MIT Technology Review

COMMUNICATIONS NEWS

4 COMMENTS

“Massive” Cyberattack Wasn’t Really So Massive

A decade-old fix could have easily stopped this weekend’s attack on an anti-spam company, but the truth is many Web companies simply ignore such fixes.

By David Talbot on March 29, 2013

An attack that disrupted Internet service over the past week would have been stopped by a simple Web server configuration fix that’s been understood for a decade but is widely ignored by Web companies, experts say.

The prolonged assault targeted Spamhaus, a European nonprofit that reports where spam is coming from and publishes a list of implicated Web servers. The apparent flashpoint was the addition of CyberBunker, a Dutch data-storage company, to its roster.

The unidentified attackers used a botnet – a network of infected ordinary computers – to attack Spamhaus’s website and then the servers of CloudFlare, a content-delivery company that stepped in to help Spamhaus manage the influx of traffic. The attack also affected regional Internet servers that are transit points for not only the two targeted companies but also many others.

While some observers have suggested that the scale of the attack was smaller than most reports indicated, according to a blog by the Austrian Computer Emergency Response Team (CERT), the attack caused “disruption in some parts of the Internet.”

The kind of attack that occurred is called “distributed denial-of-service” because many computers are tricked into sending chunks of data at one target, overwhelming it. This attack took advantage of a weakness in domain-name servers, or DNS servers, where typed Web addresses are resolved into the numerical codes that correspond to the machines that hold the relevant information.

The attack involved sending DNS servers requests forged to look as if they came from the target. These DNS servers responded by overwhelming the target with data it didn’t actually ask for. The impact can be amplified because the DNS servers – depending how they are configured – can be asked to send large amounts of data.

CloudFlare said the attack was as large as 300 gigabits per second, the largest such attack ever reported, and said delays were felt by millions. Some researchers said bloggers and news sources overstated the impact of the attack, for example in posts like this one, which claimed the attack “almost broke the Internet.”

But overall, the attack did not have a profound impact, according to a post by Richard Steenbergen of nLayer Communications, one of the network providers used by CloudFlare.

“Just the production costs of CNN discussing this were probably higher than the damage this thing might have caused,” says Radu Sion, a computer scientist at Stony Brook University.

What is beyond dispute is that the attack used a trick that has been known for years. Web companies have the ability to configure their servers so that they don’t fall for this ruse, using a method spelled out in a technical document called BCP 38 written by the voluntary Internet Engineering Task Force way back in 2000. But few do.

“Misconfigurations are rampant across the Internet,” says Mike Smith, director of the computer-security response team at Akamai, the Web optimization company based in Cambridge, Massachusetts. “There are tools and configuration guides and best practices for ISPs. But people need to use them and know that this is a problem.”

The reason for the problem is that the Internet suffers from diffusion of responsibility for security. The bottom line for many Web companies is that loose settings don’t really affect them, fixing the settings costs something, and no real penalties are in place for failing to do so. “If I have an open resolver [the technical term for the problem], it doesn’t really impact me from day to day,” Smith says. “Where we

really need to focus on is awareness of these behaviors and helping companies clean up their problem."

Internet attacks not only are costly to individuals and businesses but pose a potential national-security threat (see "Why Obama's Cyber Defense Order Won't Amount to Much" and "Cybersecurity Risk High in Industrial Control Systems").

Sion says that despite the relatively tolerable outcome of this attack, the risks are real and growing. "It will get worse before it gets better, in some cases just because of the proliferation of broadband. As that last mile gets faster, the attackers' capabilities are going to be larger," he says. "Taking over an individual server allows you to send a lot more dangerous traffic than you could five years ago."

Tagged: Computing, Communications, Web, Internet, Akamai