1Z0-070 Q&As

Oracle Exadata X5 Administration

Pass Oracle 1Z0-070 Exam with 100% Guarantee

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

https://www.certbus.com/1z0-070.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

You plan to migrate a database supporting an OLTP workload to your new X5 Database Machine.

The current database instance supports a large number of short duration sessions and a very high volume of short transactions.

Which three X5 Database Machine features can improve performance for this type of workload?

A. An improved highly efficient undo and redo architecture

B. Faster optimization due to an Exadata-specific optimizer

C. Reduced I/O latency for writes due to writeback flashcache on all Exadata X5 and later models

D. Ultra high I/O performance for reads and writes when using Exadata Extreme Flash in X5 and later models

E. Reduced I/O latency for reads due to read flashcache on all Exadata X5 and later models

Correct Answer: ACD

Explanation:

A: To further accelerate OLTP workloads, the Exadata Smart Flash Cache also implements a special algorithm to reduce the latency of log write I/Os called Exadata Smart Flash Logging.

C: Use the Write-Back Flash Cache feature to leverage the Exadata Flash hardware and make Exadata Database Machine a faster system for Oracle Database Deployments.

D: Exadata X5-2 introduces Extreme Flash Storage Servers. Each Extreme Flash storage server contains eight 1.6 TB state-of-the-art PCI Flash drives. PCI flash delivers ultra-high performance by placing flash memory directly on the high speed PCI bus rather than behind slow disk controllers and directors.

References: http://www.oracle.com/technetwork/database/exadata/exadata-x5-2-ds-2406241.pdf
http://www.oracle.com/technetwork/articles/database/exadata-write-back-flash-2179184.html

QUESTION 2

Which two statements are true about Exadata storage server and database server alerts on an X5 Database Machine?

A. Metric alerts are never stateful.

B. Metric thresholds, if set, will persist across storage and database server reboots.

C. SMTP alert notifications must be sent to both ASR manager and Enterprise Manager agents.

D. SNMP alert notifications can be sent to only one destination.

E. Metrics have no thresholds set on them by default.

Correct Answer: BC

Reference:

https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadata-storage-server-software-introduction.html#GUID-3E48425A-AB8A-4E62-80C4-BACA65A1F8D3

https://docs.oracle.com/cd/E91266_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG304

https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadata-storage-server-software-introduction.html#GUID-3E48425A-AB8A-4E62-80C4-BACA65A1F8D3
https://docs.oracle.com/cd/E91266_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG304

QUESTION 3

Which three statements are true concerning InfiniBand port and subnet monitoring on an X5 Database Machine?

A. The Infiniband subnet master location can be determined by using the getmaster command run on any database server.

B. The Infiniband subnet master location can be determined by using the getmaster command run on an IB switch.

C. The InfiniBand port status may be displayed on the storage servers by using the LIST IBPORT command in the CELLCLI utility.

D. The InfiniBand port status may be displayed on the Infiniband switches by using the LIST IBPORT command in the DBMCLI utility.

E. The InfiniBand port status may be displayed on the database servers by using the LIST IBPORT command in the DBMCLI utility.

F. The InfiniBand port monitoring is automatic on the database servers and is managed by Enterprise Manager.

Correct Answer: BCE

Explanation:

B: From any InfiniBand switch in the network (leaf switch or spine switch), log in as root and run the getmaster command to obtain the location of the master SM as follows:

getmaster

This command displays the host name or IP address and the IP address of the switch where the master SM is running.

C: CellCLI> list ibport - Will display InfiniBand configuration details

E: The DBMCLI utility is the command-line administration tool for configuring database servers, and

managing objects in the server environment.

The LIST IBPORT command displays attributes for InfiniBand ports determined by the specified attributes

and filters.

References:

http://docs.oracle.com/cd/E80920_01/DBMMN/exadata-dbmcli.htm
https://docs.oracle.com/cd/E18476_01/doc.220/e18478/GUID-9FF8B5B0-3481-4B73-89D3-108CBD7EB989.htm

http://docs.oracle.com/cd/E80920_01/DBMMN/exadata-dbmcli.htm
https://docs.oracle.com/cd/E18476_01/doc.220/e18478/GUID-9FF8B5B0-3481-4B73-89D3-108CBD7EB989.htm

QUESTION 4

Identify three valid reasons for creating multiple griddisks on a single harddisk-based celldisk.

A. to implement storage realms so that storage that can be reserved for specific resource consumer groups in the same database

B. to enable the creation of normal or high redundancy ASM diskgroups

C. to segregate storage into multiple pools with different performance characteristics

D. to enable disk mirroring for the system area

E. to implement storage realms so that storage can be reserved for specific databases

F. to implement storage realms so that storage that can be reserved for specific Grid Infrastructure clusters

Correct Answer: BCF

Explanation:

Creating multiple grid disks per cell disk allows you to create multiple pools of storage on the same

Exadata Storage Server. The multiple grid disks can be assigned to separate ASM diskgroups, which can

be provisioned to different databases.

Note: Griddisk is a logical disk that can be created on a celldisk. In a standard Exadata deployment we

create griddisks on hard disk based celldisks only. While it is possible to create griddisks on flashdisks, this is not a standard practice.

F: After you complete the cell configuration, you can perform the following optional steps on the storage cell:

1.

Add the storage cell to the Exadata Cell realm

2.

Configure security on the Oracle Exadata Storage Server grid disks

References: https://docs.oracle.com/cd/E80920_01/SAGUG/SAGUG.pdf

QUESTION 5

You are planning your deployment of Enterprise Manager to monitor all the components of an X5 Database Machine.

A part of the requirement is to provide for high availability of the monitoring infrastructure.

If the host running the agent that has Database Machine targets bound to it fails, the monitoring of these

targets must be done by another agent.

Which three statements are true regarding the configuration used to support this requirement?

A. Database Machine plug-ins must be deployed to at least two Enterprise Manager agents.

B. Fail back to the original agent when the host is restarted is done automatically.

C. Fail over to any secondary agent is done automatically.

D. Fail over to any secondary agent must be done manually.

E. Database machine plug-ins must be deployed to all Enterprise Manager agents.

F. Fail back to the original agent when the host is restarted must be done manually.

Correct Answer: BCE

QUESTION 6

You issued these commands to all Exadata Storage Servers in an X6 Exadata Database Machine using dcli:

alter iormplan objective = low_latencyalter iormplan active

There are no database or category plans defined.

You are encountering disk I/O performance problems at certain times, which vary by day and week.

DSS and Batch workloads perform well some of the time.

Further investigation shows that at times, the workloads are all OLTP I/Os, at other times all batch I/Os, and sometimes a bit of each.

You wish to have disk I/O managed so that performance will be optimized for all workloads.

Which statements would you issue to all Exadata Storage Servers to achieve this?

A. alter iormplan objective=high_throughput

B. alter iormplan objective=balanced

C. alter iormplan objective=low_latency

D. alter iormplan objective=auto

E. alter iormplan objective=‘ ’

Correct Answer: D

Explanation:

The supported IORM objectives are auto, low_latency, balanced, and high_throughput. The recommended objective option is auto which allows IORM to continuously monitor the workloads, and select the best mode based on the active workloads currently on the cells.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-iorm.htm

QUESTION 7

Which two must be true for a Smart Scan to occur on a table?

A. Sessions querying the table must set cell_offload_processing = true.

B. It must be stored in an ASM diskgroup with a 4 MByte AU size.

C. It must be heap organized.

D. It must be accessed using direct path reads.

E. The table must not be compressed.

Correct Answer: AD

QUESTION 8

Identify two valid reasons for executing an X5 Exadata storage server rescue procedure.

A. Accidental loss of all data from all griddisks in a storage server

B. Corruption in the /(root) filesystem

C. Corruption in a normal or high redundancy ASM diskgroup

D. The failure of both physical disks 0 and 1

E. Only the failure of physical disk 1

F. Only the failure of physical disk 0

Correct Answer: BD

Explanation:

The rescue procedure is necessary when system disks fail, the operating system has a corrupt file system, or there was damage to the boot area. If only one system disk fails, then use CellCLI commands to

recover. In the rare event that both system disks fail simultaneously, you must use the Exadata Storage

Server rescue functionality provided on the Oracle Exadata Storage Server Software CELLBOOT USB

flash drive.

Incorrect Answers:

E, F: If only one system disk fails, then use CellCLI commands to recover.

References: http://docs.oracle.com/cd/E80920_01/DBMMN/maintaining-exadata-storageservers.htm#GUID-710814E7-4691-49EE-95AD-726D2D6C5BFE

QUESTION 9

As part of planning for database consolidation, you grouped your databases into two categories based on different technical and business objectives.

Which three statements are true about possible configurations for your Exadata X6 Database Machine fabric?

A. The storage grid may be partitioned when deploying a single-rack Database Machine configuration using virtualization.

B. The database grid may be partitioned when deploying a multirack Database Machine configuration.

C. A single database cluster benefits from accessing multiple storage grids.

D. Multiple database clusters may access the same storage grid.

E. Multiple database clusters never benefit from accessing a single storage grid.

Correct Answer: BCD

QUESTION 10

Which three are true about Exadata Smart Flash Log?

A. Databases on the Database Machine use Exadata Smart Flash Log by default.

B. I/O Resource Manager database plans can be used to enable or disable Exadata Smart Flash Log for individual databases.

C. LGWR will not wait for writes to Exadata Smart Flash Log if the write to a disk-based logfile completes first.

D. I/O Resource Manager category plans can be used to enable or disable Exadata Smart Flash Log for different I/O categories.

E. The use of Exadata Smart Flash Logs is mandatory for support of production databases.

Correct Answer: ABC

Explanation:

A: Exadata I/O Resource Manager (IORM) has been enhanced to enable or disable Smart Flash Logging for the different databases running on the Database Machine, reserving flash for the most performance critical databases.

B: The Exadata I/O Resource Manager (IORM) has been enhanced to enable or disable Smart Flash Logging for the different databases running on the Database Machine.

C: Smart Flash Logging works as follows. When receiving a redo log write request, Exadata will do parallel writes to the on-disk redo logs as well as a small amount of space reserved in the flash hardware. When either of these writes has successfully completed the database will be immediately notified of completion. If the disk drives hosting the logs experience slow response times, then the Exadata Smart Flash Cache will provide a faster log write response time. Conversely, if the Exadata Smart Flash Cache is temporarily experiencing slow response times (e.g., due to wear leveling algorithms), then the disk drive will provide a faster response time. This algorithm will significantly smooth out redo write response times and provide overall better database performance.

Incorrect Answers:

D: Category plans are configured and enabled using the CellCLI utility on the cell. Only one category plan can be enabled at a time.

References: http://structureddata.org/2011/10/12/exadata-smart-flash-logging-explained/

QUESTION 11

You installed ASR Manager on a stand-alone server and configured Auto Service Request (ASR) for your X5 Database Machine and its assets.

Which three statements are true about this configuration?

A. Simple Network Management Protocol (SNMP) traps are used to send notifications from database servers to the ASR Manager.

B. Simple Network Management Protocol (SNMP) traps are used to send notifications from storage servers to the ASR Manager.

C. When a component fault occurs, fault telemetry is securely transmitted to Oracle via Simple Network Management Protocol (SNMP).

D. When a component fault occurs, fault telemetry is securely transmitted to Oracle via HTTPS.

E. Simple Network Management Protocol (SNMP) traps are used to send notifications from the Enterprise Manager to the ASR Manager.

F. Simple Network Management Protocol (SNMP) traps received by ASR Manager are forwarded to the Enterprise Manager.

Correct Answer: BCF

Explanation:

B: Oracle ASR Manager only processes SNMP traps that are sent from IP addresses that Oracle ASR Manager recognizes. Example of Exadata Storage Server SNMP Trap This example shows the SNMP trap for an Exadata Storage Server disk failure. The corresponding hardware alert code has been highlighted.

2011-09-07 10:59:54 server1.example.com [UDP: [192.85.884.156]:61945]: RFC1213-MIB::sysUpTime.0 = Timeticks: (52455631) 6 days, 1:42:36.31 SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: SUN-HW-TRAP-MIB::sunHwTrapHardDriveFault SUN-HW-TRAP-MIB::sunHwTrapSystemIdentifier = STRING: Sun Oracle Database Machine Etc.

C (not D): The ASR Manager uses the SNMP GET protocol to query ASR assets for additional fault

information.

To configure fault telemetry, choose one of the following three options:

Add SNMP Trap Destinations Using OneCommand (recommended for new installations)

Add SNMP Trap Destinations for Multiple Servers Using the dcli Utility

Add SNMP Trap Destinations for a Single Server

References:

http://docs.oracle.com/cd/E80920_01/ASXQI/toc.htm

https://docs.oracle.com/cd/E37710_01/install.41/e18475/ch5_troubleshooting.htm#ASRUD331

QUESTION 12

Examine this IORM plan: Which two are true concerning this plan?

```
CellCLI> list iormplan detail

name:       dmorlcel_IORMPLAN
catPlan:    name= interactive, level=1, allocation=90
            name= batch, level=2, allocation=80
            name=maintenance, level=3, allocation=50
            name=other, level=3, allocation =50
dbPlan:     name=sales, level=1, allocation =45, limit=60,
            flashcache=on, flashlog=on
            name= finance, level=1, allocation=45, limit=60,
            flashcache=on, flashlog=off
            name=other, level=1, allocation=10,
            flashcache=off, flashlog=on
objective:  balanced
status:     active
```

A. The Finance database can use at least 45%, but never more than 60%, of the total Flash Cache capacity.

B. I/Os from the finance database are guaranteed to get a minimum of 45% of the I/O bandwidth if the demand exists,

and a maximum of 60% of the I/O bandwidth even if no other databases are doing I/O to the cell, and the demand from the finance database exceeds 60% of the maximum I/O rate of the cell.

C. If I/Os come from the HR database only, then they may get up to 100% of the I/O bandwidth on the cell.

D. I/Os from the finance database are guaranteed to get 45% of the I/O bandwidth if the demand is at least 60% of the maximum I/O rate of the cell, but may get 100% of the I/O bandwidth if no other databases are performing I/O to the cell.

E. If I/Os come from the HR database only, then they may get up to 10% of the I/O bandwidth on the cell.

Correct Answer: BE

Explanation: The IORM plan can be configured using the ALTER IORMPLAN command on command-line interface (CellCLI) utility on each Exadata storage cell. It consists of two parameters - dbplan and catplan. While the "dbplan" is used to create the I/O resource directives for the databases, the "catplan" is used to allocate resources by workload category consolidated on the target system.

allocation/share - Specify the resource allocation to a specific database in terms of percentage or shares. limit - Specify maximum limit of disk utilization for a database.

Incorrect Answers:

A: IORM plans configures % I/O resources, not % of Flash Cache.

References: https://community.oracle.com/docs/DOC-998939

QUESTION 13

You are evaluating the performance of a SQL statement that accesses a very large table.

You run this query: Identify two reasons why the “physical read total bytes” statistic is greater than the “cell physical IO bytes eligible for predicate offload” statistic.

```
SQL> SELECT  s.name, m.value/1024/1024 MB FROM V$SYSSTAT s, V$MYSTAT m
2 WHERE s.statistic# = m.statistic# AND
3 (s.name LIKE 'physical%total bytes' OR s.name LIKE 'cell phys%'
4 OR s.name LIKE 'cell IO%');
```

NAME	MB
physical read total bytes	19047.2266
physical write total bytes	0
cell physical IO interconnect bytes	4808.85828
cell physical IO bytes pushed back due to excessive CPU on cell	0
cell physical IO bytes saved during optimized file creation	0
cell physical IO bytes saved during optimized RMAN file restore	0
cell physical IO bytes eligible for predicate offload	18005.6953
cell physical IO bytes saved by storage index	0
cell physical IO interconnect bytes returned by smart scan	3767.32703
cell IO uncompressed bytes	18005.6953

A. There is an index on the column used in the WHERE clause, causing “cell multiblock physical reads” to be requested by the database instance, resulting in additional I/O.

B. The table is an IOT and has an overflow segment, causing “cell multiblock physical reads” to be requested by the database instance, resulting in additional I/O.

C. There is an uncommitted transaction that has modified some of the table blocks, causing some “cell single block physical reads” to be requested by the database instance, resulting in additional I/O.

D. The table is an index clustered table, causing “cell single block physical reads” to be requested by the database instance, resulting in additional I/O.

E. There are migrated rows in the table, causing some “cell single block physical reads” to be requested by the database instance, resulting in additional I/O.

Correct Answer: BE

Explanation: Note:

1.

physical read total bytes: the size of the segment to read is known by the database, and must be read entirely from the database\\'s perspective.

2.

cell physical IO bytes eligible for predicate offload: this statistic shows the amount of data which the cell server is able to process on behalf of the database, instead of the database processing and the cell server just delivering blocks.

3.

Cell physical IO bytes eligible for predicate offload --- This number should be high

4.

The higher the number more MB/GB is filtered out at the cell level itself rather sending it to the buffer cache to filter the rows.

5.

In this case, all bytes are processed on the cellserver (cell physical IO bytes eligible for predicate offload=physical read total bytes)

Cell Offloading:

The storage cells are intelligent enough to process some workload inside them, saving the database

nodes from that work. This process is referred to as cell offloading.

QUESTION 14

Which two statements are true about Database Server Metrics and Alerts on an X5 Database Machine?

A. Database Server Metrics are collected my MMON.

B. Database Server Metrics are collected by the MS process and stored in an internal disk-based repository.

C. Database Server hardware sensor alerts are detected by the MS process on the database server.

D. Database Server Metrics are collected by the RS process.

E. Database Server Threshold Alerts may be forwarded by the MS process as SNMP traps.

Correct Answer: BE

QUESTION 15

You wish to determine if the I/O resource management plan that you created has helped improve the performance of OLTP category I/Os on your X6 Exadata Database Machine.

You decide to examine the relevant metrics on all the cells, to see whether the I/O rate has improved for this category compared to last week, and whether waits and wait time have been reduced.

You issue this command on the first cell:

```
cellcli -e list metriccurrent attributes name, metricValue, metricType
where objectType = IORM_CATEGORY and and merticObjectName = OLTP and name
like 'CT_IO.*'
```

You examine the output from the first cell which contains:

CT_IO_BY_SEC	0 MB/sec	Instantaneous
CT_IO_LOAD	1	Instantaneous
CT_IO_RQ_LG	1,172 IO requests	Cumulative
CT_IO_RQ_LG_SEC	0.0 IO/sec	Rate
CT_IO_RQ_SM	360,325 IO requests	Cumulative
CT_IO_RQ_SM_SEC	22.7 IO/sec	Rate
CT_IO_UTIL_LG	0%	Instantaneous
CT_IO_UTIL_SM	0%	Instantaneous
CT_IO_WT_LG	0 ms	Cumulative
CT_IO_WT_LG_RQ	0.0 ms/ request	Rate
CT_IO_WT_SM	0 ms	Cumulative
CT_IO_WT_SM_RQ	0.0 ms/ request	Rate

Which two sets of metrics would you use to determine whether the I/O performance has improved for the OLTP category for the duration of the one-hour measurement period?

A. CT_IO_RQ_SM, CT_IO_RQ_LG, CT_IO_RQ_SM_SEC, and CT_IO_RQ_LG_SEC

B. CT_IO_UTIL_SM and CT_IO_UTIL_LG

C. CT_IO_RQ_SM_SEC and CT_IO_RQ_LG_SEC D. CT_IO_WT_SM, CT_IO_WT_LG, CT_IO_WT_SM_RQ, and CT_IO_WT_LG_RQ

Correct Answer: D

Explanation: CT_IO_RQ_SM The cumulative number of small I/O requests issued by the category for hard disks. A large value indicates a heavy I/O workload from this category. CT_IO_RQ_LG The cumulative number of large I/O requests issued by the category for hard disks. A large value indicates a heavy I/O workload from this category. CT_IO_WT_SM_RQ The average IORM wait time per request for small I/O requests issued to hard disks by an IORM category. CT_IO_WT_LG_RQ The average IORM wait time per request for large I/O requests issued to hard disks by an IORM category.

Incorrect Answers:

A: CT_IO_RQ_SM_SEC

This metric is derived from CT_IO_RQ_SM. It specifies the rate of small I/O requests issued by the

category for hard disks. Its units are number of I/O requests per second. A large value indicates a heavy I/

O workload from this category in the past minute.

B: CT_IO_UTIL_SM

The percentage of disk resources utilized by small requests from this category.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-monitoring.htm