PassQueen
a.k.a dumpsqueen.com

Pure Storage FlashBlade Certified Architect Professional (FBAP_002) Exam

Pure Storage FBAP_002

Version Demo

Total Demo Questions: 10

Total Premium Questions: 105

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, FlashBlade Architecture	59
Topic 2, FlashBlade Installation and Configuration	14
Topic 3, FlashBlade Administration	3
Topic 4, FlashBlade File System Management	9
Topic 5, FlashBlade Object Store Management	10
Topic 6, FlashBlade Data Protection	4
Topic 7, FlashBlade Troubleshooting	6
Total	105

QUESTION NO: 1 - (SIMULATION)

An oil company moves geological data between their offices in Dubai and Houston for redundancy. The data is accessed via NFS mounts on a cluster of three 10GB connected Linux servers at each location.

A series of cron jobs are leveraged to replicate data from data storage pool A in Dubai to data storage B in Houston and from data storage A in Houston to data storage pool B in Dubai.

Data in storage pool A is read-write, updated by 50 cellular connected data collector running in the field, and deleted after 30 days.

Data in the storage pool B is read-only, queried by the firm's analytics team, and copied to a cloud-based archive repository every 30 days.

The total amount of data in pool A in Dubai is 1.5PB

The total amount of data in pool B in Houston is 500TB

The data is not expected to grow in the next 12 months and is compressed by the application.

Which configuration should the architect present to meet this customer needs?

A . A Multi-Chassis FlashBlade with 52TB blades for each site focusing on simplicity.

B. A Multi-Chassis FlashBlade with 17TB blades for each site focusing on performance.

C. A Multi-Chassis FlashBlade with 52TB blades for each site and third-party replication software. ***

D. A Multi-Chassis FlashBlade with 17TB blades for each site and third-party replication software.

ANSWER: See the explanation for the answer

Explanation:

Correct answer: C. A Multi-Chassis FlashBlade with 52TB blades for each site and third-party replication software.

The Dubai writable pool alone contains **1.5 PB**. A single FlashBlade chassis cannot provide this capacity, so a **multi-chassis** configuration is required.
52 TB blades provide the appropriate capacity density for PB-scale data. The data is already application-compressed, so capacity planning should not depend on additional storage data-reduction assumptions.
The workload uses NFS and requires data movement between independent sites in Dubai and Houston. The required scheduled, bidirectional replication workflow should therefore be implemented with **third-party replication software**.
The smaller 17 TB blades would require substantially more blades/chassis to meet the capacity requirement without a stated performance need that justifies that design.

Therefore, deploy `Multi-Chassis FlashBlade + 52 TB blades + third-party replication software` at each site.

QUESTION NO: 2

A Linux client is mounting an NFS export from FlashBlade for a workload that requires higher per-client throughput. The client has multiple network paths available to FlashBlade.

Which two actions can help the client use parallel NFS connections? (Choose two.)

A. Configure the NFS mount with the nconnect option

B. Ensure the client has sufficient network-interface bandwidth

C. Create additional Object Store access keys

D. Use a single 1GbE client NIC

E. Disable link aggregation on the upstream switches

ANSWER: A B

Explanation:

Configure the NFS mount with the nconnect option and **ensure the client has sufficient network-interface bandwidth** are correct. The `nconnect` mount option allows a Linux NFS client to establish multiple TCP connections for an NFS mount, increasing the opportunity to use available parallel network paths. The client network interfaces must also provide enough bandwidth to carry the desired throughput.

Adding Object Store access keys does not affect NFS transport behavior, and a single constrained client NIC remains a per-client bandwidth limit. The upstream network must also be designed with sufficient bandwidth and appropriate pathing to support the expected traffic. See the Linux nfs(5) manual page and the Pure Storage FlashBlade product page.

QUESTION NO: 3

A current FlashBlade customer has a chip design process that is not meeting their SLAs. Although the customer is using fast remove, the IOPS on their 7x17TB FlashBlade have plateaued around 700K IOPS.

What should the architect recommend adding?

A. Blades ***

B. VIPs

C. LAG

D. Subnets

ANSWER: A

Explanation:

Blades are the appropriate addition because FlashBlade is designed to scale performance horizontally as blades are added to the chassis. Each blade contributes compute, flash capacity, network connectivity, and processing resources to the shared scale-out architecture. A seven-blade configuration can reach a point at which its aggregate processing and I/O resources no longer provide enough performance headroom for a demanding chip-design workflow. Adding blades increases the system's available parallelism and aggregate IOPS capability, allowing more filesystem operations and client I/O to be serviced concurrently while retaining the existing FlashBlade architecture and operational model.

Fast remove addresses the efficiency of deleting large directory trees and files by accelerating removal behavior; it does not add hardware resources or raise the platform's sustainable IOPS ceiling. Where the observed constraint is an aggregate IOPS plateau and the workload SLA requires higher throughput, the scalable resource to expand is the FlashBlade blade count. The architect should size the expansion using the customer's workload profile, including metadata intensity, read/write mix, concurrency, and required performance headroom. Pure Storage describes FlashBlade as a scale-out platform that delivers performance by adding blades: Pure Storage FlashBlade. Additional product architecture and scaling information is available in the Pure Storage FlashBlade documentation.

QUESTION NO: 4

A customer is standing up a new 30x17TB FlashBlade system.

The customerâ€™s environment is as follows:

1.The upstream switches are located 50 feet away

2.The upstream switches have the required number of 100Gb/s interfaces.

3.The customer is conscious about data availability.

4.The customer requires the use of 4 LAGs

What is the minimum number of cables required?

A. 4 passive QSFP + to QSFP + cables

B. 8 passive QSFP + to QSFP + cables

C. 4 active QSFP + to QSFP + cables

D. 8 active QSFP + to QSFP + cables ***

ANSWER: D

Explanation:

8 active QSFP + to QSFP + cables * ** is correct. FlashBlade networking is designed with redundant uplinks for availability. A LAG uses two physical network connections, allowing the system to maintain connectivity when an individual cable, switch port, or network path is unavailable. Therefore, four required LAGs need two cables per LAG, for a minimum of eight physical cables.

The approximately 50-foot (15-meter) separation between FlashBlade and the upstream switches also drives the cable type selection. Active direct-attach copper cables are appropriate for this distance, whereas passive DAC cabling is intended for substantially shorter runs. The upstream switches already provide the needed 100 Gb/s interfaces, so the design can use the required high-speed uplinks without reducing link count or bandwidth. This combination of four dual-link LAGs and active cabling provides the requested availability while meeting the physical-distance requirement. Pure Storage installation guidance requires network connections to be planned according to the platform's supported port, cable, and distance specifications; see the Pure Storage FlashBlade Installation Guide and the FlashBlade networking overview.

QUESTION NO: 5

For auditing and security purposes, a customer needs to lock down the 'pureuser'. Each storage admin must log in with their own account.

Which command should the architect run first after all external preparations have been completed?

A. Pureadmin create

B. Puremastership setattr

C. pureds enable

D. Pureconfig list

ANSWER: C

Explanation:

`pureds enable <directory-service options>` is the appropriate first command because it enables and configures FlashBlade directory-service integration after the required external identity infrastructure has been prepared. This integration allows the array to use centrally managed directory identities, such as LDAP or Active Directory users and groups, for authentication and authorization. Individual administrators can then be granted the necessary array roles through their own authenticated identities, creating clear accountability for administrative actions and supporting audit requirements. Establishing the directory service first is essential because subsequent user and role configuration depends on FlashBlade being able to communicate with and query the configured directory service. The command is run with the connection and directory parameters appropriate to the customer environment, such as the directory-service URI, search base, bind credentials where required, and transport-security settings. Once directory authentication is operational and access has been validated, the customer can complete the process of restricting use of the shared local `pureuser` account in accordance with its security policy. Pure Storage FlashBlade documentation and CLI references are available through the Pure Storage Support documentation portal; see also the FlashBlade product documentation resources.

QUESTION NO: 6

A customer map application is built with a containerized micro-service infrastructure that runs on Windows and accesses storage using SMB.

Each container is stateless, but every end user interaction leaves a set of small text and image files. Each directory keeps all daily transaction files. A process exists to manage file life cycle and periodically delete a directory.

The customer notices that as the customer base grows, the background deletion process backlog also grows. Which FlashBlade technology should be used to resolve this issue?

A. Enabling concurrent data access for NFS and SMB.

B. Enabling concurrent data access for NFS and HTTP.

C. Enabling fast-remove for NFS and SMB. ***

D. Reducing resource utilization for FlashBlade garbage collection.

ANSWER: C

Explanation:

Enabling fast-remove for NFS and SMB. *** is the appropriate FlashBlade technology because the workload creates large numbers of small files within daily transaction directories and then deletes those directories as a unit. Traditional recursive deletion requires the client and file system to process individual directory entries and files, so deletion time and metadata work grow with the number of objects. As the application's customer base expands, this per-file processing can cause the lifecycle-management process to fall behind.

FlashBlade Fast Remove is designed for high-scale directory removal. When enabled for a supported file protocol, it allows a directory tree to be removed efficiently without requiring the client to wait for a conventional, synchronous traversal and deletion of every file. The namespace change is made available quickly, while the underlying cleanup is handled by FlashBlade. This directly fits a microservice application in which containers are stateless, transaction data is grouped by day, and entire daily directories are periodically expired. SMB support is particularly relevant because the Windows-based containers access the data through SMB. Review the FlashBlade documentation available through the Pure Storage Support Knowledge Base and the product overview at Pure Storage FlashBlade for current feature availability and protocol requirements.

QUESTION NO: 7

A customer has enabled S3 object versioning on a FlashBlade bucket. An application overwrites an existing object.

What happens to the prior version of the object?

A. It is retained as a noncurrent version

B. It is immediately deleted

C. It is converted to a filesystem snapshot

D. It is copied automatically to another bucket

ANSWER: A

Explanation:

The prior version is retained as a noncurrent version is correct. When versioning is enabled, an overwrite creates a new current version of the object instead of replacing the existing object permanently. The prior version remains available and can be retrieved by specifying its version ID.

This behavior supports recovery from unintended overwrites and enables historical access to object content. Retention of noncurrent versions should be included in capacity planning and governed with lifecycle policies where appropriate. See the Amazon S3 Versioning documentation and the Pure Storage FlashBlade product page.

QUESTION NO: 8

A customer is designing a FlashBlade environment for an AI training workflow. The data science team will use hundreds of GPU workers to read millions of image files concurrently.

Which three characteristics make FlashBlade a good fit for this workload? (Choose three.)

A. High aggregate read throughput

B. Large number of concurrent clients

C. High metadata operation rates

D. Sub-millisecond transactional write latency

E. Fibre Channel block-volume access

ANSWER: A B C

Explanation:

High aggregate read throughput, **a large number of concurrent clients**, and **high metadata operation rates** are the relevant characteristics. AI training workers can generate substantial parallel read demand while repeatedly opening, enumerating, and reading many small training files. FlashBlade is designed to provide scalable performance for parallel file and object workloads.

The workload is driven by aggregate bandwidth, concurrency, and namespace activity. It is not primarily a low-latency transactional database workload, and a requirement for Fibre Channel block volumes does not align with FlashBlade file and object services. See the Pure Storage FlashBlade product page and the Pure Storage AI solutions page.

QUESTION NO: 9

A customer is deploying a FlashBlade into a datacenter where the top of rack switches cannot be configured into a single LAG.

Which two actions are necessary to configure the FlashBlade to use both switches? (Choose two.)

A. Create multiple exports

B. Create multiple VIPs

C. Create multiple LAGs

D. Use round robin load balancing

E. Use a multi-chassis configuration.

ANSWER: C E

Explanation:

Create multiple LAGs and **Use a multi-chassis configuration** provide the required design for attaching FlashBlade to two top-of-rack switches when they cannot operate as one conventional single-switch LAG endpoint. The FlashBlade network configuration must have link-aggregation connectivity defined for the available switch paths, rather than treating all physical connections as a single ordinary LAG on one switch. Multiple LAG definitions allow the required aggregated paths to be represented across the two switching devices.

A multi-chassis configuration is the complementary switch-side design. It coordinates the two physical switches so they can provide resilient, active connectivity to the FlashBlade while maintaining the forwarding and loop-prevention behavior required for aggregated links. The exact multi-chassis technology and configuration syntax are switch-vendor specific, but the design goal is consistent: both switches participate in the FlashBlade data network and a failure of a link or one switch does not unnecessarily remove all storage-network connectivity. Pure Storage documentation describes FlashBlade network configuration requirements and supported link-aggregation designs; the physical switch configuration must also follow the

selected vendor's multi-chassis LAG guidance. See the Pure Storage FlashBlade network configuration documentation and Pure Storage technical documentation.

QUESTION NO: 10

A customer is reviewing a FlashBlade network design before installation.

Which two design practices help provide resilient data connectivity? (Choose two.)

A. Connect data paths to two upstream switches

B. Use LAGs with redundant member links

C. Use a single unmanaged switch

D. Use one physical uplink for all client traffic

E. Place all data interfaces only on the management network

ANSWER: A B

Explanation:

Connect data paths to two upstream switches and **use link aggregation groups with redundant member links** are correct. These practices provide multiple physical paths between FlashBlade and the client network, allowing connectivity to continue after an individual link or switch-path failure when the upstream network is configured appropriately.

The two upstream switches must support the selected multi-chassis design when a logical aggregation spans both switches. A single unmanaged switch and a single physical uplink introduce avoidable single points of failure. Management-network connectivity is important operationally, but it does not replace redundant data paths. See the Pure Storage FlashBlade Support documentation and the Pure Storage FlashBlade product page.