

Dell Validated Design for HPC Digital Manufacturing with Ansys and 4th Generation Intel Xeon Scalable Processors

May 2023

H19593

White Paper

Abstract

This white paper presents the Dell Validated Design for HPC Digital Manufacturing with performance benchmarking results for Ansys Fluent, LS-DYNA and Ansys Mechanical using 4th Generation Intel Xeon Scalable processors.

Dell Technologies Solutions



Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2023 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA 05/23 White Paper H19593.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

- Introduction4**
- Business challenges.....4**
- Solution overview.....5**
- Solution architecture5**
- Reference system.....7**
- Ansys Fluent performance9**
- LS-DYNA performance.....11**
- Ansys Mechanical performance.....14**
- Conclusion.....16**

Introduction

Executive summary

The Dell Validated Design for HPC Digital Manufacturing is designed specifically for Computer Aided Engineering (CAE) applications that are commonly used for virtual product development and engineering. This design uses a flexible building block approach, where individual building blocks can be combined to build High-Performance Computing (HPC) systems, which are optimized for specific workloads and use cases.

Document purpose

The purpose of this document is to provide guidance for designing HPC systems for use with various CAE software, including Ansys Fluent, LS-DYNA, and Ansys Mechanical.

Audience

This document is intended for decision makers influencing, purchasing, or managing CAE resources for product design and engineering organizations.

Business challenges

Market environment

Manufacturers are facing increased pressure to accelerate product development to serve a broad, rapidly changing market. Coupled with this demand, complying with the wide range of regional regulations with a global customer base has added to the complexity of the product development process. Manufacturers are increasingly relying on CAE to improve the speed and quality of their product development. As HPC integrates new technologies to meet this demand, Dell Technologies has focused on creating holistic, integrated HPC CAE solutions to assist these customers. Dell Validated Designs provide assurance that customers can keep pace with their virtual product development needs, focusing less on IT, and more on the business of bringing products to market.

Solution overview

Introduction

The Dell Validated Design for HPC Digital Manufacturing is designed specifically for CAE applications, which are commonly used for virtual product development and engineering. The solution is designed using a flexible building block architecture. This architecture allows an HPC system to be optimally configured for specific end-user requirements, while still using domain-specific system recommendations.

The building blocks consist of infrastructure servers, storage, networking, and compute servers. Configuration recommendations are provided for each of the building blocks, which are appropriate for Ansys Fluent, LS-DYNA and Ansys Mechanical software applications.

With this flexible building block approach, appropriately sized HPC clusters can be designed based on specific workloads and use-case requirements.

Solution architecture

Infrastructure servers

Infrastructure servers are used to administer the system and provide user access. They are not typically involved in computation, but they provide services that are critical to the overall HPC system. These servers are used as the management nodes and the login nodes. For small-sized clusters, a single physical server can provide the necessary system management functions. Infrastructure servers can also be used to provide storage services, in which case they must be configured with additional disk drives or an external storage array. One management server is necessary for an HPC system to deploy and manage the system. If high-availability (HA) management functionality is required, two management servers are necessary. Login nodes are optional, and one login server per 30 to 100 users is recommended.

A recommended base configuration for infrastructure servers includes:

- Dell PowerEdge R660 server
- Intel Xeon Gold 6430 processor
- 256 GB of RAM (16 x 16 GB 4800 MTps DIMMs)
- PERC H355 RAID controller
- 2 x 480 GB Mixed-Use SATA SSD RAID 1
- Dell iDRAC Enterprise
- 2 x 1100 W power supply units (PSUs)
- NVIDIA ConnectX-7 InfiniBand HCA (optional)

The PowerEdge R660 server is suited for this role. Typical HPC clusters only use a few infrastructure servers; therefore, density is not a priority, but manageability is important. The Intel Xeon Gold 6430 processor, with 32 cores per socket, is recommended for this role. If the infrastructure server is used for CPU-intensive tasks, such as compiling software or processing data, a more capable processor might be appropriate. Configuring the server with 256 GB of RAM provided by 16 16 GB DIMMs provides sufficient memory capacity, with minimal cost per GB, while also providing good memory bandwidth. These servers are not expected to perform considerable I/O, so mixed-use SATA SSDs

configured with RAID 1 are sufficient for the operating system. For most systems, NDR InfiniBand is the data interconnect of choice. It provides a high-throughput, low-latency fabric for node-to-node communications or to access file storage. Including an InfiniBand HCA in the infrastructure server is optional depending on the system design.

Compute servers Compute servers provide the computational resources for the HPC system. These servers are used to run the engineering workloads such as Ansys software applications. The best configuration for the compute servers depends on the specific applications being used and the simulation requirements. The application performance section of this white paper discusses the relevant criteria to consider when selecting a compute server configuration for specific workloads. The following table provides example configuration options for the compute servers:

Table 1. Compute server configuration options

Component	Details
Platforms	<ul style="list-style-type: none"> Dell PowerEdge R660 Dell PowerEdge C6620
Processors	<ul style="list-style-type: none"> Dual Intel Xeon Gold 6438Y+ (32 cores per socket) Dual Intel Xeon Gold 6545S (32 cores per socket) Dual Intel Xeon Platinum 8462Y+ (32 cores per socket)
Memory options	<ul style="list-style-type: none"> 256 GB (16 x 16 GB 4800 MTps DIMMs) 512 GB (16 x 32 GB 4800 MTps DIMMs) 1024 GB (16 x 64 GB 4800 MTps DIMMs)
Storage options	<ul style="list-style-type: none"> PERC H355 or H755 RAID controller 2 x 480 GB Mixed-Use SATA SSD RAID 0 4 x 480 GB Mixed-Use SATA SSD RAID 0
iDRAC	<ul style="list-style-type: none"> iDRAC Enterprise
Power supplies	<ul style="list-style-type: none"> 2 x 1100 W PSU (R660) 2 x 2800 W PSU (C6600)
Networking	<ul style="list-style-type: none"> NVIDIA ConnectX-7 NDR200 InfiniBand adapter NVIDIA ConnectX-7 NDR InfiniBand adapter

Storage

Dell Technologies offers a wide range of general-purpose and HPC storage solutions. For a general overview of the Dell HPC solution portfolio, go to www.dell.com/hpc.

For a general-purpose NAS storage solution, the Dell PowerScale A300 is recommended. PowerScale A300 provides raw storage capacity of 120 TB to 75 PB and sufficient performance for modest-sized HPC systems. For larger HPC systems, the Dell PowerScale F600 all-flash NAS, with raw storage capacity of 138 TB to 186 PB, provides a significant performance benefit compared to the PowerScale A300. For HPC systems requiring a high-performance parallel file system, the Dell Validated Design for BeeGFS storage in the high-performance configuration is recommended. This solution provides 25.6 TB of raw storage capacity per storage server and can be expanded to any required storage capacity.

System networks Most HPC systems are configured with two networks—an administration network, and a high-speed, low-latency switched fabric. The administration network is typically Gigabit Ethernet that connects to the onboard LOM of every server in the cluster. This network is used for provisioning, management, and administration. On the compute servers, this network is also used for BMC management. For infrastructure and storage servers, the iDRAC Enterprise ports may be connected to this network for OOB server management. The management network typically uses the Dell PowerSwitch N3248TE-ON Ethernet switch. If there is more than one Ethernet switch in the system, multiple switches can be stacked with 10-Gigabit Ethernet cables.

A high-speed, low-latency fabric is recommended for clusters with more than four servers. The current recommendation is an NDR InfiniBand fabric. The fabric is typically assembled using NVIDIA QM9790 64-port NDR InfiniBand switches. The number of switches required depends on the size of the cluster and the blocking ratio of the fabric.

Cluster management software

Bright Cluster Manager (BCM) is the recommended cluster management software for installing and monitoring the HPC system.

Reference system

System configuration

Performance benchmarking was performed in the Dell HPC & AI Innovation Lab using the system configurations listed in the following table:

Table 2. Benchmark system configurations

Building block	Quantity
<ul style="list-style-type: none"> • Computational server • PowerEdge R660 server • Dual Intel Xeon Platinum 8452Y CPUs • Air Cooling • 512 GB RAM 16x32GB 4800 MTps DIMMs • NVIDIA ConnectX-6 HDR adapter 	8
<ul style="list-style-type: none"> • Computational server • PowerEdge C6620 server • Dual Intel Xeon Platinum 8480+ CPUs • Direct Liquid Cooling • 512 GB RAM 16x32GB 4800 MTps DIMMs • NVIDIA ConnectX-7 NDR200 adapter 	8
<ul style="list-style-type: none"> • NVIDIA QM9790 InfiniBand Switch 	1
<ul style="list-style-type: none"> • NVIDIA QM8790 InfiniBand Switch 	1

The following table lists the BIOS configuration options used for the reference system:

Table 3. BIOS configuration

BIOS Option	Setting
Logical processor	Disabled
System profile	Performance Profile
SubNumaCluster	4-way

The following table lists the software versions used for the reference system:

Table 4. Software versions

Component	Version
Operating system	Red Hat Enterprise Linux 8.6
Kernel	4.18.0-372.43.1.el8_6.x86_64
OFED	MLNX_OFED_LINUX-5.9-0.5.6.0
Bright Cluster Manager	9.2
Ansys Fluent	2023 R1
Ansys Mechanical	2023 R1
LS-DYNA	R13.1.0 Single Precision

Benchmark disclaimer

Note the following constraints:

- Benchmark results are highly dependent on workload, specific application requirements, and system design and implementation. Relative system performance varies as a result of these and other factors. Do not use this workload as a substitute for a specific customer application benchmark when critical capacity planning and product evaluation decisions are contemplated.
- All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments might vary significantly. Dell Technologies does not warrant or represent that a user can or will achieve similar performance results.

Ansys Fluent performance

Ansys Fluent is a Computational Fluid Dynamics (CFD) software application commonly used for a wide range of CFD and multiphysics simulations. CFD applications typically scale well across multiple processor cores and servers, have modest memory capacity requirements, and typically perform minimal disk I/O while solving. However, some simulations might have greater I/O demands, such as large transient analysis. We evaluated 15 benchmark problems from the Fluent benchmark suite on the reference system. For a description of the benchmark cases, see the [Ansys Benchmarks Overview](#).

The Solver Rating metric, which is the number of 25 iteration solves that can be completed in a day, is used to report the Fluent benchmark performance. The metric is calculated using (total seconds in a day)/(25 iteration solve time in seconds). A larger value represents better performance.

Single-server performance

The following figure shows measured performance for five of the Ansys Fluent benchmarks using Ansys Fluent 2023 R1 on a single server:



Figure 1. Ansys Fluent—Single-server relative performance

The results are plotted relative to the performance of a single C6520 server configured with 32-core Intel Xeon Platinum 8358 processors. Larger values indicate better performance. These results show the performance advantage available with 4th Generation Intel Xeon Scalable processors.

Multiserver scalability

The following figure shows the parallel scalability of the Fluent benchmark models using up to four R660 servers configured with Intel Xeon Platinum 8452Y processors. The performance is presented relative to the performance of a single server.

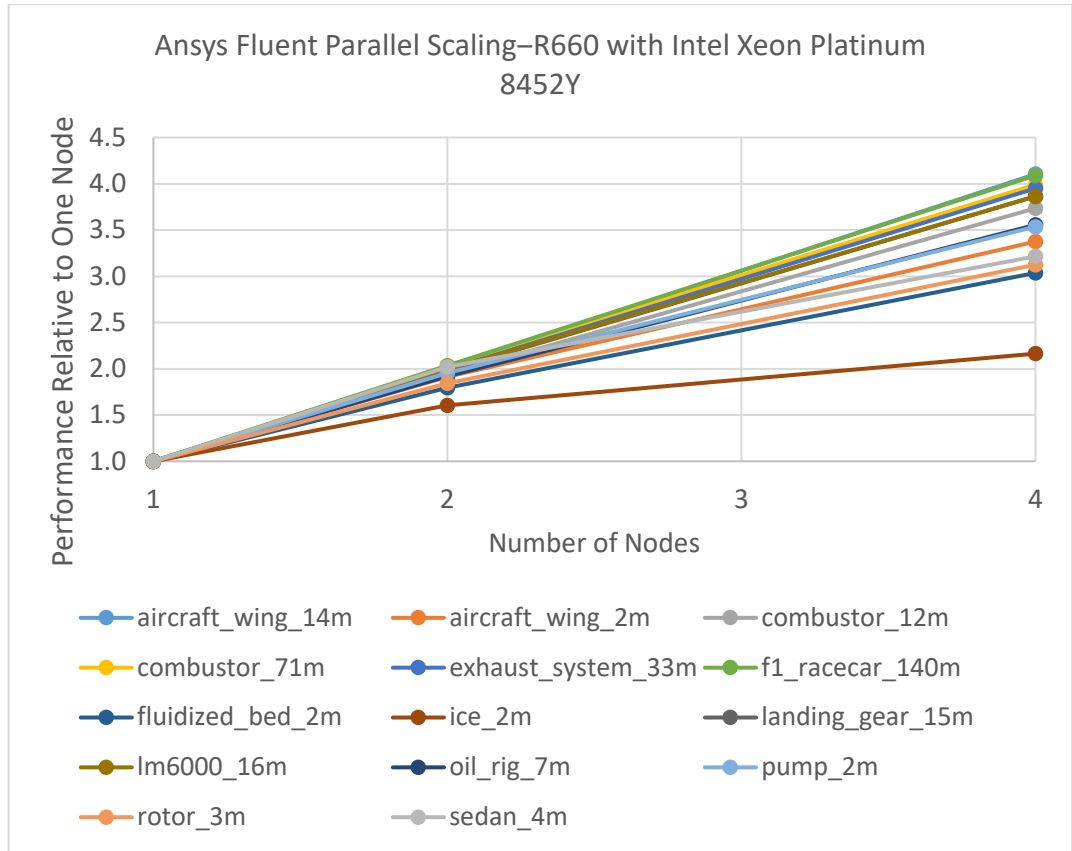


Figure 2. Ansys Fluent parallel scaling—R660 with Intel Xeon Platinum 8452Y

The parallel scalability for most of these benchmark models is good, with the models scaling as expected. The ice_2m benchmark does not scale as well as the other benchmarks. However, this limitation is expected as it is a small model that includes dynamic mesh and combustion simulation.

LS-DYNA performance

LS-DYNA is an industry-leading explicit simulation software used for applications such as drop tests, impact, penetration, crash simulation, occupant safety, and many other disciplines. The two benchmark problems presented here use the LS-DYNA explicit FEA solver, which typically scales much more efficiently than the implicit FEA solver.

The car2car benchmark is a simulation of a two-vehicle collision. This benchmark model contains 2.4 million elements. The ODB-10M benchmark is a simulation of a vehicle colliding into an offset deformable barrier. This benchmark model contains 10.6 million elements. For the ODB-10M benchmarks, the simulation end time was set to 0.02 sec. The performance for LS-DYNA is measured using the Elapsed Time metric. This metric is the total elapsed time in seconds as reported by LS-DYNA, so a smaller elapsed time is better performance.

Single-server performance

The measured performance for the LS-DYNA benchmarks using LS-DYNA MPP single-precision R13.1.0 on a single server are shown in Figure 3 and Figure 4.




Figure 3. LS-DYNA car2car-Single server relative performance




Figure 4. LS-DYNA ODB-10M—Single server relative performance

The results in Figure 3 and Figure 4 are plotted relative to the performance of a single server configured with 32-core Intel Xeon Gold 8358 processors. Larger values indicate better overall performance. These results show the performance advantage available with 4th generation Intel Xeon Scalable processors. For these two benchmark cases, the AVX2 and AVX-512 versions of LS-DYNA perform similarly and provide a performance advantage compared with the SSE2 version.

Multiserver scalability

Figure 5 presents the parallel scalability when running the AVX2 version of LS-DYNA using up to four R660 servers configured with Intel Xeon Platinum 8452Y processors. The performance is presented relative to the performance of a single server.

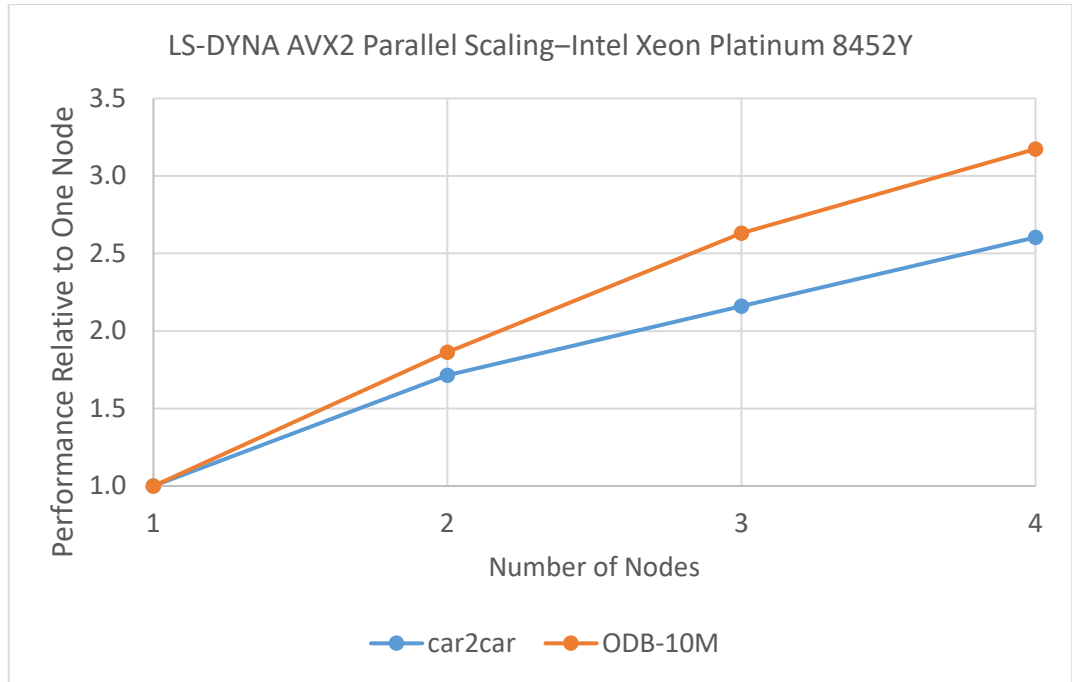


Figure 5. LS-DYNA AVX2 Parallel Scaling–Intel Xeon Platinum 8452Y

The parallel scalability for these models is as expected, with the larger ODB-10M model demonstrating better scalability than the smaller Car2Car benchmark model.

Ansys Mechanical performance

Ansys Mechanical is a finite element analysis (FEA) software with structural, thermal, acoustics, transient and nonlinear capabilities. Implicit FEA solvers often place large demands on the memory and disk I/O subsystems, particularly for out-of-core solutions, where the problem is too large to fit into the available system RAM. Because of these characteristics, the performance of any specific system configuration is highly dependent on the workload.

The specific Ansys Mechanical workload determines the appropriate processor, memory, and disk I/O configuration. Therefore, it is difficult to provide a general system configuration which is optimized for all Ansys Mechanical use cases. Critical considerations for system configuration for Ansys Mechanical include memory capacity and local disks for scratch I/O. Due to these dependencies, you must size systems for Ansys Mechanical based on the precise workloads.

Single-server performance

The performance of six standard Ansys Mechanical benchmark cases were evaluated on the reference system. The benchmark cases all run in-core with the 512 GB of RAM that is available per compute server on the reference system. So, the local disk configuration has minimal performance impact on the standard benchmarks. Two types of solvers are available with Ansys Mechanical: Distributed Memory Parallel (DMP) and Shared Memory Parallel (SMP). In general, the default DMP solver offers equivalent or better performance than the SMP solver, particularly when all cores on a processor are used. As such, only results from the DMP solver are presented. Performance is measured using the Core Solver Rating metric. This metric represents the performance of the solver core which excludes any pre- and postprocessing.

Figure 6 shows the measured performance of the Ansys Mechanical benchmarks on a single server. The results are plotted relative to the performance of a single server configured with Intel Xeon Platinum 8358 processors. Larger values indicate better overall performance. These results show the performance advantage available with 4th generation Intel Xeon Scalable processors.

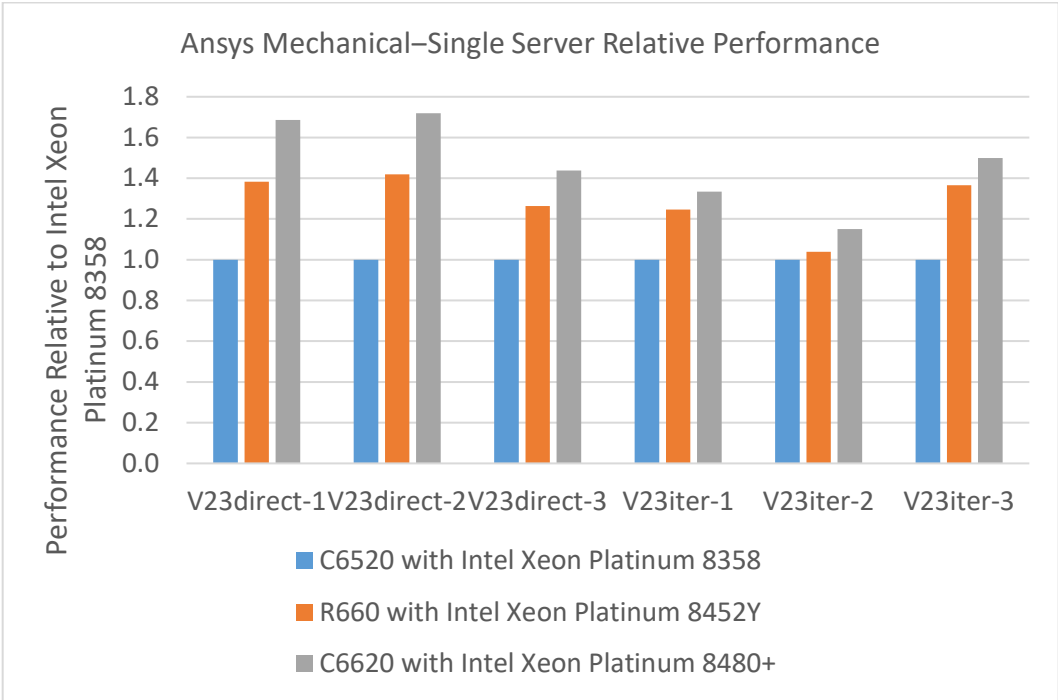


Figure 6. Ansys Mechanical-Single server relative performance

Single-server scalability

Figure 7 shows the scaling behavior of the Ansys Mechanical benchmarks on a single server. The benchmark results are plotted relative to the benchmark performance using one processor core. Each data point on the graph records the performance of the specific benchmark dataset using the number of processor cores marked on the x-axis relative to the one-core result.

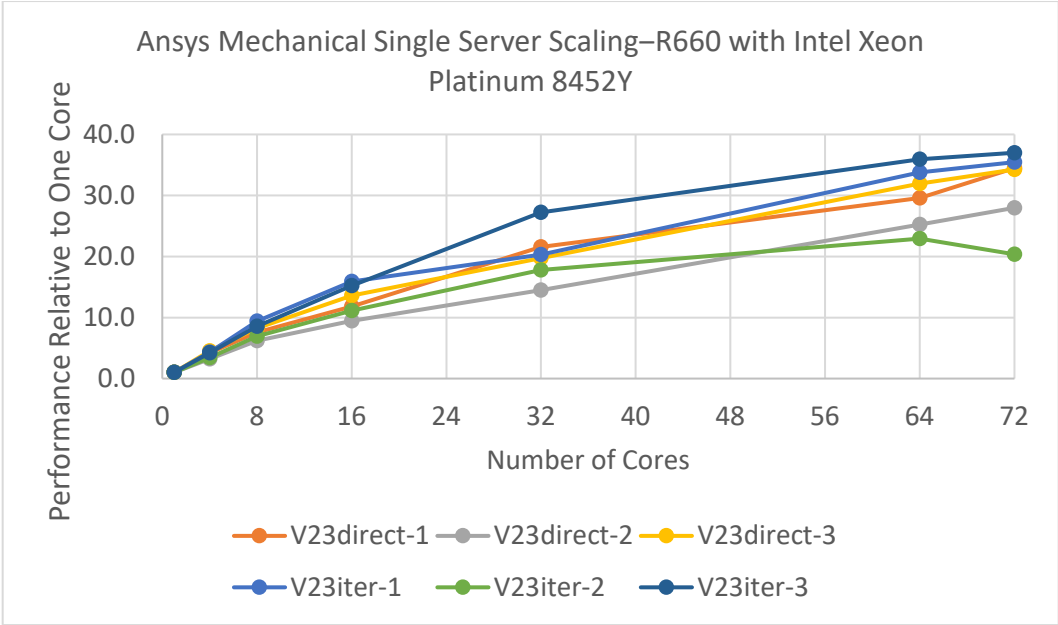


Figure 7. Ansys Mechanical Server scaling-R660 with Intel Xeon Platinum 8452Y

The results in Figure 7 demonstrate that the Ansys Mechanical DMP solver scales reasonably well from one to 72 processor cores.

Multiserver scalability

Performance results for the Ansys Mechanical solver using multiple servers are shown in Figure 8. The results are plotted relative to the benchmark performance using a single server. Each data point on the graph records the performance of the specific benchmark dataset using the number of servers marked on the x-axis relative to the single-server result.

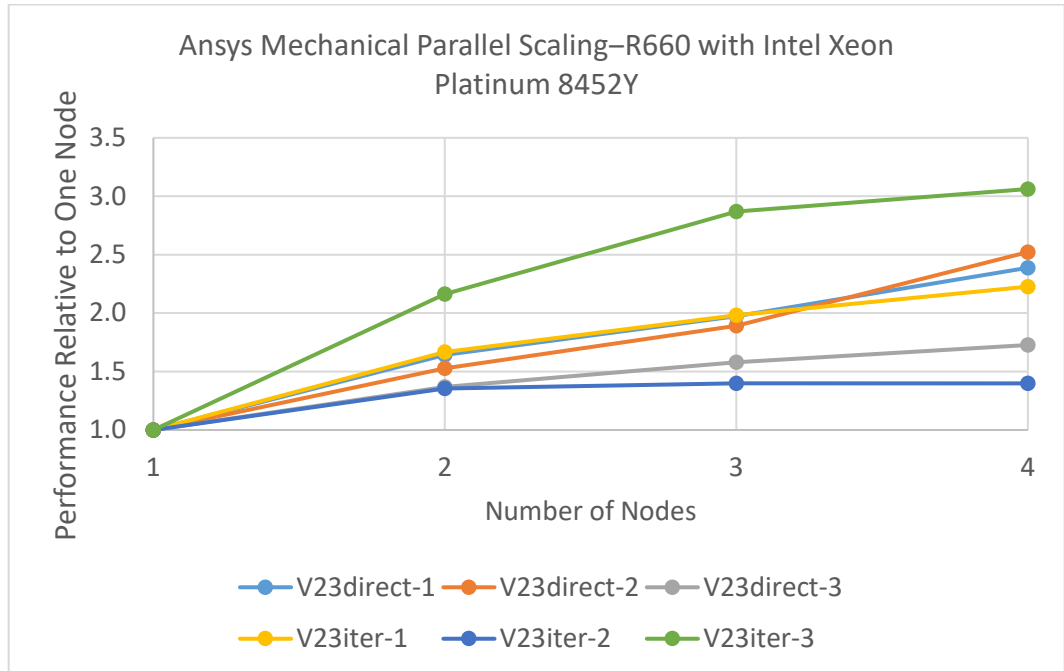


Figure 8. Ansys Mechanical Parallel scaling—R660 with Intel Xeon Platinum 8452Y

The standard benchmark set being used for this testing is primarily intended for single-server benchmarks. As such, some of the benchmark cases do not scale well beyond two nodes. However, the results also show that some of the cases do scale to four nodes, which demonstrates that the scalability performance is highly dependent on the specific use case.

Conclusion

This technical white paper presents the Dell Validated Design for HPC Digital Manufacturing with Ansys Fluent, LS-DYNA, and Ansys Mechanical performance benchmarks. This Dell Validated Design is architected for a specific purpose—to provide a comprehensive HPC solution for CAE software. The building block approach allows customers to easily deploy an HPC system customized for specific workload requirements. The design addresses computation, storage, networking, and software requirements, and provides a solution that is easy to install, configure and manage, with installation services and support readily available. The performance benchmarking substantiates the solution design, demonstrating the performance of the solution with Ansys Fluent, LS-DYNA, and Ansys Mechanical software.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#).