


DPDK Performance Report


Release 16.11

Test Date: Nov 16th 2016



Revision History

Date	Revision	Comment
Nov 16 th , 2016	1.0	Initial document for release



Contents

Audience and Purpose.....	4
Test setup:	4
Intel® Xeon® Processor E5-2699 v4 (55M Cache, 2.20 GHz)	7
Hardware & Software Ingredients.....	7
Test Case 1 - RFC2544 zero packet loss test.....	7
Test Case 2 - RFC2544 zero packet loss test.....	8
Test Case 3 - Single core performance test.....	8
Intel® Xeon® Processor E5-2699 v3 (45M Cache, 2.30 GHz)	10
Hardware & Software Ingredients.....	10
Test Case 4 - RFC2544 Zero Packet Loss test	11
Test Case 5 - Single core performance	11

Audience and Purpose

The primary audience for this test report are 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 setup:

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.

* Other names and brands may be claimed as the property of others.

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.

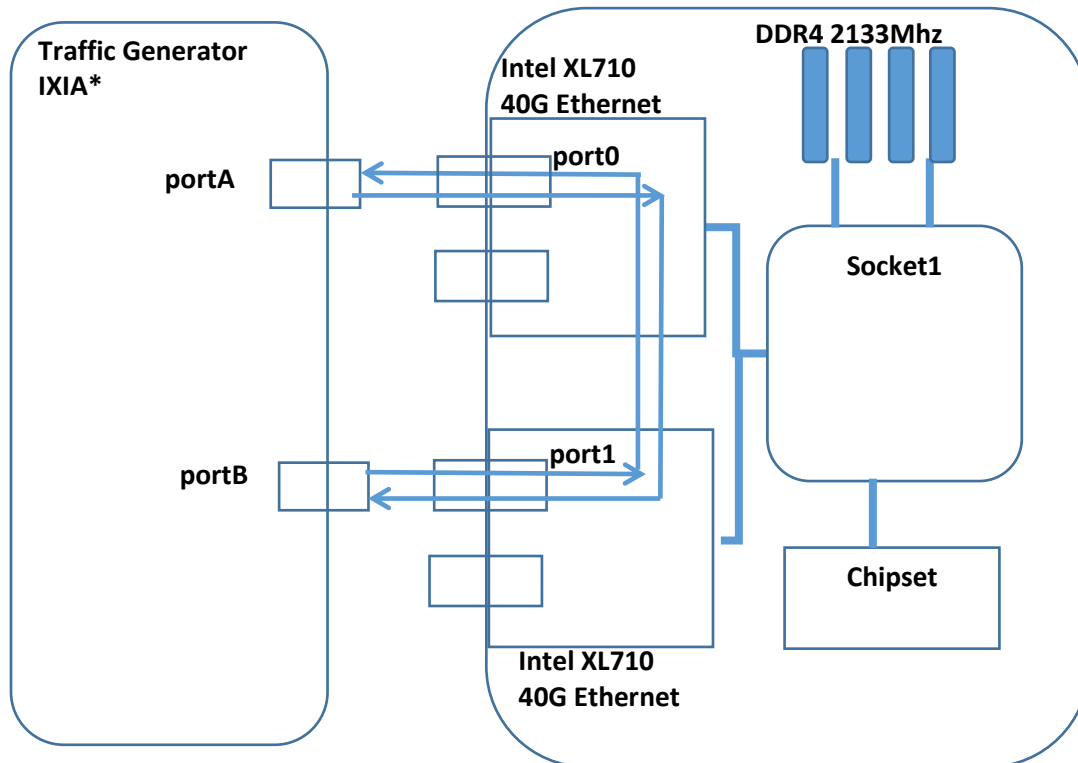


Figure1 DPDK NIC performance test setup



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* 8.12 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 4x 10GbE ports for the ixgbe DPDK Poll Mode Driver (PMD) and 2x 40GbE 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



Intel® Xeon® Processor E5-2699 v4 (55M Cache, 2.20 GHz)

Hardware & Software Ingredients

Item	Description
Server Platform	Intel® Server Board S2600CWR Intel® Server Board S2600CW Family
CPU	Intel® Xeon® Processor E5-2699 v4 (55M Cache, 2.20 GHz) http://ark.intel.com/products/91317/Intel-Xeon-Processor-E5-2699-v4-55M-Cache-2_20-GHz Number of cores 22, Number of threads 44.
Memory	Total 65536 MBs over 8 channels @ 2133 MHz
PCIe	3 x PCIe Gen3 x8 slots
NICs	Intel® Ethernet Converged Network Adapter X710-DA4 (4x10G) 2 x Intel® Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
Operating System	Ubuntu 16.04 LTS
BIOS	SE5C610.86B.01.01.0016.033120161139
Linux kernel version	4.4.0-34-generic
GCC version	gcc (Ubuntu 5.4.0-6ubuntu1~16.04.2) 5.4.0 20160609
DPDK version	16.11

Boot and BIOS settings

Item	Description
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 hugepagesz=2M hugepages=2048 isolcpus=1-11,22-33 nohz_full=1-11,22-33 rcu_nocbs=1-11,22-33 Note: nohz_full and rcu_nocbs is to disable Linux* kernel interrupts, and it's important for zero-packet loss test. Generally, 1G huge pages are used for performance test.
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Disabled Enhanced Intel® Speedstep® Tech Disabled Turbo Boost Disabled
DPDK Settings	Build L3fwd: (in l3fwd/main.c) #define RTE_TEST_RX_DESC_DEFAULT 2048 #define RTE_TEST_TX_DESC_DEFAULT 2048

Test Case 1 - RFC2544 zero packet loss test

Item	Description
Test Case	RFC2544 zero packet loss test for 10G NIC i40e
NIC	Intel® Ethernet Converged Network Adapter X710-DA4 (4x10G)



Driver	i40e DPDK PMD
Device ID	0x1572
Device Driver/ Firmware	Driver version: 1.4.25-k firmware-version: 5.05 0x800028a6 0.0.0
Test configuration	1 NIC card attached to the second processor. 1 port assigned per logical core using 1 queue. Totaling 4 logical cores, 4 queues for 4ports
Command line	./l3fwd -c 0x7800000 -n 4 -w 81:00.0 -w 81:00.1 -w 81:00.2 -w 81:00.3 -- -p 0xf --config '(0,0,23),(1,0,24),(2,0,25),(3,0,26)'

Test Result:

Packet Size(Bytes)	Throughput(Mpps)	Line rate%
64	36.09	60.63
128	33.78	100
256	18.11	100

Note: All packet sizes over 128B reach line rate.

Test Case 2 - RFC2544 zero packet loss test

Item	Description
Test Case	RFC2544 zero packet Loss test for 40G NIC i40e
NIC	2 x Intel® Ethernet Converged Network Adapter XL710-QDA2 (2X40G)
Driver	i40e DPDK PMD
Device ID	0x1583
Device Driver/Firmware	Driver version: 1.4.25-k firmware-version: 5.05 0x8000289d 0.0.0
Test configuration	2 NIC cards attached to the second processor only 1 st 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
Command line	./l3fwd -c 0x7800000 -n 4 -w 82:00.0 -w 85:00.0 -- -p 0x3 --config '(0,0,23),(0,1,24),(1,0,25),(1,1,26)'

Test Result:

Packet Size (Bytes)	Throughput (Mpps)	Line rate%
64	71.34	59.92
128	67.56	100
256	36.23	100

Note: All packet sizes over 128B reach line rate.


Test Case 3 - Single core performance test



Item	Description
Test Case	Single core performance test for i40e
NIC	Intel® Ethernet Converged Network Adapter XL710-QDA2 (2X40G), Quantities: 2 NICs
Driver	i40e DPDK PMD
Device ID	0x1572
Device Firmware	Driver version: 1.4.25-k firmware-version: 5.05 0x800028a6 0.0.0
Test configuration	Change the config/common_base file and rebuild DPDK. “CONFIG RTE_LIBRTE_I40E_16BYTE_RX_DESC=y” . 2 NIC cards attached to the first processor only 1 st port used of each NIC. Each port has 1 queue assigned. All queues are assigned to the same logical core. The first core defined is for the CLI of the Testpmd example due to the <code>-i</code> argument. I40E Vector PMD and IO forwarding is used for the test. Test duration is 60 seconds.
Command line	<code>./testpmd -c 0x180000 -n 4 -- -i --txqflags=0xf01 --txrst=32 --txfreet=32 --txd=128</code> Testpmd>start

Test Result:

Packet Size (Bytes)	Throughput (Mpps)	Line Rate%	CPU freq(Ghz)	Cycles per packet= CPU_freq/Throughput
64	44.07	37.02	2.2	50



Intel® Xeon® Processor E5-2699 v3 (45M Cache, 2.30 GHz)

Hardware & Software Ingredients

Item	Description
Server Platform	Intel® Server Board S2600CWR Intel® Server Board S2600CW Family
CPU	Intel® Xeon® Processor E5-2699 v3 (45M Cache, 2.30 GHz) http://ark.intel.com/products/81061/Intel-Xeon-Processor-E5-2699-v3-45M-Cache-2_30-GHz Number of cores 18, Number of threads 36.
Memory	Total 65536 MBs over 8 channels @ 2133 MHz
PCIe	4x PCIe Gen3 x8 slots
NICs	4x Intel 82599ES 10-Gigabit SFI/SFP+ Network Connection
Operating System	Fedora 23
BIOS	SE5C610.86B.11.01.0044.082520150603
Linux kernel version	4.6.4-201.fc23.x86_64
GCC version	gcc (GCC) 5.3.1 20160406 (Red Hat 5.3.1-6)
DPDK version	16.11

Boot and BIOS settings

Item	Description
Boot settings	default_hugepagesz=1G hugepagesz=1G hugepages=16 hugepagesz=2M hugepages=2048 isolcpus=1-29 nohz_full=1-29 rcu_nocbs=1-29 Note: nohz_full and rcu_nocbs is to disable Linux system interrupts, and it's important for zero-packet loss test. Generally, 1G huge pages are used for performance test.
BIOS	CPU Power and Performance Policy <Performance> CPU C-state Disabled CPU P-state Disabled Enhanced Intel® Speedstep® Tech Disabled Turbo Boost Disabled
DPDK Settings	Build L3fwd: (in l3fwd/main.c) #define RTE_TEST_RX_DESC_DEFAULT 2048 #define RTE_TEST_TX_DESC_DEFAULT 2048



Test Case 4 - RFC2544 Zero Packet Loss test

Item	Description
Test Case	RFC2544 Zero Packet Loss test for 10G NIC
NIC	2 x Intel® 82599ES 10 Gigabit Ethernet Controller
Driver	ixgbe DPDK PMD
Device ID	0x10fb
Device Firmware	0x61bf0001
Test configuration	2 NIC cards attached to the second processor, using one port per NIC. 1 port assigned per logical core using 1 queue. Totally 2 logical cores, 2 queues for 2ports.
Command line	<code>./l3fwd -c 0xc0000 -n 4 -w 83:00.0 -w 87:00.0 -- -p 0x3 --config '(0,0,18),(1,0,19)'</code>

Test Result:

Packet Size (Bytes)	Throughput (Mpps)	Line Rate%
64	29.55	99.30
128	16.77	99.30
256	8.99	99.30
512	4.66	99.30
1024	2.39	100

Note: All packet sizes over 1024B reach line rate.

Test Case 5 - Single core performance


Item	Description
Test Case	Single core performance for ixgbe
NIC	4 x Intel® 82599ES 10 Gigabit Ethernet Controller, using 4 NIC cards
Driver	ixgbe
Device ID	0x10fb
Device Firmware	0x61bf0001
Test configuration	4 NIC cards attached to the second processor only 1 st port used of each NIC. Each port has 1 queue assigned. All queues are assigned to the same logical core. The first core defined is for the CLI of the Testpmd example due to the <code>-i</code> argument. Ixgbe Vector PMD and IO forwarding is used for the test. Test duration is 60 seconds.



Command line	<pre>./testpmd -c 0x180000 -n 4 -- -i --txqflags=0xf01 --txrst=32 --txfreet=32 -- txd=128 Testpmd>start</pre>

Test Result:

Packet Size (Bytes)	Throughput (Mpps)	Line Rate%	CPU freq(Ghz)	Cycles per packet= CPU_freq/Throughput
64	48.78	81.9	2.3	47



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.

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/>**

Copyright © 2016 Intel Corporation. All rights reserved.