

POLIWALL® PERFORMANCE TESTING

IMPLEMENTATION OF COMPLEX FILTERING POLICIES



WWW.TECHGUARD.COM

July 16, 2012

Table of Contents

I. INTRODUCTION	2
II. TEST EQUIPMENT	3
III. TEST METHODOLOGY	3
A. Packet Generation	3
B. Baseline Configuration.....	3
C. Pre-compiled Exception List Configuration.....	4
D. DDoS SYN Flood Configuration	4
IV. TEST RESULTS	5
A. M10 Performance Results.....	5
B. V01 Performance Results.....	7
C. G01 Performance Results.....	9
D. M10 SYN Flood Results	11
V. CONCLUSION	13

I. INTRODUCTION

Today’s cyber-security environment requires the ability to implement complex security policies to defend against an increasingly resilient and distributed threat. The PoliWall is a specialized perimeter security device that blocks traffic to and from foreign countries and can process large IP block lists that can include tens of millions of IPs. By stopping needless traffic and malicious IPs at the perimeter, the workload on systems deeper inside the network is decreased, thereby increasing the performance of these devices.

The PoliWall cyber-defense tool is engineered to implement large complex filtering requirements while maintaining high throughput and low network packet latency. The PoliWall does not require any external devices to implement these filtering capabilities, and does not send any connection related meta-data to external sources for processing. This paper reports the findings of a series of test performed to evaluate the effectiveness of the PoliWall, measured in packet latency, packet throughput, and TCP connection rate.

In addition to the raw performance testing, we performed a test to determine how the PoliWall would perform while under attack. For this test we simulated a DDoS attack by generating a SYN flood from millions of IP addresses. This test was performed with country filtering enabled and 30 million IP addresses blocked by the filtering policy.

II. TEST EQUIPMENT

Three PoliWall models were tested for this report. The entry level M10 model with 100Mbps/seconds throughput, the enterprise level G01 model designed for gigabit environments, and the virtual V01 designed for cloud deployments.

Each model was tested by sending TCP connections through the filtering interface with a payload size of 128 bytes. This value was chosen to ensure that packet rates and TCP connection rates could be sufficiently tested without being constrained by the maximum throughput limitations of each models physical interfaces.

All TCP traffic was generated and measured by a Breaking Point Storm running version 2.2.7. The hardware configuration used for this testing is capable of generating 15 million concurrent sessions, 8Gbps traffic and 750,000 sessions per second.

III. TEST METHODOLOGY

A. Packet Generation

The packets were generated using the Session Sender (SS) test. The Session Sender test component measures a device's ability to set up and maintain a large number of TCP sessions over a period of time. Each session uses a unique combination of source addresses, destination addresses, source ports, and destination ports. The following settings were used for each device

Device	Bandwidth	Sessions/sec
M10	100 Mb/s	20,000
V01	1000 Mb/s	20,000
G01	1000 Mb/s	150,000

There are three phases within a Session Sender test: ramp up, steady-state, and ramp down. The graphs below show the steady-state phase of each test, where the Breaking Point device is actively opening connections, sending data, and closing connections.

B. Baseline Configuration

Each PoliWall was tested initially with a policy that only allowed IP traffic from IP addresses from U.S. network providers. At the time of the testing, the U.S. ISP address space was comprised of 14,999 distinct IP subnets. The PoliWall assigns a category to every packet passing through the filtering interfaces and compares the category to the defined policy. The time required to assign the category is constant and independent of the filtering policy, resulting in identical measurements for an "allow all" policy and a "U.S. Only" policy. Therefore, the "U.S. Only" policy was used as the baseline.

C. Pre-compiled Exception List Configuration

Pre-compiled exception lists (PCEL®) give the PoliWall the ability to protect against millions of IP addresses associated with malicious behavior. Typical lists provided by commercial and open-source IP reputation providers contain thousands to millions of addresses, with the larger lists approaching 10,000,000 entries. To test performance while protecting against these known threats, the PoliWall was configured with three PCEs – one with 5 million entries, one with 10 million entries, and one with 15 million entries. All three lists were bound to the active filtering policy, requiring that every connection be processed against 3 lists containing a total of 30 million individual IP addresses. The IP addresses on the PCEs were chosen to not overlap with the IP addresses used by the Breaking Point packet generator to ensure that each generated IP packet was processed by all stages in the PoliWall security engine.

D. DDoS SYN Flood Configuration

The PoliWall M10 was configured with a filtering policy that allows only IP traffic from the U.S., and blocks 30 million IP addresses with the PCEL configuration described above in section C. The Breaking Point traffic generator was tuned to generate 10,000 new TCP connections per second. This number was chosen based on the maximum number of connections that could be handled on a 100 megabit connection with an average transaction size of 1000 bytes. The SYN flood was started 10 seconds into the test, with an average of 66,485 SYN packets transmitted per second.

IV. TEST RESULTS

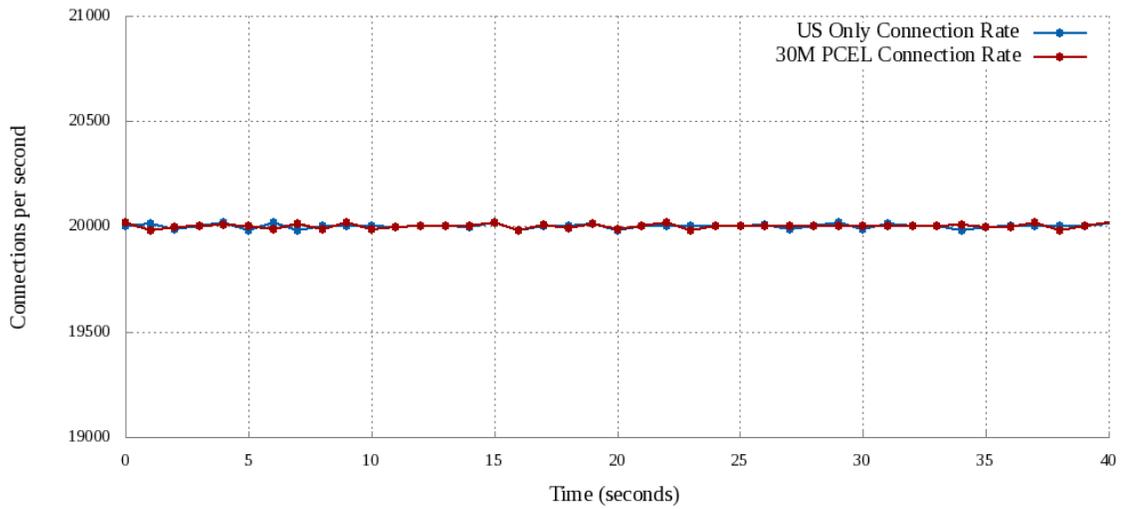
A. M10 Performance Results

The M10 was tested with an attempted connection rate of 20,000 TCP connections per second. With the baseline filtering policy allowing only the 14,999 IP subnets assigned to U.S. Internet service providers, the M10 was able to sustain an average connection rate of 20,000 connections per second, with a maximum variance of 0.1 percent. After adding 30 million blocked IP address to the filtering policy, the M10 maintained an average TCP connection rate of 20,000 connections per second, with a maximum variance of 0.1 percent. The performance variation between the baseline and 30M PCEL configuration was 0.0 percent.

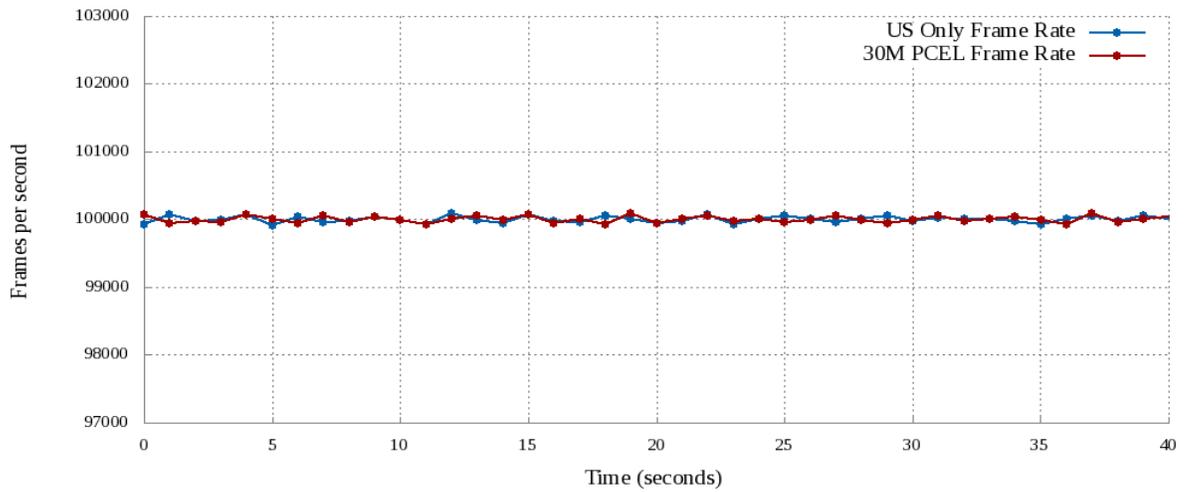
The packet throughput averaged 100,000 packets per second with both the baseline and the 30M PCEL configuration. The variance between the two tests was 0.0 percent.

Packet latency for the baseline test averaged 135.5 microseconds, with a relative standard deviation of 7.6 percent. After loading the 30M PCEL configuration, the average latency was 138.6 microseconds, with a relative standard deviation of 6.9 percent. The average latency induced by added 30 million blocked IP address was only 3.1 microseconds.

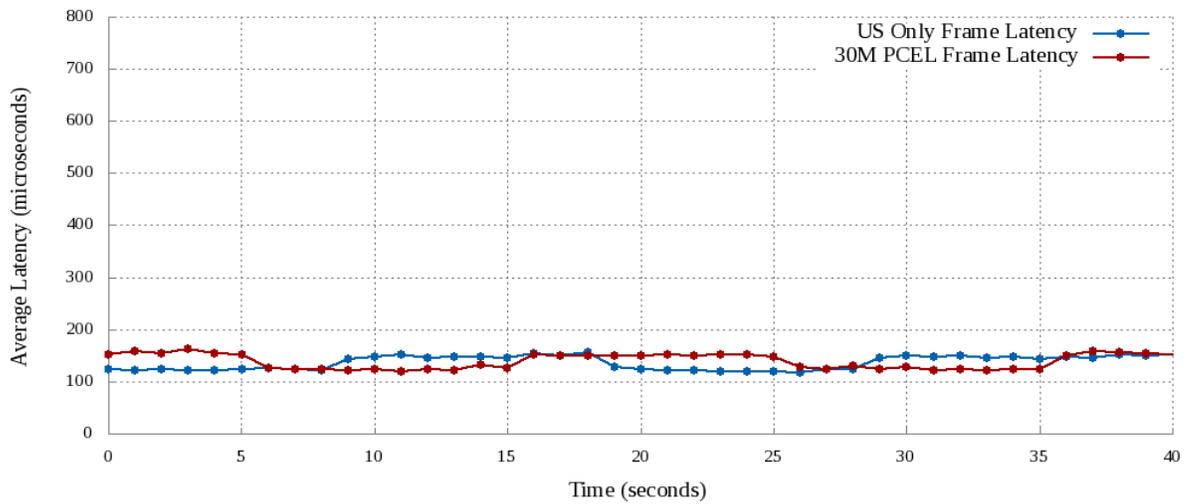
PoliWall M10 TCP Connection Rate



PoliWall M10 Frame Rate



PoliWall M10 Average Frame Latency (Microseconds)



B. V01 Performance Results

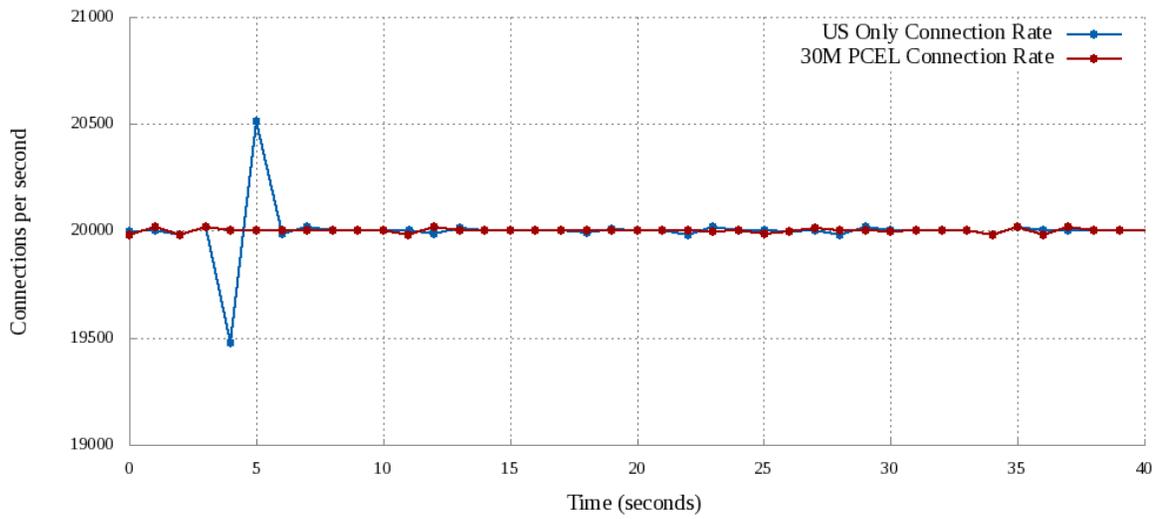
The V01 was tested with an attempted connection rate of 20,000 TCP connections per second. With the baseline filtering policy allowing only the 14,999 IP subnets assigned to U.S. Internet service providers, the V01 was able to sustain an average connection rate of 20,000 connections per second, with a maximum variance of 2.6 percent. After adding 30 million blocked IP address to the filtering policy, the V01 maintained an average TCP connection rate of 20,000 connections per second, with a maximum variance of 0.1 percent. The performance variation between the baseline and 30M PCEL configuration was 0 percent.

The packet throughput averaged 100,009 packets per second for the baseline and 99,998 packets per second for the 30M PCEL configuration. The variance between the two tests was 0.01 percent.

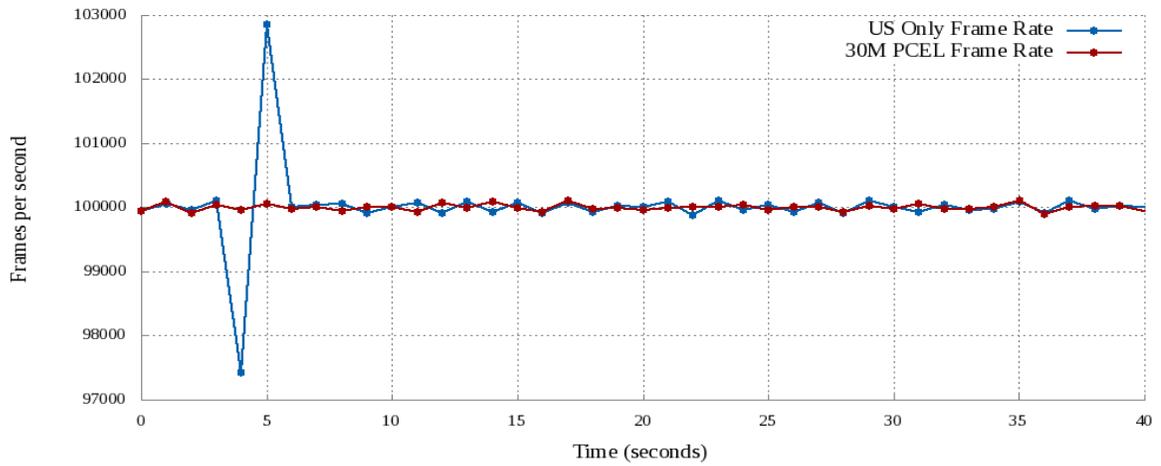
Packet latency for the baseline test averaged 380.5 microseconds, with a relative standard deviation of 1.0 percent. After loading the 30M PCEL configuration, the average latency was 396.1 microseconds, with a relative standard deviation of 1.7 percent. The average latency induced by added 30 million blocked IP address was only 15.6 microseconds.

Note: the drop and then subsequent spike in the TCP Connection Rate and Frame Rate graphs appears to be caused by schedulers on the VM hypervisor. In second 4, the PoliWall was unable to process all packets, but then in second 5 it was able to 'catch up' and process all the packets that arrived in second 5 plus the packets that weren't processed in second 4. The effect of this is also seen in the Frame Latency graph with the spike at seconds 4 and 5.

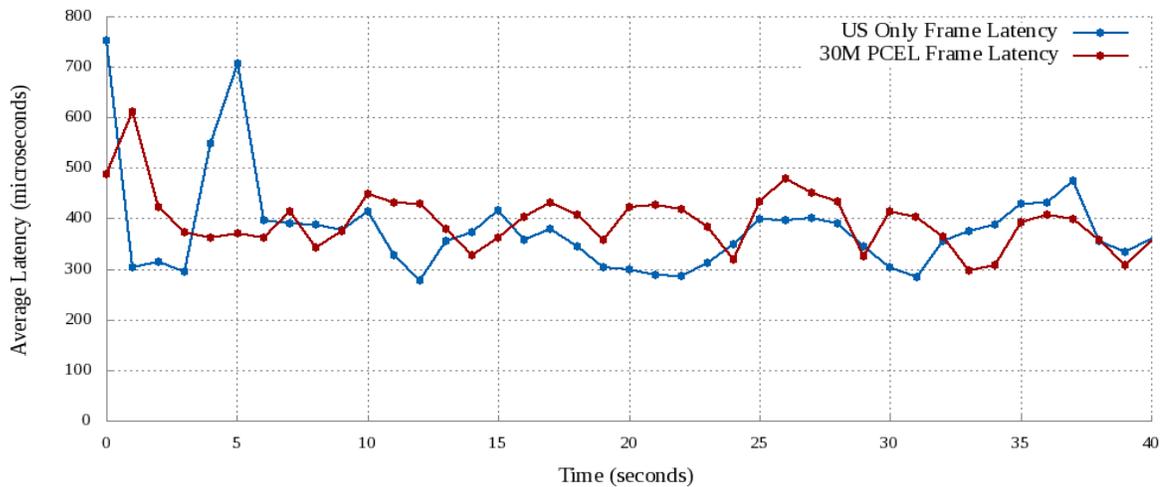
PoliWall V01 TCP Connection Rate



PoliWall V01 Frame Rate



PoliWall V01 Average Frame Latency (Microseconds)



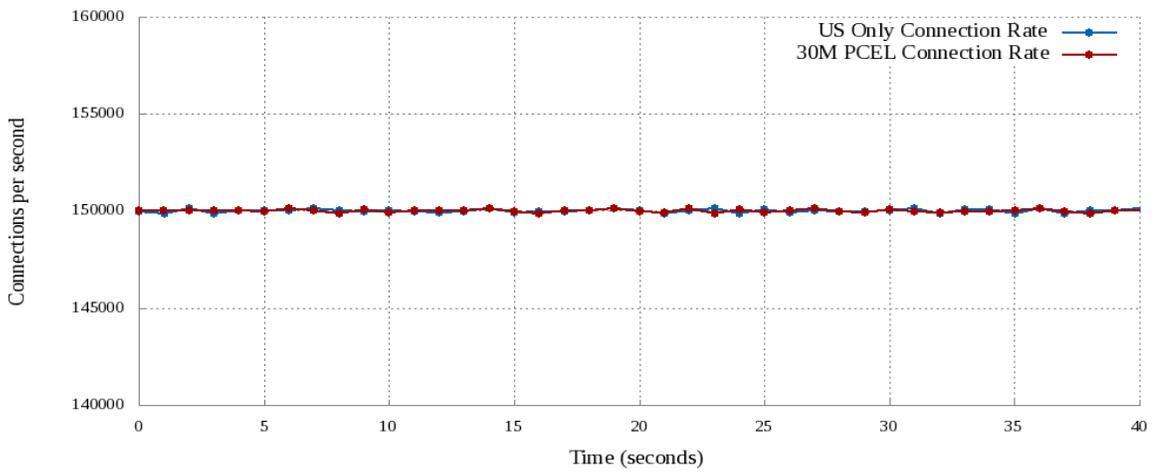
C. G01 Performance Results

The G01 was tested with an attempted connection rate of 150,000 TCP connections per second. With the baseline filtering policy allowing only the 14,999 IP subnets assigned to U.S. Internet service providers, the G01 was able to sustain an average connection rate of 149,996 connections per second, with a maximum variance of 0.09 percent. After adding 30 million blocked IP address to the filtering policy, the G01 maintained an average TCP connection rate of 150,000 connections per second, with a maximum variance of 0.1 percent. The performance variation between the baseline and 30M PCEL configuration was 0.0 percent.

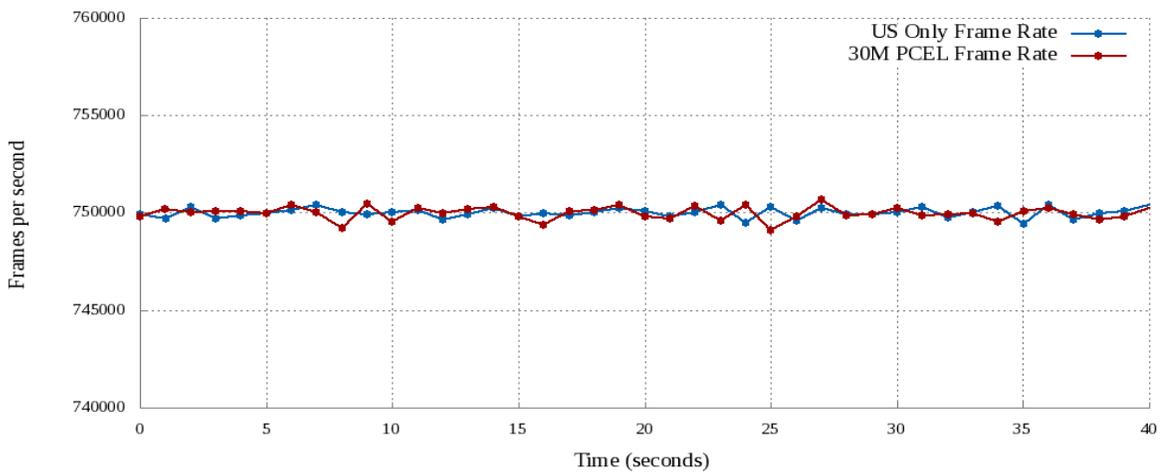
The packet throughput averaged 749,990 packets per second for the baseline and 749,996 packets per second for the 30M PCEL configuration. The variance between the two tests was 0.0 percent.

Packet latency for the baseline test averaged 112.0 microseconds, with a relative standard deviation of 4.8 percent. After loading the 30M PCEL configuration, the average latency was 124.7 microseconds, with a relative standard deviation of 4.9 percent. The average latency induced by added 30 million blocked IP address was only 12.7 microseconds.

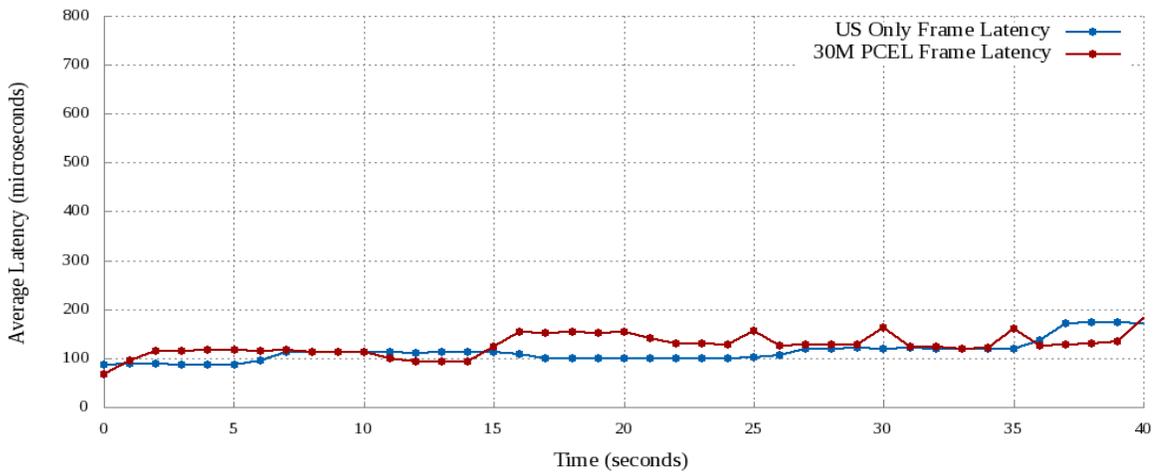
PoliWall G01 TCP Connection Rate



PoliWall G01 Frame Rate



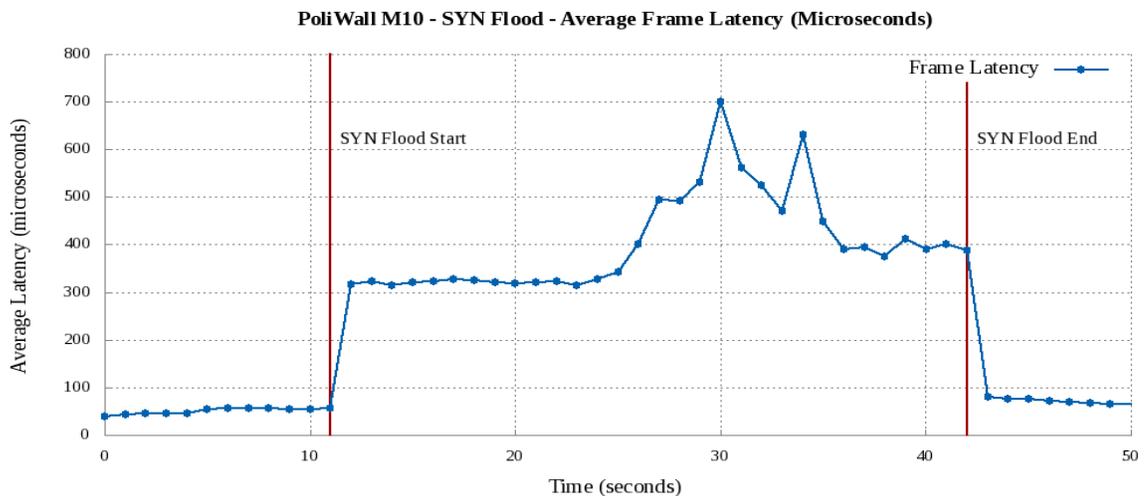
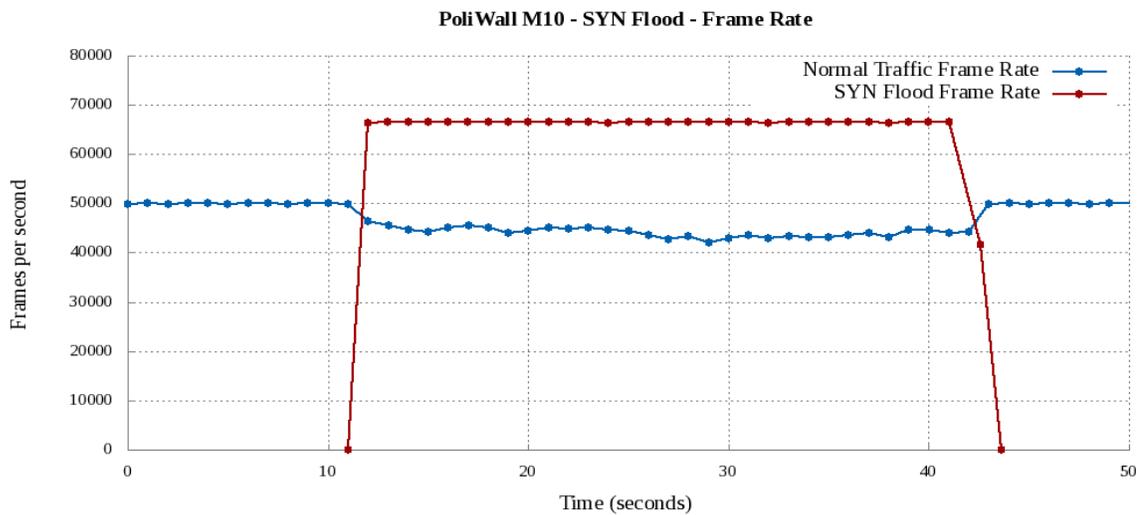
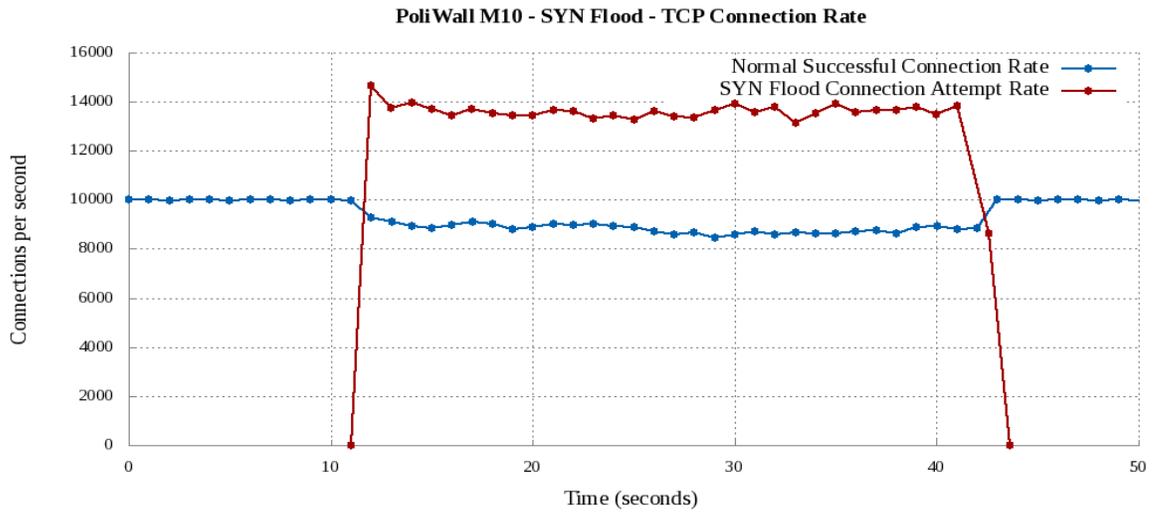
PoliWall G01 Average Frame Latency (Microseconds)



D. M10 SYN Flood Results

The SYN flood test was run with the PoliWall configured to allow only U.S. traffic and to block all traffic on the precompiled exception lists described in the test configuration section. The test was run with a baseline connection rate of 10,000 attempted TCP connections per second. After 10 seconds, 50 Mbps of TCP SYN packets were added to the existing 10,000 per second. The addition added 66,485 TCP SYN packets per second. These connections were not intended to complete and did not contribute to the measured TCP connection rate.

The PoliWall maintained an average of 10,000 connections per second before the SYN flood attack started. During the attack, the connection rate dropped to an average of 8,829 connections per second, an 11.7 percent drop from the pre-attack number. The frame rate rose from 49,991 per second to 110,628 per second during the attack. The frame latency rose from an average of 50.4 microseconds to 404.1 microseconds, with a maximum measured latency of 6,637 microseconds. The increase in latency was attributed to the attempted connection rate of 75,000 per second. Each attempt, including the packets that were part of the SYN flood attack were processed by the PoliWall's state handling code. The M10s one million entry state table allowed storing state information for new connections long enough to allow legitimate connection requests to complete before timing out.



V. CONCLUSION

All PoliWall models demonstrated exceptional performance while processing all IP traffic against a country restriction policy that allowed only packets from the U.S. and that restricted access to 30 million individual IP addresses. All models maintained an average packet latency less than 400 micro-seconds while enforcing the complex filtering policy, without requiring any external devices or sending data to any Internet processing center.

The PoliWall was able to maintain 88% of the baseline TCP connection rate while under an aggressive SYN flood attack. The latency increased slightly due to the high volume of connection attempts, but still remained an average latency under 500 microseconds.