

100% Money Back Guarantee

Vendor: IBM

Exam Code: C4040-332

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

Version: Demo

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` logfile 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 bootaddresses, 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 -Fstatus

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_fsevent` 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 cluncmd 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 topologyservices 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

Question No : 25

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 failedetection rate to delay.

Answer: B

Question No : 26

An administrator created a 2-node PowerHA cluster. Each node is in a separate location, and each location has its own network. After starting the cluster, the network interfaces did not come up. Which action will correct this problem?

```

node1 :/ # ping 4.11.1.2
PING 4.11.1.2 (4.11.1.2): 56 data bytes
64 bytes from 4.11.1.2: icmp_seq=0 ttl=255 time=5 ms

node1 :/ # mping -r -v -a 226.0.0.101
mping version 1.0
Localhost is node1, 4.11.1.1
Listening on 226.0.0.101/4098:

node2 :/ # mping -s -v -a 226.0.0.101
mping version 1.0
Localhost is node2, 4.11.1.2
mpinging 226.0.0.101/4098 with ttl=32:
Discarding sender packet
Discarding sender packet

node1 :/ # netstat -nr
Route tree for Protocol Family 2 (Internet):
default      4.11.1.1      UG      507      163388 en0
224.0.0/4    4.11.1.1      UG        0         22 en0

node2 :/ # netstat -nr
Route tree for Protocol Family 2 (Internet):
default      4.11.1.2      UG      501      143132 en1
224.0.0/4    4.11.1.2      UG        0         29 en1

```

- A. Add IP addresses to the /etc/cluster/rhosts file.
- B. Replace IP address 226.0.0.100 with a non-reserved IP address.
- C. Create an IP alias for each node and make it the default route.
- D. Allow multicast network packets to flow between the two locations.

Answer: D

Question No : 27

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

Answer: C

Question No : 28

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	Processors	Memory	I/O	Virtual Adapters	Power Controlling	Settings
Detailed below are the current processing settings for this partition profile.						
Processing mode						
<input type="radio"/> Dedicated <input checked="" type="radio"/> Shared						
Processing units						
Total managed system processing units : 48.00						
Minimum processing units :			<input type="text" value="0.1"/>			
Desired processing units :			<input type="text" value="0.2"/>			
Maximum processing units :			<input type="text" value="0.4"/>			
Shared processor pool:			<input type="text" value="DefaultPool (0)"/>			
Virtual processors						
Minimum processing units required for each virtual processor : 0.10						
Minimum virtual processors :			<input type="text" value="1.0"/>			
Desired virtual processors :			<input type="text" value="1.0"/>			
Maximum virtual processors :			<input type="text" value="2.0"/>			

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

Answer: C

Question No : 29

Which operation can be done by DARE?

- A. Change node name
- B. Change cluster name
- C. Change application server name
- D. Change failover policy of resource group

Answer: C

Question No : 30

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.

Answer: C

Question No : 31

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.

Answer: C

Question No : 32

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 Power770. 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.

Answer: B

Question No : 33

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

Answer: B

Question No : 34

In a PowerHA 6 environment, what is recommended to ensure I/O disk-write buffers are flushed and to reduce chance of deadman switch timeouts?

- A. Change "syncd" frequency to 10 seconds
- B. Set the CAA timeout parameter to "auto"
- C. Configure fast failure detection on shared disks
- D. Change all cluster node disks "hcheck_Interval" attributes to 60 Seconds

Answer: A

Question No : 35

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

Answer: A

Question No : 36

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 pathname to data

Answer: A

Question No : 37

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

Answer: C

Question No : 38

A resource group is configured with a startup policy of "Online on First Available Node" and a service IP address which uses IPAT via aliasing. An administrator needs a custom event to run only during a failover.

Which pre-defined event would be best to specify the custom event?

- A. Rg_move
- B. node_down_remote
- C. acquire_service_address
- D. acquire_takeover_address

Answer: D

Question No : 39

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

Answer: D

Question No : 40

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.

Answer: D

Question No : 41

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

Answer: B

Question No : 42

The database is running on node1 and application is running on node2. The application cannot run unless database is active and the application must be stopped before the database.

What actions will ensure the application and database run correctly?

- A. Configure Resource Group Processing Ordering and set Serial Acquisition Order.
- B. Configure Online on Different Nodes Dependency and set Priority Resource Group.
- C. Configure Parent/Child Dependency and set Parent/Child Dependency between Resource Groups
- D. Configure Start After Resource Group Dependency and set Start After Resource Group Dependency.

Answer: C

Question No : 43

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:

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

Which action will resolve the problem?

- A. Remove the file cldare_lock in /usr/es/sbin/cluster/etc/objrepos/staging and re-try the DARE operation.
- B. Use the PowerHA problem determination function to "ReleaseLocks 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.

Answer: B

Question No : 44

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.

Answer: D

Question No : 45

Given the following cluster topology information, why is this PowerHA 6 cluster configuration unsuitable for migration to PowerHA 7?

Adapter	Type	Network	Net Type	Attribute	Node	IP Address
jessica_hdisk4_01	service	net_dhb_01	diskhb	serial	jessica	/dev/hdisk4
jessica	boot	tok_net_01	ether	public	jessica	19.19.51.193
ha_svc	service	tok_net_01	ether	public	jessica	19.19.51.100
jess_boot1	standby	tok_net_01	ether	public	jessica	192.168.100.1
peyton_hdisk4_01	service	net_dhb_01	diskhb	serial	peyton	/dev/hdisk4
peyton	boot	tok_net_01	ether	public	peyton	19.19.51.194
ha_svc	service	tok_net_01	ether	public	peyton	19.19.51.100
peyton_boot1	standby	tok_net_01	ether	public	peyton	192.168.100.2

- A. Multi-node diskhb is not supported.

- B. Token Ring networks are not supported.
- C. IPAT via replacement is not supported.
- D. Disk heartbeat (diskhb) is not supported.

Answer: C

Question No : 46

An administrator is using PowerHA 7 to define a new cluster using SMIT option "Setup a Cluster, Nodes and Networks" and encountered the following error message:

```
Command: failed      stdout: yes      stderr: no

Before command completion, additional instructions may appear below.

Warning: There is no cluster found.
c1lsc1str: No cluster defined
c1lsc1str: Error reading configuration
```

What is the root cause of the problem?

- A. The nodes were not defined in the DNS
- B. The /etc/cluster/rhosts are not populated correctly
- C. The CAA repository disk is not accessible on all nodes
- D. The CAA cluster was not defined before defining the PowerHA Cluster

Answer: B

Question No : 47

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

Answer: C

Question No : 48

Which are names of default file collections in PowerHA 7?

- A. Cluster_config
OS_Config
- B. Systemjiles
PowerHA_Files
- C. HACMP_Files
Configuration_Files
- D. /etc/hosts
/etc/inet.d
/etc/services
/etc/passwd

Answer: C

Question No : 49

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

Answer: A

Question No : 50

Given the information from hacmp.out below, why is the cluster in config_too_long status?


```

Feb 1 16:48:26 EVENT START: stop_server APP1_as
APP1_rg:stop_server[48] [[ high = high ]]
APP1_rg:stop_server[48] version=1.4.1.10
APP1_rg:stop_server[49] APP1_rg:stop_server[49] cl_get_path
HA_DIR=es
APP1_rg:stop_server[51] STATUS=0
APP1_rg:stop_server[55] [ ! -n ]
APP1_rg:stop_server[57] EMULATE=REAL
APP1_rg:stop_server[60] PROC_RES=false
APP1_rg:stop_server[64] [[ 0 != 0 ]]
APP1_rg:stop_server[68] set -u
APP1_rg:stop_server[71] ALL_SERVERS=All_servers
APP1_rg:stop_server[72] [ REAL = EMUL ]
APP1_rg:stop_server[77] cl_RMupdate resource_releasing All_servers stop_server
Reference string: Sun.Feb.1.16:48:26.GMT.2009.stop_server.All_servers.APP1_rg.ref
APP1_rg:stop_server[88] controlm_rg:stop_server[88] cut -d: -f3
APP1_rg:stop_server[88] cllsserv -cn APP1_as
STOP=/usr/local/hacmp/APP1.stop
APP1_rg:stop_server[90] PATTERN=node1 APP1_as
APP1_rg:stop_server[97] [-x /usr/local/hacmp/APP1.stop ]
APP1_rg:stop_server[99] [ REAL = EMUL ]
APP1_rg:stop_server[104] /usr/local/hacmp/APP1.stop
APP1_rg:stop_server[104] ODMDIR=/etc/objrepos
Shutting down APP1 application
/APP1/serv/scripts/shut_down: There is a request to a device or address that does not exist.
Feb 1 16:54:26 EVENT START: config_too_long 360 /usr/es/sbin/cluster/events/rg_move.rp
app server still up
:config_too_long[64] [[ high = high ]]
:config_too_long[64] version=1.11
:config_too_long[65] :config_too_long[65] cl_get_path
HA_DIR=es
:config_too_long[67] NUM_SECS=360
:config_too_long[68] EVENT=/usr/es/sbin/cluster/events/rg_move.rp
:config_too_long[70] HOUR=3600
:config_too_long[71] THRESHOLD=5
:config_too_long[72] SLEEP_INTERVAL=1
:config_too_long[78] PERIOD=30
:config_too_long[81] set -u
:config_too_long[86] LOOPCNT=0
:config_too_long[87] MESSAGECNT=0
:config_too_long[88] :config_too_long[88] cllsclstr -c
:config_too_long[88] cut -d: -f2
:config_too_long[88] grep -v cname
CLUSTER=testc
:config_too_long[89] TIME=360
:config_too_long[90] sleep_cntr=0
:config_too_long[95] [-x /usr/lpp/ssp/bin/spget_syspar ]
APP1 server still up
WARNING: Cluster testc has been running recovery program '/usr/es/sbin/cluster/events/rg_move.rp' for 360 seconds.
Please check cluster status.

```

- A. The script /APP1/serv/scripts/shut_down does not exist.
- B. The APP1 stop_server script exited with an non zero return code.
- C. The application did not shut down correctly, which caused the rg_move operation to exceed the time limit to complete.
- D. The rg_move event is sleeping for anything between 360 seconds and 3600 hours until the stop_server script is complete.

Answer: C

Question No : 51

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

Answer: A

Question No : 52

Which action can an administrator do when the cluster has mixed versions within a rolling migration?

- A. Create a filesystem using C-SPOC function.
- B. Verify and synchronize the cluster configuration.
- C. Check the cluster status using View Current State from SMIT dialog.
- D. Use cluncmd from the same node that was specified as the target node.

Answer: A

Question No : 53

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

Answer: C

Question No : 54

An administrator wants to configure a new 3-node PowerHA 7 cluster for 3 instances of an application. Each node hosts 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

Answer: C

Question No : 55

An administrator wants to change the nodename of one node while a PowerHA 6cluster is running. Which actions will achieve this?

- A. Change name of node.
Verify and synchronize cluster.
- B. Stop cluster services on the node.
Change the name of the node.
Verify and synchronize cluster.
Start cluster services.
- C. Move all resource groups from the node to be changed.
Change the name of node.
Move resource groups to the hosts with the new name.
- D. Change the hostname in AIX
Update /etc/hosts with the new name of the node
Use C-SPOC "Update HACMP Communication Interface with AIX Settings".

Answer: B

Question No : 56

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

- A. clstat
- B. cldisp
- C. cldump
- D. clcheck_server

Answer: B

Question No : 57

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.

Answer: B

Question No : 58

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. Clusterservices must be restarted on all nodes.
- D. The filesystem must be mounted on a node with an online resource group.

Answer: A

Question No : 59

When PowerHA SystemMirror 7.1 is installed on AIX 7.1, what RSCT component does Cluster Aware AIX (CAA) replace?

- A. Group Services
- B. Resource Manager
- C. Topology Services
- D. Resource Monitoring and Control (RMC)

Answer: C

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Trying our product !


- ★ **100%** Guaranteed Success
- ★ **100%** Money Back Guarantee
- ★ **365 Days** Free Update
- ★ **Instant Download** After Purchase
- ★ **24x7** Customer Support
- ★ Average **99.9%** Success Rate
- ★ More than **69,000** Satisfied Customers Worldwide
- ★ Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.	 Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.	 Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.
---	---	--

Guarantee & Policy | Privacy & Policy | Terms & Conditions

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.