

Dell EMC PowerMax: Family Overview

December 2021

H17118.4

White Paper

Abstract

This white paper provides an overview of the Dell EMC PowerMax family, an NVMe-based, mission-critical data storage offering. The paper details the theory of operation, packaging, and features that make PowerMax an ultra-performing, all-flash storage product.

Dell Technologies

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 © 2018-2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA December 2021 H17118.4.

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

Contents

- Executive summary.....4**
- Introduction5**
- PowerMax overview8**
- PowerMax architecture overview9**
- Reliability, availability, and serviceability.....29**
- PowerMax Data Services30**
- Streamlined software packaging.....41**
- PowerMax system deployments43**
- Summary.....49**
- Technical support and resources49**

Executive summary

Overview

The Dell EMC PowerMax family is the first Dell EMC hardware platform that uses an end-to-end Non-Volatile Memory Express (NVMe) architecture for customer data. NVMe is a set of standards which define a PCI Express (PCIe) interface used to efficiently access data storage volumes based on Non-Volatile Memory (NVM) media, which includes modern NAND-based flash along with higher-performing Storage Class Memory (SCM) media technologies. The NVMe-based PowerMax was specifically created to fully unlock the bandwidth, IOPS, and latency performance benefits that NVM media offers to host based applications which are unattainable using the previous generation of all-flash storage arrays.

Revisions

Table 1. Revisions

Date	Description
April 2018	Initial release
October 2018	Updated for 5978 Service Release
September 2019	Incorporated major Q3 2019 platform updates
September 2020	Incorporated major PowerMaxOS Q3 2020 platform updates
December 2021	Updated template

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: James Salvatore

Note: For links to other documentation for this topic, see the [PowerMax Info Hub](#).

Introduction

PowerMax introduction and benefits

The Dell EMC PowerMax family offers unprecedented levels of performance and scale using next-generation Storage Class Memory (SCM) and high-speed SAN infrastructures. PowerMax is powerful, simple, and trusted storage without compromise. It is built for mission-critical apps of today and tomorrow with end-to-end NVMe, next-gen storage media (SCM), real-time machine learning and inline deduplication and compression, while also delivering the features and data services businesses require.

Primary benefits

The primary benefits that PowerMax platforms offer to customers include the following:

- A powerful end-to-end NVMe storage architecture that delivers:
 - Up to 15M IOPS, 350 GBps throughput (187 K IOPS per rack unit)
 - Industry-standard NVMe-Based flash and SCM drives
 - Native NVMe Drive Array Enclosures (DAEs)
 - Large-scale workload consolidation in which Open Systems and mainframe block storage can co-exist with file storage on the same platform
- Integrated, real-time machine-learning engine for automatic data placement
 - Automated I/O recognition and data placement across flash and SCM media to maximize performance with no management overhead
 - Elimination of high-performance silo and consolidation of all mission-critical workloads and secondary applications
- Enterprise grade storage security and protection
 - End to end efficient data Encryption
 - FIPS 140-2 validated Data at Rest Encryption (D@RE)
 - Secure snapshots, role-based authentication, and tamper-proof audit logs
- Enterprise levels of reliability with proven six-nine's availability in a single array
- Investment protection with the Future Proof Program
- Global inline data deduplication and enhanced compression with virtually no performance impact
 - Works with all data services
 - Provides a 3.5:1 data reduction guarantee through Dell's Future Proof Program
- Powerful data services which help protect, manage, and move customer data on the array. These data services include remote replication with SRDF, High Availability with SRDF/Metro, Local replication with TimeFinder SnapVX, and Cloud Mobility
- A comprehensive easy to use API covering storage provisioning, all configurable data services, array configuration, performance monitoring and alerting.

Note: For more information about these PowerMax features and their value propositions, see the [Dell EMC PowerMax Family web page](#).

PowerMaxOS Q3 2020 release

The newest enhancements featured in the PowerMax Q3 2020 release expand the PowerMax offering by adding cloud mobility, data resiliency, enabling SRDF replication for VMware vSphere Virtual Volumes (vVols), continuous high availability for SRDF/Metro configurations, and continued security hardening of the system. Some of the key features in the Q3 2020 release key features are shown in the following table:

Table 2. Q3 2020 PowerMaxOS release core features

Feature	Value proposition and benefits
Cloud Mobility for Dell EMC PowerMax	Extends PowerMax storage to the cloud (public or private) for long-term retention. Snapshots can be shipped to object stores AWS, Microsoft Azure, and Dell EMC ECS.
VMware vVols replication with SRDF	Combine the gold standard in storage replication (SRDF) with VMware vVols for mission-critical operations - orchestrated using VMware Site Recovery Manager (SRM).
SRDF/Metro Smart DR	A space and bandwidth efficient metro-aware Disaster Recovery for a SRDF/Metro region.
Snapshot policies	Automated, compliance-based snapshot protection for applications.
25 GbE I/O module	Enhancing the PowerMax Ethernet SAN offerings with a new 4 port 25 GbE I/O module used for iSCSI and SRDF connectivity.
End-to-end efficient encryption	Provides complete encryption protection when data is written from the host to PowerMax storage media (drives). This solution has the added benefit of incorporating up to 5:1 data reduction resulting in a highly secure, highly efficient offering for our customers.
IBM Transparent Cloud Tiering (TCT) *	A licensed IBM function that offloads all data movement processing-related workloads from the mainframe host when moving data to or from cloud repositories.

* Supported on PowerMax 8000 only through RPQ process and requires MFE 8.5.

Terminology

The following table provides definitions for some of the terms that are used in this document:

Table 3. Terminology

Term	Definition
Automated Data Placement (ADP)	ADP is the system's ability to intelligently manage data placement between two different drive technologies in the same array.
DAE24	DAE24 refers to the drive array enclosure that is used to store up to 24 NVMe drives in PowerMax arrays.
PowerMax 2000	PowerMax 2000 is the entry NVMe scale-out array sold with the Essentials and Pro software packages.
PowerMax 8000	PowerMax 8000 is the flagship NVMe scale-out array sold with the Essentials and Pro software packages.
Disk group	A disk group is a collection of hard drives sharing the same technology, size, and performance characteristics.

Term	Definition
Drive Array Enclosure (DAE)	DAE refers to the drive array enclosure used to store flash drives and SCM drives in PowerMax.
Essentials	The Essentials package is the default PowerMax software package.
Flash capacity pack	A flash capacity pack includes NVMe flash drive capacity (TBu) that can be added to a PowerMax array.
Inline compression	Inline compression refers to the intelligent compression technology used with PowerMax arrays.
Inline deduplication (dedupe)	Inline deduplication refers to the deduplication technology used with PowerMax arrays.
Non-Volatile Memory Express (NVMe)	NVMe is a command set and its associated storage-interface standards that specify efficient access to data storage devices and systems based on Non-Volatile Memory (NVM).
NVMe flash drives/NAND	NVMe/PCIe-connected flash drives are the latest flash devices used to store capacity in PowerMax arrays.
NVMe over Fabric (NVMe-oF)	NVMe-oF defines a common architecture that supports a range of storage networking fabrics for NVMe block storage protocol
NVMe over Fibre Channel (FC-NVMe)	FC-NVMe extends the NVMe block storage protocol and its benefits over data-center fabrics using high-speed Fibre Channel as the fabric transport.
PowerMax Brick	A Brick is the building block for a PowerMax array. It includes an engine, two DAEs, and a fixed TBu of capacity.
PowerMax zBrick	zBrick is the PowerMax mainframe building block consisting of an engine, two DAEs, and fixed TBu of capacity.
PowerMaxOS/ PowerMaxOS 5978	PowerMaxOS is the PowerMax storage operating system. The PowerMaxOS 5978 release supports PowerMax NVMe arrays, dedupe, and other software enhancements. It can also be installed on legacy VMAX All Flash arrays.
Pro	The Pro package is a comprehensive software offering sold as a PowerMax appliance.
RAID group	A RAID group is the minimum quantity of hard drives that comprise a specific RAID protection scheme.
Scale out	Scale out refers to adding Bricks to grow performance and expansion on PowerMax systems.
Scale up	Scale up refers to adding flash capacity packs to a PowerMax array.
Service level	Service levels identify a specific performance tier in PowerMax arrays.
Smart RAID	Smart RAID provides active/active shared RAID support for PowerMax arrays.
Storage Memory Class (SCM)	SCM is a new hybrid storage/memory tier that has read and write performance characteristics which are significantly faster than traditional flash drives

Term	Definition
Storage Resource Pool (SRP)	An SRP is a collection of hard drives which constitute a specific usable capacity for customer application data. Most PowerMax arrays consist of a single SRP. An SRP's usable capacity can consist of both NAND Flash and SCM disks. Automated Data Placement across the SCM and NAND Flash drives within the SRP is managed by the internal PowerMax machine learning engine. Each SRP is protected by a single RAID protection scheme and can provide capacity for FBA, file, and mainframe.
TBe	TBe is effective capacity in terabytes. It includes the benefits of thin provisioning, inline compression, deduplication, and space-efficient copies.
TBu	TBu is usable capacity in terabytes. It refers to the amount of hard drive capacity available in the array, taking into account the RAID efficiency of the RAID type in use.
Unisphere	Unisphere for PowerMax is a user interface that enables management and monitoring of PowerMax arrays along with legacy VMAX All Flash, VMAX3, and VMAX 1 or 2 arrays. Unisphere for PowerMax also provides REST API interface for managing and controlling Dell EMC arrays.
zEssentials	The zEssentials package is the software package sold as a PowerMax appliance for mainframe.
zFlash	A zFlash capacity pack includes NVMe flash drive capacity (TBu) that can be added to a PowerMax array for mainframe.
zPro	The zPro package is the comprehensive software offering sold as a PowerMax appliance for mainframe.

PowerMax overview

PowerMax family [The Dell EMC PowerMax family](#) is built using a 100% end-to-end Non-Volatile Memory Express (NVMe) storage architecture, allowing it to reach unprecedented I/O densities and performance by eliminating the flash media choke points found using traditional SAS and SATA interfaces. The PowerMax array opens the door for customers to deploy innovative applications in the areas of real-time analytics, machine learning, and big data that demand lower latency and higher performance.

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

Both PowerMax arrays have at their foundation the trusted Dynamic Virtual Matrix architecture and internal system software specifically written for the NVMe platform called PowerMaxOS 5978. PowerMaxOS can run natively on both PowerMax systems and on legacy VMAX All Flash systems as an upgrade. As with the previous-generation VMAX All Flash, PowerMax systems are true all-flash arrays – products specifically targeted to meet the storage capacity and performance requirements of the all-flash enterprise data center. The PowerMax products are feature-rich, all-flash offerings with specific capabilities

designed to take advantage of ultra-high performing Storage Class Memory (SCM) and higher capacity NVMe flash drives to create the densest storage configuration possible. PowerMax 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 more traditional storage workloads. In addition, PowerMax now allows customers to deploy applications such as real-time analytics, machine learning, and big data that demand the lower storage latency and higher IOPS densities previously unattainable with legacy all flash offerings.



PowerMax 2000

- Up to 2.7M IOPS and 80 GBps per System
- Up to 96 NVMe Flash and SCM Devices
- Up to 64 Front End Ports
- 1 -2 Bricks per system
- Open Systems Workloads Only



PowerMax 8000

- Up to 15M IOPS and 350 GBps per System
- Up to 288 NVMe Flash and SCM Devices
- Up to 256 Front End Ports
- 1 -8 Bricks / zBricks per system
- Open Systems / Mainframe / Mixed Workloads

Figure 1. PowerMax family

PowerMax architecture overview

Architecture introduction

Although the PowerMax platform uses many of the technologies and data services found in legacy VMAX All Flash, PowerMax provides customers with a differentiating value as it is designed from the ground up to be the first platform in the industry to take full advantage of end-to-end FC-NVMe connectivity and emerging data storage media such as SCM. The following sections detail the key PowerMax architectural value propositions for customers.

Designed for NVMe

PowerMax is a technology leader providing a full end-to-end NVMe flash storage architecture for storing customer data. The PowerMax NVMe architecture provides:

- **I/O Density with Predictable Performance** – PowerMax has been designed to deliver extreme I/O density – capable of delivering approximately 187 K IOPS per rack unit (U) or up to 15M IOPS in a two-rack system (two floor tiles), regardless of workload and storage capacity utilization.
- **NVMe Storage Density** – Using commercially available, high capacity, dual-ported enterprise NVMe flash drives, PowerMax delivers outstanding NVMe TB per floor tile. PowerMax support for high capacity commercially available NVMe flash and SCM drives provides a differentiated capability as compared to many other all-flash alternatives which use a proprietary flash drive design. This allows PowerMax to

leverage the increases in flash drive densities, performance, economies of scale, and fast time-to-market delivered by the industry flash drive suppliers.

- **Future Proof Design** - The PowerMax NVMe design is future proof as it can seamlessly implement current and future SCM and NAND Flash drive capacities and allow for hosts to connect through current and future high-speed NVMe-oF SANs.

Expandable modular architecture: PowerMax Brick

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

There are two types of Bricks available for PowerMax:

- The open systems **Brick** supports configurations with Fibre Channel, FC-NVMe, or iSCSI connectivity with FBA device formatting. The Brick also can be configured for file storage using embedded NAS.
- The mainframe **zBrick** supports configurations with FICON connectivity and CKD device formatting.

Note: In this document, the term Brick is used when discussing features and functions applicable to both the open systems and the mainframe. When discussing features specific to mainframe, the term zBrick is specifically called out.

The initial system Brick includes a single engine consisting of two directors, two system power supplies (SPS), and two 24-slot 2.5" NVMe Drive Array Enclosures (DAE24) pre-configured with an initial total usable capacity.

The Brick concept allows PowerMax to scale up and scale out. Customers can scale up by incrementally adding Flash Capacity Packs. Each Flash Capacity Pack for the PowerMax 8000 has 13 TBu or 15 TBu of usable storage, and either 11 TBu, 13 TBu, or 15 TBu for the PowerMax 2000 model, depending upon the RAID protection type selected. PowerMax scales out by aggregating up to two Bricks for the PowerMax 2000, and up to eight for the PowerMax 8000. Scaling out a PowerMax system by adding additional Bricks produces a predictable, linear performance improvement regardless of the workload.

Note: For detailed information about available PowerMax Brick configurations, see the [PowerMax Family Specification Sheet](#).

Engines

The core of the Brick is the engine. The engine is the central I/O processing unit, redundantly built for high availability. Each Brick consists of:

- Redundant directors that contain multi-core CPUs and memory modules
- Interfaces to universal I/O modules, such as front-end, back-end, InfiniBand, and flash I/O modules

The communication backbone of the Brick is the trusted Dynamic Virtual Matrix Architecture. Fundamentally, the virtual matrix enables inter-director communications over redundant internal InfiniBand fabrics. The InfiniBand fabric provides a foundation for a

highly scalable, extremely low latency, and high-bandwidth backbone which is essential for an all flash array. This capability is also essential for allowing the PowerMax to scale upwards and scale outwards in the manner that it does.




Figure 2. Brick engine director

Brick CPU core configurations

Each Brick engine has two directors, with each director having dual CPU sockets which support multi-core, multi-threaded Intel processors. The following table details the engine CPU core layout for each PowerMax model:

Table 4. Brick engine CPU cores

PowerMax model	Engine CPU type	CPU cores	Cores per Brick engine	Max cores per system
PowerMax 2000	Dual Intel Broadwell, 12 core, 2.5 GHz	24	48	96 (2 Bricks max.)
PowerMax 8000	Dual Intel Broadwell, 18 core, 2.8 GHz	36	72	576 (8 Bricks max.)

The Brick engine uses a core pooling mechanism which can dynamically load-balance the cores by distributing them to the front end, back end, and data services (such as SRDF, eNAS, and embedded management) running on the engine. The core pools can be tuned to shift the bias of the pools at any time to front-end heavy or back-end heavy workloads to further optimize the solution for a specific use case.

Note: Due to the advanced cooling dynamics of the PowerMax engine, the Intel CPUs primarily run in Turbo mode, providing additional performance capabilities.

Brick cache configurations

Every director has 16 memory slots which can be populated with 32 GB and 64 GB DDR4 DIMMS to achieve up to 1 TB cache per director (2 TB cache maximum per Brick engine).

Table 5. Brick cache configurations

PowerMax model	Cache per Brick	Max cache per system
PowerMax 2000	512 GB, 1 TB, 2 TB	4 TB (2 Bricks max.)
PowerMax 8000	1 TB or 2 TB	16 TB (8 Bricks max.)

On single-engine PowerMax 2000 systems, cache is mirrored within the engine across the directors. This is also true for multi-engine PowerMax 2000 systems and single-engine PowerMax 8000 systems. On multi-engine PowerMax 8000 systems, cache is mirrored across directors in different engines for added redundancy.

Both the PowerMax 2000 and PowerMax 8000 can support engine configurations with differing cache sizes (mixed cache). For dual-engine PowerMax 2000 models, the system can use engines with different cache sizes between the engines which are one cache size smaller or larger than the other engine in the system. For example, cache on engine 1 can be 1 TB while the cache on engine 2 is 512 GB. This would yield a total cache size of 1.5 TB for the system. Valid mixed cache configurations for the PowerMax 2000 are shown in the following table:

Table 6. Supported PowerMax 2000 mixed cache configurations

Number of ricks in system	Smallest engine cache size	Largest engine cache size	Total system cache
2	512 GB	1 TB	1.5 TB
2	1 TB	2 TB	3 TB

Mixed cache configurations are available on the PowerMax 8000; but require a minimum of four Bricks or zBricks in the system. The following table details the supported mixed cache configurations available for the PowerMax 8000:

Table 7. Supported PowerMax 8000 mixed-cache configurations

Number of Bricks in system	Smallest engine cache size	Largest engine cache size	Total system cache
4	2 engines at 1 TB	2 engines at 2 TB	6 TB
5	2 engines at 1 TB	3 engines at 2 TB	8 TB
5	3 engines at 1 TB	2 engines at 2 TB	7 TB
6	2 engines at 1 TB	4 engines at 2 TB	10 TB
6	4 engines at 1 TB	2 engines at 2 TB	8 TB
7	2 engines at 1 TB	5 engines at 2 TB	12 TB
7	5 engines at 1 TB	2 engines at 2 TB	9 TB
7	3 engines at 1 TB	4 engines at 2 TB	11 TB
7	4 engines at 1 TB	3 engines at 2 TB	10 TB
8	2 engines at 1 TB	6 engines at 2 TB	14 TB
8	6 engines at 1 TB	2 engines at 2 TB	10 TB
8	4 engines at 1 TB	4 engines at 2 TB	12 TB

Note: Cache within an engine can be upgraded (capacity added), but cache cannot be downgraded (capacity removed).

PowerMaxOS

Each PowerMax engine comes with PowerMaxOS 5978 installed. PowerMaxOS is derived from the trusted and proven HYPERMAX OS used by the legacy VMAX3 and VMAX All Flash arrays; however, PowerMaxOS has been re-written to take advantage of NVMe architectures. PowerMaxOS continues to provide industry-leading high availability, I/O management, quality of service, data integrity validation, data movement, and data security within an open application platform. PowerMaxOS uses a real-time, non-disruptive storage hypervisor that manages and protects embedded services by extending high availability to services that traditionally would have run external to the array. The primary function of PowerMaxOS is to manage the core operations performed on the array, which include:

- Processing I/O from hosts
- Implementing RAID protection
- Optimizing performance by allowing direct access to hardware resources
- Managing and monitoring the system

Drive array enclosures

Each Brick comes with two 24-slot, dual-ported, 2.5" PCIe NVMe DAEs (DAE24). These DAEs use redundant, hot-swappable Link Control Cards (LCCs) which provide PCIe I/O connectivity to the NVMe flash drives. Aside from redundant LCCs, the DAE24 features redundant power supplies with separate power feeds, providing N+1 power and cooling, resulting in an energy-efficient consumption of up to 25 watts per drive slot. The DAE24 is 2U high and 19" deep.



Figure 3. Brick NVMe DAE24

The directors are connected to each DAE through a pair of redundant back-end I/O modules. The back-end I/O modules connect to the DAEs at redundant LCCs. Each connection between a back-end I/O module and an LCC uses an independent cable assembly. Within the DAE, each NVMe drive has two ports, each of which connects to one of the redundant LCCs.

The dual-initiator feature ensures continuous availability of data in the unlikely event of a drive management hardware failure. Both directors within an engine connect to the same drives using redundant paths. If the sophisticated fencing mechanisms of PowerMaxOS detect a failure of the back-end director, the system can process reads and writes to the drives from the other director within the engine without interruption.

Drive options and configurations

Both PowerMax 2000 and PowerMax 8000 support 1.92 TB, 3.84 TB, 7.68 TB, and 15.36 TB NVMe flash drive capacities as well as 750 GB and 1.5 TB SCM drives. All the drive sizes are 2.5" and feature a dual ported U.2 form factor PCIe interface. These drive capacities can be intermixed on the system.

Table 8. PowerMax family supported capacity specifications

	PowerMax 2000	PowerMax 8000
Capacity and drives		
Max Capacity per Array (Open) ¹	1.2 PBe	4.5 PBe
Base Capacity per Brick (Open)	13.2 TBu ³	54 TBu
Base Capacity per Brick (mainframe)	N/A	13.2 TBu
Incremental Flash Capacity Packs ⁵	13.2 TBu ³	13.2 TBu
Max Drives per Brick	44 Usable + spare(s)	32 Usable + spare(s)
Max Drives per Array	96 ²	288
NVMe drives		
NVMe Drives Supported (2.5")	1.92 TB, 3.84 TB, 7.68 TB, 15.36 TB	1.92 TB, 3.84 TB, 7.68 TB, 15.36 TB
SCM drives		
SCM Drives Supported (2.5")	750 GB, 1.5 TB	750 GB, 1.5 TB
BE Interface	NVMe over PCIe	NVMe over PCIe
RAID Options Supported	RAID 5 (7+1) (default) RAID 5 (3+1) RAID 6 (6+2) RAID 1 (Mirrored) ^{4,5}	RAID 5 (7 +1) (default) RAID 6 (6+2) RAID 1 (Mirrored) ^{4,5}

¹ Max capacity per array based on over provisioning ratio of 1.0.

² 192 drives can be supported in a single cabinet when two systems are packaged in the same rack.

³ 13.2 TBu Brick and Flash capacity pack usable capacities are based on RAID 5 (7+1). 11.3 TBu base capacity and Flash capacity pack increments possible with RAID 5 (3+1) on PowerMax 2000.

⁴ RAID 1 will be offered on both PowerMax 2000 and PowerMax 8000 systems post GA of the 3Q 2020 release. Contact your local Dell Technologies sales team for more details and availability.

⁵ Incremental capacity packs under RAID 1 are 15.0TBu minimum for both the PowerMax 2000 and PowerMax 8000.

Note: For detailed information about available PowerMax Brick drive configurations, see the [PowerMax Family Specification Sheet](#).

PowerMax Storage Resource Pools overview

In PowerMax, all physical storage capacity is combined into Storage Resource Pools (SRPs). At the lowest levels, SRPs consist of Disk Groups which contain of a collection of hard drives sharing the same technology and performance characteristics. The hard drives in each disk group are split into individual back-end data device segments called TDATs. The TDATs are placed into an associated Storage Tier.

An SRP is the collection of the total capacity of all its Storage Tiers – regardless of the underlying disk technology which the storage tiers are associated with. This physical capacity stored within an SRP is referred to its usable capacity (TBu). This usable capacity is accessed by hosts using thinly provisioned front-end storage devices called TDEVs. TDEVs are virtual representation of the SRP physical capacity which also considers overprovisioning and data reduction efficiencies. For example, an array with a single SRP which has 26 TBu, could be provisioned for 78 TB of host facing TDEV

capacity when a data reduction ratio of 3:1 is applied. This 78 TB of virtualized host facing TDEV capacity is referred to be the effective capacity (TBe) of the SRP. When a PowerMax is sized, both the usable capacity and effective capacity are considered. The total usable capacity (TBu) is the primary driver for sizing hard-drive-layout configurations. The effective capacity (TBe) is a primary driver when sizing PowerMax cache.

Host provisioned TDEVs to are placed into a storage group and assigned a Service Level. When a host writes application data to its provisioned TDEVs, this data is distributed across all the storage tiers within the SRP. Which storage tier the data is placed on within the SRP is governed by the Automated Data Placement (ADP) utility. ADP uses the PowerMax internal machine learning engine to employ predictive analytics and pattern recognition algorithms to place the data at the optimal physical location to ensure that the response time requirements for the assigned service level are met.

The following diagram illustrates the key components involved with a PowerMax SRP:

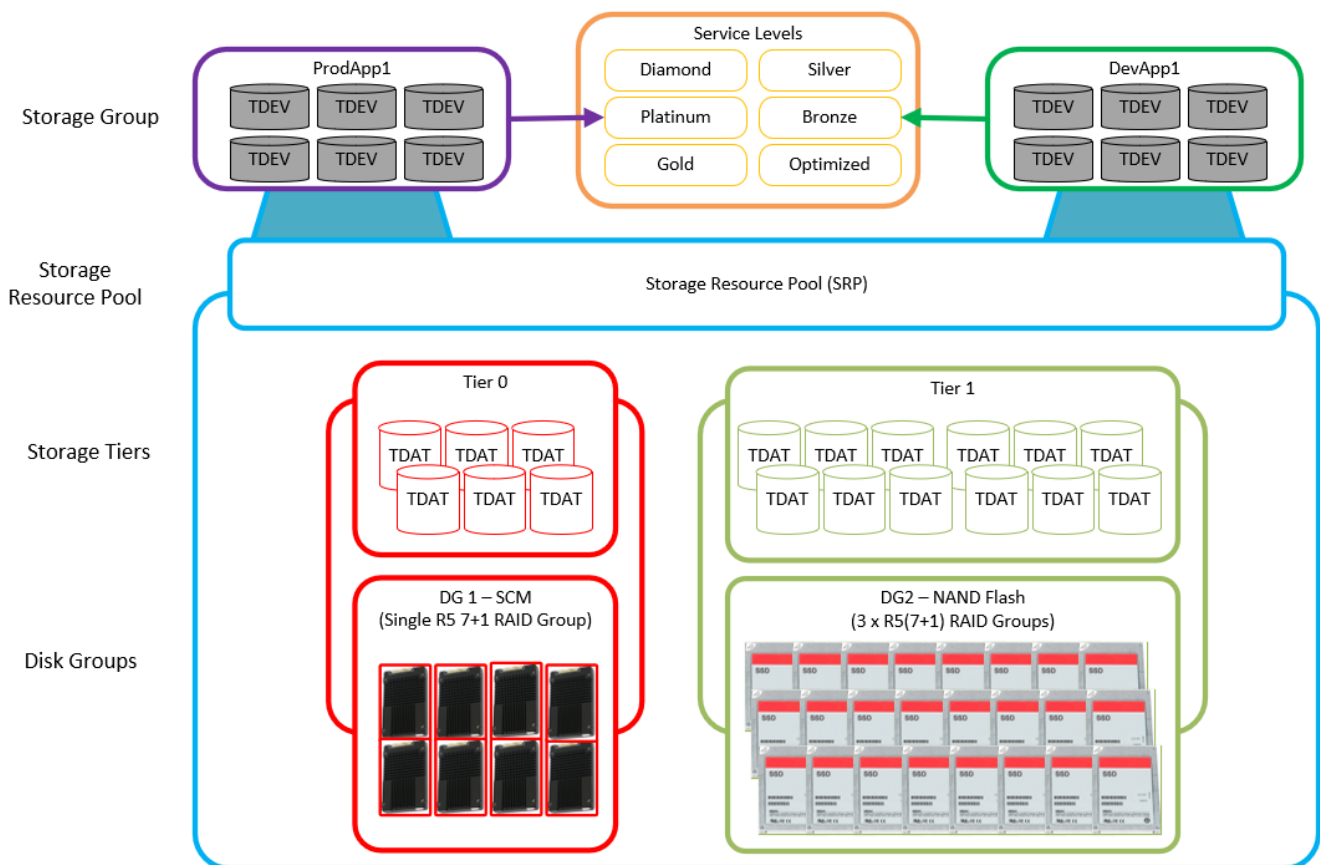


Figure 4. Typical components found with a PowerMax SRP with example of disk-group RAID-protection schemes

Note: The following points are specific notes regarding PowerMax SRPs.

- A PowerMax 8000 can now be configured so that both mainframe CKD and Open Systems FBA data can share a single SRP.

- PowerMax 8000 systems that will offer mixed FBA and CKD capacity must be born as a mixed system in the factory. CKD capacity cannot be added to an existing FBA system and conversely.
- Only a single RAID protection scheme can be used within the PowerMax SRP. The use of multiple RAID protection schemes is not supported within the PowerMax SRP.
- Dell Technologies recommends that all PowerMax systems be configured as a single SRP system so that customer data has access to the maximum amount of system resources as possible.
- While multiple PowerMax SRP systems are supported through RPQ, we do not recommend the use of multiple SRPs in a single PowerMax system for performance and manageability reasons.

Configuring SCM drives on PowerMax

SCM drives can be used on both the PowerMax 2000 and PowerMax 8000 systems. A way to think of NVMe connected SCM drives is as an extension of physical server memory where the application data stored on these SCM drives requires the highest levels of performance - typically reserved for server-based Non-Volatile DIMM (NVDIMM). While NVMe SCM does not provide the same levels of performance as NVDIMM, its economics from a \$/GB and \$/IOPS perspective make it an attractive lower cost alternative to server based NVDIMM for in memory applications such as SAP HANA.

PowerMax systems using SCM drives can be configured to have the SCM drives intermixed with traditional NAND flash drives in the DAEs. On these intermixed systems (known as “SCM as a Tier” systems – as shown in [Figure 4](#)), the devices carved from the SCM drives will be placed into “Tier 0” where the most active data on the system will reside.

Also, to ensure the highest levels of performance on the intermixed systems, the data on the SCM tier 0 is never compressed; however, it can be deduped. As said earlier, the system uses ADP’s predictive analytics and pattern recognition algorithms to ensure that the data is placed on and removed from Tier 0 in the most timely and efficient manner. Storage groups assigned the “Diamond” service level will be given priority for Tier 0 placement. Storage groups assigned as either “Silver” or “Bronze” are not eligible for Tier 0 placement and will always reside on NAND flash.


Note: The following are some other general configuration notes regarding SCM-as-a-tier PowerMax arrays.

- For optimum cost per performance, Dell Technologies recommends that the total usable capacity (TBu) of SCM Tier 0 be between 3% – 12% of the desired effective capacity (TBe) of the system.
- Up to three RAID groups of SCM (PowerMax 8000) and 4 RAID groups of SCM (PowerMax 2000) can be configured per engine as a tier 0.
- All engines must be configured identically with respect to SCM, for I/O balance (if an engine is configured with one R5 7+1 SCM RAID group, then all other engines in the system must be configured with one R5 7+1 SCM RAID group).

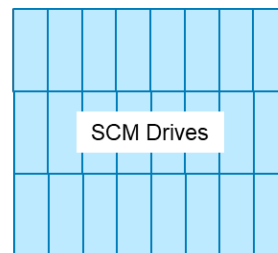
- While multiple SRPs are supported on PowerMax, only one SRP can contain SCM and this SRP must see the SCM storage as a tier (the SRP cannot be 100% SCM).
- Data is never compressed in SCM tier unless the system is comprised of 100% SCM drives.
- Data in SCM may be part of a dedupe set.
- Mixed SCM configurations using 750 GB and 1.5 TB SCM Drives are supported.
- SCM storage can use either RAID 1 (Mirrored), RAID 5 (3+1 and 7+1), RAID 6 (6+2) protection on PowerMax 2000.
- SCM storage can use either RAID 1 (Mirrored), RAID 5 (7+1) or RAID 6 (6+2) protection on PowerMax 8000.
- SCM storage must be of the same RAID type of the NAND flash in the system.
- Systems with SCM are configured with one SCM spare per engine. The SCM spare must match the largest capacity of SCM drive in the system.

PowerMax can also be configured as a 100% SCM system. In these systems (known as “SCM Bricks”), data can be both compressed and deduplicated. Activity-based compression rules apply where approximately 20% of the effective capacity of the SCM Brick is left uncompressed. The minimum capacity and incremental capacity configuration for an SCM Brick is 21 TBu consists of 17 (16 data + 1 spare) x 1.5 TB SCM drives configured into two RAID 5 (7+1) RAID groups. RAID 5 (7+1) protection using 1.5 TB drives is the only supported RAID configuration for SCM Bricks. SCM Bricks can have only a single SRP which consists of 100% SCM drives. NAND flash drives cannot be added to an SCM Brick.

The following figure shows the key differences between the two types of PowerMax SCM configurations:



SCM as a tier



SCM Brick

- Sizer will recommend 3-12% SCM by effective capacity (best for cost/performance)
- Up to 3 RAID groups of SCM (PowerMax 8000) 4 RAID groups of SCM (PowerMax 2000) can be configured per engine as a tier
- All engines must be configured identically with respect to SCM, for I/O balance
- Data is not compressed in SCM tier, no matter how large or small.
- Data in SCM may be part of a dedupe set
- PowerMax can be configured with 100% SCM
- In this case data can be compressed and deduplicated
- Activity Based Compression rules apply (~20% of total uncompressed)
- NAND drives cannot be added to 100% SCM systems
- Only one SRP supported

Figure 5. Supported PowerMax SCM configurations

PowerMax Universal Sparing

PowerMaxOS supports Universal Sparing to automatically protect a failing drive with a spare drive. Universal Sparing increases data availability of all volumes in use without loss of any data capacity, transparently to the host, and without user intervention.

When PowerMaxOS detects a drive is failing, the data on the faulty drive is copied directly to a spare drive attached to the same engine. If the faulty drive has failed, the data is rebuilt onto the spare drive through the remaining RAID members. When the faulty drive is replaced, data is copied from the spare to the new drive.

PowerMax systems have one spare drive for each drive type in each engine. The spare drives reside in dedicated DAE slots. If the system is a mixed NAND Flash and SCM system, then it will need one spare for the NAND Flash drives and one for the SCM drives. SCM Bricks only need one spare SCM drive. The spare drive type is the same as the highest capacity and performance class as the other drives in the engine.

For example, if a system uses both 3.84 TB and 7.68 TB NAND Flash drives in the configuration, only one 7.68 TB drive needs to be configured as a spare as it can replace either the 3.84 TB or 7.68 TB drives.




Figure 6. Universal Sparing example

PowerMax Smart RAID

PowerMax uses an active/active RAID group accessing scheme called Smart RAID. This allows RAID groups to be shared across directors, giving each director active access to all drives on the Brick or zBrick.

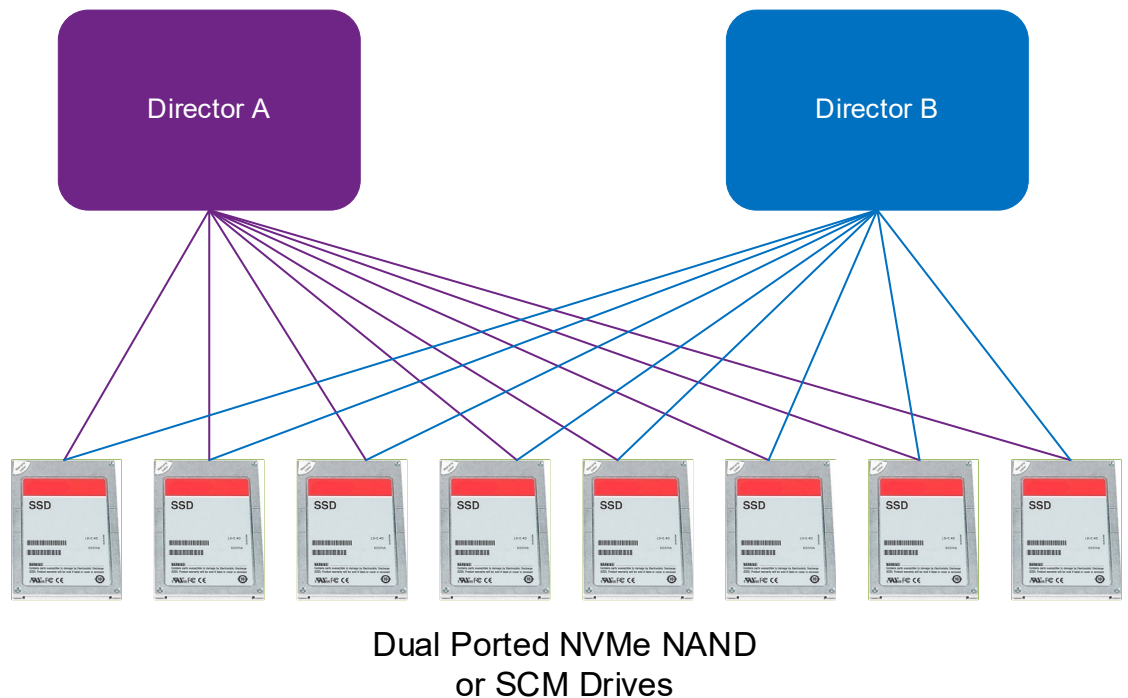


Figure 7. PowerMax Smart RAID

The use of Smart RAID on PowerMax provides customers with performance benefits as both directors on an engine can drive I/O to all the flash drives. This creates balanced configurations in the system regardless of the number of RAID groups. Smart RAID also allows for increased flexibility and efficiency as customers can order PowerMax 8000 systems with a single RAID group allowing for a minimum of 9 drives per engine with RAID 5 (7+1) and 1 spare or RAID 6 (6+2) and 1 spare; 2 drives and one spare with RAID 1 (Mirrored); and 5 drives per system for a PowerMax 2000 with RAID 5 (3+1) and 1 spare. This leaves 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.

PowerMax 2000 DAE connectivity and drive allocation schemes

Smart RAID and Universal Spare allow flexible connectivity and drive allocation schemes to occur with the PowerMax DAE. With the PowerMax 2000, each engine director has two NVMe I/O Modules. Each I/O module has two redundant paths. One path connects to either Link Control Card (LCC) A or LCC B in DAE 1 while the other path connects to either LCC A or LCC B in DAE 2. Each path from the NVMe I/O module to the LCC is a four lane PCIe Gen3 connection (4 GB/sec).

The following diagrams detail the DAE connectivity layout and drive allocation schemes for the PowerMax 2000.

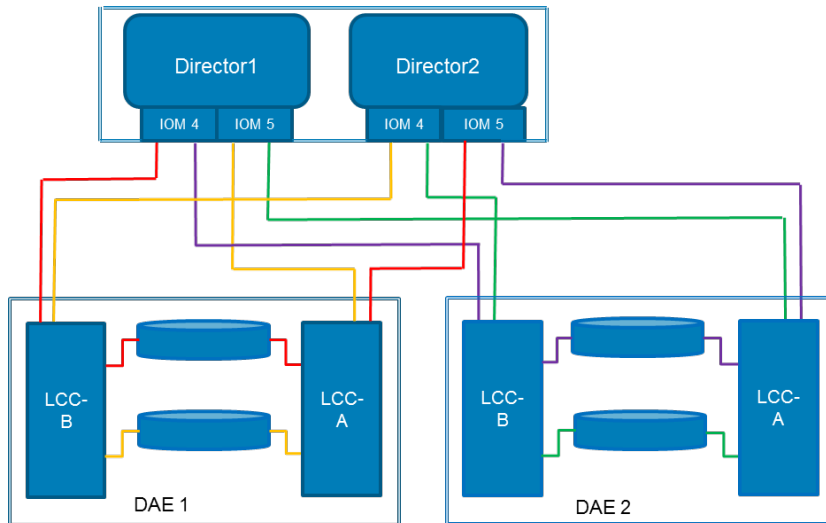


Figure 8. PowerMax 2000 single engine DAE connectivity

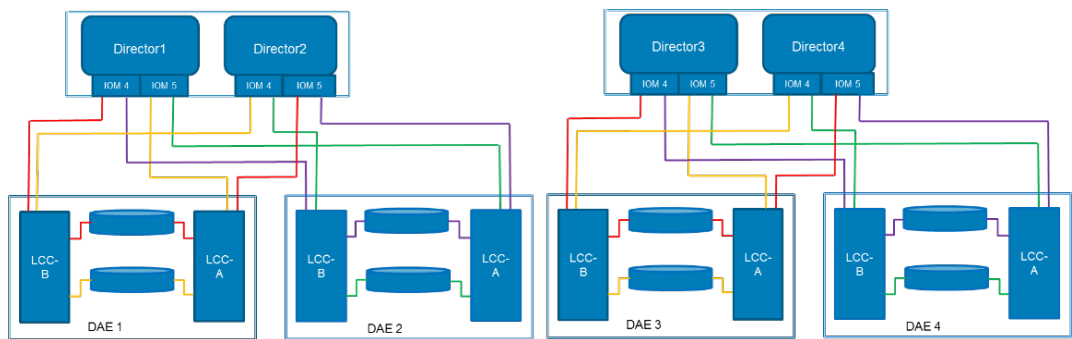


Figure 9. PowerMax 2000 dual engine DAE connectivity

The PowerMax 2000 can use the RAID 1 (Mirrored), RAID 5 (3+1), RAID 5 (7+1), or RAID 6 (6+2) protection schemes. Only one RAID protection scheme can be applied on the system. When populating the PowerMax 2000 DAEs, each engine requires a minimum of 1 RAID group including spare drives. There are two spare drive slots in a PowerMax 2000 system (slot 24 in each DAE); however, there can be only one spare drive for each Brick. When populating the drives into the system, the drives are alternately placed in DAE1 and DAE2.

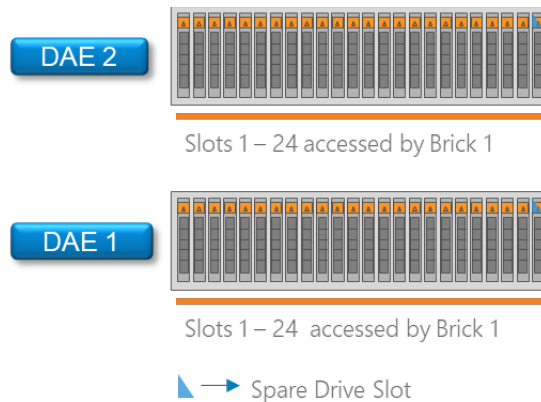


Figure 10. PowerMax 2000 DAE drive slot allocations for a single Brick

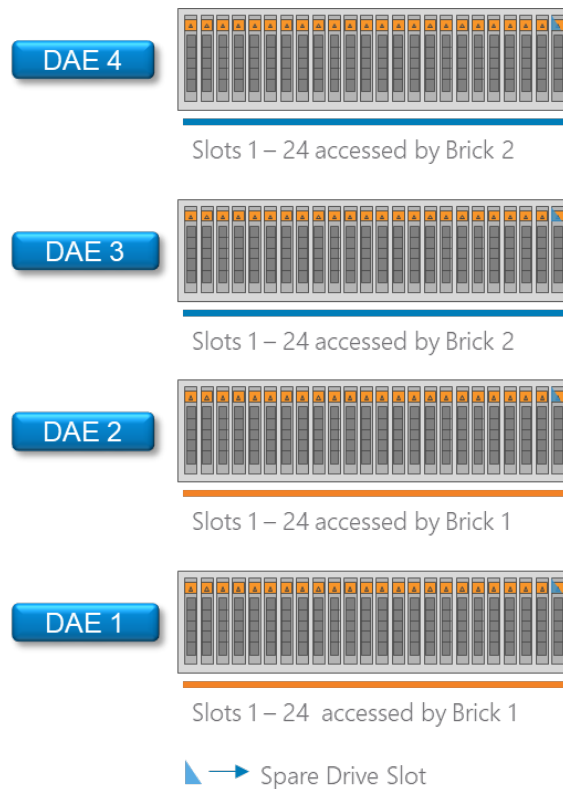


Figure 11. PowerMax 2000 DAE drive slot allocations for a dual Brick

The maximum number of usable drives which can be used with a single PowerMax 2000 Brick is 40 plus 1 spare drive for RAID 5 (7+1) or RAID 6 (6+2) configurations; and 44 usable drives plus 1 spare using a RAID 5 (3+1) configuration or RAID 1 (Mirrored).

Note: See the following list for details on PowerMax 2000 DAE and drive allocation.

- Mixed drive sizes can be used in the system for both NAND Flash and SCM. Drive sizes need to be one size increment apart (for example, 1.92 TB and 3.84 TB, or 3.84 TB and 7.68 TB).
- Only one spare drive per Brick is required. The spare needs to be the same size as the largest drive size used in the system.
- Every PowerMax 2000 system requires at least one RAID group.
- DAEs are not shared by the engines in a dual Brick PowerMax 2000 configuration.
- RAID groups are associated with a single Brick engine.
- Only one RAID protection scheme per PowerMax 2000 system is allowed.
- RAID 5 (3+1) requires a minimum of 4 drives plus 1 spare.
- RAID 5 (7+1) and RAID 6 (6+2) require a minimum of 8 drives plus 1 spare.
- RAID 1 (Mirrored) requires a minimum of 2 drives plus 1 spare.

PowerMax 8000 DAE connectivity and drive allocation schemes

The PowerMax 8000 uses Smart RAID and Universal Sparring to achieve the densest possible engine and flash drive capacity configurations in the industry. To achieve these

high densities, the PowerMax 8000 uses different DAE connectivity and drive allocation schemes from those used in the PowerMax 2000. In systems using a single Brick, the DAE connectivity is like the PowerMax 2000; however, drive slots 15 to 24 in the DAE 2 are reserved for future scale out of a second Brick.

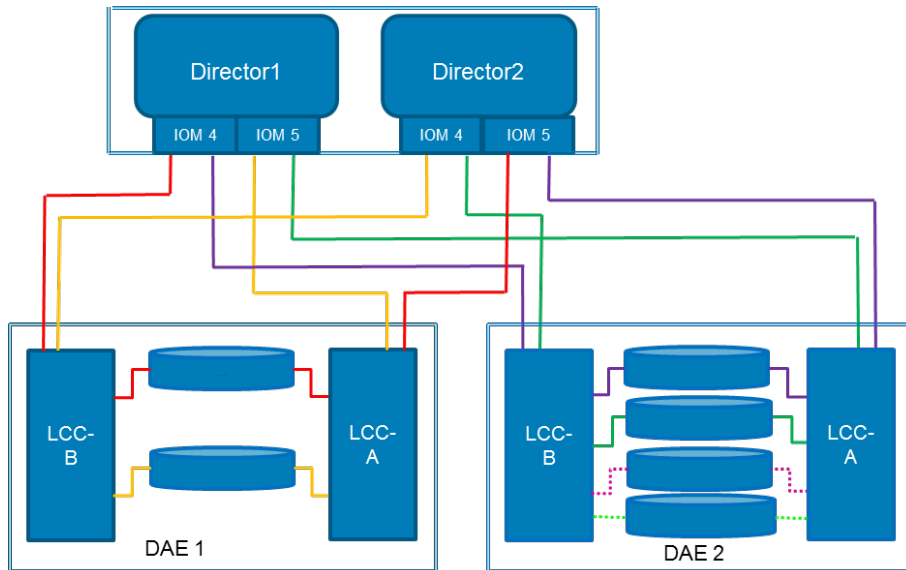


Figure 12. PowerMax 8000 single engine DAE connectivity

When a second Brick is added into the system, a third DAE is also added, and drive slots 15 to 24 of the DAE 2 on the first Brick can be populated and accessed by the second Brick. This is made possible as the 3rd and 4th Mini-SAS HD PCIe I/O ports on the LCCs in DAE 2 are used by the second Brick as shown in the following diagram:

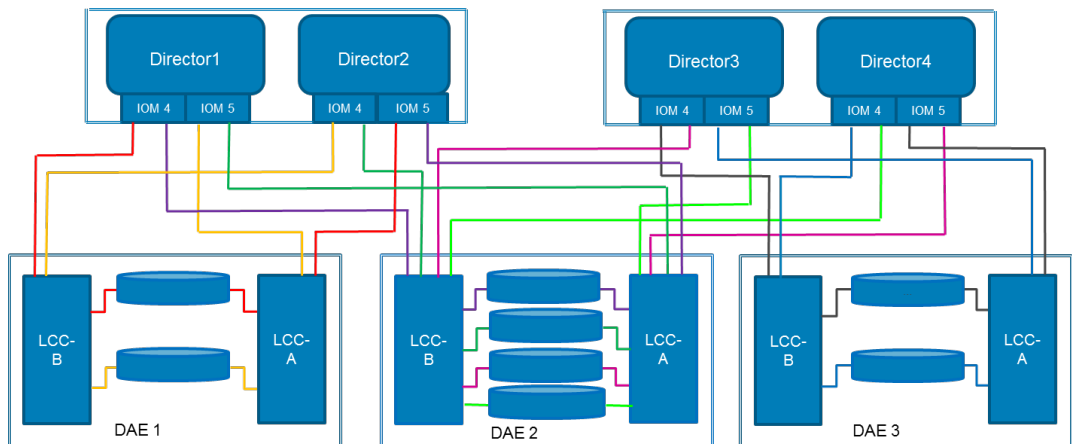


Figure 13. PowerMax 8000 dual engine DAE connectivity

The PowerMax 8000 can use the RAID 1 (Mirrored), RAID 5 (7+1), or RAID 6 (6+2) protection schemes. Like the PowerMax 2000, only one RAID protection scheme can be applied on the system, even on systems that have multiple SRPs. When populating the PowerMax 8000 DAEs, each Brick engine must have at least 1 RAID group including spare drives. For single Brick configurations, drives can be added in slots 1 to 24 of DAE 1, and in slots 1 -12 on DAE 2. Slots 13 and 14 in DAE 2 are reserved for spare drives.

This results in a maximum of 32 usable drive slots plus spares in a single Brick system. As with the PowerMax 2000, only one spare drive is required per Brick.

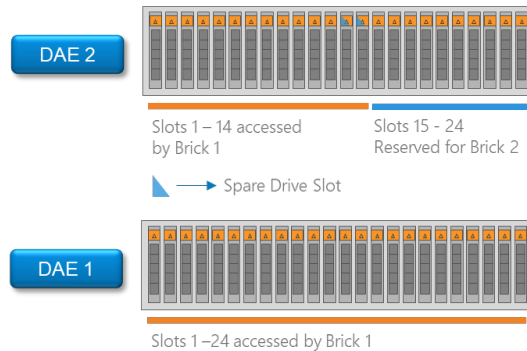


Figure 14. PowerMax 8000 drive slot allocations for a single Brick

A third DAE (DAE 3) is added to the system when adding a second Brick into the system. The second Brick uses slots 1 to 24 of DAE 3 and shares DAE 2 with the first Brick, using slots 17 to 24 in DAE 2. Slots 15 and 16 in DAE 2 are reserved for the second Brick spare drives. The following figure shows how drive slots are allocated in a dual Brick PowerMax 8000 system:

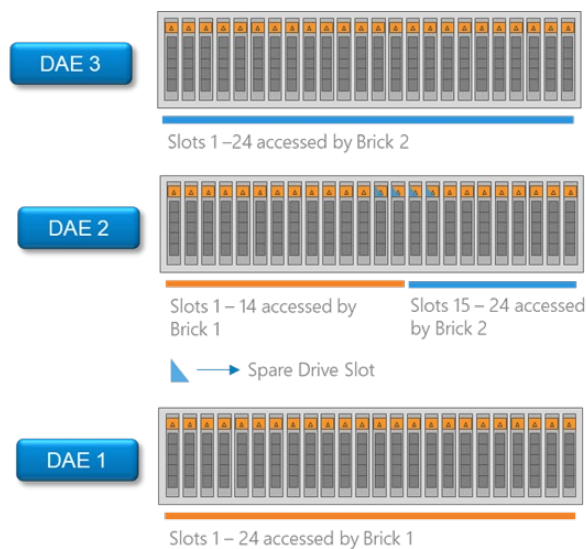


Figure 15. PowerMax 8000 drive slot allocations for dual Bricks

A PowerMax 8000 can be configured for open systems, mainframe, or mixed open systems and mainframe workloads.

Note: The following list includes PowerMax 8000 DAE and drive allocation notes.

- Every PowerMax 8000 system requires a minimum of one RAID group.
- Only one RAID protection scheme is allowed per PowerMax 8000 system
- Mixed drive sizes can be used in the system for both NAND Flash and SCM. Drive sizes need to be one size increment apart (for example, 1.92 TB and 3.84 TB, or 3.84 TB and 7.68 TB).

- Only one spare drive per Brick is required. The spare needs to be the same size as the largest drive size used in the system.
- RAID groups are associated to a single Brick engine.
- RAID 5 (7+1) and RAID 6 (6+2) protection schemes require a minimum of 8 drives plus 1 spare. RAID 1 (Mirrored) requires a minimum of 2 drives plus 1 spare
- Every even-numbered Brick will share a DAE with the previous odd-numbered Brick.
- Odd-numbered Bricks will have 24 plus 12 drives. Even-numbered Bricks will have 24 plus 10 drives.

Flash optimization

All flash-based storage systems demand the highest levels of performance and resilience from the enterprise data storage platforms that support them. The foundation of a true all flash array is an architecture that can fully leverage the aggregated performance of modern high-density flash drives while maximizing their useful life. Many features are built into the architecture of PowerMax to maximize flash drive performance and longevity. This section discusses these features in detail.

Cache architecture and caching algorithms

PowerMax is built upon a very large, high-speed DRAM cache-based architecture, driven by highly complex and optimized algorithms. These algorithms accelerate data access by avoiding physical access to the back end whenever possible. Dell Technologies has spent many years developing and optimizing caching algorithms. The algorithms used by PowerMax optimize reads and writes to maximize I/Os serviced from cache and minimize access to back-end flash drives. The system also monitors I/O patterns and proactively populates cache based on access to increase the chances of cache hits.

Some of the techniques used by the cache algorithms to minimize disk access are:

- 100% of host writes are cached.
- More than 50% of reads are cached.
- Recent data is held in cache for long periods, as that is the data most likely to be requested again.
- Intelligent algorithms destage in a sequential manner.

PowerMax write amplification reduction

Write amplification must be properly controlled to ensure the longevity of NAND flash and SCM storage devices. Controlling write amplification is one of the greatest strengths of PowerMax and what truly sets it apart from its competitors. Aside from intelligent caching algorithms, which keep data in cache as long as possible, the PowerMax uses additional methods to minimize the number of writes to flash. These methods are:

- **Write Folding** – Write Folding avoids unnecessary drive I/Os when hosts rewrite to an address range. This rewritten data is simply replaced in cache and not repeatedly written to the drive. Write folding can reduce writes to the NAND flash and SCM drives by up to 50%.
- **Write Coalescing** – Write Coalescing merges subsequent small random writes from different times into one large sequential write. These larger writes to the

storage drives align much better with the page sizes within the storage drive itself. Using write coalescing, PowerMax can take a highly random write host I/O workload and make it appear as a sequential write workload to the NAND flash and SCM drives.

- **Advanced Wear Analytics** – PowerMax also includes advanced drive wear analytics optimized for high capacity storage drives to make sure writes are distributed across the entire storage tier to balance the load and avoid excessive writes and wear to particular drives. Not only does this help manage the drives in the storage tier, but it also makes it easy to add and rebalance additional storage into the system.

All the write amplification reduction techniques used by PowerMax result in a significant reduction in writes to the back end, which in turn significantly increases the longevity of the NAND flash and SCM drives used in the array.

Boosting flash performance with PowerMaxOS FlashBoost

Dell Technologies is always striving to improve performance in its products. With every new hardware platform and release of software, the company makes strong efforts to remove potential bottlenecks which can impede performance in any way. One feature that Dell Technologies introduced and has made standard as a part of PowerMaxOS is FlashBoost.

FlashBoost maximizes PowerMaxOS efficiency by servicing read requests directly from the back-end flash drives. This approach eliminates steps required for processing I/O through global cache and reduces the latency for reads, particularly for flash drives. Customers with heavy read miss workloads residing on flash can see up to 100% greater IOPS performance. FlashBoost works with both NAND Flash and SCM storage.

Director slot layout and connectivity options

The Brick engine architecture uses a series of hot swappable modules that plug into slots in the engine directors. These modules include:

- Engine cooling fans and power supplies in slots accessible from the front of the engine director.
- I/O modules, management modules, and control stations in slots accessible from the rear of the engine director.

The following table describes the module components used in a Brick engine director:

Table 9. PowerMax engine director components

Director component	Quantity per director	Purpose
Power Supply	2	Provides redundant power to director.
Fan	5	Provides director cooling.
Management Module	1	Manage environmental functionality.
NVMe Flash I/O Module	Up to 4	The flash I/O modules use NVMe technology to safely store data in cache during the vaulting sequence (800 GB).

Director component	Quantity per director	Purpose
Front-end I/O Module	Up to 4	Provide front-end connectivity to the array. There are different types of front-end I/O modules that allow connectivity to various interfaces including Fibre Channel SCSI, Fibre Channel NVMe, iSCSI, FICON, SRDF, and embedded NAS (eNAS).
NVMe PCIe Back-end I/O Module	2	Dual-ported PCIe 4x Gen3 interface to the NVMe storage (8 GB/sec).
Data Reduction Module *	1	Performs inline data compression and deduplication, as well as SRDF compression.
Fabric I/O Module	1	Provides connectivity between directors. In multi-engine PowerMax 8000 systems, the fabric I/O modules are connected to an internal InfiniBand switch.

* An additional data reduction module is required for E2EE and will occupy a front-end I/O module slot.

The following diagram shows the director module layouts for the PowerMax 2000:




Figure 16. PowerMax 2000 director module layout by slot number

Both single-engine and multi-engine PowerMax 2000 systems use the same director module layout. Both configurations use two NVMe flash modules residing in slots 0 and 6 on each director. Slots 7 houses the data reduction module. Slots 2, 3, 8, and 9 are used for front-end connectivity modules. Slots 4 and 5 contain the NVMe PCIe back-end connectivity modules. Slot 10 houses the fabric modules. Slot 1 is reserved for future use.

The following diagrams detail the director module layouts for single-engine and multi-engine PowerMax 8000 systems:

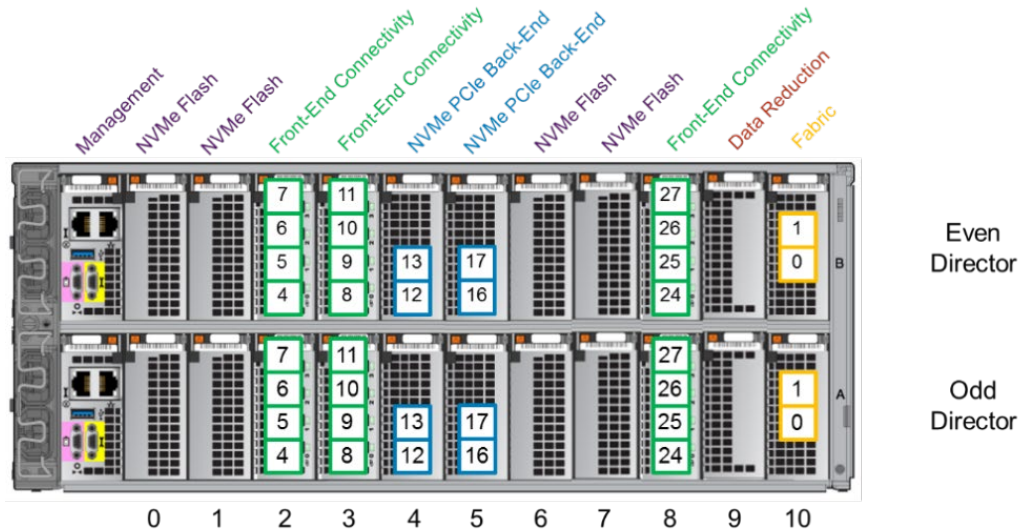


Figure 17. PowerMax 8000 director module layout by slot number: Single-engine system




Figure 18. PowerMax 8000 director module layout by slot number: Multiple-engine system

Unlike the PowerMax 2000, there are differences in the director module layouts between single-engine and multi-engine PowerMax 8000 systems. Single-engine PowerMax 8000 systems use four NVMe Flash modules. These modules occupy director slots 0, 1, 6, and 7. The data reduction module resides in slot 9. Slots 2, 3, and 8 are used for front-end connectivity modules.

Multi-engine PowerMax 8000 systems use three NVMe flash modules, occupying slots 0, 1, and 6. The data reduction module occupies slot 7. This leaves an additional slot for a front-end connectivity module allowing multi-engine PowerMax 8000 systems to have four front-end connectivity modules, occupying director slots 2, 3, 8, and 9.

Note: The following list includes director slot and connectivity notes.

- For PowerMax 8000 systems that only had a single engine originally, the single-engine configuration of three slots available for front-end modules is applied to each additional engine added to the system when the system is scaled out. When

additional engines are added to PowerMax 8000 systems that were originally multi-engine systems, these engines can have up to four slots available for front-end modules.

- On multi-engine systems, the compression module must use the same director slots on each engine.
- Data compression and deduplication are not available on the mainframe PowerMax 8000, but SRDF compression is available. On mainframe PowerMax 8000 systems (zBricks) which use SRDF compression only, place a compression module on the director with ports configured for SRDF. On single-engine configuration systems, place the SRDF compression module in slot 9; while on multi-engine configuration systems, place the SRDF compression module in slot 7.

Both the PowerMax 2000 and the PowerMax 8000 provide multiple front-end connections that implement several protocols and speeds. The following table highlights the various front-end connectivity modules available for a PowerMax system:

Table 10. Supported Brick front-end connectivity modules

Connectivity type	Module type	Number of ports	Mix with protocols	Supported speeds (Gbps)
Fibre Channel	32 Gbps FC	4	FC-NVMe / SRDF	8/16/32
Fibre Channel	16 Gbps FC	4	SRDF	4/8/16
SRDF	25 GbE	4	iSCSI	25
SRDF	10 GbE	4	iSCSI	10
iSCSI	25 GbE	4	SRDF	25
iSCSI	10 GbE	4	SRDF	10
FICON *	16 Gbps FICON	4	Single/Multi Mode	4/8/16
eNAS	10 GbE	2	None	10
eNAS	10 GbE (Copper)	2	None	10
eNAS Tape Backup	8 Gbps FC	4	None	2/4/8

* Supported on PowerMax 8000 only

Note: The following list includes other PowerMax connectivity notes.

- Each Brick engine has at least one front-end module pair (one front-end module per director).
- Since the number of front-end modules used in the Brick engine depends on the customer's requirements, some director slots may not be used.
- Front-end modules for Fibre Channel support multi-mode (MM). Front-end modules for FICON support both multi-mode (MM) and single-mode (SM). Front-end modules for 25 GbE/10 GbE support only MM optics.
- 25 GbE front-end modules will not auto-negotiate to 10 GbE.

Reliability, availability, and serviceability

PowerMax arrays are based on a revolutionary design and include key enhancements that improve the reliability, availability, and serviceability of the new systems – ideal choices for critical applications and 24x7 environments demanding uninterrupted access to information.

PowerMax systems use components that have a mean time between failure (MTBF) of several hundred thousand to millions of hours for a minimal component failure rate. A redundant design allows systems to remain online and operational during component repair. All critical components are fully redundant, including director boards, global memory, internal data paths, power supplies, battery backup, and all NVMe back-end components. Periodically, the system tests all components. PowerMaxOS reports errors and environmental conditions to the host system as well as to the Customer Support Center.

PowerMaxOS validates the integrity of data at every possible point during the lifetime of the data. From the point at which data enters an array, the data is continuously protected by error detection metadata. This protection metadata is checked by hardware and software mechanisms anytime data is moved within the subsystem, allowing the array to provide true end-to-end integrity checking and protection against hardware or software faults.

PowerMaxOS supports Industry standard T10 Data Integrity Field (DIF) block cyclic redundancy code (CRC) for track formats. For open systems, this enables host generated DIF CRCs to be stored with user data and used for end-to-end data integrity validation. Other protections include address/control fault modes for increased levels of protection against faults. These protections are defined in user definable blocks supported by the T10 standard and provide address and write status information in the extra bytes in the application tag and reference tag portion of the block CRC.

PowerMax reliability, availability, and serviceability (RAS) make it the ideal platform for environments requiring always-on availability. These arrays are designed to provide six-nines of availability in the most demanding, mission-critical environments. Some of the key PowerMax RAS features are as follows:

- No single points of failure — all components are fully redundant to withstand any component failure.
- Completely redundant and hot-pluggable field-replaceable units (FRUs) ensure repair without taking the system offline.
- Choice of RAID deployment options to provide the highest level of protection as desired.
- Mirrored cache, where the copies of cache entries are distributed to maximize availability.
- PowerMaxOS Flash Drive Endurance Monitoring – The nature of flash drives is that their NAND flash cells can be written to a finite number of times. This is referred to as flash drive endurance and is reported by drive firmware as a “percentage of life used”. PowerMaxOS periodically collects and monitors this information and uses it to trigger alerts back to Dell Support when a drive is nearing its end of useful life.

- Vault-to-flash with battery backup allows for cache-destage-to-flash and an orderly shutdown for data protection in the event of a power failure.
- Active/active remote replication using SRDF/Metro with read/write access to both Site A and Site B ensures instant data access during a site failure.
- Fully non-disruptive upgrades, including loading PowerMaxOS software from small updates to major releases.
- Continuous system monitoring, call-home notification, and advanced remote diagnostics.
- End to end efficient encryption and Data at Rest Encryption (D@RE) with integrated key manager, FIPS 140-2 validated to meet stringent regulatory requirements.
- T10 DIF data coding, with extensions for protection against lost writes.
- Detailed failure mode effects analysis (FMEA) during design of each component to ensure failure conditions can be handled gracefully.
- Extensive fault detection and isolation, allowing early wear-out detection and preventing the passing of bad data as good.
- Service defined and scripted to ensure success, including color-coded cabling, cable positioning, scripted steps, and checks of key parameters in those scripts.
- All flash cache data vault capable of surviving two Powerkey failures, ensuring that the system comes back even when something has failed before the vault and something else fails when returning from the power cycle.
- Support for thermal excursions with graceful shutdown if, for example, a data center loses air conditioning.
- Integrated data protection using Dell EMC PowerProtect Storage Direct backup and rapid restore, combining the gold standards in backup with SRDF replication technology.

Note: For more information about PowerMax RAS capabilities, see [Dell EMC PowerMax Reliability, Availability, and Serviceability](#).

PowerMax Data Services

Introduction

PowerMax Data Services help protect, manage, and move customer data on the array. These services run natively or embedded inside the PowerMax itself using the PowerMaxOS hypervisor to provide a resource abstraction layer. This allows the data services to share array resources — CPU cores, cache, and bandwidth. Doing this optimizes performance across the entire system and reduces complexity in the environment as resources do not need to be dedicated. Some of the most sought-after data services that are offered with the PowerMax product line are:

- Advanced data reduction using inline compression and inline deduplication
- Enterprise Grade Security
- Cloud Mobility

- Remote replication with SRDF
- High Availability with SRDF/Metro
- Local replication with TimeFinder SnapVX
- Embedded NAS (eNAS)
- Embedded Unisphere for PowerMax (eManagement)

Advanced data reduction

In PowerMax data storage systems, data reduction combines the proven Adaptive Compression Engine (ACE) and inline deduplication to provide a high-performing, space-efficient platform. Data reduction allows users to present more front-end effective capacity to lower back-end usable capacity. Compression and dedupe are two different functions that work together. Compression reduces the size of data sets and dedupe identifies identical data sets and stores a single instance. Performing both functions at the same time allows the system to be capacity efficient and deliver exceptional capacity savings.

Data reduction using the Adaptive Compression Engine

PowerMax employs inline hardware compression using the Adaptive Compression Engine (ACE). ACE provides PowerMax and VMAX All Flash customers with a data reduction method that provides negligible performance impact while delivering the highest space saving capability. The following design factors make the Dell EMC Adaptive Compression Engine unique:

- **Intelligent compression algorithms** – Intelligent compression algorithms determine the best compression ratios to be used and provide the ability to dynamically modify storage backend layout for the highest data compression efficiencies.
- **Inline hardware data compression** – Inline hardware data compression greatly inhibits the compression function from consuming critical PowerMax system core resources.
- **Activity Based Compression** – Activity Based Compression (ABC) focuses the compression function on the least busy data in the system, while allowing the most active data in the system to bypass the compression workflow. This ensures that all data in the system receives the appropriate compression focus while maintaining optimal response time.
- **Fine Grain Data Packing** – Fine Grain Data Packing which includes a zero reclaim function that prevents the allocation of buffers with all zeros or no actual data.
- **Enhanced Compression** – There is an additional compression algorithm found in PowerMaxOS called Enhanced Compression (EC). The EC algorithm scans already compressed data which the system determines to have been not accessed for a long period of time. The EC algorithm then tries to further reduce this data, to a larger compression ratio, to realize additional capacity savings.

The Adaptive Compression Engine is available to all open systems PowerMax and VMAX All Flash customers at no additional charge. ACE is not currently available for mainframe environments. Enhanced Compression is only supported on the PowerMax 2000 and 8000.

Data reduction using inline deduplication

PowerMax employs inline hardware deduplication to identify repeated data patterns on the array and store those repeated patterns using a single instance in the array's usable capacity. Depending on customer workloads, inline deduplication along with inline compression give PowerMax the ability to achieve an industry leading data reduction ratio with negligible performance impact. The following are the important design factors for deduplication on PowerMax:

- **Inline hardware data deduplication** – Inline hardware-based data deduplication prevents the consumption of critical PowerMax system core resources, limiting performance impact. The deduplication and compression functions are performed on the same hardware module in the PowerMax system.
- **Deduplication Algorithm** – PowerMax uses the SHA-2 algorithm to perform the deduplication function. The SHA-2 algorithm produces a unique data identifier for each item of data which is processed through the compression engine. These unique identifiers are stored in a Hash ID table on the PowerMax system.
- **Hash ID Table** – The Hash ID table stores all of the Hash IDs for the data processed through the compression engine. When a new write enters the compression engine, the Hash ID created for the write is compared to the Hash IDs already in the table. If the Hash ID is found to already exist in the Hash ID table, the write is not written to the storage.
- **Dedupe Management Object (DMO)** – The DMO is a 64-byte object which serves as the connection (pointer) between the devices and the single instance of actual data. DMOs are stored in the PowerMax global cache.

Inline deduplication is available to all open systems PowerMax customers at no additional charge. Deduplication is not currently supported for mainframe environments or on VMAX All Flash systems.

Enterprise-grade security

In modern data centers, data security is of paramount concern as it is estimated that the total cost of worldwide data breaches will exceed \$5 trillion by 2024, according to this [report](#) by Juniper Research. Across the industry, enterprises are looking at ways to secure their data by protecting it from various forms of breaches and cyberattacks that can steal data, make it inaccessible, alter data, and makes the data unreliable. For data storage, these measures are focused on securing the data path from host to the array, the storing of data internal the array, and implementing comprehensive user access controls that prevent unauthorized access to the array.

The PowerMax is known for providing the highest levels of data security for its customers as it provides hardened security measures for addressing the data path to the PowerMax, the data within the PowerMax system, and provides comprehensive user access controls that prevent unauthorized access to the PowerMax. The following bullets provide some detail on how this is done:

- Data path security
 - End to End Efficient Encryption (E2EEE) using powerful Thales Data Encryption. The Thales solution provides host to PowerMax data-in-flight encryption and once arriving at the array, the data is decrypted on internal

PowerMax components which can additionally provide up to a 5:1 data reduction ratio.

- Internal to the array
 - Data-at-Rest Encryption (D@RE): internal controller-based D@RE which is FIPS 140-2 certified.
 - PowerMax D@RE has zero performance impact and works with all data services
- User Controls
 - Secure user access with strict Role Based Authentication Controls (RBAC), plus the ability to implement traditional LDAP/AD controls
 - Tamper-proof audit logs which prevent the ability of an intruder to manipulate the historical records of what operations were performed on the system, at what time, and by whom.
 - Secure Snapshots that cannot be deleted by users or by other intrusive measures. Secure Snaps are proven to provide the ability to recover the user data environment from ransomware and other malicious attacks.

Dell Technologies continues to enhance PowerMax security-related testing and certification to include the latest Security Technical Implementation Guides (STIGs) required by the US Federal Government and Department of Defense.

- For more information about PowerMax Security see the *PowerMax Family Security Configuration Guide*.
- For more detailed information about E2EEE, see the document *Dell EMC PowerMax: End to End Efficient Encryption* on [Dell.com/StorageResources](https://www.dell.com/storage/resources).

Remote replication with SRDF

SRDF is one of the most popular data services in the enterprise data center because it is considered the gold standard for remote replication. Up to 70% of Fortune 500 companies use this tool 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.

PowerMax runs an enhanced version of SRDF specific for all flash use cases. This version uses multi-core, multi-threading techniques to boost performance; and powerful write folding algorithms to greatly reduce replication bandwidth requirements along with source and target array back-end writes to flash.

There are four types of SRDF:

- **SRDF Synchronous (SRDF/S)** – SRDF/S delivers zero data loss remote mirroring between data centers separated by up to 120 miles (200 km).
- **SRDF Asynchronous (SRDF/A)** – SRDF/A delivers asynchronous remote data replication between data centers up to 8000 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 world's 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). SRDF/Metro allows for storage array clustering, enabling even more resiliency, agility, and data mobility. SRDF/Metro allows hosts or host clusters access to LUNs replicated between two different sites. The hosts can see both views of the Metro Replicated LUN (R1 and R2), but it appears to the host operating system as if it were the same LUN. The host can then write to both the R1 and R2 devices simultaneously. This use case accounts for automated recovery and the seamless failover of applications thus avoiding recovery scenarios altogether. Other key features of SRDF Metro are:
 - It provides concurrent access of LUNS /storage groups for non-stop data access and higher availability across metro distances
 - It delivers simpler and seamless data mobility
 - It supports stretch clustering which is ideal for Microsoft and VMware environments
- **SRDF/Metro Smart DR** – SRDF/Metro can provide additional levels of data resiliency and efficiency by incorporating intelligent software (**Smart DR**) that allows both primary PowerMax sites (R1 and R2) to send a single copy of data to a shared DR site by asynchronous replication, resulting in the following benefits:
 - Increased resiliency: If one of the primary PowerMax arrays becomes unavailable, DR operations continue with the surviving primary PowerMax array
 - Increased efficiency: A single DR copy for both SRDF/Metro arrays and data that is only sent from one side to the DR array. This results in space and bandwidth efficiency.

SRDF software is included in the PowerMax Pro and zPro software packages, with no capacity-based licensing. It can be ordered as an addition to the Essentials and zEssentials software packages. Any hardware needed to support SRDF must be purchased separately.

Local replication with TimeFinder SnapVX

Every PowerMax array comes with the local replication data service TimeFinder SnapVX, which is included as part of the Essentials and zEssentials packages. SnapVX creates very low-impact snapshots. SnapVX supports up to 256 snapshots per source volume, up to 1024 snapshots per source using Snapshot Policies or zDP, and up to 65 million total snapshots per array. Users can assign names to identify their snapshots, and they can set automatic expiration dates on each snapshot.

SnapVX provides the ability to manage consistent point-in-time copies for storage groups with a single operation. Up to 1024 target volumes can be linked per source volume, providing read/write access as pointers or full-copy clones.

Local replication with SnapVX starts out as efficiently as possible by creating a snapshot: a pointer-based structure that preserves a point-in-time view of a source volume. Snapshots do not require target volumes.

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.

Each snapshot has a user-defined name and can optionally have an expiration date, both of which can be modified later. Management interfaces provide the user with the ability to take a snapshot of an entire storage group with a single command.

A point-in-time snapshot can be accessed from a host by linking it to a host-accessible volume referred to as a target. Target volumes are standard thin volumes. Up to 1024 target volumes can be linked to the snapshot or snapshots of a single source volume. This limit can be achieved either by linking all 1024 target volumes to the same snapshot from the source volume, or by linking multiple target volumes to multiple snapshots from the same source volume. However, a target volume may only be linked to a single snapshot at a time.

By default, targets are linked in a no-copy mode. This no-copy linked target functionality greatly reduces the number of writes to the back-end flash drives as it eliminates the requirement of performing a full volume copy of the source volume during the unlink operation in order to continue to use the target volume for host I/O. This saves the back-end flash devices from enduring a large amount of write activity during the unlink operation, further reducing potential write amplification on the PowerMax array.

Note: For more information about PowerMaxOS local replication options, see [Dell EMC PowerMax and VMAX All Flash: TimeFinder SnapVX Local Replication](#).

Cloud Mobility for Dell EMC PowerMax

PowerMax Cloud Mobility offers the seamless and transparent movement of data from on-premises to cloud, enabling PowerMax customers to leverage public cloud for agile and economical storage. Archiving and long-term retention are primary examples of how PowerMax customers can leverage cloud services such as Amazon Web Services (AWS), Microsoft Azure, and Dell EMC ECS for low-cost storage. PowerMax data can be recovered back to the source PowerMax if needed. Archiving to the cloud frees up capacity for on-premises PowerMax arrays to support higher priority applications on-premises—extending the useful life of PowerMax.

PowerMax data stored in the cloud can be made available to an AWS system for secondary processing. For example, a Linux image can run Oracle in AWS which in turn can mount a PowerMax database copy and perform reporting, analytics, or development/test on that database. When the secondary processing is complete, the data can be exported, and the infrastructure can be removed, allowing the customer to realize the inherent cost savings of a flexible IaaS public cloud consumption model.

Key capabilities of Cloud Mobility:

- Connect PowerMax to public cloud and hybrid cloud with ease by creating policy-based, automated snapshots in seconds
- Copy PowerMax data securely and efficiently to the cloud with in-flight encryption, in-flight compression, and flexible network bandwidth management (QoS)
- Leverage public cloud for cost-effective archive and long-term retention for up to 50% cost reduction (\$/GB)
- Run applications on AWS infrastructure with PowerMax data for reporting, analytics, test/dev

Note: For more information about PowerMax Cloud Mobility, see [Cloud Mobility for Dell EMC PowerMax](#).

PowerMaxOS Quality of Service features

In modern data-center environments, applications and workloads may require different performance envelopes which need to be delivered using a SAN environment which could be comprised of multiple generations of equipment (HBAs, switches, and storage arrays). These mixed environments produce challenges when trying to deliver the consistent performance levels enterprise applications require. In order to help deliver a consistent performance level for applications deployed in these diverse environments, PowerMaxOS employs powerful Quality of Service (QoS) features in the following ways:

- **Service levels** provide open systems customers with the ability to separate applications based on performance requirements and business importance. PowerMaxOS provides the ability to set specified service levels to ensure the highest priority application response times are not impacted by lower priority applications. Service levels address the requirements of customers to ensure that applications have a predictable, and consistent, level of performance while running on the array. The available service levels are defined in PowerMaxOS and can be applied to an application's storage group at any time. This allows for the Storage Administrator to initially set, as well as change, the performance level of an application as needed.
- **Host I/O limits** is a feature that can be used to limit the amount of front-end (FE) bandwidth and I/O operations per second (IOPS) that can be consumed by a set of devices over a set of director ports. Setting host I/O limits allows a user to define front-end port performance limits on a storage group. These front-end limits can be set by IOPS, host MB per host, or a combination of both. Host I/O limits can be set on a storage group that has a specified service level to throttle IOPS on applications that are exceeding the expected service-level desired performance.
- **Initiator Bandwidth Limits** is a feature which can be used to mitigate a well-known class of performance problem inherent to all lossless storage transport protocols called **slow drain**. Slow drains on a fabric can occur for a variety of reasons, but often they stem from a mismatch between the maximum link speeds supported by an initiator and target. This mismatch in link speeds is often seen in SAN fabrics which use multiple generations of equipment, such as when a new 32 Gb FC storage array is provisioned to legacy hosts which use slower speed 8 Gb FC HBAs. In this case, the data coming out of the new 32 Gb FC storage port could quickly overrun the processing ability of the 8 Gb HBA. Because of inherent Fibre Channel flow controls, the 32 Gb FC storage port would stop transferring frames until the 8 Gb HBA had cleared enough so it could start receiving frames again. Other hosts which are provisioned storage from the 32 Gb FC port would see a degradation in throughput and overall performance during this time period. PowerMaxOS Initiator Bandwidth Limits are designed to address this problem. Initiator Bandwidth Limits throttles the amount of throughput a PowerMax storage port can deliver to a host initiator so that the storage port will not over run the initiator's capabilities to process the incoming data. Initiator bandwidth limits are placed on an initiator group and only affect the initiators within the group. Other initiators which are using the storage port are unaffected and will still receive data at unthrottled speeds.

All PowerMax QoS features can be applied using traditional PowerMax management tools (Unisphere for PowerMax, REST API, and Solutions Enabler). PowerMaxOS QoS features are available at no additional cost for both PowerMax systems and VMAX All Flash systems which are running PowerMaxOS 5978.

Consolidation of block and file storage using eNAS

Embedded NAS (eNAS) data service extends the value of PowerMax to file storage by enabling customers to leverage vital enterprise features including flash level performance for both block and file storage, as well as simplify management, and reduce deployment costs. PowerMax, with the eNAS data service, becomes a unified block and file platform, using a multi-controller, transactional NAS solution. It is designed for customers requiring hyper consolidation for block storage combined with moderate capacity, high performance file storage in mission-critical environments. Common eNAS use cases include running Oracle on NFS, VMware on NFS, Microsoft SQL on SMB 3.0, home directories, and Windows server consolidation.

eNAS uses the hypervisor provided in PowerMaxOS to create and run a set of virtual machines within the PowerMax array. These virtual machines 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.

Aside from performance and consolidation, some of the benefits that PowerMax with eNAS can provide to a customer are:

- Scalability – easily serve over 6000 active SMB connections
- Metadata logging file system ideally suited for an all flash environment
- Integrated asynchronous file level remote replication with File Replicator
- Integration with SRDF/S
- Small attack surface – not vulnerable to viruses targeted at general purpose operating systems

The eNAS data service is included in the Pro software package. It can be ordered as an additional item with the Essentials software package. All hardware required to support eNAS on PowerMax must be purchased separately.

Non-Disruptive Migration

Data migrations have always been challenging in an enterprise environment. The complexity and size of very large data storage environments makes planning for, scheduling, and performing migrations extremely difficult. Migrations also often involve applications that cannot be taken offline, even briefly, for cutover to a new data storage array. Dell EMC Non-Disruptive Migration (NDM) allows customers to perform online data migrations that are simple and completely non-disruptive to the host and application.

NDM is designed to help automate the process of migrating hosts and applications to a new PowerMax array with no downtime. Non-Disruptive Migration leverages SRDF replication technologies to move the application data to the new array. It also uses auto-provisioning, in combination with PowerPath or a supported host multipathing solution, to manage host access to the data during the migration process.

NDM offers PowerMax customers the following benefits:

- Allows migration from VMAX or VMAX3 to PowerMax or VMAX All Flash arrays with hosts and applications completely online
- Comprehensive support for most FBA Host Stacks. Some older operating systems could require a host reboot as part of the migration process
- Designed for ease of use with control operations that automate the setup and configuration of the migration environment
- Managed by familiar and simple user interfaces using Solutions Enabler and Unisphere
- Migrations can be easily canceled and failed back to the source array for any reason before commit
- Completely built in and requires no additional software or licensing costs

Note: Migrations should take place during low I/O activity to minimize performance impact. NDM currently does not support mainframe CKD devices.

Embedded management using Unisphere for PowerMax

PowerMax customers can take advantage of simplified array management using embedded Dell EMC Unisphere for PowerMax. Unisphere for PowerMax is an HTML5 based management interface that allows IT managers to maximize productivity by dramatically reducing the time required to provision, manage, and monitor PowerMax data storage assets.

Embedded Unisphere enables customers to simplify management, reduce cost, and increase availability by running PowerMax management software directly on the array. Embedded management (eManagement) is configured in the factory to ensure minimal setup time on site. The feature runs as a container on a director, eliminating the need for a customer to allocate their own equipment to manage their arrays. Aside from Unisphere, other key elements of the eManagement data service include Solutions Enabler, Database Storage Analyzer, REST API and SMI-S management software.

Unisphere for PowerMax delivers the simplification, flexibility, and automation that are key requirements to accelerate the transformation to the all-flash data center. For customers who frequently build up and tear down storage configurations, Unisphere for PowerMax makes reconfiguring the array even easier by reducing the number of steps required to delete and repurpose volumes. With PowerMax, storage provisioning to a host or virtual machine is performed with a simple four-step process using the default Diamond class storage service level. This ensures all applications will receive sub-ms response times. Using Unisphere for PowerMax, a customer can set up a multi-site SRDF configurations in a matter of minutes. In addition, Unisphere for PowerMax provides a full REST API, enabling customers to fully automate the delivery, monitoring, and protection of storage services from their enterprise storage. REST API also enables organizations to integrate their PowerMax storage with their own DevOps environment or with third-party tools.

Embedded Unisphere for PowerMax is a great way to manage a single PowerMax array; however, for customers who need to view and manage their entire data center, Dell Technologies provides Unisphere 360. Unisphere 360 aggregates and monitors up to 200 PowerMax, VMAX All Flash, and legacy VMAX arrays across a single data center. This solution is a great option for customers running multiple PowerMax and VMAX All Flash

arrays with embedded management (eManagement) who are looking for ways to facilitate better insights across their entire data center. Unisphere 360 provides storage administrators the ability to view site-level health reports for every PowerMax and legacy system VMAX or coordinate compliance to code levels and other infrastructure maintenance requirements. Customers can leverage the simplification of PowerMax management at data center scale.

Embedded Unisphere and Database Storage Analyzer are available on every PowerMax array as they are included in the Essentials and zEssentials software packages. Unisphere 360 is included in the Pro and zPro software packages or can be ordered with the Essentials and zEssentials software packages. Unisphere 360 does not run in an embedded environment and requires additional customer-supplied server hardware.

Advanced data analytics with CloudIQ

CloudIQ is a cloud-based monitoring and storage analytics application that can be used to proactively monitor PowerMax arrays. The value of CloudIQ is centered on its ability to give users new and valuable insights into the health of the storage system. It proactively monitors and measures overall health using intelligent, comprehensive, and predictive analytics—and that makes it easier for IT to identify storage issues quickly and accurately. These analytics (which admins can access from anywhere through a web interface or mobile app) can drive business decisions that could lower the organization's total cost of ownership associated with the array. CloudIQ delivers several key values to customers:

- **Reduce Total Cost of Ownership:** CloudIQ provides an easy single pane of glass from which you can monitor your Unity and SC systems, all from the web so you can access anytime, anywhere.
- **Expedite Time to Value:** Because it is deployed from the EMC Cloud, customers can simply log in to their CloudIQ account and immediately access this valuable information. There is nothing to set up, no licenses, no burdens.
- **Drive Business Value:** The CloudIQ Proactive Health Score provides an easy way to identify and understand potential vulnerabilities in the storage environment. With these proactive and targeted guidelines, the result is a more robust and reliable storage environment, resulting in higher uptime and optimized performance and capacity.

CloudIQ is free and can be used with all PowerMax and VMAX All Flash arrays.

PowerMax storage integration with IT automation tools

To properly manage the modern data center, IT organizations need to focus on problem solving and not worry about routine and repeatable tasks that can be automated; furthermore, IT Operations Automation cannot be limited to simple scripting tasks just to save a few clicks. Automation needs to be well thought through and designed in a way that can scale across organizations, processes, and a hybrid cloud infrastructure. Dell Technologies offers a range of solutions to integrate with automation tools using the PowerMax REST API that are quickly becoming industry standards.

PowerMax plug-ins for VMware vRealize Orchestrator

VMware vRealize Orchestrator (vRO) is an IT process automation tool that allows automated management and operational tasks across both VMware and third-party applications. With vRO, IT admins can create automation routines for elaborate workflows using a simple drag-and-drop approach. The workflows can span across different pieces of infrastructure and scale to very large task volume. Dell Technologies offers vRO plug-ins for PowerMax that brings an extensive range of storage functionality in the form of programmable blocks that can be dropped into a workflow's process map. Functionality includes:

- Storage provisioning
- Scheduled and on-demand snapshots
- Remote replication and high availability
- Access to various storage objects specific to the array
- VMware integrated storage operations

The different tasks under the above categories form the basic building blocks that can then be used together with any of the compute and networking tasks to form an end-to-end operation that is easy to read, maintain, and build on.

VMware vRealize Automation

VMware vRealize Automation (vRA) can make PowerMax vRO workflow automation recipes into a self-serve catalog (anything-as-a-service) covering the entire IT Operations ecosystem including Service Delivery, Infrastructure Management across multiple cloud environments and accelerated DevOps processes. Workflows that are automated in vRO can be used as these self-serve catalog items seamlessly.

Ansible modules for PowerMax

Ansible is a very popular configuration management platform to automate IT Operations and modern DevOps workflows. Ansible Playbooks have the various tasks coded in an easy to read YAML language. A Playbook consists of task sets called Plays corresponding a given host or set of hosts that the Play is run on. A Play in turn can have multiple tasks. Each task makes calls to Ansible library functions called modules that are coded in Python programming language. With the growing popularity of Ansible, a large number of modules have been contributed that cover a wide variety of infrastructure management and application deployment. This makes Ansible a great choice to rapidly build PowerMax automation routines.

Dell Technologies has developed several storage modules to automate storage administration, provisioning, and consumption for the PowerMax storage platform. These modules enable the operations teams to rapidly provision storage infrastructure with accuracy to respond to the fast-paced needs of application developers.

Docker, Kubernetes, and the Container Storage Interface driver specification

Docker is the Container Engine that is used to extract container images of applications and run the container image. Kubernetes is one of the most popular container deployment and management platforms. Any application that needs to persist information about a database needs access to storage. The Container Storage Interface (CSI) driver is rapidly evolving as a standard to provision and manage storage from a storage array. The Dell

EMC PowerMax storage platform now supports CSI drivers to seamlessly run containerized workloads. The CSI driver is the interface between the logical volumes in the Kubernetes environment called Persistent volumes and the PowerMax storage volumes or LUNs. Storage Classes specify a set of parameters for the different characteristics unique to the underlying storage arrays.

Note: For more information about using Dell EMC storage functionality through third-party tools and REST APIs, go to Dell.com/StorageResources.

Dell Technologies Future-Proof Program

The Dell Technologies Future-Proof Program gives customers additional peace of mind with guaranteed satisfaction and investment protection for future technology changes. This program includes the entire Dell EMC Storage Portfolio including the flagship PowerMax, VMAX All Flash, XtremIO X2, SC Series, Dell EMC Unity, Data Domain, Integrated Data Protection Appliance (IDPA), Isilon, and Elastic Cloud Storage (ECS) appliance. This program provides customers with the following benefits:

- 3 Year Satisfaction Guarantee – Dell Technologies guarantees 3 years of storage and data protection appliance satisfaction
- Hardware Investment Protection – Trade in existing or competitive systems for credit towards next generation Dell EMC data storage systems, data protection appliances, or Hyper Converged Infrastructure product offerings
- Predictable Support Pricing – consistent and predictable maintenance pricing and services for your storage appliances
- Storage Efficiency Guarantee – PowerMax introduces even greater efficiency with inline deduplication and enhanced compression and comes with a **3.5:1 data reduction guarantee with the future-proof program**.
- Never-Worry Data Migrations – Use integrated data-migration tools with seamless upgrades to move to next generation data storage systems

Note: For more information about the Dell Technologies Future Proof Program, contact Dell Technologies sales.

Streamlined software packaging

Software for the PowerMax is available in different packages, each with additional, optional features. There are two packages for open systems named the Essentials package and the Pro package. These are available on all PowerMax arrays. There are two additional packages available for PowerMax 8000: the zEssentials and zPro packages. These packages are available for mainframe environments only.

The open system packages are shown in the following table:

Table 11. PowerMax open systems software packaging options

Feature	Essentials package included	Essentials package options	Pro package included	Pro package options	Notes
PowerMaxOS	✓		✓		Includes Migration Tools, vVols, QoS ³
Embedded Management	✓		✓		Includes Unisphere for PowerMax, Database Storage Analyzer, Solutions Enabler, REST APIs, SMI-S
Advanced Data Reduction	✓		✓		Includes inline compression and deduplication
Local Replication	✓		✓		Includes TimeFinder SnapVX and Snapshot Polices
iCDM Basic (AppSync)	✓		✓		
Remote Replication Suite ¹		✓	✓		Includes SRDF/S/A/STAR
SRDF/Metro ¹		✓	✓		Includes SRDF/Metro SmartDR
Unisphere 360		✓	✓		
D@RE ²		✓	✓		Includes End to End Encryption
End to End Efficient Encryption ⁴		✓	✓		Included with D@RE license
eNAS ^{1,2}		✓	✓		
SRM		✓	✓		
PowerPath		✓	✓		Includes 75 host licenses
iCDM Advanced (AppSync)		✓	✓		
PowerProtect Storage Direct (formerly ProtectPoint)		✓		✓	
RecoverPoint		✓		✓	
Dell EMC Storage Analytics	✓		✓		
CloudIQ		✓		✓	Free for all VMAX and PowerMax customers by download

¹ Software packages include Dell EMC software licensing. Required hardware ordered separately.

² Factory configured. Must be enabled during the ordering process.

³ Includes service levels, host I/O limits, and initiator bandwidth limits.

⁴ This feature is orderable through RPQ only. The required Thales software licenses and hardware are ordered separately. Data at Rest Encryption (D@RE) must be enabled to use this feature.

The mainframe software packages and options are shown in the following table:

Table 12. PowerMax mainframe software packaging options (PowerMax 8000 only)

Feature	zEssentials package included	zEssentials package options	zPro package included	zPro package options	Notes
PowerMaxOS	✓		✓		
Embedded Management	✓		✓		Includes Unisphere for PowerMax REST APIs, SMI-S
Local Replication	✓		✓		Includes TimeFinder SnapVX, Compatible Flash (FlashCopy support)
Mainframe Essentials	✓		✓		Includes Compatible High Performance FICON (zHPF) and Compatible PAV (Dynamic, Hyper, and SuperPAV) support
Remote Replication Suite ^{1,3}		✓	✓		Includes SRDF/S/A/STAR, Mirror Optimizer, Compatible Peer (PPRC)
Unisphere 360		✓	✓		
AutoSwap		✓	✓		
D@RE ²		✓	✓		
zDP		✓	✓		
Mainframe Essentials Plus		✓		✓	zBoost PAV Optimizer
GDDR ³		✓		✓	

¹ Software packages include software licensing. Order any additional required hardware separately.

² Factory configured. Must be enabled during the ordering process.

³ Use of SRDF/STAR for mainframe requires GDDR.

Note: For the up-to-date PowerMax software packaging information, see the [PowerMax Product Guide](#).

PowerMax system deployments

Introduction

The Dell EMC PowerMax family offers customers an all-NVMe storage platform designed to provide industry-leading IOPS density per system in a single- and dual-floor-tile footprint. This section describes the deployable system layouts for the PowerMax 2000 and PowerMax 8000 systems. For information about available drive configurations and system usable capacities, see [Expandable modular architecture: PowerMax Brick](#).

PowerMax 2000 system configurations

The PowerMax 2000 brings unmatched efficiency and flexibility to the data center, providing customers with over 2.7 million IOPS (8 K RRH) and up to 1 PB of effective capacity in just 20U of total space.

The PowerMax 2000 can be configured using either one or two Bricks in a single standard Dell EMC Titan rack. Each Brick consumes 10U of rack space (20U max for dual-Brick

PowerMax 2000 systems). The initial Brick occupies the bottom 10U of the rack when shipped from Dell Technologies manufacturing. The second Brick occupies the 10U directly above the initial Brick. This is applicable for systems ordered as dual Bricks or scale-out systems. An additional PowerMax 2000 system can be added into the remaining 20U in the rack.




Figure 19. Single and dual Brick PowerMax 2000 configurations

The PowerMax 2000 does not feature a system tray, KVM, or internal Ethernet or InfiniBand switches. It uses direct InfiniBand connections between engines on dual Brick systems.

Note: The PowerMax 2000 can be installed in third-party racking. For more information about PowerMax 2000 third-party racking options, see the Dell EMC PowerMax Family Site Planning Guide.

PowerMax 8000 system configurations

PowerMax 8000 is the flagship of the PowerMax family and provides customers with unmatched scalability, performance, and IOPS density. It can consolidate disparate workloads on a mass scale as eight Bricks, can support over 15 million IOPS (8 K RRH), and can provide up to 4 PB of effective capacity in just two floor tiles of space.

The PowerMax 8000 is a highly configurable data storage array that can support configurations from one to eight Bricks within two standard Dell EMC Titan racks. Each rack can support up to four Bricks. Bricks 1 to 4 always occupy a single rack. PowerMax 8000 only requires a second rack when the Brick count is greater than four.

Single-rack PowerMax 8000 configurations

The following diagram shows PowerMax 8000 single Brick and dual Brick configurations:

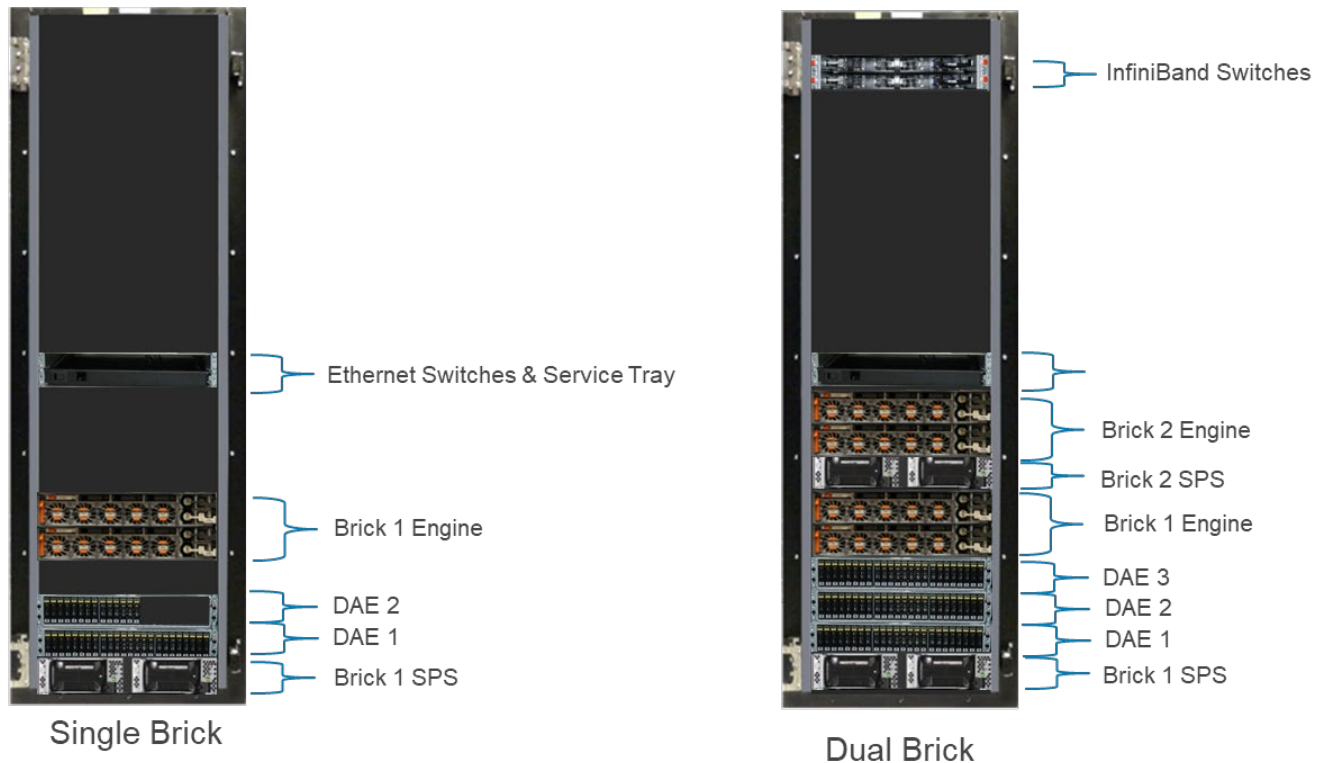


Figure 20. PowerMax 8000 single Brick and dual Brick configurations

The PowerMax 8000 uses redundant 16-port Dell EMC Networking X1018 Ethernet switches for the internal management network. This network connects to every engine and to the two internal InfiniBand fabric switches. The InfiniBand switches are required when two or more Bricks are configured in the system. The redundant 18-port InfiniBand fabric switches connect to every director in the system.

DAE 3 is added with the second Brick. As mentioned previously, DAE 2 is shared by Brick 1 and Brick 2. In DAE 2, drive slots 1 to 14 are used by Brick 1 while slots 15 to 24 are used by Brick 2. A PowerMax 8000 configuration best practice is that every even-numbered Brick shares a DAE with the previous odd-numbered Brick.

Note: For more information about PowerMax 8000 third-party racking options, see the Dell EMC PowerMax Family Site Planning Guide.

The following diagram shows three Brick and four Brick configurations for the PowerMax 8000:

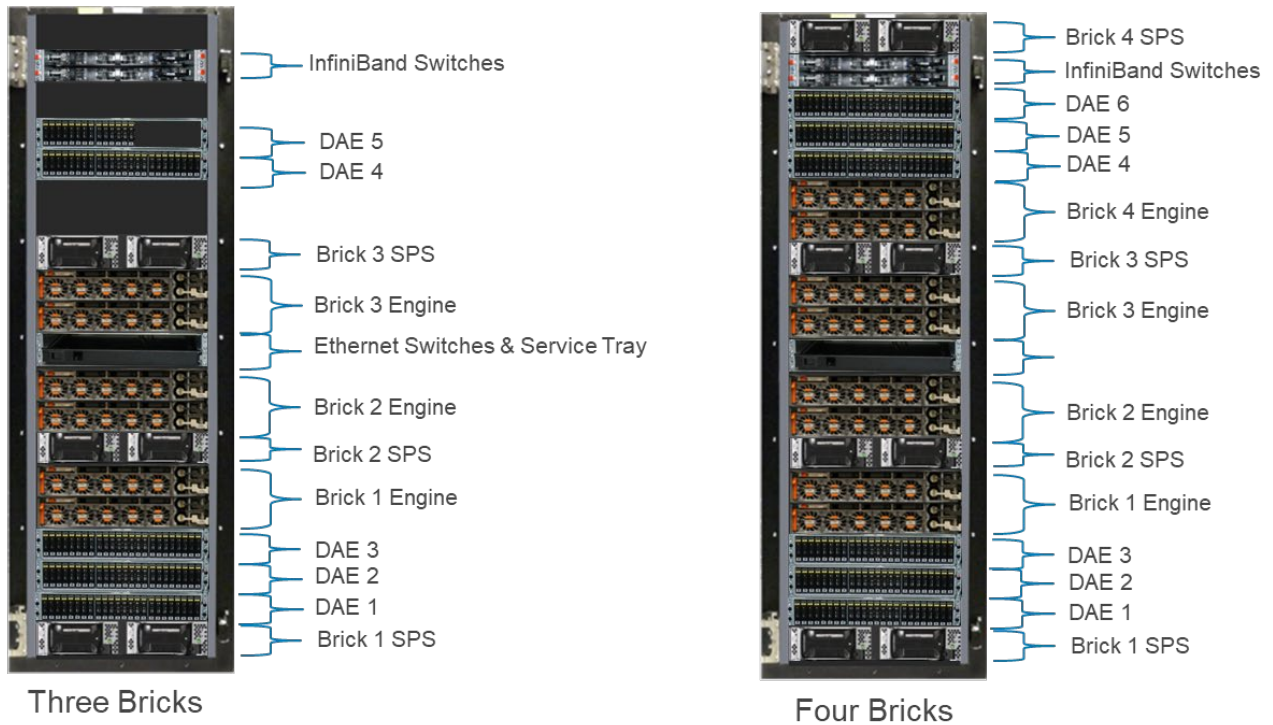


Figure 21. PowerMax 8000 three Brick and four Brick configurations

Dual-rack PowerMax 8000 configurations

PowerMax 8000 systems that have more than four engines require a second rack (system bay). Bricks are added to the second rack in the same manner and order as Bricks 1 to 4 are added in the first rack. The engines for the Bricks in the second rack are connected to the Dell X1018 Ethernet management switches. The Brick engine directors in the second rack are cabled to the InfiniBand switches in the first rack also. No additional InfiniBand or Dell X1018 switches are required for the second rack.

The PowerMax 8000 supports rack 2 dispersion configurations of up to 82 feet (25 meters) from the fabric switches in rack 1. Dispersed configurations require optical connections between the InfiniBand switches in rack 1 and the Brick engine directors in rack 2. Adjacent rack configurations can use Copper for connections between the Brick engines in rack 2 and the InfiniBand switches in rack 1.

The following diagrams show the various PowerMax 8000 two-rack configurations:

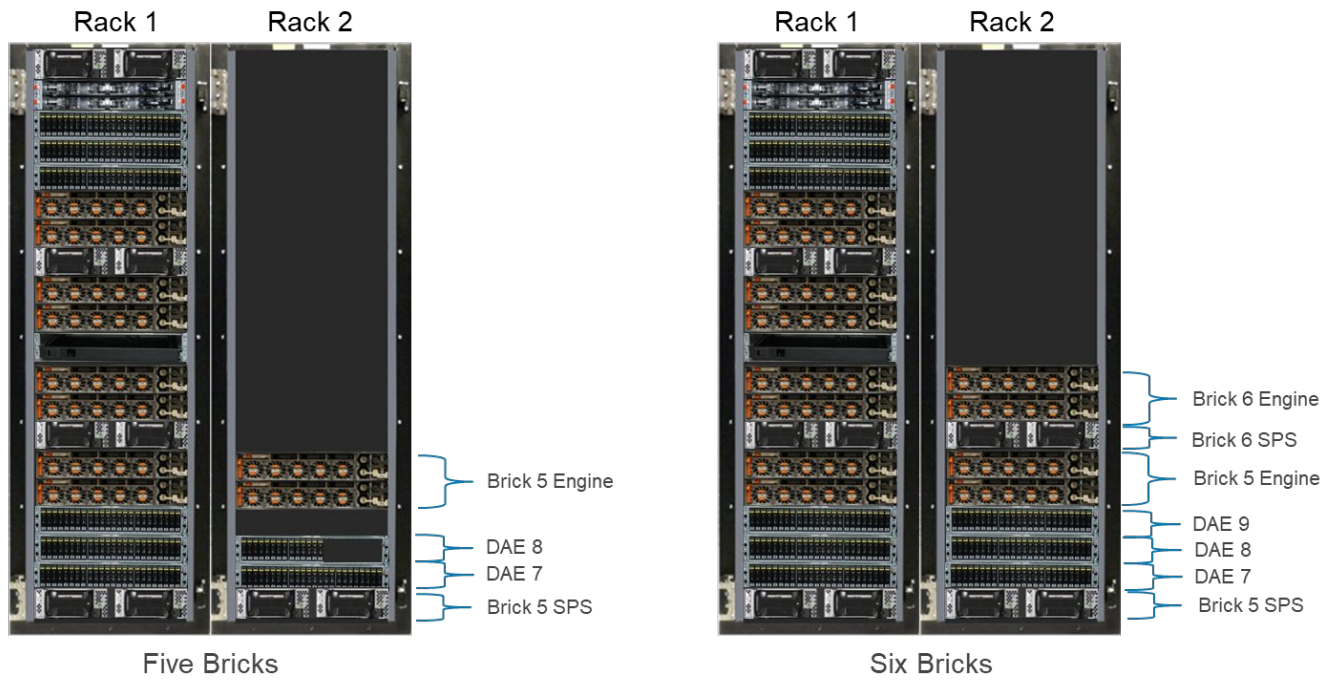


Figure 22. PowerMax 8000 five Brick and six Brick configurations

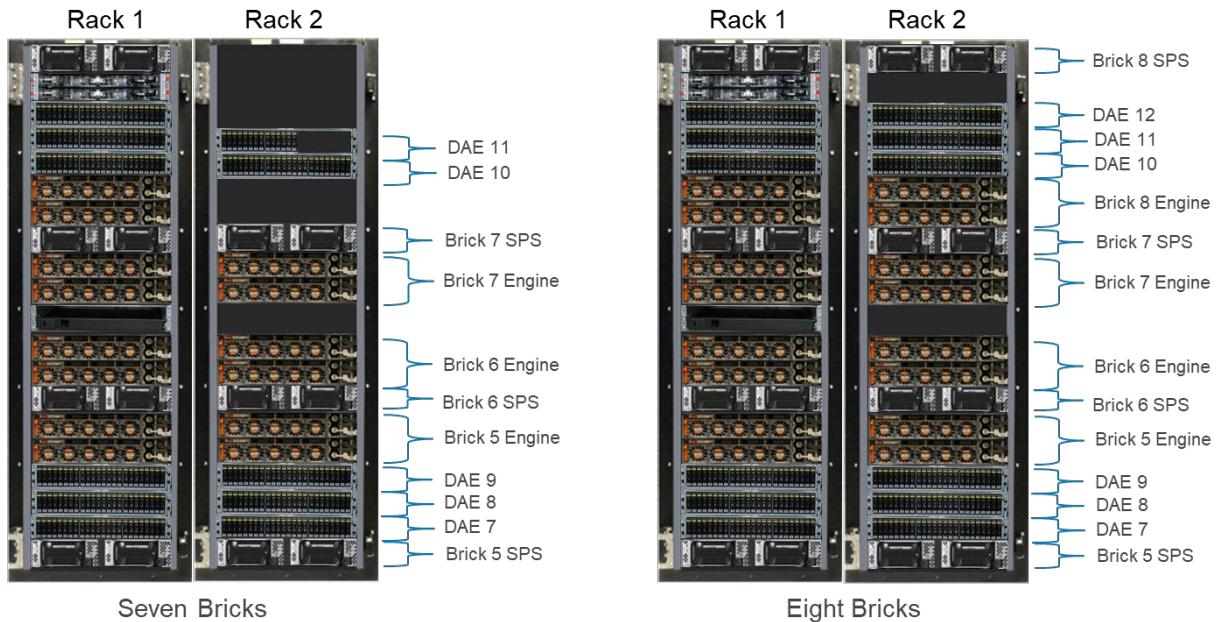


Figure 23. PowerMax 8000 seven Brick and eight Brick configurations

**Dell EMC
PowerOne with
PowerMax**

The setup, deployment, configuration, and maintenance of data center components (compute, storage, virtualization, and network) often create bottlenecks for organizations that wish to increase their business agility. It has been found that organizations need a solution that will enable simplified on-demand provisioning of these resources for their VMware virtualized data center infrastructure.

Dell EMC PowerOne is a Dell Technologies Cloud platform that brings together compute, storage, networking, virtualization, and data protection from across the Dell EMC Power portfolio in a single fully engineered end to end system. The Dell EMC PowerOne solution uses the Dell MX system for compute and Dell Networking for the system fabric. When configured with PowerMax for storage, the PowerOne solution offers its customers all the performance, reliability, data services, and security elements that enterprise-level mission-critical applications require.

Dell EMC PowerOne encompasses a system life-cycle management solution that automates and orchestrates the deployment, provisioning, and administration of a VMware virtualization platform running on Dell EMC components.




Figure 24. Dell EMC PowerOne solution featuring the PowerMax 2000

Dell EMC PowerOne, with its automation engine, removes the bottlenecks by automating and orchestrating most of the manual and repetitive tasks associated with provisioning and configuring compute, storage, virtualization, and networking resources for a given workload. It has been found that PowerOne can significantly reduce the amount of manual intervention typically required when installing and configuring a system of compute and storage resources to support an organization's workloads. The reduction in manual intervention can subsequently simplify processes related to installation, configuration, decommissioning, and maintaining these resources, thus reducing overall operational expenses.

Note: For more information about the PowerOne solution, go to <https://www.delltechnologies.com/en-us/converged-infrastructure/powerone.htm>.

Summary

The PowerMax family is the first Dell EMC data storage system to fully use NVMe technology for customer application data. Innovative PowerMax storage is built using a 100% NVMe end-to-end storage architecture, allowing it to reach unprecedented IOPS densities by eliminating the flash media choke points found using traditional SAS and SATA interfaces.

There are two PowerMax models:

- **PowerMax 2000** is designed to bring unmatched efficiency and flexibility to the data center, providing our customers with 2.7 million IOPS (8 K RRH) and up to 1.2 PB of effective capacity in just 20U total space.
- **PowerMax 8000** is designed to provide our customers unmatched scalability, performance, and IOPS density. It can consolidate disparate workloads on a mass scale as 8 Bricks can support 15 million IOPS (8 K RRH) and provide up to 4.5 PB of effective capacity in just two floor tiles of space.

The design of PowerMax is also built to provide not only the highest levels of performance, but also to deliver the highest levels of enterprise class reliability, availability, and serviceability. PowerMax customers can consolidate their workloads as block, file, and mainframe can all run on the platform. PowerMax comes complete with optimized data services which include advanced data reduction using inline deduplication and compression, the industry standard remote and local replication technologies of SRDF and TimeFinder SnapVX, and embedded array management with Unisphere for PowerMax. Cloud Mobility extends PowerMax storage to public or private clouds for long-term-retention use cases.

Technical support and resources

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage and data protection technical white papers and videos](#) provide expertise that helps to ensure customer success with Dell EMC storage and data protection products.

The following table provides a list of related resources:

Table 13. Related resources

Title	Collateral type	Part number
Dell EMC PowerMax Family Product Guide	Product Guide	
Dell EMC PowerMax Family Security Configuration Guide	Configuration Guide	
PowerMax Data Sheet	Data Sheet	H16891
Dell EMC PowerMax Software	Data Sheet	H16748
PowerMax 2000 / 8000 Specification Sheet	Specification Sheet	H16739
Dell EMC PowerMax Reliability, Availability, and Serviceability Technical White Paper	White Paper	H17064

Title	Collateral type	Part number
Dell EMC PowerMax and VMAX All Flash: GDPS and Advanced Copy Services Compatibility	White Paper	H16124
Dell EMC PowerMax Family Site Planning Guide	Technical Guide	
Dell-EMC-PowerMax and VMAX All Flash Snapshot-Policies Best Practices	White Paper	H18432
Dell EMC TimeFinder SnapVX Local Replication Technical Note	Technical Guide	H13697
Dell EMC SRDF/Metro Overview and Best Practices	Technical Guide	H14556
Dell EMC PowerMax End to End Efficient Encryption	White Paper	H18483
Cloud Mobility for Dell EMC PowerMax	White Paper	H18510
Transform SAP Landscape Challenges with Dell EMC PowerMax	Solution Overview	H17093
Consolidate Microsoft SQL Server with Dell EMC PowerMax	Solution Overview	H17092
Accelerate and Simplify Oracle Databases with Dell EMC PowerMax	Solution Overview	H16732
Top Ten Reasons Why Customers Deploy PowerMax	Handout/FAQ	H16738
Top Ten Reasons Why Customers Deploy Dell EMC PowerMax for Microsoft SQL Server	Top Reasons Handout	H17091
Top Ten Reasons Why Customers Deploy Dell EMC PowerMax for SAP Landscapes	Top Reasons Handout	H17090
Top Ten Reasons Why Customers Deploy Dell EMC PowerMax for VMware	Top Reasons Handout	H17074
Top 10 Reasons Why Dell EMC PowerMax for Oracle	Top Reasons Handout	H16725