

## OpenFOAM® Performance and Scalability with Cornelis Networks™ OPX on AMD EPYC™ 7003-series Processors

### Article

Cornelis Networks is the leading independent provider of purpose-built, open-source, scale-out interconnects for high-performance computing, artificial intelligence, and high-performance data analytics. The Cornelis Omni-Path Express™ (OPX) high-performance fabric delivers class-leading throughput, latency, and scalability allowing customers to deploy solutions which enable faster time to solution and improved workload scalability combined with leading price-performance.

To highlight the performance and price-performance capabilities of the Omni-Path fabric, this paper compares the performance of an industry standard OpenFOAM benchmark case on clusters interconnected by both Cornelis OPX and NVIDIA InfiniBand HDR fabrics.

OpenFOAM is a suite of solvers, pre- and post-processing utilities developed to solve modelling and simulation problems in the domain of CFD. In this paper, the industry standard motorbike tutorial case, visualized in Figure 1<sup>1</sup> is meshed at a resolution of 20M and 42M cells and is used to demonstrate how the fabric affects application run time and scalability.




Figure 1. Visualization of OpenFOAM Motorbike<sup>1</sup>.

These simulations model a low speed (incompressible) air flow around a motorcycle and rider. The simulation solves for a steady-state solution which is commonly used in many scientific fields where metrics such as skin friction and drag are a primary design factor. While smaller workloads can be performed in a single compute node where CPU performance and memory bandwidth are key, larger models spanning multiple nodes require a high-performance and cost-effective fabric.

Cornelis OPX is designed specifically for high-performance, parallel computing environments. It is built utilizing a unique link-layer architecture and a highly optimized OFI libfabric provider<sup>2</sup> delivering higher message rates and lower latencies than competing interconnects and with leadership price-performance.

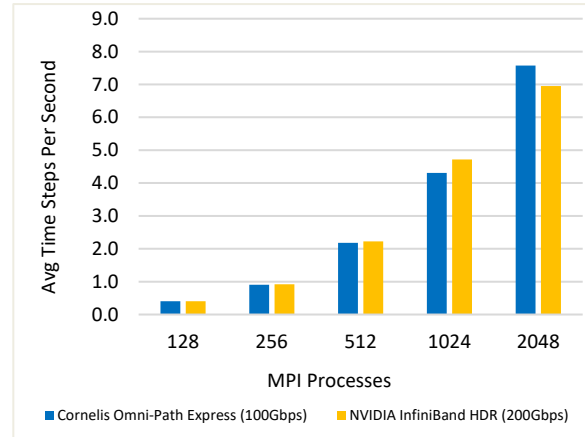
## Scalability

Figure 2 compares the performance scalability of the benchmark using up to 16 AMD EPYC™ 7713 dual-socket nodes, for a total of 2048 cores, connected with Cornelis OPX fabric using a single rail operating at 100Gbps and the same nodes connected with an NVIDIA InfiniBand HDR fabric operating at 200Gbps. OpenMPI is used - version 4.1.4 for Cornelis Omni-Path and version 4.1.5a1

<sup>1</sup> OpenFOAM runTimePostProcessing visualization; <https://www.openfoam.com/news/main-news/openfoam-v3.0/post-processing>

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

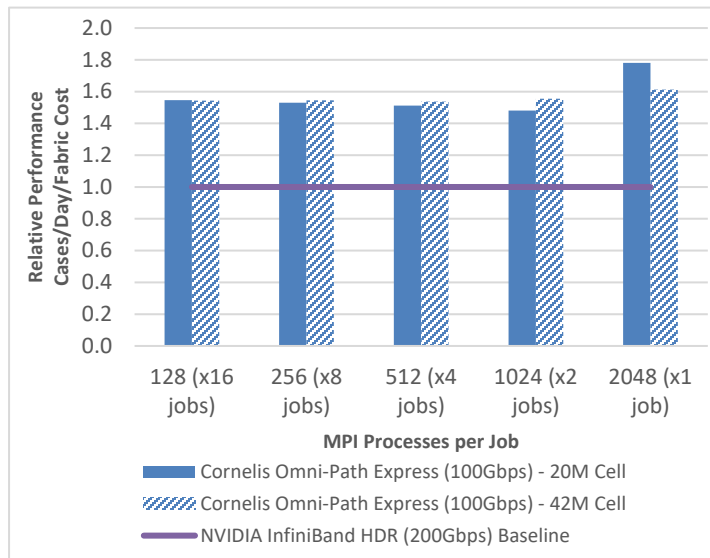
as released with NVIDIA HPC-X version 2.15. Cornelis Omni-Path measurements are performed with the OPX provider from libfabric 1.18.0. Complete configuration is given at the end of the paper. The results show that a single 100Gbps rail of Cornelis OPX performs up to 9% faster than NVIDIA InfiniBand HDR running at 200Gbps in a 16-node cluster. At 16 compute nodes (2048 MPI Processes), Cornelis OPX continues to scale beyond the ability of NVIDIA InfiniBand HDR. Each data point was run five times, eliminating the minimum and maximum performance, and averaging the middle three.



**Figure 2. Scalability of the OpenFOAM Motorbike 20M cell model.**

## Price-Performance

In addition to performance, another important consideration in fabric selection is price. For this second comparison, MSRP pricing was used<sup>3</sup> to build a 16-node cluster consisting of a single edge switch, 16 cables, and 16 host adapters.



**Figure 3. Job Throughput (Cases/Day) normalized by Fabric Cost for Motorbike 20M and Motorbike 42M cell models.**

Performance is shown in terms of job throughput per day on a fully utilized 16-node cluster normalized by the cost of the fabric.

In Figure 3, the results show that a Cornelis OPX connected cluster delivers an average of 1.57x higher job throughput per fabric cost running the Motorbike 20M and 42M test cases compared to the NVIDIA InfiniBand HDR cluster with a maximum advantage of over 1.78x throughput per fabric cost using the Motorbike 20M cell model at a job size of 16 nodes. This means users can obtain peak OpenFOAM performance with a lower budget, or they can deploy more nodes with the same budget to increase computational capacity and/or shorten the time to results.

In conclusion, the OpenFOAM software combined with Cornelis Networks OPX fabric delivers leadership performance and up to a 1.78x better return on investment. Cornelis Networks Omni-Path (100-series) hardware is available now, contact [sales@cornelisnetworks.com](mailto:sales@cornelisnetworks.com) to get started!

<sup>3</sup> MSRP Pricing obtained on 7/11/2023 from <https://store.nvidia.com/en-us/networking/store>. Mellanox MCX653105A-HDAT \$1628 per adapter. Mellanox MQM8700-HS2F managed HDR switch, \$25555. MCP1650-H002E26 2M copper cable - \$281. Cornelis OPX MSRP pricing as of 7/11/2023. Cornelis 100HFA016LSN 100Gb HFI \$880 per adapter. Cornelis Omni-Path Edge Switch 100 Series 48 port Managed switch 100SWE48QF2 - \$19750. Cornelis Networks Omni-Path QSFP 2M copper cable 100CQQF3020 - \$147. Exact pricing may vary depending on vendor and relative performance per cost is subject to change.

## System Configuration

Tests performed on 2 socket AMD EPYC 7713 64-Core Processor-based servers. 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.

OpenFOAM v22.06 SimpleFoam compiled with gcc 10.2. Example run command: `mpirun -np ${NP} --map-by ppr:${PPN}:node -hostfile ${HOSTS} simpleFoam -parallel. blockMeshDict 100x40x40 (20M) and 130x52x52 (42M), decomposeParDict - scotch decomposition`

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

NVIDIA HDR: OpenMPI 4.1.5a1 as provided by hpcx-v2.15-gcc-MLNX\_OFED\_LINUX-5-redhat8-cuda12-gdrCOPY2-nccl2.17-x86\_64, UCX version 1.15.0. Additional run flags: `-mca mtl ucx -x UCX_NET_DEVICES=mlx5_0:1`

## Legal Disclaimer

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Cornelis Networks products described herein. You agree to grant Cornelis Networks a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All product plans and roadmaps are subject to change without notice.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cornelis Networks technologies may require enabled hardware, software, or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Cornelis, Cornelis Networks, Omni-Path, Omni-Path Express, and the Cornelis Networks logo belong to Cornelis Networks, Inc. Other names and brands may be claimed as the property of others.

This offering is not approved or endorsed by OpenCFD Limited, producer and distributor of the OpenFOAM software via [www.openfoam.com](http://www.openfoam.com), and owner of the OPENFOAM® and OpenCFD® trademarks.

Copyright © 2024023, Cornelis Networks, Inc. All rights reserved.