Last year's Stuxnet virus attack represented a new kind of threat to critical infrastructure.

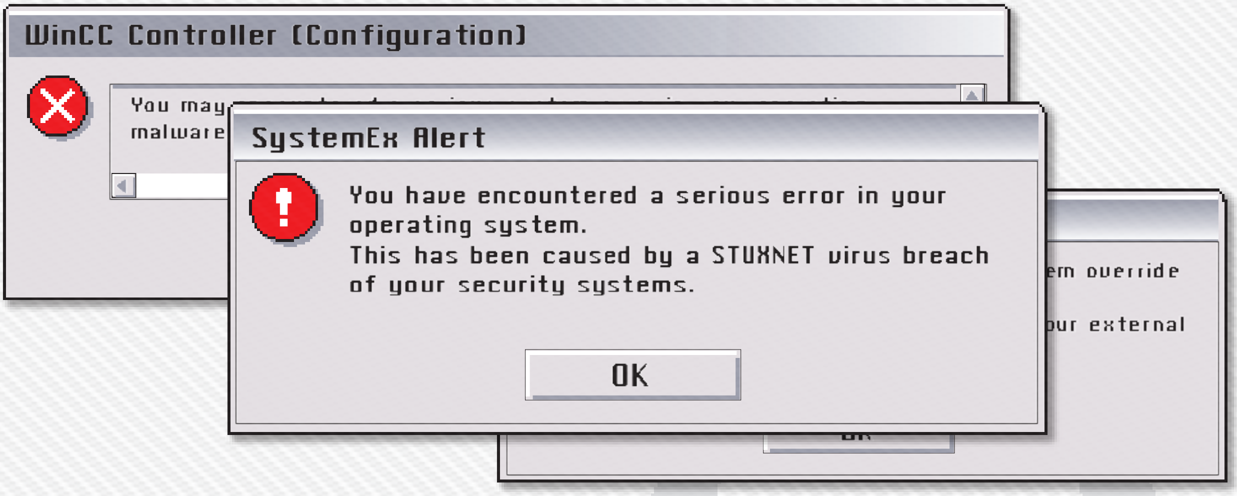

You have encountered a serious error in your operating system.
This has been caused by a STUXNET virus breach of your security systems.

IS THIS THE START OF CYBERWARFARE?

by Sharon Weinberger

Just over a year ago, a computer in Iran started repeatedly rebooting itself, seemingly without reason. Suspecting some kind of malicious software (malware), analysts at VirusBlokAda, an antivirus-software company in Minsk, examined the misbehaving machine over the Internet, and soon found that they were right. Disturbingly so: the code they extracted from the Iranian machine proved to be a previously unknown computer virus of unprecedented size and complexity.

On 17 June 2010, VirusBlokAda issued a worldwide alert that set off an international race to track down what came to be known as Stuxnet: the most sophisticated computer malware yet found and the harbinger of a new generation of cyberthreats. Unlike conventional malware, which does its damage only in the virtual world of computers and networks, Stuxnet would turn out to target the software that controls pumps, valves, generators and other industrial machines.

"It was the first time we'd analysed a threat that could cause real-world damage, that could actually cause some machine to break, that might be able to cause an explosion," says Liam O Murchu, chief of security response for the world's largest computer-security firm, Symantec in Mountain View, California.

Stuxnet provided chilling proof that groups or nations could launch a cyberattack against a society's vital infrastructures for water and energy. "We are probably just now entering the era of the cyber arms race," says Mikko Hypponen, chief research officer for F-Secure, an antivirus company based in Helsinki.

Worse yet, the Stuxnet episode has highlighted just how inadequate are society's current defences — and how glaring is the gap in cybersecurity science.

Computer-security firms are competitive in the marketplace, but they generally respond to a threat such as Stuxnet with close collaboration behind the scenes. Soon after VirusBlokAda's alert, for example, Kaspersky Lab in Moscow was working with Microsoft in Redmond, Washington, to hunt down the vulnerabilities that the virus was exploiting in the Windows operating system. (It was Microsoft that coined the name Stuxnet, after one of the files hidden in its code. Technically, Stuxnet was a 'worm', a type of malware that can operate on its own without needing another program to infect. But even experts often call it a 'virus', which has become the generic term for self-replicating malware.)

One of the most ambitious and comprehensive responses was led by Symantec, which kept O Murchu and his worldwide team of experts working on Stuxnet around the clock for three months. One major centre of operations was Symantec's malware lab in Culver City, California, which operates like the digital equivalent of a top-level biological containment facility. A sign on the door warns visitors to leave computers, USB flash drives and smart phones outside: any electronic device that passes through that door, even by mistake, will stay there. Inside the lab, the team began by dropping Stuxnet into a simulated networking environment so that they could safely watch what it did. The sheer size of the virus was staggering: some 15,000 lines of code, representing an estimated 10,000 person hours in software development. Compared with any other virus ever seen, says O Murchu, "it's a huge amount of code".

Equally striking was the sophistication of that code. Stuxnet took advantage of two digital certificates of authenticity stolen from respected companies, and exploited four different 'zero day vulnerabilities' — previously unidentified security holes in Windows that were wide open for hackers to use.

Then there was the virus's behaviour. "Very quickly we realized that it was doing something very unusual," recalls O Murchu. Most notably, Stuxnet was trying to talk to the programmable logic controllers (PLCs) that are used to direct industrial machinery. Stuxnet was very selective, however: although the virus could spread to almost any machine running Windows, the crucial parts of its executable code would become active only if that machine was also running Siemens Step7, one of the many supervisory control and data acquisition (SCADA) systems used to manage industrial processes.

Many industrial control systems are never connected to the Internet, precisely to protect them from malware and hostile takeover. That led to another aspect of Stuxnet's sophistication. Like most other malware, it could spread over a network. But it could also covertly install itself on a USB drive. So all it would take was one operator unknowingly plugging an infected memory stick into a control-system computer, and the virus could explode into action.

MURKY MOTIVES

It still wasn't clear what Stuxnet was supposed to do to the Siemens software. The Symantec team got a clue when it realized that the virus was gathering information about the host computers it had infected, and sending the data back to servers in Malaysia and Denmark — presumably to give the unknown perpetrators a way to update the Stuxnet virus covertly. Identifying the command and control servers didn't allow Symantec to identify the perpetrators, but they were able to convince the Internet service providers to cut off the perpetrators' access, rerouting the traffic from the infected computers back to Symantec so that they could eavesdrop. By watching where the traffic to the servers was coming from, O Murchu says, "we were able to see that the majority of infections were in Iran" — at least 60% of them. In fact, the infections seemed to have been appearing there in waves since 2009.

The obvious inference was that the virus had deliberately been directed against Iran, for reasons as yet unknown. But the Symantec investigators couldn't go much further by themselves. They were extremely knowledgeable about computers and networking, but like most malware-protection teams, they had little or no expertise in PLCs or SCADA systems. "At some point in their analysis they just couldn't make any more sense out of what the purpose of this thing was, because they were not able to experiment with the virus in such a lab environment," says Ralph Langner, a control-system security consultant in Hamburg, Germany.

Langner independently took it upon himself to fill that gap. Over the summer, he and his team began running Stuxnet in a lab environment equipped with Siemens software and industrial control systems, and watching how the virus interacted with PLCs. "We began to see very strange and funny results immediately, and I mean by that within the first day of our lab experiment," he says.

> "We are probably just now entering the era of the cyber arms race."

Those PLC results allowed Langner to infer that Stuxnet was a directed attack, seeking out specific software and hardware. In mid-September 2010, he announced on his blog that the evidence supported the suspicion that Stuxnet had been deliberately directed against Iran. The most likely target, he then believed, was the Bushehr nuclear power plant.

INDUSTRIAL SABOTAGE

Speculative though Langner's statements were, the news media quickly picked up on them and spread the word of a targeted cyberweapon. Over the next few months, however, as Langner and others continued to work with the code, the evidence began to point away from Bushehr and towards a uranium-enrichment facility in Natanz, where thousands of centrifuges were separating the rare but fissionable isotope uranium-235 from the heavier uranium-238. Many Western nations believe that this enrichment effort, which ostensibly provides fuel for nuclear power stations, is actually aimed at producing a nuclear weapon. The malware code, according to Langner and others, was designed to alter the speed of the delicate centrifuges, essentially causing the machines to spin out of control and break.

That interpretation is given credence by reports from the International Atomic Energy Agency (IAEA) in Vienna, which document a precipitous drop in the number of operating centrifuges in 2009, the year that many observers think Stuxnet first infected computers in Iran.

True, the evidence is circumstantial at best. "We don't know what those machines were doing" when they weren't in operation, cautions Ivanka Barszashka, a Bulgarian physicist who studied Iranian centrifuge performance while she was working with the Federation of American Scientists in Washington DC. "We don't know if they were actually broken or if they were just sitting there." Moreover, the Iranian government has officially denied that Stuxnet destroyed large numbers of centrifuges at Natanz, although it does acknowledge that the infection is widespread in the country. And IAEA inspection reports from late 2010 make it clear that any damage was at most a temporary setback: Iran's enrichment capacity is higher than ever.

However, if Natanz was the target, that does suggest an answer to the mystery of who created Stuxnet, and why.

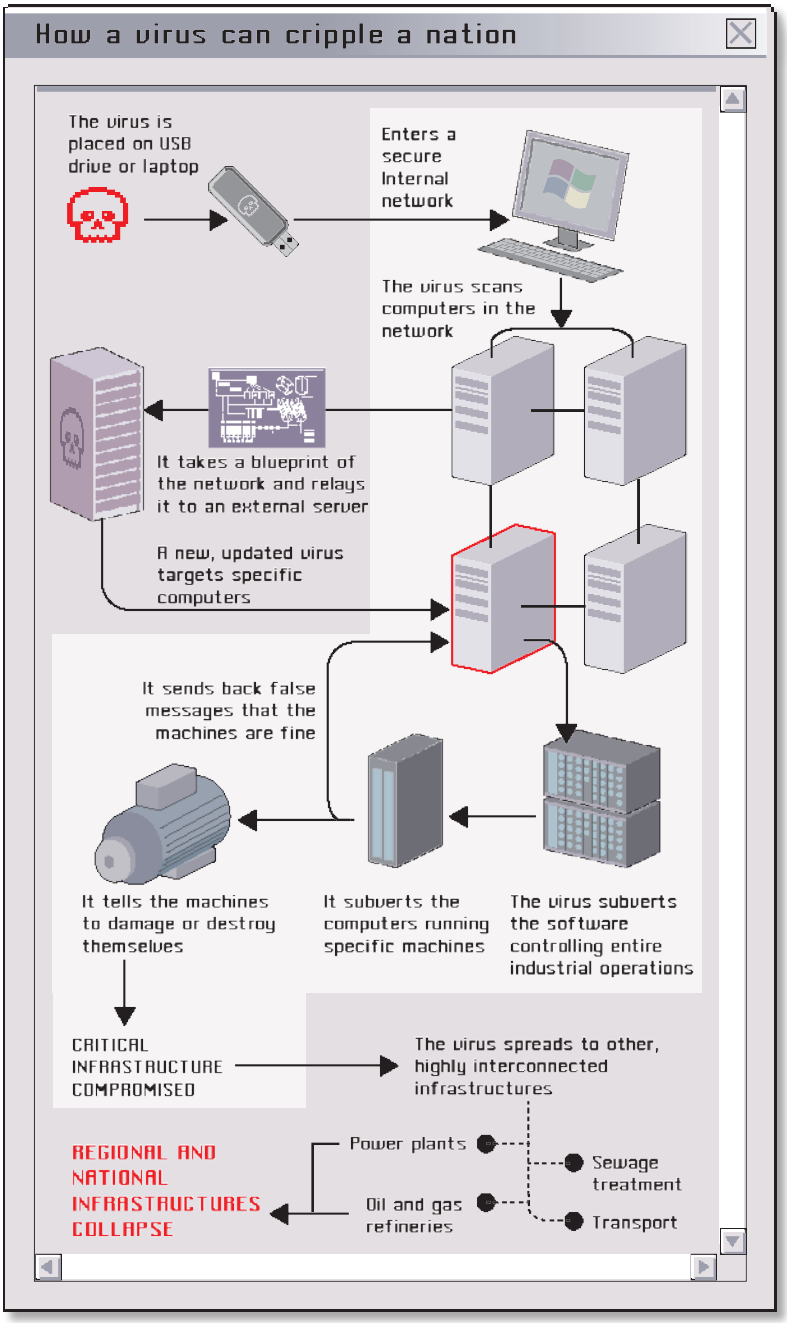
How a virus can cripple a nation
The virus is placed on USB drive or laptop
Enters a secure Internal network
The virus scans computers in the network
It takes a blueprint of the network and relays it to an external server
A new, updated virus targets specific computers
It sends back false messages that the machines are fine
It tells the machines to damage or destroy themselves
It subverts the computers running specific machines
The virus subverts the software controlling entire industrial operations
CRITICAL INFRASTRUCTURE COMPROMISED
The virus spreads to other, highly interconnected infrastructures
Power plants
Oil and gas refineries
Sewage treatment
Transport
REGIONAL AND NATIONAL INFRASTRUCTURES COLLAPSE

Given the knowledge required — including expertise in malware, industrial security and the specific types and configurations of the industrial equipment being targeted — most Stuxnet investigators concluded early on that the perpetrators were backed by a government.

Governments have tried to sabotage foreign nuclear programmes before, says Olli Heinonen, a senior fellow at the Belfer Center for Science and International Affairs at Harvard University in Cambridge, Massachusetts, and former deputy director-general of the IAEA. In the 1980s and 1990s, for example, Western governments orchestrated a campaign to inject faulty parts into the network that Pakistan used to supply nuclear technology to countries such as Iran and North Korea. Intelligence agencies, including the US Central Intelligence Agency, have also made other attempts to sell flawed nuclear designs to would-be proliferators. “Stuxnet,” says Heinonen, “is another way to do the same thing.”

Langner argues that the government behind Stuxnet is that of the United States, which has both the required expertise in cyberwarfare and a long-standing goal of thwarting Iran’s nuclear ambitions. Throughout the summer of 2010, while Langner, Symantec and all the other investigators were vigorously trading ideas and information about Stuxnet, the US Department of Homeland Security maintained a puzzling silence, even though it operates Computer Emergency Readiness Teams (CERTs) created specifically to address cyberthreats. True, the CERT at the Idaho National Laboratory outside Idaho Falls, which operates one of the world’s most sophisticated testbeds for industrial control systems, did issue a series of alerts. But the first, on 20 July 2010, came more than a month after the initial warning from Belarus and contained nothing new. Later alerts followed the same pattern: too little, too late. “A delayed clipping service,” said Dale Peterson, founder of Digital Bond, a SCADA security firm in Sunrise, Florida, on his blog.

“There is no way that they could have missed this problem, or that this is all a misunderstanding. That’s just not possible,” says Langner, who believes that the Idaho lab’s anaemic response was deliberate, intended to cover up the fact that Stuxnet had been created there.

But even Langner has to admit that the evidence against the United States is purely circumstantial. (The US government itself will neither confirm nor deny the allegation, as is its practice for any discussion of covert activity.) And the evidence against the other frequently mentioned suspect, Israel, is even more so. Symantec, for example, points out that a name embedded in Stuxnet’s code, Myrtus, could be a reference to a biblical story about a planned massacre of Jews in Persia. But other investigators say that such claims are beyond tenuous. “There are no facts” about Israel, declares Jeffrey Carr, founder and chief executive of Taia Global, a cybersecurity consulting company in Tysons Corner, Virginia.

THE AFTERMATH

The ‘who?’ may never be discovered. Active investigation of Stuxnet effectively came to an end in February 2011, when Symantec posted a final update to its definitive report on the virus, including key details about its execution, lines of attack and spread over time. Microsoft had long since patched the security holes that Stuxnet exploited, and all the antivirus companies had updated their customers’ digital immune systems with the ability to recognize and shut down Stuxnet on sight. New infections are now rare — although they do still occur, and it will take years before all the computers with access to Siemens controllers are patched.

If Stuxnet itself has ceased to be a serious threat, however, cybersecurity experts continue to worry about the larger vulnerabilities that it exposed. Stuxnet essentially laid out a blueprint for future attackers to learn from and perhaps improve, say many of the investigators who have studied it. “In a way, you did open the Pandora’s box by launching this attack,” says Langner of his suspicions about the United States. “And it might turn back to you guys eventually.”

Cybersecurity experts are ill-prepared for the threat, in part because they lack ties to the people who understand industrial control systems. “We’ve got actually two very different worlds that traditionally have not communicated all that much," says Eric Byres, co-founder and chief technology officer of Tofino Industrial Security in Lantzville, Canada. He applauds Symantec, Langner and others for reaching across that divide. But the effort required to make those connections substantially delayed the investigation.

The divide extends into university computer-science departments, say Byres, himself an ex-academic. Researchers tend to look at industrial-control security as a technical problem, rather than an issue requiring serious scientific attention, he says. So when graduate students express interest in looking at, say, cryptography and industrial controls, they are told that the subject is not mathematically challenging enough for a dissertation project.

"I'm not aware of any academic researchers who have invested significantly in the study of Stuxnet," agrees Andrew Ginter, director of industrial security for the North American group of Waterfall Security Solutions, based in Tel Aviv, Israel. Almost the only researchers doing that kind of work are in industrial or government settings — among them a team at the Idaho National Laboratory working on a next-generation system called Sophia, which tries to protect industrial control systems against Stuxnet-like threats by detecting anomalies in the network.

One barrier for academics working on cybersecurity is access to the malware that they must protect against. That was not such a problem for Stuxnet itself, because its code was posted on the web shortly after it was first identified. But in general, the careful safeguards that Symantec and other companies put in place in secure labs to protect the escape of malware may also inadvertently be a barrier for researchers who need to study them. "If you're doing research into biological agents, it's limited groups that have them and they are largely unwilling to share; the same holds true for malware," says Anup Ghosh, chief scientist at the Center for Secure Information Systems at George Mason University in Fairfax, Virginia. "To advance the field, researchers need access to good data sets," says Ghosh, who was once a programme manager at the US Defense Advanced Research Projects Agency, and is now working on a malware detector designed to identify viruses on the basis of how they behave, rather than on specific patterns in their code, known as signatures.

Academic researchers are also inhibited by a certain squeamishness about digital weaponry, according to Herb Lin, chief scientist at the Computer Science and Telecommunications Board of the US National Research Council in Washington DC. He points out that to understand how to guard against cyberattacks, it may help to know how to commit them. Yet teaching graduate students to write malware is "very controversial", he says. "People say, 'What do you mean: you're training hackers?'"

PREPARING FOR THE NEXT ATTACK

A study last year by the JASON group, which advises the US government on science and technology matters, including defence, found broad challenges for cybersecurity (JASON *Science of Cyber-Security*; MITRE Corporation, 2010). Perhaps most important was its conclusion that the field was "underdeveloped in reporting experimental results, and consequently in the ability to use them".

Roy Maxion, a computer scientist at Carnegie Mellon University in Pittsburgh, Pennsylvania, who briefed JASON, goes further, saying that cybersecurity suffers from a lack of scientific rigour. Medical professionals over the past 200 years transformed themselves from purveyors of leeches to modern scientists with the advent of evidence-based medicine, he notes. "In computer science and in computer security in particular, that train is nowhere in sight."

Computer science has developed largely as a collection of what Maxion calls "clever parlour tricks". For example, at one conference, the leading paper showed how researchers could read computer screens by looking at the reflections off windows and other objects. "From a practical point of view, anyone in a classified meeting would go, 'pooh'," he says. "In places where they don't want you to know [what's on the computer screen], there are no windows. Yet, that was the buzz that year."

Maxion sees an urgent need for computer-science and security curricula to include courses in traditional research methods, such as experimental design and statistics — none of which is currently required. "Why does it matter?" he asks. "Because we don't have a scientific basis for investigating phenomena like Stuxnet, or the kind of defences that would be effective against it."

National Security Warning

Central banking and health-care systems have failed. National security measures are no longer operational.
Wipe and reboot all systems immediately.

Also troubling for many of the Stuxnet investigators was the US government's lacklustre response to the virus (assuming that it was not the perpetrator). Stuxnet represents a new generation of cyberweapon that could be turned against US targets, but there is no evidence that the government is making the obvious preparations for such an attack — for example, plans for a coordinated response that pools resources from academia, research institutes and private business.

Other countries seem to be taking the threat more seriously. Some of China's universities and vocational colleges have reportedly forged strong connections with the military to work on cybersecurity, for example. And Israel also seems to be exploiting its computing expertise for national security. A few months before the discovery of Stuxnet, Yuval Elovici, a computer scientist and director of Deutsche Telekom Laboratories at Ben-Gurion University of the Negev in Beersheba, Israel, told *Nature* that he was working closely with the country's Ministry of Defense on cybersecurity. He presciently warned that the next wave of cyberattacks would be aimed at physical infrastructures. "What would happen if there were a code injection into SCADA? What if someone would activate it suddenly?" Elovici asked. He and other experts have been warning for several years now that such an attack on SCADA systems controlling the electricity grid could spark nationwide blackouts, or that the safety systems of power plants could be overridden, causing a shutdown or a serious accident. Similar disruptions could hit water and sewage systems, or even food processing plants.

Such attacks, Elovici warned, are both realistic and underestimated. Asked how bad one would be, Elovici was unequivocal. "I think," he said, "it would be much stronger than the impact of setting several atomic bombs on major cities." ■ SEE EDITORIAL P.127

Sharon Weinberger *is an Alicia Patterson Foundation fellow based in Washington DC.*