

**C4040-332.Examcollection.Premium.Exam.94q**

Number: C4040-332  
Passing Score: 800  
Time Limit: 120 min  
File Version: 35.2



**Exam Code: C4040-332**

**Exam Name: High Availability for AIX - Technical Support and Administration -v2**



**Actualanswers****QUESTION 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

**Correct Answer: D**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 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



**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 3**

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.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 4**

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

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:



#### **QUESTION 5**

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.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 6**

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/ignore restrict to specify which HBAs should be excluded.
- D. Add another HBA to each VIOS server and add it as avSCSI adapter for this purpose.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

### QUESTION 7

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 statusOn NodeB run the command /usr/bin/putvgda-V dhbvg -a status=1
- B. On NodeA run the command /usr/sbin/rsct/bin/dhb\_transmit hdisk10On 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-rOn 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"

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

### QUESTION 8

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.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 9**

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

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 10**

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 clcruncmd 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.

**Correct Answer:** A

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 11**

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

**Correct Answer:** D

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 12**

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

**Correct Answer:** D

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 13**

What does the gscsvmd 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

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 14**

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

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 15**

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

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 16**

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.

**Correct Answer: B****Section: (none)****Explanation****Explanation/Reference:**

Explanation:

**QUESTION 17**

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.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 18**

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



**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 19**

A PowerHA 6 cluster uses IPAT with aliasing and disk heartbeat. The administrator noticed the standby node crashes during a failover test. There are entries related to the Dead Man Switch (DMS) timeout in the AIX error log.

Which action is recommended to reduce the risk of DMS timeouts?

- A. Set Failure detection rate of Ethernet to normal.
- B. Set Failure detection rate of diskhb network module to slow.

- C. Create a RS232 heartbeat and set failure detection rate to delay.
- D. Dedicate a set of disks for disk heartbeat and set failure detection rate to delay.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 20**

To provide a secure cluster environment, an administrator will use VPN tunnels for communication between the nodes in a cluster. What must be defined on the nodes for communication across the VPN tunnel?

- A. Kerberos
- B. SSL shared keys
- C. Persistent IP Labels
- D. Public key encryption

**Correct Answer:** C

**Section:** (none)

**Explanation**



**Explanation/Reference:**

Explanation:

#### **QUESTION 21**

Consider the following PowerHA DLPAR configuration for a cluster, and LPAR profile on the HMC for a standby node in that cluster:

Change/Show Dynamic LPAR and CoD Resources for Applications

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

```
[TOP]                                     [Entry Fields]
* Application Controller Name             db2
* Minimum number of processing units     [ 0.50]
* Desired number of processing units     [ 0.90]
* Minimum number of CPUs                 [1]
* Desired number of CPUs                 [1]
```

General	<b>Processors</b>	Memory	I/O	Virtual Adapters	Power Controlling	Settings
---------	-------------------	--------	-----	------------------	-------------------	----------

Detailed below are the current processing settings for this partition profile.

**Processing mode**

Dedicated  
 Shared

**Processing units**

Total managed system processing units : 48.00  
Minimum processing units :   
Desired processing units :   
Maximum processing units :   
Shared processor pool:

**Virtual processors**

Minimum processing units required for each virtual processor : 0.10  
Minimum virtual processors :   
Desired virtual processors :   
Maximum virtual processors :

How should the standby node profile be reconfigured so that an application runs with the minimum hardware resources?

- A. Minimum and Desired processing units = 0.1 and Maximum processing units = 0.6
- B. Minimum and Desired processing units = 0.2 and Maximum processing units = 0.5
- C. Minimum processing units = 0.6 and both Desired and Maximum processing units = 1.5
- D. Minimum processing units = 0.5 and both Desired and Maximum processing units = 0.9

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 22**

After migrating to PowerHA 7 and restarting the cluster for the first time, the following error is reported:

Starting Cluster Services on node: Jessica

This may take a few minutes. Please wait...

Jessica: cl\_rsh had exit code = 13, see cspoc.log and/or clcomd.log for more information

What is the most common cause of this error?

- A. Last C-SPOC operation failed.
- B. CAA cluster services are not running.
- C. /etc/cluster/rhosts is incorrectly configured.
- D. The hacmp group does not have proper permissions or ID number.

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 23**

An administrator wants to assign a tape drive to a resource group for ease of management and control.

The tape drive is already attached to two cluster nodes. Node1 and Node2.

On Node1, the tape drive has a logical device name of rmt1. Node1 also has a tape drive, rmt0, that is not controlled by the cluster.

On Node2 the tape drive has a logical device name of rmt0. Which of the following must the administrator do before adding the tape drives to the cluster resource group?

- A. Remove the device files for rmt1 on Node1.
- B. Assign a device alias of rmt1 to the device on Node2.
- C. Redefine the logical device name rmt0 as rmt1 on Node2.
- D. Add tape rmt0 on Node1 and assign control to the cluster.

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 24**

Company policy requires logical volumes to be mirrored across separate disks and controllers for resilience. An administrator needs to create a 2-node cluster with the existing partitions, LPAR1 and LPAR2. The cluster will host an application with data on a shared volume group.

LPAR1 and LPAR2 are currently hosted on a single Power 770. Each partition is assigned two integrated SAS controllers, each of which has internal hard disk drives (HDDs) that are currently used for rootvg.

Which group of actions will create a shared volume group that complies with company storage policy?

- A. Add HDDs to LPAR1 existing SAS controller disk bays. Create shared volume group on new HDDs, ensuring the logical volume copies are across separate HDDs. Import shared volume on LPAR2.
- B. Assign LPAR1 & LPAR2 two fibre adapters each. Zone disk from two independent storage servers to each LPAR fibre adapter. Create shared volume group on new HDDs, ensuring the logical volume copies are across separate HDDs & storage servers.
- C. Assign LPAR1 SAS controllers to dual VIO Servers. Add HDDs to SAS controller disk bays. Create virtual SCSI device for new disks and assign to both LPAR1 & LPAR2. Create shared volume group, ensuring the logical volume copies are across separate HDDs and virtual SCSI devices.
- D. Assign to LPAR1 as "desired" the existing LPAR2 SAS controller. Add HDDs to LPAR2 existing SAS controller disk bays. Create shared volume group on new HDDs, ensuring the logical volume copies are across separate HDDs. Dynamically assign LPAR2 SAS controllers to LPAR1 and update profile with adapters as "desired." Import shared volume group on LPAR1.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 25**

An administrator unintentionally changed the name of an online resource group on one node of a 2- node cluster. The cluster has not been synchronized since the change was made.

How can the cluster configuration revert to the previous configuration without affecting the online resource group?

- A. Stop the cluster and restore from cluster snapshot.
- B. Restore HACMP configuration database from active configuration.
- C. Synchronize the cluster from the second node which still has the original name.
- D. Synchronize cluster from the first node, change the resource group name back, synchronize cluster again

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 26**

A 2-node cluster configuration has an application server named appl. The appl server has custom application monitors named appmon\_a and appmon\_b.

The setting of appmon\_a and appmon\_b follows:

appmon\_a: monitoring process=A, monitor interval=20, stabilization interval=60, restart count=1  
appmon\_b: monitoring process=B, monitor interval=10, stabilization interval=60, restart count=0  
If process B fails 30 seconds after process A fails and successfully restarts, what is the most amount of time it will take for a failover event to begin?

- A. about 10 seconds
- B. about 30 seconds
- C. about 60 seconds
- D. about 80 seconds

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 27**

Which application dependency can cause problems if the application is installed on a shared disk?

- A. An application which depends on CPU ID
- B. An application which depends on the service IP label
- C. An application which needs a specific TCP port number
- D. An application which needs a specific path name to data

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 28**

9.12.4.11 9.12.4.13

!REQD!ALL 100.12.7.9

!REQD!ALL 100.12.7.10

Given the netmon.cf file above, a single-adapter configuration must recognize a network\_down event if a node loses access to\_\_\_\_\_:

- A. 9.124.11 OR 9.12.4.13
- B. 9.124.11 AND 9.124.13
- C. 100.127.9OR 100.12 7.10
- D. 100.12.7.9 AND 100.12.7.10

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 29**

A customer has completed installing PowerHA 7 filesets. Which task, or tasks, should be completed next?

- A. Configure cluster topology
- B. Configure cluster resources
- C. Reboot each PowerHA cluster node
- D. Verify and synchronize the cluster

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 30**

An active, active (mutual takeover) 2-node PowerHA cluster is currently up and stable. Both nodes are located in separate buildings, one kilometer apart. Node2 in SiteB suddenly powers off. On reboot of the failed node, the cluster administrator checks the error log and discovers the node had been halted. The network team confirms there was a temporary loss of the IP network between sites.

## Error log

```
BA431EB7      0512150309 P S SRC      SOFTWARE PROGRAM ERROR
BC3BE5A3      0512150309 P S SRC      SOFTWARE PROGRAM ERROR
CAD234BE      0512150309 U H LVDD     QUORUM LOST, VOLUME GROUP CLOSING
DB14100E      0512150309 P U LIBLVM   Group Services detected a failure
12081DC6      0512150309 P S haemd    SOFTWARE PROGRAM ERROR
9DEC29E1      0512150309 P O grpsvcs  Group Services daemon exit to merge doma
AB59ABFF      0512150109 U U LIBLVM   Remote node Concurrent Volume Group fail
AB59ABFF      0512150109 U U LIBLVM   Remote node Concurrent Volume Group fail
173C787F      0512150109 I S topsvcs  Possible malfunction on local adapter
173C787F      0512150109 I S topsvcs  Possible malfunction on local adapter
```

Why did Node2 halt?

- A. Node2 was automatically set to halt on loss of the IP network.
- B. On loss of the IP network the grpsvcs daemon crashed, causing Node2 to halt.
- C. On loss of the IP network, Node1 sent a halt signal to Node2 via the disk heartbeat network.
- D. The cluster became partitioned. On recovery of the IP network, Node1 sent a halt signal to Node2.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 31**

What must be done before defining a pre-event command for a predefined cluster event?

- A. Populate the PRE\_EVENT\_CMD variable in the predefined event script
- B. Define a custom event command using the SMIT "Add a Custom Event" dialog
- C. Populate the cluster events directory on each node with the script/executable
- D. Specify the absolute path to the command on the "Change/Show Pre-defined PowerHA Event" SMIT dialog

**Correct Answer:** B

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 32**

An administrator performed a DARE operation in their cluster configuration. During this operation one of the cluster nodes halted. The node halt was an isolated problem, however, as a result the DARE operation did not complete.

When the failed node was re-integrated to the cluster, the following error was received when trying to perform the DARE operation again:

cl dare: A lock for a Dynamic Reconfiguration event has been detected.

Which action will resolve the problem?

- A. Remove the file `cl dare_lock` in `/usr/es/sbin/cluster/etc/objrepos/staging` and re-try the DARE operation.
- B. Use the PowerHA problem determination function to "Release Locks Set By Dynamic Reconfiguration" and re-try the DARE operation
- C. Verify and synchronize the cluster with "Automatically correct errors found during verification" set to "yes" which will complete the DARE operation.
- D. Perform the PowerHA problem determination function to "Restore HACMP Configuration Database from Active Configuration" which will complete the DARE operation.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 33**

While creating a new file collection named "custom" an administrator gets an "invalid name" error. What is the most likely reason for this error?

- A. The description field was left blank.
- B. The name "custom" is a reserved word.
- C. The name must have at least 8 characters.
- D. The administrator did not include required numerals.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 34**

What type of IP address configuration is required to enable Cluster Aware AIX (CAA) monitoring features?

- A. An IPV6 address
- B. A dynamic IP address
- C. A multicast IP address
- D. A reserved private address in the 10.0.0.0 to 10.255.255.255 address range

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 35**

Which are names of default file collections in PowerHA 7?

- A. Cluster\_configOS\_Config
- B. SystemjilesPowerHA\_Files
- C. HACMP\_FilesConfiguration\_Files
- D. /etc/hosts/etc/inet.d/etc/services/etc/passwd

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 36**

What are the two types of application monitors that can be created?

- A. Process Monitors and Custom Monitors
- B. Predefined Monitors and Custom Monitors
- C. System Monitors and User Defined Monitors
- D. Process Monitors and User Defined Monitors

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

### QUESTION 37

What will allow an administrator to mirror shared volume groups between two sites that are only linked by TCP/IP?

- A. GLVM Disk
- B. MetroMirror Disk
- C. PowerHA 7 Network Share Disk
- D. Cluster Aware AIX (CAA) Share Disk

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:



### QUESTION 38

What cluster service settings can be defined in SMIT on one node of a cluster and automatically updated on all other nodes?

- A. BROADCAST message at startup
- B. Start HACMP at system restart
- C. Verify Cluster Prior to Startup
- D. Startup Cluster Information Daemon

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

### QUESTION 39

An administrator wants to configure a new 3-node PowerHA 7 cluster for 3 instances of an application. Each node host an instance under normal production mode. The administrator wants a node failover to be based on which cluster node has the most free real memory available.

Which standard or user-defined dynamic node priority (DNP) attribute can be used to accomplish this?

- A. cl\_most\_free\_mem
- B. cl\_highest\_freespace
- C. cl\_highest\_udscript\_rc
- D. cl\_lowest\_nonzero\_udscript\_rc

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### QUESTION 40

Which command will show a summary of application servers in the cluster?

- A. clstat
- B. cldisp
- C. cldump
- D. clcheck\_server



**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### QUESTION 41

An administrator is planning to upgrade from HACMP 5.5 and AIX 5.3 to PowerHA 7 and AIX 7. After migration installing AIX, what must be done first?

- A. Install CAA fileset.
- B. Migrate to PowerHA 7.
- C. Run the clmigcheck program.
- D. Create shared VG for use by CAA.

**Correct Answer:** B

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 42**

An administrator added a filesystem on a shared volume group using C-SPOC. The cluster is active and the resource group is on line. What is the effect of the administrator action?

- A. The filesystem is available immediately.
- B. The cluster must be verified and synchronized.
- C. Cluster services must be restarted on all nodes.
- D. The filesystem must be mounted on a node with an online resource group.

**Correct Answer:** A

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 43**

An administrator is gathering requirements to plan a highly available environment between two buildings on the same campus. The environment must allow for a complete failure of all equipment in either building. Workloads must be able to be moved between the buildings during live processing, as well as during a failure.

Given these requirements, what must be included in the solution?

- A. Redundant VIO Server configuration
- B. Redundant serial networks for heartbeating
- C. Redundant network adapters for PowerHA networks
- D. Redundant storage servers with LVM mirroring between them

**Correct Answer:** D

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 44**

An administrator has a 4-node PowerHA cluster which was successfully verified and synchronized on Sunday.

On Wednesday the cluster was changed by a different administrator.

The cluster must be returned to the Sunday configuration and resynchronized.

How would an administrator restore the cluster to the original configuration when the changes made on Wednesday are not known?

- A. Restore the most recent mksysb backup.
- B. Restore the manual cluster snapshot which is 3 months old.
- C. Rebuild the cluster configuration manually, then verify and synchronize.
- D. Restore from the cluster snapshot automatically generated at the last synchronize.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### QUESTION 45

A PowerHA 7 cluster is running on two fully virtualized LPARs. An administrator simulates a network failure on the cluster by running the "ifconfig enX down" command on all network interfaces on one node. The resource group moved, and clstat shows the cluster as stable on both nodes.

What prevented the resource group from being active on both nodes?

- A. The nodes are able to connect to each other via vsa interface.
- B. The nodes are able to connect to each other via vscsi interface.
- C. The nodes are able to connect to each other via dpcom interface.
- D. The nodes are able to connect to each other via diskhb interface.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### QUESTION 46

In which default file can an administrator specify special options for nfs exported filesystems in PowerHA?

- A. /etc/exports
- B. /etc/filesystems
- C. /usr/es/sbin/cluster/etc/exports
- D. /usr/es/sbin/cluster/nfs\_exports

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 47**

The addition of a service IP address to a resource group fails to pass verification and synchronization.

What is the most likely reason?

- A. IPAT via Replacement is being used and the service IP address is in the same subnet as boot IP addresses.
- B. A route has not been defined for the service IP address on one or more of the nodes in the resource group.
- C. The service IP address is incorrectly defined or is missing in/etc/hosts on at least one node in the resource group.
- D. Service IP address distribution policy is set to Anti-Collocation and each interface is already hosting a service IP address.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 48**

An administrator plans to migrate a PowerHA 6 cluster with AIX 7 nodes to PowerHA 7. Which migration method should be used to minimize downtime?

- A. Offline migration
- B. Rolling migration
- C. Snapshot migration
- D. Non-disruptive migration

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 49**

The service IPs VLAN is the only routable network in a 2-node HA cluster with multiple network interfaces.

What must done to enable administrative tasks to be performed remotely on the cluster nodes?

- A. Associate a permanent Service IP address for each HA node.
- B. Configure the boot IP Address on the same VLAN as the service IP Address.
- C. Configure a persistent IP Address for each node on the service IP network.
- D. Create a route on the switch and translate the Service IP address to the boot IP address using NAT.

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 50**

An administrator wants to run the cluster test tool automated procedure. Part of this will perform a test to kill the clstrmgrES subsystem on a node, causing that node to halt.

How can the administrator avoid manual intervention to restart the node?

- A. Set the environment variable "CL\_KILL\_CLSTRMGR=FALSE" and make sure clexit.rc is executable before running the test
- B. Run the cl\_testtool\_auto command with the option "-K n" to specify that the clstrmgrES subsystem should not be killed
- C. Edit the file /usr/es/sbin/cluster/cl\_testtool\_auto\_cluster\_kill" replace the action CLSTRMGR\_KILL with CLSTRMGR\_SOFT\_KILL
- D. Modify and make executable /etc/cluster/hacmp.term" to change the action taken by the system upon an abnormal termination of clstrmgrES

**Correct Answer: D**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 51**

In multi-node disk heartbeat configurations running PowerHA 6, what is the recommended minimum number of LUNS to be shared between the nodes?

- A. 1
- B. 2
- C. 3
- D. 4

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 52**

After an administrator populates the hosts and rhosts files on the new and existing nodes, what are the high level steps to add a new node to an existing PowerHA 7 cluster?

- A. Add the existing cluster name on new node PowerHA SMIT.Configure the network topology on the new node PowerHA SMIT.Synchronize the cluster.Start Cluster Services on the new node.
- B. Take a snapshot from the existing cluster.Restore the snapshot on the new node.Update the node name on the cluster configuration.Synchronize the cluster.Start Cluster Services on the new node
- C. Add the node in to the cluster from one of existing cluster node.Configure the network topology for the new node from an existing node.Synchronize the cluster.Start Cluster Services on the new node.
- D. Shutdown the existing cluster.Use a snapshot to perform offline migration on the new node.Update the node name and network configuration on the new node.Synchronize the cluster.Start Cluster Services on the new node

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 53**

What does the following notification from node Jessica mean?

EVENT COMPLETED: server\_down Jessica 12 testapp 0

- A. Jessica has failed
- B. Jessica was taken offline

- C. An application on Jessica has failed
- D. Jessica is reporting that another node is down

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 54**

During a rolling migration to PowerHA 7, when is the Cluster Aware AIX (CAA) cluster created?

- A. After upgrading AIX on last node
- B. After running clmigcheck on the last node
- C. After starting cluster services on the last node
- D. After the first synchronization of the upgraded cluster

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:



#### **QUESTION 55**

Complete the following statement:

When planning to create PowerHA 7 nodes on Virtual I/O client partitions \_\_\_\_\_.

- A. Shared volume groups using virtual SCSI must be non-concurrent.
- B. Disks can be shared between client partitions using vSCSI and direct access.
- C. C-SPOC must be used when configuring or changing volume groups on shared disks.
- D. "reserve\_policy" attribute of shared disks through the VIO Servers must be set to "no\_reserve"

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 56**

In a PowerHA 7 cluster, an automatic error notification is created for detecting the loss of which of the following?

- A. roofrg
- B. NFS Mounts
- C. Default gateway
- D. Service IP address

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 57**

An administrator is planning a PowerHA implementation for a database and application. The application is dependent on the database being available before the application is started.

If the database is moved from one node to another, the application will need to be re-started. The database and application will be in different resource groups.

How can this be achieved within PowerHA?

- A. The application can be controlled from a post event to the rg\_move\_finish event.
- B. The application resource group can be configured as a child of the database resource group.
- C. The application resource group can be configured as the parent of the database resource group.
- D. A sequential processing order for the resource groups can be set, specifying the database and then the application

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 58**

Which type of heartbeating is valid in PowerHA 7?

- A. FDDI

- B. RS232
- C. Target mode SSA
- D. SAN-based communication

**Correct Answer:** D

**Section:** (none)


**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 59**

Based on the following output from a PowerHA 7 cluster node, what else can be done to provide higher availability?



```
lspv
hdisk2      00cee73e15f559b9      rootvg      active
hdisk3      00cee73ec9ebcc3e      None
hdisk4      00cee73ee4df96d3      datavg
hdisk5      00cee73ee4dfad5c      datavg
hdisk6      00cee73ee4dfc4e2      None
hdisk7      00cee73ee4dfe0d0      caavg_private  active

lsdev -Cc disk
hdisk2 Available 0A-08-01-5,0 16 Bit LVD SCSI Disk Drive
hdisk3 Available 0A-08-01-8,0 16 Bit LVD SCSI Disk Drive
hdisk4 Available 0B-08-01      MPIIO Other DS4K Array Disk
hdisk5 Available 0B-08-01      MPIIO Other DS4K Array Disk
hdisk6 Available 0C-08-01      MPIIO Other DS4K Array Disk
hdisk7 Available 0C-08-01      MPIIO Other DS4K Array Disk
```

- A. Enable fast failure detection.
- B. Mirror rootvg across local disk.
- C. Mirror caavg\_private across shared disk.

D. Change non mirrored volume group disks to hot spare using the chpv command

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 60**

A shared volume group is added to a resource group and verification of the change produces an error. What is the most likely reason?

- A. The auto vary on flag is set to "yes".
- B. The volume group is varied off on all of the nodes.
- C. The volume group major number is different on some nodes in the cluster.
- D. The names of one or more disks in the volume group do not match on all the nodes in the resource group.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 61**

Given the following output, why is node Jessica down?



```
#>/cllsif
Adapter                Type          Network      Net Type    Attribute  Node      IP Address
jessica_hdisk4_01     service      net_diskhb_01 diskhb      serial    jessica   /dev/hdisk4
jessica                boot         net_ether_01 ether       public    jessica   19.19.51.193
ha_svc                 service      net_ether_01 ether       public    jessica   192.168.100.10
jess_boot1            boot         net_ether_02 ether       public    jessica   192.168.100.1
peyton_hdisk4_01     service      net_diskhb_01 diskhb      serial    peyton    /dev/hdisk4
peyton                boot         net_ether_01 ether       public    peyton    19.19.51.194
ha_svc                 service      net_ether_01 ether       public    peyton    192.168.100.10
peyton_boot1         boot         net_ether_02 ether       public    peyton    192.168.100.2

#>lspv
hdisk0                00cee73e15f559b9      rootvg      active
hdisk1                00cee73ec9ebcc3e      altinst_rootvg
hdisk2                00cee73ee4df96d3      None
hdisk3                00cee73ee4dfad5c      None
hdisk4                00cee73ee4dfc4e2      datavg

#>clshowres
Resource Group Name    certtestrg
Participating Node Name(s) peyton jessica
Startup Policy         Online On First Available Node
Failover Policy        Fallover To Next Priority Node In The Li
Fallback Policy        Never Fallback
Service IP Label       ha_svc
Volume Groups          datavg
```

```
Mar 29 07:36:34 Peyton user:notice HACMP for AIX: EVENT START: network_down Jessica net_diskhb_01
Mar 29 07:36:34 Peyton user:notice HACMP for AIX: EVENT COMPLETED: network_down Jessica net_diskhb_01
Mar 29 07:36:34 Peyton user:notice HACMP for AIX: EVENT START: network_down_complete Jessica net_diskhb_01
Mar 29 07:36:34 Peyton user:notice HACMP for AIX: EVENT COMPLETED: network_down_complete Jessica
net_diskhb_01 0
Mar 29 07:36:37 Peyton user:notice HACMP for AIX: EVENT START: rg_move_release Jessica 1
Mar 29 07:36:37 Peyton user:notice HACMP for AIX: EVENT START: rg_move Jessica 1 RELEASE
Mar 29 07:36:37 Peyton user:notice HACMP for AIX: EVENT COMPLETED: rg_move Jessica 1 RELEASE 0
Mar 29 07:36:37 Peyton user:notice HACMP for AIX: EVENT COMPLETED: rg_move_release Jessica 1 0
Mar 29 07:42:37 Peyton local0:crit clstrmgrES[13893838]: Thu Mar 29 07:42:37 config_too_long: Timer type is 0
Mar 29 07:42:37 Peyton local0:crit clstrmgrES[13893838]: Thu Mar 29 07:42:37 config_too_long: Submitted event
config_too_long, PID=5177366
Mar 29 07:42:37 Peyton user:notice HACMP for AIX: EVENT START: config_too_long 360 TE_RG_MOVE
```

- A. Failover happened because disk heartbeat network was down on node Jessica.
- B. An rg\_move event happened and node Peyton has failed to acquire resource group.
- C. Node Jessica lost access to shared volume groups and releasing resource group has failed.
- D. Selective failover on network loss has happened and resource group was moved to node Peyton.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### QUESTION 62

An administrator tries to replace a hot-plug Ethernet adapter using the C-CPOC > Communication Interfaces SMIT menu. When prompted, the administrator selects the incorrect adapter and does not wish to continue. When prompted to confirm adapter replacement, he selects "no".

What is the first action that the administrator must take to restore cluster to the pre-maintenance condition?

- A. Migrate the resource groups using the adapter back to the node.
- B. Run the drslot command to take the PCI slot out of the "removed" state.
- C. Run Cluster Verification and Synchronization to reset the adapter state.

D. Run the `cfgmgr` command to re-detect the adapter and create the device special files

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### QUESTION 63

A PowerHA cluster is being planned with the following characteristics:

Nodes = 3

Networks = 2

Adapters per network = 3

IPAT via Aliasing = Yes

Service IP Address = 1

What is the minimum number of subnets that should be requested to ensure the highest availability?

A. 6

B. 7

C. 9

D. 12



**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### QUESTION 64

What will allow a customer with PowerHA 6 to enable fast disk takeover in their resource group named `rgA`?

A. Use C-SPOC to enable fast disk takeover for `rgA`.

B. Ensure that `rgA` has an Online on All Available Nodes startup policy.

C. Ensure that all volume groups participating in `rgA` are enhanced concurrent capable.

D. Use lazy update to convert participating volume groups in `rgA` to use fast disk takeover.

**Correct Answer:** C

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 65**

The `/usr/es/sbin/cluster/netmon.cf` file contains the following lines:

```
IREQD host1.ibm 100.12.7.9  
IREQD host1.ibm host4.ibm
```

What is the effect of this configuration?

- A. The node will only be considered "up" if it can ping at least one address on each line.
- B. The interface will only be considered "up" if it can ping 100.12.7.9 OR the address to which host4.ibm resolves.
- C. The interface will only be considered "up" if it can ping 100.12.7.9 AND the address to which host4.ibm resolves.
- D. The node will only be considered "up" if it can ping all addresses on the first line AND all addresses on the second line

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 66**

An administrator is configuring a 4-node cluster to host 3 instances of their production application in an N+1 failover design. Each instance will be in a separate resource group, and all 4 nodes will be participating nodes in each resource group.

Under normal circumstances each node only hosts one instance of the application. Whichever are the first three nodes to join the cluster become production, and the fourth node will be the standby.

How must the resource group startup policy be configured to accomplish this?

- A. Online On Home Node Only
- B. Online On All Available Nodes
- C. Online On First Available Node
- D. Online Using Node Distribution Policy

**Correct Answer:** D

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 67**

An administrator is migrating from PowerHA 5.4.1 to PowerHA 6.1. Which migration option creates the least amount of downtime for the clustered application?

- A. Rolling
- B. Snapshot
- C. Concurrent
- D. Non-disruptive

**Correct Answer:** D

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 68**

A company plans to deploy a 2-node PowerHA cluster using IP address takeover via IP aliasing. The HA nodes are POWER7 LPARs supported by VIO Servers. The nodes are located on different sites. Each site has its own network.

What must the company implement in order to use IP address takeover via IP aliasing?

- A. Create a common VLAN for both nodes by the network team
- B. Connect the switches on both sites using Inter Switch Link via DWDM
- C. Create a Shared Ethernet Adapter with the same VLAN ID on the all VIO Servers
- D. Create an EtherChannel adapter on both nodes that connects to the same switches

**Correct Answer:** A

**Section:** (none)

**Explanation****Explanation/Reference:**

Explanation:

**QUESTION 69**

As a part of PowerHA customization an administrator wants to change the location of hacmp.out log file. What is the correct way to accomplish this?

- A. Use clogdir PowerHA command.
- B. Use Extended Cluster Service Settings SMIT menu.
- C. Use Change/Show a Cluster Log Directory SMIT menu.
- D. Use text editor to edit /etc/objrepos/HACMPlogs file.

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 70**

An administrator has installed and configured AIX and PowerHA 7 on multipath vSCSI disks. There are two VIO Servers, and both are accessing the same SAN storage system. The administrator has modified setting for the rootvg system event in the cluster. The cluster is synchronized and stable.

How can the administrator test the rootvg system event to ensure the change behaves as expected?

- A. Delete the rootvg disk devices using the rmdev command.
- B. On both VIO Servers .remove all rootvg disk mappings for one node.
- C. Perform a quick system shutdown using the halt -q command on one node.
- D. Use the HMC to halt one of the VIO Servers which supports the primary MPIO path

**Correct Answer: B**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 71**

Which utility can be used while a rolling migration is in progress?

- A. Resource Groups -> Add a Resource Group
- B. System Management (C-SPOC)-> HACMP Services
- C. Extended Configuration -> Extended Verification and Synchronization
- D. Extended Configuration -> Snapshot Configuration -> Create a Snapshot of the Cluster Configuration

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

### QUESTION 72

A cluster consists of 2 base boot addresses per node. Each boot address can contact the other node in the cluster but is not externally routable. There is one service IP address in cluster which is externally routable.

The external network experienced a problem where the service IP could not be contacted by the users, but the base boot addresses were still able to send heartbeat packets between the cluster nodes. The resource groups did not failover to the other node where there was no network problem.

How can the administrator ensure that the resource groups failover when the service address is not contactable on a node?

- A. Create a custom application monitor.
- B. Create a process application monitor.
- C. Configure a persistent IP address on each node.
- D. Configure a post-event script to the fail\_interface event.

**Correct Answer:** A

**Section:** (none)

**Explanation**



**Explanation/Reference:**

Explanation:

### QUESTION 73

When designing a PowerHA 6 cluster that will use virtual networks, what action will ensure proper detection of an adapter or network failure?

- A. Configuring Etherchannel when using virtual Ethernet adapters
- B. Including hosts on multiple physical networks in the /etc/hosts file
- C. Including targets outside of the virtual network in the netmon.cf file
- D. Defining at least two virtual Ethernet adapters on each node on different VLANs

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 74**

When verifying a cluster, where will the clverify.log file be stored?

- A. On all cluster nodes
- B. On the cluster topology manager node
- C. On each cluster node with an active resource group
- D. On the cluster node where the operation was initiated

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 75**

An architect is designing a PowerHA cluster and plans to use a virtualized environment to maximize hardware utilization. Each cluster node is to be configured as a partition on a separate Power 770.

Which of the following options provides a supported configuration with the minimum single points of failure?

- A. Implement dual VIO Servers on each Power system. Each VIO Server will have multiple paths to the physical storage.
- B. Implement dual IVM partitions on each Power system. Each IVM partition will have multiple paths to the physical storage.
- C. Implement a single IVM partition on each Power system. Configure paths to physical storage active / passive mode on each cluster node with one path using a physical adapter and one path using a virtual adapter.
- D. Implement a single VIO server on each Power system. Configure paths to physical storage in active / passive mode on each cluster node with one path using a physical adapter and one path using a virtual adapter.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 76**

What is the main consideration when planning for GLVM with synchronous mirroring?

- A. Network bandwidth and latency.
- B. Disk architecture must be identical at each site.

- C. All GLVM volume groups must be created as enhanced concurrent.
- D. Capacity plan disk space as dynamic expansion is not possible at later date.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 77**

Both nodes of a 2-node PowerHA 7 cluster have been powered off for an extended time. An administrator attempted to start PowerHA on one of the nodes after system power on, and received the message:

clhare: A communication error prevents obtaining the VRMF from remote nodes.

The administrator is able to start PowerHA after restoring HACMP configuration database from the active configuration.

What is the most likely reason the error was received?

- A. The dcomdES and clstrmgr are not running on the other node.
- B. The PowerHA configuration is corrupted on the starting node.
- C. There is an unsynchronized configuration change in the cluster.
- D. The starting node is unable to verify the PowerHA Version Release Maintenance Fix information from the other node.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 78**

After installing and configuring PowerHA 7 in AIX 7, an administrator discovers that the clstat and cldump utilities do not work. Which corrective action will resolve this problem?

- A. Update to AIX v7.1 TL1 SP1.
- B. Convert from SNMPv3 to SNMPvt
- C. Verify and synchronize the cluster.
- D. Start the Cluster Aware AIX (CAA) daemons.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 79**

An administrator is concerned about disk I/O and the impact this may have on a PowerHA cluster. The cluster nodes have virtual disks presented from dual VIO Servers.

What action will improve PowerHA behavior during large disk writes?

- A. Tune syncd to 60 seconds
- B. Tune cluster node disks qdepth to 16
- C. Tune hcheck\_interval to 60 on all cluster node shared volumes
- D. Tune I/O pacing with an initial high water mark of 33 and low of 24

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:



#### **QUESTION 80**

An administrator created a new volume group on one local node using the mkvg command. The volume group is not known on the other nodes in the cluster, however the disks are available and have PVIDs defined.

How can this volume group be added into an online resource group, and the volume group information get updated on the remote node(s), with the least amount of effort?

- A. Set: Use forced varyon for volume groups, if necessary, to "true."
- B. Change the resource group option: Automatically Import Volume to "true."
- C. Synchronize the cluster using the auto-correct verification errors set to "yes."
- D. Allow the daily automatic cluster verification to auto-correct the volume group definition

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 81**

The administrator has written a script "correct\_problem.sh" that needs to be run if the "release\_vg\_fs" event does not complete successfully.

What should be done to ensure the correct\_problem.sh is called a maximum of 3 times if the event failed?

- A. Include a post event of correct\_problem.sh and a retry counter of 3
- B. Include a custom event of correct\_problem.sh and a restart limit of 3
- C. Include a retry command of correct\_problem.sh and an event counter of 3
- D. Include a recovery command of correct\_problem.sh and a recovery counter of 3

**Correct Answer: D**
**Section: (none)**
**Explanation**
**Explanation/Reference:**

Explanation:

**QUESTION 82**

Consider the following configuration:



```
#>cllsif -p
```

Adapter	Type	Network	Net Type	Attribute	Node	IP Address
ha_svc	service	net_ether_01	ether	public	jessica	19.168.100.100
jess_boot1	boot	net_ether_01	ether	public	jessica	19.168.100.1
ha_svc	service	net_ether_01	ether	public	jordan	19.168.100.100
jordan_boot1	boot	net_ether_01	ether	public	jordan	19.168.100.2

To use a service IP address as a source IP address when a PowerHA 7 cluster initiates a connection to an external server, an administrator should configure cluster topology by \_\_\_\_\_:

- A. Using IP Address Takeover via replacement.
- B. Making a route so that service IP address is used as a source IP address.
- C. Separating a network segment used for boot IP address from the service network segment.

D. Configuring "Heartbeating over IP Aliases" and separating the network segment used for heartbeat from the service network segment.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 83**

A 2-node PowerHA 6 cluster has Ethernet and RS232 heartbeat networks.

The RS232 connection fails What is the effect of losing the RS232 heartbeat connection?

- A. All network heartbeats are lost.
- B. The Ethernet heartbeat interval is reduced.
- C. Cluster services are stopped on affected nodes.
- D. An RS232 heartbeat network down event is recorded.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 84**

Consider the following Cluster Service Settings:

```
Cluster Startup Settings

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
Start PowerHA SystemMirror at system restart?    false      +
BROADCAST message at startup?                   true       +
Startup Cluster Information Daemon?              true       +
Verify Cluster Prior to Startup?                 true       +
```

An administrator changes cluster configurations and starts one node in the cluster. The node being started fails. What is the root cause of the failure?

- A. The administrator did not synchronize the cluster.
- B. PowerHA detected errors during startup cluster verification.
- C. Cluster synchronization failed during startup cluster verification.
- D. PowerHA failed to correct errors during startup cluster synchronization.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 85**

What can the cluster test tool be used to test?

- A. Pre or post events
- B. Failed volume group
- C. Rootvg system event
- D. Dynamic cluster reconfiguration

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:


**QUESTION 86**


Given the following information, how many processing units will be available when migtestrg is started on the standby node?

```
Add Dynamic LPAR and CoD Resources for Applications

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Application Contoller Name      testapp
* Minimum number of processing units [ 0.00]
* Desired number of processing units [ 0.00]
* Minimum number of CPUs          [3] #
* Desired number of CPUs          [5] #
* Minimum amount of memory (in megabytes) [20480] #
* Desired amount of memory (in megabytes) [40960] #
* Use CUoD if resources are insufficient? [no] +
* I agree to use CUoD resources        [no] +
  (Using CUoD may result in extra costs)
```





```
Resource Group Name      migtestrg
Participating Node Name(s)  jessica jordan
Startup Policy           Online On Home Node Only
Failover Policy          Fallover To Next Priority Node In The List
Fallback Policy          Never Fallback
Site Relationship         ignore
Dynamic Node Priority
Service IP Label         ha_svc
Volume Groups            datavg
Application Servers       testapp
```

- A. 3
- B. 4
- C. 5
- D. 6

**Correct Answer:** B

**Section:** (none)

**Explanation**



**Explanation/Reference:**

Explanation:

#### **QUESTION 87**

What is required when configuring a disk heartbeat device in a PowerHA 6 cluster?

- A. Must be on a shared data volume group disk
- B. Must be an enhanced concurrent volume group
- C. Must be on a shared disk dedicated for diskhb
- D. Must be a resource in a resource group for both nodes

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 88**

What application specific requirement will prevent it from being used in a PowerHA 7 cluster?

- A. Application must be used with NFSv3
- B. Application must be installed on rootvg
- C. Application must use a specific hostname
- D. Application must be installed on sharedvg

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 89**

A customer notices the following event in hacmp.out.

EVENT START: cluster\_notify NODE1 1 Sep 15 00:01:53 NOTIFY COMMAND:/usr/es/sbin/cluster/utilities/clissuepage NODE1 clverify has detected cluster configuration errors on node NODE1.

Why did this event occur?

- A. Automatic error notification has detected a cluster error and notified the dstrmgrES.
- B. Cluster verification has been initiated by the administrator and errors have been logged to hacmp.out.
- C. Automatic cluster verification has been configured and an error was detected with the cluster configuration.
- D. Notification has been sent that the cluster test tool automated procedure completed and has been logged in hacmp.out.

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 90**

Which fileset must be installed to enable message authentication and encryption?

- A. rsct.crypt
- B. rsct.core.sec

- C. cluster.es.cfs
- D. cluster.es.crypt

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 91**

A shared volume group (VG) was created in a PowerHA 7 environment. Given the following example, what must be changed in the VG definition?

```
root@lpar101[/]# lsvg datavg01
VOLUME GROUP:      datavg01          VG IDENTIFIER:  00c08bb600004c00000001367e2d80fb
VG STATE:          active              PP SIZE:        16 megabyte(s)
VG PERMISSION:    read/write         TOTAL PPs:      639 (10224 megabytes)
MAX LVs:          256                FREE PPs:       627 (10032 megabytes)
LVs:              2                  USED PPs:       12 (192 megabytes)
OPEN LVs:         0                  QUORUM:         2 (Enabled)
TOTAL PVs:        1                  VG DESCRIPTORS: 2
STALE PVs:        0                  STALE PPs:      0
ACTIVE PVs:       1                  AUTO ON:        no
MAX PPs per VG:   32512              MAX PVs:        32
MAX PPs per PV:   1016              AUTO SYNC:      no
LTG size (Dynamic): 256 kilobyte(s)  BB POLICY:      non-relocatable
HOT SPARE:        no                  INFINITE RETRY: no
PV RESTRICTION:   none
root@lpar101[/]#
```

- A. Set QUORUM to 1(Disabled).
- B. Set BB POLICY to relocatable.
- C. Set concurrent capable to enhanced concurrent.
- D. Set hot spare characteristics to yes(one to one).

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

#### **QUESTION 92**

The db\_group resource group become unmanaged after cluster service on node1 is forced down. Which action will change the state of db\_group resource group from unmanaged to Online on node2?

- A. Bring db\_group Online directly on node2
- B. Move resource group db\_group from node1 to node2
- C. Bring db\_group offline and then bring db\_group Online on node2
- D. Run clchres utility to change the state of db\_group resource to Online on node2

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:



#### **QUESTION 93**

After changing a PowerHA 6 cluster to reflect new IP addresses, an administrator attempts to synchronize a cluster and receives the following message:

clhare: A communication error prevents obtaining the VRMF from remote nodes. A DARE cannot be run until this is corrected. Please ensure clcomd is running.

How can the administrator correct this error?

- A. Include the new IP addresses in /.rhosts on all cluster nodes.
- B. Include the new IP addresses in /etc/hosts.equiv on all cluster nodes.
- C. Include the new IP addresses in /usr/es/sbin/cluster/etc/rhosts on all cluster nodes.
- D. Include the new IP addresses in /usr/es/sbin/cluster/etc/clhosts on all cluster nodes.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

Explanation:

**QUESTION 94**

A new application is added to the cluster and verification produces an error. What is the most likely reason?

- A. The application is running on one of the cluster nodes.
- B. Some application dependencies are missing from one or more nodes in the resource group.
- C. The application start script is not executable on one of the nodes in the resource group.
- D. The size of the application start script is not the same on all of the nodes in the resource group.

**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**