



Report

Google Cloud Next Generation Firewall Enterprise Validation Report

Test Period: 26 February 2024 – 28 May 2024

Last Revision: 18 July 2024 | Commissioned by: Google

Report Contents

1.	Introduction	2
2.	Summary	3
2.1.	Policy Enforcement.....	3
2.2.	Stateful Inspection	3
2.3.	Security Efficacy.....	4
2.4.	TLS/SSL Support:	4
2.5.	Test Infrastructure.....	5
2.6.	Scenarios Tested & Validation Procedure	6
3.	Policy Enforcement	8
3.1.	Application-Level Inspection	8
3.2.	TOR Exit Node Blocking	9
3.3.	IP Reputation Validation.....	9
3.4.	Geo Based Blocking.....	9
3.5.	Granular Policy Enforcement	9
4.	Stateful Inspection	10
5.	Cloud Workload Security Efficacy Results	11
5.1.	Cloud-Hosted Application Threats	11
5.2.	Cloud Native Application Threats	12
5.3.	Post-Exploitation Activity	12
5.4.	APT Activity	12
6.	Operational Accuracy Validation.....	12
7.	Protocol Misuse Support Validation.....	13
8.	SSL Support Validation	14
9.	Operational Efficiency Validation.....	16
10.	Conclusion.....	18
11.	Contact Information.....	19
12.	Copyright and Disclaimer	19

1. Introduction




Figure 1. Google Cloud Next Generation Firewall Enterprise Security and Operational Summary Results

Cloud firewalls virtually shield cloud resources from external and internal threats. Cloud firewalls safeguard both web and non-web traffic for all users, applications, and locations.

This validation report evaluates Google Cloud Next Generation Firewall Enterprise, the advanced tier of Google Cloud Firewall Service. The Google Cloud Firewall leverages the Next-Generation Firewall (NGFW) capabilities, including intrusion prevention service (IPS) and transport layer security (TLS) inspection. Figure 1 highlights SecureQLab's overall findings.

Our testing covered the deployment and configuration of Google Cloud Next Generation Firewall Enterprise, including the evaluation of features like hierarchical firewall policies and dynamic workload scaling. The test scenarios used to validate Google's firewall examined cloud traffic flows between the enterprise workloads and from the internet.

Additionally, SecureQLab evaluated the firewall's operational efficiency, scalability, and performance while verifying that it secured Google Cloud workloads against various cyber threats, such as spyware, protocol misuse, data exfiltration, and command-and-control attacks.

Google commissioned these tests to measure their Firewall solution's overall security efficacy and key operational efficiency parameters. These tests were also intended to evaluate the relative resiliency of the firewall's TLS/SSL support & protection and the firewall's ability to protect workloads when well-known protocols are abused to evade detection.

In writing this report, SecureQLab has made extensive efforts to guarantee the accuracy of the results while straightforwardly presenting them. However, the test results are necessarily aggregated and simplified to present them in a summary format.

Google Cloud Next Generation Firewall Enterprise performed well in our testing. The test results discussed in this report highlight the overall organizational value of the Google Cloud Next Generation Firewall Enterprise in security, effectiveness, operational efficiency, and feature set.

2. Summary

Test Categories	Google Cloud Next Generation Firewall Enterprise Result
Policy Enforcement	Pass
Stateful Inspection	Pass
Cloud Workload Security Efficacy	99.2%
Ease of Deployment	100%
Analytics	High
Additional Operational Efficiency Metrics	86%
TLS/SSL Support	Pass

Table 1. Google Next Generation Firewall Enterprise Results

This 2024 SecureQLab Google Next Generation Firewall Enterprise CyberRisk Validation Report provides test results for the Google Cloud Next Generation Firewall Enterprise.

Table 1 summarizes the overall validation results of the Google Cloud Next Generation Firewall Enterprise for the 7 categories tested: policy enforcement, stateful inspection, cloud workload security efficacy, ease of deployment, analytics, additional operational efficiency metrics, and TLS/SSL support.

As can be seen from the table, the Google firewall generally performs well. One of its biggest strengths is deployment ease and workload security efficacy. There is an opportunity for improvement in some of the additional operational efficiency metrics such as firewall policy scheduling and security recommendations based on security incidents.

2.1. Policy Enforcement

Hierarchical firewall policies streamline firewall rule management in Google Cloud Next Generation Firewall Enterprise, ensuring consistent and centralized enforcement across the organization. They offer granular control, support advanced features like layer 7 inspection, and allow targeting specific network components for enhanced security.

Google Cloud Next Generation Firewall Enterprise passed all the different types of policy enforcement category testing while effectively managing critical rules and promoting a secure and compliant network environment. Details of these test categories are provided in section 3.

2.2. Stateful Inspection

Firewalls enforce security policies for traffic between the internal network and the internet. Stateful firewalls, such as Google Cloud Next Generation Firewall Enterprise, maintain a connection state table to track connections and provide enhanced security and logging compared to packet filtering firewalls.

SecureQLab tested the stateful inspection capability using the 5-tuple method. The test setup involved adding an egress firewall rule with Layer 7 inspection enabled for all protocols and ports, simulating malicious traffic, and observing the firewall's response. Google Cloud Next Generation Firewall Enterprise successfully blocked ongoing and new malicious connections, passing the stateful inspection tests.

2.3. Security Efficacy




Figure 2. Google Cloud Next Generation Firewall Enterprise Security Score

The security efficacy test stressed measuring the effectiveness of the firewall's next-generation capability. As such, this test focused on the firewall's intrusion prevention system. The attacks launched consisted of exploits targeting lift-and-shift applications, cloud-native application-centric exploits, post-exploitation activity, and advanced persistent threat (APT) activity.

As seen in Figure 2, the overall security efficacy score of the Google Cloud Next Generation Firewall Enterprise was 99.2%, which is exceptional. Details regarding Google's *Cloud Workload Security Efficacy* are available in section 5.

2.4. TLS/SSL Support:

The Google Cloud Next Generation Firewall Enterprise fared well in terms of supporting the standard cipher suites supported by TLSv1 through TLSv1.3 and was able to scale effectively under load. The Google Cloud Next Generation Firewall Enterprise passed all the Ciphers tested that it supported.

2.5. Test Infrastructure

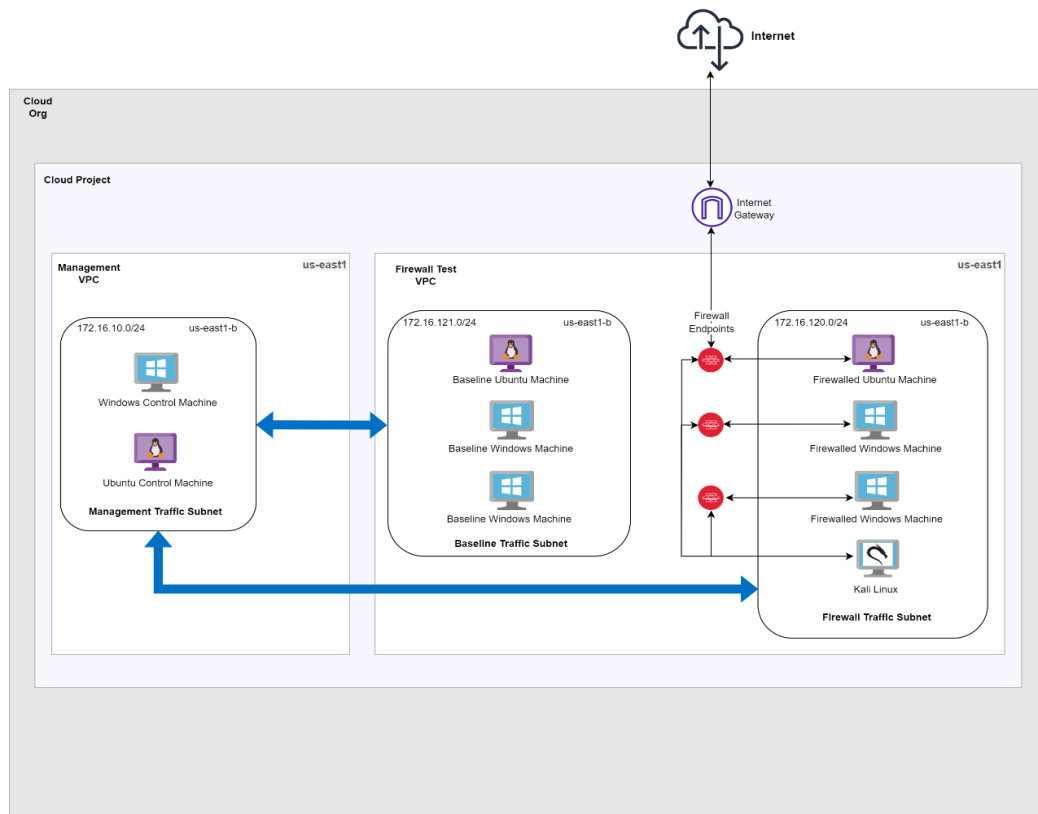


Figure 3. Test Architecture

The test architecture was created and deployed within the Google Cloud. Figure 3 provides a diagram of the test architecture. Table 2 highlights certain components of the infrastructure used during testing. Figure 3 above provides an overview of SecureQLab's test architecture. SecureQLab ensured that the best practice configurations as prescribed by Google Cloud were followed around the following components:

- Firewall Policy
- Firewall Endpoint
- Security Profile and Security profile groups
- TLS inspection policy
- VM Instances
- All VM instances are based on n2-standard-4 with 4vCPUs and 16GB of RAM

Operating Systems	Number of Virtual Machines
Ubuntu 20.04 LTS	3
Windows 10 x64 Enterprise Edition	5
Kali Linux (SSL Scaling Test)	4
Kali Linux (Attacker Machine)	1

Table 2. Infrastructure Used During Test

Publicly available best-practice documentation was referred to confirm that the firewall was at least minimally configured to best-practice specifications for all security features/modules (“best-practice or better”). For instance, the default policy for “alert only” configurations was modified to “alert and deny” to ensure the Google firewall was blocking bad activity. Because product performance is generally highly configuration-dependent, using different settings might yield different results. True positive testing was then performed to confirm the functionality of all configured security policies.

False positive testing was also performed to conservatively tune the policies to what would be appropriate for a typical enterprise. For example, the ability to browse and render general popular websites (e.g., Amazon, Bing, CNN, MSNBC, and Wikipedia) was tested.

The test architecture was designed to allow evaluation of key traffic flows that a cloud workload will be subjected to. 2 virtual private clouds (VPCs) were deployed; one of those is used for management traffic to collate the test results, and the other VPC has two subnets. One subnet has untagged Virtual Machines with no L7 inspection enabled, and the other has tagged virtual machines with L7 inspection enabled and tied to a security profile with the default options. This paired VPC setup allowed SecureQLab to separate the management traffic and prevent it from tainting the attack traffic. Additionally, multiple VPCs were used during deployment to check that the Google firewall’s org-level policy enforcement covers all the projects and VPC networks within the organization, ensuring comprehensive security coverage across all assets and environments.

The attacker machine was placed on the same subnet as the tagged Virtual Machines, and insider attacks were carried out by the attacker machine. Tests of known vulnerabilities like browser exploits, obfuscated JavaScript, cross-site scripting, and advanced threats were designed to examine the firewall’s ability to inspect traffic between workloads.

2.6. Scenarios Tested & Validation Procedure

A cloud environment can include traffic that is both bi-directional and unidirectional. Such traffic may further contain mixtures of applications and protocols coming in and out of VPC, project, and organizational levels. Such traffic may originate from a protected cloud environment to the outside.

In this test, the emphasis was on protecting cloud workloads. Hence, the tests were centered around evaluating the firewall’s ability to do the same. Benign traffic was performed on cloud instances, such as environment provisioning, troubleshooting, routine operations such as browsing, and connectivity checks.

The term “cloud workloads” describes the different jobs, procedures, or software programs utilized by cloud computing infrastructure. Application development, machine learning, artificial intelligence, web hosting, data processing, analysis, storage, and more can all be included in these workloads. The basis of these workloads is the ability to provision and scale virtual machines and virtual networks.

The test cases used to assess the Google Cloud fully distributed firewall emphasize various traffic flows between the workloads and from the internet.

The scenarios tested are:

Policy Enforcement: At the most basic level, a cloud firewall should be able to enforce a firewall policy that would restrict or allow a particular traffic flow based on the configured firewall rules. Additionally, the Google Cloud Next Generation Firewall Enterprise also inspects traffic on specific tagged resources. This test was performed on the Google Cloud Next Generation Firewall Enterprise to ensure that it effectively enforced firewall policies, allows or restricts traffic flows according to configured rules, and inspects traffic on specific tagged resources as specified. Policy enforcement was validated at the organizational, project, and VPC levels.

Traffic Inspection: As a first step in traffic inspection evaluation, SecureQLab considered the traffic flows between the workloads within the network because it is a lucrative target for insider attacks. During the test, different types of application layer traffic were run between the workloads to evaluate the L7 traffic inspection capabilities of the Google Cloud Next Generation Firewall Enterprise solution.

Vulnerability Exploits: Exploits and browser-based vulnerabilities were used to evaluate the feasibility of obtaining unauthorized access to the system. An internal attacker machine and various vulnerable versions of Google Chrome and Mozilla Firefox were used to execute the attacks.

Obfuscated JavaScript Exploits: Attackers use obfuscated JavaScript to make the malicious JavaScript code more difficult to detect and analyze by security tools. Hence, the test included a set of payloads that were obfuscated to evaluate the firewall's ability to identify and block malicious obfuscated script payloads.

Cloud Native Workload: Cloud Native Workload is an abstraction of four key components. They are development process (DevOps), Application Architecture (Microservices), deployment & packaging (Containerization), and Application infrastructure (Cloud). SecureQLab tested cloud-native workloads relying upon development applications packaged in containers. SecureQLab tested a container self-hosted in the Google Cloud platform rather than using a managed container service because self-hosted containers are more prevalent than cloud-provided container services.

Protocol Misuse Techniques: Once the initial compromise of the machine is complete, an attacker looks for ways to establish persistence in the network by looking for ways to establish interrupted access to the workload; this is particularly problematic in case of insider threats because of the leniency of the firewall rules inside the network. These sets of tests evaluated the firewall's capabilities in identifying such connections that use advanced evasion techniques like http protocol misuse, payload encryption, bash scripts, and TCP misuse.

Advanced Persistent Threats: APTs are another set of attacks that impact cloud workloads. These attacks are usually carried out from external networks. A set of tests was run with popular APT tools like Covenant and PowerShell Empire to test the firewall's ability to identify and block ongoing communications with the C2C server.

Post-Exploitation Techniques: Post-Exploitation Techniques aim to exfiltrate the data from the target cloud workloads. To counter this, behaviors that indicate a potential exfiltration must be identified and stopped by the cloud firewall.

False Positives: False positive testing was also performed to conservatively tune the policies to what would be appropriate/acceptable for a typical enterprise. For example, the ability to browse and render general popular websites (e.g., Amazon, Bing, CNN, MSNBC, and Wikipedia) was tested.

3. Policy Enforcement

Policy Test Category	Policy Enforcement Criteria	Enforcement Direction	Result
Application-Level Inspection Policy	L7 ingress inspection	External -> Internal	Pass
	L7 egress inspection	Internal -> External	Pass
TOR Exit Node Blocking Policy	TOR Exit nodes	External -> Internal	Pass
IP Reputation Validation Policy	Known Malicious IPs	External -> Internal	Pass
	Known Malicious IPs	Internal -> External	Pass
Geo-based Protection Policy	Geo-location-based blocking	External -> Internal	Pass

Table 3. Policy Enforcement Detailed Results

A new approach to cloud workload security is required to counter the emergence of advanced attack vectors, attack strategies, and new cloud attack surfaces. This approach must be designed to thwart complex and customized threats. As firewalls evolved to protect against more advanced threats, they naturally moved up the stack. In turn, Firewalls are evolving to respond. In addition to being "Layer 7 aware," firewalls have added capabilities to inspect packet content and search for Indicators of Attack (IOA) within the data stream.

As attacks become more complex, detection and protection capabilities with lower false positive rates become increasingly important, and cloud firewall policy-based enforcement is the foundation of securing cloud workloads.

The policy enforcement implemented on Google Cloud Next Generation Firewall Enterprise Solution can be applied to a specific project or at the org level within the Google Cloud platform. If the policy is enforced at the org level, then the policy rule will be applied to all nodes that are part of the target environment.

For our testing, we initially considered the firewall policy association at the Virtual Private Cloud (VPC) level, where SecureQLab tagged the cloud workloads with specific network tags and secure tags to test the firewall's ability to apply specific rules to those tags. Furthermore, we configured an org-level policy to include best practice rules, which are then inherited by all the projects under the org.

The test setup involved two scenarios; first, the org-level policy was configured in such a way that it applies to all the projects and VPC networks in the org. Second, the org-level policy was configured to be applied to specific projects and networks; a test simulation was run to check the granular control that the org-level policy provides to configure the firewall.

Table 3 shows that Google successfully provided policy enforcement.

3.1. Application-Level Inspection

Application-level traffic inspection capabilities of the firewall were evaluated to test the firewall's ability to detect and block application layer-based attacks. Various tests are run across multiple categories, such as Malware download over

HTTPS, Browser exploits, cross-site scripting attacks, Obfuscated JavaScript attacks, advanced evasion techniques, HTTP evasion, and post-exploitation.

This test validated layer 7 inspection capabilities which included vectors that a typical user usually interacts with. Google Cloud Next Generation Firewall Enterprise successfully defended against cyber threats that target layer 7 in the OSI reference model and passed the evaluation.

3.2. TOR Exit Node Blocking

The Google Cloud's Threat Intelligence has a default firewall rule that is deployed when a new firewall policy is created; one of the rules in the policy is aimed at blocking connections from the TOR exit nodes. During the test period, we looked for any connection attempts to the firewalled VMs that would be blocked by this rule. The results show that of 1369 such attempts, the firewall was able to block all the connections (100%).

3.3. IP Reputation Validation

For cloud workloads, IP blocking must be implemented to improve security. Organizations can lower their risk of malware infections, data breaches, and other security problems by restricting access to particular IP addresses or ranges. IP blocking enforces access control policies and protects sensitive data, which assists enterprises in meeting regulatory compliance requirements. Additionally, reducing the routes that attackers can exploit lowers the attack surface of cloud systems. Furthermore, IP blocking allows businesses to personalize access control policies, efficiently monitor incoming traffic, and react quickly to security events.

This test validated IP reputation inspection capabilities. Google Cloud Next Generation Firewall Enterprise successfully defended against IP addresses that exhibit malicious behavior by hosting suspicious content leveraged by cyber threats passing the evaluation.

3.4. Geo Based Blocking

Geo-location filtering for cloud workloads is essential because it helps organizations control where their data is accessed and stored, ensuring compliance with regulations and laws governing data privacy and sovereignty. By restricting access based on geographic location, organizations can reduce the risk of security breaches from known high-risk regions and optimize performance by directing traffic to nearby servers. It also allows customized access control, enabling organizations to tailor security policies to specific areas or business requirements. This is a default rule that is part of the firewall policy setup and is aimed at blocking the connections from and to a specific IP address based on Geo location.

We modified the set of locations in this rule and checked if a connection can be made to the IP addresses in the blocked geo-location and found that the firewall has "PASSED" this test.

3.5. Granular Policy Enforcement

SecureQLab tested the ease of deployment and easy policy push that can be applied at different levels (for example, organization, project, and VPC). The test result was that the Google Cloud Next Generation Firewall Enterprise provided a seamless experience. The hierarchical nature of the policy helps users easily maintain a good security posture and reduces the additional overhead activity that can directly contribute towards better operational efficiency.

4. Stateful Inspection

Test Categories	Google Cloud Next Generation Firewall Enterprise Result
Stateful Inspection	Pass

Table 4. Stateful Inspection Results

Firewalls are designed to enforce security policies. Firewalls enforce security policies for all traffic traversing the internal network to the internet and vice-versa. A stateful firewall enforces security policy by creating a directory of the connections (E.g.TCP connections). A connection state table is maintained by keeping track of the source address, source port, destination address, destination address, destination port, and connection state.

A stateful firewall provides better security than a packet filtering firewall by providing more logging functionality and is more robust to attacks and exploits that take advantage of TCP/IP. Google Cloud Next Generation Firewall Enterprise is a stateful firewall.

SecureQLab tested stateful inspection capability via the 5-tuple method. The outcome of the stateful inspection test demonstrated Google Firewall's capability.

The test setup for stateful inspection involved adding an egress firewall rule with L7 inspection enabled to all protocols and ports without any explicit ingress rule. Given the statefulness of the firewall, irrespective of having an explicit ingress rule, SecureQLab expected that the firewall would block malicious connections.

To test this scenario, initially, we modified the security group profile to allow all kinds of traffic irrespective of the criticality level. Then, we executed the PowerShell Empire command line tool to simulate malicious traffic to the tagged virtual machine guarded by the firewall. Once we were sure that the connection was established and we could run the commands on the target machine, we changed the security profile rules to their defaults and applied the update. After the update was completed, the ongoing communication between the victim machine and the attacker machine was RESET by the firewall, and any new attempts to make a connection were also actively blocked by the firewall.

As demonstrated in Table 4, Google Cloud Firewall successfully passed the Stateful Inspection tests.

5. Cloud Workload Security Efficacy Results




Figure 4. Google Cloud Next Generation Firewall Enterprise Security Efficacy

Cloud firewalls should be designed to protect cloud-based resources and applications.

The Google Cloud Next Generation Firewall Enterprise was tested against 160 attack types within four threat categories. Figure 4 above presents an overview of the SecureQLab findings during the security effectiveness validation and reporting of the Google Cloud Next Generation Firewall Enterprise. This includes the cloud-hosted application, cloud-native applications, post-exploitation, and APT activity score.

5.1. Cloud-Hosted Application Threats

Test Categories	Google Cloud Next Generation Firewall Enterprise Result
Exploits	100.0%
Plugin Control	100.0%
Active Threats	100.0%
Script Obfuscation	100.0%

Table 5. Cloud-Hosted Application Threats

The cloud-hosted application threat average scores are calculated by averaging the scores for each threat type together within their respective attack category. This section included attacks targeted to applications, the ability to provide plugin control, active threats used in current cyberattacks, and scripts designed to circumvent detection.

Table 5 shows how Google scored 100% across all test categories for the Cloud Hosted Application Threats.

5.2. Cloud Native Application Threats

Test Categories	Google Cloud Next Generation Firewall Enterprise Result
Cross-Site Scripting in Application	100%
Application Delivered via Containers	87.5%

Table 6. Cloud Native Application Threats

Cloud firewalls must extend their protective capabilities to counter cloud-native application threats. A robust cloud firewall should possess threat detection capabilities, be able to identify suspicious network traffic patterns and have the capacity to block in real time.

The Google Cloud Next Generation Firewall Enterprise was tested against two attack categories classified as cloud-native application threats. Table 6 provides the results from these tests.

The Cloud-Native Application Threat Score for the Google Cloud Next Generation Firewall Enterprise is calculated from the percentage of the total blocked attacks to the total number of attacks for both test categories. **Missing coverage has been identified and addressed.**

5.3. Post-Exploitation Activity

Test Categories	Google Cloud Next Generation Firewall Enterprise Result
Post-Exploitation Activity Blocked	100.0%

Table 7. Post-Exploitation Activity

The term post-exploitation refers to any activity performed by threats after the successful exploitation of a resource. While there are various activities that threats could choose to perform, such as data exfiltration, data destruction, and infrastructure lockdown, these activities are performed via various well-known protocols. However, this malicious activity using these protocols is difficult to detect because it looks benign.

The post-exploitation activity average score for the Google Cloud Next Generation Firewall Enterprise is calculated by averaging the scores for each threat type within their respective attack category. Table 7 showcases that the Google Cloud Next Generation Firewall Enterprise scored 100% in this category by detecting the malicious activity

5.4. APT Activity

Test Categories	Google Cloud Next Generation Firewall Enterprise Result
APT Activity Blocked	100%

Table 8. APT Activity Blocked

APT, an acronym for advanced persistent threat, is the category of attack that relies upon novel methods that use apparently benign activity to penetrate an organization's infrastructure; SecureQLab used PowerShell Empire and Covenant to run APT attack scenarios.

The cloud APT activity score for the Google Cloud Next Generation Firewall Enterprise is calculated by averaging the scores for each threat type within their respective attack category. Table 8 highlights Google blocked 100% of the APT Activity.

6. Operational Accuracy Validation

Operational Accuracy Test	Google Cloud Next Generation Firewall Enterprise Result
Resistance to False Positives	100.0%

Table 9. Operational Accuracy

False positive testing was included in the scope of this test because a Cloud Firewall that prevents 100% of malicious attacks but also prevents legitimate (non-malicious) actions can be vastly disruptive. SecureQLab used appropriate tools and techniques to ensure that the tested firewall products do not generate significant numbers of alerts with legitimate applications and processes in an enterprise environment. This section of the methodology was performed simultaneously and as part of the validation workflow wherever feasible. It was also performed in other independent sections. The goal of this test was to ensure that the firewall products do not prevent malicious traffic at the expense of operational accuracy.

The Google Cloud Next Generation Firewall Enterprise was tested for operational accuracy under real-world scenarios during the entire test cycle, and Table 9 above provides the results from these tests.

Note: For the false positive tests to be run on the machines with a firewall endpoint attached and TLS inspection enabled, SecureQLab had to forward the system time to 30 minutes to prevent SSL certificate-related errors.

7. Protocol Misuse Support Validation

Test Type	Google Cloud Next Generation Firewall Enterprise Result
HTTP Protocol Misuse	Pass
TCP Protocol Misuse	Pass

Table 10. Protocol Misuse Results

Attackers commonly abuse legitimate network protocols. SecureQLab used HTTP and TCP to test for protocol misuse. TCP is the most common transport protocol, and HTTP is the most common application layer protocol. HTTP is so complicated that different implementations of the protocol behave differently when they receive certain HTTP replies. These interpretation discrepancies can be exploited to get beyond a firewall's protection if it acts differently than the browser it is meant to guard.

The Protocol Misuse tests were designed to see how browsers behave when encountering unusual or invalid HTTP replies. The test shows how the presence of a firewall or proxy between the browser and the test server influences the outcomes and whether it is feasible to circumvent the security. Test results show that the firewall can identify HTTP protocol misuse payloads and generate alerts for the system administrator to review.

Google Cloud Next Generation Firewall Enterprise was able to alert on all the payloads delivered via HTTP. SecureQLab used different segment sizes, differing amounts of time to wait before sending out data after a connection was established, and the XOR mechanism to test TCP protocol abuse. Table 10 demonstrates that the Google Cloud Next Generation Firewall Enterprise was able to alert on all these TCP protocol abuses.

8. SSL Support Validation

SSL/TLS Threat Efficacy Test	Google Cloud Next Generation Firewall Enterprise Result
Overall SSL/TLS Security Efficacy	Pass

Table 11. SSL/TLS Threat Efficacy

Hypertext Transfer Protocol (HTTP) and its secure analogue HTTPS have long been crucial internet protocols. HTTPS uses encryption to enhance browsing safety by establishing secure connections between browsers, applications, and websites. To identify the attacks or threats in encrypted connections, the firewall must inspect the encrypted traffic using SSL/TLS ciphers and the associated techniques for managing encrypted connections. To detect a threat concealed within encrypted data, the firewall must decrypt packets, inspect the content, and take necessary action.

SecureQLab tested 10 of the TLS v1.2 ciphers, and 3 of TLS v1.3 ciphers. The testing included combinations of ciphers between clients and servers to analyze firewall behavior with weak ciphers. This was done to assess how the firewall dealt with communication using different ciphers and to evaluate the ability of the Google Cloud Next Generation Firewall Enterprise to fall back or enforce secure ciphers during communication.

The Google Cloud Next Generation Firewall Enterprise was tested for overall SSL/TLS threat efficacy and SSL/TLS Scaling Capabilities under load and its ability to protect against attacks delivered through the supported ciphers in real-world scenarios. These results are showcased in Table 11 above.

Google Cloud Next Generation Firewall Enterprise was tested against the cipher suites for TLS v1.2, as highlighted in Table 12. This test scenario was conducted with attack traffic that consisted of a variant of the Mirai malware file that the firewall had already identified as malicious.

TLSv1.2 Ciphers	TLS Support	TLS Handshake Result	Attack Traffic Prevention within TLS Session	Operational Validation with Operational TLS Traffic
ECDHE-ECDSA-CHACHA20-POLY1305	Yes	Successful	Pass	Pass
ECDHE-RSA-CHACHA20-POLY1305	Yes	Successful	Pass	Pass
ECDHE-ECDSA-AES128-GCM-SHA256	Yes	Successful	Pass	Pass
ECDHE-RSA-AES128-GCM-SHA256	Yes	Successful	Pass	Pass
ECDHE-ECDSA-AES256-GCM-SHA384	Yes	Successful	Pass	Pass
ECDHE-RSA-AES256-GCM-SHA384	Yes	Successful	Pass	Pass
AES128-GCM-SHA256	Yes	Successful	Pass	Pass
AES256-GCM-SHA384	Yes	Successful	Pass	Pass
AES128-SHA	Yes	Successful	Pass	Pass
AES256-SHA	Yes	Successful	Pass	Pass

Table 12. TLS v1.2 Cipher Support

The Google Cloud Next Generation Firewall Enterprise passed all TLS v1.2 ciphers tested that it supported. As can be seen from the above, the Google Cloud Next Generation Firewall Enterprise was able to successfully handle packet decryption and inspection.

The cipher suites for TLS v1.3, as highlighted in Table 13 below, were tested for the Google Cloud Next Generation Firewall Enterprise.

TLSv1.3 Ciphers	TLS Support	TLS Handshake Result	Attack Traffic Prevention within TLS Session	Operational Validation with Operational TLS Traffic
AES-256-GCM-SHA384	Yes	Successful	Pass	Pass
CHACHA20-POLY1305-SHA256	Yes	Successful	Pass	Pass
AES-128-GCM-SHA384	Yes	Successful	Pass	Pass

Table 13. TLS v1.3 Cipher Support

The Google Cloud Next Generation Firewall Enterprise passed all TLS v1.3 ciphers tested that it supported. As can be seen from the above, the Google Cloud Next Generation Firewall Enterprise was able to successfully handle packet decryption and inspection.

The cipher suites for TLS v1.1, as highlighted in Table 14 below, were tested for the Google Cloud Next Generation Firewall Enterprise.

TLSv1/TLSv1.1 Ciphers	TLS Support	TLS Handshake Result	Attack Traffic Prevention within TLS Session	Operational Validation with Operational TLS Traffic
ECDHE-ECDSA-AES256-SHA	Yes	Successful	Pass	Pass
ECDHE-RSA-AES256-SHA	Yes	Successful	Pass	Pass
ECDHE-ECDSA-AES128-SHA	Yes	Successful	Pass	Pass
ECDHE-RSA-AES128-SHA	Yes	Successful	Pass	Pass
AES256-SHA	Yes	Successful	Pass	Pass
AES128-SHA	Yes	Successful	Pass	Pass

Table 14. TLSv1/TLS v1.1 cipher Support

The Google Cloud Next Generation Firewall Enterprise passed all TLS v1/v1.1 ciphers tested that it supported. As can be seen from the above, it was able to handle packet decryption and inspection successfully.

9. Operational Efficiency Validation




Figure 5. Google Operational Efficiency Comparison with Group Average from Recent Public Test

Google Cloud Next Generation Firewall Enterprise's operational efficiency measures the tested operating burden and complexity of setup and use. As such, the Operational Efficiency Score measures both the ability of the Google Cloud Next Generation Firewall Enterprise to detect and respond to cyber-attacks appropriately and ease of use. The operational efficiency was evaluated by considering factors such as:

- The ease of tuning the Google Cloud Next Generation Firewall Enterprise security policy and configuration (pre-and-post deployment).
- The solution's incident response and management intuitiveness from a policy and security configuration perspective.
- Customer support.
- Enhanced security metrics reporting capabilities.
- The ease of managing and controlling assets and business continuity with appropriate configuration and policy backup (with restoration).

In the analysis, the Google Cloud Next Generation Firewall Enterprise was rated high, medium, or low across Ease of Deployment (Figure 6), Analytics (Figure 7), and six additional operational efficiency categories, as identified in Table 15 below and Figure 9. For more details on each of the categories, please [contact SecureQLab](#).

The Google Cloud Next Generation Firewall Enterprise was tested for operational efficiency, considering real-world enterprise procurement, deployment, and active scenarios throughout the test life cycle. Figure 5 above presents an impressive overall operational efficiency score of 92%.




Figure 6. Ease of Deployment




Figure 7. Analytics




Figure 8. Additional Operational Efficiency Metrics

The Google Cloud Next Generation Firewall Enterprise had high operational efficiency capabilities in the six validated categories of operational efficiency in addition to Ease of Deployment and Analytics, as highlighted in Table 14 below.

Table 15 shows that the Google Cloud Next Generation Firewall Enterprise performed better around Security Metrics reporting, Security Policy Configuration, Security Policy management, Compliance management and Customer Support.

SecureQLab validated Google Cloud Next Generation Firewall Enterprise operational efficiency in six areas of validation with a total of 20 features and functions validated.

The features and functions within each category are awarded scores based on their capabilities. These scores are then tallied together to form a rating of high, med, or low. The Operational Efficiency Rating is equal to the total number of points scored respectively by the Google Cloud Next Generation Firewall Enterprise operational efficiency validation over the total number of points. Category Scores were calculated by aggregating earned points and then dividing this number by the total number of possible points to find a percentage. Points (integers 0 – 10) are earned for each feature within a category as follows:

- High or Yes = 10 Points
- Medium = 6 Points
- Low = 3 Points
- NA/No = 0 Points

Excellent operational efficiency can increase security by decreasing complexity.

Operational Efficiency Metric	Test Case	Google Cloud
Security Policy Configuration	Recommend default out-of-the-box policy	Yes
	Pre-canned security profiles	Yes
	Grouped firewall policies	Yes
	Speed to push Security Profile settings	Medium
	Speed to push new configuration	High
	Number of policies that can be created	High
Security Policy Management	Clonable Security Policy	Yes
	Broad support for hierarchical management	Yes
	API managed Security Policy	Yes
	Centralized management platform	Yes
	Global security management	Yes
	Security recommendations based on incidents	No
Asset Management	Firewall policy scheduling	No
	Effectiveness of firewall asset visibility	Medium
Compliance Management	Can the firewall map protected assets	Yes
	User audit trail support	Yes
Security Metrics Reporting	Policy audit trail	Yes
	24-hr security dashboard visibility	High
Customer Support	Security dashboard event details	Yes
	Efficiency of support	High

Table 15. Google Additional Operational Efficiency Details

When Google Cloud Next Generation Firewall Enterprise is compared in operational efficiency metric against the recently published Advanced Cloud Firewall Comparative report of the group average, Google Cloud Next Generation Firewall Enterprise's score is substantially higher than the group average.

Furthermore, in areas such as Security Policy Configuration and Asset Management, the Google Cloud Next Generation Firewall Enterprise demonstrated better efficiency than many of the established players in the marketplace. Google Cloud Next Generation Firewall Enterprise also performed better in areas such as Security Policy Management and Customer Support than other IaaS provider offerings.

The outstanding test results in the operational efficiency metric indicate Google's commitment to simplifying the complex and delivering operationally efficient products in the marketplace.

10. Conclusion

Overall test results show that the Google Cloud Next Generation Firewall Enterprise can detect attacks carried out on workloads when exploiting known vulnerabilities in the workload.

The default security policy provided by Google Cloud Next Generation Firewall Enterprise is an excellent starting point for creating more complex firewall rules. In other words, the default firewall rules could be adequate for businesses that don't have complex user, application, or device traffic if the firewall rules are used with appropriate enforcement mechanisms such as alert or block. For more dynamic scenarios, the default firewall rules provided by the firewall are a good starting point from users can leverage Google Cloud's Threat intelligence to block connections from Malicious IPs and Geo locations.

Finally, when experimenting with changes to security profile settings, the average update time of the security profile options was between 5 and 15 minutes. Considering the distributed nature of the firewall and the asynchronous updates, the delay is within an acceptable limit.

11. Contact Information

SecureQLab, LLC.
9600 Great Hills Trail Suite #150W
Austin, TX 78759 USA
+1.512.575.3457
www.secureqlab.com
info@secureqlab.com

12. Copyright and Disclaimer

Copyright © 2024 SecureQLab, LLC. All rights reserved. The content of this report is protected by United States and international copyright laws and treaties. You may only use this report for your personal, non-commercial, informational purposes. Without SecureQLab's prior written consent, you may not: (i) reproduce, modify, adapt, create derivative works from, publicly perform, publicly display, or distribute this report; or (ii) use this report, the SecureQLab name, or any SecureQLab trademark or logo as part of any marketing, promotion, or sales activities. THIS REPORT IS PROVIDED "AS IS," "AS AVAILABLE" AND "WITH ALL FAULTS." TO THE MAXIMUM EXTENT PERMITTED BY LAW, SECUREQLAB EXPRESSLY DISCLAIMS ALL WARRANTIES AND REPRESENTATIONS, EXPRESS OR IMPLIED, INCLUDING: (a) THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE; AND (b) ANY WARRANTY WITH RESPECT TO THE QUALITY, ACCURACY, CURRENCY OR COMPLETENESS OF THE REPORT, OR THAT USE OF THE REPORT WILL BE ERROR-FREE, UNINTERRUPTED, FREE FROM OTHER FAILURES OR WILL MEET YOUR REQUIREMENTS. WITHOUT LIMITING THE GENERALITY OF THE FOREGOING SENTENCE, YOU ACKNOWLEDGE AND AGREE THAT THE QUALITY, ACCURACY, CURRENCY AND COMPLETENESS OF THE REPORT DEPEND UPON VARIOUS FACTORS, INCLUDING FACTORS OUTSIDE OF SECUREQLAB'S CONTROL, SUCH AS: (1) THE QUALITY, ACCURACY, CURRENCY OR COMPLETENESS OF INFORMATION AND MATERIALS PROVIDED BY OTHER PARTIES THAT ARE RELIED UPON BY SECUREQLAB IN PERFORMING PREPARING THE REPORT; AND (2) THE UNDERLYING ASSUMPTIONS MADE BY SECUREQLAB IN PREPARING THE REPORT REMAINING TRUE AND ACCURATE. YOU ARE SOLELY RESPONSIBLE FOR INDEPENDENTLY ASSESSING THE QUALITY, ACCURACY, CURRENCY AND COMPLETENESS OF THE REPORT BEFORE TAKING OR OMITTING ANY ACTION BASED UPON THE REPORT. IN NO EVENT WILL SECUREQLAB BE LIABLE FOR ANY LOST PROFITS OR COST OF COVER, OR DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, PUNITIVE OR CONSEQUENTIAL DAMAGES, INCLUDING DAMAGES ARISING FROM OR RELATING TO ANY TYPE OR MANNER OF COMMERCIAL, BUSINESS OR FINANCIAL LOSS, EVEN IF SECUREQLAB HAD ACTUAL OR CONSTRUCTIVE KNOWLEDGE OF THE POSSIBILITY OF SUCH DAMAGES AND REGARDLESS OF WHETHER SUCH DAMAGES WERE FORESEEABLE.

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

SecureQLab (July 2024)