
ORACLE®

ORACLE
BIG DATA

Big Made Simple

Oracle Big Data SQL

Hands-on Lab

Marty Gubar - Big Data PM

ORACLE

Database Strategy for Big Data

Conventional view of Data Management

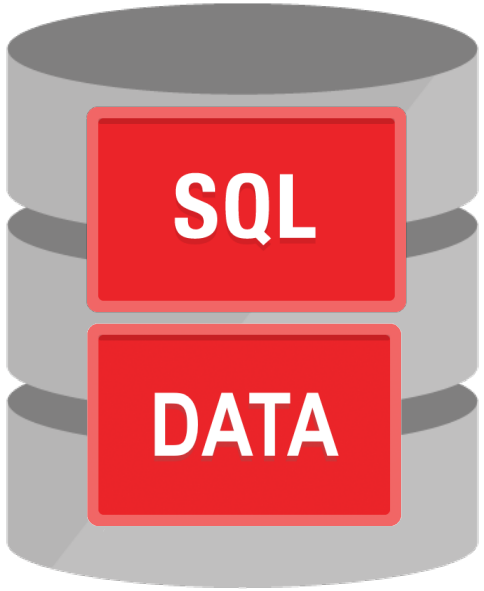

Emerging view of Data Management

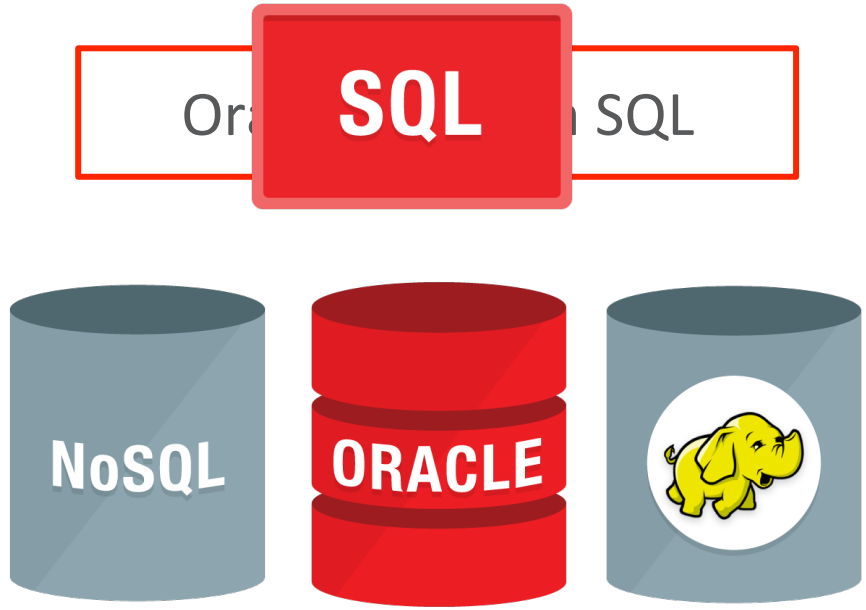

Big Data SQL Configurations

Engineered Systems

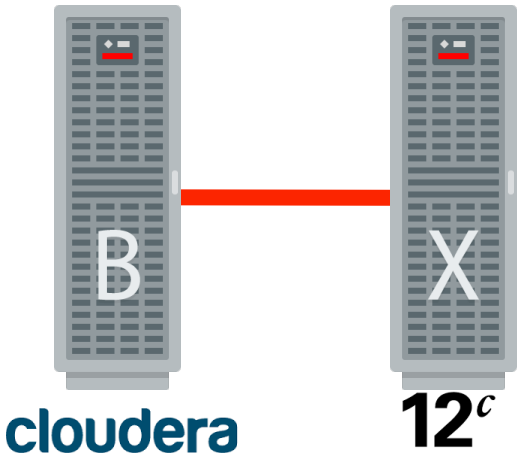

Engineered Systems*

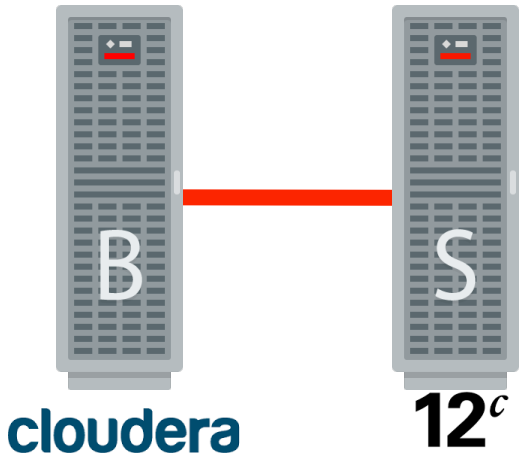

Oracle Cloud*

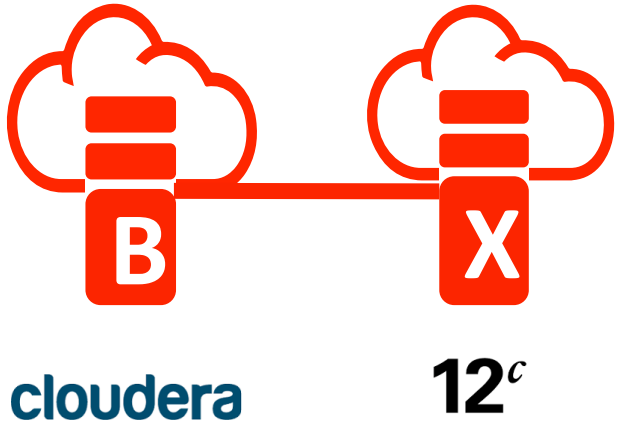

Commodity Servers

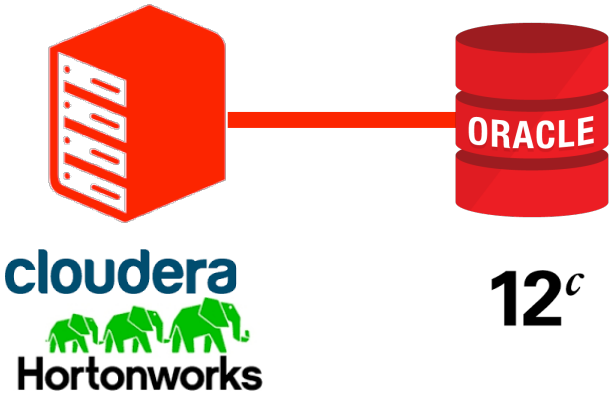

Mixed Deployment

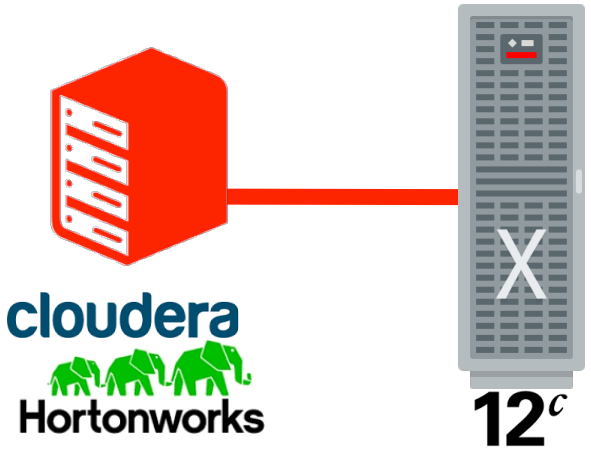

Mixed Deployment

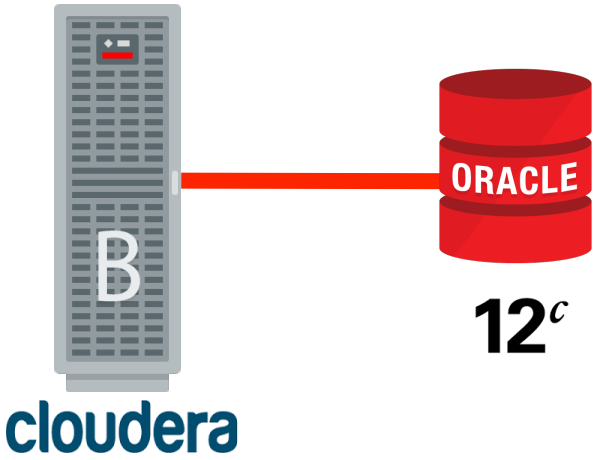

* Coming Soon!

Metadata: Extend Oracle External Tables

```
CREATE TABLE movielog (
  click VARCHAR2(4000))
ORGANIZATION EXTERNAL (
  TYPE ORACLE_HIVE
  DEFAULT DIRECTORY DEFAULT_DIR
  ACCESS PARAMETERS
  (
com.oracle.bigdata.tablename logs
com.oracle.bigdata.cluster mycluster
  ))
REJECT LIMIT UNLIMITED;
```


- New types of external tables
  - ORACLE_HIVE (leverage hive metadata)
  - ORACLE_HDFS (specify metadata)
- Access parameters used to describe how to identify sources and process data on the hadoop cluster

Oracle Big Data SQL

Massively Parallel SQL Query across Oracle, Hadoop and NoSQL

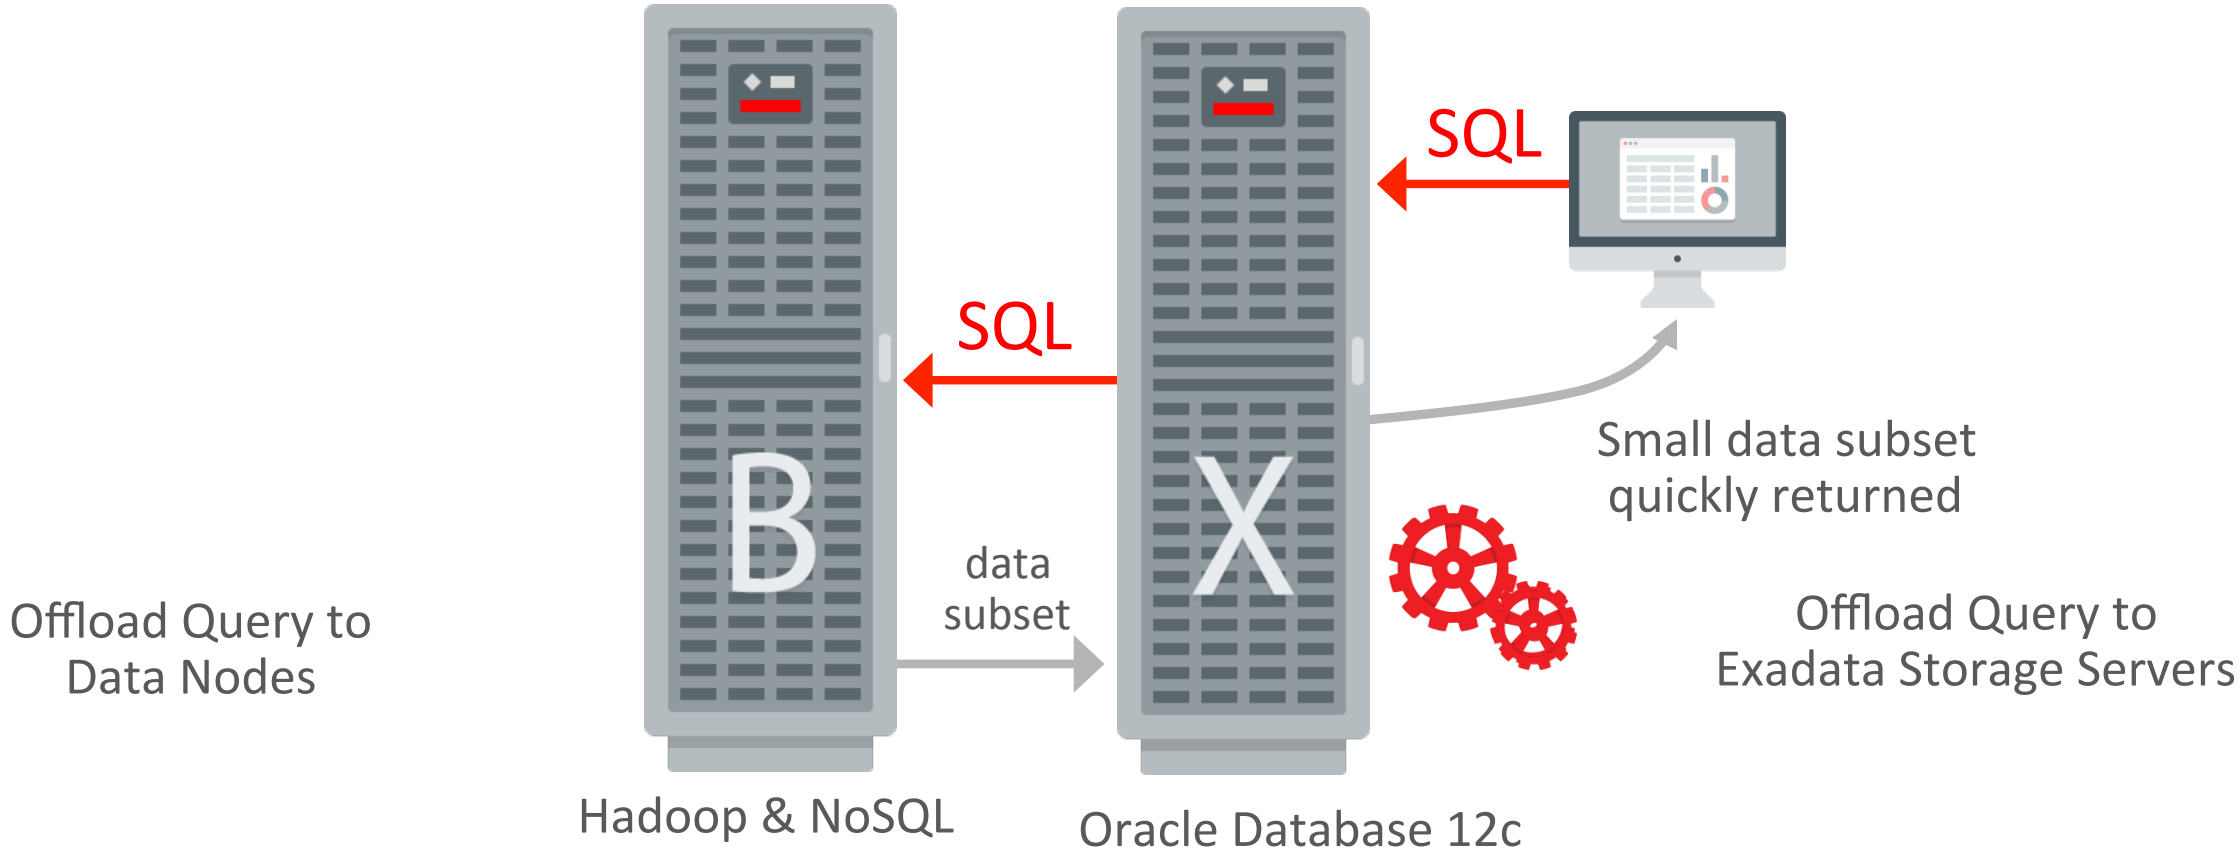

Big Data SQL: Another Hadoop Processing Engine

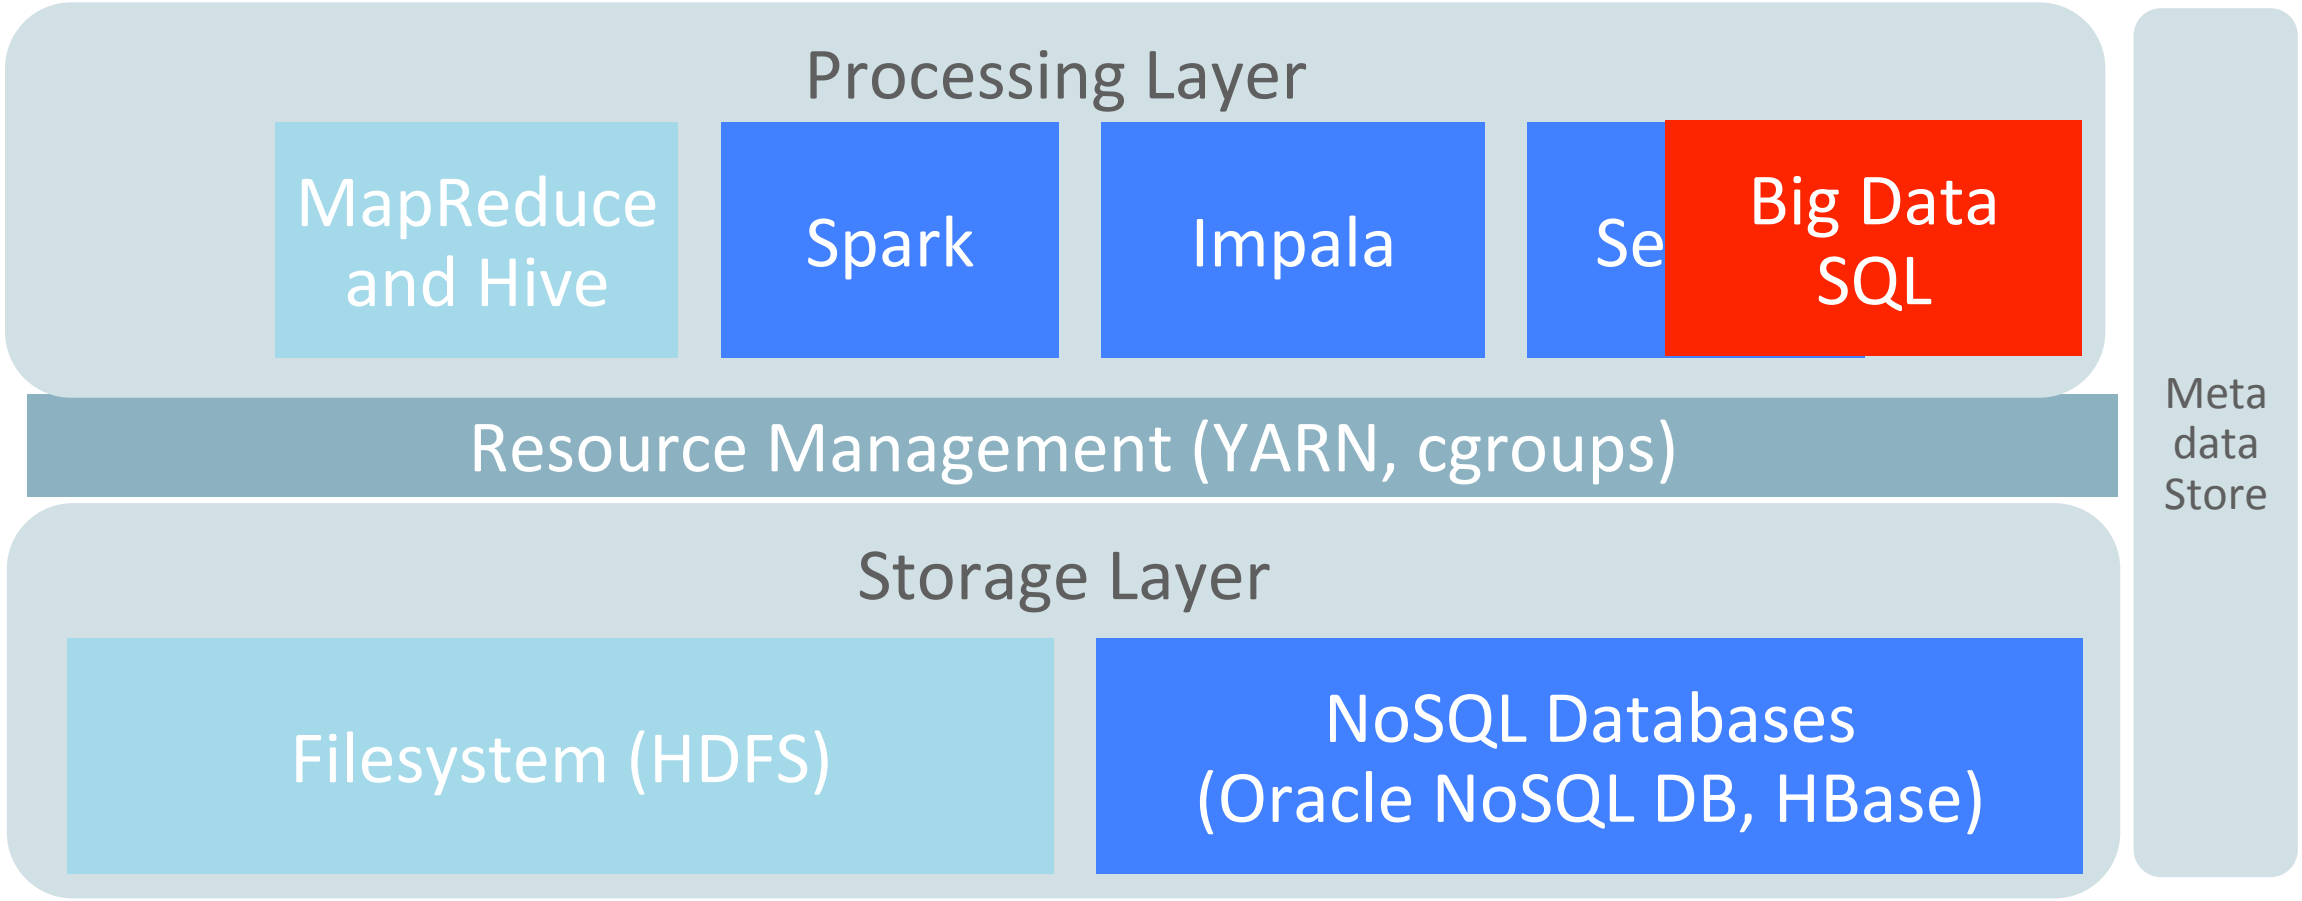

Big Data SQL Performance Features

IO Reduction Features Deliver Compound Results

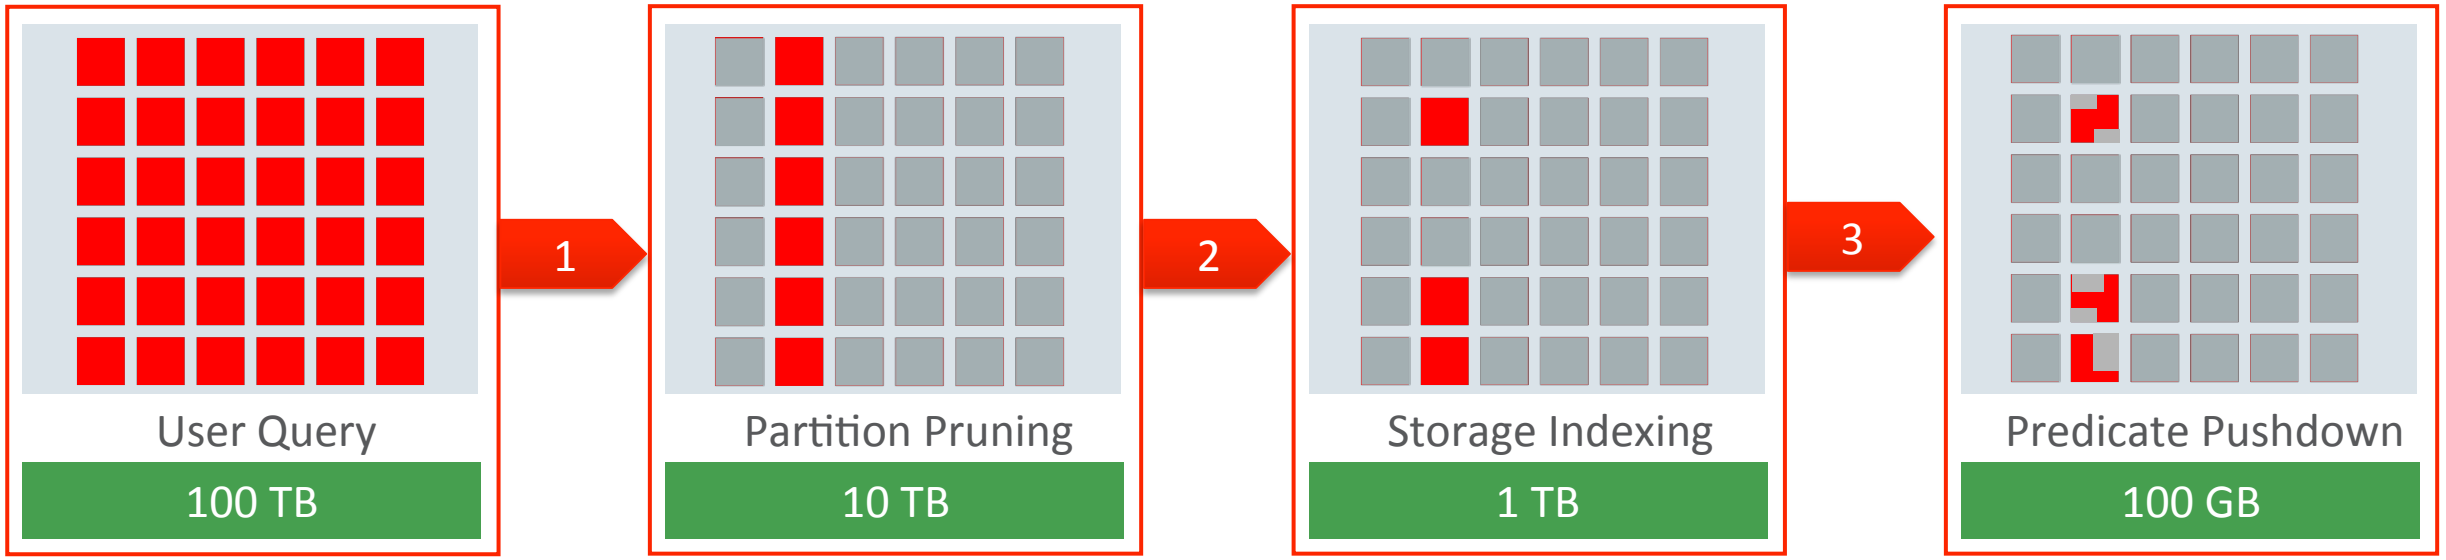

Scenario: On-line Movie Streaming Site

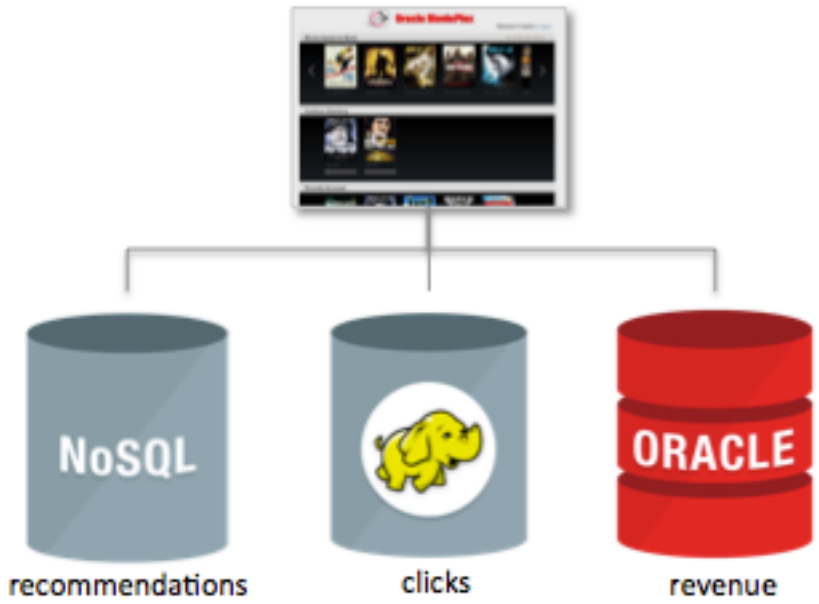

- Data sources
  - Recommendations in NoSQL DB
  - User behavior in HDFS
  - Data Warehouse in Oracle Database
- Gain value from all data!

List of workshops

- Part 1 - Configuring Oracle Big Data SQL
- Part 2 - Create Oracle Table Over Application Log
- Part 3 - Leverage the Hive Metastore to Access Data in Hadoop & Oracle NoSQL Database
- Part 4 - Applying Oracle Database Security Policies Across the Big Data Platform

List of workshops

- Part 5 - Using Oracle Analytic SQL Across All Your Data
- Part 6 - Introduction to SQL Pattern Matching
- Part 7 - Checking the Pattern Matching Process
- Part 8 - Creating a More Useful Data Set
- Part 9 - Other Useful 12c Analytical SQL Features

Ravello accelerates the move from data-center to public cloud

Run data-center based **VMware** or **KVM** workloads on public cloud

without ***any changes***

Same VMs, networking, storage

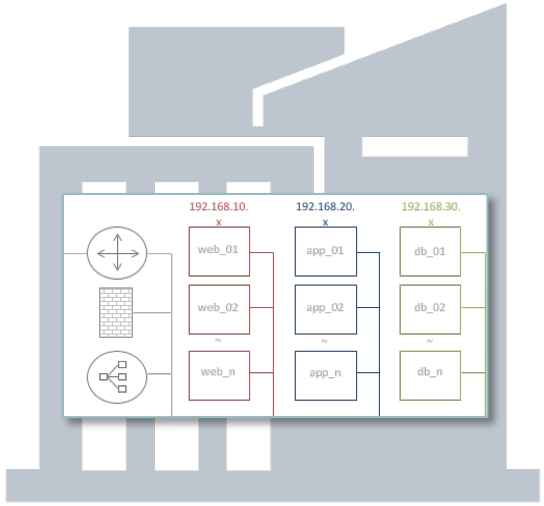

Data Center

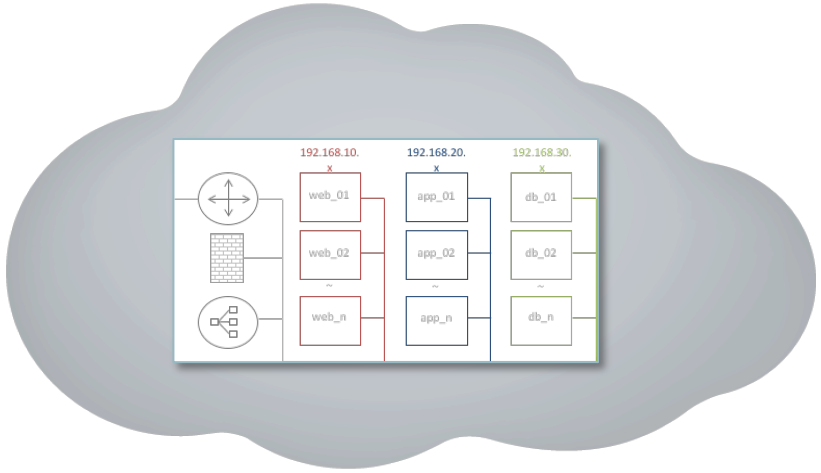

ORACLE | ravello

No app changes → Data-center to cloud *in hours*, and *NOT months*

ORACLE

Connecting

- You will be assigned an IP Address
- Use VNC:
  - vnc <ip-address>:1
  - password: welcome1
- or SQL Developer
  - ip-address:1521:orcl
  - moviedemo / welcome1

Starting the Hands-on Lab

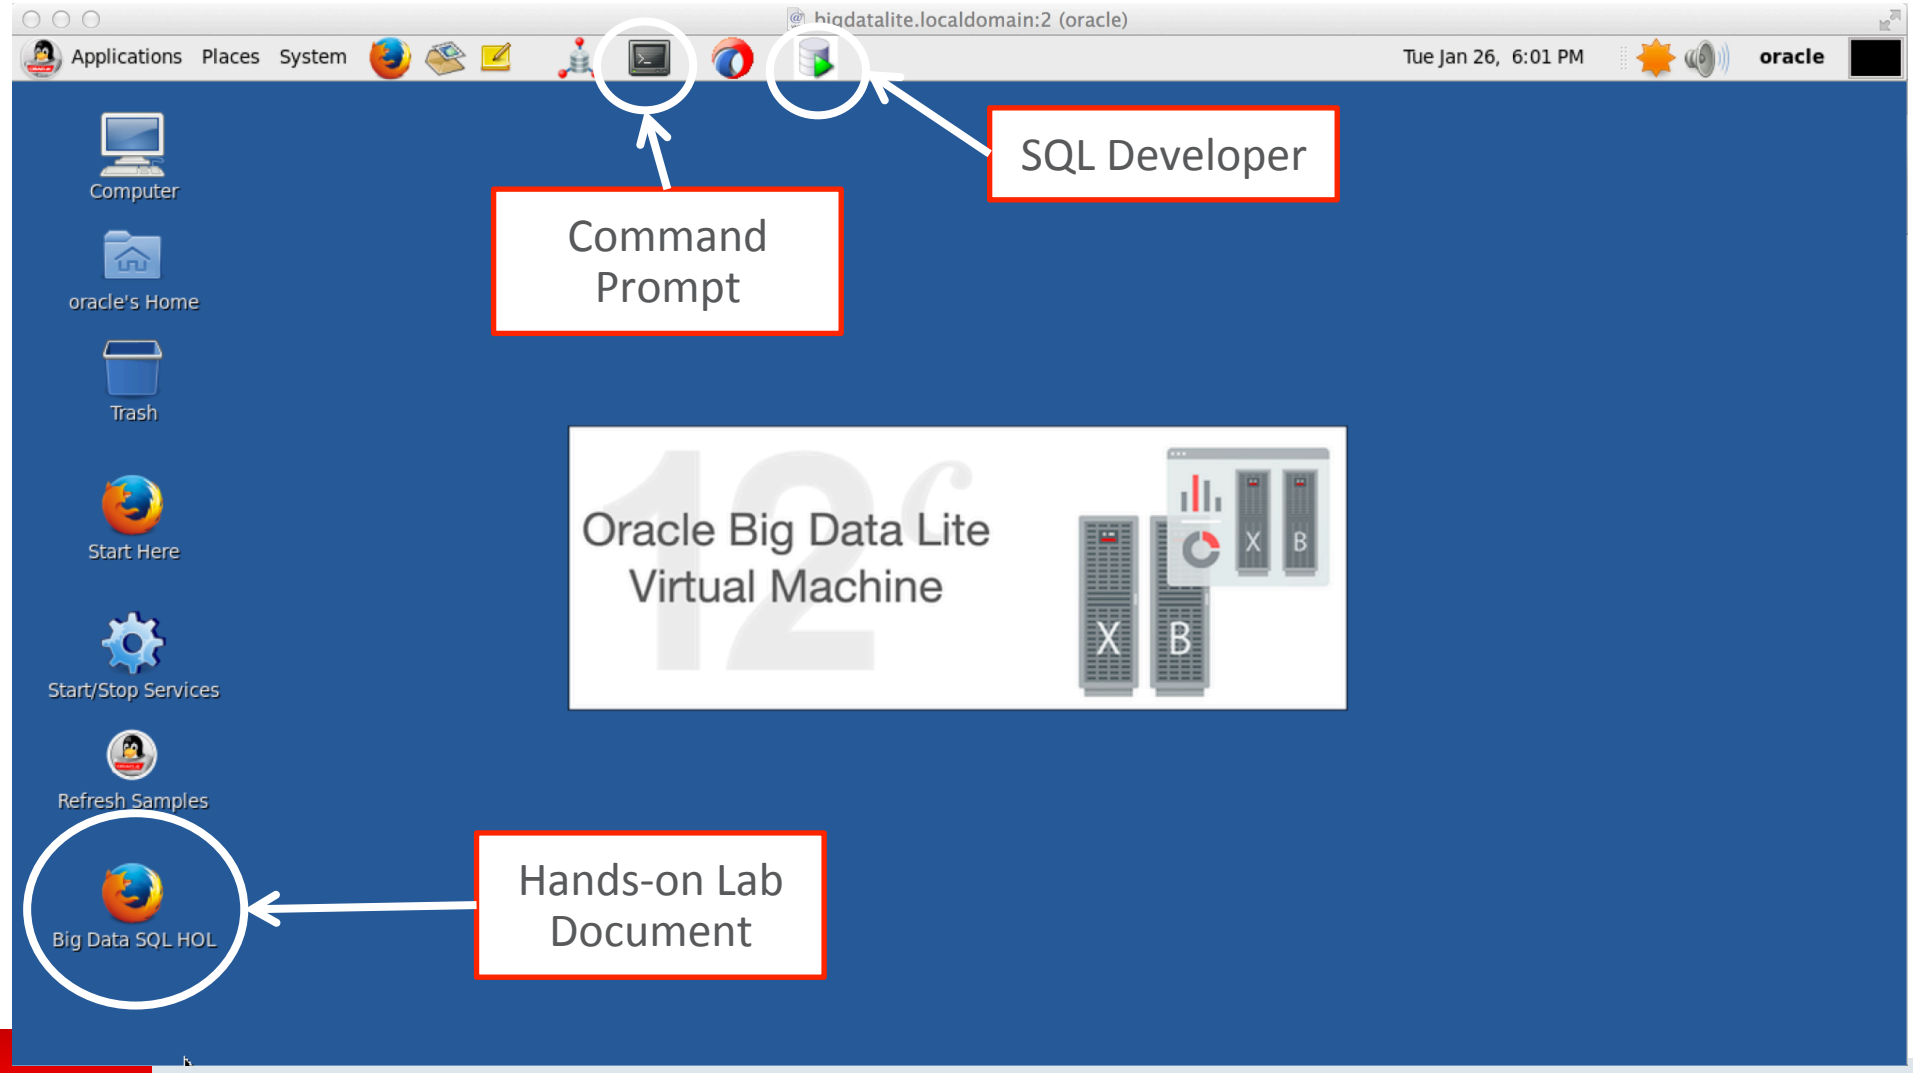

For More Information

- Get started with the software using Oracle Big Data Lite: http://www.oracle.com/technetwork/database/bigdata-appliance/oracle-bigdatalite-2104726.html
  - (simply google “Big Data Lite”)
  - Link to HOL is on that page
- Oracle.com:
  - Big Data SQL: https://www.oracle.com/database/big-data-sql/index.html
  - BDA: https://www.oracle.com/engineered-systems/big-data-appliance/index.html
- Data Warehouse Insider Blog: https://blogs.oracle.com/datawarehousing/

Integrated Cloud

Applications & Platform Services


ORACLE®