
ORACLE®

Oracle Big Data SQL Deep Dive

Subtitle

Marty Gubar
Big Data SQL PM
Oracle Corporation

ORACLE OPEN WORLD

October 1–5, 2017
SAN FRANCISCO, CA

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.

Big Data SQL Goals

Easily access any data across big data stores

Provide a unified security model across the sources

Analyze all data using Oracle's rich SQL dialect

Fast performance using Big Data SQL Smart Scan

Flexible Deployment Options

Engineered Systems	Oracle Cloud	Commodity
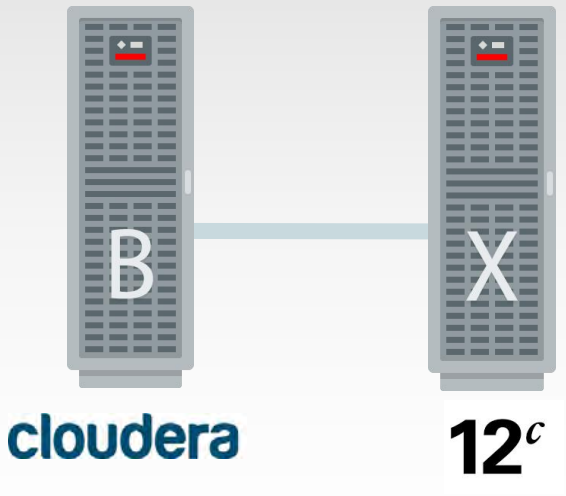	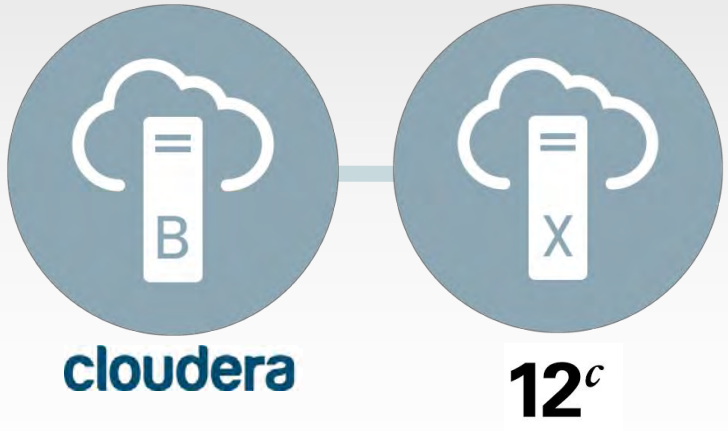	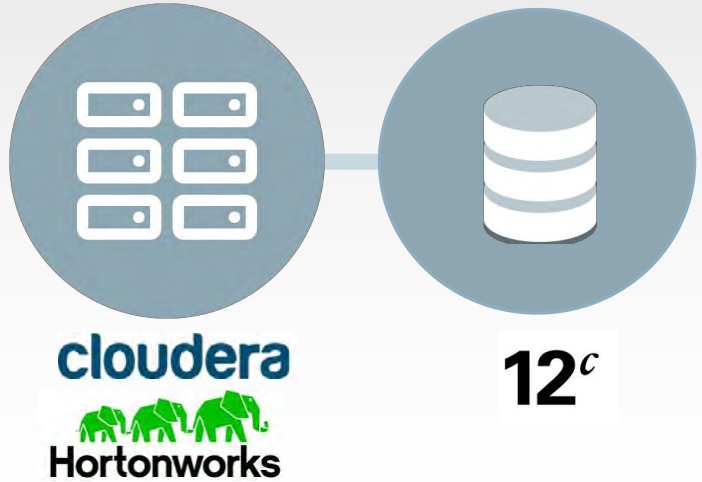

Big Data SQL Architecture

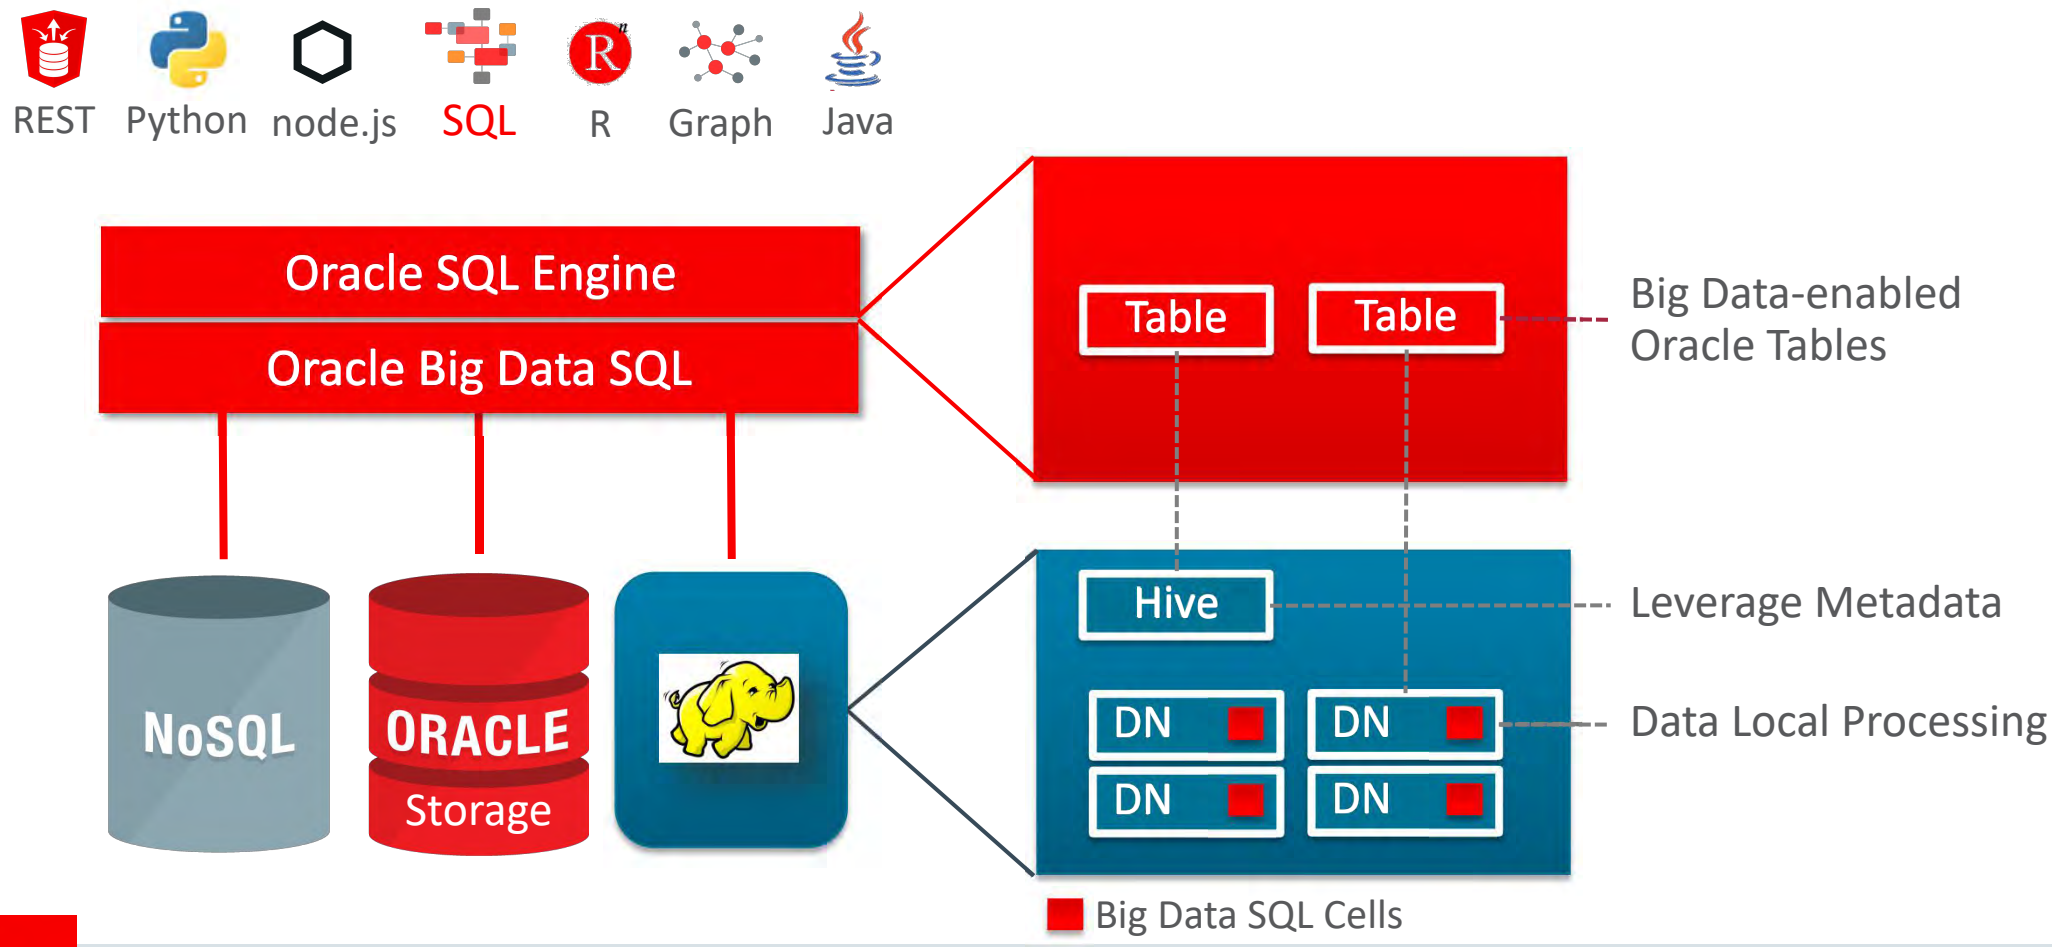

How Data is Stored in HDFS

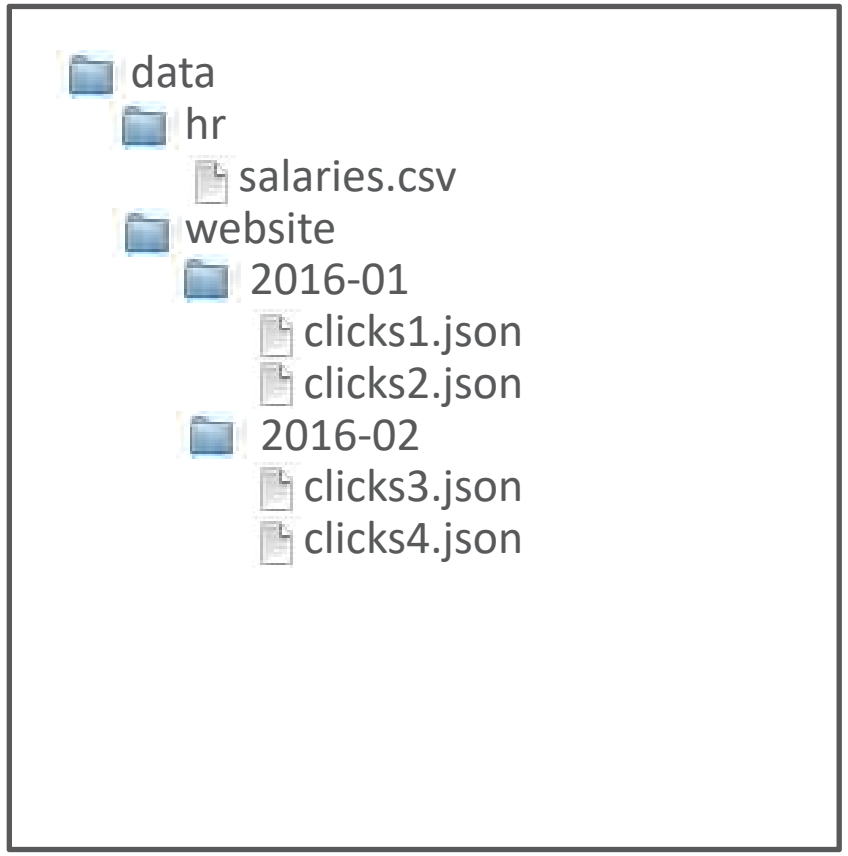

- Data stored in files and organized into folders
  - Can be any file type
  - Replicated 3x across cluster
- Schema on Read
  - The tool reading the data interprets the data as it sees fit

How Data is Stored in HDFS

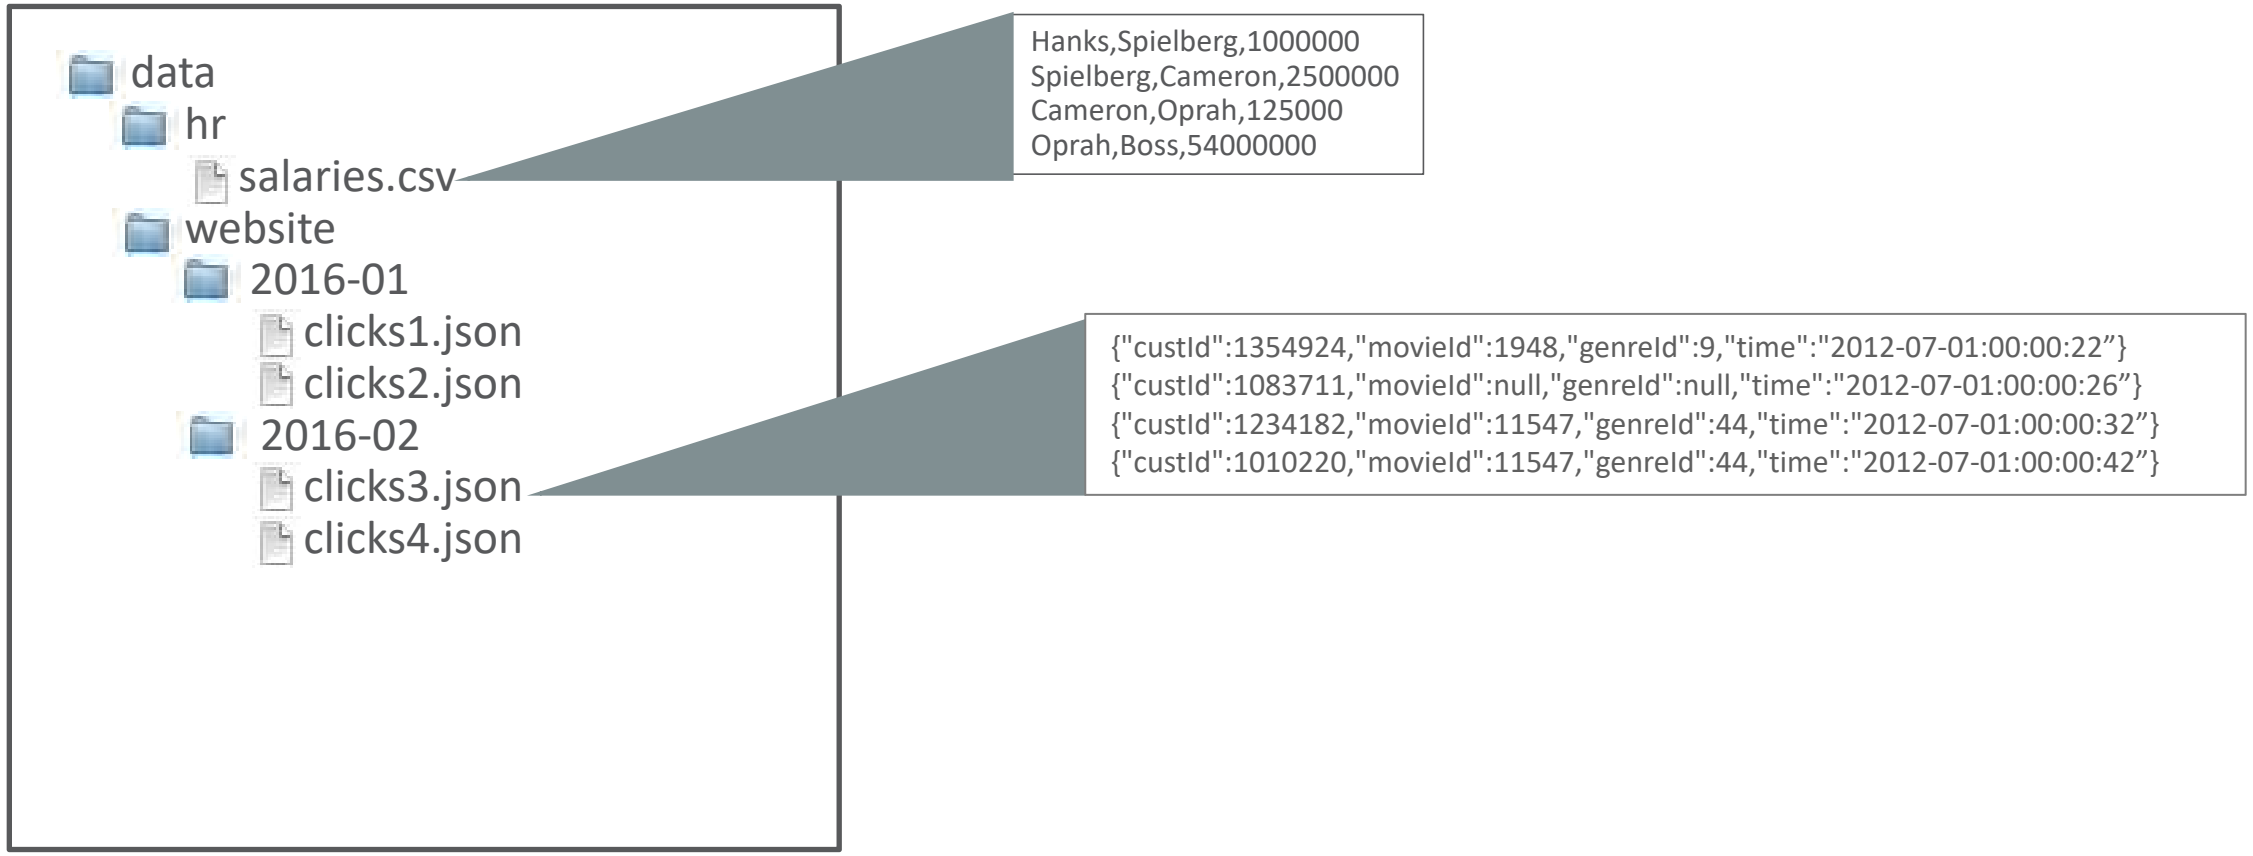

Demonstration: Accessing Data in HDFS

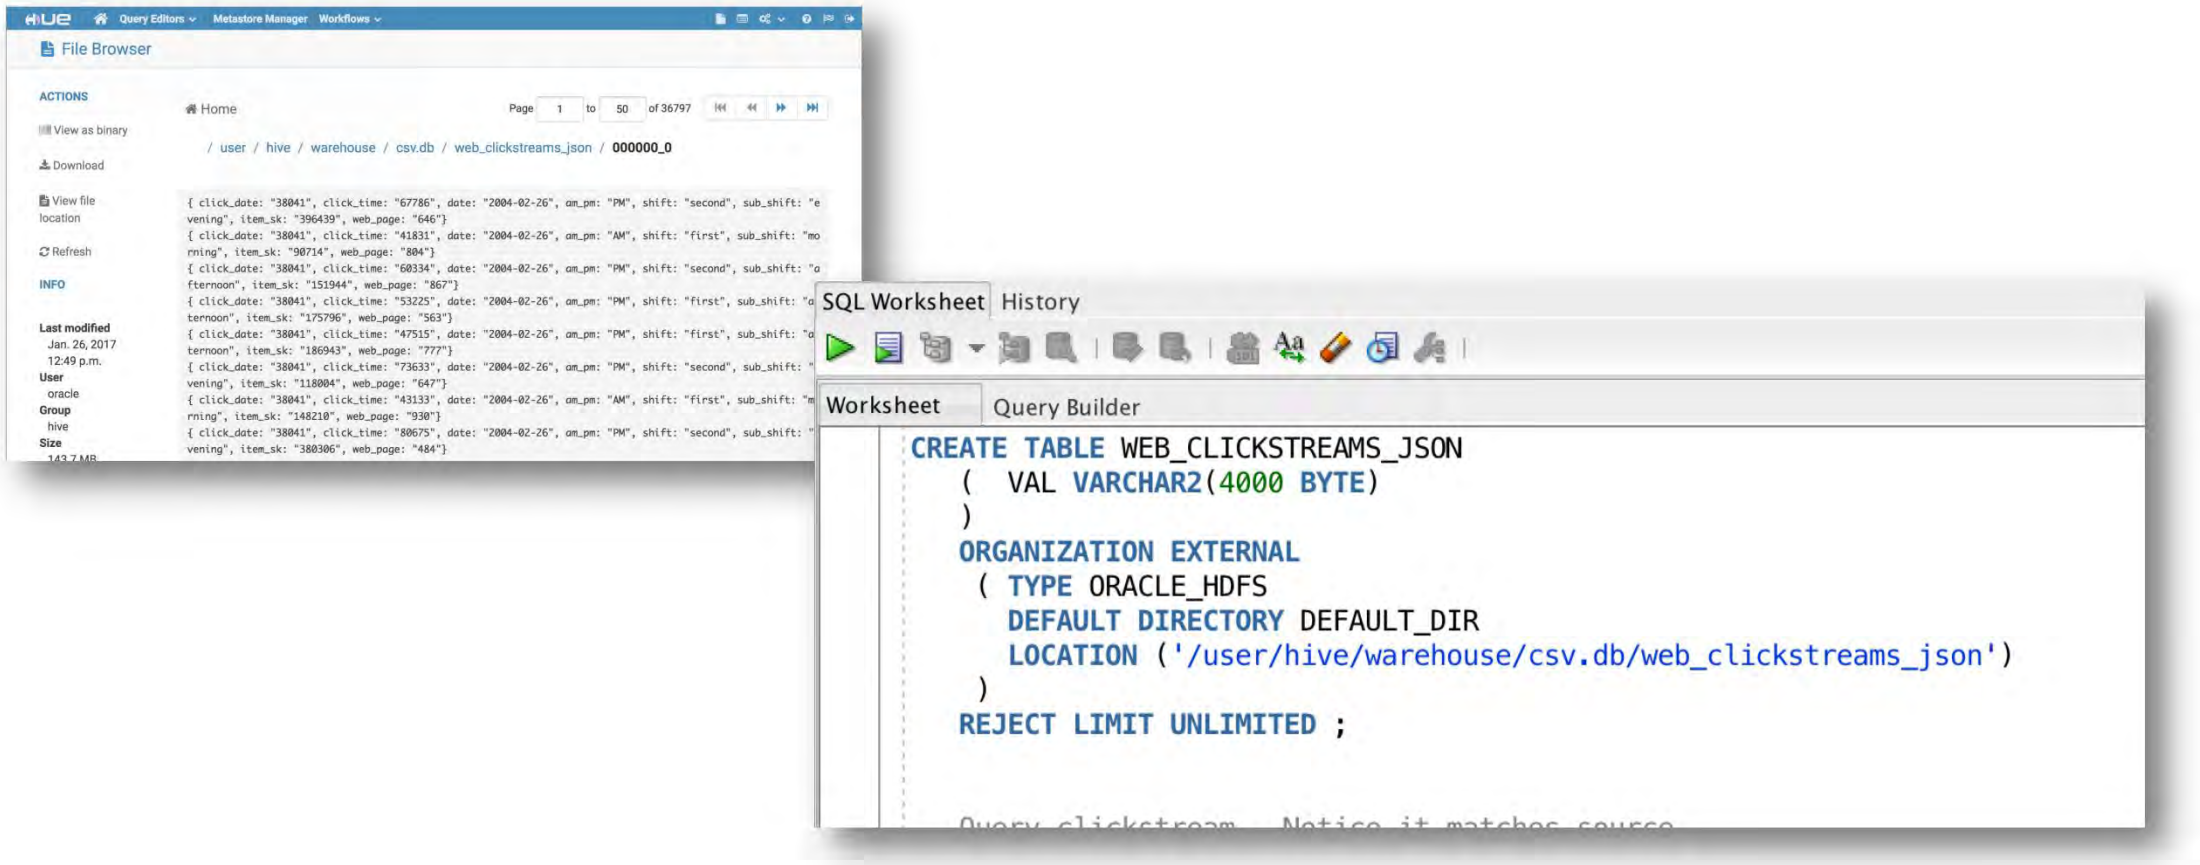

Organize and Describe Data with Hive

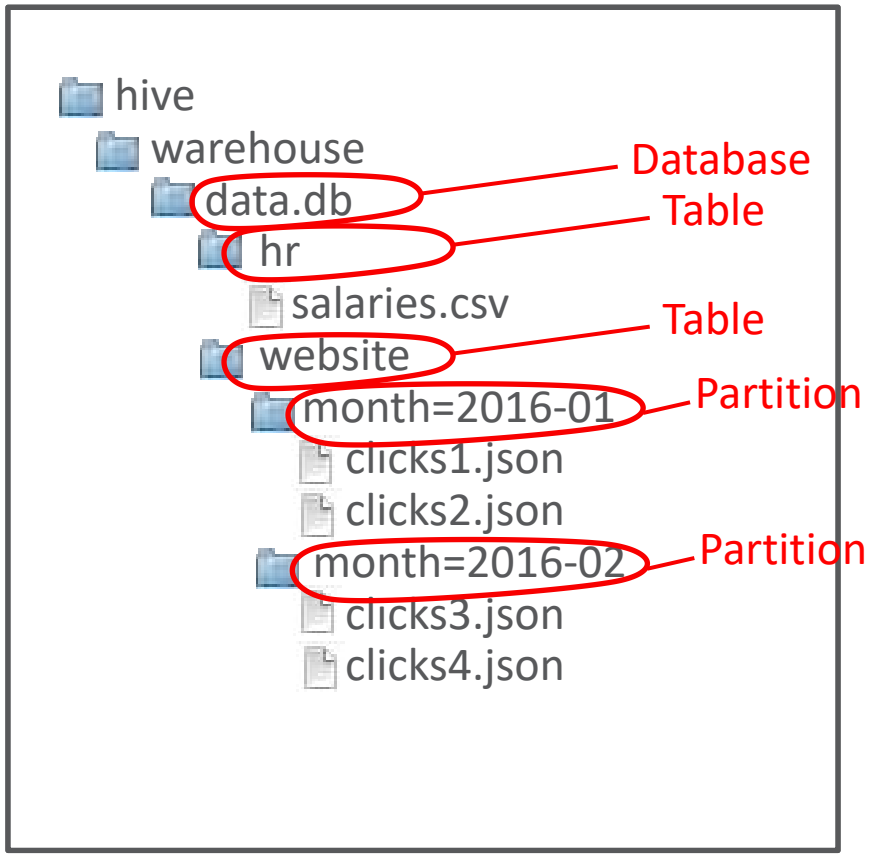

- Information is captured in **Hive Metastore**
- HDFS Folders become:
  - Databases
  - Tables
  - Partitions
- **Table** includes metadata for parsing files using Java classes
  - InputFormat defines chunks called splits based on file type
  - RecordReader creates rows out of splits
  - SerDe creates columns

Recommended Approach

Use ORACLE_HIVE When Possible

- Oracle Database query execution accesses Hive metadata at describe time
  - Changes to underlying Hive access parameters will not impact Oracle table (one exception… column list)
- Metadata an enabler for performance optimizations
  - Partition pruning and predicate pushdown into intelligent sources
- Utilize tooling for simplified table definitions
  - SQL Developer and DBMS_HADOOP packages

Demonstration: Big Data SQL Leverages Hive Metadata

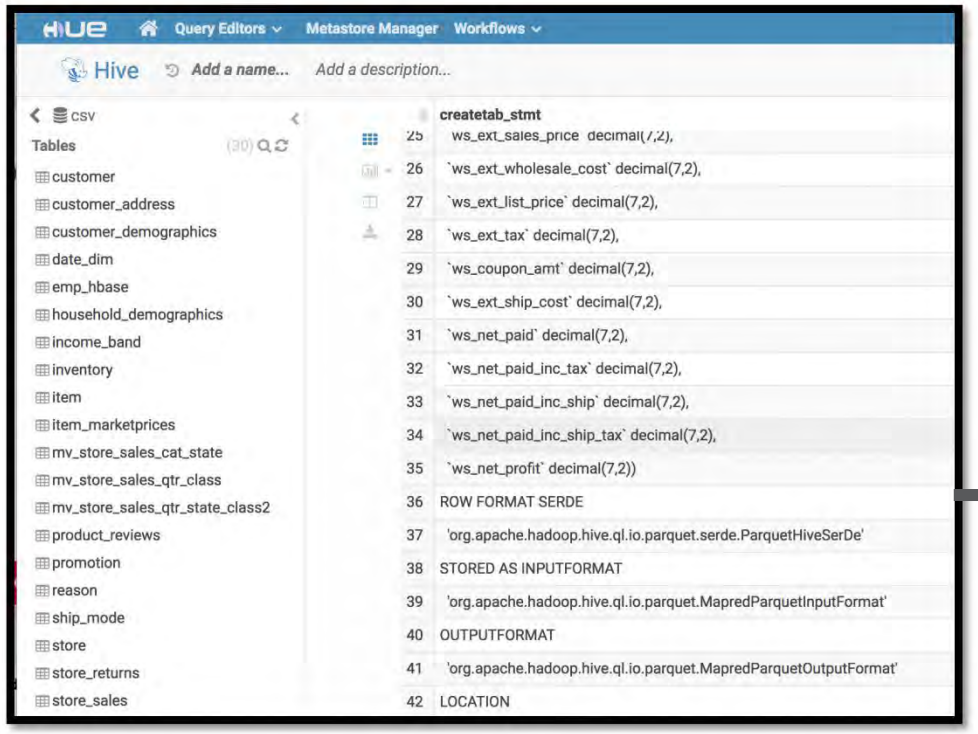

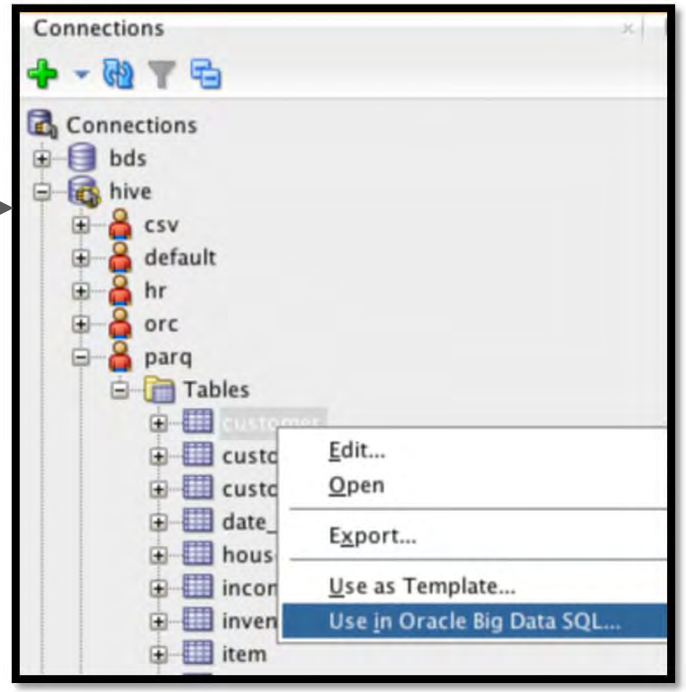

Big Data SQL Performance Features

Compound IO Reduction thru Smart Scans

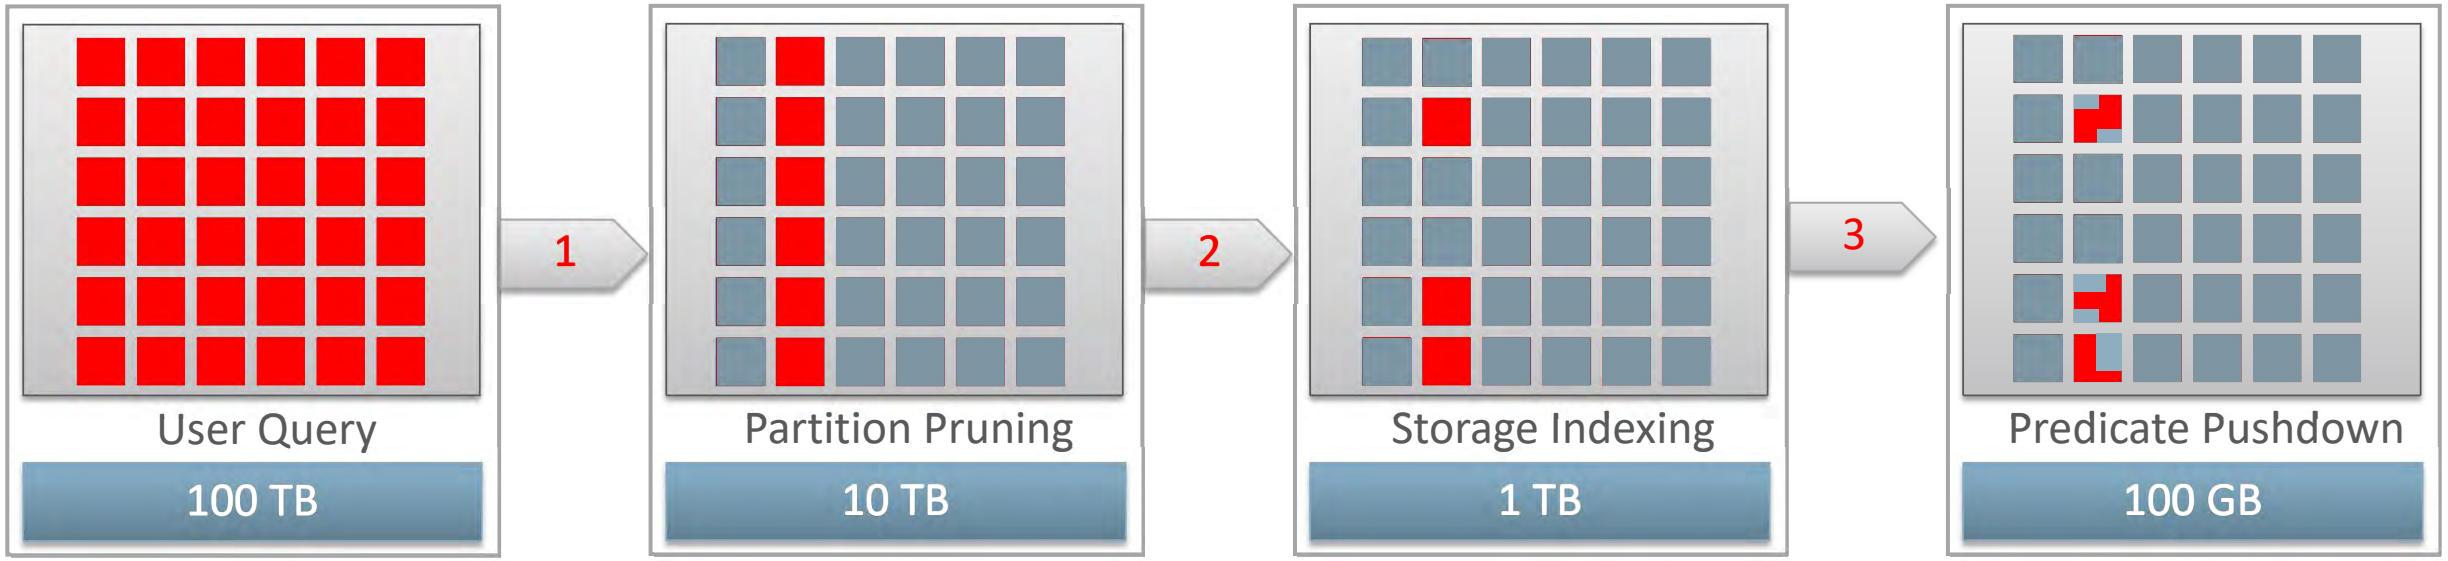

Big Data SQL Performance Features

Smart Scan – Execute Joins as Bloom Filters on Hadoop Nodes

Example: Total movie sales for customer segment

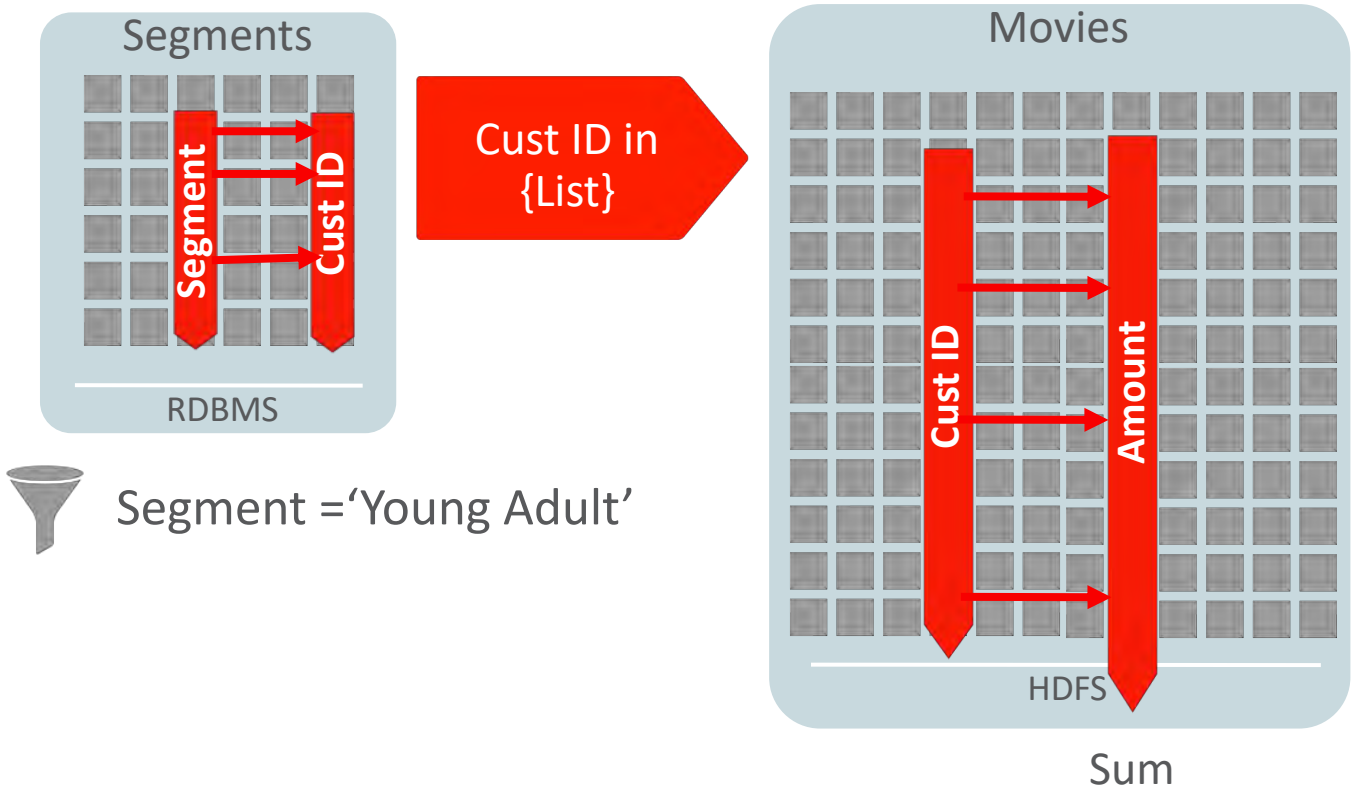

- Converts joins of data in multiple tables into scans
- Result:
  - Scans are pushed down to Hadoop nodes and executed locally
  - No data moved to Database to process joins
  - Massive speed up of query
- Works with data spanning DB and Hadoop as well as data in two Hadoop data sets

Big Data SQL Storage Index

Example: Find revenue for movies in a category **9** (Comedy)

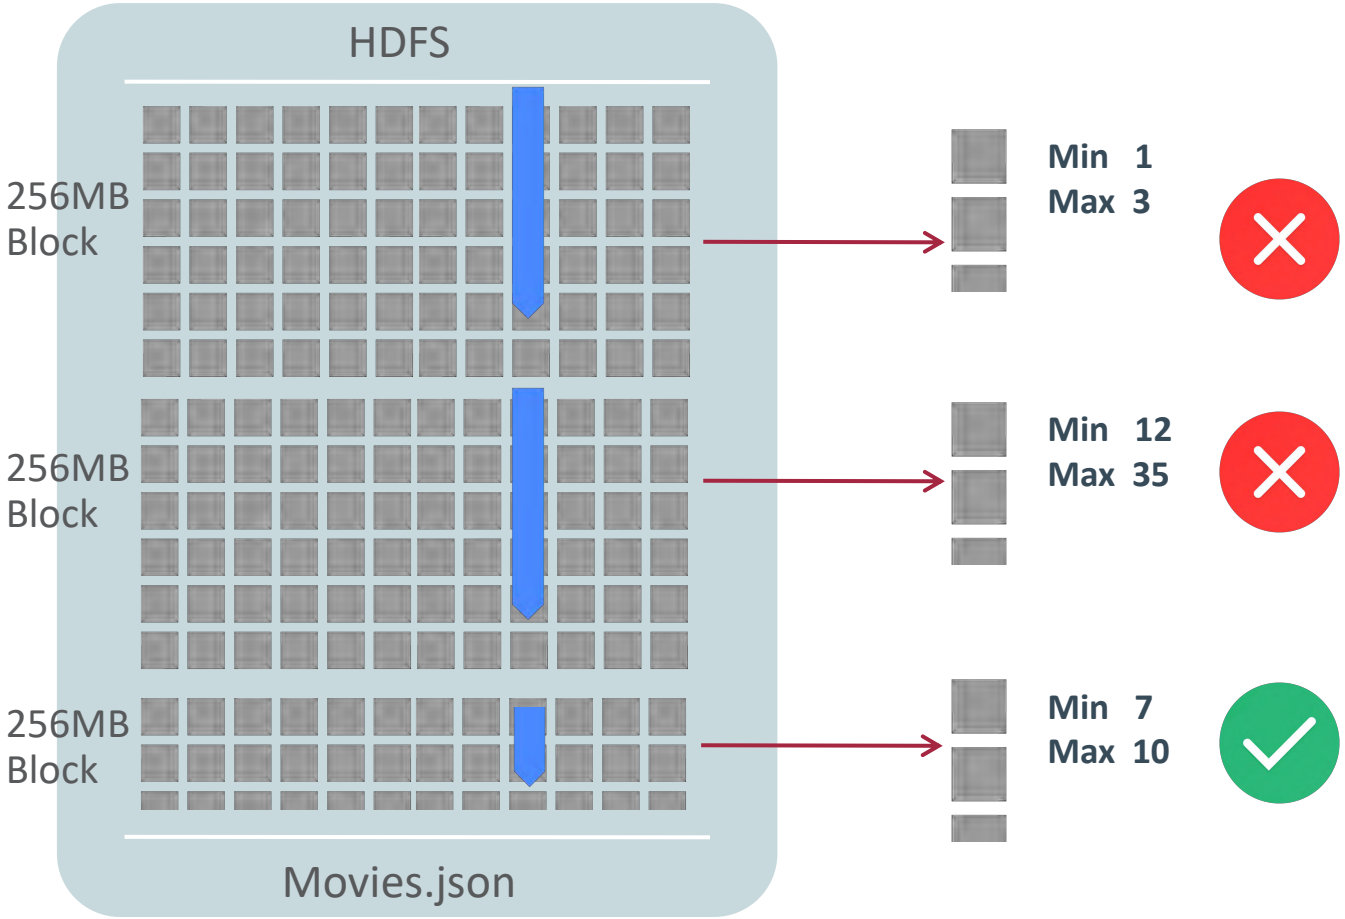

- Storage index provides query speed-up through transparent IO elimination of HDFS Blocks. It's a *negative index*
- Min / max value is recorded for columns included in a storage index (max # of colums = 32)
- Storage index provides partition pruning like performance for un-modeled data sets

Parquet: Columnar Database File in HDFS

Big Data SQL Pushes Predicates to Parquet

Schema on Write

Parquet implements a database storage structure with metadata and parsed data elements

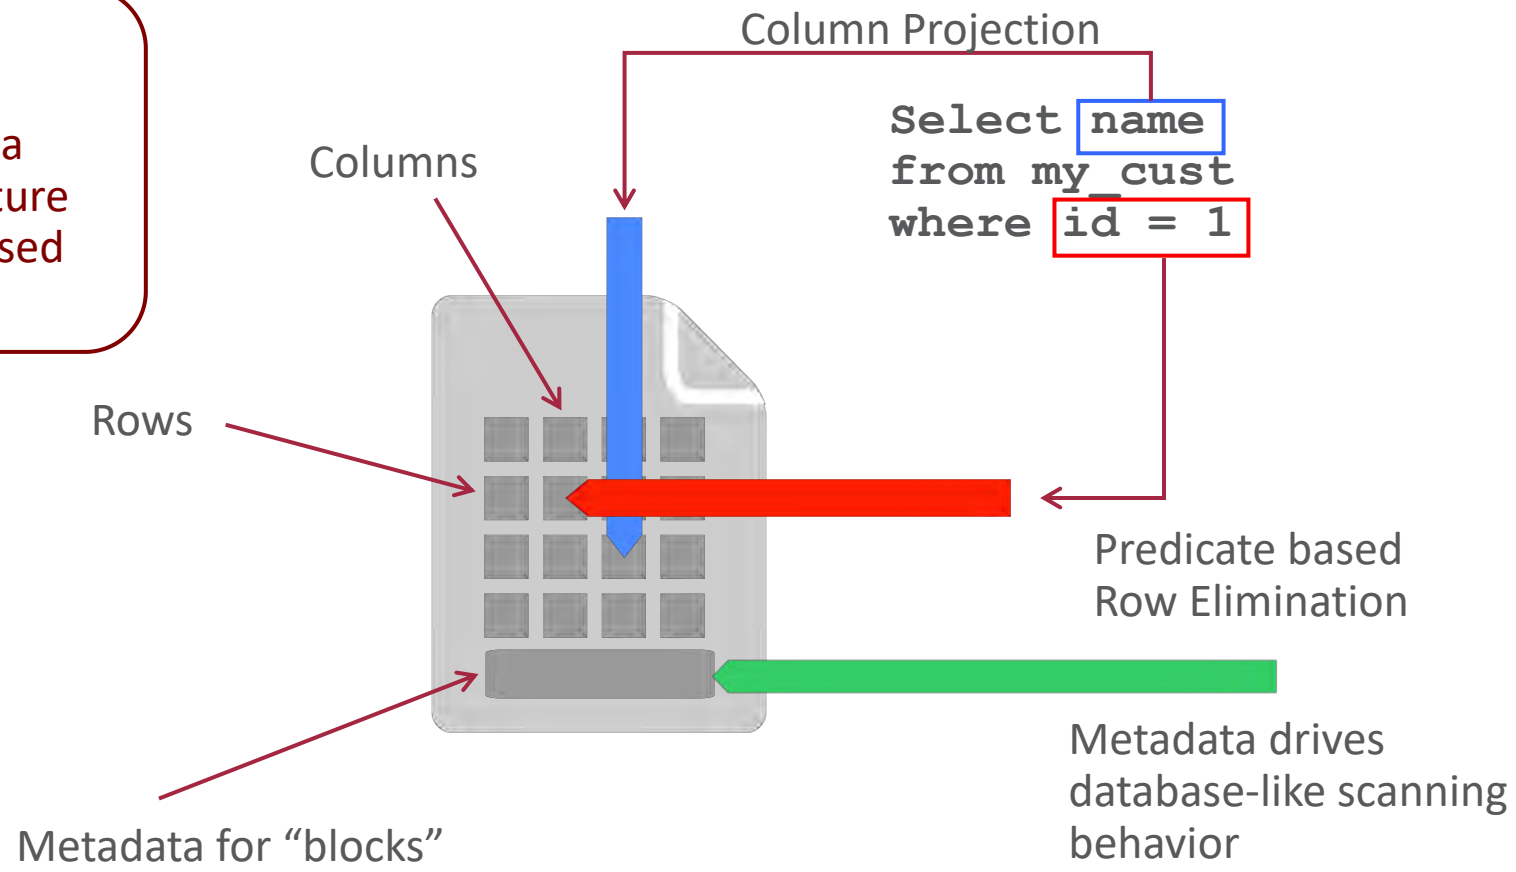

Enhance Performance with Automatic Query Rewrite

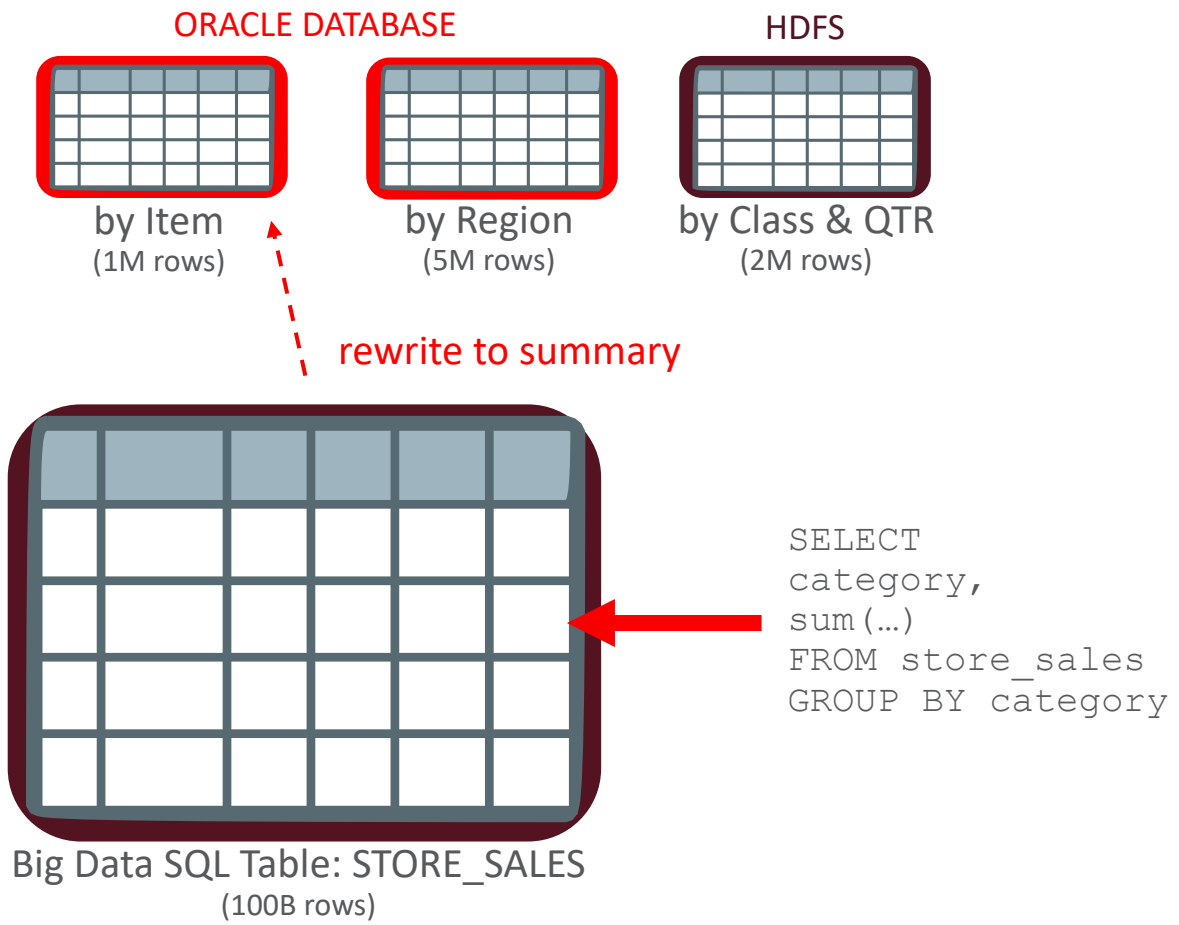

- **Orders of magnitude** performance improvement
- **Materialized view** query rewrite automatically redirects detail query to appropriate summary data
  - Store summaries in Oracle Database
  - If available, use existing summaries in HDFS
- No changes to query required

Performance Features

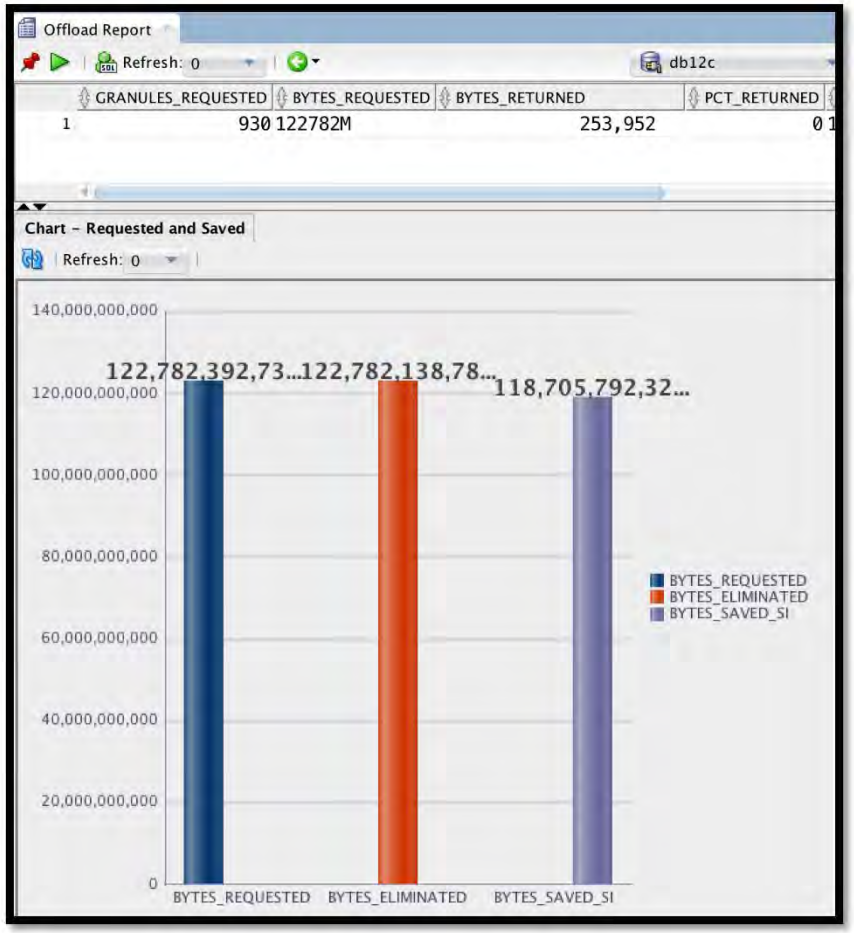

Demonstration Scenario

Analyzing Real-time Streams

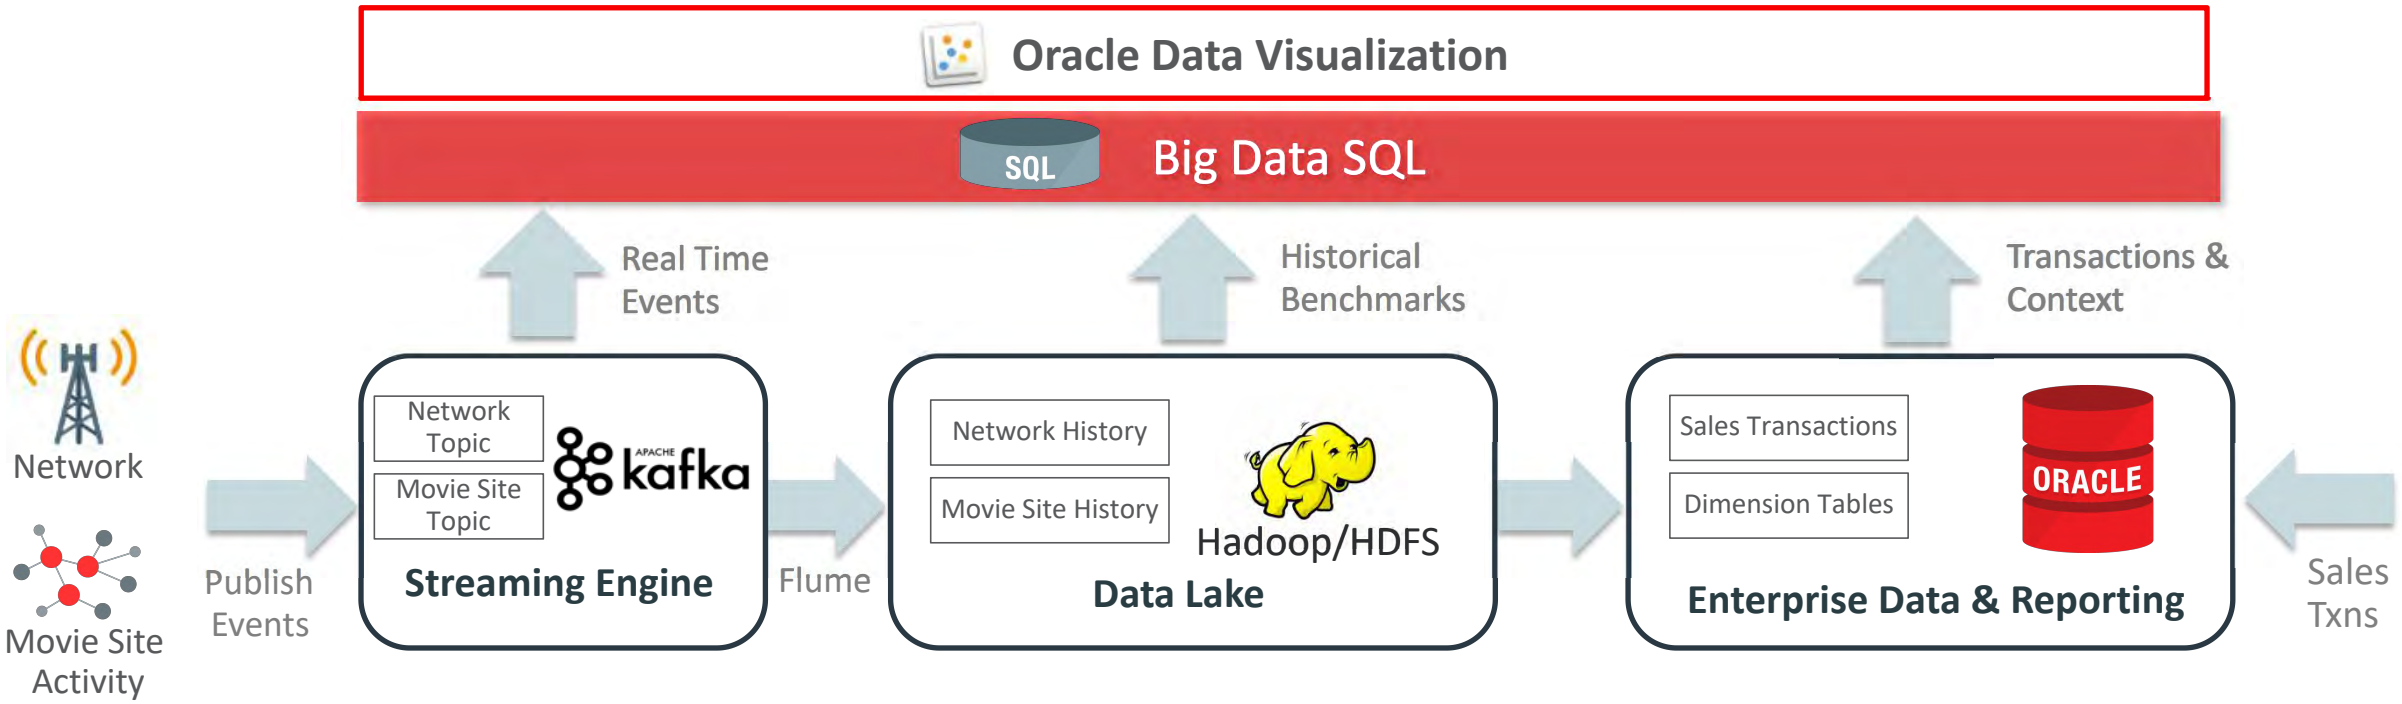

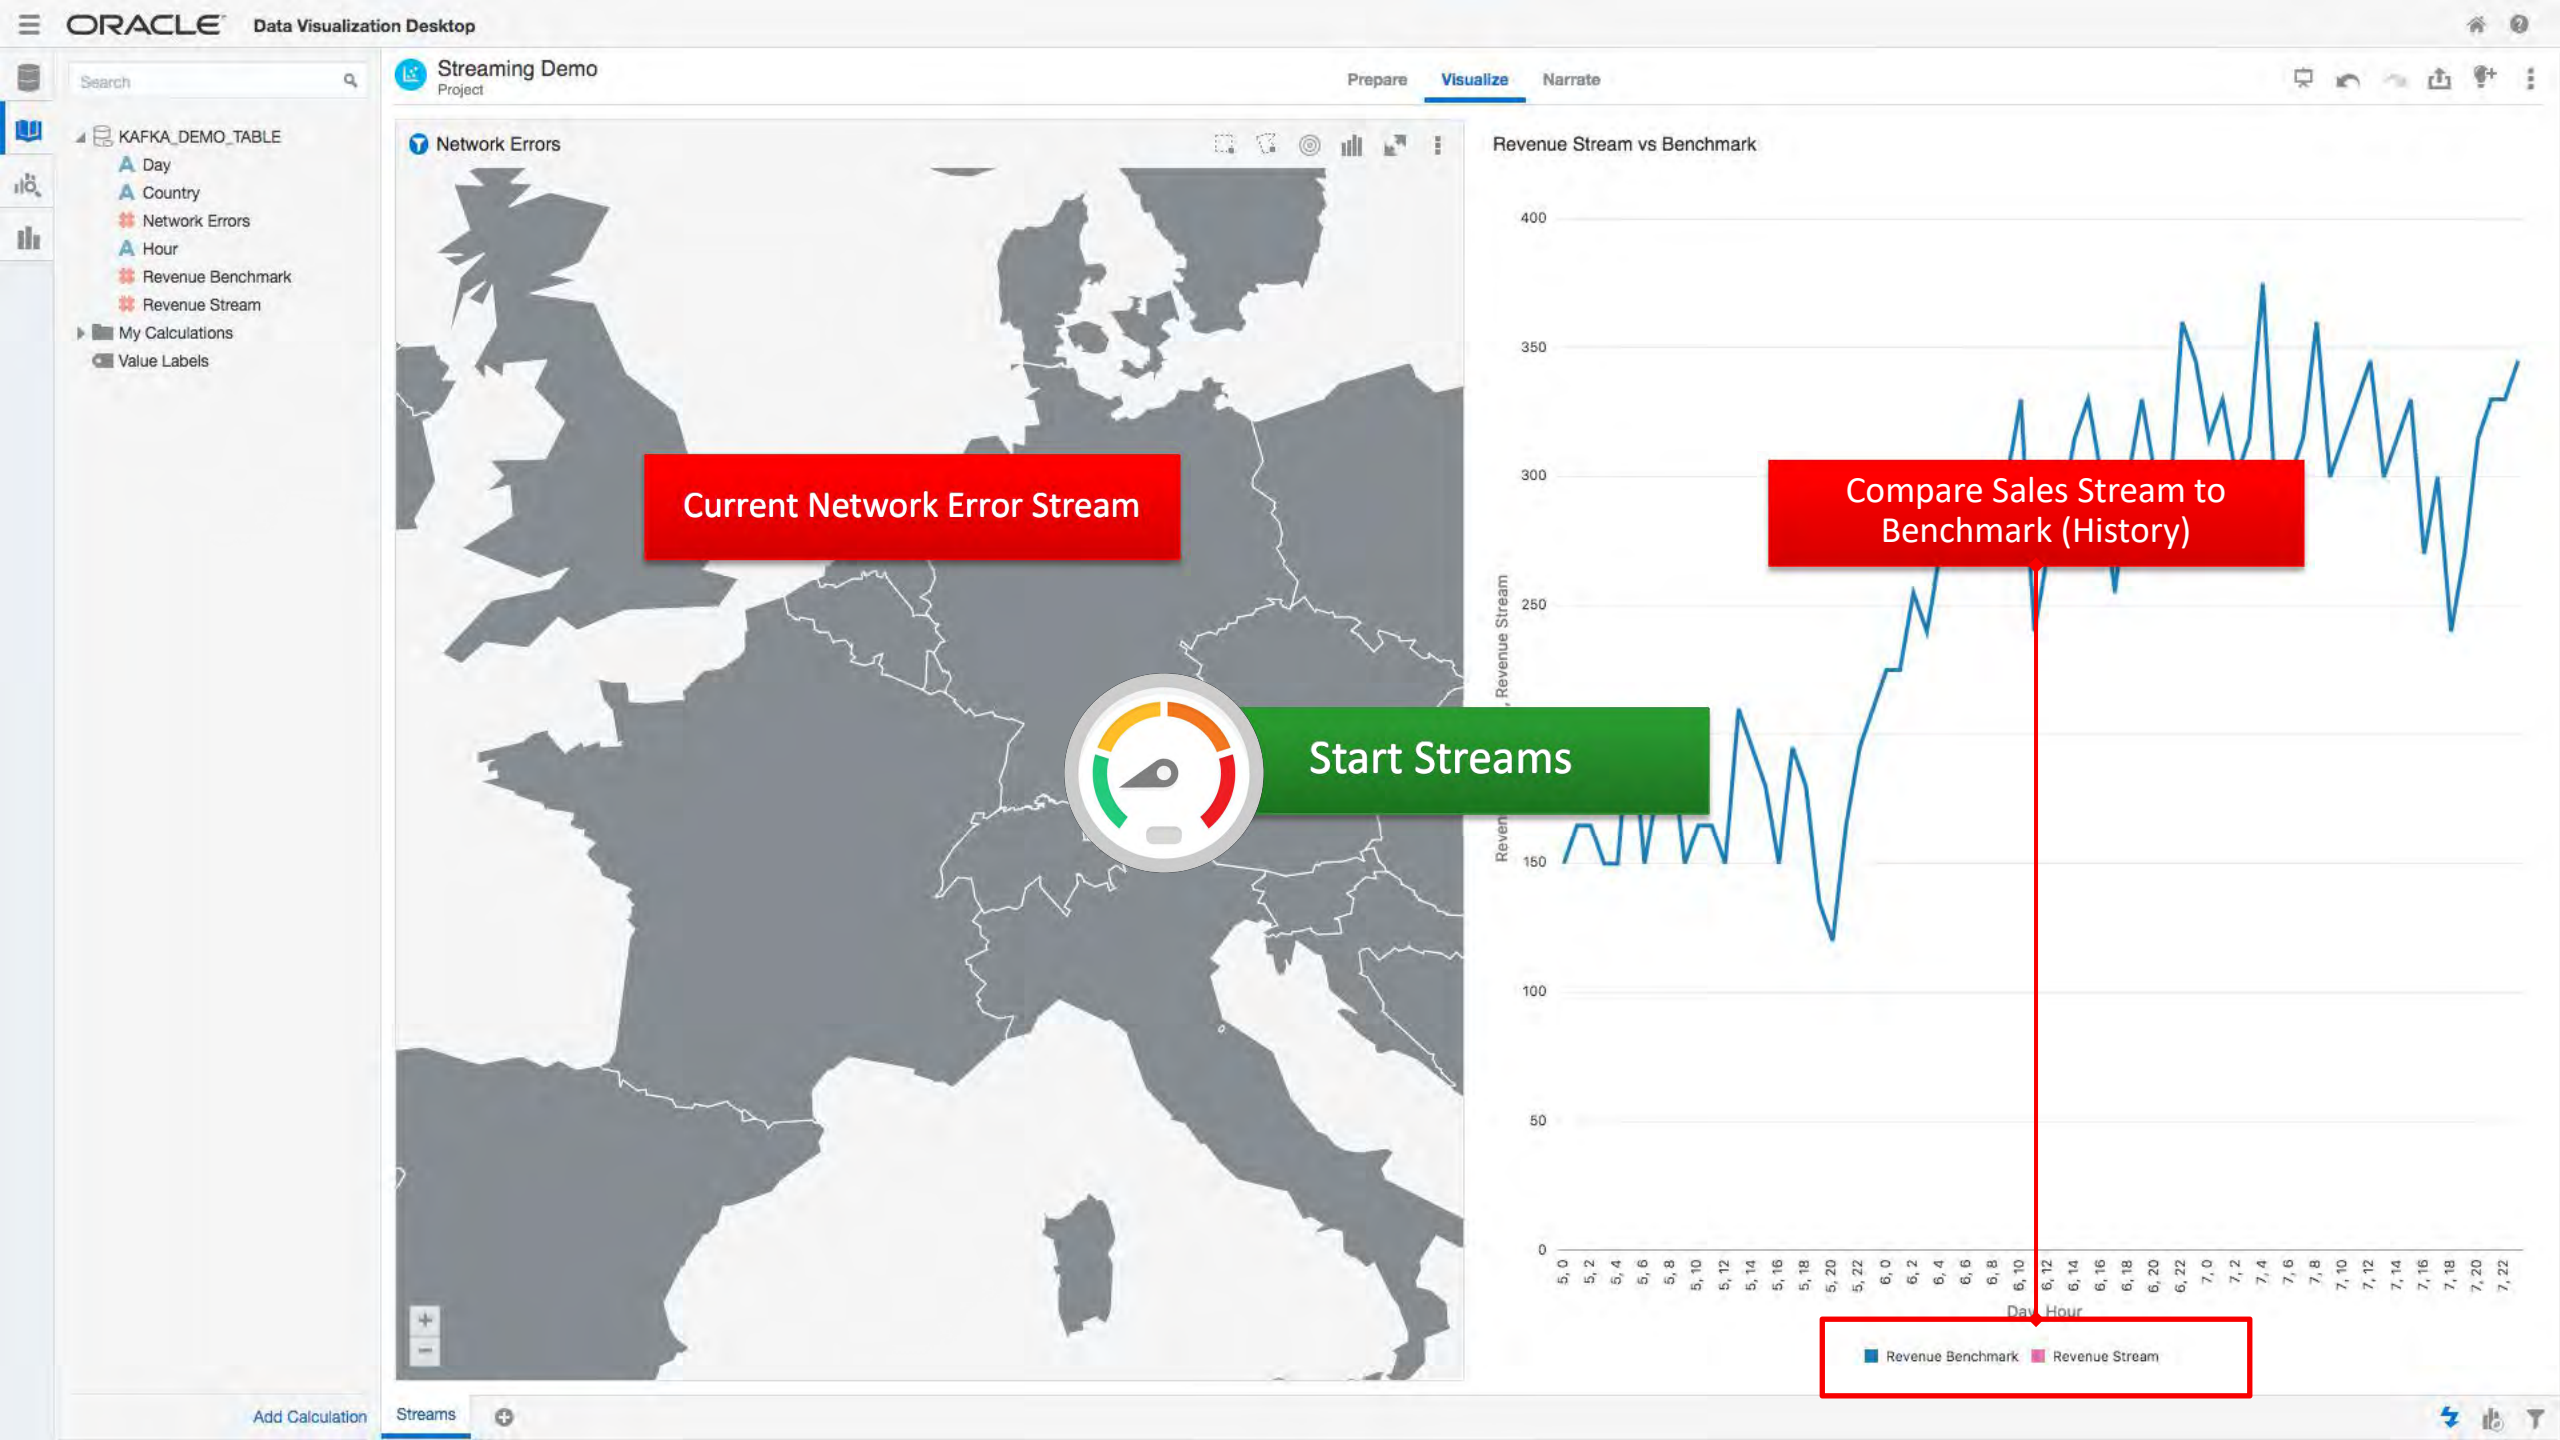
ORACLE Data Visualization Desktop
Search
KAFKA_DEMO_TABLE
Day
Country
Network Errors
Hour
Revenue Benchmark
Revenue Stream
My Calculations
Value Labels
Add Calculation
Streaming Demo
Project
Prepare
Visualize
Narrate
Network Errors
Current Network Error Stream
Start Streams
Revenue Stream vs Benchmark
Compare Sales Stream to Benchmark (History)
Revenue Stream
Day, Hour
Revenue Benchmark
Revenue Stream
Streams

Search

- KAFKA_DEMO_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Add Calculation

Streaming Demo
Project

Prepare | Visualize | Narrate

Network Errors

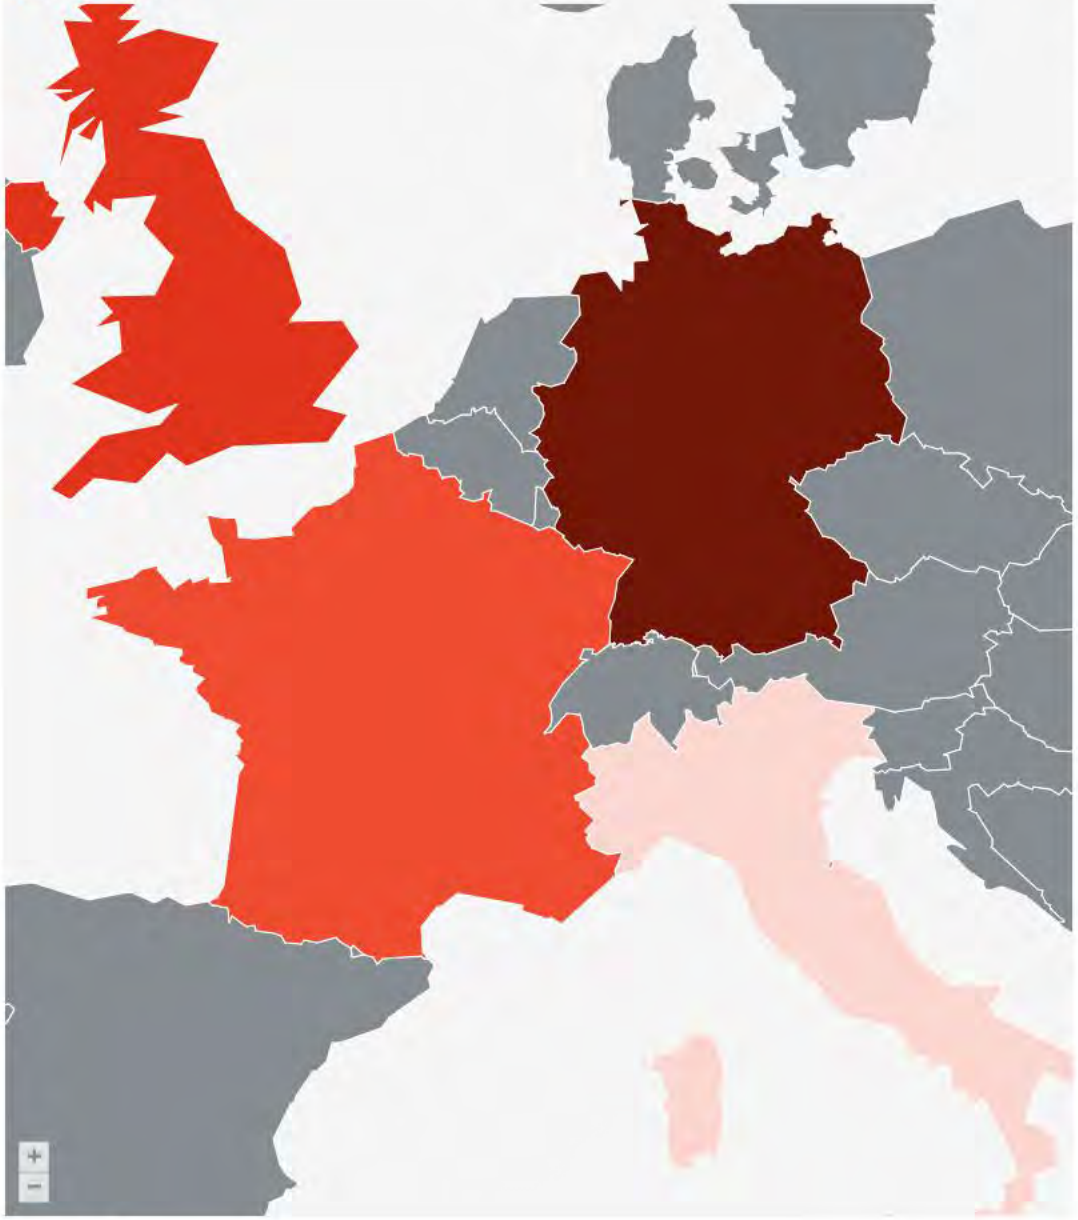

Revenue Stream vs Benchmark

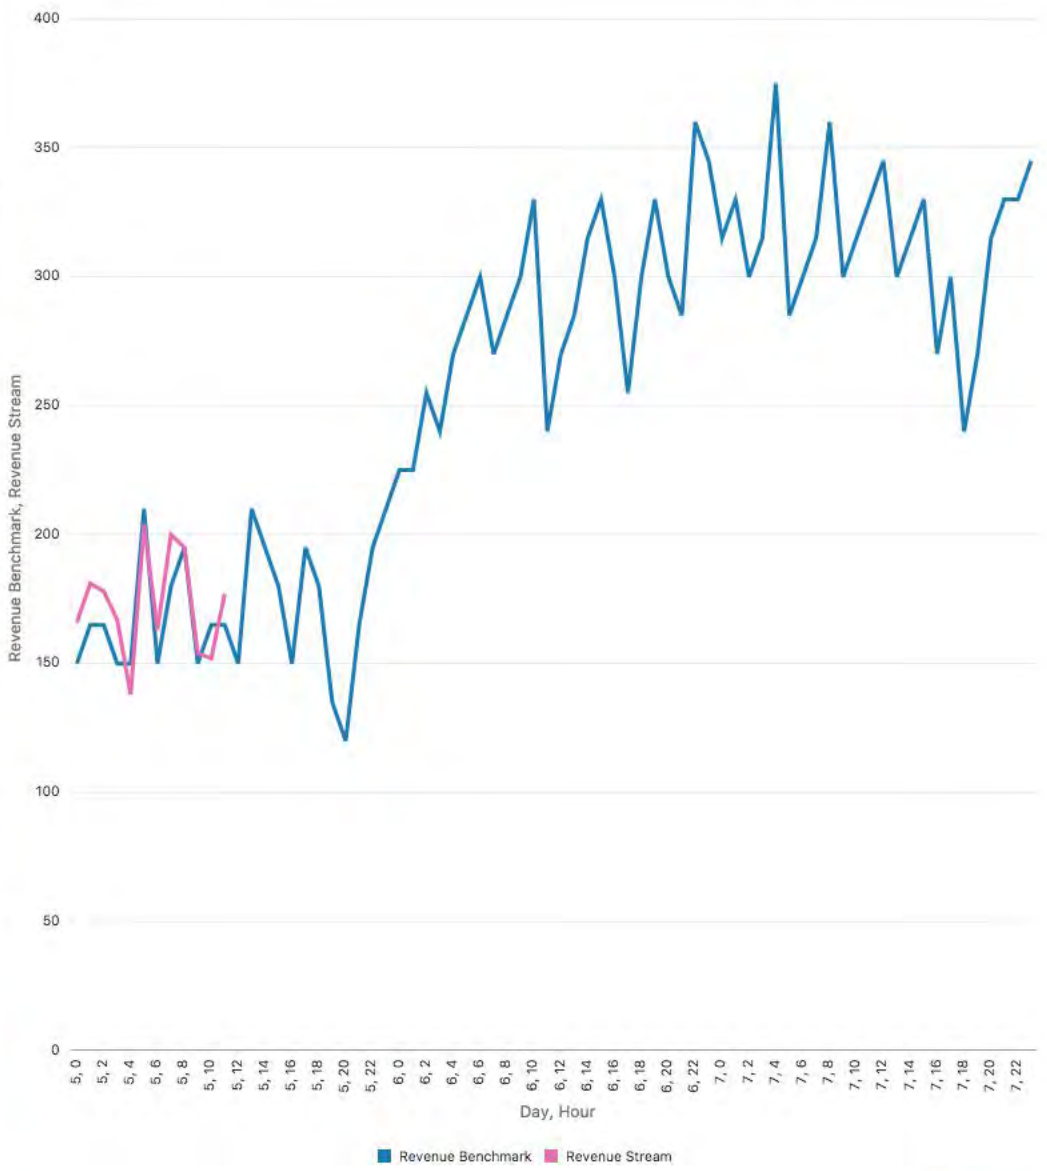

Streams

Search

- KAFKA_DEMO_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Add Calculation

Streaming Demo
Project

Prepare | Visualize | Narrate

Network Errors

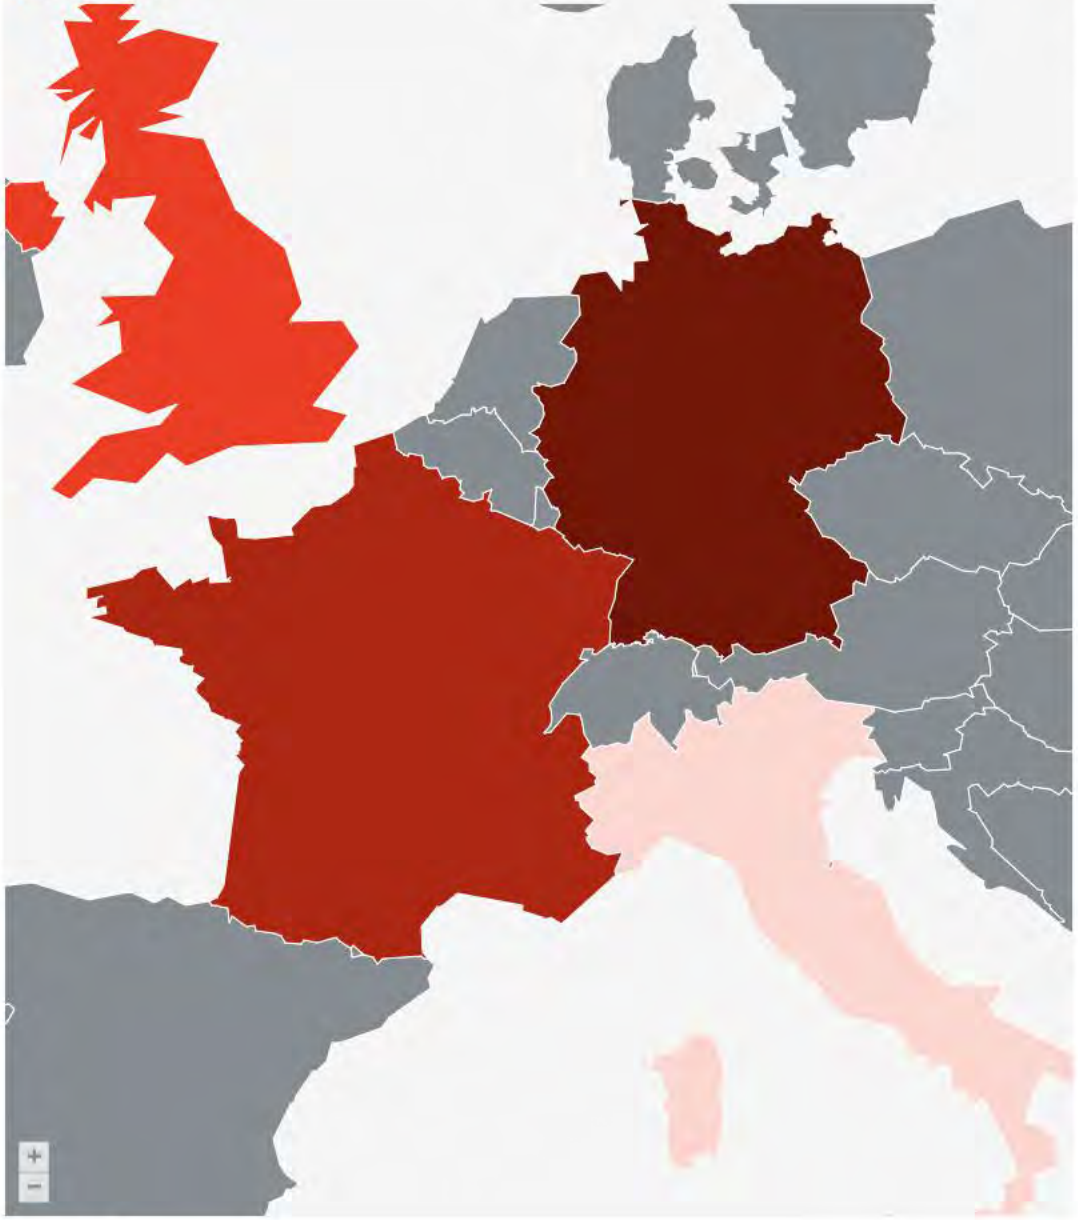

Revenue Stream vs Benchmark

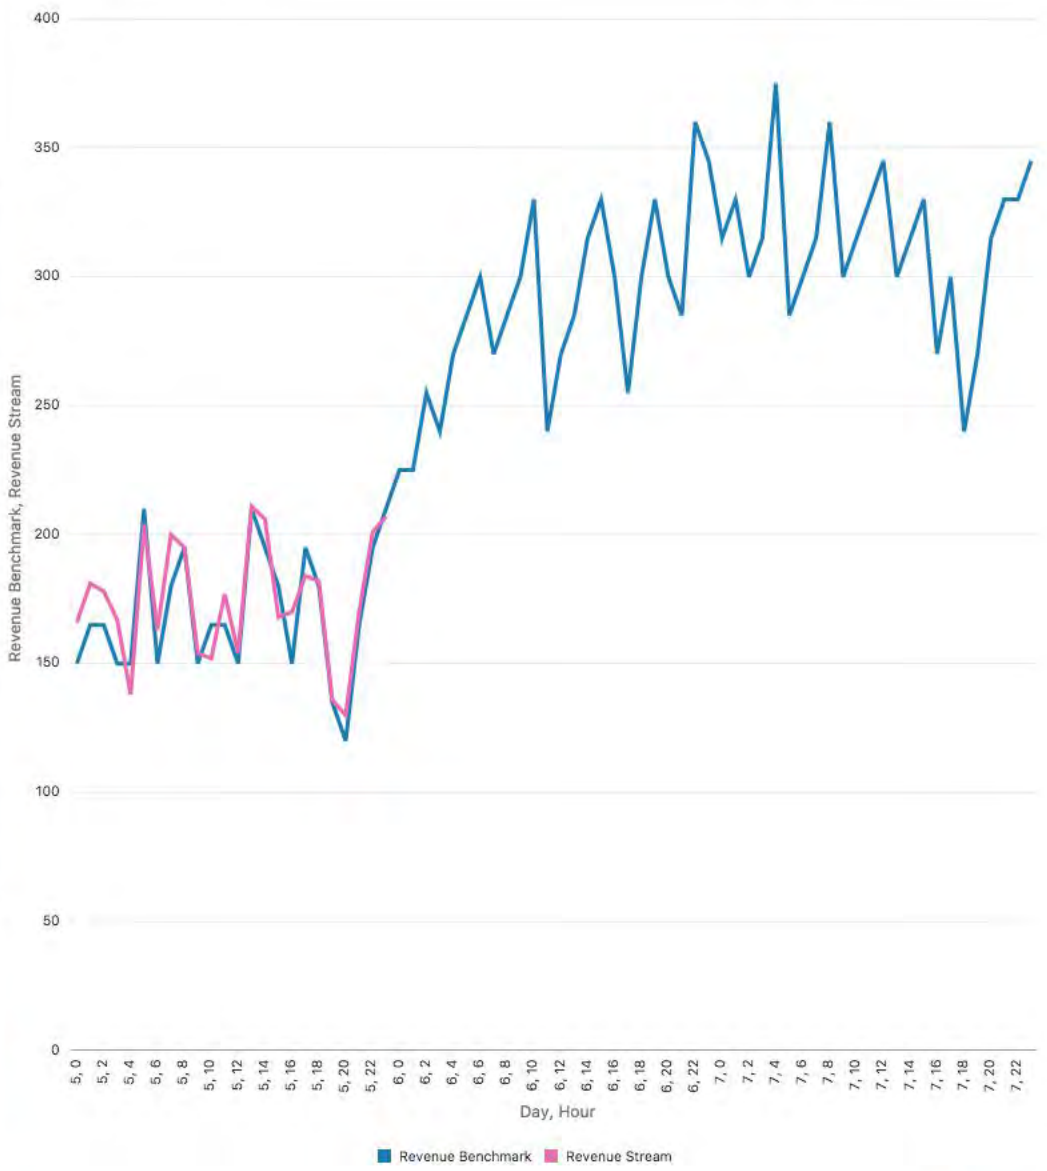

Streams

ORACLE Data Visualization Desktop

Search

- KAFKA_DEMO_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Add Calculation

Streaming Demo
Project

Prepare | Visualize | Narrate

Network Errors

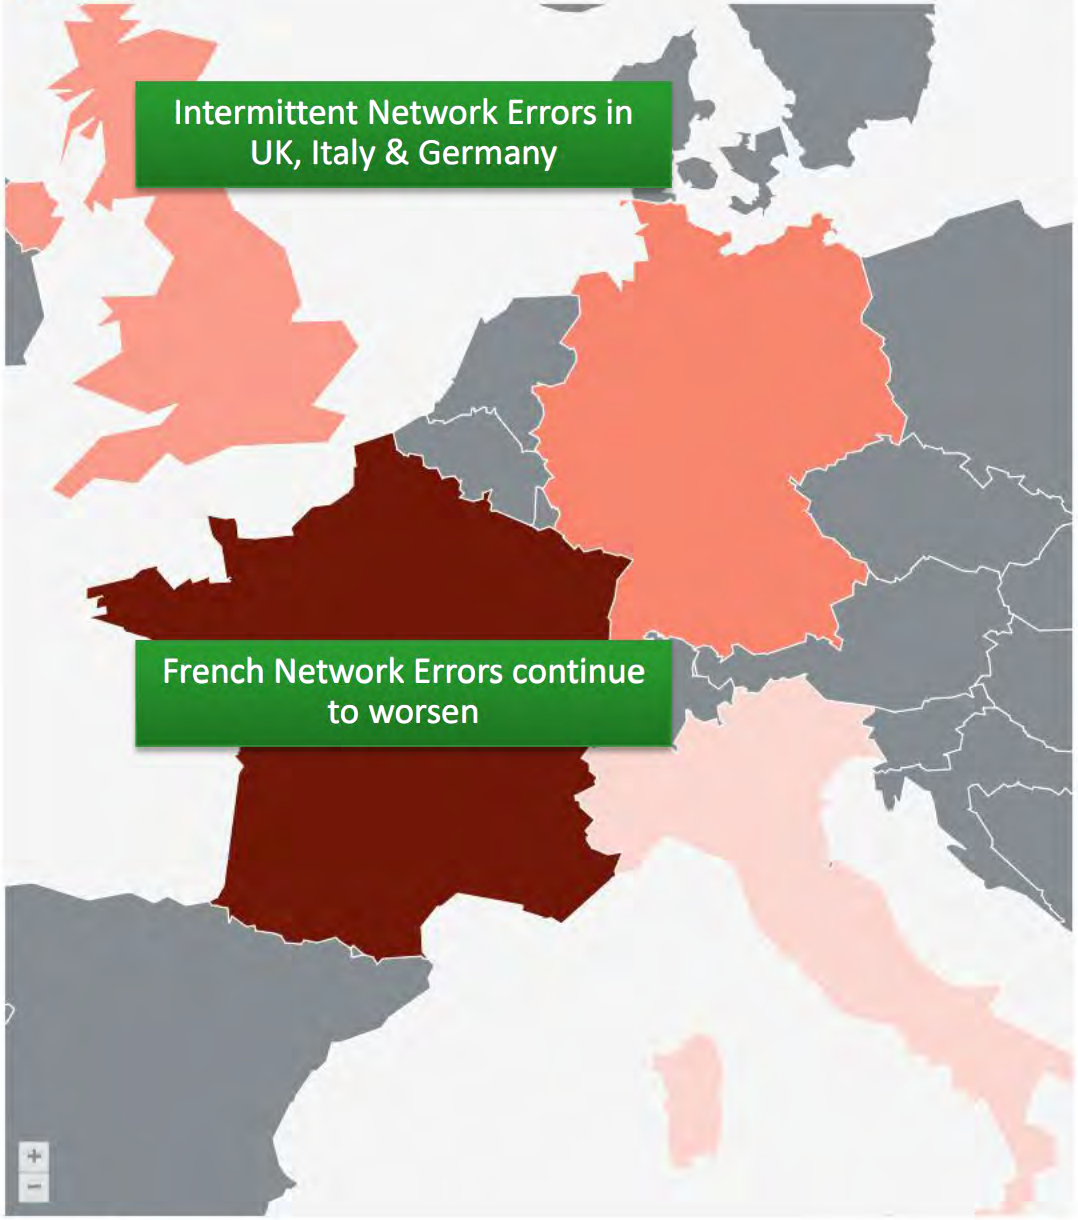

Revenue Stream vs Benchmark

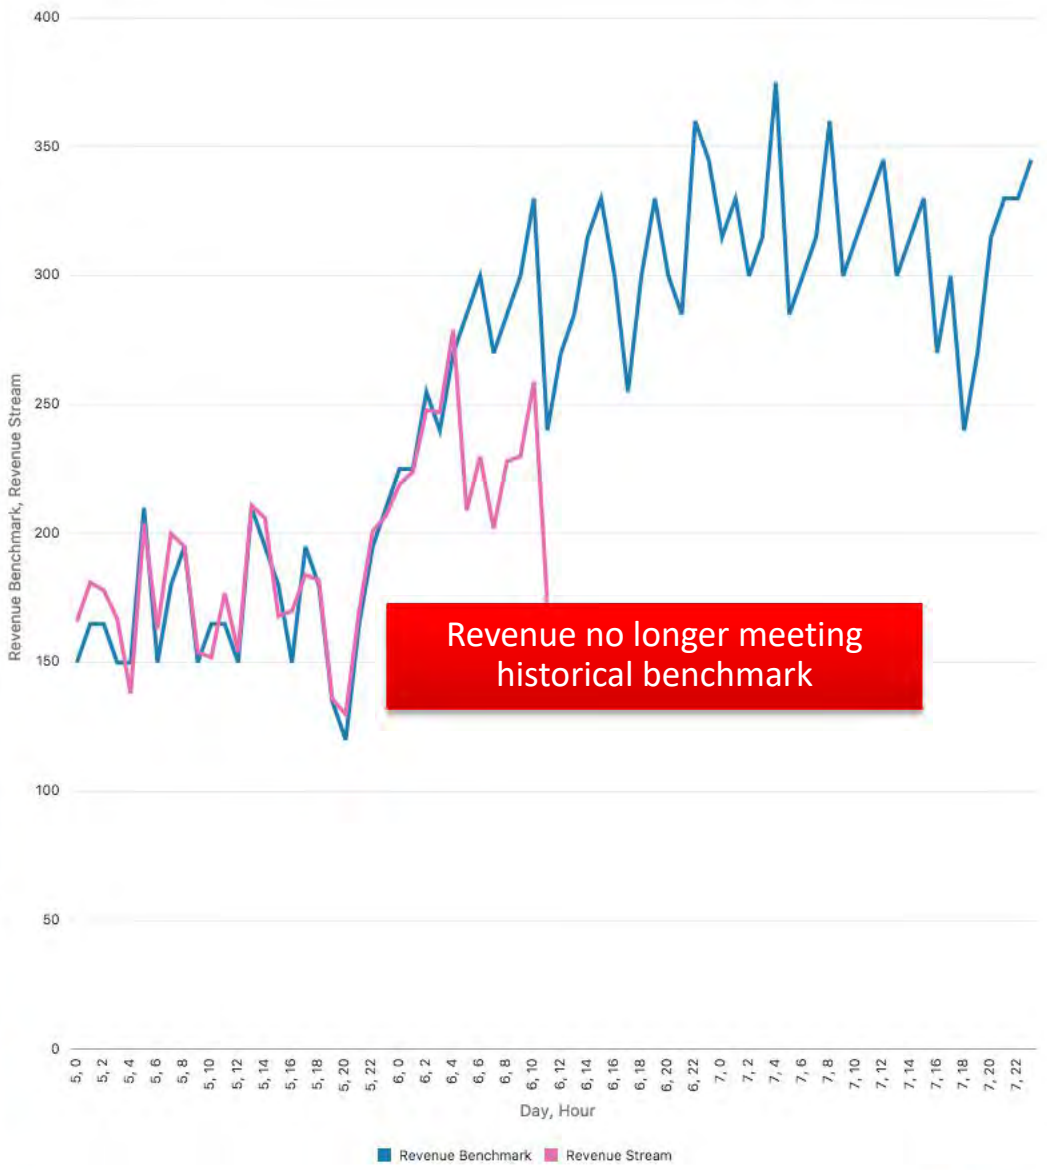

Streams

Search

- KAFKA_DEMO_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Add Calculation

Streaming Demo
Project

Prepare | Visualize | Narrate

Network Errors

Revenue Stream vs Benchmark

Streams

Search

- KAFKA_DEMO_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Add Calculation

Streaming Demo
Project

Prepare | Visualize | Narrate

Network Errors

Revenue Stream vs Benchmark

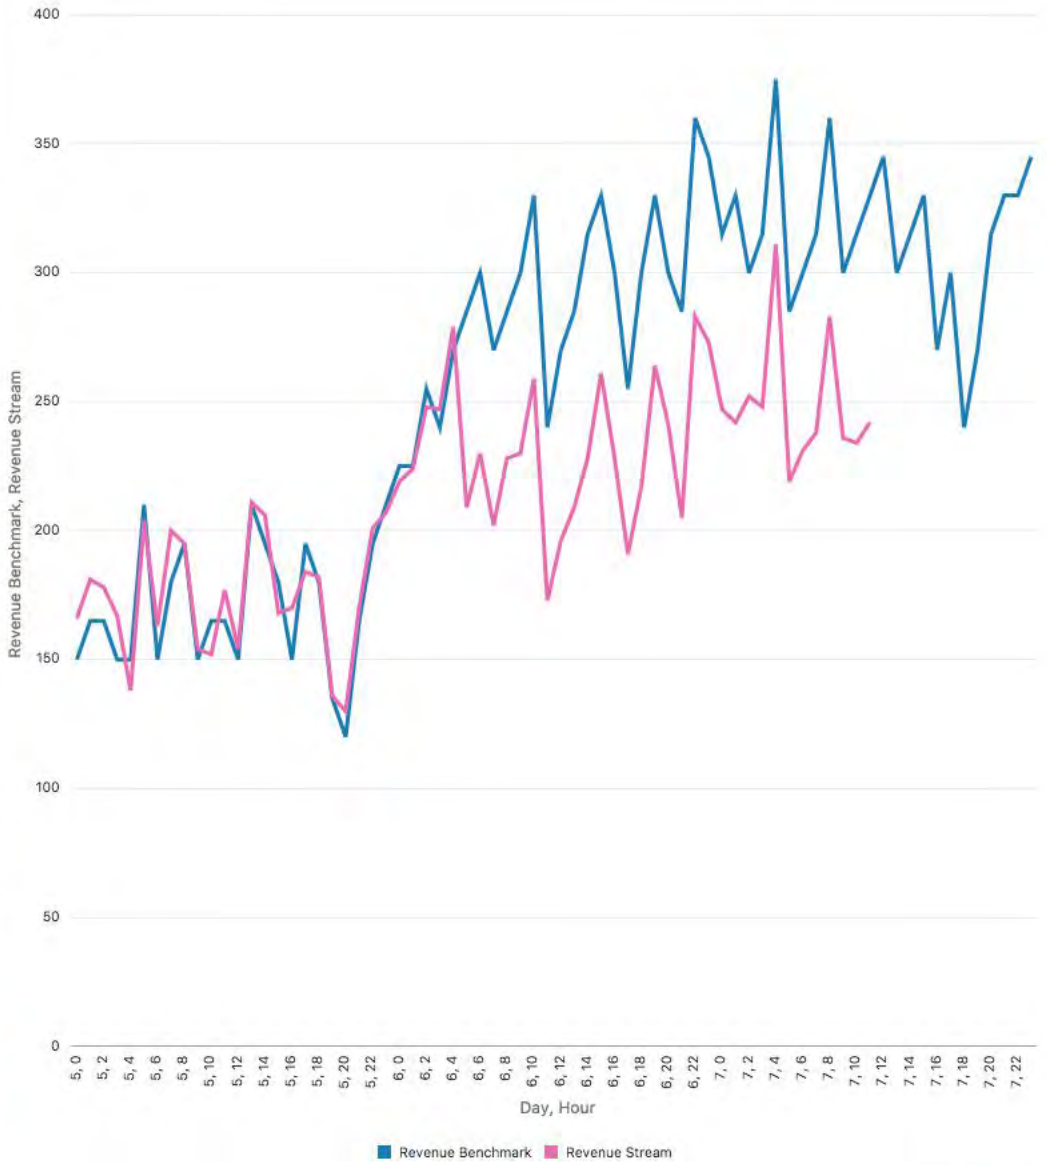

Streams

Search

- KAFKA_DEMO_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Add Calculation

Streaming Demo
Project

Prepare | Visualize | Narrate

Network Errors

Revenue Stream vs Benchmark

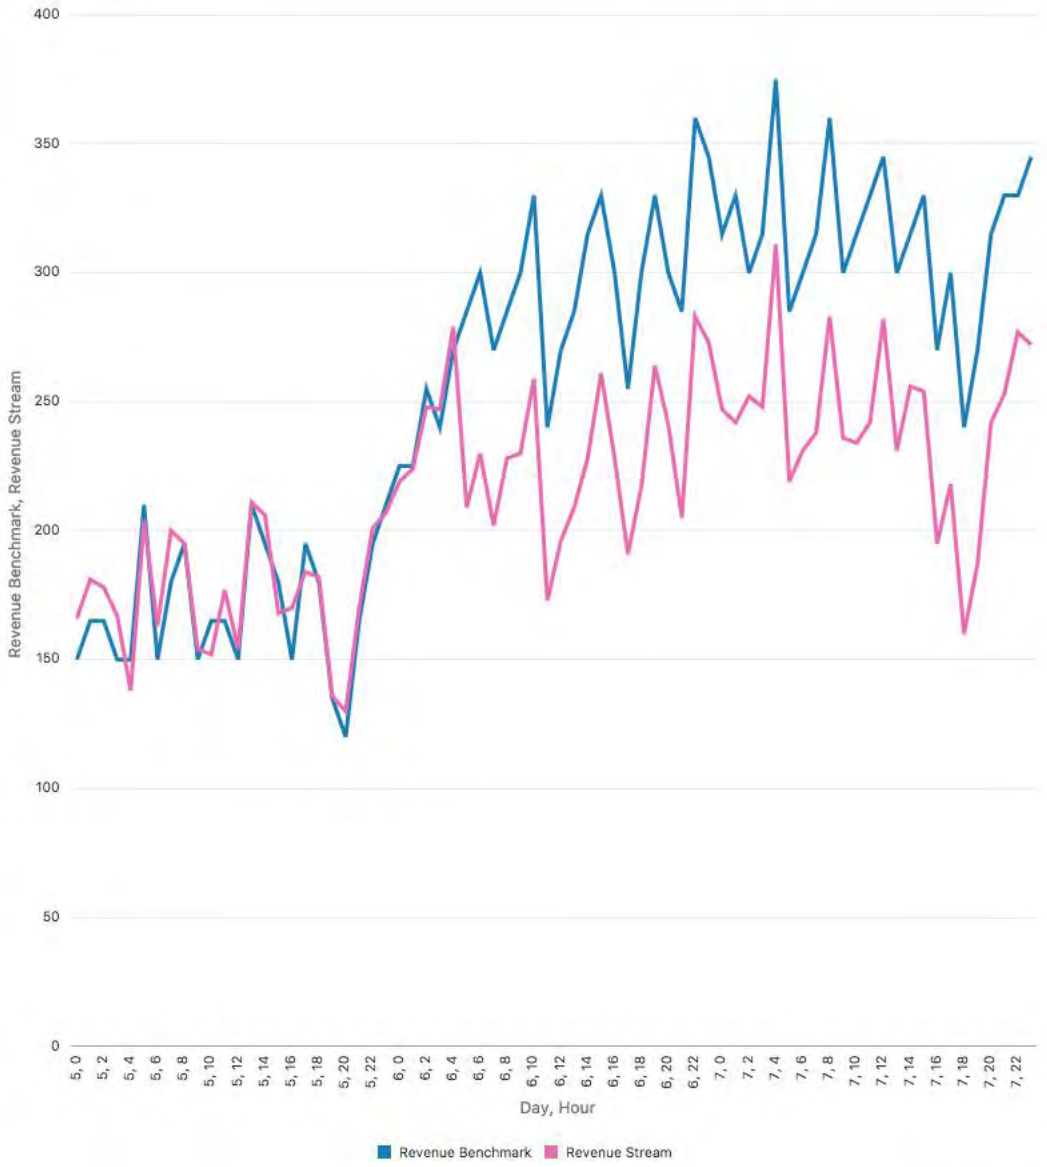

Streams

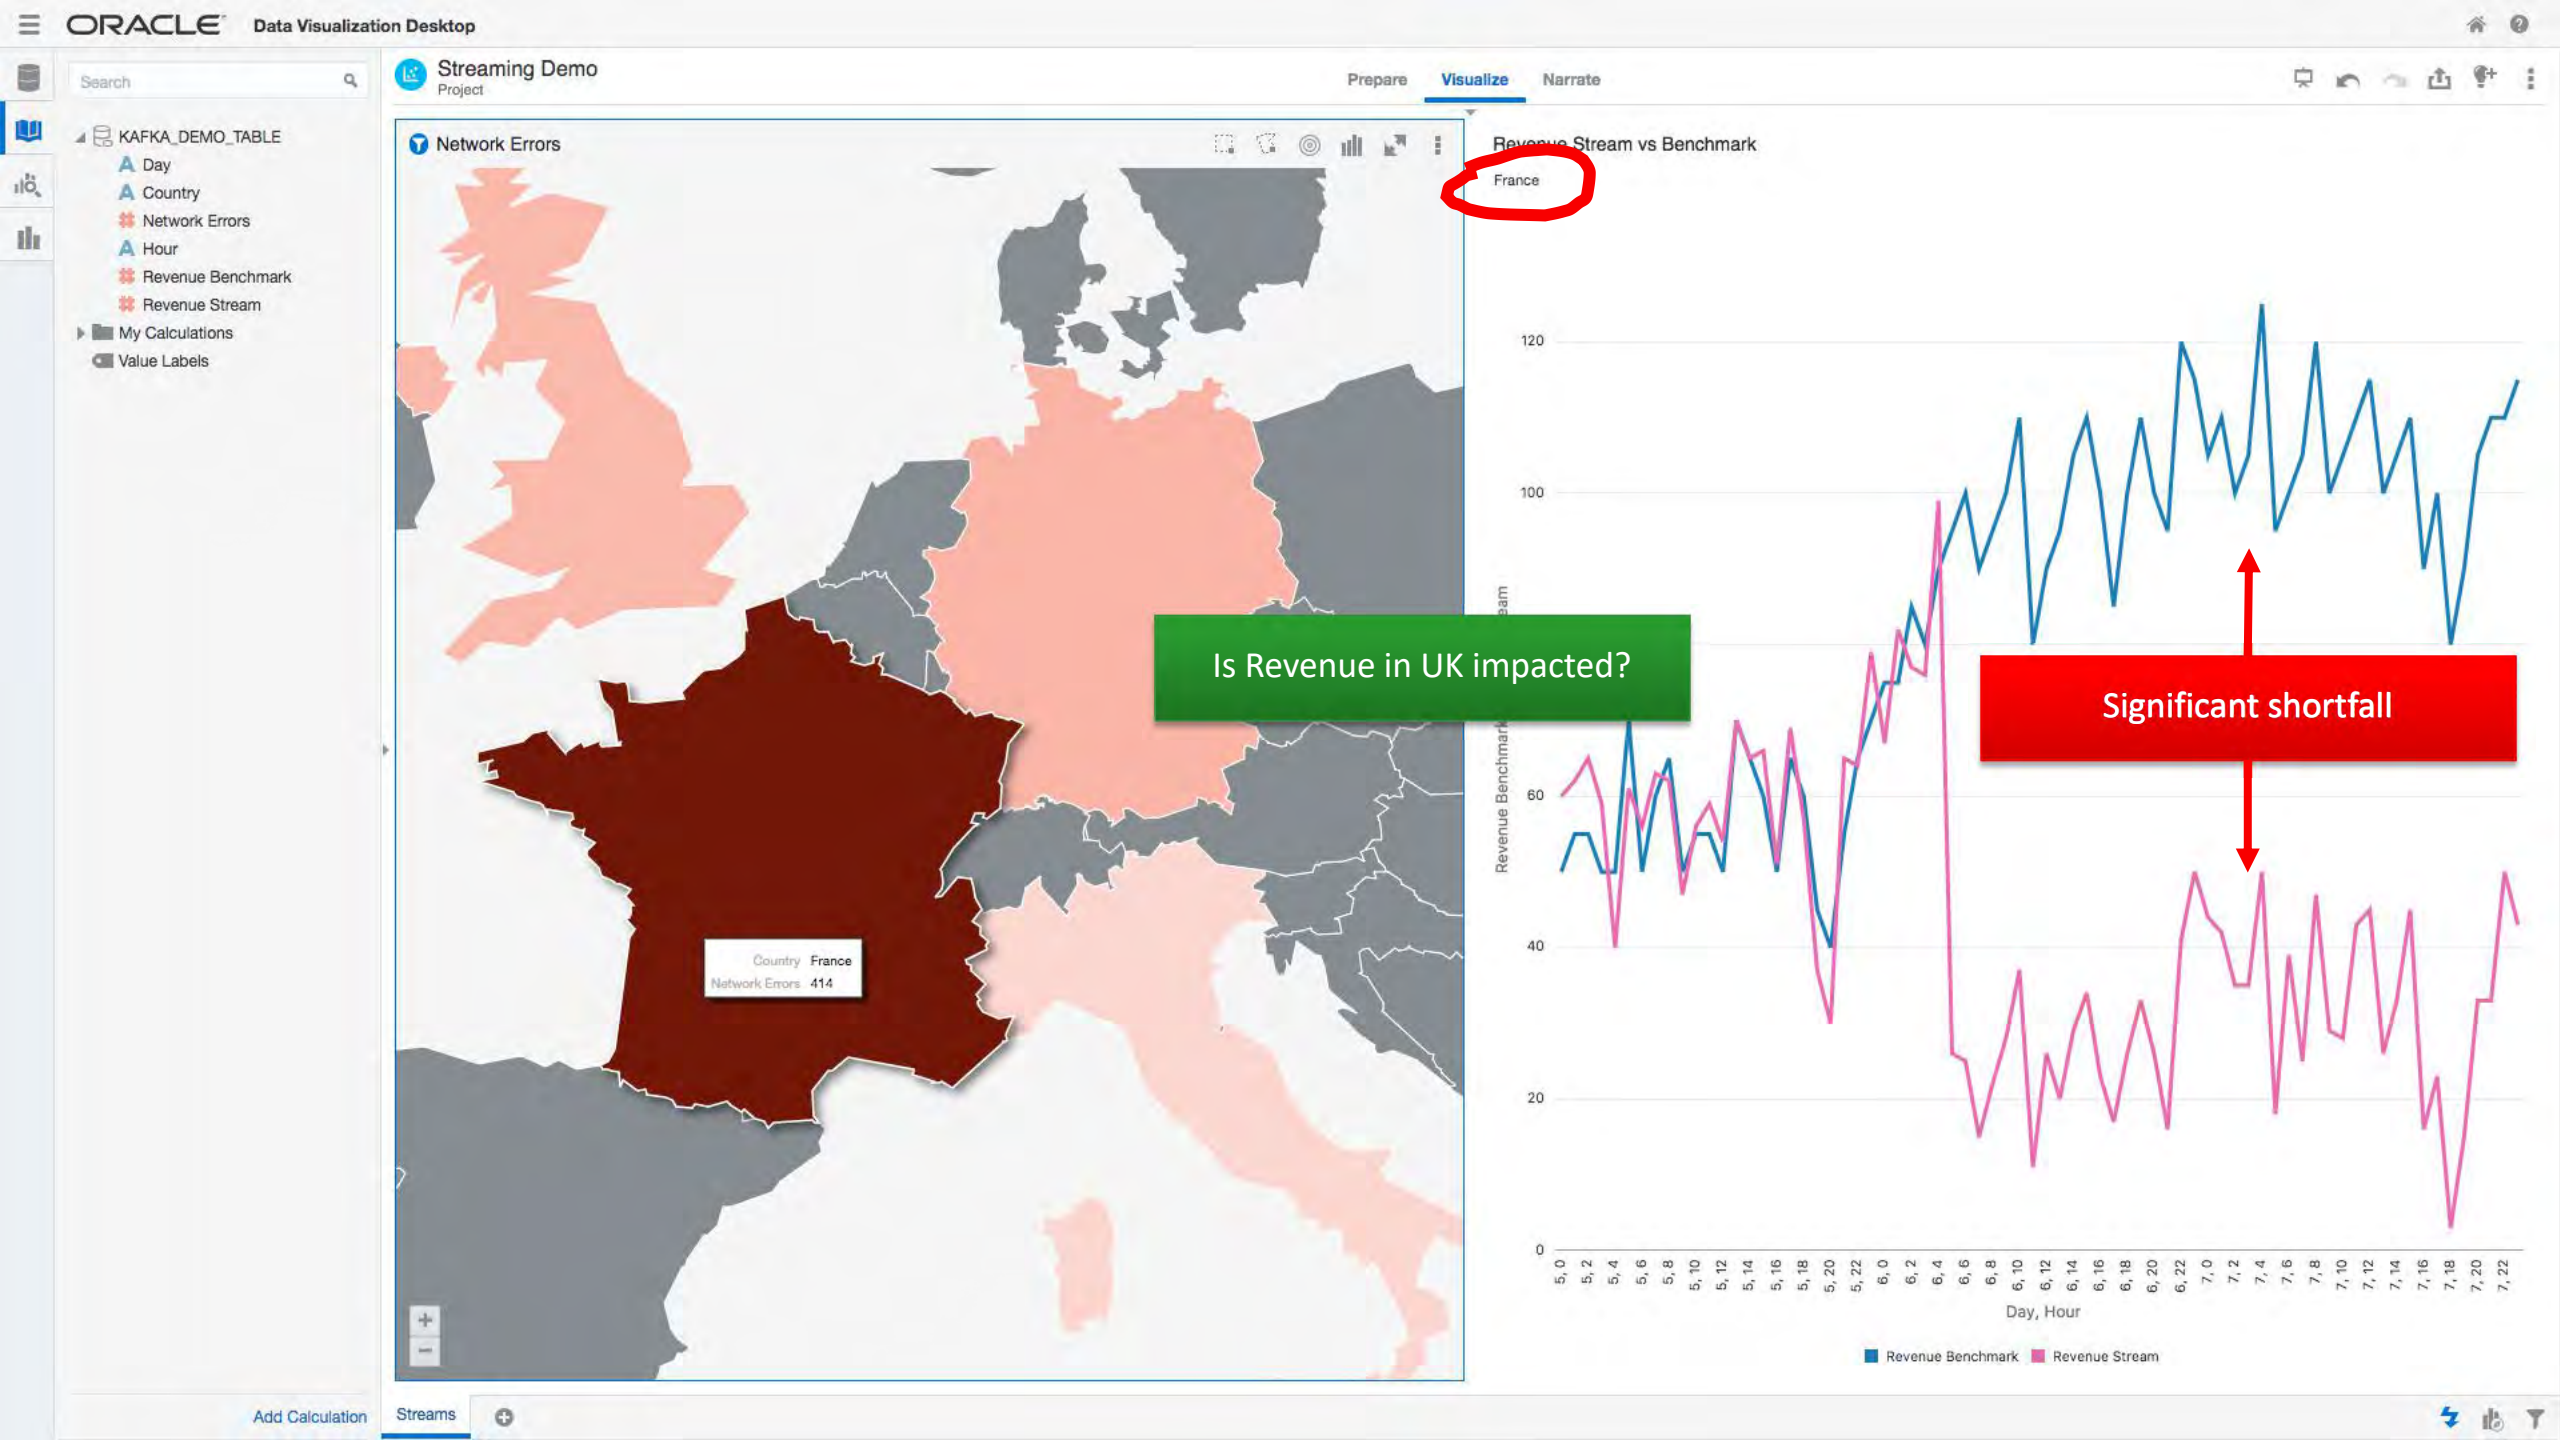
ORACLE Data Visualization Desktop
Search
KAFKA_DEMO_TABLE
Day
Country
Network Errors
Hour
Revenue Benchmark
Revenue Stream
My Calculations
Value Labels
Add Calculation
Streaming Demo
Project
Prepare
Visualize
Narrate
Network Errors
Country France
Network Errors 414
Is Revenue in UK impacted?
Revenue Stream vs Benchmark
France
Significant shortfall
Revenue Benchmark
Day, Hour
Revenue Benchmark
Revenue Stream
Streams

Search

- KAFKA_DEMO_TABLE
  - Day
  - Country
  - Network Errors
  - Hour
  - Revenue Benchmark
  - Revenue Stream
- My Calculations
- Value Labels

Add Calculation

Streaming Demo
Project

Prepare | Visualize | Narrate

Network Errors

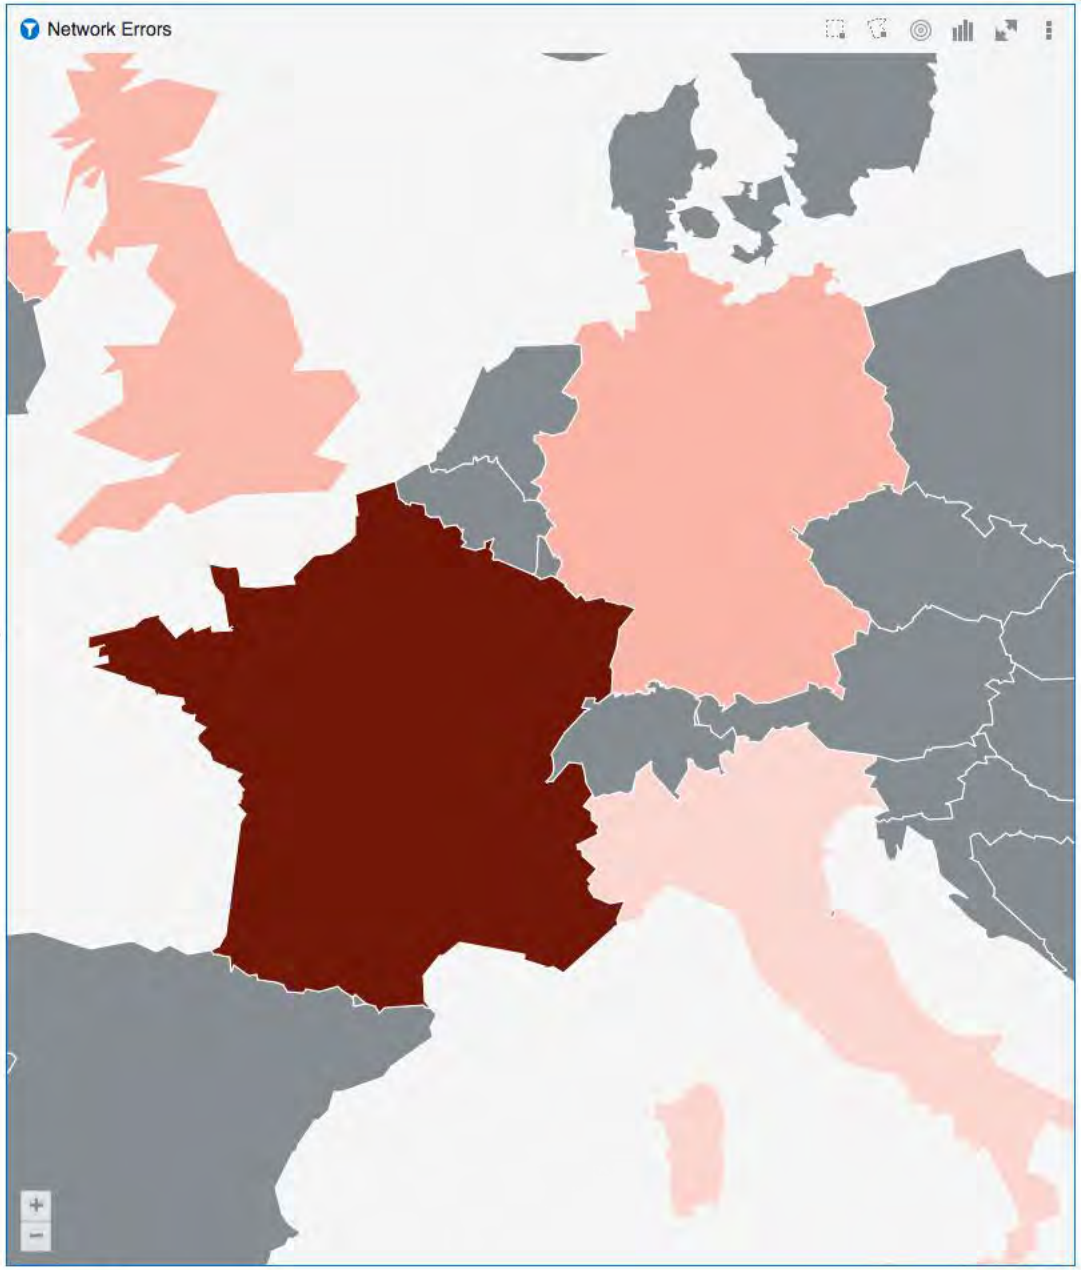

Revenue Stream vs Benchmark

United Kingdom

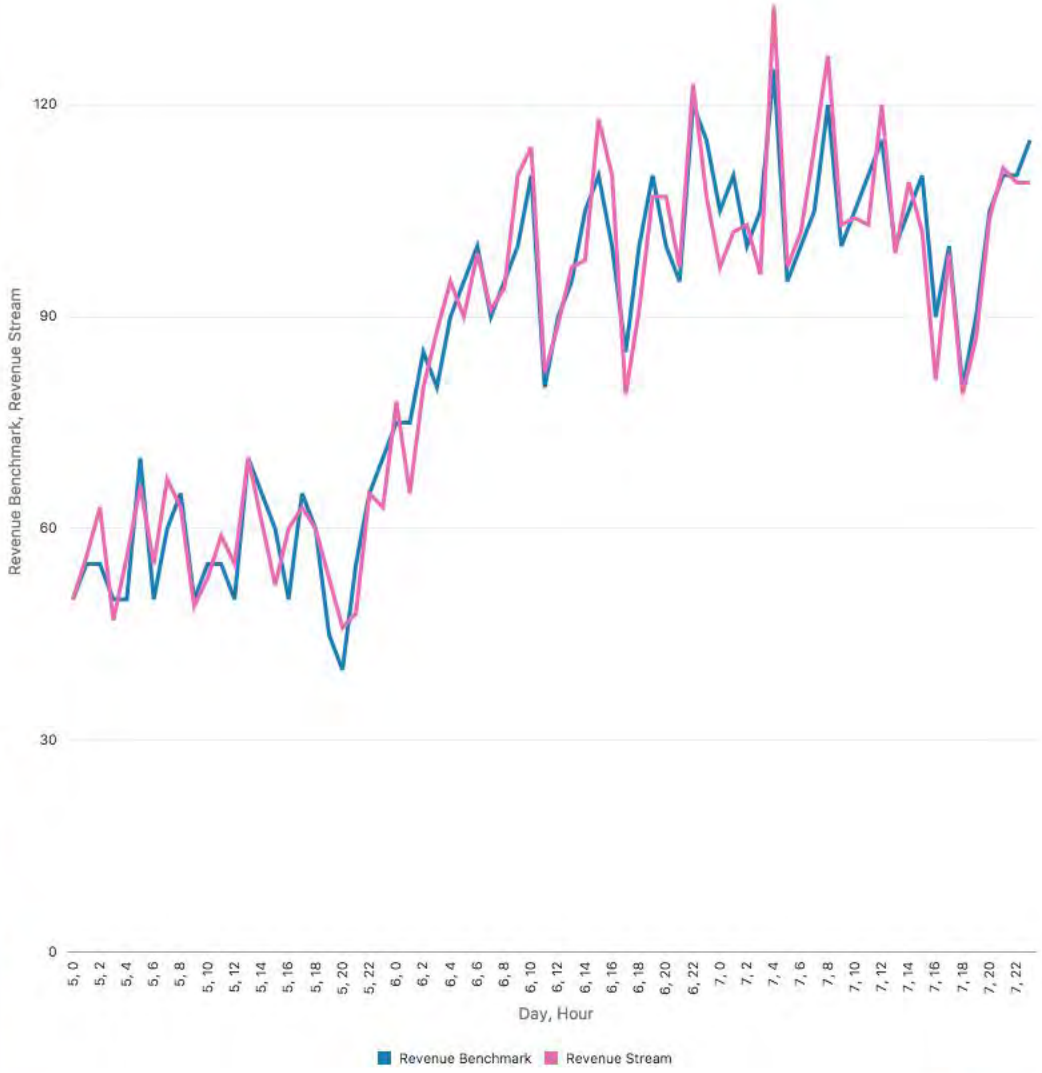

Streams

Insights Achieved with Simplicity

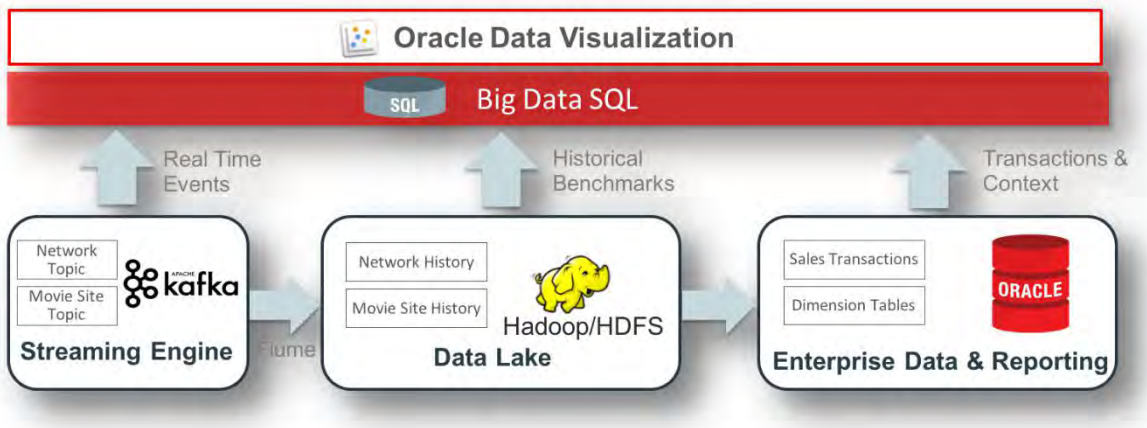

- Easily blend real time streams with history, benchmarks and context
  - Are we running at peak performance?
  - What is the opportunity cost of our current network latency?
- Any application realizes benefit
  - Use Oracle SQL and APIs over all data
- Ensure data is secure
  - Leverage Oracle advanced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

Big Data SQL Security Features

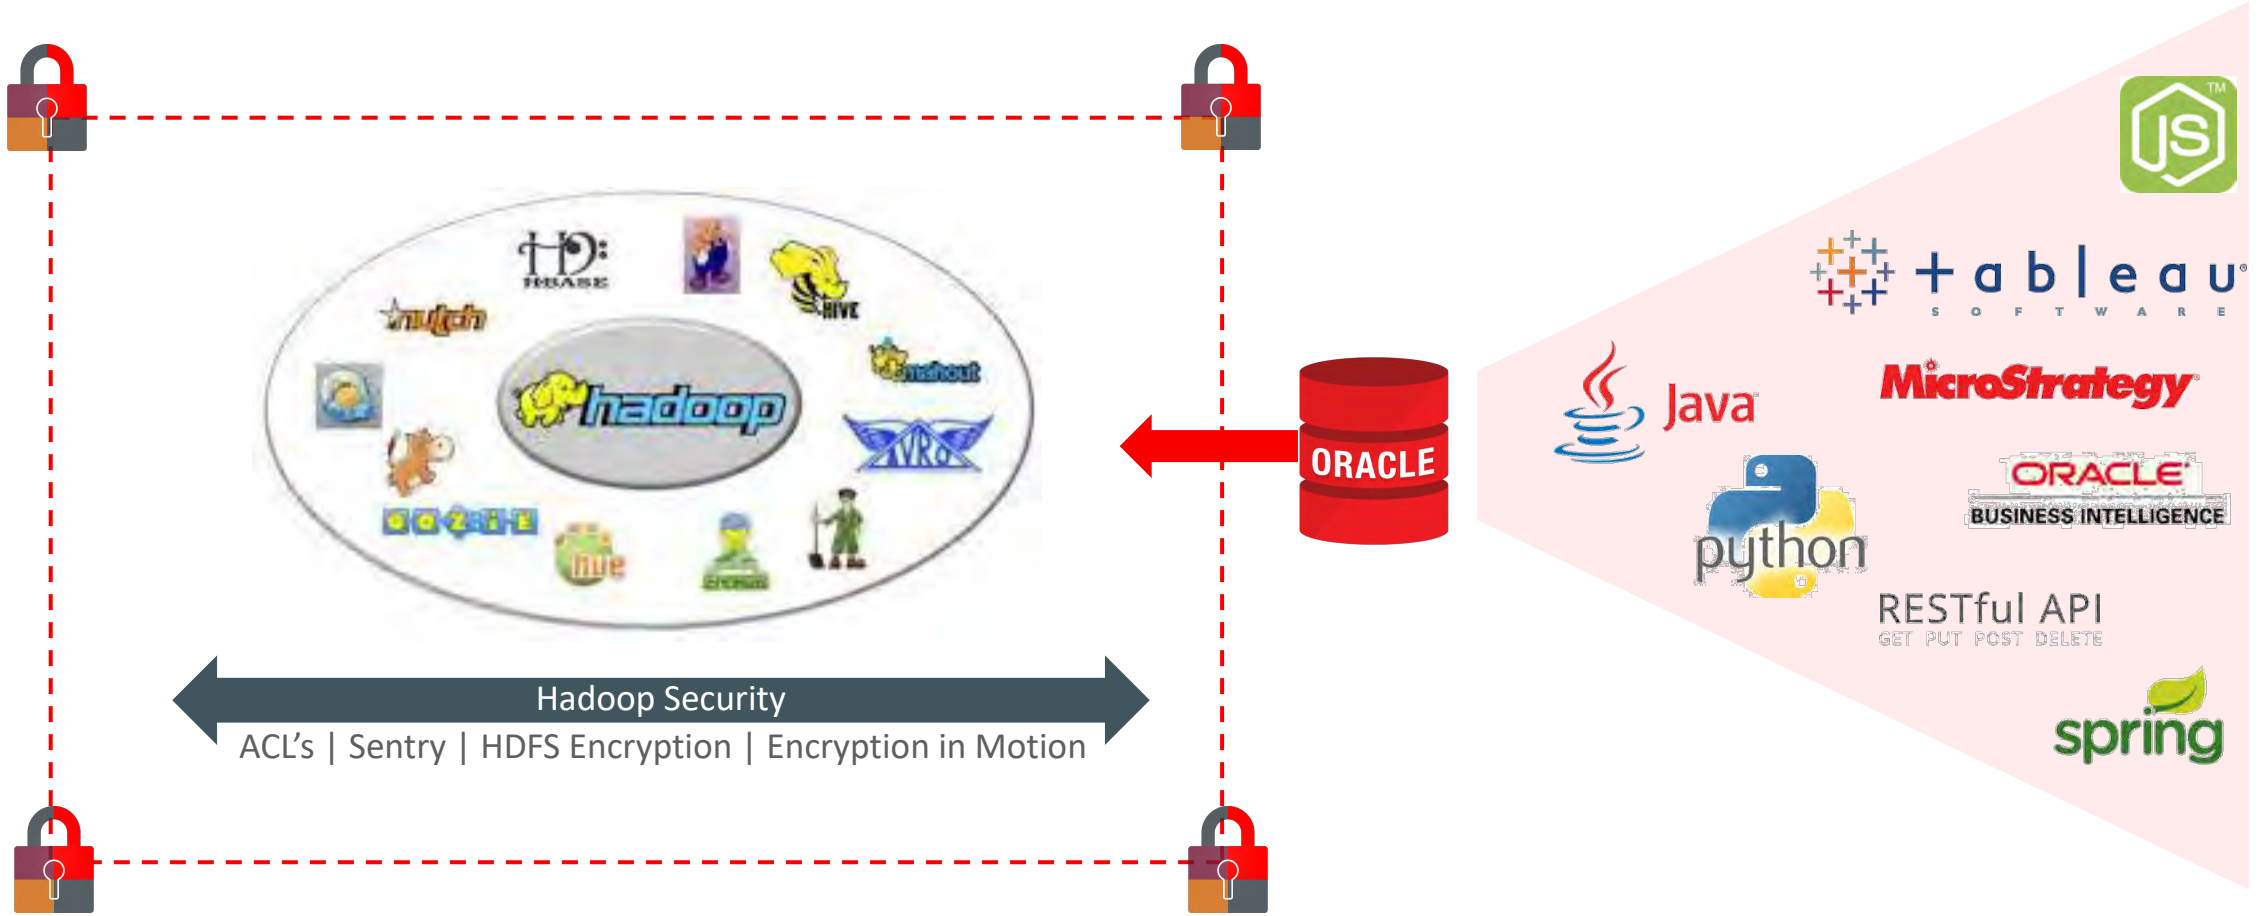

Big Data SQL Security Features

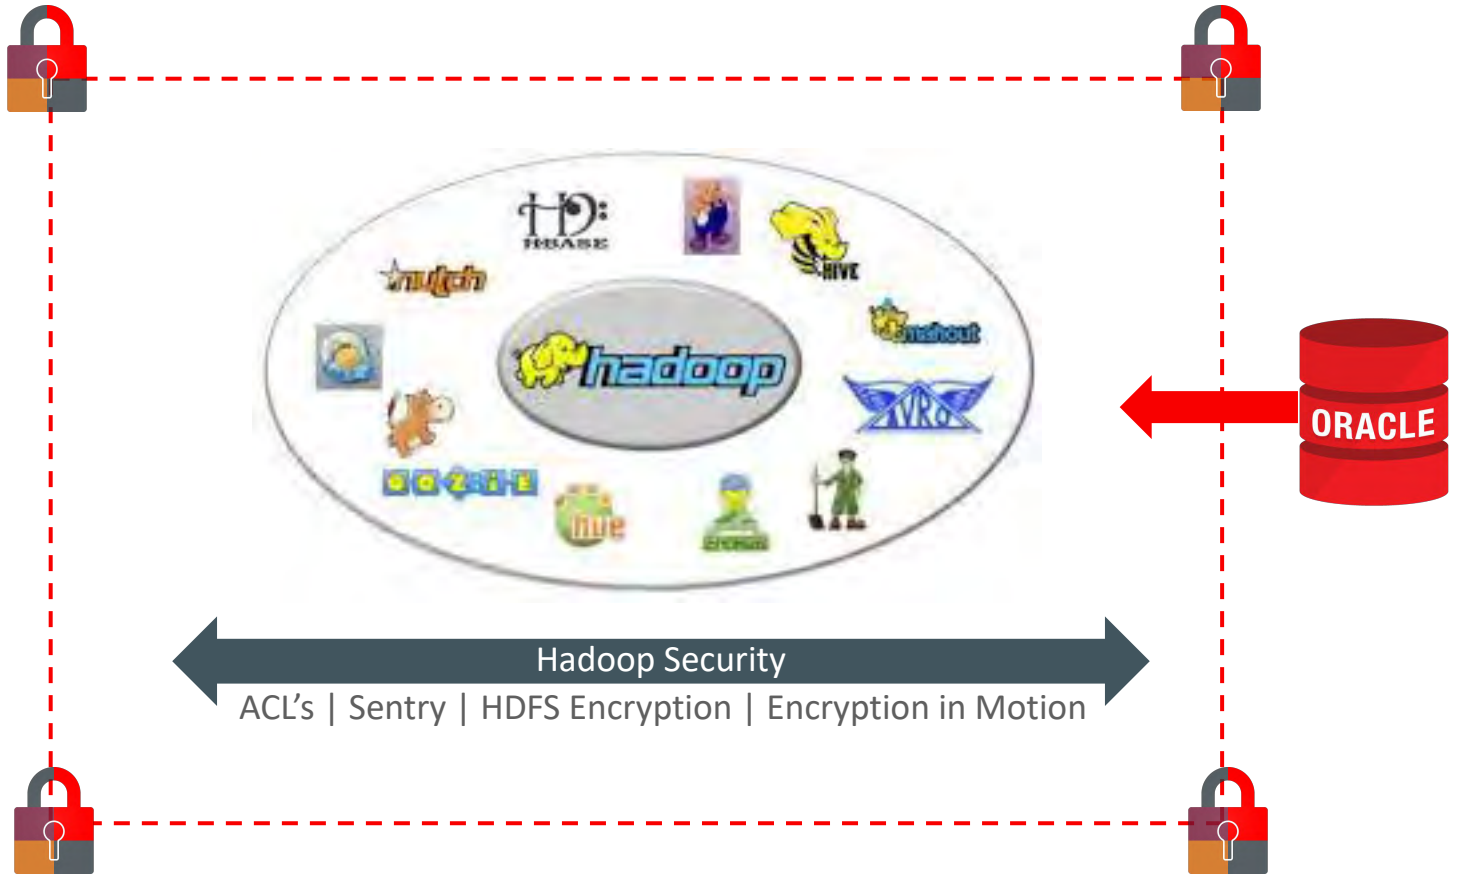

- Same security models apply to a wider range of data stores
- Advanced features such as data redaction can now be applied enabling joins between disparate sources
- Oracle security layers on top of existing Hadoop functionality

Big Data SQL Summary

Easily access any data across big data stores

Provide a unified security model across the sources

Analyze all data using Oracle's rich SQL dialect

Fast performance using Big Data SQL Smart Scan

More Information

- OTN: Big Data Lite Virtual Machine (a free sandbox environment to get started): http://www.oracle.com/technetwork/database/bigdata-appliance/oracle-bigdatalite-2104726.html
- Oracle.com: https://www.oracle.com/big-data/index.html
- Blog: (technical examples and tips): https://blogs.oracle.com/datawarehousing/