

Fighting the worms of mass destruction

SAN FRANCISCO

Hooligans are trashing our online space. How can they be stopped?

WHEN Microsoft released its latest monthly batch of software patches on November 11th, it included one designed to repair a previously unknown flaw in Windows 2000. Such an event often acts as a tip-off to the writers of computer worms and viruses, who know that new patches are never applied very widely or very quickly. It is possible that this new flaw could herald a series of computer failures at least as damaging as those seen earlier in the year.

Bill Gates, the chairman of Microsoft, once made a habit of using his keynote speech at Comdex, the computer industry's top annual trade show, to launch his company's "next big thing". Not all of these innovations succeeded, though at the time of their unveiling they all contained something to excite the industry. But times have changed. Mr Gates began his speech at the Las Vegas show this month by unveiling a dull bit of software that manages the distribution of security patches on a network. He followed this with an almost equally dreary firewall and a new spam-filtering initiative. These, rather than glitzy product announcements, are the industry's new priorities. Closing loopholes exploited by viruses, worms and hackers, said Mr Gates, is "the largest thing we are doing".

Eradicating spam is a top priority for the American government too. The Can Spam Act made comfortable progress through Congress this week, the first piece of federal legislation to attempt to reduce the amount of unsolicited electronic garbage passing over the internet. Opinion is divided as to how effective the new law will be. But if it works at all, it will also help to improve internet security. Spam is often the transmitter of computer viruses.

Cyber-louts

The biggest fear is that viruses and worms will be used by terrorists to hold societies to ransom. Last year, American spies found a shack in Pakistan where it appeared that al-Qaeda had been training hackers to break into the computer systems of dams, power grids and nuclear plants. Computer failures may have played a role in the vast power black-outs in north-eastern America and parts of Canada that occurred at the same time.

However, according to Bruce Schneier, a leading expert on network security, only one instance so far deserves to be called cyber-terrorism. In 2000, a hacker named Vitek Boden broke into the computers of an Australian sewage plant and leaked raw effluent into rivers and parks, killing fish but no people. However, Mr Boden was no

ordinary terrorist. Not only had he helped to design and install the system that he attacked, but even with his inside knowledge he had considerable difficulty breaking in.

Terrorists may try more sinister acts. Nonetheless, the internet is a surprisingly difficult medium for them. Malicious code has the potential to cause huge annoyance and disruption. But for people intent on carnage and terror, rather than disruption, blowing oneself up or similar low-tech methods remain far more attractive.

A better word for the threat of internet crime is therefore "cyber-hooliganism", says Mr Schneier. Less than 1% of recent computer attacks originated in countries that America considers breeding grounds for terrorists; the vast majority came from inside America itself. Hackers are more likely to be geeky teens on an ego trip, or greedy crooks hoping to steal money online, than Islamic fundamentalists.

Gone phishing

The promise of the internet knows few bounds: economists think it can boost productivity, efficiency and prosperity much further; entrepreneurs are still excited by its facilitation of online commerce; and more and more consumers prefer it to shops. To realise its full potential, however, the net has to become more trustworthy.

Yet it is rapidly becoming less so. The Blaster worm and SoBig virus that attacked this summer caused estimated losses of \$35 billion. Attacks are getting more frequent, as well as more insidious, relying less often on viruses (which require human action, such as double-clicking on an e-mail attachment) and more often on worms (which propagate by ▶▶

▶ themselves through any unprotected connections on the network). This means that the threat can only grow as “always-on” broadband connections to the internet replace dial-up access, and as ever more devices in addition to PCs are connected.

Attacks are also happening faster. A few years ago, it typically took virus writers a year to exploit a software vulnerability announced by a vendor. This gap between disclosure of a flaw and attack has been shrinking. For the Slammer worm in January it was six months, and for Blaster in August a mere three weeks. It is almost three weeks now since Microsoft brought out its patch for Windows 2000.


Attacks are also more intense and brief. Slammer infected 90% of vulnerable computers within ten minutes. Future attacks, says Gerhard Eschelbeck, the technology boss of Qualys, a network-security monitoring firm, will do their damage within a couple of minutes. Qualys says that it takes organisations an average of one month to patch their known vulnerabilities.

Viruses and worms, moreover, are only one form of internet crime. Brightmail, the world’s market leader in filtering e-mails for fraud and spam, recently found that 10% of all e-mails were scams of one sort or another. Nigerian letters are probably notorious enough by now to be more comical than dangerous. But a lot of fraud is cunning. This includes brand spoofs—e-mails that pretend to come from famous and trusted consumer companies—fake web pages, phoney press releases, and “phishing”, which tricks recipients into giving out sensitive information, such as credit-card numbers.

The gizmos fight back

The resulting anxiety naturally suits vendors of protection technologies, whose sales have been rising sharply. Sometimes the vendors seem to be peddling fear, and it is working. Most companies and governments nowadays use firewalls (devices to keep malicious code out of their internal networks), intrusion-detection systems (which analyse what gets past the firewalls) and similar technologies. Consumers also increasingly have anti-virus software on their computers, though many of them fail to keep it up-to-date.

These gizmos work up to a point. Jerry Ungermann, the president of Check Point, the world’s largest vendor of firewalls, boasts that none of his customers was affected by Blaster because Check Point was so quick to put the appropriate defences into its products. Rival vendors of anti-virus software often compete fiercely in their marketing, but share information as soon as a new virus appears. VeriSign, a company that manages the domain-name systems for .com and .net, is evolving into a sort of CIA of the net, spotting suspicious traffic early and warning those at risk.



Protective “good” code, however, is not by itself enough to fight off incoming “evil” code. As with crime in the physical world, the efforts to fend off break-ins need the support and sanctions of the law. Lawrence Lessig, a professor at Stanford University and an expert on cyberlaw, says that when it comes to cyberspace, policymakers have so far shown themselves to be consistently “stupid and bribable”. How else, he asks, to explain the curious hierarchy of their current priorities. Online copyrights come at the top because of the powerful lobbying of music companies, which are better described as firms faced with a rapidly eroding business model than as victims of crime. Near the bottom comes the online privacy of millions of consumers.

Though more government action will undoubtedly be needed, caution is also in order when considering new laws against cybercrime, lest they make matters worse. This is especially important because most of the experts who advise the lawmakers are not disinterested parties. Qualys’s Mr Eschelbeck, for instance, thinks Congress should pass a law requiring companies to subscribe to automated audits of their systems, which happens to be the service provided by Qualys.

All roads lead to Microsoft

The issue of commercial interests interfering with sound responses becomes especially acute when the debate turns to Microsoft, the world’s largest software company. Ask, for instance, Dan Geer, an expert on software security and a top executive of @Stake, a security consulting firm. In September, he led a group that wrote a report blaming Microsoft’s virtual “monoculture” in operating systems for the internet’s frailty. No sooner was the report published than he found himself out of a job. @Stake, which counts Microsoft among its customers, “fired me by press release, retroactively and in public,” he says.

The gist of Mr Geer’s argument is that

Microsoft has over the years created “unacceptable levels of complexity” in its computer code. It has done so because its main objective has been to lock users into its software by tying the Windows operating system together with applications such as Word, Explorer and Outlook. Complexity is “the enemy of security”, says Mr Geer’s report, since “the defender has to counter all possible attacks; the attacker only has to find one unblocked means of attack.” Moreover, complexity feeds on itself since “fixing a known flaw is likely to introduce a new, unknown flaw.”

On this analysis, many of today’s problems stem from Microsoft’s success in creating a virtual monopoly. Some 94% of PCs run on Windows. So nearly all the computers on the periphery of the internet, where the users are lay people rather than professional network-administrators, rely on the same software, which happens to be of Byzantine complexity. This practically invites hackers to attack these machines. A single good hit at Windows could take down the whole system.

Not surprisingly, Microsoft bristles at this line of thought. The only reason the firm has been bundling the operating system with applications is that customers want it to, says Mike Nash, a Microsoft executive in charge of security issues. He finds it “personally insulting that people think our motivation is anything else.”

Mr Nash also denies that Windows’ code is less secure than other operating systems’, such as Linux or Apple’s Mac OS X. Scott Charney, another Microsoft executive, goes further and defends the monoculture. If one operating system is dominant, he says, companies can save costs by training IT staff only once, and security updates are easier since there is only one source of the patches that mend flaws.

But the patches often create more security problems than they fix, and there is a fear that Microsoft might use such regular access to desktops to keep rival software-makers away, thus reinforcing the source of the original problem, its monoculture. “If you don’t trust us to download our patch, then you shouldn’t be running our software,” counters Mr Charney, as if consumers had a real choice.

Nonetheless, even if Microsoft, with its disproportionate share of the market, constitutes a disproportionate share of the problem, it is not clear what to do about it. Many of the arguments sound tediously reminiscent of the American government’s prolonged antitrust case against the firm in the late 1990s. Even Mr Geer, for instance, is not advising that Microsoft be broken up. Instead, he wants Microsoft to make its applications run on any rival platform, and to publish the interface protocols that will allow rival applications to spring up and survive. This might lead to some biodiversity of code. ▶▶

► Mr Schneier, one of the authors of the report submitted by Mr Geer, proposes a more fundamental solution. Cybercrime, he argues, is “not a technological problem; it’s an economic problem: the incentives aren’t there for smart people to solve the problem.” The culprit, in other words, is the licences that require buyers of new software to click their assent that the vendor is not liable for any flaws in its software. As long as software vendors—and this is not specific to Microsoft—cannot be held liable for security issues, Mr Schneier says, the economic incentives are stacked toward adding bells and whistles and shipping upgrades fast, rather than toward writing simpler, safer software.

Changing the law so that liability does rest at least in part with vendors, he argues, would align the incentives properly and lead to other good things as well. Software companies, just like firms in other industries, would buy product-liability insurance. Insurance companies would respond by pricing the risk, in effect voting on the security of each product. Just as companies that install sprinklers in their warehouses pay lower premiums and have a competitive edge over rivals that do not, software companies that write safer code would have an economic advantage.

No responsibility without liability?

In what could become a precedent, the first lawsuit against Microsoft on product-liability grounds was filed in a court in Los Angeles in October, accusing the company of violating California’s consumer-protection laws by selling shoddy software. Legally, the approach may be controversial. Suing Microsoft over a Windows virus is not quite analogous to suing, say, a car-maker for selling vehicles that tip over while being driven. In the first case, a third party, the hacker, is committing a crime by exploiting a weakness in the product; in the latter case, the product fails without outside criminal intervention. A better analogy may be suing a maker of bullet-proof vests whose products fail to protect their wearers against bullets.

Some argue that the cost of insuring against product liability might stifle software innovation. Not so, says Stanford’s Mr Lessig. A small upstart company making a small operating system would not present much of a target to hackers, and would thus pay negligible premiums. In any case, even if caution did lead to a few programs not being written, says Mr Schneier, so what? America’s Food and Drug Administration can be said to stifle innovation too in so far as it leads to the marketing of fewer but safer drugs. In software, the risks are now simply too great not to make a similar trade-off, he says.

Microsoft argues that the constant attacks against its software—4,000 so far against Windows, according to Symantec,

an anti-virus company—are threatening its brand and business prospects even without litigation. The argument that without product liability companies won’t pay attention to security “is just not true,” says Mr Charney. Microsoft has already pulled out all the stops, he argues, and is retraining its programmers, reviewing their code and changing its entire culture. Unfortunately, security has to be built into software from the beginning—patches are just what their name suggests.

First find the users and abusers

Concentrating entirely on the accountability of software vendors is like fighting burglary by leaning on the makers of alarm systems. A parallel approach to the problem of internet insecurity is, therefore, to focus on the internet’s users, discouraging bad behaviour and ensuring that criminals can be traced. Legally, however, that could become as controversial as product liability. Mr Lessig suggests using a bounty system to catch hackers, which might involve enlisting those most able to catch them—namely, other hackers. “I’d bet my job that it works,” he says.

The issue boils down to the question of how much anonymity society can tolerate on the internet. Drivers’ licences and registration plates dramatically reduce the incidence of hit-and-run accidents. Crack cocaine is never bought by credit card. If everybody on the internet were easily traceable, people would think twice about hacking. “I’m kind of a fan of eliminating anonymity,” says Alan Nugent, the chief technologist at Novell, a software company, “if that is the price for security.”

The internet is heading in this direction already. Enrique Salem, Brightmail’s chief executive, says that all e-mail in future will either be authenticated or be sent into a quarantined in-box where few will dare to click. The sender’s authentication may

well be tied to a driving licence, social-security number or passport. An entire industry has sprung up to work on other forms of identification, such as the biometric scanning of irises or hands.

All this may not be pleasing to libertarians, who envisioned the internet as offering individuals the cover of relative obscurity. What use is the network to dissidents in China if the Communist Party is watching everything they do online? And what use is the internet, whose whole point was to connect people, if it is balkanised into separate, walled subnets?

The reality, however, is that the internet is already balkanised. Companies and governments have intranets, where users’ privileges depend on their log-in. Virtual private networks (VPNs) traverse the public internet like guarded convoys. For example, employees at Merrill Lynch, an investment bank, cannot check their Hotmail or Yahoo! e-mail accounts while surfing the internet at work.

The proper analogy for what the internet might evolve into, says Novell’s Mr Nugent, is a public library, a place where readers can browse in relative anonymity, but only until they take a book out, at which point they have to identify themselves. The degree of traceability varies with what one does in such a place.

To preserve freedom further, suggests Mr Lessig, anonymity could be replaced by pseudonymity. It might become legal, for instance, to have credit cards for online transactions under different names, as long as these could still be traced to the individual owner. The challenge is to set the legal hurdles for online search warrants high enough so that governments cannot abuse their power. But at the same time to keep them low enough so that criminals can be found and stopped. In this respect, the online world should be no different from the real one. ■

