imperva

WHITEPAPER

How Imperva DDoS Protects at Scale

Introduction

Today's customers need the ability to access your site and applications, always. No delays, no wait time, no issues, regardless of where they are in the world. But attackers are constantly honing their approaches, finding ways to take services offline via DDoS attacks, using volumetric, protocol or application attacks to divert attention while they implant malware, steal credentials or exfiltrate sensitive data. Only Imperva can ensure that customers are not impeded by slow-loading sites, unavailable applications or any of the other side-effects of a DDoS attack.

There are three key issues when it comes to protecting at scale: time to mitigation, latency and capacity. As discussed in this report, Imperva optimizes all three by making use of its global network of full-stack PoPs to provide the first line of defense, constant availability and comprehensive visibility. The only company to offer a 3-second DDoS attack end-to-end mitigation SLA, Imperva ensures your customers have fast access without additional latency, potentially unnecessary CAPTCHA screens or slow load times.

Fastest time to mitigation

Time to mitigation (TTM) is one of the most important measures in a DDoS attack: it's the time required to start scrubbing traffic after the first DDoS attack packets hit your system. Longer TTM means more opportunity for the attack to succeed; with the average cost of DDoS downtime reaching $40,000 per hour, every second counts.

Many anti-DDoS vendors claim fast TTM, but few actually quantify it. Imperva even improved its already industry-leading TTM from 10 seconds to a groundbreaking three seconds, backed up by an SLA.

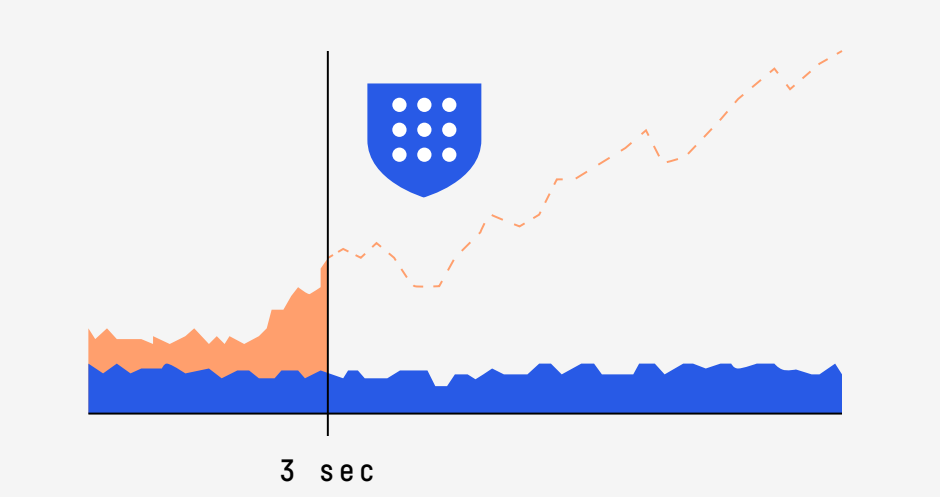

Figure 1: Imperva's industry-leading 3-second DDoS mitigation SLA

BEWARE OF HOLLOW CLAIMS

It's common to see big claims from vendors, but it's worthwhile taking a closer look. For example, one company claims zero-second mitigation but it's a boast that comes with a big caveat: it only applies when the attack is based on known traffic profiles and matches historical signatures or well-known traffic patterns.

Most recent DDoS attacks – in fact all of the largest and most damaging ones – wouldn't even be covered, because they were caused by new, previously unknown attacks. In addition, these approaches are based on proactive manual policy implementation that requires thousands of security operations center resources, and it cannot scale. Imperva's 3-second DDoS mitigation SLA covers attacks, even those that exploit a vulnerability for which no patch has yet been released (zero-day exploits). All done automatically.

This includes the time required to identify that an attack is taking place before traffic scrubbing can begin. The Imperva 3-second SLA applies to all types of attacks, at all layers of the network. Imperva DDoS Protection for Websites, IPs and Networks are always-on services, detecting attacks of any type automatically without manual intervention. Imperva quickly detects the attacks and creates directive for scrubbing in just seconds. Whether volumetric Layer 3 and 4 attacks measured in Mbps (megabits per second) or packet rate attacks measured in Mpps (thousands of packets per second), or Layer 7 attacks measured in RPS (requests per second) that attempt to crash or hang web applications, Imperva employs a multi-faceted approach that can handle the largest attacks without impacting the user experience.

Lowest latency

With DDoS, speed is the name of the game. Users expect minimal latency and have become accustomed to sites loading almost instantly. Studies show that for every second of page load time, conversion rates drop by 7%. Avoiding website downtime by mitigating DDoS attacks requires a smart, global network that can respond immediately to any threat. Imperva deploys Super PoPs inside Internet connectivity hotspots to provide on-demand DDoS mitigation power with minimal latency. Customers in need of always-on DDoS protection can reroute all website traffic through Imperva's network, with its content delivery network (CDN). This provides scalability when it is needed to absorb volumetric attacks, while simultaneously minimizing latency. The Imperva mesh network of data centers enables round-trip times of less than 50 milliseconds to more than 90% of the world, ensuring smooth performance even when servers are under attack.

Figure 2: DDoS scrubbing, WAF, bot protection, caching and load balancing supported at every location.

Other anti-DDoS vendors have smaller networks, fewer PoPs, and a mix of scrubbing and non-scrubbing PoPs. Although some have data centers around the world, many of them do not perform any scrubbing. Some use third-party CDNs, so there is no integration and latency is greatly increased. This means that when a DDoS attack happens, it might take time to reach the nearest scrubbing center, and some vendors require their customers to first notify them that an attack is underway, further increasing the latency.

Imperva customers benefit from the software-defined network all the time, not just during DDoS attacks. They get the benefits of content caching and dynamic content acceleration in Imperva's CDN for widespread, positive performance improvement. With PoPs strategically located to meet user demands, a broad peering footprint, and an entire network tuned for DDoS mitigation, customers experience high-quality connections. This translates into an even better experience for clients' end users and increased conversion rates for e-commerce websites. With reserve capacity for the biggest attacks, amplified by a mesh network, Imperva offers all customers the aggregate capacity of the entire network.

Highest capacity

Network capacity is the third key element to evaluate when looking for an effective DDoS mitigation solution. One measure of capacity is throughput, or the amount of bandwidth that can travel through the connection. Measured in Gbps (gigabits per second) or Tbps (terabits per second), network capacity determines how scalable the response to a DDoS attack can be. At Imperva, 44 DDoS-resilient data centers with more than 6 Tbps of scrubbing capacity, can process 65 billion packets per second.

Another measure of capacity is processing capability, or the number of network packets that can be processed. This is measured in forwarding rates, in Mpps. With many attacks peaking at more than 50 Mpps, and some as high as 580 Mpps, your DDoS solution needs to be able to continue processing without affecting legitimate users. Many on-premises appliances and appliances in scrubbing centers were built to withstand high throughput rates, but not similarly high forwarding rates. In fact they are often useless during a volumetric attack, since the traffic won't even reach them: it will be blackholed by the service provider. Nor can on-premises appliances offer any protection for cloud-hosted applications.

While assaults with high rates of Gbps are the norm, it is worthwhile ensuring your DDoS solution provider can handle 650+ Mpps as well. Imperva can handle even massive attacks – see below.

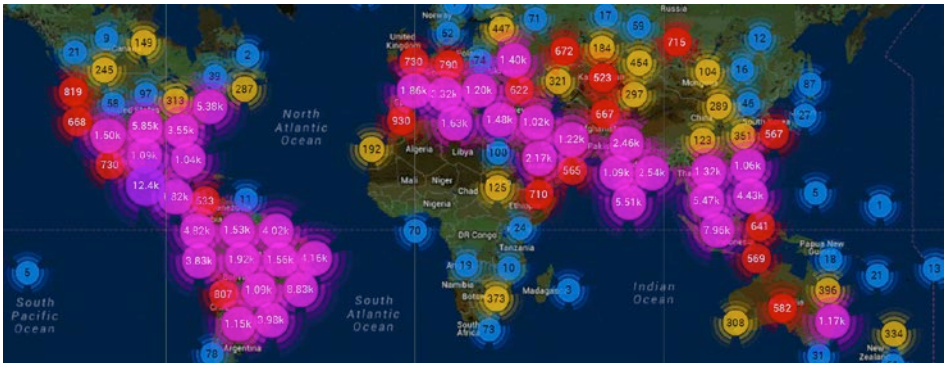

Figure 3: Imperva mitigates a massive HTTP flood: 690,000,000 DDoS requests from 180,000 botnet IPs

Imperva in action

Some examples of DDoS attacks mitigated by Imperva – in less than three seconds:

Mega Mpps attacks

In a single week in May 2019, Imperva mitigated nine of the largest-ever DDoS attacks against our customers. All were more than 500 Mpps, with the largest being a record 652 million packets per second. These were not small attacks, either – they peaked at 713 Gbps. Imperva was able to automatically halt all these attacks, without human intervention, and with no impact on the network or infrastructure. This is due to a new service called the SD-NOC (Software-Defined Network Operations Center) which automatically shifts the attack traffic toward the most suitable scrubbing centers for 100% attack mitigation, optimal performance and latency. Rather than waiting for SOC teams to constantly perform manual monitoring, determine an attack was taking place and notify the customer, Imperva was able to take action automatically.

Largest L7 / brute force attacks

In July 2019, one of our customers saw a DDoS attack led by a botnet that coordinated more than 400,000 different IPs, lasted 13 days and directed a peak flow of 292,000 requests per second. Imperva mitigated the attack during its entire 13-day span –and the customer suffered absolutely no downtime. While mitigating the attack, however, our security team investigated the attack and found that the first stage of the attack was to compromise IoT devices; using default credentials botnet attack then used the DDoS attack to upload malware to the device and connect it to a Command and Control server to receive commands. In fact, our recently launched Account Takeover Protection helps protect against these types of attacks including brute force and credential-stuffing attacks.

How can Imperva do this?

A number of factors work together to ensure the industry's fastest DDoS mitigation.

- First, Imperva provides a variety of DDoS solutions, for websites, networks, DNS servers and individual IPs. This means protection for all aspects of the business, including cloud-hosted workloads where you don't control an entire subnet.

- Second, Imperva's sheer processing power is unrivalled. Its Behemoth 2 scrubbers are fully automated mitigation appliances capable of sub-second detection and mitigation. Located within the Imperva global network of full-stack PoPs, they stand ready to mitigate even the most aggressive attack.

- Third, Imperva benefits from a network of real-time synchronization servers; these prioritize traffic data and instantly alert all other servers when an attack is detected.

- Fourth, and key to Imperva's speed and efficiency in mitigation, the ability to see all the attacks that are taking place, in context. Rather than reacting to a single attack with a multitude of alerts, Imperva uncovers situations where the attacker is using multiple vectors to carry out a sophisticated attack – perhaps using a DDoS attack as a smokescreen for stealing credentials, implanting malware or escalating privileges. Imperva Attack Analytics spares the security team from being overwhelmed by a barrage of alerts with high false-positive rates, instead presenting a single actionable narrative so they can concentrate on the real threat.

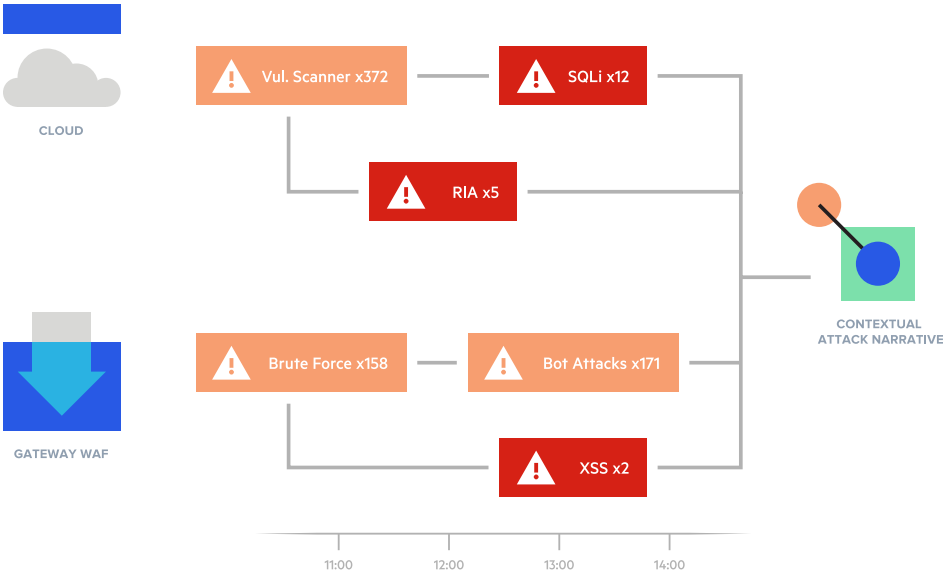

Figure 4: Imperva Attack Analytics distills the big picture into an actionable narrative

Conclusion

When it comes to customers, it's all about making the experience as satisfying as possible. Of course, this means stopping DDoS attacks in their tracks, before they can slow down or deny access to websites or applications. But it also means ensuring that content is delivered as quickly as possible to legitimate users and shielding them from potentially unnecessary challenges like CAPTCHA and others. But we can't do this at the expense of the business team who are tasked with ensuring maximum uptime and protecting the corporate brand and assets. Imperva has managed to strike the right balance – protecting customers while simultaneously making life a little easier for the security and web teams.

ABOUT IMPERVA APPLICATION SECURITY

Imperva Application Security mitigates risk for your business with full stack defense-in-depth, providing protection wherever you choose to deploy - in the cloud, on-premises, or via a hybrid model.

Imperva offers advanced analytics to quickly identify the threats that matter, Web Application Firewall (WAF) solutions which block the most critical web application security risks, DDoS protection with a 3-second mitigation SLA, API Security that integrates with leading API management vendors, Bot Management for protection against all OWASP automated threats, Runtime Application Self-Protection (RASP) for security by default against known and zero-day vulnerabilities, and a developer-friendly Content Delivery Network (CDN) for the utmost performance.

Through FlexProtect, our unique licensing model, you can deploy Imperva Application Security solutions how and when you need them. FlexProtect helps protect your applications wherever they live — in the cloud, on-premises or in a hybrid configuration.

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

+1 (866) 926-4678
imperva.com

imperva