

HPC Digital Manufacturing with Siemens Simcenter STAR-CCM+

Dell EMC PowerEdge Servers with 3rd Gen Intel® Xeon® Scalable
Processors and options for storage

May 2021

White Paper

Abstract

This document describes the design and reviews performance benchmarking results and analysis for Siemens' Simcenter STAR-CCM+™ on Dell EMC PowerEdge servers with 3rd-generation Intel® Xeon® Scalable processors, PowerSwitch and InfiniBand networking and options for storage.

Dell Technologies Validated Design

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 05/21 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

 Executive summary4

System building blocks4

 Overview4

 Infrastructure Servers5

 Compute Servers.....6

 Storage.....6

 System Networks.....8

 Cluster Management Software8

 Services and Support9

Reference System9

 Components9

 BIOS.....9

 Software10

Simcenter STAR-CCM+ Performance10

 Overview10

 Results10

Conclusion.....12

 We value your feedback12

Introduction

Executive summary

This technical paper discusses the performance of Siemens' Simcenter STAR-CCM+™ on the Validated Design for HPC Digital Manufacturing with 3rd generation Intel Xeon Scalable processors. This Validated Design for HPC was designed specifically for digital manufacturing workloads, where computer-aided engineering (CAE) applications are critical for virtual product development. The Validated Design 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 optimized for specific workloads and use cases.

The Validated Design for HPC Digital Manufacturing is one of many solutions in the Dell Technologies HPC solution portfolio. Please visit <https://delltechnologies.com/hpc> for a comprehensive overview of the available HPC solutions offered by Dell Technologies.

The architecture of the Validated Design 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 Technologies Validated Design for HPC Digital Manufacturing. Section 4 presents benchmark performance for Simcenter STAR-CCM+.

System building blocks

Overview

The Validated Design for HPC Digital Manufacturing is designed using a flexible building block architecture. This 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 and networking. Configuration recommendations are provided for each of the building blocks that 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 specific workloads and use-case requirements. Figure 1 shows three example HPC clusters designed using the Validated Design for HPC Digital Manufacturing architecture.

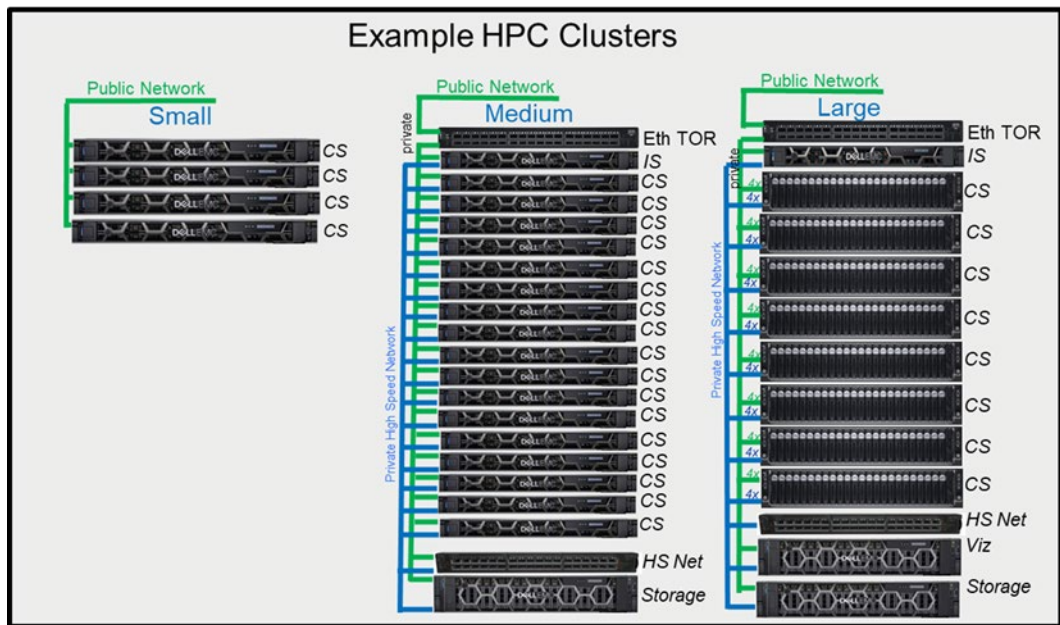


Figure 1. Example Validated Design for HPC Digital Manufacturing

Infrastructure Servers

Infrastructure servers are used to administer the system and to 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 head 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 head node is mandatory for an HPC system to deploy and manage the system. If high-availability (HA) management functionality is required, two head 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 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 higher bin 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 Validated Designs for HPC Storage solutions.

Compute Servers

Compute servers provide the computational resources for the HPC system. These servers are used to run the engineering analysis workloads such as Simcenter STAR-CCM+. 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 chapter 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 Technologies offers a wide range of general purpose and HPC storage solutions. For a general overview of the Dell Technologies HPC solution portfolio, please visit <https://delltechnologies.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 that 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 computer-aided engineering (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 the HPC environment.

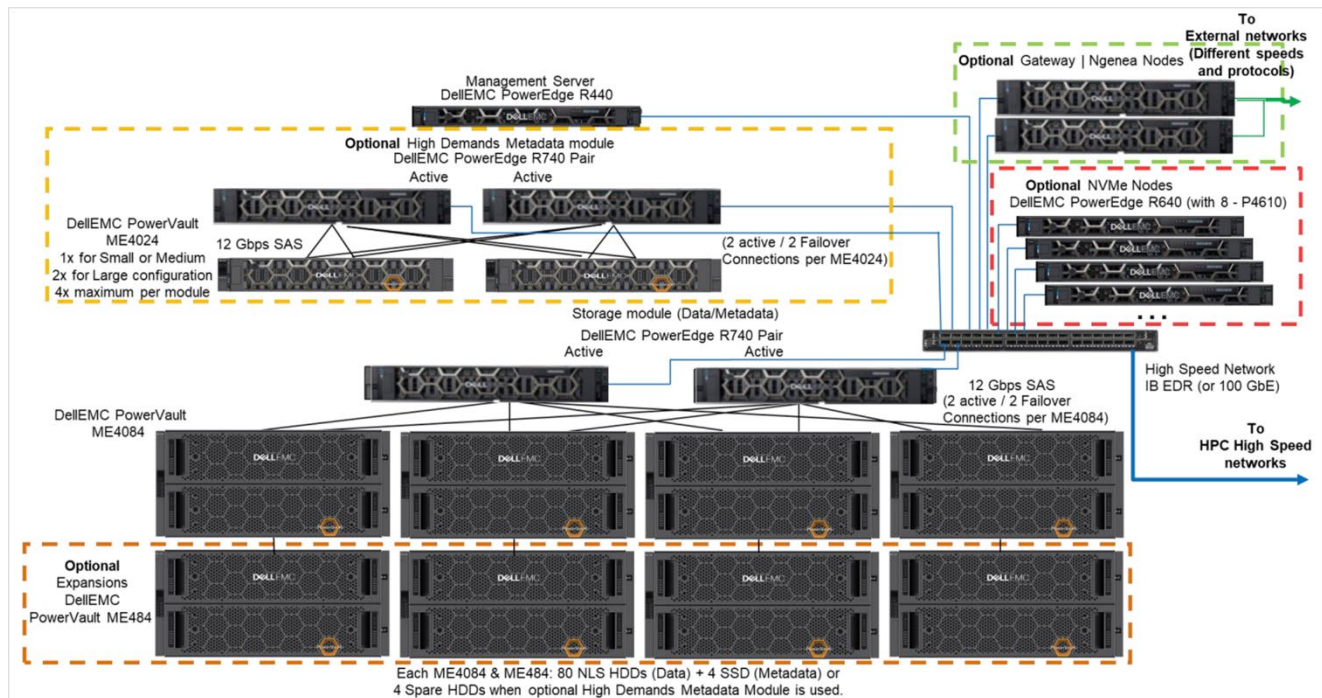
Operational storage is typically sized based on the number of expected users. For fewer than 30 users, a single NFS storage server, such as the Dell EMC PowerEdge R740xd is often an appropriate choice. A suitably equipped storage server may be:

- 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 PowerScale A200 scale-out NAS may be appropriate.

For customers desiring a shared high-performance parallel filesystem, the Validated Design for HPC PixStor Storage solution shown in Figure 2 is appropriate. This solution can scale up to multiple petabytes of 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 out-of-band (OOB) server management. The management network typically uses the Dell EMC PowerSwitch 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 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

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

Services and Support

The Validated Design for HPC Digital Manufacturing is available with full hardware support and deployment services.

Reference System

Components

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

Table 2. Benchmark system configuration

Building Block	Quantity
Computational Server <ul style="list-style-type: none"> PowerEdge C6520 Dual Intel Xeon Gold 6346 512GB RAM 16x32GB 3200 MTps DIMMs NVIDIA ConnectX-6 HDR100 adapter 	1
Computational Server <ul style="list-style-type: none"> PowerEdge C6520 Dual Intel Xeon Gold 6342 512GB RAM 16x32GB 3200 MTps DIMMs NVIDIA ConnectX-6 HDR100 adapter 	1
Computational Server <ul style="list-style-type: none"> PowerEdge C6520 Dual Intel Xeon Platinum 6338 512GB RAM 16x32GB 3200 MTps DIMMs NVIDIA ConnectX-6 HDR100 adapter 	1
Computational Server <ul style="list-style-type: none"> PowerEdge C6520 Dual Intel Xeon Platinum 8358 512GB RAM 16x32GB 3200 MTps DIMMs NVIDIA ConnectX-6 HDR100 adapter 	6
NVIDIA QM8790 InfiniBand switch	1

BIOS

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

BIOS Option	Setting
Sub NUMA Cluster	2-Way

Software

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
Simcenter STAR-CCM+	2021.1.1 mixed precision

Simcenter STAR-CCM+ Performance

Overview

Simcenter STAR-CCM+ is a multiphysics application used to simulate a wide range of products and designs under a wide range of conditions. The benchmarks reported here mainly use the computational fluid dynamics (CFD) and heat transfer features of Simcenter STAR-CCM+. 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 may have greater I/O demands, such as transient analysis.

Results

The benchmark problems from the standard Simcenter STAR-CCM+ benchmark suite were evaluated on the reference system. Simcenter STAR-CCM+ benchmark performance is measured using the Average Elapsed Time metric which is the average elapsed time per solver iteration. A smaller elapsed time represents better performance. Figure 3 shows the relative performance for a selection of Simcenter STAR-CCM+ benchmarks on a single server.




Figure 3. Single-server relative performance—Simcenter STAR-CCM+ 2021.1.1

The results in Figure 3 are plotted relative to the performance of a single compute server configured with dual 16-core Intel Xeon Gold 6346 processors. Larger values indicate better overall performance. These results show the performance advantage available with 3rd generation Intel Xeon Scalable processors. The 24-core Intel Xeon Gold 6342 provides very good performance for these benchmarks. The 32-core Intel Xeon Gold 6338 and Platinum 8358 processors provide on average 7% and 14% better performance respectively than the Gold 6342.




Figure 4. Simcenter STAR-CCM+ parallel scaling—Intel Xeon Platinum 8358

Figure 4 presents the parallel scalability of the Simcenter STAR-CCM+ benchmark models using up to six computational servers configured with Intel Xeon Platinum 8358 processors. The performance is presented relative to the performance of a single node (64 cores total).

The parallel scalability for most of these benchmark models is good, with the system demonstrating nearly linear parallel scaling. Some of the smaller models do not scale as well as the larger benchmark cases, but this is to be expected as communication overhead limits the parallel scalability of smaller cases.

Conclusion

This document presents the Validated Design for HPC STAR CCM+ with 3rd generation Intel Xeon Scalable processors. The detailed analysis of the compute server configurations demonstrate that the system is architected for a specific purpose—to provide a comprehensive HPC solution for the digital manufacturing or computer-aided engineering domain. Use of this 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 services and support readily available. The performance benchmarking bears out the solution design, demonstrating the performance of the solution with Siemens' Simcenter STAR-CCM+.

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](#).

Authors: Joshua Weage, Martin Feyereisen