



SUPERMICRO® SUPERBLADE® POWERED BY AMD EPYC™ PROCESSORS EXCELS FOR ANSYS SIMULATIONS

Balanced compute resources, memory bandwidth, and network bandwidth deliver outstanding performance and scalability for Ansys simulation software.



Contents

Executive Summary	1
Why run Ansys applications on HPC systems?	1
Ansys Simulation Applications	2
Supermicro SuperBlade: Where Performance Meets Efficiency for Today's Complex Applications	3
Supermicro SuperBlade for Ansys Applications	6
Running Ansys on AMD Powered SuperBlade Benefits:.....	16
Conclusion.....	16

Executive Summary

Supermicro continues our legacy, pushing the boundaries of compute, network, and storage to offer innovative solutions for today's complex workloads. Supermicro® SuperBlade® servers provide exceptional memory bandwidth, floating-point performance, scalability, and density to organizations needing an optimal balance for their technical computing workloads. These servers are optimized for High Performance Computing (HPC).

They are particularly useful for Ansys users who require fast processing speeds and large memory capacities for running complex simulations across various industries to help solve complex real-world problems. This white paper discusses Ansys simulation software performance benchmark results obtained on Supermicro SuperBlade servers powered by AMD EPYC™ 7003 Series Processors with AMD 3D V-Cache™ technology and AMD Instinct™ MI210 GPUs.

Why run Ansys applications on HPC systems?

Organizations running Ansys applications are investing in modern HPC infrastructure with highly performant processors to realize the full potential value of those applications. These investments help reduce constraints on simulation models' number, size, and complexity while delivering faster time to results. They also help engineers improve design quality and

prototype performance and have the potential to optimize the Total Cost of Ownership (TCO) by using fewer servers that consume less power compared to legacy solutions.

Ansys workloads challenge even modern systems for reasons that include:

- Inadequate processor frequency and/or core density that can require massive, often expensive scale-out solutions.
- Insufficient memory capacity and bandwidth and low ratios of cache per core which limit compute performance.
- Poorly optimized I/O.
- Lack of data security during computation.

A typical production environment has Ansys users running multiple parallel simulation jobs to fully utilize the 20 SuperBlade nodes.

Ansys Simulation Applications

We are witnessing a disruption in product development as the physical and digital worlds merge to generate unprecedented product innovation. By definition, engineers are a forward-looking group creating tomorrow's world today by continually pushing current boundaries to make the next discovery. Emerging trends such as 5G, autonomous vehicles, and electrification are radically changing the product development landscape, making it difficult for companies to retain engineers who can keep pace. Ansys simulation applications help engineers of all levels maximize productivity by intuitively providing the physics and materials properties needed to simulate complete product lifecycles. Engineers can test, iterate, and simulate before making the first prototype.

Supermicro, AMD, and Ansys have an ongoing technology partnership that delivers exceptional functionality and performance to our customers. Ansys software empowers innovators across industries to leverage the predictive power of simulation to drive technological leaps in areas such as sustainable transportation, advanced semiconductors, satellite systems, and life-saving medical devices.

Ansys® Fluent®

Ansys® Fluent® is a general-purpose computational fluid dynamics (CFD) and multi-physics tool that empowers engineers to optimize product performance further and faster than ever before. Fluent software contains the broad physical modeling capabilities needed to model flow, turbulence, heat transfer, and reactions for industrial applications. Fluent covers a broad reach, including unique capabilities for modeling in-cylinder combustion, aero-acoustics, turbomachinery, and multiphase systems.

Ansys® CFX®

Ansys® CFX® is a high-performance computational fluid dynamics (CFD) software tool that delivers reliable and accurate solutions quickly and robustly across a wide range of CFD and multi-physics applications. Ansys CFX is recognized for its outstanding accuracy, robustness, and speed with rotating machinery such as pumps, fans, compressors, and gas and hydraulic turbines.

Ansys® Mechanical™

Ansys® Mechanical™ is a finite element analysis (FEA) software application that models and simulates mechanical and structural designs. It offers a wide range of functionalities for stress analysis, thermal analysis, vibration analysis, and much more. Engineers use Ansys Mechanical to evaluate the structural integrity of a design, optimize its performance, and reduce the need for expensive physical testing.

Engineers can create complex 3D models and simulate the loads and stresses the design may encounter during use. It analyzes a wide range of materials, including composites and metals. Ansys Mechanical can also help the design optimization process by suggesting modifications to increase performance and reduce production costs. Ansys mechanical is widely used in industries such as aerospace, automotive, and manufacturing, where simulations can help predict the behavior of a design before it's built. This powerful tool can save time and money during the design and testing process while enhancing product safety and reliability.

Ansys® LS-DYNA®

Ansys® LS-DYNA® is a simulation application specifically designed to analyze dynamic events. It is widely used in the aerospace, automotive, and defense industries to study the behavior of complex structures and materials in extreme scenarios such as impacts, explosions, and crashes. This application can assess design safety by accurately predicting deformation, deformation rate, and event duration. Ansys LS-DYNA is based on the finite element method, which allows the creation of detailed models and provides detailed insight into how complex structures behave under dynamic loading. The tool is also highly customizable and can be used to predict how materials behave with complex material models, such as composites or metals with strain-rate dependence. Overall, Ansys LS-DYNA is a powerful simulation tool that is extensively used in industries where safety and reliability are critical considerations.

Supermicro SuperBlade: Where Performance Meets Efficiency for Today's Complex Applications

Compute requirements are increasing while data center space is not. Supermicro's new-generation blade portfolio helps optimize the TCO of key data center components, such as cooling, power efficiency, node density, and networking management. Supermicro SuperBlade servers powered by 3rd Gen AMD EPYC processors are built for the most demanding workloads that require high CPU density and the fastest networking available today in a trusted platform that meets enterprise customer demands for on-prem private/hybrid cloud deployments. This solution combines incredible core density with a rich feature set by marrying innovative server design with innovative CPU architecture to deliver the performance needed to run today's and tomorrow's workloads. Some key features include:

- **Form Factor:** 20 x 1P nodes in a single 8U rack enclosure with up to 1280 CPU cores.
- **Performance:** 3rd Gen AMD EPYC 7003 processors with up to 64 cores/128 threads and 768 MB of L3 cache per CPU, up to 2 TB of DDR4 memory, and two hot-swap drives (4 NVMe and 2 SATA3) per node for the most demanding workloads.
- **Density:** Incredibly high density with up to 20 hot-swappable nodes in an 8U form factor that includes flexible and robust 1G, 10G, 25G, or 100G Ethernet or 200G InfiniBand I/O options and up to 2 additional low-profile PCI-E 4.0 x16 expansion slots per node.

- **Efficiency:** Redundant 2200W titanium-level high-efficiency (96%) power supplies shared across all SuperBlade nodes provide full redundancy based on configuration and application load while helping optimize power usage and costs. Shared cooling further optimizes power consumption.
- **Accessibility:** OCP-compliant front IO.



Figure 1 - Supermicro SuperBlade 8U Enclosure with 20 blades



Figure 2 - SBA-4114S-C2N/T2N
OCP 3.0 Mezzanine Card



Figure 3 - SBA-4119SG
GPU/PCIe Cards

Flexible Configurations

Supermicro offers three SuperBlade models powered by AMD EPYC processors, which can be mixed and matched in a single 8U enclosure (as shown in Figure 1):

- **SAS/SATA:** SuperBlade SAS/SATA models (as shown in Figure 3) support AIOM for front I/O, which extends the Open Compute Project 3.0 specification to support a wide range of networking options in a small form factor.
- **GPU-accelerated:** Each 8U SuperBlade (as shown in Figure 2) can support up to 40 single width GPUs or 20 double width GPUs

Each SuperBlade enclosure contains at least one Chassis Management Module (CMM) for remote monitoring and managing server blades, power supplies, cooling fans, and networking switches. SuperCloud Composer (SCC) is a composable cloud management platform that provides a unified dashboard for administering software-defined data centers. SCC can orchestrate cloud workloads via the streamlined industry-standard Redfish API. SCC can also monitor and manage a broad portfolio of multi-generation Supermicro servers from a single pane of glass, including the Supermicro SuperBlade.

3rd Gen AMD EPYC processors with AMD 3D V-Cache Technology

The Supermicro SuperBlade is powered by AMD EPYC 7003 Series Processors with AMD 3D V-Cache technology. AMD EPYC 7003 Series Processors are built around the "Zen3" core and contain up to 64 cores per socket. AMD 3D V-Cache technology triples the total L3 cache from a maximum of 256 MB in standard 3rd Gen AMD EPYC processors to a maximum of 768 MB per socket. This extra-large cache works in concert with high frequencies, high core counts, and high memory bandwidth to unleash exceptional HPC performance. Supermicro and AMD have set multiple world records for performance.

3rd Gen AMD EPYC processors with AMD 3D V-Cache technology consist of up to eight Core Complex Dies (CCD) connected to a central IO die (see Figure 4 below). Each CCD includes 32MB of L3 cache, with a further 64MB of AMD 3D V-Cache technology layered on top for a total of 96MB of L3 cache per CCD and up to 768MB of L3 cache per CPU. Tripling the size of the L3 cache relative to standard 3rd Gen AMD EPYC processors can improve the performance of memory-bound applications by storing more of the data closer to the cores using that data, thereby reducing the need to access memory and/or storage.




Figure 4- 3rd Gen AMD EPYC with AMD 3D V-Cache Technology SOC Architecture


AMD Instinct MI210 GPUs for AI and HPC

The Supermicro SuperBlade SBA-4119SG supports 3rd Gen AMD EPYC processors with AMD 3D V-Cache technology and AMD Instinct MI210 GPUs. These GPU accelerated Supermicro SuperBlade servers are ideal for running converged AI Inference and HPC workloads.

Supermicro SuperBlade for Ansys Applications

Supermicro SuperBlade servers powered by 3rd Gen AMD EPYC processors with up to 64 cores per socket can deliver high throughput per node for Computational Fluid Dynamics (CFD) applications such as Ansys Fluent and Ansys CFX that benefit from multicore parallelism.

Crash applications such as Ansys LS-DYNA benefit from Supermicro SuperBlade systems with medium-core-count 3rd Gen AMD EPYC processors with high frequencies and high L3 cache-per-core that offer high performance while helping efficiently utilize per-core software licenses.



AMD EPYC 7003 Series Processors with AMD 3D V-Cache™ Technology
200G HDR InfiniBand Support
768 MB of Large L3 Cache per Node, Front I/O, OCP compliant
20x 1-socket nodes in an 8U Enclosure (1280 Cores) +
up to 2TB of Memory

Figure 5 - AMD powered Supermicro SuperBlade is the platform of Innovation for Ansys

SuperBlade Reference System Configuration for Benchmarks

Table 1 shows the system configuration used to test the performance benchmarks described in this whitepaper.

SUPERBLADE SERVER CONFIGURATION	
Compute nodes	Supermicro SuperBlade Servers SBA-4114S-T2N, SBA-4119SG
CPU	AMD EPYC 7473X (24 cores/48 threads, 2.8 GHz base frequency, 240W TDP) 3D V-Cache Technology
GPU	AMD Instinct MI210 (64GB, 300W TDP)
Memory	8x 64GB DDR4 3200 (512GB)
Network	SuperBlade InfiniBand (Mellanox CX-6) HDR 200Gb/s
Storage	1 x 1.8TB NVMe M.2
Software	Ansys v23.1
OS	RHEL 8.6 (4.18.0-372.9.1.el8)
Mellanox OFED Driver	MLNX_OFED_LINUX-5.8-1.1.2.1-rhel8.6-x86_64
MPI Version	Intel MPI 2021
Fluent version	23 release 1
OpenMPI Version	>4.0.1
WRF Version	>3.8.1
Compiled with ICC	2019.3.222
OS Settings Network	Transparent Huge Pages=disabled, Swappiness=disabled, Governor=Performance

Table 1 Supermicro SuperBlade Configuration

Ansys Fluent Benchmarks Overview and SuperBlade Performance

Measuring Ansys Fluent performance uses sets of benchmark problems selected to represent typical real-world scenarios. These benchmark cases range in size from a few hundred thousand cells to over 100 million cells. The suite contains both pressure-based (segregated and coupled) and density-based implicit solver cases using a variety of cell types and a range of physics.

Ansys Fluent Benchmark Description

Fluent Benchmark	Description
ICE_2m	• ~2 million mixed elements • K-e Turbulence model • Partially Premixed Combustion • Dynamic mesh model • Green Gauss for gradient calculations • Unsteady solver
Landing_gear_15 m	• ~15 million cells • LES, Acoustics model • Pressure based coupled solver

	- Least Squares cell based - Unsteady solver
Fluidized_bed_2m	- Circulating fluidized bed with ~2 million mixed elements - Smooth, multiphase model - Pressure based segregated solver - Least Squares cell based - Unsteady solver
F1_racecar_140m	- External flow over a Formula One racecar with ~140 million cells - Hex-Core - Realizable K-e Turbulence - Pressure based coupled solver - Least Squares cell based - Pseudo transient solver
aircraft wing_14m	- External flow over an aircraft wing with ~14 million hexahedral cells - K-e Turbulence model - Coupled implicit solver
Aircraft wing_2m	- External flow over an aircraft wing with ~2 million hexahedral cells - K-e Turbulence model - Coupled implicit solver
Open race care_280m	- External flow over an open-wheel racecar with ~280 million Hex-core cells - K-e Turbulence model - Pressure based coupled solver - Cell based - Pseudo transient solver
Sedan_4m	- External flow over a passenger sedan with ~4 million mixed elements - K-e turbulence model - Pressure based coupled solver
Combustor_12m	12 Million Polyhedral cells - K-e turbulence mode - Pressure based coupled solver - Least squares for Gradient Calculation - Pseudo transient
Combustor 71	71 Million Hex-core cells - LES for turbulence modeling - Pressure based coupled solver - Non premixed combustion - Least squares for Gradient calc - Unsteady
Rotor_3m	3 Million hexahedral cells - SST K-omega Turbulence model - Least squares for Gradient calc - Pressure based coupled solver - Sliding mesh interface - Pseudo transient

Exhaust _system_33m	• 33 Million mixed element cells • SST K-omega Turbulence model • Least Squares cell based • Steady state solver
Oil rig 7m	• 7 Million hexahedral cells • Volume of Fluid • SST K-omega Turbulence model • Pressure based segregated • Green-Gauss cell based for gradient calc • Unsteady solver

SuperBlade Ansys Fluent Single and Multi-Node Benchmark Results




Figure 6 - Strong super linear scaling as node count increases from 1 node to 16 nodes

Figure 6 shows a strong, superlinear scaling of Ansys Fluent performance with Supermicro SuperBlade powered by an AMD EPYC processor as the node count increases from 1 to 16. Adding more computational nodes to a technical computing cluster

reduces the portion of the dataset being processed by each node. Enough reduction allows each portion of the dataset to fit entirely within the L3 cache in each compute node, which causes a sudden performance boost called super linear scaling. Figure 6 shows that Supermicro SuperBlade can handle industrial scale simulations with large mesh sizes and deliver super-linear scaling as the node count increases.




Figure 7 - Ansys Fluent Scaling Performance from 1 to 16 Nodes

In Figure 7, the dotted line shows linear scaling, while the solid lines representing different test cases show super linear scaling boosts as the cluster expands from 1 to 16 nodes, which is directly attributable to the large 768 MB L3 cache offered by the AMD EPYC 7473X CPUs.

Ansys CFX Benchmarks Overview and SuperBlade Performance:

The Ansys CFX benchmark cases cover mesh sizes ranging from approximately 500 thousand to over 100 million nodes and include a variety of element types and physics models. Hardware vendors use these cases to benchmark their hardware systems, and the resulting performance data is then included in these pages.

Ansys CFX Benchmark Description

Benchmark	Description
Flow in an Automotive Pump	• 597,252 Nodes • Mixed tetrahedra, prisms, and pyramids • k-epsilon Turbulence, Multiple Frames of Reference • Coupled Implicit solver
External Flow Over a LeMans Car	• 1.8 Million tetrahedra • k-epsilon Turbulence, Heat Transfer • Coupled Implicit Solver
External Flow Over an Airfoil 10M	• 9.9 Million hexahedral • SST Turbulence, Heat Transfer • Coupled Implicit Solver
External Flow Over a50Mn Airfoil	• 48 Million hexahedral • SST Turbulence, Heat Transfer • Coupled Implicit Solver




Figure 8 – Ansys CFX Scaling Performance from 1 to 8 Nodes

Figure 8 shows the scaling performance of Supermicro SuperBlade powered by AMD EPYC 7473X CPUs from 1 to 8 nodes running Ansys CFX benchmarks. The "ansys_cfx_airfoil_10m" and "ansys_cfx_airfoil_100m" show superlinear scaling as the node count increases from one to four nodes. This scaling is efficient and super-linear for some of these large CFX benchmark models. The external flow over an airfoil 100M use case shows super-linear scaling at eight nodes.

Ansys Mechanical Benchmarks Overview and SuperBlade Performance:

Ansyes Mechanical is a finite element analysis (FEA) solver with structural, thermal, acoustics, transient, and nonlinear capabilities.-The Ansys Mechanical benchmark cases range from 650,000 to over 14 million degrees of freedom. The suite contains both PCG and Sparse equation solver benchmark cases covering a wide variety of analysis types, including static, transient, modal, cyclic symmetry modal, and harmonic. Both thermal and structural physics are also covered.

Ansyes Mechanical Benchmark Description

Benchmark	Description
Energy block	Static structural analysis of an engine block without the internal components
Gear box	Modal analysis of a transmission housing without the internal components
Peltier Cooling Box	Static nonlinear thermal-electric coupled field analysis of a Pelletier cooling block.
Power Supply Module	Steady state thermal analysis of a power supply module
Radial Impeller	Cyclic symmetric modal analysis of a single blade of an impeller
Semi-submersible	Transient nonlinear structural analysis of a submersible drilling rig
Speaker	Harmonic structural analysis of a speaker and its surroundings
Tractor Rear Axle	Static structural analysis of a farm tractor rear axle assemble
Turbine	Static nonlinear structural analysis of a turbine blade as found in aircraft engines
Engine	Static structural analysis of an engine block without the internal components
Power Supply Module	Steady state thermal analysis of a power supply module
Ball Grid Array	Transient nonlinear structural analysis of an electronic ball grid array




Figure 9 - SuperBlade for Ansys Mechanical Benchmark performance with and without AMD Instinct MI210 GPUs

Figure 9 shows the measured performance of Ansys Mechanical benchmarks on a single SuperBlade server powered by both a single 24-core AMD EPYC 7473X processor using multi-threading and an AMD Instinct™ MI210 GPU. This run was performed with eight threads per MPI rank across 8 SuperBlade nodes with and without MI210 GPUs. The AMD Instinct MI210 GPU delivers a significant performance boost of up to ~2-3x for Ansys Mechanical applications using both sparse direct and iterative solvers.

The benchmark results are presented in two panels: the left panel shows benchmarks with a direct sparse solver, and the right panel shows benchmarks with an iterative solver. Both show improvements when a GPU is added.

Ansys LS-DYNA Benchmarks Overview and SuperBlade Performance:

Ansyes LS-DYNA is a general-purpose finite element program used to simulate complex real-world problems, such as metals, composites, and soils, and can also simulate structures subject to impacts, explosions, and other extreme loading conditions. It is widely used in the automotive, aerospace, and manufacturing industries.

Ansyes LS-DYNA Benchmark Description:

Benchmark	Description
LS-DYNA Neon with distributed memory parallel simulation	Frontal crash with initial speed at 31.5 miles/hour. The model was created by the National Crash Analysis Center (NCAC) at George Washington University. One of the few publicly available models for vehicle crash analysis, based on the 1996 Plymouth Neon.




Figure 10 - Supermicro SuperBlade for Ansys LS-DYNA Benchmark Simulation Scales Linearly

Figure 10 demonstrates how Ansys LS-DYNA Neon performance scales up efficiently through eight SuperBlade nodes. The large L3 cache of the AMD EPYC 7003 processor is shared among eight cores. For hybrid (MPI+OpenMP) applications, the best performance is achieved by adjusting the number of threads per rank to 4 or 8 by assigning one (or two) MPI ranks per CCD and 4 (or 8) threads per rank. The LS-DYNA Neon benchmark demonstrates better scaling when multi-threads are enabled. Multi-threads enable large memory required by explicit solvers used in the models. The figure also shows close to linear scaling from 2 to 8 nodes.

For the Neon dataset, Supermicro, with its low latency and high bandwidth network, showed scaling up to eight nodes for the given size of the benchmark data.

Most industrial jobs perform parametric studies to explore the design space. The throughput of the system is a crucial factor in the design process. SuperBlade is equipped with 20 nodes. It is possible to deploy several instances of the same application simultaneously to achieve high throughput and speed up the analysis.

With SuperBlade, Ansys users can run simultaneous jobs on multiple nodes in parallel and leverage the performance features of all 20 nodes. SuperBlade is so flexible that an Ansys user can launch, say, a four-node parallel job on four nodes and simultaneously two 8-node parallel jobs on the remaining 16 nodes to utilize all 20 nodes of the SuperBlade. Alternatively, the users can run five 4-node parallel jobs on the 20 nodes or ten 2-node parallel jobs on the 20 nodes. This shows that SuperBlade gives flexibility to the Ansys Users to run parallel jobs and gives the best performance and throughput required for a shorter design cycle.

Running Ansys on AMD Powered SuperBlade Benefits:

- Significant per-core performance: AMD EPYC™ 7003 series processors offer high frequencies and large L3 caches and support up to 2TB of DDR4 3200 RAM and up to 128 lanes of PCIe® Gen 4 I/O that deliver superb performance for Ansys simulations.
- Each Supermicro SuperBlade includes up to 20 nodes with up to 1280 processor cores in a compact 8U rackmount enclosure.
- Supermicro SuperBlade servers equipped with AMD Instinct MI210 GPUs deliver blazing performance, simulating large mechanical models using the sparse direct solver.
- 200G InfiniBand offers low latency and high bandwidth to help boost multi-node performance.
- Supermicro's SuperBlade can deliver super-linear scaling performance on some Ansys workloads.
- Ansys users can run parallel jobs across the SuperBlade 20 nodes in a single 8U enclosure for optimum performance and throughput.
- Reduce Platform Complexity: Remove bottlenecks and replace outdated hardware with our building block solutions and the high performing Supermicro SuperBlade.
- Perform complex structural simulations that drive higher quality, more efficient designs for cars, planes, and other products while meeting tight deadlines.

Conclusion

Modern meshing, analytics, and simulation applications are demanding workloads that require balanced memory bandwidth, floating-point performance, and network I/O. Supermicro SuperBlade, powered by 3rd Gen AMD EPYC processors with AMD 3D V-Cache technology and AMD Instinct GPUs, delivers an optimal architecture and excellent performance for Ansys Fluent, Mechanical, CFX, and LS-DYNA that helps organizations optimize costs while improving product quality and accelerating time to market. All of this is available in a single 8U rackmount enclosure with a high core density that optimizes the data center footprint while optimizing power and cooling requirements, all of which can reduce TCO.

For More Information

www.supermicro.com/superblade

SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements. Visit www.supermicro.com

ANSYS

Our Mission: Powering Innovation that Drives Human Advancement™

When visionary companies need to know how their world-changing ideas will perform, they close the gap between design and reality with Ansys simulation. For more than 50 years, Ansys software has enabled innovators across industries to push boundaries by using the predictive power of simulation. From sustainable transportation to advanced semiconductors, from satellite systems to life-saving medical devices, the next great leaps in human advancement will be powered by Ansys. Visit www.ansys.com

AMD

For more than 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies. Billions of people, leading Fortune 500 businesses and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work and play. AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit www.amd.com