

Featured this month:

De-perimeterisation: inevitable and essential

In a world where employees are mobile, businesses are highly interconnected, and the internet is king, the difference between 'secure' private networks and 'insecure' public networks is blurring. To tackle the security challenges of this new way of doing business, security professionals need to de-perimeterise, focusing on data rather than the network, and phasing out traditional perimeter security.

Key elements of de-perimeterisation are hardening every endpoint, web-enabling applications, encrypting data at rest and in transit, and enforcing data-level authentication and authorisation. The new, de-centralised, distributed model has been proven to support the secure exploitation of the potential of ebusiness.

Go to Page 7...

A window to the world?

Web 2.0, cloud computing, AJAX, and web applications represent the future of internet use. As a result, web content has undergone an explosion as businesses integrate data and services and present complex internal IT systems to users on the outside of the network perimeter. This represents a risk to the network, as a poorly configured CMS means that network security

can be quickly bypassed, data assets corrupted, user information accessed, and the network compromised.

Tamas Biro, co-founder of Sense/Net, looks at potential threats from a badly managed corporate website, and how to take steps to minimise the issues.

Go to Page 11...

Cybersecurity: Obama takes care of business

President Obama treated federal cybersecurity as one of his first points of business after taking office last month. Within a month of his inauguration, he set out a cybersecurity policy, and appointed a former key national intelligence executive to conduct a wide-ranging review of federal cybersecurity.

The president appointed Melissa Hathaway, former top cybersecurity advisor to the director of national intelligence under former president Bush, as acting senior director for Cyberspace for the National Security and Homeland Security Councils. She will develop a strategic framework to ensure that cybersecurity initiatives are properly integrated and

resourced, said a White House statement, and will also co-ordinate cybersecurity efforts between the federal government, Congress, and the private sector.

Hathaway's appointment makes her a likely candidate for permanent 'cyberczar', in which she would oversee the federal Government's entire cybersecurity effort. Obama laid out the position when he transferred his cybersecurity strategy from his election campaign site to the White House site.

In addition to a national cybersecurity advisor, he also vowed to promote research and development into safe computing, with an emphasis on "a new generation of

Contents

News

Cybersecurity: Obama takes care of business	1
Kyrgyzstan comes under cyberattack	2
Cost of breaches rises	3

Features

Password protection: the next generation

Passwords and access codes are the core foundations to protect access to computer systems. This article looks at how these passwords are constructed, the underlying technical details and lessons to enhance them as security threats increase.	4
--	---

De-perimeterisation: inevitable and essential

Mobile working and ebusiness have blurred the boundaries of the corporate network. Alastair Broom of Dimension Data explains how networks must adapt to cope with new security requirements.	7
--	---

A window to the world?

New web publishing techniques are improving consumer experience but when built on insecure platforms or principles, are they wide open doors into corporate networks?	11
---	----

From vulnerability to patch: the window of exposure

The gap between the identification of a vulnerability and the release of a patch is a crucial period. Aditya Sood explains how researchers and product vendors navigate it.	14
---	----

Three tips for your network

Bruce Potter outlines three must-do tips for securing your network.	16
---	----

The death of MD5

Existing research had already rendered the MD5 hashing algorithm suspect. More recent work has confirmed its inadequacy. Dario Forte discusses what researchers found – and what happens next.	18
--	----

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
Petterson, 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:

€1052 for all European countries & Iran

US\$1177 for all countries except Europe and Japan

¥139 600 for Japan

(Prices valid until 31 December 2009)

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 da-
mage 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

secure hardware and software". The agenda
also calls for new cybersecurity standards,
anti-corporate cyber-espionage systems,
and a cohesive strategy to fight cybercrime.

He also promised to mandate stand-
ards to secure personal data, along with
data breach notification rules.

Kyrgyzstan comes under cyberattack

Security experts were musing over
the perpetrator of a DDoS attack on
the former soviet state of Kyrgyzstan
last month, following a flood of traffic
from Russian networks that pounded
the country's relatively underdevel-
oped network infrastructure. High-up
officials in the Kyrgyz ISP community
reported the attack on 18 January to
intelligence groups in the West.

Don Jackson, senior security researcher
at SecureWorks, was convinced that the
Russian government was behind the
attack, which he theorised was an attempt
to block western access to independent
media sites in the country. The opposi-
tion to the Kyrgyz Government had been
using the sites to oppose the suggested
closure of a US airbase in the country,
which had been debated as a sweetener
for a Russian investment deal.

Jeffrey Carr, CEO of private sec-
tor cyber-intelligence firm GreyLogic,
believed that the Kyrgyz Government
could have stopped the attack at any time.
Instead, he suggested that it had paid
Russian hackers to instigate the attack,
with the same suggested motive: to quash
the opposition's voice.

Rafal Rohozinski, principal investigator
for the OpenNet Initiative, argued that nei-
ther theory could be proven, and suggested

that many different groups and motives
may have perpetrated the attack, for any
number of personal or commercial reasons.

The DDoS traffic abated during the
first week of February, a few days before
a meeting between the Kyrgyz president
and Russian officials to thrash out the
details of the investment deal.

Cost of breaches rises


The average cost of data breaches is
rising, according to a report from
the Ponemon Institute, which says that
lost business is the biggest expense
for companies that have their data
pilfered.

The Ponemon Report, *2008 Annual
Study: Cost of a Data Breach*, also revealed
that data breaches as a result of outsourced
functions had increased over previous
years, and that first-time victims suffered
greater costs than those that had already
experienced a breach. The per-victim cost
of a data breach was \$243, compared to
\$192 for those who had suffered breaches
in the past.

The 69% of data breach costs stem-
ming from lost business was reflected in an
increase in the average customer churn rate
suffered by companies that were victims of
data loss. The 2008 churn rate was 3.6%,
up from 2.67% last year. The healthcare
industry was the most affected, with a 6.5%
churn rate, followed by finance at 5.5%.

The percentage of breaches down to
third-party organisations such as contrac-
tors, customers, and outsourcers increased
to 44% in 2008, over 40% in 2007, and
29% in 2006.

The survey examined the experiences of
43 US companies across 17 different sec-
tors that had suffered from data losses.



Average cost of a data breach (Source: Ponemon 2008 annual study: Cost of a data breach).

In brief

Symantec spots media Trojan

Symantec warned of a Trojan infecting files used by Windows Media Player. Trojan.Brisv.A. infects .asf, .mp2, .mp3, .wma, and .wmv movie and music files, sending the program to a malicious URL when they are played. It also converts .mp2 and .mp3 files found on the system to the .wma format, making it hard to clean them by disguising the malicious commands as digital rights management code.

ITGI wants risk input

The IT Governance Institute is asking the public to comment on its proposed IT risk framework, based on COBIT. The comment period for *Enterprise Risk: Identify, Govern and Manage IT Risk: The Risk IT Framework* ends on 13 March. www.itgi.org

Google flags entire net as bad

Google temporarily flagged the entire internet as dangerous in its search results for a short time in late January. For 40 minutes one Saturday, anything showing up in search results was flagged as potentially harmful. The mistake happened due to human error when updating a malicious site list provided by Stopbadware.org, said a representative.

Opal standards to lock down drives

The Trusted Computing Group has outlined specifications for native storage encryption standards. The Opal standards cover storage in PCs and data centres, and for interoperability with protocols including SCSI, ATA and Fibre Channel. They can also be combined with the group's Trusted Platform Module for increased security.

Nexor updates content protection

Information protection firm Nexor has developed two additions to its family of information guard appliances. Its Sentinel 2 product, which ensures the compliance of outbound and inbound data, features SNMP monitoring, SSH, S/MIME, multiple simultaneous content types, structured security labels and the ability to use third party content filters. The company also ships Sentinel 3 in March 2009. It runs on the SELinux security extensions founded as a development project by the NSA.

ID fraud rising, costs falling

US identity fraud victims soared by a fifth to 9.9 million adults last year, according to Javelin Strategy & Research's *2009 Identity Fraud Survey* report. Nevertheless, the average cost of identity fraud per consumer dropped by 31% to \$496 per incident, as fraud is discovered and stopped earlier.

Microsoft patches remote execution vulns

Microsoft released patches for three vulnerabilities with an exploitability index of 1 (meaning

that consistent exploit code is likely). Two of the vulnerabilities allowed remote code execution when malicious web sites were viewed using IE7. Another allowed remote code execution via an SQL injection attack in SQL Server.

Sophos finds new polymorphic virus

Sophos found a new polymorphic mid-infecting virus called W32/Scribble-A. The malware infects the host file at arbitrary locations in its executable code, rather than simply targeting the entry-point, says the firm. It can also modify htm, html, php and asp files with an iFrame that points to a malicious site.

Online game no fun for CISOs

NCC Group found that one third of 14 000 senior executives in blue chip companies, were willing to open and play an online game without knowing its origin. Many people also forwarded the game, spreading it to 19 different countries. The Flash game could have been used to direct a browser to any web site, said the firm, or to pop up a fake Windows login screen to grab passwords.

McAfee puts head in cloud

McAfee has launched a software-as-a-service (SaaS) security unit to house all of its cloud-based services.

Netgear launches SMB security appliances

Netgear launched a range of security appliances aimed at small businesses. The first three launched under its ProSecure brand offer web and email protection for different numbers of users, up to a maximum 600. They are licensed on functionality rather than on a per-user basis.

Porn sites hit with innovative DNS attack

The Register reports that attackers have been fooling DNS servers into bombarding porn sites with root nameserver information by spoofing malformed DNS requests. The attack causes a flood of information to be sent to the victim's server with minimal traffic outlay, said Don Jackson, senior security researcher at SecureWorks. Researchers are now waiting for the attack to make its way into commercial exploit kits.

HP patches printer flaw

HP has issued a security firmware update for selected LaserJet printers. It fixes a vulnerability that could allow hackers to access files stored in its memory. <http://bit.ly/hIkVU>

Geeks.com not so geeky

Compgeeks.com, which operates the Geeks.com web site, has settled a case with the FTC, which complained that the company did not implement proper security measures to protect private customer information.

Windows 7 security flaw isn't, really

Microsoft officials have responded that a 'flaw' in Windows 7 isn't a vulnerability at all. Concerns were raised that digitally signed code could execute elevated-privilege operations in Windows 7 without being prompted for approval by the user. Microsoft argues that the code would have to be running in the first place, and that Windows 7 includes several features to stop that happening. The default setting for the User Access Control (UAC) feature in Windows 7 had been changed to make it less annoying that it is in Vista.

Nessus adds database feature

Tenable Network Security has expanded Nessus with SQL database configuration auditing capabilities. Only Security Center users and ProfessionalFeed subscribers can access this feature.

IE surfing safety tool unveiled

CyberPatrol is offering a free online web-surfing safety assessment tool. It includes the firm's Threat Detector tool, which helps analyse surfing behaviour to look for inappropriate or unsafe sites. Unfortunately, it only works with IE. <http://bit.ly/CEBd1>

Spam domains highly concentrated

Anti-spam outfit Knujon has says that 83% of domains used by spammers can be tied to just ten domain registrars:

- XIN NET
- eNom
- Network Solutions
- Register.com
- PLANETONLINE
- RegTime
- OnlineNIC
- SpotDomains
- Wild West
- HICHINA Web Solutions

Data Breach Watch

The **US Federal Aviation Administration** lost over 45 000 records to hackers after a computer was illegally accessed.

An individual was arrested with a computer file containing the records of 29 000 **Kaiser Permanente** employees. Employee name, address, phone number, Social Security number, and date of birth information was in the file. Only a handful of employees have reported identity theft, said the firm, which is providing one year of free credit monitoring to employees whose information was included.

An employee at US retailer **Best Buy** was arrested by the Secret Service after using a card skimmer to steal the credit card details of 4000 customers.

Password protection: the next generation

Tom Rowan, security consultant, Magirus



Since the dawn of computing, passwords and access codes have been used to protect access to computer systems. These passwords protect user accounts. Systems ask for a username and password before access is granted. As some user accounts have increased privileges on systems, the security of these accounts is often given more thought than normal user accounts. However, obtaining the password of even a lowly unprivileged account can be used to gain a toehold onto a system. This can provide a potential springboard to escalate privileges and gain administration access.

This article will discuss some of the problems associated with this security model. It will cover the storage of passwords, methods used to attempt obtaining them and some of the mechanisms used to protect them from discovery. It will not cover the techniques needed to obtain a password database from a computer system. This subject is too complex, varied and even contentious for discussion here. This article will discuss passwords in general terms that apply equally to all general-purpose operating systems; some are better at password security and some are worse. The principles are the same regardless.

Forcing open the cookie jar

Passwords are stored in some form of database. This database might be as complex as a directory service such as Microsoft ActiveDirectory or OpenLDAP or as simple as a flat text file. Regardless of the operating system concerned, the database provides a one-to-one mapping between a user account and its related password. The database might also contain other pertinent metadata including system UIDs, group memberships, authorised login times and contact details.

Obviously, it is essential to ensure the security and integrity of the information inside the database from outside attack. To this end, passwords are very rarely

stored as a simple text string. Usually, a one-way mathematical function is used to convert the password into a value that can be stored more safely. This function is known as a one-way hash function, and is designed so that a unique value is derived from a unique password; no two passwords should ever hash to the same value.

When a user attempts to log into the system, their password attempt is run through the same function and the result is compared with the stored value. If they are the same, then access is granted.

This is a sample UNIX password file:

```
bob:$6$JN2efjEK$N.zOr891yppdkW
YEyUe4SfliuesdcRIYBpkBmtJR7Sz9vR
Dc4Oyi14.q08GPrafj4DiTYgqRD4cshY
G2popTg0:14273:0:99999:7:::
alice:$6$EBBjg1me$VCN/
WxruiSq4P/USUHx2Zyflvd5GpzuR
SHnm1Socx4jRJBAPTVpsg8a3Yq3eelii
8ziWH6uZd5cJfXcBX61A/:14273:0:99
999:7:::
```

If the database was stolen from a system, the passwords inside would not be immediately apparent to the attacker. Moreover, if the hash function is sufficiently robust it will not be possible to reverse-engineer a password from the hashed version. However, the attacker can recover the clear text password by using a brute force technique. The attacker generates every possible password in turn, calculates its hash value,

and compares it with the captured target value. Eventually, the clear-text password would be found that produces a match.

“With a short password limited to a small alphabet there is only a very small table space to check”

For example, consider a password that must be four characters long and must be numeric. This means that the alphabet is ten characters in size (0–9). The stored password for the user account is, and has been hashed using the MD5 algorithm. The attacker simply needs to generate each of the possible passwords in turn, apply the MD5 function and compare to the stored value. The table demonstrates this in practice. Naturally, the process stops as soon as all the passwords in the database are found.

With a short password limited to a small alphabet there is only a very small table space to check. The number of possible passwords for this example can be calculated using the equation $num = alph^{len}$, where num is the number of passwords, $alph$ is the alphabet size and len is a fixed length.

In this case $n = 10^4 = 10\ 000$. This is a fairly trivial problem for a modern computer and would be solved in a few seconds at most. More complex password hashes will require much more computing power. But, in theory, given enough

time it would be possible to crack any password in this way eventually.

Read the dictionary from cover to cover

Rather than systematically moving through the entire possible key space, a more conservative search can be performed. A dictionary-based approach can improve the efficiency of the attack by searching only where success is more likely. This method works by running each word from a predetermined list into the hash function and comparing the result against the target password hash. The 'dictionary' can be as simple as a short list of common passwords through to the whole English dictionary.

Often, customised dictionaries are useful such as substitution tables where numbers are substituted into dictionary words, for example p4ssw0rd or p@55w0rd, password123. Other possible substitutions include capitalisation (Password, PAssword, PASsword), insertion of symbols (_password, +password) or combinations of all of these (_P4ssw0rd).

These generated dictionaries are used alongside unmodified words to improve the chances of obtaining a password. Most people do not produce sufficiently strong passwords, so a dictionary attack using the words in the English Dictionary will generally pick up a large percentage of passwords. Appending numbers onto names, such as jessica1966 or arsenal99, usually finds the majority of the rest.

The traditional way of running either a brute force or dictionary attack against a set of captured hashes involved the following cycle:

1. Choose password guess
2. Hash
3. Compare
4. Repeat

Nothing is stored for longer than required to complete this attack run. This means that the process has to be repeated in its entirety for each set of target hashes. Each run is separate from any previous run.

At the end lies a pot of gold

A more modern approach is to pre-compute the entire list of possible hashes just once. The resulting table of hashes/password pairs is then indexed. It can then be searched for the target values. Then it is a simple matter of reading off the corresponding clear-text password from the table.

"Software is available freely to create rainbow tables on a Windows or Linux machine if enough local computing time is available"

The disk space required to store these tables can be very large, but with the decreasing cost of terabytes of storage space, it is now very cost effective to expend the processing power once in the production of tables for entire key spaces. These tables are known as rainbow tables and are generally attributed to Philippe Oechslin, who developed the concept in a Doctoral paper in 2003.¹ A table must be produced for each key space that is to be searched. For instance, different tables would be needed to search MD5, SHA-1 and NTLM key spaces.

Software is available freely to create rainbow tables on a Windows or Linux machine if enough local computing time is available.² Tables are generally produced that include either a 'full' ASCII dictionary or a limited subset. Typical dictionaries would be lowercase alphanumeric, full alphanumeric, and for the truly dedicated full alphanumeric plus symbols.

The tables produced vary in size from a few hundred megabytes to tens of gigabytes.³ These large tables can take a while to search for a matching hash value, even when cleverly indexed. A search can take several minutes to complete depending on the speed of the CPU and the disk IO subsystem. This has been proven from experience to be very painful on an external hard drive accidentally located on a USB v1.1 bus!

For those that cannot afford to wait, there are commercial downloadable

tables for most useful password hashing algorithms available from a number of organisations. It pays to shop around if a purchase is being contemplated; some tables are available for free, but their commercial counterparts may come with better performing compression and indexing software. Equally, they may not.

There are even online hash submission sites. Using these sites is as simple as submitting a password hash (usually unsalted) into a web form and clicking submit.⁴ In the background, the website's servers are searching rainbow tables for a match. Queue times can be very long on busy sites and the user should be aware of the dangers of submitting passwords into essentially unknown hands. If a test is to be attempted before, say, embarking on a purchase of a set of tables, it is humbly suggested that the reader's best password is not given away!

Rub salt in the wounds

Ironically, operating system security features such as forcing a user to comply with password complexity schemes can actually decrease the required search space slightly. If an attacker knew that users had to choose a password made up of alphanumeric characters with a minimum number of symbols added, then they would not have to search the simple alphanumeric part of the key space at all. While this sounds counterproductive, in practice it is far better to have passwords with a bigger possible alphabet. Security is improved greatly by forcing users to choose from a bigger alphabet than to enforce longer passwords. This can be

Alphabet	Length	Number
10	4	10 000
10	5	100 000
10	6	1 000 000
26	4	456 976
26	5	11 881 376
26	6	308 915 776

Table 1: Comparative complexity of password length versus alphabet size.

proven by revisiting the simplified equation for showing a key space size: $num = alph^n$. Table 1 shows how the key space grows when each of these attributes is changed.

To complicate matters for the attacker, a 'salt' is usually added to the hash function. This is used to improve the encryption of the password by using both the password and some random bits of data (the salt) as inputs to the hash function. The implementation of this may be as simple as concatenation of the password and the salt, such as: (mypassword+'salt') -> hash -> encrypted_passwd. For each binary bit of salt added, the key space to be searched is increased massively.

"To complicate matters for the attacker, a 'salt' is usually added to the hash function"

This phenomenon is best illustrated with an example. This example assumes that the salt value is kept secret from the attacker, and the dictionary to be checked is a mere ten words long – perhaps a list of known pet's names. A random 16-bit salt would mean that the salted value could be one of $2^{16} \times 10$. That is $65\,535 \times 10 = 655\,350$. Increasing the salt to a 17 bits value would increase this to some $2^{17} \times 10 = 1\,310\,720$. Remember that without any salt at all, the attacker would only have to perform ten hash calculations. On the contrary, with seventeen bits of salt added, they have to perform up to 1.3 million tries against the target hash value in order to exhaust all possible values of salt and password in combination.

The Oxford English Dictionary contains full definitions for 171 476 words. That means that even an attacker who knows that a user's password is a simple English dictionary word must perform up to $2^{17} \times 171\,476 = 22\,475\,702\,272$ (nearly 22.5 billion) hash and compare operations. The combinatorial explosion accelerates as the salt values increase.

It is, of course, possible to pre-compute all these values as well. However, the computation time and disk space

requirements are enormous. The scale of the problem is obvious when it is considered that a table will be required for one bit of salt, then two bits and so on. The tables will increase in size exponentially as the salt value increases.

Working together at last

One method that can be employed to either generate rainbow tables or perform 'live' cracks against a set of target hashes is to utilise parallel or grid computing. This involves dividing the task into smaller chunks and sending these to a network of computers to solve. This is akin to the search for extraterrestrial intelligence at home (SETI@Home) project, where data chunks are analysed using the CPU idle time on millions of computers across the globe. There are a number of applications for password recovery that utilise this technique. In this scenario, providing the computers perform in a similar manner, the time to complete the task can be approximately divided by the number of machines being used.

Taking this concept further, a fairly recent development in the password cracking field is the use of dedicated graphics processors to power through the calculations. The graphics processing unit (GPU) vendors have been trying to promote their hardware as a general-purpose processor offload calculation engine for some time. These highly complex chips are designed to solve very specific mathematical problems very quickly, such as vector set up calculations, texture mapping and alpha blending. It happens that the mathematics behind these operations can also be used to seriously accelerate password cracking attempts.

If hash functions can be recoded to utilise API calls provided by the software development kits for GPUs, then they can be run very quickly indeed. Running these operations in dedicated hardware adds orders of magnitude to the number of operations that can be carried out per second. Elcomsoft, the Russian firm that first proposed a commercial product using this technology

in October 2007, reports a decrease in recovery time in the order of a factor of 20 to 25.⁵

Whereas a single core top end Xeon processor can handle in the region of three to five million MD5 operations per second, a 25-fold increase means that a GPU could handle 125 million.

"It is predicted that salt values will be increased massively in future password obfuscation and encryption libraries"

We found this figure after running a lab test using md5crack-12.exe on Windows 2003 Server, running on VMware ESX 3.5, 1 Virtual CPU, Xeon 3.1 GHz. Result: 4.5 million MD5 operations per second over three tests of full English Dictionary. This is an impressive figure. Suddenly, the massive computational problems caused by large salt values and complex dictionaries become less of a worry for the attacker.

The example used earlier with seventeen bits of salt and a password being chosen from the Oxford English dictionary resulted in a maximum requirement for 22.5 billion tests. On a single CPU at even an optimistic five million checks per second, this would take in the order of 75 minutes to complete. Using the hardware power of a GPU, the test should be some 25 times quicker – that is in the order of three minutes.

As discussed previously, even a rainbow table search can take several minutes and relies upon gigabytes of stored tables. The use of a GPU to bring the time requirements of live cracking attempt down to a fraction of the previous is a major step forward. It is particularly interesting to note that multiple, compatible GPUs can be used by some multi-threaded software implementations to provide a clustered attack engine. Already, figures in the billions of password tries per second are being quoted for these scenarios.⁶

It is predicted that salt values will be increased massively in future password obfuscation and encryption libraries. Also, new hash algorithms may well be developed which are designed to be

less susceptible to these forms of attack. However, even if these changes were made today, the majority of systems in the wild would not be running them for some time. It may be a decade or more before key systems running on corporate networks have been migrated away from operating systems that employ vulnerable password database systems. The opportunities for password theft at high speed will remain worryingly prevalent.

References

1. Dr Philippe Oechslin's web page, EPFL <<http://lasecwww.epfl.ch/philippe.shtml>>
2. Project RainbowCrack web page, accessed February 2009 <http://project-rainbowcrack.com/>
3. Rainbow Tables web page, Shmoo Group, accessed February 2009 <<http://rainbowtables.shmoo.com/>>
4. Password.ru website, accessed February 2009 <<http://passcracking.com>>
5. Andrew Brandt, "Password cracking chip causes security concerns", *New Scientist*, October 2007 <www.newscientist.com/article/dn12825>
6. Andrey Belenko, "Modern password recovery: GPUs and supercomputers" <www.elcomsoft.com/presentations/GPUs_and_supercomputers.pdf>

De-perimeterisation: inevitable and essential

Alastair Broom, UK security line of business director, Dimension Data

My business's employees are mobile, our applications are highly interconnected and we use ecommerce platforms, extranets and the web to communicate with suppliers, partners, and customers. So where is the perimeter of my network? And how do I protect it if I cannot define it?

In 2003, a group of CISOs from global corporations such as BP and ICI set out to answer these questions and in January 2004 they formed the Jericho Forum. Their aim was to determine the best way to enable secure e-business in a world without boundaries where the internet is king. They concluded that hard perimeters separating 'secure' corporate networks from 'insecure' public networks, inside from outside, and trusted from untrusted were not conducive to this new way of doing business. So they coined the term 'de-perimeterisation' to reflect a fundamental shift in the way the corporate network in general and network security in particular should be conceptualised, designed, constructed, and managed. With de-perimeterisation, the corporate network and the internet are synonymous and the focus for security moves from the network to the data.

Why de-perimeterise?

An open and highly interconnected environment where business is

conducted over the internet brings numerous benefits including improved business agility, global reach, and cost-effective communications. These days, business communication and businesses themselves cannot prosper without engaging to an extent with this new milieu. Therefore, de-perimeterisation and the corresponding erosion of the traditional constraints of hard perimeters are inevitable.

For example, an organisation's ecosystem of partners and suppliers requires degrees of trust rather than a simple trusted/untrusted dichotomy. Employees, suppliers, and business partners alike will have different requirements for accessing corporate resources and a hard perimeter limits the ability to provide this. The perimeter is designed to keep the bad guys out and to protect the corporate network from infiltration and abuse. It is not designed to enable collaborative working within an extranet environment.

Perimeter security is also becoming less able to fulfil its fundamental purpose which is to protect the corporate

network from external threats. New threats emerge daily, requiring the perimeter to be continually updated, added to, or supplemented with complementary technology. Among other things, this creates a management headache. Modern threats are content-rather than network-based and evade traditional security technology.

For instance, you might receive an email containing a graphic comprised of pictures and text. The text invites you to click on an embedded URL that takes you to a website. You are then prompted to click another button and download an application or otherwise run some kind of script on the site that subsequently infects your machine and your network.

Firewalls do nothing to protect against this threat. Anti-spam engines that rely on keywords and text classification fail to identify this as spam because the text is within a graphic and the email subject appears harmless or is disguised using random text, foreign characters, or number and letter combinations. Content filters also fail to pick

it up as malicious as they are unable to interpret the graphical content. At best they will classify it as 'suspicious' if they cannot verify the reputation of the sender.

The effectiveness of this architecture is further reduced by the rise in mobile working which enables employees to simply walk past the security perimeter with their laptops, PDAs, and portable media devices and plug into corporate and public networks at will.

But what can de-perimeterisation offer?

While ebusiness provides numerous benefits, it is inherently insecure. De-perimeterisation addresses this security challenge so that the benefits can be realised without compromising network security. It promises highly flexible business communications where a new business partner or customer simply needs an internet connection and the proper authorisation to access a company's applications and information. No firewall reconfiguration, no lead times for private circuits, and no specialised client software are required.

De-perimeterisation does this by taking a data-centric rather than a network-centric approach to security. This means that every endpoint must be hardened, all data must be encrypted, and all users and applications must be given access rights to only the information that their privileges allow. Security moves from a centralised to a de-centralised, distributed model.

Securing every endpoint, encrypting all data, and authenticating at the data level allows organisations to relax or completely remove the security perimeter. This in turn can provide a highly flexible environment where new applications, users, and third party connections can be rolled out quickly and easily without the need for firewall changes, VPN configuration or new client software.

So what is the catch?

The inevitability of de-perimeterisation is backed up by findings of a survey

carried out at the Jericho Forum Conference 2007, which found that over 90% of European security professionals believe networks will be de-perimeterised within five years.¹ However, a similar survey carried out in the US revealed that businesses would still be reliant on a hardened perimeter in five years. While this difference is probably due to EU organisations having a better understanding of the benefits of de-perimeterisation, it nevertheless highlights that for many the jury is still out.

"Over 90% of European security professionals believe networks will be de-perimeterised within five years"

The reasons behind this include practical issues such as the significant up-front cost and investment of human resources that a de-perimeterisation project requires as well as an associated disruption to business. To achieve the de-perimeterisation dream all applications must be web-enabled, all data must be encrypted end-to-end – in transit and at rest – and data-level authentication and authorisation must be enforced. This brings its own challenges in terms of the reach and technical capability of the IT department. New technology may need to be provisioned to every endpoint, data must be carefully classified, and access policies defined. All of this takes money, specialist expertise, and time.

The up-front cost, however, is offset by a reduction in ongoing operational management costs. De-perimeterisation makes moves, adds, and changes to the network easier and faster to implement, and network management is less resource intensive by virtue of a flatter network architecture with fewer security devices. WAN bandwidth costs will also be reduced as the internet replaces private networks.

Other factors impeding the adoption of de-perimeterisation are less concrete, but just as powerful. It is easy for a CISO to argue the case for implementing a strong security perimeter. The concept is analogous to the physical

security world and easy for CEOs and CFOs to understand. It is the norm today, tried and tested, and although not flawless it works reasonably well. Compare this with the de-perimeterisation argument: remove the perimeter, connect all endpoints directly to the internet and transact all business over the public network. Justifying the move from the familiar to something new, particularly in the security space, is a tough call. This requires a shift in mind-set which even security professionals find challenging.

But is the idea of de-perimeterisation so radical? Every time you log on to your internet banking site, you are exploiting the bank's move to de-perimeterisation. You are accessing data stored on a banking system served via an internet-facing web server. You authenticate yourself via your bank card or other credentials and have access only to your data on the system. Your transactions are secured using VPN technology built into your web browser and if you use hard disk or file encryption and endpoint security software already, you are almost in total compliance with the de-perimeterisation blueprint.

In the consumer world, de-perimeterisation is ubiquitous. It is not risk-free, but it is an accepted way of conducting highly sensitive transactions. And with the general consumerisation of IT, users will naturally begin to expect the same level of access and freedom in the corporate world as they enjoy from their armchair. De-perimeterisation supports this approach and helps bridge the expectation gap.

So the business case for de-perimeterisation is challenging but logical. Security professionals must make the case for capital expenditure versus operational savings and increased agility and convince the business that de-perimeterisation provides at least as good security as traditional models.

It is important to note that de-perimeterisation may not suit every organisation. Small companies with little third party integration, largely office-based staff, or limited ebusiness operations may find the security perimeter entirely

adequate. If they do not, however, de-perimeterisation will deliver significant benefits.

How do you go about de-perimeterising?

As mentioned above, full de-perimeterisation assumes that the corporate network and the internet are essentially the same thing. Endpoints connect directly to the internet and use encryption to protect the confidentiality of information passing between them. Encryption based on SSL VPNs rather than IPsec is a logical choice, providing end-to-end encryption, application-level control, and ubiquitous support. This implies that all applications in the de-perimeterised environment must be web-enabled.

As corporate information is now exposed to the internet, it is critical that authentication is reliable and robust at both the user and application level and authorisation is controlled at the data level. In addition, users will not want to authenticate to every application or file they wish to access, so single sign-on (SSO) is an important technology for improving the user experience and minimising any additional complexity introduced by de-perimeterisation. De-perimeterisation promises improved business agility and this must not happen at the expense of user agility.

With perimeter security in place, it is relatively easy to control access at the user, system, application, or service level. Without the perimeter, all systems are exposed to all users and access control needs to be much more granular and, most importantly, data-centric rather than network-centric. To enable this level of access control requires some form of data classification or, ideally, digital rights management technology (DRM). DRM enables individual files or data sets to be transmitted with their associated privileges (e.g. ability to print, copy, or forward) on a per-user or per-group basis. (It is important to note though that DRM technology is still relatively new and involves more than just the security department. Its

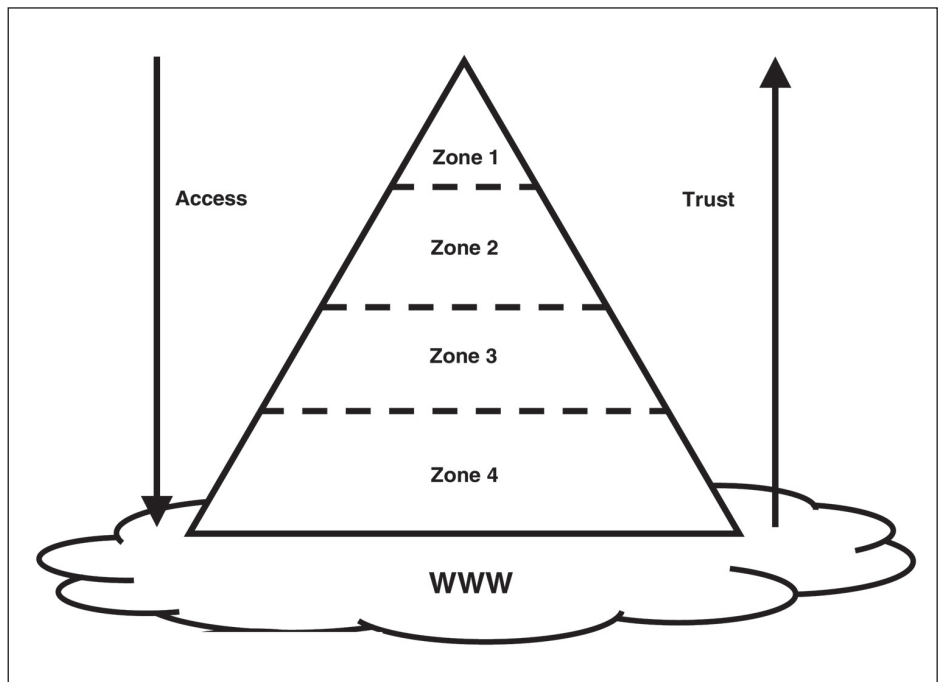


Figure 1: Transition to de-perimeterisation.

implementation is a significant undertaking and the associated cost should form part of the de-perimeterisation business case.)

“With perimeter security in place, it is relatively easy to control access at the user, system, application or service level”

As all corporate PCs and laptops are now effectively internet nodes rather than connecting via a private network, it is imperative to protect the endpoint via robust endpoint security (e.g. firewall, anti-x, and intrusion prevention) and encrypt all data at rest to prevent data leakage either accidentally or as a result of a compromised machine. Encryption should be extended to portable media devices such as memory sticks, external drives, and CD/DVDs and desktop security policies should be robust enough to enforce this.

What does this mean for existing security controls?

Where does de-perimeterisation leave the firewall, traditionally the

cornerstone of any security implementation? Can we throw this out with the other perimeter technologies and remove the perimeter completely? “Not yet,” is the short answer. The implementation of de-perimeterisation in practice does not follow the utopian route at the moment, and while full de-perimeterisation is the ultimate goal, the perimeter will remain in place for some time yet. What is happening is a gradual migration of the perimeter security technologies to the endpoint, while the perimeter is also being shored up against the evolving threats.

For example, the threats which network security seeks to counter, such as viruses and spyware, will not disappear with de-perimeterisation. Content filtering, secure proxies, anti-virus, anti-spyware, anti-phishing, and a host of other technologies are already required in addition to the firewall and intrusion detection and prevention system to mitigate these risks. Much of this is currently deployed at the desktop, so partial de-perimeterisation is happening even within the traditional perimeter model.

As organisations de-perimeterise, a new architecture will emerge over time rather than overnight. For example, branch offices will be connected directly

to the internet rather than via costly private circuits into the head office for the purposes of shielding behind the corporate perimeter. Firewall, intrusion prevention, and VPN functionality will move from the network boundary out to the endpoint.

The role of the traditional perimeter firewall is evolving in the short to medium term to provide protection between defined security 'zones' within the corporate network. Each zone represents a trust domain with a clear demarcation boundary around it. This 'zoning' essentially creates multiple perimeters and provides a transition path to a fully de-perimeterised network by allowing zones to be de-perimeterised in phases.

Figure 1 illustrates the transition to de-perimeterisation. Zone 1 may represent simple internet access and clients using this zone connect directly to the internet. Zones with higher security requirements sit within increasingly secure perimeters providing demarcation from the internet. Over time, as de-perimeterisation becomes more mature and widely accepted, the model 'sinks' into the internet like an ice lolly melting in the sun. Boundaries are gradually removed and more sensitive applications and data are exposed directly to the internet as confidence grows and adoptions become more mainstream.

However, this transitional model adds complexity, and adding complexity to security inevitably increases risk. There are more things to break and more chance of configuration errors. Businesses need to weigh up the risks associated with this transitional approach against the cost and disruption of a wholesale move to de-perimeterisation.

Risk in a de-perimeterised world will be largely governed by the effectiveness of the encryption, authentication and endpoint security solutions. The same threats exist in both the perimeterised and de-perimeterised environments and the confidentiality, integrity, and availability (CIA) security mantra remain. Confidentiality will be determined by the strength of the encryption used, digital certificates should ensure integrity is maintained, and availability will be

improved by the very nature of the pervasive internet and the flatter, less complex architecture that de-perimeterisation provides.

What is happening in the real world?

In the UK, several companies have already started down the de-perimeterisation route. BP now has around 30 000 users communicating securely over the internet rather than via a corporate network. ICI has taken steps to dismantle the perimeter, initially via an in-the-cloud service agreement with ScanSafe. The service cleans all web traffic before it gets to the user, irrespective of where he or she is in the world, thereby removing the need for web security at the network perimeter. The service also saves on WAN bandwidth costs, provides users with better response times, and the business with global visibility of the users – who they are and what they are doing.

In 2007, Dutch airline KLM gave 4000 employees allowances to buy PCs, which they were to manage themselves. The move was a nod to the Jericho Forum and resulted in reduced support costs. The PCs used the same security suite as corporate counterparts and connected to the corporate network through an IPSec VPN using two-factor authentication.

Conclusion

The move towards de-perimeterisation is inevitable in a world with only one network – the internet – and where businesses seek more flexible working environments and closer engagement with their business partners, customers, and suppliers. New communication channels are now simply a question of access rights and authorisation.

In the corporate environment, the significant investment in traditional perimeter security, coupled with a natural paranoia towards any radical changes to security architecture, are undoubtedly barriers to early, large-scale adoption of de-perimeterisation. What is evident, however, is that businesses and governments are now realising that a data-

centric approach to security is essential and deployment of technologies to support this approach will support the move to de-perimeterisation over time. While the current economic climate does not support large-scale capital investment projects (another potential barrier to adoption), companies who have already de-perimeterised will find that the flexibility and agility this affords will provide competitive advantage in the turmoil of mergers and acquisitions that inevitably accompanies these market conditions.

As a security professional, the most important thing is to understand de-perimeterisation, accept it, and plan for it. If you do not, you will be left behind as everyone around you finds ways to circumvent the perimeter in the name of business agility. As the traditional perimeter continues to crumble and its effectiveness diminishes in the face of new threats, mobile working, and e-commerce, it is time to prepare for an alternative solution that improves security while creating the optimum environment for ebusiness.

About the author

Alastair Broom is UK security line of business director for Dimension Data. Alastair has a complex portfolio of responsibilities that includes providing business development and presales support and resources to the UK sales team, developing and implementing the company's UK security strategy, and management of its portfolio of security products and services. He also manages vendor relationships and the development of new products and services, and maintains Dimension Data's security accreditations. Alastair has built a twenty-year career in IT, specialising in security over the latter half of his career, and has various vendor accreditations, including being a certified ISO 9001 auditor.

Reference

1. "Jericho Forum survey reveals growing awareness of need for de-perimeterization solutions among global security leaders." Jericho Forum. 27 August 2007, 8 November 2008 <https://www.opengroup.org/jericho/uploads/40/14529/JF_USlaunch_release_070828.pdf>

A window to the world?

Tamas Biro, co-founder, Sense/Net

Organisations must share information that is public and protect information that is private. This seems to be a simple enough concept for the network manager. However, with Gartner research showing worldwide security software sales topping \$10 billion in 2007 and Forrester predicting 21% of IT security budgets growing in 2009, security continues to be an afterthought – a bolt-on to network infrastructures rather than central to information strategy.

As a result, protecting networks has become a game of cat and mouse. New technologies spur criminals to find innovative ways to access corporate data which then forces corporations to call for more complex protection and so on.

The growth in security holes and fixes has paralleled that of web usage for business, not just for marketing and communication but as a platform for transactional and dynamic ebusiness. Software-as-a-service and cloud-based computing, increasingly capable web applications, and savvy consumers have caused an explosion in web content and usage.

Now websites have moved from being mere shop windows or static, web-based catalogues to deep, integrated elements of overall business IT. Looking at stock levels of a book or CD on an ecommerce site is not second-hand information

that someone has updated – it is live information taken from the warehouse system's database. But choosing that level of integration has created new challenges in keeping networks secure.

Smash and grab

When planning a new website or the build of an online portal, large organisations traditionally outsource the work to third-party contractors or web agencies. These third parties are highly skilled at communication and design but less adept at securing complex technology infrastructures. They often create custom web applications in which the concentration on functionality is so great that they simply incorporate 'visual security' – security implemented as an afterthought and not as an integral part of the architecture.

A simple login feature, for instance, with selective functions made available after authentication is no security at all. It is a 'nod' to security that leaves a gaping hole in the network and one that more and more people are focused on exploiting.

The content management risk

Unless a company opts for one official portal platform that is standardised for use throughout all departments, the balance will remain in favour of small websites, usually one per project, campaign, or product. Every new website creates a potential new window for attack.

How does an enterprise eliminate risk? IT managers need to protect content from various well known threats, such as destruction, unauthorised alteration, and data being compromised. What kind of system can protect against these threats? The problem lies in this question. Security is not 'a' system, it is 'the' system and a portal is only secure as its weakest link. Many IT professionals overlook this seemingly obvious fact and this results in serious vulnerabilities.

Ways in over the web

1. SQL injection

SQL injection can be as simple as it sounds; a mere piece of code used in place of a username when attempting login on a website. The SQL command then runs a routine behind the site, allowing the attacker to enter the site itself and either manipulate code or change/corrupt data. Although code injection can get more complex, the reason it is so shocking is the frequency of its success.

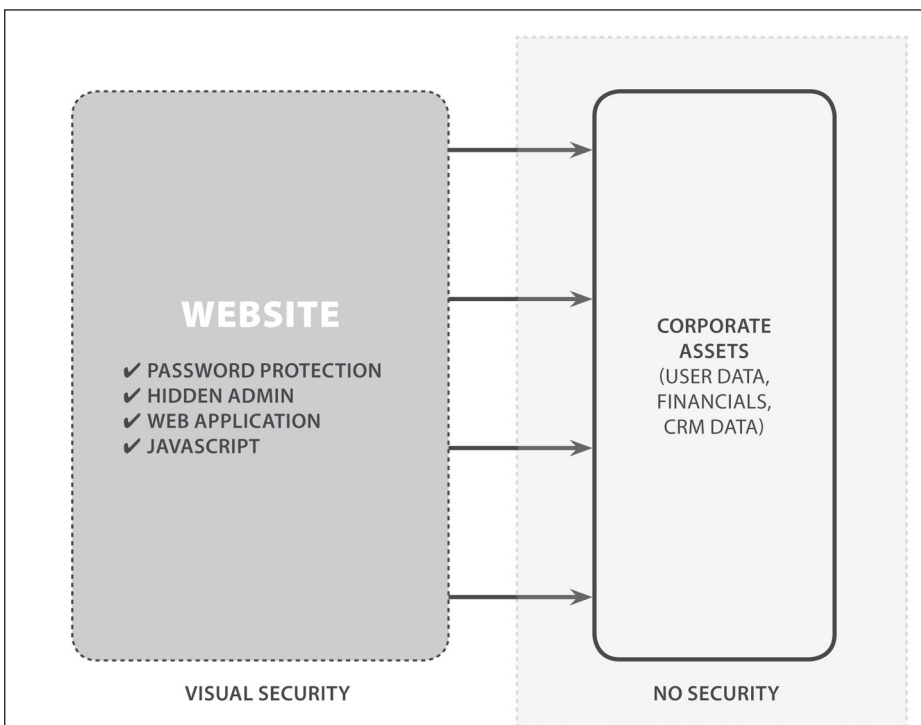


Figure 1: Visual security is little more than lip service.

If the database driving the website is vulnerable to accepting code via online forms or URLs, the issues might not stop at login. Any form that communicates with the database, such as a search box or 'forward to a friend' field may be a way to inject code past the security measures and into the network.

2. Hijacked links

An apparently innocuous website URL hides a multitude of ways in which an attacker can sufficiently manipulate users or systems in order to access or change data behind a firewall. One way would be SQL injection via the URL in order to embed commands in the URL line that instruct the web server and tell the site to display user data, passwords, etc. Another would be rogue, embedded URLs in forums or on social media sites that, if clicked on, can execute local commands, steal cookie information, or change other data.

"An apparently innocuous website URL hides a multitude of ways in which an attacker can sufficiently manipulate users or systems in order to access or change data behind a firewall"

While cross-site scripting (XSS) attacks are more challenging to deploy than SQL injection methods, they are less easy to trace due to the fact that a great deal of software and web servers do not even spot anything untoward.

3. Passwords and authorisation cracking

Brute-forcing a website password is common and can range from knowing the user and guessing that his or her dog's name is the password, through to sophisticated methods of automated password cracking. Another well known example of password hacking is through the use of default administration details. For this method to work, attackers have to be aware of the default settings of content management or security systems and those settings must not have been changed since implementation.

Passwords can also be bypassed through simple tweaks to HTML code, such as viewing source, removing

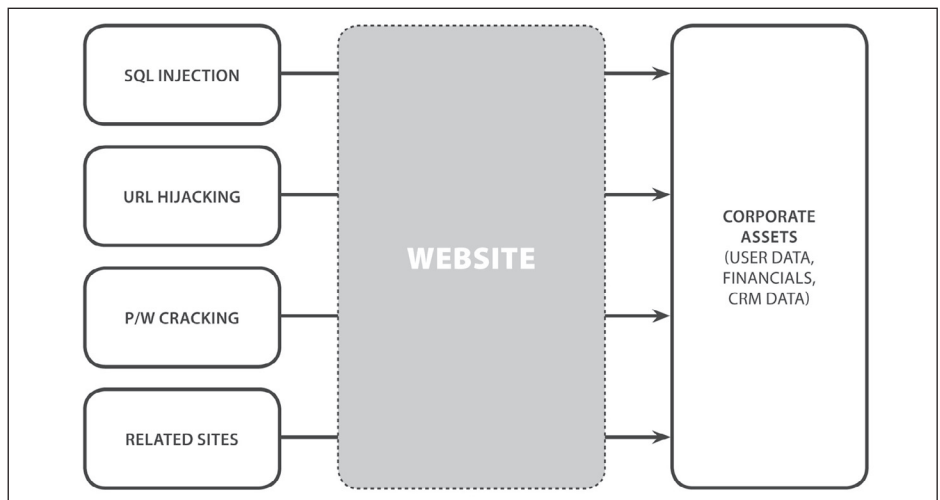


Figure 2: The website can be a welcoming way in.

authentication JavaScript, and running from a local drive. Once more, poorly configured content management will enable an attacker to log in and potentially view and manipulate data.

4. Related sites

You might have impeccable security on your site, including well configured authentication, a regularly patched OS on your web server, intrusion detection, and the digital equivalent of barbed wire running around and across your databases to avoid code injection. What about that trusted partner site, though, that communicates with your sales database? You have no control over its configuration, but have opened your data to the site in order to ensure a smooth and seamless supply chain.

These sites represent an unseen and difficult-to-control area of risk. If there is poor access configuration, they could provide another way in for malicious code or way out for your information.

Ways to fight back

1. Improved permission/security management

An organisation must first look at how it manages its content and, more importantly, how employees manage security. If security management is difficult in any way, it is guaranteed that employees will not manage it. Similarly, centralising permission management usually results in unmanaged security. IT becomes overloaded with setting requests, resulting in

users feeling that they have to circumvent the system to save time. This invariably leads to chaos and security risk.

The key is to choose security software where security management is nearly invisible – not an extra, but a natural way of handling content. Monitoring and auditing are very important, in order to double check users and make sure they are using the system effectively. Moreover, if employees are aware that their adherence to procedure is being observed, they will be more careful.

2. Monitor the gates

An unpatched web server is the most common factor in successful attacks. It goes without saying that a network infrastructure must be secure. More importantly, though, it must be securely operated – not by shutting the gates completely, but by monitoring them. Network managers should resist the temptation to block communication altogether.

The key is to perform the necessary audits and checks needed for secure platform management and consistently review, check traffic, and tweak as necessary. While most operating systems are secure enough to protect software from low-level and network-based attacks and although a well maintained firewall can help to enforce global security rules, neither can replace one-by-one machine configurations.

3. Review the platform itself

Another consideration is the software used to manage content on a portal. If CMS or portal software can be hacked

into easily, an organisation is vulnerable no matter how rigorous the security. How does this occur? Malicious users can take advantage of weaknesses in the software through SQL command injection and poorly written code means web servers and the OS could be compromised.

The ongoing debate about web server and OS security product names continues to rage in developer circles. Which is best? Linux or Windows, Apache or IIS? All the brand names are comparably good and a fondness for one over the other is largely personal taste. None is more secure than the rest, however, and security really depends on two things; the competence of the developers and the operation. The two are not mutually exclusive, both are necessary.

Security as standard

While many of the security issues mentioned tend to put the onus of security firmly in the lap of the developer, development can be done in such a way that weaknesses do not develop. There is no doubt that expertise is a factor when configuring a portal or the CMS software that drives it. No matter how well written the software is, if the operational teams configure it badly, there will be vulnerabilities. Competent sysops and well organised patch management are crucial to a secure portal and CMS.

Diversity is also a factor. Protecting one system is always more effective than protecting a multitude of different systems running on different software and OS. Moves like the Content Management Interoperability Services (CMIS) standard are extremely useful from a developer's perspective but this standard encourages the longevity of many systems. This is not a positive thing for security or integrity of networks which rely on regular updates in the game of cat and mouse.

There is no silver bullet. Setting up an enterprise CMIS standard has many advantages from a security perspective, as well as others such as licensing, training, and support costs.

Another security factor lies in the building of the portal or CMS-based site. I have already pointed out the

challenges inherent to third-party website creation, but there is also the cost of security. By their very size and nature, web applications are as expensive to make secure as security is hard to sell.

Is open source secure?

When it comes to platforms for running websites, open source ECM and portal software is becoming more popular thanks mainly to the fact that it is a more cost-effective solution. Open source CMS can be trialled for free (in stark contrast to proprietary ECM software) and the licensing model means a more flexible budgeting process which is also handy when businesses are unsure where the economy is going next.

Another advantage of the open source approach is the ability to examine the source code for security issues and vulnerabilities. Well written open source, checked and tested by many people, may be better than a big vendor's proprietary system. At the same time, a new vulnerability or method of attack does not create the necessity for one vendor to issue a patch – multiple developers will pick up on the issue quickly and updates to the platform will arise quickly. If you have a development team on board, they will be able to access and update the code directly.

However open source does not have a great reputation when it comes to security. For some, it seems that the larger vendors of proprietary software have more resources dedicated to security research and implementation. This is a fair point but can also mean companies are sluggish and the rigorous tests may ultimately hinder new versions from hitting the market.

On balance, both approaches have upsides and downsides. Open source technology can mean quicker, more fluid updates in the face of security risks. The proprietary model will mean fewer updates but heavier applications and slower release cycles.

A recent survey revealed that 87% of business managers have cited their proprietary software security as a cause for concern. They state that they are unaware of what is happening within the information environment and they do

not have the tools to monitor or control the platform.

Then the solution to the above problems would involve establishing risk-based governance and controls to ensure long-term compliance with business policy and industry regulations. But attempts to regulate users by imposing controls will only incite some to circumvent them.

Less is more

The fact that many company websites are built by design agencies rather than developers was cited as an issue at the beginning of this article. A greater concern, however, could be the growing use of complex, public-facing integrated web applications. In essence, web 2.0 is allowing ecommerce sites to communicate directly with the databases that run businesses. Whereas websites were windows in the past, they are now doors.

In order to ensure these doors do not let unwanted visitors into the network, managers need to approach ecommerce websites with the same inherent cynicism as the network edge. Adding website issues and potential vulnerabilities to the network security checklist would be a good start to managing the issues.

Cisco's suggested network security checklist, for example, poses questions such as: 'Who has access to your digital assets?' 'Who controls those assets?' 'Where do the assets reside?'

When applied to an overlooked, unprotected and poorly configured corporate website, the answers to those questions could be: 'Everyone, everywhere.'

About the author

Tamas Biro was born in Hungary in 1974. He received his primary education there and studied at Nottingham Trent University before graduating in 1994 with a degree in integrated engineering. In 1995, at the age of 21, he started his first and only job to date when he co-founded CMS provider Sense/Net with Sandor Kiss. He was appointed technical director for the company, which has now released the world's first open source .NET-based CMS platform, Sense/Net 6. For more information on Sense/Net go to www.sensenet.hu.

From vulnerability to patch: the window of exposure

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



Aditya K. Sood

It is difficult to believe that there will be no security weaknesses in the software of a developed product. From the time a piece of software is released, a new 'vulnerability war' starts in which attackers and security researchers inspect the software to find flaws.

This building and breaking process is beneficial; it creates a positive feedback loop. But it carries inevitable dangers. The time period between the date of finding a new vulnerability and the release of a security patch is known as the window of exposure. The danger of exploitation is still significant, even in the event of responsible disclosure, where the researcher has disclosed the vulnerability only to the vendor for the affected product. However, this danger becomes critical if the researcher discloses the vulnerability directly to the public.

The black hat community doesn't stop trying to exploit the vulnerability after this window has closed and a patch has been issued. The recent spread of the Downadup worm shows just how beneficial it can be to exploit vulnerabilities patched by the vendors. However, the danger becomes especially acute during this window of exposure, when there is no patch available.

WOE dependency – overview

Many factors influence the effectiveness of the vulnerability mitigation process. These include the type of information released to the public, along with the number of users that deploy a patch once it has been released.

Software companies usually make every effort to start working on vulnerabilities as soon as possible. Once a patch has

been released, though, it still takes time to deploy it. Most systems now seem to have auto update services that scan the software for required updates. Although this crucial security practice checks software for missing patches, user intervention still plays a major role, making the resolution process subject to human error.

Vulnerability reporting by security companies is good practice, and security companies work with researchers to resolve vulnerabilities for vendors. Although they pay researchers for this process, the major aim is to reduce the window of exposure. An example is the Zero Day Initiative (ZDI) founded by TippingPoint. A well defined ZDI program for vulnerability reporting exists, and its advantage is that as soon as a CVE is issued for a vulnerability and public disclosure is made, the vulnerability is normally already patched.

“User intervention still plays a major role, making the resolution process subject to human error”

If the vulnerability is not patched, a disclosure is still made by the ZDI in which only an overview of the vulnerability is released, and not its full details. This lowers the exposure to a great extent.

WOE exceptions

A system without exceptions would be rare, and this holds true for the security world too. Certain exceptions impact directly on a WOE. The software product business model is such an exception. Software contains vulnerabilities that have the potential to affect a vendor's business model. Companies find it difficult to

work on vulnerabilities because changing code patterns in a number of products is not an easy task and it may even result in a business loss. For example, if a vulnerability is revealed in a centralised component of all of a company's products, the severity of the threat in relation to the vulnerability would be very great.

Business processes play a generic role in crafting WOE. For example, the Windows Server service is vulnerable to remote execution due to a flaw in handling RPC service requests. History shows that a variant of this vulnerability was released long ago. This vulnerability can be used for worm infection in networks running Windows XP. If the variant of this flaw was known long ago, then why did it take so long to patch? Only temporary patches were made.

The answer to this question is that technology changes with the passage of time, for example, Windows 2000 changed to XP and then Vista and now we are expecting Windows 7. In addition, internal platforms remain the same to some extent but most system-specific requirements change. This has a serious impact on security, especially when a temporary patch is released for vulnerabilities specific to that platform. Companies do however provide retrospective compatibility features to run codes from old platforms.

A complex process may result in the same vulnerability which was previously patched – with the same component but a different platform. An analysis might reveal that even though the exposure window of the vulnerability was partially reduced, it opened again as soon as the environment changed. This type of event may arise from a hard core development process. Microsoft has taken a major step and

introduced a new methodology, the secure development lifecycle (SDL), which makes security a process and not a static entity. This concept has lowered the exploitation of routine vulnerabilities to a great extent.

Infinite WOE – website attack surface

The conditions surrounding WOE change when the discussion moves to web-based vulnerabilities like cross-site scripting (XSS). Web-based vulnerabilities are different from product-based ones and their affects are critical in products that provide a web-based interface for versatile functioning. Most web-based vulnerabilities are related to development problems, which raises serious concerns. XSS is one of the greatest nightmares in the web application world because it is easy to initiate but it is usually hard to trace. Applications are prone to XSS-driven affects in terms of security, and XSS acts as a base for a number of web application attacks. The complexity of triggering XSS in an application depends on several factors.

“Quality assurance in security does not depend on vendor-specific approaches but more on how individuals secure their machines”

A predictability element is hard to trace in a complex environment. Whereas most automated tools fail to track XSS in a complex environment, manual testing seems to help. Why is XSS an area of complex predictability? A pen tester experienced in web applications may have noticed that XSS sometimes occurs in unexpected parts of an application. This is the nature of XSS prevalence in the web world. The WOE for web-based vulnerabilities is infinite. You can find flaws in one or another way. The WOE lasts longer than the threat cloud.

There is no standard or reasonable approach to identifying vulnerabilities. Developers are dependant on application testers to spot persistent vulnerabilities in software or applications. Secure coding is one of the most important aspects of application development. When designing business applications, logic has to be

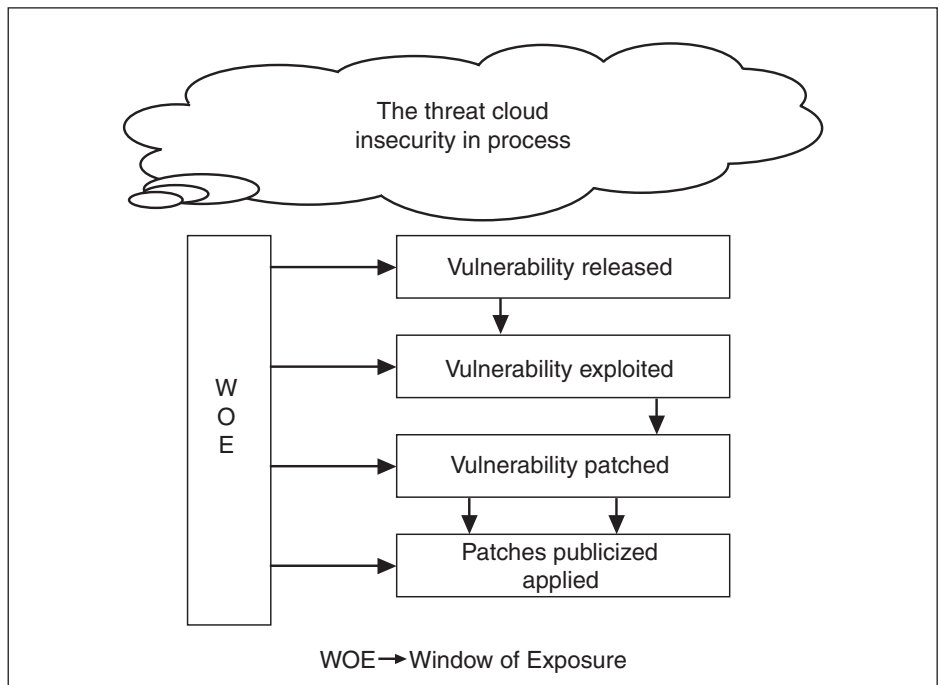


Figure 1: These are the factors on which a window of exposure depends.

accurately constructed. Logic is considered the core, and if it goes wrong the application becomes vulnerable to logic flaws that cannot be circumvented by secure coding. If any wrong parameter is constructed, the business may be on the verge of severe loss.

We have mentioned the exceptional WOE changes that occur when moving from the computer system to a website. In general, however, insecurity remains. The Open Web Application Security Project (OWASP) described WOE as the cyclic process shown in figure 2.

Figure 2 explains many elements in the cycle of vulnerability finding and

patching. Environmental conditions play a crucial role in this process. The WOE is a kind of loophole that may be either reduced or fully closed by mitigating the information flow and making users firmly security oriented. This clarifies the process that should be followed.

Security assurance and time prevalence

Quality assurance in security does not depend on vendor-specific approaches but more on how individuals secure their machines. The process is not restricted to smaller attack surfaces but includes

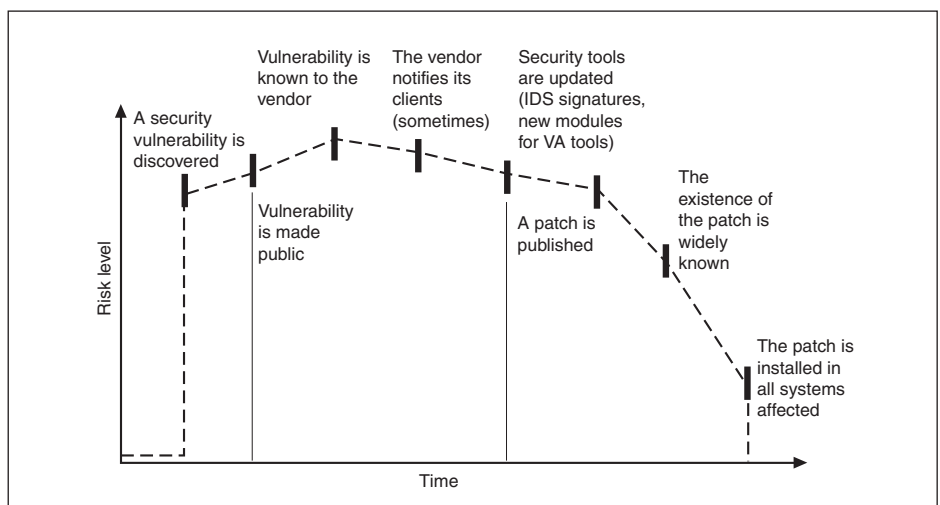


Figure 2: Window of exposure – risk/time statistics.

the whole paradigm. The WOE can be reduced by following a well designed security assurance practice. This should become a process, not a one-time approach.

Since assurance is considered a process, a cycle should be established to track each change. Time is a critical factor in security operations, as evidenced by the fact that WOE is the time lapse during which a product is made secure. Time plays a critical role in vulnerability exposure too and it should be reduced in order to avoid the exploitability index of a vulnerability that has been released. Software companies should speed up patch processing and release to ensure that exposure time is nil and the window is actually closed. Companies could even consider this process as a general detection and response mechanism. Response time is important. As the response time lengthens so does the WOE, which increases risk. The best way to reduce

a WOE is to treat the whole cycle as a process. Mitigation steps should be taken and applied within the time frame in order to avoid attacks and business losses due to technology compromise.

Conclusion

Security must be a continuous process of combating existing attacks and reducing the attack surface. The lifecycle of good software design defines the overall approach to vulnerability detection, response and mitigation. The WOE should be controlled in the best possible way in order to minimise that attack surface. The more complex software becomes, the more vulnerabilities we can expect to find.

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

www.secniche.org
<http://zeroknock.blogspot.com>

Three tips for your network

Bruce Potter, CTO, Ponte Technologies

Often, IT security is discussed as if it exists in a vacuum. Buying security products such as firewalls and anti-virus, responding to security incidents, and the creation of security policies can seem completely disconnected from the rest of the enterprise. Security professionals deal with obscure and sometimes challenging problems that can cause them to put blinkers on and only focus on the issues directly in front of them.

However, having security be a disconnected entity is a sure way to fail. It may seem cliché, but your security organisation needs to be integrated into all your IT and (arguably) your entire business. Further, the business needs to be aware and in tune with IT security needs and goals.

Of all the IT organisations that IT security interacts with, probably the most important is the group that controls the networking infrastructure. Historically the network and security groups are tightly intertwined, at least from a functionality perspective. The

tools and products each organisation uses bump right in to each other and both can have a profound impact on the user population. Both network and security groups have products (routers and firewalls) that can deny access to large user groups with a single catastrophic failure.

However, as much as the functionality of the network and security groups are closely woven, the actual architecture and operations of the two groups can be miles apart. The best efforts of an IT security organisation can be sunk by disconnects and miscommunication with

the networking organisation. Looking beyond the basic network connections for firewalls and IP address allocation, there are a number of things a networking organisation can do to make the network more secure and more conducive to the efforts of the security group.

Build a defensible network

One of the biggest keys to success to the network security program of an enterprise is actually building a network that works in conjunction with the security



Bruce Potter

needs of the organisation rather than against it. The network architecture and design needs to provide the proper hooks and design characteristics to enable the security products and processes to work.

“Of all the IT organisations that IT security interacts with, probably the most important is the group that controls the networking infrastructure”

As an example, think about the state of networks in the mid-1990s. Prior to wide-scale deployments of firewalls and network security products, networks were very flat, with few routers and choke points. As external threats started to develop, firewalls became a fact of life. Unfortunately, the networks weren't architected in a manner that facilitated good firewall placement. Routers had to be reconfigured to bring on new subnets, networks had to be 'unflattened', and security boundaries started to mess with routing protocols. But over time, including choke points close to your network boundaries has become common practice for network engineers.

Fast forward to today, and we have a similar situation. The contemporary threat is against the end user's workstation through spearphishing and drive-by attacks. Even the best firewall in the world can't protect the infrastructure from these dangers. What we are faced with today is a shift in emphasis from defending the network to detecting successful attacks and containing their blast radius. Unfortunately, the running gag about network security for the last decade has been "hard outer shell, soft squishy innards." We have embraced the firewall but left internal security and segregation as an exercise for later. As much as we do have some security capability running on end systems, the overall internal network architecture remains largely unchanged; big VLANs spanning multiple buildings, and logical groups of employees with huge trunks between facilities.

To address the current threat, network architectures need to change dramatically. For example, ideally users are segregated on the network based on their role

and access needs. This makes it much easier to apply policy for fine-grained network access control at LAN boundaries. Also, having clearly delineated and segregated networks facilitates incident response allowing security and network engineers to isolate small parts of the network to contain attackers or to protect critical parts of the network.

For some network this architecture requires port-by-port VLAN configuration. For other networks, NAC/NAP solutions can have a hand in network segregation. Neither of these options are easy to retrofit into an existing network. It can mean large-scale upgrades to the switching infrastructure, upgrades to the cable plant, and even the elimination of personal-use hubs users often plug in at their desk. However if architectural concepts like this are considered in long term planning issues, networks can evolve into a defensible posture.

One other issue that is becoming an issue in more and more enterprises is the need for network taps. Network taps are valuable troubleshooting tools for network teams, and a critical component of a variety of security products. As more and more organisations make a shift to detecting adversarial activity deep within the network with flow-based and packet capturing tools, the need for taps will continue to increase. Unfortunately in modern networks with complex VLAN architectures, creating tap ports that have proper visibility into the right parts of the network can be a huge challenge. This, again, is an issue that needs to be addressed at the architectural level as VLANs, trunks, and switches must all be configured properly to facilitate tap ports, not get in the way.

Secure network administration

Looking past how a network is built, you must consider how a network is administered. Even the best architected and designed network can be broken because of poor administration practices. Network engineers need to be educated on how their actions and management practices impact the security of the enterprise.

One of the most important issues to be addressed in network administration is removing all clear text interactive management mechanisms such as telnet. Nearly every router and switch made in the last few years has support for secure shell (SSH) access. Network engineers are sometimes hesitant to make the jump to SSH. Sometimes that's due to simple habits such as typing 'telnet' instead of 'ssh' at the command prompt. Other times there is concern that SSH induces some extra overhead on the end system, which quite simply is not true. There is no legitimate reason not to use encrypted communications to interact with network devices.

“Network taps are valuable troubleshooting tools for network teams, and a critical component of a variety of security products”

Taking things a step further, if possible, moving to SNMP v3 from v1 is preferred. While the migration from v1 to v3 involves a much larger change than moving from telnet to SSH, it is critical for organisations using SNMP to actively configure devices, not just receive traps and monitoring data.

One last consideration is controlling where administration occurs. It is remarkable how many network engineers log in remotely to routers and switches from places like conferences and public cafes using clear text protocols. Through policy and through filters, external administration without the protection of a VPN should simply not be allowed. Especially with the widespread use of wireless networks, it potentially places the credentials needed to compromise an entire enterprise out in the open for anyone to sniff.

Network documentation

Even the best-designed and run network can't be kept secret. Eventually the network group needs to have architectural artefacts that can be given to other organisations. The security group in particular needs to have an understanding of the network architecture for

placement of infrastructure and containment of incidents.

While for the network engineers, a whiteboard full of scribbling may serve as a valid network diagram, it doesn't work as a tool to use during a large scale break-in. And even complicated architectural issues like the creation of SPAN ports and insertion of network taps can become brain surgery without valid diagrams. Calling network documentation a critical item in the integrity of an enterprise may seem a bit overblown, but it's shocking when things really go bad and a broad cross section of people need to understand the network quickly, good documentation can make the difference between stopping a crisis before it starts, or a full time disaster.

One thing to consider when building network diagrams is to know what you are attempting to represent before you

build the diagram. For example, if you're interested in the physical plant layout, don't worry about routers and higher-level devices. Or, if you're creating a logical diagram of network boundaries, only include devices that serve as LAN or enclave boundaries such as routers and firewalls. As a general rule of thumb, do not intermingle OSI layers on single diagram. Don't have routers and switches on the same piece of paper. When attempting to represent the logical segregation (ie: routers) and the switching infrastructure in a single diagram is a challenge for the engineer to create and an even bigger challenge for the reader to understand.

Parting shots

The security of your network infrastructure is a key aspect of the overall security of your enterprise. However,

the security of your network goes beyond simply having good passwords and a couple of firewalls at the network borders. A network needs to be built and maintained in a manner that is conducive to the overall security goals of an organisation. This means not just meeting the needs of the security organisation today, but building the foundation to assist the security organisation in addressing future threats and meeting future goals.

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. Bruce is also the founder of Ponte Technologies, a company specialising in wireless security, IT security operations, and advanced network defence techniques.

The death of MD5

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

MD5 has been used widely so far. It is still heavily used in areas such as banking and even computer forensics. However, a recent security research paper seems to have reaffirmed fears that this algorithm is no longer suitable.

Everything started in 2004, when collisions were announced in SHA-0, MD4, MD5, HAVAL-128, and RIPEMD. French researcher Antoine Joux presented the collision in SHA-0, and a group of collisions against MD4, MD5, HAVAL-128, and RIPEMD were found by the Chinese researcher Xiaoyun Wang with co-authors Dengguo Feng, Xuejia Lai, and Hongbo Yu. After that, in February 2005, the same Xiaoyun Wang, Lisa Yiqun Yin, and Hongbo Yu found collisions in SHA-1 using 2^{69} hash computations.

Collision versus pre-image attacks

At that time, the same group of researchers clarified what they had

found. In particular, a differentiation between collision attack and a pre-image attack has been defined. This is very important, especially for those working in the forensic industry. While a pre-image attack would enable an attacker to find an input message that causes a hash function to produce a particular output, a collision attack finds two messages with the same hash, but the attacker can't pick what the hash is.

It is important to define what kind of exploit an attacker could carry out with a collision attack. It would typically involve the construction of two messages with the same hash. At that time, the attacks announced were collision attacks, not pre-image attacks.



Dario Forte

Alexander Sotirov, Marc Stevens, Jacob Appelbaum, Arjen Lenstra, David Molnar, Dag Arne Osvik, Benne de Weger identified a vulnerability in the public key infrastructure (PKI) adopted to issue digital certificates for secure websites.¹ The attack is not easy to perform, so the researchers published also a proof of concept (PoC) executing a practical attack scenario and successfully creating a rogue certification authority (CA) certificate trusted by all common web browsers. This gave them a certificate that could be used to impersonate any internet website, including banking and ecommerce sites secured using the HTTPS protocol.

The attack takes advantage of a weakness in the MD5 cryptographic hash

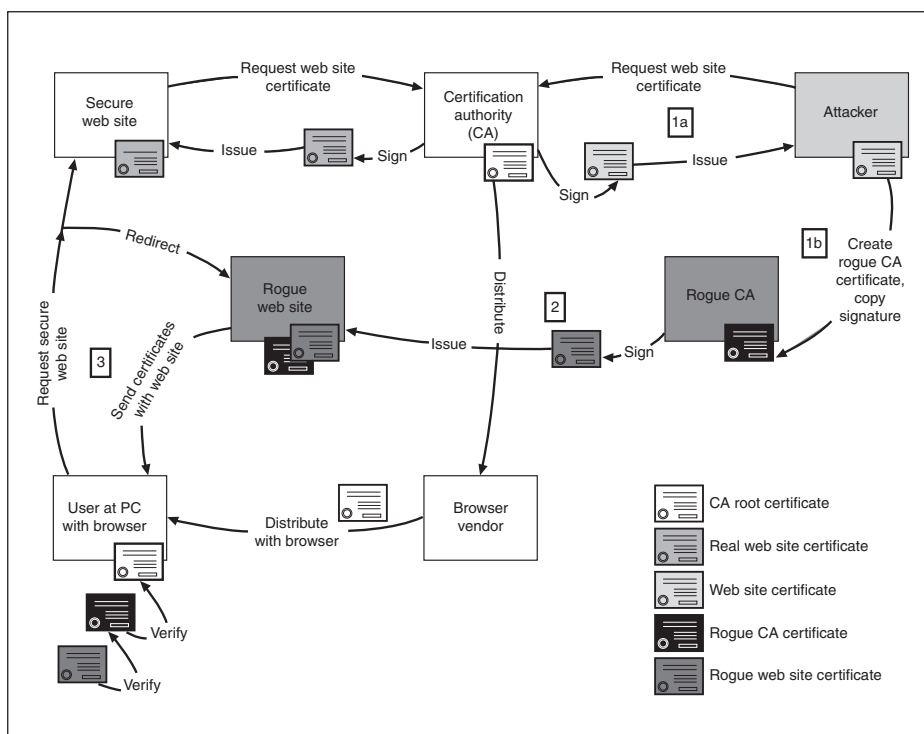


Figure 1: The basic attack scheme. Courtesy of Sotirov et al.

function that allows the construction of different messages with the same MD5 hash. This is known as an MD5 collision. Basically speaking, it is the practical application of one of the scenarios drawn by the Chinese researchers during Crypto 2004.

The real-world effect

So, we have had two scientific researches at two different points in history, plus a PoC attack. And the worst thing is that this successful PoC demonstrates that the certificate validation performed by browsers can be subverted and malicious attackers might be able to monitor or tamper with data sent to secure websites. This would be an extremely useful tool for phishers.

The internet PKI is meant to prevent those two types of attack. The research showed that known weaknesses in the MD5 hash function can be exploited in

realistic scenarios. What's worse is that even after years of warnings about the lack of security of MD5, some root CAs have still been using this broken hash function.

The reaction of the security community

Banks and other institutions were well aware of the problem, but many appear to have accepted it as an operational risk, waiting for somebody else to mitigate the problem before they took action themselves. The same is true in the vendor community.

Verisign declared that it had solved the issue. Tim Callan, vice president of VeriSign's product marketing, has said that "within hours of last week's news that researchers had uncovered a devastating weakness in secure sockets layer certificates issued by VeriSign subsidiary RapidSSL, the



Figure 2: The group of researchers who published the Proof of Concept.

company made changes to ensure all its SSL products were immune to the attacks."

But one company moving to solve the problem is not enough. All vendors are required to patch. The need to patch the problem is urgent. Both vendors and service providers could face legal liability for problems that emerged because they were using insecure certificate mechanisms.

There are also other potential side effects. For example, almost all financial institutions worldwide and their partners must be certified to the PCI standard developed by the major credit card companies. Not addressing the MD5 issue could result in PCI-related losses if companies lost their compliance with this industry standard.

"Banks and other institutions were well aware of the problem, but many appear to have accepted it as an operational risk"

The digital investigation and computer forensics community also face some challenges in the light of the latest research. The news moved across several specialised forensic practitioners' forums, and raised concerns over the possible legal impact on forensic procedures. In my view, the impact on forensic cases is very limited and only



A SUBSCRIPTION INCLUDES:

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

www.networksecuritynewsletter.com

potential, not actual. The applicability of the attack mentioned above on the digital forensic environment is, at this stage, very limited.

Some good news

Cylab is a security research centre at Carnegie Mellon University. Its technical team immediately covered the MD5 issue, and it has already developed a tool called Perspectives that is able to thwart such an attack at the browser level.

Richard Power, now research fellow at CMU Cylab, said: "Available as a Firefox plug-in, Perspectives has an estimated 30 000 users. When the browser opens an HTTPS connection (and thus establishes an SSL/TLS connection with the web server), Perspectives contacts several globally distributed notary servers, which keep a history of servers' SSL/TLS public keys. If the key history from the network notaries does not match the received server key, Perspectives warns the user about a possible attack."

"Not addressing the MD5 issue could result in PCI-related losses if companies lost their compliance with this industry standard"

Of course it is not a completely automated tool. According to Cylab's Security advisory, to allow Perspectives to detect the rogue certificate attack, the user must instruct it to contact notaries for all HTTPS sites (not only for self-signed certificates, which is the default setting), even if your browser considers the certificate valid. To do that, select Tools->Add-ons->Perspectives and then click on the 'Preferences' tab and select the option 'Contact notaries for all HTTPS sites' to enable such verification.

The Firefox plug-in works for Windows, Mac and Linux, and can be downloaded and installed from www.cs.cmu.edu/~perspectives/firefox.html#install.

Conclusion

The question is simple: we had about four years to solve the issue. Why must we always allow the horse to run from the stable before bolting the door? The reasons are many. Budget is a critical issue. Repairing this type of vulnerability requires a significant investment in time and money. Thus, many entities prefer to keep watching what happens to the others, instead of being proactive.

Availability is a second issue. Recovering from such a problem requires infrastructures to be repaired, and consequently stopped for a while. This is often unacceptable to business customers. Pushing too much could result in the usual fight between security and business people.

In any case, time has come. It is not possible to wait. MD5 is dead, and we must look to other algorithms for our salvation.

About the author

Dario Forte, CFE, CISM, former police detective and founder of DFLabs 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.

www.dflabs.com

Reference

1. Alexander Sotirov, Marc Stevens, Jacob Appelbaum, Arjen Lenstra, David Molnar, Dag Arne Osvik, Benne de Weger: "MD5 considered harmful today." December 2008 <www.win.tue.nl/hashclash/rogue-ca>

EVENTS CALENDAR

23 February 2009

1st Workshop on the Theory and Practice of Provenance

Location: San Francisco, CA, USA

Website: www.usenix.org/events/tapp09

23–24 February 2009

ASIA–Pacific CACS – Computer Audit, Control and Security (CACS) Conference

Location: Kyoto, Japan

Website: <http://bit.ly/c750w>

8–12 March 2009

SAC 2009, 24th ACM Symposium on Applied Computing

Location: Honolulu, Hawaii, USA

Website: www.dmi.unict.it/~giamp/sac/09cfp.html

11 March 2009

IDC IT Security Roadshow

Location: Athens, Greece

Website: www.idc-cema.com/?showproduct=34362

16–19 March 2009

ARES 2009 – The International Dependability Conference: Fourth International Conference on Availability, Reliability and Security

Location: Fukuoka, Japan

Website: www.ares-conference.eu/conf

17–18 March 2009

Secure World Expo

Location: Atlanta, Georgia, USA

Website: <http://secureworldexpo.com/events/index.php?id=266>

25–26 March 2009

Infosecurity, Belgium

Location: Brussels, Belgium

Website: www.infosecurity.be/sites/www_infosecurity_be/en/index.asp