PassQueen
a.k.a dumpsqueen.com

Data ONTAP 7-Mode Administrator

Network Appliance NS0-155

Version Demo

Total Demo Questions: 22

Total Premium Questions: 220

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Storage System Fundamentals	21
Topic 2, Data ONTAP Fundamentals	16
Topic 3, Storage Provisioning	14
Topic 4, Storage Efficiency	8
Topic 5, Data Protection	96
Topic 6, High Availability	15
Topic 7, NAS	19
Topic 8, SAN	30
Topic 9, Mix Questions	1
Total	220

QUESTION NO: 1

When using a Protection Manager policy to manage Open Systems SnapVault backups on a UNIX server, which three are valid objects to include in the data set? (Choose three.)

A. A file

B. A qtree

C. A directory

D. The entire client

ANSWER: A C D

Explanation:

For Open Systems SnapVault (OSSV), Protection Manager defines backup content from the UNIX host's filesystem namespace. A data set can therefore identify an individual file, a directory and its contained data, or the entire client. These selections support appropriately scoped backup policies: a single file can be protected when only a specific artifact is required, a directory can represent an application or filesystem subtree, and the entire client provides host-wide coverage for the configured OSSV client. The essential point is that OSSV is host based: the SnapVault client running on the UNIX system collects and transfers the selected host data to the SnapVault destination according to the Protection Manager policy. This lets administrators apply retention, scheduling, and vaulting behavior to data identified at the client level rather than requiring the source object to be a storage-system namespace object. NetApp documentation describes OSSV as software that enables data on open systems hosts to be backed up to NetApp storage, while SnapVault documentation provides the broader disk-to-disk backup and retention context. See NetApp SnapVault backup documentation and NetApp data-protection concepts.

QUESTION NO: 2

Which three configurations would increase data transfer rates in a volume SnapMirror environment for traditional volumes?

A. The destination volumes are flexible volumes.

B. The destination and source contain disks of the same size.

C. The destination and the source contain RAID groups of the same size.

D. The destination and source contain the same number of raid groups.

E. The destination and source storage systems are no father than 30 meters apart.

ANSWER: B C D

Explanation:

For SnapMirror replication between traditional volumes, transfer performance is improved when the destination is configured to parallel the source volume's underlying disk and RAID layout. "The destination and source contain disks of the same size" supports comparable disk characteristics and avoids introducing a substantially different disk-performance profile at the destination. "The destination and the source contain RAID groups of the same size" preserves similar RAID-group geometry, which helps the destination absorb replicated writes at a rate aligned with source read activity. "The destination and source contain the same number of raid groups" likewise maintains a comparable degree of disk parallelism. Together, matching disk size, RAID-group size, and RAID-group count reduces structural imbalance between the source and destination traditional volumes and allows SnapMirror to use the available storage resources more consistently during baseline and update transfers. This reflects the general NetApp SnapMirror design principle that replication throughput depends on both network capacity and the ability of the source and destination storage to read and write the replicated data efficiently. See NetApp's SnapMirror disaster-recovery overview and SnapMirror configuration and best-practices technical report.

QUESTION NO: 3

There are three phases of Non Disruptive Volume Movement (NDVM). What is the correct sequence of these phases?

A. Setup Phase, Mirror Phase, Cutover Phase

B. Initialization phase, Copy Phase, Migrate Phase

C. Begin Phase, Move Phase, Complete Phase

D. Setup Phase, Data Copy Phase, Cutover Phase

ANSWER: D

Explanation:

The correct sequence is **Setup Phase, Data Copy Phase, Cutover Phase**. In clustered Data ONTAP terminology, non-disruptive volume movement is performed by the `volume move` operation. The setup phase prepares the destination aggregate and establishes the relationships and resources required for the move. During the data copy phase, ONTAP transfers the volume's data to the destination while the source volume remains online and continues serving client I/O. ONTAP performs repeated updates so that the destination stays synchronized with changes made during copying. The cutover phase is the final, brief transition in which ONTAP redirects the volume from its source aggregate to its destination aggregate. This design allows the volume to be relocated without requiring clients to remount it or substantially interrupt normal data access. The three phase names therefore describe both the operational order and the mechanism that permits a non-disruptive migration: prepare the move, synchronize the data, then switch the active location. NetApp's ONTAP documentation describes volume movement as a nondisruptive operation and documents the monitoring and cutover behavior of volume moves. See NetApp ONTAP: Move a volume and NetApp ONTAP: Monitor a volume move operation.

QUESTION NO: 4

What are three true statements about physical reallocation of data added in ONTAP 7.3? (Choose three.)

A. Use reallocate -p pathname to run the physical reallocate.

B. Use aggr reallocate -p pathname to run the physical reallocate.

C. Reinitialize any SnapMirror relationship for the new physical geometry.

D. Volume Snapshots do not need to be deleted for the reallocate to run effectively.

E. Dont run it on files, volumes, and LUNs that are in an aggregate created by a version of ONTAP earlier than 7.2.

ANSWER: A D E

Explanation:

Physical reallocation, introduced with the `-p` switch, is run with `reallocate -p pathname`. This mode rearranges user data at the physical-block level in the aggregate while preserving the logical block locations seen by the flexible volume. Preserving those logical locations is particularly useful when snapshots exist: snapshots do not need to be deleted before the operation, and physical reallocation can reduce the additional space otherwise associated with reallocating a volume containing Snapshot copies. The same characteristic can also reduce the amount of changed data that a subsequent SnapMirror update must transfer when the source volume has been physically reallocated.

Physical/aggregate reallocation has an aggregate-version prerequisite. It is not supported for aggregates created by Data ONTAP releases earlier than 7.2. Therefore, files, volumes, and LUNs residing in such an aggregate should not be targeted for this operation. These behaviors and restrictions are documented in the Data ONTAP `reallocate` manual page, including the definition of the `-p` option and its Snapshot and SnapMirror considerations. See the NetApp reallocate manual page and the NetApp data protection documentation.

QUESTION NO: 5

Protection Manager has policies that allow for tape-based protection as well as disk-based protection.

A. True

B. False

ANSWER: A

Explanation:

True is correct. In the NetApp 7-Mode DataFabric Manager and Protection Manager architecture, protection policies are used to define and automate how dataset members are protected according to recovery and retention requirements. Protection Manager supports disk-based protection workflows, such as replication to secondary storage using NetApp Snapshot and replication technologies, and it also supports tape-oriented protection through archive policies and tape backup operations. This enables an administrator to use a single policy-driven framework for operational recovery from disk and longer-term retention on tape.

Policies are associated with datasets and specify the protection behavior that Protection Manager applies to the dataset's primary data. The policy model can therefore combine the operational advantages of disk protection, including faster recovery and replication management, with tape's role in archival retention and offsite data handling. This distinction is important in 7-Mode environments because Protection Manager is intended to coordinate protection services rather than limit a dataset to one media type. NetApp's ONTAP documentation also describes tape backup as a supported data-protection capability and explains the role of policy-based replication in protecting data. See NetApp ONTAP tape backup documentation and NetApp ONTAP data protection documentation.

QUESTION NO: 6

Which statement applies to the Open Systems SnapVault Free Space Estimator utility?

A. This utility runs automatically after each transfer.

B. This utility must be enabled on the secondary storage system.

C. This utility determines if there is sufficient disk space available on the primary to perform an OSSV transfer.

D. This utility determines if there is sufficient disk space available on the secondary to perform an OSSV transfer.

ANSWER: D

Explanation:

The correct statement is: "This utility determines if there is sufficient disk space available on the secondary to perform an OSSV transfer." In an Open Systems SnapVault (OSSV) relationship, the primary system contains the source data and the secondary storage system receives and retains the SnapVault backup data. Before a transfer, the Free Space Estimator evaluates the expected capacity requirement at the destination and helps establish whether the secondary has enough available disk space for the incoming transfer. This is important because a successful OSSV update depends not only on source-side data availability and connectivity, but also on the secondary system's ability to accommodate the data and Snapshot-related overhead created by the replication operation. Estimating destination free space in advance helps administrators plan capacity and avoid interrupted or unsuccessful backup transfers caused by insufficient space on the backup target. Although OSSV is a legacy Data ONTAP 7-Mode technology, the same destination-capacity principle remains central to NetApp replication planning. See NetApp's data protection documentation and the Snapshot capacity management documentation for related replication and space-planning concepts.

QUESTION NO: 7

Which set of commands should be used for troubleshooting Ethernet connectivity problems from the storage system console?

A. ifstat, netdiag

B. perfstat, tracert

C. sysconfig - cr, sysstat

D. dish shm_stats, sysstat

ANSWER: A

Explanation:

ifstat, netdiag is the appropriate command set for troubleshooting Ethernet connectivity directly from a NetApp 7-Mode storage-system console. The `ifstat` command displays device-level network-interface statistics, including packet counts, errors, collisions where applicable, and utilization-related information. These counters help an administrator identify symptoms of physical-link, duplex, cabling, or interface-level problems and determine whether errors are increasing while traffic is being sent or received.

The `netdiag` command performs network diagnostic operations from the storage system. It is intended to help validate network connectivity and diagnose communication behavior from the controller itself, which is important because the relevant path is between the storage system and its Ethernet network rather than only between external hosts. Used together, these commands provide both interface-statistics visibility and active network diagnostic capability, making them a practical first-line toolset for Ethernet troubleshooting in Data ONTAP 7-Mode.

NetApp's legacy command reference describes `ifstat` as displaying device-level statistics for network interfaces and documents the associated diagnostic tools: ifstat(1) command reference and netdiag(1) command reference.

QUESTION NO: 8

To manage SnapVault relationships with an NDMP management application, you must specify a:

A. directory path

B. volume and qtree

C. SnapVault license

D. username and password

ANSWER: D

Explanation:

username and password is correct because an NDMP management application must authenticate to the NetApp storage system before it can issue NDMP-related operations associated with SnapVault management. In Data ONTAP 7-Mode, this is normally accomplished by creating a backup user with the required Backup Operators group membership and then generating an NDMP-specific password for that user with the `ndmpd password` command. The management application uses this user name and generated password when it connects to the storage system's NDMP service. These credentials identify the account authorized for the backup-management workflow and allow the application to establish the authenticated NDMP session needed to coordinate the relationship and backup activity. The password generated for NDMP is specifically intended for the external NDMP application configuration; it is not merely a generic storage-system login detail. NetApp's SnapVault documentation describes adding the backup user and generating its NDMP password before configuring the management station. NetApp's NDMP documentation also explains that NDMP clients require credentials to access the storage system NDMP service. See the NetApp SnapVault documentation and NetApp NDMP documentation.

QUESTION NO: 9

When viewing the output of the sysstat –f command, which three fields are part of the output? (Choose three)

A. FCP

B. FTP

C. CIFS

D. NFS

E. iSCSI

ANSWER: A C D

Explanation:

The `sysstat -f` command selects the FCP-oriented default display format in Data ONTAP 7-Mode. Its heading combines controller-utilization and protocol-activity information, allowing an administrator to view workload indicators in one recurring statistics display. The documented heading includes `CPU`, `NFS`, `CIFS`, `FCP`, network throughput, disk throughput, FCP throughput, and cache-age information. Therefore, **FCP**, **CIFS**, and **NFS** are all fields in the output of this command. FCP represents Fibre Channel Protocol activity; CIFS and NFS identify client protocol activity reflected in the system statistics. The accompanying throughput columns provide related read, write, inbound, and outbound rate values, but the requested fields are explicitly present in the principal format header. This is useful when assessing whether workload pressure is associated with NAS protocols, Fibre Channel access, or broader controller and disk activity. The 7-Mode `na_sysstat` manual documents the `-f` switch and shows the complete default FCP-format header, including these fields: na_sysstat(1) manual page. NetApp's documentation portal is available at NetApp Documentation.

QUESTION NO: 10

In a tape environment after a SnapRestore reversion of a volume, incremental backup and restore operations on the file or volume cannot rely on the timestamps to determine what data needs to be backed up or restored. Which course of action ensures correct incremental backups?

A. Perform a base-level backup of the volume after you restore it.

B. Reboot the storage system to reset timestamps back to their original times before the reversion.

C. Take a new snapshot of the volume to create new base-level timestamps for the incremental backup.

D. Take a new snapshot of the volume to create new base-level timestamps for the incremental backup, and then perform a base-level backup of the volume after you restore it.

ANSWER: A

Explanation:

Perform a base-level backup of the volume after you restore it. is the required action because a SnapRestore reversion returns the volume to the state preserved in a Snapshot copy, including file-system metadata that backup software may use to identify changes. Following the reversion, the relationship between the tape backup history and the volume's current data and timestamps is no longer reliable for determining the changed data set. A base-level backup, also called a baseline or full backup, creates a new known recovery point that accurately represents the reverted volume. Subsequent incremental tape backups can then compare the current volume state against that newly established baseline and reliably capture changes made after it. This procedure also ensures that a later restore sequence has a valid full backup on which its associated incremental backups depend. NetApp documentation describes SnapRestore as a rapid reversion of a volume or files to the contents of a Snapshot copy and notes the operational considerations that apply after reverting data. See the NetApp Snapshot restore documentation and the NetApp data protection documentation for the role of consistent baseline recovery points in protection workflows.

QUESTION NO: 11

To restore a qtree protected with SnapVault from the secondary storage system to the primary storage system, you can _____. (Choose two.)

A. restore to a non existing qtree with DOT 7. 2

B. add the new qtree to the primary storage system with DOT 7.2

C. overwrite the existing qtree on the primary storage system with DOT 7.3

D. delete the existing qtree from the secondary storage system with DOT 7.3

ANSWER: A C

Explanation:

SnapVault qtree restore behavior differs between these Data ONTAP releases. With Data ONTAP 7.2, a SnapVault restore to the primary system requires a destination qtree that does not already exist. Therefore, *restore to a non existing qtree with DOT 7. 2* is a supported recovery method: the restore operation creates and populates the destination qtree from the backed-up data on the secondary system.

Data ONTAP 7.3 added support for restoring SnapVault data to an existing primary qtree. This supports both baseline and incremental restore workflows, enabling the backed-up secondary copy to replace the contents of the established destination qtree as appropriate for the selected restore operation. Consequently, *overwrite the existing qtree on the primary storage system with DOT 7.3* is also valid. These release-specific restore capabilities are important when planning qtree recovery, because the destination qtree's prior existence determines which restore procedure is available. NetApp's archived Data ONTAP documentation describes SnapVault qtree restore behavior and the 7.3 enhancement for restores to existing qtrees: SnapVault restore documentation. For current SnapVault concepts and replication context, see NetApp ONTAP data protection documentation.

QUESTION NO: 12

Which three statements apply to the SnapMirror visibility_interval argument? (Choose three.)

A. The default interval value is five minutes.

B. The smallest interval value supported is 30 seconds.

C. This variable controls the view of the data on the destination system.

D. By setting a value, you specify the amount of time before an automatic snapshot is created on the source volume that is synchronously mirrored.

ANSWER: B C D

Explanation:

The statements "The smallest interval value supported is 30 seconds," "This variable controls the view of the data on the destination system," and "By setting a value, you specify the amount of time before an automatic snapshot is created on the source volume that is synchronously mirrored." accurately describe the 7-Mode SnapMirror `visibility_interval` argument. This setting applies to synchronous and semi-synchronous SnapMirror relationships. It determines how frequently Data ONTAP creates the automatic Snapshot copy on the source volume that makes replicated changes visible at the destination. Although writes may already have been synchronously replicated, the destination's visible point-in-time view advances when this automatic Snapshot copy is created and propagated. The default interval is three minutes, and the documented minimum is 30 seconds. Administrators should choose the interval carefully: a shorter interval provides more frequent destination visibility, but it also increases Snapshot-related processing and can affect mirror performance. The argument accepts a number with an optional seconds, minutes, or hours suffix; without a suffix, the value is interpreted in seconds. It is not used for asynchronous mirrors. See the archived snapmirror.conf(5) manual page and NetApp's SnapMirror configuration and best-practices technical report.

QUESTION NO: 13

Which statement is true about expanding an aggregate from 32-bit to 64-bit in place?

A. All aggregates are automatically converted from 32-bit to 64-bit with the Data ONTAP 8.1 upgrade.

B. The expansion is triggered by an aggr convert command.

C. The expansion is triggered by adding disks to exceed 16 TB.

D. The 32-bit aggregates are degraded and must be Volume SnapMirrored to a new 64-bit aggregates with Data ONTAP 8.1 upgrade.

ANSWER: C

Explanation:

The statement “The expansion is triggered by adding disks to exceed 16 TB.” is correct. In Data ONTAP 7-Mode, a 32-bit aggregate can use an in-place expansion process to become a 64-bit aggregate when its physical size grows beyond the 16 TB addressing limit of a 32-bit aggregate. The administrator triggers this conversion operationally by adding sufficient disks to increase the aggregate past that threshold; Data ONTAP then begins the conversion automatically rather than requiring a separate aggregate-convert command.

This behavior is important when planning capacity growth for older 32-bit aggregates. The conversion is not merely a metadata flag change: Data ONTAP progressively expands the aggregate’s block-number addressing capability while keeping the aggregate online. As a result, the aggregate can continue to grow beyond the 32-bit limit and support the larger address space provided by 64-bit aggregate formatting. Administrators should plan disk additions and available capacity carefully because the conversion is initiated by crossing the size threshold and may take time to complete depending on aggregate size and system activity.

NetApp’s 7-Mode transition documentation describes the distinction between 32-bit and 64-bit aggregates and their capacity implications: NetApp documentation: 32-bit and 64-bit aggregates. Additional background on the in-place expansion behavior is available at How in-place expansion works.

QUESTION NO: 14

Which two statements apply to the fs_size_fixed volume option? (Choose two.)

A. The default setting is on.

B. This option is automatically set to be on when a volume becomes a SnapMirrored volume.

C. This option turns to off after the snapmirror break command is issued for the relationship.

D. This option causes the SnapMirror destination volume to remain the same size and not grow or shrink when a SnapMirrored volume relationship is broken.

ANSWER: B D

Explanation:

The `fs_size_fixed` setting is automatically enabled when a volume becomes a SnapMirror destination. It preserves the destination file system’s existing size even when the underlying volume has additional space, which is important for maintaining the size relationship expected by SnapMirror. When the SnapMirror relationship is broken, this setting remains enabled rather than allowing the file system to expand or contract automatically. Consequently, the former destination volume remains at its established file-system size and can later be resynchronized back to its original source without requiring the source volume to be expanded merely to match filesystem growth on the destination. This behavior is particularly useful after a disaster-recovery failover, because it prevents an unintended file-system-size change while the destination is writable. The default value for `fs_size_fixed` is off for ordinary volumes; SnapMirror enables it as part of configuring the destination volume. NetApp’s legacy Data ONTAP documentation and 7-Mode transition guidance describe the preservation of SnapMirror relationship characteristics and destination-volume behavior: NetApp 7-Mode Transition documentation and NetApp ONTAP 7-Mode Transition documentation.

QUESTION NO: 15

Which three statements describe FlexClone volumes in Data ONTAP 7-Mode? (Choose three)

A. A FlexClone volume is created from a Snapshot copy.

B. A FlexClone volume is writable.

C. A FlexClone volume must reside in the same aggregate as its parent volume.

D. A FlexClone volume requires a SnapMirror baseline transfer before it can be used.

E. A FlexClone volume immediately reserves space equal to the parent volume size.

ANSWER: A B C

Explanation:

A FlexClone volume is created from a Snapshot copy of a parent FlexVol volume. The clone initially shares the parent volume's data blocks, which makes creation rapid and space efficient. Although it is based on a Snapshot copy, the resulting FlexClone volume is writable and can be presented or exported independently of the parent.

A FlexClone volume must reside in the same aggregate as its parent because the shared WAFL blocks are located in that aggregate. It does not require a separate SnapMirror baseline, and it does not initially reserve capacity equal to the full size of the parent volume. See the NetApp FlexClone volume documentation.

QUESTION NO: 16

Which two are true when creating a SyncMirror volume? (Choose two.)

A. The disks selected for Plex 0 and Plex 1 must come from the same disk shelf.

B. The disks selected for Plex 0 and Plex 1 must come from the same disk pools.

C. The number of disks added must be an even number and equally divided between the two plexes.

D. The disks are selected first on the basis of equivalent bytes per sector (bps) size, then on the basis of the size of the disk. If there is no equivalent-sized disk, Data ONTAP takes a larger capacity disk and right-sizes it.

ANSWER: C D

Explanation:

SyncMirror protects an aggregate by maintaining two mirrored plexes, with each plex containing an equivalent RAID layout. Consequently, the same number of disks or array LUNs is required in each plex. When disks are being added as part of SyncMirror creation, the total number must therefore be divisible evenly between the plexes; this ensures that each side of the mirror has the same disk count and matching usable layout.

Data ONTAP also follows a specific disk-selection sequence to maintain compatible plexes. It first seeks disks with an equivalent bytes-per-sector size and then considers disk capacity. If an identically sized disk is unavailable, Data ONTAP can select a larger-capacity disk and limit its usable size so that it matches the corresponding disk. This right-sizing behavior allows the plexes to remain symmetric even where available physical disk capacities differ.

These requirements support the fundamental SyncMirror design goal: two equivalent copies of aggregate data distributed across separate disk pools for continued availability after a pool-level failure. NetApp documentation on aggregate and disk management is available in the ONTAP Disks and Aggregates documentation and the NetApp ONTAP systems setup documentation.

QUESTION NO: 17

For each Open Systems platform directory to be backed up to the SnapVault secondary storage system you must execute _____.

A. An initial baseline copy

B. A temporary copy

C. An incremental copy

D. A scheduled update copy

ANSWER: A

Explanation:

An initial baseline copy is required for every Open Systems SnapVault directory relationship because it establishes the first complete protected image of that directory on the SnapVault secondary system. During this baseline operation, the Open Systems platform transfers the specified directory's initial data set to its designated qtree location on the secondary storage system. This creates the foundation against which subsequent SnapVault updates can identify and transfer only changed data. In Open Systems SnapVault, the source is an Open Systems host rather than a primary Data ONTAP volume, so the baseline is the essential first transfer that initializes the relationship and provides the secondary copy from which later protection updates proceed. The initial baseline must be performed separately for each directory selected for backup, ensuring that each directory has an initialized destination and an established reference point for later updates. NetApp's 7-Mode SnapVault documentation describes this process as the transfer of initial base images of specified directories to qtrees on the secondary system. See the NetApp documentation for Open Systems SnapVault backup behavior and the SnapVault replication concept.

QUESTION NO: 18

Which tool is available on an OSSV client to assist in starting or stopping the OSSV service, modifying NDMP settings, enabling debugging, and capturing trace files?

A. Operations Manager

B. Protection Manager

C. Svconfigurator GUI

D. svdebug.exe

ANSWER: C

Explanation:

Svconfigurator GUI is the Open Systems SnapVault client-side configuration and troubleshooting utility. It provides a graphical interface for common OSSV administrative tasks, including controlling the OSSV service, configuring NDMP-related parameters used for backup and restore integration, enabling diagnostic logging, and collecting trace information for support analysis. These functions are grouped in this utility so that an administrator does not need to manually alter configuration files or run separate service and tracing commands for routine OSSV maintenance. In particular, trace capture and debug settings are important when investigating replication, connection, or data-transfer issues between an OSSV client and a NetApp storage system. NetApp's Open Systems SnapVault documentation describes the OSSV configuration utility and its role in managing client settings and diagnostics. See the NetApp Technical Report TR-3466 and NetApp's SnapVault replication documentation for background on SnapVault data-protection operations.

QUESTION NO: 19

Which are three characteristics of an iSCSI or FCP SAN implementation? (Choose three.)

A. A LUN can be converted into a VLD.

B. LUNs can be accessed using the FCP protocol.

C. LUNs can be accessed using the iSCSI protocol.

D. Block services are associated with an iSCSI or FCP SAN environment.

ANSWER: B C D

Explanation:

An iSCSI or Fibre Channel Protocol SAN implementation provides block-level storage to hosts through LUNs. A LUN is a logical unit of storage presented by the NetApp system to an initiator or host; the host then treats that LUN much like a locally attached block device and can create a supported filesystem or use it with an application. Fibre Channel Protocol is a SAN transport that carries SCSI commands over a Fibre Channel fabric, so LUNs can be accessed through the FCP protocol. iSCSI similarly carries SCSI block commands over TCP/IP networks, enabling hosts to access LUNs through iSCSI initiators and properly configured storage targets. Therefore, both protocols are valid methods for delivering LUN-based storage. Because iSCSI and FCP expose raw blocks rather than a shared file namespace, block services are the defining storage model for these SAN environments. NetApp documentation describes SAN as block-based storage using LUNs and identifies both iSCSI and FCP as supported SAN protocols. See the NetApp SAN administration documentation and the NetApp SAN configuration documentation.

QUESTION NO: 20

Which set of protocols provide block-level access to NetApp storage?

A. SNA and RPC

B. iSCSI and FCP

C. HTTP and FTP

D. CIFS and NFS

ANSWER: B

Explanation:

iSCSI and FCP provide block-level access to NetApp storage. Both are SAN protocols that transport SCSI commands between a host and NetApp LUNs. A LUN is presented to the host as a raw block device, much like a locally attached disk. The host operating system can then partition the device, create a filesystem, or use it directly for applications such as databases.

iSCSI carries SCSI block-storage commands over standard TCP/IP networks, allowing hosts with Ethernet connectivity to access SAN LUNs. FCP, meaning Fibre Channel Protocol, carries SCSI commands across a Fibre Channel SAN fabric. In NetApp 7-Mode environments, both protocols are configured for SAN access and use initiator identities and LUN mapping to control which hosts can access specific LUNs.

This distinction is fundamental to NetApp storage administration: block protocols expose storage as addressable blocks to the host, while file-serving protocols expose files and directories through a NAS namespace. NetApp documentation describes SAN environments as using LUNs accessed through protocols including iSCSI and Fibre Channel. See the NetApp SAN administration documentation and the NetApp SAN configuration documentation.

QUESTION NO: 21

Is NetApp storage Encryption supported in Data ONTAP 8.1.1 Cluster-Mode?

A. No, but you can file a PVR to request support.

B. No, it is targeted for a future release of Data ONTAP.

C. Yes, only with a special license installed.

D. Yes, it has been supported since 8.0.1.

ANSWER: B

Explanation:

No, it is targeted for a future release of Data ONTAP. is correct for the specific platform and release named in the question. Although NetApp Storage Encryption (NSE) was introduced for Data ONTAP 8.0.1, its initial availability applied to the 7-Mode implementation; it did not mean that NSE was available in every Data ONTAP personality or deployment model. Data ONTAP 8.1.1 Cluster-Mode had a separate feature-support lifecycle, and NSE support had not yet been delivered for that environment. Therefore, an administrator using 8.1.1 Cluster-Mode could not deploy NSE merely by obtaining a license, installing self-encrypting drives, or configuring external key management. The required Cluster-Mode integration was planned for a later Data ONTAP release. This distinction is important when reviewing historical NetApp feature matrices: the release number alone is insufficient unless the applicable operating mode is also considered. NetApp's encryption-at-rest documentation describes NSE as a self-encrypting-drive solution and documents its platform-specific requirements: NetApp ONTAP encryption at rest documentation. NetApp also provides product information for NSE and its key-management model at NetApp data encryption.

QUESTION NO: 22

Which two actions can occur when a volume reaches its Snapshot reserve? (Choose two)

A. New Snapshot copies can use volume free space.

B. Snapshot creation can fail when no volume space remains.

C. The aggregate automatically expands by the size of the reserve.

D. The volume is automatically converted to read-only.

ANSWER: A B

Explanation:

When the Snapshot reserve is full, new Snapshot copies can consume free space in the active file system portion of the volume. If both the Snapshot reserve and the volume free space are exhausted, Snapshot creation can fail. Existing Snapshot copies can also be deleted manually or automatically through configured Snapshot autodelete policies to reclaim blocks. Snapshot reserve consumption does not automatically increase the aggregate size, and it does not make the active file system read-only merely because the reserve has filled. See the NetApp Snapshot reserve documentation.