The CISO's Guide to Beating Web DDoS Attacks

The first half of 2023 has seen an unparalleled rise in DDoS attack activity. DDoS attacks have grown in frequency, size, and sophistication, and the audacity of attackers seems only to be growing.

In particular, web DDoS attacks, have emerged as a particularly devastating threat to organizations across every vertical and geography. A new generation of attack tools is enabling attackers to launch new, and disruptive web DDoS tsunami attacks with ease. As a result, organizations are faced with the urgent need to enhance their defenses against these sophisticated and potent attacks.

The purpose of this guide is to outline some of the key trends in the modern DDoS threat landscape, explain why web DDoS attacks are such a threat, and show how organizations can protect themselves against such attacks.

Part I: The Modern DDoS Threat Landscape

DDoS attack activity has reached unprecedented heights over the course of 2022 and the first half of 2023. According to Radware's Threat Hub, which tracks attack activity across its global network, 2022 has seen a year-over-year (YoY) increase of 152% in blocked DDoS events compared to 2021. Similarly, the total blocked volume of DDoS attacks rose by 32% YoY between 2021 and 2022, and the largest DDoS attack observed by Radware in 2022 was 1.46 Tbps in size, a 2.8x increase compared to the largest attack observed in 2021.

Moreover, attackers have grown in audacity and scale of the targets attacked. Over the course of the past 18 months, attackers have launched targeted DDoS attacked against public sector and healthcare organizations, civilian airports and air traffic controls, educational institutions, Cloud and service providers, and more. This has led to a far greater impact of DDoS attacks than ever before.

This rise in DDoS attack size, frequency and sophistication indicates a fundamental shift in the threat landscape of DDoS attacks. This change has been governed by a convergence of three main factors.

Three Trends Reshaping the DDoS Threat Landscape

Factor #1: Rise of State Actors

Perhaps the single biggest change in the attack landscape is that political motivation has superseded financial motivation as the primary driver for DDoS attacks.

This shift began on February 24, 2022: the day when Russia invaded Ukraine.

Unlike previous conflicts, where cyber activity was limited to the periphery of the conflict, for the first time cyberattack activity was lock-in-step with the ground action. In fact, the first wave of cyberattacks began even before the main invasion, as part of a series of preliminary attacks and special operations.

Since then, cyberattacks have been an essential part of Russian war plans and are executed either directly by the Russian military, or through a myriad of independent state-supported groups such as Killnet, Passion, Zarya, NoName057(16), and others. These groups operate independently, and continuously form, re-form, merge and splinter with each other.

As the war has drawn-out, the activities of these groups have extended to target not only Ukraine, but also Ukraine's key allies in Europe and North America.

The shift from criminally-motivated hacker rings to state-sponsored hacktivist groups – even if in some cases they are the same – is significant in three respects:

- **Resources and capabilities:** State-backed organizations have more resources to create bigger and more sophisticated attack tools.
- **Profile of targeted organizations:** Organizations and networks which previously may not have been targeted, could now be targets as a result of political motivation.
- **Remedies against attackers:** Backing by the state effectively provides hacker groups with immunity. They will not be arrested, prosecuted or extradited for their activities.

This trend, however, has not been confined to Russia. Over time, the rise of Russian state-sponsored hacktivist groups has given rise to a new wave of similar ‘hacktivist’ groups, unrelated to Russia or the war in Ukraine, that are more religiously motivated, such as Anonymous Sudan, Team Insane PK, Eagle Cyber, Mysterious Team and others. As a result, targets in the US, Israel, India, Australia, Sweden, and other countries have been targeted by political and religious hacktivist groups.

Factor #2: Attacks Grow in Size and Complexity

Another key factor has been the growth in the attack size and complexity.

Typically, DDoS attacks are based on globally-distributed botnets, which use known attack tools, and are re-used by different groups across botnets. One of the effects of the rise in state-sponsored hacking groups mentioned above, is the development of new attack tools, which not only lead to larger attacks, but also to more sophisticated ones.

One of the ways in which this sophistication manifests itself is in the development of new attack vectors, and usage of multiple attack vectors within a single attack. Instead of prolonged barrages by a single attack vector, attackers increasingly mix-it-up with short, small bursts of one vector, before switching to a different one, and doing so continuously.

While such attacks have been known for several years, they have now become commonplace. According to Radware data, attacks between 1-10 Gbps, on average, leveraged two dissimilar attack vectors per attack, attacks between 10-100 Gbps leveraged, on average, four attack vectors, and attacks over 100 Gbps utilized on average more than nine different vectors.

Factor #3: Attacks Shift to Application Layer

The third significant change is the shift of DDoS attacks to the application layer. While this trend has been ongoing for several years, recent attack waves have brought it to new heights.

The first major botnet to make extensive use of application-layer HTTP/S attacks was the Mirai botnet in 2017. However, new botnets developed in aftermath of the conflict in Ukraine have greatly enhanced those capabilities.

Moreover, they also use make use of newer web DDoS attack tools such as Blood, MHDDoS, Saphyra, and others, which not only offer multiple web DDoS attack vectors, but also provide mitigation bypass techniques such as header randomization, CAPTCHA solving, IP spoofing, cookie harvesting, and more. This is making web DDoS attack particularly difficult to mitigate for traditional DDoS mitigation tools.

As more and more internet services shift to web applications, and application-layer attacks are getting more sophisticated, web DDoS attacks are quickly becoming the premier approach for DDoS attacks.

Part II: What are Web DDoS Attacks and Why They are Hard to Mitigate

The result of the trends redefining the DDoS threat landscape is the emergence of web DDoS attacks the premier – and most potent – vector for modern DDoS attacks.

A Web DDoS attack leverages the application-layer HTTP or HTTPS protocols to launch a DDoS attack. Under such an attack, attackers direct large amounts of HTTP requests towards a web application in order to overload target servers with requests. In this respect, this attack is similar to network-level DDoS flood attacks, except that this one operates at the application layer, using HTTP GET or POST requests.

However, as majority of the internet traffic today is encrypted, most HTTP flood attacks today are, in fact, HTTPS floods. This adds a layer of complexity for the mitigation of such attacks, since DDoS defenses usually cannot inspect the contents of the HTTPS requests without fully decrypting all traffic.

It should be noted that web DDoS attacks are not a new phenomenon, but have been known for many years. However, the combination of more powerful state actors, more sophisticated attack tools, and a general migration to application-layer DDoS attacks has led to a surge in this type of attack. Indeed, Radware research shows that in 2022, 81.4% of L7 DDoS attacks were HTTP/S floods.

To illustrate, here are a few examples of last Web DDoS attacks that Radware automatically and accurately mitigated:

An EMEA-based governmental agency was hit by repeated waves of attacks by a hacktivist group. The attack was characterized by large attacks waves with high request-per-second count, peaking at nearly 800,000 requests-per-second (RPS) over a prolonged duration of several hours.

Figure 1: Attack wave peaking at nearly 800k RPS, across many hours

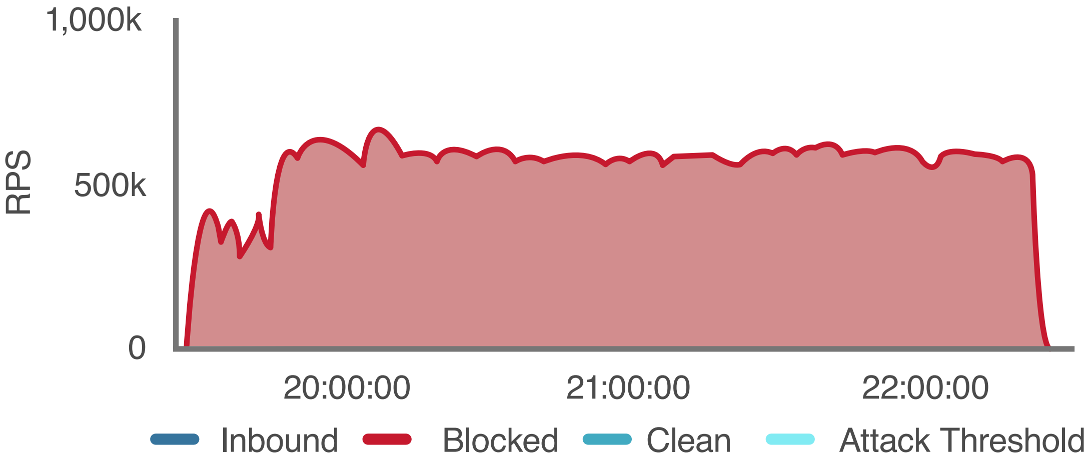

Another example is a **DDoS campaign against major Online Retail Business**, which was targeted by a hacktivist group. This organization was attacked by a complex web DDoS attack, with massive attack waves, peaking at nearly 2 million Requests per Second (RPS). The campaign lasted for nearly two weeks, with multiple attack waves. Some of the attack waves lasted for hours, reaching over 10 billion requests in aggregate.

Figure 2: Attack wave peaking at nearly 2 million RPS

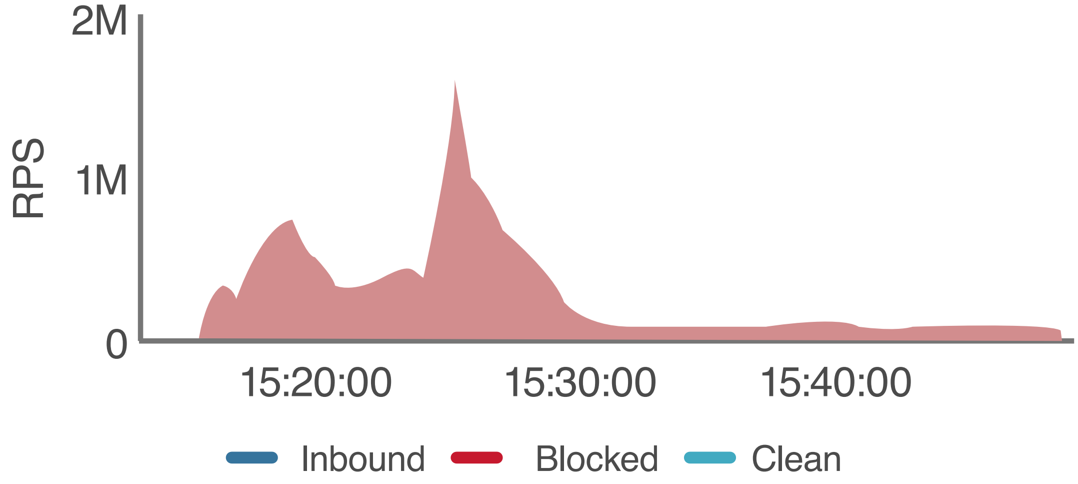

Why Web DDoS Attacks Are So Difficult to Mitigate?

The immense challenges of mitigating web DDoS attacks, and the damages they incur, raise the fundamental question of why Web DDoS attacks are so hard to mitigate?

The reasons for this are attributed to the specific characteristics of HTTPS-based web traffic, and their impact on the ability of traditional DDoS defenses to identify and mitigate application-layer attacks. .

Here are some of the key reasons that make web DDoS attacks particularly hard to mitigate:

- **Asymmetric Processing Requirements:** SSL/TLS are asymmetrical protocols which require more computing resources from the server than from the host initiating the request. This means attackers can generate massive attacks with a relatively small number of requests.
- **Payload is Encrypted:** Most of the web traffic today is encrypted using the HTTP Secure (HTTPS) protocol. This means that, by default, the payload of attack traffic is encrypted, making it immune to inspection by traditional network-layer DDoS defenses.

- **Attacks the Application Logic:** Application-layer attacks often appear, at first glance, similar to legitimate requests. This means that accurately mitigating application-layer attacks requires a deep understanding of the target application, the context of the request, and the ability to detect abnormal behavior indicative of an attack.
- **New Advanced Tools:** Attackers are continuously evolving the tools at their disposal, and utilize new advanced techniques to carry out DDoS attacks. New generations of attack tools are using methods such as attack randomization, IP spoofing, cookie harvesting and more to evade DDoS protections.

As most internet traffic nowadays is encrypted and based on HTTPS, attackers can easily hide behind the built-in encryption and logic of the application layer. This makes web DDoS attacks to be very attractive for potential offenders to use this approach and launch massive flood attacks.

New Attack Tools Are Built to Bypass Traditional Defenses

The landscape of cyber-attacks continues to evolve, with attackers developing tools and strategies that can outsmart traditional defense mechanisms. These tactics pose significant challenges for cybersecurity defenses, as they make it difficult to differentiate between legitimate and malicious traffic, trace the origin of attacks, and identify and block malicious activities effectively. Here are some of the tools and tactics:

- **Attack Vector Randomization:** The HTTPS Flood also use randomized attack vectors so no two requests are identical. By altering these arguments with each request, they aim to evade signature-based detection mechanisms that rely on static patterns or known attack signatures.
- **Randomized Request Arguments:** Attackers inject random arguments, such as query parameters, into their HTTPS requests. This makes it challenging for DDoS defenses to differentiate between legitimate and malicious traffic. Moreover, the dynamic structure of the request means it will not be cached by the customer’s CDN, but go directly to the origin server.
- **IP / XFF Spoofing:** Attackers employ IP spoofing techniques to mask the source of their attack traffic. They forge or spoof the source IP addresses of their requests, making it difficult to identify the actual origin of the attack. Additionally, attackers may manipulate the X-Forwarded-For (XFF) headers, which contain client IP information, to further obfuscate the source of attack.

- **Harvesting Application Cookies Using Headless Browsers:** In some cases, attackers utilize headless browsers or automated scripts to mimic human behavior. They use these headless browsers to navigate pages, simulate form submissions, and harvest legitimate user cookies. By obtaining valid session cookies, attackers can bypass certain security measures and make their requests appear authentic. This method allows attackers to bypass protections that rely on client-side cookies for identification.
- **Anonymous Proxies:** Attackers often leverage anonymous proxies or Tor networks to route their attack traffic through multiple intermediate nodes, making it difficult to trace the origin of the attacks. By hiding behind proxy services, attackers obfuscate their real IP, making it challenging for DDoS defenses to identify and block the malicious traffic effectively.

Beyond the technological novelty of these attack tools, the problem with these new randomized attack techniques is that traditional DDoS protection techniques were not designed to deal with such complex attacks. In essence, these attack techniques render traditional DDoS protection methods obsolete.

Traditional DDoS Mitigation Tools Cannot Protect Against New Attacks

Traditional DDoS mitigation tools are struggling to defend against the ever-evolving landscape of new attacks. The rapidly changing attack techniques render these tools ineffective, highlighting the need for advanced and adaptive defense mechanisms to combat the emerging threats:

- **Access Control Lists (ACLs)** are useful when the attack is coming from a single source or a small number of known sources/IPs. However, if requests are distributed across large numbers of distinct IP addresses (real or spoofed), blocking those sources using fixed access control lists is impossible due to the large number of sources.
- **Geo-blocking** can be effective if most attacks are coming from a specific country or region of the world. However, if attacks are based on large botnets (or spoofed IPs) spanning multiple countries and regions, this type of filtering is ineffective because it would mean blocking out large portions of the world.
- **Rate limits** are popular with some DDoS mitigation services as a means of limiting the number of requests - good or bad - reaching the application. A significant downside of this approach is that it does not distinguish between good and bad connections. This means that legitimate transactions are also blocked, leading to high rates of false positives.

- **Static Signatures** are used against known attacks. However, since new attack tools use built-in randomization techniques, each attack looks different, and static signatures are not enough to identify the attack patterns or distinguish between malicious and legitimate traffic.
- **CAPTCHA** is a tool often used to distinguish between human users and automated bots, which are frequently used to launch DDoS attacks. However, CAPTCHA is a very ‘noisy’ tool, which interferes with the user experience and causes delays. Moreover, it is an increasingly ineffective tool, as new generations of bots learn to mimic user behavior and bypass CAPTCHA challenges.

As traditional DDoS mitigation tools fail in the face of new and newer attack techniques, a new approach is needed to tackle this challenge; one that does not rely on a one-size-fits all approach to DDoS mitigation, but rather is focused specifically on the complex and dynamic characteristics of the new generation of attack.

Behavioral-based Protection Adapts to Dynamic Attack Patterns

The mirror image of traditional brute force methods such as geo-blocking and rate limiting is the behavioral-based approach. While requiring a higher degree of technological know-how, it allows for more granular detection and mitigation of malicious traffic, without impacting legitimate connections.

The behavioral approach is based on a ‘positive’ security model, utilizing advanced traffic learning to understand what normal user behavior looks like, and create a baseline for legitimate traffic patterns. Thus, any traffic which falls outside of the parameters of normal legitimate traffic will be automatically blocked, without impacting normal users.

This approach has several key advantages to it:

- **Adapting to Dynamic Attack Patterns:** HTTP/S flood attacks often exhibit dynamic and fast-changing attack patterns. Attackers constantly modify their attack vectors, traffic characteristics, and request parameters to evade static detection mechanisms. By monitoring traffic behavior in real-time, behavioral-based systems can adapt and identify emerging attack patterns, even if they have not been encountered before.
- **Zero-Day Attack Detection:** Traditional signature-based detection methods rely on known attack signatures or patterns to identify and block malicious traffic. However, HTTP/S flood attacks often involve new attack vectors that do not have pre-existing signatures. Behavioral-based detection techniques do not rely on static signatures and instead focus on analyzing traffic behavior in real-time. This allows them to identify unusual traffic patterns, excessive request rates, irregularities in user behavior, or unexpected traffic spikes indicative of an ongoing attack.

- **Minimizing False Positives:** One of the challenges in DDoS mitigation is avoiding false positives, which would block legitimate user traffic and impact user experience. By establishing their decisions on behavioral analysis rather than relying only on predefined rules or signatures, these systems can effectively differentiate between malicious and legitimate traffic, reducing the likelihood of blocking legitimate users.
- **Adjusting to Changing Traffic Patterns:** Behavioral-based mitigation systems are designed to adapt and learn from the traffic patterns over time. They establish a baseline of normal behavior by continuously analyzing legitimate traffic. This allows them to dynamically adjust detection thresholds, rules, and mitigation strategies based on real-time traffic behavior, enhancing their ability to detect and mitigate known and emerging attack vectors.

Overall, behavioral-based detection and mitigation techniques offer a dynamic, adaptive, and intelligent approach to counter HTTP/S flood attacks or web DDoS attacks. By monitoring and analyzing traffic behavior in real-time, these systems can identify and mitigate attacks that employ sophisticated evasion techniques, ensuring the availability and security of web applications and services.

Part III – Radware's Cloud Web DDoS Protection

To combat a new generation of powerful web DDoS flood, Radware has developed a new engine for accurately detecting and mitigating web DDoS tsunami attacks, based on Radware's patented behavioral-based DDoS mitigation technologies.

Radware's Cloud Web DDoS Protection leverages dedicated, behavioral-based algorithms with advanced learning capabilities designed to quickly detect and surgically block web DDoS attacks, while minimizing false positives and ensuring legitimate user traffic is not blocked.

This behavioral-based web DDoS protections can accurately distinguish between a legitimate surge in traffic (such as flash crowd or holiday traffic surges) and a malicious traffic generated by botnets.

Highlights of Radware's Solution

Automated, Accurate Detection and Mitigation
Behavioral-based algorithms with advanced learning capabilities to accurately distinguish flash crowd vs. flood attack.

Widest L7 DDoS Attack Coverage
Large-scale, sophisticated Web DDoS Tsunami attacks, smaller, sophisticated attacks, new L7 attack tools and vectors.

Best Protection for Web DDoS Tsunami
Combines automated algorithms and high-scale infra to accurately protect against high-RPS, and complex L7 DDoS threats.

How Radware's Cloud Web DDoS Protection Works

- During peace time, Radware collects both rate variant and rate invariant parameters, to learn application and user behavior under normal conditions. This learning is used to create a granular and detailed baseline of legitimate traffic patterns. The normal baseline created support time of day traffic changes and also application legitimate changes.
- When active attack happens, Radware's web DDoS behavioral detection engine identifies any traffic patterns which fall outside of the known legitimate traffic patterns.
- Real-time L7 signature creation patterns automatically map-out malicious traffic, and create a granular signature which is tailored to the specific characteristics of attack traffic. This signature is dynamic and adaptive, so it can adjust to changing attack traffic patterns and randomized parameters.
- Using the custom signature, attack traffic is stopped, without any impact to legitimate user activity.

Figure 3: Radware's behavioral-based detection engine

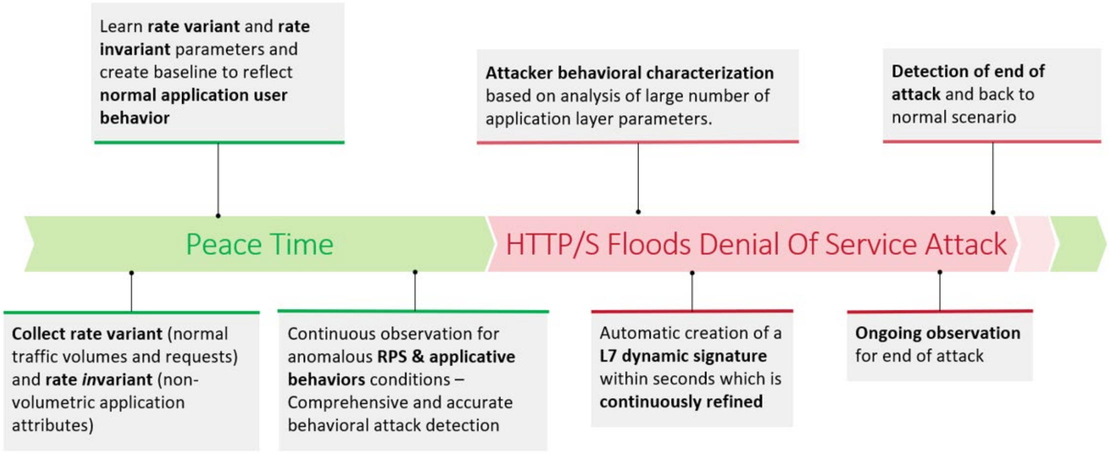
Learn rate variant and rate invariant parameters and create baseline to reflect normal application user behavior
Attacker behavioral characterization based on analysis of large number of application layer parameters.
Detection of end of attack and back to normal scenario
Peace Time
HTTP/S Floods Denial Of Service Attack
Collect rate variant (normal traffic volumes and requests) and rate invariant (non-volumetric application attributes)
Continuous observation for anomalous RPS & applicative behaviors conditions – Comprehensive and accurate behavioral attack detection
Automatic creation of a L7 dynamic signature within seconds which is continuously refined
Ongoing observation for end of attack

Summary

The past 18 months have seen unprecedented growth in DDoS attack activity, which have increased in size, frequency, and sophistication. This growth has been driven by a combination of factors. While each of these factors stands on its own, they coalesced into a fundamental shift in the threat landscape, which is more dangerous than ever before.

Of these changes, web DDoS tsunami attacks have emerged as a uniquely devastating threat to organizations, threatening the availability of mission-critical applications and services. Traditional DDoS protection methods, however, are incapable of providing adequate protection against these attacks, calling for a new approach to DDoS protection.

Radware's behavioral-based Cloud Web DDoS Protection provides real-time, automated, and accurate protection against web DDoS attacks. By combining rate-based and non-rate-based parameters, Radware's algorithms can granularly distinguish between legitimate and attack traffic, and block malicious traffic without impacting legitimate users.

If you are facing a Web DDoS attack, contact Radware immediately for emergency onboarding to our DDoS protection services

Radware Under Attack Contact Page:

https://www.radware.com/underattack/