NEUSTAR SECURITY

Q1, 2019 Cyber Threats & Trends Report

neustar

Table of Contents

CAN NETWORK DEFENSE BE AUTOMATED?

Ever since the arguably "first" Denial of Service (DoS) attack in 1974, attackers have been looking for ways to make these incursions more powerful and more difficult to recognize. In the last forty-five years there have certainly been new threat vectors that have been discovered, but the vast majority of DoS and Distributed Denial of Service (DDoS) attacks that we've seen continue to use tried-and-true methodology. The biggest difference is the power that Artificial Intelligence (AI) and Machine Learning can add – to both attack and defense.

On the attack side, the rise of AI provides new ways to use existing attacks more efficiently, particularly when considering the use of Internet of Things (IoT) devices. Another troubling vector is the uptick in AI-controlled devices and services themselves. Not only could this result in bad actors manipulating AI programs, it could enable them to use AI to learn and exploit vulnerabilities in their targets' defenses.

There is unquestionably a place for AI and Machine Learning on the defense side, as well. At Neustar, we are building methods that will work together across our product portfolio. The ability to share information across different network entry points, as well as to build Machine Learning into the process, will make every Neustar service work more effectively. We are committed to using our capabilities at every point in the security process where it can help to speed defense or foreshadow issues.

It is vital to realize, however, that while AI is important, human ingenuity is also key in delivering security. That is why our Security Operations Center (SOC) engineers are available to help our customers at any time, 24/7/365. As you will see throughout this report, our SOC is involved not only in mitigating attacks to which our infrastructure alerts us, but also to working collaboratively with customers as they face security issues. The power of engineering expertise and security infrastructure together form a highly effective cybersecurity solution.

It is vital to realize, however, that while AI is important, human ingenuity is also key in delivering security.

Rodney Joffe, Senior Vice President, Senior Technologist and Fellow

Rodney Joffe serves as a Neustar Senior Vice President and is a Senior Technologist and Fellow. His accomplishments include founding the first commercial Internet hosting company, Genuity, as well as the first outsourced and cloud-based Domain Name System (DNS) company, UltraDNS, where he invented Anycast Technology for DNS. Joffe has served on a number of the U.S. government's cybersecurity intelligence panels and was the leader of the groundbreaking Conficker Working Group. He is one of the first civilians to receive the Federal Bureau of Investigation (FBI) Director's Award due in no small part to his role in uncovering and taking down the Butterfly Botnet. He has also been honored with the Mary Litynski Lifetime Achievement Award from M3AAWG, the global Messaging, Malware and Mobile Anti-Abuse Working Group, and was most recently publicly recognized for his years of work and dedication in helping protect against cybercrime, winning The Computing Security Award for his contribution to Cyber Security in 2018.

Joffe is also the chairman of the Neustar International Security Council (NISC), which is comprised of an elite group of cybersecurity leaders across industries and companies who meet regularly to discuss the latest cyberattack trends.

A LOOK BACK AT 2018

Cybersecurity, particularly around DDoS attacks, saw new highs in 2018. One of biggest headlines was the 1.35 Tbps attack against GitHub's website in March of last year. Around the same time, a 1.7 Tbps attack against an unnamed entity was cited, but not reported on in detail. Both attacks used the memcached (pronounced "mem-cash-dee") amplification vector.

This free, open-source memory caching system was designed to speed up dynamic database-driven websites. Memcached was designed to be internal to the organizations using it; unfortunately, some websites were left open to the Internet, which created the DDoS opportunity. A small query to a memcached system can result in a very large response, and this attribute was used to launch the large attacks mentioned above. According to the German Federal Office of Information Security's Computer Emergency Response Team (CERT-Bund) Report, "Memcached servers openly accessible from anywhere on the Internet via UDP are abused for DDoS reflection attacks against third parties on a regular basis. This way, extremely high amplification factors can be achieved which pose a serious security threat."

It is interesting to note that while there is the occasional spike, the majority of DDoS attacks seen across the industry are relatively small; a trend mirrored in Neustar's SOC in Q1, 2019. What these DDoS incursions might lack in size, however, they make up for in the number of vectors employed in each attack. When considering the highest bandwidth attacks in the last quarter, for example, over half contained at least 3 different vectors.

As attacks have become more sophisticated, they have become a greater concern to IT executives surveyed by NISC. NISC includes over 170 senior security experts, primarily C-Suite executives and senior decision makers, representing both small and large companies. In the most recent survey of these professionals, most indicated that they believe that DDoS is the top threat to their network (figure 1).

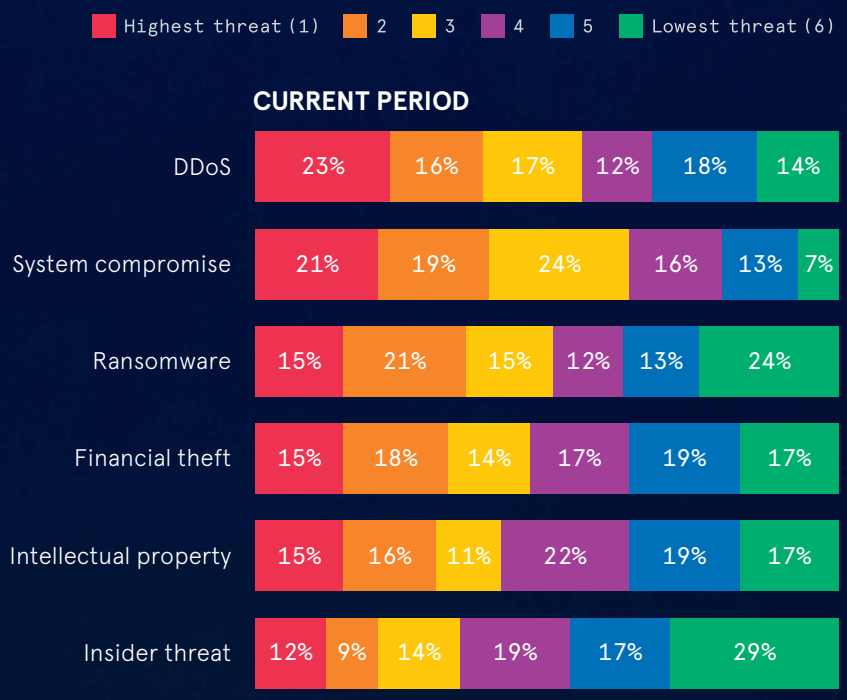

Figure 1. Cyberthreats ranked in order of concern, January 2019 NISC report

Forty eight percent of executives surveyed believe that DDoS attack threats have increased (figure 2), while 42% believe that they had increased their ability to respond to such threats (figure 3).

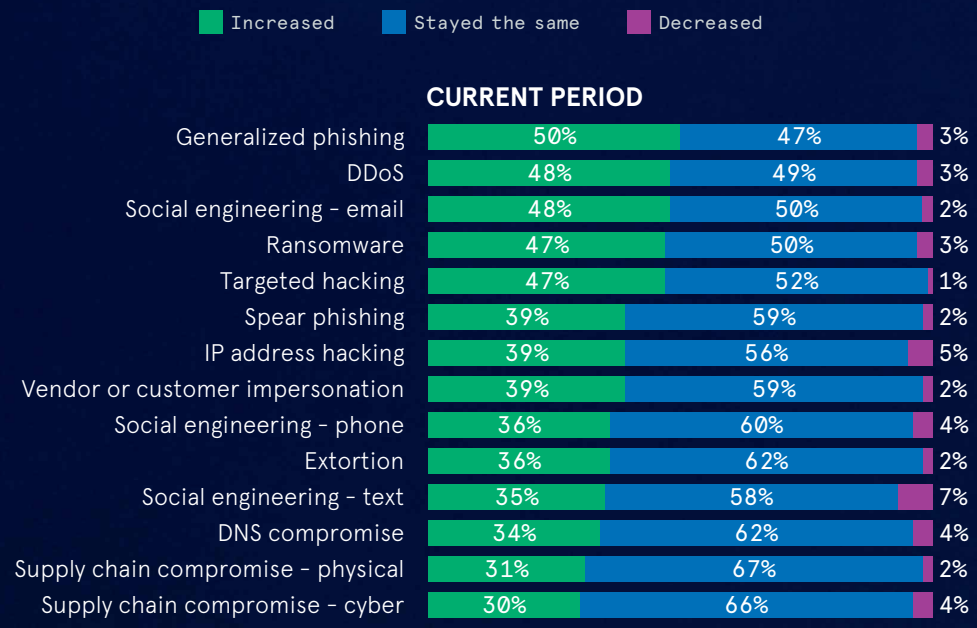

Figure 2. Threat of attack by various vectors has increased, January 2019 NISC report

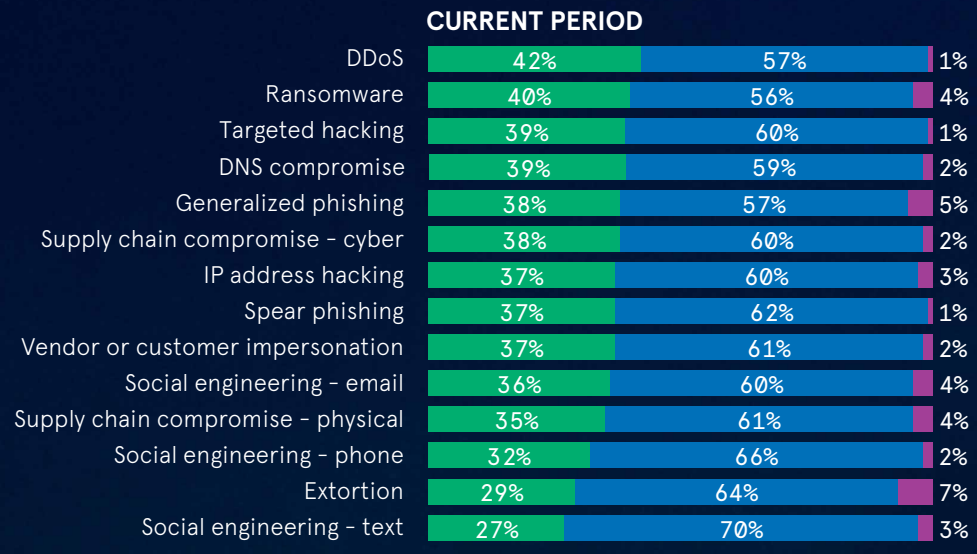

Figure 3. How organizations' ability to respond to threats has changed, January 2019 NISC report
NOTE: these figures reported represent an increase of 10% over the 10-month average

IoT botnets, which first entered common parlance with the Mirai attacks that brought down the Krebs on Security website and DNS provider Dyn in 2016, continue to grow. These botnets, which some criminals have made available by the hour, take advantage of poor security features and default security credentials, which are seldom changed by the consumers that use them. This trend is increasingly turning up in enterprises as well. Internet-connected devices are typically Linux or UNIX-based and available 24/7/365. A skilled bot herder can muster sufficient numbers to launch a DDoS attack within minutes. It is therefore no surprise to see that IT executives believe that the largest risks are now posed by criminals (figure 4).

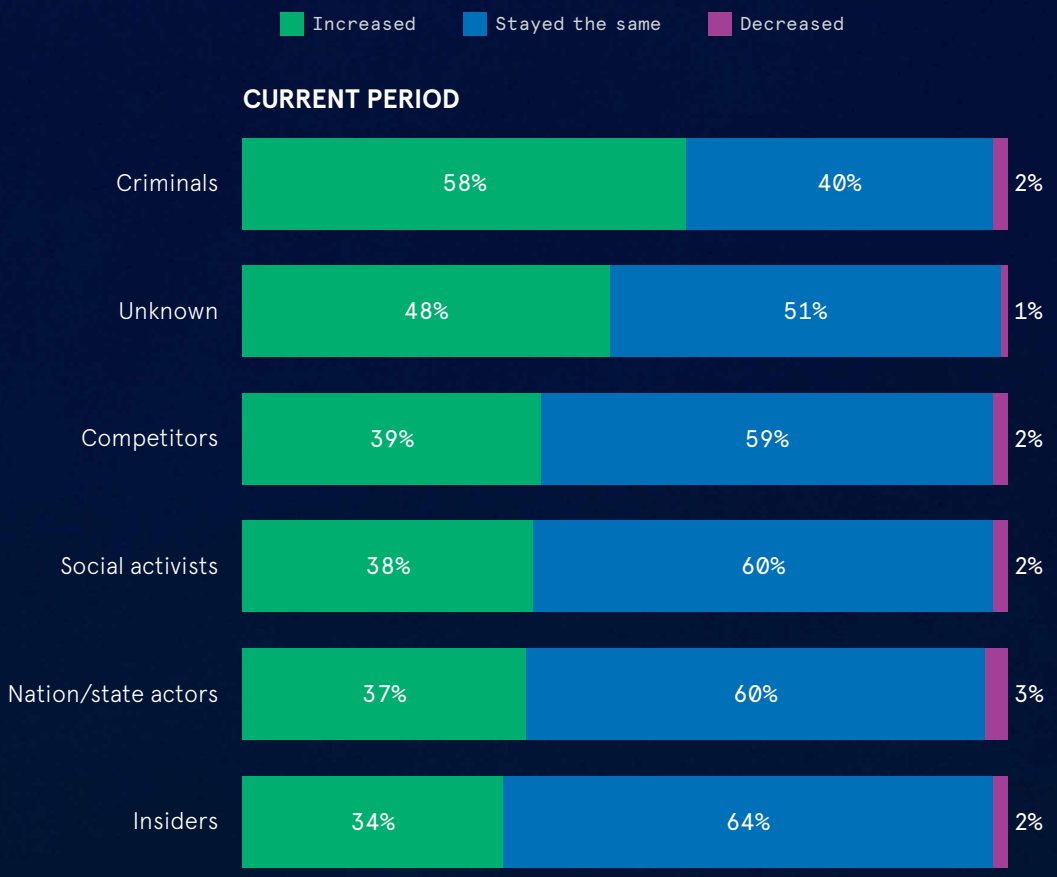

Figure 4. How the risk of attack from various actors has changed, January 2019 NISC report

Q1, 2019 THREATS AND TRENDS

This section contains the observations and insights derived from DDoS attack mitigations enacted on behalf of, and in cooperation with, customers of Neustar DDoS Protection Service during Q1, 2019[1]. This report offers a unique view into the attack trends unfolding online, including attack statistics and behavioral trends for Q1, 2019.

Attack Summary

Comparing Q1, 2019 vs. Q1, 2018, the number of attacks on directly provisioned customers has increased 200%.

The largest attack size observed in Q1, 2019 was 587 Gbps in volume. The largest attack size observed in Q1, 2018 was 345 Gbps in volume. The longest duration for a single attack was nearly a day and a half.

Comparing the number of attacks in Q1, 2019 with the number of attacks in the same time period of 2018, Neustar observed a 257% increase in attacks 5 Gbps and below, but, more significantly, observed a 967% increase in attacks 100 Gbps and higher (figure 5).

200%
Increase in the number of attacks comparing Q1, 2019 to Q1, 2018

587 Gbps
Largest attack size Q1, 2019

>70%
Increase in the largest attack size comparing Q1, 2019 to Q1, 2018

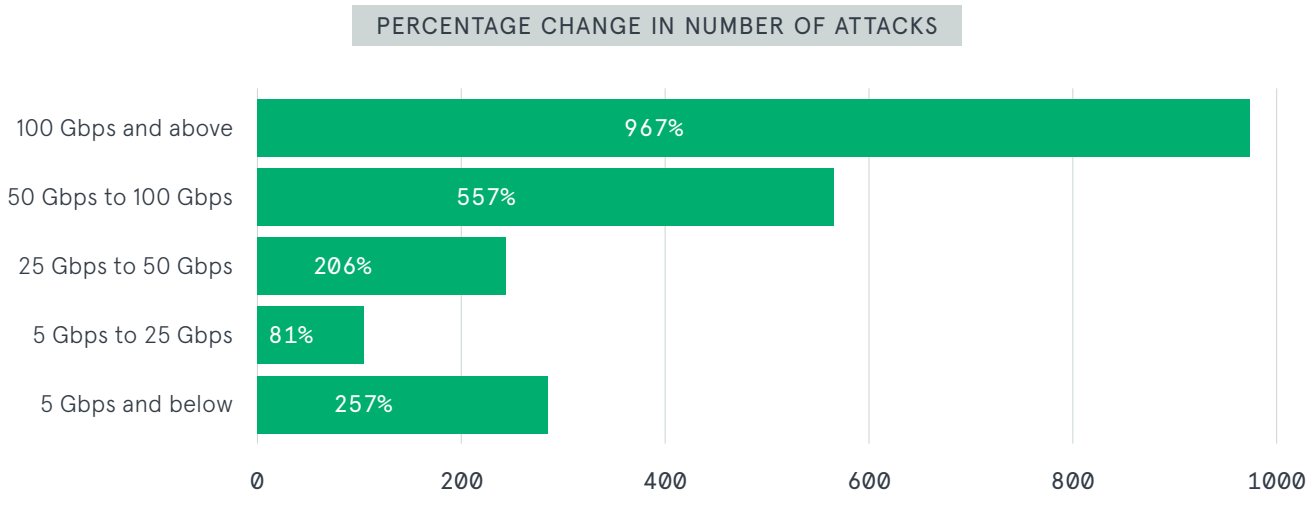

Figure 5 - Percent change in number of attacks by size range Q1, 2019 vs. Q1, 2018

Attack Volume

In Q1, 2019, almost 58% of all attacks mitigated by Neustar were 5 Gbps or less.

While there is an incidence of attacks between 25 and 100+ Gbps in each time period considered, the majority of attacks in both periods were 25 Gbps and below.

Comparing Q1, 2019 to Q1, 2018, attacks within specified size ranges decreased in only one area; 5 Gbps to 25 Gbps. Between Q1, 2019 and Q1, 2018, the highest percent of growth was observed in attacks 100 Gbps and above.

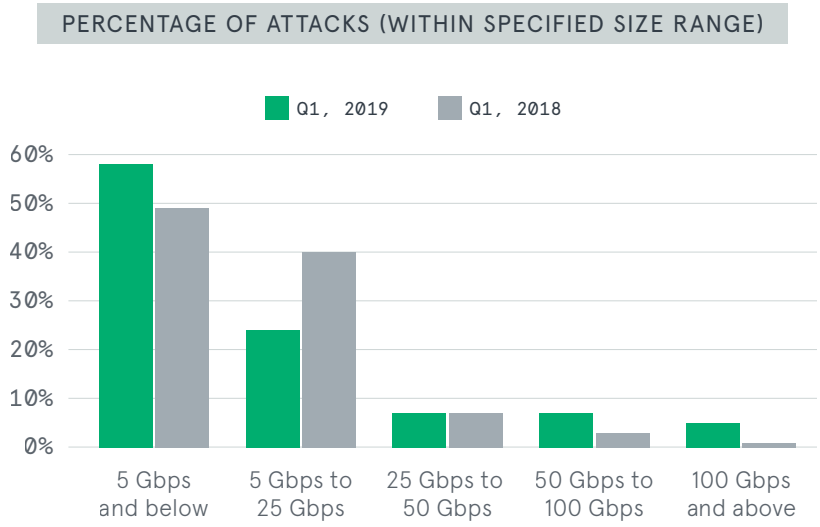

Figure 6 - Comparison of attacks by size Q1, 2019 vs. Q1, 2018

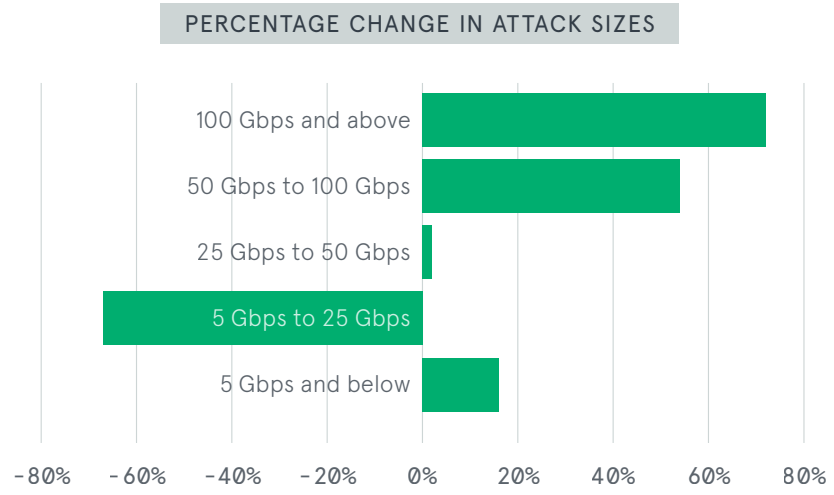

Figure 7 - Change in attacks sizes in specified ranges, Q1, 2019 vs. Q1, 2018

24.5 Gbps
Average attack size Q1, 2019

14.1 Gbps
Average attack size Q1, 2018

73%
increase in average attack size comparing Q1, 2019 to Q1, 2018

Attack Intensity

Comparing the intensity of attacks in Q1, 2019 vs. the intensity of attacks in the same period for Q1, 2018, Neustar observed a 122% increase in attack volume.

193 Mpps Largest packets per second in an attack Q1, 2019

87 Mpps Largest packets per second in an attack Q1, 2018

122% increase in peak attack rate between Q1, 2018 and Q1, 2019

3.7 Mpps Average packets per second in attacks in Q1, 2019

3.9 Mpps Average packets per second in attacks in Q1, 2018

Attack Vectors

In Q1, 2019, 77% of all attacks mitigated by Neustar used two or more vectors; none of the top attacks mitigated by Neustar used only a single vector.

NUMBER OF THREAT VECTORS PER ATTACK

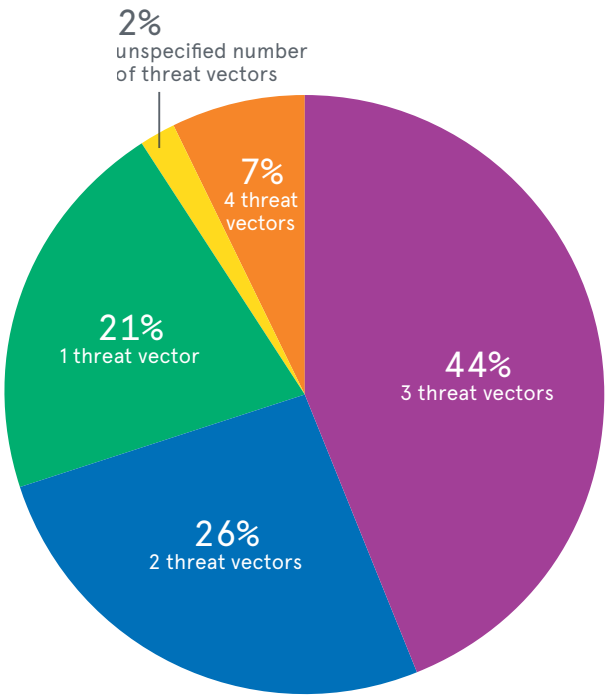

Figure 8 - Vectors per attack, Q1, 2019

DDOS ATTACK PRIMER

When most people think of DDoS attacks, the first thing that comes to mind are those volumetric Layer 3 or 4 exploits that use a high traffic load to saturate the bandwidth of the target site, such as the attacks on GitHub, Dyn, and others. But those are not the only flavor of DDoS exploit. Another attack type, called protocol attacks, are designed to overwhelm intermediate devices such as routers, firewalls, or load balancers within the target's network, rather than to saturate the link itself. Additionally, some hackers take aim still further into the target's network with an application layer attack aimed at a specific server, application, or Application Programming Interface (API).

Most troubling of all, many of today's attacks feature a combination of these three vectors. When considering the largest attacks of this quarter, we found that 77% used two or more vectors. That is roughly the same percentage observed in Q1, 2018.

In this section, we will take a closer look at these DDoS attack types and describe some of the examples that have passed through Neustar's SOC in the last quarter. Along the way, we will consider how to strike the right balance between the speed of AI and the ingenuity of human experience to counter such attacks.

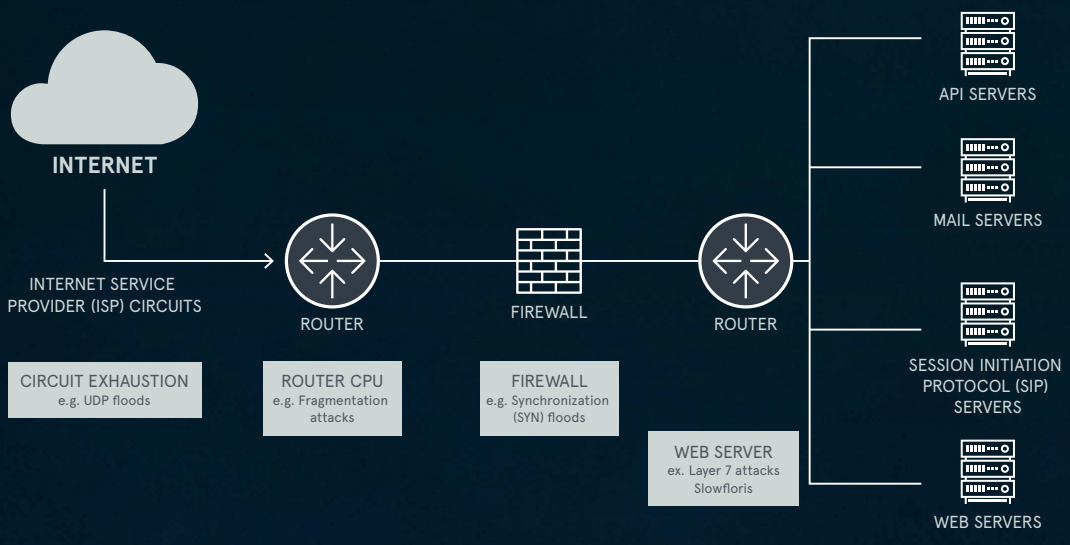

VOLUMETRIC OR LAYER 3 OR 4 ATTACKS

An Oldie and Still a Baddie

Volumetric attacks, which typically operate at Layer 3 or 4, are the best-known type of DDoS attack, with elements like UDP or Internet Control Message Protocol (ICMP) floods. The word "flood" is well chosen, as it can require quite a lot of traffic to fill today's enterprise network "pipe." In order to succeed with such an approach, the attacker must achieve several goals, including:

- Obscure attacker's identity
- "Disguise" the traffic
- Turn up the volume

Attackers often achieve the first two goals by taking advantage of inherent internetworking processes. The Internet was originally designed to support several types of connection requests. For example, when a request for the IP address of a hostname comes to a recursive DNS resolver, the resolver will look for and return that address to the requestor in the format specified. As you will see in the section about DNS amplification, the requested format can be significant to the efficacy of the attack.

This process can be misused by substituting the attacker's IP address with the address of the target, so that the information returned by the DNS resolver goes to the spoofed (target's) address. This is referred to as a reflection attack, and it accomplishes two things:

- The attacker's identity is obscured (at least at first glance)
- The traffic—which is, after all, perfectly legitimate—is disguised (at least for a while)

The next step is to turn up the volume. The best way to do that is by increasing both the overall number of requests and the amount of data sent per request. The easiest way for an attacker to increase the number of requests is by using a botnet, which can not only increase the number of queries, but also further remove the attacker from the process. Using a botnet no longer requires an attacker to create one; current reports show that botnets can be rented for less than $20 USD/day.[2] Many botnets now put poorly secured IoT devices into service, and the number of devices that could be compromised to take part in

a botnet is growing as well. Gartner forecasts that 14.2 billion connected "things" will be in use in 2019, and that the total will reach 25 billion by 2021.[3]

The next step in amplification is to craft a query that returns a large response. This was the secret behind last year's 1.3 Tbps attack against GitHub, which used requests to open memcached servers to generate responses that were over 51,000 times the size of the query. Because there are a limited (and hopefully shrinking) number of open memcached servers, and because the traffic, which runs on UDP port 11211, is relatively distinctive, this vector has become easier to spot. It is important to note that this vector is still being used; Neustar's SOC mitigated a number of memcached attacks in Q1, 2019.

Due to its size and port numbers, the memcached amplification attack was distinct, but that is not the case for many other amplification attacks. Certain types of legitimate DNS requests, for example, can result in very large responses; a favorite is the ANY request, which asks for all available information about a zone and can return records up to 80 times the size of the request. And in an ironic twist, a DNSSEC (Domain Name System Security Extension) request, designed to improve security, can result in an even larger response packet.

Botnets can be rented for less than $20 USD/day.

How to Protect Your DNS from DDoS Attacks (or Being Used in One!)

DNS servers themselves have become a compelling target for DDoS in recent years, as we saw with the attack on Dyn in 2016. If your company uses a managed DNS service, you need to know that your service has robust protections. Interestingly, Neustar originally built its DDoS mitigation service, which currently features over 10 Tbps scrubbing capacity, specifically to protect its own managed DNS service, UltraDNS.

But the threat to DNS service does not stop with the servers being protected. It is also important to ensure that DNS servers cannot be used in an amplification attack. Neustar's UltraDNS and UltraRecursive services monitor incoming traffic and have built-in rate limiting. If traffic flow increases above a set threshold, the system will be put into mitigation using Neustar's DDoS mitigation service.

"CARPET BOMBING"

A Twist on the Familiar

Volumetric attacks that feature reflection and amplification vectors are comparatively easy to recognize and mitigate. This is true in part because these attacks are typically aimed at a specific IP address or groups of addresses. A new take on this attack method showed up in the Neustar SOC early this quarter. Rather than aiming at a single IP address, this attack was instead directed at complete Classless Inter-Domain Routing (CIDR) blocks, or subnets.

Many organizations believe that working with a larger, Tier-1 ISP will protect them from DDoS attacks...This approach can be effective in the case of smaller attacks but can also provide a false sense of security.

Our ISP Just Dumped Us!

Many organizations believe that working with a larger, Tier-1 ISP will provide them sufficient DDoS attack protection; in fact, some ISPs offer DDoS attack mitigation as part of their service. This approach can be effective in the case of smaller attacks but can also provide a false sense of security. This was true in the case of one new Neustar customer who was hit by a CIDR block attack.

This company came to Neustar with an immediate issue: their Tier 1 ISP could not handle the volume of the attacks that were coming in and was starting to drop traffic. This would be alarming for any company, but for this customer, being offline translated to being out of business. The attacks themselves were interesting, as they were about 100 Gbps each, and moved throughout the customer's entire netblock, with up to six different hosts under attack from different vectors at the same time.

- The Neustar SOC logged over sixty different attacks in a forty-eight hour period
- Average attack duration in the first two days of the attack was 40 minutes; attacks ranged from four minutes to over 2.5 hours.
- The vast majority of the attacks featured three vectors; all contained more than two.

The severity of this attack, combined with the fact that being online was paramount to the company's business, resulted in the need to get this customer onboarded quickly. The Neustar SOC engineers were in constant communication with the customer throughout the process, handling mitigations around the clock.

PROTOCOL ATTACKS

Taking Advantage of the Rules (at Layer 3 or 4)

Unlike volumetric attacks, which make the target inaccessible by saturating bandwidth, protocol attacks work by consuming the processing power of the target itself, or that of critical infrastructure such as a router, firewall, web server, load balancer, or Virtual Private Network (VPN) concentrator between the Internet and the target. While these are technically similar to volumetric attacks, the difference is that the resource being exhausted is router CPU cycles or firewall state tables, rather than raw network bandwidth. This lowers the amount of traffic required to achieve denial of service, making such attacks harder to spot. These incursions can exploit vulnerabilities in the Layer 3 or 4 protocols themselves, which can easily be seen in the two examples on the following pages.

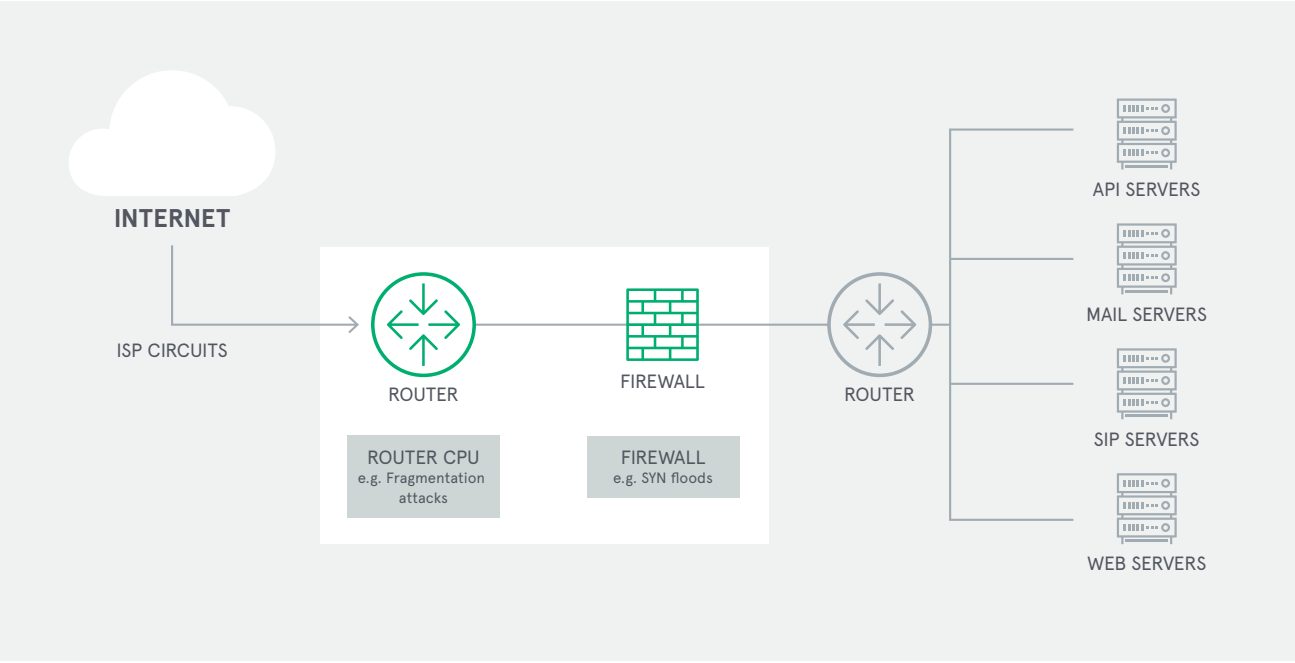


ROUTER

IP Fragmentation

IP fragmentation attacks are a good example of how traffic can become congested or stopped without link saturation. There are several ways in which IP packets can be used to fill up device CPUs, and they all take advantage of the fact that IP packets must often be fragmented in order to be transmitted. This is because every data transfer system has its own Maximum Transmission Unit (MTU), which signifies the largest packet that can be sent. One example of this attribute being used in an attack is to send an IP packet large enough to require fragmentation, but to add a bit that instructs the receiving router "Don't Fragment." In that case, the router will need to use resources to return an ICMP message to the sender noting that the destination is unreachable and including a code that notes that the "Don't Fragment" bit is set. This requires router CPU cycles and with enough packets sent, could degrade performance.

Another method takes advantage of the fact that the MTU is not limited to the payload portion of the transmission; each packet must also include all the information required to reassemble the original message when it is received. That reassembly happens at the end device, such as a web server. Because of the nature of Internet transmissions, packets are not always received in the order that they are sent, so the receiving device "holds" the packets it has received until it gets the complete set of packets indicated. Fragmentation attacks can include packets that indicate the start of a transmission but never complete it, or they can include malformed packets that make it impossible to reassemble the entire transmission. Such methods can disable or degrade the web server's performance. Some enterprises provide protection to the server by having reassembly done at an interim device, such as a router or proxy, which will be left holding the incomplete transmission until it times out. During this period, assets downstream from that interim device will not receive traffic normally.


FIREWALL

SYN Floods

SYN floods were one of the first DoS methods to be employed. They take advantage of the three-way handshake that starts every data transmission. The client begins the process by sending a packet with a synchronization (SYN) bit to the server to establish the sequence of communications between the two. In the next step, the server sends a packet that acknowledges receipt of the request (SYN/ACK). In the final step, the client sends a packet acknowledging the receipt of the acknowledgement (ACK), and transmissions begin.

It's not hard to see how this process could go awry if the client was an attacker. In the first use of SYN floods, attackers simply never sent the final ACK packet, leaving the server holding the port open until timeout. Many enterprises responded by putting a stateful firewall between the attacker and the target server and establishing rules that prohibited sending communications to the target until the full handshake was complete. Unfortunately, the result can be a situation in which the firewall's table fills up with incomplete handshakes. This creates a choke point for traffic.

Protocol layer attacks are often sized to avoid detection which allows them to wreak unknown damage. Because they do not attract immediate attention, intermediate attacks can go on for quite a long time before they are detected—if they are detected at all.

Protocol layer attacks are often sized to avoid detection which allows them to wreak unknown damage.

The Case of the Dead Load Balancer

It was a relatively quiet weekend when the Neustar SOC got a call from a puzzled customer. The customer's load balancer appeared to be dead. The appliance itself was fine, but it was not passing any traffic. The SOC could see traffic coming in, but nothing of sufficient volume to knock down these devices at first glance. A closer look, however, revealed an attack targeting the load balancers specifically, with a barely noticeable stream of encrypted traffic coming in on port 443. The attack featured malformed packets that were expertly crafted to take out this critical infrastructure, without saturating any circuits. This exploit could have caused extensive damage and required extensive skill to mitigate; in fact, this was technically a zero-day attack. Not only is this a good example of a protocol-based attack on an interim device, it serves to emphasize the value of communicating directly with SOC experts. The Neustar team worked directly with the customer to develop new, custom countermeasures. After this attack was effectively mitigated, the attacker switched to a series of other, well-known vectors, including SYN floods and Network Time Protocol (NTP) amplifications which were mitigated as well.

APPLICATION LAYER OR LAYER 7 ATTACKS

Lower Volume, Higher Value

Layer 7 attacks are directed at a specific server, application, or API, and are generally highly targeted. As a result, application layer attacks require even less overall traffic to accomplish their goals. In addition, many application layer exploits take advantage of how networking is designed to work, allowing such attacks to penetrate to the target server or app without being detected. As we saw with protocol attacks, the volume required to exhaust the resources of any single piece of hardware is much lower than that needed to completely saturate a link.

Slowloris

Although the code for Slowloris was originally released almost a decade ago, it remains a good example of how application layer DDoS attacks work. Slowloris took advantage of vulnerabilities in specific web servers, including Apache, many of which have been patched since the attack was released. As you would expect, Slowloris worked by exhausting the resources required for equipment to function; unlike some other attacks, however, this attack required very low bandwidth and can be used to take out only the targeted Hypertext Transfer Protocol (HTTP) service without affecting other processes on the server itself. That is because Slowloris established a genuine Transmission Control Protocol (TCP) connection, completing the three-way handshake that we talked about in the SYN flood example above. In the case of a SYN flood, the server (or, more likely, the firewall) consumes resources waiting for the client to complete the transaction.

Slowloris established a valid TCP connection, then initiated—but did not complete—a HTTP communication with the web server. The web server would then hold the connection open, waiting for the rest of the HTTP communication. The goal in this case was to fill up all available connections with partial HTTP requests, resetting the server's timeout counter repeatedly by adding bogus data just before the connection was dropped. If the targeted server was set up for high traffic, the attack could take time to succeed. There have since been patches to web servers and a variety of other fixes that have taken Slowloris off the radar, but it remains a good example of an attack in which the link was not saturated and other services worked perfectly, even as the performance of individual applications or servers was disabled or degraded.

Application layer attacks require even less overall traffic to accomplish their goals.

Credential Stuffing

A credential stuffing attack is not a denial of service in the classic sense. The purpose of a credential stuffing attack is to establish the validity of a credential set rather than to saturate a link or to disable/degrade the performance of a resource. A side effect of credential stuffing is often high traffic volume, however, which makes it look like a DDoS attack.

A credential stuffing attack begins when an attacker purchases a block of credentials, usually obtained as part of a data breach. These credentials are then fed into a botnet, which will attempt to use them to log in to a target. If the volume of credentials being tested is sufficiently high, it will sometimes be detected and mitigated in much the same way that a DDoS attack is handled.

Who Do You Call When Your Web Server Disappears?

That's the question one of Neustar's clients had in mind when calling into the SOC. The overall network traffic was not sufficient to put the client into mitigation, but when Neustar engineers took a closer look, the problem became obvious. The server itself was working; the problem was with the login page, which was being hit by a brute force credential stuffing attack. Upon further examination, engineers noticed that the majority of the traffic was coming from countries which did not match the profile of the customer's clients. A combination of geo-fencing and rate limiting got the customer back into business quickly.

A similar situation was encountered by an insurance customer who called in to say that the company's online rate quoting tool had gone down. This was puzzling since the customer had carefully scaled the application to handle traffic spikes as prospects went through the process of entering their information along with the details of what they were looking for. Neustar engineers found that the front-end application, which had been sized to handle the bulk of the inquiries, was indeed working well. The problem appeared further in the process, at the point where the app made a database call to get the quote itself. An attacker had created a way to skip ahead in the information flow and was overloading the database in an attempt to scrape the insurance company's quoted rates.

Neither of these cases represent classic DDoS scenarios and neither case featured sufficient traffic to automatically put the customer into mitigation, although both resulted in critical resources becoming unavailable. Once again, the ability to communicate directly with SOC engineers made the difference.

CONCLUSION

The first major DDoS attack was launched in 2000 by a 15-year-old Canadian known as "Mafiaboy." Not to be confused with the 1974 DoS attack, which targeted Programmed Logic for Automatic Teaching Operations (PLATO) terminals in a single classroom, the attack in 2000 was a sizable, distributed attack that targeted high-profile sites with enormous reach. It's important to note that even then this DDoS attack wasn't believed to be a particularly sophisticated piece of work. What this attack did succeed in, however, was to set into motion a constant see-saw effect that continues today, as organizations ratchet up their DDoS attack protections and hackers increase the sophistication of their attacks to outmaneuver those protections.

Today, DDoS attacks continue unabated. The latest attacks typically present more than a single vector, and morph over time, using a variety of ports and protocols. The trend of targeting subnets and CIDR blocks presents familiar threats in a worrisome new form.

Today's AI and machine learning technologies enable us to identify anomalous traffic and patterns, correlate data across systems, and perform behavior analytics on users and entities. Neustar's 10+ Tbps of scrubbing capacity and variety of offerings are world class, and we have more power than ever to defend against DDoS attacks. But it's important to remember our most powerful defense: people. None of these systems amount to very much without people who know how to deploy them, interpret their data, identify the existence and location of problems, and mitigate them. In every situation we saw this quarter, Neustar's SOC worked consultatively with the customer to determine the real issue and come up with an effective solution—usually against the clock.

With DDoS attacks on the rise, it's good to know that you have advanced protections. But if you're under attack, it's great to know that there's a team of experts you can call to explain what is happening and to help you eradicate the threat. The Neustar SOC team is ready.

GLOSSARY

ACK - Acknowledgement

AI – Artificial Intelligence

API – Application Programming Interface

CERT – Computer Emergency Response Team

CIDR - Classless Inter-Domain Routing

CPU – Central Processing Unit

DDoS – Distributed Denial of Service

DoS – Denial of Service

DNS – Domain Name System

DNSSEC – Domain Name System Security Extensions

FBI – Federal Bureau of Investigation

Gbps – Gigabits per second

HTTP - Hypertext Transfer Protocol

ICMP – Internet Control Message Protocol

IoT – Internet of Things

IP – Internet Protocol

ISP – Internet Service Provider

IT – Information Technology

M3AAWG - Messaging, Malware and Mobile Anti-Abuse Working Group

Mpps – Million packets per second

MTU - Maximum Transmission Unit

NISC - Neustar International Security Council

NTP – Network Time Protocol

PLATO - Programmed Logic for Automatic Teaching Operations

SOC – Security Operations Center

SYN - Synchronize

SYN/ACK - Synchronize-Acknowledgement

Tbps – Terabits per second

TCP - Transmission Control Protocol

UDP – User Datagram Protocol

VPN – Virtual Private Network

REFERENCES

[1]Statistics provided are exclusive of Neustar Security Operations Center-as-a-Service (SOC-as-a-Service) mitigations

[2]https://securityboulevard.com/2018/08/heres-how-anyone-with-20-can-hire-an-iot-botnet-to-blast-out-a-week-long-ddos-attack/

[3]https://www.gartner.com/en/newsroom/press-releases/2018-11-07-gartner-identifies-top-10-strategic-iot-technologies-and-trends

[4]https://thehackernews.com/2018/02/memcached-amplification-ddos.html

About Neustar.

Neustar, Inc. is a leading global information services provider driving the connected world forward with responsible identity resolution. As a company built on a foundation of Privacy by Design, Neustar is depended upon by the world's largest corporations to help grow, guard and guide their businesses with the most complete understanding of how to connect people, places and things. Neustar's unique, accurate and real-time identity system, continuously corroborated through billions of transactions, empowers critical decisions across our clients' enterprise needs.

More information is available at

www.home.neustar

WP-SS-168108-04.17.2019