

1z0-485 Dumps

Oracle Exadata Database Machine 2014 Implementation Essentials

<https://www.certleader.com/1z0-485-dumps.html>



NEW QUESTION 1

When using IORM, which statement accurately describes when redo log file writes take place?

- A. Based on the priority of the user
- B. Immediately
- C. Based on the resource group of the user
- D. Based on the intradatabase resource plan

Answer: B

Explanation: Whenever a transaction is committed, the corresponding redo entries temporarily stored in redo log buffers of the system global area are written to an online redo log file.

Note: I/O Resource Manager =IORM http://download.oracle.com/docs/cd/A58617_01/server.804/a58397/ch5.htm

NEW QUESTION 2

Consider the following software changes that are performed manually on a Linux server:

1. Changes for Linux kernel firewall configuration
2. Changes for custom performance monitoring tools
3. Changes for security scan tools
4. Changes for Linux system performance optimization

Which of the software changes listed are permitted on Exadata Storage Servers?

- A. 1, 2, 3, and 4
- B. only 3
- C. none
- D. only 2
- E. only 1 and 2
- F. only 2 and 3

Answer: E

Explanation: 1: The Storage Server Patch is responsible for keeping our cell nodes always up-to-date, fixing possible problems, and this patch includes different component patches, like kernel patches, firmware, operation system, etc... for the Storage Server.

Incorrect:

3,4: security scan tools changes and Linux system performance optimization changes would be on the database server.

NEW QUESTION 3

What are two choices that a customer must make that impact diskgroup creation?

- A. What is the level of redundancy required?
- B. What OS will be run?
- C. Where will disk backups be written?
- D. How many databases will run on the cluster?

Answer: AB

Explanation: B: There are a number of ASM disk group attributes that you can set when creating your disk groups, but the following are the most important:

* (B) compatible.rdbms: Set this to the software version of your RDBMS home.

* au_size: Set this to 4 MB.

* compatible.asm: Set this to the software version of your Grid Infrastructure home.

* cell.smart_scan_capable: Set this to TRUE. If this attribute is set to FALSE, Smart Scan will be disabled to segments that reside in the disk group.

* disk_repair_time: Leave this defaulted to 3.6 hours unless you're performing maintenance on a call and know that your outage window will be greater than 3.6 hours.

A:

Once you identify candidate grid disks, use the CREATE DISKGROUP command to create your ASM disk groups.

Here are some of the more important considerations to think about when creating ASM disk groups on Exadata:

* (A) When capacity planning, take your redundancy specification into consideration.

Normal

redundancy will have the effect of reducing your usable storage to half the raw capacity, and

high redundancy will shrink it to a third of your raw disk capacity.

* Simplicity is best on Exadata. Using wild-carded CREATE DISKGROUP syntax not only offers the most terse command syntax, but also ensures your ASM disk groups are spread evenly across your Exadata Storage Server disks.

* Take the time to plan grid disk prefix names and overall grid disk configuration in the context of your desired ASM disk group design.

* Make sure to set the appropriate compatible.asm and compatible.rdbms attributes when creating ASM disk groups.

* Whenever possible, use a 4 MB extent size when creating disk groups on ASM storage.

NEW QUESTION 4

Which two attributes describe key benefits of the InfiniBand network?

- A. All Exadata database servers have a direct path link to each Exadata Storage Server.
- B. Cell-to-cell communication uses Reliable Datagram Sockets (RDS) over InfiniBand to achieve low latency.
- C. Expanding from two Full racks to four only requires adding an external InfiniBand switch to be at the top of the fat-tree topology.
- D. Each InfiniBand link provides 10 Gigabits of bandwidth.

E. Oracle's interconnect protocol uses direct memory access (DMA) to eliminate buffer copies and reduce CPU use.


Answer: CE

Explanation: C: Oracle Exadata is architected to scale-out to any level of performance. To achieve higher performance and greater storage capacity, additional database servers and Exadata cells are added to the configuration – e.g., Half Rack to Full Rack upgrade. As more Exadata cells are added to the configuration, storage capacity and I/O performance increases near linearly.

E: Oracle's interconnect protocol uses direct data placement (DMA – direct memory access) to ensure very low CPU overhead by directly moving data from the wire to database buffers with no extra data copies being made. The InfiniBand network has the flexibility of a LAN network, with the efficiency of a SAN.

Incorrect:

Not A: The architecture of the Exadata solution includes components on the database server and in the Exadata cell. The software architecture for a Quarter Rack configuration is shown below.



Not B: No cell-to-cell communication is ever done or required in an Exadata configuration. Not D: Each InfiniBand link provides 40 Gigabits of bandwidth –

NEW QUESTION 5

You get a Host Unreachable error when you attempt to connect to a server through a network terminal command line. What are two other ways in which you can connect?

- A. Use the ILOM Web GUI.
- B. Use the dcli command at the root prompt on a database node.
- C. Attach a terminal device to the back panel of the server with a serial cable.
- D. Connect by using SQL *Plus.
- E. Log in as root on the database node using the Net1 IP address.

Answer: AC

Explanation: In addition to gaining shell access via SSH to manage your Exadata servers, you can also access them from the Integrated Lights Out Management (ILOM) console or KVM console.

and should typically not require modifications unless you have changed network information inside your database machine.

Note: A KVM switch (with KVM being an abbreviation for "keyboard, video and mouse") is a hardware device that allows a user to control multiple computers from one or more[1] keyboard, video monitor and mouse. Although multiple computers are connected to the KVM, typically a smaller number of computers can be controlled at any given time

NEW QUESTION 6

You are measuring the I/O savings provided by storage indexes for Table A. One of your scripts displays the I/O savings as a result of the storage indexes. Which two statements are true?

```
SQL> select count (*) from transmap.mymap_comp 2 where map_id between 400 and 500;
COUNT(*)
```

103000

Elapsed: 00:00:00.08

```
SQL> select b.name, a.value/1024/1024 value from v$smystat a, v$statname b,
where b.statistic# = a.statistic#
```

and b.nafne in ('cell physical IO bytes eligible for predicate offload',

'cell physical IO interconnect bytes',

'cell physical IO bytes saved by storage index',

'cell physical IO interconnect bytes returned by smart scan') order by 1;

Statistic Value (MB)

cell physical IO bytes eligible for predicate offload 2,255.09 cell physical IO bytes saved by storage index .00

cell physical IO interconnect bytes 1.72

cell physical IO interconnect bytes returned by smart scan 1.68

- A. The storage indexes were disabled using the disable_storage_INDEX Parameter.
- B. Storage Index data is not on the Cell Server's region index memory structures yet because this predicate has not been used.
- C. Since the database was started, no queries were run against this table with the same predicate.
- D. The storage indexes reduced the amount of physical I/O bytes significantly for this query.

Answer: BC

Explanation: * cell physical I/O bytes eligible for predicate offload Number of bytes eligible for predicate offload, an indication of smart scan
* (not D) cell physical I/O bytes saved by storage index Here 0.
Number of bytes saved by a storage index; this is a reflection of how many physical disk I/O bytes (and by proxy, requests) were saved due to a storage index eliminating candidate blocks

NEW QUESTION 7

Which two actions are permitted with Exadata Database Machine?

- A. replacing the Ethernet switch with an equivalent 1U 48-port Gigabit Ethernet switch
- B. installing a second Ethernet switch in the Exadata rack for client access network connectivity
- C. replacing the Sun Data Center InfiniBand Switch 36-spine switch with an InfiniBand Gateway Switch
- D. configuring a Linux active-active channel bonding on the database servers by using two 10-Gigabit Ethernet port
- E. configuring a Fibre Channel over Ethernet (FCoE) protocol on database servers

Answer: AC

Explanation: There are total two category of network switches used to prepare computing environment inside the rack.
* InfiniBand Switches - two models used depending on requirements Sun Oracle 36-port InfiniBand Switch
Sun Oracle InfiniBand Gateway Switch
* Ethernet Switch - primarily for management purposes Cisco Catalyst 4948

	External IB Ports	IB Signal Bitrate	IB Port Labels	Ethernet Ports
Sun Oracle 36-port InfiniBand Switch	36 QSFP+	40Gbps	0A-17A 0B-17B	
Sun Oracle InfiniBand Gateway Switch	36-4 =32 QSFP+	40Gbps	0A-15A 0B-15B	EoIB Two QSFP+ 10Gbps per port 0A-ETH-[1 to 4] 1A-ETH-[1 to 4]
Cisco Catalyst 4948				48 [1-48] 10/100/1000 Base-T

Common information that applies to both of these InfiniBand switches Form Factor: One rack unit (1U) height
Power Supplies: Two Cooling Fans: Five
IB Subnet Management: Yes Firmware Upgradeable: Yes
Command Line Access: Yes. Via ssh and usb-serial access Web Based Management: Yes
SNMP Access: Yes
As you might have figured out by now that the IB Gateway switch is almost like a super set of 36-port switch in terms of features and capabilities.
Incorrect:
Not E: Since the Exadata hardware cannot be modified, it is not supported to add HBA cards to any of the Exadata servers. It is supported to present storage via the network ports on the database servers via NFS or iSCSI, although Fibre Channel over Ethernet (FCoE) is not supported.

NEW QUESTION 8

Identify the relevant steps in the correct order for activating an Auto Service Request (ASR) configuration.

1. Add SNMP traps manually or using OneCommand.
2. Install ASR Manager.
3. Activate a node on ASR Manager.
4. Validate the configuration.
5. Register ASR Manager with Oracle.

- A. 2, 5, 1, 3, and 4
- B. 2, 4, 1, 3, and 5

- C. 5, 2, 4, 1, and 3
- D. 5, 4, 2, 1, and 3

Answer: A

Explanation: 2.Install Oracle Auto Service Request (ASR) Packag 5.Register the ASR Manager
1.Add SNMP Trap Destinations for Multiple Servers Using the dcli Utility
3. Activate Node on the ASR Manager

NEW QUESTION 9

When should you use Hybrid Columnar Compression?

- A. always
- B. on large active tables where deeper compression is desired
- C. on tables or partitions that have fairly static data
- D. on every table where Advanced Compression is not used

Answer: C

Explanation: It is recommended that HCC be enabled on tables or partitions with no or infrequent DML operations

NEW QUESTION 10

Which statement is true about Enterprise Manager 12c for Exadata?

- A. EM Agents are installed on each database and storage node.
- B. EM Agents should not be installed on Exadat
- C. Remote monitoring is the best practice.
- D. EM Agents can communicate to storage cells via SSH.
- E. The Exadata Plug-In for monitoring switches is installed in each EM Agent.
- F. Storage cells communicate directly to the EM Management Server via SNMP and/or SMTP.

Answer: D

Explanation: All the code required for monitoring the various Oracle Exadata hardware components has been bundled into the Oracle Exadata plug-in that is pushed to the agents running on the Oracle Exadata Compute (or Database) nodes. Once the plug-in is deployed on the compute nodes, the database machine is ready for discovery with Oracle Enterprise Manager.

Incorrect:

Not A: Since agents only run on the compute nodes it is important to assign specific agents to different Oracle Exadata components. Not B: Oracle Exadata Discovery in Oracle Enterprise Manager 12c

NEW QUESTION 10

Consider the following setup:

User A1 belongs to resource group High on Database A. User B2 belongs to resource group Low on Database B. User C3 is a user on Database C without any DBRM setup.

DBRM setup:

Database A: Resource group High gets 80% and Low gets 20%. Database B: Resource group High gets 60% and Low gets 40%.

IORM setup:

Database A: Share=20, limit=5 Database B: Share=30, limit=10 Database C: 5 shares

Total number of shares in the IORM setup = 100

What percent of I/O will each database user theoretically be using when the Exadata storage unit I/O throughout is used 100% and no other databases but A, B, and C are running?

- A. A1 = 36%, B2=18%, and C3=9%
- B. A1 = 33%, B2=33%, and C3=33%
- C. A1 = 10%, B2=5%, and C3=20%
- D. A1 = 8%, B2=12%, and C3=5%
- E. A1 = 5%, B2=10%, and C3=85%

Answer: E

Explanation: IORM setup limits Database A to 5%, and Database B is limited to 10%, while Database C has not IORM limit. Not that the resource groups are for CPU allocation.

NEW QUESTION 11

When running OS Watcher, which two data outputs are valid for Exadata storage cell performance analysis? Select the two correct choices that apply?

- A. iostat
- B. mpstat
- C. netstat
- D. pkginfo
- E. label

Answer: AC

Explanation:

Example:

```
# A few find examples
# \(-name "*.*" \) \
# \(-name "*vmstat*.*" -o -name "*iostat*.*" \) \
# \(-name "*vmstat*.*" -o -name "*iostat*.*" \) -mmin -60 \
# \(-name "*vmstat*.*" -o -name "*iostat*.*" -o -name "*netstat*.*" \) -mtime -8 \
```



NEW QUESTION 15

High-disk IOPS (I/Os per second) rates are leading to longer execution times. Which three approaches would you use to improve this scenario?

- A. Enable Write-Back Flash Cache due to heavy DBWR usage.
- B. Enable Write-Back Flash Cache due to heavy LGWR usage.
- C. Swap hard disks to High Capacity disks.
- D. Tune the application to reduce I/O requests.
- E. Leverage IORM to give priority to critical workloads.

Answer: BDE

Explanation: B (not A): Smart Flash Logging is a recent Exadata enhancement that provides physical disk redo write relief in times of high LGWR activity; alone, probably not a driver for Exadata.

NEW QUESTION 16

Exadata Database Machine offers an Intelligent Platform Management Interface for the various components in the Exadata product. Which option is true?

- A. IPMI can be used to remotely start and stop servers.
- B. IPMI can be used to remotely log in to the operating system.
- C. IPMI commands can be executed by using SNMP traps.
- D. IPMI settings can be secured by backing up the Oracle Linux installation on the storage system.

Answer: A

Explanation: IPMI – short for Intelligent Platform Management Interface - is an interface standard that allows remote management of a server from another using standardized interface. The servers in the Exadata Database Machine follow that. It's not an Exadata command but rather a general Linux one.

To power on a cell or database server, issue this from another server:

```
# ipmitool -H procl01-ilom -U root chassis power on
```

To stop a server, use the shutdown command. To stop immediately and keep it down, i.e. not reboot, execute: # shutdown -h -y now

NEW QUESTION 18

Identify four significant changes when a backup of Exadata compute nodes must be performed.

- A. application of operating system patches
- B. before shutdown to preserve storage indexes
- C. application of Oracle patches
- D. reconfiguration of significant operating system parameters
- E. installation or reconfiguration of significant non-Oracle software
- F. storage server rebalancing
- G. addition of an Exadata storage expansion rack

Answer: ACDG

NEW QUESTION 23

Which two statements are true about troubleshooting failed patching activities?

- A. Dependency issues found during yum updates require rolling back to a previous release before retrying.
- B. Bundle patches applied using opatch auto cannot roll back only the database or the grid infrastructure home.
- C. Failed OS patches on database servers can be rolled back.
- D. Failed storage cell patches are rolled back to the previous release automatically.
- E. Database server OS updates can be rolled back using opatch auto -rollback.
- F. Dependency issues found during yum updates should be ignored using the force option.

Answer: AE

Explanation: * Oracle has shifted the strategy to patching the exadata in 11.2.3.2.0 onwards to using Yum as the method of patching.

* Database servers are patched using yum; there is a yum channel for each Exadata image version. Recently, this functionality replaced the "minimal pack."

* In the README for each storage server patch, Oracle provides detailed rollback instructions that are to be followed in the event of a patch failure.

NEW QUESTION 24

Your customer is hesitant to install the Oracle Configuration Manager in their environment. Give them three ways in which it will benefit their Exsdata Database Machine support experience and potentially resolve some of the issues they are having with the length of time it is taking Oracle to process their Exadata Service Requests (SRs).

- A. Host information can be gathered and sent to Oracle for license compliance.
- B. Potential issues can be addressed before they impact operations.
- C. Priority handling can be extended for SRs, with attached configuration.
- D. Exadata patching cannot be done successfully without the Oracle Configuration Manager.
- E. Root cause analysis can be accelerated.

Answer: BCE

Explanation: Oracle Configuration Manager is used to personalize the support experience by collecting configuration information and uploading it to the Oracle repository.

When customer configuration data is uploaded on a regular basis, customer support representatives can analyze this data and provide better service to the customers. For example, when a customer logs a service request, he can associate the configuration data directly with that service request (C). The customer support representative can then view the list of systems associated with the customer and solve problems accordingly.

Some of the benefits of using Oracle Configuration Manager are as follows:

- / Reduces time for resolution of support issues (E)
- / Provides pro-active problem avoidance (B)
- / Improves access to best practices and the Oracle knowledge base
- / Improves understanding of customer's business needs and provides consistent responses and services

NEW QUESTION 29

Which table compression technique will not work on Exadata without licensing the Oracle Advanced Compression Option?

- A. COMPRESS
- B. COMPRESS BASIC
- C. COMPRESS FOR OLTP
- D. COMPRESS FOR QUERY
- E. COMPRESS FOR ARCHIVE

Answer: C

Explanation: OLTP Table Compression is a part of the Oracle Advanced Compression option, which requires a license in addition to the Enterprise Edition.

NEW QUESTION 33

Why is ASM High redundancy an important configuration choice when rolling Exadata Storage Server patching is planned?

- A. High redundancy protects from partner disk failure while a cell is offline being updated.
- B. High redundancy forces ASM rebalance before allowing disks to be taken offline.
- C. High redundancy speeds up ASM fast mirror resync when a cell is brought back online after patching.
- D. Normal redundancy provides the same protection during rolling patching, so High redundancy is not important in this case

Answer: A

Explanation: To ensure redundancy during a rolling upgrade of the Exadata Storage Server Software, MAA recommends ASM high redundancy disk groups.

NEW QUESTION 37

After migrating from legacy disk-based configuration, which three approaches would you use to evaluate the efficiency of Exadata Flash Cache?

- A. Review the Flash Hit rate via cellcli metrics.
- B. Compare Optimized Physical Reads and Total Read Requests in Automatic Workload Repository (AWR).
- C. Review the IOSTAT data that is gathered from each compute node by OS Watcher.
- D. Evaluate the Smart Flash Logging efficiency metrics via CellCLI.
- E. Check I/O latency on large I/O to Temp in AWR.

Answer: ACD

Explanation: A: You wish to determine which database objects are currently cached in Smart Flash Cache.

Use the list flashcachecontent CellCLI command to report the objects

currently stored in Smart Flash Cache and map these to database object names. Using dcli or cellcli from a storage cell, run the following command:

```
[oracle@cm01dbm01 ~]$ dcli -g ./cell_group cellcli -e list flashcachecontent \  
> attributes dbUniqueName,hitCount,missCount,cachedSize,objectNumber cm01cel01: EDW 0 2 98304 3  
cm01cel01: DWPRD 0 0 57344 8  
cm01cel01: VISY 0 0 8192 8  
cm01cel01: EDW 9 15 729088 18  
cm01cel01: DWPRD 0 0 16384 18
```

Output omitted for brevity

C: IOSTAT can be used to get both DISK and FLASH performance data.

D: Using Storage Cell Metrics to Measure IORM (Exadata IO Resource Manager) Performance Impact

Exadata provides a number of performance metrics for each of the Category, Interdatabase, and Intradatabase IORM types. You can find these metric and their

descriptions using
CellCLI commands.

With a representative database workload running and after your IORM plan has been created, use the list metriccurrent or list metrichistory CellCLI command to report your current or historical IORM metrics.

NEW QUESTION 39

You are calculating the storage capacity that your customer will require on their ZFS Backup Appliance. Identify the set of factors that will affect this calculation.

- A. database size, full/incremental cycle, RPO (Recovery Point Objective), and RTO (Recovery Time Objective)
- B. database size, retention period, RPO (Recovery Point Objective), and compression
- C. database size, full/incremental cycle, compression, and retention period
- D. database size, compression, and backup window
- E. full/incremental cycle, retention period, and ZFS Backup Appliance model

Answer: C

Explanation: Note: The Sun ZFS Backup

Appliance has extremely fast backup and restore throughputs, ensuring that backup windows and recovery time objectives (RTOs) are met by providing timely recovery in the event of a disaster

Incorrect:

Not A, Not E: compression is a factor

NEW QUESTION 43

Which three migration options are available when you migrate a database from the Big Endian format system to Exadata?

- A. Data Pump Export and Import
- B. transportable database
- C. transportable tablespaces
- D. Data Guard
- E. Insert as Select

Answer: BCD

Explanation: Methods at a Glance:

* Data Guard Physical Standby

* Transportable Database (Note:732053.1)

* Transportable Tablespace (TTS)

NEW QUESTION 46

The Exadata compute node, exadbs04, has been having issues for some time. The server had to be repaired and you just received a new server. Your LVM snapshot was corrupt and now the compute node must be re-image with an Oracle Support installation image. Which three tasks would you need to perform to re-image the compute node?

- A. Mount the LVM backup image across the network and recover the system.
- B. Remove exadbs04 from the RAC cluster and all associated Oracle Homes from the central inventory.
- C. Create a CELLBOOT USB Flash Drive using an external USB drive.
- D. Reboot the new server using a bootable recovery image from the external USB drive.
- E. Reconfigure all site-specific settings (host name, IP addresses, NTP server, and so on).

Answer: BCE

Explanation: C: Exadata provides a way for you to create your own CELLBOOT USB. All you have to is plug-in a USB disk which should be at least 1 GB and run the tool:

```
$ /opt/oracle.SupportTools/make_cellboot_usb
```

Exadata turns your USB disk into a bootable disk of active image of your system.

NEW QUESTION 49

Which Exadata feature eliminates unnecessary data transfers between database nodes and storage?

- A. database views
- B. InfiniBand networking
- C. Flash Cache
- D. high performance SAS2 disk drives
- E. cell offloading

Answer: C

Explanation: The Exadata Smart Flash Cache feature of the Exadata Storage Server Software intelligently caches database objects in flash memory, replacing slow, mechanical I/O operations to disk with very rapid flash memory operations. .

NEW QUESTION 52

The mpstat output from OS Watcher shows a database node as being 90% idle on an average. What would you do to get a full picture of CPU utilization on the entire Exadata RAC cluster?

- A. Average the mpstat id1 output from all the nodes.
- B. Ask application users if they have noticed a slowdown in screen response.
- C. Look for an increase in batch job servicing times.
- D. A & B above

Answer: A

NEW QUESTION 53

What would be the best way to ensure that batch Jobs do not impact performance of online queries that use the same database?

- A. Configure IORM to disable Flash Cache usage for batch workload.
- B. Build a Database Resource Management Plan, giving priority to query workload.
- C. Deploy Instance Caging to control batch jobs.
- D. Configure and enable Parallel Query.
- E. Configure resource plans by using the I/O Resource Manager on the storage cells, giving priority to query workload.

Answer: B

Explanation: Using the Database Resource Manager, you can distribute available processing resources by allocating percentages of CPU time to different users and applications. In a data warehouse, a higher percentage may be given to ROLAP (relational on-line analytical processing) applications than to batch jobs.
Incorrect: not C:

In addition to controlling I/O utilization on the Exadata storage cells, Oracle also allows you to limit CPU resource utilization on the compute nodes using a technique called instance caging.

NEW QUESTION 57

Which is a best practice for High Availability (HA) in an Exadata environment?

- A. Checksums performed on the Exadata Storage Server ensure logical consistency of block content.
- B. When a standby database is in place, DB_BLOCK_CHECKING is not required and thus not recommended.
- C. Oracle Exadata Storage Server Software HARD checks operate transparently after enabling DB_BLOCK_CHECKSUM on a database.
- D. Due to the different architecture that Exadata has with storage servers, compared to regular Storage Area Network, DB_LOST_WRITE_PROTECT need not be set.
- E. Stretched RAC is the best alternative for Exadata that combines both HA and DR into one.

Answer: C

Explanation: Exadata storage cells include Oracle Hardware Assisted Resilient Data (HARD) to provide a unique level of validation for Oracle block data structures such as data block address, checksum, and magic numbers prior to allowing a write to physical disks. HARD validation with Exadata is automatic (setting DB_BLOCK_CHECKSUM is required to enable checksum validation). The HARD checks transparently handle all cases including Oracle ASM disk rebalance operations and disk failures.

Incorrect:

Not A: The checksum is used to validate that a block is not physically corrupt, detecting corruptions caused by underlying disks, storage systems, or I/O systems. Checksums do not ensure logical consistency of the block contents.

Not B: MAA recommends that you set DB_BLOCK_CHECKING=MEDIUM or FULL on the physical standby as a minimum practice to prevent the standby from various logical block corruptions.

Not D: The recommendation is:

On the primary database: DB_LOST_WRITE_PROTECT=TYPICAL (default TYPICAL on Exadata) On the Data Guard Physical Standby Database:

DB_LOST_WRITE_PROTECT=TYPICAL Not E: Protection from a broad range of, but not all disasters

While not a full disaster recovery (DR) solution, an Extended Distance

Oracle RAC or Oracle RAC One Node deployment will provide protection from a broad range of

disasters. For a full DR protection Oracle recommends deploying an Oracle RAC together with a local and a remote Oracle Data Guard setup as described in the Maximum Availability Architecture (MAA).

NEW QUESTION 61

Which two statements are true about enabling write-back flash cache?

- A. When enabling write-back flash cache in a non rolling manner, it is important to ensure that asmdeactivationoutcome is set to YES and asmModestatus is set to ONLINE for all grid disks.
- B. Before using write-back flash cache, you need to verify the minimum required versions.
- C. Before write back-flash cache is enabled, you need to drop the Flash Cache first.
- D. The setting flashCacheMode should be set to writeback by updating cellinit.ora and restarting cellsv.
- E. When enabling write-back flash cache in a rolling manner, dcli should be used to inactivate the grid disks on all cells first.

Answer: BC

Explanation: B: Exadata storage version 11.2.3.2.1 is the minimum version required to use this write back flash cache option.

C: Steps for Enabling Write back flash cache:

First of all, you don't need the stop CRS or database (This is ROLLING method) ,you can do it cell by cell.

* drop flashcache

* Be sure asmdeactivationoutcome is YES is before disabling grid disk

NEW QUESTION 65

Identify two reasons why using external tables for high performance data loads on Exadata is more advantageous than SQL*Loader.

- A. Transformations can be applied directly on the file data using SQL or PL/SQL constructions.

- B. External tables allow transparent parallelization inside the database.
- C. Parallelizing loads with external tables enables more efficient space management.
- D. External tables can take advantage of storage indexes to speed up data loads.

Answer: C

Explanation: Parallelizing loads with external tables enables a more efficient space management compared to SQL*Loader, where each individual parallel loader is an independent database sessions with its own transaction.

NEW QUESTION 66

What are three ways to be informed of an Exadata Storage Server disk failure?

- A. review of the output of list alerthistory
- B. review of the output of select status from vSdatafile
- C. review of the ILOM log
- D. alert notifications through email when SMTP notification is configured
- E. alert notifications in the Enterprise Manager Exadata Plug-In
- F. review of database server iostat output

Answer: ADE

Explanation: A: A disk status change would be associated with an entry in the storage cell alerthistory.
D: The Cell Alert Delivery Configuration Worksheet allows you to provide SMTP details to allow for e-mail communication of various cell alerts and failures.
Incorrect:
Not F: Use the iostat command to report statistics about disk input and output, and to produce measures of throughput, utilization, queue lengths, transaction rates, and service time.

NEW QUESTION 67

Which two statements are true about CellCLI?

- A. Using a CellCLI command, you can change the SNMP subscribers or add multiple SNMP subscribers.
- B. Using CellCLI, you can stop and start the ASM instances that are using the grid disks on the Exadata storage cell.
- C. With a simple CellCLI command, you can synchronize changes to all the storage cells in the Exadata realm.
- D. Using a CellCLI command, you can see whether grid disks are being used by ASM or not.
- E. You cannot use CellCLI to gather the serial numbers of physical disks.

Answer: BD

Explanation: B:
* Making a grid disk inactive effectively offlines its associated ASM disk.
CellCLI> alter griddisk PRORECO_CD_11_cell01 inactive GridDisk PRORECO_CD_11_cell01 successfully altered
* The command will wait until the ASM disk becomes offline. If you want the prompt to come back immediately without waiting, you can use nowait clause.
CellCLI> alter griddisk PRORECO_CD_11_cell01 inactive nowait You can make it active again:
CellCLI> alter griddisk PRORECO_CD_11_cell01 active GridDisk PRORECO_CD_11_cell01 successfully altered Incorrect:
not E: CellCLI> list physicaldisk detail output include:
physicalSerial: XXXXXX

NEW QUESTION 68

Which statement is true about Oracle compression?

- A. A non-partitioned table can use Advanced Compression and Hybrid Columnar Compression concurrently.
- B. A partitioned table can define the use of Advanced Compression or Hybrid Columnar Compression for each partition.
- C. Hybrid Columnar Compression can be defined for a single column.
- D. A partitioned table can use only Advanced Compression or only Hybrid Columnar Compression.

Answer: CD

Explanation: C: • Tables are organized by column and compressed. This makes it much easier to get similar values together, which enhances the compression greatly. D: It is recommended that HCC be enabled on tables or partitions with no or infrequent DML operations. If frequent DML operations are planned on a table or partition, then the Oracle Advanced Compression Option is better suited for such data.

NEW QUESTION 73

Your customer would like to use DBFS in their Exadata environment. They are asking you for the key characteristic of DBFS on Exadata.

- A. DBFS in an Exadata environment is faster than an NFS mount system because of the Smart Scan performance gains.
- B. Tens of thousands of files are the perfect use case for DBFS on Exadata.
- C. DBFS offers tremendous I/O bandwidth.

Answer: B

NEW QUESTION 77

Storage indexes are unique to the Exadata Database Machine and their primary goal is to reduce the amount of I/O required to service I/O requests for Exadata Smart Scan. Put the following steps in order:

1. The Exadata cell services software conducts I/O requests on I MO storage regions.
2. cellsrv checks the high and low values, and determines the storage region does not contain any values meeting the predicate.
3. The database is started.
4. Physical I/O to the region is bypassed if the query selection falls outside the high/low storage index storage.
5. A query is issued against the MYOBJ_CTRL table that has a predicate OBJECT_ID=1500.
6. A subsequent query is Issued against the MYOBJ_CTRL table that has a predicate OBJECT_ID=2234.
7. MYOBJ_STATE'S region index is populated with high and low values for the OBJECT_ID column during the I/O request.

- A. 3, 7, 5, 2, 1, 6, and 4
- B. 3, 7, 1, 5, 6, 2, and 4
- C. 3, 6, 1, 7, 2, 5 and 4
- D. 3, 5, 1, 7, 6, 2. and 4

Answer: D

Explanation: 3. The database is started.

5. A query is issued against the MYOBJ_CTRL table that has a predicate OBJECT_ID=1500.

The Exadata cell services software conducts I/O requests on I MO storage regions.

7. MYOBJ_STATE'S region index is populated with high and low values for the OBJECT_ID column during the I/O request.

6. A subsequent query is Issued against the MYOBJ_CTRL table that has a predicate OBJECT_ID=2234.

2. cellsrv checks the high and low values, and determines the storage region does not contain any values meeting the predicate.

4. Physical I/O to the region is bypassed if the query selection falls outside the high/low storage index storage.

Note: Example:

Step 1 (step 5 in answer): The first time that each cell's cell services software issued an I/O request to access extents

from the D14.MYOBJ_UNCOMP table, Exadata populated a region index for each storage index with the high and low values found for the OBJECT_ID column, based on the WHERE

OBJECT_ID BETWEEN 100 AND 200 query predicate.

Step 2 (step 6 in answer) : Subsequent queries against this table generated an iDB message instructing Exadata's

cell services software to read the same sets of extents as the first query, but in this case the

region indexes would have been populated based on the I/Os read from the first query. Step 3: (step 2 in answer) Prior to issuing a physical disk read, cellsrv

checks the high and low values stored in the region index and if it determines that a storage region does not contain any values meeting the query predicate, bypasses the physical I/O to the region.

NEW QUESTION 81

When an Exadata Storage Server hard disk failure alert is received, what manual action must you take to restore the system to full redundancy?

- A. Replace the disk and run MegaCLI to rebuild the degenerate mirror.
- B. No manual action is required because Automatic Storage Management (ASM) fast mirror resync is automatic
- C. No manual action is required because ASM rebalancing is automatic.
- D. Replace the disk and manually copy the mirror extents to the new drive.
- E. Run RMAN REPAIR FAILURE.

Answer: C

Explanation: As soon as the Hard Disk failure is noticed by the MS (Management Server) background process on the Cell, it will raise an alert that will also be published to Grid Control, if configured. Immediately, due to Pro-Active Disk Quarantine, the ASM-, Grid- and Celldisks get dropped. ASM rebalancing is triggered. You as the responsible Admin notice the alert and order a replacement Disk resp. use a Spare Disk to plug it into the Cell after you plugged out the damaged one. The Cell can stay online, because the Hard Disks are hot-pluggable.

No further administrative work to be done, typically.

NEW QUESTION 85

Your customer has purchased their brand new Exadata Database Machine X3-2 Full rack to achieve 20-TB-per-hour backups on their 300 TB database. You have been called to determine why they are not able to achieve even half that rate. Use the information in the image below to identify two reasons their backups are still slow.

FLOW	COMPONENTS	QUANTITY	ESTIMATED RATE (GB/SEC)	THROUGHPUT RATE (GB/SEC)
1	Exadata Cell	14	1 ⁴	14
2	Database Server	8	2.0 ⁵	16
3	Network to Media Server:			
	a) Media Server InfiniBand HCA using TCP/IP	2	2	4
	b) Media Server 10GigE Active/Active NICS	4	1	4
	c) Media Server GigE NICS	4	0.12	0.48
4	Media Server to Tape Library SAN ⁶ Links	4	0.8	3.2
5	Tape Drives (LTO4)	14	0.17	2.3

- A. The Media Server to Tape Library transfer is causing a bottleneck.

- B. The number of tape drives is causing a bottleneck.
- C. Exadata backups are being performed using RMAN over the 10 Gigabit Ethernet network.
- D. The number of database servers is causing a bottleneck.

Answer: BC

Explanation: The following bottlenecks will be reached first if using InfiniBand fabric to media server:

1. Media server to tape library SAN transfer rate
2. Number of tape drives

NEW QUESTION 89

Which CellCLI command is used to conclude that a storage server can be taken offline without impacting database availability?

- A. LIST GRIDDISK ATTRIBUTES name, asmDeactivationOutcome;
- B. LIST CELLDISK ATTRIBUTES name, raidLevel;
- C. LIST CELL ATTRIBUTES cellsrvStatus, flashCacheMode;
- D. LIST CELLDISK ATTRIBUTES name, raidLevel, asmModeStatus;
- E. LIST GRIDDISK ATTRIBUTES name, status;

Answer: A

Explanation: ASMDeactivationOutcome – recall that grid disks can be deactivated, which is effectively taking them offline. Since ASM mirroring ensures that the data is located on another disk, making this disk offline does not lose data. However, if the mirror is offline, or is not present, then making this grid disk offline will result in loss of data. This attribute shows whether the grid disk can be deactivated without loss of data. A value of “Yes” indicates you can deactivate this grid disk without data loss.

```
CellCLI> list griddisk attributes name, ASMDeactivationOutcome, ASMModeStatus DBFS_DG_CD_02_cell01 Yes ONLINE  
DBFS_DG_CD_03_cell01 Yes ONLINE DBFS_DG_CD_04_cell01 Yes ONLINE  
... output truncated ...
```

NEW QUESTION 90

Identify two tasks that the Database Resource Manager can perform, which the I/O Resource Manager cannot.

- A. Manage I/O based on the application that is connected to the database.
- B. Manage the number of parallel sessions for a query.
- C. Manage I/O and CPU between databases on the same cluster or physical database host.
- D. Terminate database sessions when certain limits have been reached.
- E. Manage the throughput of an I/O-bound application based on the service name used for the connection.

Answer: BD

Explanation: IORM is similar to Oracle Database Resource Manager (DBRM) in that it provides a means for controlling allocation of system resources. Where DBRM's primary goals are to control CPU resources, limit the degree of parallelism, and impose resource consumption constraints for different types of sessions within an Oracle database, IORM's goal is to govern I/O resource allocations between databases on a shared storage infrastructure. When consolidating Oracle databases on Exadata, IORM can be used to ensure that I/O is controlled between databases as well as classifications of consumers that utilize the same ASM disk infrastructure and, as such, provide resource control capabilities beyond what DBRM provides within a database.

/ Using the Database Resource Manager, you can:

- * Guarantee certain users a minimum amount of processing resources regardless of the load on the system and the number of users
- * Distribute available processing resources by allocating percentages of CPU time to different users and applications. In a data warehouse, a higher percentage may be given to ROLAP (relational on-line analytical processing) applications than to batch jobs.
- * Limit the degree of parallelism of any operation performed by members of a group of users etc

NEW QUESTION 93

What are the three customer options for hosting the Platinum Services Advanced Support Gateway?

- A. Install on Exadata Engineered System.
- B. Provide individual x86 64-Bit gateway hardware.
- C. Install in Oracle Virtual Machine with required hardware.
- D. Install on Oracle Database Appliance.
- E. Purchase the recommended x86 64-Bit gateway hardware from Oracle.

Answer: ACE

Explanation: Oracle provides a flexible model for deployment of the Oracle Advanced Support Gateway via no charge software that can be installed in three configurations:

1. On x86 customer provided servers;
2. On x86 servers purchased from Oracle;
3. Installed on Oracle Exalogic Elastic Cloud or Oracle SuperCluster systems.

Note:

Oracle Platinum Services provides enhanced support for high availability and performance to Oracle Premier Support customers running certified configurations of Oracle Exadata Database Machine, Oracle Exalogic Elastic Cloud, or Oracle SuperCluster, at no additional cost.

NEW QUESTION 97

You are asked to enable Write Back Flash Cache for one of your customers X3-2 Full Rack in a rolling fashion. Put the following steps in the right order:

- A) Set flashCacheMode to writeback.
- B) Inactivate all grid disks and shut down cellsrv.

- C) Drop the Flash Cache.
- D) Verify that cell for flashcacheMode is writeback.
- E) Create the Flash Cache.
- F) Start up cellsvr and activate all grid disks.
- G) Check griddisk asmdeactivationoutcome and asmmodestatus for all grid disks to ensure that all grid disks on all cells are set to Yes and online, respectively.

- A. B, C, A, F, E, D, and G
- B. G, A, B, C, F, E, D, and G
- C. G, C, B, A, F, E, D, and G
- D. G, E, C, B, F, A, D, and G

Answer: C

Explanation: Steps for Enabling Write back flash cache:

First of all, you don't need the stop CRS or database (This is ROLLING method) ,you can do it cell by cell.

* (C) drop flashcache

```
CellCLI> drop flashcache
```

```
Flash cache owtcel03_FLASHCACHE successfully dropped
```

* (G) Be sure asmdeactivationoutcome is YES is before disabling grid disk

```
</pre>
```

```
CellCLI> list griddisk attributes name,asmmodestatus,asmdeactivationoutcome
```

* (B)

```
CellCLI> alter griddisk all inactive
```

NEW QUESTION 102

Your customer wants to increase the size of the DATA diskgroup on the Exadata systems. The customer is currently using 600 GB disks. Which two are the best options that you would recommend?

- A. adding a High Capacity Storage expansion rack
- B. expanding the Exadata rack from a Half Rack to a Full Rack
- C. moving underutilized grid disks from the RECO diskgroup to DATA
- D. adding a ZFS storage appliance
- E. adding a High Performance Storage expansion rack

Answer: AB

Explanation: A: Oracle Exadata Storage Expansion Rack X4-2 enables you to grow the Oracle Exadata storage capacity and bandwidth of Oracle Exadata Database Machine X4- 2 and X3-8 and Oracle SuperCluster. It is designed for database deployments that require very large amounts of data, including historical or archive data; backups and archives of Oracle Exadata Database Machine data; documents; images; file and XML data; LOB's; and other large unstructured data.

NEW QUESTION 105

How would you execute CellCLI commands and scripts?

- A. using SQL*Plus on database nodes
- B. by CellCLI commands executed on the database nodes
- C. using third party tools after installing the CellCLI RPM plug-in
- D. directly executing the commands and scripts on the Exadata storage cell
- E. remotely by connecting to Port 1521 using SSL

Answer: D

Explanation: The storage cells in Exadata Database Machine are managed via two tools called CellCLI and DCLI.

NEW QUESTION 110

Which statement is true about Exadata Storage Servers?

- A. The Exadata Storage Server automatically deletes old diagnostic and metric files.
- B. Exadata requires a running database instance on all storage servers and database servers.
- C. Redundancy for user data stored in a database that is running on Exadata is achieved with RAID5.
- D. Communication between a database and an Exadata storage flows over low latency 10 Gigabit Ethernet.
- E. Exadata uses network affinity to determine which storage server data is written.

Answer: C

Explanation: Incorrect:

Not B: Only on database servers.

not D: Exadata Storage Servers have dual 40 Gigabit InfiniBand links that provide connectivity many times faster than traditional storage or server networks.

NEW QUESTION 115

Identify three Exadata Storage Server software processes and their purpose?

- A. CELLSRV: The Cell Server is responsible for servicing disk I/O and predicate processing offload.
- B. CELLSRV: The Cell Server is responsible for balancing workload to other storage servers.

- C. MS: The Management Server is responsible for storage cell management and configuration.
- D. MS: The Management Server is responsible for starting a local Enterprise Manager agent.
- E. RS: The Restart Server is responsible for Automatic Storage Management (ASM) instance restart.
- F. RS: The Restart Server is responsible for CELLSRV and MS monitoring and restart.

Answer: ACF

Explanation: A: CELLSRV (Cell Services) is the primary component of the Exadata software running in the cell and provides the majority of Exadata storage services. CELLSRV is multi-threaded software that communicates with the database instance on the database server, and serves blocks to databases based on the iDB protocol. It provides the advanced SQL offload capabilities, serves Oracle blocks when SQL offload processing is not possible, and implements the DBRM I/O resource management functionality to meter out I/O bandwidth to the various databases and consumer groups issuing I/O.

C: The MS is the primary interface to administer, manage and query the status of the Exadata cell. It works in cooperation with the Exadata cell command line interface (CLI) and EM Exadata plug-in, and provides standalone Exadata cell management and configuration. For example, from the cell, CLI commands are issued to configure storage, query I/O statistics and restart the cell. Also supplied is a distributed CLI so commands can be sent to multiple cells to ease management across cells.

F: Restart Server (RS) ensures the ongoing functioning of the Exadata software and services. It is used to update the Exadata software. It also ensures storage services are started and running, and services are restarted when required.

NEW QUESTION 118

Consider this CellCLI command:

```
CellCLI> CREATE GRIDDISK ALL HARDDISK PREFIX=data, size=423G;
```

Which two statements describe what happens when you execute this command?

- A. It creates one 423 GB grid disk on the first available cell hard disk.
- B. It creates one 423 GB grid disk on each available cell hard disk.
- C. It creates grid disks on the outermost 423 GB that is available on each hard disk.
- D. It creates grid disks on the innermost 423 GB that is available on each hard disk.
- E. It creates an Exadata Smart Flash Cache on all flash drives.

Answer: BC

Explanation: * Example:

```
CellCLI> create griddisk all harddisk prefix=temp_dg, size=570G
```

This command will create 12 Griddisks, each of 570G in size from the outer (fastest) sectors of the underlying Harddisks. It fills up the first 2 Celldisks entirely, because they have just 570G space free – the rest is already consumed by the OS partition.

NEW QUESTION 120

Identify a recommended configuration to set up Auto Service Request (ASR) for Exadata.

- A. Install ASR Manager on Exadata Database Server.
- B. Install ASR Manager on Exadata Storage Server.
- C. ASR is not recommended for Exadata; the Oracle Configuration Manager is preferred.
- D. Install ASR Manager on a Standalone Server.

Answer: D

Explanation: The recommended configuration is to install the ASR Manager, which receives fault telemetry information from the servers in Oracle Exadata Database Machine, on an external standalone server. This server must run Solaris or Linux as the operating system.

NEW QUESTION 124

A customer has three databases named CC, FIN, and DW. The CC database is for their CallCenter. Even a slight decrease in the response time of the database would mean more people "on hold" in their data center. The orders received through the CallCenter are stored in the Finance (FIN) database. Both databases, CC and FIN, serve as sources for the Data Warehouse (DW) database. All databases use the same Automatic Storage Management (ASM) diskgroup and therefore, the same physical disks on Exadata storage. The customer wants to dynamically control the resources that are available for the CallCenter database because this has direct impact on their operations?

Which option should be implemented?

- A. DBRM on the CC database
- B. DBRM on all the databases
- C. IORM, because DBRM cannot be changed for an existing session
- D. IORM and DBRM

Answer: B

Explanation: Using the Database Resource Manager, you can:

Distribute available processing resources by allocating percentages of CPU time to different users and applications. In a data warehouse, a higher percentage may be given to ROLAP (relational on-line analytical processing) applications than to batch jobs.

NEW QUESTION 129

You are conceh is the best location to point your customer to, for finding the latest Exadata prned about how recovery from a failed Exadata Storage Server would work. Which statement is true about the Exadata CELLBOOT USB?

- A. Exadata automatically copies OS binaries and configuration files from another cell during patching to sync the internal USB.
- B. Regular Exadata Storage Server patches make sure that the internal USB is updated so that it can be used for recovery.
- C. Changes to network configuration files are possible without using ipconf.pl as long as all relevant files are updated.
- D. Some space from the CELLBOOT USB is used for cache metadata when write back flash cache is enabled.
- E. Each patch that is performed on Exadata storage servers requires manual resync of the CELLBOOT USB.

Answer: AC

Explanation: A (not B, not E): Oracle automatically performs backups of the operating system and cell software on each Exadata Storage Server. The contents of the system volumes are automatically backed up and require no Oracle DMA intervention or operational processes. Oracle assumes responsibility for backing up the critical files of the storage cells to an internal USB drive called the CELLBOOT USB Flash Drive.

C:

Note: The ipconf utility is installed in /opt/oracle.cellos on both the Exadata Storage Servers and Compute Nodes and is symbolically linked to /usr/local/bin/ipconf. ipconf is called at system startup time to set and validate your Exadata server network information.

NEW QUESTION 131

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 1z0-485 Exam with Our Prep Materials Via below:

<https://www.certleader.com/1z0-485-dumps.html>