APRIL 2008

virus BULLETIN

Fighting malware and spam

CONTENTS

2 **COMMENT**
Political DDoS around the world

3 **NEWS**
VB2008 conference programme revealed
Sullied site stats

3 **VIRUS PREVALENCE TABLE**

4 **MALWARE ANALYSIS**
Your computer is now stoned (...again!)

9 **FEATURE**
Anti-stealth fighters: testing for rootkit detection and removal

13 **COMPARATIVE REVIEW**
Windows Vista Business Edition SP1

28 **END NOTES & NEWS**

IN THIS ISSUE

STEALTH RISING

Mebroot – the MBR rootkit – is one of the most advanced and stealthiest pieces of malware seen to date. It operates in the lowest levels of the operating system, uses many undocumented tricks and relies heavily on unexported functions and global variables. Elia Florio and Kimmo Kasslin track the rise of the MBR rootkit.
page 4

VB100 ON VISTA SP1

John Hawes wipes the sweat from his brow after completing a comparative review of 40 anti-malware products for Vista. With polymorphic trip-ups, false positives and stability issues in the fray it proved to be a tough month for the products involved.
page 13

vbSpam supplement

This month: anti-spam news and events, and Tyler Moore questions the wisdom of crowds when assessing phishing websites.

ISSN 1749-7027

'We have tracked tens of thousands of DDoS attacks ... A subset of [them] appear to be politically motivated.'

Jose Nazario, Arbor Networks

POLITICAL DDoS AROUND THE WORLD

DDoS attacks are designed to overwhelm a target network with resource requests, leaving the victim unable to handle legitimate requests. These can come in many forms, but typically we see traffic floods that consume bandwidth rather than application resources. DDoS attacks are not new, and have grown in intensity and popularity in the past ten years with the rise of botnets.

Botnets provide the needed firepower behind a DDoS attack – bandwidth and computers – as well as the infrastructure to manage such an attack. In measurements conducted in 2006 we found that approximately half of all of the botnets we monitored launched at least one DDoS attack. Traditional botnets are not the only source of these attacks, though, as we are increasingly seeing specialized kits being deployed to launch and control DDoS attacks.

Our own research over the years has shown a steady increase in the severity of DDoS attacks. Based on surveys with tier-1 ISP operators, we found that the largest observed DDoS attacks in the wild top over 40 Gbps.

Motivations for DDoS attacks are often related to retaliation or anger, and sometimes include extortion or punitive attacks. In the past few years we have tracked tens of thousands of these sorts of attack across the globe and have found that no network is immune to such an event. Most frequently we see small attacks against broadband subscribers or small e-commerce sites. Larger, more sophisticated attacks involve extorting major online businesses. Some attacks have caused businesses significant financial problems through the loss of the ability to handle customers or bandwidth charges.

At present, we are witnessing a series of DDoS attacks against online gambling sites. These are orchestrated by a small set of attackers and may be related to extortion schemes. In these attacks, several poker and casino sites have been hit with sustained attacks lasting days and, in some cases, weeks. These can cripple the victim's site – directly impacting on the business.

A subset of DDoS attacks appear to be politically motivated. In one of the most high-profile events recently, Estonian government and national infrastructure sites were hit with several weeks' worth of DDoS attacks. These attacks coincided with the staging of street protests over Russia's history in Estonia. Many people assumed that Russian authorities had orchestrated the attacks, although no evidence was found to support that claim. We found that botnets as well as manual coordination were behind most of the DDoS attacks, with Russian-language forums used in the organization of the attacks.

More attacks were staged in the winter of 2007 against Estonian newspaper *DELFI*, during its coverage of the trials of several Russians charged with street-level crimes during the protests earlier in the year.

Other politically motivated DDoS attacks we have seen recently include those against Russian politician Gary Kasparov and his party during the run up to the winter 2008 elections.

Political DDoS events are not limited to Russian and European networks. Most of the attacks we measure through our ATLAS system are sourced from the US, and the majority of the attacks we see target US victims. This makes sense given the amount of address space located in the US. In the past we have also seen DDoS attacks related to Indian and Pakistani conflicts, and recently against Iranian targets.

As international tensions rise and the number and size of botnets continue to increase, we expect this specific attack motivation to continue. It will be interesting to see how geopolitical events unfold online in the coming months and years.

NEWS

VB2008 CONFERENCE PROGRAMME REVEALED

VB has revealed the conference programme for VB2008, Ottawa.

Once again, the three-day conference programme boasts an exceptional line-up of anti-malware and anti-spam expert speakers and caters for both technical and corporate audiences. Presentations will cover subjects including: sample sharing, anti-malware testing, automated analysis, rootkits, spam and botnet tracking techniques, corporate policy, business risk and more.

In addition to the 40-minute presentations already scheduled, a portion of the technical stream is set aside for brief (20-minute) technical presentations, dealing with up-to-the-minute specialist topics. Proposals for the 'last-minute' presentations must be submitted by 5 September 2008 (details of how to submit proposals will be announced in due course). The schedule for the last-minute presentations will be announced shortly before the start of the conference.

VB2008 takes place 1–3 October 2008 in Ottawa, Canada. Online registration is now available. For the full programme see http://www.virusbtn.com/conference/vb2008/programme/.

SULLIED SITE STATS

Key findings of a study by *ScanSafe* of the more than 80 billion corporate web requests it scanned and 800 million web threats it blocked in 2007 include that viruses, trojans, password stealers and other forms of malware are becoming more prevalent, that an increasing number of legitimate sites are unwittingly hosting malware, and that compromised sites are remaining infected for longer – in some cases for more than two months.

The news come just weeks after reports of a major new outbreak of website infections, with as many as 20,000 legitimate sites thought to have been hit with a single wave of malicious iframe insertion attacks. Unfortunately it seems that not even anti-malware vendors are immune, with *Trend Micro* having to issue a warning on its Japanese site last month that some of its web pages had been infected – and in late December security firm *CA* was among thousands of legitimate websites to have been infected by hackers taking advantage of an SQL injection vulnerability. All of the above should serve as a reminder to web administrators to ensure that their web servers are properly secured.

Prevalence Table – February 2008

Malware	Type	%
NetSky	Worm	29.18%
Cutwail/Pandex/Pushdo	Trojan	27.04%
Mytob	Worm	12.54%
Mydoom	Worm	6.03%
Bagle	Worm	5.64%
Small	Trojan	2.72%
Agent	Trojan	2.68%
Grew	Worm	2.26%
Zafi	Worm	1.81%
Virut	Virus	1.45%
Blebla	Worm	1.14%
Mywife/Nyxem	Worm	0.89%
Stration/Warezov	Worm	0.83%
ExploreZip	Worm	0.72%
Bugbear	Worm	0.51%
Badtrans	Worm	0.35%
Frethem	Worm	0.35%
Sality	Virus	0.35%
Gibe	Worm	0.34%
Benjamin	Worm	0.29%
Goner	Worm	0.25%
Zlob/Tibs	Trojan	0.24%
Grum	Worm	0.23%
Bagz	Worm	0.18%
VB	Worm	0.17%
Klez	Worm	0.16%
Doombot	Worm	0.12%
Aliz	Worm	0.11%
Delf	Trojan	0.10%
Nuwar/Peacomm/Zhelatin	Trojan	0.10%
Fleming	Worm	0.08%
Higuy/Tettona	Worm	0.08%
Oror/Roron	Worm	0.07%
Sdbot	Worm	0.07%
Others[1]		0.92%
Total		100.00%

[1]Readers are reminded that a complete listing is posted at http://www.virusbtn.com/Prevalence/.

MALWARE ANALYSIS

YOUR COMPUTER IS NOW STONED (...AGAIN!)

Elia Florio
Symantec Security Response, Ireland

Kimmo Kasslin
F-Secure Security Lab, Malaysia

We can trace the first evidence of 'Mebroot' (the MBR rootkit) back to the end of 2007. According to the PE timestamp of the oldest sample seen, it was compiled during early November 2007 and distributed multiple times over several weeks at the end of the year. The timeline of Mebroot's evolution (as shown in Figure 1) was first outlined by Matt Richard from *iDefense* [1], who discovered the first sample in the wild together with the *GMER* team [2].

We know that during November 2007 a malicious domain (hxxp://gfeptwe.com) – used in the past to distribute and install variants of Trojan.Anserin (a.k.a. Sinowal) – began to serve copies of the MBR rootkit for a limited period of time. The malware was installed via drive-by exploits using a set of old *Microsoft* vulnerabilities, probably in an attempt to stay under the radar during this 'beta' release stage.

Two waves of related drive-by attacks took place between December 2007 and January 2008. These attacks were followed by a period of calm before finally, in February 2008, the steady flow of attacks installing Mebroot resumed [3]. The whole timeline seems like a development and QA project; in fact all the variants released in the initial period have close PE timestamps and very small changes in the code.

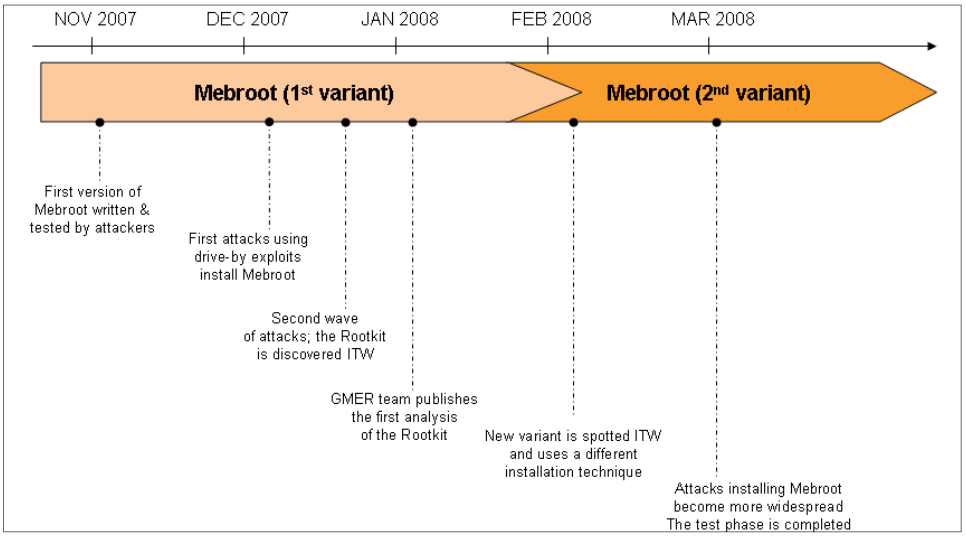

Figure 1: Timeline of Mebroot evolution from 'beta' to final release.

We do not know how long it has taken the authors to develop and write the code of this sophisticated threat, but the idea of malicious code that modifies a system's MBR is not new (even ignoring DOS attacks of old), having first been discussed some years ago. In 2004, Greg Hoglund wrote about MBR attacks in his book *Exploiting Software* [4], while the most notable research in the area of MBR rootkits was undertaken by Derek Soeder of *eEye* during 2005 [5]. Soeder created BootRoot, a proof-of-concept MBR rootkit able to target *Windows XP* and *2000*. Finally, researchers Nitin and Vipin Kumar of *NVLabs* recently published a paper [6] about a new type of MBR rootkit called Vbootkit, designed expressly to work on *Windows Vista*.

It is quite obvious that Mebroot's authors have benefited from other people's research, and this fact is confirmed by a quick comparison of the MBR code of Mebroot and BootRoot. A large area of Mebroot's MBR loader is almost identical to the BootRoot code published by *eEye*. Mebroot's MBR code hooks INT 13 at boot exactly as BootRoot does, with the intent of patching the OSLOADER image (part of the NTLDR file) when it is loaded. This patch is done on the fly with the same static signature as that used by BootRoot (8BF085F6742?803D). The signature is patched with a CALL DWORD[addr] instruction that passes control to the second-stage payload of the malware.

RAW DISK ACCESS UNDER WINDOWS

Mebroot arrives with an EXE installer that is typically between 250 KB and 350 KB and which takes control of the system by overwriting the MBR. This attack is possible because some versions of *Windows* allow programs to overwrite disk sectors (including the MBR) directly and without proper restrictions.

The initial reports about this MBR attack were slightly confused, so let's clarify some points to understand when the attack is possible. On *Windows 2000*, *XP* and *2003* systems, raw access to disk is possible for any user-mode program running in ring-3 (no need to go into ring-0!), but this requires

Administrator privileges [7]. The fact that most users run *Windows* as Administrator clearly makes them vulnerable to this type of rootkit.

The issue has been known about for some time in the *2K/XP* families, and *Vista* was partially secured in 2006 (with Release Candidate 2) following a successful attack demonstration by Joanna Rutkowska [8]. In fact, the attack is now mitigated on *Vista* by User Account Control (UAC), which blocks raw access to disks. The table below summarizes which OSs can be infected by Mebroot:

Windows OS	*Can MBR be infected?*	*Is rootkit active?*
Windows 2000 (user is Administrator)	YES	YES
Windows XP (user is Administrator)	YES	YES
Windows 2003 (user is Administrator)	YES	YES
Windows Vista (UAC disabled)	YES	NO
Windows Vista (UAC enabled)	NO	NO

It is important to clarify that Mebroot can only infect the *Vista* MBR if UAC is disabled and that, even after a successful infection, the rootkit will not be able to load itself at boot because it targets specific signatures of the *Windows* kernel that are not present on *Vista*. In this scenario *Vista* users may live with an infected MBR that boots up the operating system normally, without any rootkit activity, because the malware is never loaded in memory. In addition to this, *Vista* is secure because its boot process is completely different from any previous OS. It is possible that future variants of the threat may be built to overcome this limitation though.

OWNING THE MBR

All the rootkit installer samples are encrypted with a custom 'spaghetti-like' packer that has already been seen in the Anserin/Sinowal family. This nasty packer scrambles the execution flow of a program by interleaving valid opcodes with JMP or JMP DWORD[addr] instructions. The result is a piece of polymorphic code that is difficult to trace and analyse, but which retains its functionality. Unpacking is trivial and requires just a breakpoint on the VirtualFree API.

The installer file (mat[n].exe) seems to have a double entry point because it is designed to run either as an EXE or as a DLL module. In fact, after the infection of the MBR, the malware patches itself to become a DLL and runs a second time using the regsvr32 command to delete itself and reboot the machine. Some installers also have a delayed infection routine with a random timer. The infection starts after a delay of a few minutes to avoid automated analysis by honeypot systems and to fool quick black box analysis.

Mebroot tries to infect the first 16 disk drives connected to the machine with a loop that calls CreateFile() on \\.\PhysicalDrive[n] several times. A side effect of this behaviour is that in some cases the rootkit also infects external USB disks and hard drives. Infected external disks do not have active infections because typically they are not used to boot the operating system, but the disks will still contain traces of the malware on some sectors. One of the latest variants of Mebroot (mat25.exe) uses a different approach to perform raw operations on disk. Instead of

```
00403810  $vFF25 E0104000    JMP DWORD PTR DS:[4010E0]          mat25.00403852
00403816  . 5A               POP EDX
00403817  .vFF25 90104000    JMP DWORD PTR DS:[401090]          mat25.004039E2
0040381D  .v0F8F F2000000    JG mat25.00403915
00403823  .vFF25 50104000    JMP DWORD PTR DS:[401050]          mat25.004039A1
00403829    83               DB 83
0040382A    EC               DB EC
0040382B  . 02FF             ADD BH,BH
0040382D  . 25 88104000      AND EAX,401088
00403832  .v76 73            JBE SHORT mat25.004038A7
00403834  .vFF25 D0104000    JMP DWORD PTR DS:[4010D0]          mat25.00403964
0040383A  . BE 9F4A4000      MOV ESI,mat25.00404A9F
0040383F  .vFF25 BC104000    JMP DWORD PTR DS:[4010BC]          mat25.004038D2
00403845  . 54               PUSH ESP
00403846  .vFF25 20104000    JMP DWORD PTR DS:[401020]          mat25.00403987
0040384C  >vFF25 04104000    JMP DWORD PTR DS:[401004]          mat25.00403867
00403852    60               DB 60                              CHAR '`'
00403853  .vFF25 DC104000    JMP DWORD PTR DS:[4010DC]          mat25.004038F9
00403859  . 85C0             TEST EAX,EAX
0040385B  .vFF25 A4104000    JMP DWORD PTR DS:[4010A4]          mat25.004039C7
00403861  >vFF25 08104000    JMP DWORD PTR DS:[401008]          mat25.0040384C
00403867    61               DB 61                              CHAR 'a'
00403868  .vFF25 00104000    JMP DWORD PTR DS:[401000]          mat25.00403889
0040386E  . 81FB 9F4A4000    CMP EBX,mat25.00404A9F
00403874  .^FF25 D4104000    JMP DWORD PTR DS:[4010D4]          mat25.00403832
0040387A  . 59               POP ECX
0040387B  .^FF25 A8104000    JMP DWORD PTR DS:[4010A8]          mat25.00403859
```

Figure 2: Mebroot 'spaghetti' packer in action. Too many jumps...

```
0040172B 68 00 00 07 00                  push    IOCTL_DISK_GET_DRIVE_GEOMETRY ; dwIoControlCode
00401730 FF 75 08                        push    [ebp+hPhysicalDriveX] ; hDevice
00401733 FF 15 20 10 40 00               call    ds:DeviceIoControl
00401733
00401739 85 C0                           test    eax, eax
0040173B 0F 84 F9 03 00 00               jz      exit_EAX_0
0040173B
00401741
00401741                          checkHD_BytesPerSector:
00401741 BB 00 02 00 00                  mov     ebx, 512
00401746 39 5D D8                        cmp     [ebp+outbuf_DiskGeom.BytesPerSector], ebx
00401749 0F 85 EB 03 00 00               jnz     exit_EAX_0
00401749
0040174F 56                              push    esi             ; lpOverlapped
00401750 8D 45 F8                        lea     eax, [ebp+NumberOfBytesRead]
00401753 50                              push    eax             ; lpNumberOfBytesRead
00401754 53                              push    ebx             ; 512
00401755 8D 85 C4 FD FF FF               lea     eax, [ebp+buf_OriginalMBR]
0040175B 50                              push    eax             ; lpBuffer
0040175C FF 75 08                        push    [ebp+hPhysicalDriveX] ; hFile
0040175F FF 15 18 10 40 00               call    ds:ReadFile
0040175F
00401765 85 C0                           test    eax, eax
00401767 0F 84 CD 03 00 00               jz      exit_EAX_0
00401767
0040176D 39 5D F8                        cmp     [ebp+NumberOfBytesRead], ebx
00401770 0F 85 C4 03 00 00               jnz     exit_EAX_0
00401770
00401776
00401776                          check_BootMagic:
00401776 66 81 7D C2 55 AA               cmp     [ebp+buf_OriginalMBR.BootMagicSignature], 0AA55h
0040177C 0F 85 B8 03 00 00               jnz     exit_EAX_0
```

Figure 3: Raw access to disk simply requires the CreateFile and ReadFile/WriteFile APIs.

using CreateFile() on \\.\PhysicalDrive, it installs and loads a driver that works as a 'wrapper' for the system driver disk.sys. Essentially, the new installer uses its own driver to communicate with the OS disk driver and to perform low-level read/write actions using IRP. This strategy can probably bypass protection systems that block raw access to disk.

During installation, the malware first reads the current disk MBR and checks some characteristics of the drive such as the number of bytes per sector (it expects 512 bytes), the signature 0x55AA at the end of the MBR, and whether the drive has already been infected (the infection marker is the DWORD 0xAD022C83 at offset 0x16 of the MBR). Next, it parses the partition table to find the physical end of the disk and verifies that there is enough unpartitioned slack space at the end for it to write its own malicious code. The installer usually needs at least 650 free sectors to store the main rootkit driver. This strategy is clever for two reasons: first, the driver is not stored as a file on the system, but in raw disk sectors. Secondly, writing the malicious driver after the end of the disk means that it requires some forensic expertise to extract samples from infected machines.

The installer makes note of the sector in which the rootkit executable is stored and then adjusts in memory the Payload Loader shellcode that will load the SYS driver at the next reboot. Finally, it overwrites three sectors immediately before the beginning of the first partition. On *Windows 2000* and *XP* with single partitions, Mebroot typically overwrites sectors 60, 61 and 62. These sectors may be different on systems with multiple OS and disk partitions.

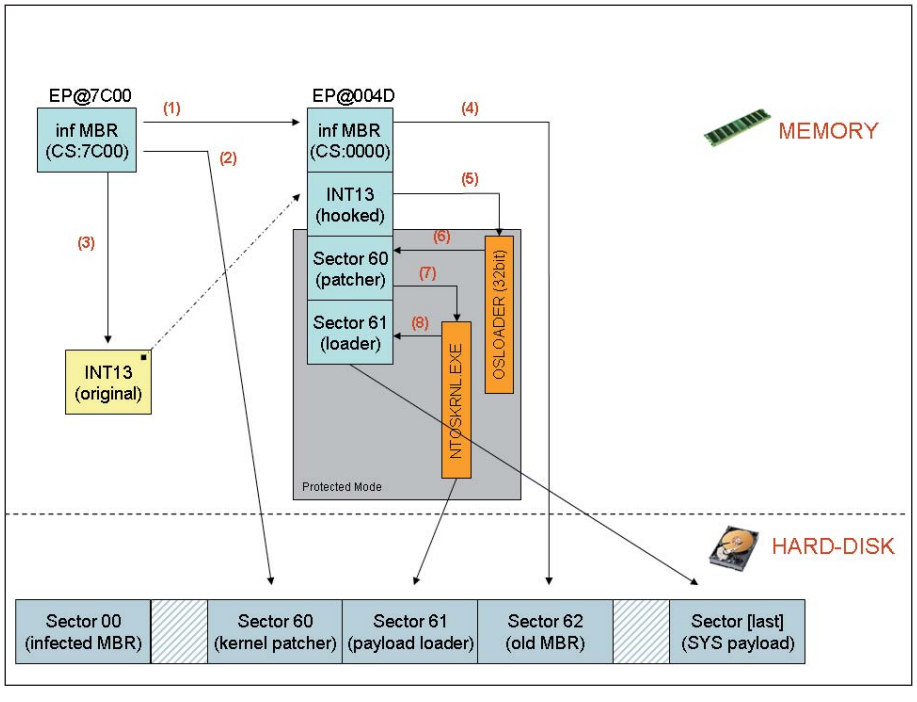

Figure 4: Mebroot loading process: how to own the system from the boot.

OWNING THE SYSTEM FROM THE BOOT

The complete scheme of the Mebroot loading process is shown in Figure 4.

The following is a step-by-step description of the rootkit boot process and kernel infection:

1. The infected MBR reserves 2 KB of conventional memory and relocates itself from 0x7C00 to 0x0000.
2. Next, it reads payloads from sectors 60 (kernel patcher) and 61 (payload loader) into memory blocks adjacent to the relocated code.
3. The MBR code hooks INT 13 and passes control to relocated code at 0x004D.
4. It reads sector 62 (old MBR) back to 0x7C00 memory and passes control to it; the OS starts booting up normally while INT 13 is hooked by the threat.
5. The hooked code intercepts all disk-reading operations and patches the OSLOADER module (part of NTLDR) when it is loaded from disk.
6. The patched OSLOADER calls the Kernel Patcher shellcode in memory (sector 60).
7. This shellcode scans and patches the NTOSKRNL.EXE image near 'CALL nt!IoInitSystem'.
8. The modified NTOSKRNL.EXE calls the Payload Loader shellcode (sector 61), which loads and runs the rootkit driver stored in the last sector of the disk.

To minimize the footprint and traces in memory, the loader shellcode deletes itself by filling the memory area in which it is stored with zeroes. This detail leads us to believe that nothing is left to chance and the authors of this nasty piece of code are skilled and meticulous malware programmers.

NEW DISK STEALTH TECHNIQUES

Analysis of the final rootkit driver loaded in memory requires

some extra effort. Some rootkit variants have an extra packing layer that unpacks the real kernel driver using scrambled spaghetti code. In this case, a good breakpoint on ExAllocatePoolWithTag will do the job and allow us to dump the final unpacked driver.

Since the rootkit SYS driver is loaded by its own loader in an unusual way, the module does not expect the normal parameters passed to *Windows* drivers. In fact, it receives three parameters passed by the Payload Loader: the kernel ImageBase of the unpacked driver, a pointer to PsLoadedModuleList (used to resolve imports) and the ImageBase of the packed driver. The rootkit resolves all NTOSKRNL and HAL imports with its own routine and also deletes from memory the packed driver image when it is no longer needed. Later on, even the MZ header of the unpacked driver is deleted from memory to minimize its footprint, leaving in the kernel space only random traces of code in executable memory pages.

The rootkit hides itself by hooking the disk.sys driver. It finds DeviceObject for \Device\HardDisk[N]\DR0 and reads the old MBR from sector 62 into an allocated pool that will be used as a 'cached copy' of the old MBR to improve the performance of stealth operations. Since the rootkit does not have files, process or registry keys to hide, the stealth functionalities are limited to intercepting read/write operations on raw disk sectors. This is done by hooking the dispatch handlers of the \Driver\Disk for IRP_MJ_READ and IRP_MJ_WRITE routines. When a program tries to read the MBR (sector 00) or any other sector used by the rootkit (60, 61, 62 or sectors after the end of the disk) the hooked code will return a fake image of the sector, showing the old MBR or an empty sector filled with zeroes in the other cases. In a similar way, the rootkit will protect itself by blocking all write operations to its sectors.

The rootkit needs to maintain hook-free versions of the IRP_MJ_READ and IRP_MJ_WRITE functions, so it uses a special trick: it generates a random DWORD value used as a 'magic key'. Later, the rootkit is able to perform normal read/write operations with the original dispatch routines simply by calling the disk.sys driver with an IRP packet that contains this magic key at offset 0x40.

NEW FIREWALL-BYPASSING TECHNIQUES

Analysing the rootkit driver's network code becomes even more difficult. The majority of its functions are still heavily obfuscated, even after successful unpacking. The fastest approach to bypass the obfuscation is to use code tracing and custom scripts to clean up the trace logs of extra garbage. After a lot of frustration and some breakthroughs, we now know that Mebroot's firewall-bypassing technique is similar to, but goes one step further than that used by Srizbi [9].

Like Srizbi, Mebroot operates in the NDIS layer, but it uses a different approach to gain access to the internal NDIS structures. Whereas Srizbi installed a dummy protocol, Mebroot uses the unexported ndisMiniportList which points to an existing miniport described by the _NDIS_MINIPORT_BLOCK structure. To gain exclusive access to the list it acquires the ndisMiniportListLock spinlock, which is also unexported.

From the miniport block the code uses a similar approach to that used by Srizbi to find a suitable adapter that is bound to either the PSCHED or the TCPIP protocol. Finally, it finds the address of the lowest-level send handler function and hooks four NDIS handler functions.

To send packets it uses the following handler function:

```
NDIS!_NDIS_M_DRIVER_BLOCK

 +0x020 MiniportCharacteristics : _NDIS51_MINIPORT_
CHARACTERISTICS

   +0x040 SendPacketsHandler : 0xf9adf332 void pcntpc
i5!LanceSendPackets+0
```

To get a notification after the send operation has completed it uses the following hook:

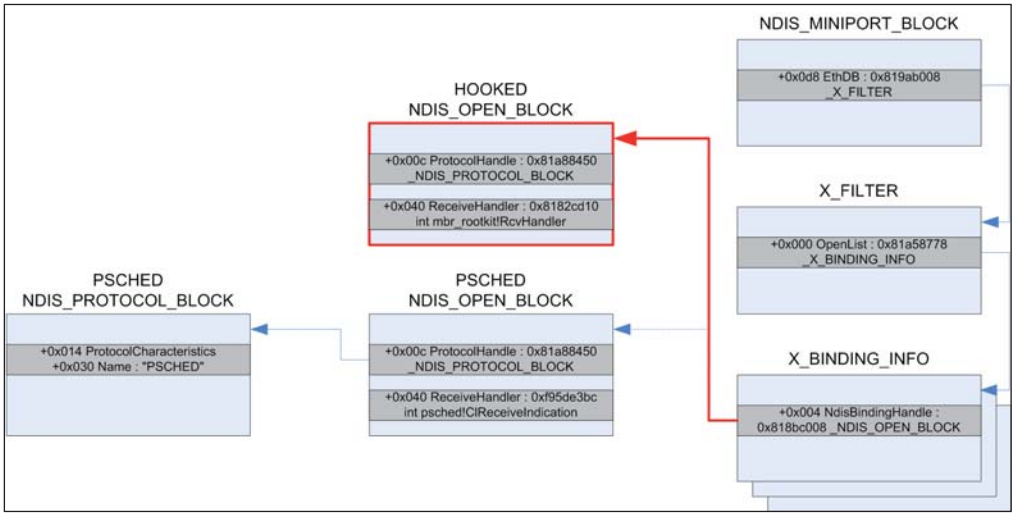

Figure 5: Mebroot activates the full set of its hooks only when it needs them.

```
NDIS!_NDIS_MINIPORT_BLOCK
 +0x0ec SendCompleteHandler : 0x81825bb0 void mbr_
rootkit!Hook_SndCompHdlr
```

To receive packets it uses the following hooks:

```
NDIS!_NDIS_OPEN_BLOCK
 +0x040 ReceiveHandler : 0x8182cd10 int mbr_
rootkit!Hook_RcvHdlr
 +0x050 ReceivePacketHandler : 0x8182e400 int mbr_
rootkit!Hook_RcvPcktHdlr
```

Mebroot's network code is advanced in many ways. It is stealthy – only a single pointer is hooked at any one time. The rest of the hooks in the selected protocol's _NDIS_OPEN_BLOCK structure are in use only when the rootkit is sending packets. It accomplishes this by creating a copy of the original open block structure which is then hooked. When it needs to send a packet it replaces a single pointer from the _X_BINDING_INFO structure to point to its private open block structure to make sure the packets received from that point onwards will be processed by its own handler functions. Once the packets have been processed the original pointer is put back. This process is illustrated in Figure 5.

Another example of Mebroot's stealth is the way it ensures that none of the NDIS API functions it relies on are hooked by firewalls. Instead of just copying the original ndis.sys from disk into allocated memory and using it as its private module, it uses a 'code pullout' technique to load only the relevant parts of the code into memory. This technique was first described by Alexander Tereshkin, a.k.a. 90210, at rootkit.com [10]. Once the relevant code blocks have been copied into one continuous block of memory it is prepared for execution. Finally, the code patches its own import address table to make sure all imported NDIS API functions point to the code that was pulled out.

CONCLUSIONS

Mebroot is the most advanced and stealthiest malware we have analysed so far. It operates in the lowest levels of the operating system, uses many undocumented tricks and relies heavily on unexported functions and global variables. We did not encounter a single blue screen while examining the latest samples that were distributed after February 2008. This is a clear sign of the level of professionalism of today's malware authors.

It is also evident that the author of Mebroot is following closely the research carried out by individuals who have presented their findings at Black Hat conferences and on rootkit.com. Mebroot's MBR code was almost identical to BootRoot's, while the firewall-bypassing code closely resembles the most advanced ideas presented by Tereshkin at Black Hat USA 2006 [11]. In addition, after we successfully deciphered some of the code used to perform the code pullout it became clear that some of the functions were one-to-one with functions that are part of the phide2 source code. Maybe the next malware from Mebroot's author will use virtualization to make it even more difficult to detect and remove – at least proof of concept source code is already available for this [12].

Elia Florio and Kimmo Kasslin will present an extended and more detailed look at the MBR rootkit at VB2008 later this year. VB2008 takes place 1–3 October 2008 in Ottawa, Canada. See http://www.virusbtn.com/conference/vb2008/ for the full programme and registration details.

REFERENCES

[1] Master Boot Record timeline. http://isc.sans.org/diary.html?storyid=3820.

[2] Stealth MBR rootkit (Jan 2nd, 2008). GMER team. http://www2.gmer.net/mbr/.

[3] The Flow of MBR Rootkit Trojan Resumes. http://www.symantec.com/enterprise/security_response/weblog/2008/02/the_flow_of_mbr_rootkit_trojan.html.

[4] Hoglund, G.; McGraw, G. Exploiting software: how to break code. 2004, p.429.

[5] eEye BootRoot. http://research.eeye.com/html/tools/RT20060801-7.html.

[6] BOOT KIT: Custom boot sector based Windows 2000/XP/Vista subversion. http://www.nvlabs.in/?q=node/11.

[7] INFO: Direct Drive Access Under Win32. Microsoft. http://support.microsoft.com/kb/q100027.

[8] Rutkowska, J. Subverting Vista Kernel for fun and profit. 2006. http://www.invisiblethings.org/papers/joanna%20rutkowska%20-%20subverting%20vista%20kernel.ppt.

[9] Kasslin, K.; Florio, E. Spam from the Kernel. Virus Bulletin. November 2007, pp.5–9. http://www.virusbtn.com/vba/2007/11/vb200711-srizbi.

[10] Phide2. http://rootkit.com/vault/90210/phide2.zip.

[11] Tereshkin, A. Rootkits: Attacking Personal Firewalls. Black Hat USA 2006. http://www.blackhat.com/presentations/bh-usa-06/BH-US-06-Tereshkin.pdf.

[12] Blue Pill Project. http://bluepillproject.org/.

FEATURE

ANTI-STEALTH FIGHTERS: TESTING FOR ROOTKIT DETECTION AND REMOVAL

Andreas Marx & Maik Morgenstern
AV-Test.org, Germany

'Most people don't even know what a rootkit is, so why should they care about it?' (Thomas Hesse, President, Global Digital Business, Sony BMG [2005])

Malware is becoming more and more complex every day. The number of newly discovered malware samples is skyrocketing, but that's not the only challenge for the AV industry. In most cases, we're looking at malware that is built in a modular way, with plug-ins that support new features such as hiding the malware's presence from the user and from AV products. While it is easy for a good signature-driven product to find a known sample that has not yet been activated, it is becoming increasingly challenging to detect the sample once it is running and trying to hide itself and other malicious components. On the *Windows* platform the hidden objects usually include services and processes, registry keys and values, as well as directories and files.

Shortly after the infamous '*Sony* rootkit' was released in 2005 [1], *AV-Test.org* started testing for rootkit detection. At that time, most AV programs could easily be fooled. As soon as the rootkit was running, the system was reported to be clean – even if a hidden piece of malware was running in the background, sending out junk emails and attempting to infect further computers. Until now, our anti-rootkit test results have only been published in certain hard copy magazines and in German. In an attempt to close this information gap we have decided to present the results of two recent tests here in *VB*.

The first part of our research, a dedicated anti-rootkit test covering 27 products on *Windows XP* (32-bit, SP2) and *Vista* (32-bit, RTM), was published in the German *ComputerBild* magazine [2]. The second part, a small-scale anti-rootkit review as part of a comprehensive AV test of 17 tools on *Windows Vista* (32-bit, RTM), was performed for the German *c't* magazine [3].

STEP 1: SELECTING THE SAMPLES

Before a review can start, samples of standalone rootkits and malware using rootkit technologies must be selected to test against. The manual and automated analysis of such samples is tricky and good reverse-engineering skills are required. For a less comprehensive basic check it might be sufficient just to compare the system in a clean state (without any malware), in an infected state (with the activated rootkit running on the system) and in the state in which the malware-infected system has been booted from a known clean installation (so no files and registry entries are hidden, as the rootkit is not active). For a good review, further analysis needs to be performed to check for other hidden objects on the infected system. This might take several hours per sample.

For the part of the *ComputerBild* review that focused on *Windows XP*, we used a total of 60 samples, including two versions of the *Sony* rootkit (XCP/First4Internet rootkit) found on CDs and one copy of the Alpha DVD (Settec) rootkit used on the German DVD *Mr. and Mrs. Smith* [4]. Malware samples included several variants of Agent, Delf, Dragonbot, Feebs, Fuzen, Graybird, Hacker Defender, Haxdoor, Hider, Hupigon, iBill, Kenfa, Klone, Madtol, Maslan, NsAnti, NT Illusion, NT Rootkit, Nuwar, Pakes, PC Client, QQPass, Rontokbro, Small, Tibs, Wopla and X-Shadow. Some of the malware listed is included on the WildList. The exact samples used for the test have already been shared with the tested AV companies. The *Windows Vista* test for *ComputerBild* was performed with a much smaller set of samples and will not be discussed in detail.

The *c't* review on *Windows Vista* included just six samples which run well on *Vista*, covering the two aforementioned CD rootkits, two versions of Hacker Defender, as well as one copy of NT-Illusion and a copy of Vanquish. These rootkits are a little older, but still work well on *Vista* as long as User Account Control (UAC) has been switched off (a step that was performed prior to testing).

We only used 'real' PCs (equipped with a Core 2 Duo 6600 processor, 2 GB RAM and a 400 GB NTFS-formatted hard disk) for the tests. The reason for this is that a lot of malware checks for the presence of virtualization products such as *VMware* or *Virtual PC*, and in such cases the malicious software might behave differently. Besides this, the helper tools installed on a guest operating system might be incompatible or cause problems with the rootkits, as they also try to hook critical system functions.

STEP 2: TESTING FOR DETECTION OF INACTIVE ROOTKITS

It is important to check whether the AV products are able to detect the rootkits before installation when they are easy to identify using standard AV techniques such as signature scanning. This will demonstrate the products' ability to block the malware before it can harm the system. This test should be performed both with the on-demand scanner and especially using the on-access guard. If the guard cannot

prevent the download and installation of the rootkit, a proper detection will be much more difficult.

Preparation for this test is straightforward and does not differ from any other tests: one only needs to install the test product on a known clean system (e.g. from a *Windows XP SP2* image file), update the product to the latest available version (this might involve a few reboots), and create an image of this system. Once this has been done, the PC will not need to be connected to the Internet again, and will only be used in a secure test lab environment. The selected samples will then be used to test the products. This only takes a few minutes per tool, including proper documentation and the creation of report files.

Testing web-based online scanners (usually implemented as ActiveX controls or Java applets) is a bit trickier, as these tools require a working Internet connection and update themselves regularly. Special precautions must be taken, such as limiting the Internet connection (so that only the required IP addresses from the AV company's servers and ports can be accessed). Furthermore, the tests of online scanners have to be performed at almost the same time, in order for the products to be in a comparable state. In order to be able to reproduce the test results at a later time, it is a good idea not only to create image files of the system, but also to capture all the Internet traffic and to create screenshots or videos of the entire test, showing each detection and miss in detail.

STEP 3: TESTING FOR DETECTION AND REPAIR OF ACTIVELY RUNNING ROOTKITS

The testing of products against active rootkit samples is actually the 'real' rootkit test, showing how well the products handle hidden objects, not only regarding detection but also with regard to disabling the rootkit and removing all of its components.

There are many different possible scenarios in which a rootkit could enter a system. One is that the computer is not running an AV solution, and another is that the AV product on the system is outdated or doesn't have signatures for the specific version of the rootkit in its database. For our testing, we used the scenario that the AV product is up to date and working, but the on-access protection is turned off, so the rootkit can be installed without any warning messages from the guard. This way, we do not need to install and update the product again and again, which saves a lot of time. Besides this, we can use images of the products for testing, thus making the reproduction of the results (when required) a lot easier, as the same version is used in all cases.

After the malware is executed on the test system, it is important to check whether the rootkit has installed properly and is running. This includes checking that all the files and registry entries that should be present according to our previous analysis are actually present, and that the objects that should be hidden are hidden.

We then turn the guard back on and if anything is detected we let the product perform its cleaning routine (if any). We then perform an on-demand scan using the default settings of the product. Again, if anything is detected we let the product perform the suggested repair routine (if any). The system is rebooted if the tool prompts for this to complete the cleaning operation.

Straight after this, we need to determine whether the rootkit (and the related malware) is still active and find out which components have been removed (or renamed) and which have not been handled. Of course, the job of the AV tool should include the removal of all active traces of the malware, but it should not be considered a fault if some inactive traces, such as harmless text files, are left on the system. Scanner report files and snapshots created before and after the malware execution and cleaning are a good way of documenting the actions of the tool, but we have to be sure that these tools deal properly with the rootkits used during the testing.

For every test run, only one product should be checked against a single rootkit, and afterwards the system must be restored from a clean image file before the next test can start. Testing against active samples usually requires around 20 to 30 minutes per sample, depending on the documentation and quality requirements of the test. So the test of a single product against 60 samples can easily take 20 to 30 hours. As performing such tests requires quite a lot of knowledge and experience, they cannot easily be automated. However, the tasks of the tester can be supported by various self-developed helper tools to make the work easier to perform.

As with the test against inactive rootkits, the testing of online scanners against active samples is more problematic than testing standalone AV products. Once again, the problems include the reproducibility of the results and the fact that a system with actively running malware needs to be connected to the Internet for a short amount of time.

LOOKING AT THE RESULTS

In the case of the *ComputerBild* review on *Windows XP*, all products (in their most current versions) were updated and then frozen on 25 October 2007. The only exceptions were the online scanners, which were tested on 25 October and 2 November 2007.

ComputerBild review (Windows XP Home Edition, 32-bit, SP2) [1]							
Product	**Version**	**Detection of inactive samples**	**Detection of actively running rootkits**	**Detection of malware hidden by rootkits**	**Removal of inactive samples**	**Removal of actively running rootkits**	**Removal of malware hidden by rootkits**
	Reference (max) ->	30	30	30	27	30	30
INTERNET SECURITY SUITES							
Avira AntiVir Premium Security Suite	7.06.00.168	28	29	30	25	7	7
BitDefender Internet Security 2008	11.0.13	30	28	29	27	23	27
Bullguard Internet Security Suite	7.0.0.27	30	7	10	27	4	0
G DATA InternetSecurity 2008	18.0.7227.533	30	9	4	27	7	0
Kaspersky Internet Security 7.0	7.0.0.119	28	24	28	25	22	25
Kaspersky Personal Security Suite V	6.0.2.621	28	21	27	25	19	17
Norton Internet Security 2008	15.0.0.60	25	18	25	25	18	25
WEB-BASED ONLINE SCANNER							
BitDefender Online Scanner	1.0 Build 2422	30	5	3	27	2	0
F-Secure Online Virus Scanner	3.2 Beta (1.0.64)	24	27	26	24	23	23
Kaspersky Online Scanner	5.0.98.1	28	6	21	25	0	0
Microsoft Windows Live Safety Scanner	1.1.3007.0	20	17	25	19	10	8
Panda Security ActiveScan	5.54.01	28	25	26	27	15	26
Trend Micro HouseCall	6.6 (1103-1060)	27	8	5	27	7	1
SPECIALIZED ANTI-ROOTKIT TOOLS							
AVG Anti-Rootkit Free	1.1.0.42	n/a	30	29	n/a	26	27
Avira RootKit Detection	1.0.1.17 Beta	n/a	28	30	n/a	23	28
BitDefender RootKit Uncover	1.0 Beta 2	n/a	24	28	n/a	16	12
F-Secure Blacklight	2.2.1064.0 Beta	n/a	28	28	n/a	20	27
GMER	1.0.13.12551	n/a	30	28	n/a	19	26
IceSword	1.2.2.0	n/a	25	26	n/a	10	6
McAfee Rootkit Detective	1.1.0.0	n/a	26	29	n/a	21	28
Microsoft Rootkit Revealer	1.71.0.0	n/a	15	14	n/a	n/a	n/a
Panda Security Anti-Rootkit	1.07.00	n/a	24	28	n/a	22	27
Rootkit Unhooker LE	3.7.300.509	n/a	30	30	n/a	22	28
Safe'n'Sec Pro	3.0.0.4104	n/a	18	9	n/a	7	3
Sophos Anti-Rootkit	1.3.1 (1.07)	n/a	26	26	n/a	17	24
System Virginity Verifier	2.3	n/a	15	3	n/a	10	3
Trend Micro Rootkit Buster	1.6 Beta	n/a	30	29	n/a	20	24

We first checked the products' on-demand detection and removal of inactive samples. This already revealed some missing signatures in the scanners' databases. The results of the on-access scanning were identical to the on-demand results, so they are not listed separately in the results table. None of the dedicated anti-rootkit tools we tested had an integrated on-demand scanner, so no results are available in this category. The maximum number of samples the tools could detect was 30 dedicated rootkits, and no more than 27 rootkits could be removed because we used the original (and thus, write-protected) CD and DVD media with the three 'commercial' rootkits.

The test with 30 active rootkits and 30 items of other malware using rootkit technologies was a lot more challenging both for the testers and the products. On average, the specialized anti-rootkit utilities were able to detect around 80% of the test samples. The security suites detected a little more than 66% of the rootkit infections and the online scanners performed the worst, with a detection rate of just 53%. We encountered significant problems in several cases in which the tools either crashed or hung during or after finishing a scan (in these cases we counted the rootkit as not detected).

Rootkit removal proved even more problematic. Once again the specialized tools performed the best on average, with a disinfection score of a little below 66% of the samples. However, the security suites were not able to clean more than 50% of the infections and the online scanners were almost useless, with a disinfection rate of only around 32%.

We also saw a good number of crashes and related problems in this section, but sometimes the rootkit was gone after a bluescreen and one or two reboots. Tools like *Avira RootKit Detection* sometimes removed the *Windows* explorer.exe file, so the system could not be started after a 'successful' disinfection run. *McAfee Rootkit Detective* renamed the original *Internet Explorer* iexplore.exe file in two cases. Sporadically, *AVG Anti-Rootkit Free* also tried to remove some system files, leaving the system in an unbootable state. (Note: this list of problems is not comprehensive.)

The *c't* magazine review on *Windows Vista* only included 'pure' anti-virus programs. The tools were last updated and frozen on 2 October 2007. To our surprise, the detection rate of inactive samples reached just 90% on average, even though most of the rootkits used were

c't review (Windows Vista Ultimate Edition, 32-bit, RTM) [2]				
Product	**Version**	**Detection of inactive samples**	**Detection of actively running rootkits**	**Removal of actively running rootkits**
	Reference (max) ->	6	6	6
Avast! Antivirus Professional Edition	4.7.1043 (000778-1)	6	3	3
AVG Anti-Malware	7.5.488 (269.13.37 / 1042)	6	0	0
Avira Antivir PersonalEdition Premium	7 Build 308 (7.06.00.18)	4	6	3
G Data AntiVirus 2008	18.0.7227.533 (8434 / 393)	6	3	3
BitDefender Antivirus 2008	11.0.0.25 (7.15077)	6	5	5
CA Anti-Virus Plus 2008	4.0.0.130 (31.1.0 / 5178)	6	6	4
ClamWin Free Antivirus	0.91.2 (4 / 4452)	3	3	1
Dr Web Antivirus für Windows	4.44.0.09170	2	2	2
F-Secure Anti-Virus 2008	6.80.2610.0 (2007-10-02_01)	6	6	6
Ikarus virus.utilities	1.0.60 (1.1.13)	6	2	1
Kaspersky Anti-Virus	7.0.0.119	6	6	2
McAfee VirusScan 2008	11.2.121 (5100-5131)	6	2	2
Microsoft Windows Live OneCare	1.6.2111.32 (1.1.2803.0)	5	1	1
Eset Nod32 Antivirus	2.70.39.0 (10902)	5	5	5
Norton Antivirus 2008	15.0.0.58	6	6	6
Panda Security Antivirus 2008	3.00.00 (2.1.29.0)	6	6	6
Trend Micro Antivirus + Antispyware 2008	16.00.1413 (8.500-4.752.90)	6	5	5

released during 2005 and 2006. Only four of the six installed rootkits could be detected by an average tool and the cleaning rate was even lower with 54%. *AVG* (with one of the best standalone tools on *Windows XP*) performed poorly with no detection or cleaning of running rootkits on *Vista*. Tools from *Microsoft*, *Ikarus* and *Doctor Web* also demonstrated the need for some significant improvements on this platform.

CONCLUSION

Tests of the active rootkit detection and cleaning features of anti-malware products are rather time consuming and require a lot of resources to perform. However, programmers and testers should dedicate more attention to these features, as most AV tools still perform poorly in this area. Without proper anti-rootkit features a protection program may give the user the wrong impression about the status of his PC.

A step in the right direction could be to focus on providing bootable rescue media, too: this might be the product installation CD or a CD or disk that a user can create and update himself [5, 6]. When the system is started from this media, the rootkit cannot be activated on the system, so a scanner would be able to see all files and registry entries which would usually be hidden. This way, the scanner could detect and delete all rootkit and malware components as long as the signature database is up to date and comprehensive.

REFERENCES

[1] Russinovich, M. Inside Sony's rootkit. Virus Bulletin, December 2005, p.11. http://www.virusbtn.com/virusbulletin/archive/2005/12/vb200512-sonys-rootkit.

[2] Melfsen, T.; Badenius, F.; Pursche, O. Wurzelbehandlung (Root Treatment). ComputerBild 26/2007, pp.74–95, Axel Springer Verlag, Hamburg, Germany. http://www.computerbild.de/.

[3] Knop, D.; Schmidt, J. Auf der Pirsch: 17 Antivirenlösungen unter Windows Vista und XP. (Going Hunting: 17 AV solutions for XP and Vista). c't 01/2008, pp.92–103, Verlag Heinz Heise, Hanover, Germany. http://www.heise.de/.

[4] Florio, E. Stories from the DRM world: the Settec case. Virus Bulletin, April 2006, p.10. http://www.virusbtn.com/virusbulletin/archive/2006/04/vb200604-smith.

[5] Marx, A. Rescue Me: Updating Anti-Virus Rescue Systems. Virus Bulletin May 2002, pp.10–12. http://www.virusbtn.com/pdf/magazine/2002/200205.pdf.

[6] Marx, A. Rescue Me 2: Disinfection With Bootable Rescue Media. Virus Bulletin March 2004, pp.14–16. http://www.virusbtn.com/pdf/magazine/2004/200403.pdf.

COMPARATIVE REVIEW

VB100 MARCH 2008 – WINDOWS VISTA BUSINESS EDITION SP1

John Hawes

Windows Vista makes another appearance on the *VB* test bench just over a year after its debut (see *VB*, February 2007, p.14) and eight months after a slightly less well-attended review of products on its 64-bit version (see *VB*, August 2007, p.16). During that time the platform has failed to make enormous headway in the marketplace, with most estimates reckoning it resides on at most 10% of the world's desktops, languishing far behind the dominant *XP*, which is thought still to hold sway on around 75% of systems.

The release of the first service pack for an operating system is often seen as a sign of maturity though, and SP1 for *Vista* could signal an upturn in the uptake of the platform. The upgrade promises a raft of improvements to performance and general functionality.

SP1 was released shortly after the deadline for product submissions for this review. With at least some participants not yet able to get their hands on a copy for testing, even more than the usual number of bugs and unpredictable behaviours were expected, and with one of the largest sets of products yet seen on the VB100 test bench, I anticipated an arduous slog through the tests this month.

Initially a total of 40 products were submitted on the February 28th deadline. Many of these were new to the tests, including a pair of products based on the open-source *ClamAV* detection technology, which promised to provide some interesting results. Several others returned after lengthy absences, and all the usual suspects were also present. With such a huge number of products to get through I decided that any which could not be made to provide usable results after the standard three installs would have to be shelved.

I also decided to streamline the results reporting process somewhat, if only to make the figures readable on the page. Thus, exact numbers of missed samples will not be reported in this month's comparative. As always, the percentages listed in the detection table represent the number of variants covered (or not) by the products, rather than the number of unique samples.

PLATFORM AND TEST SETS

Installation and set-up of *Vista* has become a fairly familiar process, and while the prettification of the install process itself always impresses me, my first act on seeing the garish, flashy interface is invariably to roll it back to its 'classic' appearance, which seems much less intrusive over long periods of exposure. Doubtless the glitzy 'Aero' look can be tweaked into something less nauseating, but this was not a process I was ready to spend any time on.

Application of the service pack was less painful than I had feared. Once installed, it seemed to make no obvious difference to the way things ran, although I did notice considerably fewer crashes and blue-screens than in previous tests. As in those earlier runs a user with minimal rights was created, to measure integration with user access controls and so on, but I expected to have to switch to the admin user or even disable UAC entirely for some products.

The clean test sets saw a fairly substantial enlargement, with most of the sets used for speed measurement now fast approaching a good size for freezing. This will enable more useful comparisons of product speeds over time as platforms are revisited (assuming the test hardware manages to withstand the heavy usage). An especially large number of additions were made to the polymorphic sets, thanks to there having been numerous misses in those areas in recent months – the variants of W32/Virut which have already been known to cause widespread problems are now more fully represented, and this should continue to tax the participants' detection abilities to the utmost.

The deadline for the test sets was February 25th, and with the January WildList having been available for a few weeks by then the WildList set was synchronized with that list. This meant a fairly large number of additions, but as the previous list saw some clearing out of older items the total number of variants in the core set remained around the usual level. The additions were dominated as usual by worms and bots such as W32/Ircbot and W32/Agent, with another handful of variants of W32/Virut and other nasty file infectors also joining the set for the first time.

I had hoped that this month would see the inclusion of a preliminary set of non-replicating malware, having spent some time gathering, validating and categorising new samples for this purpose. As a step towards this significant change, much of the older part of the test set has been moved aside into a 'legacy' set, as a precursor to the eventual retirement of all older items. With the unexpectedly high turnout of products, and some unforeseen delays in getting the lab ready for testing, the final stages of implementing the new trojan set had to be put on hold. Pared-down selection was included in the test out of interest and to gather data for name-referencing, and a further expanded version of the set should be ready in time for the next review. Even without this new challenge, the potential both for false positives and incomplete detection of polymorphic items was increased considerably and it

Detection rates on demand (OD) and on access (OA)	WildList viruses		Worms & bots		File infector viruses		Polymorphic viruses		Legacy samples		Clean sets	
	OD	OA	OD	OA	OD	OA	OD	OA	OD	OA	FP	Susp.
AEC Trustport Antivirus	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	86.39%	86.39%	97.62%	97.62%		
Agnitum Outpost Security Suite Pro	99.99%	99.99%	100.00%	100.00%	99.21%	99.21%	79.29%	79.29%	99.98%	99.98%		2
Ahnlab V3	100.00%	99.90%	99.76%	99.76%	99.21%	99.21%	88.06%	88.06%	96.77%	96.77%		
Alwil avast!	99.83%	99.83%	100.00%	100.00%	100.00%	100.00%	86.20%	86.20%	97.02%	97.02%	1	
AVG	100.00%	100.00%	100.00%	100.00%	98.43%	98.43%	73.89%	73.89%	97.63%	97.63%		
Avira AntiVir	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	84.68%		
BitDefender AntiVirus 2008	99.996%	99.996%	100.00%	100.00%	98.95%	98.95%	100.00%	100.00%	99.59%	99.59%		
Bullguard	99.996%	99.996%	100.00%	100.00%	98.95%	98.95%	100.00%	100.00%	99.59%	99.59%		
CA eTrust Antivirus	100.00%	100.00%	100.00%	100.00%	99.84%	99.84%	99.70%	99.70%	99.10%	99.10%		
CA Internet Security	100.00%	100.00%	100.00%	100.00%	99.84%	99.84%	99.70%	99.70%	99.10%	99.10%		
Check Point Zone Alarm	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.88%	99.88%	100.00%	100.00%		
Doctor Web Dr.Web	95.21%	95.21%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%		
ESET NOD32 Antivirus	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%		
Fortinet FortiClient	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%		
Frisk F-PROT Antivirus	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%		
F-Secure Client Security	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.88%	99.88%	100.00%	100.00%		2
G DATA AntiVirus 2008	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.88%	99.88%	100.00%	100.00%		
Hauri ViRobot	99.50%	99.50%	98.42%	98.42%	94.49%	94.49%	94.96%	94.96%	65.00%	65.00%		
Ikarus Virus Utilities	99.84%	99.84%	99.80%	99.80%	95.67%	95.67%	72.85%	72.85%	71.30%	71.30%	6	
K7 Total Security	99.93%	99.93%	99.53%	99.53%	97.32%	97.32%	59.45%	59.45%	77.91%	77.91%	2	2
Kaspersky Anti-Virus	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.88%	99.88%	100.00%	100.00%		
Kingsoft Internet Security 2008	100.00%	100.00%	35.00%	35.00%	75.00%	75.00%	48.00%	48.00%	56.00%	56.00%		
McAfee VirusScan Enterprise	99.998%	99.998%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%		
Microsoft Forefront Client Security	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	95.00%	95.00%	99.99%	99.99%		
Microsoft Windows Live OneCare	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	95.00%	95.00%	99.99%	99.99%		
MWTI eScan Internet Security	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.88%	99.88%	100.00%	100.00%		
Norman Virus Control	100.00%	100.00%	100.00%	100.00%	99.05%	98.43%	73.77%	70.23%	98.95%	98.95%	1	
PC Tools AntiVirus	99.99%	99.99%	100.00%	100.00%	99.21%	99.21%	79.29%	79.29%	99.98%	99.98%		2
Quick Heal Anti-Virus Lite	100.00%	100.00%	100.00%	100.00%	98.43%	98.03%	83.86%	83.86%	96.40%	96.40%	2	
Redstone Redprotect	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.88%	99.88%	100.00%	100.00%		
Rising Antivirus	99.97%	99.97%	99.73%	99.73%	93.70%	93.70%	44.63%	44.63%	56.40%	56.40%	1	
Security Coverage PC Live	84.35%	76.00%	56.00%	53.00%	97.20%	93.51%	54.00%	49.00%	47.00%	43.00%	1	
Sophos Anti-Virus	99.997%	99.997%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.80%	99.80%		
Symantec Norton AntiVirus	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%		
Trend Micro Internet Security	99.99%	99.99%	100.00%	100.00%	99.21%	99.21%	80.41%	80.41%	98.85%	98.04%	2	
VirusBuster Professional	99.99%	99.99%	100.00%	100.00%	99.21%	99.21%	79.29%	79.29%	99.98%	99.98%		2
Webroot Spy Sweeper with Antivirus	99.997%	99.997%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	99.98%	99.98%		

looked likely to be a tough month for the large crowd of products I dragged into the *VB* test lab.

AEC Trustport Antivirus 2.8.0.3001

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	97.62%
File infector	100.00%	**Polymorphic**	86.39%
False positives	0		

After a lengthy spell atop the scoreboard with near immaculate detection rates thanks to an intensive multi-engine design, *Trustport*'s performance dipped a little in recent tests which coincided with the product having a new set of engines under its bonnet. The version provided for this test was different again, with only the *Norman* and *AVG* engines in use this time. The look and feel of the product also seemed a little different, with the interface laid out in a pleasant and logical fashion, with access to the full range of configurations users of such a serious product should expect. The installation process was straightforward, requiring the administrator password and a reboot to complete.

In default mode, the 'intensive' settings option was selected, with all files and most archive types checked in both modes. As a result, scanning speeds were less than stellar and on-access overheads were on the high side. However, the product remained stable even under heavy bombardment and the system remained responsive throughout. Detection rates were solid, and with the WildList covered without problems *Trustport* regains its VB100 certified status.

Agnitum Outpost Security Suite Pro 6.0.2282.253.0485

ItW	99.99%	**Worms & bots**	100.00%
ItW (o/a)	99.99%	**Legacy**	99.98%
File infector	99.21%	**Polymorphic**	79.29%
False positives	0		

Agnitum's suite features a diverse set of security functions, and under *Vista* the installer requires both the administrator password and confirmation that the user does indeed want to install the software. After the obligatory reboot a message alerted me that a driver had failed to load, warning that functionality may be impaired, and the on-access scanner did seem to be working erratically. After a second reboot the error and related instability disappeared and all seemed to operate properly.

With a colourful and easy-to-use interface and a decent selection of options, testing proved a pleasure. Reporting was a little odd at times though, not least in the set where the Eicar test file was used to measure the depth of archive scanning – while in most types of archive the test file was described as 'malware', when stored in .tgz format it was labelled a virus and given a higher risk rating. This minor quibble aside, detection was pretty good throughout, with a few suspect packed files pointed out in the clean sets but no false positives. In the WildList set, however, a handful of samples from one of the newly expanded W32/Virut sets were missed, denying *Agnitum* a VB100 award this time.

Ahnlab V3 Internet Security 7.0 Platinum Enterprise

ItW	100.00%	**Worms & bots**	99.76%
ItW (o/a)	99.90%	**Legacy**	96.77%
File infector	99.21%	**Polymorphic**	88.06%
False positives	0		

Ahnlab's latest suite is a good-looking beast – glossy without being flashy – and installs simply, with the administrator password required but no reboot necessary. The main interface is similarly straightforward, offering some basic configuration – enough at least to get through the *VB* testing process without difficulty.

Logging seemed a little odd until I realised it had a hard-coded size limit for the log file, which meant that large chunks of precious detection records were being discarded before they could be saved. A somewhat laborious process of splitting the test sets up into several smaller chunks and scanning each separately got around this problem.

The product showed fairly good detection rates. A smattering of misses were recorded on access in the WildList set, thanks to some file extensions which are not scanned by default but are used by some worm variants for spreading. Thus, despite achieving full detection in manual scans and managing to avoid false positives across the expanded clean sets, *Ahnlab* does not reach the required standard for a VB100 award on this occasion.

Alwil avast! 4.7.1098

ItW	99.83%	**Worms & bots**	100.00%
ItW (o/a)	99.83%	**Legacy**	97.02%
File infector	100.00%	**Polymorphic**	86.20%
False positives	1		

Alwil's submission was troubled by some missing data, the problem being given away by the product running suspiciously quickly through the on-access tests. Once the correct files were in place, after an installation process which required no password for the normal user but did insist on a reboot, quite the opposite result was obtained. Speeds remained excellent over the clean sets but slowed right down when faced with large numbers of infected files. The entire system became bogged down in the process, and after leaving it to finish overnight another reboot was needed to pull itself together. Some investigation hinted that this was perhaps a result of the product deleting each infected file without prompting, despite the 'interactive' setting being selected. On demand things were less troublesome, with the action dialog appearing as usual with its friendly 'apply to all' option, and these scans sped through at a lightning pace.

The dual interface system has never been a favourite of mine, but here it seemed steady and responsive except during the on-access incident mentioned above, which does not reflect any likely real-world situation. Settings seemed plentiful, with defaults ignoring archives but full scanning available for those who want it, and detection rates were reasonable as ever. However, a set of file infectors recently added to the WildList set were missed, and a single item in the clean set was mistakenly flagged as malware, and thus *Alwil* misses out on a VB100 by a whisker.

AVG 7.5.516

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	97.63%
File infector	98.43%	**Polymorphic**	73.89%
False positives	0		

AVG (formerly known as *Grisoft*) provided the shiny new version of its product for a recent standalone review (see *VB*, March 2008, p.18), and it was with some disappointment that the earlier incarnation was received for this month's test, the new interface having impressed me considerably. This one required the administrator password to install as well as when changing settings, which

seemed a sensible way to go about things. The available configuration, once I had managed to refamiliarise myself with its somewhat arcane layout, appeared to offer the option to scan all file types on access. However, this did not seem to cover archives, or even files with unusual extensions, which remained resolutely undetected. The settings also seemed to revert after a reboot, and having shut down the system over a long weekend I found, much to my annoyance, that the same scan I had run previously was now eating up my test collection.

Perhaps aided by the fact that it ignored many file types, speeds were excellent, and detection rates were also pretty good. With nothing missed in the Wild and no false positives, *AVG* earns its first VB100 award under its new company name.

Avira AntiVir Windows Workstation 7.06.00.507

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	100.00%
False positives	0		

Avira's *AntiVir* is another product that is familiar from many previous tests, this one offering a much more rational interface and a great flexibility of configuration. Defaults seemed sensible and increasing the range of items scanned produced impressively thorough results, without adding greatly to the excellent speeds.

Detection was similarly splendid, with not much missed across all the sets on demand. A strange anomaly on access left large swathes of the legacy set not spotted, but the items were detected without problems on demand and the oddity did not extend to the new sets. The WildList was fully covered and no false positives were generated in the clean sets, thus *Avira* wins another VB100 award.

BitDefender AntiVirus 2008

ItW	99.996%	**Worms & bots**	100.00%
ItW (o/a)	99.996%	**Legacy**	99.59%
File infector	98.95%	**Polymorphic**	100.00%
False positives	0		

BitDefender's 2008 product is another that has been subjected to an in-depth review in *VB* (see *VB*, September 2007, p.17) as well as repeated entries in *VB* comparative reviews, and thus proved familiar and simple to use. The interface achieves a happy balance between friendly straightforwardness for the less advanced user, displaying large and comforting green ticks and assurances that the system is protected, and in-depth configuration for those with more individual requirements.

Zipping through the speed tests in good time and getting near-perfect scores in the infected sets, it was only a pair of samples from the expanded set of W32/Virut variants that stood between *BitDefender* and a VB100 award.

Bullguard 8.0

ItW	99.996%	**Worms & bots**	100.00%
ItW (o/a)	99.996%	**Legacy**	99.59%
File infector	98.95%	**Polymorphic**	100.00%
False positives	0		

Bullguard requires the administrator password and a reboot before presenting its colourful interface and friendly 'Welcome to your Bullguard' sense of security. The version submitted for testing was presumably a trial version without a licence provided, and trying to run an on-demand scan brought up another demand for the admin password followed by the sad announcement that the scan could not continue as the licence had expired. Tests were able to proceed however, thanks to a right-click scanning option which mercifully remained functional.

With more thorough defaults on access than *BitDefender*'s own implementation of the same engine, speeds were pretty similar and detection levels likewise highly impressive with little missed in any of the sets. With no false positives only the two missed polymorphic samples in the WildList set prevented *Bullguard* earning itself another VB100 award.

CA eTrust Antivirus 8.1.637.0

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	99.10%
File infector	99.84%	**Polymorphic**	99.70%
False positives	0		

CA's corporate product remains unchanged once again, presenting the same interface for yet another test despite a few minor changes behind the scenes. The installation went through the traditional run of lengthy EULAs, personal data gathering and complex activation codes, but with the benefit of much experience the product was quickly up and running.

Experience is of little help once the fiddly and frustratingly slow-to-respond interface comes into play, but the tests proceeded despite a lack of available configuration, and at the usual excellent speed. Although the interface appeared to provide the option to scan inside archives there was no evidence of this actually happening.

Logging was awkward, and trying to view logs of any significant size from within the product brought about a grinding freeze. As usual the logs were taken off the system and converted into normal text, dropping much of their extraneous content, for parsing. Results showed an absence of false positives and pretty decent detection, and with nothing missed in the WildList *eTrust* adds another VB100 award to its tally.

CA Internet Security 4.0.0.172

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	99.10%
File infector	99.84%	**Polymorphic**	99.70%
False positives	0		

CA's home-user product also eschews the admin password and simply requires a 'continue' button to be clicked to install, as well as each time the configuration is changed. This was not often though, as little configuration was available, and again archives could not be scanned internally on access; speeds were extremely impressive.

Detection also kept up with the corporate product, and false positives were absent again. After several detections in a small test scan, identical alert popups appeared repeatedly – over a dozen times for a single detection – but this behaviour did not recur, thankfully, while scanning the full test sets. With nothing missed in the wild and no false positives, *CA* notches up its second award of the month.

Check Point Zone Alarm

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	99.88%
False positives	0		

Since I joined *VB* some time ago, *Zone Alarm* has been one of the highest-profile products not to appear in our tests. Frequent queries from readers have diligently been followed up with attempts to contact the vendor, and at last the good people at *Check Point* have seen fit to submit a product for testing. I installed it with excitement.

The installation process was a little complex, thanks to the test lab being isolated from the web to ensure all products are tested on a level field. Once the installer had been run, using the admin password, and detection data had been replaced in safe mode, the product finally presented an interface that was little changed from the one I had grown accustomed to many years ago when running the free version of the firewall on a home system. In the modern setting of *Windows Vista* (even without the fancy Aero stylings), it looks a little dated and perhaps in need of a restyle, but there's much to be said for sticking with the tried and trusted.

The suite includes intrusion prevention, email filtering and spyware monitoring as well as the anti-virus and firewall, and little room has been left in the interface for in-depth configuration of the virus scanning. There was, however, enough control to get through our tests, and detection, provided by the *Kaspersky* engine, was as excellent as one would expect.

Scanning times were somewhat slow on demand, quite spectacularly so over the archive set, but not unacceptable on access, and the product seemed to run stably with no noticeable impact on system performance. With nothing missed in the WildList set, and precious little elsewhere, and nothing more than a few warnings of possibly unwanted items in the clean set, *Check Point*'s product wins itself its first VB100 award at the first attempt, without even breaking into a sweat.

Doctor Web Dr.Web 4.44.4

ItW	95.21%	**Worms & bots**	100.00%
ItW (o/a)	95.21%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	100.00%
False positives	0		

Doctor Web's product requires its installer to be run with full administrator rights, and requires the password again for updating and changing settings. The product is split into two quite separate parts, the scanner and the on-access component, 'SpIDer Guard'. The latter seemed to be having some difficulties operating on this occasion, but disabling *Windows Defender*, the security tool which runs by default under *Vista*, put an end to this problem.

Scanning was thorough rather than speedy, with thoroughness particularly evident on the archive sets. Detection rates were as good as ever, and false positives absent, but once again a handful of the latest additions to the WildList set were not detected, and *Doctor Web* misses out on a VB100 award.

ESET NOD32 Antivirus 3.0.642.0

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	100.00%
False positives	0		

The recent overhaul of *ESET*'s product, both cosmetically and at a deeper level, has added considerably to *NOD32*'s charms. The installation was one of the most straightforward, with a simple 'continue' prompt and no reboot, and with ample configuration and typically zippy scanning speeds, testing took barely any time at all.

Detection was as flawless as usual, with nothing missed across any of the sets, and *ESET* continues its lengthy run of success with another VB100 award.

Fortinet FortiClient Host Security 3.0.470

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	100.00%
False positives	0		

FortiClient has a rather more complex installation process thanks to its multi-function nature, and requires confirmation of numerous drivers as the process chugs along. Once it is installed, much of the configuration is greyed out for the normal user and requires the interface be opened with the 'Run as administrator' option to allow the settings to be changed, although no password is required to access this.

Considering the thoroughness of the scanning, speeds were surprisingly good. At the end of one large scan the product froze and had to be shut down forcibly, but this was the only significant problem encountered – detection was splendid, false positives absent, and *Fortinet* wins itself another VB100 award.

Frisk F-PROT Antivirus for Windows 6.0.8.1

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	100.00%
False positives	0		

The version of *F-PROT* submitted for testing described itself as a beta version and advised would-be users only to install it on test systems. This didn't put me off, and I proceeded merrily with the install, pausing only to enter the administrator password. The interface presented is a simple, pared-down little thing, not heavy on the configuration options, but it runs solidly and smoothly. My only quibble with it would be that the scanning page seems to snag in one place after scans, and must be clicked away from and back to in order to access the controls again.

Speeds were splendid and detection excellent, and with no significant misses in the infected sets and no false positives, *Frisk* qualifies for a VB100 award this month.

F-Secure Client Security 7.11.107

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	99.88%
False positives	0		

F-Secure's installation process has the full range of requirements – passwords, licence codes, a reboot, and admin rights are required to run the offline update. It even seems to limit the configuration controls for less privileged users, and it offers the option of installing a standalone or remotely managed version from the off.

Scanning seemed fairly fast, at least until more thorough options are selected, but there did seem to be some noticeable slowing of the system. This impact was made up for by the superb detection which, coupled with a lack of false positives earns *F-Secure* a VB100 award once again.

G DATA AntiVirus 2008 18.4.8051.821

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	99.88%
False positives	0		

G DATA seems to have dropped the cosy 'AVK' from its latest product names, but little else has changed in this multi-engine powerhouse. Installation needed no password, but a 'continue' prompt appeared during the install as well as each time the

options page was visited and whenever a manual scan was initiated. After the install, the product requested a reboot, but with the prompt hidden behind another window it wasn't noticed for some time; this didn't seem to affect the running of the product in any way, however, and testing progressed with much success before it was spotted.

The interface is attractive and well designed, with none of my specialist requirements absent or hard to find. Logging, however, has long posed a problem, with details spread over multiple lines requiring some extra tinkering to extract data. This time the headache was increased by the log file taking a long time to open. Once acquired and processed however, detection proved to have been almost impeccable, and with nothing missed in the wild and no false positives *G DATA* once again qualifies for a VB100 award.

Hauri ViRobot Desktop 5.5

ItW	99.50%	**Worms & bots**	98.42%
ItW (o/a)	99.50%	**Legacy**	65.00%
File infector	94.49%	**Polymorphic**	94.96%
False positives	0		

Hauri has been absent from the *VB* test bench for some time, its last appearance also being the last test conducted by my predecessor (see *VB*, June 2006, p.11). After this lengthy break the *ViRobot* product returns to the fold, bravely, just in time for one of the toughest tests for a while.

The installation, which required no password, mentioned that some technology provided by *BitDefender* was included in the product – which promised good things. The interface looked well designed and ran in a solid and stable manner; configuration was ample and well presented. One option which was conspicuously absent was a control to deactivate the beep made each time a detection is recorded on access – I had to flee the lab to escape the barrage of sounds as the full test set flooded past the scanner.

Archive detection seemed sensible, with no scanning on access and a maximum of five levels on demand, and speeds were comfortably in the mid-range. Logging took an enormous amount of time to deal with. Saving the log file left me gazing forlornly at an egg timer for ages only for a log to be produced in the most bizarre format I have ever encountered, and one which required considerable hacking to render it readable. Scanning results revealed an absence of false positives, but several misses in the WildList set and less than perfect coverage elsewhere. Consequently, *Hauri* does not quite reach the required standard for a VB100 award this time.

Ikarus Virus Utilities 1.0.61

ItW	99.84%	**Worms & bots**	99.80%
ItW (o/a)	99.84%	**Legacy**	71.30%
File infector	95.67%	**Polymorphic**	72.85%
False positives	6		

Plucky *Ikarus* continues to fling itself at the walls of the VB100 fortress, despite repeated knock-backs in recent tests. Each time the product has seemed to improve in stability and detection rates, but this trend slowed somewhat this month. After a solid start, with some sensible integration into the UAC system requiring administrative rights to install and alter settings, the interface seemed reluctant to open over several attempts, and once running provided the usual dearth of options.

The .NET-based GUI also suffered some shakiness during scanning, with flickering and ghostly whiting out not uncommon, especially when scanning large sets. A limitation on the size of logs meant on-demand scans had to be carried out in smaller chunks, and once all the information presented was processed the results showed several file-infecting viruses not fully covered in the WildList set. With a smattering of false positives as well, *Ikarus*' quest for VB100 certification must continue another day.

K7 Total Security 9.5.0502

ItW	99.93%	**Worms & bots**	99.53%
ItW (o/a)	99.93%	**Legacy**	77.91%
File infector	97.32%	**Polymorphic**	59.45%
False positives	2		

K7 achieved VB100 certification at its first attempt a little under a year ago, but then disappeared from our radar for some time, missing out on several less arduous tests only to return in time for a tricky platform. The product handled the new surroundings with some ease however, requiring the admin password and a reboot to get running, and asking for a username and email address to keep in touch with users, before presenting a clear and stable interface.

UAC was again integrated with the product, with on-access controls disabled for unprivileged users, although oddly anyone has the power to disable on-access scanning completely. Scanning speeds were excellent and though configuration was not available in any great depth the defaults seemed sensibly chosen. Detection was no more than reasonable across the sets, however, with a handful of tricky polymorphic items missed in the WildList set and a clutch of false positives in the clean set spoiling *K7*'s immaculate VB100 record.

Kaspersky Anti-Virus 7.0.1.325

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	99.88%
False positives	0		

Kaspersky has a much longer and more illustrious record, producing consistently excellent detection rates and unimpeachable standards of design and implementation. As usual the product proved thorough rather than speedy, but still produced perfectly acceptable times even over archive sets, and powered through the infected sets with little difficulty.

With nothing missed in the WildList and no false positives, *Kaspersky* easily wins another award for its collection.

Kingsoft Internet Security 2008.2.22.11

ItW	100.00%	**Worms & bots**	35.00%
ItW (o/a)	100.00%	**Legacy**	56.00%
File infector	75.00%	**Polymorphic**	48.00%
False positives	0		

Kingsoft has a single VB100 award under its belt, gained in last year's 64-bit *Vista* test. The product has shown some decent detection levels in the newer sets, and presents a slick and professional-looking interface, but has on occasion been a little inconsistent in its scanning behaviour.

This was another of those occasions, with an initial install seeming to miss well over half the samples in most sets. Assuming there were some problems, the installation was re-run with full admin rights (a password had been required as a normal user) and things seemed to go somewhat more smoothly. After several re-runs and re-scans, the product managed to squeeze out some reliable results, with the WildList samples covered in full but a surprising number of misses still evident in the set of worms, many of which have been detected by the product in the past. Nevertheless, *Kingsoft* scrapes its way to a second VB100 award.

McAfee VirusScan Enterprise 8.5.0i

ItW	99.998%	**Worms & bots**	100.00%
ItW (o/a)	99.998%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	100.00%
False positives	0		

McAfee is another old-timer with a long history of success in VB100 testing. The product remains little changed and its simple, dependable style always makes it a welcome visitor to the test lab. The UAC integration is solid, as one would expect from an enterprise-class product, with passwords required to install and on-access controls pared down for non-administrative users.

Everything ran solidly and well with no difficulties caused by the new platform. Speeds were decent and detection rates dependably excellent, until another single sample of the W32/Virut strain which has caused a few upsets already this month reared its ugly head, and since it was in the WildList set was enough to deny *McAfee* a VB100 award this time.

Microsoft Forefront Client Security 1.5.1937.0

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	99.99%
File infector	100.00%	**Polymorphic**	95.00%
False positives	0		

Another giant company, *Microsoft* provides two anti-malware products that ooze professional attention to detail and solidity, and should have been well tested on the new service pack for *Vista*. Both products contrast starkly with the previous offerings however, in their minimal flexibility. Configuration is barely present – the few choices that are available require the administrator password to access them.

Logging is another area which is kept to a minimum. An uncooperative 'History' page lingered regularly for long periods before opening and often showed no detections despite having recently scanned large infected test sets. All data had to be extracted, with some difficulty, from the *Windows* event log, but when finally checked over it showed detection was very good. With the WildList fully covered *Forefront* qualifies for a VB100 award.

Microsoft Windows Live OneCare 2.0.2500.22

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	99.99%
File infector	100.00%	**Polymorphic**	95.00%
False positives	0		

On Demand Throughput (MB/S)	Archive Files				Binaries and System Files				Media and Documents				Other File Types			
	Default Settings		All Files		Default Settings		All Files		Default Settings		All Files		Default Settings		All Files	
	Time	Through-put	Time	Through-put	Time	Through-put	Time	Through-put	Time	Through-put	Time	Through-put	Time	Through-put	Time	Through-put
AEC Trustport Antivirus	1258	3.1	1258	3.1	639	6.1	639	6.1	147	12.2	147	12.2	220	4.2	220	4.2
Agnitum Outpost Security Suite Pro	1777	2.2	1777	2.2	448	8.7	448	8.7	122	14.7	122	14.7	100	9.3	100	9.3
Ahnlab V3	1350	2.9	1350	2.9	606	6.4	606	6.4	82	21.8	82	21.8	144	6.4	144	6.4
Alwil avast!	32	123.8	1066	3.7	325	11.9	335	11.6	39	45.9	96	18.7	33	28.0	74	12.5
AVG	1828	2.2	1828	2.2	576	6.7	576	6.7	97	18.5	97	18.5	125	7.4	125	7.4
Avira AntiVir	920	4.3	920	4.3	142	27.3	142	27.3	45	39.8	45	39.8	32	28.9	32	28.9
BitDefender AntiVirus 2008	1506	2.6	1506	2.6	526	7.4	526	7.4	89	20.1	89	20.1	95	9.7	95	9.7
Bullguard	1967	2.0	1967	2.0	560	6.9	560	6.9	99	18.1	99	18.1	112	8.3	112	8.3
CA eTrust Antivirus	841	4.7	841	4.7	103	37.6	103	37.6	46	38.9	46	38.9	38	24.3	38	24.3
CA Internet Security	1257	3.2	N/A	N/A	130	29.8	130	29.8	53	33.8	53	33.8	44	21.0	44	21.0
Check Point Zone Alarm	12271	0.3	12271	0.3	308	12.6	308	12.6	598	3.0	598	3.0	657	1.4	657	1.4
Doctor Web Dr.Web	5226	0.8	5226	0.8	210	18.5	210	18.5	169	10.6	169	10.6	181	5.1	181	5.1
ESET NOD32 Antivirus	1322	3.0	1322	3.0	801	4.8	801	4.8	42	42.7	42	42.7	59	15.7	59	15.7
Fortinet FortiClient	561	7.1	561	7.1	708	5.5	708	5.5	39	45.9	39	45.9	71	13.0	71	13.0
Frisk F-PROT Antivirus	287	13.8	287	13.8	455	8.5	455	8.5	43	41.7	43	41.7	36	25.7	36	25.7
F-Secure Client Security	3192	1.2	3362	1.2	342	11.3	340	11.4	49	36.6	110	16.3	33	28.0	126	7.3
G DATA AntiVirus 2008	2947	1.3	3416	1.2	543	7.1	544	7.1	133	13.5	160	11.2	123	7.5	133	7.0
Hauri ViRobot	811	4.9	811	4.9	735	5.3	735	5.3	103	17.4	103	17.4	131	7.1	131	7.1
Ikarus Virus Utilities	186	21.3	N/A	N/A	436	8.9	436	8.9	66	27.1	66	27.1	116	8.0	116	8.0
K7 Total Security	276	14.4	N/A	N/A	185	21.0	185	21.0	30	59.7	30	59.7	30	30.8	30	30.8
Kaspersky Anti-Virus	2658	1.5	2658	1.5	613	6.3	613	6.3	115	15.6	115	15.6	122	7.6	122	7.6
Kingsoft Internet Security 2008	465	8.5	465	8.5	735	5.3	735	5.3	37	48.4	37	48.4	61	15.2	61	15.2
McAfee VirusScan Enterprise	60	66.0	988	4.0	493	7.9	494	7.8	87	20.6	86	20.8	102	9.1	110	8.4
Microsoft Forefront Client Security	1469	2.7	1469	2.7	867	4.5	867	4.5	78	23.0	78	23.0	69	13.4	69	13.4
Microsoft Windows Live OneCare	1247	3.2	1247	3.2	658	5.9	658	5.9	107	16.7	107	16.7	74	12.5	74	12.5
MWTI eScan Internet Security	2683	1.5	2683	1.5	637	6.1	637	6.1	468	3.8	468	3.8	472	2.0	472	2.0
Norman Virus Control	999	4.0	999	4.0	2120	1.8	2120	1.8	81	22.1	81	22.1	231	4.0	231	4.0
PC Tools AntiVirus	683	5.8	963	4.1	355	10.9	357	10.9	77	23.3	77	23.3	77	12.0	77	12.0
Quick Heal Anti-Virus Lite	1133	3.5	1217	3.3	94	41.3	95	40.8	77	23.3	83	21.6	57	16.2	73	12.7
Redstone Redprotect	2089	1.9	2089	1.9	554	7.0	554	7.0	274	6.5	274	6.5	273	3.4	273	3.4
Rising Antivirus	2359	1.7	2359	1.7	1158	3.3	1158	3.3	278	6.4	278	6.4	138	6.7	138	6.7
Security Coverage PC Live	[8000+]	[>0.5]	[8000+]	[>0.5]	1074	3.6	1074	3.6	[3600+]	[>0.5]	[3600+]	[>0.5]	130	7.1	130	7.1
Sophos Anti-Virus	51	77.7	2166	1.8	376	10.3	411	9.4	71	25.2	95	18.9	137	6.8	129	7.2
Symantec Norton AntiVirus	406	9.8	406	9.8	562	6.9	562	6.9	165	10.9	165	10.9	146	6.3	146	6.3
Trend Micro Internet Security	527	7.5	661	6.0	316	12.3	319	12.2	94	19.1	94	19.1	103	9.0	103	9.0
VirusBuster Professional	598	6.6	1265	3.1	353	11.0	384	10.1	33	54.3	69	26.0	22	42.1	68	13.6
Webroot Spy Sweeper with Antivirus	1109	3.6	1109	3.6	504	7.7	504	7.7	61	29.4	61	29.4	61	15.2	61	15.2

Like its big brother *Forefront*, *Microsoft*'s home-user product looks slick and smooth, integrates sensibly with the user access controls and offers precious little by way of configuration options. Its insistent pestering to be allowed to disinfect items, and habit of scanning 'additional locations' after a scan of a selected area (for several hours in one instance) slowed testing down somewhat, but actual speeds were quite good and detection once again impressive.

Sharing technology and detection data with *Forefront*, it was no surprise to see *OneCare* achieving similarly high detection rates, covering the WildList in full including the tricky sets of polymorphic samples, and joining its stablemate on the VB100 award podium.

MWTI eScan Internet Security 9.0.779.1

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	99.88%
False positives	0		

Microworld's implementation of the *Kaspersky* scanning engine has numerous additional bells and whistles, and requires several extras be installed including some C++ components and, rather unexpectedly, an update from the *Microsoft* knowledgebase. All of these are thoughtfully provided and only need

confirmation to proceed. A management console for multiple installs is also offered.

Once up and running, the product provides an excellent interface with all the controls one could ever need. Default settings are turned up to the max and scanning was a little languid but results impeccable. With excellent detection throughout, and no false positives, *eScan* earns yet another VB100 award.

Norman Virus Control 5.90

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	98.95%
File infector	99.05%	**Polymorphic**	73.77%
False positives	1		

Norman once again required the admin password to get started but after that provided a very speedy and simple installation, with just four or five hits on the enter key, the offer of an update (which was declined), and a few seconds pause before everything was up and running. A rather bizarre error message suggesting I had no hard drives appeared briefly but seemed to have no impact on functionality, and tests plodded along merrily.

Scanning options were somewhat limited, but a range of common archive types were scanned by default on demand, adding considerably to the scanning time over the archive set. Other scans were mostly pretty speedy, but the set of clean executables took rather a long time on demand. Detection was at its usual high level, with no problems posed by the W32/Virut sets, although somewhat annoyingly the samples were deleted or disinfected despite specifying explicitly the 'please-don't-destroy-my-test-set' option. In the clean sets a single item was mislabelled as malware, leaving *Norman* a fraction short of the required standard for VB100 certification.

PC Tools AntiVirus 2008 4.0.0.25

ItW	99.99%	**Worms & bots**	100.00%
ItW (o/a)	99.99%	**Legacy**	99.98%
File infector	99.21%	**Polymorphic**	79.29%
False positives	0		

The *PC Tools* product evolved from the company's anti-spyware speciality and closely resembles its *Spyware Doctor* flagship offering. This version has a few tweaks which provide an experience pretty similar to any other anti-virus product aimed squarely at the home-user market: bright and colourful, with simple controls and limited configuration. The installer offered to include a toolbar from *Google*, which I declined, but was otherwise fairly straightforward.

Adjustments made by the developers to the default settings caused some early difficulties when it was found that the on-access scanning is now in most cases only activated when fully executing files. Viewing the files in *Explorer* also sparked some detection, but with large numbers of folders needing checking it was felt simplest to adjust the setting to scan files when opened via the testing script. The option to provide this seemed not to function at first, but after a reinstallation everything was fine and testing continued without too much trouble.

The system was a little unresponsive at times, and after scanning the full collection the product interface – along with the rest of the screen – got rather snarled up and couldn't be used until after a reboot. Eventually results were gathered showing good detection and decent speeds, but as with other products based on the *VirusBuster* engine, a few of those nasty Virut samples were missed in the WildList set and a VB100 award is just out of reach for *PC Tools* this month.

Quick Heal Anti-Virus Lite 9.50

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	96.40%
File infector	98.43%	**Polymorphic**	83.86%
False positives	2		

Quick Heal was listed in my submission set under the company's former name, *CAT*, and thus was run rather early on in the test prior to the many frustrations and annoyances which built up in the second half of the product set. It fitted better here anyway, providing amply for my needs with a simple, flexible and highly stable product.

Scanning proceeded at tip-top speeds, with the product's usual decent level of detection and only the older items bringing the scores below excellent. The WildList was ably covered despite the many tricky items, but in the clean set yet more false positives appeared, spoiling *Quick Heal*'s chance of a VB100 for a second time in a row.

Redstone Redprotect 0.5.3.0

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	99.88%
False positives	0		

Redprotect is still a relatively young product with a little development to go before it settles into full stability. The

File Access Lag time (S/MB)	Archive Files				Binaries and System Files				Media and Documents				Other File Types			
	Default Settings		All Files		Default Settings		All Files		Default Settings		All Files		Default Settings		All Files	
	Time	Lag	Time	Lag	Time	Lag	Time	Lag	Time	Lag	Time	Lag	Time	Lag	Time	Lag
AEC Trustport Antivirus	1259	0.32	1259	0.32	666	0.17	666	0.17	178	0.08	178	0.08	241	0.23	241	0.23
Agnitum Outpost Security Suite Pro	79	0.02	N/A	N/A	583	0.14	583	0.14	231	0.11	231	0.11	219	0.21	219	0.21
Ahnlab V3	136	0.03	N/A	N/A	674	0.17	N/A	N/A	101	0.04	N/A	N/A	86	0.06	N/A	N/A
Alwil avast!	102	0.02	1080	0.27	343	0.08	430	0.10	174	0.08	207	0.10	88	0.07	93	0.07
AVG	29	0.01	N/A	N/A	208	0.05	212	0.05	52	0.01	60	0.02	41	0.02	71	0.05
Avira AntiVir	37	0.01	108	0.03	149	0.03	178	0.04	64	0.02	82	0.03	40	0.01	80	0.06
BitDefender AntiVirus 2008	331	0.08	775	0.19	532	0.13	562	0.14	114	0.05	225	0.11	124	0.10	129	0.11
Bullguard	1098	0.28	1098	0.28	568	0.14	568	0.14	104	0.04	104	0.04	117	0.10	117	0.10
CA eTrust Antivirus	26	0.01	N/A	N/A	119	0.02	119	0.02	72	0.02	72	0.02	68	0.04	68	0.04
CA Internet Security	32	0.01	N/A	N/A	139	0.03	139	0.03	79	0.03	79	0.03	69	0.05	69	0.05
Check Point Zone Alarm	61	0.01	N/A	N/A	322	0.08	322	0.08	141	0.06	141	0.06	134	0.12	134	0.12
Doctor Web Dr.Web	760	0.19	4705	1.19	748	0.19	1074	0.27	162	0.07	209	0.10	161	0.14	230	0.22
ESET NOD32 Antivirus	14	0.00	N/A	N/A	86	0.02	86	0.02	72	0.02	72	0.02	80	0.06	80	0.06
Fortinet FortiClient	429	0.11	429	0.11	728	0.18	728	0.18	55	0.02	55	0.02	95	0.07	95	0.07
Frisk F-PROT Antivirus	83	0.02	N/A	N/A	484	0.12	484	0.12	67	0.02	67	0.02	57	0.03	57	0.03
F-Secure Client Security	54	0.01	2096	0.53	349	0.08	609	0.15	84	0.03	249	0.12	58	0.03	202	0.19
G DATA AntiVirus 2008	293	0.07	862	0.22	518	0.13	541	0.13	241	0.12	252	0.12	175	0.16	181	0.17
Hauri ViRobot	256	0.06	N/A	N/A	536	0.13	536	0.13	107	0.04	107	0.04	117	0.10	117	0.10
Ikarus Virus Utilities	192	0.05	215	0.05	456	0.11	747	0.19	91	0.04	99	0.04	138	0.12	139	0.12
K7 Total Security	66	0.02	N/A	N/A	221	0.05	221	0.05	63	0.02	63	0.02	67	0.04	67	0.04
Kaspersky Anti-Virus	18	0.00	171	0.04	110	0.02	294	0.07	114	0.05	143	0.06	82	0.06	132	0.11
Kingsoft Internet Security 2008	36	0.01	N/A	N/A	751	0.19	751	0.19	54	0.01	54	0.01	73	0.05	73	0.05
McAfee VirusScan Enterprise	55	0.01	808	0.20	509	0.12	499	0.12	92	0.04	93	0.04	97	0.08	97	0.08
Microsoft Forefront Client Security	127	0.03	N/A	N/A	589	0.15	589	0.15	95	0.04	95	0.04	86	0.06	86	0.06
Microsoft Windows Live OneCare	149	0.04	N/A	N/A	596	0.15	596	0.15	96	0.04	96	0.04	94	0.07	94	0.07
MWTI eScan Internet Security	1413	0.36	1413	0.36	394	0.10	394	0.10	166	0.08	166	0.08	160	0.14	160	0.14
Norman Virus Control	40	0.01	N/A	N/A	365	0.09	365	0.09	104	0.04	104	0.04	150	0.13	150	0.13
PC Tools AntiVirus	5	0.00	N/A	N/A	189	0.04	189	0.04	209	0.10	209	0.10	160	0.14	160	0.14
Quick Heal Anti-Virus Lite	16	0.00	N/A	N/A	102	0.02	102	0.02	69	0.02	69	0.02	39	0.01	39	0.01
Redstone Redprotect	70	0.02	N/A	N/A	387	0.09	387	0.09	209	0.10	209	0.10	203	0.19	203	0.19
Rising Antivirus	92	0.02	419	0.10	645	0.16	696	0.17	160	0.07	161	0.07	176	0.16	179	0.16
Security Coverage PC Live	232	0.06	232	0.06	405	0.10	405	0.10	106	0.04	106	0.04	96	0.08	96	0.08
Sophos Anti-Virus	43	0.01	1391	0.35	383	0.09	412	0.10	68	0.02	88	0.03	68	0.04	117	0.10
Symantec Norton AntiVirus	30	0.01	N/A	N/A	254	0.06	254	0.06	73	0.02	73	0.02	67	0.04	67	0.04
Trend Micro Internet Security	84	0.02	610	0.15	298	0.07	290	0.07	74	0.03	105	0.04	82	0.06	93	0.07
VirusBuster Professional	36	0.01	N/A	N/A	375	0.09	375	0.09	59	0.02	90	0.03	44	0.02	89	0.07
Webroot Spy Sweeper with Antivirus	10	0.00	N/A	N/A	46	0.01	N/A	N/A	67	0.02	N/A	N/A	61	0.04	N/A	N/A

offering is part of a managed security setup intended to be controlled from a web interface supervising a large number of systems, but the developers thoughtfully provided a testing tool to access controls without having to delve into the registry (in a fully operational setup such changes made by the end-user would quickly be reverted to the master settings).

An initial run proved puzzlingly ineffective, until I remembered that the trial licence for the *Kaspersky* engine powering the product was set to expire halfway through the test period. The addition of an updated key file quickly had things moving along nicely. Although limitations to the settings meant the test collection had to be deleted on access, things still moved along at a pleasingly rapid pace with no major hiccups. The .NET-based interface presented the flakiness I have come to expect from such things, occasionally failing to open or to run a scan when requested, but generally responding well. Results reflected the solid engine at the heart of things, with splendid detection and no false positives earning *Redstone* another VB100 award.

Rising Antivirus 20.33.10

ItW	99.97%	**Worms & bots**	99.73%
ItW (o/a)	99.97%	**Legacy**	56.40%
File infector	93.70%	**Polymorphic**	44.63%
False positives	1		

Beijing-based *Rising* had its first stab at the VB100 at the end of last year and missed out on certification by a whisker; it was good to see the company bravely back in

the saddle. An initial attempt at installing the product went somewhat awry – after demanding the admin password, things seemed to be going well until after the requested reboot, when a message popped up insisting that another reboot was required. Switching to full admin user, I clicked on the desktop shortcut set up during the install, only to find an uninstallation process initiated. Bewildered, I switched to a fresh machine and tried installing with full administrative rights and the UAC controls disabled.

This time everything ran smoothly, eventually bringing up a nice shiny interface and a cute little cartoon lion which lurked in the corner of my desktop occasionally performing stunts, and whipping out a magnifying glass and peering around when scans were run.

The scans seemed to go smoothly too, at a fairly leisurely pace but with very thorough default settings on demand – the on-access controls offered an option to enable archive scanning which, although slowing things down a fraction, seemed to have no effect on detection rates. Disabling the on-access component brought up a CAPTCHA for confirmation, presumably to prevent infiltrations from switching it off.

Detection was fairly good on more recent items, but a small number of polymorphic items in the WildList set were missed and a single false positive was flagged in the clean set. As a result, *Rising* does not quite make it to a VB100 award this time.

Security Coverage PC Live

ItW	84.35%	**Worms & bots**	56.00%
ItW (o/a)	76.00%	**Legacy**	47.00%
File infector	97.20%	**Polymorphic**	54.00%
False positives	1		

PC Live marks the first appearance of the open-source *ClamAV* engine on the VB100 test bench, though it was not the only product based on this engine to be submitted (more on this later). Like many commercial products in the open-source world, *PC Live* is provided as a free tool funded by paid-for support, and the interface is bright and colourful, apparently aiming at the PC novice market with its 1950s soap powder stylings.

As a result, configuration is less than ideal for my requirements, with logging particularly awkward, but things seemed to run pretty smoothly and stability was not a problem. Scanning speeds were rather impressive on access, but less so on demand, with some of the scans of the clean sets abandoned after having been left to run longer than any of the rest of the field, hinting at some kind of snagging issue. Detection results were not excellent, with a pronounced difference between on-demand and on-access scores and with large numbers of clean files blocked on access with no explanation. This implies that the on-access side of things also needs a little improvement.

With the *ClamAV* technology now backed by a commercial firm this could well prove a contender for a VB100 award in the near future, but for now a number of WildList misses and a false positive deny *Security Coverage* its first VB100 award.

Sophos Anti-Virus 7.0.7

ItW	99.997%	**Worms & bots**	100.00%
ItW (o/a)	99.997%	**Legacy**	99.80%
File infector	100.00%	**Polymorphic**	100.00%
False positives	0		

Sophos is a much more familiar product, and another that has changed little for some time. It has a solid simplicity about it which has become an increasingly welcome sight on the test bench, providing respite from the strange and bewildering array of newcomers. Installing the product and opening the interface brings up a confirmation prompt but no password is needed.

As an enterprise-level product there is, of course, a full range of configuration options, making testing a breeze. Speeds were impressive, even with all files and archive types enabled on access, and detection was pretty good across the sets. Once again, however, a couple of samples of the trickiest of those Virut variants were not detected, and *Sophos* misses out on a VB100 award this month.

Symantec Norton Anti-Virus 15.5.0.23

ItW	100.00%	**Worms & bots**	100.00%
ItW (o/a)	100.00%	**Legacy**	100.00%
File infector	100.00%	**Polymorphic**	100.00%
False positives	0		

Despite having tested a wide range of anti-malware products on a regular basis for some years, and despite *Symantec*'s *Norton AntiVirus* being one of the biggest and most widely used products on the market, this is the first time our paths have crossed in any serious fashion.

Initial impressions were not good, the product presenting a rather sickly look with its gaudy yellow on a background of shimmering black, and

Archive scanning		ACE	CAB	JAR	LZH	RAR	TGZ	ZIP	ZIP-SFX	EXT*
AEC Trustport Antivirus	OD	X	√	√	√	√	√	√	√	√
	OA	X	√	X	√	√	√	√	√	√
Agnitum Outpost Security Suite Pro	OD	1	√	√	√	√	√	√	√	√
	OA	X	X	X	X	X	X	X	X	√
Ahnlab V3 Internet Security	OD	X	9	X	9	9	X	9	X	√
	OA	X	X	X	X	X	X	X	X	X
Alwil avast!	OD	X/√	X/√	X/√	X/√	X/√	X/√	X/√	X/√	X/√
	OA	X/8	X/√	X/√	X/√	X/√	X/√	X/√	X/√	X/√
AVG	OD	X	√	1	X	√	X	√	√	X
	OA	X	X	X	X	X	X	X	X	X
Avira AntiVir	OD	√	√	√	√	√	√	√	√	√
	OA	√*	√	√	√	√	√	√	√	√
BitDefender AntiVirus 2008	OD	√	√	√	√	√	8	√	8	√
	OA	X/√	X/√	√	X/√	X/√	X/8	1/√	X/8	√
Bullguard	OD	√	√	√	√	√	8	√	8	√
	OA	√	√	√	√	√	8	√	8	√
CA eTrust Antivirus	OD	X	√	√	√	√	√	√	X	√
	OA	X	X	1	X	X	X	1	X	√
CA Internet Security	OD	X	√	√	√	√	√	√	√	√
	OA	X	X	1	X	X	X	1	X	√
Check Point Zone Alarm	OD	√	√	√	√	√	√	√	√	√
	OA	X	X	X	X	X	X	X	X	√
Doctor Web Dr.Web Antivirus	OD	X	√	√	√	√	X/4	X/9	X/8	√
	OA	X	X/√	X/√	X/√	X/√	X/√	X/√	X/√	√
ESET NOD32 Antivirus	OD	X	√	√	√	√	5	√	√	√
	OA	X	X	X	X	X	X	X	X	√
Fortinet FortiClient	OD	X	√	√	√	√	√	4	√	√
	OA	X	√	√	√	√	√	4	√	√
Frisk F-PROT Antivirus	OD	X	√	√	√	√	√	√	√	√
	OA	X	X	2	X	X	X	2	2	√
F-Secure Client Security	OD	X/√	X/5	X/5	X/5	X/5	X/2	X/5	X/5	X/√
	OA	X/√	X/5	X/5	X/5	X/5	X/2	X/5	X/5	X/√
G DATA AntiVirus 2008	OD	√	√	√	√	√	√	√	√	√
	OA	√	√	√	√	√	8	8	4	√
Hauri ViRobot Desktop	OD	5	5	5	5	5	5	5	5	√
	OA	X	X	X	X	X	X	X	X	√
Ikarus Virus Utilities	OD	2	3	3	3	3	1	3	1	√
	OA	2	3	3	3	3	1	3	1	√
K7 Total Security	OD	X	1	1	1	1	X	1	X	√
	OA	X	X	X	X	X	X	X	X	√
Kaspersky Anti-Virus	OD	√	√	√	√	√	√	√	√	√
	OA	X/4	X/4	X/4	X/4	X/5	X/1	X/2	X/1	√
Kingsoft Internet Security 2008	OD	X	X	√	√	√	√	√	X	√
	OA	X	X	X	X	X	X	X	X	√
McAfee VirusScan Enterprise	OD	2/X	X/√	X/√	X/√	X/√	X/√	X/√	X/√	√
	OA	2/X	X/√	X/√	X/√	X/√	X/√	X/√	X/√	√
Microsoft Forefront Client Security	OD	√	√	√	√	√	√	√	√	√
	OA	X	X	1	X	X	X	1	1	√
Microsoft Windows Live OneCare	OD	√	√	√	√	√	√	√	√	√
	OA	X	X	1	X	X	X	1	1	√
MWTI eScan Internet Security	OD	√	√	√	√	√	√	√	√	√
	OA	√	√	√	√	√	√	√	√	√
Norman Virus Control	OD	X	X	√	√	X	√	√	X	√
	OA	X	X	X	X	X	X	X	X	√
PC Tools AntiVirus	OD	1/2	1/√	1/√	X	1/√	X/√	1/√	1/√	√
	OA	1	1	1	X	1	X	1	1	√
Quick Heal Anti-Virus Lite	OD	X/2	2/5	2/5	X	2/5	1	2/5	X	X/√
	OA	X	X	X	X	X	X	X	X	X
Redstone Redprotect	OD	√	√	√	√	√	√	√	√	√
	OA	X	X	X	X	X	X	X	X	√
Rising Antivirus	OD	√	√	√	√	√	√	√	√	√
	OA	X	X	X	X	X	X	X	X	√
Security Coverage Pc Live	OD	X	X	√	X	√	√	√	√	√
	OA	X	X	5	X	5	X	5	X	√
Sophos Anti-Virus	OD	X	5	5	5	5	5	5	5	√
	OA	X	X/5	X/5	X/5	X/5	X/5	X/5	X/5	X/√
Symantec Norton AntiVirus	OD	X	9	9	9	9	9	9	9	√
	OA	X	X	X	X	X	X	X	X	√
Trend Micro Internet Security	OD	3/6	3/6	3/6	3/6	3/6	1/3	3/6	X	√
	OA	X/6	X/6	X/6	X/6	X/6	1/3	X/6	X	√
VirusBuster Professional	OD	1	√	X/√	X	√	√	√	√	X/√
	OA	X	X	X	X	X	X	X	X	X/√
Webroot Spy Sweeper with Antivirus	OD	X	√	5	√	√	6	√	5	√
	OA	X	X	X	X	X	X	X	X	√

presenting a few error messages about the 'service framework' having stopped working as well as some nondescript internal errors.

In spite of this the product seemed to run pretty solidly with no problems in detection – this took some time to work out though, as the product defaults to removing or disinfecting items, with virtually no configuration options available to the user. This meant that on-access results had to be gathered by means of checking remaining files for changes, and the on-demand scan needed to be left overnight to complete its lengthy operations. It then needed to be re-run as I had missed the unobtrusive button needed to save the scanning log, which is not available from the history screen.

At the end of all this everything seemed okay, with splendid detection rates and no false positives. As a result, *Norton AntiVirus* gains a VB100 award, but does not make any new friends.

Trend Micro Internet Security 16.10.1063

ItW	99.99%	**Worms & bots**	100.00%
ItW (o/a)	99.99%	**Legacy**	98.85%
File infector	99.21%	**Polymorphic**	80.41%
False positives	2		

Trend Micro's product is another slick and professional piece of work from an industry giant, and again seemed to exhibit some distinctly flaky behaviour.

After a complex installation process, thanks to the test lab's lack of an outside connection and the need to update various components manually, the on-demand scanning seemed not to work, both from the 'custom scan' area of the main interface and from the context menu option. This oddity was quickly resolved by starting a full system scan, which worked without a hitch, then stopping it; the other scanning options suddenly became properly responsive.

Configuration was pleasingly thorough, although the custom scan did insist on checking the memory, registry and system areas each time, which became rather tiresome when running a series of small quick scans as part of the speed test, and rendered gathering times for these scans somewhat inexact.

Despite this, scanning speeds turned out to be fairly good and detection rates were reasonable, but a small number of file infectors were missed in the WildList set and a couple of items in the clean set were labelled as 'TROJ_Generic'. As a result *Trend* does not qualify for the VB100 award on this occasion.

VirusBuster Professional 6.0.191

ItW	99.99%	**Worms & bots**	100.00%
ItW (o/a)	99.99%	**Legacy**	99.98%
File infector	99.21%	**Polymorphic**	79.29%
False positives	0		

A much more pleasant experience was had with *VirusBuster*. The installation process sped through in a few steps from the admin password, via a licence code to full operation in a few moments. The interface is a tried and trusted one – not a favourite thanks to a little complexity in the on-demand task design, but perfectly usable and logically designed. Scanning was remarkably speedy and configuration plentiful, although the addition of 'all files' to the on-access mode apparently did not cover archive types.

The manual scans were a little difficult to monitor, presenting no information on their progress; the only way to tell when one was finished was to watch the greyed-out buttons for a return to normal.

Detection was at its usual fairly high level, but as expected from testing a few other products using the same detection technology this month, a small handful of samples from one of the expanded sets of W32/Virut samples – a different strain this time from those causing trouble for most other products – went undetected, and *VirusBuster* misses out on a VB100 award.

Webroot Spy Sweeper with AntiVirus 5.5.7.124

ItW	99.997%	**Worms & bots**	100.00%
ItW (o/a)	99.997%	**Legacy**	99.98%
File infector	100.00%	**Polymorphic**	100.00%
False positives	0		

Webroot is by tradition an anti-spyware product and thus operates in a slightly different style from the majority of other products on test; like some of the other products its on-access detection is not always sparked by simple file access, and had to be measured by copying the collections to the system. Some scanning of file accesses did seem to be happening though, judging by the slight delays running over the clean sets. These scanning speeds are recorded for interest, but do not represent full scanning.

On demand things were pretty speedy, and scans were completed without much difficulty despite there being a shortage of configuration options. Once the logs had been tracked down, the results were processed and tallied pretty closely with those of the *Sophos* engine powering the

product. Unfortunately the similarity extended to the pair of missed Virut samples which also upset *Webroot*'s efforts to earn another VB100.

UNTESTED PRODUCTS

With the exceptionally large number of products taking part this month, a few problems along the way were only to be expected. Most products were eventually coaxed into producing some usable results, but a handful were left by the wayside after taking up more than their fair share of testing time.

A regular participant in recent tests, *iolo AntiVirus* was provided initially as a full version with licence code; unfortunately this version needed access to the Internet to further validate the installation. A plea to the developers brought forth a second version with the option of running in temporary trial mode, and this at least reached the installation stage. Once up and running, however, it insisted that the trial allowance had expired in mid-2007 (despite the build apparently dating from late February 2008); repeated retries failed to persuade the product to enable its protection features, and with many more products jostling for space on the test bench I decided regretfully to spend no further time on *iolo* this month.

As mentioned earlier, a second *ClamAV*-based product was also entered for the test, a fully open-source project called *Moon Secure Anti-Virus* with the delightful tagline 'Anti-Virus from Space'. The product proved well designed, reasonably stable and easy to use, but had been set up such that its on-access scanning would activate only on full execution. The hard-working developers quickly provided patches to enable full on-read scanning, but time was closing in and when these could not be made to work after a couple of attempts it was decided that *Moon Secure* would have to wait a couple of months before its full debut in the VB100 comparative – which should allow the developers sufficient time to perfect their product.

PC Tools submitted two products to this month's test. Alongside the straight anti-virus product was submitted the better-known *Spyware Doctor* with additional anti-virus functionality. This seemed to operate in a similar fashion to its stablemate and *Moon Secure*, with on-access scanning activated only when fully executing samples, and again the options to adjust this did not seem to function properly. As running each sample in all our infected and clean sets would be an enormously complex and time-consuming task, and as *PC Tools* already had a product successfully put through the testing process, it was decided that *Spyware Doctor* should also be dropped from the test, having taken up enough testing time unprofitably.

As these products fell into the category of 'untestable' no judgement can be given on their performance, and they will not be counted as entries in the VB100 history listings.

CONCLUSIONS

Another bumper crop of products and another draining month of intensive testing to a tight deadline.

Again a large number of issues have emerged from the test, most of them caused by a few variants of polymorphic file infectors, some of which have been resident on the WildList for some months and are now represented in greater numbers in the test set to provide a more accurate indication of detection rates. These sets, some of which previously contained only a handful of samples, now contain at least 50 and in some cases 100 or more; I would like to have each polymorphic item in the test sets represented by at least 500 samples, and over time will continue to expand the sets until they reach this sort of level. With many of these items continuing to trip up a variety of products, it seems sensible to have them as thoroughly represented in the test sets as possible.

Issues of stability and unexpected behaviour also caused problems this month, many of which seem likely to be a result of the updated platform. Hopefully many of these issues will be resolved as the service pack becomes more widely implemented by users and more fully tested by developers.

For me, the biggest headache of all was the sheer scale of the test, with a massive 40 products taking part. While a few of these didn't quite make it to the final line-up, they still took a hefty share of testing time, each having been given several retries before being deemed untestable. With a final total of 37, the field shows no signs of shrinking, and the test on *Windows XP* coming up in the summer looks likely to be even more heavily subscribed than this one – unless the developers are scared off by the rash of failed products this time around.

Technical details:

Tests were run on identical machines with *AMD Athlon64* 3800+ dual core processors, 1 GB RAM, 40 GB and 200 GB dual hard disks, DVD/CD-ROM and 3.5-inch floppy drive, all running *Microsoft Windows Vista Business Edition SP1* (32-bit).

Any developers interested in submitting products for VB's comparative reviews (and VB100 certification) should contact John Hawes on john.hawes@virusbtn.com. The current schedule for forthcoming VB comparative reviews can be found at http://www.virusbtn.com/vb100/about/schedule.xml.

END NOTES & NEWS

Forrester's Security Forum will be held 2–3 April 2008 in Amsterdam, the Netherlands. *Forrester* is offering *Virus Bulletin* readers a 15% discount on the registration fee, which can be claimed by downloading the brochure from http://www.forrester.com/imagesV2/uplmisc/Forrester_Virus_Bulletin_Security_Brochure.pdf or calling +31 (0)20 305 4848 and quoting the code 'Virus Bulletin reader'.

RSA Conference 2008 takes place 7–11 April 2008 in San Francisco, CA, USA. This year's theme is the influence of Alan Mathison Turing, the British cryptographer, mathematician, logician, philosopher and biologist, often referred to as the father of modern computer science. Online registration is now available. See http://www.rsaconference.com/2008/US/.

Infosecurity Europe takes place 22–24 April 2008 in London, UK. For more information and to register interest in attending see http://www.infosec.co.uk/virusbulletinevents.

A meeting of the Anti-Malware Testing Standards Organization (AMTSO) will take place on 30 April 2008 in Amsterdam, the Netherlands. For information see http://www.amtso.org/.

The 2nd International CARO Workshop will be held 1–2 May 2008 in Hoofddorp, the Netherlands. The focus of this year's workshop will be on the technical aspects and problems caused by packers, decryptors and obfuscators in the broadest sense For details see http://www.datasecurity-event.com/.

EICAR 2008 will be held 3–6 May 2008 in Laval, France. See http://www.eicar.org/conference/ for the full details.

The 5th Information Security Expo takes place 14–16 May 2008 in Tokyo, Japan. For more details see http://www.ist-expo.jp/en/.

The 9th National Information Security Conference (NISC) will be held 21–23 May 2008 in St Andrews, Scotland. For full details and registration information see http://www.nisc.org.uk/.

Hacker Halted USA 2008 takes place 1–4 June 2008 in Myrtle Beach, SC, USA. The conference aims to raise international awareness towards increased education and ethics in information security. Hacker Halted USA delegates qualify for free admission to the Techno Security Conference which runs concurrently. For more details see http://www.hackerhalted.com/.

The 20th annual FIRST conference will be held 22–27 June 2008 in Vancouver, Canada. The five-day event comprises two days of tutorials and three days of technical sessions where a range of topics of relevance to teams in the global response community will be discussed. For more details see http://www.first.org/conference/.

The 17th USENIX Security Symposium will take place 28 July to 1 August 2008 in San Jose, CA, USA. A two-day training programme will be followed by a 2.5-day technical programme, which will include refereed papers, invited talks, posters, work-in-progress reports, panel discussions, and birds-of-a-feather sessions. For details see http://www.usenix.org/events/sec08/cfp/.

Black Hat USA 2008 takes place 2–7 August 2008 in Las Vegas, NV, USA. Online registration is now open and a call for papers has been issued (deadline 1 May). For details see http://www.blackhat.com/.

VB2008 will take place 1–3 October 2008 in Ottawa, Canada. For the full conference programme including abstracts for all papers and online registration, see http://www.virusbtn.com/conference/vb2008.

Black Hat Japan 2008 takes place 7–10 October 2008 in Tokyo, Japan. For full details see http://www.blackhat.com/.

The SecureLondon Workshop on Computer Forensics will be held 21 October 2008 in London, UK. For further information see https://www.isc2.org/cgi-bin/events/information.cgi?event=58.

RSA Europe 2008 will take place 27–29 October 2008 in London, UK. For full details see http://www.rsaconference.com/2008/Europe/.

SUBSCRIPTION RATES

Subscription price for 1 year (12 issues):

- Single user: $175
- Corporate (turnover < $10 million): $500
- Corporate (turnover < $100 million): $1,000
- Corporate (turnover > $100 million): $2,000
- *Bona fide* charities and educational institutions: $175
- Public libraries and government organizations: $500

Corporate rates include a licence for intranet publication.

See http://www.virusbtn.com/virusbulletin/subscriptions/ for subscription terms and conditions.

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 Web: http://www.virusbtn.com/

vb Spam supplement

CONTENTS

S1 **NEWS & EVENTS**

S1 **FEATURE**

How wise are crowds when assessing phishing websites?

NEWS & EVENTS

CHINESE MOBILE SPAM INVESTIGATED

A large-scale SMS spam attack is being investigated in China after unwanted text messages were sent by seven advertising companies to more than 200 million mobile phone users through the *China Mobile* and *China Unicom* networks.

The mass spamming was highlighted in an investigation by the state-run China Central Television which was timed to coincide with World Consumer Rights Day.

China's State Council has promised to carry out a thorough investigation into the spamming, while *China Mobile* says it will now block SMS messages originating from the seven firms involved. The deputy head of the State Council Office for Rectifying Malpractice encouraged the parties involved to reflect on their actions, saying: 'We urge the parties concerned to beef up self-scrutiny to correct their wrongdoing, which is profit driven in defiance of public interests.' China's Ministry of Information Industry is said to be working alongside other departments to introduce legislation that will clamp down on online and text advertisements.

EVENTS

The 13th general meeting of the Messaging Anti-Abuse Working Group (MAAWG) will be held in Heidelberg, Germany, 10–12 June 2008. The meeting is open to MAAWG members only. The 14th general meeting (also members only) will take place 22–24 September 2008 in Harbour Beach, FL, USA. See http://www.maawg.org/.

CEAS 2008 will take place 21–22 August 2008 in Mountain View, CA, USA. CEAS is soliciting non-spam email for use in its 2008 spam challenge. Non-sensitive legitimate email can be donated at http://ceas.klika.eu/ceas/. For more information about the event see http://www.ceas.cc/2008/.

FEATURE

HOW WISE ARE CROWDS WHEN ASSESSING PHISHING WEBSITES?

Tyler Moore
University of Cambridge, UK

Phishing is the process of enticing people to visit fraudulent websites and persuading them to enter personal information such as usernames and passwords. The information is harvested and used to impersonate the victims in order to empty their bank accounts, run fraudulent auctions, launder money, and so on. New fraudulent websites are set up as quickly as the existing ones are removed.

Maintaining an updated feed of new phishing websites requires constant vigilance and demands significant resources. Most banks and specialist take-down companies maintain their own feeds. One group, called *PhishTank* [1], has created an open-source list of phishing URLs powered by end-user participation. Users can contribute in two ways. First, they submit reports of suspected phishing websites. Second, they examine suspected websites and vote on whether or not they believe them to be phishing sites. Figure 1 shows a screenshot of *PhishTank*'s online voting interface. *PhishTank* relies on the so-called 'wisdom of crowds' [2] to pick out incorrect reports (perhaps pointing to a legitimate bank) and confirm correct reports of malicious websites. Each report is only confirmed (and subsequently disseminated to anti-phishing mechanisms) following the vote of a number of registered users.

PhishTank is part of a growing trend in which web-based participation plays a part in the implementation of security mechanisms, from aggregating spam to tracking malware. Together with my colleague Richard Clayton, I have studied participation in *PhishTank* in order to gain a better understanding of the effectiveness of crowd-based security[1]. We have identified several problems with *PhishTank* which leave the system vulnerable to manipulation. Unfortunately, these weaknesses are not limited to *PhishTank*, but reflect fundamental difficulties that can arise whenever security decisions are taken as a result of mass participation.

[1]A complete technical paper is available [3].

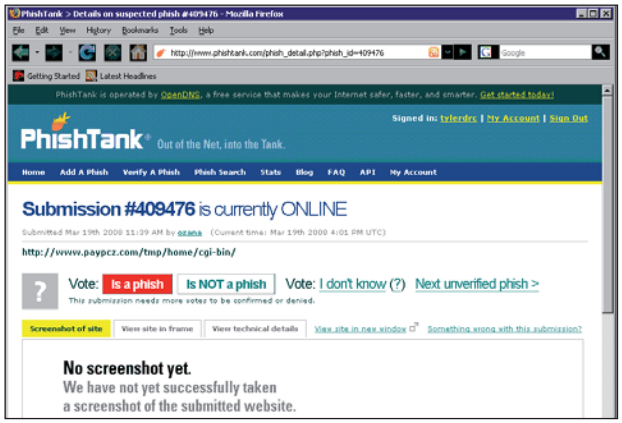

Figure 1: PhishTank user interface.

DATA COLLECTION AND ANALYSIS

We examined completed reports from 176,366 phishing URLs submitted to *PhishTank* between February and September 2007. A total of 3,798 users participated by submitting reports and/or voting. In all, 881,511 votes were cast, suggesting an average of 53 submissions and 232 votes per user. In reality, however, a small number of users are responsible for the majority of submissions and votes. The top two submitters, adding 93,588 and 31,910 phishing records respectively, are actually two anti-phishing organizations that have contributed their own, unverified, feeds of suspect websites. The top verifiers have voted over 100,000 times, while most users vote only a few times.

Many of the leading verifiers have been invited to serve on *PhishTank*'s panel of 25 moderators. Moderators are assigned additional responsibilities such as cleaning up malformed URLs from submissions. Collectively, moderators cast 652,625 votes, or 74% of the total. So while the moderators are doing the majority of the work, a significant contribution is made by the large number of 'regular' users.

In fact, the distributions of user submissions and votes in *PhishTank* are each characterized by a power law. Power-law distributions appear in many real-world contexts, from the distribution of city populations to the number of academic citations to BGP routing topologies. Power-law distributions have highly skewed populations with 'long tails' – that is, a limited number of large values appear several orders of magnitude beyond the much smaller median value. In the case of *PhishTank*, while most users submit and vote only a handful of times, a few users participate many thousands of times.

The intuitive argument put forth in favour of the robustness of 'crowd-sourced' applications like *PhishTank* is that the opinions of many users can outweigh the occasional mistake, or even the views of a malicious user. However, when the rate of participation follows a power-law distribution, a single, highly active user's actions can impact greatly a system's overall accuracy – one subversive participant might undermine the system. This brittleness can lead to big problems if phishers decide to manipulate *PhishTank*.

PhishTank asks its users to vote on every unique URL submitted, which imposes a very large and unnecessary burden on its volunteers. The 'rock-phish' gang is a group of criminals who perpetrate phishing attacks on a massive scale [4]. Instead of compromising machines for hosting fake HTML in an ad-hoc manner, the gang first purchases a number of domains with meaningless names like 'lof80.info'. They then send email spam containing a long URL of the form 'http://www.bank.com.id123.lof80.info/vr'. This URL includes a unique identifier; all variants are resolved to a particular IP address using wild-card DNS. Up to 25 banks are impersonated within each domain. For a more complete description of rock-phish attacks see [5].

Transmitting unique URLs trips up spam filters looking for repeated links, and also fools collators like *PhishTank* into recording duplicate entries. Consequently, voting on rock-phish attacks becomes very repetitive. We observed 3,260 unique rock-phish domains in *PhishTank*. These domains appeared in 120,662 submissions, 60% of the overall total. Furthermore, 893 users voted a total of 550,851 times on these domains! This is a dreadfully inefficient allocation of user resources, which could instead be directed to speeding up verification times, for example.

TESTING THE ACCURACY OF PHISHTANK'S DECISIONS

We now examine the correctness of *PhishTank* users' contributions. We first describe common causes of inaccuracy and discuss their prevalence. We then demonstrate that inexperienced users are far more likely to make mistakes than experienced ones. Finally, we show that users with bad voting records 'cluster' by often voting together.

Miscategorization in PhishTank

The vast majority of submissions to *PhishTank* are indeed phishing URLs. Of 176,654 verified submissions, just 5,295, or 3%, are voted down as invalid. Most appear to be honest mistakes. Some users submit all URLs from their spam, while others add URLs for other types of malicious websites, such as those involved in advanced fee fraud (419 scams). Sometimes, though, carefully crafted phishing websites and legitimate non-English websites are miscategorized. Most commonly, an obscure credit union or bank that uses a different domain name for its online

banking may be marked as a phish. Even moderators make mistakes: 1.2% of their submissions are deemed invalid.

In addition to invalid submissions that are correctly voted down, submissions that are incorrectly classified present a significant worry. Identifying false positives and negatives is hard because *PhishTank* rewrites history without keeping any public record of changes. By periodically re-checking all *PhishTank* records for reversals, we identified 39 false positives – legitimate websites incorrectly classified as phishing sites – and three false negatives – phishing websites incorrectly classified as legitimate. Twelve of these classifications were initially agreed upon unanimously.

Of the false positives, 30 were legitimate banks, and the remaining nine were other scams miscategorized as phishing. Several popular websites' primary domains were voted as phish, including *eBay* (ebay.com, ebay.de), *Fifth Third Bank* (53.com) and *National City* (nationalcity.com). Minimizing these types of false positive is essential for *PhishTank* because even a small number of false categorizations could undermine its credibility.

Unsurprisingly, there are many more false positives than false negatives since the vast majority of submitted phishes are valid. Most noteworthy was the fact that a URL for the rock-phish domain eportid.ph was incorrectly classified as innocuous. Five other URLs for the same domain were submitted to *PhishTank* prior to the false negative, with each correctly identified as a phish. Thus, requiring users to vote for the same rock-phish domain many times is not only inefficient, it is unsafe.

Experience influences user accuracy

Where do these mistakes come from? It is reasonable to expect occasional users to commit more errors than those who contribute often. Indeed, we find strong evidence for this in the data. Figure 2 plots the rates of inaccuracy for submissions and votes grouped by user participation rates.

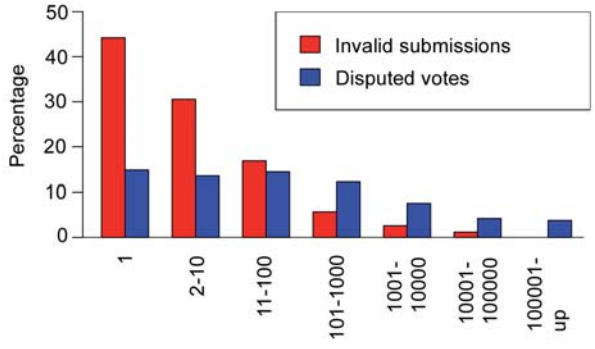

Figure 2: Inaccuracy of user submissions and votes according to the total number of submissions and votes per user, respectively.

For instance, 44% of URLs from users who submit just once are voted down as invalid. Accuracy rate improves with frequency of submissions (30% of submissions from users who submit between two and 10 URLs are invalid; only 17% are invalid for users submitting between 11 and 100 times), with the top submitters incorrect just 1.2% of the time.

A similar, albeit less drastic, difference can be observed in voting accuracy. Users voting fewer than 100 times are likely to disagree with their peers 14% of the time. This improves steadily for more active users, with the most active voters in conflict just 3.7% of the time, which is in line with the overall average. These results suggest that the views of inexperienced users should perhaps be assigned less weight when compared to highly experienced users.

Users with bad voting records vote together

We also found evidence that bad decisions reinforce themselves. Users with bad voting records are more likely to vote on the same phishing reports than would be expected if their votes were independent. For 186 of the 1,791 users who have voted, over half of their votes were disputed. These high-conflict voters voted on the same phishing URLs approximately one thousand times more frequently than would be the case if there were no connection between how they selected their votes.

What are the implications? While it is possible that these high-conflict users are deliberately voting incorrectly together (or are the same person), the more likely explanation is that incorrect decisions reinforce each other. When well-intentioned users vote incorrectly, they have apparently made the same mistakes.

DISRUPTING PHISHTANK'S VERIFICATION SYSTEM

Recently, a number of anti-phishing websites were targeted by a denial-of-service attack, severely hindering their work in removing malicious sites [6]. Hence, there is already evidence that phishers are motivated to disrupt the operations of groups like *PhishTank*. But even if enough bandwidth is provisioned to counter these attacks, *PhishTank* remains susceptible to vote rigging, which could undermine its credibility. Any crowd-based decision mechanism is susceptible to manipulation. However, as we will see, certain characteristics of user participation make *PhishTank* particularly vulnerable.

Attacks and countermeasures

We anticipate three types of attack on *PhishTank*: (1) the submitting of invalid reports accusing legitimate websites,

(2) the voting of legitimate websites as phish, and (3) the voting of malicious websites as legitimate. A selfish attacker seeks to protect their own phishing websites by voting down any accusatory report as invalid (attack type 3). A selfish attacker may be prepared to implicate the websites of other phishers in order to protect their own sites. An undermining attacker takes a wider view by going after the credibility of *PhishTank*, which is best achieved by combining attacks 1 and 2: submitting URLs for legitimate websites and promptly voting them to be phish. This attacker may also increase confusion by attempting to create false negatives, voting phishing websites as legitimate.

Detecting and defending against these attacks while maintaining an open submission and verification policy is hard. Many of the straightforward countermeasures can be sidestepped by a smart attacker. We consider a number of countermeasures in turn, demonstrating their inadequacy.

One simple countermeasure is to place an upper limit on the number of actions any user can take. This is unworkable for *PhishTank* due to its power-law distribution: some legitimate users participate many thousands of times. In any case, an enforced even distribution is easily defeated by a Sybil attack [7], where users register many identities. Given that many phishing attackers use botnets, even strict enforcement of 'one person, one vote' can probably be overcome.

The next obvious countermeasure is to impose voting requirements. For example, a user must have participated 'correctly' *n* times before their opinion is weighed. This is ineffective for *PhishTank*, though the developers tell us that they do implement this countermeasure. Since 97% of all submissions are valid, an attacker can quickly boost their reputation by voting for a phish slightly more than *n* times. A savvy attacker can even minimize their implication of real phishing websites by voting only for rock-phish domains or duplicate URLs. Indeed, the highly stylized format for rock-phish URLs makes it easy to automate correct voting at almost any desired scale.

What about ignoring any user with more than *n* invalid submissions or incorrect votes? After all, a malicious user is unlikely to force through all of his bad submissions and votes. Unfortunately, the power-law distribution of user participation causes another problem. Many active users who do a lot of good also make a lot of mistakes. For instance, the top submitter, *antiphishing*, is also the user with the highest number of invalid submissions (578). An improvement would be to ban users who are wrong more than *x*% of the time. Nevertheless, attackers can simply pad their statistics by voting randomly, or by voting for duplicates and rock-phish URLs.

Moderators already participate in nearly every vote, so it would not be unreasonable to insist that they were the submitter or voted with the majority. However, we know that even moderators make mistakes – over 1% of moderators' submissions were voted down as invalid. Nonetheless, perhaps the best strategy for *PhishTank* is to use trusted moderators exclusively if there is any suspicion that the organization is under attack. Given that the 25 moderators already cast 74% of *PhishTank's* votes, silencing the whole crowd to root out the attackers may sometimes be wise, even if it contradicts the principles of open participation.

Lessons for secure crowd-sourcing

After examining the *PhishTank* data we can draw several general lessons about applying the open-participation model to security tools.

Lesson 1: The distribution of user participation matters. There is a natural tendency for highly skewed distributions, even power laws, in user-participation rates. While there may certainly be cases that are not as skewed as *PhishTank*, security engineers should check the distribution for wide variance when assessing the risk of leveraging user participation.

Skewed distributions can create security problems. First, corruption (or simply the absence) of a few high-value participants can completely undermine the system. Second, because good users can participate extensively, bad users can too. This can frustrate simple rate-limiting countermeasures.

Lesson 2: Crowd-sourced decisions should be difficult to guess. Any decision that can reliably be guessed can be automated and exploited by an attacker. The underlying accuracy of *PhishTank*'s raw data (97% phish) makes it easy for an attacker to improve their reputation by voting all submissions blindly as phish.

Lesson 3: Do not make users work harder than necessary. Requiring users to vote multiple times for duplicate URLs and rock-phish domains is not only an efficiency issue. It becomes a security liability since it allows an attacker to build up reputation without making a positive contribution.

COMPARING OPEN AND CLOSED PHISHING FEEDS

PhishTank is not the only organization tracking and classifying phishing websites. Other organizations do not follow *PhishTank*'s open submission and verification policy; instead, they gather their own proprietary lists of suspicious websites and employees determine whether they are phishing. We have obtained a feed from one such company. This has enabled us to compare the feeds for completeness and speed of verification.

We compared the feeds during a four-week period in July and August 2007. We first examined ordinary phishing

websites, excluding rock-phish URLs. *PhishTank* reported 8,296 unique phishing URLs, while the other company identified 8,730. The two feeds shared 5,711 reports in common. For rock-phish URLs the difference is more stark. *PhishTank* identified 586 rock-phish domains during the sample, while the other company detected 1,003 – nearly twice as many. Furthermore, the other company identified 78% of the rock-phish domains found in *PhishTank*, along with an additional 544 missed by *PhishTank*. Venn diagrams for the feeds are presented in Figure 3.

It is noteworthy that both feeds include many phishing websites which do not appear on the other. This observation supports the case for a universal feed shared between the banks and the various anti-phishing organizations.

Prompt identification and removal of phishing websites is critical, so a feed's relevance depends upon how quickly it is updated. Requiring several users to vote introduces significant delays. On average, *PhishTank* submissions take approximately 46 hours to be verified. A few instances take a very long time to be verified, which skews the average. The median, by contrast, is around 15 hours.

We compared the submission and verification times for URLs appearing in both feeds. On average, *PhishTank* saw the submissions first, by around 11 minutes, but after an average delay of just eight seconds the other company had verified them. *PhishTank*'s voting-based verification meant that it did not verify the URLs (and therefore did not disseminate them) until 16 hours later. For the rock-phish URLs, we compared the earliest instance of each domain, finding that overlapping domains appeared in *PhishTank*'s feed 12 hours after they appeared in the other company's feed, and were not verified for another 12 hours.

CONCLUSION

End-user participation is an increasingly popular resource for carrying out information security tasks. Having examined one such effort to gather and disseminate phishing information, we conclude that while such open approaches are promising, they are currently less effective overall than the more traditional closed methods. Compared to a data feed collected in a conventional manner, *PhishTank* is less complete and less timely. On the positive side, *PhishTank*'s decisions appear mostly accurate: we identified only a few incorrect decisions, all of which were later reversed. However, we found that inexperienced users make many mistakes and that users with bad voting records tend to commit the same errors. So the 'wisdom' of crowds sometimes descends into folly.

We also found that user participation varies greatly, raising concerns about the ongoing reliability of *PhishTank*'s decisions due to the risk of manipulation by small numbers of people. We have described how *PhishTank* can be undermined by a phishing attacker bent on corrupting its classifications, and furthermore how the power-law distribution of user participation makes attacks simultaneously easier to carry out and harder to defend against.

Despite these problems, we do not advocate against leveraging user participation in the design of all security mechanisms. Rather, we believe that the circumstances must be examined more carefully for each application, and furthermore that threat models must address the potential for manipulation.

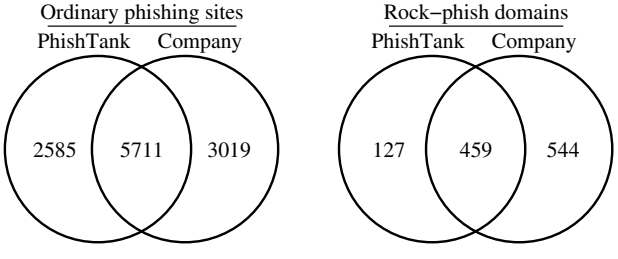

Figure 3: Venn diagram comparing coverage of phishing websites identified by PhishTank and a take-down company.

REFERENCES

[1] PhishTank. http://www.phishtank.com/.

[2] Surowiecki, J. The wisdom of crowds: why the many are smarter than the few. Doubleday, New York (2004).

[3] Moore, T.; Clayton, R. Evaluating the wisdom of crowds in assessing phishing websites. 12th International Financial Cryptography and Data Security Conference (FC). LNCS, to appear. Springer (2008).

[4] McMillan, R. 'Rock Phish' blamed for surge in phishing. InfoWorld, 12 Dec 2006. http://www.infoworld.com/article/06/12/12/HNrockphish 1.html.

[5] Moore, T.; Clayton, R. Examining the impact of website take-down on phishing. Anti-Phishing Working Group eCrime Researcher's Summit, pp.1–13. ACM Press, New York (2007).

[6] Larkin, E. Online thugs assault sites that specialize in security help. PC World, 11 Sep 2007. http://www.pcworld.com/businesscenter/article/137084/online thugs assault sites that specialize in security help .html.

[7] Douceur, J.R. The Sybil attack. 1st International Workshop on Peer-to-Peer Systems. Lecture Notes in Computer Science (LNCS), vol. 2429. Springer (2002) 251–260.