



Accelerating Ansys Fluent® CFD Simulations with Cornelis® CN5000

White Paper

Cornelis Networks CN5000 Omni-Path solutions in AMD EPYC™ 9005 processor-based servers enable up to 30% higher cluster throughput and 129% scaling efficiency on Ansys Fluent CFD workloads vs. InfiniBand, delivering more simulations per day, shorter design cycles, and clear economic value.

Cornelis Networks delivers purpose-built interconnects for AI, HPC, and data-intensive computing. The CN5000 Fabric is engineered to eliminate communication bottlenecks, enabling customers to scale accelerated infrastructure with unmatched predictability, efficiency, and performance. Built on the Omni-Path Architecture and delivered with our open-source OPX Software stack (OFI/libfabric), the 400 Gbps CN5000 Fabric provides lossless, congestion-free transport with ultra-low latency, exceptional message rate, and scalable performance to hundreds of thousands of endpoints—purpose built for tightly coupled workloads like computational fluid dynamics (CFD).

AMD EPYC™ Processors power many of the world's leading supercomputers that are used in CFD, climate modeling, genomics, and physics research. The 5th Generation AMD EPYC Processors have outstanding core performance, compute density, and memory bandwidth to deliver the speed, scale, and efficiency to make them an ideal choice to pair with CN5000 for running CFD workloads.

Ansys Fluent fluid simulation software iterations trigger millions of small MPI messages; the bottleneck is message rate and tail latency, not raw link bandwidth. Conventional 400 GB networks optimized for bulk bandwidth can struggle with message rate and tail latency in tightly coupled solvers. CN5000 is optimized for precisely these workloads.

CN5000's architecture sustains far higher messaging throughput with consistently low latency, restoring linear scaling on CPU clusters. In the testing presented in this paper, CN5000 delivered up to 30% faster Fluent runtimes than InfiniBand NDR 400G and up to 129% scaling efficiency at 8 nodes—achieving 2.5x the MPI message rate with ~35% lower latency.

The result is more simulations per day, shorter design cycles, and better economics without adding compute nodes. The sections that follow outline the networking characteristics that govern Fluent performance and present methodology, results, and guidance for deploying CN5000 for AMD EPYC-based Fluent workloads.

CN5000 vs. NDR400: Technical Advantage

Scaling out performance is founded on the architectural strengths of CN5000, presenting superior empirical network performance—message rate, latency, and congestion behavior—that underpins its application-level benefits. CN5000’s architectural and design innovation ensures tight coupling between compute nodes, manifesting in differentiated network performance.

1. Message Rate

CN5000 sustains over 800 million bi-directional messages per second. Higher message rates directly benefit workloads that rely on many small MPI calls per second—characteristic of tightly-synchronized solvers like Fluent.

2. Latency

CN5000 delivers end-to-end MPI latencies below one microsecond including a switch hop, with a 35% reduction compared to NDR400—both average and tail latencies. In CFD workloads, this translates to fewer idle cores waiting on remote data, thereby improving time-to-solution.

3. Congestion Control

By integrating credit-based flow control and dynamic adaptive routing, CN5000 eliminates fabric congestion organically—unlike NDR400, which relies on larger buffers and recovery mechanisms. This ensures reliable and predictable performance even under high network load and complex communication patterns.

Metric	Network Performance	Benefits to Fluent Workloads
Message Rate	2.5x Higher vs. NDR 400G	Faster MPI exchanges, driving tighter solver loops
Latency	35% Lower vs. NDR 400G	Reduced idle wait time, improving core utilization
Lag and Congestion	Congestion Free	Maintaining performance at scale
Scaling Efficiency	115% Avg on 14m+ models	Delivering linear scaling, enabling predictable throughput

In the following section, we’ll demonstrate that these improvements enable faster Fluent simulations and linear scaling with job size.


Methodology and Ansys Fluent Results

We evaluated CN5000 vs. InfiniBand NDR 400G on identical 8-node AMD EPYC 9005 CPU-based systems running Fluent 2025 R1 with Open MPI, measuring wall-clock time per iteration on large, tightly coupled models. Each node was run at 128 processes per node (ppn), corresponding to 64 cores per socket. While these CPUs expose up to 128 hardware cores per socket, we standardized to 64 to relieve any constraints on memory-bandwidth and locality, while also reflecting volume core counts deployed in production today.

Both CN5000 and NDR 400G fabrics were installed on the same hardware paired with system software to measure the impact of the fabric on application performance. All tests run multiple iterations to ensure statistical stability. Details of measurement configurations are available in the [Appendix](#).

We focused on large, tightly coupled CFD simulations designed to take advantage of full cluster scaling. These correspond to Ansys Fluent models with high cell counts, and the results are summarized in the charts below.

Figure 1. CN5000 Performance Relative to InfiniBand NDR—8 Nodes



As shown in the results, CN5000 consistently outperforms NDR 400G across all large-scale models. Fluent workloads are dominated by numerous small MPI exchanges. The higher message rate of CN5000 allows more of these fine-grained, inter-node communications to complete within the same time frame, while its lower latency reduces idle waiting periods for data—collectively leading to shorter overall runtimes.

The performance gains underscore how CN5000 unlocks consistent solver acceleration and efficiency benefits, enabling tangible value for Ansys Fluent users.

- **Throughput Uplift**

Clusters with CN5000 can generate ~20–30% more Fluent CFD iterations by runtime savings alone—equivalent to adding extra compute nodes without extra hardware investment.

- **Time-to-Insight Reduction**

Faster turnaround eliminates delays in design iteration cycles—meaning critical decisions can be made more rapidly and designs validated sooner.

- **Competitive Differentiation**

For teams running large models on AMD EPYC 9005 CPU-based servers (for example, external aerodynamics, combustion, structural flow), CN5000 provides a clear and consistent performance edge over mainstream InfiniBand fabrics.

Figure 2. CN5000 Performance Scaling




Table 1. CN5000 8-Node Scaling Efficiency on Ansys Fluent Models

aircraft_wing_14m	landing_gear_15m	exhaust_system_33m	f1_racecar_140m
125%	129%	106%	99%

* 8-node performance scaling efficiency calculated relative to 1-node performance

Maintaining performance on distributed infrastructure is critical for meeting the growing demands of complex CFD simulations. Strong scaling efficiency is the key metric, ensuring that investments in additional compute nodes translate directly into faster results rather than diminishing returns from network bottlenecks. In our testing, CN5000 delivered exceptionally strong and consistent scaling across all models, averaging 115% efficiency from one to eight nodes. The 99% scaling efficiency on the largest model (f1_racecar_140m) is particularly noteworthy, as it validates that CN5000 preserves near-linear performance even on the most demanding workloads. This result, combined with the direct runtime uplift of up to 30%, confirms CN5000's ability to maximize cluster utilization and deliver predictable performance for large-scale Fluent workloads on AMD EPYC 9005 CPU-based servers.

About Cornelis Omni-Path

Cornelis CN5000 is built on the Omni-Path Architecture to specifically optimize for HPC and AI performance. Key architectural pillars of Omni-Path include:

- **Credit-Based Flow Control:** Ensures packets are sent only when receive buffers have available credits, preventing overrun and eliminating pause frames. Credit-based flow control allows per-virtual-lane flow regulation, enabling fair allocation across competing workloads.
- **Enhanced Congestion Avoidance:** Uses forward and backward-propagating telemetry to dynamically adjust pacing at the source. This avoids the classic cascading congestion seen in oversubscribed topologies and maintains consistent throughput even under hot-spot pressure.
- **Fine-Grained Adaptive Routing:** Real-time telemetry builds a global heatmap of switch buffer utilization, enabling path selection based on current network state rather than static tables. It automatically bypasses congested paths, reducing long-tail latency and improving collective efficiency.
- **Dynamic Lane Scaling and Link-Level Replay:** Each link can degrade gracefully if individual lanes fail, maintaining connectivity instead of triggering a full path teardown. Link-level replay catches forward-error-correction (FEC) misses at the hardware layer, eliminating costly end-to-end retransmissions.
- **Open and Interoperable:** Open software stack (OPX) built on OFI/libfabric, ensuring low-overhead integration with MPI stacks (for example, Open MPI), and ease of deployment in heterogeneous environments.

Conclusion

As Fluent CFD models grow in scale and complexity and as design iteration cycles accelerate, the performance of the underlying network fabric plays an increasingly vital role in cluster efficiency. This white paper has demonstrated that:

1. The CN5000 Fabric, rooted in Omni-Path Architecture, delivers a 2.5X increase in MPI message rate, 35% lower latency, and a lossless, congestion-free fabric design—parameters essential for tightly-coupled CPU-based Fluent workloads.
2. Testing across four large-scale Fluent models (up to 140 million cells) on an 8-node AMD EPYC 9005 CPU-based cluster produced consistent performance gains, up to 30% faster runtime.
3. The CN5000 design delivers superior scaling efficiency, up to 129%, enabling meaningful throughput increases and efficient utilization of compute resources.

With CN5000 in AMD EPYC 9005 CPU-based systems, Ansys Fluent simulation users will benefit from maximized throughput, driving faster design cycles with existing hardware, predictable performance guaranteeing consistent results in real project deployments, and cost-effective scaling.

With an Ansys technical endorsement, we invite you to explore CN5000 and Ansys Fluent models. Elevate your CPU-based Fluent simulations on AMD EPYC 9005 CPU-based clusters—unlock faster design cycles, predictable scalability, and improved compute utilization with CN5000.

Appendix

Hardware and software configurations detailed below:

Application:

Ansys Fluent 2025 R1. Select models.

8-Node Cluster Platform:

2 sockets using 2 x AMD EPYC 9755 processors.

SMT=disabled, NPS=1, 128ppn.

24 x 32 GB, 768 GB total, Memory Speed: 5600 MT/s.

Turbo enabled with acpi-cpufreq driver.

OS:

Rocky Linux 9.5 (Blue Onyx). 5.14.0-503.33.1.el9_5.x86_64 kernel.

CN5000 Network Parameters and Software:

Cornelis CN5000 Omni-Path SuperNIC

Cornelis CN5000 Omni-Path Switch

Cornelis OPX Software 12.0.2.0.16. hfi1 driver parameters: num_user_contexts=0,128

num_vls=4 num_sdma=8

Open MPI 5.0.8.

NVIDIA NDR InfiniBand Parameters and Software:

Mellanox Technologies MT2910 Family [ConnectX-7]

MQM9700-NS2F Quantum 2 switch

UCX and Open MPI 4.1.7rc1 as packaged in hpcx-v2.23.



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 © 2025, Cornelis Networks, Inc. All rights reserved.