

Dell EMC Ready Solution for SAP Designed for SAP Landscape Consolidation

With PowerEdge Servers, Dell Networking components, PowerMax storage, and Data Domain storage protection

January 2019

H17550

Design Guide

Abstract

This design guide provides best-practice guidelines for designing and sizing an infrastructure solution that consolidates and simplifies IT for SAP landscapes.

Dell EMC 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 © 2019 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 01/19 Design Guide H17550.

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**
- Solution components.....6**
- Solution architecture19**
- Data protection21**
- Infrastructure-level high availability22**
- Solution configuration best practices27**
- Sizing and performance.....36**
- Conclusion.....47**
- References.....48**
- Appendix: SAPS values for PowerEdge R840 and R740/R740xd.....49**

Introduction

Executive summary

The digital era offers exceptional opportunities for businesses, while also bringing formidable challenges. IT departments struggle to derive fresh insights from complex, isolated systems. SAP S/4HANA and BW/4HANA, as replacements for classic SAP Enterprise Resources Planning (ERP), SAP Business Warehouse (BW), and SAP Business Suite, address those concerns. SAP S/4HANA and BW/4HANA are enabled with real-time insights, machine learning, and AI, bringing corporate business data together with big data and the Internet of Things (IoT).

The challenge of enabling SAP applications to run in a digital economy requires more than a quick investment and switching the applications on and off. The majority of SAP customers need to maintain traditional SAP ERP and SAP BW systems running on storage-based databases, while transitioning to an SAP HANA in-memory database and making the ultimate move to SAP S/4HANA. While every organization's path to transforming IT for SAP is unique, most have some common objectives, starting with sustaining and invigorating classic SAP deployments while migrating to SAP HANA and S/4HANA.

The following figure depicts this migraton path:

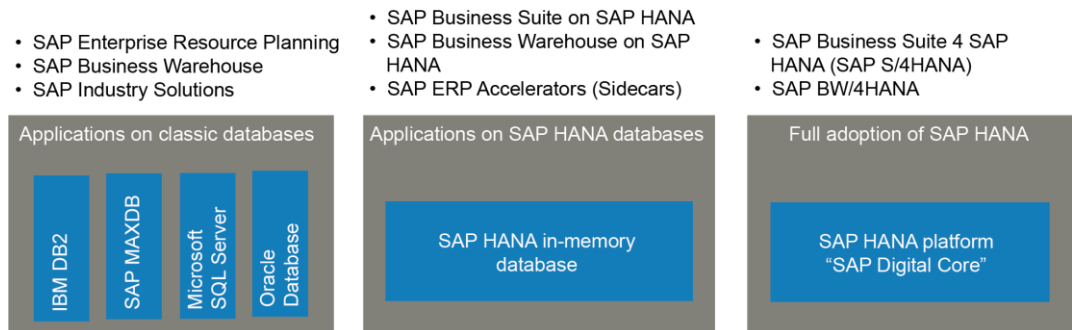


Figure 1. Transforming IT for SAP

The purpose of this Ready Solution is to consolidate and simplify IT and to lower the total cost of ownership (TCO), while delivering increased performance and productivity by running traditional SAP ERP landscapes on infrastructure that is SAP HANA-certified.

Solution overview

This Dell EMC Ready Solution for SAP designed for SAP landscape consolidation incorporates the following Dell EMC products: PowerEdge R840 and R740/R740xd servers, Dell Networking, Connectrix Fibre Channel (FC) switches, the PowerMax family of storage arrays, and Data Domain storage protection systems. The solution encompasses a variety of design configurations and deployment options. Customers use sizing tools for SAP systems on Dell EMC infrastructure to determine the requirements of the deployment and work with Dell EMC representatives to configure and deploy the solution.

Key benefits

PowerEdge servers that are paired with the PowerMax storage array and Data Domain storage protection power this Ready Solution. Customers implementing the solution can expect the following benefits:

- **Investment protection**—Complex traditional SAP landscapes are consolidated on infrastructure that is ready for SAP HANA.
- **Agility**—A modern SAP landscape management experience delivers automated provisioning capabilities for SAP applications and shortens the time-to-value.
- **Engineering**—Compute, networking, and storage are integrated with the defined prerequisites, and dependencies have been tested to deliver a seamless solution experience.
- **Optimization**—Design and deployment guides highlight proven performance, automation, and resiliency best practices for SAP landscapes.

Document purpose

This design guide provides validated best practices for sizing and designing a virtualized SAP landscape using Dell EMC servers, storage, and networking products. Customers planning to use Dell EMC infrastructure for SAP systems determine the solution design according to their sizing requirements and work with Dell EMC representatives to configure and deploy the solution.

This Ready Solution focuses on traditional SAP landscapes while using SAP HANA-certified Dell EMC servers and storage. The solution provides a certified performance platform for SAP landscapes and helps you prepare and chart a clear migration path to SAP HANA and SAP S4/HANA.

Audience

This guide is for SAP Basis administrators, system administrators, storage administrators, and presales architects who design mission-critical SAP landscapes to use on Dell EMC infrastructures. Readers should have some knowledge of Dell EMC PowerEdge servers, storage, and networking products as well as SAP NetWeaver and VMware virtualization technologies.

We value your feedback

Dell EMC and the authors of this guide welcome your feedback on the solution and the solution documentation. Contact the Dell EMC Solutions team by [email](#) or provide your comments by completing our [documentation survey](#).

Authors: Donagh Keeshan, Jill Orofino, Alexander Almeida, James Salvadore, Kevin Locklear, Ritika Saini, James Whalen, Aighne Kearney

Note: The following page of the Everything SAP at Dell EMC space on the Dell EMC Communities website provides links to additional documentation for this solution:
[SAP Info Hub for Ready Solutions](#)

Solution components

Overview

SAP is one of the largest vendors of ERP applications in the world. SAP landscapes are the application environments that manage the progress of changes such as patches, upgrades, new features, and reports from initial staging to a production implementation. SAP landscapes include development (DEV), quality assurance (QAS), and production (PRD) systems. This Dell EMC Ready Solution Designed for SAP Landscape Consolidation encompasses a variety of design configurations and deployment options that are tailored to the requirements and preferences of individual organizations.

This Ready Solution comprises:

- [Dell EMC components](#)
- [VMware components](#)
- [SAP landscapes](#)

Dell EMC components

Dell EMC offers a wide range of enterprise products and solutions, including servers, storage, networking, data protection, and services. These products and solutions are designed for reliability and scalability and are engineered to handle highly demanding business applications and workloads.

PowerEdge servers

PowerEdge servers, which are built to support the work of IT departments, include features to better run applications such as databases, ERP, business intelligence (BI), high performance computing (HPC), and data warehousing. Customers running SAP landscapes have the benefit of knowing that this Ready Solution includes high-performing servers, as documented in the [SAP Standard Application Benchmarks](#). At the time of publication of this guide, the PowerEdge servers with four 28-core Intel Xeon Platinum 8180 processors have achieved the best-of-published four-socket SAP Sales and Distribution (SD) two-tier benchmark results.

PowerEdge R840 server

The PowerEdge R840 server delivers unprecedented performance density for in-database analytics. Exceptional reliability and integrated security ensure maximum uptime in the enterprise data center. Automated intelligence simplifies IT management so that the focus is on key business results.

The R840 server with four 28-core Intel Xeon Platinum 8180 processors achieved an SAP Application Performance Standard (SAPS) total of 359,620. This result was published on April 30, 2018: [Certification #2018015](#).

Based on these SAP SD benchmark results, for a single SAP system, the PowerEdge R840 server offers the following scalability:

- In a two-tier SAP architecture on a single server: 359,620 SAPS
- In a three-tier SAP architecture with five servers, assuming for OLTP a 4:1 ratio of application SAPS to database SAPS (four application servers and one database server): 1,798,100 SAPS

For more information about SAPS values, see [SAPS overview](#) on page 37.

The following table shows the technical specifications for R840 servers.

Table 1. Technical specifications for PowerEdge R840 servers

Feature	Technical specifications
Processor	Up to 4 Intel Xeon SPs, up to 28 cores
Memory	48 DDR4 DIMM slots; supports RDIMM /LRDIMM, up to 2,666 MT/s, 6 TB maximum Up to 12 NVDIMM, 192 GB maximum Supports registered ECC DDR4 DIMMs only
Acceleration	Supports up to 2 DW FH GPUs or 2 FPGAs
Storage controllers	Internal controllers: PERC S140, H330, H740P Boot Optimized Storage Solution (BOSS) subsystems, HBA330 External HBAs (RAID): H840 External HBAs (non-RAID): 12 Gb/s SAS HBA
Drive bays	Front drive bays: Up to 24 x 2.5 in. SAS/SATA (HDD/SSD) with up to 12 NVMe PCIe SSD, maximum 122 TB 24 x 2.5 in. NVMe drives, maximum 153 TB Rear drive bays: 2 x 2.5 in. SAS/SATA option available
Power supplies	Platinum: 750W, 1100W, 1600W, 2000W, 2400W Titanium: 750W DC: 1100W Mix Mode/HVDC: 750W, 1100W Hot plug power supplies with full redundancy option Hot plug fans with N +1 redundancy
Dimensions	Form factor: Rack (2U) Chassis depth: 812 mm
I/O and ports	Network daughter card options: 4 x 1 Gb 4 x 10 Gb 2 x 10 Gb + 2 x 1 Gb or 2 x 25 Gb Front ports: Video, 2 x USB 2.0, 1 dedicated iDRAC Direct Micro-USB (optional: 1 x USB 3.0 offered on the 8-drive configuration) Rear ports: Video, serial, 2 x USB 3.0 Video card: VGA Up to 6 Gen3 slots (4 x 16 slots or 2 x 16 + 4 x 8 slots)
Supported operating systems ¹	Canonical Ubuntu LTS Citrix XenServer Microsoft Windows Server with Hyper-V Red Hat Enterprise Linux (RHEL) SUSE Linux Enterprise Server VMware ESXi Specifications and interoperability details at: Server & Storage Operating System Support

¹Only ESXi, SUSE, and RHEL are SAP HANA-certified operating systems

The PowerEdge R840 server's scalable business architecture can process the most business-critical workloads. With up to 6 TB of memory and 24 NVMe front storage combined with 6 PCIe Gen3 slots, the R840 server has all the resources to maximize application performance and scale for future demands. The R840 server:

- Maximizes storage performance with up to 24 NVMe drives and ensures that application performance scales easily
- Frees up storage space using internal, boot-optimized M.2 solid-state drives (SSDs)
- Eliminates bottlenecks with up to 6 TB of memory in 48 DIMMS, 12 of which can be NVDIMMs
- Minimizes latency through direct-attached NVMe configurations
- Accelerates application performance with GPUs

PowerEdge R740/R740xd server

The PowerEdge R740/R740xd two-socket server is a general-purpose platform with highly expandable memory (up to 3 TB) and I/O capability to match. The R740 can handle demanding workloads and applications such as data warehouse, databases, e-commerce, and HPC.

The R740 server with two 28-core Intel Xeon Platinum 8180 processors has achieved a total of 175,230 SAPS. This result was published on July 11th, 2017: [Certification #2017017](#).

Based on these SAP SD benchmark results, for a single SAP system, the PowerEdge R740/740xd two-socket server offers the following scalability:

- In a two-tier SAP architecture on a single server: 175,230 SAPS
- In a three-tier SAP architecture with five servers, assuming for OLTP a 4:1 ratio of application SAPS to database SAPS (four application servers and one database server): 876,150 SAPS

For more information about SAPS values, see [SAPS overview](#) on page 37.

The following table shows the technical specifications for these servers.

Table 2. PowerEdge R740 and R740xd technical specifications

Feature	Technical specifications
Processor	Up to two Intel Xeon SPs, up to 28 cores per processor
Memory	<ul style="list-style-type: none"> • 24 DDR4 DIMM slots; supports RDIMM/LRDIMM, up to 2,666 MT/s, 3 TB maximum • Up to 12 NVDIMM, 192 GB maximum • Supports registered ECC DDR4 DIMMs only

Feature	Technical specifications
Storage controllers	<ul style="list-style-type: none"> Internal controllers: PERC H730p, H740p, HBA330, software RAID (SWRAID) S140 External PERC (RAID): H840 External HBAs (non-RAID): 12 Gb/s SAS HBAs Boot Optimized Storage Solution (BOSS) subsystem: HWRAID 2 x M.2 SSDs 120 GB or 240 GB
Drive bays	<ul style="list-style-type: none"> Front bays: Up to 24 x 2.5 SAS/SATA (HDD/SSD), NVMe SSD maximum 153 TB or up to 12 x 3.5 in, SAS/SATA HDD maximum 120 TB Mid bay: Up to 4 x 2.5in, maximum 25 TB SAS/SATA (HDD/SSD), NVMe SSD maximum 25 TB or up to 4 x 3.5 in, maximum 40 TB Rear bays: Up to 4 x 2.5 in, maximum 25 TB SAS/SATA (HDD/SSD), NVMe SSD or up to 2 x 3.5 in, maximum 20 TB SAS/SATA HDD
Power supplies	<ul style="list-style-type: none"> Titanium: 750 W Platinum: 495 W, 750 W, 1,100 W, 1,600 W, and 2,000 W 48 VDC 1,100 W, 380 HVDC 1,100 W, 240 HVDC 750 W Hot-plug power supplies with full redundancy Up to 6 hot-plug fans with full redundancy
Dimensions	<ul style="list-style-type: none"> Form factor: Rack (2U) Max depth: 715.5 mm
I/O and ports	<ul style="list-style-type: none"> Network daughter card options: 4 x 1 GbE or 2 x 10 GbE + 2 x 1 GbE or 4 x 10 GbE or 2 x 25 GbE Front ports: VGA, 2 x USB 2.0, dedicated integrated Dell Remote Access Controller (iDRAC) Direct Micro-USB Rear ports: VGA, Serial, 2 x USB 3.0, dedicated iDRAC network port Video card: VGA Riser options with up to 8 PCIe Gen3 slots, maximum of 4 x 16 slots
Supported operating systems	<ul style="list-style-type: none"> Canonical Ubuntu LTS Citrix XenServer Microsoft Windows Server with Hyper-V Red Hat Enterprise Linux SUSE Linux Enterprise Server VMware ESXi Specifications and interoperability details: Server & Storage Operating System Support

Dell EMC PowerMax storage overview

The Dell EMC PowerMax product family consists of two models: the PowerMax 2000 array and the PowerMax 8000 array. The PowerMax 2000 array provides customers with efficiency and maximum flexibility in a 20U stack height, while the PowerMax 8000 array is designed for massive scale, performance, and IOPS density within a two-floor-tile

footprint. Both PowerMax models are based on the Dell EMC Dynamic Virtual Matrix architecture and a version of HYPERMAX OS that was rewritten for the NVMe platform. This HYPERMAX OS version is called PowerMaxOS.

Like the VMAX All Flash arrays, PowerMax arrays are specifically designed to meet the storage capacity and performance requirements of an all-flash enterprise data center. PowerMax products take advantage of the higher-capacity NVMe flash drives that are used in the densest configuration possible. PowerMax technology offers enterprise customers trusted data services along with the simplicity, capacity, and performance that their highly virtualized environments demand, while still meeting the economic needs of traditional storage workloads. Furthermore, the PowerMax platform enables you to deploy applications such as real-time analytics, machine learning, and big data, which require lower storage latency and higher IOPS densities than what was attainable with previous all-flash offerings.

The following figure shows the essential features of the PowerMax 2000 and 8000 arrays.



PowerMax 2000

- 1.7M IOPS (8K RRH)
- 64 FC/iSCSI Ports
- 1 -2 Bricks per system
- Up to 1 PBe per system of NVMe storage
- All in 1/2 Rack (20U)
- Open Systems Workloads Only

PowerMax 8000

- 10M IOPS (8K RRH)
- 256 FC/FICON/iSCSI Ports
- 1 -8 Bricks per system
- Up to 4 PBe per system of NVMe storage
- All in 2 Racks (2 floor tiles)
- Open Systems / Mainframe / Mixed Workloads

Figure 2. PowerMax 2000 and PowerMax 8000 system features

The main benefits that PowerMax platforms offer customers are:

- An end-to-end NVMe storage design providing industry-leading IOPS density in a system with a small footprint
- Readiness for next-generation data storage media, such as Storage Class Memory (SCM) and NVMe over Fabric (NVMe-oF) infrastructure
- Efficient workload consolidation for block, file, and mainframe workloads on a single platform
- Ultrahigh reliability, availability, and serviceability for the enterprise
- Excellent data services, including advanced data reduction using inline deduplication and compression

System specifications

The following table shows the system specifications for PowerMax 2000 and PowerMax 8000 arrays.

Table 3. PowerMax system specifications

PowerMax model	Bricks per system	Brick CPU	Brick cache	Maximum cache	Maximum FE modules per Brick	Maximum number of drives	Maximum capacity
2000	1-2	12 core, 2.5 GHz Intel Broadwell (2 per director) Total: 48 cores per Brick	512 GB, 1 TB, and 2 TB (DDR4)	4 TB	8 (32 total FE ports per Brick) Total: 64 FE ports per system	96	1 PBe
8000	1-8	18 core, 2.8 GHz Intel Broadwell (2 per director) Total: 72 cores per Brick	1 TB and 2 TB (DDR4)	16 TB	6 (24 total FE ports per Brick – open systems / mixed) Total: 256 FE ports per system	288	4 PBe

Expandable modular architecture

PowerMax configurations consist of modular building blocks called Bricks. The modular Brick architecture reduces complexity and enables easier system configuration and deployment. This architecture also enables the system to scale while continuing to deliver predictable high performance.

Two types of Brick blocks are available for PowerMax configurations:

- The open systems Brick block supports configurations with FC or iSCSI connectivity with fixed-block architecture (FBA) device formatting. You can also configure the Brick block for file storage by using embedded NAS.
- The mainframe zBrick block supports configurations with fiber connectivity (FICON) and count key data (CKD) device formatting.

The initial Brick block includes a single engine consisting of two directors, two system power supplies (SPS), and two 24-slot 2.5 in. NVMe Drive Array Enclosures (DAE24) preconfigured with an initial total usable capacity. The Brick 2000 block comes with an initial capacity of 11 or 13 TBu, depending on the RAID configuration. The PowerMax 8000 array model comes with an initial capacity of 53 TBu for open systems, 13 TBu for mainframe, and 66 TBu for mixed systems. Each Brick block comes preloaded with PowerMaxOS software.

The Brick concept makes it possible for the PowerMax platform to scale up and out. Customers can scale up by adding flash capacity packs, which include NVMe flash drive capacity that can be added to a PowerMax array. Each flash capacity pack for the

PowerMax 8000 system has 13 TBu of usable storage. For the PowerMax 2000 model, each pack has 11 TBu or 13 TBu, depending on the RAID protection type selected. Scale out a PowerMax system by aggregating up to two Brick blocks for the PowerMax 2000 array and up to eight blocks for the PowerMax 8000 array in a single system with fully shared connectivity, processing, and capacity resources. Scaling out a PowerMax system by adding additional Brick blocks produces a predictable, linear performance improvement regardless of the workload.

The PowerMax provides multiple front-end connections that implement several protocols and speeds. The following table shows the various front-end connectivity modules that are available for a PowerMax system.

Table 4. Connectivity modules for PowerMax systems

Connectivity type	Module type	Number of ports	Mix with protocols	Supported speeds (Gbps)
FC	16 Gbps FC	4	SRDF	2/8/16
SRDF	10 GigE	4	iSCSI	10
iSCSI	10 GigE	4	SRDF	10
eNAS	10 GigE	2	None	10
FICON ¹	16 Gbps FICON	4	Single/Multi Mode	4/8/16
eNAS	10 GigE (Copper)	2	None	10
eNAS Tape Backup	8 Gbps FC	4	None	2/4/8

¹ Supported on the PowerMax 8000 array only

PowerMax software options

All PowerMax arrays come preloaded with PowerMaxOS software. Software for the PowerMax system is available in different packages, each with additional optional features. For up-to-date information about these options, see the [Dell EMC PowerMax Family Overview White Paper](#).

Non-Volatile Memory Express (NVMe)

PowerMax arrays provide a full NVMe flash storage backend for storing customer data.

The PowerMax NVMe architecture provides:

- **I/O density with predictable performance**—PowerMax arrays are designed to deliver industry-leading I/O density. They are capable of delivering over 10 million IOPS in a two-rack system (two floor tiles), regardless of workload and storage capacity utilization.
- **NVMe storage density**—PowerMax technology delivers industry-leading NVMe TB per floor tile. While other all-flash alternatives use a proprietary flash drive design, PowerMax systems support high-capacity, commercially available, dual-ported enterprise NVMe flash drives and can therefore leverage increases in flash drive densities, economies of scale, and reduced time-to market.

- **Future-proof design**—The PowerMax NVMe design is ready for Storage Class Memory (SCM) flash and future NVMe-oF SAN connectivity options, which include 32 Gb FC and high-bandwidth converged Ethernet: RoCEv2 over 25 GbE/50 GbE/100 GbE.

Smart RAID

PowerMax uses a new active/active RAID group accessing scheme called smart RAID. Smart RAID allows the sharing of RAID groups across directors, giving each director active access to all drives on the Brick or zBrick block. Both directors on an engine can drive I/O to all the flash drives, creating balanced configurations in the system whatever the number of RAID groups and providing performance benefits.

Smart RAID also offers increased flexibility and efficiency. Customers can order PowerMax systems with a single RAID group, allowing for at least nine drives per engine with RAID 5 (7+1) or RAID 6 (6+2 and 1 spare) and as few as five drives per system for a PowerMax 2000 with RAID 5 (3+1 and 1 spare). These configurations leave more drive slots available for capacity upgrades in the future. When the system is scaled up, customers have more flexibility because flash capacity pack increments can be a single RAID group.

Remote replication with SRDF

The Symmetrix Remote Data Facility (SRDF) is considered a gold standard for remote replication in the enterprise data center. Up to 70 percent of Fortune 500 companies use SRDF to replicate their critical data to geographically dispersed data centers throughout the world. SRDF offers customers the ability to replicate tens of thousands of volumes, with each volume being replicated to a maximum of four different locations globally.

SRDF is available in three types:

- **SRDF Synchronous (SRDF/S)**—SRDF/S delivers zero data loss remote mirroring between data centers separated by up to 60 miles (100 km).
- **SRDF Asynchronous (SRDF/A)**—SRDF/A delivers asynchronous remote data replication between data centers up to 8,000 miles (12,875 km) apart. SRDF/S and SRDF/A can be used together to support three or four site topologies as required by the most mission-critical applications.
- **SRDF/Metro**—SRDF/Metro delivers active-active high availability for non-stop data access and workload mobility within a data center, or between data centers separated by up to 60 miles (100 km).

Local replication with TimeFinder SnapVX

Every PowerMax array includes the local replication data service TimeFinder SnapVX, which creates low-impact snapshots. Local replication with SnapVX starts by creating a snapshot, a pointer-based structure that preserves a point-in-time (PIT) view of a source volume. Snapshots do not require target volumes. Instead, they share back-end allocations with the source volume and other snapshots of the source volume, and only consume additional space when the source volume is changed. A single source volume can have up to 256 snapshots.

You can access a PIT snapshot from a host by linking it to a host-accessible volume called a target. Target volumes are standard thin volumes. Up to 1,024 target volumes can be linked to the snapshots of a single source volume. By default, targets are linked in a no-copy mode. This functionality significantly reduces the number of writes to the back-

end flash drives. It eliminates the need to perform a full-volume copy of the source volume during the unlink operation to continue using the target volume for host I/O.

Advanced data reduction using inline compression and deduplication

PowerMax arrays use the Adaptive Compression Engine (ACE) for inline hardware compression. The ACE data-reduction method provides a negligible performance impact with the highest space-saving capability. PowerMax technology uses inline hardware-based data deduplication. This method identifies repeated data patterns on the array and stores those that are repeated once only, thus preventing the consumption of critical PowerMax system core resources and limiting performance impact.

Embedded NAS

The embedded NAS (eNAS) data service extends the value of PowerMax to file storage. eNas enables customers to use vital enterprise features such as flash-level performance for both block and file storage, as well as to simplify management and reduce deployment costs.

eNAS uses the hypervisor in PowerMaxOS to create and run a set of virtual machines (VMs) within the PowerMax array. These VMs host two major elements of eNAS software: data movers and control stations. The embedded data movers and control stations have access to shared system resource pools so that they can evenly consume PowerMax resources for both performance and capacity.

With the eNAS data service, PowerMax becomes a unified block-and-file platform that uses a multicontroller, transactional NAS solution.

For more detailed information about the PowerMax product family, see the [Dell EMC PowerMax Family Overview White Paper](#).

PowerMax and SAP HANA TDI

SAP has certified Dell EMC PowerMax storage arrays as meeting all performance and functional requirements for SAP HANA TDI deployments. Therefore, customers can use PowerMax arrays for SAP HANA TDI deployments in a fully supported environment using their existing data center infrastructures.

Dell EMC Engineering performed tests on a PowerMax 2000 single engine using the SAP hwcct tool for HANA-HWC-ES 1.1 certification. Based on the test results, the following table provides scalability guidelines for the maximum possible number of SAP HANA production hosts that can be connected to a given PowerMax array:

Table 5. PowerMAX All Flash scalability

PowerMax model	Number of PowerMax Bricks	Number of SAP HANA worker hosts
2000	1	16
	2	24
8000	1	26
	2	42

PowerMax model	Number of PowerMax Bricks	Number of SAP HANA worker hosts
	3	62
	4	82
	5	102
	6	122
	7	142
	8	162

For more information, see the [Storage Configuration Best Practices for SAP HANA TDI on Dell EMC PowerMax Arrays Validation Guide](#).

Dell EMC Connectrix SAN switches

The following Connectrix Fibre Channel (FC) SAN switches provide cross-sectional bandwidth:

- **Connectrix DS-6610B**—A switch with a base configuration of eight ports with 16 GB SFP+ (upgradable to 32 GB by replacing the SFP+). You can expand the configuration up to 24 ports by applying two Ports On Demand (POD) upgrade licenses of eight ports each. The Connectrix DS-6610B is an entry-level, highly cost-effective FC SAN switch to configure with SAP landscapes.
- **Connectrix DS-6620B**—A 32 GB-capable enterprise switch with 24 ports at 32 GB in the base. You can expand it to 48 ports by applying two POD upgrade licenses of 12 ports each. This switch can include two 128 GB ports that can also be mapped to 16 x 32 GB ports using breakout quad small form-factor pluggable (QSFP).

For enterprise-scale FC SAN solutions, the Connectrix ED-DCX6 directors are available in both four-slot and eight-slot chassis models, providing the foundation for FC SANs that scale to 9,000 ports.

The Connectrix B Series products provide features for validation, monitoring, and alerting as well as remediation of the storage network infrastructure. These features are collectively referred to as Fabric Vision. Fabric Vision technology combines built-in capabilities in Connectrix Gen 5 and Gen 6 platforms, Fabric OS, and Connectrix Manager features.

For information about how to configure Fabric Vision for your SAP landscape deployment, see the [Dell EMC Ready Bundle for SAP Deployment Guide](#).

Dell EMC Networking SAN switch

The Dell EMC Networking S4148U-A switch utilizing OS10 Enterprise Edition with 48 ports provides FC/ISCI connectivity at a top-of-rack (ToR) scale:

- 24 SFP+ ports of 10 Gb Ethernet connectivity
- 24 SFP+/SFP28 unified ports providing 10 GB Ethernet or FC8/FC16 connectivity
- 2 QSFP+ ports of 40 Gb connectivity
- 4 QSFP28 unified ports providing 10 Gb, 25 Gb, 40 Gb, 50 Gb, 100 Gb or FC8, FC16, or FC32 connectivity

Dell EMC data protection appliances and software

Dell EMC data protection purpose-built appliances (PBBAs, our Data Domain series) and software solutions can be used separately or together to provide data protection and storage protection for large enterprise workloads. Using the protection appliance system and software together enables you to more effectively consolidate your backup, archive, and disaster recovery operations. With the protection appliance system, you can protect more data faster to meet aggressive SLAs and reduce the risk of data loss. With the data protection software and its built-in integration with SAP, you can enable DBA self-service via SAP BR*Tools and make cost savings from the implementation of direct data paths to Data Domain.

The Data Domain 9300 data protection storage system works seamlessly with a range of backup, archive, and enterprise applications. The system is configurable for high availability—you can add a standby or a second controller as needed.

Key benefits of the Data Domain 9300 system include:

- Throughput of 20 TB per hour increasing up to 41 TB per hour using Dell EMC DD Boost software, which means the 9300 system can back up 328 TB in under 8 hours.
- Support for 810 concurrent backup streams.
- Up to 36 PB local logical capacity. This capacity expands to 108 PB with optional Data Domain Cloud Tier software, supporting long-term data retention.

For more information, see the following documents:

- [*Dell EMC Data Domain Deduplication Storage Systems Spec Sheet*](#)
- [*Dell EMC Data Domain Boost for Enterprise Applications and ProtectPoint Application Agent Installation and Administration Guide*](#)

VMware components

VMware vSphere

VMware vSphere uses virtualization to transform individual data centers into aggregated computing infrastructures that include CPU, storage, and networking resources. VMware vSphere manages these infrastructures as a unified operating environment and gives you the tools to administer the data centers that participate in that environment.

vSphere support for existing and next-generation applications includes:

- Simplified customer experience for automation and management at scale
- Comprehensive built-in security for protecting data, infrastructure, and access
- Universal application platform for running any application anywhere

With vSphere you can run, manage, connect, and secure your applications in a common operating environment across clouds and devices.

For additional and relevant information about vSphere, see the following documentation:

- SAP applications with vSphere: [SAP on VMware vSphere](#)

- SAP HANA and SAP HANA-related applications with vSphere: [SAP HANA on VMware vSphere](#)
- SAP HANA in virtualized environments on vSphere with Dell EMC storage: [VMware Virtualized SAP HANA with Dell EMC Storage Deployment Guide](#)

SAP landscapes

SAP landscapes support the implementation of an SAP system consisting of isolated and stable environments that provide a solid "promote to production" change-management and control process for all configuration and development. SAP recommends this scenario for all SAP applications and any system that is based on SAP NetWeaver.

The typical SAP environment is a three-system landscape with one DEV system, one QAS system, and one PRD system. Many customers supplement a three-system landscape with additional environments, such as a sandbox (SBX) that is used for testing, learning, and destructive testing.

Note: A three-system landscape (DEV, QAS, PRD) is different to a three-tier SAP architecture system consisting of a database layer, an application layer, and a presentation layer.

The following figure shows a typical SAP system landscape for a single application.




Figure 3. SAP system landscape

Most SAP customers have more than one SAP system landscape in their total landscape environment, including applications such as ERP, BW, customer relationship management (CRM), and supply chain management (SCM). If you replicate a three-system landscape approach for every system that you deploy in a landscape, the number of individual systems quickly inflates, requiring additional compute and storage resources. Managing the TCO and the expanding total landscape is a common challenge for all customers who have more than one system landscape.

Solution architecture

Overview

This section describes the design principles of this Ready Solution and the design of the infrastructure subsystems.

The following figure shows the solution architecture.

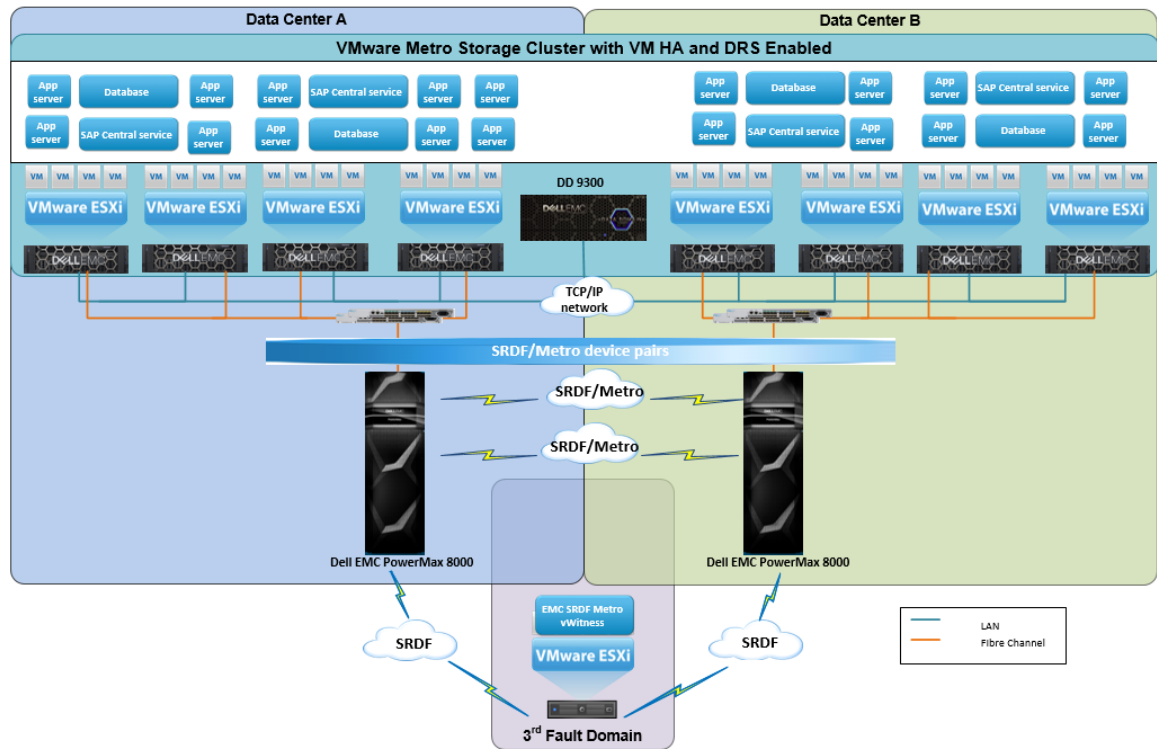


Figure 4. Solution architecture

Server architecture

In a data center environment, consolidating multiple components to reduce the data center footprint and cost is important. Published levels of best SAPS performance enable customers using PowerEdge servers to achieve greater server consolidation ratios because fewer servers are required to achieve the application or database SAPS requirements of an SAP customer.

Server virtualization using VMware vSphere software enables optimal resource utilization. This solution uses server virtualization to make it possible to run multiple operating systems and SAP applications on the same server simultaneously, increasing consolidation, flexibility, and scalability while creating significant cost savings.

Network architecture

This Ready Solution provides end-to-end I/O connectivity. It employs two Emulex LightPulse LPe32000 16 GB FC adapters per host for FC connectivity. One FC port per HBA is connected to a pair of redundant Connectrix DS-6610B FC switches to provide high availability.

Note: The solution uses Connectrix 16 GB- and 32 GB-capable switches for FC SAN connectivity. Switches are configured to provide redundant A and B fabrics. For the highest availability, Fabric Vision is enabled to monitor the FC SAN health to detect and remediate network congestion that results from devices behaving incorrectly and other network issues.

Key differentiators of Dell EMC Networking data center switch solutions include:

- **Open networking offerings**—More choice with a broad offering of Dell EMC and third-party operating system solutions from partners such as Big Switch Networks, Cumulus Networks, IP Infusion, and Pluribus Networks.
- **Cost-effective port connectivity**—Dense port counts in 1RU and 2RU form factors with 24 to 128 ports ranging from 1 G, 10 G, 25 G, 40 G, and 50 G to 100 G.
- **Large scale operations**—A leaf-spine 100 G fabric architecture that is massively scalable and easily automated at scale.

The Dell EMC Networking S5200-ON family, which comprises the S5248, S5296, and S5232 switches, provides multi-line rate (100/25/10 GbE) network connectivity between the compute cluster and the rest of the data center. For more information, see the [Dell EMC Networking S5200-ON Series Switches Specification Sheet](#).

The Dell EMC S4100-ON family, comprising the S4112, S4124, and S4148 switches, can provide 10/1 GbE connectivity in either SFP(+) configuration or in 10Gbase-T configuration. The 10Gbase-T configuration allows for a 10/1 GbE Cat 6 or 7 connection for data or management. For more information, see the [Specification Sheet](#).

The Dell EMC S3048-ON switch provides 1 GbE connectivity. For more information, see the [Dell EMC Networking S3048-ON Specification Sheet](#).

To provide high availability for network connections, multiple network cards in each server are connected to redundant network switches. Two 1 GB NICs connect to a pair of redundant S4100-ON series or S3048-ON switches, while two 10 or 25 GB NICs also connect to redundant S5200F-ON series switches. In accordance with SAP best practices and Dell EMC infrastructure design principles, each application is deployed as a separate workload VLAN. The VLAN is defined in the data center core network—for example, as DEV, QAS, or PRD.

Storage architecture

This Ready Solution uses the PowerMax system with the following profile:

- NVMe flash drives
- RAID 5 3+1, 7+1 (PowerMax 2000) , RAID 5 7+1 (PowerMax 8000)
- Data reduction enabled
- Remote replication using SRDF/Metro

Data protection

Backup, recovery, and archiving

Whether your company is a small-to-midsize or a large enterprise, the smooth operation of your SAP environment is crucial to your daily business and long-term competitiveness. Application performance, availability, and agility are critical. As SAP landscapes become more complex, how you protect your SAP data becomes increasingly important. Using the wrong data protection system slows down operations and limits growth. This section describes how the demands created by cloud or big data projects require a new way of providing SAP data protection.

Redefining SAP protection

Today, SAP protection includes:

- Continually maximizing SAP performance and availability
- Improving the visibility and control of SAP data protection
- Swiftly and adeptly responding to changing business SLAs
- Empowering database administrators and app owners to use their own native tools for backup and recovery

Meeting customers' business challenges involves a consolidated approach to data protection using a protection storage architecture. Dell EMC Data Domain systems enable you to redefine SAP backup, archive, and availability with deduplication and consolidated data protection. Data Domain systems work seamlessly with a range of backup, archive, and enterprise applications.

Proven results with Data Domain and Data Protection Suite for Applications

Using Data Domain systems in SAP environments lets you:

- Consolidate backup and archive workloads, eliminating "one-off" protection architectures
- Improve visibility and control of data protection, and enable application owners and DBAs to back up and restore SAP environments
- Reduce your protection footprint by employing deduplication and data reduction across your SAP landscape
- Simplify system management
- Protect multiple database copies efficiently and reliably
- Provide network-efficient replication and reliable recovery
- Lower TCO
- Complement SAP HANA's protection capabilities

When you pair Data Domain with Data Protection Suite for Applications, you can:

- Take advantage of direct data paths to Data Domain that offer faster backup and recovery compared to traditional approaches
- Protect SAP without relying on traditional backup infrastructure
- Reduce or even remove impact on SAP application servers

Infrastructure-level high availability

Overview

SAP applications are responsible for most business processes and unplanned outages can represent a substantial loss in revenue. Customers require solutions with:

- Resiliency, enabling them to recover from temporary failures through automated services restarts or corrections
- Agility, which requires highly available servers, network paths, switches, and storage
- Redundancy, which must be built into any solution supporting business-critical applications

To ensure high availability and a minimal component failure rate, we recommend using components that have a high mean time between failure (MTBF).

This Ready Solution combines a redundant design, RAID for storage drives, dual controllers, and multiple paths to provide storage high availability. Multiple FC HBAs and switches connect to the solution infrastructure and the storage network, building resiliency in network connectivity and enabling systems to remain online and operational during component repair.

Because the solution is also virtualized, high availability is built in at the vSphere level to support both planned and unplanned downtime.

vSphere vMotion enables live migration of virtual machines (VMs) from ESXi server to ESXi server so that you can perform planned maintenance with zero application downtime. In addition, vSphere HA, which provides high availability for VMs, reduces unplanned downtime by using multiple vSphere ESXi hosts. The hosts are configured as a cluster to provide rapid recovery from outages and cost-effective high availability for applications that are running in VMs.

To ensure application availability, vSphere HA:

- Reacts to hardware failure and network disruptions by restarting VMs on active hosts within the cluster
- Detects operating system failures by continuously monitoring a VM and restarting it as required
- Provides a mechanism to react to application failures
- Provides the infrastructure to protect all workloads within the cluster, in contrast to other clustering solutions

Because high availability protects all workloads, installing additional software within the application or VM is not necessary. After configuration, any VMs that are added to a cluster are automatically protected with no further action required.

vSphere HA is often used with vSphere Fault Tolerance (FT). Fault tolerance enables a VM to survive a physical server failure by creating an exact replica VM on another host that can take over at the time of failure. Fault tolerance provides protection for extremely critical VMs, where any loss of service is intolerable.

You can combine high availability with vSphere Distributed Resource Scheduler (DRS) to protect against failures and provide load balancing across the hosts within a cluster.

The advantages of vSphere HA over traditional failover solutions include:

- Minimal setup
- Reduced complexity—no need for quorum disks
- Reduced hardware and configuration costs
- Increased application availability without the expense of additional idle failover hosts or the complexity of maintaining identical hosts for failover pairs
- DRS and vMotion integration
- Automatic or manual mode configuration to prevent unwanted live migrations

SRDF replication This Ready Solution uses Dell EMC Symmetrix Remote Data Facility (SRDF) software for disaster recovery and data mobility. SRDF software replicates data between two, three, or four arrays located in the same room, on the same campus, or thousands of kilometers apart. SRDF replication modes include:

- **SRDF synchronous (SRDF/S)**—SRDF/S replication maintains a real-time copy at arrays that are located within 200 kilometers. The local array acknowledges writes from the production host when they are written to cache at the remote array.
- **SRDF asynchronous (SRDF/A)**—SRDF/A replication maintains a dependent write-consistent copy at arrays at unlimited distances. The local array immediately acknowledges writes from the production host, and, therefore, replication has no impact on host performance. Data at the remote array is typically only seconds behind the primary site.

SRDF/Metro This solution focuses on SRDF/Metro, which significantly changes the behavior of SRDF/S mode with respect to remote device availability to better support host application in high availability environments. With SRDF/Metro, the secondary SRDF device (the R2) is also read/write-accessible to the host and takes on the federated personality of the primary device (the R1)—geometry, device WWN, and so on.

Because they provide a federated personality on the R2 device, both R1 and R2 devices can appear as a single virtual device across the two SRDF paired arrays for host presentation. The host (or hosts, in a cluster) can read and write to both devices. SRDF/Metro ensures that each copy remains current and consistent and addresses any write conflicts that might occur between the paired SRDF devices.

The following figure shows two SRDF/Metro configurations. The top half shows an SRDF/Metro configuration with a standalone host. The host has visibility to both the R1 and R2 arrays by using multipathing software such as Dell EMC PowerPath technology to enable parallel reads and writes to each array. The lower half of the graphic shows a clustered host environment, where each cluster node has dedicated access to an individual PowerMax array.

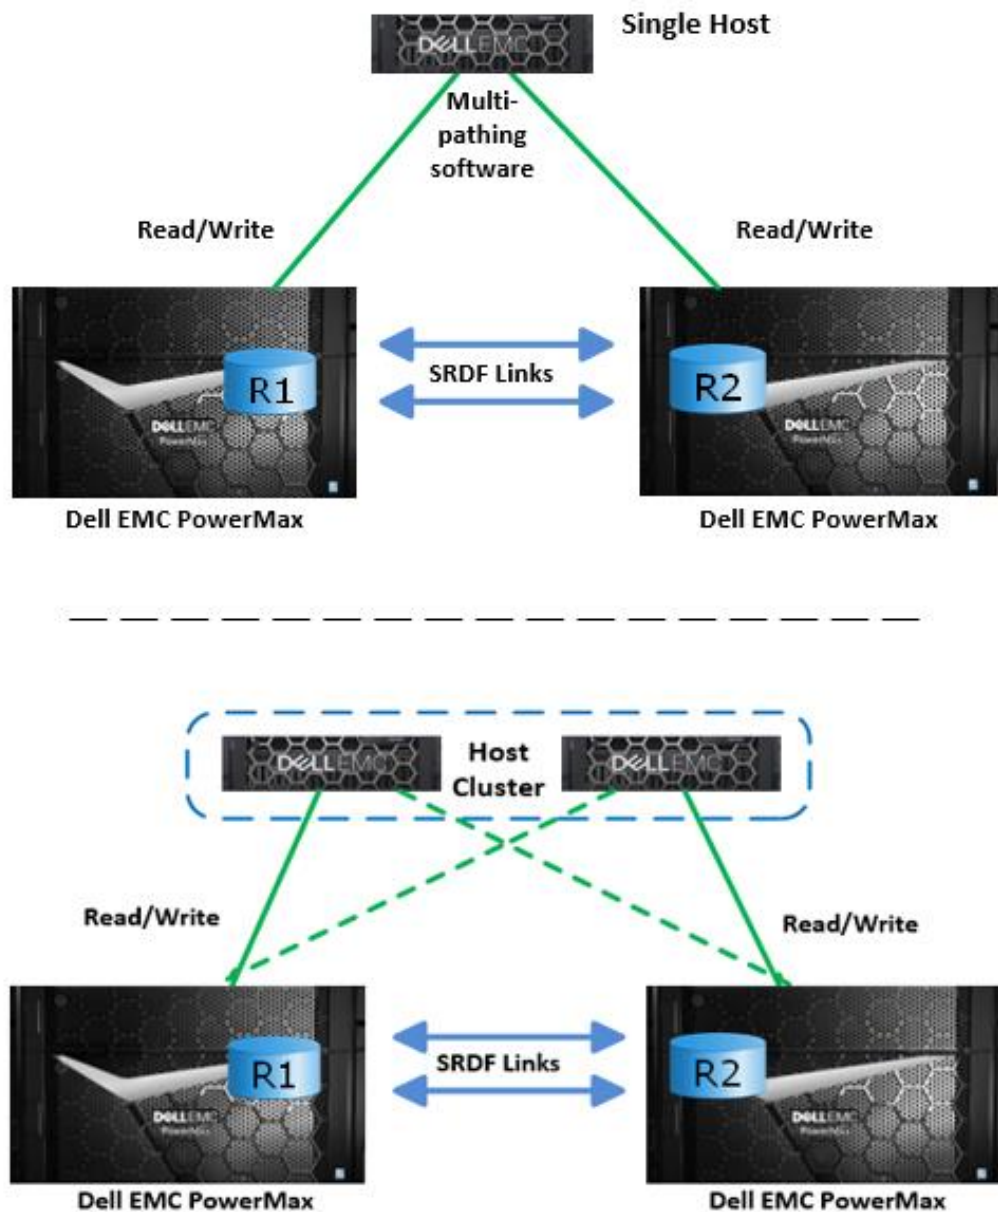


Figure 5. SRDF/Metro configurations: standalone and clustered hosts

In both the standalone host environment and the clustered host environment, writes to the R1 or R2 devices are synchronously copied to their SRDF paired devices. If a conflict occurs between writes to paired SRDF/Metro devices, the conflicts are internally resolved to ensure that a consistent image between paired SRDF devices is maintained to the individual host or host cluster.

You can select and manage SRDF/Metro through Dell EMC Solutions Enabler SYMCLI or Dell EMC Unisphere for PowerMax (Unisphere) client software. You must install a separate SRDF/Metro license on each of the PowerMax systems to be managed.

SRDF/Metro uses the SRDF link between the two sides of the SRDF device pair to ensure consistency of the data on the two sides. If the status of the SRDF device pair becomes Not Ready (NR) on the SRDF link, SRDF/Metro must respond. SRDF/Metro chooses one side of the SRDF device pair to remain accessible to the host or hosts, while making the other side of the SRDF device pair inaccessible. The bias and witness options enable this choice. Both these options prevent data inconsistencies between the two sides of the SRDF device pair.

For more information, see the following documentation:

- [The Dell EMC PowerMax Family Overview](#)
- [Dell EMC SRDF/Metro Overview and Best Practices](#)

Bias

Bias is a function of the two PowerMax systems in the SRDF/Metro configuration and is a required and integral component of the configuration. The Create Pair operation places an SRDF device pair into an SRDF/Metro configuration and preconfigures the bias to the R1 side of the pair. You can change the bias when all SRDF device pairs in the SRDF group have reached the active-active SRDF pair state.

If the Remote Data Facility (RDF) device pair becomes NR on the RDF link, the nonbias side is inaccessible to the hosts while the bias side remains accessible to them.

Witness

Witness is an optional physical or virtual component of SRDF/Metro. The physical witness requires a third PowerMax array or a PowerMax array with the correct code. It is enabled through a special witness SRDF group that is created on each SRDF/Metro array and paired with the third witness array.

This Ready Solution uses a virtual Witness (vWitness), a virtual appliance (vApp) running on an ESXi host. The vApp maintains multiple IP connections to redundant management guests on both the primary and secondary SRDF/Metro-managed arrays. These IP connections use TLS/SSL to ensure connectivity between the vWitness and SRDF/Metro paired arrays, as shown in the following figure. When configured, a witness system supersedes the bias functionality unless a failure situation requires specific knowledge of the bias system. In such cases, the system defaults to the bias functionality.

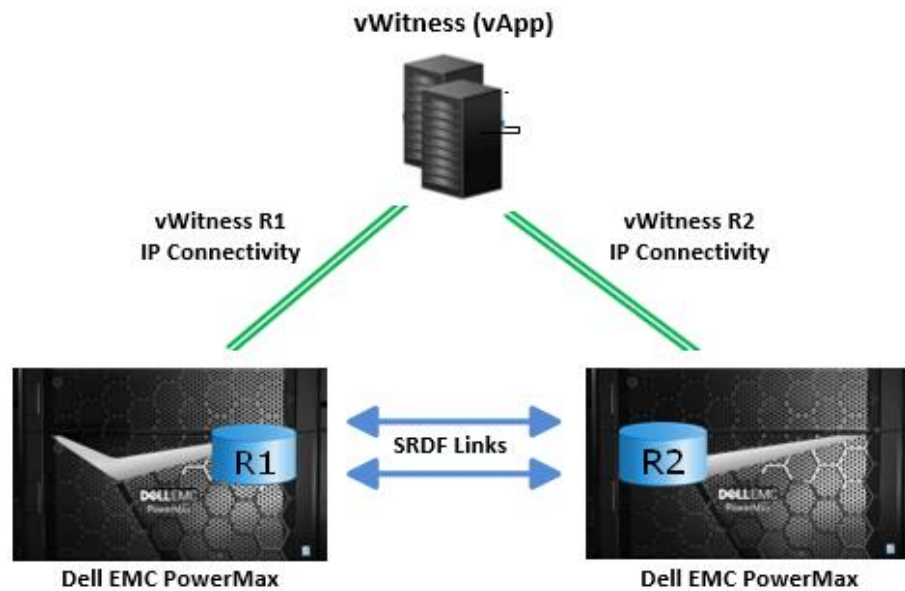


Figure 6. vWitness for SRDF Metro

VMware vSphere Metro Storage Cluster

A VMware vSphere Metro Storage Cluster (vMSC) configuration is a VMware-certified solution that combines replication with array-based clustering. These solutions are typically deployed in environments where the distance between data centers is limited—often metropolitan or campus environments.

A vMSC configuration requires what is in effect a single storage subsystem that spans both sites. In this design, a specified data store must be accessible (that is, be able to be read from and written to) simultaneously from both sites. In addition, when problems occur, the ESXi hosts must be able to continue to access data stores from either array transparently and with no impact to ongoing storage operations.

The storage subsystem for a vMSC must be able to write to and be read from both locations simultaneously. All disk writes are committed synchronously at both locations to ensure that data is always consistent, regardless of the location from which it is being read. This storage architecture requires significant bandwidth and very low latency between the sites in the cluster. Increased distances or latencies cause delays in writing to disk and a dramatic decline in performance. They also preclude successful vSphere vMotion migration between cluster nodes that reside in different locations.

VMware vMSC and SRDF/Metro

The VMware vMSC and SRDF/Metro environment that this design guide describes is a nonuniform configuration for vMSC with SRDF/Metro, meaning that ESXi hosts in each data center only see their respective PowerMax array. Because SRDF/Metro maintains the cache state on each array, an ESXi host in either data center detects the virtual device as local. Even when two VMs reside on the same data store but are located in different data centers, they write locally without any performance impact to either of them. Each host assesses and is served by its own array, with SRDF/Metro ensuring consistency.

For more information, see [VMware vSphere Metro Storage Cluster Recommended Practices](#).

Solution configuration best practices

Overview

This section describes the best practices that Dell EMC recommends for configuring the Ready Solution Designed for SAP Landscape Consolidation.

Server configuration

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

- Use the latest BIOS and firmware.
- Use the latest vendor-subscribed HBA drivers.
- Enable Intel Hyper-threading.

Local area network

Dell EMC recommends the following best practices for configuring the LAN:

- Use two 25/10 GbE S5200-ON series network switches for end-to-end client traffic to prevent a single point of failure (SPOF).
- Use a minimum of two network cards to provide high availability and prevent an SPOF. In our laboratory testing, we used qualified 25/10 GbE cards to enable connectivity between the SAP host and the external network.
- Use VLAN tagging to sustain switch failure and NIC teaming to provide multitier network high availability.
- Use NIC teaming to provide high availability at the NIC port level and for virtual and physical network connections.

Storage area network

Dell EMC recommends the following best practices for designing the storage area network (SAN):

- Use redundant Connectrix 32 Gb/s-capable switches or directors (DS-6610B, DS-6620B, ED-DCX6) to prevent a SPOF.
- Include the optional Enterprise Bundle with Fabric Vision for the DS-6610B and DS-6620B switches (Enterprise Bundle is included with the ED-DCX6 directors).
- Use two FC HBA cards for high availability.
- Enable high availability by making crisscross cable connections between redundant FC HBA cards and redundant FC switches.
- Use name zoning of end devices (FC HBAs) and PowerMax ports while creating zones. Name zoning of end devices provides better flexibility because it does not tie them to specific physical ports on either the switch or the storage array.
- In general, for best availability and performance, Dell EMC recommends that you zone or mask each server HBA port to one or two PowerMax ports. Where possible, ensure that the PowerMax ports are on different engines and directors and avoid creating a bandwidth mismatch between a host and the engine it accesses, such as a single 16 Gb/s FC host initiator accessing four 16 Gb/s FC ports on one engine. In this type of configuration, the host might simultaneously transmit I/O requests to all four storage ports and be overwhelmed by the response

from the targets, leading to spreading congestion and a potential negative impact on overall SAN performance.

- Use Brocade fabric performance impact (FPI) monitoring and their Monitoring and Alerting Policy Suite (MAPS) to monitor the SAN for slow-drain events. For more information, see the [Brocade Monitoring and Alerting Policy Suite Configuration Guide](#).

Storage configuration

When configuring Dell EMC storage, consider the following business needs:

- Performance and capacity, including future growth
- IOPS and MB/s performance
- Preferred server connectivity type

This section provides guidance and best practices for meeting such business needs.

NVMe flash drive capacity and RAID protection service-level objectives

A service-level objective (SLO) defines an expected average response-time target for a storage group. The following SLOs are available in the PowerMax storage array: the diamond SLO with defined, expected average response time targets, and the optimized SLO with no explicit response time target.

By default, the optimized SLO manages all the data in the PowerMax storage array. Because the PowerMax contains all NVMe flash devices, the optimized and the diamond settings are effectively the same. The actual response time of an application that is associated with each SLO varies based on the observed workload. Response time depends on average I/O size, read/write ratio, and the use of local or remote replication, along with the availability of other resources within the array. For more information, go to [Dell EMC Online Support](#) and search for Dell EMC FAST and VMAX documentation.

Volume provisioning and mapping

Thin provisioning enables you to increase capacity utilization by presenting more storage to a host than is physically consumed and by allocating storage only as needed from a shared virtual pool. The storage allocation operations are performed in small units of storage (of one track or 128 KB) called virtually provisioned device extents.

Dell EMC recommends the following best practices for creating thin LUNs to obtain optimal performance:

- Place data and log files on separate LUNs.
- Separate production and nonproduction LUNs into different storage groups.
- Leave room for data growth and avoid exceeding 80 percent capacity of the LUNs for database files.
- Carve LUNs for data files on SSD drives.
- Use the latest storage array firmware.
- Configure end-to-end multipathing for redundancy and throughput optimization.
- Ensure that I/O is distributed to all front-end ports.

After creating the LUNs, create a server cluster for high availability and map the volumes to the server cluster by using auto provisioning. With auto provisioning:

- Related initiators (HBAs) are grouped by an initiator group.
- Related front-end ports are grouped by a port group.
- Related devices are grouped by a storage group.
- A masking view associates the above initiator, port, and storage groups.

At the time the masking view is created, a single operation automatically performs all required mapping and masking. This approach dramatically simplifies storage allocation.

RAID specification

Selecting the appropriate RAID type is an important decision for a successful implementation. Consider elements such as performance, capacity, and availability.

Dell EMC storage systems support RAID 5, and RAID 6. Each RAID type provides different performance, capacity, and protection levels, as follows:

- **RAID 5**—Stripes data across disks in large stripe sizes. RAID 5 stores parity information across all disks so that data can be reconstructed, which protects against a single disk failure. With its high write penalty, RAID 5 is most appropriate in environments with mostly read I/O.
- **RAID 6**—Stripes data across disks in large stripe sizes. RAID 6 stores two sets of parity information across all disks so that data can be reconstructed. RAID 6 can accommodate the simultaneous failure of two disks without data loss.

Customers select the RAID type and disk type combination based on their requirements. In a FAST for Virtual Pools (FAST VP) environment, Dell EMC recommends RAID 5 for flash and RAID 6 for SATA. The following table shows how RAID levels compare in read, write, sequential read, and sequential write scenarios.

Table 6. Comparison of RAID levels

RAID level	Random read	Random write	Sequential read	Sequential write
RAID 5	Excellent	Good	Good	Moderate
RAID 6	Good	Poor	Good	Moderate

Global memory

Global memory is a crucial component in the solution architecture. All read and write operations are transferred to or from global memory. Transfers between the host processor and channel directors can be processed at much greater speeds than transfers involving physical drives. PowerMaxOS uses complex, intelligent, statistical pre-fetch algorithms, which can adjust to workload on the array by constantly monitoring, evaluating, and optimizing cache decisions.

You can configure PowerMax arrays with up to 2 TB of mirrored memory per engine and up to 16 TB mirrored memory per array (PowerMax 8000)

Host connectivity

Configuring multiple connections from the host servers (direct connect) or FC switch (SAN connect) to the system provides channel redundancy. With SAN connectivity through FC switches, each front-end port can support multiple host attachments, enabling storage consolidation across a large number of host platforms. The multiple connections are distributed across separate directors to ensure uninterrupted access in the event of a channel failure. A minimum of two connections per server or SAN provides full redundancy.

Host connectivity to the front-end director ports is spread across physical components for the most efficient form of redundancy.

Dell EMC recommends the following best practices for connecting a host or cluster:

- Configure two to four front-end paths in the port group for masking and establishing zones to the host (single initiator zoning is recommended).

For cabling options, one approach is to connect all even-numbered ports to FabricA and all odd-numbered ports to FabricB.

- In single-engine systems, select two I/O ports spanning both SAN fabrics on each director.

Example: Port 4 and 24 on both directors 1 and 2

- In a multi-engine system, select an odd and even port on at least two different directors spanning different engines to spread the load for performance and ensure fabric redundancy.

Example: Port 4 on directors 1, 2, 3, and 4.

The following figure shows an example of redundant SAN connectivity to a PowerMax array with a single Brick block.




Figure 7. SAN connection to a PowerMax array with a single Brick

Setting I/O paths

FC zoning partitions the FC fabric into multiple subsets. A zone includes FC initiators and targets and enables communication between them. Consider the following when setting the I/O paths:

- Single initiator/multiple target zoning creates each FC zone with a single initiator (HBA port) and multiple targets (PowerMax front-end ports). Therefore, each HBA port requires a separate FC zone with PowerMax front-end ports. Create independent zones for each HBA installed in the host.
- WWN zoning contains only the host HBA port and the PowerMax front-end primary ports. In most cases, you do not have to include the PowerMax front-end reserve ports because they are not used for volume mapping. Dell EMC recommends creating zones by using a single initiator host port and multiple PowerMax ports.
- Multipathing specifies how many of the Dell EMC storage front-end ports that the system allows the volume to be mapped through. Dell EMC recommends that you set this option to **OS Default**.
- When using a mixed fabric (that is, a fabric containing two or more switches running different release levels of Fabric OS), use the switch with the latest Fabric OS level to perform zoning tasks.
- Perform zoning using the core switch rather than an edge switch.

PowerMax inline data reduction

PowerMax inline data reduction significantly increases the effective capacity of the PowerMax array by reducing the physical footprint of the dataset and, therefore, the number of NVMe flash disks required. Inline data reduction also intelligently optimizes system resources to ensure that the system is delivering the best balance of performance and efficiency at all times.

Data reduction is managed at the storage group level. You can easily enable or disable data reduction to target the workloads that would benefit most. You can also apply data reduction to existing data. Each storage group reports the effective data reduction ratio (CR) related to the data that is specific to that storage group. A PowerMax array typically consists of a variety of storage groups and SAP landscapes with production and nonproduction databases, each of which has its own achieved CR. The overall system CR is a mix of the various underlying storage group ratios. With a normal mix of SAP workloads, you can expect to see an approximately 2:1 system CR. This ratio could be higher or lower depending on the data mix. In our test environment, we saw a CR of 2.7:1 across our SAP landscape. When inline data reduction is combined with other PowerMax space-saving capabilities (such as Dell EMC Virtual Provisioning technology, zero space reclaim, and space-efficient snapshots), an overall efficiency rate of 4:1 is achievable.

For more information, see the [PowerMax Data Reduction White Paper](#).

Embedded NAS

Embedded NAS (eNAS) extends the value of PowerMax to file storage. PowerMax eNAS uses the hypervisor that is provided in PowerMaxOS to create and run a set of virtual machines (VMs) on PowerMax directors. These VMs host the two major elements of eNAS: software data movers and control stations. The VMs are distributed based on the

PowerMax mirrored pair architecture to evenly consume resources for both performance and capacity.

Distributed SAP NetWeaver systems architecture requires that global file systems are physically shared from the ABAP SAP Central Services (ASCS) instance using the NFS protocol to all SAP NetWeaver Application Servers (AS) belonging to the same SAP system. For example:

- Global directories such as `/sapmnt/SID` and its subdirectories
- The global transport directory, `/usr/sap/trans`

Two methods are used to physically share these global file systems:

- Server-based NFS
- Storage-based NFS

PowerMax eNAS can provide storage-based NFS to satisfy the SAP global file system requirements.

Customers should plan, implement, and configure eNAS in PowerMax storage arrays to obtain maximum operating efficiencies from the eNAS component of the array. For more information, see the [Dell EMC PowerMax eNAS Best Practices White Paper](#).

SRDF/Metro

Dell EMC recommends the following best practices for SRDF/Metro configuration:

- Install the SRDF/Metro license on both arrays that are to be managed.
- Ensure that the R1 and R2 are identical in size.
- Do not set Geometry Compatibility Mode (GCM) on any devices.
- Do not set User Geometry on any devices.

Note the following:

- Online device expansion is not supported.
- The following operations apply to all devices in the SRDF group: `createpair -establish`, `establish`, `restore`, `suspend`.
- Control of devices is not supported in an SRDF group that contains a mixture of R1s and R2s.
- vWitness configurations require Embedded Element Management (EEM or eMgmt) on each SRDF/Metro paired array.
- vWitness vApp requires VMware ESX 4.0 or later and Dell EMC Solutions Enabler or Unisphere 8.3 or later.

Data Domain configuration

Dell EMC recommends the following best practices for configuring your Data Domain system:

- Use 10 GbE instead of 1 GbE for the cabling.
- Use Data Domain Boost (DDBoost) for the backup and restore operations.

- Try to keep the Data Domain system less than 85 percent full. File system cleaning and other operations are faster and more efficient when the system has enough disk space available to perform them.
- Schedule file system cleaning to run during times when active backups are not running, to avoid resource contention and extended cleaning times.
- If replicating to a Data Domain system at another site, do not schedule the replication to overlap with your active backup window. Both processes require substantial resources and less time if they run separately rather than concurrently.
- Never use encryption, multiplexing, pre-compression, or client-side deduplication from the client backup application. These features greatly reduce the compression factor that is obtained on the Data Domain system. Perform these activities on the Data Domain system only. Ensure that none of these features are enabled by default for your application.
- Avoid selecting the highest available values for configuration options—number of streams, throttling, replication, and so on. Use of more moderate configurations often offers the best overall performance.
- Configure multiple dedicated network connections between the Oracle RMAN server and the Data Domain system to maximize throughput and provide redundancy.
- Note that many factors determine capacity requirements, including the size of the database being protected, rate of change, type of backups being performed, and retention requirements for the backups. A knowledgeable Data Domain system engineer can perform a detailed analysis using tools and models that have been developed over thousands of real-world deployments.
- Use the Oracle backup method 'RMAN backup' to obtain a faster backup speed. The RMAN channel set follows the "1 CPU with one channel" rule for the system.
- Restrict both NFS and CIFS shares that are accessed directly by Oracle RMAN to the servers that are running Oracle RMAN operations. Backup security to a Data Domain system is usually maintained by controlling access to the share points.

For information about configuring Data Domain for SAP landscapes, see the [Dell EMC Ready Systems for SAP Landscapes Deployment Guide](#).

Running SAP on VMware

Dell EMC recommends that you adhere to the following VMware sizing rules and considerations when you run your SAP system on VMware.

- Enable Hyper-threading.
- Disable Power Saving mode.
- Try to size the VM to fit within the NUMA node (non-wide).
- Install VMware Tools and perform the configuration using the VMXNET3 network adapter.
- Spread the virtual disks across all four of the available virtual SCSI controllers to a VM to maximize performance.
- Spread the database files across multiple LUNs and LUN queues to maximize I/O performance.

- Separate database log and data files into different LUNs to minimize I/O contention.
- Use the Virtual Machine File System (VMFS) where possible to increase operational management efficiency.
- Use eager zeroed thick format for all virtual disks.
- Note that the maximum number of vCPUs in VMware is 128 for a single VM.
- Use vSphere HA for all SAP instances.
- Install Central Services instances (ASCS/SCS) in a standalone VM and ensure that the VM is protected with VMware Fault Tolerance, which supports up to four vCPUs.
- Ensure that you use the latest SAP host agent.

For more information, see [SAP on VMware Best Practices](#).

The following figure shows our storage design layout for VMware VMs.

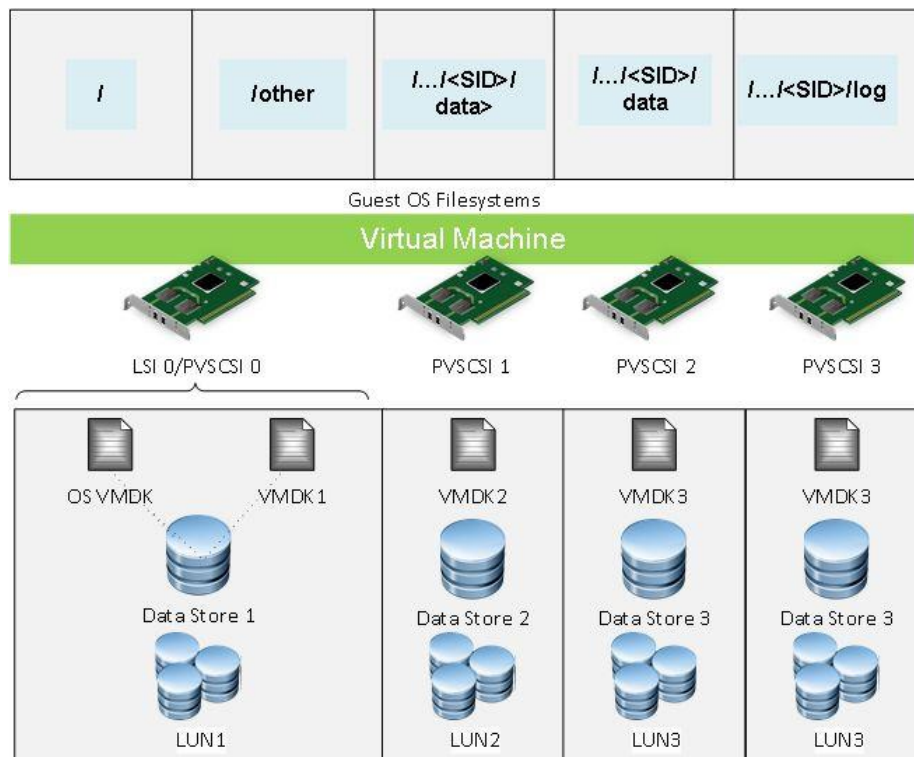


Figure 8. Storage design layout for VMware virtual machines

Note: If you are planning to deploy SAP HANA and SAP HANA-related applications with VMware vSphere, see [SAP HANA on VMware vSphere](#) for support and sizing information.

vSphere HA and SRDF/Metro

When configuring your Ready Solution with SRDF/Metro, follow these Dell EMC- and VMware-recommended best practices:

- Configure the front end with a stretched layer 2 network so that when a VM moves between sites, its IP address stays the same.

- If you want site high availability, configure two or more ESXi hosts at each site. Ensure that vSphere HA and Distributed Resource Scheduler (DRS) are active and that VM monitoring is activated.

Note: If an ESXi host fails, vSphere HA powers on the SAP VMs according to the VM restart priority. ASCS has the highest priority followed by the database, followed in turn by the application servers.

- Use host affinity rules. Host affinity rules keep VMs running in the preferred site as long as possible, only moving them to the nonpreferred site if they cannot run in the preferred site. Also use affinity rules to control VM placement and ensure that the SAP database and ASCS instances are running on separate hosts at all times.
- Use affinity rules to ensure that Oracle database VMs are running on processor cores that are fully licensed for Oracle in VMware environments. For more information, see [Understanding Oracle Certification, Support and Licensing for VMware Environments](#) on the VMware website.
- Ensure that the server resources exist to fail over to a single site. Configure admission control of vSphere HA for 50 percent CPU and 50 percent memory.
- Configure heartbeat mechanisms, which vSphere HA uses to validate the state of a host, as follows:
 - **Network (primary)**, which pings the default gateway—Specify a minimum of two additional isolation addresses. Configure these addresses in the **Advanced Options** section using the name `das.isolationAddress.x` (`x` is incremented for each address).
 - **Datastore (secondary)**—Increase the number of heartbeat data stores from two to four in a stretched cluster environment. Configure the data stores in the **Advanced Options** section.
- Change the **Disk.PathEvalTime** on each ESXi host from 300 s to 30 s to avoid the need to manually rescan after presenting the R2 devices.
- Enable VM Component Protection (VMCP) to take advantage of new capabilities around All Paths Down (APD) and Permanent Device Loss (PDL).
To enable VMCP, check the **Protect against Storage Connectivity Loss** option. Then, in the **Failure conditions and VM response** list box under the **vSphere HA Cluster setting**, select **Power off and restart VMs** for:
 - Response for Datastore with PDL
 - Response for Datastore with APD
- Use a witness instead of a bias with SRDF/Metro to provide the best availability.

For more information, see [VMware vSphere Metro Storage Cluster Recommended Practices](#) on the VMware website.

SAP system architecture

Because SAP ERP is one of the most important transactional systems in a typical enterprise IT environment, the system architecture must consider both performance and availability. Dell EMC highly recommends a distributed system architecture, where each of the main components resides on its own VM, as follows:

- **ABAP central services (ASCS) instance**—ASCS comprises a message server and an enqueue server that are both SPOFs. Separating ASCS from an application server instance in a central system architecture can minimize the impact of other work processes. A lower chance of failure also provides for the highest level of protection by using VMware fault tolerance. The SAP shared file systems `/sapmnt/<SID>` and `/usr/sap/<SID>` can be stored on this instance and shared to all other SAP instances within the same system.
- **Database instance**—A dedicated database instance has full command of its VM resources and is isolated from any other possible threat to the stability of the database. Because the network traffic between the database instance and the application server instances is usually high and RAM state change within the VM is frequent, avoid using VMware FT to protect the database instance. Use operating system/database-specific clustering tools or vSphere HA to provide a higher level of protection, but do not use both together.
- **Additional application server (AAS) instances**—AAS is a scale-out architecture that performs most of the computational tasks when transactions and background jobs are executed. You can add AAS at any time for additional performance and availability. Login groups (T-Code: SMLG) usually manage access to provide flexibility and increase availability. If one application server instance fails, connected users lose connection and reconnect to other AAS instances, and transaction-in-flight is rolled back. Standard vSphere HA is sufficient to provide a quick restart from ESXi or OS failures.

Sizing and performance

Sizing considerations for SAP landscape design

Business and technology considerations influence the sizing of the hardware infrastructure. Work with Dell EMC representatives 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 Ready Solution.

Table 7. Infrastructure sizing steps

Step	Activity	Performed by
1	Determine the number of SAP production systems you need.	Dell EMC completes these steps when collecting your business requirements.
2	Define the SAP system landscapes. Typically, each SAP production system (ERP, BW, CRM, and so on) has its own SAP system landscape, consisting of a DEV, QAS, SBX, and PRD environment.	
3	Determine if virtualization is to be used.	

Step	Activity	Performed by
4	Determine the high availability requirements.	
5	Determine the disaster recovery requirements and the number of data centers or sites involved.	
6	Consider data protection requirements for backing up the SAP systems.	
7	Determine an expected annual data growth and the number of years of maintenance required.	
8	Size each of the SAP systems and system landscapes using the SAP Quick Sizer tool on your production systems.	You complete this step and then provide the results to Dell EMC.
9	Calculate the total compute requirements and determine the number and models of servers (CPU/memory/storage) that are required to support the workloads.	Dell EMC completes these steps using the information that you provided in step 8.
10	Calculate the total storage requirements, based on capacity or IOPS or both, and use the PowerMax Sizer tool to size the PowerMax array.	
11	Determine the backup capacity requirements for Data Domain.	

SAP Quick Sizer

SAP Quick Sizer is a web-based tool that calculates hardware requirements. The Quick Sizer tool is available in two versions. Use the SAP HANA version to size an SAP HANA database and the classic version if you want to size a non-SAP HANA database.

The tool bases its calculations on functional parameters such as the number of users working with the different SAP application systems, throughput, and other inputs.

SAP has identified different and independent sizing models with different advantages and disadvantages. The Quick Sizer tool incorporates both of the following sizing models:

- **User-based sizing**—In this model, SAP identified three types of active users who work with the system in different ways. Note that the number of users provides little information about the throughput that these users produce.
- **Throughput-based sizing**—This model relies on actual or expected throughput. Business assumptions such as the number of order line items per year must be cross-checked against individual installations.

Quick Sizer presents the results in SAPS. Hardware vendors provide their SAPS for a particular server configuration by running [SAP Benchmark](#) tests and posting the results on the SAP website. For more information about SAPS sizing, see the Quick Sizer documentation on the [SAP Service Marketplace](#) (access requires login credentials).

SAPS overview

SAPS is a hardware-independent unit of measure that describes the performance of a system configuration in the SAP environment. It is derived from the SD benchmark, where 100 SAPS is defined as 2,000 fully business-processed order line items per hour.

In the SD benchmark, fully business-processed means the full business process of an order line item, as follows:

- Creating the order
- Creating a delivery note for the order
- Displaying the order
- Changing the delivery
- Posting a goods issue
- Listing orders
- Creating an invoice

This throughput is achieved by processing 2,400 SAPS transactions, 6,000 dialog steps (screen changes), or 2,000 postings per hour in the SD benchmark. SAPS is divided into requirements for the database layer (database SAPS) and application layer (application SAPS). Database SAPS are more relevant for sizing your storage requirements.

Using SAPS for sizing

The design approach for this Ready Solution uses the published SAPS values for the PowerEdge R840 and R740/R740xd servers based on official SAP SD Standard Application Benchmark Results. The certifications are based on high-performance Intel Xeon Platinum 8180 processors. Dell EMC performed internal testing with Standard Performance Evaluation Corporation (SPECint) benchmarks to extrapolate and determine values for SAPS and SAPS per core. Calculations are for all platinum, gold, silver, and bronze processors available in the R840 and R740/R740xd PowerEdge server range. You can calculate server sizing and storage capacity requirements by using the output from SAP Quick Sizer projects and for the R840 or R740/R740xd PowerEdge server model by using the SAPS values.

The following table provides details of certified SAPS values for PowerEdge servers. For a full list of PowerEdge R840 and R740/R740xd systems with extrapolated SAPS values, see [Appendix: SAPS values for PowerEdge R840 and R740/R740xd](#).

Table 8. SAPS values for PowerEdge servers

PowerEdge server (CPU model)	Number of cores	Number of sockets	Number of cores per socket	Certified SAPS	SAPS per core
PowerEdge R840 (Intel Xeon Platinum 8180, 2.50 GHz)	112	4	28	359,620	3,211
PowerEdge R740 or R740xd (Intel Xeon Platinum 8180M, 2.50 GHz)	56	2	28	175,230	3,129

Note: The Quick Sizer tool bases calculations on 65 percent utilization. Therefore, additional calculations to account for overhead are not required to achieve predictable server behavior. You can compare your Quick Sizer SAPS result with existing certified benchmark results. However, you can expect a 10 percent performance degradation in virtualized environments.

Greenfield and brownfield sizing

We use the SAP QuickSizer tool to estimate greenfield (new) sizing for SAP standard solutions if you are planning either of the following:

- A new installation of SAP software
- Addition of SAP software components to an existing SAP environment (delta-sizing)

We use brownfield (post-go-live) sizing if you are currently running SAP but want to expand your hardware capacity or introduce new infrastructure. For brownfield sizing, we collect data and metrics from the system on the existing hardware and use the information to extrapolate sizing as follows:

- SAP Early Watch reports provide the utilized SAPS value.
- SAP application monitors, such as DB02, ST06, and ST03N, provide results for database capacity, database growth, workload performance, and server or system utilization.
- The [Dell EMC Live Optics](#) tool gathers core requirements such as disk I/O, throughput, capacity, memory, and CPU utilization and produces an in-depth view of the server workload and capacity requirements.

Interpreting Quick Sizer results

This section provides information about how to interpret Quick Sizer results for CPU, disk, disk I/O, and memory requirements.

Sizing the CPU for database and application servers

CPU sizing results provide requirements for the relevant layers, primarily the database layer and the ABAP application layer. The results are specified in SAPS, which, in CPU terms, is a performance unit that describes the throughput power of a server. Two-thousand fully business-processed order line items per hour equates to 100 SAPS. If a customer provides the total SAPS requirement but not the breakdown between database SAPS and application SAPS, we apply general rules for the ratio of database SAPS to application SAPS:

- For transactional (OLTP) systems such as ERP, we apply a 1:4 ratio. For example, an ERP system with a total SAPS requirement of 15,000 translates to 3,000 database SAPS and 12,000 application SAPS.
- For analytical (OLAP) systems such as BW, we apply a ratio of 1:2 database SAPS to application SAPS. For example, a BW system with a total SAPS requirement of 15,000 translates to 5,000 database SAPS and 10,000 application SAPS.

Sizing the disk and disk I/O

Disk sizing is the general sizing determined for the database by the SAP Quick Sizer tool, using the following assumptions:

- The customer provides the disk size result.
- For throughput sizing, disk growth is determined by the number of business objects per year, their size, and the length of time they will remain in the database.

Some factors that might influence disk size, such as small tables, indexes, backups, dumps, and exports, are not considered or are considered only as exceptions. Calculating disk I/O is linked to database SAPS based on the analytical or transactional consumption of SAPS and disk I/O.

The following general rule applies for translating SAPS to IOPS:

- **OLTP workloads such as Business Suite**—Database SAPs x 0.6 = IOPS requirement
- **OLAP workloads such as BW**—Database SAPs x 0.9 = IOPS requirement

Note: Disk space calculations from Quick Sizer are for database space only. Also consider storage sizing for OS, swap, database software binaries, and SAP software binaries on each host.

Sizing the system memory

If a customer does not provide a memory size result from the SAP Quick Sizer project, we use the following general rule from the SAPS requirement: estimate 3 GB of memory for every 1,000 SAPS in most throughput-based sizing.

Note: Quick Sizer results are assumption-based estimates. Contact your Dell EMC representative to obtain solution sizing for your specific requirements.

Single system versus system landscape considerations

Quick Sizer results are for a single application system. Sizing for SAP must be performed against a system landscape rather than a single system for each application that is implemented. A basic SAP system landscape comprises at least three systems: PRD, QAS, and DEV. Many customers have up to five or seven systems per SAP application, which can result in a high number of systems in a landscape, with each system demanding resources.

Virtualization sizing considerations

VMware has provided a direct comparison between virtualized and bare-metal systems with the same hardware configuration and the same SAP benchmark workload. The overhead is less than 6 percent. When sizing a virtualized SAP system or landscape with VMware vSphere, consider an overhead of 10 percent. Accordingly, if you are using VMware virtualization, you can calculate the virtual SAPS as 90 percent of the certified server SAPS value. In addition, follow these guidelines:

- Do not overcommit CPU for production systems. (This does not apply to nonproduction systems.)
- Never overcommit memory with SAP environments, whether production or nonproduction.

The maximum number of vCPUs in VMware is 128 for a single VM, which might limit the SAPS value that is achievable in a single VM.

For more information, see the following documents:

- [SAP on VMware Best Practices](#) (on the VMware website)
- [SAP Sizing and Virtualization Guide](#) on the [SAP Service Marketplace](#) (access requires login credentials)

High availability

With a VMware ESXi cluster, you define two or more physical machines that are to provide resources for the hosts (or resource pools) that are assigned to the cluster. High availability and load balancing of VMs (VMware DRS) is achievable by using ESXi clusters. Both these features use the vSphere VMotion tool to move these virtual guests from one physical host to another.

When you size clusters, the amount of spare compute resources that you require depends on the number of physical hosts in the cluster. Spare compute resources must be available to take over the workload for at least one physical host in the cluster to ensure continuous availability.

Disaster recovery and multisite data centers

Failure and downtime of mission-critical SAP environments can bring an entire organization or large parts of it to a halt. The business costs associated with SAP downtime can be high. If your primary site suffers a disaster (such as flooding, fire, or a major earthquake), the disaster recovery plan must have your normal operations up and running in as short a time (recovery time objective, or RTO) as possible and with minimal data loss (recovery point objective, or RPO).

With multisite data centers, the secondary site must have enough available compute, network, and storage resources to take over the workloads from the failed site. Therefore, the compute, network, and storage requirements are doubled when you size for disaster recovery.

Sizing for future database growth

SAP systems grow year on year over time, affecting the size of the database and the number of users. Dell EMC recommends factoring an annual growth rate into the sizing of the storage requirements over the term of the maintenance support for the infrastructure.

For example, if a database sizing is 600 GB and you expect an annual growth rate of 8 percent with maintenance for 4 years, calculate the requirement as follows:

$$600 \text{ GB} \times 1.08 \times 1.08 \times 1.08 \times 1.08 = 816 \text{ GB}$$

SAP HANA TDI Phase 5

SAP HANA TDI Phase 5 introduced some important changes and benefits for customers:

- SAP HANA QuickSizer and SAP HANA sizing reports were enhanced to provide separate CPU (SAPS) and RAM sizing results, enable customers to fine-tune their system configurations for their specific workload, and purchase systems with the optimal number of cores and memory.
- SAP supports the complete Intel Broadwell E7 processor family as well as the Skylake Platinum, Gold, and Silver processor family (8+ cores per processor).

This combination of workload-driven SAPS-based sizing with an increased variety of CPU processors gives customers more flexibility when choosing SAP HANA compute nodes from the wide range of R840 and R740/R740xd systems to cost-optimize the servers for their specific workload requirements.

For a full list of PowerEdge R840 and R740/R740xd systems with extrapolated SAPS values, see [Appendix: SAPS values for PowerEdge R840 and R740/R740xd](#).

Sizing example

The following sizing example uses sample customer inputs to show how you can determine compute and storage requirements for SAP landscapes.

A customer has three PRD systems, ERP, CRM, and BW, each with a system landscape consisting of DEV, QAS, and SBX systems. Based on its Quick Sizer projects, the customer provided the information in the following table for each PRD system.

Table 9. Customer inputs from a Quick Sizer project

Application	Application type	Total SAPS	Disk size (GB)
ERP	OLTP	96,000	4,000
CRM	OLTP	32,000	1,000
BW	OLAP	60,000	2,000

The customer has two data centers that are 30 miles apart and has chosen to use VMware virtualization for server consolidation and high availability.

The customer advised that the quality system should be sized the same as production, with a VMware CPU overcommit factor of 2. Both the DEV and SBX system should be sized at 50 percent of production, with a VMware CPU overcommit factor of 4.

SRDF/Metro is included for storage continuous availability and disaster recovery with storage replication for the PRD and QAS systems. A Data Domain system is included for data protection and backups. The annual growth forecast is 7 percent with a maintenance term of 4 years.

The following table shows the SAP landscapes that we derived from the customer inputs.

Table 10. SAP landscapes and customer inputs from Quick Sizer

System	Application	Type	Description	Site	Replication site	SAPS (input)	Disk size (input) (GB)	Annual disk growth (input)
BM1	ERP	OLTP	PRD	A	B	96,000	4,000	7 %
QM1	ERP	OLTP	QAS	B	A	96,000	4,000	7 %
DM1	ERP	OLTP	DEV	A	—	48,000	2,000	7 %
SBX	ERP	OLTP	SBX	B	—	48,000	2,000	7 %
CRM	CRM	OLTP	PRD	B	A	32,000	1,000	7 %
CRQ	CRM	OLTP	QAS	A	B	32,000	1,000	7 %
CRD	CRM	OLTP	DEV	B	—	16,000	500	7 %
CRS	CRM	OLTP	SBX	A	—	16,000	500	7 %

System	Application	Type	Description	Site	Replication site	SAPS (input)	Disk size (input) (GB)	Annual disk growth (input)
BWP	BW	OLAP	PRD	A	B	60,000	2,000	7 %
BWQ	BW	OLAP	QAS	B	A	60,000	2,000	7 %
BWD	BW	OLAP	DEV	A	—	30,000	1,000	7 %
BWS	BW	OLAP	SBX	B	—	30,000	1,000	7 %

Sizing the compute requirements

Next, we used the customer inputs from Table 10 to calculate the compute and memory requirements, as shown in the following table.

Table 11. CPU and memory calculations for each system in the landscapes

System	Site	SAPS (input)	Database SAPS ¹	Application SAPS ¹	VMware CPU overcommit factor	Database SAPS (with VMware overcommit) ²	Application SAPS (with VMware overcommit) ²	Database memory (GB) ³	Application memory (GB) ³	Number of application servers ¹
BM1	A	96,000	19,200	76,800	1	19,200	76,800	57.6	230.4	5
DM1	A	48,000	9600	38,400	4	2,400	9,600	28.8	115.2	2
CRQ	A	32,000	6,400	25,600	2	3,200	12,800	19.2	76.8	3
CRS	A	16,000	3,200	12,800	4	800	3200	9.6	38.4	1
BWP	A	60,000	20,000	40,000	1	20,000	40,000	60	120	3
BWD	A	30,000	10,000	20,000	4	2,500	5,000	30	60	1
QM1	B	96,000	19,200	76,800	2	9600	38,400	57.6	230.4	5
SBX	B	48,000	9600	38,400	4	2,400	9,600	28.8	115.2	2
CRM	B	32000	6,400	25,600	1	6,400	25,600	19.2	76.8	3
CRD	B	16000	3,200	12,800	4	800	3,200	9.6	38.4	1
BWQ	B	60,000	20,000	40,000	2	10,000	20,000	60	120	3
BWS	B	30,000	10,000	20,000	4	2,500	5,000	30	60	1

¹ Ideally, the customer provides a breakdown of the database and application SAPS from a Quick Sizer project. Otherwise, we apply a general rule to calculate the ratio of database SAPS to application SAPS using total SAPS. For OLTP, we use 1:4 and for OLAP we use 1:2.

² The database and application SAPS values for each system are adjusted to include the VMware CPU overcommit factor.

³ To calculate the required memory for each database and application server, we use the general rule for throughput-based sizing, where we estimate 3 GB of memory for each 1,000 SAPS. We calculate this memory against the SAPS values without the CPU overcommit factor.

Note: In virtual environments, do not overcommit memory for SAP PRD and non-PRD systems.

For each site that is described in Table 10, we calculated the total SAPS and memory requirements, as shown in the following table.

Table 12. Total compute requirements for Site A and Site B

Site	Total database SAPS	Total application SAPS	Total SAPS	Total database memory (GB)	Total application memory (GB)	Total memory (GB)
A	48,100	147,400	195,500	205.2	640.8	846
B	31,700	101,800	133,500	205.2	640.8	846

To calculate the number of physical servers that are required, we compared our total compute requirements for each site with the physical capabilities of the SAPS value and the memory of the chosen server models.

Additional compute resources for high availability and disaster recovery are included on each site. That way, if a server or site is lost, the remaining site has the resources to take over the failed sites workloads.

In this example, the result was calculated using both the R840 and R740/R740xd models. The following table shows the physical compute requirements for the solution using the R840 server on each site.

Table 13. Physical compute requirements with the PowerEdge R840 server

Site	Total SAPS required	Total memory required (GB)	Virtual SAPS	Cores (4 x 28)	RAM (GB) ¹	Number of servers SAPS bound ²	Number of servers memory bound ³	Number of servers required ⁴	Number of servers required with HA ⁵
A	195,500	846	323,658	112	768	1	2	2	3
B	133,500	846	323,658	112	768	1	2	2	3

¹ In this sizing example, we are using PowerEdge R840 servers with 768 GB of RAM.

² The PowerEdge R840 has the capabilities of up to 323,658 virtual SAPS, and the SAPS requirement in the example is 195,500 for Site A and 133,500 for Site B. Therefore, one server on each site is needed to satisfy the requirement.

³ The PowerEdge R840 server in the example has 768 GB of RAM and the compute requirement is 846 GB of RAM on both site A and site B. Therefore, two servers on each site are needed to satisfy the requirement.

⁴ The higher value of ² and ³ is the number of servers required.

⁵ We added a server to the value of ⁴ for high availability.

The following table shows the physical compute requirements for the solution using the R740 or R740xd server on each site.

Table 14. Physical compute requirements with the PowerEdge R740 or R740xd server

Site	Total SAPS required	Total memory required (GB)	Virtual SAPS	Cores (2 x 28)	RAM (GB) ¹	Number of servers SAPS bound ²	Number of servers memory bound ³	Number of servers required ⁴	Number of servers required with high availability ⁵
A	195,500	846	157,707	56	384	2	3	3	5
B	133,500	846	157,707	56	384	1	3	3	5

¹ In this sizing example, we are using PowerEdge R740 servers with 384 GB of RAM.

² The PowerEdge R740 has the capabilities of up to 157,707 virtual SAPS, and the SAPS requirement in the example is 195,500 for Site A and 133,500 for Site B. Therefore, the requirement is for two servers on site A and one server on site B.

³ The PowerEdge R740 server in the example has 384 GB of RAM and the compute requirement is 846 GB of RAM on both site A and site B. Therefore, three servers are needed to satisfy the requirement.

⁴ The higher value of ² and ³ is the number of servers required.

⁵ We added a server to the value of ⁴ for high availability.

Sizing storage requirements

The following table presents the storage requirements for Site A and Site B.

Table 15. Customer storage requirements for Site A and Site B

System	Site	Replicati on site	Disk size (input) (GB)	DB disk final size (GB) ¹	Database SAPS (with VMware overcommit)	IOPS (calculated) ²	Number of app servers	Total OS size (GB) ³	Total exe size (GB) ³	Dump size (GB) ⁴	Total storage require-ments ⁵
BM1	A	B	4,000	5,243	19,200	11,520	5	768	768	5,243	12,022
DM1	A	—	2,000	2,622	2,400	1,440	2	384	384	2,622	6,011
CRQ	A	B	1,000	1,311	3,200	1,920	3	512	512	1,311	3,646
CRS	A	—	500	655	800	480	1	256	256	655	1,823
BWP	A	B	2,000	2,622	20,000	18,000	3	512	512	2,622	6,267
BWD	A	—	1,000	1,311	2,500	2,250	1	256	256	1,311	3,134
QM1	B	A	4,000	5,243	9600	5,760	5	768	768	5,243	12,022
SBX	B	—	2,000	2,622	2,400	1,440	2	384	384	2,622	6,011
CRM	B	A	1,000	1,311	6,400	3,840	3	512	512	1,311	3,646
CRD	B	—	500	655	800	480	1	256	256	655	1,823

System	Site	Replicati on site	Disk size (input) (GB)	DB disk final size (GB) ¹	Database SAPS (with VMware overcommit)	IOPS (calculated) ²	Number of app servers	Total OS size (GB) ³	Total exe size (GB) ³	Dump size (GB) ⁴	Total storage requirements ⁵
BWQ	B	A	2,000	2,622	10,000	9,000	3	512	512	2,622	6,267
BWS	B	—	1,000	1,311	2,500	2,250	1	256	256	1,311	3,134

¹ Future database disk size = disk size x annual growth of 7 percent over 4 years

² To estimate the required IOPS, calculate the database SAPs based on the analytical or transactional consumption of SAPs. See [Sizing the disk and disk I/O](#) on page 39. Disk I/O for OLTP applications = 60 percent of database SAPs. Disk I/O for OLAP applications = 90 percent of database SAPs.

³ 128 GB disk capacity for OS and 128 GB for exe on each DB and application server of a system

⁴ Storage capacity for a database dump, exports, or logs

⁵ Total system storage is the value of the disk + OS + exe columns added.

With storage replication and SRDF/Metro across both sites, the systems being replicated must also be included to size the storage on each site.

The following table presents the total storage requirements for Site A. The total represents requirements for each SAP system residing primarily on Site A as well as the systems from Site B that are being replicated to Site A.

Table 16. Storage requirements for Site A

Site A	Storage (GB)	IOPS
SAP systems	34,726	35,610
Replicated SAP systems	25,581	18,600
Total Site A requirements	60,307	54,210

Using the total storage and total IOPS, the Dell EMC representatives use the PowerMax Sizer tool to provide the configuration of the PowerMax array for Site A.

The following table presents the total storage requirements for Site B.

Table 17. Storage requirements for Site B

Site B	Storage (GB)	IOPS
SAP systems	32,903	22,770
Replicated SAP systems	21,935	31,440
Total Site B requirements	54,838	54,210

Using the total storage and total IOPS, your Dell EMC representatives use the PowerMax Sizer tool to provide a detailed configuration of the PowerMax array for Site B.

Conclusion

This Dell EMC Ready Solution for SAP is designed to consolidate and simplify IT for complex SAP ERP and Business Suite landscapes by lowering the TCO while delivering increased performance and productivity on infrastructure that is SAP HANA-certified. The solution enables application owners to accelerate application lifecycle management because it provides the integration and features to provision and manage SAP systems more efficiently, helping to decrease costs.

Customer goals for enhancing an SAP infrastructure include:

- Streamlining the system and reducing costs
- Consolidating SAP landscapes onto a single storage platform while delivering significant performance increases
- Preparing for the digital economy and modernizing SAP landscapes

This Ready Solution addresses these challenges by providing:

- Storage consolidation that provides for substantial savings through increased data reduction, higher density drives, and the ability to start small and grow.
- Consistently improved performance for a consolidated SAP landscape with consistent scale-out performance and close to 3 PB effective capacity in one array.
- Highly efficient data replication that is metadata-aware and therefore data reduction-aware. Less data on the wire means smaller bandwidth requirements per unit of written or replicated data.
- Ability to bring other workloads to the PowerMax array.
- Sizing calculations for SAP landscapes to ensure predictable performance.
- Simplified and repeatable standard for ordering and sizing infrastructure for SAP landscapes.

References

Dell EMC documentation

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

- [*Dell EMC Ready Bundle for SAP Landscapes Deployment Guide*](#)
- [*Using SRDF/Metro in a VMware Metro Storage Cluster Running Oracle E-Business Suite and 12C RAC White Paper*](#)
- [*The PowerMax Storage Family White Paper*](#)
- [*Dell EMC PowerMax Family*](#)
- [*Dell EMC Data Domain Deduplication Storage Systems Spec Sheet*](#)
- [*Dell EMC Data Domain Boost for Enterprise Applications and ProtectPoint Application Agent Installation and Administration Guide*](#)
- [*Dell EMC Data Domain Operating System Administration Guide*](#)
- [*Dell EMC Live Optics—Overview*](#)
- [*Dell EMC Networking S5200-ON Series Switches Specification Sheet*](#)
- [*Dell EMC Networking S4100-ON Specification Sheet*](#)
- [*Dell Networking S3048-ON Specification Sheet*](#)

VMware documentation

The following VMware documentation provides additional and relevant information:

- [*Architecture Guidelines and Best Practices for Deployments of SAP HANA on VMware vSphere*](#)
- [*Virtualize Business Critical Applications*](#)
- [*SAP and VMware Virtualization--Overview*](#)
- [*SAP Solutions on VMware Best Practices Guide*](#)
- [*VMware vSphere Metro Storage Cluster Recommended Practices*](#)
- [*Understanding Oracle Certification, Support and Licensing for VMware Environments*](#)

SAP documentation

The following SAP website provides access to additional and relevant information: [SAP Service Marketplace](#). Access to the website requires valid login credentials.

Under [SAP Quick Sizer documentation](#) search for the following documents:

- *Classic Quick Sizer for Beginners V46*
- *Classic Rough Guide V46*
- *Quick Sizer Best Practices*

Appendix: SAPS values for PowerEdge R840 and R740/R740xd

The following tables show the estimated SAPS values for Dell EMC PowerEdge R840 and R740/R740xd servers, respectively. The SAPS values for the top-end platinum CPU model, the Intel Xeon Platinum 8180, are certified with SAP. These values are presented in bold in the tables. Dell EMC extrapolated the SAPS values for all other PowerEdge server models from these results using SPECint. All the listed models are SAP HANA-certified with [SAP HANA tailored data center integration \(HANA TDI\) phase 5](#). For more information, see [SAP HANA TDI Phase 5](#) on page 41.

Note: Dell EMC performed internal testing with the SD benchmarks to validate the extrapolated SAPS values across the platinum, gold, silver, and bronze CPU range. The results showed that the extrapolations were accurate to within a 5 to 10 percent range in our laboratory environment. "M" editions of a CPU model have the same extrapolated SAPS values as non-M editions of that model.

Table 18. PowerEdge R840 systems with values for SAPS and SAPS per core

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)	40	4	10	124,652	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)	56	4	14	164,540	2,938
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)	64	4	16	198,819	3,107
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)	48	4	12	177,629	3,701
PowerEdge R840 (Intel Xeon Gold 6138, 2.00 GHz)	80	4	20	225,619	2,820
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)	32	4	8	133,377	4,168
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 6150, 2.70 GHz)	72	4	18	240,578	3,341
PowerEdge R840 (Intel Xeon Gold 6152, 2.10 GHz)	88	4	22	249,927	2,840
PowerEdge R840 (Intel Xeon Gold 6154, 3.00GHz)	72	4	18	262,392	3,644
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)	96	4	24	271,117	2,824
PowerEdge R840 (Intel Xeon Platinum 8164, 2.00 GHz)	104	4	26	274,234	2,637
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)	112	4	28	306,020	2,732

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R840 (Intel Xeon Platinum 8180, 2.50GHz)	112	4	28	359,620	3,211

Table 19. PowerEdge R740 systems values for with SAPS and SAPS per core

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R740 (Intel Xeon Bronze 3106, 1.70 GHz)	16	2	8	2,7933	1,746
PowerEdge R740 (Intel Xeon Silver 4108, 1.80 GHz)	16	2	8	40,801	2,550
PowerEdge R740 (Intel Xeon Silver 4109T, 2.00 GHz)	16	2	8	41,447	2,590
PowerEdge R740 (Intel Xeon Silver 4110, 2.10 GHz)	16	2	8	45,068	28,17
PowerEdge R740 (Intel Xeon Silver 4114, 2.20 GHz)	20	2	10	5,7936	2,897
PowerEdge R740 (Intel Xeon Silver 4116, 2.10 GHz)	24	2	12	67,247	2,802
PowerEdge R740 (Intel Xeon Gold 5115, 2.40 GHz)	20	2	10	62,850	3,143
PowerEdge R740 (Intel Xeon Gold 5118, 2.30 GHz)	24	2	12	73,713	3,071
PowerEdge R740 (Intel Xeon Gold 5120, 2.20 GHz)	28	2	14	82,765	2,956
PowerEdge R740 (Intel Xeon Gold 6126, 2.60 GHz)	24	2	12	82,765	3,449
PowerEdge R740 (Intel Xeon Gold 6130, 2.10 GHz)	32	2	16	97,637	3,051
PowerEdge R740 (Intel Xeon Gold 6132, 2.60 GHz)	28	2	14	97,637	3,487
PowerEdge R740 (Intel Xeon Gold 6134, 3.20 GHz)	16	2	8	66,600	4,163
PowerEdge R740 (Intel Xeon Gold 6136, 3.00 GHz)	24	2	12	93,758	3,907
PowerEdge R740 (Intel Xeon Gold 6138, 2.00 GHz)	40	2	20	11,3803	2,845
PowerEdge R740 (Intel Xeon Gold 6140, 2.30 GHz)	36	2	18	113,803	3,161
PowerEdge R740 (Intel Xeon Gold 6142, 2.60 GHz)	32	2	16	109,276	3,415
PowerEdge R740 (Intel Xeon Gold 6148, 2.40 GHz)	40	2	20	126,088	3,152
PowerEdge R740 (Intel Xeon Gold 6150, 2.70 GHz)	36	2	18	124,148	3,449
PowerEdge R740 (Intel Xeon Gold 6152, 2.10 GHz)	44	2	22	127,381	2,895
PowerEdge R740 (Intel Xeon Gold 6154, 3.00 GHz)	36	2	18	132,554	3,682
PowerEdge R740 (Intel Xeon Platinum 8153, 2.00 GHz)	32	2	16	85,352	2,667
PowerEdge R740 (Intel Xeon Platinum 8160, 2.10 GHz)	48	2	24	138,374	2,883
PowerEdge R740 (Intel Xeon Platinum 8164, 2.00 GHz)	52	2	26	144,193	2,773
PowerEdge R740 (Intel Xeon Platinum 8164, 2.00 GHz)	52	2	26	139,020	2,673
PowerEdge R740 (Intel Xeon Platinum 8168, 2.70 GHz)	48	2	24	157,125	3,273
PowerEdge R740 (Intel Xeon Platinum 8170, 2.10 GHz)	52	2	26	147,426	2,835
PowerEdge R740 (Intel Xeon Platinum 8176, 2.10 GHz)	56	2	28	155,185	2,771

Appendix: SAPS values for PowerEdge R840 and R740/R740xd

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R740 (Intel Xeon Platinum 8180, 2.50 GHz)	56	2	28	175,230	3,129