

Intel Ethernet' s Performance Report with DPDK 22.07

Test Date: July 25th, 2022

Author: Intel DPDK Validation team



Revision History

Date	Revision	Comment
July 25th, 2022	1.0	Initial document for release

Contents

Audience and Purpose	4
Test Description	4
Test#1 RFC2544 zero packet loss test on 1x Intel® Ethernet Converged Network Adapter X710-DA4	7
Test Results	8
Test#2 RFC2544 zero packet loss test on 2x Intel® Ethernet Network Adapter XXV710-DA2	9
Test Results	10
Test#3 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter XXV710-DA2	11
Test Results	12
Test#4 i40e PMD Single core performance test	13
Test Results	15
Test#5 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-CQDA2	16
Test Results	17
Test#6 ice PMD Single core performance test	18
Test Results	20
<i>Subcase #1: Turbo boost comparison</i>	20
<i>Subcase #2: Instruction comparison</i>	21
<i>Subcase #3: Hyper-threading comparison</i>	22
Test#7 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-CQDA2 using SR-IOV over KVM	23
Test Results	26
Test#8 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-XXVDA4	27
Test Results	28
Test#9 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-2CQDA2	29
Test Results	30
Test#10 Zero packet loss test on Intel server of m6id.metal AWS cloud instance	31
Test Results	33

Audience and Purpose

The primary audience for this test report is architects and engineers implementing the Data Plane Development Kit (DPDK). This report provides information on packet processing performance testing for the specified DPDK release on Intel® architecture. The initial report may be viewed as the baseline for future releases and provides system configuration and test cases based on DPDK examples.

The purpose of reporting these tests is not to imply a single “correct” approach, but rather to provide a baseline of well-tested configurations and procedures with reproducible results. This will help guide architects and engineers who are evaluating and implementing DPDK solutions on Intel® architecture and can assist in achieving optimal system performance.

Test Description

The device under test (DUT) consists of a system with an Intel® architecture motherboard populated with the following;

- A single or dual processor and PCH chip, except for System on Chip (SoC) cases
- DRAM memory size and frequency (normally single DIMM per channel)
- Specific Intel Network Interface Cards (NICs)
- BIOS settings noting those that updated from the basic settings
- DPDK build configuration settings, and commands used for tests

Connected to the DUT is an IXIA, a hardware test and simulation platform to generate packet traffic to the DUT ports and determine the throughput at the tester side. The IXIA is used to implement RFC2544 on the DUT.

Benchmarking a DPDK system requires knowledge of networking technologies including knowledge of network protocols and hands-on experience with relevant open-source software, such as Linux, and the DPDK. Engineers also need benchmarking and debugging skills, as well as a good understanding of the device-under-test (DUT) across compute and networking domains.

DPDK L3fwd Test Case: Documentation may be found at http://www.dpdk.org/doc/guides/sample_app_ug/l3_forward.html.

The application demonstrates the use of the hash and LPM libraries in the DPDK to implement packet forwarding. The lookup method is either hash-based or LPM-based and is selected at compile time. In this test case, the selected lookup method is the default LPM-based without modification, an LPM object is used to emulate the forwarding stage for IPv4 packets. The LPM object is used as the routing table to identify the next hop for each input packet at runtime.

The LPM lookup key is represented by the Destination IP Address field read from the input packet. The ID of the output interface for the input packet is the next hop returned by the LPM lookup. The set of LPM rules used by the application is statically configured and loaded into the LPM object at initialization time.

In the sample application, hash-based forwarding supports IPv4 and IPv6. LPM-based forwarding supports IPv4 only.

DPDK Testpmd Test Case: Documentation may be found at http://www.dpdk.org/doc/guides/testpmd_app Ug/index.html.

The testpmd application can be used to test the DPDK in a packet forwarding mode and also to access NIC hardware features. Note in the Testpmd example if the `-i` argument is used, the first core is used for the command language interface (CLI).

Below is an example setup topology for the performance test. Generally, NIC ports in DUT connect to the traffic generator ports directly. And we will send bi-directional traffic from traffic generators and calculate the aggregated throughputs at the traffic generator side for the overall throughput. Please note the NIC port connections to traffic generator would vary from case to case.

Figure 1: DPDK NIC performance test setup (1 port per NIC)

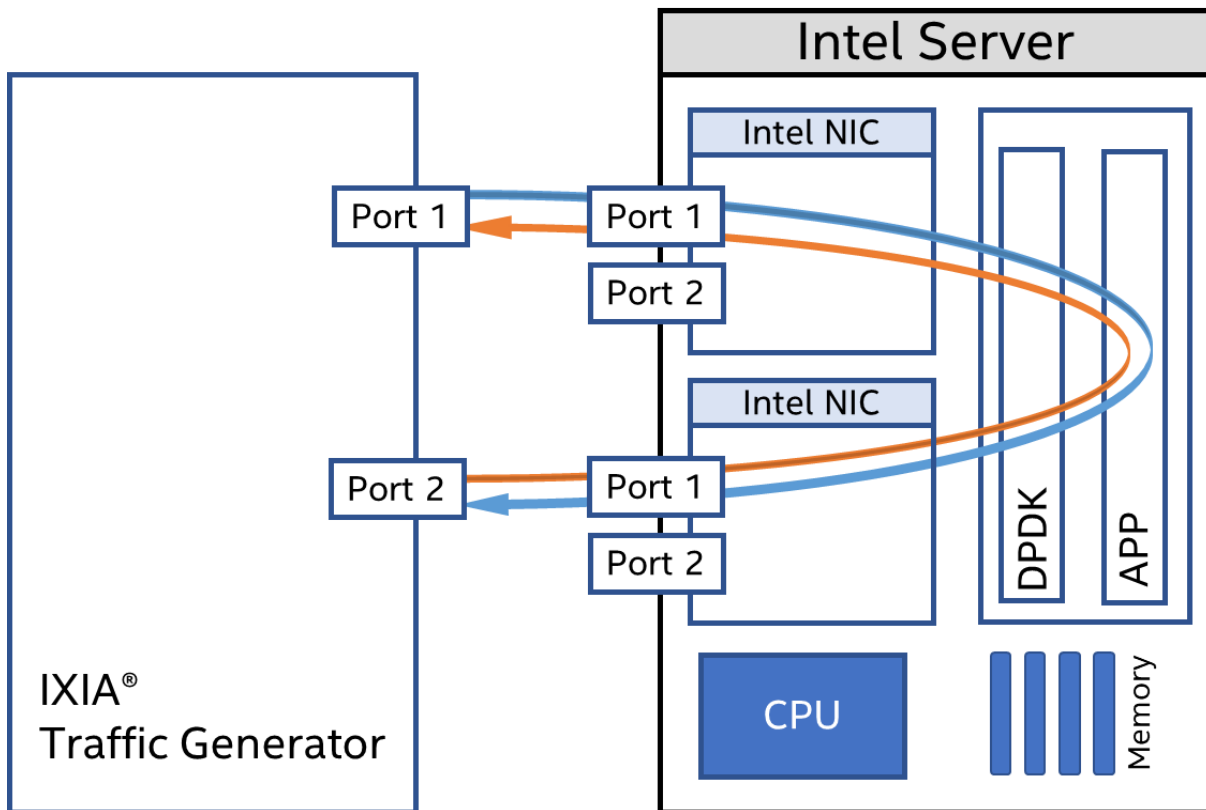
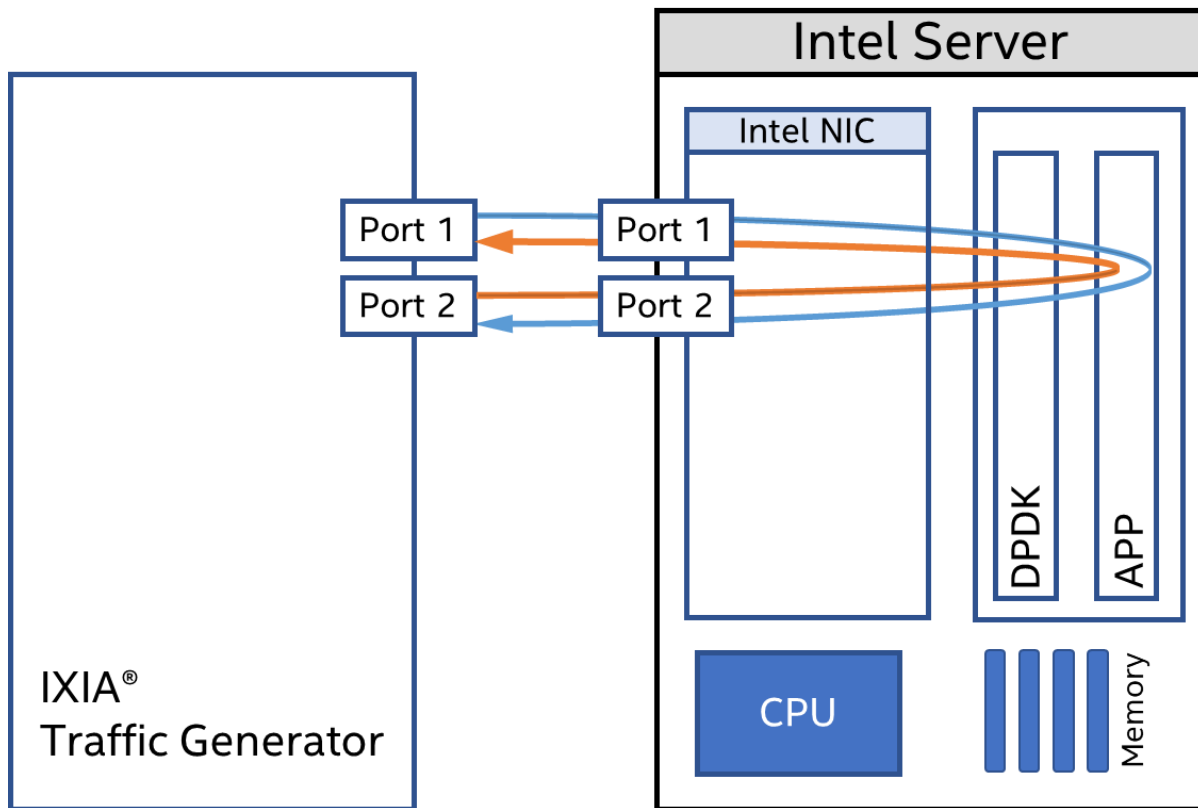


Figure 2: DPDK NIC performance test setup (2 ports on 1NIC)



RFC2544 Zero packet loss test case: Used to determine the DUT throughput as defined in RFC1242(<https://www.ietf.org/rfc/rfc1242.txt>). Note RFC6201 <https://www.ietf.org/rfc/rfc6201.txt> has updated RFC2544 and RFC1242. Please check the link for more details. In this report, RFC2544 test uses DPDK l3fwd as test application.

Procedure: Send a specific number of frames at a specific rate through the DUT and then count 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 offered stream is reduced and the test is rerun. 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 test equipment.

Test settings: IxNetwork 9.0 is used here for the test execution. The duration for each round is 60 seconds, and the acceptable packet loss rate is 0. The traffic is 256 flows per port.

Single core test case: Used to check the maximum IO throughput for a single core. The case requires at least 1x 100GbE ports for the ice DPDK Poll Mode Driver (PMD) and 3x 25GbE ports for the i40e DPDK PMD. In this case, we don't calculate the packet loss rate but record the average throughput within 60 seconds.

More details about the DPDK setup and BIOS information can be found at http://dpdk.org/doc/guides/linux_gsg/nic_perf_intel_platform.html

Test#1 RFC2544 zero packet loss test on 1x Intel® Ethernet Converged Network Adapter X710-DA4

Table 1: Test #1 Setup

Item	Description
Test Case	Test#1 RFC2544 zero packet loss test on 1x Intel® Ethernet Converged Network Adapter X710-DA4
Server Platform	Intel® Server Board S2600WFT
CPU	Intel® Xeon® Platinum 8180 CPU @ 2.50GHz 28 CPU cores * 2 NUMA nodes
Memory	48GB: 8GB x 6 DIMMs x 2 NUMA nodes @ 2666MHz
NIC	1x Intel® Ethernet Converged Network Adapter X710-DA4
BIOS	SE5C620.86B.00.01.0013.030920180427
Microcode	0x2006b06
Operating System	Ubuntu 20.04.4 LTS
Linux kernel version	5.13.0-44-generic
GCC version	gcc version 9.4.0 (Ubuntu 9.4.0-1ubuntu1~20.04.1)
Intel NIC Device ID	0x1572
Intel NIC Firmware version	8.70
DPDK version	22.07
Test configuration	1 NIC card attached to the first processor. 1 port assigned per logical core using 1 queue. Totaling 4 logical cores, 4 queues for 4ports

Figure 3: Test #1 Setup – 1x Intel® Ethernet Converged Network Adapter X710-DA4 connected to IXIA (4 ports on 1 NIC)




Table 2: Test #1 Setting


Item	Description
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Disabled Turbo Boost Disabled
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=1-21,28-48 nohz_full=1-21,28-48 rcu_nocbs=1-21,28-48
DPDK Settings	CC=gcc meson -Dlibdir=lib -Dexamples=l3fwd -Dc_args=- DRTE_LIBRTE_I40E_16BYTE_RX_DESC --default-library=static x86_64-native- linuxapp-gcc
Command line	dpdk-l3fwd -l 5-8 -n 6 -a 0000:18:00.0 -a 0000:18:00.1 -a 0000:18:00.2 -a 0000:18:00.3 -- -p 0xf --config '(0,0,5),(1,0,6),(2,0,7),(3,0,8)' --tx-queue-size 2048 -- rx-queue-size 2048

Test Results

Table 3: Test #1 Result

Frame Size (Bytes)	Line Rate[4x10G] (Mpps)	Frame Rate (Mpps)	% Line Rate
64	59.52	36.51	61.33
128	33.78	33.78	100
256	18.12	18.12	100
512	9.40	9.40	100
1024	4.79	4.79	100
1280	3.85	3.85	100
1518	3.25	3.25	100

Figure 4: Test #1 Result - RFC2544 zero packet loss test on 1x Intel® Ethernet Converged Network Adapter X710-DA4



Test#2 RFC2544 zero packet loss test on 2x Intel® Ethernet Network Adapter XXV710-DA2

Table 4: Test #2 Setup

Item	Description
Test Case	Test#2 RFC2544 zero packet loss test on 2x Intel® Ethernet Network Adapter XXV710-DA2
Server Platform	Inspur® Server Board NF5280M6
CPU	Intel® Xeon® Platinum 8380 CPU @ 2.30GHz 40 CPU cores * 2 NUMA nodes
Memory	128GB: 32GB x 2 DIMMs x 2 NUMA nodes @ 3200MHz
NIC	2x Intel® Ethernet Network Adapter XXV710-DA2; PCIe 3.0 x8
BIOS	05.01.01
Microcode	0xd0002a0
Operating System	Ubuntu 20.04.4 LTS
Linux kernel version	5.13.0-41-generic
GCC version	gcc version 9.4.0 (Ubuntu 9.4.0-1ubuntu1~20.04.1)
Intel NIC Device ID	0x158b
Intel NIC Firmware version	8.70
DPDK version	22.07
Test configuration	2 NIC cards attached to the first processor only 1st port used of each NIC. Each port has 2 queues assigned for a total of 4 queues 1 queue assigned per logical core. Totaling 4 logical cores, 4 queues for 2 ports

Figure 5: Test #2 Setup – 2x Intel® Ethernet Network Adapter XXV710-DA2 connected to IXIA (1 port per NIC)

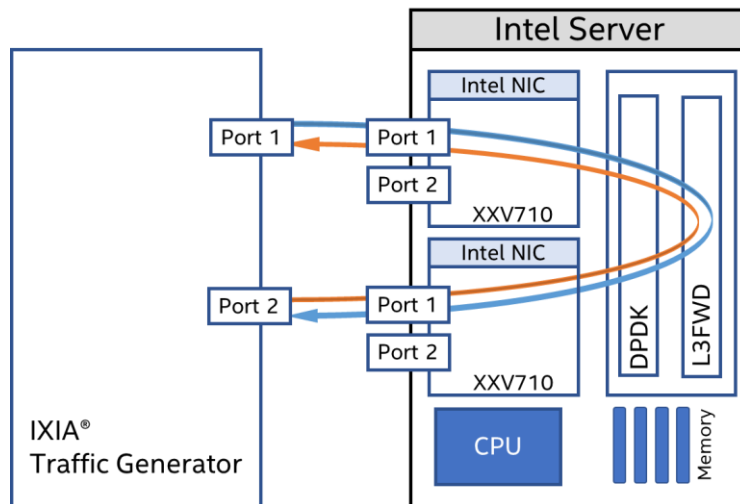


Table 5: Test #2 Setting


Item	Description
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Disabled Intel® Hyper-Threading Tech Enabled Turbo Boost Disabled
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=1-35,81-115 nohz_full=1-35,81-115 rcu_nocbs=1-35,81-115
DPDK Settings	CC=gcc meson -Dlibdir=lib -Dexamples=l3fwd -Dc_args=- DRTE_LIBRTE_140E_16BYTE_RX_DESC --default-library=static x86_64-native- linuxapp-gcc
Command line	dpdk-l3fwd -l 6,7,86,87 -n 6 -a 18:00.0 -a 1a:00.0 -- -p 0x3 --config '(0,0,6),(0,1,7),(1,0,86),(1,1,87)' --tx-queue-size 2048 --rx-queue-size 2048

Test Results

Table 6: Test #2 Result

Frame Size (Bytes)	Line Rate[2x25G] (Mpps)	Frame Rate (Mpps)	% Line Rate
64	74.40	74.40	100
128	42.23	42.23	100
256	22.65	22.65	100
512	11.75	11.75	100
1024	5.99	5.99	100
1280	4.81	4.81	100
1518	4.06	4.06	100

Figure 6: RFC2544 zero packet Loss test for 2x Intel® Ethernet Network Adapter XXV710-DA2



Test#3 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter XXV710-DA2

Table 7: Test #3 setup

Item	Description
Test Case	Test#3 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter XXV710-DA2
Server Platform	Inspur® Server Board NF5280M6
CPU	Intel® Xeon® Platinum 8380 CPU @ 2.30GHz 40 CPU cores * 2 NUMA nodes
Memory	128GB: 32GB x 2 DIMMs x 2 NUMA nodes @ 3200MHz
NIC	1x Intel® Ethernet Network Adapter XXV710-DA2; PCIe 3.0 x8
BIOS	05.01.01
Microcode	0xd0002a0
Operating System	Ubuntu 20.04.4 LTS
Linux kernel version	5.13.0-41-generic
GCC version	gcc version 9.4.0 (Ubuntu 9.4.0-1ubuntu1~20.04.1)
Intel NIC Device ID	0x158b
Intel NIC Firmware version	8.70
DPDK version	22.07
Test configuration	1 NIC cards attached to the first processor and 2 ports of the NIC are used. Each port has 1 queue assigned for a total of 2 queues 1 queue assigned per logical core. Totaling 2 logical cores, 2 queues for 2 ports

Figure 7: Test #3 Setup – 1x Intel® Ethernet Network Adapter XXV710-DA2 connected to IXIA (2 ports used on 1 NIC)

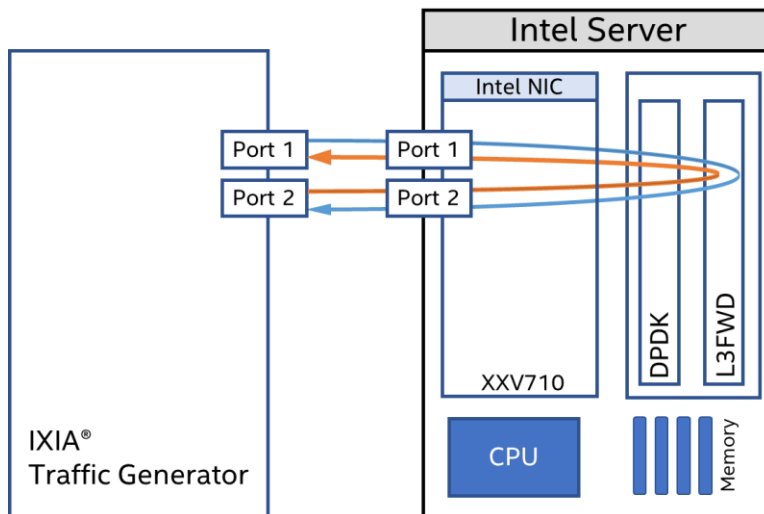


Table 8: Test #3 Setting

Item	Description
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Disabled Intel® Hyper-Threading Tech Enabled Turbo Boost Disabled
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=1-35,81-115 nohz_full=1-35,81-115 rcu_nocbs=1-35,81-115
DPDK Settings	CC=gcc meson -Dlibdir=lib -Dexamples=l3fwd -Dc_args=- DRTE_LIBRTE_140E_16BYTE_RX_DESC --default-library=static x86_64-native- linuxapp-gcc
Command line	dpdk-l3fwd -l 6,86 -n 6 -a 18:00.0 -a 18:00.1 -- -p 0x3 --config '(0,0,6),(1,0,86)' --tx- queue-size 2048 --rx-queue-size 2048


Test Results

Table 9: Test #3 Result

Frame Size (Bytes)	Line Rate[2x25G] (Mpps)	Frame Rate (Mpps)	% Line Rate
64	74.40	37.26	50.08
128	42.23	34.51	81.72
256	22.65	21.69	95.78
512	11.75	11.42	97.21
1024	5.99	5.86	97.88
1280	4.81	4.71	97.98
1518	4.06	3.98	96.88

Note: Due to hardware limitation, the throughput at large packets can only reach 95%-97% line rate.

Figure 8: RFC2544 zero packet Loss test for 1x Intel® Ethernet Network Adapter XXV710-DA2



Test#4 i40e PMD Single core performance test

Table 10: Test #4 setup

Item	Description
Test Case	Test#4 i40e PMD Single core performance test
Server Platform	Intel® Server Board NF5280M6
CPU	Intel® Xeon® Platinum 8380 CPU @ 2.30GHz 40 CPU cores * 2 NUMA nodes
Memory	128GB: 32GB x 2 DIMMs x 2 NUMA nodes @ 3200MHz
NIC	4x Intel® Ethernet Network Adapter XXV710-DA2; PCIe 3.0 x8
BIOS	05.01.01
Microcode	0xd0002a0
Operating System	Ubuntu 20.04.4 LTS
Linux kernel version	5.13.0-41-generic
GCC version	gcc version 9.4.0 (Ubuntu 9.4.0-1ubuntu1~20.04.1)
Intel NIC Device ID	0x158b
Intel NIC Firmware version	8.70
DPDK version	22.07
Test configuration	4 NIC cards attached to the same processor and 1 port used of each NIC. Each port has 1 queue assigned for a total of 4 queues. All queues are assigned to the same logical core. The first core defined is for the CLI of the Testpmd example due to the -i argument. I40e Vector PMD and IO forwarding is used for the test. Test duration is 60 seconds.

Figure 9: Test #4 Setup – 4x Intel® Ethernet Network Adapter XXV710-DA2 connected to IXIA (1 port per NIC)




Table 11: Test #4 Setting

Item	Description
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Enabled Intel® Hyper-Threading Tech Disabled Turbo Boost Enabled
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=1-35,81-115 nohz_full=1-35,81-115 rcu_nocbs=1-35,81-115
DPDK Settings	CC=gcc meson -Dlibdir=lib -Dexamples=l3fwd -Dc_args=- DRTE_LIBRTE_I40E_16BYTE_RX_DESC --default-library=static x86_64-native-linuxapp-gcc
Command line	dpdk-testpmd -l 33,34 -n 6 -a 0000:86:00.0 -a 0000:af:00.0 -a 0000:b1:00.0 -a 0000:da:00.0 -- -i --txd=2048 --rx=2048 --forward=io -a add "--force-max-simd-bitwidth=512" if using AVX512, default is AVX2


Test Results

Table 12: Test #4 Result

Frame Size (Bytes)	Line Rate [4*25G] (Mpps)	AVX2		AVX512		AVX512 vs. AVX2 %
		CPU Freq	3.0	CPU Freq	2.9	
		Frame Rate (Mpps)	Cycles per Packet=CPU_freq/Frame_rate	Frame Rate (Mpps)	Cycles per Packet=CPU_freq/Frame_rate	
64	148.81	74.67	40	81.45	36	9.07
128	84.46	74.51	40	81.54	36	9.43
256	45.29	45.29	-	45.29	-	-
512	23.5	23.50	-	23.50	-	-
1024	11.97	11.97	-	11.97	-	-
1280	9.62	9.62	-	9.62	-	-
1518	8.13	8.13	-	8.13	-	-

Note: All packet sizes over 256B reach line rate, so the single core is not fully stressed.

Figure 10: i40e PMD Single core performance on 4x Intel® Ethernet Network Adapter XXV710-DA2



Test#5 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-CQDA2

Table 13: Test #5 Setup

Item	Description
Test Case	Test#5 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-CQDA2
Server Platform	Supermicro® Server Board SYS-730A-I
CPU	Intel® Xeon® Platinum 8380 CPU @ 2.30GHz 40CPU cores * 2 NUMA nodes
Memory	256GB: 16GB x 8 DIMMs x 2 NUMA nodes @ 3200MHz
NIC	1x Intel® Ethernet Network Adapter E810-CQDA2; PCIe 3.0/4.0 x16
BIOS	1.1a
Microcode	0xd0002c1
Operating System	Ubuntu 20.04.4 LTS
Linux kernel version	5.13.0-41-generic
GCC version	gcc version 9.4.0 (Ubuntu 9.4.0-1ubuntu1~20.04.1)
Intel NIC Device ID	0x1592
Intel NIC Firmware version	4.0
Intel NIC Package	1.3.30.0
DPDK version	22.07
Test configuration	1 NIC card attached to the processor, and 1 port used. Each port has 4 queues assigned for a total of 4 queues 1 queue assigned per logical core. Totaling 4 logical cores, 4 queues for 1 port.

Figure 11: Test #5 Setup – 1x Intel® Ethernet Network Adapter E810-CQDA2 connected to IXIA (Single port)

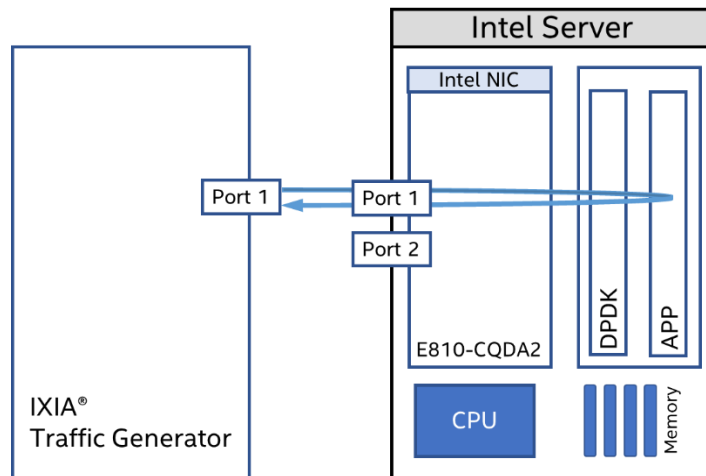


Table 14: Test #5 Setting


Item	Description
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Enabled Intel® Hyper-Threading Tech Disabled Turbo Boost Enabled
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=1-21 nohz_full=1-21 rcu_nocbs=1-21 idle=poll nmi_watchdog=0 audit=0 nosoftlockup processor.max_cstate=0 intel_idle.max_cstate=0 hpet=disable mce=off tsc=reliable numa_balancing=disable
DPDK Settings	CC=gcc meson -Dlibdir=lib -Dexamples=l3fwd -Dc_args=- DRTE_LIBRTE_ICE_16BYTE_RX_DESC --default-library=static x86_64-native- linuxapp-gcc
Command line	dpdk-l3fwd -l 5-8 -n 8 --force-max-simd-bitwidth=512 -- -p 0x1 --config '(0,0,5),(0,1,6),(0,2,7),(0,3,8)'

Test Results

Table 15: Test #5 Result

Frame Size (Bytes)	Line Rate[100G] (Mpps)	Frame Rate (Mpps)	% Line Rate
64	148.81	116.37	78.20
128	84.46	83.87	99.30
256	45.29	44.97	99.29
512	23.5	23.5	100
1024	11.97	11.97	100
1280	9.62	9.62	100
1518	8.13	8.13	100

Figure 12: RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-CQDA2



Test#6 ice PMD Single core performance test

Table 16: Test #6 Setup

Item	Description
Test Case	Test#6 ice PMD Single core performance test on 2x Intel® Ethernet Network Adapter E810-CQDA2
Server Platform	Supermicro® Server Board SYS-730A-I
CPU	Intel® Xeon® Platinum 8380 CPU @ 2.30GHz 40 CPU cores * 2 NUMA nodes
Memory	256GB: 16GB x 8 DIMMs x 2 NUMA nodes @ 3200MHz
NIC	2x Intel® Ethernet Network Adapter E810-CQDA2; PCIe 3.0/4.0 x16
BIOS	1.1a
Microcode	0xd0002c1
Operating System	Ubuntu 20.04.4 LTS
Linux kernel version	5.13.0-41-generic
GCC version	gcc version 9.4.0 (Ubuntu 9.4.0-1ubuntu1~20.04.1)
Intel NIC Device ID	0x1592
Intel NIC Firmware version	4.0
Intel NIC Package	1.3.30.0
DPDK version	22.07
Test configuration	2 NIC cards attached to the same processor and 1 port used of each NIC. Each port has 1 queue assigned for a total of 2 queues All queues are assigned to the same core.

Figure 13: Test #6 Setup – 2x Intel® Ethernet Network Adapter E810-CQDA2 connected to IXIA




Table 17: Test #6 Setting

Item	Description
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Disabled
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=1-21 nohz_full=1-21 rcu_nocbs=1-21 idle=poll nmi_watchdog=0 audit=0 nosoftlockup processor.max_cstate=0 intel_idle.max_cstate=0 hpet=disable mce=off tsc=reliable numa_balancing=disable
DPDK Settings	CC=gcc meson -Dlibdir=lib -Dexamples=l3fwd -Dc_args=- DRTE_LIBRTE_ICE_16BYTE_RX_DESC --default-library=static x86_64-native- linuxapp-gcc

Test Results

Subcase #1: Turbo boost comparison

Advanced setting:


Item	Description
BIOS	Intel® Hyper-Threading Tech Disabled Turbo Boost changed in the test
Command line	dpdk-testpmd -l 5,6 -n 8 --force-max-simd-bitwidth=512 -a 0000:31:00.0 -a 0000:4b:00.0 -- -i --portmask=0x1 --rxq=1 --txq=1 --txd=1024 --rxd=1024 --forward=io -a

Table 18: Test #6 Subcase #1 Result

Frame Size (Bytes)	Line Rate [2*100G] (Mpps)	Turbo Boost Disabled		Turbo Boost Enabled		Disable vs. Enable %
		CPU Freq	2.3	CPU Freq	2.9	
		Frame Rate (Mpps)	Cycles per Packet=CPU _freq/Frame_rate	Frame Rate (Mpps)	Cycles per Packet=CPU _freq/Frame_rate	
64	297.62	109.30	21	122.78	24	12.34
128	168.92	109.01	21	122.56	24	12.43
256	90.58	90.58	-	90.58	-	-
512	46.99	46.99	-	46.99	-	-
1024	23.95	23.95	-	23.95	-	-
1280	19.23	19.23	-	19.23	-	-
1518	16.26	16.26	-	16.26	-	-

Note: All packet sizes over 256B reach line rate, so the single core is not fully stressed.

Figure 14: ice PMD Single core performance on 2x Intel® Ethernet Network Adapter E810-QCDA2 with turbo boost comparison



Subcase #2: Instruction comparison

Advanced setting:


Item	Description
BIOS	Intel® Hyper-Threading Tech Disabled Turbo Boost Enabled
Command line	dpdk-testpmd -l 5,6 -n 8 -a 0000:31:00.0 -a 0000:4b:00.0 -- -i --portmask=0x1 --rxq=1 --txq=1 --txd=1024 --rxd=1024 --forward=io -a add "--force-max-simd-bitwidth=512" if using AVX512, default is AVX2

Table 19: Test #6 Subcase #2 Result

Frame Size (Bytes)	Line Rate [2*100G] (Mpps)	AVX2		AVX512		AVX512 vs. AVX2 %
		CPU Freq	3.0	CPU Freq	2.9	
		Frame Rate (Mpps)	Cycles per Packet=CPU _freq/Frame_rate	Frame Rate (Mpps)	Cycles per Packet=CPU _freq/Frame_rate	
64	297.62	94.79	32	122.78	24	29.53
128	168.92	94.66	32	122.56	24	29.48
256	90.58	90.58	-	90.58	-	-
512	46.99	46.99	-	46.99	-	-
1024	23.95	23.95	-	23.95	-	-
1280	19.23	19.23	-	19.23	-	-
1518	16.26	16.26	-	16.26	-	-

Note: All packet sizes over 256B reach line rate, so the single core is not fully stressed.

Figure 15: ice PMD Single core performance on 2x Intel® Ethernet Network Adapter E810-QCDA2 with instruction comparison



Subcase #3: Hyper-threading comparison

Advanced setting:


Item	Description
BIOS	Intel® Hyper-Threading Tech changed in the test Turbo Boost Enabled
Command line	1C1T: dpdk-testpmd -l 5,6 -n 8 --force-max-simd-bitwidth=512 -a 0000:31:00.0 -a 0000:4b:00.0 -- -i --portmask=0x1 --rxq=1 --txq=1 --txd=1024 --rxd=1024 --forward=io -a 1C2T: dpdk-testpmd -l 5,6,86 -n 8 --force-max-simd-bitwidth=512 -a 0000:31:00.0 -a 0000:4b:00.0 -- -i --portmask=0x1 --rxq=1 --txq=1 --txd=1024 --rxd=1024 --forward=io -a - -nb-cores=2 Note: 6 and 86 come from the same physical core.

Table 20: Test #6 Subcase #3 Result

Frame Size (Bytes)	Line Rate [2*100G] (Mpps)	1C1T	1C2T	1C2T vs. 1C1T %
		Frame Rate (Mpps)	Frame Rate (Mpps)	
64	297.62	122.78	161.96	31.91
128	168.92	122.56	161.32	31.62
256	90.58	90.58	90.58	-
512	46.99	46.99	46.99	-
1024	23.95	23.95	23.95	-
1280	19.23	19.23	19.23	-
1518	16.26	16.26	16.26	-

Note: All packet sizes over 256B reach line rate, so the single core is not fully stressed.

Figure 16: ice PMD Single core performance on 2x Intel® Ethernet Network Adapter E810-QCDA2 with hyper-threading comparison



Test#7 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-CQDA2 using SR-IOV over KVM

Table 21: Test #7 Setup

Item	Description
Test Case	Test#7 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-CQDA2 using SRIOV over KVM
Server Platform	Inspur® Server Board NF5280M6
CPU	Intel® Xeon® Platinum 8380 CPU @ 2.30GHz 40 CPU cores * 2 NUMA nodes
Memory	256GB: 16GB x 8 DIMMs x 2 NUMA nodes @ 3200MHz
NIC	1x Intel® Ethernet Network Adapter E810-CQDA2; PCIe 3.0/4.0 x16
BIOS	1.1a
Microcode	0xd0002c1
Hypervisor	Ubuntu 20.04.4 LTS QEMU emulator version 4.2.1 (Debian 1:4.2-3ubuntu6.17)
Hypervisor kernel version	5.13.0-41-generic
Hypervisor Intel Driver	ice-1.9.5
Guest Operating System	Ubuntu 20.04 LTS
Guest kernel version	5.4.0-52-generic
Guest GCC version	gcc version 9.3.0 (Ubuntu 9.3.0-10ubuntu2)
Intel NIC Device ID	0x1592
Intel NIC Firmware version	4.0
Intel NIC Package	1.3.30.0
DPDK version	22.07
Test configuration	1 NIC card attached to the processor and 1 port used to generate 1 VF. VF has 4 queues, 1 queue per logical core. Totaling 4 logical cores, 4 queues for 1VF.

Figure 17: Test #7 Setup – 1x Intel® Ethernet Network Adapter E810-CQDA2 connected to IXIA (Single port)

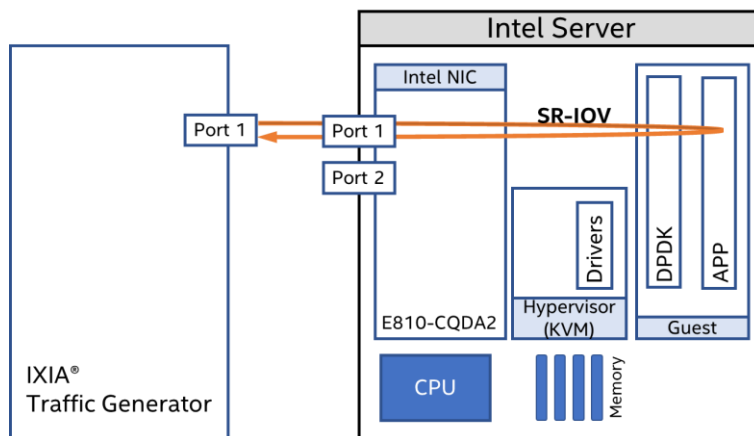


Table 22: Test #7 Setting

Item	Description
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Enabled Intel® Hyper-Threading Tech Disabled Turbo Boost Enabled Intel® Virtualization Technology Enabled Intel® VT for Directed I/O Enabled
Hypervisor Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=1-21 nohz_full=1-21 rcu_nocbs=1-21 intel_pstate=disable idle=poll
Hypervisor settings	<pre> 1, creating VF echo 1 >/sys/bus/pci/devices/0000\:18\:00.0/sriov_numvfs 2, setting VF MAC ip link set ens785f0 vf 0 mac A6:9C:BB:07:55:C0 3, binding VF to vfio-pci ./usertools/dpdk-devbind.py -b vfio-pci 0000:18:01.0 4, mount hugepage mount -t hugetlbfs nodev /mnt/huge 5, launch qemu qemu-system-x86_64 -name vm0 -enable-kvm --monitor stdio -monitor unix:/tmp/vm0_monitor.sock,server,nowait -device e1000,netdev=nttstp1 -netdev user,id=nttstp1,hostfwd=tcp:127.0.0.1:6013-:22 -device vfio- pci,host=0000:18:01.0,id=pt_0 -cpu host -m 32768 -object memory-backend- file,id=mem,size=32768M,mem-path=/mnt/huge,share=on -numa node,memdev=mem -mem-prealloc -smp cores=10,sockets=1 -chardev socket,path=/tmp/vm0_qga0.sock,server=on,wait=off,id=vm0_qga0 -device virtio-serial -device virtserialport,chardev=vm0_qga0,name=org.qemu.guest_agent.0 -vnc :10 -drive file=/home/image/ubuntu2004.img,if=virtio,index=0,media=disk 6, vcpu pinning 6.1 get vcpu thread ids by qemu monitor (qemu) info cpus * CPU #0: thread_id=6106 CPU #1: thread_id=6107 CPU #2: thread_id=6108 CPU #3: thread_id=6109 CPU #4: thread_id=6110 CPU #5: thread_id=6111 CPU #6: thread_id=6112 CPU #7: thread_id=6113 CPU #8: thread_id=6114 CPU #9: thread_id=6115 6.2 pin thread ids to logical cores taskset -pc 9 6115 taskset -pc 8 6114 taskset -pc 7 6113 taskset -pc 6 6112 </pre>

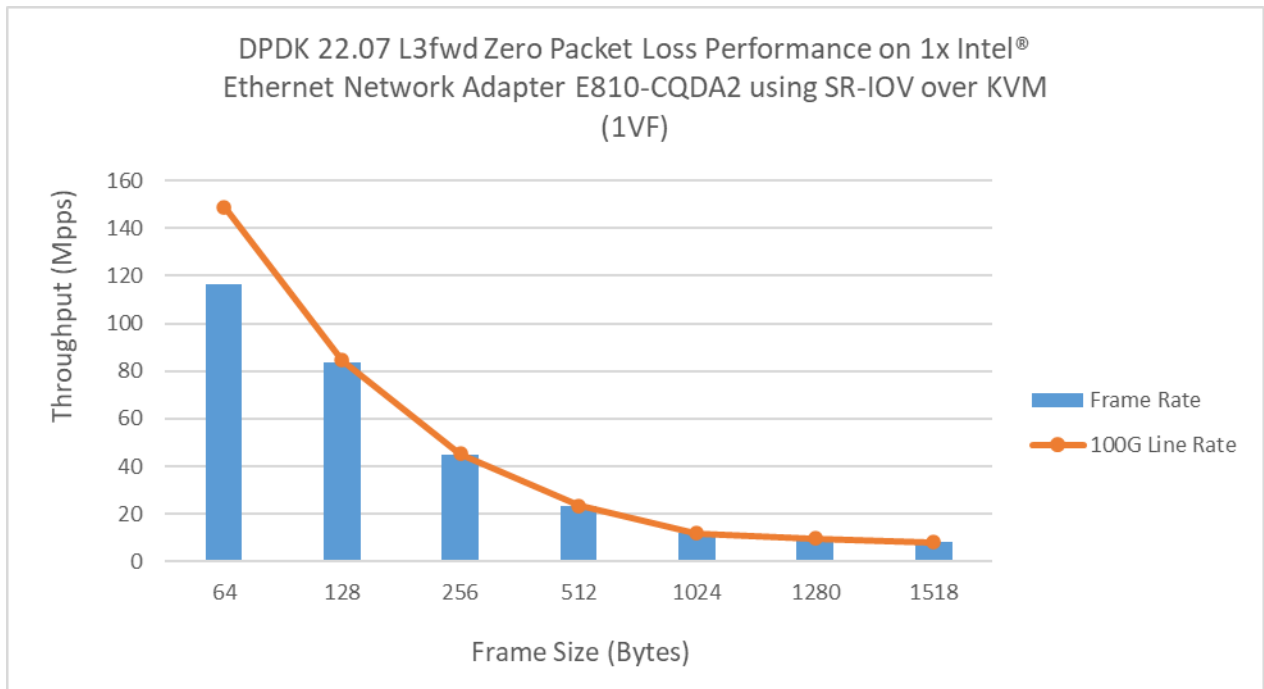
	<pre>taskset -pc 5 6111 taskset -pc 4 6110 taskset -pc 13 6109 taskset -pc 12 6108 taskset -pc 11 6107 taskset -pc 10 6106</pre>
Other optimizations on Hypervisor	<pre>echo -1 > /proc/sys/kernel/sched_rt_period_us echo -1 > /proc/sys/kernel/sched_rt_runtime_us echo 10 > /proc/sys/vm/stat_interval echo 0 > /proc/sys/kernel/watchdog_thresh</pre>
Guest Boot settings	<pre>default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=4-9 nohz_full=4-9 rcu_nocbs=4-9 intel_pstate=disable idle=poll processor.max_cstate=0 intel_idle.max_cstate=0</pre>
Other optimizations on Guest	<pre>echo 1 > /sys/module/vfio/parameters/enable_unsafe_noiommu_mode echo -1 > /proc/sys/kernel/sched_rt_period_us echo -1 > /proc/sys/kernel/sched_rt_runtime_us echo 10 > /proc/sys/vm/stat_interval echo 0 > /proc/sys/kernel/watchdog_thresh</pre>
DPDK Settings on Guest	<pre>CC=gcc meson -Dlibdir=lib -Dexamples=l3fwd --default-library=static x86_64-native-linuxapp-gcc</pre>
Command line on Guest	<pre>dpdk-l3fwd -l 4-7 -n 6 -a 00:04:0 -- -p 0x1 --config '(0,0,5),(0,1,6),(0,2,7),(0,3,4)' --parse-ptype</pre>

Test Results

Table 23: Test #7 Result

Frame Size (Bytes)	Line Rate[100G] (Mpps)	Frame Rate (Mpps)	% Line Rate
64	148.81	116.37	78.20
128	84.46	83.87	99.30
256	45.29	44.97	99.29
512	23.5	23.5	100
1024	11.97	11.97	100
1280	9.62	9.62	100
1518	8.13	8.13	100

Figure 18: RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-CQDA2 using SRIOV over KVM



Test#8 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-XXVDA4

Table 24: Test #8 Setup

Item	Description
Test Case	Test#8 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-XXVDA4
Server Platform	Supermicro® Server Board SYS-730A-I
CPU	Intel® Xeon® Platinum 8380 CPU @ 2.30GHz 40 CPU cores * 2 NUMA nodes
Memory	256GB: 16GB x 8 DIMMs x 2 NUMA nodes @ 3200MHz
NIC	1x Intel® Ethernet Network Adapter E810-XXVDA4; PCIe 3.0/4.0 x16
BIOS	1.1a
Microcode	0xd0002c1
Operating System	Ubuntu 20.04.4 LTS
Linux kernel version	5.13.0-41-generic
GCC version	gcc version 9.4.0 (Ubuntu 9.4.0-1ubuntu1~20.04.1)
Intel NIC Device ID	0x1593
Intel NIC Firmware version	4.0
Intel NIC Package	1.3.30.0
DPDK version	22.07
Test configuration	1 NIC card attached to the processor and 4 ports used. Each port assigned 1 core and using 1 queue. Totaling 4 logical core, 4 queues for 4 ports.

Figure 19: Test #8 Setup – 1x Intel® Ethernet Network Adapter E810-XXVDA4 connected to IXIA (4 ports of 1 NIC)

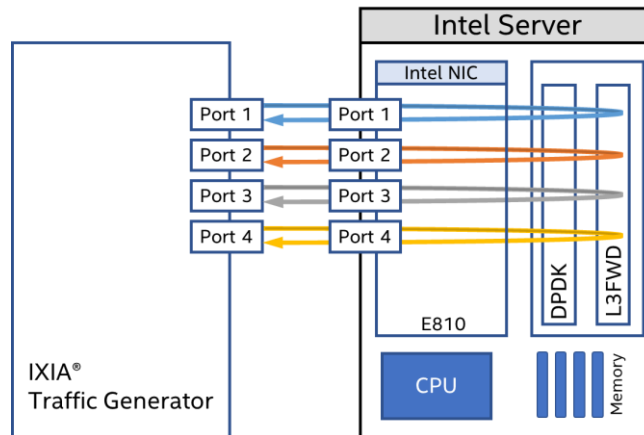


Table 25: Test #8 Setting


Item	Description
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Enabled Intel® Hyper-Threading Tech Disabled Turbo Boost Enabled
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=1-21 nohz_full=1-21 rcu_nocbs=1-21 idle=poll nmi_watchdog=0 audit=0 nosoftlockup processor.max_cstate=0 intel_idle.max_cstate=0 hpet=disable mce=off tsc=reliable numa_balancing=disable
DPDK Settings	CC=gcc meson -Dlibdir=lib -Dexamples=l3fwd -Dc_args=- DRTE_LIBRTE_ICE_16BYTE_RX_DESC --default-library=static x86_64-native- linuxapp-gcc
Command line	dpdk-l3fwd -l 5-8 -n 8 --force-max-simd-bitwidth=512 -- -p 0xf --config '(0,0,5),(1,0,6),(2,0,7),(3,0,8)'

Test Results

Table 26: Test #8 Result

Frame Size (Bytes)	Line Rate[4*25G] (Mpps)	Frame Rate (Mpps)	% Line Rate
64	148.81	116.37	78.20
128	84.46	84.46	100
256	45.29	45.29	100
512	23.5	23.5	100
1024	11.97	11.97	100
1280	9.62	9.62	100
1518	8.13	8.13	100

Figure 20: RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-XXVDA4



Test#9 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-2CQDA2

Table 27: Test #9 Setup

Item	Description
Test Case	Test#9 RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-2CQDA2
Server Platform	Supermicro® Server Board SYS-730A-I
CPU	Intel® Xeon® Platinum 8380 CPU @ 2.30GHz 40 CPU cores * 2 NUMA nodes
Memory	128GB: 16GB x 4 DIMMs x 2 NUMA nodes @ 3200MHz
NIC	1x Intel® Ethernet Network Adapter E810-2CQDA2; PCIe 4.0 x 16
BIOS	1.1a
Microcode	0xd0002c1
Operating System	Ubuntu 20.04.4 LTS
Linux kernel version	5.13.0-41-generic
GCC version	gcc version 9.4.0 (Ubuntu 9.4.0-1ubuntu1~20.04.1)
Intel NIC Device ID	0x1592
Intel NIC Firmware version	4.0
Intel NIC Package	1.3.30.0
DPDK version	22.07
Test configuration	1 NIC cards and 2 ports of the NIC are used. Each port has 4 queues assigned, 1 queue per logical core. Totaling 8 logical cores, 8 queues for 2 ports

Figure 21: Test #9 Setup – 1x Intel® Ethernet Network Adapter E810-2CQDA2 connected to IXIA (2 ports used on 1 NIC)

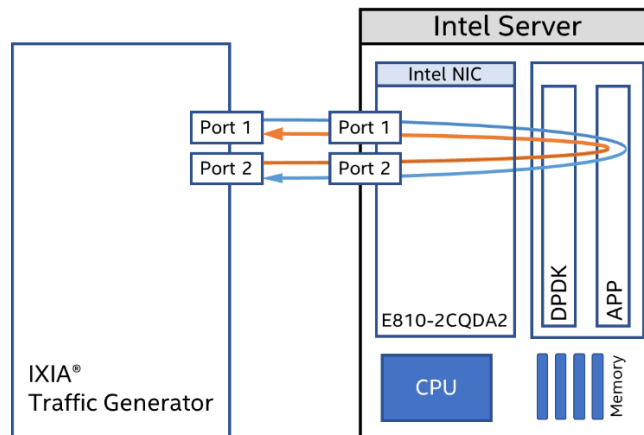


Table 28: Test #9 Setting


Item	Description
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Enabled Intel® Hyper-Threading Tech Disabled Turbo Boost Enabled
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 intel_iommu=on iommu=pt isolcpus=1-21 nohz_full=1-21 rcu_nocbs=1-21 idle=poll nmi_watchdog=0 audit=0 nosoftlockup processor.max_cstate=0 intel_idle.max_cstate=0 hpet=disable mce=off tsc=reliable numa_balancing=disable
DPDK Settings	CC=gcc meson -Dlibdir=lib -Dexamples=l3fwd -Dc_args=- DRTE_LIBRTE_ICE_16BYTE_RX_DESC --default-library=static x86_64-native- linuxapp-gcc
Command line	dpdk-l3fwd -l 5-12 -n 8 --force-max-simd-bitwidth=512 -- -p 0x3 --config '(0,0,5),(1,0,6),(0,1,7),(1,1,8),(0,2,9),(1,2,10),(0,3,11),(1,3,12)'

Test Results

Table 29: Test #9 Result

Frame Size (Bytes)	Line Rate[2x100G] (Mpps)	Frame Rate (Mpps)	% Line Rate
64	297.62	228.56	76.80
128	168.92	152.29	90.15
256	90.58	90.58	100.00
512	46.99	46.99	100.00
1024	23.95	23.95	100.00
1280	19.23	19.23	100.00
1518	16.26	16.26	100.00

Figure 22: RFC2544 zero packet loss test on 1x Intel® Ethernet Network Adapter E810-2CQDA2-



Test#10 Zero packet loss test on Intel server of m6id.metal AWS cloud instance

Table 30: Test #10 Setup

Item	Description
Test Case	Test#10 Zero packet loss test on 1x Amazon.com, Inc. Elastic Network Adapter & Intel® Xeon® Platinum 8375C CPU
Server Platform	Amazon EC2 m6id.metal
CPU	Intel® Xeon® Platinum 8375C CPU @ 2.90GHz
Memory	512GB (16x32GB DDR4 3200MT/s [3200MT/s])
vCPU	128
NIC	Elastic Network Adapter (5.13.0-1031-aws)
BIOS	1.0
Microcode	0xd000363
Operating System	Ubuntu 20.04.1 LTS
Linux kernel version	5.13.0-1031-aws
GCC version	gcc version 9.4.0 (Ubuntu 9.4.0-1ubuntu1~20.04.1)
NIC Device ID	0xec20
DPDK version	22.07
Test configuration	Interface 1 of each AWS instance's NIC is connected, and bandwidth is 50Gbps.

Figure 23: Test #10 Setup – 1x Elastic Network Adapter connected to TREX on both Intel server of m6id.metal AWS cloud instances

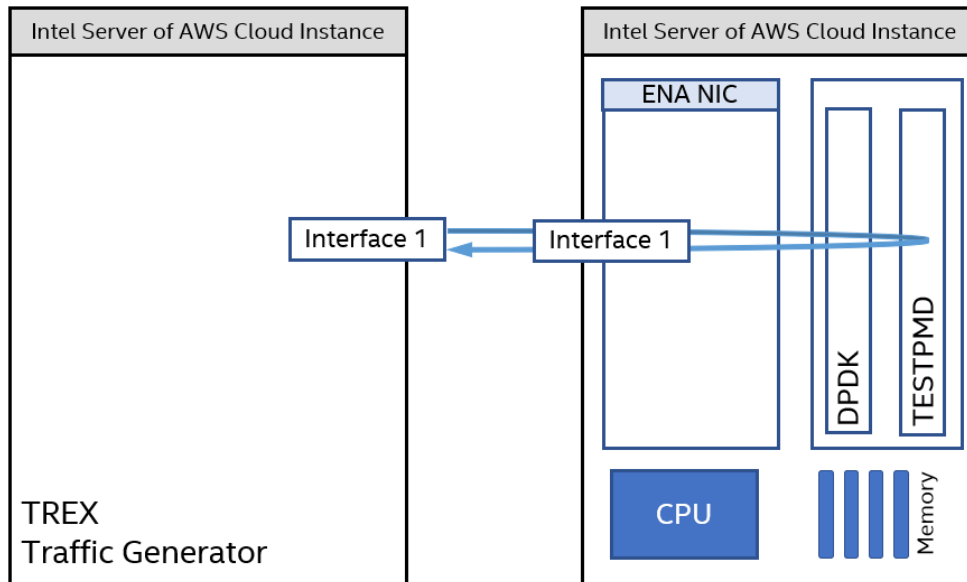


Table 31: Test #10 Setting

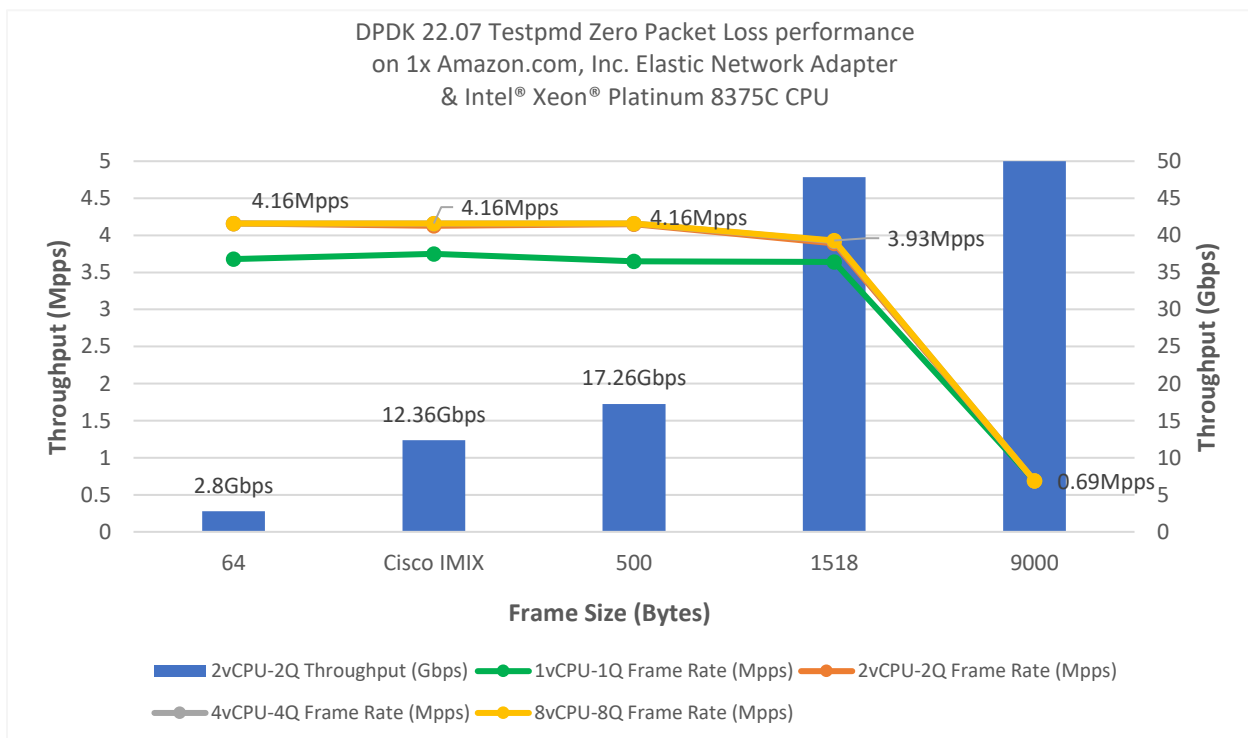
Item	Description
BIOS	Hyper-Threading Enabled Turbo Boost Enabled
Boot settings	intel_iommu=on iommu=pt default_hugepagesz=1G hugepagesz=1G hugepages=16
Command line	<p>1vCPU-1Q: dpdk-testpmd -l 0,1 -a 0000:31:00.0 -- -i --forward-mode=mac --txd=1024 --rxd=2048 --txq=1 --rxq=1 --eth-peer=0, XX:XX:XX:XX:XX:XX --nb-cores=1 --rss-udp</p> <p>2vCPU-2Q: dpdk-testpmd -l 0,1,2 -a 0000:31:00.0 -- -i --forward-mode=mac --txd=1024 --rxd=2048 --txq=2 --rxq=2 --eth-peer=0, XX:XX:XX:XX:XX:XX --nb-cores=2 --rss-udp</p> <p>4vCPU-4Q: dpdk-testpmd -l 0,1,2,3,4 -a 0000:31:00.0 -- -i --forward-mode=mac --txd=1024 --rxd=2048 --txq=4 --rxq=4 --eth-peer=0, XX:XX:XX:XX:XX:XX --nb-cores=4 --rss-udp</p> <p>8vCPU-8Q: dpdk-testpmd -l 0,1,2,3,4,5,6,7,8 -a 0000:31:00.0 -- -i --forward-mode=mac --txd=1024 --rxd=2048 --txq=8 --rxq=8 --eth-peer=0, XX:XX:XX:XX:XX:XX --nb-cores=8 --rss-udp</p> <p>Note: vCPU 0 is used for control path, remained vCPUs are used for data path. --eth-peer=0,XX:XX:XX:XX:XX:XX, set the MAC address XX:XX:XX:XX:XX:XX of the peer interface 1, NIC doesn't support promiscuous mode.</p>

Test Results

Table 32: Test #10 Result (1vCPU-1Q, 2vCPU-2Q, 4vCPU-4Q, 8vCPU-8Q)

Frame Size (Bytes)	1vCPU-1Q Frame Rate (Mpps)	2vCPU-2Q Frame Rate (Mpps)	4vCPU-4Q Frame Rate (Mpps)	8vCPU-8Q Frame Rate (Mpps)	2vCPU-2Q Throughput (Gbps)	% Line Rate
64	3.68	4.16	4.16	4.16	2.8	5.6%
Cisco IMIX	3.75	4.13	4.16	4.16	12.36	24.7%
500	3.65	4.15	4.16	4.16	17.26	34.5%
1518	3.64	3.89	3.92	3.93	47.86	95.7%
9000	0.69	0.69	0.69	0.69	50	100%

Figure 24: Test #10 Result – Zero packet loss test on 1x Amazon.com, Inc. Elastic Network Adapter & Intel® Xeon® Platinum 8375C CPU



Note: Throughputs for 2vCPU-2Q, 4vCPU-4Q and 8vCPU-8Q are almost same, so lines are overlapped. Throughput is limited by the AWS instance configuration on ENA NIC driver especially for small packets (~4.16 Mpps).

DISCLAIMERS

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

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.

Performance results are based on testing as of July 25th and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. For more information go to <http://www.intel.com/performance>

Intel® AES-NI requires a computer system with an AES-NI enabled processor, as well as non-Intel software to execute the instructions in the correct sequence. AES-NI is available on select Intel® processors. For availability, consult your reseller or system manufacturer. **For more information, see <http://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni/>**

© 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.

Copyright © 2022 Intel Corporation. All rights reserved.