Dell EMC VMAX All Flash: Family Overview

Abstract

This document provides an in-depth overview of the Dell EMC™ VMAX™ All Flash family. It includes details about the theory of operation, packaging, and the unique features which make it the premier all-flash storage product for the modern data center.

September 2020

Revisions

Date	Description
December 2017	Initial release
March 2018	Content update
April 2018	Update for PowerMaxOS
September 2019	Updates for PowerMaxOS Q3 2019 SR
September 2020	Updates for PowerMaxOS Q3 2020 release

Acknowledgments

Author: James Salvadore

Table of contents

D∆LL Technologies

D&LL Technologies

Executive summary

All flash arrays are accelerating the pace of business transformation as IT professionals search for the most relevant technologies to modernize their operation and drive down operational and capital expenditures. Dell EMC™ VMAX™ All Flash arrays are architected to solve the CIO challenge of embracing a modernized flash-centric data center for mission-critical applications while simultaneously simplifying, automating, and consolidating IT operations. VMAX All Flash data services are engineered to use the latest high-density flash technology to deliver high performance at attainable economics. These data services address the new requirements of the modern data center while continuing to deliver the reliability and mission-critical availability that our customers have relied on for years.

DELLTechnologies

1   Introduction

The primary hardware platforms in VMAX All Flash are the Dell EMC VMAX 250F and Dell EMC VMAX 950F. Both platforms use a modular brick-building concept which provides the following:

The VMAX 250F platform provides:
- 2 x 12 core, 2.5 GHz Intel® CPUs yielding 48 cores per engine
- 512 GB, 1 TB, or 2 TB of DDR4 cache per engine
- 1 to 2 V-Bricks per system
- Up to 64 FE ports per system (open systems only)
- Up to 1 PBe capacity per system in a half rack using SAS-3 all-flash storage
- Third-party rack support
- 1 M IOPS (8 K RRH)

The VMAX 950F platform provides:
- 2 x 18 core, 2.8 GHz, Intel® CPUs yielding 72 cores per engine
- 1 TB, 2 TB DDR4 cache per engine
- 1 to 8 V-Bricks/zBricks per system
- Up to 256 FE ports per system
- Up to 4 PBe capacity per system using SAS-2 all-flash storage
- Support for Open Systems, mainframe, or both
- 6.7 M IOPS (8 K RRH)

1.1   Primary VMAX All Flash benefits

The primary benefits that VMAX All Flash offers to our customers are as follows:

- Scalable performance and economics
- Leverage advanced multi-core/multi-threading algorithms and a flash- optimized design to meet strict SLAs for high-demand online transaction processing (OLTP), virtualized applications, and high-growth Oracle and SQL databases
- Scale out performance and scale up capacity to achieve millions of IOPS, PBs of capacity and predictable performance while meeting economic requirements using all flash storage.
- Mission-critical availability – Mission-critical six-nines availability architecture with advanced fault isolation, robust data integrity checking, proven non-disruptive hardware and software upgrades, and data services.
- Hyper consolidation availability – Achieve massive consolidation with support for mixed open, mainframe, IBM® i, and file storage on the same system, simplifying management and significantly lowering overall TCO.

DELLTechnologies

1.2 PowerMaxOS Q3 2020 release

The newest enhancements featured in the PowerMaxOS Q3 2020 release expand the VMAX All Flash offering. Some of the key features in the Q3 2020 release specific to the VMAX All Flash Family are shown in the table below:

Table 1   Q3 2020 PowerMaxOS release core features

Q3 2020 PowerMaxOS release core feature	Value proposition and benefits
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.
IBM Transparent Cloud Tiering (TCT)[1]	A licensed IBM function that offloads all data movement processing-related workloads from the mainframe host when moving data to or from cloud repositories.

1.   Supported on VMAX 950F only through RQP process and requires MFE 8.5.

2 VMAX All Flash overview

To meet the emerging requirements of the enterprise storage environment, Dell Technologies is pleased to provide an all-flash offering called VMAX All Flash. VMAX All Flash differs from legacy VMAX hybrid arrays as they are true all-flash arrays being specifically targeted to meet the storage capacity and performance requirements of the all-flash enterprise data center. VMAX All Flash products are feature-rich all-flash offerings with specific capabilities designed to take advantage of the new higher-capacity flash drives used in the densest configuration possible. VMAX All Flash arrays offer enterprise customers the trusted VMAX data services highly virtualized environments demand, while remaining economically comparable to more traditional storage workloads.

2.1 Terminology

The following table provides a detailed list of VMAX All Flash-related terms and their definitions:

Table 2   Key VMAX All Flash terms and definitions

Term (first use in document)	Equivalent term (after first use in document)	Definition
VMAX All Flash	VMAX All Flash	VMAX All Flash refers to the all-flash VMAX offering.
VMAX 250F	VMAX 250F	VMAX 250F is the entry VMAX All Flash array. The F appliance-based package includes base software titles.
VMAX 250FX	VMAX 250FX	VMAX 250FX ships with the FX package, adding additional software titles to the F bundle.
VMAX 950F	VMAX 950F	VMAX 950F is a scalable VMAX All Flash array. The F appliance-based package includes base software titles.
VMAX 950FX	VMAX 950FX	VMAX 950FX ships with the FX package, adding additional software titles to the F bundle.
V-Brick	V-Brick	A V-Brick is the building block for constructing a VMAX All Flash array. It includes a VMAX engine, two DAEs, and a fixed TBu of capacity.
Flash Capacity Pack	Flash Capacity Pack	A Flash Capacity Pack includes a fixed amount of capacity that can be added to a VMAX All Flash array.
PowerMaxOS 5978	PowerMaxOS	The PowerMaxOS 5978 release supports PowerMax NVMe arrays, dedupe, and other software enhancements and is offered with VMAX All Flash arrays.
Drive Array Enclosure	DAE	DAE refers to the drive array enclosure used to store flash drives in VMAX All Flash

DELLTechnologies

Term (first use in document)	Equivalent term (after first use in document)	Definition
Unisphere for PowerMax V9.0	Unisphere	Unisphere for PowerMax V9.0 enables management and monitoring of PowerMax and VMAX All Flash arrays.
DAE25	DAE25	DAE25 refers to the drive array enclosure used to store flash drives in the VMAX 250F (12 Gb/sec SAS)
inline compression	compression	Inline compression refers to the intelligent data efficiency technology used with VMAX All Flash arrays.
usable capacity (in Terabytes)	TBu	The amount of unique, noncompressed data that can be written into the array.
effective capacity (in Terabytes)	TBe	Includes the benefits of virtual provisioning, inline compression, and space-efficient copies.
non-disruptive migration	NDM	Provides simplified data migration from VMAX1/VMAX2 arrays to VMAX All Flash/V3 systems.
virtual witness	vWitness	Virtual witness is a key component to the SRDF/Metro DR offering, arbitrating between SRDF/Metro sites.
zBrick	zBrick	zBrick is the VMAX All Flash mainframe building block consisting of an engine, two DAEs, and capacity.
zCapacity Pack	zCapacity Pack	zCapacity Pack is the storage unit used to add capacity to VMAX All Flash mainframe systems.
zF software package	zF package	zF software package is the base software offering for VMAX All Flash mainframe arrays.
zFX software package	zFX package	zFX software package is a comprehensive software suite offered with VMAX All Flash mainframe arrays.
scale up	scale up	Scale up refers to adding Flash Capacity Packs to a VMAX All Flash array.
scale out	scale out	Scale out refers to adding V-Bricks to grow performance and expansion for VMAX All Flash systems.
service level	service level	Service level will be used with VMAX All Flash arrays. This term replaces SLOs reflecting the single service level used in this platform.

2.2   The VMAX All Flash family

The Dell EMC VMAX All Flash family offers two primary models: the VMAX 250F and the VMAX 950F. The VMAX All Flash products are feature-rich all-flash offerings with specific capabilities designed to take advantage of the new higher capacity flash drives used in the densest configuration possible. The VMAX All Flash arrays offer enterprise customers the trusted VMAX data services, the improved simplicity, capacity, and performance which their highly virtualized environments demand, while still meeting the economics of the more traditional storage workloads.

VMAX 250F

- 1M IOPS (8K RRH)
- 64 FC/iSCSI Ports
- 1 – 2 V-Bricks per system
- Up to 1 PBe per system
- All in ½ rack (20U)
- Open system workloads only

VMAX 950F

- 6.7M IOPS (8K RRH)
- 256 FC/iSCSI/FICON Ports
- 1 – 8 V-Bricks per system
- Up to 4 PBe per system
- Open system / Mainframe / Mixed workloads

Figure 1   The VMAX All Flash family

3   VMAX All Flash key value propositions

VMAX All Flash is designed for scalable performance, mission critical availability, hyper-consolidation, and to support the densest flash configuration possible by using industry standard high-capacity flash drives. Below is a summary of some of the essential value propositions that VMAX All Flash offers customers:

- **Performance** – Regardless of workload and storage capacity utilization, VMAX All Flash is designed to provide reliable high performance to the enterprise data center, delivering up to 6.7 million IOPS with consistent sub-ms latency at 150 GB/sec bandwidth.
- **High availability and resiliency** – VMAX All Flash is built with trusted architecture featuring no single points of failure, and it has a proven six nines of availability track record. The ability to use SRDF gives customers full multisite replication options for disaster recovery and rapid restart.
- **Inline compression -** Compression is a space-saving function designed to allow the system to manage capacity in the most efficient way possible. Compression is performed inline on VMAX All Flash using multiple compression ranges in order to achieve a 4:1 data reduction average for the system when other data reduction methods are employed as well. Compression is available for open systems workloads. Compression is not supported with mainframe workloads.
- **Non-Disruptive Migration (NDM)** - NDM is designed to help automate the process of migrating hosts and applications to a new VMAX All Flash array with no downtime at all.
- **Non-Volatile Memory Express (NVMe) Flash** – VMAX All Flash is the technology leader using NMVe flash for storing critical system configuration data.
- **Enhancing flash drive endurance** –VMAX All Flash has unique capabilities for greatly minimizing write amplification on the flash drives. It employs large amounts of cache to store writes and then uses intelligent destaging algorithms to coalesce the writes into a larger sequential write, minimizing random write I/O to the back-end. VMAX All Flash also employs proven write-folding algorithms which drastically reduce the amount of write I/O to the back-end.
- **Flash density** – Using standard high-capacity flash drives, VMAX All Flash delivers the highest IOPS/TB/floor tile in the industry. This support for industry standard flash drives provides a differentiated capability as compared to many all-flash alternatives as this allows VMAX All Flash to leverage increases in flash drive densities, economies of scale, and faster time-to-market.
- **Scalability** – VMAX All Flash configurations are built with modular building blocks called "bricks." A brick includes an engine and two drive DAEs preconfigured with an initial total usable capacity. Brick capacity can be scaled up in specific increments of usable capacity called Flash Capacity Packs.
- **Data services** – Full support for the industry's gold standards in remote replication with SRDF and local replication with TimeFinder™ SnapVX. VMAX All Flash integrates fully with Dell EMC AppSync™ for easier local replication management of critical applications. Dell EMC RecoverPoint is also available for all VMAX All Flash models.
- **Consolidation** – VMAX All Flash systems are the only all flash storage products in the industry which can consolidate open systems, mainframe, and IBM i block as well as file storage onto a single floor tile. VMAX All Flash supports many front-end connectivity options including Fibre Channel, iSCSI, and FICON for mainframe.
- **Streamlined packaging** - The VMAX All Flash family features "F" and "FX" model options. The difference in the models is specifically related to the greatly simplified software packaging for the VMAX All Flash product line. The VMAX All Flash base models are the VMAX 250Fand 950F. The base "F" models include an entry-level software packaging with features such as embedded Unisphere. The "FX" models include the entry-level "F" packaging, plus more advanced software offerings such as SRDF.
- **Ease of management** – Embedded Unisphere for PowerMax V9.0 is provided in both the F and FX package. Unisphere's intuitive management interface allows IT managers to maximize productivity by

D∕ELL Technologies

dramatically reducing the time required to provision, manage, and monitor VMAX All Flash storage assets. The fact that Unisphere is embedded within VMAX All Flash allows for this simplicity of management without the need for additional servers and hardware. The FX package also includes Unisphere 360, which lets storage administrators view site-level health reports for every VMAX in the data center, and also coordinate compliance to code levels and other infrastructure maintenance requirements.

3.1 Expandable modular architecture with the V-Brick and zBrick

VMAX All Flash employs a simplified appliance-based modular building block architecture to reduce complexity and simplify configuration and deployment. Building blocks called "bricks" allow VMAX All Flash to scale to deliver predictable high performance where needed.

Two types of bricks are available for the VMAX All Flash:
- The **V-Brick** supports open systems configurations with Fibre Channel or iSCSI connectivity and FBA device formatting.
- The **zBrick** supports mainframe configurations with FICON connectivity and CKD device formatting.

Note: In this document, the term "brick" will be used when discussing features and functions applicable to both the open systems V-Brick and the mainframe zBrick. The zBrick will be discussed in more detail in the VMAX All Flash mainframe support portion of this document.

Each brick has the following components:

- One engine using the dynamic virtual matrix architecture running PowerMaxOS 5978
- Fully redundant hardware with multiple power supplies and interconnecting fabrics

    – No single points of failure architecture
    – Proven six nines availability

- 2 x 2.5" drive slot Drive Array Enclosures (DAEs)

    – VMAX 250F has 2 x 25 slot 2.5" drive,
    – VMAX 950F has 2 x 120 slot 2.5" drive
    – VMAX All Flash starter brick configuration has a set amount of usable capacity
    – Additional V-Brick storage capacity is added in defined increments called "Flash Capacity Packs," while additional zBrick storage capacity is added in defined increments called "zCapacity Packs."

- Up to 32 ports of front-end (FE) connectivity for the VMAX 250F. Up to 24 ports for VMAX 950F (32 for mainframe)
- Up to 2 TB of Cache per brick

The following table details the various VMAX All Flash model brick specifications.

Table 3   Brick specifications by VMAX All Flash model

Component	Specification	VMAX 250F	VMAX 450F[5]	VMAX 850F[5]	VMAX 950F
System Layout	Floor Tile Space Required	1	1–2	1–4	1–4
Compute	Bricks per System	1–2	1–4	1–8	1–8
	Support for mainframe zBrick	No	Yes	Yes	Yes
	Maximum cores per system	96	128	384	576
Cache	Cache per Brick Options	512 GB, 1 TB, and 2 TB	1 TB and 2 TB		
	Mixed cache support	Yes	No	No	No
Ports and Modules	Maximum FE modules per V-Brick	8 (32 total FE ports per V-Brick)	6 (24 total FE ports per V-Brick – open systems/mixed workloads)		
	Maximum FE modules per zBrick	NA	6–8 (24 or 32 total FICON ports per zBrick)[1,2]		
	Maximum FE ports per system	64	96 (OS), 128 (MF)	192 (OS), 256 (MF)	192 (OS), 256 (MF)
Drives and Capacity	Brick DAE Type and QTY	2 x 25 slot, 2.5" (DAE25)	2 x 120 slot, 2.5"		
	Maximum drives per system	100	960	1920	1920
	Maximum open systems effective capacity per system[3, 4]	1 PBe	2 PBe	4 PBe	4 PBe
	Maximum mainframe usable capacity per system[3, 4]	NA	800 TBu	1.7PBu	1.7 PBu
	Starter Brick usable capacity	11 or 13 TBu[3]	53 TBu		53 TBu (OS) 13 TBu (MF)
	Flash Capacity Pack increment size	11 or 13 TBu[3]	13 TBu		
	RAID Options	RAID 5 (3+1), RAID 5 (7+1) RAID 6 (6+2)	RAID 5 (7+1), RAID 6 (14+2)		

Component	Specification	VMAX 250F	VMAX 450F[5]	VMAX 850F[5]	VMAX 950F
	Supported V-Brick Flash Drive Sizes	960 GB, 1.92 TB, 3.84 TB, 7.68 TB, 15.36 TB	960 GB, 1.92 TB, 3.84 TB		960 GB, 1.92 TB, 3.84 TB, 7.68 TB, 15.36 TB
	Supported zBrick Flash Drives Sizes	NA	960 GB, 1.92 TB, 3.84 TB		960 GB, 1.92 TB, 3.84 TB, 7.68 TB, 15.36 TB

1. Default zBrick comes with 2 FICON modules. Extra FICON modules can be ordered separately
2. A zBrick can support up to 32 FE ports if SRDF compression is not used in configuration
3. The VMAX 250F starter V-Brick and capacity increments can be either 11 TBu (RAID 5 3+1) or 13 TBu (RAID 6 6+2, RAID 5 7+1)
4. Dell Technologies uses PBu (and TBu) to define usable storage capacity in the absence of compression, referring to the amount of usable physical storage in the box. Dell Technologies uses PBe (and TBe) to define effective storage capacity in the presence of compression. For example, if a customer has 50 TBu of physical storage, and it is compressible on a 2:1 basis, then the customer has 100 TBe (effective storage).
5. While the VMAX 450F and 850F are still part of the VMAX All Flash portfolio, they are now considered legacy systems having been superseded by the VMAX 250F and VMAX 950F. The VMAX 450F and VMAX 850F will no longer be available for ordering after November 2019.

The brick concept allows VMAX All Flash to scale up and scale out. Customers can scale up by adding Flash Capacity Packs. Each Flash Capacity Pack has a multiple of 13 TBu of usable storage for the VMAX 950F models, and 11 TBu or 13 TBu for the VMAX 250F model, depending upon the RAID protection type selected. VMAX All Flash scales out by aggregating up to two bricks for the VMAX 250F, and up to eight bricks for the VMAX 950F in a single system with fully shared connectivity, processing, and capacity resources. Scaling out a VMAX All Flash system by adding additional bricks produces a predictable, linear performance improvement regardless of the workload.

3.1.1 Engines

The core of the brick is the engine. The engine is the central I/O processing unit, redundantly built for high availability. It consists of redundant directors that each contain multi-core CPUs, memory modules, and attach 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 also allows VMAX All Flash to scale up and out.

DELLTechnologies
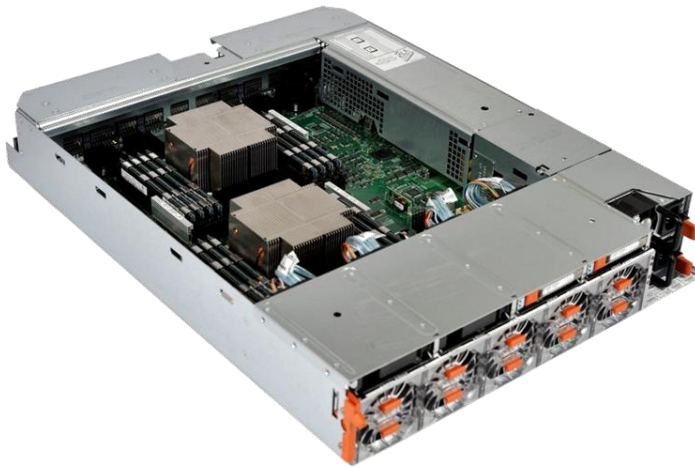
Figure 2   VMAX All Flash Engine Director (2 per engine)

3.1.1.1   Brick Engine CPU core configurations

Each brick engine has two directors. Each director has dual CPU sockets that can support multi-core, multi-threaded Intel processors. The following table details the engine CPU core layout for each VMAX All Flash model:

Table 4   Brick Engine CPU Cores per VMAX All Flash Model

VMAX All Flash Model	Engine CPU type	Cores per director	Cores per Brick	Max Cores per system
250F (V-Brick Only)	Dual Intel Broadwell 12 core	24	48	96 (2 bricks max.)
950F	Dual Intel Broadwell 18 core	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 running on the engine (such as SRDF, eNAS, and embedded management). The core pools can be dynamically 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.

Aside from being able to dynamically adjust the core pools, VMAX All Flash can also implement advanced Quality of Services (QoS), such as setting the maximum amount of IOPS for a particular storage group. This is extremely helpful at properly managing system core consumption so that a "noisy" virtual machine or host cannot overly consume system resources. QoS helps ensure that all connected hosts and virtual machines receive an evenly distributed amount of resources to deliver the maximum performance possible in terms of IOPS and throughput.

3.1.1.2　Brick cache configurations

Every brick 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 Engine Cache Configurations per VMAX All Flash model

VMAX All Flash Model	Cache per Brick	Max Cache per System
250F (V- Brick Only)	512 GB, 1 TB, 2 TB	4 TB (2 bricks max.)
950F	1 TB or 2 TB	16 TB (8 bricks max.)

For dual-brick VMAX All Flash 250F models, the system can use engines with differing cache sizes (mixed memory). For example, the cache for the engine on brick A can be 1 TB while the cache for the engine for brick B can be 512 TB. This would yield a total system cache size of 1.5 TB. The difference in cache sizes between the engines has to be one capacity size smaller or larger. Valid mixed memory configurations for the VMAX 250F are shown in the following table:

Table 6　　　Supported PowerMax 2000 mixed cache configurations

Mixed Cache Configuration	Smallest Engine Cache Size	Largest Engine Cache Size	Total System Cache
Configuration 1	512 GB	1 TB	1.5 TB
Configuration 2	1 TB	2 TB	3 TB

Note: The VMAX 950F model does not support mixed cache sizes between engines. In these systems, the cache size between engines must be equal.

VMAX All Flash supports Dynamic Cache Partitioning (DCP) on the system engines. DCP is a QoS feature which allows specific amounts of cache to be set aside for particular environments such as "production" or "development." Another example would be the separation of cache resources for "file data" from "block data" on systems which are using eNAS services. Being able to fence off and isolate cache resources is a key enabler for multitenant environments,

3.1.1.3　PowerMaxOS

Each VMAX All Flash engine comes preloaded with PowerMaxOS 5978. PowerMaxOS is derived from the trusted and proven HYPERMAX OS used by the legacy VMAX3 and VMAX All Flash arrays; however, PowerMaxOS has been rewritten to take advantage of the NVMe architectures used by the PowerMax family. PowerMaxOS continues to provide industry-leading high availability, I/O management, QoS, data integrity validation, data movement, and data security within an open application platform. Like the previous generation HYPERMAX OS, 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 such as:

- Processing I/O from hosts
- Implementing RAID protection
- Optimizing performance by allowing direct access to hardware resources
- System Management and Monitoring

3.1.2 Drive Array Enclosures and drive configurations

3.1.2.1 VMAX 250F DAEs

Each brick for the VMAX 250F comes with two 25 slot, 2.5" drive, 2U front loading Drive Array Enclosures (DAEs) along with 11 TBu of preconfigured initial capacity using RAID 5 (3+1) protection, or 13 TBu using RAID 5 (7+1) or RAID 6 (6+2) protection. The VMAX 250F requires a minimum configuration of two (2) RAID groups.

The VMAX 250F DAE supports 12 Gb/sec SAS connectivity and requires 12 Gb/sec SAS flash drives. Flash drives which use 6 Gb/sec SAS connectivity are not supported in the VMAX 250F. The VMAX 250F DAE has dual-ported drive slots and dual power zones for high availability.

Figure 3   Fully populated VMAX 250F DAE

Additional scale up capacity is added to the VMAX 250F system using flash capacity pack increments of 11 TBu, scaling up to a maximum of effective capacity of 500 TBe per brick. A dual-brick VMAX 250F can scale up to a total capacity of 1 PBe using a half rack (20U) within a single floor tile footprint.

3.1.2.2 VMAX 950F DAEs

Each brick for the VMAX 950F comes with two 120-slot, 2.5" drive, 4U drawer DAEs. The V-Brick comes with 53 TBu of preconfigured initial capacity that can use either RAID 5 (7+1) or RAID 6 (14+2) protection. The mainframe zBrick comes preconfigured with an initial capacity of 13 TBu for RAID 5 (7+1) or 26 TBu for RAID 6 (14+2).

Figure 4   Fully populated VMAX 950F Brick DAE

Each VMAX 950F DAE uses dual-ported drive slots and four separate power zones to eliminate any single points of failure. The VMAX 950F DAE layout is shown in the diagram below:

Figure 5   VMAX 950F DAE Layout – drive slots and power zones

Additional scale up capacity is added to the system in multiples of 13 TBu increments called "Flash Capacity Packs" for V-Bricks and "zCapacity Packs" for zBricks. The capacity pack concept allows for considerable internal capacity growth over the lifespan of the VMAX 950F arrays – especially when higher capacity flash drives are used. Each VMAX 950F V-Brick can start small with 53 TBu of capacity. The mainframe zBrick capacity can start as small as 13 TBu for the VMAX 950F. As capacity requirements grow, additional capacity can be added into the empty slots of the DAEs. This allows for easy expansion as no extra DAEs would need to be added into the system bay. When a two-brick VMAX 950F system bay is deployed with 2 TB cache engines using 2:1 compression, a customer can have up to 1 PBe of effective flash capacity on a single floor tile while using only 500 TB of physical storage. Using an average of 2:1 compression, VMAX 950F systems can scale to 4 PBe.

3.1.2.3   Important notes about VMAX All Flash DAE capacity

- VMAX All Flash arrays will use a single RAID protection scheme for the entire system. The specific protection scheme is determined by the initial usable capacity of the system. All follow-on capacity and brick additions will use the same RAID protection scheme as the initial usable capacity regardless of the drive size used by the additional flash capacity pack.
- Multiple flash drive sizes can co-exist within the brick DAE for all VMAX All Flash products.
- Brick RAID groups span across both DAEs
- Spare drive requirements are calculated with 1 spare per 50 drives of a particular size, on a per engine basis.

- Mixed System VMAX 950F notes:

  – The VMAX 950F system must be "born" as a mixed system. This means that the system must be originally ordered from the factory to support mixed workloads. For example - it cannot be ordered specifically for open systems workloads, and then be upgraded to support mainframe workloads at a later time, and conversely.
  – Every engine in the mixed workload VMAX 950F will be a V-Brick with 66 TBu of initial capacity (53 TBu for FBA and 13 TBu for CKD) and shipped with a pair of compression modules.
  – The mixed VMAX 950F system will employ two SRPs – one specifically for FBA and the other specifically for CKD. Both SRPs will use the same RAID protection scheme. Each SRP will require its own spare drives as the mixed VMAX 950F SRPs cannot share the same spares even if the drives in each SRP are the same size and type. The same sparing rules apply to the mixed VMAX 950F systems as with the Open System specific and mainframe specific systems.
  – The system can scale up in capacity by adding 13 TBu Flash Capacity, zCapacity packs, or a mix of both. The system can scale out by adding additional V-Bricks (up to eight in total).

3.1.3   VMAX All Flash Director and connectivity options

The VMAX All Flash engine architecture uses a series of hot swappable modules that plug into various slots in the engine directors. These modules include engine cooling fans and power supplies which are accessed by slots from the front of the engine director, along with I/O modules, management modules, and control stations which are accessed by slots from the rear of the engine directors. The following table details the module components used in an engine director:

Table 7   VMAX All Flash engine director components (for V-Brick, zBrick, and Mixed Systems)

Director component	Qty. per Director	Purpose
Power Supply	2	Provides redundant power to director
Fan	5	Provides director cooling
Management Module	1	Manage environmental functionality
Vault to NVMe Flash Module	Up to 4	The flash modules use NVMe technology to safely store data in cache during the vaulting sequence (4 x 800 GB for VMAX 950F, 3 x 400 GB or 800 GB for VMAX 250F)
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, iSCSI, FICON, SRDF, and embedded NAS (eNAS)
NVMe PCIe Back-end I/O Module	Up to 2	Back-end SAS connection to DAEs (1 x 12 Gbps for VMAX 250F, 2 x 6 Gbps for VMAX 950F)
Data Reduction Module	1	Performs inline data compression as well as SRDF compression for the V-Brick. The data reduction module performs SRDF compression only for the zBrick.
Fabric I/O Module	1	Provides connectivity between directors. In multi-engine VMAX All Flash 950 systems, the fabric I/O modules are connected to internal InfiniBand switches

3.1.3.1 Open Systems V-Brick

The VMAX 250F uses up to 3 pairs of Vault to NVMe Flash modules (six modules total) while the VMAX 950F systems use up to 4 pairs of Vault to NVMe Flash modules (eight modules total). The extra flash module required for the VMAX 950F systems is due to the larger usable capacities that these systems can scale to. The vault to flash modules usually will occupy slots 0, 1, and 6 on the VMAX 250F V-Brick engine director while the vault to flash modules will usually occupy slots 0, 1, 6, and 7 on the and VMAX 950F V-Brick engine directors.

The Data Reduction module performs all operations for the Adaptive Compression Engine (ACE) as well SRDF compression for the VMAX All Flash systems. These operations are performed inline on the module. This results in an offloading of the compression task to the module rather than using engine CPU core cycles. Each V-Brick engine will use a pair of data reduction modules (one per each V-Brick director). The compression modules are usually located in director slot 7 on the VMAX 250F and director slot 9 on the VMAX 950F.

The following diagram depicts a typical director module for the VMAX 250F V-Brick engine:
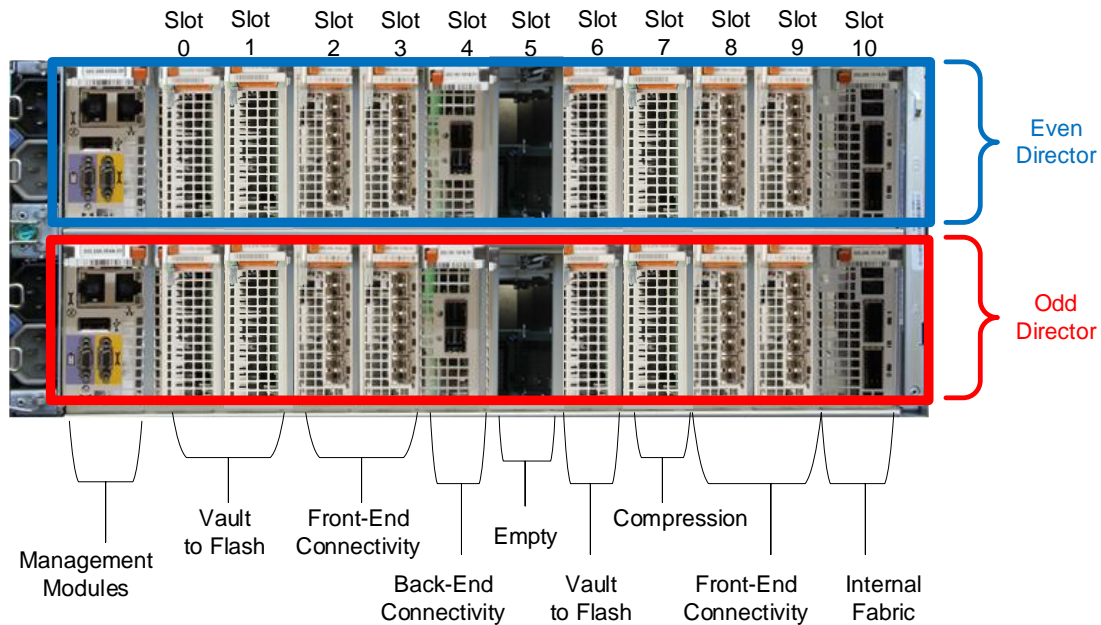
Figure 6   Typical VMAX 250F V-Brick Engine layout

Note: On the VMAX 250F, director slot 5 is left empty (unused).

The following figure depicts a typical director module layout for a VMAX 950F V-Brick engine:
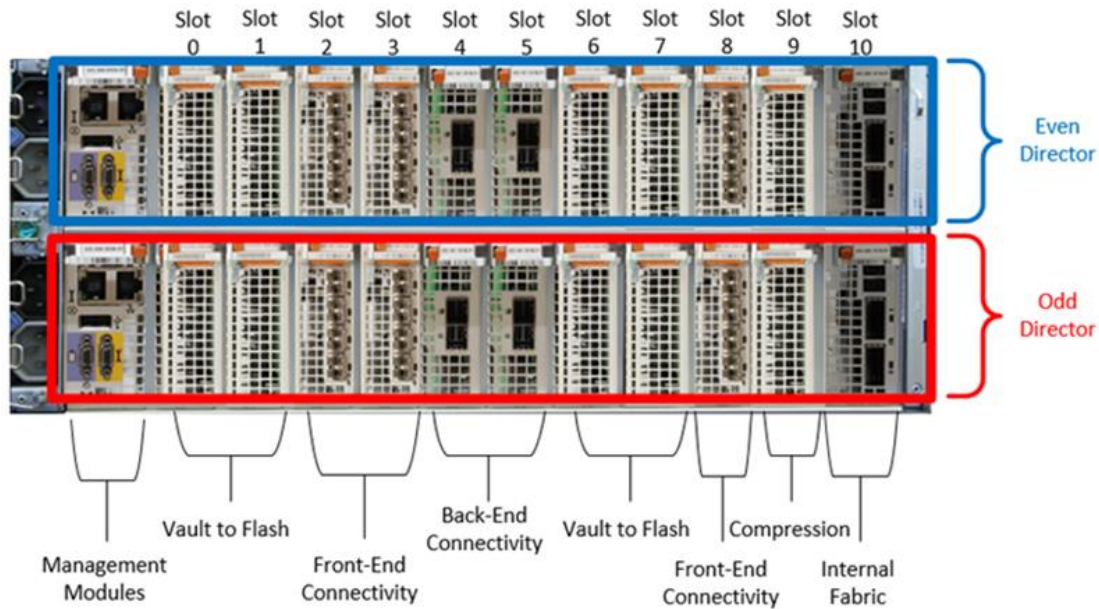
Figure 7　　Typical VMAX 950F V-Brick Engine layout

Multiple supported V-Brick front-end connections are available to support several protocols and speeds. The table below highlights the various front-end connectivity modules available to the VMAX All Flash V-Brick:

Table 8　　VMAX All Flash Open Systems V-Brick front-end connectivity modules

Connectivity Type	Module Type	Number of Ports	Mix With Protocols	Supported Speeds (Gbps)
Fibre Channel	8 Gbps FC	4	SRDF	2 / 4 / 8
Fibre Channel	16 Gbps FC	4	SRDF	2 / 8 / 16
SRDF	10 GigE	4	iSCSI	10
SRDF	GigE	2	None	1
iSCSI	10 GigE	4	SRDF	10
Cloud Array (CA)	8 Gbps FC	4	FC, SRDF	2 / 4 / 8
eNAS	10 GigE	2	None	10
eNAS	10 GigE (Copper)	2	None	10
eNAS Tape Backup	8 Gbps FC	4	None	2 / 4 / 8

3.1.3.2　Mainframe zBrick

For the mainframe zBrick, engine cooling fans and power supplies can be accessed from the front, while the I/O modules, management modules, and control stations can be accessed from the rear. Since the number of universal I/O modules used in the zBrick engine depends on the customer's required functionality, some slots can remain unused.

The zBrick supports FICON and SRDF front-end connectivity. The table below highlights the various front-end connectivity modules available to the VMAX All Flash zBrick:

Table 9   VMAX All Flash mainframe zBrick front-end connectivity modules

Connectivity Type	Module Type	Number of Ports	Mix With Protocols	Supported Speeds (Gbps)
FICON	16 Gbps FICON	4	Single / Multi Mode	4 / 8 / 16
SRDF	16 Gbps Fibre Channel	4	None	4 / 8 / 16
SRDF	8 Gbps Fibre Channel	4	None	4 / 4 / 8
SRDF	10 GigE	4	None	10
SRDF	GigE	2	None	1

Note: The VMAX 250F does not support mainframe workloads.

The quantity of zBrick front-end ports scales to a maximum of 32 when SRDF is not being used. When SRDF is used in the configuration, one of the front-end slots is taken by the SRDF compression module on each engine director. This limits the number of available zBrick front-end ports to 24. By default, each zBrick comes with two FICON modules.

When SRDF is used in the configuration, each zBrick will use a pair of data reduction modules which only perform SRDF compression (one per each zBrick director). The SRDF compression modules are usually located in engine director slot 9 on the VMAX 950F.

The following figure depicts a typical VMAX 950F zBrick engine that is configured for SRDF:
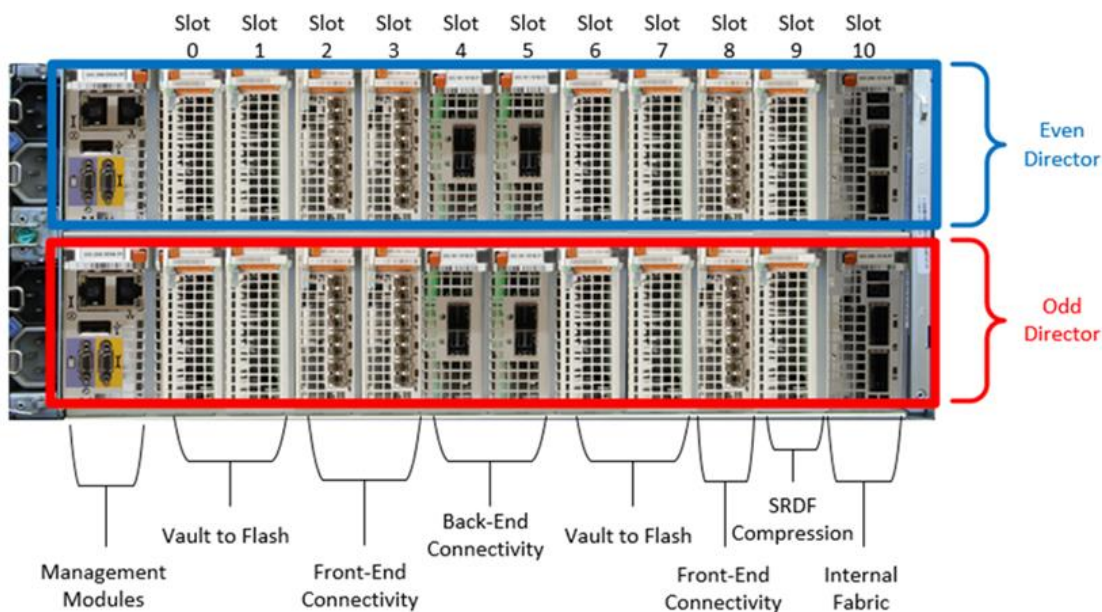

Figure 8   Typical VMAX 950F zBrick Engine Layout with SRDF

D&LL Technologies

When SRDF is not used in the zBrick configuration, a front-end module can be placed into slot 9 providing additional front-end connectivity. The following diagram shows a typical non-SRDF zBrick engine configuration:
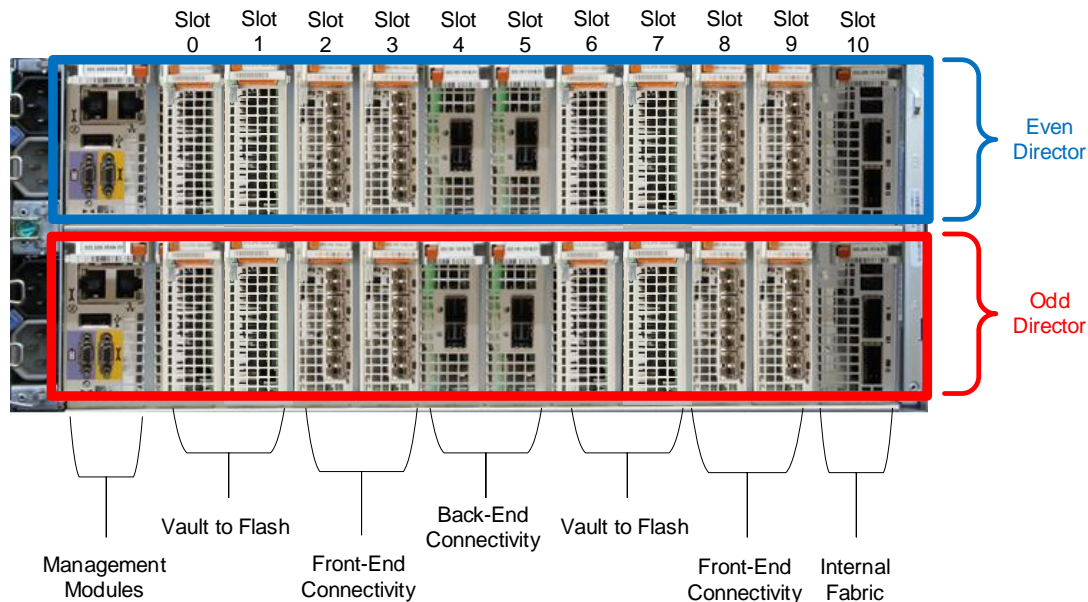
Figure 9   Typical VMAX 950F zBrick Engine Layout without SRDF

The VMAX 950F systems use up to of four pairs NVMe vault to flash modules. The extra flash module pair is required due to the larger usable capacities available in these systems. The NVMe vault to flash modules will usually occupy slots 0, 1, 6, and 7 on the zBrick engine directors.

3.1.3.3   Mixed open systems and mainframe (VMAX 950F only)

Mixed Open Systems and mainframe workloads are allowed only on the VMAX 950F systems, which have specific configuration requirements for the support of mixed workloads:

- The mixed VMAX 950F system must be "born" as a mixed system. This means that the system must be originally ordered from the factory to support mixed workloads. It cannot be ordered specifically for Open Systems workloads, and then be upgraded to support mainframe workloads at a later time, and conversely.
- Every engine in the mixed workload VMAX 950F must be a V-Brick with 66 TBu of initial capacity (53 TBu for FBA and 13 TBu for CKD) and must be shipped with a pair compression modules.
- The mixed VMAX 950F system uses two SRPs – one specifically for FBA and the other specifically for CKD. Both SRPs use the same RAID protection scheme. Each SRP requires its own spare drives as the mixed VMAX 950F SRPs cannot share the same spares even if the drives in each SRP are the same size and type. The same sparing rules apply to the mixed VMAX 950F systems as with the Open System-specific and mainframe-specific systems.
- The system can scale up in capacity by adding 13 TBu Flash Capacity and/or zCapacity packs. The system can scale out by adding additional V-Bricks (up to eight in total).
- The system requires a consistent software packaging level (F and zF or FX and zFX).

The front-end connectivity options for the mixed-system VMAX 950F incorporate the connectivity options for both V-Bricks and zBricks, even though it is a V-Brick engine. The mixed system engine requires four pairs of NVMe vault to flash modules, typically occupying slots 0, 1, 6, and 7 on each engine director. The data

reduction modules will occupy slot 9 on each director. This leaves slots 2, 3, and 8 available for I/O modules on the mixed system V-Brick. Any supported open-systems or mainframe I/O module can be used in these slots. An example mixed VMAX 950F configuration is shown in the diagram below:
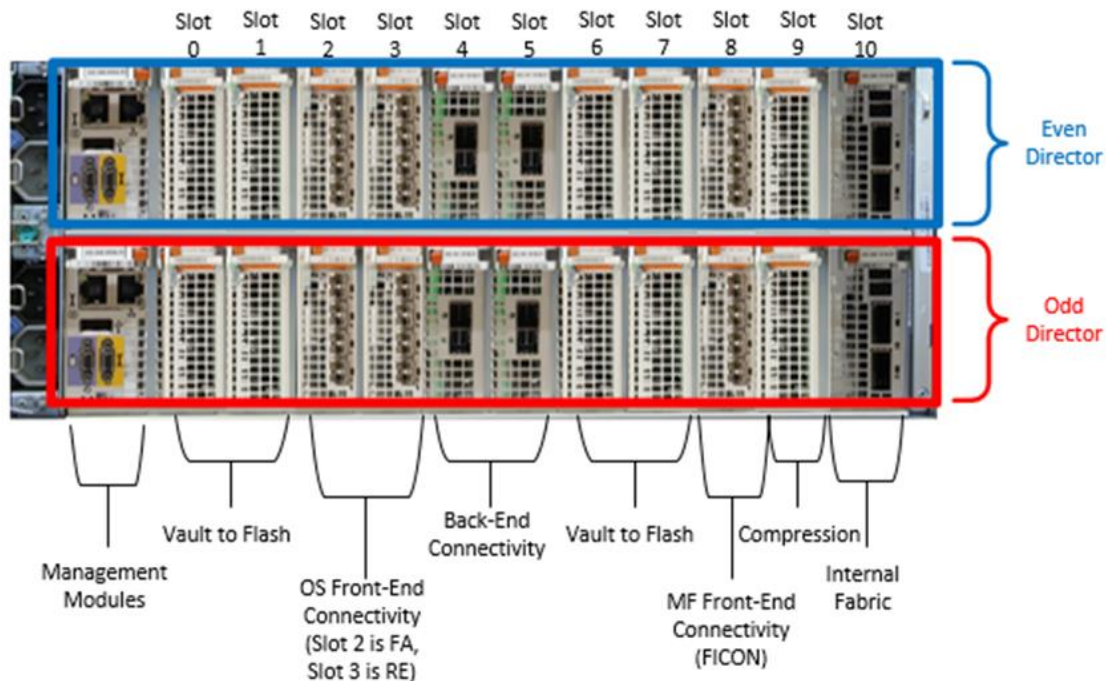
Figure 10   Possible mixed system VMAX 950F V-Brick Engine Layout

3.2   Flash optimization

All-flash-based storage systems demand the highest levels of performance and resiliency from the enterprise storage platforms that support them. The foundation of a true all-flash array is an architecture which can fully leverage the aggregated performance of modern high-density flash drives while maximizing their useful life. VMAX All Flash has several features built into the architecture specifically designed to maximize flash drive performance and longevity. This section will discuss these features in detail.

3.2.1   Cache architecture and caching algorithms

VMAX All Flash 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 cache algorithms used by VMAX All Flash 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

3.2.2 Understanding flash cell endurance

Write cache management is essential to improving performance, but it is also a key part of how VMAX All Flash helps extend the endurance of flash drives. Flash drive longevity and endurance are most impacted by writes, particularly small block random writes. Writing to a flash cell requires that the cell is first erased of any old data and then programmed with the new data. This process is called the Program and Erase Cycle (P/E Cycle). Each flash cell has a finite number of P/E Cycles that it can endure before it wears out (can no longer hold data). Most modern flash cells can endure several thousand P/E Cycles.

One of the peculiarities of flash is that writes are spread out across a flash page (typically KBs in size); however, prior to the write operation, the existing data in entire flash block (typically MBs in size) that the page is located in must be erased. Prior to erasing the page, the flash controller chip finds an empty (erased) location on the drive and copies (writes) any existing data from the page to that location. Because of how flash writes data, a simple 4 KB write from a host could result in many times that amount of data being written internally on the drive, causing P/E cycling on a large number of cells. This write-multiplying effect is called "write amplification" and is detrimental to flash cell endurance. This effect is even more dramatic with small block random write workloads. In this situation, a large number of small block random writes tends to "buckshot" across the drive, impacting an even greater number of cells and invoking P/E cycling on a much larger cell area. Write amplification is not nearly as significant with larger sequential writes as this data is written sequentially local to a single flash block, thereby aligning better with flash page sizes and containing the P/E cycling to a smaller area.

3.2.3 VMAX All Flash write amplification reduction

Write amplification must be properly controlled and mitigated in order to ensure the longevity of flash devices as uncontrolled write amplification is the number one reason for premature wear out of flash storage. Controlling flash cell write amplification is one of VMAX All Flash's greatest strengths and is what sets it truly apart from other competitors' flash arrays. Aside from using intelligent caching algorithms that keep data in cache as long as possible, VMAX All Flash employs additional methods to minimize the amount of writes to flash. These methods are:

- **Write folding** – Write folding avoids unnecessary drive I/Os when hosts rewrite to a particular address range. This rewritten data is simply replaced in cache and never written to the flash drive. Write folding can reduce writes to the flash 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 flash drives align much better with the page sizes within the flash drive itself. Using write coalescing, VMAX All Flash can take a highly random write host I/O workload and make it appear as a sequential write workload to the flash drives.
- **Advanced wear analytics** - VMAX All Flash also includes advanced drive wear analytics optimized for high capacity flash drives to make sure writes are distributed across the entire flash pool to balance the load and avoid excessive writes and wear to particular drives. Not only does this help manage the flash drives in the storage pools, it makes it easy to add and rebalance additional storage into the system.

All of the write amplification reduction techniques used by VMAX All Flash result in a significant reduction in writes to the back-end, which in turn significantly increases the longevity of the flash drives used in the array.

3.2.4 Boosting flash performance with PowerMaxOS FlashBoost

Dell Technologies strives to improve performance in its products. With every new hardware platform and release of software, we try to remove potential bottlenecks that can impede performance. One feature that Dell Technologies introduced and has made standard as a part of PowerMaxOS is FlashBoost, which

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.

3.3 Reliability, availability, and serviceability

VMAX All Flash 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.

VMAX All Flash 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 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 any time data is moved within the subsystem, allowing the array to provide true end-to-end integrity checking and protection against hardware or software faults.

The protection metadata is appended to the data stream, and contains information describing the expected data location as well as CRC representation of the actual data contents. The expected values to be found in protection metadata are stored persistently in an area separate from the data stream. The protection metadata is used to validate the logical correctness of data being moved within the array any time the data transitions between protocol chips, internal buffers, internal data fabric endpoints, system cache, and system disks.

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. Additional protections for address/control fault modes are defined in user definable blocks supported by the T10 standard, with address and write status information in the extra bytes in the application tag and reference tag portions of the block CRC.

VMAX All Flash 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 VMAX All Flash RAS features are summarized below:

- No single points of failure – all components are fully redundant to withstand any component failure
- Completely redundant and hot-pluggable field-replaceable units (FRUs) to ensure repair without taking the system offline
- Choice of RAID 5 or RAID 6 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 particular drive is nearing its end of useful life.

- Vault to flash with battery backup to allow for cache destage to flash and an orderly shutdown for data protection in the event of a power failure
- Active/active remote replication over 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
- Data at Rest Encryption (D@RE) with integrated RSA key manager, FIPS 140-2 compliant 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 key 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 VMAX All Flash RAS capabilities, see the document Dell EMC VMAX All Flash Reliability, Availability, and Serviceability.

3.4   Data Services

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

- Advanced data reduction using inline compression
- Remote replication with SRDF
- Local replication with TimeFinder SnapVX
- Embedded NAS (eNAS)
- eManagement – embedded Unisphere for PowerMax

3.4.1   Data reduction using the Adaptive Compression Engine

VMAX All Flash employs inline hardware compression using the Adaptive Compression Engine (ACE). ACE provides VMAX All Flash and its customers with a data reduction method which delivers the highest space

saving capability with negligible performance impact. The following design factors make the Dell EMC Adaptive Compression Engine unique in the industry:

- **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 reduces the compression function from consuming critical VMAX All Flash 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 busiest (hot) data in the system to bypass the compression workflow. This ensures that all data in the system will receive the appropriate compression focus while maintaining optimal response time.
- **Fine grain data packing** - Fine grain data packing includes a zero reclaim function that prevents the allocation of buffers with all zeros or no actual data.

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

3.4.2   Remote replication with SRDF

SRDF is perhaps the most popular data service in the enterprise data center because it is considered a 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 replicate tens of thousands of volumes to a maximum of four different locations globally.

VMAX All Flash runs an enhanced version of SRDF specific for all-flash use cases. This version uses multi-core, multithreading 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 three types of SRDF:

- **SRDF Synchronous (SRDF/S)** – SRDF/S delivers zero data loss remote mirroring between data centers separated by up to 60 miles (100 km).
- **SRDF Asynchronous (SRDF/A)** – SRDF/A delivers asynchronous remote data replication between data centers up to 8,000 miles (12,875 km) apart. SRDF/A can be used 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

D∕ELL Technologies

- **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 VMAX All Flash FX and zFX software package**s**, with no capacity-based licensing. It can be ordered a la carte as an addition to the F and zF software package**s**. Any hardware needed to support SRDF must be purchased separately.

3.4.3   Local replication with TimeFinder SnapVX

Every VMAX All Flash 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 1,024 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, share back-end allocations with the source volume and other snapshots of the source volume, and only consume additional space when the source volume is changed. A single source volume can have up to 256 snapshots.

Each snapshot has a user-defined name and can optionally have an expiration date, both of which can be modified later. New 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 by linking it to a host-accessible volume referred to as a target. The target volumes are standard thin LUNs. Up to 1,024 target volumes can be linked to the snapshots of a single source volume. This limit can be achieved either by linking all 1,024 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 can only be linked to one snapshot at a time.

By default, targets are linked in a no-copy mode. This no-copy linked target functionality greatly reduces the amount of writes to the back-end flash drives because this eliminates the need to perform a full volume copy of the source volume during the unlink operation in order 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 VMAX All Flash array.

3.4.4   Consolidation of block and file storage using eNAS

eNAS data service extends the value of VMAX All Flash 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 by up to 33%. VMAX All Flash with the eNAS data service becomes a unified block and file platform, using a multi-controller, transaction 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.

Embedded NAS (eNAS) uses the hypervisor provided in PowerMaxOS to create and run a set of virtual machines within the VMAX All Flash 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 VMAX All Flash resources for both performance and capacity.

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

- Scalability – easily serve over 6000 active SMB connections
- Meta-data 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 FX software package. It can be ordered as an additional item with the F software package. All hardware required to support eNAS on VMAX All Flash must be purchased separately.

3.4.5   Embedded Management (eManagement) using Unisphere for PowerMax

VMAX All Flash customers can take advantage of simplified array management using embedded 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 VMAX All Flash storage assets.

Embedded Unisphere enables customers to simplify management, reduce cost, and increase availability by running VMAX All Flash management software directly on the array. Embedded management 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, and SMI-S management software.

Unisphere for PowerMax delivers the simplification, flexibility, and automation required 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 VMAX All Flash, 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 that all applications will receive sub-millisecond response times. Using Unisphere for PowerMax, a customer can set up a multisite SRDF configuration in a matter of minutes.

Embedded Unisphere for PowerMax is a great way to manage a single VMAX All Flash array; however, for customers who need to view and manage their entire data center, there is Unisphere 360. Unisphere 360 software 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, VMAX All Flash, and legacy system VMAX or coordinate compliance to code levels and other infrastructure maintenance requirements. Customers can leverage the simplification of VMAX All Flash management, at data center scale.

Embedded Unisphere and Database Storage Analyzer are available on every VMAX All Flash array as they are included in the F and zF software packages. Unisphere 360 is included in the FX and zFX software packages or can be ordered separately with the F and zF software packages. Unisphere 360 does not run in an embedded environment and requires additional customer supplied server hardware.

3.4.6   PowerMaxOS service levels and host I/O limits

Service Levels for PowerMaxOS provide customers with the ability to separate applications based on performance requirements and business importance. PowerMaxOS will give the ability to set specified Service Levels to ensure the highest priority application response times are not impacted by lower priority application IOPs.

Service Levels address the requirements of customers to ensure the 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 set the initial, as well as, change the performance level of an application as needed. A Service Level can be applied to a storage group using the PowerMax management tools (Unisphere for PowerMax, REST API, Solutions Enabler, and SMIS-S).

Service Levels can be used along with Host I/O Limits to make application performance more predictable while enforcing a specified service level. 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 storage group that has a specified Service Level to throttle IOPs on applications that are exceeding expected Service Level desired performance.

There are six PowerMaxOS service levels to choose from:

Table 10   PowerMaxOS Service Levels

Service Level	Expected Average Response Time[1]	Penalizes
Diamond (Highest Priority)	0.6 ms	Platinum, Gold, Silver, Bronze
Platinum	0.8 ms	Gold, Silver, Bronze
Gold	1 ms	Silver, Bronze
Silver	3.6 ms	Bronze
Bronze (lowest Priority)	7.2 ms	
Optimized	-	

1. Diamond, Platinum, and Gold service levels have an upper limit but no lower limit, ensuring IO will be serviced as fast as possible. Silver and Bronze service levels have both an upper and lower limit designed to allow higher priority IOPs to be unaffected. Storage groups that are set to Optimized will be throttled to allow high priority IOPS on all service levels with the exception of Bronze.

Note: For more information about PowerMaxOS Service Levels, see the document Dell EMC PowerMax Service Levels for PowerMaxOS.

PowerMaxOS Service Levels and Host IO Limits are available at no additional cost for both PowerMax systems and VMAX All Flash systems which are running PowerMaxOS 5978.

3.4.7   Non-Disruptive Migration

Data migrations have always been challenging in an enterprise environment. The complexity and size of very large 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 storage array. Dell EMC VMAX Non-Disruptive Migration (NDM) allows customers to perform online open systems data migrations that are simple and completely non-disruptive to the host and application.

NDM is available for VMAX All Flash storage arrays running PowerMaxOS, HYPERMAX OS 5977 Q3 2016 Service Release or later with Solutions Enabler 8.3 or later. NDM is designed to help automate the process of migrating hosts and applications to a new PowerMax or VMAX All Flash 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 multi-pathing solution, to manage host access to the data during the migration process.

 NDM offers these benefits:

- Allows migration between VMAX to VMAX3 to newer VMAX All Flash or PowerMax arrays with hosts and applications completely online
- 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 prior to 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.

Note: NDM currently does not support mainframe CKD devices

3.5 The Dell Technologies Future-Proof Loyalty Program

The Dell Technologies Future-Proof Loyalty 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 PowerMax, VMAX All Flash, XtremIO™ X2, SC Series, Dell EMC Unity, Data Domain™, Integrated Data Protection Appliance (IDPA), Isilon™, and ECS™ appliance. This program provides our customers with the following benefits:

- Full Portfolio Program:

  – 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 storage, data protection appliances, or HCI product offerings
  – Predictable Support Pricing – consistent and predictable maintenance pricing and services for your storage appliances

- Select Storage Program Offers – specific to VMAX All Flash

  – 4:1 All Flash Storage Efficiency Guarantee – Get effective logical capacity at least 4X the purchased physical capacity
  – Never-Worry Data Migrations – Use integrated data migration tools with seamless upgrades to move to next generation storage
  – All-Inclusive Software with the FX and zFX packages

For more information about the Dell Technologies Future-Proof Loyalty Program, contact Dell Technologies sales.

3.6 Streamlined software packaging

VMAX All Flash arrays are built for simplicity and ease of ordering, with appliance-based packaging that combines both hardware and software elements. VMAX All Flash offers two open systems packages called the "F Software Package" and the "FX Software Package."

Two mainframe-specific software packages are offered with the VMAX All Flash arrays that support mainframe. These mainframe software packages are called the "zF Software Package" and "the zFX Software Package."

3.6.1 VMAX All Flash open systems software packaging

In order to simplify the software ordering and management processes, VMAX All Flash offers two different software packages for the VMAX 250F/950F in open systems environments.

The first option is known as the "F package" which can be considered a starter package. The F package includes PowerMaxOS, Embedded Management, SnapVX, and an AppSync starter pack. Any software title

supported on VMAX All Flash can be added to the F package as a separate software addition. The VMAX All Flash models that use the F package are the VMAX 250F and the VMAX 950F.

The second option – the more encompassing package – is known as the "FX package." The FX package includes everything t in the F package plus SRDF/S, SRDF/A, SRDF/STAR, SRDF/Metro, CloudIQ, D@RE, eNAS, Unisphere 360 and SRM. The FX is priced to offer a bundled discount over an equivalent F option which has many a la carte titles as additions. Customers can add any title supported on VMAX All Flash to the FX in an a la carte fashion as well, including Dell EMC PowerProtect Storage Direct (formerly known as ProtectPoint), the full AppSync suite, and Dell EMC Storage Analytics. VMAX All Flash models that use the FX package are the VMAX 250FX and the VMAX 950FX.

The following table details what software is included in each VMAX All Flash package.

Table 11   VMAX All Flash Systems Software Packaging Options

Feature	F package included	F package a la carte	FX package Included	FX package a la carte	Notes
PowerMaxOS	✓		✓		Includes Migration Tools, VMware® vSphere® Virtual Volumes™ (vVols), Service Levels, QoS[3]
Embedded Management	✓		✓		Includes Unisphere for PowerMax, Database Storage Analyzer, Solutions Enabler, REST APIs, SMI-S
Advanced Data Reduction	✓		✓		Includes inline compression only
Local Replication	✓		✓		Includes TimeFinder SnapVX and Snapshot Policies
AppSync iCDM Starter Bundle	✓		✓		
Non-Disruptive Migration	✓		✓		
Remote Replication Suite[1]		✓	✓		Includes SRDF/S/A/STAR
SRDF/Metro[1]		✓	✓		Includes SRDF/Metro SmartDR
Unisphere 360		✓	✓		
D@RE[2]		✓	✓		
eNAS[1, 2]		✓	✓		
SRM		✓	✓		
PowerPath		✓	✓		Includes 75 host licenses
AppSync Advanced		✓		✓	

D&LLTechnologies

Feature	F package included	F package a la carte	FX package Included	FX package a la carte	Notes
Storage Direct (ProtectPoint)		✓		✓	
RecoverPoint		✓		✓	
Dell EMC Storage Analytics		✓		✓	

1. Software packages include software licensing. Required hardware must be ordered separately.
2. Factory configured. Must be enabled during the ordering process.
3. Includes host I/O limits.

3.6.2   VMAX All Flash mainframe software packaging

Software for mainframe support comes in two packages:

- zF – the basic package
- zFX – a larger bundle with more advanced features

Additionally, many software features for mainframe can be ordered separately. The packages are different from the standard all-flash packages and represent the core functionalities used by the mainframe customer. The following table highlights the VMAX All Flash for mainframe software packaging:

Table 12   VMAX All Flash mainframe Software Packaging Options (not available for VMAX 250F)

Feature	zF Package Included	zF Package a la Carte	zFX Package Included	zFX Package a la Carte	Notes
PowerMaxOS	✓		✓		Includes Migration Tools, QoS
Embedded Management	✓		✓		Includes Unisphere for PowerMax, Database Storage Analyzer, Solutions Enabler, REST APIs, SMI-S
Local Replication	✓		✓		Includes TimeFinder SnapVX
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
Unisphere 360		✓	✓		
AutoSwap		✓	✓		
D@RE[2]		✓	✓		
zDP		✓	✓		

DELLTechnologies

Feature	zF Package Included	zF Package a la Carte	zFX Package Included	zFX Package a la Carte	Notes
Mainframe Essentials Plus		✓		✓	zBoost PAV Optimizer
GDDR[3]		✓		✓	

1. Software packages include software licensing. Any additional required hardware must be ordered separately.
2. Factory configured. Must be enabled during the ordering process.
3. Use of SRDF/STAR for mainframe requires GDDR.

Note: For mixed open systems and mainframe configurations, the VMAX All Flash array requires a consistent software packaging level (F and zF or FX and zFX).

D∕∕LL Technologies

4 VMAX All Flash deployment models

The Dell EMC VMAX All Flash family is designed to offer customers a flexible all-flash storage platform that provides extreme IOPS density per system in a one to four floor-tile footprint. This section describes the deployable system layouts for the VMAX All Flash systems. It does not discuss available drive configurations and system usable capacities as these topics were previously discussed in the DAE overview section of this document.

4.1 VMAX 250F system configurations

The VMAX 250F is the entry point into the VMAX All Flash family. It is ideally suited for customers with modest capacity requirements who still want to take advantage of the enterprise capabilities of VMAX All Flash at midrange prices.

4.1.1 VMAX 250F configurations

The VMAX 250F can be configured using either one or two V-Bricks in a single standard Dell EMC Titan rack. Each V-Brick consumes 10U of rack space (20U maximum for dual V-Brick VMAX 250F systems). The initial V-Brick will occupy the bottom 10U of the rack when shipped from Dell Technologies manufacturing. The second V-Brick will occupy the 10U directly above the initial V-Brick. This is applicable for systems ordered as dual V-Bricks or scale-out systems. The remaining rack space can be occupied by an additional VMAX 250F system or by customer hardware such as servers and switches. The VMAX 250F does not feature a system tray, KVM, or internal Ethernet or InfiniBand switches. It uses direct InfiniBand connections between engines on dual PowerBrick systems
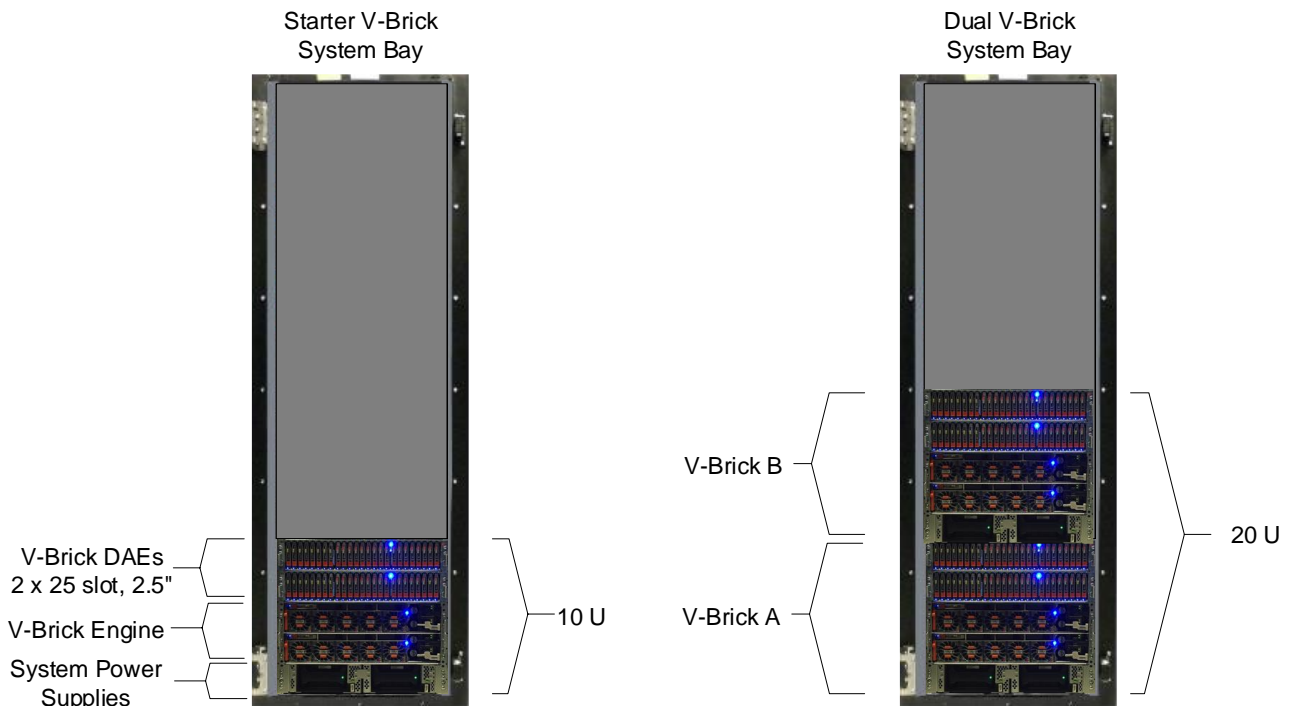

Figure 11   Single and dual V-Brick VMAX 250F configurations

Note: The VMAX 250F supports the use of third party racking. The third-party rack must be a standard NEMA 19-inch rack and meet Dell Technologies standards for cable access and cooling. For more information about supported VMAX 250F third-party racking options, see the Dell EMC VMAX All Flash Site Planning Guide.

4.2 VMAX 950F

The VMAX 950F is designed to provide Dell Technologies customers with the highest levels of scalability, performance, and mission-critical availability. It can consolidate disparate open systems, mainframe, IBM i, and file workloads into a single all-flash system. It can scale to support over 6.7 million IOPS (8K RRH) and provide up to 4 PB of effective capacity in four floor tiles of space.

4.2.1 VMAX 950F system configurations

The VMAX 950F is a highly configurable storage array that can support configurations from 1 to 8 V-Bricks within four standard Dell EMC Titan racks (system bays). Each rack can support up to two V-Bricks with up to 480 flash drives in a rack (1 PBe per rack).

Notes: The system bays on the VMAX 950F can be separated by up to 82 feet (25 meters) using optical connectors.

The following diagram shows the 950F starter V-Brick system bay configuration along with a dual V-Brick system bay configuration:
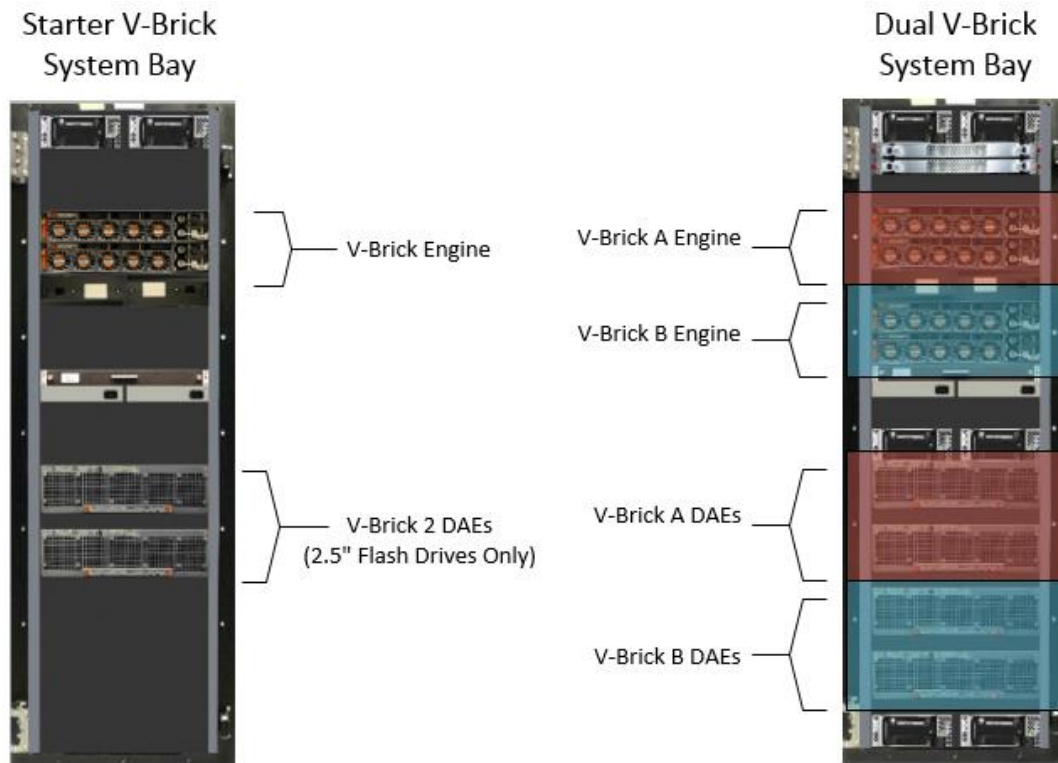

Figure 12   VMAX 950F Starter V-Brick and dual V-Brick system bay configurations

D&LLTechnologies

5 Summary

VMAX All Flash is an all-flash array designed for the most demanding and critical workloads in the enterprise data center. Its unique modular architecture allows it to massively scale while delivering predictable high performance regardless of the workload. Built into the array are complex algorithms which maximize flash performance while greatly enhancing flash drive endurance. Its unique data services and highly available architecture make it a premier choice for the enterprise environment where ease-of-use coupled with trusted dependability is an essential requirement.

D∕ELL Technologies

A Technical support and resources

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.

A.1 Related resources

Document Title	Collateral Type	Part Number
Dell EMC VMAX All Flash Product Guide	Product Guide	
Dell EMC PowerMax: Service Levels for PowerMaxOS	White Paper	H17108
Dell EMC PowerMax and VMAX All Flash: Embedded Management	White Paper	H16856
Dell EMC PowerMax: Data Reduction	White Paper	H17072
Dell EMC PowerMax: Reliability, Availability, and Serviceability	White Paper	H17064
Dell EMC PowerMax and VMAX All Flash: GDPS and Advanced Copy Services Compatibility	White Paper	H16124
Dell EMC VMAX3 and VMAX All Flash Quality of Service Controls for Multitenant Environments	White Paper	H15498
Dell EMC VMAX All Flash Site Planning Guide	Technical Guide	
Dell EMC PowerMax and VMAX All Flash: Timefinder SnapVX Local Replication	Technical Guide	H13697
Dell EMC PowerMax and VMAX All Flash: SRDF/Metro Overview and Best Practices	Technical Guide	H14556
Oracle Ready Bundle Validation Guide	Technical Guide	H16631
Dell EMC Ready Bundle for Oracle	Solution Overview	

D&LL Technologies