

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

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

Joshua Weage
Martin Feyereisen

Dell EMC HPC Innovation Lab

January 2018

Revisions

Date	Description
January 2018	Initial release

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 © 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [2/12/2018] [Technical White Paper]

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

Table of contents

- Revisions.....2
- 1 Introduction.....4
- 2 System Building Blocks5
 - 2.1 Infrastructure Servers5
 - 2.2 Explicit Building Blocks.....6
 - 2.3 Implicit Building Blocks6
 - 2.4 Dell EMC NSS-HA Storage7
 - 2.5 Dell EMC IEEL Storage8
 - 2.6 System Networks.....8
 - 2.7 Cluster Software9
 - 2.8 Services and Support9
- 3 Reference System.....10
- 4 Fluent Performance.....12
- 5 CFX Performance.....15
- 6 Mechanical Performance.....17
- 7 Conclusion.....20
- References.....21

Introduction

This technical white paper describes the performance of ANSYS® Fluent®, ANSYS® Mechanical™, and ANSYS® CFX® on the Dell EMC Ready Bundle for HPC Digital Manufacturing, which 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 Bundle for HPC Digital Manufacturing uses a flexible building block approach, where individual building blocks can be combined to build HPC systems which are ideal for customer specific work-loads and use cases. The individual building blocks are configured to provide good performance for specific application types and workloads common to this industry.

The architecture of the Dell EMC Ready Bundle 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 Bundle for HPC Digital Manufacturing. Section 4 quantifies the capabilities of the system and presents benchmark performance for ANSYS Fluent. Section 5 contains the performance data for CFX, and section 6 contains the performance information for ANSYS Mechanical.

2 System Building Blocks

The Dell EMC Ready Bundle for HPC Digital Manufacturing is assembled by using preconfigured building blocks. The available building blocks are infrastructure servers, storage, networking, and application specific compute building blocks. These building blocks are preconfigured to provide good performance for typical applications and workloads within the manufacturing domain. The building block architecture allows for a customizable HPC system based on specific end-user requirements, while still making use of standardized, domain-specific building blocks. This section describes the available building blocks along with the rationale of the recommended system configurations.

2.1 Infrastructure Servers

The infrastructure servers are used to administer the system and provide user access. They are not typically involved in computation or storage, but they provide services that are critical to the overall HPC system. Typically these servers are the master nodes and the login nodes. For small sized clusters, a single physical server can provide these functions. The infrastructure server can also be used for storage, by using NFS, in which case it must be configured with additional disk drives or an external storage array. One master node is mandatory and is required to deploy and manage the system. If high-availability (HA) 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 memory, 12 x 16GB 2667 MT/s DIMMS
- PERC H330 RAID controller
- 1 x 800GB Mixed-use SATA SSD
- 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 here. The PowerEdge R640 server is suited for this role. A cluster will have only a small number of infrastructure servers; therefore, density is not a concern, but manageability is more important. The Intel Xeon 3106 processor, with 8 cores per socket, is sufficient for this role. 192 GB of memory provided by 12x16 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 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-node communications, or access to a Dell EMC NFS Storage Solution (NSS) or Dell EMC Intel Enterprise Edition for Lustre (IEEL) storage solution.

2.2 Explicit Building Blocks

Explicit Building Block (EBB) servers are typically used for Computational Fluid Dynamics (CFD) solvers such as ANSYS Fluent®. These software applications typically scale well across many processor cores and multiple servers. The memory capacity requirements are typically modest and these solvers perform minimal disk I/O while solving. In most HPC systems used for Digital Manufacturing, the large majority of servers are EBB servers.

The recommended configuration for EBBs is:

- Dell EMC PowerEdge C6420 server
- Dual Intel® Xeon® Gold 6142 processors
- 192 GB of memory, 12 x 16GB 2667 MT/s DIMMS
- PERC H330 RAID controller
- 2 x 480GB Mixed-use SATA SSD in RAID 0
- Dell EMC iDRAC9 Enterprise
- 2 x 1600 W power supply units per chassis
- Mellanox EDR InfiniBand™ (optional)

The recommended configuration for the EBB servers is described here. Because the largest percentage of servers in the majority of systems will be EBB servers, a dense solution is important; therefore, the PowerEdge C6420 server is selected. The Intel Xeon Gold 6142 processor is a 16-core CPU with a base frequency of 2.6 GHz and a maximum all-core turbo frequency of 3.3 GHz. 32 cores per server provides a dense compute solution, with good memory bandwidth per core, and a power of two quantity of cores. The maximum all-core turbo frequency is important because EBB applications are typically CPU bound. This CPU model provides the best balance of CPU cores and core speed. 192 GB of memory using 12x16GB DIMMs provides sufficient memory capacity, with minimal cost per GB, while also providing good memory bandwidth. Relevant applications typically perform limited I/O while solving; therefore, the system is configured with two disks in RAID 0 using the PERC H330 RAID controller, which leaves the PCIe slot available for an EDR InfiniBand HCA. The compute nodes do not require extensive out-of-band (OOB) management capabilities; therefore, an iDRAC9 Express is sufficient. For small systems (four nodes or less), an Ethernet network may provide sufficient 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-node communications, or access to an NSS or IEEL storage solution.

2.3 Implicit Building Blocks

Implicit Building Block (IBB) servers are typically used for implicit FEA solvers such as ANSYS Mechanical. These applications typically have large memory requirements and do not scale to as many cores as the EBB applications. File system I/O performance can also have a significant effect on application performance.

The recommended configuration for IBB servers is:

- Dell EMC PowerEdge R640 server
- Dual Intel® Xeon® Gold 6136 processors
- 384 GB of memory, 24 x 16GB 2667 MT/s DIMMS
- PERC H740P RAID controller
- 4 x 480GB Mixed-use SATA SSD in RAID 0
- Dell EMC iDRAC9 Express
- 2 x 750 W power supply units (PSUs)
- Mellanox EDR InfiniBand™ (optional)

The recommended configuration for the IBB servers is described here. Typically, a smaller percentage of the system will be comprised of IBB servers. Because of the memory and drive recommendations explained here, a 1U PowerEdge R640 server is a good choice. The Intel Xeon Gold 6136 processor is a twelve-core CPU with a base frequency of 3.0 GHz and a max all-core turbo frequency of 3.6 GHz. A memory configuration of 24 x 16 GB DIMMs is used to provide the larger memory capacities needed for these applications. While 384GB is typically sufficient for most CAE workloads, customers expecting to handle very large production jobs should consider increasing the memory capacity to 768GB. IBB applications 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 require extensive OOB management capabilities; therefore, an iDRAC9 Express is recommended. InfiniBand is not typically necessary for IBBs because most uses cases only require running applications on a single IBB; however, an InfiniBand HCA can be added to enable multi-server analysis or to access an NSS or IEEL storage solution.

2.4 Dell EMC NSS-HA Storage

The Dell EMC NFS Storage Solution (NSS) provides a tuned NFS storage option that can be used as primary storage for user home directories and application data. The current version of NSS is NSS7.0-HA with options of 240 TB or 480 TB raw disk space. NSS is an optional component and a cluster can be configured without NSS.

NSS-HA is a high performance computing network file system (NFS) storage solution from Dell EMC, optimized for performance, availability, resilience, and data reliability. The best practices used to implement this solution result in better throughput compared to non-optimized systems. A high availability (HA) setup, with an active-passive pair of servers, provides a reliable and available storage service to the compute nodes. The HA unit consists of a pair of Dell EMC PowerEdge R730 servers. A Dell EMC PowerVault MD3460 dense storage enclosure provides 240 TB of storage for the file system with 60 x 4 TB, 7.2K near-line SAS drives. This unit can be extended with a PowerVault MD3060e to provide an additional 240 TB of disk space for the 480 TB solution. Each of the PowerVault arrays is configured with 6 virtual disks (VDs). Each VD consists of 10 hard drives configured in RAID6 (8+2).

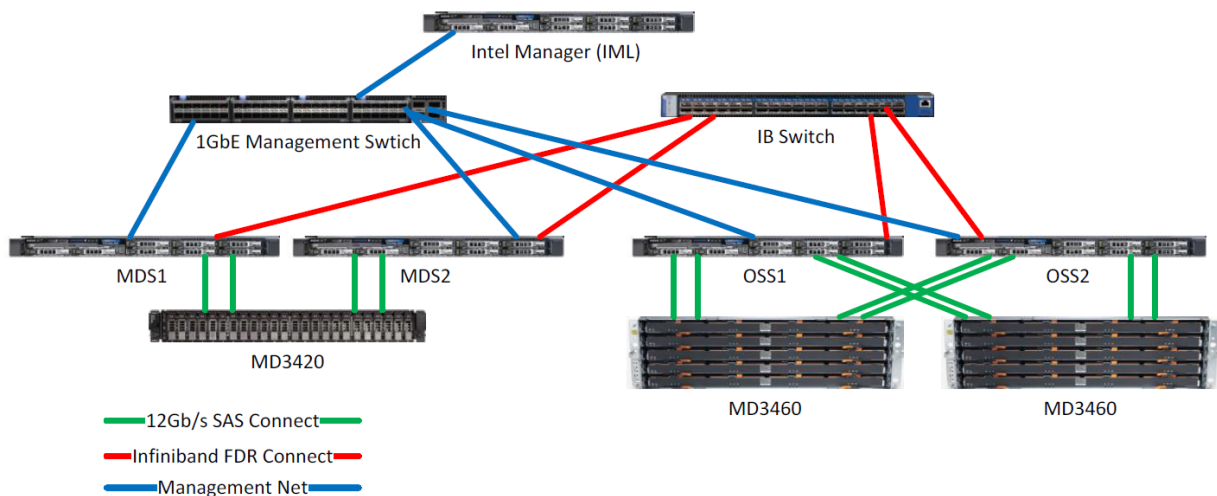
The NFS server nodes are directly attached to the dense storage enclosures by using 12 Gbps SAS connections. NSS7.0-HA provides two network connectivity options for the compute cluster: EDR InfiniBand and 10 Gigabit Ethernet. The active and passive NFS servers run Red Hat Enterprise Linux (RHEL) 7.2 with Red Hat's Scalable File System (XFS) and Red Hat Cluster Suite to implement the HA feature.

2.5 Dell EMC IEEL Storage

Dell EMC IEEL storage is an Intel Enterprise Edition for Lustre (IEEL) based storage solution consisting of a management station, Lustre metadata servers, Lustre object storage servers, and the associated backend storage. The management station provides end-to-end management and monitoring for the entire Lustre storage system.

The Dell EMC IEEL storage solution provides a parallel file system with options of 480 TB or 960 TB raw storage disk space. This solution is typically used for scratch space for larger clusters.

Figure 1 Overview of the Dell EMC IEEL Components and Connectivity



2.6 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 EBB and IBB 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 heartbeat ports for NSS-HA and IEEL Ethernet management ports may also be connected to this network. The management network typically uses the Dell Networking S3048-ON Ethernet switch. If there is more than one switch in the system, multiple switches will be stacked with 10 Gigabit Ethernet stacking 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.

2.7 Cluster Software

The Cluster Software is used to install and monitor the system's compute servers. Bright Cluster Manager (BCM) is the recommended cluster software.

2.8 Services and Support

The Dell EMC Ready Bundle for HPC Digital Manufacturing is available with full hardware support and deployment services, including NSS-HA and IEEL deployment services.

3 Reference System

The reference system was assembled in the Dell EMC HPC Innovation Lab using the building blocks described in Section 2. The building blocks used for the reference system are listed in Table 1.

Table 1. Reference System Configuration

Building Block	Quantity
Infrastructure Server	1
Explicit Building Block with EDR InfiniBand	8
Implicit Building Block	1
Dell Networking S3048-ON Ethernet Switch	1
Mellanox SB7790 EDR InfiniBand Switch	1

The BIOS configuration options used for the reference system are listed in Table 2.

Table 2. BIOS Configuration Options

BIOS Option	Setting
Logical Processor	Disabled
Virtualization Technology	Disabled
System Profile	Performance Optimized

The software versions used for the reference system are listed in Table 3.

Table 3. Software Versions


Component	Version
Operating System	RHEL 7.3
Kernel	3.10.0-514.el7.x86_64
OFED	Mellanox 3.4-2.0.0.0
Bright Cluster Manager	7.3 with RHEL 7.3 (Dell version)
ANSYS Fluent	18.2
ANSYS CFX	18.2
ANSYS Mechanical	18.1

4 Fluent Performance

ANSYS Fluent software is a Computational Fluid Dynamics (CFD) tool commonly used across a very wide range of CFD and multiphysics 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. For these types of application characteristics, the explicit building block servers are appropriate. Fifteen benchmark problems from the Fluent benchmark suite v18 were evaluated on the EBB servers in the reference system.



The results for Fluent are presented by using the Solver Rating metric which counts 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 2 shows the relative performance of the two compute building block types for eight of the ANSYS Fluent benchmarks. For this comparison, all processor cores in the individual building blocks are utilized while running ANSYS Fluent. This comparison demonstrates that for Fluent, application performance is primarily determined by processor performance. The Intel Xeon Gold 6142 processor used in the Explicit BB is a good choice for Fluent.



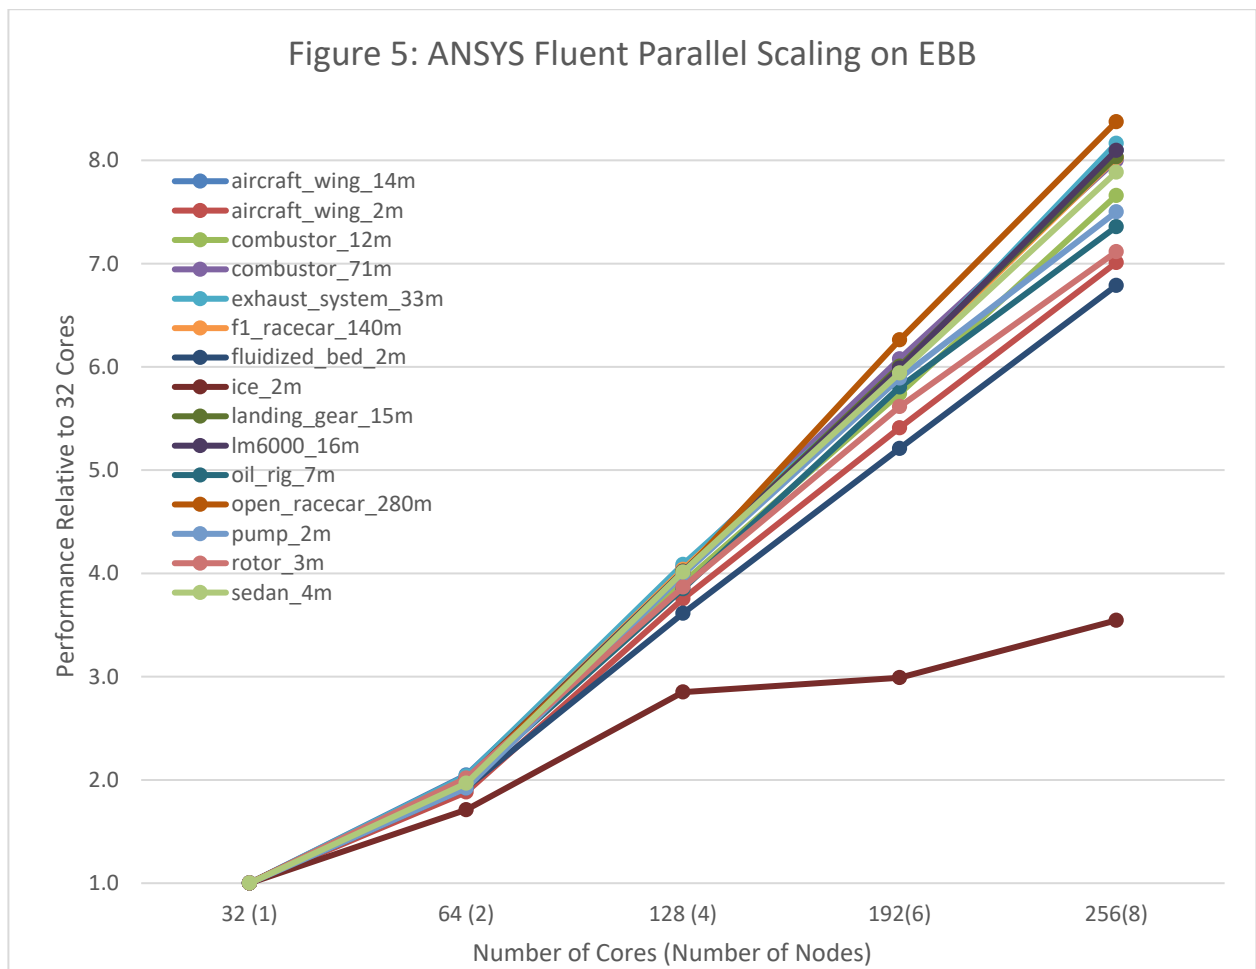
The charts in Figure 3 and Figure 4 show the measured performance of the reference system, from one to eight EBBs, using 32 to 256 cores. Each data point on the graphs records the performance of the specific benchmark data set by using the number of cores marked on the horizontal axis in a parallel simulation. The results are divided into two charts for easy readability—the scale for Solver Rating is large and some models

run much faster than others depending on the number of cells in the model, type of solver used and physics of the problem.




The benchmark models f1_racecar_140m and open_racecar_280m are large models that require two or more servers for sufficient memory capacity. The results for these cases start with the first valid result obtained for the specific problem.

Figure 5 presents the same performance data, but plotted relative to the “32-cores (1 node)” result. It makes it easy to see the scaling of the solution—the performance improvement as more cores are used for the analysis. Please note that a minimum of two nodes was required to run the f1_racecar_140M, and 4 nodes were required to run the open_racecar_280m model. To plot on the same chart with the other data, it is assumed that these models exhibit linear scaling up to 64-way (2-nodes) and 128-way (4-nodes) respectively. Problem scalability depends on the cell count and physics being simulated for the specific benchmark problem. Notably, the scaling on the small model, Ice_2m in not very good. The other benchmarks models all scaled well to 8 nodes (256-way parallel).




5 CFX Performance

ANSYS CFX software is a Computational Fluid Dynamics (CFD) tool similar in many aspects to ANSYS Fluent and recognized for its accuracy, robustness and speed with rotating machinery applications. Figure 6 shows the measured performance of various CFX standard V16 benchmarks.



Here the performance is measured based on the elapsed solver time with the reference of 1.0 for the Dell EMC 13G Ready Bundle IBB (Dell EMC R630 with dual Intel Xeon E5-2667v4 8-core processors). For comparison, the figure also includes the performance data for the Dell EMC 13G Ready Bundle EBB (Dell EMC C6320 with dual Intel Xeon E5-2697Av4 16-core processors). Higher Results indicate better overall performance. There was not enough memory to carry out the airfoil_100M benchmark on the 13G EBB. These results show a steady performance advantage of the new 14G based EBBs.

Figure 7 presents the 14G EBB parallel performance data when run on multiple nodes (up to 8). The figure uses the performance reference of 1.0 for a single node (32 cores total).




The overall parallel scalability for these models is good, with the parallel scalability for the largest two models actually super-linear, which is related to the “cache effect”, where more of the problem dataset is able to be kept in cache as the number of nodes used increases.

6 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. Given the varying system requirements for different types of FEA problems, benchmarking for ANSYS Mechanical was performed using the Implicit and Explicit building block systems. The ten benchmark problems from the ANSYS Mechanical v18.0 benchmark suite were evaluated on the test system.

The performance results for individual building block types are presented in Figure 8 below both for 13G and 14G systems. 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 of the cores on a processor are used. As such, only results from the DMP solver are presented, using all available cores on each server. The results are presented using the Core Solver Rating metric. This metric represents the performance of the solver core which excludes any job pre and post-processing.



These results are shown relative to the performance of the 13G IBB, where higher indicates better overall performance. The performance of the 14G building blocks are noticeably better than their 13G counterparts. Typically the EBB outperform the IBB, but this can be problem dependent.

Figure 9 and Figure 10 present the performance data for the 14G IBB and EBB servers, where each data point on the graphs records the performance of the specific benchmark data set using the number of cores marked on the x-axis relative to the single core result.

Figure 9: ANSYS Mechanical Scaling on a single IBB






Figure 10: ANSYS Mechanical Scaling on a single EBB



These figures make it easy to see the scaling of the solution, i.e. the performance improvement as more cores are used for the analysis. Problem scalability depends on many factors including the number of degrees of freedom in the problem and on the particular solution type and solver being used. The overall behavior for both servers are similar. The LN-2 and CG-3 datasets scale well up to all available cores for both servers, while the other benchmark cases typically only scale up to about 16 cores per server. Limits in parallel scalability for utilizing cores on a single server can arise from both inherent limits in the parallel scalability of the solver algorithms and contention for data from memory can limit the ability to effectively utilize all of the cores for some datasets.

The performance results for the ANSYS Mechanical solver on multiple EBBs are shown in Figure 11. Each data point on the graphs 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 (32-core) result.



The DMP solver continues to scale up to 256 processor cores for some of the benchmarks. Like the single server results, problem scalability depends on many factors including the number of degrees of freedom in the problem and on the particular solution type and solver being used. For these benchmark cases, parallel scalability above four nodes is limited. However, it is common to see good parallel scalability up to eight nodes for larger data sets.

Conclusion

This technical white paper presents a validated architecture for the Dell EMC Ready Bundle 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. The design takes into account 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 system design, providing actual measured results on the system using ANSYS software.

References

- [ANSYS engineering simulation software](#).
- [Bright Computing](#), Bright Cluster Manager.
- [Dell EMC Ready Solutions for High Performance Computing](#).
- [Dell EMC Intel Enterprise Edition for Lustre Storage Solution](#).
- [Dell EMC NFS Storage Solution](#).
- [Dell HPC System for Manufacturing—System Architecture and Application Performance](#), Technical White Paper, July 2016.