

ORACLE®

Oracle Database 12c – Built for Data Warehousing

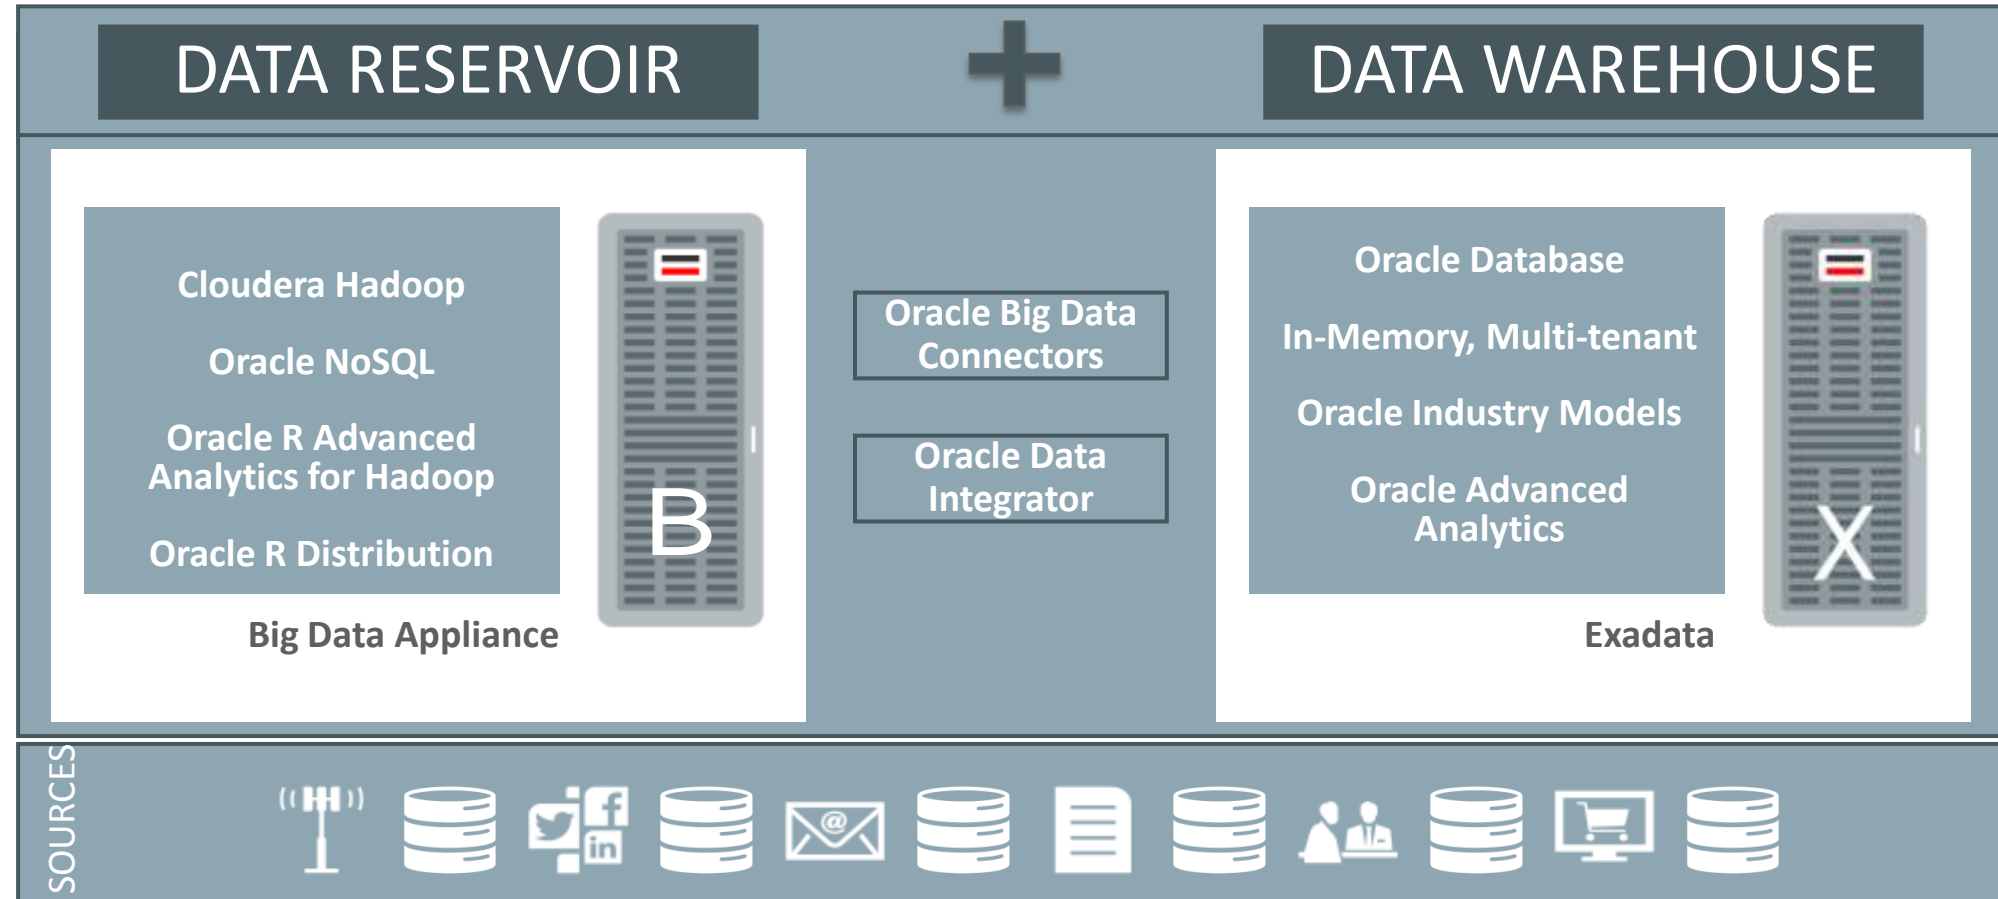
Fekete Zoltán

Agenda

- 1 The Oracle Big Data Management System
- 2 Oracle Database 12c for Data Warehousing
- 3 The Data Warehouse as Foundation for Big Data

Oracle Big Data Management System

Comprehensive Information Management



Oracle Big Data Management System

Comprehensive Information Management

Oracle Database 12c and
Oracle Exadata for
Data Warehousing

The Data Warehouse as
Foundation for Big Data



Agenda

- 1 The Oracle Big Data Management System
- 2 Oracle Database 12c for Data Warehousing
- 3 The Data Warehouse as Foundation for Big Data

Why Oracle for Data Warehousing?

Built to address your Information Management needs

- Extreme performance
 - Optimized, intelligent hardware and software
 - From disk to memory
- Efficient database management
 - Optimized data management
 - One platform for all data warehousing applications
- Accelerated analytical performance across all data
 - All your analysis
 - All your data

Extreme Performance

Optimized Hardware for Data Warehousing

Oracle Exadata Database Machine

Oracle's strategic database platform



- One integrated system
 - Servers, network, storage
 - Built and optimized for the Oracle Database
- Intelligent, scale-out storage
 - Offloaded database processing
 - Smart storage hierarchy – memory, flash, disk
- Industry best data compression
- Superior performance, price-performance, availability, supportability

Highly Engineered and Standardized

Less risk, better results



- Hundreds of Engineer years optimizing and hardening Exadata end-to-end
 - Database, OS, drivers, firmware all optimized together
 - Frees IT talent to focus on business needs
- Less deployment risk
 - Delivered assembled, debugged, and ready-to-run
- Single-vendor support: from database to disk
- Runs all existing Oracle database workloads

Elastically Scale-Out from Eighth-Rack to Multi-Rack

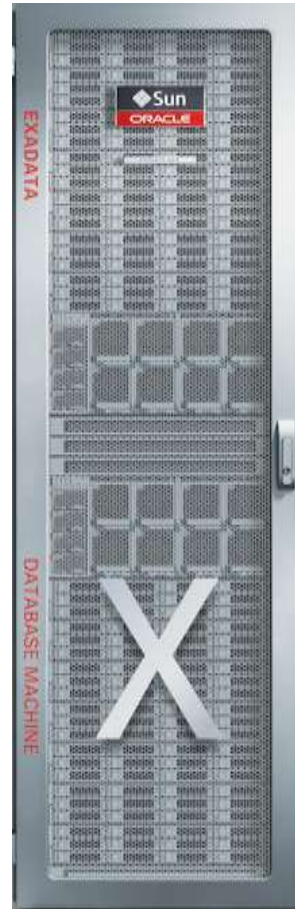


- Start with 2 Database Servers and 3 Storage Servers
 - Add database or storage servers online as needed
- Can expand older machines with new generation servers

Max 22 servers or
38U of servers per rack

Unique Software Capabilities for Data Warehousing

- Smart storage
 - Processing offload in storage
 - Smart Flash
 - Storage Indexes
 - 140GB/sec SQL data throughput per rack for High Capacity DW Machine
 - Over 4 million flash read IOPS
- Hybrid Columnar compression
 - Optimized for query or archival
 - 10x data reduction



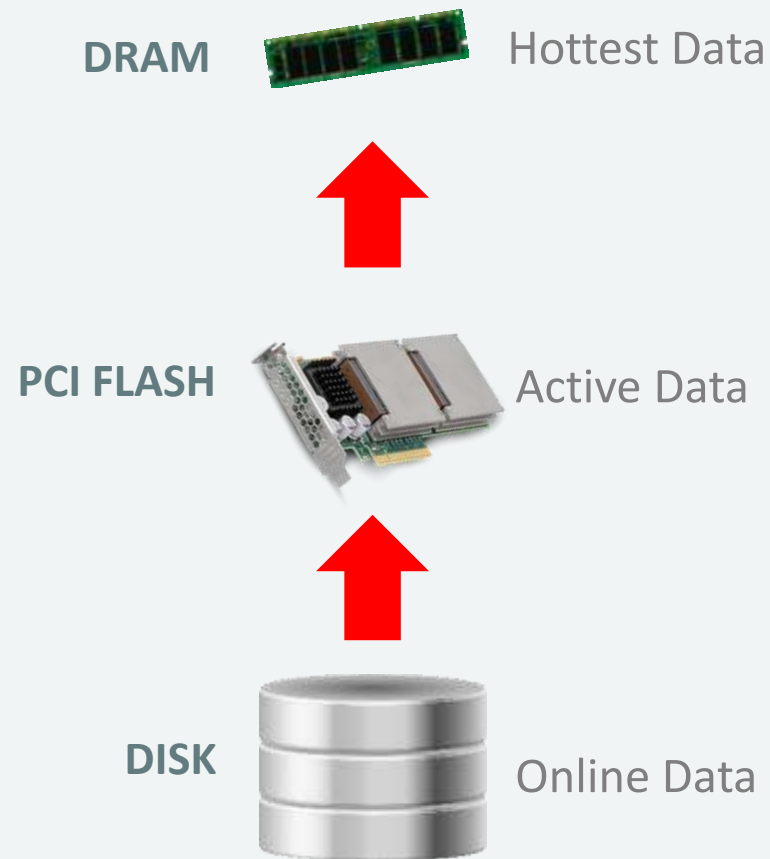
- Optimized for strategic and operational data warehousing
 - Ad-hoc queries over detail data
 - Near instantaneous tactical, short-running queries
 - Real-time, concurrent query and data load with managed quality of service
- Optimized for large data sets
 - Transparent compression, partitioning and zone maps

Runs your data warehousing workload 10x faster

Extreme Performance

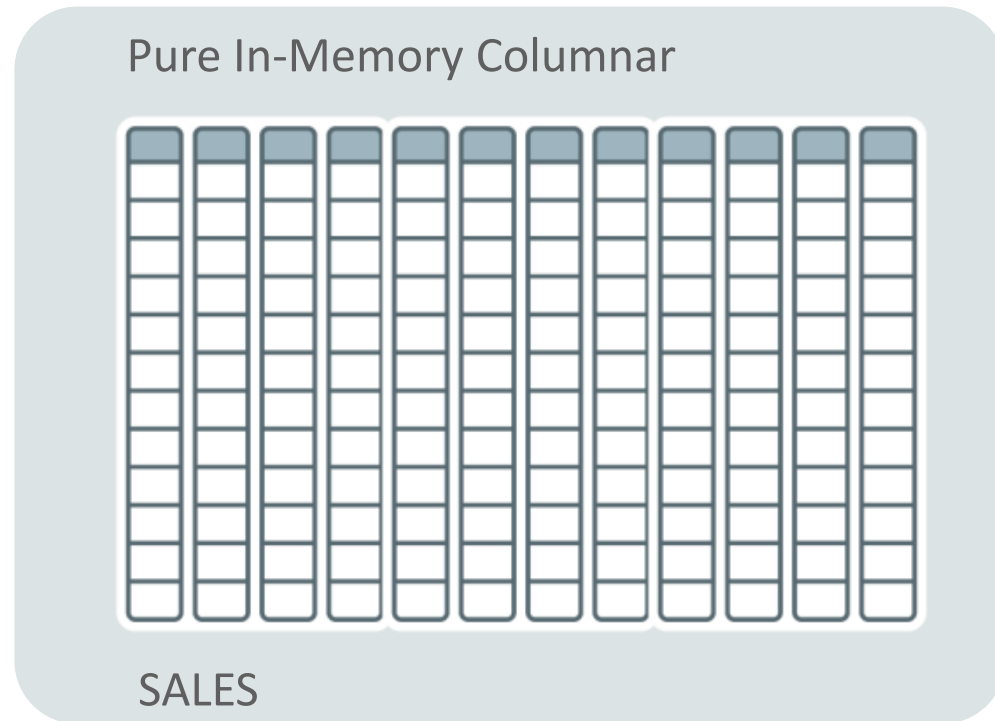
Real-time Analysis With Oracle Database In-Memory

Optimized from Disk to Memory



- Oracle Database In-Memory 100's GB/sec
- Exadata Smart Flash Cache 50-100 GB/sec
- Exadata Storage Servers 10's GB/sec

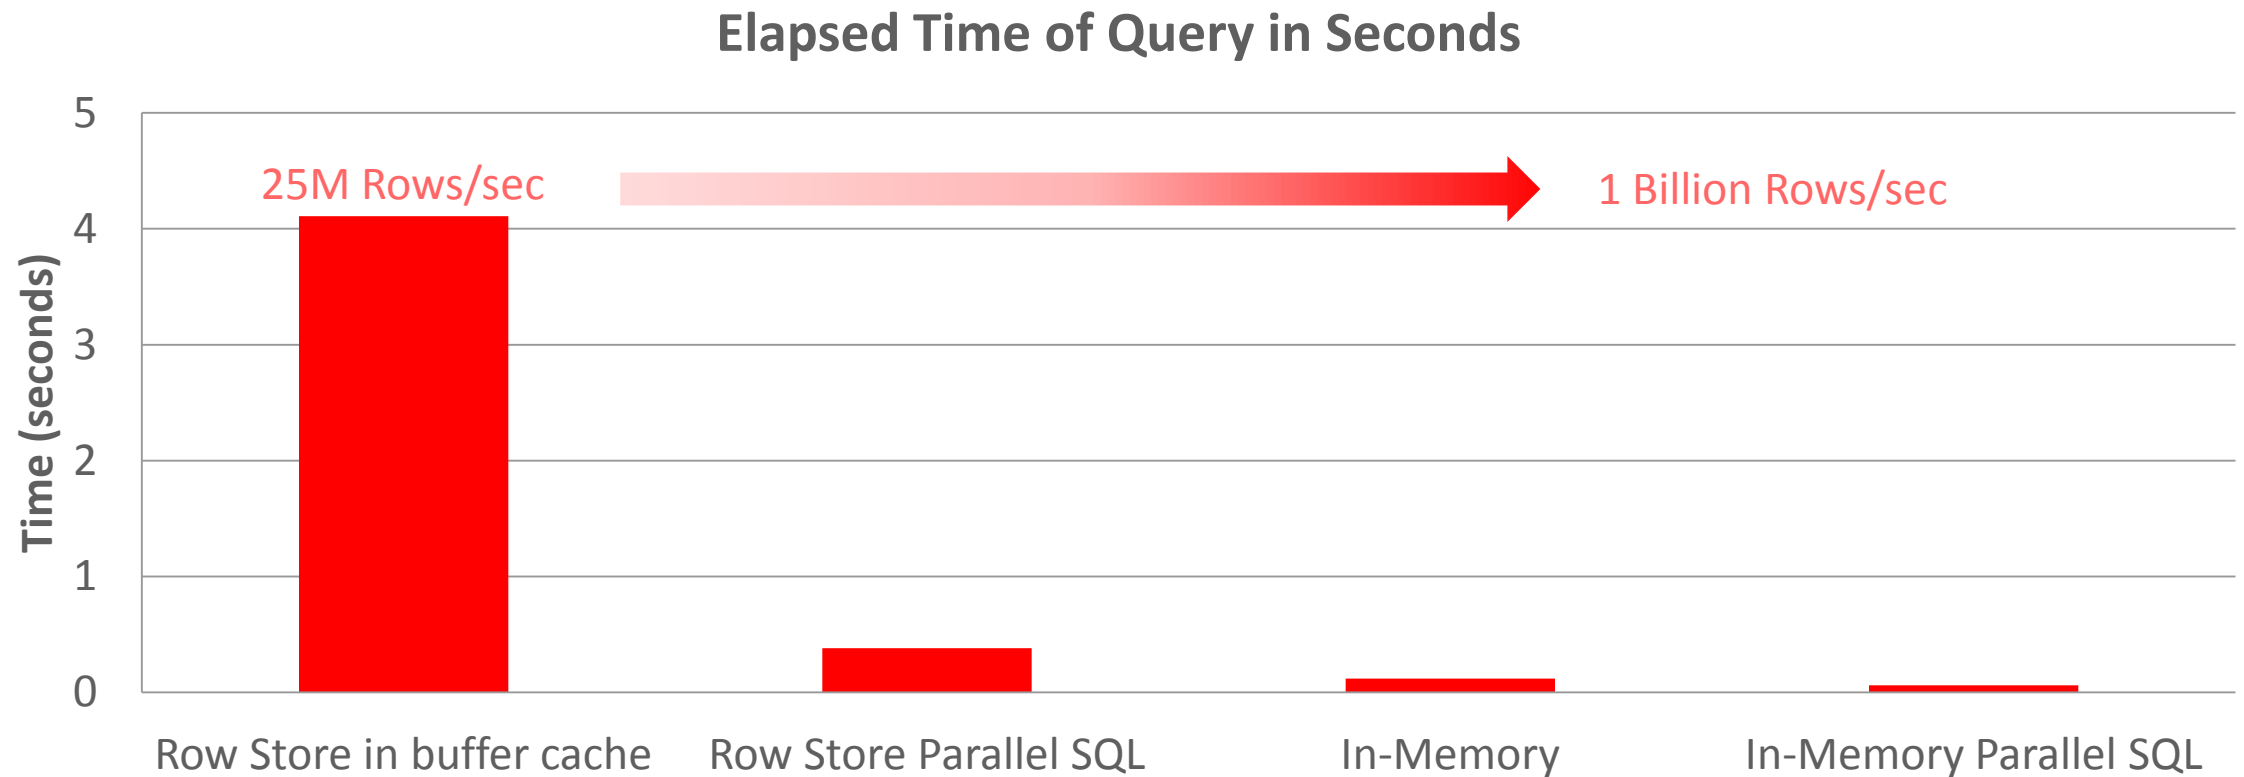
Oracle In-Memory Columnar Technology



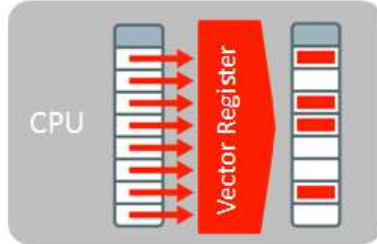
- Pure in-memory column format
- Not persistent, and no logging
- Quick to change data: fast OLTP
- 2x to 20x compression
- Enabled at table or partition level
- Available on all hardware platforms

In-Memory Scan Performance Example

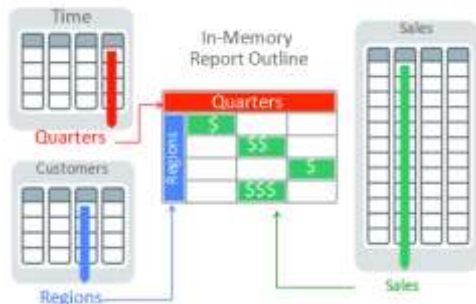
Scenario: scan 100 million rows to find a single row, no indexes



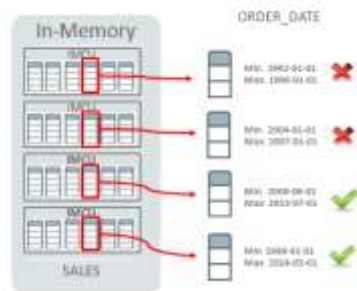
In-Memory: Much More Than Columnar Storage



- Faster scan, join and compare with vector (SIMD) CPU instructions

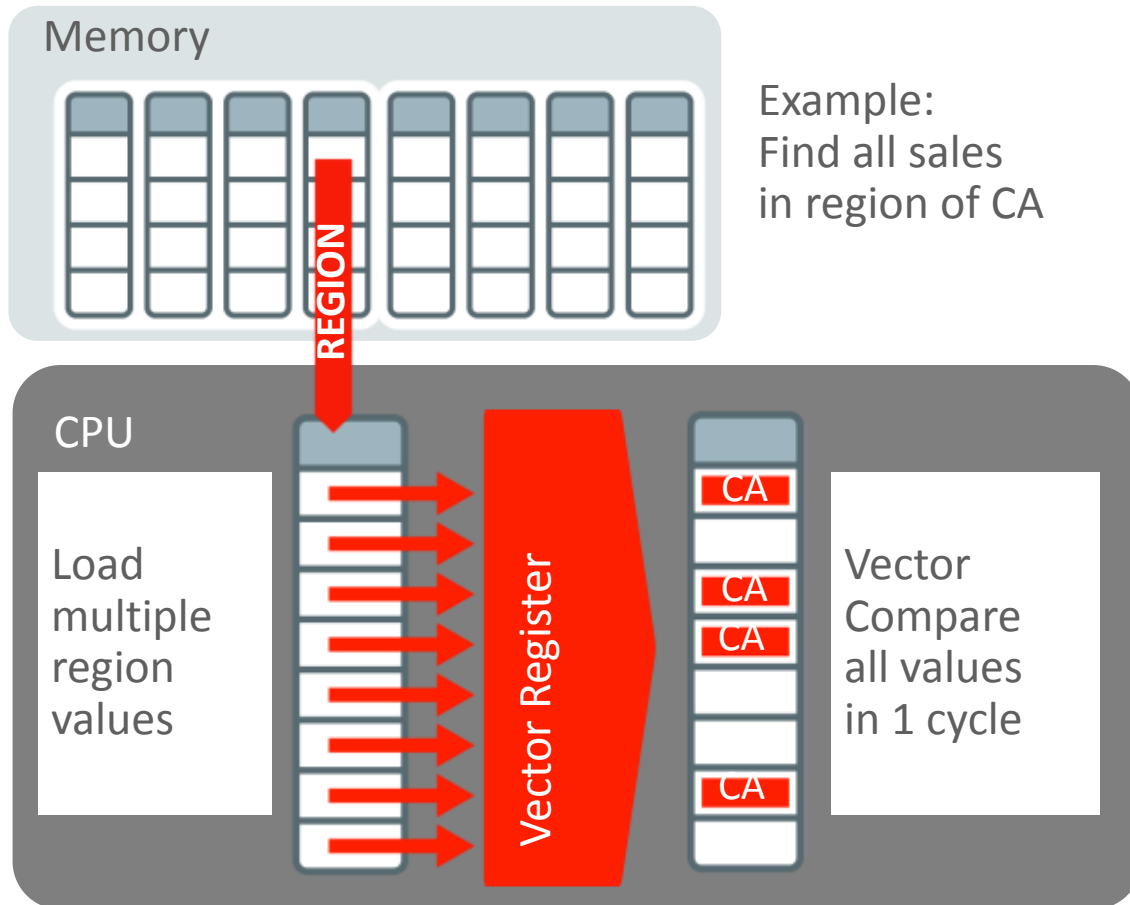


- Optimized star query processing with in-memory aggregation



- Optimized data access with in-memory storage indexes

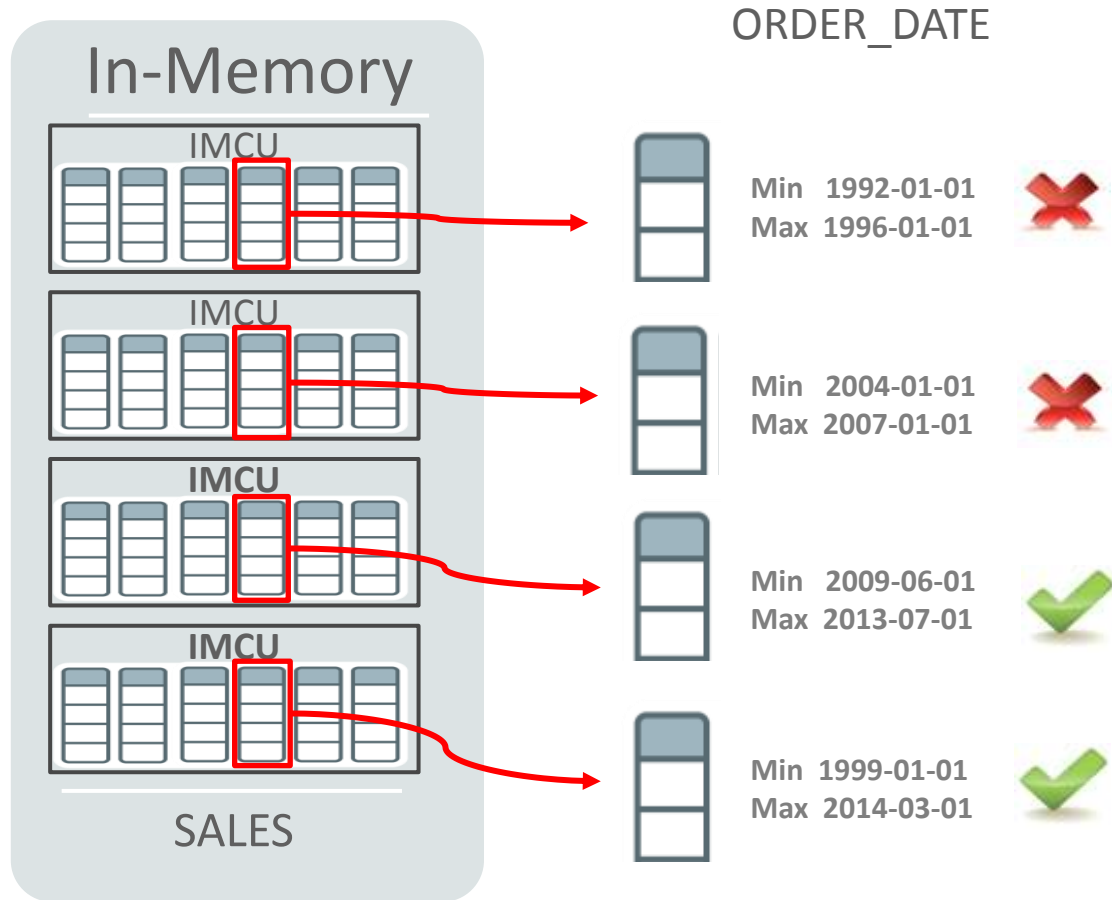
Scans Billions of Rows per Second per CPU Core



> 100x Faster

- Each CPU core scans local in-memory columns
- Scans use super fast “Single Instruction multiple Data Values” (SIMD) vector instructions
 - Originally designed for graphics & science
- **Billions of rows/sec** scan rate per CPU core

In-Memory Column Store Storage Index



- Data stored in “In-Memory Compression Units” (IMCU’s)
- A **storage index** records min/max values for each column unit
- Storage indexes allow IMCU pruning

```
select *  
from SALES  
where ORDER_DATE between  
      '2013-01-01' and  
      '2014-01-01'
```

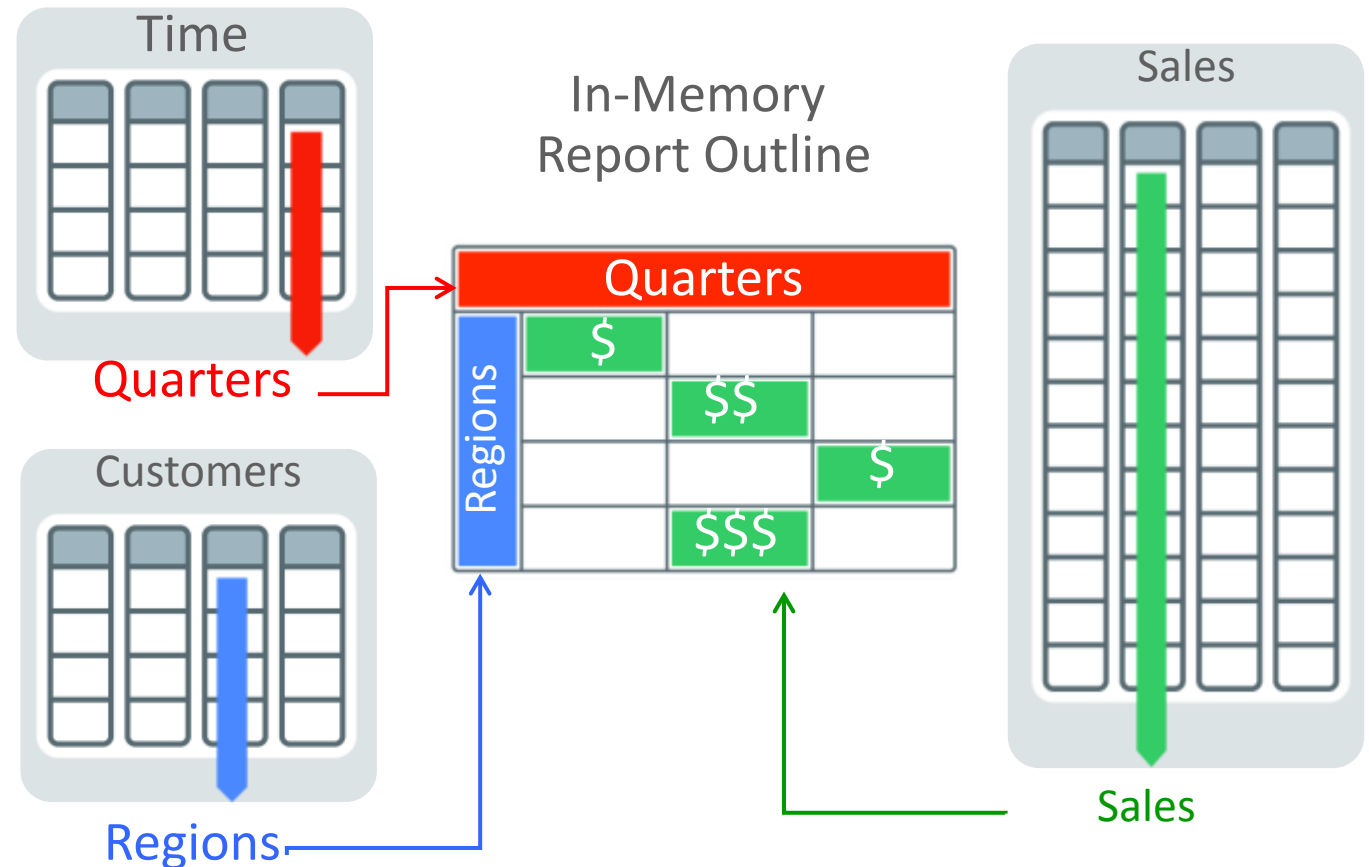
In-Memory Aggregation

New optimized algorithm for star query processing

Processing steps

- Transform joins into scan of the fact table
 - Fast in-memory scan with array lookups
- Optimize aggregation using in-memory arrays
 - Optimized in-memory data structures
- Late joins to dimension data
 - Minimizes data movement in the execution plan

Example: Report sales by Quarter and Region



Extreme Performance

Oracle Database - The Best Query Processing Engine

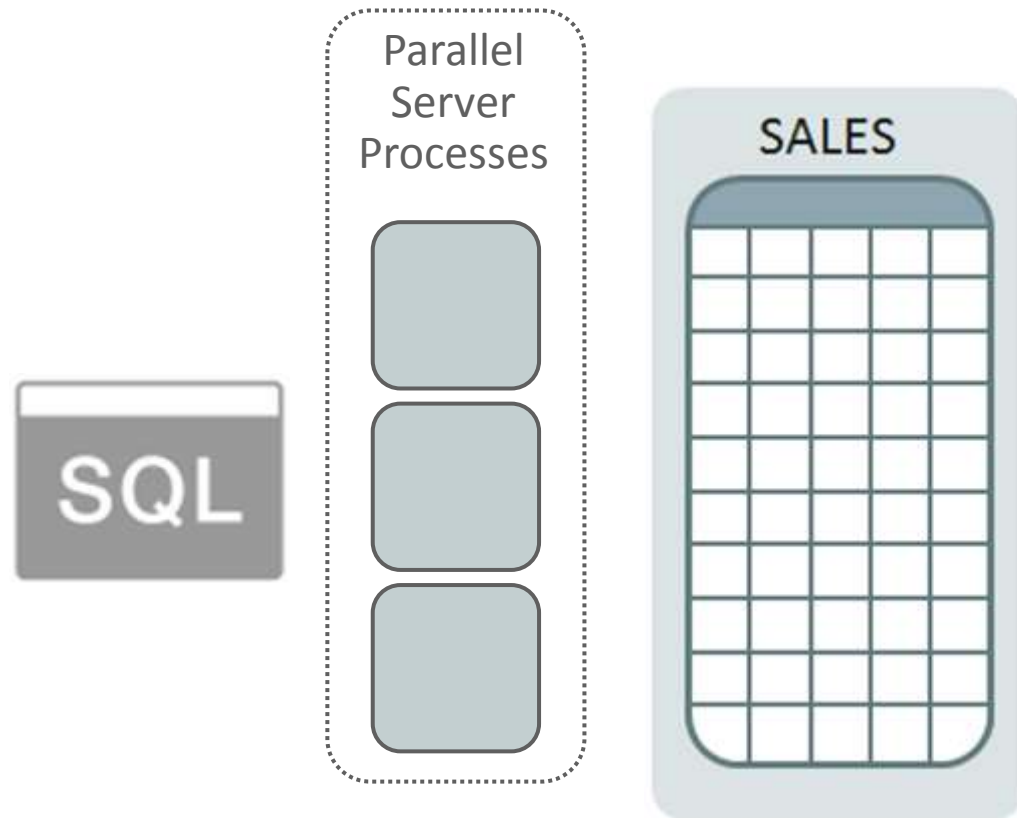
Best Query Processing Engine

Optimized processing of any SQL statement, with any data model

- Long-running strategic deep analysis
 - Ad-hoc queries over large volumes of data
 - Processed with high degrees of parallelism
- Near instantaneous tactical, short-running queries
 - Fast response times
 - Predefined reports – but potentially complex queries
- Concurrent data loading without compromises
 - Batch and trickle feeds
 - Ongoing query processing

Strategic Deep Analysis

Ad-hoc queries over large volumes of data

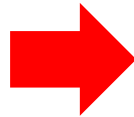


- Higher Performance and Scale-Out with Parallel Execution
 - Independent of partitioning
 - Fully integrated, “parallel everything”
 - Application transparent
 - Query, batch, data load, backup ...
 - Queries, DML, DDL, PL/SQL
 - Can leverage Real Application Clusters (RAC) for scale-out
 - Can leverage partitioning for performance

Strategic Deep Analysis

Ad-hoc queries over large volumes of data

What is the total sales figure for May 1-2?



SALES

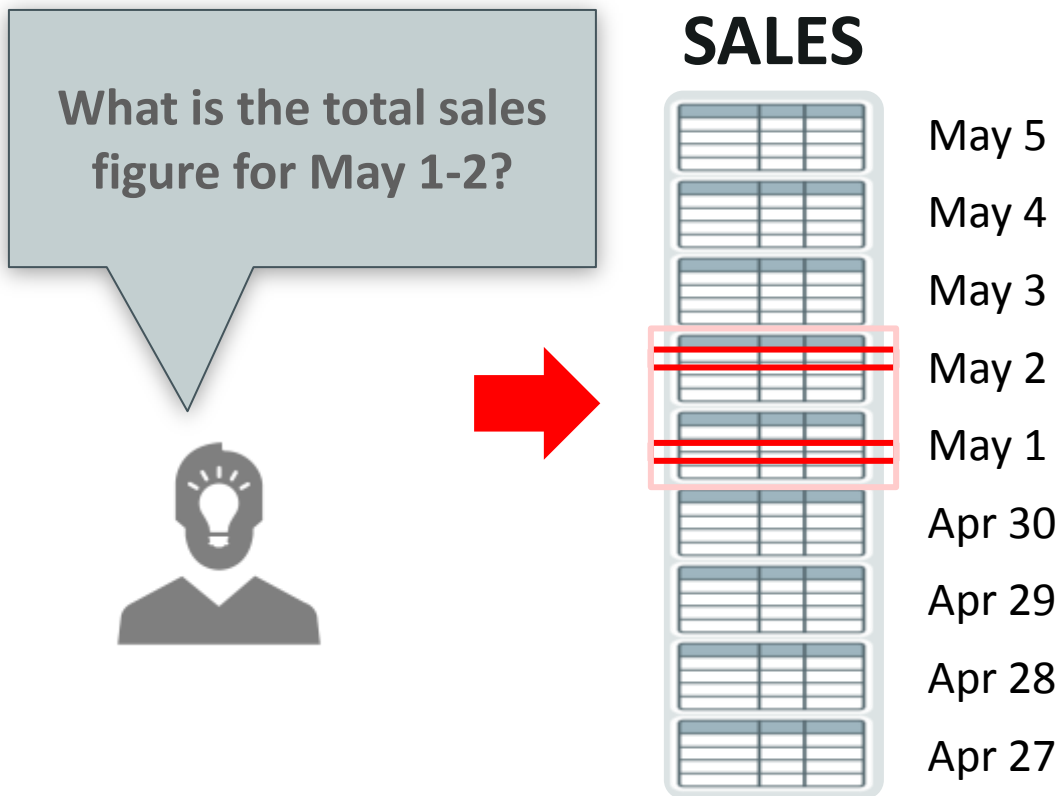
			May 5
			May 4
			May 3
			May 2
			May 1
			Apr 30
			Apr 29
			Apr 28
			Apr 27

- Partition elimination

- Dramatically reduces amount of data retrieved from storage
- Performs operations only on relevant partitions
- Transparently improves query performance and optimizes resource utilization

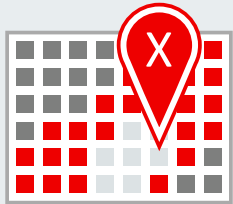
Strategic Deep Analysis

Ad-hoc queries over large volumes of data



- Partition elimination
 - Dramatically reduces amount of data retrieved from storage
 - Performs operations only on relevant partitions
 - Transparently improves query performance and optimizes resource utilization
- Zone map pruning
 - Further reduces amount of data retrieved from storage
 - For Oracle Engineered Systems
 - Transparently improves query performance and optimizes resource utilization

Zone Maps and Attribute Clustering



Zone maps

Stores min/max of specified columns per zone

Used to filter un-needed data during query execution



Attribute Clustering

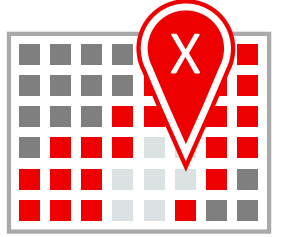
Orders data so that columns values are stored together on disk

Combined Benefits:

- Improved query performance and concurrency
 - Reduced physical data access
 - Significant IO reduction for highly selective operations
- Optimized space utilization
 - Less need for indexes
 - Improved compression ratios through data clustering
- Full application transparency
 - Any application will benefit

Zone Maps

Persisted storage index



- Stores minimum and maximum of specified columns
- Analogous to a coarse index structure
 - Much more compact than an index
 - Zone maps filter out what you don't need, indexes find what you do need
- Significant performance benefits with complete application transparency
 - IO reduction for table scans with predicates on the table itself or even a joined table using join zone maps (a.k.a. “hierarchical zone map”)
 - Partition pruning for every column of a partitioned table, not only the partition key columns
- Benefits are most significant with ordered data
 - Used in combination with attribute clustering or data that is naturally ordered



Attribute Clustering

Concept and Benefits

- Orders data so that it is in close proximity based on selected columns values: “attributes”
 - Attributes can be from a single table or multiple tables
 - e.g. from fact and dimension tables
 - Able to cluster data during MOVE PARTITION
- Benefits
 - Significant IO pruning when used with zone maps
 - Also, reduced block IO for table lookups in index range scans
 - Improved performance for queries that sort and aggregate pre-ordered data
 - Improved compression ratios
 - Ordered data is likely to compress more than unordered data

Zone Maps and Attribute Clustering

Example

- CLUSTERING BY INTERLEAVED ORDER (category, country)
- Zone map benefits are most significant with ordered data

		Country			
Category	10	11	14	15	
	AR WOMEN	JP WOMEN	SA WOMEN	US WOMEN	
	8	9	12	13	
	AR MEN	JP MEN	SA MEN	US MEN	
	2	3	6	7	
	AR GIRLS	JP GIRLS	SA GIRLS	US GIRLS	
	0	1	4	5	
	AR BOYS	JP BOYS	SA BOYS	US BOYS	

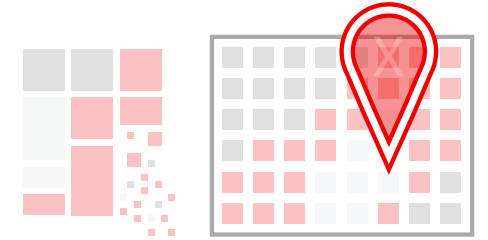
INTERLEAVED ORDER

Pruning with:

```
SELECT ..  
FROM table  
WHERE category = 'BOYS' ;
```

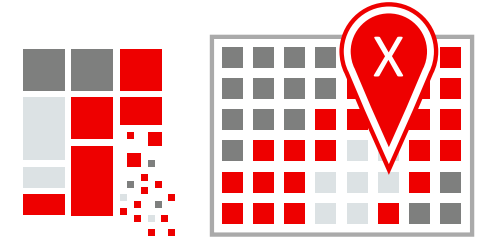
```
SELECT ..  
FROM table  
WHERE country = 'US'
```

```
SELECT ..  
FROM table  
WHERE category = 'BOYS' ;  
AND country = 'US'
```



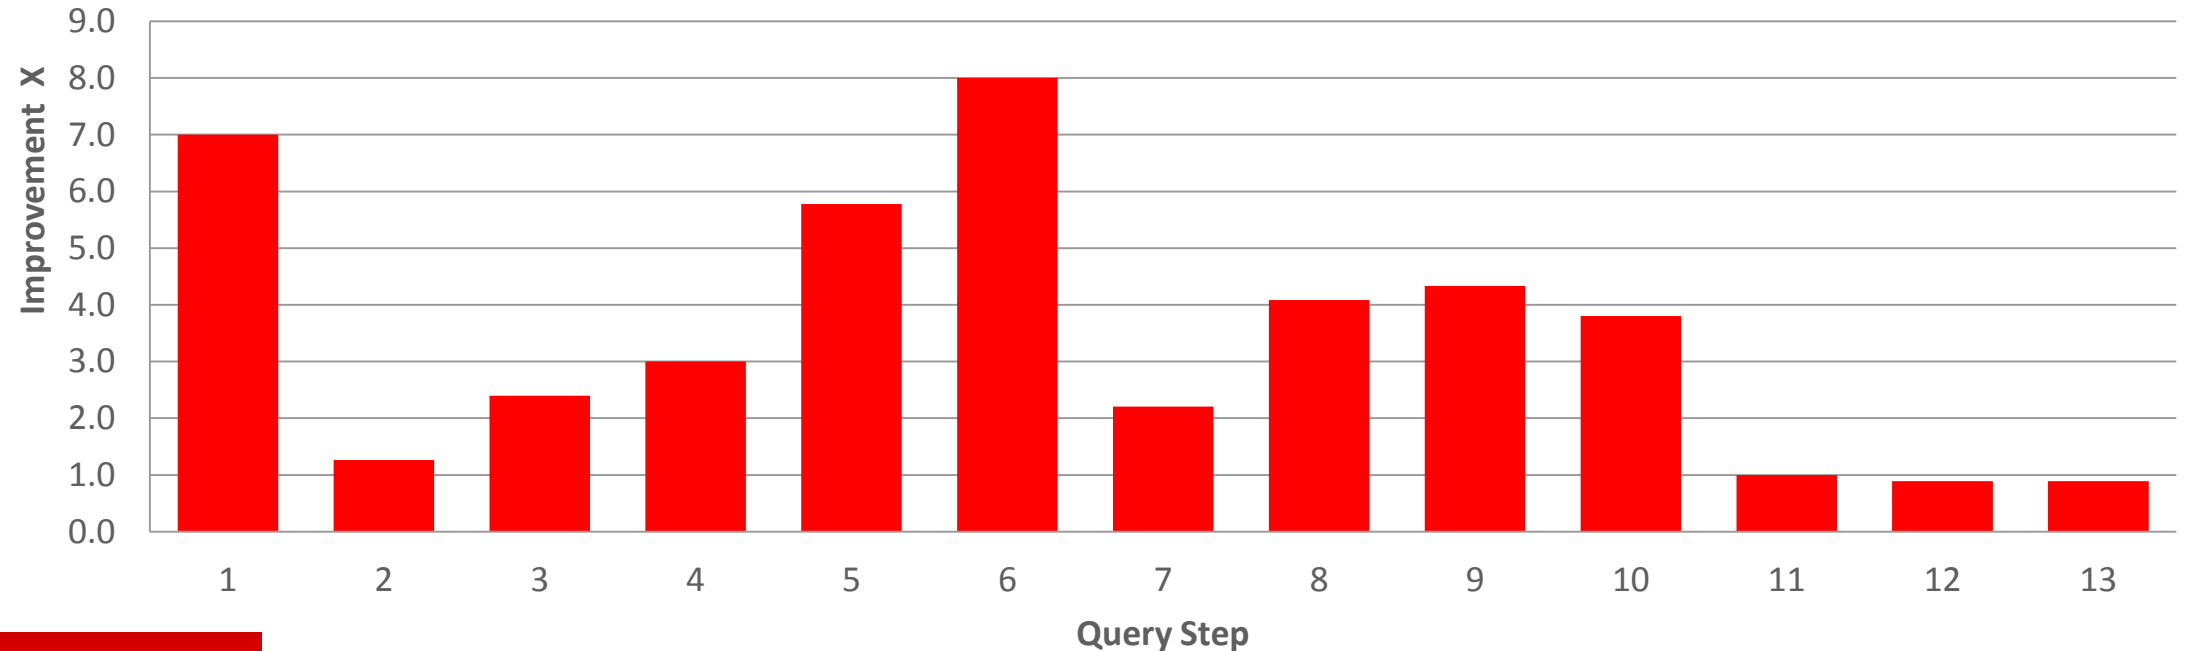
Zone Maps with Attribute Clustering

Star Schema Benchmark

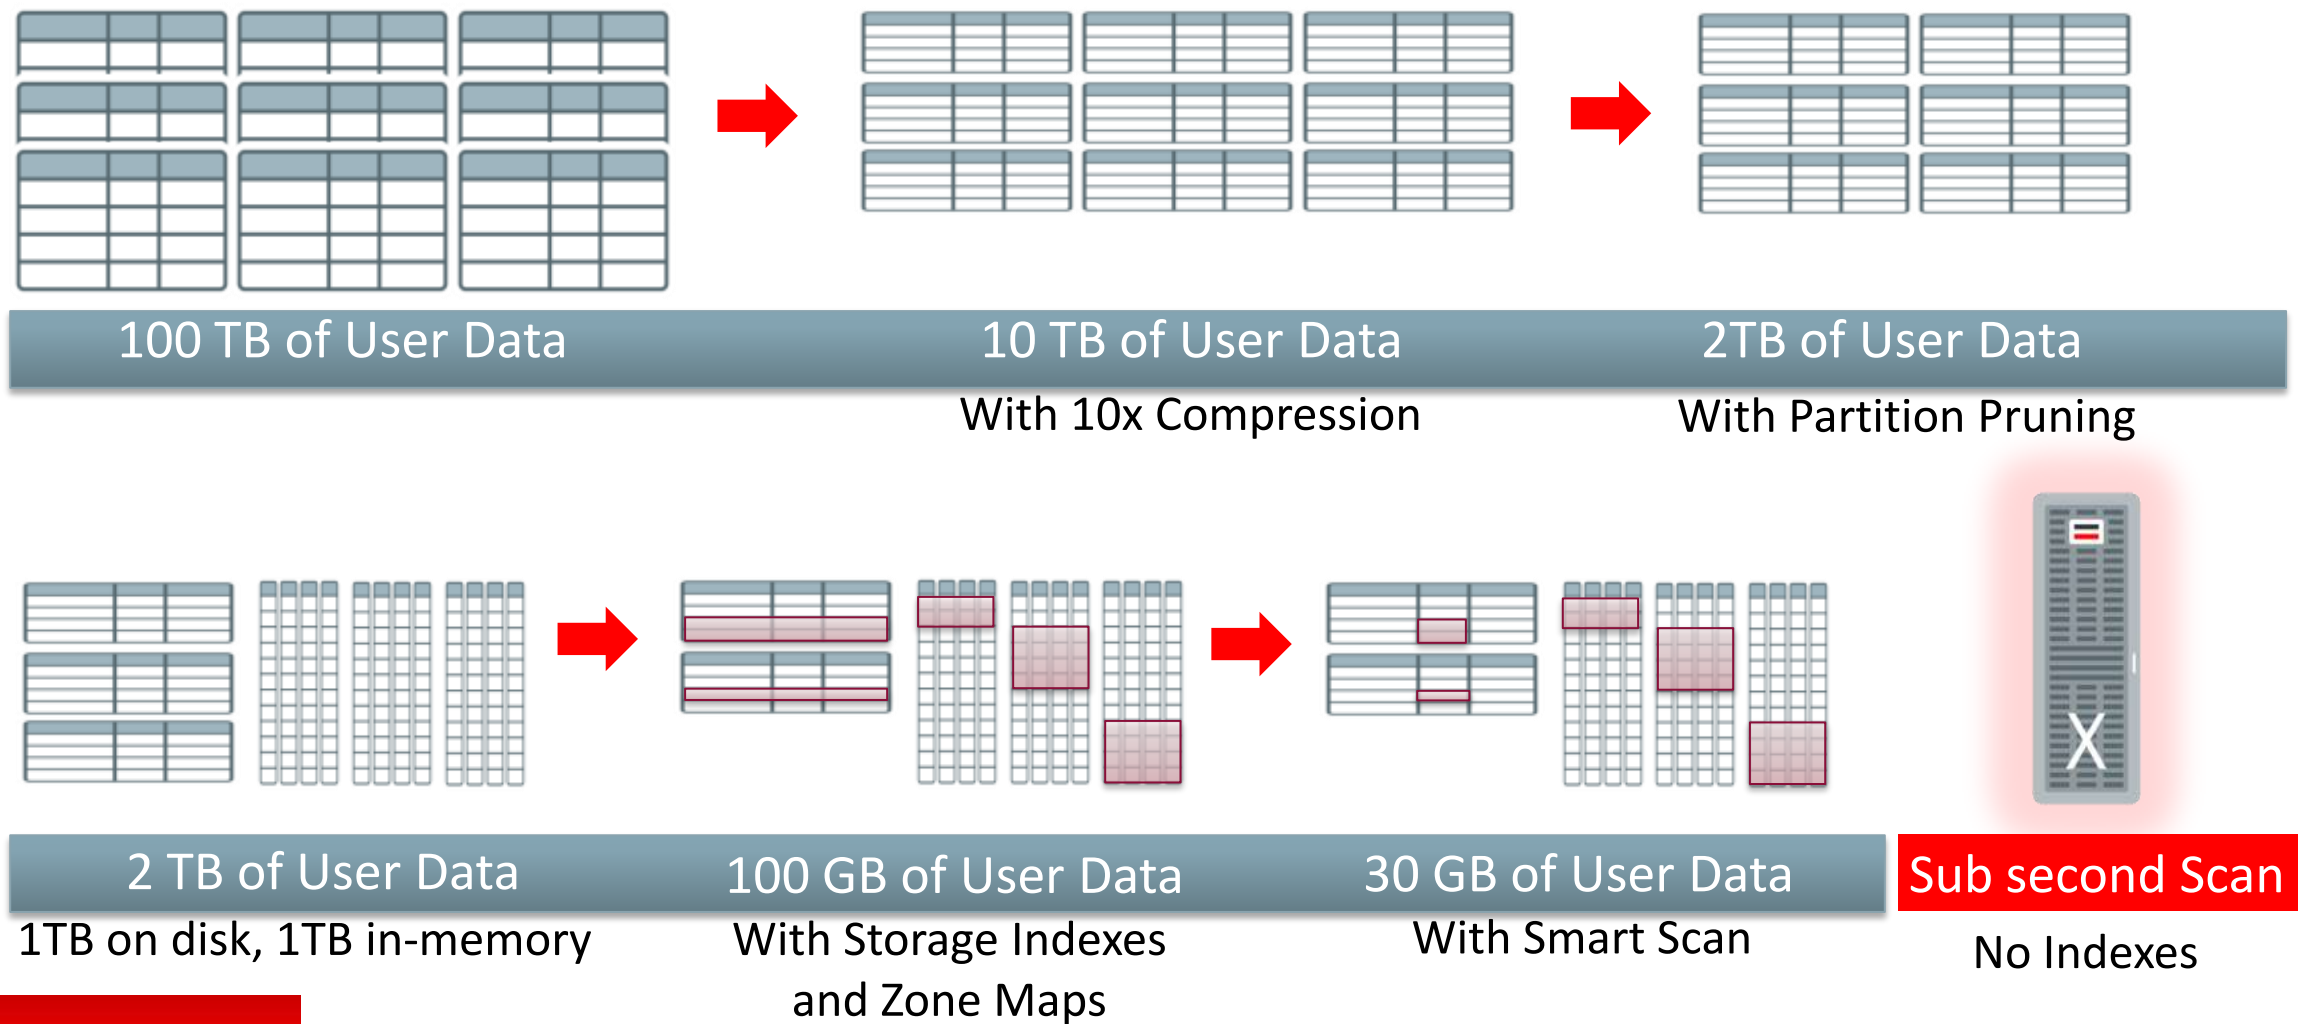


- Overall, 2.6X elapsed time improvement over baseline
 - Comparing with and without zone map and attribute clustering

Query Elapsed Time Improvements



Performance Features Multiply the Benefits



Operational Data Warehousing

Near instantaneous tactical, short-running queries

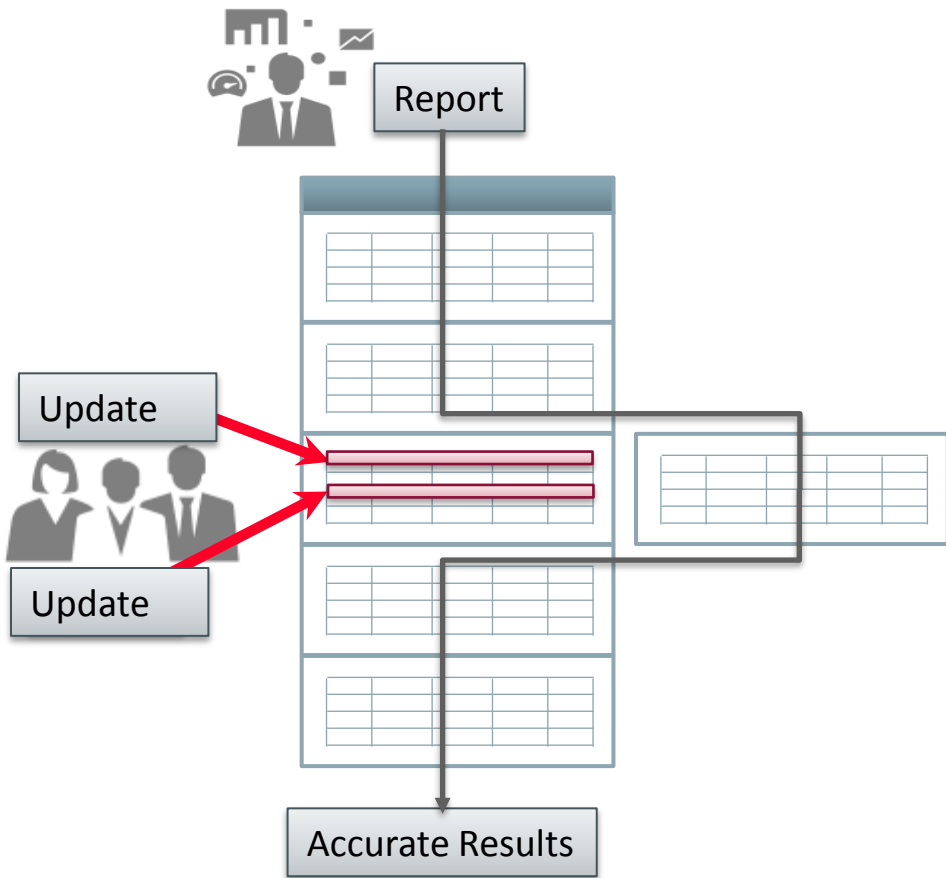


- Aggregate management
 - Query re-write to transparently utilize materialized views
- Dimensional schema optimization
 - Star-schema query optimization with bitmap indexing
- Adaptive query optimization
 - Adjusts query execution at run time

Real-World Loading and Access

- Round-the-clock loading with concurrent querying
- Strategic and operational warehousing requires ongoing data refresh
 - Trickle and bulk feed
- Data ingestion must not impact data usage
 - Writers block readers is unacceptable
- Key technologies
 - Continuous read-consistent, non-blocking access
 - Efficient high-performant data ingestion and ETL capabilities

Oracle Multi-Version Read Consistency



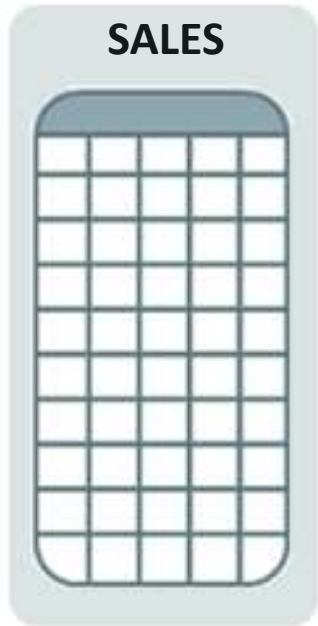
- Unique architecture built for concurrency
 - Readers don't block writers
 - Writers don't block readers
- Maximum throughput with correct results
 - See only committed data via Multi-Versioning
 - No waiting and no dirty reads!
- No lock escalation
 - Scalable row level lock management

Efficient Database Management

Petabyte Data Management

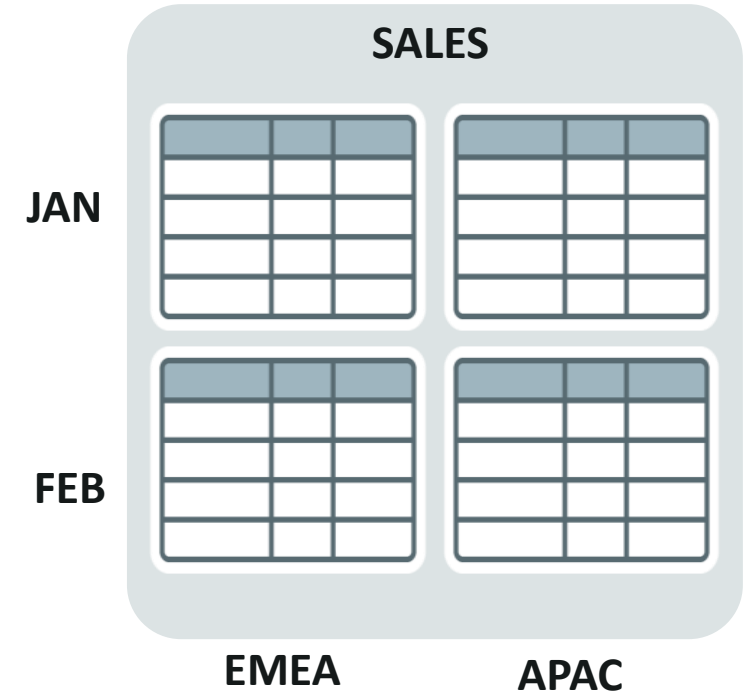
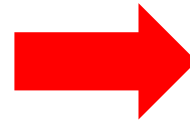
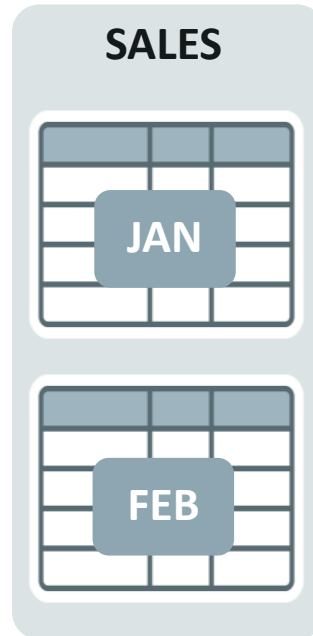
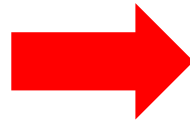
Maintain Consistent Performance as Database Grows

Oracle Partitioning



Challenges:

Large tables are difficult to manage

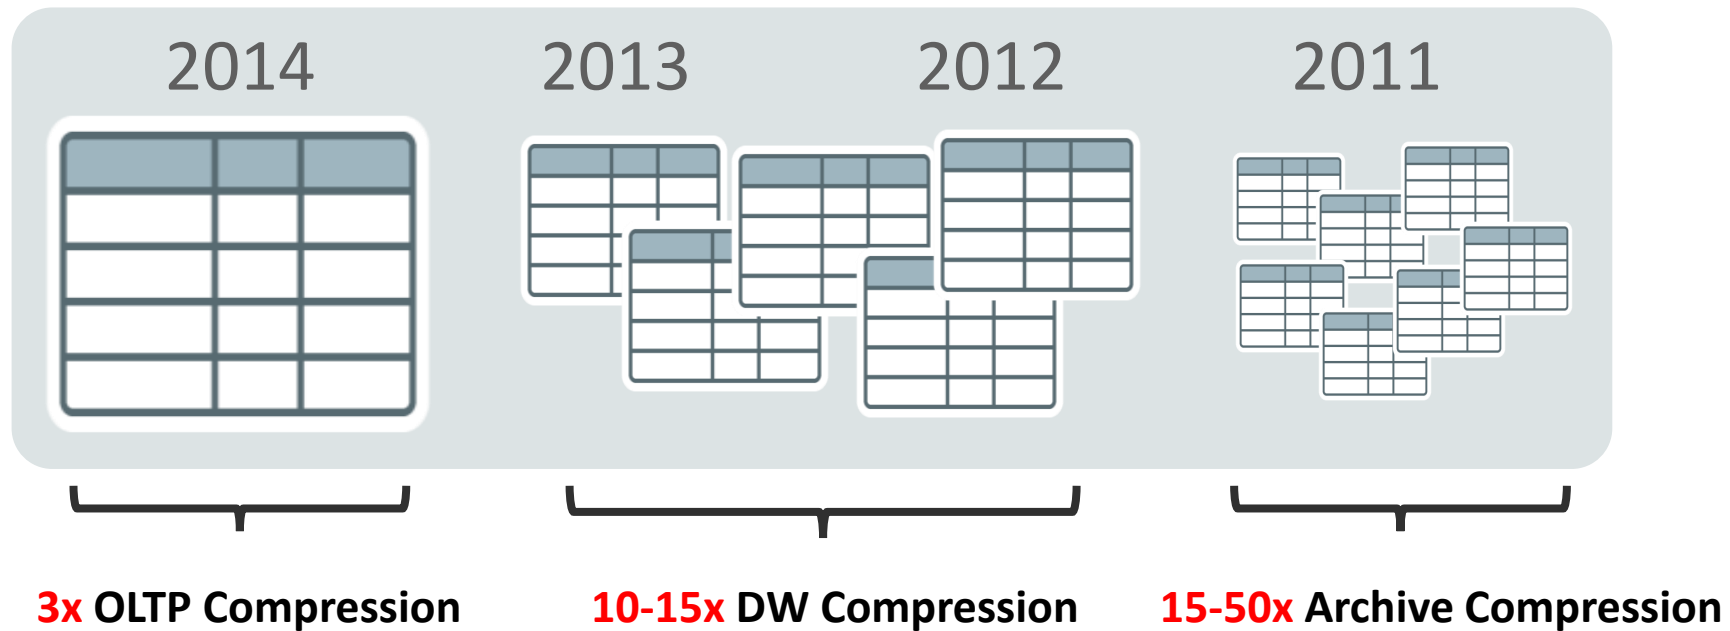


Solution: Partitioning

- Divide and conquer
- Easier data management
- Improve performance

Optimized Data Storage

Distribute partitions across multiple compression tiers



Free up storage space and execute queries faster

Automatic Data Optimization

Usage based data compression

01110101010010
10000100010101

Hot Data

10101010111010100110101110
00010100010110111010101001
01001001000010001010101101
00101101001110000101001001
01000010010000100010101011
10011010

3X

Advanced Row Compression

Warm Data

1010101011101010011010111000010100010110
11101010100101001001000010001010110100
1011010011100001010010010100001001000010
0010101011010010
1000010100100101001010110111000010
1110010100100101001010110111011010

10X

Column Query Compression

Archive Data

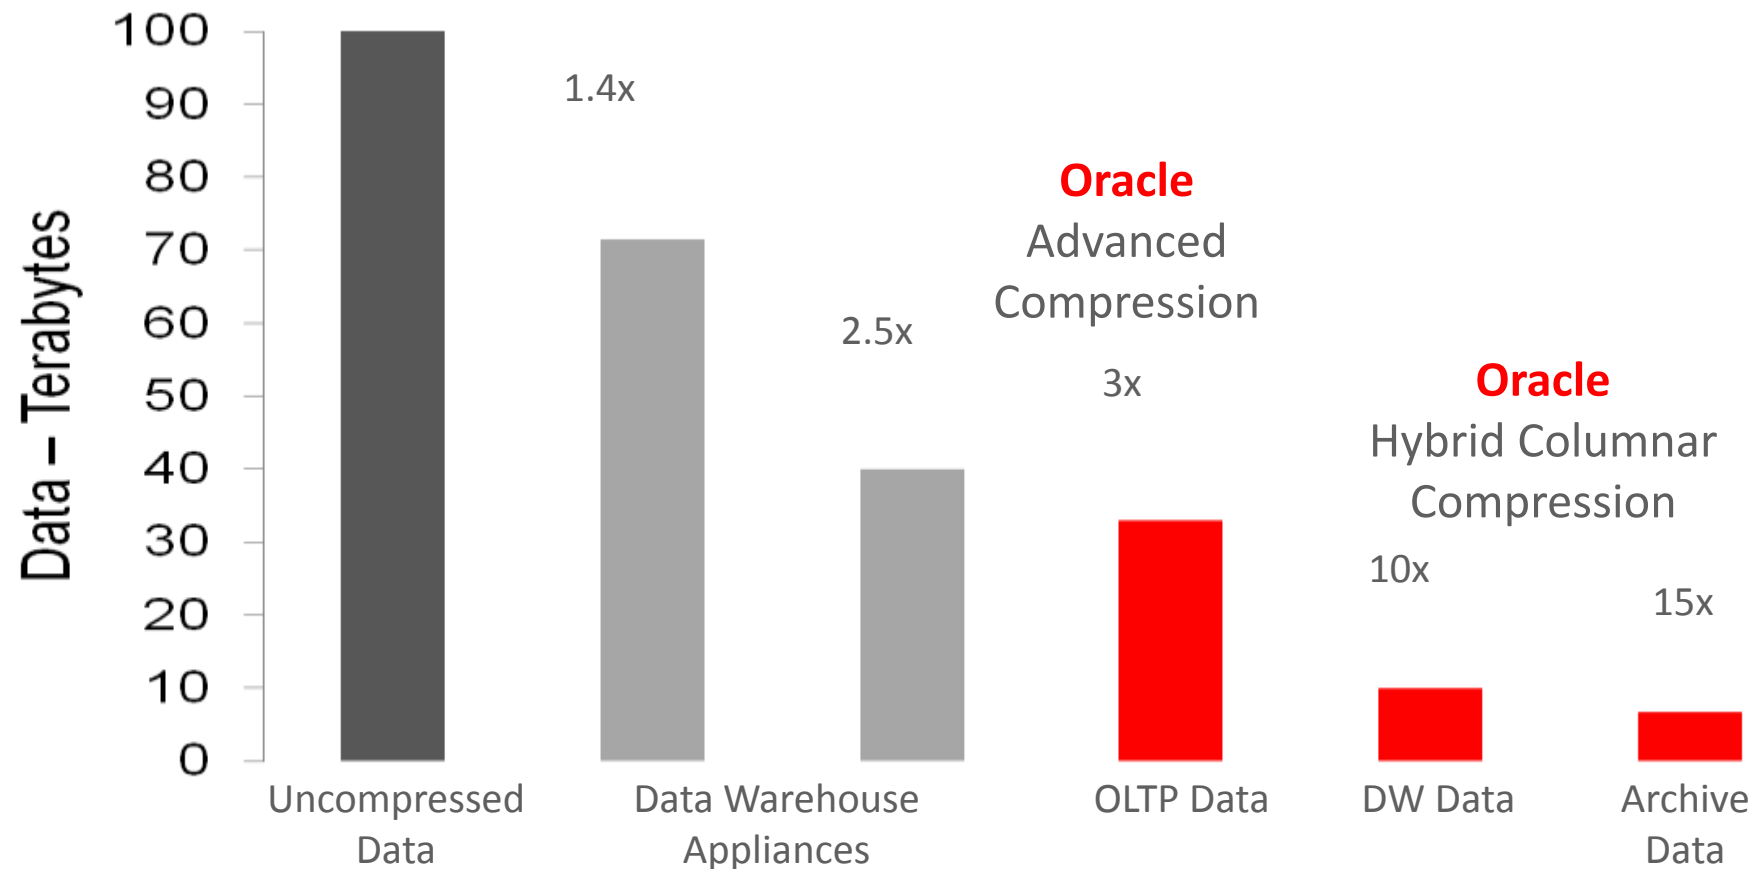
1010101011101010011010111000010100010110110
101010010100100100001000101010110100101100
11100001010010010100001001000010001010
101010101110101001101011100001010001011011
101010101110101001101011100001011101011001

15X

Column Archive Compression

Compression

Reduce Disk Space Requirements

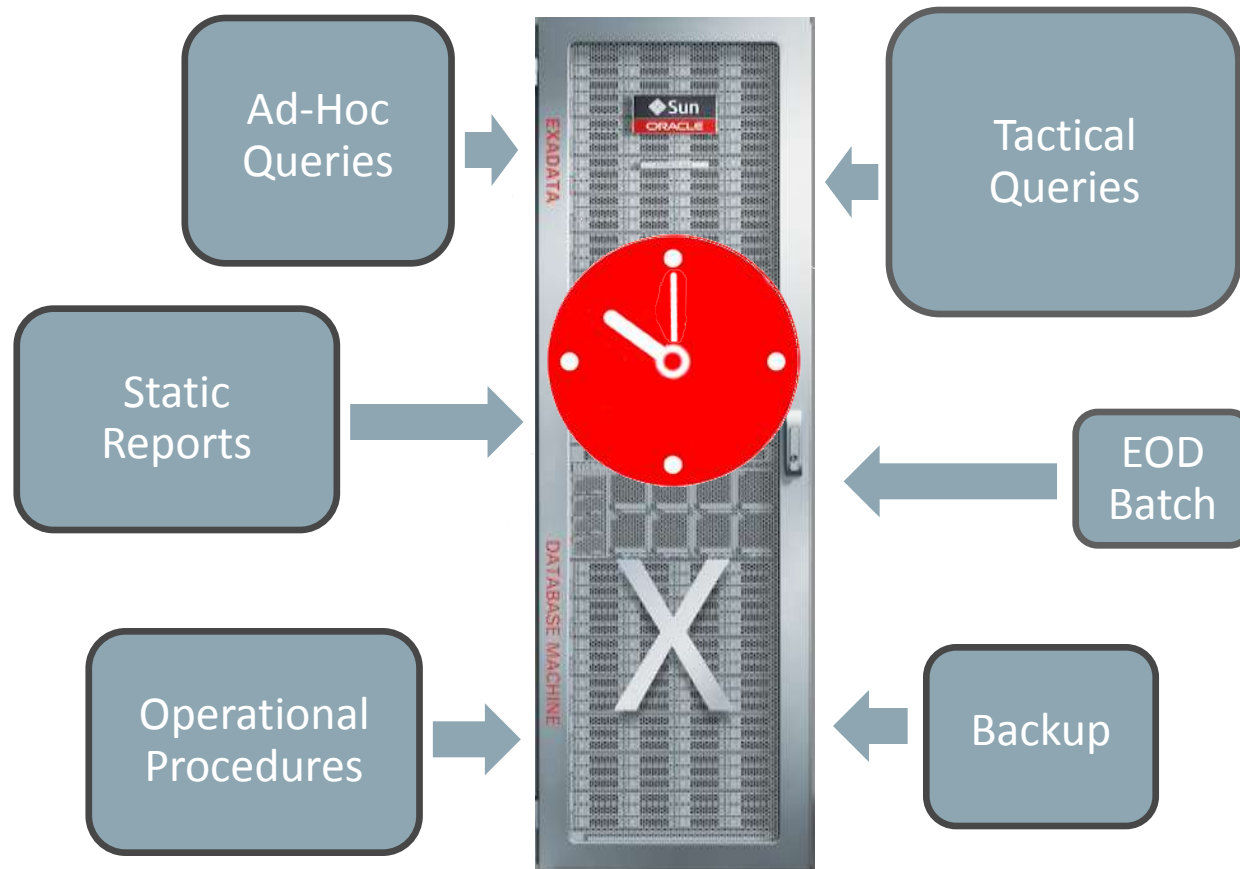


Efficient Database Management

Automated Workload Management

Enterprise-Class Workload Management

Database Resource Manager (DBRM) and IO Resource Manager (IORM)

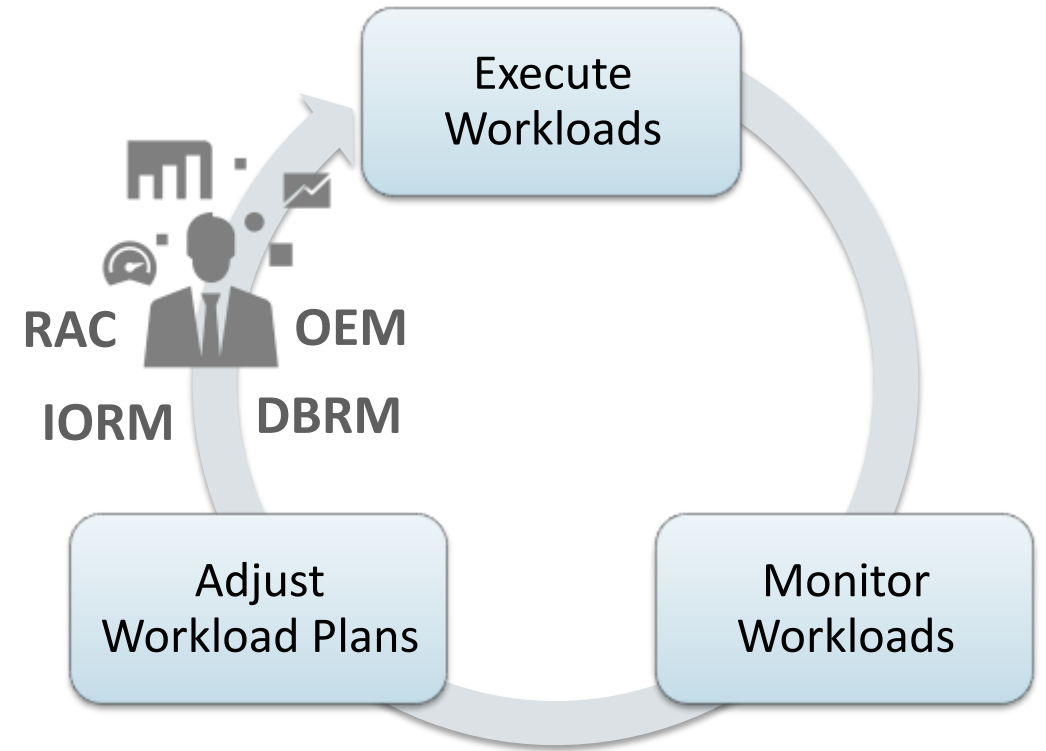


- Manage Competing Resources
 - Resource guarantees
 - Usage limits
- Resource Managers Control:
 - CPU
 - Exadata disk I/O
 - Runaway queries
 - Parallel execution
 - Idle time

Enterprise-Class Workload Management

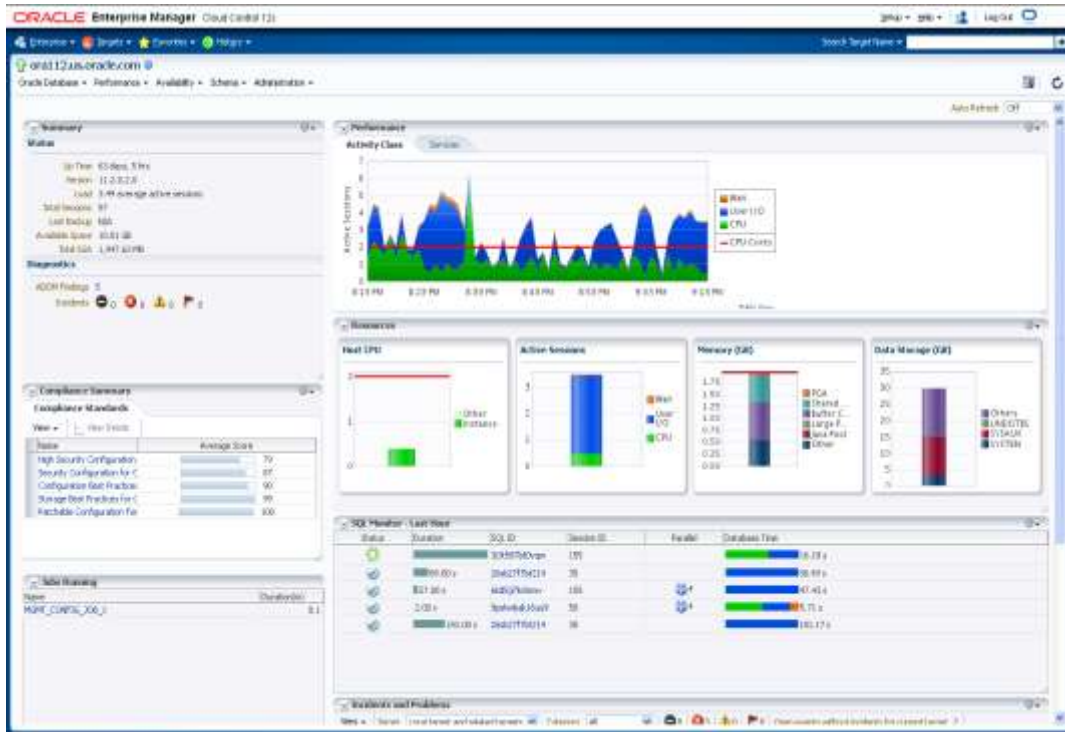
Continuous Optimization

- Workload based resource allocation
 - Workloads should use critical system resources according to their priority
- Proper prioritization
 - Tactical workload must run with guaranteed resource allocation
- Full utilization of critical resources
 - Avoid idle resources
- Flexible resource allocation
 - E.g. priority of ETL is based on time of day
- Monitor Workloads

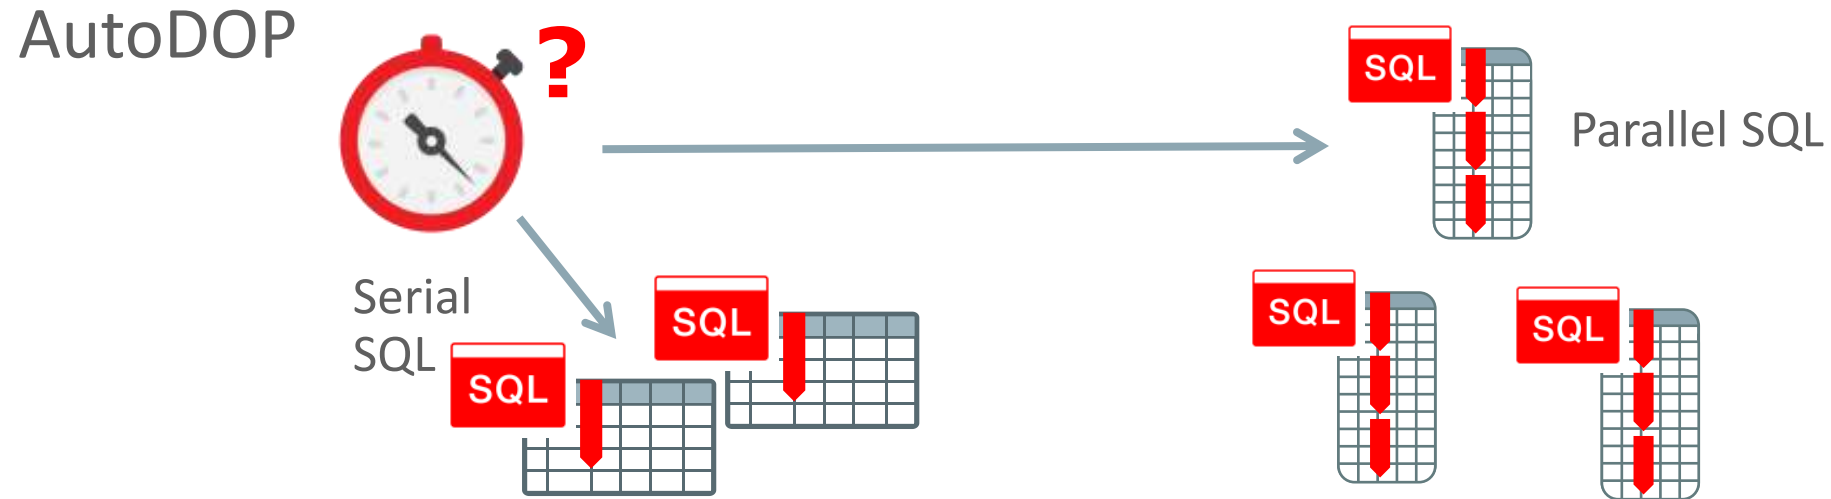


Comprehensive Workload Management and Monitoring

- Workload management; monitoring and configuration
- Monitoring application workloads
 - ETL operations, quarter end close jobs
- Real-time monitoring, driven by application specified tagging
- Historic resource usage
- Visibility of top SQL, system and session performance metrics



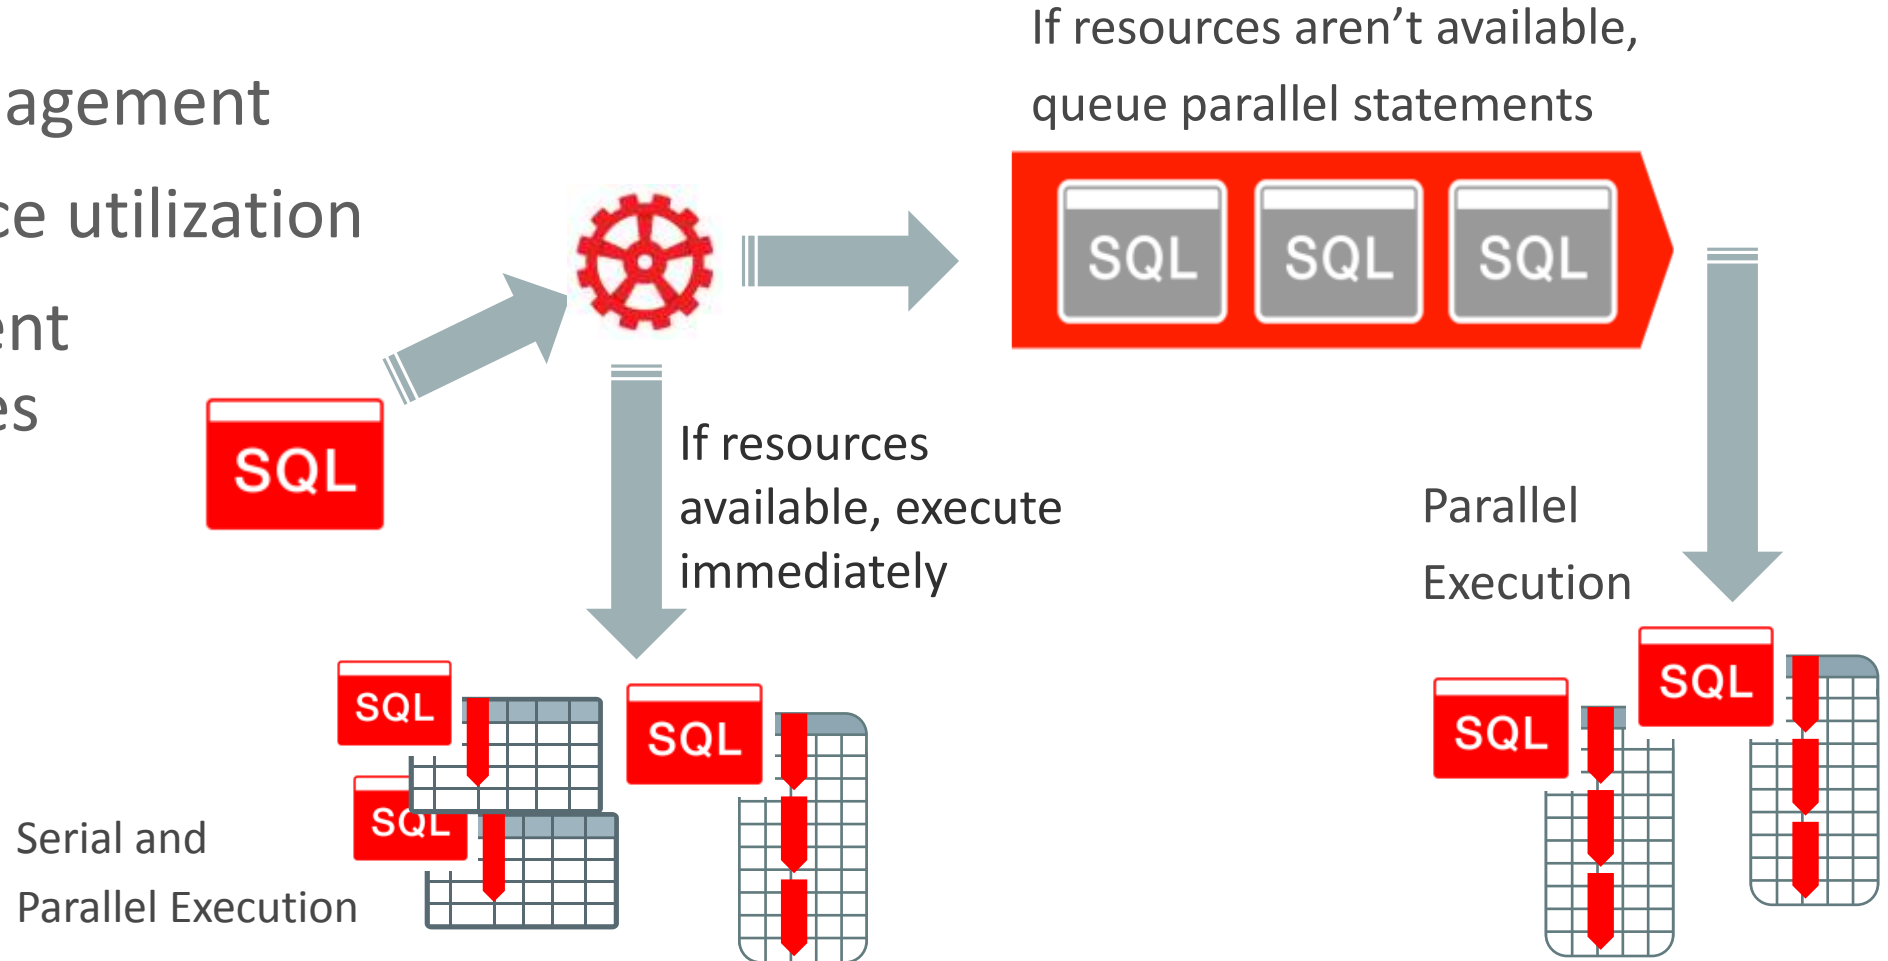
Automatic Degree of Parallelism (Auto DOP)



- Optimizer derives the best degree of parallelism (DOP)
- Based on resource requirements and linear scalability
- Less DBA management, better resource utilization

Parallel Statement Queuing

- Less DBA management
- Better resource utilization
- More consistent response times

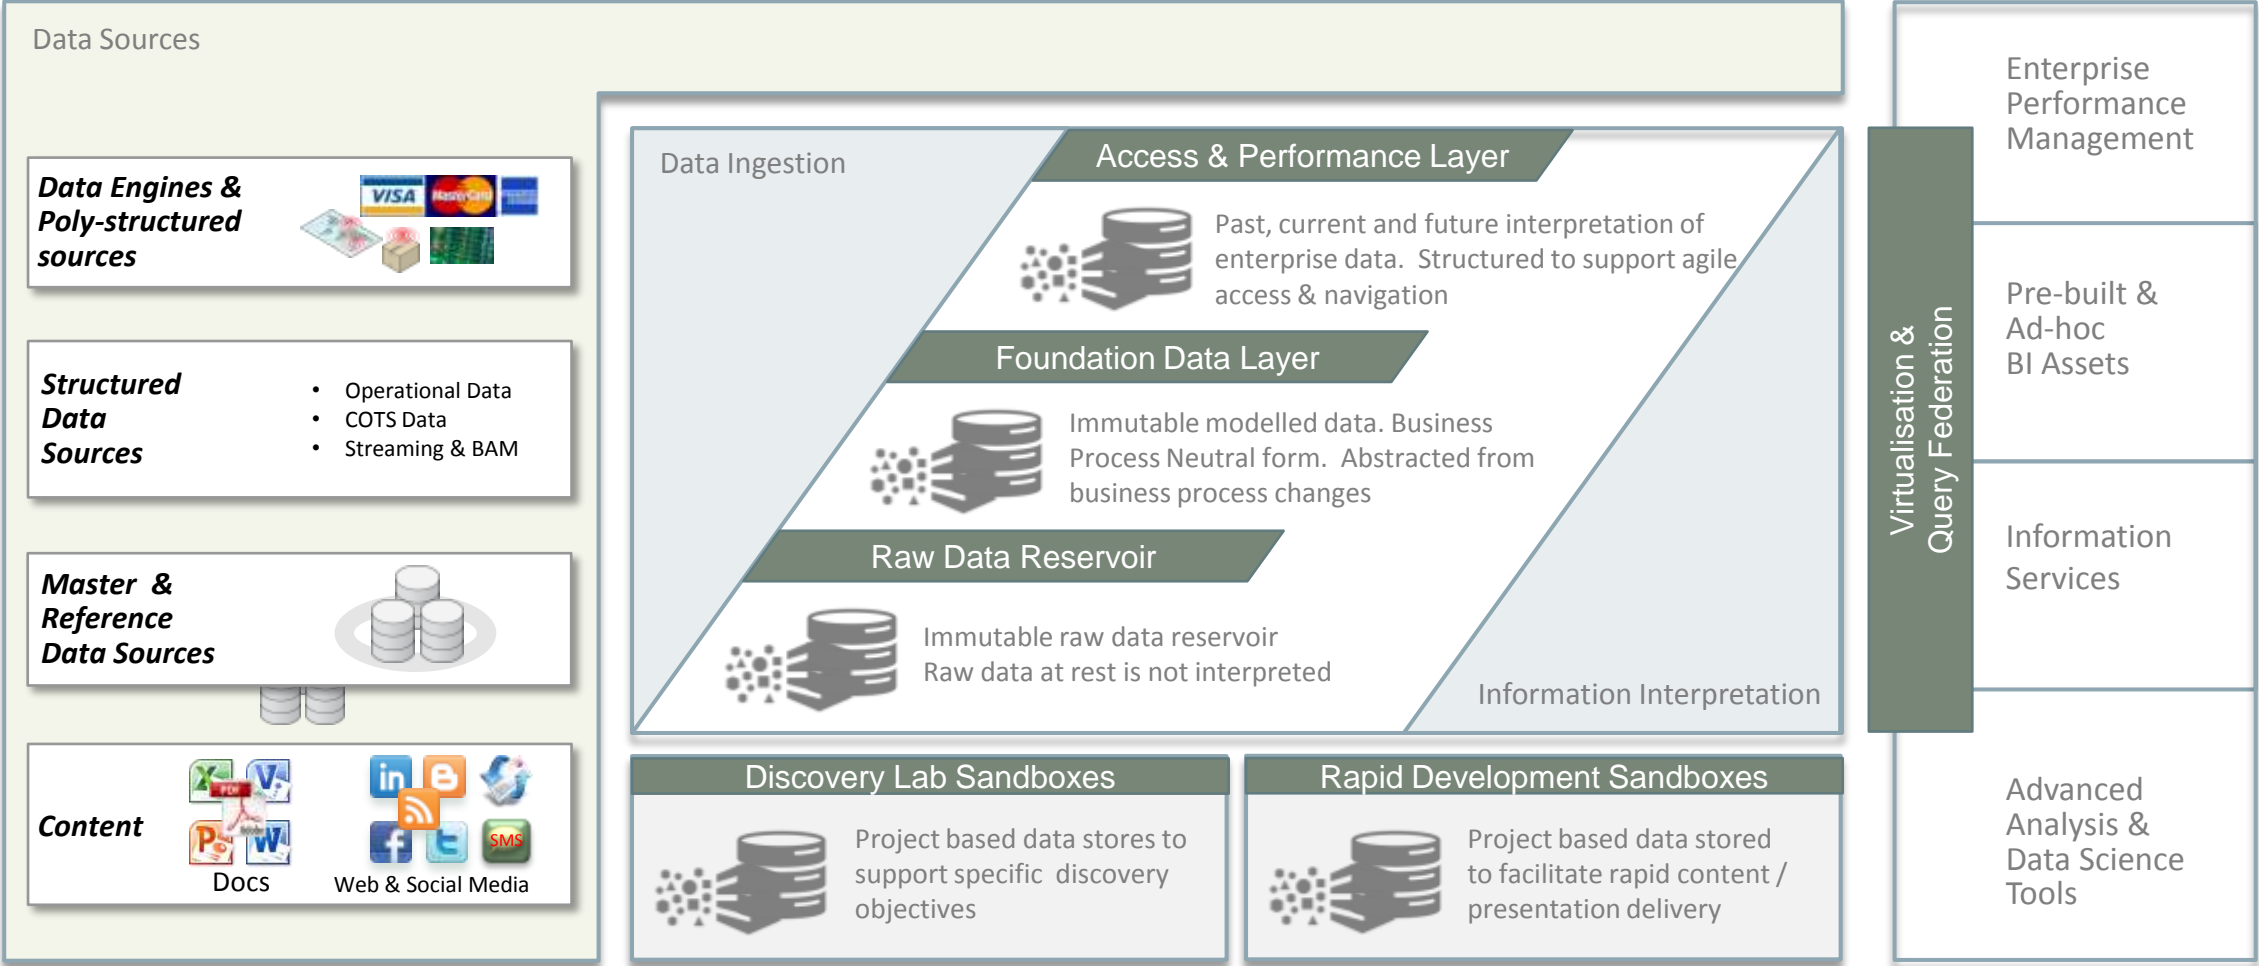


Efficient Database Management

One Platform for all your Data Warehousing Applications

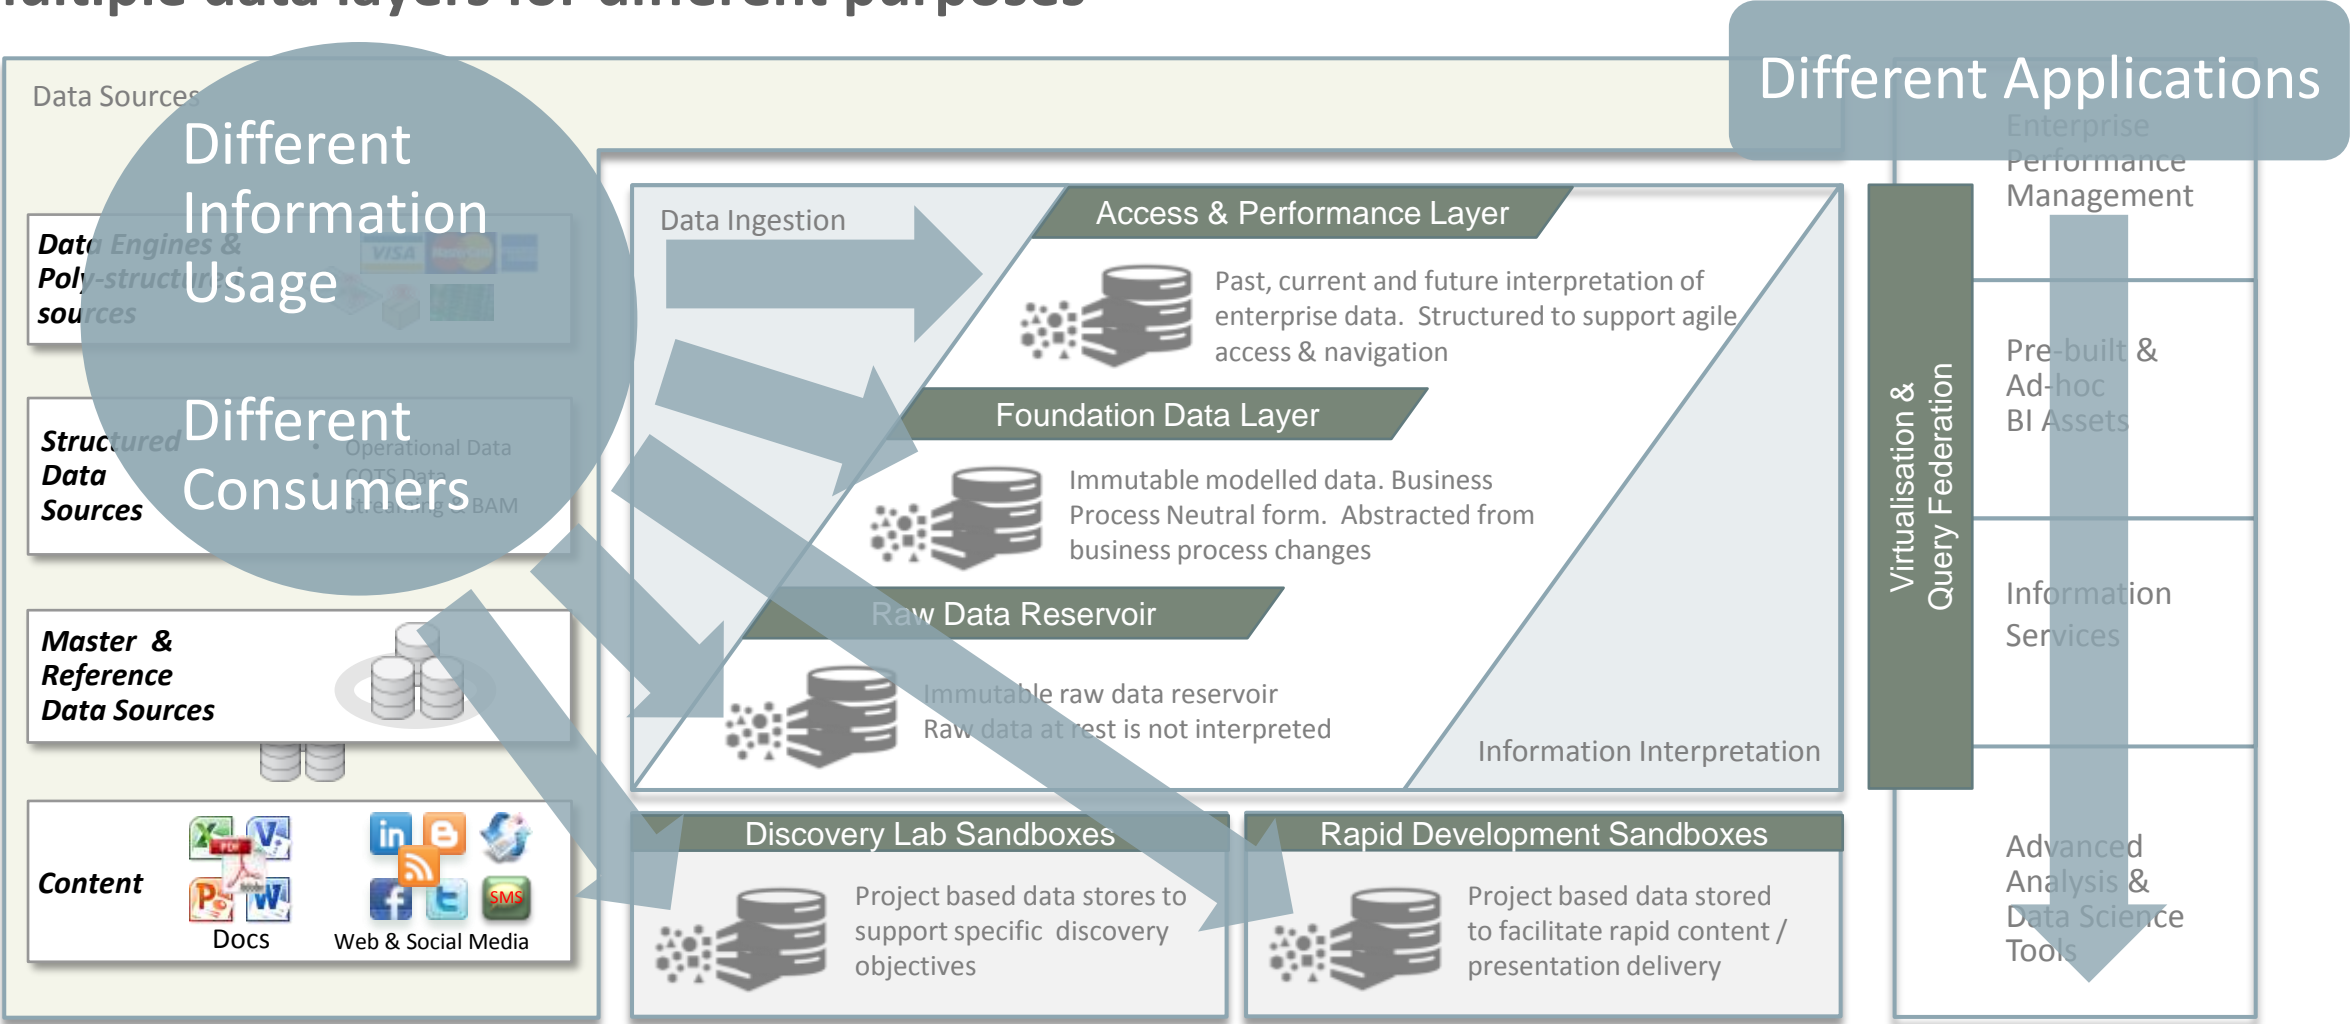
Data Warehousing Information Management Architecture

Multiple data layers for different purposes



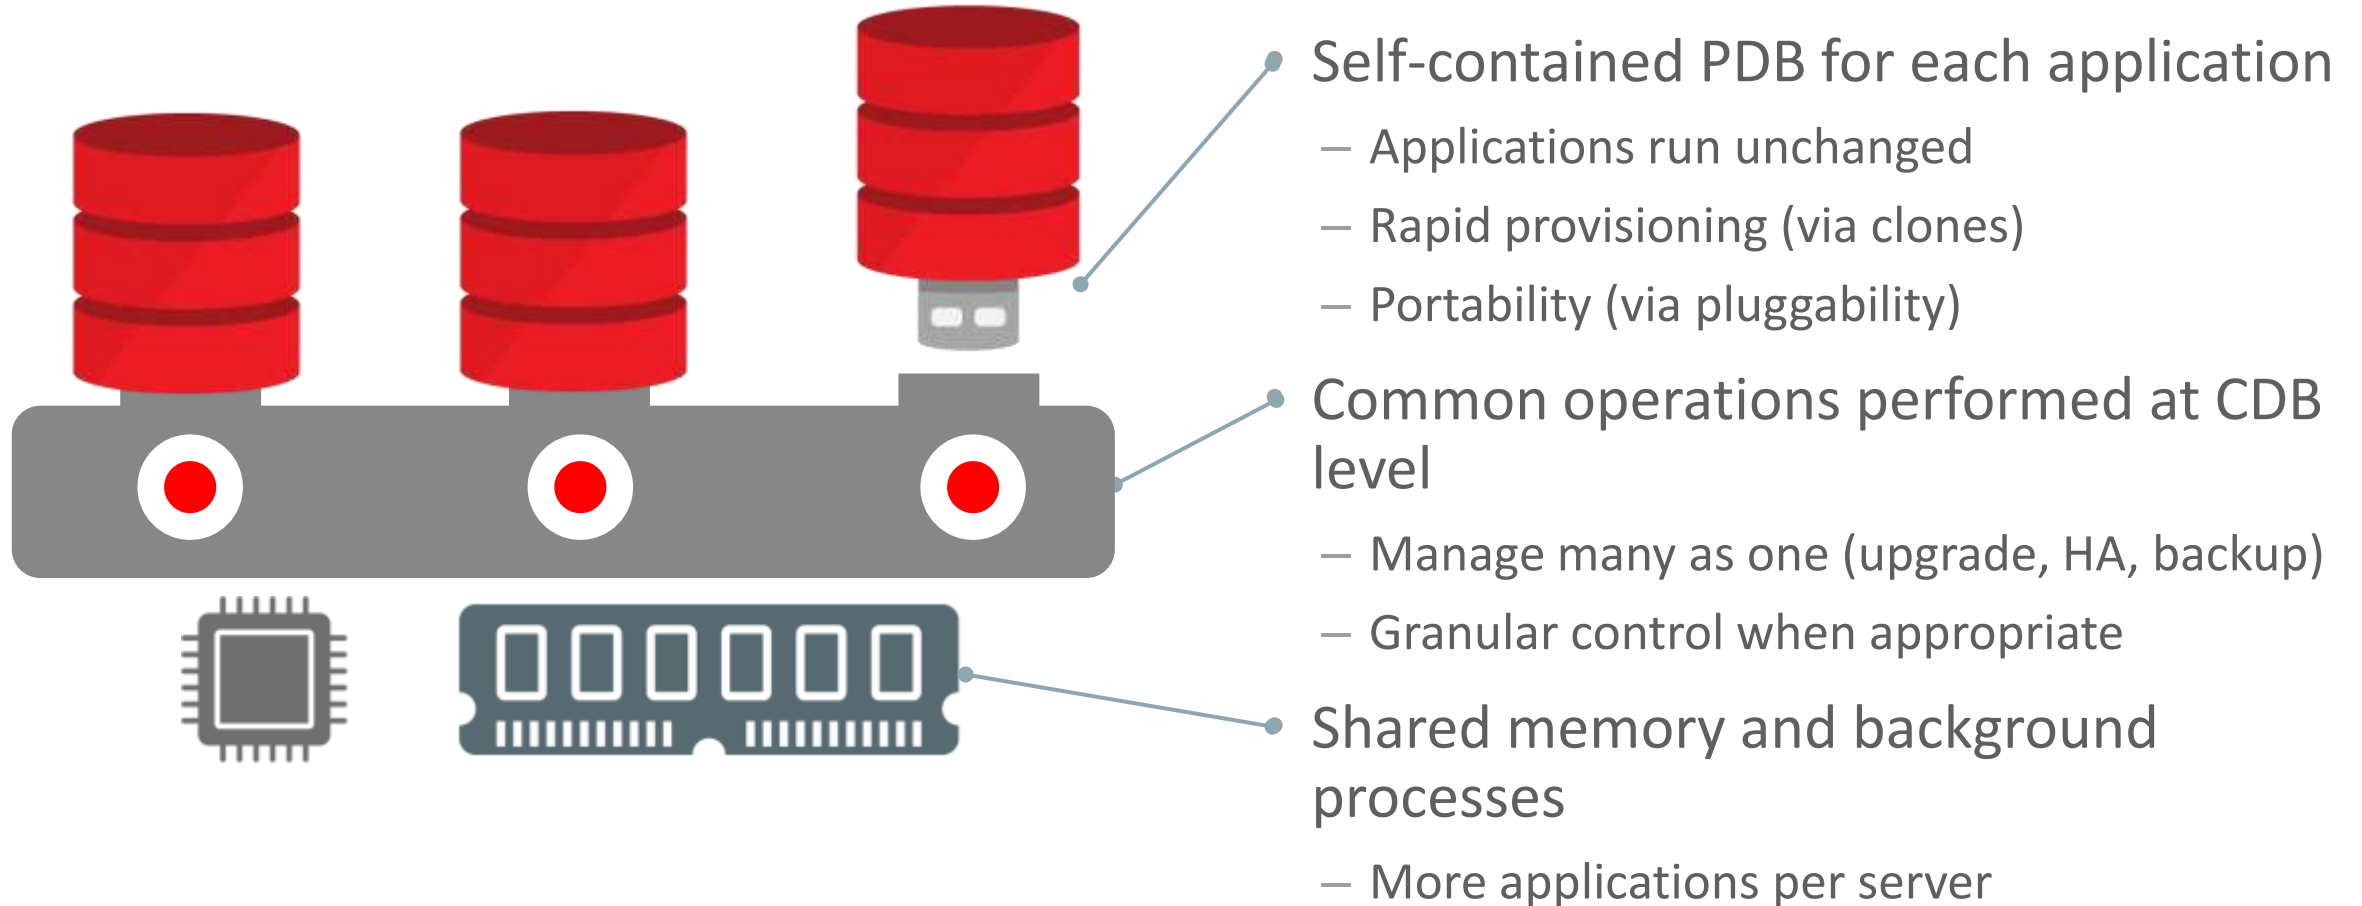
Data Warehousing Information Management Architecture

Multiple data layers for different purposes



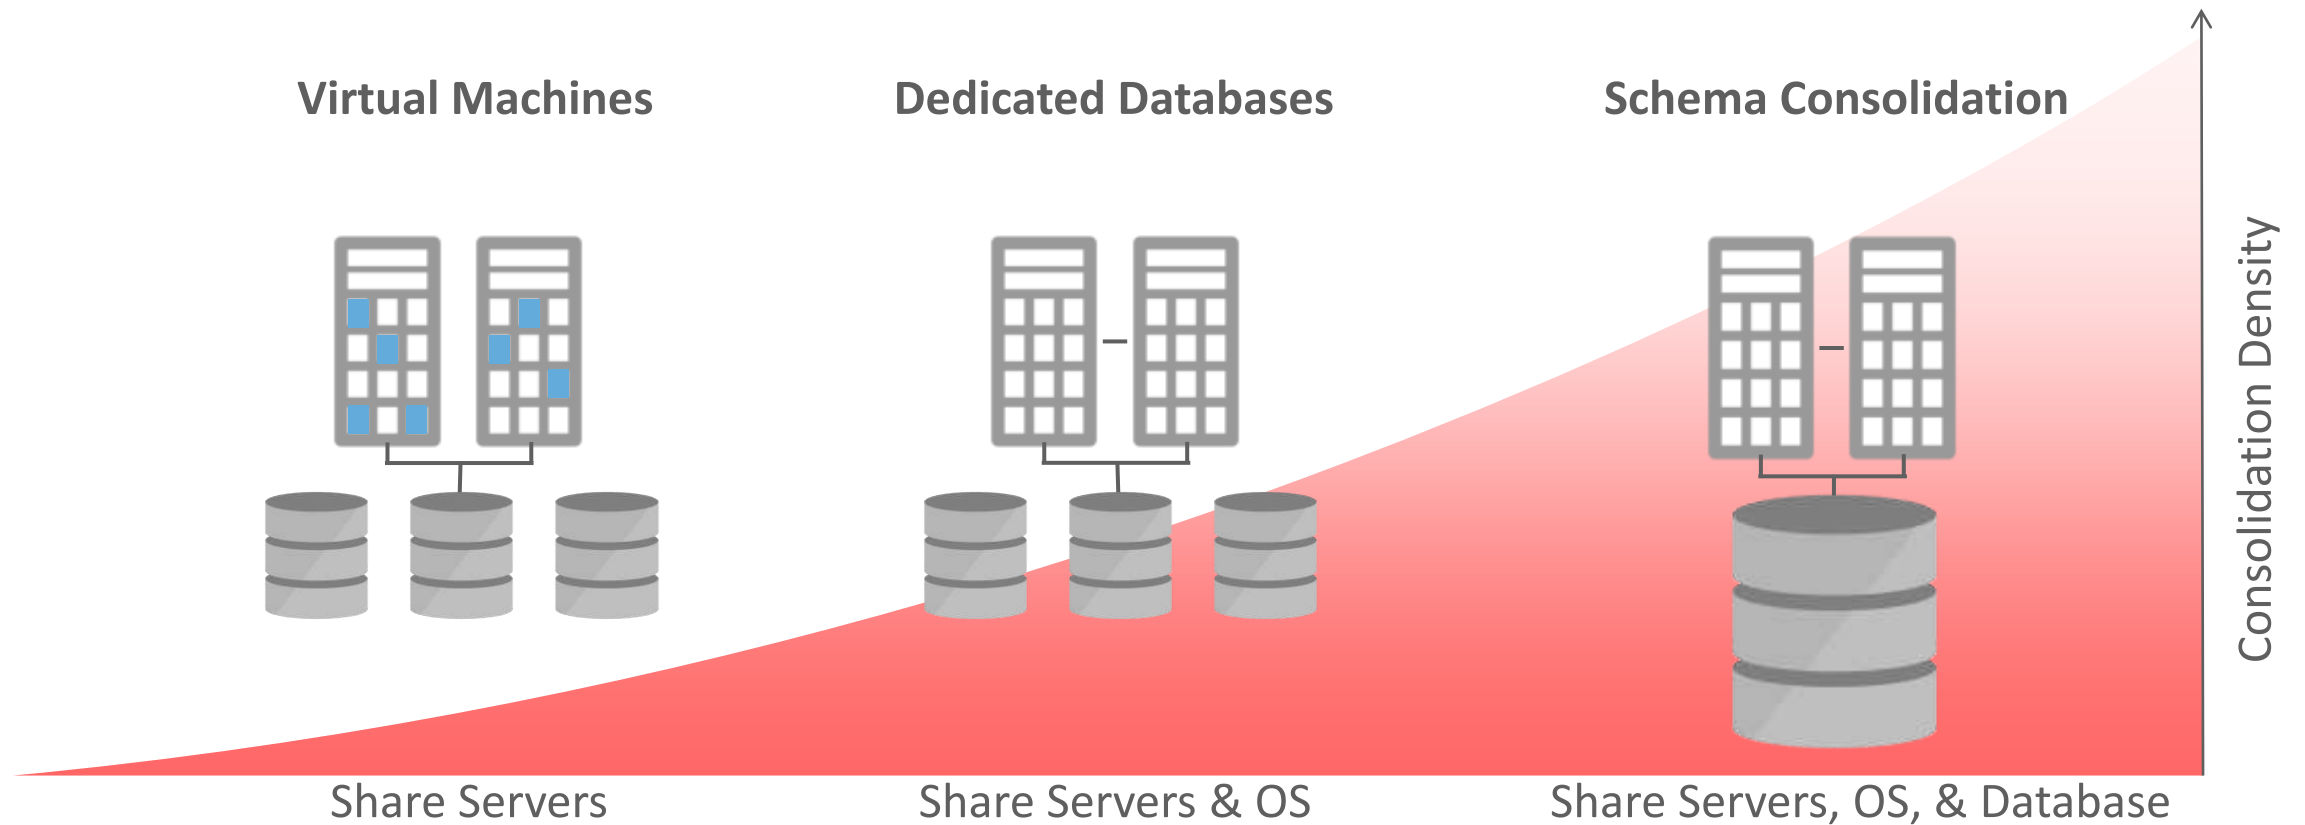
Oracle Multitenant

Consolidate databases and applications - simplify operations



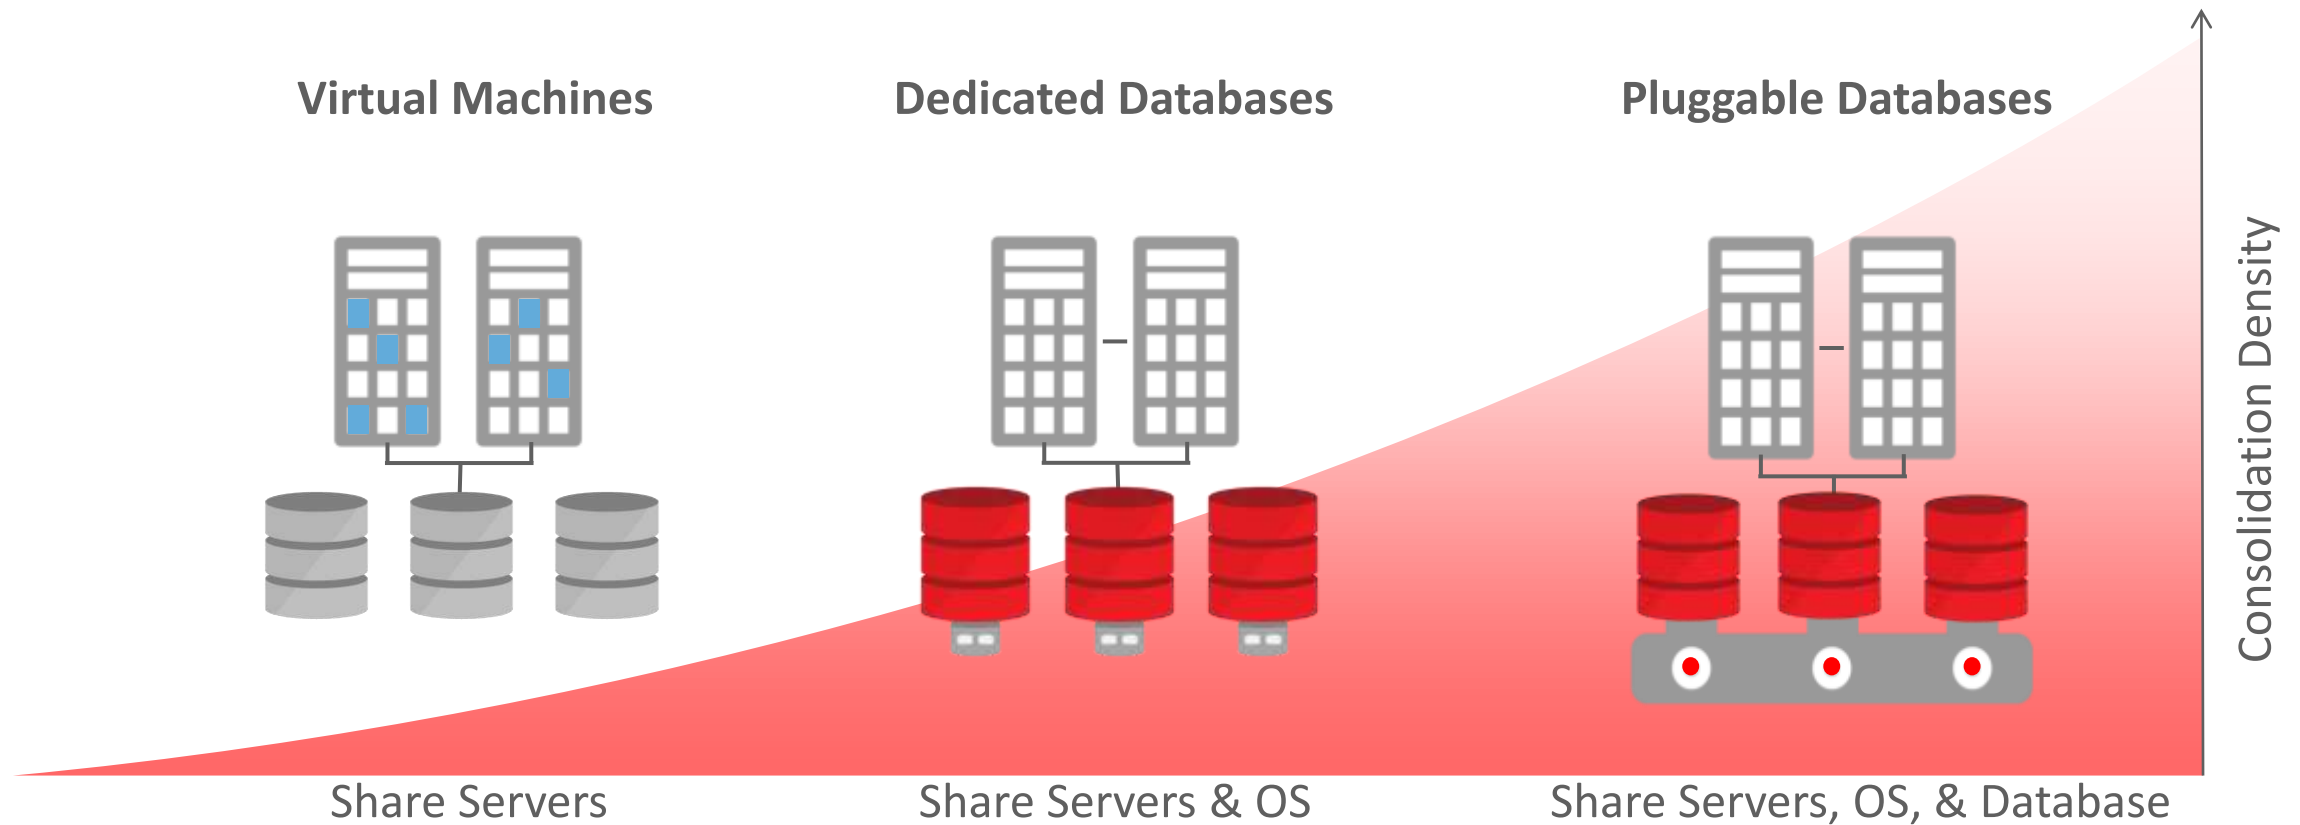
Database Consolidation

Traditional consolidation methods



Oracle Multitenant

High consolidation density, transparent to existing applications



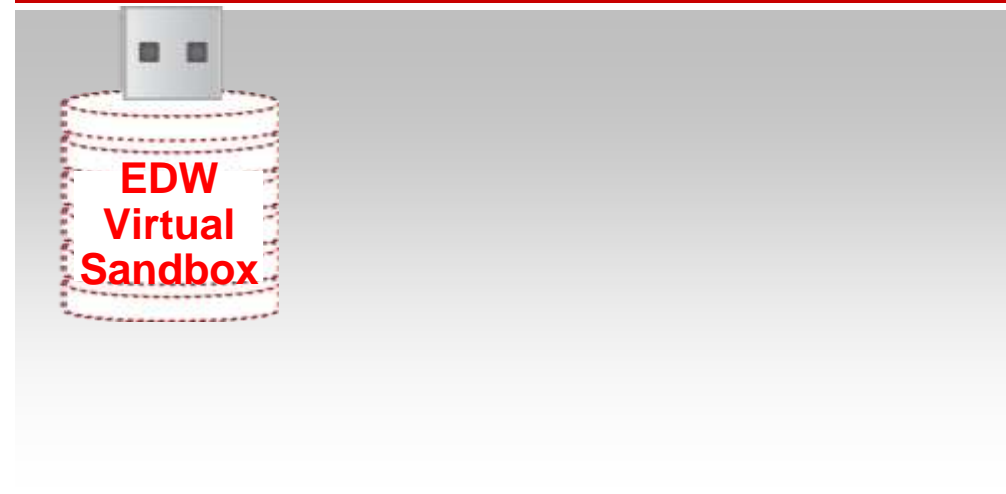
Sandboxes for Data Scientists

Pluggable databases provide virtual, resource managed sandboxes

Production Data Warehouse



Analytical Sandboxes



Accelerated Analytical Performance Across All Data

In-depth analysis

The Power of Oracle SQL

For a wide variety of 'Big Data' types



- Structured data
 - Numeric, string, date, ...
 - Row and column formats
- Unstructured data
 - LOBs
 - Text
 - XML
 - JSON
 - Spatial
 - Graph
 - Multimedia

Evolution of Analytical SQL

- Introduction of “window” functions



- Enhanced window functions (percentile, etc)
- Rollup, grouping sets, cube

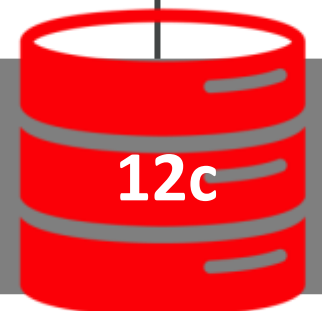
- Statistical functions
- SQL model clause
- Partition outer Join
- In-database data mining



- SQL Pivot
- Recursive WITH
- ListAgg, Nth value window



- Pattern matching
- Top N clause
- Approximate Count distinct
- JSON support

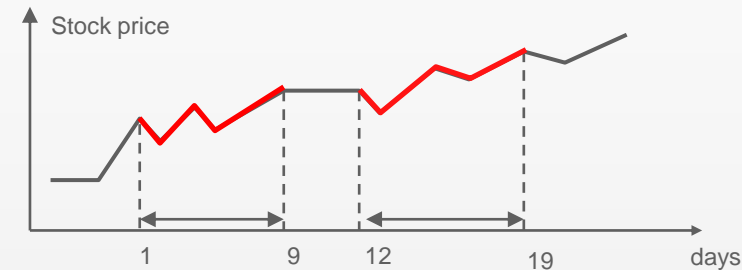


SQL Pattern Matching

Scalable discovery of business event sequences

Find event A (“privilege revoked”) followed by 3 or more occurrences of event B (“attempted login”) within 1 minute

Find 10-day periods where a stock price has “double-bottomed”



- SQL Pattern Matching provides expressive syntax and fast performance for pattern matching
- New SQL construct: MATCH_RECOGNIZE
- Define patterns using regular expression syntax

Approximate Count Distinct

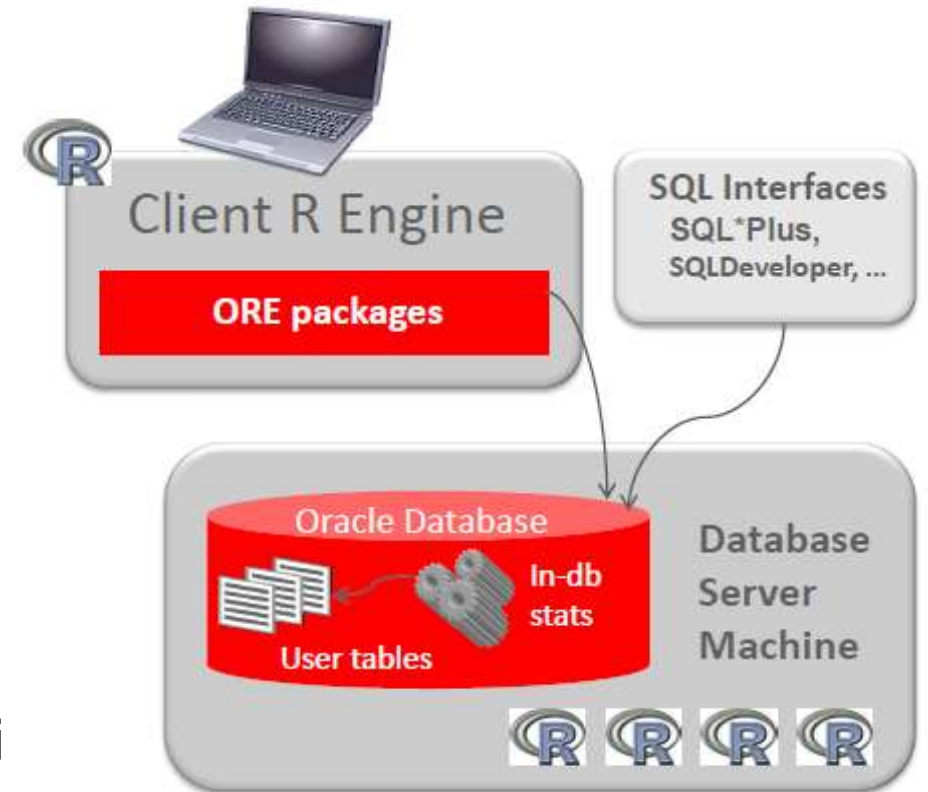
1,2,3...

- Not every query requires a completely accurate result
 - “How many distinct individuals visited our website last week?”
- New SQL function for approximate results for COUNT DISTINCT aggregates
 - APPROX_COUNT_DISTINCT()
- Approximate results can be significantly faster and use less resources than exact calculations
 - 5x to 50x ++ times faster (depending upon number of distinct values and complexity of SQL)
 - Accuracy > 97% (with 95% confidence)

Oracle Advanced Analytics

Fastest way to deliver scalable enterprise-wide predictive analytics

- Powerful
 - Scalable, parallel in-database processing
 - Mining algorithms and open source R algorithms
- Easy to Use
 - Range of GUI and IDE options for business users to
- Enterprise-wide
 - Integrated feature of the Oracle Database
 - Seamless support for enterprise analytical applicati



Agenda

- 1 The Oracle Big Data Management System
- 2 Oracle Database 12c for Data Warehousing
- 3 The Data Warehouse as Foundation for Big Data

Extending Data Management...

Big Data = Hadoop + NoSQL + Relational

Hadoop



- Change the Business
 - Disrupt competitors
 - Disintermediate supply chains
 - Leverage new paradigms
 - Exploit new analyses

NoSQL



- Scale the Business
 - Meet mobile challenges
 - Accelerate developer agility
 - Scale-out economically
 - Serve data faster



Relational

- Run the Business
 - Integrate existing systems
 - Support mission-critical tasks
 - Protect existing expenditures
 - Insure skills relevance

Oracle Big Data SQL – A New Architecture

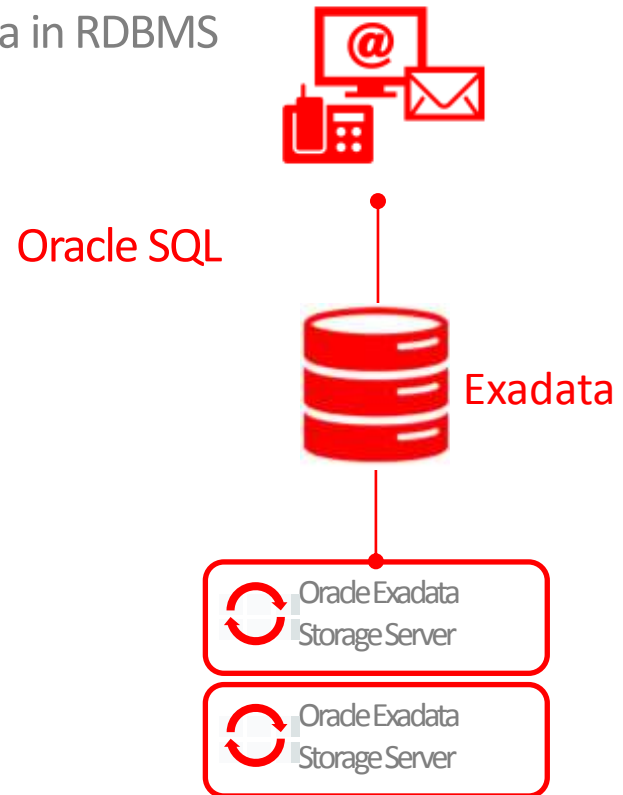
- Powerful, high-performance SQL on Hadoop
 - Full Oracle SQL capabilities on Hadoop
 - SQL query processing local to Hadoop nodes
- Simple data integration of Hadoop and Oracle Database
 - Single SQL point-of-entry to access all data
 - Scalable joins between Hadoop and RDBMS data
- Optimized hardware
 - High-speed Infiniband network between Hadoop and Exadata



What gives Exadata **extreme** performance?

Exadata: Applies SmartScan Close to the Data

Query Data in RDBMS



Oracle Big Data SQL

Exadata & Big Data SQL: Applies SmartScan Close to All Data

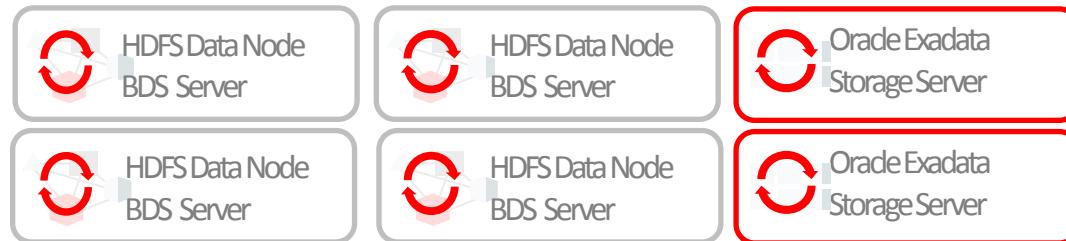
Query Data in RDBMS
and Hadoop



Oracle SQL



Exadata



Big Data Appliance

 **Fast**

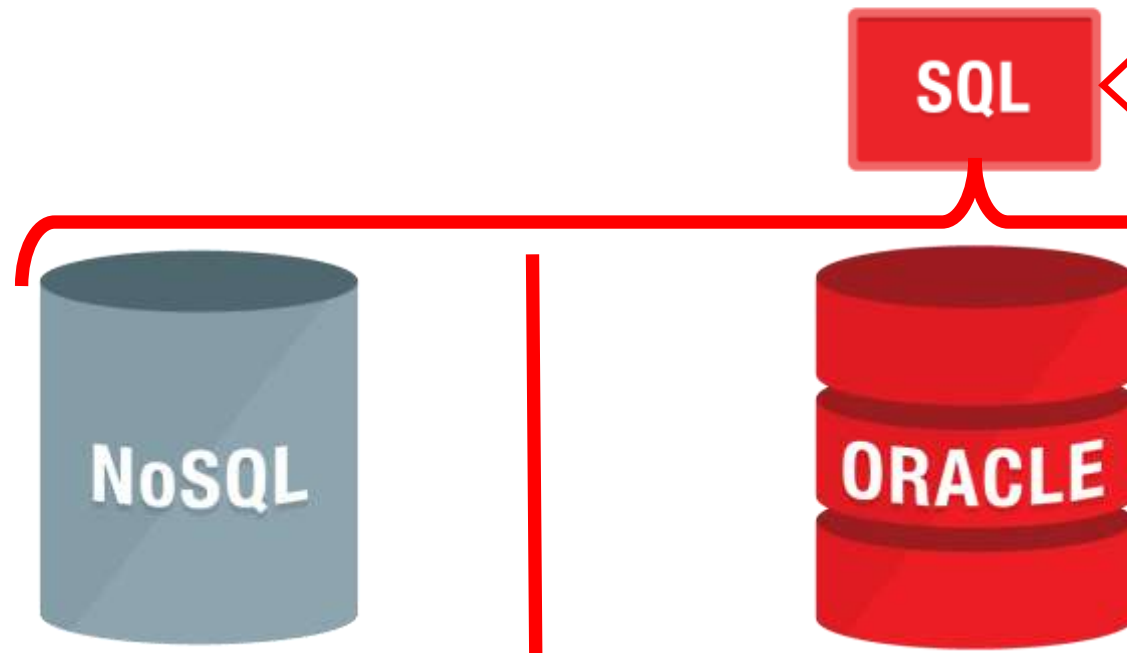
Massive Parallelism

Filtered Locally

Minimized Data Movement

Big Data Management

Hadoop + NoSQL + Relational...



The Power of Oracle SQL

Wide variety of 'Big Data' types

Structured data

Numeric, string, date, ...

Unstructured data

LOBs, Text, XML, JSON, Spatial,
Graph, Multimedia

Rich SQL Analytic Functions

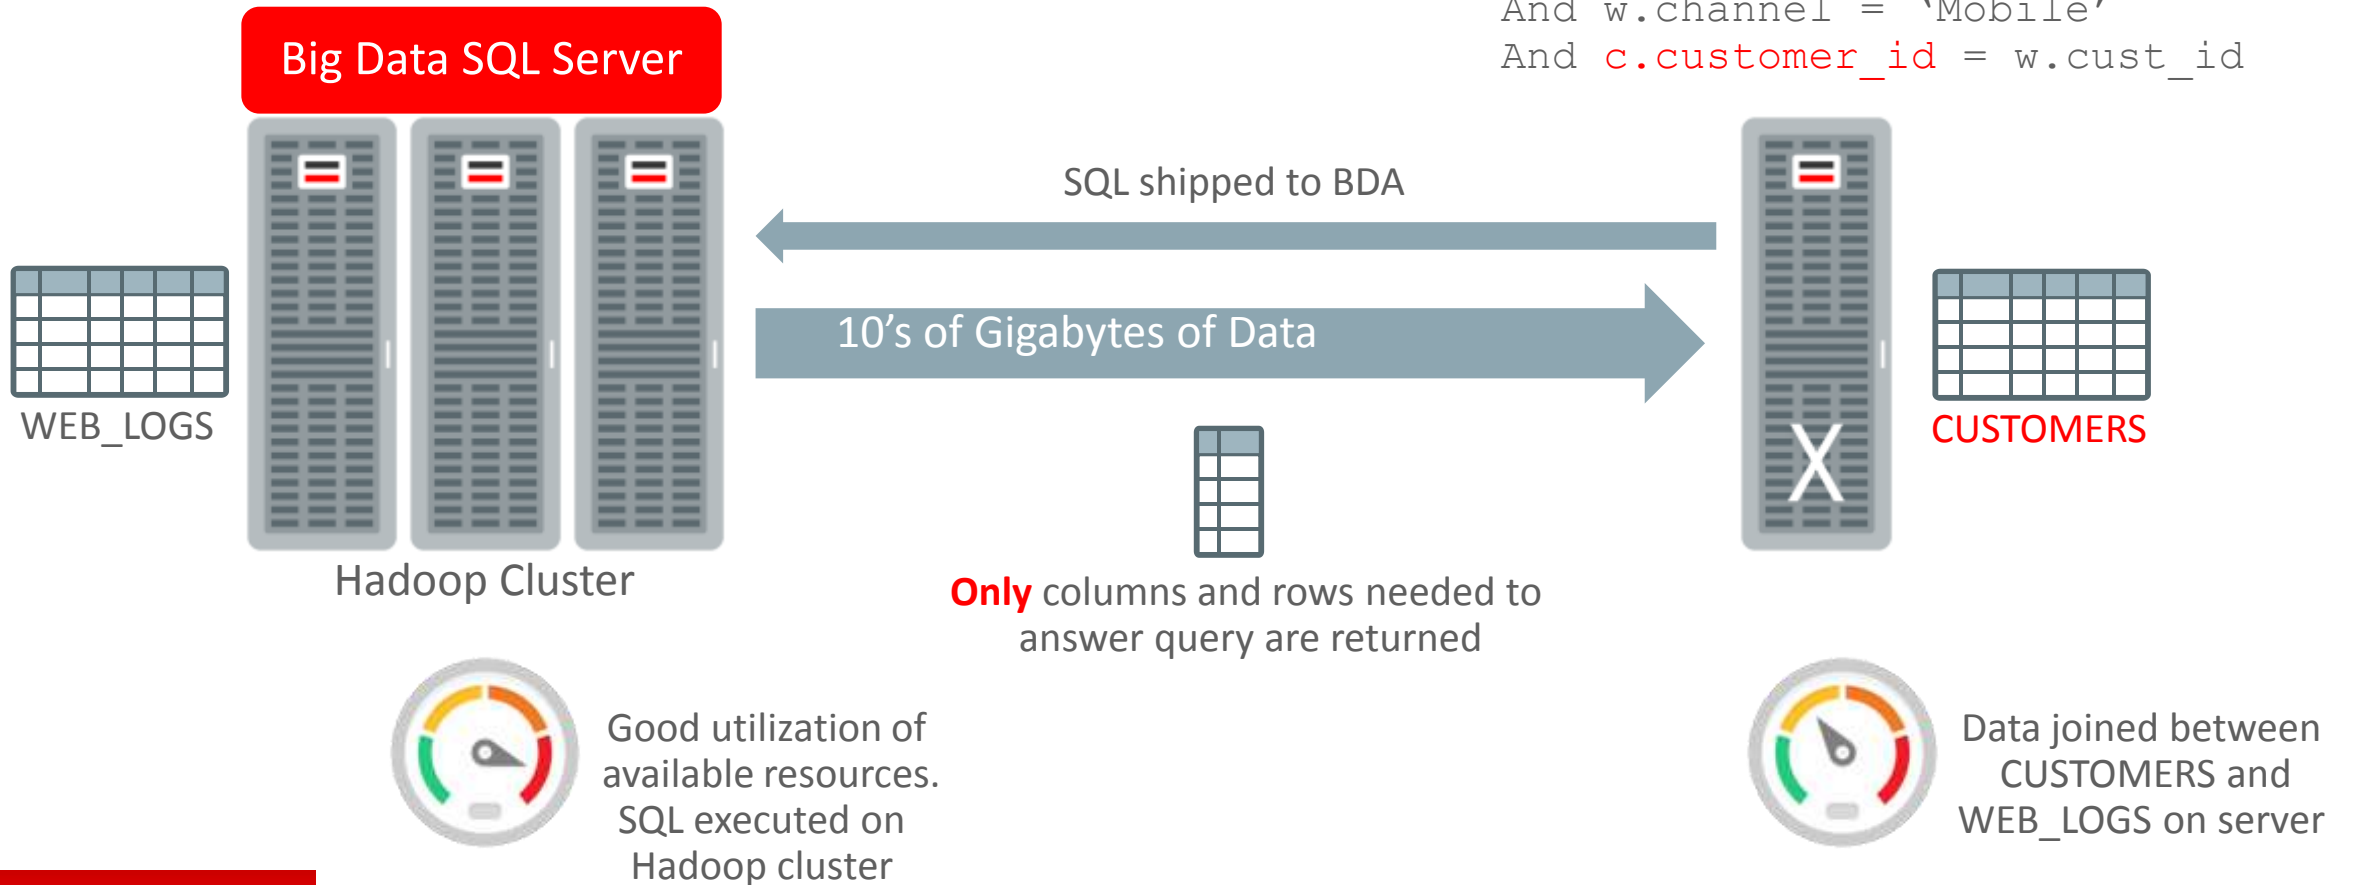
Ranking, Windowing, LAG/LEAD,
Aggregate, Statistical, Linear
Regression, Correlations, Cross
Tabs, Hypothesis Testing,
Distribution Fitting, ...

Big Data SQL Push Down

With SQL Push Down

```
Select w.sess_id,  
       w.cust_id,  
       w.page_id,  
       c.name
```

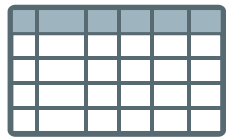
```
From web_logs w, customers c  
Where w.source_country = 'Brazil'  
And w.category = 'TV'  
And w.channel = 'Mobile'  
And c.customer_id = w.cust_id
```



Big Data SQL Push Down

With SQL Push Down

```
Select w.sess_id,  
       w.cust_id,  
       w.page_id,  
       c.name  
From web_logs w, customers c  
Where w.source_country = 'Brazil'  
And w.category = 'TV'  
      = 'Mobile'  
      id = w.cust_id
```



WEB_LOGS

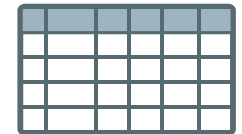
Big Data



Hadoop

SQL Push Downs supported by Big Data SQL

- Hadoop scans (InputFormat, SerDe)
- JSON parsing
- WHERE clause evaluation
- Column projection
- Bloom filters for faster join



CUSTOMERS



Good utilization of
available resources.
SQL executed on
Hadoop cluster



Data joined between
CUSTOMERS and
WEB_LOGS on server

Oracle for Data Warehousing and Big Data

Best platform for your Information Management needs

- Extreme performance
 - Optimized, intelligent hardware and software from disk to memory
- Efficient database management
 - One optimized platform for all data warehousing applications
- Accelerated analytical performance
 - All your analysis across all your data



Hardware and Software

Engineered to Work Together

ORACLE®