

Scaling Ansys® LS-DYNA® with Cornelis Omni-Path Express

High Performance Computing (HPC) is an increasingly important tool in the design process in many sectors of manufacturing. Leveraging HPC tools for engineering simulations and ultimately digital twins reduces time to results, increases accuracy and improves overall design efficiency.

The high performance of simulations is critical because it means that either more simulations can be completed at a given resolution in a design study, or higher fidelity simulations can be completed before strict deadlines. When simulations are performed using more than one server (compute node), the communication network (fabric) between those nodes can impact performance. The simulation progresses only as fast as the fabric enables them to communicate.

Ansys LS-DYNA^[1] is an industry leading Finite Element Analysis (FEA) solver for all types of physics such as drop tests, impacts and penetration, crashes, and more. In this article, the performance of Ansys LS-DYNA is shown on Dell PowerEdge R7525 platforms with AMD EPYC™ 7713 CPUs using two different fabric technologies - 100 Gbps **Cornelis Omni-Path Express (OPX)** and 200Gb HDR InfiniBand. The focus is on the compute and communication performance using the Intel® MPI Library and how the performance scales up to a cluster of 8 Dell Technologies dual socket AMD EPYC™ 7713 nodes, each with 32x16GB 3200MHz DDR4. Complete configuration is at the end of the article.

The standard topcrunch car2car benchmark^[2] (a two vehicle collision) is used to compare the scaling performance of Cornelis 100Gb OPX versus 200Gb HDR InfiniBand. Time to run the simulation is shown in Figure 1 for up to 8 nodes/1024 cores using the average of three runs, as shown by the columns - lower is better. The scaling efficiency normalized to one node is shown on the secondary axis. The performance of OPX matches that of HDR for this application at 8 nodes, where both networks scale with roughly 60% efficiency. This is a modest sized workload for Ansys LS-DYNA which pushes scaling limits, while other larger workloads will scale with better efficiency.




Figure 1: Performance of the Ansys LS-DYNA car2car benchmark up to 8 servers - lower is better

The advantage of OPX is clear when considering its substantially lower cost compared to HDR. To perform this comparison, fabric pricing was obtained from public sources^[3] for OPX and HDR to build an 8 node cluster, built from 8 host adapters, 2M copper cables and one managed edge switch. Performance is defined as the number of benchmark cases that can run in one year, and for each node count the performance is normalized by the cost of the 8 node fabric.

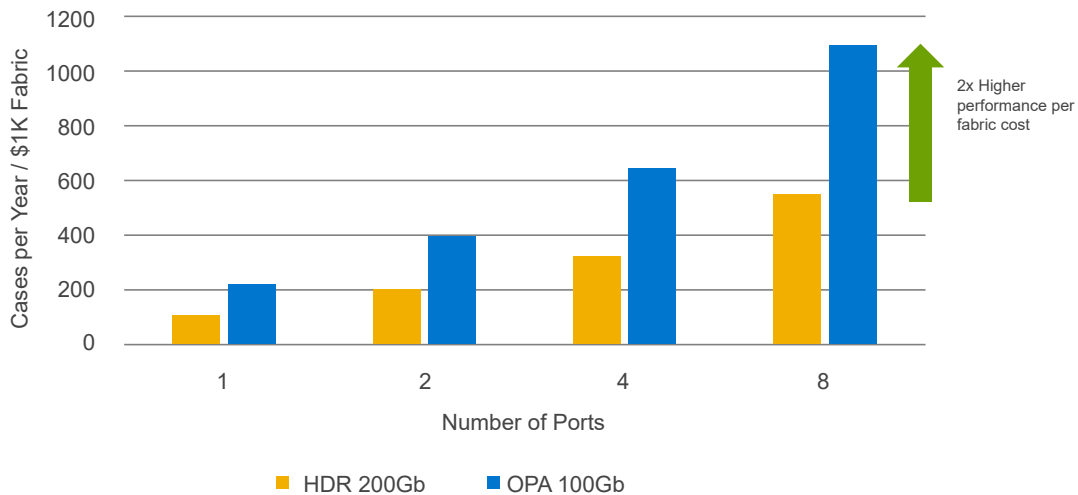


Figure 2: Performance normalized by 8 node fabric cost

As seen in Figure 2, **Omni-Path Express delivers 2x higher performance per fabric cost** for Ansys LS-DYNA when running the car2car benchmark. This means that customers can deploy an equally performant cluster for less and redirect unused budget to additional application licenses or compute power.

The new **Cornelis Omni-Path Express (OPX)** provider for libfabric unleashes additional performance on existing 100Gb Omni-Path hardware and is the provider of choice for the next generation CN5000 architecture. Thanks to a streamlined code path designed specifically for libfabric, it provides small message latency and message rate advantages over the traditional PSM2 provider. OPX is a drop-in replacement for PSM2 and can be evaluated for performance side-by-side with PSM2. Obtain OPX either through the Cornelis Omni-Path Express Suite (OPXS) software download at www.cornelisnetworks.com, or build your own using libfabric version 1.16.1 or newer from [github](https://github.com)^[4].

Cornelis Omni-Path Express is ready to ship with minimal lead time, contact your vendor of choice today to start experiencing HPC with OPX!

System configuration

Tests performed on Dell PowerEdge R7525 platforms, 2 socket AMD EPYC™ 7713 64-Core Processors. Rocky Linux 8.4 (Green Obsidian). 4.18.0-305.19.1.el8_4.x86_64 kernel. 32x16GB, 256 GB total, 3200 MT/s. BIOS: Logical processor: Disabled. Virtualization Technology: disabled. NUMA nodes per socket: 4. CCXAsNumaDomain: Enabled. ProcTurboMode: Enabled. ProcPwrPerf: Max Perf. ProcCStates: Disabled.

```
lsdyna_sp_mpp.e v221 i=c2c.inp memory=2000M memory2=200M.Intel® MPI Library 2019.6, 64 processes per node.
I_MPI_FABRICS=shm:ofi I_MPI_PIN=yes I_MPI_PIN_PROCESSOR_LIST=0-3,8-11,16-19,24-27,32-35,40-43,48-51,56-59,64-67,
72-75,80-83,88-91,96-99,104-107,112-115,120-123. Cornelis Omni-Path 10.11.1.1, hfi1.conf driver settings: rcvhdrCnt=8192
num_user_contexts=64. HFI installed on NUMA 4. I_MPI_OFI_PROVIDER=opx, FI_OPX_RELIABILITY_SERVICE_USEC_MAX=300,
FI_OPX_RELIABILITY_SERVICE_PRE_ACK_RATE=64, pre-production OPX rc1 internal build. HDR InfiniBand: InfiniBand:
MLNX_OFED_LINUX-5.4-2.4.1.3. mlx5_core version 5.4-2.4.1, HPC-X 2.11 with UCX 1.13.0. HCA installed on NUMA 14.
I_MPI_OFI_PROVIDER=mlx, UCX_NET_DEVICES=mlx5_0:1,UCX_TLS=ud,sm,self, I_MPI_OFI_EXPERIMENTAL=1.
```

References

- <https://www.ansys.com/products/structures/ansys-ls-dyna>
- <https://ftp.lstc.com/anonymous/outgoing/topcrunch/car2car/car2car.tgz>
- Pricing obtained on 1/27/2023 from <https://www.colfaxdirect.com/store/pc/home.asp>. Mellanox MCX653105A-HDAT \$1267.50 per adapter. Mellanox MQM8700-HS2F managed HDR switch, \$19,910.50. MCP1650-H002E26 2M copper cable - \$248. Omni-Path Express pricing obtained on 1/27/2023 from <https://www.nextwarehouse.com/>. Cornelis 100HFA016LSDCELL 100Gb HFI \$493.21 per adapter. Cornelis Omni-Path Edge Switch 100 Series 48 port Managed switch - \$11,366.78. Cornelis Networks Omni-Path QSFP 2M copper cable - \$101.26. Exact pricing may vary depending on vendor and relative performance per cost is subject to change.
- <https://github.com/ofiwg/libfabric>