

Storage Networking Foundations Exam

SNIA S10-110

Version Demo

Total Demo Questions: 10

Total Premium Questions: 108

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Storage Networking Concepts	9
Topic 2, Storage Technologies	23
Topic 3, Storage Networking Technologies	36
Topic 4, Storage Virtualization	9
Topic 5, Storage Management	2
Topic 6, Data Protection	20
Topic 7, Storage Security	9
Total	108

QUESTION NO: 1

What is the purpose of a host bus adapter (HBA)?

- A. to connect a switch port to front end storage
- B. to provide host backups to tape media
- C. to connect a host to a switch
- D. to use a host CPU to interface with storage controllers

ANSWER: C

Explanation:

The purpose of **to connect a host to a switch** is correct because a host bus adapter is the hardware interface installed in, or integrated with, a server that attaches the host to a storage network. In a Fibre Channel SAN, the HBA supplies the host-side Fibre Channel port and connects by cable to a Fibre Channel switch port. This attachment enables the server to communicate with storage resources presented through the SAN fabric. The HBA handles storage-network connectivity and protocol functions, such as Fibre Channel frame processing, so that the host can discover and access logical storage devices through its SAN connection. Depending on the design, an HBA may also be used for a direct host-to-storage attachment rather than a switched fabric; however, connecting the host to a switch accurately describes its common role in a Fibre Channel SAN. SNIA defines an HBA as an I/O adapter that connects a host system to a storage network, which directly supports this answer. See the [SNIA Online Dictionary definition of host bus adapter](#) and Cisco's overview of [Fibre Channel SAN switching](#).

QUESTION NO: 2

Which two statements about jumbo Ethernet frames are correct? (Choose two.)

- A. A jumbo frame is an Ethernet frame with a payload greater than 1500 bytes.
- B. Jumbo frames enable the transmission of iSCSI packets on the IP network.
- C. Jumbo frames reduce per-frame protocol overhead and host processing for large data transfers.

ANSWER: A C

Explanation:

A jumbo frame is an Ethernet frame with a payload greater than 1500 bytes. Standard Ethernet commonly uses a maximum IP packet payload of 1,500 bytes, whereas jumbo-frame configurations permit a larger payload (often configured near 9,000 bytes, subject to the adapter, switch, and network design). This larger payload capacity is the defining characteristic of jumbo Ethernet frames; the precise maximum is implementation-dependent rather than a single universal Ethernet value.

Jumbo frames reduce per-frame protocol overhead and host processing for large sequential transfers. When the same amount of application data is carried in fewer Ethernet frames, endpoints and network devices process fewer headers and perform fewer frame-handling operations. This can improve throughput efficiency and reduce CPU overhead in appropriately engineered storage traffic paths, including iSCSI environments. To obtain that benefit reliably, the selected maximum transmission unit must be supported and configured consistently across the complete path between communicating endpoints. SNIA's terminology resources provide storage-networking context, while Cisco's Ethernet documentation describes the operational impact of larger frames: [SNIA Online Dictionary](#) and [Cisco: Jumbo Frames](#).

QUESTION NO: 3

When presenting a single volume to multiple hosts, which type of service needs to be running on the hosts?

- A. clustering services
- B. DNS services

C. directory services

D. DHCP services

ANSWER: A

Explanation:

clustering services are required when the same volume is presented to multiple hosts because the hosts must coordinate ownership and access to the shared storage. A cluster service provides the control-plane functions that let member hosts recognize one another, monitor node health, manage which node owns a clustered workload, and perform controlled failover when a node or application becomes unavailable. For shared writable volumes, clustering software also supplies the coordination mechanisms needed to prevent simultaneous, uncoordinated metadata updates and resulting file-system corruption. These mechanisms can include distributed locking, cluster-aware file-system access, reservations, and fencing of failed or unhealthy nodes. The exact implementation depends on the operating system and cluster stack, but the fundamental requirement is that all hosts accessing the shared volume participate in a supported cluster configuration. SNIA describes storage networking concepts and shared-storage terminology in its [Online Dictionary](#). Microsoft's [Failover Clustering overview](#) likewise explains that clustered servers work together to provide high availability for applications and services, with shared storage as a common cluster resource.

QUESTION NO: 4

Which two statements about asynchronous replication are correct? (Choose two.)

- A. The host write can be acknowledged before the remote copy is updated.
- B. It can support replication over higher-latency links than synchronous replication.
- C. It guarantees zero data loss after a primary-site failure.
- D. Every write must be committed remotely before host acknowledgment.

ANSWER: A B

Explanation:

Asynchronous replication acknowledges a write to the host after it has been committed at the primary storage system, before the data is necessarily committed at the remote replica. Replication updates are then transmitted to the secondary system according to the replication policy and available connectivity.

Because the remote write is not in the host write-completion path, asynchronous replication can support greater distances and tolerate more network latency than synchronous replication. However, a failure at the primary site can result in loss of data that had not yet reached the remote copy. See the [SNIA Online Dictionary](#) for replication terminology.

QUESTION NO: 5

Which two statements regarding RAID 10 are correct? (Choose two.)

- A. It combines mirroring and striping.
- B. It requires a minimum of four disks.
- C. It uses distributed parity to protect data.
- D. It provides the usable capacity of all member disks.

ANSWER: A B

Explanation:

RAID 10 combines mirroring and striping. Data is first mirrored to provide redundant copies and then striped across mirrored pairs to improve I/O performance. Because it uses mirrors, RAID 10 can continue operating after the failure of one disk in each mirrored pair, provided both disks in the same pair do not fail.

RAID 10 requires at least four disks because it needs at least two mirrored pairs to implement both mirroring and striping. Approximately half of the raw capacity is usable for data because the remaining capacity is used for mirror copies. See the [SNIA RAID overview](#) for RAID terminology and protection concepts.

QUESTION NO: 6

What should a valid iSCSI qualified name include?

- A. the iSCSI string
- B. date code, in yyyy-mm format
- C. the WWN string
- D. initials of company name

ANSWER: B

Explanation:

A valid iSCSI Qualified Name (IQN) includes a date code in yyyy-mm format. The IQN naming authority format begins with the literal prefix `iqn.`, followed immediately by the year and month in which the naming authority acquired its domain name. For example, an initiator or target might use `iqn.2024-06.com.example:storage.target1`. The date is not a device creation date or an arbitrary identifier; it identifies the point at which the organization obtained the domain name used as the naming authority. This convention helps preserve global uniqueness even if ownership of a domain name changes later. After the date, the IQN contains the organization's reversed domain name and can include an optional, authority-defined suffix to uniquely identify a particular iSCSI node. The iSCSI specification defines this `iqn.yyyy-mm` construction as the required structure for names using the IQN format. See the iSCSI naming-format requirements in [RFC 7143, section 4.2](#) and SNIA's [storage networking overview](#).

QUESTION NO: 7

According to SNIA, what are two valid storage tiering types? (Choose two.)

- A. storage cache tiering
- B. host tiering
- C. application tiering
- D. cloud tiering

ANSWER: A D

Explanation:

Storage cache tiering and cloud tiering are valid storage-tiering approaches. Storage cache tiering places a high-performance cache layer, commonly flash-based, in front of a lower-performance capacity tier. Frequently accessed, or "hot," data can be served from that faster layer, while the authoritative data remains in the capacity tier. This applies the fundamental tiering principle of matching data-access requirements to the performance and cost characteristics of available storage media.

Cloud tiering extends the same principle across storage locations by moving or copying data between local storage and cloud storage tiers according to policy. For example, data that must remain immediately accessible may stay on a local performance tier, while older or less frequently used data can be placed in a cloud-based capacity or archive tier. Such policies can account for access frequency, age, retention requirements, and cost. SNIA describes storage tiering as

assigning data to storage resources with different performance, capacity, availability, and cost attributes; cache and cloud implementations are practical forms of that concept. See the [SNIA Online Dictionary](#) and SNIA's [storage networking educational resources](#).

QUESTION NO: 8

Which device would communicate to a DAS?

- A. NIC
- B. switch
- C. HBA
- D. DB9 serial cable

ANSWER: C

Explanation:

An HBA is the device that communicates with directly attached storage (DAS). A host bus adapter is an I/O adapter installed in, or integrated with, a server or workstation that provides the host-side connection to storage devices. Depending on the DAS implementation, it may use protocols and physical interfaces such as SAS, SCSI, Fibre Channel, or NVMe/PCIe. The HBA initiates and manages storage I/O commands between the operating system and the directly connected disks, disk enclosure, or storage subsystem. This direct host-to-storage relationship is the defining characteristic of DAS: the storage is attached to a specific host rather than accessed through a storage network. In practical designs, the HBA and its supported cable/interface type must be compatible with the DAS enclosure and installed drives. SNIA describes DAS as storage directly connected to a server or workstation, and its storage terminology identifies a host bus adapter as the interface that connects a host system to an I/O bus. See the [SNIA overview of Direct-Attached Storage](#) and the [SNIA Dictionary](#).

QUESTION NO: 9

Your customer has a Fibre Channel SAN environment and has properly zoned servers for storage with front end controller ports. When an administrator goes to initialize a provisioned volume, it is not seen on the server.

What would cause this problem?

- A. an unplugged cable
- B. a multi-path I/O failure
- C. a corrupted volume
- D. improper LUN masking

ANSWER: D

Explanation:

improper LUN masking is the cause because Fibre Channel zoning and LUN masking perform separate access-control functions. Zoning establishes which Fibre Channel initiators and storage target ports can communicate across the fabric. It does not, by itself, grant a particular server access to every logical volume presented by an array target port. After a volume is provisioned, the storage array must map or mask that volume to the appropriate host or host group, usually using the server's initiator WWPNs and an assigned LUN number. If the host definition, initiator association, host group membership, volume mapping, or LUN presentation is missing or configured incorrectly, the server can successfully reach the controller ports but will not discover the provisioned volume during a rescan or initialization attempt. Correcting the array-side host-to-volume mapping and ensuring the intended initiator WWPN is included in the masking configuration makes the LUN visible to the server. SNIA describes LUN masking as an access-control mechanism that restricts host access to logical units; its terminology is available through the [SNIA Online Dictionary](#). For additional Fibre Channel SAN access-control context, see the [Cisco Fibre Channel zoning documentation](#).

QUESTION NO: 10

What are two benefits of using multipathing software with shared block storage? (Choose two.)

- A. It provides an alternate I/O path after a path failure.
- B. It can distribute I/O across available paths.
- C. It replaces the need for storage-system snapshots.
- D. It grants host access to otherwise masked LUNs.

ANSWER: A B

Explanation:

Multipathing software manages multiple physical paths between a host and a shared storage device and presents them to the operating system as one logical device. If one path fails because of an HBA, cable, switch, or controller-port failure, the software can direct I/O through an available alternate path.

Multipathing can also use multiple active paths to distribute I/O, depending on the configured policy and storage-array support. It is not a replacement for backup or array-level access controls. See the [Microsoft Multipath I/O overview](#) for multipath failover and load-balancing concepts.