


DDOS REPORT | 2019

Global DDoS Threat Landscape



Contents

01 Highlights

02 Overview

03 Network layer attacks

- Attack sizes

- Top attacked countries

- Most attacked industries

- Attack duration

- Attack persistence

- Attack vectors

- Multi-vector attacks

- FAQs

04 Application layer attacks

- Attack sizes

- Top attacked industries

- Attack persistence

- Source countries

- Destination countries

- FAQs

05 Methodology

- Terminology

06 What's next

07 About Imperva Research Labs

Highlights



Imperva recorded the largest ever network and application attacks

April saw a network layer DDoS attack that reached 580 million packets per second (PPS). Later in the year, a separate application layer attack lasted for 13 days and peaked at 292,000 RPS.




Persistence became the norm

The vast majority of network attacks were persistent and aimed at the same targets, a quarter of whom were hit 10 times or more.



An even split between single- and multi-vector attacks may be deceptive

Network attacks were split evenly between single- and multi-vector attacks. But this split might not have been exactly what it seemed, with many attacks accompanied by an unintended byproduct.



Asia was both the target and the source of many attacks

With India topping the list for the first time ever, the vast majority of countries targeted by network attacks were located in East Asia. We also saw that almost half of application layer attack requests originated in the Philippines and China.

Overview

2019 saw the biggest attack ever recorded but, overall, attacks were smaller, shorter, and more persistent.

For some time now, carrying out a DDoS attack has no longer required any technical skills. With a myriad of “stressing” or DDoS-for-Hire services available to carry out both network and application layer attacks, such skills can be replaced by little more than a couple of dollars. This was evident in the DDoS statistics we collected throughout the year.

Imperva did observe a few massive network attacks during 2019 - including the largest ever recorded, which reached 580 Mpps at its peak.¹ But these only accounted for a small fraction of the total number, most of which didn't exceed 50 Mpps/Gbps.

Most network attacks were short, too. 51% were less than 15 minutes in length - and, at between 100 and 1,000 RPS, application layer attacks were generally mid-sized.

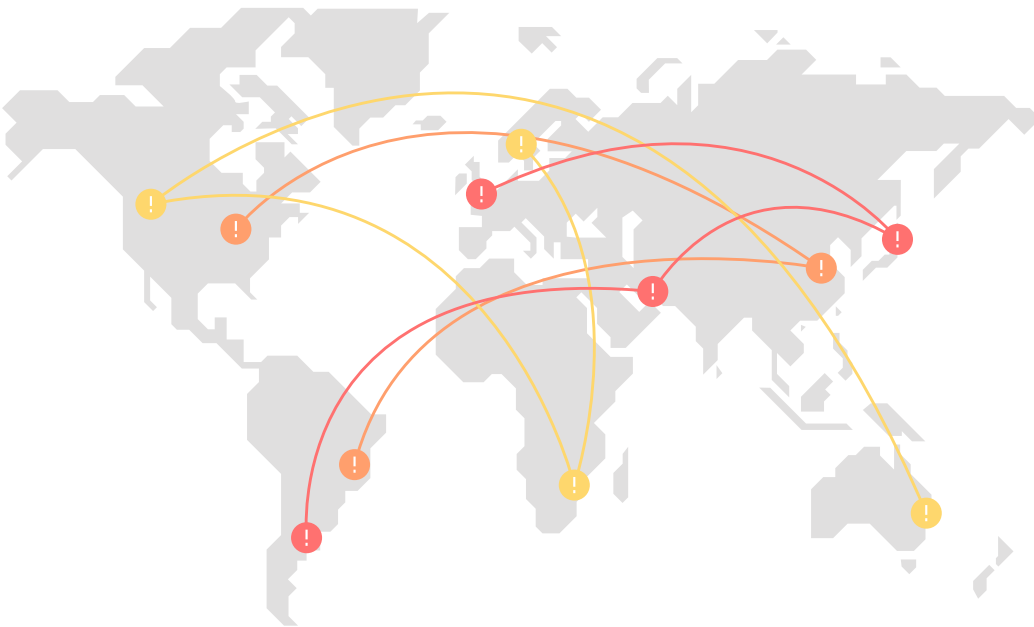
This makes sense. After all, most attackers would rather not waste time and resources on achieving their proof of impact or - in the case of targets protected by Imperva - lack of it.

They were persistent, though. We saw that attacks were either carried out in short streaks - two-thirds of targets were attacked up to five times - or were ultra-persistent, with a quarter of targets attacked 10 times or more. This might appear contradictory to the findings above, given that an attacker might not repeat an unsuccessful attack, but it can be largely dependent on the profile of the target.

Some businesses are more prone to multiple attacks by different attackers, while others are hit more randomly. And if a target has a particularly high profile, a single attacker may consider it worth trying several attack techniques before admitting defeat.

51%

of network attacks were less than 15 mins in length.



¹This DDoS Attack Unleashed the Most Packets Per Second Ever. Here's Why That's Important - Imperva blog - April 30, 2019.

Network layer attacks

1-3%

of network attacks hit an incredible 300+ Gbps and 200+ Mpps.

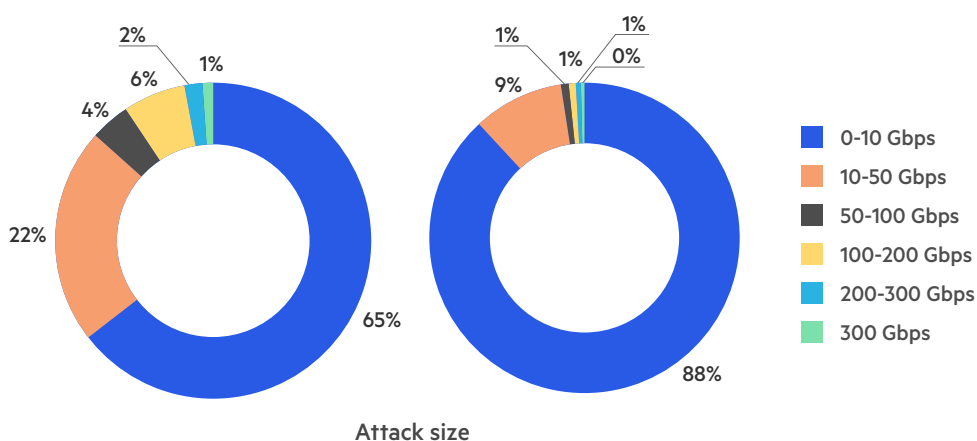
Attack sizes

As in previous years, the size and rate of most of the attacks that took place in 2019 were relatively small. The majority of attacks reached no more than 50 Gbps (87%) / 50 Mpps (97%), a fact that could be attributed to the use of “stresser” or DDoS-for-hire services, which typically offer unlimited quick and small attacks.

However, between almost 1% and 3% of attacks hit an incredible 300+ Gbps and 200+ Mpps respectively.

In April, for example, we saw an attack that peaked at 580 Mpps - the highest attack rate ever recorded - and reached 680 Gbps. But, impressive as this latter figure appears, attacks of this bandwidth have become increasingly common since the discovery of the Memcached DDoS amplification method in 2018.²

Fortunately, despite its size, our DDoS Protection service was able to quickly escalate and mitigate this particular attack.




Top attacked countries

Until now, India had never featured as one of the most attacked countries. During 2019, however, it topped the list. Hong Kong, which previously held the top spot, has now become the country in which most targets were attacked.


Interestingly, as shown in the graphs on the next page, the top four countries in each category are located in East Asia, making it the most dangerous region both in terms of the number of attacks and the likelihood of being attacked.

²Inside a New DDoS Amplification Attack Vector via Memcached Servers - Imperva blog - March 1, 2018

Attacked countries by number of attacks



Attacked countries by number of targets




Most attacked industries

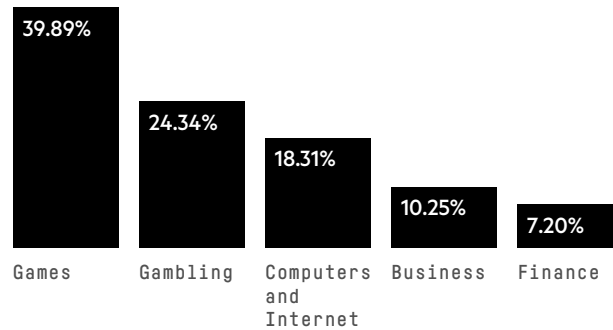
As has often been the case in the past, most DDoS attacks in 2019 were directed toward companies operating in the Gaming and Gambling sectors. When you consider that these are hugely competitive industries, both of which involve a high level of risk and where some players refuse to follow the rules, this is perhaps unsurprising.

Interestingly, businesses in the Computing and Internet sector were ranked third when measured in terms of both Gbps and Mpps. Organizations such as ISPs, web-hosting services, and domain providers are often prone to attack due to their high value - particularly those that host sites belonging to high-risk businesses such as gaming and gambling companies and crypto-currency traders.

Top attacked industries, according to number of attacks



Top attacked industries, according to number of targets



Attack duration



At Imperva, we consider the entire span of attacks with a burst-peace-burst pattern to be a single attack. This provides us with a fairly strict definition of when an attack ceases and when we can stop scrubbing traffic.

Even so, most of the attacks we recorded were relatively short, lasting less than an hour. Furthermore, within a one-hour distribution, the most common attacks were those that only lasted for 10 minutes or less.

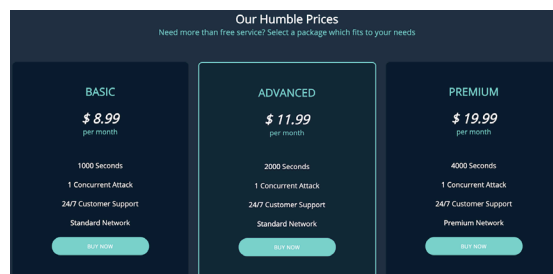
This, combined with our observation that most attacks were low in terms of both volume and rate, suggests that these short-term, weak attacks were most likely performed by DDoS-for-hire - or stresser - groups - as illustrated in the screenshots below - whose limited resources tend to be spread thin in order to service as many customers as possible.

The most common attacks lasted for less than **10 mins**

Attack duration



Examples of DDoS-for-Hire pricing, demonstrating the unlimited, short and small attacks offered by stressers




Attack persistence

Looking at the number of attacks faced by individual targets showed that most targets were attacked up to five times each. However, looking at the total number of attacks presented a very different picture.

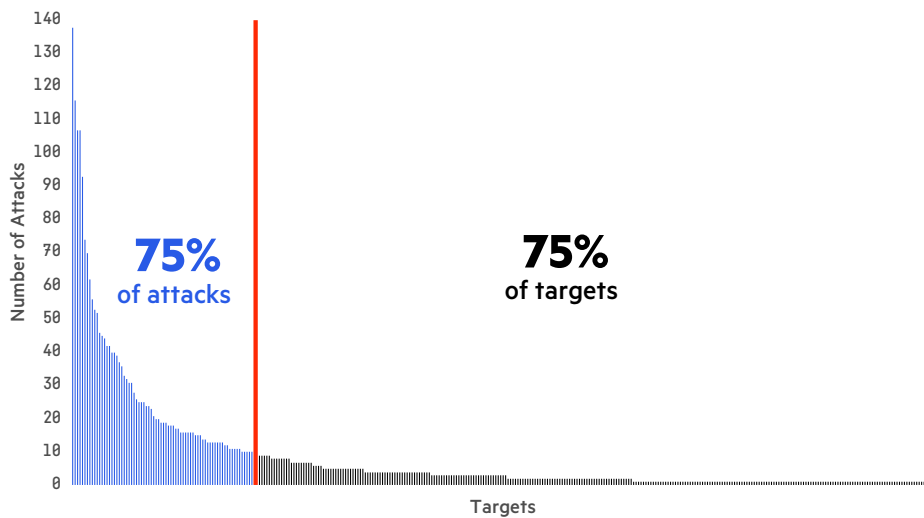
While most targets sustained only a handful of attacks, most of the attacks were aimed at a small amount of targets. If we assume that these attackers knew their targets would have DDoS protection in place, it begs the question of why they'd continue to launch attacks against them.

Given that most attacks lasted less than 10 minutes, one possible explanation is that the attackers were aiming to cause as much disruption as possible before the protection kicked in. It's worth noting that such a tactic would prove uniquely ineffective with Imperva's DDoS protection where, unlike with other solutions, the warm-up time is negligible, at just three seconds.

Attack persistence



Attacks on targets



The number of attacks each account sustained, out of a representative sample of accounts. 75% of the attacks are left of the red line.

Another possibility is that these were attempts to exhaust the targets' resources, forcing them to keep using top notch DDoS protection. Indeed, when you consider how cheap DDoS-for-hire can be - sometimes less than \$5 - this could be a viable technique for some malicious actors.


Economy	Deluxe	Ultimate
600 Seconds (10 Minutes)	1800 Seconds (30 Minutes)	3600 Seconds (60 Minutes)
500 Mbps	1500 Mbps	3000 Mbps
1 Month (25% Off)	1 Month (25% Off)	1 Month (25% Off)
\$5.00 \$3.75 USD (Save \$1.25 USD)	\$15.00 \$11.25 USD (Save \$3.75 USD)	\$30.00 \$22.50 USD (Save \$7.50 USD)
Add To Cart	Add To Cart	Add To Cart

Example of low prices for DDoS-for-Hire services

Attack vectors

UDP (User Data Protocol) was the most popular attack vector in 2019, probably because it's easy to spoof, is employed in virtually all amplification attacks, and is widely used in high-risk industries such as gambling and gaming. Following this came SYN floods, DNS Response, and TCP attacks, all of which are considered to be hard to mitigate. SSDP, Large SYN, and DNS attacks - which are more easily mitigated - were the least popular.

In the graph below, you can see that we split the attack vectors into the most popular. This means that, even though some vectors such as SSDP and NTP are performed over UDP, they're displayed separately. Splitting them out in this way highlights some important differences. For example, although technically the same, large SYN attacks weren't especially popular while SYN attacks were ranked second. This is likely due to the fact that, while large SYNs aren't legitimate packets, SYN's often are, making SYN floods harder to mitigate.




Multi-vector attacks

The diagram below shows an even split between single- and multi-vector attacks. These statistics may be deceiving, however, as many attacks came with an unintended byproduct - another attack vector entirely.

For example, SSDP might be accompanied by UDP packets from various different sources, while spoofed/amplification attacks were accompanied by ICMP packets from servers that had received an unexpected packet from the attacker or victim.

Furthermore, in some cases attackers would accidentally malfom a packet while in the process of randomizing it, thereby making it appear a different protocol than originally intended.



Network layer attacks - FAQs

What's the difference between network and application layer DDoS attacks?

The main difference between the two DDoS attack types is that they target different resources. A network attack attempts to clog network pipes, while an application layer attack seeks to deplete resources, such as CPU and RAM.

This translates into further differences in the ways these attacks are executed. It also means that mitigating each of these threats requires a significantly different set of security methods and skills.

In fact, outside of some superficial similarities, network and application layer attacks are two very different types of threats.

How are network layer DDoS attacks measured?

Network layer DDoS attacks are measured in Mpps (million packets per second) and Gbps (gigabits per second).

What's the difference between Mpps and Gbps?

Mpps measures the rate at which packets are delivered (a.k.a. forwarding rate) while Gbps measures the total load placed on a network (a.k.a. throughput).

From a mitigation point-of-view, it's important to be aware of both metrics, as they can each be bottlenecked by DDoS traffic.

For example, if your mitigation solution has the capacity to handle 80 Gbps and process packets at a rate of 10 Mpps, a 40 Gbps DDoS attack at a rate of 20 Mpps can still bring down your network, even if it doesn't surpass your total capacity.

Why are some countries targeted more than others?

Generally, for-profit DDoS perpetrators are interested in targeting wealthy countries with developed digital markets.

A lack of anti-cybercrime legislation or enforcement is also a contributing factor, as some for-profit and nonprofit attackers go after local targets. Finally, countries that serve likely-to-be-targeted industries, such as gambling, are more prone to attack.

Why are some industries targeted more than others?

Attacker motivation typically determines why a specific industry is frequently targeted by DDoS perpetrators.

Motivations can be broken down into the following categories:

- **Business competition** – In competitive industries, such as gambling, a DDoS attack can be used to take down a rival's website.
- **Extortion** – Certain industries, such as ecommerce, are very dependent on their online presence and are easy prey for perpetrators extorting money in exchange for keeping a specific website online.
- **Hactivism** – Hacktivists typically target political, media or corporate websites to protest against their actions.
- **Vandalism** – Cyber vandals, typically disgruntled users or random offenders, often attack gaming services or other high-profile targets.

What influences the duration of a network layer attack?

The length of a DDoS attack is largely determined by the resources at a perpetrator's disposal.

Shorter attacks are typically associated with DDoS-for-hire services (a.k.a. booters or stressers) that can be rented to launch short-lived attacks, usually lasting under 30 minutes. Longer attacks are almost always the work of more professional bad actors who use their own botnets to carry out persistent assaults.

Are short attacks a real threat?

Yes. There's no correlation between the length of an attack and the duration of a site's downtime. While a website (or web service) can be taken down in minutes, it usually takes hours for it to recover.

Additionally, a short attack might be part of a repeat assault, in which a target is hit with multiple short bursts. This method is commonly used to bypass mitigation solutions that rely on manual activation or are otherwise slow and cumbersome to deploy.

Why do perpetrators continue attacking a protected target?

There are a number of reasons to repeatedly attack a protected target. For example:

- It's common for perpetrators to change methods and try different attack vectors in an attempt to break through a site's defenses.
- The price of executing an attack is extremely low. If a first attempt fails, a perpetrator can try again (and again), even if their chances of success are slim.
- Repeat assaults are part of the MO of certain perpetrators, such as those executing pulse wave attacks.

What types of enterprises are more likely to be targeted by persistent attacks?

Generally speaking, large organizations are more likely to be the targets of persistent attacks, which are often initiated by competitors or skilled extortionists.

Why would a perpetrator use different attack vectors?

DDoS offenders will often switch between different attack payloads (i.e., different types of network packets) in an attempt to bypass a network's filtering mechanisms.

What's the difference between amplified and non-amplified attack vectors?

Amplified attacks vectors, such as DNS and NTP, are executed through a third party, such as an open DNS server. Conversely, non-amplified attacks are executed using a perpetrator's botnet.

Why do perpetrators launch multi-vector attacks?

In a multi-vector attack, different streams of payloads (network packets) are simultaneously sent to a target. This can help a perpetrator bypass an enterprise's security mechanisms, which are typically not equipped for complex filtering and might allow some of these streams to reach their target.

What do multi-vector attacks tell us about a perpetrator?

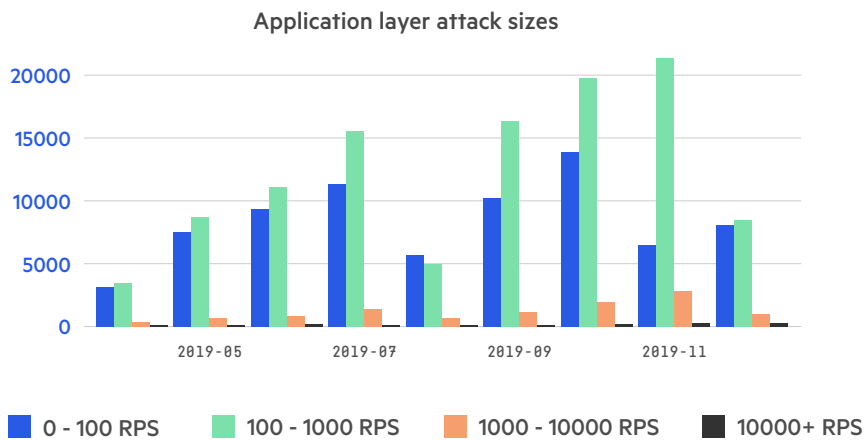

A multi-vector assault requires more resources and skill than a single-vector attack. The more sophisticated a bad actor is, the more likely they are to employ such techniques in their assaults.

Application layer attacks

Attack sizes

At no more than 1,000 requests per second (RPS), the vast majority of attacks mitigated by our WAF in 2019 were small to medium in size. The typical explanation for this is the wealth of “stresser” services offering attacks of this size for a relatively low budget, with higher throughput attacks usually only accessible to more expert attackers who herd their own botnets.


However, some of the exceptionally large attacks we observed actually beat the all-time record for attacks seen in our systems.



As the graph above illustrates, we saw a clear peak in the number of medium- large- and huge-sized attacks in November. This may have been related to the fiscal year end, where attackers - or those who've hired attackers - have leftover budget to spend. It's also worth noting that victims tend to be most susceptible at this time, when the year-end sales are on.

No significant correlation was found between the size and duration of attacks.

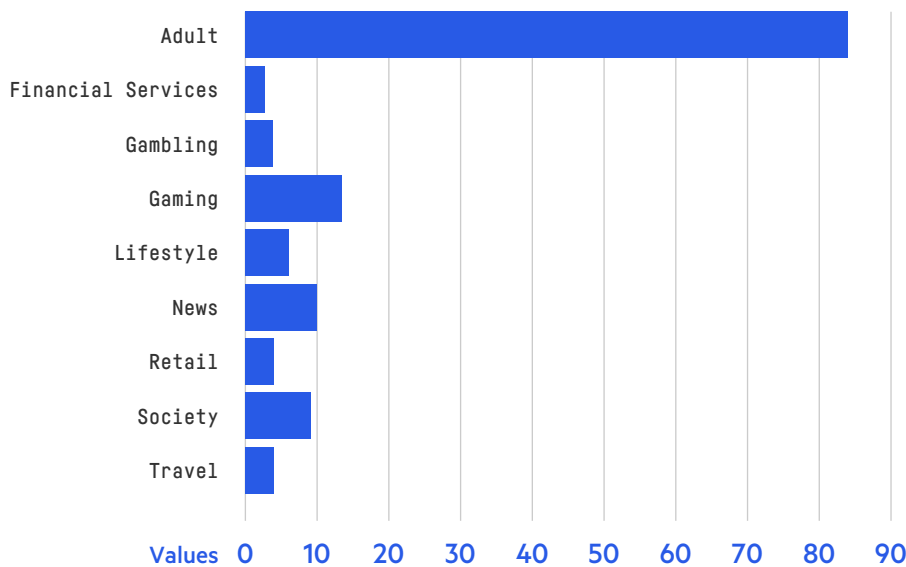
Attacks by duration



Top attacked industries

Perhaps unsurprisingly, given that it may be worth more to the US economy than Netflix, the NBA, or the NFL³, we found that the Adult industry was the most attacked by far. Every Adult site we tracked over the course of the year experienced an average 84 attacks which, between May and December, equates to 10.5 attacks per site each month.


Average number of attacks per site, by category



³Porn could be worth more to the US economy than Netflix, the NBA or the NFL - Quartz Index - June 21, 2018

Attack persistence


Interestingly, we saw no correlation between the types of industries attacked and the persistence of the attacks targeting them.



Source countries

The majority of attacks in 2019 originated in China and the Philippines. This doesn't suggest anything about the location of the attackers themselves, however, but rather the location of the machines used to carry out the attacks - these could be compromised devices, virtual or physical private servers owned (or, again, compromised) by attackers, cloud assets etc.

It's worth noting that, in the past, we observed how geographical bias in source countries was related to a vulnerability found in IoT devices widely used in those particular countries.⁴

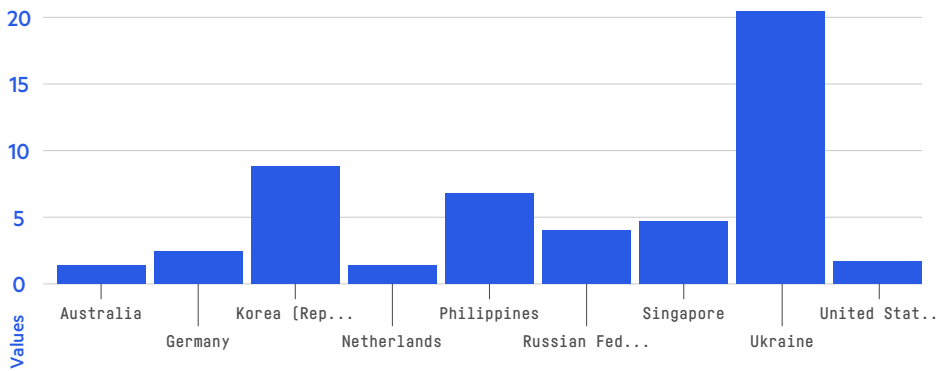


⁴New Variant of Mirai Embeds Itself in TalkTalk Home Routers - Imperva blog - December 7, 2016

Destination countries

The United States topped the list of attack targets for 2019. But it also topped the list for the number of Imperva customers. So, to avoid bias, we calculated the average number of attacks per site in each of the countries where our service is active.

As the graph below shows, this approach revealed that Ukraine was the most hit country, with an average of more than 20 DDoS incidents per site.



Application layer attacks - FAQs

How are application layer DDoS attacks measured?

Application layer DDoS attacks are measured in RPS (requests per second).

How many RPS does an attack need to take down a website?

An application layer attack's success depends on the amount of workload a single request can force on a target server. For example, a request that downloads an image file is far less resource-intensive than a request that initiates a string of API calls.

That said, many websites work on relatively low operating margins and can be taken offline by just a few dozen well-placed requests. There aren't many that can handle an additional 10,000 RPS, which is equal to 36 million requests an hour.

What influences the duration of an application layer attack?

Similar to network layer attacks, the duration of an application layer attack largely depends on the resources at a perpetrator's disposal. That said, application layer assaults are easier to execute and sustain, as even a sizable attack of several thousand RPS can be launched from a single computer.

Why do perpetrators continue attacking a protected target?

Similar to network layer attacks, perpetrators will repeatedly attack a protected target because it's so cheap - many offenders see no point in quitting, even if the chances of success are slim. Additionally, launching application layer attacks is easy and can even be done from a home PC or a very small amount of botnet devices.

Methodology

Our analysis of attacks is based on data from 3,643 network layer DDoS attacks on websites using Imperva services from January 1, 2019 through December 31, 2019, and on 42,390 application layer attacks on websites using Imperva services from May 1, 2019 through December 31, 2019,

Terminology

Application layer attack

An assault occurring on OSI layer 7. Its goal is to bring down a server by exhausting its processing resources (e.g., CPU or RAM) with a high number of requests. It's measured in requests per second (RPS)—the number of processing tasks initiated per second. Such attacks are executed by DDoS bots able to establish a TCP handshake to interact with a targeted application.

Botnet

A cluster of compromised, malware-infected devices remotely controlled by an offender. Device owners are unaware of their system participation.

DDoS attack

A persistent, distributed denial of service event against the same target (e.g., IP address or domain). A single attack is preceded by a quiet (attack free) period of at least three hours.

DDoS bot

A malicious software application (script) used by a perpetrator. So-called bad bots only come into play in application layer attacks, where a TCP connection is established. They typically masquerade as browsers (human visitors) or legitimate bots (e.g. search engine crawlers) to bypass security solutions.

Network layer attack

An assault against either the network or transport layers (OSI layers 3 and 4). Its goal is to cause network saturation by deleting much of the available bandwidth and CPU required to handle the packets. It's typically measured in gigabits per second (Gbps) and mega packets per second (Mpps), respectively referring to the amount of bandwidth it can consume, and the number of packets being sent per second.

Payload

In the context of this study, a payload is a packet type used in a network layer assault. It's fabricated by an attack script and can often be altered on the fly. In many cases, multiple payload types are used simultaneously during the course of a single event.

Stresser/Booter/DDoS-for-hire

A service that offers to launch DDoS attacks against targets for a fee. Although these services can be found using a simple Google search, they're mostly illegal. These services usually own a Botnet which they use to launch the attacks.

About Imperva Research Labs

Imperva Research Labs is a premier research organization for security analysis, vulnerability discovery and compliance expertise. The organization provides round-the-clock research into the latest security vulnerabilities and is comprised of some of the world's leading experts in data and application security. Imperva Research Labs combines extensive lab work with hands-on testing in real world environments to ensure that Imperva's products, through advanced data and application security technology, deliver up-to-date threat protection and unparalleled compliance automation. Incorporating exceptional insight, Imperva Research Labs publishes reports on a quarterly basis like the Global DDoS Threat Landscape Report and Bad Bot Report that provide insight and guidance on the latest security threats and how to mitigate them.

WHAT'S NEXT

For more information about Imperva DDoS protection services, visit imperva.com/products/ddos-protection-services/

And, visit the DDoS Mitigation section of our Resource Library for more DDoS-related content.



Imperva is an analyst-recognized, cybersecurity leader championing the fight to secure data and applications wherever they reside.

+1 [866] 926-4678
imperva.com