

Exadata Monitoring and Management Best Practices

Mughees A. Minhas
Oracle
Redwood Shores, CA, USA

Keywords:

Oracle, Exadata, monitoring, management, database, performance, monitor, Enterprise Manager, EM

Exadata Monitoring and Management

Oracle Enterprise Manager 13c provides a unified view of Oracle Exadata hardware and software with deep insight into the health and performance of all components such as the database servers, InfiniBand switches, Oracle Exadata storage cells, Oracle Databases, Automatic Storage Management (ASM), etc.

For hardware targets, you now have a photorealistic view that displays all component details, including any open incidents. Infrastructure components are now available as new target types which enable drill down to alerts and notifications for specific components like IB switch, power distribution unit (PDU), Switches etc.

Critical incidents in the hardware layer can now be published to a 3rd party ticketing system using the connector framework, something that Oracle hardware customers have asked for a long time.

As Oracle releases new hardware and software, the Oracle Exadata plug-in is updated, tested, and certified to accommodate the new products and 13.2 currently would support the X6-2 configuration of Exadata

Key Enterprise Manager New Capabilities in 13c

EM 13c allows you to monitor and manage exadata storage both at the individual cell level and at the grid level. The storage cells would consist of flash and hard disks.

The Grid and Cell home pages along with the performance pages can be used to get answers to questions like

- I. How is I/O distributed between flash and hard disk?
- II. How much benefit are we getting from Flash Cache?
- III. How much disk I/O is due to Flash Cache?
- IV. How does I/O correlate to database single block reads?

Based on your analysis, it also allows you to do administrative tasks like setting up IORM or run an `exacli` or `cellcli` command. The support for `exacli` is new in EM 13c

There are two main reasons to use ExaCLI.

- Many companies, especially government organizations, require root access or the SSH service on the cells to be disabled as part the compliance requirements.
- Service providers that host Exadata machines provide access to a virtual machine running on the compute node. Customers are not allowed to ssh into the cell to manage the cell using CellCLI.

For the use cases above, the only way to manage the cell is by using ExaCLI running on the compute node.

Enterprise Manager 13c allows you to monitor the extreme flash only configuration that was announced with the X5-2 configuration of exadata. One of the typical questions that customers have is how much of flash throughput would an OLTP consume? Benchmarks at oracle say that OLTP peak loads would consume 20% of flash throughput which means that we can use the rest of it for scans.

EM allows us to monitor and track the throughput usage, utilization percentage from a throughput perspective and also to monitor the latency and Flash IORM waits. This gives us a very good visual interpretation of the workload distribution and basis which we can chose to change your IORM strategy.

New Charts on Cell and Grid Home Page

Starting with ESS 12.1.2.1.0 I/O Resource Management (IORM) now manages flash drive I/Os in addition to disk drive I/Os to control I/O contention between databases, pluggable databases, and consumer groups. Because it is very rare for Oracle Exadata environments to be limited by OLTP I/Os, IORM automatically prioritizes OLTP flash I/Os over smart scan flash I/Os, ensuring fast OLTP response times with little cost to smart scan throughput.

EM now allows us to monitor and track the throughput usage, utilization percentage from a throughput perspective and also to monitor the flash latency and IORM waits in addition to the Disk latency and IORM waits. You can also drill down to the performance page to see a holistic view of I/O distribution across all the databases and identify the top consumers.

Cell and Grid Performance Page

You can further drill down to the performance page to see a holistic view of I/O distribution across all the databases. To identify the top consumers, you can switch to navigate to the workload distribution page across all the databases. This page allows you to configure IORM plans and to allocate a fixed amount of space in the flash cache for a database. The value specified in `flashcachesize` is a hard limit, which means that the database cannot use more than the specified value. This is different from the `flashcachelimit` value, which is a "soft"

maximum: databases can exceed this value if the flash cache is not full. Scrolling down the page gives you visual charts on the IORM distribution, IO Requests and throughput across the databases and thereby to identify the rouge consumers.

Infiniband Network Management

EM 13c provides enhanced IB switch management with enriched dashlets on the home page covering important performance and configuration metrics. This allows you to detect and notify and configuration changes or best practice violations. You also get to see a photo realistic view the hardware component.

You could do a variety of administrative tasks like enable/disable port , clear error counters and set up SNMP subscription which is new in 13c.

Monitoring: ILOM, Cisco Switch, PDU

You have similar capabilities for compute node ILOM, Switches, PDU etc. with enriched dashlets on the home page covering a variety of performance and configuration metrics like failures and speed control by looking at various views like Energy View, Network connectivity ,Service processor (SP) configuration etc

Patching Automation Support of Exadata Stack

Oracle releases Quarterly Full Stack Patch Download (QFSDP) every quarter. This is part of a new strategy for patching Exadata that reduces the number and frequency of patches released to customers even more.

Primary components are:

- Database (Database, Clusterware).
- Infrastructure includes Exadata Storage Server, InfiniBand Switch, and PDU.

With EM 13c, there is full support for patching the entire Exadata stack, including:

- Compute nodes: Firmware and operating system (OS).
- Exadata Storage Server cells: Firmware and cell software.
- InfiniBand network: Switch firmware.

The patching application offers the facility to run the pre-flight checks and monitor the logs from a single place (imagine having to manually monitor the patch execution logs for grid infrastructure, operating system, storage for all the compute and storage cells in a rack)

Exadata Virtual Provisioning

Provisioning involves repeatable, reliable, automated, unattended, and scheduled mass deployment of a RAC Cluster including virtual machines (VMs), Oracle Database (DB), Grid Infrastructure, and ASM on Virtualized Exadata.

With the Exadata plug-in's virtualization provisioning functionality, you can:

- Define the size and shape of the Virtual Machines including compute, memory and networking components.
- Create and delete a RAC Cluster including VMs, Oracle Database (DB), Grid Infrastructure, and ASM.
- Scale up and scale down a RAC Cluster by adding or removing VMs including DB, Grid Infrastructure, and ASM.

Provisioning and Deployment Challenges

Let us now focus on deployment challenges from a consolidation perspective. While deploying databases on to exadata and also plan to consolidate into Oracle Public cloud, you will need to have answers to many questions like

- I. Which apps to consolidate?
- II. What are the risks?
- III. How much cloud compute will I need?
- IV. Will I have performance improvement?
- V. How to validate performance improvements?

The new Consolidation Workbench provides an end-to-end workflow that provides three distinct steps:

- What-if analysis on various consolidation scenarios: commodity to engineered systems, non-multitenant to multitenant databases and on-prem to Oracle Cloud.
- The actual enactment of the consolidation by integrating with the provisioning features
- Post-consolidation testing using the SQL Performance Analyzer

Benefits of the DB Consolidation Workbench

The DB Consolidation Workbench helps you migrate to Oracle Exadata Cloud Service along with multi-tenant for greater consolidation. It accurately estimate required Oracle cloud compute size and shape for Exadata Cloud Service Instances. The workbench relies on AWR data to give consolidation advice by identifying candidate databases for the designated consolidation platform.

Health Check and Diagnostics Best Practices

Exachk – OraChk is a non-intrusive tool to conduct a comprehensive Exadata health check on your Exadata Database Machine to validate your hardware, firmware, and configuration which also checks for impactful reoccurring problems across Oracle customer base.

You can audit various important configuration settings within an Oracle Database Machine Exadata System - Database Servers, Storage Servers and Infiniband Switches. It produces two reports, summary and detailed. The detailed report will contain Benefit/Impact, Risk and Action/Repair information.

With EM 13c, integration of Exacheck into the Compliance framework of Enterprise Manager is complete; this would let you generate automated notifications and reports for any violation the Exadata configurations. You can install setup and upgrade the exachk utility end to end through EM.

Diagnose and repair Hardware Faults using ASR

Oracle Auto Service Request which is part of Oracle Premier Support can be used to get notification of hardware failures and get rectified. ASR is now integrated with EM 13c

EM infrastructure is used to configure ASR, there is no need to identify a separate infrastructure and separate set up steps for this as before. You can view ASR Incidents in EM Incident Manager with minimal configurations

Contact address:

Mughees Minhas

Oracle
400 Oracle Parkway
Redwood City, CA 94065
USA

Phone: +1 650 506 8661
Email mughees.minhas@oracle.com
Internet: <http://www.oracle.com>