NOVEMBER 2004

virus BULLETIN

The International Publication on Computer Virus Prevention, Recognition and Removal

CONTENTS

ISSN 0956-9979

IN THIS ISSUE

CHICAGO, CHICAGO

With new 'pre-conference' sessions, a new conference stream, new faces on the speaker line up and record attendance levels, VB2004 was a big success and Chicago proved to be *VB*'s kind of town. Dare we wonder whether the 'curse of *VB*' has finally been broken ...?

page 4

Virus Bulletin thanks the sponsors of VB2004:

Computer Associates

eset

Microsoft

TREND MICRO

COMPARATIVE REVIEW

With a whopping 25 products submitted for this month's comparative review, Matt Ham had his work cut out putting them all through their paces on *Windows Server 2003*. Find out which of the submissions earned their VB 100% awards.

page 12

Nov 2004
100% VIRUS BULLETIN
www.virusbtn.com

vbSpam supplement

The *VB Spam Supplement* celebrates its first anniversary this month. As usual, the supplement contains anti-spam news and events and a summary of the ASRG mailing list. Also in this edition, Matthew Prince looks at why anti-spam laws haven't worked … yet.

‘I was even quite pleased to have lasted in the AV industry long enough to be a “dinosaur”.’

Peter Morley, McAfee, UK

THE DINOSAURS LIVE ON

I was delighted to read Nick Scales’s comment ‘Definition-based AV software is dead’ last month (see *VB*, October 2004, p.2). I was even quite pleased to have lasted in the AV industry long enough to be a ‘dinosaur’. Unlike Scales, however, I believe the AV dinosaurs who surround me will not be extinct before the decade is out, and that they may last some further ten years or more.

Prevention is better than cure, and *XP Service Pack 2* has made great strides. However, I have implemented *XP Service Pack 2*, and if I run without anti-virus software, it keeps popping up and reminding me. Obviously the authors are well aware that they haven’t killed it yet, and that they have more work to do.

So, where do we go from here, to improve still further? Prevention comes from four sources.

First, the Operating System provider. Bill Gates stressed two years ago that security is a prime consideration for *Microsoft*. He associated security with *Longhorn*, his project for the next OS, and promised to concentrate on it. Since then, the initial *Longhorn* implementation has been watered down, and the diluted version is scheduled for 2005. *Microsoft* has since taken over *GeCAD*, an excellent AV vendor. I conclude that Gates believes there is still a place for the conventional AV strategy (detect after the attack) for some time to come.

The first version of *Longhorn* will have to be highly backward compatible or it will not take off. It will also, at some stage, support the new hardware security requirements about which some vendors (including *IBM*) are arguing. All this adds up to more of a delay in getting *Longhorn* bedded down.

The second source of prevention is the hardware providers. In his comment, Scales mentioned that, by 2007, anti-virus will be built into the chipsets of the latest computers and devices. He is right, and several vendors (including *McAfee*) have started providing the means to do it. However, I think it will be several years before it becomes really effective.

User policy enforcement techniques represent the third source of prevention. The implementation of these is not easy. Over the next five years, volumes and Internet usage are set to explode further. There is also the integration of both the communication and entertainment industries into the computer industry. (You doubt it? May I remind you of *Sony-MGM*.)

Last but not least, the surviving AV companies will remain responsible for excluding the known ‘nasties’ where possible. They may be replaced by a new, shining ‘*MSAV*’, but I doubt it. Some unknown nasties *will* still get through, and someone will need to respond as quickly as possible. Who, other than the surviving anti-virus vendors, will prevent them from continuing to get through?

Of the forthcoming Trojans, I am sure there will be some which get in, wait up to three months, do something horrible, and then delete themselves. This raises the classic subject of backup. Most large-scale users will have to improve their ability to retreat to a working system, and repeat the essential transactions since.

What will happen to reviewers during the next ten years? I suspect they will fade slightly, as the number of field nasties declines, but I don’t believe they will fade out completely until about 2015. Bear in mind that Chinese and other Far East users are several years behind the game, but growing very quickly, and that their reviewers expect that anything which was ever detected remains detectable.

Finally, some big AV customers have their own virus collections, and expect that anything which has ever infected or attacked them will continue to be detected.

Scales’s predictions are right, but the extinction of the AV dinosaurs will happen later. Perhaps much later.

[*See this month’s Letters page (p.10) for some different reactions to last month’s comment, ‘Definition-based AV software is dead’ – Ed.*]

Editor: Helen Martin

Technical Consultant: Matt Ham

Technical Editor: Morton Swimmer

Consulting Editors:
Nick FitzGerald, *Independent consultant, NZ*
Ian Whalley, *IBM Research, USA*
Richard Ford, *Florida Institute of Technology, USA*
Edward Wilding, *Data Genetics, UK*

NEWS

PHISHY GOINGS ON

According to *Commtouch Software* the US, UK, Brazil and Romania led the world in sending phishing emails during September 2004. But Brazil may be able to be taken out of the equation for the time being, since Brazilian federal police arrested more than 50 people in connection with phishing scams last month. Reports suggest that the majority of those arrested across four states in northern Brazil were under the age of 25 and involved in creating Trojans for use in phishing scams.

Meanwhile, *CipherTrust* revealed the findings of its analysis of its customers' email messages last month. The company says it has evidence that fewer than five zombie network operators are responsible for all Internet phishing attacks worldwide. During the first half of October 2004, researchers found that less than one per cent of email messages were phishing attacks. They also discovered that phishing attacks on the Internet were delivered each day via a different set of 1,000 zombie machines, 70 per cent of which were also used to send spam.

Finally, *Google* fixed a phishing vulnerability in its site last month, just under two days after its discovery by *Netcraft* – the vulnerability would have allowed attackers to place their own content on *Google*'s site, giving it the appearance of official content published by *Google*. The vulnerability was located in the application used to search the *Google* site itself.

DIAL A DETECTION

UK telecoms watchdog the Independent Committee for the Supervision of Standards of Telephone Information Services (ICSTIS) has issued a leaflet for consumers which provides guidance on how to deal with rogue Internet diallers ('porn diallers'). The guide explains how to detect the difference between legitimate and rogue diallers and how consumers can make a complaint to ICSTIS if they believe they have been stung by a rogue dialler.

Meanwhile, at the end of September, AV company *Sophos* was reported to be taking legal advice on its detection of the Coulomb dialler. The AV vendor suspended detection of the dialler following a complaint from its manufacturer, *Coulomb Ltd*. *VB* reported in December 2002 on the difficulties AV vendors face in making the decision of whether or not to detect porn diallers as malware (see *VB* December 2002, p.12), with German AV vendor *H+BEDV* having encountered significant legal problems in 2002 after having included detection of diallers in its product. The chief executive of *Coulomb Ltd* told *The Register* that a number of AV scanners detect the firm's dialler, and *Sophos* had simply been the first AV company it had contacted.

Prevalence Table – September 2004

Virus	Type	Incidents	Reports
Win32/Netsky	File	173,443	80.52%
Win32/Bagle	File	33,916	15.75%
Win32/Zafi	File	1,402	0.65%
Win32/Dumaru	File	970	0.45%
Win32/Mydoom	File	907	0.42%
Win32/Funlove	File	822	0.38%
Win32/Klez	File	559	0.26%
Win32/Mabutu	File	511	0.24%
Win32/Lovgate	File	338	0.16%
Win32/Valla	File	320	0.15%
Win32/Mimail	File	307	0.14%
Win32/Bugbear	File	267	0.12%
Win32/MyWife	File	230	0.11%
Win95/Spaces	File	175	0.08%
Win32/Swen	File	145	0.07%
Redlof	Script	108	0.05%
Win32/Fizzer	File	108	0.05%
Win32/Parite	File	92	0.04%
Win32/Mota	File	67	0.03%
Win32/Hybris	File	60	0.03%
Win32/Yaha	File	59	0.03%
Win32/Sobig	File	58	0.03%
Win32/Magistr	File	49	0.02%
Win32/Evaman	File	38	0.02%
Win32/Elkern	File	32	0.01%
Win32/Korgo	File	32	0.01%
Win32/BadTrans	File	30	0.01%
Win32/Kriz	File	26	0.01%
Win32/Nachi	File	24	0.01%
Laroux	Macro	23	0.01%
Win32/Plexus	File	19	0.01%
Win32/Pate	File	15	0.01%
Others[1]		252	0.12%
Total		215,404	100%

[1] The Prevalence Table includes a total of 252 reports across 50 further viruses. Readers are reminded that a complete listing is posted at http://www.virusbtn.com/Prevalence/.

CONFERENCE REPORT

ONE TOWN THAT WON'T LET YOU DOWN ...

Helen Martin

'Chicago is one town that won't let you down', sang Frank Sinatra in *My Kind of Town*. Well, it certainly didn't let *VB* down – VB2004 ran smoothly (not a hurricane, infectious disease, customs delay or terrorist attack in sight), with record delegate numbers and was described by a number of delegates as the best *Virus Bulletin* conference they had attended.

WON'T YOU PLEASE COME TO CHICAGO

The Fairmont Chicago provided a luxurious venue for the 14th *Virus Bulletin* conference. The hotel's positively palatial conference rooms were ideal for what was the largest audience the *VB* conference has ever had, with more than 330 in attendance. And such was the hotel's impeccable service that barely an eyelid was batted when the request was made for the video-taping, and later screening, of the US presidential debate. Delegates were able to sit back and enjoy Thursday's gala dinner, safe in the knowledge that they would be able to keep up to date with politics later in the evening.

BORN IN CHICAGO

VB2004 saw the *VB* conference's first ever 'pre-conference' sessions, on Wednesday afternoon.

Each of the four conference sponsors was invited to make a 50-minute presentation on a topic of their choice. Representing *Trend Micro*, David Perry kicked off the afternoon's proceedings with a look at the players and faces of the anti-virus industry. Andrew Lee followed, with a look at 'The *Eset* virus radar', and the afternoon was rounded off with sessions from Sam Curry, who presented an overview of the ever-evolving threat environment for *Computer Associates*, and *Microsoft*'s Matthew Braverman, who spoke about the role of security in the company's vision of seamless computing. All sessions were well attended and each of the sponsor representatives must be congratulated for steering well clear of marketing babble.

Also taking place on Wednesday afternoon was the 'AVIEWS Live!' discussion forum. Andrew Lee hot-footed it from his *Eset* presentation to chair the session, in which a panel of five AVIEN members each introduced a subject of interest then opened it up for discussion. Such was the popularity of the AVIEWS forum that some attendees were overheard expressing disappointment at the brevity of the *two-hour* session. Indeed, the feedback from delegates on all of the pre-conference sessions was overwhelming – you liked them and you want more!

After a gentle warm-up for the conference on Wednesday afternoon, the evening's drinks reception provided the traditional opportunity for delegates to have a couple of drinks, catch up with acquaintances and make some new ones – indeed, the organisers of the conference were delighted that this year's conference saw an influx of new faces as well as the regulars.

YOU'LL LOSE THE BLUES IN CHICAGO

Over the years, the *Virus Bulletin* conference has become well known for two things: mishaps and great entertainment. Given the former, some might say it was a brave, or even reckless, decision to engage an entertainment act for the gala dinner comprising a husband and wife team who shoot at each other with 125-pound cross bows.

But world record holders Ross and Elisa Hartzell were faultlessly professional and their astounding skills had everyone on the edge of their seats (for some reason delegates kept their distance from the stage). The jaw-dropping finale of the act involved a William Tell-style performance in which Ross Hartzell fired a single arrow to trigger a rally of shots which ended in the simultaneous impaling of apples balanced on each of the couple's heads.

As something of a relief from the tension aroused by the first act of the evening, the raucous Blooze Brothers rounded off the evening by playing into the night in true Chicago style – and even managed (eventually) to persuade a respectable number of delegates to abandon their seats for a makeshift dance floor at the front of the room.

To add (further) to the excitement of the evening, *VB* decided to hold a charity auction of special, limited edition 'VB2004 Chicago Virus Expert' baseball caps. Delegates were

invited to submit sealed bids, the idea being that the top 30 bids would each win one of the highly sought after caps. Somewhat foolishly, *VB* had overlooked the mischievous japes and capers that tend to arise as a direct result of plying delegates with alcohol – and by the end of the evening certain conference attendees had *apparently* pledged more than $30,000. Luckily, the sharp eyes of the *VB* crew members managed to sort the real from the bogus, and a total of $828 was donated to Geekcorps, a division of the International Executive Service Corps which places technical volunteers in developing nations.

SWEET HOME CHICAGO

This year's conference programme saw the first *VB* conference stream dedicated to spam. A series of four presentations relating to spam and anti-spam techniques took place on Friday morning. John Graham-Cumming, author of email sorting program POPFile, provided an overview of the trends in content trickery in spam. John Morris and Chris Lewis gave us an insight into the anti-spam infrastructure at *Nortel Networks*, and described the lessons that have been learned over the five years since its deployment. Steen Pedersen looked at the Sender Policy Framework (SPF), and Phyllis Schneck focused on the epidemiology of spam.

Also on the programme, of course, were the old regulars the corporate and technical streams. Past editor of *VB* Richard Ford, now of Florida Institute of Technology (FIT), presented Gatekeeper II, a generic virus prevention system developed by researchers at FIT. Richard's students Jason Michalske and Matt Wagner gave a live demonstration of some of the spin off tools of the system, including Gatekeeper Yo Yo – a tool which undoes all the changes made by an application – as well as Gatekeeper's viral behaviour detection capabilities.

Eric Chien introduced *Microsoft Shell*, a scripting platform currently in beta which is due to ship with *Longhorn*. After introducing the architecture and language syntax of *MS Shell*, Eric gave a series of demonstrations of *MS Shell*'s functionality, looking specifically at the functionality of which he belives worms, viruses and other miscreants are likely to take advantage.

Steve Garfink and Mary Landesman's presentation started with an AV game show, 'The Virus Price is Right', in which volunteers from the audience were asked to guess the correct answer to 'How big is Sobig?'. Of course all of the choices, ranging from $50 million to $36 billion were correct, each having been quoted by various analyst firms in the media. Steve and Mary went on to describe how a malware cost forecasting system can be used to provide more useful figures for the cost of virus attacks, on an individual organisation basis.

John Lyons provided a fascinating overview of what the UK's National Hi Tech Crime Unit is doing towards crime reduction and its intelligence regarding organised crime on the Internet, in particular phishing and DDoS attacks.

A panel discussion on malware threats to mobile devices took place on Friday afternoon – just 24 hours after the first confirmed reports of SymbOS/Cabir in the Wild. Panellists Vanja Svajcer, Mikko Hyppönen, Randy Brown, Chris Lewis and John Alexander agreed that, while we have not seen any really concerning malware for mobile devices yet, they are likely to become prime targets for malware in the near future – both in terms of malware coming from 'traditional' virus writers/script kiddies and malware written with the specific aim of collaborating with spammers.

David Perry led the closing panel session of the conference, 'What is an infection?'. David and *ICSA*'s Larry Bridwell are about to embark upon a project which, through surveys, ballots and open discussion, will attempt to define 16 AV terms over the course of one year. Panel members Jeannette Jarvis, Andrew Lee, Steve Christie, Nick FitzGerald and Richard Ford each described some of the problems they encounter with the lack of clarity in AV terminology (in their roles as customer, vendor, government representative, 'elder statesman and curmudgeon' and academic, respectively). In general, there was agreement that the lack of clarity in AV terminology is a problem, but there were few concrete suggestions as to how to solve the problem. David himself admitted that he and Larry think they 'should completely be able to fail entirely to [define the 16 terms] in one year'. Watch this space!

VB2005: THE IRISH ROVER

After three years in North America, the time has come for *VB* to visit European shores once again. VB2005 takes place 5–7 October 2005 in Dublin, Ireland. You can book your place for VB2005 now at at http://www.virusbtn.com/. Put it in your diaries and join us next year to experience the legendary craic in Dublin!

VIRUS ANALYSIS

LET THEM EAT BRIOCHE

Peter Ferrie
Symantec Security Response, USA

In 2003 I wrote: 'A virus using the manual reconstruction technique seems unlikely, since the underlying structures in *.NET* are extremely complex and contain many interdependencies' (see *VB*, April 2003, p.5). However, in 2004 we received one that did it: MSIL/Impanate.

Written by the virus writer known as 'roy g biv', a specialist in proof-of-concept viruses (most recently, the first 64-bit viruses on the *Win64* platform: W64/Rugrat on *IA64*, [see *VB*, June 2004, p.4] and W64/Shruggle on *AMD64*), Impanate is the first known parasitic, entry point obscuring appender for the *.NET* platform.

SIGN OF THE TIMES

Impanate searches in the current directory for files which do not contain a zero in the Second field of the LastWriteTime field. Impanate sets the Second field to zero in every file it examines, which serves both as an infection marker and as a means to avoid re-examining uninfectable files.

The use of the timestamp field is a speed optimization method, since it can be queried without incurring the performance penalty of opening the file. In addition, the LastWriteTime field is the only time field that is never changed when a file is copied to another location.

FILTRATION DEVICE

As with all viruses produced by this virus writer, files are infected only if they pass a strict set of filters. The conditions include that the file must be a character-mode or GUI application for the *.NET* framework, that the file is not a DLL, that the file contains no digital certificates, and that it has no bytes outside the image.

The virus avoids files that contain StrongNameSignatures or VTableFixups. StrongNameSignatures are used for digital signing of *.NET* files, so it is clear why the virus avoids files which contain them. However, it is not clear why the virus avoids VTableFixups.

The virus avoids files whose last section is writable, because the virus wants to place its code in the last section of the host, but the *.NET* framework will not allow code to execute from within a writable section.

In addition, the virus supports both 32-bit and 64-bit files, and will infect them both correctly, using a tiny piece of code trickery.

SLIPSTREAM

The virus parses the Metadata root header manually, searching for the streams that it requires. The streams are named '#~', '#Strings' and '#Blob'. The streams may appear in any order – most tools produce a constant order – but the virus will reorder them when it infects a file.

The virus is also aware of several undocumented characteristics of the *.NET* file format, including the extra data fields that can appear in the header and the flags that control the size of the stream references.

The '#~' stream contains information that is of interest to the virus. Specifically, the virus requires that the host contains 16-bit references to the '#Blob', '#GUID' and '#Strings' streams, which make the '#~' stream easier to parse, and that the host contains the following elements: TypeRefs, MemberRefs, StandAloneSigs, AssemblyRefs, Assemblies and Methods.

The virus parses the stream manually to find the TypeRefs, MemberRefs, StandAloneSigs, AssemblyRefs and Methods. The virus is not interested in the Assemblies as such, but simply requires that some are present.

SOME ASSEMBLY REQUIRED

The TypeRefs contain pointers into the '#Blob' stream of the descriptions (the types of parameters to be passed, if any, and the type of the return value, if any) of the library functions used by the host. The virus appends its own TypeRefs to those of the host and updates the references in the '#Blob' stream.

The MemberRefs contain pointers into the '#Strings' stream of the names of the library functions and properties used by the host. The virus appends its own MemberRefs to those of the host and appends the MemberRef names to the '#Strings' stream.

The StandAloneSigs contain the number and type of variables in each method. The virus chooses randomly from the StandAloneSigs of the host, duplicates one of them and appends the StandAloneSigs of the virus to it.

The AssemblyRefs contain pointers into the '#Strings' stream of the names of external assemblies that contain the functions used by the host. The virus requires two particular assemblies to be referenced in order to replicate.

The first assembly the virus requires is 'mscorlib', which is the assembly that contains many core functions, and which is roughly equivalent to 'kernel32' for *Windows* applications.

The second assembly the virus requires is 'System', which the virus uses to access the process memory, in order to copy the virus code to a local buffer, for modification prior to placing it in the host.

The virus does not alter the AssemblyRefs collection, perhaps because it would mean updating each method of the host, resulting in many changes to the file.

METHOD ACTOR

The methods contain the host code. The virus finds the first method that uses the StandAloneSigs that the virus chose earlier, and which supports the use of local variables. The virus also requires that the method contains no exception handling information. The most likely reason for this is that the process of updating the exception handling information is extremely complicated.

Having found a suitable method, the virus duplicates it, then appends the virus code to it. After the host method has run it would normally return to the caller; now, the virus will begin to execute at that time, before returning to the caller.

After appending the virus code, the virus parses it manually to update the references to the local variables and functions. The virus contains code to calculate the length of each instruction in the MSIL instruction set, and it knows which instructions need to be processed specially.

After updating the code, the virus updates the size of the last section and the host image size, and recalculates the file checksum, if required.

EXPECT THE UNEXPECTED

And so it comes to pass that, in the hands of a skilled programmer, the unlikely can become the ordinary. At least I didn't say that it could not be done because it was too difficult – anything is possible for those who have enough patience.

MSIL/Impanate

Size:	7539 bytes.
Type:	Direct action, parasitic, entry point obscuring appender.
Infects:	Microsoft *.NET* files.
Payload:	None.
Removal:	Delete infected files and restore them from backup.

FEATURE

MALWARE IN A PIG PEN – PART 2

Martin Overton
Independent Researcher, UK

In the first part of this feature (see *VB*, October 2004, p.10) I covered the use of SNORT to detect malware using simple binary and MIME strings. This part will cover more complex malware and SNORT signatures/rules to detect them. It will also cover some other parts of the SNORT rule structure (or nomenclature), as well as some of the other directives including some non-malware related rules/signatures that show how the directive/keyword is used. Finally, the article will look at some of the things that can go wrong when signatures are chosen poorly.

FLAVOURSOME PACKETS

As well as the use of signatures/rules that act on TCP packets, SNORT can act on UDP, ICMP and IP packets. Other packet types (such as IGMP) may be supported in future versions of SNORT.

THIS LITTLE PIGGY WENT TO …

The rule/signature for W32/Netsky.p described in part one of this article showed how to detect infected email coming from an external network to your internal network. But what if you want to reverse this test, or even test both directions at the same time?

You simply change the original part of the signature/rule from:

```
$EXTERNAL_NET any -> $HOME_NET any
```

which is used to detect inbound packets (from an IP not on our internal network), to:

```
$HOME_NET any -> $EXTERNAL_NET any
```

This is used to detect outbound packets (from an IP on our internal network). To reverse the direction of the test you cannot just use '<-', as this is not supported by SNORT.

Alternatively, the following is what you would use if you wanted to test data going in either direction (both inbound and outbound) with a single signature/rule:

```
$EXTERNAL_NET any <> $HOME_NET any
```

GO WITH THE FLOW

A useful keyword is the 'flow' directive. This can be used to limit rules to client or server traffic. For example:

```
alert tcp $EXTERNAL_NET 110 -> $HOME_NET any
(msg:″VIRUS Klez Incoming″;
flow:to_server,established; dsize:>120;
content:″MIME″; content:″VGhpcyBwcm9";
classtype:misc-activity; sid:1800; rev:2;)
```

This rule/signature will trigger only once a client has connected to a server (in this case a POP3 server), and will act on the data received from the server.

PIGGIN' WEB CONTENT

Let us imagine that you want SNORT to alert on web traffic that meets a specific signature. The following rule will trigger when a URL contains '/readme.eml' in any letter case (upper, lower or mixed – this is specified by the 'nocase' keyword):

```
alert tcp $HOME_NET any -> $EXTERNAL_NET $HTTP_PORTS
(msg:″WEB-MISC readme.eml download attempt″; flags:A+;
uricontent:″/readme.eml″; nocase; classtype:attempted-
user; sid:1284; reference:url,www.cert.org/advisories/
CA-2001-26.html; rev:8;)
```

This is an ideal solution for handling malware that downloads components and updates from the web, such as W32/Bagle.az@MM (see http://vil.nai.com/vil/content/v_128582.htm) or Downloader-PU (see http://vil.nai.com/vil/content/v_128464.htm).

TELL SID!

The 'sid' keyword is used as a unique identification of a specific rule/signature. However, before starting to number your own signatures you must bear in mind the following:

- Numbers <100 are reserved for future use.
- Numbers 100–1,000,000 are for use *only* for rules included with SNORT (i.e. 'official' rules).
- Numbers >1,000,000 can be used for local rules on a free-for-all basis.

MULTIPLE CONTENT

The first part of this article showed a SNORT signature/rule that contained one 'content' section (signature) to be matched against incoming data. However, SNORT signatures/rules are not limited to single 'content' sections and you can even mix content types, such as binary and text strings. For example:

```
alert udp $EXTERNAL_NET any -> $HOME_NET 1434
(msg:″W32.SQLEXP.Worm propagation (1434)″; con-
tent:″|68 2E 64 6C 6C 68 65 6C 33 32 68 6B 65 72
6E|″; content:″|04|″; offset:0; depth:1;)
```

You can even mix content types in the same content section, for example:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET 139
(msg:″NETBIOS SMB ADMIN$access″;
flow:to_server,established; content:″\\ADMIN$|00 41
3a 00|″; reference:arachnids,340;
classtype:attempted-admin; sid:532; rev:4;)
```

OBFUSCATED AND ENCRYPTED SAMPLES

Obfuscated samples

These are samples that are often packed (as many as ten different packers may be used) or that display some mild form of polymorphism, such as adding random text or other garbage to the file in order to fool MD5 and other hash functions.

Encrypted samples

In the context of this article, encrypted samples refer to the password-protected zips that have been seen in many of the Bagle variants – both those with a plain text password in the body of the email and those that use the graphic password trick to try to slow down or stop the scanner from scanning the file held in the password-protected zip.

WHEN SNORT TELLS PORKIES

This section covers some of the possible problems you may encounter when using SNORT. These are not issues with SNORT, but issues you may encounter with the signatures/rules themselves.

FALSE POSITIVES

A false positive occurs when a rule is triggered on a file that is not malicious, but is flagged as if it were. For example, a beast that grunts like a pig and eats like a pig might be flagged as a pig, but actually be a frog (*Rana grylio*, aka the pig frog).

As with anti-virus products, false positives do occur, especially when signatures are selected in haste and are not sufficiently tested. To date I have found very few false positive issues with the signatures/rules I have created. I attribute this to the level of testing I carry out before making the signatures available.

If you use the 'Flexible Response' features (flexresp) in SNORT you could end up with a self-inflicted DoS, so do be careful when using this feature.

FALSE NEGATIVES

False negatives occur when a rule is not triggered on a file that is malicious. For example a beast that grunts like a pig, eats like a pig, and *is* a pig, might be misclassified due to the fact that it is kept as a house pet (*Sus scrofa domestica,* aka the Vietnamese potbellied pig).

This is a more serious problem, as it means that the signature is flawed and misses 'real' infected files/content that should have been identified. In some cases this is difficult to resolve, especially with complex obfuscated or encrypted malware. Resolving this issue usually requires multiple signatures/rules to be created or a different approach, such as using header information rather than MIME body data.

MAIL HEADERS

Let us now look at a different way of detecting worms, not by the attachment, but by the manufactured headers.

PCRE (PERL-COMPATIBLE REGULAR EXPRESSIONS)

As mentioned, there are sometimes other ways to detect obfuscated or encrypted malware emails reliably, by looking at the manufactured mail headers they use (or do not use).

For example, both the MyDoom and Bagle families use manufactured headers which can be a reliable method of detecting them, without the need to create signatures to detect the attachment.

The following will detect MyDoom-constructed emails, even if they are corrupted, non-viable or truncated:

```
alert tcp $EXTERNAL_NET any -> any any (msg:”MyDoom
Mail Header Match/PCRE”; pcre:”/X-MIMEOLE\: Produced
By Microsoft MimeOLE V6\.00\.2600\.0000/”; pcre:”/
boundary=[“][-]{4}\=\_NextPart\_000\_\d{4}\_.{8}\
..{8}/”; pcre:”/filename=[“]\S{1,}[.](bat|scr|com|
cmd|exe|pif|zip)/”;classtype:misc-activity; rev:1;)
```

The following will reliably detect Bagle-constructed emails under the same conditions as above:

```
alert tcp $EXTERNAL_NET any -> any any (msg:”Bagle
Mail Header Match/PCRE”; pcre:”/Message-
ID\:\W{1,}[<][a-z]{19}[@]/”; pcre:”/boundary=[“][-
]{8}[a-z]{20}/”; classtype:misc-activity; rev:1;)
```

You can also use this technique to detect/block unwanted attachments in email:

```
alert tcp $EXTERNAL_NET any -> any any (msg:”Bad
Extensions Match/PCRE”; pcre:”/
attachment\;\W{1,}filename=[“]\S{1,}[.](scr|com|exe|cpl|pif|
hta|vbs)/”; classtype:misc-activity; rev:1;)
```

The signature above does not include all recommended 'bad extensions' to block, just a small subset – so feel free to add any you want to include.

The signature/rule below will usually trigger only on password-protected zip files created by email worms:

```
alert tcp $EXTERNAL_NET any -> any any
(msg:”Encrypted PKZip - SUSPECT/PCRE”;
flow:to_server,established; pcre:”/
UEsDBAoAA\S{10,}[A]{4,}/”; classtype:misc-activity;
rev:1;)
```

To date, this signature has not triggered on password-protected zips that contain samples sent to me from other researchers. However, this is still a 'test' rule and it should be used with care.

NETWORK WORMS, BLASTER, ETC.

SNORT is extremely useful for detecting, tracing and blocking many of the network worms that have become part of the background noise on the Internet, as well as the vast array of bot families and their numerous offspring.

EXPLOIT ME!

Although the SNORT maintainers no longer supply (or support) the 'virus.rules' signature set for the product, they do offer signatures that can be used to identify the use of most of the exploits upon which a reasonable percentage of worms, viruses and bots depend to allow them to auto-run when previewed in *Outlook*, or to get onto a system via a known exploit in, say, DCOM, LSASS or GDI.

```
alert tcp any any -> any 135 (msg:”DCOM Exploit
(MS03-026) targeting Windows XP SP1"; content:”|BA 26
E6 77 CC E0 FD 7F CC E0 FD 7F|”; classtype:attempted-
admin; sid:1100007; reference:URL,www.microsoft.com/
security/security_bulletins/ms03-026.asp;
reference:URL,jackhammer.org/rules/1100007; rev:1;)
```

BLOCKING INSTEAD

In the first part of this article I covered only the 'alert' directive, which will send an alert to the SNORT logs, Syslog, database or other configured storage options when a signature is matched.

There are other options for what action to take when a signature is matched. These include the ability to terminate the session, at either the originator or destination end of the conversation, or both at the same time.

The advantage of this is that you can stop an infection attempt dead in its tracks.

Below is an example:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET any
(msg:”Backdoor.GoBot.p [KAV] - SMB”; content: “|B6 B9
ED ED CD 77 5E 11 75 1B 8B BB 01 7E 05 29 54 BF 0D B6
F0 83 7B 0C 3F 44 64 EB 96 0A 8B 72|”; classtype:
misc-activity; resp:rst_all;)
```

This would terminate the connection between the source and destination IP addresses when the signature was matched. Obviously, this 'power' should be used with caution as it can cause problems with some applications, especially if there is a false positive problem with the signature itself.

However, using this feature effectively turns SNORT into a so-called IPS (Intrusion Prevention System), rather than an IDS (Intrusion Detection System).

SNIFFING WITH SNORT

SNORT is not just an IDS; it can also be used as a Sniffer – simply run SNORT in one of the following ways to achieve this:

```
Snort -v
```

This will show TCP/IP packet headers (TCP, UDP and ICMP) on the console.

```
Snort -vd
```

This will show TCP/IP packet headers and application data on the console.

```
Snort -vde
```

This will show TCP/IP packet headers, application data and data link layer headers on the console.

SNORT SIGS BOARD

Readers who would like access to the latest malware rule/signatures that I maintain for SNORT might like to create an account on my SNORT Sigs Board. This can be found at http://arachnid.homeip.net/cgi-bin/blah/Blah.pl. (Please note: those who do not supply the requested information when signing up will not be granted access.)

CONCLUSIONS

I hope that I have whet your appetite and shown that SNORT does indeed have its place in the anti-malware toolbox. This is increasingly true when we consider the merging of many technologies between the spammers, scammers, malware and hacking (cracking) communities.

This two-part article should not be considered as an exhaustive or complete look at SNORT. I have merely scratched the surface of the pig – and there is plenty more goodness under the crackling … dig in and pig out! (*Samson the pig appears courtesy of Farm Sanctuary, http://www.farmsanctuary.com.*)

LETTERS

A QUESTION OF POLICY

Is Nick Scales, author of the October 2004 comment article (see *VB*, October 2004, p.2), correct when he asserts that many organizations apply a policy rule which allows active code only when it is both signed and from a known, expected source?

Ignoring Scales's technical oversimplification (being able to prove the origin of a piece of code is not the same as being able to vouch that it is safe), I suggest that the policy he describes would absolutely require the deactivation of JavaScript both for email and for web browsing. Yet almost all organizations seem neither to enforce such a restriction, nor even to regard it as desirable.

Ironically, many of the anti-virus companies which Scales goes out of his way to diss in his article have had policy-based features in their software for years. The problem is not so much in offering these features (such as the ability to quarantine all executable email attachments), but in getting them accepted as normal business practice.

If you look beyond Scales's rather tired conspiracy-theoretical analysis, I am confident that you will find an anti-virus industry that *does* care about getting businesses to adopt safer IT practices. The problem is not so much enforcing such practices, but evangelising enough that they are written into IT policies in the first place.

Aluta continua.

Paul Ducklin
*Sophos, Australi*a

A QUESTION OF USER MENTALITY

I have been working in the AV industry for more than 16 years. Starting from year three, every now and then somebody has predicted the 'death' of known virus scanning. So, it was with little surprise (but a certain level of exasperation) that I read Nick Scales's article 'Definition-based AV software is dead' in the October 2004 issue of *Virus Bulletin*.

As have all the other such oracles, Mr Scales has completely missed the point.

At various times, various AV vendors have desperately tried to promote products, which provide protection *before* a virus has attacked. Integrity checkers, behaviour blockers, heuristic analysers, encrypted partitions, whatever. Even the great Dr Fred Cohen (who introduced the concept of computer viruses to the scientific world a couple of decades ago) tried to market such a product. Unsuccessfully, of course – the product died quickly, while scanners live and enjoy widespread use.

The reason is because generic AV products are simply too difficult for the average user to use and understand. A scanner tells the user ‘no, you do not have a virus’ or ‘yes, you have the Foo.A virus, do you want me to remove it?’. Generic products, on the other hand, tell the user ‘Process Foo is trying to write to the file Bar.exe’, ‘The file Bar.exe has been modified’, or ‘The file Foo.exe might contain a virus’.

But is it a virus that is trying to write to the file, or is it the compiler? Did the file change because it was infected, or because the *Windows Update* has changed it? Does the file contain a virus or not? What does ‘might’ mean? A generic AV product cannot answer these questions in a satisfactory way and it confuses the user (who usually has no clue what exactly is happening in their computer). Contrary to that, the report from a scanner is clear and unambiguous: you either have a virus or you don’t. (Of course, either report can be *wrong*, but that’s a different issue.) The users want certainty, not technical gibberish.

Furthermore, most generic anti-virus products have to be installed *before* a virus strikes. It doesn’t make sense to create a database of file checksums of your computer if it is already infected – it won’t help you determine which files are infected and it won’t help you clean them. And, sadly, many people reach for the AV software only *after* they suspect that their machine has become infected.

It is true that known virus scanning is the weakest line of defence against viruses – any honest person in the AV industry would admit that. But it is also the kind of defence that the users understand and are actually willing to use.

Policy enforcement that Mr Scales advocates is a great thing when it *can* be enforced. Unfortunately, a fascist policy which ensures that a virus is not able to infect the protected computers makes these computers all but unusable – little more than intelligent terminals. And, as experience shows, few people are willing to accept that.

Furthermore, policy enforcement can work only in a corporate environment. It is meaningless in the home computer environment, where the user is both owner and administrator of the machine, often without the necessary qualification for being the latter. A large number of virus infections, spam, etc. come from compromised home machines.

The reason why the AV vendors are not setting aside huge budgets for development of generic AV products is simple: we don’t like to waste money. Two decades of experience has shown us that the users simply aren’t willing to use (i.e. buy) such products. As much as we would like to see the viruses go away and see users use the strongest possible anti-virus defences (I assure you that all the honest people in my profession want just that), it is the *users* who refuse to comply.

And since, sadly, we have to eat too, we are forced to make products that have a fighting chance of being bought and used. Every company that has tried to promote a generic anti-virus defence (and there have been quite a few – how often has *VB* mentioned AV products that claimed to detect ‘all viruses – past, present and future’?) has eventually failed (or switched to providing a scanner too) and in many cases not because their product was bad, but because the users simply wouldn’t buy it.

I very much doubt that the users will change their mentality by the end of the decade. In fact, I would go on record to predict that, contrary to Mr Scales’s prediction, by 2007 the known virus scanners will still be alive and kicking. They might evolve, of course – just like nowadays we are not using the same scanning methods we used 15 years ago – but the concept itself will still be with us. Unfortunately.

As for the subscription-based model being financially successful – well, yes, it is. Which is why I expect it to be *more* widely used in the future, not less. After all, the famous *Windows Update* and *Office Update* are not very different from the updates used in AV software – a new vulnerability becomes known; a patch against it is created; the vulnerable computers are updated. I won’t be surprised if *Microsoft* starts asking this kind of service to be paid for by those who use it, just like the anti-virus companies do.

Dr Vesselin Bontchev
FRISK Software International, Iceland

BACK TO THE MACS

I was surprised to read Pete Sergeant’s response (see *VB*, October 2004, p.16) to my comment piece about Macs in the August issue of *Virus Bulletin* (see *VB*, August 2004, p.2). Not because the arguments were unfamiliar (I’ve had this debate before), but because I did not expect to find exactly the kind of reflexive defensiveness that prompted the comment piece, or a (mild, admittedly) attack on my competence as a researcher specialising in Mac virus issues.

Contrary to Mr Sergeant’s assessment of my acquaintance with OS X, there are two Macintoshes in my office, both of which currently run versions of OS X (one of them running out-of-the-box as supplied). Neither offers automatic patch updates and both default to logging in as root. (I do not need any advice on how to change this, these are test rigs.)

Mr Sergeant may, of course, argue that I must have installed OS X incorrectly – to which I can only say that if an experienced Mac support and Unix administration professional can install incorrectly by accepting installation defaults, less experienced Mac users must also be at risk of a configuration less secure than Mr Sergeant's.

However, this debate is beside the point. Being fully patched is *not* the same as being invulnerable to malware. System patches are, or should be, effective against the exploits they were intended to patch, but many malicious programs do not rely on system vulnerabilities. I would like to address some of the other fallacies raised in the letter:

- Legacy systems are not confined to the PC-using community. Mac users are notoriously reluctant to get rid of older systems that are still fit for use. There are plenty of Macs in use that are not capable of running any version of OS X, even with a memory upgrade.
- I suggested that not logging in as root would mitigate some exploits. It certainly does not confer automatic invulnerability. There are many potential exploits (viral and otherwise) that do not rely on logging in with administrator privileges.
- Automatic update functionality is available for *Me* and *W2K*, not just *XP*, and it is not difficult to schedule checks for other *98* versions. Of course, many users will not do this, but the same applies to Macs that are not configured as well as Mr Sergeant's. Of course, patches for unsupported OS versions are rather rare, but then so are new exploits, and many current threats rely on a more recent version of *Windows* to work at all.
- So that leaves users with pirated copies of *XP,* and a consequent inability to run *Windows Update,* as the primary source of secondary infection. An interesting hypothesis, but this is not supported by my data.
- There are many possible reasons as to why there is so little Mac malware: there are too few Macs around to excite the blackhats, a lack of kudos in Mac-hacking, more PC/*Windows*/*Linux*-specific blackhat resources available to skiddies, etc. To suggest that it is all down to Apple's exemplary patch management is specious. Did I mention that a high proportion of malware does not rely on system vulnerabilities?
- I have probably been too critical of *Microsoft* in the past ever to hope to work there, but the company is in a no-win situation on patch management. If it tests comprehensively, it's too slow. If a single system exhibits patch problems, BugTraq and *The Register* are all over it. What have zero day exploits to do with it?

I agree with one assertion: OS X *is* a very capable, powerful and user-friendly OS. But it is not invulnerable.

David Harley
Independent writer and researcher, UK

COMPARATIVE REVIEW

WINDOWS SERVER 2003

Matt Ham

Windows 2003 Server is now an environment which can be considered mature. Furthermore, there were no major problems encountered during the last comparative review to be carried out on this platform (see *VB*, November 2003, p.13). With these factors in mind I had my hopes set on what might be a more relaxing review period than usual. Sadly, however, my hopes were dashed by the arrival of the test sets.

THE TEST SETS

The test sets were based on the most recent version of the (RealTime) WildList available on 6 October 2004, the deadline for product submission having been 8 October. However, three months had passed since the last comparative review (and the last maintenance of the test set), and that was sufficient time for close to 90 new worms to have been added to the In the Wild (ItW) category. The preponderance of all-but-identical worms in, for example, the W32/Sdbot, W32/Rbot, W32/Agobot and W32/Korgo families made replication a particularly mind-numbing process.

These additions are sufficiently irritating to name that the WildList Organization has taken to using checksum values to describe versions. They also add very little, if anything, to the difficulty of their detection. Although many are packaged in layer upon layer of obfuscating archive, the files themselves are easily recognisable. With this in mind, a bumper crop of VB 100% awards was expected.

Alwil avast! 4.5.286

ItW Overall	100.00%	**Macro**	99.56%
ItW Overall (o/a)	100.00%	**Standard**	99.36%
ItW File	100.00%	**Polymorphic**	93.58%

The review of *avast!* began with a sinking feeling, since the on-access scanner refused to load. This turned out to be a result of it starting as a service under the local administrator account. *Windows* refused to allow this to happen since the default image used for *Windows 2003* testing has no password. This was easily remedied by changing to the system account. The problem can be discounted as an issue in the real world – except for administrators who have no passwords on their servers. Such folk, however, are likely to

Nov 2004
100%
VIRUS BULLETIN
www.virusbtn.com

On-access tests	ItW File		ItW Boot		ItW Overall	Macro		Polymorphic		Standard	
	Number missed	%	Number missed	%	%	Number missed	%	Number missed	%	Number missed	%
Alwil avast!	0	100.00%	0	100.00%	100.00%	18	99.56%	112	93.58%	18	99.17%
Authentium Command	0	100.00%	0	100.00%	100.00%	0	100.00%	1	99.95%	5	99.58%
BLC Win Cleaner	0	100.00%	0	100.00%	100.00%	87	97.92%	1087	92.85%	506	71.49%
CA eTrust Antivirus (InoculateIT)	0	100.00%	0	100.00%	100.00%	3	99.93%	2	99.78%	3	99.69%
CA eTrust Antivirus (Vet)	1	99.73%	0	100.00%	99.73%	12	99.82%	2	99.87%	5	99.60%
CA Vet Anti-Virus	1	99.73%	0	100.00%	99.73%	0	100.00%	2	99.87%	5	99.60%
CAT Quick Heal	0	100.00%	0	100.00%	100.00%	87	97.92%	1087	92.85%	506	71.49%
DrWeb DrWeb	2	99.45%	0	100.00%	99.46%	0	100.00%	0	100.00%	3	99.69%
Eset NOD32	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	1	99.82%
Fortinet FortiClient	0	100.00%	0	100.00%	100.00%	201	95.52%	5658	61.51%	86	97.58%
FRISK F-Prot Antivirus	0	100.00%	0	100.00%	100.00%	0	100.00%	1	99.95%	4	99.60%
F-Secure Anti-Virus	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	7	99.49%
GDATA AntiVirusKit	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	0	100.00%
Grisoft AVG	0	100.00%	0	100.00%	100.00%	23	99.44%	757	83.64%	34	98.17%
H+BEDV AntiVir	0	100.00%	0	100.00%	100.00%	6	99.84%	271	98.74%	24	98.93%
Hauri ViRobot	0	100.00%	0	100.00%	100.00%	8	99.80%	8	99.69%	10	99.54%
Kaspersky KAV	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	14	99.51%
McAfee VirusScan	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	3	99.79%
MicroWorld eScan	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	0	100.00%
Norman Virus Control	0	100.00%	0	100.00%	100.00%	2	99.95%	181	91.03%	11	99.63%
SOFTWIN BitDefender	0	100.00%	0	100.00%	100.00%	9	99.78%	6	99.73%	23	99.05%
Sophos Anti-Virus	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	15	99.30%
Symantec SAV	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	0	100.00%
Trend ServerProtect	0	100.00%	0	100.00%	100.00%	0	100.00%	215	95.77%	8	99.66%
UNA UNA	69	89.10%	4	0.00%	88.14%	1993	52.72%	14267	20.14%	584	72.98%
VirusBuster VirusBuster	0	100.00%	0	100.00%	100.00%	0	100.00%	101	91.45%	16	99.17%

find this the least of their problems. A more concerning issue was observed when the on-access scanner claimed to have crashed while scanning the test sets. However, the failure appeared to have been non-critical since the remainder of the test set was scanned with no problems. With minor glitches as the only moments of note, it will come as no surprise that *avast!* receives a VB 100 % award on this occasion.

Authentium Command AntiVirus 4.92.1

ItW Overall	100.00%	**Macro**	100.00%
ItW Overall (o/a)	100.00%	**Standard**	99.72%
ItW File	100.00%	**Polymorphic**	99.95%

There is far less to comment upon where *Command AntiVirus* is concerned. All misses are those which will be

On-demand tests	ItW File		ItW Boot		ItW Overall	Macro		Polymorphic		Standard	
	Number missed	%	Number missed	%	%	Number missed	%	Number missed	%	Number missed	%
Alwil avast!	0	100.00%	0	100.00%	100.00%	18	99.56%	112	93.58%	15	99.36%
Authentium Command	0	100.00%	0	100.00%	100.00%	0	100.00%	1	99.95%	2	99.72%
BLC Win Cleaner	0	100.00%	0	100.00%	100.00%	80	98.05%	1087	92.85%	169	92.91%
CA eTrust Antivirus (InoculateIT)	0	100.00%	0	100.00%	100.00%	3	99.93%	0	100.00%	1	99.82%
CA eTrust Antivirus (Vet)	1	99.73%	0	100.00%	99.73%	13	99.78%	2	99.87%	3	99.72%
CA Vet Anti-Virus	1	99.73%	0	100.00%	99.73%	0	100.00%	2	99.87%	3	99.72%
CAT Quick Heal	0	100.00%	0	100.00%	100.00%	80	98.05%	1087	92.85%	506	71.49%
DrWeb DrWeb	1	99.73%	0	100.00%	99.73%	0	100.00%	0	100.00%	0	100.00%
Eset NOD32	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	1	99.82%
Fortinet FortiClient	0	100.00%	0	100.00%	100.00%	201	95.52%	5658	61.51%	86	97.58%
FRISK F-Prot Antivirus	0	100.00%	0	100.00%	100.00%	0	100.00%	1	99.95%	2	99.72%
F-Secure Anti-Virus	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	0	100.00%
GDATA AntiVirusKit	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	0	100.00%
Grisoft AVG	0	100.00%	0	100.00%	100.00%	20	99.51%	257	85.97%	26	98.74%
H+BEDV AntiVir	0	100.00%	0	100.00%	100.00%	6	99.84%	271	98.74%	24	98.87%
Hauri ViRobot	0	100.00%	0	100.00%	100.00%	8	99.80%	4	99.78%	14	99.17%
Kaspersky KAV	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	0	100.00%
McAfee VirusScan	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	3	99.79%
MicroWorld eScan	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	0	100.00%
Norman Virus Control	1	99.73%	0	100.00%	99.73%	2	99.95%	180	91.24%	5	99.69%
SOFTWIN BitDefender	0	100.00%	0	100.00%	100.00%	9	99.78%	6	99.73%	22	99.23%
Sophos Anti-Virus	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	15	99.30%
Symantec SAV	0	100.00%	0	100.00%	100.00%	0	100.00%	0	100.00%	0	100.00%
Trend ServerProtect	0	100.00%	0	100.00%	100.00%	0	100.00%	215	95.77%	9	99.72%
UNA UNA	64	89.62%	4	0.00%	88.65%	1712	58.90%	14246	21.08%	537	75.21%
VirusBuster VirusBuster	0	100.00%	0	100.00%	100.00%	0	100.00%	102	91.45%	13	99.31%

painfully familiar to regular readers of the *VB* comparatives – a selection of samples missed entirely as a result of choices made by the developer based on product efficiency, rather than the product being unable to detect them. One problem that did occur here, however, was in the production of logs, since the original rtf log was mysteriously truncated. Results were therefore obtained by deletion. The award of a VB 100% duly followed.

Nov 2004
100%
VIRUS BULLETIN
www.virusbtn.com

BLC Win Cleaner 7.02

ItW Overall	100.00%	**Macro**	98.05%
ItW Overall (o/a)	100.00%	**Standard**	92.91%
ItW File	100.00%	**Polymorphic**	92.85%

Business Logic Corporation is a name that is new to the VB 100% testing roll call, though its pedigree is instantly recognisable when installed. The product is both functionally and, in all but a few strategically placed logos, visually identical to the *CAT* product from which *Win Cleaner* (*WC*) is derived. Despite its rather unfortunate acronym, *WC* denied any opportunity for jokes at its expense by detecting all viruses in the ItW test set. With no false positives, *Win Cleaner* earns a VB 100% award on its first appearance.

CA eTrust Antivirus 7.1.192

ItW Overall	99.73%	**Macro**	99.78%
ItW Overall (o/a)	99.73%	**Standard**	99.72%
ItW File	99.73%	**Polymorphic**	99.87%

It has been noted on several occasions that *eTrust* can operate with either the *InoculateIt* or *Vet* engines, both being supplied in a standard installation. On this occasion both engines were tested, with the intention of comparing their performance (see box). Currently the default installation is the *Vet* engine, which missed one of the W32/Agobot samples in the ItW test set. This was enough to deny *eTrust* a VB 100% award when used with the *Vet* engine. The logging facility of the product in either incarnation remains an affront to sanity, there being no real means to obtain logs which are readable to either machine or human.

CA Vet Anti-Virus 10.64.0

ItW Overall	99.73%	**Macro**	100.00%
ItW Overall (o/a)	99.73%	**Standard**	99.72%
ItW File	99.73%	**Polymorphic**	99.87%

This offering from *CA* contains the same engine as the previous offering, yet has a very different interface. Scanning here was in most cases slightly slower than the product's *eTrust* counterpart – except on the zipped OLE files, where the *Vet* product was considerably speedier. With the same engine inside the product, it should come as no surprise that the scanning results were the same for both *Vet*-based products and, of course, the miss of the W32/Agobot sample in the ItW test set denies *Vet* a VB 100% on this occasion.

CAT Quick Heal 7.02

ItW Overall	100.00%	**Macro**	98.05%
ItW Overall (o/a)	100.00%	**Standard**	71.49%
ItW File	100.00%	**Polymorphic**	92.85%

CA eTrust Antivirus 7.1.192 (InoculateIT engine)

ItW Overall	100.00%	**Macro**	99.93%
ItW Overall (o/a)	100.00%	**Standard**	99.82%
ItW File	100.00%	**Polymorphic**	100.00%

As noted in the text, the *eTrust* product can operate with either the *InoculateIt* or *Vet* engines, and on this occasion both engines were tested, with the intention of comparing their performance.

The *InoculateIT* engine, which is not currently the standard installation, performed much as expected. This included missing samples of W97M/Pain.A (a strange miss considering the otherwise full detection of macro viruses). Despite this, detection was, in general, very good and a VB 100% award would be obtained easily with the product using the *InoculateIT* engine.

As far as scanning speed is concerned, *eTrust* is marginally faster when using the *Vet* engine than when using the *InoculateIT* engine.

One other item of note was observed while testing: it seems to be possible to operate *eTrust* with one engine operating on demand and the other operating on access.

With a derived product (*Win Cleaner*) having already obtained a VB 100% in this comparative, it will come as no shock to learn that *Quick Heal* also earns a VB 100% award this month. Strangely, despite an otherwise identical performance, the *CAT* product was slightly slower than the *Business Logic* version.

Doctor Web Dr.Web 4.32a

ItW Overall	99.73%	**Macro**	100.00%
ItW Overall (o/a)	99.46%	**Standard**	100.00%
ItW File	99.73%	**Polymorphic**	100.00%

Dr.Web is now produced by Russian company *Doctor Web* rather than *DialogueScience*. Uncharacteristically, *Dr.Web* missed a sample of W32/Flopcopy, a sample located in the ItW test set, and was thus prevented from earning a VB 100% award. The slightly better news was that the product generated no false suspicious files, which has not been the case for a while.

Eset NOD32 1.889

ItW Overall	100.00%	**Macro**	100.00%
ItW Overall (o/a)	100.00%	**Standard**	99.82%
ItW File	100.00%	**Polymorphic**	100.00%

Hard Disk Scan Rate	Executables			OLE Files			Zipped Executables		Zipped OLE Files	
	Time (s)	Throughput (kB/s)	FPs [susp]	Time(s)	Throughput (kB/s)	FPs [susp]	Time (s)	Throughput (kB/s)	Time(s)	Throughput (kB/s)
Alwil avast!	99	5524.6		12	6611.1		21	7591.3	9	8289.7
Authentium Command	114	4797.7		5	15866.8		50	3188.3	5	14921.5
BLC Win Cleaner	62	8821.5		15	5288.9		48	3321.2	18	4144.9
CA eTrust Antivirus (InoculateiT)	132	4143.4		4	19833.4		58	2748.6	9	8289.7
CA eTrust Antivirus (Vet)	141	3879.0		4	19833.4		66	2415.4	11	6782.5
CA Vet Anti-Virus	144	3798.1		6	13222.3		68	2344.4	3	24869.2
CAT Quick Heal	72	7596.3		15	5288.9		50	3188.3	25	2984.3
DrWeb DrWeb	194	2819.2		15	5288.9		63	2530.4	12	6217.3
Eset NOD32	49	11161.9		7	11333.4		29	5497.1	8	9325.9
Fortinet FortiClient	77	7103.0	1	12	6611.1		21	7591.3	13	5739.0
FRISK F-Prot Antivirus	139	3934.8		5	15866.8		55	2898.5	4	18651.9
F-Secure Anti-Virus	129	4239.8		15	5288.9		86	1853.7	23	3243.8
GDATA AntiVirusKit	672	813.9	[1]	18	4407.4		305	522.7	20	3730.4
Grisoft AVG	145	3771.9		8	9916.7		59	2702.0	9	8289.7
H+BEDV AntiVir	420	1302.2		14	5666.7		210	759.1	17	4388.7
Hauri ViRobot	536	1020.4	20 [2]	14	5666.7		-	-	28	2664.6
Kaspersky KAV	164	3335.0		15	5288.9		77	2070.3	18	4144.9
McAfee VirusScan	93	5881.0		8	9916.7		64	2490.9	15	4973.8
MicroWorld eScan	286	1912.4		23	3449.3		115	1386.2	46	1621.9
Norman Virus Control	345	1585.3		5	15866.8		144	1107.1	6	12434.6
SOFTWIN BitDefender	591	925.4	[1]	8	9916.7		242	658.7	8	9325.9
Sophos Anti-Virus	56	9766.6		10	7933.4		42	3795.6	11	6782.5
Symantec SAV	149	3670.7		21	3777.8		69	2310.4	21	3552.7
Trend ServerProtect	74	7391.0		8	9916.7		32	4981.8	10	7460.7
UNA UNA	80	6836.7	2 [1]	19	4175.5		110	1449.2	34	2194.3
VirusBuster VirusBuster	185	2956.4		7	11333.4		120	1328.5	14	5329.1

Coming very close to full detection of all samples in all test sets, *NOD32* continues to be entitled to quote its unblemished record of ItW detection on its marketing materials. If a failure in this area does ever occur, I am sure that the printers of *Eset*'s marketing materials will be as happy as *Eset* will be sad. With no incidents of note during testing, I can only congratulate Eset on another VB 100% award.

Nov 2004 100% VIRUS BULLETIN www.virusbtn.com

Fortinet FortiClient 1.2.130

ItW Overall	100.00%	**Macro**	95.52%
ItW Overall (o/a)	100.00%	**Standard**	97.58%
ItW File	100.00%	**Polymorphic**	61.51%

This product's detection has improved in leaps and bounds since first being submitted for *VB* testing a few months ago.

On this occasion all samples from the ItW test set were detected, the only real weaknesses in detection lying in the polymorphic test set. However, the improvement in detection has not come without a further false positive, which is sufficient grounds to deny *FortiClient* a VB 100% by the narrowest of margins. One suspects that it is merely a matter of when, rather than if, this situation will change for the better.

FRISK F-Prot Antivirus 3.15 b

ItW Overall	100.00%	**Macro**	100.00%
ItW Overall (o/a)	100.00%	**Standard**	99.72%
ItW File	100.00%	**Polymorphic**	99.95%

A product with a far longer history behind it, logical readers will have already been able to guess much about *F-Prot*'s performance from the performance of *Command* earlier in the testing. Indeed, like *Command*, *FRISK*'s product is eligible for another VB 100% award.

F-Secure Anti-Virus 5.50

ItW Overall	100.00%	**Macro**	100.00%
ItW Overall (o/a)	100.00%	**Standard**	100.00%
ItW File	100.00%	**Polymorphic**	100.00%

Since this product also makes use of *FRISK*-derived detection, the fate of *F-Secure Anti-Virus* is also fairly easy to predict – full detection and no false positives mean that the product earns a VB 100% award.

GDATA AntiVirusKit 14.0.8

ItW Overall	100.00%	**Macro**	100.00%
ItW Overall (o/a)	100.00%	**Standard**	100.00%
ItW File	100.00%	**Polymorphic**	100.00%

Another example of repackaged engines, *AVK* is one of the older players in this area. One concern about the use of two engines might be an increase in the likelihood of false positives. On this occasion the product did alert on a clean file, although it was identified only as suspicious, rather than being a full blown false positive. The combination of *BitDefender* and *Kaspersky* engines in *AVK* seems a good choice; on this occasion all samples in all test sets were detected and *AVK* receives a VB 100% award.

Grisoft AVG 7.0.275

ItW Overall	100.00%	**Macro**	99.51%
ItW Overall (o/a)	100.00%	**Standard**	98.74%
ItW File	100.00%	**Polymorphic**	85.97%

Returning to products which are tested in only one incarnation, *Grisoft*'s *AVG* is the next in line. Misses here were, as ever, in the more complex variety of polymorphic virus. These polymorphics do tend, however, to be restricted to zoo collections rather than breaking into the wild. This does not, therefore, make a dent in the product's ItW detection rate. False positives were the cause of a temporary glitch in *AVG*'s performance a few months ago, but this seems very much consigned to history now. As a result, a VB 100% award wings its way towards *Grisoft*.

H+BEDV AntiVir Windows Server 6.28.0.101

ItW Overall	100.00%	**Macro**	99.84%
ItW Overall (o/a)	100.00%	**Standard**	98.87%
ItW File	100.00%	**Polymorphic**	98.74%

AntiVir's GUI is distinctive, in that is seems to have been designed for server use at the expense, in certain aspects, of user-friendliness. Since scheduled scans are stressed, which can be run in the background, there is little in the way of immediate user feedback on scans, for example. When scanning the clean test set, several files were flagged as 'possibly destroyed by a virus' but not considered to be suspicious or infected in any way. Detection was full in the ItW test sets, thus earning a VB 100% award for *H+BEDV*.

Hauri ViRobot Advanced Server

ItW Overall	100.00%	**Macro**	99.80%
ItW Overall (o/a)	100.00%	**Standard**	99.17%
ItW File	100.00%	**Polymorphic**	99.78%

Hauri's detection has been improving over recent tests and this occasion was no different. On the other side of the equation, however, the new detection has come at a cost. A full 20 false positives were noted along with two suspicious files in the clean test set. Most of these were for HLLC.Fataller, a name I have heard far more often during false positive testing than on any other occasion. Given that one false positive is causing much of the problem, it seems likely that this issue will be resolved soon, but in the meantime *ViRobot* is denied a VB 100% award.

Kaspersky KAV 4.5.0.97

ItW Overall	100.00%	**Macro**	100.00%
ItW Overall (o/a)	100.00%	**Standard**	100.00%
ItW File	100.00%	**Polymorphic**	100.00%

Kaspersky AntiVirus (*KAV*) is one of those products where different names for different components are the order of the day. The version number given above is that for the main scanner – other components all being in the 4.5.0.9x region. *KAV* continues to behave smoothly and without any other cause for major comment. With full detection In the Wild and no false positives a VB 100% is awarded to the *Kaspersky* product.

McAfee VirusScan Enterprise 8.0.0 4396

ItW Overall	100.00%	**Macro**	100.00%
ItW Overall (o/a)	100.00%	**Standard**	99.79%
ItW File	100.00%	**Polymorphic**	100.00%

McAfee's product is another which has been reviewed a sufficient number of times for no surprises to be expected. Indeed, all requirements for a VB 100% award were reached without problems. However, strange matters arrived to pique the interest somewhat. In this case it was the log file which perplexed, since all viruses detected seemed to have been detected twice.

MicroWorld eScan 2003

ItW Overall	100.00%	**Macro**	100.00%
ItW Overall (o/a)	100.00%	**Standard**	100.00%
ItW File	100.00%	**Polymorphic**	100.00%

This product is a rebadge of *AVK*, thus like *AVK* it is derived from *BitDefender* and *Kaspersky* engines. *eScan*'s results rarely differ extensively from those expected as a result of its ancestry. Scanning here was notably faster on executables than *AVK*'s scanning speed, though the difference was reversed on OLE files. It was also notable that *AVK*'s declaration of a suspicious file was not mirrored here, suggesting that tweaks have been made behind the scenes. The differences were not, however, continued into the area of detection.

With full detection of all files and no false positives a VB 100% is a sure result for *MicroWorld*.

Norman Virus Control 5.70

ItW Overall	99.73%	**Macro**	99.95%
ItW Overall (o/a)	100.00%	**Standard**	99.69%
ItW File	99.73%	**Polymorphic**	91.24%

Norman Virus Control is another of those products which usually presents no problems at all, though on this occasion it elicited at least one surprise. Unfortunately this was not a particularly pleasant one for the developers, since it was a miss of BAT/Mumu. This is a particularly surprising miss, considering its relative age and its location in the ItW test set. This, of course, prevents *Norman* from obtaining a VB 100% award.

SOFTWIN BitDefender 8 Professional Plus

ItW Overall	100.00%	**Macro**	99.78%
ItW Overall (o/a)	100.00%	**Standard**	99.23%
ItW File	100.00%	**Polymorphic**	99.73%

One of the components of *AVK*, it came as no surprise that *BitDefender* declared a suspicious file in exactly the same location as that product – though, again, this was not one serious enough to negate the possibility of a VB 100% award. *BitDefender* did miss slightly more samples than its hybrid offspring, but none of these were likely to become an issue In the Wild. Not unexpectedly, a VB 100% was earned for this combination of detection and lack of false detection.

Sophos Anti-Virus 3.83

ItW Overall	100.00%	**Macro**	100.00%
ItW Overall (o/a)	100.00%	**Standard**	99.30%
ItW File	100.00%	**Polymorphic**	100.00%

There was much rejoicing when reviewing *Sophos Anti-Virus* on this occasion, since the perennially irritating log format seems at last to have been brought up to date – simplifying log parsing immensely.

Sophos's detection rate is approaching full in all categories too. With no problems with regard to detection or false positives, *SAV* obtains a VB 100% award – and I regard the product with somewhat less antipathy.

Symantec SAV 9.0/0.338

ItW Overall	100.00%	Macro	100.00%
ItW Overall (o/a)	100.00%	Standard	100.00%
ItW File	100.00%	Polymorphic	100.00%

Symantec's *SAV* continues to be one of the slower scanners when faced with infected files, the volume of its log files potentially bearing some responsibility for this. With each virus report occupying an average of 230 bytes, test sets numbering several tens of thousands of samples tend to imply vast log files. Such logs were sufficient, in fact, to crash the client when scanning had completed. Despite this (admittedly lab-specific) behaviour, the rates of detection and lack of false positives for *SAV* remained at their usual high levels such that a VB 100% award is appropriate.

Trend ServerProtect 5.58(1060)

ItW Overall	100.00%	Macro	100.00%
ItW Overall (o/a)	100.00%	Standard	99.72%
ItW File	100.00%	Polymorphic	95.77%

Trend Micro's product is among the more complex to install, since it is inextricable from its management software. A certain degree of fiddling is required to install update files and it tends to lose, among administrative options, the simple commands required during comparative testing. However, these complaints would be irrelevant to a real-world administrator who might be expected to concentrate on the grand scale as well as detection and false positives. Detection remained slightly under par on polymorphic samples, though full in the ItW test set, justifying the award of a VB 100% to *Trend*.

UNA UNA 1.83 Kernel 255

ItW Overall	88.65%	Macro	58.90%
ItW Overall (o/a)	88.14%	Standard	75.21%
ItW File	89.62%	Polymorphic	21.08%

Still a relative newcomer to the *VB* tests, the *UNA* product seems to have improved markedly in its ease of testing – though this may simply be a function of extra practice. False positive rates have certainly become less of a problem and new detections have been added in the test sets. Though there is still a considerable way to go until the product will achieve a VB 100% award, *UNA*'s developers have shown that this might be possible in time.

VirusBuster VirusBuster 4.7 build 18

ItW Overall	100.00%	Macro	100.00%
ItW Overall (o/a)	100.00%	Standard	99.31%
ItW File	100.00%	Polymorphic	91.45%

Last, but in the way of time-honoured cliches, by no means least, *VirusBuster*'s product leaves me scraping the barrel for worthwhile comments once more. At times such as these it pays to be called Aardvark Antivirus for sure. *VirusBuster* easily qualifies for a VB 100% award, with no false positives generated and misses being noticeable only among the more complex polymorphic samples.

CONCLUSION

The theory that the new worm samples included in the test sets would cause few problems turned out, by and large, to be correct – though there were a few surprising exceptions for usually steadfast products. In many similar cases in the past this has turned out to be due to the developers having a sample which they believe to be In the Wild and which their product can detect, while in fact a different sample is generally considered to be In the Wild. Whether this is the case here remains to be seen.

The lack of stability issues in *Windows 2003* that was seen in last November's comparative followed through on this occasion. *Microsoft* has been working ever more closely with anti-virus developers over the last few years and this could well be the reason behind the added stability. Platform stability certainly simplifies the matter of testing and can hardly be a bad thing as far as the real world is concerned either. The optimist in me dares to hope that this will be the case ever more as new operating systems are created, though the pessimist still tells me that major unforeseen disasters will be in store.

Technical details

Test environment: Identical 1.6 GHz Intel Pentium machines with 512 MB RAM, 20 GB dual hard disks, DVD/CD-Rom and 3.5-inch floppy drive running Windows Server 2003 Web Edition V5.2 Build 3790.

Virus test sets: Complete listings of the test sets used are at http://www.virusbtn.com/Comparatives/Win2K/2004/test_sets.html.

A complete description of the results calculation protocol is at http://www.virusbtn.com/Comparatives/Win95/199801/protocol.html.

END NOTES & NEWS

The 31st Annual Computer Security Conference and Expo takes place 8–10 November 2004 in Washington, D.C., USA. 14 tracks will cover topics including wireless, management, forensics, attacks and countermeasures, compliance and privacy and advanced technology. For details see http://www.gocsi.com/.

The ISACA Network Security Conference will be held 15–17 November 2004 in Budapest, Hungary. Presentations will discuss the technologies, and the best practices in designing, deploying, operating and auditing them. See http://www.isaca.org/.

The seventh Association of Anti-Virus Asia Researchers International conference (AVAR2004) will be held 25–26 November 2004 at the Sheraton Grande Tokyo Bay hotel in Tokyo, Japan. The theme for the conference is 'Eutaxy or chaos? Network security: present and future'. For details see http://www.aavar.org/.

Infosec USA will be held 7–9 December 2004 in New York, NY, USA. For details see http://www.infosecurityevent.com/.

The SANS Cyber Defensive Initiative East takes place 7–14 December 2004 in Washington, D.C., USA. The event offers both extended bootcamp and hands-on sessions. Focused training disciplines include security, legal, operations, managerial and audit. For more information see http://www.sans.org/.

Computer & Internet Crime 2005 will take place 24–25 January 2005 in London, UK. The conference and exhibition are dedicated solely to the problem of cyber crime and the associated threat to business, government and government agencies, public services and individuals. For more details and online registration see http://www.cic-exhibition.com/.

The 14th annual RSA Conference will be held 14–19 February 2005 at the Moscone Center in San Francisco, CA, USA. For more information, including online registration and the conference agenda, see http://www.rsaconference.com/.

The E-crime and Computer Evidence conference ECCE 2005 takes place at the Columbus Hotel in Monaco from 29–30 March 2005. ECCE 2005 will consider aspects of digital evidence in all types of criminal activity, including timelines, methods of evidence deposition, use of computers for court presentation, system vulnerabilities, crime prevention etc. A reduced daily registration rate of 150 euros per delegate applies until 21 November 2004. For more details see http://www.ecce-conference.com/.

The first Information Security Practice and Experience Conference (ISPEC 2005) will be held 11–14 April 2005 in Singapore. ISPEC is intended to bring together researchers and practitioners to provide a confluence of new information security technologies, their applications and their integration with IT systems in various vertical sectors. For more information see http://ispec2005.i2r.a-star.edu.sg/.

The 14th EICAR conference will take place from 30 April to 3 May 2005 in Saint Julians, Malta. Authors are invited to submit papers for the conference. The deadlines for submissions are as follows: non-academic papers 26 November 2004; academic papers 14 January 2005; poster presentations 18 February 2005. For full details of the conference see http://conference.eicar.org/.

The sixth National Information Security Conference (NISC 6) will be held 18–20 May 2005 at the St Andrews Bay Golf Resort and Spa, Scotland. For details of the agenda (which includes a complimentary round of golf at the close of the conference) or to register online, see http://www.nisc.org.uk/.

The third International Workshop on Security in Information Systems, WOSIS-2005, takes place 24–25 May 2005 in Miami, USA. For full details see http://www.iceis.org/.

NetSec 2005 will be held 13–15 June 2005 in Scottsdale AZ, USA. The program covers a broad array of topics, including awareness, privacy, policies, wireless security, VPNs, remote access, Internet security and more. See http://www.gocsi.com/events/netsec.jhtml.

The 15th Virus Bulletin International Conference, VB2005, will take place 5–7 October 2005 in Dublin, Ireland. For registration and more details see http://www.virusbtn.com/.

SUBSCRIPTION RATES

Subscription price for 1 year (12 issues) including first-class/airmail delivery: £195 (US$310)

Editorial enquiries, subscription enquiries, orders and payments:

Virus Bulletin Ltd, The Pentagon, Abingdon Science Park,
Abingdon, Oxfordshire OX14 3YP, England

Tel: +44 (0)1235 555139 Fax: +44 (0)1235 531889
Email: editorial@virusbtn.com www.virusbtn.com

No responsibility is assumed by the Publisher for any injury and/or damage 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.

vb Spam supplement

CONTENTS

NEWS & EVENTS

SPAM GETS THE SNIFFLES

Not only has the shortage of flu vaccine been something of a political hot potato in the run up to the US Presidential elections, but now spammers have seized the opportunity to cash in on the problem. Canadian email security firm *Vircom* says it has seen a surge of emails touting flu vaccinations, generally at highly inflated prices. It looks like Viagra may have to take a back seat until the flu season has passed.

STORMS PUT THE WIND UP SPAMMERS

Email security firm *FrontBridge Technologies Inc.* reported a significant decline in the volume of spam messages seen in the days immediately following the three hurricanes that hit the South East coast of the US in September.

Hurricanes Frances, Ivan and Jeanne each had the same apparent effect on spam levels, with the day the hurricane hit and the day following the hurricane showing the most significant dips in spam traffic. Hurricane Frances appeared to be the cause of spam volumes falling from 89 per cent to below 82 per cent; Hurricane Ivan saw spam fall from 91 per cent to less than 84 per cent; and Hurricane Jeanne saw the number of spam messages drop from 89 per cent to 83 per cent. However, despite the wobbles of hurricane season, *FrontBridge* also recorded a record peak in spam volumes of 91 per cent in September, with a monthly average of 85 per cent – an increase of three per cent from August.

SPAM BECOMES A COLLECTORS' ITEM

Just in case you hadn't already seen enough spam in your inbox, or in case your spam filter is so efficient that you find yourself missing the stuff, a British man has set up his own Museum of Spam. Stephen Newton considers spam to be 'as much a part of contemporary culture as just about anything you care to name', and feels that it is worthy of preservation for posterity. None of the content of the museum is solicited – all the spam messages on display have been collected purely as a result of posting the email address stephennewton.mofs@blogger.com on various websites.

Should you feel the need for a spam fix, you can check out the daily spam exhibits, along with six months of archived spam messages, at http://www.spammuseum.blogspot.com/.

If you prefer something a little more cerebral, you could pop out to your local bookstore to buy a copy of *I Am Spam*, a poetry collection inspired by spam. Each of the poems, written by Chicago poet Larry O. Dean, uses the subject line of a spam message as its title – these include 'Stop paying too much', 'Be your own boss' and 'Emilee is that you?'. Dean says he came up with the idea after being inundated by spam and finding many of the subject lines amusing, provocative or odd-sounding. The book already has a fan base – Dean claims, 'People tell me they are not so angry over their spam now and are beginning to see its more ludicrous and light-hearted potential.' *I Am Spam* is published by Fractal Edge Press.

EVENTS

The Federal Trade Commission (FTC) and the US Department of Commerce's National Institute of Standards and Technology will host an Email Authentication Summit on November 9 and 10, 2004. The Summit will bring together technologists and other interested parties to discuss the market's development, testing, evaluation and deployment of domain-level authentication systems. For details see http://www.ftc.gov/.

INBOX East takes place 17–19 November 2004 in Atlanta, GA, USA. The event will feature over 50 sessions across five tracks: systems, solutions, security and privacy, marketing and 'The Big Picture'. See http://www.inboxevent.com/.

The Second Conference on Email and Anti-Spam (CEAS 2005), will be held in summer 2005 (date and venue yet to be announced). More details will be announced at http://www.ceas.cc/.

FEATURE

WHY SPAM LAWS HAVEN'T WORKED ... YET

Matthew Prince
Unspam, LLC and John Marshall Law School, USA

More than 75 governments around the world have passed anti-spam laws. They have tried opt-in and opt-out regimes. They have required emails to carry labels or valid return addresses. They have mandated truthful subject lines or insisted that senders confirm their identities. They have passed every combination of regulations to express the sentiment that they don't like spam. Unfortunately, the one thing these regulations have in common is that they have all failed.

That is not an overstatement. The number of successful prosecutions against spammers since 1997 (when the state of Nevada in the United States passed the world's first anti-spam law) can be counted on your fingers. That's damning, but worse still is the fact that spam laws have proved to have virtually no deterrent effect.

You would expect that as the governments of the world came together to put an industry out of business, the number of participants in that industry would decrease. The empirical evidence, however, is that with each new law the number of spammers and spam messages has increased.

The most telling statistic is this: when the US Congress passed the CAN-SPAM Act, the compliance rate among spammers was around three per cent, according to a survey by anti-spam vendor *MX Logic*. That is pathetically low, but six months later when the survey was taken again, compliance was down to a scant one per cent.

The US is not alone in its failure. While many in the anti-spam community trumpeted the European directive requiring marketers to get permission from recipients before sending them messages – a so-called 'opt-in' anti-spam regime, arguably much stricter than the United States opt-out approach – the practical effect has been virtually no prosecutions and a dramatic rise in European-based spammers. Spam is like the Jerry Springer style of TV programme: a vulgar, but very profitable business that started in the US but is quickly taking root throughout the rest of the world.

WHY

Why have anti-spam laws failed? Try putting yourself in the position of a prosecutor. You are under-funded. You are under-staffed. Your caseload is overwhelming, and spammers rank far below some of your other priorities: bank robbers, murderers and rapists.

Don't get me wrong; prosecutors would love to go after spammers. Pleas to stop spam are among the most common consumer complaints prosecutors hear regularly. Having interviewed countless officials charged with enforcing anti-spam laws around the world, I can tell you that they would fall over themselves to put a few spammers behind bars. But today, when they weigh the potential costs of these cases against the benefits, the balance falls consistently against prosecution.

Spam cases are not easy. Imagine yourself again as a prosecutor. If you go after a spammer, not only do you have to understand and explain the technical minutiae of email – something many prosecutors are ill-equipped to do – you are also basing your case on a fragile premise: that the messages sent really were unsolicited. Every night you lose sleep worrying over the risk that the next day the spammer will produce a document, forged or real, purportedly proving that the recipient asked for the messages. And, with that, your case will be scuttled.

This is not a baseless worry. Studies have shown that Internet users regularly forget what they have signed up for. As a result, under current anti-spam laws, the line between 'spammer' and 'online merchant' becomes blurred almost beyond distinction. Take, for example, the recent case brought by the New York Attorney General Elliot Spitzer against alleged mega-spammer Scott Richter. After much initial gusto, the case was settled by Spitzer's office for a relatively trivial sum after the Attorney General was reportedly unable to prove that Richter's recipients had not in fact opted-in.

The practical result of this uncertainty is a dramatic increase in prosecutors' costs. For instance, a major US-based ISP experienced in anti-spam litigation has said that these cases start at $200,000 if defendants don't put up much of a fight, and can quickly escalate to $2 million if they do.

These extraordinary costs come from a number of sources. First, and easiest to identify, is the cost of tracking down the spammer and preserving the evidence in a manner acceptable for trial. On top of this, the prosecutor must bear the substantial cost of bringing the case to trial. These costs must be multiplied by the likelihood of success at trial. For example, if only 20 per cent of cases are won by a prosecutor, then the costs of suing a single spammer successfully are effectively multiplied five times.

These costs are inherent to every anti-spam law in effect today. While passing a European-style opt-in law may express a stronger sentiment that a community does not approve of spam, it does little to decrease the costs prosecutors face when they bring a case. Therefore, even opt-in laws are unlikely to do any good in deterring spammers.

An old lawyers' adage states: without enforcement there is no law. The moral of anti-spam regulation to date is that, regardless of what approach a government takes, these laws fail because they have not recognized the real costs prosecutors face. Lawmakers have focused on expressing the sentiment that they don't like spam, rather than empowering prosecutors with real tools to do something about it. If anti-spam laws are ever to succeed, they must move from merely expressing a sentiment to empowering real action.

GETTING THE MESSAGE

Governments are beginning to get the message. A couple of recently enacted laws give me hope that lawmakers are beginning to move from mere sentiment to real action.

Australia's law, for instance, should be a model of how to create a good anti-spam law. While on the surface the law looks little different from the Europeans' opt-in approach, the key to the Australian law's effectiveness is that its drafters focused on how to decrease the burdens placed on law enforcers. Cases are streamlined, the burden of proof is lowered, clear lines are drawn, and prosecutors are given real tools to pursue spammers effectively. These choices, more than the choice of opt-in versus opt-out, appear to be essential to making an anti-spam law effective.

While Australian law enforcement agencies have yet to bring a single case against a spammer, the law alone has had enough of a deterrent effect that most Australia-based spammers have decided to get out of their current line of business. Unfortunately, because of the law's otherwise traditional approach, Australian prosecutors still face substantial costs if they ever do go to trial. Over time, if the law cannot be enforced regularly, spammers may creep back into the country.

A more revolutionary approach is being taken by two U.S. states. Both Michigan and Utah recently enacted laws that focus not on eliminating all spam, but on targeting the worst aspect of the problem: inappropriate messages being sent to children. A significant challenge when drafting an anti-spam law is defining what constitutes 'spam'. Laws are often watered down as part of a compromise to ensure their spam definition is not too expansive. As a result, when prosecutors bring cases under these laws, the watered-down definitions prove difficult, and therefore costly, to enforce.

While it may be difficult to agree on a comprehensive definition of 'spam', nearly everyone can agree that messages advertising pornography, alcohol, tobacco, firearms, gambling services or prescription drugs should not be sent to addresses to which children have access. Any such messages targeting children – solicited or unsolicited – are an easily identifiable and reprehensible form of spam.

Michigan and Utah are in the process of implementing registries of children's addresses that are off-limits to inappropriate messages. Legitimate marketers sending these messages will pay a small fee to scrub their lists of any forbidden addresses.

Inevitably, spammers will ignore the lists. The beauty of these laws, however, is that they provide prosecutors real resources and a comparatively easy case to go after any marketers who continue to prey on children. First, money generated from the list-scrubbing fees can be used to help fund prosecutions. This offsets the costs law enforcers face when they have to bring these cases and eliminates insufficient funding as a reason for law enforcers to avoid prosecuting spammers.

Second, compared with traditional spam cases, the costs are substantially lower and the likelihood of success is substantially higher. Gone is much of the required explanation and analysis of the technical minutiae of email headers. Instead, prosecutors need only prove three things: 1) that the message was sent to an address accessed by a child and listed on the registry, 2) that the message contained or linked to inappropriate content and 3) that the defendant sent, or hired someone else to send, the message.

Whether a message is solicited or unsolicited, the worry that keeps prosecutors awake at night under traditional laws becomes irrelevant. Just as a tavern owner should be liable if a minor is served a beer on his premises, senders of potentially inappropriate messages should be held to a standard of care that requires them to prevent their messages from landing in the inboxes of children.

NEXT GENERATION

The next generation of anti-spam laws must focus on lowering enforcement costs and empowering real action by prosecutors. The approaches of Australia, Michigan and Utah provide a good starting point for lawmakers, but more legislative experimentation is needed. The one thing we know for certain is that simply repeating the traditional approach will not work. For anti-spam laws to have any deterrent effect, we must move beyond merely expressing that we do not approve of spam, to an action-oriented approach that allows law enforcers to do something about it. Until then, spam law will rightfully continue to be viewed as a failure.

SUMMARY

ASRG SUMMARY: OCTOBER 2004

Helen Martin

A large proportion of this month's postings to the ASRG mailing list centred around a discussion that was started last month, about filtering by detecting 'anti-Bayesian' elements.

Brian Azzopardi said that it is a trivial task for anyone to create 'nonsense' words – pointing out that, with a typical word length of six characters and 26 letters of the alphabet, a random token generator could produce 26^6 combinations. Brian called for more research into the representation of the message that is passed to the filter. Whereas most current filters use a direct representation, where the filter is simply fed the tokens found, Brian argued that it does not have to be like this – for example, tokens longer than 12 characters could be given a different token, such as 'BAYESIAN_TOKEN_TOO_LONG', which would then be fed to the filter.

Laird Breyer agreed that it all comes down to representation – and pointed out that a full and complete representation is not necessarily better than a simple and incomplete one. 'Sometimes,' he said, 'extra information only confuses the decision procedure (not unlike the saying "too many cooks spoil the broth")' – moreover, the best representation depends on the algorithm being used, and vice versa.

Markus Stumpf posted a link to John Graham-Cumming's *Spammers' Compendium* – an online collection of the tricks spammers use to beat spam filters (http://www.jgc.org/tsc/). Markus commented that, recently, he had not seen the use of characters displayed repeatedly to build up a 'graphical' representation of a word that is readable to the human eye, but not to an automated filter – such as:

```
                                      .o.

                                       888
ooo. .oo.    .ooooo.  oooo oooo   ooo 888
‘888P”Y88b   d88' ‘88b  ‘88. ‘88. .8’   Y8P
  888  888  888   888   ‘88..]88..8’   ‘8’
  888  888  888   888    ‘888’’888’    .o.
  o888o o888o ‘Y8bod8P’     ‘8’  ‘8’    Y8P
```

Laird Breyer said that what the filter sees in such cases is untypical of ordinary language and therefore, like nonsense words, it stands out clearly. He asked 'what percentage of legitimate messages do you receive that don't contain the word "the"?' – to which a flood of replies came pointing out that non-native English speakers receive rather a lot. Jose Marcio Martins da Cruz pointed out that many legitimate senders use this kind of composition as a footer to their messages, therefore the automatic recognition of such compositions by statistical filters would cause a lot of false positives. Laird agreed that this would be the case in the short term, but pointed out that, with training, the footer will be recognized. 'As a general rule,' he said, 'tokens which occur commonly in both ham and spam have little effect on a filtering decision. The decisions depend much more on the presence of extreme tokens which (statistically) only occur in spam, or only occur in ham.'

On the subject of different languages, Markus Stumpf reported that, for a lot of his users, for whom less than five per cent of legitimate email communication is in English, SpamAssassin works 'like a charm' – however with the sudden rise in German language spam that has been seen recently, these success rates look likely to change before long.

On a different note, Markus made a plea for greater hierarchy in DNS. As part of plans to implement greylisting, he and his colleagues have been trying to establish a whitelist of 'well known' mail servers. As a starting point they took logfiles of around a million connections and noticed that 'even mailserver farms are named as brain damaged and hierarchy breaking.' Markus asked, 'Why does it have to be mail-[0-9][0-9].iinet.net.au instead of host[0-9][09-].mail.iinet.net.au?'

He said it would be more correct, and much easier to anti-greylist, '.mail.domain.tld' rather than adding 20 records (mail-smtp-01.domainl.tld … mail-smtp-20.domainl.tld). Jochen Topf said the reason was because, most of the time, admins within an organisation would use only the hostname and not the fully qualified domain name – and won't want to have to ask whether 'host165' is 'host165.mail' or 'host165.web' etc.

Douglas Campbell agreed that ISPs should corral all their official mailservers under special subdomains, but pointed out that, in the absence of specification, ISPs have implemented their own internal methods of determining official mailservers and will not change without significant reason to do so.

Finally, Phillip Hallam-Baker relayed to the group a report of a new level of perfidy to which spammers have sunk. Spammers subscribe to a *Yahoo!* group, then send out their messages to the group with a notice saying that if anyone has a problem with the spam they should unsubscribe from the group. Jim Witte was quick to point out that this behaviour almost certainly violates *Yahoo!* group policy and, should the spammers be caught, the company would be in a strong position to sue them.

[*An archive of all messages posted to the ASRG mailing list is at http://www1.ietf.org/mail-archive/web/asrg/current/*].