

Dell EMC Ready Bundle for SAP Landscapes with Dell EMC Unity Storage Arrays

Infrastructure Design for SAP Landscapes

April 2018

H16947.1

Design Guide

Abstract

This design guide provides best-practice guidelines for designing and sizing an infrastructure solution for SAP landscapes. The solution incorporates Dell EMC PowerEdge servers, Unity storage systems, Connectrix Fibre Channel SAN switches, and Data Domain storage protection.

Dell 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 © 2018 Dell Inc. or its subsidiaries. All Rights Reserved. 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 the property of their respective owners. Published in the USA 04/18 Design Guide H16947.1.

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 architecture14**
- Data protection16**
- Infrastructure-level high availability17**
- Solution configuration best practices19**
- Sizing and performance.....28**
- Conclusion.....38**
- References.....39**
- Appendix: SAPS values for PowerEdge R940 and R740/R740xd.....40**

Introduction

Executive summary

Modern companies are looking for technologies that transform their IT departments into agile business units capable of delivering continuous application availability through updates and upgrades. Well-managed change can enable greater efficiencies along with more reliable and secure delivery of services.

Performance is central to an excellent user experience. Maintaining performance while scaling an application ecosystem ensures that response times meet service-level agreements (SLAs) for financial and supply chain management, among other functions. Application administration teams must have the confidence to deploy copies of Enterprise Resource Planning (ERP) systems while providing a highly responsive experience for end users, developers, and other groups. The automation of routine tasks such as provisioning ERP systems means a faster time-to-value, an increase in operational efficiencies, and delivery of a more reliable product.

The modern application system must include a broad range of resiliency capabilities, high availability (HA), and protection solutions that provide optimal application uptime.

Business case

Traditionally, IT teams have selected the components of an application infrastructure separately. Because the approach does not always deliver the expected results, many IT teams now look for integrated solutions that have been pretested and precertified and can be accurately sized to meet their business requirements. The Dell EMC™ Ready Bundle for SAP Landscapes solution delivers compute, networking, and storage in one system. This solution incorporates laboratory-validated configuration best practices for running SAP landscapes on the following Dell EMC products: PowerEdge™ servers, Connectrix™ switches, Unity storage systems, and Data Domain™ storage protection systems.

Key benefits

PowerEdge servers that are paired with the Unity storage system power this Dell EMC Ready Bundle for SAP Landscapes solution. Customers implementing the solution can expect the following benefits:

- **Agility**—A modern SAP landscape management experience that delivers automated provisioning capabilities for SAP applications provides faster time to value.
- **Engineering**—Compute, networking, and storage are integrated with the 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 describes how to plan and prepare an IT infrastructure consisting of Dell EMC servers and storage for your SAP landscape. The guide provides validated best practices for sizing and designing a virtualized SAP landscape using PowerEdge servers and Connectrix switches with Unity data storage and Data Domain storage protection.

Terminology

Table 1 defines some of the terms that are used in this guide.

Table 1. Terminology

Term	Definition
Application lifecycle management (ALM)	A change process that covers all aspects of development, testing, deployment, and maintenance for the entire life of the product
Enterprise resource planning (ERP)	The collective term for the functional and technical modules that enable enterprises to manage business processes through a unified system
Business intelligence (BI)	The applications and technologies that are used by enterprises for the analysis of business information
High performance computing (HPC)	The application of computing power to solve large-scale business problems in enterprises

Audience

This design 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 of this guide should have technical experience with virtualization, servers, operating systems, networking, and storage solutions.

We value your feedback

Dell EMC and the authors of this guide welcome your feedback on the solution and the solution documentation. Contact EMC.Solution.Feedback@emc.com with your comments.

Authors: Donagh Keeshan, Fergal Murphy, Aighne Kearney

Solution components

Overview

SAP is one of the world's largest vendors of ERP applications. 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 systems (DEV), quality assurance systems (QAS), and production systems (PRD). The Dell EMC Ready Bundle for SAP Landscapes solution encompasses a variety of design configurations and deployment options that are tailored to individual business requirements and preferences. 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 Bundle for SAP Landscapes 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, software, and services. These products and solutions are designed for reliability and scalability and are engineered to handle highly demanding business applications and workloads.

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

Dell EMC PowerEdge servers

PowerEdge servers, which are built to support the work of IT organizations, include features to better run applications such as databases, ERP, BI, high performance computing (HPC), and data warehousing. Customers running SAP landscapes have the benefit of knowing that this solution includes high-performing servers, as documented in the [SAP Standard Application Benchmarks](#).

At the time of publication of this guide, the PowerEdge R940 and R740/R740xd servers have achieved the following:

- **R940 server** (with four 28-core Intel Xeon Platinum 8180 processors)— Best-of-published four-socket SAP Sales and Distribution (SD) 2-tier benchmark results on Linux (62,500 benchmark users, 341,100 SAP Application Performance Standard [SAPS]). [Result published July 11, 2017](#). Certification [#2017018](#).
- **R740/R740xd server** (with two 28-core Intel Xeon Platinum 8180 processors)— Best-of-published two-socket SAP SD 2-tier benchmark results on Linux (32,085 benchmark users, 175,230 SAPS). [Result published July 11, 2017](#). Certification [#2017017](#).

Based on these SAP SD benchmark results, for a single SAP system, the PowerEdge R940 and R740/R740xd servers offer the following scalability:

- In a two-tier SAP architecture on a single server:
 - R940 server: 341,100 SAPS
 - R740/R740xd servers: 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):
 - R940 server: 1,705,500 SAPS
 - R740/R740xd servers: 876,150 SAPS

For an explanation of SAPS values, see [SAPS overview](#).

PowerEdge R940 server

The PowerEdge R940 server is designed to support mission-critical applications and real-time decisions. With four sockets and up to 12 Non-Volatile Memory Express (NVMe) drives, the R940 provides scalable performance in just 3U.

Table 2 provides the technical specifications for R940 servers.

Table 2. Technical specifications for PowerEdge R940 servers

Feature	Technical specifications
Processor	Up to 4 Intel Xeon SPs, up to 28 cores
Memory	<ul style="list-style-type: none"> • 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
Storage controllers	<ul style="list-style-type: none"> • Internal controllers: PERC S140, H330, H740P Boot Optimized Storage Solution (BOSS) subsystems • 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, for a maximum of 122 TB
Power supplies	<ul style="list-style-type: none"> • Platinum: 1100 W, 1600 W, 2000 W, 2400 W, 336 to 380 VDC • Gold: 1100 W, 48 VDC • Hot-plug power supplies with full redundancy option, 8 hot-plug fans with full redundancy
Dimensions	<ul style="list-style-type: none"> • Form factor: Rack (3U) • Chassis depth: 777.05 mm

Feature	Technical specifications
I/O and ports	<ul style="list-style-type: none"> Network daughter card options: <ul style="list-style-type: none"> 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 3.0, dedicated indirect Dell Remote Access Controller (iDRAC) Direct USB Rear ports: Video, serial, 2 x USB 3.0, dedicated iDRAC network port Video card: VGA Up to 13 x Gen3 slots (3 x 8 + 10 x 16)
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 <p>For specifications and interoperability details, see Operating Systems.</p>

The PowerEdge R940 server's scalable business architecture can process the most business-critical workloads. Combined with up to 6 TB of memory and 13 PCIe Gen3 slots, the R940 server has all the resources that are needed to maximize application performance and scale for future demands. The R940 server:

- Maximizes storage performance with up to 12 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

PowerEdge R740/R740xd server

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

Table 3 shows the technical specifications for R740 and R740xd servers.

Table 3. Technical specifications for PowerEdge R740 and R740xd servers

Features	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

Features	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, 1100 W, 1600 W, and 2000 W 48 VDC 1100 W, 380 HVDC 1100 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 <p>For specifications and interoperability details, see Operating Systems.</p>

Dell EMC Unity storage systems

This Dell EMC Ready Bundle for SAP Landscapes solution uses Unity arrays for data storage. Unity technology enables affordable all-flash performance solutions and cost-effective hybrid solutions by using a midrange system that is optimized for simplicity across the storage lifecycle—that is, the system is simple to acquire, deploy, manage, and

service. The Unity storage system is ideal for midsize deployments, remote office/branch office (ROBO) locations, and cost-sensitive, mixed-workload environments.

Dell EMC Unity technology is designed for all-flash arrays and is available for converged deployment, as a hybrid system, and as a software-defined virtual edition.

Figure 1 shows the Unity product family.




Figure 1. Dell EMC Unity product family

You can order and configure the Unity platform as an All Flash or Hybrid system. The All Flash models (350F to 650F) and Hybrid models (300 to 600) are identical in form factor and connectivity but scale correspondingly in processing and memory capabilities.

Note: Unity Hybrid configurations support a maximum of 50 percent total flash capacity. If you require more capacity, consider a Unity All Flash configuration.

The All Flash models feature additional processor cores and more system memory. The system limits change depending on the Unity model. The Unity platform uses the latest technology in multicore processors, providing greater performance capabilities. Dell EMC recommends maximizing the use of flash drives in all Unity systems to optimize their performance potential. For more information, see the data and specification sheets on the [Dell EMC Unity Technical Documentation](#) web page.

Table 4 provides a detailed comparison of the Unity Hybrid models.

Table 4. Unity Hybrid model comparison

Feature	Unity 300	Unity 400	Unity 500	Unity 600
Processor	Intel E5-2603	Intel E5-2630	Intel E5-2660	Intel E5-2680
Memory per storage processor (SP)	24 GB	48 GB	64 GB	128 GB
Maximum number of drives	150	250	350	500
Maximum capacity (raw)	1.5 PB	2.5 PB	3.5 PB	5.0 PB

Table 5 provides a detailed comparison of the Unity All Flash models.

Table 5. Unity All Flash model comparison

Feature	Unity 350F	Unity 450F	Unity 550F	Unity 650F
Processor	Two 6-core, 1.7 GHz	Two 10-core, 2.2 GHz	Two 14-core, 2.0 GHz	Two 14-core, 2.4 GHz
Memory per SP	48 GB	64 GB	128 GB	256 GB
Maximum number of drives	150	250	500	1,000
Maximum capacity (raw)	2.4 PB	4.0 PB	8.0 PB	16.0 PB

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 Landscapes with Dell EMC Unity Storage Arrays Deployment Guide*, which you can download from the [Dell EMC website](#).

Dell EMC Data Domain

Dell EMC Data Domain deduplication storage systems enable fast, reliable disk backup, archiving, and disaster recovery with high-speed, inline deduplication. The new generation of large-enterprise and midsize Data Domain systems are powered with a flash solid-state drive (SSD). By consolidating backup and archive data on a Data Domain system, you can reduce storage requirements by 10 to 55 times, making disks cost-effective for onsite retention and highly efficient for network-based replication to disaster recovery sites.

Data Domain data protection systems work seamlessly with a range of backup, archive, and enterprise applications. Because the Dell EMC DD Boost™ library sends only unique data (instead of all the data in the database) to the Data Domain system, using Data Domain with DD Boost for Databases and Applications (DDBDA) provides the following benefits:

- Greater throughput
- Significantly reduced network bandwidth requirements

Table 6 shows the specifications for Data Domain deduplication storage systems.

Table 6. Data Domain controller performance and capacity specifications

Feature	DD3300	DD6300	DD6800 ²	DD9300 ²	DD9800 ²
Maximum throughput per hr	4.2 TB	8.5 TB	14 TB	20 TB	31 TB
Maximum throughput (DD Boost) per hr	7 TB	24 TB	32 TB	41 TB	68 TB
Logical capacity ¹	200 TB to 1.6 PB	1.8 to 8.9 PB	2.8 to 14.4 PB	7.2 to 36 PB	10 to 50 PB
Maximum usable capacity	Up to 32 TB	Up to 178 TB	Up to 288 TB	Up to 720 TB	Up to 1 PB

¹ A mix of typical enterprise backup data (file systems, databases, email, developer files). The low end of the capacity range represents a full weekly or monthly backup or an incremental daily or weekly backup to system capacity. The top end of the range represents a full daily backup to system capacity. All capacity values are calculated using Base10 (that is, 1 TB = 1,000,000,000,000 bytes).

² Supports HA active/standby configuration.

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 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 6.5 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 6.5, you can run, manage, connect, and secure your applications in a common operating environment across clouds and devices.

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, one QAS, and one PRD system.

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

Many customers supplement a three-system landscape with additional environments, such as a sandbox (SBX) that is used for testing, learning, and destructive testing.

Figure 2 shows a typical SAP system landscape for a single application.




Figure 2. 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 total cost of ownership (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 the Dell EMC Ready Bundle for SAP Landscapes solution with Unity storage and the design of the infrastructure subsystems.

Figure 3 shows the solution architecture.

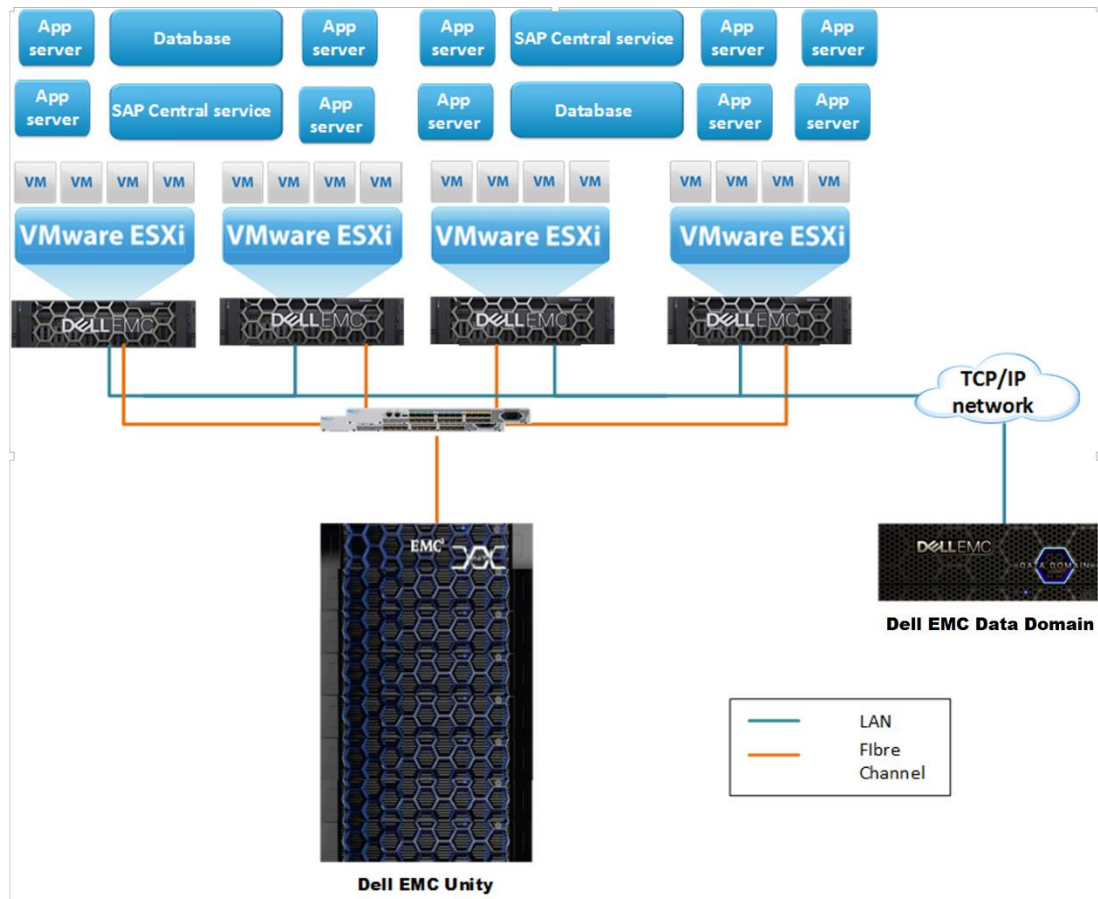


Figure 3. Solution architecture of Ready Bundle for SAP Landscapes with Unity Storage Array

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

The Ready Bundle for SAP Landscapes 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 HA.

Note: The Ready Bundle for SAP Landscapes 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.

Dell EMC Networking S4048-ON switches provide 10 GbE network connectivity between the compute cluster and the rest of the data center, while the Dell EMC S3148P switch provides 1 GbE connectivity.

To provide HA 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 S3148P switches, while two 10 GB NICs also connect to redundant S4048-ON 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 datacenter core network, for example, as DEV, QAS, or PRD.

Storage architecture

Unity All Flash storage systems are based on Intel E5-2600 processors. The systems implement an integrated architecture for block, file, and VMware Virtual Volumes (VVols), with concurrent support for native network-attached storage (NAS), iSCSI, and FC protocols. Each system uses dual SPs, full 12 Gb SAS back-end connectivity, and Dell EMC's patented multicore operating environment to deliver high performance and efficiency. Disk-array enclosures provide additional storage capacity, and online controller upgrades are available for added performance.

The Ready Bundle for SAP Landscapes uses the Unity All Flash array with the following storage profile:

- Two SPs, for HA and load balancing
- All-flash drives
- Shared physical memory
- Thin provisioning
- RAID 5
- Dynamic pools
- Compression enabled

Data protection

Backup, recovery, and archiving

Whether your company is a small-to-midsize or a large enterprise, the smooth functioning 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

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.

Proven results with Data Domain

Using Data Domain systems in SAP environments enables you to:

- 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 compression across your SAP landscapes
- Simplify system management
- Protect multiple database copies efficiently and reliably
- Provide network-efficient replication and reliable recovery
- Lower TCO

Infrastructure-level high availability

Overview

SAP applications are responsible for most business processes, and any 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, requiring highly available servers, network paths, switches, and storage
- Redundancy, which must be built into any solution supporting business-critical applications

To ensure HA and a minimal component failure rate, Dell EMC recommends using components that have a high mean time between failure (MTBF).

The Ready Bundle for SAP Landscapes solution combines a redundant design, RAID for storage drives, dual controllers, and multiple paths to provide storage HA. 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 Ready Bundle for SAP Landscapes solution is also virtualized, HA 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 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 HA 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 HA 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). FT 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. FT provides protection for extremely critical VMs, where any loss of service is intolerable.

You can combine HA 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

Solution configuration best practices

Overview

This section describes the best practices that Dell EMC recommends for configuring the Ready Bundle for SAP Landscapes solution.

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 hyperthreading.

Local area network

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

- Use two 10 GbE S4048-ON 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 HA and prevent a SPOF. In our testing, we used two Broadcom 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 HA.
- Use NIC teaming to provide HA 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 HA.
- Enable HA by making crisscross cable connections between redundant FC HBA cards and redundant FC switches.
- Use name zoning of end devices (FC HBAs) and Unity 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 Unity ports. Where possible, the Unity ports should be on different SPs. Where possible, avoid creating a bandwidth mismatch between a host and the SPs it accesses, such as a single 16 Gb/s FC host initiator accessing four 16 Gb/s FC SPs. 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 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

The Dell EMC Unity array is a virtually provisioned, flash-optimized storage system that is designed for ease of use. This section introduces specific configuration recommendations to enable you to obtain good performance from your Unity array.

The main design principles supporting good performance are:

- Maximize flash capacity
- Distribute workloads across available storage resources
- Simplify the configuration
- Design for resilience
- Maintain the latest released Unity Operating Environment (OE) version

Flash drives

Dell EMC storage arrays use multiple types of drives with varying capacities and performance levels. The arrays try to match data with the optimal drive type based on storage profiles. Typically, mission-critical data is stored on high-performance drives, while less-critical data is moved to high-capacity drives.

Unity all-flash technology supports high-capacity flash drives that provide very high IOPS per terabyte (IOPS/TB). By taking advantage of the increases in flash memory densities and economies of scale that industry-standard flash-drive technology provides, Unity all-flash technology offers a differentiated capability with respect to many all-flash alternatives.

Storage provisioning

As a fully unified storage system, the Unity array offers both block and file provisioning in the same enclosure. The system provisions disks into storage pools that can host both block and file data and supports connectivity for both block and file protocols. Dell EMC recommends the following best practices:

- **Block connectivity**—Use iSCSI or FC to access LUNs, consistency groups, VMware datastores, and VMware VVols.
- **File connectivity**— Use NAS servers to host file systems that are accessed using Server Message Block (SMB) shares or Network File System (NFS) shares. NAS servers can also host VMware NFS datastores.

Dynamic pools

Unity OE version 4.2 on Unity All Flash models supports dynamic pools. Dynamic storage pools apply RAID to groups of drive extents from drives within the pool and allow for greater flexibility in managing and expanding the pool. Dynamic pools must be all-flash pools and cannot be built with HDDs (SAS and NL-SAS).

Dynamic pools do not require hot spares. A dynamic pool automatically reserves spare space in the pool at a rate of 1 drive's worth of space for every 32 drives. If a drive failure

occurs, the data that was on the failed drive is rebuilt into the spare capacity on the other drives in the pool.

RAID specification

The Unity array applies RAID protection to the drives in a storage pool. 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 1, RAID 5, and RAID 6. Each RAID type provides different performance, capacity, and protection levels, as follows:

- **RAID 1**—Provides data protection by mirroring data onto another physical disk, producing better performance with minimal or no performance impact in the event of a disk failure.
- **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 Dell EMC FAST™ for Virtual Pools (FAST VP) environment, Dell EMC recommends RAID 5 for flash, RAID 1 for FC, and RAID 6 for SATA. Table 7 shows how RAID levels compare in random read, random write, sequential read, and sequential write scenarios.

Table 7. Comparison of RAID levels

RAID level	Random read	Random write	Sequential read	Sequential write
RAID 1	Excellent	Excellent	Excellent	Excellent
RAID 5	Excellent	Good	Good	Moderate
RAID 6	Good	Poor	Good	Moderate

Unity global memory

Global memory is a crucial component in the Ready Bundle for SAP Landscapes solution architecture. All read and write operations are transferred to or from global memory. The processing of transfers between the host processor and SPs occurs at much greater speeds than transfers involving physical drives. Unity OE uses complex, intelligent, statistical pre-fetch algorithms, which adjust to workload on the array by constantly monitoring, evaluating, and optimizing cache decisions.

Unity All Flash systems share a number of similarities in form factor and connectivity but scale in processing and memory capabilities from 48 GB per SP to 256 GB per SP.

Front-end connectivity

The Unity array provides multiple options for front-end connectivity through on-board ports directly on the disk processor enclosure (DPE) and through optional I/O modules. In general, Dell EMC recommends that you:

- Connect and configure all required front-end ports symmetrically across the two SPs to facilitate HA and continued connectivity in case of SP failure.
- For best performance, use all front-end ports that are installed in the system so that workload is spread across as many resources as possible.

For example, when configuring the four-port FC I/O module, zone different hosts to different ports so that all eight ports across the two SPs are used. Do not simply zone all hosts to the first port of each I/O module.

Fibre Channel

For FC, configure Dell EMC Unity CNA ports and I/O module ports with 8 Gb or 16 Gb small form-factor pluggable (SFP) transceivers. All FC ports can negotiate to lower speeds. Dell EMC recommends 16 Gb FC for the best performance.

Dell EMC recommends that you use single-initiator zoning when creating zone sets. For HA purposes, zone a single host initiator to at least one port from SP A and one port from SP B. For load balancing on a single SP, zone the host initiator to two ports from SP A and two ports from SP B. When zoning additional host initiators, where possible, zone them to different SP ports to spread the load across all available SP ports.

On hosts connected with FC, use multipathing software such as the Dell EMC PowerPath™ solution. The PowerPath software coordinates with the Unity system to provide path redundancy and load balancing.

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 (Unity front-end ports). Therefore, each HBA port requires a separate FC zone with Unity front-end ports. Create independent zones for each HBA installed in the host.
- WWN zoning contains only the host HBA port and the Unity front-end primary ports. In most cases, you do not have to include the Unity 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 Unity 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.

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.

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

- Place data and log files on separate LUNs.
- 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 HA and map the volumes to the server cluster. Related initiators (HBAs) are grouped by host, while related devices are grouped by storage pool.

NAS servers

NAS servers host file systems on the Unity storage system. NAS servers use virtual interfaces to enable host connectivity to SMB, NFS, and multiprotocol file systems, as well as VMware NFS datastores and VVols. Depending on what is enabled on the NAS server, you can create SMB file systems and NFS file systems separately or in a multiprotocol configuration. File systems and virtual interfaces are isolated to a single NAS server, allowing for multitenancy over multiple NAS servers. NAS servers are hosted on an SP and automatically fail over if the SP becomes faulted. Any associated file systems also fail over.

Unity file systems

File systems are file-based storage resources that are hosted on NAS servers. They are accessed through shares, which can be provisioned for SMB or NFS access, providing access for Windows or UNIX hosts. Depending on what protocol is selected for your file system, only shares of that type can be provisioned. For example, for an SMB file system, you can provision SMB shares, but not NFS shares. To provision both SMB and NFS shares from a file system, use a multiprotocol file system.

File systems can be extended and shrunk in size, and allocated space is automatically taken and reclaimed based on file system usage patterns. Finally, file systems can be configured with quotas to better regulate file system space usage. Tree quotas and user quotas are supported on file systems and can be used in combination.

Global file systems

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

- 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 and storage-based NFS. Unity file systems can provide storage-based NFS to satisfy the SAP global file system requirements.

For more information, see [DELL EMC UNITY: NAS CAPABILITIES](#).

Unity compression

To help reduce the total cost of ownership (TCO) and increase the efficiency of a Unity storage system, Dell EMC added compression in Unity OE version 4.1, with initial support for "thin block" storage resources. Unity OE version 4.2 and later versions include support for "thin clones" and "thin file" storage resources.

With compression, a smaller amount of space is required to store a dataset for compression-enabled storage resources. This space savings reduces the amount of physical storage that is required to store a dataset, which can yield cost savings. Space savings are also realized on the storage resources.

Snapshots and thin clones

Snapshots and thin clones inherit the compression setting of the source storage resource, which helps increase the space savings that they can provide.

For more information, see the [Dell EMC Unity: Compression White Paper](#).

Competing workloads

When adding workloads to the Unity system, consider the reported CPU utilization rates as well as the capacity and IOPS sizing. Brief spikes of high utilization are normal and expected on any Unity system. For information about sustained CPU utilization and the operating ranges that influence whether the system can accept additional workloads, see the [Dell EMC Unity: Best Practices Guide](#).

In highly consolidated environments, databases and applications compete for storage resources. The Host I/O Limits quality of service (QoS) feature controls "noisy neighbors" and protects the performance of the production system. You can apply host I/O limits on LUNs, VMware vStorage, VMware NFS datastores, and their associated snapshots to limit the number of IOPS serviced by the Unity system.

Use host I/O limits to limit incoming host activity on the basis of IOPS, bandwidth, or both. You can enforce limits on individual resources, or you can share a limit among a set of resources.

Host I/O limits can be effective if an overload on the storage resources affects the performance of critical applications such as SAP production installations. Consider configuring host I/O limits on LUNs, datastores of non-SAP applications, or SAP nonproduction systems to limit the total IOPs or bandwidth, which might otherwise consume a large portion of the system's resources and affect the performance of SAP production systems.

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 DD Boost 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, precompression, 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. Some backup applications enable these features by default, so ensure that none of them are enabled for your application.
- Avoid selecting the highest available values for configuration options (number of streams, throttling, replication, and so on). Using 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 Oracle RMAN accesses directly 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 Bundle for SAP Landscapes with Dell EMC Unity Storage Arrays Deployment Guide*, which you can download from the [Dell EMC website](#).

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 (nonwide).

- 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 FT, which supports up to four vCPUs.
- Ensure that you use the latest SAP host agent.
- If you want site HA, configure two or more ESXi hosts in the vSphere cluster.
- Ensure that vSphere HA and 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 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.
- 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

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

Figure 4 shows our storage design layout for VMware VMs.

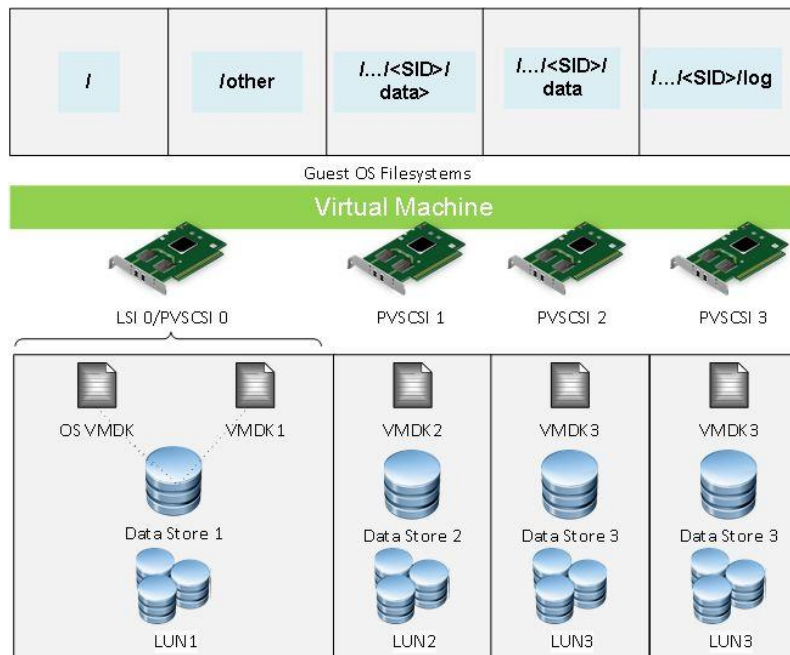


Figure 4. Storage design layout for VMware VMs

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 FT. 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 operating system failures.

Sizing and performance

Sizing considerations for SAP landscape design

Business and technology considerations influence the sizing of the hardware infrastructure. When planning to use Dell EMC infrastructure, work with Dell EMC representatives to determine a solution design that includes business requirements for performance, response times, availability, data protection, and DR for the SAP systems.

Table 8 describes the sequence of steps that are involved in sizing the infrastructure requirements for your Ready Bundle for SAP Landscapes solution.

Table 8. 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.	
4	Determine the HA requirements.	
5	Determine the DR requirements and the number of data centers or sites involved.	
6	Consider data protection requirements for backing up the SAP systems.	
7	Determine an expected annual data growth and the number of years' 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 Unity Sizer tool to determine the model and size of the Unity storage system.	
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 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](#) (SAP Marketplace access is required).

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 box 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 Bundle for SAP Landscapes solution uses the published SAPS values for the PowerEdge R940 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 R940 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 the chosen R940 or R740/R740xd PowerEdge server model by using the SAPS values.

Table 9 provides details of certified SAPS values for PowerEdge servers. For a full list of PowerEdge R940 and R740/R740xd systems with extrapolated SAPS values, see [Appendix: SAPS values for PowerEdge R940 and R740/R740xd](#).

Table 9. 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 R940 (Intel Xeon Platinum 8180, 2.50 GHz)	112	4	28	341,100	3,046
PowerEdge R740 or R740xd (Intel Xeon Platinum 8180, 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 the operating system, 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 Dell EMC 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. Therefore, 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](#) (search for “Sizing Guidelines” on the [SAP Service Marketplace](#); SAP Marketplace access is required)

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. You can achieve HA and load balancing of VMs with vSphere DRS 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 DR plan must have your normal operations up and running in as short a recovery 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 DR.

Sizing for future database growth

SAP systems will 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:

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

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 system has a system landscape consisting of DEV, QAS, and SBX non-PRD systems. Based on its Quick Sizer projects, the customer provided the information in Table 10 for each PRD system.

Table 10. 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 uses VMware virtualization for server consolidation and HA.

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.

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.

Table 11 shows the SAP landscapes that we derived from the customer inputs.

Table 11. SAP landscapes and customer inputs from Quick Sizer

System	Application	Type	Description	SAPS (input)	Disk size (input) (GB)	Annual disk growth (input)
BM1	ERP	OLTP	PRD	96,000	4,000	7%
QM1	ERP	OLTP	QAS	96,000	4,000	7%
DM1	ERP	OLTP	DEV	48,000	2,000	7%
SBX	ERP	OLTP	SBX	48,000	2,000	7%
CRM	CRM	OLTP	PRD	32,000	1,000	7%
CRQ	CRM	OLTP	QAS	32,000	1,000	7%
CRD	CRM	OLTP	DEV	16,000	500	7%
CRS	CRM	OLTP	SBX	16,000	500	7%
BWP	BW	OLAP	PRD	60,000	2,000	7%
BWQ	BW	OLAP	QAS	60,000	2,000	7%

System	Application	Type	Description	SAPS (input)	Disk size (input) (GB)	Annual disk growth (input)
BWD	BW	OLAP	DEV	30,000	1,000	7%
BWS	BW	OLAP	SBX	30,000	1,000	7 %

Sizing the compute requirements

We used the inputs from Table 11 to calculate the compute and memory requirements, as shown in Table 12.

Table 12. CPU and memory calculations for each system in the landscapes

System	SAPS (input)	Database SAPS ¹	Application SAPS ¹	VMware CPU over-commit 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	96,000	19,200	76,800	1	19,200	76,800	57.6	230.4	5
QM1	96,000	19,200	76,800	2	9,600	38,400	57.6	230.4	5
DM1	48,000	9600	38,400	4	2,400	9,600	28.8	115.2	2
SBX	48,000	9600	38,400	4	2,400	9,600	28.8	115.2	2
CRM	32000	6,400	25,600	1	6,400	25,600	19.2	76.8	3
CRQ	32,000	6,400	25,600	2	3,200	12,800	19.2	76.8	3
CRD	16000	3,200	12,800	4	800	3,200	9.6	38.4	1
CRS	16,000	3,200	12,800	4	800	3,200	9.6	38.4	1
BWP	60,000	20,000	40,000	1	20,000	40,000	60	120	3
BWQ	60,000	20,000	40,000	2	10,000	20,000	60	120	3
BWD	30,000	10,000	20,000	4	2,500	5,000	30	60	1
BWS	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.

Using the customer inputs in Table 10, we calculated the total SAPS and memory requirements. Table 13 shows the results of our calculations.

Table 13. Total compute requirements

Total database SAPS	Total application SAPS	Total SAPS	Total database memory (GB)	Total application memory (GB)	Total memory (GB)
79,800	249,200	329,000	410.4	1,281.6	1,692

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

Additional compute resources for HA are included. Therefore, if a server is lost, sufficient resources are available to take over the failed server workloads.

In this example, we calculated the result using both the PowerEdge R940 and PowerEdge R740/R740xd models. Table 14 shows the physical compute requirements for the solution using the R940 server on each site.

Table 14. Physical compute requirements with the PowerEdge R940 server

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
329,000	1,692	306,990	112	768	2	3	3	4

¹ In this sizing example, we are using PowerEdge R940 servers with 768 GB of RAM.

² The PowerEdge R940 has the capabilities of up to 306,990 virtual SAPS, and the SAPS requirement in the example is 329,000. Therefore, two servers are needed to satisfy the requirement.

³ The PowerEdge R940 server in the example has 768 GB of RAM and the compute requirement is 1,692 GB of RAM. 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 HA.

Table 15 shows the physical compute requirements for the solution using the R740 or R740xd server on each site.

Table 15. Physical compute requirements with the PowerEdge R740 or R740xd server

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 HA ⁵
329,000	1,692	157,707	56	384	3	5	5	6

¹ In this sizing example, we are using PowerEdge R940 servers with 384 GB of RAM.

² The PowerEdge R940 has the capabilities of up to 157,707 virtual SAPS, and the SAPS requirement in the example is 329,000. Therefore, three servers are needed to satisfy the requirement.

³ The PowerEdge R940 server in the example has 384 GB of RAM and the compute requirement is 1,692 GB of RAM. Therefore, five 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 HA.

Sizing storage requirements

Table 16 presents the storage requirements for each production and nonproduction SAP system.

Table 16. Customer storage requirements

System	Disk size (input) (GB)	DB disk final size (GB) ¹	Database SAPS (with VMware overcommit)	IOPS (calculated) ²	Number of app servers	Total operating system size (GB) ³	Total .exe file size (GB) ³	Dump size (GB) ⁴	Total storage requirements ⁵
BM1	4,000	5,243	19,200	11,520	5	768	768	5,243	12,022
QM1	4,000	5,243	9,600	5,760	5	768	768	5,243	12,022
DM1	2,000	2,622	2,400	1,440	2	384	384	2,622	6,012
SBX	2,000	2,622	2,400	1,440	2	384	384	2,622	6,012
CRM	1,000	1,311	6,400	3,840	3	512	512	1,311	3,646
CRQ	1,000	1,311	3,200	1,920	3	512	512	1,311	3,646
CRD	500	655	800	480	1	256	256	655	1,822
CRS	500	655	800	480	1	256	256	655	1,822
BWP	2,000	2,622	20,000	18,000	3	512	512	2,622	6,268
BWQ	2,000	2,622	10,000	9,000	3	512	512	2,622	6,268
BWD	1,000	1,311	2,500	2,250	1	256	256	1,311	3,134
BWS	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 From calculation of database SAPS based on the analytical or transactional consumption of SAPS. See [Sizing the disk and disk I/O](#). 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 operating system and 128 GB for .exe file on each DB and application server of a system

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

⁵ Sum of the DB final disk size + operating system size + .exe file size + dump size values per system.

Table 17 presents the total storage requirements.

Table 17. Total SAP storage requirements

SAP Landscapes	Storage (GB)	IOPS
Total SAP systems	65,805	58,380

We use the storage and IOPS totals in the Unity Sizer tool to determine the model and configuration of the Unity All Flash array.

Conclusion

The Dell EMC Ready Bundle for SAP Landscapes is designed for an SAP ecosystem. The solution enables application owners to accelerate application lifecycle management because it provides the integration and features to provision and manage SAP ecosystems 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

The Ready Bundle for SAP Landscapes 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
- Ability to bring other workloads to the Unity storage system
- 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 documentation on [Dell EMC.com](#) and [Dell EMC Online Support](#) provides additional, relevant information. Access to documents on Online Support 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 with Dell EMC Unity Storage Arrays Deployment Guide*
- [Dell EMC Unity: Introduction to the Platform](#)
- [Dell EMC Unity: Best Practices Guide](#)
- [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](#)

VMware documentation

The following documentation provides additional, relevant information:

- [Architecture Guidelines and Best Practices for Deployments of SAP HANA on VMware vSphere](#)
- [SAP and VMware Virtualization--Overview](#)
- [SAP Solutions on VMware Best Practices Guide](#)
- [Understanding Oracle Certification, Support and Licensing for VMware Environments](#)

SAP documentation

The following website provides access to additional, 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 R940 and R740/R740xd

Table 18 and Table 19 show the SAPS values for Dell EMC PowerEdge R940 and R740/R740xd servers, respectively. These SAPS values were certified with SAP for the top-end platinum CPU model, the Intel Xeon Platinum 8180 (presented in bold in Table 18 and Table 19). Dell EMC extrapolated the SAPS values for all other PowerEdge server CPU models from these results. All the listed models are SAP HANA-certified with [SAP HANA tailored data center integration \(HANA TDI\) phase 5](#).

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 R940 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 R940 (Intel Xeon Gold 5115, 2.40 GHz)	40	4	10	120,697	3,017
PowerEdge R940 (Intel Xeon Gold 5118, 2.30 GHz)	48	4	12	144,312	3,006
PowerEdge R940 (Intel Xeon Gold 5120, 2.20 GHz)	56	4	14	161,367	2,882
PowerEdge R940 (Intel Xeon Gold 6126, 2.60 GHz)	48	4	12	165,302	3,444
PowerEdge R940 (Intel Xeon Gold 6130, 2.10 GHz)	64	4	16	191,541	2,993
PowerEdge R940 (Intel Xeon Gold 6132, 2.60 GHz)	56	4	14	192,197	3,432
PowerEdge R940 (Intel Xeon Gold 6134, 3.20 GHz)	32	4	8	133,160	4,161
PowerEdge R940 (Intel Xeon Gold 6136, 3.00 GHz)	48	4	12	187,605	3,908
PowerEdge R940 (Intel Xeon Gold 6138, 2.00 GHz)	80	4	20	219,091	2,739
PowerEdge R940 (Intel Xeon Gold 6140, 2.30 GHz)	72	4	18	221,059	3,070
PowerEdge R940 (Intel Xeon Gold 6142, 2.60 GHz)	64	4	16	215,155	3,362
PowerEdge R940 (Intel Xeon Gold 6148, 2.40 GHz)	80	4	20	244,674	3,058
PowerEdge R940 (Intel Xeon Gold 6150, 2.70 GHz)	72	4	18	243,362	3,380
PowerEdge R940 (Intel Xeon Gold 6152, 2.10 GHz)	88	4	22	250,577	2,847
PowerEdge R940 (Intel Xeon Gold 6154, 3.00 GHz)	72	4	18	265,008	3,681
PowerEdge R940 (Intel Xeon Platinum 8153, 2.00 GHz)	64	4	16	171,862	2,685
PowerEdge R940 (Intel Xeon Platinum 8158, 3.00 GHz)	48	4	12	188,261	3,922
PowerEdge R940 (Intel Xeon Platinum 8160, 2.10 GHz)	96	4	24	272,880	2,843
PowerEdge R940 (Intel Xeon Platinum 8164, 2.00 GHz)	104	4	26	280,752	2,700

PowerEdge server	Number of cores	Number of sockets	Number of cores per socket	SAPS	SAPS per core
PowerEdge R940 (Intel Xeon Platinum 8168, 2.70 GHz)	96	4	24	312,894	3,259
PowerEdge R940 (Intel Xeon Platinum 8170, 2.10 GHz)	104	4	26	289,935	2,788
PowerEdge R940 (Intel Xeon Platinum 8180, 2.50 GHz)	112	4	28	341,100	3,046

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	27,933	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	2,817
PowerEdge R740 (Intel Xeon Silver 4114, 2.20 GHz)	20	2	10	57,936	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	113,803	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

Appendix: SAPS values for PowerEdge R940 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 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
PowerEdge R740 (Intel Xeon Platinum 8180, 2.50 GHz)	56	2	28	175,230	3,129