



Lesser-Known Tips & Tricks of Exadata

LuxOUG - 2022

Fernando Simon

Senior Database Specialist



SAFE HARBOR STATEMENT

- *I am here as Fernando Simon, I do not represent my current employer/company.*
- *I am sharing my experience, so, is based on environments in that I worked. So, be aware that they, maybe, do not represent that same as yours.*
- *“The postings on this document are my own and don’t necessarily represent my actual employer positions, strategies or opinions. The information here was edited to be useful for general purpose, specific data and identifications were removed to allow reach the generic audience and to be useful for the community.”*



500+ technical experts helping peers globally

The **Oracle ACE Program** recognizes and rewards community members for their technical contributions in the Oracle community

3 membership tiers



ORACLE
ACE Director



ORACLE
ACE



ORACLE
ACE Associate

For more details on Oracle ACE Program:
bit.ly/OracleACEProgram



Nominate

yourself or someone you know:

acenomination.oracle.com

Connect:  oracle-ace_ww@oracle.com

 Facebook.com/oracleaces

 [@oracleace](https://twitter.com/oracleace)




AGENDA

- Who I am and what I will be talking about.
- Project for Exadata, consolidation.
- Exadata and features.
- Resource management and monitoring.
- Patching.
- QA.

ABOUT ME

- Senior Database Specialist at Luxembourg.
 - OCP, OCE RAC, OCI Architect, Autonomous Specialist.



- Contacts:

fernando.simon.br@gmail.com

<https://www.fernandosimon.com/blog/>

<https://twitter.com/FSimonDBA>

<https://www.linkedin.com/in/fernando-simon/>

ABOUT ME

- DBA since 2004:
 - Oracle, PostgreSQL, SQLServer, and DB2.
 - Head and DBA Team Manager at Court of Justice:
 - Exadata since 2010:
 - Exadata V2, X2, X4 (Full), X5 (Full EF), and X6.
 - ZDLRA since 2014/2015:
 - [OOW 2015 - MAA Project, Multi-Site, RAC+RAC, DG, ZDLRA.](#)
 - Contributing with Oracle Brazil community since 2010.
- Luxembourg October/2017.
 - European institutions:
 - LCM (Life Cycle Management) to the Oracle Products.
 - Financial Institution:
 - MAA Infrastructure, ExaCC/Exadata and ZDLRA.

WHAT I WILL (AND NOT) TALK ABOUT

- **I will talk about:**

- Exadata:
 - What I learned using it since 2010.
- Exadata Cloud@Customer.
- Hidden details for some features:
 - Smart Scan, IORM, HCC.
- Monitoring:
 - Metrics. Why, how to use, and benefits.
- Patching.

- **I will not talk about:**

- How to use/configure features:
 - Smartsan.
 - HCC.
- How to install it.
- Prices.

SESSION



The screenshot shows the Cambridge Dictionary interface. At the top, there is a navigation bar with the Cambridge Dictionary logo, a menu icon, and links for Dictionary, Translate, and Grammar. Social media icons for Facebook, Instagram, and Twitter are also present, along with a 'Log in / Sign up' link. Below the navigation bar is a search bar containing the text 'lesser-known'. To the right of the search bar, there are icons for 'English', a search magnifying glass, and a dropdown menu showing 'English-German' and 'German-English'. The main content area displays the meaning of 'lesser-known' in English, with a play button icon. The word 'lesser-known' is prominently displayed in a large font, followed by the word 'adjective'. Below this, the UK and US pronunciations are provided: UK /ˌles.əˈnəʊn/ and US /ˌles.əˈnoʊn/. A yellow button with a plus sign and a list icon is located to the right of the pronunciation. The definition is given as 'not as popular or famous as something else:'. An example sentence is provided: '• We stayed on one of the lesser-known Greek islands.'

EXADATA SINCE 2010



[Overview](#) » [Exadata Cloud@Customer](#) » [Exadata Infrastructure](#) » Infrastructure Details






ACTIVE

EXACC: 

[Manage Contacts](#) [Scale Infrastructure Storage](#) [Deployment A](#)

[Infrastructure Information](#) [Tags](#)

General Information

Compartment:  (root)/ExaCC
OCID:  [Show](#) [Copy](#)
Created: Thu, Sep 2021 
Time Zone: UTC
Shape: Full Rack
Exadata System Model: Exadata Cloud@Customer X9M-2
Lifecycle State: Active

EXADATA SINCE 2010

- Tribunal de Justiça de Santa Catarina (TJSC) – 2000 to 2010:

- More than 100 databases (8i, 9i, and 10c).


- Paper process:

- If something is missing recover from paper.

- Client server application.

- Lack of support.

- Several outages.



EXADATA SINCE 2010

- Exadata/ZDLRA/MAA Project – 2010-2017:

- Centralized in eight core databases.
 - More than 60TB of data.
 - More than 100.000 IOPS (explained later).
 - Monthly grows – 14tapes LTO6/month
- Digital process, paperless.
- No SPOF.
- Huge increase in reliability.
- Always-on database/data.



EXADATA SINCE 2010

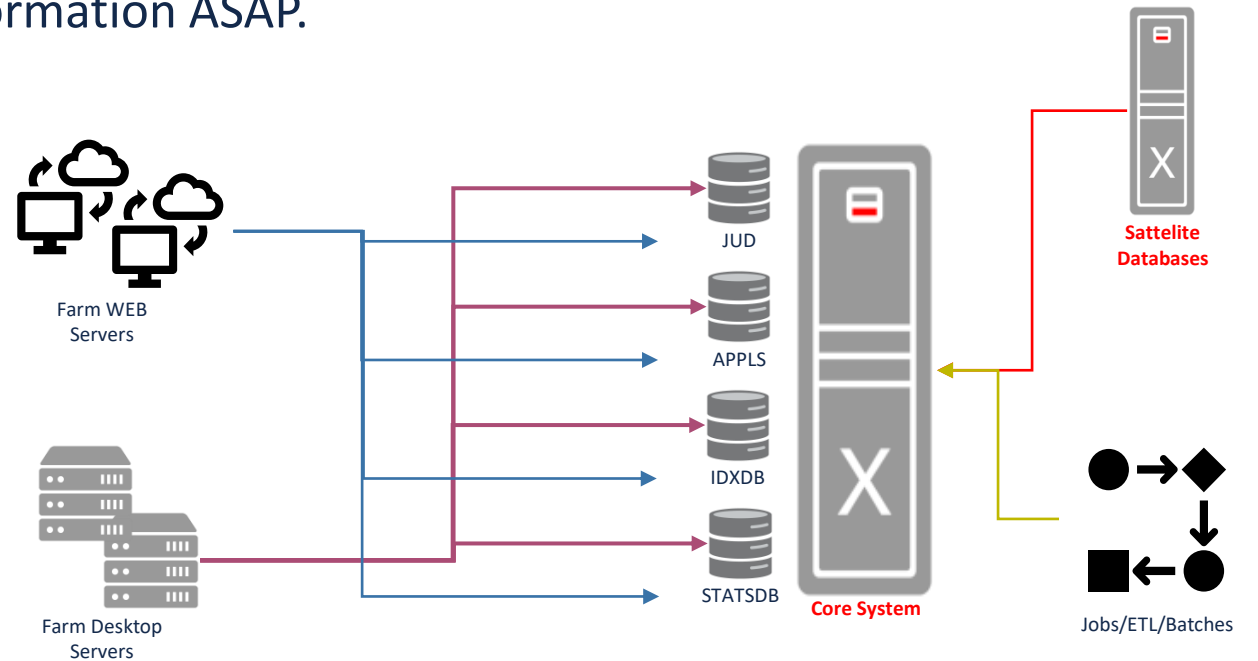
- Exadata/ZDLRA/MAA Project – 2010-2017:

- +13.000 employees:
 - +9.500 internal users.
 - +200.000 web requests per week.
- Internal and public services:
 - Courts, Process, Appeals, Taxes.
 - Lawyers and external agents.
 - Used by Police daily.
- Nearly 14 million processes subjects:
 - +750.000 just in 2015.
- Core databases:
 - Desktop, WEB, Sync, Batch, ETL process.
 - Satellite databases.



EXADATA SINCE 2010

- Basic rules:
 - #1 – Desktop/Internal Users.
 - #2 – Web/External access needs to retrieve information ASAP.
 - #3 - Sync between DB's needs to be done ASAP.
 - #4 – Possible to disable one group ASAP.



EXADATA



EXADATA

- My opinion, Exadata have 3 pillars:
 - Smart Scan and Storage Index:
 - Saves I/O.
 - Plug-in feature.
 - Resource manager (IORM + DB RM):
 - You can control everything.
 - Metrics.
- But not just that:
 - HCC (“save” space).
 - Easy to sustain.



EXADATA – SMART SCAN

- Smart Scan, key points that you need to know:
 - A Full Table Scan or an Index Fast Full Scan (in general) operation must occur.
 - The segment must be big enough to fire a direct path read operation:
 - Can tune the Small Table Threshold (STT) and Medium Table Threshold (MTT)
 - Do you know one OLTP/DSS system that does not execute FTS?
- TJSC Example:
 - Core database made 100.000 IOPS, **BUT** this was 95% already saved by smart scan
 - In the case of returning to the traditional environment, one storage that delivers more than 1 Million IOPS was needed (just for one DB).

EXADATA – SMART SCAN

- SLOB example (made in 2015):

- Update 0% - Work Unit 65536 - Schemas 128 – Threads 64 - RAC 8 Nodes – SGA 256 GB

physical read IO requests	938,935,189	2,173,938.24	9,389,351.89
physical read bytes	7,691,763,163,136	17,808,916,172.16	76,917,631,631.36
physical read total IO requests	938,951,812	2,173,976.73	9,389,518.12

- Update 0% - Work Unit 8296 - Schemas 128 – Thread 16 - RAC 8 Nodes – SGA 256 GB

physical read IO requests	758,823,542	2,471,792.93	758,823,542.00
physical read bytes	6,243,449,634,816	20,337,422,071.31	6,243,449,634,816.00
physical read total IO requests	758,841,843	2,471,852.54	758,841,843.00

- Why not all Exadata 4.4 million of IOPS?

- SLOB (at that time) never executed the FTS queries, only indexed ones.
- **This is optimal FOR storage. This kills Smart Scan.**
- So, is important to understand the features, how to use them, and how to troubleshoot them.

EXADATA – RESOURCE MANAGER

- Resource manager:
 - Better way to learn is to read [Chapter 6 of Exadata User Guide](#).
 - IORM:
 - **CATPLAN**: Priority over the others. Based on categories coming from databases (different DB's can have the same categories).
 - **DBPLAN**: Priority is based on database names.
 - **CLUSTERPLAN**: Priority is based on cluster names. New at 21 version of Exadata Software.
 - Database Resource Manager:
 - Multi levels plans (inside PDB or 11g) allow more tuning.
 - Shared plans are more tricky to monitor due to the volatile usage of the CPU.
 - Map your consumers by service name or username is good to better control.


EXADATA – RESOURCE MANAGER

- Tips:
 - Create the same categories at different databases.
 - Map your consumer groups for the desired categories, and RAC services with the desired consumer group.
 - Enable the plan to allow have metrics values and allow monitoring. (at the storage level).
 - Multi-level plans (PDB and IORM), resource not used (CPU, I/O) always go to the next level.
 - Never defines 100% for one level (besides the one for SYS) and always define SYS as the first level or highest share.
 - **More detailed the plan is, more complex and difficult to sustain.**

EXADATA – RESOURCE MANAGER

- There is no comparison with/at the traditional environment, you can prioritize I/O base in connections (independently from which database are coming).

- BUT....



EXADATA – RESOURCE MANAGER

- **Oracle is cutting the IORM:**
 - *“Starting with Oracle Exadata System Software release 21.2.0, the category plan is deprecated and a warning message is issued when a category plan is set.”*
 - Without catplan, Exadata I/O prioritization lost fine tune and will lose important key point.

EXADATA – RESOURCE MANAGER

- Creating categories inside of DB:

```
BEGIN
  DBMS_RESOURCE_MANAGER.CREATE_PENDING_AREA();
  DBMS_RESOURCE_MANAGER.CREATE_CATEGORY(CATEGORY => 'COR_HIGH_PRIORIDADE', COMMENT => 'Category for HIGH priotiry');
  DBMS_RESOURCE_MANAGER.CREATE_CATEGORY(CATEGORY => 'COR_MEDIUM_PRIORIDADE', COMMENT => 'Category for MEDIUM priotiry');
  DBMS_RESOURCE_MANAGER.CREATE_CATEGORY(CATEGORY => 'COR_LOW_PRIORIDADE', COMMENT => 'Category for LOW priotiry');
  DBMS_RESOURCE_MANAGER.SUBMIT_PENDING_AREA();
END;
```

- IORM plan:

```
ALTER IORMPLAN
catPlan=( (name=COR_HIGH_PRIORIDADE, level=1, allocation=70),
          (name=COR_HIGHMED_PRIORIDADE, level=2, allocation=70),
          (name=COR_MEDIUM_PRIORIDADE, level=3, allocation=60),
          (name=COR_LOW_PRIORIDADE, level=4, allocation=80),
          (name=OTHER,level=5, allocation=100 )
        ),
dbPlan=( (name=JUD, level=1, allocation=60),
         (name=APPLS, level=2, allocation=50),
         (name=IDXDB, level=2, allocation=30),
         (name=STATSDB, level=3, allocation=80),
         (name=ADVSC, level=4, allocation=90),
         (name=OTHER, level=5, allocation=100, flashCache=OFF)
       ),
objective=low_latency
```

EXADATA – RESOURCE MANAGER - FUNNY HISTORY

Oracle Support- 11+ years ago [Notes]

Generic Note

Hi ,

We see that there's 13 categories ...there's a simple workaround to reduce the categories to 10 for the workaround of [bug# 11818623](#). Please provide a good time so that we can talk.

IORM state dump

Time: 03-24-2011 22:21:36.335209000

IORM Enabled: Active Plan, Throttling

Current IORM Plans

Category Plan

number of categories: 13

id 0: PG5_ALTA_PRIORIDADE

EXADATA – RESOURCE MANAGER

Help Center Search

Engineered Systems / Exadata Database Machine

User's Guide

Table of Contents

- [Title and Copyright Information](#)
- [Preface](#)
- [1 Introducing Oracle Exadata System Software](#)
- [2 Configuring Oracle Exadata System Software](#)
- [3 Administering Oracle ASM on Exadata](#)
- [4 Administering Oracle Database on Exadata](#)
- [5 Maintaining Oracle Exadata System Software](#)
- [6 Managing I/O Resources](#)**
- [7 Monitoring Exadata](#)
- [8 Using the CellCLI Utility](#)
- [9 Using the dcli Utility](#)
- [10 Setting up Oracle Exadata Storage Snapshots](#)

Oracle® Exadata System Software

User's Guide
22.1
F29254-21
September 2022

[+ Title and Copyright Information](#)

EXADATA - METRICS



EXADATA - METRICS

- Exadata metrics are recorded observations about the health of Exadata Software:
 - Flash cache.
 - Cell and Grid disks.
 - CPU, Memory.
 - IOPS.
 - Many others...
- Divided into three types:
 - **Cumulative** metrics are statistics that accumulate over observed time.
 - **Instantaneous** metrics contain the current value at the time of the metric observation.
 - **Rate** metrics are computed statistics where the value is averaged over time.

EXADATA - METRICS

- **Some isolated/per cell metrics:**

- CL_CPU: The CPU usage.
- CL_MEMUT: The percentage of total physical memory used.

- N_HCA_MB_RCV_SEC: The MB/s received by the InfiniBand.
- N_HCA_MB_TRANS_SEC: The MB/s transmitted by the InfiniBand.

- FC_BY_ALLOCATED: The size of megabytes allocated in flash cache.
- FC_IO_BY_R_SEC: The MB/s read from the flash cache.
- FC_IO_BY_W_SEC: The MB/s written at the flash cache.
- FC_IO_BY_R_SKIP_SEC: The MB/s read from disks for I/O requests that bypass flash cache.
- FC_IO_BY_R_MISS_SEC: The MB/s read from disks because no data was in the flash cache.

- SIO_IO_OF_RE_SEC: The MB/s returned by smart I/O.

EXADATA - METRICS

- **Some Categories/Database/Consumer Groups Metrics:**

- DB_FC_IO_BY_SEC: The MB/s for a specific database to flash cache.
- DB_FC_IO_RQ_LG_SEC: The number of large IOPS issued by a database to flash cache.
- DB_FC_IO_RQ_SM_SEC: The number of small IOPS issued by a database to flash cache.


- DB_IO_BY_SEC: The MB/s for a specific database to hard disks.

- DB_IO_RQ_LG_SEC: The large IOPS made at hard disks by a database.
- DB_IO_RQ_SM_SEC: The small IOPS made at hard disks by a database.
- DB_IO_TM_LG_RQ: The average latency of (reading or writing) large blocks at each request by a database from hard disks.
- DB_IO_TM_SM_RQ: The average latency of (reading or writing) small blocks at each request by a database from hard disks.

- DB_IO_UTIL_LG: The percentage of hard disk resources utilized by large requests from this database.
- DB_IO_UTIL_SM: The percentage of hard disk resources utilized by small requests from this database.

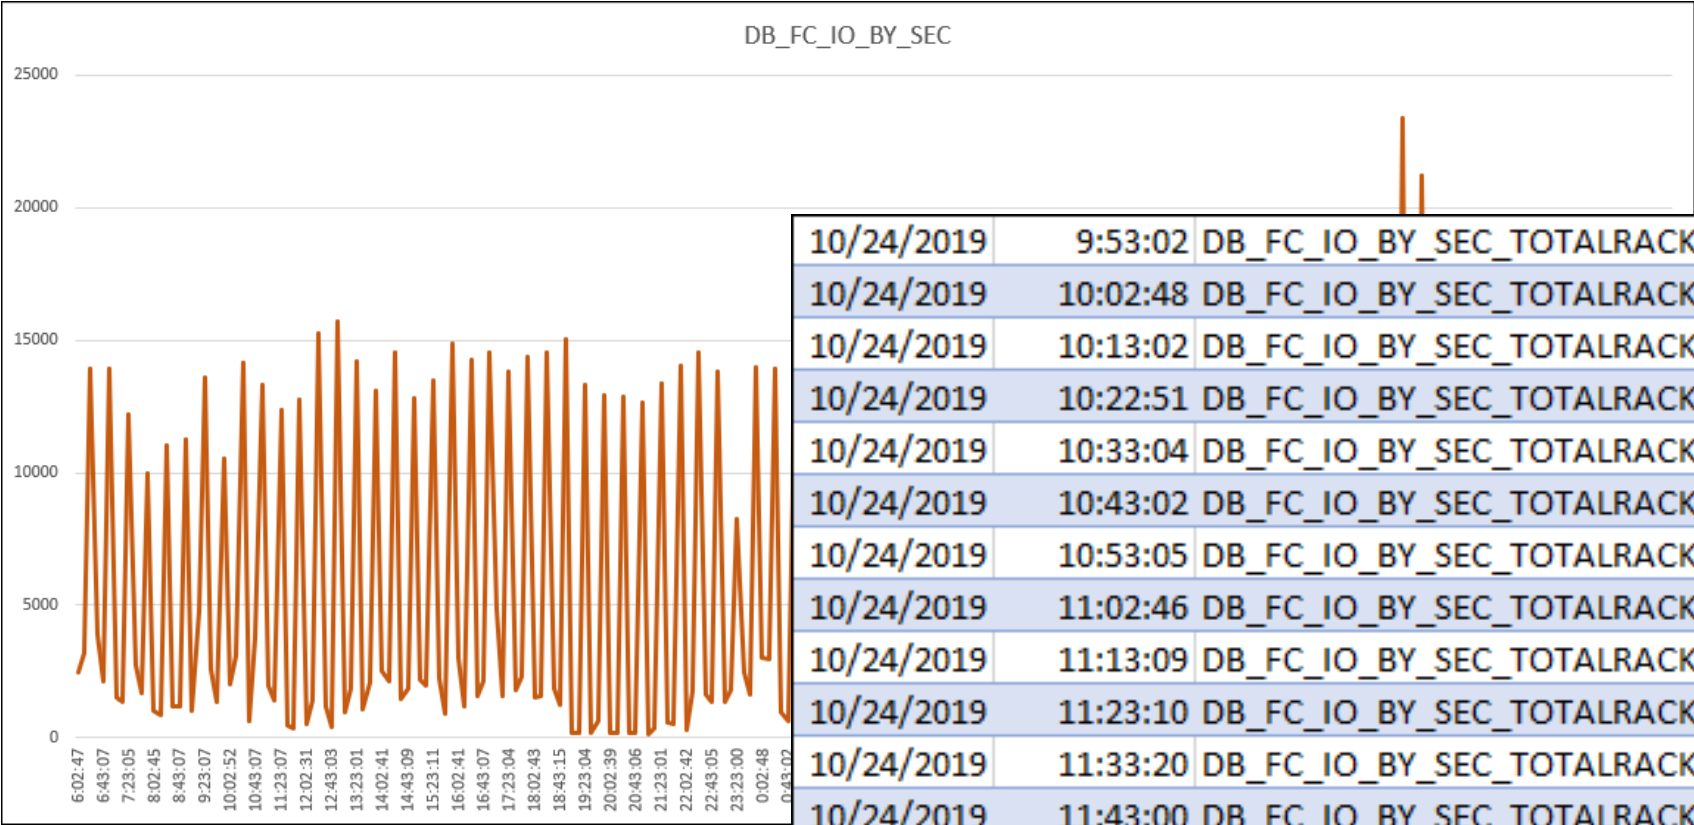
EXADATA - METRICS

- What do you think about this database:



- Something wrong?

EXADATA - METRICS



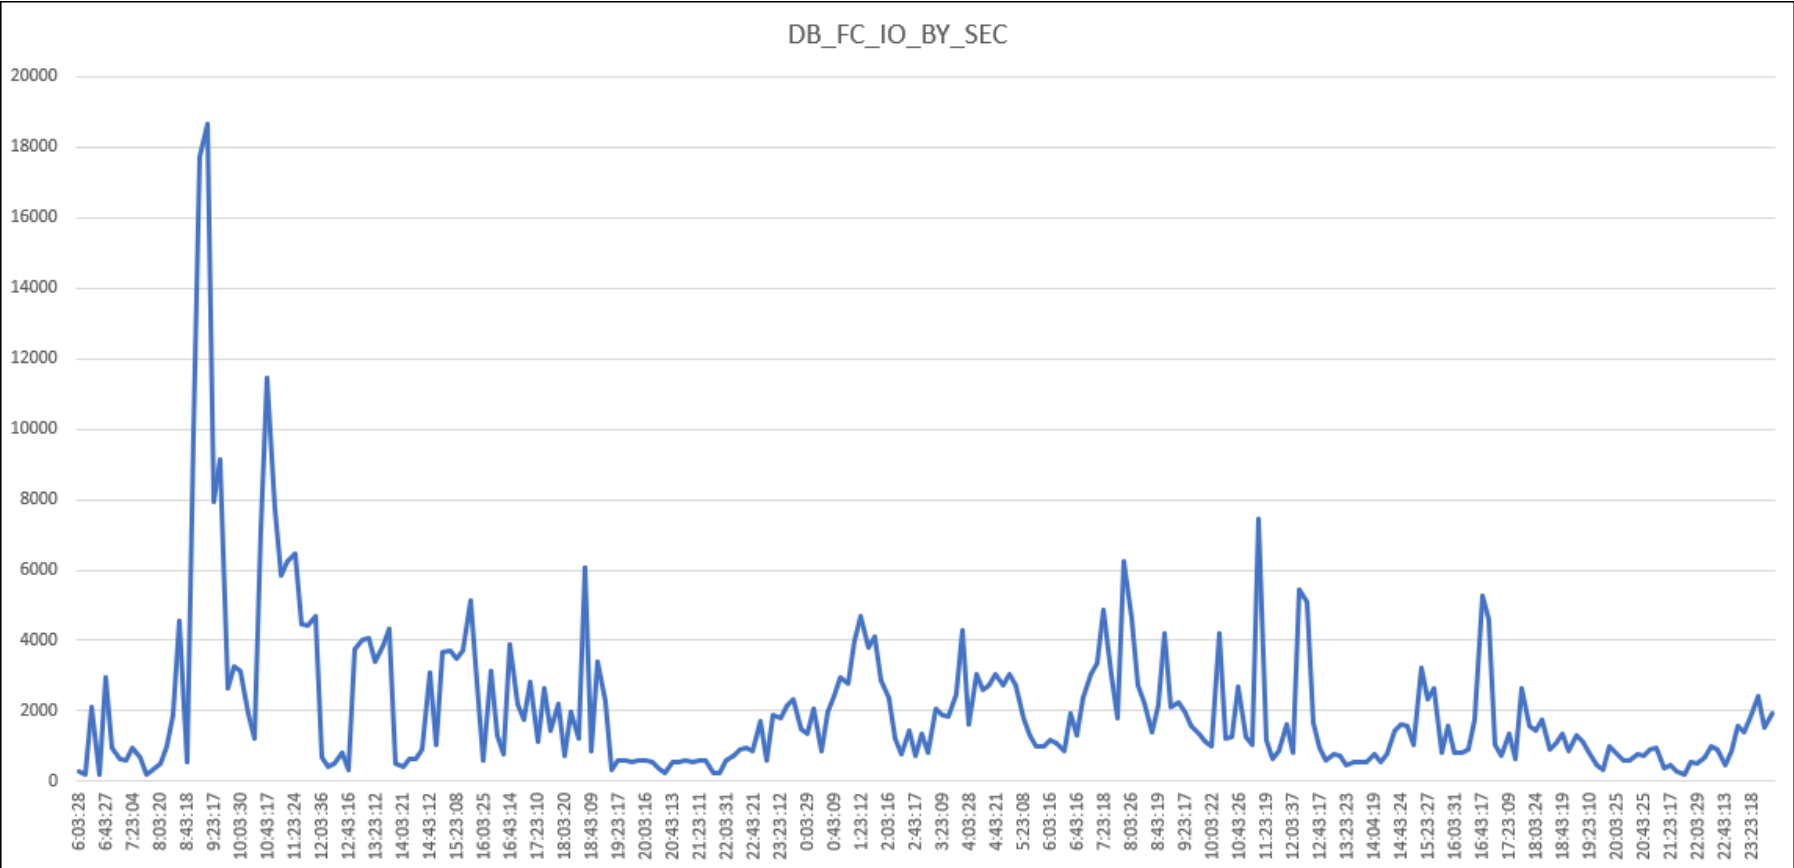
10/24/2019	9:53:02	DB_FC_IO_BY_SEC_TOTALRACK	13099.45	MB/s
10/24/2019	10:02:48	DB_FC_IO_BY_SEC_TOTALRACK	317.25	MB/s
10/24/2019	10:13:02	DB_FC_IO_BY_SEC_TOTALRACK	227.3	MB/s
10/24/2019	10:22:51	DB_FC_IO_BY_SEC_TOTALRACK	12194.95	MB/s
10/24/2019	10:33:04	DB_FC_IO_BY_SEC_TOTALRACK	2567.95	MB/s
10/24/2019	10:43:02	DB_FC_IO_BY_SEC_TOTALRACK	401.35	MB/s
10/24/2019	10:53:05	DB_FC_IO_BY_SEC_TOTALRACK	14969.65	MB/s
10/24/2019	11:02:46	DB_FC_IO_BY_SEC_TOTALRACK	1441.3	MB/s
10/24/2019	11:13:09	DB_FC_IO_BY_SEC_TOTALRACK	1390.6	MB/s
10/24/2019	11:23:10	DB_FC_IO_BY_SEC_TOTALRACK	13821.25	MB/s
10/24/2019	11:33:20	DB_FC_IO_BY_SEC_TOTALRACK	2057.6	MB/s
10/24/2019	11:43:00	DB_FC_IO_BY_SEC_TOTALRACK	1949.6	MB/s
10/24/2019	11:53:15	DB_FC_IO_BY_SEC_TOTALRACK	13714.65	MB/s
10/24/2019	12:02:59	DB_FC_IO_BY_SEC_TOTALRACK	1932.2	MB/s

EXADATA - METRICS

- The report says that DB_FC_IO_BY_SEC (MB/s read from FC) got a peak at 25GB/s (for the entire rack):
 - This is below what Exadata can deliver, but the metric shows some strange behavior.
- After some research over running SQL's hit two notes:
 - *Slow Performance Of DBA_AUDIT_SESSION Query From "Failed logins" Metric in Enterprise Manager (Doc ID 1178736.1).*
 - *EM 12c, 11g: High Cpu Utilization from Enterprise Manger Agent perl Process Running failedLogin.pl for audit_failed_logins metric (Doc ID 1265699.1).*
- Basically, the EM/CC collection for failed login attempts was running and consuming a lot of resources.

EXADATA - METRICS

- After disabling EM check:



EXADATA - METRICS

- Tips:
 - Prefer to read the METRICHISTORY.
 - METRICCURRENT can be not “complete” when you read it because some metrics are cumulative.
 - Some metrics have correlation, mainly missed flash reads from smart scan counts at disk reads.
 - You can drill down over metrics.
 - Disk > Flash Cache > Smart Scan > Category > Database > PDB.
 - Older versions of the docs are better:
 - [22.1: User's Guide \(oracle.com\)](#).
 - [19.3: User's Guide \(oracle.com\)](#).
 - Some personal posts about Exadata Metrics: <https://www.fernandosimon.com/blog/tag/metrics/>

EXADATA - METRICS

FC_IO_BY_R_MISS_SEC: MB/s from disks because some of the requested data was not in the flash cache.

SIO_IO_RD_HD_SEC: MB/s read from hard disk by smart IO.

```
CellCLI> list metristory where name = 'FC_IO_BY_R_MISS_SEC' and collectionTime > '2021-10-20T10:00:00+02:00' and collectionTime < '2021-10-20T10:30:00+02:00' and metricvalue > 2;
```

FC_IO_BY_R_MISS_SEC	FLASHCACHE	13.360 MB/sec	2021-10-20T10:11:58+02:00
FC_IO_BY_R_MISS_SEC	FLASHCACHE	18.565 MB/sec	2021-10-20T10:12:58+02:00
FC_IO_BY_R_MISS_SEC	FLASHCACHE	26.359 MB/sec	2021-10-20T10:13:58+02:00
FC_IO_BY_R_MISS_SEC	FLASHCACHE	14.020 MB/sec	2021-10-20T10:14:58+02:00
FC_IO_BY_R_MISS_SEC	FLASHCACHE	2.111 MB/sec	2021-10-20T10:15:58+02:00
FC_IO_BY_R_MISS_SEC	FLASHCACHE	2.019 MB/sec	2021-10-20T10:20:59+02:00
FC_IO_BY_R_MISS_SEC	FLASHCACHE	3.539 MB/sec	2021-10-20T10:21:59+02:00

```
CellCLI> list metristory where name = 'SIO_IO_RD_HD_SEC' and collectionTime > '2021-10-20T10:00:00+02:00' and collectionTime < '2021-10-20T10:30:00+02:00' and metricvalue > 2;
```

SIO_IO_RD_HD_SEC	SMARTIO	4.305 MB/sec	2021-10-20T10:02:57+02:00
SIO_IO_RD_HD_SEC	SMARTIO	4.026 MB/sec	2021-10-20T10:10:58+02:00
SIO_IO_RD_HD_SEC	SMARTIO	8.375 MB/sec	2021-10-20T10:11:58+02:00
SIO_IO_RD_HD_SEC	SMARTIO	13.266 MB/sec	2021-10-20T10:12:58+02:00
SIO_IO_RD_HD_SEC	SMARTIO	15.171 MB/sec	2021-10-20T10:13:58+02:00
SIO_IO_RD_HD_SEC	SMARTIO	15.689 MB/sec	2021-10-20T10:14:58+02:00
SIO_IO_RD_HD_SEC	SMARTIO	14.997 MB/sec	2021-10-20T10:24:59+02:00

```
CellCLI>
```

EXADATA - METRICS

The screenshot shows a web browser window with the address bar containing the URL: <https://www.oracle.com/technetwork/pt/articles/oem/metricas-storage-servers-exadata-2550748-ptb.html>. The page is from the Oracle Technology Network (OTN) and is in Portuguese. The article title is "Lendo métricas dos Storage Servers a partir de um Database Server em um Oracle Exadata". The author is Fernando Simon (OCE) and Deiby Gómez, and it was posted in May 2015. The article is reviewed by Marcelo Pivovar, a Solution Architect. The content discusses Oracle Exadata solutions and metrics, mentioning that there are many solutions in use and that administrators have various tools to monitor performance. It notes that logging into each Storage Server to check metrics is a lot of work, but the article will show how to do it without logging into each server. The article is structured into five steps: 1. Creation of the script, 2. Testing the script, 3. External Table, 4. Data adjustment, and 5. Reading the data.

ORACLE | Fazer Login/Inscrever-se | Ajuda | País | Comunidades | Eu sou... | Eu quero... | Pesquisar

Produtos | Soluções | Downloads | Loja | Suporte | Treinamento | Parceiros | Sobre | OTN

Oracle Technology Network > Articles > Enterprise Management

Lendo métricas dos Storage Servers a partir de um Database Server em um Oracle Exadata

Por Fernando Simon (OCE) e Deiby Gómez
Postado em Maio 2015

Revisado por Marcelo Pivovar - Solution Architect

Atualmente existem inúmeras soluções Oracle Exadata em uso nas mais diversas áreas e os seus administradores tem diversas ferramentas para monitorar o desempenho e verificar o que está acontecendo nos Storage Servers, Database Servers e rede de intercomunicação. De qualquer forma, existem inúmeras métricas do Exadata Software presente nos Storage Servers com informações sobre o ambiente e que geralmente não são utilizadas e muitas nem são bem compreendidas.

Nós podemos utilizar estas métricas dos Storage Servers para entender a carga de trabalho sobre o ambiente, mas infelizmente você tem que logar em cada um dos Storage Servers para ver as métricas (nem o Cloud Control entrega todas elas). Se você tem um Exadata Full Rack terá 14 Storage Servers para fazer o login, é bastante trabalho! Neste artigo vamos mostrar como ler estas métricas sem precisar fazer login em cada um dos Storage Server, após seguir os passos deste artigo você conseguirá consultar as métricas através de simples comandos SQL.

Basicamente os passos que iremos abordar seguem o seguinte fluxo:

A partir de um SELECT em um Database Server, a instância utilizará uma external table para chamar um script que irá extrair as informações das métricas de cada um dos Storage Servers.

O que veremos:

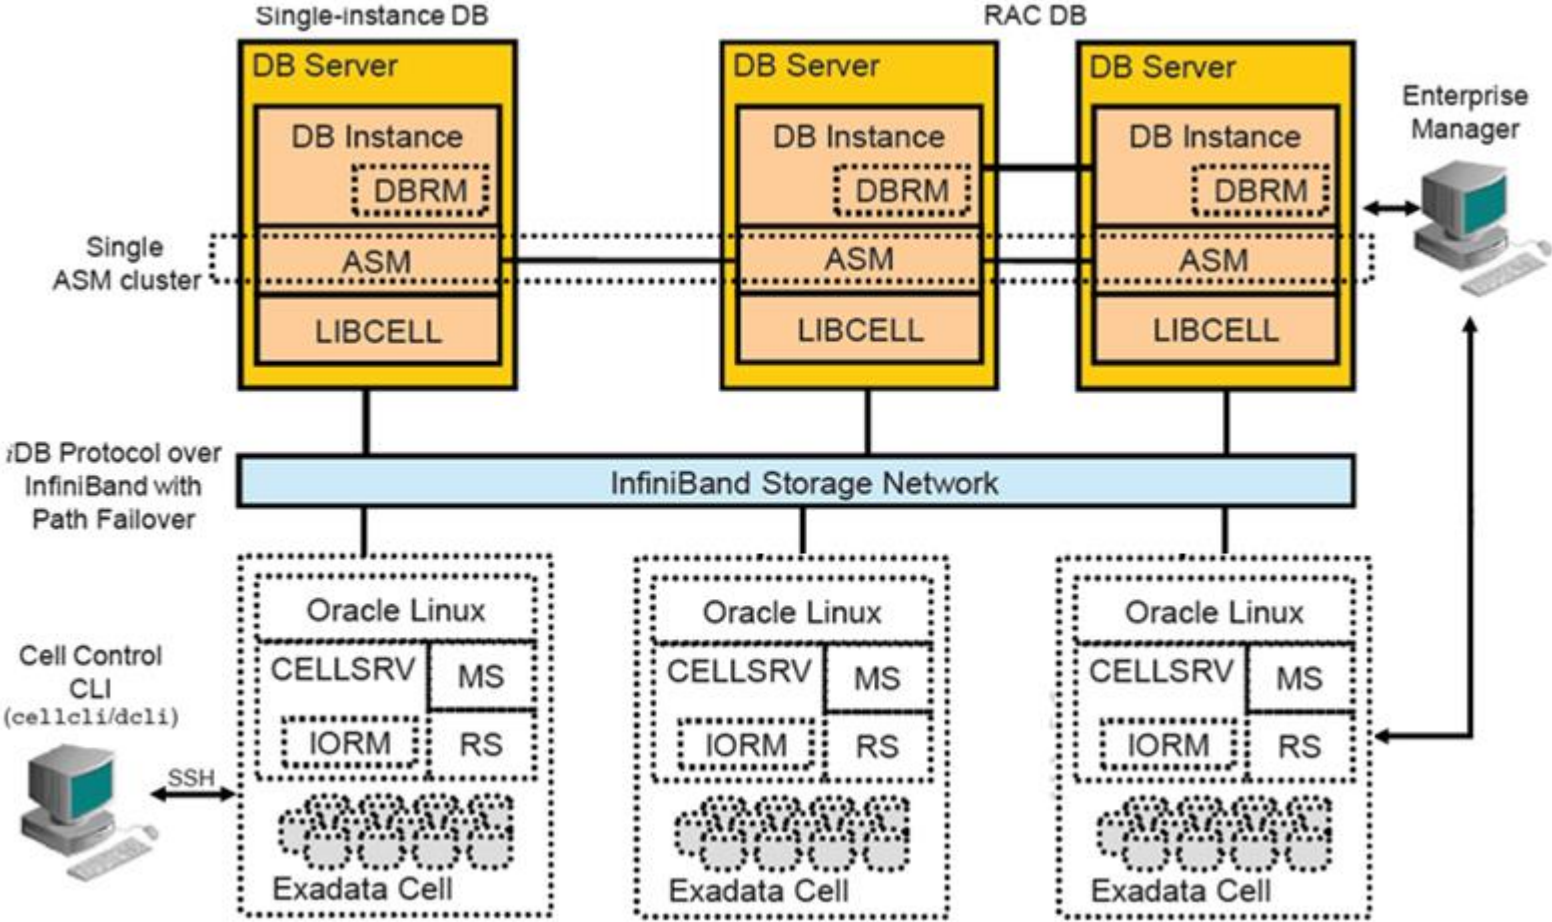

- Passo 1 – Criação do Script
- Passo 2 – Testando o Script
- Passo 3 – External Table
- Passo 4 – Ajuste de dados
- Passo 5 – Lendo os dados

Passo 1 – Criação do Script




EXADATA - METRICS

- With 22.1 version you can have the Real-Time Insight:
 - Your metrics can be read by external metric collection (and visualization) platform(Grafana).
 - Needs to be enabled (and tuned).
 - Some key metrics already are enabled by default.
 - You can control the interval of the push events.
 - Gavin Parish Exadata PM Blog Post:
 - [Exadata Real-Time Insight \(oracle.com\)](https://www.oracle.com/technetwork/middleware/exadata/real-time-insight-2312811.html).

EXADATA - PATCH



EXADATA - PATCH



EXADATA - PATCH

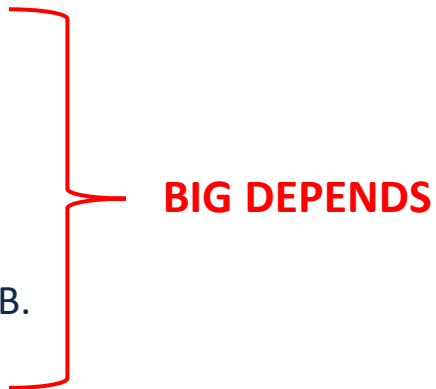
- Oracle Software:
 - Oracle Database and GI.
 - Exadata Software:
 - Cellsrv:
 - MS/RS/CS.
 - Sun X series machine:
 - ILOM.
 - RoCE/InfiniBand.
 - Controller/FW.
- Oracle Enterprise Linux:
 - Storage nodes.
 - Database nodes.
- Fabric Switch:
 - Sun InfiniBand Switch 36.
 - RCoE switch (ExaM machines).
 - OFED/iDB/RDS/ZDP.
- Cisco switch:
 - Admin network.

EXADATA - PATCH

- **WHY** do you need to patch?
 - New versions for Operational System for DBNodes.
 - Fix firmware bugs.
 - Fix Exadata Software Bugs:
 - **Exadata Critical Issues (Doc ID 1270094.1).**

Exadata Storage Server		
#	Applies to	Issue
EX70	Storage servers running Exadata version 21.2.0 through 21.2.4 with flashCacheMode=WriteThrough	Bug 33380819 - Storage servers running Exadata version 21.2.0 through 21.2.4 with flash cache configured in WriteThrough mode do not use the flash cache, which results in hard disks servicing all database I/O operations, causing severe performance degradation and possible instance crash.
EX69	Database servers (non-virtual or KVMHOST) and storage servers running Exadata version 21.2.0 through 21.2.2, or 20.1.0 through 20.1.12, or 19.3.0 through 19.3.20	Bug 33042327 - Soft-offlined memory pages may be reused while marked poisoned when DIMM correctable errors (CE) occur. Exadata systems affected by this issue may experience node crashes, memory corruption, instance crashes, or process kills.
EX68	High Capacity (HC) storage servers being updated from Exadata <=19.2 to Exadata >=19.3.6	Bug 32843838 - During or immediately after software update from Exadata 19.2, or lower, to Exadata 19.3.6, or higher, on High Capacity (HC) storage servers, the database may experience high write IO latency, which may cause severe database performance degradation, or database process crash, or database instance crash when a process like LGWR is affected.
EX67	Database servers on RoCE-based systems running Exadata 21.2.0 through 21.2.2, 20.1.0 through 20.1.12, or 19.3.10 through 19.3.20	Bug 32404239 - Database instances may crash or database processes may fail due to an I/O buffer memory corruption.
EX66	X8M, X8, X7, X6, or X5 database servers updated to 2020-Sep, 2020-Oct, or 2020-Nov Exadata	Bug 32218558 - In rare circumstances the disk controller on X8M, X8, X7, X6, or X5 database

EXADATA - PATCH

- Create a plan:
 - **888828.1.**
 - Check Incompatibilities:
 - Apply needed RU's/PSU's before Exadata patch.
 - Bare Metal/OVM/KVM:
 - Rolling:
 - Switch Rac OneNode to other node/Patch Storage Node By Node/Patch DB.
 - Non-Rolling:
 - Shutdown Everything/Patch Storage/Patch DB.
 - OVM/KVM.
- 
- BIG DEPENDS**

EXADATA - PATCH

- Two modes:
 - Rolling:
 - “No Downtime”.
 - Time consuming (**up to 2 hours per cell**).
 - Problems appear ASAP and stop the patch.
 - Non-Rolling mode:
 - Shutdown everything, patch all.
 - **Up to 2 hours to patch everything.**
 - Ooops... Can kill everything:
 - EX54.
- Before Patch:
 - Always reset ILOM for DBs and CELLS because long-running ILOM’s can hide HW errors:
 - But, **NEVER** reset ILOM for SW since it is the own O.S.
 - After ILOM reset, always check for new HW errors.
 - **Never continue with a machine that have HW errors.**
- **Each patch is different.**

EX54	High capacity storage servers with 8 TB or 10 TB disk drives updated to Exadata version 12.2.1.1.8, or to version >= 18.1.6.
------	--

Bug 29645783 - shortly after software upgrade hard disks may fail on High Capacity storage servers with 8 TB or 10 TB disk drives. If this issue occurs on partner disk drives on multiple storage servers when performing non-rolling upgrade, then one or more ASM disk groups may be forcibly dismantled and lost, requiring restore from backup.
--

EXADATA - PATCH

- Plan the plan:
 - **Where** you are:
 - Which version are you running?
 - **Where** you want to go:
 - Which version do you need to patch?
 - **README, RELEASE NOTES and know issues.**
 - **What** you need to do:
 - What do I need to patch?
 - Am I compatible?
 - Will I stay compatible?
- Plan how to do:
 - **Which** mode I will use?
 - Rolling/Non-Rolling mode.
 - **What** is the expected downtime.
 - **What** it is running:
 - OVM.
 - Baremetal.
 - **What** are the licensing:
 - RAC/RAC One Node/Single.

EXADATA - PATCH

Document 888828.1

support.oracle.com/epmos/faces/DocumentDisplay?_afLoop=402357541242339&id=888828.1&_afWindowMode=0&_adf.ctrl-state=1d3qkyne7f_4#aref_section333

ORACLE MY ORACLE SUPPORT PowerView is Off

Last Login: October 21, 2021 9:45 PM CEST Switch to Cloud Support Fernando (Available) (0) Contact Us Help

Dashboard Knowledge Service Requests Patches & Updates Community Certifications Systems Advanced Customer Services Proactive Hardware Services Settings

Copyright (c) 2021, Oracle. All rights reserved. Oracle Confidential.

★ Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)

To Bottom

In this Document

- [Purpose](#)
- [Scope](#)
- [Details](#)
- [Latest Releases and Software Maintenance News](#)
- [Exadata Software and Hardware Maintenance Planning](#)
- [Software Certification, Release Frequency, and End Date](#)
- [Recommended Versions](#)
- [Policy to Obtain Older Exadata System Software](#)
- [Critical Issues](#)
- [Security-Related Guidance](#)
- [Latest Releases](#)
- [Guidelines for Related Products and Features](#)
- [Version Compatibility and Requirements](#)
- [Exadata System Software Requirements for Exadata Database Machine Hardware](#)
- [Oracle Database](#)
- [Minimum Exadata, Grid Infrastructure, and Database Requirements for Exadata Feature Usage](#)
- [Database Server Operating System](#)
- [Oracle Linux x86-64 on 2-socket and 8-socket Intel x86-64 Database Servers](#)
- [Patch Release History for Exadata Database Machine Components](#)
- [Quarterly Full Stack Download Patch \(QFSDP\)](#)
- [Exadata Software](#)
- [Image Version and Date Code](#)
- [Exadata 21](#)
- [Exadata 20](#)
- [Exadata 19.3](#)
- [Exadata 19.2, 19.1](#)
- [Exadata 18](#)
- [Exadata 12.2](#)
- [Exadata 12.1](#)

Was this document helpful?

Yes

No

Document Details

Type:	REFERENCE
Status:	PUBLISHED
Last Major Update:	Oct 21, 2021
Last Update:	Oct 21, 2021
Language:	English

Related Products

- Oracle Database Cloud Service
- Oracle Database Backup Service
- Oracle Cloud Infrastructure - Database Service
- Exadata Database Machine X2-2 Hardware
- Gen 1 Exadata Cloud at Customer (Oracle Exadata Database Cloud Machine)

Show More

Information Centers

Document References

Recently Viewed

- Exadata Critical Issues [1270094.1]
- Exadata Database Machine and Exadata Storage Server Supported Versions [888828.1]

EXADATA - PATCH

Step 1.2.3.1 - Grid Infrastructure Upgrade to 19c

Required patches in the target Grid Infrastructure Home:

From Release	Fix required in target GI home	Description	Additional comments
18c	none		
12.2.0.1	none		
12.1.0.2	30272995	PRE UPGRADE CLUVFY FAILS TO DETECT PATCHES APPLIED IN ENV	Required for upgrades to 19.8.0.0.0.200714 and 19.9.0.0.0.201020

Required patches in the source Grid Infrastructure Home:

From Release	Fix required in source GI home	Description	Additional comments
18c	none		
12.2.0.1	none		
12.1.0.2	21255373	CSSD : DUPLICATE RESPONSE IN GROUP DATA UPDATE	Fix included in 12.1.0.2.190115

Step 1.2.3.2 - Database Upgrade to 19c

From Release	Fix required in source DB home	Fix required in target DB home	Fix required in GI home	Description	Additional comments
18c	28553832			LNx-191-EXA:CSSD PROCESS CANNOT GET REAL-TIME PRIORITY	
		29213893		DBMS_STATS FAILING WITH ERROR ORA-01422 WHEN GATHERING STATS FOR USER\$ TABLE	Only required if the following query returns multiple rows (select count(*) from obj\$ where name='USER\$;')
12.2.0.1	28553832			LNx-191-EXA:CSSD PROCESS CANNOT GET REAL-TIME PRIORITY	
		29213893		DBMS_STATS FAILING WITH ERROR ORA-01422 WHEN GATHERING STATS FOR USER\$ TABLE	Only required if the following query returns multiple rows (select count(*) from obj\$ where name='USER\$;')
12.1.0.2	21255373			CSSD : DUPLICATE RESPONSE IN GROUP DATA UPDATE	Fix included in 12.1.0.2.190115
	28553832			LNx-191-EXA:CSSD PROCESS CANNOT GET REAL-TIME PRIORITY	
		29213893			
11.2.0.4			31561819	INCOMPATIBLE MAXMEMBERS AT CSSD LEVEL CAUSING DATABASE INSTANCE NOT ABLE TO START	Fix included in GI 19.9.0.0.0
		29213893		LNx-191-EXA:CSSD PROCESS CANNOT GET REAL-TIME PRIORITY	Only required if the following query returns multiple rows (select count(*) from obj\$ where name='USER\$;')

EXADATA - PATCH

☆ Exadata 21.2.5.0.0 release and update (33455086) (Doc ID 2804134.1)

APPLIES TO:

Oracle Exadata Storage Server Software - Version 21.2.0.0.0 to 21.2.5.0.0 [Release 21]
x86 64 bit

DETAILS

Highlights

Review "[What's New in Oracle Exadata Database Machine 21.2.0](#)" in the [Oracle Exadata Database Machine System Overview](#) before updating to this release.

Software Release Requirements

Exadata Storage Server Software 21.2.5.0.0 supports the following minimum Oracle Grid Infrastructure and Database software releases:

- Oracle Grid Infrastructure:
 - 19.4.0.0.0.190716 *
 - 18.7.0.0.0.190716 *
 - 12.2.0.1.0.190716 *
 - 12.1.0.2.0.190716 *

- Oracle Database:
 - 19.4.0.0.0.190716 *
 - 18.7.0.0.0.190716 *
 - 12.2.0.1.0.181016
 - 12.1.0.2.0.180831
 - 11.2.0.4.0.180717

EXADATA - PATCH

Document 2804134.1

support.oracle.com/epmos/faces/DocumentDisplay?_afLoop=402909686528060&parent=DOCUMENT&sourceId=888828.1&id=2804134.1&_afWindowMode=0&_adf.ctrl-state=1d3qkyne7f_507

Exadata 21.2.5.0.0 compatibility

- Database servers using Xen virtualization require the following:
 - All user domains (domUs) must be updated to Exadata 12.2 or higher before updating the management domain (dom0) to Exadata 12.2 or higher.

Refer to [Note 888828.1](#) for additional details relating to the following:

- Patch-level requirements and recommendations for Oracle Database and Grid Infrastructure releases and patch sets
- Minimum version requirements to support specific Exadata features and patching methods
- References to guidelines for products related to Exadata (e.g. Oracle SuperCluster, Oracle Enterprise Manager, Oracle ZFS Storage Appliance)

Exadata 21.2.5.0.0 Software and Image files

- Software and Image files for upgrade
 - [Patch 33455086](#) - Storage server software (21.2.5.0.0.211013)
 - [Patch 33306814](#) - Admin/RDMA network switch (7.0(3)I7(9)) and InfiniBand network switch (2.2.16-2) software
 - [Patch 33306815](#) - Database server bare metal / KVM / Xen domU ULN exadata_dbserver_21.2.5.0.0_x86_64_base OL7 channel ISO image (21.2.5.0.0.211013)
 - [Patch 33306816](#) - Database server Xen dom0 ULN exadata_dbserver_dom0_21.2.5.0.0_x86_64_base OVM channel ISO image(21.2.5.0.0.211013)
- Image files for new deployment or re-imaging
 - [Patch 33306817](#) - Xen DomU System.img OS image for 21.2.5.0.0 VM creation on upgraded Xen dom0 (21.2.5.0.0.211013)
 - [Patch 33306818](#) - KVM guest System.img OS image for 21.2.5.0.0 VM creation on upgraded KVM Host (21.2.5.0.0.211013)
 - [Patch 33306820](#) - Storage cell USB image (21.2.5.0.0.211013)
 - [Patch 33306821](#) - Database server USB image (21.2.5.0.0.211013)
 - Part V1008909-01 - Storage cell ISO/PXE image (21.2.5.0.0.211013)
 - Part V1008910-01 - Database server ISO/PXE image (21.2.5.0.0.211013)
- Image file for Secure Eraser
 - [Patch 33306824](#) - Secure Eraser image (21.2.5.0.0.211013)
- Image file for diagnostics
 - [Patch 33306825](#) - Diagnostic ISO image (21.2.5.0.0.211013)

Known Issues

This section describes known issues and workarounds. Review this note for the latest issues, workarounds and tips for this patch. This note is updated regularly after the patch release with any new issues encountered and solutions or workarounds for them.

Issues Discovered in This Release

None

Exadata 21.2.5.0.0 release and update [2804134.1]
Critical Patch Update (CPU) Program Jul 2021 Patch Availability Document (PAD) [2773670.1]
Critical Patch Update (CPU) Program Oct 2021 Patch Availability Document (PAD) [2796575.1]
Oracle Database 19c Release Update & Release Update Revision October 2021 Known Issues [19202110.9]
19c Grid Infrastructure and Database Upgrade steps for Exadata Database Machine running on Oracle Linux [2542082.1]
Show More

EXADATA - PATCH

- Execution:
 - Check **from where you will call** the patch:
 - Can be from one dbnode inside Exadata or external machine.
 - You will need SSH equivalence for root to everything (switches, storage, dbnodes).
 - Runs exachk/AHF before patch.
 - What will be the source:
 - ISO.
 - Repository:
 - Needs to create local YUM/OEL repository:
 - Needs specific registration to download all the packages - It is slow!!!
 - Good for huge environments.

EXADATA - PATCH

Oracle Exadata Database Machine

Help Center

Engineered Systems / Exadata Database Machine

Maintenance Guide

Table of Contents

- Title and Copyright Information
- Preface
- 1 General Maintenance Information
- 2 Maintaining Exadata Database Servers
- 3 Maintaining Oracle Exadata Storage Servers
- 4 Maintaining Other Exadata Database Machine Components
- 5 Managing Oracle VM User Domains
- 6 Managing Oracle Linux KVM Guests
- 7 VM_MAKER Command Reference
- 8 Updating Exadata Software**
- 9 Using the DBMCLI Utility
- 10 Using the ExaCLI Utility
- 11 Using the exadcli Utility

Oracle® Exadata Database Machine

Maintenance Guide

22.1

F29250-25

September 2022

Title and Copyright Information

EXADATA - PATCH

Patchmgr Syntax for Database Servers

```
./patchmgr --dbnodes database_node_file
{ --backup [--rolling] [--unkey] |
  --precheck { --yum_repo base_URL | --iso_repo zipped_iso_file } [ --
target_version version ] [--unkey] |
  --upgrade { --yum_repo base_URL | --iso_repo zipped_iso_file } [ --
target_version version ] [--rolling] [--unkey] |
  --complete [ --target_version version ] [--unkey] |
  --rollback [--rolling] [--unkey] |
  --cleanup [--unkey] }
[ --log_dir { log_directory | auto } ]
```

Exadata 21.2.5.0.0 Software and Image files

- Software and Image files for upgrade
 - [Patch 33455086](#) - Storage server software 21.2.5.0.0.211013
 - [Patch 33306814](#) - Admin/RDMA network switch (7.0(3)17(9)) and InfiniBand network switch (2.2.16-2) software
 - [Patch 33306815](#) - Database server bare metal / KVM / Xen domU ULN exadata_dbserver_21.2.5.0.0_x86_64_base OL7 channel ISO image (21.2.5.0.0.211013)
 - [Patch 33306816](#) - Database server Xen dom0 ULN exadata_dbserver_dom0_21.2.5.0.0_x86_64_base OVM channel ISO image(21.2.5.0.0.211013)
- Image files for new deployment or re-imaging
 - [Patch 33306817](#) - Xen DomU System.img OS image for 21.2.5.0.0 VM creation on upgraded Xen dom0 (21.2.5.0.0.211013)
 - [Patch 33306818](#) - KVM guest System.img OS image for 21.2.5.0.0 VM creation on upgraded KVM Host (21.2.5.0.0.211013)
 - [Patch 33306820](#) - Storage cell USB image (21.2.5.0.0.211013)
 - [Patch 33306821](#) - Database server USB image (21.2.5.0.0.211013)
 - Part V1008909-01 - Storage cell ISO/PXE image (21.2.5.0.0.211013)
 - Part V1008910-01 - Database server ISO/PXE image (21.2.5.0.0.211013)

Version	Software
21.2.5.0.0	Patch 33455086 - Storage server software 21.2.5.0.0.211013 Patch 33306814 - Admin/RDMA network switch (7.0(3)17(9)) and InfiniBand network switch (2.2.16-2) software Patch 33306815 - Database server bare metal / KVM / Xen domU ULN exadata_dbserver_21.2.5.0.0_x86_64_base OL7 channel ISO image (21.2.5.0.0.211013) Patch 33306816 - Database server Xen dom0 ULN exadata_dbserver_dom0_21.2.5.0.0_x86_64_base OVM channel ISO image (21.2.5.0.0.211013)

EXADATA - PATCH

- Patch Cell from 19.2.19.0.0.201013 to 20.1.8.0.0.210317

```
[DOM0 - root@dbm0dbadm01 patches]$ ./patch_20.1.8.0.0.210317/patchmgr -cells /root/all_cell -patch -rolling
...
2021-04-07 12:27:08 +0200 1 of 5 :Working: Initiate patch on cells. Cells will remain up. Up to 5 minutes ...
...
2021-04-07 13:03:35 +0200 Done dbm0celadm01 :SUCCESS: Execute plugin check for Post Patch.
...
2021-04-07 15:14:05 +0200          :INFO      : Exiting.

[DOM0 - root@dbm0dbadm01 patches]$
```

- Patch cell from 18.1.12.0.0.190111 to 21.2.3.0.0.210817

```
[root@dbm0dbadm01 storage]# ./patch_21.2.3.0.0.210817/patchmgr -cells /root/cell_group -patch -rolling
...
2021-09-18 20:39:57 -0300 1 of 5 :Working: Initiate patch on cells. Cells will remain up. Up to 5 minutes ...
...
2021-09-18 21:25:23 -0300 Done dbm0celadm01 :SUCCESS: Execute plugin check for Post Patch.
...
2021-09-18 22:52:36 -0300          :INFO      : Exiting.

[root@dbm0dbadm01 storage]#
```

- Patch time depends how old you are.

EXADATA - PATCH

- Patch DB from 19.2.19.0.0.201013 to 20.1.8.0.0.210317

```
[DOM0 - root@dbm0dbadm02 patches]$ ./dbserver_patch_20.210314/patchmgr -dbnodes /root/dbm0dbadm01 -upgrade --iso_repo /patches/p32459081_201000_Linux-x86-64.zip --target_version 20.1.8.0.0.210317
...
2021-04-08 10:39:43 +0200      :INFO      : Checking hosts connectivity via ICMP/ping
...
2021-04-08 11:33:46 +0200      :SUCCESS: Completed run of command: ./dbserver_patch_20.210314/patchmgr -dbnodes /root/dbm0dbadm01 -upgrade --iso_repo /patches/p32459081_201000_Linux-x86-64.zip -
-target_version 20.1.8.0.0.210317
...
2021-04-08 11:33:47 +0200      :INFO      : Exiting.

[DOM0 - root@dbm0dbadm02 patches]$
```

- Patch DB from 18.1.12.0.0.190111 to 21.2.3.0.0.210817

```
[root@dbm0dbadm02 db]# ./dbserver_patch_21.210820/patchmgr -dbnodes /root/dbm0dbadm01 --upgrade --iso_repo /patches/p33129611_212000_Linux-x86-64.zip --target_version 21.2.3.0.0.210817
...
2021-09-19 02:12:35 -0300      :INFO      : Checking hosts connectivity via ICMP/ping
...
2021-09-19 04:17:15 -0300      :SUCCESS: Completed run of command: ./dbserver_patch_21.210820/patchmgr -dbnodes /root/dbm0dbadm01 --upgrade --iso_repo /patches/p33129611_212000_Linux-x86-64.zip
--target_version 21.2.3.0.0.210817
...
2021-09-19 04:17:15 -0300      :INFO      : Exiting.

[root@dbm0dbadm02 db]#
```

- Again, patch time depends how old you are.

EXADATA - PATCH - RESUME

- **Plan and Preparation:**

1. Check HW errors.
2. Download Patches.
3. Reset Hardware/ILOM.
4. Check HW errors.
5. Pre-check.

- **Executing:**

1. Apply needed RU/PSU.
2. Patch Switches.
3. Patch Storage Cells.
4. Patch DB Nodes:
 - Bare Metal.
 - Patch OVM/KVM:
 1. Patch domU.
 2. Patch dom0.



EXADATA CLOUD@CUSTOMER - EXACC



EXACC – THE GOOD THE BAD AND THE UGLY

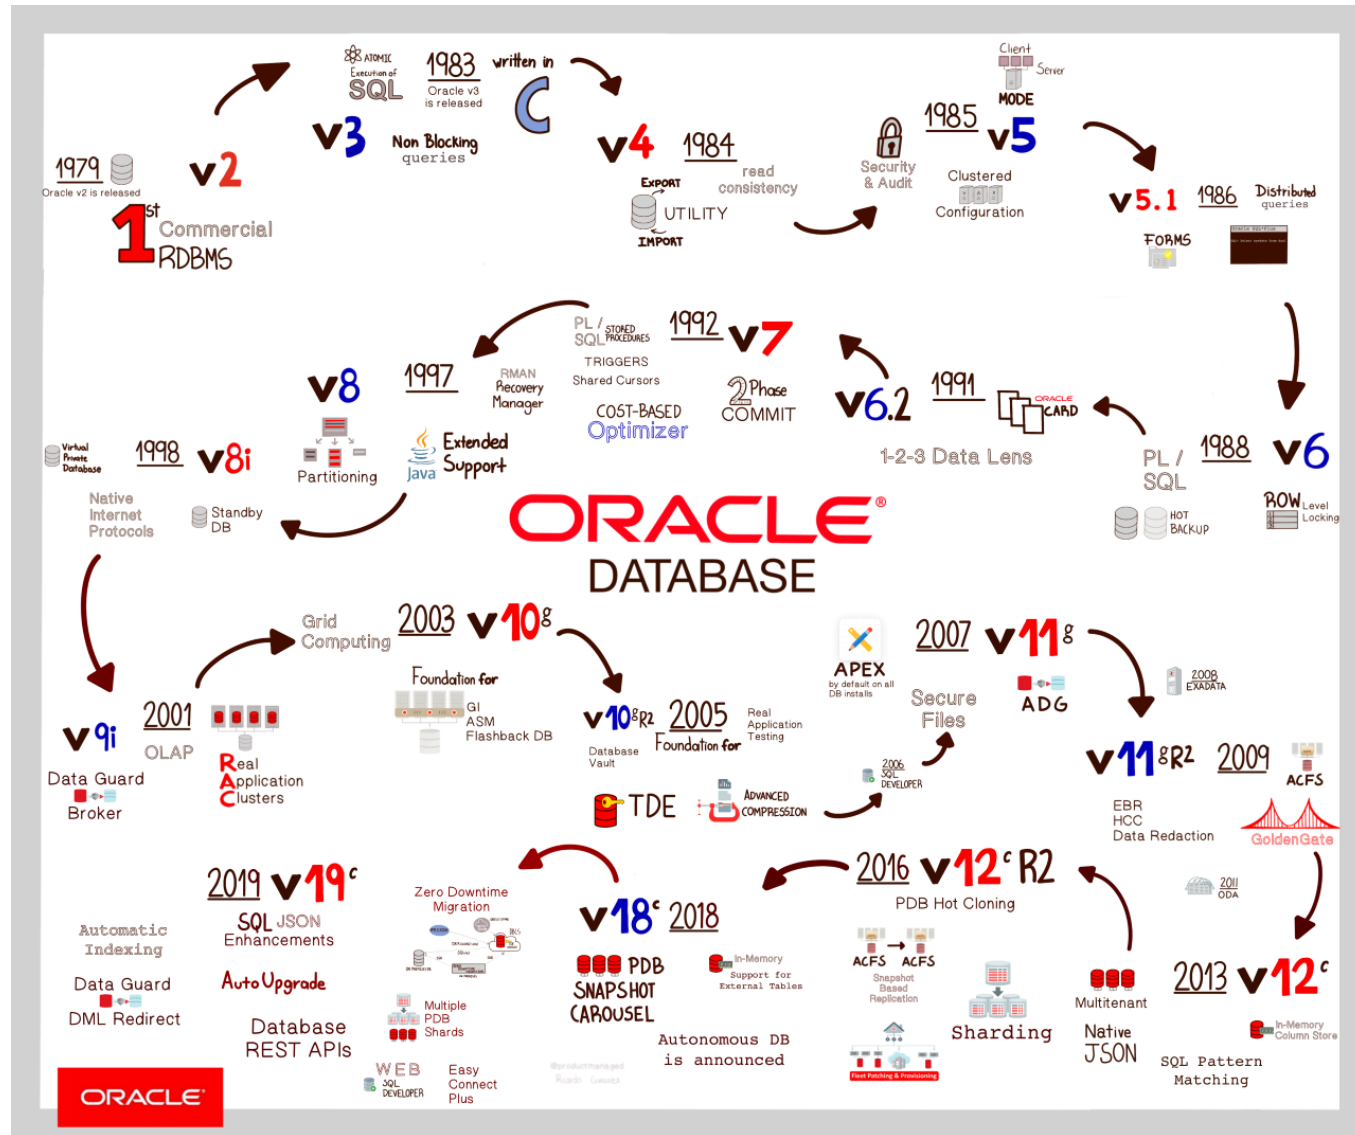
- Good for automatization and daily use:
 - Helps to reduce bureaucratic activities like patches.
- Good integration and variety of tools:
 - DBAAS tools/options cover most of the daily activities.
- Be careful with CPU scale up/down.

EXACC – THE GOOD THE BAD AND THE UGLY

- **Check your requirements**, maybe your migration will not be transparent:
 - If you have multiple DB_DOMAIN's, this can block you.
- Not support the same (AS IS) what you have at On-Prem:
 - Block size is one example.
 - Different templates can be tricky to change.
- Needs improvement:
 - DBAASPI deserves better documentation.
 - REST API is not good (yet) to create databases.

EXACC - CLOUD AT CUSTOMER

- What you need to take care of:
 - Cost calculated by CPU usage.
 - Size correctly your ExaCC:
 - Everything is based on the number of dbnodes.
 - CPUs (that you need to increase) * number of the nodes
 - Node subsetting released in 2021
 - Be careful with multi-VLAN:
 - Just accept one client VLAN per cluster.
 - So maybe you will need more clusters, and more clusters, more CPU allocation.
 - Read the documentation and the way that databases are created:
 - Difficult to escape from OCI tools.



Credit: Ricardo Gonzalez - Senior Principal Product Manager at Oracle - <https://www.linkedin.com/posts/activity-6719433973749166080-VGoR/>