



WHITEPAPER

 SRC CYBER SOLUTIONS LLP

THREATX

Application Security Evolved: How to Use Attacker-Centric Security to Fix the WAF

Introduction



Application security is in the midst of a transformation. Virtually all enterprise applications and assets have become web-facing whether in the form of a traditional web-application, cloud applications, APIs, microservices, or legacy apps accessed through a web interface. These applications are being continuously developed and delivered at unprecedented speed, making them prime targets for constant attacks by human and automated threats.

In order to keep pace with the expanding web-by-default enterprise, security teams need a new approach to the web application firewall (WAF). WAFs have long been difficult to manage, often requiring constant tuning and work to keep pace with alerts and avoid false positives. Next-Generation WAFs (NG-WAFs) have begun a much-needed shift from the old rules and signatures approach to a new model based on behavioral monitoring of the applications themselves. And while this is a necessary improvement and makes the WAF easier to manage, it often fails to stop subtle or advanced threats, and still relies heavily on manual analysis after the damage has occurred.

ThreatX delivers a fundamentally new approach, extending the evolution of the NG-WAF to add real-time attacker-centric analysis in addition to application-centric analysis. By

correlating across the unique attributes of attackers and actively engaging suspicious behavior, ThreatX can find and stop true threats in real time, while avoiding false positives that require additional work from busy security teams. This paper details core concepts, drivers, and benefits of this attacker-centric approach. It is important to note, however, that attacker-centric technology is only one element of the ThreatX platform.

ThreatX brings a truly innovative approach to the modern WAF, combining protection from traditional threats (e.g. OWASP Top 10), application-centric protection, threat-centric protection, with DDoS protection and application caching services.

Learn more at www.threatx.com

The New Threat and Application Landscape

Over the past decade, most organizations have fundamentally changed how they deliver and access applications. Internally-hosted applications have given way to web and cloud-based applications. Monolithic application architectures have given way to microservice architectures that may publish APIs. Even older legacy applications are often accessed via a web-based portal or run as microservices. While these changes bring immense value to the enterprise, they also introduce new pressures to adapt to them.

The Shifting Security Model

Web-facing applications are by nature accessible from anywhere -- by anyone and anything. Legitimate or nefarious. The immense access means a greater likelihood that weaknesses will be exploited.

This increased exposure leaves very little margin for error in threat prevention. Historically, internally-hosted applications could be protected by multiple layers of enterprise defense that relied on only a few trusted users. Modern web-based applications don't enjoy such luxury, yet security must be delivered at scale and in real time. Today, security teams will typically get one chance to defend an exposed application. In this architecture, the WAF simply must do its job -- reliably and repeatedly.

Traditional WAFs have historically been time-consuming and complex to maintain. As a result, most organizations staff to support only their critical applications. As more applications move to the web, more and more of the enterprise's valuable assets are exposed to the Internet. Security teams have typically spent 90% of their resources supporting 10% of their applications, and this problem threatens to intensify with the migration of applications to the web.

The Age of DevOps and Continuous Delivery

The development and delivery of applications has also fundamentally changed. Highly agile DevOps and Continuous Integration / Continuous Delivery (CI/CD) models are delivering application updates faster than ever before. While these processes are highly beneficial to the organization, the speed of change makes it virtually impossible to predict and detect every security issue prior to deployment.

This has led to a tighter integration between DevOps and Security teams. Make no mistake, coordination between DevOps and Security is a very good thing and leads to safer applications.



Safe application development has never been a replacement for real-time threat prevention for any type of software, whether an operating system, a database, a desktop application, or a web application. The goal is always to build or buy the most secure code you can, and then build defenses to protect it from inevitable attack.

However, when it comes to the protection phase, the constantly evolving nature of applications can make things tricky to say the least. If tuning signatures and rules was painful in the old model, it becomes nearly impossible when the application itself can be updated on a daily basis. The more an application (and its baseline) changes, the more important it is to reliably distinguish attackers from normal users in real time.

The Threat Landscape

As the speed and scope of application development accelerates, the threat landscape also continues to grow more complex. In addition to the traditional threats, such as SQL injection, XSS, XSRF, and others, organizations must also deal with a new wave of ever-evolving threats. This includes automated attacks driven by bots, DDoS attacks, Layer 7 attacks, and more.

Reliance on open-source libraries for application development means that when a new vulnerability is discovered, the issue can instantly expose hundreds of thousands of applications across the Internet. High profile vulnerabilities in Apache Struts or the infamous Heartbleed are just a few examples, but new vulnerabilities make their rounds every day. Once those vulnerabilities are seen, organizations are in a race to patch their apps before attackers automatically scan and exploit them.

This is the reality facing enterprise security teams. More applications, more complexity, faster development, and a vastly sophisticated and ever-changing threat landscape. All this volatility puts high expectations on the WAF, one of the most critical layers of real-time defense. A WAF must be real-time and highly effective, but also easy enough to use that it can be utilized to protect all enterprise applications. WAF technologies must evolve quickly.

The Evolution of the WAF

In order for enterprises to move to a web-by-default approach to applications, they must be able to deliver protection for all of their assets. This has proven to be virtually impossible using existing technologies that rely on signatures and rules or application anomaly identification. And, they're hard to manage and miss threats.

First Generation WAFs

Traditional WAFs have largely relied on highly complex signature and rules engines. While signatures can certainly deliver value, they lead to the perpetual process of updating and tuning. As a result, an organization using a traditional WAF will often use substantial staffing resources to protect only the top 10% to 20% of their web assets from the most obvious threats like the OWASP Top 10.

This is a good start, but savvy attackers continually layer new evasion and encoding techniques to thwart traditional signatures. Protecting a sliver of existing assets against the most basic threats simply doesn't scale to support the broad set of cloud, legacy, and other web assets that exist in an organization. The very real fear of false positives blocking valid users also means that many organizations are often reticent to block anything or block solely based on the most reliable signatures. For organizations that do use signatures to block traffic, every application update requires corresponding security and WAF work to ensure signatures don't inadvertently do damage to the apps they're trying to protect. As a result, many organizations simply accept false negatives (missed threats) as an acceptable tradeoff. Ironically, traditional WAFs have proved to be both hard to use and not particularly effective.

Next Generation WAFs

Next-Generation WAFs have taken a new approach to apply machine-learning to profile an application and understand deviations in behavior. This allows an automated system to learn what is normal and allowed within an application, freeing staff from the Sisyphean task of managing rules. The problem is that detections are not always conclusive around whether anomalies are truly malicious or simply abnormal.

This is a common challenge of many anomaly-based solutions and not necessarily limited to WAFs. Behavioral analysis and machine-learning algorithms can detect problems in ways that signatures can't, but often require manual analysis to distinguish a true threat from a benign anomaly. To avoid the need to investigate a large number of potentially benign anomalies, most solutions set thresholds high and only alert on the most egregious anomalies. This allows subtle and low-and-slow attacks to succeed without detection.



And this contains a critically important point - a behavioral anomaly in an application is often a symptom of an attack. It's a sign that something is amiss, but we don't yet know the cause. As an analogy, is the patient's cough inconsequential or a sign of something serious? To go from symptom to actionable diagnosis we need to be able to identify the threat. Modern WAFs need the ability to test and actively detect specific threats in the same way a physician runs tests to confirm a diagnosis. In security terms, we refer to this as being attacker-centric, and it is often the difference between an application being compromised and not.

Analysis from a recent Ponemon study found that while 73% of organizations use a WAF, 80% of them still report compromised applications. Even more concerning is that the 80% of those companies saw an increase in compromises compared to the same analysis in 2015. Not only are the top-level numbers alarming, the industry is also trending in the wrong direction. At this point, the next-generation WAF has solved only part of the problem: it's much easier to use, but still not reliably effective.

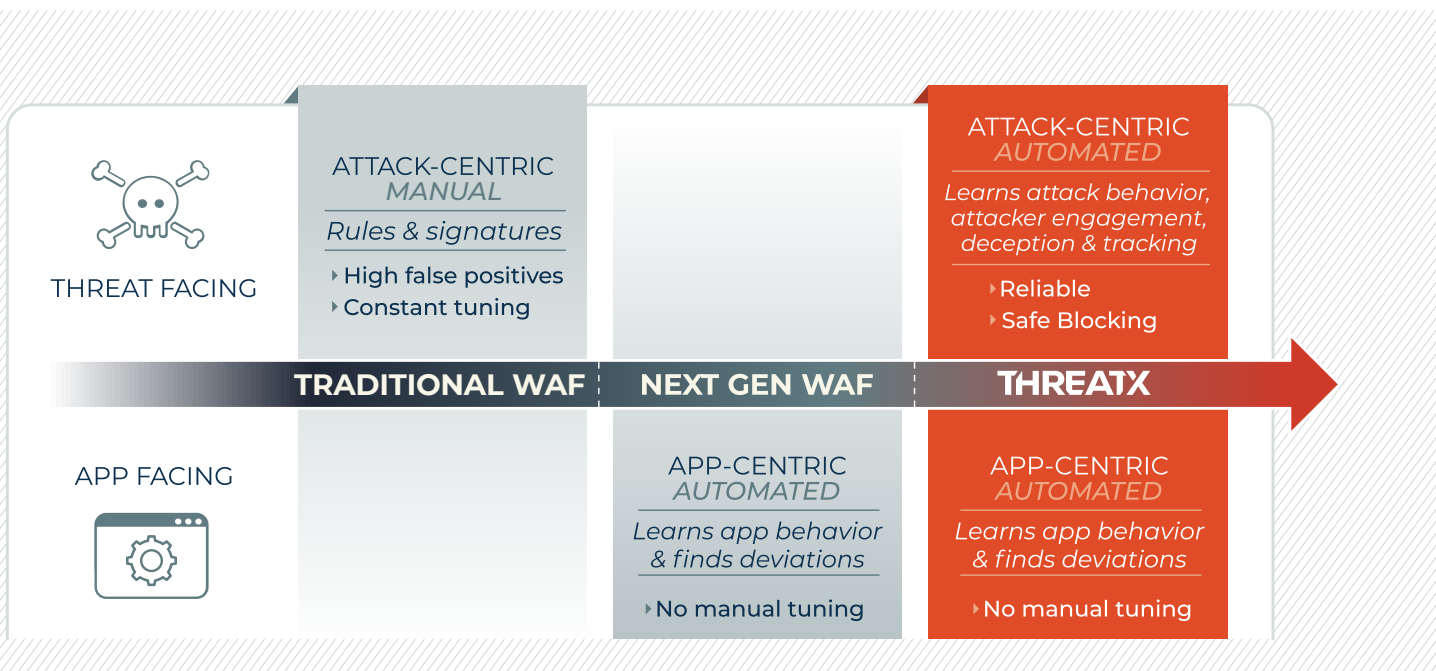
ThreatX WAF Generation

ThreatX completes the evolution of the next-generation WAF technology by adding innovative attacker-centric approaches to application-centric models and capabilities. This approach focuses on detecting the fundamental behaviors and techniques of attackers and actively engaging them to identify and track them over time.

Instead of simply waiting for signs of abuse in the application, ThreatX actively identifies bad actors before a malicious attack impacts the application. Machine-learning models are used to uniquely identify attacker behaviors and traits. The platform persistently monitors across the lifecycle of an attack and can track an attacker over long periods of scanning or reconnaissance. This enables the ability take preventative, blocking action before the attackers progress to more damaging phases of attack.

ThreatX actively interrogates suspicious behaviors and uses deception to confirm a threat, and then persistently tracks the attacker across the Internet. This allows ThreatX to enumerate and follow attackers even as they traverse across different IP addresses and domains. Threats can be verified without employing staff resources. Accumulated knowledge of attacker activities grows over time. ThreatX then leverages this intelligence across all customers so that users can quickly identify bad actors based on cross-site behaviors.

By focusing on the attacker and combining multiple corroborating indicators of suspicious activity, ThreatX builds a progressive risk profile of intent to make precisely accurate blocking decisions based on real behaviors.



The Critical Role of Attacker-Centric Security

The evolution toward attacker-centric security is not limited to web applications. For the past several years, the industry has been struggling to apply machine learning and artificial intelligence models to the job of threat detection. The reasoning was always pretty straightforward - sophisticated attackers were able to avoid traditional, signature-based controls, such as legacy antivirus or IPS, and organizations had a wealth of data that could be used for threat detection.

The Search for the Magic Algorithm

Originally, SIEMs were going to answer all security questions through the analysis of logs. Next came security analytics solutions, various machine learning detections models, and user behavior analysis. The challenge with all of these technologies was not that they didn't work. The problem was that they were rarely conclusive individually. They would detect anomalies, deviations, or "possible" threats. But more often than not, it fell to a human analyst to look at the anomalies and determine if an event was benign or malicious.



This approach is unacceptable for application security. First, organizations simply don't have the human resources to perform reactive investigations every time something strange happens on an application. Even more importantly, by the time an investigation is completed, the damage from a web application attack is probably already done.

Unlike investigating a malware infection on a laptop, the impact of an application attack is immediate and severe. A slow IR-style of response process doesn't work for web-facing applications, which heavily relies on a compressed, real-time layer of defense. For these applications, the only defense that counts is the defense that is automated.

The ThreatX Approach

ThreatX brings together complementary technologies and services to ensure security postures remain actionable and reliable in real time:



Attacker-Centric Threat Detection - Track and model the activity of users from a behavior perspective - whether human or automated - to identify those that are potentially malicious. This unique approach results in high true positives and true negatives.



Risk-Based Protection - Monitor and track risk of user activity. Risk changes based on the content, intensity and frequency of user requests.



24-7 Expertise - Built-in access to managed application protection services. From onboarding new application sites to assisting with day-to-day monitoring and management and adding configurations and updates, ThreatX's security experts are available 24/7. Extra brains (and hands) to ensure peace of mind -- and security of applications.



Ubiquitous & Cloud-Native Deployment - Support varied types of enterprise deployments, including public or private cloud, container-based or on-premises. The ThreatX solution is cloud native, allowing it to horizontally auto-scale with customer requirements -- anywhere in the world.



Holistic Protection - A single solution that protects against OWASP top 10, automated bots, application DDoS, and volumetric DDoS. Can provide protection of all types of external applications, including web-based applications, APIs leveraging REST, GraphQL and web-sockets



Attack Visibility - Visibility into the details of the attack provides necessary telemetry to perform attack validation and initiate incident response. Summaries that illuminate attack classifications over time provide ongoing data to associate with security risk.

This approach brings a wealth of benefits to a security team. First, it provides a progressive way to identify threats conclusively and automatically, while separating low-risk events from those that require real-time response. Anomalies or suspicious activity can be automatically interrogated and active deception can conclusively distinguish a benign user from an attacker. Ongoing profiling and risk scoring gives teams complete control over how they want to handle enforcement.




Most importantly, the ability to enumerate and track attackers across the Internet, allows security teams to escape the seemingly endless game of attacker whack-a-mole. Instead of an attacker being able to retreat to the Internet and start a new attack, ThreatX can immediately recognize an attacking entity returning to a site and apply appropriate controls.

Conclusion

The evolution of attacker-centric security is a critical step for defending modern applications. As the majority of applications transition to cloud deployments that are web-facing and API-backed, application security must be automated, reliable, and provide real-time protection. To deliver on this goal a WAF must have the inward-looking intelligence to understand the application as well as the outward-looking intelligence to continually understand evolving threats. Both of these approaches need to work without static signatures.

And when algorithms alone are not enough, technology needs challenge suspicious behavior, interrogate, and even use code-level deception in order to automatically distinguish a true threat with confidence. This ability to challenge and verify means that organizations can finally avoid false positives that require time-consuming manual tuning and proactively and instantaneously stop even the most advanced threats, before any damage is done.

If you would like to learn more or see the solution in action, please reach out to our team.

www.srccybersolutions.com | +91 120 232 0960 / 1 | sales@srccybersolutions.com   

ABOUT SRC CYBER SOLUTIONS LLP

At SRC Cyber Solutions LLP, we provide Next Generation, Automated and User-Friendly solutions in partnership with AUTOMOX for Patch and Endpoint Management, IRONSCALES for Comprehensive Email Security and Anti-Phishing Protection, THREATX for WAAP (WAF++) with an Attack-Centric approach for Web Application and API protection and Project Ares for Online Gamified Simulation-based Cyber Security Training.