



NetXtreme-E/NetXtreme-S DPDK 20.08 Performance Report

Test Report

Broadcom, the pulse logo, NetXtreme, Connecting everything, Avago Technologies, Avago, and the A logo are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries, and/or the EU.

Copyright © 2020 Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com.

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 PS225 2x 25G Bi-directional Zero Frame Loss – A72 IO-FWD	4
2 P225p 2x 25G Bi-directional Zero Frame Loss	6
3 P425G 4x 25G Bi-directional Zero Frame Loss	8
4 P2100G 1x 100G Bi-directional Zero Frame Loss	10
5 P425G 4x 25G Bi-directional Zero Frame Loss – PCIe 4.0	12
6 P2100G 1x 100G Bi-directional Zero Frame Loss – PCIe 4.0	14
7 P2100G 2x 100G Bi-directional Zero Frame Loss – PCIe 4.0	16
Revision History	18

1 PS225 2x 25G Bi-directional Zero Frame Loss – A72 IO-FWD

RFC2544 Zero Frame Loss Test on Broadcom NetXtreme-S Series PS225 with 2x 25G Bi-directional traffic – A72 IO-FWD.

Table 1: Setup of PS225 2x 25G Bi-directional Zero Frame Loss – A72 IO-FWD

Item	Description
Server	Dell PowerEdge R730
NIC	Eight 64-bit ARM Cortex-A72 cores at 3.0 GHz
SNIC CPU	16 GB memory on-board at 2400 MHz
SNIC RAM	Broadcom PS225-H16 NetXtreme-S Dual-port 25Gb PCIe SmartNIC
Operating System	Broadcom Yocto Linux
Kernel Version	4.14.174
Broadcom Firmware Version	218.1.63.0
Test Topology	One NIC and two ports are used Each port receives a stream of 256 IP flows from IXIA Frames received on one port is forwarded to the other port by testpmd Two TX/RX queue pairs (one pair per port) are assigned to a single A72 CPU core Data points are collected using one, two, and four A72 CPU cores

Figure 1: Topology of PS225 2x 25G Bi-directional Zero Frame Loss – A72 IO-FWD

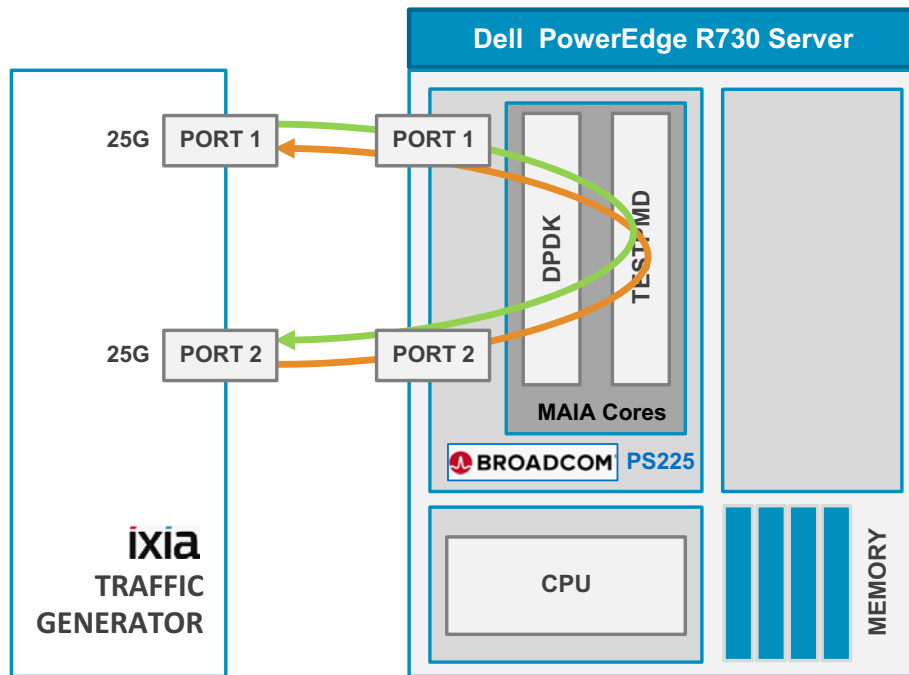


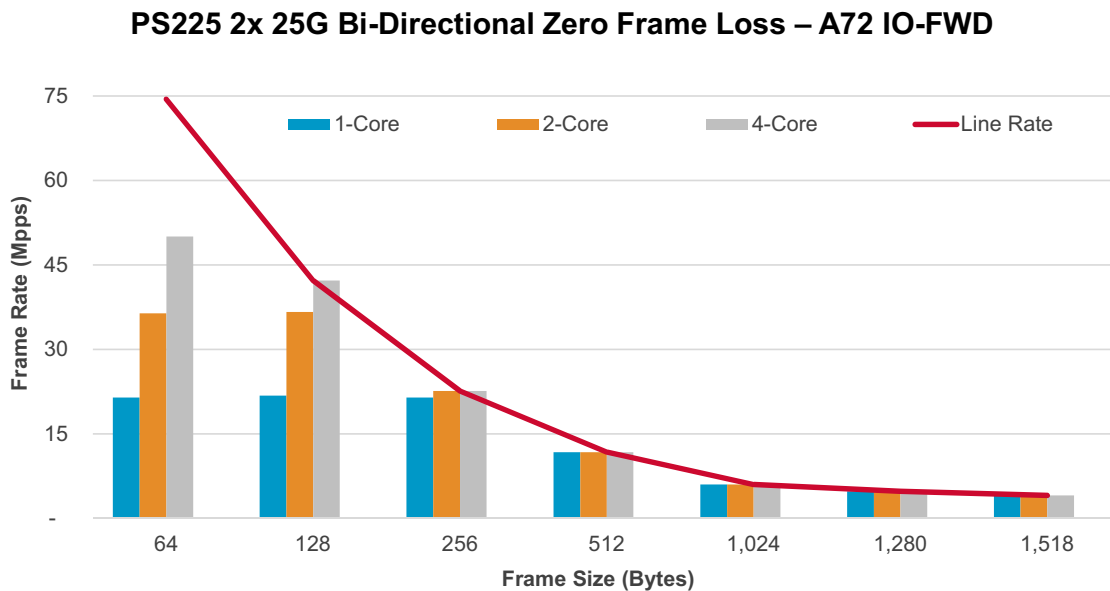
Table 2: Configuration of PS225 2x 25G Bi-directional Zero Frame Loss – A72 IO-FWD

Item	Description
BIOS	Set System Profile to Performance. Disable Virtualization and SR-IOV.
Boot Setting for A72	Execute the following command under boot shell: set extraarg "isolcpus=1,2,3,4,5,6,7 nohz_full=1-7 rcu_nocbs=1-7" nvsync
Command Line	<pre>./testpmd -l 1,2 --master-lcore 1 -w 0008:01:00.0 -w 0008:01:00.1 -n4 -- -- txq=1 --rxq=1 --rxd=256 --txd=256 --nb-cores=1 -i -a ./testpmd -l 1,2,3 --master-lcore 1 -w 0008:01:00.0 -w 0008:01:00.1 -n4 -- -- txq=2 --rxq=2 --rxd=256 --txd=256 --nb-cores=2 -i -a ./testpmd -l 1,2,3,4,5 --master-lcore 1 -w 0008:01:00.0 -w 0008:01:00.1 -n4 - --txq=4 --rxq=4 --rxd=256 --txd=256 --nb-cores=4 -i -a</pre>

Table 3: Results of PS225 2x 25G Bi-directional Zero Frame Loss – A72 IO-FWD

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)		
		1-core (A72 Core)	2-core (A72 Cores)	4-core (A72 Cores)
64	74.40	21.45	36.40	50.00
128	42.23	21.82	36.67	42.23
256	22.64	21.46	22.64	22.64
512	11.75	11.75	11.75	11.75
1,024	5.99	5.99	5.99	5.99
1,280	4.81	4.81	4.81	4.81
1,518	4.06	4.06	4.06	4.06

Figure 2: Results of PS225 with 2x 25G Bi-directional Traffic – A72 IO-FWD



2 P225p 2x 25G Bi-directional Zero Frame Loss

RFC2544 Zero Frame Loss Test on Broadcom NetXtreme-E Series P225p with 2x 25G Bi-directional traffic.

Table 4: Setup of P225p 2x 25G Bi-directional 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	128 GB:16 GB × eight DIMMS × two NUMA nodes at 2666 MHz
NIC	Broadcom NetXtreme-E Series P225p 2x 25G PCIe Gen3 x8
Operating System	Red Hat Enterprise Linux Server release 7.6
Kernel Version	3.10.0-957.el7.x86_64
Broadcom Firmware Version	214.4.114.0
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 TX/RX queue pairs (one per port) are assigned to a single physical CPU core Data points are taken with one and two physical CPU cores

Figure 3: Topology of P225p 2x 25G Bi-directional Zero Frame Loss Test

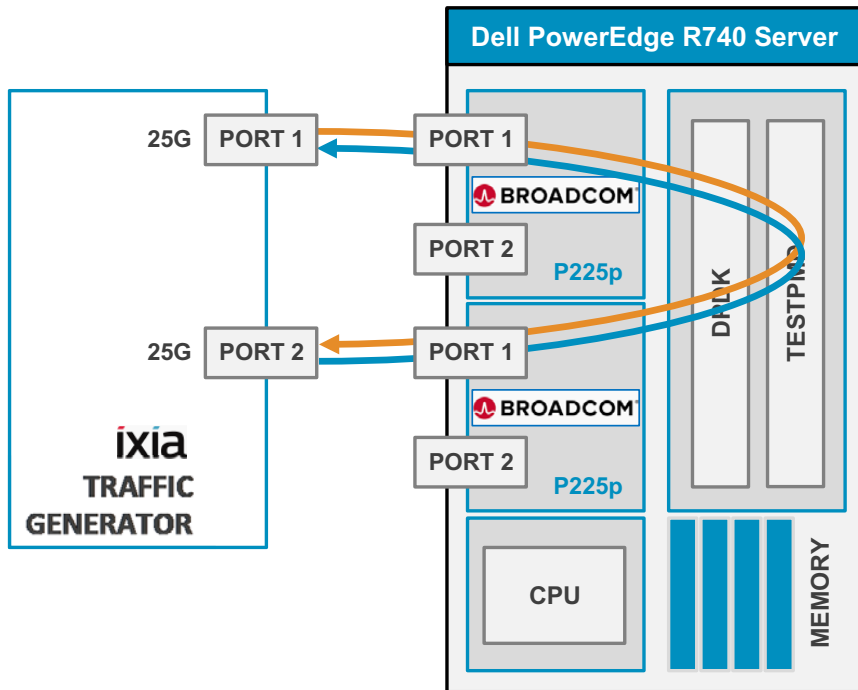



Table 5: Configuration of P225p 2x 25G Bi-directional Zero Frame Loss Test

Item	Description
BIOS	Set System Profile to Performance Disable Virtualization and SR-IOV Disable Hyper Threading (Logical Cores)
Boot Settings	isolcpus=3-7 nohz_full=3-7 rcu_nocbs=3-7 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 ./testpmd -l 3,5 --master-lcore 3 -n4 --socket-num=1 --txq=1 --rxq=1 --rxd=512 --txd=512 --nb-cores=1 -i --burst=32 -a chrt -r 1 ./testpmd -l 3,5,7 --master-lcore 3 -n4 --socket-num=1 --txq=2 --rxq=2 --rxd=512 --txd=512 --nb-cores=2 -i --burst=32 -a

Table 6: Results of P225p 2x 25G Bi-directional Zero Frame Loss Test

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)	
		1-Core (Physical Cores)	2-Core (Physical Cores)
64	74.40	42.55	60.00
128	42.23	42.23	42.23
256	22.64	22.64	22.64
512	11.75	11.75	11.75
1,024	5.99	5.99	5.99
1,280	4.81	4.81	4.81
1,518	4.06	4.06	4.06

Figure 4: Results of P225p 2x 25G Bi-directional Zero Frame Loss Test



3 P425G 4x 25G Bi-directional Zero Frame Loss

RFC2544 Zero Frame Loss Test on Broadcom NetXtreme-E Series P425G with 4x 25G Bi-directional traffic.

Table 7: Setup of P425G 4x 25G Bi-directional 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	128 GB:16 GB × eight DIMMS × two NUMA nodes at 2666 MHz
NIC	Broadcom NetXtreme-E Series P425G 4x 25G, PCIe Gen3/4 x16 (Gen3 was used in the test)
Operating System	Red Hat Enterprise Linux Server release 7.6
Kernel Version	3.10.0-957.el7.x86_64
Broadcom Firmware Version	216.4.74.0
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 (vice versa) by testpmd Frames that are received on Port 3 are transmitted on Port 4 (vice versa) by testpmd Eight TX/RX queue pairs for four ports (two per port) are assigned to four physical CPU cores Data points are taken with four physical CPU cores

Figure 5: Topology of P425G 4x 25G Bi-directional Zero Frame Loss Test

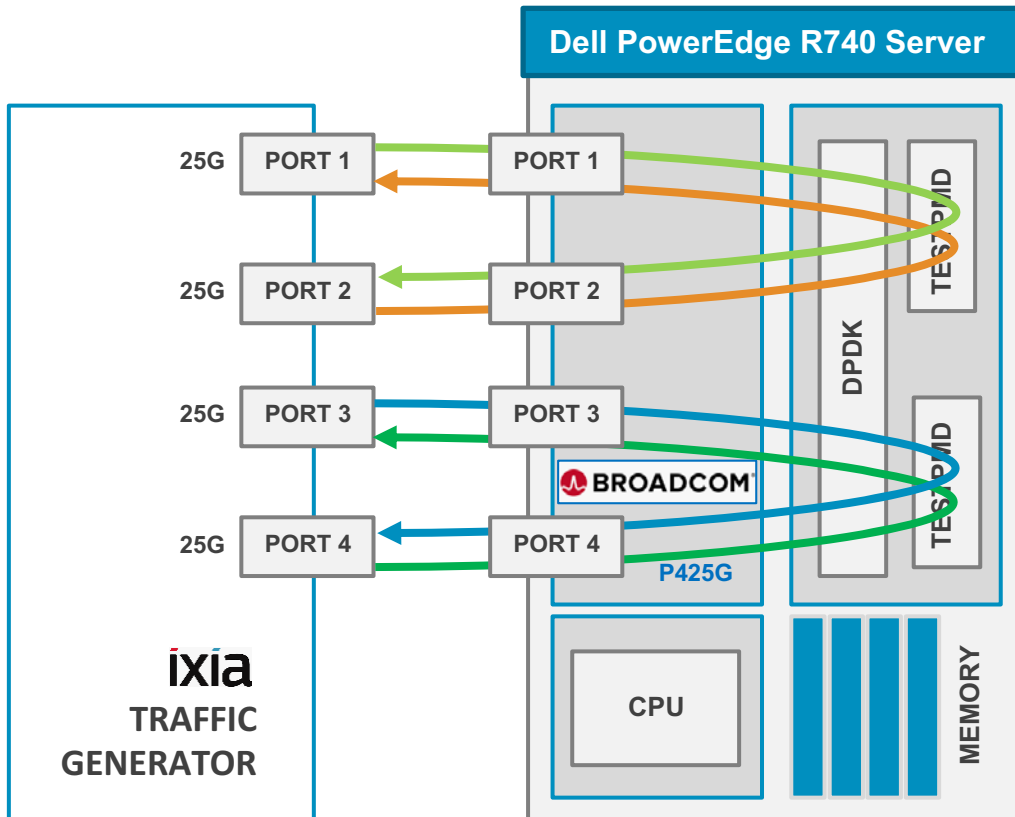


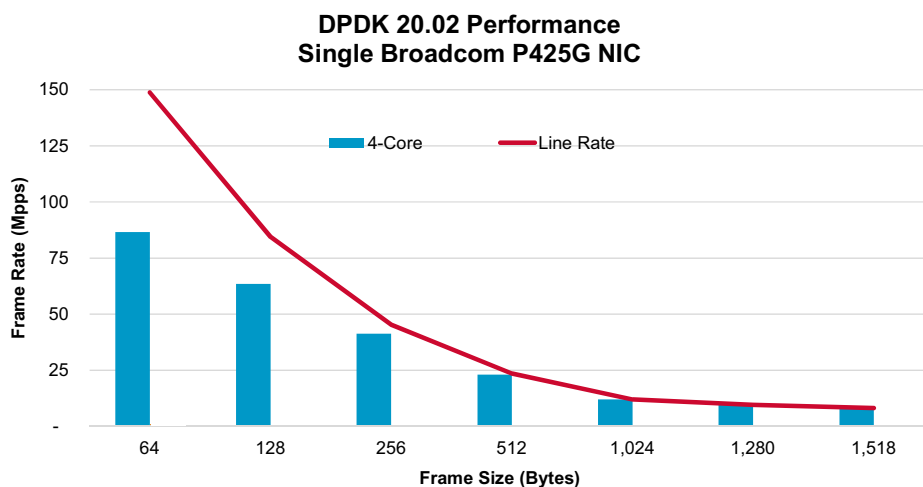
Table 8: Configuration of P425G 4x 25G Bi-directional Zero Frame Loss Test

Item	Description
BIOS	Set System Profile to Performance. Disable Virtualization and SR-IOV. Disable HyperThreading (Logical Cores).
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 ./testpmd -l 3,5,7 --master-lcore 3 -n4 -w 0000:d8:00.0 -w 0000:d8:00.1 -- --socket-num=1 --txq=2 --rxq=2 --rxd=2048 --txd=2048 --nb- cores=2 -i -a chrt -r 1 ./testpmd -l 9,11,13 --master-lcore 9 -n4 -w 0000:d8:00.2 -w 0000:d8:00.3 --file-prefix pg2 -- --socket-num=1 --txq=2 --rxq=2 --rxd=2048 - -txd=2048 --nb-cores=2 -i -a

Table 9: Results of P425G 4x 25G Bi-directional Zero Frame Loss Test

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		4-Core (Physical Cores)
64	148.81	87.49
128	84.46	63.41
256	45.29	41.17
512	23.50	23.00
1,024	11.97	11.97
1,280	9.62	9.62
1,518	8.13	8.13

Figure 6: Results of P425G 4x 25G Bi-directional Zero Frame Loss Test



4 P2100G 1x 100G Bi-directional Zero Frame Loss

RFC2544 Zero Frame Loss Test on Broadcom NetXtreme-E Series P2100G with 1x 100G Bi-directional traffic.

Table 10: Setup of P2100G 1x 100G Bi-directional 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	128 GB:16 GB × eight DIMMS × two NUMA nodes at 2666 MHz
NIC	Broadcom NetXtreme-E Series P2100G 2x 100G, PCIe Gen3/4 x16 (Gen3 was used in the test)
Operating System	Red Hat Enterprise Linux Server release 7.6
Kernel Version	3.10.0-957.el7.x86_64
Broadcom Firmware Version	216.4.74.0
Test Topology	One NIC and two ports are used Port 1 receives a stream of 8,192 IP flows from IXIA Frames are received on Port 1 and transmitted on Port 2 Two TX/RX queue pairs (one per port) are assigned to a single physical CPU core Data points are taken with one, two, and four physical CPU cores

Figure 7: Topology of P2100G 1x 100G Bi-directional Zero Frame Loss Test

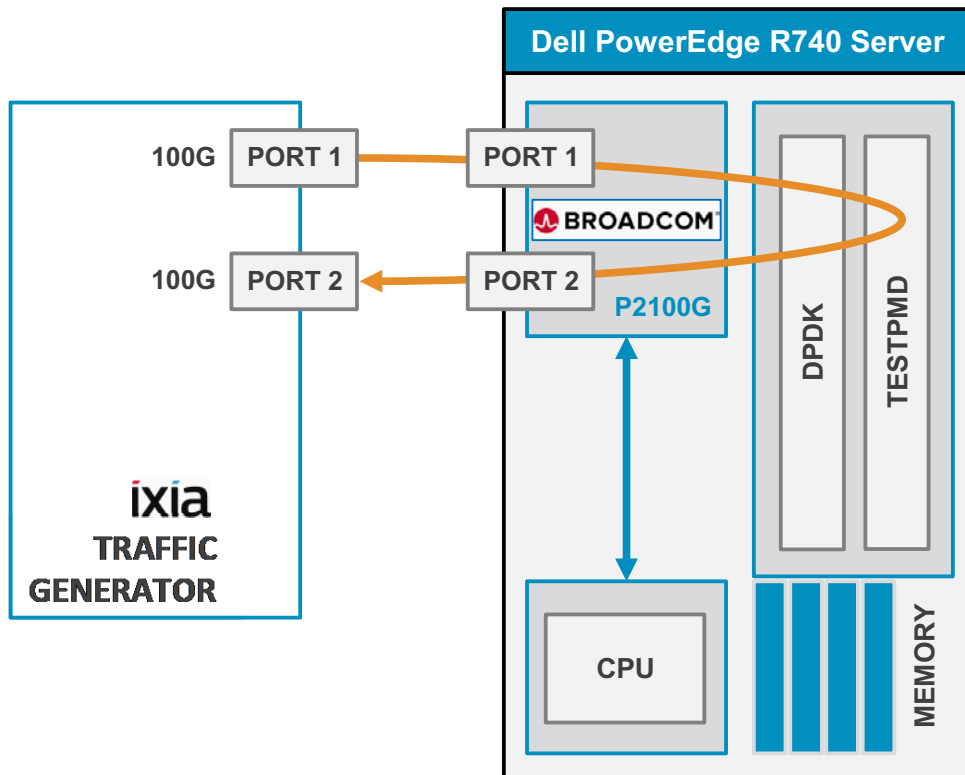


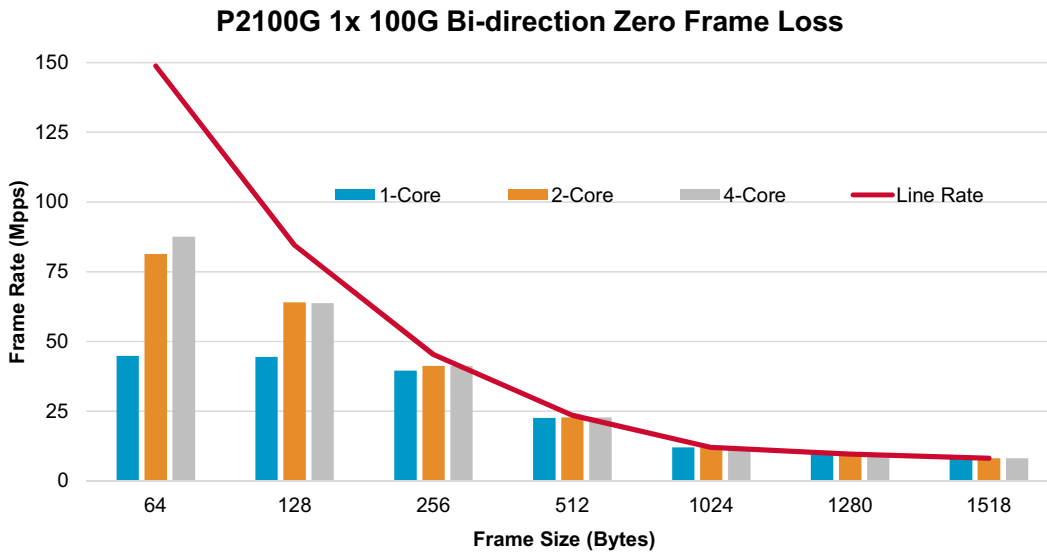
Table 11: Configuration of P2100G 1x 100G Bi-directional Zero Frame Loss Test

Item	Description
BIOS	Set System Profile to Performance Disable Virtualization and SR-IOV Disable HyperThreading (Logical Cores)
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 ./testpmd -l 4,6 --master-lcore 4 -n4 -- --socket-num=0 --txq=1 -- rxq=1 --rxd=2048 --txd=2048 --nb-cores=1 -i --burst=32 -a chrt -r 1 ./testpmd -l 4,6,8 --master-lcore 4 -n4 -- --socket-num=0 --txq=2 -- rxq=2 --rxd=2048 --txd=2048 --nb-cores=2 -i --burst=32 -a chrt -r 1 ./testpmd -l 4,6,8,10,12 --master-lcore 4 -n4 -- --socket-num=0 -- txq=4 --rxq=4 --rxd=2048 --txd=2048 --nb-cores=4 -i --burst=32 -a

Table 12: Results of P2100G 1x 100G Bi-directional Zero Frame Loss Test

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)		
		1-Core (Physical Core)	2-Core (Physical Core)	4-Core (Physical Core)
64	148.81	44.98	82.29	87.95
128	84.46	44.48	64.23	63.85
256	45.29	41.62	41.64	42.41
512	23.50	22.51	23.12	23.08
1,024	11.97	11.97	11.97	11.97
1,280	9.62	9.62	9.62	9.62
1,518	8.13	8.13	8.13	8.13

Figure 8: Results of P2100G 1x 100G Bi-directional Zero Frame Loss Test



5 P425G 4x 25G Bi-directional Zero Frame Loss – PCIe 4.0

RFC2544 Zero Frame Loss Performance on Broadcom NetXtreme-E Series P425G with 4x 25G Bi-directional – PCIe 4.0.

Table 13: Setup of P425G 4x 25G Bi-directional Zero Frame Loss Test – PCIe 4.0

Item	Description
Server	AMD Rome Server Reference Platform
CPU	Single-Socket AMD EPYC 7742 64-Core Processor at 2.25 GHz (running at 3.38 GHz)
RAM	128 GB : 16 GB x 8 DIMMs at 3200 MHz
NIC	Broadcom NetXtreme-E Series P425G 4x 25G, PCIe Gen4 x16
Operating System	Red Hat Enterprise Linux Server release 7.6
Kernel Version	Linux Upstream 5.3.4
Broadcom Firmware Version	216.4.74.0
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 is transmitted on Port 2 (vice versa) by testpmd Frames that are received on Port 3 is transmitted on Port 4 (vice versa) by testpmd Eight TX/RX queue pairs for four ports (two per port) are assigned to eight logical CPU cores Data points are taken with eight logical CPU cores

Figure 9: Topology of P425G 4x 25G Bi-directional Zero Frame Loss Test – PCIe 4.0

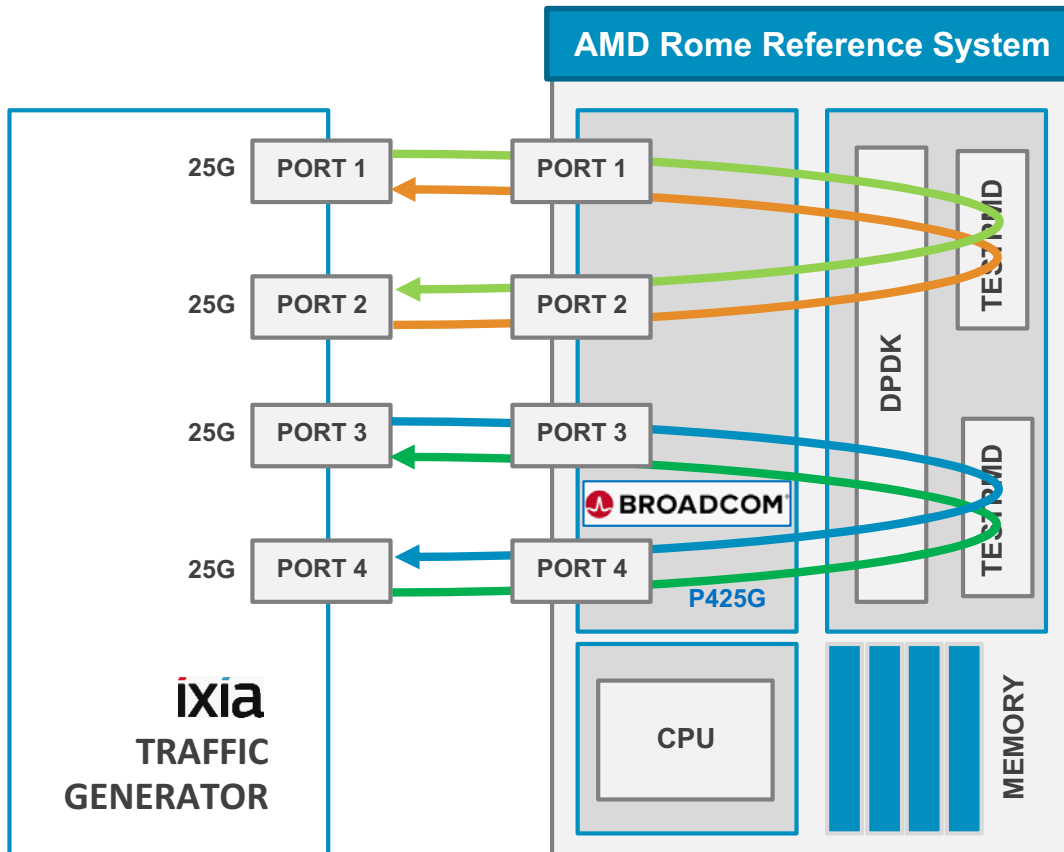


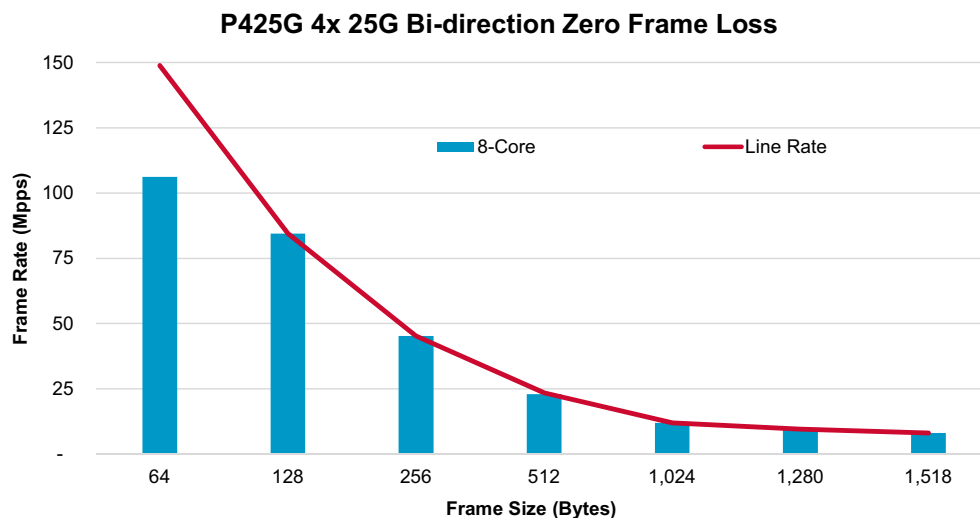
Table 14: Configuration of P425G 4x 25G Bi-directional Zero Frame Loss Test – PCIe 4.0

Item	Description
BIOS	NPS=1, Set Preferred IO: Enable L3 Cache as NUMA: Disabled, PCIE 10 Bit Tag: Enable APIC Mode: X2APIC Deterministic Slider: Performance SMT: Enabled APBDIS-> 1 Enhanced Preferred I/O: Auto (Enabled)
Other Settings	MRRS = 4K (default), PCIe Relaxed Ordering: Enabled
Boot Settings	rcu_nocbs=16-63 isolcpus=16-63 amd_iommu=on iommu=pt default_hugepagesz=1G hugepagesz=1G hugepages=64 nohz=off nosoftlockup selinux=0 numa_balancing=disable processor.max_cstate=0
Command Line	<pre>./testpmd -l 16,17,18,19,24 --master-lcore 24 -n4 -w 0000:41:00.0 -w 0000:41:00.1 -- --txq=2 --rxq=2 --rxd=1024 --txd=1024 --nb-cores=4 -i --burst=32 -a ./testpmd -l 20,21,22,23,25 --master-lcore 25 -n4 -w 0000:41:00.2 -w 0000:41:00.3 --file-prefix pg2 -- --txq=2 --rxq=2 --rxd=1024 --txd=1024 --nb-cores=4 -i --burst=32 -a</pre>

Table 15: Results of P425G 4x 25G Bi-directional Zero Frame Loss Test – PCIe 4.0

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		8-Core (Logical Cores)
64	148.81	106.16
128	84.46	84.46
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 10: Results of P425G 4x 25G Bi-directional Zero Frame Loss Test – PCIe 4.0



6 P2100G 1x 100G Bi-directional Zero Frame Loss – PCIe 4.0

RFC2544 Zero Frame Loss Performance on Broadcom NetXtreme-E Series P2100G with 1x 100G Bi-directional – PCIe4.0.

Table 16: Setup of P2100G 1x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0

Item	Description
Server	AMD Rome Server Reference Platform
CPU	Single-Socket AMD EPYC 7742 64-Core Processor at 2.25 GHz (running at 3.38 GHz)
RAM	128 GB : 16 GB x 8 DIMMs at 3200 MHz
NIC	Broadcom NetXtreme-E Series P2100G 1x 100G, PCIe Gen4 x16
Operating System	Red Hat Enterprise Linux Server release 7.6
Kernel Version	Linux Upstream 5.3.4
Broadcom Firmware Version	216.4.74.0
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 Eight TX/RX queue pairs for two ports (four per port) are assigned to eight logical CPU cores Data points are taken with eight logical CPU cores

Figure 11: Topology of P2100G 1x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0

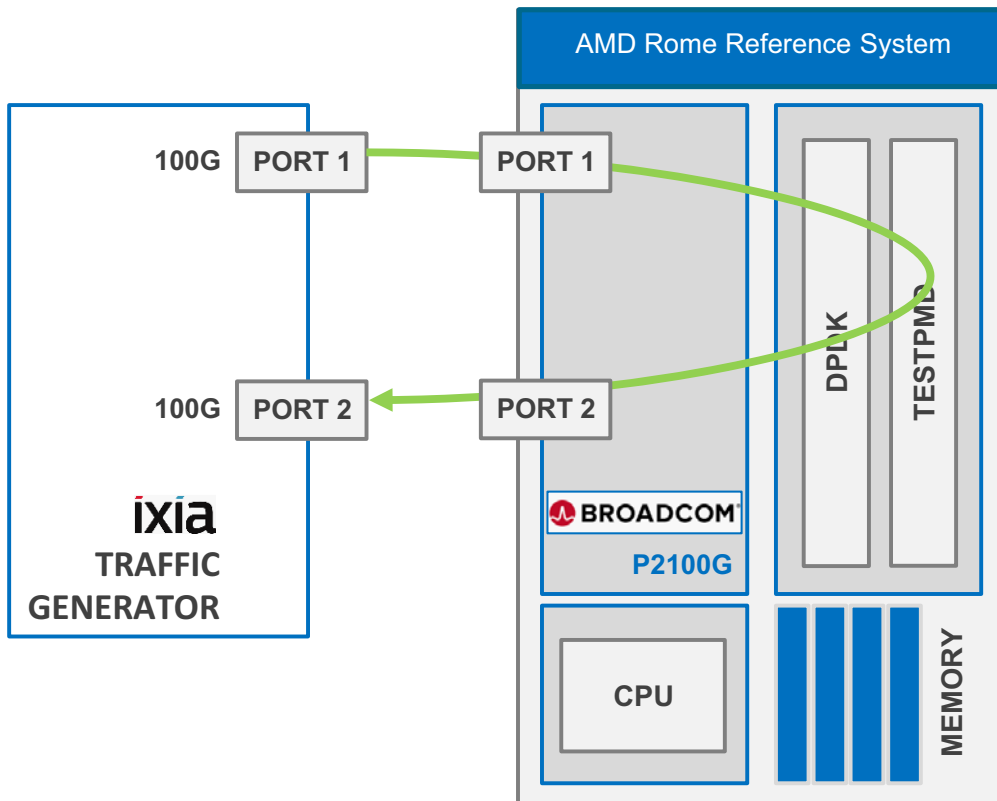


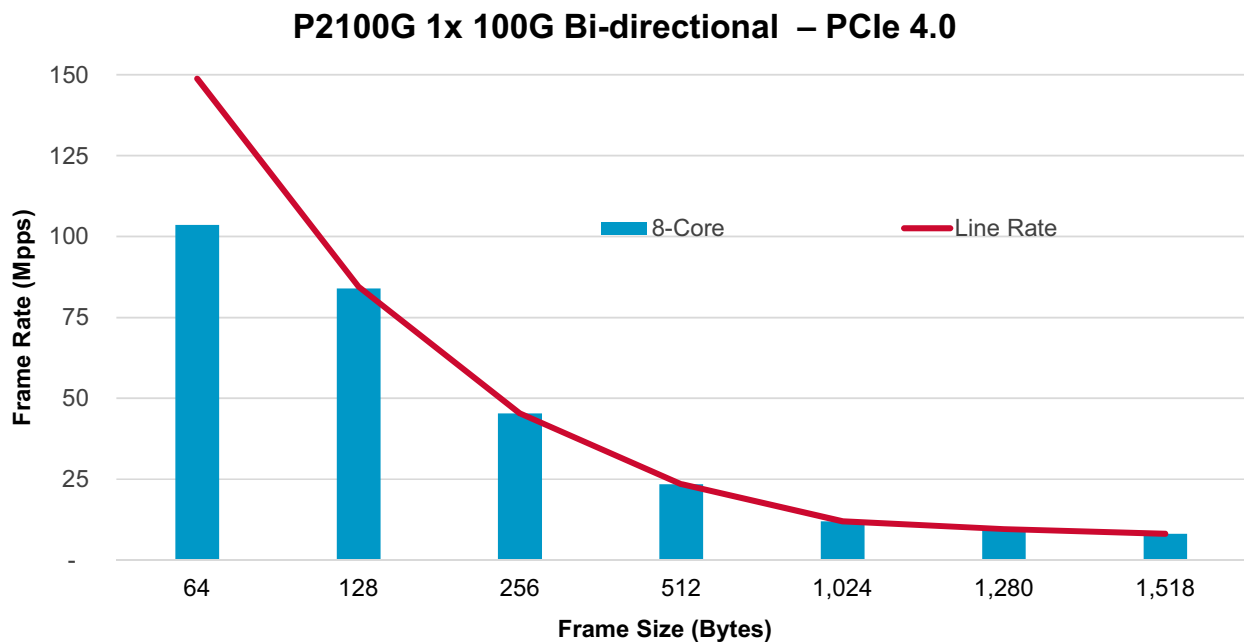
Table 17: Configuration of P2100G 1x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0

Item	Description
BIOS	NPS=1, Set Preferred IO: Enable L3 Cache as NUMA: Disabled, PCIE 10 Bit Tag: Enable APIC Mode: X2APIC Deterministic Slider: Performance SMT: Enabled APBDIS-> 1 Enhanced Preferred I/O: Enabled
Other Settings	MRRS = 1K, PCIe Relaxed Ordering: Enabled
Boot Settings	rcu_nocbs=16-63 isolcpus=16-63 amd_iommu=on iommu=pt default_hugepagesz=1G hugepagesz=1G hugepages=64 nohz=off nosoftlockup selinux=0 numa_balancing=disable processor.max_cstate=0
Command Line	testpmd -l 32,33,34,35,36,37,38,39,63 -n 4 --socket-mem=4096 --master-lcore 63 -- --txq=8 --rxq=8 --rxd=4096 --txd=4096 --nb-cores=8 -i

Table 18: Results of P2100G 1x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		8-Core (Logical Cores)
64	148.81	103.61
128	84.46	84.46
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 12: Results of P2100G 1x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0



7 P2100G 2x 100G Bi-directional Zero Frame Loss – PCIe 4.0

RFC2544 Zero Frame Loss Performance on Broadcom NetXtreme-E Series P2100G with 2x 100G Bi-directional traffic.

Table 19: Setup of P2100G 2x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0

Item	Description
Server	AMD Rome Server Reference Platform
CPU	Single-Socket AMD EPYC 7742 64-Core Processor at 2.25 GHz (running at 3.38 GHz)
RAM	128 GB : 16 GB x 8 DIMMs at 3200 MHz
NIC	Broadcom NetXtreme-E Series P425G 4x 25G, PCIe Gen4 x16
Operating System	Red Hat Enterprise Linux Server release 7.6
Kernel Version	Linux Upstream 5.3.4
Broadcom Firmware Version	216.4.74.0
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 is forwarded to the other port by testpmd Eight TX/RX queue pairs for two ports (four per port) are assigned to eight logical CPU cores Data points are taken with eight logical CPU cores

Figure 13: Topology of P2100G 2x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0

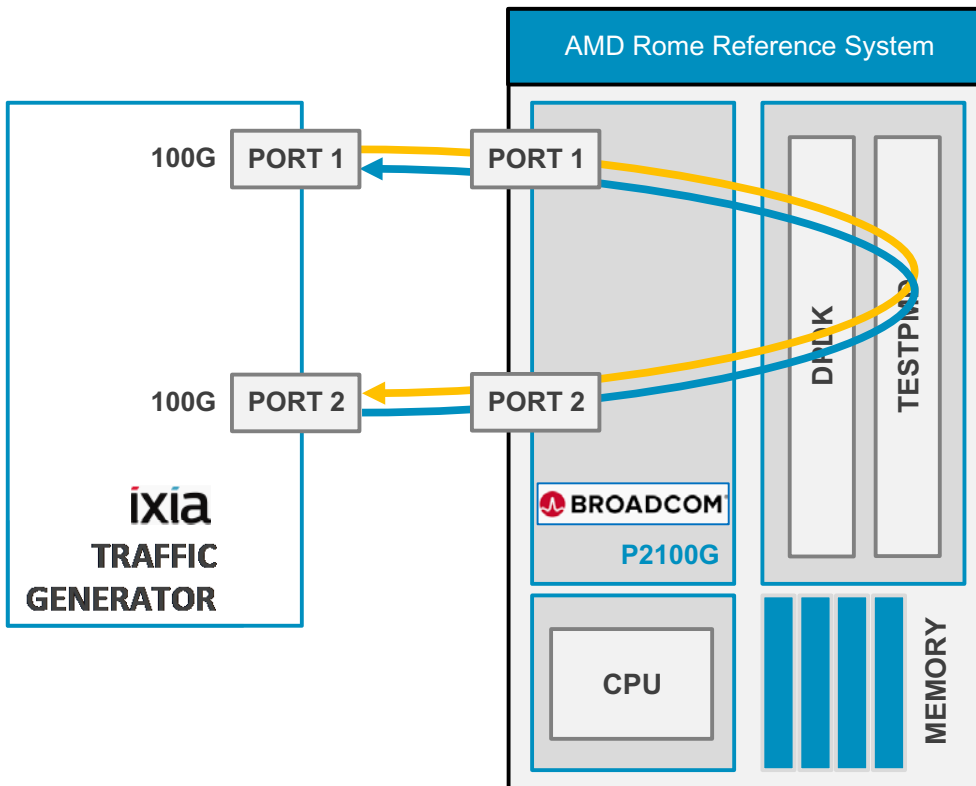



Table 20: Configuration of P2100G 2x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0

Item	Description
BIOS	NPS=1, Set Preferred IO: Enable L3 Cache as NUMA: Disabled, PCIE 10 Bit Tag: Enable APIC Mode: X2APIC Deterministic Slider: Performance SMT: Enabled APBDIS-> 1 Enhanced Preferred I/O: Enabled
Other Settings	MRRS = 1K (default), PCIe Relaxed Ordering: Enabled
Boot Settings	rcu_nocbs=16-63 isolcpus=16-63 amd_iommu=on iommu=pt default_hugepagesz=1G hugepagesz=1G hugepages=64 nohz=off nosoftlockup selinux=0 numa_balancing=disable processor.max_cstate=0
Command Line	testpmd -l 32,33,34,35,36,37,38,39,40,41,42,43,44,45,46,47,48,49,50,51,52,53,54,55,56,57 ,58,59,60,61,62,63 -n 4 --socket-mem=4096 --master-lcore 63 -- --txq=8 --rxq=8 --rxd=4096 --txd=4096 --nb-cores=8 -i -a

Table 21: Results of P2100G 2x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0

Frame Size (Bytes)	Line Rate (Mpps)	Frame Rate (Mpps)
		8-Core (Logical Cores)
64	297.62	106.23
128	168.92	104.54
256	90.58	80.26
512	46.99	42.11
1,024	23.95	23.29
1,280	19.23	19.23
1,518	16.25	16.25

Figure 14: Results of P2100G 2x 100G Bi-directional Zero Frame Loss Test – PCIe 4.0



Revision History

NetXtreme-DPDK-20-08-Performance-TR100; October 20, 2020

Initial release.

