

Dell EMC Ready Solution for HPC Digital Manufacturing—ANSYS® Performance

Abstract

This Dell EMC technical white paper discusses performance benchmarking results and analysis for ANSYS® CFX®, Fluent®, and Mechanical™ on the Dell EMC Ready Solution for HPC Digital Manufacturing.

April 2019

Revisions

Date	Description
January 2018	Initial release with Intel® Xeon® Scalable processors (code name Skylake)
February 2019	Added Basic Building Block
April 2019	Revised with 2 nd generation Intel Xeon Scalable processors (code name Cascade Lake)

Acknowledgements

This paper was produced by the following:

Authors: Joshua Weage
Martin Feyereisen

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 © 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [4/23/2019] [Technical White Paper]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
1 Introduction.....	4
2 System Building Blocks.....	5
2.1 Infrastructure Servers.....	5
2.2 Compute Building Blocks.....	6
2.3 Basic Building Blocks	7
2.4 Storage	8
2.5 System Networks.....	10
2.6 Cluster Management Software.....	10
2.7 Services and Support	11
3 Reference System.....	12
4 ANSYS CFX Performance	14
5 ANSYS Fluent Performance.....	16
6 ANSYS Mechanical Performance	18
7 Basic Building Block Performance	24
8 Conclusion.....	27

1 Introduction

This technical white paper discusses the performance of ANSYS® CFX®, Fluent® and Mechanical™ on the Dell EMC Ready Solution for HPC Digital Manufacturing. This Dell EMC Ready Solution for HPC was designed and configured specifically for Digital Manufacturing workloads, where Computer Aided Engineering (CAE) applications are critical for virtual product development. The Dell EMC Ready Solution for HPC Digital Manufacturing uses a flexible building block approach to HPC system design, where individual building blocks can be combined to build HPC systems which are optimized for customer specific workloads and use cases.

The Dell EMC Ready Solution for HPC Digital Manufacturing is one of many solutions in the Dell EMC HPC solution portfolio. Please visit www.dell EMC.com/hpc for a comprehensive overview of the available HPC solutions offered by Dell EMC.

The architecture of the Dell EMC Ready Solution for HPC Digital Manufacturing and a description of the building blocks are presented in Section 2. Section 3 describes the system configuration, software and application versions, and the benchmark test cases that were used to measure and analyze the performance of the Dell EMC HPC Ready Solution for HPC Digital Manufacturing. Section 4 presents benchmark performance for ANSYS CFX. Section 5 contains the performance data for ANSYS Fluent, and section 6 contains the performance information for ANSYS Mechanical. Section 7 describes the performance of the Basic Building Block solution for modest HPC deployments.

2 System Building Blocks

The Dell EMC Ready Solution for HPC Digital Manufacturing is designed using preconfigured building blocks. The building block architecture allows an HPC system to be optimally designed for specific end-user requirements, while still making use of standardized, domain-specific system recommendations. The available building blocks are infrastructure servers, storage, networking, and compute building blocks. Configuration recommendations are provided for each of the building blocks which provide good performance for typical applications and workloads within the manufacturing domain. This section describes the available building blocks along with the recommended server configurations.

With this flexible building block approach, appropriately sized HPC clusters can be designed based on individual customer workloads and requirements. Figure 1 shows three example HPC clusters designed using the Dell EMC Ready Solutions for HPC Digital Manufacturing architecture.

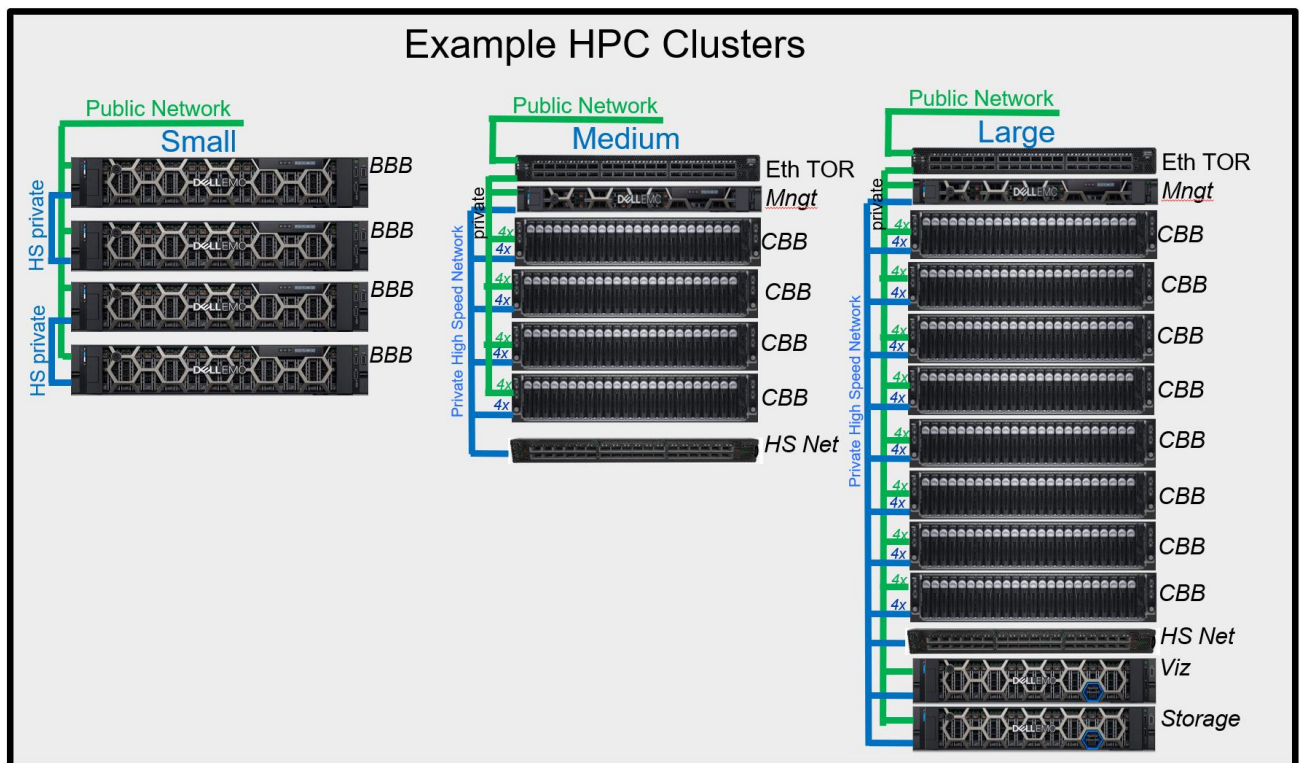


Figure 1 Example Ready Solutions for HPC Digital Manufacturing

2.1 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 master 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 master node is mandatory for an HPC system to deploy and manage the system. If high-availability (HA) management functionality is required, two master nodes 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 R640 server
- Dual Intel® Xeon® Bronze 3106 processors
- 192 GB of RAM (12 x 16GB 2667 MTps DIMMs)
- PERC H330 RAID controller
- 2 x 480GB Mixed-Use SATA SSD RAID 1
- Dell EMC iDRAC9 Enterprise
- 2 x 750 W power supply units (PSUs)
- Mellanox EDR InfiniBand™ (optional)

The recommended base configuration for the infrastructure server is described as follows. The PowerEdge R640 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 Bronze 3106 processor, with 8 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. 192 GB of memory provided by twelve 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 a single mixed-use SATA SSD 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, EDR InfiniBand is likely to be the data interconnect of choice, which provides a high-throughput, low-latency fabric for node-to-node communications or access to a Dell EMC Ready Solution for HPC NFS Storage solution or a Dell EMC Ready Solution for HPC Lustre Storage solution.

2.2 Compute Building Blocks

Compute Building Blocks (CBB) provide the computational resources for most HPC systems for Digital Manufacturing. These servers are used to run the ANSYS CFX, Fluent or Mechanical simulations. The best configuration for these servers depends on the specific mix of applications and types of simulations being performed by each customer. Since the best configuration may be different for each customer, a table of recommended options are provided that are appropriate for these servers. The specific configuration can then be selected based on the specific system and workload requirements of each customer. Relevant criteria to consider when making these selections are discussed in the application performance chapters of this white paper. The recommended configuration options for the Compute Building Block are provided in Table 1.

Table 1 Recommended Configurations for the Compute Building Block

Platforms	Dell EMC PowerEdge R640 Dell EMC PowerEdge C6420
Processors	Dual Intel Xeon Gold 6242 (16 cores per socket) Dual Intel Xeon Gold 6248 (20 cores per socket) Dual Intel Xeon Gold 6252 (24 cores per socket)
Memory Options	192 GB (12 x 16GB 2933 MTps DIMMs) 384 GB (12 x 32GB 2933 MTps DIMMs) 768 GB (24 x 32GB 2933 MTps DIMMs, R640 only)
Storage Options	PERC H330, H730P or H740P RAID controller 2 x 480GB Mixed-Use SATA SSD RAID 0 4 x 480GB Mixed-Use SATA SSD RAID 0
iDRAC	iDRAC9 Enterprise (R640) iDRAC9 Express (C6420)
Power Supplies	2 x 750W PSU (R640) 2 x 2000W PSU (C6400)
Networking	Mellanox® ConnectX®-5 EDR InfiniBand™ adapter

2.3 Basic Building Blocks

Basic Building Block (BBB) servers are selected by customers to create simple but powerful HPC systems. These servers are appropriate for smaller HPC systems where reducing the management complexity of the HPC system is important. The BBB is based on the 4-socket Dell EMC PowerEdge R840 server.

The recommended configuration for BBB servers is:

- Dell EMC PowerEdge R840 server
- Quad Intel Xeon Gold 6142 processors
- 384 GB of RAM (24 x 16GB 2666 MTps DIMMS)
- PERC H740P RAID controller
- 2 x 240GB Read-Intensive SATA SSD RAID 1 (OS)
- 4 x 480GB Mixed-Use SATA SSD RAID 0 (scratch)
- Dell EMC iDRAC9 Enterprise
- 2 x 1600W power supply units (PSUs)
- Mellanox ConnectX-5 EDR InfiniBand (optional)
- Mellanox 25 GbE (optional)

The R840 platform is used to minimize server count and provide good compute power per server. Each server can contain up to four Intel Xeon processors, where each BBB is essentially two CBB's fused into a single server. The Intel Xeon Gold 6142 processor is a sixteen-core CPU with a base frequency of 2.6 GHz and a max all-core turbo frequency of 3.3 GHz. With four processors, a BBB contains 64 cores, a natural number for many CAE simulations. A memory configuration of 24 x 16GB DIMMs is used to provide balanced performance and capacity. While 384GB is typically sufficient for most CAE workloads, customers expecting to handle larger production jobs should consider increasing the memory capacity to 768GB. Various CAE applications, such as implicit FEA, often have large file system I/O requirements and four Mixed-use SATA

SSD's in RAID 0 are used to provide fast local I/O. The compute nodes do not normally require extensive OOB management capabilities; therefore, an iDRAC9 Express is recommended.

Additionally, two BBB's can be directly coupled together via a high-speed network cable, such as InfiniBand or Ethernet, without need of an additional high-speed switch if additional compute capability is required for each simulation run (HPC Couplet). BBB's provide a simple framework for customers to incrementally grow the size and power of the HPC cluster by purchasing individual BBBs, BBB Couplets, or combining the individual and/or Couplets with a high-speed switch into a single monolithic system.

Performance testing for BBB's has been done using both Linux and Windows Server 2016. In general, Linux provides better overall performance, and an easier path to combining BBB's to create larger, more capable HPC clusters. We have tested up to two BBB's with Linux using both 25 Gigabit Ethernet and EDR InfiniBand adapters/cable in the two-node couplet configuration. With Linux, the EDR couplet gave the best overall performance across our tests. While the performance of the 25GbE based couplet was often comparable to the EDR based couplet, we saw little reason not to use EDR to ensure better overall performance at a similar cost and complexity. However, for customers not wishing to deploy InfiniBand in their environment, choosing a 25GbE based couplet is a suitable alternative.

For Windows testing, we tested only a couplet with a 25 GbE network. Support for InfiniBand on Windows is not currently feasible for most customers. Customers wishing for the highest level of performance, and potential cluster expansion would be advised to use Linux as an operating system.

2.4 Storage

Dell EMC offers a wide range of HPC storage solutions. For a general overview of the entire 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 of time. 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 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® Bronze 4110 processors
- 96 GB of memory, 12 x 8GB 2667 MT/s DIMMS
- PERC H730P RAID controller
- 2 x 250GB Mixed-use SATA SSD in RAID-1 (For OS)
- 12 x 12TB 3.5: nSAS HDDs in RAID-6 (for data)
- Dell EMC iDRAC9 Express
- 2 x 750 W power supply units (PSUs)
- Mellanox EDR 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 Ready Solution for HPC NFS Storage (NSS), shown in Figure 2, with up 840 TB of raw storage capacity may be appropriate:

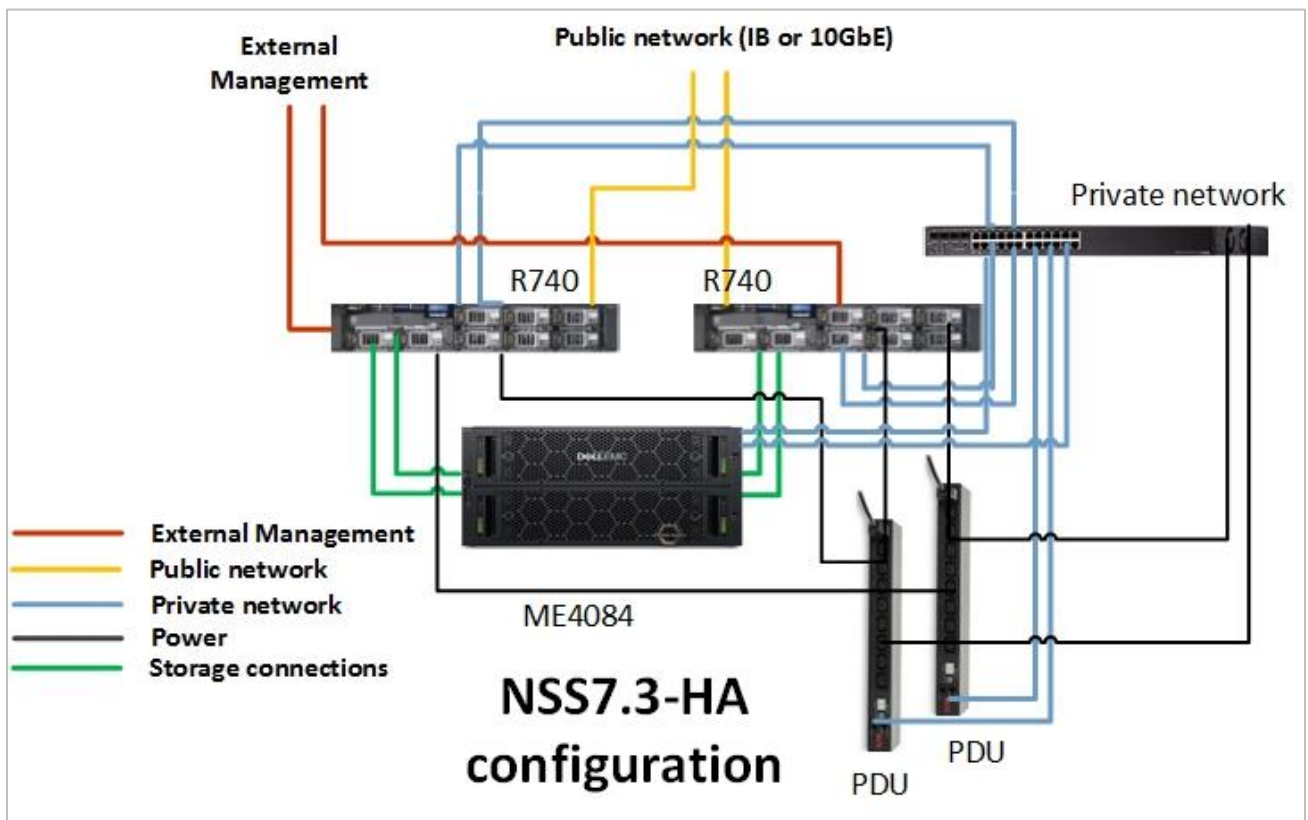


Figure 2 NSS7.3-HA Storage System Architecture

For customers desiring a shared high-performance parallel filesystem, the Dell EMC Ready Solution for HPC Lustre Storage solution shown in Figure 3 is appropriate. This solution can scale up to multiple petabytes of storage.

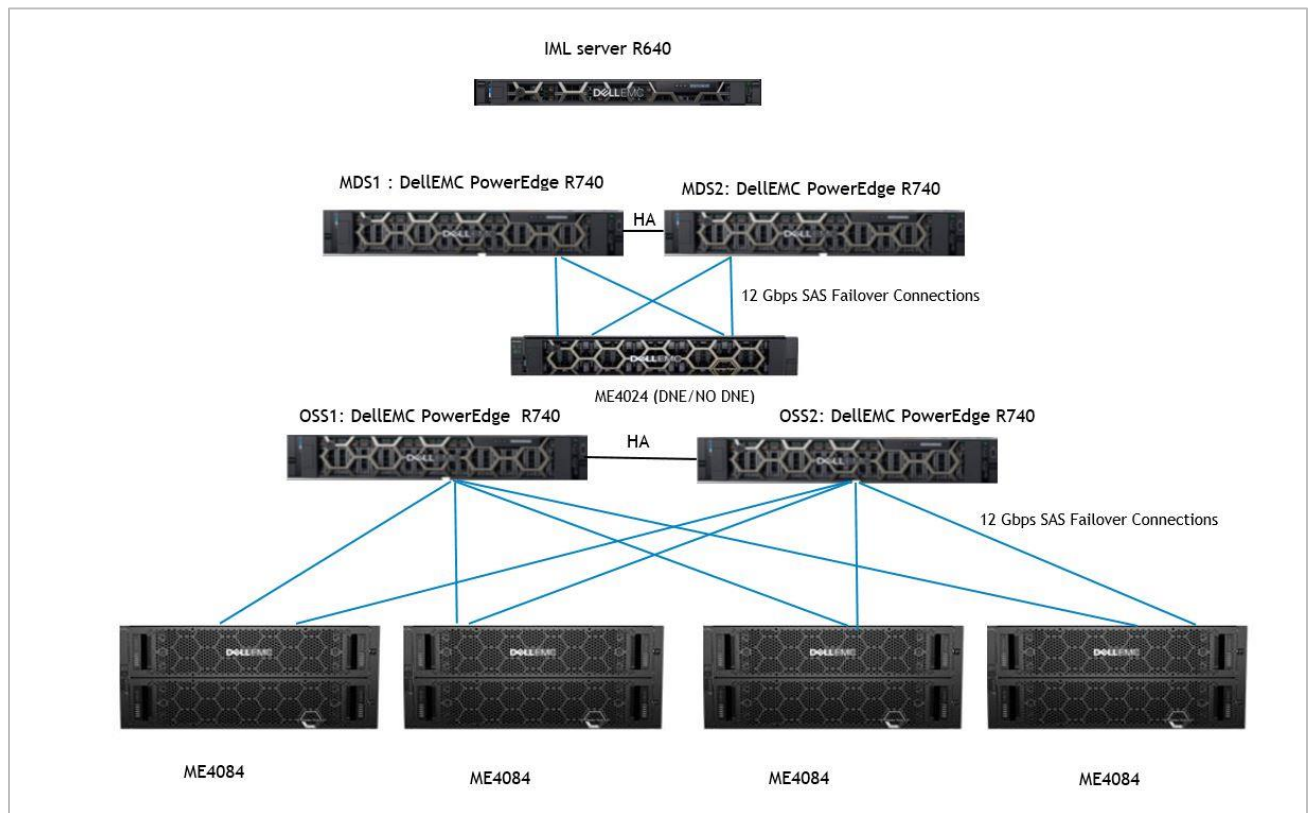


Figure 3 Dell EMC Ready Solution for Lustre Storage Reference Architecture

2.5 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 CBB servers, this network will also be used for IPMI hardware 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 should 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 EDR InfiniBand fabric. The fabric will typically be assembled using Mellanox SB7890 36-port EDR InfiniBand switches. The number of switches required depends on the size of the cluster and the blocking ratio of the fabric.

2.6 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.

2.7 Services and Support

The Dell EMC Ready Solution for HPC Digital Manufacturing is available with full hardware support and deployment services, including additional HPC system support options.

3 Reference System

The reference system was assembled in the Dell EMC HPC and AI Innovation Lab using the building blocks described in section 2. The building blocks used for the reference system are listed in Table 2.

Table 2 Reference System Configuration

Building Block	Quantity
Infrastructure Server	1
Computational Building Block (CBB) PowerEdge C6420 Dual Intel Xeon Gold 6242 192GB RAM 12x16GB 2933 MTps DIMMs Mellanox ConnectX-5 EDR adapter	2
Computational Building Block (CBB) PowerEdge C6420 Dual Intel Xeon Gold 6252 192 GB RAM 12x16GB 2933 MTps DIMMs Mellanox ConnectX-5 EDR adapter	8
Basic Building Block	2
Dell Networking S3048-ON Ethernet Switch	1
Mellanox SB7700 EDR 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
System Profile	Performance Profile
Sub NUMA Cluster	Enabled (CBB) Disabled (BBB)

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

Table 4 Software Versions

Component	Version
Operating System	RHEL 7.6 Windows Server 2016 (BBB)
Kernel	3.10.0-957.el7.x86_64
OFED	Mellanox 4.5-1.0.1.0
Bright Cluster Manager	8.2
ANSYS CFX	2019R1 (19.2 for Windows)
ANSYS Fluent	2019R1 (19.2 for Windows)
ANSYS Mechanical	2019R1 (19.2 for Windows)

4 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 in the solver section. However, some simulations, such as large transient analysis, may have greater I/O demands. Figure 4 shows the measured performance of the standard CFX v16 benchmarks using CFX 2019R1 on a single server. The performance for each benchmark is measured using the solver elapsed time.




Figure 4 Single Server Relative Performance—ANSYS CFX 2019R1

The results in Figure 4 are plotted relative to the performance of a single server configured with Intel Xeon Gold 6242 processors. For comparison, the figure also includes performance data for prior generations of the Ready Solution for HPC Digital Manufacturing. This includes the 14G CBB with Intel Xeon Gold 6142 processors and the 13G CBB with 16-core Intel Xeon E5-2697A v4 processors. Higher results indicate better overall performance. The result for the Airfoil_100M benchmark is not provided for the E5-2697A v4 processor as there was insufficient memory capacity to run this benchmark. These results show the performance advantage available with 14G servers with 2nd generation Intel Xeon Scalable processors (code name Cascade Lake). These results also show that the CFX benchmarks can make use of the higher core count processors such as the Intel Xeon Gold 6252.




Figure 5 ANSYS CFX Parallel Scaling—Intel Xeon Gold 6252

Figure 5 presents the parallel scalability when running CFX using up to eight CBB nodes configured with Intel Xeon Gold 6252 processors. The performance is presented relative to the performance of a single node (48 cores total).

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 problem dataset fits in cache as the number of nodes used increases.

5 ANSYS Fluent Performance

ANSYS Fluent is a Computational Fluid Dynamics (CFD) application commonly used across a very wide range of CFD and multi-physics 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 in the solver section. However, some simulations, such as large transient analysis, may have greater I/O demands. Fifteen benchmark problems from the Fluent benchmark suite v18 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, $(\text{total seconds in a day}) / (25 \text{ iteration solve time in seconds})$. A higher value represents better performance. Figure 6 shows the measured performance for six of the ANSYS Fluent benchmarks on a single server.

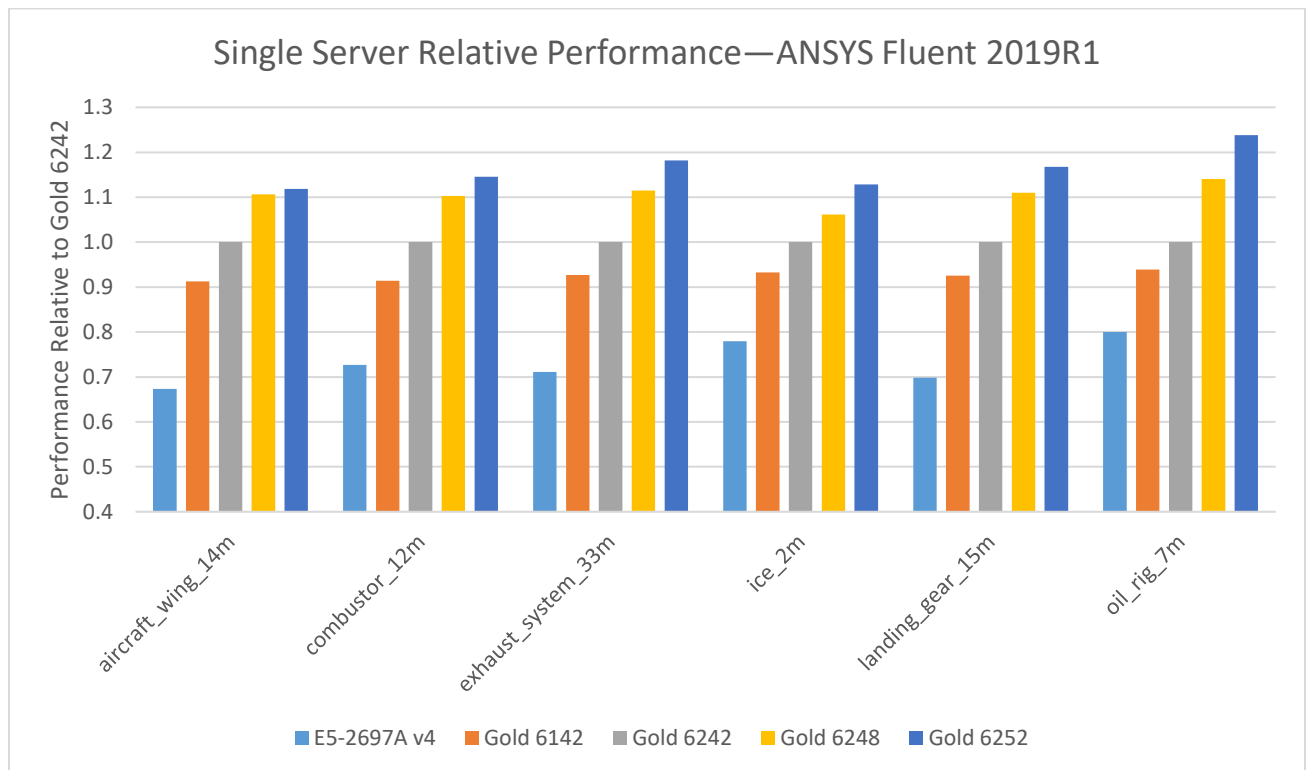


Figure 6 Single Server Relative Performance—ANSYS Fluent 2019R1

The results in Figure 6 are plotted relative to the performance of a single server configured with Intel Xeon Gold 6242 processors. For comparison, the figure also includes performance data for prior generations of the Ready Solution for HPC Digital Manufacturing. This includes the 14G CBB with Intel Xeon Gold 6142 processors and the 13G CBB with 16-core Intel Xeon E5-2697A v4 processors. Higher results indicate better overall performance. These results show the performance advantage available with 14G servers with 2nd generation Intel Xeon Scalable processors (code name Cascade Lake). These results also show that most of the Fluent benchmarks can make use of the higher core count processors such as the Intel Xeon Gold 6252.

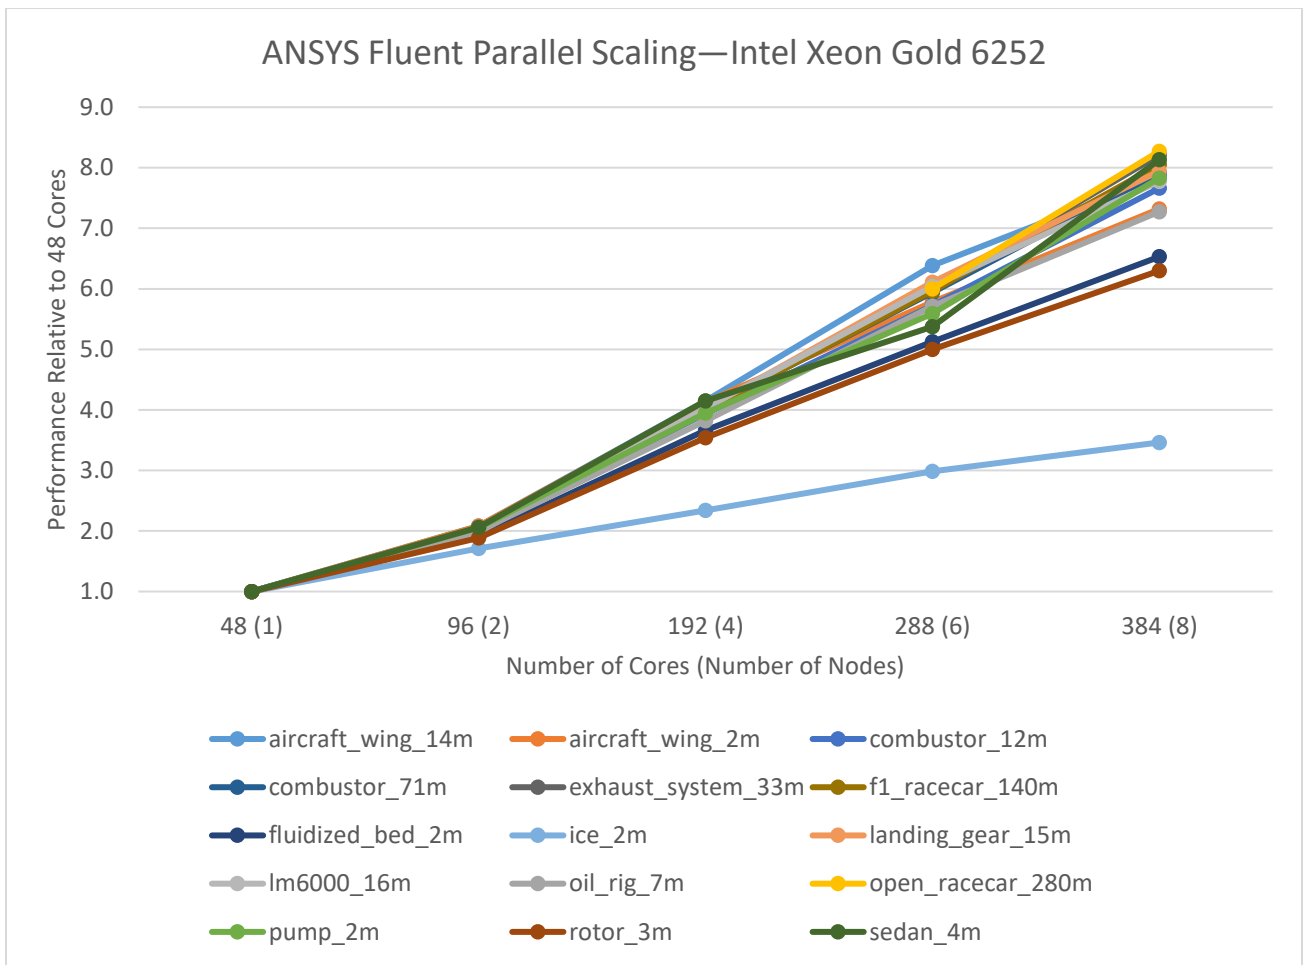


Figure 7 ANSYS Fluent Parallel Scaling—Intel Xeon Gold 6252

Figure 7 presents the parallel scalability of the Fluent benchmark models using up to eight CBB nodes configured with Intel Xeon Gold 6252 processors. The performance is presented relative to the performance of a single node (48 cores total).

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 normal as it is a small model which includes dynamic mesh and combustion simulation.

6 ANSYS Mechanical Performance

ANSYS Mechanical is a multi-physics Finite Element Analysis (FEA) software commonly used in multiple engineering disciplines. Depending on the specific problem types, FEA codes may or may not scale well across multiple processor cores and servers. 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 Building Block (CBB) 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 384GB Memory CBB Configuration for ANSYS Mechanical

Platform	Dell EMC PowerEdge C6420
Processor	Dual Intel Xeon Gold 6242 (16 cores per socket)
Memory	12 x 32GB 2933 MTps DIMMS (384GB)
Storage	PERC H730P RAID controller 4 x 480GB Mixed-Use SATA SSD RAID 0
iDRAC	iDRAC9 Express
Power Supplies	2 x 2000W PSU
Networking	Mellanox® ConnectX®-5 EDR InfiniBand™ adapter

A sample CBB 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 768GB and using the more advanced PERC H740P RAID controller.

Table 6 Example 768GB Memory CBB Configuration for ANSYS Mechanical

Platform	Dell EMC PowerEdge R640
Processor	Dual Intel Xeon Gold 6242 (16 cores per socket)
Memory	24 x 32GB 2933 MTps DIMMS (768GB)
Storage	PERC H740P RAID controller 4 x 480GB Mixed-Use SATA SSD RAID 0
iDRAC	iDRAC9 Enterprise
Power Supplies	2 x 750W PSU
Networking	Mellanox® ConnectX®-5 EDR InfiniBand™ adapter

The performance of the ten standard ANSYS Mechanical v19 benchmark cases were evaluated on the reference system. The ten benchmark cases all run in-core with the 192GB of RAM that is available per compute node 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 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, using all available cores on each server. Performance is measured using the Core Solver Rating metric. This metric represents the performance of the solver core which excludes any analysis pre and post-processing.

Figure 8 and Figure 9 show the measured performance of the ANSYS Mechanical benchmarks on a single CBB server. The results are plotted relative to the performance of a single CBB configured with Intel Xeon Gold 6242 processors. For comparison, the figures also include performance data for prior generations of the Ready Solution for HPC Digital Manufacturing. This includes the 14G CBB with Intel Xeon Gold 6142 processors and the 13G CBB with 16-core Intel Xeon E5-2697A v4 processors. Higher results indicate better overall performance. These results show the performance advantage available with 14G servers with 2nd generation Intel Xeon Scalable processors (code name Cascade Lake). The results with the higher core count processors are mixed, which indicates that the 16-core Intel Xeon Gold 6242 processor is a good choice for general use with ANSYS Mechanical.

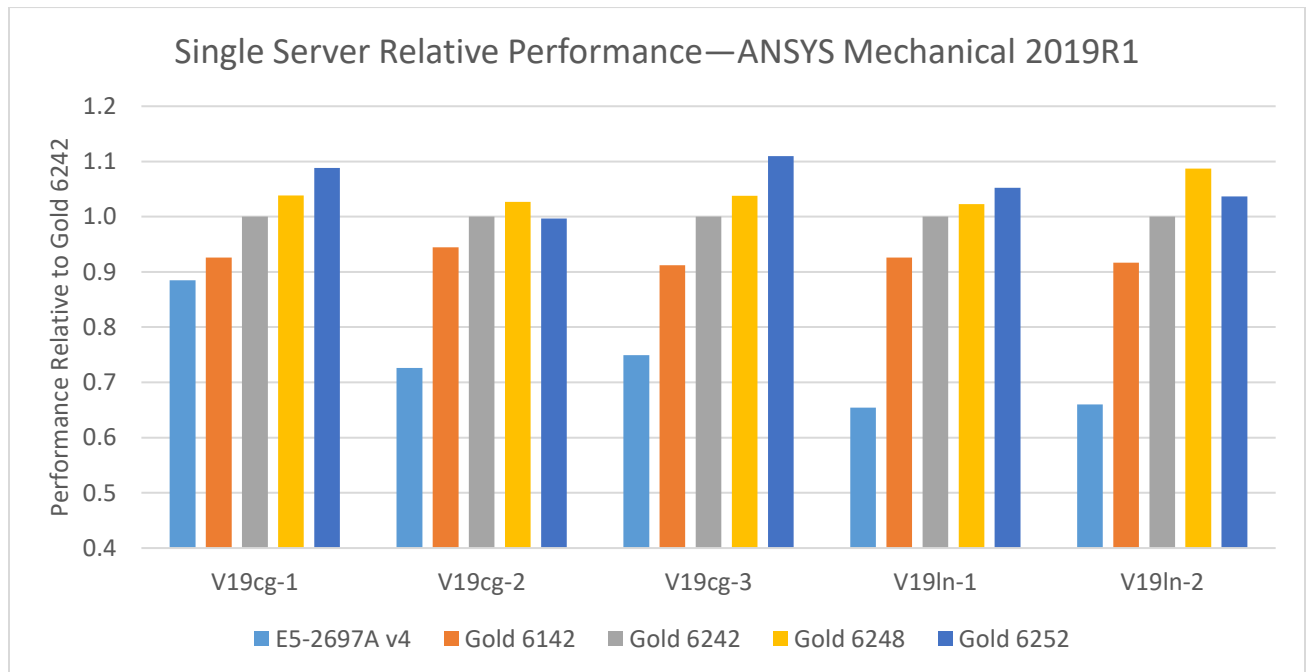


Figure 8 Single Server Relative Performance—ANSYS Mechanical 2019R1




Figure 9 Single Server Relative Performance—ANSYS Mechanical 2019R1

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 when using a single processor core. 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 single core (1-core) result.




Figure 10 ANSYS Mechanical Single Node Scaling—Intel Xeon Gold 6252

The results in Figure 10 demonstrate that the ANSYS Mechanical DMP solver scales reasonably well from one to 32 processor cores. Using more than 32 cores on a single server results in minimal performance

increases for most of the benchmark models. This data shows that using 32 cores per node works well for the standard benchmark models.

Performance results for the ANSYS Mechanical solver using multiple servers are shown in Figure 11. The results are plotted relative to the benchmark performance when using a single server. 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 single node (48-core) result.




Figure 11 ANSYS Mechanical Parallel Scaling—Intel Xeon Gold 6252

The DMP solver scales up to eight nodes or 384 processor cores 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.

6.1.1 GPGPU Performance


ANSYS Mechanical supports using NVIDIA® GPUs to accelerate some of the analysis types that are available with ANSYS Mechanical. Each GPU is counted as one core for software licensing purposes so there is interest in utilizing GPUs to accelerate solutions while minimizing software license usage. Please refer to the ANSYS Mechanical APDL Parallel Processing Guide for further details on supported GPUs, solver behavior when using GPUs and supported analysis types.

A PowerEdge C4140 configured with Intel Xeon Gold 6148 processors and four NVIDIA TESLA 32GB V100 SXM2 GPUs was used to obtain performance benchmark results with GPUs. The system configuration is shown in Table 7.

Table 7 PowerEdge C4140 System Configuration

Platform	Dell EMC PowerEdge C4140
Processor	Dual Intel Xeon Gold 6148
Memory	12 x 16GB 2666 MTps DIMMS (192 GB)
GPU	4 x NVIDIA® Tesla® V100 32GB SXM2
OS	Red Hat Enterprise Linux Server 7.4
CUDA Toolkit	9.2.88

The base software license for ANSYS Mechanical allows using up to four CPU cores and/or GPUs. For example, 4 CPU cores or 2 CPU cores plus 2 GPUs would both consume one software license. The results in Figure 12 show the benchmark performance when using a single software license for various combinations of CPU cores and GPUs. The performance results are mixed, with some cases performing very well when utilizing GPUs and others performing worse with GPUs.

**Figure 12 Base License Relative Performance—ANSYS Mechanical 2019R1**

The results in Figure 13 show the benchmark performance when using all CPU cores in the server with zero to four NVIDIA Tesla V100 GPUs. The results are plotted relative to the performance of the CPU only result. These results are also mixed; however, most of the benchmarks show some performance improvement when utilizing GPUs.




Figure 13 Server Performance with NVIDIA Tesla V100 GPUs—ANSYS Mechanical 2019R1

These benchmark results demonstrate that GPUs may provide some performance benefit for ANSYS Mechanical; however, the performance with GPUs is dependent on the specific solution types being utilized and some solution types are unable to use GPUs. The exact performance benefit of GPUs is highly dependent on the specific problems being simulated. Because of this, it would be incumbent upon each user to determine whether utilizing GPUs would be appropriate for their simulations.

7 Basic Building Block Performance

The performance of systems created with Basic Build Blocks (BBB's) was tested for ANSYS CFX, Fluent and Mechanical. Performance was measured on both a single BBB and a BBB couplet composed of two BBB's with a direct high-speed network connection. Testing was performed using both RHEL 7 and Windows 2016 Enterprise Edition. For Linux, the results were obtained using an EDR InfiniBand fabric and for Windows a 25 GbE network, since this is how we envision customers using these systems.

Figure 14 shows the results obtained for the Fluent standard benchmarks with the system configurations mentioned above.

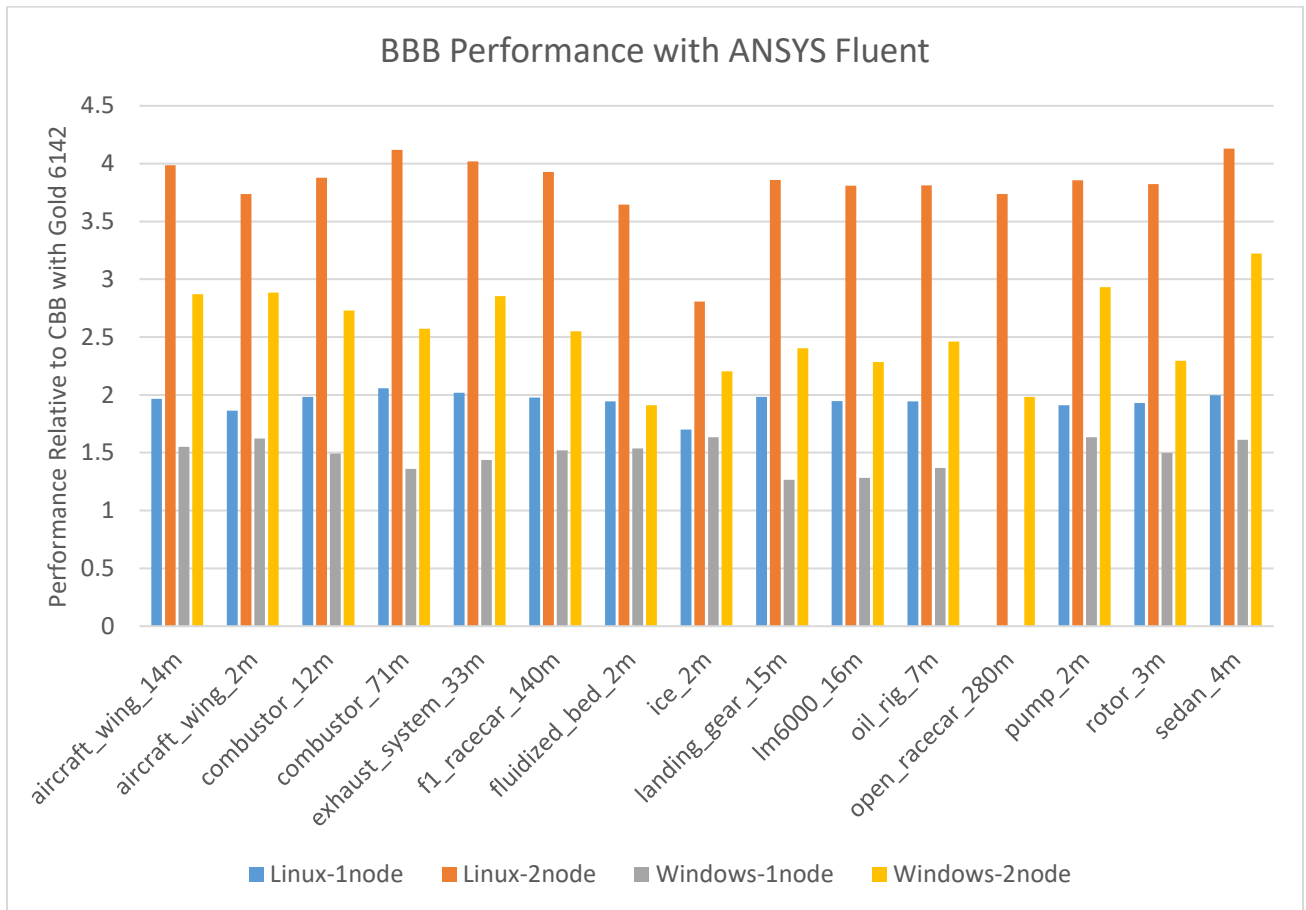


Figure 14 BBB Performance with ANSYS Fluent

The vertical axis in this figure represents the measured performance relative to a single 32-core CBB configured with Intel Xeon Gold 6142 processors. The benchmark performance was measured in jobs/day, so for this figure, higher is better. Nominally, we would expect a single BBB to performance to be about twice as fast as a single CBB, since it has four Intel Xeon 16-core 6142 processors compared to two of the same processors in the CBB. Of course, this presumes that the benchmarks models are capable of linear performance scaling from 32-cores to 64-cores. Likewise, we would expect a 2-node couplet to have a nominal performance equivalent to 4 CBBs. Overall, the data is consistent with the performance expectations, except for the ice_2m benchmark, which has limited parallel scalability as shown previously in Figure 7.

The results using Windows are not quite as good as with Linux. On average, the performance of a single Windows based BBB is about 1.5X that of the Linux based CBB, and the Windows 25GbE based couplet delivers about 2.5X the baseline CBB performance. However, for customers comfortable with managing Windows systems, this offers a significant performance potential over most existing Windows based workstations.

The results for a similar comparison for ANSYS CFX are shown in Figure 15.

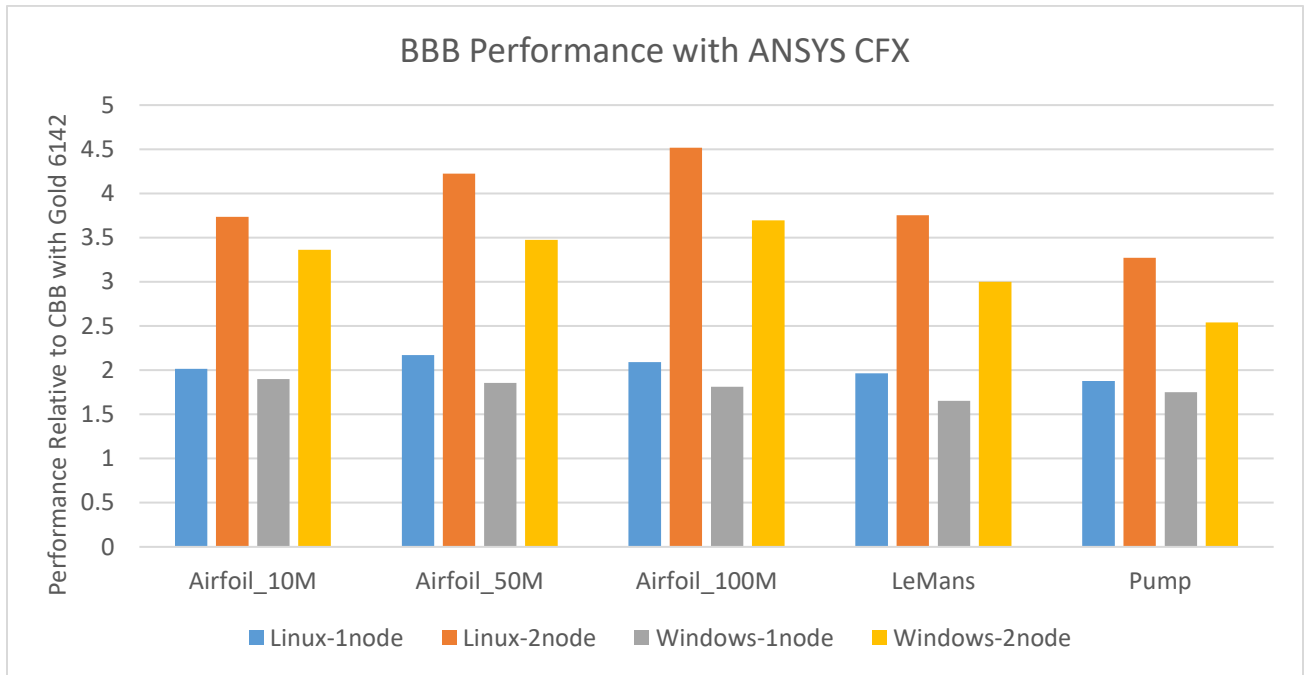


Figure 15 BBB Performance with ANSYS CFX

Again, the vertical axis denotes performance in terms of a single Linux based CBB, where the single BBB server performance would be expected to be about 2.0. Overall the single BBB node performance is close to 2X for both the Linux and Windows based solutions. For the two node couplets, the Linux based performance is on par with the CBB cluster performance, whereas the Windows couplet offers noticeably less overall performance compared to the Linux based counterpart. However, the Windows based couplet offers an overall high level of solution performance.

Figure 16 shows the BBB performance for the Mechanical standard v19 benchmarks.




Figure 16 BBB Performance with ANSYS Mechanical

Unlike the Fluent and CFX performance figures, the ANSYS Mechanical performance figure shows the performance of only a single BBB server. As shown in Figure 16 above, the standard benchmarks included with ANSYS Mechanical tend to exhibit modest parallel scalability. For these single node benchmarks, benchmarks were carried out using either 32 cores or 64 cores of the BBB. The results for the Linux based BBB were essentially equivalent to those obtained with either a single CBB or two CBBs connected via an EDR InfiniBand network switch, as expected. For the Windows system, both the 32-core and 64-core benchmarks were noticeably slower than the Linux based results. Further testing needs to be carried out to better understand this performance gap. While the Windows results demonstrate poorer performance than the Linux solutions, they still offer customers the potential for significantly better performance than what can be obtained on a two-socket Windows based workstation.

Overall, these benchmarks display the significant performance potential that can be obtained by customers desiring to improve their CAE capabilities without installing and maintaining a traditional Linux based HPC cluster. Customers could incrementally increase their solution capabilities by adding more BBBs, aggregating them into more powerful couplets, re-provision them from Windows to Linux, and eventually aggregate them with an InfiniBand switch into a single HPC cluster based on BBBs, all while preserving their hardware investment, and minimizing production downtime.

8 Conclusion

This technical white paper presents the Dell EMC Ready Solution for HPC Digital Manufacturing. The detailed analysis of the building block configurations demonstrate that the system is architected for a specific purpose—to provide a comprehensive HPC solution for the manufacturing domain. Use of this building block approach allows customers to easily deploy an HPC system optimized for their 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 system performance with ANSYS software.