

Delivering Leadership Performance with OpenRadioss™ and Cornelis™ Omni-Path Express™

Article

The open-source software model has proven that software developed and tested by the community delivers more reliable, functional, and performant solutions to users. In a pioneering strategy to accelerate innovation around the industry-proven Altair® Radioss® finite element analysis solver, Altair has released an open-source version named OpenRadioss¹.

Cornelis Networks has likewise embraced the open-source community with the release of Cornelis Omni-Path Express, an enhanced version of the high-performance Omni-Path interconnect based on a new open-source software stack.

To highlight the performance and price/performance capabilities of these open-source solutions, this article compares the job turnaround time of multi-node OpenRadioss simulations when they are run on a Cornelis Omni-Path Express fabric or an NVIDIA InfiniBand HDR fabric.

With Cornelis Omni-Path Express, users will experience better OpenRadioss performance and up to 2.3x performance per fabric cost than NVIDIA InfiniBand HDR.

OpenRadioss software simulates how materials interact and deform based on outside influences, such as a car crash, bridge deformations under heavy load, or even a cell phone dropping on a kitchen floor. These simulations model how millions of elements react to external forces through each millisecond of an event. For small workloads that can be performed in a single compute node, CPU and memory bandwidth are key to performance. However, as the simulation size grows beyond the capability of a single node, the fabric becomes a critical consideration.

Cornelis Omni-Path Express is designed specifically for high-performance, parallel computing environments. It is built utilizing a unique link-layer architecture and a highly optimized OFI libfabric provider² that delivers higher message rates and lower latencies than competing interconnects with a leadership price/performance value proposition.

In this paper, the 10-million cell taurus model (TAURUS_A05_FFB50³) shown in **Figure 1** is used to demonstrate how the network fabric affects application run time. Consistent with other studies^{4,5}, the model was shortened to two milliseconds to increase turn-around time.




Figure 1. Taurus 10M Cell Model.

Figure 2 compares the performance of the benchmark using up to 8 AMD EPYC 7713 dual-socket nodes, for a total of 1024 cores, connected with a 100Gbps Cornelis Omni-Path Express fabric and the same nodes connected with a 200Gbps NVIDIA InfiniBand HDR fabric. OpenMPI is used and OpenRadioss was compiled with gcc10.2 using the default build flags. The performance shown is job throughput (the number of cases able to run in one day if they were executed back-to-back without down time). Since the simulation time was shortened by a factor of 60 from the original model, the

¹ Industry-Proven Altair Radioss Finite Element Analysis Solver Now Available as Open-Source Solution, <https://www.altair.com/newsroom/news-releases/industry-proven-altair-radioss-finite-element-analysis-solver-now-available-as-open-source-solution> and www.openradioss.org

² <https://ofiwg.github.io/libfabric/>

³ HPC Benchmark Models, <https://openradioss.atlassian.net/wiki/spaces/OPENRADIOSS/pages/47546369/HPC+Benchmark+Models>

⁴ AMD EPYC™ & Altair Radioss™ Powering the Future of HPC, <https://www.amd.com/system/files/documents/amd-epyc-with-altair-radioss-powering-hpc.pdf>

⁵ Assuring Scalability: Altair Radioss™ Delivers Robust Results Quickly for Crash-Safe Vehicle Designs, https://www.altair.com/docs/default-source/resource-library/hw_radioss_whitepaper_letter_052620.pdf

number of cases per day is reduced by the same factor of 60 to represent the original model. The results show that Cornelis Omni-Path Express performs up to 9% faster than NVIDIA InfiniBand HDR in an 8-node cluster. The best performance for both fabrics was at 8-nodes achieved using 32 MPI ranks per node with 4 OpenMP threads per rank, leveraging the hybrid parallelization mode offered with OpenRadioss.

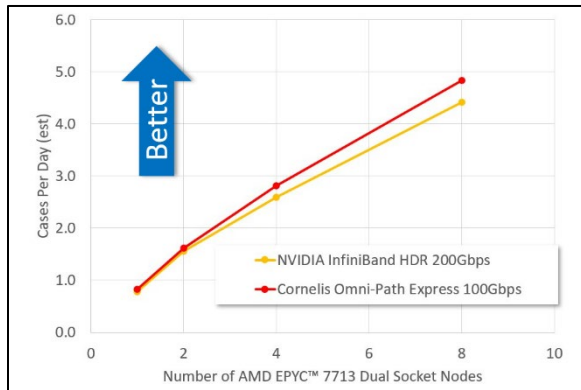


Figure 2. Performance of the OpenRadioss Taurus T10m benchmark.

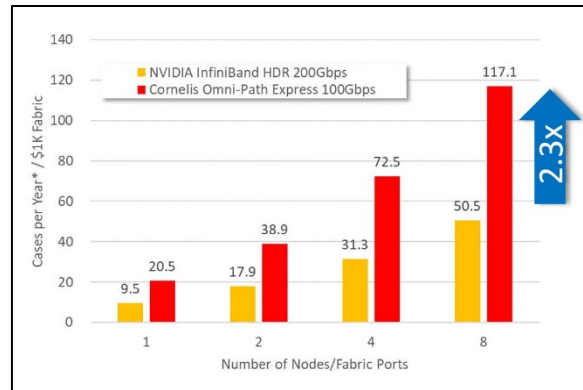


Figure 3. Performance of the OpenRadioss Taurus T10m benchmark, per fabric cost.

In addition to performance, another important consideration in fabric selection is price. For this second comparison, fabric pricing was obtained from public sources⁶ to build an 8-node cluster consisting of a single edge switch, 8 cables, and 8 host adapters. Performance is shown in terms of job throughput per year normalized by the cost of the 8-node fabric.

As seen in **Figure 3**, the Cornelis Omni-Path Express cluster delivers up to 2.3x better job throughput per fabric cost compared to the NVIDIA InfiniBand HDR cluster. This means users can obtain higher OpenRadioss performance with a lower budget, or they can deploy more nodes with the same budget to shorten time to results. In conclusion, the OpenRadioss open-source software combined with Cornelis Omni-Path Express delivers leadership performance with a significantly higher return on investment when compared to NVIDIA InfiniBand HDR.

System Configuration

Tests performed on 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.

OpenRadioss-latest-20230209 compiled with gcc 10.2. Example run command: mpirun -np 256 --map-by numa:PE=4 -x OMP_PLACES=cores --bind-to core -x OMP_NUM_THREADS=4 -mca btl self,vader -x OMP_STACKSIZE=400m -hostfile 8nodes ./engine_linux64_gf_ompi -i TAURUS_A05_FFB50_0001.rad -nt 4.

Cornelis Omni-Path Express: Open MPI 4.1.4 compiled with gcc 10.2. libfabric 1.16.1 compiled with gcc 10.2. Additional run flags: -x FI_PROVIDER=opx -mca mtl ofi -x FI_OPX_HFI_SELECT=0

NVIDIA HDR: OpenMPI 4.1.5a1 as provided by hpcx-v2.13.1-gcc-MLNX_OFED_LINUX-5-redhat8-cuda11-gdrcopy2-nccl2.12-x86_64, UCX version 1.14.0. Additional run flags: -x UCX_NET_DEVICES=mlx5_0:1 -mca coll_hcoll_enable 0 (performance was lower with hcoll enabled).

⁶ 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 100HFA016LS 100Gb HFI \$558.88 per adapter. Cornelis Omni-Path Edge Switch 100 Series 48 port Managed switch 100SWE48QR2 - \$9,996.64. 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.