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.pass2lead.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

Which three are true concerning Storage Indexes?

A. A maximum of eight table columns for any table are indexed per storage region.

B. The use of Storage Indexes for particular categories of I/O can be displayed by using an I/O Resource Manager Category Plan.

C. Storage Indexes persist across any Exadata storage server reboots.

D. The use of Storage Indexes for a particular database can be disabled by using an I/O Resource Manager Database Plan.

E. A Storage Index is automatically maintained by CELLSRV based on the filter columns of the offloaded SQL.

F. Different storage regions may have different columns indexed for the same table.

Correct Answer: ADE

Explanation:

A: Each disk in the Exadata storage cell is divided into equal sized pieces called storage regions (default 1MB). There is an index entry for every storage regions (1MB of data stored on disk). Each entry contains the minimum and maximum value for columns seen in ‘where’ clause predicates. Information for up to 8 columns can be stored. The index is then used to eliminate disk IO by identifying which storage regions don\\'t match the ‘where’ clause of a query.

QUESTION 2

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 3

Which two statements are true regarding the use of Auto Service Request (ASR) with an X6 Database Machine?

A. The database server ILOMs must use SMTP over the management network for notifications to ASR Manager.

B. The database server ILOMs must have SNMP traps configured to use the management network for notifications to ASR Manager.

C. The storage server ILOMs must have SNMP traps configured to use the management network for notifications to ASR Manager.

D. The database server MS process must have SNMP traps configured to use the management network for notifications to ASR Manager.

Correct Answer: BC

Explanation:

B: Database Server ILOM plug-in

Monitoring databases and their instances, ASM environments, the Grid Infrastructure, and the host software environment are done by Enterprise Manager in the usual way as these are standard targets. But monitoring the hardware for the database servers requires the ILOM plug-in, as there is no Management Server (MS) on the database servers to receive SNMP traps from the ILOM. The plug-in will receive sensor state and availability data from the ILOM including alerts based on pre-set ILOM thresholds.

C: Exadata Storage Server plug-in extends the monitoring of exadata cells in addition to providing a GUI interface. The plug-in uses an SSH connection to the cellmonitor user on the cells and uses list commands only. This is for interactive monitoring. One may also set thresholds using the plug-in which are distinct from any thresholds set using cellcli utility as the celladmin user. For alerts to be sent to the plug-in, SNMP traps are used as follows:

Cell ILOM alerts are sent to the cell Management Server (MS) via an SNMP trap. The MS then send

SNMP notifications onward to the plug-in.

Cell alerts flagged by MS itself, such as cell thresholds being exceeded, or ADR software alerts, are sent

to the plug-in using SNMP.

References:

https://dbatrain.wordpress.com/2011/06/

http://docs.oracle.com/cd/E21659_01/html/E21660/z40015671004046509.html

QUESTION 4

Which two statements are true about the X5 Exadata storage server rescue procedure?

A. The rescue procedure can be executed from the CELLBOOT USB flash drive.

B. An Exadata storage server automatically enters the rescue environment when it cannot boot from the system area.

C. The rescue procedure can be used to repair corruption in an ASM diskgroup.

D. The rescue procedure can be used to restore a corrupt system area.

E. The rescue procedure must be used to recover from a failed Exadata storage server software upgrade.

Correct Answer: AD

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.

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

QUESTION 5

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

alter iormplan objective = offalter iormplan active

There are no database or category plans defined.

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

DSS and batch workloads perform well, but your OLTP workloads suffer poor response times when running at the same time as these other workloads.

You must ensure that DSS and batch workloads retain relatively high throughput without causing excessive degradation of OLTP performance.

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

A. alter iormplan objective = low_latency

B. alter iormplan objective = high_throughput

C. alter iormplan objective = balanced

D. alter iormplan objective = auto

Correct Answer: A

Explanation:

When the objective option is set to basic, the database resource plan maximum utilization limits are not

enforced. For stricter plan conformance, and enforcement of maximum utilization limits, the objective

option must be set to something other than basic. The supported IORM objectives are auto, low_latency,

balanced, and high_throughput.

The objective may also be changed to low_latency in which case the latency for critical I/Os is very good

but there is a significant degradation in scan throughput when both workloads are running concurrently.

Incorrect Answers:

B: The objective can be changed to high_throughput which will increase the scan throughput at the cost of critical I/O latency.

C, D: 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. The objective values of auto and balanced have the same behavior.

QUESTION 6

Which three statements are true concerning the configuration of SNMP on an Exadata x5 or x6 Database Machine?

A. All SNMP notifications from the storage server ILOMs are sent directly to the Automatic Service Request (ASR) manager.

B. Some SNMP notifications may be configured using the alter cell command on the storage servers.

C. The Database Machine internal Cisco Switch communicates directly with the Enterprise Manager agent via SNMP.

D. Some SNMP notifications from a storage server ILOM are sent directly to the management server (MS) process on the same storage servers.

E. All SNMP notifications from a database server ILOM are sent directly to the Automatic Service Request (ASR) manager.

F. All SNMP notifications from a database server ILOM are sent directly to the management server (MS) process on the same database server.

Correct Answer: BDE

Explanation:

B: You can configure SNMP Trap Destinations for the Exadata Storage Servers with the following

command:

cellcli -e “alter cell snmpSubscriber=(host =\\'ASR-Manager-name-or-IPaddress\\',port=162,community=public,type=asr)”

D: Exadata Storage Server plug-in

This extends the monitoring of exadata cells in addition to providing a GUI interface. The plug-in uses an

SSH connection to the cellmonitor user on the cells and uses list commands only. This is for interactive

monitoring. One may also set thresholds using the plug-in which are distinct from any thresholds set using

cellcli utility as the celladmin user. For alerts to be sent to the plug-in, SNMP traps are used as follows:

Cell ILOM alerts are sent to the cell Management Server (MS) via an SNMP trap. The MS then send

SNMP notifications onward to the plug-in.

Cell alerts flagged by MS itself, such as cell thresholds being exceeded, or ADR software alerts, are sent

to the plug-in using SNMP.

E: Database Server ILOM plug-in

Monitoring databases and their instances, ASM environments, the Grid Infrastructure, and the host software environment are done by Enterprise Manager in the usual way as these are standard targets. But monitoring the hardware for the database servers requires the ILOM plug-in, as there is no Management Server (MS) on the database servers to receive SNMP traps from the ILOM. The plug-in will receive sensor state and availability data from the ILOM including alerts based on pre-set ILOM thresholds.

References:

https://dbatrain.wordpress.com/2011/06/http://docs.oracle.com/cd/E21659_01/html/E21660/z40015671004046509.html

QUESTION 7

Which four statements are true about the configuration of Auto Service Request (ASR) for use with an X5 Database Machine?

A. ASR Manager software may be installed on one of the database servers in the Database Machine.

B. ASR can connect to the Internet using Simple Network Management Protocol (SNMP), using a proxy server.

C. ASR Manager can connect to the Internet directly with HTTPS.

D. ASR Manager can connect to the Internet with HTTPS using a proxy server.

E. ASR can connect to the Internet directly using Simple Network Management Protocol (SNMP).

F. ASR Manager software may be installed on a server that is not part of the Database Machine.

Correct Answer: ACDF

Explanation: C: If you need to use HTTPS for security purposes, you can set up HTTPS/SSL for the ASR Manager HTTP receiver.

D: As part of the registration process for ASR Manager, you can optionally set the ASR Manager to access the internet through a proxy server.

References: https://docs.oracle.com/cd/E37710_01/install.41/e18475/ch2_asr_manager.htm#ASRUD137

QUESTION 8

You plan to migrate a very large database supporting a DSS workload to your new X5 Database Machine.

It will be the only database on this full rack.

Which three statements are true about Database Machine features that improve performance for the DSS

workload?

A. Smart Storage operations can improve the performance of joins.

B. Smart Storage operations can improve the performance of scans.

C. Hybrid Columnar Decompression overheads can be offloaded from the database servers for index full scans.

D. Full table scan operations can improve due to the default Smart Flash Cache implementation.

E. Hybrid Columnar Compression can reduce the amount of physical I/O required to scan large tables.

Correct Answer: BCD

Reference: http://www.informit.com/articles/article.aspx?p=2418151&seqNum=3

Reference: http://www.informit.com/articles/article.aspx?p=2418151andseqNum=3

QUESTION 9

Which statement is true about operating systems in an X5 Database Machine multirack configuration consisting of two full racks and one Exadata storage expansion rack?

A. All Exadata storage servers used by the same virtual cluster nodes must run the same O/S but Exadata Storage Servers in different clusters may run different operating systems.

B. All Exadata storage servers must run the Oracle Solaris O/S and all database servers within the same cluster must run Oracle Linux.

C. All Exadata storage servers may run Oracle Virtual Machine (OVM).

D. All Exadata storage servers must run Oracle Linux.

E. All Exadata storage servers must run the Oracle Linux O/S and all database servers within the same cluster must run the same version of Oracle Virtual Machine (OVM).

Correct Answer: D

Explanation:

On both physical and virtual deployments, Exadata systems use minimal Linux distributions to ensure that

just the RPMs needed to run Oracle database, are installed and enabled.

References: http://www.oracle.com/technetwork/database/exadata/exadata-x5-2-ds-2406241.pdf

QUESTION 10

Which three statements are true about Recovery Manager (RMAN) daily differential incremental backup strategies on an X5 Database Machine for a database having 25% or more of its blocks modified each day and which has an 8 k block size?

A. Fast incremental backups when 50% or more of the blocks have changed since the last backup, will run as slowly as normal incremental backup.

B. Enabling Block Change Tracking (BCT) on the database can result in reduced consumption of storage network bandwidth.

C. Enabling Block Change Tracking (BCT) on the database can result in a reduction of physical I/O on the cells during incremental backups.

D. For level-1 backups, Block Change Tracking (BCT) is most beneficial when more than 25 percent of the blocks have changed since the last backup.

E. For level-0 backups, Block Change Tracking (BCT) is most beneficial when more than 25 percent of the blocks have changed since the last backup.

F. cellsrv returns only blocks that have changed since the last backup.

Correct Answer: ACF

Explanation:

A: Fast Incremental backups is possible with Block change tracking, which is initially introduced from version 10.2 onwards, by this tool it\\'s very useful to reduce the RMAN incremental backup duration. If the changes are something around 20% then in this situation BCT helps a lot.

C: Exadata Storage Server offload capability combined with RMAN block change tracking will efficiently perform large

I/Os at the storage-tier level, returning only individual changed blocks for incremental backups and increasing the backup performance of the system.

Note: Level 1 backup: A level 1 backup includes only those blocks that have been changed since the "parent" backup was taken. Remember a parent backup may be either a level 0 or a level 1 backup. Block change tracking allows indeed the highest benefit for databases where the changes are not so high,

Level 0 backup: A level 0 incremental backup is physically identical to a full backup and it includes every data block in the file except empty blocks. The only difference is that the level 0 backup is recorded as an incremental backup in the RMAN repository, so it can be used as the parent for a level 1 backup.

References: http://www.dba-oracle.com/t_rman_backup_types.htm
http://www.oracle.com/technetwork/database/availability/maa-tech-wp-sundbm-backup-11202-183503.pdf
https://www.toadworld.com/platforms/oracle/w/wiki/11124.fast-incremental-backups-active-data-guard

QUESTION 11

Which four statements are true concerning the configuration or use of Enterprise Manager Cloud Control to monitor and manage Exadata Database Machine components?

A. Computer nodes forward their SNMP traps to the Management Server process on the same computer node.

B. Storage nodes forward their SNMP traps to the Management Server process on the same storage node.

C. Cloud control can monitor and manage a Database Machine compute node expansion rack.

D. Oracle management agents must only be installed on one storage node in a Database Machine to enable monitoring and management of all storage nodes.

E. Cloud Control cannot monitor partitioned Exadata Database Machine systems.

F. Cloud Control can monitor and manage a Database Machine storage Expansion rack.

G. Oracle management agents must only be installed on one compute node in a Database Machine to enable monitoring and management of all compute nodes and storage nodes.

Correct Answer: BCFG

Explanation:

B (not A): Oracle Exadata components—including database and storage servers, switches, and power

distribution units (PDUs)—use SNMP to raise alerts and report monitoring information.

Oracle\\'s CellCLI interface is a Java-based framework delivered by the storage cell\\'s management server

process (MS) that provides administrative capabilities to your storage server entities.

C: Cloud control can monitor the Exadata Database Machine.

Figure: Oracle Enterprise Manager Cloud Control 12c monitoring architecture. Note the Management Server process (MS) on the Storage servers.

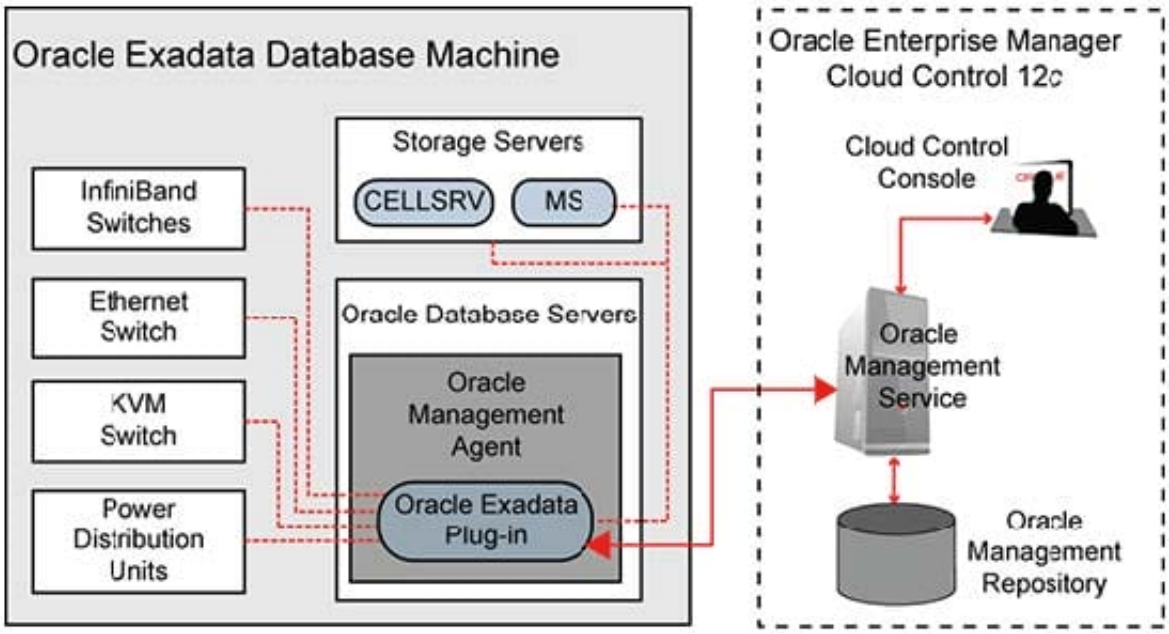

G (not D): Before using Oracle Enterprise Manager Cloud Control 12c with Oracle Exadata, an Oracle Management Agent and Oracle Exadata plug-in must be installed on every Oracle Exadata database server (compute node). This agent monitors software targets, such as the database instances and Oracle Clusterware resources, on the database servers. The plug-in enables monitoring of other hardware components in Oracle Exadata, including the storage servers, switches, and power distribution units.

References: http://www.oracle.com/technetwork/articles/servers-storage-admin/monitor-exadata-em122291964.html

QUESTION 12

Which four are true about Exadata features?

A. Storage Indexes persist across Exadata storage server reboots.

B. Data Warehouse workloads will benefit from Smart Flash Cache configured in Write-Through mode.

C. Hybrid Columnar Compressed tables can be compressed and decompressed on Exadata storage servers.

D. Hybrid Columnar Compressed tables can be compressed and decompressed on the database servers.

E. OLTP workloads will benefit from Smart Flash Cache configured in Write-Back mode.

F. Storage Indexes persist across database server reboots.

Correct Answer: CDEF

Explanation:

CD: That data remains compressed not only on disk, but also remains compressed in the Exadata Smart

Flash Cache, on Infiniband, in the database server buffer cache, as well as when doing back-ups or log

shipping to Data Guard.

E (not B): 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. Write-through cache mode

is slower than write-back cache mode. However, write-back cache mode has a risk of data loss if the

Exadata Storage Server loses power or fails.

F: Storage indexes are not stored on disk; they are resident in the memory of the storage cell servers.

Incorrect Answers:

A: Storage indexes are not stored on disk; they are resident in the memory of the storage cell servers. They are created automatically after the storage cells receive repeated queries—with predicates-for columns. No user intervention is needed to create or maintain storage indexes. And because they are memory-resident structures, they disappear when the storage cells are rebooted.

References:

QUESTION 13

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 14

You plan to monitor the ASM configuration on an X5 Database Machine as part of your role supporting Exadata-based ASM diskgroups.

You want to check for potential space problems that take ASM mirroring requirements into account.

Which two values would you monitor from V$ASM_DISKGROUP or by using the ASMCMD LSDG command?

A. cold_used_mb

B. total_mb

C. required_mirror_free_mb

D. free_mb

E. usable_file_mb

Correct Answer: BD

Explanation:

Determine the Amount of Available Space

To increase the size of the disks in a disk group you must either have unallocated disk space available, or

you have to reallocate space currently used by a different disk group.

Example: View the space currently used by the disk groups.

SELECT name, total_mb, free_mb, total_mb - free_mb used_mb, round(100*free_mb/total_mb,2) pct_free

FROM v$asm_diskgroup

ORDER BY 1;

```
SELECT name, total_mb, free_mb, total_mb - free_mb used_mb, round(100*free_mb/total_mb
FROM v$asm_diskgroup
ORDER BY 1;

NAME                            TOTAL_MB    FREE_MB    USED_MB   PCT_FREE
------------------------------ ---------- ---------- ---------- ----------
DATAC1                           68812800    9985076   58827724      14.51
RECOC1                           94980480   82594920   12385560      86.96
```

The example above shows that the DATAC1 disk group has only about 15% of free space available while the RECOC1 disk group has about 87% free disk space. The PCT_FREE displayed here is raw free space, not usable free space. Additional space is needed for rebalancing operations.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-administeringasm.htm#SAGUG20526

QUESTION 15

You are checking the status of the ports on one of the InfiniBand switches on your X5 Database Machine.

You run the getportstatus command as shown:

[root@exampsw-ib2 bin]# getportstatus 17bPort status for connector 17B Switch Port 2 Adminstate:.....Enabled LinkWidthEnabled: .....1X or 4X LinkWidthSupported:.....1X or 4X LinkWidthActive.....4X

LinkSpeedSupported: . ....2.5 Gbps or 5.0 Gbps or 10.0 GbpsLinkState: .....Active PhysLinkState: .....LinkUpLinkSpeedActive: .....10.0 GbpsLinkSpeedEnabled: .

...2.5 Gbps or 5.0 Gbps or 10.0 Gbps

You have 36 ports to check.

A. Use the ibqueryerrors.pl script after logging in to the InfiniBand switch as root.

B. Use Enterprise Manager.Cloud Control.

C. Use Enterprise Manager Express.

D. Create user-defined metrics for the InfiniBand switch.

Correct Answer: AB

Explanation:

There are two approaches for monitoring Oracle Exadata Storage Servers: using a command-line interface

(CLI) or using the graphical interface provided by the Oracle Enterprise Manager Cloud Control 12c console.

You can manually monitor the InfiniBand Switch Ports with the ibqueryerrors.pl command.

Run the ibqueryerrors.pl command to report on switch port error counters and port configuration

information using the command:

ibqueryerrors.pl -rR -s RcvSwRelayErrors,XmtDiscards,XmtWait,VL15Dropped

References: http://www.oracle.com/technetwork/articles/servers-storage-admin/monitor-exadata-em12-2291964.html

References: http://www.oracle.com/technetwork/articles/servers-storage-admin/monitor-exadata-em12-2291964.html