


HPE ProLiant DL380 Gen11 and Intel® 4th Gen Xeon® processors achieve ~1.4 Tbps with DPDK RFC 2544

Up to ~1.4 Tbps throughput on a HPE ProLiant DL380 powered by Intel® 4th Gen Xeon Scalable Gold 6454S Processors with a DPDK packet forwarding RFC 2544 zero packet loss test



**Hewlett Packard
Enterprise**

Authors

HPE:

Edith Chang

Lee Roberts

Shemal Shah

Jason Fleischli

Intel:

Sarita Maini

Georgii Tkachuk

Brad Chaddick

Introduction

This document describes the outstanding performance of the HPE ProLiant DL380 Gen11 server with Intel® 4th Gen Xeon Scalable Processors (formerly code-named Sapphire Rapids) and Intel® Ethernet Network Adapter E810-2CQDA2 NICs, which are included in the Intel 800 Series Network Adapters (formerly code-named Columbiaville) running DPDK RFC 2544.

The HPE ProLiant DL380 Gen11 server is a scalable 2U dual socket rack-mountable server that delivers exceptional compute performance expandability and scalability for diverse workloads and environments. Powered by 4th Gen Intel® Xeon® Scalable Processors which are available with up to 60 physical cores, increased memory bandwidth, and high-speed PCIe Gen5 I/O, the HPE ProLiant DL380 Gen11 server is an ideal dual-socket, 2U/2P, scalable solution. It is designed for data-intensive, large-capacity workloads.

This HPE server has the expanded PCIe capability to hold up to 14 E810-CAM1 controllers mounted on six E810-2CQDA2 network adapters and two E810-CQDA1 for OCP 3.0 network adapters. This makes it possible to achieve ~1.4 Tbps of lossless network throughput in a single server. Utilizing the Xena 400GbE per port hardware traffic generator RFC 2544 with zero loss tolerance, this test produced up to 1.399 Tbps, with 99.9% line rate convergence @ 256 Bytes or greater payload packet sizes and achieved 1.5 billion low latency packets per second utilizing DPDK.

Each of the two Intel® Xeon® Gold 6454S processors used in this performance test have 32 cores, 64 logical cores when hyper-threading is turned on, with a processor base frequency of 2.20 GHz and a max turbo frequency of 3.40 GHz, 60 MB cache, Intel® UPI speed of 16 GT/s, maximum number of UPI links is 4 and the TDP is 270W.

The 100GbE Intel® Ethernet Network Adapter 810 supports speeds up to 100GbE per port. This HPE ProLiant DL 380 server has a top-notch best-in-class capacity of 14 x 100 GbE ports in a single dual-socket server, which enabled such a high I/O throughput. Each socket was populated with 7 x 100 GbE ports.

Table of Contents

Introduction.....	1
RFC 2544 Test.....	2
Test Cases.....	2
Test Details.....	4
Benchmarks.....	5
Summary	5

RFC 2544 Test

The Testpmd application for RFC 2544 is a standard DPDK test, included in every release. The XENA traffic generator sends a specific number of frames at a specific rate through the HPE server which is the Device Under Test (DUT) in this case and then counts the frames that are transmitted by the DUT. If the count of offered frames is not equal to the count of received frames, the rate of the traffic stream is automatically reduced, and the test is rerun as needed until zero packet loss is achieved. The throughput is the fastest rate at which the count of test frames transmitted by the DUT is equal to the number of test frames sent to it by the Xena hardware traffic generator.

Intel® Xeon® Gold processors are optimized for demanding network workloads. With support for higher memory speeds

and enhanced memory capacity, these processors deliver performance gains, better memory efficiency and capability and workload acceleration, reduce total cost of ownership and power consumption with increased performance per watt. The 4th Gen Intel® Xeon® Scalable Processors have the most built-in accelerators of any CPU in the market to improve performance in AI, networking, data analytics, storage and HPC.

The test setup consisted of the following layers:

Hardware: HPE ProLiant DL380 Gen11 server with Intel® Xeon® Gold 6454S processors

Operating System: Red Hat Enterprise Linux release 9.2

Application: DPDK Testpmd

Test Cases

The following test scenarios were executed with the RHEL 9.2 OS, kernel 5.14.0-284.11.1.el9_2.x86_64 using DPDK 21.11.5 and the dpdk-testpmd application which is included in each DPDK release. Both sockets of the HPE ProLiant DL380 Gen11 server were populated with Intel® Xeon® Gold 6454S processors. The best peak performance for all packet sizes was seen with the

2c2q test scenario. Two 32-core processors, with 14 out of a total of 64 cores were used for test cases 1 and 2, and 28 out of 64 cores were used for test case 3. The highest line rate throughput was achieved with packet sizes of 256B, 512B, 1024B, 1280B and 1518B.

Test Case #1 - 1c1q

Test #1 RFC 2544 zero packet loss test on 14 x 100 GbE ports. The ports were in 3 (2x100 GbE) PCIe NICs and 1 (1x100GbE) OCP NIC attached to each of the two NUMA nodes. Each NIC

port was assigned 1 core and 1 Tx/Rx queue pair resulting in a total use of 14 cores out of 64 cores.

Test Case #2 - 1c2q

Test #2 RFC 2544 zero packet loss test on 14 x 100 GbE ports. The ports were in 3 (2x100 GbE) PCIe NICs and 1 (1x100GbE) OCP NIC attached to each of the two NUMA nodes. Each NIC

port was assigned 1 core and 2 Tx/Rx queue pairs resulting in a total use of 14 cores out of 64 cores.

Test Case #3 - 2c2q

Test #3 RFC 2544 zero packet loss test on 14 x 100 GbE ports. The ports were in 3 (2x100 GbE) PCIe NICs and 1 (1x100GbE) OCP NIC attached to each of the two NUMA

nodes. Each NIC port was assigned 2 cores and 2 Tx/Rx queue pairs resulting in a total use of 28 cores out of 64 cores.

Test Details

Linux kernel arguments:

```
# cat /proc/cmdline
BOOT_IMAGE=(hd2,gpt2)/vmlinuz-5.14.0-284.11.1.el9_2.x86_64 root=/dev/mapper/rhel_en113-root ro crashkernel=auto resume=/dev/mapper/rhel_en113-swap rd.lvm.lv=rhel_en113/root rd.lvm.lv=rhel_en113/swap rhgb quiet default_hugepagesz=2M hugepagesz=2M hugepages=8192 hugepagesz=1G hugepages=8 intel_iommu=on iommu=pt isolcpus=1-15,17-63 nohz_full=1-15,17-63 rcu_nocbs=1-15,17-63 nmi_watchdog=0 audit=0 nosoftlockup hpet=disable mce=off tsc=reliable numa_balancing=disable memory_corruption_check=0 workqueue.power_efficient=false init on alloc=0 module_blacklist=ast,iavf
```

```
modprobe.blacklist=ice,qat_4xxxvf,qat_4xxx
```

PCIe devices:

```
# lshw -c net -businfo
Bus info          Device          Class           Description
-----
pci@0000:11:00.0  ens3            network         Ethernet Controller E810-C for QSFP
pci@0000:14:00.0  enp20s0         network         Ethernet Controller E810-C for QSFP
pci@0000:23:00.0  ens2            network         Ethernet Controller E810-C for QSFP
pci@0000:26:00.0  enp38s0         network         Ethernet Controller E810-C for QSFP
pci@0000:35:00.0  ens14           network         Ethernet Controller E810-C for QSFP
pci@0000:59:00.0  ens1            network         Ethernet Controller E810-C for QSFP
pci@0000:5c:00.0  enp92s0         network         Ethernet Controller E810-C for QSFP
pci@0000:89:00.0  ens5            network         Ethernet Controller E810-C for QSFP
pci@0000:8c:00.0  enp140s0        network         Ethernet Controller E810-C for QSFP
pci@0000:9c:00.0  ens6            network         Ethernet Controller E810-C for QSFP
pci@0000:9f:00.0  enp159s0        network         Ethernet Controller E810-C for QSFP
pci@0000:af:00.0  ens15           network         Ethernet Controller E810-C for QSFP
pci@0000:c2:00.0  ens4            network         Ethernet Controller E810-C for QSFP
pci@0000:c5:00.0  enp197s0        network         Ethernet Controller E810-C for QSFP
```

DPDK testpmd commands:

1c1q:

```
$DPDK_BUILD/app/dpdk-testpmd -l 0,1-2 --main-lcore 0 -n 8 --force-max-simd-
bitwidth=512 -a 0000:59:00.0 -a 0000:5c:00.0 --socket-mem=512 --file-prefix max-ens1 --
--socket-num=0 --nb-cores=2 --rxq=1 --txq=1 -i --burst=32 --rxd=1024 --txd=1024 --
forward-mode=io --portmask=0x3 -a
```

1c2q:

```
$DPDK_BUILD/app/dpdk-testpmd -l 0,1-2 --main-lcore 0 -n 8 --force-max-simd-
bitwidth=512 -a 0000:59:00.0 -a 0000:5c:00.0 --socket-mem=512 --file-prefix max-ens1 --
--socket-num=0 --nb-cores=2 --rxq=2 --txq=2 -i --burst=32 --rxd=1024 --txd=1024 --
forward-mode=io --portmask=0x3 -a
```

2c2q:

```
$DPDK_BUILD/app/dpdk-testpmd -l 0,1-2,33-34 --main-lcore 0 -n 8 --force-max-simd-
bitwidth=512 -a 0000:59:00.0 -a 0000:5c:00.0 --socket-mem=512 --file-prefix max-ens1 --
--socket-num=0 --nb-cores=4 --rxq=2 --txq=2 -i --burst=32 --rxd=1024 --txd=1024 --
forward-mode=io --portmask=0x3 -a
```

CPU:

```
# lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Address sizes:          52 bits physical, 57 bits virtual
Byte Order:             Little Endian
CPU(s):                 64
On-line CPU(s) list:   0-63
Vendor ID:              GenuineIntel
BIOS Vendor ID:         Intel(R) Corporation
Model name:             Intel(R) Xeon(R) Gold 6454S
BIOS Model name:        Intel(R) Xeon(R) Gold 6454S
CPU family:             6
Model:                  143
Thread(s) per core:    1
Core(s) per socket:    32
Socket(s):              2
```

```

Stepping:           8
CPU max MHz:       3400.0000
CPU min MHz:       800.0000
Virtualization features:
Virtualization:    VT-x
Caches (sum of all):
L1d:               3 MiB (64 instances)
L1i:               2 MiB (64 instances)
L2:                128 MiB (64 instances)
L3:                120 MiB (2 instances)
NUMA:
NUMA node(s):      2
NUMA node0 CPU(s): 0-15,32-47
NUMA node1 CPU(s): 16-31,48-63
Vulnerabilities:
Itlb multihit:     Not affected
Lltd:              Not affected
Mds:               Not affected
Meltdown:          Not affected
Mmio stale data:   Not affected
Retbleed:          Not affected
Spec store bypass: Mitigation; Speculative Store Bypass disabled via prctl
Spectre v1:        Mitigation; usercopy/swaps barriers and __user pointer
sanitization
Spectre v2:        Mitigation; Enhanced IBRS, IBPB conditional, RSB filling,
PBRBS-eIBRS SW sequence
    
```

Test Setup

Figure 1 below shows the hardware setup used for the RFC 2544 test:




Figure 1. Test Setup

RFC 2544 Benchmarks

The results of the RFC 2544 DPK Packet Forwarding benchmarks on a HPE ProLiant DL380 Gen11 server with Intel Xeon Gold 6454S processors, showed a network throughput of up to 1399.59 Gbps (~ 1.4 Tbps) with 0% packet loss. The RFC recommended packet sizes (frame sizes) of 64 bytes, 128 bytes, 256 bytes, 512 bytes, 1024 bytes, 1280 bytes and 1518 bytes were used in the testing. A line rate of 99.99% was measured with packet sizes of 256 bytes and larger, with the 2c2q test case providing the peak performance.

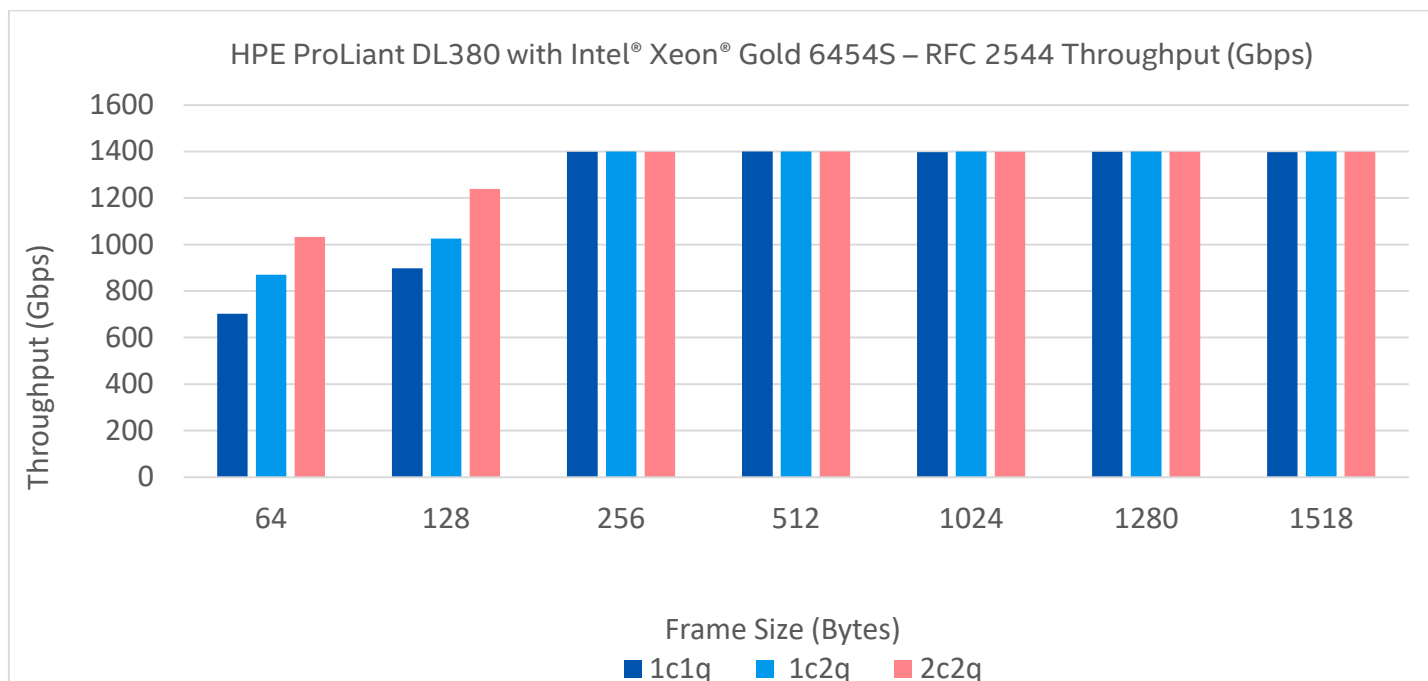


Figure 2. RFC 2544 DPK L3 Forwarding Benchmarks^{1,2}

Summary

HPE ProLiant DL380 Gen11 system with Intel® Xeon® Gold 6454S processors and Intel® Ethernet Network Adapters delivered an impressive 1399.59 (~1.4 Tbps) throughput with 36 physical cores left for control and user plane workloads or other computational tasks. The packet loss was 0%, which demonstrated the performance determinism required for network workloads.

The HPE ProLiant Gen11 servers offer twice the I/O bandwidth compared to the previous generation servers, increased memory bandwidth, high energy efficiency and optimized performance for all workloads. This high performing server combined with the 4th Gen Intel® Xeon® Scalable processors, provides performance leadership, as demonstrated by the exceptional benchmarks highlighted in this paper.



¹ Notices & Disclaimers

Performance varies by use, configuration, and other factors. Learn more on the Performance Index site. Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure. Your costs and results may vary. Intel technologies may require enabled hardware, software or service activation. © Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. *Other names and brands may be claimed as the property of others.

² Configuration

Test by HPE as of 10/25/23: System – HPE ProLiant DL380 Gen11, 2x Intel® Xeon® Gold 6454S, 32 cores, HT Off, Turbo On, Total Memory 512GB (16x32GB DDR5 4800 MT/s [4800 MT/s]), BIOS 1.4, microcode 0x2b0004b1, 14x Ethernet Controller E810-C for QSFP, 2x 447.1G MK000480GZXRA, 2x 745.2G KCM6XVUL800G, 1x 447.1G HPE NS204i-u Gen11 Boot Controller, Red Hat Enterprise Linux 9.2 (Plow), 5.14.0-284.11.1.el9_2.x86_64, ice driver 1.11.17.1, DPK: dpdk-stable-21.11.5, test app – testpmd,