

## Featured this month:

### Defending Windows servers

**T**he Windows operating system has historically been attacked for being one of the most insecure and complicated systems to harden. Microsoft has invested a considerable amount of time and effort in improving this situation and has made measurable progress towards producing a secure 'by default' operating system. It has also gone a long way to helping system administrators build and deploy systems that are secure and resilient to attack.

However the Windows operating system is still a very popular target for attack and the process of defending the operating system is still not a trivial one; hackers are always looking to find a way to defeat your defence mechanisms. Implementing all of the best practices to the letter can often lead to an unusable or difficult to manage system in many real world environments.

*Go to Page 4...*

### How bad is it?

**I**T security architects have great faith in their security products. Some new research from Google gives architects, and end users, something to worry about.

Recently, Google researchers examined millions of web sites in an effort to determine the number of sites attempting drive-by attacks against users. Drive-by attacks can be particularly

dangerous because they are designed to evade detection and strike users when they least expect it; when they're surfing the web. Unsurprisingly, Google found a large number of malicious sites are working every day to compromise users' data. What's even worse, however, is that their research paints a bleak picture of the overall state of IT security.

*Go to Page 18...*

### McAfee snaps up Secure Computing

**M**cAfee continued along the acquisition path last month, announcing its plan to purchase email and web security company Secure Computing for approximately \$465m.

Secure Computing makes firewall, mail and web security products. McAfee's VP of worldwide solutions and competitive marketing Vimal Solanki said that the firm's product base will flesh out McAfee's offering in three of

the six segments of network security. McAfee is stronger in network intrusion prevention, network access control and data protection at the gateway, he said, arguing that although McAfee has web and email security products, they focus mostly on the small to medium business market. Secure Computing's market lies mostly among larger companies.

*Continued on Page 2...*

## Contents

### NEWS

- McAfee snaps up Secure Computing 1  
Security experts: US trailing on cybersecurity 2

### FEATURES

- Defending Windows servers** 4  
David Hartley of Activity Information Management explains the basic techniques of Windows operating system hardening. He explores the general principles, focusing on hardening Windows server environments, and demonstrates some common attack techniques that attempt to subvert operating system controls and defence mechanisms.

- Characterising intrusion detection sensors, part 2** 8

Researchers from Cranfield and York Universities evaluate the costs involved in deploying and monitoring sensors across networks as part of an intrusion detection system.

- Hiding a knife behind a smile – OBS hacking threats** 12

The attack surface for financial institutions is increasing day by day. Integrated online banking services (OBS) now pose a potential threat to the clients that depend on them. Aditya Sood discusses the effect of web-based insecurities and how the use of integrated services by third parties can exacerbate them.

- An overview of the best virtualisation solutions** 16

Modern organisations will find it difficult to escape the virtualisation craze. Dario Forte looks at some of the available solutions from a security perspective.

- How bad is it?** 18

Security products may make you feel secure, but how genuine is that sense of security? Bruce Potter investigates.

### REGULARS

- News in brief 3  
Events 20

#### Photocopying

**Editorial Office:**

Elsevier Ltd, The Boulevard, Langford Lane  
Kidlington, Oxford OX5 1GB, United Kingdom

**Programme Editor:** Steve Barrett

Tel: +44 (0)1865 843239

Fax: +44 (0)1865 853933

Email: s.barrett@elsevier.com

Web: www.networksecuritynewsletter.com

**Editor:** Danny Bradbury

Email: danny@itjournalist.com

**Senior Editor:** Sarah Gordon**International Editorial Advisory Board:**

Dario Forte, Edward Amoroso, AT&T Bell Laboratories;  
Fred Cohen, Fred Cohen & Associates; Jon David, The  
Fortress; Bill Hancock, Exodus Communications; Ken Lindup,  
Consultant at Cylink; Dennis Longley, Queensland University  
of Technology; Tim Myers, Novell; Tom Mulhall; Padget  
Pettersen, Martin Marietta; Eugene Schultz, Hightower;  
Eugene Spafford, Purdue University; Winn Schwartau, Inter.Pact

**Production Editor:**

Lin Lucas

**Subscription Information**

An annual subscription to Network Security includes 12  
printed issues and online access for up to 5 users.

**Prices:**

€992 for all European countries & Iran

US\$1110 for all countries except Europe and Japan

¥131 700 for Japan

(Prices valid until 31 December 2008)

To subscribe send payment to the address above.

Tel: +44 (0)1865 843687/Fax: +44 (0)1865 834971

Email: commsales@elsevier.com,

or via www.networksecuritynewsletter.com

Subscriptions run for 12 months, from the date payment is  
received. Periodicals postage is paid at Rahway, NJ 07065,  
USA. Postmaster send all USA address corrections to: Network  
Security, 365 Blair Road, Avenel, NJ 07001, USA

Permissions may be sought directly from Elsevier Global Rights  
Department, PO Box 800, Oxford OX5 1DX, UK; phone: +44 1865  
843830, fax: +44 1865 853333, email: permissions@elsevier.com. You  
may also contact Global Rights directly through Elsevier's home page  
(www.elsevier.com), selecting first 'Support & contact', then 'Copyright  
& permission'. In the USA, users may clear permissions and make  
payments through the Copyright Clearance Center, Inc., 222 Rosewood  
Drive, Danvers, MA 01923, USA; phone: +1 978 750 8400, fax: +1 978  
750 4744, and in the UK through the Copyright Licensing Agency Rapid  
Clearance Service (CLARCS), 90 Tottenham Court Road, London W1P  
0LP, UK; tel: +44 (0)20 7631 5555; fax: +44 (0)20 7631 5500. Other  
countries may have a local reprographic rights agency for payments.

**Derivative Works**

Subscribers may reproduce tables of contents or prepare lists of arti-  
cles including abstracts for internal circulation within their institutions.  
Permission of the Publisher is required for resale or distribution outside  
the institution. Permission of the Publisher is required for all other  
derivative works, including compilations and translations.

**Electronic Storage or Usage**

Permission of the Publisher is required to store or use electronically  
any material contained in this journal, including any article or part of  
an article. Except as outlined above, no part of this publication may  
be reproduced, stored in a retrieval system or transmitted in any form  
or by any means, electronic, mechanical, photocopying, recording or  
otherwise, without prior written permission of the Publisher. Address  
permissions requests to: Elsevier Science Global Rights Department, at  
the mail, fax and email addresses noted above.

**Notice**

No responsibility is assumed by the Publisher for any injury and/or dam-  
age to persons or property as a matter of products liability, negligence  
or otherwise, or from any use or operation of any methods, products,  
instructions or ideas contained in the material herein. Because of  
rapid advances in the medical sciences, in particular, independent  
verification of diagnoses and drug dosages should be made. Although  
all advertising material is expected to conform to ethical (medical)  
standards, inclusion in this publication does not constitute a guarantee  
or endorsement of the quality or value of such product or of the claims  
made of it by its manufacturer.

02158

Printed by Mayfield Press (Oxford) Limited

Solanki also pointed to the cloud computing component of Secure Computing's product offering. "They bring in this whole reputation-based service, looking not just at content, but at where it's coming from," he said. "So security in the cloud is again complementary between the two companies."

Cloud computing, in which content is analysed using internet-based servers, rather than at a customer's premises, has gained traction recently, with firms like Prevx and Panda offering their own cloud-based services. Benefits linked to the concept include the ability to devote more processing power to content analysis, simpler implementation for customers, and the collection of data from the entire customer base to help refine security decisions. McAfee recently launched its own cloud computing service called Artemis.

This is the latest in a long line of purchases for McAfee. Starting off as purely an anti-virus vendor, it expanded its operations through acquisitions, including Foundstone, which it bought for \$85m in 2004, and Preventsys, which aggregated security reporting data from different vendors' tools, in 2006. SiteAdvisor, which alerts surfers to dangerous sites, also merged with McAfee that year. Last October, it snapped up encryption firm SafeBoot for \$350m, and this July it agreed to buy data loss prevention firm Reconnex for \$46m.

Secure Computing made a loss of \$35.06m in its last financial year, on revenues of \$237.9m. The company had roughly 22 000 enterprise customers, which McAfee will add to its roughly 100m customers.

## Security experts: US trailing on cybersecurity

**The US is lagging behind in its cyberdefence policy and urgently needs improvement, according to Congressional testimony given this month. A hearing (<http://tinyurl.com/42zvzx>) held by the House Permanent Select Committee on Intelligence found inadequate links between public sector homeland security efforts**

**and private sector components of the critical national infrastructure.**

Paul Kurtz, former senior director for Critical National Infrastructure Protection for the White House Homeland Security Council, praised the Comprehensive National Cyber Initiative (CNCI), formed to help protect federal government computer systems, but said that its work was belated, and had been hindered by bureaucratic infighting at the Department of Homeland Security.

"The CNCI represents a beginning, but we have [fallen] woefully behind and are vulnerable. US government leadership, vision, and commitment are critical," he said.

Amit Yoran, the CEO of NetWitness Corporation, which provides network monitoring technology to the US intelligence community alongside private clients, strongly criticised the DHS. Yoran predicted the outcome of a report due in the next two months from the Center for Strategic and International Studies (CSIS) Commission on Cyber Security for the 44th Presidency. The Commission was convened to provide security recommendations for the next president.

"The Department of Homeland Security lacks the personnel, capability, authority and culture required to do the job entrusted to them by the President and Congress," said Yoran, also quoting Robert Stephan, DHS assistant secretary for infrastructure protection.

"Most of the time, I spend most of the bullets in my single 30-round magazine that I bring to work every day shooting into the backs of our own bureaucracy trying to clear a field of fire," Stephan reportedly said to Yoran. "So, I have one bullet left to either pump at al Qaeda or save it for me because the bureaucracy is about to overwhelm me."

Yoran also called for a more comprehensive security research programme, rather than relying on private-sector research, and suggested a national oversight policy to ensure that privacy was properly considered when creating federal security systems.

## In brief

### US tops attack source chart

According to Atlanta-based research firm SecureWorks, the US has been the greatest source of internet-based attacks against its clients so far this year. At 20.6m attempted attacks, it beat runner-up China (with 7.7m attempts) almost three-fold.

### Intercage caged

Atrivo, the ISP formerly known as Intercage, has disappeared from the web. The company, which had been linked to multiple cases of online scams and hosted malware, lost access when its upstream provider, Pacific Internet Exchange, pulled the plug.

### Intel ships security-conscious chip

Intel released version 3.0 of its vPro technology, which includes a facility for PCs to remotely ask for help if they find themselves operating unusually. It also features an access monitor to keep track of system activity, and embedded support for Microsoft's network access protection technology.

### GSX updates server monitor

GSX released version 9 of its server monitoring and reporting system, which lets administrators check server infrastructures for risks in real time. <http://tinyurl.com/3vhy2j>

### VMware patches buffer holes

VMware patched two vulnerabilities in its Openvman system management tool, used in the VMware Management Console and in its ESXi hypervisor product. The vulnerabilities were both buffer overflows, and could have led to system exploitation by attackers without correct login credentials, it said. <http://tinyurl.com/3lsp9r>

### Motorola puts finger verification on enterprise phones

Motorola announced a biometric attachment for its MC70 and MC75 enterprise digital assistants. The attachment supports the firm's Mobile Automated Fingerprint Identification System, alongside third party applications.

### Astaro ships email gateway protection

Astaro has entered the mail gateway market with a product focused on small to medium businesses. Its system includes two separate anti-virus scanners, and built-in email encryption.

### Kaspersky tackles SMB market

Kaspersky has launched a malware protection pack for small to medium-sized businesses. Small Office Security features licenses for five PCs and a file server, and an expansion pack is available for five additional users. The software

features anti-virus, intrusion prevention, anti-rootkit and anti-spam protection.

### IBM upgrades UTM

IBM upgraded its Proventia unified threat management system to include an SSL VPN.

### Fortinet moves into database defence

Another UTM vendor, Fortinet, has entered the database security market. Its new FortiDB-1000B V appliance scans databases and detects flaws in areas such as password settings, access privileges and configuration settings. The device recommends remediation measures for administrators.

### Intego claims Quicktime flaw

Security software firm Intego says that it has found a bug in Quicktime Player 7.5.5. The software handles long strings incorrectly in its `<? quicktime type= ?>` tag, says the firm, which can lead to a heap overflow in any programs using it (such as iTunes). Apple released the new version of Quicktime in September.

### Telco remotely manages phone security

Nordic telco TeliaSonera is launching a mobile device management service for SME customers. The service will use Nokia's Intellisync Device Management and F-Secure's Mobile Security software to remotely manage mobile phone security.

### NIST embraces CPE naming

The National Institute of Standards and Technology has upgraded its vulnerability database to support the Common Platform Enumeration (CPE). The CPE is Mitre's official naming scheme for vulnerabilities.

### Symantec offers secure remote access

Symantec has added another service to the Symantec Protection Network (SPN) that it unveiled in February. The SPN, a software platform designed to support a range of managed online services, is now offering a beta version of an access service that will let individuals use their applications remotely.

### Security vendors get virtual

McAfee launched Total Protection for Virtualisation, its tool designed to protect virtual machines. Symantec also launched version 12.5 of its Backup Exec tool, designed to support backups from virtual machines alongside 'real' ones.

### Malicious Javascript moves into third party apps

Finjan has spotted malware writers using third-party application support for Javascript as a vector for malicious software delivery.

Criminals are using obfuscated Javascript in PDF files to infect victims' machines, said the firm, in its Malicious Page of the Month report in September. It fears that Flash files could also be used to deliver obfuscated Javascript by writing malicious script to web pages.

### Los Alamos security criticised

The Government Accountability Office has criticised security in unclassified networks used within the Los Alamos National Laboratory. LANL, one of the labs used in the US nuclear weapons programme, suffered problems in identifying users, encrypting sensitive information and compliance auditing, said the report. <http://tinyurl.com/4zwuky>

### Phorm to be trialled again

BT is said to be planning a third trial of the controversial web tracking system Phorm. The system, which assigns unique IDs to web surfers and then tracks their movements at the ISP level, is designed to make it easier to serve relevant advertisements to users. This trial will be conducted on an opt-in basis, unlike previous trials, say reports.

### States get tough on data encryption

The state of Nevada requires businesses transmitting personal information electronically to begin encrypting it this month. In January next year, the Commonwealth of Massachusetts will require any portable data, such as that held on USB keys, to be encrypted as well.

### Data Breach Watch

Personal data on up to 50 000 current and former military personnel was lost when three USB-connected hard drives went missing from **RAF Innsworth** in Gloucester. Names, service numbers, addresses and dates of birth were among the information on the drives. Patrick Kaljoni, an employee of offshore gambling website **BetOnSports.com**, pleaded guilty to stealing the private information of at least 150 victims. Kaljoni worked with a large identity theft ring, passing the data to members after stealing it from the company's credit department, where he worked. <http://tinyurl.com/4k7ok6>. And the Korea Information Security Agency found that personal information on 164 536 individuals were exposed in **Google** searches this year. The Agency developed a software tool to probe the search engine, looking for residence registration numbers (South Korea's version of a national ID or social security number). Googling Excel files containing 'residence registration number' revealed almost 7 000 hits, most of them containing personal details. They were mostly accessible, even though those files had been removed from the internet.

# Defending Windows servers

Dave Hartley, senior security consultant, Activity Information Management

Windows is the most widely deployed operating system for enterprise environments. Nearly all organisations deploy at least some Windows server and desktop operating systems within their IT environments. It is true that Linux, OSX and other Unix derivatives are making substantial headway into the server market place; but most of the world requires interoperability with the Microsoft platform. This article therefore focuses primarily on the Windows platform, although many of the principles discussed apply to other operating systems.

There have been many books written and dedicated to the subject matter of hardening the Windows operating system. Security guides are also available from Microsoft, NSA and NIST that provide detailed instructions and recommendations to help strengthen the security of Windows deployments. This article aims to simply introduce a few of the best practice principals and perhaps introduce a few subversion techniques of interest.

The configuration and build of any operating system (OS) should undergo a hardening process to increase its resilience to attack. The purpose of hardening an OS is to prevent an attacker from achiev-

ing their goal of compromising the system, help protect the system resources and to maintain the integrity of the system. There are a number of techniques, tools and security best practices and principles at your disposal that can aid in preventing the compromise of the OS and in the event of a compromise, help mitigate the level of access that can be leveraged by an attacker.

A typical hardening process consists of:

- Physically protecting the system
- Removing unnecessary programs/set up files
- Patching the system

- Setting appropriate file system permissions
- Removing unnecessary network protocols and bindings
- Securing well known/default user accounts
- Hardening the TCP/IP stack against DoS attacks
- Defining an audit/event log policy
- Configuring appropriate user rights and ensuring that communications are authorised and only required channels are available.

Policies and procedures should also be defined that detail the hardening process applied and to manage the update process for keeping current with the latest vendor patches and future configuration changes. Following these basic principles will help decrease the risk exposure and increase the security posture of the host and serve to increase the performance and stability as a by-product.

Any hardened OS should then be independently assessed and benchmarked in order to provide assurances that the configuration and deployment of the system is in line with industry best practices and conforms to any regulatory requirements.

## Let's get physical

The first step in protecting the system is to physically secure it. If an attacker can gain physical access to the system and is able to boot from removable media such as a USB stick or CD-ROM, it is possible for them to gain access to the operating system by either cracking the passwords for the local user accounts or resetting the user credentials. Windows stores users' passwords within the Security Accounts Manager (SAM). The SAM does not contain the passwords themselves. The password is hashed, and then the hash is encrypted and stored within the SAM. The encryption key is stored on the server along with the hashes that the key encrypts, and the hashing algorithm is very well known.

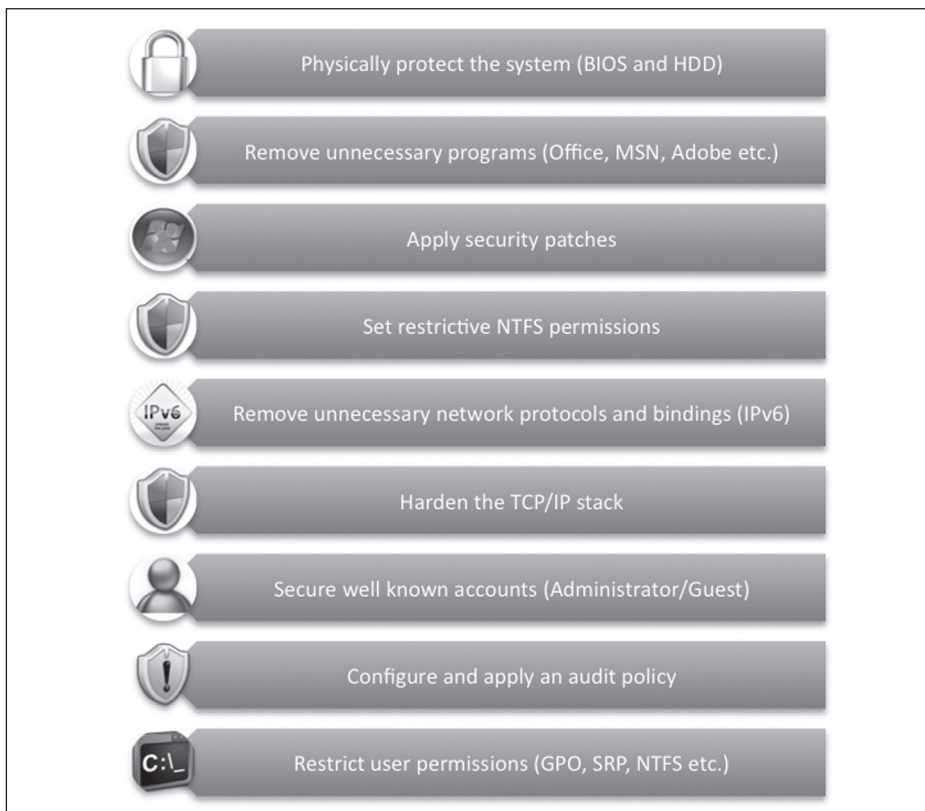


Figure 1: Steps to hardening a Windows server.

The Ophcrack<sup>1</sup> Project has released a Linux Live-CD that can be used to trivially retrieve and crack passwords. Ophcrack is a Windows password cracker based on a time-memory trade-off using pre-computed rainbow tables. The Live-CD can be used to boot from the CD-ROM (or USB stick) and load an OS into RAM. It then locates, retrieves and cracks the password hashes, without user interaction. Alternatively an attacker could simply change the password for a valid account on WinNT or Win2000 systems using a utility called chntpw, which is a Linux utility to reset the password of any user that has a valid local account on a system.<sup>2</sup> It works by modifying the password in the registry's SAM file.

It is possible to set a password to protect the BIOS and disable the ability to boot from other media sources; however many BIOS manufacturers have provided backdoor passwords that can be used to access the BIOS. Some of these are well published on the internet. Some BIOS protection mechanisms can also be bypassed through key combinations (such as holding the shift key), repeatedly pressing the mouse buttons or by overloading keyboard buffers by continuously pressing keys.

Many motherboards feature a set of jumpers or dipswitches that will clear the CMOS and wipe all of the custom settings including BIOS passwords. A small battery that is attached to the motherboard buffers the CMOS settings on most systems. On older systems removing the CMOS battery or chip for 10-15 minutes can clear the BIOS settings. A lot of newer systems store their BIOS passwords in a manner that does not require continuous power, so removing the CMOS battery may not work.

A password can be set on a HDD in an attempt to prevent an attacker from removing the HDD and mounting it from a custom operating system. The password hash is written to the firmware tracks on the HDD. If the HDD is removed from the system, in theory, it should not be possible to access the data. However, there are several tools available on the internet that allow for the HDD password to be reset and unlocked.

Use a product such as the free and open source TrueCrypt and encrypt the HDD<sup>3</sup>. The product creates a boot loader on the boot sector of the HDD. In order to boot

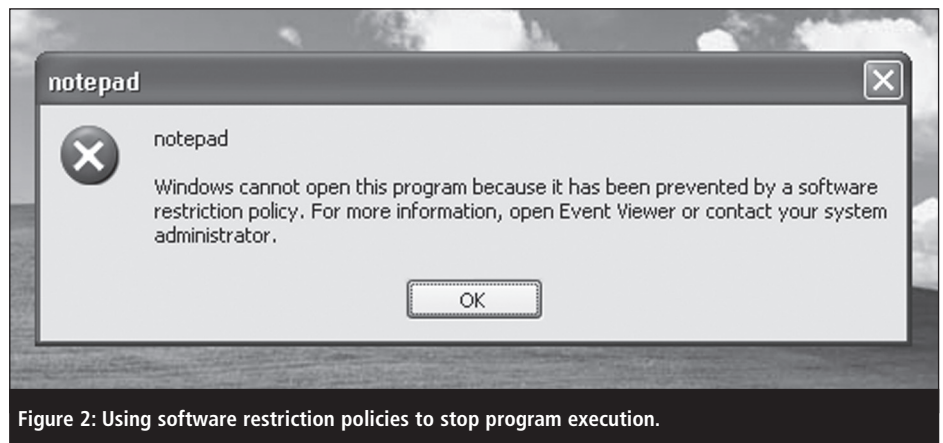


Figure 2: Using software restriction policies to stop program execution.

into the file system or access its contents the decryption key needs to be supplied. Alternatively you can use BitLocker Drive Encryption, a data protection feature available in Windows Vista Enterprise and Ultimate and in Windows Server 2008 that works in a much similar way.<sup>4</sup> If an attacker can gain physical access to the system, bypass the BIOS and HDD password they cannot easily read the data on the drive.

## Lose some weight, fat boy!

Reduce the attack surface. Systems should run with a minimised operating system for security and ease of manageability. Non-essential packages, applications and other files should be removed to reduce possible security exposures. Local services should be limited to only those required for business operation. Also separate critical components of your infrastructure so that each server only hosts one critical service i.e. web servers only running web service(s), mail servers running only mail service(s). This provides a degree of isolation and provides layered protection against attacks. It is strongly recommended that network and application server services and critical resources be distributed as a multi-tiered deployment, as part of a defence in depth strategy. For example do not install a database server on the same host as a web server or domain controller. Do not install MSN Messenger, media players and Adobe readers on servers. Remove or disable any unused protocols. By default, Windows Vista and Windows Server 2008 installs the standard TCP/IP version 4 and 6 protocols. If IPv6 is not in use on your networks, disable it.

The Security Configuration Wizard (SCW) and the Server Manager (Windows 2008) utilities can also be used to help secure systems. The SCW is an attack-surface reduction tool. SCW guides administrators in creating security policies based on the minimum functionality required for a server's role or roles. Server Manager, when used to install server roles on a Windows Server 2008 server, enforces the Microsoft-recommended security settings by default, the SCW can then be used to fine-tune the security configuration.

## Patch, update, patch, repeat

Attackers will look for the easiest way into a system and often security vulnerabilities that are present in resident services or applications are the first port of call. After installing the operating system, use Windows Update to install the latest security and system updates on the system. As a proactive initiative, security patch management is the primary line of defence for protecting a corporate computing infrastructure. A formal security patch management process should be implemented and enforced to ensure that all identified software updates are in place, thereby eliminating vulnerabilities from the environment and mitigating the risk of compromise. Develop internal processes for the continued management of the system; ensuring that critical patches are applied to a system no more than four weeks after issue from the vendor.

Use the Microsoft Baseline Security Analyzer (MBSA).<sup>5</sup> The tool analyses Windows systems for security vulnerabilities and misconfigurations before benchmarking them against Microsoft security recommendations, and offers specific remediation

guidance. Use the IIS Lockdown tool if deploying an IIS web server or Exchange server.<sup>6</sup> The tool functions by turning off unnecessary features. To provide in-depth defence or multiple layers of protection against attackers, UrlScan has been integrated into the tool.<sup>7</sup> All of the default security-related configuration settings in IIS versions 6.0 and 7.0 meet or exceed the security configuration settings made by the tool. Therefore, you do not need to run this tool on web servers running IIS version 6.0 or 7.0. The output of these programs should provide sufficient information to fine tune the configuration. Commercial vulnerability scanning and security management tools such as GFI Languard Network Security Scanner and Nessus can be used to perform authenticated configuration, patch and policy checks against systems.<sup>8,9</sup>

## If your name's not down, you're not coming in!

If an attacker has failed to exploit a vulnerable application or service, the attacker will move onto trying to guess, steal or brute force the credentials for exposed services that require authentication, such as SMB, NetBios, RDP, SSH etc.

One of the most common methods of enumerating valid users on a Windows system is to use a null session. Null sessions are a type of Windows Server Message Block (SMB) communication that provides

the foundation of network file and print sharing services. Null sessions are used for various unauthenticated networking processes including maintaining the browse list and down level trusts within a windows domain. An attacker can use null sessions to perform such activities as enumerating usernames, groups, shares, permissions, policies, services and more. Disabling support for null sessions is not a trivial task for enterprise environments. Steps outlining the process are detailed in Microsoft KB articles (KB143474 and KB246261); however please note that the required configurations could have a detrimental effect on legacy software applications. Making the changes reduces the functionality of the computer. A balance of usability/functionality and security needs to be found.

If null sessions are not available attackers will target accounts that they are reasonably confident are present, such as the default administrator account. This is why it is important to disable and remove extraneous and legacy user accounts. Ensure that the guest account, HelpAssistant and the administrator accounts are renamed and disabled. Create individual user accounts for users and assign administrative permissions through the use of groups. This allows user actions to be tracked and audited. User accounts that are installed as part of an installation of third party software should have their credentials changed from their defaults to match a strong password policy.

Default usernames and password combinations are easily available on the internet and are often used by attackers to compromise systems.

If an attacker obtains valid authentication credentials and compromises a user account, the game is far from over. You can take steps to limit their actions and restrict their movements. It is important to harden the file system and restrict access to sensitive locations such as %ProgramFiles%, %SystemDrive%, %SystemRoot% and %WinDir%. A user who can write to a %ProgramFiles% location that contains a common executable that is launched by other users of the system, could replace it with a piece of custom code that elevates the user's privileges. If it is possible to write to the %SystemRoot% and the BIOS has been protected to prevent booting from removable media, an attacker can still boot into another OS. GRUB for DOS is a universal boot loader and can be used to boot from a CD-ROM<sup>10</sup>. An attacker could install this tool into the %SystemRoot%, edit the boot.ini file to execute the boot loader and load the custom OS.

## Where's your tool?

Remove unnecessary system administration tools, because attackers can make use of these tools. Have administrators load a custom tool set from the network or removable media when local to the system. Do not rely on a Group Policy Object to prevent access to tools such as cmd.exe, regedit.exe, regedt32.exe, taskmgr.exe etc. If this GPO is set, the binary that is executed will check for the existence of a registry key. The execution continues if the value of this key authorises it do so. If an attacker has write access to the binary or can copy the binary to an area of the file system with write and execute permissions, it is possible to patch the binary to change the location of the registry key. The patched binary will look for a key that doesn't exist, and thus continue execution. Other methods include injecting a DLL that intercepts calls to the registry API and filtering the return values. Limited users can inject DLLs into their own processes. Both methods are published on the internet and tools exist to automate the process.

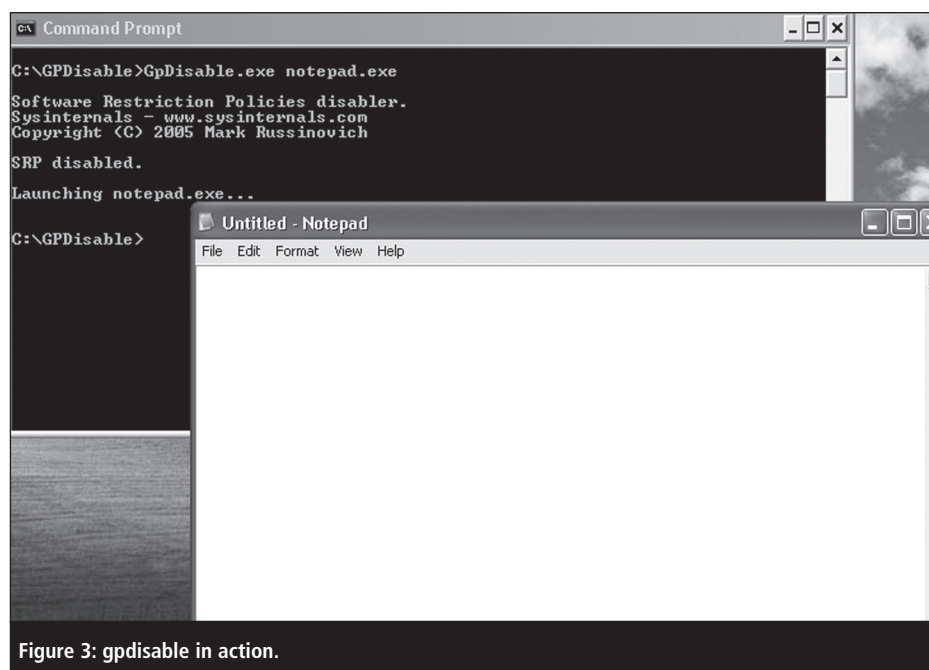


Figure 3: gpdisable in action.

## Computer says no!

Software restriction policies (SRPs) can be used to help protect against unauthorised software.<sup>11</sup> An SRP is a policy-driven mechanism for controlling what programs can and cannot execute. In an ideal world the policy should be configured with a list of trusted applications that can be executed (whitelist). This is a cumbersome laborious task and is also very difficult to manage, but will provide the best level of protection. An alternative and more manageable approach which is often chosen is to specify the locations on the system where programs can be executed from, and use NTFS access control lists to protect the file system and executables, ensuring that if a location on the file system requires users to have write access, then they are unable to execute from within that location and vice versa.

Software restriction policies do not apply to any program run by the SYSTEM account or Macros executed inside of Microsoft Office documents. An attacker who can execute macros within an office document could use the Windows API to execute code.<sup>12</sup> A motivated attacker can use the CreateProcess function from within a VBA (Office) macro with CREATE\_SUSPENDED flag in the process creation options.<sup>13,14</sup> The WriteProcessMemory function can then be used to inject an arbitrary payload after the SRP checks have been performed before calling ResumeThread on the handle.<sup>15,16</sup>

SRPs can also be bypassed if a user can execute arbitrary executables. When an SRP is enforced the executed binary queries a registry key value, which if present and non-zero indicates that execution is restricted. Mark Russinovich wrote an application called Gpdisable.<sup>17,18</sup> The program hooks the Windows API used for querying the registry and intercepts any calls and returns the error value STATUS\_OBJECT\_NAME\_NOT\_FOUND. This results in the binary being able to execute. Didier Stevens wrote a tool called bpmk to achieve the same goal.<sup>19</sup> The tool renames the value of the registry key in the process memory of the programs that impose SRPs. It is possible to restrict the scripting capabilities of Microsoft Office to prohibit the execution of macros embedded within documents,

but this setting does not prevent users from creating and executing their own macro. Ultimately if the Microsoft Office suite is installed and macro security restrictions are in place, SRPs can be bypassed.

The best solution is to use a combination of NTFS permissions to prevent the loading of arbitrary executables, spend the time to develop a comprehensive white listed SRP and extend the SRP to restrict not only executables but also .dll files. This has the unwanted side effect of making the system very slow to use, as every time an executable or .dll is executed, the SRP will need to be checked and this will still not prevent an attacker from using the Win32 APIs directly to manipulate process memory via Microsoft Office's VBA functions.

## Go deep

A hardened windows server is not going to withstand attack on its own. I have illustrated a number of defence mechanisms at your disposal and an attacker with sufficient skills and motivation under the right environmental conditions can probably circumvent all of them. That is why, in addition to server hardening I would always recommend a defence in depth strategy across the underlying network infrastructure.

Producing a secure operating system build is only one of many steps in all encompassing security strategy, it is a great start and will definitely not make the attacker's job an easy one, it will slow attackers down considerably and will hopefully force them into doing something that alerts you to their presence or frustrates them enough so that they give up and find an easier target. Introducing security products such as network firewalls, intrusion detection and prevention systems, host intrusion detection and prevention systems, anti-virus, web application proxy firewalls, web proxy content filters and so on into the IT environment will help prevent, alert and restrict an attacker's movements. The right security products in the right place, when configured in the right manner, can complement a hardened computing platform. It can also cost a lot of time and money, so a balance of cost and security is required, as is a common sense approach. It is important to have an independent risk assessment performed

and engage with a trusted security services provider to ensure that the solution adequately mitigates any identified risks, meets the business goals of the organisation and is cost effective in the protection it provides. In addition the solution should also be audited and assessed to ensure that at a technical level your defences cannot be trivially circumvented.

## About the author

*Dave has been working in the IT Industry since 1998. Dave is an experienced and competent security consultant and has worked on engagements at clients in a wide range of industry and government sectors internationally. Dave is responsible for the delivery of Activity's technical auditing program as well as the development and alignment of technical auditing methodologies, internal delivery support systems and the mentoring and education of the security services team members. Dave has performed many diverse and varied technical auditing services under the CESC CHECK scheme and independently of the scheme. Dave is also a CREST certified consultant.*

## References

1. Ophcrack <<http://ophcrack.sourceforge.net>>
2. Chntpw <<http://freshmeat.net/projects/chntpw>>
3. TrueCrypt product page, accessed Sept 2008. <[www.truecrypt.org](http://www.truecrypt.org)>
4. BitLocker product page, Microsoft, accessed September 2008. <[www.microsoft.com/windows/windows-vista/features/bitlocker.aspx](http://www.microsoft.com/windows/windows-vista/features/bitlocker.aspx)>
5. MBSA Microsoft Baseline Security Analyser product page, Microsoft, accessed Sept 2008. <<http://technet.microsoft.com/en-us/security/cc184924.aspx>>
6. IIS Lockdown tool product page, Microsoft, Oct 2002. <[www.microsoft.com/technet/Security/tools/locktool.msp](http://www.microsoft.com/technet/Security/tools/locktool.msp)>
7. UrlScan product page, Microsoft, accessed Sept 2008. <<http://technet.microsoft.com/en-us/security/cc242650.aspx>>
8. GFI NSS LANGuard Network Security Scanner product page, GFI, accessed Sept 2008. <[www.gfi.com/landing/lansseff](http://www.gfi.com/landing/lansseff)>

- asp?adv=69&loc=7>
9. Nessus product page, Tenable Network Security, accessed Sept 2008. <[www.nessus.org/nessus](http://www.nessus.org/nessus)>
  10. Grub4dos project page, Sourceforce, Oct 2007. <<https://sourceforge.net/projects/grub4dos>>
  11. Using Software Restriction Policies to Protect Against Unauthorized Software, Microsoft, May 2004. <<http://technet.microsoft.com/en-gb/library/bb457006.aspx>>
  12. Patrick Ogenstad, Breaking out of Jail with Microsoft Word, Ogenstad.net, Oct 2006. <<http://ogenstad.net/2006/10/04/breaking-out-of-jail-with-microsoft-word>>
  13. CreateProcess function description page, Microsoft, Sept 2008. <<http://msdn.microsoft.com/en-us/library/ms682425.aspx>>
  14. Process creation flags description page, Microsoft, Sept 2008. CREATE\_SUSPENDED <[http://msdn.microsoft.com/en-us/library/ms684863\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/ms684863(VS.85).aspx)>
  15. WriteProcessMemory function description page, Microsoft, Sept 2008. <<http://msdn.microsoft.com/en-us/library/ms681674.aspx>>
  16. ResumeThread function description page, Microsoft, Sept 2008. <[http://msdn.microsoft.com/en-us/library/ms685086\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/ms685086(VS.85).aspx)>
  17. Mark Russinovich, Circumventing Group Policy as a Limited User, Microsoft, December 2005. <<http://blogs.technet.com/markrussinovich/archive/2005/12/12/circumventing-group-policy-as-a-limited-user.aspx>>
  18. GPdisable tool download, hacker-soft.net, accessed Sept 2008. <[www.hacker-soft.net/download.php?id=9476&url=1](http://www.hacker-soft.net/download.php?id=9476&url=1)>
  19. bpmk Didier Stevens, "Introducing the basic process manipulation toolkit", personal blog, February 2008. <<http://blog.didierstevens.com/2008/02/28/introducing-the-basic-process-manipulation-tool-kit/>>

# Characterising intrusion detection sensors, part 2

Siraj A. Shaikh, research officer, Department of Informatics and Sensors, Cranfield University, UK

Howard Chivers, professor, Department of Informatics and Sensors, Cranfield University, UK

Philip Nobles, lecturer, Department of Informatics and Sensors, Cranfield University, UK

John A. Clark, professor, Department of Computer Science, University of York, UK

Hao Chen, research associate, Department of Computer Science, University of York, UK

**Last month, part one of this article emphasised the need for examining various features of intrusion detection sensors. It provided a broad definition of a sensor, as a device that is callable of signalling a suspicious event, and characterised sensors in terms of location and response, which are key factors in determining the effectiveness of sensor deployments. Part two characterises the various costs involved in deploying and monitoring, using such sensors, and shows how a diverse range of sensors are distinguished using such criteria.**

The value of an intrusion detection sensor inevitably lies in its ability to detect most suspicious events. Essentially two factors come into play here. First is the efficient detection of events of interest. Sensors placed at the right location with appropriate capabilities are what are needed. Excessive disruption at the site of monitoring and inadequate extra features do not add value for this purpose. Secondly, there is the added dimension

of cost: to procure, deploy, and monitor, using sensors. Such costs need to be thoroughly appreciated to make effective choices.

In practice there are no clear rules to obey when it comes to either of these two factors. The openness and connectedness of modern day networks, along with the increasingly innovative nature of intrusions, makes it very difficult to assess what sensors are to be placed

where and with what capabilities. In most cases it is a regular struggle and is unavoidable. Most sensors are freely available. Real challenges, however, lie in minimising costs associated with various operational overheads: the unique expertise required for deploying and configuring the sensor infrastructure, along with continuous updating of signature databases and rule sets, and unavoidable detection failures all incur a considerable cost.

Different sensors involved in monitoring of networked systems incur various different costs. Such costs are very important in considering the cost-effectiveness of a site-wide IDS deployment. This includes the cost of procurement, deployment across the monitored site and the cost of actual monitoring using the sensors.

## Availability

Sensors could be available as either open source or proprietary. Open source sensors, since the source code is freely available, are not only free for use but also

free to improve or modify as desired. Sensors available as proprietary, on the other hand, usually come with dedicated support that may make it cheaper to maintain and therefore offset some of the maintenance costs discussed below.

## Deployment costs

These costs include the installation and overhead costs associated with the actual deployment of a sensor. A variety of factors come into play here, including disruption caused to existing infrastructures, installation and configuration of sensors, and operational costs such as maintenance.

The cost of disruption will vary from site to site, depending on the configuration and design of the network, operating systems in use, and types of applications and services running.

***“Open source sensors, since the source code is freely available, are not only free for use but also free to improve or modify as desired”***

In general, network-based monitors such as Tcpdump or Arpwatch require no changes to network configuration and cause minimal disruption.<sup>1,2</sup> Inline sensors on the network such as firewalls and inline Snort would require some rearrangements to allow all traffic to go through it.<sup>3</sup> Host-based sensors, on the other hand, are most likely to cause greatest disruption. Enabling audit daemons, such as auditd, is likely to require rebooting and configuration changes on every host (FreeBSD versions earlier than 7.0 require auditing support to be explicitly compiled into the kernel).<sup>4,5</sup> Another example is the host-based IDS, such as OSSEC, which requires client software to be installed and manually configured on each host machine.<sup>6</sup> Such installation and configuration, particularly in busy large networks, is likely to interrupt access and operations, and is therefore costly.

An effective deployment of a sensor requires setting it up and configuring it

so that it allows efficient detection of suspicious events as it is designed to do so. Sensors need to be configured to handle a variety of different tasks including:

- Design and tuning of filters to detect particular events of interest
- Remote or local storage of any notifications or output generated
- Continuous updating of rules or signatures
- Provision of appropriate active response
- Manual steering from time to time to deal with any special incidents

Most of the above involves sensor-specific expertise, which is costly. Some sensors by their very nature are complicated to set up. An example of this is the host-based IDS OSSEC, which is based on a client-server architecture and also allows active response against possible intrusions as means of defence. Setting it up and running involves a range of configuration tasks to ensure:

- Client (software) agents are installed with appropriate rights on every client host
- Clients and server are bootstrapped to facilitate encrypted communication between them (which involves manual keying for every client agent)
- Desired files, processes and logs are monitored
- Notification and alerts are set up to be communicated through appropriate means
- Appropriate defence is carried out in response to intrusions

There are other sensors for which configuration is very manual but relatively easy to do. An example of this is the Cisco IOS port security mechanism, which is available on a range of Cisco network switches and allows individual switch ports to be limited to specified MAC addresses. Any device with a MAC address other than the specified ones attempting to access the port results in either dropping of all packets or a port shutdown, along with logging of such an attempt. Each switch port is statically assigned an entry for a valid MAC

address. Though tedious and difficult to administer, it is straightforward to set up and provides a very useful security measure that prevents ineligible devices from accessing the network and warns of such attempts very early on (at the point of first physical contact). Deployment costs for such a mechanism could be lower if the number of switch ports to configure is small and physical nodes remain mostly immobile.

Maintenance costs take into account routine tasks such as patching (of sensor software), upgrading due to functionality and compatibility, and reporting and replacement in case of any failures. Some sensors require more maintenance than others in this respect. Honeypots, for example, are designed to attract intrusive activity and malicious code that may be potentially zero-day in nature. Systems used as decoy for this purpose are therefore at a high risk of compromise and failure and, are likely to require frequent rebuilding and configuration. Virtual honeypots, such as *Honeyd* and *Nepenthes*, help to avoid this as they only emulate potential vulnerabilities and prevent intruders from exploiting services that a real host may offer.<sup>7,8</sup> Systems hosting such virtual solutions do nevertheless require more attention than others, given the high risk and unexpected nature of such activity.

## Monitoring costs

Monitoring costs are concerned with the actual use of a sensor to detect potentially suspicious events. Our characterisation of monitoring costs takes into account two factors, including:

- The level of manual engagement a sensor requires for the purposes of monitoring
- The performance of a sensor to accurately detect events of interest.

Monitoring involves various judgments to be made at different levels; ideally most such decisions would be automated. In practice, however, manual intervention and judgements are imperative and incur a high cost. A well-conceived deployment of sen-

sors should offset such cost as much as possible.

Some sensors are cheaper to monitor by their very nature. An example is the HP Virus Throttling technology, which is part of an enterprise-wide security solution and is installed on individual hosts as part of a software agent (client).<sup>9</sup> It is designed to throttle multiple unique outbound connection attempts which indicate potential worm propagation. By its nature, it is designed to prevent worm scanning while requiring minimal manual intervention. No connections are permanently blocked and in case of a rare false positive, genuine connections are easily reinitiated.

Another consideration is the suitability of sensors for various type of detection. So, for example, to detect suspicious patterns in network traffic it is easier to use Wireshark that provides a graphical user interface and traffic statistics, than Tcpdump; both are otherwise very similar in the visibility they provide and the cost to deploy.<sup>10,11</sup> A better alternative for such monitoring however is Snort, which

provides a range of options, tools and techniques specifically for this purpose, though it is much more costly to deploy, given the level of configuration involved.

***“False negatives that may result in missed activity and even entire attacks are more worrisome; the costs of which are incalculable”***

The severity of a network security policy is also a factor here. If, for example, the Cisco IOS port security mechanism is configured to shut a port down every time it detects an ineligible MAC address, then it will have to be restored manually. If however it is configured to just drop traffic and log the attempt in every such instance, then it would save restoring the port manually each time. In this sense, the cost of monitoring using a sensor is affected by the type of events monitored and also what response is activated.

Monitoring for potentially suspicious events by its nature gives way to potential misdetection. Although unavoidable, it brings another dimension to the cost assessment of sensors. Poor design of sensors is likely to increase the number of false positives (and negatives). Consistent and accurate detection, on the other hand, is certain to add value to monitoring.

False alarms incur a high cost as some of the valuable time and effort is spent dealing with what is essentially a non-event and as a result is diverted away from other potential suspicious events. Besides this, the disruption caused to eligible activity (such as dropped connections, blocked addresses) may also cost, particularly if the sensors are configured to provide response. While good sensor design is important in reducing false positives, resulting costs are also somewhat relative to deployment costs. For sensors such as IPFW, for example, a careful fine-tuning of its rule set, often an on-going effort, is bound to reduce the number of resulting false positives

Sensor	Location	Layers	Defense available	Cost		
				Availability	Deployment	Monitoring
Tcpdump*	I, B, S	P, N, T, A	No	Open source	Low	High
Cisco IOS Port Security	I, S	P	Yes	Proprietary	Medium	Medium
HP Virus Throttling	H	N, T, A	Yes	Proprietary	High	Low
Tripwire	H	A	No	Open source	Medium	Medium
Auditd	H, K	N, T, A	No	Open source	Medium	High
Snort*	I, B, S	P, N, T, A	Yes	Open source	High	Low
Honeyd*	B, S	N, T	No	Open source	High	Medium
Nepenthes*	B, S	A	No	Open source	High	Medium
OSSEC HIPS	H	N, T, A	Yes	Open source	High	Low
IPFirewall (IPFW)	I, B, S	N, T	Yes	Open source	Medium	Medium
Arpwatch	S	P	No	Open source	Low	Medium
Wireshark (Ethereal)*	I, B, S	P, N, T, A	No	Open source	Low	Medium

Table 1: A list of intrusion detection sensors (in no particular order) with the location they are deployed at, the layers at which they operate, and the availability of defence and various costs associated with them. Location is abbreviated to (I)nline, (B)ackbone, (S)egment, (H)ost, (K)ernel and layers are abbreviated to (P)hysical, (N)etwork, (T)ransport, (A)pplication.

\* Location for these sensors refers to one or more possible places where they could be placed.

(and negatives).<sup>11</sup> Comparatively, Snort is characterised to cost more to deploy, given the vast number of configuration options that it provides. Although, relatively, it does cost less to monitor using Snort, given that it is designed and dedicated to detecting and (in an inline setup) preventing intrusions.

False negatives that may result in missed activity and even entire attacks are more worrisome; the costs are incalculable. Some of the missed activity could, for example, devoid knowledge of internal compromise, which could later be used as a stepping stone for further severe compromise.

Understanding various cost factors associated with intrusion detection sensors is important to assessing how cost-effective a particular deployment is on the whole. However, there are costs other than those of sensors that need to be taken into account here. The cost of damage caused by intrusions is one which we have not examined; and it comes first and foremost since a low cost of damage may not justify deploying an IDS in the first place. Such costs simply do not go away once an IDS is deployed, however, considering factors such as false negatives which are unavoidable.

## Some common sensors

A broad definition of intrusion detection sensors allows us to compare and contrast a variety of sensors commonly deployed in table 1. We also make note of some of their features that have been discussed. A quick review of some of the common sensors deployed for the purposes of intrusion detection tells us that sensors range widely, from kernel-based daemons, such as auditd, to dedicated products such as Snort.<sup>5,3</sup> They operate at a variety of layers and in a variety of locations with some of them providing defensive capabilities. For the purposes of relative comparison, we assign high, medium or low to deployment and monitoring costs considering the factors highlighted in the discussion above.

Honeypots, such as Honeyd and Nepenthes, could be configured to provide a range of responses.<sup>7,8</sup> They are typically located in some dark part of the network which inhibits their visibility of

traffic. It is possible to deploy them in a grey fashion such that their presence is lightly scattered throughout the network, and in more than one or even all physical segments. Such a deployment could prove to be much more effective.<sup>12</sup>

File integrity checkers, such as Tripwire, are placed inside of hosts. They are limited to changes in file systems only and designed to be passive.<sup>13</sup>

Traffic monitors like firewalls, such as IPFW, are placed inline and work to block connections, close ports or drop packets if need be.<sup>11</sup> Their location falls somewhere between the backbone and the physical segment. Though they could be outside of a physical segment, their inline placement does allow them to be effective in monitoring all traffic to and from a possible source of a suspicious event.

Sensors operating at the physical layer are placed in individual segments to be effective. Address-pairing monitors that track IP-MAC address pairs, such as Arpwatch, are designed to log Address Resolution Protocol (ARP) entries and detect any attempts at spoofing ARP entries.<sup>2</sup> Another example of a sensor at this layer is the switch port security mechanism, such as the one provided by Cisco, designed to prevent devices from accessing the network, other than those configured to do so, by detecting ineligible MAC addresses. Such a mechanism is designed to be active (to drop packets or shut down ports if necessary) and inline by location.

## Conclusion

Devising efficient deployment of intrusion detection sensors that maximise the visibility of suspicious events at a minimum cost offers a real challenge. With a possibly variable location and typically a range of configuration and analysis options, along with the possibilities to respond to potential intrusions, such sensors offer a range of intricate issues that need to be considered if one is to achieve an optimal deployment. This article has tried to facilitate a clearer understanding of diverse characteristics of sensors, which should serve to inform and align the type and position of sensors required for IDS deployments.

## Acknowledgment

This effort is a result of a collaborative project between Cranfield University and the University of York funded by the EPSRC (EP/E028268/1) to study system level approaches to intrusion detection.

## References

1. Tcpdump. 3 August 2008 <[www.tcpdump.org/](http://www.tcpdump.org/)>
2. Arpwatch. 3 August 2008 <[www.nrg.ee.lbl.gov/](http://www.nrg.ee.lbl.gov/)>
3. Beale, Jay. Snort 2.1 Intrusion Detection. Syngress, 2004.
4. Chapter 17, Security Event Auditing of the FreeBSD Handbook. FreeBSD. 3 August 2008 <[www.freebsd.org/doc/en/books/handbook/audit.html](http://www.freebsd.org/doc/en/books/handbook/audit.html)>
5. Section 17.3 Installing audit support. Section 17.3. FreeBSD, 3 August 2008 <[www.freebsd.org/doc/en/books/handbook/audit-install.html](http://www.freebsd.org/doc/en/books/handbook/audit-install.html)>
6. OSSEC. 3 August 2008 <[www.ossec.net/](http://www.ossec.net/)>
7. Neils Provos and Thorsten Holz. "Virtual honeypots: From botnet tracking to intrusion detection". Addison-Wesley, 2007.
8. Nepenthes. 3 August 2008 <<http://nepenthes.mwcollect.org/>>
9. Williamson, Matthew W. "Throttling viruses: Restricting propagation to defeat malicious mobile code". 18th Annual Computer Security Applications Conference (ACSAC 2002), Las Vegas, NV, USA. Washington: IEEE Computer Society, 2002.
10. Wireshark. 3 August 2008 <[www.wireshark.org/](http://www.wireshark.org/)>
11. IPFW. Section 28.6. Chapter 28, Firewalls of the FreeBSD Handbook. 3 August 2008 <[www.freebsd.org/doc/en/books/handbook/firewalls-ipfw.html](http://www.freebsd.org/doc/en/books/handbook/firewalls-ipfw.html)>
12. Warren Harrop and Grenville Armitage. "Defining and evaluating greynets (sparse darknets)". Proceedings of the IEEE Conference on Local Computer Networks 30th Anniversary (LCN'05). 2005. Washington: IEEE Computer Society.
13. Tripwire. 3 August 2008 <[www.tripwire.com/](http://www.tripwire.com/)>

# Hiding a knife behind a smile – OBS hacking threats

Aditya K. Sood, independent security researcher and founder, SecNiche Security



The implementation of online business services (OBS) involves a number of parties who are dependant on one another and whose transactions are processed by a centralised entity. An OBS usually works as a business ecosystem containing various elements that each play a critical role and affect the business in one way or another.

## Persisting web OBS ecosystems

OBS companies design environments in which clients obtain services from a provider. This process generates mutual business transactions and, as a result, business niches are created.

Within each niche, service providers should bear the sole responsibility for security, functioning, and reliability for the different entities of the environment. They do not always do this, however. Some banks, for example, have outsourced online banking services to other service providers, which means that transactions occur on the web sites of these service providers. In addition, client information is updated in the client repository belonging to the outsourced service providers. Because financial transactions are involved, the dependency of the relationship is critical; it defines the customised business model involving both parties.

Web services are often developed using a service oriented architecture (SOA). In SOAs the interface layer services can be accessed without knowing the background processes. The interface defines the working functionality needed by the client to manage transactions. Most services use WSDL and SOAP as communication protocols. XML is used for transactional data along different platforms, in order to ensure there is absolutely no problem with interoperability. The overall ecosystem consists of organisational and customer entities that are made functional by an interface. The OBS model is based on the layout shown in figure 1.

Figure 1 shows the operating detail of an OBS. The aim of this discussion is to understand the insecurity that exists in centralised web sites, creating problems for clients.

## The real threat – the OBS ripple effect

We have looked at how OBS peripherals work, but we are more interested in the security aspect of OBS that may result in attacks. Web-based threats consist of XSS, CSRF, phishing, and others.

Integrated services are based on distributed technology with a centralised environment. Web clients have databases on centralised servers, which is a generic fundamental concept of distributed web functioning. What we want to focus on, though, is the prob-

lem of the OBS ripple effect, which is entirely based on the web business ecosystem and insecurities present in the web sites of the centralised online service providers. The OBS ripple effect was revealed during real testing conducted against a vulnerable service provider portal. Clarification of the testing outcomes will help us to better understand this effect. Let's have a look at the model first (figure 2).

Given service provider S with a number of clients, C(i):

$$C(i) = \{C1, C2, C3, \dots, Cn\}$$

and for every client there is a set of web sites, W(i):

$$W(i) = \{W1, W2, W3, \dots, Wn\}$$

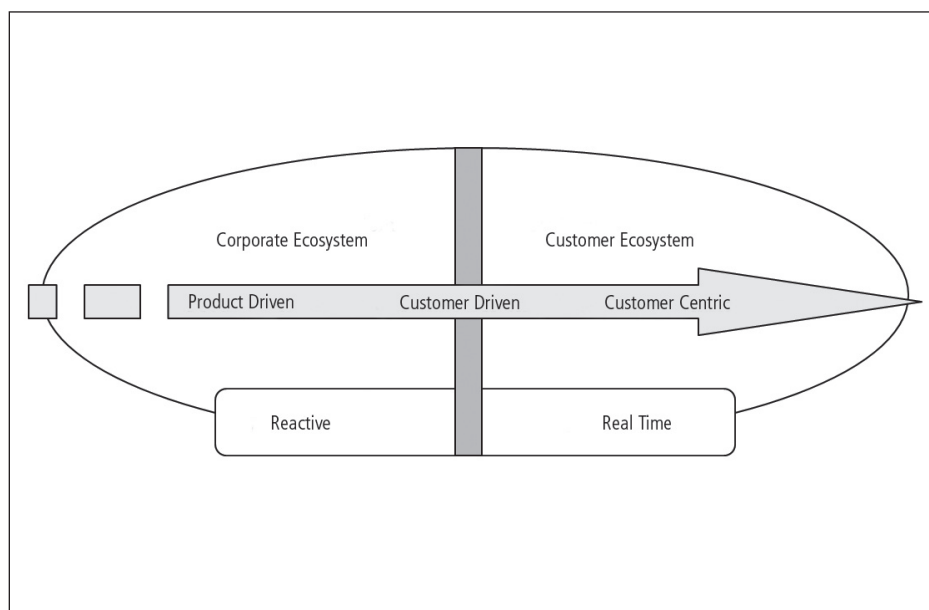


Figure 1: Business ecosystem overview.

where transactions are carried out per click,  $T(i)$ :

$$T(i) = \{T1, T2, T3, \dots, Tn\}$$

the client  $C(i)$  transaction is carried out on server  $S$  through  $W(i)$  by transaction per click  $T(i)$ . The user generally clicks the integrated service object on the client ( $C1$ ) web site ( $W1$ ). After clicking the online service object, the browser will open the window for the centralised server transaction page, i.e.  $S$ . The client ( $C1$ ) has a proper account with credentials to authenticate him or herself. The required URL shows the client code which is recognised by the server, and acts as a referrer check.

The description above is of a conventional layout, where the dependency factor is clear. The threat cloud defines the environment in which web attacks may flourish. The web site of the service provider possesses flaws and vulnerabilities that may lead to attacks and cause it to be a threat to its online clients. A flaw in a service provider's web site affects business in two ways:

- First, the intrinsic business of the service provider is affected. In online business systems the transaction process is usually based on 'click-to-continue' and 'pay-per-click' systems. Service providers charge for every click that is redirected to the web site, and for the transactions performed by users or clients.
- Second, the ripple effect may occur, defined as the problem that occurs when an effect is diversified from the source to the dependant elements. For instance, the web site of a service provider has a flaw, which seriously alters the working functionality for users and thereby affects a number of its clients, one of which is a bank. Let us say a user is redirected from the bank's web site to this integrated web service's web site where the transaction process occurs and, due to a vulnerability in the web site, an attacker is able to force the user to move to a malicious web site.

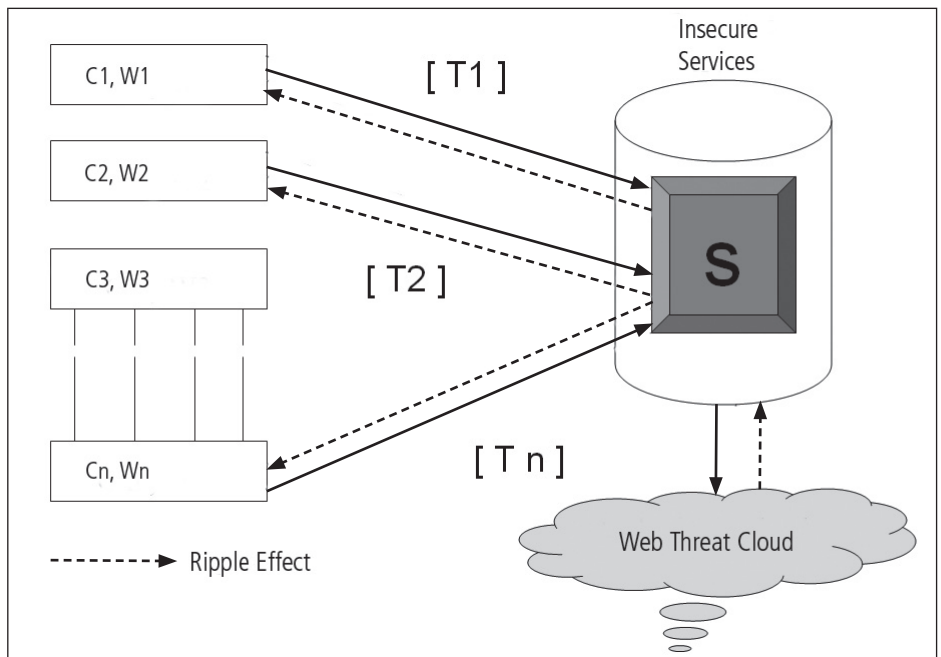


Figure 2: The ripple effect.

In this case the client's (bank's) web site has no vulnerability, but the service provider's web site has, and it is causing problems for the client. In other words, the effect is being transferred from the parent to the child. This state of affairs impacts the business behavior of both parties and the vulnerability in the service provider persists, causing a threat. Therefore, both factors are destructive.

### Human error framework – OBS perspective

The overall human error framework needs to be analysed. The human error framework analyses the impact of user action on a business. Developers are always on the bleeding edge, but people in the framework are also responsible to some extent.

***"An optimum level of user knowledge lowers the threat level to some extent"***

The effective use of technology requires an element of applied security from both parties in a transaction. Are users aware of potential threats? An optimum level of user knowledge lowers the threat level to some extent. For example: How many users take note of the URLs they visit?

According to a research study conducted at Carleton University, 45% of users do not even look at the URL. Approximately 35% of users do not understand the significance of https in a URL, and this plays a large part in creating an attack surface. Services used by clients need to be understood from a security perspective in order for those services to function effectively.

### Applied security – is it worthwhile?

Banking web sites have incorporated a number of security features for authorisation and authentication. However, there are number of other attacks of a different class defined in the Web Application Security Framework (WASF). Security for banking web sites and integrated web sites are affected by the following factors:

- Web site vulnerabilities
- Traffic manipulation
- Identity-specific parameters

The integrated security factors of password management, account monitoring for e-service and threat factors are emphasised with the aim of protecting user information disclosure.

When it comes to banking, security is applied to protect passwords, personal identification numbers, customer IDs, etc. This applied security translates well into user functioning and there is constant talk of phishing scams, Trojan horses, worms, etc., that devastate networks to steal information and affect machines. It should be clarified, however, that all these attacks occur as a result of web site vulnerabilities. The vulnerabilities in banking and service providers' web sites give rise to attack surfaces. These are battlegrounds which allow attackers to exploit web site vulnerabilities such as:

- Parameters in the host web site
- Third-party web site parameters – to launch attacks on the target
- Browser vulnerabilities
- Protocol discrepancies and communication models
- Flaws in the web server
- Insecure coding.

All these vulnerable elements affect financial institutions.

## Case study of a real-time OBS hacking threat

This case study will illustrate the concept of OBS integrated-service flaws. The study is the result of real penetration testing on a Canadian bank. Although details of the study cannot be revealed, the scenario may be used for education purposes. The scenario is described below.

The bank in question is one of the finest in Canada and provides normal online services, which it instituted so that users could manage their accounts online. It offers all types of transactions to users in a well defined manner and the processes are implemented by integrated services. These services are not hosted on the same domain as the bank's web site. Users are therefore redirected to a service provider's web site when the service tag is clicked, which means users are completing their transactions on the service provider's web site in the context of the bank, having supplied proper credentials to prove their authenticity. In other words, the service provider is

hosting a number of financial clients and credit union parties directly on its server. The credit union clients number 300, which must result in a large number of money transactions occurring on the centralised web site.

The aim of the study was to track a bug and analyse the ripple effect it caused. The most basic part of the testing was analysis of the threat. In this type of scenario, two conditions are normally possible:

- Vulnerabilities exist on the main banking web site
- Vulnerabilities exist on the service provider's web site.

Both conditions may result in drastic effects. The ripple effect will emerge if web vulnerabilities are found on a service provider's web site; and the effect may be in the form of a SQL injection, persistent or reflected XSS, session fixation, or other.

The testing was performed in a hierarchical manner. When the bank's web site was tested against attack parameters, it

	Win 98	Win 2000	Win XP	Vista	Mac 10.3.x	Mac 10.4
IE 7.0*	n/a	n/a	Level 1	Level 1	n/a	n/a
IE 6.0*	n/a	Level 2	Level 1	n/a	n/a	n/a
IE 5.5*	n/a	Level 3	Level 3	n/a	n/a	n/a
IE 5.0*	n/a	Level 3	Level 3	n/a	n/a	n/a
Netscape 8.1	n/a	Level 2	Level 2	Level 2	n/a	n/a
Netscape 4.7 & up	Level 3	Level 3	Level 3	n/a	Level 3	Level 3
Firefox 2.0 & 1.5	Level 3	Level 2	Level 1	Level 1	Level 2	Level 1
Firefox 1.0	Level 3	Level 2	Level 2	Level 3	Level 2	Level 2
Mozilla Derivatives	Level 3	Level 3	Level 3	Level 3	Level 3	Level 3
Opera 9	Level 3	Level 2	Level 2	Level 2	Level 2	Level 2
Opera pre 9	Level 3	Level 3	Level 3	Level 3	Level 3	Level 3
Safari 2.0	n/a	n/a	n/a	n/a	n/a	Level 1
Safari 1.3 & up	n/a	n/a	n/a	n/a	Level 2	n/a

Table 1: Browser support statistics.

was properly structured and security was implemented in a precise way. When the browser specification was checked it was found that the browser usage was structured as illustrated in table 1:

Level 1: quality assurance test is performed

Level 2: quality assurance test is not conducted but browser is supported

Level 3: no specific test is conducted  
n/a: not applicable

\* indicates biometric support for login.

The overall scenario was quite clear from a testing perspective. The following tests were performed on the main web site of the bank:

- Going through client-side code to understand data flow
- Input data validation checks
- Finding entry and exit points in the application
- Testing context precision of an application
- Stress testing for application robustness
- Automated vulnerability scanning including XSS checks.

All the tests conducted did not yield any useful results. This is because the web site was properly secured and well constructed. The next step was to test the online integrated service web site. The web site had a login page and two peripheral pages for help standards and contacts. The following tests were conducted by looking at the available entities:

- The URL that manipulates the traffic from bank web site to that of service provider was disassembled

- Referrer context was analysed
- Input points were checked and rogue tests performed
- Blind SQL injection testing was performed
- Manual approach was used for finding deadly XSS.

A lot of information was provided by the first four, tests which favoured the manual testing approach and that is where we hacked. We'll now discuss the web site vulnerability that lead to the OBS hacking.

***“The attack surface depends on the existing level of vulnerability and the way it is triggered”***

The https protocol was well defined. After scrutinising the URL it was found that it was possible to inject some input. The URL was:

```
https://isp_example.com /business/default.jsp?referrer = /home/bank/&app=business
```

So rogue input was subjected to this link as:

```
https://isp_example.com /business/default.jsp?referrer = '><
```

The response was displayed on the web site screen and the login panel was changed.

The truncated URL that appeared on the screen was:

```
&lang=en&action=goto&from=default.jsp&usecase=Password&step=&origin=default.jsp%3Finst%3D%27%253
```

```
E&requestURI=/business/' title="View Help for this Page" target="_blank" tabindex='300'>Help Print
```

The web site was therefore showing vulnerability signs due to insecure coding; but this did not help our study as such. After testing the application with a number of XSS attacks, the result was null. For example, the following failed completely:

```
https://isp_example.com /business/default.jsp?referrer = '%3Cscript%3Ealert('Pimped');%3C/script%3E
```

After analysing the cause of the error we found that we could inject an argument with properties in the middle of code. The URL for the error was as shown in figure 3.

Voila! The online click event was there. Without wasting time an onclick event was injected as the following XSS code:

```
https://isp_example.com /business/default.jsp?referrer = 'onmouseover=%22alert(document.cookie)%22
```

This test was successful; it acted like a poison. The code was injected as a help object and when a user moved the mouse over the help button in order to click it, the domain cookie was extracted. This is somewhat persistent in nature and the attack pattern resulted in number of different attacks.

**Case study – ripple effect**

The case study shows that there is XSS vulnerability present in the main internet service provider's web page

```
<a onclick="openNewWindow('/direct/info_help.jsp?inst='>&lang=en&action=goto&from=default.jsp&usecase=Password&step=&origin=default.jsp%3Finst%3D%27%253E&requestURI=/business/' ,width=500,height=300,top=100,left=100,toolbar=0,location=0,directories=0,status=0,menubar=0,scrollbars=1,resizable=0');return false;" class='button' id='NavigateHelp' href='/direct/info_help.jsp?inst='>&lang=en&action=goto&from=default.jsp&usecase=Password&step=&origin=default.jsp%3Finst%3D%27%253E&requestURI=/business/' title="View Help for this Page" target="_blank" tabindex='300'><span class='text'>Help</span></a>
<a class='button' id='NavigatePrint' href='javascript:void(window.print());' title="Print this Page" target='_self' tabindex='300'><span class='text'>Print</span></a></div>
```

Figure 3: Code used to inject input into an integrated service provider serving a Canadian credit union.

on which login is done; consequently the centralised transaction web site is prone to attacks. There are probably a number of clients whose traffic is purged into this main web site. The attack surface and its effect is diversified. The attacker exploits this vulnerability widely as the attacks are not limited to one client; they affect all clients using online services. Some of the attacks are listed below:

- Cross-site request forgery attacks (these can be performed very easily)
- Traffic redirection to a fake destination
- Malicious script injections to steal identities
- Phishing.

The ripple effect is evident. The results are disastrous from a financial perspective and exploitation is very high. The cloud encapsulates in its vicious sphere all clients that are dependent on the service provider.

## Conclusions

The service provider industry is growing day by day, and a secured mechanism to carry out financial and monetary transactions is needed. Internet threats should be properly modelled. The attack surface depends on the existing level of vulnerability and the way it is triggered. The business ecosystem is a very important factor in this scenario – different entities have different impacts on business. It could even be called business logic, and the OBS ripple effect is an example of this. The vulnerability does not exist in the main web site but a single dependency on an insecure service provider may cause panic. Security is a dual process when it comes to the dependency of clients and service providers. If this not applied correctly then you are on the edge of disaster.

### About the author

*Aditya is an independent security researcher and founder of SecNiche Security.*

*He is a leading author for the Hakin9 group, writing security and hacking papers. His research has been featured in the Usenix magazine ;login:. Aditya's academic background includes a BE and MS in cyber law and information security from the Indian Institute of Information Technology (IIIT-A). Aditya has spoken at conferences such as EuSecWest, XCON, OWASP, and CERT-IN.*

*Aditya's other projects include Mlabs, CERA and TrioSec. He has written security papers released at Packet Storm security, Linux security, InfoSecWriters, XSSed portal, etc. and he has advised various leading companies.*

*At present Aditya is working as an IT advisor at KPMG IT Advisory Services.*

### Resources

SecNiche Security <[www.secniche.org](http://www.secniche.org)>  
Aditya's personal blog, 0Kn0ck <<http://zeroknock.blogspot.com>>

# An overview of the best known virtual solutions

Dario Forte, CFE, CISM, founder and CEO, DFLabs

**Our overview of virtual solutions will begin with the numerous available open source options. Xen is one of those currently attracting the most attention, thanks to its performance and zero cost. The Xen virtual machine component is part of the kernel, which makes it faster than others of the same type. Most system administrators who have begun working with Xen understand this type of virtualisation to be more commonly used in a Linux environment. Xen has, however, been developing for a while toward Windows.**

Currently, the most widely used commercial alternative is VMware – in all its versions. VMware has been the undisputed leader in the field for a variety of reasons. It has survived the criticisms levelled by many of its users regarding its price and it has now released a functional player at zero cost, leaving it able to invest itself in more important tasks which are better suited to its expertise.

## Application environments

Although laboratory tests on the compatibility of virtualisation products with the Mac Pro architecture long ago demonstrated positive results, analysts are still curious to know which security applications can be generally implemented with the support of virtualisation.



Dario Forte

***“Xen is one of those currently attracting the most attention, thanks to its performance and zero cost”***

The management of endpoint PCs is one of the most common starting points in this type of environment. The IT department must be able to centrally manage PCs that are owned by third parties (e.g. consultants and outsourcers)

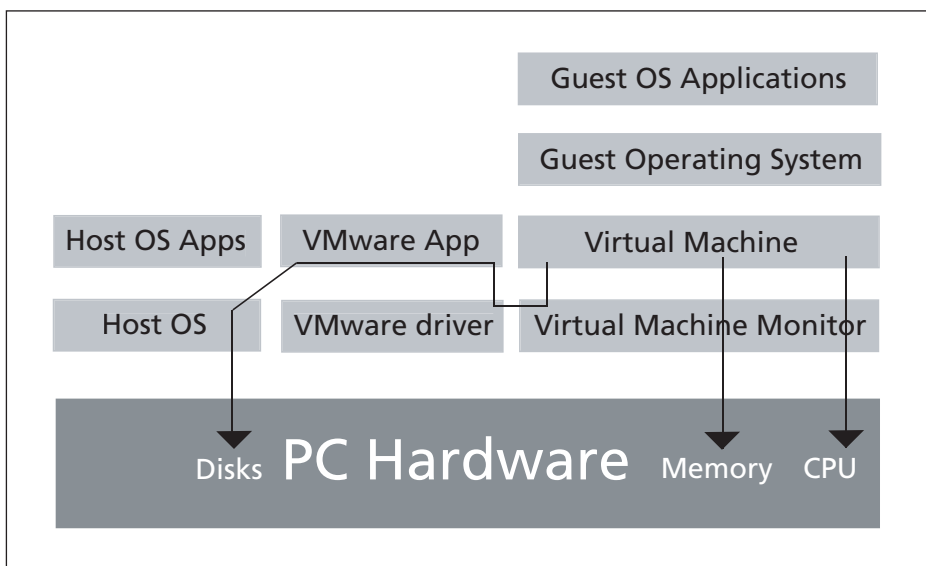


Figure 1: A typical structure of a VMWare virtual machine.

who need to access company resources. An endpoint PC is thus a virtual machine configured according to the IT environment in which the external user needs to operate and it is delivered to that user, who will use it to access company IT resources.

Security of data on PCs is another requirement that must perforce be managed centrally. One reason for this is the need to comply with privacy legislation. The most common solution for this requirement is file system encryption. Using virtual machines, it is possible to use the same type of encryption while maintaining a strong barrier against attackers in the event of a violation of the host machine. Some vendors actually interface this method with the use of token hardware.

Another commonly requested application that virtualisation can provide is remote access management. This allows users to access particular company network resources in a security cell, in which every action can be traced. This kind of application also comes in handy when consolidating several workstations under one, or under a server, the latter being the better option.

## Vendors in support of virtualisation

Virtualisation involves both hardware and software. For some time, Intel and AMD have been developing native vir-

tualisation support (in the CPU), which will result in a reduction in latency time and greater productivity. In the meantime, producers of automated incident response architectures are also working on the issue. They are introducing tools for managing the imaging process of compromised machines directly from a central workstation. This will further decrease response times and thus the incident impact (which is directly proportional to the length of the response and remediation time). A VMWare feature, called 'encapsulation', has been available for some time. Encapsulation

not only completely isolates the virtual machine (and its memory) but also allows the VM file to be used for additional incident response analyses using live system analysis (LSA).

## Upcoming events

According to the web site [www.virtualization.com](http://www.virtualization.com), Mass High Tech is reporting on a new virtualisation startup, which is currently in stealth mode and preparing for a public launch in the field of laptop virtualisation.<sup>1</sup>

The software would be designed to isolate a laptop's four major components – hardware, operating system, applications, and user data – and create versions on the laptop of those components that operate remotely and without a network connection. The software would operate transparently for the user and it would enable wholesale security updates and other types of upgrades when connected to the network. Isolating the portions of the system enables it to limit access to work-related applications while enabling users to install their own personal preferences on the device.

Nortel has also unveiled a business solution that gives road warriors an 'office on a stick', using a specially formatted USB key for Windows PCs. The

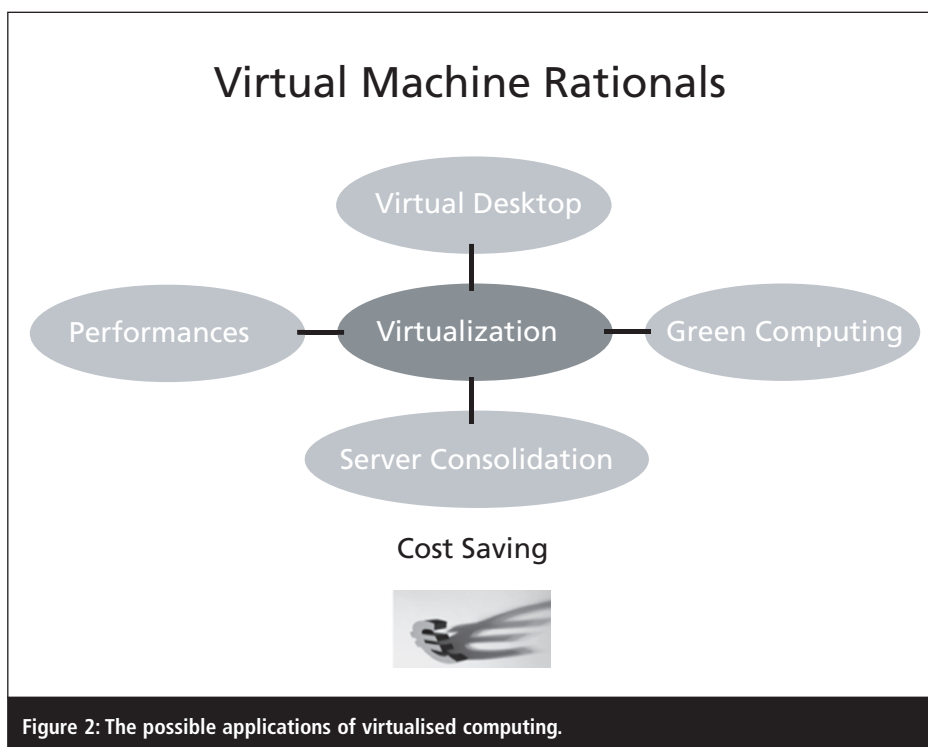


Figure 2: The possible applications of virtualised computing.

stick automates business network access when inserted and protects information and applications by completely removing them when the USB key is removed.

### **“Virtualisation is not the exclusive domain of system administrators and computer geeks”**

The two examples mentioned above demonstrate that virtualisation will not be taken as a definitive panacea but as an IT solution with few security implications. If a similar trend had been introduced five or six years ago, many sceptical ‘gurus’ would have said there was no future for it. Now a more constructive and positive approach is suggested: “if you don’t solve the problem, you are part of it.” For this reason virtualisation offers a good opportunity to improve global security within an organisation.

### **Is it just for geeks?**

Virtualisation is not the exclusive domain of system administrators and computer geeks. Many end users may benefit from it as well; especially advanced mobile users, who,

with their classic two gigabytes of RAM, are able to operate a number of environments for their work applications vs. technical applications. The intrinsic limitations of the spread of the malicious code that is characteristic of virtual machines should also be taken into account. Even in these situations, attacks can be managed via containment. Compromising a virtual machine (when possible) has less serious consequences, in theory, than a ‘conventional’ attack, due to the sandbox architecture that contains the virtual machine. This condition, however, should not lull the security department into a false sense of security.

In conclusion, virtualisation is an investment that may achieve more than one purpose; IT as well as security management, for example. For this reason its cost may be spread across several centres, thus ensuring an increase in the general level of protection and a more tangible return on investment.

One factor remains to be addressed: digital investigations. A virtualised solution should be implemented with a close look at the recovery/investigative phase. This is not always done

and, if not, it could be the cause of failure in case of security incidents.

### **About the author**

*Dario Forte, CFE, CISM, former police detective and founder of DFLabs (www.dflabs.com) has worked in information security since 1992. He has been involved in numerous international conferences on information warfare, including the RSA Conference, Digital Forensic Research Workshops, the Computer Security Institute, the US Department of Defense Cybercrime Conference, and the US Department of Homeland Security (New York Electronic Crimes Task Force). He was also the keynote speaker at the Black Hat conference in Las Vegas. He provides security consulting, incident response and forensics services to several government agencies and private companies.*

### **Reference**

1. “Stealth ‘laptop virtualization’ company virtual computer readying launch”. Virtualization.com. 8 September 2008 <<http://virtualization.com/rumors/2008/08/29/stealth-laptop-virtualization-company-virtual-computer-readying-launch/>>

# How bad is it?

**Bruce Potter, CTO, Ponte Technologies**

**Security products are a commodity. The use of products such as firewalls, anti-virus systems, web filters, and IDS are well understood and documented. The purchasing of security products is well integrated into the IT spending of most organisations, complete with technology refresh cycles and very competitive procurement processes. The whole process can make an enterprise feel relatively secure based on the sheer volume of time invested and dollars spent.**

But how effective are these products? It can be hard for an enterprise to judge. Sometimes attackers manage to get through the defences due to either weaknesses in the products or blind spots in your network. However, it is hard to take the few successful intrusions and many

failed attacks that happen every day inside of a small to midsize network and extrapolate the real world ramifications. Few networks have a broad enough view of the internet and the complete attack space to be able to gauge the real effectiveness of security products.



**Bruce Potter**

### **Google research**

At the 2008 USENIX Security conference in San Jose, California, two researchers from Google presented research they had performed with researchers from John Hopkins University, using Google’s large data-

base of URLs and vast computing power.<sup>1</sup> Niels Provos and Panayiotis Mavrommatis examined the content on millions of web pages looking for suspicious or malicious activity on the page. By leveraging many virtual machines and the Google database, the team was able to look at more than 66 million distinct web pages. The team sought to specifically identify pages that were doing drive-by attacks versus pages that were merely hosting malicious content.

### **“Of the 66 million pages examined, 3.4 million of them contained drive-by attacks”**

In order to determine if a page was malicious, content in the Google database was examined using a custom developed machine learning process that scored each page for likelihood of containing drive-by attacks. If a URL scored high enough, it was subjected to more detailed analysis. The Google team created a farm of virtual machines that were safely sandboxed in order to create a secure, observable environment. Then, using the virtual machines, they loaded pages and watched for meaningful changes such as registry access, unexpected filesystem modification, and application crashing.

Not only was the page loaded, but all linked IFRAMES and SCRIPTS were loaded as well, even if those objects came from a remote system. If a virtual machine was changed in a way known to be malicious, then it was determined to be contain a drive-by attack. Some pages changed the virtual machine, but didn't match known patterns. These pages were deemed suspicious but not necessarily malicious.

The results of the research are staggering. Of the 66 million pages examined, 3.4 million of them contained drive-by attacks. Further, another 3.4 million

pages contained activity that was suspicious but didn't match known malicious behaviour. These pages may also actually be malicious, but it was time-prohibitive for the team to examine them.

Contrary to conventional wisdom, the malicious pages weren't mostly hosted on the seedier parts of the internet such as adult and gambling websites. While there were a large number of drive-by infections on adult sites, the majority of the malicious data is hosted on sites whose categorisation is more mundane such as finance, home and garden, and business. Also, Google discovered that on average between .4% and 1.4% of all URLs returned for Google queries were for malicious sites.

## **Changing presumption**

The real kicker in the Google research was the work it did to determine the detection rate for security products in the face of this onslaught of drive-by attacks. Google ran each malicious page through three different major anti-virus products to see if a given product would find and stop the malicious activity. Google did not identify the products used except to say that they were “well known anti-virus engines.”

The AV products didn't fair well. The top performer stopped, on average, 70-80% of the malicious software from compromising the virtual machine. The next best performer only managed about 50% successful detection. And the worst product only stopped about 30% of attempted attacks.

These numbers do not paint a pretty picture of the state of enterprise security. According to Google, 1% of the sites your users are likely to run into while searching the internet will attempt to shove malware onto your systems. Further, your AV products probably have about a 50% chance of stopping

the infection before it can successfully compromise your system. In an enterprise consisting of thousands of users, it's likely a new infection happens every day and you won't even know it.

Ultimately, the Google research helps reinforce the trends seen in the broad IT security space. It used to be sufficient to buy a few security products, stay current on your pages and signature files, and feel that most of the attackers would be kept outside the wall. The latest attack trends combined with Google's depiction of the state of drive-by attacks show that even with the basic security infrastructure in place, attackers will still get in. Dedicated adversaries with targeted attacks have historically blown through security perimeters. But now even the generic, broad scale attacks that used to be stopped by AV and web filtering are likely to break through your security perimeter as well.

## **Product response**

Developing a response to this situation is difficult. Ideally, rather than changing the way your security organisation is run, you could buy products that are better geared towards stopping modern attacks. Some existing product vendors have attempted to add more functionality to their current capabilities.

For instance, Symantec's anti-virus offering has expanded to include a more holistic concept of endpoint protection. The general idea is that this new breed of AV product doesn't just look for signatures of known virus but also looks at the overall integrity of a system, such as changing files, processes, and users, to determine if an attack is underway.

Similarly, firewall vendors such as Check Point are now including more and more functionality in their firewalls. Beyond simple port filtering, Check Point's SmartDefense enables administrators to



## **A SUBSCRIPTION INCLUDES:**

- 12 printed issues
- Online access for 5 users
- A three-year archive of back issues
- Free delivery

[www.networksecuritynewsletter.com](http://www.networksecuritynewsletter.com)

turn on web filtering, content inspection, protocol validation, and proxies for common protocols. The idea is that if the attack is allowed through the firewall by policy, maybe one of these other sub-systems will inspect the flow of data and find the malicious activity.

For other product vendors, the situation is an opportunity to introduce entire new classes of security products to the broad IT enterprise security market. While some of the products starting to take hold today are not new, they are nonetheless becoming more applicable than in the last few years. Network behavioural analysis (NBA) uses network sniffing and network flow data to look for anomalies in traffic patterns. Rather than dealing with specific patterns, these tools look for data that is different from the norm. NBA products have been around for years (even under different names) but due to the lack of effectiveness in current security products, they are getting renewed attention.

***“Many organisations have put their emphasis on defending against attacks in an effort to keep them from being successful, but fewer have also put emphasis on responding to successful attacks”***

Other product groups such as document loss prevention, secure virtualisation, and high speed proxies are also seeing a bump due to the current trends. It's at times like this that administrators and engineers find creative solutions to hard problems. IT departments will find the best tool for the job, not necessarily the ones that are pitched by the product vendors. Don't be surprised if the killer app for preventing drive-by attacks isn't actually a security-specific tool.

## Operational ramifications

The sad fact is, networks are getting broken in to. You will have to deal with this fact. Many organisations have put their emphasis on defending against attacks in an effort to keep them from being successful, but fewer have also put emphasis on responding to successful attacks. As enterprises begin

to deal with the increasing success of their attackers, more money will have to be spent on detection and response capabilities.

Employees need to be better trained to deal with incident response. Incident response needs to occur as rapidly and efficiently as possible. This means trained employees with the right tools that can identify the scope of the attack, remove the compromised systems, and restore service as quickly as possible.

## Parting shots

Those who are defending networks from malicious attackers are finding themselves at a distinct disadvantage. The tools they have counted on for years are becoming ineffective. Further, things as mundane as surfing the net looking for sports scores have hidden dangers that can put your users and your network at risk. Your network may be infected by malicious code right now even though all your security systems report everything is OK.

These problems are not easily solved. Through creative use of existing products, deploying products that are just now emerging on the scene, and tuning your security operations, you can try and get back in front of this problem. Unfortunately, the way the trend is going, the attackers will continue to have the upper hand for some time to come.

## About the author

*Bruce Potter is the founder of The Shmoo Group of security, crypto, and privacy professionals. He helps organise the yearly ShmooCon security conference held each winter in Washington DC. Mr. Potter is also the founder of Ponte Technologies, a company specializing in wireless security, IT security operations, and advanced network defence techniques.*

## Reference

1. Niels Provos, Panayiotis Mavrommatis, Moheeb Abu Rajab and Fabian Monrose. 'All your iFrames point to us'. USENIX, June 2008. <[www.usenix.org/events/sec08/tech/provos.html](http://www.usenix.org/events/sec08/tech/provos.html)>

## EVENTS CALENDAR

7–8 October 2008

### Sector 2008

Location: Toronto, ON, Canada

Website: <http://sector.ca>

20–22 October 2008

### IHTCIA 2008 International Training Conference & Exposition

Location: Atlantic City, NJ, USA

Website: [www.htciaconference.org](http://www.htciaconference.org)

22 October 2008

### Security Canada Central 2008

Location: Toronto, ON, Canada

Website: [www.securitycanadaexpo.com/en/attendees/central/index.html](http://www.securitycanadaexpo.com/en/attendees/central/index.html)

26–30 October 2008

### Second Workshop on Information Credibility

Location: Napa Valley, CA, USA

Website: [www.dl.kuis.kyoto-u.ac.jp/wicow2](http://www.dl.kuis.kyoto-u.ac.jp/wicow2)

28–30 October 2008

### Conference on Risks and Security of Internet and Systems

Location: Tozeur, Tunisia

Website: [www.redcad.org/crisis2008](http://www.redcad.org/crisis2008)

5–7 November 2008

### 6th ACM Conference on Embedded Network Sensor Systems

Location: Raleigh, NC, USA

Website: <http://senssys.acm.org/2008>

10–12 November 2008

### 7th International Workshop on Digital Watermarking

Location: Busan, Korea

Website: <http://multimedia.korea.ac.kr/iwdw2008>