EANTC

EANTC Independent Test Report

F5 BIG-IP Advanced Firewall Manager on HPE ProLiant Servers Performance Benchmarking

June 2019

intel

Hewlett Packard Enterprise

f5

Red Hat

Introduction

EANTC has been commissioned by Intel to evaluate the performance of a virtualized SGi firewall provided by F5 Networks and hosted on HPE ProLiant DL360 Gen10 servers provided by HPE.

Mobile service providers (MSPs) have always protected their mobile core infrastructure against threats from the Internet. It is standard practice to deploy firewall functions between the mobile core network and the external Packet Data Network (PDN), also known as SGi interface. Figure 1 depicts the typical placement of the SGi firewall in LTE networks.

Recently, three types of challenges are changing the market for SGi firewalls:

1. MSPs virtualize mobile core infrastructure; thus, firewalls need to be virtualized as well.
2. Mobile network traffic volumes are growing strongly; firewalls are expected to grow alike and to perform adequately.
3. Distributed Denial of Service (DDoS) attacks and cyber-security threats are becoming more advanced. Firewalls need to implement advanced detection and protection mechanisms.

In this context, manufacturing a virtualized SGi firewall is a non-trivial undertaking. EANTC set out to validate a few key cornerstones of data-plane performance of the F5 Gi-LAN firewall solution.

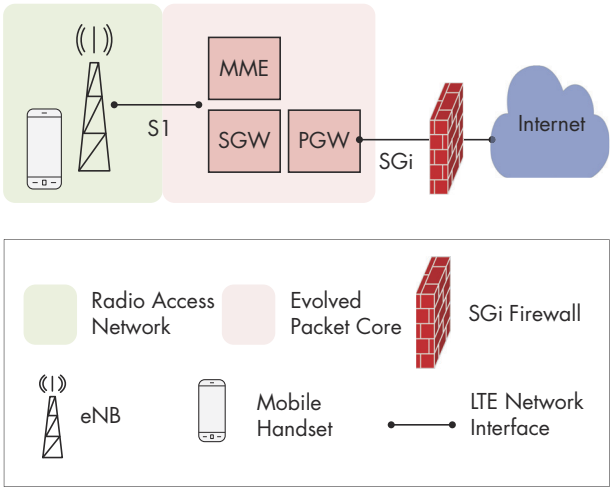

Figure 1: SGi Firewall Placement in LTE Networks

Test Scenario

Traffic between the mobile network and the Internet will pass through the SGi firewall and will be inspected by the pre-configured security Layer 3 and Layer 4 access rules and the DDoS attack detection module.

Test Highlights

- 40 Gbit/s HTTP throughput with a mix object size
- 800,000 TCP connection per second with 1 KByte object size
- 18 million concurrent TCP connection capacity

A major portion of the Internet traffic are web-based applications using HTTP/HTTPS protocols which are transported over TCP. F5 configured the firewall to analyze traffic on the TCP layer only; consequently, we emulated HTTP/TCP throughout in this test.

Figure 2 shows the logical network topology of the test bed.

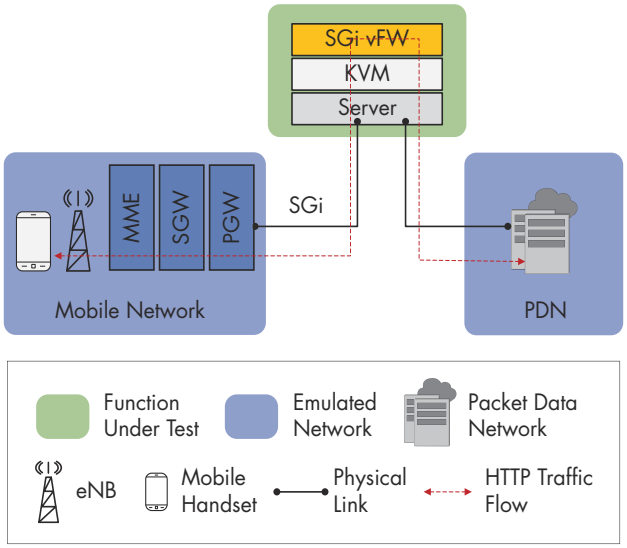

Figure 2: Test Bed Logical Network Topology

Executive Summary

In the EANTC test, the F5 BIG-IP Advanced Firewall Manager (AFM) was able to handle up to 18 million TCP sessions. It sustained a TCP connection setup rate of 800,000 connections per second with 1 KByte object sizes. With the assistance of network acceleration Single Root Input/Output Virtualization (SR-IOV), traffic throughput results obtained are considered as competent as compared to the Physical Network Function (PNF). Throughput was tested both in an IPv4-only scenario and with a 50:50 mix of IPv4 and IPv6 traffic.

Detailed Test Setup

Component	Description
HPE ProLiant DL360 Gen10	2 x Sockets Intel® Xeon® Gold 6152 CPU@ 2.10GHz (microcode: 0x200005e) 12x 32GB DDR4 RAM 2 x HPE 480GB SATA 6G RI SFF/ LFF SC DS SSD 2 x HPE 1.92TB NVMe x4 RI SFF SCN DS SSD 1x HPE Ethernet 40Gb-2port 565QSFP+ Adapter (Intel XL710 controller) 1x HPE Ethernet 10Gb 2-port 562FLR-SFP+ Adapter Server Platform Services (SPS) Firmware: 4.0.4.393 System ROM: U32 v1.42 Innovation Engine (IE) Firmware: 0.1.6.1
Host OS	Red Hat Enterprise Linux 7.6
KVM Hypervisor	QEMU emulator v2.5.0
HPE Ethernet 40Gb-2port 565QSFP+ Adapter (Intel XL710 controller)	Driver: i40e v2.7.29 Firmware: 6.80
F5 BIG-IP Advanced Firewall Manager (AFM)	BIG-IP-14.1.0-0.0.116
OpenStack	Red Hat OpenStack Platform13
HPE 5900 Series Switch JC772A	7.1.045, Release 2422P01
Spirent Avalanche Commander C100-S3	chassis OS v4.96.0172
Spirent Avalanche	v4.96 build 1306 32bit

Table 1: Test Bed Components and Description

The test bed infrastructure consisted of three HPE ProLiant DL360 Gen 10 Servers. Out of the three servers, one was setup as Red Hat OpenStack controller node, second one as director node and the third server was utilized as a compute node to host the VNF under test.

Each HPE ProLiant DL360 Gen 10 Server was equipped with two Intel® Xeon® Gold 6152 processors, one HPE Ethernet 40Gb-2port 565QSFP+ Adapter (Intel XL710 controller) and one 2-port 10GbE HPE built-in 562FLR-SFP+ adapter.

Red Hat Enterprise Linux 7.6 was running on all three servers as the host Operating System (OS). Kernel-based Virtual Machine (KVM) hypervisors were installed as the hypervisor. Based on F5's recommendation for NFV performance and security optimization, hyper-threading was disabled on the compute node. The compute node and the traffic generator were connected through a HPE 5900 switch series JC772A.

Red Hat OpenStack Platform 13 was deployed on all three HPE ProLiant DL360 Gen 10 servers. Figure 3 shows the Red Hat OpenStack architecture. The undercloud is the main director node, which allows managing and provisioning for the compute nodes that form the Network Functions Virtualization Infrastructure (NFVI) cluster. The controller node provides administration and networking for the OpenStack environment. The compute node provides computing resources for the OpenStack environment.

The Function Under Test (FUT) was a BIG-IP AFM High-Performance virtualized edition (VE) solution provided by F5 Networks. According to F5 Networks, Big-IP AFM is an integrated security solution which provides L3/L4 security access control and defends against DDoS attacks. The F5 BIG-IP AFM was instantiated based on the configurations that are listed in Table 2.

Configuration Parameter	Value
Assigned CPU Cores	16
Assigned Memory Volume	32 GB
Assigned Physical Network Interfaces	2x40GbE with SR-IOV support
Assigned Virtual Function Network Interfaces	2 Virtual function in total 2 Physical function in total
Assigned Storage Space	100 Gigabyte
Firewall Deployment Mode	Routed
L4 Access Control List (ACL) Rules	1,000

Table 2: Configuration Parameters

Flood	Bad Header	Other
IPv4/IPv6 Fragment	TTL <= tunable	Host Unreachable
ICMP Fragment	IP Option Frames	TIDCMP
ICMPv4/ ICMPv6 flood	IPv6 Hop Count	Sweep
TCP SYN/ SYN ACK/ SYN Oversize	IPv6 Extension Header Too Large	
Non TCP Connection	IPv6 Extended Header Frames	
TCP RST/ PUSH/ HALF OPEN/ Window Size	IPv6 Too Many Extension Headers	
UDP	Unknown TCP Option Type	
	Option Present With Illegal Length	
	TCP Option Overruns TCP Header	
	TCP Flags-Bad URG	

Table 3: DDoS Enabled Vectors

The DDoS attack mitigation was enabled on the F5 BIG-IP AFM with fully automated detection and threshold monitoring during all test cases. As listed in Table 3, the DDoS configurations included 27 different vectors that specifically protect the network from flood, bad header and other attacks.

Test Equipment

Tests were conducted using Spirent Communications C100-S3 high-performance appliance hardware with Avalanche software. According to Spirent, Avalanche provides an assessment framework to test and stress next-generation firewalls and other networking solutions with stateful application and encrypted traffic on interfaces from Gigabit Ethernet to 100 Gigabit Ethernet. The C100-S3 also supports Spirent's CyberFlood assessment solution for advanced mixed traffic, attack, malware and NetSecOPEN methodologies.

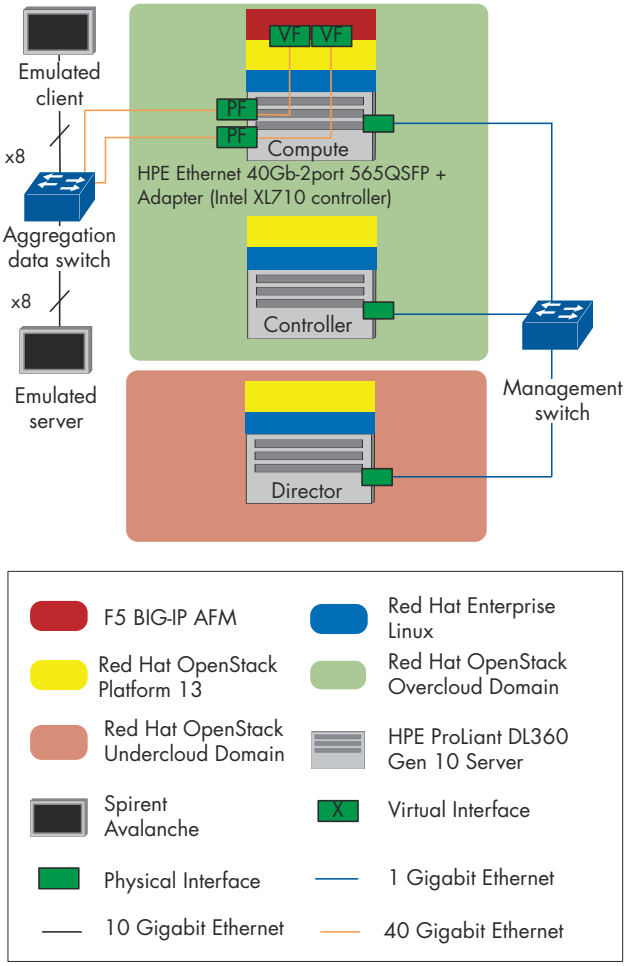

Figure 3: Physical Test Bed Topology

Test Results

One of the guiding factors for any EANTC benchmark is reproducibility: We aim to share sufficient details to enable readers to reproduce our test setup and results independently. In the firewall test area, a large number of test configuration and methodology parameters needs to be defined to achieve this goal. Fortunately, the IETF is in the process of standardizing terminology and methodology for Next Generation Firewall benchmarking. Our tests followed the IETF draft-ietf-bmwg-ngfw-performance-00 (work in progress). Any configuration details not mentioned here can be referenced from the IETF draft.

The configuration of F5 BIG-IP AFM remained identical during all performance tests. The tests consisted of three different phases: ramp up, steady and ramp down.

In the ramp up phase, the traffic ramped up slowly to reach the target key performance indicator (KPI) values. The traffic was continually flowing and stable during the steady phase. All KPIs measured in the steady phase were documented and used as main

source of the result. The traffic then ramped down slowly to finalize the test and bring the FUT back into idle mode.

For the throughput and CPS tests, ramp up time was 180 s, steady state time was 600 s, and ramp down time was 180 s. EANTC executed each test case once.

HTTP/TCP Throughput

The purpose of this test case was to measure the maximum Layer 4 throughput performance of the F5 BIG-IP AFM. It was configured to analyze traffic on the TCP layer only; consequently, we emulated HTTP/TCP throughout. The maximum sustained throughput helps to dimension the network capacity efficiently and design the network adequately.

The selected object sizes were 1 KByte and a MIX object size. The MIX object size used for the test is derived from IETF (draft-ietf-bmwg-ngfw-performance-00) and listed in Table 4 below. Each TCP connection had 10 transactions; it closed with FIN immediately after these 10 transactions.

Object Size	Weight
0.2 KByte	1
6 KByte	1
8 KByte	1
9 KByte	1
10 KByte	1
25 KByte	1
26 KByte	1
35 KByte	1
59 KByte	1
347 KByte	1

Table 4: MIX Object Size Distribution

Table 5 shows the achieved throughput results and other KPIs during the steady phase with 1 KByte sized objects; Table 6 shows the respective throughput for mixed object sizes.

KPI	Minimum	Average	Maximum
Throughput	24.3 Gbit/s	24.7 Gbit/s	25.2 Gbit/s
HTTP transactions per second	2.08 Million	2.11 Million	2.15 Million
TCP conns per second	207,559	211,460	215,491
CPU utilization	75%	-	85%
Memory utilization	16.2%	16.2%	16.2%

Table 5: TCP Throughput Test Results with HTTP 1 KByte Object Size

KPI	Minimum	Average	Maximum
Throughput	38.6 Gbit/s	39.7 Gbit/s	40.0 Gbit/s
HTTP transactions per second	83,432	85,664	86,381
TCP conns per second	8,339	8,568	8,648
CPU utilization	48%	-	55%
Memory utilization	16.2%	16.2%	16.2%

Table 6: TCP Throughput Test Results with HTTP Mixed Object Size

TCP Connections Per Second (CPS)

In this test case, we measured the maximum connection setups per second performance of the F5 BIG-IP AFM. Each TCP connection had only one transaction. The TCP connection closed with FIN immediately after this single transaction. This test was executed in two iterations, first with 1 KByte object size and second with 64 KByte object size.

Table 7 represents the results of CPS and other KPIs during the steady phase with 1 KByte object size; Table 8 represents the respective KPIs for 64 KByte object size.

KPI	Minimum	Average	Maximum
TCP Connection Per Second	746,887	810,480	839,539
Throughput	10.99 Gbit/s	11.93 Gbit/s	12.35 Gbit/s
TCP Time to First Byte	0.188 ms	0.79 ms	2061.65 ms
F5 BIG-IP AFM CPU utilization	93%	-	97%
F5 BIG-IP AFM memory utilization	16.2%	16.2%	16.2%

Table 7: TCP Connection Per Second Test with 1 KByte HTTP object size

KPI	Minimum	Average	Maximum
TCP Connection Per Second	67,031	69,111	70,364
Throughput	37.97 Gbit/s	39.14 Gbit/s	39.84G bit/s
TCP Time to First Byte	0.214 ms	1.24 ms	2000.99 ms
F5 BIG-IP AFM CPU utilization	55%	-	65%
F5 BIG-IP AFM memory utilization	16.2%	16.2%	16.2%

Table 8: TCP Connection Per Second Test with 64 KByte HTTP object size

TCP Concurrent Connection Capacity

In this test case, we verified the maximum number of Concurrent Connections (CC) that was supported by the F5 BIG-IP AFM. The maximum sustained CC gives a better understanding about capacity limits of the F5 BIG-IP AFM based on the assigned compute resources.

Each TCP connection had 10 transactions. The object size used was 1 KByte. We added think time between each transaction so that all TCP connections were kept open during the steady phase.

For the CC measurements, we established only the TCP sessions during the ramp up phase. The average CPS rate was around 400,000 connections/second.

The F5 BIG-IP AFM showed 18 million active sessions in the session table at the end of the ramp up phase. All the sessions remained open during the whole steady phase. No session was opened or closed during the steady phase. The traffic was continually flowing and stable during the steady phase.

Table 9 shows the maximum concurrent connection capacity and other KPIs during the stable phase.

KPI	Minimum	Average	Maximum
Throughput	4.76 Gbit/s	4.78 Gbit/s	4.8 Gbit/s
Concurrent TCP connections	18,000,000	18,000,000	18,000,000
Application transaction latency	<0.001 ms	0.651 ms	1,007 ms
Application transactions per second	398,466	400,233	401,970
F5 BIG-IP AFM CPU utilization	25%	-	40%
F5 BIG-IP AFM memory utilization	76.4%	76.4%	76.4%

Table 9: TCP concurrent connection capacity with 1 KByte HTTP object size

Conclusion

The data sheet for the F5 BIG-IP AFM VE claims 40 Gbit/s throughput performance when 1,000 firewall rules are configured and the DoS function is enabled. EANTC independently verified that these claims are correct: In our test, the FUT demonstrated 40 Gbit/s throughput with mixed HTTP object sizes. In addition, we measured 800,000 connections per second and (separately) 18 million CC, both with 1 KByte HTTP object size. We used Red Hat OpenStack Platform 13 as an orchestrator and SR-IOV as the virtualization acceleration technology.

Throughout the testing process, the F5 solution performed excellently and showed readiness for large-scale deployment scenarios. EANTC looks forward to facilitate additional tests of the BIG-IP AFM VE security protection functions and resilience capabilities in the future.

About EANTC

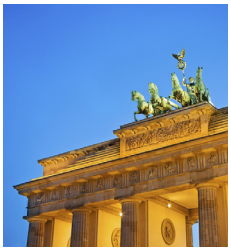

EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies.

Based in Berlin, the company offers vendor-neutral consultancy and realistic, reproducible high-quality testing services since 1991. Customers include leading network equipment manufacturers, tier 1 service providers, large enterprises and governments worldwide. EANTC's Proof of Concept, acceptance tests and network audits cover established and next-generation fixed and mobile network technologies.

Tests were conducted by European Advanced Networking Test Center (EANTC). Hardware configurations: three servers with dual Intel Xeon Gold 6152 processors running at 2.1 GHz with 22 cores, 384 Gigabits of RAM, and 40 GbE connections provided by one HPE® Ethernet 40Gb-2port 565QSFP+ Adapter and by one HPE® Ethernet 1Gb 4-port 331i Adapter. Software configurations: F5 BIG-IP AFM 14.1.0-0.0.116, Red Hat OpenStack Platform 13 and Red Hat Enterprise Linux 7.6.* Simulation of application protocol conducted using 2x Spirent Avalanche® C100-S3 appliances using 4xDual-port 10Gbps adapters.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Performance results are based on testing as of May 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/
[v1 20190619]