
Exam : FBAA_004

Title : Pure Storage FlashBlade
Architect Associate

Version : DEMO

1. A batch application reading data from FlashBlade is changed to issue fewer, larger reads. During its next run, the engineer observes lower read IOPS but higher read throughput. The application owner asks whether storage performance has deteriorated.

Which two additional measurements are most useful for assessing these observations? (Choose two.)

- A. The array's data-reduction ratio during each run
- B. The average size of completed read requests during each run
- C. The physical storage utilization at the end of each run
- D. The read I/O response times during each run
- E. The number of files retained after each run

Answer: B, D

Explanation:

- B: Request size helps explain why fewer operations can transfer more data per second. A lower IOPS value by itself does not establish a regression when the workload changes.
- D: Response times show whether individual reads are taking longer, adding a responsiveness measure to the observed transfer rate. Interpret this alongside the application's requirements; the measurements alone do not establish a root cause.
- A: Data reduction describes storage efficiency, not the size and completion time of the observed reads.
- C: Physical utilization may matter in a broader investigation, but it does not directly resolve the IOPS-versus-throughput comparison.
- E: File count does not measure the data transferred by each read or its response time.

Authoritative reference: Arista and Pure Storage — Deploying NVMe-oF for Enterprise Leaf-Spine Architecture, page 10, "Test Results"; verified 2026-09-21. Used for the general relationship among request size, IOPS, throughput, and latency—not as evidence of FlashBlade NVMe-oF support.

2. A team is evaluating FlashBlade//S for two applications. One requires NFS access to its working files. The other writes a separate data set through an S3 API. The team wants to keep both application interfaces and consolidate storage without introducing a protocol gateway.

Which proposed design best fits this requirement?

- A. Place both data sets in NFS file systems and translate the S3 application's requests through a gateway.
- B. Place both data sets in S3 buckets and modify the file application to use object requests.
- C. Use native NFS file services for one application and native S3 buckets for the other on FlashBlade//S.
- D. Deploy separate FlashBlade//S systems because file and object services require different storage platforms.

Answer: C

Explanation:

- C: FlashBlade//S supports native file and object access, allowing these separate workloads to share a platform.
- A: This introduces the gateway the customer wants to avoid.
- B: This requires the application-interface change the customer has ruled out.
- D: Separate arrays are not required merely because one workload uses NFS and the other S3.

This is a protocol-suitability decision, not a completed sizing recommendation. It does not imply that an NFS file is also exposed as an S3 object.

Authoritative reference: Everpure — FlashBlade//S Data Sheet, ID19249-02-en-04/26, page 1,

“Advanced hardware and software working together”; verified 2026-09-21. The linked product address returned the five-page official datasheet when checked.

3. During a FlashBlade procurement discussion, the finance team requires ownership of the storage hardware. It also wants capacity consumption to determine subscription charges, subject to an agreed reserve commitment, rather than purchasing the full capacity entitlement upfront.

Which Evergreen model best matches these requirements?

- A. Evergreen//One
- B. Evergreen//Flex
- C. Evergreen//Forever

Answer: B

Explanation:

- B: Evergreen//Flex combines customer hardware ownership with a consumption-based subscription. Its reserve commitment means this is not an unrestricted zero-minimum payment model.
- A: Evergreen//One provides storage as a service rather than the customer-owned hardware model requested here.
- C: Evergreen//Forever supports traditional ownership with ongoing modernization, but does not provide the requested Flex consumption structure.

Specific configuration eligibility, minimum commitments, and prices still require a commercial quote; none is inferred from this model-selection question.

Authoritative references:

- Pure Storage — Data Storage That Gives You Flexibility with Control, PS2374-02, September 2023, pages 3–4 and 8, portfolio comparison; verified 2026-09-21.
- Evergreen//Flex FAQs, “What is a flexible storage ownership subscription?” and “What is the reserve commit and on-demand usage?”; live document checked 2026-09-21.
- Pure Storage — Own and Pay as You Go with Evergreen//Flex, official transcript, 03:20–04:13, distinguishing provider-owned service from customer ownership; verified 2026-09-21.

4. What distinguishes FlashBlade//S DirectFlash media management from an architecture that relies on the internal controllers of individual commodity SSDs?

- A. Purity//FB coordinates flash-media management across the system.
- B. Each DirectFlash Module independently sets the system's media-management policy.
- C. Pure1 performs flash-media management remotely for every client write.
- D. Client file-system drivers coordinate media management across the modules.

Answer: A

Explanation:

- A: Purity//FB manages resources globally, providing system-level media coordination.
- B: This incorrectly assigns the global policy to isolated module-level management.
- C: Pure1 monitoring is not the local flash-media-management function.
- D: Client drivers do not take over Purity//FB's media-management role.

Authoritative reference: Everpure — FlashBlade//S Data Sheet, ID19249-02-en-04/26, page 2, “Unmatched performance efficiency” and “Take control of your unstructured data”; verified 2026-09-21.

5. A FlashBlade//S customer expects retained data to grow much faster than processing demand. An

initial budget assumes that every capacity increase must include a proportional increase in compute resources. The customer asks whether this is an architectural requirement.

Which response is most accurate?

- A. That assumption is required because blade count determines capacity and compute in a fixed ratio.
- B. That assumption can be avoided by using data compression as a substitute for capacity planning.
- C. That assumption is required for NFS workloads but can be avoided by using object storage.
- D. That assumption can be revised because the modular design separates capacity and performance growth.

Answer: D

Explanation:

- D: FlashBlade//S supports independent capacity and performance scaling; validate the actual expansion against configuration rules.
- A: A universal fixed ratio contradicts that modular design.
- B: Compression does not establish the resources future workloads need.
- C: The architectural distinction is not conditional on choosing S3 instead of NFS.

Authoritative reference: Everpure — FlashBlade//S Data Sheet, ID19249-02-en-04/26, page 2, “Scale compute and storage independently,” and page 4, “Storage that gets better over time”; verified 2026-09-21.