

Dell Technologies Validated Design for Ansys with 3rd Generation Intel Xeon Scalable Processors

Ansys® CFX®, Fluent®, LS-DYNA® and Mechanical™, Ansys, Inc.

December 2021

White Paper

Abstract

This white paper describes the Dell Technologies Validated Design for Ansys with 3rd Generation Intel® Xeon® Processors. The design is presented and validated with Ansys CFX, Fluent, LS-DYNA and Mechanical software applications.

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 © 2021 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 12/21 White Paper .

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 architecture6**
- Reference System10**
- Ansys CFX Performance.....12**
- Ansys Fluent Performance14**
- LS-DYNA Performance.....16**
- Ansys Mechanical Performance.....18**
- Conclusion.....22**

Introduction

Executive summary

The Dell Technologies Validated Design for Ansys with 3rd Generation Intel Xeon Processors is designed specifically for Ansys Computer Aided Engineering (CAE) applications which 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 HPC systems which are optimized for specific workloads and use cases.

Document purpose

The purpose of this document is to provide guidance on designing HPC systems for use with Ansys CAE software.

Audience

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

Business challenges

Market environment

The recent developments in the world-wide economy have placed increased pressure on manufacturers to accelerate product development to serve a broad, rapidly changing market. Coupled with this demand, complying with the wide range of regulations required with a global customer base has added to the complexity of the product development process. Manufacturers are increasingly relying on computer aided engineering (CAE) to improve the speed and quality of their product development. As High-Performance Computing (HPC) integrates new technologies to meet this demand for I/T resources, Dell Technologies has focused on creating holistic, integrated HPC CAE solutions to assist these customers. Dell Technologies Validated Designs provide assurance that customers can keep pace with their virtual product development needs, focusing less on I/T and more on the business of bringing products to market.

Solution overview

Introduction

The Dell Technologies Validated Design for Ansys with 3rd Generation Intel Xeon Processors is designed specifically for Ansys Computer Aided Engineering (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 making use of standardized, 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 software applications.

With this flexible building block approach, appropriately sized HPC clusters can be designed based on specific workloads and use-case requirements. Figure 1 shows three example HPC clusters designed using the Dell Technologies Validated Design for Ansys architecture.

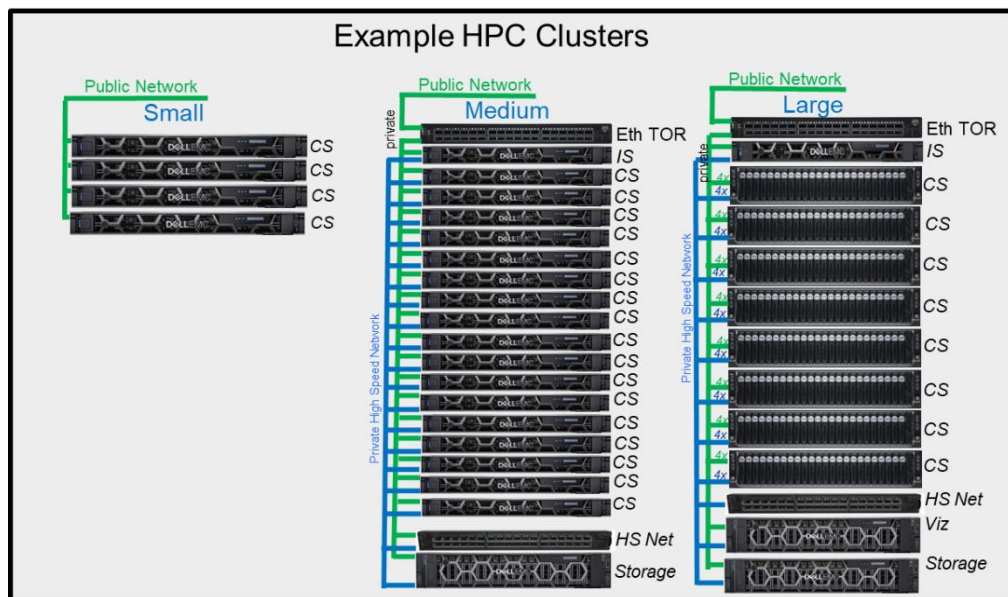


Figure 1. Example Dell Technologies Validated Design for Ansys Solutions

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, by using NFS, in which case they must be configured with additional disk drives or an external storage array. One management server is mandatory 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 30-100 users is recommended.

A recommended base configuration for infrastructure servers is:

- Dell EMC PowerEdge R650 server
- Dual Intel Xeon Silver 4314 processors
- 256 GB of RAM (16 x 16GB 3200 MTps DIMMs)
- PERC H345 RAID controller
- 2 x 480GB Mixed-Use SATA SSD RAID 1
- Dell EMC iDRAC Enterprise
- 2 x 750 W power supply units (PSUs)
- NVIDIA® ConnectX-6 InfiniBand™ HCA (optional)

The recommended base configuration for the infrastructure server is described as follows. The PowerEdge R650 server is suited for this role. Typical HPC clusters will only use a few infrastructure servers; therefore, density is not a priority, but manageability is important. The Intel Xeon Silver 4314 processor, with 16 cores per socket, is a basic recommendation for this role. If the infrastructure server will be used for CPU intensive tasks, such as compiling software or processing data, then a more capable processor may be appropriate. 256 GB of RAM provided by sixteen 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 much I/O, so mixed-use SATA SSDs configured with RAID 1 should be sufficient for the operating system. For small systems (four nodes or less), an Ethernet network may provide sufficient application performance. For most other systems, HDR InfiniBand is likely to be the data interconnect of choice, which provides a high-throughput, low-latency fabric for node-to-node communications or to access Dell Technologies Validated Design Storage solutions.

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 is dependent on the specific applications in use and the simulation requirements. Since the best configuration may be different for each use case, a table of recommended options are provided that are appropriate for these servers. A specific configuration can be selected to match the requirements of the workloads and use cases. Relevant criteria to consider prior to selecting a compute server configuration are discussed in the application performance section of this white paper. The recommended configuration options for the compute servers are provided in Table 1.

Table 1. Compute Server Configuration Options

Platforms	Dell EMC PowerEdge C6520 Dell EMC PowerEdge R650 Dell EMC PowerEdge R750
Processors	Dual Intel Xeon Gold 6346 (16 cores per socket) Dual Intel Xeon Gold 6342 (24 cores per socket) Dual Intel Xeon Gold 6338 (32 cores per socket) Dual Intel Xeon Gold 8358 (32 cores per socket)
Memory Options	256 GB (16 x 16GB 3200 MTps DIMMs) 512 GB (16 x 32GB 3200 MTps DIMMs) 1024 GB (16 x 64GB 3200 MTps DIMMs)
Storage Options	PERC H345, H745 or H755 RAID controller 2 x 480GB Mixed-Use SATA SSD RAID 0 4 x 480GB Mixed-Use SATA SSD RAID 0
iDRAC	iDRAC Enterprise (R650 and R750) iDRAC Express (C6520)
Power Supplies	2 x 750W PSU (R640 and R750) 2 x 2400W PSU (C6400)
Networking	NVIDIA® ConnectX®-6 HDR100 InfiniBand™ adapter NVIDIA® ConnectX®-6 HDR InfiniBand adapter

Storage

Dell EMC offers a wide range of general purpose and HPC storage solutions. For a general overview of the Dell EMC HPC solution portfolio please visit www.dell.com/hpc. There are typically three tiers of storage for HPC: scratch storage, operational storage, and archival storage, which differ in terms of size, performance, and persistence.

Scratch storage tends to persist for the duration of a single simulation. It may be used to hold temporary data which is unable to reside in the compute system's main memory due to insufficient physical memory capacity. HPC applications may be considered "I/O bound" if access to storage impedes the progress of the simulation. For these HPC workloads, typically the most cost-effective solution is to provide sufficient direct-attached local storage on the compute nodes. For situations where the application may require a shared file system across the compute cluster, a high-performance shared file system may be better suited than relying on local direct-attached storage. Typically, using direct-attached local storage offers the best overall price/performance and is considered best practice for most CAE simulations. For this reason, local storage is included in the recommended configurations with appropriate performance and capacity for a wide range of production workloads. If anticipated workload requirements exceed the performance and capacity provided by the recommended local storage configurations, care should be taken to size scratch storage appropriately based on the workload.

Operational storage is typically defined as storage used to maintain results over the duration of a project and other data, such as home directories, such that the data may be accessed daily for an extended period. Typically, this data consists of simulation input and results files, which may be transferred from the scratch storage, typically in a sequential manner, or from users analyzing the data, often remotely. Since this data may persist for an extended period, some or all of it may be backed up at a regular interval, where the interval chosen is based on the balance of the cost to either archive the data or regenerate it if need be. Archival data is assumed to be persistent for a very long term,

and data integrity is considered critical. For many modest HPC systems, use of the existing enterprise archival data storage may make the most sense, as the performance aspect of archival data tends to not impede HPC activities. Our experience in working with customers indicates that there is no 'one size fits all' operational and archival storage solution. Many customers rely on their corporate enterprise storage for archival purposes and instantiate a high-performance operational storage system dedicated for their HPC environment.

Operational storage is typically sized based on the number of expected users. For fewer than 30 users, a single NFS storage server, such as the Dell PowerEdge R740xd is often an appropriate choice. A suitably equipped storage server may be:

- Dell EMC PowerEdge R740xd server
- Dual Intel® Xeon® Silver 4210 processors
- 96 GB of memory, 12 x 8GB 2666 MTps DIMMS
- PERC H740P RAID controller
- 2 x 480GB Mixed-use SATA SSD in RAID-1 (For OS)
- 12 x 12TB 3.5: NLSAS HDDs in RAID-6 (for data)
- Dell EMC iDRAC9 Express
- 2 x 750 W power supply units (PSUs)
- ConnectX-6 HDR100 InfiniBand Adapter
- Site specific high-speed Ethernet Adapter(optional)

This server configuration would provide 144TB of raw storage. For customers expecting between 25-100 users, an operational storage solution, such as the Dell EMC Isilon A200 scale-out NAS may be appropriate.

For customers desiring a shared high-performance parallel filesystem, the Dell Technologies Validated Design for PixStor Storage solution shown in Figure 2 is appropriate. This solution can scale up to multiple petabytes of storage.




Figure 2. Dell Technologies Validated Design for PixStor Storage

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/NDC of every server in the cluster. This network is used for provisioning, management, and administration. On the compute servers, this network will also be 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 Networking S3048-ON Ethernet switch. If there is more than one 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 HDR InfiniBand fabric. The fabric will typically be assembled using NVIDIA QM8790 40-port HDR InfiniBand switches. The number of switches required depends on the size of the cluster and the blocking ratio of the fabric.

Cluster Management Software

The cluster management software is used to install and monitor the HPC system. Bright Cluster Manager (BCM) is the recommended cluster management software.

Reference System

System Configuration

Performance benchmarking was performed in the Dell EMC HPC and AI Innovation Lab using the system configurations listed in Table 2.

Table 2. Benchmark System Configurations

Building Block	Quantity
Computational Server PowerEdge C6520 Dual Intel Xeon Gold 6346 512GB RAM 16x32GB 3200 MTps DIMMs NVIDIA ConnectX-6 HDR100 adapter	1
Computational Server PowerEdge C6520 Dual Intel Xeon Gold 6336Y 512GB RAM 16x32GB 3200 MTps DIMMs NVIDIA ConnectX-6 HDR100 adapter	1
Computational Server PowerEdge C6520 Dual Intel Xeon Platinum 8352Y 512GB RAM 16x32GB 3200 MTps DIMMs NVIDIA ConnectX-6 HDR adapter	1
Computational Server PowerEdge C6520 Dual Intel Xeon Platinum 8358 512GB RAM 16x32GB 3200 MTps DIMMs NVIDIA ConnectX-6 HDR adapter	8
NVIDIA QM8790 InfiniBand Switch	1

The BIOS configuration options used for the reference system are listed in Table 3.

Table 3. BIOS Configuration

BIOS Option	Setting
Logical Processor	Disabled
Virtualization Technology	Disabled
Snoop Holdoff Timer	Roll2kCycles
System Profile	Performance Profile
Sub NUMA Cluster	2-Way

The software versions used for the reference system are listed in Table 4.

Table 4. Software Versions

Component	Version
Operating System	RHEL 8.3
Kernel	4.18.0-240.22.1.el8_3.x86_64
OFED	NVIDIA Mellanox 5.2-2.2.0.0
Bright Cluster Manager	9.0
Ansys CFX	2021 R2
Ansys Fluent	2021 R2
LS-DYNA	R11.2.2 S AVX2
Ansys Mechanical	2021R2

Benchmark Disclaimer

- Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, this workload should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or 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 may vary significantly. Dell Technologies does not warrant or represent that a user can or will achieve similar performance results.

Ansys CFX Performance

Ansys CFX software is a Computational Fluid Dynamics (CFD) application recognized for its accuracy, robustness, and speed with rotating machinery applications. 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 may have greater I/O demands, such as large transient analysis.

Single Server Performance

Figure 3 shows the measured performance of the standard CFX v16 benchmarks using CFX 2021 R2 on a single server. The performance for each benchmark is measured using the solver elapsed time.




Figure 3. Ansys CFX 2021 R2—Single Server Relative Performance

The results in Figure 3 are plotted relative to the performance of a single server configured with Intel Xeon Gold 6248 processors. Larger values indicate better overall performance. These results show the performance advantage available with 3rd generation Intel Xeon Scalable processors (code name Ice Lake). These results also show that the CFX benchmarks can make use of the higher core count processors such as the 32-core Intel Xeon Platinum 8358.

Multi-Server Scalability

Figure 4 presents the parallel scalability when running CFX using up to eight nodes configured with Intel Xeon Platinum 8358 processors. The performance is presented relative to the performance of a single node.

The overall parallel scalability for these models is good, with the parallel scalability of the largest model demonstrating super-linear scaling. This is related to the “cache effect”, where more of the active dataset fits in cache as the number of nodes used increases.




Figure 4. Ansys CFX 2021 R2 Parallel Scaling-Intel Xeon Platinum 8358

Ansys Fluent Performance

Ansys Fluent is a Computational Fluid Dynamics (CFD) software application commonly used for a wide range of CFD and multi-physics 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 may have greater I/O demands, such as large transient analysis. Fifteen benchmark problems from the Fluent benchmark suite were evaluated on the reference system.

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

Single Server Performance

The measured performance for six of the Ansys Fluent benchmarks using Ansys Fluent 2021 R2 on a single server is shown in Figure 5.

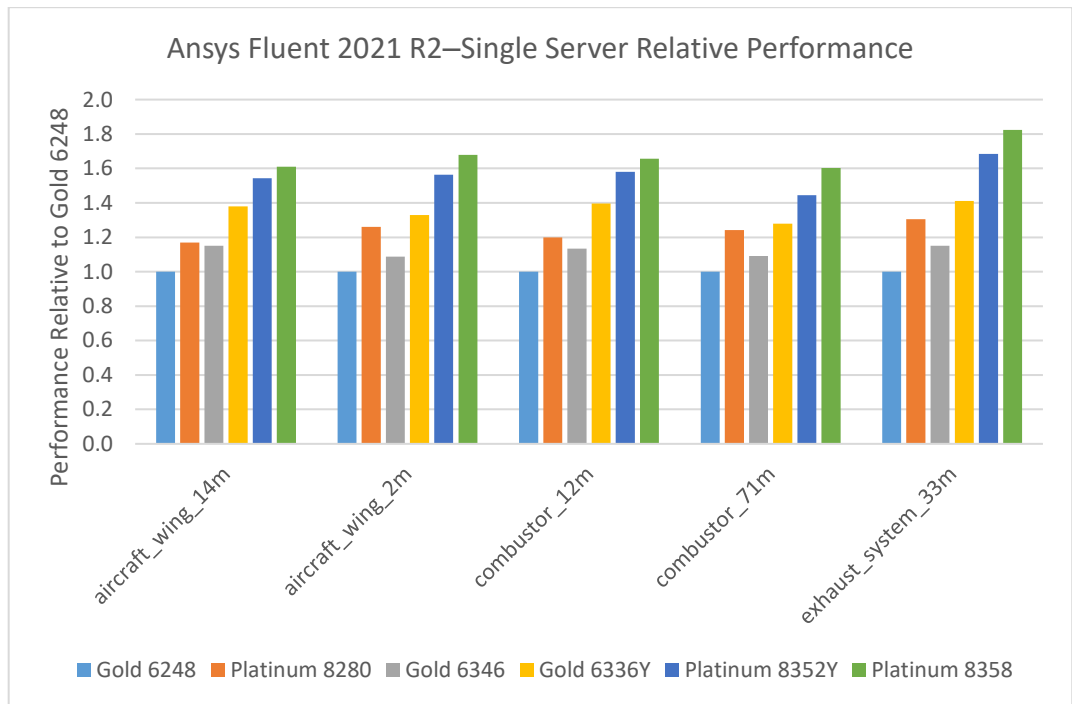


Figure 5. Ansys Fluent 2021 R2—Single Server Relative Performance

The results in Figure 5 are plotted relative to the performance of a single server configured with Intel Xeon Gold 6248 processors. Larger results indicate better performance. These results show the performance advantage available with 3rd generation Intel Xeon Scalable processors (code name Ice Lake). These results also show that most of the Fluent benchmarks can make use of the higher core count processors such as the 32-core Intel Xeon Platinum 8358.

Multi-Server Scalability

Figure 6 presents the parallel scalability of the Fluent benchmark models using up to eight servers configured with Intel Xeon Platinum 8358 processors. The performance is presented relative to the performance of a single server.

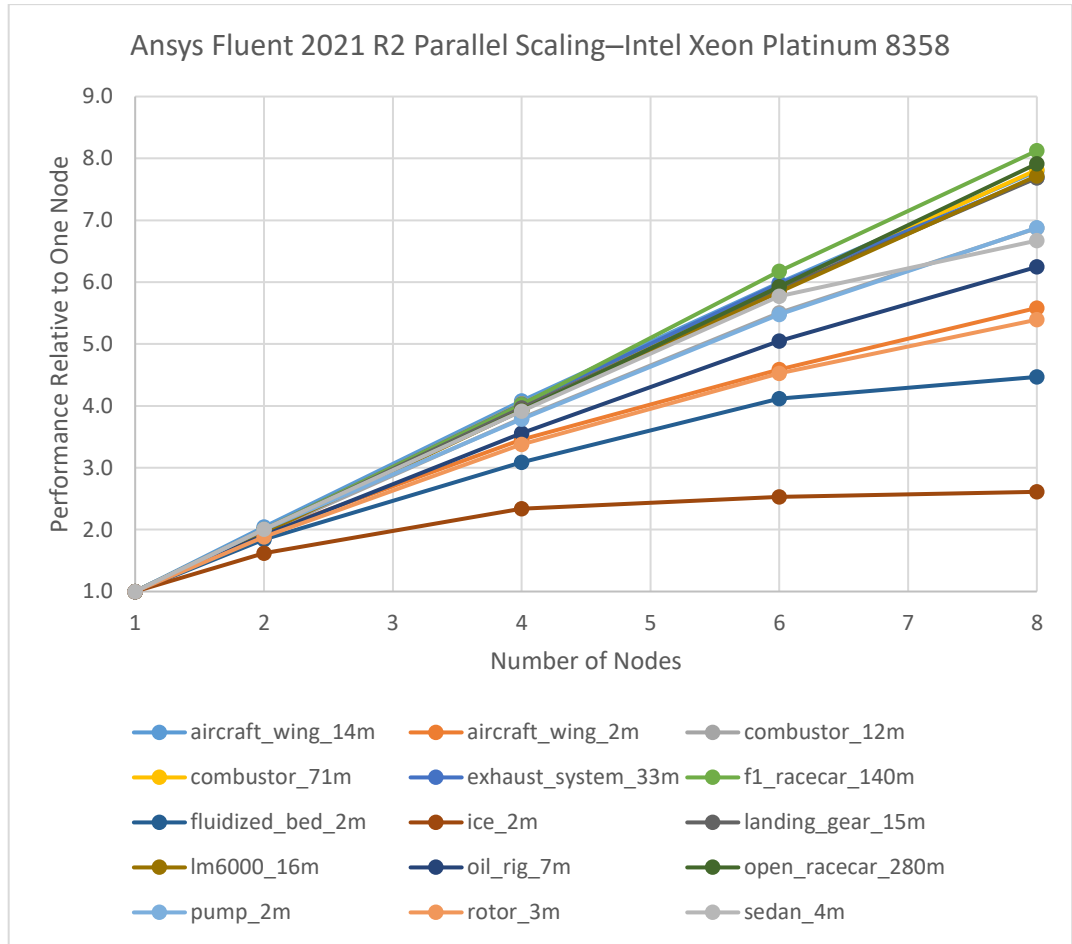


Figure 6. Ansys Fluent 2021 R2 Parallel Scaling—Intel Xeon Platinum 8358

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, but this is expected as it is a small model which includes dynamic mesh and combustion simulation.

LS-DYNA Performance

LS-DYNA is the 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. The benchmark model used here contains 2.4 million elements, which is relatively small compared to current automotive industry usage. The ODB-10M benchmark is a simulation of a vehicle colliding into an offset deformable barrier. This benchmark model contains 10.6 million elements, which is similar to current automotive industry usage. For these 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 R11.2.2 AVX2 on a single server is shown in Figure 7.




Figure 7. LS-DYNA R11.2.2 AVX2—Single Server Relative Performance

The results in Figure 7 are plotted relative to the performance of a single server configured with Intel Xeon Gold 6248 processors. Larger values indicate better overall performance. These results show the performance advantage available with 3rd generation Intel Xeon Scalable processors (code name Ice Lake). These results also show that these benchmarks can make use of the higher core count processors such as the 32-core Intel Xeon Platinum 8358.

Multi-Server Scalability

Figure 8 presents the parallel scalability when running LS-DYNA using up to eight servers configured with Intel Xeon Platinum 8358 processors. The performance is presented relative to the performance of a single server.

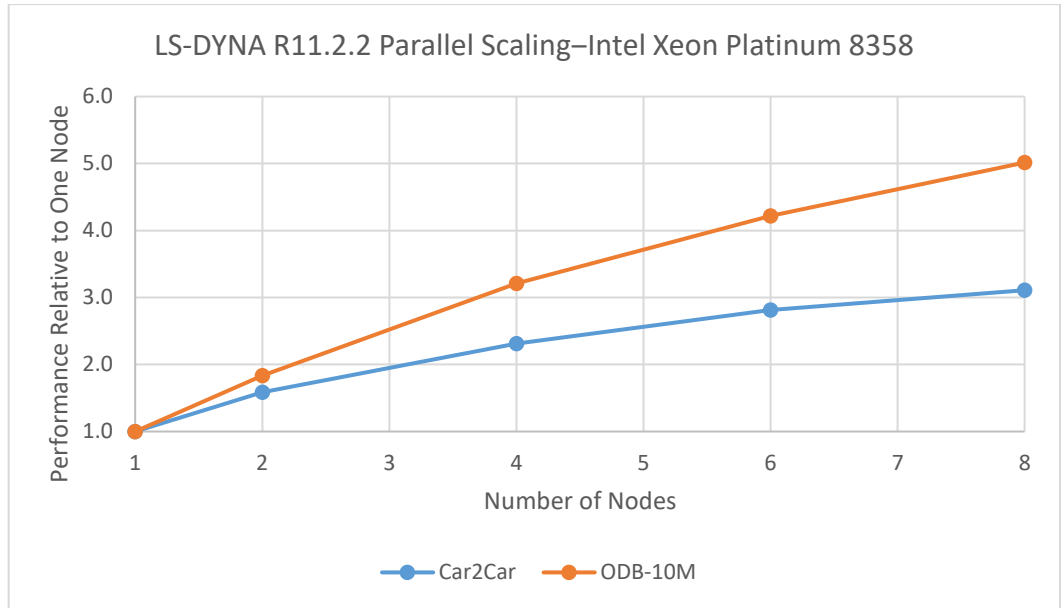


Figure 8. LS-DYNA R11.2.2 AVX2 Parallel Scaling–Intel Xeon Platinum 8358

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) solver with structural, thermal, acoustics, transient and nonlinear capabilities. Implicit FEA problems often place large demands on the memory and disk I/O sub-systems, 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.

Since the specific Ansys Mechanical workload determines the appropriate processor, memory, and disk I/O configuration, it is difficult to provide a general system configuration which is optimized for all Ansys Mechanical use cases. However, a typical general use compute server could be configured as shown in Table 5. This system configuration balances compute, memory and local disk I/O performance and would work well for many common Ansys Mechanical use cases.

Table 5. Example 512GB RAM Compute Server Configuration for Ansys Mechanical

Platform	Dell EMC PowerEdge C6520
Processor	Dual Intel Xeon Gold 6336Y (24 cores per socket)
Memory	16 x 32GB 3200 MTps DIMMs (512GB)
Storage	PERC H745 RAID controller 4 x 480GB Mixed-Use SATA SSD RAID 0
iDRAC	iDRAC9
Power Supplies	2 x 2400W PSU
Networking	NVIDIA® ConnectX®-6 HDR InfiniBand™ adapter

An example configuration which would provide better performance for relatively large Ansys Mechanical simulations is shown in Table 6. The primary changes for this configuration relative to the previous configuration are increasing the RAM capacity to 1TB, using the higher performance PERC H755 RAID controller, and increasing the local disk capacity to 3.84TB.

Table 6. Example 1TB RAM Compute Server Configuration for Ansys Mechanical

Platform	Dell EMC PowerEdge R650
Processor	Dual Intel Xeon Gold 6336Y (24 cores per socket)
Memory	16 x 64GB 3200 MTps DIMMs (1TB)
Storage	PERC H755 RAID controller 4 x 960GB Mixed-Use SATA SSD RAID 0
iDRAC	iDRAC9 Enterprise
Power Supplies	2 x 1100W PSU
Networking	NVIDIA® ConnectX®-6 HDR InfiniBand™ adapter

Single Server Performance

The performance of the ten standard Ansys Mechanical benchmark cases were evaluated on the reference system. The ten benchmark cases all run in-core with the 512GB 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 post-processing.

Figure 9 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 Gold 6248 processors. Larger values indicate better overall performance. These results show the performance advantage available with 3rd generation Intel Xeon Scalable processors (code name Ice Lake). The results demonstrate that processors with core counts between 24 and 32 cores are good choices for general use with Ansys Mechanical.




Figure 9. Ansys Mechanical 2021 R2-Single Server Relative Performance

Single Server Scalability

Figure 10 shows the scaling behavior of the Ansys Mechanical benchmarks on a single server. The benchmark results are plotted relative to the benchmark performance using two processor cores. Each data point on the graph records the performance of the specific benchmark data set using the number of processor cores marked on the x-axis relative to the two-core result.




Figure 10. Ansys Mechanical 2021 R2 Single Server Scaling–Intel Xeon Platinum 8352Y

The results in Figure 10 demonstrate that the Ansys Mechanical DMP solver scales reasonably well from two to 48 processor cores. Using more than 48 cores on a single server results in modest performance increases for most of the benchmark models. This data shows that using 48 cores per node works well for the standard benchmark models.

Multi-Server Scalability

Performance results for the Ansys Mechanical solver using multiple servers are shown in Figure 11. 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 data set using the number of servers marked on the x-axis relative to the single server result.




Figure 11. Ansys Mechanical 2021 R2 Parallel Scaling–Intel Xeon Platinum 8352Y

The DMP solver scales up to eight servers for some of the benchmarks. Problem scalability depends on many factors including the number of degrees of freedom and the solution type and solver being used. For these benchmark cases, parallel scalability above four nodes is limited; however, it is possible to see good parallel scalability up to eight nodes for larger models.

Conclusion

This technical white paper presents the Dell Technologies Validated Design for Ansys with 3rd generation Intel Xeon Scalable processors. The detailed analysis of the server configurations demonstrate that the system is architected for a specific purpose—to provide a comprehensive HPC solution for Ansys software. Use of the building block approach allows customers to easily deploy an HPC system optimized 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 bears out the solution design, demonstrating the performance of the solution with Ansys 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](#) or provide your comments by completing our [documentation survey](#).