



PURE STORAGE CERTIFIED

PURE STORAGE ARCHITECT PROFESSIONAL **STUDY GUIDE**

EXAM NUMBER: FAP_001




TABLE OF CONTENTS

OBJECTIVE	3
AUDIENCE	3
FLASHARRAY ARCHITECTURE	4
STUDY RESOURCES.....	5
GUI.....	6
DATA REDUCTION	8
STUDY RESOURCES.....	9
PERFORMANCE FACTORS.....	9
ARRAY SIZING AND PHYSICAL SPECIFICATIONS	10
STUDY RESOURCES.....	10
ACTIVECLUSTER SYNCHRONOUS REPLICATION.....	10
STUDY RESOURCES.....	12
QoS.....	13
STUDY RESOURCES.....	13
PURITY.....	14
STUDY RESOURCES.....	14
FLASHARRAY CONFIGURATION.....	15
SUPPORTABILITY	17
STUDY RESOURCES.....	18
PURE1® META.....	18
STUDY RESOURCES.....	19
EVERGREEN™ MODEL.....	19
STUDY RESOURCES.....	20
SUPPORT	21



OBJECTIVE

This study guide is designed to help you prepare for the Pure Storage FlashArray Architect Professional Exam, exam number FAP_001.



AUDIENCE

This self-study guide is intended for those who wish to undertake self-study or review activities before taking the actual Pure Storage FlashArray Architect Professional exam. The guide is not intended as a substitute for training or hands-on time with Pure Storage products.

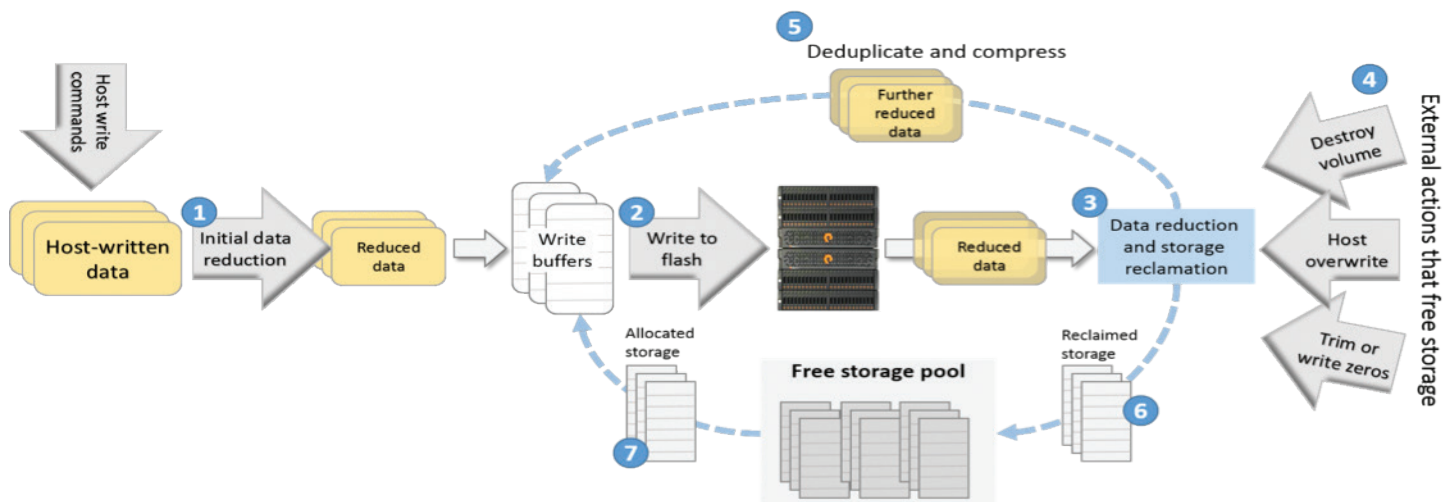
How to make the most of this study guide: The study guide summarizes the key topics on the Pure Storage FlashArray Architect Professional exam for you in an easy to use format. It is organized closely around the exam objectives but does not cover all potential questions from the exam and does not guarantee success on the actual exam. We suggest this guide be used in conjunction with our free online or in-person training covering FlashArray concepts and administration.

We hope you find this guide useful in your journey towards Pure Storage Certification, and we welcome your feedback by sending an email to certification@purestorage.com.

FLASHARRAY ARCHITECTURE

A FlashArray controller contains the processor and memory complex that runs the Purity//FA software, buffers incoming data, and interfaces to storage shelves, other controllers, and hosts. FlashArray controllers are stateless, meaning that all metadata related to the data stored in a FlashArray is contained in storage shelf storage. Therefore, it is possible to replace the controller of an array at any time with no data loss.

Data stored in a FlashArray undergoes continuous reorganization to improve physical storage utilization and reclaim storage occupied by data that has been superseded by host overwrite or deletion.



As an IO traverses the fabric into any port of the active/active Pure Storage controllers it is then buffered into DRAM and then staged to two NVRAM devices outside the controllers. Once completed, a write acknowledgement is sent back to the host to validate data is safe, as the data is now fully persisted. There is no need for any battery backup or standby power supply.

- **COMPLETELY STATELESS CONTROLLERS:** All configuration information and data in transit is stored on redundant NVRAM devices. That said, you can lose a controller, even lose both controllers, and you can replace them and be back up and running in minutes. Any state held in DRAM within the controller can be replayed from the data within NVRAM when the system comes back up. Simple, easy, and nothing for you to manage or worry about.
- **STAGING FOR EFFICIENCY TECHNOLOGY:** NVRAM allows us to complete dedupe, compression, pattern removal, and thin provisioning before the data is written to the FlashModules, extending the life the flash media.
- **BUFFERS WRITES AWAY FROM THE SSDS**
- **NO NEED FOR BATTERY BACKUP OR STANDBY POWER SUPPLY**

DirectFlash™ implements an elegant (but technically difficult) approach: instead of using SSDs, it takes raw NAND flash, wires it up with fast networking (an enhanced NVMe over PCIe in FlashBlade™’s case), and enables the flash to talk directly to our smart storage software.

The DirectFlash Module is a very simple piece of hardware, whose only job is to connect a large pool of flash over massively-parallel NVMe pipes to the FlashArray. From that point – all the magic’s in the software – the DirectFlash Software, which implements all the intelligence for flash management that used to live in the SSD globally across the entire flash pool.

The architectural advantage here is that the parallelism enables each core of each processor in our controllers to have a dedicated queue for each DirectFlash module. For comparison, in our traditional SSD-based flash modules we enable a queue depth of 8, so DirectFlash Modules provides a 32X improvement in parallelism. Beyond parallelism, IO to the DFM is deterministic – it’s bit addressable and there’s consistent access time to each flash block – eliminating the flash latency guessing game.

Furthermore, the DFM is 100% provisioned, meaning that Purity and the DirectFlash Software can “see” 100% of the flash in the system. A traditional consumer SSD has ~8% over-provisioning, and performance enterprise SAS and NVMe SSDs can have up to 50% over-provisioning, or flash that’s simply hidden from the system. By combining this 100% provisioning with the DirectFlash Software’s more efficient global flash management, DFM delivers between 14-36% more effective capacity from the same raw flash.

NVMe over Fabrics is the basis for our new DirectFlash Shelf. Exactly the same DirectFlash modules used in the FlashArray//X chassis are used in the DirectFlash shelf. The shelf connects to FlashArray//X via RoCE running over 50 Gb/sec Ethernet. Our new DirectFlash Shelf Controller translates between NVMe over Fabrics and NVMe (over PCIe). The translation is simple, because the commands and responses stay NVMe throughout. The shelf controller offloads data movement to RoCE hardware. Our software is highly efficient – built with an advanced architecture that is lockless, multicore, and runs in polling mode with no context switches.

RAID-HA - FlashArray’s dynamic multi-level scheme for protecting against data loss due to uncorrectable read errors and device failures. RAID-HA minimizes the impact of read error recovery, and automatically adjusts protection parameters based on the nature of stored data and conditions within an array.

Focus Areas

- Review components that comprise the FlashArray
- Understand the impact to metadata from stateless controllers
- Understand the background grooming process
- Understand how write IOs are processed
- Identify the role NVRAM plays in IO writes
- Understand the hardware and software design principles of DirectFlash
- Identify the protocols used by DirectFlash Shelf
- Review how FlashArray protects data’s availability

STUDY RESOURCES

- [Purity Reliability](#)

GUI

ANALYSIS

By default, Purity//FA displays the performance details for the entire array. To analyze the performance details of specific volumes, click the Volumes sub-tab along the top of the Performance page, select Volumes from the drop-down list, and select the volumes you want to analyze. To analyze the performance details of volumes within specific volume groups, click the Volumes sub-tab along the top of the Performance page, select Volume Groups from the drop-down list, and select the volume groups you want to analyze. You can analyze up to five volumes and volume groups at one time.

Click Clear All to clear the selections and display the performance details of all volumes again. The Analysis > Performance page includes Latency, IOPS, and Bandwidth charts. The point-in-time pop-ups in each of the performance charts display the following values:

LATENCY

The Latency chart displays the average latency times for various operations.

- SAN - Average time, measured in milliseconds, required to transfer data between the initiator and the array. SAN times are only displayed in graphs of one I/O type, such as Read or Write.
- Read Latency (R) - Average arrival-to-completion time, measured in milliseconds, for a read operation
- Write Latency (W) - Average arrival-to-completion time, measured in milliseconds, for a write operation
- Mirrored Write Latency (MW) - Average arrival-to-completion time, measured in milliseconds, for a write operation. Represents the sum of writes from hosts into the volume's pod and from remote arrays that synchronously replicate into the volume's pod
- Queue Depth - Average number of queued I/O requests for all volumes

IOPS

The IOPS (Input/output Operations Per Second) chart displays I/O requests processed per second by the array. This metric counts requests per second, regardless of how much or how little data is transferred in each.

- Read IOPS (R) - Number of read requests processed per second.
- Write IOPS (W) - Number of write requests processed per second.
- Mirrored Write IOPS (MW) - Number of write requests processed per second. Represents the sum of writes from hosts into the volume's pod and from remote arrays that synchronously replicate into the volume's pod.

BANDWIDTH

The Bandwidth chart displays the number of bytes transferred per second to and from all file systems. The data is counted in its expanded form rather than the reduced form stored in the array to truly reflect what is transferred over the storage network. Metadata bandwidth is not included in these numbers.

- Read Bandwidth (R) - Number of bytes read per second.
- Write Bandwidth (W) - Number of bytes written per second.
- Mirrored Write Bandwidth (MW) - Number of bytes written into the volume's pod per second. Represents the sum of writes from hosts into the volume's pod and from remote arrays that synchronously replicate into the volume's pod.

CAPACITY

The Array Capacity chart displays the amount of usable physical storage on the array and the amount of storage occupied by data and metadata. The data point fluctuations represent changes in physical storage consumed by a volume.

In the Array Capacity chart, the point-in-time pop-up displays the following metrics:

- Empty Space: Unused space available for allocation.
- System: Physical space occupied by internal array metadata.
- Shared Space: Physical space occupied by deduplicated data, meaning that the space is shared with other volumes and snapshots as a result of data deduplication.
- Snapshots: Physical space occupied by data unique to one or more snapshots.
- Volumes: Physical space occupied by volume data that is not shared between volumes, excluding array metadata and snapshots.
- Used: Physical storage space occupied by volume, snapshot, shared space, and system data.
- Usable Capacity: Total physical usable space on the array.
- Data Reduction: Ratio of mapped sectors within a volume versus the amount of physical space the data occupies after data compression and deduplication. The data reduction ratio does not include thin provisioning savings.

The Host Capacity chart displays the provisioned size of all selected volumes. In the Host Capacity chart, the point-in-time pop-up displays the following metrics:

- Provisioned: Total provisioned size of all volumes. Represents storage capacity reported to hosts. The Settings > System page displays and manages the general attributes of the FlashArray array.

Focus Areas

- Identify the details provided by the Analysis section of the GUI
- Identify the metrics tracked for Latency, IOPs, Bandwidth and Capacity

DATA REDUCTION

CAPACITY EFFICIENCY TECHNOLOGIES technologies either *avoid* or free the “unused” capacity within a storage volume so that it is available for use by other volumes, thereby increasing storage efficiency. Examples of capacity efficiency technologies include Thin Provisioning, Zero Detection, and Unmap. Results of Capacity efficiency technologies are NOT included in the Data Reduction number shown in the GUI.

DATA REDUCTION TECHNOLOGIES *reduce* the actual size of the data. For example, 10TBs of data may be reduced down to 2TBs with 5-to-1 data reduction. Examples of data reduction technologies include Deduplication, Compression, Pattern Removal, Deep Reduction (performed during Garbage Collection), and Copy Reduction (for clones and xCopy commands, not snapshots). Data reduction results in two important benefits: it allows you to address your usable capacity needs with a smaller amount of raw capacity, thereby lowering the \$/GB usable, and it maximizes the lifespan of flash by reducing the write IOs to flash. The goal for performance testing is to approximate your production workload as nearly as possible. Performance testing with synthetic workloads using non-reducible data may show up to 85% fewer IOPS than with reducible workloads. Our experience with customers is that almost all datasets and workloads are reducible by a Pure Storage FlashArray.

PATTERN REMOVAL

Pattern removal identifies and removes repetitive binary patterns, including zeroes. In addition to capacity savings, pattern removal reduces the volume of data to be processed by the dedupe scanner and compression engine. The software treats zero-filled sectors as if they had been trimmed—no space is allocated for them.

DEDUPLICATION

High-performance, inline deduplication operates on a 512-byte aligned, variable block size range from 4 - 32K. Purity//FA computes a hash value for each incoming sector and attempts to determine whether another sector with the same hash value is stored in the array. If so, the sector is read and compared with the incoming one to avoid the possibility of aliasing. Instead of storing the incoming sector redundantly, Purity//FA stores an additional reference to the single data representation. Purity//FA deduplicates data globally (across an entire array), so if an identical sector is stored in an array, it is a deduplication candidate, regardless of the volume(s) with which it is associated. Only unique blocks of data are saved on flash – removing even the duplicates that fixed-block architectures miss. Best of all, these savings are delivered without requiring any tuning.

COMPRESSION

Inline compression reduces data to use less capacity than the original format. Append-only write layout and variable addressing optimize compression savings by removing the wasted space that fixed-block architectures introduce. Combined with Deep Reduction, compression delivers 2 - 4x data reduction, and is the primary form of data reduction for databases.

DEEP REDUCTION

Purity Reduce doesn't stop at inline compression – additional, heavier-weight compression algorithms are applied post-process that increase the savings on data that was compressed inline. Most other all-flash products lack the use of multiple compression algorithms, and simply miss these savings.

COPY REDUCTION

Copying data on a FlashArray only involves metadata! Leveraging the data reduction engine, Purity provides instant pre-deduplicated copies of data for snapshots, clones, replication, and xCopy commands.

Focus Areas

- Understand how Data Reduction technologies are used in FlashArray
- Identify what capacity efficiency technologies are included in data reduction tracking
- Identify the components of Data Reduction used in FlashArray

STUDY RESOURCES

- [Purity Reduce](#)

PERFORMANCE FACTORS

In the latency graph, you'll find **Queue Depth**. Queue depth is defined as the number of queued SCSI commands pending processing. Another, less strict definition, would be the number of pending IO requests. Queue depth is managed at a number of points along a SAN path, the initiator OS, the HBA, the SAN switches, and the array. By itself, queue depth isn't terribly helpful, it's about context. But, do know that high queue depth can simply mean that things are busy. When queue depth roughly exceeds 100, the Pure array will tally higher latencies as it queues more commands.

The IOPS graph reports **Avg IO Size**. Keep in mind that this is "Average" IO size. If you have 3 volumes, average might mean something. If you have 100 or more volumes, it's more helpful to view it as a moving average. In general, it takes longer to write a large IO than a small IO. 8KB IOs should be fast to write, 1MB IOs will obviously need more time to be written out. For example, by default SQL Server will use jumbo sized IOs (typically of 1MB) to perform backup and restore operations. These can adversely affect the IO latency for other operations on your array, unless leveraging QoS.

Bandwidth is what it's all about. Bandwidth is one way of measuring how long it takes to do a job. If you have a job that requires 60 minutes to complete on a disk-based array, when you upgrade to Flash, this job time will plummet. Less time means you have more bandwidth. Bandwidth is the sanity checker. If your users are complaining about latency in a VDI environment, check the bandwidth. At times of high load, bandwidth should be very impressive. At times higher latencies (and queue depth) might follow. Bandwidth should always be viewed as the general status of the array, as high bandwidth could mean higher numbers in all other metrics.

Focus Areas

- Understand the impact of Queue Depth, Avg IO Size and Bandwidth has to performance

ARRAY SIZING AND PHYSICAL SPECIFICATIONS

	CAPACITY	CONNECTIVITY	PHYSICAL
//M10	Up to 30 TB / 28 TiB effective capacity** 5 – 10 TB / 4.7 – 9.3 TiB raw capacity	16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI 1 Gb/s Management & Replication ports	3U 575 – 625 Watts (nominal – peak) 95 lbs (43.1 kg) 5.12" x 18.94" x 29.72" chassis
//M20	Up to 250+ TB / 230+ TiB effective capacity** 5 – 80 TB / 4.7 – 74.4 TiB raw capacity	16 Gb/s Fibre Channel 10/40 Gb/s Ethernet iSCSI 1/10 Gb/s Replication ports 1 Gb/s Management ports	3U – 5U 600 – 950 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded 5.12" x 18.94" x 29.72" chassis
//M50	Up to 500+ TB / 450+ TiB effective capacity** 20 – 176 TB / 18.6 – 162.8 TiB raw capacity	16 Gb/s Fibre Channel 10/40 Gb/s Ethernet iSCSI 1/10 Gb/s Replication ports 1 Gb/s Management ports	3U – 7U 650 – 1280 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded + 44 lbs per expansion shelf 5.12" x 18.94" x 29.72" chassis
//M70	Up to 1,500 TB / 1,360 TiB effective capacity** 42 – 512 TB / 39.6 – 474.6 TiB raw capacity	16 Gb/s Fibre Channel 10/40 Gb/s Ethernet iSCSI 1/10 Gb/s Replication ports 1 Gb/s Management ports	5U - 7U 1230 – 1760 Watts (nominal – peak) 97 lbs (44.0 kg) fully loaded + 44 lbs per expansion shelf 5.12" x 18.94" x 29.72" chassis
//X70	Up to 1,050 TB / 950 TiB effective capacity** † 22 – 366 TB / 19.2 – 332.7 TiB raw capacity†	16 Gb/s Fibre Channel 10/40 Gb/s Ethernet iSCSI 1/10 Gb/s Replication ports 1 Gb/s Management ports	3U 1050 – 1320 Watts (nominal – peak) 97 lbs (44 kg) fully loaded 5.12" x 18.94" x 29.72" chassis

* Stated //M specifications are applicable to //M R2 versions. //X70 specifications are preliminary until GA.

** Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1.

† 1PB scale is achieved with 18.3TB DirectFlash Modules, which will be available post-GA of //X.

STUDY RESOURCES

- [Capacity Sizing](#)
- [Performance Analysis](#)

ACTIVECLUSTER SYNCHRONOUS REPLICATION

We wanted administrators that manage a FlashArray to simply keep managing the array in the same way they were used to, regardless of whether or not the array is participating in an ActiveCluster. To that end all administrative tasks such as creating volumes, creating snapshots, creating clones, managing asynchronous replication, and managing snapshot offload to NFS and the cloud, is all done the same way regardless of whether or not you are managing volumes on a single array or managing volumes in ActiveCluster. We only introduced one new command, the purepod command, in our management model to enable the configuration of ActiveCluster.

CONFIGURING ACTIVECLUSTER can be done in 4 easy steps. 3 of which are the same commands any FlashArray administrator is already familiar with.

STEP 1: Connect the Two FlashArrays - Connecting two arrays for ActiveCluster is done in the same way that we have done for asynchronous replication. We simply introduced a new connection type “Sync Replication.” Jumbo frames (MTU size 9000) are supported for replication. If using jumbo frames, ensure the network devices between the arrays have been configured to support the larger MTU size. Connecting arrays to each other requires access to port 443 and 8117. Ports 443 and 8117 are used for connect and disconnect operations. Replication data transfer uses only port 8117.

STEP 2: Create and Stretch a Pod - A pod defines a set of objects that can synchronously exist on two arrays simultaneously as though there is only one stretched instance of each object. This allows the objects in the pod to be managed as one entity rather than two, but from either array. Pods can contain volumes, snapshots, clones, protection groups (for snapshot scheduling and asynchronous replication) and other configuration information such as which volumes are connected to which hosts, and performance statistics and security audit log information. The pod acts as a consistency group, ensuring that multiple volumes within the same pod remain write order consistent. Pods are required for ActiveCluster.

STEP 3: Create a volume – Simply prepend the pod name to the volume name and separate them with :: this means that any scripts or REST calls that manage objects on the FlashArray work in the same way, by simply using the full volume name of pod1::vol1. An administrator can also move existing volumes into a pod (which is 100% non-disruptive) and then stretch that pod between two FlashArrays. When configuring existing volumes for ActiveCluster, our asynchronous replication technology is used in the background to perform the initial baseline copy of data on between the two arrays. Our asynchronous replication engine is data reduction aware, meaning we maintain compression and data deduplication as we transfer data from arrayA to arrayB.

STEP 4: Connect Hosts - Hosts can be connected to and do reads and writes into the same volume on both arrays.

ActiveCluster uses the preferred array setting to ensure that hosts have the best possible performance by exposing optimized paths to hosts that prefer that array and non-optimized paths to hosts that do not prefer that array. Hosts will then distribute IOs across optimized paths according to the host’s path selection policy, either round robin (RR) or least queue depth (LQD). The host will not use the non-optimized paths for front-end IO to the array unless the optimized paths are not available.

The Pure Storage solution provides this automatically by utilizing an integrated cloud based mediator. The Pure1® Cloud Mediator provides two main functions:

Prevent a split brain condition from occurring where both arrays are independently allowing access to data without synchronization between arrays. Determine which array will continue to service IO to synchronously replicated volumes in the event of a replication link outage. The Pure1 Cloud Mediator has the following advantages over a typical voter or witness component:

SAAS OPERATIONAL BENEFITS

As with any SaaS solution the operational maintenance complexity is removed: nothing to install onsite, no hardware or software to maintain, nothing to configure and support for HA, no security patch updates, etc.

AUTOMATICALLY A 3RD SITE

The Pure1 Cloud Mediator is inherently in a separate failure domain from either of the two arrays.

AUTOMATIC CONFIGURATION

Arrays configured for ActiveCluster will automatically connect to and use the Pure1 Cloud Mediator.

NO MIS-CONFIGURATION

Automatic and default configuration helps ensure that the mediator is correctly configured.

NO HUMAN INTERVENTION

A significant number of issues in active/active synchronous replication solutions, particularly those related to accidental split brain, are related to human error. An automated non-human mediator ensures properly driven safe failover in the event of an array outage or replication link outage.

PASSIVE MEDIATION

Continuous access to the mediator is not required for normal operations. The arrays will monitor the availability of the mediator, however if the arrays lose connection to the mediator they will continue to synchronously replicate and serve data as long as the replication link is active. Failover mediation for ActiveCluster can also be provided using an on-premises mediator distributed as an OVF file and deployed as a VM. Failover behaviors are exactly the same as if using the Pure1 Cloud Mediator, the on-premises mediator simply replaces the role of the Pure1 Cloud Mediator during failover events.

HOW TRANSPARENT FAILOVER OCCURS

In the event that the arrays can no longer communicate with each other over the replication interconnect, both arrays will briefly pause I/O and reach out to the mediator to determine which array can stay active for each sync replicated pod. The first array to reach the mediator is allowed to keep its synchronously replicated pods online. The second array to reach the mediator must stop servicing I/O to its synchronously replicated volumes, in order to prevent split brain. The entire operation occurs within standard host I/O timeouts to ensure that applications experience no more than a pause and resume of I/O.

Focus Areas

- Understand the steps needed to enable and configure ActiveCluster
- Identify the components that comprise an ActiveCluster configuration
- Understand the purpose of the Pure1 Cloud Mediator

STUDY RESOURCES

- [ActiveCluster – Simple Stretch Clustering For All](#)
- [FlashRecover Replication Configuration and Best Practices Guide](#)
- [Purity Replication Requirements and Interoperability Matrix](#)
- [Synchronous Replication for FlashArrays Technical Report 170101](#)

QoS

ALWAYS-ON QoS

Purity constantly monitors incoming I/O and is able to determine when the array is 'too busy' to deliver sub-ms predictable performance and automatically detects which volumes are consuming more than their fair share of resource and then throttles that workload to avoid performance degradation on all other workloads.

PERFORMANCE CLASSES

Performance Classes are built on top of Always-On QoS and provides you the opportunity to specify a Performance Class (Gold/Silver/Bronze) to a single volume or a workload/tenant. This approach empowers you to consolidate workloads and allocates performance to applications when the system approaches 100% performance utilization. This implementation is ideal for consolidating applications and tagging some as mission-critical, permitting throttling of less-critical applications in the event of system performance contention.

PERFORMANCE LIMITS

Performance limits introduce throughput limits on a per volume or workload/tenant basis. This is a hard limit which will determine the max performance a given workload can consume before being throttled.

Limits can be used in conjunction with Performance Classes to enable interesting use cases where Service Providers can create a catalog to meet minimums, maximums and bursts. Performance classes will ensure that if the system reaches performance saturation (and there are enough less important workloads to throttle to satisfy Gold classes performance demand) minimums are met for higher classes workloads and when the system is below performance saturation allow all workloads to burst. For this use case limits can be configured to represent peak value for burst.

Focus Areas

- Identify the performance classes for QoS
- Understand how to set throughput limits

STUDY RESOURCES

- [QoS FAQ](#)

PURITY

PURITY SECURE

Purity for FlashArray delivers an effective technology platform for compliance regulations, including data-protection and privacy regulations such as GDPR, and has earned an NIAP / Common Criteria Certification. Data security is provided by always-on FIPS 140-2 validated AES-256 data-at-rest encryption, KMIP integration, and rapid data locking for smart card-based instant locking of the array.

PURITY RUN

Now you can run VMs, containers, or custom applications on FlashArray! Windows File Services is the first data service built on Purity Run, adding SMB support to FlashArray. Eco system partners and customers can also utilize Purity Run for custom apps.

RBAC

Role-based access control requires the use of external accounts in an external directory service. RBAC is achieved by configuring groups in the directory that correspond to the following permission groups (roles) on the array and then assigning users to those groups within the directory.

ARRAY ADMIN GROUP: Array Admin users have all the privileges of Storage Admin users, plus the ability to perform array-wide changes. In other words, Array Admin users can perform all FlashArray operations.

STORAGE ADMIN GROUP: Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering volumes, hosts, and host groups. Storage Admin users cannot perform operations that deal with global and system configurations.

READ ONLY GROUP: Read Only users have read-only privileges to run commands that convey the state of the array. Read Only users cannot alter the state of the array.

Focus Areas

- Identify the different security features in Purity
- Understand the use cases for Purity Run
- Identify the group types configurable using Role Based Access Control

STUDY RESOURCES

- [Purity Features](#)
- [Rest API](#)
- [Role Based Access Control](#)

FLASHARRAY CONFIGURATION

VOLUMES

FlashArrays eliminate drive-oriented concepts such as RAID groups and spare drives that are common with disk arrays. Purity//FA treats the entire storage capacity of all flash modules in an array as a single homogeneous pool from which it allocates storage only when hosts write data to volumes created by administrators. Creating a FlashArray volume, therefore, only requires a volume name, to be used in administrative operations and displays, and a provisioned size.

1. Select **Storage > Volumes**.
2. In the Volumes panel, click the menu icon and select **Create...** The Create Volume dialog box appears.
3. In the Container field, select the root location, pod, or volume group to where the volume will be created.
4. In the Name field, type the name of the new volume.
5. In the Provisioned Size field, specify the provisioned (virtual) size number and size unit. The volume size must be between one megabyte and four petabytes. The provisioned size is reported to hosts.
6. Click **Create**.

Creating a volume creates persistent data structures in the array but does not allocate any physical storage. Purity//FA allocates physical storage only when hosts write data. Volume creation is therefore nearly instantaneous. Volumes do not consume physical storage until data is actually written to them, so volume creation has no immediate effect on an array's physical storage consumption.

RESIZE an existing volume to change the virtual capacity of the volume as perceived by the hosts. The volume size changes are immediately visible to connected hosts. If you decrease (truncate) the volume size, Purity//FA automatically takes an undo snapshot of the volume. The undo snapshot enters a 24-hour eradication pending period, after which time the snapshot is destroyed. During the 24-hour pending period, the undo snapshot can be viewed, recovered, or permanently eradicated through the Destroyed Volumes folder. Increasing the size of a truncated volume will not restore any data that is lost when the volume was first truncated.

COPY a volume to create a new volume or overwrite an existing one. After you copy a volume, the source of the new or overwritten volume is set to the name of the originating volume.

DESTROY a volume if it is no longer needed. When you destroy a volume, Purity//FA automatically takes an undo snapshot of the volume. The undo snapshot enters a 24-hour eradication pending period. During the 24-hour pending period, the undo snapshot can be viewed, recovered, or permanently eradicated through the Destroyed Volumes folder. Eradicating a volume completely obliterates the data within the volume, allowing Purity//FA to reclaim the storage space occupied by the data. After the 24-hour pending period, the undo snapshot is completely eradicated and can no longer be recovered.

HOSTS

The host organizes the storage network addresses - the Fibre Channel worldwide names (WWNs) or iSCSI qualified names (IQNs) - that identify the host computer initiators. The host communicates with the array through the Fibre Channel or iSCSI ports. The array accepts and responds to commands received on any of its ports from any of the WWNs or IQNs associated with a host.

Purity//FA will not create a host if:

- The specified name is already associated with another host in the array.
- Any of the specified WWNs or IQNs is already associated with an existing host in the array.
- The creation of the host would exceed the limit of concurrent hosts, or the creation of the WWN or IQN would exceed the limit of concurrent initiators.

Purity//FA will not delete a host if:

- The host has private connections to one or more volumes.

Purity//FA will not associate a WWN or IQN with a host if:

- The creation of the WWN or IQN would exceed the maximum number of concurrent initiators.
- The specified WWN or IQN is already associated with another host on the array. Hosts are configured through the GUI (**Storage > Hosts**) and CLI (**purehost** command).

Host-volume connections are performed through the GUI (**Storage > Hosts** and **Storage > Volumes**) and CLI (**puregroup connect**, **purehost connect** and **purevol connect** commands).

The Connections page displays connectivity details between the Purity//FA hosts and the array ports. The Host Connections pane displays a list of hosts, the connectivity status of each host, and the number of initiator ports associated with each host. Connectivity statuses range from “None”, where the host does not have any paths to any target ports, to “Redundant”, where the host has the same number of paths from every initiator to every target port on both controllers.

Host connections and target ports are displayed through the GUI (select **Health > Connections**) and CLI (**pureport list**, **purehost list --all**, and **purevol list --all** commands).

Focus Areas

- Understand host and array fiber channel configuration requirements.
- Understand how the handshake process between the host and the array.
- Identify configuration conditions that can lead to the array aborting host creation.
- Identify conditions that would prevent Purity from deleting a host.
- Identify conditions that would prevent Purity from associating a WWN or IQN with a host.

- Identify how to conduct a host-volume connection via the GUI.
- Identify how to view host connection status.
- Identify what host connectivity status states exist.

SUPPORTABILITY

An alert is triggered when there is an unexpected change to the array or to one of the Purity//FA hardware or software components. Alerts are categorized by severity level as critical, warning, or informational. Alerts are displayed in the GUI and CLI. Alerts are also logged and transmitted to Pure Storage Support via the phone home facility. Furthermore, alerts can be sent as messages to designated email addresses and as Simple Network Management Protocol-based (SNMP) traps and informs to SNMP managers.

PHONE HOME FACILITY

The phone home facility provides a secure direct link between the array and the Pure Storage Support team. The link is used to transmit log contents and alert messages to the Pure Storage Support team. Optionally configure the proxy host for HTTPS communication.

The phone home facility is managed through the GUI (**Settings > System**) and CLI (**purearray** command).

The audit trail represents a chronological history of the Purity//FA GUI, Purity//FA CLI, or REST API operations that a user has performed to modify the configuration of the array. Audit trails are displayed through the GUI (**Settings > Users**) and the CLI (**puremessage** command).

REMOTE ASSIST

Remote assist sessions are controlled by the array administrator, who opens a secure channel between the array and Pure Storage Support, making it possible for a technician to log in to the array. The administrator can check session status and close the channel at any time. Sessions not administratively closed will automatically close after 48 hours.

Remote assist sessions are opened and closed through the GUI (**Settings > System**) and CLI (**purearray remoteassist** command).

Configure the network connection attributes, including the interface, netmask, and gateway IP addresses, and the MTU. Ethernet and bond interface IP addresses and netmasks are set explicitly, along with the corresponding netmasks. DHCP mode is not supported. The array requires a minimum of three IP addresses configured: one for each physical controller management port and one for the multi-homed VIP.

NOTE: Network interfaces and DNS settings are configured through the GUI (**Settings > Network**) and CLI (**purenetwork** command for network interfaces, and **puredns** for DNS settings). Ethernet interfaces support one of three services: replication, management, or iSCSI.

Focus Areas

- Understand the functionality and use of the phone home facility feature.

- Identify where to configure the phone home facility.
- Understand what details are captured in the audit trail capability.
- Understand the functionality and use of the remote assist feature.
- Identify how to initiate a remote assist session.
- Identify how to configure the remote assist feature.
- Identify configuration requirements for remote assist.

STUDY RESOURCES

- [FlashArray User Guide](#)

PURE1® META

GLOBAL SENSOR NETWORK

Pure1 Meta receives a rich set of telemetry data from each of the 1000s of Pure Storage arrays currently deployed. Sensors at multiple levels – from the array itself to external, connected devices – provide more than 1 trillion data points per day, which to date has created a data lake of more than 7+ PB.

ARTIFICIAL INTELLIGENCE ENGINE

Pure1 Meta includes machine learning artificial intelligence built on a massive collection of storage array performance data. Applying predictive analytics to the data lake makes possible both a white glove customer support experience and breakthrough capabilities like accurate performance forecasting.

WORKLOAD PLANNER

With visibility into the performance of more than 100,000 workloads, Pure1 Meta is able to generate “workload DNA” – workload profiles based on deep analytics of 1000s of performance characteristics. From this continuously-refined set of profiles, Meta is able to give customers critical information on the performance and capacity requirements of their own workloads, as well as recommendations on which of their workloads will perform well together on the same array.

REAL-TIME SCANNING

Pure1 Meta continuously scans rich telemetry data to protect Pure Storage arrays from known vulnerabilities and to alert customers to potential issues beyond the array that may have an effect on performance. More than that, Meta learns – and becomes even more effective – over time.

GLOBAL PREDICTIVE INTELLIGENCE

Pure1 Support develops “issue fingerprints” – a set of data points that are uniquely a predictor of an issue – and Pure1 Meta then uses these fingerprints to continuously scan our global installed base of arrays. When Meta finds a match, the customer is notified, and Pure1 Support automatically opens a ticket and remediates the situation, even before an issue has occurred.

STUDY RESOURCES

- [Introducing Pure1 Meta: Pure's AI Platform To Enable Self-Driving Storage](#)

EVERGREEN™ MODEL

Evergreen Subscriptions give you all-inclusive access to our software, today and into the future, without having to pay extra for your existing storage footprint. It's as simple as that. Our customers have already benefited tremendously as we've rolled out major enhancements over the years, and there are plenty more to come.

Tired of big expenses hitting your budget at inconvenient times? Or having to deal with increasingly higher costs to own aging storage? Evergreen Subscriptions give you predictably flat maintenance rates over time. They'll never go up, and they may even go down. And, we'll replace any problematic hardware or software with like or better for as long as you are under our subscription.

Evergreen Subscriptions give you a full complement of high end services to make sure your array runs fast and smoothly. Our Global Insight Engine in the Pure1® cloud powers our Predictive Analytics, which are optimized to find and fix potential issues – before they become real issues. Our experts are keeping tabs on your arrays at all times, helping you through upgrades, responding within a mere 15 minutes for any Sev 1 incident, and ready to notify you if we need your assistance. If you do call us, we'll be standing by with instant access to L2 support – no L1 hassles and repeated explanations that waste time and money. It's practically like a managed service included in your subscription.

Our standard Evergreen Gold subscription completes the experience with a range of programs to modernize both controllers and flash media – without re-buys. Evergreen Gold includes next-generation controllers with every three-year renewal. With Evergreen Gold's Upgrade Flex bundles, anytime you expand FlashArray with a qualifying capacity purchase you receive trade-in credit for your existing controllers towards our most modern controllers – even if they are next-generation. Either way, your controllers will stay modern. And for flash media, any capacity expansions not tied to an Upgrade Flex qualify for trade-in credit for a portion of your existing, less dense flash. The end result? No repurchases of TBs you already own even as your whole FlashArray modernizes. That's what we mean by Evergreen.



Evergreen™ Storage

		EVERGREEN GOLD SUBSCRIPTION	EVERGREEN SILVER SUBSCRIPTION
SOFTWARE SUBSCRIPTION	All-Inclusive Software	✓	✓
	Future Software Features	✓	✓
HW & SW MAINTENANCE SUBSCRIPTION	Flat & Fair Maintenance	✓	✓
	Evergreen Maintenance	✓	✓
WHITE GLOVE SUPPORT SUBSCRIPTION	Predictive Analytics & Services	✓	✓
	Instant L2 Support Access	✓	✓
	Sev 1 Response SLA	15 min	15 min
	Support Managed Upgrades	✓	✓
	On-Site Break/Fix SLAs	4HR & NBD	4HR & NBD
	Mobile/SaaS-Based Monitoring	✓	✓
MEDIA & CONTROLLERS SUBSCRIPTION	Evergreen Architecture	✓	✓
	Right Size & 30-Day Guarantees	✓	✓
	Free Every Three	✓	X
	Upgrade Flex Controllers**	Trade-in Credit	No Credit
	Capacity Consolidation**	Trade-in Credit /w 4X Purchase	No Credit

* See Evergreen Storage program terms ** Requires purchase of qualifying capacity

STUDY RESOURCES

- [Evergreen Efficiency](#)
- [Evergreen Program](#)

SUPPORT

For information on Pure Storage's certification programs, visit www.examslocal.com. To contact us, please send an email to certification@purestorage.com.



Pure Storage, Inc.
Twitter: [@purestorage](https://twitter.com/purestorage)
www.purestorage.com

650 Castro Street, Suite #260
Mountain View, CA 94041

T: 650-290-6088
F: 650-625-9667

Sales: sales@purestorage.com
Support: support@purestorage.com
Media: pr@purestorage.com

© 2018 Pure Storage, Inc. All rights reserved.
ps_sg_fasupport-pro-employee-exam2018_04