Pure Storage FlashArray

Overview

Pure Storage FlashArray systems are all solid-state systems, providing NVMeoF-RoCE (called DirectFlash Fabric), FC, and iSCSI block storage access and file access with SMB and NFS. All data is both deduplicated and compressed prior to being written. Pure Storage claims there is no performance penalty for deduplicating and compressing data in-line, which enables an effective capacity ranging between 73 TB and 3.3 PB for their //X products, up to 5.5 for //XL models and up to 8.9 PB for their //C systems.

With FlashArray //X, Pure Storage has updated their second generation //m systems that continue to utilize purpose-built controller hardware and a custom NVRAM module. The system has dual controllers with either SAS attached SSDs or DirectFlash modules internal to the controllers attached via PCIe and add-on device enclosures using NVMeoF-RoCE for attach. In 2019, Pure Storage introduced the FlashArray//C system as a high-capacity flash array utilizing QLC flash. In late 2021, Pure introduced //XL models, XL130 and XL170, with improved performance and capacity as a scaled-up alternative to //X systems.

The Futurum Group Comment: Pure Storage marketing calls the controllers "active/active" but they are really active/passive in that both controllers cannot be actively transferring data between devices and hosts. The "expanded" definition of active/active used is because the standby controller keeps the host bus adapters active and allows selections to pass to the other controllers.

The FlashArray //X is an update of the //m controller hardware and utilizes NVMe protocol and PCIe interfaces for custom designed flash modules called DirectFlash. DirectFlash can replace the use of standard SSDs in FlashArray //X, FlashArray //C, and previous //m systems are upgradable according to the guarantee made in the fall of 2016. Pure Storage claims a 4x performance improvement with the use of NVMe DirectFlash

Highlights

- All Flash storage
- Capacity-oriented //C model available
- NVMe DirectFlash custom modules
- DirectMemory SCM
- NVMeoF – RoCE, FC & iSCSI block storage
- Thin Provisioning, Snapshots, replication, encryption
- VMware VAAI & vCenter support
- HA - active / passive controllers
- Compression & deduplication
- Purpose built hardware with custom NVRAM module
- Always-On QoS and QoS on LUNs
- CloudSnap to AWS and Azure
- Continuous async replication -ActiveDR
- FlashArray File Services – NAS

modules vs SSDs. The //X70 and //X90 can utilize a further performance increase by using Pure Storage's DirectMemory cache.

Usage

The Pure Storage //X systems are designed as a general-purpose, high-performance storage system competing with other mid and high range storage systems. Target applications are those requiring high performance without a large capacity requirement.

The //C arrays are focused on high capacity storage tier 2 storage applications. Use cases for the FlashArray//C include policy-based tiering, disaster recovery, and snapshot consolidation. The embedded Purity OS adds advanced features, Fibre Channel and iSCSI host access, NVME-oF support via RoCE, and active / passive HA controller nodes. The systems are highly available, providing all hot-swap components and symmetric LUN access. This enables non-disruptive access during component failures and upgrades, including for host access paths or storage controllers.

VMware support enables users to perform storage allocation along with performance and capacity management from within the VMware console. Additionally, support for VMware VAAI provides enhanced performance with VMware based workloads, including thin reclamation, block zeroing and virtual machine cloning operations.
The addition of file services allow FlashArray to be utilized as a flexible unified block and file system.

- Characteristics
  - Performance – High performance with all flash Direct Flash Modules and NVMe. Read latency can be increased with the use of DirectFlash SCM. //XL170 claims 150 microsecond latency and up to 36 GB/s throughput.
  - Availability – With two controllers, the FlashArray has the high availability characteristic including the failover and failback required.
  - Replication for BC/DR – Both periodic asynchronous and synchronous remote replication is available. ActiveDR is continuous async replication.
  - Data reduction – inline compression and deduplication is always on with the system.
- Applications
  - The Pure Storage FlashArray //X is targeted at applications that use block I/O with high performance requirements. The FlashArray //C targets high-capacity needs and tier 2 applications. FlashArray//E is economics focused and is a competitor to disk based systems.
  - The features of snapshot and remote replication give customers high-end system capabilities. The performance and features allow applications to be deployed in a high availability environment.
- System environments

  - The system is targeted for SAN storage with NVMe over Fabric – RoCE, Fibre Channel, and iSCSI environments.
  - NAS functionality available with NFSv3 and SMB3.
- Deployment and Administration
  - The system is intended for installation and management by the IT generalist and includes wizards and quick start configurations to allow for simple deployments.
  - The system reports to Pure Storage web software for continuous analysis of log data and the ability to use a SaaS application for monitoring and reporting.

Key Capabilities

Architecture

Pure Storage FlashArrays are configured with stateless commodity controller designs that support multiple interfaces with pluggable modules for advanced interface connections including 32/16 Gb/s Fibre Channel and multiple Ethernet speeds. With Ethernet attachment for hosts, iSCSI and NVMeoF-RoCE are supported. NVMeoF-RoCE requires use of RDMA capable NICs on servers and DCB switches.

The FlashArray system is an enclosure, featuring two RAID active/passive controllers with NVRAM modules connected via PCIe/NVMe, host interface ports, DirectFlash modules, and redundant power and cooling. A DirectFlash expansion shelf was added in 2017, which connects over Ethernet using NVMeoF with RoCE. The //XL models, introduced in 2021, no longer require NVRAM modules, using instead NVRAM on each of the DirectFlash Modules. In this case, NVRAM is implemented using SLC Flash as the non-volatile buffer area for data to be replicated to the other controller.

This stateless controller design means Pure is separating the configuration and user data from the compute power, which can make swapping out controllers for later generations easier to do.

NVRAM

Incoming data is written to NVRAM, and mirrored to the partner controller in HA configurations. Next, a write acknowledgement is sent to the client. A unique hash is created from the data and then stored at the address of its computed hash. This two-tiered write caching helps to alleviate host visible write delays encountered from SSD devices when performing garbage collection.

Purity OS

The embedded "Purity OS" now at version 6 (FA6) has many advanced features including encryption, snapshots, and remote replication along with a vCenter web plugin and VMware VAAI support. Version 4.8 added a fairness algorithm called Always On QoS to eliminate the noisy neighbor problem and NPIV support for upgrades to redirect I/Os to ports. Version 5 added VVOL support, improved compression, active-active stretched cluster support, the ability

to run VMs and containers on the storage platform, and portable snapshots including CloudSnap. Version 6 introduced continuous asynchronous replication called ActiveDR and FlashArray File Services for SMB and NFS access.

Unified File Services

FlashArray was initially launched as a block storage only solution. In 2019, Pure acquired Compuverde with the intent of integrating file capabilities into FlashArray. In 2020 file services were added with support for NFS and SMB. In 2023, Pure announced what they call a truly unified block and file architecture. In its unified solution, FlashArray utilizes a single global storage pool for both block and file storage. With this architecture, the entire global storage pool of both blocks and files can take advantage of FlashArray features such as global deduplication and data-at-rest encryption, as well as unified management. Pure's unified approach also added functionality for VM-aware NFS stores for granular management and protection of VMs.

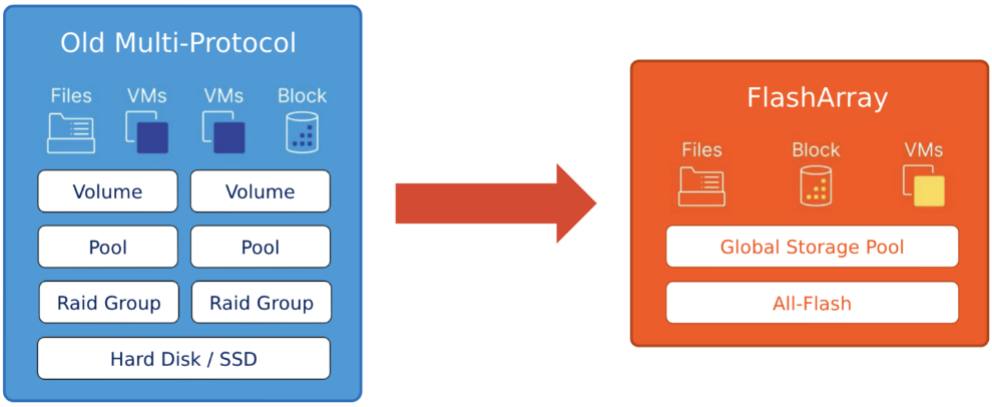

Figure 1 – Pure FlashArray Unified Block and File Storage (Source: Pure Storage)

FlashArray//XL

The FlashArray//XL models are high end 5U systems focused on both higher capacity and performance than //X systems. The feature support for both DirectFlash and a DirectMemory cache. The //XL systems DirectFlash includes distributed NVRAM. The controllers in //XL systems support 9 PCIe slots per controller.

Figure 2 - Pure Storage FlashArray //XL Chassis

Figure 3 - Pure Storage FlashArray //XL Bezel (Source: Pure Storage)

FlashArray//C

The FlashArray//C capacity-oriented models make a trade off of the latency achieved in the //X models, for a higher capacity. The FlashArray//C uses QLC (quad level cell) flash chips, capable of storing four bits on a cell, within its DirectFlash modules to achieve this high capacity. The FlashArray//C currently has three models, the FlashArray//C40, FlashArray//C60 and FlashArray//C90, which are available in configurations offering up to 494 TB, 1.9PB, or 2.3PB of raw storage respectively.

FlashArray//E

The FlashArray//E models are FlashArray systems that have been modified to optimize efficiency and economics. The E model starts with a 3U base chassis and grows 4PB raw with DFM expansion shelves that utilize 75 TB DFMs. FlashArray//E systems are targeted at replacing disk based systems and tier 2 workloads where efficiency and cost optimization are the priority.

DirectMemory

In September 2019, Pure Storage announced the ability to use an SCM cache, known as DirectMemory Cache. DirectMemory utilizes the high-performance technology of storage class memory to bridge the gap between NAND flash and DRAM for read caching. DirectMemory is available for the //X70 and //X90 models as well as //XL systems, allowing them improved read latency. The DirectMemory is available in a 3 TB option, consisting of four 750 GB DirectMemory Optane modules, and a 6 TB option, consisting of eight 750 GB DirectMemory Optane modules.

Capacity Optimization

Pure's inline deduplication, compression, and thin provisioning services are called FlashReduce. The deduplication engine can perform at variable block sizes down to 512B, performs pattern removal, and uses a global namespace for the FlashArray. FlashArray//X, C, and XL feature DirectCompress accelerators that provides deeper inline compression.

For continued capacity optimization, FlashArray offers Adaptive Data Reduction to ensure continuous background optimization, and constantly scours the array looking for ways to optimize the data layout. Anytime data is moved during this process it is put through the

deduplicate process again for further deduplication processing. When new media is added this process will rebalance the data through its process and therefore there is no need for a separate re-balancing feature.

Snapshots

Purity snapshots are reservation-less, and the number is based on memory constraints. With Purity 6.2.3, 100,000 snapshots are supported. Note that limitations for CloudSnap are different. Pure Storage claims FlashArray's snapshot creation capabilities have no impact on capacity or performance. The Purity OE uses Redirect-on-Write snapshots to update metadata associated with blocks of a volume at the time of snapshot creation and redirects all overwrites to new blocks.

Replication

Purity's FlashRecover Replication combines snapshots and policy-based automation for disaster recovery, and is an IP-based, asynchronous replication that enables multi-site long-distance data protection with low recovery point objective (RPO). FlashRecover Replication uses snapshots to create point-in-time replicas of multiple LUNs. These snapshots can be replicated simultaneously to multiple remote sites over long distances.

ActiveCluster

ActiveCluster is the name for the active-active stretched cluster implementation using synchronous replication. The stretched cluster supports a 150-mile distance or 5ms of delay. Pure1 Cloud Mediator which is an 'as a Service' offering can be used as a third party witness. For customers that have 'dark sites' where they do not allow the external web connection over the maintenance Ethernet, a local VM version of the Mediator can be installed. The ActiveCluster feature is available for no additional cost.

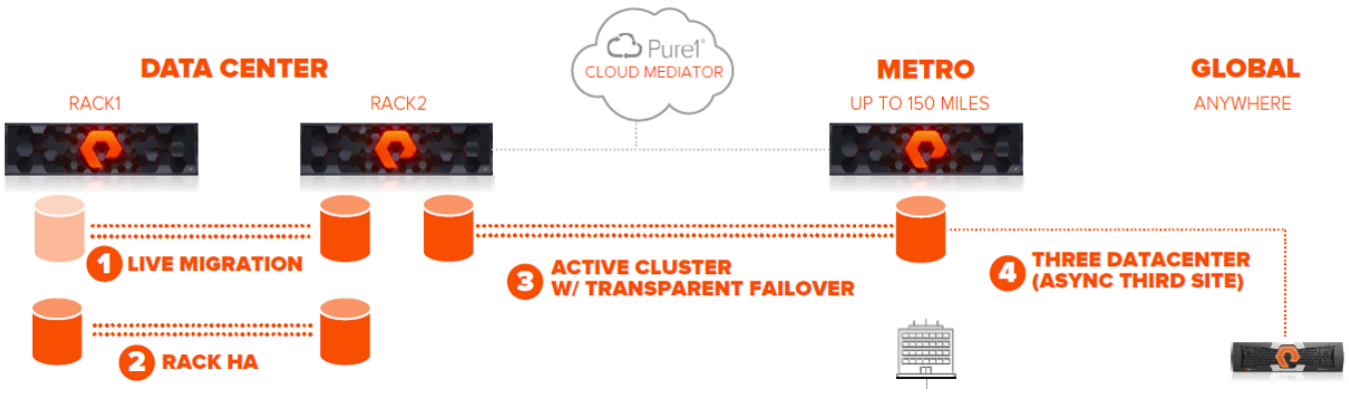

Figure 4: Stretched Clusters

PureFusion

Pure Fusion is a self-service provisioning and automated workload management tool built on the Pure1 cloud. Pure Fusion works across FlashArray//X, FlashArray//C, and Pure Cloud Block Store in AWS and Azure to provide a cloud-like experience with storage provisioning across availability zones. Users can provision storage through the Pure Fusion SaaS management plane by choosing storage classes and protection policies, and Pure Fusion will automatically place data on what it determines as the best device within the appropriate availability zone. Pure Fusion additionally continues to optimize workload placement by automatically rebalancing workloads as needed.

Significant Announcements

- March 2023 – FPGAs added to new //XL systems to increase compression. The FPGAs provide a deeper inline compression, with a second deeper compression done later.
- April 2023 – Pure announced unified file and block services from a single, global storage pool.
- June 2023 – Pure Accelerate Announcements
    - FlashArray//X and //C update
        - Sapphire rapids processors, PCIe Gen4, DDR5 Memory
        - DirectCompress compression accelerator
    - New //C90 model
    - 75TB FCM
    - New //E systems
    - Ransomware Recovery SLA
    - Ransomware anomaly detection
        - Monitors data reduction rate for possible anomalies

Futurum Group EvaluScale – Data Protection Software

The Futurum Group product review methodology "EvaluScale" assesses each product within a specific technology area. The evaluation of each product is based on its capabilities, with

capabilities for each technology segment grouped into distinct categories. For the NAS and SAN EvaluScales, products are evaluated based on the following 4 criteria categories:

- Performance / Capacity
- Basic Functionality
- Advanced Capabilities
- Ability to Execute

The full NAS EvaluScale can be found here.

The full Object Storage EvaluScale can be found here.

The Futurum Group Opinion and Outlook for Pure Storage FlashArray

As one of the early flash systems to market, Pure Storage has done well with systems that provide good performance with modest capacity, increased by software-implemented compression and deduplication. Pure Storage has also continued to develop the system with improved features and adding support for NVMe and PCIe attached devices, with a 4x performance improvement according to Pure Storage information. FlashArray supports the majority of the reliability, data protection, and advanced features to meet most enterprise needs.

The system is a traditional dual controller design with active/passive for data transfer between devices and hosts. Pure Storage claims active/active because the HBAs on the standby controller remain active for selection but do not transfer data. It is unfortunate that the claim of active/active is made because it gives competitors the opportunity to point out how misleading Pure Storage is and detracts from the real story of how valuable the products have been. Pure Storage has been very competitive with price and has performance to meet most customer demands. Pure has also been successful with its Evergreen program that allows for in place controller upgrades, further improving the systems economics.

The expansion of the FlashArray line to include the //C and //XL systems has widened the market for Pure, with //C adding a more capacity centric model, and //XL adding greater performance. The development of these new models, as well as continued development of DirectFlash Module and DirectMemory devices has been greatly beneficial to FlashArray.

The addition of file services was a worthwhile development and a good competitive move by Pure. By offering both block and file storage, FlashArray is able to expand its footprint and use as a NAS device. It also provides a more competitive offering against several other vendors that offer unified block and file storage. Pure has implemented its file services in a way that creates a simple, unified offering available from entry level to high for a wide range of applications.

A key area for development is FlashArray's scalability. FlashArray is not a scale-out system in any of the X, C, or XL variations, which means for increasing performance in parallel with capacity, a new controller pair (another model) must be installed. Other competitive systems have implemented scale-out functionality for a more measured growth. The Futurum Group believes that eventually Pure Storage will continue advancements to counter the competitors that have scale-out systems. Other developments such as stretched clusters and further integration of other systems software have made the Pure Storage system more widely applicable as a solution for primary storage. The transition to NVMe and PCIe attached devices will give Pure Storage a competitive lead in performance.