Researchers figure out how to hack tens of thousands of servers
Timothy B. Lee

(Photo by Frédéric BISSON

Security researchers at the University of Michigan have found a potentially devastating security vulnerability that afflicts at least 40,000 servers on the Internet. The researchers say the flaw could allow hackers to compromise certain servers manufactured by Supermicro from anywhere on the Internet. Tens of thousands of servers produced by other vendors could also be at risk.

In the past, certain administrative tasks, such as reinstalling an operating system, required an administrator to travel to the physical location of the server. But more recently, administrators have demanded sophisticated remote management capabilities that allow them to do their jobs without ever setting foot in their data centers.

To address that demand, an Intel-led consortium developed a standard called the Intelligent Platform Management Interface (IPMI). It essentially provides an entire second computer, called the baseboard management controller (BMC), that's built into the server. It has the ability to reboot the server, reinstall software and perform other administrative tasks.
Because the BMC is so powerful, having it compromised would be catastrophic, giving a hacker complete control of the server. And that's a real risk because the BMC has a user-friendly, Web-based administrative interface that makes it a server in its own right. Yet the software on the BMC has not gotten the kind of rigorous security analysis traditionally applied to server software. Until now.

"There's just so much that's broken"

A team from the University of Michigan purchased a brand of server manufactured by Supermicro, a SYS-5017C-LF rack-mounted system. Anthony Bonkoski, the Michigan graduate student who performed the research, describes the security problems he found as "pretty awful." He says that "I don't know where to begin because there's just so much that's broken."

Bonkoski presented his findings on Tuesday at the Usenix security conference in Washington. He says that after the presentation, he "looked over and saw someone getting enraged, saying, 'These are elementary mistakes that we teach our undergrads not to make.'"

A basic principle of Web security is that a server should check the integrity of information submitted by a user before using it. For example, in an attack known as a buffer overflow, a malicious user submits a text string that's longer than the space set aside for storing it. This causes the data to "overflow" its assigned space and overwrite other data that the user shouldn't be able to modify. A clever hacker can use this technique to modify a server's software and thereby take control of the machine.

Preventing buffer overflows isn't hard. You just have to check the length of strings before using them. But the programmers who developed the remote-management software on Supermicro's server didn't do it. Or at least they didn't do it correctly. The Michigan team found that limits on string lengths are enforced only by the user's browser. And that makes it trivial for an attacker to override those limits.

One place where this buffer-overflow vulnerability existed was in the "username" and "password" fields used to log into the BMC's Web-based management interface. And that means that if the BMC is connected to the Internet, anyone on the Internet can attack it, get control of the BMC and thereby gain complete control of the server.

To guard against these risks, good security practices dictate that a management interface like IPMI should be available only on a private management network, not the public Internet. But Bonkoski found that Supermicro's user instructions didn't warn of this concern. To the contrary, the company's user guide "provides detailed instructions on which firewall ports to open to allow remote connections."

Thousands of vulnerable servers

To figure out how widespread this problem was, Bonkoski and his team scanned the Internet for servers running Supermicro's implementation of the IPMI management software. They found 41,545 servers that appeared to be running Supermicro's insecure BMC software. They didn't try to compromise any of these systems for obvious ethical reasons. But based on their research, they "conservatively estimate that it would take less than an hour to launch successful parallel attacks" against all of those 41,545 servers.

That's not all. During their scan, they also found 40,413 servers running Dell's implementation of IPMI and 23,376 servers running HP's IPMI implementation. Bonkoski hasn't examined the Dell and HP servers to see if they're vulnerable to similar attacks. But he argues that given the potential for vulnerabilities and the catastrophic consequences if they're compromised, it's irresponsible for these interfaces to be connected to the public Internet at all.

"I don't want to pin too much blame on Supermicro because it was really a third party vendor" that provided flawed software to Supermicro, Bonkoski says. Still, he says, Supermicro's name was attached to the devices so it will get a share of the blame.

Contacted by e-mail, a Supermicro spokesman wrote that "we've been aware of this report and have published support information on our IPMI page." The support information is a PDF document advising customers to ensure that servers' IPMI interfaces are not connected to the public Internet. It also suggested changing default passwords on the BMC and taking other precautions to make the kinds of attacks the Michigan researchers describe more difficult.

Yet it's not clear if the 40,000 Supermicro customers who have already exposed their servers to the Internet will hear about the vulnerability in time to avoid having their servers compromised.

The challenge of embedded security

Why is the security of Supermicro's IPMI implementation so flawed? Bonkoski believes the issue is largely cultural. For more than a decade, mainstream programmers have been trained to make security a high priority. That training has been reinforced by practical experience: When they didn't take security seriously, their software got hacked and their employers lost millions of dollars.

But Bonkoski says the BMC's software was written by programmers from a different community, known as embedded computing. Embedded device programmers produce software for small, low-powered chips that run in our cars, appliances and industrial equipment. These devices are not traditionally connected to the Internet, and so the embedded software community hasn't learned the painful lessons of the last decade.

"I would compare it to the 1990s, when we started connecting PCs to the Internet and everything broke," Bonkoski says. Back then, PC software developers were as oblivious to security considerations as embedded software developers are today.

Improving the security of embedded computing devices is a hot topic in security research right now. Last month, security researchers revealed that two popular car models were vulnerable to hacking. While these attacks required physical access to the vehicles, the increasingly sophisticated wireless capabilities of future cars will create the risk of hackers sabotaging our vehicles from thousands of miles away, with potentially deadly consequences.

But there's probably no turning back. The advantages of being connected to the Internet, for both cars and server management interfaces, makes the pressure to do so irresistible. Remote management interfaces like IPMI create security risks, but Bonkoski believes they're too useful for systems administrators to live without.

"If you're going to have large data centers, you're going to have a system like this eventually," he says. "From a management perspective, you want something that can install an operating system. You want something that can power-cycle the machine and do this monitoring. If you're going to try to automate this, you have to have that hardware/software layer in there." And so, embedded device vendors are going to face a steep learning curve. They're going to have to learn to anticipate the security problems that inevitably crop up when previously offline devices are suddenly connected to the Internet. And they're going to have to change their engineering culture so that security is a top priority, not an afterthought.

Fonte: The Washington Post, Washington, 14 Aug. 2013, International.