fastly

The state of IoT security

Dr. Jose Nazario
Director of Security Research, Fastly

The state of IoT security

Jose Nazario, Director of Security Research, Fastly
Dr. Jose Nazario is the Director of Security Research at Fastly, and is a recognized expert on cyberthreats to ISPs, network subscribers, and enterprises from cybercrime and malware.

He was previously the Research Director for Malware Analysis at Invincea Labs. Before his work at Invincea Labs he was the senior Manager of Security risk at Arbor Networks. In this capacity he was the head of Arbor security and engineering Response Team. From 2002-2008 he was the ASERT Engineer where he led the creation of Arbor's malware analysis capability. Additionally, he was the QA Engineer at Arbor.

Prior to his work at Arbor, he was an Independent consultant at OpenBSD project focusing on operational security, including forensics, installations and configuration, and hardening for local businesses. He has published several works, most notably his books: "Secure Architectures with OpenBSD" and "Defense and Detection Strategies against Internet Worms."

The Internet of Things (IoT) has created a bounty of opportunities for various industries, and has fundamentally changed the way consumers interact with their devices. But with opportunity also comes risk: in the past year, nefarious groups have used these devices to conduct distributed denial of service (DDoS) attacks — most famously against journalist Brian Krebs and Dyn, which were launched by a massive IoT botnet of hundreds of thousands of infected devices. IoT security remained a trending topic in July 2017, both at hacker convention Defcon and in the news with a German court's conviction in an IoT-based, DDoS-for-hire scheme.

Internet-connected devices are being churned out of factories and infected by malware, or malicious code, at an alarming rate. Armies of compromised IoT devices immediately try to enroll new devices, join a botnet, and participate in large-scale DDoS attacks. As a result, we've recently seen some of the biggest DDoS attacks in history.

The Fastly security team is focused on leveraging our network intelligence to proactively defend the modern web. We took a look at some of the more recent (and troubling) threats on the internet, and found that the emerging IoT market is under attack.

Just how big of a problem is this? We did an analysis of the anatomy of an IoT botnet attack, all the way down to the individual device level – and exposed some interesting data. Read on to learn what we found, and how the IoT industry is responding.

Anatomy of an IoT botnet attack

Here's a breakdown of what we learned:

- On average, an IoT device was infected with malware and had launched an attack within 6 minutes of being exposed to the internet.
- Over the span of a day, IoT devices were probed for vulnerabilities 800 times per hour by attackers from across the globe.
- Over the span of a day, we saw an average of over 400 login attempts per device, an average of one attempt every 5 minutes; 66 percent of them on average were successful.

The majority of attacks were automated attacks run by malicious scripts targeting common IoT devices such as DVRs, IP cameras, and NVRs (network video recorders). The most common malware dropped was intended for IoT and other devices, including processors, as well as hardware platforms used by the automotive industry, electronic meters, healthcare, and more. The scope of these attacks goes far beyond IP cameras and home routers.

Meanwhile, attackers were distributed around the world, with the top 5 locations being:

- 13.5% coming from China
- 9.9% coming from Brazil
- 8.6% coming from Republic of Korea
- 7.1% coming from Vietnam
- 5.8% coming from India

The recent Mirai attacks have focused attention on the threat that IoT places on the broader internet. As 6.4 billion devices come online, that's a lot of firepower. Presently, thinking in the security community is that

IoT vendors created this mess with fully capable Linux-based computers on those devices together with a handful of default usernames and passwords that the bots simply guess at. This enables attackers and now malware to log in, upload arbitrary botnet code, and begin attacks.

Companies and consumers who are running these devices, or anyone deploying the devices, such as broadband providers, need to take some responsibility to keep the hardware from being used in attacks. They need to change the default passwords and disable logins from the open internet. In the long term, however, security standards will need to come into play. To accomplish this, the industry will need to establish requirements for devices to be sold or installed. Big broadband equipment vendors and industry groups like CableLabs are a natural place to work together to address this issue. If not, we'll possibly see the FCC take a role, although enforcing rules on millions of devices won't be an easy task, if it comes to that. Such regulations are also under discussion in Europe. But we've got to figure it out. If we don't do something to keep attackers from turning all these devices into DDoS weapons we'll see more sites go dark. And nobody wants that.

In the next section, we'll take a look at how manufacturers have responded to these threats — we took a couple of IoT devices for a spin to see how they fared on the internet.

The industry's response to emerging threats

Following our first round of research, we did more robust vulnerability research on IoT devices that have been found vulnerable in the past (smart cameras, baby monitors, and light bulbs) and concluded that while malicious probes are constant, manufacturers have taken action to update their firmware and address security holes. An example was the Chinese device manufacturer that recalled a good chunk of their product line for insecure configurations. Some feel that this is insufficient, however. For instance, Bruce Schneier has recently been calling for government policy to regulate IOT device security.

In our honeypot, all login attempts came through telnet, which is what Mirai uses to hack devices — it isn't built to hack HTTP, UDP, etc., although later variants began to take advantage of vendor-specific bugs when they were widespread, such as the TR-069 exploits that were merged into some of these botnets, abusing a broadband forum protocol to gain entry to the devices.

Here's some of what we found:

- 62 username and password combinations in the Mirai source code are used to attempt to infect devices. This isn't surprising, as this list has been used successfully by many embedded device botnets since then.
- Given the size of the botnet, it would take less than 6 minutes for the Mirai botnet to scan the entire IPv4 space for hackable IoT devices. This means that a new exploit could be deployed rapidly, or the Mirai botnet could be used as a reconnaissance platform by third parties, not just denial of service (DoS) attacks.
- Attacks to the IoT devices we tracked were constant, every second, for 7 days. This ceaseless activity is a testament to the prevalence of this problem. Despite a massive fracturing of the Mirai botnet into various competing botnets, compromised IoT devices number so many that this has become constant internet background noise.
- The Telnet login requests from a single infected host come in at an average of every 3.7 seconds — faster than you can log in and patch a device, download a patch, or update the password. wFor a typical end user, this means that the window to update a vulnerable IOT device is virtually

non-existent, reminding me of the worst days with a vulnerable Windows box on the open internet: before you could patch you were already compromised.

IoT devices expose a lot, reflecting their engineers' focus on quickly getting to market and enabling people to get online easily, rather than building with security best practices in mind — undoing a significant amount of security work from the past 15 years. This work included convincing platform vendors such as Microsoft, Apple, RedHat, and others to take security seriously and to make security defaults a reality. Examples include Windows XP SP2's default-on local firewall and exploit mitigation technologies, RedHat's configuration changes to mail and web servers in default installations, and Apple's inclusion of address randomization to defeat various attacks. These efforts took years of effort by a whole cast of characters, but clearly needs a new audience in IoT vendors, some of whom are coming to internet-enabled devices and security risks for the first time.

The large size of the Mirai botnet makes it an internet-scale issue — the fact that they can scan the entire web in under six minutes makes it a concern for the entire internet community as noted above. These botnets enable widespread secondary attacks by providing stepping stones and overlay networks for more sophisticated attacks, for example. But not every IoT device is a ticking time bomb. Many vendors, including Cisco, Philips, and Apple, have strengthened their default, out-of-the box experience to provide ease of use married to security. For the average consumer it's relatively easy to defend against these sorts of issues for end users with IoT devices by employing basic hygiene on a home network behind a firewall.