and may require customer information derived from a cookie stored on the client’s PC or other device. Thus, the elements of encryption and identification, important components of the “insurance model” paradigm, are also incorporated into “virtual trust.”

THE PRACTICE OF INFORMATION SECURITY, DEVELOPED DURING THE PAST FOUR DECADES, HAS CREATED ADDITIONAL TOOLS AND CONCEPTS THAT MAY USEFULLY BE INCORPORATED INTO THE VIRTUAL TRUST MODEL.

The practice of information security, developed during the past four decades, has created additional tools and concepts that may usefully be incorporated into the virtual trust model. For example, the logging and monitoring of security-related events is a significant and necessary effort; auditors currently expect, as a matter of due diligence, that information security staff will capture and retain electronic or hardcopy evidence of events occurring within networks and systems. Logging and monitoring tools could perform similarly valuable services for the virtual trust function. For example, customers would be able to review reports describing recent online transactions or receive automated alerts concerning possible unauthorized activity. Current reporting mechanisms are usually implemented on an ad hoc basis only—businesses voluntarily determine that transactions will be confirmed via email. Similarly, customers are not automatically notified if a transmitted digital certificate has expired; a conscious effort to discover this information must be made. However, increasing reliance upon virtual trust seems to require that logging and monitoring tools must be more accessible and meaningful to customers. Buyers need not beware that their money has been spent and that records of the transaction are nonexistent.

“Layered security” is another concept that may beneficially be borrowed from the “insurance model” paradigm. The idea of “layering” refers to the process of implementing numerous security controls, some of which perform

seemingly redundant functions, in order to prevent unauthorized access. “Layered security” is frequently applied to networks, especially those that connect internal applications and systems to the untrusted Internet. Properly configured routers, firewalls, and intrusion detection and prevention systems are all integral elements of a layered approach to network security.

This form of security is expensive, because of the many software purchases involved, and it is time-consuming to test, implement, and monitor. However, it is deemed necessary in order to prevent unauthorized intrusions into a network that provides access to important data. “Layered security” does not, unfortunately, offer a seamless and centralized means to protect network resources; it is often a rather messy patchwork of unconnected software.

The virtual trust paradigm of information security will, in the not-distant future, confront the problem of whether or not to adopt a layered approach to its promise of delivering trusted transactions. It is tempting, for example, to envision handheld computing devices that embed security controls capable of automating the processes of connectivity, identification, authentication, authorization, and encryption. The user simply accesses an e-commerce site via the browser, the customer and the business are mutually authenticated, and transactions can occur with no further ado. Similarly, it is an appealing prospect to

imagine a universal connectivity standard that eliminates the need for storing multiple cookies, or receiving numerous digital certificates, or the remembering of innumerable passwords.

As mentioned previously, several retail banks have, in fact, recently implemented a form of virtual trust that authenticates user transactions in an efficient, seamless manner. The customer simply holds his or her bank card to a point-of-sale terminal and, when read by an RFID device, the customer's account is debited and a purchase made. The messiness of "layering" has been tidily cleared: there is no password to remember, no PIN to enter, no signature to write or authenticate. "Virtual trust" has become very simple indeed.

However, has the process of authentication become so simple that trust is eroded? Is the mere act of holding a card that is read electronically equivalent to virtual trust? Banks that adopt this system are clearly aware that the card may be "held" by an unauthorized person; perhaps a decision has been made that the benefits associated with ease-of-use outweigh the costs associated with fraud. At any rate, these banks have chosen to implement virtual trust without layered security.

Electronic commerce transacted via the Web often represents an attempt to incorporate virtual trust into a layered security environment. Digital certificates provide authentication, the web site usually requires identification and further authorization of the user, cookies supply certain details concerning customer preferences, and an SSL session offers an encrypted session. Several components, some of which are apparently redundant, comprise the totality of virtual trust. These components are analogous—if perhaps less messy—than the elements comprising traditional "layered" network security. However, in the Web-based context, layering is intended to ensure that a high level of trust is present before money is exchanged and goods purchased.

Does the layered approach adopted by most major websites truly guarantee a high level of trust? Or is the point-of-sale bank card, absent of multiple security controls, a good-enough method of authentication? If, as mentioned previously, legislation will increasingly

focus upon the need for strong authentication when funds are transferred electronically, it seems that "layered security" may emerge as a required element of the virtual trust paradigm. However, as exemplified by easy point-of-sale authentication based upon RFID technology, businesses may be reluctant to encumber customers with security controls. It seems that the "trust" element of "virtual trust" will be subject to increasing scrutiny and possible redefinition.

Beyond the Metrics of Loss Prevention

Economists have proposed three major methods for measuring the dollar value of benefits derived from implementing security controls. As summarized in the annual CSI/FBI Computer Crime and Security Survey, these methods include Return on Investment (ROI), Net Present Value (NPV), and Internal Rate of Return (IRR). These metrics essentially consist of comparing the costs associated with information security—especially salaries, software licensing fees, hardware purchases—to the estimated value of money saved by preventing loss. Proponents of these metrics acknowledge that assigning dollar amounts to potential losses is an highly subjective matter—how, for example, is preventing loss due to damage of corporate reputation assigned a monetary value? However, the metrics of loss prevention remain our only means of conducting cost/benefit analysis for the information security function.

The virtual trust paradigm introduces a new perspective from which to view metrics. Because virtual trust creates the possibility of commerce conducted from remote locations, and because this trust is established for the explicit purpose of generating cash flow, it may be possible to develop metrics that recognize information security as an enabler of business. Such measures would, assumedly, involve a comparison of relevant costs—such as expenses associated with the purchase of digital certificates and the support of additional identification and authentication systems—to the value of commerce generated by virtual trust. Although this value is subject to the same subjective interpretation as the metrics involved with loss prevention, quantitative research has been conducted concerning the likelihood of customers to patronize

online sites that provide a high degree of trust. For example, the online brokerage firm E*TRADE has estimated that the implementation of strong authentication has increased its trading volume by 30%.

As the volume of electronic commerce continues to expand, and as this expansion is increasingly dependent upon virtual trust, it seems that metrics intended to quantify the economic value of information security cannot ignore the income generated as a result of establishing trust.

“Virtual Trust” as an Over-the-Counter Product

Less than ten years ago, many businesses seeking to engage in electronic commerce

hired specialists to design, develop, implement, and maintain websites. The specialists frequently included HTML and Java programmers, graphic designers, and network professionals. These developers comprised the burgeoning “dot com” industry. However, as websites proliferated and became increasingly central to Internet communication and to e-commerce, the tools required to develop these sites emerged as over-the-counter products. Computer stores in local malls now invite would-be online merchants to purchase software that greatly simplifies the web development effort. With a bit of ingenuity and patience, anyone wishing to establish an electronic storefront can create their own e-business presence.

IF THE PRESENT EXPANSION OF ELECTRONIC COMMERCE CONTINUES AT ITS CURRENT RATE, THE ROLE OF VIRTUAL TRUST WILL ASSUME EVEN GREATER CENTRALITY.

Based on this democratization of web design, it seems reasonable to assume that “virtual trust” is likely to experience a similar destiny as an over-the-counter product intended to empower amateurs. Currently, the issuance of digital certificates largely remains the business of private “authorities.”

Similarly, the enabling of SSL sessions and the design and transmission of cookies have remained the responsibility of technical specialists. However, as virtual trust becomes an increasingly critical component of the e-commerce arena, the elements that comprise this trust will be more readily available to the buying public. Indeed, this trend has already begun: Windows XP includes a feature to permit digital certificates, or electronic signatures, to accompany email transmissions; eBay has now entered the business of selling digital certificates to its participating merchants. However, there are no technical or legal obstacles to the packaging of virtual trust as a readily purchased software product.

Aspiring e-commerce entrepreneurs will be provided the tools required to use digital certificates and signatures, to design and trans-

mit cookies, and to provide encrypted sessions.

The Future of “Virtual Trust”: Opportunities to Grasp and Lessons to Learn

If the present expansion of electronic commerce continues at its current rate, the role of virtual trust will assume even greater centrality. New products and services—such as virtual keys for automobiles and the notarizing of legal documents via electronic signatures—will, doubtless, emerge. However, this future is not necessarily limitless; serious obstacles remain that will, and should, serve to constrain the reliance upon virtual trust as a critical enabler of business. As mentioned previously, cultural barriers—including diverse languages and privacy regulations—cannot be ignored as potential brakes upon the momentum currently experienced by the virtual trust paradigm. Similarly, the paradigm must confront the issue of “layered security” and determine if weak authentication alone is sufficient to guarantee consumer trust, or if additional security controls must be borrowed from the tools currently associated with the “insurance model.”

VII. Enabling trust may become the dominant paradigm of information security

We believe that the new virtual trust model may become the dominant paradigm of information security. Businesses exist in order to generate revenue and, ultimately, profit. Protection of assets is simply a cost of doing business, and commercial enterprises wish to decrease expenses whenever possible. However, virtual trust focuses upon the enablement of business to generate more cash and, hence, increased profit; the objective of protecting assets, while an integral component, is not the primary goal of virtual trust. It seems that profit-oriented enterprises, while seeking to gain and maintain a competitive advantage, will increasingly adopt and measure the benefits of information security as presented from the perspective of virtual trust.

This perspective asserts that some security controls—such as digital certificates and signatures—actually create the possibility of doing business. A visit to any major e-commerce website quickly reveals that information security is no longer concerned merely with protecting assets and providing insurance against loss. FUD has become an obsolete rationale for the existence of information security.

There are additional reasons to anticipate the increasing dominance of the virtual trust model. First, as mentioned previously, the continued expansion of commerce conducted remotely—an expansion encouraged by the ubiquity of multi-functional mobile devices—is dependent upon the establishment of trust. Second, the vested interests of businesses and information security professionals will promote the significance of virtual trust. Commercial enterprises will seek new methods to provide secure, yet also efficient, relationships between themselves and their customers; information security professionals will strive to develop these methods. Third, as a result of the demand for processes and products that establish trust, new employment possibilities will be created since security will be pushed out into a new space. Finally, information security professionals will have a more positive and more important role within the organization; they will be viewed as creators of cash flow, revenue, and profit.

The enabling trust function will promote information security as a critical driver of business, not merely a system of controls that pleases auditors, satisfies regulators, and prevents loss. The virtual trust model of information security is not based upon selling fear; it envisions security as a creator and driver of business.

Kenneth F. Belva (CISSP, CEH, CISM) is currently employed at Credit Industriel et Commercial (New York) where he manages the Information Technology Risk Management Program. He reports directly to the Senior Vice President and Deputy General Manager. He is currently on the Board of Directors for the New York Metro Chapter of the Information Systems Security Association. He has presented on topics such as patch management as well as moderated a panel discussion on corporate governance. He taught as an Adjunct Professor in the Business Computer Systems Department at the State University of New York at Farmingdale. Mr. Belva is credited by Microsoft and IBM for discovering vulnerabilities in their software. He is the author of the chapter “Encryption in XML” in *Hackproofing XML* published by Syngress. He can be reached at www.ftusecurity.com and www.bloginfosec.com.

Sam H. DeKay, PhD. (CISM) has worked in field of information security for more than twenty years. Dr. DeKay is currently responsible for developing information security policies and standards at The Bank of New York. Prior to this, he served as Manager of Information Security at Empire Blue Cross/Blue Shield and at ABN Bank (now ABN AMRO), New York. As a CISM, he has written and edited study materials intended for security professionals planning to take the examination leading to CISM certification. Dr. DeKay has also been appointed a member of the Generally Accepted Information Security Principles (GAISP) Project, under sponsorship of the Information Systems Security Association. He holds PhD degrees from Columbia University and Fordham University.

Stop gambling your safety...

hakin9
Hard Core IT Security Magazine



www.en.hakin9.org