

WHITE PAPER

PURE STORAGE FLASHARRAY//X VS. DELL EMC POWERMAX

AN ARCHITECTURAL COMPARISON

TABLE OF CONTENTS

- OVERVIEW** 3
- STORAGE EFFICIENCY** 4
 - Dell EMC PowerMax 4
 - Pure FlashArray//X 4
 - Comparison 4
- STORAGE MANAGEMENT AND AUTOMATION** 5
 - Dell EMC PowerMax 5
 - Pure FlashArray//X 6
 - Comparison 7
- HIGH AVAILABILITY** 7
 - Dell EMC PowerMax 7
 - Pure FlashArray//X 8
 - Comparison 8
- ARCHITECTURE** 8
 - Dell EMC PowerMax 8
 - Pure FlashArray//X 9
 - Comparison 9
- TOTAL COST OF OWNERSHIP (TCO)** 10
 - Dell EMC PowerMax 10
 - Pure FlashArray//X 10
 - Comparison 11
- SUMMARY** 12

OVERVIEW

This whitepaper provides Pure's take on an architectural comparison of two industry-leading enterprise-class [all-flash storage solutions](#): the Pure Storage® FlashArray//X and the Dell EMC PowerMax.

The Pure FlashArray™ family consists of five different models: the //X10, //X20, //X50, //X70, and //X90 – scaling from 20TBe all the way to 3PBe in 6U using a highly available, dual-controller architecture with the ability to non-disruptively scale from one model to the next as performance and capacity requirements increase. From its inception, the Pure Storage architecture was purpose-built as an [all-flash storage array](#), and is engineered to be able to non-disruptively upgrade with data in place and without the need for data migrations.

The Dell EMC [PowerMax](#) family consists of two models, PowerMax 2000 and 8000, which support maximum capacities of 500TBe and 4PBe respectively, in half of a full rack for the PowerMax 2000 and up to two full racks for the PowerMax 8000. The PowerMax 8000 scales from one to eight dual-controller engines, now called Bricks, while the PowerMax 2000 supports either one or two Bricks with slightly slower and fewer CPU cores than the PowerMax 8000.

The PowerMax evolved from EMC's previous VMAX system, which supported a mix of disk-based and flash storage, and which in turn evolved from EMC's Symmetrix and Symmetrix DMX systems; these originally supported high-latency HDD media. EMC's first use of flash technology was with the addition of flash drives to DMX4 systems in 2008. While EMC's original Symmetrix architecture was designed to overcome the latencies of magnetic, rotating HDDs, in contrast to Pure's FlashArray, which is now in its fourth generation using [NVMe](#), the Dell EMC PowerMax architecture has been retrofitted, and now also uses NVMe to communicate to its back-end storage devices – the first generation of its architecture to do so.

Although both architectures provide an all [NVMe flash storage](#) system that is both highly available and performant, there are many key differences to consider. This paper will examine 5 essential criteria to consider when evaluating an enterprise-class all-flash storage array:

- Storage Efficiency
- Storage Management and Automation
- High Availability
- Environmentals
- Total Cost of Ownership (TCO)

STORAGE EFFICIENCY

Dell EMC PowerMax

Dell has recently acknowledged that the previous VMAX all-flash systems have achieved an average data reduction of 1.98:1 through the use of compression.¹ The PowerMax uses a new and different compression algorithm, as well as an additional new post-ingestion data compression algorithm, so it is reasonable to expect that compression on the PowerMax should be improved over VMAX.

Dell EMC, through its Future-Proof Loyalty Program, guarantees a 4:1 storage efficiency across its range of all-flash storage arrays, based on a combination of technologies that include compression, deduplication, thin provisioning, and snapshots. However, when considering data reduction alone, which includes compression and deduplication, Dell EMC is [claiming an average 3:1 data reduction ratio](#) for PowerMax.

The PowerMax uses a 128KB comparison size for deduplication, which means that more granular data is unlikely to get deduplicated. With many applications writing at an 8KB block size, customers may not achieve significant data reduction from the PowerMax deduplication feature unless they are using very large block sizes (greater than 128KB). Additionally, Dell EMC notes that [up to the most “active” 20% of data will not receive data reduction at all](#). Data reduction can be optionally enabled or disabled for each storage group.



Pure FlashArray//X

The Pure FlashArray//X is engineered to provide always-on data deduplication, compression, thin provisioning, encryption, snapshotting, ActiveCluster multi-site synchronous mirroring, Quality of Service (QoS) without any performance impact. The FlashArray//X delivers an average [data reduction](#) ratio (compression and deduplication only) of 5:1 data reduction. When combined with thin provisioning, the ratio increases to an over [10:1 average total efficiency rate](#). Pure Storage provides a simple, all-inclusive [storage software](#) licensing model.



Comparison

Dell EMC is [claiming](#) an average data reduction of 50% improvement over the previous generation VMAX 950F, and “up to” 5:1 for PowerMax, but is willing to only guarantee up to 4:1 when including additional potential efficiencies from thin provisioning and snapshots. 50% improvement over the VMAX 950F’s average of 1.98:1 yields an expectation of 2.97:1, which is consistent with Dell EMC’s claimed expectation of 3:1.

¹ <https://www.youtube.com/watch?v=gBvdXY0WnEg>

The average 3:1 data reduction rate achieved by the PowerMax should be compared to the FlashArray//X's average data reduction rate of 5:1, which provides a 67% improvement over the PowerMax. Similarly, Pure's 10:1 average total efficiency should be directly compared to Dell's 4:1 storage efficiency guarantee claim.

In general, the FlashArray//X can deliver [total efficiency](#) that is typically 2x more efficient than the PowerMax when measured against comparable [data efficiency technologies that include deduplication, compression, and thin provisioning](#). In part, this is due to the five different ways in which Pure performs data reduction. With Pure, the lookup data unit size is 4KB, the lookup data unit alignment is 512B, and for data comparison Pure uses the matched 4KB data unit as an anchor point to which it then extends the match comparison before and after the anchor point in 512B increments until a unique segment found.

Comparing Pure FlashArray//X to the Dell EMC PowerMax:

- The FlashArray//X lookup size is 4KB vs. 128KB. **FlashArray//X should find more matches.**
- The FlashArray//X lookup alignment is 512 Bytes vs. 128KB. **FlashArray//X is adaptive to various real-world workloads.**
- The FlashArray//X incremental match granularity is 512B vs. 128KB fixed-match. **FlashArray//X should be able to deduplicate more data.**

While the PowerMax provides the customer with the ability to turn data reduction on or off per storage group, which gives the customer control over what data will be made more efficient, Pure Storage believes that since performance is not impacted, customers should aim to optimize economics, which are achieved when data reduction is built-in, and always-on. This also has the benefit of reducing the amount of management.

The differences between the two systems become obvious in [mission-critical enterprise applications](#) such as those that use the Oracle database. For example, if a user is storing Oracle databases on a PowerMax, Pure believes that customers will not see any benefits from deduplication for primary [Oracle databases](#) on PowerMax, since each Oracle block has a unique header that is most likely to misalign with the PowerMax 128KB fixed block lookup. In contrast, Pure FlashArray//X can typically provide [deduplication value for Oracle databases](#) due to its more granular and variable comparison size.

STORAGE MANAGEMENT AND AUTOMATION

Dell EMC PowerMax

Unisphere is the primary software set of tools for provisioning, managing, and monitoring PowerMax. It can be run on a dedicated server or run embedded directly in a container on the PowerMax OS hypervisor. In addition to extensive monitoring features, Unisphere provides a web-based interface to manage many system features, including managing user accounts, role-based access control, creating thin volumes, masking volumes, setting storage attributes, setting volume attributes, setting port flags, managing replication and backup operations, managing service levels, workload planning, performance troubleshooting, and much more.

Users have a high degree of control over each of the features available on the PowerMax, and more often than not, users are required to provide their preferences for each of these features as part of the initial configuration and on-going optimization of the system.

To best understand the features and management of the PowerMax, Dell EMC has a 16-hour [online learning training](#).²

Pure FlashArray//X

Pure1® [storage management software](#) is a SaaS-based provisioning, management, and monitoring solution that integrates with Pure's proactive support. It leverages machine-learning and predictive analytics to help advise customers on optimization and what-if situations, including [capacity planning, and performance simulation](#).

Pure1 capabilities are built on a global predictive intelligence engine called Pure1 Meta™ that leverages the accumulated data from the thousands of FlashArrays currently deployed. Pure1 Meta is the AI engine within Pure1 that provides the [intelligence to manage, automate, and proactively support the FlashArray](#). Pure1 Meta collects more than a trillion telemetry data points of performance data per day. Part of the intelligence of Pure1 comes by way of its ability to recognize usage patterns. Pure1 identifies known patterns that may affect the optimal operations of FlashArrays, and notifies other FlashArrays with similar usage patterns of the concern. This way, customers are aware of potential impacts to their arrays and can proactively take preventive measures.

Pure1 is browser-based and can be run without additional software to license, install, upgrade, manage, or change security policies. As new features become available, customers have immediate access to them. And with Pure1, organizations can [manage, monitor, and analyze their storage](#) from anywhere with any device, including mobile devices. Pure1 is consistent with Pure's mantra of effortless operations, and results in a [data platform](#) that is self-managing and requires no tuning.

In fact, for customers who prefer a command-line interface (CLI), there is usually no need to go beyond what is printed on the business card-sized Getting Started guide included with each FlashArray//X. The guide even includes all the necessary steps to establish a stretched [multi-site replication cluster](#).



FIGURE 1. Pure FlashArray Getting Started guide

² https://education.emc.com/content/emc/en-us/home/widgets_template4.external.html

Comparison

Unisphere for PowerMax is highly configurable with many features to leverage. While Dell EMC has added several automated tools and tasks to speed up provisioning and management, the overall solution requires significant manual administrator involvement – not only for initial setup, but for ongoing operational support. The basic difference between management of the two platforms in Pure’s opinion is that Unisphere give you extensive tools to perform management manually and Pure automates or simply does not require many of those manual interventions.

The [Pure software management solution](#) can be set up in four easy steps, and once it’s set up the predictive intelligence and built-in automation makes the solution self-managing and self-healing. In the event that there is an issue, the [Pure1 proactive support team](#) will typically know about it and fix it before it becomes a problem.

HIGH AVAILABILITY

Dell EMC PowerMax

Although PowerMax is a new system, “PowerMAX is the next generation of the VMAX systems”, “It is the same scale-out architecture [as VMAX]”.³ The VMAX in turn evolved from Symmetrix DMX systems, which evolved from the original Symmetrix architecture. All of these previous generations of EMC systems established a proven degree of high-availability suitable for mission-critical application environments. Some of the world's most demanding and critical application environments run in production on VMAX, and they almost always deploy multi-site replication with either SRDF or VPLEX to achieve the desired level of high-availability.

Dell EMC claims that PowerMax arrays are [architected for six-nines \(99.9999%\) availability](#). Either SRDF/Metro, or VPLEX/Metro are usually configured if the customer desires continuous operations in the event of an array failure. Both SRDF/Metro licenses and VPLEX appliances and licenses, along with the associated professional services requires a significant upfront financial commitment from the customer, along with the associated software maintenance necessary to sustain the license.

One of the promises of the [Dell EMC Future-Proof Loyalty Program](#)⁴ is “Never Worry Migrations”. PowerMax supports “Non-Disruptive Migration” (NDM) for migrations from VMAX to PowerMax. The NDM process flow is shown in Figure2 at right.

In the same documentation, Dell EMC warns that “non-disruptive does not always equal no impact. NDM is designed to migrate a storage group (SG) at a time – not the entire system”.

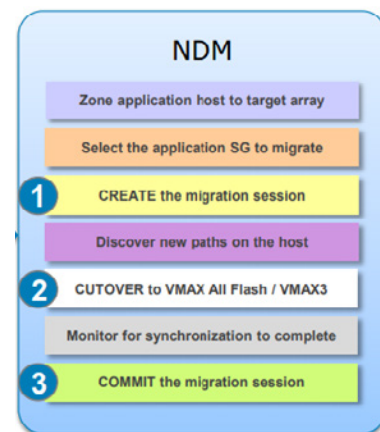


FIGURE 2. PowerMax Non-Disruptive Migration (NDM) process flow⁵

³ <http://tupodcast.blubrry.net/2018/09/11/ep08-dell-emc-powermax-nvme-all-flash-arrays-storage-class-memory-and-machine-learning-with-vince-westin/>

⁴ <https://www.dell.com/en-us/storage/future-proof-loyalty-program.htm>

⁵ Source: [Google](#)

Pure FlashArray//X

Pure FlashArray//X and its previous generations have established a proven 99.9999% uptime suitable for [mission-critical high-availability](#) with a single system, without the requirement of remote replication to achieve that high degree of uptime. With Pure's ActiveCluster, customers have the ability to extend that level of uptime even further.

The Purity//FA operating environment that powers the FlashArray//X includes ActiveCluster at no additional cost. ActiveCluster provides the ability to link two different data center sites up to 150 miles apart in an [active-active stretch cluster](#) with transparent failover, zero recovery point objective (RPO) and zero recovery time objective (RTO).

Pure's ActiveCluster solution includes Pure1® Cloud Mediator, which is a software-based third entity that monitors the link between the two sites, and declares which site becomes the primary site, should the link fail. [Pure1 Cloud Mediator](#) runs in the Cloud, so no extra software and no extra hardware, and its associated maintenance, is needed. It can be used to provide rack-level active clustering inside a data center as well as linking separate data centers. A remote third data center can also be added for asynchronous replication, which is accessible and live for replication from both of the primary arrays in the [ActiveCluster](#).

Comparison

Uptime and resiliency is a critical factor for mission-critical Tier 1 production applications. Both the VMAX/PowerMax and FlashArray//X have proven that they can deliver 99.9999% uptime. However, the configuration needed for the VMAX/PowerMax will typically be significantly larger and more complex, and has an associated cost in management, tuning, and operations.

ARCHITECTURE

Dell EMC PowerMax

PowerMax engines or "Bricks" each have two redundant active-active controllers. The PowerMax 2000 supports up to two Bricks per system, and the PowerMax 8000 supports up to eight Bricks. Bricks own their own Disk Array Enclosures (DAEs) which house the NAND flash drives.

The PowerMax 2000 can fit up to two Bricks in half a standard 19" rack, while the PowerMax 8000 can fit up to four Bricks in a single rack, and up to eight Bricks across two full racks taking up two full floor tiles.

From whatever starting point a PowerMax system is initially configured, it can be grown to its maximum configuration by adding capacity and DAEs to existing Bricks, or in a "scale-out" fashion by keeping all of the existing Bricks and DAEs, and incrementally adding additional Bricks and DAEs, thus potentially increasing the footprint, as well as needing additional requirements for power and cooling.

Historically, any upgrade between generations of VMAX systems required a full "forklift"-style upgrade and a complete data migration from the old system to the new system, as no two generations of engines/Bricks were able to coexist on

the same system. Pure believes that the same situation has been carried forward to PowerMax, and we are unaware of any claims by Dell to the contrary.

Although the PowerMax 2000 is based upon the same architecture as the PowerMax 8000, it is not capable of delivering the same “scale-out” features as the PowerMax 8000, since it is limited to a system maximum of only two Bricks. In a single-brick configuration the PowerMax 2000 is simply a conventional dual-controller system, and in a two-Brick configuration it cannot “scale out” any further than its starting point, other than to possibly add capacity or cache.

Pure FlashArray//X

Pure FlashArray//X uses a scale-up storage architecture which allows for a simpler and more flexible upgrade path. [Pure's Evergreen™ subscription to innovation](#) makes it easy to scale from one model to the next (i.e. //M series → //X family) while keeping capacity in place, upgrading controllers non-disruptively, and even while in full-performance production mode.

Due to Pure's [scale-up architecture](#), customers do not require forklift upgrades because Pure's array controllers can be dynamically upgraded. Pure has a proven track record of not only enabling [non-disruptive, data-in-place upgrades](#) without performance impact or needing a data migration among systems of the same generation, but even between systems of different generations. For example, customers have upgraded from our FA-400 series, which leveraged two distinct controllers, to our current generation of integrated controller design – never having to move data, migrate data, or suffer any loss of accessibility or performance.

Comparison

Although the PowerMax 8000 can possibly support up to 25% more effective capacity than FlashArray//X, it does so in massively more physical space, with significantly more power required, and with significantly more cooling required. PowerMax 8000 requires on average over 90% more system components to achieve roughly equivalent capacity deployments.

CRITERIA	PURE FLASHARRAY//X90	DELL POWERMAX 8000	FLASHARRAY VS. POWERMAX
MAXIMUM WEIGHT	185 LBS.	3,195 LBS.	94% less weight for Pure
NUMBER OF RACK UNITS	6	80	92% less rack space for Pure
MAXIMUM CAPACITY PBe	3PBe	4PBe	25% more PBe for PowerMax
MAXIMUM POWER REQUIRED	2,176 kVA	16,315 kVA	87% less power for Pure
MAXIMUM CONTROLLERS	2	16 (8 BRICKS)	Much simpler configuration for Pure
HEAT DISSIPATION	7,435 BTU/HR	55,667 BTU/HR	87% less heat generated for Pure requires much less cooling

TABLE 1. System component comparison

While the architectures of the PowerMax and FlashArray//X differ significantly, they have both demonstrated their ability to [host mission-critical Tier 1 applications](#) that drive the continued success of many organizations, including life-critical healthcare systems, government services, and large financial enterprises.

The distinction in approach to the architecture is reflective of the origins of each of the products – with the origins of the PowerMax focused on absorbing the latencies of magnetic media with the use of a large cache, as opposed to the FlashArray//X that was engineered for, and with, NAND flash media in mind.

As such, whether a system is “scale-out” or “scale-up”, does not impact its ability to deliver the highest level of availability, performance, or enterprise critical features.

TOTAL COST OF OWNERSHIP (TCO)

Dell EMC PowerMax

PowerMax is still based on legacy EMC retrofit technology, as outlined in the High Availability section above, requiring disruptive and costly forklift upgrades. Upgrades usually entail a re-purchase of most (if not all) of the storage system, including controllers, flash media, and associated components. These forklift upgrades contribute the most to the TCO of a storage system – and avoiding them, by contrast, represents the single biggest TCO savings opportunity.

PowerMax is covered by the Dell Future-Proof Loyalty Program. The Future-Proof Loyalty Program for PowerMax consists of the following components: 3 Year Satisfaction Guarantee, 4:1 Storage Efficiency Guarantee, Never Worry Data Migrations, Hardware Investment Protection, All-Inclusive Software, Cloud Consumption, Cloud Enabled, and Clear Price.

Under Dell EMC’s Future-Proof Loyalty Program customers are required to purchase a three-year ProSupport agreement, which is a re-branded version of Dell EMC’s maintenance agreement. The details of the agreement are fairly ambiguous, and we would recommend customers and prospects ask Dell EMC directly for their unpublished definition of each offering, effective date, and the detailed terms and conditions associated with each.

Pure FlashArray//X

Pure FlashArray//X was engineered from the ground up to be Evergreen, enabling [non-disruptive upgrades](#) for many generations, with data in place, no migrations, and with no performance impact. The key to this Evergreen expandability and upgradability is [FlashArray’s modular, stateless architecture](#). By leveraging a chassis-based design and customizable modules, FlashArray enables both capacity and performance to be independently improved and upgraded over time while customer data remains in place, so there is no need for a time-consuming, risky, expensive data migration.

Pure’s Evergreen Gold Subscription support and maintenance offerings provide new controllers every three years, to bring customers to the latest generation of controllers, as well as trade-in credits, giving customers the ability to future-

proof their investment. This subscription to innovation is a commitment on the part of Pure to ensure customers never need to repurchase capacity or performance that they have already purchased.

Comparison

The bottom line is that, compared to Dell EMC's PowerMax, Pure's Evergreen architecture and [subscription-based upgradability significantly lower storage TCO](#) – typically by up to 55%.

While Dell EMC's Future-Proof Loyalty Program is designed to give customers peace of mind that their product offering is upgradable and expandable, their product upgrades have two notable requirements that are major concerns for most customers – the need to perform data migrations with each upgrade, and the need to upgrade their architecture holistically to remain operational.

The "Hardware Investment Protection" component of the Loyalty Program claims to offset a portion of the capital outlay of a complete storage refresh. However, the fact that the system was never designed to be upgraded incrementally as well as non-disruptively – combined with EMC's own history of complex and disruptive upgrade cycles – suggests that customers will still need to repurchase many of their components. For example, will the current PowerMax flash media work with whatever next-generation director, processor, and interface technologies Dell EMC offers in three to five years, or will those components need to be repurchased as well? History clearly suggests the need to repurchase.

Another cost that the Loyalty Program program doesn't fully offset is data migrations. As examined in the High Availability section of this white paper, data migrations can be risky, can degrade performance and cause downtime, are complicated to perform, and warrant often-expensive professional services engagements. The risk associated with data migrations can sometimes expose enterprises to data corruptions or, even worse, data loss.

The "Never Worry Migrations" component provides credits toward Dell migration services, but that's a far cry from delivering true data-in-place upgrades that require minimal planning, and impose no performance impact or downtime. These costs can be considerable, especially when they involve an organization's own resources for planning downtimes and upgrade schedules – especially involving [critical application workloads](#) – regardless of any services offered by a storage vendor.

As PowerMax systems expand or "scale-out" to increase performance, it adds the need to increase the number of components (directors, processors, and numerous I/O modules for NVMe Flash, PCIe, Fabric, and compression/deduplication), which adds complexity and the need for additional space, power, and cooling.

By contrast, Pure's approach to investment protection and TCO is a partnership with their customers. We believe in taking a customer-first approach by making [storage array technology](#) innovation (hardware and software) immediately available via (and included in) Evergreen subscriptions. And, by eliminating complexity using automation such as AI, self-driving, and on-demand technologies, Pure delivers even greater TCO savings in storage management and also avoids the need for costly professional services consulting implementations.

The modular FlashArray//X is the key component in Pure's Evergreen ownership model, which is specifically designed for data-in-place in- and across-generation upgrades. The goal of Pure's Evergreen model is to [eliminate forklift upgrades with infrastructure](#) that expands the life of their storage platform from a few years to possibly a decade or more. Controller upgrades are included in Evergreen Gold subscriptions and are included every three years; all array software upgrades inclusive of new and emerging technologies are included; and a trade-in credit is provided for existing controllers – such as when purchasing additional capacity that necessitates a controller upgrade, or anytime a customer wants upgraded controllers – before the three-year mark.

SUMMARY

Both the Pure FlashArray//X and the Dell EMC PowerMax have proven themselves to be [enterprise-class storage systems suitable](#) for mission-critical application environments. Both are capable of meeting and exceeding most customer requirements for performance, capacity, and availability. The key differences between the two systems fall into three major categories.

1. **Effortless:** Pure FlashArray//X will require less initial set-up and configuration, and less ongoing operational management, thus reducing operational costs, improving time-to-production for ongoing use and changes, and reducing the risk of issues associated with potential human error. The end result is accelerated time-to-value and time-to-innovation that allows enterprises to compete more effectively and efficiently, and allows IT staff to focus on value creation activities that generate revenue and profits, as opposed to administrative and management operations that are costly and negatively impact the bottom line.
2. **Efficient:** Pure FlashArray//X will require less physical storage, less physical space, less power, less cooling, provide less complexity, and fewer components that can fail. This translates to a better TCO and a higher availability for all workloads.
3. **Evergreen:** Pure FlashArray//X will result in more predictable real-world costs, easier technology upgrades with less risk and effort, and greater architectural compatibility. This gives enterprises an infrastructure platform from which they can grow organically, inexpensively, and take advantage of the latest technologies as they become available and relevant to the organization.

© 2019 Pure Storage, Inc. All rights reserved.

Pure Storage, FlashArray, FlashBlade, Pure1, and the Pure Storage Logo are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. Other company, product, or service names may be trademarks or service marks of their respective owners.

The Pure Storage products described in this documentation are distributed under a license agreement restricting the use, copying, distribution, and decompilation/reverse engineering of the products. The Pure Storage products described in this documentation may only be used in accordance with the terms of the license agreement. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any. Pure Storage may make improvements and/or changes in the Pure Storage products and/or the programs described in this documentation at any time without notice.

ps_wp12p_pure-vs-emc-architectural-comparison_ltr_01

SALES@PURESTORAGE.COM | 800-379-PURE | @PURESTORAGE