



[HP HP0-914](#)

Exam Name: HP-UX High Availability

Q & A : 149 Q&As

[Pdf Demo](#)

Quality and Value for the HP0-914 Exam

[Just4Exams Practice Exams](#) for HP HP CSA HP0-914 are written to the highest standards of technical accuracy, using only certified subject matter experts and published authors for development.

100% Guarantee to Pass Your HP0-914 Exam

If you do not pass the HP CSA HP0-914 exam on your first attempt using our Just4Exams **HP0-914 testing engine and pdf study guide**, we will give you a FULL REFUND of your purchasing fee.

Downloadable, Interactive HP0-914 Testing engines and PDF Version

Our Exam Preparation Material provides you everything you will need to take a [HP CSA certification](#) examination. Details are researched and produced by [HP Certification](#) Experts who are constantly using industry experience to produce precise, and logical.

Free HP0-914 Exams:

This is demo only, this pdf do not include the questions and answers picture

Exam : HP HP0-914

Title : HP-UX High Availability

1. In a two node HA cluster, two identical servers (same CPU, memory, and I/O card configurations) are designed to host the database tier of a 3-tier application. They are connected to a large external disk array through two fibre channel links. Under normal operations the production database is hosted on server A and a batch reporting tool is hosted on server B. What are the best-practice remedies for slow performance of the production database during non-failure operation? (There are three correct answers. Select any two.)

- A. Implement PVlinks on both servers.
- B. Activate available iCAP CPUs on server A.
- C. Use I/O load balancing software with dynamic multipathing.
- D. Terminate batch reporting on server B to reduce disk array load.
- E. Terminate gWLM and Perfview middleware to free up additional CPU.

Answer: ABCDE

2. Which statements are true regarding Serviceguard? (Select two.)

- A. There are two high-level user access roles available; root and non-root.
- B. Setting up the `./rhosts` file is required to install and configure Serviceguard and cluster nodes.
- C. Setting up the file `cmclnodelist` is recommended to configure or add Serviceguard cluster nodes.
- D. The `cmclnodelist` is used permanently in a Serviceguard cluster to authorize communication between cluster nodes.

Answer: AC

3. A cost sensitive customer wants to utilize Oracle Disk Manager (ODM) for an Oracle single instance database within a Serviceguard cluster. Which HP Serviceguard Storage Management Suite bundles should you recommend?

- A. Bundle 1: HP Serviceguard Storage Management
- B. Bundle 2: HP Serviceguard Storage Management Premium
- C. Bundle 3: HP Serviceguard Storage Management for Oracle
- D. Bundle 4: HP Serviceguard Storage Management for Oracle Premium
- E. Bundle 5: HP Serviceguard Cluster File System
- F. Bundle 6: HP Serviceguard Cluster File System for Oracle
- G. Bundle 7: HP Serviceguard Cluster File System for RAC

Answer: C

4. When a Serviceguard package fails over in a highly available vPar environment, WLM integrates with Serviceguard and disaster tolerant solutions by doing which process?

- A. WLM can move processors from one virtual partition to another.
- B. `cmcl` can move processors from one virtual partition to another.
- C. WLM creates a new vPar to assume the load of the failed partition.
- D. `cmcl` creates a new vPar to assume the load of the failed partition.
- E. WLM can move a Serviceguard package from one virtual partition to another.

Answer: A

5. Which components are the minimum required to implement a Metrocluster solution for two data centers with XP arrays? (Select three.)

- A. a Quorum server residing in one of the data centers
- B. configuration of Business Copies (BCs) on each XP
- C. physical data replication links between XP storage arrays
- D. configuration of LVM mirroring between the two XP arrays
- E. redundant network connection between the two data centers
- F. a standard Serviceguard installation on all potential cluster nodes

Answer: CEF

6. Serviceguard allows a choice of volume managers for data storage. Which volume managers are supported by Serviceguard. (Select three.)

- A. JFS (Journal File System)
- B. CFS (Cluster File System)
- C. LVM (Logical Volume Manager)
- D. SVM (Shared Volume Manager)
- E. CVM (Cluster Volume Manager)
- F. VxVM (Veritas Volume Manager)

Answer: CEF

7. In the context of high availability, which phrase describes Mean Time Between Failure (MTBF)?

- A. the amount of time to restart a failed service
- B. the amount of time to restart a service after a failure
- C. the amount of time a system can provide service without failure
- D. the amount of time between stopping and starting a service after a failure

Answer: C

8. In a two node cluster, two identical servers (same CPU, memory, and I/O card configurations) are designed to host the application and database servers of a 3-tier application. Under normal operations the database is hosted on server A and the application server is hosted on server B; server A will load at 50% of capacity and server B will load at 75% of capacity. What is likely to occur in the event of a single server failure, after both applications have migrated to the surviving node?

- A. Both applications run in degraded mode.
- B. The application server runs in degraded mode.

- C. Nothing; both applications will perform as usual.
- D. The database application runs in degraded mode.

Answer: A

9. Which statements are true for a server running Quorum service? (Select three.)

- A. It cannot be used for any other purpose.
- B. It can be a Red Hat or SuSE Linux system.
- C. It can be a Microsoft Windows 2003 system.
- D. It can be an HP 9000 System running HP-UX 11.x.
- E. There cannot be a Serviceguard package on the Quorum server.
- F. It should reside outside the cluster for which it provides the service.

Answer: BDF

10. Which components are NOT considered single points of failure within a single HP-UX server? (Select two.)

- A. CPUs
- B. FC HBAs
- C. backplane
- D. LAN interfaces
- E. memory DIMMs
- F. mirrored root disks
- G. redundant power supply

Answer: FG

11. Which statement is true regarding Metrocluster?

- A. Metrocluster is actually two clusters with manual failover from one cluster to another.
- B. Metrocluster is actually two clusters with automatic failover from one cluster to another.
- C. Metrocluster can utilize a Quorum service or arbitrator nodes to prevent split brain syndrome between locations.
- D. Metrocluster is one cluster spread across long geographic distances utilizing dual lock disks to avoid split brain syndrome.

Answer: C

12. What are the advantages of RAC cluster? (Select three.)

- A. user scalability
- B. hardware scalability
- C. application scalability
- D. larger database size
- E. better database security
- F. operating system scalability
- G. shorter package failover time

Answer: BCF

13. Which statements are true in a standard Serviceguard cluster environment? (Select two.)

- A. The cluster coordinator node exchanges heartbeat packets with the cluster member nodes.
- B. A cluster lock disk ensures that the LVM cluster volume group is activated only on one node.
- C. Redundant LAN interfaces can be configured as standby LANs for data and heartbeat traffic.
- D. The "vgchange -a e /dev/vg10" command activates the volume group /dev/vg10 in shared mode.
- E. The cmcld process runs on each node to ensure that the LVM cluster volume group is activated in exclusive mode.

Answer: AC

14. What is the purpose of the node interconnect in a Serviceguard cluster with Oracle RAC implementation?

- A. Nodes coordinate their online/offline status.
- B. Nodes coordinate the status of their hardware.
- C. Nodes coordinate the consistency of their Serviceguard cluster membership.
- D. Nodes coordinate the consistency of their individual memory resident database buffer pools.

Answer: D

15. In a Serviceguard CFS cluster, which statement is true in regards to disk I/O?

- A. Only the CFS primary can update the mirrored data.
- B. All CFS secondaries can write user data directly to disk.
- C. Only the CFS primary can write user data directly to disk.
- D. All CFS secondaries read data via the cluster interconnect.

Answer: B

16. In the context of high availability, what is the effect of split-brain syndrome?
- A. different partitions accessing the same disks without knowledge of each other
 - B. two sub-clusters initiated from a single cluster without knowledge of each other
 - C. different applications accessing the same disk without knowledge of each other
 - D. two clusters initiating an application that accesses the same disk simultaneously

Answer: B

17. Which statements are correct regarding a Serviceguard cluster with SGeRAC and Oracle RAC? (Select two.)

- A. If CFS is used then RAW devices are not allowed.
- B. Each Oracle instance must have its own unique database.
- C. The shared disk devices can be a SLVM RAW device or a CFS.
- D. SGeRAC cannot be configured together with Cluster File System (CFS).
- E. You can have more than one Oracle instance accessing the same database at the same time.

Answer: CE

18. What is the main feature of OLAR?

- A. the ability to replace a cell board on an HP-UX system that supports OLAR
- B. the ability to repair a PCI I/O card on an HP-UX system that supports OLAR
- C. the ability to online replace a PCI I/O card on an HP-UX system that supports OLAR
- D. the ability to online replace an HP-UX system that supports OLAR while it is configured in a Serviceguard cluster

Answer: C

19. What are the advantages for a customer using Serviceguard in an HA environment? (Select two.)

- A. Application database corruptions will be minimized.
- B. An upgrade of the operating system will only require a package switch.
- C. A power outage in the customer's computer room will not affect the users.
- D. Rolling upgrades of the application can be performed in a Serviceguard environment.
- E. A failure in a redundant power supply will only cause a package switch to the adoptive node.

Answer: BD

20. Which statements are true regarding High Availability (HA)? (Select three.)

- A. HA can be implemented just by purchasing expensive hardware.
- B. HA is a technical high-performance environment that does not consider business needs.
- C. Downtime is any amount of time when the application is unavailable (planned or unplanned).
- D. HA systems provide a disaster recovery solution to recover from the loss of an entire data center.
- E. HA is a combination of environmental, process, software and computing hardware enhancements that are made to maximize the time that the application and system are available.
- F. Highly available systems typically contain data redundancy through mirroring or RAID, some kind of backup system hardware, a UPS for power outage protection, and multiple LANs.

Answer: CEF

More [HP0-914 Braindumps](#) Information

Related HP0-914 Exams

HP0-091	HP0-095	HP0-914	HP0-G13	HP0-094
HP0-096				

Other HP Exams

HP0-918	HP0-767	HP0-656	HP0-S14	HP0-216
HP0-335	HP0-J15	HP0-Y18	HP0-D05	HP0-069
HP0-661	HP0-786	HP0-087	HP0-P17	HP0-704
HP0-634	HP0-S16	HP0-Y16	HP0-A01	HP0-505