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.

ORACLE REAL-WORLD PERFORMANCE

Real-World Challenges with Cloud Migrations and Proof-of-Concept Projects

ORACLE OPEN WORLD

October 1–5, 2017
SAN FRANCISCO, CA

Michael Hallas
Andrew Holdsworth

Real-World Performance

October 4, 2017

ORACLE REAL-WORLD PERFORMANCE

RWP Sessions @ OOW17 Oct 4th Rm 3012

When	Id	Topic
11am	CON6560	Optimizing Table Scans in Today's Cloud Platforms
12pm	CON6561	Migrating On-Premises Applications to the Cloud: Examining the Connection Strategy
1pm	CON6629	Real-World Challenges with Cloud Migrations and Proof-of-Concept Projects
2pm	CON6660	Applying Oracle Database 12*c* and Real-World Performance Techniques to SAP

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

The Real World Performance Perception Problem

Where database user look for performance improvements

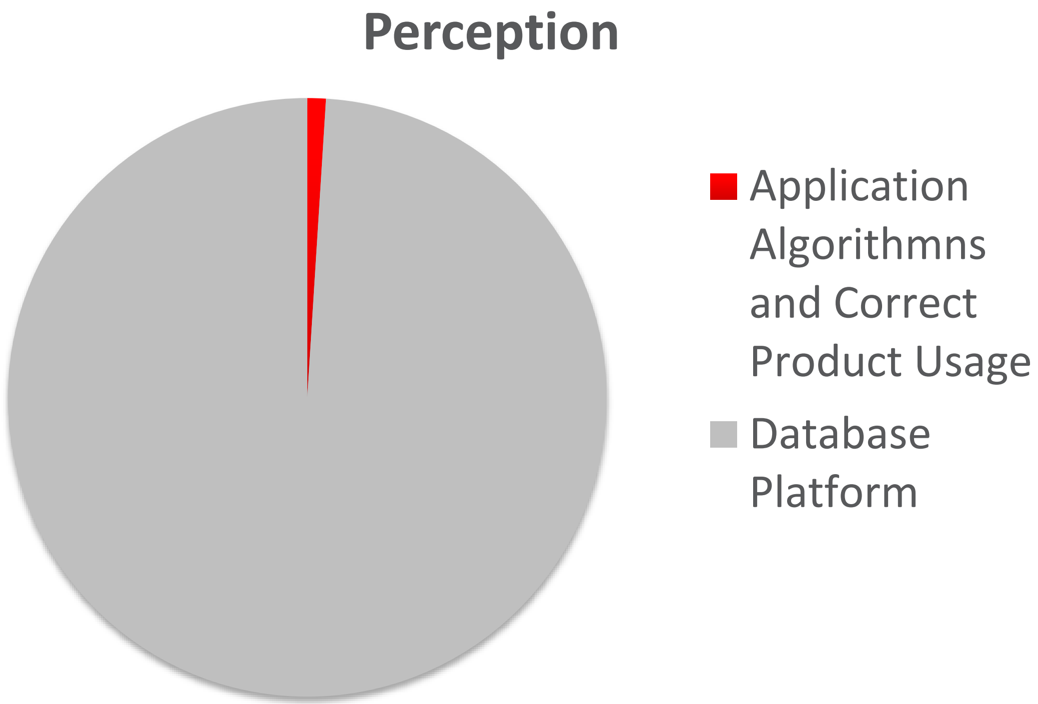

The best place to look for performance Improvements

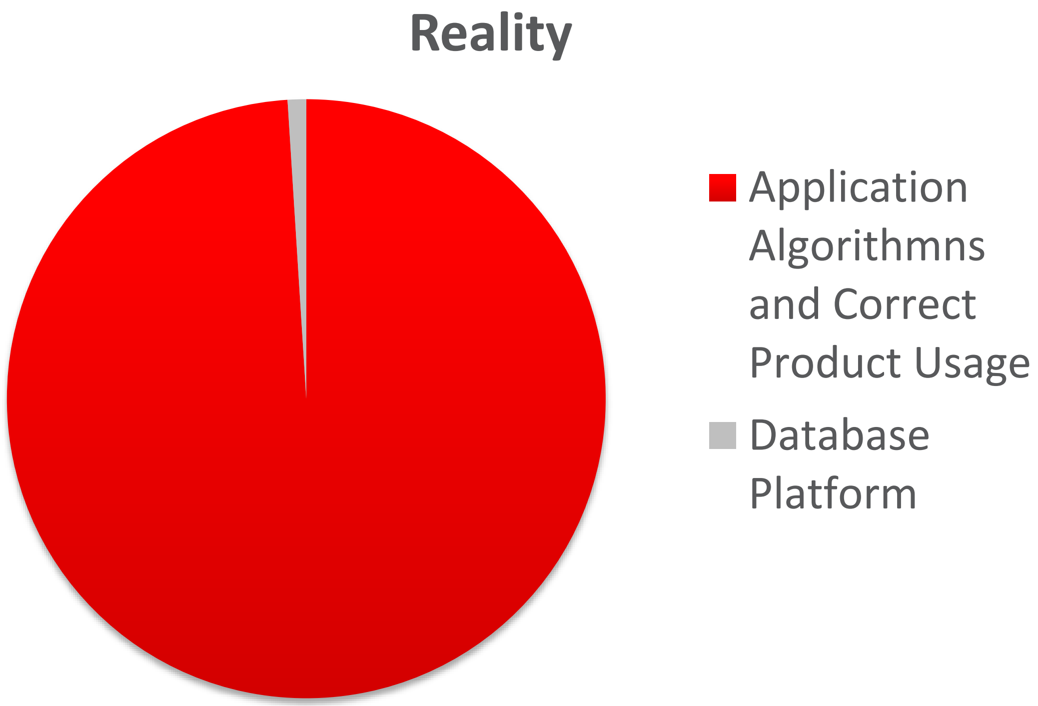

Introduction

- We understand the arguments in favour of Cloud
- Existing systems are good candidates for lift-and-shift migration to Cloud
- Cloud gives us access to lots of new and fast hardware
- How hard can it be? What could possibly go wrong?

Performance Trajectories

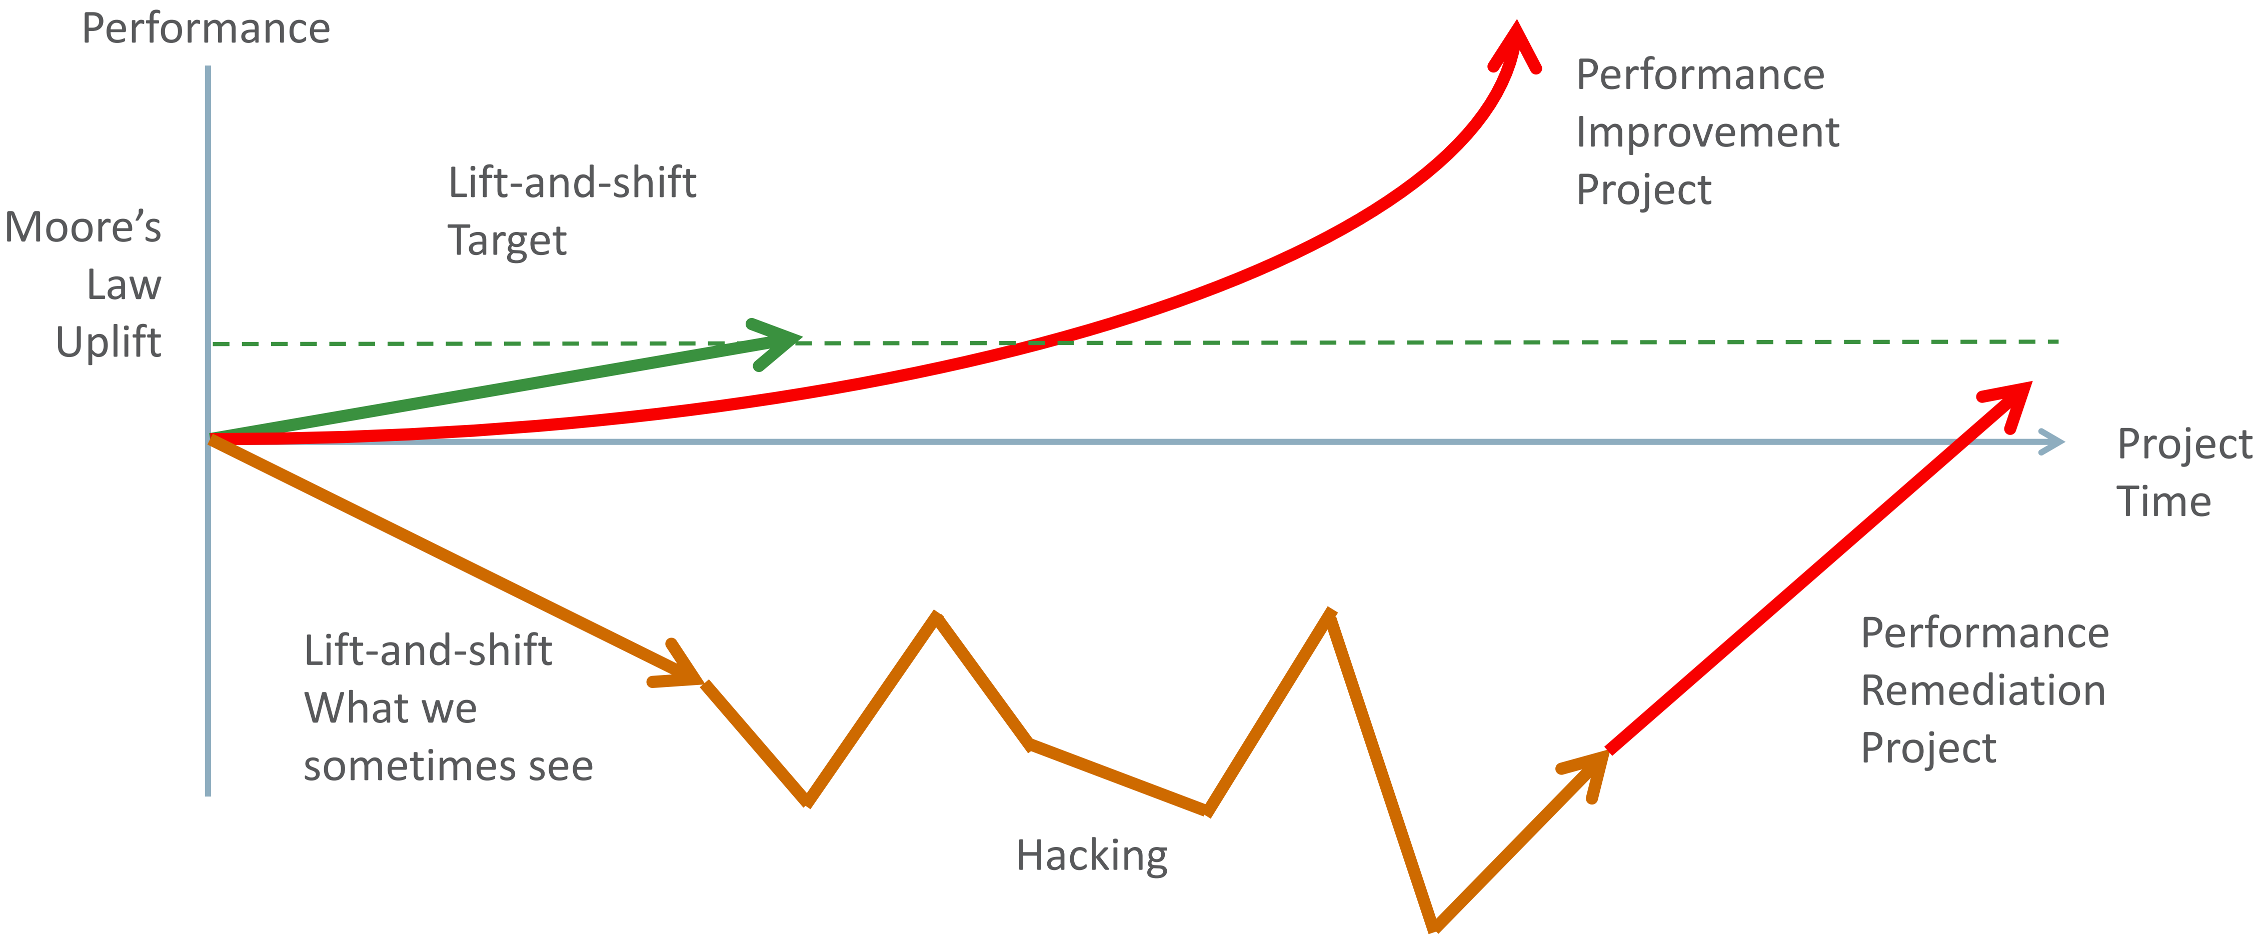

Program Agenda

1. Cloud Platform
2. Optimizer
3. Network
4. Virtualization

Hardware Components

DB System in Oracle Cloud Infrastructure

CPU

- 2 sockets
- 36 cores
- Allocated in increments of 2
- 2 – 36 cores
- 4 – 72 threads

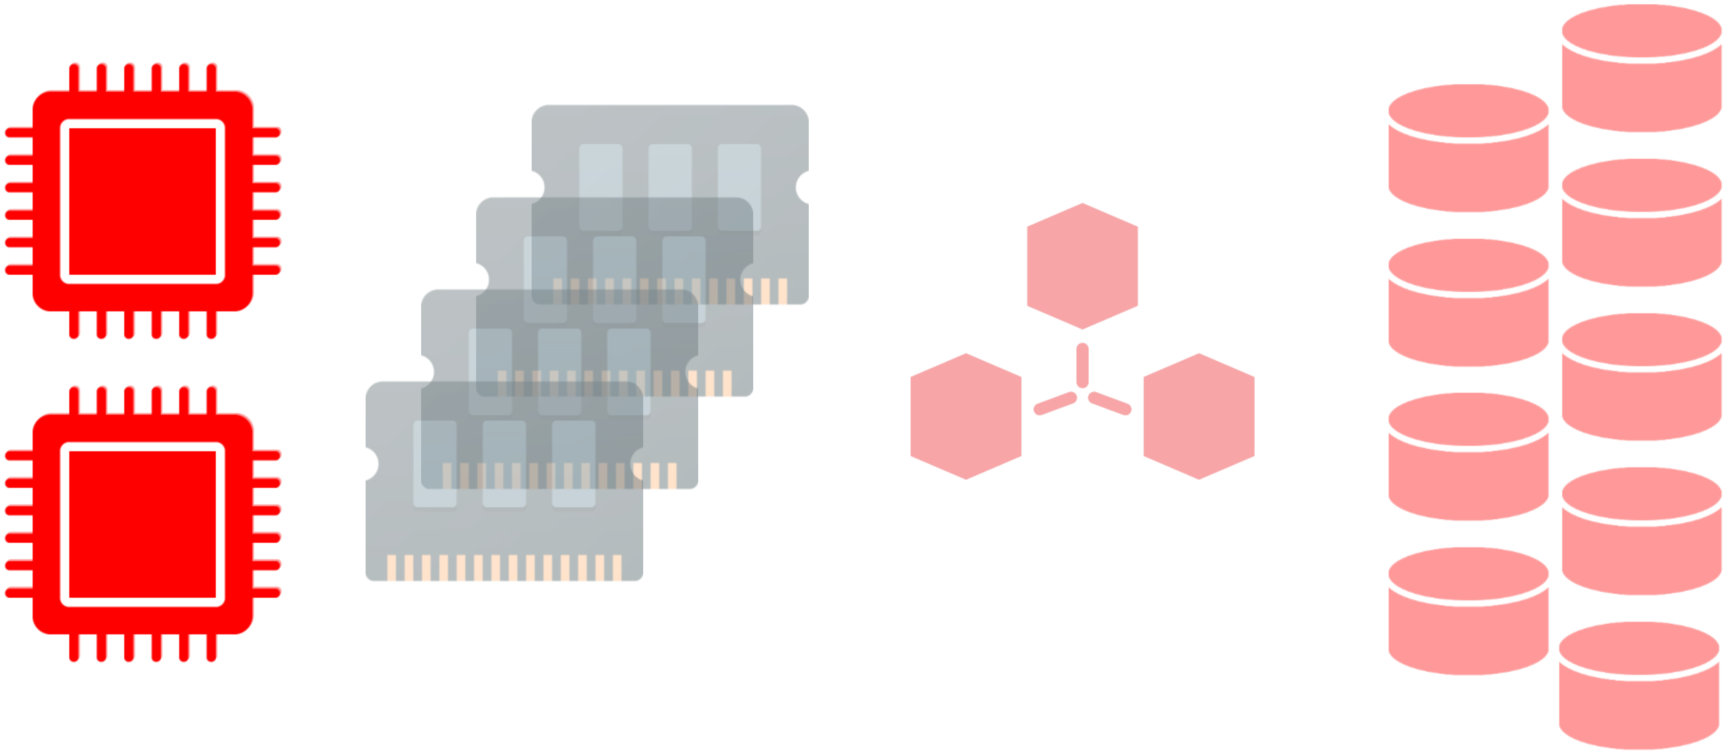

REAL-WORLD PERFORMANCE

Hardware Components

DB System in Oracle Cloud Infrastructure

Memory

- 512 GByte RAM

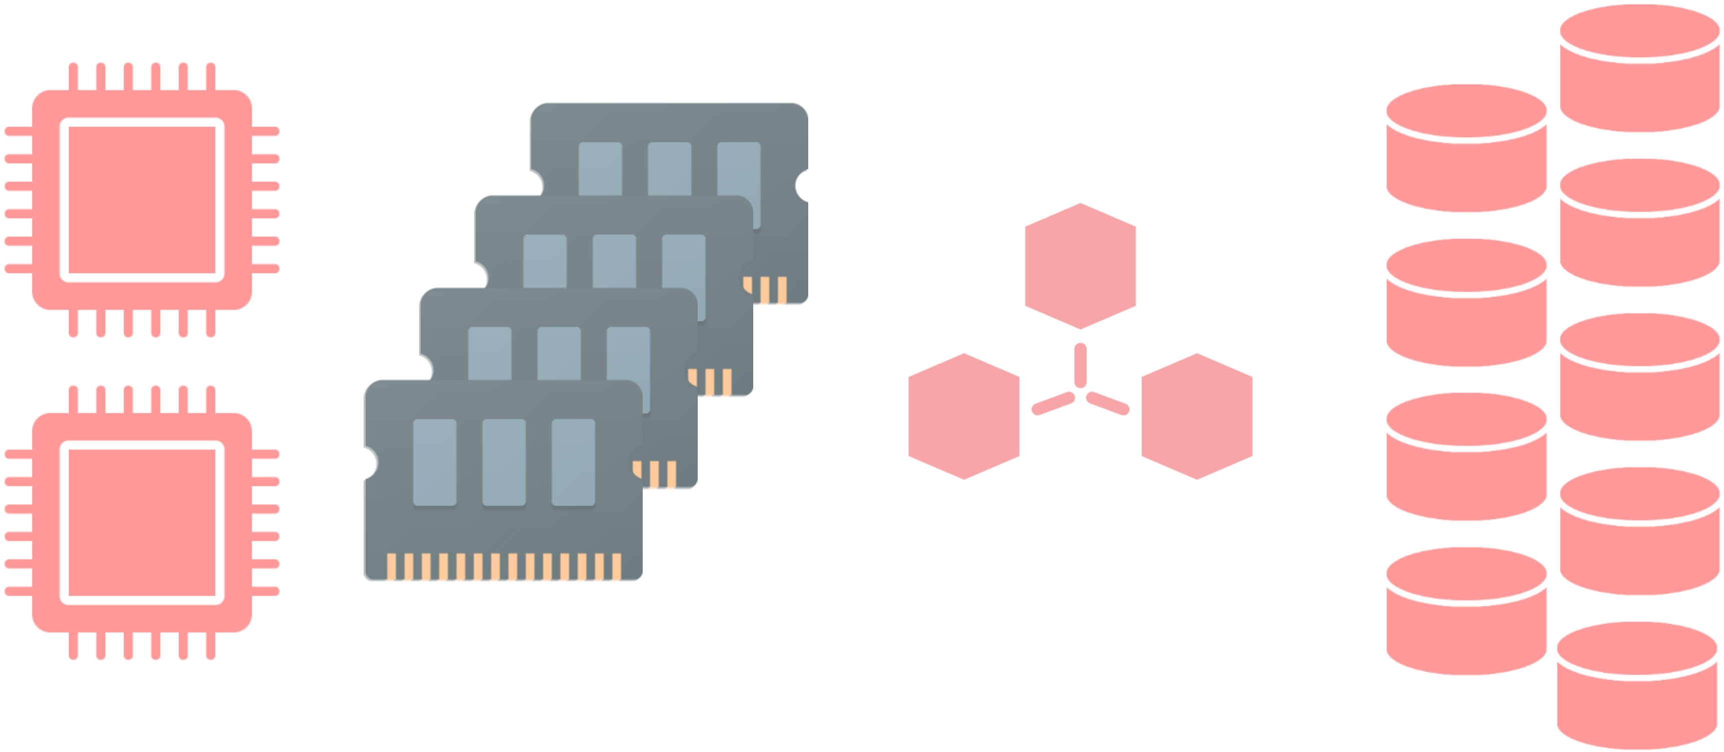

Hardware Components

DB System in Oracle Cloud Infrastructure

Network

- 10 GbE network

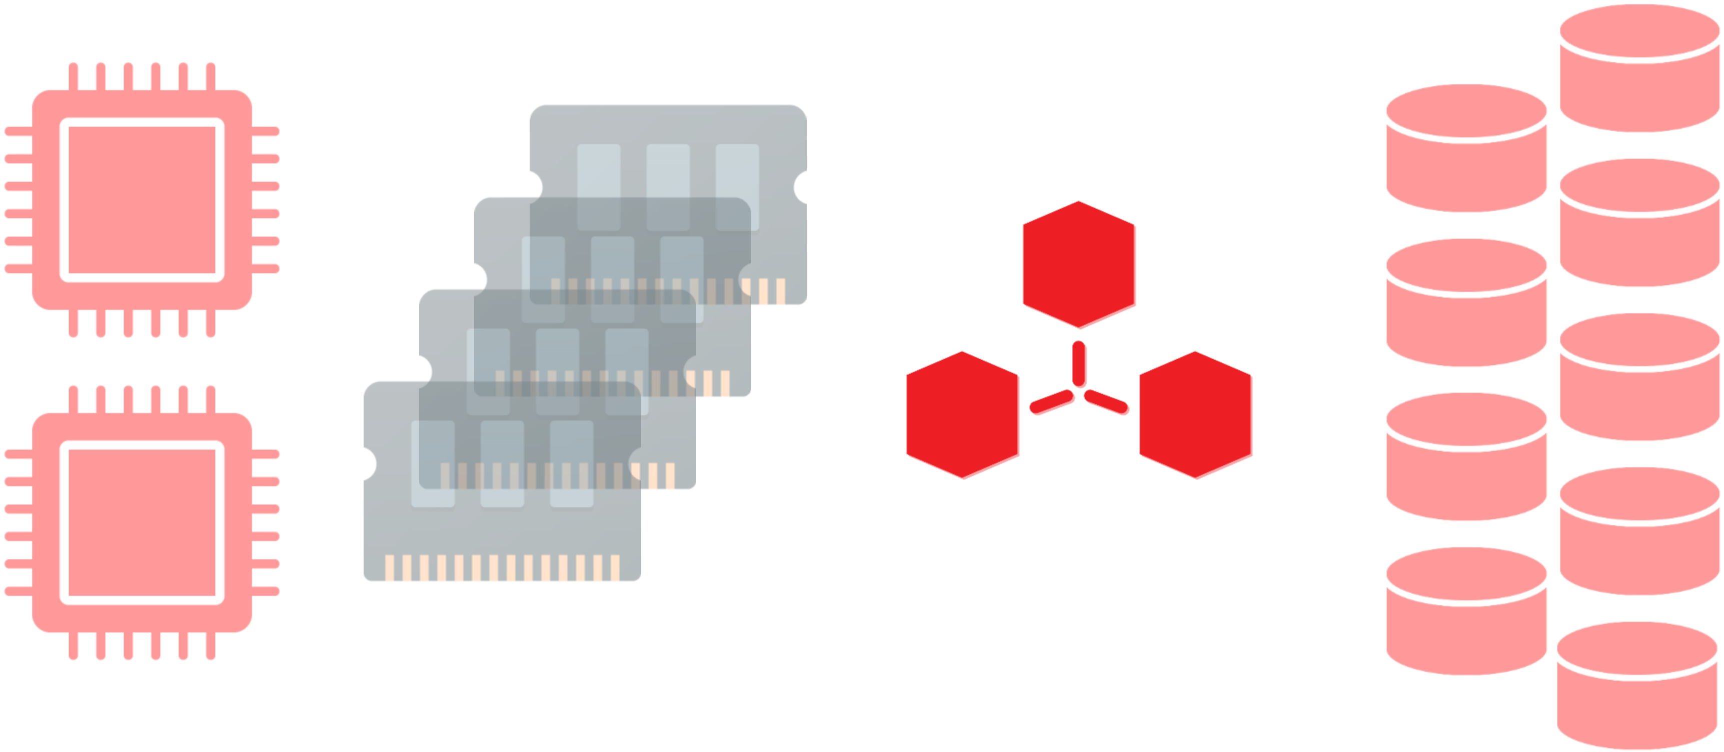

Hardware Components

DB System in Oracle Cloud Infrastructure

Storage

BM.HighIO1.36
- 4 x 3.2 TByte NVMe solid state disks

BM.DenseIO1.36
- 9 x 3.2 TByte NVMe solid state disks

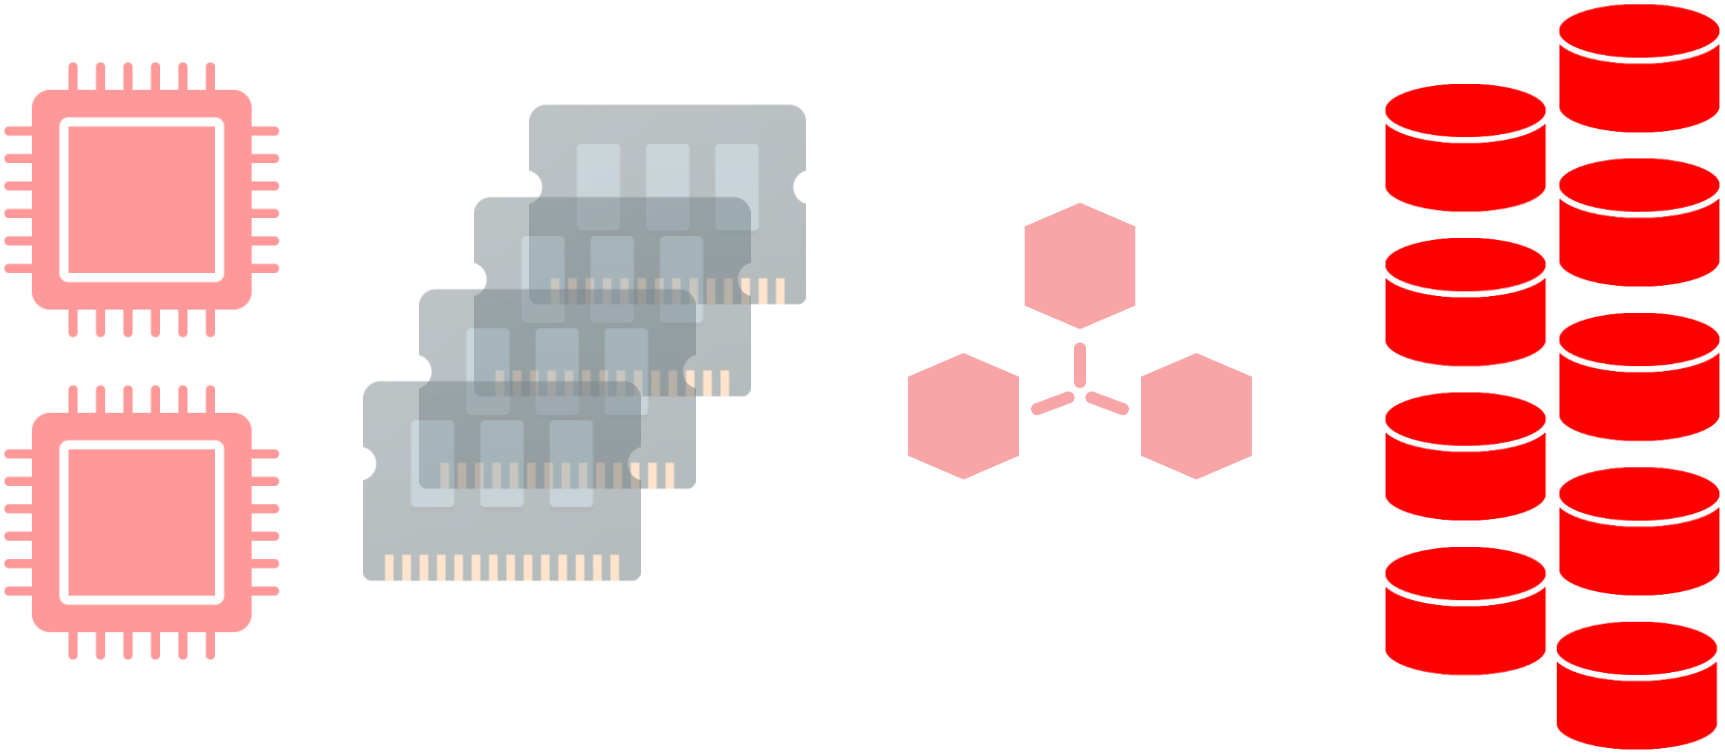

Hardware Components

DB System in Oracle Cloud Infrastructure

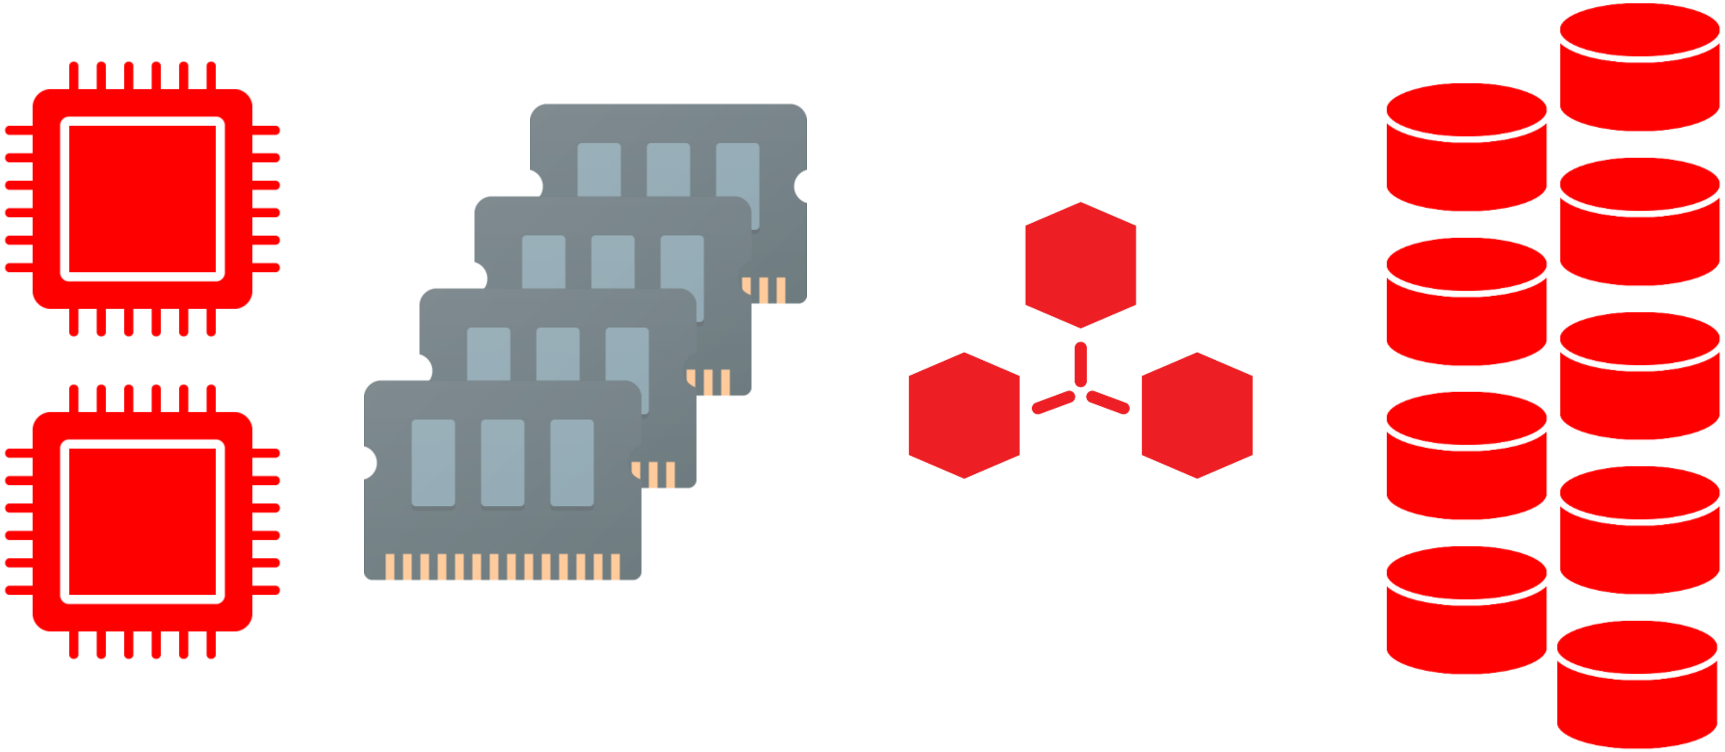

Storage Performance

	9 x Magnetic Disks [Estimated]	9 x NVMe SSDs Reported in AWR
Single block read rate	~ 1 thousand / s	Many hundreds of thousands / s
Corresponding single block read time	~ 5 ms	< 0.15 ms
Multiblock read rate	~ 1 GByte / s	> 12 GByte / s
Multiblock write rate	~ 1 GByte / s	> 5 GByte / s

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Lift and Shift

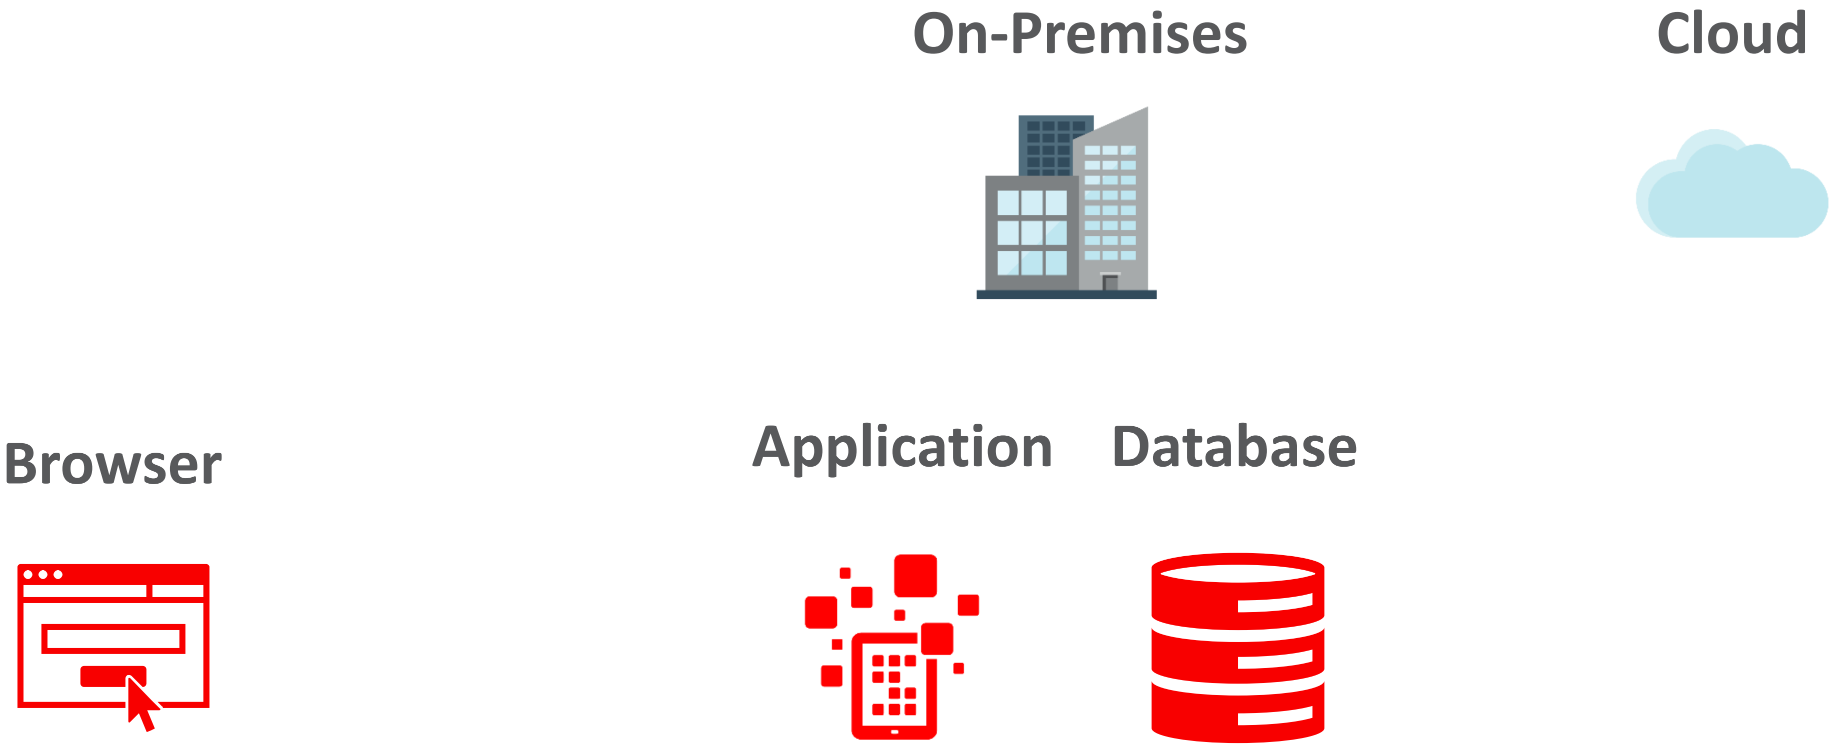

On-Premises to Cloud

Is Your Database Ready for Cloud?

	On-Premises	Cloud
Hardware	Custom	Standard
Operating System	Custom	Standard
Application Server	Custom	Standard
Database	Custom	Standard

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

On-Premises to Cloud

Is Your Database Ready for Cloud?

- Does your database workload require
  1. Custom patches or configuration of the operating system?
  2. Custom patches for Oracle Database?
  3. Custom configuration parameters for Oracle Database?
     1. Basic?
     2. Advanced?
     3. Undocumented?

ORACLE

ORACLE REAL-WORLD PERFORMANCE

On-Premises to Cloud

Is Your Database Ready for Cloud?

- The configurations chosen for Cloud are designed to work well for the majority of workloads
  - There is no magic involved
  - It is impossible to provide an optimal configuration for every workload
- Typical configurations are relatively simple
  - Few deviations from the defaults
  - Most non-default values are chosen with reliability in mind
- This is **boring** by design

ORACLE

ORACLE REAL-WORLD PERFORMANCE

On-Premises to Cloud

Is Your Database Ready for Cloud?

- Configurations are flexible but do not cater well for **exciting** databases
- Common signs that your database is too exciting
  - Very high values of `processes` or `sessions` parameters
  - Very high values of `open_cursors` parameter
  - Undocumented parameters

ORACLE

ORACLE REAL-WORLD PERFORMANCE

On-Premises to Cloud

Is Your Database Ready for Cloud?

- Cloud does not eliminate the usual project challenges
  - Not enough testing
  - Too many changes
- Review and remove any unnecessary undocumented parameters
- Document the changes you make
  - To the standard Cloud configuration
  - From your customized configuration
- Simple fix up scripts often help to bridge the gaps

Program Agenda

1. Cloud Platform
2. Optimizer
3. Network
4. Virtualization

On-Premises to Cloud

History Lesson

- Exadata was first introduced in 2009
- At the time, most customers were using Oracle Database 10g Release 2 (or earlier)
- Exadata only supported Oracle Database 11g
- Many migrations to Exadata involved two major changes at the same time
  - Exadata
  - Database
- Changing the database release sometimes proved more challenging than changing the platform

ORACLE

ORACLE REAL-WORLD PERFORMANCE

On-Premises to Cloud

Query Execution Plans

- Let's assume that the performance of your existing application is *good enough*
- Changes in query execution plans can have more impact on performance than any benefit from changes in hardware
  - Both positive and negative
  - You are more likely to notice the negative
- For lift-and-shift, the goal is a stable system with performance that is *good enough*

On-Premises to Cloud

Is Your Database Ready for Cloud?

- Many databases running on-premises still use Oracle Database 11g Release 2 (or earlier)
- Most Cloud configurations use Oracle Database 12c
- Database upgrade is a common trigger for new problem execution plans
- How can you reduce the risk of problem execution plans?

ORACLE
ORACLE REAL-WORLD PERFORMANCE

Adaptive Optimizer in Oracle Database 12c Release 1

Goals

- Adaptive Query Optimization was introduced in Oracle Database 12c Release 1
- The goal is to increase the likelihood of good plans first time
- Adaptive Query Optimization makes use of execution statistics
  1. During first execution
  2. As feedback for subsequent executions
  3. As feedback to trigger dynamic sampling and to guide the gathering of statistics
- The biggest change since the introduction of the cost-based optimizer?

ORACLE ORACLE REAL-WORLD PERFORMANCE

Adaptive Optimizer in Oracle Database 12c Release 1

Reality

- Multiple feedback loops are involved
  - It takes time for statistics to evolve as the optimizer learns from the workload
  - Diagnosing problems can be challenging if not impossible
- When problems occur for specific SQL statements
  - Query execution times can be longer or more variable
  - The number and duration of hard parses may increase
  - Symptoms can include latch and mutex wait events
- In some cases the side-effects can impact the stability of the system

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

Oracle Database 12c Release 1

Controlling Adaptive Features

OPTIMIZER_ADAPTIVE_FEATURES

Optimizer Adaptive Features

Change plans at runtime	Learn from previous executions

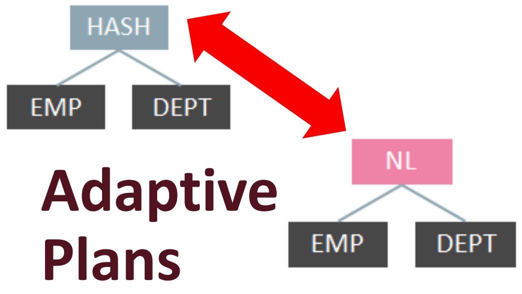

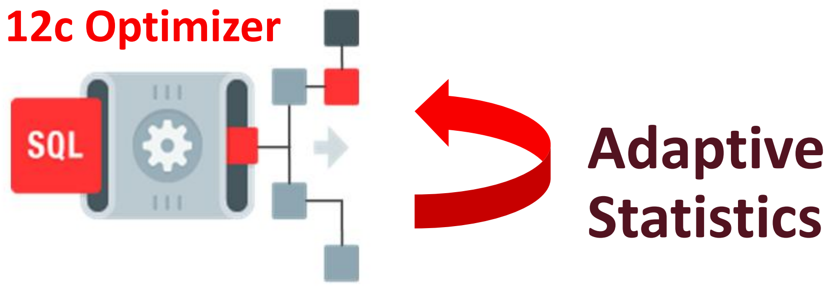

From Oracle Database 12c Release 2

Finer Control Using New Parameters

~~OPTIMIZER_ADAPTIVE_FEATURES~~ Obsolete

Optimizer Adaptive Features

OPTIMIZER_ADAPTIVE_PLANS

Change plans at runtime

OPTIMIZER_ADAPTIVE_STATISTICS

Learn from previous executions

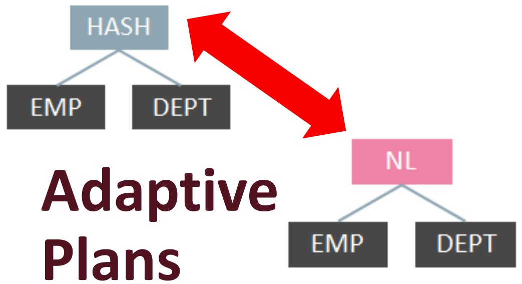

Adaptive Plans

12c Optimizer

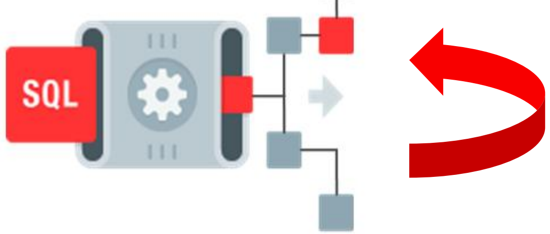

Adaptive Statistics

From Oracle Database 12c Release 2

New Defaults

Optimizer Adaptive Features

OPTIMIZER_ADAPTIVE_PLANS (TRUE)

OPTIMIZER_ADAPTIVE_STATISTICS (FALSE)

Change plans at runtime

Learn from previous executions

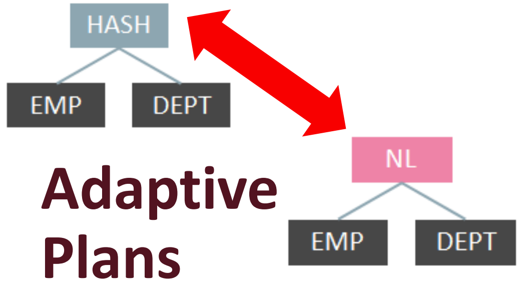

Adaptive Plans

Adaptive Statistics

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Adaptive Optimizer in Oracle Database 12c Release 1

Real-World Performance Recommendation

- Adopt the default behaviour of Oracle Database 12c Release 2
  - Adaptive Plans enabled
  - Adaptive Statistics disabled
- How?
  - Review Doc ID 2187449.1 for recommended patches, currently
    - 22652097
    - 21171382
  - Check Oracle Optimizer blog post by Nigel Bayliss for more details

Controlling the Variables

Optimizer Parameters

- Parameters influencing plan selection include
  - `optimizer_*` parameters
  - fix controls
  - undocumented optimizer parameters
- General position
  - Real-World Performance does not recommend changing optimizer parameters, except where directed by an application vendor
- In practice
  - If your database currently sets optimizer parameters, do not introduce an unwanted variable by using the defaults

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Controlling the Variables

System Statistics

- General position
  - Real-World Performance does not recommend gathering system statistics
- Specific position
  - Real-World Performance does not recommend gathering system statistics in Cloud
- In practice
  - If your database does not have system statistics, do not introduce a variable by gathering system statistics
  - If your database does have system statistics, use the existing system statistics, and do not introduce another variable by re-gathering or removing system statistics

Controlling the Variables

Moving your Data to Cloud

- An Oracle database upgrade will preserve existing schema metadata (e.g. statistics) and SQL metadata (e.g. SQL Profile)
- Some methods of moving data to Cloud can lose information that is important for the optimizer
  - Data Pump with EXCLUDE=STATISTICS
  - Classic Export and Import
  - Flat files
- Simply gathering fresh statistics may introduce many unwanted variables

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Moving Your Database to Cloud

Table statistics
Column statistics
Histograms
Extended statistics
Column usage information

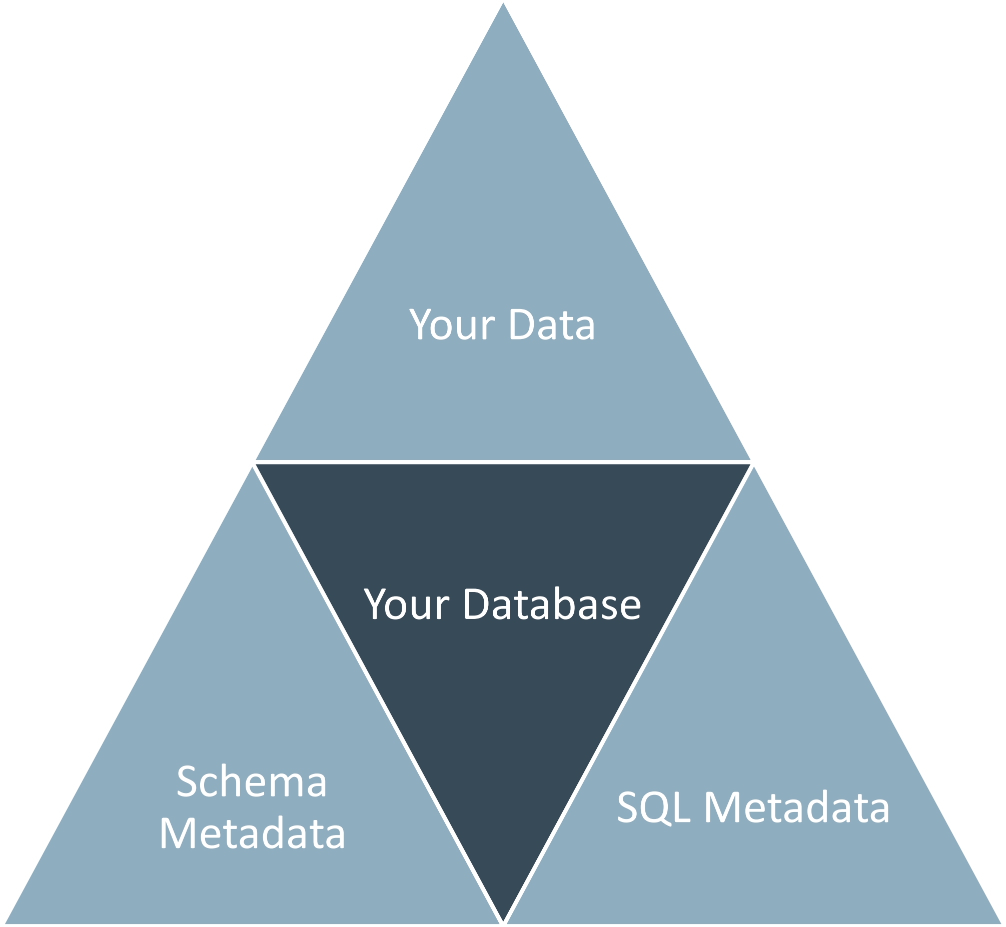

SQL Plan Management
SQL Profile
SQL Patch
Stored outline

Program Agenda

1. Cloud Platform
2. Optimizer
3. Network
4. Virtualization

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Reality Check

- Some components will not move to Cloud
- Few enterprise applications exist in isolation
  - Almost all enterprise applications have multiple integration points
  - Not all components will move to Cloud simultaneously, or even any time soon
- A longer network hop will be introduced

One Data Centre

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

Two Data Centres

Two Cloud Providers

Database in Cloud

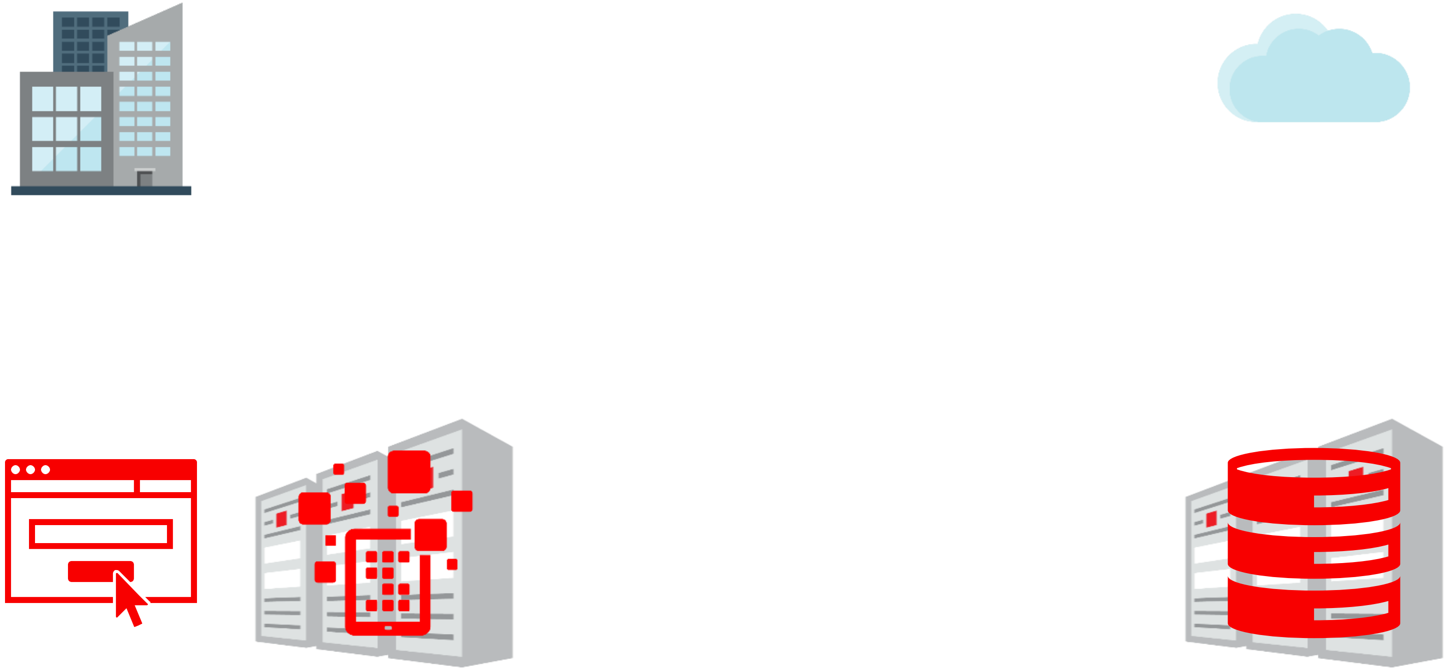

Round Trip Time with Distance

100 ms
- International
- 6000 miles / 10000 Kilometres

10 ms
- Regional
- 600 miles / 1000 Kilometres

1 ms
- Metropolitan
- 60 miles / 100 Kilometres

< 1ms
- Local
- Propagation delay is small relative to other costs in the network stack

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Estimating the Impact

Measurement	How?	Value
Increase in round trip time	ping	t
Number of Oracle Net round trips	**SQL*Net roundtrips to/from client** per second in AWR	n
Estimate the **Additional Time** from the application perspective	t * n	
Database Time	**DB Time** per second in AWR	

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Estimating the Impact

Low Risk High

Additional Time << DB Time

Additional Time >> DB Time

Estimating the Impact

- Measure the increase in round trip time (e.g. using ping)
- Find the number of Oracle Net round trips from AWR
- Estimate the increase in time from the application perspective
- Compare with existing Database Time from AWR

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Estimating the Impact

Worked Example

- A ping between an application server and an adjacent database server in the same data centre in California reports 0.1 milliseconds
- A ping between the application server in California and a database server in a Cloud data centre in Texas reports 42 milliseconds.
- This is faster than the blink of an eye. Is it important?

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Estimating the Impact

Response Time

- When the servers are in the same data centre in California, an AWR report indicates 2366 roundtrips per second
- SQL*Net roundtrips to/from client is a **counter**, not a wait event
- If the database server is in Texas, the average response time for each transaction will increase
  - 8 roundtrips per transaction x 0.042 s = 0.32 s
  - Would you notice?

Statistic	Total	per Second	per Trans
SQL*Net roundtrips to/from client	80,009	2,365.87	8.00

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Estimating the Impact

Throughput

- When the servers are in the same data centre in California, an AWR report indicates DB Time is 0.7 s per second
- How long would it take to perform 2366 roundtrips?
  - 2366 x 0.042 ~ 100 s
- Would you notice if DB Time increased from less than 1 s to 100 s
  - This is the effect from the perspective of the application
  - Not what actually happens in the database

	Per Second
DB Time(s):	0.7

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Estimating the Impact

Problem Scenarios

- What happens if the application was designed to use a single worker process?
- What happens if the application was designed to use multiple worker processes but has a serialization point?
  - Perform some SQL
  - `UPDATE CTRL SET N = N + 1`
  - Perform some more SQL
  - `COMMIT`
- You may be able to buy bandwidth. The speed of light is non-negotiable.

The Impact of Increased Round Trip Time

- From the application perspective, the database appears slower
- From the database perspective, database sessions wait longer for work
- Locks may be held for longer
  - Reduces application scalability
- More database sessions will usually be needed to support the same throughput
  - Impacts efficiency and stability

Program Agenda

1. Cloud Platform
2. Optimizer
3. Network
4. Virtualization

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

Basic Concepts

- Virtualization is commonplace in Cloud but the two are not the same
  - You might already be using virtualization on-premises
  - Most Cloud services use virtualization, while some do not
- Can you trust your usual tools?

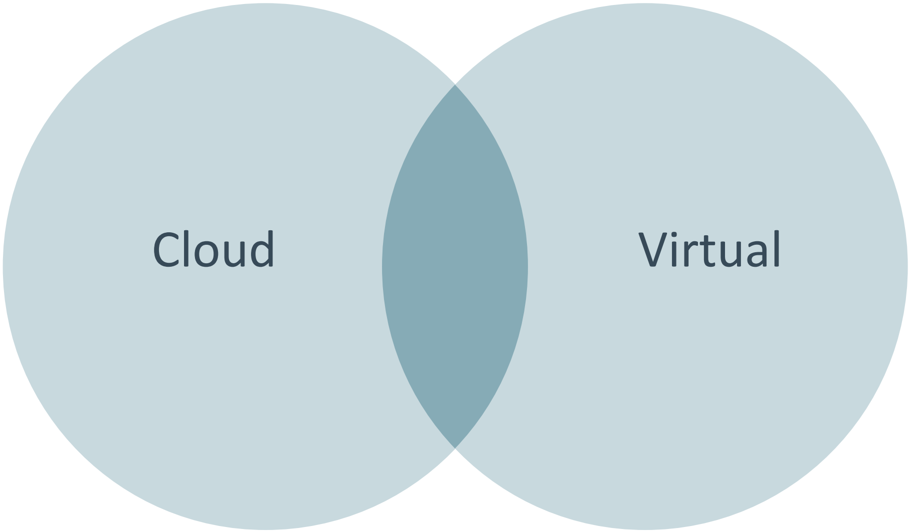

ORACLE
ORACLE REAL-WORLD PERFORMANCE

Virtual CPUs

- A single x86 64-bit core with 4 threads? How can that be?
  - Oracle relies on Linux
  - Linux relies on the virtualization layer
  - It is a tough problem to solve
- In many cases, only the number of virtual CPUs is reliable
- Rules of thumb need to be adjusted accordingly

Host Name	Platform	CPUs	Cores	Sockets
myvm	Linux x86 64-bit	4	1	1

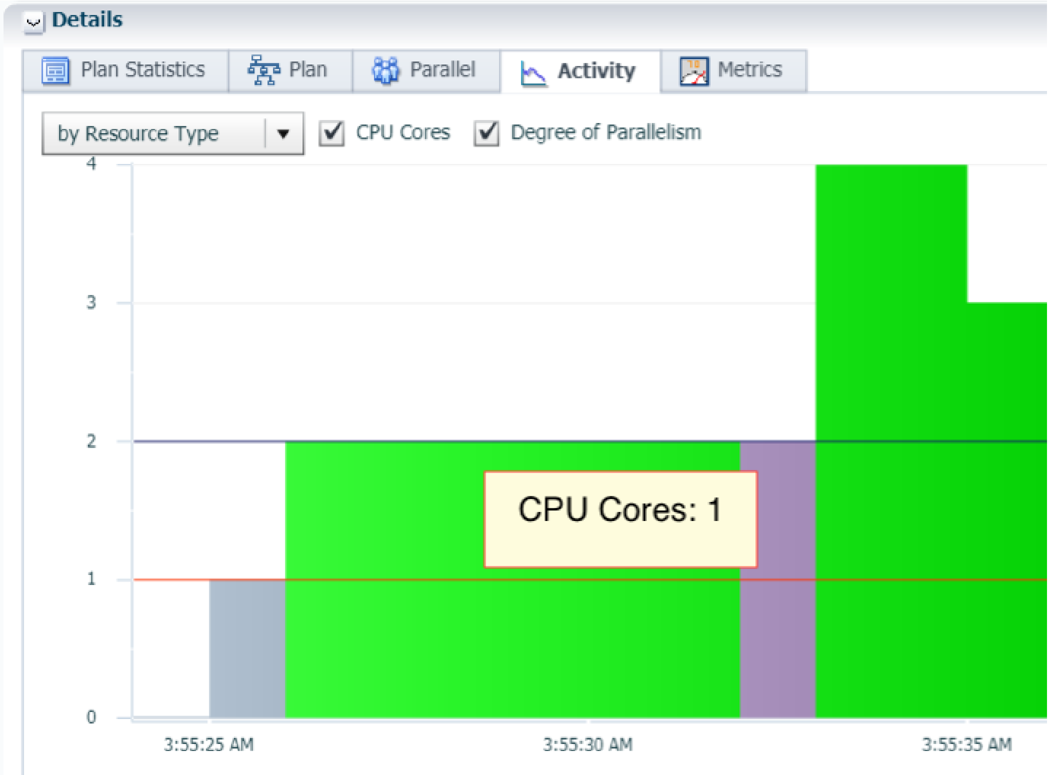

Virtual CPUs

- Question
  - With four virtual CPUs, how much CPU power is available to the database?
- Answer
  - It depends
- The answer is important for performance and consistency
  - Sometimes you might receive **more** than you pay for
  - At other times you might get only what you pay for
- Oversubscription will magnify the variation

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Performance

- Virtualization is well established, so let's assume we all understand the arguments in favour of virtualization
- What is the impact on performance?
  - The difference in performance relative to bare metal continues to improve
  - There is still a difference
- The overhead is usually incurred when Oracle Database interacts with the operating system
  - Some system calls take longer and consume more CPU time
  - The additional cost is often tiny for an individual call
  - The overall impact is magnified by frequent calls.

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

Virtualization Considerations

Instrumentation

- The instrumentation to provide DB Time is enabled by default
- Oracle Database obtains its information from the operating system
  - For wall-clock time `clock_gettime()` and `gettimeofday()`
  - For resource consumption `getrusage()` and `times()`
- High value in return for a small overhead

ORACLE
ORACLE REAL-WORLD PERFORMANCE

Virtualization Considerations

Instrumentation

- In most operating systems, a user-mode library call is used to obtain the wall-clock time without the overhead of a system call
- The wall-clock time is derived in a few instructions based on a hardware clock maintained by the processor
- Some virtualized configurations perform an expensive system call instead
  - The overall impact is usually small for real-world workloads
  - It can be important when using performance diagnostic tools such as SQL trace or `GATHER_PLAN_STATISTICS` hint

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Instrumentation

Worked Example

- Snippet of strace output for a nested loop join with SQL trace not enabled
- System calls are used to obtain wall-clock time
- However, the frequency is small, so the impact is small

```
    calls  syscall
---------  ----------------
       53  clock_gettime
       24  getrusage
       14  gettimeofday
```

ORACLE

ORACLE REAL-WORLD PERFORMANCE

Instrumentation

Worked Example

- Snippet of strace output for the same nested loop join with SQL trace enabled
- The frequency is now very much higher for `clock_gettime()`, so the impact is much greater

```
    calls  syscall
---------  ----------------
  6031904  clock_gettime
      564  getrusage
       17  gettimeofday
```

Cost of Instrumentation

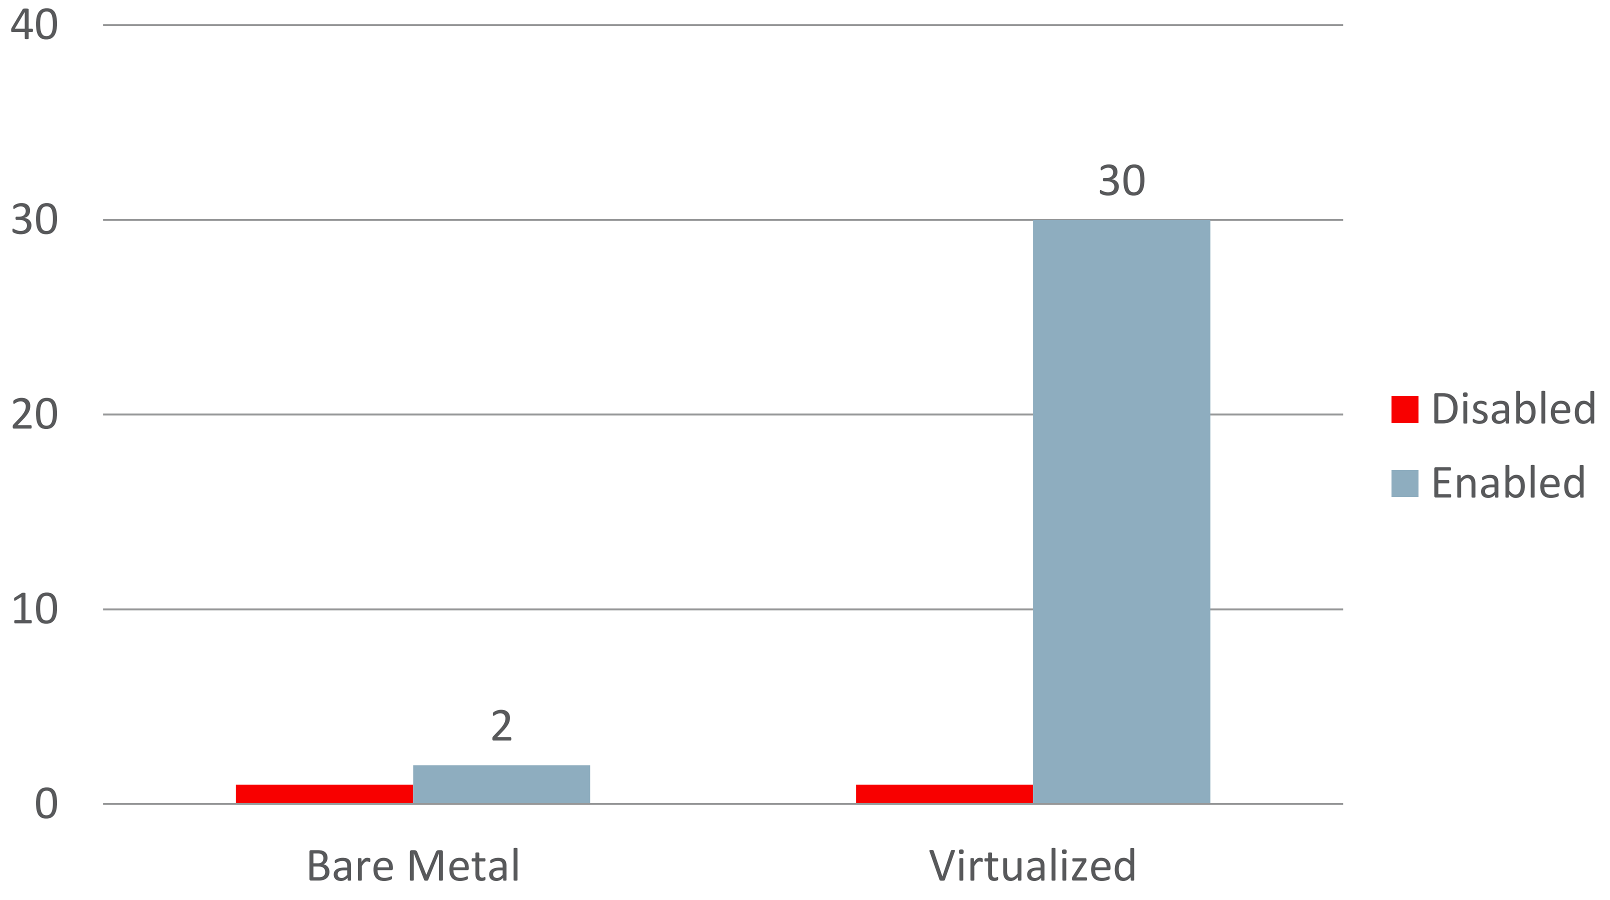

Relative cost of SQL Trace on bare metal and virtualized platforms

Same data

Same SQL statement

Same execution plan

The extra overhead might cause you to make the wrong decision when optimizing a specific statement

Using SQL Monitor avoids this behaviour

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

Summary

If performance problems do occur, remember that the Cloud Platform may be just one of many changes

Eliminate common sources of poor query execution plans by taking care with schema metadata and SQL metadata

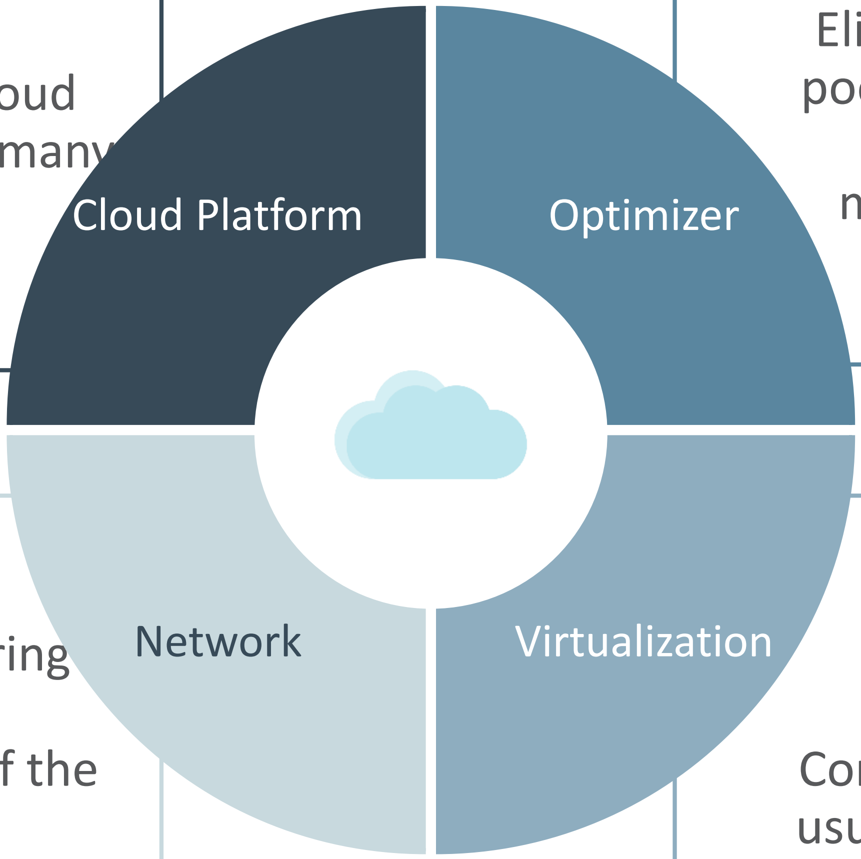

Quantify the risk by considering the number of Oracle Net roundtrips and the latency of the network

Consider the reliability of your usual tools and rules of thumb

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

Real-World Performance

Online Video Series

- Real-World Performance Engineers discussing and demonstrating performance issues, root causes and when to apply the correct techniques
  - The Optimizer
  - Core DB Performance
  - Extreme OLTP
  - Extreme DW
- http://www.oracle.com/goto/oll/rwp

Safe Harbor Statement

The preceding 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.

ORACLE

ORACLE
REAL-WORLD PERFORMANCE

Integrated Cloud

Applications & Platform Services