Krebs on Security

In-depth security news and investigation

About the Author
Blog Advertising

25
Sep 16

The Democratization of Censorship

John Gilmore, an American entrepreneur and civil libertarian, once famously quipped that "the Internet interprets censorship as damage and routes around it." This notion undoubtedly rings true for those who see national governments as the principal threats to free speech.

However, events of the past week have convinced me that one of the fastest-growing censorship threats on the Internet today comes not from nation-states, but from super-empowered individuals who have been quietly building extremely potent cyber weapons with transnational reach.

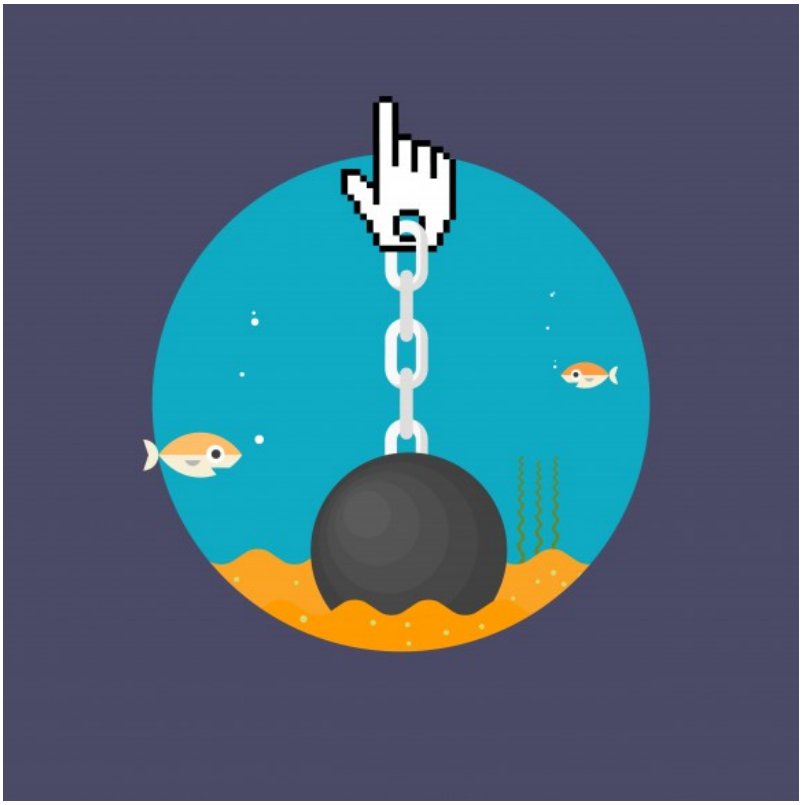

More than 20 years after Gilmore first coined that turn of phrase, his most notable quotable has effectively been inverted — "Censorship can in fact route around the Internet." The Internet can't route around censorship when the censorship is all-pervasive and armed with, for all practical purposes, near-infinite reach and capacity. I call this rather unwelcome and hostile development the "The Democratization of Censorship."

Allow me to explain how I arrived at this unsettling conclusion. As many of you know, my site was taken offline for the better part of this week. The outage came in the wake of a historically large distributed denial-of-service (DDoS) attack which hurled so much junk traffic at Krebsonsecurity.com that my DDoS protection provider **Akamai** chose to unmoor my site from its protective harbor.

Let me be clear: I do not fault Akamai for their decision. I was a pro bono customer from the start, and Akamai and its sister company Prolexic have stood by me through countless attacks over the past four years. It just so happened that this last siege was nearly twice the size of the next-largest attack they had ever seen before. Once it became evident that the assault was beginning to cause problems for the company's paying customers, they explained that the choice to let my site go was a business decision, pure and simple.

Nevertheless, Akamai rather abruptly informed me I had until 6 p.m. that very same day — roughly two hours later — to make arrangements for migrating off their network. My main concern at the time was making sure my hosting provider wasn't going to bear the brunt of the attack when the

shields fell. To ensure that absolutely would not happen, I asked Akamai to redirect my site to 127.0.0.1 — effectively relegating all traffic destined for KrebsOnSecurity.com into a giant black hole.

Today, I am happy to report that the site is back up — this time under Project Shield, a free program run by **Google** to help protect journalists from online censorship. And make no mistake, DDoS attacks — particularly those the size of the assault that hit my site this week — are uniquely effective weapons for stomping on free speech, for reasons I'll explore in this post.

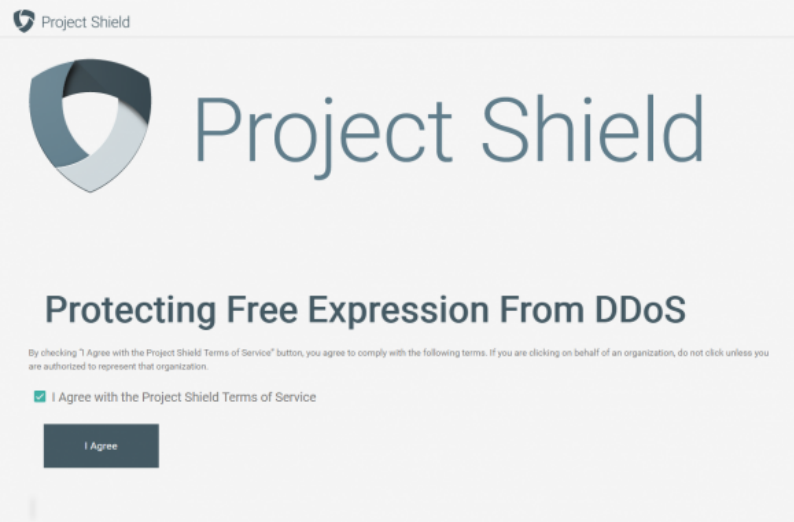

Google's Project Shield is now protecting KrebsOnSecurity.com

Why do I speak of DDoS attacks as a form of censorship? Quite simply because the economics of mitigating large-scale DDoS attacks do not bode well for protecting the individual user, to say nothing of independent journalists.

In an interview with *The Boston Globe*, Akamai executives said the attack — if sustained — likely would have cost the company millions of dollars. In the hours and days following my site going offline, I spoke with multiple DDoS mitigation firms. One offered to host KrebsOnSecurity for two weeks at no charge, but after that they said the same kind of protection I had under Akamai would cost between $150,000 and $200,000 per year.

Ask yourself how many independent journalists could possibly afford that kind of protection money? A number of other providers offered to help, but it was clear that they did not have the muscle to be able to withstand such massive attacks.

I've been toying with the idea of forming a 501(c)3 non-profit organization — 'The Center for the Defense of Internet Journalism', if you will — to assist Internet journalists with obtaining the kind of protection they may need when they become the targets of attacks like the one that hit my site. Maybe a Kickstarter campaign, along with donations from well-known charitable organizations, could get the ball rolling. It's food for thought.

CALIBRATING THE CANNONS

Earlier this month, noted cryptologist and security blogger **Bruce Schneier** penned an unusually alarmist column titled, "Someone Is Learning How to Take Down the Internet." Citing unnamed sources, Schneier warned that there was strong evidence indicating that nation-state actors were actively and aggressively probing the Internet for weak spots that could allow them to bring the entire Web to a virtual standstill.

"Someone is extensively testing the core defensive capabilities of the companies that provide critical Internet services," Schneier wrote. "Who would do this? It doesn't seem like something an activist, criminal, or researcher would do. Profiling core infrastructure is common practice in espionage and intelligence gathering. It's not normal for companies to do that."

Schneier continued:

"Furthermore, the size and scale of these probes — and especially their persistence — points to state actors. It feels like a nation's military cyber command trying to calibrate its weaponry in the case of cyberwar. It reminds me of the US's Cold War program of flying high-altitude planes over the Soviet Union to force their air-defense systems to turn on, to map their capabilities."

Whether Schneier's sources were accurate in their assessment of the actors referenced in his blog post is unknown. But as my friend and mentor **Roland Dobbins** at Arbor Networks eloquently put it, "When it comes to DDoS attacks, nation-states are just another player."

"Today's reality is that DDoS attacks have become the Great Equalizer between private actors & nation-states," Dobbins quipped.

"Today's reality is that DDoS attacks have become the Great Equalizer between private actors & nation-states," Dobbins quipped.

UM…YOUR RERUNS OF 'SEINFELD' JUST ATTACKED ME

What exactly was it that generated the record-smashing DDoS of 620 Gbps against my site this week? Was it a space-based weapon of mass disruption built and tested by a rogue nation-state, or an arch villain like SPECTRE from the James Bond series of novels and films? If only the enemy here was that black-and-white.

No, as I reported in the last blog post before my site was unplugged, the enemy in this case was far less sexy. There is every indication that this attack was launched with the help of a botnet that has enslaved a large number of hacked so-called "Internet of Things," (IoT) devices — mainly routers, IP cameras and digital video recorders (DVRs) that are exposed to the Internet and protected with weak or hard-coded passwords. Most of these devices are available for sale on retail store shelves for less than $100, or — in the case of routers — are shipped by ISPs to their customers.

Some readers on Twitter have asked why the attackers would have "burned" so many compromised systems with such an overwhelming force against my little site. After all, they reasoned, the attackers showed their hand in this assault, exposing the Internet addresses of a huge number of compromised devices that might otherwise be used for actual money-making cybercriminal activities, such as hosting malware or relaying spam. Surely, network providers would take that list of hacked devices and begin blocking them from launching attacks going forward, the thinking goes.

As KrebsOnSecurity reader **Rob Wright** commented on Twitter, "the DDoS attack on @briankrebs feels like testing the Death Star on the Millennium Falcon instead of Alderaan." I replied that this maybe wasn't the most apt analogy. The reality is that there are currently millions — if not tens of millions — of insecure or poorly secured IoT devices that are ripe for being enlisted in these attacks at any given time. And we're adding millions more each year.

I suggested to Mr. Wright perhaps a better comparison was that ne'er-do-wells now have a virtually limitless supply of Stormtrooper clones that can be conscripted into an attack at a moment's notice.

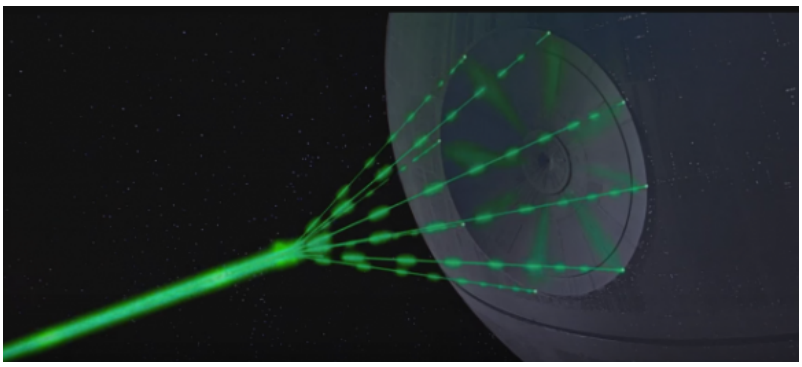

A scene from the 1977 movie Star Wars, in which the Death Star tests its firepower by blowing up a planet.

SHAMING THE SPOOFERS

The problem of DDoS conscripts goes well beyond the millions of IoT devices that are shipped insecure by default: Countless hosting providers and ISPs do nothing to prevent devices on their networks from being used by miscreants to "spoof" the source of DDoS attacks.

As I noted in a November 2015 story, The Lingering Mess from Default Insecurity, one basic step that many ISPs can but are not taking to blunt these attacks involves a network security standard that was developed and released more than a dozen years ago. Known as BCP38, its use prevents insecure resources on an ISPs network (hacked servers, computers, routers, DVRs, etc.) from being leveraged in such powerful denial-of-service attacks.

Using a technique called traffic amplification and reflection, the attacker can reflect his traffic from one or more third-party machines toward the intended target. In this type of assault, the attacker sends a message to a third party, while spoofing the Internet address of the victim. When the third party replies to the message, the reply is sent to the victim — and the reply is much larger than the original message, thereby amplifying the size of the attack.

BCP38 is designed to filter such spoofed traffic, so that it never even traverses the network of an ISP that's adopted the anti-spoofing measures. However, there are non-trivial economic reasons that many ISPs fail to adopt this best practice. This blog post from the Internet Society does a good job of explaining why many ISPs ultimately decide not to implement BCP38.

Fortunately, there are efforts afoot to gather information about which networks and ISPs have neglected to filter out spoofed traffic leaving their networks. The idea is that by "naming and shaming" the providers who aren't doing said filtering, the Internet community might pressure some of these actors into doing the right thing (or perhaps even offer preferential treatment to those providers who do conduct this basic network hygiene).

A research experiment by the **Center for Applied Internet Data Analysis** (CAIDA) called the "Spoofer Project" is slowly collecting this data, but it relies on users voluntarily running CAIDA's software client to gather that intel. Unfortunately, a huge percentage of the networks that allow spoofing are hosting providers that offer extremely low-cost, virtual private servers (VPS). And these companies will never voluntarily run CAIDA's spoof-testing tools.

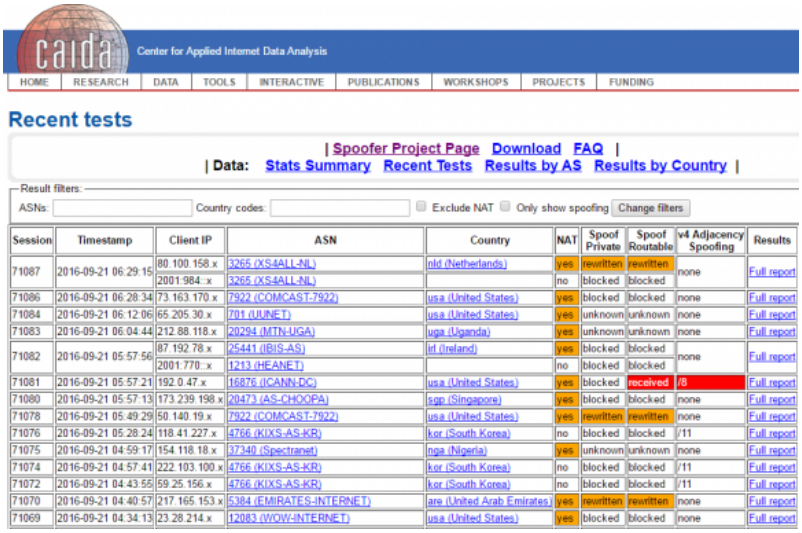

caida Center for Applied Internet Data Analysis

HOME | RESEARCH | DATA | TOOLS | INTERACTIVE | PUBLICATIONS | WORKSHOPS | PROJECTS | FUNDING

Recent tests

| Spoofer Project Page Download FAQ |

| Data: Stats Summary Recent Tests Results by AS Results by Country |

Result filters:

ASNs: Country codes: Exclude NAT Only show spoofing Change filters

Session	Timestamp	Client IP	ASN	Country	NAT	Spoof Private	Spoof Routable	v4 Adjacency Spoofing	Results
71087	2016-09-21 06:29:15	80.100.158.x	3265 (XS4ALL-NL)	nld (Netherlands)	yes	rewritten	rewritten	none	Full report
		2001:984::x	3265 (XS4ALL-NL)		no	blocked	blocked		
71086	2016-09-21 06:28:34	73.163.170.x	7922 (COMCAST-7922)	usa (United States)	yes	blocked	blocked	none	Full report
71084	2016-09-21 06:12:06	65.205.30.x	701 (UUNET)	usa (United States)	yes	unknown	unknown	none	Full report
71083	2016-09-21 06:04:44	212.88.118.x	20294 (MTN-UGA)	uga (Uganda)	yes	unknown	unknown	none	Full report
71082	2016-09-21 05:57:56	87.192.78.x	25441 (IBIS-AS)	irl (Ireland)	yes	blocked	blocked	none	Full report
		2001:770::x	1213 (HEANET)		no	blocked	blocked		
71081	2016-09-21 05:57:21	192.0.47.x	16876 (ICANN-DC)	usa (United States)	yes	blocked	received	/8	Full report
71080	2016-09-21 05:57:13	173.239.198.x	20473 (AS-CHOOPA)	sgp (Singapore)	yes	blocked	blocked	none	Full report
71078	2016-09-21 05:49:29	50.140.19.x	7922 (COMCAST-7922)	usa (United States)	yes	rewritten	rewritten	none	Full report
71076	2016-09-21 05:28:24	118.41.227.x	4766 (KIXS-AS-KR)	kor (South Korea)	no	blocked	blocked	/11	Full report
71075	2016-09-21 04:59:17	154.118.18.x	37340 (Spectranet)	nga (Nigeria)	yes	unknown	unknown	none	Full report
71074	2016-09-21 04:57:41	222.103.100.x	4766 (KIXS-AS-KR)	kor (South Korea)	no	blocked	blocked	/11	Full report
71072	2016-09-21 04:43:55	59.25.156.x	4766 (KIXS-AS-KR)	kor (South Korea)	no	blocked	blocked	/11	Full report
71070	2016-09-21 04:40:57	217.165.153.x	5384 (EMIRATES-INTERNET)	are (United Arab Emirates)	yes	rewritten	rewritten	none	Full report
71069	2016-09-21 04:34:13	23.28.214.x	12083 (WOW-INTERNET)	usa (United States)	yes	blocked	blocked	none	Full report

CAIDA's Spoofer Project page.

As a result, the biggest offenders will continue to fly under the radar of public attention unless and until more pressure is applied by hardware and software makers, as well as ISPs that are doing the right thing.

How might we gain a more complete picture of which network providers aren't blocking spoofed traffic — without relying solely on voluntary reporting? That would likely require a concerted effort by a coalition of major hardware makers, operating system manufacturers and cloud providers, including **Amazon**, **Apple**, **Google**, **Microsoft** and entities which maintain the major Web server products (**Apache**, **Nginx**, e.g.), as well as the major **Linux** and **Unix** operating systems.

The coalition could decide that they will unilaterally build such instrumentation into their products. At that point, it would become difficult for hosting providers or their myriad resellers to hide the fact that they're allowing systems on their networks to be leveraged in large-scale DDoS attacks.

To address the threat from the mass-proliferation of hardware devices such as Internet routers, DVRs and IP cameras that ship with default-insecure settings, we probably need an industry security association, with published standards that all members adhere to and are audited against periodically.

The wholesalers and retailers of these devices might then be encouraged to shift their focus toward buying and promoting connected devices which have this industry security association seal of approval. Consumers also would need to be educated to look for that seal of approval. Something like Underwriters Laboratories (UL), but for the Internet, perhaps.

THE BLEAK VS. THE BRIGHT FUTURE

As much as I believe such efforts could help dramatically limit the firepower available to today's attackers, I'm not holding my breath that such a coalition will materialize anytime soon. But it's probably worth mentioning that there are several precedents for this type of cross-industry collaboration to fight global cyber threats.

In 2008, the **United States Computer Emergency Readiness Team** (CERT) announced that researcher **Dan Kaminsky** had discovered a fundamental flaw in DNS that could allow anyone to intercept and manipulate most Internet-based communications, including email and e-commerce applications. A diverse community of software and hardware makers came together to fix the vulnerability and to coordinate the disclosure and patching of the design flaw.

In 2009, Microsoft heralded the formation of an industry group to collaboratively counter Conficker, a malware threat that infected tens of millions of Windows PCs and held the threat of allowing cybercriminals to amass a stupendous army of botted systems virtually overnight. A group of software and security firms, dubbed the Conficker Cabal, hashed out and executed a plan for corralling infected systems and halting the spread of Conficker.

In 2011, a diverse group of industry players and law enforcement organizations came together to eradicate the threat from the DNS Changer Trojan, a malware strain that infected millions of Microsoft Windows systems and enslaved them in a botnet that was used for large-scale cyber fraud schemes.

These examples provide useful templates for a solution to the DDoS problem going forward. What appears to be missing is any sense of urgency to address the DDoS threat on a coordinated, global scale.

That's probably because at least for now, the criminals at the helm of these huge DDoS crime machines are content to use them to launch petty yet costly attacks against targets that suit their interests or whims.

For example, the massive 620 Gbps attack that hit my site this week was an apparent retaliation for a story I wrote exposing two Israeli men who were arrested shortly after that story ran for allegedly operating vDOS — until recently the most popular DDoS-for-hire network. The traffic hurled at my site in that massive attack included the text string "freeapplej4ck," a reference to the hacker nickname used by one of vDOS's alleged co-founders.

Most of the time, ne'er-do-wells like Applej4ck and others are content to use their huge DDoS armies to attack gaming sites and services. But the crooks maintaining these large crime machines haven't just been targeting gaming sites. **OVH**, a major Web hosting provider based in France, said in a post on Twitter this week that it was recently the victim of an even more massive attack than hit my site. According to a Tweet from OVH founder **Octave Klaba**, that attack was launched by a botnet consisting of more than 145,000 compromised IP cameras and DVRs.

I don't know what it will take to wake the larger Internet community out of its slumber to address this growing threat to free speech and ecommerce. My guess is it will take an attack that endangers human lives, shuts down critical national infrastructure systems, or disrupts national elections.

But what we're allowing by our inaction is for individual actors to build the instrumentality of tyranny. And to be clear, these weapons can be wielded by anyone — with any motivation — who's willing to expend a modicum of time and effort to learn the most basic principles of its operation.

The sad truth these days is that it's a lot easier to censor the digital media on the Internet than it is to censor printed books and newspapers in the physical world. On the Internet, anyone with an axe to grind and the willingness to learn a bit about the technology can become an instant, self-appointed global censor.

I sincerely hope we can address this problem before it's too late. And I'm deeply grateful for the overwhelming outpouring of support and solidarity that I've seen and heard from so many readers over the past few days. Thank you.

Tags: Akamai, BCP38, CAIDA, censorship, Center for Applied Internet Data Analysis, Conficker Cabal, DDoS, DNS, DNS Changer Trojan, google, internet of things, IoT, John Gilmore, Project Shield, Prolexic, Rob Wright, SPECTRE, spoofer project, United States Computer Emergency Readiness Team

This entry was posted on Sunday, September 25th, 2016 at 7:58 am and is filed under Other. You can follow any comments to this entry through the RSS 2.0 feed. You can skip to the end and leave a comment. Pinging is currently not allowed.

194 comments

1. *Alex*
   September 27, 2016 at 12:44 pm

   In this situation what did your attackers accomplish? They censored you for a few days? But look at all the attention this is getting now! If I had to guess, I bet the number of your readers have since boosted. This attack has been covered by many outlets, and sending a lot of readers your way.

   I'm not sure how much of your time and money was lost due to this attack, but I'm guessing the attackers used up more of their own resources and even put themselves at greater risk of discovery.

   Congrats on coming back, I'm glad the censorship failed.

   Reply

   - *David Moore*
     September 30, 2016 at 8:58 am

     Agreed… I read Krebs… but when this happened, you can be rest assured I read a lot more to get a grip on the mitigation methods used. Good job Brian, Akamai, and Google for handling this properly. Redirecting to local host…. nice… but a shame you had to go offline.

     Reply

2. *Mark Ratledge*
   September 27, 2016 at 1:41 pm

   Brian,

   Glad you're back. I'd really be interested in "listening in" at Google to see them working to keep your site up and how Project Shield works, but I know few of us will ever be privy to that information.

   Mark

   Reply

3. *Daniel W*
   September 27, 2016 at 4:28 pm

   It appears this attack did not rely on spoofing. According to the NetworkWorld article: http://www.networkworld.com/article/3123672/security/largest-ddos-attack-ever-delivered-by-botnet-of-hijacked-iot-devices.html

Akamai chief security officer Andy Ellis stated: “The attack didn’t use reflection or amplification, so all the traffic consisted of legitimate http requests to overwhelm Krebs’s site.”

If so, shaming ISP’s for not preventing something attackers no longer depends on, would, at best, be a waste of resources.

Reply

- *BrianKrebs*
  September 27, 2016 at 4:39 pm

  Daniel,

  Who do you think reported that information first?

  https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/

  Also, this story you’re commenting on quite clearly delineates the attack on my site from other types of DDoS attacks that employ spoofing, amplification and reflection. So, yes, naming and shaming is a big part of the solution, as spoofing attacks are still very common.

  Reply

  - *Daniel W*
    September 29, 2016 at 10:29 am

    Brian, I know you did. I quoted Network World because they named the Akamai source. I also know spoofing is *currently* popular.

    My point is that stopping spoofing is, as demonstrated by this attack, no longer an efficient use of resources for the particular purpose of stopping DDOS, as it would just make all other DDOS operators switch tactic too, and carry on as usual.

    Hounding ISPs to stop spoofing would thus mainly consume resources they could spend on measures that still would be efficient against DDOS.

    For example, DDOS operators can obviously recognize the IoT devices involved. Reasonably ISPs could learn this too, and react when a webcam seemingly starts reading blogs over and over again.

    Reply

  - *Xander*
    September 29, 2016 at 12:45 pm

    Something Akamai may have missed as well, is some methods that technically fall under reflection or amplification attacks, do not require spoofing at all. over a number of hours i was hit with 10TB of DDoS from 5500 unique ips. I was on the system durring the attack analyzing it. It was thousands of wordpress sites, doing legitimate http requests in the form of pingbacks, as commanded by an IP elsewhere. That ip was listed in the useragent data, and it was the same in each request, no matter what wordpress site the hit came from. These wordpress sites were being remote controlled most likely by use of an un-patched or un-discovered vulnerability.

    I was able to discover that the source was instabooter.com in my analysis. I even recorded a log of all the IPs used to hit me, and provided them to a contact i have who does cyber-warfare activities.

    Brian, i have screenshots and logs if you want them 😉

    Reply

    - *Cedric Knight*
      October 21, 2016 at 3:52 pm

      The WordPress pingback attack *is* a reflected attack (servers spoof your site’s URL with calls to xmlrpc.php), but it’s a application-layer reflection. So, true, BCP-38 wouldn’t help, but something analogous at the reflector (rate-limiting, checking the IP address corresponds to the site) *would* prevent reflection and amplification. So that’s not an argument against BCP-38 for layer 3 attacks.

      Yes, you can have be overwhelmed by a directed Layer 3 attack from thousands of sources, but the ability to amplify attacks through spoofed DNS requests, half-open SYNs and so on means that the cost to the attacker is less than the cost to the defender. ISOC has been trying to promote BCP-38 for years, and there is an argument for naming and shaming.

      Reply

4. *Phillipp*
   September 27, 2016 at 5:12 pm

   I wondered if it wouldn’t be possible to *deflect* DDoS attack by observing the attacker IPs and using them against each other so that a major part of the attack bandwith is used against other attackers.

   At least in theory you would expect a 50% drop if the attacker uses half of his bandwith to attack other attackers. Asymmetric links may affect that, but it might be worth a shot.

Kind of an internet Aikido…

Reply

- *alex kent*
  September 28, 2016 at 7:28 am

  The problem with the idea of pitting attackers against each other is the same as the problem "why not just block all the attacking ip addresses?". It's really hard to figure out who the attackers are.
  Looking at a massive amount of incoming requests and realising you're being DDOS'd is easy, figuring out which of those requests are attackers and which are real users is very hard.

  Reply

5. *Christenson*
   September 27, 2016 at 6:47 pm

   The fundamental issue is that computers are not under the control of their nominal owners to various degrees.

   Until that becomes the case, I can always borrow millions of IOT devices to ask Brian's site for pages in whatever pattern I like…effectively DDOS'ing it. And blocking those devices from the net will be politically and technically impossible.

   Politically, how are you going to explain to someone that they are off the internet, or part of it, because their IOT device is browsing the web?

   Technically, the pattern of traffic will be made to look more and more like just a huge amount of human fans reading Brian's site.

   If Brian's website is to have more capacity, he (and we readers) will need to trust more entities that carry the data to us, and those trust relationships have a nonzero financial cost.

   It's a hard problem. As of this week, the Linux kernel maintainers have recognized the need to upgrade the fundamental security of Linux itself.

   ***
   My ideal IOT device won't talk to anything it hasn't been "next to" and had a chance to "friend", for some definition of physical proximity and exchange of addresses. It might share its "friends" list with the modem/router in my house.

   That would go a long way…

   Reply

6. *Ryan C*
   September 27, 2016 at 10:12 pm

   I find it hard to believe that there's DPI, but no way to figure out that traffic is coming from an IP address that an ISP doesn't own.

   Reply

7. *Laurent*
   September 27, 2016 at 10:47 pm

   Hello Brian,

   Is your site still under DDoS Attack? Is Google providing details on the attack and how they mitigate?

   Laurent

   Reply

8. *Mackke*
   September 28, 2016 at 3:53 am

   Brian, thank you for all you do in the ongoing fight
   against Internet aholes and parasites.

   Reply

9. *tz*
   September 28, 2016 at 11:51 pm

   Note that krebsonsecurity.com was thrown offline, but if your content was mirrored it would he hard to censor. Google search cache, minimally, but if many sites copied the articles whole, the attackers could not bring down one site.

   Something worse and more subtle, censorship does and will occur and Google will support it because they consider some things "hate speech" or at least politically incorrect so won't protect it. Right now they are demonetizing You-Tube when the videos are controversial, and isn't censorship (which I have to point out) but illustrates the problem.

   Most people don't want or care about "free speech". They will have their pets and clients which they will support, but if they say anything they don't like they will be purged.

Worse, whatever you think of Milo (@Nero until banned), social media is the most effective censor yet.

Reply

10. *Alan Polinsky*
September 29, 2016 at 6:31 am

Brian:

It's wonderful having you back and hope your new provider is able to properly protect yo against future DDoS attacks. I for one would be willing to contribute to a Kickstarter (or similar) campaign, if you start one.

Reply

11. *Jon*
September 29, 2016 at 7:51 am

Would lawsuits against manufacturers of devices that lack sound security be a mechanism to force all manufacturers to protect their machines? Government requirements on IOT devices?

One limitation of the voluntary approaches—the damage from hijacking is largely an externallity. A customer with a hacked camera that is part of a botnet may suffer little or no damage; it is the victims of the botnet that experience the cost.

Thus in situations where the customer does not value any damage incurred by hacking an IOT device, then he or she has little incentive to purchase a safe one.

Reply

12. *Robert LaValley*
September 29, 2016 at 10:12 am

Mr. Krebs, I don't know whether this is material, but, when I heard about the "largest DD0S attack in history" being directed against OVH, I scratched my head and thought, "Who the hell is 'OVH'?" Well, it turns out, OVH hosts WikiLeaks.

Also, I have seen stories over the past week that Ecuador is being increasingly pressured by Sweden to end its asylum of Julian Assange. Now, I don't like Assange, and I think he deserves whatever he gets, but, I am curious WHY he's being targeted so aggressively, now, SO CLOSE TO THE U.S. ELECTION.

Maybe I'm crazy.

Thanks for your time 🙂
Bob

Reply

13. *Betan Testravosky*
September 29, 2016 at 12:43 pm

There's no conspiracy here … it was a proof of concept test.

Of Akamai …

… which obviously wasn't up to the task.

I like Krebs for the reading entertainment value as a security blogger. I like Brian.

But Krebs as Public Enemy Number One for large scale DDOS'ers … eh, there are hundreds if not thousands of security bloggers just like Krebs.

So the true target was Akamai itself … and what it's response would be. Krebs was picked on because, let's face it, Krebs is going to post what the outcome was more than just a simple one line blog entry about being DDOS'd and sorry about the interruption.

And sure enough, Akamai folded like a cheap suit when it got just too hot and too much of a nuisance. Unlike Project Shield … Akamai and AWS and CenturyLink TS etc operate essentially for profit only. They throw a few free bones out to folks like Krebs which is basically to get good PR. But when those free rides start digging into their wallet, the free riders get the unceremonious boot.

A proof of concept attack that once and for all proved IoT devices are lousy at security, even large services like Akamai can shoved to its knees and kicked in the teeth, and Brian Krebs can get muzzled by a Smart Waffle Iron in the kitchen of some farmhouse in Kansas.

The real question is (aside from the above) … is why do we have Internet Connected Waffle Irons that can be zombied and botnetted in the first place? Why do I need a IoT Toothbrush? Colgate is even experimenting with disposable IoT Toothpaste tubes supposedly that can reorder toothpaste from Amazon when you've used 3/4s of the tube. Meaning Krebs could in the future be attacked by the same product that attacks tartar.

Reply

- *Roger Bradley*
October 1, 2016 at 6:17 am

Excellent. You hit the nail right on the head.

Reply

- *Dror Harari*
  October 3, 2016 at 5:00 am

  Betan Testravosky's argument is very compelling. I also view the planted reference to the hacker name as a ruse.

  However, just like this being a kind of proof-of-concept test on Akamai, it also goes the other way – both as a teaching moment and a wakeup call.

  One thing that works very well in the real world is financial incentives. If such a DDoS traffic is legislated as prohibited from being compensated for (meaning, if an ISP egress a 50Gbbs DDoS traffic, it would be illegal for that traffic to be used in the traffic balance between ISPs). This may help put the incentive at the right point to ensure the bright minds are put to work on the problem.

  Ignoring attacks based on bugs and protocol deficiencies for now, the main reason why the recent attacks are difficult to address is because the individual request seems rather legitimate. Still, when zooming out, the pattern becomes clear and much harder to obfuscate. An ISP could easily identify IP addresses egressing attack traffic (especially if some global alert system is set up to coordinate dissemination of information about active attacks).

  This is just a hint at a direction but with the right incentives and with smart people putting their minds to it, the Internet will be a much safer and resilient fabric.

  Last word on censorship — censoring sites is easy and has been done by governments for a long time — be it the great firewall of China, the blocking of bit-torrent pirated media site, censoring of 'irrelevant personal data' in the EU, etc. The kind of censorship we're talking about here is more of the underground type. It goes against the infrastructure of the Internet and that's why the biggest actors – namely states – should be very interested in being able to protect against this phenomenon.

  Reply

  - *Alex Moen*
    October 17, 2016 at 11:04 am

    Dror,

    First, full disclosure: I work for an ISP. Your idea of laying all of the blame and cost of DDOS attacks at the feet of ISPs is not well thought out. There is no process, appliance, or service that can effectively or easily "identify IP addresses egressing attack traffic". There is no national dynamic database that identifies attackers. I wish there was, and I wish we could purchase it right now, because the fallout of these attacks are crushing to our business and to our customers! But, it just does not exist.

    And actually, you're blaming the victim here. ISPs are the ones who bear the brunt of these attacks, and attempting to mitigate them from affecting our networks is, simply put, impossible. What are we supposed to do: tell our paying customers that they are not allowed to use a certain brand of IP surveillance camera, because the company who produces that camera doesn't care about security? Tell a customer that their $300.00 firewall that they just purchased from Best Buy (who assured them that it is the best router on the market) needs to be disconnected because it's participating in a botnet? "What do *you* know, the Best Buy guy told me it's the one to have!" is the response we hear. So, do you seriously think that will work? The customer is paying for his connection, and in his mind, he can use it for anything that that he wants. And, rightfully so.

    Look at the amount of bandwidth Netflix is consuming: we aren't allowed to touch that issue, because "net neutrality" (again, full disclosure: I am also a Netflix subscriber). How long before IoT devices are included in the net neutrality rules, because they are considered by the current administration at the time as a basic right of the modern American? ISPs have been screaming about security for years; no one listens. We're effectively told "Quiet down: you're simply providing a pipe, you have nothing to say about it." and "Don't touch my traffic, because if you do, it's censorship and/or discrimination."

    Security is the very last thing considered in the manufacture of devices or writing of code or protocols. It has been the bane of networks since the inception of networks. Look at protocols like Zigbee and their security flaws, even though they are extensively used in strategic infrastructure, like the US power grid. Even those extremely important protocols, which should have had security in the forefront of their development, fall far short of the mark.

    This issue must be dealt with by the manufacturers of the "infected" devices and the weak software. Until there is some kind of industry and consumer backlash against these manufacturers, nothing will be done about it, and consumers will continue to purchase them by the thousands, to be used in botnets and participating in DDOS attacks. It's not the ISP's fault, and the ISP is powerless to prevent it.

    Reply

14. *Brad Regan*
    September 29, 2016 at 3:28 pm

    "Democratization of Censorship" is just another term for mob rule. The mob reacts to rumor, scant facts, conspiracy theories, etc. and its reactions are multiplied by shear numbers. A justice system attempts to gather facts and present both sides in a safe atmosphere relatively free of emotional reaction so that cooler head can prevail. Unfortunately, the Internet, social media, email and other mass communications methods facilitate the rule of the mob. Not sure of a solution to this but we are definitely entering a period in which reason, logic, and fairness are losing out to ignorance, corruption, and violence.

    Reply

15. *Stan Stahl*
September 29, 2016 at 7:23 pm

Glad to see you back online. The last few days has been a long intermittent sequence of disappointments as you've been offline.

Good article. You're so right that the pieces are in place to meet the DDoS challenge; all that's required is the coordinated will to act.

I'm beginning to see good signs of of "coordinated will." Here in LA, banks and the FBI are talking to organizations about cyber. And of course, on the technical side, there's the Cyber Threat Alliance and the various ISACs. It only through these kinds of "coordinated will" — including those like you described — that we'll meet the DDoS and other cybersecurity challenges.

Like you, though, I wonder if it will be enough soon enough; or if it will take a major cyber-catastrophe before we rally together.

Reply

16. *Industrial IT*
September 30, 2016 at 12:09 pm

It's great that you're back up and running. I'd love to be able to go to Google to find out exactly how Project Shield works and watch how they keep your site safe in the future. Although I know this is very unlikely to happen!

Good work!

Reply

17. *JimV*
September 30, 2016 at 9:29 pm

Brian, you've gotten a nice pat on the back from the Wall Street Journal who dug a little further into some of the business aspects involved:

http://www.wsj.com/articles/hackers-infect-army-of-cameras-dvrs-for-massive-internet-attacks-1475179428?tesla=y

Sorry you bore the brunt of this nastiness, but keep up the excellent work and thanks for all you do!

Reply

18. *Ian*
October 5, 2016 at 11:47 am

Could software be written that searches and detects bot net hijacks and re-directs the DDos back to the hacker?

(Like the ping -sting counter attack in the James Bond movie — the one with the Russians and the guy shooting "I am invincible!")

Reply

19. *Hallowed Hill*
October 6, 2016 at 9:34 am

Hey Brian, another great read.

After reading the Verisign report, one of the facts that jumped out was that the average peak attack size dropped significantly from Q1 to Q2 of 2016 (From 19.37 to 17.37 Gbps).

As you suggest, this seems to corroborate that the system has been calibrated, and true potential threat remains to be seen.

Is there any evidence that the technology community is reducing the threat by voluntarily implementing more secure practices?

Reply

20. *cctv dvr 8 channel*
October 6, 2016 at 12:37 pm

Finally, the CCTV systems that these security companies setup are already fully integrated and complete with security lighting, access control, as well as alarms. The majority of CCTV cameras available today usually are for surveillance and security purposes. cctv dvr 8 channel dvr security Wireless CCTV cameras will be the upgraded version of CCTV system.

It all hangs on the harddrive capacity that is built in to the DVR.
IP CCTV is simply a system of surveillance cameras wired to a single network which enables businesses to continuously monitor every part of the premises.

Reply

21. *Clayton E. Cramer*
October 12, 2016 at 9:09 am

Did not even know who you are. Just added you to my blogroll.

Reply

22. *offshore vps*
October 19, 2016 at 12:05 am

Why hosting companies allow ddos. There is so many company which blocked ddos scripts.

Reply

← Older Comments

Leave a comment

Name (required)

Email (required)

Website

Comment

Submit Comment

- ## My New Book!

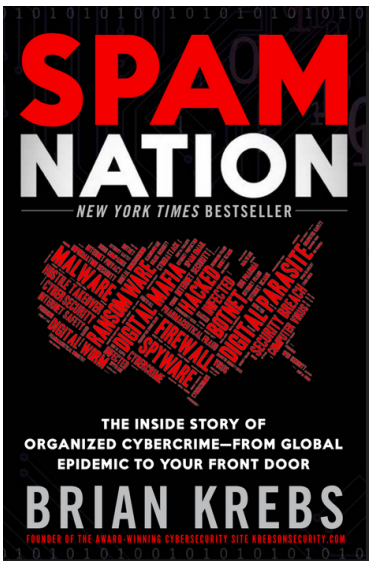

A New York Times Bestseller!

- Buy at Amazon
- Donate with PayPal

- ## Recent Posts

  - Hacked Cameras, DVRs Powered Today's Massive Internet Outage
  - DDoS on Dyn Impacts Twitter, Spotify, Reddit
  - Spreading the DDoS Disease and Selling the Cure
  - Hackers Hit U.S. Senate GOP Committee
  - Self-Checkout Skimmers Go Bluetooth

- ## Subscribe by email

Please use your primary mailbox address, not a forwarded address.

Your email:

Enter email address...

Subscribe Unsubscribe

- ## All About Skimmers

Click image for my skimmer series.

- ## The Value of a Hacked PC

Badguy uses for your PC

- ## Tools for a Safer PC

Tools for a Safer PC

- ## The Pharma Wars

Spammers Duke it Out

- ## Badguy Uses for Your Email

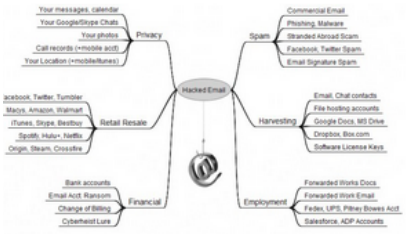

Your email account may be worth far more than you imagine.

- ## eBanking Best Practices

eBanking Best Practices for Businesses

- ## Most Popular Posts

  - Online Cheating Site AshleyMadison Hacked (798)
  - Sources: Target Investigating Data Breach (620)
  - Cards Stolen in Target Breach Flood Underground Markets (445)
  - Reports: Liberty Reserve Founder Arrested, Site Shuttered (416)
  - Was the Ashley Madison Database Leaked? (377)
  - True Goodbye: 'Using TrueCrypt Is Not Secure' (363)
  - Who Hacked Ashley Madison? (360)
  - Following the Money, ePassporte Edition (353)
  - U.S. Government Seizes LibertyReserve.com (315)
  - Extortionists Target Ashley Madison Users (310)

- ## Category: Web Fraud 2.0

Innovations from the Underground

- 

ID Protection Services Examined

- **Is Antivirus Dead?**

The reasons for its decline

- **The Growing Tax Fraud Menace**

File 'em Before the Bad Guys Can

- **Inside a Carding Shop**

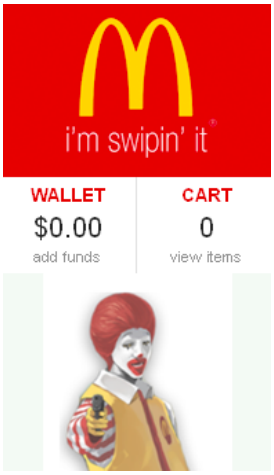

A crash course in carding.

- **Beware Social Security Fraud**

Sign up, or Be Signed Up!

- **How Was Your Card Stolen?**

Finding out is not so easy.

- **Krebs's 3 Rules…**

...For Online Safety.
