



# HP0-J66<sup>Q&As</sup>

HP Storage Migration

**Pass HP HP0-J66 Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/hp0-j66.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by HP Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





### QUESTION 1

A customer plans to migrate data from a third-party storage system to an HP 3PAR solution. The customer has leased data center space. Which HP 3PAR Storage solution must be integrated into an HP 3PAR rack?

- A. HP 3PAR StoreServ 10800
- B. HP 3PAR StoreServ 10400
- C. HP 3PAR StoreServ F200
- D. HP 3PAR StoreServ 7400

Correct Answer: A

---

### QUESTION 2

A company is planning a large client Virtualization environment featuring an HP BladeSystem with ProLiant servers, HP Virtual Connect Flex-10, and HP StoreVirtual 4800 Storage Systems. Which storage advantage does this architecture provide over competitors' client Virtualization solutions?

- A. Storage traffic that enters the Virtual Connect module is prioritized automatically.
- B. Storage traffic remains within the Virtual Connect domain for better Performance.
- C. Each SAN provides dedicated storage for each hypervisor.
- D. Each SAN is its own cluster, eliminating data segmentation across multiple clusters.

Correct Answer: A

HP Virtual Connect FlexFabric interconnects pair with HP FlexFabric 10Gb 2-port onboard or mezzanine based adapters and provide FlexNIC functionality with Converged Network Adapters (CNAs) capable of either iSCSI, FCoE or native Fibre Channel storage protocols. With Virtual Connect 4.00, this flexible bandwidth capability is extended further with minimum and maximum bandwidth settings configurable from within the Virtual Connect management interface. This allows you to prioritize certain types of traffic over other types if contention occurs as well as insuring that certain networks have guaranteed bandwidth.

---

### QUESTION 3

A customer needs to migrate data to a new database environment. They need to make a physical copy of their data that can be presented to the new environment with the highest level of availability. Which form of replication should the customer use?

- A. backups to disk
- B. snapshots
- C. clones



D. asynchronous replication

Correct Answer: C

Incorrect: Not B: A storage snapshot is a set of reference markers, or pointers, to data stored on a disk drive, on a tape, or in a storage area network (SAN). A snapshot is something like a detailed table of contents, but it is treated by the computer as a complete data backup. Snapshots streamline access to stored data and can speed up the process of data recovery. There are two main types of storage snapshot, called the copy-on-write (or low-capacity) snapshot and the split-mirror snapshot. Utilities are available that can automatically generate either type.

#### QUESTION 4

A company needs a storage solution that provides four host bus adapter (HBA) slots and 64 GB of cache. Which HP storage solution meets these requirements?

A. HP 3PAR StoreServ 7400

B. HP StoreOnce 6200

C. HP 3PAR StoreServ 7200

D. HP StoreVirtual 4000

Correct Answer: A

HP 3PAR StoreServ 7400 4-node base with two 2U enclosures, four controller nodes, 64 GB cache, eight 8Gb/sec FC ports, four adapter slots, and 48 SFF drive bays.

Incorrect:

Not C: 3PAR StoreServ 7200 only has 32 GB cache.

#### QUESTION 5

A storage administrator needs to perform a gradual migration of large backup volumes and jobs to a remote D2D system over time. Because this data is being replicated for the first time, the administrator needs to populate the target device with all the relevant hash codes.

Which method should the administrator use to achieve this goal?

A. using co-location

B. using removable media

C. seeding data over the WAN link

D. setting initializer to migration

Correct Answer: C

\*

prior to being able to replicate only unique data between source and target D2D, we must first ensure that each site has



the same hash codes or "bulk data" loaded on it ?this can be thought of as the reference data against which future backups are compared to see if the hash codes exist already on either source or target. The process of getting the same bulk data or reference data loaded on the D2D source and D2D target is known as "seeding".

\*

Seeding is generally is a one-time operation which must take place before steady-state, low bandwidth replication can commence. Seeding can take place in a number of ways:

/ Over the WAN link ?although this can take some time for large volumes of data / Using co-location where two devices are physically in the same location and can use a GbE replication link for seeding. After seeding is complete, one unit is

physically shipped to its permanent destination.

/ Using a form of removable media (physical tape or portable USB disks) to "ship data" between sites.

[Latest HP0-J66 Dumps](#)

[HP0-J66 VCE Dumps](#)

[HP0-J66 Practice Test](#)