DUMPSBOSS.

Administration of Veritas InfoScale Storage 7.3 for UNIX/Linux

Veritas VCS-261

Version Demo

Total Demo Questions: 10

Total Premium Questions: 105

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Describe Veritas InfoScale Storage	28
Topic 2, Configure Veritas InfoScale Storage	13
Topic 3, Administer Veritas InfoScale Storage	29
Topic 4, Use Veritas InfoScale Storage features	31
Topic 5, Troubleshoot Veritas InfoScale Storage	4
Total	105

QUESTION NO: 1

The following command adds two snapshot mirrors to the volume voll on disks mydg10 and mydg11.

vxsnap –g mydg addmir voll nmirror=2 alloc=mydg10,mydg11 Which plex state indicates the snapshot is fully synchronized?

A. SNAPREADY

B. SNAPTMP

C. SNAPDONE

D. SNAPATT

ANSWER: C

Explanation:

SNAPDONE indicates that a snapshot plex created by `vxsnap addmir` has completed synchronization with its source volume. The `addmir` operation creates snapshot mirrors as additional plexes for the specified volume. Before a snapshot can be used reliably as a point-in-time copy, Volume Manager must copy and synchronize the relevant volume data to those snapshot plexes. Once that copy process is complete, the plex state changes to `SNAPDONE`, confirming that the snapshot mirror is fully synchronized and ready for subsequent snapshot processing, such as preparing and creating the snapshot volume.

This state is important operationally because it verifies that the snapshot mirror contains a consistent, complete copy of the source volume at the completion of mirror synchronization. Administrators can inspect plex status with commands such as `vxprint` while monitoring snapshot preparation. The Veritas InfoScale Storage documentation set includes the Volume Manager and snapshot command references for this release, including guidance for the `vxsnap` workflow and snapshot plex lifecycle: Veritas InfoScale Storage 7.3 documentation and Veritas SORT documentation portal.

QUESTION NO: 2

Which two commands can an administrator use to enable and disable Veritas File System disk quotas? (Choose two.)

A. vxquotaon

B. vxquotaoff

C. vxquotacheck

D. fsadm

E. vxfsstat

ANSWER: A B

Explanation:

`vxquotaon` enables VxFS disk quotas for a specified mounted file system. Once quotas are enabled, VxFS enforces the configured user or group quota limits for that file system.

`vxquotaoff` disables quota enforcement for the specified file system. It is used when an administrator needs to suspend quota checking, such as during quota maintenance or a configuration change. Quota limits and usage can be reviewed with quota-reporting utilities, but those utilities do not enable or disable quota enforcement. Refer to the InfoScale Storage Administrator's Guide for VxFS quota administration information.

QUESTION NO: 3

Which two commands can an administrator use to remove the datadg02 disk from the datadg disk group based on the following configuration? (Choose two.)

Evaluation only.
Created with Aspose.Imaging.
Copyright 2010-2019 Aspose Pty Ltd.

```
ault      dcfault    80000      1086905353.6.trcassun40
Disk_0      auto       20351      143328960 -
dm  datadg02    Disk_8      auto       20351      143328960 -

v   data        -           ENABLED    ACTIVE     6146304   ROUND     -        gen
pl  data-01     data        ENABLED    ACTIVE     6146304   CONCAT    -        RW
sd  datadg01-05 data-01     datadg01   22539839   6146304   0         Disk_0   ENA
pl  data-02     data        ENABLED    ACTIVE     6146304   CONCAT    -        RW
sd  datadg02-03 data-02     datadg02   22539840   6146304   0         Disk_8   ENA
```

A. vxdg

B. vxdiskunsetup

C. vxplex

D. vxdisk

E. vxvol

ANSWER: A B

Explanation:

The `vxdg` command can remove a disk from a specified Veritas Volume Manager disk group by using its `rmdisk` operation, such as `vxdg rmdisk datadg datadg02`. This is the standard disk-group-level administrative method when the disk is eligible for removal. Before removing a disk that contains allocated subdisks, an administrator must first evacuate or otherwise relocate its data so that removing the disk does not affect configured volumes.

The `vxdiskunsetup` command also supports removing a disk from its current disk group and returning it to an unconfigured state. It is useful when the goal is to unconfigure the disk after it has been detached from `datadg`. As with disk-group removal, the disk must not contain remaining allocated storage that needs to be preserved. Both commands therefore provide supported administration paths for removing `datadg02` from the `datadg` configuration, with the appropriate preparation dictated by the disk's allocation state.

Veritas documents Volume Manager disk-group and disk administration in the Veritas InfoScale Storage Administrator's Guide and provides product documentation access through the Veritas Documentation portal.

QUESTION NO: 4

Which two Veritas Volume Manager volume layouts provide data redundancy? (Choose two.)

A. mirror

B. RAID-5

C. concat

D. stripe

E. simple

ANSWER: A B

Explanation:

The **mirror** layout provides redundancy by maintaining complete copies of volume data in separate plexes. If one plex becomes unavailable, the volume can continue to use another synchronized plex.

The **RAID-5** layout provides redundancy by distributing parity information across the stripe columns. The parity data enables VxVM to reconstruct data following the loss of one disk in the RAID-5 plex. A concat or stripe layout alone does not provide redundant copies or parity protection. See the Veritas Volume Manager Administrator's Guide for volume-layout information.

QUESTION NO: 5

Which command is used to display current file system I/O parameters for a Veritas file system?

A. vxtune

B. vxtunefs

C. vxstat

D. fsadm

ANSWER: B

Explanation:

`vxtunefs` is the Veritas File System (VxFS) utility used to display and modify tunable file-system parameters, including parameters that control I/O behavior. When run against a mounted VxFS file system with the appropriate display-oriented syntax, it reports the current values of the relevant VxFS tunables. These settings can include I/O-related behavior such as read-ahead and other file-system performance controls, allowing an administrator to inspect the effective configuration before deciding whether tuning is required. The command operates at the VxFS layer, so it is the appropriate utility when the question concerns parameters of a Veritas file system rather than general volume, disk, or host statistics. Veritas documents `vxtunefs` as the command for tuning and viewing VxFS file-system parameters in the InfoScale Storage administration documentation. See the Veritas InfoScale Storage 7.3 documentation and the InfoScale Storage documentation portal for VxFS administration and command-reference material.

QUESTION NO: 6

Which two commands are used for thin reclamation? (Choose two.)

A. fsadm

B. vxdctl

C. vxdisk

D. vxtune

E. vxdmpadm

ANSWER: A C

Explanation:

`fsadm` and `vxdisk` are used together in InfoScale Storage thin-reclamation workflows. The `fsadm` command performs file-system-level reclamation by identifying unused space in a supported file system and issuing the required discard or unmap operations for those freed ranges. This lets the storage stack know that blocks no longer contain active file-system data. The `vxdisk` command performs the corresponding disk or array-facing reclamation operation, allowing Veritas Volume Manager to reclaim eligible thin-provisioned extents and communicate the release of those blocks to storage that supports thin provisioning and reclamation. In practice, reclamation depends on support throughout the stack: the file system, volume configuration, multipathing path, and storage array must support the relevant thin-reclamation capabilities. Using these commands helps return no-longer-used capacity to the thin-provisioned pool, improving utilization and avoiding unnecessary allocation growth. Veritas documents thin reclamation as part of Volume Manager administration and disk-management operations; see the InfoScale Storage documentation and the Veritas Volume Manager Administrator's Guide.

QUESTION NO: 7

An administrator uses the command shown below to create volume volzebra in the disk group mydg:

vxassist -b - g mydg make volzebra 10g layout=atripe

What is the default number of stripe columns for the volume?

A. 2

B. 3

C. 4

D. 5

ANSWER: A

Explanation:

2 is correct. For a VxVM volume created with the striped layout and without an explicit `nstripe` attribute, `vxassist` uses two stripe columns by default. A stripe column is a plex component allocated on a separate disk, and VxVM distributes successive stripe units across those columns. Therefore, the resulting 10-GB striped volume is built with two columns unless the administrator specifically requests another count, such as by including `nstripe=4` in the `vxassist make` command. This default provides a basic striped layout while allowing VxVM to select suitable disks in the target disk group. The layout keyword shown in the question is evidently intended to denote the standard striped layout; the applicable behavior is the documented default for stripe creation. Veritas documents `vxassist` volume-layout attributes and striped-volume configuration in the InfoScale Storage administration documentation. See the InfoScale Storage 7.3 Administrator's Guide and the Veritas InfoScale Storage documentation.

QUESTION NO: 8

Which Veritas InfoScale Storage feature is used when migrating Veritas File System data from thick to thin LUNs?

A. SmartIO

B. SmartMove

C. Compression

D. SmartTier

ANSWER: B

Explanation:

SmartMove is the Veritas File System feature designed for efficient migration of file-system data between storage volumes, including a migration from traditionally provisioned (thick) LUNs to thin-provisioned LUNs. Rather than copying every block in the source volume, SmartMove identifies and copies only the blocks that are actually allocated and contain file-system data. This is particularly important when the target is a thin LUN: unallocated regions on the source do not consume capacity on the destination merely because the source device has a larger apparent size.

By avoiding the transfer of unused blocks, SmartMove can reduce the amount of physical thin-pool capacity consumed, decrease I/O during the migration, and shorten the migration window. It is commonly used with Veritas volume migration operations where preserving efficient thin provisioning is required. The feature operates with VxFS awareness, so its copy behavior is based on file-system allocation information rather than a simple block-for-block device copy. See Veritas InfoScale Storage documentation in the Veritas Documentation portal and the Veritas knowledge base for SmartMove and thin-reclamation guidance.

QUESTION NO: 9

Which two changes can the Veritas File System File Change Log record when it is enabled? (Choose two.)

A. file creation

B. file deletion

C. storage array controller replacement

D. DMP path failure

E. disk group import

ANSWER: A B

Explanation:

The File Change Log records file-system changes, including file creation and file deletion. This change information can be used by supported features to identify files that require policy evaluation without requiring a complete scan of the file system.

The File Change Log does not record storage-array configuration events or Dynamic Multi-Pathing path failures. Those events occur below the VxFS file-system layer and are not file-system changes. See the InfoScale Storage Administrator's Guide for VxFS File Change Log administration.

QUESTION NO: 10

Which command should an administrator use to guarantee data is cached when SmartIO is configured?

A. vxcache set autocache

B. sfcache preload

C. vxassist cache pin

D. sfcache pin

ANSWER: D

Explanation:

The `sfcache pin` command is used to guarantee that selected data remains cached in a SmartIO cache area. SmartIO normally applies caching policies that can populate and evict data according to workload behavior and available cache capacity. Pinning explicitly identifies the data that must be retained, so it is the appropriate administrative action when an application requires predictable cache residency rather than relying on automatic caching decisions. The command is part of the SmartIO `sfcache` command interface and is used after SmartIO caching has been configured for the relevant volume or file system. By pinning frequently accessed or performance-sensitive files, an administrator can help ensure reads are served from the faster cache tier while the pinned content remains eligible for that guarantee. This is especially useful for critical database files, application binaries, or other consistently high-demand data sets. Veritas product documentation for InfoScale Storage and its SmartIO functionality is available through the Veritas InfoScale Storage documentation and the Veritas SORT documentation portal.