

Achieve Highly Efficient Network Scaling with Ansys Fluent[®] on Intel[®] Xeon[®] 6 and Cornelis[®] CN5000 Omni-Path[®] 400G

White Paper

Computational Fluid Dynamics (CFD) simulations are among the most demanding workloads in high-performance computing (HPC), requiring both substantial compute performance and efficient inter-node communication. As engineering organizations increasingly rely on simulation-driven design, scalable HPC infrastructure has become essential for reducing time-to-solution and accelerating product development cycles.

This paper evaluates the performance and scalability of Ansys Fluent fluid simulation software running on Intel Xeon 6 processor-based clusters interconnected with the Cornelis CN5000 Omni-Path 400G fabric. Nine representative CFD workloads ranging from 24 million to 280 million cells were evaluated across configurations ranging from 1 to 16 nodes.

Results demonstrate excellent application scalability, achieving a geometric mean speedup of 13.93 \times at 16 nodes while maintaining 87.1% geometric mean parallel efficiency. Across all workloads, application throughput increased consistently as the cluster size increased, demonstrating the ability of Intel Xeon 6 processors and the CN5000 fabric to support large-scale CFD simulations efficiently.

Test Environment

Benchmark testing was conducted using a cluster comprised of Intel Xeon 6 processor-based compute nodes interconnected using the Cornelis CN5000 Omni-Path 400G fabric.

Hardware Configuration

Processors	Dual socket Intel Xeon 6787P
Processor Cores	86 cores per CPU
Memory	16 x 32 GB 6400 MT/s DDR5
OS	SLES 15 SP7, kernel: 6.4.0-150700.53.40-default
Interconnect	Cornelis CN5000 Omni-Path, 12.1.1.0.18 host software release
Network Speed	400 Gbps
MPI Library	Intel MPI 2021.17 (Build 20251215)
Node Counts Tested	1, 2, 4, 8, 16

Methodology

Performance was measured using Millions of Iterative Updates Per Second (MIUPS), the primary performance metric reported by Ansys Fluent. MIUPS measures solver throughput and provides a direct indication of simulation performance. Higher MIUPS values correspond to faster time-to-solution and increased engineering productivity.

Application scalability was evaluated using both speedup and parallel efficiency. Parallel efficiency measures how effectively additional compute resources contribute to application performance and is calculated as achieved speedup divided by ideal speedup.

Performance Results and Scaling Analysis

Benchmark Workloads


Nine representative CFD workloads were selected to evaluate both application performance and scalability.

Workload	Cell Count (Millions)
Aircraft_wing_14m	14
combustor_24m	24
exhaust_system_33m	33
winsor_41m	41
drivaer_50m	50
airfoil_80m	80
f1_racecar_140m	140
drivaer_250m	250
open_racecar_280m	280

All workloads demonstrate substantial throughput improvements as additional nodes are added. The second largest benchmark, DrivAer 250M, achieves 572.78 MIUPS at 16 nodes.

As shown in Figure 1, performance increases consistently as additional nodes are added, with all workloads demonstrating substantial throughput gains through 16 nodes, highlighting the scalability of Intel Xeon 6 processors interconnected with the Cornelis CN5000 Omni-Path 400G fabric.

Figure 1. Ansys Fluent Performance Measured in MIUPS Across Eight CFD Workloads



The MIUPS results demonstrate strong throughput scaling across all workloads as cluster size increases.

Several workloads exhibited substantial performance gains:

Workload	1 Node MIUPS	16 Node MIUPS	Speedup (16 Node vs. 1 Node)
combustor_24m	17.09	192.30	11.25x
exhaust_system_33m	11.08	157.68	14.23x
winsor_41m	31.35	402.21	12.83x
drivaer_50m	35.23	488.67	13.87x
airfoil_80m	4.89	66.67	13.63x
f1_racecar_140m	7.81	124.78	15.98x
drivaer_250m	34.05	572.78	16.82x
open_racecar_280m	10.73	175.42	16.35x
Geometric Mean	14.22	197.43	13.93x

The three largest workloads, f1_racecar_140m, drivaer_250m, and open_racecar_280m, achieved 16x or greater increase in throughput between 1 and 16 nodes.

Parallel Scaling Efficiency


Parallel scaling evaluates how effectively application performance improves as additional compute resources are added. Parallel efficiency measures the percentage of ideal scaling achieved as node count increases.

The benchmark suite maintains near-ideal efficiency through 16 nodes, demonstrating effective utilization of distributed compute resources and low communication overhead. The largest workloads exhibit particularly strong scalability, with `drivaer_250m`, `open_racecar_280m`, and `f1_racecar_140m` demonstrating periods of super-linear scaling. This behavior is commonly observed in HPC applications when the working dataset becomes more effectively distributed across the aggregate memory and cache resources of the cluster, reducing memory-access overhead and improving computational efficiency.

At the other end of the spectrum, the smallest workloads, including `aircraft_wing_14m` and `combustor_24m`, begin to exhibit reduced parallel efficiency at 16 nodes. Although Ansys Fluent is highly scalable and can efficiently utilize configurations approaching 10,000 cells per core, these smaller models eventually reach a point where inter-node communication and synchronization overhead become a larger fraction of total runtime. As a result, the benefits of adding additional compute resources diminish for these workloads at larger node counts. Despite this expected behavior, both models continue to realize significant performance gains through 16 nodes.

As shown in Figure 2, efficiency is normalized to the single-node baseline (100%). The benchmark suite maintains approximately 100% geometric mean efficiency through eight nodes and achieves approximately 89% efficiency at 16 nodes, demonstrating excellent scalability and efficient utilization of distributed compute resources.

Figure 2. Parallel Efficiency Across Representative Ansys Fluent CFD Workloads, Including the Geometric Mean



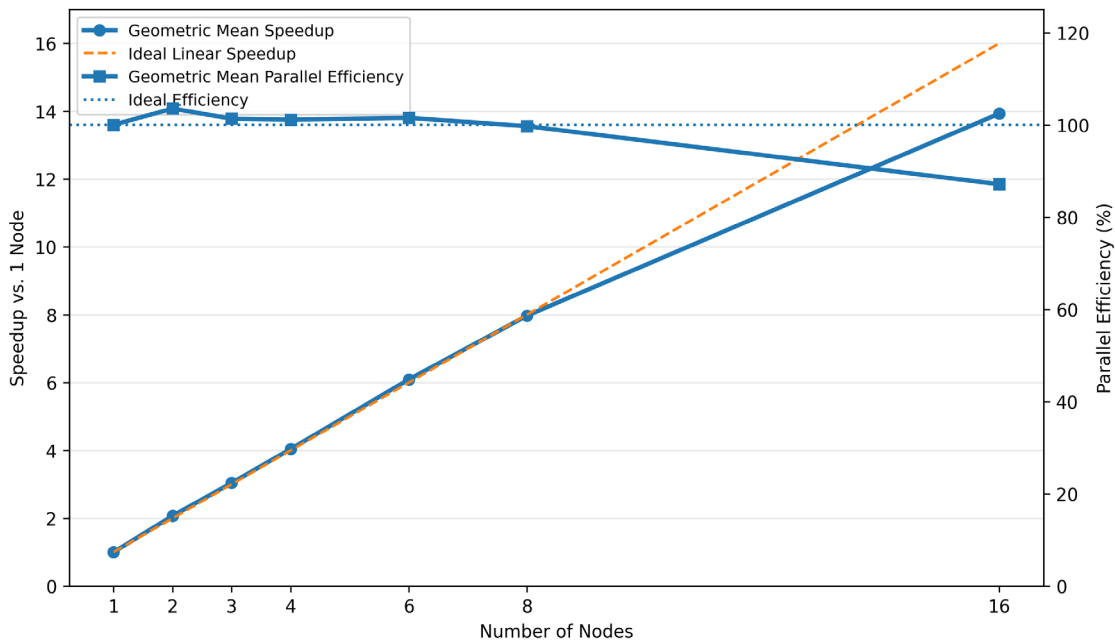
Geometric Mean Speedup and Scaling Summary

To provide a workload-independent view of overall application scalability, geometric mean speedup and parallel efficiency were calculated across all eight CFD workloads. The geometric mean is commonly used in HPC performance analysis because it gives equal weight to each benchmark and avoids allowing larger workloads to disproportionately influence the overall result. Together, speedup and parallel efficiency provide insight into both the absolute scaling achieved and how effectively additional compute resources are utilized as cluster size increases.

Nodes	Geometric Mean Speedup	Geometric Mean Efficiency
1	1.00x	100.0%
2	2.07x	103.5%
3	3.04x	101.3%
4	4.04x	101.1%
6	6.09x	101.5%
8	7.97x	99.7%
16	13.93x	87.1%

As shown in Figure 3, CN5000 achieves near-linear scaling through eight nodes, delivering a geometric mean speedup of 7.97x while maintaining 99.7% parallel efficiency. At 16 nodes, the benchmark suite reaches a geometric mean speedup of 13.93x with 87.1% efficiency, demonstrating the CN5000’s ability to scale communication-intensive CFD workloads across distributed clusters efficiently.

Figure 3. Geometric Mean Speedup and Parallel Efficiency Across Nine Ansys Fluent CFD Workloads



Conclusion and Key Findings

The combination of Intel Xeon 6 processors and Cornelis CN5000 Omni-Path 400G delivers excellent performance and scalability for Ansys Fluent CFD workloads.

Key findings include:

- Geometric mean speedup of 13.93× at 16 nodes
- Approximately 89% geometric mean parallel efficiency at 16 nodes
- Near-linear scaling through eight nodes
- Consistent performance gains across all benchmark workloads
- Efficient support for CFD models ranging from 14 million to 280 million cells
- Geometric mean performance increased from 14.22 MIUPS on a single node to 197.43 MIUPS at 16 nodes

These results demonstrate that Intel Xeon 6 processors and the Cornelis CN5000 fabric provide a highly scalable platform for engineering organizations seeking to accelerate CFD simulation workflows while reducing overall time-to-solution.

Cornelis 5000 + Xeon test clusters are available from Intel and Dell for you to benchmark this solution on your own workloads. For more information or to request a proof-of-concept demo, contact Cornelis Networks at sales@cornelisnetworks.com.

Solution Overview

Ansys Fluent

Ansys Fluent is a market-leading computational fluid dynamics (CFD) application that enables engineers to model fluid flow, heat transfer, turbulence, and combustion processes with high accuracy. Designed to leverage modern HPC systems, Fluent scales efficiently across distributed clusters, allowing organizations to solve increasingly large and complex engineering problems in automotive, aerospace, energy, and industrial applications.

For more information, visit the [Ansys Fluent](#) product page.

Intel Xeon 6 Processors

Intel Xeon 6 processors are designed to deliver exceptional performance for compute-intensive workloads, combining high core counts, increased memory bandwidth, and advanced I/O capabilities. Built to accelerate modern HPC, AI, and engineering simulation applications, Intel Xeon 6 processors provide the computational foundation required to efficiently scale demanding workloads such as Ansys Fluent across distributed cluster environments.

For more information, visit the [Intel Xeon 6](#) product page.

Cornelis CN5000 Omni-Path 400G

Cornelis CN5000 Omni-Path is an end-to-end high-performance networking solution engineered to deliver low-latency, high-message-rate communication for large-scale HPC and AI workloads.

At the core of the solution is the 400G CN5000 SuperNIC, designed to accelerate communication-intensive applications while maximizing scalability across large clusters. The CN5000 architecture combines advanced congestion management, adaptive routing, and lossless communication mechanisms to ensure predictable application performance at scale.

Key Technical Advantages

Lossless Data Transmission: CN5000 employs hardware-based, credit flow control to eliminate packet loss and retransmissions, ensuring predictable performance for communication-intensive HPC applications.

Dynamic Lane Scaling and Link-Level Replay: The architecture provides resiliency against cable and transceiver failures while automatically correcting transmission errors locally without impacting application performance.

Fine-Grained Adaptive Routing: CN5000 continuously monitors fabric conditions and dynamically routes traffic around congestion points, minimizing latency variation and improving scalability for tightly-coupled applications such as CFD.

Open Standards-Based Architecture: Built on the OpenFabrics Alliance (OFA) Libfabric software framework, CN5000 provides vendor-neutral interoperability across modern CPU and GPU platforms without proprietary software dependencies.

These capabilities make CN5000 particularly well-suited for latency-sensitive applications such as Siemens Simcenter STAR-CCM+.

Figure 4. Cornelis CN5000 Omni-Path 400G Portfolio



For more information, visit the [Cornelis CN5000](#) product page.



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.

Copyright © 2026, Cornelis Networks, Inc. All rights reserved.