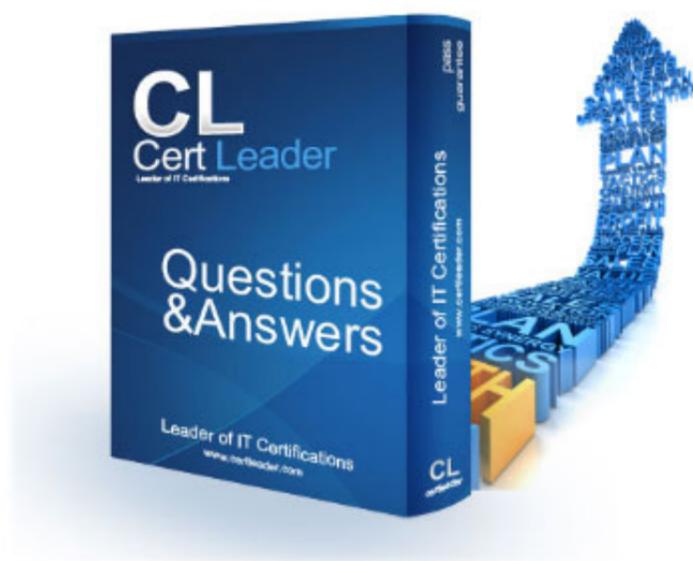


1z0-485 Dumps

Oracle Exadata Database Machine 2014 Implementation Essentials

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



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

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 3

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.

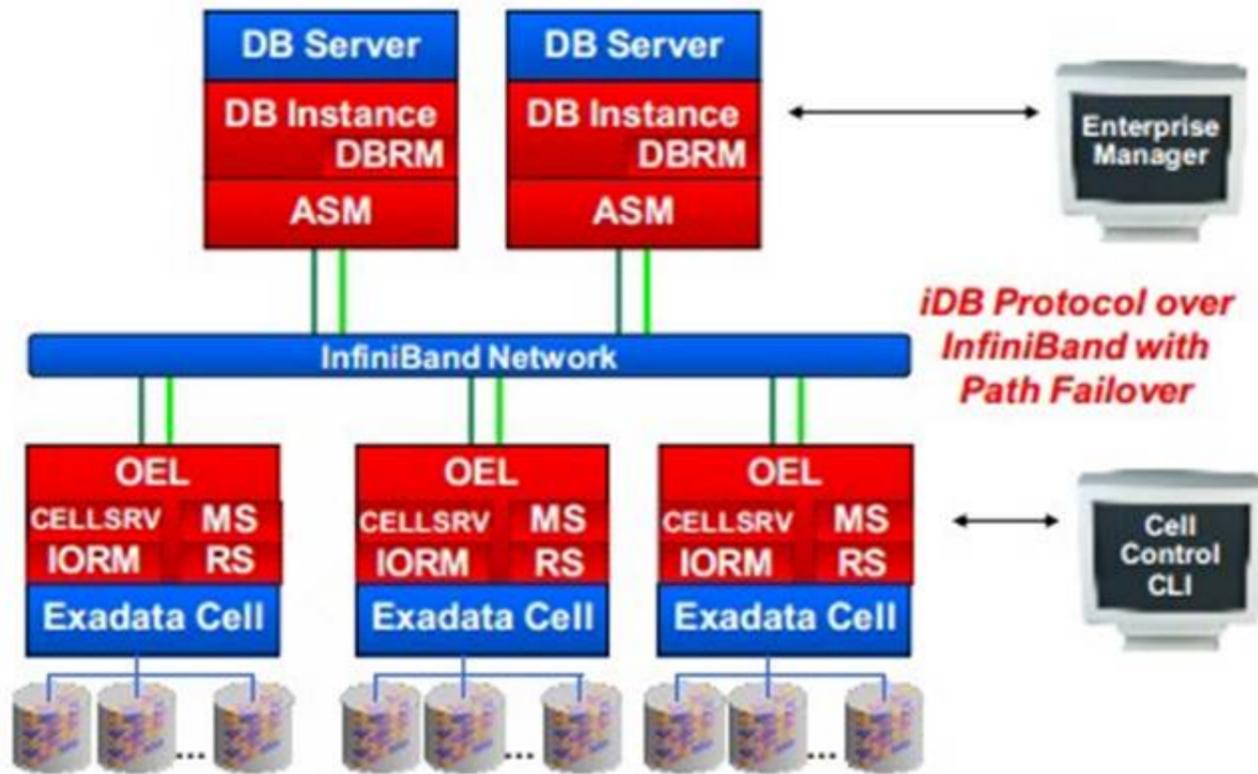
Answer: CE

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 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

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 7

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 8

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 9

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 10

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 10

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 11

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 14

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 19

A control file tracks the physical components of a database. It is the root file that the database uses to find all the other files used by the database. Because of the importance of the control file, Oracle recommends that the control file be ?

- A. size constrained, to prevent size overruns
- B. multiplexed, or have multiple identical copies
- C. stored on an HA USB Drive, to ensure that it does not get overwritten
- D. managed by EM12c, to ensure proper maintenance

Answer: D

NEW QUESTION 23

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 27

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 30

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 31

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 34

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 37

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 38

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 41

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 GgE 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 44

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 49

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 53

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 55

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 59

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 62

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 67

You are concerned about how to recover from a failed Exadata Storage Server. What is the best way to recover from a storage cell failure when content in both system disks is lost?

- A. Database Machine Administrators should be prepared by creating a backup using dbserver_backup.sh.
- B. Database Machine Administrators should be prepared by creating a backup to an NFS file location as described in the owners guide.
- C. Any USB can be plugged in to create a backup.
- D. Be sure that you have a backup by running make_cellboot_usb.sh in /opt/oracle.cellos.
- E. Use the Exadata Storage Server rescue functionality that is provided on the CELLBOOT USB flash drive.
- F. Re-imaging the cell is the only option.

Answer: E

Explanation: Recovering Storage Cells from Loss of a System Volume Using CELLBOOT Rescue

Problem

You have either corrupted your system volume images or suffered from simultaneous loss of the first two disk drives in your Exadata storage cell, and you wish to use the internal CELLBOOT USB drive to recover from failure.

Solution

At a high level, these are the steps you should take:

- * Understand the scope of the failure
- * Contact Oracle Support and open a Service Request
- * Boot your system from the internal CELLBOOT USB image
- * Recover your storage cell using the cell rescue procedure
- * Perform post-recovery steps and validation

NEW QUESTION 69

.....

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>