

The Threats You Can't See

No matter how diligently you patch your PC, it's still vulnerable to **zero-day** attacks. Here's what you should know about these hazards and the security holes they exploit.

BY RYAN SINGEL



YOU'RE NO SECURITY slouch. You keep your programs up-to-date, and you have antivirus installed. You're careful about where you surf and what you install on your computer.

But last September, if you had visited a blog hosted by Host-Gator, a top-tier provider based in Florida, your PC's browser would have been summarily redirected to an infected Web site that exploited a vulnerability in an old Microsoft image format. ▶

Within seconds, a payload of malware would have invaded your computer.

Had this happened, you'd have fallen victim to a zero-day exploit—an attack against a software flaw that occurs at a time when no patch to correct the problem exists. The term originally described a vulnerability that was exploited “in the wild” (that is, outside a research lab) on the same day that a patch became available for it, leaving IT staffs zero intervening days to close the hole.

Today, the value of zero-day exploits to online criminals is skyrocketing precisely because the attacks can break into up-to-date, well-maintained systems. Last December, for instance, Trend Micro chief technology officer Raimund Genes noticed a sales pitch scrolling by in an Internet chat room: A hacker wanted to sell an undisclosed vulnerability in a beta version of Windows Vista for a staggering \$50,000, though Genes was unable to determine if anyone bought the code.

“There’s much more of an organized undercurrent now,” says Dave Marcus, security research manager for McAfee. “[The criminals] have figured out they can make money with malware.”

MALWARE FOR MONEY

THE MALWARE COULD be a “bot,” for example, capable of forcing your PC to relay spam or participate in denial-of-service attacks that push Web sites offline. “It’s a lot easier than knocking some old lady over the head and stealing her pocketbook,” Marcus says. “It’s very anonymous, and [a criminal could] do it from the safety and comfort of Starbucks.”

Thanks to features such as improved scanning that doesn’t rely on signatures, McAfee’s antivirus and other security programs are becoming more nimble at protecting against unknown threats. And a wide array of new free and commercial programs supply proactive protection against zero-day assaults by limiting a successful attack’s destructive power.

The right security setup can protect you 99 percent of the time, says Jeff Moss, who founded the annual BlackHat secu-

rity conference. But targeted attacks can sometimes sneak through anyway. “You can go and buy a lot of firewalls and software and equipment,” he says, “but if the right zero-day exists in the right component, it’s almost like all that extra fanciness doesn’t make a difference.”

DRIVE-BY DOWNLOAD

A ZERO-DAY ATTACK MARCHES ON



LAST SEPTEMBER Sunbelt Software discovered attacks against a vulnerability in Vector Markup

Language graphics, which are rarely seen but still supported in Windows. Within a week, criminals infected thousands of sites with poisoned images capable of inflicting a drive-by-download attack on any hapless user who viewed the image.

September 18, 2006: The first VML image attacks are reported on a Russian Web site.

September 19, 2006: Microsoft issues an advisory with a workaround, and says a patch will follow on October 10.

September 20, 2006: Symantec reports that the malicious code has been included in an easy-to-use exploit kit sold in Eastern Europe.

September 22, 2006: Zeroday Emergency Response Team releases an unofficial patch. Thousands of legitimate but compromised HostGator sites redirect visitors to sites with the VML attack code.

September 26, 2006: Microsoft releases a fix two weeks ahead of schedule.

January 16, 2007: iDefense confirms that a similar zero-day attack is exploiting another critical VML hole.

The most dangerous varieties of pre-patch attackware permit drive-by downloads, where simply browsing a poisoned page or reading an infected HTML e-mail can trigger an invasion capable of stuffing your PC full of spyware, Trojan horses, or other malware. Between the end of 2005 and the end of 2006, online thugs used at least two such zero-day assaults to attack millions of people by exploiting holes in a rarely used Microsoft image format.

In the case of the HostGator debacle involving the Windows image flaw, the exploit took advantage of a long-unnoticed vulnerability in Internet Explorer’s handling of the Vector Markup Language (VML), an infrequently used standard for creating 3D graphics.

The threat was first reported in September by security company Sunbelt Software, which found it on a pornographic Russian Web site. By itself, the hole was bad enough: If you browsed a site containing a booby-trapped image, you could be hit by a drive-by download. But opportunistic attackers recognized how to magnify the damage.

By targeting a second unknown hole in cPanel, a Web site management interface, crooks hijacked thousands of sites maintained by HostGator. Visitors to these legitimate but compromised Web sites were redirected to malicious sites that contained the VML exploit.

Microsoft products such as Internet Explorer, Office, and the Windows operating system itself are common targets of zero-day (and other) attacks, in part because they dominate the software landscape. But Microsoft’s failure in the past to adequately integrate security into its product development has contributed to its products’ status as popular (and easy) targets. Vista, on the other hand, is getting high marks for security, at least early on; see “Zero-Day Defense in Microsoft’s New Operating System” on page 122.

In 2006 alone, four different zero-day exploits attacked Internet Explorer 6, directly or indirectly. The year began with continuing attacks that capitalized on a flaw discovered in December 2005, in the

Windows Metafile image format; the hole was in an underlying part of Windows that IE used to render a WMF image.

Once the attacks became publicly known, Microsoft first said that it would include a patch to fix the hole weeks later, as part of its normal patch cycle—but as exploits and the public outcry against them escalated, the company released an out-of-cycle fix in early January.

The patch didn't end the attacks, however, demonstrating that zero-day exploits can have long-term effects. Like the VML flaw, the Metafile exploit opened the door to drive-by-downloads, which criminals love because victims don't have to click an infected image to be hit. If you installed Microsoft's patch via Automatic Updates, you were fine. But clearly, many Windows users didn't.


In July a malicious banner ad for Deck-outyourdeck.com made its way onto sites like MySpace and Webshots via an ad distribution network serving thousands of sites. The malware hidden in the banner downloaded a Trojan horse onto victims' PCs, and it in turn installed adware and spyware. Informed observers put the number of victims—seven months after a patch was available—in the millions.

TARGETED OFFICE ATTACKS

UNLIKE THE MOST dangerous of the IE-targeted zero-day threats, those directed against Word and other Office apps can't employ drive-by downloads. Instead, they typically rely on getting a victim to double-click an e-mail attachment—and when they are paired with orchestrated attacks against particular companies, even wary users can accidentally click.

By sending the employees of a targeted firm a faked (or "spoofed") e-mail message that appears to come from a coworker or other source within the company, a malicious hacker has a much better chance of persuading recipients to open an attached Word document than if the message appeared to come from a random sender.

In mid-December, Microsoft confirmed that Word contained two such vulnerabilities that crackers exploited to launch



**IF THE RIGHT
ZERO-DAY
EXISTS...IT'S
ALMOST LIKE ALL
THAT EXTRA
[SECURITY]
DOESN'T MAKE A
DIFFERENCE.**

Jeff Moss, BlackHat conference founder

"very limited and targeted attacks," following manipulation of a string of similar flaws in Excel and PowerPoint. The company now warns users to be wary not only of e-mail attachments in messages from unknown senders, but also of unsolicited attachments from known senders.

Microsoft products may be the most popular zero-day targets, but other common software has provided equally dangerous avenues of attack. In January a researcher announced discovery of a flaw in QuickTime's handling of streaming video that would have allowed an attacker to effectively commandeer a victim's computer. In late November 2006 a zero-day vulnerability in the Adobe ActiveX

browser control introduced a similar risk.

The rise in zero-day incidents mirrors a major increase in the annual number of reported software vulnerabilities. In 2006 software makers and researchers catalogued some 7247 vulnerabilities; that's 39 percent more than in 2005, according to Internet Security Systems Xforce.

Most of these bugs don't lead to a zero-day exploit, however. Software companies often receive reports of bugs and crashes from their users, leading to discovery of security holes, which the companies then patch before any attacker can exploit them. When outside security researchers discover a flaw, they (for the most part, anyway) adhere to a set of practices, known as "ethical disclosure," specifically designed to avoid zero-day attacks.

Under ethical disclosure, researchers first contact the software vendor confidentially to report their findings. The company doesn't announce the problem until it has a patch ready, at which ▶

time it publicly credits the original researchers with having uncovered the flaw.

But sometimes researchers, frustrated by the slow pace of a software maker's investigation, go public with details about the vulnerability while it is still unresolved. Some experts consider this tactic a necessary evil to force recalcitrant companies to issue a fix; others decry it as an unethical breach of industry practices.

People who advocate going public argue that if a researcher knows about the flaw, criminals might, too—and smart criminals will keep their attacks small and targeted so as to avoid the maker's attention, and a fix. Unfortunately, all too often, public disclosures of a vulnerability prompt public zero-day attacks.

Another controversial practice is bounty hunting. Some organizations, including iDefense and 3Com's Zero Day Initiative, pay researchers to report zero-day exploits to them. iDefense, for example, offers an \$8000 bounty for information about vulnerabilities in IE 7 and Vista.

OFFICE

FOCUSED DOC STRIKE



MAY 21, 2006: Targeted attacks are launched from Taiwan and China, exploiting a Microsoft

Word bug (one of many zero-day Office flaws reported during 2006), to strike an unnamed company. According to the Internet Storm Center, the attacks mimic an internal company e-mail, increasing the odds that an unsuspecting employee will open a poisoned attachment.

The security companies then communicate their discoveries confidentially to the software companies. While not universally admired, these programs do put cash in researchers' pockets—an outcome that many prefer to a simple public pat on the back from software vendors.

Perhaps more important, security company bounties compete with the growing black market for zero-day exploits. The Vista vulnerability seller that Trend Micro's Genes observed in a chat room may or may not have found a buyer at the \$50,000 asking price, but reports from eWeek.com and security companies say that the Windows Metafile attacks began immediately after a sale of relevant bug details for the tidy sum of \$4000.

To find marketable flaws, researchers and criminals use automated tools called fuzzers to locate places where a program accepts input, and then systematically feed them bizarre combinations of data. Frequently this testing turns up an exploitable flaw called a buffer overflow.

Software companies, including Microsoft, commonly use the tools to find flaws in their own products proactively. But so do the crooks: BlackHat organizer Moss and many other experts say that Eastern European organized criminals, disciplined groups of Chinese hackers, ▶

VISTA

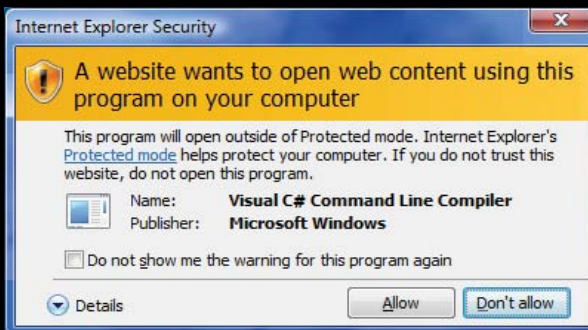
ZERO-DAY DEFENSE IN MICROSOFT'S NEW OPERATING SYSTEM

HOW WELL WILL the much-touted new security features in Windows Vista protect you from a zero-day attack? Better than you might think.

The key new feature may be User Account Control, which changes user account permissions in Vista. As a matter of convenience—since administrator privileges are required for many common system tasks—almost every home user runs Windows XP

under the administrator setting. But attackers can take advantage of the setting's carte-blanche rights to make major modifications to a system, such as by installing malware-hiding rootkits.

In contrast, the default Vista user account occupies a middle ground between an anything-goes administrator



VISTA'S NEW SECURITY features, such as IE's Protected Mode, offer additional protection—and more pop-up dialog boxes.

account and a hands-tied guest pass. Microsoft has tried to make the change more palatable by authorizing standard account holders to perform some routine system tasks such as printer driver installation, but power users are already complaining about having to click through too many User Account Control

prompts that demand an administrator password.

Also, Internet Explorer by default runs in a protected mode with the fewest permissions possible. This arrangement limits the havoc that a zero-day exploit capable of hijacking IE (such as the WMF or VML exploit) could wreak on your PC.

Finally, Vista ships with Windows Defender, which can block malware attempts

to add entries to the startup folder—for example, in the guise of baseline anti-spyware. The operating system also shuffles the locations at which libraries and programs load into memory, so malware that attempts to find and change the system's most important processes will have to hit moving targets.

and other miscreants use fuzzers to find valuable zero-day exploits. The discoverers can use the exploit to attack on their own or, as in the case of the WMF exploit, can sell it on the black market.

When a software maker can patch a security flaw before any attacks occur, both company IT staff and home users have time to update their software and stay ahead of the curve. But as soon as a zero-day attack commences, the clock is ticking—and sometimes it ticks for a while before the needed patch arrives.

During the first half of 2006, according to Symantec's September 2006 Internet Threat Security Report, Microsoft tied Red Hat Linux for the fastest patch development time for commercial operating systems: an average time of 13 days.

BROWSER PATCH DELAYS

BUT IN UPDATING browsers—especially when a zero-day attack was in progress—Microsoft trailed Apple, Mozilla, and Opera in providing fixes. On average, IE patches appeared ten days after the flaw was reported, while Opera, Mozilla, and Safari browsers were patched, on average, in two, three, and five days, respectively. (For a timeline that tracks the zero-day VML attack and patch cycle, see “A Zero-Day Attack Marches On” on page 120.)

Adam Shostack, a program manager for Microsoft's security development life-cycle team, says that sometimes a longer time frame is only to be expected, given the complexity of Microsoft's user base.

“We have to test security updates to make sure they will work with 28 different languages and every OS that supports the application,” Shostack says. “We really work to balance quality with speed.”

Security software helps protect against unknown threats during the dangerous interval that separates an initial attack from the release of an effective patch, but traditional antivirus programs rely on having an identified signature for an attack in order to guard against it. This pits malware writers against security companies in a constant cat-and-mouse game, with malicious hackers sending out a

IMAGE ATTACK

MYSPACE INVADED



DECEMBER 28, 2005: In this precursor to the similar VML attack, a flaw in Microsoft's little-used WMF image type permits drive-by downloads to occur if a user views a page containing a poisoned image. Microsoft releases an early patch on January 5, but in July a malicious banner ad infects millions of as-yet-unpatched PCs visiting MySpace, Webshots, and other sites.

steady stream of designer Trojan horses and the like that they have tweaked just enough to defeat signature recognition.

Heuristics and behavior-based analysis can move beyond this evolutionary pattern to give security programs an edge. Such scans use algorithms, rather than signatures, to look for abnormal behavior or files. Heuristic analysis checks the contents of potential malware for things such as a suspect method of working with memory. Behavioral analysis, meanwhile, watches programs for conduct typical of malware (such as starting an e-mail relay server), trying to identify unwanted interlopers by what they do rather than by what they contain.

Today most major antivirus products incorporate one or both types of analysis. Last year, *PC World* tests that used one-month-old signatures yielded success rates of between 20 and 50 percent.

Heuristics and behavior analysis are susceptible to false positives, though. A security program may not be able to distinguish between a keylogger and a game that asks for direct access to the keyboard to shorten response time. As a result, the security software may needlessly bother a user with pop-up alerts and questions.

BlackHat's Jeff Moss estimates that this kind of detection won't be genuinely useful on its own for another five years. “The false positive and false negative

rates are too high. Everyone is coming up with novel ways of detecting misuse; but as soon they try to deploy it, users revolt.”

OTHER APPROACHES

MEMBERS OF ANOTHER class of security products try to resist new threats by changing the user's computing environment to limit damage from a successful invasion. Some (like GreenBorder Pro) create a “sandbox,” or virtually walled-off environment, for frequently targeted programs such as Web browsers and e-mail clients. An attack might break through IE, for instance, but any attempt to install spyware or make other malicious changes would not escape the sandbox.

Other programs, in lieu of creating a virtual environment, modify users' rights so as to remove an application's ability to make deep system changes. This category of utility includes the free DropMyRights applet from Microsoft.

Still other types of programs, such as the free VMWare Player, install a distinct, encapsulated operating system that includes its own browser. The cordoned-off browser is completely detached from your regular computing environment. For more information on all these types of damage-mitigation security programs, see Erik Larkin's “Disarm Net Threats” (find.pcworld.com/56458).

Windows Vista introduces several security updates that work along some of the same lines. But no one thinks software vulnerabilities or zero-day exploits are going to disappear. Unfortunately, the established black market for stolen data and unwitting spam-senders all but guarantees that criminals will continue to find ways to profit from malware misery.

Nevertheless, David Perry, global director of education for Trend Micro, remains cautiously optimistic about the future state of Internet safety. “I believe eventually we will get the Web to the point where threats are just a nuisance,” he says. “But that isn't coming this year.” ■

Ryan Singel covers computer security as a San Francisco-based freelance author.