PassQueen
a.k.a dumpsqueen.com

# Administration of Veritas InfoScale Availability 7.1 for UNIX/Linux

## Veritas VCS-256

**Version Demo**

**Total Demo Questions: 11**

**Total Premium Questions: 118**

**Buy Premium PDF**

https://passqueen.com

support@passqueen.com

# Topic Break Down

| Topic | No. of Questions |
|---|---|
| Topic 1, Overview of InfoScale Availability | 12 |
| Topic 2, Configuring a basic cluster | 5 |
| Topic 3, Managing service groups | 25 |
| Topic 4, Managing cluster resources | 23 |
| Topic 5, Managing cluster communication | 9 |
| Topic 6, Managing cluster membership | 6 |
| Topic 7, Managing cluster operations | 16 |
| Topic 8, Configuring I/O fencing | 13 |
| Topic 9, Configuring notification | 1 |
| Topic 10, Managing global clusters | 8 |
| Total | 118 |

QUESTION NO: 1

An administrator needs to move the websg service group from node1 to node2 in a controlled manner.

Which command should the administrator use?

**A.** hagrp -switch websg -to node2

**B.** hagrp -online websg -sys node2

**C.** hasys -switch node1 node2

**D.** hares -switch websg -to node2

**ANSWER: A**

**Explanation:**

**hagrp -switch websg -to node2** is correct because the `hagrp -switch` command performs a controlled service group switchover to the specified target system. VCS takes the service group offline on its current system and brings it online on node2, processing the configured resource dependencies in the appropriate order.

The target system must be an eligible system in the service group's `SystemList` and must be capable of bringing the group online. A switch is an administrative relocation, unlike an automatic failover that occurs after VCS detects a fault. See the VCS command reference for `hagrp` command usage.

**QUESTION NO: 2**

Which two requirements must be meet to allow an administrator to configure LLT heartbeats over UDP? (Select two.)

**A.** the LLT private links must be on the same subnet

**B.** each NIC must have an IP address configured before configuring LLT

**C.** the maximum transmission unit (MTU) must be set to the lowest value supported by the NICs

**D.** each link must have a unique UDP port

**E.** broadcast must be enabled for the heartbeat subnet

**ANSWER: B D**

**Explanation:**

LLT-over-UDP transports Low Latency Transport heartbeat traffic through the host IP networking stack rather than directly through raw Ethernet private links. Therefore, *each NIC must have an IP address configured before configuring LLT*. LLT needs a usable IP endpoint on every interface that is declared as a UDP transport link, and those addresses must be configured and available before the LLT configuration is brought online.

In addition, *each link must have a unique UDP port*. The UDP port is part of the LLT link definition and identifies the local UDP socket used for that heartbeat path. Assigning a distinct port to each configured link prevents port-binding conflicts when multiple LLT UDP links are configured on the same system. The corresponding nodes must use the matching link definitions so that heartbeat packets are directed to the intended LLT transport endpoint.

These settings are made in the LLT configuration, commonly in `/etc/llttab`, before starting or restarting the VCS/InfoScale Availability stack. Veritas documents LLT configuration and UDP transport considerations in its InfoScale Availability administration documentation: Veritas InfoScale Availability documentation and the Veritas SORT documentation portal.

**QUESTION NO: 3**

A cluster has two high-priority LLT links and one low-priority link. The low-priority link is used for public network traffic. What is the state of the cluster if both high-priority links fail simultaneously?

**A.** All nodes remain in regular GAB membership, but are also in jeopardy membership.

**B.** The cluster is in a split-brain condition until at least one other LLT link is working.

**C.** The fencing race is initiated and service groups are failed over the winning nodes.

**D.** Service groups are unable to failover when resource faults occur, but can failover if a system fault occurs.

**ANSWER: A**

**Explanation:**

All nodes remain in regular GAB membership, but are also in jeopardy membership. LLT normally communicates over high-priority links, reserving low-priority links—often configured on public-network adapters—for use only when every high-priority link is unavailable. When both high-priority links fail, LLT can still exchange its heartbeat traffic through the configured low-priority link. Therefore, the nodes retain communication and GAB can maintain its regular membership.

Because the cluster has lost all of its preferred private heartbeat paths, LLT reports a jeopardy condition. GAB establishes jeopardy membership in addition to regular membership, allowing VCS to recognize that cluster communication is operating through its fallback path. This state is specifically intended to preserve cluster operation while signaling reduced network redundancy. The low-priority link is not treated as ordinary high-priority private-link connectivity, even though it prevents the complete loss of inter-node communication. Veritas documents LLT link priorities and the use of low-priority links as backup heartbeat links in the InfoScale Availability documentation; see the Veritas InfoScale Availability documentation and the LLT configuration and administration guidance.

**QUESTION NO: 4 - (SIMULATION)**

Match all five Veritas Cluster Server global cluster options on the left to the corresponding definitions on the right.

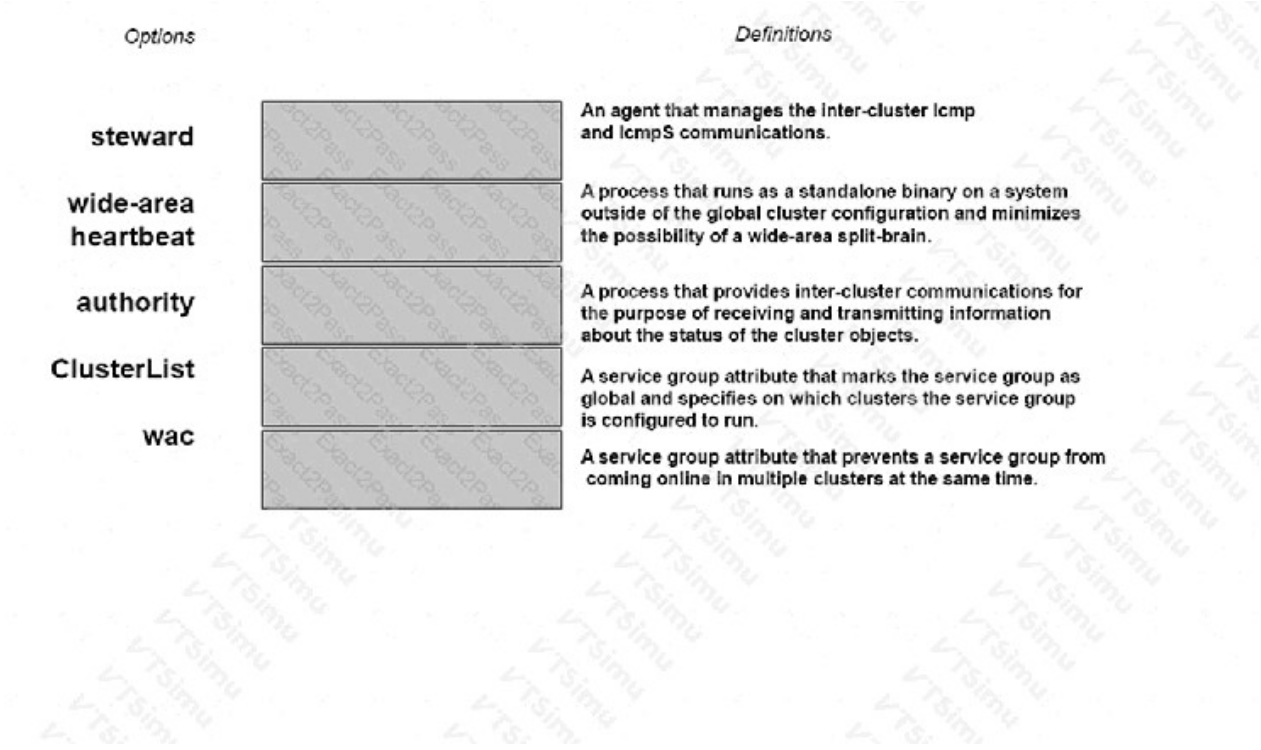


**ANSWER: See the explanation for the answer**

**Explanation:**

**Correct matches:**

The `ClusterList` and `Authority` attributes control global-service-group placement and exclusivity. The steward helps prevent split-brain, while wide-area heartbeat and WAC provide the inter-cluster communication functions.

## QUESTION NO: 5

What does the AdaptiveHA feature of Veritas Cluster Server (VCS) provide?

**A.** It uses the Priority attribute of service groups to take lower priority service groups offline when a higher priority service group fails over to the system.

**B.** It selects the biggest available target system to fail over an application when the service group attribute FailOverPolicy is set to BiggestAvailable.

**C.** It switches a service group with FailOverPolicy set to BiggestAvailable to less loaded systems in the cluster when the local system becomes overloaded.

**D.** It provides increased availability to the applications on the single node VCS cluster in virtual environments for planned maintenance events.

## ANSWER: D

**Explanation:**

It provides increased availability to the applications on the single node VCS cluster in virtual environments for planned maintenance events. AdaptiveHA is designed for a VCS deployment in which the protected application runs in a virtual machine and traditional clustering across multiple VCS nodes is not available. During a planned maintenance operation affecting the virtual machine's host, such as host maintenance initiated through the virtualization platform, AdaptiveHA coordinates with the virtualization environment to preserve application availability. VCS can take the appropriate action to stop and restart the application under control while the virtual machine is relocated or restarted, helping avoid an unnecessarily prolonged application outage. This extends VCS application monitoring and service-group management to the planned-maintenance workflow in virtualized environments, rather than merely reacting to an unexpected application or system fault. The feature is therefore specifically associated with improving availability for applications protected by a single-node VCS cluster in supported virtual environments. Veritas documents AdaptiveHA as part of the VCS high-availability capabilities and its integration with virtualized infrastructure. See the Veritas InfoScale Availability documentation and the Veritas AdaptiveHA technical information.

## QUESTION NO: 6

During a Veritas Cluster Server startup procedure, GAB fails and the administrator receives the following message:

vcs:11032 registration failed. Exiting

Which two files should the administrator check to ensure proper configuration? (Select two.)

**A.** main.cf

**B.** gabtab

**C.** inittab

**D.** vxfentab

**E.** llttab

## ANSWER: B E

**Explanation:**

The files to verify are **gabtab** and **llttab**. GAB depends on the Low Latency Transport (LLT) layer for cluster-node communications. The `llttab` file defines the LLT configuration, including the local node identity and the network links used

for private cluster communication. If LLT is not configured correctly, is unable to start, or cannot communicate across its configured links, GAB cannot form or join the cluster membership needed by VCS.

The `gabtab` file controls GAB startup configuration, including the command used to configure GAB and the expected cluster membership parameters. A mismatch in node-count expectations, an invalid GAB startup entry, or an incorrect configuration can prevent GAB from registering its required ports. When VCS starts, it must register with GAB; therefore, the reported registration failure warrants checking both the GAB configuration and its LLT dependency before attempting to restart the cluster services.

Veritas documents LLT and GAB as the foundational communications and membership components for InfoScale Availability clusters. Consult the Veritas InfoScale Availability documentation and the Veritas support documentation search for release-specific configuration and troubleshooting procedures.

## QUESTION NO: 7

Which utility should an administrator use to propagate changes made to the Solaris Zone configuration from one node to other cluster nodes?

**A.** hasync

**B.** hazoneconfigsync

**C.** haconfigdistribute

**D.** hamigrate

## ANSWER: B

**Explanation:**

`hazoneconfigsync` is the VCS utility designed specifically to synchronize Solaris Zone configuration changes across the nodes of an InfoScale Availability cluster. When an administrator modifies a zone configuration on one node—for example, by changing zone attributes, resource controls, or network-related settings—the cluster requires equivalent zone configuration information on the other potential service-group nodes. Running `hazoneconfigsync` propagates the relevant configuration so that the zone can be consistently managed, started, and failed over by VCS throughout the cluster.

This synchronization is important because VCS relies on a uniform zone definition when monitoring and controlling zone resources. Applying the change only on the local node can leave peer nodes unable to bring the application service group online correctly after a migration or failover. The command is part of the Solaris Zones support delivered with Veritas InfoScale Availability and should be used after making supported zone-configuration changes, following the operational procedures for the deployed InfoScale release. Veritas documentation for InfoScale Availability and its agent configuration guidance is available at Veritas InfoScale Availability documentation and the Veritas Support knowledge base.

## QUESTION NO: 8

An administrator performs an orderly shutdown of a healthy cluster by issuing `hastop -all`.

Which two actions occur during this operation? (Select two.)

**A.** Online service groups are taken offline.

**B.** HAD is stopped on all cluster systems.

**C.** All service groups are switched to the lowest numbered system.

**D.** Online service groups continue to run without VCS management.

**E.** LLT links are deleted from the cluster configuration.

## ANSWER: A B

**Explanation:**

**Online service groups are taken offline** is correct because an orderly cluster shutdown processes managed service groups before stopping the VCS engine. VCS takes the groups offline in dependency order so that applications and their supporting resources are stopped in a controlled manner.

**HAD is stopped on all cluster systems** is also correct because the `-all` option directs VCS to stop the High Availability Daemon throughout the cluster after the orderly shutdown processing completes.

The command does not migrate every service group to one node, and it does not leave normally managed online service groups running. Administrators should use forced local shutdown options only when they specifically intend to stop HAD without normal service-group processing. Veritas documents `hastop` operations in the VCS command reference.

## QUESTION NO: 9

What are two effects of a failed resource in the service group configuration below? (Select two.)

```
group websg {
      SystemList = { sys1 = 0, sys2 =1 }
      }

Apache httpd_Server {
      Critical = 0
      httpDir = "/apache/bin"
      HostName = vcssol1
      Port = 999
      SecondLevelMonitor = 1
      ConfigFile = "/apache/config/httpd.conf"
      }

Diskgroup Apache_dg {
      Critical = 0
      DiskGroup = apc1
      }

IP Apache_ip {
      Critical = 0
      Device = bde0
      Address = "11.123.99.168"
      NetMask = "255.255.254.0"
      }

Mount Apache_mnt {
      Critical = 0
Mount Apache_mnt {
      Critical = 0
      MountPoint = "/apache"
      BlockDevice = "/dev/vx/dsk/apc1/apcvol1"
      FSType = vxfs
      FsckOpt = "-y"
      }
```

**A.** Any resource fault leads to a faulted state for the resource.

**B.** Any resource fault leads to the service group failing over.

**C.** Any resource fault leads to the service group being taken offline.

**D.** Any resource fault leaves the service group partially online.

**E.** Any resource fault leaves the service group fully online.

**ANSWER: B C**

**Explanation:**

In the configuration shown, a resource fault is handled as a service-group failure condition and initiates the configured failover behavior. Therefore, *Any resource fault leads to the service group failing over.* During this process, VCS stops the service group on the node where the failure was detected and attempts to bring it online on the next eligible system in the service group's SystemList. This protects service availability by moving the workload away from the affected node.

The associated local action is that *Any resource fault leads to the service group being taken offline.* VCS takes the group offline on the current system as part of orderly failover processing, including taking dependent resources offline in the proper dependency order. After the local offline operation, VCS brings the group online on the selected target node. The combination of local offline processing and failover ensures the service is not left active in an inconsistent state while it is relocated.

Veritas describes service-group fault handling, critical-resource behavior, and failover processing in the InfoScale Availability documentation. Additional product documentation is available from the Veritas InfoScale Availability documentation portal.

## QUESTION NO: 10 - (DRAG DROP)

An administrator wants to enable administration of Veritas high availability from the vSphere Web Client in VMware environments. Click and drag the five (5) actions from the Action List on the left to the Configuration Process on the right in the order needed to accomplish the task.

**Select and Place:**

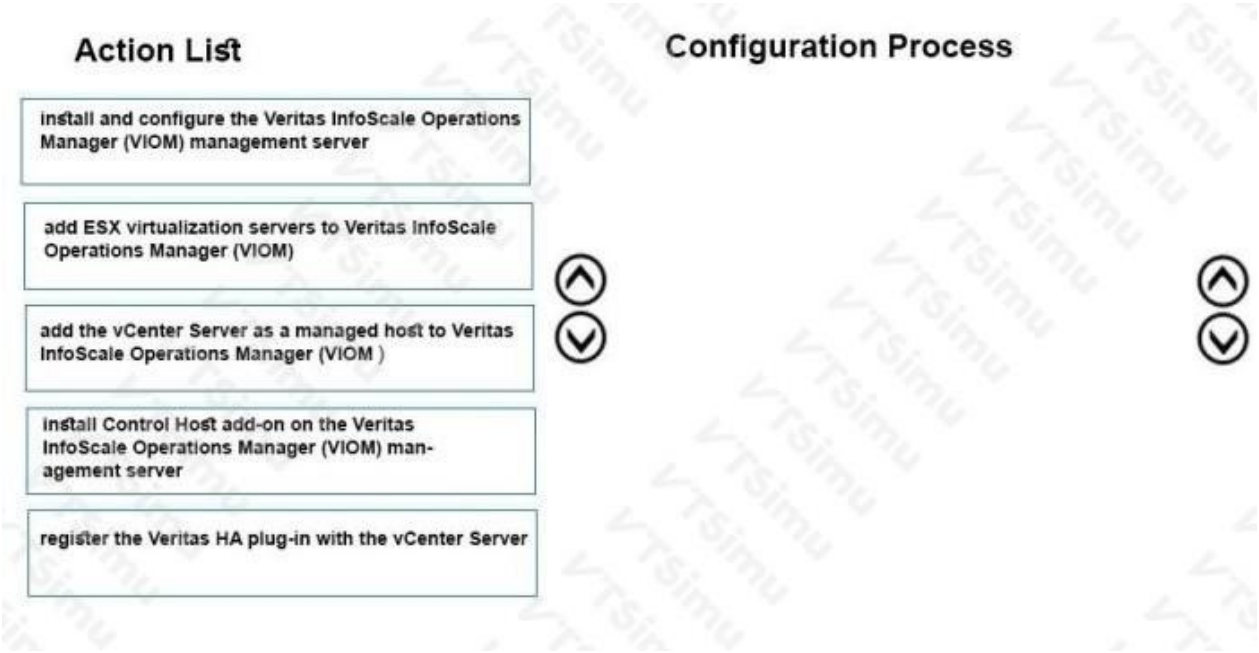


**ANSWER:**

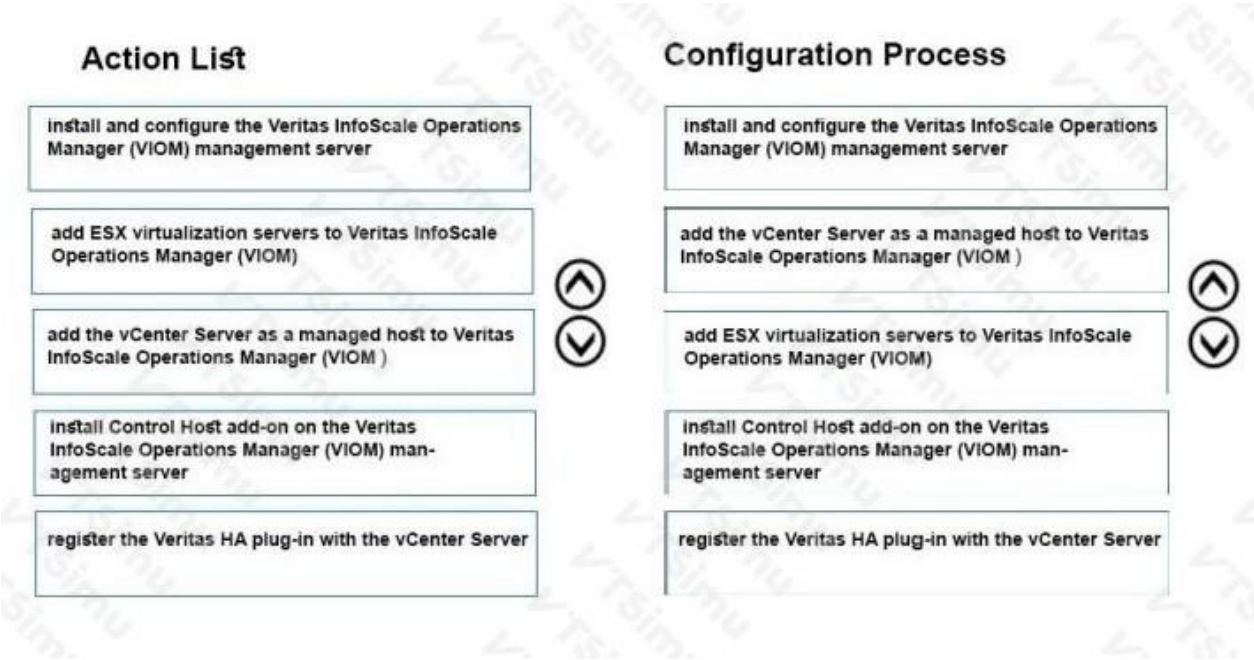


**Explanation:**

Administration of Veritas high availability from the vSphere Web Client is enabled by building the Operations Manager and VMware integration in dependency order. First, the Veritas InfoScale Operations Manager management server must be installed and configured. This establishes the central management service that stores the managed-host inventory and provides the services needed by the later integration components. Once that service is available, the vCenter Server is added as a managed host to Operations Manager. Operations Manager can then use the vCenter connection to discover and manage the VMware environment.

After vCenter is under management, the ESX virtualization servers are added to Veritas InfoScale Operations Manager. This associates the virtualization infrastructure with the management environment and makes the required host-level context available. The Control Host add-on is then installed on the Veritas InfoScale Operations Manager management-agent server. The add-on provides the integration functionality that supports the HA operations exposed through the VMware-facing interface. Finally, the Veritas HA plug-in is registered with the vCenter Server. Registering it at the end makes the plug-in available in the vSphere Web Client after the Operations Manager, vCenter, ESX hosts, and Control Host integration prerequisites have all been established.

This sequence ensures that every action has the service, inventory, and integration information it requires before the next action begins. Veritas documentation for InfoScale Availability and Operations Manager is available through the Veritas InfoScale documentation portal, and VMware documents vCenter plug-in administration in its vSphere documentation.

## QUESTION NO: 11

Which behavior can be specified using the PreferredFencingPolicy attribute?

**A.** The fencing driver gives preference to high-capacity networks.

**B.** The fencing driver gives preference to high-capacity systems.

**C.** The fencing driver gives preference to Oracle type resources.

**D.** The fencing driver gives preference to disk groups with solid state devices.

## ANSWER: A

**Explanation:**

The fencing driver gives preference to high-capacity networks. The `PreferredFencingPolicy` attribute is a cluster-level setting used by Veritas InfoScale Availability in fencing decisions involving a network partition. It enables an administrator to define a preference for the partition that has the greater network capacity. This preference helps VCS make a consistent and intentional fencing decision when connectivity between groups of cluster nodes is disrupted, supporting continued availability

for the partition with the preferred network characteristics while preserving the protection against simultaneous access to shared storage. The setting is relevant to the cluster's I/O-fencing and split-brain protection behavior; it is not a resource-selection or storage-device-performance setting. Administrators should configure fencing and its related cluster attributes carefully, because correct fencing configuration is essential to maintaining data integrity during failures and partitions. Veritas documents the VCS cluster configuration attributes and fencing administration procedures in its InfoScale Availability documentation. See the Veritas InfoScale Availability documentation and the Veritas knowledge base guidance for I/O fencing.