

Architecting High Performance Ansys CAE Solutions with Intel Xeon 6 and Lenovo ThinkSystem Servers

Planning / Implementation

Importance of HPC in CAE

Computer-Aided Engineering (CAE) has become a foundational capability for organizations developing complex products across industries such as automotive, aerospace, energy, and manufacturing. Advances in simulation fidelity—driven by finer meshes, nonlinear material models, transient physics, and increasingly coupled multiphysics workflows—have enabled engineers to evaluate designs earlier and more comprehensively in the development cycle.

These advances, however, have significantly increased the computational demands placed on underlying infrastructure. Modern CAE workloads require high-performance computing (HPC) platforms that can deliver sustained performance, high memory bandwidth, and predictable scalability rather than peak theoretical capability alone. As simulation sizes grow and solver concurrency increases, limitations in memory subsystem performance, power delivery, and thermal management increasingly determine achievable time-to-solution.

As a result, HPC is no longer reserved for only the largest simulations or centralized clusters; it has become an essential component of day-to-day CAE workflows. Selecting and architecting HPC platforms that align with solver behavior is critical to maximizing simulation throughput, improving engineering productivity, and ensuring that infrastructure investments scale effectively as CAE workloads continue to evolve.

Synopsys: A Major Player in the CAE Landscape and HPC

Through the acquisition of Ansys, Synopsys now complements its leadership in Electronic Design Automation (EDA) with one of the world's most comprehensive CAE portfolios, offering industry-leading physics-based simulation tools across computational fluid dynamics, structural and explicit dynamics, electromagnetics, optics, photonics, and multiphysics analysis. Ansys CAE applications from Synopsys span a broad range of numerical methods and solver behaviors, each placing different demands on compute, memory, and system architecture. While these solvers are often deployed within the same CAE environment, their performance characteristics vary substantially depending on the underlying physics, discretization approach, model resolution, and parallelization strategy. Understanding these differences is essential for interpreting benchmark results and for selecting processor, memory, and platform configurations that align with specific workload requirements.

Topics in this section:

- [Fluent – Computational Fluid Dynamics Software](#)
- [LS-DYNA – Explicit Finite Element Analysis Software](#)
- [Mechanical – Implicit Finite Element Analysis Software](#)

Fluent – Computational Fluid Dynamics Software

Fluent is a general-purpose computational fluid dynamics (CFD) software package used to model fluid flow, heat transfer, turbulence, and related physical phenomena across a wide range of industries. Fluent simulations commonly employ structured or unstructured meshes ranging from millions to hundreds of millions of cells, with solution methods that rely on iterative solvers, sparse linear algebra, and frequent data movement through the memory hierarchy.

From a system perspective, Fluent software exhibits a balanced performance profile influenced by CPU frequency, core count, cache capacity, and memory bandwidth. Smaller and moderately sized models tend to benefit from strong per-core performance and sustained CPU frequency, while larger models increasingly become constrained by aggregate memory bandwidth as working sets exceed cache capacity. As a result, increases in core count alone do not guarantee proportional performance gains unless sufficient memory throughput is available to keep cores fully utilized.

Fluent software generally scales efficiently within a single node and across multiple nodes when memory bandwidth and interconnect performance are well balanced. At higher core counts, the ratio of MPI ranks to available memory bandwidth per core becomes a critical tuning parameter. Platforms that provide higher aggregate memory bandwidth—through additional memory channels or higher-speed memory technologies—tend to deliver more consistent and predictable performance improvements for large-scale Fluent workloads.

LS-DYNA – Explicit Finite Element Analysis Software

LS-DYNA is an explicit finite element analysis (FEA) software package widely used for highly dynamic simulations such as crashworthiness, impact, drop tests, and forming processes. Besides an implicit solver, LS-DYNA can also advance the solution using explicit time integration with very small time steps, resulting in a large number of repeated computations over the course of a simulation.

From a computational perspective, LS-DYNA explicit workloads are dominated by sustained arithmetic throughput and parallel efficiency rather than complex global matrix operations. Because each time step performs relatively simple calculations that are executed repeatedly, LS-DYNA software benefits strongly from increased core counts and the ability to maintain stable all-core frequency over long runtimes. As a result, explicit dynamics workloads often exhibit predictable scaling behavior as computational resources increase.

Memory bandwidth remains an important consideration for large LS-DYNA models with high element counts, but explicit solvers are generally less sensitive to memory latency than implicit FEA workloads. Platform characteristics such as thermal management and power delivery therefore play a critical role in determining achievable performance, as simulations often run at high utilization for extended periods with minimal idle time.

Mechanical – Implicit Finite Element Analysis Software

Mechanical is an implicit FEA software package used for structural, thermal, and coupled multiphysics simulations, including static and transient structural response, contact, and nonlinear material behavior. Many Mechanical workloads rely heavily on sparse linear algebra, where performance is shaped by how efficiently the system can assemble and solve large sparse systems of equations.

From a system perspective, Mechanical performance is often dominated by memory behavior—particularly memory bandwidth, latency, and cache efficiency—rather than raw floating-point throughput. As model sizes grow and working sets exceed cache capacity, solver phases such as matrix assembly and solution can become increasingly constrained by data movement and synchronization overhead, leading to diminishing returns from additional cores unless the memory subsystem scales accordingly.

Mechanical performance also depends strongly on the linear solver strategy. Direct solvers (factorization-based) tend to be more memory-intensive and can become bandwidth-limited, while iterative solvers (preconditioned Krylov methods) reduce memory footprint but introduce sensitivity to convergence behavior and global synchronization. Sustained CPU frequency under continuous load remains important for both approaches, particularly for nonlinear analyses that require repeated solution phases.

Implications for System Architecture

The distinct computational characteristics of Fluent, Mechanical, and LS-DYNA underscore the need for a balanced, solver-aware system architecture. No single hardware configuration is optimal across all CAE workloads: CFD and explicit dynamics benefit from increased core density and memory bandwidth, while implicit FEA places greater emphasis on memory behavior, cache efficiency, and sustained per-core performance.

These differences highlight why processor architecture, memory technology, and platform-level design must be considered together rather than in isolation.

As CAE workloads scale, limitations related to memory bandwidth, power delivery, and thermal management increasingly determine achievable performance and scalability. Effective CAE infrastructure therefore requires platforms that can expose processor capabilities consistently under real-world solver behavior, not just peak theoretical specifications.

Topics in this section:

- [Intel Xeon Processors: The CAE Compute Foundation](#)
- [Memory Technology as a First Order Performance Factor](#)
- [Lenovo ThinkSystem Servers for Ansys CAE Solutions](#)

Intel Xeon Processors: The CAE Compute Foundation

Intel® Xeon® processors have long served as the computational workhorse for CAE workloads, providing a balanced combination of core density, sustained performance, memory scalability, and a mature HPC software ecosystem. This balance has made Xeon processors a dependable foundation for engineering simulation across a wide range of use cases, from departmental CAE environments to large-scale HPC deployments.

Intel Xeon 5th Generation processors represent the established baseline for many CAE infrastructures, delivering balanced performance across CFD, implicit FEA, and explicit dynamics workloads.

Intel Xeon 6 processors extend this foundation with architectural enhancements that increase core density, improve sustained all-core performance, and expand memory bandwidth capabilities, enabling higher performance as solver concurrency and model complexity increase.

The Xeon 6 family spans SP-class processors optimized for balanced deployments, and AP-class processors designed for maximum throughput in the most demanding CAE workloads. Rather than focusing on product specifications in isolation, CAE performance is best understood through the architectural attributes that most directly influence solver behavior: core count and sustained frequency, memory bandwidth and latency, and cache capacity.

These characteristics determine how effectively solvers utilize available compute resources and define the requirements that system platforms must satisfy to deliver consistent, production-grade performance under real-world CAE workloads.

Memory Technology as a First Order Performance Factor

As CAE solvers continue to scale in model size and concurrency, memory subsystem performance has become a first-order design consideration rather than a secondary optimization.

For many modern workloads, particularly CFD and large finite element analyses, solver execution is increasingly dominated by data movement rather than floating-point computation. In these regimes, limitations in memory bandwidth and latency can constrain performance even when sufficient compute resources are available. Traditional DDR5 RDIMM configurations provide a balanced and widely deployed

solution, but may become a bottleneck as core counts increase, and working sets exceed cache capacity.

High-bandwidth memory technologies are therefore critical to sustaining performance as solver concurrency grows. Without sufficient memory throughput, additional cores cannot be effectively utilized, leading to diminishing returns from processor scaling alone. Multiplexed Rank DIMMs (MRDIMMs) address this challenge by increasing effective memory bandwidth per channel and reducing contention as core counts rise.

For memory-intensive CAE workloads—particularly large Fluent models and selected LS-DYNA cases—higher memory throughput can deliver performance gains that rival or exceed those achieved through increases in CPU frequency or core count. The impact of MRDIMM is workload-dependent but becomes increasingly significant as model sizes grow and memory traffic dominates solver execution.

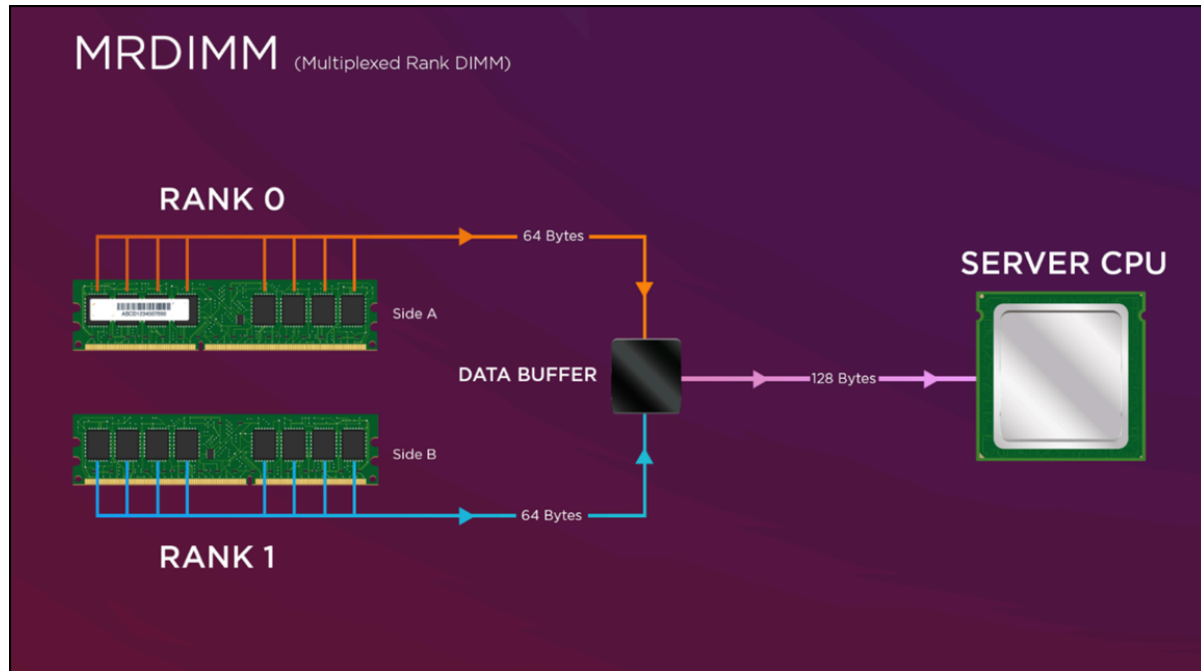


Figure 1. MRDIMM Multiplex Functionality

Lenovo ThinkSystem Servers for Ansys CAE Solutions

Ansys CAE solvers, Intel® Xeon® processors, and high-bandwidth memory technologies together define the computational foundation for modern engineering simulation. Realizing their full potential, however, requires a system platform engineered to sustain performance under real-world solver behavior rather than peak theoretical conditions.

Lenovo ThinkSystem servers integrate compute, memory, cooling, and power delivery into a balanced platform, enabling a complete, production-ready solution for Ansys CAE workloads ranging from single-node simulations to large-scale, memory-intensive HPC environments.

Topics in this section:

- [ThinkSystem SR630 V4 Server: The Traditional CAE Foundation](#)
- [Transitioning Direct Water-Cooling Technology with ThinkSystem SC750 V4 Servers](#)

ThinkSystem SR630 V4 Server: The Traditional CAE Foundation

Ansys CAE solvers, Intel® Xeon® processors, and high-bandwidth memory technologies together define the computational foundation for modern engineering simulation. Realizing their full potential, however, requires a system platform engineered to sustain performance under real-world solver behavior rather than

peak theoretical conditions.

Lenovo ThinkSystem servers integrate compute, memory, cooling, and power delivery into a balanced platform, enabling a complete, production-ready solution for Ansys CAE workloads ranging from single-node simulations to large-scale, memory-intensive HPC environments.



Figure 2. Lenovo ThinkSystem SR630 V4 Server – 2-socket, 1U, Air-Cooled Rack Server

Transitioning Direct Water-Cooling Technology with ThinkSystem SC750 V4 Servers

As CAE workloads scale in model size, solver concurrency, and memory bandwidth demand, processor power envelopes and thermal requirements increasingly exceed what air-cooled platforms can consistently sustain.

The Lenovo ThinkSystem SC750 V4 addresses these challenges through direct water cooling (DWC), removing thermal constraints that can otherwise limit sustained CPU frequency and memory performance under continuous load. By enabling stable operation of high-core-count Intel Xeon 6 AP-class processors and high-bandwidth memory configurations such as MRDIMMs, the SC750 V4 allows systems to operate closer to their architectural limits for extended runtimes.

This capability is especially critical for large-scale CFD, implicit FEA, and explicit dynamics workloads, where sustained performance directly translates to improved scalability and reduced time-to-solution.



Figure 3. Lenovo ThinkSystem N1380 Neptune Chassis – 13U8T Enclosure of 19” Rack Cabinets




Figure 4. Lenovo ThinkSystem SC750 V4 Neptune Server Tray

Performance Results

The performance results presented in this section are based on application-level benchmarking of Fluent, Mechanical, and LS-DYNA software using representative workloads and production-class server configurations.

Testing was conducted on Lenovo ThinkSystem platforms powered by Intel® Xeon® 5th and 6th Generation processors to evaluate generation-over-generation performance, platform behavior, the impact of memory technology, and scaling characteristics under realistic CAE usage scenarios.

Topics in this section:

- [Test Methodology](#)
- [Benchmark Systems](#)
- [Fluent Software Performance Results](#)
- [LS-DYNA Software Performance Results](#)
- [Mechanical Software Performance Results](#)

Test Methodology

Benchmarks were designed to reflect common CAE deployment models and include both single-node performance measurements and multi-node scaling studies. Single-node testing focuses on isolating CPU and memory subsystem behavior, highlighting the effects of processor architecture, core count, sustained frequency, cache capacity, and memory bandwidth without the influence of network variability.

Multi-node scaling tests evaluate how solver performance evolves as workloads are distributed across multiple compute nodes, characterizing parallel efficiency, communication behavior, and sensitivity to system balance as aggregate core counts and memory bandwidth increase.

All systems were configured with consistent software environments, including operating system, compilers, MPI libraries, and Ansys solver versions, to ensure comparability across platforms. Performance results are reported as relative performance, normalized to an Intel Xeon 5th Generation baseline configuration. This approach highlights architectural and platform differences while avoiding dependence on absolute timing values, which can vary with model configuration and solver settings.

Benchmark Systems

The benchmark systems are detailed in the table below.

Table 1. Specification of Benchmark Systems

Processor	Intel Xeon 5th Gen 8562Y+	Intel Xeon 5th Gen 8592+	Intel Xeon 6732P	Intel Xeon 6960P	Intel Xeon 6972P	Intel Xeon 6980P
Core per socket	32	64	32	72	96	128
Max turbo frequency	4.1GHz	3.9GHz	4.3GHz	3.9GHz	3.9GHz	3.9GHz
Base frequency	2.8GHz	1.9GHz	3.8GHz	2.7GHz	2.4GHz	2.0GHz
Cache	60MB	320MB	144MB	432MB	480MB	504MB
Thermal design power (TDP)	300W	350W	350W	500W	500W	500W
Ultra Path Interconnect (UPI) speed	20GT/s	20GT/s	24GT/s	24GT/s	24GT/s	24GT/s
Memory Channels	8	8	8	12	12	12
Memory Type	DDR5	DDR5	DDR5	DDR5	DDR5	DDR5 & MRDIMM
Memory Speed	5600MT/s	5600MT/s	6400MT/s	6400MT/s	6400MT/s	6400MT/s & 8800 MT/s
Installed Memory	512GB	512GB	1024GB	768GB	768GB	1536GB & 768GB
ThinkSystem Server(s)	SD650 V3	SD650 V3	SR630 V4	SC750 V4	SC750 V4	SC750 V4
Network Interconnect	NVIDIA NDR400	NVIDIA NDR400	NVIDIA NDR400	NVIDIA NDR400	NVIDIA NDR400	NVIDIA NDR400

The benchmark systems used the following configuration:

- Operating System: Rocky Linux 9.6
- Kernel Level: 5.14.0-570.58.1.el9_6.x86_64
- Intel OneAPI: 2025.2.1
- Ansys: 2025.2
- LS-DYNA: 16.1.1

Fluent Software Performance Results

Fluent workloads represent a broad class of CFD simulations with varying computational characteristics depending on model size, mesh resolution, and physical complexity. Fluent exhibits a balanced dependence on CPU frequency, core count, cache capacity, and memory bandwidth. The following results illustrate how these factors interact across Intel Xeon 5 and Intel Xeon 6 platforms.

Topics in this section:

- [Single Node Performance](#)
- [Impact of Memory Configuration](#)
- [Multi-Node Scaling Performance](#)

Single Node Performance

Single-node Fluent software performance results highlight the combined impact of processor architecture and memory subsystem capability. For smaller and moderately sized models, performance improvements are driven primarily by per-core efficiency and sustained CPU frequency, while larger models increasingly become constrained by aggregate memory bandwidth.

Across the evaluated configurations, Intel Xeon® 6 SP-class processors deliver approximately 8–12% higher per-core performance compared to Intel Xeon® 5th Generation processors for Fluent workloads (Figure 6). This improvement reflects architectural gains in IPC, cache behavior, and sustained frequency rather than increases in core count alone.

Moving to Xeon® 6 AP-class processors, single-node Fluent software performance increases by approximately 60–110% relative to Intel Xeon® 5th Generation platforms, depending on model size (Figure 7). These gains are driven by a combination of higher core density, substantially larger cache capacity, and increased aggregate memory bandwidth. Larger CFD models show the highest uplift, reinforcing the sensitivity of Fluent software to memory throughput as working sets exceed cache capacity.

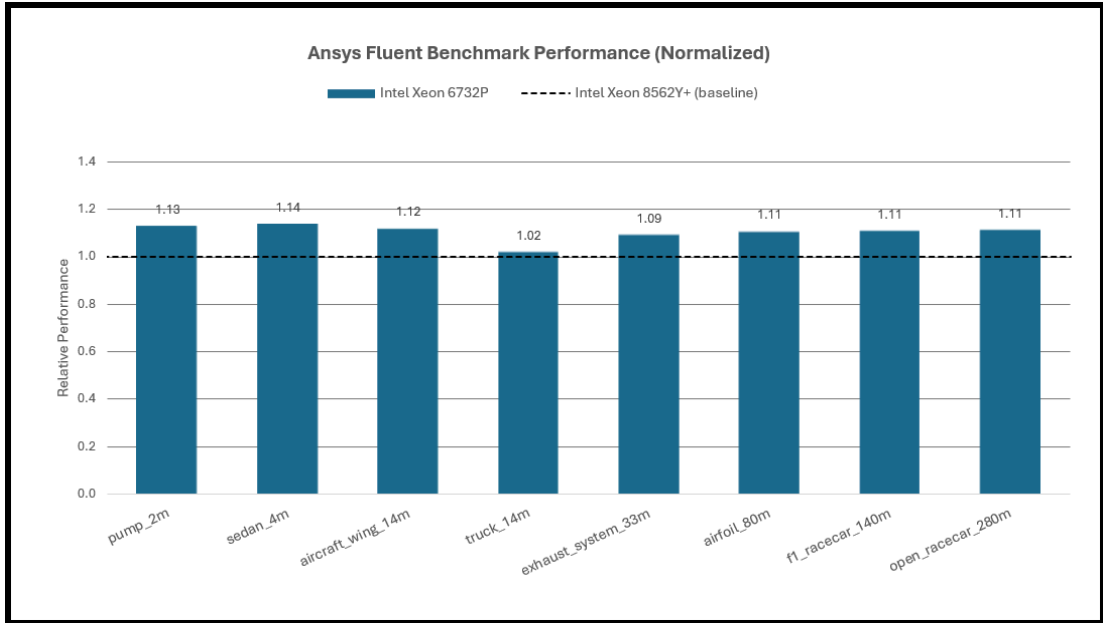


Figure 5. Ansys Fluent Performance: Intel Xeon 5th Gen Processor & 6th Gen SP Processor

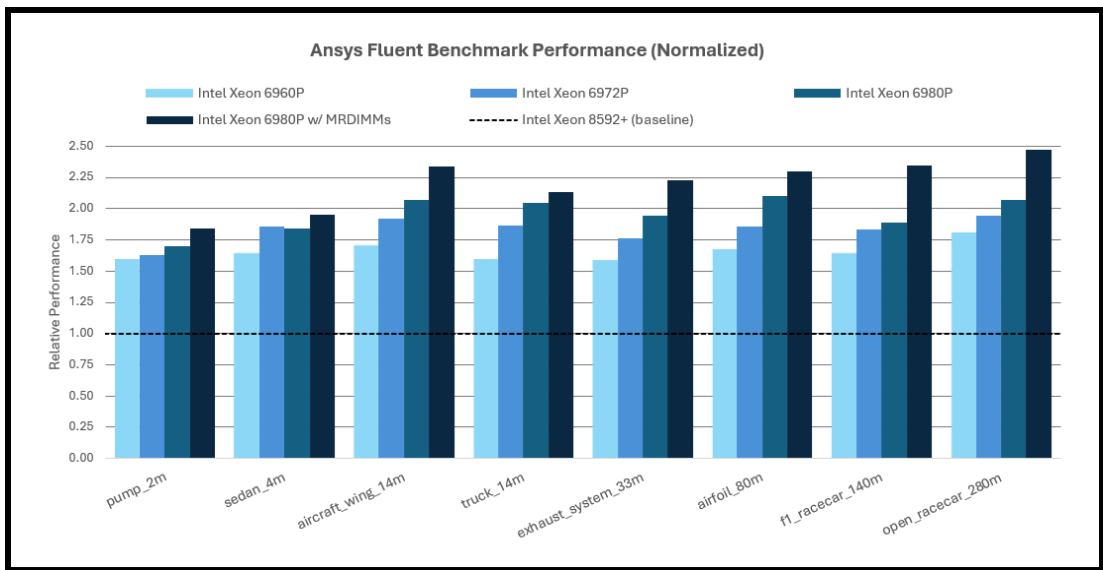


Figure 6. Ansys Fluent Performance: Intel Xeon 5th Gen Processor & 6th Gen AP Processor

Impact of Memory Configuration

Memory configuration plays a significant role in single-node Fluent software performance, particularly for large and memory-intensive models. When comparing DDR5 RDIMM and MRDIMM configurations on Intel Xeon® 6 AP platforms, MRDIMMs deliver an additional 6–10% improvement for smaller Fluent cases and up to 20–25% improvement for the largest models.

This behavior reflects the increasing reliance of Fluent software on memory bandwidth as mesh sizes grow. For smaller workloads, performance remains partially cache-resident and sees limited benefit from higher memory throughput, while large CFD cases experience substantial gains once memory traffic becomes the dominant limiting factor.




Figure 7. Ansys Fluent Performance: Impact of Memory Technology

Multi-Node Scaling Performance

Multi-node scaling results demonstrate the ability of Fluent software to efficiently distribute large CFD workloads across multiple compute nodes when sufficient compute and memory resources are available. For the largest models evaluated, Fluent software exhibits near-linear scaling up to 4,096 cores, achieving approximately 15× performance improvement when scaling from 256 to 4,096 cores.

These results highlight the suitability of Fluent software for clustered HPC environments and underscore the importance of balanced node design. High per-node performance combined with high-bandwidth memory and a low-latency interconnect enables Fluent software to maintain strong parallel efficiency at scale, directly translating into reduced time-to-solution for large CFD simulations.




Figure 8. Ansys Fluent Performance: Scaling on Intel Xeon6 AP Processors

LS-DYNA Software Performance Results

LS-DYNA software advances the solution through very small time steps using explicit integration methods, resulting in a large number of repeated computational operations. These characteristics make LS-DYNA particularly sensitive to core count, sustained CPU throughput, and parallel efficiency, while still benefiting from adequate memory bandwidth for large models.

Topics in this section:

- [Single Node Performance](#)
- [Impact of Memory Configuration](#)
- [Multi-Node Scaling Performance](#)

Single Node Performance

Single-node LS-DYNA software performance is driven primarily by available core count and sustained all-core frequency. Because explicit solvers perform relatively simple calculations per time step but execute them repeatedly, throughput scales well with additional cores as long as processors can maintain stable operating frequencies under continuous load.

Across the evaluated configurations, Intel Xeon® 6 SP-class processors deliver approximately 10–15% higher per-core performance compared to Intel Xeon® 5th Generation processors for LS-DYNA workloads. This improvement reflects architectural efficiency and sustained frequency gains rather than increases in core count.

Moving to Xeon® 6 AP-class processors, single-node LS-DYNA software performance increases by approximately 45–70% relative to Intel Xeon® 5th Generation platforms, depending on the workload. These gains are primarily driven by significantly higher core density and the ability to sustain performance under long-running, high-utilization simulations.




Figure 9. Ansys LS-DYNA Performance: Intel Xeon 5th Gen Processor & 6th Gen SP Processor

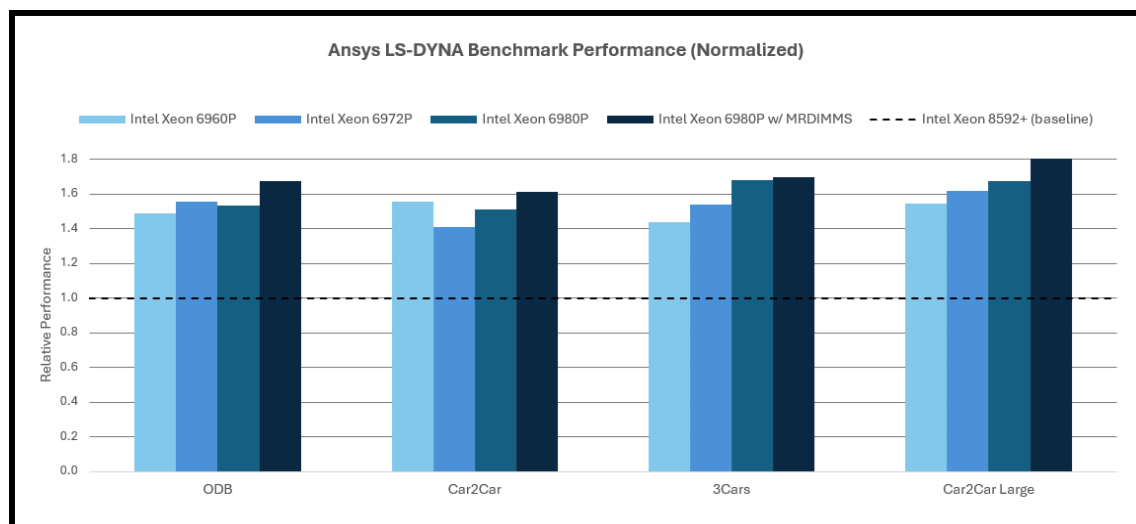


Figure 10. Ansys LS-DYNA Performance: Intel Xeon 5th Gen Processor & 6th Gen AP Processor

Impact of Memory Configuration

The impact of memory configuration on LS-DYNA software performance is typically secondary to core count and frequency but becomes increasingly relevant as model sizes grow. Larger explicit simulations generate higher memory traffic due to frequent state updates and data movement across elements.

When comparing DDR5 RDIMM and MRDIMM configurations on Intel Xeon®6 AP platforms, MRDIMMs provide an additional 3–6% improvement for smaller LS-DYNA cases and up to 8–10% improvement for the largest models. While these gains are more modest than those observed for CFD workloads, they confirm that memory bandwidth can still influence performance once explicit simulations reach sufficient scale.




Figure 11. Ansys LS-DYNA Performance: Impact of Memory Technology

Multi-Node Scaling Performance

The impact of memory configuration on LS-DYNA software performance is typically secondary to core count and frequency but becomes increasingly relevant as model sizes grow. Larger explicit simulations generate higher memory traffic due to frequent state updates and data movement across elements.

When comparing DDR5 RDIMM and MRDIMM configurations on Intel Xeon®6 AP platforms, MRDIMMs provide an additional 3–6% improvement for smaller LS-DYNA cases and up to 8–10% improvement for the largest models. While these gains are more modest than those observed for CFD workloads, they confirm that memory bandwidth can still influence performance once explicit simulations reach sufficient scale.

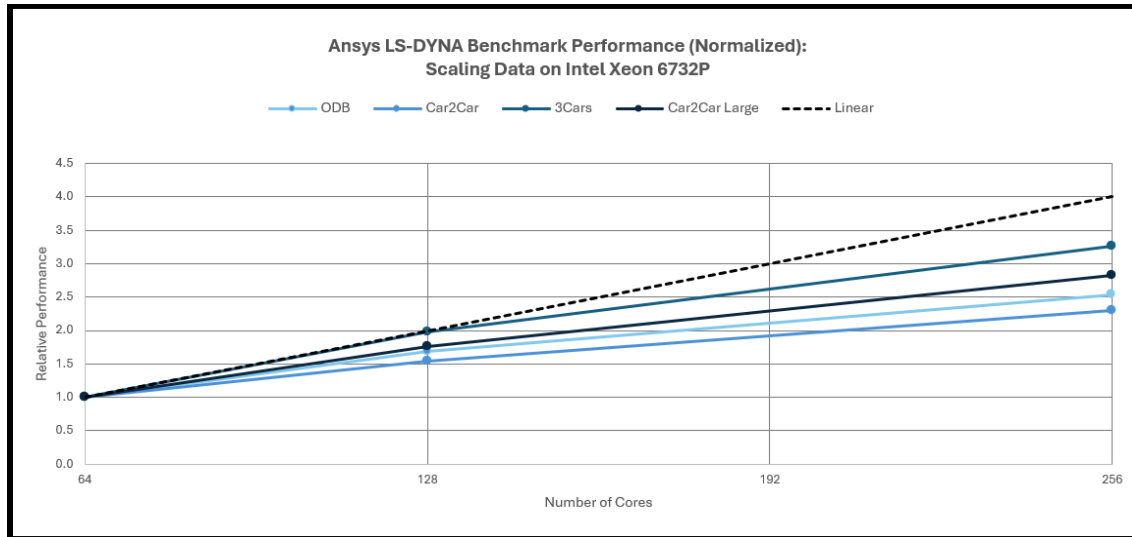


Figure 12. Ansys LS-DYNA Performance: Scaling on Intel Xeon6 AP Processors

Mechanical Software Performance Results

Mechanical workloads represent implicit FEA simulations whose performance is dominated by sparse linear algebra operations and solver synchronization behavior. As a result, Mechanical software performance is often governed by memory subsystem efficiency, cache behavior, and sustained CPU frequency, and it can exhibit diminishing returns from additional cores once memory or synchronization limits are reached.

Topics in this section:

- [Single Node Performance: Core-for-Core, SP-class](#)
- [Single Node Performance: Throughput and Scaling Effects, AP-class](#)

Single Node Performance: Core-for-Core, SP-class

The core-for-core comparison (64 cores/node vs 64 cores/node) between Intel Xeon 5th Gen 8562Y+ (baseline) and Intel Xeon 6 SP 6732P shows the following results. For the direct solver cases, the Xeon 6 SP processor shows very large gains:

- V25direct-4: ~+110%
- V25direct-5: ~+95%
- V25direct-6: 1~+13%

These results indicate that the direct solver benefits strongly from architectural improvements that increase effective memory throughput and cache efficiency. Direct solver test case named V25direct-6 has more performance uplift drop off because the problem size exceeds cache and stresses main memory more heavily and memory bandwidth is more saturated.

For the iterative solver cases, the same core-for-core comparison shows little to no per-core improvement, and in these specific tests even a slight reduction:

- V25iter-4: ~-3%
- V25iter-5: ~-5%
- V25iter-6: ~-7%

This contrast is consistent with solver fundamentals: direct solvers are typically dominated by memory bandwidth/cache efficiency, while iterative solvers are more limited by convergence behavior and synchronization overhead, which does not necessarily improve with per-core architectural uplift alone.

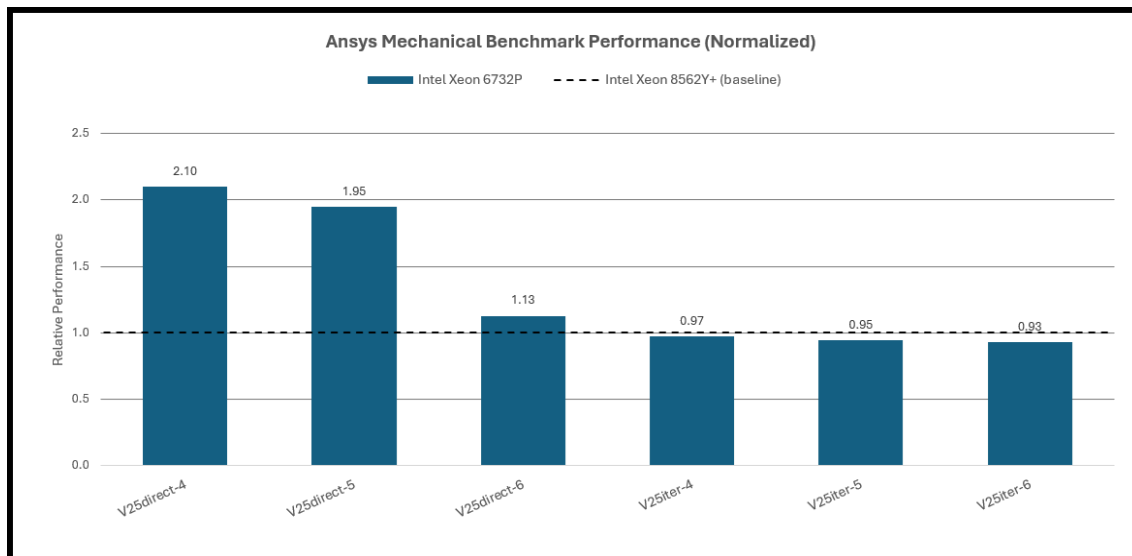


Figure 13. Ansys Mechanical Performance: Intel Xeon 5th Gen Processor and 6th Gen SP Processor

Single Node Performance: Throughput and Scaling Effects, AP-class

A performance comparison between Intel Xeon 6 AP-class processors (6960P / 6972P / 6980P) to an Intel Xeon 5 baseline (8592+) was conducted. These are not core-matched comparisons (core counts differ significantly), so results represent platform throughput gains rather than pure per-core uplift. For the direct solver cases, the Xeon 6 AP processors show very large gains:

- V25direct-4: ~+79% to ~+150%
- V25direct-5: ~+30% to ~+118%
- V25direct-6: ~+49% to ~+131%

These results show that the direct solver workloads in Mechanical software can realize very large throughput gains on Xeon 6 AP platforms, especially for the direct-4 and direct-5 cases. The lower relative uplift for direct-6 compared to direct-4/5 is consistent with a case that more heavily saturates main memory bandwidth and/or exceeds cache capacity, reducing the marginal benefit of additional architectural improvements. For the iterative solver cases, the Xeon 6 AP processors' improvements are much smaller and more variable than for direct solver workloads:

- V25iter-4: ~+32% to ~+51%
- V25iter-5: ~+14% to ~+22%
- V25iter-6: ~+6% to ~+23%

This is consistent with iterative solver behavior: while additional cores can increase throughput, overall efficiency is often constrained by convergence characteristics and synchronization (e.g., global reductions), which limits the per-core benefit and reduces scaling efficiency at higher core counts.

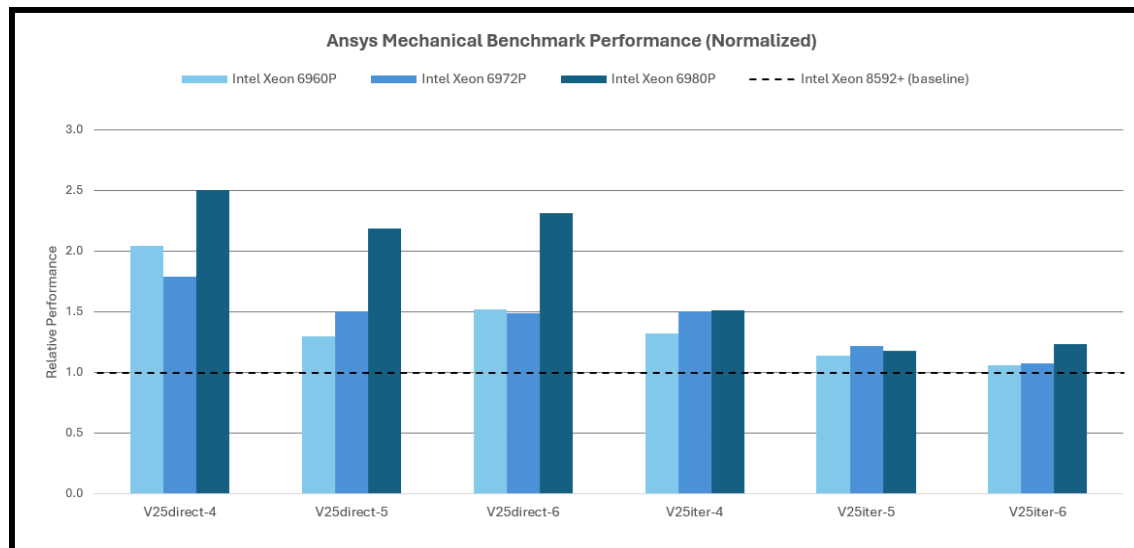


Figure 14. Ansys Mechanical Performance: Intel Xeon 5th Gen Processor and 6th Gen AP Processor

Results Summary and Key Takeaways

The performance results presented in this paper demonstrate that achieving optimal CAE performance requires aligning processor architecture, memory technology, and platform design with the computational characteristics of individual solvers. While Intel® Xeon®6 processors deliver clear generation-over-generation improvements compared to Intel Xeon5, the magnitude and source of performance gains vary significantly by solver and workload scale.

Several consistent themes emerge across Fluent, Mechanical, and LS-DYNA workloads:

Solver behavior strongly influences performance outcomes

Fluent software exhibits a balanced dependence on compute capability and memory bandwidth, with larger CFD models increasingly constrained by memory throughput as working sets exceed cache capacity. LS-DYNA software benefits most directly from increased core density and sustained CPU throughput, reflecting the repetitive, explicit time-integration nature of the solver. In contrast, Mechanical software performance depends heavily on solver mode: direct solvers are dominated by memory behavior, while iterative solvers are more limited by algorithmic convergence and synchronization.

Memory bandwidth has become a first-order performance factor

Across the evaluated workloads, memory subsystem performance plays a critical role in determining achievable solver performance. Higher-bandwidth memory configurations consistently improve performance for memory-intensive CFD workloads and selected explicit dynamics cases, while offering more limited benefit for iterative implicit solvers. These results reinforce that increasing core count alone is insufficient unless accompanied by adequate memory throughput.

Single-node performance remains critical, even in scaled environments

Strong single-node performance improves overall scalability by reducing the number of nodes required to achieve a given time-to-solution. This effect is particularly important for implicit FEA workloads, where multi-node scaling efficiency is more limited, and per-node performance directly influences solver efficiency. Balanced node design therefore remains essential, even for large, clustered deployments.

Platform design enables sustained performance

Modern CAE workloads often operate at high utilization for extended durations, making sustained performance under continuous load a key differentiator. Platforms that can maintain stable CPU frequency and memory performance without thermal or power-related constraints deliver more consistent and predictable results, particularly for high-core-count and memory-intensive configurations.

There is no single “best” configuration for all CAE workloads

Optimal CAE infrastructure depends on solver mix, model size, and scaling requirements. CFD, explicit dynamics, and implicit FEA place different demands on compute, memory, and system design, requiring solver-aware platform selection rather than a one-size-fits-all approach. Balanced configurations that align hardware capabilities with solver behavior are essential to avoid bottlenecks and maximize return on investment.

Conclusion

As CAE workloads continue to increase in scale and complexity, achieving optimal simulation performance requires more than incremental improvements in processor capability. The results presented in this paper demonstrate that modern CAE performance is shaped by the interaction of processor architecture, memory subsystem design, and platform-level capabilities, all of which must be aligned with solver-specific behavior to deliver consistent and scalable results.

Across Fluent, Mechanical, and LS-DYNA workloads, Intel® Xeon® 6 processors provide clear generation-over-generation performance improvements compared to prior platforms. However, the source of these gains varies by solver. CFD workloads increasingly depend on memory bandwidth as model sizes grow, explicit dynamics benefit primarily from increased core density and sustained CPU throughput, and implicit FEA performance is governed by memory behavior and solver strategy rather than core count alone. These distinctions reinforce that CAE infrastructure decisions must be solver-aware, rather than driven by generic performance metrics.

Memory subsystem performance has emerged as a critical factor for many CAE workloads. As core counts rise and working sets exceed cache capacity, insufficient memory bandwidth can constrain solver efficiency even on advanced processor architectures. High-bandwidth memory technologies therefore play an increasingly important role in enabling scalable performance for large CFD and selected explicit dynamics simulations, while offering more limited benefits for algorithm-constrained iterative solvers.

Equally important, platform design determines whether processor and memory capabilities can be fully realized in practice. CAE workloads often operate at high utilization for extended durations, making sustained performance under continuous load a key requirement. Platforms that can maintain stable CPU frequency and memory performance without thermal or power-related constraints deliver more predictable time-to-solution, particularly for high-core-count and memory-intensive configurations.

By offering both air-cooled and direct-water-cooled ThinkSystem platforms, Lenovo enables CAE customers to deploy infrastructure that scales with their workloads—from traditional departmental environments to high-density HPC systems—while maintaining a consistent software and operational model. When combined with Intel Xeon 6 processors, high-bandwidth memory technologies, and validated Ansys solver stacks, Lenovo ThinkSystem servers provide a flexible and future-ready foundation for modern CAE simulation.

In summary, this study underscores the importance of a balanced, system-level approach to CAE infrastructure design. Matching processor class, memory technology, and platform capabilities to solver characteristics allows organizations to maximize performance, improve scalability, and reduce time-to-solution as engineering simulations continue to evolve.

Authors

Kevin Dean is the Senior Manager of the HPC and AI Performance and Operations Team within the Infrastructure Solutions Group at Lenovo. The role consists of leading the HPC performance engineering process, operations, and strategy as well as providing CAE application performance support and leading the customer support for the manufacturing vertical. Kevin has 8 years of experience in HPC and AI application performance support at Lenovo plus 12 years of aerodynamic design and computational fluid dynamics experience in the US defense and automotive racing industries. Kevin holds an MS degree in Aerospace Engineering from the University of Florida and a BS degree in Aerospace Engineering from Virginia Polytechnic Institute and State University.

Special thanks to Special thanks to Wim Slagter at Synopsys and the Intel team for providing input and feedback to this paper..

Related product families

Product families related to this document are the following:

- [High Performance Computing](#)
- [ThinkSystem SC750 V4 Server](#)
- [ThinkSystem SD650 V3 server](#)
- [ThinkSystem SR630 V4 Server](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2026. All rights reserved.

This document, LP2442, was created or updated on June 2, 2026.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2442>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2442>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

Neptune®

ThinkSystem®

The following terms are trademarks of other companies:

Intel®, the Intel logo and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Dynamics and Fluid are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.