How Cloudflare's Architecture Can Scale to Stop the Largest Attacks

1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com

Executive Summary

New Mirai-based IoT botnets are used to stage the largest ever cyber-attacks. In a recent attack against Dyn, huge patches of the Internet were affected with a large number of high profile sites experiencing service disruptions and outages. Traditional hardware based DDoS mitigation services based in a few scrubbing locations cannot scale to win in the arms race against distributed and essentially free botnets. Cloudflare believes that architecture matters and that the only solution against massively distributed botnets is a massively distributed network. Cloudflare has based its service on this architectural approach, and while there are limits to any service, so far Cloudflare has not been impacted by any IoT botnet attack.

New Mirai-based IoT Botnets Pose a New Degree of Cyber-Attacks

On 21 October 2016 a massive and sustained distributed Denial-of-Service (DDoS) attack impacted huge parts of the Internet, interrupting or bringing down high profile web sites and services, such as Airbnb, Amazon.com, BBC, CNN, Comcast, DirecTV, Fox News, Netflix, The New York Times, Paypal, Pinterest, Reddit, Tumblr, Twitter, Verizon Communications, Visa, The Wall Street Journal, Yelp, Zillow and many others (a more comprehensive list can be found at https://en.wikipedia.org/wiki/2016_Dyn_cyberattack). The direct target of the attack was Dyn, a DNS service provider, which maps domain names to their Internet Protocol (IP) addresses so that traffic can be routed to a specific site. Dyn managed to resolve the incident after 11 hours, bringing the services of all impacted websites back to normal.[1]

> *"Assuming a device is publicly accessible, the chance of being hacked is probably 100 percent. The IPv4 address space just isn't that big. You can now run a scan across that entire space in hours, especially if you have a big botnet. The scans for vulnerability are continuous, and if anything, have accelerated over the last couple of years.*
>
> - Matthew Prince
> CEO of Cloudflare

The hackers compromised 100,000 connected devices, such as home routers, smart home gadgets, security cameras or video recorders, to create an Internet-of-Things (IoT) botnet, which was used to launch the attack with potentially as much as 1.2 Tbps of traffic.

The devices were hacked using a malware called Mirai. Mirai first scans the Internet for devices, which still have the factory default username and password settings, making it easy to then log in, take control of the device and infect it with the Mirai malware. The owner of the devices will not notice that the device was compromised, other than occasional sluggish performance.

The Mirai source code was recently uploaded to hacker forums and is now accessible to anyone to carry out the next attack.

IoT Botnet Attacks Against Cloudflare

During the last few weeks Cloudflare has also experienced DDoS attacks appearing to come from IoT botnets. DDoS attacks against Cloudflare have historically been volumetric in nature, targeting the network (3) and transport (4) layers of the OSI model, using the familiar DDoS amplification, DDoS flood and DDoS SYN flood attack approaches. The new types of attacks Cloudflare is seeing are heavily focused on the application layer (7). While it is not possible for Cloudflare to investigate all of the infringing devices, it is fair to say that these attacks came from an Internet-of-Things (IoT) category of botnet. Multiple hints confirm this theory:

- All of the attacking devices have port 23 (telnet) open (closing connection immediately) or closed. The traffic is never filtered. This is a strong hint that the malware disabled the telnet port just after it installed itself
- Most of the hosts from the Vietnamese networks look like connected CCTV cameras. Multiple of them have port 80 open while presenting a "NETSurveillance WEB" page
- The Ukrainian devices are a bit different, and most have port 80 closed, making it harder to identify
- Cloudflare had noticed one device with port 443 open serving a valid TLS cert issued by Western Digital, handling domain device-xxxx.wd2go.com suggesting it was a Network Attached Storage device

Although the most recent attacks have mostly involved Internet-connected cameras, there's no reason to think that they are the only device of choice for future DDoS attacks. As more and more devices—fridges, fitness trackers, sleep monitors, etc. are brought online they too will likely be unwilling participants in future attacks. There is no doubt of the existing arms race with attackers. Thus far Cloudflare has managed to withstand even the latest IoT botnet attacks without any service interruption to customers.

Cloudflare's Architecture—Designed to Stop What Comes Next

While there are limits to any service, ultimately only the right architecture can stop whatever comes next. Before delving into Cloudflare's architecture, it is worth taking a second to think about another analogous technology problem that is better understood: scaling databases.

Analogy: How Databases Scaled

From the mid-1980s, when relational databases started taking off, through the early 2000s, the way companies scaled their database was by buying bigger hardware. Companies were purchasing the biggest database servers they could afford, filling it with data, and hoping that hardware companies would release a bigger and still affordable server before the existing ones ran out of room. Hardware companies responded with more and more exotic, database-specific hardware.

At some point, the bounds of a box couldn't contain all of the data some organizations wanted to store. Google is a famous example. Back when the company was a startup, Google didn't have the resources to purchase the largest database servers. Even if they could the largest servers still could not store everything Google wanted to index—which was, literally, everything.

Rather than going the traditional route, Google wrote software allowing many cheap, commodity servers to work together as if they were one large database. As Google continued to develop more services, the software became

efficient at distributing large loads across all the machines in Google's network, maximizing the utilization of network, compute, and storage. As needs proceeded to grow, Google was in a position to simply add more commodity servers, scaling their network linearly.

Legacy DNS and DDoS Mitigation

The analogy of scaling databases can be extended to the way legacy DNS and DDoS services mitigate attacks. Traditionally, the way to stop an attack was to buy or build a big box and use it to filter incoming traffic. Most legacy DDoS mitigation service vendors use hardware from companies like Cisco, Arbor Networks, and Radware clustered together into "scrubbing centers."

Just like in the old database world, there were tricks to get these behemoth mitigation boxes to work together, but it was kludgy. Physical limits on the number of packets a single box could absorb became the effective limit on the total volume that could be mitigated by a service provider. In very large DDoS attack situations, most of the attack traffic will never reach the scrubbing center because, with only a few locations, upstream ISPs become the bottleneck.

The expense of the equipment meant that it is not cost effective to distribute scrubbing hardware broadly. How often would a DNS provider get attacked? How could this provider justify investing in expensive mitigation hardware in every one of the data centers? It was typical for legacy DDoS vendors to only provision their service when a customer came under attack; it never made sense to have capacity beyond a certain margin over the largest attack previously seen. It seemed rational that any investment beyond that was a waste, but that conclusion is proving ultimately fatal to the traditional model.

The Future Doesn't Come in a Box

Cloudflare views its infrastructure much more like Google sees its database. In Cloudflare's early days, traditional DDoS mitigation hardware vendors pitched Cloudflare to use their technology. Cloudflare even considered building mega boxes themselves, using them just to scrub traffic. It seemed like a fascinating technical challenge, but Cloudflare realized that it would never be a scalable model.

Instead, Cloudflare started with a very simple architecture. Cloudflare's first racks had only three components: router, switch, server. Today the rack is even simpler, often dropping the router entirely and using switches that can also handle enough of the routing table to route packets across the geographic region the data center serves.

Rather than using load balancers or dedicated mitigation hardware, which could become bottlenecks in an attack, Cloudflare wrote software that uses Border Gateway Protocol (BGP), the fundamental routing protocol of the Internet, to distribute load geographically and also within each data center in the network. For Cloudflare's model to work properly, it's critical that every server in every rack is able to answer every type of request. Cloudflare's software dynamically allocates traffic load based on what is needed for a particular customer at a particular time. That means that Cloudflare automatically spreads load across literally tens of thousands of servers during large attacks.

It also means that Cloudflare can cost-effectively continue to invest in its network. If Frankfurt needs 10 percent more capacity, Cloudflare can ship it 10 percent more servers, rather than having to make the step-function decision of whether to buy or build another Colossus Mega Scrubber™ box.

Since every core, in every server, in every data center can help mitigate attacks, each new data center Cloudflare brings online makes the service better and more capable of stopping attacks nearer to the source. In other words, the solution to a massively distributed botnet is a massively distributed network. This is how the Internet was meant to work: distributed strength, not focused brawn, within a few scrubbing locations.

How Cloudflare Made DDoS Mitigation Essentially Free

The efficient use of resources doesn't only come with capital expenditures, but also operating expenditures. Because Cloudflare uses the same equipment and networks to provide all of its functionality, Cloudflare rarely has any additional bandwidth costs associated with stopping an attack. In order to understand this, it is helpful to understand how Cloudflare buys bandwidth.

Cloudflare pays for bandwidth from transit providers on an aggregated basis, billed monthly, at the 95th percentile of the greater of ingress vs. egress. Ingress is network speak for traffic being sent into the Cloudflare network. Egress is traffic being sent out of the network.

In addition to being a DDoS mitigation service, Cloudflare also offers other functionalities, including caching. The nature of a cache is that there should always be more traffic going out from the cache than coming in. In Cloudflare's case, during normal circumstances, there is many times more egress (traffic out) than ingress (traffic in).

Large DDoS attacks drive up Cloudflare's ingress but don't affect the egress. Even in a very large attack it is extremely rare that ingress exceeds egress. Because Cloudflare only pays for the greater of ingress vs. egress, and because egress is always much higher than ingress, Cloudflare effectively has an enormous amount of zero-cost bandwidth to soak up attacks.

As Cloudflare's services continue to expand, the capacity to stop attacks increases proportionately. People wonder how Cloudflare can provide DDoS mitigation at a fixed cost to customers, regardless of the size of the attack; the answer is because attacks don't increase the largest of Cloudflare's unit costs. And, while legacy providers have stated that their offering pro bono DDoS mitigation would cost them millions, Cloudflare is able to protect politically and artistically important sites against huge attacks for free, through Project Galileo (https://www.cloudflare.com/galileo/), without it breaking the bank.

Winning the Arms Race

Cloudflare is the only DNS provider that was designed, from the beginning, to mitigate large scale DDoS attacks. Just as DDoS attacks are by their very nature distributed, Cloudflare's DDoS mitigation system is also distributed across its massive global network.

There's no doubting the existence of an arms race with attackers. Cloudflare is positioned technically and economically to win that race. Against most legacy service providers, attackers have an advantage: providers' costs are high because they have to buy expensive boxes and bandwidth, while attackers' costs are low because they use an overwhelming number of hacked devices. That's why Cloudflare's secret sauce is the software that allocates the load across Cloudflare's massively distributed network of commodity hardware. By keeping the costs low, Cloudflare is able to continue to expand its capacity efficiently and stay ahead of attacks.

Today, Cloudflare believes it has more capacity to stop attacks than the publicly announced capacity of all its competitors—combined. And Cloudflare continues to expand, opening nearly a new data center a week. The good news for Cloudflare's customers is that Cloudflare was designed in such a way that capacity can continue to be cost effectively scaled as attacks grow. There are limits to any service, and Cloudflare remains ever vigilant for new attacks, but is confident that its architecture is ultimately the right way to stop whatever comes next.

Takeaways

Protect yourself against all DDoS attacks, including the new, large IoT-botnet based DDoS attacks by setting up Cloudflare. **The setup is very simple and usually takes less than 5 minute to get up and running.**

Check out the plans, ranging from free to enterprise at www.cloudflare.com/plans.

References

[1] **Dyn status updates** https://www.dynstatus.com/incidents/nlr4yrr162t8

Dyn analysis summary http://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/?AID=12159285&PID=7596969&SID=66960X1514734X65bc674770b60a094e1f9cf00dfe6008&utm_source=CJ&utm_medium=Affiliate&utm_term=Skimlinks&utm_content=7596969&utm_campaign=InboundCJ

Wikipedia, 2016 Dyn cyberattack https://en.wikipedia.org/wiki/2016_Dyn_cyberattack

1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com