

ORACLE®

Big Data – Are You Ready?

Mark Townsend, VP Product Management, Database and Exadata

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 Buzz

“Why big data
is a big deal”

InfoWorld – 9/1/11

“The challenge—
and opportunity—
of big data”

McKinsey Quarterly—5/11

“Ten reasons why
Big Data will
change the travel
industry”

Tnooz -8/15/11

“Keeping Afloat
in a Sea of 'Big
Data”

ITBusinessEdge – 9/6/11

“Getting a Handle
on Big Data with
Hadoop”

Businessweek-9/7/11

“The promise of
Big Data”


Intelligent Utility-8/28/11



Big Data Use Cases

Today's Challenge	New Data	What's Possible
Healthcare Expensive office visits	Remote patient monitoring	Preventive care, reduced hospitalization
Manufacturing In-person support	Product sensors	Automated diagnosis, support
Location-Based Services Based on home zip code	Real time location data	Geo-advertising, traffic, local search
Public Sector Standardized services	Citizen surveys	Tailored services, cost reductions
Retail One size fits all marketing	Social media	Sentiment analysis segmentation


What Makes it Big Data?



VOLUME



VELOCITY



VARIETY



VALUE



Why Is Big Data Important?

US HEALTH CARE

Increase industry value per year by

\$300 B

MANUFACTURING

Decrease dev., assembly costs by

-50%

GLOBAL PERSONAL LOCATION DATA

Increase service provider revenue by

\$100 B

EUROPE PUBLIC SECTOR ADMIN

Increase industry value per year by

€250 B

US RETAIL

Increase net margin by


60+%

ORACLE

Source: * McKinsey Global Institute: Big Data – The next frontier for innovation, competition and productivity (May 2011)




Big Data in Action



**Make
Better
Decisions
Using
Big Data**




Big Data in Action



Acquire all
available data



Acquiring Big Data Challenge





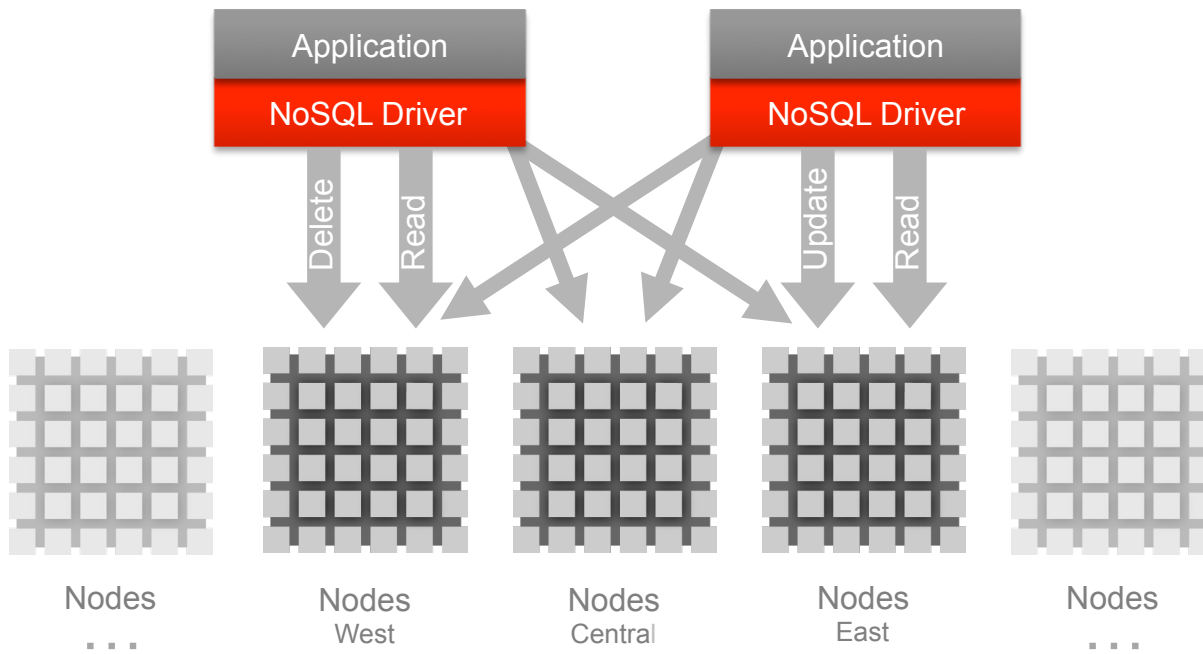
Acquiring Big Data Challenge

Need to process
high volume, low-
density information

Application will need
to change
frequently

Must scale out to
meet aggressive
roll out plan

Oracle NoSQL Database



- Key value pair database
- Dynamic data model
- Highly scalable, available
- Transparent load balancing
- Built using BerkeleyDB




Oracle NoSQL Database






Big Data in Action



Oracle NoSQL Database




Big Data in Action



Organize and **distill** big data using massive parallelism



Organizing Big Data Challenge





Organizing Big Data Challenge


Have existing
Oracle data
warehouse

Also want to
perform analysis
on big data

Can't negatively
impact data
warehouse SLAs



Analysis Sandbox

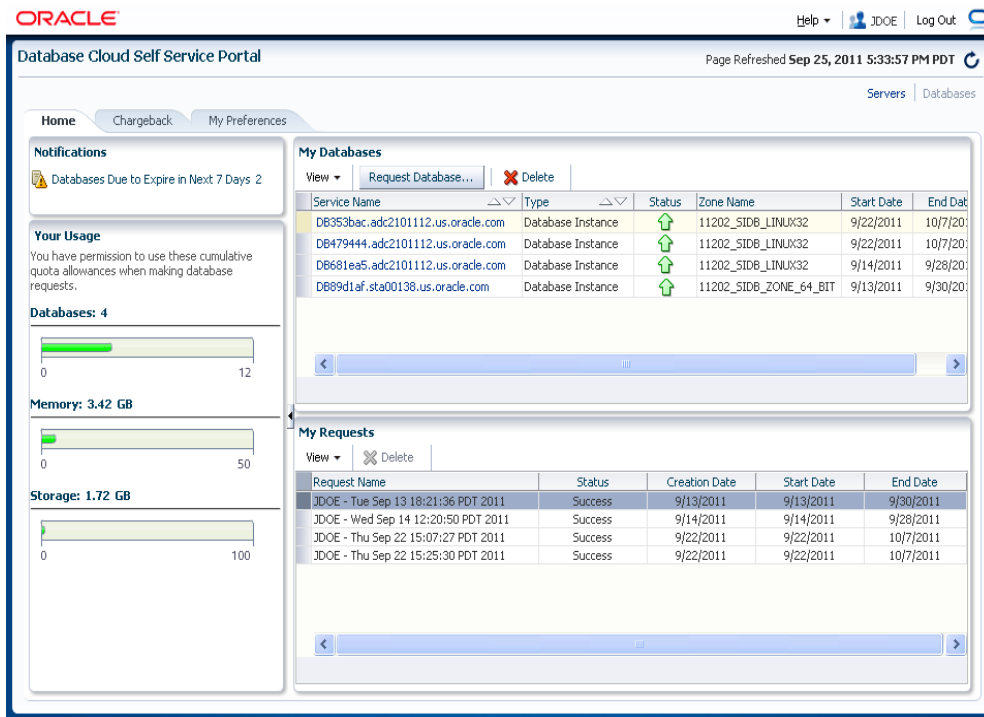


Provides analysis workspace

Controlled access to resources and data

Doesn't impact production system

Sandboxing with Oracle Enterprise Manager



The screenshot displays the Oracle Database Cloud Self Service Portal interface. The top navigation bar includes the Oracle logo, user information (JDOE), and a Log Out button. The main content area is divided into several sections:

- Notifications:** A warning icon indicates "Databases Due to Expire in Next 7 Days 2".
- Your Usage:** A section explaining cumulative quota allowances for database requests.
- Databases: 4:** A progress bar showing usage from 0 to 12.
- Memory: 3.42 GB:** A progress bar showing usage from 0 to 50.
- Storage: 1.72 GB:** A progress bar showing usage from 0 to 100.
- My Databases:** A table listing database instances with columns for Service Name, Type, Status, Zone Name, Start Date, and End Date.
- My Requests:** A table listing requests with columns for Request Name, Status, Creation Date, Start Date, and End Date.

Service Name	Type	Status	Zone Name	Start Date	End Date
D6353bac.adc2101112.us.oracle.com	Database Instance	↑	11202_SIDB_LINUX32	9/22/2011	10/7/2011
DB479444.adc2101112.us.oracle.com	Database Instance	↑	11202_SIDB_LINUX32	9/22/2011	10/7/2011
D6681ea5.adc2101112.us.oracle.com	Database Instance	↑	11202_SIDB_LINUX32	9/14/2011	9/28/2011
D689d1af.sta00138.us.oracle.com	Database Instance	↑	11202_SIDB_ZONE_64_BIT	9/13/2011	9/30/2011

Request Name	Status	Creation Date	Start Date	End Date
JDOE - Tue Sep 13 18:21:36 PDT 2011	Success	9/13/2011	9/13/2011	9/30/2011
JDOE - Wed Sep 14 12:20:50 PDT 2011	Success	9/14/2011	9/14/2011	9/28/2011
JDOE - Thu Sep 22 15:07:27 PDT 2011	Success	9/22/2011	9/22/2011	10/7/2011
JDOE - Thu Sep 22 15:25:30 PDT 2011	Success	9/22/2011	9/22/2011	10/7/2011

Simple to set up

Efficient server utilization

Secure and scalable

Accountable via charge back

Ideal for Oracle Exadata




Sandboxing with Oracle Enterprise Manager





Big Data in Action




Oracle NoSQL Database

Oracle Enterprise Manager



Organizing and Distilling Big Data Challenge





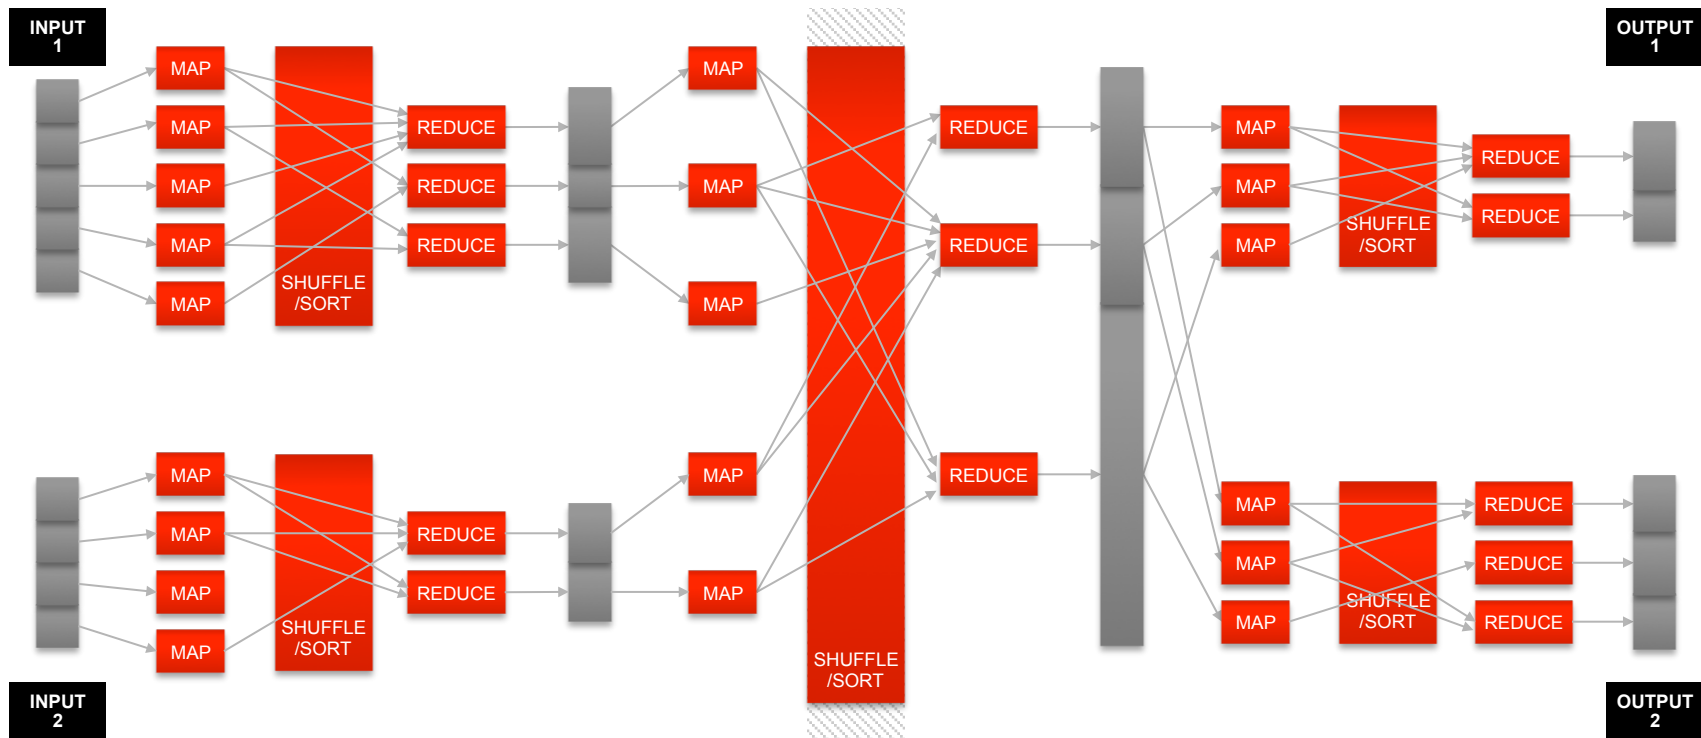
Organizing and Distilling Big Data Challenge

Must transform big data into something easily analyzed

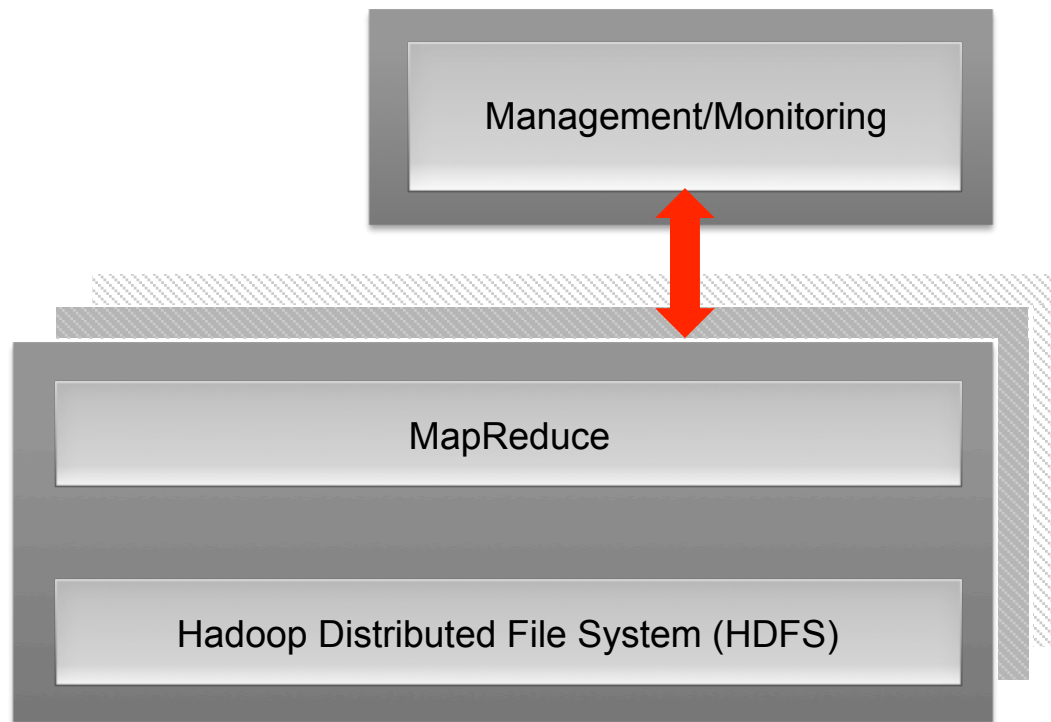
Want to avoid writing lots of Hadoop code

Need to load data quickly into Oracle Data Warehouse

A Map/Reduce Pipeline



Hadoop Architecture



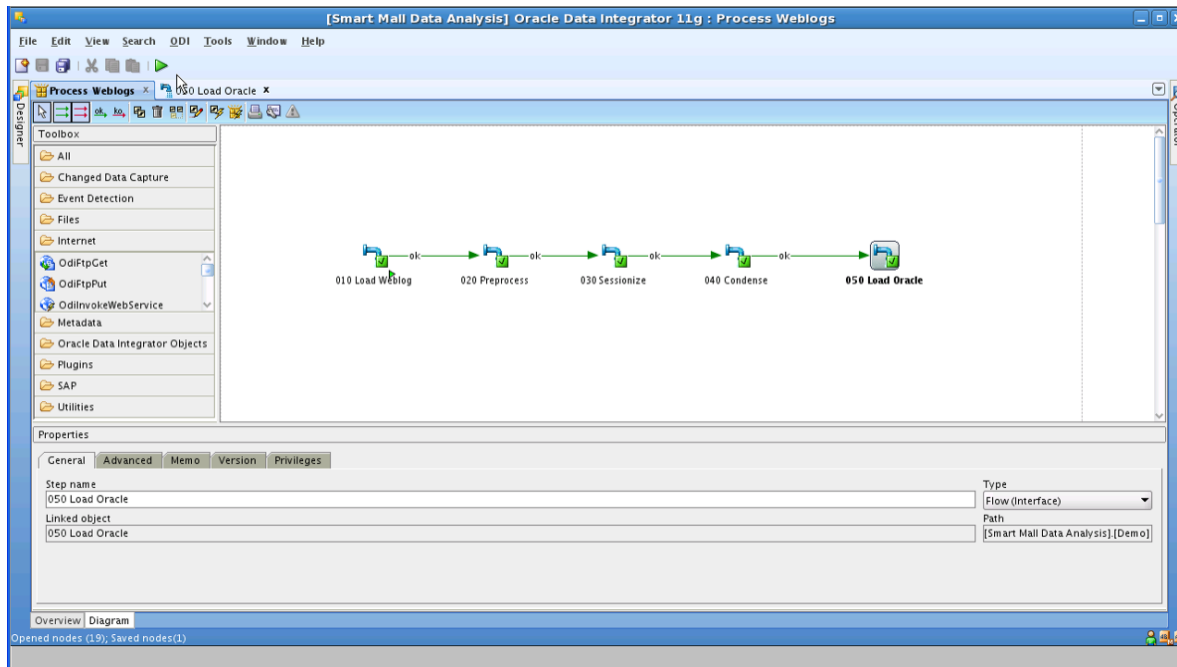
Distributed file system with redundant storage

Map/Reduce programming paradigm

Highly scalable data processing

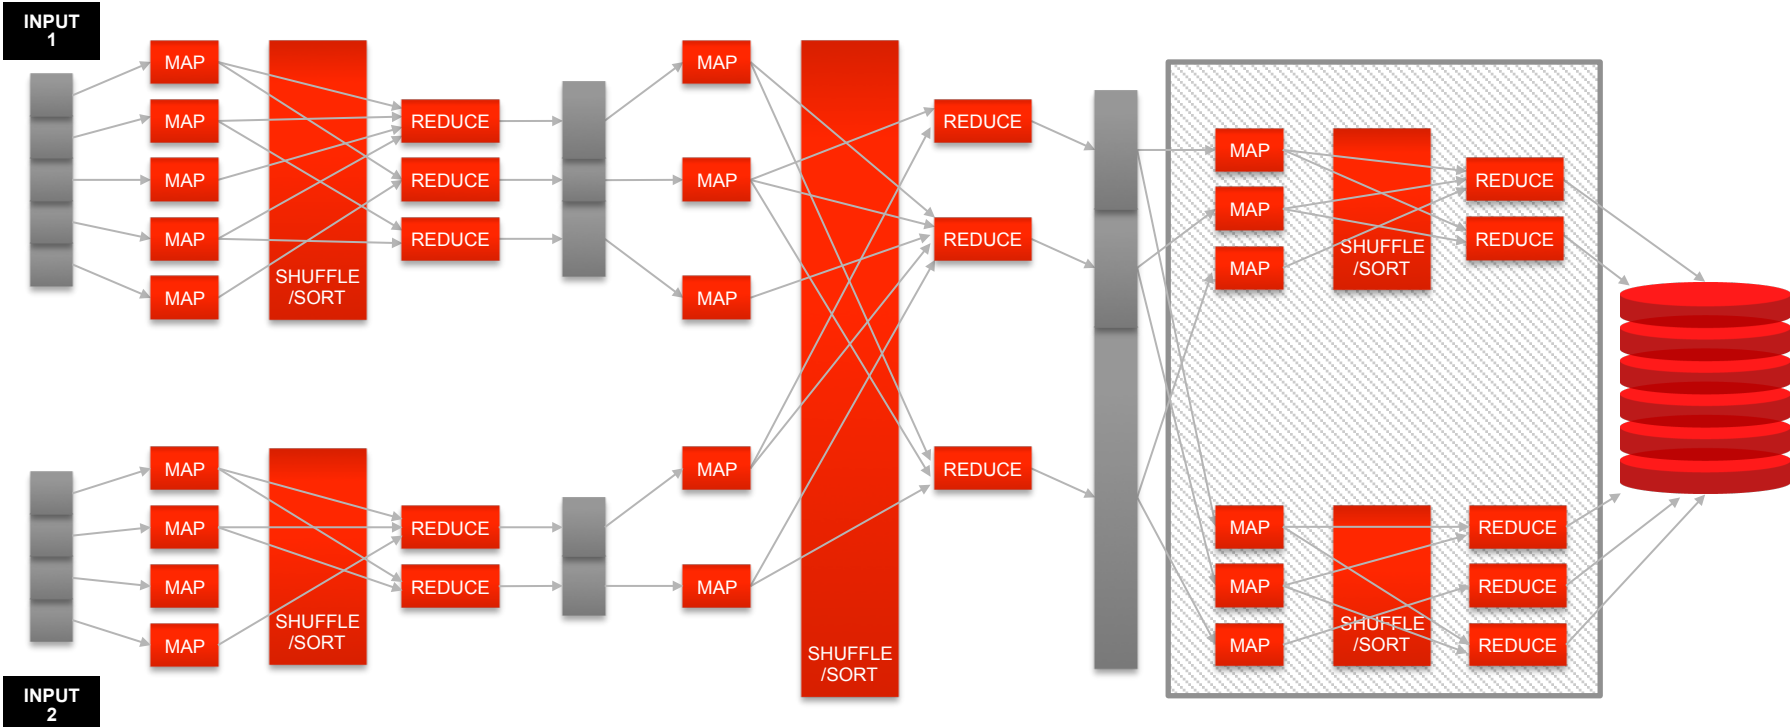
Cost-effective model for high volume, low density data

Oracle Data Integrator



Reduces Hadoop complexities through graphical tooling

Oracle Loader for Hadoop






Oracle Data Integrator, Oracle Loader for Hadoop






Big Data in Action



- Oracle NoSQL Database
- Oracle Enterprise Manager
- Oracle Data Integrator
- Oracle Loader for Hadoop




Big Data in Action



Analyze all your data, at once



Analyzing Big Data Challenge





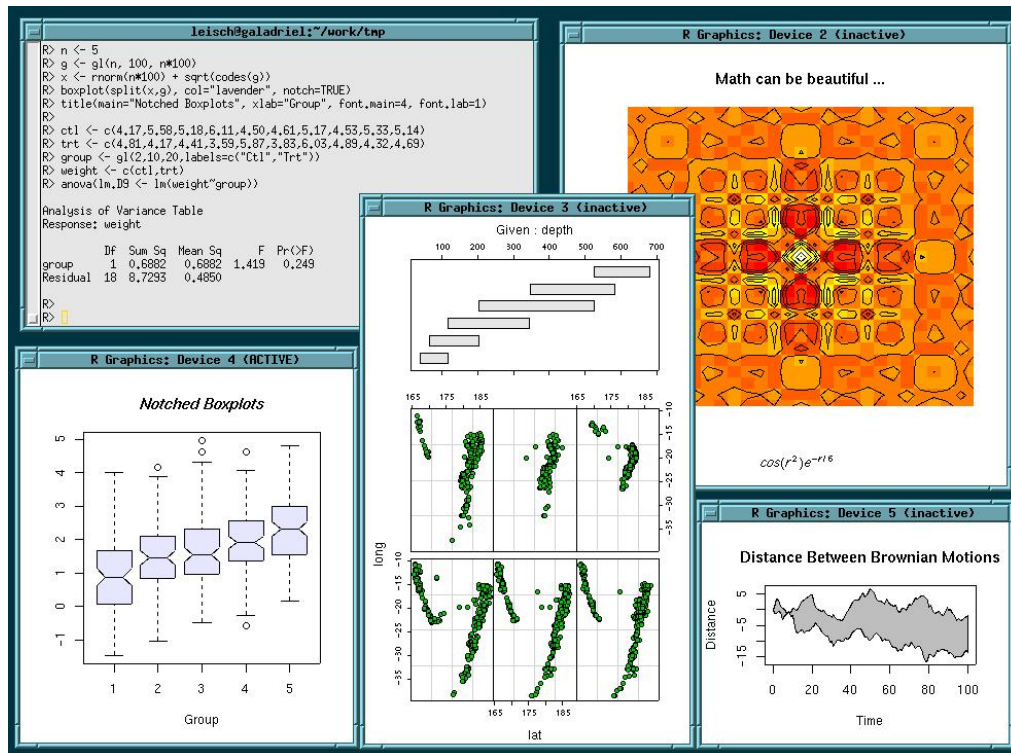
Analyzing Big Data Challenge

Require access
to all data

Want to perform
statistical analysis
using R

Doing analysis on a
laptop is slow and
not secure

R Statistical Programming Language



Open source language and environment

Used for statistical computing and graphics

Strength in easily producing publication-quality graphs

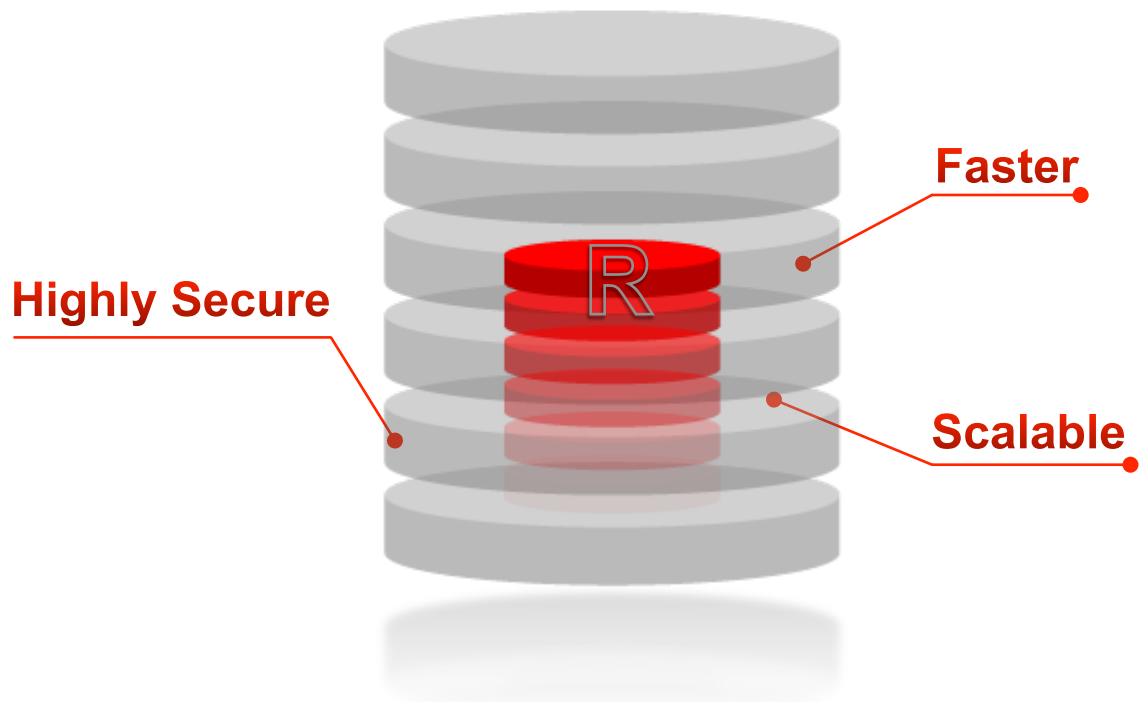
Highly extensible

Why R Wasn't Ready for the Enterprise



Small data models
only are stored and
run on user's laptop

Oracle R Enterprise Approach



Models run in-database
Processes large data sets
Uses the power of Oracle Database 11g and Exadata
Same code, much faster




Oracle R Enterprise






Big Data in Action



- Oracle NoSQL Database
- Oracle Enterprise Manager
- Oracle Data Integrator
- Oracle Loader for Hadoop
- Oracle R Enterprise**




Big Data in Action



Decide based
on real-time
big data



Making Decisions Based on Big Data Challenge





Making Decisions Based on Big Data Challenge

Big data has been transformed into actionable insight

Want to add new insights into BI dashboard

How do we quickly integrate R analytics into dashboard?

Dashboard Analytics

- Oracle Business Intelligence Enterprise Edition
 - Advanced dashboard visualization
 - Runs BI and EPM applications
- Integrating R Analytics
 - Embed R script's web interface in BI dashboard
 - Graphics will stream to BI dashboard



Oracle Exalytics Hardware

Engineered for extreme analytics

- 40 Intel processor cores
- 1 Terabyte main memory
- 40 Gb InfiniBand connection to Oracle Exadata



Oracle Exalytics Software

- Oracle TimesTen In-Memory Database
 - Adaptive in-memory caching of analytics
 - In-memory columnar compression
 - Tightly integrated with Oracle Exadata
 - Enables speed-of-thought visualization
- Oracle Business Intelligence Foundation Suite

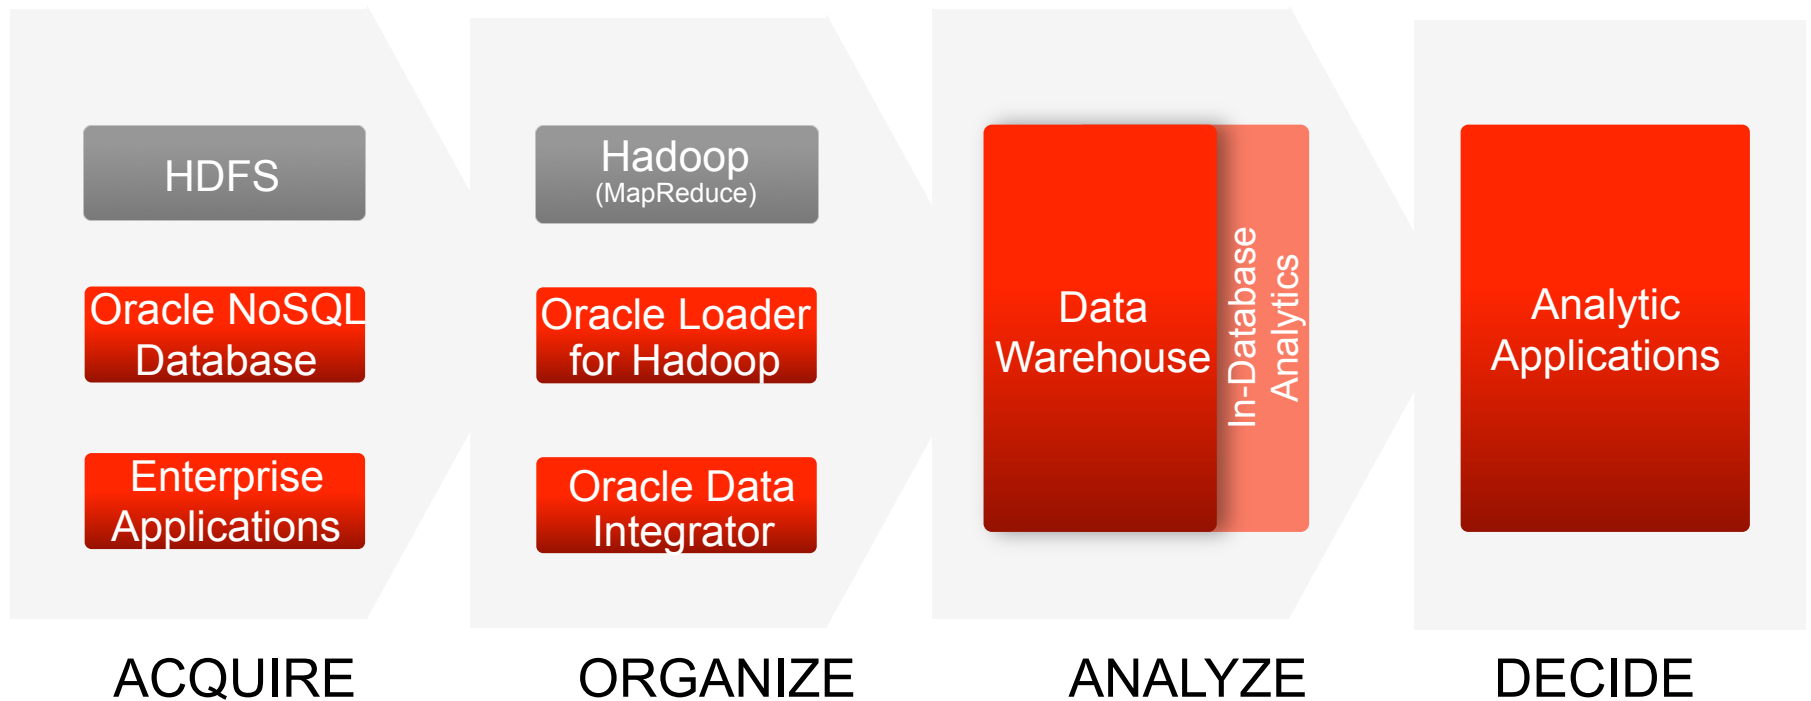




Oracle Exalytics



Oracle Integrated Solution Stack for Big Data



Oracle Integrated System Stack



ORACLE
ACQUIRE
EXALYTICS



ORACLE
ORGANIZE
EXADATA ANALYZE



ORACLE
DECIDE
BIG DATA APPLIANCE

Oracle Big Data Appliance System Hardware

- 18 Sun X4270 M2 Servers
 - 48 GB memory per node = 864 GB memory
 - 12 Intel cores per node = 216 cores
 - 24 TB storage per node = 432 TB storage
- 40 Gb p/sec InfiniBand
- 10 Gb p/sec Ethernet



Oracle Big Data Appliance Software

- Oracle Linux
- Java Hotspot VM
- Apache Hadoop Distribution
- R Distribution
- Oracle NoSQL Database
- Oracle Data Integrator for Hadoop
- Oracle Loader for Hadoop



Maximizing the Value of Enterprise Big Data

- Hardware and software for Big Data
- Integrates all enterprise data
 - Structured and unstructured
 - SQL and NoSQL
- Fastest time-to-market
- Single vendor support



Hardware and Software

ORACLE®

Engineered to Work Together

ORACLE®