

Exam Questions 1z0-485

Oracle Exadata Database Machine 2014 Implementation Essentials

<https://www.2passeasy.com/dumps/1z0-485/>



NEW QUESTION 1

Identify three best practices for applying asmdeactivationoutcome es on Exadata Database Servers and Exadata Storage Servers?

- A. Backing up database servers and storage cells is not recommended before performing planned maintenance.
- B. Database server updates can be rolled back using the the "yum downgrade" procedure.
- C. Bundle patches do not require testing before being installed on a production system.
- D. It is recommended that Exadata systems with Data Guard configured use the "Standby First" patching approach.
- E. Patching should never be interrupted due to a connection dro
- F. It is therefore recommended that you use VNC or the screen utility.
- G. Before patching cells in a rolling manner, you must check asmdeactivationoutcome amModestatus and make sure that cells on all disks are online and that disks can be deactivated.

Answer: DEF

NEW QUESTION 2

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 3

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.name 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 4

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 5

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 6

Which two migration will result in the least down time for a physical or logical migration of a database with active users during the migration?

- A. using GoldenGate
- B. using SQL Developer migration tools
- C. using Data Guard Physical Standby
- D. using cross-platform transportable tablespaces
- E. using incremental cross-platform RMAN restore

Answer: AE

Explanation: A: Zero Downtime Migration to Oracle Exadata using Oracle GoldenGate

- Reduce downtime for any method
- Zero data loss fallback
- Phased migration

E: Cross Platform Incremental Backups

- Reduce downtime for Transportable Tablespaces

NEW QUESTION 7

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 8

Which of the following software changes that are performed manually on a Linux system?

- A. owner's guide
- B. Patch database in MyOracle Support
- C. MyOracle Support note 888828.1
- D. MyOracle Support for database patches, ULN for OS patches, and Sun Support for Server and InfiniBand patches.

Answer: B

Explanation: Before starting, we would like to share and note here two documents from My Oracle Support, aka metalink. These notes must be the first place that you need to go to review before patching the Exadata environment.

* (B) Database Machine and Exadata Storage Server 11g Release 2 (11.2) Supported Versions (Doc ID. 888828.1)

- This is for the second and third generation (V2 and X2) for Oracle Exadata, using Sun hardware.

* Database Machine and Exadata Storage Server 11g Release 1 (11.1) Supported Versions (Doc ID. 835032.1)

* - This is for the first generation (V1) for Oracle Exadata, using HP hardware.

NEW QUESTION 9

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 Exadata
- 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

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:

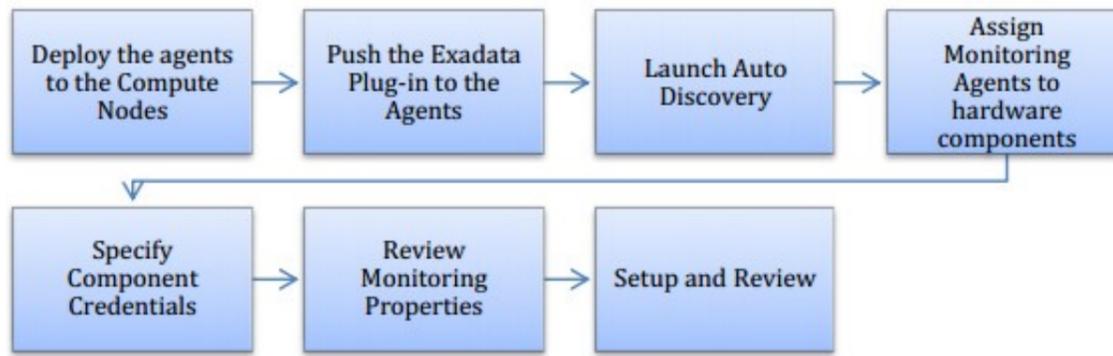
```
# A few find examples
```

```
# \(-name "*.*" \) \
```

```
# \(-name "*vmstat*.*" -o -name "*iostat*.*" \) \
```

```
# \(-name "*vmstat*.*" -o -name "*iostat*.*" \) -mmin -60 \
```

```
# \(-name "*vmstat*.*" -o -name "*iostat*.*" -o -name "*netstat*.*" \) -mtime -8 \
```



NEW QUESTION 10

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 15

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 19

When would be the best time to run an Exadata health check (exachk)?

- A. before patching, before upgrades, before backups, and on a regular basis
- B. after patching, after upgrades, and after backups
- C. only when advised by Oracle Support
- D. before and after patching, when advised by Oracle Support, and on a regular basis
- E. only after a hardware failure
- F. monthly and after a hardware failure

Answer: D

Explanation: #1: Check for updates frequently.
 #2: Execute before & after system changes.
 #3: Make part of regular planned maintenance

NEW QUESTION 24

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 25

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 26

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 27

Which two DML operations will add rows compressed by Hybrid Columnar Compression (HCC) to a table that is created to use HCC?

- A. INSERT
- B. insert with an append hint
- C. UPDATE
- D. CREATE TABLE AS SELECT

Answer: BD

Explanation: To maximize storage savings with Hybrid Columnar Compression, data must be loaded using data warehouse bulk loading techniques. Examples of bulk load operations commonly used in data warehouse environments are:

- * Insert statements with the APPEND hint
- * Parallel DML
- * Direct Path SQL*LDR
- * Create Table as Select (CTAS)

Incorrect:

Not A, Not C: DML operations (INSERT/UPDATE) against a Hybrid Columnar Compressed table/partition can reduce the overall compression savings over time since data INSERTED/UPDATED via DML operations will not be compressed to the same ratio as data that is bulk loaded.

NEW QUESTION 28

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 33

Identify the three values to which the cell_flash_cache attribute can be set.

- A. read
- B. WRITE
- C. KEEP
- D. DEFAULT
- E. NONE

Answer: CDE

Explanation: There are three values the CELL_FLASH_CACHE attribute can be set to.

DEFAULT specifies the cache used for a DEFAULT object is automatically managed as described in the previous section. NONE specifies that the object will never be cached. KEEP specifies the object should be kept in cache.

Note: The CELL_FLASH_CACHE can be assigned to a database table, index, partition and LOB column.

NEW QUESTION 37

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 38

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 cellsvr.
- 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 42

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 47

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 49

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 54

Which two statements appropriately describe the backup and recovery procedures on Exadata?

- A. Cell software backup is automatic and cell recovery is accomplished by using a rescue procedure.
- B. Data stored in a database is backed up and restored by using CellCLI commands.
- C. Database server software backup is automatic and recovery is accomplished by using a rescue procedure.
- D. Data stored in a database is backed up and restored by using RMAN commands.
- E. Cell software backup is performed through CellCLI and cell recovery is accomplished by using a rescue procedure.

Answer: AD

Explanation: A: Your Exadata Storage Servers do not need to be backed up; Exadata automatically backs these up using a self-maintained internal USB drive. D: Backup your Oracle databases using Oracle RMAN

NEW QUESTION 55

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. cellsvr 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. cellsvr 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, cellsvr

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 59

In looking to improve query performance for your Data Warehouse, select the best way that Exadata's Flash Cache feature can be leveraged?

- A. Enable Smart Flash Log.
- B. Execute alter table ... cell_flash_cache=keep on heavily scanned tables.
- C. Enable Write Back Flash Cache.
- D. Create an ASM diskgroup on Flash Cache and move the indexes from disk to Flash.

Answer: B

Explanation: In earlier releases, database administrators had to mark an object as KEEP to have it cached in flash cache for large scan workloads. This feature primarily benefits table scan intensive workloads such as Data Warehouses and Data Marts. Random I/Os such as those performed for Online Transactional Processing (OLTP)

continue to be cached in the flash cache the same way as in earlier releases.

Note: With the Exadata Storage Server Software 11.2.3.3.0 and above, the Exadata Smart Flash Cache software automatically caches objects read, in large table scans, in the flash cache based on how frequently the objects are read. Previously, the default behavior was to bypass the flash cache for such large sequential scans.

NEW QUESTION 64

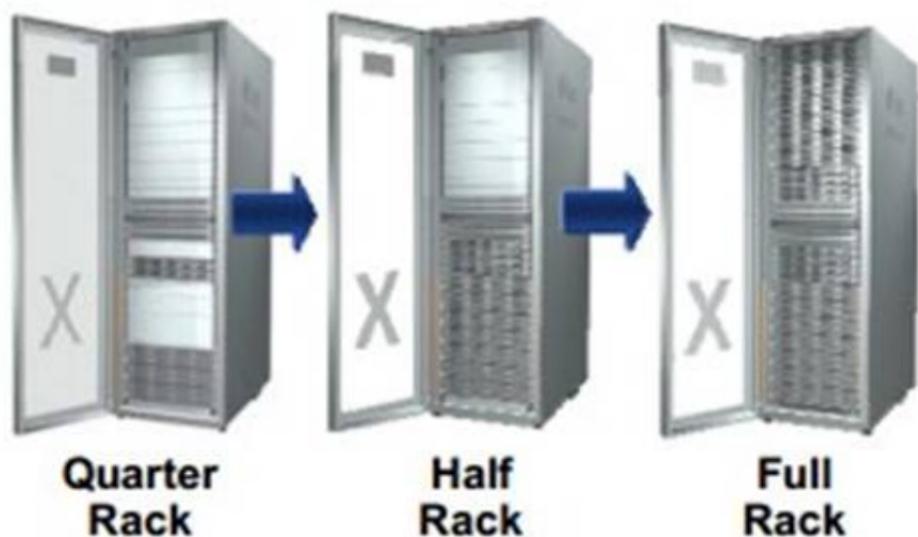
Identify the correct steps required to upgrade a 1/8th rack of Exadata to a 1/4th rack.

- A. Remove the doors of the 1/8th rack. Add the new servers. Cable the Database Servers in the rack. Cable the Exadata Storage Servers in the rack. Close the rack.
- B. Activate all cores by using a software command. Activate all hard drives and flash drives. Create the grid disks. Expand the Oracle ASM diskgroups. Review the new configuration.
- C. Remove the doors of the 1/8th rack. Add an additional Sun Datacenter InfiniBand Switch. Add the new servers. Cable the Database Servers in the rack. Cable the Exadata Storage Servers in the rack. Close the rack. Configure the new hardware.
- D. Activate all cores by using a software command. Activate all hard drives and flash drives.

Answer: B

Explanation: Not A: no need to remove doors. Not C: No additional Infiband Switch needed. Note:

* Database Machine Upgrades



* Database Machine Components and Upgradeability

	Database Machine Full Rack	Database Machine Half Rack	Database Machine Quarter Rack
Database Servers	8	4	2
Exadata Storage Servers	14	7	3
InfiniBand Switches	3	2	2
Upgradability	Connect multiple Full Racks via the included InfiniBand fabric	Field upgrade from Half Rack to Full Rack	Field upgrade from Quarter Rack to Half Rack

NEW QUESTION 67

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)	THROUGH-PUT RATE (GB/SEC)
1	Exadata Cell	14	1 ⁴	14
2	Database Server	8	2.0 ³	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 ⁴ 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 69

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 73

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 cellsv.
- C) Drop the Flash Cache.
- D) Verify that cell for flashcacheMode is writeback.
- E) Create the Flash Cache.
- F) Start up cellsv and activate all grid disks.
- G) Check grdidksk 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 75

Which two statements describe correct network configuration for Exadata Database Machine?

- A. The InfiniBand network subnet manager runs on all database servers to achieve High Availability.
- B. Oracle Clusterware communication is configured to use the management network.
- C. The InfiniBand network interfaces on Linux servers are configured using active-passive bonding.
- D. Database connections to the SCAN listener route through the Ethernet switch in the Exadata rack.
- E. Database servers are deployed with three logical network interfaces configured: management, client access, and private.

Answer: CD

Explanation: Incorrect:

Not A: The InfiniBand switches use an OpenSMInfiniBand subnet manager to manage the switch configuration.

NEW QUESTION 77

How does partition pruning, a 11gR2 database feature, help to improve performance on an Exadata Database Machine?

- A. Partition pruning allows you to reduce a 5 TB I/O operation to a much smaller I/O operation and therefore, enables a much faster return of information.
- B. Using partition pruning can force the optimizer to use that method instead of Smart Sca
- C. Partitioning in Exadata should be used wisely,
- D. Partition pruning is not a supported option in the 11gR2 database code that runs on an Exadata Machine, due to issues with Cost Based Optimizer plans.
- E. Partition pruning allows you to reduce the size of the physical table by reducing fragmentation within the individual tables.
- F. Partition pruning allows you to reduce a 5 TB I/O operation to a larger I/O operation and therefore, forces the use of the Smart Scan technology.

Answer: A

Explanation: Partition pruning allows you to reduce a 5TB I/O operation to a much smaller I/O operation and therefore much faster return of the information.

NEW QUESTION 82

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 84

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 85

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 86

CORRECT TEXT

Which two statements are true about migrating your database to Exadata?

- A. Because Exadata uses InfiniBand, in order to migrate your database to Exadata, you must have InfiniBand on the system that you are migrating from.
- B. Using Data Guard Physical Standby to migrate from an 11.1 database to Exadata is beneficial because it allows you to adopt HCC during migration.
- C. ASM and database best practice configuration supplied during Exadata deployment should be retained during and after migration,
- D. Logical migration methods allow more flexibility than physical methods to change the database structure for best performance.
- E. All indexes should be dropped when migrating to Exadata.

Answer: CD

Explanation: Databases on Exadata use ASM. Incorrect:

Not A: 3 network choices: 10 Gb/s Ethernet

40 Gb/s InfiniBand

1 Gb/s Ethernet (No fibre channel)

NEW QUESTION 88

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 90

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 1z0-485 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 1z0-485 Product From:

<https://www.2passeasy.com/dumps/1z0-485/>

Money Back Guarantee

1z0-485 Practice Exam Features:

- * 1z0-485 Questions and Answers Updated Frequently
- * 1z0-485 Practice Questions Verified by Expert Senior Certified Staff
- * 1z0-485 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * 1z0-485 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year