



Ethernet NIC Controllers DPDK 23.07 Performance Report

Test Report

Copyright © 2023 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to www.broadcom.com. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Table of Contents

1 P225p 2x 25G Zero Frame Loss	4
2 P425G 4x 25G Zero Frame Loss – PCIe 4.0	6
3 P2100G 1x 100G Zero Frame Loss – PCIe 4.0	9
4 P2100G 2x 100G Zero Frame Loss – PCIe 4.0	12
5 P2100G 2x 100G Zero Frame Loss (Icelake) – PCIe 4.0	15
6 P2100G 1x 200G Zero Frame Loss (Icelake) – PCIe 4.0	18
Revision History	21

1 P225p 2x 25G Zero Frame Loss

RFC2544 Zero Frame Loss Test on Broadcom Ethernet NIC P225p with 2x 25G traffic.

Table 1: Setup of P225p 2x 25G Zero Frame Loss Test

Item	Description
Server	Dell PowerEdge R740
CPU	Intel Xeon Gold 6154 CPU at 3.00 GHz. All cores used are on the local socket.
RAM	128GB: 16GB × 4 DIMMS × 2 NUMA nodes at 2666 MHz
NIC	Broadcom Ethernet NIC controller P225p 2x25G PCIe Gen3 x8
Operating System	Red Hat Enterprise Linux release 9.2
Kernel Version	5.14.0-284.25.1.el9_2.x86_64
Broadcom Firmware Version	227.0.127.0
Test Topology	Two NICs each using one port. Each port receives a stream of 256 IP flows per port from IXIA. Frames are transmitted and received on the first port of each of the NICs. Two queue pairs (one per port) are assigned to a single core.

Figure 1: Topology of P225p 2x 25G Zero Frame Loss Test





Table 2: Configuration of P225p 2x 25G Zero Frame Loss Test

Item	Description
BIOS	Set System Profile to Performance Disable Hyperthreading
Boot Settings	isolcpus=3-13 nohz_full=3-13 rcu_nocbs=3-13 default_hugepagesz=1G hugepagesz=1G hugepages=64 intel_idle.max_cstate=0 processor.max_cstate=0 intel_pstate=disable rcu_nocb_poll audit=0 nosoftlockup intel_iommu=off mce=ignore_ce
Other Optimizations	Disable Linux realtime throttling: echo -1 > /proc/sys/kernel/ sched_rt_runtime_us
Command Line	chrt -r 1 ./dpdk-testpmd -l 3,5 --main-lcore 3 -n4 -- --socket-num=1 --txq=1 --rxq=1 --rxd=512 --txd=512 --nb-cores=1 -i --burst=64 -a

Table 3: Results of P225p 2x 25G Zero Frame Loss Test

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		1-Core
64	74.40	68.18
128	42.23	42.23
256	22.64	22.64
512	11.75	11.75
1,024	5.99	5.99
1,280	4.81	4.81
1,518	4.06	4.06

Figure 2: Results of P225p 2x 25G Zero Frame Loss Test

2 P425G 4x 25G Zero Frame Loss – PCIe 4.0

RFC2544 Zero Frame Loss Performance on Broadcom Ethernet NIC P425G with 4x 25G – PCIe 4.0.

Table 4: Setup of P425G 4x 25G Zero Frame Loss Test – PCIe 4.0

Item	Description
Server	Dell PowerEdge R7515
CPU	Single-Socket AMD EPYC 7713 64-Core Processor at 2.5 GHz (running at 3.1 GHz)
RAM	128GB: 16GB x 8 DIMMs at 3200 MHz
NIC	Broadcom Ethernet NIC controller P425G 4x25G, PCIe Gen4 x16
Operating System	Red Hat Enterprise Linux release 9.2
Kernel Version	5.14.0-284.25.1.el9_2.x86_64
Broadcom Firmware Version	227.0.131.0
Test Topology	One NIC and four ports are used. Each port receives a stream of 256 IP flows from IXIA. Frames that are received on Port 1 are transmitted on Port 2 (and vice versa) by testpmd. Frames that are received on Port 3 are transmitted on Port 4 (and vice versa) by testpmd. Eight queue pairs for four ports (two pairs per port) are assigned to two cores.

Figure 3: Topology of P425G 4x 25G Zero Frame Loss Test – PCIe 4.0





Table 5: Configuration of P425G 4x 25G Zero Frame Loss Test – PCIe 4.0

Item	Description
BIOS	NPS=1 L3 Cache as NUMA: Disabled IOMMU: Enabled, x2APIC Mode: Enabled Determinism Slider: Performance Determinism MADT Core Enumeration: Linear, Logical Processor: Disabled Algorithm Performance Boost Disable (ApbDis): Enabled, ApbDis Fixed Socket P-State: P0 Root Complex LCLK Frequency: 593 MHz Set Preferred IO: Enable, Enhanced Preferred I/O: Enabled
NVM Settings	bnxtnvm --dev=p3p1 setoption=pcie_relaxed_ordering#0x01
Boot Settings	isolcpus=1-36 rcu_nocbs=1-36 amd_iommu=on iommu=pt default_hugepagesz=1G hugepagesz=1G hugepages=64 nohz=off selinux=0 processor.max_cstate=0 nosoftlockup rcu_nocb_poll audit=0 mce=ignore_ce
Other Optimizations	Disable Linux realtime throttling: echo -1 > /proc/sys/kernel/sched_rt_runtime_us
Command Line	chrt -r 1 ./dpdk-testpmd -l 4,5,6 --main-lcore 4 -n 4 --socket-mem=4096 -- -i --nb-cores=2 --nb-ports=4 --rxq=2 --txq=2 --rxd=1024 --txd=1024 --burst=64 -a testpmd> set portlist 0,1,2,3

Table 6: Results of P425G 4x 25G Zero Frame Loss Test – PCIe 4.0

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		2-Core
64	148.81	106.53
128	84.45	84.45
256	45.29	45.29
512	23.50	23.50
1,024	11.97	11.97
1,280	9.62	9.62
1,518	8.13	8.13

Figure 4: Results of P425G 4x 25G Zero Frame Loss Test – PCIe 4.0



3 P2100G 1x 100G Zero Frame Loss – PCIe 4.0

RFC2544 Zero Frame Loss Performance on Broadcom Ethernet NIC P2100G with 1x 100G traffic – PCIe4.0.

Table 7: Setup of P2100G 1x 100G Zero Frame Loss Test – PCIe 4.0

Item	Description
Server	Dell PowerEdge R7515
CPU	Single-Socket AMD EPYC 7713 64-Core Processor at 2.5 GHz (running at 3.1 GHz)
RAM	128GB: 16GB x 8 DIMMs at 3200 MHz
NIC	Broadcom Ethernet NIC controller P2100G 2x100G, PCIe Gen4 x16
Operating System	Red Hat Enterprise Linux release 9.2
Kernel Version	5.14.0-284.25.1.el9_2.x86_64
Broadcom Firmware Version	227.0.131.0
Test Topology	One NIC and two ports are used. Port 1 receives a stream of 8,192 IP flows from IXIA. Frames received on port 1 are forwarded to port 2 by testpmd. Four queue pairs for two ports (two pairs per port) are assigned to two cores.

Figure 5: Topology of P2100G 1x 100G Zero Frame Loss Test – PCIe 4.0





Table 8: Configuration of P2100G 1x 100G Zero Frame Loss Test – PCIe 4.0

Item	Description
BIOS	NPS=1, L3 Cache as NUMA: Disabled IOMMU: Enabled, x2APIC Mode: Enabled Determinism Slider: Performance Determinism MADT Core Enumeration: Linear, Logical Processor: Disabled Algorithm Performance Boost Disable (ApbDis): Enabled, ApbDis Fixed Socket P-State: P0 Root Complex LCLK Frequency: 593 MHz Set Preferred IO: Enable, Enhanced Preferred I/O: Enabled
NVM Settings	bnxtnvm --dev=ens4f0np0 setoption=pcie_relaxed_ordering#0x01
Boot Settings	isolcpus=1-36 rcu_nocbs=1-36 default_hugepagesz=1G hugepagesz=1G hugepages=64 nosoftlockup processor.max_cstate=0 amd_iommu=on iommu=pt numa_balancing=disable selinux=0 nohz=off rcu_nocb_poll audit=0 mce=ignore_ce
Other Optimizations	Disable Linux realtime throttling: echo -1 > /proc/sys/kernel/sched_rt_runtime_us
Command Line	chrt -r 1 ./dpdk-testpmd -l 4,5,6 --main-lcore 4 -n 4 -- -i --nb-cores=2 --nb-ports=2 --rxq=2 --txq=2 --rxd=2048 --txd=2048 --socket-num=0 --burst=64 -a

Table 9: Results of P2100G 1x 100G Zero Frame Loss Test – PCIe 4.0

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		2-Core
64	148.81	102.15
128	84.45	84.45
256	45.29	45.29
512	23.50	23.50
1,024	11.97	11.97
1,280	9.62	9.62
1,518	8.13	8.13

Figure 6: Results of P2100G 1x 100G Zero Frame Loss Test – PCIe 4.0



4 P2100G 2x 100G Zero Frame Loss – PCIe 4.0

RFC2544 Zero Frame Loss Performance on Broadcom Ethernet NIC P2100G with 2x 100G traffic.

Table 10: Setup of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0

Item	Description
Server	Dell PowerEdge R7515
CPU	Single-Socket AMD EPYC 7713 64-Core Processor at 2.5 GHz (running at 3.1 GHz)
RAM	128GB: 16GB x 8 DIMMs at 3200 MHz
NIC	Broadcom Ethernet NIC controller P2100G 2x100G, PCIe Gen4 x16
Operating System	Red Hat Enterprise Linux release 9.2
Kernel Version	5.14.0-284.25.1.el9_2.x86_64
Broadcom Firmware Version	227.0.131.0
Test Topology	One NIC and two ports are used. Each port receives a stream of 8,192 IP flows from IXIA. Frames received on one port are forwarded to the other port by testpmd. Four queue pairs for two ports (two pairs per port) are assigned to two cores.

Figure 7: Topology of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0

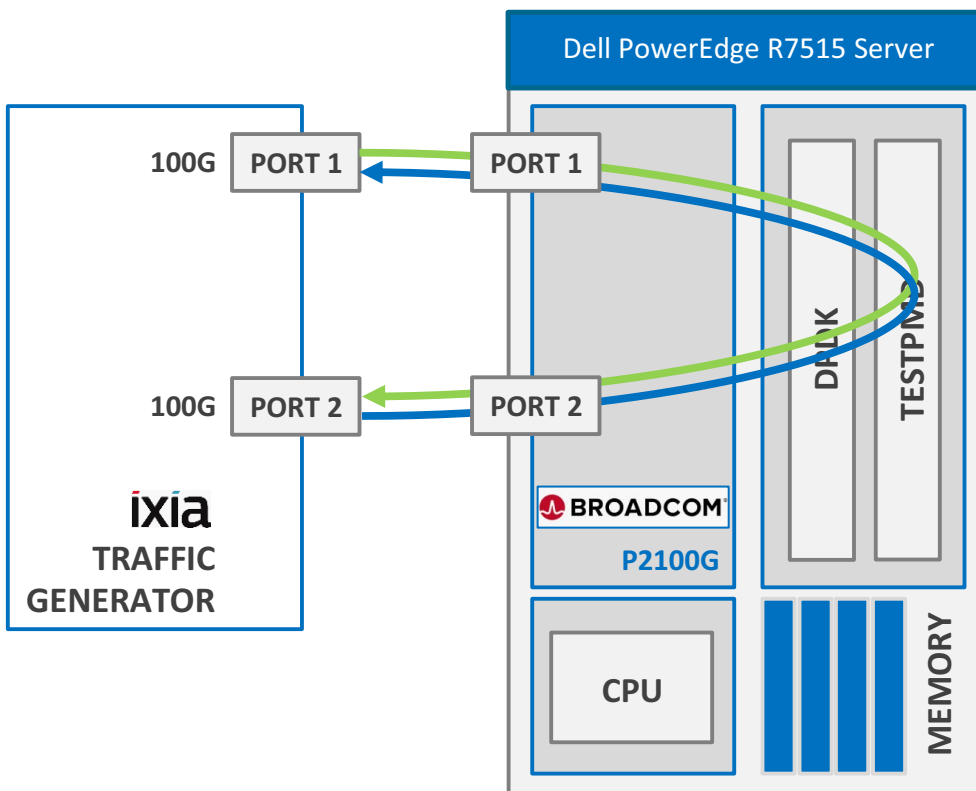



Table 11: Configuration of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0

Item	Description
BIOS	NPS=1, L3 Cache as NUMA: Disabled IOMMU: Enabled, x2APIC Mode: Enabled Determinism Slider: Performance Determinism MADT Core Enumeration: Linear, Logical Processor: Disabled Algorithm Performance Boost Disable (ApbDis): Enabled, ApbDis Fixed Socket P-State: P0 Root Complex LCLK Frequency: 593 MHz Set Preferred IO: Enable, Enhanced Preferred I/O: Enabled
NVM Settings	bnxtnvm --dev=ens4f0np0 setoption=pcie_relaxed_ordering#0x01
Boot Settings	isolcpus=1-36 rcu_nocbs=1-36 default_hugepagesz=1G hugepagesz=1G hugepages=64 nosoftlockup processor.max_cstate=0 amd_iommu=on iommu=pt numa_balancing=disable selinux=0 nohz=off rcu_nocb_poll audit=0 mce=ignore_ce
Other Optimizations	Disable Linux realtime throttling: echo -1 > /proc/sys/kernel/sched_rt_runtime_us
Command Line	chrt -r 1 ./dpdk-testpmd -l 4,5,6 --main-lcore 4 -n 4 -- -i --nb-cores=2 --nb-ports=2 --rxq=2 --txq=2 --rxd=2048 --txd=2048 --socket-num=0 --burst=64 -a

Table 12: Results of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		2-Core
64	297.62	106.79
128	168.92	106.96
256	90.58	82.73
512	46.99	46.99
1,024	23.94	23.94
1,280	19.23	19.23
1,518	16.25	16.25

Figure 8: Results of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0



5 P2100G 2x 100G Zero Frame Loss (Icelake) – PCIe 4.0

RFC2544 Zero Frame Loss Performance on Broadcom Ethernet NIC P2100G with 2x 100G traffic.

Table 13: Setup of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0

Item	Description
Server	Dell PowerEdge R750
CPU	Intel Xeon Platinum 8358 CPU at 2.60 GHz (running at 3.3 GHz)
RAM	128GB: 16GB x 8 DIMMs at 3200 MHz
NIC	Broadcom Ethernet NIC controller P2100G 2x100G, PCIe Gen4 x16
Operating System	Red Hat Enterprise Linux release 9.2
Kernel Version	5.14.0-284.25.1.el9_2.x86_64
Broadcom Firmware Version	227.0.131.0
Test Topology	One NIC and two ports are used. Each port receives a stream of 8,192 IP flows from IXIA. Frames received on one port are forwarded to the other port by testpmd. Two/four queue pairs for two ports (one/two pairs per port) are assigned to one/two cores.

Figure 9: Topology of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0

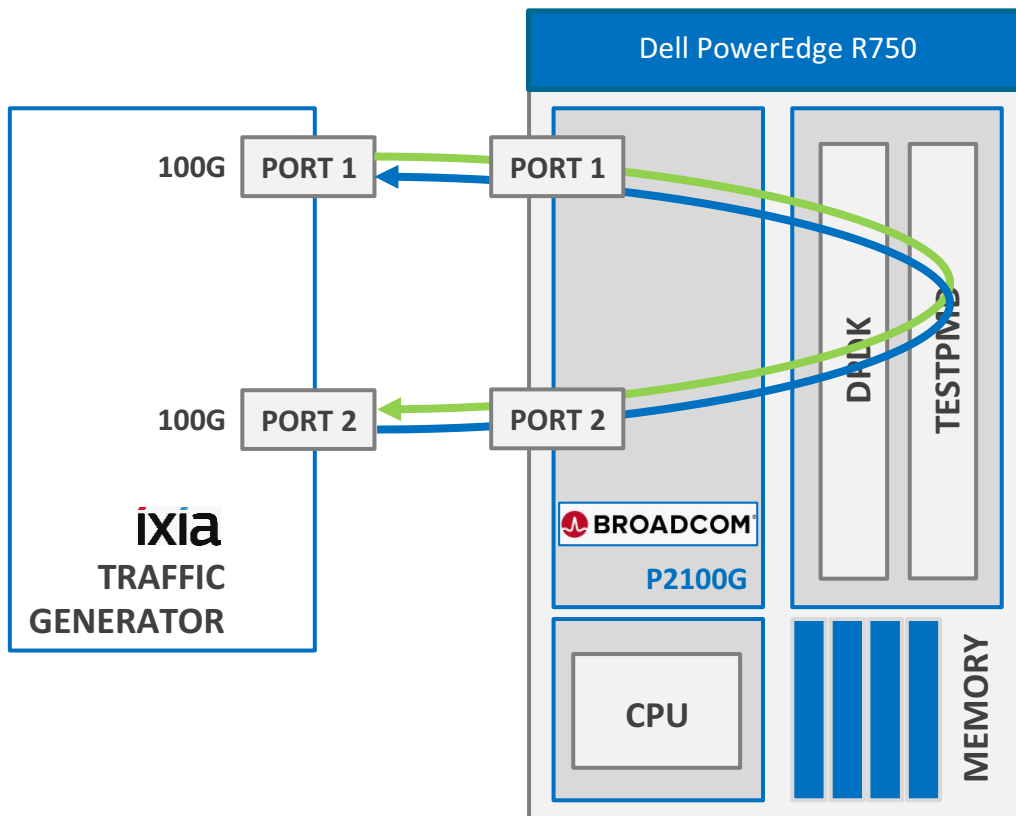



Table 14: Configuration of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0

Item	Description
BIOS	Logical Processor: Disabled Sub NUMA Cluster: 2-Way Clustering MADT Core Enumeration: Round Robin I/O Snoop Hold Off Response: 2K Cycles System Profile: Performance
NVM Settings	-
Boot Settings	default_hugepagesz=1G hugepagesz=1G hugepages=64 intel_iommu=on iommu=pt intel_idle.max_cstate=0 processor.max_cstate=0 intel_pstate=disable rcu_nocb_poll audit=0 nosoftlockup mce=ignore_ce isolcpus=4,6,7,8,10,11,12,14,15,16,18,19,20,22,23,24,26,27,28,30,31,32,34,35, 36,38,39,40,44,48,52,56,58,60,62 nohz_full=4,6,7,8,10,11,12,14,15,16,18,19,20,22,23,24,26,27,28,30,31,32,34,35 ,36,38,39,40,44,48,52,56,58,60,62 rcu_nocbs=4,6,7,8,10,11,12,14,15,16,18,19,20,22,23,24,26,27,28,30,31,32,34,35 ,36,38,39,40,44,48,52,56,58,60,62
Other Optimizations	Disable Linux realtime throttling: echo -1 > / proc/sys/kernel/sched_rt_runtime_us
Command Line	chrt -r 1 ./dpdk-testpmd -l 4,8 --main-lcore 4 -n 4 -a 17:00.0 -a 17:00.1 -- - -i --nb-cores=1 --nb-ports=2 --rxq=1 --txq=1 --rxd=2048 --txd=2048 --socket- num=0 --burst=64 --mbcache=512 --record-core-cycles --record-burst-stats -a chrt -r 1 ./dpdk-testpmd -l 4,8,12 --main-lcore 4 -n 4 -a 17:00.0 -a 17:00.1 - - -i --nb-cores=2 --nb-ports=2 --rxq=2 --txq=2 --rxd=2048 --txd=2048 --socket- num=0 --burst=64 --mbcache=512 --record-core-cycles --record-burst-stats -a

Table 15: Results of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		1-Core
64	297.62	105.13
128	168.92	102.60
256	90.58	84.11
512	46.99	46.99
1,024	23.94	23.94
1,280	19.23	19.23
1,518	16.25	16.25

Figure 10: Results of P2100G 2x 100G Zero Frame Loss Test – PCIe 4.0



6 P2100G 1x 200G Zero Frame Loss (Icelake) – PCIe 4.0

RFC2544 Zero Frame Loss Performance on Broadcom Ethernet NIC P2100G with 1x 200G traffic.

Table 16: Setup of P2100G 1x 200G Zero Frame Loss Test – PCIe 4.0

Item	Description
Server	Dell PowerEdge R750
CPU	Intel Xeon Platinum 8358 CPU at 2.60GHz (running at 3.3 GHz)
RAM	128GB: 16GB x 8 DIMMs at 3200 MHz
NIC	Broadcom Ethernet NIC controller P2100G 1x200G, PCIe Gen4 x16
Operating System	Red Hat Enterprise Linux release 9.2
Kernel Version	5.14.0-284.25.1.el9_2.x86_64
Broadcom Firmware Version	227.0.131.0
Test Topology	One NIC and one port are used. Port receives a stream of 8,192 IP flows from IXIA. Frames received on one port are transmitted back on the same port by testpmd. Two queue pairs a port assigned to two cores.

Figure 11: Topology of P2100G 1x 200G Zero Frame Loss Test – PCIe 4.0

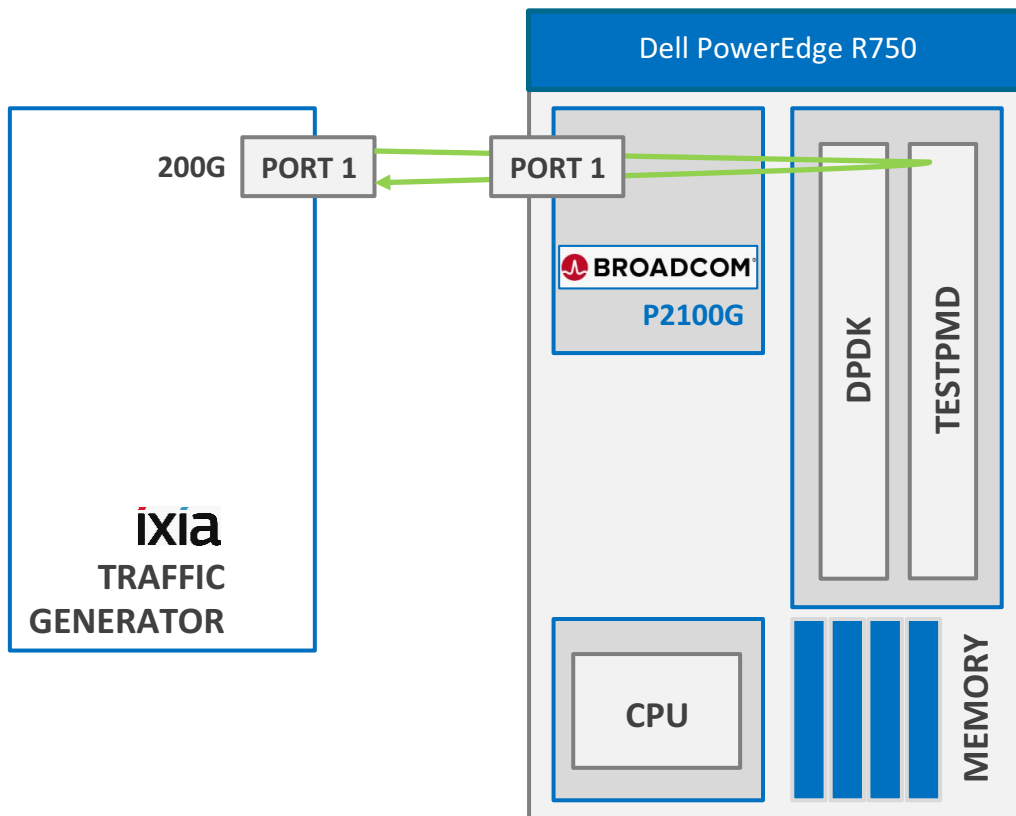



Table 17: Configuration of P2100G 1x 200G Zero Frame Loss Test – PCIe 4.0

Item	Description
BIOS	Logical Processor: Disabled Sub NUMA Cluster: 2-Way Clustering MADT Core Enumeration: Round Robin I/O Snoop HoldOff Response: 2K Cycles System Profile: Performance
NVM Settings	bnxtnvm -dev=ens2f0np0 setoption=firmware_link_speed_d0:0#7 bnxtnvm -dev=ens2f0np0 setoption=firmware_link_speed_d3:0#7 bnxtnvm -dev=ens2f0np0 setoption=fwd_err_correct:0#16 bnxtnvm -dev=ens2f0np0 setoption=link_training:0#1 bnxtnvm -dev=ens2f0np0 setoption=port_hide#1
Boot Settings	default_hugepagesz=1G hugepagesz=1G hugepages=64 intel_iommu=on iommu=pt intel_idle.max_cstate=0 processor.max_cstate=0 intel_pstate=disable rcu_nocb_poll audit=0 nosoftlockup mce=ignore_ce isolcpus=4,6,7,8,10,11,12,14,15,16,18,19,20,22,23,24,26,27,28,30,31,32,34,35, 36,38,39,40,44,48,52,56,58,60,62 nohz_full=4,6,7,8,10,11,12,14,15,16,18,19,20,22,23,24,26,27,28,30,31,32,34,35, 36,38,39,40,44,48,52,56,58,60,62 rcu_nocbs=4,6,7,8,10,11,12,14,15,16,18,19,20,22,23,24,26,27,28,30,31,32,34,35, 36,38,39,40,44,48,52,56,58,60,62
Other Optimizations	Disable Linux realtime throttling: echo -1 > /proc/sys/kernel/sched_rt_runtime_us
Command Line	chrt -r 1 ./ dpdk-testpmd -l 6,10,14 --main-lcore 6 -n 4 -a 4b:00.0 -- -i --nb- cores=2 --nb-ports=1 --rxq=2 --txq=2 --rxd=2048 --txd=2048 --socket-num=1 -- burst=64 --mbcache=512 --record-core-cycles --record-burst-stats --port- topology=chained -a testpmd> set fwd macswap

Table 18: Results of P2100G 1x 200G Zero Frame Loss Test – PCIe 4.0

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		2-Core
64	297.62	101.36
128	168.92	100.04
256	90.58	83.52
512	46.99	46.99
1,024	23.94	23.94
1,280	19.23	19.23
1,518	16.25	16.25

Figure 12: Results of P2100G 1x 200G Zero Frame Loss Test – PCIe 4.0



Revision History

EtherNIC-Ctrl-DPDK-23-07-Performance-TR100; September 7, 2023

Initial release.

