



Image credit: [leowolfert](#)



William McKnight, Jake Dolezal
Oct 29, 2021

High Performance Application Security Testing v1.0

Product Evaluation: NGINX App Protect WAF vs.
ModSecurity (plus AWS Web Application Firewall and Azure
Web Application Firewall)

Security, Security & Risk

High Performance Application Security Testing

Product Evaluation: NGINX App Protect WAF vs. ModSecurity (plus AWS Web Application Firewall and Azure Web Application Firewall)

Table of Contents

- 1 Summary
- 2 API Security in the Cloud
- 3 GigaOm API Workload Test Setup
- 4 Test Results
- 5 Conclusion
- 6 Appendix: Recreating the Test
- 7 Disclaimer
- 8 About NGINX
- 9 About William McKnight
- 10 About Jake Dolezal
- 11 About GigaOm
- 12 Copyright

1. Summary

Data, web, and application security have evolved dramatically over the past few years. Just as new threats abound, the architecture of applications—how we build and deploy them—has changed. We've traded monolithic applications for microservices running in containers and communicating via application programming interfaces (APIs)—and all of it deployed through automated continuous integration/continuous deployment (CI/CD) pipelines. The frameworks we have established to build and deploy applications are optimized for time to market—yet security remains of utmost importance.

The challenge of securing and innovating is profound and requires a lightweight and integrated security solution that won't impede performance and delivery. For example, DevOps teams need security controls that work across distributed environments without invasively slowing down or burdening the release cycle. The maturation of these controls and processes ultimately transitions into the realm of DevSecOps, where security is built into the CI/CD pipeline.

The multitude of deployed apps, APIs, and microservices produces a constant flow of communication and data among applications that require active management—both internal and external. Apps themselves can vary greatly in the protocols, allowed methods, authorization/authentication schemes, and usage patterns. Perhaps most important, IT departments need granular control over the entire application ecosystem to prevent security breaches and attacks, be they man-in-the-middle, distributed denial of service, or script/code/SQL injection attacks.

While security is of utmost importance, the pace of modern business demands high performance, and this is especially true in application- and microservice-enabled enterprises. The conventional approach—deploying a perimeter Web Application Firewall (WAF) to protect applications by filtering and monitoring traffic between the app and the internet—is no longer enough. Even internal communication between apps and microservices on the trusted corporate network can be compromised and must be addressed. A defense-in-depth strategy is needed with multiple WAFs.

This report focuses on web application security mechanisms deployed in the cloud and closer to your apps. The cloud enables enterprises to differentiate and innovate with microservices at a rapid pace and allows microservice endpoints to be cloned and scaled in a matter of minutes. The cloud also offers elastic scalability compared to on-premises deployments, enabling faster server deployment and application development and less costly compute. However, the cloud is just as vulnerable, if not more so, to attacks and breaches as on-premises APIs and apps are.

Our focus is specifically on approaches to securing apps, APIs, and microservices that are tuned for

high performance and availability. We define “high performance” as companies that experience workloads of **more than 1,000 transactions per second (tps)** and require a **maximum latency below 30 milliseconds** across the landscape.

For many organizations, performance is a big deal—they need to ensure secured transactions at rates that keep pace with the speed of their business. A WAF or application security solution *cannot* be a performance bottleneck. Many of these companies seek a solution that can load balance across redundant microservices and enable high transaction volumes.

The numbers add up. If a business experiences 1,000 transactions per second, that translates into 3 *billion* API calls in a month. And it is not uncommon for large companies with high-end traffic levels to experience 10 billion or more API calls in a 30-day period. Make no mistake, performance is a critical factor when choosing an API security solution.

In this report, we performance tested security mechanisms on NGINX, AWS, and Azure: **ModSecurity**, **NGINX App Protect WAF**, **AWS Web Application Firewall (WAF)**, and **Azure WAF**. This last product was tested as a fully managed security offering. Note, ModSecurity is commercially distributed by NGINX and will be referred to as “ModSecurity” throughout the rest of this report.

In our benchmarks, NGINX App Protect WAF outperformed ModSecurity at all tested attack rates. NGINX App Protect WAF produced 4.7x lower latency than NGINX running ModSecurity at the 99th percentile at 1,000 transactions per second (tps) on the 5% bad request test. In our tests, the latencies for App Protect and ModSecurity diverged at the higher percentiles, becoming pronounced at the 95th percentile and above.

For fully managed offerings, NGINX App Protect WAF produced 128x lower latency than AWS WAF at 1,000 tps on the 5% bad request test at the 99th percentile. Also, NGINX App Protect WAF produced 82x lower latency than Azure WAF at 1,000 tps on the 5% bad request test at the 99th percentile. Since AWS and Azure’s WAF is fully managed, we do not know what underlying compute resources are working behind the scenes, which makes an apples-to-apples performance comparison difficult. Once again, latency differences were minimal until the 90th percentile, with a significant difference witnessed at the 99th percentile and above.

On a single small 2 CPU and 5.25GB of RAM EC2 instance, we captured the maximum transaction throughput achieved with 100% success (no 5xx or 429 errors) and less than 30ms maximum latency. NGINX App Protect WAF produced about 5,000 requests per second, compared to only 2,000

requests per second with ModSecurity. App Protect provides the same level of throughput as hitting the API directly without a WAF in between.

Testing hardware and software in the cloud is very challenging. Configurations may favor one vendor over another in feature availability, virtual machine processor generations, memory amounts, storage configurations for optimal input/output, network latencies, software and operating system versions, and the workload itself. Even more challenging is testing the fully managed, as-a-service offerings where the underlying configurations (processing power, memory, networking, and the like) are unknown. Our testing demonstrates a narrow slice of potential configurations and workloads.

As the sponsor of the report, NGINX opted for a default NGINX installation and API gateway configuration out of the box—the solution was not tuned or altered for performance. GigaOm selected identical hardware configurations for both App Protect and ModSecurity. The fully managed AWS WAF and Azure WAF were used “as-is,” since, by virtue of being fully managed, we have no access, visibility, or control over their infrastructure.

We leave the issue of fairness for the reader to determine. We strongly encourage you to look past marketing messages and discern for yourself what is of value. We hope this report is informative and helpful in uncovering some of the challenges and nuances of security architecture selection.

We have provided enough information in the report for anyone to reproduce this test. You are encouraged to compile your own representative workloads and test compatible configurations applicable to your requirements.

2. API Security in the Cloud

The landscape of API security in the cloud varies greatly based on an organization's need and underlying architecture. API security solutions offer either build-your-own or fully managed cloud deployment styles, and a few offer both. While there are many ways to secure APIs, we are interested specifically in enabling security while maintaining high performance. Again, for this report, we define **“high performance”** as companies that experience workloads of **more than 1,000 transactions per second** and need a **maximum latency below 30 milliseconds** across back-end APIs and microservices.

Furthermore, in terms of high performance, we focus particularly on latency results at the 99th percentile and above. At first glance, this might seem like an outlier case. However, in our experience, these measures are extremely important in latency results, which tend to be multi-modal over time, with the tops of the spikes representing “hiccups” in response times.

These hiccups matter. If the median response time or latency is less than 30 milliseconds, but there are “hiccups” with latencies above one second, the cumulative effect will impact subsequent user experiences. For example, if you visit a fast food drive-through where the median wait time for food is one minute, you probably think that was a good customer experience. However, what if the customer in front of you has a problem with their order, and it takes 10 minutes to resolve? Your wait time would actually be 11 minutes. Because your request came in line after the “hiccup,” the 99.99th percentile's delay becomes your delay too.

This report aims to explore vendor API security options to better support this high-end performance use case.

NGINX App Protect WAF

NGINX Open Source was first released in 2004 as a reverse proxy load balancer, web server, mail proxy, and HTTP cache. NGINX is offered as free-to-use, open-source software. Its popularity is evident in a March 2020 report by Netcraft¹ that found that 37% of all public websites use NGINX, compared to 24% that use Apache (usage of which peaked at 70% of all websites over a decade ago). NGINX was acquired by F5 Networks in 2019.

NGINX Plus is a commercial offering built atop NGINX Open Source and was first introduced in 2013. It quickly grew into a robust load balancer, Kubernetes Ingress controller, API gateway, and sidecar proxy solution, with hundreds of millions of instances deployed worldwide.

NGINX App Protect WAF combines the proven effectiveness of F5's advanced WAF technology with NGINX's agility and performance. App Protect runs natively on NGINX Plus to address security challenges facing modern DevOps environments.

For App Protect, we installed the latest version of Attack Signatures and Threat Campaigns to enable the highest level of protection for underlying applications and APIs.

ModSecurity

ModSecurity was developed in 2002 to protect the Apache HTTP web server. In 2010, ModSecurity was acquired by Trustwave SpiderLabs and released as an open-source, cross-platform web application firewall (WAF) module. Known as the "Swiss Army Knife" of WAFs, it enables web application defenders to gain visibility into HTTP traffic and provides a rules language and API to implement advanced protections.

For this test, we employed the OWASP ModSecurity Core Rule Set (CRS). CRS is a set of generic attack detection rules for use with ModSecurity and compatible web application firewalls. This ruleset aims to protect APIs from a wide range of attacks, including the OWASP Top Ten, with a minimum of false alerts.

Amazon Web Services Web Application Firewall

Amazon Web Application Firewall (WAF) is a fully managed cloud service that provides a web application firewall to help protect web applications and APIs against common web exploits that may affect availability, compromise security, or consume excessive resources.

For AWS WAF, we installed the `AWSManagedRulesCommonRuleSet`, which is also the OWASP implementation of the Core Rule Set for AWS WAF.

Azure Web Application Firewall

Azure Web Application Firewall (WAF) is also a fully managed cloud service that provides a WAF to help protect Azure back ends against the same exploits that can also impact service and application availability, compromise security, or consume excessive resources.

For AWS WAF, we installed the Azure OWASP 3.1 policy.

¹ <https://news.netcraft.com/archives/category/web-server-survey/>

3. GigaOm API Workload Test Setup

API Workload Test

The GigaOm API Workload Field Test is a simple workload designed to attack an API or an API management worker node (or a load balancer in front of a cluster of worker nodes) with a barrage of identical GET requests at a constant number of requests per second.

To perform the attacks, we used the HTTP load testing tool [Vegeta](#), a free-to-use workload test kit available on GitHub. The Vegeta tool returns a results bin file that contains the latencies and status code of every request. The attacker measured latency as the elapsed time between the points when an individual API request was made and when the API response back was received. Thus, if we tested 1,000 requests per second for 60 seconds, the attack tool recorded 60,000 latency values. We used that data to compile and interpret the results of the test.

The test also requires a back end API that can listen and respond to requests. In this case, our back end API listens for a GET request, such as:

```
http://ipaddress/
```

The API would respond with a string of 1024 pseudo random Unicode characters, such as:

```
taZ3psgHkQ...
```

Thus, for these tests, we used a request payload size of 1 KB.

The back-end API we used is further documented in the Appendix.

We completed three attempts per test on each platform, configuration, and request rate. We started with an attack rate of 1,000 requests per second (rps) and incremented up to attack rates of 2,000 rps, 3,000 rps, and 5,000 rps. We ran each test for 60 seconds. We captured the latencies at the 50th, 90th, 95th, 99th, 99.9th, 99.99th percentiles and the maximum latency seen during the test run. We recorded the test run that resulted in the lowest maximum latency or the highest success rate in the event of errors. Error status codes included HTTP status codes 429 Too Many Requests or any 5xx codes, most often 500 Internal Server Error. A success rate of 100% meant all requests returned a 200 OK status code.

In addition to standard attacks with “good” or expected traffic, we also tested the configuration’s ability to block “bad” traffic under load. For this metric, we tested these security mechanisms with 5% and 10% bad traffic. The “bad” traffic was sent as script injections, such as:

```
http://ipaddress/?a=<script>
```

The results are shared in the Field Test Results section.

Test Environments

Selecting and sizing the compute and storage for comparison can be challenging, particularly for fully managed as-a-service vendor offerings. The figures below give a visual layout of the configurations we tested.

Configurations Used for NGINX Comparable Tests

The first configuration we tested was single node workers of NGINX with no security. We installed Vegeta on an attack node and performed the API requests directly to the API worker, which routed the requests through to a back-end API. The back-end API would respond back to the API worker node, which routed the responses back to the attack node.

We installed all test components (excluding AWS WAF) onto AWS EC2 instances. For the components vital to the test, we used the “c5n” family of EC2 instances to take advantage of enhanced networking capabilities. [According to Amazon:](#)

C5n instances are ideal for high compute applications (including High-Performance Computing (HPC) workloads, data lakes, and network appliances such as firewalls and routers) that can take advantage of improved network throughput and packet rate performance. C5n instances offer up to 100 Gbps network bandwidth and increased memory over comparable C5 instances.






We also made a single operating system-level change to all NGINX instances. We increased the user limits for soft and hard open files to 65,536. The reason for this move is that during high response per second attack API testing, you can experience errors on worker nodes when user open file limits are hit.

Configurations Used for Comparable Tests of Fully Managed Platforms

For the fully managed security mechanisms (AWS WAF and Azure WAF), we were limited on the configurations we could test and still achieve a comparable result. The “instance type” of AWS WAF or Azure WAF is unknown—it is a fully-managed serverless platform, and we do not know its CPU or memory capabilities.

As shown in **Table 1**, we also installed the attack node, back-end APIs, NGINX instances, and the AWS and Azure test components within the same placement group to ensure the closest network proximity.

Table 1. EC2 Instances Used for Test Components

 Attack Node	C5n.2xlarge
 Backend API	C5n. large
 NGINX No Security	C5n.large
 NGINX Mod Security	C5n.large
 NGINX App Protect	C5n.large

Source: Gigaom 2021

Results may vary across different configurations, and again, you are encouraged to compile your own representative workloads and test compatible configurations applicable to your requirements.

4. Test Results

This section analyzes the latencies in milliseconds from the various 60-second runs of each of the scaled GigaOm API Workload Field Tests described above. Lower latency is better—meaning API responses are coming back faster. Also, the latency reveals the response time at the 50th, 90th, 95th, 99th, 99.9th, and 99.99th percentiles and the maximum latency. These are important values for service-level agreements (SLAs) and gauging the slowest response times a user might experience.

Test Results without AWS or Azure WAF

Here we show all of the results of the App Protect, Mod Security, and No Security runs for all tests that are completed.

Test Results with No Bad Traffic

For these tests, we did not send any bad traffic through in order to get a baseline of latency. Focusing on the 99th percentile, at the 1,000 rps level, Mod Security had 5.4 times the latency of App Protect.

At the 99th percentile, App Protect response times were 40% slower than with no security at all.

Figure 1 shows the results.

At 2,000 rps (**Figure 2**), ModSecurity response times were in the 5.5-second range, while App Protect was less than 1 millisecond. At 2,000 rps and above, Mod Security was not processing the traffic at 100% success, so it is not depicted on the charts in **Figures 3** and **4**. Meanwhile, App Protect and No Security were generally very close in latency at both the 3K rps and 5K rps levels.




Figure 1. Latency Curve at 1K rps




Figure 2. Latency Curve at 2K rps




Figure 3. Latency Curve at 3K rps




Figure 4. Latency Curve at 5K rps

Test Results with 5% Bad Traffic

With 5% bad traffic, Mod Security consistently produced higher latency times than App Protect did. At 500 rps of combined good and bad traffic, shown in **Figure 5**, Mod Security latency was 4.7 times higher than App Protect at the 99th percentile. We show charts for the good traffic (475 rps) and bad traffic (25 rps) below.




Figure 5. Latency Curve at 475 rps and 5% Bad Traffic




Figure 6. Latency Curve at 25 rps and 5% Bad Traffic (Bad Traffic Only)

Test Results with 10% Bad Traffic

With 10% bad traffic, Mod Security consistently continued to have more latency than NGINX App Protect WAF. At 500 rps of combined good and bad traffic, the latency gap was 4.3 times at the 99th percentile. We show charts for the good traffic (450 rps in **Figure 7**) and bad traffic (50 rps in **Figure 8**) below.




Figure 7. Latency Curve at 450 rps and 10% Bad Traffic




Figure 8. Latency Curve at 50 rps and 10% Bad Traffic (Bad Traffic Only)

Test Results with AWS WAF and Azure WAF

Test Results with No Bad Traffic

With no bad traffic, AWS WAF and Azure WAF produced much higher latency than App Protect, with the gap expanding at the 99.9th percentile for 1K rps (**Figure 9**), 2K rps (**Figure 10**), and 3K rps (**Figure 11**). At 5K rps, the difference was dramatic from the start, as shown in **Figure 12**.




Figure 9. Latency Curve versus AWS WAF at 1K rps




Figure 10 – Latency Curve versus AWS WAF at 2K rps




Figure 11. Latency Curve versus AWS WAF at 3K rps




Figure 12. Latency Curve versus AWS WAF at 5K rps

Test Results with 5% Bad Traffic

When bad traffic was introduced, AWS WAF and Azure WAF experienced more latency at the higher percentiles, with an expanding gap at the 99.99th percentile. **Figures 13 through 16** show the results.




Figure 13. Latency Curve versus AWS WAF at 475 rps and 5% Bad Traffic




Figure 14. Latency Curve versus AWS WAF at 25 rps and 5% Bad Traffic (Bad Traffic Only)

Test Results with 10% Bad Traffic

The 5% bad traffic pattern continued at 10% bad traffic.




Figure 15. Latency Curve versus AWS WAF at 450 rps and 10% Bad Traffic




Figure 16. Latency Curve versus AWS WAF at 50 rps and 10% Bad Traffic (Bad Traffic Only)

Finally, we highlight the maximum throughput achieved with 100% success and no 5xx or 429 errors and with less than 30ms maximum latency. See **Figure 17**.




Figure 17. Max Throughput (1KB payload)

The maximum transaction throughput achieved with 100% success (no 5xx or 429 errors) and with less than 30ms maximum latency with our tiny c5n.large (2 CPU and 5.25 GB RAM) instance was approximately 37,000 requests per second for NGINX with No Security. By comparison, ModSecurity began to produce errors at the 2,000 requests per second threshold. That is, above 2,000 rps, we started to receive 500 and 429 errors back from NGINX with ModSecurity, and above 19,000 rps, we started to receive 500 and 429 errors back from NGINX with App Protect.

5. Conclusion

This report outlines the results from a GigaOm API Workload Field Test.

NGINX App Protect WAF outperformed ModSecurity at all tested attack rates. NGINX App Protect WAF had 4.7x lower latency of NGINX running ModSecurity at the 99th percentile at 1,000 tps on the 5% bad request test. The latencies for App Protect and ModSecurity diverged at higher percentiles. Although the differences are minimal until you get to the 90th percentile, the difference in latency is pronounced at the 95th percentile and above.

For fully managed offerings, NGINX App Protect WAF had 128x lower latency than AWS WAF at 1,000 tps, on the 5% bad request test at the 99th percentile and 83x lower latency than Azure WAF. Since AWS WAF and Azure WAF are fully managed, we do not know what underlying compute resources are working behind the scenes, making it difficult to perform a true apples-to-apples comparison. Although the differences are minimal until you get to the 90th percentile, the difference in latency is pronounced at the 99th percentile and above.

On a single small 2 CPU and 5.25GB of RAM EC2 instance, the maximum transaction throughput was achieved with 100% success (no 5xx or 429 errors) and with less than 30ms maximum latency was approximately 19,000 rps for NGINX App Protect WAF, compared to only 2,000 rps for ModSecurity. App Protect in the test was found to provide the same level of throughput as No Security.

For this test using this particular workload with these particular configurations, API requests came back with the lowest latencies and highest throughput on NGINX App Protect WAF compared to any of the other tested security solutions.

Keep in mind, optimizations on all platforms would be possible as the offerings evolve or internal tests point to different configurations.

6. Appendix: Recreating the Test

The complete setup and execution of the Gigaom WAF Security Field Test is documented in the following GitHub repository:

<https://github.com/GigaOM/go-wafsecurity-test>

This repository contains the exact cloud infrastructure, operating system, dependencies, installation, and deployment steps we took. Feel free to replicate this test or modify it for your unique use case and needs.

The back-end API used in this test was a simple application developed by GigaOm that leverages the open-source NGINX.

The application works by binding the API application to port 1980 with NGINX listening for GET requests, such as:

```
GET http://fqdn-or-ip-address:1980/
```

The API would respond with 1024 pseudorandom Unicode characters from [/dev/urandom](#).

The following is the NGINX configuration for the back-end API. You are free to use and modify it at your own discretion. GigaOm makes no warranty or claim for its use beyond the scope of this test or report.

```
worker_processes auto;
worker_cpu_affinity auto;
error_log /var/log/nginx/error.log;
pid /run/nginx.pid;
worker_rlimit_nofile 20480;
events {
    accept_mutex off;
    worker_connections 10620;
}

http {
    access_log off;
    server_tokens off;
    keepalive_requests 10000;
    tcp_nodelay on;
```

```

server {
listen 1980 reuseport;
location / {
return 200
"JXvkE5pBpPN3T8bknNsqaM0kKu0j8BCV0S6TkNljlpDCYi8dIdn2TL11oHvliFkJAjj8VDnEcBoJSy73
QTuCcI8oeCna3jg34beyd7n5fZ22WSZP6gynF6PF5lMKsJTRRFRlur5trPpTU4nvzJOsbGY6O1bAoeCNT
G1VpDHZXQH67wZi35mNj6flLR3glKJwkwXzdrVgbeivVbT2fOz9zjxr0U8A4SONXYRyEr4jZzCqlYG4Eu
V08X4e5unvkO410ZeRrt3larys9hwr8tuCSi4a6KUsVeA5eZ74GQMv2NByz7R5BFCHIg0BbtexFsxdE9R
Zyj2sINlqbTQHNqwuiWDRG1CSJdOrTXYNmNz98Ib9BtAGMY7ikINWTeCaH8Qjet6wsDMyLbMjDfH3TjBT
eMJDVyLItqfY6MZbblEiEV0mNVBF1G0pn3s12EP0X7DzgIfSP6vU3jVdsuEWENja6DdWG0zciTAMbe4xw
RpyG0GWLsmoUoEVAOPsWPemthsLmjKO2WBQ9vUub2XV0IyO09vZKGajMaEZnXSqhb1RrKYcknK7Is2TIg
I6o6C0iIKEql1jhdJA15iFj4VytPftb9k8qbA5QE4dr2wcjWp8b0Rw9wBx9xYUDIkJO6IdrZqgR1APvAF
9UyokXgTkHtYycEC1QG0GSUhAT61FjGxtkZU86rV4djtt8zwJaKH7B126rSwvCVWYM82SRxZVJ2RkyQ3
xOaRM9DilXg4J90LSAlYu2TUpZpkym8Uk0qOsIWPr2e9jwLkonfdh2AqRX4QS9tCrvA2pfwLEptRNxsVL
KmNb2BJpt2YQ7K5OdYmW5oLwKTYtaB2sbCKQCGXWiieLfgt70gdumDsrBM8QslALQLZhX24rfadHvQ9sU
KUrW7KW3rkAhxJ1cvvU1up8NHza167KFLtFS8bJCb22cFL6L7shynseVS9a1YxYOSroaRDhz0WX4xdW7U
yJ4GrsqE9sXd66U8iAv78IaprC3M3HnJyieqyGzewvqSkAvhcnBKj";
}
}
}

```

7. Disclaimer

Performance is important but it is only one criterion for a Web Application Firewall selection. This test is a point-in-time check into specific performance. There are numerous other factors to consider in selection across Administration, Features and Functionality, Workload Management, User Interface, Scalability, Vendor, Reliability, and numerous other criteria. It is also our experience that performance changes over time and is competitively different for different workloads. Moreover, a performance leader can hit up against the point of diminishing returns and viable contenders can close the gap quickly.

GigaOm runs all of its performance tests to strict ethical standards. The results of the report are the objective results of the application of load tests to the simulations described in the report. The report clearly defines the selected criteria and process used to establish the field test. The report also clearly states the tools and workloads used. The reader is left to determine for themselves how to qualify the information for their individual needs. The report does not make any claim regarding the third-party certification and presents the objective results received from the application of the process to the criteria as described in the report. The report strictly measures performance and does not purport to evaluate other factors that potential customers may find relevant when making a purchase decision.

This is a sponsored report. NGINX chose the competitors and the test, and the NGINX Plus configuration was the default provisioned by NGINX Controller. GigaOm chose the most compatible configurations as-is out-of-the-box, and ran the testing workloads. Choosing compatible configurations is subject to judgment. We have attempted to describe our decisions fully in this report.

8. About NGINX

NGINX, acquired by F5 in 2019, is behind the popular open source project trusted by more than 450 million sites. NGINX offers a suite of technologies for developing and delivering modern applications. NGINX software solutions enable enterprises undergoing digital transformation to modernize legacy and monolithic applications as well as deliver new, microservices-based applications. Companies like Netflix, Starbucks, and McDonalds rely on NGINX to reduce costs, improve resiliency, and speed innovation. NGINX simplifies the journey to microservices. As enterprises move to a DevOps approach to application development and delivery, the tools, stack and interoperability of it all can get very complex. NGINX software reduces this complexity by consolidating common functions like load balancing, API management, security (App Protect), and service mesh down to far fewer components, to help make application infrastructure scalable and more manageable.

9. About William McKnight

William McKnight is a former Fortune 50 technology executive and database engineer. An Ernst & Young Entrepreneur of the Year finalist and frequent best practices judge, he helps enterprise clients with action plans, architectures, strategies, and technology tools to manage information.

Currently, William is an analyst for GigaOm Research who takes corporate information and turns it into a bottom-line-enhancing asset. He has worked with Dong Energy, France Telecom, Pfizer, Samba Bank, ScotiaBank, Teva Pharmaceuticals, and Verizon, among many others. William focuses on delivering business value and solving business problems utilizing proven approaches in information management.

10. About Jake Dolezal

Jake Dolezal is a contributing analyst at GigaOm. He has two decades of experience in the information management field, with expertise in analytics, data warehousing, master data management, data governance, business intelligence, statistics, data modeling and integration, and visualization. Jake has solved technical problems across a broad range of industries, including healthcare, education, government, manufacturing, engineering, hospitality, and restaurants. He has a doctorate in information management from Syracuse University.

11. About GigaOm

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.

12. Copyright

© [Knowingly, Inc.](#) 2021 "*High Performance Application Security Testing*" is a trademark of [Knowingly, Inc.](#). For permission to reproduce this report, please contact sales@gigaom.com.