PassQueen
a.k.a dumpsqueen.com

Administration of Veritas InfoScale Availability 7.3 for UNIX/Linux

Veritas VCS-260

Version Demo

Total Demo Questions: 15

Total Premium Questions: 154

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

Which two capabilities can an administrator use to ensure availability within Docker Containers? (Select two.)

A. provide HA to the application within the container for failover and app restart

B. provide provisioning of a new container when an application within the container fails

C. first restart the application within the container and if continued app failures then provision a new container

D. provide HA to the container and move it between cluster nodes

E. provide provisioning of a new container when a container fails

ANSWER: A D

Explanation:

Veritas InfoScale Availability supports two complementary availability approaches for Docker-based workloads. “provide HA to the application within the container for failover and app restart” is correct because VCS can monitor an application running in a container and take configured recovery action when the application is no longer healthy. This enables application-level availability without requiring the administrator to treat every application fault as a complete container or host failure. The recovery policy can restart the managed application and maintain service continuity according to the service group configuration.

“provide HA to the container and move it between cluster nodes” is also correct because VCS can manage the container as a highly available workload. When the container, its supporting resources, or its current node cannot sustain the service, VCS can fail over the service group to another eligible cluster node. This combines cluster membership, resource monitoring, dependency ordering, and failover behavior to restore the containerized service on a surviving node. These capabilities reflect the application-level and container/service-group-level protection models available through InfoScale Availability’s Docker integration. See the Veritas InfoScale 7.3.1 Linux Product Guides and the Veritas InfoScale Availability documentation for configuration and agent guidance.

QUESTION NO: 2

Which two resource attributes must the administrator set correctly to prevent a fault for the service group when a new resource is added? (Choose two.)

A. AutoDisabled

B. Critical

C. Enabled

D. Frozen

E. MonitorOnly

ANSWER: B C

Explanation:

Critical and **Enabled** are the resource attributes that must be deliberately configured when introducing a resource into a Veritas Cluster Server service group. Critical controls the effect of a resource fault on the service group. When a resource is designated as critical, its inability to come online or a detected fault can cause VCS to fault the service group and initiate the configured recovery behavior. Administrators must therefore set this attribute according to whether the resource is essential to the application’s availability.

Enabled controls whether VCS actively manages the resource. A resource that is enabled participates in normal online, offline, and monitoring operations; its configuration must be suitable for the current service-group configuration and

deployment sequence. Setting Enabled appropriately avoids unintentionally managing a newly added resource before it is ready, while still allowing it to be brought under VCS control when intended. Together, these attributes determine both whether the new resource is managed and whether its failure affects service-group health. Veritas describes VCS resource configuration and service-group behavior in its InfoScale Availability documentation and publishes product documentation through the Veritas SORT documentation portal.

QUESTION NO: 3

What is the impact to unsaved configuration changes in Veritas Cluster Server when an administrator stops the cluster with a force option from the command-line interface?

A. The changes are automatically saved.

B. The changes are automatically discarded.

C. The potential corruption of the cluster configuration.

D. The changes are prompted to be discarded.

ANSWER: B

Explanation:

The changes are automatically discarded. In Veritas Cluster Server, configuration modifications are initially held in the running cluster configuration and are not persistent until they are written to the main configuration file, typically by using `haconf -dump` or an equivalent save operation. A normal cluster shutdown can provide an opportunity to address pending, unsaved configuration changes. However, using the force form of the cluster-stop command, such as `hastop -force`, bypasses that interactive safeguard so the shutdown can proceed immediately. Consequently, pending in-memory changes are not saved and are lost when VCS stops. This behavior is important operationally: before a forced shutdown, an administrator should explicitly save any intended configuration updates to ensure that the configuration is retained when the cluster is restarted. The Veritas InfoScale Availability documentation describes VCS cluster administration and the command-line procedures for managing and stopping a cluster; see the Veritas InfoScale Availability 7.3 documentation and the Veritas SORT documentation portal for platform-specific command references.

QUESTION NO: 4

While reviewing the system logs, an administrator notices the following entries on previously running cluster node.

Which two actions should the administrator take to determine and resolve the issue? (Select two.)

A. Check whether the link has become physically disconnected from the system or switch

B. Check the content of the /etc/llttab file

C. Verify the NIC agent is started

D. Forcefully restart LLT with lltconfig

E. Check the NIC/link health and replace the NIC/link if necessary.

ANSWER: B E

Explanation:

Checking the content of the `/etc/llttab` file is appropriate because LLT uses this configuration file to define the node's private-network links, including the network interfaces and associated link definitions. Reviewing it confirms that the interfaces configured for LLT are present and correspond to the intended private interconnect configuration. This is an essential investigation step when a previously functioning cluster node reports communications or link-related events.

Checking NIC/link health and replacing the NIC/link if necessary addresses the physical and operating-system network layer on which LLT depends. The administrator should verify link state, cabling or switch connectivity, interface errors, and the NIC's operational status. A failed or degraded private-network adapter, cable, transceiver, or switch path can interrupt LLT heartbeats and cause cluster communication problems. Restoring a reliable private interconnect permits LLT and the cluster's membership communications to operate normally. Veritas describes LLT as the low-latency transport used for cluster communications and documents its network-link configuration and troubleshooting considerations in the InfoScale Availability documentation. Additional product documentation and platform-specific administration guidance are available through the Veritas InfoScale Availability support portal.

QUESTION NO: 5

Refer to the exhibit.

On which node will the wefosg service group start if all nodes are powered on and come online at the same time?

A. s4

B. s2

C. si

D. s3

ANSWER: B

Explanation:

The wefosg service group will start on **s2**. For a failover service group, VCS uses the service group's configured startup placement information when the group is initially brought online. When all eligible cluster systems become available concurrently, VCS evaluates the ordered systems defined for the group rather than relying on which node happened to complete its operating-system boot first. The exhibit's configuration places s2 at the applicable preferred position for wefosg, so VCS selects it as the initial host. This deterministic placement behavior prevents simultaneous cluster startup from producing an arbitrary service-group location and ensures that application resources start on the intended system whenever that system is available and eligible. The selection is governed by the group configuration, including its system-list ordering and automatic-start settings, which administrators can inspect with VCS configuration commands such as `hagrp -display wefosg`. Veritas documents service-group startup, system lists, and automatic start behavior in the InfoScale Availability Administrator's Guide: Veritas InfoScale 7.3 Administrator's Guide. The Veritas documentation portal is also available at Veritas Product Documentation.

QUESTION NO: 6

Which two capabilities can an administrator use to ensure availability within Docker Containers?

(Choose two.)

A. provide HA to the application within the container for failover and app restart

B. provide provisioning of a new container when an application within the container fails

C. first restart the application within the container; if continued app failures, then provision a new container

D. provide HA to the container and move it between cluster nodes

E. provide provisioning of a new container when a container fails

ANSWER: A D

Explanation:

Veritas InfoScale Availability supports high availability at two complementary layers for Docker-based workloads. The capability to “provide HA to the application within the container for failover and app restart” addresses the application layer. InfoScale can monitor the application process or service running in a container and take configured recovery action when the monitor detects a fault. This provides application-focused availability rather than merely confirming that a container process still exists.

The capability to “provide HA to the container and move it between cluster nodes” addresses the container workload at the cluster layer. InfoScale can manage a container as a highly available resource, monitor its state, and fail the workload over to another eligible cluster node when local recovery is not sufficient or the hosting node becomes unavailable. The combination of in-container application monitoring and cluster-node failover provides both rapid recovery for application faults and continued service when a container host fails.

These capabilities reflect InfoScale Availability’s integration of application-aware monitoring with clustered failover orchestration for containerized services. See the Veritas InfoScale Availability documentation and the Veritas InfoScale Availability product overview.

QUESTION NO: 7

Which attributes are required for configuring the Oracle agent for Veritas Cluster Server?

A. User, EnvFile

B. SID, DBAUser, DBAPword

C. SID, Owner, Home

D. SID, Owner, MonitorOption

ANSWER: C

Explanation:

SID, Owner, Home is correct because these are the fundamental required attributes that identify and control an Oracle database instance through the Veritas Cluster Server Oracle agent. **SID** specifies the Oracle system identifier for the database instance that the resource manages. **Owner** identifies the operating-system account under which Oracle is installed and under which the agent runs Oracle administrative commands. **Home** supplies the Oracle home directory, enabling the agent to locate the appropriate Oracle binaries, libraries, and utilities needed to start, stop, monitor, and clean the database resource. Together, these values establish the database identity, execution user, and software installation location required for reliable agent operation in the cluster. The attributes are configured on the Oracle resource, typically using the Cluster Server configuration mechanisms or the Cluster Management Console, and must accurately match the Oracle instance installation on every system that can host the service group. Veritas documents the Oracle agent resource type and its attributes in the InfoScale Availability agent documentation; see the Veritas InfoScale Availability documentation and the Veritas Cluster Server Agents for Linux guide.

QUESTION NO: 8

A resource named db_res fails to come online in service group DB_SG.

Which two logs should the administrator review first to investigate the failure? (Choose two.)

A. /var/VRTSvcs/log/engine_A.log

B. the log for the resource agent type

C. /etc/VRTSvcs/conf/config/main.cf

D. /etc/lltab

E. /etc/gabtab

ANSWER: A B

Explanation:

The VCS engine log, `/var/VRTSvcs/log/engine_A.log`, records the cluster engine processing associated with service-group operations, resource state changes, faults, and agent responses. It provides the cluster-level context for the failed online operation.

The agent log for the resource type records the resource-agent activity, including the commands and error output associated with online, offline, monitor, and clean operations. For example, the log for an Oracle resource is typically named `Oracle_A.log`. Reviewing both logs helps correlate the agent-level failure with the VCS engine response. Veritas documents VCS and agent logging in the Veritas SORT documentation portal.

QUESTION NO: 9

An administrator needs to determine the current state of resource app_res on each system in the cluster.

Which two commands can provide this information? (Choose two.)

A. hares -state app_res

B. hastatus -sum

C. gabconfig -a

D. lltstat -n

E. vxfentsthdw

ANSWER: A B

Explanation:

The `hares -state app_res` command displays the state of the named resource on the systems where it is configured. It is useful for directly checking whether the resource is Online, Offline, Faulted, or in a transitional state.

The `hastatus -sum` command displays a cluster status summary that includes resource state information by service group and system. It provides a broader view that can be used to correlate the state of app_res with its dependent resources and service-group placement. See the Veritas Documentation portal for VCS status and resource command references.

QUESTION NO: 10

A multi-tier application has Oracle at the first level and WebSphere at the second level, which depends on the Oracle application. The Oracle and WebSphere service groups are configured in two separate cluster. The administrator wants the following fault behavior:

when the Oracle service group faults, the WebSphere service group should stay online
when the Oracle service group comes back online, the WebSphere service group should be taken offline and brought back online

Which type of dependency must the administrator configure to acquire the desired behavior?

A. virtual business services with soft type of fault dependency

B. virtual business services with an online global firm dependency

C. virtual business services with firm type of fault dependency

D. virtual business services with restart type of fault dependency

ANSWER: D

Explanation:

Virtual business services with restart type of fault dependency is the appropriate configuration because it provides the required recovery sequence for an inter-cluster, multi-tier application. A restart fault dependency does not force the dependent WebSphere service group offline when the Oracle service group initially faults. This allows the upper-tier service group to remain online during the Oracle outage, as required. When the Oracle service group subsequently returns to the online state, the dependency triggers a restart of the dependent WebSphere service group. VCS takes WebSphere offline and then brings it online again, allowing the application tier to reconnect cleanly to the newly available database tier and re-establish any required sessions, connections, or initialization state. Virtual business services are used to model and manage dependencies between service groups that reside in separate clusters. The restart dependency behavior is specifically intended for cases where an upstream service recovery should cause a controlled restart of a dependent service, rather than an immediate shutdown at the time of the upstream fault. See the Veritas InfoScale Availability documentation library at Veritas InfoScale Availability manuals and the Veritas SORT documentation portal for VCS virtual business service and dependency configuration guidance.

QUESTION NO: 11

Which attribute can an administrator configure to allow a service group to come back online after a

persistent resource has faulted and other nodes are unavailable?

A. RestartLimit

B. AutoRestart

C. FailOverPolicy

D. AutoStartPolicy

ANSWER: B

Explanation:

AutoRestart is the service-group attribute that controls whether Veritas Cluster Server attempts to bring a service group online again on the same system after the group faults. This is particularly relevant when a persistent resource faults and there is no eligible alternate node to which the service group can fail over. With AutoRestart enabled, VCS can make a local restart attempt after the fault condition is cleared, allowing the service group to return to an online state without requiring an administrator to manually start it or waiting for another cluster node to become available.

The setting supports availability in configurations where a temporary resource failure should not leave an otherwise usable system permanently offline. The restart remains subject to normal VCS dependency handling, resource-agent behavior, and the underlying cause of the persistent-resource fault being resolved. Administrators configure this attribute at the service-group level, so its effect applies to the group's recovery behavior rather than merely changing an individual resource's monitor or restart settings. Veritas documents service-group recovery attributes, including AutoRestart, in the InfoScale Availability Administration Guide. See the Veritas InfoScale Availability documentation and the Veritas SORT documentation portal for the applicable 7.3 administration material.

QUESTION NO: 12

An administrator is configuring disk-based I/O fencing using coordinator disks.

Which two requirements apply to coordinator disks? (Choose two.)

A. They must be accessible to all cluster nodes.

B. They must support SCSI-3 persistent reservations.

C. They must be imported as Veritas disk groups.

D. They must contain application file systems.

E. They must be configured as LLT heartbeat interfaces.

ANSWER: A B

Explanation:

Coordinator disks must be accessible from every node that participates in the fenced cluster. The fencing driver uses the coordination points to arbitrate cluster membership and determine which nodes can retain access to shared data disks. A disk that is not visible to all cluster nodes cannot provide consistent coordination information.

The coordinator disks must also support SCSI-3 persistent reservations. I/O fencing uses persistent-reservation registration keys to exclude a departed or faulted node from shared storage and protect application data from split-brain access. Veritas documents coordinator-disk and SCSI-3 persistent-reservation requirements in the Veritas SORT documentation portal.

QUESTION NO: 13

An administrator assigns the coordinator disks to a cluster, then starts I/O fencing: but the fencing driver display following error message.

```
000068 09:37:33 2bdd5815 0 ... 3066 0 VXFEN WARNING V-11-1-56
Coordinator disk has key with cluster id 47813
which does not match local cluster id 57019
```

Which action should the administrator perform to clear the error message and configure I/O fencing to the local cluster?

A. use the command vxfenswap to replace the coordinator disk group

B. use the command vxfenadm to remove the key

C. use the command vxfenconfig to configure the vxfen driver

D. use the command vxfenclearpre to clear the keys

ANSWER: D

Explanation:

Use the command `vxfenclearpre` to clear the keys. I/O fencing uses SCSI-3 persistent-reservation keys on coordinator disks to identify cluster membership and to prevent split-brain conditions. If coordinator disks retain persistent-reservation registrations from a prior cluster configuration, a previous installation, or nodes that are no longer part of the current cluster, the fencing driver can detect those existing keys when it starts. The disks must be cleaned before they can safely be used as coordinator disks by the local cluster. The `vxfenclearpre` utility is designed for this recovery and reconfiguration task: it clears pre-existing persistent-reservation keys from the specified coordinator disks, allowing the local cluster's fencing configuration to establish its own registrations. This operation should be performed carefully, after confirming that the affected disks are truly coordinator disks and are not actively used by another running cluster. Veritas documents the role of coordinator disks and SCSI-3 persistent reservations in its InfoScale Availability administration materials; see the InfoScale Availability 7.3 documentation and the Veritas SORT portal for product-specific procedures and compatibility guidance.

QUESTION NO: 14

Which log can an administrator review to determine why a resource failed to come online?

A. /var/VRTSvcs/log/vbs_state_notify.log

B. /var/VRTSvcs/log/ProcessOnOnly_A.log

C. /var/VRTSvcs/log/HostMonitor_A.log

D. /var/VRTSvcs/log/engine_A.log

ANSWER: D

Explanation:

`/var/VRTSvcs/log/engine_A.log` is the primary VCS engine log for the local cluster node and is the appropriate starting point when a resource does not come online. The High Availability Daemon writes engine-level messages here for service-group processing, resource state changes, online and offline requests, fault handling, and communications with resource agents. When an online operation fails, the log records the affected service group and resource, the attempted operation, timestamps, and the status or error returned during processing. This lets an administrator correlate the failed online request with the engine's decision-making and identify the event that prevented the resource from reaching the ONLINE state. The `_A` suffix denotes the active engine log; VCS rotates these logs as they grow, so related historical entries may also be present in rotated engine logs. Reviewing this log at the time of the failure is a standard first troubleshooting step and provides the central VCS context needed to investigate the resource event. Veritas documents VCS administration and troubleshooting resources in its InfoScale Availability documentation and provides product documentation access through the Veritas Documentation portal.

QUESTION NO: 15

An administrator wants to make a single-instance Oracle database highly available and take advantage of Veritas Cluster File System.

Which solution(s) should the administrator deploy?

A. InfoScale Enterprise

B. InfoScale Availability

C. InfoScale Storage

D. InfoScale Foundation and InfoScale Availability

ANSWER: A

Explanation:

InfoScale Enterprise is the appropriate solution because it combines the capabilities required for both parts of the requirement: highly available application failover for a single-instance Oracle database and shared, clustered file-system access through Veritas Cluster File System. Its availability component provides Veritas Cluster Server functionality, including application monitoring, service-group management, and automated failover between cluster nodes. This allows Oracle to be managed as a highly available clustered application while remaining a single database instance that runs on one node at a time. Its storage component provides the Veritas storage stack and Cluster File System capabilities needed to present and manage shared storage consistently across the cluster. With CFS, the Oracle database files can be hosted on a shared file system accessible from all participating nodes, allowing the database service to restart on another node without requiring data relocation. Veritas positions InfoScale Enterprise for business-critical applications that need both high availability and enterprise storage and file-system services. See the Veritas InfoScale product overview and the InfoScale documentation portal for product and deployment documentation.