EANTC Independent Test Report

F5 BIG-IP Local Traffic Manager on HPE ProLiant Servers Performance Benchmarking

September 2019

Introduction

Application availability is one of the significant pillars to be considered in any advanced network design. An Application Delivery Controller (ADC); sometimes called application load balancer, is the network function which ensures the application is available, scalable and secure. Usually, the ADC is deployed as a reverse proxy to handle all client HTTPS requests, then decrypt and load balance the traffic to the pool of back-end servers as shown in Figure 1.

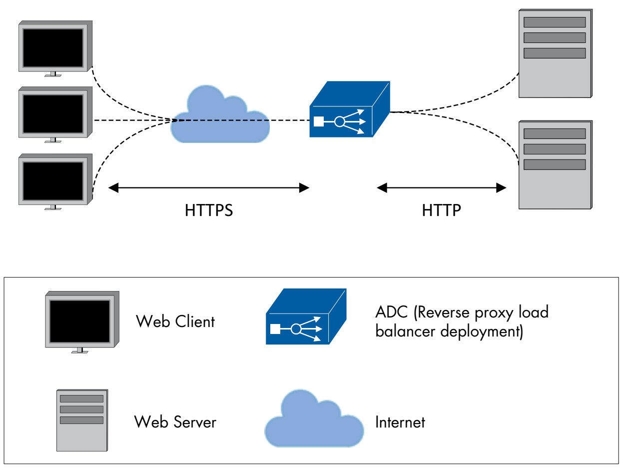

Figure 1: ADC Reverse Proxy Deployment

Virtualized versions of ADCs (vADC) have been adopted in many data center deployments where the majority of the workloads and applications are hosted in a virtualized environment. The vADC inherits a lot of the advantages from the general principles of virtualization, such as shorter deployment time, lower CapEx & OpEx, and ease of scale-in or scale-out operations. The performance of any VNF is mainly determined by the hosting Network Function Virtualization Infrastructure (NFVI) consisting of compute, storage and networking resources. The vADC is categorized as a compute- and network-intensive VNF: Compute resources represented by CPU performance and memory size have a direct influence on the performance of the VNF.

Recently, Intel released 2nd generation Intel® Xeon® Scalable processors. This Intel Xeon CPU release marks an optimization in Intel's "Process-Architecture-Optimization" model. To verify the expected performance enhancements over the previous Intel Xeon processor model, Intel commissioned EANTC to conduct this test for a reference VNF, the F5 Networks BIG-IP Local Traffic Manager (LTM) virtual edition.

Test Highlights

→ Throughput (Object size IMIX)
12.98 Gbit/s on Intel Xeon Gold 6152 CPU
14.35 Gbit/s on Intel Xeon Gold 6252N CPU
10.55% throughput improvement

→ Connections per second (Object size 64 KByte)
5,165 CPS on Intel Xeon Gold 6152 CPU
5,705 CPS on Intel Xeon Gold 6252N CPU
10.45% CPS improvement

Executive Summary

EANTC was commissioned by Intel to verify the performance gain of Intel Xeon Gold 6252N processors over Xeon Gold 6152 for a reference VNF which is F5 Networks BIG-IP LTM virtual edition. The BIG-IP LTM was configured as reverse proxy vADC to handle all the HTTPS requests from the clients, then decrypt and load balance the traffic to the pool of back end servers as shown in Figure 1.

Test Setup

Component	Description
3 x HPE ProLiant DL360 Gen10, each equipped identically as shown here	*Initial set up:* 2 x Intel® Xeon® Gold 6152 CPU@ 2.10GHz (microcode: 0x200005e) *Upgraded set up:* 2 x Intel® Xeon® Gold 6252N CPU@ 2.30GHz (microcode: 0x5000021) 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) Server Platform Services (SPS) Firmware: 4.1.4.251 System ROM: U32 v2.04 Innovation Engine (IE) Firmware: 0.2.0.11
Host OS	Red Hat Enterprise Linux 7.6 (Maipo)
KVM Hypervisor	QEMU emulator v2.12.0

Component	Description
HPE Ethernet 40Gb-2port 565QSFP+ Adapter (Intel XL710 controller)	Driver: i40e v2.9.21 Firmware: 6.80
F5 BIG-IP Local Traffic Manager (LTM)	BIGIP-14.1.0.3-0.0.6
OpenStack	Red Hat OpenStack Platform 13
HPE 5900 Series Switch JC772A	7.1.045, Release 2422P01
Spirent Avalanche Commander C100-S3	Chassis OS v4.98.0205
Spirent Avalanche	v4.98 build 1266 32bit

Table 1: Test Bed Components and Description

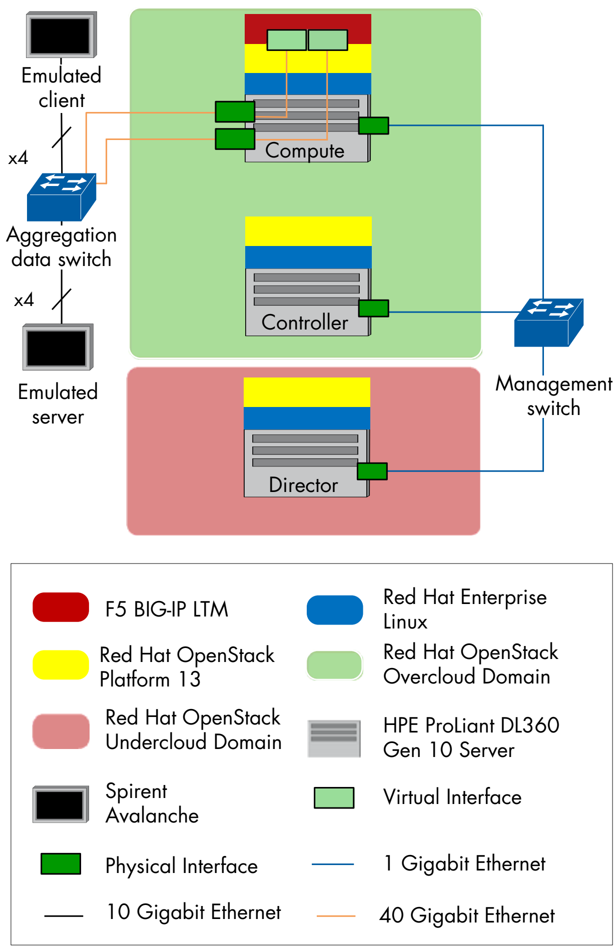

Figure 2: Test Bed Setup

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

Initially, each HPE ProLiant DL360 Gen 10 Server was equipped with two Intel® Xeon® Gold 6152 processors. For the comparison with the latest server CPU edition, we upgraded the servers using Intel® Xeon® Gold 6252N processors. All other hardware parts remained unchanged.

Red Hat Enterprise Linux 7.6 (RHEL) was running on all three servers as the host Operating System (OS). Kernel-based Virtual Machine (KVM) was installed as the hypervisor. Based on F5 Networks' recommended settings to optimize for NFV performance and security, we disabled hyper-threading and Turbo Boost on the compute node. The compute node and the traffic generator were connected through an HPE 5900 switch series JC772A.

Red Hat OpenStack Platform 13 (RHOSP) was deployed on all the three servers. Figure 2 shows the Red Hat OpenStack architecture. The undercloud is the main director node, which manages and provisions 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 compute resources for the OpenStack environment.

The Function Under Test (FUT) was a BIG-IP LTM High-Performance virtualized edition (VE) solution provided by F5 Networks. The F5 BIG-IP LTM was instantiated based on the configurations that are listed in Table 2.

Configuration Parameters	Value
Assigned CPU cores	8
Assigned Memory value	16 Gigabytes
Assigned Physical Network Interface	2x40GbE with SR-IOV enabled
Assigned Virtual Function Network Interface	2 Virtual Function in total 2 Physical Function in total (details in figure 3)
Assigned storage space	100 Gigabyte
Load balancing algorithm	Round Robin
Backend Server number	1,000 IPv4 1,000 IPv6

Configuration Parameters	Value
TCP persistence	No
Source Network Address Translation	No
Destination Network Address Translation	Yes

Table 2: F5 BIG-IP LTM Hardware and Software Configuration

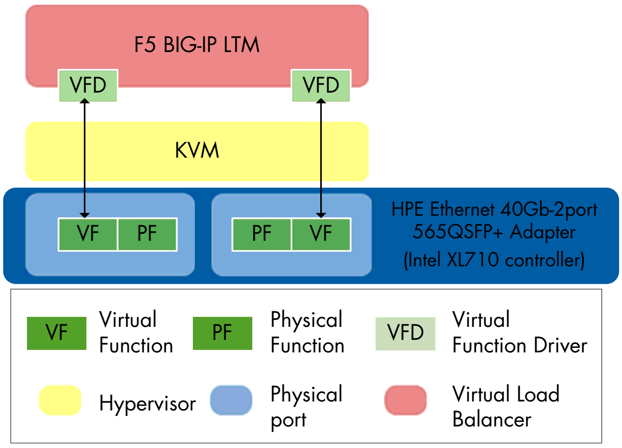

Figure 3: SR-IOV setup

Test Equipment

Tests were conducted using Spirent Communications C100-S3 high-performance appliance hardware with Avalanche software. The C100-S3 also supports Spirent's CyberFlood assessment solution for advanced mixed traffic, attack, malware and NetSecOpen methodologies. C100-S3 was used in the test execution to generate the traffic and collect performance statistics. The traffic profile included HTTP and HTTPS with a range of object sizes and application behavior settings described in detail below.

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. The configuration of F5 LTM remained identical during the whole test session. LTM received HTTPS traffic in the front end virtual server. The LTM broke the sessions and decrypted the content. Afterwards, it established new HTTP sessions and forwarded the content to the back end server with the round robin load balancing algorithm.

The TLS version is 1.2. We selected the cipher suite ECDHE-RSA-AES128-GCM-SHA256. The RSA key is RSA 2048 bits. It is applied to all the HTTPS test cases.

Each test run consisted of three 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 the main source of the results. The traffic then ramped down slowly to finalize the test and bring the FUT back into idle mode. The general ramp up, steady and ramp down time are 180 seconds, 600 seconds and 180 seconds respectively except for the Concurrent Connections (CC) test case.

HTTPS Connections Per Second (CPS)

KPI	Intel Xeon CPU	Min.	Max.	Average
TCP connections per second	Gold 6152	5,985	6,026	6,013
	Gold 6252N	6,482	6,524	6,512
	Improvement			+8.29%
Throughput (Mbit/s)	Gold 6152	196.72	204.6	200.8
	Gold 6252N	214.77	220.13	217.48
	Improvement			+8.3%
Time to TCP First Byte (msec)	Gold 6152	1.16	101.47	9.2
	Gold 6252N	1.06	109.59	7.93
	Improvement			–13.8%

Table 3: CPS Test Results with HTTPS 1 KByte Object Size

In this test case, we measured the maximum connection setup rate of the F5 BIG-IP LTM. Each HTTPS 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 3 represents the results of CPS and other KPIs during the steady phase with 1 KByte object size; Table 4 represents the respective KPIs for 64 KByte object size. We recorded the values based on minimum, maximum and arithmetic average.

KPI	Intel Xeon CPU	Min.	Max.	Average
TCP connections per second	Gold 6152	5,142	5,178	5,165
	Gold 6252N	5,677	5,717	5,705
	Improvement			+10.45%
Throughput (Mbit/s)	Gold 6152	2,913.2	3,054.9	2,988.6
	Gold 6252N	3,227.1	3,380.7	3,300.9
	Improvement			+10.44%
Time to TCP First Byte (msec)	Gold 6152	1.14	90.4	6.25
	Gold 6252N	1.06	79.12	6.37
	Regression			+1.92%

Table 4: CPS Test Results with HTTPS 64 KByte Object Size

Based on the readings of Table 3, Intel Xeon Gold 6252N processor achieved CPS and throughput enhancement of around 8.3% compared to Intel Xeon Gold 6152 processor. Moreover, the Time to TCP First Byte was reduced by 1.27 msec which gives a faster data transmission after the TCP connection is established. CPS and throughput gain is slightly better for the larger packet sizes as shown in Table 4. But, the trade-off is slower Time to TCP First Byte by 0.12 msec.

HTTPS Throughput

The purpose of this test case is to measure the maximum HTTPS throughput performance of the F5 BIG-IP LTM. 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 5 below. Each TCP connection had 10 transactions; it closed with FIN immediately after these 10 transactions.

Number of Requests/Weight	Object Size (KByte)
1 of each	0.2, 6, 8, 9, 10, 25, 26, 35, 59, 347

Table 5: MIX Object Size Distribution

Table 6 shows the achieved throughput results and other KPIs during the steady phase with 1 KByte sized objects; Table 7 shows the respective throughput for mixed object sizes.

KPI	Intel Xeon CPU	Min.	Max.	Average
Throughput (Mbit/s)	Gold 6152	777.93	802.27	789.92
	Gold 6252N	860.83	897.23	877.35
	Improvement			+11.06%
Transactions Per Second	Gold 6152	53,150	54,993	54,088
	Gold 6252N	58,787	61,633	60,074
	Improvement			+11.06%
Concurrent Connections	Gold 6152	280	477	337
	Gold 6252N	314	665	385
	Trade-off			+14.24%
TCP connections per second	Gold 6152	5,382	5,413	5,406
	Gold 6252N	5,963	6,023	6,010
	Improvement			+11.17%

Table 6: Throughput Test Results with HTTPS 1 KByte Object Size

KPI	Intel Xeon CPU	Min.	Max.	Average
Throughput (Mbit/s)	Gold 6152	12,791.3	13,153.2	12,988.9
	Gold 6252N	14,158.3	14,621.5	14,359.5
	Improvement			+10.55%
Transactions Per Second	Gold 6152	27,987	28,514	28,269
	Gold 6252N	30,922	31,520	31,256
	Improvement			+10.56%
Concurrent Connections	Gold 6152	217	593	458
	Gold 6252N	214	716	491
	Trade-off			+7.2%
TCP connections per second	Gold 6152	2,820	2,834	2,826
	Gold 6252N	3,099	3,133	3,125
	Improvement			+10.58%

Table 7: Throughput Test Results with HTTPS MIX Object Size

HTTPS Transaction Latency

In this test case, we measured the latency when the F5 BIG-IP LTM ran with 50% of the maximum workload, because 50% of the load is the normal workload when the FUT is in the production network. We split into two scenarios: 50% of the maximum CPS and 50% of maximum throughput with different object sizes.

In the CPS test, each TCP connection has only one transaction and closes the session immediately after the single transaction. In the throughput test, each TCP connection has 10 transactions and closes immediately after the 10 transactions.

Table 8 and Table 9 show latency and other KPIs when the FUT has only 50% of the maximum CPS during the steady phase. Table 10 and Table 11 show latency and other KPIs when the FUT has only 50% of the maximum throughput during the steady phase.

KPI	Intel Xeon CPU	Min.	Max.	Average
Throughput (Mbit/s)	Gold 6152	99.79	100.88	100.47
	Gold 6252N	108.37	109.25	108.89
	Improvement			+8.38%
Transactions Per Second	Gold 6152	2,984	3,026	3,006
	Gold 6252N	3,242	3,271	3,257
	Improvement			+8.34%
TCP connections per second	Gold 6152	2,993	3,013	3,006
	Gold 6252N	3,241	3,262	3,257
	Improvement			+8.34%
Time to TCP First Byte (msec)	Gold 6152	1.12	58.01	2.46
	Gold 6252N	1.03	59.36	2.31
	Improvement			–6.09%
Application transaction latency (msec)	Gold 6152	<1[a]	62	2.25
	Gold 6252N	<1	58	2.17
	Improvement			–3.55%

a. Due to the tester resolution limitation, values lower than 1 msec are shown as <1

Table 8: Transaction Latency Test Results at 50% CPS with HTTPS 1 KByte Object Size

KPI	Intel Xeon CPU	Min.	Max.	Average
Throughput (Mbit/s)	Gold 6152	1,489.5	1,502.5	1,496.4
	Gold 6252N	1,642.7	1,659.5	1,653.7
	Improvement			+10.51%

KPI	Intel Xeon CPU	Min.	Max.	Average
Transactions Per Second	Gold 6152	2,573	2,598	2,586
	Gold 6252N	2,840	2,870	2,858
	Improvement			+10.51%
TCP connections per second	Gold 6152	2,576	2,591	2,586
	Gold 6252N	2,837	2,863	2,858
	Improvement			+10.51%
Time to TCP First Byte (msec)	Gold 6152	1.12	62.09	2.31
	Gold 6252N	1.03	64.72	2.19
	Improvement			–5.19%
Application transaction latency (msec)	Gold 6152	1	73	3.45
	Gold 6252N	1	69	3.4
	Improvement			–1.44%

Table 9: Transaction Latency Test Results at 50% CPS with HTTPS 64 KByte Object Size

KPI	Intel Xeon CPU	Min.	Max.	Average
Throughput (Mbit/s)	Gold 6152	393.34	397.17	395.39
	Gold 6252N	436.72	440.73	439.07
	Improvement			+11.04%
Transactions Per Second	Gold 6152	26,932	27,190	27,071
	Gold 6252N	29,903	30,185	30,066
	Improvement			+11.06%

KPI	Intel Xeon CPU	Min.	Max.	Average
TCP connections per second	Gold 6152	2,694	2,712	2,706
	Gold 6252N	2,992	3,015	3,005
	Improvement			+11.04%
Time to TCP First Byte (msec)	Gold 6152	1.12	61.56	2.27
	Gold 6252N	1.02	56.44	2.21
	Improvement			–2.64%
Application transaction latency (msec)	Gold 6152	<1	83	1.72
	Gold 6252N	<1	66	1.7
	Improvement			–1.16%

Table 10: Transaction Latency Test Results at 50% throughput with HTTPS 1 KByte Object Size

KPI	Intel Xeon CPU	Min.	Max.	Average
Through put (Mbit/s)	Gold 6152	6,467.5	6,527	6,498.9
	Gold 6252N	7,148.9	7,230.1	7,189.6
	Improvement			+10.62%
Transactions Per Second	Gold 6152	14,096	14,179	14,145
	Gold 6252N	15,572	15,718	15,647
	Improvement			+10.61%
TCP connections per second	Gold 6152	1,408	1,419	1,414
	Gold 6252N	1,556	1,570	1,564
	Improvement			+10.6%

KPI	Intel Xeon CPU	Min.	Max.	Average
Time to TCP First Byte (msec)	Gold 6152	1.13	55.39	2.03
	Gold 6252N	1.04	53.42	1.95
	Improvement			–3.94%
Application transaction latency (msec)	Gold 6152	<1	145	2.129
	Gold 6252N	<1	107	2.126
	Improvement			–0.14%

Table 11: Transaction Latency Test Results at 50% throughput with HTTPS MIX Object Size

HTTPS Concurrent Connections Capacity

In this test case, we verified the maximum number of Concurrent Connections that was supported by the F5 BIG-IP LTM. The maximum sustained CC gives a better understanding of capacity limits of the FUT 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 3,000 connections/second with Intel Xeon Gold 6152 and 3,250 connections/second with Intel Xeon Gold 6252N. The F5 BIG-IP LTM showed 1 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 12 shows the maximum concurrent connections capacity and other KPIs during the steady phase.

KPI	Intel Xeon CPU	Min.	Max.	Average
Throughput (Mbit/s)	Gold 6152	59.14	98.28	78.93
	Gold 6252N	63.85	105.62	85.34
	Improvement			+8.12%
Concurrent Connections	Gold 6152	1,000,000	1,000,000	1,000,000
	Gold 6252N	1,000,000	1,000,000	1,000,000
	No difference			0%
Application transaction latency (msec)	Gold 6152	<1	54	1.29
	Gold 6252N	<1	57	1.28
	Improvement			–0.77%
Transactions Per Second	Gold 6152	4,521	7,468	6,012
	Gold 6252N	4,886	8,011	6,500
	Improvement			+8.11%

Table 12: Concurrent Connections Test Results with HTTPS 1 KByte Object Size

Observations

During the test, we observed random TCP reset behavior. EANTC investigated and found out that these are two different TCP reset behaviors.

- F5 is still sending packets to a closed session which causes TCP reset. This behavior happened randomly but lasted for the whole test session.
- F5 is not forwarding traffic in random session. Once the idle time out is triggered, the session ends up with a TCP reset. However, this behavior has only happened when the system workload is high (more than 90%). When the system load is moderate (around 50%), there is no such behavior.

Conclusion

We verified the performance enhancement of Intel Xeon Gold 6252N @2.3 GHz over Intel Xeon Gold 6152 @2.1 GHz is between 8-11% throughput and connections per second (refer to Table 3 and 4).

Moreover, latency was reduced between 5-6% for (Time to TCP First Byte) when the system is loaded with 50% of the max CPS (refer to Table 8 and 9).

Refer to Table 13 for a summary.

Object Size	Max CPS	Max Throughput (Mbit/s)	Time to TCP First Byte
1 KByte	+8.29%	+8.3%	–6.09%
64 KByte	+10.45%	+10.44%	–5.19%

Table 13: Summary

Generally, this gain is correlated to the processor clock frequency improvement, which technically justifies upgrading the processors of NFVI to Intel Xeon Gold 6252N to achieve a higher performance independent of VNF.

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.

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.

Tests were conducted by European Advanced Networking Test Center (EANTC). Hardware configurations (initial tests): three servers with dual Intel Xeon Gold 6152 processors (microcode: 0x200005e) running @ 2.1 GHz with 22 cores, 384 Gigabytes 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 1x Spirent Avalanche® C100-S3 appliances using 4xDual-port 10Gbps adapters.

In the second set of tests, Intel Xeon Gold 6152 processors were substituted with Intel Xeon Gold 6252N processors (microcode: 0x5000021) @ 2.3 GHz with 24 cores QS samples. The remaining hardware and software configurations remained the same.

Performance results are based on testing as of August 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/
[v1.0 20190920]