NVIDIA ConnectX-7 400 GbE Performance Report with DPDK 24.07 and Grace C1 CPU (2x200 GbE)

Rev 1.0

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

For the complete and most updated list of Mellanox trademarks, visit http://www.mellanox.com/page/trademarks.

Copyright

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051

http://www.nvidia.com

Table of Contents

Document History

Table 1 - Document History

Version	Date	Description of Change
1.0	09-01-2025	Initial report release

1 About this Report

The purpose of this report is to provide packet rate performance data for NVIDIA® Network Interface Card (NIC) (ConnectX®-7 200GbE) that has been achieved while being used with the Data Plane Development Kit (DPDK) release 24.07, running on a server with NVIDIA® Grace CPU. The report provides the measured packet rate performance as well as the hardware layout, procedures, and configurations for replicating these tests.

The document does not cover all network speeds available with the ConnectX® or BlueField® family of NICs / DPUs and is intended as a general reference of achievable performance for the specified DPDK release.

1.1 Target Audience

This document is intended for engineers implementing applications with DPDK to guide and help achieving optimal performance.

1.2 Terms and Conventions

The following terms, abbreviations, and acronyms are used in this document.

Table 2 - Terms, Abbreviations and Acronyms

Term	Description
DPU	Data Processing Unit
DOCA	Data Center Infrastructure-on-a-Chip Architecture is NVIDIA's software framework designed to unlock the capabilities of BlueField DPUs and SuperNICs, enabling offload, acceleration, and isolation of data center workloads through a unified SDK and runtime environment.
DUT	Device Under Test
MPPS	Million Packets Per Seconds
PPS	Packets Per Second
ZPL	Zero Packet Loss

2 Test Description

2.1 Hardware Components

The following hardware components are used in the test setup:

- NVIDIA sample server based on the NVIDIA Grace CPU C1
- NVIDIA ConnectX-7 Network Interface Card (NIC)
- IXIA® XM12 packet generator

2.2 Zero Packet Loss Test

Zero Packet Loss tests utilize testpmd https://doc.dpdk.org/guides/testpmd_app_ug/index.html as the test application for testing maximum throughput with zero packet loss at various frame sizes based on RFC2544 https://tools.ietf.org/html/rfc2544.

The packet generator transmits a specified frame rate towards the Device Under Test (DUT) and counts the received frame rate sent back from the DUT. Throughput is determined with the maximum achievable transmit frame rate and is equal to the received frame rate i.e. zero packet loss.

- Duration for each test is 60 seconds.
- Traffic of 8192 IP flows is generated per port.
- IxNetwork (Version 9.37EA) is used with the IXIA packet generator.

3 Test#1 NVIDIA ConnectX-7 400GbE Throughput at Zero Packet Loss (2x 200GbE)

Table 1: Test #1 Setup

Item	Description
Test #1	NVIDIA ConnectX-7 400GbE Throughput at zero packet loss
Server	NVIDIA sample server based on the NVIDIA Grace CPU C1
CPU	NVIDIA Grace C1 Neoverse-V2 CPU @ 3.4 GHz 72 CPU cores
RAM	512GB: 32 * 16GB DIMMs @ 3100MHz
BIOS	NVIDIA Grace A02 – version: 02.03.05
NIC	NVIDIA CX713106AE-HEA_QP1_Ax ConnectX-7 HHHL adapter Card; 200GbE (default mode) / NDR200 IB; Dual-port QSFP112; PCIe 5.0 x16 with x16 PCIe extension option; Crypto Enabled; Secure Boot Capable.
Operating System	Red Hat Enterprise Linux 9.5 (Plow)
Kernel Version	5.14.0-503.11.1.el9_5.aarch64
GCC version	gcc (GCC) 11.5.0 20240719 (Red Hat 11.5.0-2)
Mellanox NIC firmware version	28.42.1000
DOCA version	DOCA 2.8 (August 2024)
DPDK version	24.07
Test Configuration	1 NIC, 2 ports used on NIC; each Port has 16 queues assigned to it, 1 queue per logical core for a total of 16 logical cores. Each port receives a stream of 8192 IP flows from the IXIA

The Device Under Test (DUT) is a NVIDIA server equipped with an NVIDIA ConnectX-7 dual-port NIC, two ports are used for this test. The DUT is connected to an IXIA packet generator, which generates traffic directed to the ConnectX-7 NIC. The ConnectX-7 receives the data traffic, which is processed through the DPDK by the test application **testpmd** and the redirected in the opposite direction on the ports. IXIA measures both throughput and packet loss.

Figure 1: Test #1 Setup – NVIDIA ConnectX-7 200GbE connected to IXIA

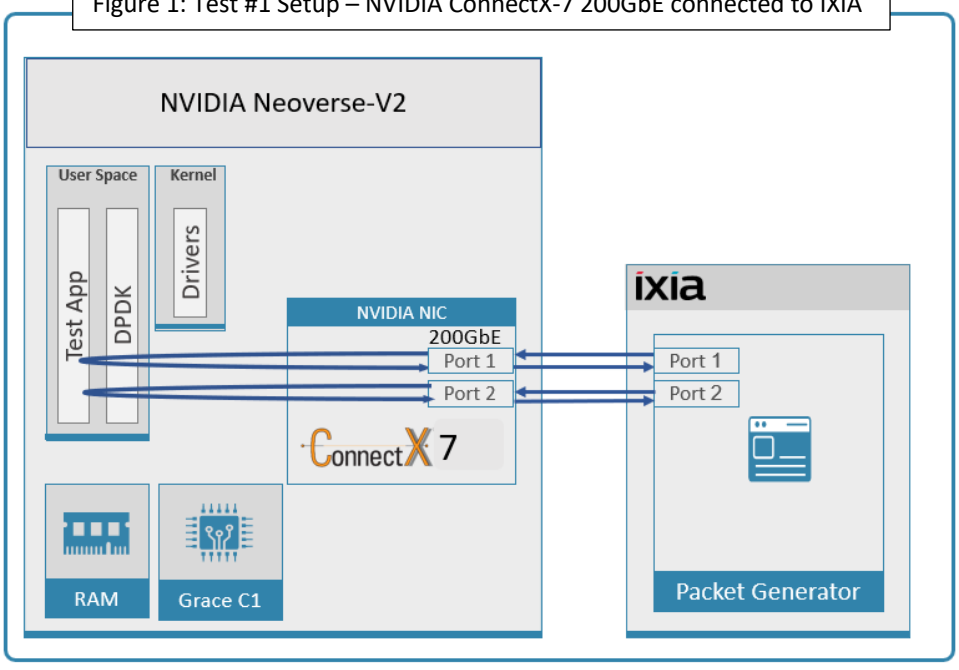

3.1 Test Settings

Table 2: Test #1 Settings

Item	Description
BIOS Settings	PCIe Max Payload Size was set to 256B BIOS settings > Grace Configuration > Socket Configuration > PCIe Configuration > Max Payload size.
BOOT Settings	ro rhgb quiet isolcpus=30-50 nohz_full=30-50 rcu_nocbs=30-50 default_hugepagesz=1G hugepagesz=1G hugepages=80
DPDK Settings	Compile DPDK using: meson <build> ; ninja -C <build> During testing, testpmd was given real-time scheduling priority.
Command Line	./build/app/dpdk-testpmd -n 6 -l 32-48 --socket-mem=4096 -a 0000:01:00.0,mprq_en=1,rxqs_min_mprq=1,mprq_log_stride_num=9,txq_inline_mpw=128,rxq_pkt_pad_en=1 -a 0000:01:00.1,mprq_en=1,rxqs_min_mprq=1,mprq_log_stride_num=9,txq_inline_mpw=128,rxq_pkt_pad_en=1 -- -i -a --burst=64 --mbcache=512 --rxd=8192 --txd=8192 --rxq=16 --txq=16 --nb-cores=16 --forward-mode=io --port-topology=loop
Other optimizations	a) Flow Control OFF: "ethtool -A $netdev rx off tx off" (for both ports) b) Memory optimizations: "sysctl -w vm.zone_reclaim_mode=0"; "sysctl -w vm.swappiness=0" c) Move all IRQs to far NUMA node: "IRQBALANCE_BANNED_CPUS=$LOCAL_NUMA_CPUMAP irqbalance --oneshot" d) Disable irqbalance: "systemctl stop irqbalance" e) Set CQE COMPRESSION to “AGGRESSIVE”: mlxconfig -d $PORT_PCI_ADDRESS set CQE_COMPRESSION=1 f) Remove DUT ports from DHCP Network management: “nmcli dev set $netdev managed no” (for both ports) g) Disable Linux realtime throttling: echo -1 > /proc/sys/kernel/sched_rt_runtime_us h) Change PCI MaxReadReq to 4096B for each port of the NIC: Run "setpci -s $PORT_PCI_ADDRESS 68.w", it will return 4 digits ABCD --> Run "setpci -s $PORT_PCI_ADDRESS 68.w=5BCD”

3.2 Test Results - Line Rate with Packet Size >= 512 Bytes

Table 3: Test #1 Results – NVIDIA ConnectX-7 2x200GbE Zero Packet Loss Throughput

Frame Size (Bytes)	Frame Rate (Mpps)	Line Rate [400G] (Mpps)	% Line Rate
64	264.82	595.24	44.49
128	186.56	337.84	55.22
256	139.98	181.16	77.27
512	93.98	93.98	100
1024	47.89	47.89	100
1280	38.46	38.46	100
1518	32.51	32.51	100

Figure 2: Test #1 Results – NVIDIA ConnectX-7 2x200GbE Throughput at Zero Packet Loss