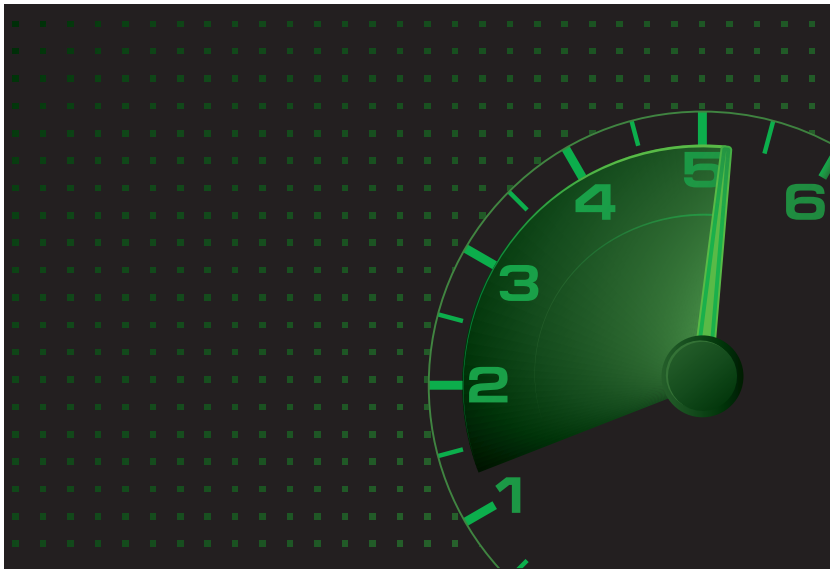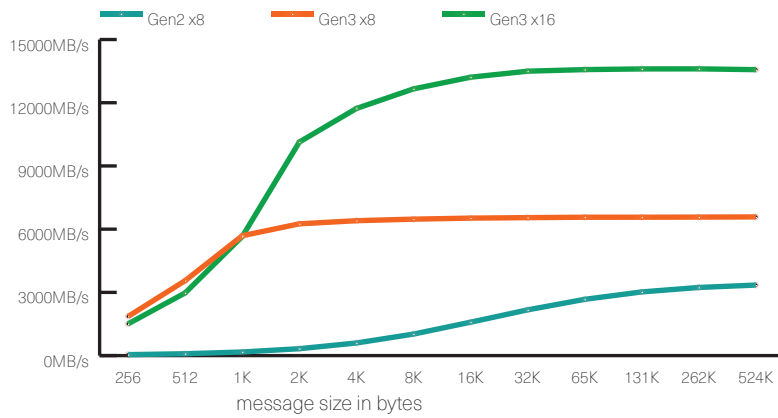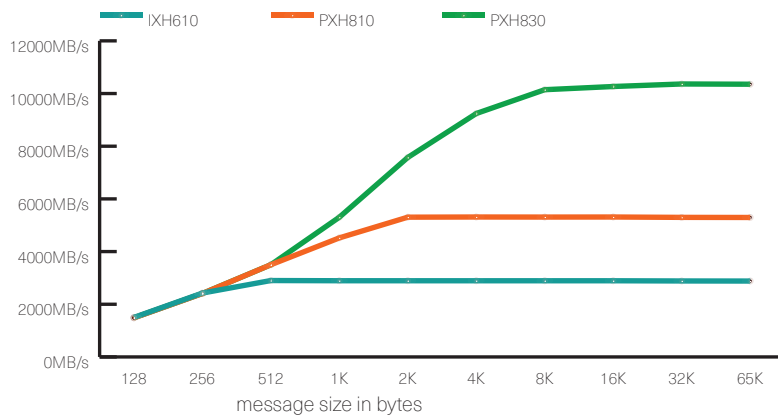# Performance/ Topologies

PCI Express provides various aspects of performance including latency, throughput, and scalability. Dolphin's products expose these aspects of performance with our eXpressWare software and high performance hardware. The combination enables building various system topologies that deliver high performance results.
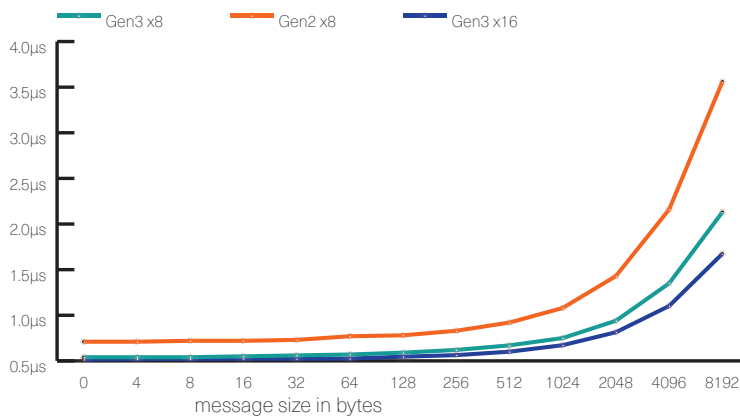
Throughput

DMA Throughput

A DMA benchmark is included with the eXpressWare suite. The graph on the left shows the throughput of different PCIe generations. The current Gen3 PXH830 can achieve 13 GBps with DMA enabled.

Throughput chart (Gen2 x8, Gen3 x8, Gen3 x16)

y-axis: 15000MB/s, 12000MB/s, 9000MB/s, 6000MB/s, 3000MB/s, 0MB/s
x-axis (message size in bytes): 256, 512, 1K, 2K, 4K, 8K, 16K, 32K, 65K, 131K, 262K, 524K
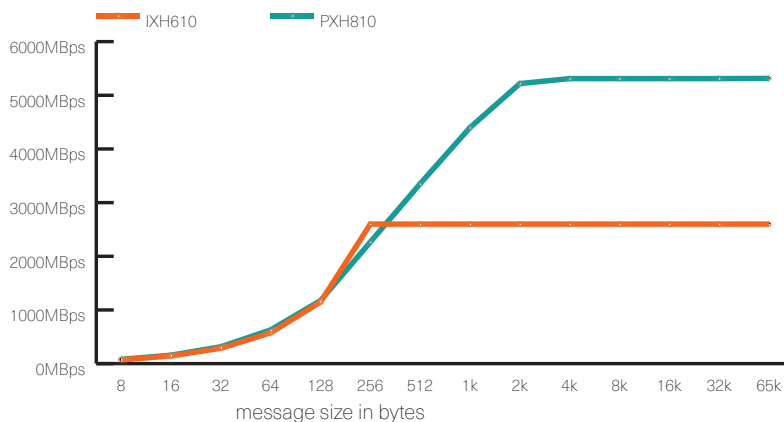
Throughput

Programmed IO (PIO) Throughput

The PIO benchmark provides a CPU only throughput analysis. This performance varies based on CPU. This chart shows that over 10 GBps can be achieved using only PIO.

Throughput chart (IXH610, PXH810, PXH830)

y-axis: 12000MB/s, 10000MB/s, 8000MB/s, 6000MB/s, 4000MB/s, 2000MB/s, 0MB/s
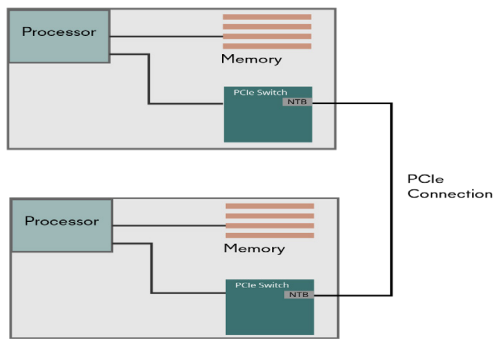x-axis (message size in bytes): 128, 256, 512, 1K, 2K, 4K, 8K, 16K, 32K, 65K

Latency

PCIe latencies are some of the lowest in the industry. This chart illustrates consistently low latency. Our latencies start at 540 ns between nodes. We test this with a loop back test, so this is actual application latency.

Latency chart (Gen3 x8, Gen2 x8, Gen3 x16)

y-axis: 4.0µs, 3.5µs, 3.0µs, 2.5µs, 2.0µs, 1.5µs, 1.0µs, 0.5µs
x-axis (message size in bytes): 0, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192

Reflective Memory

Dolphin's reflective memory solution improves performance vs older generations of reflective memory solutions. It provides higher throughput along with the lower latency of PCI Express.

Reflective Memory chart (IXH610, PXH810)

y-axis: 6000MBps, 5000MBps, 4000MBps, 3000MBps, 2000MBps, 1000MBps, 0MBps
x-axis (message size in bytes): 8, 16, 32, 64, 128, 256, 512, 1k, 2k, 4k, 8k, 16k, 32k, 65k

Two Node

PCI Express is a point-to-point protocol. Using non-transparent bridging (NTB), two systems can be connected through a PCI Express connection. This connection can be a cable or through a backplane.
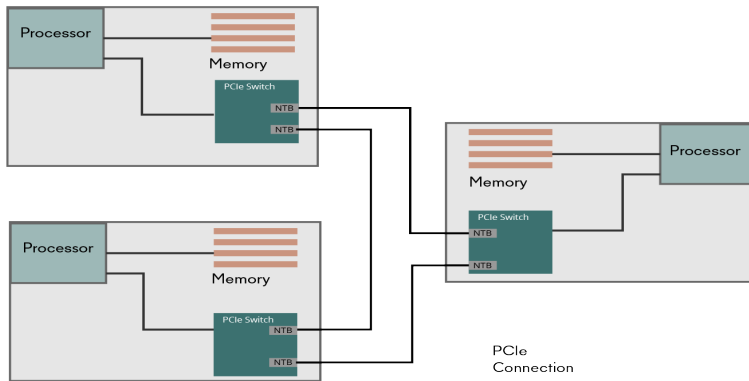

Switchless

Depending on the number of NTBs, three or more nodes can be connected without using an external switch. These topologies are high performance as they don't introduce the latency of another switch chip.
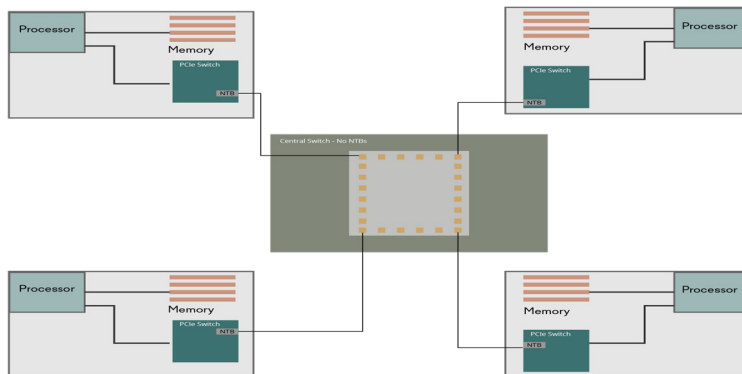
Central Switch

To scale out topologies, a external switch can be introduced. The NTBs are located on the host adapters to provide higher availability. Multiple switches can be cascaded to create large clusters.
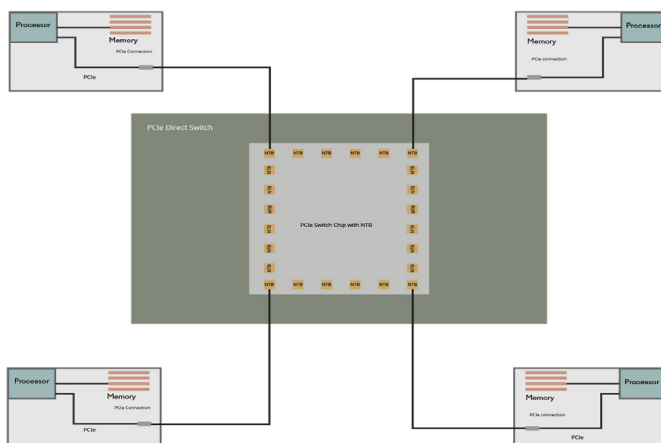
Direct Switch

In a direct switch topology, the NTBs are located in the central switch and not on the host adapters. This is ideal for backplane applications since it eliminates the need for power-on sequencing and provides better performance.

11