

SAP HANA TDI Deployments on Dell EMC PowerEdge Servers

Configuration and Deployment Guidelines

March 2020

H17840.4

Validation Guide

Abstract

This validation guide provides best-practice configuration and deployment guidelines for SAP HANA Tailored Datacenter Integration (TDI) deployments on SAP HANA certified Skylake- and Cascade Lake-based Dell EMC PowerEdge systems.

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 © 2020 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 March 2020 Validation Guide H17840.4.

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**
- Technology overview6**
- Hardware requirements9**
- Sizing considerations11**
- Recommended PowerEdge settings and configuration27**
- Conclusion.....30**
- References.....31**
- Appendix A: Designing and implementing an SAP HANA scale-out landscape using the PowerEdge MX platform32**
- Appendix B: SAPS values for SAP HANA certified PowerEdge servers40**

Introduction

Business case

In today's competitive, data-intensive world, businesses require instantaneous access to information insights. Dell EMC's SAP HANA solutions dramatically increase the speed of access to business information.

SAP HANA is a multipurpose, data-source-agnostic, in-memory database. Dell EMC's SAP HANA solutions are optimized to run SAP HANA software on Dell EMC servers, networking, and storage.

When SAP first introduced SAP HANA on dedicated appliances, IT architects could not apply many of their traditional practices and principles. The SAP HANA Tailored Data Center Integration (TDI) model changed that. Although SAP still specifies restrictions on SAP HANA infrastructure architectures, most of the key principles of IT system architecture can now be applied to SAP HANA systems. A holistic approach to IT architecture across the data center is achieved by using existing data center resources, providing a strong return-on-investment (ROI) case for SAP HANA adoption. The flexibility and cost-effectiveness of SAP HANA TDI make it the future-proof option for organizations planning to make SAP HANA mainstream in their data centers.

Solution overview

With SAP HANA TDI, you can deploy any Dell EMC PowerEdge system that is certified by SAP, as listed under **Certified Appliances** in the [Certified and Supported SAP HANA Hardware Directory](#), with any Dell EMC enterprise storage system that is listed under **Certified Enterprise Storage** in the same directory.

Effective from [SAP HANA TDI Phase 5](#), SAP supports the complete family of Intel Xeon Skylake and Cascade Lake platinum, gold, and silver processors, featuring eight+ cores per processor and a minimum of two sockets. The increased variety of CPU processor options enables customers to be more flexible in their hardware choices and to cost-optimize their Dell EMC PowerEdge systems for their specific workload requirements. SAP HANA TDI Phase Five allows for a workload-driven sizing approach by employing the SAP Application Performance Standard (SAPS) to determine the maximum RAM capacity for a specific customer application. Accordingly, customers whose workloads are less CPU-intensive might have a larger main memory capacity compared to SAP HANA appliance-based solutions that use a fixed core-to-memory sizing approach.

This Dell EMC PowerEdge Server solution for SAP HANA TDI deployments incorporates the latest Intel Xeon CPU architectures, while PowerEdge servers provide an exceptional combination of performance, resiliency, and flexibility to meet the needs of the modern enterprise data center.

Key benefits

With SAP HANA TDI, all scenarios that are applicable to Dell EMC's SAP HANA certified appliances are supported in TDI deployments using the same Dell EMC server model. PowerEdge servers are designed for business continuity and offer outstanding enterprise-class features, including:

- **Powerful technology**—Single-server configurations ranging in size from 192 GB to 6 TB provide a consistent experience and a solid base for future expansion without forcing “rip and replace” as system requirements evolve. The latest PowerEdge

servers with Intel Skylake or Cascade Lake processors are certified for SAP HANA single-server and scale-out solutions and include everything that is required to support your SAP HANA implementation. For more information, see the [Certified and Supported SAP HANA Hardware Directory](#).

- **Virtualization**—Dell EMC PowerEdge solutions are certified to run a virtualized SAP HANA implementation, which provides a more economical application development environment and is ready for deployment in production environments.
- **Enterprise-class availability**—With highly redundant hardware configurations and the use of both SAN/NAS technology and a highly available multinode design, Dell EMC scale-out solutions for SAP HANA provide resiliency and easy expandability.
- **High performance**—Automated data tiering with Dell EMC storage software provides quick access to the datasets that are most needed for analysis.
- **Backup and disaster recovery**—Dell EMC solutions for SAP HANA can provide remote disaster recovery through SAP HANA system replication or storage-based replication. Tests are performed with the [Dell EMC Data Domain appliance](#) to provide a high-performance, cost-effective solution.
- **Modular growth**—The scale-out solution can grow from 1.5 TB to 48 TB in modular increments without disrupting the existing system, preserving your investment as your system changes.
- **SAP applications in memory**—Dell EMC solutions for SAP HANA support application deployment that enables production business applications to benefit from the speed and performance of in-memory computing without the need for multiple compute environments.
- **A single point of contact**—Dell EMC is your single source for all Dell EMC hardware components of the SAP HANA solution.

Document purpose

This validation guide provides guidelines for designing and configuring an SAP HANA TDI solution for servers and the associated components that run best with SAP HANA. This guide is updated periodically with the latest information.

Audience

This guide is for SAP Basis, system, and storage administrators and architects who design SAP HANA systems for deployment on PowerEdge servers. You should have some knowledge of Dell EMC storage arrays, PowerEdge servers, and VMware vSphere Hypervisor.

We value your feedback

Dell Technologies and the authors of this guide 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: Benjamin Simroth, Dat Nguyen, Donagh Keeshan

Contributor: Aighne Kearney

Note: For links to additional documentation for this solution, see the [Dell Technologies Solutions for SAP Info Hub](#).

Technology overview

Introduction

SAP HANA is an in-memory data platform that can be deployed on-premises or in the cloud. Organizations use the SAP HANA platform to analyze large volumes of data and develop and deploy applications in real time.

SAP HANA combines SAP software components that are optimized on proven partner-provided hardware. Two models are available for on-premises deployment, as shown in the following figure:




Figure 1. SAP HANA appliance model and TDI model comparison

Appliance model

An SAP HANA appliance includes integrated storage, compute, and network components by default. The appliance is certified by SAP, built by one of the SAP HANA hardware partners, and shipped to customers with all its software components preinstalled, including the operating systems and SAP HANA software. Dell EMC provides preinstalled SAP HANA appliance solutions for a quicker time-to-market and easy integration into an existing or new SAP landscape.

The SAP HANA appliance model presents the following limitations to customers:

- Limited choice of servers, networks, and storage
- Inability to use existing data center infrastructure and operational processes
- Little knowledge and control of the critical components in the appliance
- Fixed sizes for SAP HANA server and storage capacities, leading to higher costs from lack of capacity and inability to respond rapidly to unexpected growth demands

TDI model

With the TDI model, customers can choose from a broad portfolio of SAP HANA certified servers that can be combined with certified network and storage components. The storage and network components can be shared by different workloads to optimize total cost of ownership (TCO). Customers can seamlessly integrate SAP HANA into existing data center operations such as disaster recovery, data protection, monitoring, and management, reducing the cost, time-to-value, and risk of an overall SAP HANA adoption. For more information, see the following documents:

- [SAP HANA Tailored Data Center Integration - Overview](#)
- [SAP HANA Tailored Data Center Integration - Frequently Asked Questions](#)

Dell EMC solutions for SAP HANA

Dell EMC offers a complete portfolio of SAP HANA certified solutions, including appliances and TDI solutions that are available in both bare-metal and virtual configurations. We also offer converged infrastructure (CI) and hyperconverged infrastructure (HCI) solutions for SAP HANA.

The following figure shows the Dell EMC family of SAP HANA solutions:

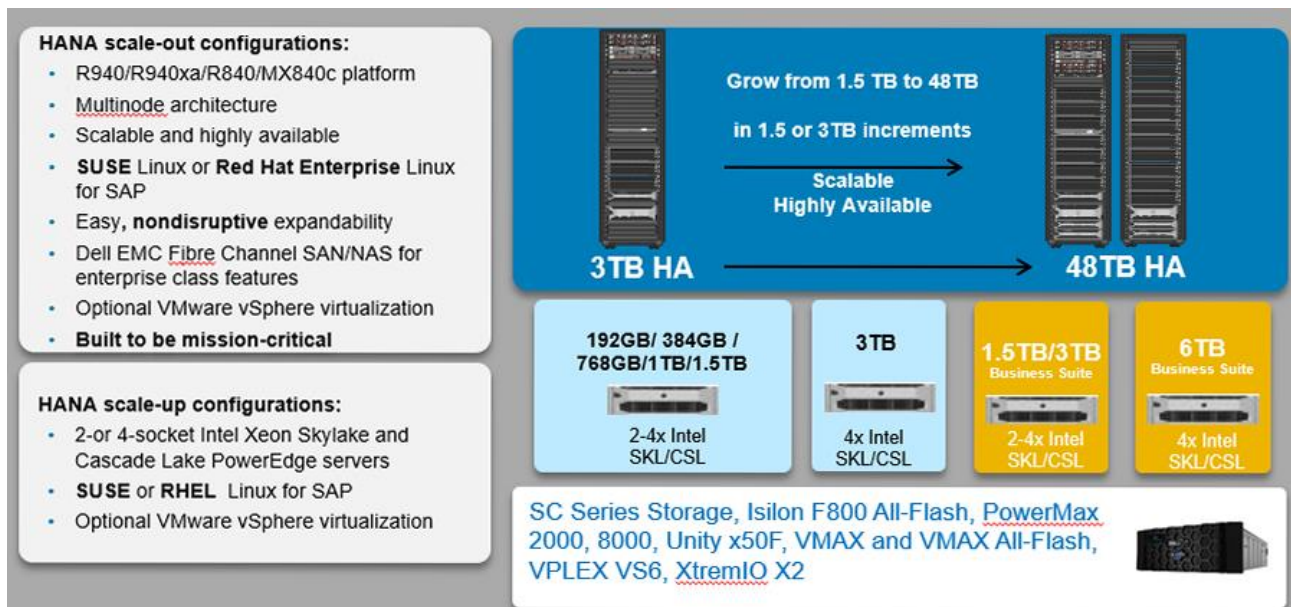


Figure 2. Dell EMC solutions for SAP HANA

Dell EMC appliances for SAP HANA

The following table shows the Dell EMC PowerEdge models that are supported for SAP HANA single-node (scale-up) appliance solutions:

Table 1. Supported PowerEdge models for preinstalled SAP HANA solutions

Intel Xeon CPU	PowerEdge server models
SP Skylake	R940, R840, R740xd
SP Cascade Lake	R940, R840

For more information about supported SAP HANA appliance types and memory sizes, see the [Certified and Supported SAP HANA Hardware Directory](#).

Dell EMC TDI solutions for SAP HANA

The following table shows the components of a Dell EMC solution for an SAP HANA TDI deployment:

Table 2. Configuration of a Dell EMC solution for SAP HANA TDI

Component	Description
Server	Dell EMC PowerEdge servers listed in Table 3
CPUs	<ul style="list-style-type: none"> 2 to 4 Intel Xeon Skylake 2 to 4 Intel Xeon Cascade Lake
Memory	<ul style="list-style-type: none"> Up to 24 DDR4 RDIMMs/LRDIMMs for a 2-CPU socket server Up to 48 DDR4 RDIMMs/LRDIMMs for a 4-CPU socket server
Operating system	SUSE Linux Enterprise Server, Red Hat Enterprise Linux, VMware vSphere (for virtual SAP HANA configurations)
Infrastructure switches	<ul style="list-style-type: none"> Ethernet fabric: 10 GbE to 200 GbE Dell EMC S-Series and Z-Series data center network switches SAN fabric: Dell EMC Connectrix Fibre Channel (FC), Dell EMC unified S4148U, or MX9116n/MX7116n modular switches
Storage configuration	<ul style="list-style-type: none"> Local storage: All-flash or a combination of serial-attached storage (SAS) and solid-state device (SSD) drives. External storage: Dell EMC storage arrays
Backup and disaster recovery	<ul style="list-style-type: none"> SAP HANA system replication Backup and recovery of SAP HANA on the Dell EMC Data Domain appliance
Solution and hardware life cycle management	Dell EMC OpenManage

If you have questions about specific licenses and services, consult your Dell EMC sales representative.

Hardware requirements

PowerEdge server models for TDI deployments

You can choose server options from a broad portfolio based on your organization's needs. For example, the PowerEdge R940xa, R940, R840, and R740xd servers have a capacity of up to 32 internal drives and provide outstanding performance.

The following table shows the PowerEdge servers that are SAP HANA certified and supported for use with Intel Skylake and Cascade Lake CPUs in TDI deployments:

Table 3. Supported PowerEdge server models for TDI deployments

Intel Xeon CPU	Supported PowerEdge server models*
Intel Xeon SP Skylake	R640, T640, R740, R740xd, R840, R940, R940xa FC640, M640, FC640, MX740c, MX840c
Intel Xeon SP Cascade Lake	R640, T640, R740, R740xd, R740xd2, R840, R940, R940xa FC640, M640, FC640, MX740c, MX840c

* R: Rack servers, T: Tower servers, FC and M/MX: Modular blade servers

For more detailed information about specific servers, see the [Dell Technologies website](#).

PowerEdge MX for SAP HANA TDI solutions

The Dell EMC PowerEdge MX platform is the latest hardware implemented for Dell EMC solutions for SAP HANA TDI deployments. Dell EMC PowerEdge MX servers can be used in SAP HANA scale-out TDI and SAP HANA scale-up (single-node) TDI deployments, including virtual configurations.

PowerEdge MX is a modular blade-based platform that integrates servers, storage, networking, and management into a single chassis that is designed for maximum density, efficiency, and manageability. As a unified, high-performance, kinetic data center infrastructure, PowerEdge MX provides the agility, resiliency, and efficiency to optimize a wide variety of traditional and new emerging data center workloads and applications. Its responsive design provides the innovation and longevity that organizations of all sizes need for their IT and digital business transformations. For more information, see the [Dell EMC PowerEdge MX](#) page on the Dell EMC website.

For configuration recommendations and a demonstration of how the PowerEdge MX infrastructure can help consolidate hardware requirements and simplify the implementation of an SAP HANA scale-out solution, see [Appendix A: Designing and implementing an SAP HANA scale-out landscape using the PowerEdge MX platform](#).

Certified enterprise storage

The following table shows the SAP HANA certified enterprise storage components for Dell EMC storage arrays:

Table 4. SAP HANA certified enterprise storage components for Dell EMC storage arrays

Storage family	Storage models	Storage connector
PowerMax	PowerMax 2000, 8000	SAN-FC

Storage family	Storage models	Storage connector
SC Series	SCv3020, SC5020/F SC7020/F, SC9000	SAN-FC
Isilon F800	F800 All-Flash	NAS-Shared file system
Unity x50F	Unity 350F, 450F, 550F, 650F	SAN-FC
VMAX and VMAX All-Flash	<ul style="list-style-type: none"> VMAX 100K, 200K, 400K VMAX All Flash 250F/FX, 450F/X, 850F/FX, 950F/FX 	SAN-FC
VPLEX	VPLEX VS6	SAN-FC
XtremIO X2	<ul style="list-style-type: none"> XtremIO X2 (-S/-R/-T) XtremIO X2-T, 1 X-Brick cluster XtremIO X2-S, 1 to 4 X-Brick cluster XtremIO X2-R, 1 to 8 X-Brick cluster 	SAN-FC

A broad range of Dell EMC storage options are available to centralize overall management needs and support easy replacement using FC or network connections. Contact your Dell EMC sales representative about the best components for your needs.

The following guides describe configuration best practices for Dell EMC enterprise storage systems that are certified for SAP HANA TDI deployments. These guides provide information about how to configure the storage to meet the SAP HANA key performance indicators (KPIs) and prepare the servers for Linux native multipathing and accessing enterprise storage from the SAP HANA nodes.

- [Storage Configuration Best Practices for SAP HANA TDI on Dell EMC PowerMax Arrays Validation Guide](#)
- [Storage Configuration Best Practices for SAP HANA TDI on Dell EMC VPLEX Systems Validation Guide](#)
- [Dell EMC SC Series Storage Configuration Best Practices for SAP HANA TDI Validation Guide](#)
- [Storage Configuration Best Practices for SAP HANA Tailored Data Center Integration on Dell EMC Compellent Storage Systems Solution Guide](#)
- [Storage Configuration Best Practices for SAP HANA TDI on Dell EMC Unity Storage Systems Solution Guide](#)
- [Storage Configuration Best Practices for SAP HANA Tailored Data Center Integration on Dell EMC VMAX Solution Guide](#)
- [Dell EMC XtremIO Storage Configuration Best Practices for SAP HANA TDI Configuration Guide](#)
- [SAP HANA TDI on Dell EMC Isilon All-Flash Scale-Out NAS Validation Guide](#)

Sizing considerations

Introduction

Before you deploy SAP HANA on a PowerEdge system, determine the number of production and nonproduction SAP HANA systems that you require and the CPU, memory, and disk capacity requirements for these systems.

SAP HANA system sizing

With [SAP HANA TDI Phase 5](#), SAP introduced workload-driven SAP HANA system sizing, in which SAPS requirements for specific customer workloads are used to determine the type and number of processors that are required to run SAP HANA. Customers use the SAP HANA Quick Sizer tool and sizing reports to determine the optimal memory sizes, number and type of CPUs, and number and type of disks for their SAP HANA environment, and then communicate the results to Dell EMC.

For more information, see [Quick Sizer](#). If you need help with sizing your SAP system, consult your Dell EMC sales representative.

With SAP HANA PowerEdge solutions, you can use internal as well as external storage for SAP HANA. Your sales representative can help you identify the best solution for the specific workload and environment.

If you are planning to install the operating system locally, consider using an [M.2 SSD](#) in combination with a BOSS controller to decouple the operating system from your workload. BOSS devices provide RAID capabilities and offer a good performance-to-price ratio. You can use Dell EMC OpenManage tools to keep the firmware on these devices up to date as well as to monitor the system and keep a hardware inventory.

CPU considerations

For information about choosing the right CPU model for your workload, see [Appendix B: SAPS values for SAP HANA certified PowerEdge servers](#) and consult your Dell EMC representative.

Memory requirements

Dell EMC recommends that SAP HANA solutions have a homogenous, symmetrical assembly of DIMMs and maximum utilization of all memory channels per processor. In accordance with this recommendation, two-socket servers must be populated with at least 12 DIMMs, and four-socket solutions must have at least 24 DIMMs.

The following figures show the optimal DIMM size and count for all possible memory permutations:

DIMM GB / RAM GB	192	384	576	768	1152	1536	2304	3072
8	24							
16	12	24						
32		12	12+12	24				
64				12	12+12		24	
128						12	12+12	24

Figure 3. Supported DIMM/memory configurations for two socket servers

DIMM GB / RAM GB	192	384	768	1152	1536	2304	3072	4608	6144
8	24	48							
16		24	48	24+24					
32			24		48	24+24			
64					24		48	24+24	
128							24		48

	Recommended
	Neutral
	Not allowed

Figure 4. Supported DIMM/memory configurations for four socket servers

Note: When selecting a DIMM size, consider future memory upgrade requirements to avoid having to replace existing DIMMs at that time.

Intel Optane memory

Dell EMC PowerEdge servers with the latest Intel Xeon CPUs in combination with Intel Optane memory can lower the TCO for SAP HANA environments and increase the overall memory capacity within a machine.

Intel Optane memory on SAP HANA enables up to 15 TB of data to be stored within one four-socket server beyond the physical capabilities of DRAM. SAP HANA supports persistent memory (PMEM) in App Direct mode, but not in Memory mode.

Memory limits are as follows:

- **Two-socket systems**—12 x 512 GB = 6 TB PMEM + 12 x 128 GB = 1.5 TB DRAM, yielding 7.5 TB.
- **Four-socket systems**—24 x 512 GB = 12 TB PMEM + 24 x 128 GB = 3 TB DRAM, yielding 15 TB.

Note: PMEM works only with SAP HANA certified Cascade Lake servers.

Currently, SAP HANA supports three operating system versions with PMEM:

- **SLES 12 SP4**—Certified for PowerEdge R940, R840, R740, and R740xd servers by Dell for Cascade Lake.
- **SLES 15**—Certified for PowerEdge R940, R940xa, R840, R740, R740xd, R640, MX740c, and MX840c servers.
- **Red Hat Enterprise Linux 7.6 and 8**—Certified for PowerEdge R940, R940xa, R840, R740, R740xd, R640, MX740c, and MX840c servers. For support details for specific versions and solutions, see SAP Note: [KB3830541](#).

For more information about Intel Optane memory, see SAP Note 2700084: [FAQ: SAP HANA Persistent Memory](#) (access requires login credentials).

CPU and memory requirements

All Cascade Lake CPUs with Xeon Gold or higher are compatible with Intel Optane memory. The specific processor needs depend on the total memory size (Intel Optane + DRAM). The CPU supports a specific memory size.

The following tables describe the minimum CPU requirements. Depending on performance needs, you might prefer a higher version, for example, Intel Xeon Platinum.

Table 5. Minimum CPU memory requirements

Processor type	Memory limit
No suffix in processor type	1 TB per socket (4 TB maximum)
M suffix in processor type	2 TB per socket (8 TB maximum)
L suffix in processor type	4.5 TB per socket (18 TB maximum)

Intel Optane for SAP HANA can be used in two-socket or four-socket systems, as described in the following tables:

Table 6. Systems with PMEM and two sockets populated

DRAM (size)	PMEM (size)	Ratio (DRAM: PMEM)/(total system capacity)	Minimum CPU model
12 x 32 GB (384 GB)	12 x 128 GB (1,536 GB)	1:4 (1,920 GB)	Intel Xeon Gold CPU with no suffix
12 x 64 GB (768 GB)	12 x 128 GB (1,536 GB)	1:2 (2,304 GB)	Intel Xeon Gold CPU with M suffix
12 x 64 GB (768 GB)	12 x 256 GB (3,072 GB)	1:4 (3,840 GB)	Intel Xeon Gold CPU with M suffix
12 x 128 GB (1,536 GB)	12 x 128 GB (1,536 GB)	1:1 (3,072 GB)	Intel Xeon Gold CPU with M suffix
12 x 128 GB (1,536 GB)	12 x 256 GB (3,072 GB)	1:2 (4,608 GB)	Intel Xeon Gold CPU with L suffix
12 x 128 GB (1,536 GB)	12 x 512 GB (6,144 GB)	1:4 (7,680 GB)	Intel Xeon Gold CPU with L suffix

Table 7. Systems with PMEM and four sockets populated

DRAM (size)	PMEM (size)	Ratio (DRAM, PMEM)/(total system capacity)	Minimum CPU model
24 x 32 GB (768 GB)	24 x 128 GB (3,072 GB)	1:4 (3,840 GB)	Intel Xeon Gold CPU with no suffix
24 x 64 GB (1,536 GB)	24 x 128 GB (3,072 GB)	1:2 (4,608 GB)	Intel Xeon Gold CPU with M suffix
24 x 64 GB (1,536 GB)	24 x 256 GB (6,144 GB)	1:4 (7,680 GB)	Intel Xeon Gold CPU with M suffix
24 x 128 GB (3,072 GB)	24 x 128 GB (3,072 GB)	1:1 (6,144 GB)	Intel Xeon Gold CPU with M suffix
24 x 128 GB (3,072 GB)	24 x 256 GB (6,144 GB)	1:2 (9,216 GB)	Intel Xeon Gold CPU with L suffix
24 x 128 GB (3,072 GB)	24 x 512 GB (12,288 GB)	1:4 (15,360 GB)	Intel Xeon Gold CPU with L suffix

Sizing recommendations for storage with PMEM

You can use any internal storage or certified Dell EMC SAP HANA external storage. For more information, see the disk sizing requirements for scale-up systems in [SAP HANA TDI-Storage Requirements](#).

Scale-out systems are more complex, as described in the following table:

Table 8. Recommended storage sizes for scale-out systems

Volume	Size
/hana/data	1.2 x anticipated net data size on disk or 1 x total main memory (DRAM + PMEM)
/hana/log	512 GB
/hana/shared	1 TB

Preparing the Intel Optane memory

Configuring the Intel Optane memory DIMMs involves two steps:

1. Creating a goal configuration from the BIOS environment
2. Creating the namespaces with an installed operating system

Creating a goal configuration

To access the memory DIMMs, you can create a goal configuration using the `ndctl` tool at operating system level, the system setup dialog, or the Integrated Dell Remote Access Controller (iDRAC).

Note: You can perform this procedure without any operating system on the machine. The namespace creation procedure requires that an SAP HANA supported operating system that includes the `ndctl` tool is installed on the machine.

To create a goal configuration using the iDRAC:

1. Log in to the iDRAC of the PMEM machine and select **Configuration > BIOS Settings**.
2. Expand **Memory Settings** and **Persistent Memory**, as shown in the following figure:

The screenshot displays the iDRAC BIOS Settings interface for Persistent Memory configuration. It is organized into several sections:

- Persistent Memory:**
 - Current Value: Non-Volatile DIMM
 - Pending Value: (empty)
 - Sanitize All NVDIMMs: Disabled
- Intel Persistent Memory:**
 - Current Value: (empty)
 - Pending Value: (empty)
 - Raw capacity: 2.9 TiB
 - App Direct capacity: 2.9 TiB
 - Memory capacity: 0 B
 - Unconfigured capacity: 0 B
- Persistent Memory DIMM Configuration:** (Collapsible section)
- Region Configuration:**
 - Persistent Memory Region 1 Information
 - Persistent Memory Region 2 Information
 - Persistent Memory Region 3 Information
 - Persistent Memory Region 4 Information
- Create goal config:**
 - Current Value: (empty)
 - Pending Value: (empty)
 - Persistent [%]: 100
 - Memory Mode [%]: 0
 - Persistent memory type: App Direct Interleaved

At the bottom of the interface, there are two buttons: **Apply** (in blue) and **Discard** (in grey).

Figure 5. Persistent Memory settings

3. Expand **Region Configuration** and **Create goal config**.

4. To create a per-socket configuration, choose **100% Persistence** and **App Direct Interleaved**.
5. Click **Apply and reboot**.

Creating the namespaces

After you complete the goal configuration, you have prepared the regions but not yet allocated the namespaces. To list all unallocated namespaces, run the `ndctl` command using the `-i` flag:

```
linux:~ # ndctl list -i

[
  {
    "dev": "namespace24.0",
    "mode": "raw",
    "size": 0,
    "uuid": "00000000-0000-0000-0000-000000000000",
    "sector_size": 512,
    "state": "disabled",
    "numa_node": 0
  },
  {
    "dev": "namespace26.0",
    "mode": "raw",
    "size": 0,
    "uuid": "00000000-0000-0000-0000-000000000000",
    "sector_size": 512,
    "state": "disabled",
    "numa_node": 2
  },
  {
    "dev": "namespace25.0",
    "mode": "raw",
```



```

    "size":0,
    "uuid":"00000000-0000-0000-0000-000000000000",
    "sector_size":512,
    "state":"disabled",
    "numa_node":1
  },
  {
    "dev":"namespace27.0",
    "mode":"raw",
    "size":0,
    "uuid":"00000000-0000-0000-0000-000000000000",
    "sector_size":512,
    "state":"disabled",
    "numa_node":3
  }

```

Creating namespaces on the configured regions

Depending on the number of populated sockets that exist in the system, you might need to repeat the create namespace procedure. Perform the following procedure two or more times, depending on the number of populated sockets in the system.

Run the following commands:

```

linux:~ # ndctl create-namespace
{
  "dev":"namespace24.0",
  "mode":"fsdax",
  "size":"744.19 GiB (799.06 GB)",
  "uuid":"8bc41612-bebc-4ead-bbca-6cc6f0b93be0",
  "raw_uuid":"3c1b2d89-31ac-4686-a6d0-ac22260d7515",
  "sector_size":512,
  "blockdev":"pmem24",
  "numa_node":0

```

Sizing considerations

```
}  
  
linux:~ # ndctl create-namespace  
  
{  
  
  "dev":"namespace26.0",  
  
  "mode":"fsdax",  
  
  "size":"744.19 GiB (799.06 GB)",  
  
  "uuid":"2785910e-a01d-42fc-990c-0b8f9563e49e",  
  
  "raw_uuid":"6f116048-9f99-4d14-b3f8-b47a8473da7e",  
  
  "sector_size":512,  
  
  "blockdev":"pmem26",  
  
  "numa_node":2  
  
}  
  
linux:~ # ndctl create-namespace  
  
{  
  
  "dev":"namespace25.0",  
  
  "mode":"fsdax",  
  
  "size":"744.19 GiB (799.06 GB)",  
  
  "uuid":"a4809f98-69c4-4fb0-b3f9-4cbb9293716f",  
  
  "raw_uuid":"84272619-6a3c-42ea-8d31-95198ff67589",  
  
  "sector_size":512,  
  
  "blockdev":"pmem25",  
  
  "numa_node":1  
  
}  
  
linux:~ # ndctl create-namespace  
  
{  
  
  "dev":"namespace27.0",  
  
  "mode":"fsdax",  
  
  "size":"744.19 GiB (799.06 GB)",  
  
  "uuid":"ea4cf63a-9910-408a-bae6-6b503e734dd8",
```

```

"raw_uuid":"feae79ac-f0ab-4bf2-a7d7-4b6591dcbe0d",
"sector_size":512,
"blockdev":"pmem27",
"numa_node":3
}

```

Configuring FS-DAX and create and mount the partitions

You now have either two or four block devices configured in the system. The names of the block devices start with `pmem` and a number. These names are displayed in the listing in the last step in the “`blockdev`” section. In our example, the names are `/dev/pmem24`, `/dev/pmem25`, `/dev/pmem26`, and `/dev/pmem27`.

Create an `xfs` file system by running the following `mkfs.xfs` command:

```

linux:~ # mkfs.xfs /dev/pmem24

meta-data=/dev/pmem24          isize=512    agcount=4,
agsize=48770944 blks
        =                       sectsz=4096  attr=2,
projid32bit=1
        =                       crc=1        finobt=1, sparse=0,
rmapbt=0, reflink=0

data      =                       bsize=4096  blocks=195083776,
imaxpct=25
        =                       sunit=0      swidth=0 blks

naming    =version 2           bsize=4096  ascii-ci=0 ftype=1

log       =internal log       bsize=4096  blocks=95255,
version=2
        =                       sectsz=4096  sunit=1 blks, lazy-
count=1

realtime  =none                extsz=4096  blocks=0,
rtextents=0

linux:~ # mkfs.xfs /dev/pmem25

meta-data=/dev/pmem25          isize=512    agcount=4,
agsize=48770944 blks
        =                       sectsz=4096  attr=2,
projid32bit=1
        =                       crc=1        finobt=1, sparse=0,
rmapbt=0, reflink=0

```

Sizing considerations

```
data      =          bsize=4096  blocks=195083776,
imaxpct=25
          =          sunit=0     swidth=0 blks
naming    =version 2  bsize=4096  ascii-ci=0 ftype=1
log       =internal log bsize=4096  blocks=95255,
version=2
          =          sectsz=4096  sunit=1 blks, lazy-
count=1
realtime  =none      extsz=4096  blocks=0,
rtextents=0
linux:~ # mkfs.xfs /dev/pmem26
meta-data=/dev/pmem26  isize=512  agcount=4,
agsize=48770944 blks
          =          sectsz=4096  attr=2,
projid32bit=1
          =          crc=1       finobt=1, sparse=0,
rmapbt=0, reflink=0
data      =          bsize=4096  blocks=195083776,
imaxpct=25
          =          sunit=0     swidth=0 blks
naming    =version 2  bsize=4096  ascii-ci=0 ftype=1
log       =internal log bsize=4096  blocks=95255,
version=2
          =          sectsz=4096  sunit=1 blks, lazy-
count=1
realtime  =none      extsz=4096  blocks=0,
rtextents=0
linux:~ # mkfs.xfs /dev/pmem27
meta-data=/dev/pmem27  isize=512  agcount=4,
agsize=48770944 blks
          =          sectsz=4096  attr=2,
projid32bit=1
          =          crc=1       finobt=1, sparse=0,
rmapbt=0, reflink=0
```

```

data      =                               bsize=4096   blocks=195083776,
imaxpct=25

          =                               sunit=0      swidth=0 blks

naming    =version 2                       bsize=4096   ascii-ci=0 ftype=1

log       =internal log                     bsize=4096   blocks=95255,
version=2

          =                               sectsz=4096   sunit=1 blks, lazy-
count=1

realtime  =none                             extsz=4096   blocks=0,
rtextents=0

```

Creating /etc/fstab entries for PMEM devices

For each PMEM device, perform the following steps:

1. Create a folder on which to mount the device by running the following command:

```
linux:~ # mkdir -p /hana/pmem/0; mkdir /hana/pmem/1; mkdir
/hana/pmem/2; mkdir /hana/pmem/3
```

2. To persist the mounting after the system reboots, add a mount entry in /etc/fstab for each namespace by running the following commands.

Note: The `x-systemd.device-timeout` parameter influences how long the system waits for the device to be ready. This example uses 20 minutes of initializing time. Decrease this value if necessary depending on your landscape needs.

```
/dev/pmem24 /hana/p
```

```
mem/0 xfs noatime,dax,x-systemd.device-timeout=1200 1 2
```

```
/dev/pmem25 /hana/pmem/1 xfs noatime,dax,x-systemd.device-
timeout=1200 1 2
```

```
/dev/pmem26 /hana/pmem/2 xfs noatime,dax,x-systemd.device-
timeout=1200 1 2
```

```
/dev/pmem27 /hana/pmem/3 xfs noatime,dax,x-systemd.device-
timeout=1200 1 2
```

3. Mount all the file systems and check that they are properly mounted by running the following commands:

```
linux:~ # mount -a -t xfs
```

```
linux:~ # df -h|egrep "File|pmem"
```

```

Filesystem                Size  Used Avail Use% Mounted on
/dev/pmem24                744G  792M  744G   1% /hana/pmem/1
/dev/pmem25                744G  792M  744G   1% /hana/pmem/2

```

```
/dev/pmem26          744G  792M  744G   1% /hana/pmem/3
/dev/pmem27          744G  792M  744G   1% /hana/pmem/4
```

Deploying and configuring SAP HANA

Next, you must make SAP HANA aware of the new Intel Optane memory DIMMs:

- Upgrade existing SAP HANA installations to SAP HANA SPS03 or later. In the [persistence] section of the `global.ini` file, provide a line with a comma-separated list of all mounted `pmem` devices by running the following command:

```
[persistence]
```

```
basepath_persistent_memory_volumes=/hana/pmem/0;/hana/pmem/1;
/hana/pmem/2;/hana /pmem/3
```

- For new SAP HANA installations, you can extend the tool `hdblcm` with two more options besides the normal installation parameters:

```
--use_pmem --pmempath=/hana/pmem
```

HDBLCM determines and uses all `pmem` devices below that `/hana/pmem` subfolder.

SAP HANA uses the persistent memory devices and loads data to them. You can also move databases and tables individually to a specific region (DRAM or PMEM).

Summary

Dell EMC conducted a test with SAP Business Warehouse on SAP HANA (BWoH), which involved running complex queries through the whole SAP HANA stack (HW,OS,DB,APP). The test results showed a similar read performance between DRAM and Intel Optane memory, indicating that this technology facilitates access to greater amounts of data while balancing TCO.

Data management with SAP HANA

The exponential growth of data has increased the demand for scaled-up storage and processing power to keep up with the complexities of enterprise information management. Businesses thirst for a data management strategy that satisfies their requirements for SAP HANA performance and data accessibility with a limited IT budget while lowering their TCO.

Mission-critical (or “hot”) data is stored in memory on the SAP HANA database for real-time processing and analytics. Less frequently accessed (or “warm”) data is saved in a lower-cost storage tier while still managed as a unified part of the SAP HANA database. Retaining and managing older data in this cost-effective way is key in maintaining data growth while minimizing the expense of hardware growth.

SAP HANA Native Storage Extension

Native Storage Extension (NSE), a native SAP HANA feature, is a warm data solution. NSE provides an improved cost-to-performance ratio that is similar to the ratio provided by other SAP HANA warm data tiering solutions, such as SAP HANA Extensions Node and SAP HANA Dynamic Tiering.

The following figure shows the difference between standard SAP HANA in-memory storage and the storage that is offered with NSE:

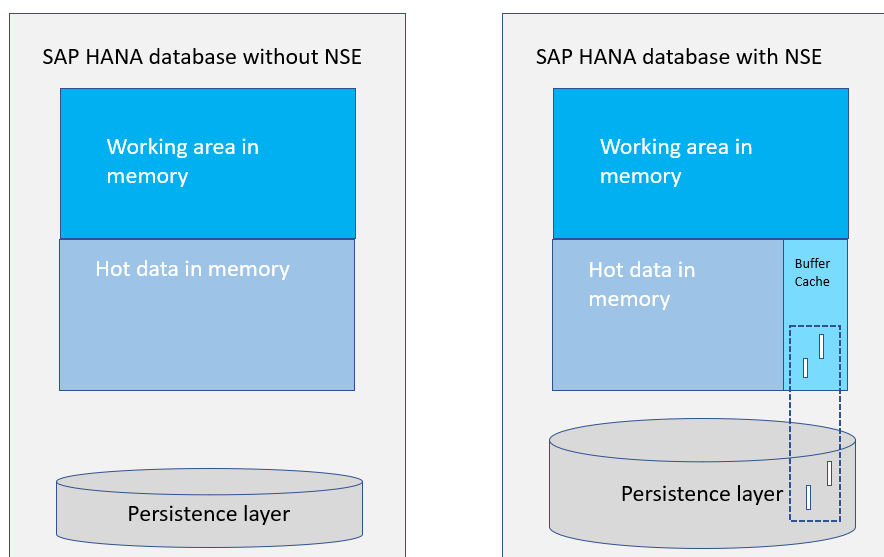


Figure 6. Standard SAP HANA database and SAP HANA database with NSE comparison

The capacity of a standard SAP HANA database is limited to the amount of main memory. Using SAP HANA NSE, customers can bypass these limits by storing warm data on a storage system. A relatively small amount of SAP HANA memory for the NSE buffer cache is needed for paging operations because the buffer cache can handle up to eight times the size of warm data on disk. For example, a 2 TB SAP HANA system without NSE equates to a 1 TB database in memory. With NSE and the addition of a 500 GB buffer cache, you can expand your 1 TB database to a 5 TB database: 1 TB of hot data, 4 TB of warm data, and a 500 GB buffer cache to page data between memory and disk.

Note: Only SAP HANA scale-up systems are currently supported with NSE.

Hot (or “column-loadable”) data resides completely in memory for fast processing and is loaded from disk in columns into SAP HANA memory. SAP HANA NSE enables you to specify certain warm data as “page loadable.” This data is loaded page by page into memory as required for query processing. Unlike column-loadable data, page-loadable data does not need to reside completely in memory.

NSE reduces the memory footprint for page-loadable data. The database is partly in-memory and partly on disk, as illustrated in the following figure:




Figure 7. SAP HANA database with NSE

NSE is integrated with other SAP HANA functional layers, such as the query optimizer, query execution engine, column store, and persistence layers,

For more information about SAP HANA NSE, SAP HANA NSE Data Sizing, and other NSE related topics, see [SAP HANA Native Storage Extension](#).

NSE is subject to certain functional restrictions. For more information, see SAP Note 2771956L: [SAP HANA NSE Functional Restrictions](#).

PowerMax storage family

PowerMax storage arrays can be used to substantially increase SAP HANA data capacity and reduce the TCO for customers.

The [PowerMax](#) family is the first Dell Technologies hardware platform to use an end-to-end Non-Volatile Memory Express (NVMe) architecture for customer data. NVMe is a set of standards which define a PCI Express (PCIe) interface that is used to efficiently access data storage volumes based on Non-Volatile Memory (NVM) media, including NAND-based flash and higher-performing Storage Class Memory (SCM) media technologies. The NVMe-based PowerMax was created to fully unlock the bandwidth, IOPS, and latency performance benefits that NVM media offers to host-based applications. These benefits are not attainable with the previous generation of all-flash storage arrays.

The PowerMax family consists of two models: the PowerMax 2000 and the flagship PowerMax 8000. The PowerMax 2000 model is designed to provide Dell Technologies customers with efficiency and maximum flexibility in a 20U footprint. The PowerMax 8000 model is designed for massive scale, performance, and IOPS density within a one or two-floor-tile footprint.

When using SCM drives, you can configure a PowerMax system to have its SCM drives intermixed with traditional NAND flash drives in the DAEs or as a 100 percent SCM system.

For more information, see [Validation Guide – PowerMax Storage Configuration Best Practices for SAP HANA TDI](#).

Network requirements

To support SAP HANA bandwidth requirements, SAP requires high-speed network connectivity for SAP HANA related networks. For more information, see [SAP HANA Network Requirements](#).

The internode network is tested to ensure that every node can connect to every other node with the minimal throughput requirement provided by SAP.

In addition to the system networks that are required for management and [vSphere for virtual HANA configurations](#), Dell EMC recommends configuring the following networks for the SAP HANA environments that will use the specified function:

- Application/client network: 10 Gb/s (recommended)
- Source/servers network: 10 Gb/s (recommended)
- Internode network: 10 Gb/s (required)
- System replication network: 10 Gb/s (recommended)
- Backup network: 10 Gb/s (recommended)
- Storage network:
 - 10 Gb/s IP (required)
 - 8 Gb/s Fibre Channel (required)
- Ethernet networks: 10 GbE to 200 GbE Dell EMC S Series and Z Series data center network switches

For SAN fabric, Dell EMC's storage networking recommendations include Dell EMC Connectrix FC or Dell EMC unified S4148U switches. The unified S4148U switch supports both Ethernet and FC networks.

MX Series modular switches are part of the PowerEdge MX server. PowerEdge MX Ethernet and FC I/O modules provide cost-effective, high-performance scalable networking—25 GbE and 32 GbE FC host connectivity with 100 GbE and 32 GbE FC uplinks. For more information about the supported hardware and the latest Open Networking OS10EE, see [MX-Series Modular Switches](#).

For information about Dell EMC's network deployment recommendations, see the following documentation:

- [Dell EMC Networking Guides](#)
- [SAP HANA Network Requirements White Paper](#)

Sizing considerations for an SAP landscape design

Business and technology considerations influence the sizing of the hardware infrastructure. Consult with your Dell EMC representative to determine a solution design that includes business requirements for performance, response times, availability, data protection, and disaster recovery for the SAP systems.

The following table describes the sequence of steps that are involved in sizing the infrastructure requirements for your TDI deployments:

Table 9. Infrastructure sizing steps

Step	Activity	Performed by
1	Determine the number of SAP production systems and SAP HANA systems that are required.	Dell EMC performs these steps when collecting your business requirements.
2	Define the SAP system landscapes. Typically, each SAP production system (enterprise resource planning or ERP, business warehouse or BW, customer relationship management or CRM, and so on) has its own SAP system landscape, consisting of a development (DEV), quality assurance (QAS), sandbox (SBX), and production (PRD) environment.	
3	Determine if virtualization is to be used.	
4	Determine the high availability requirements.	
5	Determine the DR requirements and the number of data centers or sites involved.	
6	Consider data protection requirements for backing up the SAP HANA systems.	
7	Determine an expected annual data growth and the required number of years of maintenance.	
8	Size each of the SAP HANA systems and system landscapes using the SAP Quick Sizer tool on your production systems.	Customer performs this step and then provides the results to Dell EMC.
9	Calculate the total compute requirements and determine the number and models of PowerEdge servers (SAPS, CPU, memory) that are required to support the SAP HANA systems in the TDI deployments	Dell EMC performs these steps using the information that the customer provides in step 8.
10	Calculate the total storage requirements based on capacity or IOPS, or both, and determine which Dell EMC enterprise storage system is best suited to meet the customer's business needs.	
11	Determine the backup capacity requirements.	
12	Set up the SAP HANA infrastructure and use HWCCT/HCOT to check that the SAP HANA TDI KPIs are met.	Customer or Dell EMC Professional Services performs this step.
13	Install the SAP HANA software.	Customer or Dell EMC Professional Services performs this step. Note: Only SAP HANA certified personnel can perform the installation.

Recommended PowerEdge settings and configuration

Introduction

This section describes Dell EMC's configuration recommendations for your SAP HANA scale-up, scale-out, and virtual configurations.

PowerEdge settings

For factory-installed servers, Dell EMC recommends the following settings for an SAP HANA solution:

- Memory Configuration Type: **Performance Optimized**
- BIOS and Advance System Configuration Settings: **Performance Optimized**
- Advanced System Configurations: **UEFI BIOS Boot Mode with GPT Partition**

BIOS settings

Dell EMC recommends the following performance settings in the PowerEdge System BIOS:

- Physical SAP HANA configuration:
 - Virtualization Technology: **Disabled**
- Virtual SAP HANA configuration:
 - Logical Processors: **Enabled**
 - Virtualization Technology: **Enabled**

Configuring BIOS settings on a PowerEdge server

Note: The following steps are performed on a PowerEdge R940 server for illustration purposes. The steps might vary on different server models.

You can change system BIOS settings in various different locations, as follows:

- Integrated Dell Remote Access Controller (iDRAC) UI: Set the **First Boot Device/Next boot** to **BIOS Boot Manager**.
- OpenManage: Enter the BIOS by pressing F2 during restart.
- System BIOS:
 - To access the processor settings, select **System BIOS > System BIOS Settings > Processor Settings**, as shown in the following figure::

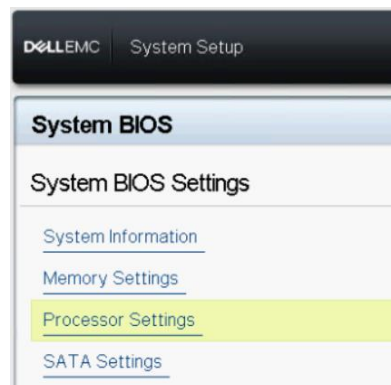


Figure 8. PowerEdge R940 System BIOS settings

The processor settings are displayed, as shown in the following figure:

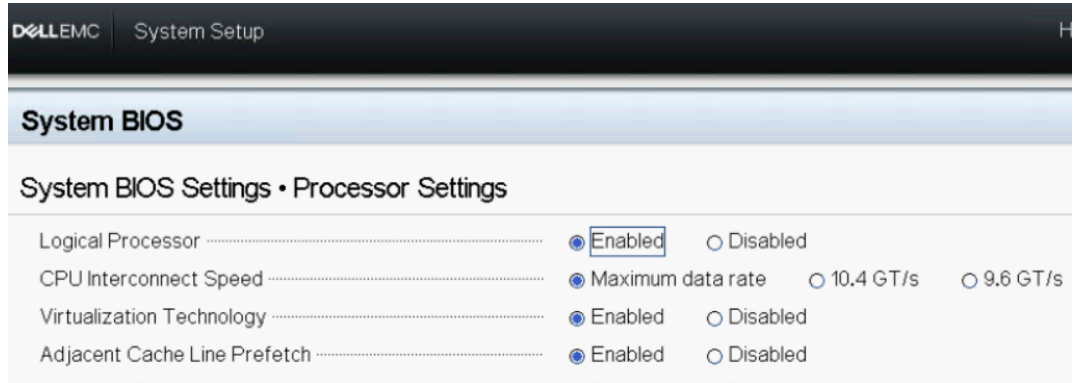


Figure 9. PowerEdge R940 Processor settings

- To access Memory Operating Mode, select **System BIOS > System BIOS Settings > Memory Settings**, as shown in the following figure:

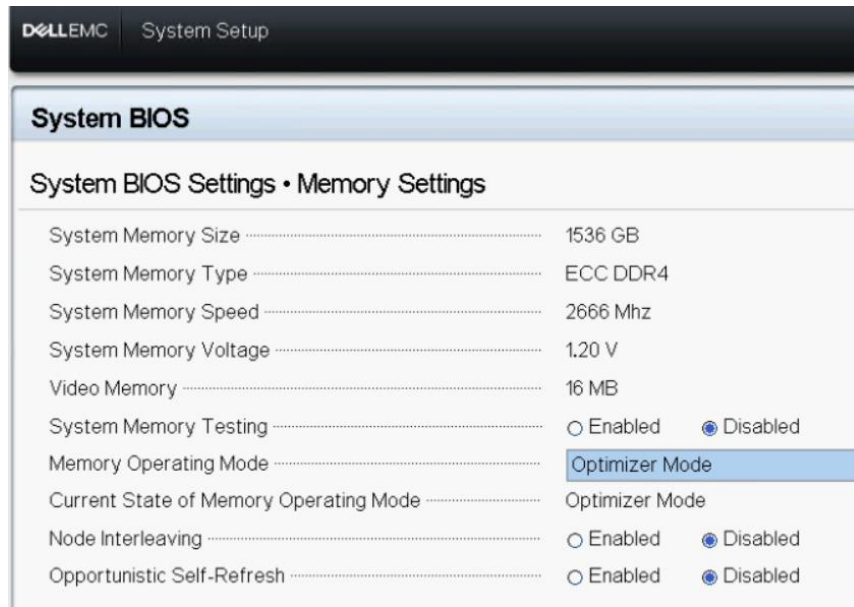


Figure 10. PowerEdge R940 Memory Operating Mode setting

Generic settings

Dell EMC recommends the following best practices for configuring the PowerEdge servers:

- Use the latest BIOS and firmware.
- Use the latest vendor-subscribed host bus adapter (HBA) drivers.
- Enable Intel Hyper-Threading.

For more information, see the [Architecture Guidelines and Best Practices for Deployments of SAP HANA on VMware vSphere: Architecture and Technical Considerations Guide](#).

For detailed deployment architecture and recommended settings for all your hardware and software components, contact your Dell EMC sales representative.

Scale-out considerations

Because Dell EMC's SAP HANA solutions are on scalable building blocks, you can easily add blocks of 1.5 TB or 3 TB to scale out to a 48 TB solution. SAP HANA certified enterprise storage components for Dell EMC storage arrays with dual redundant controllers provide the persistent storage for the SAP HANA database. All components, including storage, fabric, network, and controllers, are highly redundant.

Operating systems

The SAP HANA hardware certification program has certified hardware on specific operating system versions. For the supported combinations, see the [Certified and Supported SAP HANA Hardware Directory](#).

Supported operating systems for SAP HANA

You will need one of the following Enterprise Linux distribution products, in the version specified in the [index page](#) of the [Certified and Supported SAP HANA Hardware Directory](#), to run SAP HANA:

- Red Hat Enterprise Linux for SAP Solutions
- Red Hat Enterprise Linux for SAP HANA
- SUSE Linux Enterprise Server for SAP Applications
- SUSE Linux Enterprise Server

Note: Dell EMC strongly recommends using Red Hat Enterprise Linux for SAP Solutions or SUSE Linux Enterprise Server for SAP Applications because of their features and extended support cycles.

For information about SAP HANA support for the intended operating system version, see SAP Note [2235581 - SAP HANA: Supported Operating Systems](#) (access requires login credentials). This SAP note also provides the recommended settings for the operating system version. Ensure that you always apply these settings to the installed operating systems.

Operating system support life cycle

The following sections provide the end-of-support dates for the latest operating system versions associated with SAP HANA TDI: SUSE, Red Hat, and VMware ESXi.

SUSE Linux Enterprise Server

The following table provides end-of-support dates for SUSE operating system versions.

The information in the table is subject to change. For the most current information, see [SUSE Product Support Lifecycle](#).

Table 10. SUSE operating system versions: End of support

Operating system	End of support
SLES 12 for SAP Applications	July 1, 2019
SLES 12 SP1 for SAP Applications	May 31, 2020
SLES 12 SP2 for SAP Applications	March 31, 2021
SLES 12 SP3 for SAP Applications	June 30, 2022

Operating system	End of support
SLES 12 SP4 for SAP Applications	Not announced so far by SUSE
SLES 15 for SAP Applications	January 1, 2023

Red Hat Enterprise Linux

The following table provides end-of-support dates for Red Hat Enterprise Linux operating system versions. The information in the table is subject to change. For the most current information, see [Red Hat Enterprise Linux Life Cycle](#).

Table 11. Red Hat Enterprise Linux operating system versions: End of support

Operating system	End of support
Red Hat Enterprise Linux 7.4 EUS	August 31, 2019
Red Hat Enterprise Linux 7.4 for SAP Solutions	August 31, 2019
Red Hat Enterprise Linux 7.5 EUS	April 30, 2020
Red Hat Enterprise Linux 7.6 for SAP Solutions	October 31, 2022
Red Hat Enterprise Linux 7.6 EUS	October 31, 2020

VMware ESXi

The following table provides end-of-support dates for VMware ESXi hypervisor versions. For the latest list, see [VMware Lifecycle Product Matrix](#).

Table 12. VMware ESXi hypervisor versions: End of support

VMware version	End of support
ESXi 6.0	March 12, 2020
ESXi 6.5 and 6.7	November 15, 2021

Conclusion

Enterprises can now take advantage of the breakthrough technology of SAP HANA to achieve better business performance while keeping cost and complexity to a minimum. SAP has certified the Dell EMC PowerEdge servers for use in SAP HANA installations on production and nonproduction systems and on single-node (scale-up) and scale-out systems. By using the TDI approach with PowerEdge servers and selecting the right infrastructure architecture for your SAP HANA deployment, your organization can look ahead to faster and better IT innovation and responsiveness to meet your evolving business requirements.

References

Dell Technologies documentation

The following Dell Technologies documentation provides additional relevant information. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [SAP Info Hub for Ready Solutions](#)

Look under the following headings:

- SAP HANA TDI solutions for SAP HANA
 - White Papers
 - Case Studies
- [Dell EMC servers](#)

- [Dell EMC Networking Guides](#)

Under **MX-Series Modular Switches (PowerEdge MX)**:

- *PowerEdge MX IO Guide*
- *PowerEdge MX SmartFabric Configuration and Troubleshooting Guide*
- *MX Series Fibre Channel Storage Network Deployment with Ethernet IOMs*
- *PowerEdge MX SmartFabric Deployment – Video*
- *PowerEdge MX Network Architecture Guide*

VMware documentation

The following VMware documentation provides additional relevant information:

- [SAP Community Wiki: SAP HANA on VMware vSphere](#)
- [Architecture Guidelines and Best Practices for Deployments of SAP HANA on VMware vSphere: Architecture and Technical Considerations Guide](#)
- [Virtualize Applications: SAP HANA on vSphere—Support Status and Best Practices Summary](#)
- [SAP and VMware—Overview](#)
- [SAP Solutions on VMware Best Practices Guide](#)

SAP documentation

The following SAP documentation provides additional relevant information:

- [Quick Sizer](#)
- [SAP HANA on VMware vSphere](#)
- [SAP Help Portal](#)
- [SAP HANA Tailored Data Center Integration – Overview](#)
- [SAP HANA Tailored Data Center Integration – Frequently Asked Questions](#)

Note: The following documentation requires an SAP username and password.

- [SAP Note 2399079—Elimination of hdbparam in HANA 2](#)

Appendix A: Designing and implementing an SAP HANA scale-out landscape using the PowerEdge MX platform

This section describes how to design and implement the SAP HANA scale-out solution using PowerEdge MX kinetic infrastructure that was tested and validated in Dell EMC laboratories. The SAP HANA TDI scale-out solution described in this guide uses the following PowerEdge MX platform components and Dell EMC enterprise storage:

- **MX7000 chassis**—A 7U modular chassis that hosts flexible blocks of resources and provides outstanding efficiency through shared power, cooling, networking, I/O, and management.
- **PowerEdge MX networking modules**—Modular switches that are designed to reduce network management complexity and grow with business needs.
- **PowerEdge MX840c**—An SAP HANA certified, four-socket, double-width sled with elastic compute resources for database-driven, mission-critical applications and performance workloads.
- **Dell EMC Unity 650F**—An SAP HANA certified enterprise storage system with integrated architecture for block file with concurrent support for native NAS, iSCSI, and Fibre Channel (FC) protocols.

Network requirements

The SAP HANA scale-out solution requires numerous cables such as FC and Ethernet connections using rack servers in a fully redundant fabric configuration. The following figure shows a logical view of the minimum different network requirements:

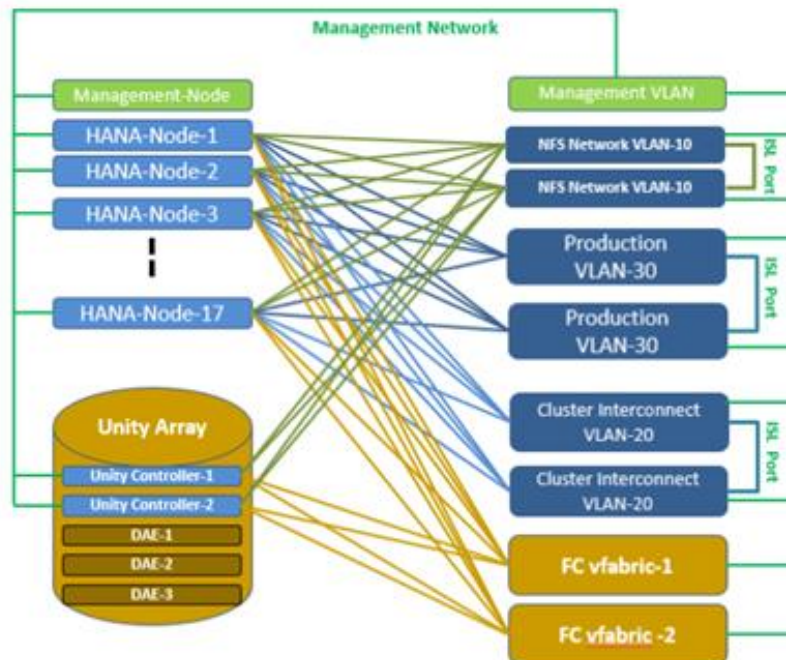


Figure 11. SAP HANA scale-out network configuration using MX9116n/MX7116n switches

A significant benefit of using the MX PowerEdge platform for an SAP HANA scale-out solution is its modular network infrastructure. This infrastructure helps to eliminate most network/fabric cabling from SAP HANA nodes to the virtual switch layer. Only switch interconnect links and network/fabric uplinks are required, as shown in the following figure:




Figure 12. Network cabling for an SAP HANA scale-out solution using two PowerEdge MX7000 chassis

The preceding illustration shows two PowerEdge MX7000 chassis, which can support up to eight MX840c servers (with four CPU sockets each), cascaded in a pair of Dell EMC modular network MX9116n/MX7116n switches. For illustration purposes, we used Dell EMC Unity array 650F and other network/fabric uplinks.

The PowerEdge MX7000 chassis is a 7U enclosure with eight front slots for blade-based servers (for up to four PowerEdge MX840c nodes) or storage sleds. The MX7000 chassis supports up to two pairs of redundant general-purpose switches or pass-through modular bays (Fabrics A and B) and a redundant pair of storage-specific switch bays (Fabric C). Up to 25Gbps Ethernet, 32Gbps FC, and 12Gbps SAS are also supported for the server's front-end I/O ports. The chassis has up to six power supply units (PSUs) as well as OpenManage Enterprise Modular Edition embedded management software running on up to two redundant MX9002m management modules.

Using Dell EMC modular network modules in PowerEdge MX

The PowerEdge MX platform introduced the concept of scalable fabric architecture. A scalable fabric spans multiple chassis, enabling them to behave like a single chassis from a networking perspective.

A scalable fabric consists of two main components: a pair of MX9116n Fabric Switch Engines (FSEs) in the first two chassis and additional pairs of MX7116n Fabric Expander Modules (FEMs) in the remaining chassis. Each MX7116n FEM connects to the MX9116n

FSE corresponding to its fabric and slot. All I/O modules (IOMs) participating in the fabric are configured in either Full Switch or SmartFabric mode.

The Dell EMC Networking MX9116n FSE operates in one of two modes:

- **SmartFabric Mode**—Switches operate as layer 2 I/O aggregation devices and are managed through the Open Manage Enterprise Modular console.
- **Full Switch Mode**—This mode is enabled by default, making all switch-specific OS10EE capabilities available.

The following figure shows three MX7000 chassis in a single scalable fabric architecture. The number of chassis is expandable to ten.




Figure 13. Scalable fabric example using Fabric A

The first two chassis each contain one MX9116n FSE and one MX7116n FEM. As shown in the diagram, chassis 3 to 10 each contain two MX7116n FEMs. All connections shown use QSFP28-DD (2x100GB) connections.

In an SAP HANA scale-out solution using the PowerEdge MX platform, one pair of MX9116n FSEs and additional FEM pairs for additional MX7000 chassis are used for all the required networks, including FC SAN. This helps to consolidate all the required network switches into one scalable fabric, keeping the fabric administration to a minimum and saving the data center footprint compared to an approach using rack servers.

The following figure shows the Dell EMC SAP HANA (16+2) scale-out solution using PowerEdge MX:

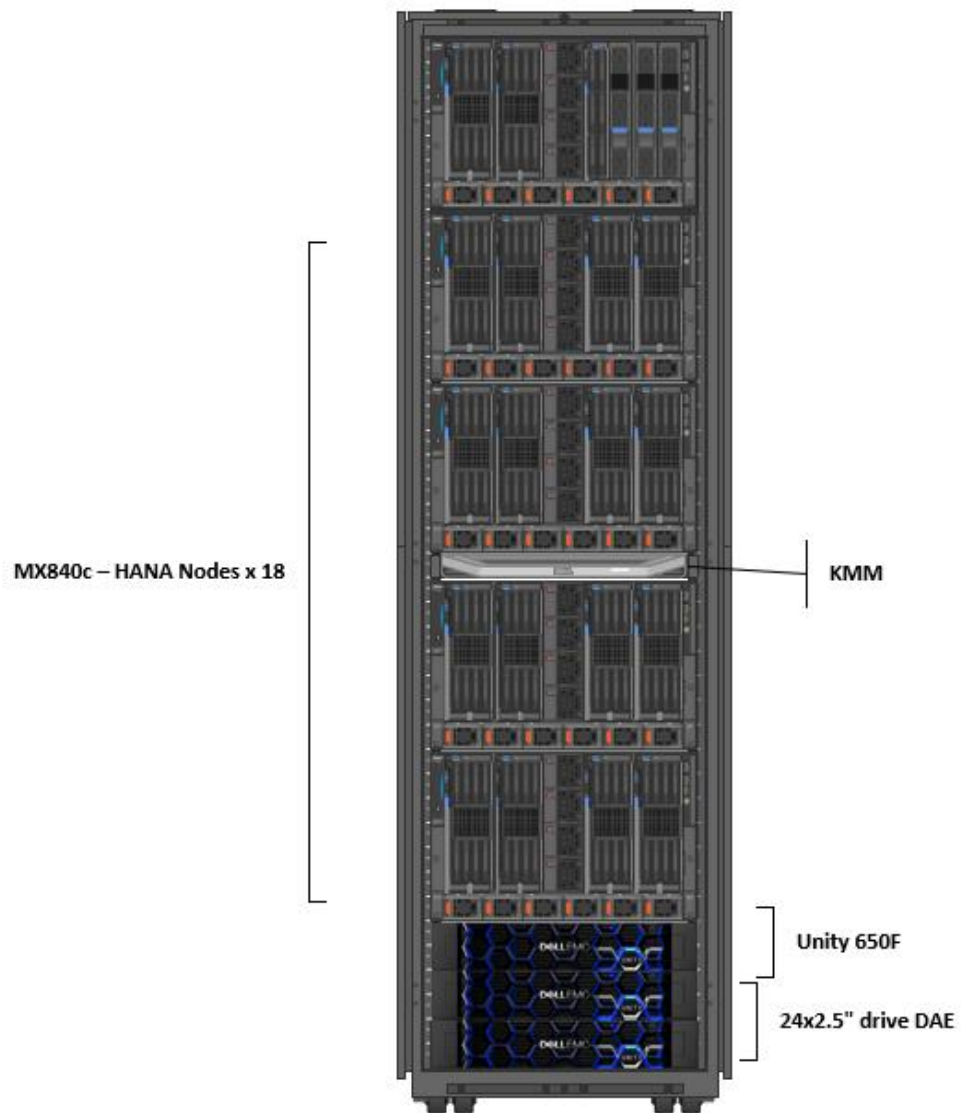


Figure 14. Dell EMC SAP HANA (16+2) scale-out solution using PowerEdge MX

The entire hardware requirement can fit in a 42U rack. The previous SAP HANA scale-out solution's rack design required double the rack space.

SAP HANA 6 TB scale-out solution

The following table shows the hardware components of an SAP HANA scale-out (2+1) solution using PowerEdge MX servers:

Table 13. Dell EMC SAP HANA 6 TB scale-out solution

Component	Configuration
SAP HANA nodes	2 x Dell EMC PowerEdge MX840c active SAP HANA nodes + 1 to 2 MX840c standby SAP HANA nodes. each with: <ul style="list-style-type: none"> SUSE Linux Enterprise Server for SAP Applications or Red Hat Enterprise Linux for SAP Solutions 2 x 25GbE dual port Mezzanine CNA cards
Storage	Any SAP HANA certified Dell EMC enterprise storage
Storage fabric and network	One pair of MX9116n modular switches to support both FC and all other Ethernet networks required by the SAP HANA scale-out landscape.

Two 25 GbE dual port mezzanine cards are used in each of the PowerEdge MX840c servers. 25 GbE ports are partitioned with specific bandwidth allocation, connected to the described fully redundant Scalable Fabric, and configured as bonded network or multipath I/O devices for Ethernet and FC fabric respectively to ensure that there is no single point of failure for the networks that an SAP HANA scale-out landscape requires.

The following figure shows the VLAN ID assignment in the deployed SAP HANA scale-out landscape:

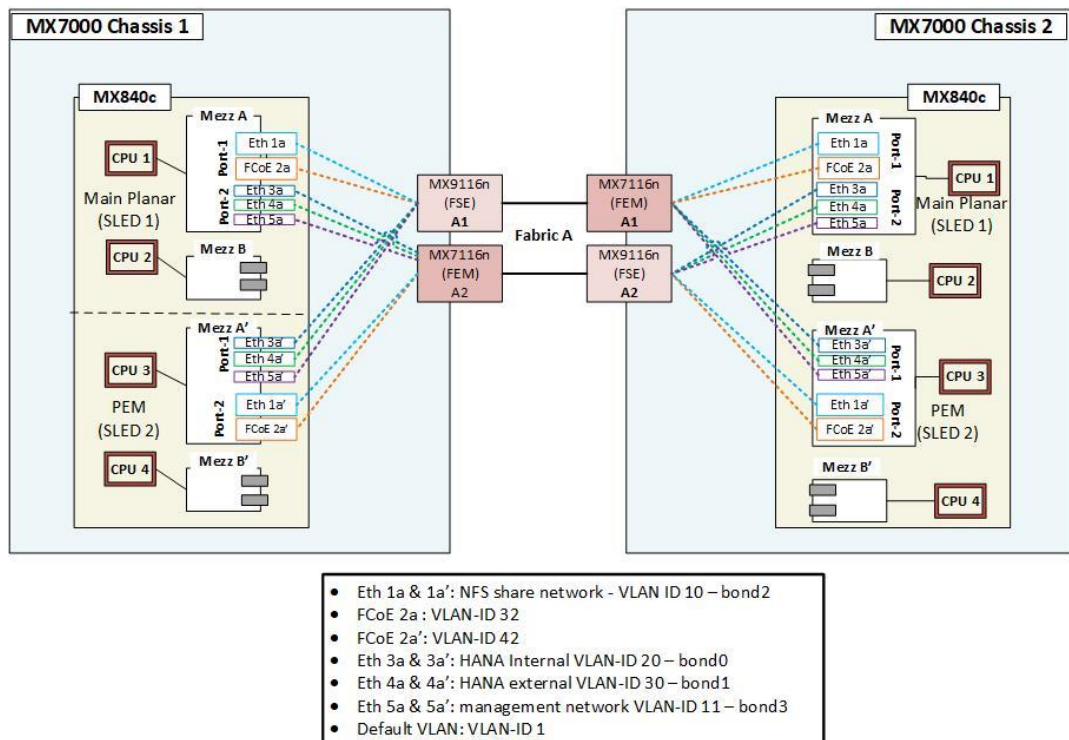


Figure 15. Network port partition on MX840c

The preceding figure and following table show how the 25GE mezzanine ports are partitioned. FC-related NIC partitions are configured as FCoE HBAs for FC SAN fabrics.

- The Unity array’s FC front-end HBAs are connected through the SAN fabric to MX9116n’s native FC uplinks.

- The MX7000 Fabric B module pair and future support of 25 GbE quad port mezzanine I/O cards can provide more network expansion capability if required.

A NIC partition such as 2a is allocated to have a minimum transfer bandwidth of 65 percent of the total available 25GE connection. Also, NIC partition 2a can have up to 100 percent of the available transfer bandwidth if NIC partition 1a is idle.

Table 14. CNA configuration within MX840c HANA database hosts

Mezz/CAN slot	Port number	Partition number/label	Partition type	Partition minimum TX bandwidth	Partition maximum TX bandwidth	Application function
Mezz 1A	Port 1	Partition 1 - 1a	NIC	35	100	NFS network
		Partition 2 - 2a	FCoE	65	100	FC SAN
	Port 2	Partition 1 - 3a	NIC	48	100	HANA internal network
		Partition 2 - 4a	NIC	48	100	HANA external network
		Partition 3 - 5a	NIC	4	100	Management network
Mezz 2A	Port 1	Partition 1 - 3a'	NIC	48	100	HANA internal network
		Partition 2 - 4a'	NIC	48	100	HANA external network
		Partition 3 - 5a'	NIC	4	100	management network
	Port 2	Partition 1 - 1a'	NIC	35	100	NFS network
		Partition 2 - 2a'	FCoE	65	100	FC SAN

Controlling, automating, and managing the SAP HANA configuration

PowerEdge MX networking switches support Dell EMC OS10 Enterprise Edition and operate in either Full Switch or SmartFabric mode. Users have more control over IOM in Full Switch mode because IOM operates as a regular switch. SmartFabric offers several key features that are crucial for the latest fabric solutions, such as automation, life cycle management, scalability, and manageability.

We implemented our SAP HANA scale-out solution in SmartFabric mode for proof of concept and certification purposes. With SmartFabric Services (SFS), customers can quickly and easily deploy and automate datacenter networking fabrics. The most burdensome tasks were creating NIC partitions from the first SAP HANA server as a template server and creating different VLAN IDs and network uplinks on the MX9116n pair. The template is easily and automatically deployed to the next SAP HANA server as well as its Ethernet and SAN fabric configuration on the switches involved. For more information, see the following documents:

- [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide.](#)
- [Dell EMC PowerEdge MX Series Fibre Channel Storage Network Deployment with Ethernet IOMs](#)

In addition to being able to assign VLANs to server profiles, SFS automates quality of service (QoS) settings based on user input. When a VLAN is created, the user selects the related traffic type—iSCSI, vMotion, and so on—and the SFS engine assigns the correct QoS setting to that VLAN.

Users can also select a “metal” such as gold, bronze, and so on to assign their own priority values to traffic. The following table shows the traffic types and related settings:

Table 15. Traffic types and QoS settings

Traffic type	Used for	QoS setting
General purpose (bronze)	Low priority data traffic	2
General purpose (silver)	Standard/default priority data traffic	3
General purpose (gold)	High priority data traffic	4
General purpose (platinum)	Extremely high priority data traffic	5
Cluster interconnect	Cluster heartbeat VLANs	5
Hypervisor management	Hypervisor management connections such as the ESXi management VLAN	5
Storage - iscsi	Used for iSCSI VLANs	5
Storage – fcoe	Used for FCoE VLANs	5
Storage - data replication	Used for VLANs supporting storage data replication such as for VMware VSAN	5
VM migration	Used for VLANs supporting vMotion and similar technologies	5
VMware ft logging	Used for VLANs supporting VMware fault tolerance	5

Appendix B: SAPS values for SAP HANA certified PowerEdge servers

Introduction

[SAP HANA TDI Phase 5](#) introduced important changes and benefits for customers who are planning for SAP HANA deployments:

- The new workload-driven sizing approach allows for SAPS-based sizing. The Quick Sizer tool and SAP HANA sizing reports have been enhanced to provide separate CPU (SAPS) and RAM sizing results, enabling customers to fine-tune their configurations for their specific workload and to purchase systems with the optimal number of cores and memory.
- The increased variety of CPU processors that are supported enables customers to be more flexible when choosing SAP HANA compute nodes from the PowerEdge Skylake- and Cascade Lake-based range of systems and to cost-optimize the servers for their workload requirements.

SAPS values for PowerEdge servers

The following tables show the estimated SAPS values for PowerEdge R940, R840, MX840c, R740, R740xd, and MX740c servers with over 90 percent utilization. The SAPS values for the top-end platinum CPU models, the Intel Xeon Platinum 8180 (Skylake) and 8280 (Cascade Lake), are certified with SAP. These certified values are presented in bold in the tables. Dell EMC extrapolated the SAPS values for all other PowerEdge server models from these results using SPECint2017. All the listed models are SAP HANA certified with [SAP HANA TDI Phase 5](#). Bronze CPU ranges and any servers which do not meet the SAP HANA minimum requirements with less than two sockets and eight cores per socket are not supported and therefore not listed in the tables.

Note: Dell EMC performed internal testing with the SD benchmarks to validate the extrapolated SAPS values across the platinum, gold, and silver CPU ranges. The results showed that the extrapolations were accurate to within a 5 to 10 percent range in our laboratory environment. "M/L" editions of a CPU model have the same extrapolated SAPS values as non-M and non-L editions of that model. The SAP Quick Sizer tool bases calculations on a 65 percent utilization with greenfield sizing. For brownfield sizing, calculate 60-65 percent utilization from the extrapolated SAPS values listed here. The extrapolated values are for bare-metal environments. You can expect a performance degradation of approximately 10 percent with virtualized environments.

Table 16. PowerEdge R940 Cascade Lake-based systems with SAPS values

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R940 (Intel Xeon Gold 5217, 3.00GHz)	32	4	8	121,048	3,783
PowerEdge R940 (Intel Xeon Gold 5218, 2.30GHz)	64	4	16	210,088	3,283
PowerEdge R940 (Intel Xeon Gold 5220, 2.20GHz)	72	4	18	230,456	3,201
PowerEdge R940 (Intel Xeon Gold 6222V, 1.80GHz)	80	4	20	229,292	2,866
PowerEdge R940 (Intel Xeon Gold 6226, 2.70GHz)	48	4	12	189,719	3,952
PowerEdge R940 (Intel Xeon Gold 6230, 2.10GHz)	80	4	20	258,390	3,230

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R940 (Intel Xeon Gold 6234, 3.30GHz)	32	4	8	147,818	4,619
PowerEdge R940 (Intel Xeon Gold 6238, 2.10GHz)	88	4	22	282,251	3,207
PowerEdge R940 (Intel Xeon Gold 6240, 2.60GHz)	72	4	18	258,972	3,597
PowerEdge R940 (Intel Xeon Gold 6242, 2.80GHz)	64	4	16	238,604	3,728
PowerEdge R940 (Intel Xeon Gold 6246, 3.30GHz)	48	4	12	209,506	4,365
PowerEdge R940 (Intel Xeon Gold 6248, 2.50GHz)	80	4	20	279,341	3,492
PowerEdge R940 (Intel Xeon Gold 6252, 2.10GHz)	96	4	24	303,201	3,158
PowerEdge R940 (Intel Xeon Platinum 8253, 2.20GHz)	64	4	16	186,227	2,910
PowerEdge R940 (Intel Xeon Platinum 8268, 2.90GHz)	96	4	24	349,758	3,643
PowerEdge R940 (Intel Xeon Platinum 8270, 2.70GHz)	104	4	26	359,651	3,458
PowerEdge R940 (Intel Xeon Platinum 8276, 2.20GHz)	112	4	28	342,193	3,055
PowerEdge R940 (Intel Xeon Platinum 8280, 2.70GHz)	112	4	28	380,020	3,393

Table 17. PowerEdge R940 Skylake-based systems with SAPS values

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R940 (Intel Xeon Gold 5115, 2.40 GHz)	40	4	10	126,850	3,171
PowerEdge R940 (Intel Xeon Gold 5118, 2.30 GHz)	48	4	12	147,270	3,068
PowerEdge R940 (Intel Xeon Gold 5120, 2.20 GHz)	56	4	14	165,833	2,961
PowerEdge R940 (Intel Xeon Gold 6126, 2.60 GHz)	48	4	12	168,927	3,519
PowerEdge R940 (Intel Xeon Gold 6130, 2.10 GHz)	64	4	16	198,628	3,104
PowerEdge R940 (Intel Xeon Gold 6132, 2.60 GHz)	56	4	14	196,772	3,514
PowerEdge R940 (Intel Xeon Gold 6134, 3.20 GHz)	32	4	8	131,800	4,119
PowerEdge R940 (Intel Xeon Gold 6136, 3.00 GHz)	48	4	12	188,109	3,919
PowerEdge R940 (Intel Xeon Gold 6138, 2.00 GHz)	80	4	20	227,711	2,846
PowerEdge R940 (Intel Xeon Gold 6140, 2.30 GHz)	72	4	18	227,711	3,163
PowerEdge R940 (Intel Xeon Gold 6142, 2.60 GHz)	64	4	16	215,954	3,374
PowerEdge R940 (Intel Xeon Gold 6144, 3.50 GHz)	32	4	8	142,319	4,447
PowerEdge R940 (Intel Xeon Gold 6148, 2.40 GHz)	80	4	20	251,843	3,148
PowerEdge R940 (Intel Xeon Gold 6150, 2.70 GHz)	72	4	18	245,037	3,403
PowerEdge R940 (Intel Xeon Gold 6152, 2.10 GHz)	88	4	22	257,413	2,925
PowerEdge R940 (Intel Xeon Gold 6154, 3.00 GHz)	72	4	18	262,982	3,653
PowerEdge R940 (Intel Xeon Platinum 8153, 2.00 GHz)	64	4	16	173,877	2,717

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R940 (Intel Xeon Platinum 8158, 3.00 GHz)	48	4	12	189,347	3,945
PowerEdge R940 (Intel Xeon Platinum 8160, 2.10 GHz)	96	4	24	280,307	2,920
PowerEdge R940 (Intel Xeon Platinum 8164, 2.00 GHz)	104	4	26	291,445	2,802
PowerEdge R940 (Intel Xeon Platinum 8168, 2.70 GHz)	96	4	24	314,959	3,281
PowerEdge R940 (Intel Xeon Platinum 8168, 2.70GHz)	96	4	24	350,848	3,655
PowerEdge R940 (Intel Xeon Platinum 8170, 2.10 GHz)	104	4	26	297,633	2,862
PowerEdge R940 (Intel Xeon Platinum 8176, 2.10 GHz)	112	4	28	313,103	2,796
PowerEdge R940 (Intel Xeon Platinum 8180, 2.50 GHz)	112	4	28	360,130	3,215

Table 18. PowerEdge R940xa Skylake-based systems with SAPS values

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R940xa (Intel Xeon Gold 5115, 2.40GHz)	40	4	10	158,976	3,974
PowerEdge R940xa (Intel Xeon Gold 5117, 2.00GHz)	56	4	14	188,386	3,364
PowerEdge R940xa (Intel Xeon Gold 5118, 2.30GHz)	48	4	12	188,386	3,925
PowerEdge R940xa (Intel Xeon Gold 5120, 2.20GHz)	56	4	14	209,848	3,747
PowerEdge R940xa (Intel Xeon Gold 6126, 2.60GHz)	48	4	12	213,822	4,455
PowerEdge R940xa (Intel Xeon Gold 6130, 2.10GHz)	64	4	16	241,643	3,776
PowerEdge R940xa (Intel Xeon Gold 6132, 2.60GHz)	56	4	14	241,643	4,315
PowerEdge R940xa (Intel Xeon Gold 6134, 3.20GHz)	32	4	8	166,925	5,216
PowerEdge R940xa (Intel Xeon Gold 6136, 3.00GHz)	48	4	12	226,540	4,720
PowerEdge R940xa (Intel Xeon Gold 6138, 2.00GHz)	80	4	20	303,644	3,796
PowerEdge R940xa (Intel Xeon Gold 6140, 2.30GHz)	72	4	18	287,746	3,996
PowerEdge R940xa (Intel Xeon Gold 6142, 2.60GHz)	64	4	16	270,259	4,223
PowerEdge R940xa (Intel Xeon Gold 6144, 3.50GHz)	32	4	8	170,104	5,316
PowerEdge R940xa (Intel Xeon Gold 6146, 3.20GHz)	48	4	12	239,258	4,985
PowerEdge R940xa (Intel Xeon Gold 6148, 2.40GHz)	80	4	20	328,285	4,104
PowerEdge R940xa (Intel Xeon Gold 6150, 2.70GHz)	72	4	18	306,823	4,261
PowerEdge R940xa (Intel Xeon Gold 6152, 2.10GHz)	88	4	22	337,029	3,830
PowerEdge R940xa (Intel Xeon Gold 6154, 3.00GHz)	72	4	18	370,413	5,145
PowerEdge R940xa (Intel Xeon Gold 6154, 3.00GHz)	72	4	18	344,182	4,780
PowerEdge R940xa (Intel Xeon Platinum 8153, 2.00GHz)	64	4	16	222,566	3,478
PowerEdge R940xa (Intel Xeon Platinum 8160, 2.10GHz)	96	4	24	345,772	3,602

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R940xa (Intel Xeon Platinum 8160, 2.10GHz)	96	4	24	372,003	3,875
PowerEdge R940xa (Intel Xeon Platinum 8164, 2.00GHz)	104	4	26	349,747	3,363
PowerEdge R940xa (Intel Xeon Platinum 8168, 2.70GHz)	96	4	24	372,003	3,875
PowerEdge R940xa (Intel Xeon Platinum 8170, 2.10GHz)	104	4	26	372,003	3,577
PowerEdge R940xa (Intel Xeon Platinum 8176, 2.10GHz)	112	4	28	391,080	3,492
PowerEdge R940xa (Intel Xeon Platinum 8180, 2.50GHz)	112	4	28	360,080	3,215

Table 19. PowerEdge MX840c Skylake-based systems with SAPS values

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge MX840c (Intel Xeon Gold 5118 CPU, 2.30GHz)	48	4	12	160,694	3,348
PowerEdge MX840c (Intel Xeon Gold 5120 CPU, 2.20GHz)	56	4	14	180,440	3,222
PowerEdge MX840c (Intel Xeon Gold 6130 CPU, 2.10GHz)	64	4	16	212,443	3,319
PowerEdge MX840c (Intel Xeon Gold 6134, 3.20GHz)	32	4	8	145,714	4,554
PowerEdge MX840c (Intel Xeon Gold 6144, 3.50GHz)	32	4	8	154,566	4,830
PowerEdge MX840c (Intel Xeon Gold 6146, 3.20GHz)	48	4	12	181,121	3,773
PowerEdge MX840c (Intel Xeon Gold 6148 CPU, 2.40GHz)	80	4	20	271,682	3,396
PowerEdge MX840c (Intel Xeon Gold 6152 CPU, 2.10GHz)	88	4	22	274,406	3,118
PowerEdge MX840c (Intel Xeon Gold 6154, 3.00GHz)	72	4	18	239,679	3,329
PowerEdge MX840c (Intel Xeon Platinum 8160 CPU, 2.10GHz)	96	4	24	298,237	3,107
PowerEdge MX840c (Intel Xeon Platinum 8168 CPU, 2.70GHz)	96	4	24	335,687	3,497
PowerEdge MX840c (Intel Xeon Platinum 8176 CPU, 2.10GHz)	112	4	28	337,049	3,009
PowerEdge MX840c (Intel Xeon Platinum 8180 CPU, 2.50GHz)	112	4	28	360,200	3,216

Table 20. PowerEdge R840 Cascade Lake-based systems with SAPS values

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R840 (Intel Xeon Gold 5218, 2.30GHz)	64	4	16	212,630	3,322
PowerEdge R840 (Intel Xeon Gold 6222V, 1.80GHz)	80	4	20	231,322	2,892
PowerEdge R840 (Intel Xeon Gold 6226, 2.70GHz)	48	4	12	189,264	3,943
PowerEdge R840 (Intel Xeon Gold 6230, 2.10GHz)	80	4	20	257,609	3,220

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R840 (Intel Xeon Gold 6238, 2.10GHz)	88	4	22	284,480	3,233
PowerEdge R840 (Intel Xeon Gold 6240, 2.60GHz)	72	4	18	257,025	3,570
PowerEdge R840 (Intel Xeon Gold 6242, 2.80GHz)	64	4	16	237,748	3,715
PowerEdge R840 (Intel Xeon Gold 6246, 3.30GHz)	48	4	12	208,541	4,345
PowerEdge R840 (Intel Xeon Gold 6248, 2.50GHz)	80	4	20	278,638	3,483
PowerEdge R840 (Intel Xeon Gold 6252, 2.10GHz)	96	4	24	300,252	3,128
PowerEdge R840 (Intel Xeon Gold 6262V)	96	4	24	276,886	2,884
PowerEdge R840 (Intel Xeon Platinum 8253, 2.20GHz)	64	4	16	197,442	3,085
PowerEdge R840 (Intel Xeon Platinum 8260M, 2.40GHz)	96	4	24	319,529	3,328
PowerEdge R840 (Intel Xeon Platinum 8268, 2.90GHz)	96	4	24	349,320	3,639
PowerEdge R840 (Intel Xeon Platinum 8270, 2.70GHz)	104	4	26	364,508	3,505
PowerEdge R840 (Intel Xeon Platinum 8276, 2.20GHz)	112	4	28	335,885	2,999
PowerEdge R840 (Intel Xeon Platinum 8280, 2.70GHz)	112	4	28	380,280	3,395

Table 21. PowerEdge R840 Skylake-based systems with SAPS values

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R840 (Intel Xeon Gold 5115, 2.40 GHz)	20	2	10	62,949	3,147
PowerEdge R840 (Intel Xeon Gold 5115, 2.40 GHz)	40	4	10	124,652	3,116
PowerEdge R840 (Intel Xeon Gold 5117, 2.00GHz)	56	4	14	147,712	2,638
PowerEdge R840 (Intel Xeon Gold 5118, 2.30 GHz)	24	2	12	74,791	3,116
PowerEdge R840 (Intel Xeon Gold 5118, 2.30 GHz)	48	4	12	147,712	3,077
PowerEdge R840 (Intel Xeon Gold 5120, 2.20 GHz)	28	2	14	83,517	2,983
PowerEdge R840 (Intel Xeon Gold 5120, 2.20 GHz)	56	4	14	164,540	2,938
PowerEdge R840 (Intel Xeon Gold 6126, 2.60 GHz)	24	2	12	85,386	3,558
PowerEdge R840 (Intel Xeon Gold 6126, 2.60 GHz)	48	4	12	167,656	3,493
PowerEdge R840 (Intel Xeon Gold 6130, 2.10 GHz)	32	2	16	95,982	2,999
PowerEdge R840 (Intel Xeon Gold 6130, 2.10 GHz)	64	4	16	198,819	3,107
PowerEdge R840 (Intel Xeon Gold 6132, 2.60 GHz)	28	2	14	100,345	3,584
PowerEdge R840 (Intel Xeon Gold 6134, 3.20 GHz)	16	2	8	66,689	4,168
PowerEdge R840 (Intel Xeon Gold 6134, 3.20 GHz)	32	4	8	130,884	4,090
PowerEdge R840 (Intel Xeon Gold 6136, 3.00 GHz)	24	2	12	92,242	3,843

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R840 (Intel Xeon Gold 6136, 3.00 GHz)	48	4	12	177,629	3,701
PowerEdge R840 (Intel Xeon Gold 6138, 2.00 GHz)	40	2	20	115,926	2,898
PowerEdge R840 (Intel Xeon Gold 6138, 2.00 GHz)	80	4	20	225,619	2,820
PowerEdge R840 (Intel Xeon Gold 6140, 2.30 GHz)	36	2	18	115,303	3,203
PowerEdge R840 (Intel Xeon Gold 6140, 2.30GHz)	72	4	18	225,619	3,134
PowerEdge R840 (Intel Xeon Gold 6142, 2.60 GHz)	32	2	16	106,577	3,331
PowerEdge R840 (Intel Xeon Gold 6142, 2.60 GHz)	64	4	16	211,908	3,311
PowerEdge R840 (Intel Xeon Gold 6144, 3.50 GHz)	16	2	8	71,675	4,480
PowerEdge R840 (Intel Xeon Gold 6144, 3.50 GHz)	32	4	8	133,377	4,168
PowerEdge R840 (Intel Xeon Gold 6146, 3.20 GHz)	24	2	12	96,605	4,025
PowerEdge R840 (Intel Xeon Gold 6146, 3.20 GHz)	48	4	12	187,601	3,908
PowerEdge R840 (Intel Xeon Gold 6148, 2.40 GHz)	80	4	20	244,317	3,054
PowerEdge R840 (Intel Xeon Gold 6148, 2.40GHz)	40	2	20	125,275	3,132
PowerEdge R840 (Intel Xeon Gold 6150, 2.70 GHz)	36	2	18	122,782	3,411
PowerEdge R840 (Intel Xeon Gold 6150, 2.70 GHz)	72	4	18	240,578	3,341
PowerEdge R840 (Intel Xeon Gold 6152, 2.10 GHz)	44	2	22	126,521	2,875
PowerEdge R840 (Intel Xeon Gold 6152, 2.10 GHz)	88	4	22	249,927	2,840
PowerEdge R840 (Intel Xeon Gold 6154, 3.00GHz)	36	2	18	130,884	3,636
PowerEdge R840 (Intel Xeon Gold 6154, 3.00GHz)	72	4	18	262,392	3,644
PowerEdge R840 (Intel Xeon Platinum 8153, 2.00 GHz)	32	2	16	87,879	2,746
PowerEdge R840 (Intel Xeon Platinum 8153, 2.00 GHz)	64	4	16	174,512	2,727
PowerEdge R840 (Intel Xeon Platinum 8160, 2.10 GHz)	48	2	24	137,740	2,870
PowerEdge R840 (Intel Xeon Platinum 8160, 2.10 GHz)	96	4	24	271,117	2,824
PowerEdge R840 (Intel Xeon Platinum 8164, 2.00 GHz)	52	2	26	143,973	2,769
PowerEdge R840 (Intel Xeon Platinum 8164, 2.00 GHz)	104	4	26	274,234	2,637
PowerEdge R840 (Intel Xeon Platinum 8168, 2.70 GHz)	48	2	24	160,177	3,337
PowerEdge R840 (Intel Xeon Platinum 8170, 2.10 GHz)	52	2	26	149,582	2,877
PowerEdge R840 (Intel Xeon Platinum 8170, 2.10 GHz)	104	4	26	291,685	2,805
PowerEdge R840 (Intel Xeon Platinum 8176, 2.10 GHz)	56	2	28	156,438	2,794
PowerEdge R840 (Intel Xeon Platinum 8176, 2.10 GHz)	112	4	28	306,020	2,732
PowerEdge R840 (Intel Xeon Platinum 8180, 2.50GHz)	56	2	28	182,615	3,261
PowerEdge R840 (Intel Xeon Platinum 8180, 2.50GHz)	112	4	28	359,620	3,211

Table 22. PowerEdge R740, R740xd, MX740c Cascade Lake Refresh-based systems with SAPS values

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R740xd (Intel Xeon Bronze 3206R, 1.90 GHz)	16	2	8	30,268	1,892
PowerEdge R740xd (Intel Xeon Silver 4210R, 2.40 GHz)	20	2	10	64,617	3,231
PowerEdge R740xd (Intel Xeon Silver 4214R, 2.40 GHz)	24	2	12	81,054	3,377
PowerEdge R740xd (Intel Xeon Silver 4215R, 3.20 GHz)	16	2	8	57,248	3,578
PowerEdge R740xd (Intel Xeon Gold 5218R, 2.10 GHz)	40	2	20	120,164	3,004
PowerEdge R740xd (Intel Xeon Gold 5220R, 2.20 GHz)	48	2	24	146,805	3,058
PowerEdge R740xd (Intel Xeon Gold 6208U, 2.90 GHz)	16	1	16	57,248	3,578
PowerEdge R740xd (Intel Xeon Gold 6226R, 2.90 GHz)	32	2	16	118,464	3,702
PowerEdge R740xd (Intel Xeon Gold 6230R, 2.10 GHz)	52	2	26	153,606	2,954
PowerEdge R740xd (Intel Xeon Gold 6238R, 2.20 GHz)	56	2	28	163,809	2,925
PowerEdge R740xd (Intel Xeon Gold 6240R, 2.40 GHz)	48	2	24	146,805	3,058
PowerEdge R740xd (Intel Xeon Gold 6242R, 3.10 GHz)	40	2	20	153,606	3,840
PowerEdge R740xd (Intel Xeon Gold 6246R, 3.40 GHz)	32	2	16	136,602	4,269
PowerEdge R740xd (Intel Xeon Gold 6248R, 3.00 GHz)	48	2	24	175,145	3,649
PowerEdge R740xd (Intel Xeon Gold 6258R, 2.70 GHz)	56	2	28	184,781	3,300

Table 23. PowerEdge R740/R740xd Cascade Lake-based systems with SAPS values

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R740 (Intel Xeon Silver 4214, 2.20GHz)	24	2	12	75,953	3,165
PowerEdge R740 (Intel Xeon Gold 5218, 2.30GHz)	32	2	16	103,727	3,241
PowerEdge R740 (Intel Xeon Gold 5220, 2.20GHz)	36	2	18	108,261	3,007
PowerEdge R740 (Intel Xeon Gold 6230N, 2.30GHz)	40	2	20	125,266	3,132
PowerEdge R740 (Intel Xeon Gold 6242, 2.80GHz)	32	2	16	115,630	3,613
PowerEdge R740 (Intel Xeon Gold 6248, 2.50GHz)	40	2	20	140,003	3,500
PowerEdge R740 (Intel Xeon Gold 6254, 3.10GHz)	36	2	18	140,003	3,889
PowerEdge R740 (Intel Xeon Platinum 8253, 2.20GHz)	32	2	16	90,690	2,834
PowerEdge R740xd (Intel Xeon Silver 4208, 2.10GHz)	16	2	8	47,159	2,947
PowerEdge R740xd (Intel Xeon Silver 4210, 2.20GHz)	20	2	10	60,082	3,004
PowerEdge R740xd (Intel Xeon Silver 4214, 2.20GHz)	24	2	12	76,520	3,188
PowerEdge R740xd (Intel Xeon Silver 4214, 2.20GHz)	24	2	12	74,819	3,117

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R740xd (Intel Xeon Silver 4215, 2.50GHz)	16	2	8	52,657	3,291
PowerEdge R740xd (Intel Xeon Silver 4216, 2.10GHz)	16	1	16	48,122	3,008
PowerEdge R740xd (Intel Xeon Silver 4216, 2.10GHz)	32	2	16	99,192	3,100
PowerEdge R740xd (Intel Xeon Gold 5215, 2.50GHz)	20	2	10	68,584	3,429
PowerEdge R740xd (Intel Xeon Gold 5217, 3.00GHz)	16	2	8	60,082	3,755
PowerEdge R740xd (Intel Xeon Gold 5218N, 2.30GHz)	32	2	16	108,261	3,383
PowerEdge R740xd (Intel Xeon Gold 5220S, 2.70GHz)	36	2	18	113,363	3,149
PowerEdge R740xd (Intel Xeon Gold 6210U, 2.50GHz)	20	1	20	70,285	3,514
PowerEdge R740xd (Intel Xeon Gold 6226, 2.70GHz)	24	2	12	93,524	3,897
PowerEdge R740xd (Intel Xeon Gold 6230, 2.10GHz)	40	2	20	126,966	3,174
PowerEdge R740xd (Intel Xeon Gold 6238, 2.10GHz)	44	2	22	138,869	3,156
PowerEdge R740xd (Intel Xeon Gold 6240, 2.60GHz)	36	2	18	128,100	3,558
PowerEdge R740xd (Intel Xeon Gold 6244, 3.60GHz)	16	2	8	75,386	4,712
PowerEdge R740xd (Intel Xeon Gold 6246, 3.30GHz)	24	2	12	102,593	4,275
PowerEdge R740xd (Intel Xeon Gold 6252, 2.10GHz)	48	2	24	150,205	3,129
PowerEdge R740xd (Intel Xeon Gold 6252N, 2.30GHz)	48	2	24	154,173	3,212
PowerEdge R740xd (Intel Xeon Platinum 8260, 2.40GHz)	48	2	24	153,606	3,200
PowerEdge R740xd (Intel Xeon Platinum 8268, 2.90GHz)	48	2	24	171,744	3,578
PowerEdge R740xd (Intel Xeon Platinum 8270, 2.70GHz)	52	2	26	180,246	3,466
PowerEdge R740xd (Intel Xeon Platinum 8276M, 2.20GHz)	56	2	28	167,777	2,996
PowerEdge R740xd (Intel Xeon Platinum 8280, 2.70GHz)	56	2	28	193,850	3,462

Table 24. PowerEdge MX740c Skylake-based systems with SAPS values

PowerEdge server	Cores	Sockets	Cores per socket	SAPS	SAPS per core
PowerEdge MX740c (Intel Xeon Silver 4108 CPU, 1.80GHz)	16	2	8	41,622	2,601
PowerEdge MX740c (Intel Xeon Silver 4110 CPU, 2.10GHz)	16	2	8	47,081	2,943
PowerEdge MX740c (Intel Xeon Silver 4114 CPU, 2.20GHz)	20	2	10	59,413	2,971
PowerEdge MX740c (Intel Xeon Silver 4116, 2.10GHz)	24	2	12	68,727	2,864
PowerEdge MX740c (Intel Xeon Gold 5117 CPU, 2.00GHz)	28	2	14	77,077	2,753
PowerEdge MX740c (Intel Xeon Gold 5118 CPU, 2.30GHz)	24	2	12	76,435	3,185
PowerEdge MX740c (Intel Xeon Gold 5120 CPU, 2.20GHz)	28	2	14	85,427	3,051
PowerEdge MX740c (Intel Xeon Gold 6130 CPU, 2.10GHz)	32	2	16	100,842	3,151

PowerEdge server	Cores	Sockets	Cores per socket	SAPS	SAPS per core
PowerEdge MX740c (Intel Xeon Gold 6134, 3.20GHz)	16	2	8	68,085	4,255
PowerEdge MX740c (Intel Xeon Gold 6144, 3.50GHz)	16	2	8	72,581	4,536
PowerEdge MX740c (Intel Xeon Gold 6146, 3.20GHz)	24	2	12	90,565	3,774
PowerEdge MX740c (Intel Xeon Gold 6152 CPU, 2.10GHz)	44	2	22	131,031	2,978
PowerEdge MX740c (Intel Xeon Platinum 8160 CPU, 2.10GHz)	48	2	24	145,804	3,038
PowerEdge MX740c (Intel Xeon Platinum 8168, 2.70GHz)	48	2	24	159,292	3,319
PowerEdge MX740c (Intel Xeon Platinum 8176, 2.10GHz)	56	2	28	160,577	2,867
PowerEdge MX740c (Intel Xeon Platinum 8180 CPU, 2.50GHz)	56	2	28	175,350	3,131

Table 25. PowerEdge R740/R740xd Skylake-based systems with SAPS values

PowerEdge server	Cores	Sockets	Cores per socket	SAPS	SAPS per core
PowerEdge R740xd (Intel Xeon Bronze 3106, 1.70 GHz)	16	2	8	26,166	1,635
PowerEdge R740xd (Intel Xeon Silver 4108, 1.80 GHz)	16	2	8	38,835	2,427
PowerEdge R740xd (Intel Xeon Silver 4109T, 2.00 GHz)	16	2	8	41,736	2,608
PowerEdge R740xd (Intel Xeon Silver 4110, 2.10 GHz)	16	2	8	43,097	2,694
PowerEdge R740xd (Intel Xeon Silver 4114, 2.20 GHz)	20	2	10	55,825	2,791
PowerEdge R740xd (Intel Xeon Silver 4116, 2.10 GHz)	24	2	12	65,119	2,713
PowerEdge R740xd (Intel Xeon Gold 5115, 2.40 GHz)	20	2	10	60,975	3,049
PowerEdge R740xd (Intel Xeon Gold 5118, 2.30 GHz)	24	2	12	71,631	2,985
PowerEdge R740xd (Intel Xeon Gold 5120, 2.20 GHz)	28	2	14	79,919	2,854
PowerEdge R740xd (Intel Xeon Gold 6130, 2.10 GHz)	32	2	16	97,087	3,034
PowerEdge R740xd (Intel Xeon Gold 6132, 2.60 GHz)	28	2	14	97,087	3,467
PowerEdge R740xd (Intel Xeon Gold 6134, 3.20 GHz)	16	2	8	63,935	3,996
PowerEdge R740xd (Intel Xeon Gold 6136, 3.00 GHz)	24	2	12	89,983	3,749
PowerEdge R740xd (Intel Xeon Gold 6138, 2.00 GHz)	40	2	20	111,887	2,797
PowerEdge R740xd (Intel Xeon Gold 6140, 2.30 GHz)	36	2	18	112,479	3,124
PowerEdge R740xd (Intel Xeon Gold 6142, 2.60 GHz)	32	2	16	104,783	3,274
PowerEdge R740xd (Intel Xeon Gold 6144, 3.50 GHz)	16	2	8	69,263	4,329
PowerEdge R740xd (Intel Xeon Gold 6146, 3.20GHz)	24	2	12	94,719	3,947
PowerEdge R740xd (Intel Xeon Gold 6150, 2.70 GHz)	36	2	18	120,175	3,338

PowerEdge server	Cores	Sockets	Cores per socket	SAPS	SAPS per core
PowerEdge R740xd (Intel Xeon Gold 6152, 2.10 GHz)	44	2	22	126,095	2,866
PowerEdge R740xd (Intel Xeon Gold 6154, 3.00 GHz)	36	2	18	130,239	3,618
PowerEdge R740xd (Intel Xeon Platinum 8153, 2.00 GHz)	32	2	16	84,063	2,627
PowerEdge R740xd (Intel Xeon Platinum 8170, 2.10 GHz)	52	2	26	145,038	2,789
PowerEdge R740xd (Intel Xeon Platinum 8176, 2.10GHz)	56	2	28	155,102	2,770
PowerEdge R740 (Intel Xeon Silver 4108, 1.80GHz)	16	2	8	38,716	2,420
PowerEdge R740 (Intel Xeon Gold 5117, 2.00 GHz)	28	2	14	71,039	2,537
PowerEdge R740 (Intel Xeon Gold 6126, 2.60GHz)	24	2	12	83,471	3,478
PowerEdge R740 (Intel Xeon Gold 6137, 3.90GHz)	16	2	8	69,263	4,329
PowerEdge R740 (Intel Xeon Gold 6148, 2.40GHz)	40	2	20	121,359	3,034
PowerEdge R740 (Intel Xeon Platinum 8160, 2.10GHz)	48	2	24	134,382	2,800
PowerEdge R740 (Intel Xeon Platinum 8164, 2.00 GHz)	52	2	26	140,302	2,698
PowerEdge R740 (Intel Xeon Platinum 8168, 2.70GHz)	48	2	24	153,326	3,194
PowerEdge R740 (Intel Xeon Platinum 8180, 2.50 GHz)	56	2	28	175,230	3,129