Crossbeam/Check Point Firewall for Mobile Networks Test Report

Version 1.3
2011-12-22

EANTC AG

Salzufer 14D
D–10587 Berlin
Germany

Tel. +49. (0)30. 318 05 95–0
Fax +49. (0)30. 318 05 95–10
E-Mail info@eantc.de
WWW http://www.eantc.de/

Table of Contents

1 Introduction

Light Reading online publication contracted EANTC in November 2011 to execute a test program to validate the Crossbeam's X80-S performance and scalability in the context of Long Term Evolution (LTE) Mobile Networks.

Crossbeam offers a consolidated approach to network security. The Crossbeam solution is a blade server that is able to ran various applications all within the same chassis saving service providers network resources such as switches, routers and load balancers.

The specific tests documented here focus on mobile service provider deployment scenarios on the SGi interface - the interface connecting the Evolved Packet Core (EPC) to the Packet Data Network (PDN).

Test Setup

The test setup is shown below.

FIGURE 1. Physical Test Setup

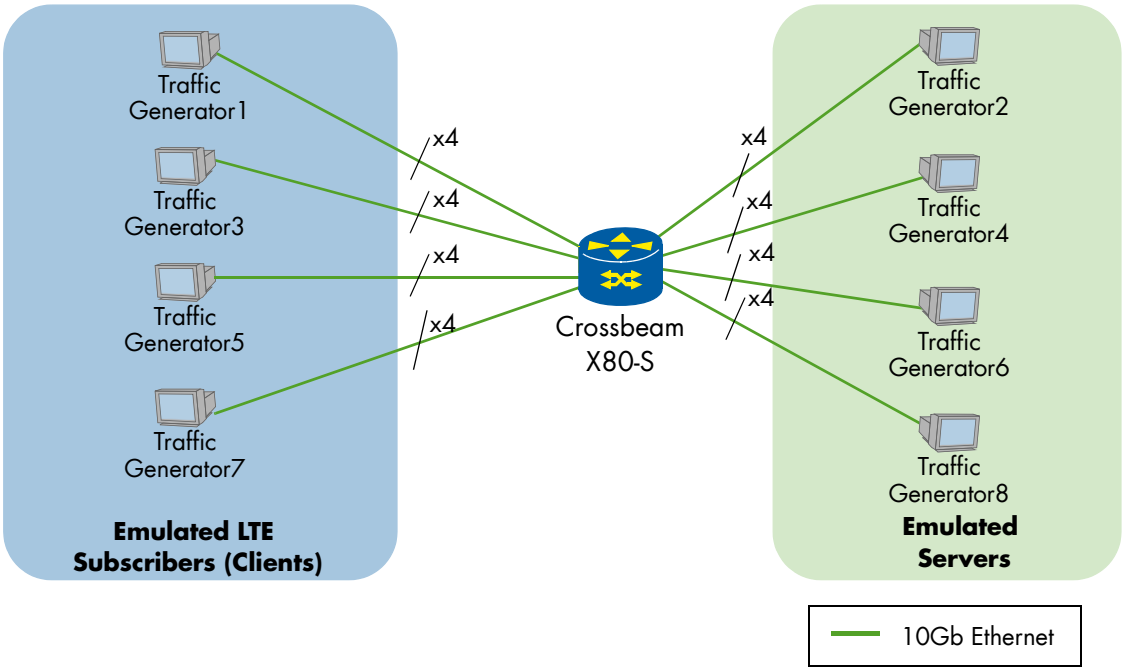

TABLE 1. Test Setup - Interface Details

IP Range	Tester Port	Function
11.10.0.0/16 to 11.13.0.0/16	Traffic Generator 1 port 1 to 4	Client Emulation
11.20.0.0/16 to 11.23.0.0/16	Traffic Generator 2 port 1 to 4	Server Emulation
51.10.0.0/16 to 51.13.0.0/16	Traffic Generator 3 port 1 to 4	Client Emulation
51.20.0.0/16 to 51.23.0.0/16	Traffic Generator 4 port 1 to 4	Server Emulation
91.10.0.0/16 to 91.13.0.0/16	Traffic Generator 5 port 1 to 4	Client Emulation
91.20.0.0/16 to 91.23.0.0/16	Traffic Generator 6 port 1 to 4	Server Emulation
131.10.0.0/16 to 131.13.0.0/16	Traffic Generator 7 port 1 to 4	Client Emulation
131.20.0.0/16 to 131.23.0.0/16	Traffic Generator 8 port 1 to 4	Server Emulation

FIGURE 2. **Logical Network - Firewall forwarding performance on SGi**

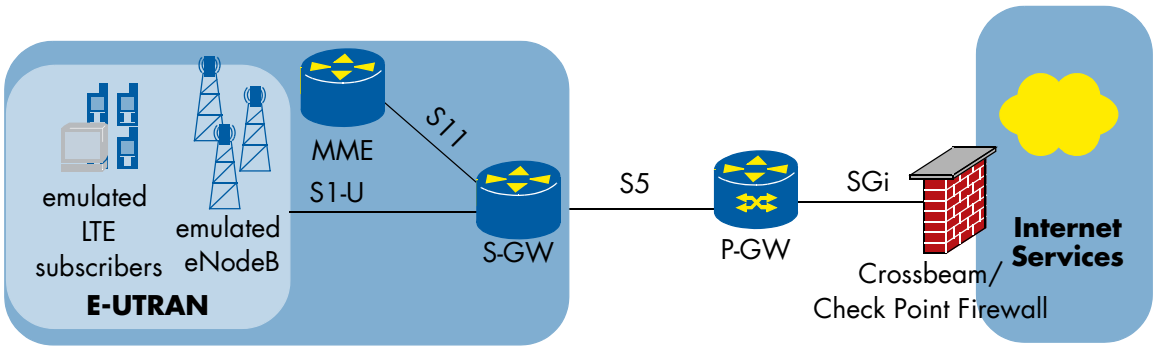

Hardware & Software Versions

The following tables describe the tester and device under test hardware and software. They will be filled in during the test execution.

TABLE 2. **Analyzer Equipment Hardware & Software Versions**

Vendor	Hardware	Software Version
Spirent	Avalanche 3100B	3.80GA

TABLE 3. **Device Under Test (DUT) Hardware & Software Versions**

Vendor	Hardware	Software Version
Crossbeam	X80-S with 4 x NPM 9650 & 8x APMs 9600 24GB RAM	Board Revision: AD Control FPGA Revision: 0x207 Focus FPGA Revision: 0x207 XOS 9.6
Check Point	Firewall-1	Major version: R75.20 version ID: 2

TABLE 4. Crossbeam X80-S Chassis Setup

Slot	Present	Module Name	Module Type	Status
1	Yes	np1	NP9650	Up
2	Yes	np2	NP9650	Up
3	Yes	np3	NP9650	Up
4	Yes	np4	NP9650	Up
5	Yes	ap3	AP9600	Active
6	Yes	ap4	AP9600	Active
7	Yes	ap5	AP9600	Active
8	Yes	ap6	AP9600	Active
9	Yes	ap7	AP9600	Active
10	Yes	ap8	AP9600	Active
11	Yes	ap9	AP9600	Active
12	Yes	ap10	AP9600	Active
13	Yes	cp1	CP9600	Up
14	No	n/a	n/a	n/a

Contacts

Crossbeam, 80 Central Street Boxborough, MA 01719, U.S.A

EANTC AG, Salzufer 14, 10587 Berlin, Germany

2 Performance and Scale Tests

2.1 SGi Throughput and Activation Rate

PURPOSE

We verified that the Crossbeam/Check Point firewall is able to support up to 107 Gbit/s throughput of IPv4 flows between emulated mobile users and emulated servers while performing the following functions: firewall, Network Address Translation (NAT) and Intrusion Prevention System (IPS).

DESCRIPTION

This test case aims to evaluate the performance of the firewall inserted in the SGi interface, upstream from the P-GW towards the Internet under realistic loads. At this interface packet streams are presented as clear IPv4 traffic (unencrypted) and without any tunnel headers (such as GTP). At this position in the network the firewall will deal with an aggregation of all mobile subscribers attached to the packet gateway - the number of subscribers attached to the Evolved Packet Core (EPC) will depend on the size of the network and the geographical organization of the EPC.

For the test, consumer IPv4 application streams are emulated (95% HTTP, 5% mix of DNS, SmartPhone OS Updates, SMTP, POP3). In addition, a number of supplementary firewall features are activated for certain test runs in order to quantify their performance impact (if at all). The features include network address translation (NAT) in 1:1 mode and intrusion prevention system (IPS).

In the test, real subscribers are being created by the traffic generator (Spirent Avalanche). The Spirent terminology of "Simulated User" is synonymous with the term subscriber. Depending on the series of actions the specific subscriber is tasked to accomplish, the relationship of Simulated Users (SimUser) to TCP Connections to Layer 7 Transactions will vary (as in a real subscriber in a production network). A subscriber is therefore a single User Equipment (UE). Each UE is responsible for more than one TCP connection.

Using the traffic distribution chosen for the test, the system will split the SimUsers according to the following logic: From an instantaneously scheduled 100 SimUsers, 95 SimUsers will be executing HTTP page requests, 2 users will be transferring SMTP messages, 1 user will be performing DNS, and 1 User will be transferring a large bulk object. The number of simulated users varies according to a loading curve and increases after a step to a sustaining period, then ramps down. The relative number by protocol will dynamically adapt within the ratios described to the number of open simulated users.

For the user performing HTTP, when the system needs to schedule a new user (Evaluated heuristically every 200 milliseconds based on real world conditions, constrains, and the desired position of the testing curve), the system instantiates a new simulated user (Simulated User Animating in the graphs below). The SimUser is assigned a MAC address, and an IP address. The Simulated user will then begin to request a new web page which will results in new TCP connection for the Level 1 URL. The SimUser will issue an HTTP GET to the first server, measuring errors and response times, as well as HTTP Error codes. The SimUser will open additional TCP session the level 2 URLs. The user will then wait 8 seconds and request a 4 Kb Large JPG object (Final Level 2 URL) and then upon receipt will close the connections. At this point the Completed SimUser count is 0, the Completed TCP Connection count is 3, and the Completed transaction count is 5. This is repeated with 2 additional servers. At the end of the Simulated user ActionList, the SimUser is deallocated. Per simulated user pass, the SimUser count is increased by 1, TCP Connection count is increased by 9, and transaction count is increased by 15.

For the SMTP Simulated user, a TCP Connection is opened, and an SMTP message is formed with an attachment of 64 Kb bytes from user1@crossbeam.com to user2@crossbeam.com, and then the message is sent to a mail server (emulated as well). In this case, per pass of this simulated user, the SimUser count is 1, the TCP connection is 1, and the transaction is 1.

For the POP3 Simulated user, the server is an emulated mail store, with a mailbox called ‘cb’. The POP3 mailbox will store 1 or 2 (Equal probability) of messages 64 bytes long when checked. On the client side, the POP3 simulated user will log in to the mailbox account, CHECK for new messages, then RETRIEVE mail. By the end of the POP3 SimUser, the SimUser count goes up by 1, the TCP connection is increased by 1.

For DNS SimUsers, the server emulated a DNS Zone called ‘Crossbeam’ with an A Record resolving www.crossbeam.com. The SimUser will issue a DNS record lookup for www.crossbeam.com. By one pass of this SimUser, the SimUser count is incremented by 1, the TCP connection is increased by 1, and the transaction count is increased by 1.

For the Bulk User Simulated User, the client will request a 2 Mb object. By the end of the SimUser, the SimUser count increases by 1, the TCP connection count increases by 1, and the Transaction count increases by 1.

TEST SETUP The following figure describes the logical test setup.

FIGURE 3. Test Setup- Firewall forwarding performance

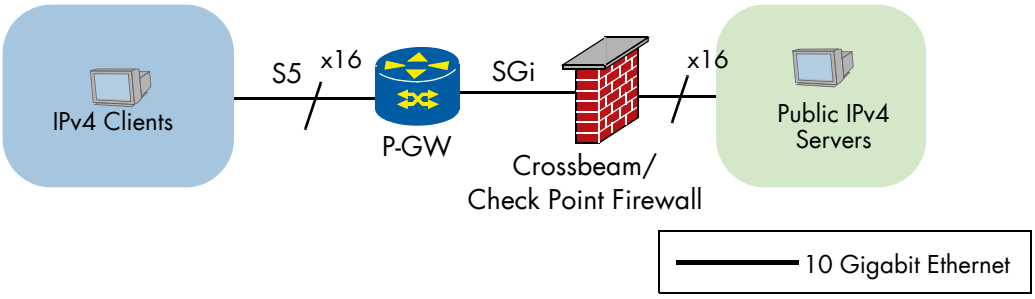

TABLE 5. Parameter Overview - Firewall forwarding performance

Parameter	Value
Traffic Profile	Ramp Up Time = 30 seconds per step, Ramp UP Repetition= 5 Stair Step Height = 208,000 Steady time per Ramp UP step = 30 second Ramp Down Time = 240 seconds Steady time per Ramp UP step = 80 seconds
NAT Type	Each client IP addresses is mapped to another IP address
Test Duration	Ramp up phase: 750 seconds Steady phase: 300 seconds Ramp Down phase: 320 seconds

TABLE 6. Total Traffic Definition for Firewall forwarding performance

Parameter	Value
HTTP Traffic	
Number of emulated clients	1,040,000
Number of emulated server	48
HTTP Type	1.1
1x level 1 URL	1,024 kByte
3x Level 2 URLs	1.60 kBytes (small icons)
1x Level 2 URL	4 kByte (large image)
Think timer	10 seconds
Percent of total traffic	95%
SMTP Traffic	
Number of emulated clients	1,040,000
Number of emulated server	48
Object size (transferred file)	64 kByte
Percent of total traffic	1%
POP3 Traffic	
Number of emulated clients	1,040,000
Number of emulated server	48
Object size (transferred file)	64 kByte
Percent of total traffic	2%
SmartPhone OS Over the Air Updates	
Traffic Type	HTTP 1.1
Number of emulated clients	1,040,000
Number of emulated server	These subscribers use the same HTTP servers as the HTTP traffic definition
Object size (transferred file)	2 MegaByte
Percent of total traffic	1%
DNS Traffic	
Number of emulated clients	1,040,000
Number of emulated server	48
Object size (transferred file)	552 Bytes
Percent of total traffic	1%

TABLE 7. Test Possibilities - Firewall forwarding performance

# Run	Test Goal	NAT	IPS
1	Throughput (Mix as stated)	Off	Off
2		On	Off
3		Off	On
4		On	On
5	Maximum connection setup rate (http)	On	On

Please note that this test case was executed with a single device under test configuration with 8 APMs and 4 NPMs. Intially the customer suggested to

run an additional condition in which all successful connections were being logged by the firewall. Due to time constraints the customer decided to drop this condition.

PROCEDURE

Two test phases were planned, each executed according to the test possibilities described in Table 7.

TCP Connection Setup Rate

Purpose: With this test we determined the maximum rate at which new TCP connections can be established.

The analyzer increased the connection setup rate until the targeted session number is reached.

Throughput and Maximum Concurrent Connections

Purpose: With this test we determined the maximum throughput performance at a high but realistic number of concurrent connections.

Using the traffic profile described in table 6, the analyzer increased the transactions and the resulting bandwidth until session timeouts and TCP failures exceed the allowed values described in Table 8.We defined the following parameters for healthy session:

TABLE 8. Per Subscriber Connection Quality

Monitored Parameter	Maximum Value Accepted
Maximum Response Time per Page	3,000 msec
HTTP Errors	0
User Aborts	0.01%
TCP Connection Error	0
TCP Timeout	0
Avalanche memory pool	<=95% of physical RAM
Avalanche CPU Utilization	<= 97%
TX L2 Server Packet rate - RX Client packet Rate	< 3% Loss

RESULTS

Figure 4 illustrates the traffic load profile configuration as described in Table 5. Specifically it shows the 5 steps, steady-state and ramp-down phases of the test.

FIGURE 4. **Traffic Load Profile - Firewall forwarding performance**

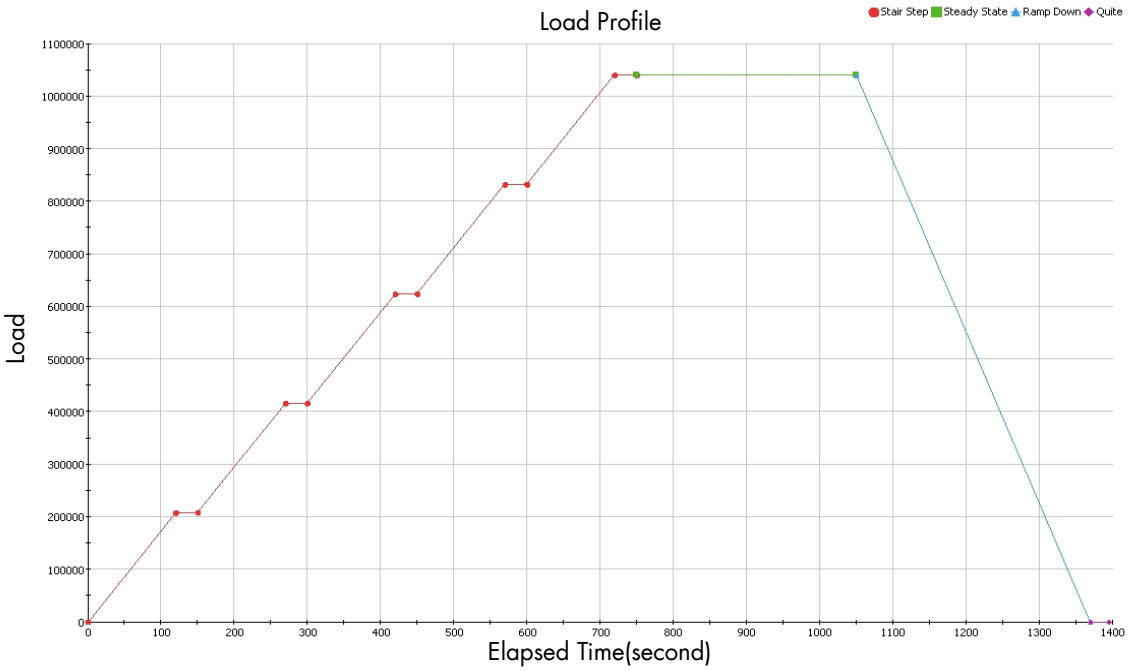

Figure 5 shows the distribution of simulated alive users load profile.

FIGURE 5. **SIM Users Alive Users- Firewall forwarding performance**

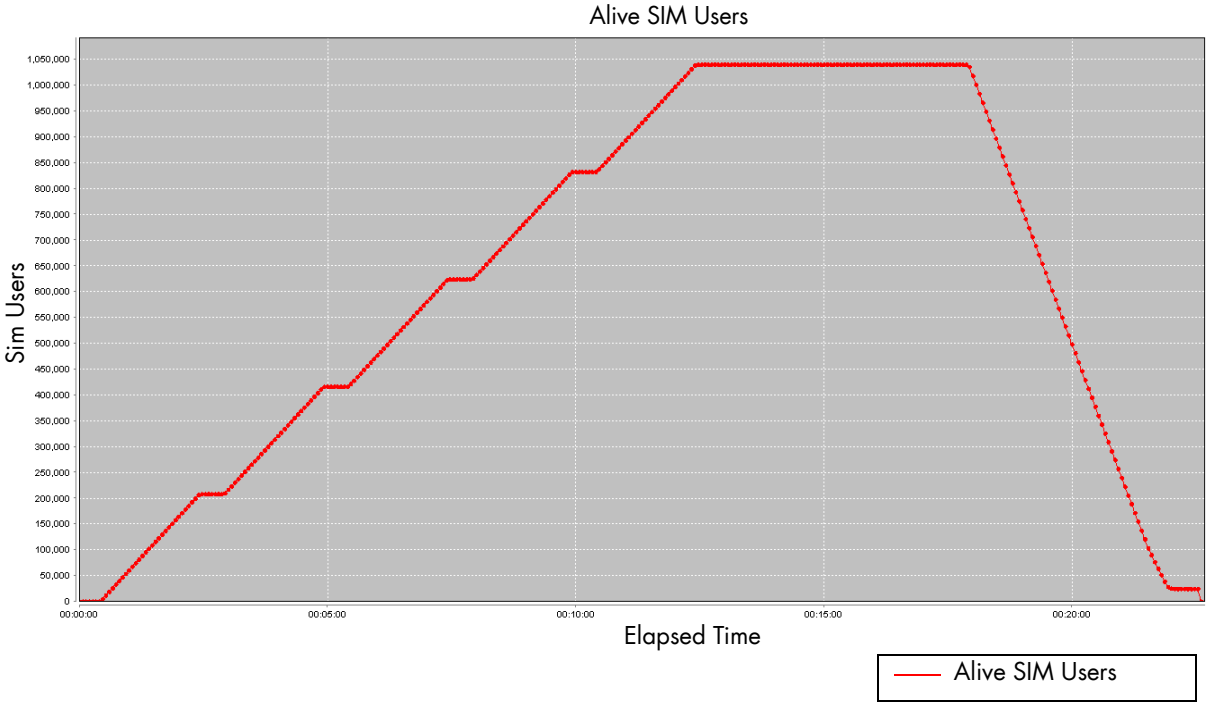

FIGURE 6. **User Stats - Firewall forwarding performance**

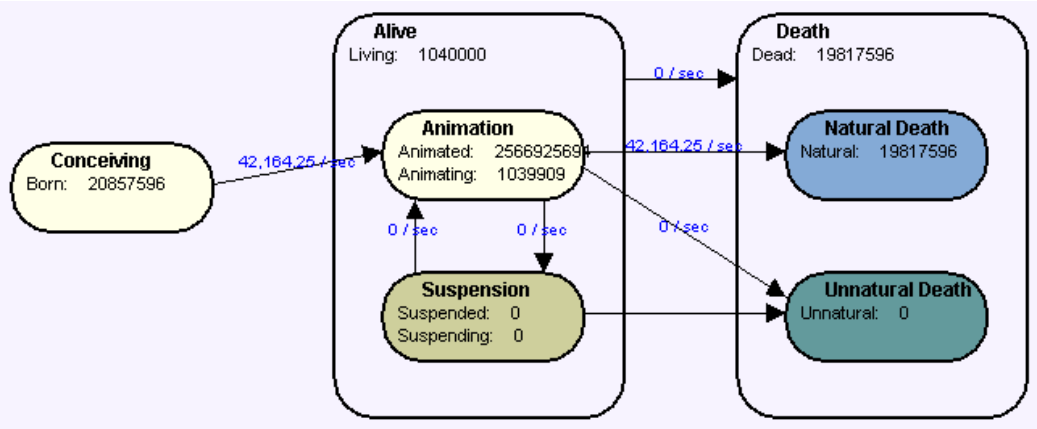

Figure 6 describes the user state during the steady phase. The average new SIMUser generated rate was approximately 42,184.25 per second.

Test Run 1(Firewall only - NAT and IPS Disabled)

In this test run only the firewall function was active on the DUT. Both NAT and IPS were deactivated. Measurements are only meaningful in the steady state phase when all TCP connections are active.

The analyzer showed approximately 108 Gbit/s average Layer two throughput (bidirectional) when maximum traffic profile load was reached.

FIGURE 7. **Throughput - Firewall forwarding performance**

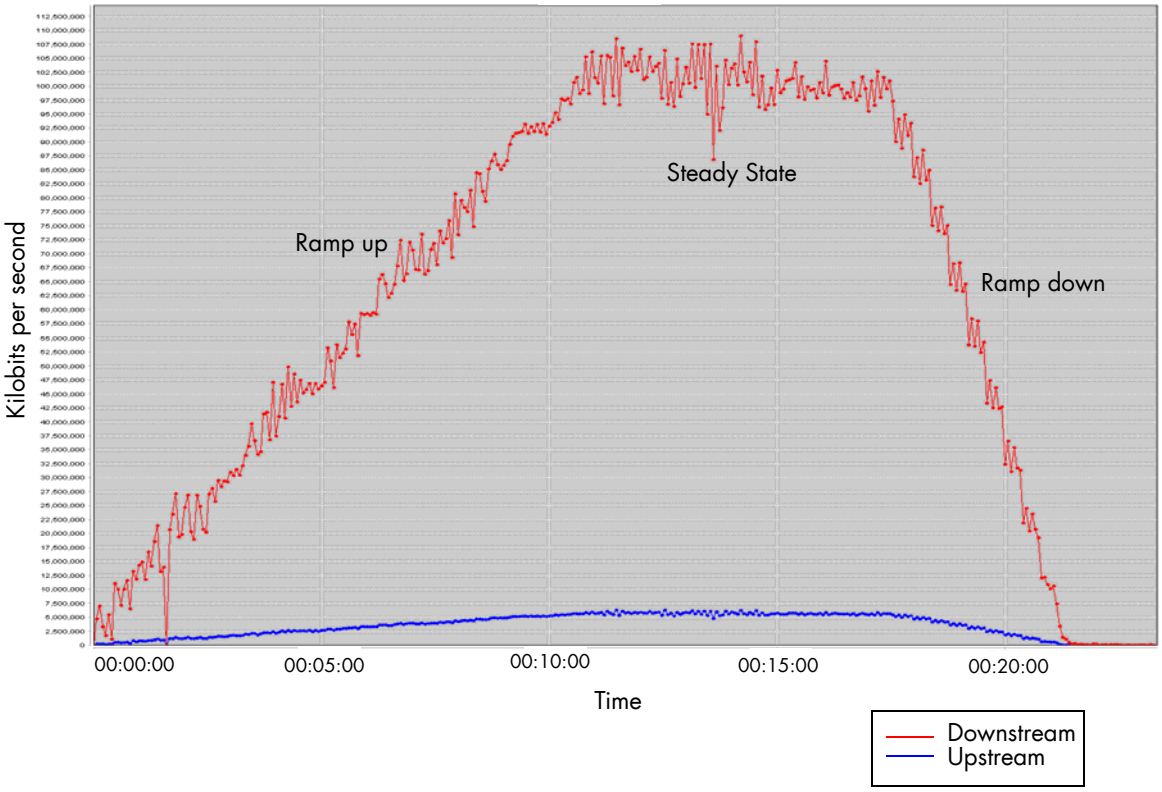

FIGURE 8. **Transaction per second - Firewall forwarding performance**

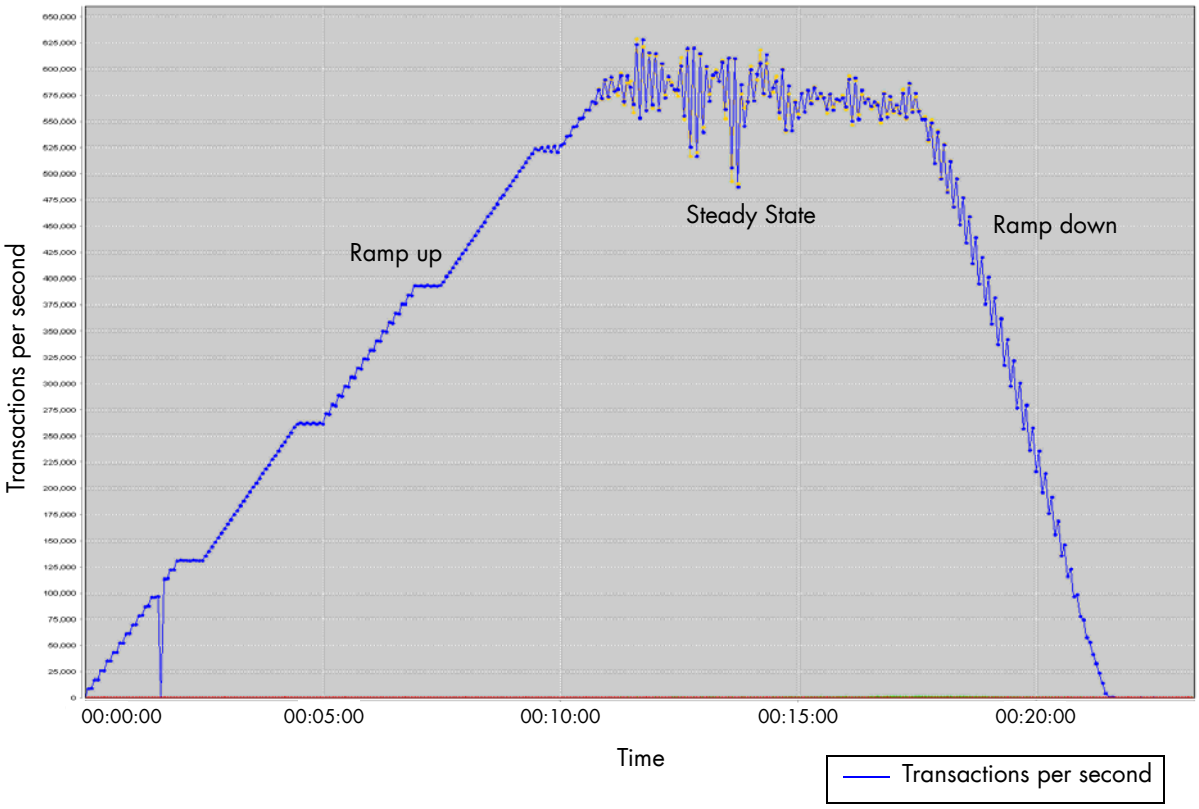

FIGURE 9. **TCP connections per second - Firewall forwarding performance**

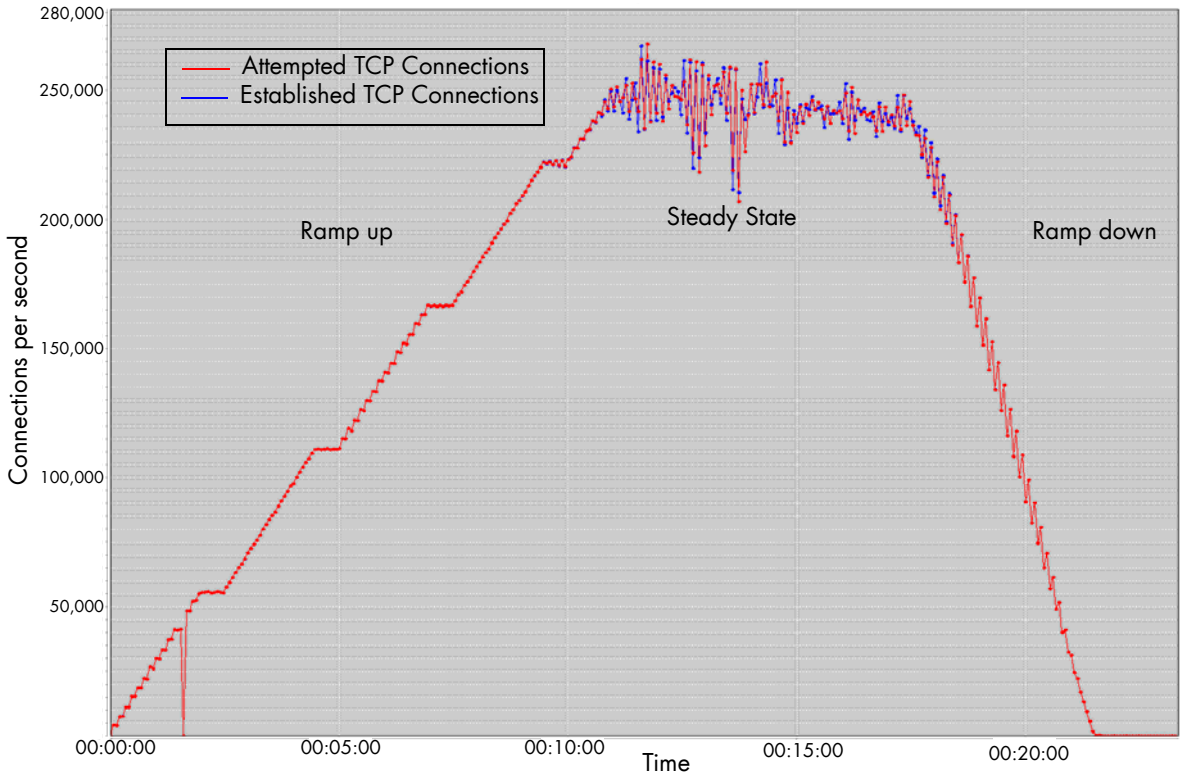

FIGURE 10. **URL Response Time - Firewall forwarding performance**

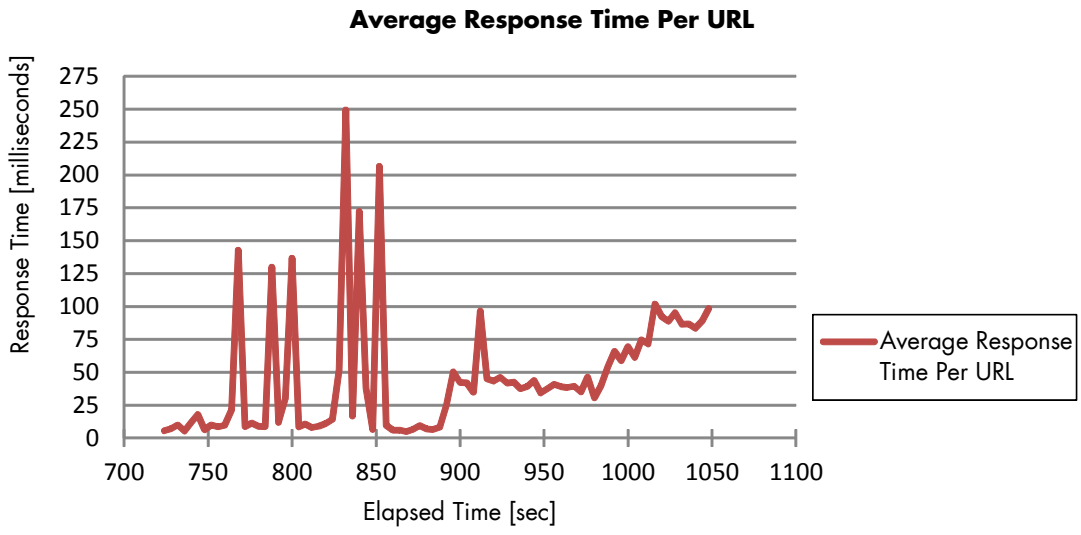

Test Run 2 (Only NAT enabled)

Figure 11 describes the throughput performance of the DUT when NAT function was enabled. The analyzer showed approximately 107 Gbit/s average Layer two throughput (bidirectional) when maximum traffic profile load was reached.

FIGURE 11. **Throughput - Firewall with NAT**

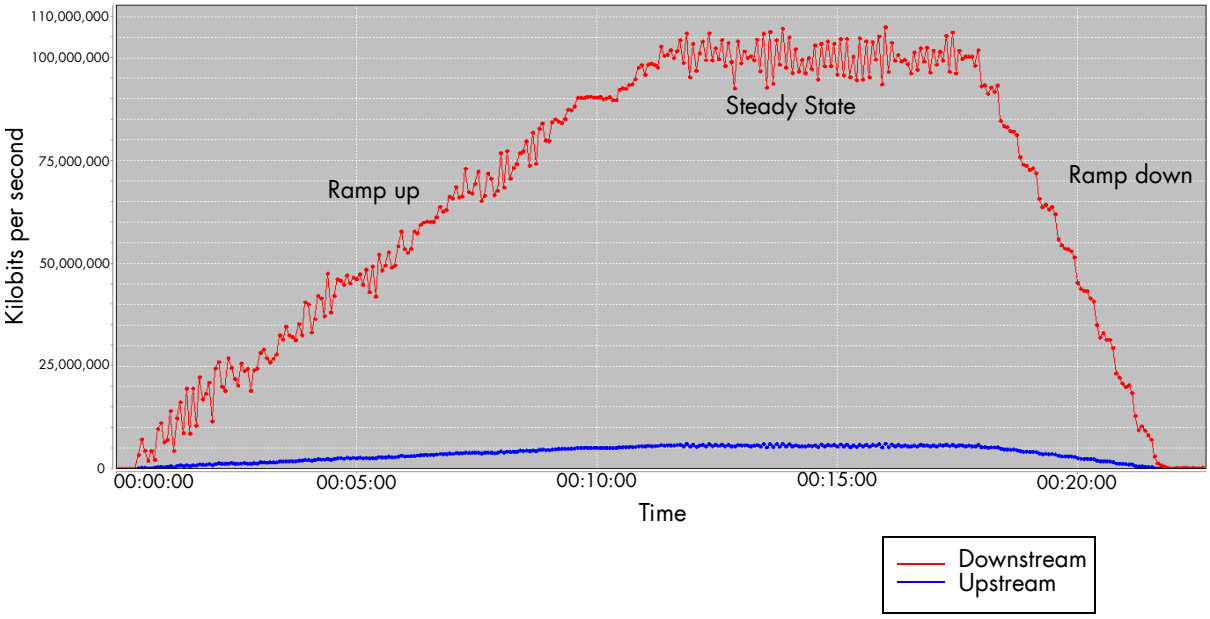

Figure 12 shows the distribution of application layer objects transaction performance of DUT. At the steady state, the DUT supported approximately 585,158 objects transactions per second.

FIGURE 12. **Transactions per second - Firewall with NAT**

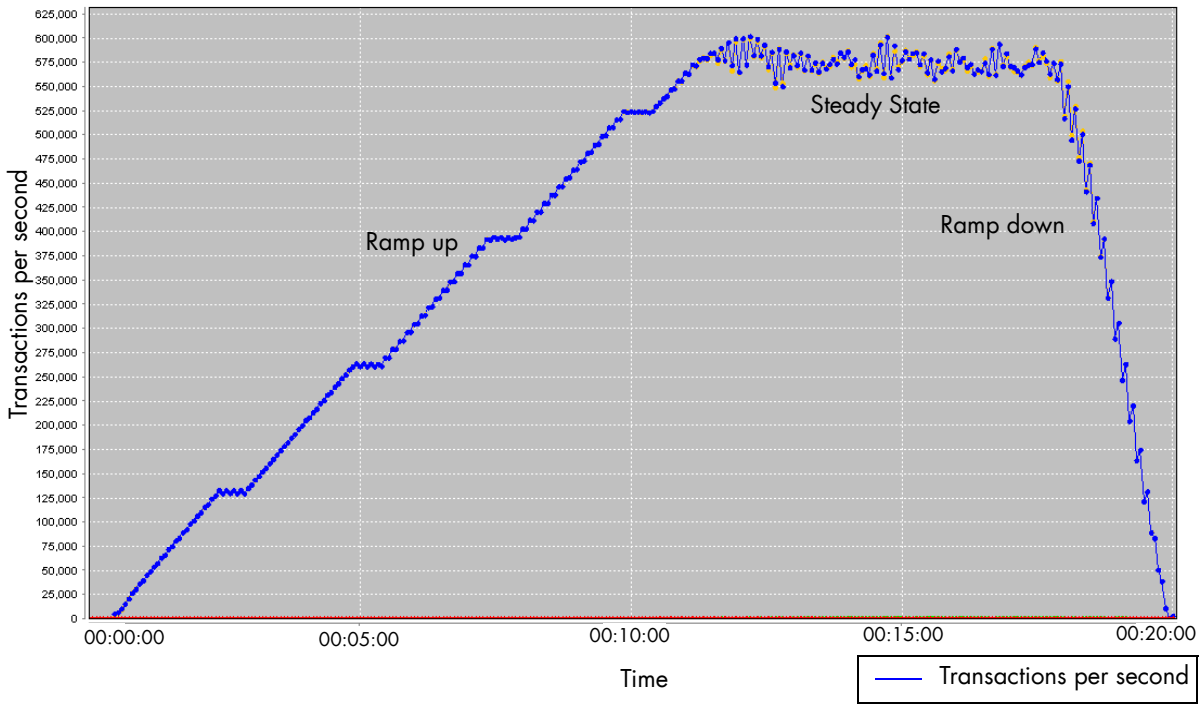

FIGURE 13. **TCP connections per second - Firewall with NAT**

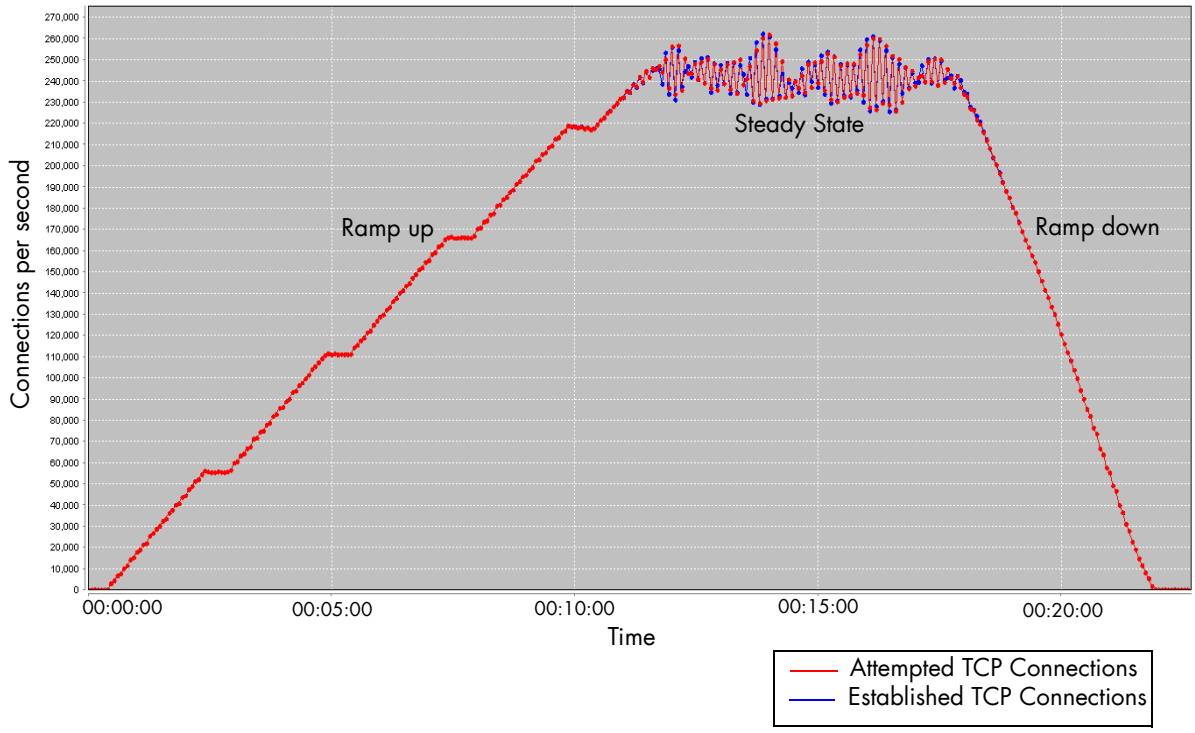

FIGURE 14. **URL Response Time - Firewall forwarding performance**

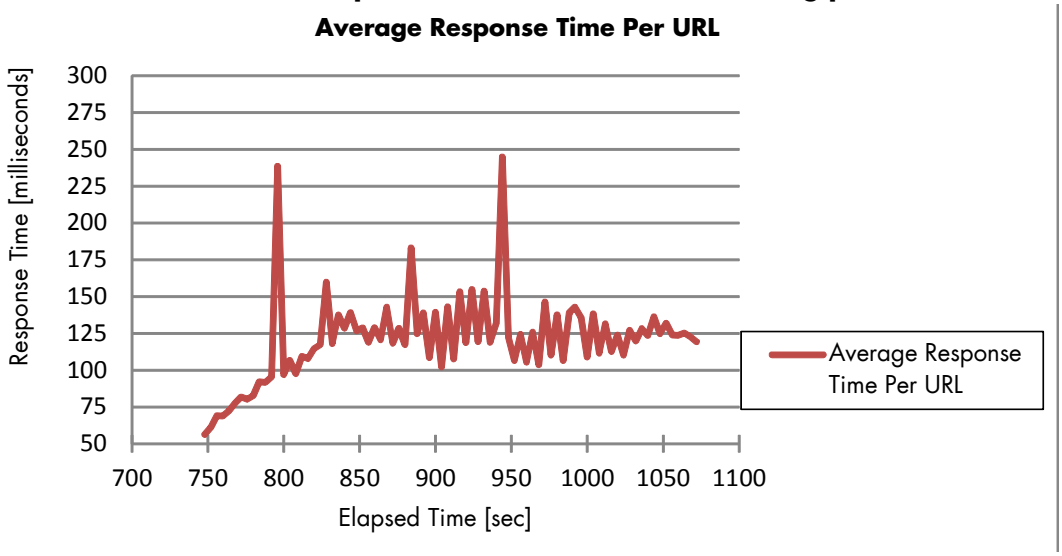

Test Run 3 (IPS only enabled)

Figure 15 describes the throughput performance of the DUT when Intrusion Prevention System (IPS) was enabled. The IPS configuration used the default option that is pre-packaged with the Check Point solution. The analyzer showed approximately 106.5 Gbit/s average Layer two throughput (bidirectional) when maximum traffic profile load was reached.

FIGURE 15. **Throughput - Firewall with IPS**

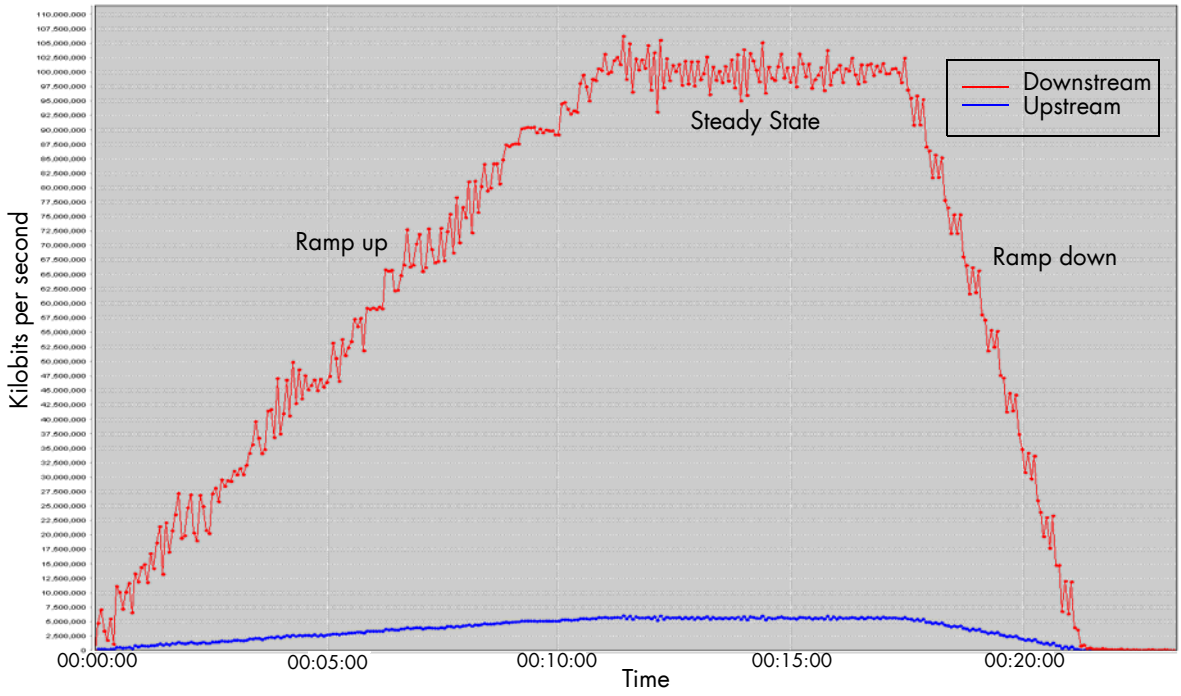

FIGURE 16. **Transactions per second - Firewall with IPS**

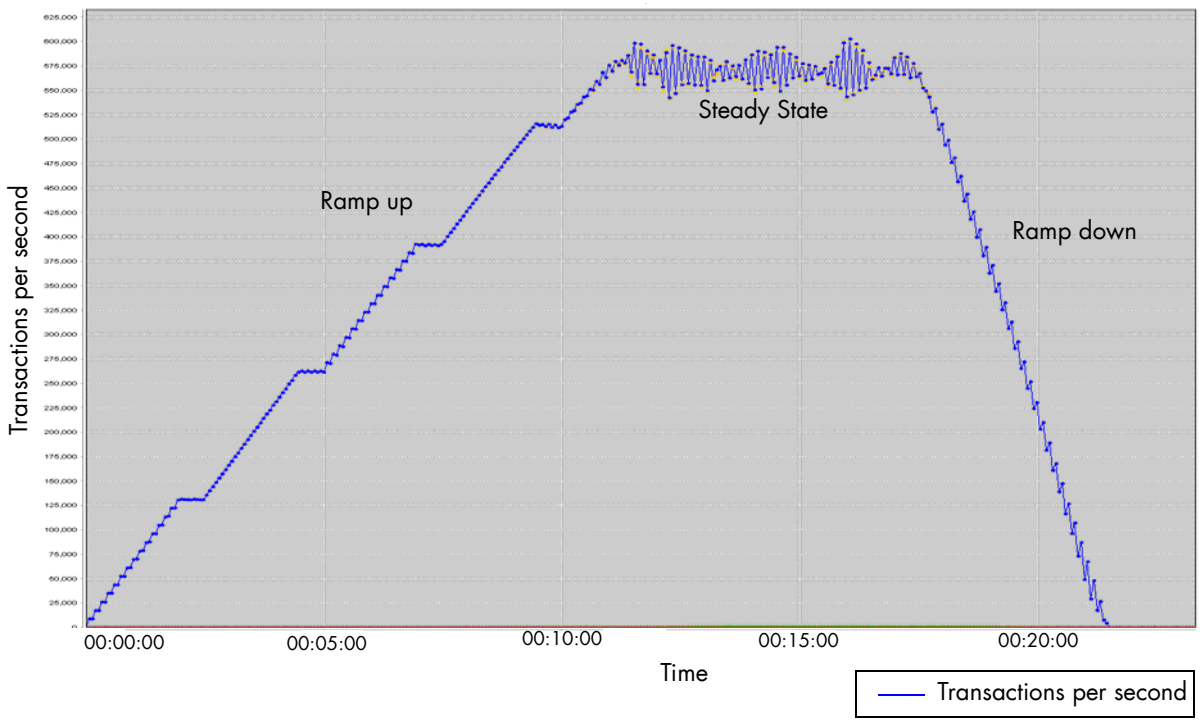

FIGURE 17. **TCP connections per second - Firewall with IPS**

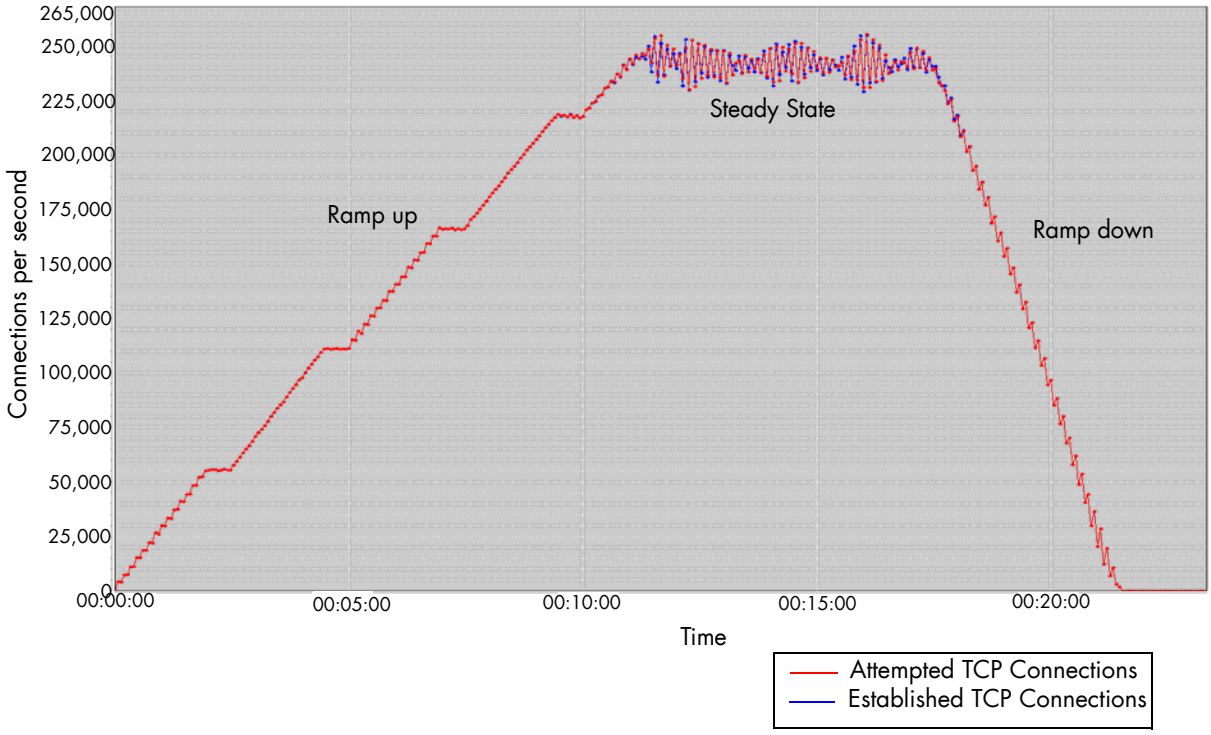

FIGURE 18. **URL Response Time - Firewall forwarding performance**

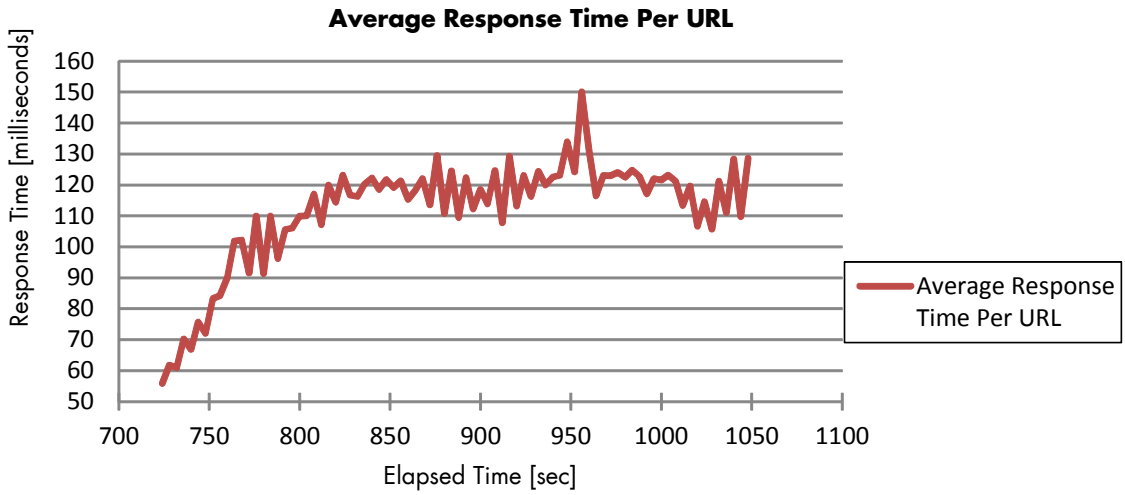

Test Run 4 (NAT and IPS Enabled)

The following figures depict the results of the same test configuration with both Network Address Translation (NAT) and Intrusion Prevention System (IPS) were activated.

Figure 19 shows the throughput performance of the DUT when Intrusion Prevention System (IPS) and Network Address Translation (NAT) were enabled. Both services were using the same configuration described in the previous two test cases. In this test run both services were active at the same time. We recorded approximately 105.5 Gbit/s of throughput during the steady state in this test.

FIGURE 19. **Throughput - Firewall with NAT and IPS**

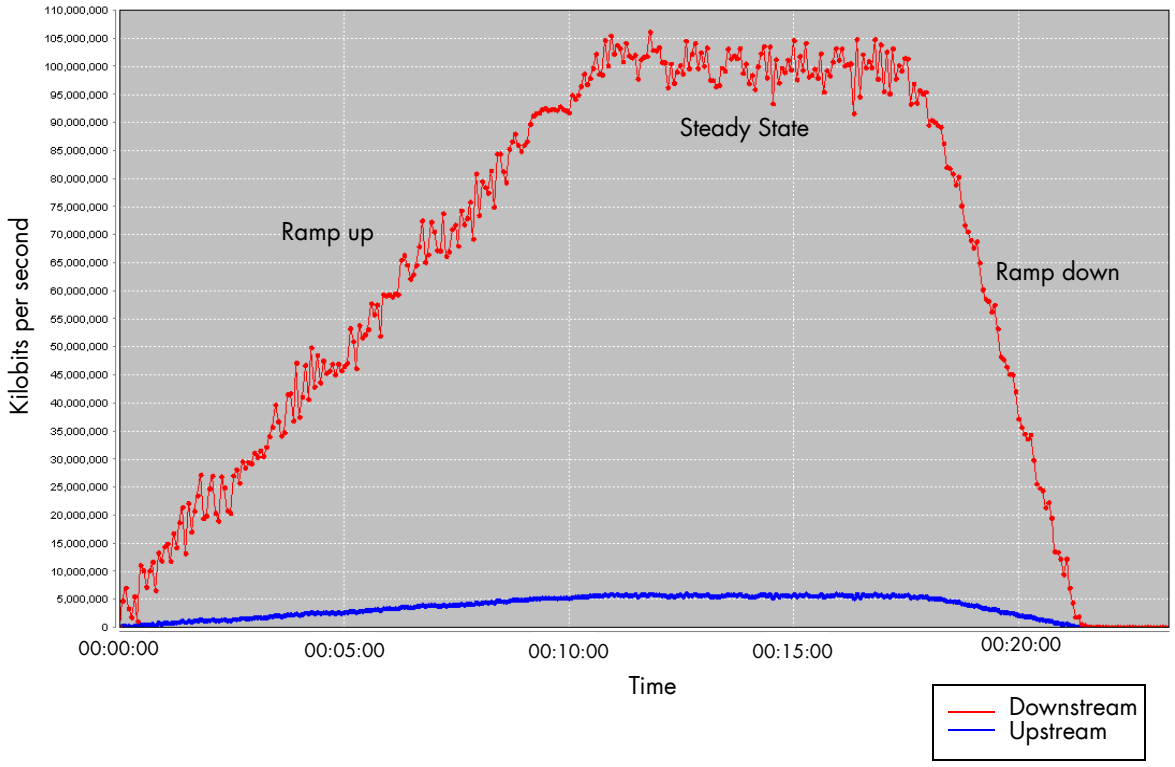

FIGURE 20. **Transactions per second - Firewall with NAT and IPS**

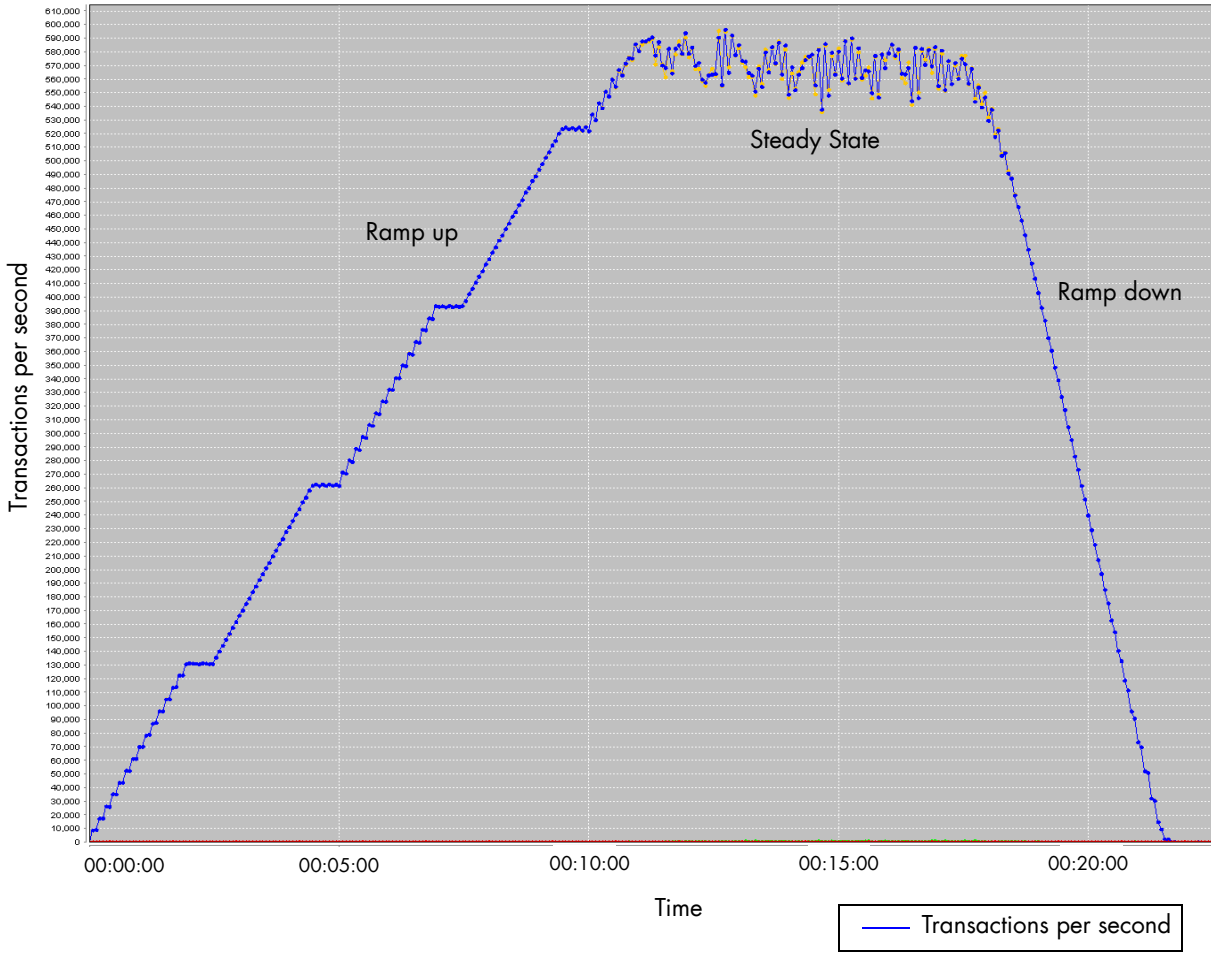

FIGURE 21. **TCP connections per second - Firewall with NAT and IPS**

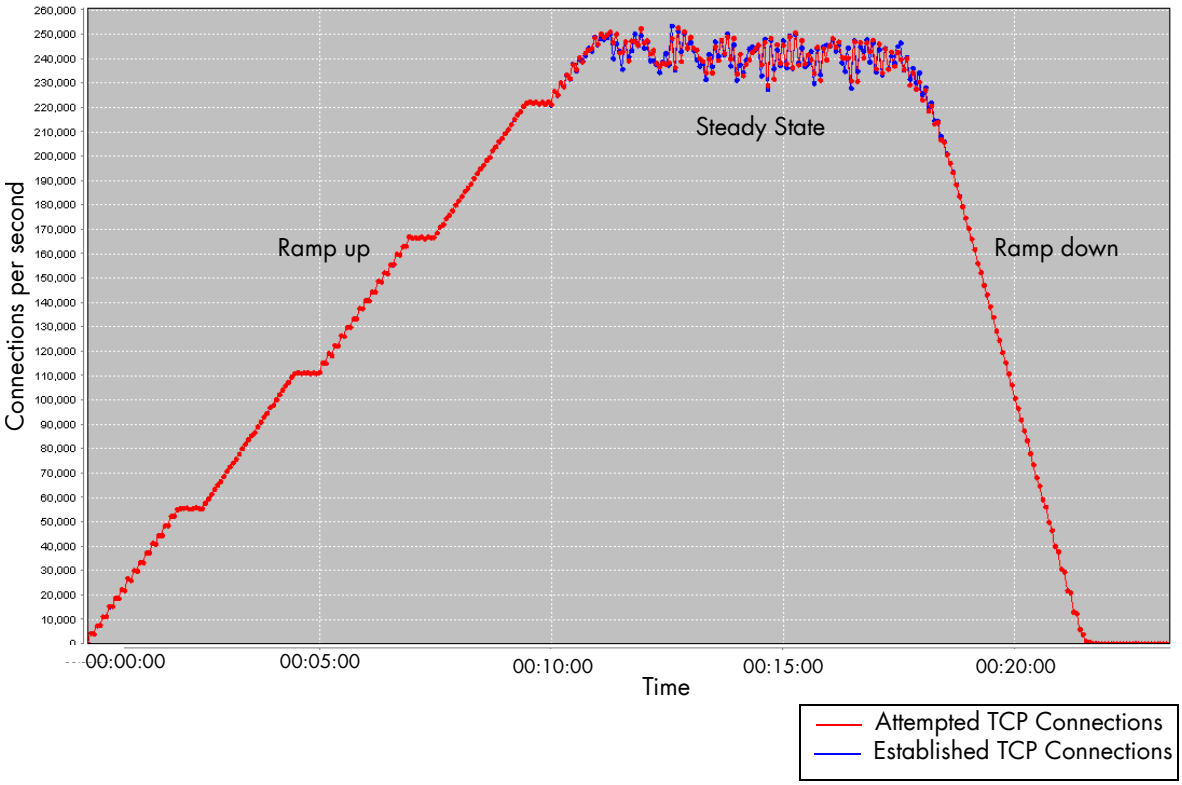

FIGURE 22. **URL Response Time- Firewall forwarding performance**

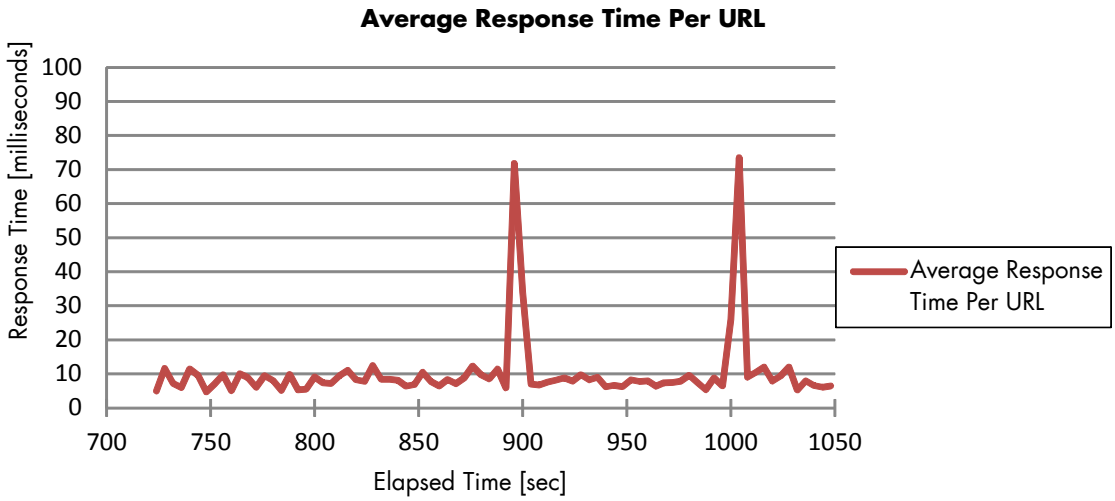

TABLE 9. Results - Firewall forwarding performance

	Run1 - Firewall Only	Run2 - Firewall and NAT	Run3 - Firewall and IPS	Run4 - Firewall NAT and IPS
Established TCP Connections	210,270,881	207,996,911	208,058,410	210,830,122
Average downstream throughput [Gbit/s][a]	101.8	101.5	100.9	100.0
Average upstream throughput [Gbit/s][a]	5.5	5.6	5.6	5.5
Total throughput [Gbit/s]	107.2	107.1	106.5	105.5
CPU utilization (APM / NPM) [%]	58 / 99	60 / 99	61 / 99	65/ 99
TCP connection Setup Rate per second[b]	244,166	247,152	243,616	241,595
Layer 7 Object Transaction per second[c]	577,263	585,158	576,686	578,151
URL Average Response Time [ms]	45.8	121.0	111.8	10.2
Maximum Concurrent connection	4,045,337	4,021,164	4,014,250	3,980,403

a. As reported by the analyzer; throughput is defined as the point in time in which the maximum desired load is equal to the current load trended with the incoming L2 and outgoing L2 traffic.

b. TCP connection setup rate was measured approximately when the traffic profile load is in steady state.

c. Layer 7 Object Transaction Rate was measured approximately when the traffic profile load was in steady state

Test Run 5 (TCP Session Setup Rate - NAT and IPS enabled)

This test case was executed in line with the first goal of tests: evaluate the TCP session setup rate that the device under test could support.

FIGURE 23. **TCP session setup rate - Firewall with NAT and IPS**

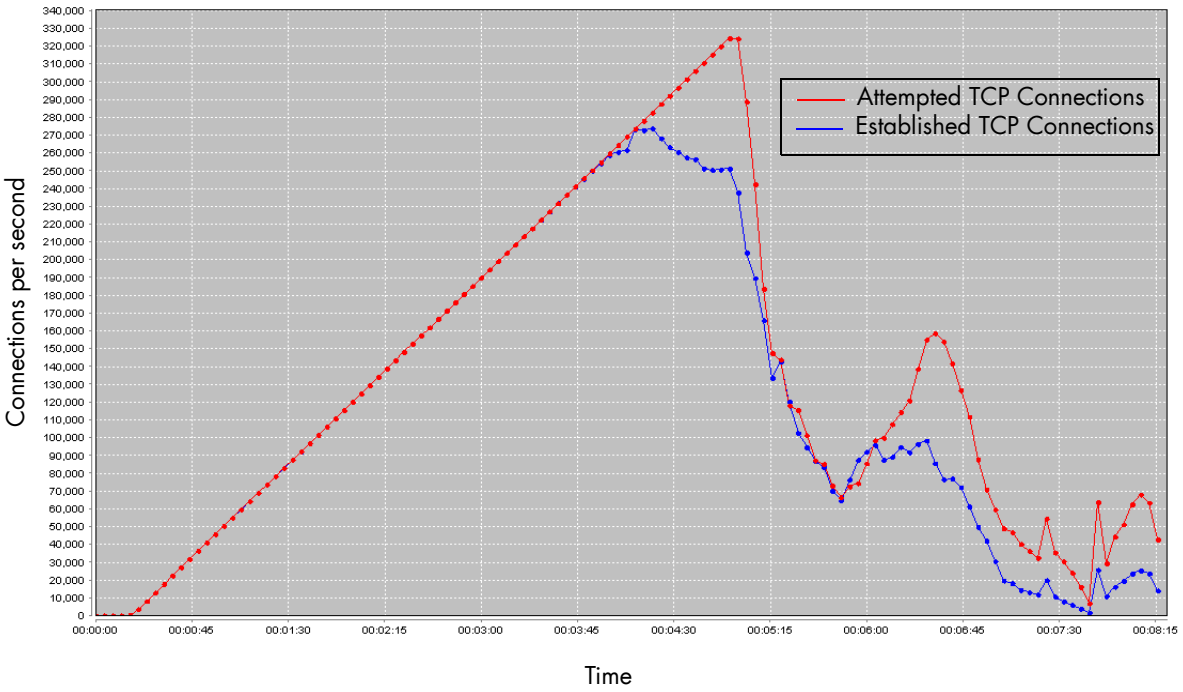

Figure 23 provides a useful visualization to the methodology such a test uses. The tester is configured to a certain value of TCP sessions rate to achieve (the goal). Typically the goal is set to a value higher than the one the vendor believes he or she could support. This is represented by the red line above. The blue line represents the rate at which TCP sessions could be setup successfully. Since the tester continues trying to increase the rate even as the device under test can not support a higher activation rate, the tester affects the DUT in such a way that the performance decreases after the maximum is reached (see in the downwards slope after 274,000 TCP sessions per second.

FIGURE 24. Total Established TCP connections

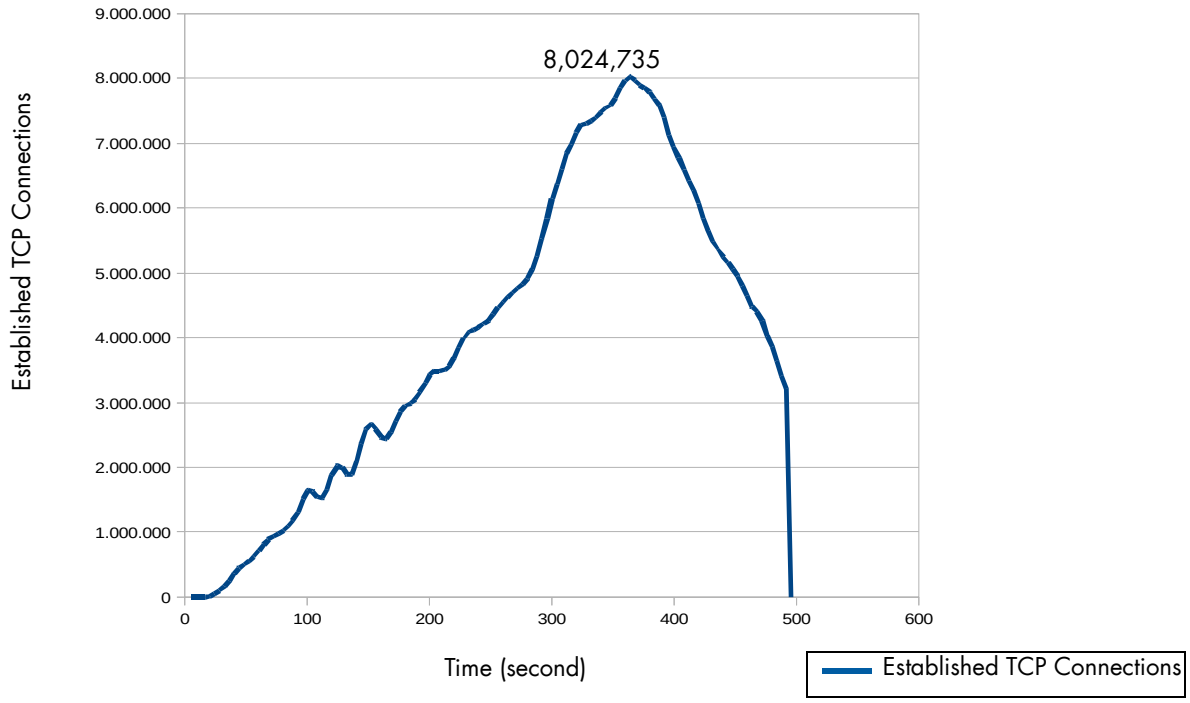

TABLE 10. Results - TCP session Setup Rate

Step	Direction	Expected Value	Observed Value
1	What was the total number of sessions establish	6,400,000	8,024,735
	maximum connection rate per second?	250,000-260,00	274,464
2	What was the DUT's CPU and memory usage observed?	Record values	APM 37% NPM 99%

COMMENTS

Initially the intention of the NAT use case was to hide the subscriber IP address behind a provider IP address. This feature is called "Hide NAT" in the firewall configuration. Port numbers were going to be used by the firewall to identify the responses destined to the subscribers.

In pre-testing we found out that due to the test setup, in which a relatively low number of web servers were configured (48 web servers), there was a small likelihood that an APM would receive traffic which would then ne dropped because that firewall process would know nothing about the connection or NAT entry. The solution to this would have been to create multiple NAT addresses on each firewall blade. Unfortunately this process would have been time consuming and would not have fit within the test schedule.

In order to bypass this issue Crossbeam configured a different NAT mode, one that would be used in a NAT64 setup - the subscriber IP would be translated into another IP address which will be presented to the Internet. While this type of NAT does not solve such problems as IP address depletion, it does allow a service provider to use IPv6 addresses internally, while still enabling subscribers to access the IPv4 world.

INTERPRETATION

For interpretation please see Heavy Reading's article.