

Oracle Cloud Performance Best Practices

Tuning, Diagnostics and Issue Resolution

Deb Bhattacharjee

Nikolai Potapov

Nathan Reynolds

PSR – Performance, Scalability and Reliability

October 3rd, 2017

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Program Agenda

- 1 Introduction
- 2 Customer Stories
- 3 Trending Mistakes
- 4 Best Practices
- 5 Monitoring and Diagnostics

Introduction

Speaker Picture



Deb Bhattacharjee
Senior Director, Oracle

Speaker Picture



Nikolai Potapov


Architect, Oracle

Speaker Picture



Nathan Reynolds

Architect, Oracle



Cloud Platform



Oracle Cloud Platform



Develop & Deploy



Integrate & Extend



Publish & Engage



Analyze & Predict



Secure & Manage

Innovate with a
**Comprehensive, Open,
Integrated and Hybrid**
Cloud Platform
that is
**Highly Scalable, Secure
and Globally Available**

Oracle Cloud Platform

Comprehensive

Open





Integrated





Hybrid

Oracle
Public Cloud



Oracle
Data
Center

-  Data Management
-  Application Development
-  Enterprise Integration
-  Data Integration

-  Analytics and Big Data
-  Content & Experience
-  Identity & Security
-  Systems Management

Oracle Cloud
at Customer



Built on High Performant Oracle Cloud Infrastructure

Oracle Cloud Platform Momentum

14,000+
Oracle
Cloud Platform
Customers



3,000+
Apps in the
Oracle Cloud
Marketplace



\$1.4 Billion
FY17 Oracle Cloud
Platform
Revenue
(60% YoY Growth)



10 PaaS
Categories where
Oracle is a **Leader**
According to
Industry
Analysts



Customer Stories

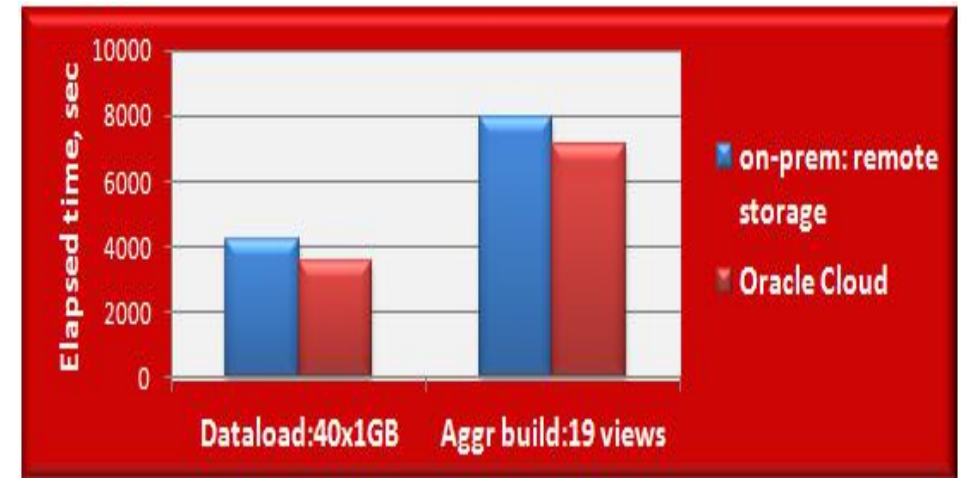
A man with glasses and a denim shirt is sitting at a table, gesturing with his hands as he speaks to a woman in a yellow top. They are surrounded by papers, a coffee cup, and a smartphone on the table. The background shows a blurred office environment.

Customer A (retail)

- Initial requirements:
 - IO: 1,200 MB/sec and/or 15K iOPS from single VM
 - Per customer other Cloud vendors could not meet same requirements
- Results:
 - Oracle Cloud surpassed I/O requirements... before SSD based filers were available.
- Conclusions: Cloud is a different platform requiring different tunings.
 - Customer tuning: 4K stripe size may be okay for on-premise. In Cloud, stripe size needs to be aligned with Cloud infrastructure. (stripe size 64K)
 - Monitor housekeeping performance and adjust process based on findings



Customer B (consumer goods)

- Initial POC requirements:
 - Performance comparable to on-prem (Exalytics) or better
 - Enhanced security
- Done:
 - High Memory shapes, latency volumes for Essbase
 - SSH two-way tunneling for VPN POC
 - Co-location
 - LUKS for data at rest
- Conclusions: Cloud can meet even extreme performance on-prem systems



Customer C (global interconnection provider)

- Initial requirements:
 - Custom applications based on Siebel CRM performance to be at par or better with on-prem and other Cloud vendors
- Done:
 - DBaaS like optimizations on IaaS VM
 - Co-location
- Conclusions: Cloud infrastructure is ready to run even complex applications



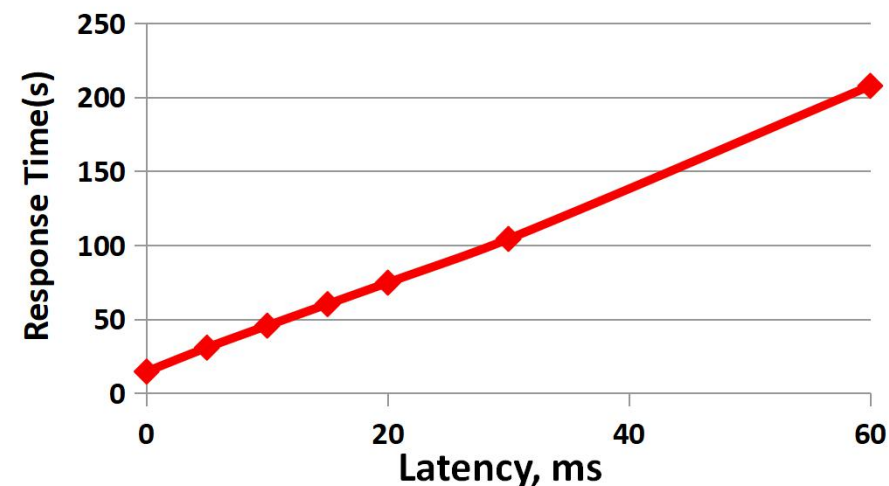
A group of people in a meeting looking at a whiteboard with sticky notes. The scene is set in a modern office environment. In the foreground, a man with a beard and a woman with long red hair are looking towards a whiteboard. The whiteboard is covered with numerous colorful sticky notes (yellow, green, pink, blue) and a hand is pointing at one of them. Other people are visible in the background, also engaged in the discussion. The overall atmosphere is collaborative and professional.

Trending Mistakes

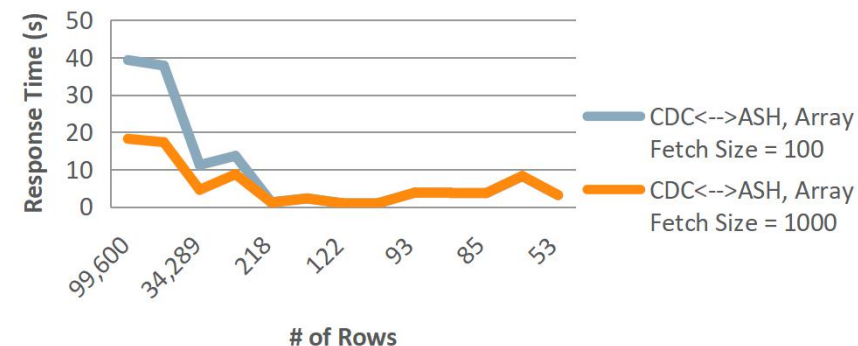
Co-location

- IaaS VM + DBaaS
 - Example: Financial applications (OLTP)
 - Even 5 ms latency caused 2x response time increase
 - Higher 30 ms latency resulted in 7x higher RT increase
 - Example: Reporting applications (OLAP)
 - Lower impact, can be mitigated by using a larger fetch array size
- PaaS VM + OSS
 - Example. DBaaS backup to remote OSS
 - 10 ms caused ~4x backup slowness
- Recommendation
 - Understand application and co-locate, if necessary
 - Even within same metro area, co-locate to same zones
 - Hybrid setup? – Move app tier closer to DB

EBS application




Response Time vs Row Count & Array Fetch Size (BICS OLAP workload)



Images

- Oracle Linux is primary OS for optimization
- Identified several kernel settings that helps I/O performance in OPC
 - Published UEK3 images have the fixes, UEK4 will have them soon
 - Older UEK3 may lack such configurations as well; Important to validate
 - Validate your VM has the known settings if you face I/O issues
- Images not based on Oracle Linux may lack such optimizations

DB custom setup on IaaS/PaaS

- Custom setups 
 - DB setup on IaaS
 - DB manual creation on DBaaS Virtual Image

- Pros
 - Full customization
 - IaaS early features availability. Example: SSD volumes
- Cons
 - **Need tuning**

- Recommendation

- If possible, use DBaaS since it is already tuned to best knowledge
- Otherwise, apply key tunings if still need custom setup
 - Put redo log and temp on latency or faster volumes
 - Init.ora: filesystemio_options = SETALL
 - Apply latest UEK tuning to improve I/O. Example - [Improving Block I/O Performance](#)

Network transfers

- Initial data transfer using single pipe
 - Example: Single SCP process
 - Even with ideal latency and bandwidth, SCP has internal limits.
- FastConnect limited by client's ISP or configuration
 - Example 1: Client had ISP bandwidth limit of 100 MB/sec.
 - Example 2: Client had egress limit due to firewall policy.
- FastConnect + Corente VPN used for bulk data uploads
 - VPN is great for daily application usage. It's not for massive data uploads.
- Recommendation
 - Use multiple pipes to upload data
 - Use OSS as yet another secure transfer option instead of VPN for bulk uploads

Storage

- Block vs. Object storage confusion

- Block storage types

- What is default or latency?
- What are SSD options?

- Recommendation.

- If I/O performance is critical, get SSD volumes
- If SSD is not available, use Latency volumes
- If need extreme I/O performance, consider High Density shapes in OCI with local SSD option (not VM failover persistent)

Update Storage Volume

Enter a name for the storage volume and specify the size (raw capacity). If you want to use this storage volume to boot your instance, select Boot Drive. [Learn more.](#)

Name: OL_6_6_UEKR3_x86_64_20170626101553_storage

Size: 12 GB (Maximum size 2,000 GB)

Storage Property: storage/default

Description: storage/default

Attach as Disk #: storage/protocol/scsi

Boot Drive: storage/snapshot/default

storage/ssd/gp1

Update Cancel



Best Practices

Image

- Windows OS
 - Change Admin password
 - Disable Guest account
 - Create 2-3 additional Admin users
 - Update PV Driver if newer version is available
 - Increase disk timeout to 300 sec
- Always use or upgrade to the latest OEL image.
- Parallelize work loads
- For high availability, place VMs on different nodes (i.e. anti-affinity)

Storage

- Logical Volume Management (LVM) for flexible resizing of volumes.
- Disable “DOS Compatible Mode”
- Throughput volumes for application binaries.
- Custom image? Apply ring buffer fix (\geq 3.8.13-68.2.2.3.el6uek.x86_64).
- Linux Unified Key Setup (LUKS) only if needed for encryption of data at rest.
- 64 KB stripe size for RAID.
- Latency volumes or SSD where I/O performance is critical.
- Bootable storage volumes (local drives are ephemeral).

Network

- Tune TCP send/receive buffers for workload.
 - Do **not** change TCP scaling (tcp_window_scaling).
- Connectivity issues? Try disabling IPv6.
- Shared network supports up to 100 subnets or IPv4 addresses.
- Unique hostnames for each VM.
- SSH Tunneling while VPN is not ready.
- Ensure all components are placed in the same data center.

Migration

- Transfer OS tunings from on-premise to VMs.
- Latency volume as a staging area.
- DB System and Program Global Area (SGA and PGA) match source DB.
- Corente VPN may slow down data transfer speed.
 - Direct transfer to Oracle Object Storage (OSS) as staging area.
- Parallel threads or processes (e.g. SCP).
- Tune TCP send/receive buffers.

Database

- Tune JDBC row prefetch size. Default 10. Keep ≤ 1000 .
- RMAN parallelism and compression to speed up backup and restores.
- Enable Linux's huge pages.
- Database as a Service (DBaaS) already tuned.
 - Use High Performance Option for production.
 - Install Diagnostics and Tuning pack.
 - Must tune custom DB.
- Assess enabling Linux's multi-page ring support.
- Place redo logs and temporary data files in latency volumes.

MySQL Database

- For heavy write workloads
 - Increase the flush thread count.
 - Disable binary logging, if not required.
- For CPU intensive workloads, use jemalloc memory allocator.
- Tune parameters innodb-io-capacity and innodb-io-capacity-max.
- Use Latin-1 over UTF-8 when possible.

Java

- GC tuning affected by number of vCPUs in the VM.
- JVM process's memory can be as high as 1.8x-2x of max heap setting.
- Keep Weblogic Admin server's VM shape and heap size to a minimum.
- Enable heap dump on OutOfMemoryError.
 - Consider configuring the JVM to exit on OutOfMemoryError.
- Compress heap dumps before transferring.
- Enable Java Flight Recorder for production profiling.

Business Intelligence Cloud Service (BICS)

- High-latency access to DB? Increase the set array size inside RPD to 1000.
- Validate datasource and init blocks.
- For Oracle Analytics Cloud, BICS and Remote Data Connector (RDC) tuning and best practices...

Oracle Analytics Cloud: Performance Best Practices for Cloud and Hybrid Access (CON7028)

Wednesday, October 4th 1:00 pm – 1:45 pm

Moscone West – Room 3009

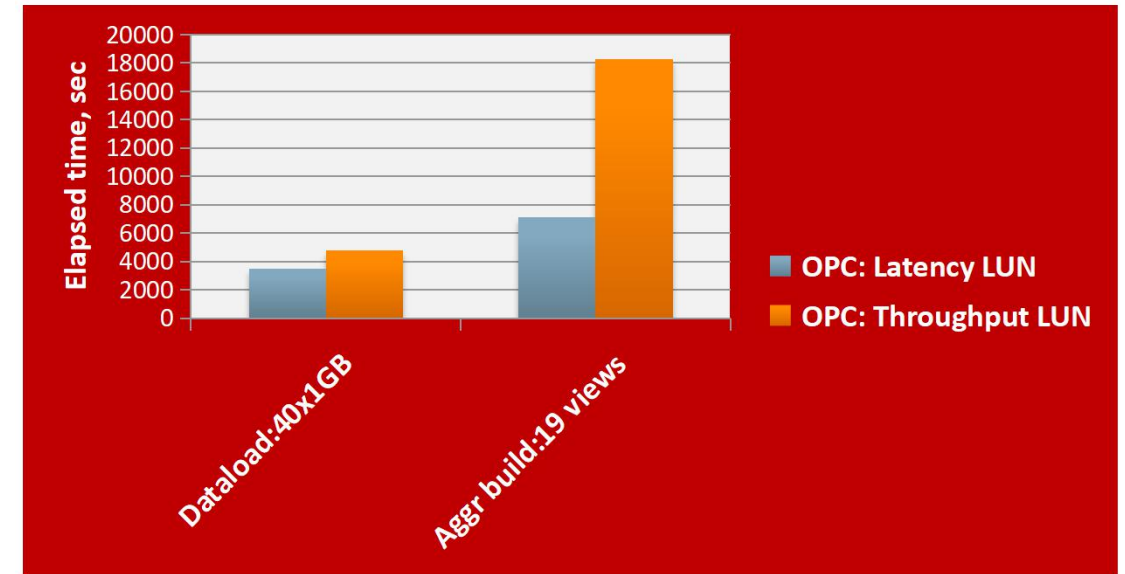
When you live in a hybrid world, it's critical to ensure that you architect your data and analytics system to accommodate where your data resides. In this session, hear about performance considerations you should plan for, in order to field the best performance analytics applications possible.

Big Data Cloud Service (BDCSCE) - SparkSQL

- Set minimum number of Executors
- Ensure Spark Executors have enough memory
 - `spark.kryoserializer.buffer=128m`
 - `spark.kryoserializer.buffer.amx=2047m`
 - `spark.executor.memory=8g`
- Reduce shuffle partitions
- Set proper timeout for long running queries
 - `spark.network.timeout=600`
- Disable Spark history, if not required

Essbase

- SSD or Latency volumes for data files.
- High Memory Shape VMs.
- Adjust # of threads to match VM shape
 - SERVERTHREADS = #vCPU
 - Same tuning for...
WORKERTHREADS, CALCPARALLEL (ASO), DLTHREADSPREPARE(ASO), EXPORTTHREADS, RESTRUCTURETHREADS, THREADS, FIXPARALLEL



Oracle Traffic Director (OTD)

- Enable response compression.
- Static files have cache control HTTP pragma.
- Adjust cache size to hold all objects.
- Enable HTTP keep alive header for client.
 - 5 seconds for REST and JET UI.
 - 61 seconds for web-based UI with few connections per user.
- Consider persistent HTTP connections to backend server.
- Max thread pool set to 20,480.
- Thread pool queue set to 3,000.

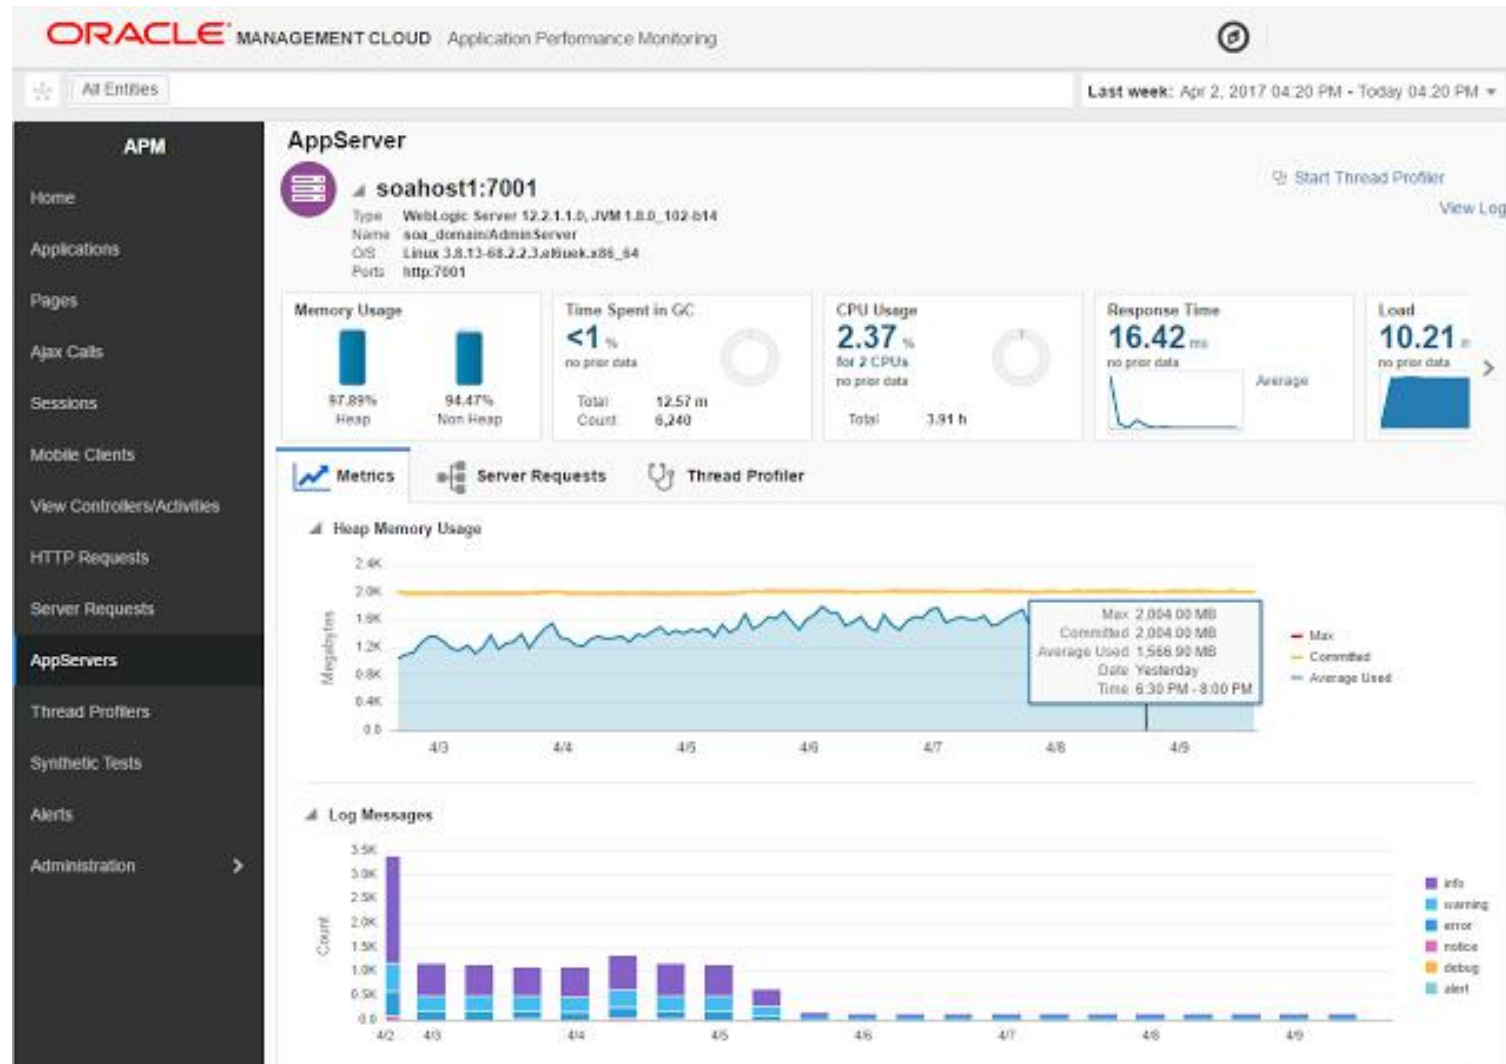
A photograph of two men in a server room. They are looking at a large computer monitor that displays a complex interface with various data points and charts. The room is dimly lit, with a desk lamp providing light. The men are wearing glasses and appear to be focused on their work.

Monitoring and Diagnostics

Monitoring Tools

- OS Watcher
- Thousand Eyes
- Oracle Management Cloud (OMC)

Application Performance Monitoring (APM)



Diagnostic Tools

- Database – swingbench
- I/O – dd, fk, vdbench
 - Example – `dd if=/dev/zero of=./testfile bs=8K count=100000 oflag=direct`
 - Example – `dd if=/dev/zero of=./testfile bs=1M count=1000`
- Network – iperf
 - Example – `iperf3 -c VM1 -i 1 -t 30 -p 8000 -O 5 -w 1M -P 8`
- CPU – SPECcpu2017, unixbench

Java

HotSpot GC Logging

- Enable GC logging
 - -XX:+DisableExplicitGC
 - -XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+PrintGCTimeStamps
 - -Xloggc:*file*
 - -XX:+UseGCLogFileRotation
 - -XX:NumberOfGCLogFiles=4
 - -XX:GCLogFileSize=5M
- Analyze GC logs with <https://gceasy.io>

Closing

Resources

- Best Practices Guide
 - http://www.oracle.com/webfolder/technetwork/tutorials/ohc/E83035_01.pdf
- Improving Block I/O Performance on DBaaS
 - <https://docs.oracle.com/en/cloud/paas/database-dbaas-cloud/csdbi/improve-block-io-performance.html>

Q&A

- What are your pain-points in OPC?
- What are your issues in OPC?

Integrated Cloud

Applications & Platform Services