SAFIRE
NEXT-GENERATION FIREWALL TESTER

XENA NETWORKS

FIREWALL TEST REPORT

SUMMARY

4421Mbps	2.479ms	12%	0%	1020
Goodput	Latency	Error Percentage	Malware Block Rate	Number of Users

NETWORK TOPOLOGY

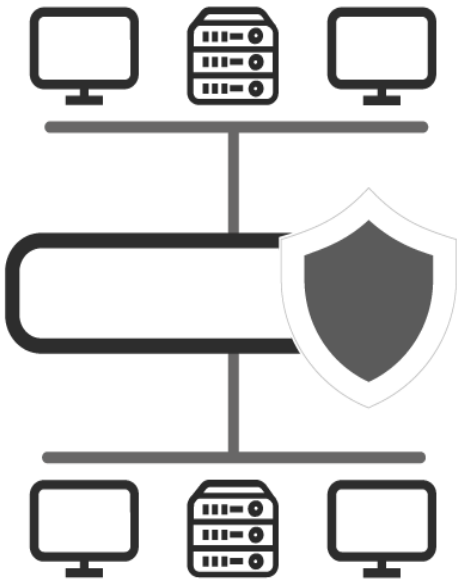

Enterprise Internal Segmentation

FIREWALL FUNCTIONS ENABLED

Firewall Model: PA-3060

Application Control | Logging and Reporting

Date:	2019-05-20
Time Start:	2019-05-20 (09:04:56)
Time Stop:	2019-05-20 (09:32:53)
Duration:	00:27:57
Safire:	v0.9.2

SAFIRE
NEXT-GENERATION FIREWALL TESTER

Contents

1. Introduction

Enterprises are investing in enterprise-grade firewalls to counteract the threats posted by the increasing network and data security breach every year. Firewalls are typically deployed at the edge of or inside enterprises to protect clients and servers from malwares, virus affection, data breach, target attacks, etc. Many security experts such as IT and security managers, and CSOs believe that they are improving their network security posture by implementing a new security solution.

However, verifying the performance of firewall and any network security device is essential to the success of the enterprise network it defends. With various advanced protection features such as application identification, intrusion prevention, threat detection, logging, etc., your firewall can easily become the performance bottleneck of the network, degrading the overall performance and user experience. Because of this trade-off, it is vital to test the performance of your firewall with specific features enabled on the firewall appliance.

Either out-of-box or firmware upgrade, firewall appliances should always be tested and evaluated before deployment in order to guarantee that new security protections do not adversely impact performance and that security shortcuts are not taken to maintain or improve the performance. The testbed should attempt to replicate the production network as close as possible, which includes network topology, network traffic that traverses through the firewall, features and policies enabled on the firewall, etc. Firewall appliance should deliver the expected performance under all circumstances.

This test report is generated by Safire, Xena's innovative enterprise firewall performance tester. Safire automatically measures and characterizes the performance of a firewall under realistic traffic conditions, with the results and key conclusions automatically being represented in a readable report format. Thus, the IT and network managers can easily assess the negative performance impact for the multiple firewall security features.

Safire makes it a simple and cost-efficient tool for:

- Comparing firewall performances from different vendors during the purchasing process.
- Validating firewall performances prior to network installation.
- Characterizing firewall performances after software updates and patches.
- Verifying firewall performances when your LAN infrastructure needs substantial changes.
- Characterizing firewall performances for specific firewall application scenarios, e.g. corporate or datacenter backup applications, and high availability application.

2. Test Information

Test Name:	PA3060 Example
Test Date Time:	2019-05-20
Test Start:	2019-05-20 (09:04:56)
Test Stop:	2019-05-20 (09:32:53)
Test Duration:	00:27:57
Network Topology:	Enterprise Internal Segmentation
Firewall Model:	Palo Alto PA-3060
Firewall Interface Addresses:	**Segment A:** 10.0.0.1, **Segment B:** 11.0.0.1
Firewall Interface Speed:	10G
Firewall Functions Enabled:	

Application Control Logging and Reporting

Comment:

This is an example.

3. Network Topology

In this test case, the network topology is configured as:
Enterprise Internal Segmentation

Firewalls are deployed in differnet network locations, i.e. inside the network or at the perimeter, and they are used to protect different devices, i.e. clients or servers. Depending on where the firewall is deployed and what the firewall protects, the traffic profile seens by the firewall varies. This section lists four main network topologies that you usually see in firewall deployment scenarios:

- Enterprise Internal Segmentation
- Enterprise Security Perimeter
- Data Center Internal Segmentation
- Data Center Security Perimeter

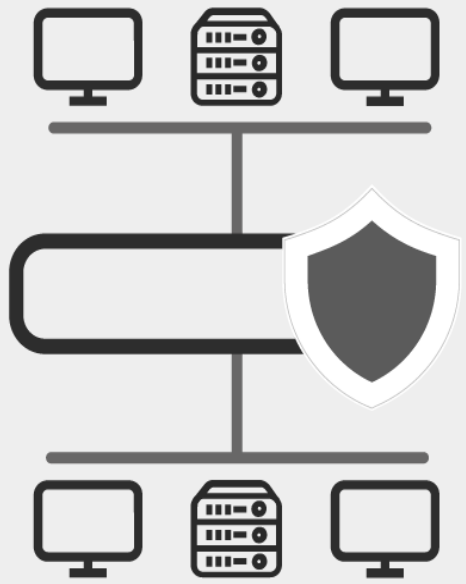

Illustration: 3.1 Enterprise Internal Segmentation

3.1 Enterprise Internal Segmentation

Firewall is placed inside the enterprise network to secure enterprise network by segmenting the corporate LAN and protecting each segment from others against malware and virus usually by means of application control, antivirus, web filtering, DNS filtering, and SSL deep inspection. It is usually referred to as "Zero Trust". Traffic characteristics are symmetric and west-east. Throughput demand is high since enterprise LAN capacities and speeds are orders of magnitudes higher than at the edge.

3.2 Enterprise Security Perimeter

Firewall is placed at the edge of the enterprise network to protect enterprise users from internet malware and virus usually by means of application control, antivirus, web filtering, DNS filtering, and SSL deep inspection. Traffic characteristics are asymmetric and north-south. Throughput is limited by the WAN interface provisioned by the ISP.

Illustration: 3.2 Enterprise Security Perimeter

Illustration: 3.3 Data Center Internal Segmentation

3.3 Data Center Internal Segmentation

Firewall is placed inside the data center network. It controls traffic flowing between servers and application tiers inside the data center usually by means of IPS, antivirus, and web filtering. Traffic characteristics are symmetric and west-east. Throughput demand is very high because intra-data center communication, such as data backup, demands bandwidth capacity of hundreds of gigabits per second.

Illustration: 3.4 Data Center Security Perimeter

3.4 Data Center Security Perimeter

Firewall is placed at the edge of the data center network. It controls traffic flowing from the internet to the data center and flowing from data center to the internet usually by means of IPS. Traffic characteristics are asymmetric and north-south. Throughput demand is high because SaaS applications demand high bandwidth and low latency in order to services SaaS users.

4. Traffic Profile

In this test report, the traffic profile is configured as follows:

Traffic Profile: **Internal Segmentation - 10G**

Users Allocation Per Application

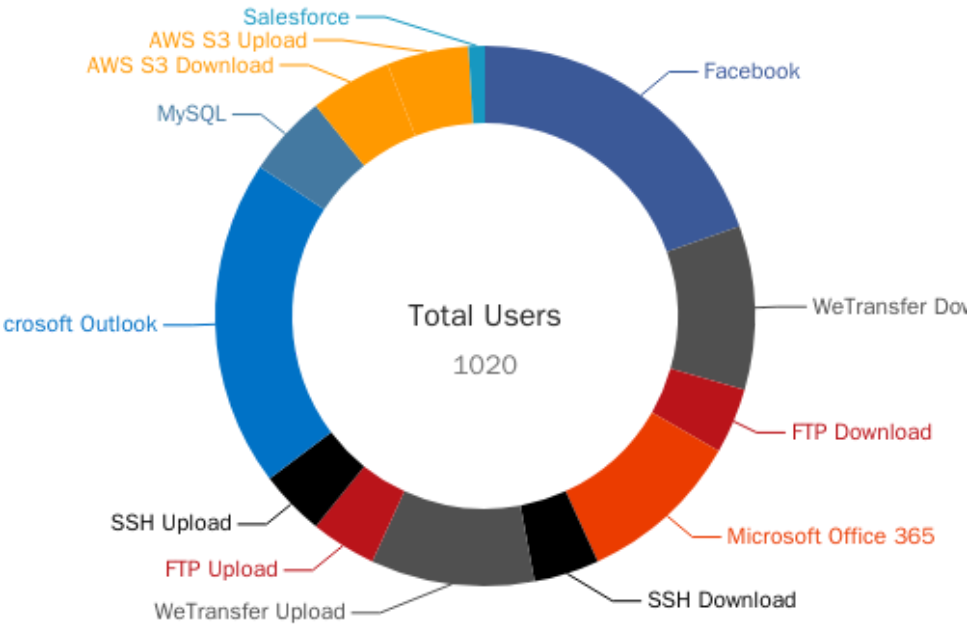

Bandwidth Allocation Per Application

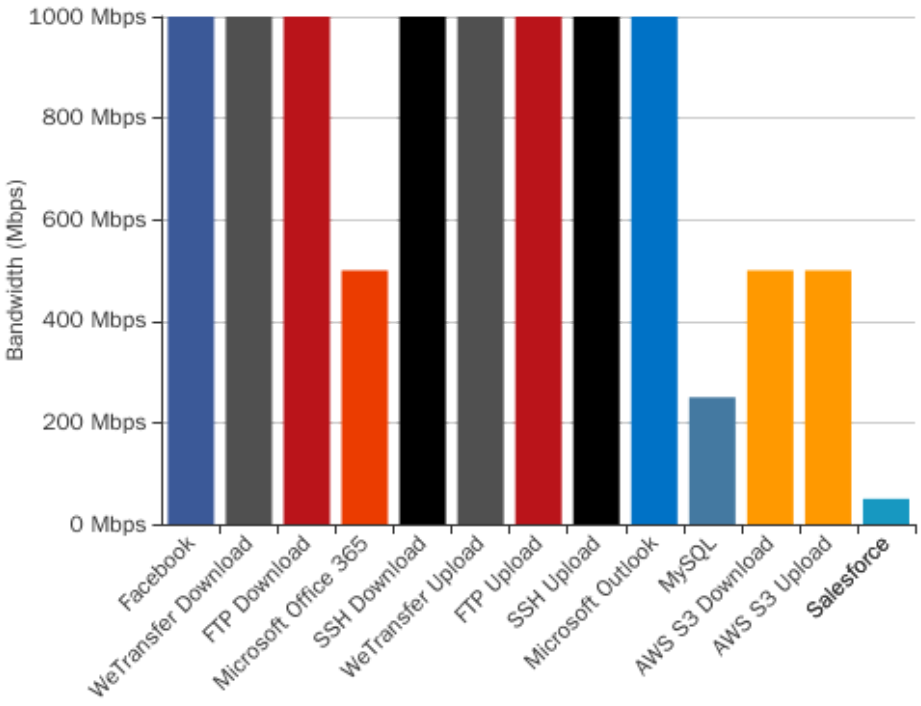

Application	Bandwidth Per User	Number of Users	Bandwidth
Facebook	5 Mbps	200	1000 Mbps
WeTransfer Download	10 Mbps	100	1000 Mbps
FTP Download	25 Mbps	40	1000 Mbps
Microsoft Office 365	5 Mbps	100	500 Mbps
SSH Download	25 Mbps	40	1000 Mbps
WeTransfer Upload	10 Mbps	100	1000 Mbps
FTP Upload	25 Mbps	40	1000 Mbps
SSH Upload	25 Mbps	40	1000 Mbps
Microsoft Outlook	5 Mbps	200	1000 Mbps
MySQL	5 Mbps	50	250 Mbps
AWS S3 Download	10 Mbps	50	500 Mbps
AWS S3 Upload	10 Mbps	50	500 Mbps
Salesforce	5 Mbps	10	50 Mbps
Total:		**1020**	**9800** Mbps

5. Test Result

Firewalls from different vendors can have large performance differences when tested with various realistic traffic profiles. Evaluating firewalls only from the datasheet is far from enough. Additionally, regular software update to the firewall requires retesting before putting into the production network again. Testing the firewall with Safire to verify how the upgrade reacts to realistic application traffic provides better foresight.

This section contains five main key performance metrics as follows:

- Goodput
- Concurrent Sessions
- Latency
- Error Percentage
- Malware Block Rate

Each section starts with the definition of the performance metric, followed by a discussion. Measurement result as a function of number of users are shown in both a chart and a table.

5.1 Goodput

Goodput (aggregated) is the application-level throughput defined as the useful amount of application data per unit of time that the application layer transports in both client-to-server (upstream) and server-to-client (downstream) directions, excluding protocol overhead bits as well as retransmitted data packets. The goodput is always lower than the layer-1 throughput (the gross bit rate that is transferred on the wire).

Factors that cause lower goodput than layer-1 throughput:

- Retransmission of lost or corrupt packets caused by bit errors or packet dropping in congested network devices, such as firewall, switches and routers.
- Transport layer flow control and congestion control.
- Protocol overhead: transport layer, network layer and data link layer protocol overhead is typically included in the throughput, but is excluded from the goodput.

The chart below shows the goodput under different numbers of users, as well as layer-1 throughput rate. The goodput increases as the number of users increases. However, after a certain limit, the goodput will start converging. When this happens, it indicates that the firewall has reached it's performance limit and cannot handle more traffic. Packets can get lost or misordered due to the congestion inside the firewall.

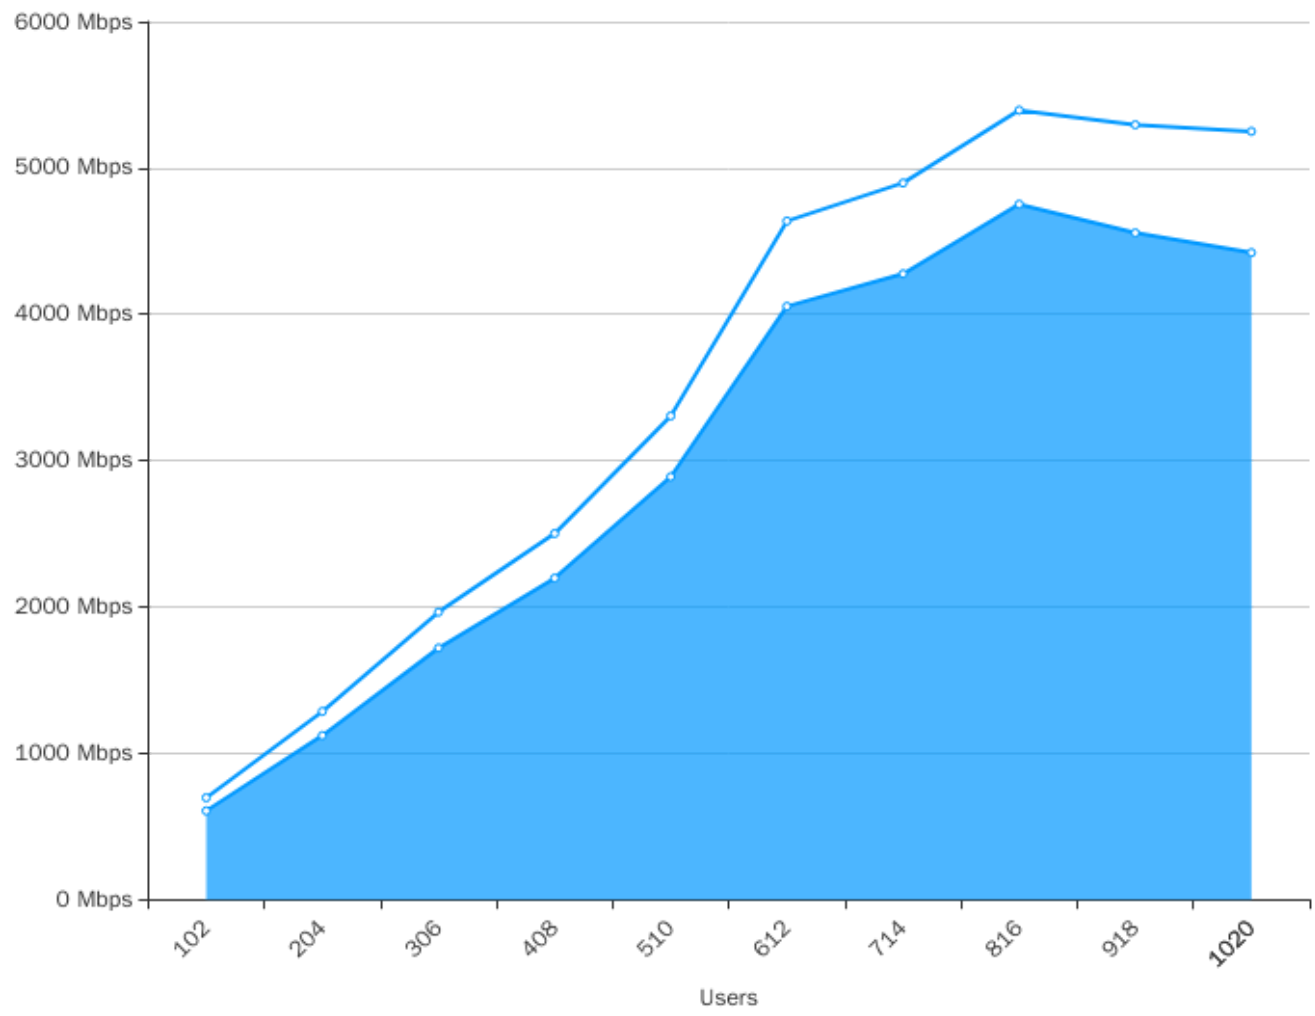

Number of Users	Goodput (Mbps)	Layer 1 Rate (Mbps)
102	605 Mbps	695 Mbps
204	1120 Mbps	1285 Mbps
306	1718 Mbps	1964 Mbps
408	2196 Mbps	2501 Mbps
510	2889 Mbps	3305 Mbps
612	4054 Mbps	4637 Mbps
714	4276 Mbps	4897 Mbps
816	4752 Mbps	5395 Mbps
918	4556 Mbps	5295 Mbps
1020	4421 Mbps	5250 Mbps

5.2 Concurrent Sessions

A session is defined by two uni-directional flows each uniquely identified by a 5-tuple key: source-address, destination-address, source-port, destination-port, and transport layer protocol. The concurrent session describes the maximum established/active sessions maintained at a given point in time by the firewall during each test.

The chart below shows the concurrent sessions under different numbers of users. The number of concurrent sessions increases as the number of users increases. However, after a certain limit, the number of concurrent sessions will start converging. When this happens, it indicates that the firewall has reached its limit and cannot handle more sessions due to the limited capacity of its session table.

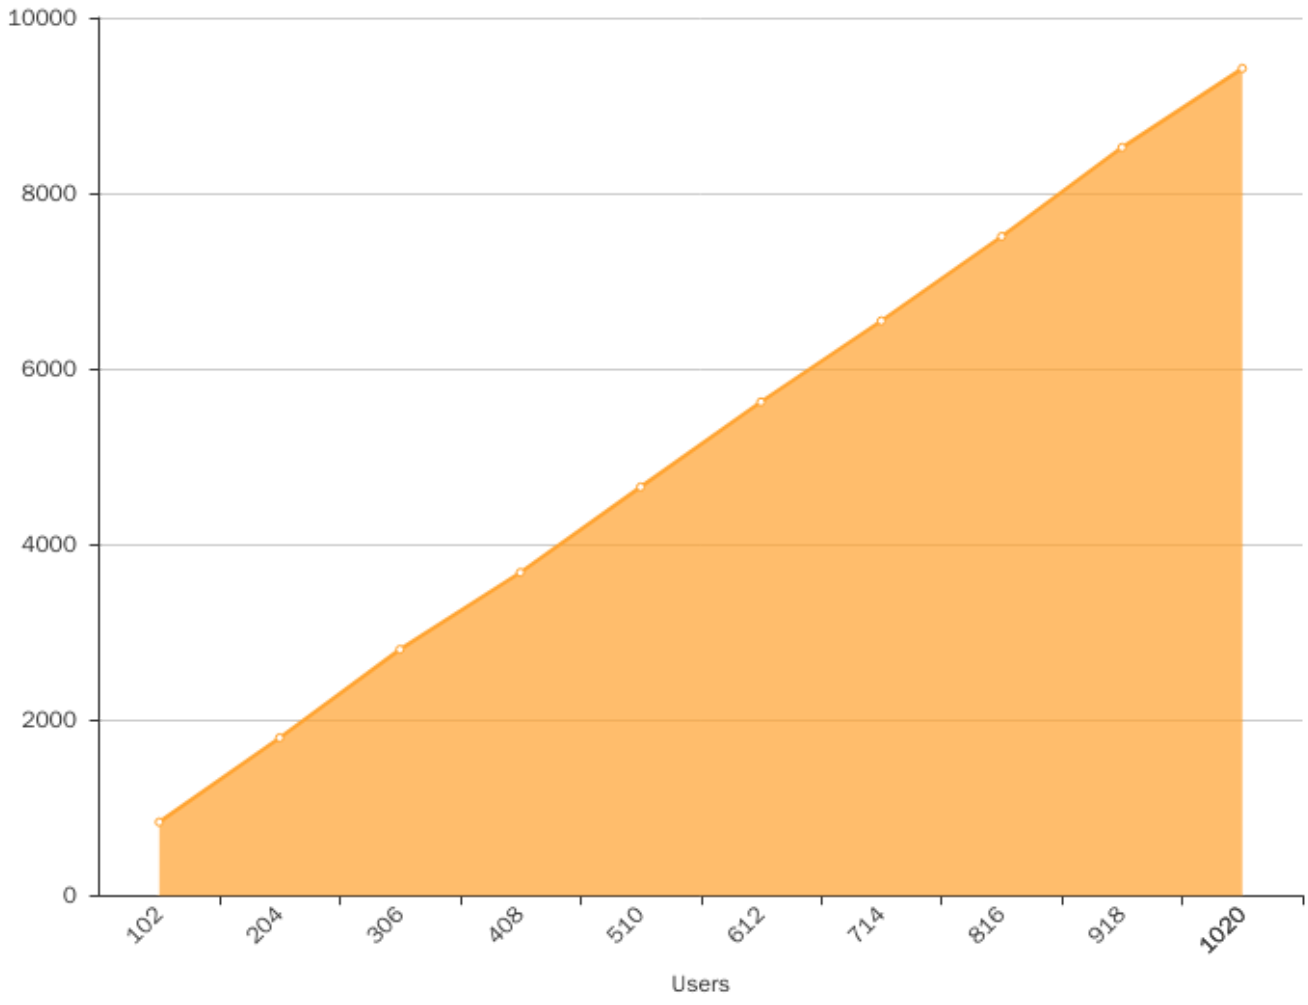

Number of Users	Concurrent Sessions
102	835
204	1795
306	2802
408	3678
510	4653
612	5622
714	6545
816	7507
918	8521
1020	9422

5.3 Latency

Latency is defined as the round-trip time (RTT) delay between the simulated clients and servers. The RTT latency measurement indicates how long for the data transmitter to receive the acknowledgement from the receiver. If packets are dropped by the firewall, TCP retransmission timeout will be triggered and will dramatically increase the RTT value.

The chart below shows the average RTT latency of the test scenario under different number of users. Latency increases with the growing number of users because the firewall needs to spend more time to process the incoming traffic. After a certain number of users, the latency may increase exponentially. This indicates that the firewall cannot handle the amount of traffic.

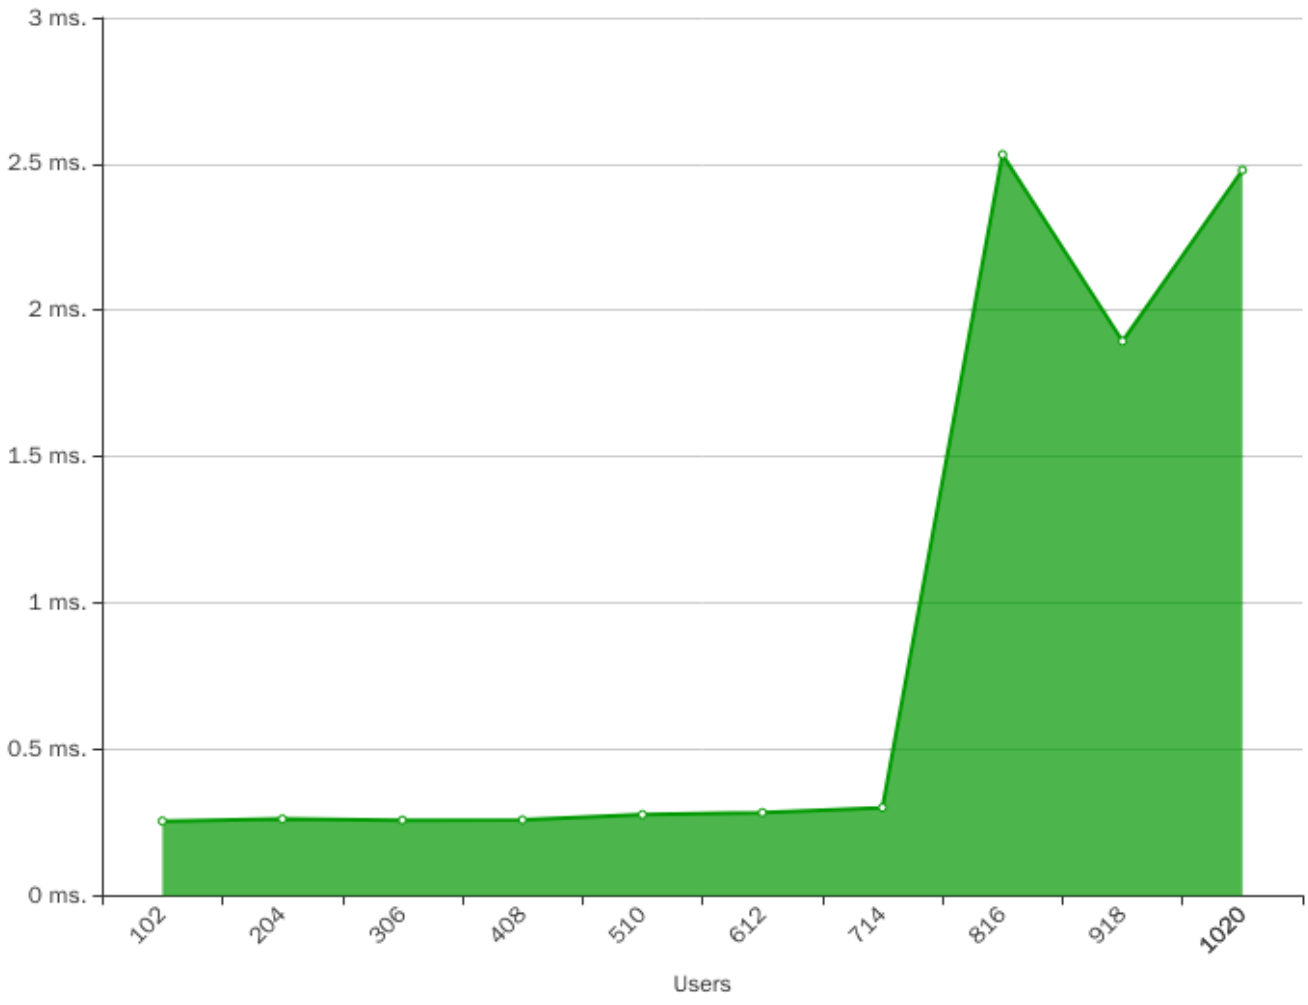

Number of Users	Latency (milisecond)
102	0.254 ms
204	0.261 ms
306	0.257 ms
408	0.258 ms
510	0.276 ms
612	0.283 ms
714	0.3 ms
816	2.532 ms
918	1.894 ms
1020	2.479 ms

5.4 Error Percentage

Error percentage is defined as the ratio between the number of retransmissions and the total number of packets transmitted. The retransmission includes both TCP SYN retransmissions, TCP fast retransmissions, FIN retransmissions, out-of-order packets, and duplicated ACKs.

TCP SYN retransmission indicates that the firewall fails to establish TCP connections before the timeout occurs. TCP fast retransmission indicates that transmitted packets are not received by the receiver due to packet dropping caused by the congested firewall.

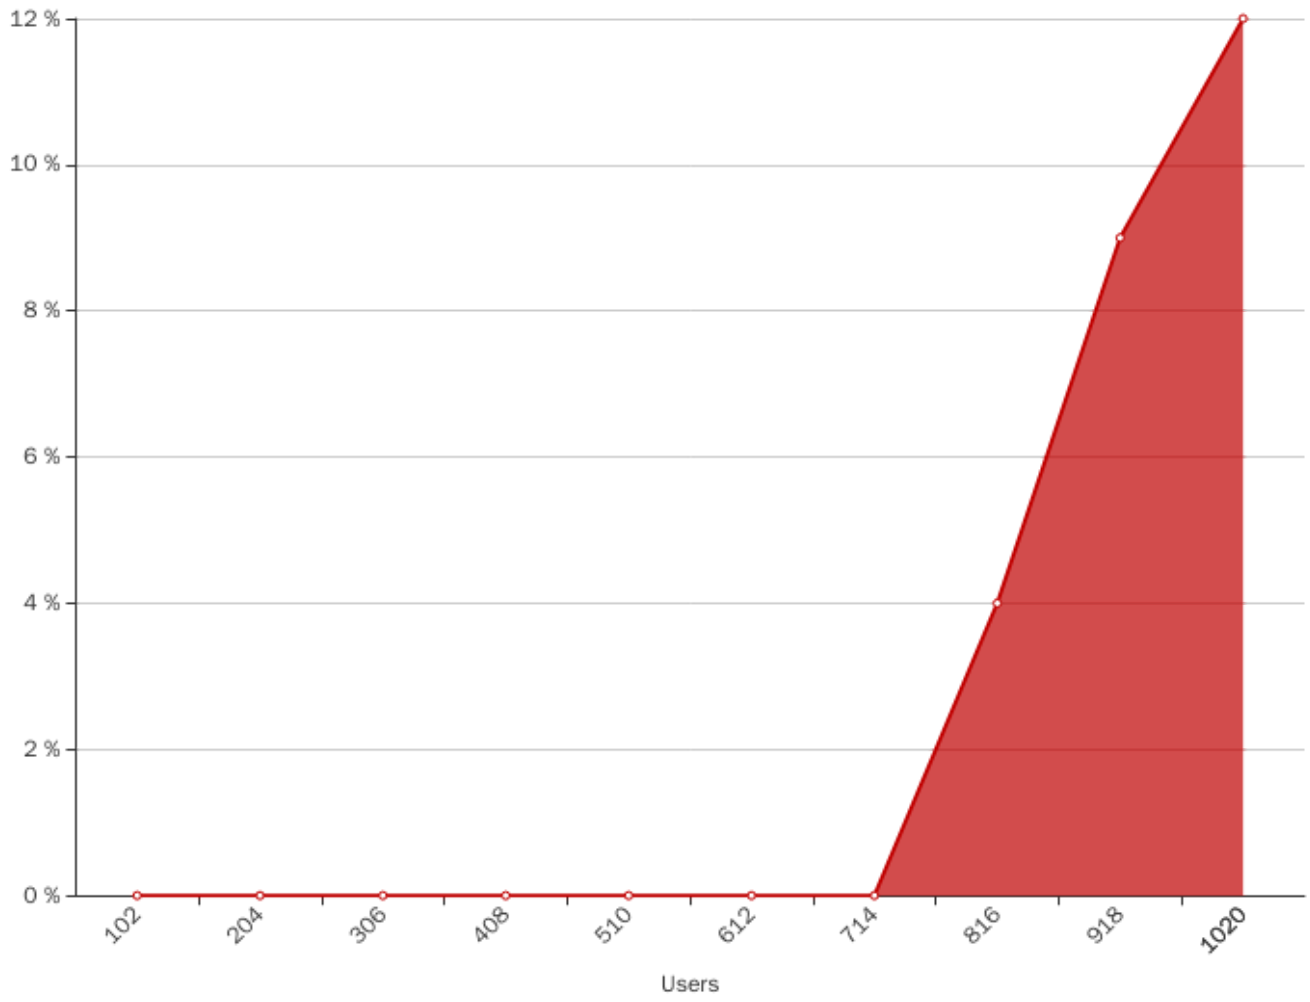

Number of Users	Error Percentage
102	0 %
204	0 %
306	0 %
408	0 %
510	0 %
612	0 %
714	0 %
816	4 %
918	9 %
1020	12 %

5.5 Malware Block Rate

Malware block rate is defined as the the ratio between number of successfully blocked malware and the total number of malware injected.

Malware (malicious software) is any software intentionally designed to cause damage to a computer, server, client or network. Malware does the damage after it is implanted or introduced in some way into a target's computer and can take the form of executable code, scripts, active content, and other software. The code is described as computer viruses, worms, Trojan horse, ransomware, spyware, adware, etc.

A firewall without an anti-malware function enabled can place high risks on the network security for the enterprise. Thus, performance testing the firewall without keeping it busy with the work it is supposed to is invalid. In order to exercise the anti-malware engine, virus injection is used together with the simulated user traffic. The goal is not to test the security efficacy of the firewall but to keep the anti-malware engine busy so that the test result is convincing.

The test malware file used by Safire is safe, because it is not a virus, and does not include any fragments of viral code. Most security products react to it as if it were a virus. The file is a legitimate DOS program, and produces sensible results when run.

Anti-malware test traffic includes both non-encrypted malware injection (plaintext) and encrypted malware injection (TLS-encrypted). It requires the firewall to have anti-malware function enabled in order to successfully block the non-encrypted malware traffic. It requires the firewall to have anti-malware and SSL deep inspection functions enabled in order to successfully block the encrypted malware traffic. This is because when malware traffic is encrypted, the firewall won't be able to identify the threat but let it pass if the decryption function is off. A firewall with anti-malware function enabled should successfully block 100% of the non-encrypted malware under all circumstances. If not, it indicates the firewall has reached its bottleneck and the security function start being unstable.

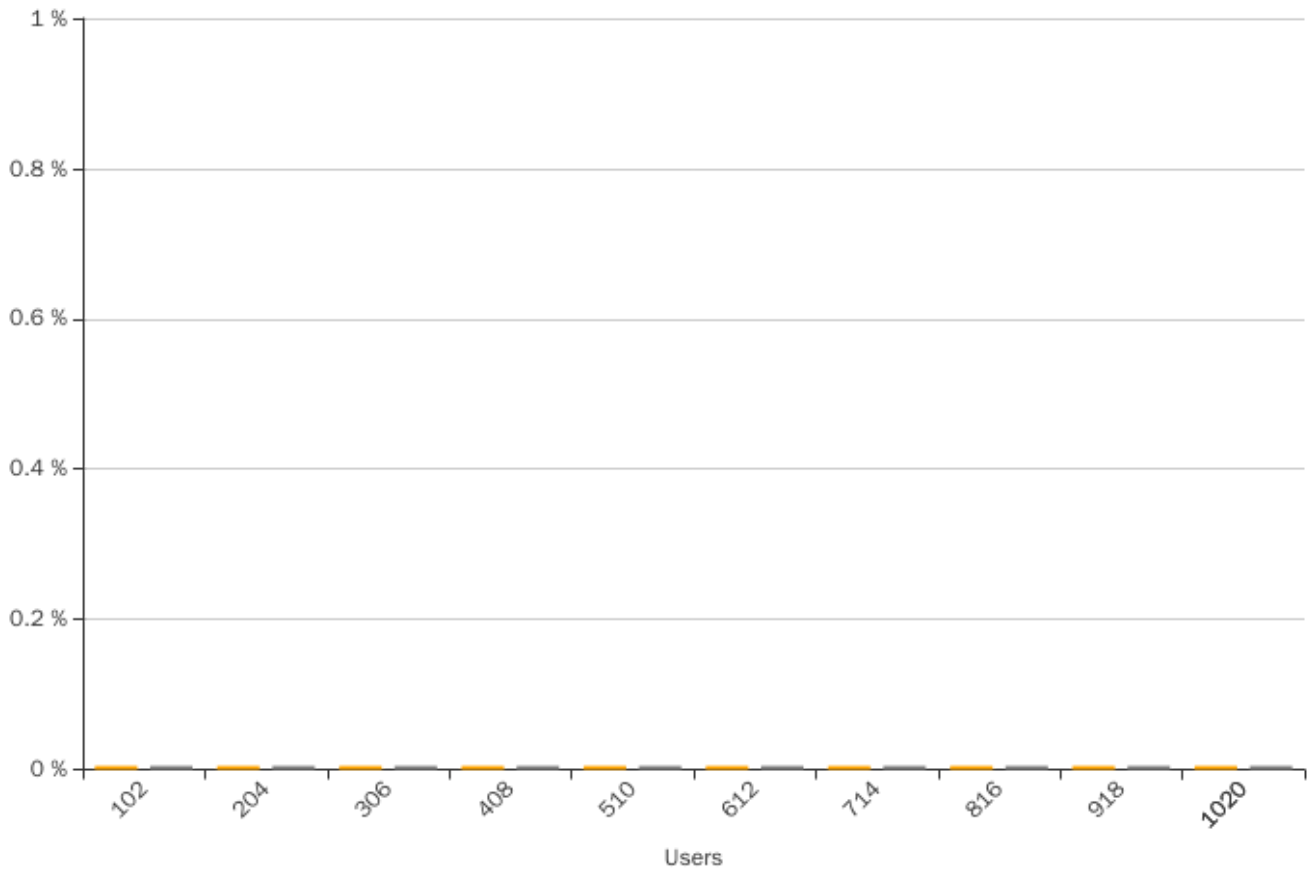

Number of Users	Malware Block Rate (non-encrypted)	Malware Block Rate (encrypted)
102	0 %	0 %
204	0 %	0 %
306	0 %	0 %
408	0 %	0 %
510	0 %	0 %
612	0 %	0 %
714	0 %	0 %
816	0 %	0 %
918	0 %	0 %
1020	0 %	0 %

SAFIRE
NEXT-GENERATION FIREWALL TESTER

Contact Information

Xena Networks
Lottenborgvej 26
2800 Lyngby
Denmark

support@xenanetworks.com
www.xenanetworks.com