Accelerating Siemens Simcenter™ STAR-CCM+™ Simulations with AMD EPYC™ 9755 Processors and Cornelis® CN5000 Omni-Path® 400G

White Paper

Computational Fluid Dynamics (CFD) simulations continue to increase in complexity as engineers pursue higher-fidelity models to improve product performance while reducing development time. As model sizes grow from hundreds of millions to more than one billion cells, efficiently scaling applications across distributed compute resources become essential for minimizing time-to-solution and maximizing engineering productivity.

This paper evaluates the performance and scalability of Siemens Simcenter STAR-CCM+™ running on dual AMD EPYC™ 9755 processor-based systems interconnected with the Cornelis® CN5000 Omni-Path® 400G fabric. Ten representative CFD workloads, spanning automotive, aerospace, industrial, and multiphase applications, were used to evaluate performance across mesh sizes ranging from 20 million to over one billion cells. The ultra-low, sub-1 microsecond MPI latency and high message rates achievable with Cornelis CN5000 enable STAR-CCM+ to scale with excellent parallel efficiencies.

Across nine representative benchmark workloads, the platform achieved a geometric mean parallel efficiency of 112.4%. Several workloads exhibited near-linear or super-linear scaling, including LeMans104M, EmpHydroCyclone, and DrivAer128M, demonstrating the ability of AMD EPYC processors and the Cornelis CN5000 interconnect to efficiently scale communication-intensive CFD applications.

Results demonstrate that the combination of Simcenter STAR-CCM+, AMD EPYC 9755 processors, and Cornelis CN5000 provides an exceptionally scalable platform for large-scale engineering simulations, enabling organizations to reduce simulation turnaround time while maximizing utilization of distributed compute resources.

Test Environment

Table 1. Hardware and Software Configuration

Component	Specification
Processor	Dual AMD EPYC 9755
Processor Cores	128 Cores per Processor
Memory	24 x 32 GB 6400 MT/s DDR5
OS	Rocky Linux 9.5 (Blue Onyx). 5.14.0-503.33.1.el9_5.x86_64 kernel
Interconnect	Cornelis CN5000 Omni-Path, 12.1.1 host software release
Network Speed	400 Gbps
MPI Library	Open MPI 5.0.6, `-mca mtl ofi -mca btl self,sm -x FI_PROVIDER=opx`
Node Counts	1, 2, 4, and 8

Performance Results and Scaling Analysis

The following results evaluate the performance and scalability of Simcenter STAR-CCM+ running on AMD EPYC 9755 processors interconnected by the Cornelis CN5000 fabric.

The primary performance metric is Average Time per Iteration (T_{avg}), as reported in the Simcenter STAR-CCM+ output files. Unlike total elapsed runtime, Average Time per Iteration removes one-time initialization overhead and reflects the steady-state computational performance experienced during production CFD simulations.

Because engineering workloads frequently execute hundreds or thousands of solver iterations, Average Time per Iteration provides a more representative measure of application performance than metrics based solely on total execution time. This methodology also aligns with how STAR-CCM+ users typically evaluate solver performance when comparing hardware platforms. For completeness, Total Elapsed Time is also reported for every benchmark configuration.

Benchmark Workloads

The benchmark suite spans a diverse range of automotive, aerospace, industrial, and multiphase CFD applications, representative of production engineering environments.

Table 2. Simcenter STAR-CCM+ Benchmark Workloads

Workload	Model Size (Million Cells)
EmpHydroCyclone	30
LeMans104M	104
AeroSUVSteadyCoupled322M	322
AeroSUVSteadySeg322M	322
AeroSUVSteadyCoupled106M	106
AeroSUVSteadySeg106M	106
CivilAirliner	20
DrivAer128M	128
VTM178M	178
LeMans1002M	1,002

Among the benchmark suite, LeMans1002M represents a large-scale production CFD workload containing approximately one billion cells. Due to memory requirements, the benchmark was executed on four-node and eight-node configurations, which represent the minimum practical cluster sizes for this model.

Solver Performance

Table 3 summarizes both the Total Elapsed Time and Average Time per Iteration measured for each benchmark workload.

Table 3. Simcenter STAR-CCM+ Solver Performance

Model	Performance Metrics	1 Node	2 Nodes	4 Nodes	8 Nodes
EmpHydroCyclone	Total Elapsed	3356.063	1566.399	664.508	304.997
	Avg. Iteration	108.011	50.302	21.123	9.509
LeMans104M	Total Elapsed	179.875	88.374	43.495	19.097
	Avg. Iteration	2.788	1.361	0.655	0.275
AeroSUVSteadyCoupled106M	Total Elapsed	198.446	113.957	65.293	43.828
	Avg. Iteration	7.339	3.795	1.764	0.948
AeroSUVSteadyCoupled322M	Total Elapsed	488.465	266.217	154.616	100.034
	Avg. Iteration	17.515	8.733	4.269	2.110
AeroSUVSteadySeg106M	Total Elapsed	84.553	48.096	30.607	21.859
	Avg. Iteration	2.704	1.319	0.616	0.292
AeroSUVSteadySeg322M	Total Elapsed	268.047	146.517	93.769	60.535
	Avg. Iteration	8.584	4.191	2.079	1.043
CivilAirliner	Total Elapsed	34.788	18.176	10.597	7.576
	Avg. Iteration	1.015	0.479	0.230	0.114
DrivAer128M	Total Elapsed	413.667	218.641	106.346	54.624
	Avg. Iteration	12.782	6.478	2.903	1.338
VTM178M	Total Elapsed	213.517	122.809	77.203	57.661
	Avg. Iteration	5.890	2.983	1.502	0.736
LeMans1002M[1]	Total Elapsed	—	—	547.457	278.463
	Avg. Iteration	—	—	8.445	4.124

NOTE:

1. LeMans1002M was only benchmarked at four and eight nodes due to memory requirements.

The average time per iteration (T_{avg}) is plotted in Figure 1. The logarithmic y-axis enables direct comparison across workloads with substantially different absolute iteration times. Lower values indicate better solver performance.

Figure 1. Average Time Per Iteration for Simcenter STAR-CCM+ Workloads Running On One, Two, Four, And Eight Compute Nodes

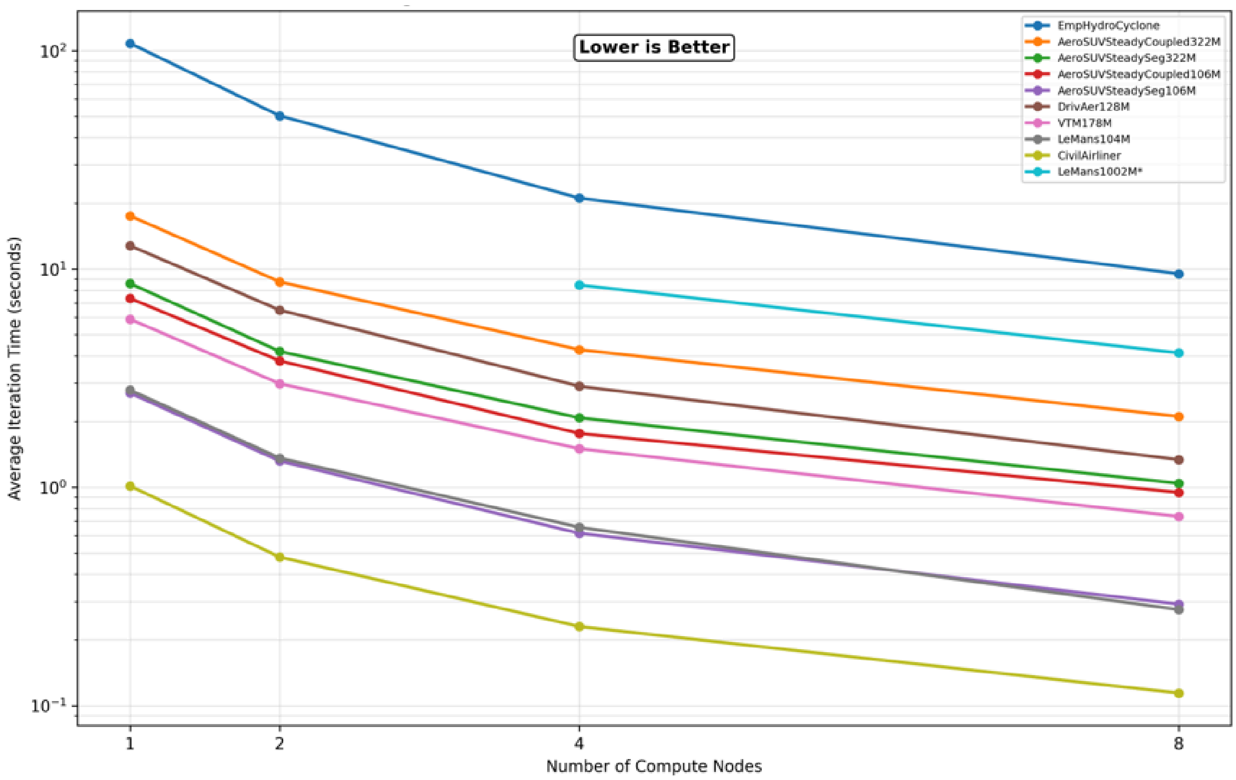

Across the benchmark suite, T_{avg} decreases consistently as node count increases, demonstrating strong scaling on AMD EPYC 9755 processors interconnected by the Cornelis CN5000 fabric. This trend is observed across a diverse set of production CFD workloads, ranging from approximately 20 million to more than one billion cells, including steady-state, coupled, segregated, aerospace, automotive, and multiphase simulations. Larger communication-intensive workloads, such as EmpHydroCyclone, DrivAer128M, and the LeMans benchmarks, achieve substantial reductions in T_{avg}, illustrating the platform's ability to efficiently scale demanding engineering simulations while minimizing communication overhead. LeMans1002M was evaluated at four and eight nodes because of its memory capacity requirements.

Speedup Relative to Single-Node Performance

Speedup measures how much faster a simulation executes relative to a single-node baseline. Because average time per iteration, T_{avg}, is the primary performance metric used in this study, speedup is calculated as the ratio of T_{avg} at 1 node to T_{avg} at each larger node count. A speedup equal to the increase in node count represents ideal linear scaling.

To provide a workload-independent view of overall application scalability, the geometric mean speedup was calculated across nine of the ten CFD workloads. The geometric mean is widely used in HPC performance analysis because it gives equal weight to each benchmark and prevents larger workloads from disproportionately influencing the overall result. LeMans1002M was excluded from the geometric mean calculation because the workload was benchmarked only on four-node and eight-node configurations due to memory capacity requirements. Since one-node and two-node baseline results were unavailable, speedup and parallel efficiency could not be calculated consistently with the remainder of the benchmark suite. For completeness, a 2.05× speedup relative to the four-node configuration is reported for the eight-node LeMans1002M benchmark, shown in Table 4.

Table 4. Simcenter STAR-CCM+ Speedup vs. One-Node Performance

Model	1 Node	2 Nodes	4 Nodes	8 Nodes
EmpHydroCyclone	1.00×	2.15×	5.11×	11.36×
AeroSUVSteadyCoupled322M	1.00×	2.01×	4.10×	8.30×
AeroSUVSteadySeg322M	1.00×	2.05×	4.13×	8.23×
AeroSUVSteadyCoupled106M	1.00×	1.93×	4.16×	7.74×
AeroSUVSteadySeg106M	1.00×	2.05×	4.39×	9.27×
DrivAer128M	1.00×	1.97×	4.40×	9.55×
VTM178M	1.00×	1.97×	3.92×	8.01×
LeMans104M	1.00×	2.05×	4.25×	10.13×
CivilAirliner	1.00×	2.12×	4.41×	8.88×
LeMans1002M[1]	—	—	1.00×	2.05×
Geometric Mean[2]	**1.00×**	**2.03×**	**4.31×**	**8.99×**

NOTES:

1. LeMans1002M was executed only on four-node and eight-node configurations because of memory capacity requirements. The reported 2.05× speedup at eight nodes is relative to the four-node baseline.
2. The Geometric Mean excludes LeMans1002M because one-node and two-node baseline measurements were unavailable.

The speedup plotted in Figure 2 was calculated using T_{avg}, with higher values indicating better parallel performance. Compared to a one-node baseline, the benchmark suite achieved a geometric mean speedup of 8.99× at eight nodes (112% efficiency), demonstrating near-linear application scaling across a diverse set of production CFD workloads.

$$Parallel\,Efficiency = \frac{Speedup}{Number\ of\ Nodes} \times 100\%$$

Figure 2. Speedup Relative to Single-Node Performance Across Simcenter STAR-CCM+ CFD Workloads, Including Geometric Mean Speedup and Ideal Linear Scaling

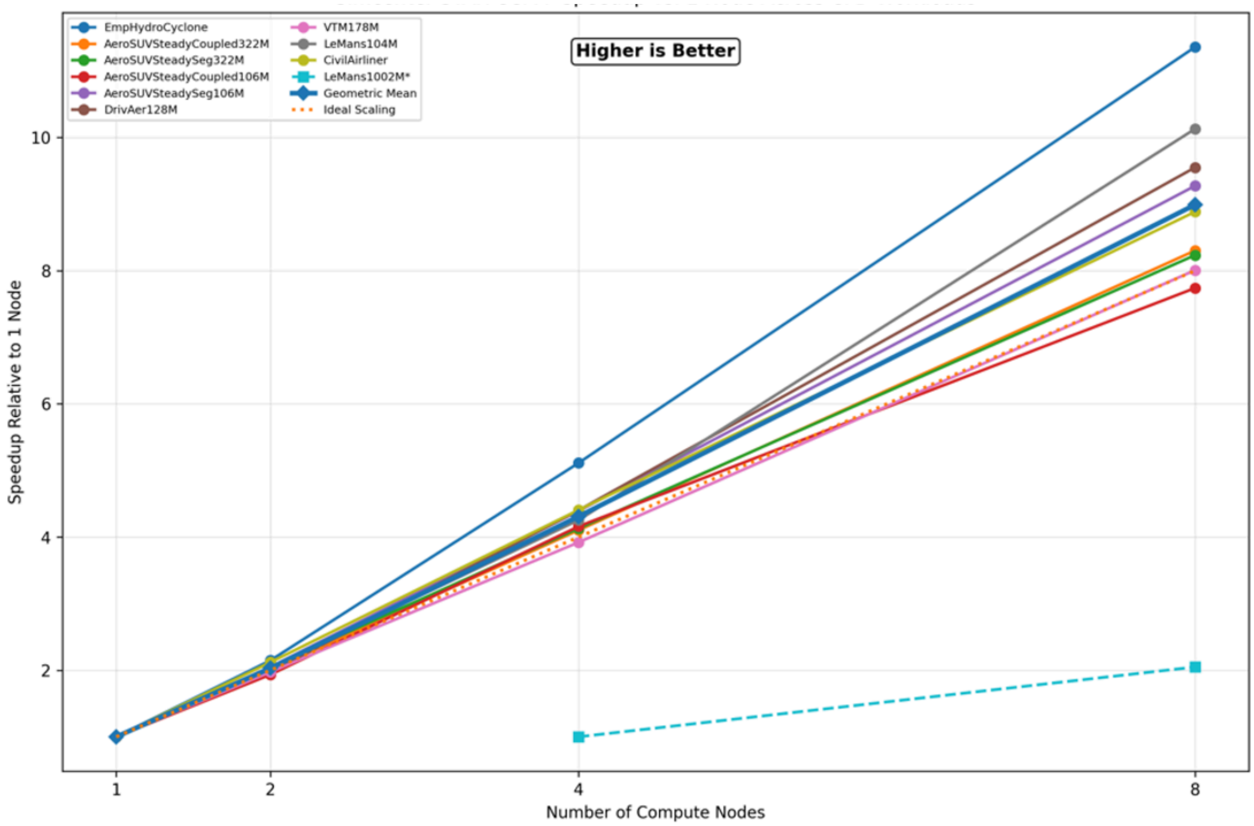

The corresponding parallel scaling efficiencies are shown in Table 5 computed as:

Table 5. Simcenter STAR-CCM+ Parallel Scaling Efficiency

Model	1 Node	2 Nodes	4 Nodes	8 Nodes
EmpHydroCyclone	100.0%	107.4%	127.8%	142.0%
AeroSUVSteadyCoupled322M	100.0%	100.3%	102.6%	103.8%
AeroSUVSteadySeg322M	100.0%	102.4%	103.2%	102.9%
AeroSUVSteadyCoupled106M	100.0%	96.7%	104.0%	96.7%
AeroSUVSteadySeg106M	100.0%	102.5%	109.8%	115.9%
DrivAer128M	100.0%	98.7%	110.1%	119.4%
VTM178M	100.0%	98.7%	98.0%	100.1%
LeMans104M	100.0%	102.4%	106.3%	126.6%
CivilAirliner	100.0%	105.9%	110.2%	111.1%
LeMans1002M[1]	—	—	100.0%	102.4%
Geometric Mean[2]	**100.0%**	**101.5%**	**107.8%**	**112.4%**

NOTES:

1. LeMans1002M was executed on four-node and eight-node configurations because of memory capacity requirements. The reported 2.05× speedup at eight nodes is relative to the four-node baseline.
3. The Geometric Mean excludes LeMans1002M because single-node and two-node baseline measurements were unavailable.

Super-linear scaling is commonly observed in large-scale CFD applications when additional compute resources improve cache utilization, memory locality, and solver efficiency. Siemens has previously highlighted the impact of CPU cache utilization on Simcenter STAR-CCM+ performance, demonstrating that improved cache residency can significantly reduce memory access overhead and accelerate CFD simulations. These effects contribute to the near-linear and super-linear scaling observed across several workloads in this benchmark suite, particularly EmpHydroCyclone, LeMans104M, DrivAer128M, and AeroSUVSteadySeg106M. For additional information on this topic, review the Siemens Simcenter blog, *CPU Cache and CFD – a Core Friendship*.

Conclusion and Key Findings

Simcenter STAR-CCM+ demonstrates excellent scalability on AMD EPYC 9755 processors connected through the Cornelis CN5000 fabric. Across a diverse set of automotive, aerospace, and industrial CFD workloads, the platform consistently delivers strong application scaling and efficient use of distributed computing resources.

Cornelis CN5000's ultra-low MPI latency, below 1 microsecond, and high message rates enable STAR-CCM+ to achieve outstanding parallel performance. Across all models, the system delivers a geometric mean speedup of 8.99×, corresponding to 112% parallel efficiency at eight nodes.

These results show that Simcenter STAR-CCM+, AMD EPYC 9755 processors, and the Cornelis CN5000 fabric together provide a highly scalable platform for production CFD simulations. The solution helps engineering teams reduce time to solution, improve resource utilization, and accelerate product development.

Solution Overview

Simcenter STAR-CCM+

Simcenter STAR-CCM+ is a leading CFD and multiphysics simulation platform used to model fluid flow, heat transfer, structural mechanics, and other coupled physical phenomena. Widely deployed across automotive, aerospace, energy, and industrial markets, Simcenter STAR-CCM+ enables engineers to accelerate product development through simulation-driven design.

For more information, visit the Siemens Simcenter STAR-CCM+ product page.

AMD EPYC 9005 Series Processors

The AMD EPYC 9005 Series processors are designed to deliver exceptional performance for HPC and AI workloads. Built on AMD's latest Zen 5 architecture, EPYC 9005 processors combine high core counts, increased memory bandwidth, PCIe Gen5 connectivity, and strong per-core performance to address the needs of modern simulation and scientific computing environments.

The AMD EPYC 9755 processor used in this study provides 128 cores per processor, enabling significant compute density while maintaining excellent performance for highly parallel engineering applications.

For more information, visit the AMD EPYC 9005 Series product page.

Cornelis CN5000 Omni-Path

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 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.

Figure 3. Cornelis CN5000 Omni-Path Portfolio

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+.

For more information, visit the [Cornelis CN5000 Omni-Path] product page.

Legal Disclaimer