

IBM

Exam 000-332

High Availability for AIX - Technical Support and Administration -v2

Version: 8.0

[Total Questions: 131]

Question No : 1

PowerHA automatic error notification for shared data volume groups traps on which specific error?

- A. LVM_SA_WRTERR
- B. LVM_SA_PVMISS
- C. LVM_SA_STALEPP
- D. LVM_SA_QUORCLOSE

Answer: D

Question No : 2

Failure to plan for which element will have a negative impact on PowerHA cluster availability,when using virtual I/O resources on the cluster nodes?

- A. Planning for NPIV capable SAN switches
- B. Planning for Etherchannel capable switches
- C. Planning for update of Virtual I/O Servers
- D. Planning for at least 2 virtual Ethernet adapters for each cluster network and each node

Answer: C

Question No : 3

What must be verified before configuring GLVM?

- A. TCP/IP communication between sites is possible.
- B. A replacement IP address is configured for XD_IP network.
- C. The PowerHA cluster resource definition is completed on the primary site.
- D. Shared volume groups for Geographically Mirrored Volume Groups are configured.

Answer: A

Question No : 4

An administrator finds that a colleague has increased the size of a file system in an Enhanced Concurrent Volume Group using base AIX commands. How can the administrator make the cluster configuration consistent again?

- A. Run `chfs` on the other cluster nodes.
- B. Run Extended Verification and Synchronization.
- C. Run the `synclvodm` command on all cluster nodes.
- D. Run Synchronize a Shared Volume Group Definition.

Answer: D

Question No : 5

In which configuration would an AIO Cache LV be required?

- A. Cross-site LVM
- B. Synchronous GLVM
- C. Asynchronous GLVM
- D. When using mirror pools

Answer: D

Question No : 6

Given the following `cluster.mmddyyyy` log file output, what has happened?

```
Mar 29 06:09:22 EVENT START: network_down -1 net_ether_01
```

```
Mar 29 06:09:22 EVENT COMPLETED: network_down -1 net_ether_01 0
```

```
Mar 29 06:09:29 EVENT START: network_down_complete -1 net_ether_01
```

```
Mar 29 06:09:29 EVENT COMPLETED: network_down_complete -1 net_ether_01 0
```

- A. Non-redundant network switch has failed
- B. One interface in `net_ether_01` has failed
- C. The service IP address was moved to another adapter
- D. Some heartbeat packets have been lost on one node in a cluster

Answer: A

Question No : 7

After installing PowerHA packages, what is the next step when configuring the cluster?

- A. Configure cluster networks.
- B. Add repository disk to the cluster.
- C. Configure persistent IP addresses on all nodes.
- D. Add cluster boot addresses, service addresses and hostnames to /etc/hosts.

Answer: D

Question No : 8

In order to utilize PowerHA 7 SAN based communications in a VIO Server environment, what task must be performed on the AIX client nodes?

- A. Add a SANhb network into the cluster topology.
- B. Add a virtual Ethernet adapter and specify VLAN 3358.
- C. Configure/etc/cluster/ifrestrict to specify which HBAs should be excluded.
- D. Add another HBA to each VIOS server and add it as avSCSI adapter for this purpose.

Answer: B

Question No : 9

An administrator is configuring a PowerHA 6 cluster to use disk heartbeat as a non IP network. A shared volume group, dhbvg, has been configured using a single disk, hdiskIO, on NodeA and NodeB.

The administrator believes that the prerequisites have been met, but would like to test connectivity of the disk heartbeat network before adding the network into the cluster.

How can the connectivity of this disk heartbeat network be tested?

- A. On NodeA run the command /usr/bin/readvgda-V dhbvg -F status

On NodeB run the command `/usr/bin/putvgda-V dhbvg -a status=1`
B. On NodeA run the command `/usr/sbin/rsct/bin/dhb_transmit hdisk10`
On NodeB run the command `/usr/sbin/rsct/bin/dhb_receive hdisk10`
C. On NodeA run the command `/usr/sbin/rsct/bin/dhb_read-p hdisk10-r`
On NodeB run the command `/usr/sbin/rsct/bin/dhb_read-p hdisk10-t`
D. On NodeA run the command `"echo 1 | /usr/bin/dd of=/dev/rhdisk10 bs=4K skip=2 count=1"`
On NodeB run the command `/usr/bin/dd if=/dev/rhdisk10 bs=4K skip=2 count=1"`

Answer: C

Question No : 10

There is a two-node cluster with Node1 and Node2. An administrator changes filesystem size on Node1 by using the `chfs` command, and moves resource the group to Node2.

The administrator finds the change of filesystem size is recognized on Node2.

Why is the filesystem size change reflected on Node2?

- A.** The `gsclvmd` daemon automatically synchronizes filesystem changes.
- B.** The filesystem size change is recognized when filesystem is mounted on Node2.
- C.** A pre-event is defined to `get_disk_vg_fs` event to reflect filesystem size change.
- D.** The shared volume group is re-imported on Node2 by lazy update when moving the resource group.

Answer: D

Question No : 11

A 2-node cluster has three multi-node disk heartbeat (MNDHB) networks and data volume groups on separate disks. Application resource groups are online on Node1. What happens if Node1 loses access to two of the three MNDHB disks?

- A.** Node1 halts
- B.** Application resource groups move to Node2
- C.** Application resource groups remain on Node1
- D.** Application resource groups on Node1 go offline

Answer: C

Question No : 12

What is the command used to check the status of all interfaces and hearbeats in a PowerHA 7 cluster?

- A. cllsif-p
- B. cltopinfo-m
- C. lscluster-i
- D. lssrc-ls topsvcs

Answer: C

Question No : 13

An administrator has a 2-node cluster and has chosen the C-SPOC option to bring a resource group offline. The resource group was stable on the primary node prior taking resource offline, however it has gone into an error state. The administrator corrects the error and issues the clrunclmd to continue cluster processing.

What will happen to the resource group?

- A. The resource group goes offline.
- B. The resource group falls over to the remaining active node.
- C. The node shuts down and the resource group goes into temporary error state.
- D. The resource group remains in error state until services on all cluster nodes are stopped and restarted.

Answer: A

Question No : 14

After upgrading one node of a 2-node cluster to PowerHA 7, an administrator discovered that the wrong disk was chosen for the repository disk. What must be modified on both nodes to correct this error without starting over from the beginning?

- A. HACMPsircol
- B. HACMPcluster
- C. cluster.conf
- D. clmigcheck.txt

Answer: D

Question No : 15

When using PowerHA 7, a shared volume group (VG) must be configured as_____.

- A. Scalable VG
- B. Cluster aware VG
- C. Concurrent capable VG
- D. Enhanced concurrent VG

Answer: D

Question No : 16

An administrator adds a new network interface to each node in a 2-node PowerHA 7 cluster, and sets the IP addresses. When is the new network interface used for heartbeating?

- A. Immediately
- B. After refreshing clcomd
- C. After the cluster is synchronized
- D. When restarting PowerHA on both nodes

Answer: A

Question No : 17

What does the gsclvmd daemon do?

- A. Enhanced concurrent mode file system names are kept synchronized across nodes
- B. New enhanced concurrent mode volume group definitions are imported to all nodes

- C. Enhanced concurrent mode logical volume sizes are kept synchronized across nodes
- D. Quorum is maintained across cluster disks for enhanced concurrent mode volume groups

Answer: C

Question No : 18

A PowerHA 6 cluster has missed too many heartbeats during a large disk write operation, and has triggered the deadman switch. Which tunable can help ensure that HA Cluster Manager continues to run?

- A. I/O Pacing and syncd frequency
- B. I/O Pacing and asynchronous I/O
- C. Syncd frequency and asynchronous I/O
- D. Asynchronous I/O and disk queue_depth

Answer: A

Question No : 19

A company is determining how they can use DLPAR operations within PowerHA to maximize the usage of their system. Which DLPAR attribute can be controlled through PowerHA?

- A. The physical I/O adapters that will need to move to an application server
- B. The LPAR CPU Shared Pool weighting required to start an application server
- C. The minimum and desired processing units required to start an application server
- D. The minimum and maximum virtual processors required to start an application server

Answer: C

Question No : 20

An administrator has configured a 2-node PowerHA 6 cluster, Node1 and Node2, with no non-IP networks.

Sometime after configuring the cluster, the administrator noticed that Node2 appeared to have been shut down.

After restarting Node2 the administrator analyzed the cluster log files in an attempt to determine the cause of the problem.

The clstrmgr.debug log file contained the following information:

```
Mon Feb 9 13:59:32 announcementCb: GRPSVCS announcement code=512; exiting
Mon Feb 9 13:59:32 CHECK FOR FAILURE OF RSCT SUBSYSTEMS (topsvcs or grpsvcs)
Mon Feb 9 13:59:32 die: clstrmgr on Node2 is exiting with code 4
```

The error log contained the following information:

```
12081DC6 0209135909 P S haemd      SOFTWARE PROGRAM ERROR
9DEC29E1 0209135909 P O grpsvcs  Group Services daemon exit to merge domains
173C787F 0209135909 I S topsvcs Possible malfunction on local adapter
90D3329C 0209135809 P S topsvcs  NIM read/write error
```

Based on the above information, which of the following could explain why Node2 was shut down?

- A. The node was halted due to a DMS_TIMEOUT.
- B. Node isolation occurred causing a domain merge.
- C. Automatic error notification detected a problem and halted the node.
- D. The node was unable to fork a new topology services process due to memory limitations.

Answer: B

Question No : 21

When implementing PowerHA in a cross site LVM configuration, which logical volume option is required?

- A. Serialize IO=yes
- B. Mirror Write Consistency=on
- C. Scheduling Policy=sequential
- D. Allocation Policy=superstrict

Answer: A

Question No : 22

A cluster administrator wants to be sure that specific resource group will failover to the node with the most available processor resources. Which configuration supports the requirement?

- A. Use dynamic node priority based on PctTotalTimeIdle.
- B. Create DLPAR script for moving CPU resources on target LPAR.
- C. Put node with highest CPU power as home node for resource group.
- D. Use the uncapped processor resources for all nodes in the cluster.

Answer: A

Question No : 23

An administrator wants to shut down a node which has a resource group from the HMC. Which shutdown option results in a node shutdown without failover?

- A. Delayed
- B. Immediate
- C. Operating System
- D. Operating System Immediate

Answer: D

Question No : 24

An administrator notices that PowerHA has moved a resource group from one node to another. He looks on the vacated node, but doesn't see any problems reported in the application controller or application monitor logs.

Why are the logs missing the pertinent information?

- A. PowerHA has been restarted.
- B. The clcycle command has been run.
- C. The cluster ran a log_truncate event.
- D. Verify and Synchronize Cluster has been run

Answer: A