

# Oracle Big Data SQL: Roadmap to High-Performance Data Virtualization

ORACLE  
OPEN  
WORLD

October 1–5, 2017  
SAN FRANCISCO, CA

Marty Gubar  
Director - Oracle Big Data PM

Gaurav Singh  
Energy Australia  
Big Data & Data Warehouse Solution Architect

October 3, 2017

ORACLE

# 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.

# Agenda

- 1 Introduction
- 2 High Value Transformation at Energy Australia
- 3 Big Data SQL Today and Development Themes

# Big Data SQL Goals



Easily access any data  
across big data stores



Provides a **unified security model** across the sources



Analyze all data using  
Oracle's rich SQL dialect



**Fast performance** using Big  
Data SQL Smart Scan

# Big Data SQL

**Gaurav Singh**

**Big Data & Data Warehouse Solution Architect  
Information & Analytics**

# Introducing EnergyAustralia

## Energy Generation

Owens and operates a multibillion dollar portfolio of Energy Generation Assets.

Has made an investment of \$1.5 billion in Renewables.



- Yallourn Power Station
- Cathedral Rocks Wind Farm
- Hallett Power Station
- Waterloo Wind Farm
- Mt Piper Power Stations
- Stony Gap Wind Farm Project
- Tallawarra Power Station
- Mallee Solar Park Project
- Pine Dale Mine

## Energy Retailing

Is one of Australia's largest and most experienced energy retailers, servicing over 1.7 million households and businesses across Australia.



Smart Hybrid  
Solar Inverter System




## **Business Capabilities**

...utilising Oracle Big Data Solution

# High-Value capabilities delivered



A query that took 23 seconds on a legacy platform takes 0.08 seconds on the new platform



Increased speed of query performance

## New Smart Meter Analytics Capabilities

Oracle R Advanced Analytics for Hadoop



Revenue assurance improvements



## Bringing data to life

Strategic BI and Analytical tools will translate data into visual dashboards that tell a compelling story



## Golden Record of Customer & Premise



## Single View of Agent



550+ Enterprise definitions endorsed

Across these domains...



Customer



Customer Service & Agent



Financial Performance



Meter



Onboarding & Billing



Payments, Credit & Collections



Product & Pricing



Sales & Offer

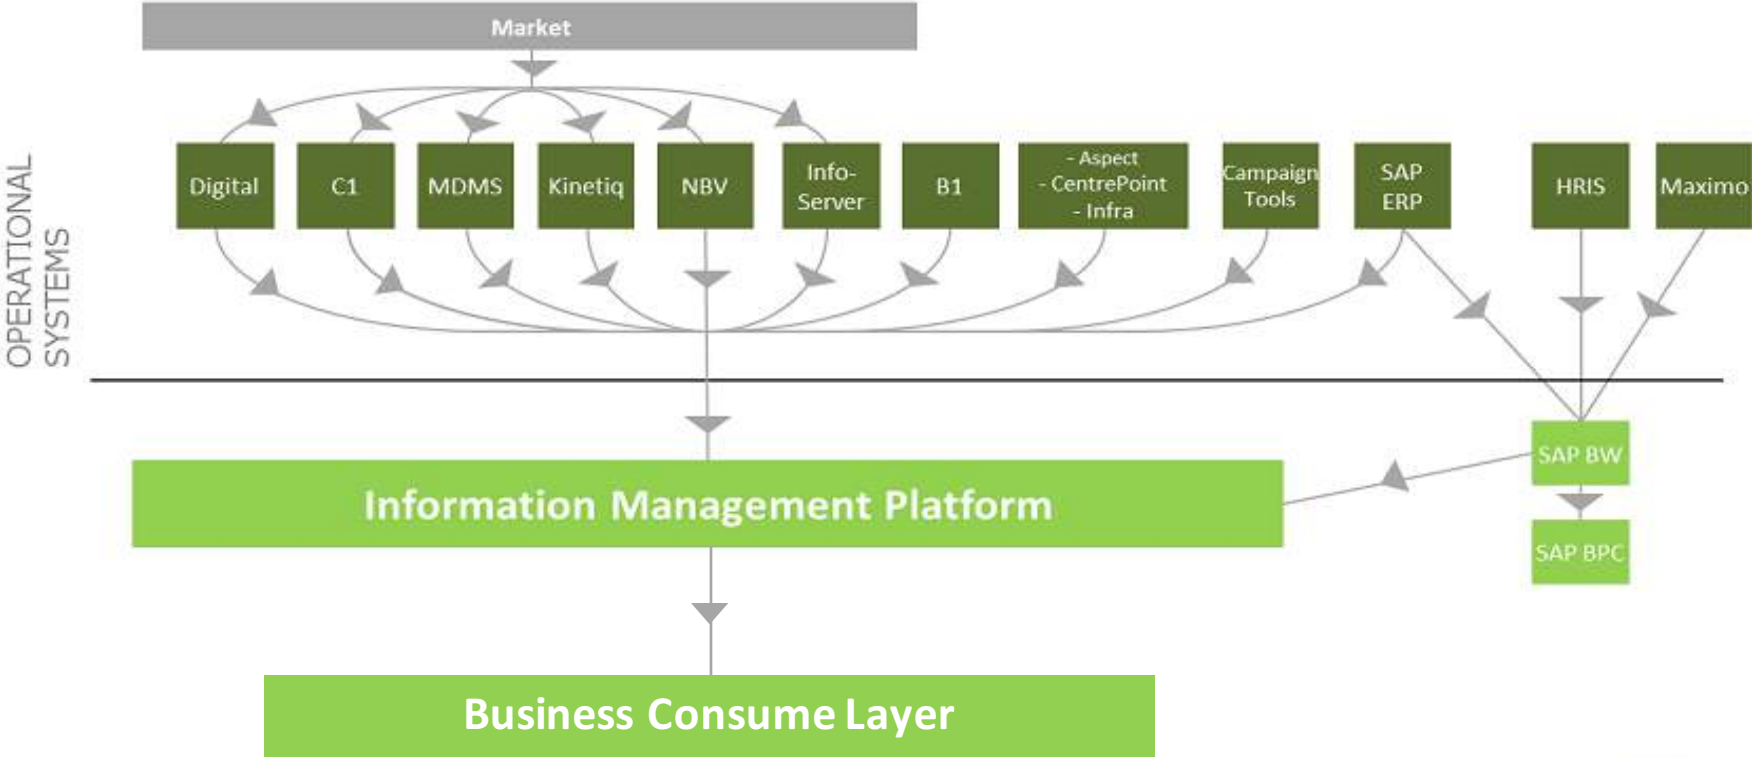


## **Past Methodology**

...utilising complex interesting  
custom built solutions

# Legacy analytics eco-system

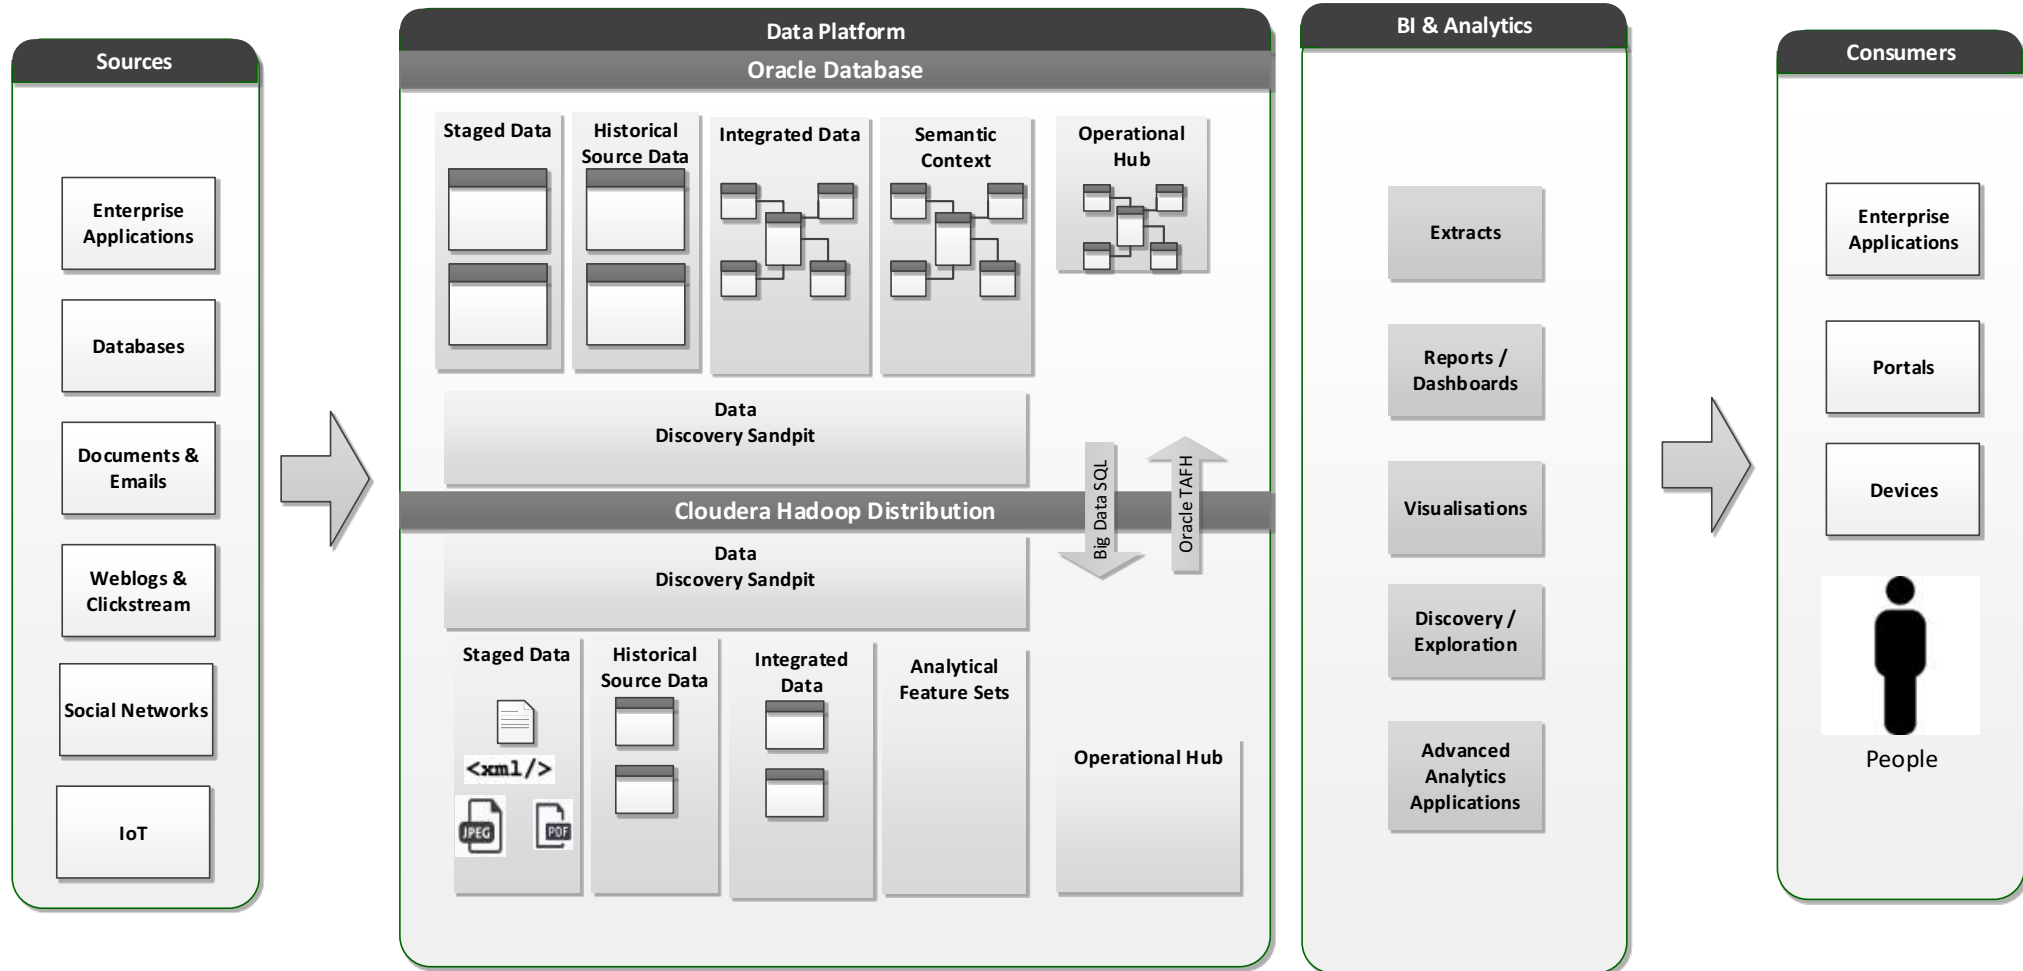
To a streamlined and efficient IM Platform



## **New Methodology**

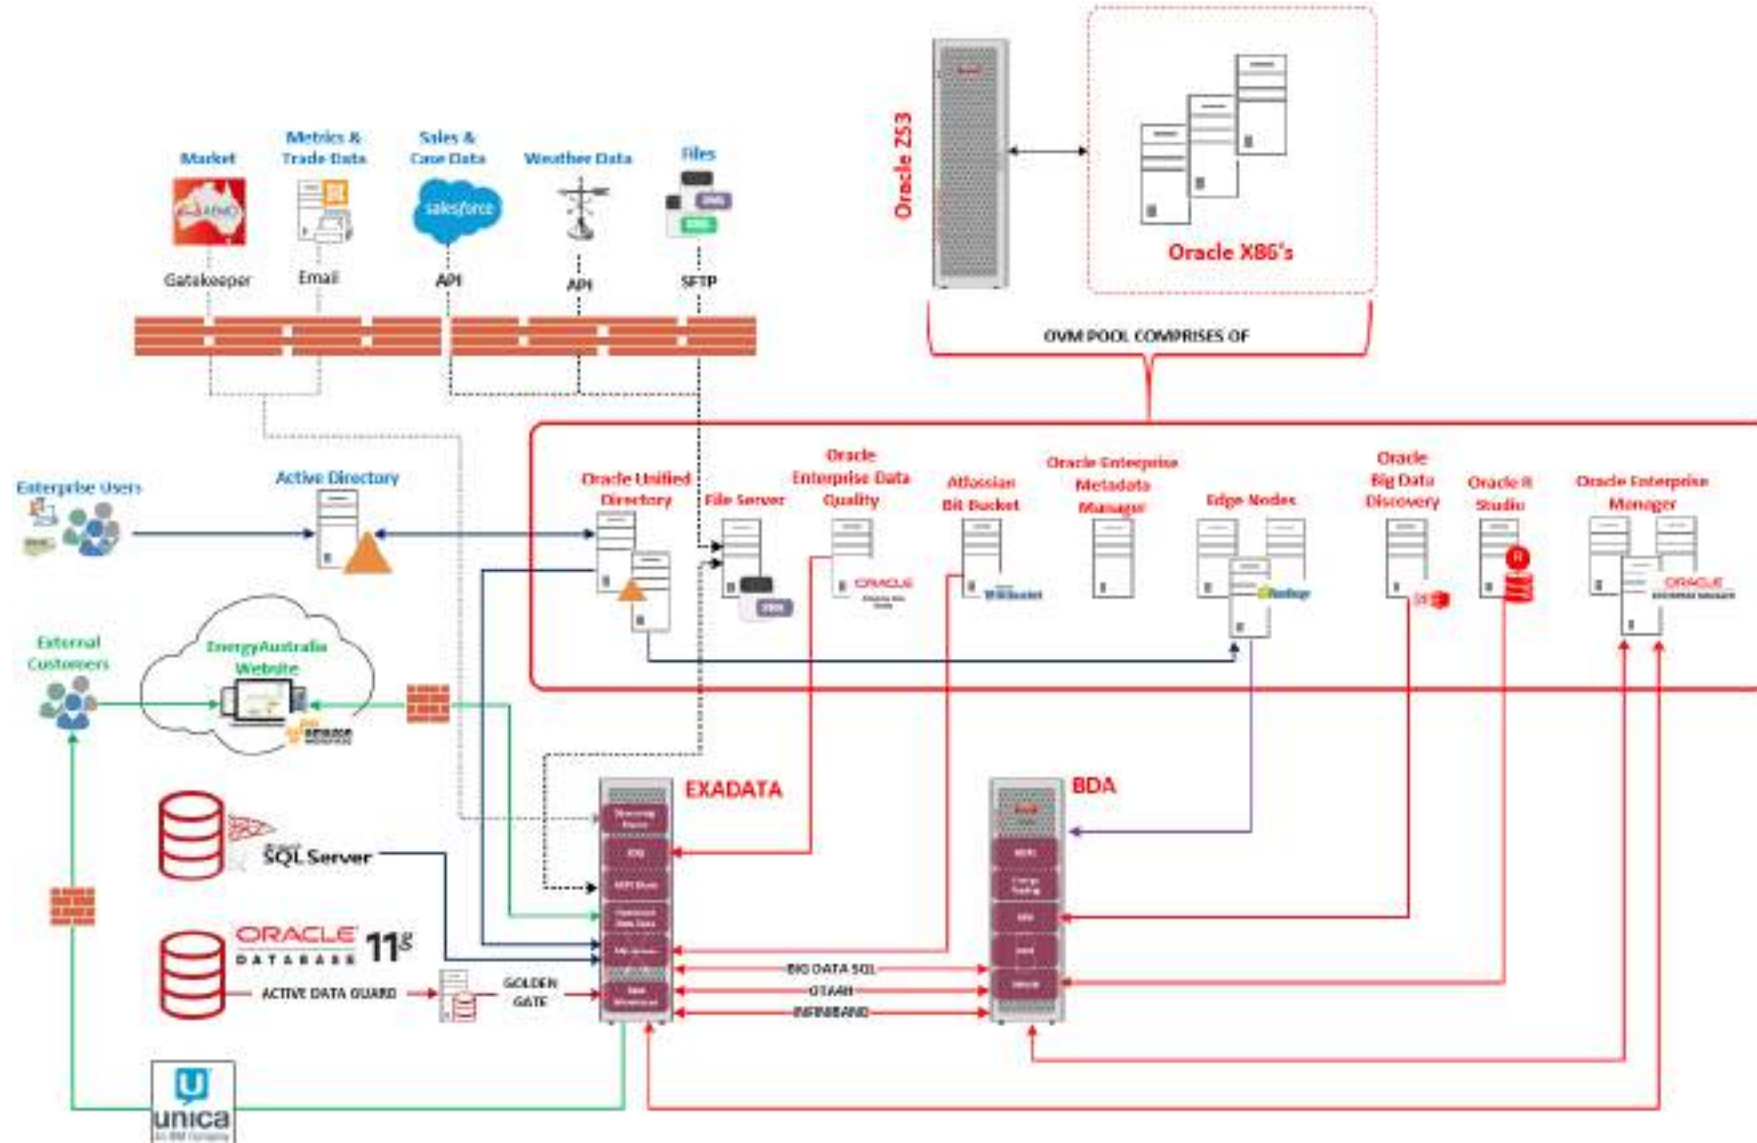
...utilising Oracle Big Data Solution

# Solution Architecture



Managed File Transfer - ELT - Database Replication - Big Data Engineering  
 Data Quality Profiling & De-duplication - Metadata Management

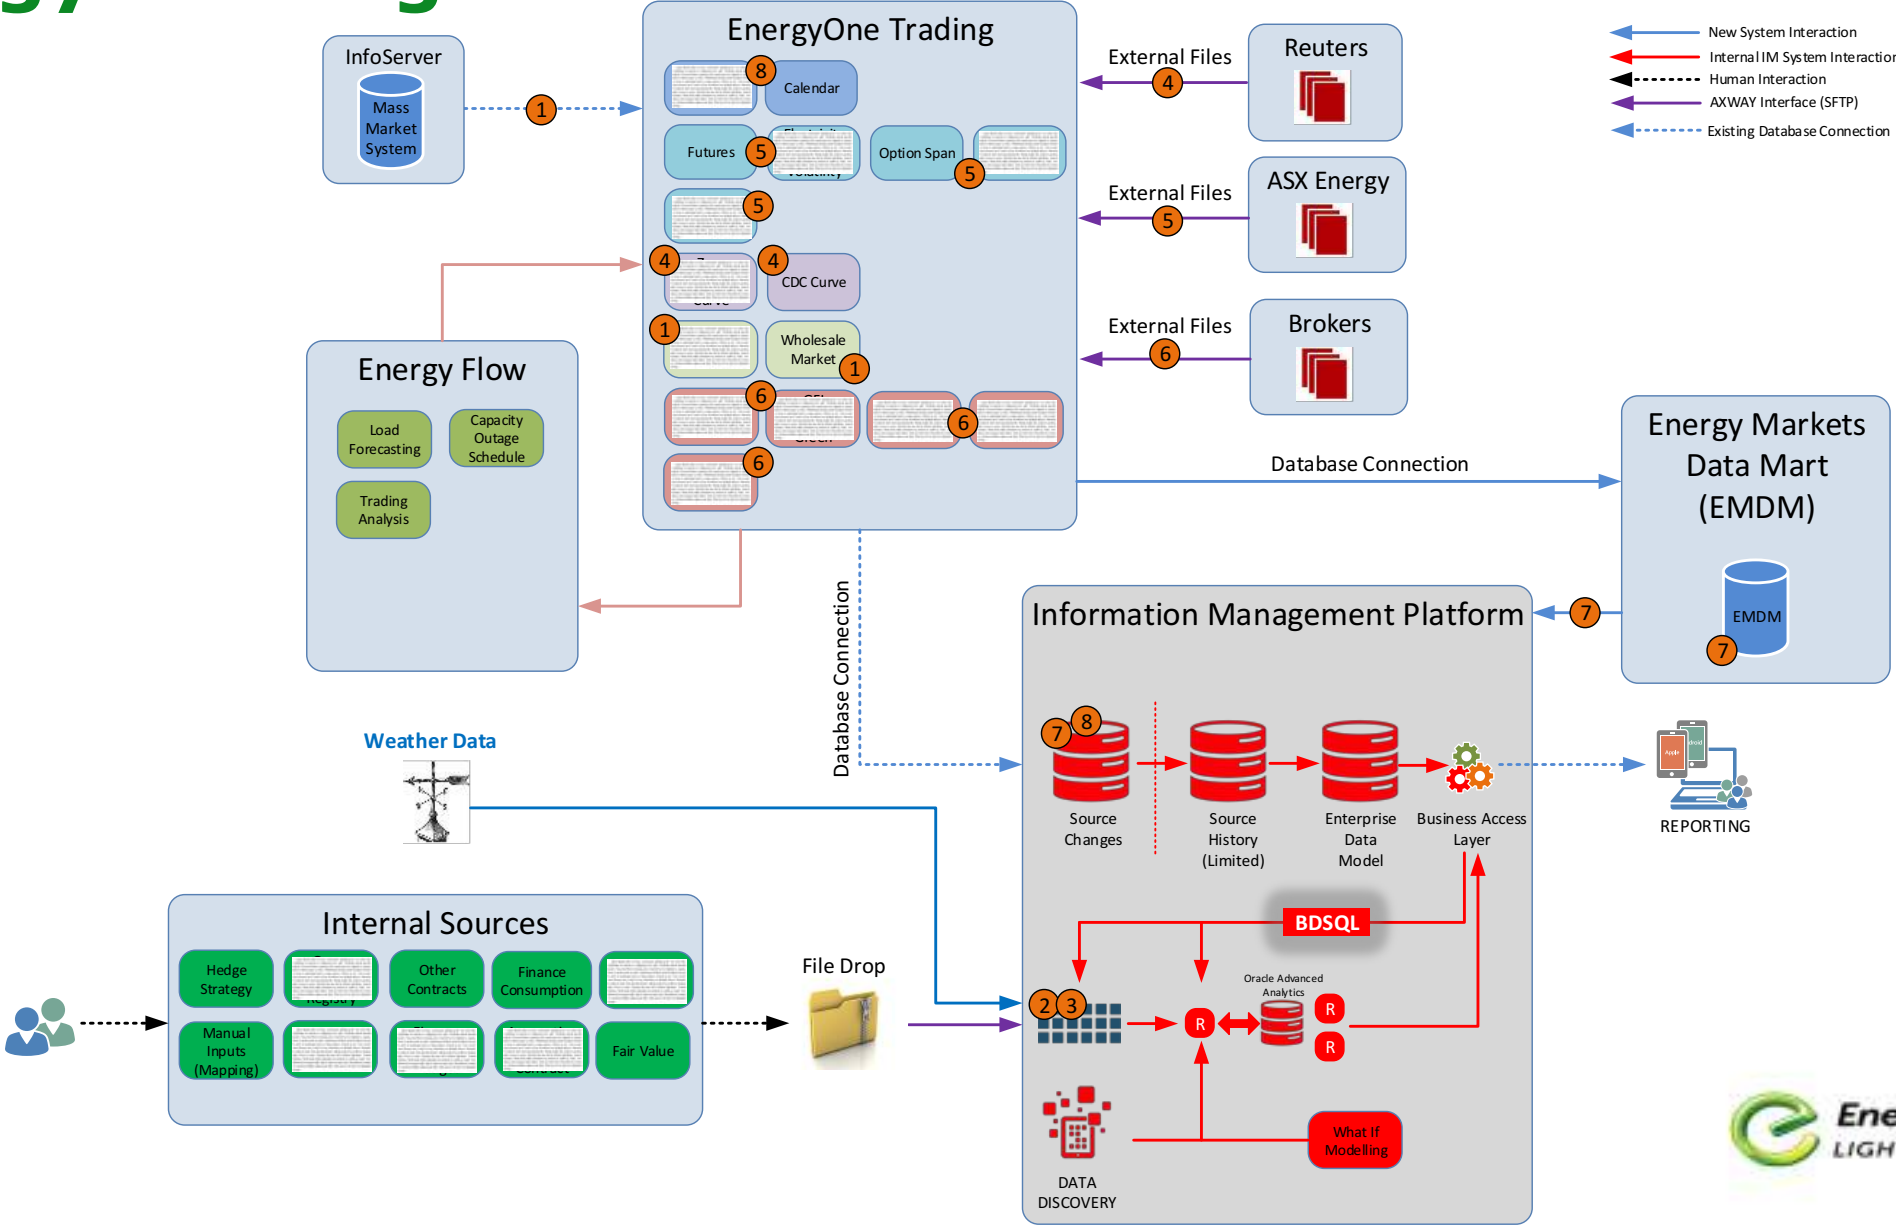
# Technical Architecture



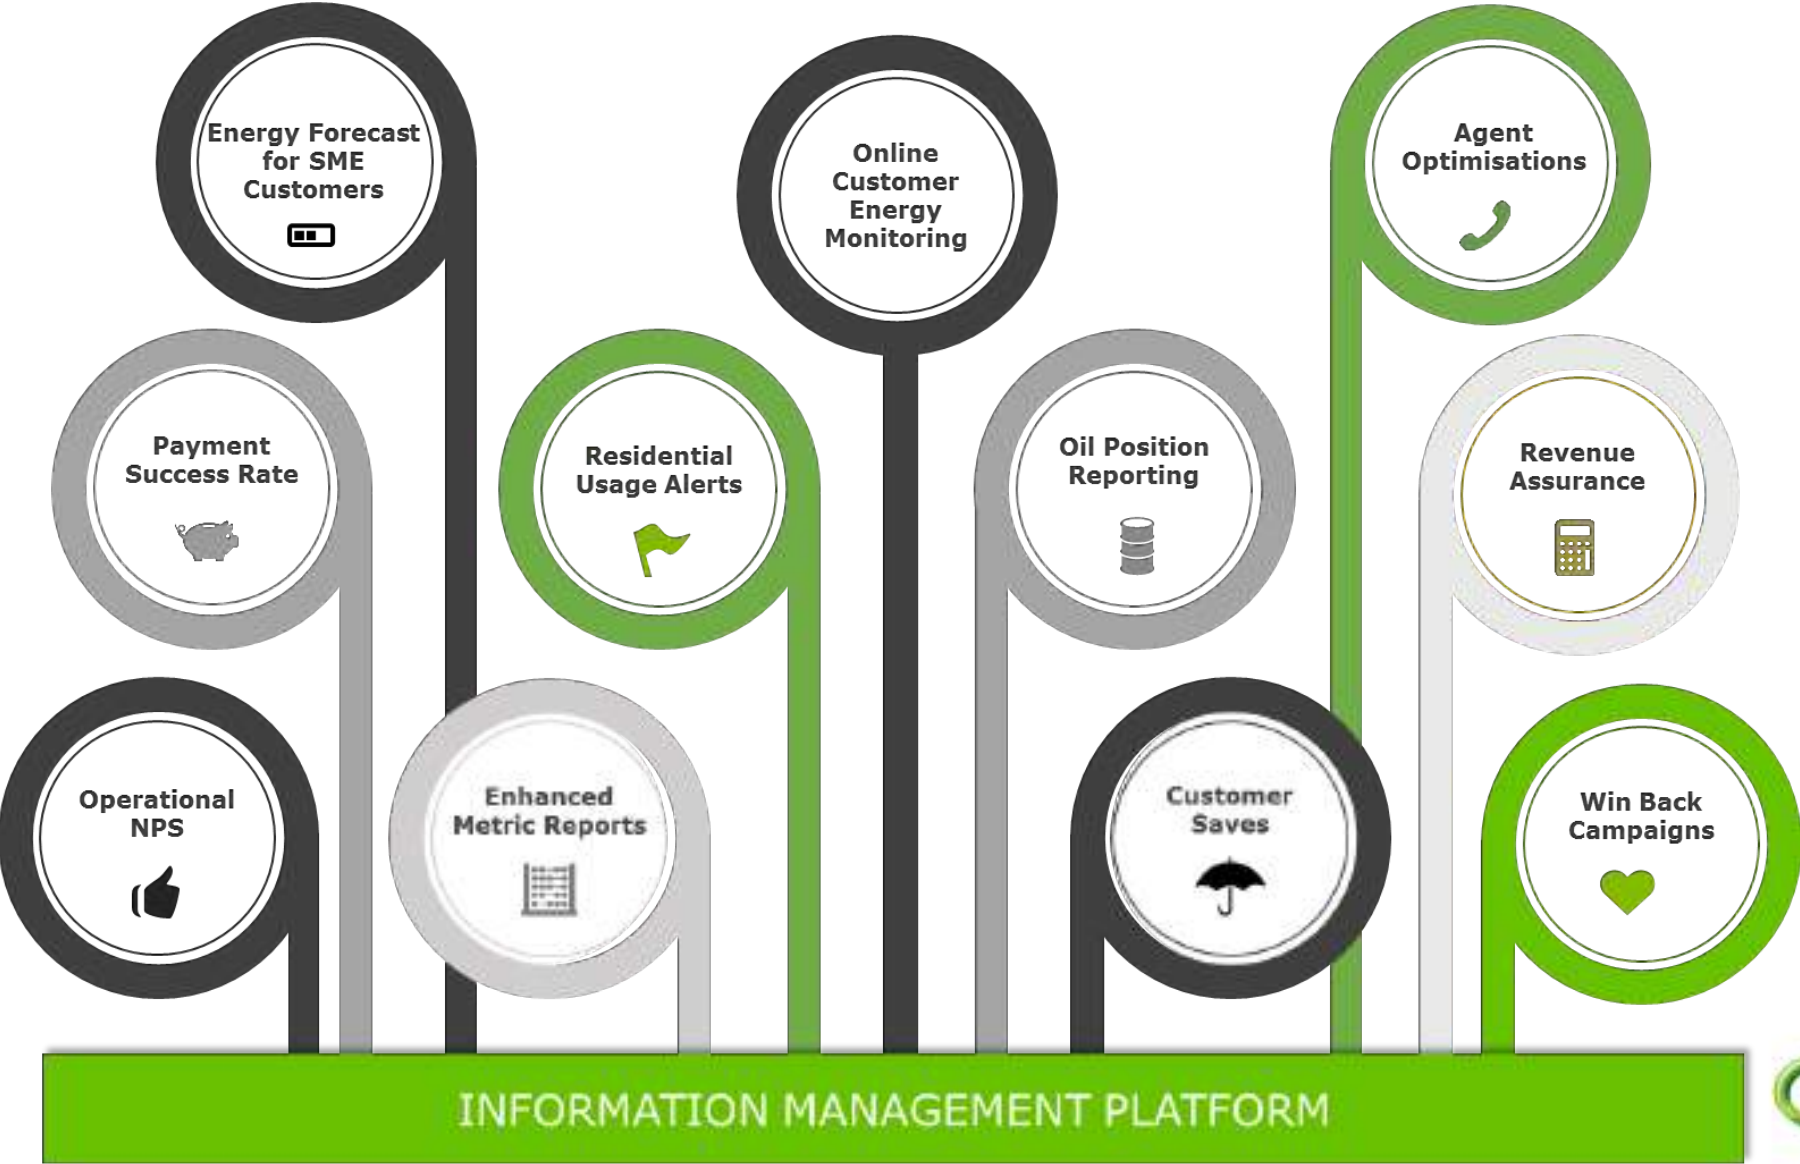
**Business Enablement**

...through Big Data SQL

# Energy Trading



# Major business initiatives delivered





---

# Learnings

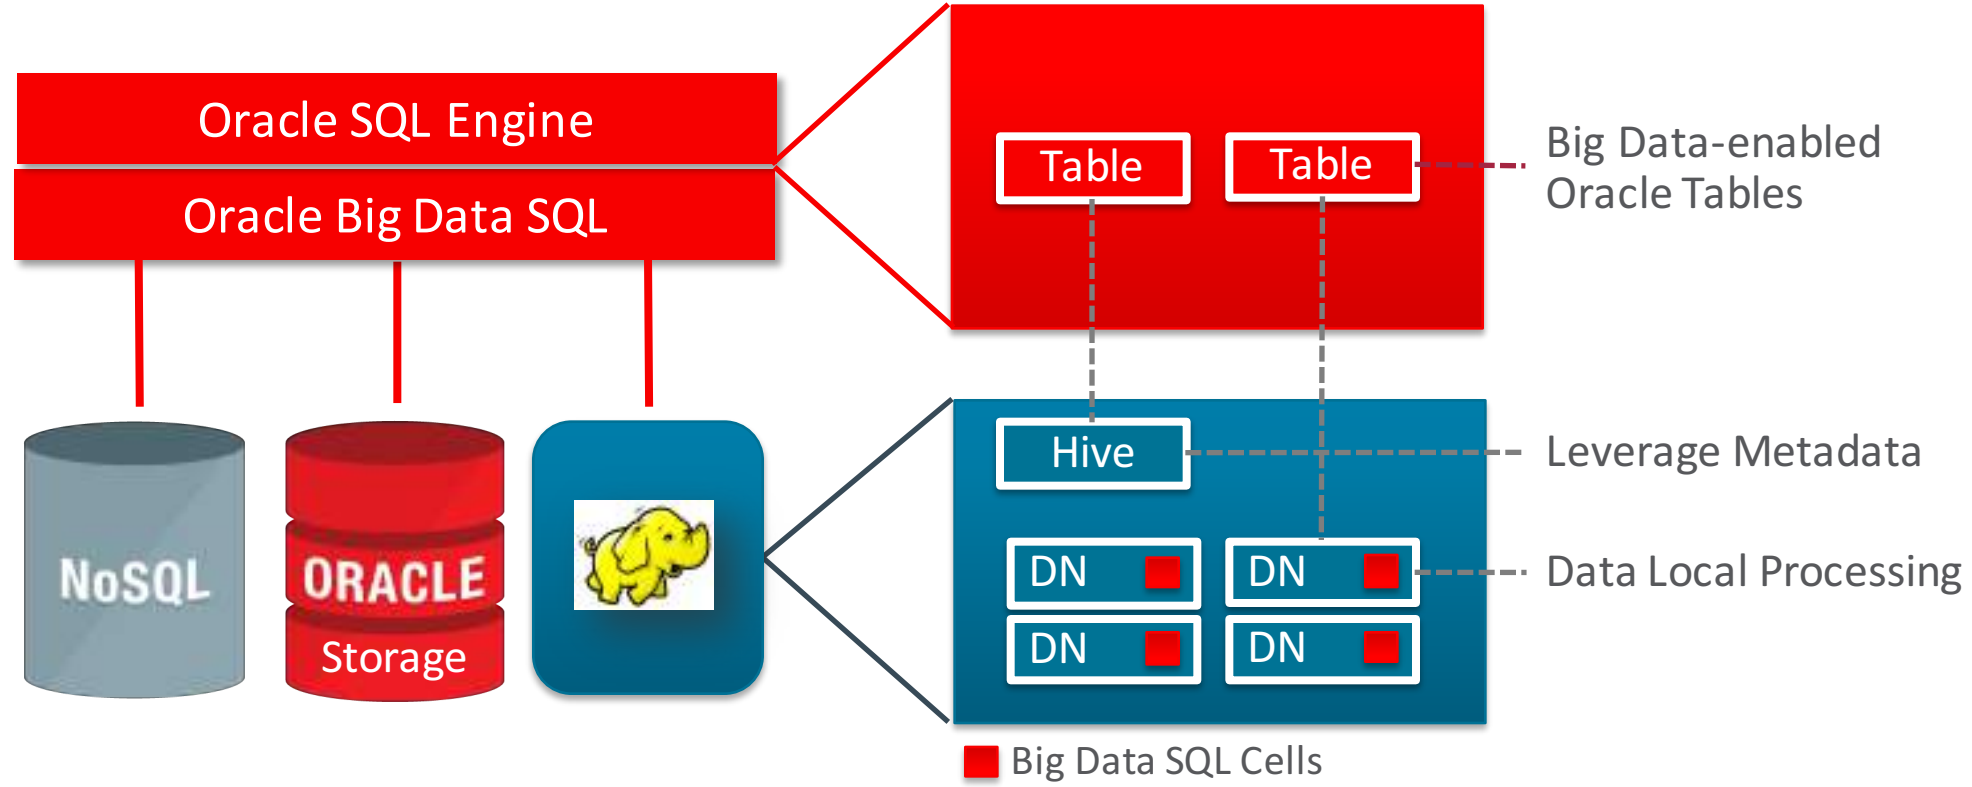
- **Big Data SQL** allows for seamless access to underlying data sets, **regardless** of source
- A **heterogenous BI ecosystem** provides the **flexibility** to complement an integrated data platform utilising **BDSQL**
- **Connect and Analyse** data and make it available to a growing pool of **business users** who do not need to learn how to Hadoop, but still **utilise the power** of Hadoop through **BDSQL**
- Running **advanced analytics in-place** helps further leverage the investment in engineered systems and makes the use case of utilising a Big Data solution **viable and beneficial** to the organisation
- **Return on investment** by a factor of **15 times**

A man with a beard and mustache, wearing a dark suit, light blue shirt, and dark tie, is looking intently at a tablet computer. The background is a blurred cityscape at night with bokeh lights. The image is overlaid with a teal geometric pattern.

# Big Data SQL Today


## Architecture and Key Features

# Big Data SQL Architecture

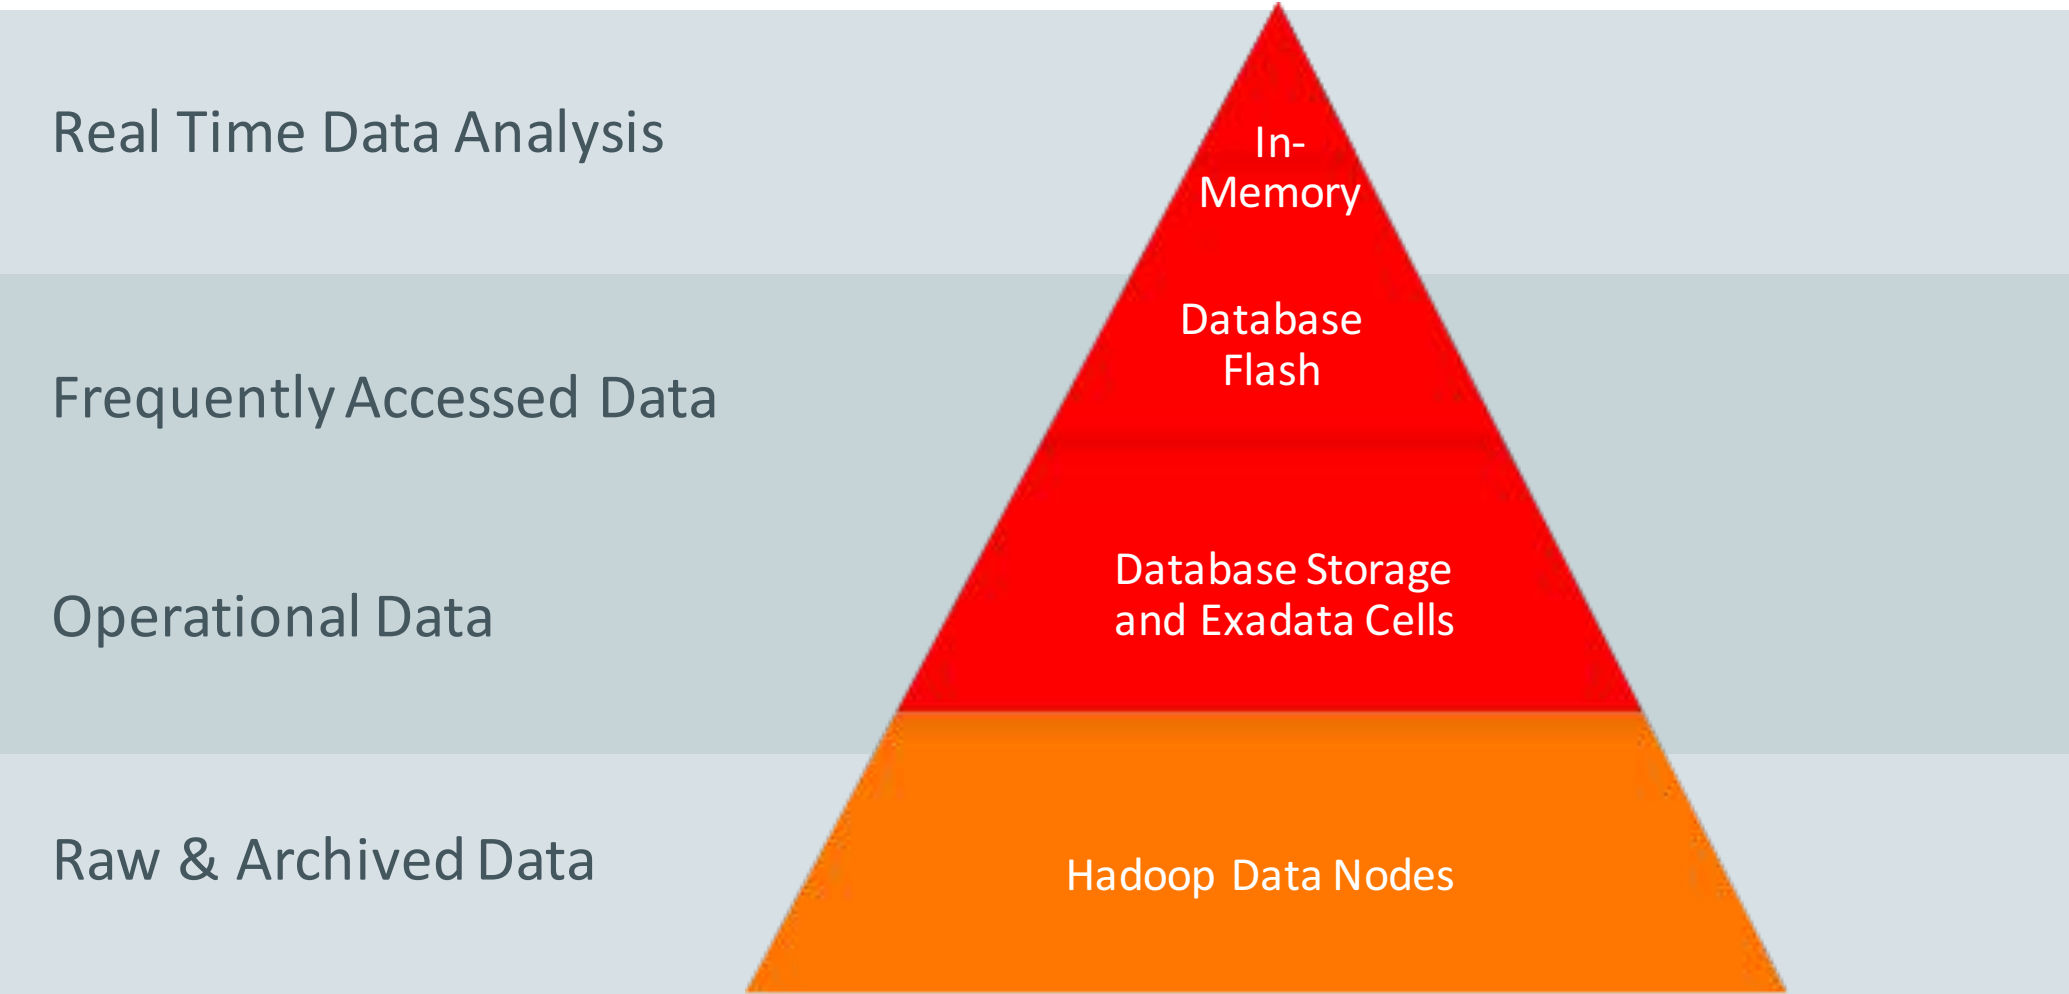


# Big Data SQL Performance Features

## Compound IO Reduction thru Smart Scans




# Big Data Use Case: Information Lifecycle Management

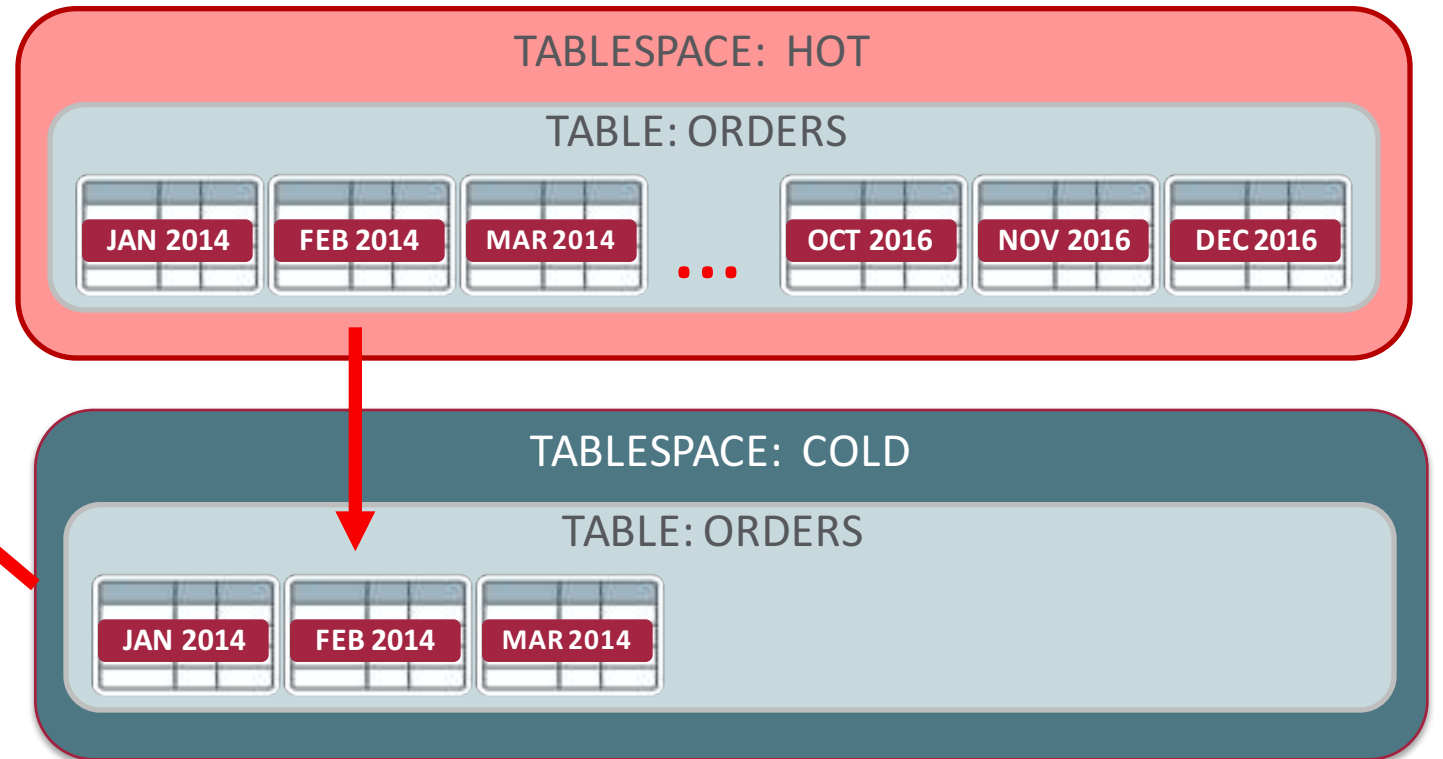


# Archive Data to “Cold” Partitions

HDFS:




Oracle Database 12c:



# Flexible Deployment Options

## Engineered Systems



cloudera

12<sup>c</sup>

## Oracle Cloud



cloudera

12<sup>c</sup>

## Commodity



cloudera



12<sup>c</sup>

A man with a beard and mustache, wearing a dark suit, light blue shirt, and dark tie, is looking down at a smartphone he is holding. The background is a blurred cityscape at night with bokeh lights. The image is overlaid with a teal geometric pattern.

# Big Data SQL Cloud Service

Cloud Simplicity



# Dashboard

Cloud Account a422306 (traditional)

## Welcome

**Create Instance**  
Provision a new service in minutes

**Account Management**  
Administer and manage your account and orders

**Customize Dashboard**  
Fine tune your cloud services overview

**What's New**  
Learn about changes in the latest releases

## Cloud Services

Start with an Exadata and Big Data Cloud Service

Important Notifications

**Exadata**  
Subscription ID: 20151217012259

Instances

**BigData**  
Subscription ID: 537050936

Instance



Big Data Cloud Service

Services

Activity

Welcome!

As of Aug 30, 2017 6:32:58 PM UTC

### Services

Create Service



[bdcs1](#)

Version: 4.8.1

Created On: Aug 28, 2017 5:00:15 PM UTC

Click to view the details for this service instance

▶ Service Create and Delete History



Overview

5 Nodes

### Service Overview

As of Aug 30, 2017 6:33:30 PM UTC

- Cloudera Manager console
- Big Data Manager console
- Start
- Stop
- Restart
- Update
- Add Association
- Service Credentials
- View Activity

Version: 4.8.1

Big Data Appliance System: bdcdevelopment

Secure Setup: false

### Resources

	Host Name: <code>cfcbv3872.us2.oraclecloud.com</code> Public IP: 129.152.232.4	Role: Permanent Hadoop Node	
	Host Name: <code>cfcbv3872.us2.oraclecloud.com</code> Public IP: 129.152.232.5	Role: Permanent Hadoop Node	
	Host Name: <code>cfcbv3873.us2.oraclecloud.com</code> Public IP: 129.152.232.6	Role: Permanent Hadoop Node	
	Host Name: <code>cfcbv3874.us2.oraclecloud.com</code> Public IP: 129.152.232.7	Role: Permanent Hadoop Node	



Overview

5 Nodes

### Add Association

Association Name: Big Data Appliance - Exadata e

Service Type: Oracle Database Cloud Service

\* Service Name: bdsdb

\* Description: This is a optional description that is used to help identify the purpose for this association

\* Cloudera Manager password:

Add Cancel

Cloudera Manager password

Public IP: 129.152.232.5	Role: Permanent Hadoop Node
Host Name: cfcbv3873.us2.oraclecloud.com	Role: Permanent Hadoop Node
Public IP: 129.152.232.6	Role: Permanent Hadoop Node
Host Name: cfcbv3874.us2.oraclecloud.com	Role: Permanent Hadoop Node
Public IP: 129.152.232.7	Role: Permanent Hadoop Node



Overview

5 Nodes

### Add Association

Association Name: Big Data Appliance - Exadata s

Service Type: Oracle Database Cloud Service

\* Service Name: bdsdb

\* Description: This is a optional description that is used to help identify the purpose for this association

\* Cloudera Manager password: .....

Add Cancel

As of Aug 30, 2017 6:33:30 PM UTC

Public IP:	129.152.232.5	Role:	Permanent Hadoop Node
Host Name:	cfclbv3873.us2.oraclecloud.com	Role:	Permanent Hadoop Node
Public IP:	129.152.232.6	Role:	Permanent Hadoop Node
Host Name:	cfclbv3874.us2.oraclecloud.com	Role:	Permanent Hadoop Node
Public IP:	129.152.232.7	Role:	Permanent Hadoop Node



Use this page to search and review activities of Cloud Services in your identity domain.

Search Activity Log

Start Time Range: 8/29/2017 6:37 PM

Search to Time

Service Name: bdcs1

Operation Status: All

Service Type: Big Data Cloud Service

Search Reset

Operation: All



With just a few clicks, Big Data SQL Cloud Service enables Exadata Cloud Service to efficiently query big data stores

- Aug 30, 2017 6:36:26 PM UTC Started operation to create association between [bdcs1] and [poc00] in Identity Domain [a422306].
- Aug 30, 2017 6:36:26 PM UTC Started Create association for tenant domain=[a422306] service=[bdcs1]
- Aug 30, 2017 6:36:43 PM UTC Started Action [createAssociation]
- Aug 30, 2017 6:36:43 PM UTC Started Action [createAssociationstatus] to check for asynchronous result
- Aug 30, 2017 6:36:43 PM UTC Finished Action [createAssociation] with return status=[Running]

Restart Service	bdcs1	Big Data Cloud S...	Succeeded	Aug 30, 2017 3:57:31 PM UTC	Aug 30, 2017 4:01:47 PM UTC	marlin.guba...
-----------------	-------	---------------------	-----------	-----------------------------	-----------------------------	----------------

# Securely Access All Data



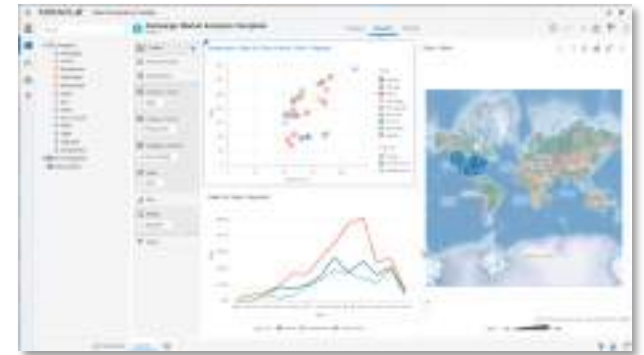
## Define

- Create big data-enabled tables
- Leverage underlying metadata
- Use tooling to simplify creation



## Secure

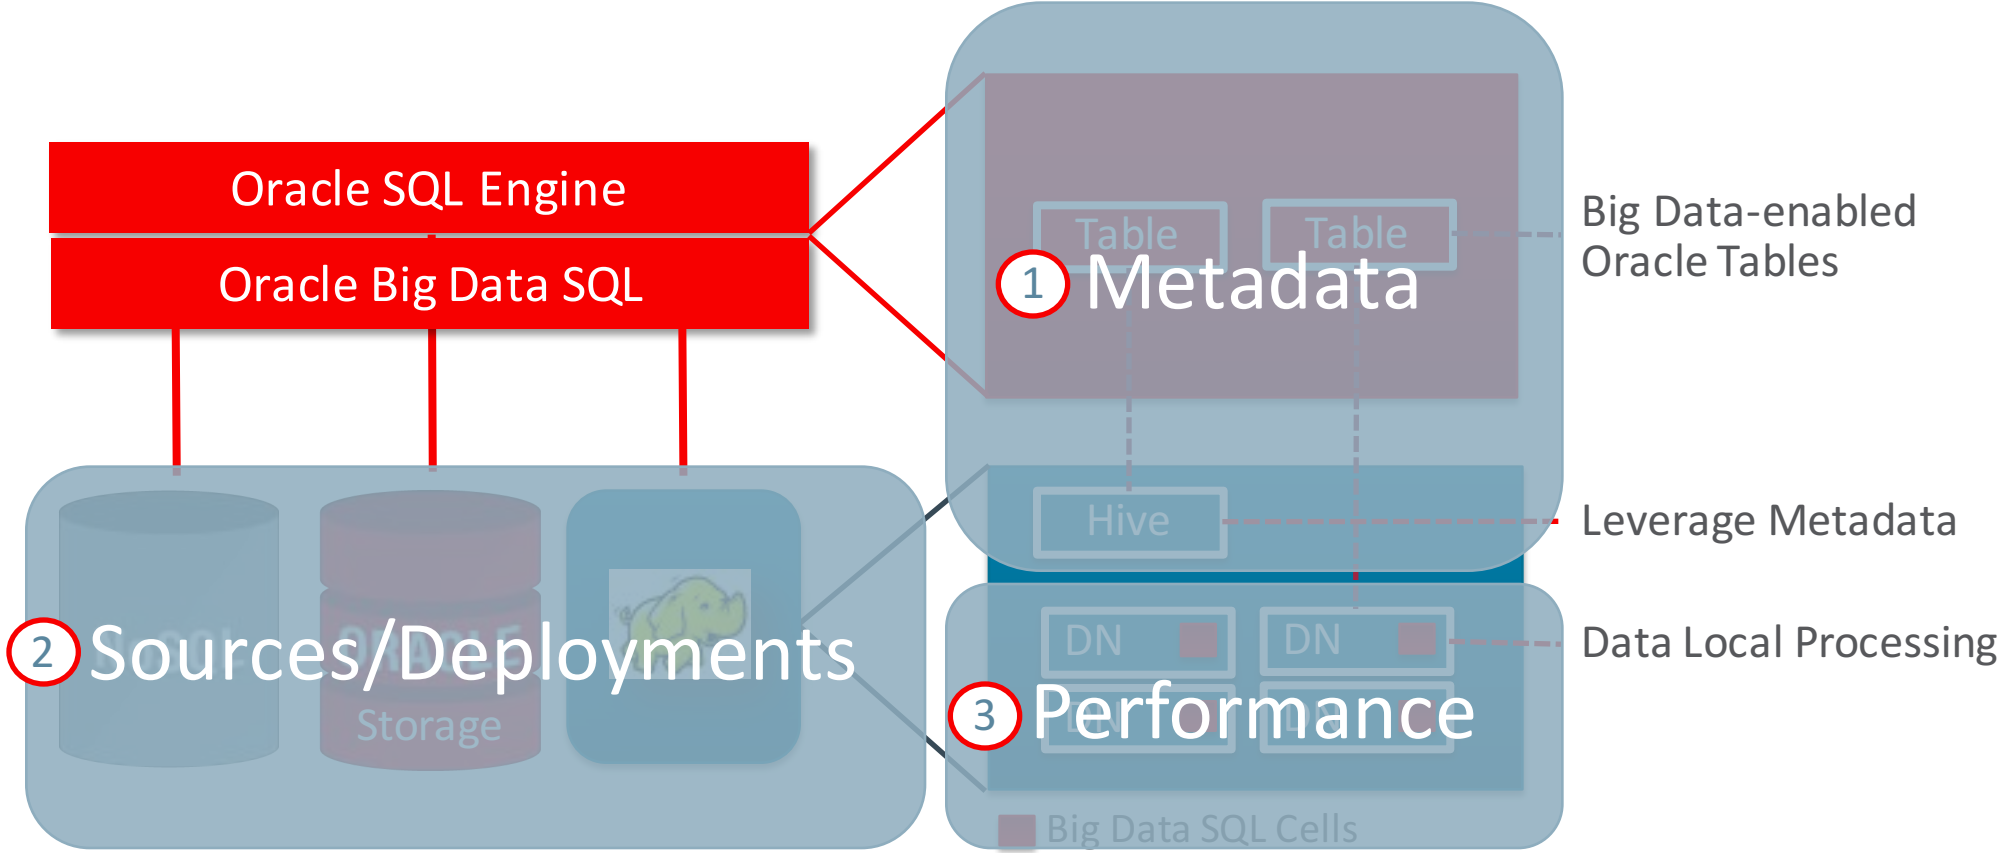
- Grant access to users/groups
- Apply Oracle advanced security policies to authorize access



## Query

- Use Oracle SQL to analyze data
- Seamlessly combine big data sources
- Use favorite tools and applications

# Big Data SQL Focus Areas



Confidential – Oracle Internal/Restricted/Highly Restricted

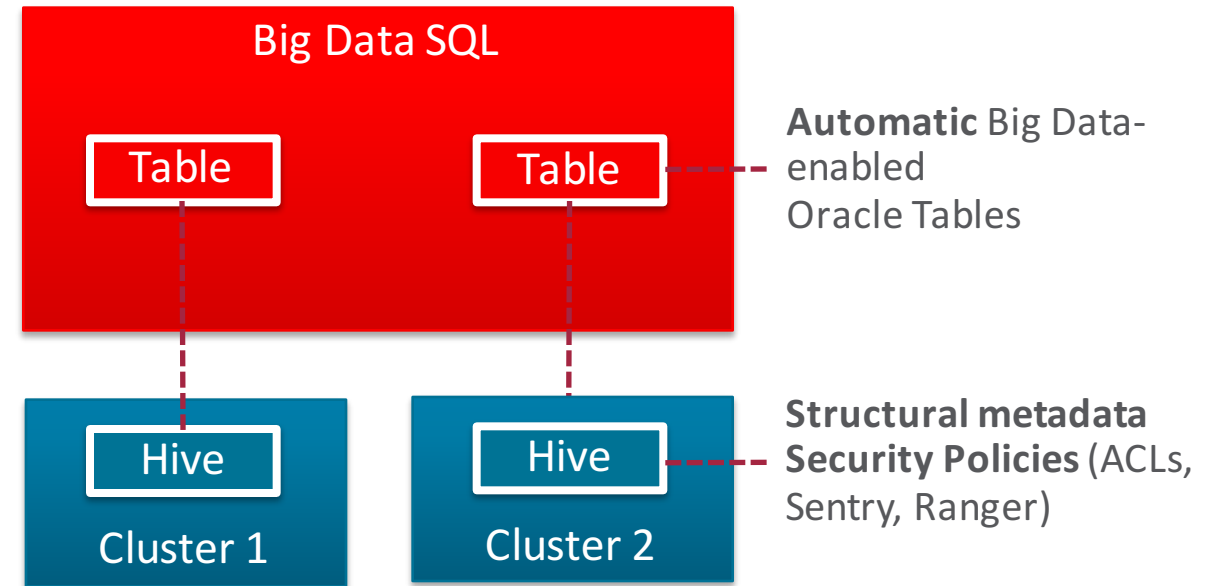




# Metadata

## Big Data SQL Aggregates Metadata from Underlying Sources

- **Simplify Access**
  - Automatic table definition when metadata is available
  - Discover metadata from structured files
- **Support Source Security Policies**
  - Use security policies defined on source
- **Extend with Unique Capabilities**
  - Advanced Oracle Security policies



# New Sources



## Apache Kafka

- Query Apache Topics for real time analysis
- Blend real time streams with other data sources



Oracle Object Store



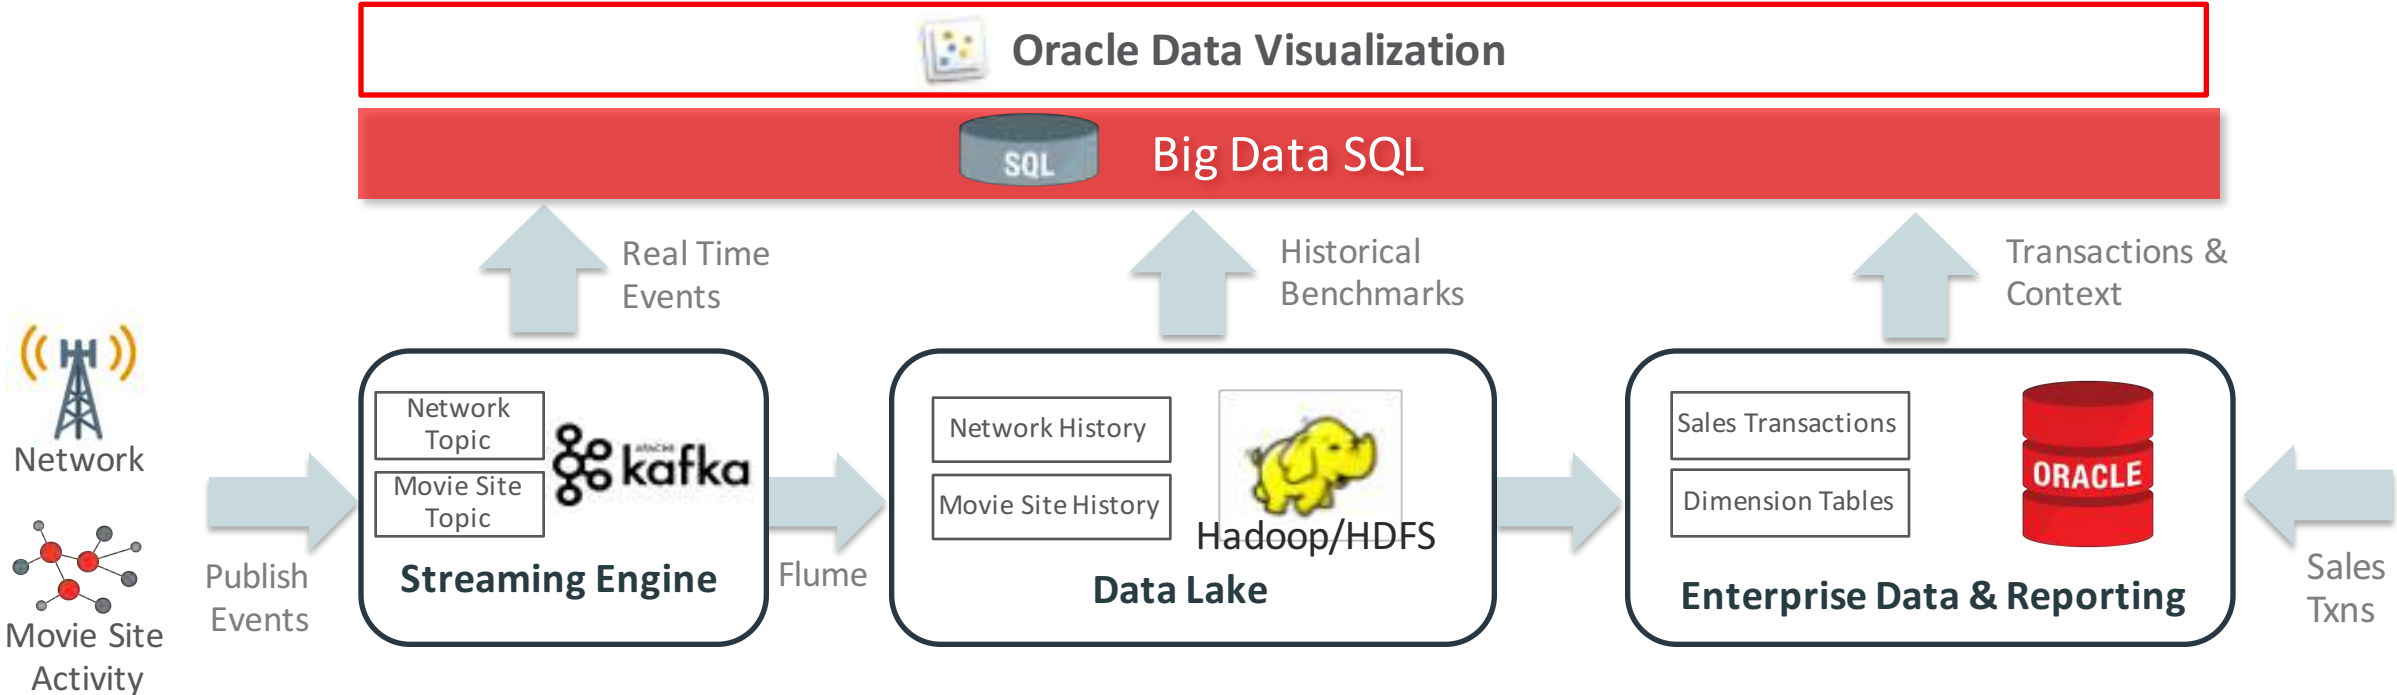
Amazon S3

## Object Stores

- Efficiently query object store data
- Support any file type

# Demonstration Scenario

## Analyzing Real-time Streams




Search

- KAFKA\_DEMO\_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Streaming Demo Project

Prepare Visualize Narrate



Search


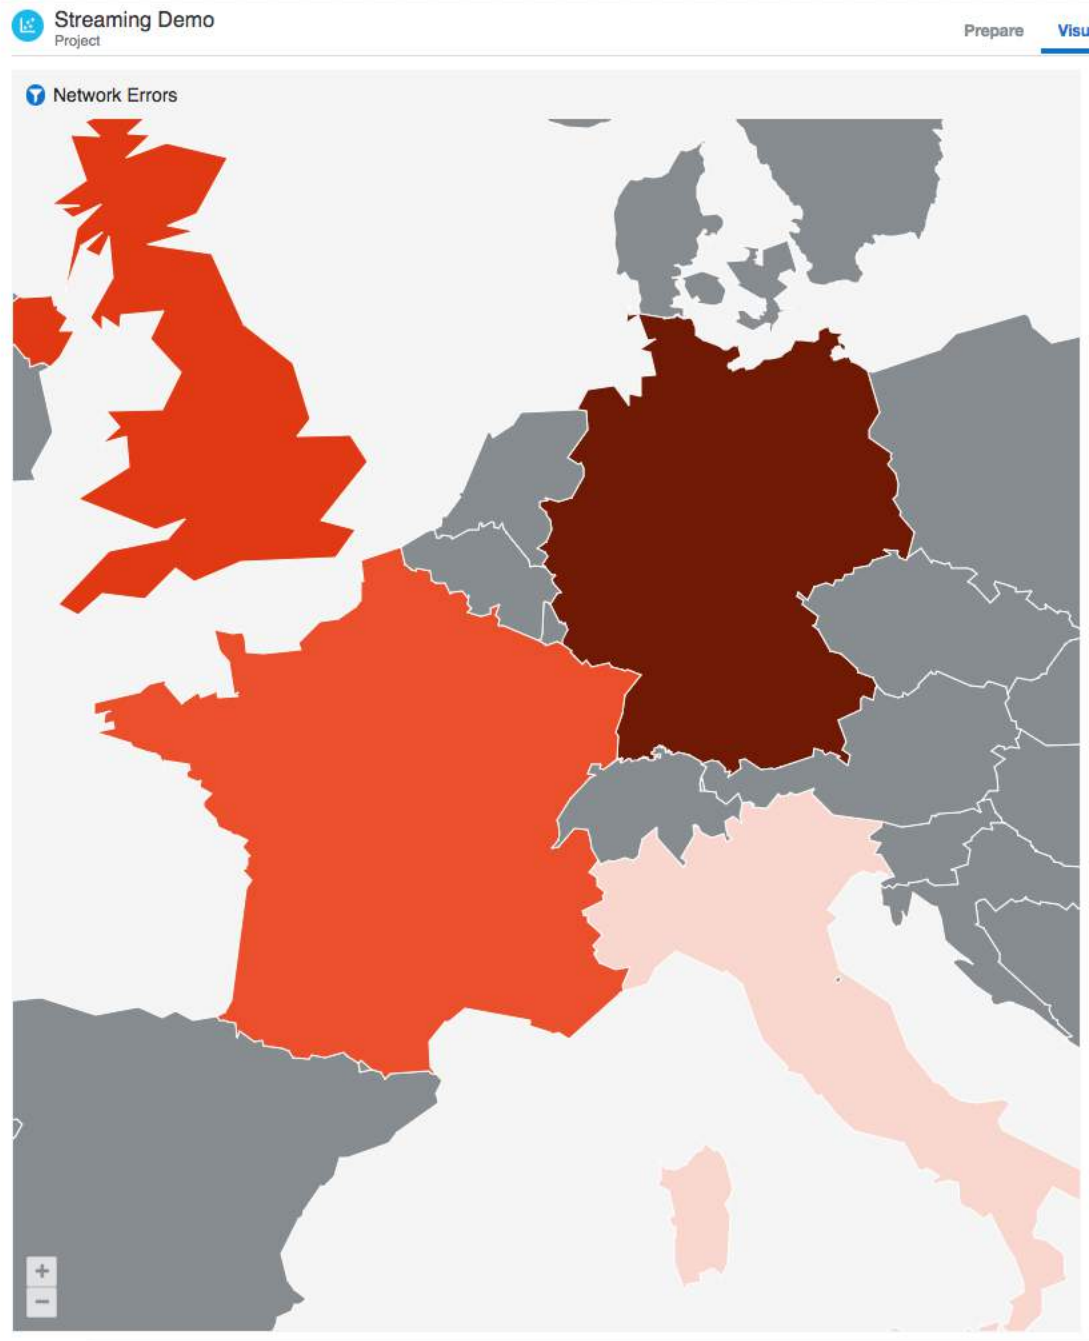
KAFKA\_DEMO\_TABLE

- Day
- Country
- Network Errors
- Hour
- Revenue Benchmark
- Revenue Stream

My Calculations

Value Labels

Add Calculation



Search

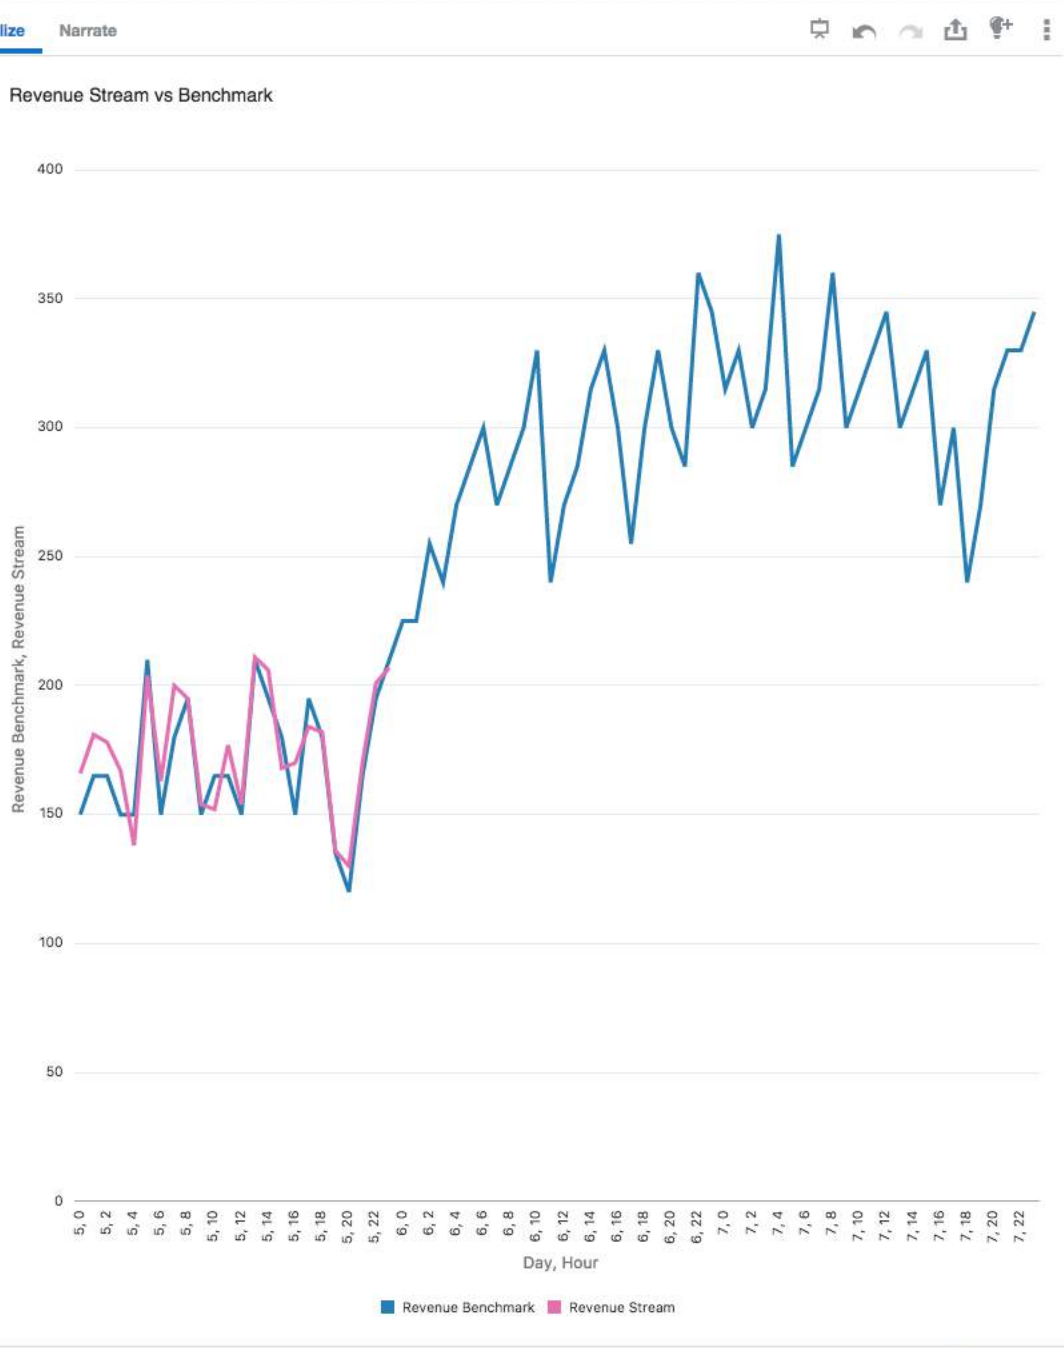
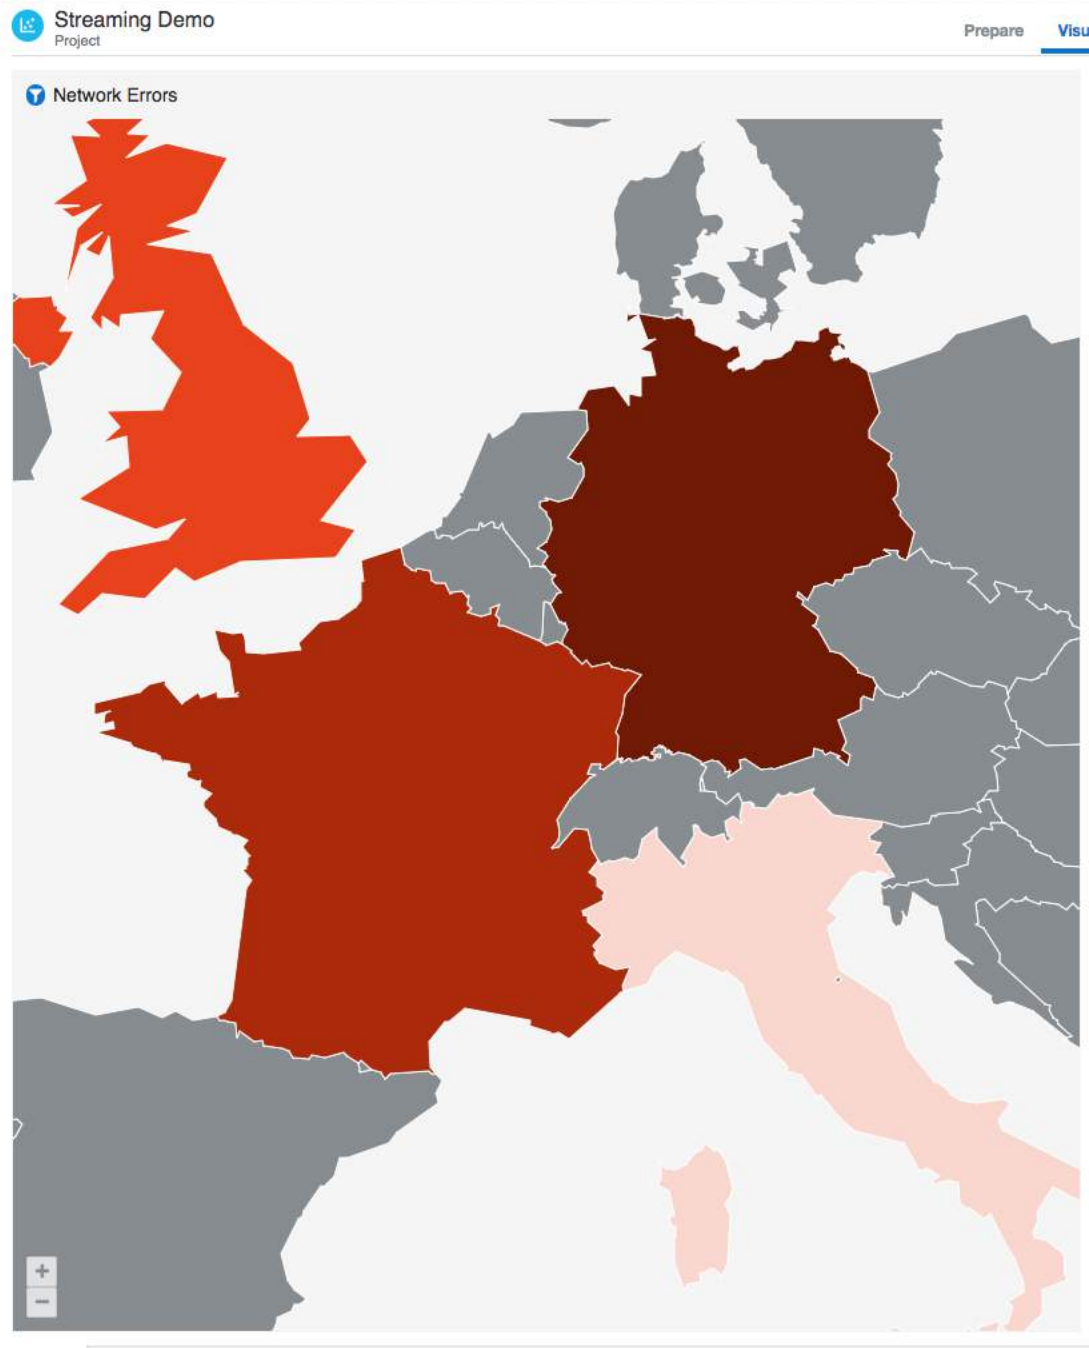
KAFKA\_DEMO\_TABLE

- Day
- Country
- Network Errors
- Hour
- Revenue Benchmark
- Revenue Stream

My Calculations

Value Labels

Add Calculation



Search

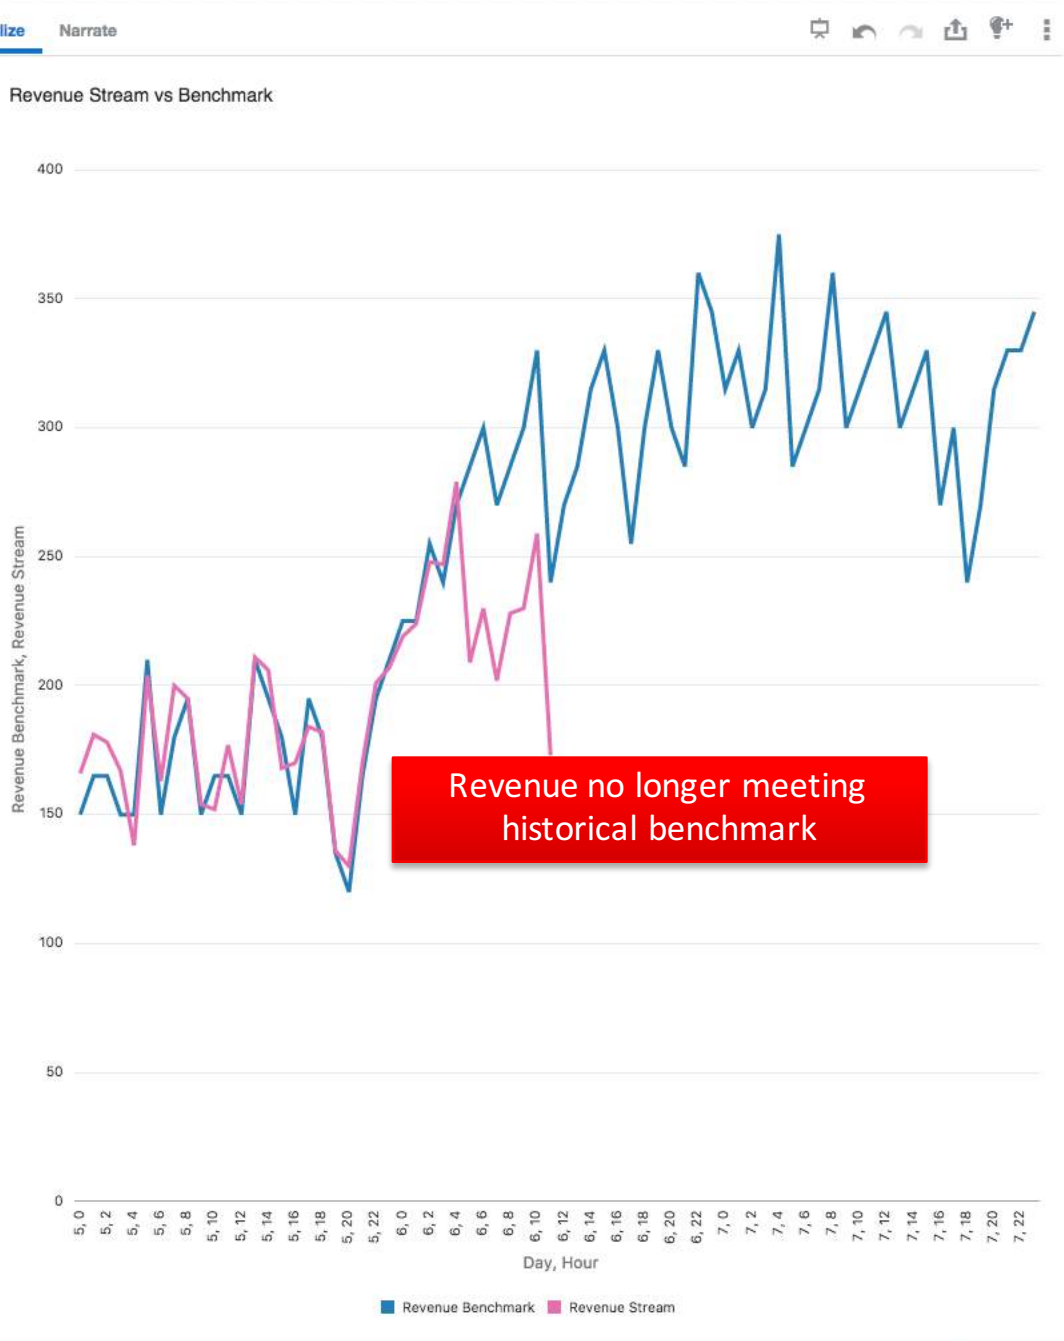
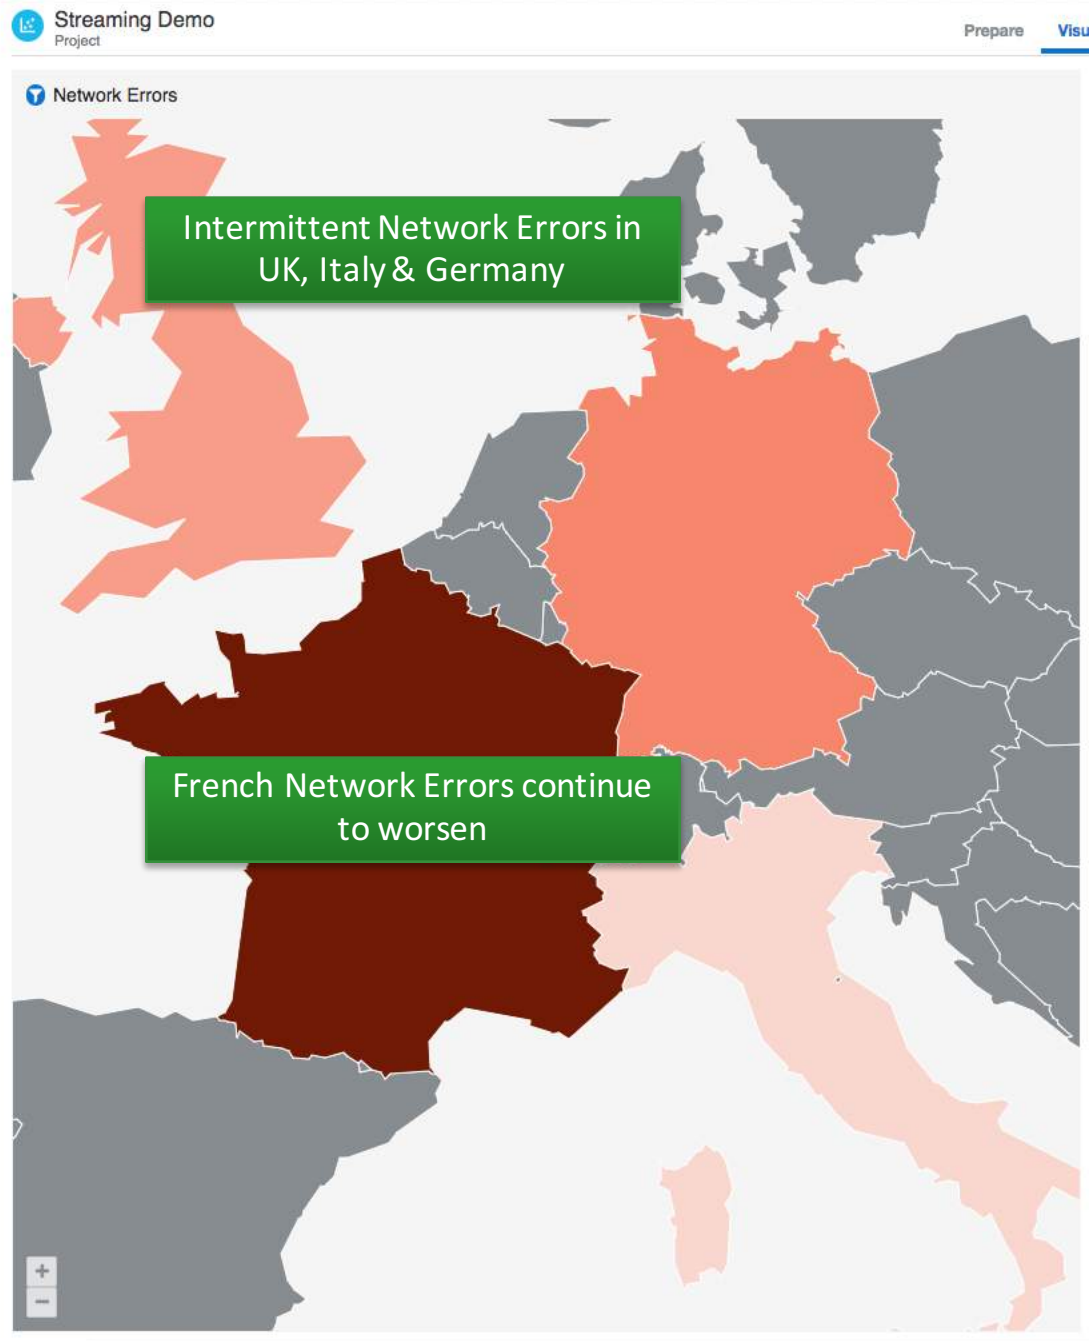
KAFKA\_DEMO\_TABLE

- Day
- Country
- Network Errors
- Hour
- Revenue Benchmark
- Revenue Stream

My Calculations


Value Labels

Add Calculation




- KAFKA\_DEMO\_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Network Errors



Revenue Stream vs Benchmark





Search


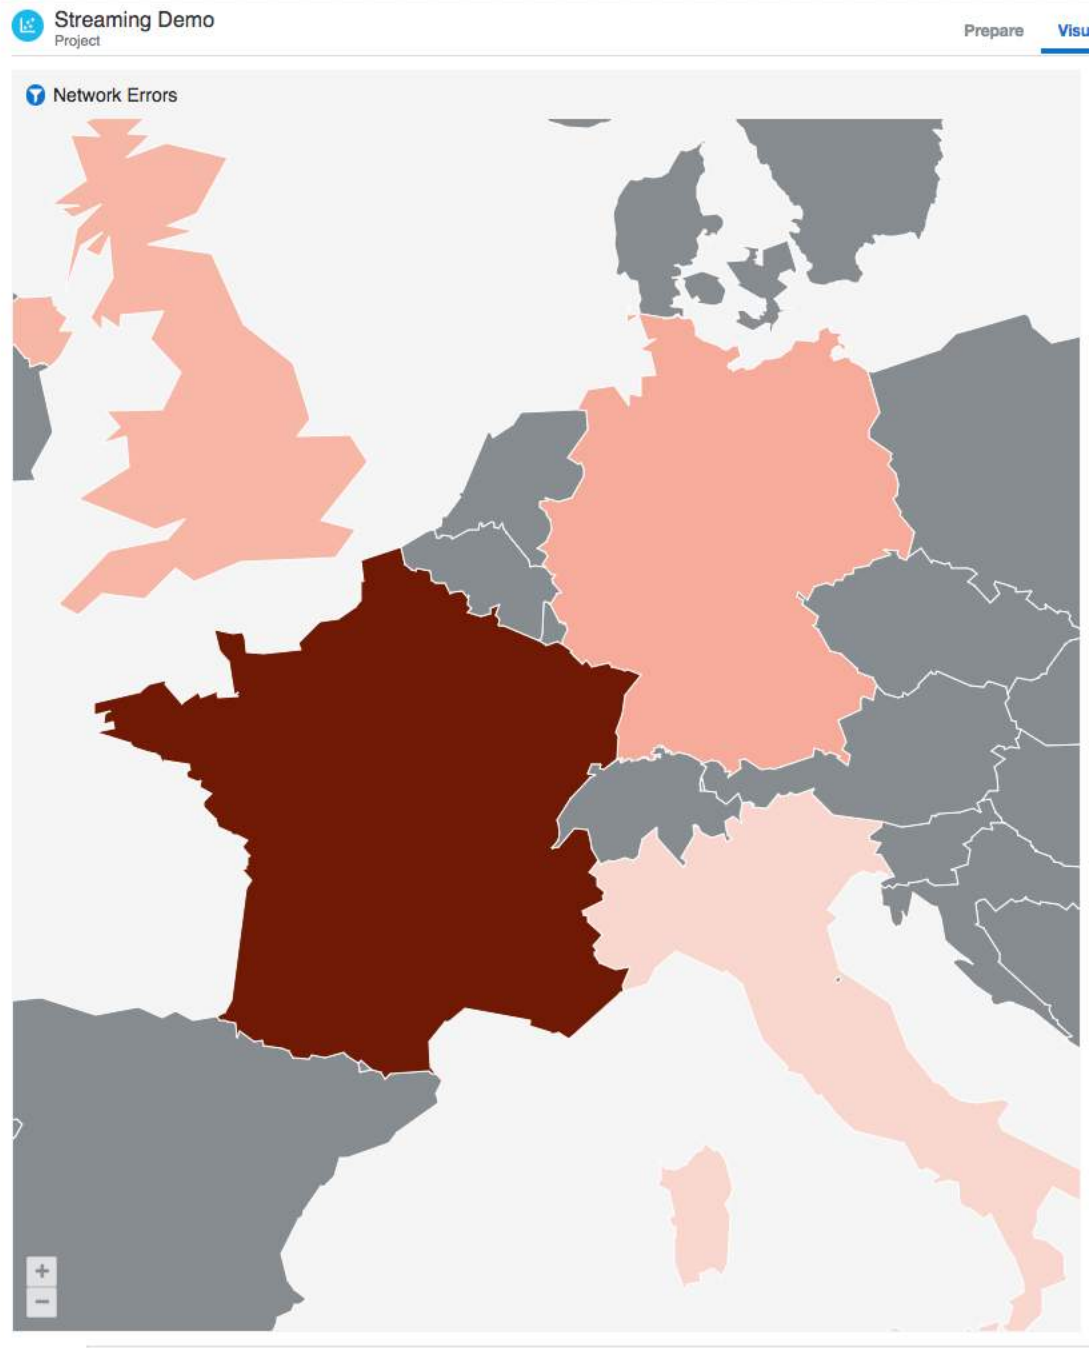
KAFKA\_DEMO\_TABLE

- Day
- Country
- Network Errors
- Hour
- Revenue Benchmark
- Revenue Stream

My Calculations

Value Labels

Add Calculation




- KAFKA\_DEMO\_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Network Errors

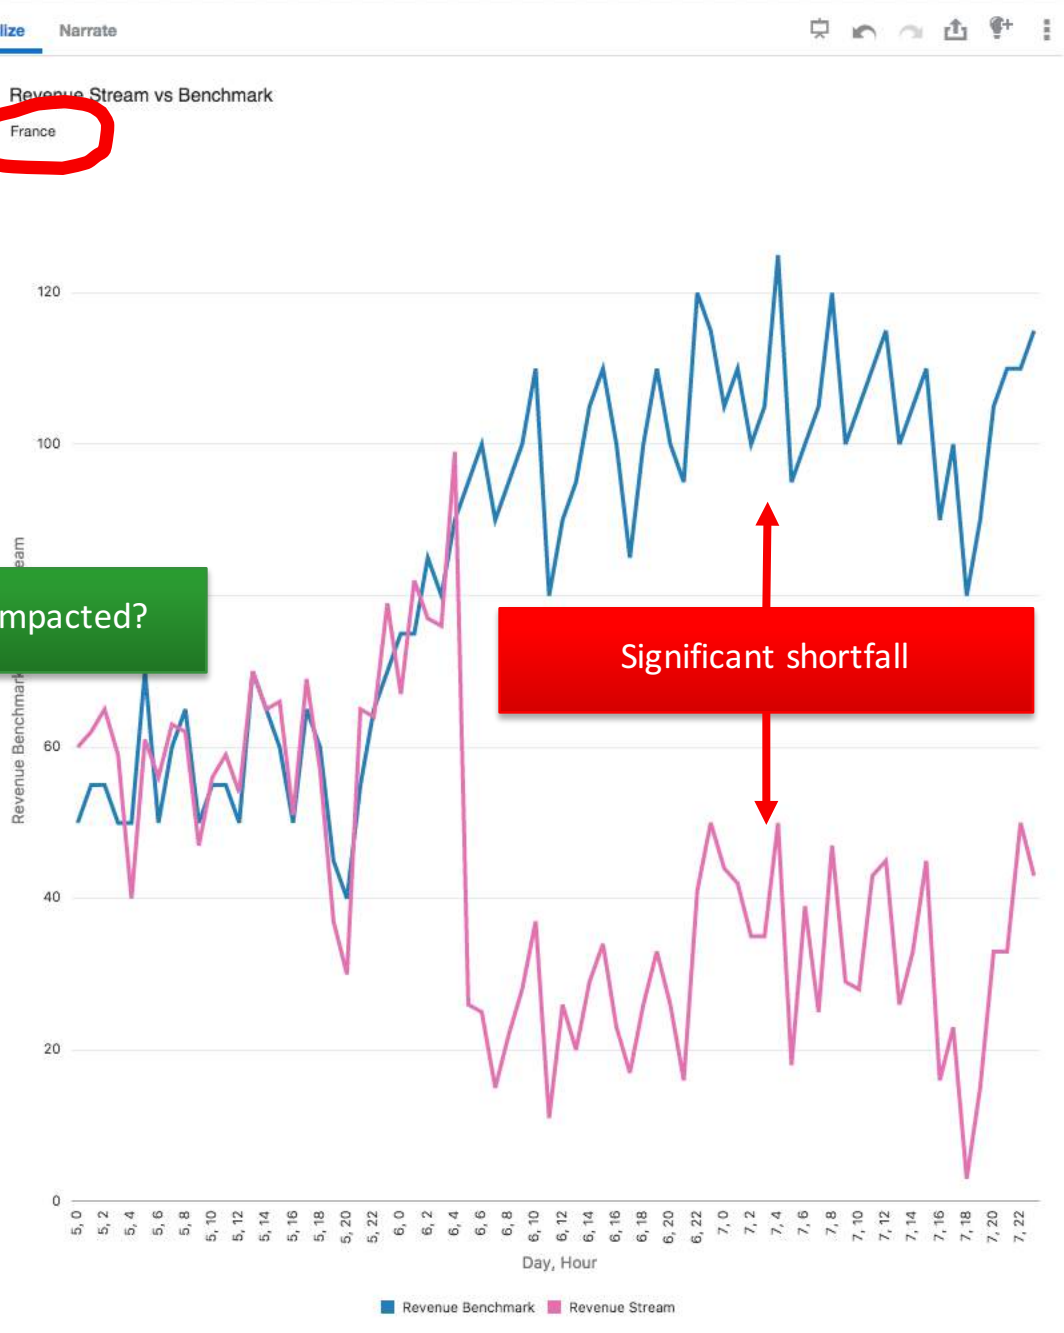
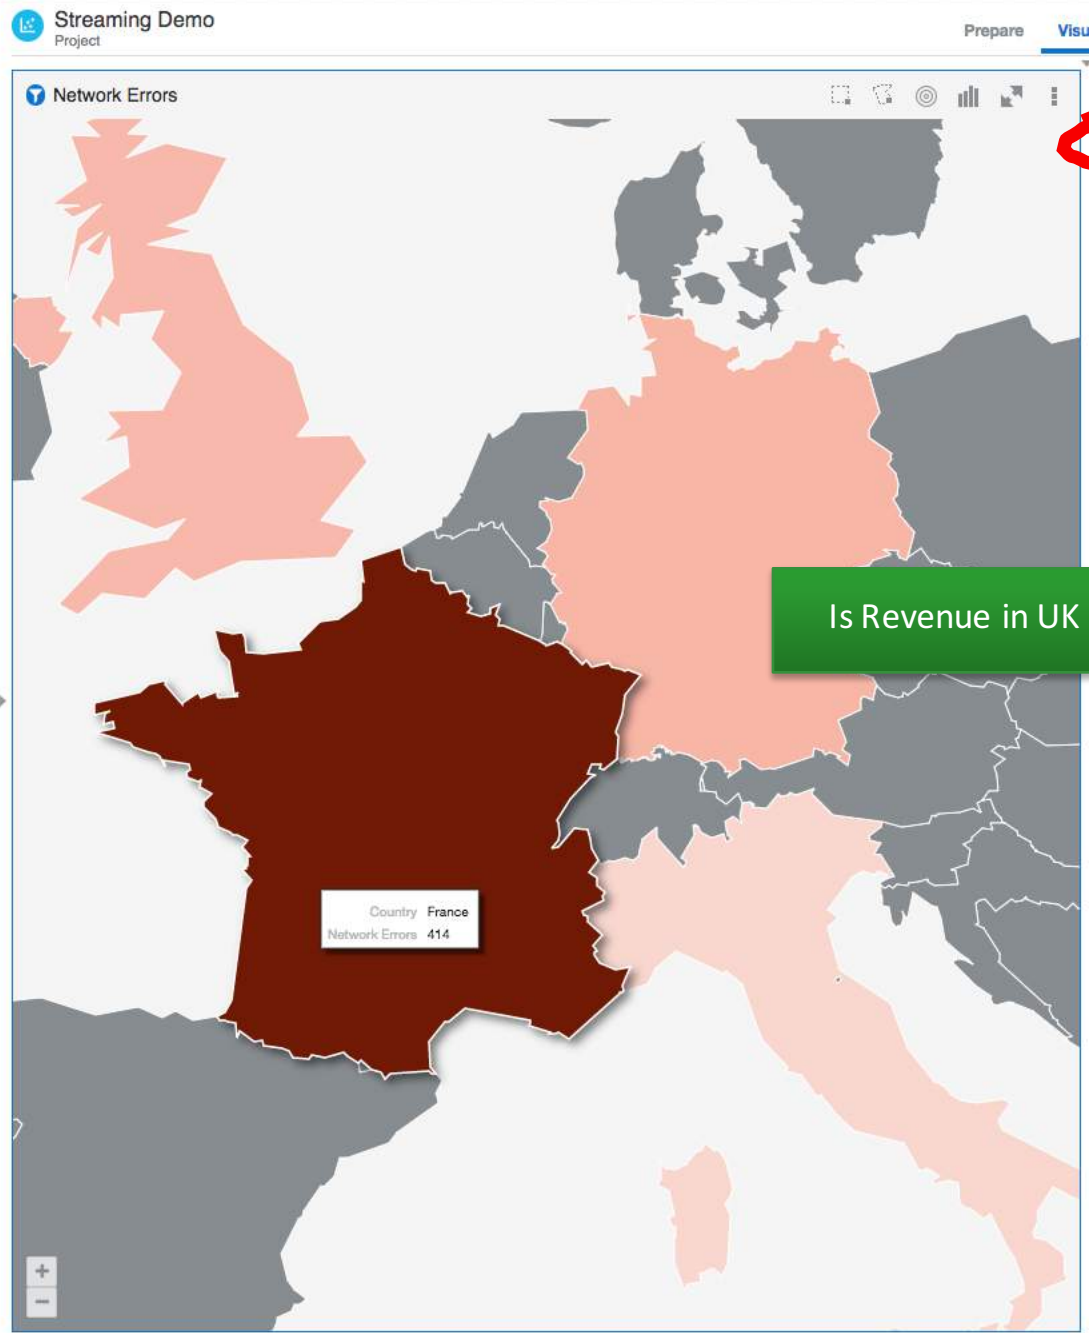


Revenue Stream vs Benchmark

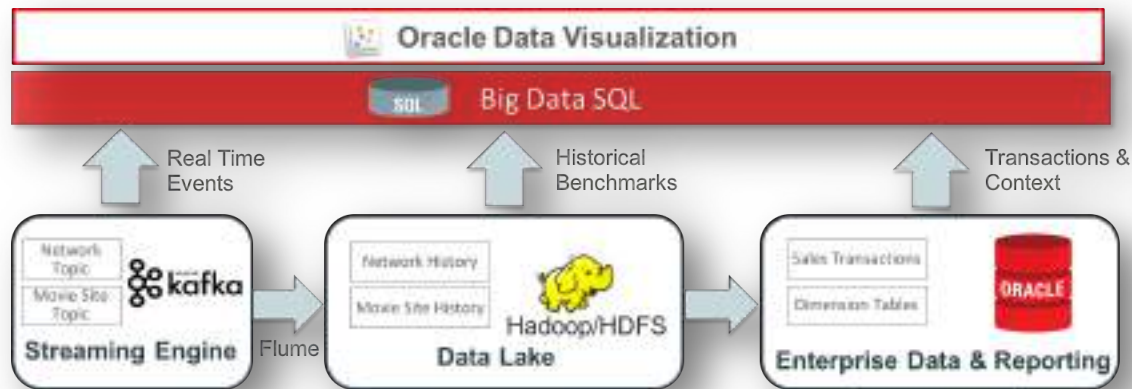


Search

- KAFKA\_DEMO\_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels



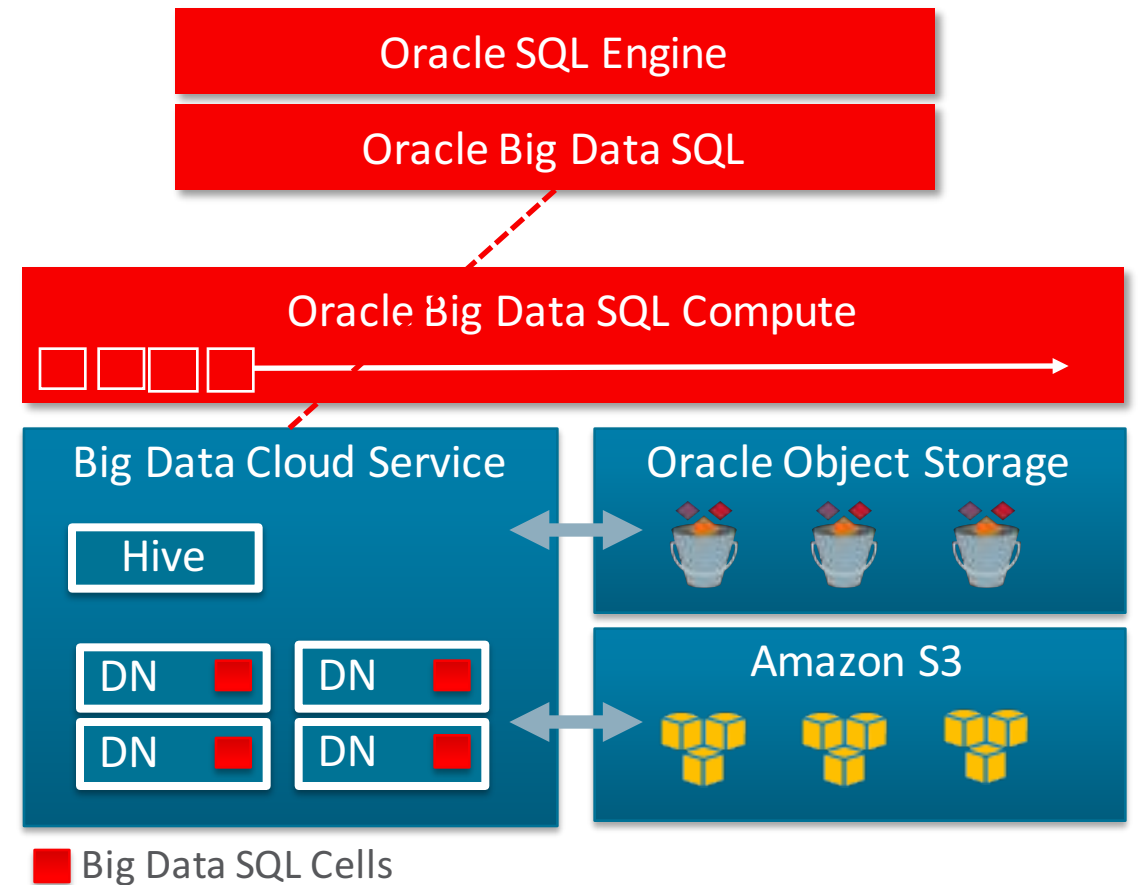
# Insights Achieved with Simplicity



- Easily blend real time streams with history, benchmarks and context
  - Are we running at peak performance?
  - What is the opportunity cost of our current network latency?
- Any application realizes benefit
  - Use Oracle SQL and APIs over all data
- Ensure data is secure
  - Leverage Oracle advanced security

# Efficiently Query Data Lake

- Data local processing for best performance
- Query Object Store directly without need for Hadoop Cluster
- Elastically scale based on workload
- Optimize ground to cloud



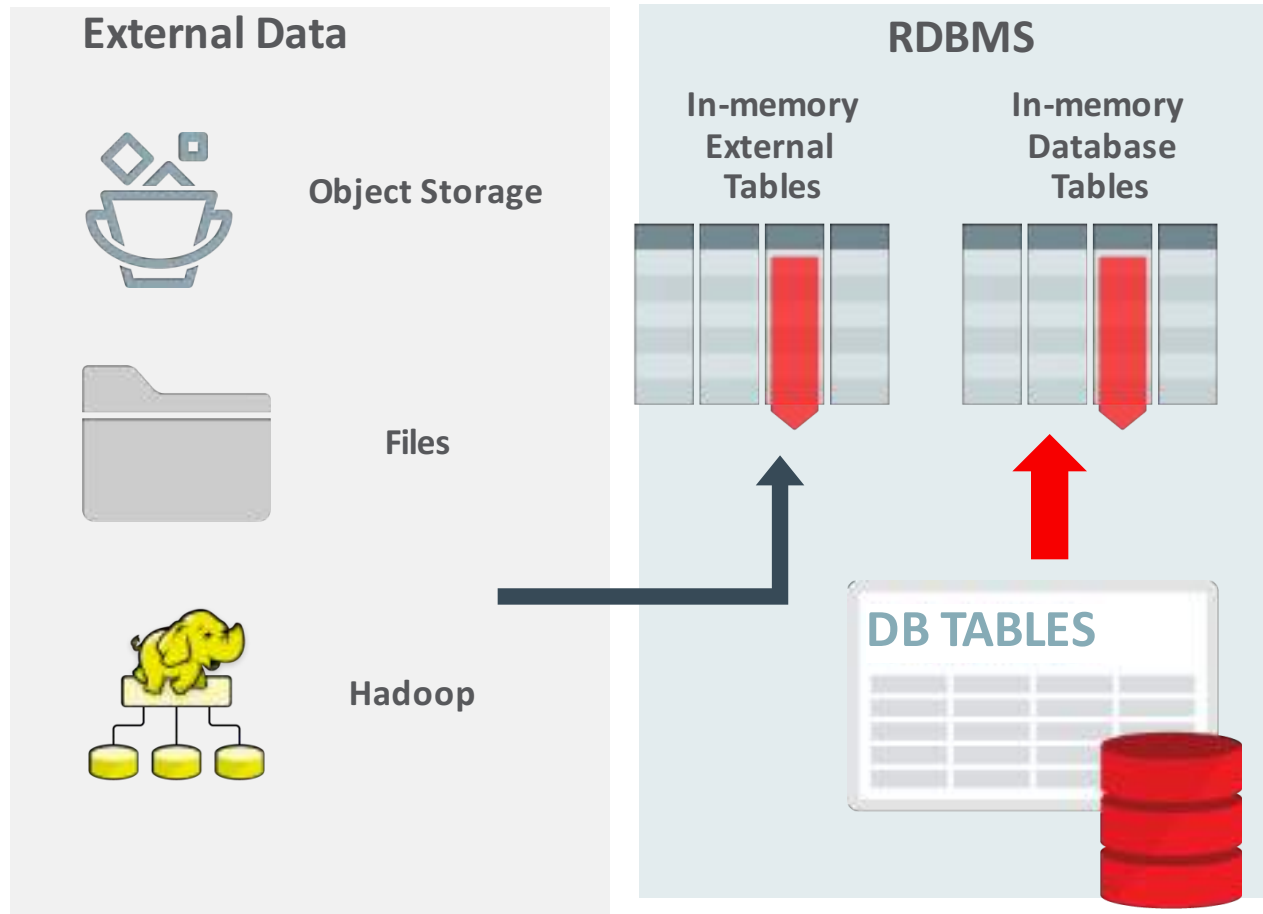
# Optimized Drivers for Common File Types

- Continue to push intelligent processing to cluster
  - Use open source when Oracle optimizations unavailable
- Examples:
  - Oracle Parquet Driver
  - C-based filter and parse of CLOBs (e.g. JSON documents)

Oracle Parquet Driver  
Up to **8x faster** than  
Hive Default Parquet Driver

# In-Memory and Columnar Processing

## Fast Analytics on Big Data Sources



- In-Memory For External Tables builds in-memory column cache of data outside the DB for ultra-fast analytics on external data
- **All In-Memory Optimizations** apply
  - Vector processing, JSON expressions extend transparently to external data
  - Enhanced distributed processing across Hadoop cluster
- Up to **100X** faster