


The background is a vertical split. The left half is a light beige color with faint, organic, brush-stroke-like shapes and a pattern of small, light grey ovals. The right half is a dark purple and black color with a complex, abstract pattern of concentric, dotted lines and a dense, woven texture of orange and yellow lines.

Healthcheck report for Oracle Database

Oracle Cloud Systems Hub
June 2021

What is the Healthcheck service?



The Oracle Healthcheck Service offers a **FREE diagnostic report for your database environment.**

With the metric data files that you provide, our Oracle Cloud Systems Hub engineers can analyze the system and database performance and overall health.

What data sources can I use for data collection?


AWR Miner

StatsPack

**AWR classic
HTML report**

**EXAchk
ORAchk**

FOUR SIMPLE STEPS



STEP 4. SHOWCASE

We share thoughts on the analysis results and plan together for the future. As trusted advisors, we come with many recommendations and proposals

STEP 3. ANALYSIS

Our Solution Engineer teams digest the data and construct a presentation for pointing out the problems and solutions for your database systems

STEP 2. COLLECTION

You collect diagnostic information from your systems via an SQL script and pass it on for analysis

STEP 1. DISCOVERY

By engaging with our sales teams by requesting this FREE service, you identify which systems are viable for a check-up

What will I receive?

An **analysis of the system** will be provided in a presentation along with relevant database **best practice recommendations** and if needed, a **proposal** for system upgrade or refresh.

The only requirement is running an Oracle certified SQL script on your database environment which collects diagnostic data from the database AWR repository, or provide your own AWR reports.

**No user or confidential data
will be queried.**



AGENDA

System Overview

A brief recap of the current infrastructure, system characteristics and resources allocated

Healthcheck Analysis

The analysis of the diagnostic data provided, structured by main areas of utilization

Recommendations

Our advice on what can be improved or needs further investigation

Solution proposal

A one-in all solution to all the problems discovered before



SYSTEM OVERVIEW

CURRENT ENVIRONMENT:

- 3 Exadata database machines:
- X4-2 Full Rack (with X6 and X7 upgrades)
- X6-2 Half Rack with virtualization
- X7-2 Eight Rack for FCCMDB instance
- Database version for all instances is aligned at 12.1.0.2.0
- Out of all databases analyzed, only 3 have serious performance issues:
DBM01, PROD, OBIEREPO
- The majority of instances have low resource consumption

Database Name	Database Version
DBM01	12.1.0.2.0
HYPERNDB	12.1.0.2.0
IBPSFSS	12.1.0.2.0
SOA	12.1.0.2.0
HYPEDEV	12.1.0.2.0
OBIEREPO	12.1.0.2.0
OFSAPROD	12.1.0.2.0
FBNOFSA	12.1.0.2.0
TESTEBS	12.1.0.2.0
PROD	12.1.0.2.0
FCCMUAT	12.1.0.2.0
FBNSIT	12.1.0.2.0
MEDATDB	12.1.0.2.0
OFSADDEV	12.1.0.2.0
EBSTST	12.1.0.2.0
OFSAUAT	12.1.0.2.0
OBIEPROD	12.1.0.2.0
DISC	12.1.0.2.0
IBPSPMH	12.1.0.2.0
AGENCYDB	12.1.0.2.0



SYSTEM OVERVIEW

MULTIPLE WAIT EVENTS IDENTIFIED

Administrative wait events regularly occurring

DATABASE VERSION UP TO DATE

RDBMS Software is 19.0.0

No database upgrade is recommended.

BACKUP STRATEGY TO BE REVISED

Most of the waits are related to Backup activity.

HARDWARE UPGRADE CONSIDERED

The system is clearly busy and during high activity, the CPU is getting close to 100%.

Memory consumption is also high and during backup processing, the database is facing real performance issues.

CPU USAGE IS PEAKING AT 100%

CPU usage varies from to 6% at low to 100% high with 27% average.

AREAS FOR POTENTIAL IMPROVEMENT

AGENDA

System Overview

A brief recap of the current infrastructure, system characteristics and resources allocated

Healthcheck Analysis

The analysis of the diagnostic data provided, structured by main areas of utilization

Recommendations


Our advice on what can be improved or needs further investigation

Solution proposal

A one-in all solution to all the problems discovered before



CPU UTILIZATION




0% - 5% Database CPU

Overall Host CPU utilization can be seen on the graphic on the left side:

- **73% / 49% / 38%** Peak CPU
- **20% - 25%** Average CPU
- **2 – 6%** Low CPU

MEMORY UTILIZATION




DATABASE MEMORY

Database Memory utilization varies between 16 and 72GB RAM as shown in the left side chart.


HOST MEMORY

Overall Host memory utilization can be seen on the graphic below:



DATABASE SPACE UTILIZATION



Database Name	Database Size (GB)
baseq044	9318,57
baseq042	1631,01
basep051	1674,02
basep085	4689,28
basep055	4572,36
basep031	6929,91
baseq046	4371,64
basei112	9065,67
basep053	9418,85
baseq018	2866,78
baseq001	7527,19



62065,26GB





*for the databases analyzed – not total on the machine

STORAGE UTILIZATION





Exadata Storage	TB
Total raw capacity	943.37
Cell Failure Coverage*	Disk
Available raw capacity	943.37
Available space after redundancy	471.68
Database used	314.83
Database free	129.57
FRA used	23.46
FRA free	22.68
DBFS allocated	3.82


AVERAGE IO UTILIZATION



PEAK IO UTILIZATION



DATABASE WAIT EVENTS COMPARED




Concurrency: A high number of wait events can be seen here due to sessions being blocked, while waiting on the same resource.

CPU: DB CPU wait events are not particularly indicating a performance degradation. Having a high percentage of DB CPU means that the database is processing database user-level calls.

User IO: wait events mainly consist of events like “dbfile sequential read” and “dbfile scattered read” which translates to index usage and full table scans. It is advised to make sure that common queries have their execution plans optimized, in order to reduce IO operations and maintain database performance. In this instance’s case, User IO wait time is very low.

Other: These wait events resulted from unusual errors and internal database engine processes being delayed. Always check alert log for internal errors and abnormal behavior after having performance problems.

DATABASE WAIT EVENT SUBCATEGORIES



BIGGEST ISSUES DETECTED:

19% backup contention


18% slow storage IO

6% direct path reads

⇒ **43%**

Total performance degradation due to storage

WAIT EVENTS – DB CPU AND USER IO




The following graphic represents the percentage of DB Time used on DB CPU and User IO.

Having a high percentage of **DB CPU** means that the database is processing database user-level calls and it is a sign of healthy database activity.

User IO is mainly consisting of the bellow events:

- “direct path read”, “direct path write”, “ direct path read temp”
https://docs.oracle.com/database/121/TGDBA/pfgrf_instance_tune.htm#TGDBA94485
- “direct path write temp” waits are typically caused by sort operations that cannot be completed in memory, spilling over to disk.
https://support.oracle.com/epmos/faces/DocumentDisplay?_afLoop=132051023512846&parent=EXTERNAL_SEARCH&sourceId=PROBLEM&id=1576956.1&_afrWindowMode=0&_adf.ctrl-state=1brk1dwyfh_4
- “db file sequential read” - Is one of the most common wait event.
https://docs.oracle.com/database/121/TGDBA/pfgrf_instance_tune.htm#TGDBA94485
- “db file scattered read” A scattered read is usually a multiblock read. It can occur for a fast full scan (of an index) in addition to a full table scan.
https://docs.oracle.com/database/121/TGDBA/pfgrf_instance_tune.htm#TGDBA94479

WAIT EVENTS – CONCURRENCY/APPLICATION



The following graphic shows the concurrency and application waits on each of the most used databases in the analysis.

These wait events are caused either by sessions running user application code (for example row level locking or explicit lock commands) or waits for internal database resources (library cache locks or latches).


Application is mainly consisting of the bellow events:

- *enq: TM – contention” - indicates the time spent waiting for a TM lock that is used to coordinate activities for many base table / partition operations. To resolve this type of event, you can follow the recommendation from the My Oracle Support [Doc ID 1905174.1](#)*
- *“enq: TX – row lock contention” - Reducing this type of waits typically involves altering application code to avoid the contention scenario/s. . To resolve this type of event, you can follow the recommendation from the My Oracle Support [Doc ID 1966048.1](#)*

Concurrency is mainly consisting of the bellow events:

- *“library cache lock” - To optimize the functioning of the database and diminish the library cache lock waits you can follow the recommendation from the My Oracle Support [Doc ID 34578.1](#)*
- *“cursor: pin S wait on X” - To optimize the functioning of the database and diminish the cursor: pin S wait on X waits you can follow the recommendation from the My Oracle Support [Doc ID 1298015.1](#)*

WAIT EVENTS – NETWORK/OTHER







Network: Waits related to network messaging. It is advised to investigate database links before looking into the network first. Usually, network waits are related to the time it takes to pack messages for the network before they are sent. 95% of the cases it is a database link issue.

Configuration: Waits caused by inadequate configuration of database or instance resources. Database initialization parameters must be sized appropriately for the resources allocated to the server. It is advised to check for undersized init parameters and use advisory statistics for guidance.

Other: Wait events resulted from unusual errors. Always check alert log for internal errors and abnormal behavior.

MAIN ACTIVITY VARIABLES



AGENDA

System Overview

A brief recap of the current infrastructure, system characteristics and resources allocated

Healthcheck Analysis

The analysis of the diagnostic data provided, structured by main areas of utilization

Recommendations

Our advice on what can be improved or needs further investigation

Solution proposal

A one-in all solution to all the problems discovered before



NOTABLE EVENTS

Recurring Network events
consist of **7%** activity on
instance basep055

Constant high commit rate of **21%**
total on instance basep053

Concurrency (pin S wait on X)
on basei112 – **5%** total
database activity

It is recommended to investigate the issues
mentioned above to prevent further database
performance degradation


Backup Restore throttle sleep
and latch: redo allocation +
concurrency events on instance
basep085 – **20%** of total
activity

Significant application
Spikes – **25%** on instance baseq001

MEMORY ADVISOR RECOMMENDATIONS


+32%

Consider increasing SGA with ~10GB on basep051 for a 32% decrease in Physical Reads.




+45%

Consider increasing SGA with ~4GB on basep085 and PGA with ~0.8GB for a 10% and 35% Physical Reads and MB Read/Written to disk.




+61%

Consider increasing SGA with ~10GB on baseq046 for a 60-61% decrease in Physical Reads.



+17%

Consider increasing PGA with ~1.1GB on baseq044 for a 17% increase in MB R/W performance.



EXACHK OBSERVATIONS

Status: CRITICAL | Type: Database Check

Message	Status On	Solution
Database parameter CLUSTER_INTERCONNECTS is not set to the recommended value	od01db01:captdb, od01db02:DEV2, od01db02:captdb	Database parameter CLUSTER_INTERCONNECTS should be a colon delimited string of the IP addresses returned from sbin/ifconfig for each cluster_interconnect interface returned by oifcfg.
One or more disk groups which contain critical files do not use high redundancy	All Databases	Solution not yet available...
Database control files are not configured as recommended	All Databases	Place the control file in DATA because both DATA and RECO are high redundancy disk groups

Status: CRITICAL | Type: OS Check

Message	Status On	Solution
System is exposed to Exadata Critical Issue EX44	All Database Servers	Solution not yet available...
System is exposed to Exadata Critical Issue DB41	All Database Servers	Solution not yet available...
One or more InfiniBand network parameters on Database Servers are not as expected	All Database Servers	Solution not yet available...
TCP Selective Acknowledgment is not enabled	od01db01	Solution not yet available...
The vm.min_free_kbytes configuration is not set as recommended	All Database Servers	If the value is incorrect in /etc/sysctl.conf and matches the value in active memory, edit the /etc/sysctl.conf file to include the line 'vm.min_free_kbytes = 9correct_value_for_your_configuration.' and reboot the database server.

SUMMARY OF FINDINGS

Findings	No of events	Exadata's proposed solution
High CPU Utilization	3	IORM, RDMA, PMEM, RoCE, Smart Flash Cache, Exadata Smart Scan
Active Sessions Surge	5	
Waits on Network	3	Exafusion , Smart Fusion Block Transfer, RDMA, PMEM, RoCE
Waits on Cluster	7	
Waits on Commit	1	RDMA, PMEM, RoCE
Increased Redo Generation Rate	4	
Increased Physical Reads	6	Exadata Smart Scan, Storage Indexes, HCC, In-Memory Columnar, Smart Flash Cache
Increased Waits on I/O	1	
Configuration Problems	2	Exadata Automatic Indexing
Other wait events	4	Exadata Smart Scan, Storage Indexes, HCC, In-Memory Columnar Cache
Concurrency Issues	14	Automatic Diagpack Upload for Oracle ASR
Application generated waits	8	Faster Performance for Large Analytic Queries and Large Loads

