

ORACLE®

# Build versus Buy: Big Data for the Enterprise

Lajos Sárecz

Principal Sales Consultant

Oracle Hungary

BIG DATA

APPLIANCE



# Agenda

- Architecting Big Data
- Big Data Products
- Build versus Buy
- Connect Existing Infra with Big Data
- Summary

# Architecting Big Data



ORACLE

# Sample of Big Data Use Cases Today

## AUTOMOTIVE

Auto sensors reporting location, problems



## HIGH TECHNOLOGY INDUSTRIAL MFG.

Mfg quality  
Warranty analysis



## OIL & GAS

Drilling exploration sensor analysis



*What is the main difference in this data?*


# Volume, Velocity, Variety



*These Characteristics Challenge your Existing Architecture*

## EDUCATION & RESEARCH

Experiment sensor analysis



## HEALTH CARE

Patient sensors, monitoring, EHRs  
Quality of care



## LAW ENFORCEMENT & DEFENSE

Threat analysis - social media monitoring, photo analysis



Clinical trials



Video advertising



MEDIA

People & career



# Building your big data architecture

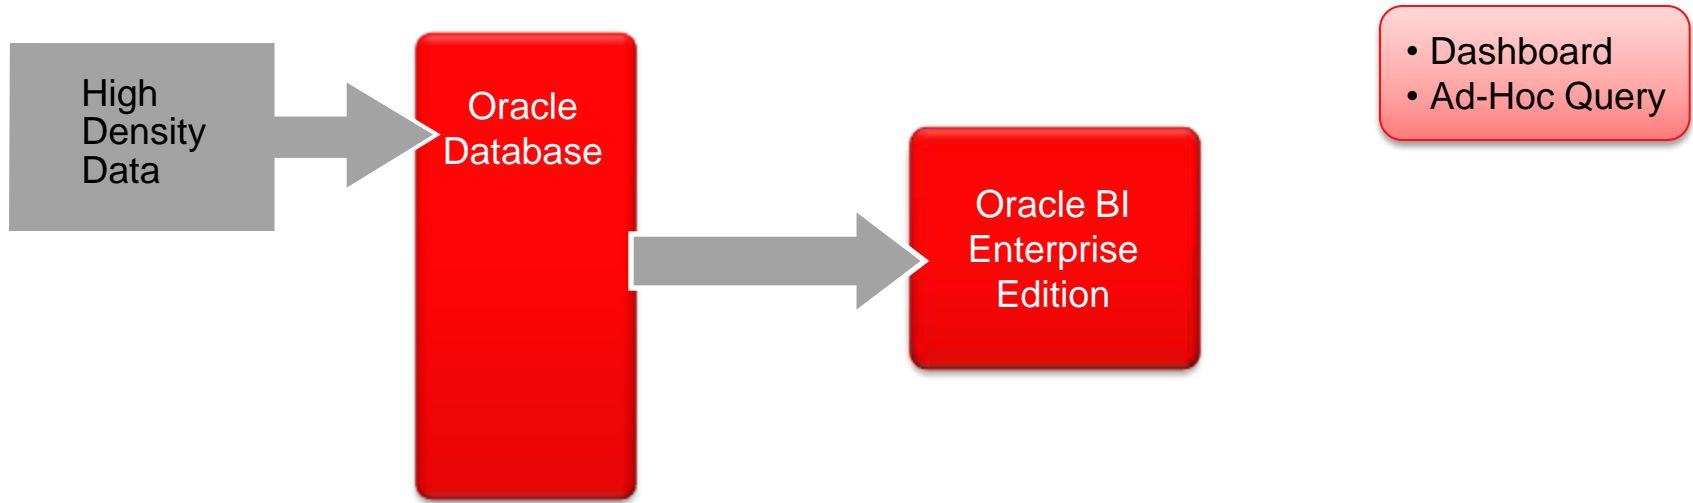
Gradually Extending your Existing Architecture for Big Data:

- Step 1: Further Analyze Current Data
- Step 2: Architect for Data Variety and Volume
- Step 3: Architect for Data Velocity
- Step 4: Discover New Patterns

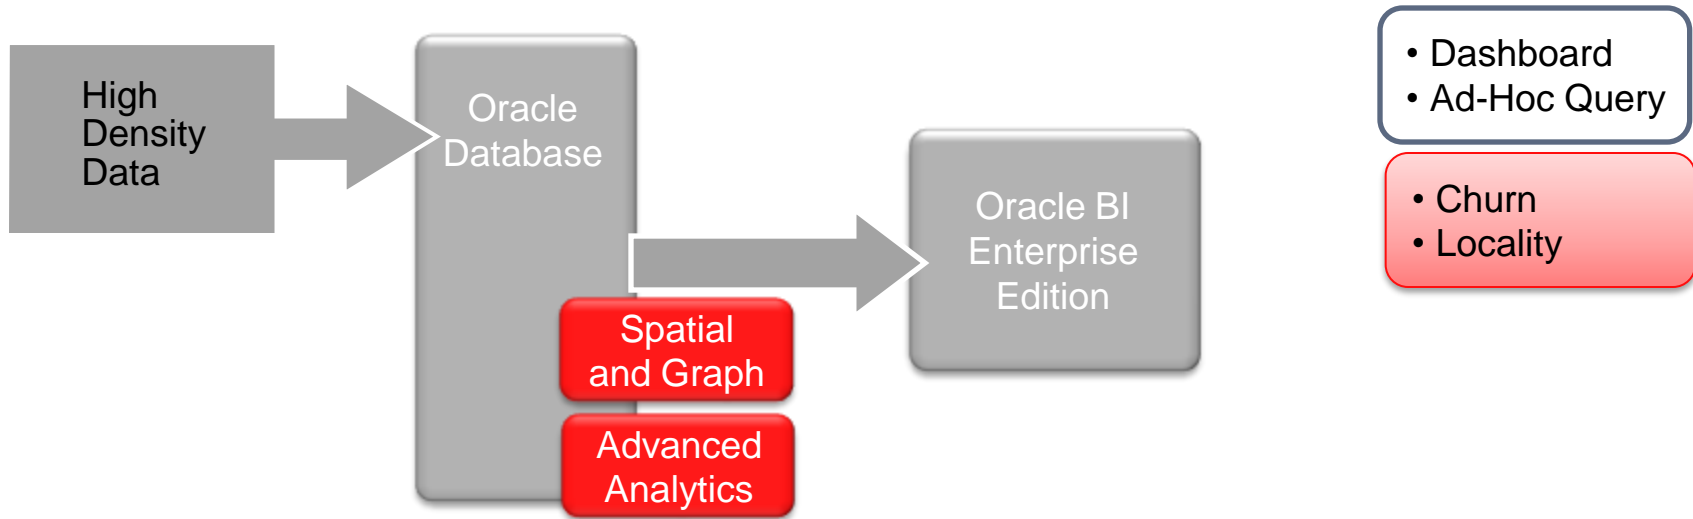


**Business Value**

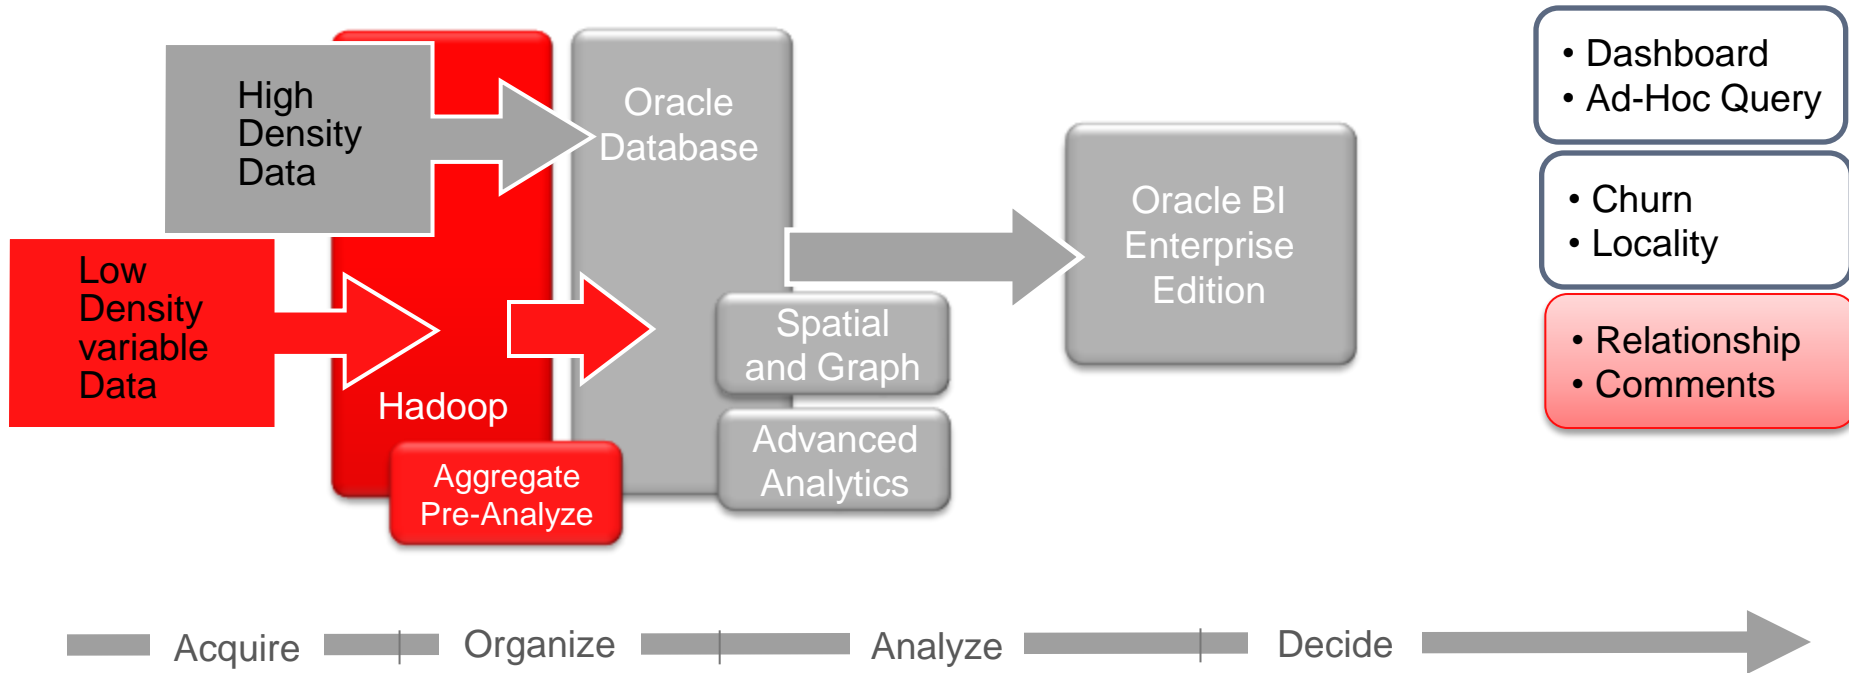
# Step 0: Data Warehouse Foundation



# Step 1: Deep Analysis of Current Data

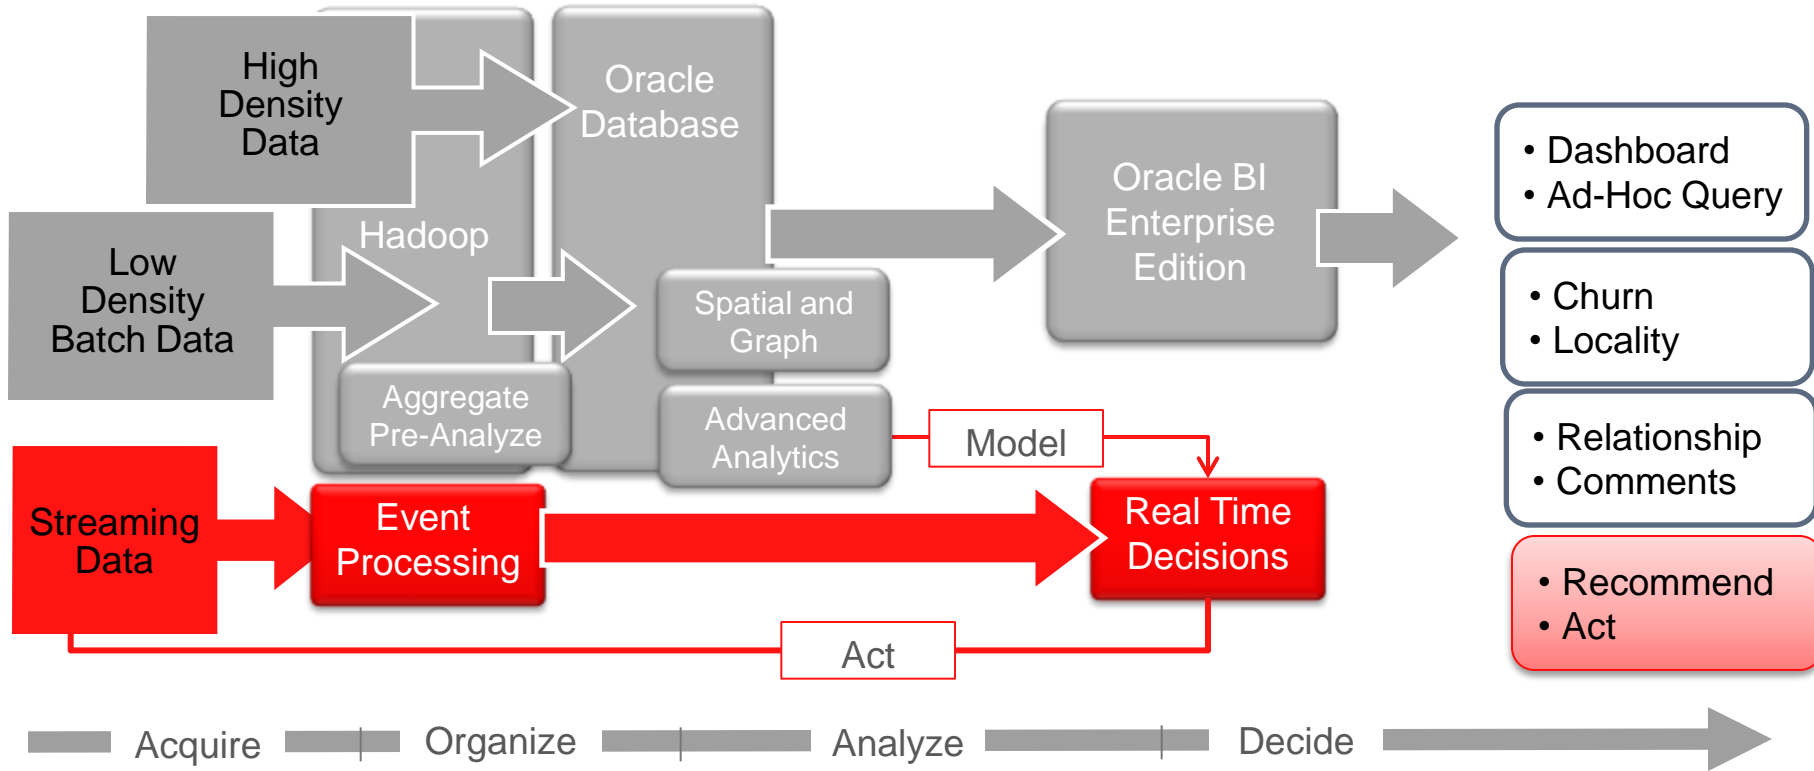


# Step 2: Architect for Volume and Variety

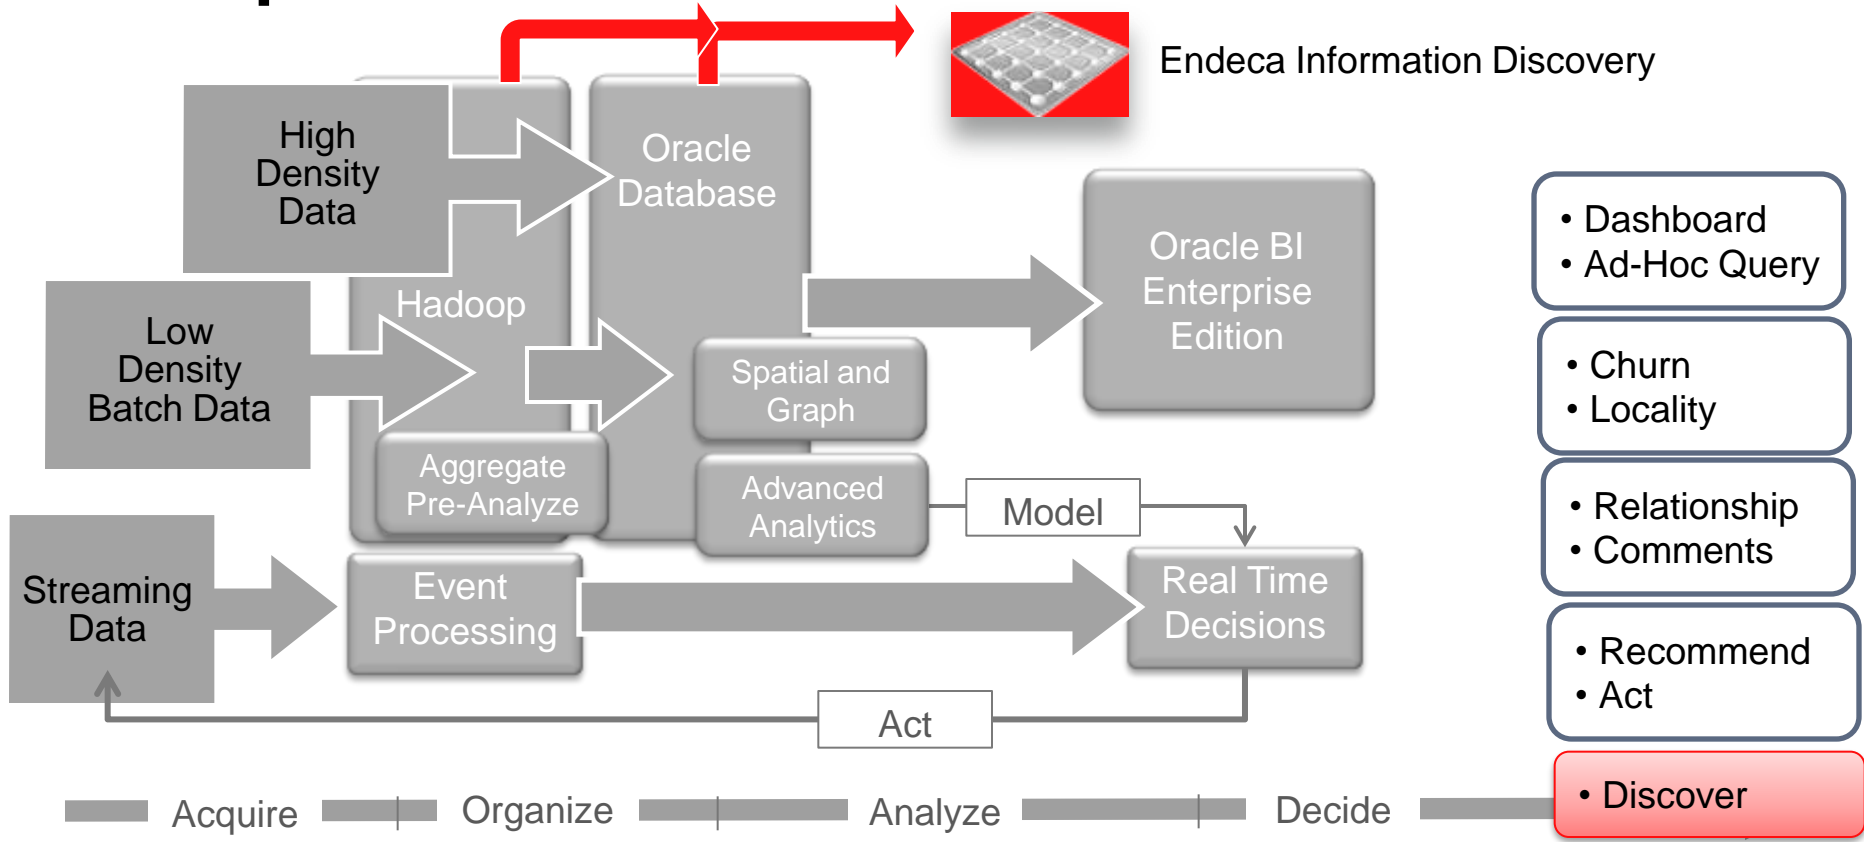




# Step 3: Architect for Velocity



# Step 4: Discover New Information



# Oracle Engineered Systems Simplify Big Data

Oracle  
Big Data  
Appliance



Oracle  
Exadata



Oracle  
Exalytics



- Dashboard
- Ad-Hoc Query

- Churn
- Locality

- Relationship
- Comments

- Recommend
- Act

- Discover

Acquire

Organize

Analyze

Decide

ORACLE

# Big Data Products

# Big Data Appliance X3-2

Sun Oracle X3-2L Servers with per server:

- 2 \* 8 Core Intel Xeon E5 Processors
- 64 GB Memory
- 36TB Disk space

Totals per Full Rack:

- 288 Processor Cores
- 1152 GB of Memory
- 648TB Available Disk space



# Big Data Appliance Software Stack

## Integrated Software:

- Oracle Linux 5.8 with UEK
- Cloudera CDH 4.2 & Cloudera Manager 4.5
- Big Data Appliance Enterprise Manager Plug-In
- Oracle R Distribution

All integrated software is supported as part of Premier Support for Systems and Premier Support for Operating Systems

## Optional Software:

- Oracle NoSQL Database 2.x
- Oracle Big Data Connectors 2.x



# Big Data Appliance Product Family



- Starter Rack is a fully cabled and configured for growth with 6 servers
- In-Rack Expansion delivers 6 server modular expansion block
- Full Rack delivers optimal blend of capacity and expansion options
- Grow by adding rack – up to 18 racks without additional switches

# Big Data Appliance X3-2 Starter Rack

Start and grow in increments of six servers

- 6 Nodes fully cabled in Starter Rack
  - 96 Intel® Xeon® E5 Processors
  - 384 GB total memory
  - 216TB total raw storage capacity
- 6 Nodes In-Rack Expansion added in-rack
  - 96 Intel® Xeon® E5 Processors
  - 384 GB total memory
  - 216TB total raw storage capacity






# Why Oracle Big Data Appliance?



- Beats DIY Clusters on:
  - Initial Cost and Time to Value
  - Performance and Scalability
- Pre-configured with leading Hadoop Distribution
  - Proven at large scale
  - Contributors across all components for better support
- Better Integration with your Oracle ecosystem with:
  - High-performance connectivity to Exadata
  - Unified analytics API (SQL, R, MapReduce etc.)
  - Single Enterprise Manager Framework

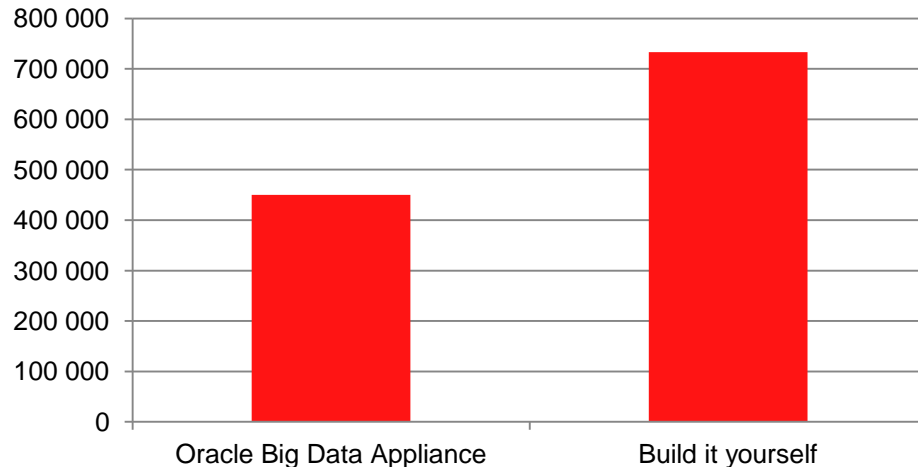
# Engineered for Quicker Time to Value at Lower Cost

## Time to Market (Weeks)



ESG believes that a "buy" versus "do-it-yourself" approach will yield roughly **one-third faster time-to-market** benefit improvement...

## Cost: Initial Infrastructure/Tasks

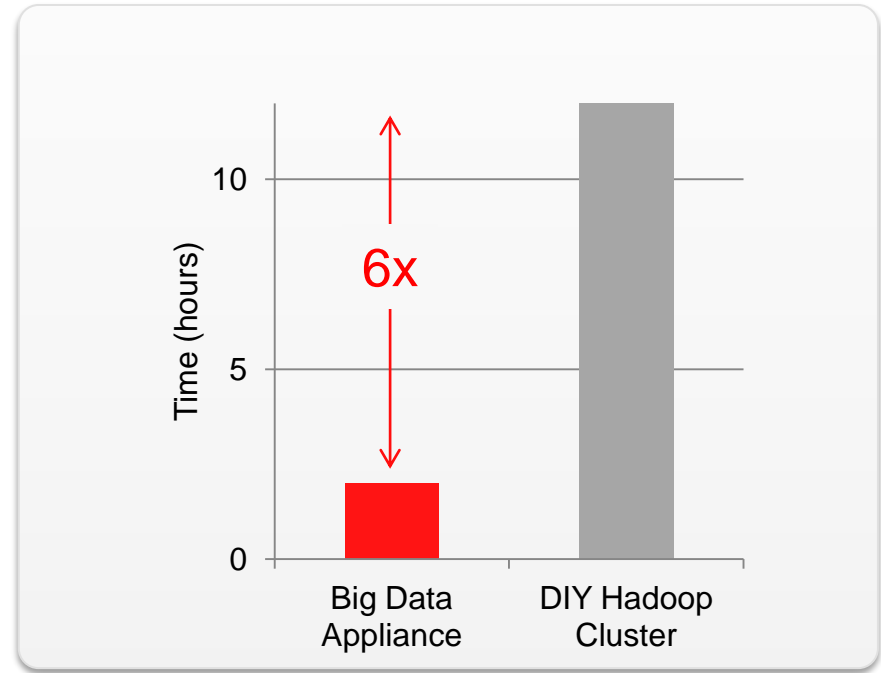


[...] nearly **40% cost savings** versus IT architecting, designing, procuring, configuring, and implementing its own big data infrastructure.

# Engineered for Performance

## Compared with a DIY Cluster

- Configured for exceptional performance on delivery
- 6x faster than custom 20-node Hadoop cluster for large batch transformation jobs
- Engineering done by Oracle and Cloudera:
  - OS and File System Tuning
  - Java Virtual Machine Tuning
  - Hadoop Configuration and Setup



# Big Data Connectors

Optimized integration of Hadoop with Oracle Database and Oracle Exadata


- Oracle Loader for Hadoop
  - Oracle SQL Connector for Hadoop Distributed File System (HDFS)
  - Oracle Data Integrator Application Adapter for Hadoop
  - Oracle R Connector for Hadoop
- 
- Does not require Big Data Appliance – can be licensed for Hadoop running on non-Oracle hardware



# Analyze Data across your Oracle Systems

## SQL Analytics on ALL data

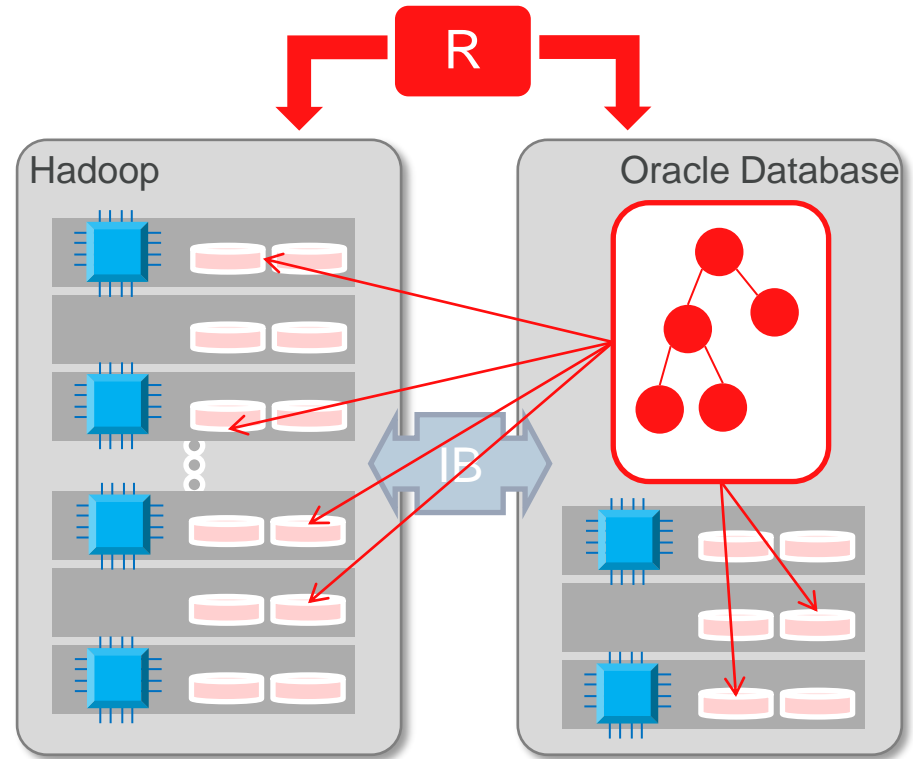
- Expand the data pool for analytics leveraging Hadoop
- Stream Hadoop resident data through Big Data Connectors for SQL processing
- Use the full power of Oracle SQL on all data
- Or use Oracle Loader for Hadoop to integrate data in Oracle Database



# Analyze Data across your Oracle Systems

## R Analytics on ALL data

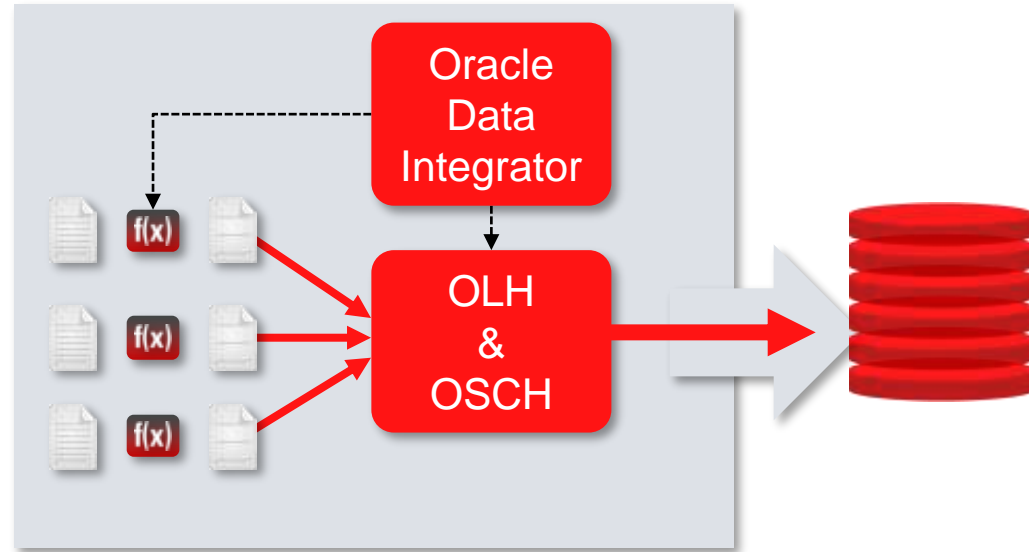
- Expand the data pool for analytics leveraging Hadoop
- Improve scalability and performance for R without changes to your programs
- Dynamically leverage Hadoop through Big Data Connectors to execute R analytics



# Oracle Data Integrator

## Simplify Map Reduce

- Automatically generates MapReduce code
- High performance loads into Data Warehouse leveraging both OLH and OSCH
- Manages the process across platforms



# Oracle Big Data Solution

Decide

Oracle Real-Time  
Decisions




Oracle BI  
Enterprise Edition



Oracle Information  
Discovery



Stream



Applications

- Complete
- Integrated
- Scalable



In-Database  
Analytics

Acquire – Organize – Analyze



# Hardware and Software

ORACLE

# Engineered to Work Together

ORACLE

ORACLE®