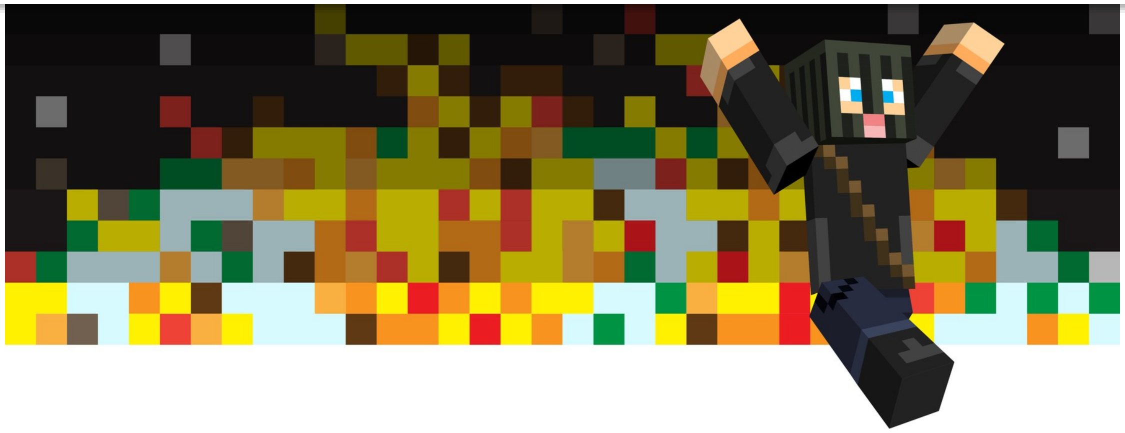

BEN BOURS/WIRED

GARRETT M. GRAFF SECURITY 12.13.17 03:55 PM

HOW A DORM ROOM *MINECRAFT* SCAM BROUGHT DOWN THE INTERNET

THE MOST DRAMATIC cybersecurity story of 2016 came to a quiet conclusion Friday in an Anchorage courtroom, as three young American computer savants pleaded guilty to masterminding an unprecedented botnet—powered by unsecured internet-of-things devices like security cameras and wireless routers—that unleashed sweeping attacks on key internet services around the globe last fall. What drove them wasn't anarchist politics or shadowy ties to a nation-state. It was *Minecraft*.

It was a hard story to miss last year: In France last September, the telecom provider OVH was hit by a distributed denial-of-service (DDoS) attack a hundred times larger than most of its kind. Then, on a Friday afternoon in October 2016, the internet slowed or stopped for nearly the entire eastern United States, as the tech company Dyn, a key part of the internet's backbone, came under a crippling assault.

As the 2016 US presidential election drew near, fears began to mount that the so-called Mirai botnet might be the work of a nation-state practicing for an attack that would cripple the country as voters went to the polls. The truth, as made clear in that Alaskan courtroom Friday—and unsealed by the Justice Department on Wednesday—was even stranger: The brains behind Mirai were a 21-year-old

Rutgers college student from suburban New Jersey and his two college-age friends from outside Pittsburgh and New Orleans. All three—Paras Jha, Josiah White, and Dalton Norman, respectively—admitted their role in creating and launching Mirai into the world.

Originally, prosecutors say, the defendants hadn't intended to bring down the internet—they had been trying to gain an advantage in the computer game *Minecraft*.

“They didn't realize the power they were unleashing,” says FBI supervisory special agent Bill Walton. "This was the Manhattan Project."

Unraveling the whodunit of one of the internet's biggest security scares of 2016 led the FBI through a strange journey into the underground DDoS market, the modern incarnation of an old neighborhood mafia-protection racket, where the very guys offering to help today might actually be the ones who attacked you yesterday.

Then, once the FBI unraveled the case, they discovered that the perpetrators had already moved onto a new scheme—inventing a business model for online crime no one had ever seen before, and pointing to a new, looming botnet threat on the horizon.

THE FIRST RUMORS that something big was beginning to unfold online came in August 2016. At the time, FBI special agent Elliott Peterson was part of a multinational investigative team trying to zero in on two teens running a DDoS attack-for-hire service known as vDOS. It was a major investigation—or at least it seemed so at the time.

VDOS was an advanced botnet: a network of malware-infected, zombie devices that its masters could commandeer to execute DDoS attacks at will. And the teens were using it to run a lucrative version of a then-common scheme in the online gaming world—a so-called booter service, geared toward helping individual gamers attack an opponent while fighting head-to-head, knocking them offline to defeat them. Its tens of thousands of customers could pay small amounts, like $5 to $50, to rent small-scale denial-of-service attacks via an easy-to-use web interface.

Yet as that case proceeded, the investigators and the small community of security engineers who protect against denial-of-service attacks began to hear rumblings about a new botnet, one that eventually made vDOS seem small.

As Peterson and industry colleagues at companies like Cloudflare, Akamai, Flashpoint, Google, and Palo Alto Networks began to study the new malware, they realized they were looking at something entirely different from what they'd battled in the past. Whereas the vDOS botnet they'd been chasing was a variant of an older IoT zombie army—a 2014 botnet known as Qbot—this new botnet appeared to have been written from the ground up.

And it was good.

“From the initial attacks, we realized this was something very different from your normal DDoS,” says Doug Klein, Peterson's partner on the case.

The new malware scanned the internet for dozens of different IoT devices that still used the manufacturers' default security setting. Since most users rarely change default usernames or passwords, it quickly grew into a powerful assembly of weaponized electronics, almost all of which had been hijacked without their owners' knowledge.

“The security industry was really not aware of this threat until about mid-September. Everyone was playing catch-up,” Peterson says. “It's really powerful—they figured out how to stitch together multiple exploits with multiple processors. They crossed the artificial threshold of 100,000 bots that others had really struggled with.”

It didn't take long for the incident to go from vague rumblings to global red alert.

Mirai shocked the internet—and its own creators, according to the FBI—with its power as it grew. Researchers later determined that it infected nearly 65,000 devices in its first 20 hours, doubling in size every 76 minutes, and ultimately built a sustained strength of between 200,000 and 300,000 infections.

“These kids are super smart, but they didn’t do anything high level—they just had a good idea,” the FBI’s Walton says. “It’s the most successful IoT botnet we’ve ever seen—and a sign that computer crime isn’t just about desktops anymore.”

Targeting cheap electronics with poor security, Mirai amassed much of its strength by infecting devices in Southeast Asia and South America; the four main countries with Mirai infections were Brazil, Colombia, Vietnam, and China, according to researchers. As a team of security professionals later concluded, dryly, “Some of the world’s top manufacturers of consumer electronics lacked sufficient security practices to mitigate threats like Mirai.”

At its peak, the self-replicating computer worm had enslaved some 600,000 devices around the world—which, combined with today’s high-speed broadband connections, allowed it to harness an unprecedented flood of network-clogging traffic against target websites. It proved particularly tough for companies to fight against and remediate, too, as the botnet used a variety of different nefarious traffic to overwhelm its target, attacking both servers and applications that ran on the servers, as well as even older techniques almost forgotten in modern DDoS attacks.

On September 19, 2016, the botnet was used to launch crushing DDoS attacks against French hosting provider OVH. Like any large hosting company, OVH regularly saw small-scale DDoS attacks—it noted later that it normally faces 1,200 a day—but the Mirai attack was unlike anything anyone on the internet had ever seen, the first thermonuclear bomb of the DDoS world, topping out at 1.1 terabits per second as more than 145,000 infected devices bombarded OVH with unwanted traffic. The company’s CTO tweeted about the attacks afterward to warn others of the looming threat.

Until then, a large DDoS attack was often considered to be 10 to 20 gigibits per second; vDOS had been overwhelming targets with attacks in the range of 50 Gbps. A follow-on Mirai attack against OVH hit around 901 Gbps.

Mirai was particularly deadly, according to court documents, because it was able to target an entire range of IP addresses—not just one particular server or website—enabling it to crush a company’s entire network.

“Mirai was an insane amount of firepower,” Peterson says. And no one had any idea yet who its creators were, or what they were trying to accomplish.

Normally, companies fight a DDoS attack by filtering incoming web traffic or increasing their bandwidth, but at the scale Mirai operated, nearly all traditional DDoS mitigation techniques collapsed, in part because the tidal wave of nefarious traffic would crash so many sites and servers en route to its main target. “DDOS at a certain scale poses an existential threat to the internet,” Peterson says. “Mirai was the first botnet I’ve seen that hit that existential level.”

Through September, the inventors of Mirai tweaked their code—researchers were later able to assemble 24 iterations of the malware that appeared to be primarily the work of the three main defendants in the case—as the malware grew more sophisticated and virulent. They actively battled the hackers behind vDOS, fighting for control of IoT devices, and instituting kill procedures to wipe competing infections off compromised devices—natural selection playing out at internet speed. According to court documents, they also filed fraudulent abuse complaints with internet hosts associated with vDOS.

“They were trying to outmuscle each other. Mirai outperforms all of them,” Peterson says. “This crime was evolving through competition.”

Whoever was behind Mirai even bragged about it on hacker bulletin boards; someone using the moniker Anna-senpai claimed to be the creator, and someone named ChickenMelon talked it up as well, hinting that their competitors might be using malware from the NSA.

Days after OVH, Mirai struck again, this time against a high-profile technology target: security reporter Brian Krebs. The botnet blasted Krebs’ website, Krebs on Security, knocking it offline for more than four days with an attack that peaked at 623 Gbps. The assault was so effective—and sustained—that Krebs’ longtime DDoS mitigation service, Akamai, one of the largest bandwidth providers on the internet, announced it was dropping Krebs’ site because it couldn’t bear the cost of defending against such a massive barrage. The Krebs attack, Akamai said, was twice the size of the largest attack it had ever seen before.

Whereas the OVH attack overseas had been an online curiosity, the Krebs attack quickly pushed the Mirai botnet to the FBI's front burner, especially as it seemed likely that it was retribution for an article Krebs had published just days earlier about another DDoS-mitigation firm that appeared to be engaged in nefarious practices, hijacking web addresses that it believed were being controlled by the vDOS team.

"This is strange development—a journalist being silenced because someone has figured out a tool powerful enough to silence him," Peterson says. "That was worrisome."

The IoT attacks began to make big headlines online and off; media reports and security experts speculated that Mirai might have the fingerprints of a looming attack on the internet's core infrastructure.

"Someone has been probing the defenses of the companies that run critical pieces of the internet. These probes take the form of precisely calibrated attacks designed to determine exactly how well these companies can defend themselves, and what would be required to take them down," wrote security expert Bruce Schneier in September 2016. "We don't know who is doing this, but it feels like a large nation-state. China or Russia would be my first guesses."

Behind the scenes, the FBI and industry researchers raced to unravel Mirai and zero in on its perpetrators. Network companies like Akamai created online honeypots, mimicking hackable devices, to observe how infected "zombie" devices communicated with Mirai's command-and-control servers. As they began to study the attacks, they noticed that many of the Mirai assaults had appeared to target gaming servers. Peterson recalls asking, "Why are these *Minecraft* servers getting hit so often?"

THE QUESTION WOULD lead the investigation deep into one of the internet's strangest worlds, a $27 game with an online population of registered users—122 million—larger than the entire country of Egypt. Industry analysts report 55 million people play *Minecraft* each month, with as many as a million online at any given time.

The game, a three-dimensional sandbox with no particular goals, allows players to construct entire worlds by "mining" and building with cartoonish pixelated blocks.

Its comparatively basic visual appeal—it has more in common with the first-generation videogames of the 1970s and 1980s than it does the polygon-intense lushness of *Halo* or *Assassin's Creed*—belies a depth of imaginative exploration and experimentation that has propelled it to be the second-best-selling videogame ever, behind only *Tetris*. The game and its virtual worlds were acquired by Microsoft in 2014 as part of a deal worth nearly $2.5 billion, and it has spawned numerous fan sites, explanatory wikis, and YouTube tutorials—even a real-life collection of *Minecraft*-themed Lego bricks.

It has also become a lucrative platform for *Minecraft* entrepreneurs: Inside the game, individual hosted-servers allow users to link together in multiplayer mode, and as the game has grown, hosting those servers has turned into big business—players pay real money both to rent "space" in *Minecraft* as well as purchase in-game tools. Unlike many massive multiplayer games where every player experiences the game similarly, these individual servers are integral to the *Minecraft* experience, as each host can set different rules and install different plug-ins to subtly shape and personalize the user experience; a particular server, for instance, might not allow players to destroy one another's creations.

As Peterson and Klein explored the *Minecraft* economy, interviewing server hosts and reviewing financial records, they came to realize how amazingly financially successful a well-run, popular *Minecraft* server could be. "I went into my boss's office and said, 'Am I crazy? It looks like people are making a ton of money,'" he recalls. "These people at the peak of summer were making $100,000 a month."

The huge income from successful servers had also spawned a mini cottage industry of launching DDoS attacks on competitors' servers, in an attempt to woo away players frustrated at a slow connection. (There are even YouTube tutorials specifically aimed at teaching *Minecraft* DDoS, and free DDoS tools available at Github.) Similarly, *Minecraft* DDoS-mitigation services have sprung up as a way to protect a host's server investment.

The digital arms race in DDoS is inexorably linked to *Minecraft*, Klein says.

“We see so many attacks on *Minecraft*. I’d be more surprised sometimes if I didn’t see a *Minecraft* connection in a DDoS case,” he says. “You look at the servers—those guys are making huge money, so it’s in my benefit to knock your server offline and steal your customers. The vast majority of these *Minecraft* servers are being run by kids—you don’t necessarily have the astute business judgment in the quote-unquote ‘executives’ running these servers.”

As it turned out, French internet host OVH was well-known for offering a service called VAC, one of the industry’s top *Minecraft* DDoS-mitigation tools. The Mirai authors attacked it not as part of some grand nation-state plot but rather to undermine the protection it offered key *Minecraft* servers. “For a while, OVH was too much, but then they figured out how to even beat OVH,” Peterson says.

This was something new. Whereas gamers had become familiar with one-off DDoS attacks by booter services, the idea of DDoS as a business model for server hosts was startling. “This was a calculated business decision to shut down a competitor,” Peterson says.

“They just got greedy—they thought, ‘If we can knock off our competitors, we can corner the market on both servers and mitigation,’” Walton says.

In fact, according to court documents, the primary driver behind the original creation of Mirai was creating "a weapon capable of initiating powerful denial-of-service attacks against business competitors and others against whom White and his coconspirators held grudges.”

Once investigators knew what to look for, they found *Minecraft* links all over Mirai: In an less-noticed attack just after the OVH incident, the botnet had targeted ProxyPipe.com, a company in San Francisco that specializes in protecting *Minecraft* servers from DDoS attacks.

“Mirai was originally developed to help them corner the *Minecraft* market, but then they realized what a powerful tool they built,” Walton says. “Then it just became a challenge for them to make it as large as possible.”

On September 30, 2016, as public attention piqued following the Krebs attack, the maker of Mirai posted the malware's source code to the website Hack Forum, in an attempt to deflect possible suspicions if he was caught. The release also included the default credentials for 46 IoT devices central to its growth. (Malware authors will sometimes release their code online to muddy investigators' trail, ensuring that even if they're found to possess the source code, authorities can't necessarily identify them as the original author.)

That release opened the tool for use by a wide audience, as competing DDoS groups adopted it and created their own botnets. All told, over five months from September 2016 through February 2017, variations of Mirai were responsible for upwards of 15,194 DDoS attacks, according to an after-action report published in August.

As the attacks spread, the FBI worked with private-industry researchers to develop tools that allowed them to watch DDoS attacks as they unfolded, and track where the hijacked traffic was being directed—the online equivalent of the Shotspotter system that urban police departments use to detect the location of gunshots and dispatch themselves toward trouble. With the new tools, the FBI and private industry were able to see a looming DDoS attack unfold and help mitigate it in real time. "We really depended on the generosity of the private sector," Peterson says.

The decision to open source Mirai also led to its most high-profile attack. The FBI says Jha, White, and Dalton were not responsible for last October's DDoS of the domain name server Dyn, a critical piece of internet infrastructure that helps web browsers translate written addresses, like Wired.com, into specific numbered IP addresses online. (The FBI declined to comment on the Dyn investigation; there have been no arrests publicly reported in that case.)

The Dyn attack paralyzed millions of computer users, slowing or stopping internet connections up and down the East Coast and interrupting service across North America and parts of Europe to major sites like Amazon, Netflix, Paypal, and Reddit. Dyn later announced that it might never be able to calculate the full weight of the assault it faced: “There have been some reports of a magnitude in the 1.2 Tbps range; at this time we are unable to verify that claim.”

Justin Paine, the director of trust and safety for Cloudflare, one of the industry’s leading DDoS mitigation companies, says that the Dyn attack by Mirai immediately got the attention of engineers across the internet. “When Mirai really came on the scene, the people who run the internet behind the scenes, we all came together,” he says “We all realized that this isn’t something that just affects my company or my network—this could put the entire internet at risk. Dyn affected the entire internet.”

“The concept of unsecured devices to be repurposed by bad guys to do bad things, that’s always been there,” says Paine, “but the sheer scale of insecure modems, DVRs, and webcams in combination with how horribly insecure they were as device really did a present a different kind of challenge.”

The tech industry began intensively sharing information, both to help mitigate ongoing attacks as well as working to backtrack and to identify infected devices to begin remediation efforts. Network engineers from multiple companies convened an always-running Slack channel to compare notes on Mirai. As Paine says, “It was real-time, we were using Slack, sharing, ‘Hey, I’m on this network seeing this, what are you seeing?’”

The power of the botnet was made even more clear as the fall unfolded and Mirai attacks targeted the African country of Liberia, effectively cutting off the entire country from the internet.

Many of these follow-on attacks also appeared to have a gaming angle: A Brazilian internet service provider saw its *Minecraft* servers targeted; the Dyn attacks also appeared to target gaming servers, as well as servers hosting Microsoft Xbox Live and Playstation servers and those associated with gaming hosting company called Nuclear Fallout Enterprises. “The attacker was likely targeting gaming infrastructure that incidentally disrupted service to Dyn’s broader customer base,” researchers later declared.

“Dyn got everyone’s attention,” says Peterson, especially as it represented a new evolution—and a new unknown player fiddling with Anna-senpai’s code. “It was the first truly effective post-Mirai variant.”

The Dyn attack catapulted Mirai to the front pages—and brought immense national pressure down on the agents chasing the case. Coming just weeks before the presidential election—one in which US intelligence officials had already warned about attempts by Russia to interfere—the Dyn and Mirai attacks led officials to worry that Mirai could be harnessed to affect voting and media coverage of the election. The FBI team scrambled for a week afterward with private-industry partners to secure critical online infrastructure and ensure that a botnet DDoS couldn’t disrupt Election Day.

The plague unleashed by Mirai’s source code continued to unfold across the internet last winter. In November, the German company Deutsche Telekom saw more than 900,000 routers knocked offline when a bug-filled variant of Mirai accidentally targeted them. (German police eventually arrested a 29-year-old British hacker in that incident.) Yet the various competing Mirai botnets undercut their own effectiveness, as an increasing number of botnets fought over the same number of devices, eventually leading to smaller and smaller—and therefore less effective and troubling—DDoS attacks.

WHAT ANNA-SENPAI DIDN’T realize when he dumped the source code was that the FBI had already worked through enough digital hoops to finger Jha as a likely suspect, and had done so from an unlikely perch: Anchorage, Alaska.

That one of the big internet stories of 2016 would end up in an Anchorage courtroom last Friday—guided by assistant US attorney Adam Alexander to a guilty plea barely a year after the original offense, a remarkably rapid pace for cybercrimes—was a signal moment itself, marking an important maturation in the FBI's national approach to cybercrimes.

Until recently, nearly all of the FBI's major cybercrime prosecutions came out of just a handful of offices like Washington, New York, Pittsburgh, and Atlanta. Now, though, an increasing number of offices are gaining the sophistication and understanding to piece together time-consuming and technically complex internet cases.

Peterson is a veteran of the FBI's most famous cyber team, a pioneering squad in Pittsburgh that has put together groundbreaking cases, like that against five Chinese PLA hackers. On that squad, Peterson—an energetic, hard-charging, college computer science major and Marine Corps adjutant who deployed twice to Iraq before joining the bureau, and now serves on the FBI Alaska SWAT team—helped lead the investigation into the GameOver Zeus botnet that targeted Russian hacker Evgeny Bogachev, who remains at large with a $3 million reward for his capture.

Often, FBI agents end up being pulled away from their core specialties as their career advances; in the years after 9/11, one of the bureau's few dozen Arabic-speaking agents ended up running a squad investigating white supremacists. But Peterson stayed focused on cyber cases even as he transferred nearly two years ago back to his home state of Alaska, where he joined the FBI's smallest cyber squad—just four agents, overseen by Walton, a longtime Russian counterintelligence agent, and partnering with Klein, a former UNIX systems administrator.

The tiny team, though, has come to take on an outsized role in the country's cybersecurity battles, specializing in DDoS attacks and botnets. Earlier this year, the Anchorage squad was instrumental in the take-down of the long-running Kelihos botnet, run by Peter Yuryevich Levashov, aka "Peter of the North," a hacker arrested in Spain in April.

In part, says Marlin Ritzman, the special-agent-in-charge of the FBI's Anchorage Field Office, that's because Alaska's geography makes denial-of-service attacks particularly personal.

"Alaska's uniquely positioned with our internet services—a lot of rural communities depend on the internet to reach the outside world," Ritzman says. "A denial-of-service attack could shut down communications to entire communities up here, it's not just one business or another. It's important for us to attack that threat."

Putting together the Mirai case was slow going for the four-agent Anchorage squad, even while they worked closely with dozens of companies and private sector researchers to piece together a global portrait of an unprecedented threat.

Before they could solve an international case, the FBI squad first—given the decentralized way that federal courts and the Justice Department work—had to prove that Mirai existed in their particular jurisdiction, Alaska.

To establish the grounds for a criminal case, the squad painstakingly located infected IoT devices with IP addresses across Alaska, then issued subpoenas to the state's main telecom company, GCI, to attach a name and physical location. Agents then criss-crossed the state to interview the owners of the devices and establish that they hadn't given permission for their IoT purchases to be hijacked by the Mirai malware.

While some infected devices were close by in Anchorage, others were further afield; given Alaska's remoteness, collecting some devices required plane trips to rural communities. At one rural public utility that also provided internet services, agents found an enthusiastic network engineer who helped track down compromised devices.

After seizing the infected devices and transporting them to the FBI field office—a low-slung building just a few blocks from the water in Alaska's most populous city—agents, counterintuitively, then had to plug them back in. Since Mirai malware exists only in flash memory, it was deleted every time the device was powered off or restarted. The agents had to wait for the device to be reinfected by Mirai; luckily, the botnet was so infectious and spread so rapidly that it didn't take long for the devices to be reinfected.

From there, the team worked to trace the botnet’s connections back to the main Mirai control server. Then, armed with court orders, they were able to track down associated email addresses and cell phone numbers used for those accounts, establishing and linking names to the boxes.

“It was a lot of six degrees of Kevin Bacon,” Walton explains. “We just kept stepping down that chain.”

At one point, the case bogged down because the Mirai authors had established in France a so-called popped box, a compromised device that they used as an exit VPN node from the internet, thereby cloaking the actual location and physical computers used by Mirai’s creators.

As it turned out, they’d hijacked a computer that belonged to a French kid interested in Japanese anime. Given that Mirai had, according to a leaked chat, been named after a 2011 anime series, Mirai Nikki, and that the author’s pseudonym was Anna-Senpai, the French boy was an immediate suspect.

“The profile lined up with someone we’d expect to be involved in the development of Mirai,” Walton says; throughout the case, given the OVH connection, the FBI worked closely with French authorities, who were present as some of the search warrants were conducted.

“The actors were very sophisticated in their online security,” Peterson says. “I’ve run against some really hard guys, and these guys were as good or better than some of the Eastern Europe teams I’ve gone against.”

Adding to the complexity, DDoS itself is a notoriously difficult crime to prove—even simply proving the crime ever happened can be extraordinarily challenging after the fact. “DDoS can happen in a vacuum, unless a company captures logs in the right way,” Peterson says. Klein, a former UNIX administrator who grew up playing with Linux, spent weeks piecing together evidence and reassembling data to show how the DDoS attacks unfolded.

On the compromised devices, they had to carefully reconstruct the network traffic data, and study how the Mirai code launched so-called “packets” against its targets —a little-understood forensic process, known as analyzing PCAP (packet capture) data. Think of it as the digital equivalent of testing for fingerprints or gunshot residue. “It was the most complex DDoS software I’ve run across,” Klein says.

The FBI zeroed in on the suspects by the end of the year: Photos of the three hung for months on the wall in the Anchorage field office, where agents dubbed them the "Cub Scout Pack,” a nod to their youthfulness. (Another older female suspect in an unrelated case, whose photo also hung on the board, was nicknamed the "Den Mother.”)

Security journalist Brian Krebs, an early Mirai victim, publicly fingered Jha and White in January 2017. Jha’s family initially denied his involvement, but on Friday he, White, and Norman all pleaded guilty to conspiracy to violate the Computer Fraud and Abuse Act, the government’s main criminal charge for cybercrime. The pleas were unsealed Wednesday, and announced by the Justice Department’s computer crimes unit in Washington, DC.

Jha was also accused of—and pleaded guilty to—a bizarre set of DDoS attacks that had disrupted the computer networks on the Rutgers campus for two years. Beginning in the first year Jha was a student there, Rutgers began to suffer from what would ultimately be a dozen DDoS attacks that disrupted networks, all timed to midterms. At the time, an unnamed individual online pushed the university to purchase better DDoS mitigation services—which, as it turns out, was exactly the business Jha himself was trying to build.

In a Trenton courtroom Wednesday, Jha—wearing a conservative suit and the dark-rimmed glasses familiar from his old LinkedIn portrait—told the court that he aimed attacks against at his own campus when they would be most disruptive—specifically during midterms, finals, and when students were trying to register for class.

“In fact, you timed your attacks because you wanted to overload the central authentication server when it would be the most devastating to Rutgers, right?” the federal prosecutor queried.

“Yes,” Jha said.

Indeed, that the three computer savants ended up building a better DDoS mousetrap isn’t necessarily surprising; it was an area of intense intellectual interest for them. According to their online profiles, Jha and White had actually been working together to build a DDoS-mitigation firm; the month before Mirai appeared, Jha’s email signature described him as “President, ProTraf Solutions, LLC, Enterprise DDoS Mitigation.”

As part of building Mirai, each member of the group had his own role, according to the court documents. Jha wrote much of the original code and served as the main online point of contact on hacking forums, using the Anna-senpai moniker.

White, who used the online monikers Lightspeed and thegenius, ran much of the botnet infrastructure, designing the powerful internet scanner that helped identify potential devices to infect. The scanner’s speed and effectiveness was a key driver behind Mirai’s ability to outcompete other botnets like vDOS last fall; at the peak of Mirai, an experiment by *The Atlantic* found that a fake IoT device the publication created online was compromised within an hour.

According to court documents, Dalton Norman—whose role in the Mirai botnet was unknown until the plea agreements were unsealed—worked to identify the so-called zero-day exploits that made Mirai so powerful. According to court documents, he identified and implemented four such vulnerabilities unknown to device manufacturers as part of Mirai’s operating code, and then, as Mirai grew, he worked to adapt the code to run a vastly more powerful network than they’d ever imagined.

Jha came to his interest in technology early; according to his now deleted LinkedIn page, he described himself as “highly self-motivated” and explained that he began to teach himself programming in seventh grade. His interest in science and technology ranged widely: The following year, he won second prize in the eighth-grade science fair at Park Middle School in Fanwood, New Jersey, for his engineering project studying the impact of earthquakes on bridges. By 2016, he listed himself as proficient in “C#, Java, Golang, C, C++, PHP, x86 ASM, not to mention web ‘browser languages’ such as Javascript and HTML/CSS.” (One early clue for Krebs that Jha was likely involved in Mirai was that the person calling themself Anna-Senpai had listed their skills by saying, “I’m very familiar with programming in a variety of languages, including ASM, C, Go, Java, C#, and PHP.)

This is not the first time that teens and college students have exposed key weaknesses in the internet: The first major computer worm was unleashed in November 1988 by Robert Morris, then a student at Cornell, and the first major intrusion into the Pentagon’s computer networks—a case known as Solar Sunrise—came a decade later, in 1998; it was the work of two California teens in concert with an Israeli contemporary. DDoS itself emerged in 2000, unleashed by a Quebec teen, Michael Calce, who went online by the moniker Mafiaboy. On February 7, 2000, Calce turned a network of zombie computers he’d assembled from university networks against Yahoo, then the web’s largest search engine. By mid-morning it had all but crippled the tech giant, slowing the site to a crawl, and in the days following, Calce targeted other top websites like Amazon, CNN, eBay, and ZDNet.

On a conference call announcing the guilty pleas Wednesday, Justice Department Acting Deputy Assistant Attorney General Richard Downing said that the Mirai case underscored the perils of young computer users who lose their way online—and said that the Justice Department planned to expand its youth outreach efforts.

"I've certainly been made to feel very old and unable to keep up," prosecutor Adam Alexander joked Wednesday.

What really surprised investigators, though, was that once they had Jha, White, and Norman in their sights, they discovered that the creators of Mirai had already found a new use for their powerful botnet: They’d given up DDoS attacks for something lower-profile—but also lucrative.

They were using their botnet to run an elaborate click-fraud scheme—directing about 100,000 compromised IoT devices, mostly home routers and modems, to visit advertising links en masse, making it appear that they were regular computer users. They were making thousands of dollars a month defrauding US and European advertisers, entirely off the radar, with no one the wiser. It was, as far as investigators could tell, a groundbreaking business model for an IoT botnet.

As Peterson says, “Here was a whole new crime that industry was blind to. We all missed it.”

Even as the case in Alaska and New Jersey wraps up—the three defendants will face sentencing later on—the Mirai plague that Jha, White, and Dalton unleashed continues online. “This particular saga is over, but Mirai still lives,” Cloudflare’s Paine says. “There’s a significant ongoing risk that’s continued, as the open source code has been repurposed by new actors. All these new updated versions are still out there.”

Two weeks ago, at the beginning of December, a new IoT botnet appeared online using aspects of Mirai’s code.

Known as Satori, the botnet infected a quarter million devices in its first 12 hours.

Garrett M. Graff (@vermontgmg) is a contributing editor for WIRED. *He can be reached at garrett.graff@gmail.com.*

This article has been updated to reflect that Mirai struck a hosting company called Nuclear Fallout Enterprises, not a game called Nuclear Fallout.