1Z0-068 Q&As

Oracle Database 12c: RAC and Grid Infrastructure Administration

Pass Oracle 1Z0-068 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/1z0-068.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

Examine this query and output:

SQL> select order_flag, cache_size, session_flag, keep_value, 2 from user_sequences where sequence_name = `SEQ1\\';

O CACHE_SIZE S K

Y 10 N N

Performance analysis revealed severe SQ enqueue contention on the SEQ1 sequence.

The SEQ1 sequence is incremented from all instances equally and is frequently used.

Which two statements should you execute to reduce SQ enqueue contention?

A. alter sequence seq1 cache 10000;

B. alter sequence seq1 order;

C. alter sequence seq1 noorder;

D. exec sys.dbms_shared_pool.keep (`SEQ1\\', `Q\\')

E. alter sequence seq1 keep;

Correct Answer: AD

Section: (none)

A: Use cache.

D: The KEEP procedure keeps an object in the shared pool. Once an object has been kept in the shared pool, it is not subject to aging out of the pool. This may be useful for frequently used large objects. When large objects are brought into the shared pool, several objects may need to be aged out to create a contiguous area large enough.

References:

https://ora600tom.wordpress.com/2015/01/09/enq-sq-contention/
https://docs.oracle.com/cd/B19306_01/appdev.102/b14258/d_shpool.htm#i999221

QUESTION 2

Which three statements are true concerning diagnostic components and requirements of Oracle 12c Clusterware?

A. There is one ologgered service for the cluster health monitor (CHM) on each cluster node regardless of cluster size.

B. The Grid Infrastructure Management Repository database must run on a hub node if Flex Clusters are used.

C. There is one osysmond service for the cluster health monitor (CHM) on each cluster node regardless of cluster size.

D. The oclumon utility may be used to get and set parameters for the cluster health monitor (CHM) repository.

E. The diagcollection.pl script must be run from the Grid home directory as the Grid infrastructure owner.

F. The clusterware log files are stored inside the Grid Infrastructure Management Repository database used by the cluster health monitor (CHM).

Correct Answer: BCD

Section: (none)

B: The Oracle Grid Infrastructure Management Repository Runs on one node in the cluster.This must be a Hub Node in an Oracle Flex Cluster configuration.

C: There is one system monitor service on every node. The system monitor service (osysmond) is a real- time, monitoring and operating system metric collection service that sends the data to the cluster logger service.

D: Use the oclumon manage command to view and change configuration information from the system monitor service.

Syntax

oclumon manage -repos {{changeretentiontime time} | {changerepossize memory_size}} | -get {key1 [key2 ...] | alllogger [-details] | mylogger [-details]}

Where changerepossize memory_size: Use this option to change the CHM repository space limit to a specified number of MB.

Note: The OCLUMON command-line tool is included with CHM and you can use it to query the CHM repository to display node-specific metrics for a specified time period. You can also use OCLUMON to perform miscellaneous administrative

tasks, such as changing the debug levels, querying the version of CHM, and changing the metrics database size. Incorrect Answers:

E: You can collect CHM data from any node in the cluster by running the Grid_home/bin/diagcollection.pl script on the node. References: https://docs.oracle.com/database/121/CWADD/troubleshoot.htm#CHDIGCEC

QUESTION 3

Which three statements are true about Global Resource Management in an Oracle 12c RAC database?

A. Lazy remastering occurs when an instance shuts with SHUTDOWN IMMEDIATE.

B. Object remastering causes all blocks in any instances\\' buffer cache from the same object to be mastered in the Global Resource Directory (GRD).

C. When a database instance fails, then some global resource masters lost from the failing instance are remastered among the surviving instances.

D. Lazy remastering instances occurs when an instance shuts with SHUTDOWN TRANSACTIONAL.

E. Global Enqueue resources are recovered after Global Cache Resources after an instance failure.

Correct Answer: BCD

Section: (none)

B: Remastering is the term used that describes the operation whereby a node attempting recovery tries to own or master the resource(s) that were once mastered by another instance prior to the failure. When one instance leaves the cluster, the GRD of that instance needs to be redistributed to the surviving nodes. RAC uses an algorithm called lazy remastering to remaster only a minimal number of resources during a reconfiguration.

D: Using the SHUTDOWN TRANSACTIONAL command with the LOCAL option is useful to shut down a particular Oracle RAC database instance. Transactions on other instances do not block this operation.

C: Recovery from instance failure is automatic, requiring no DBA intervention. In case of instance failure, a surviving instance can read the redo logs of the failed instance. For example, when using the Oracle Parallel Server, another instance performs instance recovery for the failed instance. In single-instance configurations, Oracle performs crash recovery for a database when the database is restarted, that is, mounted and opened to a new instance. The transition from a mounted state to an open state automatically triggers crash recovery, if necessary.

Incorrect Answers:

A: After a NORMAL or IMMEDIATE shutdown, instance recovery is not required. References: https://docs.oracle.com/database/121/RACAD/admin.htm#RACAD8910

QUESTION 4

You support a three-instance, policy-managed, multitenant RAC database CDB1 with two PDBs PDB_1 and PDB_2.

It runs on an eight-node cluster and a serverpool prod_pool has three servers.

Examine these commands executed on HOST01:

$ srvctl add service ?b CDB1 ?db PDB_1 ?ervice HR ?erverpool prod_pool $ srvctl start service ?b CDB1 ?ervice HR

$ srvctl stop service ?b CDB1 ?ervice HR

Which two statements are true?

A. HR is uniformly managed across all instances in prod_pool.

B. srvctl stop service only closes PDB_1 in the instance on HOST01.

C. srvctl stop service closes PDB_1 on all CDB1 instances.

D. srvctl stop service prevents logins for HR only to CDB1 on HOST01.

E. srvctl stop service prevents logins to any instance of CDB1 using service HR.

F. HR is available on a single instance of prod_pool.

Correct Answer: BF

Section: (none)

QUESTION 5

Which two statements are true concerning Oracle 12c Clusterware-managed application VIPs?

A. If an application sends messages to be displayed and sets the DISPLAY variable, then an application VIP is required.

B. An application VIP is created on the default network by the appvipcfg utility.

C. An application VIP is created on the interconnect network by the crsctl utility.

D. An application VIP can be created with the crsctl utility.

E. Application VIPs do not fail over to surviving cluster nodes when the node hosting the VIP fails.

Correct Answer: BD

Section: (none)

B: Oracle 11.2 introduced appvipcfg utility for creating VIPs. From the GRID_HOME/bin directory run the appvipcfg command to create the application VIP. Oracle Clusterware assigns this VIP to a physical server in the cluster and will migrate the VIP to a surviving node in the cluster in the event of a server failure.

Example: appvipcfg create -network=1 -ip=192.168.20.111 -vipname=MyTestVIP -user=grid

D: While you can add a VIP in the same way that you can add any other resource that Oracle Clusterware manages, Oracle recommends using the script Grid_home/bin/appvipcfg to create or delete an application VIP. Incorrect Answers:

E: When a node dies in an Oracle RAC cluster, the Virtual IP (VIP) fails over to a different node. Upon node failure application VIP fails over to a surviving node along with the protected application. It is the Application VIP that is used for accessing the application, thus in case of failure the application will be highly.

References: https://gjilevski.com/2011/11/13/build-ha-for-third-party-application-with-oracle-gi-11-2-0-3/

Latest 1Z0-068 Dumps 1Z0-068 Practice Test 1Z0-068 Exam Questions