

Administration of Veritas InfoScale Storage 7.1 for UNIX/Linux

Veritas VCS-257

Version Demo

Total Demo Questions: 12

Total Premium Questions: 128

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Introduction to InfoScale Storage	20
Topic 2, Volume Manager Configuration	18
Topic 3, Volume Manager Administration	31
Topic 4, Volume Manager Advanced Administration	20
Topic 5, File System Configuration	6
Topic 6, File System Administration	22
Topic 7, Cluster File System	8
Topic 8, SmartTier	3
Total	128

QUESTION NO: 1

What are two advantages of a RAID-5 volume layout? (Select two.)

- A. protection from a single-disk failure
- B. capacity efficiency through distributed parity
- C. elimination of write overhead
- D. protection from simultaneous failure of two disks
- E. automatic replication to a remote site

ANSWER: A B

Explanation:

A RAID-5 volume provides redundancy through distributed parity. Data and parity information are striped across the disks in the RAID-5 plex, so the volume can continue operating after the loss of one disk in the plex. This provides protection against a single-disk failure without requiring a complete second copy of all data as mirroring does.

RAID-5 also uses storage capacity more efficiently than a mirrored layout. The equivalent capacity of one disk is used for distributed parity, while the remaining capacity is available for data. The trade-off is that write operations can incur parity-update overhead, so RAID-5 is not necessarily the preferred layout for write-intensive workloads. Veritas documents supported volume layouts in the [InfoScale Storage documentation](#).

QUESTION NO: 2

Which command should an administrator use to perform an online migration of a native file system to a Veritas file system?

- A. fscdsadm
- B. vxcdsconvert
- C. fsmigadm
- D. vxvmconvert

ANSWER: C

Explanation:

`fsmigadm` is the Veritas File System (VxFS) administration command designed for online file-system migration. It supports converting an eligible native Linux file system, such as ext2 or ext3, to VxFS while the source file system remains mounted and available. This lets an administrator move to VxFS with substantially less service disruption than an offline backup, reformat, restore, and remount procedure. The command manages the migration process and maintains the operational state needed to complete the conversion safely; administrators should nevertheless verify supported source file-system types, platform requirements, free-space requirements, and perform appropriate backups before beginning a production migration. After migration, the file system uses VxFS capabilities and can be administered with the normal VxFS tools. Veritas documents VxFS administration and migration procedures in its InfoScale Storage documentation set; use the version-specific documentation for the installed release, including the [InfoScale Storage 7.1 documentation](#). Veritas product documentation and manuals are also available through the [Veritas Documentation portal](#).

QUESTION NO: 3

An administrator needs to examine Dynamic Multi-Pathing devices and the physical paths associated with them.

Which two `vxddmpadm` subcommands provide this information? (Select two.)

- A. vxddmpadm getdmpnode

- B. vxddmpadm getsubpaths
- C. vxddmpadm getplexes
- D. vxddmpadm getvolumes
- E. vxddmpadm getdiskgroups

ANSWER: A B

Explanation:

vxddmpadm getddmpnode displays DMP nodes, which are the DMP device representations that Volume Manager uses for multipathed storage. This output lets an administrator identify the DMP device and its associated characteristics.

vxddmpadm getsubpaths displays the individual subpaths associated with DMP devices. The output includes path-related information such as the path name and state, helping an administrator investigate path availability and verify multipathing configuration.

QUESTION NO: 4 - (DRAG DROP)

Click and drag three objects from the list that can be managed by the cvm service group in a Cluster File System (CFS) storage cluster to the right.

Select and Place:

Object

Shared disk groups

Agent framework

I/O fencing

Cluster file systems

CFS infrastructure

GAB ports

cvm service group manageable
(in no particular order)

⬅

➡


⬆

⬇

ANSWER:

Object

cvm service group manageable (in no particular order)



Explanation:

In a Cluster File System storage cluster, the cvm service group is the parallel service group used to bring the clustered storage stack online consistently on the cluster nodes. Its managed scope includes the shared disk groups that provide the common Volume Manager storage, the CFS infrastructure required for coordinated clustered file-system operation, and the cluster file systems that are mounted from that shared storage. Keeping these resources together in the cvm service group establishes the required operating order: the clustered-volume and disk-group layer is available first, the CFS support infrastructure is active, and then the file systems can be made available to applications across the participating nodes. This arrangement lets VCS monitor the storage stack as a coordinated unit and perform orderly online, offline, and recovery actions when necessary. The group is parallel because these clustered storage components are intended to be active on all appropriate nodes rather than owned by only one node. Veritas documentation describes the CFS environment as an integrated cluster storage solution built around clustered volume management and clustered file systems; the corresponding VCS service-group configuration is what manages those operational resources. See the [Veritas product documentation portal](#) for the InfoScale Storage and Cluster File System administration guides, including the service-group and CFS configuration material.

QUESTION NO: 5

Which two Veritas File System utilities are used to administer disk quotas? (Select two.)

- A. vxquota
- B. vxedquota
- C. vxresize
- D. vxtunefs
- E. vxfscnvert

ANSWER: A B

Explanation:

`vxquota` is used to display quota information for a Veritas File System. It can report the configured limits and current disk-space or inode usage for users and groups on a file system where quotas are enabled.

`vxedquota` is used to create or modify quota limits. It presents the quota information for editing so that an administrator can define the allowed disk-block and file-count limits for the selected user or group. These utilities are VxFS-specific quota administration commands.

QUESTION NO: 6

When should an administrator use the `-o tmplog` option while mounting a Veritas file system?

- A. to track file and directory changes when a nonpersistent file change log is required for replication
- B. to improve performance for temporary file systems by bypassing the intent log
- C. to improve performance for temporary file systems at the cost of higher data loss upon system failure
- D. to add a temporary data change log by enabling Veritas file system snapshot operations

ANSWER: C

Explanation:

The correct use is *to improve performance for temporary file systems at the cost of higher data loss upon system failure*. The `tmplog` mount option directs Veritas File System (VxFS) to use an in-memory, temporary intent log rather than a persistent on-disk intent log. Intent logging normally records metadata updates so that, after an unexpected interruption, VxFS can quickly replay the log and restore filesystem consistency. Keeping that log in memory avoids the disk I/O associated with persistent logging, which can be useful for filesystems intended only for transient data, such as scratch space, short-lived work areas, or rebuildable intermediate output.

Because the log is not persistent, its contents disappear if the host crashes or loses power. Therefore, updates that had not been safely committed to disk can be lost, and normal intent-log recovery is unavailable. An administrator should use this setting only when the performance benefit is more important than preserving recent filesystem changes and the application can tolerate recreating the temporary data. Veritas describes VxFS mount options and intent-log behavior in its InfoScale Storage documentation: [Veritas InfoScale Storage documentation](#). General VxFS administration resources are also available through the [Veritas InfoScale documentation portal](#).

QUESTION NO: 7

CORRECT TEXT

How many Veritas Volume Manager I/O kernel threads are started at boot time, by default?

- A. 16

ANSWER: A

Explanation:

16 is the default number of Veritas Volume Manager I/O kernel threads started during system boot. These threads, commonly associated with the VxVM I/O daemon infrastructure, provide the kernel-level resources that Volume Manager uses to process I/O work for managed volumes. Starting with 16 threads provides a practical default level of parallelism for

normal host workloads while avoiding unnecessary kernel-resource consumption on systems that do not require a larger I/O-thread pool.

This value is a default rather than an absolute sizing requirement. An administrator can review and tune relevant VxVM I/O settings when workload characteristics, processor availability, storage latency, or heavy concurrent volume activity justify doing so. For the exam question, however, the requested boot-time default is specifically 16. Veritas documents Volume Manager administration and its tunable operating characteristics in the InfoScale Storage documentation set. See the [Veritas InfoScale Storage documentation](#) and the [Veritas Volume Manager documentation](#) for Volume Manager configuration and tuning guidance.

QUESTION NO: 8

How should an administrator configure SmartIO to ensure that data files on a specific Veritas file system remain in SmartIO cache?

- A. Assign cache priority to the file system files with the `sfcache set` command.
- B. Flush the cache and prioritize the file system files with the `sfcache app cachearea` command.
- C. Load the file system files into cache with the `sfcache pin` command.
- D. Load the file system files into cache using the `sfcache load` command.

ANSWER: C

Explanation:

Use `sfcache pin` to pin the required files in the SmartIO cache. Pinning is the SmartIO mechanism intended for data that must be retained in cache: it marks the selected file data as pinned so that normal cache replacement activity does not evict it when cache space is needed for other I/O. This is appropriate when an administrator has identified files on a particular Veritas File System that have predictable, sustained performance importance and wants those files to remain cache-resident. The command is applied to the relevant file or file set after SmartIO caching has been configured for the file system; administrators should also ensure that the cache has sufficient capacity for the pinned working set. Because pinned data consumes cache capacity until it is unpinned or otherwise removed through an administrative action, pinning should be used selectively for the most important data rather than broadly across all files. Veritas documents SmartIO administration, including cache operations and file-based caching behavior, in the InfoScale Storage documentation set. See the [Veritas InfoScale Storage documentation](#) and the [SmartIO Blueprint for Linux](#).

QUESTION NO: 9

What is a benefit of using SmartTier?

- A. it enables one-time assignment of application files to a storage class
- B. it allows files to be automatically migrated between storage types based on access frequency
- C. it enables files to be migrated in response to application access requests
- D. it transparently moves data based on policies that implement business requirements

ANSWER: D

Explanation:

SmartTier transparently moves data based on policies that implement business requirements. In Veritas InfoScale Storage, SmartTier is a policy-based data-placement and movement capability for file-system data. Administrators define tiers of storage and create policies that express the organization's requirements, such as placing data with particular characteristics on an appropriate storage tier or relocating it as its lifecycle changes. SmartTier evaluates the applicable policies and performs the file movement without requiring users or applications to change the file path or manually manage the underlying tier. This preserves transparent file access while helping organizations align storage cost and performance with the business

value of their data. The policy-driven approach is the central benefit: it enables ongoing, automated management of data placement rather than a static placement decision. Veritas documentation for the InfoScale Storage product family and its SmartTier administration content is available through the [InfoScale Storage 7.1 documentation portal](#) and the [Veritas product documentation portal](#).

QUESTION NO: 10

What are two advantages of a properly configured striped volume layout? (Select two.)

- A. reduced SAN traffic
- B. load balancing
- C. improved performance
- D. data parity
- E. improved reliability

ANSWER: B C

Explanation:

A properly configured striped volume layout provides **load balancing** because Volume Manager divides a volume into stripe-sized regions and distributes those regions across multiple disks. Reads and writes can therefore be serviced by more than one underlying disk path rather than concentrating activity on a single disk. This distribution helps spread I/O demand across the available storage resources, particularly when the stripe width, stripe unit size, disk characteristics, and workload are selected appropriately.

It also provides **improved performance** by enabling parallel I/O. For sequential or sufficiently concurrent workloads, multiple disks can transfer portions of the requested data at the same time, increasing aggregate throughput and reducing the likelihood that one disk becomes the sole I/O bottleneck. The performance benefit is workload- and configuration-dependent, so a stripe layout should be planned with the intended I/O pattern and storage architecture in mind. Veritas describes striped volumes as a layout that distributes data across multiple disks to improve I/O performance and balance disk activity. See the [Veritas InfoScale Storage documentation](#) and the [InfoScale Storage documentation portal](#).

QUESTION NO: 11

An administrator needs to install Veritas InfoScale Storage for UNIX/Linux on a system.

Which two requirements are validated by the Veritas Services and Operations Readiness Tools (SORT) Installation and Upgrade custom report?

- A. Any attached disk arrays are supported.
- B. The LUN UDID values are unique.
- C. Sufficient file system free space is available.
- D. The SAN fabric transfer speed is sufficient.
- E. SCSI-3 persistent group reservation is supported.

ANSWER: A C

Explanation:

The SORT Installation and Upgrade custom report is intended to assess whether a target host and its planned storage environment meet the compatibility and readiness requirements for an InfoScale installation or upgrade. "Any attached disk arrays are supported." is validated because array compatibility is an essential supportability consideration for InfoScale Storage and Veritas Volume Manager. The report uses the selected product, version, operating system, and environment

details to identify applicable support and compatibility information. "Sufficient file system free space is available." is also validated because the installation readiness assessment checks host prerequisites needed to stage and install product components successfully, including available space in relevant file systems. This allows an administrator to resolve environmental prerequisites before beginning the installer workflow, reducing the risk of an interrupted or unsupported deployment. SORT is Veritas's central service for generating installation and upgrade readiness reports and for consulting product compatibility information. See the [Veritas SORT portal](#) and the [InfoScale installation documentation](#) for installation planning and prerequisite guidance.

QUESTION NO: 12 - (SIMULATION)

What is the maximum supported number of nodes in a Flexible Storage Sharing configuration?

ANSWER: See the explanation for the answer

Explanation:

The maximum supported number of nodes in a Veritas InfoScale Flexible Storage Sharing (FSS) configuration is **8 nodes**.

FSS supports one node acting as the storage server and up to seven additional client nodes sharing the exported storage.