



SecureIQlab®

2022 Cloud Web Application Firewall (WAF) CyberRisk Validation Comparative Report

Published: November 30, 2022

Revision 1.3: 12/8/2022

Report Contents:

1	INTRODUCTION	2
2	AVERAGE PERFORMANCE METRICS	3
3	SECURITY EFFICACY	4
3.1	SECURITY EFFICACY RESULTS OVERVIEW	4
3.2	SECURITY RESULTS DETAILS	5
4	OPERATIONAL EFFICIENCY	6
4.1	OPERATIONAL EFFICIENCY RESULTS OVERVIEW	6
4.2	OPERATIONAL EFFICIENCY DETAILS	7
5	RETURN ON SECURITY INVESTMENT	9
6	CYBERRISK RIPPLE	10
6.1	ANNUAL TOTAL COST OF OWNERSHIP	10
6.2	COMPLETE SECURITY SCORE VS. ROSI	11
6.3	OPERATIONAL EFFICIENCY SCORE VS. ROSI	12
6.4	2022 CLOUD WAF CYBER RISK RIPPLE	14
7	CONCLUSION	15
8	APPENDIX	16
8.1	CLOUD WAF TEST DEPLOYMENT	16
8.2	TEST EXECUTION	17
8.3	ATTACK TYPES	18
8.4	PRODUCT CONFIGURATION	18
8.5	PRODUCT RULES	18
8.6	PRODUCT PRICING	19
8.7	CLOUD WAF 2022 TEST STATISTICAL ANALYSIS	20
8.8	VENDOR PARTICIPATION	21
8.9	METHODOLOGY	21
9	CONTACT INFORMATION	22
10	COPYRIGHT AND DISCLAIMER	23

1 INTRODUCTION

The remote workforce has dissolved the network perimeter. This, along with cost savings and operational simplicity, is driving organizations to adopt cloud infrastructure. The explosive adoption of the cloud comes with associated risks. For example, web application-based vulnerabilities are among the top breach vectors¹. Cloud-based web application firewalls (WAFs) are designed to mitigate this risk by protecting applications without interrupting business operations in the cloud-first world. Cloud-based web application firewalls (WAFs) should accurately detect, prevent, and log attacks while remaining resistant to false positives. In addition, WAFs should provide auditing capabilities, meet regulatory compliance requirements, and enable business continuity.

SecureIQLab has conducted an exhaustive test of 14² enterprise class web application firewall (WAF) products to determine their security efficacy, operational efficiency, and return on security investment (ROSI). This report summarizes the comparative test results for the 14 products that were put through their paces. Individual reports for each of the 14 WAF solutions tested are also available at <https://secureiqlab.com/publications/>.

In the course of this test, the products were subjected to a battery of diverse attacks. Simple ecommerce applications and multiuser web applications were used as targets. Test cases were based upon industry frameworks and regulatory requirements, such as the OWASP Top 10³ and PCI DSS⁴. The resultant data was obtained while securing targeted cloud applications hosted on AWS, Azure, GCP, and Oracle.

The vendors and tested products were selected based on their meeting one or more of the following three criteria:

1. Market Leaders: Either in terms of revenue generated, customer numbers globally, or strong channel play
2. Analyst and Enterprise challengers: Small-mid-large enterprise security professional surveys, Direct 1:1 Inquiries and engagement with enterprises, organizations, MSP's, MSSP's and Gartner MQ, buyers guide, Forrester Wave, and IDC reports
3. New market entrants and interested participating vendors: Challengers claiming breakthrough technology offerings

¹ <https://www.verizon.com/business/resources/reports/dbir/2021/masters-guide/summary-of-findings/>

² Testing was attempted on a total of 17 cloud WAF solutions. See Appendix section 8.8 for details.

³ [Open Web Application Security Project®](#)

⁴ [Payment Card Industry Data Security Standards](#)

2 AVERAGE PERFORMANCE METRICS

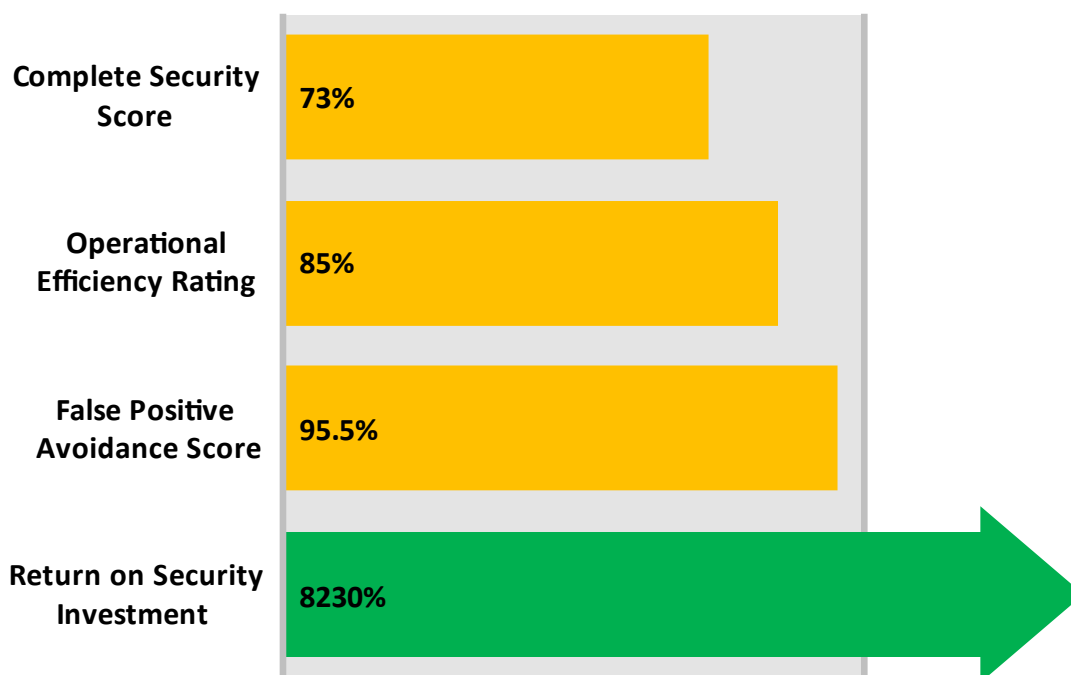


Figure 1. Overall Averages for the 14 Validated Products

Figure 1 provides the average scores for security efficacy, operational efficiency, false positive avoidance, and return on security investment (ROSI). Return on security investment is based upon the aggregation of the *Complete Security Score*, the *Operational Efficiency Rating*, the *False Positive Avoidance Score*, and the *Cost of Ownership*.

The good news in SecureIQLab's findings is that the average ROSI is a substantially positive number.

The fair news is that the average *Operational Efficiency Score* is respectable, but there's room for improvement.

The bad news is that the average *Complete Security Score*⁵, although an improvement over last year's results, falls short of where we need to be as an industry to protect critical assets. This is not to diminish the great work vendors do. However, vendors are up against very well-funded, highly motivated, and intelligent adversaries.

⁵ Products tested used default configurations without custom tuning. Tuning may increase the security efficacy of WAFs.

3 SECURITY EFFICACY

3.1 SECURITY EFFICACY RESULTS OVERVIEW




Figure 2. Comparative for Complete Security Scores

Figure 2 provides an overview comparative of the SecureIQLab findings during the security validation of the tested products. The *Complete Security Score* depicts the percentage of all attacks blocked by the WAF versus the total number of attacks tested. Equation 1 below depicts the *Complete Security Score* calculation, which is based on an unweighted percentage of all attacks blocked.

$$\text{Complete Security Score} = \frac{\left(\frac{\text{Deduped OWASP Score}}{5} + \frac{\text{Bot Attack Score}}{5} + \frac{\text{Layer 7 DoS Score}}{5} + \frac{\text{Resiliency Score}}{5} + \frac{\text{Vulnerable Environment Score}}{5} \right)}{5}$$

Equation 1. Calculation of Complete Security Score

Every cloud WAF evaluated in this test was subjected to more than 400 real-world test scenarios targeting small-to-medium businesses and enterprises alike. A grand total of over 9100 attacks were used that encompass these scenarios and categories. The depth and scope of the testing performed by SecureIQLab is a first in the cybersecurity industry. SecureIQLab will continue to add attack libraries and other relevant operational metrics in future iterations of this test.

3.2 SECURITY RESULTS DETAILS

Security efficacies were determined for five categories. Detailed explanations and results for each of these categories are provided in the individual test reports. Table 1 provides an overview of test results.

Vendor	OWASP	Bot	Layer 7 DoS	Resiliency	Vul. Web Environment	Complete Security Score
Akamai	86%	17%	40%	63%	80%	60%
AWS	79%	17%	60%	17%	65%	46%
Microsoft	96%	83%	100%	89%	85%	92%
Barracuda	75%	50%	80%	52%	75%	63%
Cloudflare	84%	17%	80%	55%	70%	63%
F5	93%	67%	100%	69%	100%	84%
Fortinet	96%	100%	100%	87%	90%	96%
Google	89%	50%	80%	63%	90%	74%
Imperva	95%	100%	100%	68%	75%	91%
Oracle	90%	17%	40%	90%	65%	67%
Prophaze	85%	33%	80%	68%	55%	68%
Stackpath	94%	100%	100%	38%	75%	84%
Sucuri	81%	33%	100%	46%	70%	69%
Wallarm	89%	17%	80%	76%	60%	70%
Average	88%	50%	81%	63%	75%	73%

Table 1. Security Efficacy Results

Security is the primary purpose of purchasing a WAF, but it isn't as easy as looking at just the complete security score to make a purchasing decision. The static scores shown above do not factor in the ability to tune WAFs to meet an organization's needs. Tuning can improve security efficacy, but there is the potential that an increase in the number of false positives may occur. Evaluating the sheer combination of the numerous tuning parameters for the tested WAFs is beyond the scope of this general test report. Organizations that wish to have their tuned products tested should contact SecureIQLab for more information.

4 OPERATIONAL EFFICIENCY

4.1 OPERATIONAL EFFICIENCY RESULTS OVERVIEW

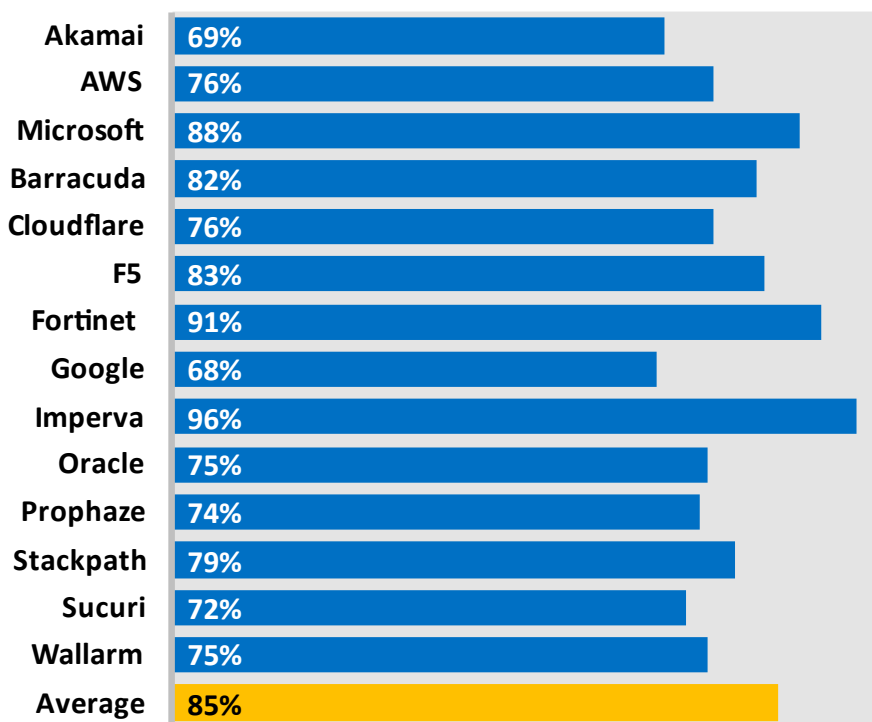


Figure 3. Comparative for Operational Efficiency Ratings

Cloud-based WAF technology allows for the creation of customizable security configurations which contribute to operational efficiency in the following ways:

- Ease of deployment and integration
- Less complex to manage
- Ease of risk management
- Scalable and elastic
- Monitoring, logging, and control capabilities
- Secure business-related transactions

All 14 products were validated in each of these areas of operational efficiency. Figure 3, above, provides a comparison of the resulting *Operational Efficiency Ratings*.

Category scores were calculated by aggregating earned points and then dividing this number by the total number of possible points to find a percentage. Detailed category results are available for each vendor in their individual report.

The *Operational Efficiency Rating* was calculated by adding together the total points for each category, then dividing this number by the maximum potential points (114) and multiplying that number by 100%. Below, Equation 2 states the *Operational Efficiency Rating* calculation.

$$\text{Operational Efficiency Rating} = \frac{\left(\frac{\text{Ease of Deployment Points}}{114} + \frac{\text{Ease of Management Points}}{114} + \frac{\text{Ease of Risk Management Points}}{114} + \frac{\text{Scalable and Elastic Points}}{114} + \frac{\text{Logging and Auditing Points}}{114} \right) \times 100\%}{114 \text{ points}}$$

Equation 2. Operational Efficiency Score Calculation

The mode for each feature validated is used to calculate the test group feature results. Group test averages were then calculated by adding the modes for each feature and then dividing this number by the total number of possible points to find a percentage.

4.2 OPERATIONAL EFFICIENCY DETAILS

Operational efficiency was determined for five different areas of use. Operational efficiency testing also incorporated a false positive avoidance test discussed separately after these five tests. Detailed explanations and results for each of these five tests are provided in the individual test reports⁶. Table 2 provides an overview of the results of our five operational efficiency tests.

Vendor	Ease of Deployment	Ease of Management	Ease of Risk Management	Scalable and Elastic Capabilities	Logging and Auditing Capabilities	Operational Efficiency Rating
Akamai	67%	70%	67%	73%	71%	69%
AWS	63%	83%	62%	93%	86%	76%
Microsoft	85%	87%	76%	100%	95%	88%
Barracuda	96%	90%	67%	67%	76%	82%
Cloudflare	85%	67%	86%	73%	71%	76%
F5	85%	90%	71%	73%	90%	83%
Fortinet	89%	97%	86%	87%	95%	91%
Google	56%	63%	57%	93%	86%	68%
Imperva	96%	97%	95%	100%	95%	96%
Oracle	63%	83%	62%	93%	81%	75%
Prophaze	100%	63%	71%	73%	57%	74%
Stackpath	81%	63%	90%	73%	90%	79%
Sucuri	85%	73%	57%	80%	62%	72%
Wallarm	70%	80%	71%	87%	71%	75%
Average	78%	87%	81%	93%	90%	85%

Table 2. Operational Efficiency Results

⁶ Individual test reports available at <https://secureiqlab.com/publications/>




Figure 4. False Positive Avoidance Scores

WAFs need to allow business-related transactions while blocking malicious activity. False positives interfere with the operation of the business and can result in decreased revenue and reduced employee efficiency. False positives also increase alerts for security teams who are already stretched thin, thereby contributing to alert fatigue. Alert fatigue can reduce the security efficacy of a product as critical alerts may be lost in noise. Policies need to be tuned to minimize false positives; however, care must be taken so as not to adversely impact security efficacy. In an ideal world, products would have a perfect 100% *False Positive Avoidance Score* while maintaining excellent security efficacy. In the real-world this is not always the case. Properly tuned security devices should not block benign traffic. SecureIQLab used over 300 different false positive test cases to validate that the product under test (PUT) would not block simulated consumer usage. Blocking legitimate user activity constitutes a false positive, increases the operational burden for the enterprise and would require additional tuning to correct. The results for the False Positive Avoidance testing are found above in Figure 4.

Table 3 shows how the *Security Efficacy Scores* and *False Positive Avoidance Scores* correlate with the *Operational Efficiency Scores*. Table 3 is sorted by operational efficiency scores from high to low.

Vendor	Operational Efficiency	Security Efficacy	False Positive Avoidance
Imperva	96%	91%	97.7%
Fortinet	91%	96%	87.3%
Microsoft	88%	92%	95.1%
F5	83%	84%	81.2%
Barracuda	82%	63%	95.8%
StackPath	79%	84%	98.1%
Cloudflare	76%	63%	100.0%
AWS	76%	46%	92.5%
Wallarm	75%	70%	99.0%
Oracle	75%	67%	98.7%
Prophaze	74%	68%	98.4%
Sucuri	72%	69%	99.7%
Akamai	69%	60%	94.8%
Google	68%	74%	98.1%

Table 3. Operational Efficiency, Security Efficacy and False Positive Avoidance Correlation

5 RETURN ON SECURITY INVESTMENT




Figure 5. Return on Security Investment

Return On Security Investment (ROSI) differs from the traditional Return On Investment (ROI) in that ROSI is calculated on the bases of prevented losses and not generated income. The Security Effectiveness (SE), Operational Efficiency (OE), *Annual Product Cost* (APC), and *Annual Loss Expectancy* (ALE) are all used in the calculation of the ROSI in Figure 5. Equation 3 demonstrates how ROSI is calculated by SecureIQLab.

- *Security Effectiveness* (SE): Security solutions with higher security efficacies will stop more threats and prevent more loss. *Complete Security Scores* in decimal format are used as SE values in our ROSI calculations.
- *Annual Total Cost of Ownership* (TCO): The total cost of the security solution is the APC plus the *Annual Operational Expense* for the solution. See Appendix, section 8.6 for TCO calculation.
- *Annual Loss Expectancy* (ALE): Anticipated annual financial loss related to security incidents. This is unique to each organization and may be extrapolated from historical losses. SecureIQLab uses an ALE equal to \$4.24M⁷ in calculating ROSI.

$$ROSI = \frac{ALE \times SE - TCO}{TCO}$$

Equation 3. ROSI Calculation

⁷ <https://www.ibm.com/reports/data-breach>

6 CYBERRISK RIPPLE⁸

Now it is time to start putting it all together.

6.1 ANNUAL TOTAL COST OF OWNERSHIP

First, let's look at *Annual Total Cost of Ownership*, sorting products into high (large markers), medium (medium sized markers) and low categories (small markers). An *Annual Total Cost of Ownership* less than \$35,000⁹ is in the low category, an TCO less than \$46,000¹⁰ is in the medium category and all TCOs above \$46,000 are in the high category. This leads to the following Table 4:





























Vendor	TCO	Complete Security Score vs. ROSI Graph Marker	Operational Efficiency Rating vs. ROSI Graph Marker
Akamai	High		
AWS	Low		
Microsoft	Low		
Barracuda	Med		
Cloudflare	Low		
F5	Med		
Fortinet	Med		
Google	Low		
Imperva	Low		
Oracle	Low		
Prophaze	Low		
Stackpath	High		
Sucuri	Low		
Wallarm	High		

Table 4. Annual Total Cost of Ownership¹¹

The round marker is used for the subsequent Figure 6 and Figure 8 graphs. The triangular marker is used for the subsequent Figure 7 and Figure 8 graphs. Because each product was evaluated in both categories, each product has both a circle and a triangle.

TCO is essential because it helps organizations budget appropriately for the optimal products for their security and operational requirements.

⁸ All colors in figures and tables in this section do not indicate quality.

⁹ Determined by Average TCO – STDEV/2, rounded to the nearest \$1000.

¹⁰ Determined by Average TCO + STDEV/2, rounded to the nearest \$1000.

¹¹ See Appendix 8.6 for additional details on Total Cost of Ownership

6.2 COMPLETE SECURITY SCORE VS. ROSI

The *Complete Security Score* relates to the Return on Security Investment (ROSI). This leads to the following graph. See Figure 6.



Figure 6. CyberRisk Security Efficacy Ripple

Four categories for Security Efficacy are derived from the results of the *Operational Efficiency Scores*, *Security Efficacy Scores* and ROSI¹². These four categories are:

- Leaders:** These Cloud WAF solutions demonstrated a combination of superior security and ROSI. That is, these solutions provide stronger security technology at competitive pricing. *Security Efficacy Leaders* have an above average ROSI and a *Complete Security Score* greater than the average of the *Operational Efficiency* and *Security Efficacy Scores*.
 - 🔒 F5, Fortinet, Imperva, and Microsoft are *Security Efficacy Leaders*.
- Contenders:** These Cloud WAF solutions demonstrated excellent prevention and detection capabilities, delivering with an attractive ROSI, befitting small-to-medium enterprises and businesses. *Security Efficacy Contenders* have a ROSI value greater than the average ROSI – 1 standard deviation and a *Complete Security Score* greater than one standard deviation below the average of the *Operational Efficiency* and *Security Efficacy Scores*.
 - 🔒 Google, Prophaze, StackPath, and Sucuri are *Security Efficacy Contenders*.

¹² A Security Efficacy Leader ranking for a product is no guarantee that the product will meet your specific security requirements.

- **Visionaries:** These Cloud WAF solutions demonstrated either excellent security or ROSI. That is, solutions in this category were priced competitively or provided better than average security. *Security Efficacy Visionaries* have a ROSI value greater than the average ROSI – 2 standard deviations and a *Complete Security Score* greater than two standard deviations below the average of the *Operational Efficiency* and *Security Efficacy Scores*.

🔒 Barracuda, Cloudflare, and Oracle, Wallarm are *Security Efficacy Visionaries*.

- **Upcomers:** These Cloud WAF solutions demonstrated lower *Security Efficacy* standards which contributed to lower ROSI. *Security Efficacy Upcomers* have a ROSI value less than the average ROSI – 2 standard deviations or a *Complete Security Score* less than two standard deviations below the average of the *Operational Efficiency* and *Security Efficacy Scores*.

🔒 Akamai and AWS are *Security Efficacy Upcomers*¹³.

6.3 OPERATIONAL EFFICIENCY SCORE VS. ROSI

Third, we look at the comparison between *Operational Efficiency* and ROSI. The Y-axis labels are found on the right of the graph because we are going to combine this graph with the prior graph when we synthesize the results.








Figure 7. CyberRisk Operational Efficiency Ripple

¹³ Products ranked as *Security Efficacy Upcomers* may still meet business-specific security use case requirements.

Figure 7 allows us to determine that cloud web application firewalls as a group have a more mature operational efficiency than overall security efficacy. In addition, four categories for *Operational Efficiency* are derived from the results of the *Operational Efficiency Scores*, *Security Efficacy Scores* and ROSI¹⁴. These four categories are:

- **Leaders:** These cloud WAF solutions demonstrated a combination of high-grade operational efficiency and superior ROSI. That is, these solutions combine ease of deployment, integration and resource management at competitive pricing. *Operational Efficiency Leaders* have an above average ROSI and an *Operational Efficiency Score* greater than the average of the *Operational Efficiency* and *Security Efficacy Scores*.
 F5, Fortinet, Imperva, and Microsoft are *Operational Efficiency Leaders*.
- **Contenders:** These Cloud WAF solutions demonstrated excellent operational efficiency when it came to ease of deployment, integration and strike a good balance between the technology and resource management with an attractive ROSI. *Operational Efficiency Contenders* have a ROSI value greater than the average ROSI – 1 standard deviation and an *Operational Efficiency Score* greater than one standard deviation below the average of the *Operational Efficiency* and *Security Efficacy Scores*.
 AWS, Barracuda, Cloudflare, Google, Oracle, Prophaze, StackPath, and Sucuri are *Operational Efficiency Contenders*.
- **Visionaries:** These Cloud WAF solutions demonstrated good operational efficiency standards when it came to ease of deployment, integration or an excellent ROSI. *Operational Efficiency Visionaries* have a ROSI value greater than the average ROSI – 2 standard deviations and an *Operational Efficiency Score* greater than two standard deviations below the average of the *Operational Efficiency* and *Security Efficacy Scores*.
 Wallarm is an *Operational Efficiency Visionary*.
- **Upcomers:** These Cloud WAF solutions demonstrated lower operational efficiency standards when it came to ease of deployment, integration and struck a lower balance between the technology and resource management, delivering a lower ROSI. *Operational Efficiency Upcomers* have a ROSI value less than the average ROSI – 2 standard deviations or an *Operational Efficiency Score* less than two standard deviations below the average of the *Operational Efficiency* and *Security Efficacy Scores*.
 Akamai is an *Operational Efficiency Upcomer*¹⁵.

¹⁴ An Operational Efficiency Leader ranking is no guarantee that the product will meet your specific operational requirements.

¹⁵ Products ranked as *Operational Efficiency Upcomers* may still meet business-specific operational use case requirements.

6.4 2022 CLOUD WAF CYBERRISK RIPPLE

Assembling the previous data into one figure yields:




Figure 8. SecureIQLab Cloud WAF CyberRisk Ripple

The SecureIQLab Cloud WAF CyberRisk Ripple, Figure 8, is the culmination of an exhaustive and detailed approach to cloud WAF testing and validation research. This chart plots the Complete Security Score, circular markers, and the Operational Efficiency Score, triangular markers, versus ROSI. The upper right corner highlights the best operational efficiency, security efficacies and ROSI. Small markers indicate lower *Total Cost of Ownership*, medium sized markers indicate medium *Total Cost of Ownership*, and larger markers indicate a higher *Total Cost of Ownership*.

7 CONCLUSION

Summarizing the data from the previous section we get Table 5.

Vendor	Security Efficacy	Operational Efficiency
Akamai	Upcomer	Upcomer
AWS	Upcomer	Contender
Microsoft	Leader	Leader
Barracuda	Visionary	Contender
Cloudflare	Visionary	Contender
F5	Leader	Leader
Fortinet	Leader	Leader
Google	Contender	Contender
Imperva	Leader	Leader
Oracle	Visionary	Contender
Prophaze	Contender	Contender
Stackpath	Contender	Contender
Sucuri	Contender	Contender
Wallarm	Visionary	Visionary

Table 5. SecureIQLab's Cloud WAF CyberRisk Ripple Results

Testing evolves with the changing IT and threat landscapes. Each test iteration provides additional insight into how the WAF industry and individual vendors evolve to address operational business requirements. The 2021 Cloud WAF CyberRisk testing provided a baseline for cloud WAF industry performance. The 2022 Cloud WAF CyberRisk test provides readers with a second evaluation of security efficacy, operational efficiency, and the resultant ROSI for popular cloud WAF solutions.

Apart from this comparative report which highlights the overall comparative metrics, SecureIQLab's individual test reports offer greater details for each of the vendors tested. Still, given that every organization's attack surface, business requirements, and risk mitigation strategy are unique, thorough evaluation of cloud WAF technologies before deployment is recommended.

8 APPENDIX

8.1 CLOUD WAF TEST DEPLOYMENT




Figure 9. Amazon Machine Image WAF Deployment




Figure 10. AWS WAF Deployment

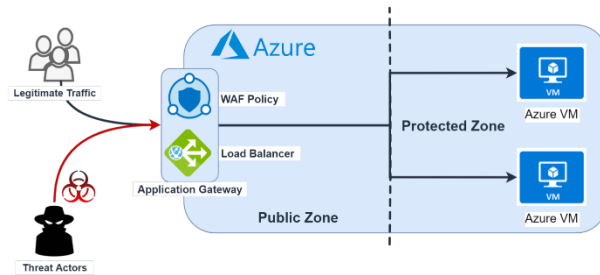


Figure 11. Azure WAF Deployment

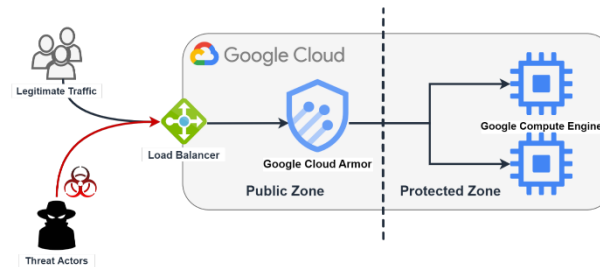


Figure 12. Google WAF Deployment




Figure 13. Oracle WAF Deployment




Figure 14. SaaS WAF Deployment

The cloud WAFs were deployed with default policies¹⁶ with elastic load balancers to protect the web-applications on AWS, Azure, GCP or OCI see Figure 9 through Figure 14. All web-application transactions were inspected by the cloud WAFs. In doing so, the cloud WAFs were expected to provide protection against threats that were originated by malicious actors while allowing intended usage of application resources.

During deployment, our engineers noted the time it took to deploy with out of the box controls and the complexity of the deployment. Also noted was whether our engineering team was required to contact the WAF vendor's support team to successfully complete the WAF deployment. See Table 2 for an overview of deployment findings.

8.2 TEST EXECUTION

SecureQLab performed security validation using crafted attacks that are relevant to today's cloud application hosted on cloud and cloud native applications. SecureQLab carefully curated such attacks via research generated by our own red team as well as the attacks that are prevalent in the wild. Open-source tool kits were also utilized while performing this assessment.

Before the testing was conducted, SecureQLab validated that the cloud WAF solution was in an operational state by verifying the following:

Connection Validation:

1. Before any test is conducted, SecureQLab ensures that the Cloud WAF can be accessed by the administrator and is passing normal application traffic. This is to ensure that any dynamic content such as IP blacklist protection can be updated on regular basis by the cloud WAF.

Logging:

2. SecureQLab understands that logging is a critical and crucial component while running a cloud WAF. SecureQLab verifies that the cloud WAF being tested will have sufficient administrative as well as attack logging to ensure Security Analyst can troubleshoot and fix issues as required.

¹⁶ Please see Appendix 8.4 Product Configuration and Appendix 8.5 Product Rules for deployment details.

Updates:

3. Protocol updates in the form of rules, signatures and reputations will be applied as they become generally available. SecureQLab will make best effort to apply these updates to the products prior to the evaluation.

The above processes were repeated wherever applicable throughout the test. Once the deployment of the WAF solution and baseline testing were completed, the security validation testing began.

The first phase of attack was to gather information and perform reconnaissance against the application. The was done to gather as much information as possible to be utilized when penetrating the target during the vulnerability assessment and exploitation phases. SecureQLab performed vulnerability analysis using automated tools such as Burpsuite and Nessus in addition to performing manual analysis. The main objective of vulnerability analysis is to discover flaws in the systems and applications which can be leveraged by an attacker. These flaws ranged anywhere from host and service misconfiguration to insecure application design. Vulnerability Analysis was based on:

1. Active Scan: Active scan involves direct interaction with the component being tested for security vulnerabilities.
2. Passive Scan: Passive scan involves meta-data analysis and traffic monitoring.

Once information gathering and reconnaissance was completed, we began exploitation as the next phase in this process. Penetration testing was critical in the evaluation of cloud WAF technologies.

Once exploited, “post-exploitation” was undertaken. Post-exploitation refers to the actions taken after the initial compromise of a system or device. It often describes the methodical approach of using privilege escalation or pivoting techniques—which allowed SecureQLab, in this case, to establish a new source of attack from the new vantage point in the system—to gain additional access to systems or network resources. We demonstrate the risk presented by exploitable systems and what post-exploitation may likely occur with web applications.

Additionally, defense evasion is an important tool in an attacker’s arsenal. This allows old methods and techniques to be repurposed to evade protection against attacks which might otherwise get blocked by the Cloud WAF.

The testing demonstrates the effectiveness of the product under test (PUT) to protect vulnerable assets from targeted threats and exploitation. This asset/target and threat-based approach forms the basis from which the PUT security effectiveness is measured.

8.3 ATTACK TYPES

The SecureQLab threat and attack suite contains attacks (including mutations of the same underlying attacks) and proprietary exploits harvested through our test harness or crafted by our threat research team. SecureQLab has a number of complex web applications which have also been constructed to include known vulnerabilities and coding errors. Based on the intended attack, groups of exploits are carefully selected from this library to test. Each exploit has been validated to impact the target vulnerable host(s) by compromising the asset, which can range from the web server, the web application, to the site. The level of compromise can include instigating a denial-of-service (DoS) condition, providing administrator/root access to the host server, allowing malicious users to amend system parameters or application data before submission, browsing and/or retrieving files stored on the host server, or escalating user privileges.

8.4 PRODUCT CONFIGURATION

Cloud WAF products were deployed and configured according to the default instructions found on the vendors’ websites and, where applicable, on cloud infrastructure providers’ marketplaces.

8.5 PRODUCT RULES

Cloud WAF products were configured according to the default instructions found on the vendors’ websites and, where applicable, on cloud infrastructure providers’ marketplaces. Default rule sets were used for the products in

this test. However, any “Detect Only” mode settings that were part of default configurations were modified to “Block” mode, with default rulesets used as applicable.

Any further required tuning to deploy in the test environment was performed according to publicly available vendor recommendations to replicate the customer experience during the deployment and management of the product. (Enterprises are advised to exercise due diligence during this process to avoid impacting business.)

8.6 PRODUCT PRICING¹⁷

Normalized pricing for products tested is listed in Table 6. Product pricing models varied in complexity, from one monthly fee to an à la carte feature set, prorated hourly, model.¹⁸ There were also significant variations in pricing structure and feature set combinations¹⁹ between all 14 vendors tested.²⁰ In order to best compare the variety of options tested, the pricing options used are normalized²¹ to support a minimum of:

- 1 TB of data per month
- 40 million requests monthly
- 1 site being protected

Normalized pricing was also calculated based on the product being enabled for the entire month, fees being paid monthly where required²², and protection against botnet attacks²³ being included. Lastly, cloud infrastructure costs for the virtual deployment models are included in the normalized pricing.

Vendor	Normalized Price	Annual TCO
Akamai	\$34,800	\$69,197
AWS	\$1,122	\$33,128
Microsoft	\$5,142	\$32,028
Barracuda	\$11,075	\$39,717
Cloudflare	\$2,400	\$32,005
F5	\$4,195	\$37,580
Fortinet	\$7,984	\$36,306
Google	\$576	\$34,305
Imperva	\$1,522	\$25,511
Oracle	\$216	\$30,611
Prophaze	\$2,400	\$33,300
Stackpath	\$20,400	\$49,433
Sucuri	\$240	\$31,584
Wallarm	\$25,402	\$55,705
Average	\$8,391	\$38,601

Table 6. Normalized Pricing and Annual Total Cost of Ownership

The Annual Total Cost of Ownership (TCO) is the *Annual Product Cost* (APC) of the WAF solution plus the *Annual Operational Expense* (AOE) for the product. The APC is the *Normalized Price*. The *Annual Operational Expense* is the cost of one full-time employee that spends 10% of their time supporting the WAF product, divided by the decimal format for the *Operational Efficiency Rating* (OE) and the *False Positive Avoidance Score* (FPA). Equation 4

¹⁷ Support tiers and offerings vary by vendor and may need to be considered for budgetary purposes. Prices for support are not included in pricing.

¹⁸ Pricing will vary.

¹⁹ Vendor standard feature sets may change with time.

²⁰ Pricing is dynamic.

²¹ Pricing normalization may cause pricing variance with vendor list pricing.

²² Some vendors offer incentives for annual and multi-year commitments.

²³ Botnet protection is not a standard feature for all products.

depicts the TCO calculation.

$$TCO = \frac{AOE}{OE * FPA}$$

Equation 4. Calculation of TCO

The cost of the full-time employee is based on a \$75/hr rate and includes overhead. For the purpose of this report, the total cost of the employee comes to \$225k/year.

8.7 CLOUD WAF 2022 TEST STATISTICAL ANALYSIS

Statistical analysis of the Cloud WAF 2022 test results were performed at the conclusion of testing. The Matthews correlation coefficient, precision, and recall for the tests results are provided below in Table 7.

Vendor	MCC	Precision	Recall
Akamai	0.20	0.95	0.07
AWS	0.14	0.93	0.05
Microsoft	0.50	0.95	0.29
Barracuda	0.21	0.96	0.08
Cloudflare	0.23	1.00	0.08
F5	0.30	0.81	0.15
Fortinet	0.59	0.87	0.42
GWAF	0.28	0.98	0.11
Imperva	0.49	0.98	0.27
Oracle	0.24	0.99	0.09
Prophaze	0.25	0.98	0.09
Stackpath	0.37	0.98	0.17
Sucuri	0.26	1.00	0.10
Wallarm	0.26	0.99	0.10
Average	0.27	0.96	0.11

Table 7. Cloud WAF 2022 Test Statistics

To better understand these performance metrics, Table 8 provides definitions for the variables used in the calculation of these metrics.

Variable	Meaning	Definition
TP	True Positive	# Benign Allowed
FP	False Positive	# Benign Blocked
FN	False Negative	# Attacks Missed
TN	True Negative	# Attacks Blocked

Table 8. Definition of Variables

Calculation of the Matthews correlation coefficient is provided in Equation 5 below.

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP) \times (TP + FN) \times (TN + FP) \times (TN + FN)}}$$

Equation 5. Matthews Correlation Coefficient Calculation

Calculation of precision and recall are provided in Equation 6 and Equation 7, respectively.

$$Precision = \frac{TP}{TP + FP}$$

Equation 6. Precision Calculation

$$Recall = \frac{TP}{TP + FN}$$

Equation 7. Recall Calculation

The statistical analysis of the 2022 Cloud WAF test results indicate utility in incorporating further false positive testing as an avenue for future research.

8.8 VENDOR PARTICIPATION

Table 9 lists the test status for each of the 17 vendors that SecureQLab attempted to evaluate in this test iteration.

Vendor	Test Status
Akamai	Testing Completed
AWS	Testing Completed
Microsoft	Testing Completed
Barracuda	Testing Completed
Citrix	Evaluation terminated due to product issues ²⁴
Cloudflare	Testing Completed
F5	Testing Completed
Fastly	Evaluation terminated by vendor ²⁴
Fortinet	Testing Completed
Google	Testing Completed
Imperva	Testing Completed
Oracle	Testing Completed
Prophaze	Testing Completed
Radware	Evaluation terminated by vendor ²⁴
StackPath	Testing Completed
Sucuri	Testing Completed
Wallarm	Testing Completed

Table 9. Vendor Test Status

8.9 METHODOLOGY

The test was conducted in accordance with the standards of the Anti-Malware Testing Standards Organization ([AMTSO](#)). The test used version 2.0 of the SecureQLab [Cloud WAF CyberRisk Validation Methodology](#) (AMTSO Test ID: AMTSO-LS1-TP054).

Tests were performed utilizing black-box and gray-box testing. Black-box testing assumes that the internal code structure of the product being tested is unknown to the tester. For this testing approach, testers are not required to know a system's implementation details. Gray-box testing assumes that part of the product's internal code structure is known to the tester.

²⁴ Contact SecureQLab for details.

Default configurations and rule sets were used for the majority of the products in this test. However, any “Detect Only” mode settings that were part of default configurations were modified to “Block” mode, with default rulesets used as applicable.

Tuning a WAF can be complex. Tuning was based on industry and marketplace expectations that these solutions will require minimal to no tuning during provisioning, deployment, and management phases, which translates to lower operational expenses and increased revenue for the targeted audience, i.e., SMBs, managed service providers (MSPs), and managed security service providers (MSSPs).

Further, any required tuning was performed according to publicly available vendor recommendations, to replicate the customer experience during the deployment and management of the product. (Enterprises are advised to exercise due diligence during this process to avoid impacting business.) Browsing the WAF-protected applications was performed using standard user transactions that included form submissions, comment writing, ecommerce transactions, and other transactions.

The Anti-Malware Testing Standards Organization (AMTSO) is an international nonprofit association that focuses on addressing the global need for improvement in the objectivity, quality, and relevance of anti-malware testing methodologies. SecureQLab is a member of AMTSO.

The Open Web Application Security Project (OWASP) is a nonprofit foundation that works to improve the security of software. “Through community-led open-source software projects, hundreds of local chapters worldwide, tens of thousands of members, and leading educational and training conferences, the OWASP Foundation is the source for developers and technologists to secure the web.”²⁵ It publishes the OWASP Top 10 Report. SecureQLab has no affiliation with OWASP.

More detailed information about our testing methods is contained in version 2.0 of the SecureQLab [Cloud WAF CyberRisk Validation Methodology](#) (AMTSO Test ID: AMTSO-LS1-TP054).

SecureQLab is a cybersecurity testing lab that was founded in 2019. SecureQLab works with enterprises, governments, and security vendors to bridge the applied intelligence gap that exists between market and technology research. SecureQLab also provides services to operationalize security and the metrics to help organizations improve their return on security investments.

Change

9 CONTACT INFORMATION

SecureQLab, LLC.
6001 W. Parmer Lane Ste 370, #970
Austin, TX 78704 USA

+1.512.575.3457

www.secureiqlab.com
info@secureiqlab.com

²⁵ <https://owasp.org/>

10 COPYRIGHT AND DISCLAIMER

This publication is Copyright © 2022 by SecureQLab®. Any use of the results, etc., in whole or in part, is ONLY permitted after the explicit written agreement of SecureQLab prior to any publication. SecureQLab cannot be held liable for any damage or loss, which might occur as result of, or in connection with, the use of the information provided in this paper. We take every possible care to ensure the correctness of the basic data, but a liability for the correctness of the research results cannot be taken by any representative of SecureQLab. We do not give any guarantee of the correctness, completeness, or suitability for a specific purpose of any of the information/content provided at any given time. No one else involved in creating, producing or delivering research results shall be liable for any indirect, special or consequential damage, or loss of profits, arising out of, or related to, the use or inability to use, the services provided by the website, research documents or any related data.

For more information about SecureQLab and the testing methodologies, please visit our website.

SecureQLab (November 2022)

- 1.1 Corrected tables and figures, added additional pricing info.
- 1.2 Corrected Imperva pricing and related figures and tables.
- 1.3 Typo correction.