PassQueen
a.k.a dumpsqueen.com

# Storage Networking Management and Administration

## SNIA S10-210

**Version Demo**

**Total Demo Questions: 13**

**Total Premium Questions: 135**

**Buy Premium PDF**

https://passqueen.com

support@passqueen.com

# Topic Break Down

| Topic | No. of Questions |
|---|---|
| Topic 1, Storage Networking Technologies | 64 |
| Topic 2, Storage Networking Management | 18 |
| Topic 3, Storage Networking Administration | 25 |
| Topic 4, Storage Networking Troubleshooting | 28 |
| Total | 135 |

QUESTION NO: 1

Which two high availability elements can be achieved by remote storage replication? (Choose two.)

**A.** Recoverability

**B.** Fault tolerance

**C.** Robustness

**D.** Redundancy

**E.** Serviceability

**ANSWER: A D**

**Explanation:**

Remote storage replication supports **recoverability** by maintaining a copy of application data at a separate storage system or location. If the primary storage system, site, or its data becomes unavailable, the replicated copy can be used as the basis for recovery and for resuming application operations. The achievable recovery point and recovery time depend on whether replication is synchronous or asynchronous, as well as on the failover and operational procedures in place.

It also provides **redundancy**, because the data is deliberately stored in more than one independent location or storage system. This additional copy reduces dependence on a single storage instance and is a core architectural component of disaster-recovery designs. Replication therefore contributes to availability by pairing a redundant data copy with a defined recovery path. SNIA describes replication as copying data between storage systems for data protection and recovery purposes; its storage-dictionary terminology also distinguishes redundancy and recoverability as availability-related concepts. See the SNIA storage definitions and the SNIA What Is Storage? overview.

**QUESTION NO: 2**

You need to add an iSCSI host and storage array to your environment. Your manager wants at least two options that allow all iSCSI initiators to query a central management server to locate the address of all iSCSI resources (E. g., iSCSI initiators, targets, switches, and management stations).

Which two protocol methods would be used? (Choose two.)

**A.** Directory Name Service protocol (DNS)

**B.** Simple Network Management Protocol (SNMP)

**C.** Internet Storage Name Server protocol (iSNS)

**D.** Service Location Protocol (SLP)

**ANSWER: C D**

**Explanation:**

Internet Storage Name Server protocol (iSNS) is designed specifically for discovery and management of IP storage resources. An iSNS server maintains a registration database containing entities such as iSCSI initiators, iSCSI targets, storage nodes, and management services. Registered clients can query that central service to discover accessible resources and obtain the network addressing information needed to establish storage connectivity. iSNS also supports discovery domains, which allow visibility and access relationships to be managed across a larger IP SAN. The iSNS architecture and its registration, query, and discovery functions are defined in RFC 4171.

Service Location Protocol (SLP) is also a standards-based service-discovery method that can provide centralized or directory-agent-assisted discovery. In an iSCSI environment, service agents advertise available services and user agents query for services of interest; an SLP Directory Agent can collect registrations and answer client queries. This enables initiators to locate advertised iSCSI target services without requiring each resource address to be manually configured. SLP's service-registration and service-request model is specified in RFC 2608, while iSCSI's use of SLP-based discovery is described in RFC 4018.

## QUESTION NO: 3

A solution architect needs to know how your switches enforce zoning. The enforcement method and identification type being used in your current FC switches must provide the highest level of security.

Which zoning enforcement should you use?

**A.** Session-based without hardware support

**B.** Frame-based and hardware support

**C.** Software-based with hardware support

**D.** Frame-based and software enforced

## ANSWER: B

**Explanation:**

**Frame-based and hardware support** provides the strongest Fibre Channel zoning enforcement because the switch applies zoning checks to individual frames in switching hardware. This is commonly called hard zoning or hardware-enforced zoning. When a frame arrives, the switch validates the source and destination against the active zoning configuration before forwarding it. As a result, prohibited communication is blocked in the data path itself, rather than relying solely on a host, management function, or session-establishment behavior to honor the zone definition.

This approach is particularly appropriate when the architect's requirement is the highest level of security. Hardware frame filtering provides consistent enforcement at line rate and prevents unauthorized initiators and targets from exchanging Fibre Channel frames even if an endpoint attempts to bypass normal discovery or login expectations. In practice, strong FC designs pair this enforcement with carefully controlled zone membership, typically using persistent identifiers such as WWPNs where operational requirements favor them, and with change control for the active zoneset.

SNIA defines the relevant storage-networking terminology in its Online Dictionary. Cisco's FC zoning guidance also describes zoning as an access-control mechanism enforced by the SAN switch: Cisco NX-OS SAN Switching Configuration Guide.

## QUESTION NO: 4

You have twelve 500 GB drives and you need to provide two LUNs using RAID 6 protection for a customer.

What will be the size of each LUN?

**A.** 1.5 TB

**B.** 1.75 TB

**C.** 2 TB

**D.** 2.25 TB

## ANSWER: C

**Explanation:**

**2 TB** is the correct size for each LUN when the twelve 500 GB drives are allocated evenly to create two RAID 6-protected LUNs. Each LUN is backed by a six-drive RAID 6 set. RAID 6 reserves capacity equivalent to two drives for dual distributed parity, leaving the equivalent of four drives for user data. Therefore, the usable capacity per six-drive set is calculated as (6 − 2) × 500 GB = 2,000 GB, or 2 TB. Creating two such sets consumes all twelve drives and provides equal capacity and RAID 6 protection to both LUNs. RAID 6 is designed to tolerate failure of up to two member drives in a RAID group while maintaining data availability, which is why two drive-equivalents are excluded from usable capacity. SNIA defines RAID 6 as a RAID level that uses dual parity, and its storage terminology distinguishes raw capacity from the usable capacity available after protection overhead. See the SNIA definition of RAID 6 and the SNIA Online Dictionary.

## QUESTION NO: 5

Click the Exhibit button.

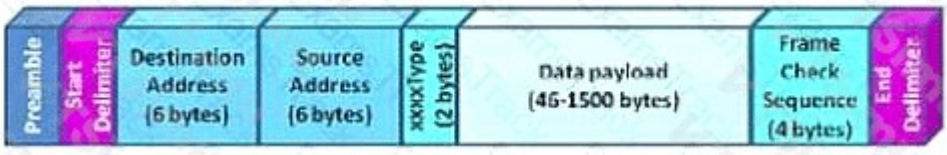


Which protocol is represented in the exhibit?

**A.** iSCSI

**B.** IPv6 Ethernet

**C.** IPv4 Ethernet

**D.** FCoE

## ANSWER: C

**Explanation:**

**IPv4 Ethernet** is correct because the exhibit represents an Ethernet frame carrying an Internet Protocol version 4 packet. In an Ethernet II encapsulation, the EtherType field identifies the protocol transported in the Ethernet payload. The value assigned to IPv4 is hexadecimal `0x0800`, which distinguishes an IPv4 payload at Layer 3 while Ethernet provides the Layer 2 framing, including destination and source MAC addresses. This is commonly described as IPv4 over Ethernet and is a fundamental encapsulation used for IP-based storage networking traffic, management connectivity, and conventional LAN communications.

RFC 894 specifies the standard encapsulation of IPv4 datagrams on Ethernet networks and identifies the Ethernet type value `0x0800` for IP. The IPv4 packet structure itself, including its version field and header format, is defined by RFC 791. Together, the Ethernet framing information and IPv4 protocol identification shown in the exhibit establish that the represented protocol is IPv4 Ethernet. See RFC 894: IP Datagrams over Ethernet Networks and RFC 791: Internet Protocol.

## QUESTION NO: 6

You are designing a security solution for a storage array which stores customer credit card information. Data on the storage array must be encrypted and if necessary, you want to be able to revoke access to data on demanD.

In addition to encryption, which solution would allow you to do this in a secure and automated fashion?

**A.** Use a manual record deletion solution.

**B.** Physically destroy hard drives holding the data

**C.** Store data indefinitely.

**D.** Use the key management solution.

## ANSWER: D

**Explanation:**

Use the key management solution is correct because encryption protects stored credit-card data only while the cryptographic keys remain available to authorized systems. A centralized key-management solution securely generates, stores, distributes, rotates, archives, and revokes the encryption keys used by the storage array. To revoke access on demand, an administrator can disable, delete, or otherwise make the relevant key unavailable according to the organization's approved key-destruction and retention policy. This makes the ciphertext on the array inaccessible, a capability commonly called cryptographic erasure or crypto-shredding. It is substantially faster and more automatable than attempting to locate and overwrite every physical copy of data across disks, snapshots, replicas, and decommissioned media. Key management also

supports separation of duties: the storage platform manages encrypted data while a dedicated key-management system controls authorization to use the keys. For cardholder data, this approach helps support PCI DSS requirements to protect stored account data with strong cryptography and to manage cryptographic keys securely throughout their lifecycle. NIST guidance similarly identifies key destruction as a key-management lifecycle function and notes that destroying keys can render encrypted data unrecoverable. See PCI DSS document library and NIST SP 800-57 Part 1, Revision 5.

## QUESTION NO: 7

A storage administrator is reviewing the backup process for critical application data. The backup jobs complete successfully, but the organization needs assurance that protected data can be recovered when required.

Which two actions should be performed? (Choose two.)

**A.** Perform periodic restore tests.

**B.** Review backup-job completion and error reports.

**C.** Increase the file system block size.

**D.** Disable backup encryption.

**E.** Delete older backup catalogs.

## ANSWER: A B

**Explanation:**

**Perform periodic restore tests** and **review backup-job completion and error reports** are correct. A restore test validates that backup data, catalogs, media, credentials, and recovery procedures can actually be used to recover the required files, volumes, or applications. A successful backup job alone does not prove that the backup is readable or that the recovery process will meet business requirements.

Reviewing job reports identifies failed jobs, incomplete backups, media errors, missed clients, and other conditions that can leave data unprotected. These operational checks should be correlated with retention requirements and documented recovery objectives. NIST contingency-planning guidance emphasizes testing backup and recovery capabilities in NIST SP 800-34 Rev. 1.

## QUESTION NO: 8

You have been hired to consult with storage administrators that are working with zoning and naming convention consistency issues. There are numerous types of zoning methods available using standards that also include aliasing. You want to use a method that guarantees unique global port identification.

What accomplishes this goal?

**A.** WWPN

**B.** domain_id:port_#

**C.** Alias

**D.** Zone set

## ANSWER: A

**Explanation:**

WWPN is the appropriate identifier because a World Wide Port Name is a globally unique, persistent identifier assigned to an individual Fibre Channel port. It is 64 bits in length and is normally assigned by the hardware manufacturer using an IEEE-registered organizational identifier combined with a vendor-controlled unique value. Because it identifies the port itself rather than its current connection location, a WWPN remains the reliable basis for consistent zoning and naming even when

a device is moved to a different switch port or a fabric is reconfigured. Storage administrators commonly use WWPN-based zoning to define access between host bus adapter ports and storage target ports in a way that is portable and unambiguous across the SAN. This directly meets the requirement for guaranteed global port identification and helps avoid dependence on fabric-specific addressing or locally administered labels. SNIA defines a World Wide Name as a unique 64-bit identifier used in Fibre Channel environments; a WWPN is the port-level form of that identifier. See the SNIA Online Dictionary and IBM's overview of World Wide Port Names.

## QUESTION NO: 9

What describes the FCoE transport protocol?

**A.** Fibre Channel frames encapsulated within Ethernet frames, transmitted on a Fibre Channel infrastructure

**B.** SCSI frames encapsulated within Ethernet frames, transmitted on an Ethernet infrastructure

**C.** Fibre Channel frames encapsulated within FCoE frames, transmitted on an Ethernet infrastructure

**D.** Fibre Channel frames encapsulated within Ethernet frames, transmitted on an Ethernet infrastructure

## ANSWER: D

**Explanation:**

Fibre Channel frames encapsulated within Ethernet frames, transmitted on an Ethernet infrastructure correctly describes Fibre Channel over Ethernet (FCoE). FCoE is a storage transport technology that preserves native Fibre Channel frame structure while carrying those frames across an Ethernet network. The Fibre Channel frame is placed into the payload of an Ethernet frame, allowing storage traffic to use a converged Ethernet infrastructure rather than requiring a separate physical Fibre Channel network for every connection. To support the lossless, low-latency characteristics traditionally expected by Fibre Channel workloads, FCoE commonly uses Ethernet enhancements collectively known as Data Center Bridging, such as priority-based flow control. FCoE does not translate storage commands into a different protocol; it transports Fibre Channel traffic over Ethernet. This distinction enables Fibre Channel hosts, switches, and storage systems to retain familiar Fibre Channel services and addressing concepts while using Ethernet links. SNIA's storage networking dictionary defines FCoE as the encapsulation of Fibre Channel frames over Ethernet, and the INCITS FC-BB-5 standard specifies the Fibre Channel over Ethernet mapping and encapsulation. See the SNIA FCoE overview and the FC-BB-5 specification.

## QUESTION NO: 10

A storage administrator is implementing role-based access controls for a storage management platform.

Which two practices support least-privilege administration? (Choose two.)

**A.** Assign only the permissions required for each administrator's job function.

**B.** Use separate individual administrator accounts.

**C.** Provide all administrators with a shared superuser account.

**D.** Grant storage administrators unrestricted network access from all subnets.

**E.** Use one permanent vendor account for all support incidents.

## ANSWER: A B

**Explanation:**

**Assign only the permissions required for each administrator's job function** and **use separate individual administrator accounts** are correct. Least privilege limits each administrator to the permissions necessary to perform assigned duties. For example, an operator responsible for monitoring capacity should not automatically receive authority to delete volumes, alter replication relationships, or modify security settings.

Individual accounts provide accountability by allowing audit records to associate administrative actions with a specific person. Shared administrator accounts reduce traceability and make it difficult to remove access for a single individual without affecting others. NIST access-control guidance addresses least privilege and account management in NIST SP 800-53 Rev. 5.

## QUESTION NO: 11

A storage administrator is investigating intermittent Fibre Channel link errors between an HBA and an FC switch.

Which two physical components should be inspected first? (Choose two.)

**A.** The optical transceivers

**B.** The fibre-optic cable

**C.** The storage array LUN masking configuration

**D.** The host file system permissions

**E.** The backup retention policy

## ANSWER: A B

**Explanation:**

**The optical transceivers** and **the fibre-optic cable** are correct. Fibre Channel link errors can be caused by a failing or incompatible SFP transceiver, contamination on optical connectors, excessive attenuation, a damaged cable, or an unsupported cable type. Inspecting and cleaning the optical path, confirming compatible components, and replacing suspect parts are appropriate first steps when errors are isolated to a particular FC link.

Zoning and LUN masking affect access control, but they do not create physical-layer transmission errors. Administrators should also review port error counters before and after corrective action to verify that the error condition has stopped. The Fibre Channel Industry Association provides Fibre Channel technology information at FCIA: Fibre Channel, and SNIA terminology is available in the SNIA Online Dictionary.

## QUESTION NO: 12

Click the Exhibit button.

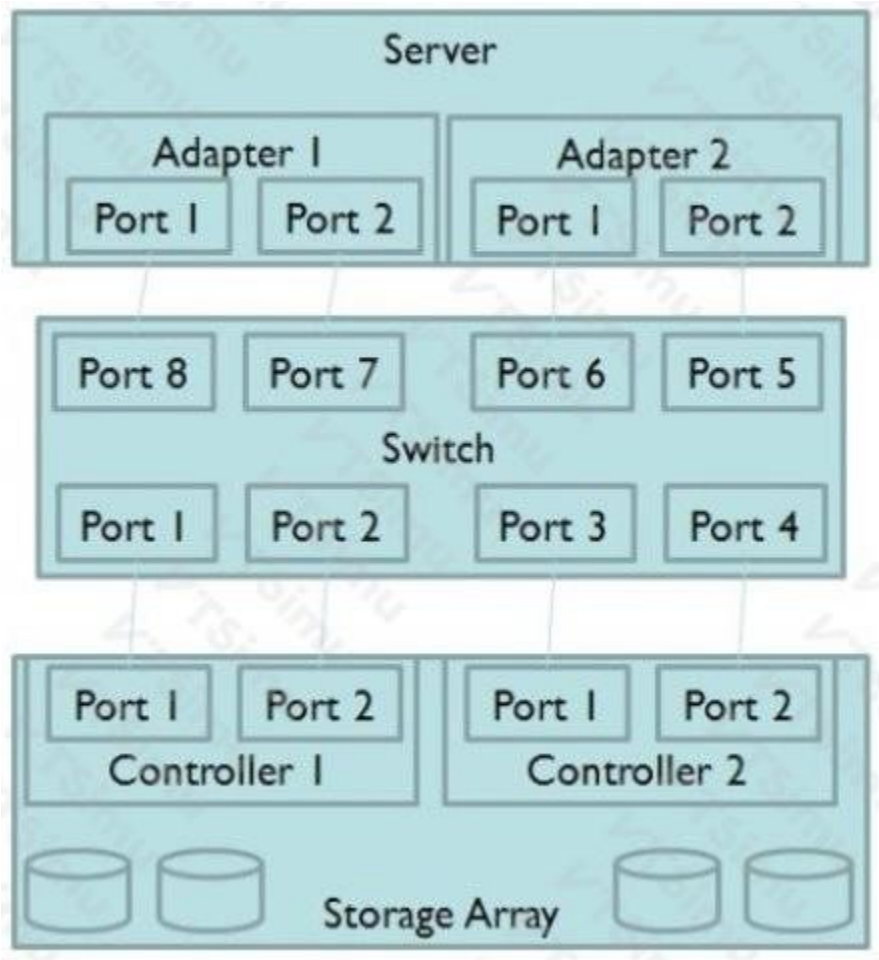


Referring to the exhibit, what are the two largest risks to a highly available solution? (Choose two.)

**A.** Only two adapters per controller

**B.** Only two adapters on the server

**C.** Only one server

**D.** Only one switch

**ANSWER: C D**

**Explanation:**

“Only one server” is a major availability risk because that server is a single point of failure for the hosted application or service. A hardware failure, operating-system fault, maintenance event, or loss of that server’s local connectivity can make the service unavailable even when storage remains online. High-availability designs normally use at least two application hosts and a clustering, failover, or load-balancing mechanism so a surviving host can continue providing the workload.

“Only one switch” is also a major risk because every path between the server and the storage environment depends on one network device. A switch hardware failure, software issue, power interruption, or required maintenance can therefore interrupt all storage access at once. Storage networks designed for availability commonly use independent redundant fabrics or switches, with hosts and storage/controllers connected across both fabrics. This removes the switch as a single point of failure and permits continued I/O through the remaining path when one fabric is unavailable.

These two conditions affect broad service availability rather than merely reducing component-level path redundancy. See Cisco’s guidance on designing resilient networks with redundancy at Cisco High Availability and NetApp’s SAN configuration guidance at ONTAP SAN configuration.

**QUESTION NO: 13**

Which standard would be used to increase interoperability between solutions of different cloud storage providers?

**A.** SES

**B.** CDMI

**C.** CWDM

**D.** DWDM

**ANSWER: B**

**Explanation:**

CDMI, the Cloud Data Management Interface, is the SNIA standard designed to improve interoperability between cloud storage providers and the applications that consume their services. It specifies a RESTful interface for creating, retrieving, updating, and deleting data stored in the cloud, along with standardized mechanisms for handling metadata, containers, access control, capabilities, and data-management requirements. By giving clients a common way to interact with cloud-storage services, CDMI reduces dependence on provider-specific APIs and helps support data portability between compatible implementations. This is particularly useful when an organization wants to manage data consistently across multiple providers or move workloads without redesigning all of its storage integration logic. SNIA describes CDMI as an international standard that enables cloud-storage clients to discover storage-service capabilities and manage cloud data through a standardized interface. The standard therefore directly addresses the interoperability objective stated in the question. See the SNIA CDMI standards page and the SNIA CDMI overview for its purpose, scope, and specifications.