


ORACLE®

Big Data – The Big Story

Jean-Pierre Dijcks

Big Data Product Mangement

1



Friday, February 22, 13

Agenda

- What is Big Data?
- Architecting Big Data
- Building Big Data Solutions
- Oracle Big Data Appliance and Big Data Connectors
- Customer Stories


What is Big Data?



ORACLE

Big Data

React to an Event




ORACLE

* Oracle Profit Magazine, Volume 17, Number 1


Friday, February 22, 13

Big Data

React to an Event



Pro-Actively Change Outcomes




ORACLE

* Oracle Profit Magazine, Volume 17, Number 1


Friday, February 22, 13

Big Data

React to an Event



Pro-Actively Change Outcomes



*“Technology presents the opportunity to transform business”**
Mark Hurd, President, Oracle

ORACLE

* Oracle Profit Magazine, Volume 17, Number 1

Friday, February 22, 13

Sample of Big Data Use Cases Today

AUTOMOTIVE

Auto sensors reporting location, problems



COMMUNICATIONS

Location-based advertising



Retail / CPG

Sentiment analysis
Hot products
Optimized Marketing




FINANCIAL SERVICES

Risk & portfolio analysis
New products



EDUCATION & RESEARCH

Experiment sensor analysis



HIGH TECHNOLOGY / INDUSTRIAL MFG.

Mfg quality
Warranty analysis



LIFE SCIENCES

Clinical trials
Genomics



MEDIA / ENTERTAINMENT

Viewers / advertising effectiveness
Cross Sell



ON-LINE SERVICES / SOCIAL MEDIA

People & career matching
Web-site optimization



HEALTH CARE

Patient sensors, monitoring, EHRs
Quality of care



OIL & GAS

Drilling exploration
sensor analysis



Games

Adjust to player behavior
In-Game Ads



TRAVEL & TRANSPORTATION

Sensor analysis for optimal traffic flows
Customer sentiment



UTILITIES

Smart Meter analysis for network capacity,



LAW ENFORCEMENT & DEFENSE

Threat analysis - social media monitoring, photo analysis



ORACLE

Friday, February 22, 13

Every industry has examples where the improved precision that can be provided with Big Data could be valued.

Machines, equipment, people

Sample of Big Data Use Cases Today

What is the main difference in this data?

Volume, Velocity, Variety

These Characteristics Challenge your Existing Architecture

AUTOMOTIVE

Auto sensors reporting location, problems



HIGH TECHNOLOGY INDUSTRIAL MFG.

Mfg quality
Warranty analysis




OIL & GAS

Drilling exploration sensor analysis



EDUCATION & RESEARCH

Experiment sensor analysis



HEALTH CARE

Patient sensors, monitoring, EHRs
Quality of care



LAW ENFORCEMENT & DEFENSE

Threat analysis - social media monitoring, photo analysis



Friday, February 22, 13

Every industry has examples where the improved precision that can be provided with Big Data could be valued.

Machines, equipment, people

Big Data Extends the Breadth and Speed of Data

Information
Architectures

Today:

Decisions based on
database data

Transactions



ORACLE

Big Data Extends the Breadth and Speed of Data

Big Data:
Decisions based on
all your data

Video and Images



Documents



Social Data



Machine-Generated Data



Information
Architectures

Today:
Decisions based on
database data

Transactions





ORACLE

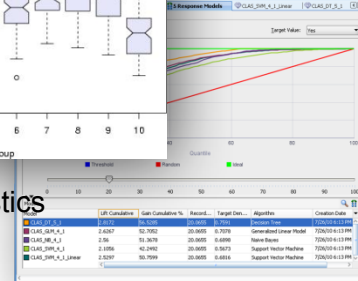
Big Data Extends the **Depth of Analytics**



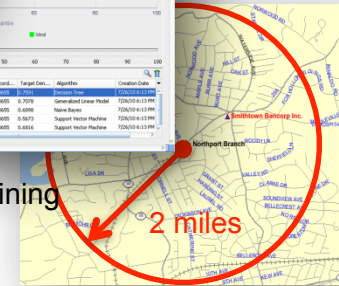
Query and Reporting



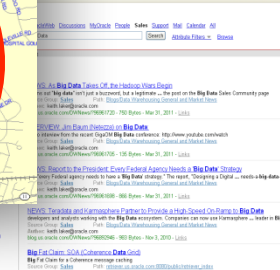
Statistics



Data Mining



Spatial Analytics



Text Analytics
ORACLE

Architecting Big Data



ORACLE

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

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

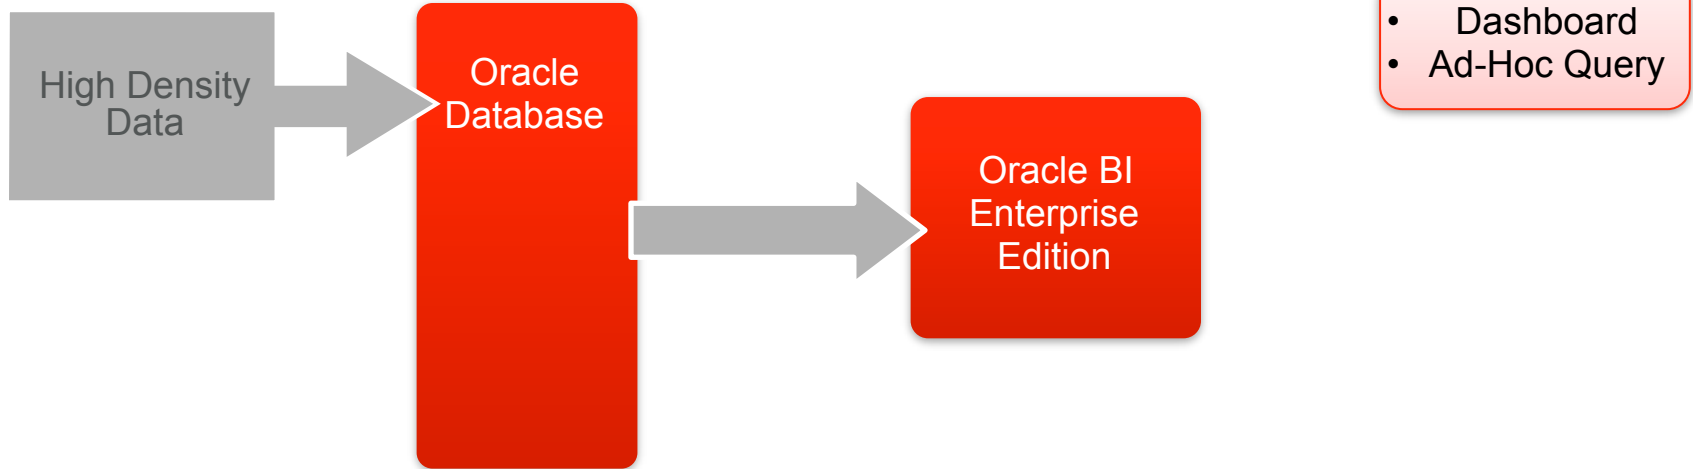
Step 3: Architect for Data Velocity

Step 4: Discover New Patterns



Business Value

Step 0: Data Warehouse Foundation




ORACLE

Friday, February 22, 13

Who are my most valuable customers?

What is last Months Revenue?

Step 1: Deep Analysis of Current Data



- Dashboard
- Ad-Hoc Query

- Churn
- Locality



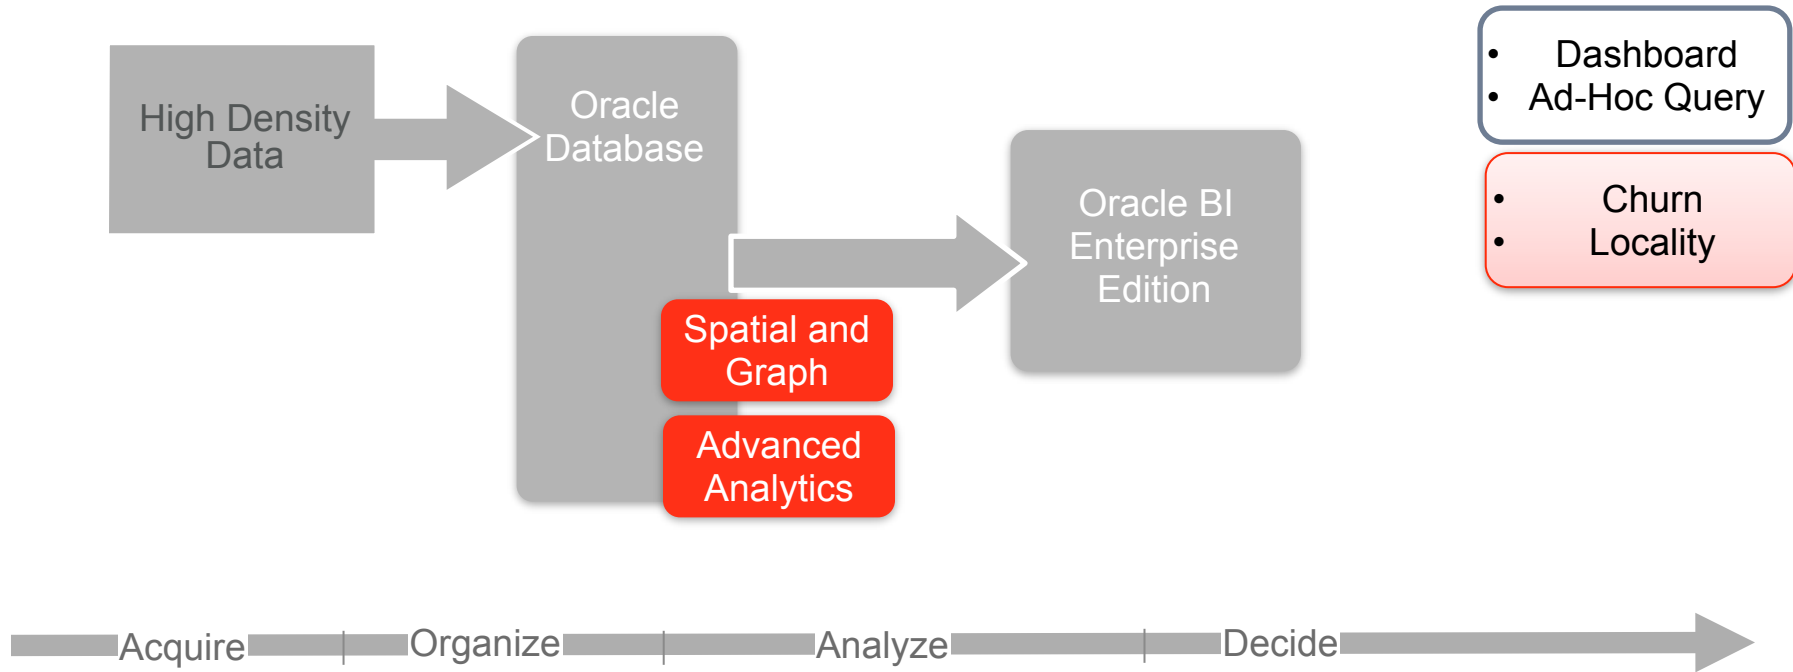
ORACLE

Friday, February 22, 13

Where are my most valuable customers located?

Who is most likely to churn in the next month?

Step 1: Deep Analysis of Current Data



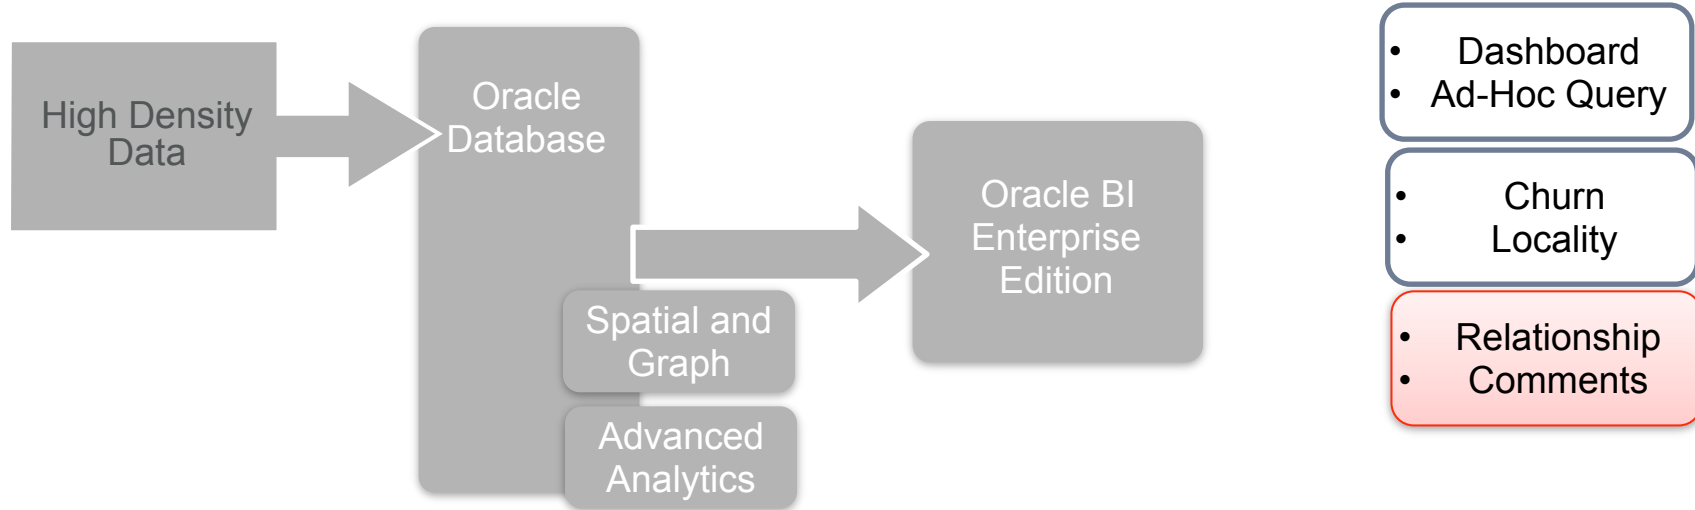
ORACLE

Friday, February 22, 13

Where are my most valuable customers located?

Who is most likely to churn in the next month?

Step 2: Architect for Volume and Variety



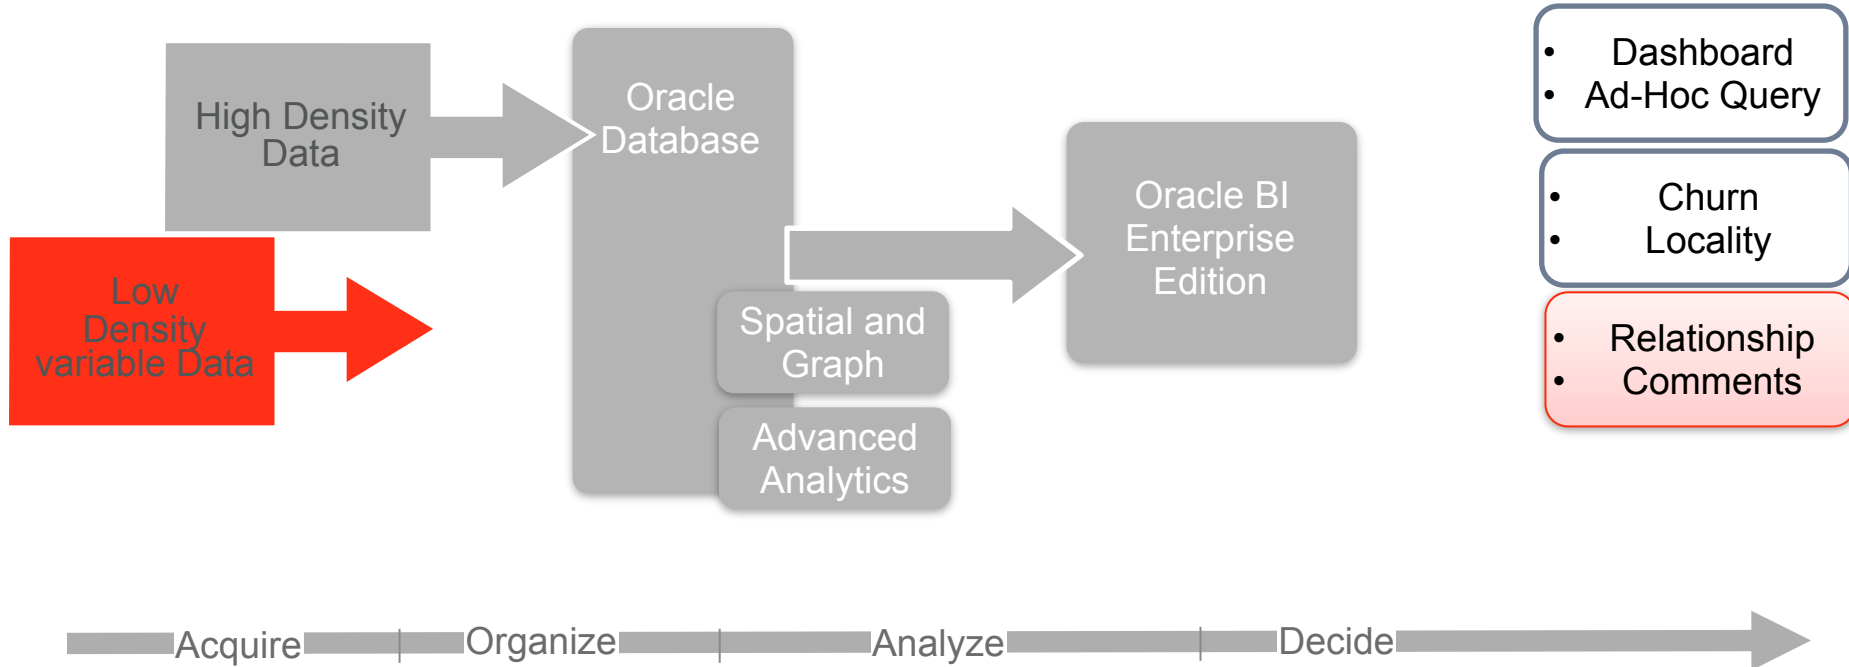
ORACLE

Friday, February 22, 13

What opinions do people have?

Why do people buy these products?

Step 2: Architect for Volume and Variety

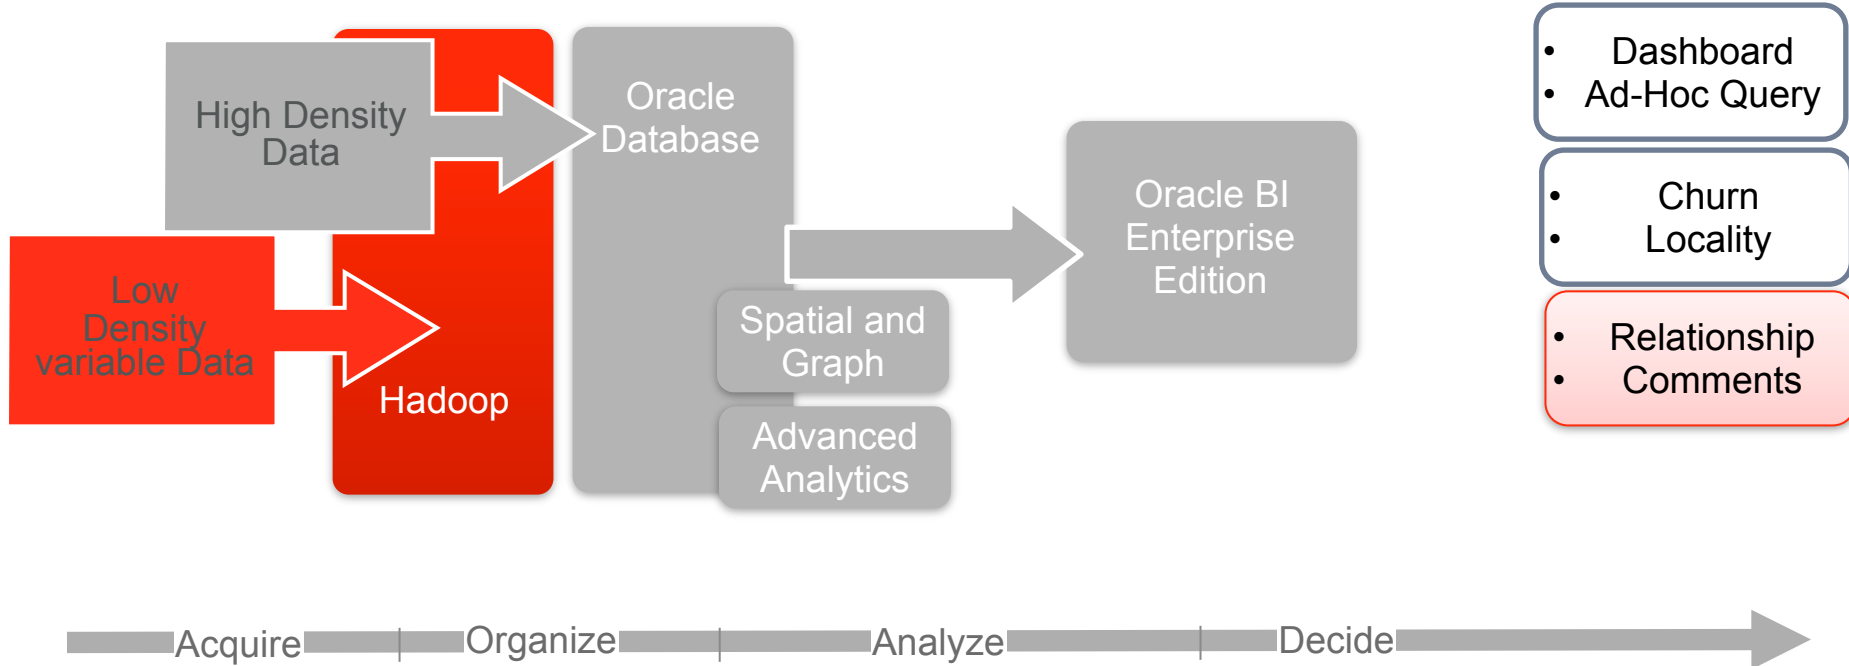


Friday, February 22, 13

What opinions do people have?

Why do people buy these products?

Step 2: Architect for Volume and Variety



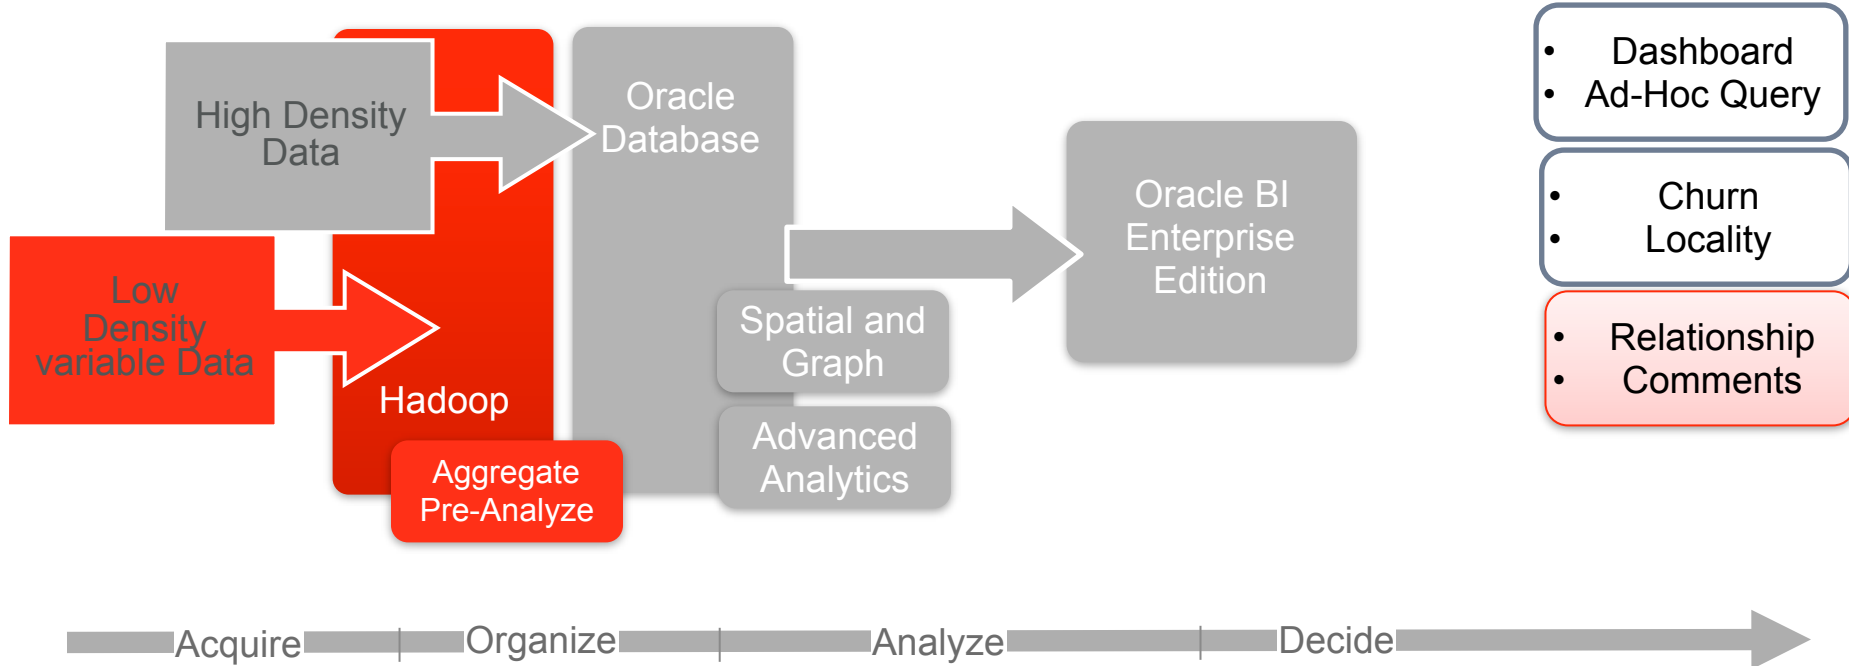
ORACLE

Friday, February 22, 13

What opinions do people have?

Why do people buy these products?

Step 2: Architect for Volume and Variety

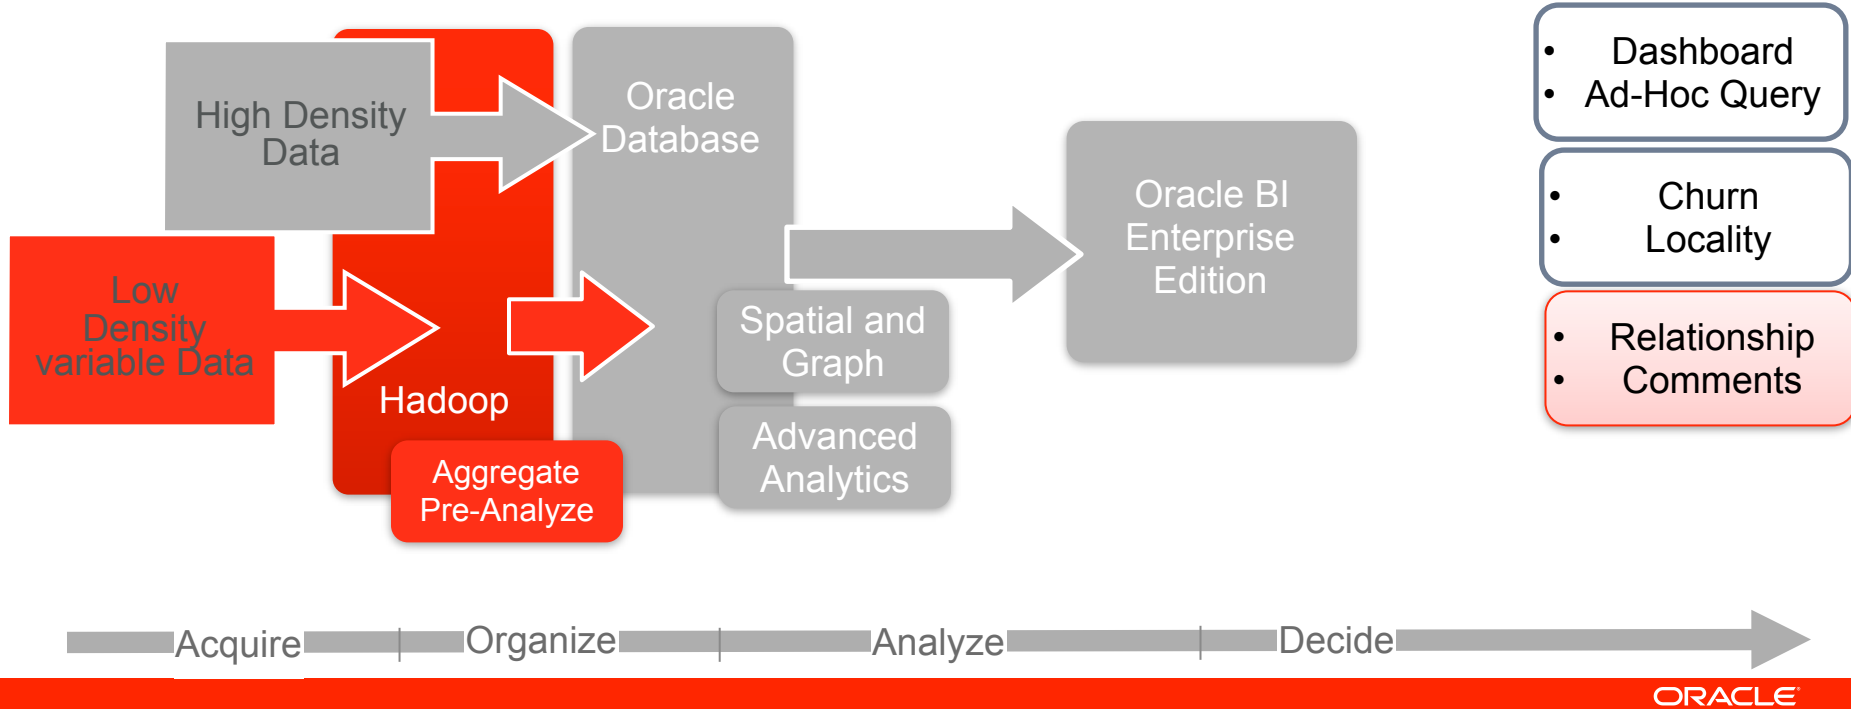


Friday, February 22, 13

What opinions do people have?

Why do people buy these products?

Step 2: Architect for Volume and Variety

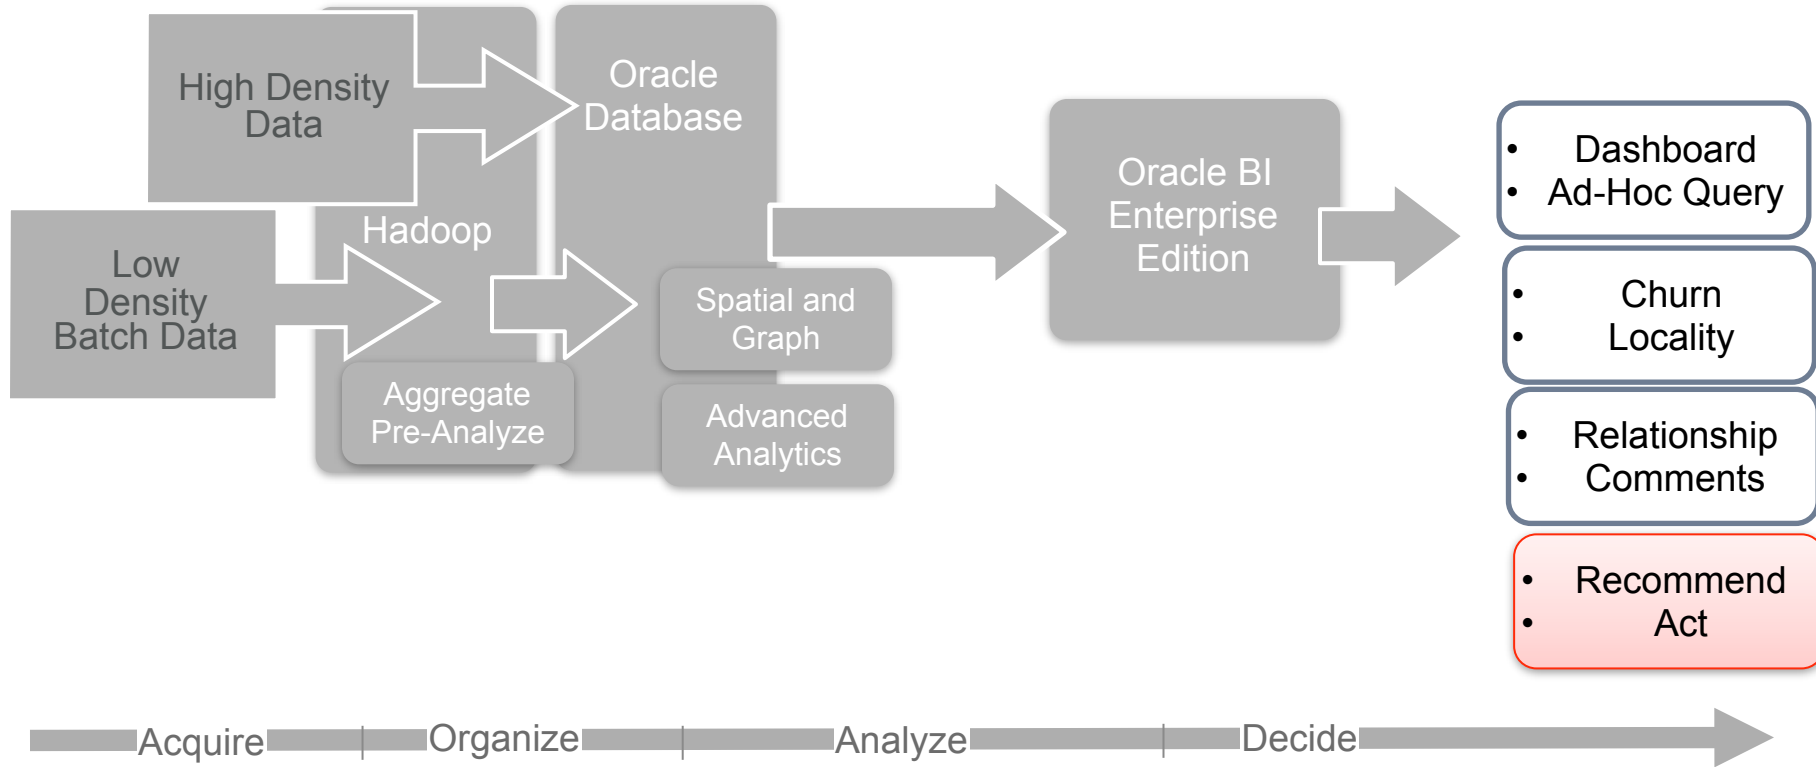


Friday, February 22, 13

What opinions do people have?

Why do people buy these products?

Step 3: Architect for Velocity



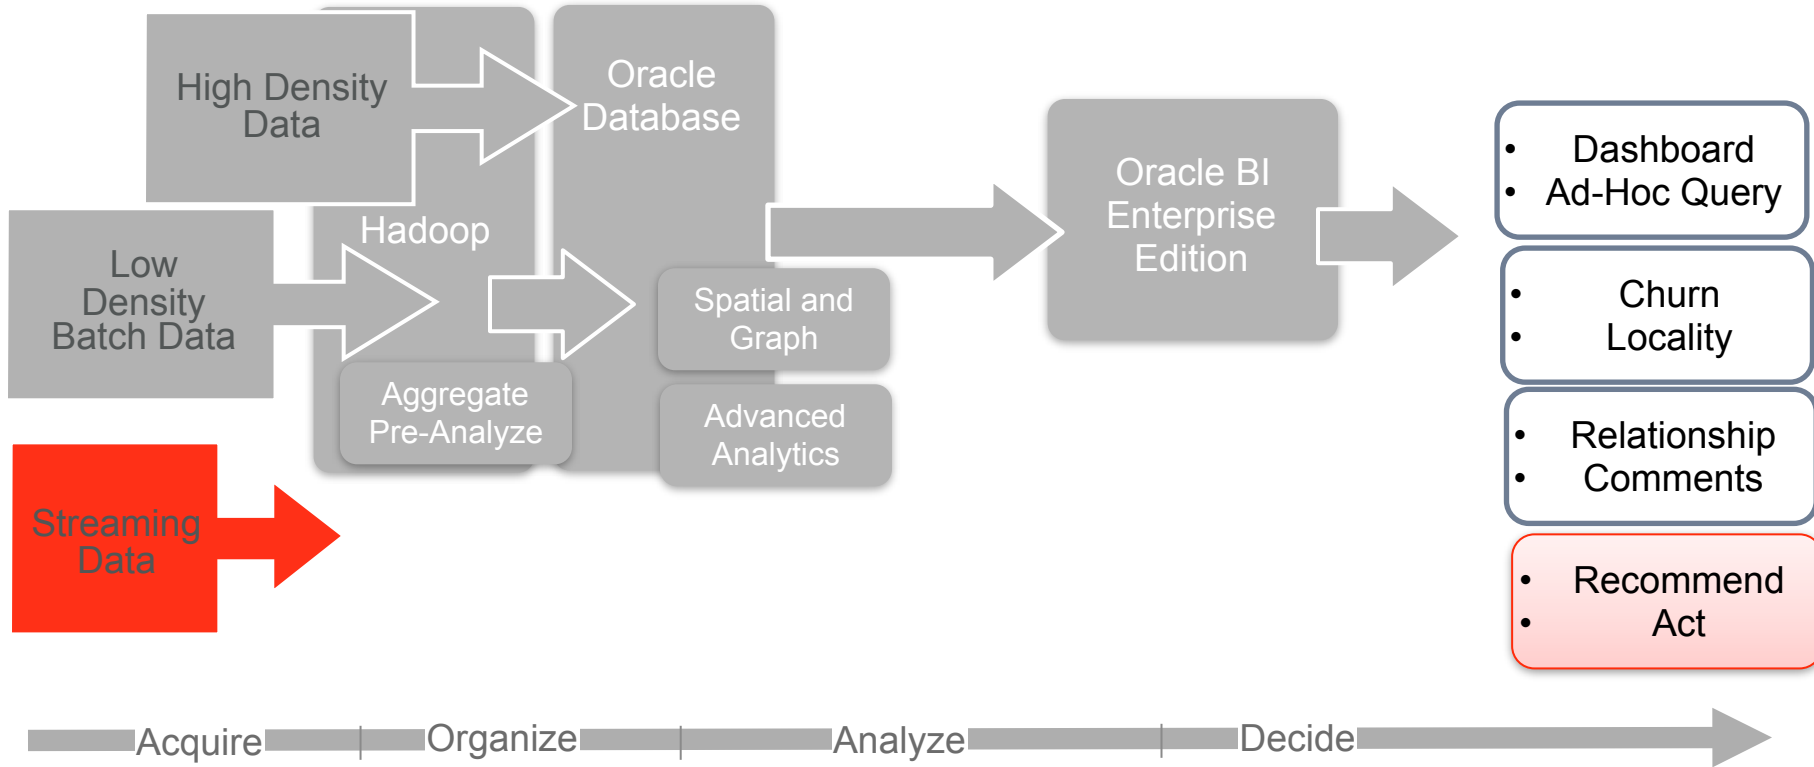
ORACLE

Friday, February 22, 13

How do I influence buying decisions?

How do I prevent serious issues from happening?

Step 3: Architect for Velocity



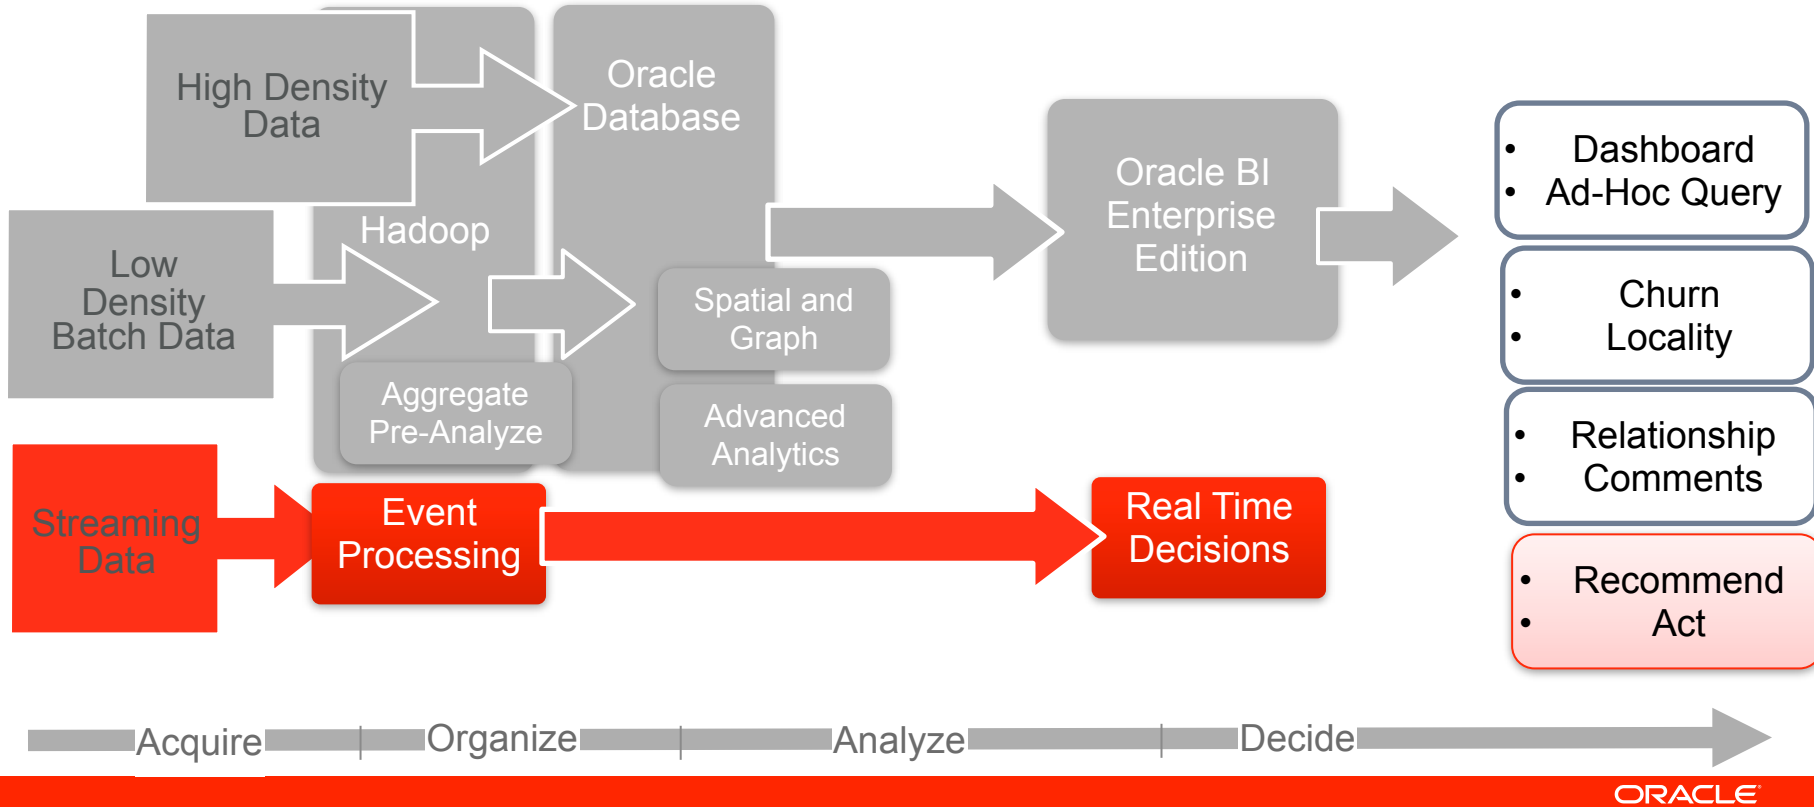
ORACLE

Friday, February 22, 13

How do I influence buying decisions?

How do I prevent serious issues from happening?

Step 3: Architect for Velocity

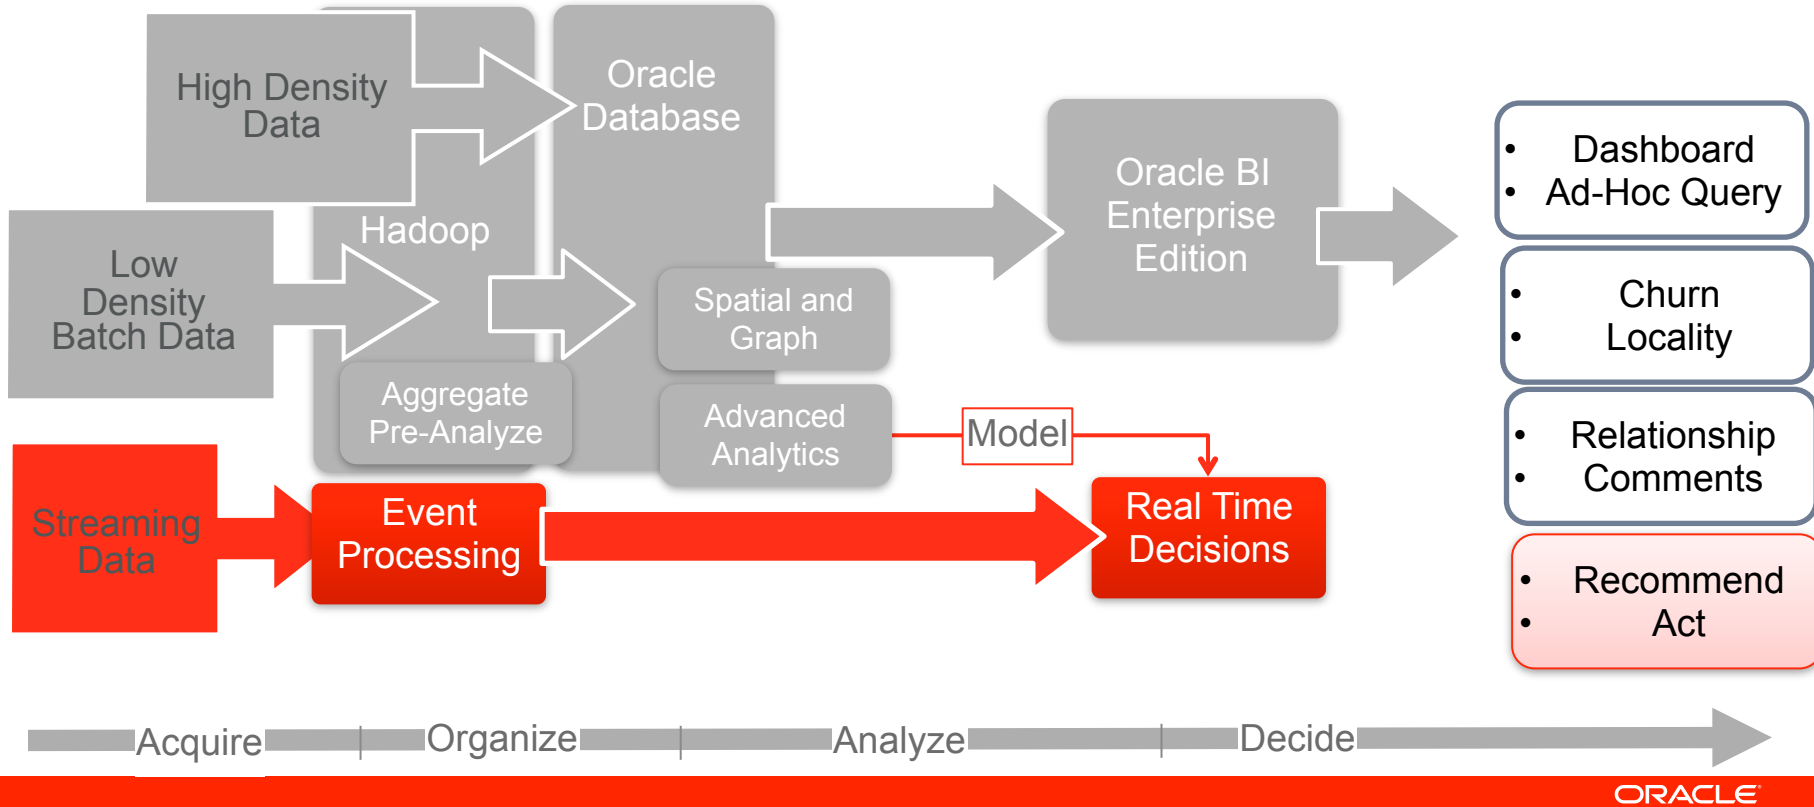


Friday, February 22, 13

How do I influence buying decisions?

How do I prevent serious issues from happening?

Step 3: Architect for Velocity

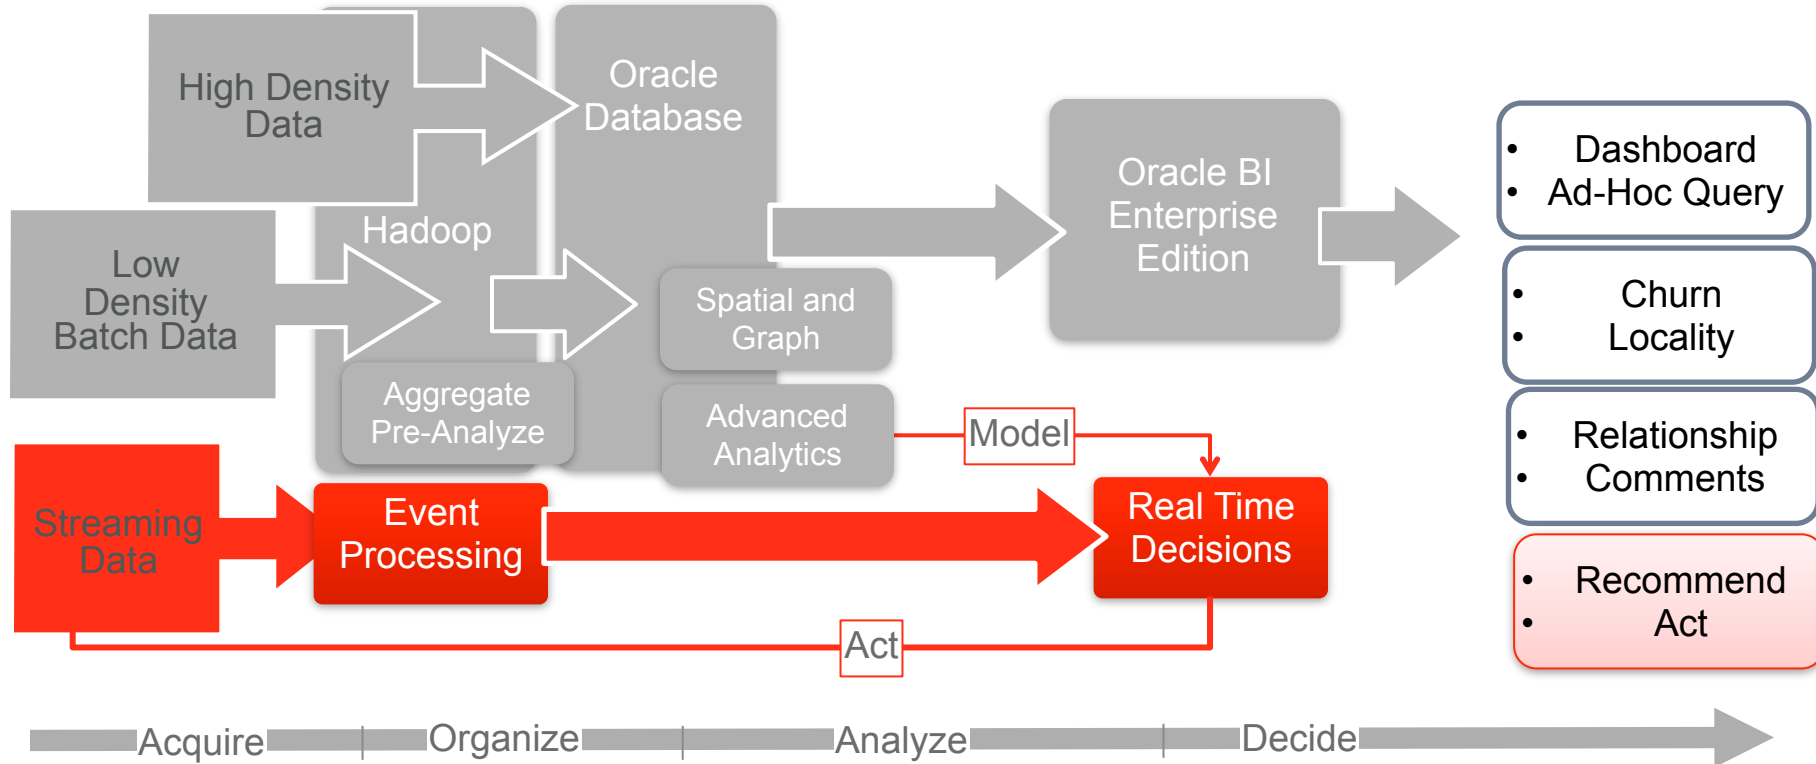


Friday, February 22, 13

How do I influence buying decisions?

How do I prevent serious issues from happening?

Step 3: Architect for Velocity

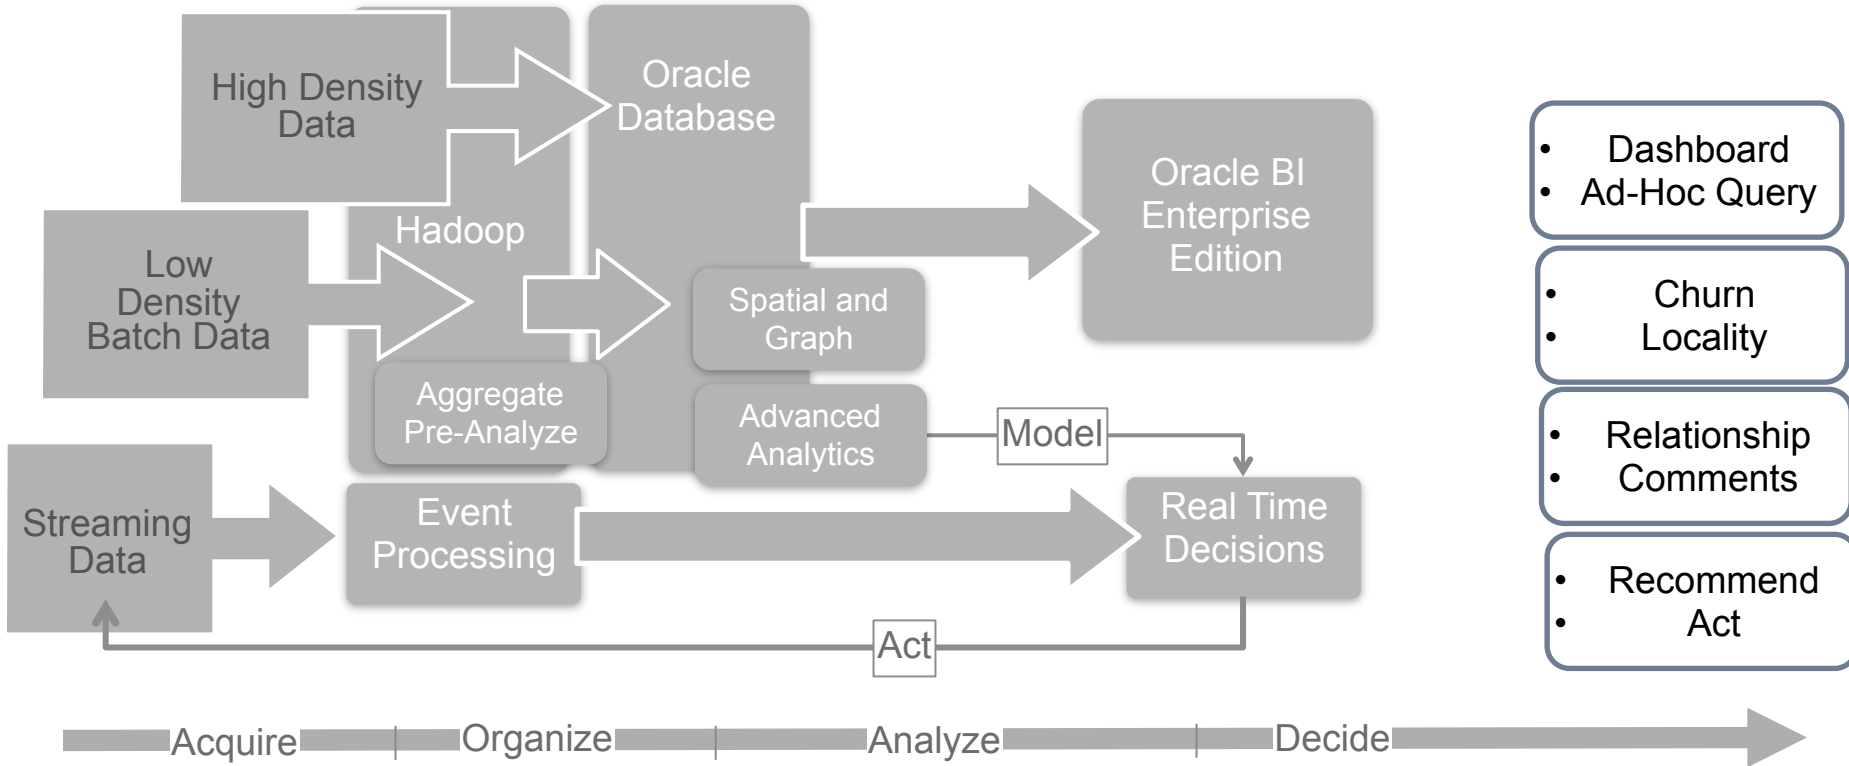


Friday, February 22, 13

How do I influence buying decisions?

How do I prevent serious issues from happening?

Step 4: Discover New Information



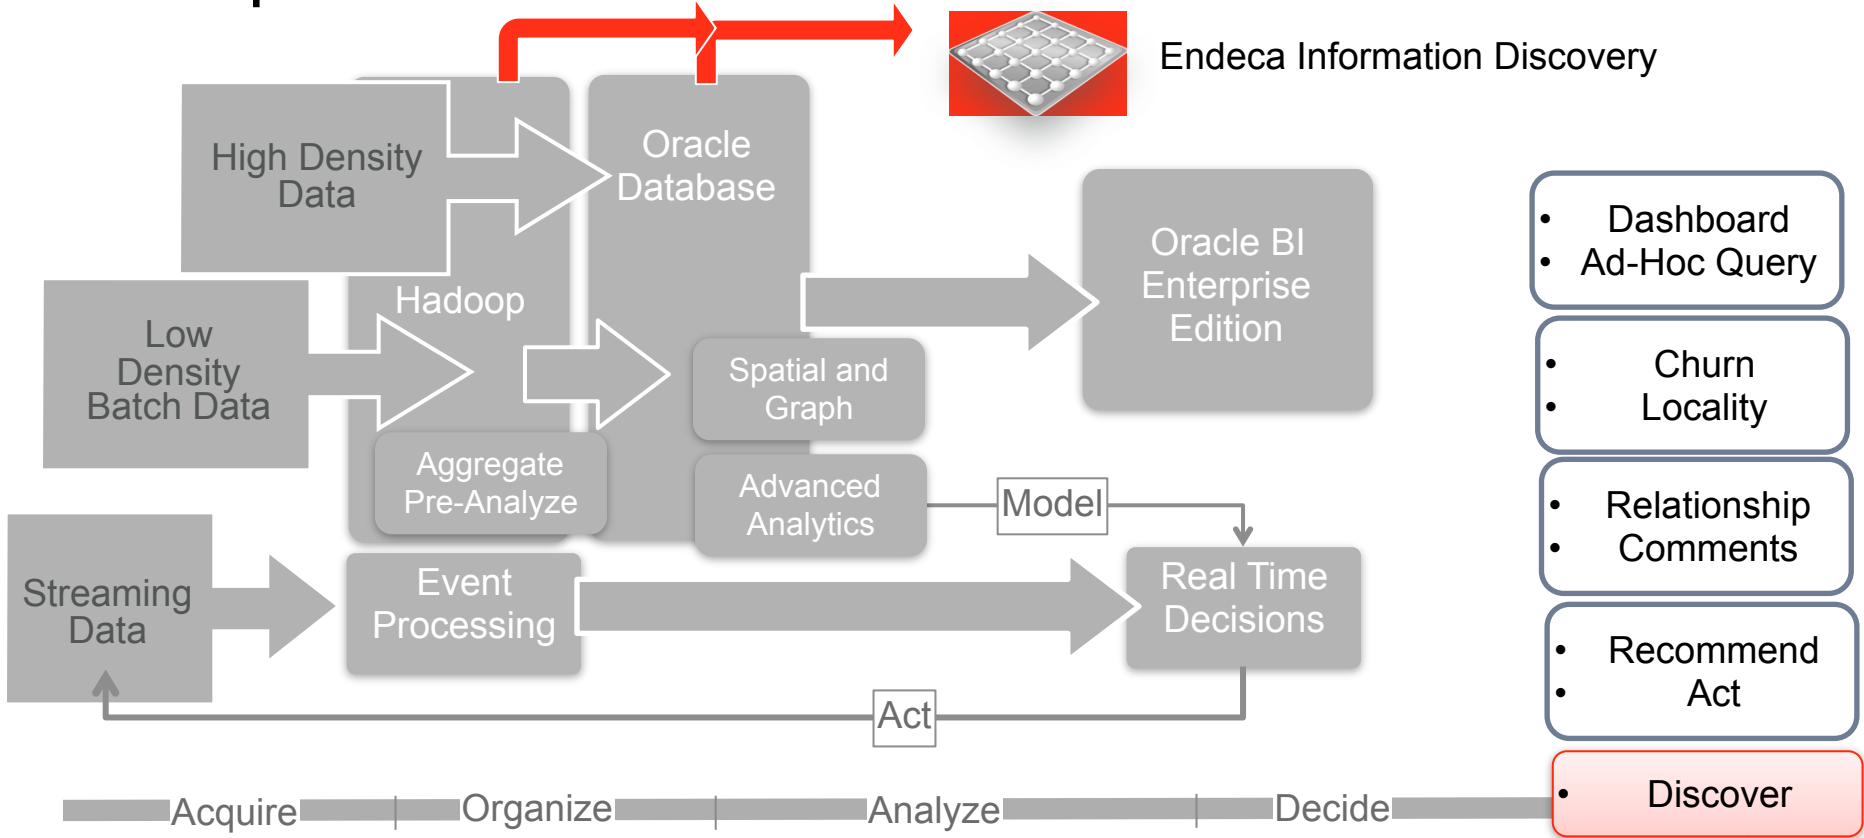
- Dashboard
- Ad-Hoc Query
- Churn
- Locality
- Relationship
- Comments
- Recommend
- Act

Friday, February 22, 13

How do I influence buying decisions?

How do I prevent serious issues from happening?

Step 4: Discover New Information

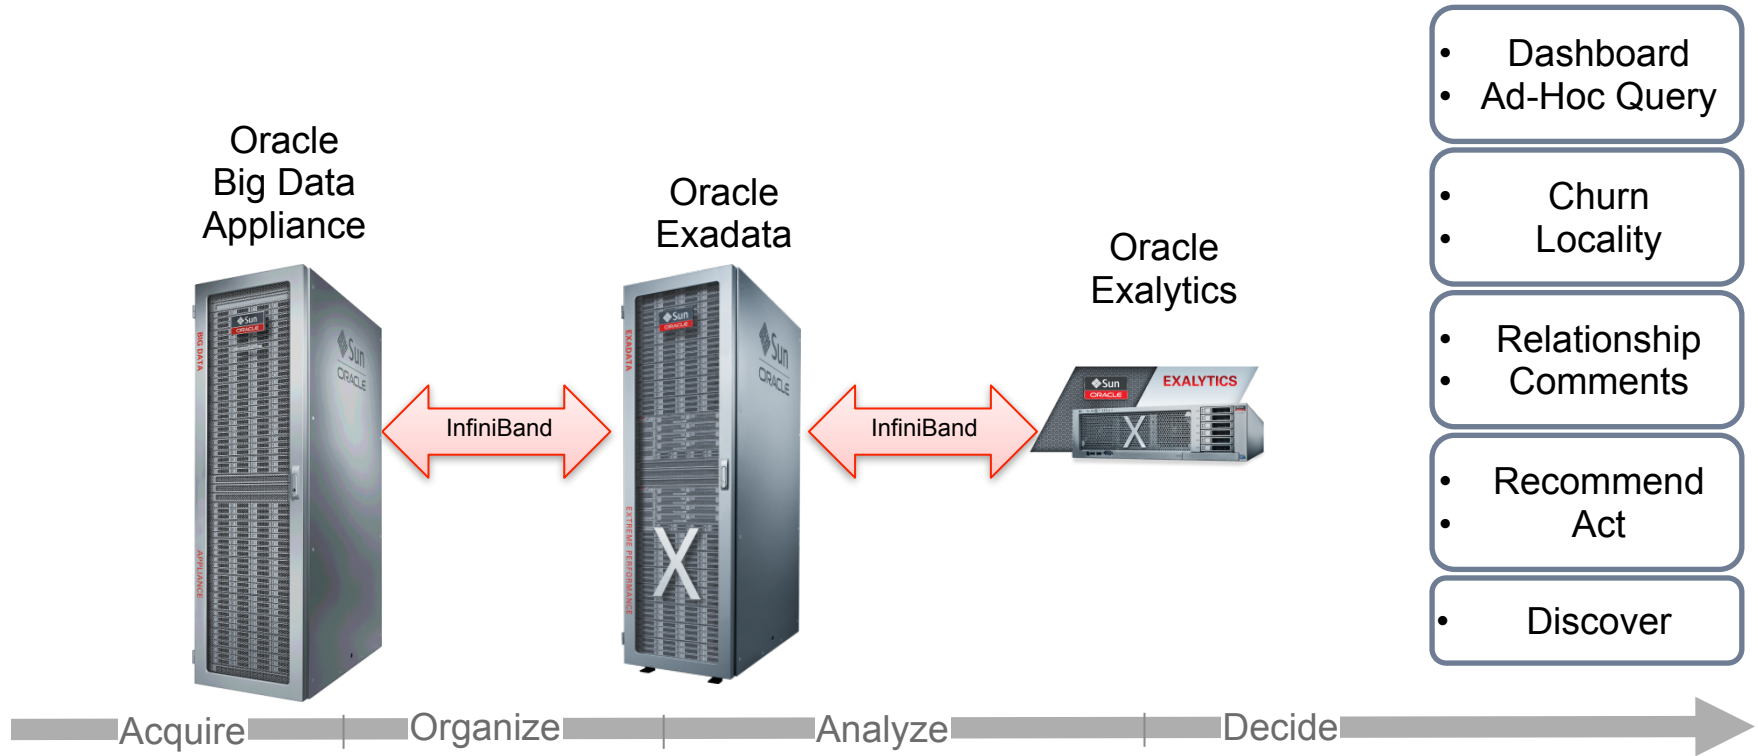


Friday, February 22, 13

How do I influence buying decisions?

How do I prevent serious issues from happening?

Oracle Engineered Systems



ORACLE


Friday, February 22, 13

How do I influence buying decisions?

How do I prevent serious issues from happening?

Oracle Engineered Systems

Simplify IT – Simplify Big Data



- Dashboard ✓
- Ad-Hoc Query ✓
- Churn ✓
- Locality ✓
- Relationship ✓
- Comments ✓
- Recommend ✓
- Act ✓
- Discover ✓

Friday, February 22, 13

How do I influence buying decisions?

How do I prevent serious issues from happening?

Building Big Data Solutions



ORACLE

Building a Big Data Solution

Customer Retention and Social Media

Develop Data Mining models on data in the Data Warehouse

Add Social Data (Twitter) and Relationships (Facebook)

Analyze Social Data and Relationships

Deploy New Retention Policy into the Data Stream

Building a Big Data Solution

Customer Retention and Social Media

Develop Data Mining models on data in the Data Warehouse

Add Social Data (Twitter) and Relationships (Facebook)

Analyze Social Data and Relationships

Deploy New Retention Policy into the Data Stream




Business Value

Predict Customer Churn

Find at-risk customers

Predictive models

Widely used



ORACLE

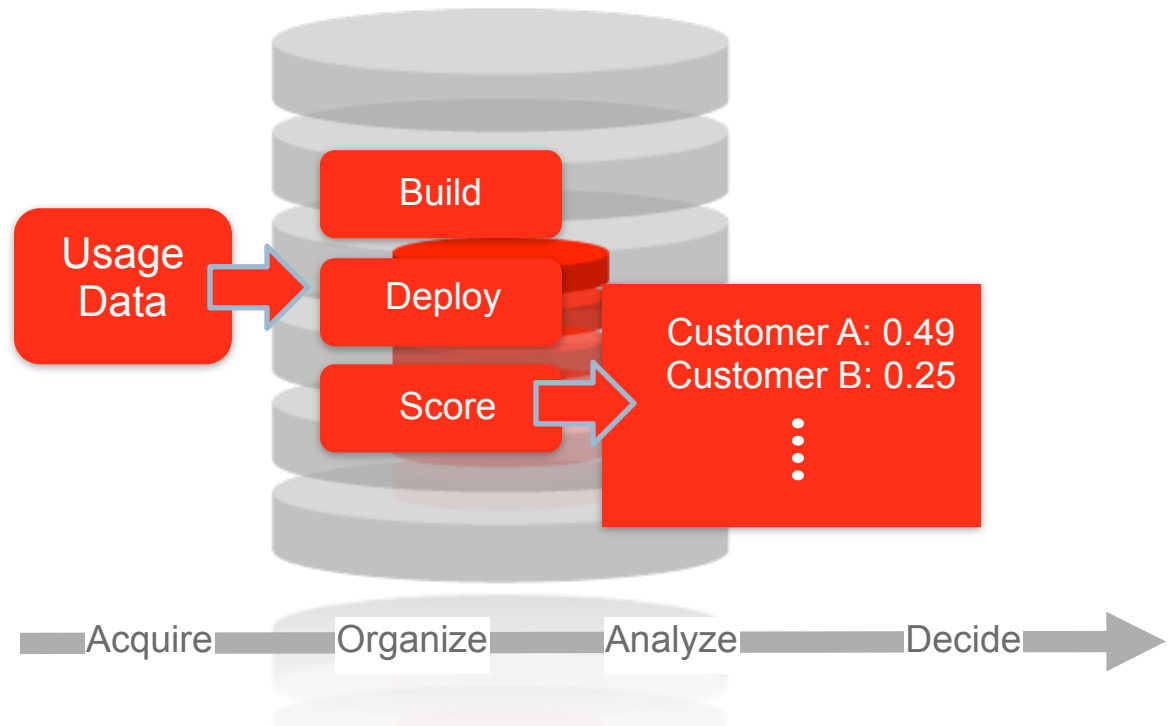
Oracle Advanced Analytics

Oracle Data Mining

12 Algorithms

Extreme performance

GUI for Data Scientists




ORACLE

Understand Your Customers

Find interesting data sets

Process the data to represent meaningful actions

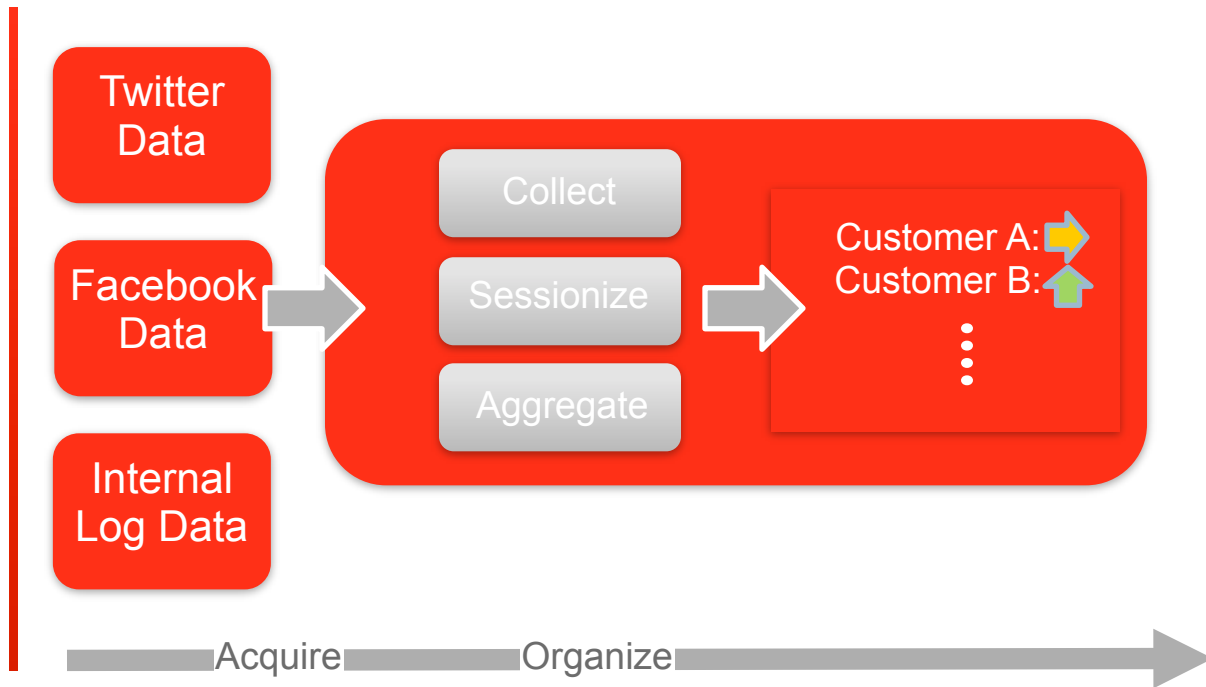
Uncover new relationships



Oracle Big Data Appliance

Big Data Acquisition and Organization

Pre-integrated cluster
Sun Oracle hardware
Cloudera CDH
Massive scale



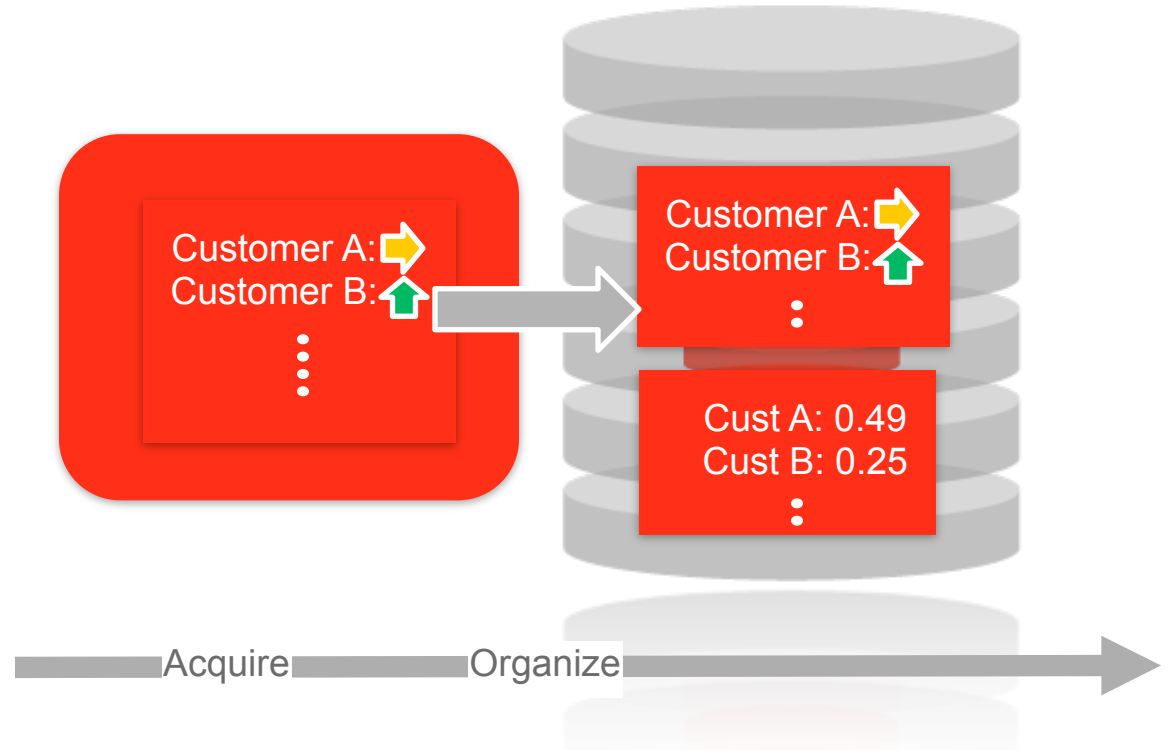
ORACLE

Oracle Big Data Connectors

High performance

Secure

Efficient

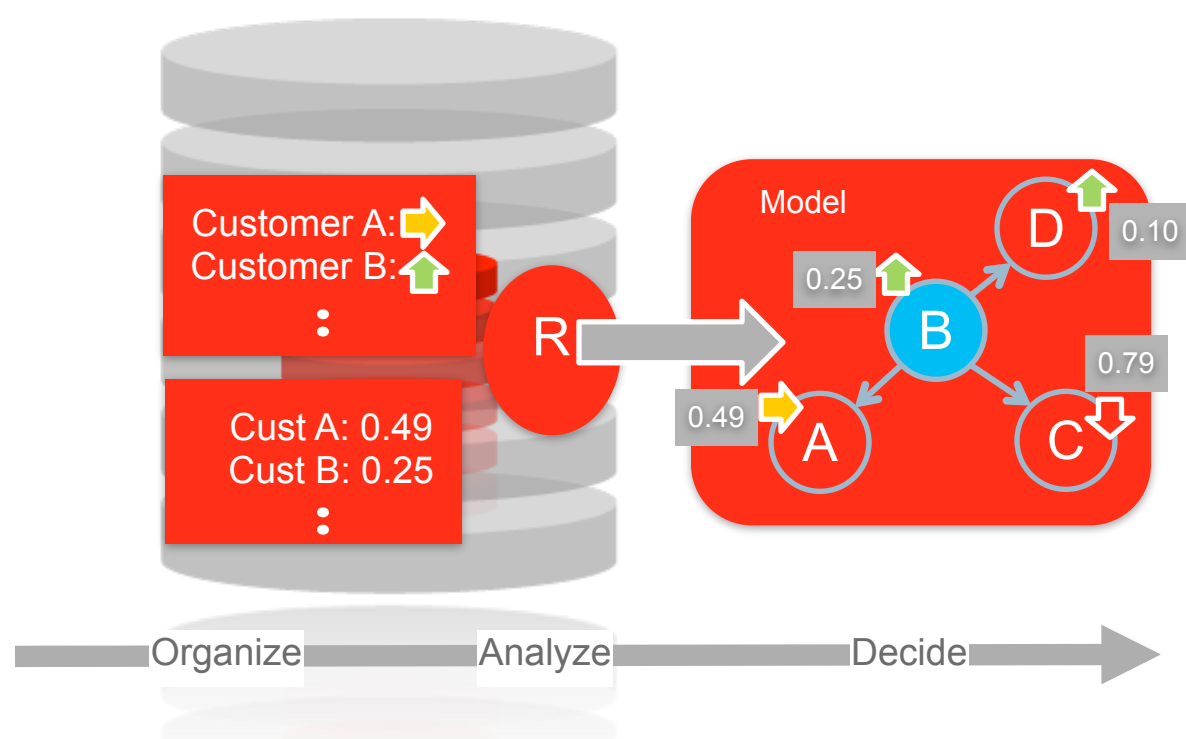


ORACLE

Oracle Advanced Analytics

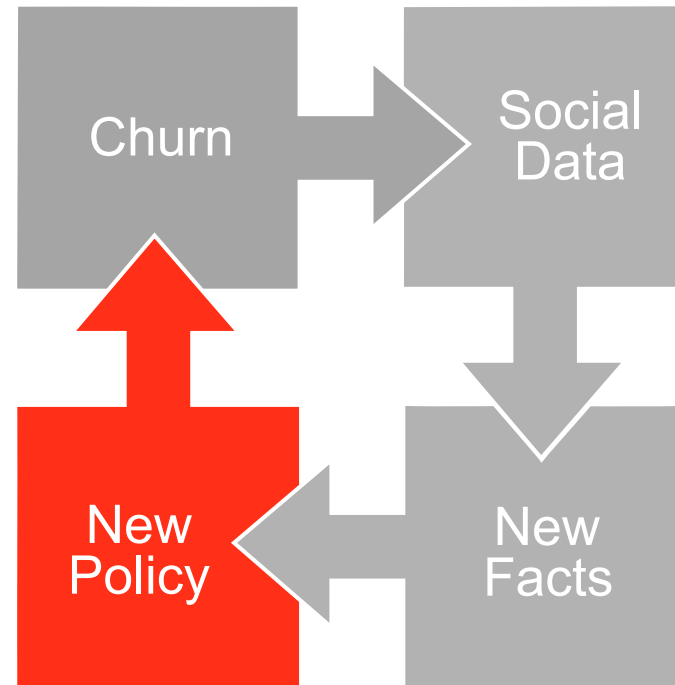
Oracle R Enterprise

- Open source compatible
- Scales to huge data sets
- Extreme performance
- Widely used



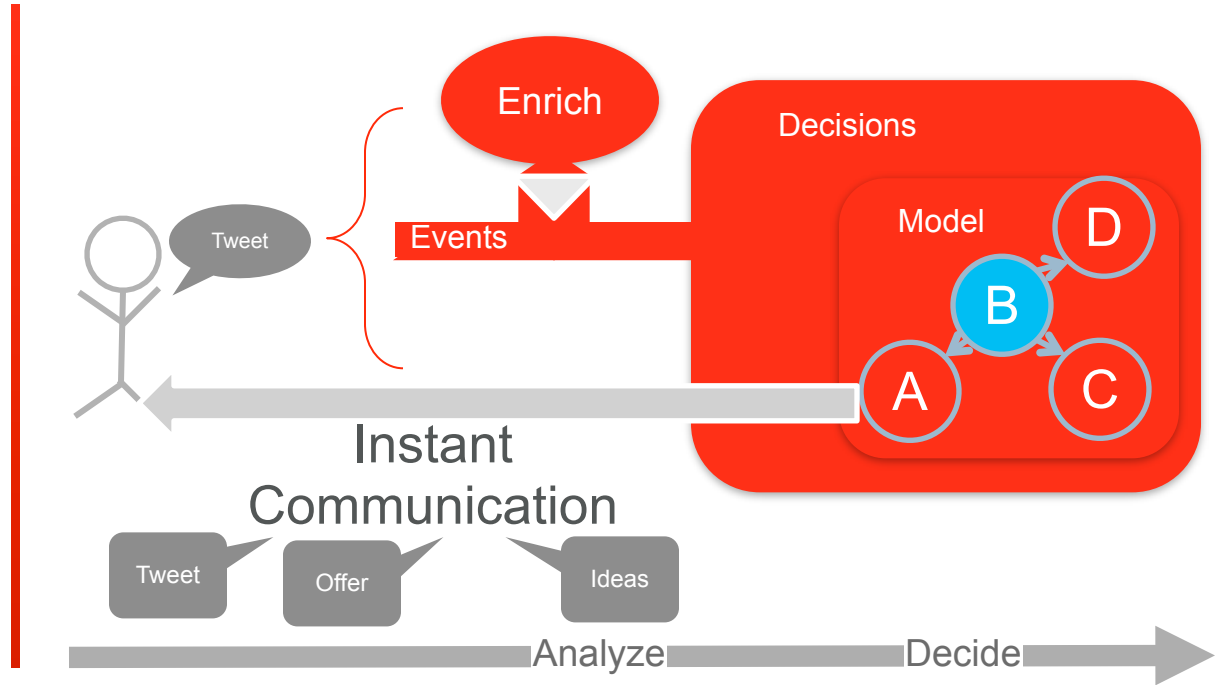
Real-Time Retention Decisions

React to real-time events
Influence key customers
Adjust in real-time



Oracle RTD and Oracle NoSQL Database

Self learning
Fast profile lookups
Integrated



Big Data Products



ORACLE

Oracle Big Data Appliance

Get up and Running
Quickly

Improve Performance of
Hadoop

Integrated with Exadata

Lower TCO for Big Data



ORACLE

Big Data Appliance

Hardware:

- 288 CPU cores with 1152 GB RAM
- 648 TB of raw disk storage
- 40 Gb/s InfiniBand

Integrated Software:

- Oracle Linux
- Oracle Java VM
- Cloudera Distribution of Apache Hadoop (CDH)
- Cloudera Manager
- Open-source distribution of R
- NoSQL Database Community Edition

All integrated software (except NoSQL DB CE) is supported as part of Premier Support for Systems and Premier Support for Operating Systems



ORACLE

Price comparison

Oracle Big Data Appliance

	Year 1	Year 2	Year 3	Total
BDA Cost	\$450,000			
Support Cost	\$54,000	\$54,000	\$54,000	
On-site Installation	\$14,150			
Total	\$518,150	\$54,000	\$54,000	\$626,150

“Build-Your-Own” – HP hardware and Cloudera


	Year 1	Year 2	Year 3	Total
Servers and switches	\$428,220			
Support Cost	\$136,233	\$72,000	\$72,000	
Installation & configuration not included				
Total	\$564,453	\$72,000	\$72,000	\$708,453

Full details at https://blogs.oracle.com/datawarehousing/entry/price_comparison_for_big_data


ORACLE

Big Data Appliance performance comparisons

6x faster than custom 20-node Hadoop cluster for large batch transformation jobs



2.5x faster than 30-node Hadoop cluster for tagging and parsing text documents



Big Data Connectors

Optimized integration of Hadoop with Oracle Database and Oracle Exadata

- Oracle Loader for Hadoop
 - Oracle Direct 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



ORACLE

Customer Stories



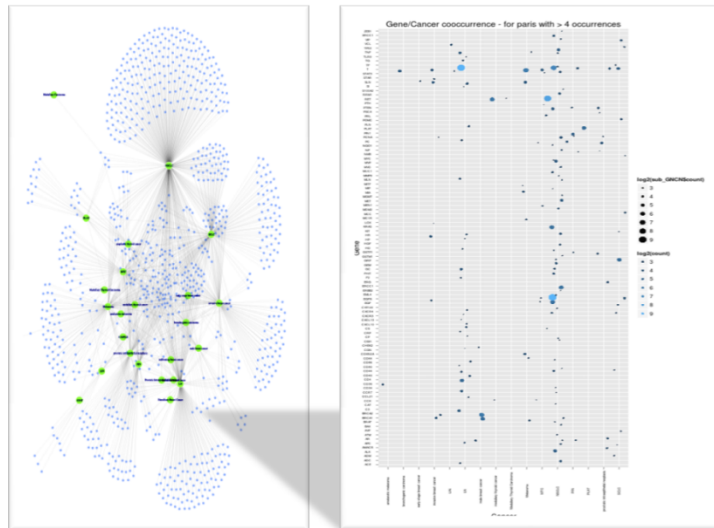
ORACLE

National Cancer Institute and Oracle

2012 Judge's Pick: National Cancer Institute, Oracle, and SAIC found relationships between Gene to Cancer interaction

- Cross-referenced the relationships between 17000 genes and five major cancer types across 20 million medical publication abstracts
- Cross-referenced genes from 60 Million patients and miRNA for a simulated 900 Million population.
- Understanding additional layers of the pathways these genes operate in and the drugs that target them is expected to help researchers in their work

CTOLabs.com



9

ORACLE

Thomson Reuters on Big Data



ORACLE

Big Data Appliance Scenarios

Sample: Current Customers and Active POCs

Travel Services	Optimize travel buying experience through web log analysis, previously analyzed customers purchases, but will additionally analyze all travel itineraries that were displayed
Information Services	Analysis of application logs (how do customers use their data services?) Interconnected with Exadata
High Technology	Analyze usage of features and applications on devices based on log data Compare BDA performance to existing build-you-own Hadoop infrastructure Interconnected with Exadata
Travel Services	Preprocess log data for loading into data warehouse Interconnected with Exadata
Financial Services	Ensure compliance and improve customer interactions by analyzing emails, call transcripts and other customer interactions Integrated with third-party NLP solution
CPG Industry	Combining POS data at the product level with external and social data to build a detailed understanding of why category revenue declines and gains occurred
Gaming	Analyzing console data streams to understand game play behavior (Re-) Design games to drive more in-game revenue

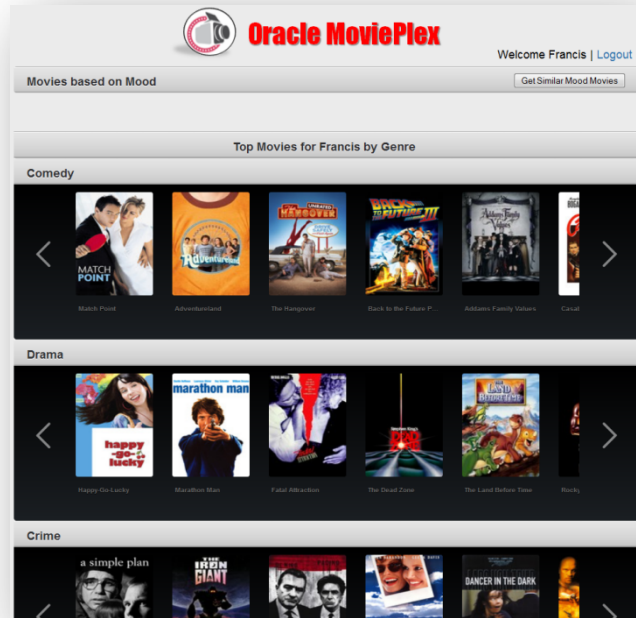
Big Data End-to-End



ORACLE

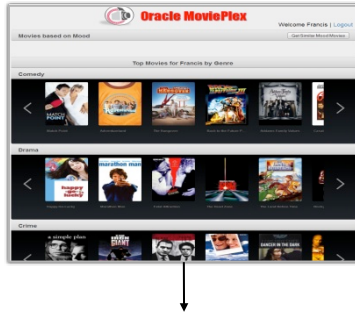
Oracle MoviePlex

Introduction



- Oracle MoviePlex is an on-line movie streaming company
- Like many other on-line stores, they needed a cost effective approach to tackle their “big data” challenges
- They recently implemented Oracle’s Big Data Platform to better manage their business, identify key opportunities and enhance customer satisfaction

Big Data Challenge



- Applications are generating massive volumes of unstructured data that describe user behavior and application performance
- Today, most companies are unable to fully capitalize on this potentially valuable information due to cost and complexity
- How do you capitalize on this raw data to gain better insights into your customers, enhance their user experience and ultimately improve profitability?

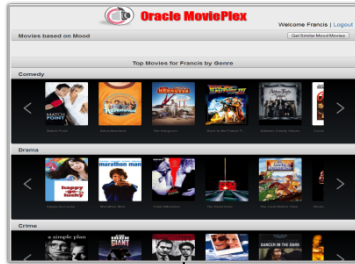
Big Data Challenge



```
{ "custId": 1185972, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:07", "recommended": null, "activity": 8 }
{ "custId": 1354924, "movieId": 1948, "genreId": 9, "time": "2012-07-01:00:00:22", "recommended": "N", "activity": 7 }
{ "custId": 1083711, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:26", "recommended": null, "activity": 9 }
{ "custId": 1234182, "movieId": null, "genreId": 44, "time": "2012-07-01:00:00:32", "recommended": "Y", "activity": 7 }
{ "custId": 1010220, "movieId": 11547, "genreId": 44, "time": "2012-07-01:00:00:42", "recommended": "Y", "activity": 6 }
{ "custId": 1143971, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:43", "recommended": null, "activity": 8 }
{ "custId": 1253676, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:50", "recommended": null, "activity": 9 }
{ "custId": 1351777, "movieId": 608, "genreId": 6, "time": "2012-07-01:00:01:03", "recommended": "N", "activity": 7 }
{ "custId": 1143971, "movieId": null, "genreId": null, "time": "2012-07-01:00:01:07", "recommended": null, "activity": 9 }
{ "custId": 1363545, "movieId": 27205, "genreId": 9, "time": "2012-07-01:00:01:18", "recommended": "Y", "activity": 7 }
{ "custId": 1067283, "movieId": 1124, "genreId": 9, "time": "2012-07-01:00:01:26", "recommended": "Y", "activity": 7 }
{ "custId": 1126174, "movieId": 16309, "genreId": 9, "time": "2012-07-01:00:01:35", "recommended": "N", "activity": 7 }
{ "custId": 1234182, "movieId": 11547, "genreId": 44, "time": "2012-07-01:00:01:39", "recommended": "Y", "activity": 7 }
{ "custId": 1067283, "movieId": null, "genreId": null, "time": "2012-07-01:00:01:55", "recommended": null, "activity": 9 }
{ "custId": 1377537, "movieId": null, "genreId": null, "time": "2012-07-01:00:01:58", "recommended": null, "activity": 9 }
{ "custId": 1347836, "movieId": null, "genreId": null, "time": "2012-07-01:00:02:03", "recommended": null, "activity": 8 }
{ "custId": 1137285, "movieId": null, "genreId": null, "time": "2012-07-01:00:03:39", "recommended": null, "activity": 8 }
{ "custId": 1354924, "movieId": null, "genreId": null, "time": "2012-07-01:00:03:51", "recommended": null, "activity": 8 }
{ "custId": 1036191, "movieId": null, "genreId": null, "time": "2012-07-01:00:03:55", "recommended": null, "activity": 8 }
{ "custId": 1143971, "movieId": 1017161, "genreId": 44, "time": "2012-07-01:00:04:00", "recommended": "Y", "activity": 7 }
{ "custId": 1363545, "movieId": 27205, "genreId": 9, "time": "2012-07-01:00:04:03", "recommended": "Y", "activity": 5 }
{ "custId": 1273464, "movieId": null, "genreId": null, "time": "2012-07-01:00:04:39", "recommended": null, "activity": 9 }
{ "custId": 1346299, "movieId": 424, "genreId": 1, "time": "2012-07-01:00:05:02", "recommended": "Y", "activity": 4 }
```

- Applications are generating massive volumes of unstructured data that describe user behavior and application performance
- Today, most companies are unable to fully capitalize on this potentially valuable information due to cost and complexity
- How do you capitalize on this raw data to gain better insights into your customers, enhance their user experience and ultimately improve profitability?

Big Data Challenge



```
{ "custId": 1185972, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:07", "recommended": null, "activity": 8 }
{ "custId": 1354924, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:22", "recommended": "N", "activity": 7 }
{ "custId": 1083711, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:26", "recommended": null, "activity": 9 }
{ "custId": 1234182, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:32", "recommended": "Y", "activity": 7 }
{ "custId": 1010220, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:42", "recommended": "Y", "activity": 6 }
{ "custId": 1143971, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:43", "recommended": null, "activity": 8 }
{ "custId": 1253676, "movieId": null, "genreId": null, "time": "2012-07-01:00:00:50", "recommended": null, "activity": 9 }
```

How can you get answers to....

```
{ "custId": 1007285, "movieId": null, "genreId": null, "time": "2012-07-01:00:01:05", "recommended": null, "activity": 5 }
{ "custId": 1377537, "movieId": null, "genreId": null, "time": "2012-07-01:00:01:58", "recommended": null, "activity": 9 }
{ "custId": 1347836, "movieId": null, "genreId": null, "time": "2012-07-01:00:02:03", "recommended": null, "activity": 8 }
{ "custId": 1137285, "movieId": null, "genreId": null, "time": "2012-07-01:00:03:39", "recommended": null, "activity": 8 }
{ "custId": 1354924, "movieId": null, "genreId": null, "time": "2012-07-01:00:03:51", "recommended": null, "activity": 9 }
{ "custId": 1036191, "movieId": null, "genreId": null, "time": "2012-07-01:00:03:55", "recommended": null, "activity": 8 }
{ "custId": 1143971, "movieId": null, "genreId": null, "time": "2012-07-01:00:04:00", "recommended": "Y", "activity": 7 }
{ "custId": 1363545, "movieId": 27205, "genreId": 9, "time": "2012-07-01:00:04:03", "recommended": "Y", "activity": 5 }
{ "custId": 1273464, "movieId": null, "genreId": null, "time": "2012-07-01:00:04:39", "recommended": null, "activity": 9 }
{ "custId": 1346299, "movieId": 424, "genreId": 1, "time": "2012-07-01:00:05:02", "recommended": "Y", "activity": 4 }
```

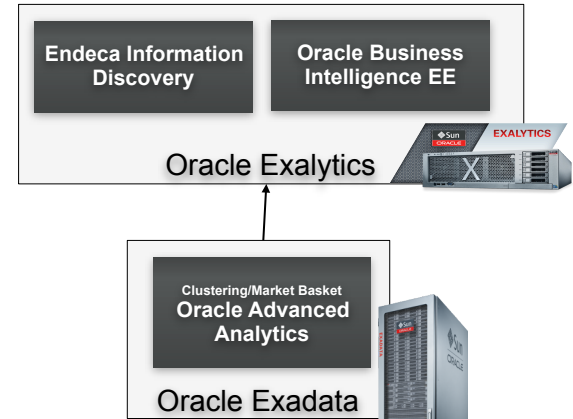
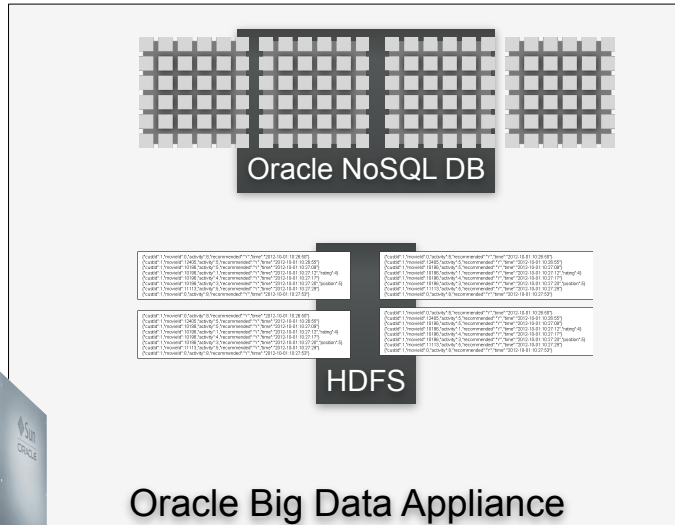
- Applications are generating massive volumes of unstructured data that describe user behavior and application performance
- Today, most companies are unable to fully capitalize on this potentially valuable information due to cost and complexity
- How do you capitalize on this raw data to gain better insights into your customers, enhance their user experience and ultimately improve profitability?

Derive Value from Big Data

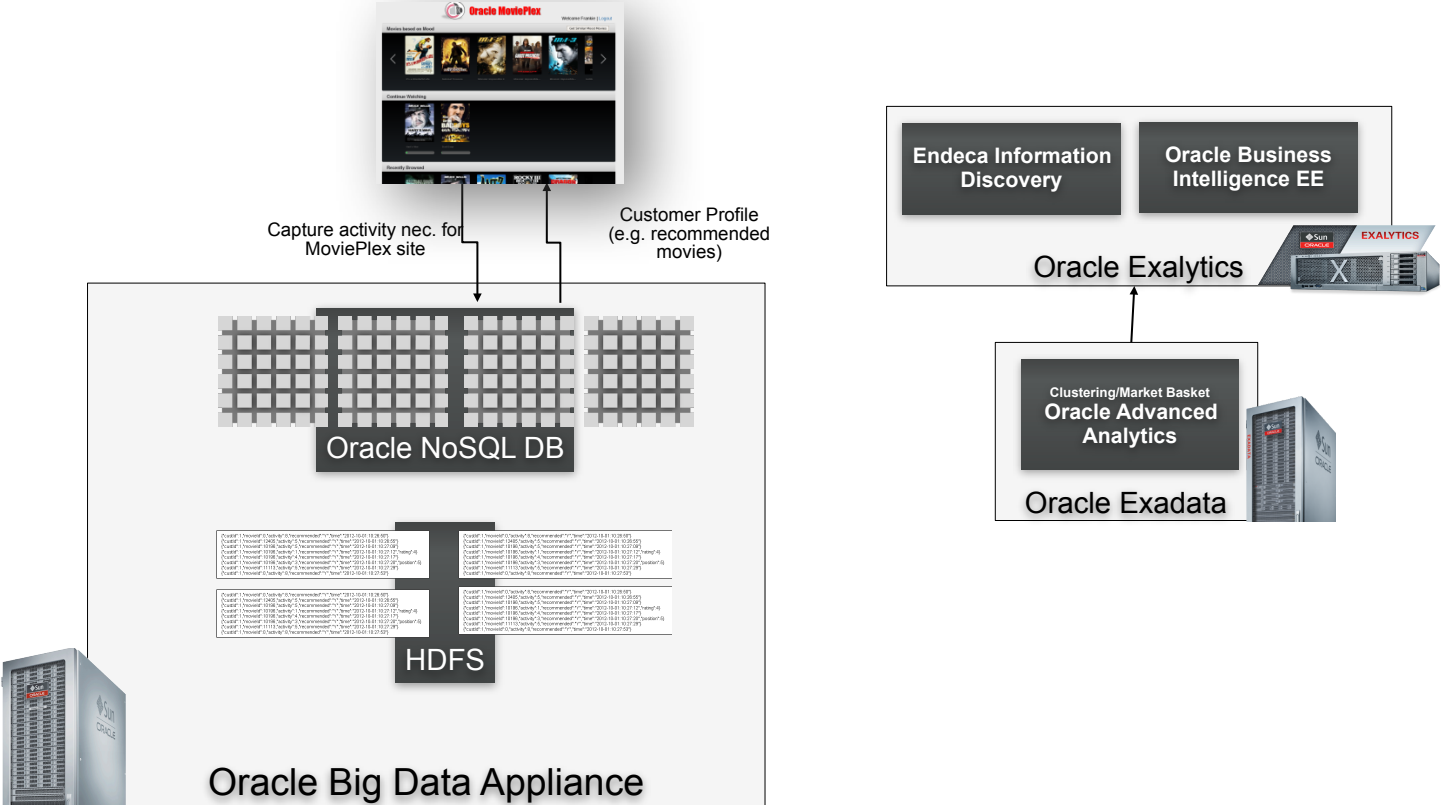
How can you

- Make the right movie offers at the right time?
- Better understand the viewing trends of various customer segments?
- Optimize marketing spend by targeting customers with optimal promotional offers?
- Minimize infrastructure spend by understanding bandwidth usage over time?
- Prepare to answer questions that you haven't thought of yet!

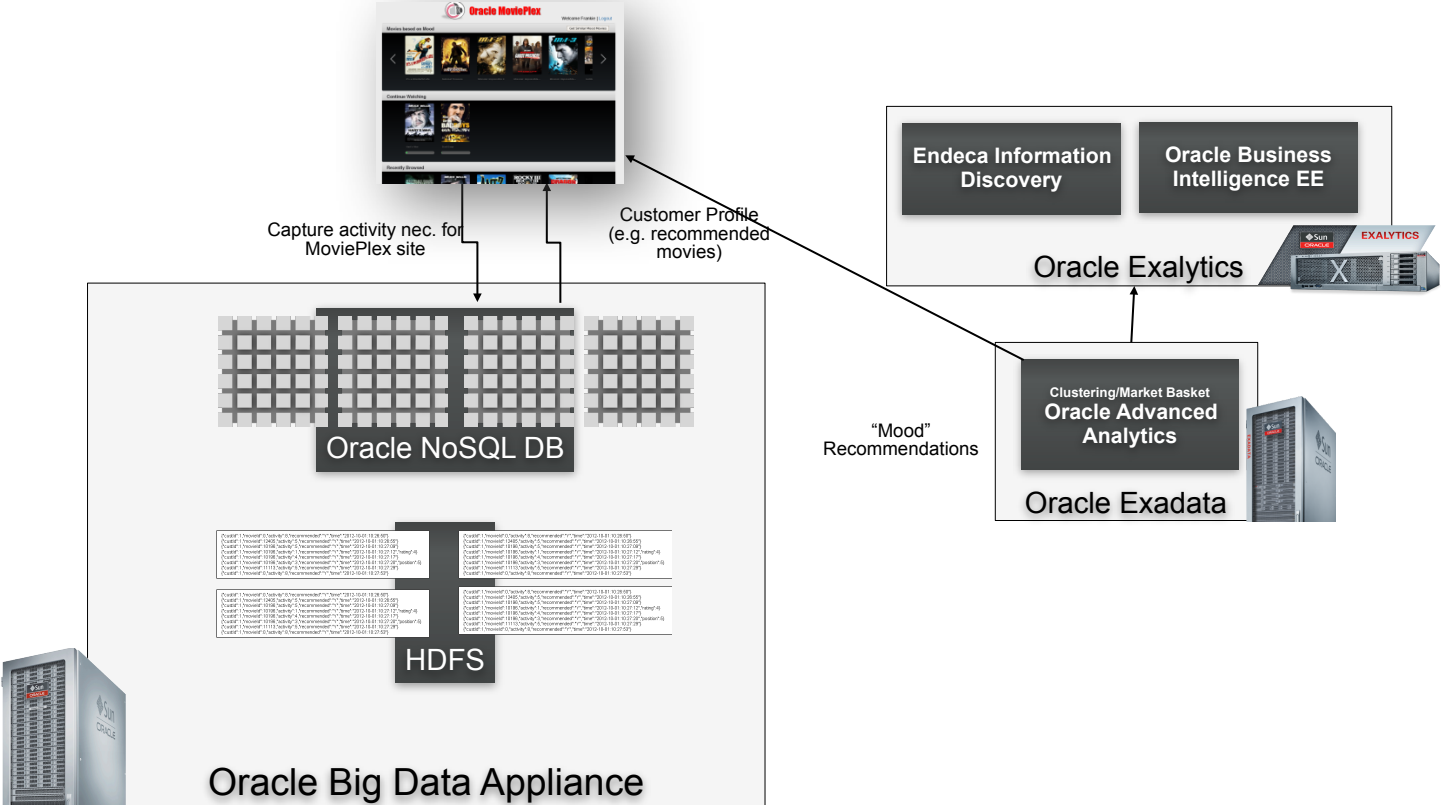
MoviePlex Architecture



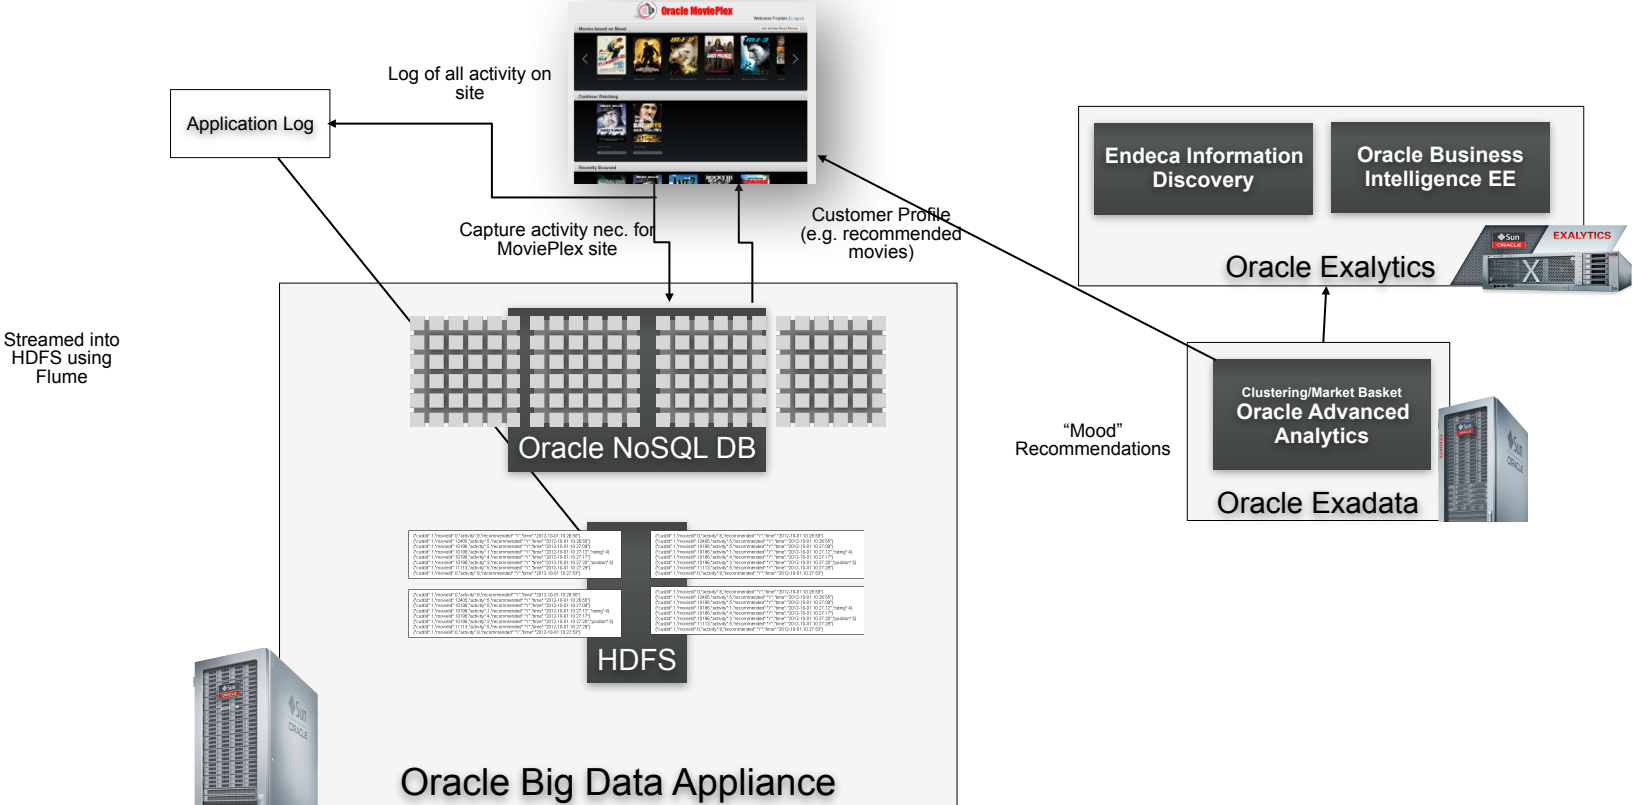
MoviePlex Architecture



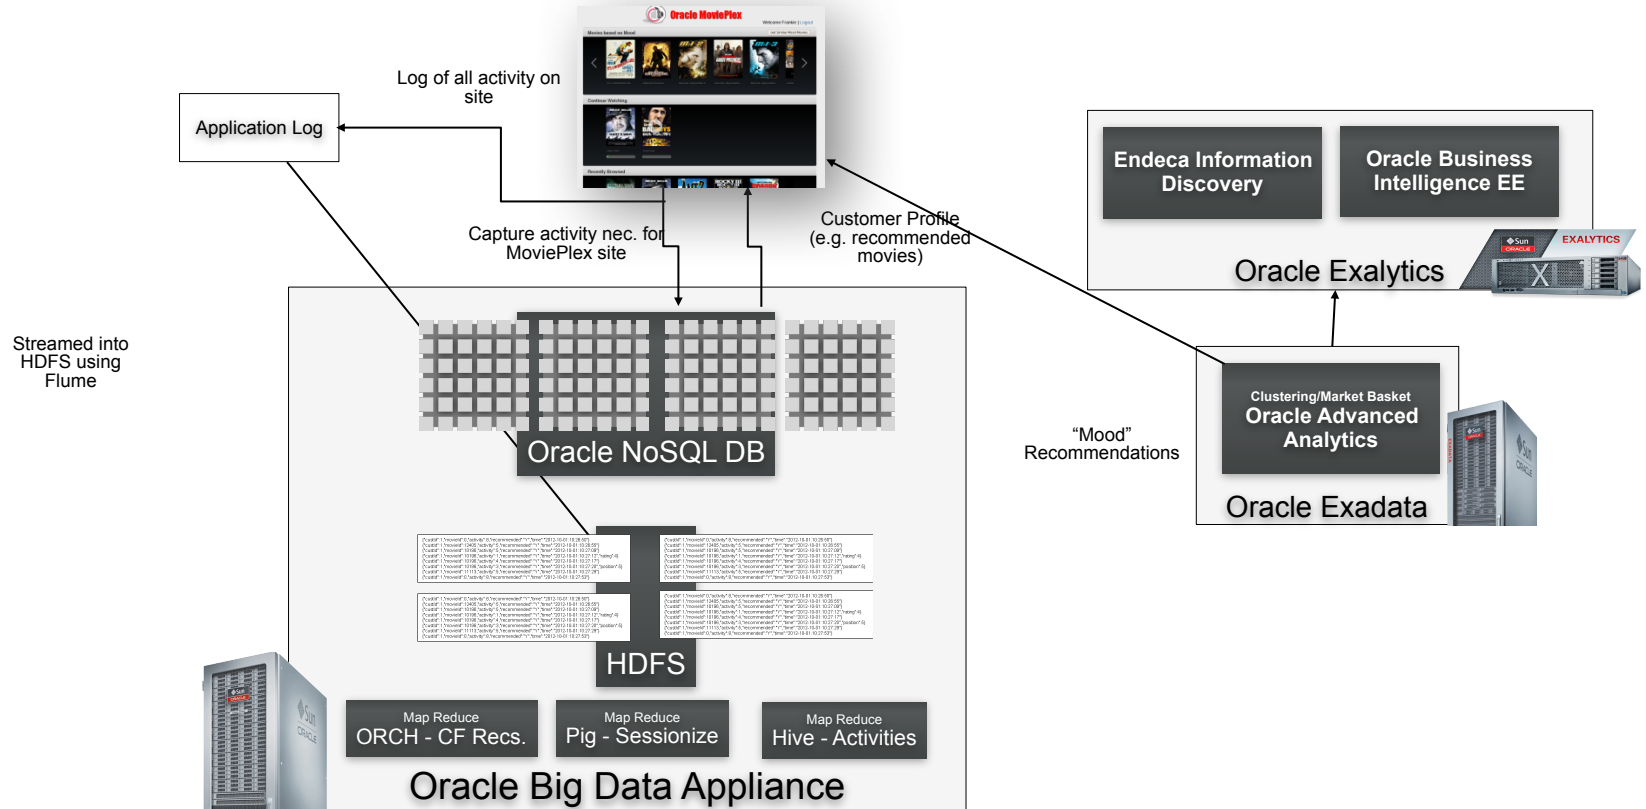
MoviePlex Architecture



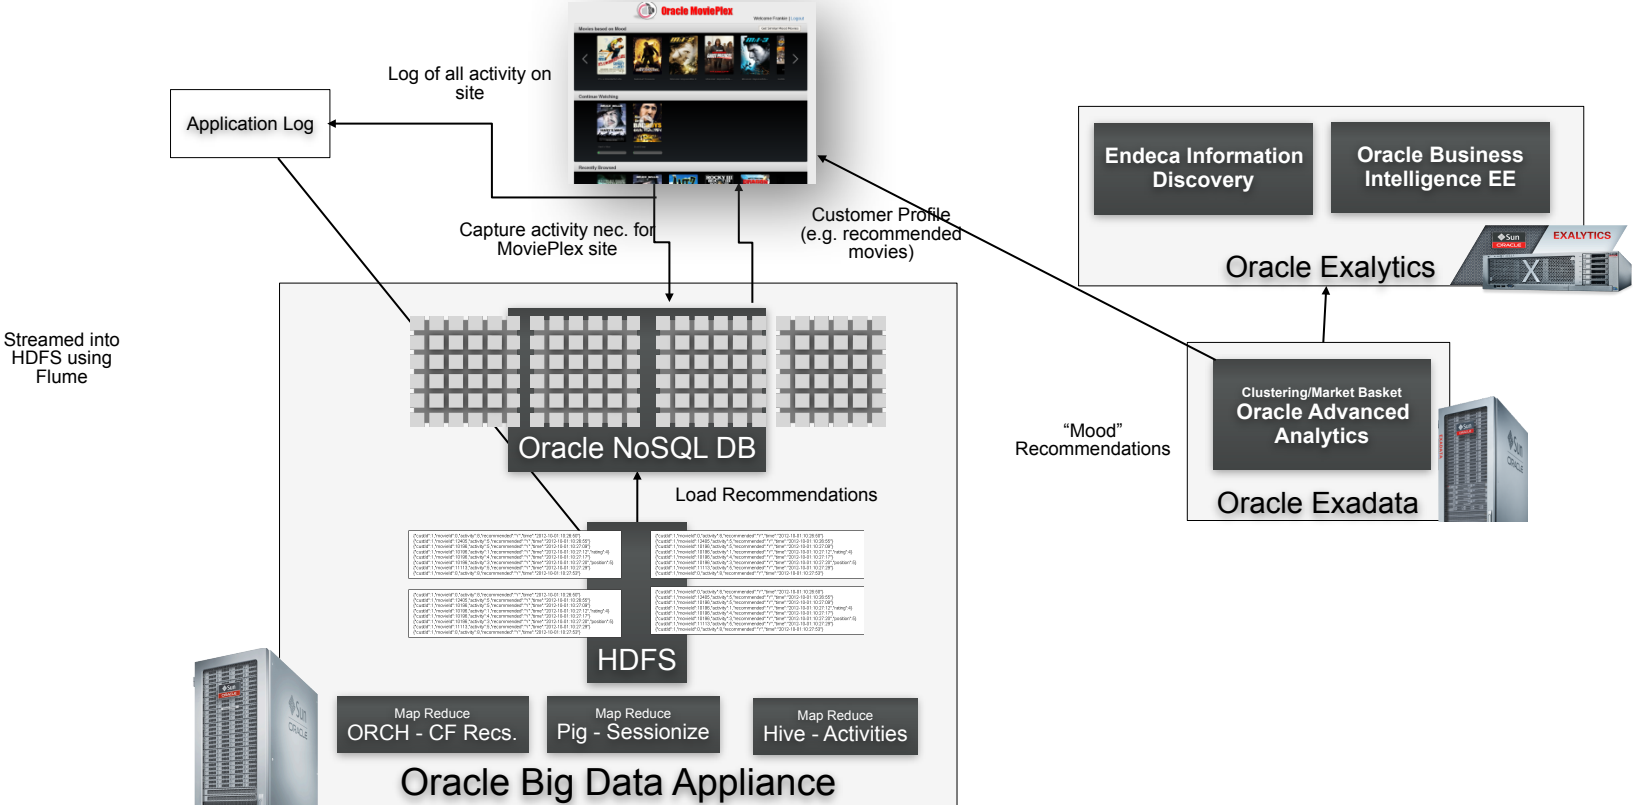
MoviePlex Architecture



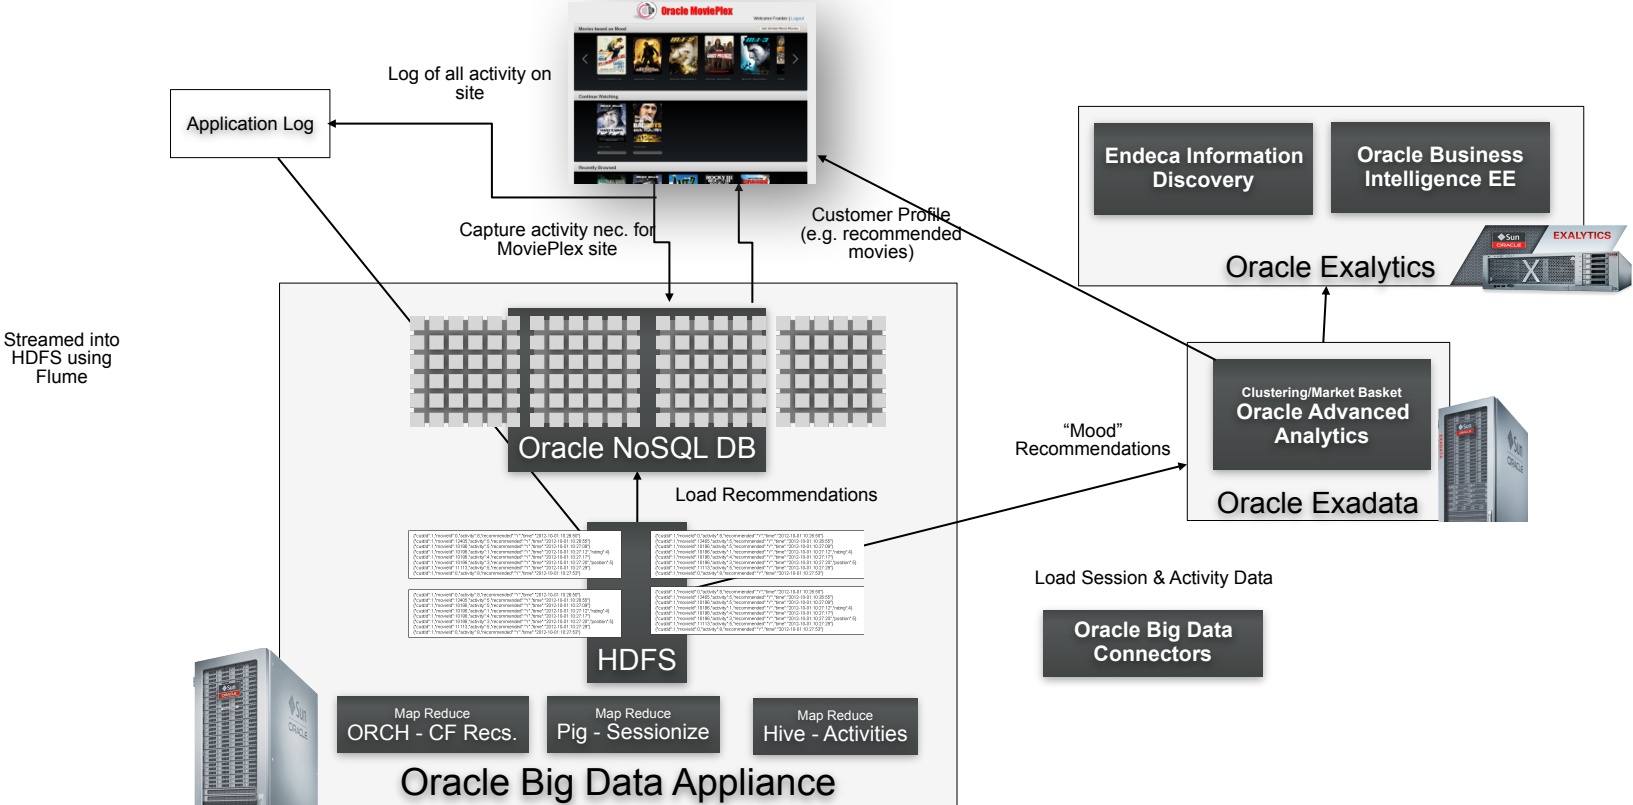
MoviePlex Architecture



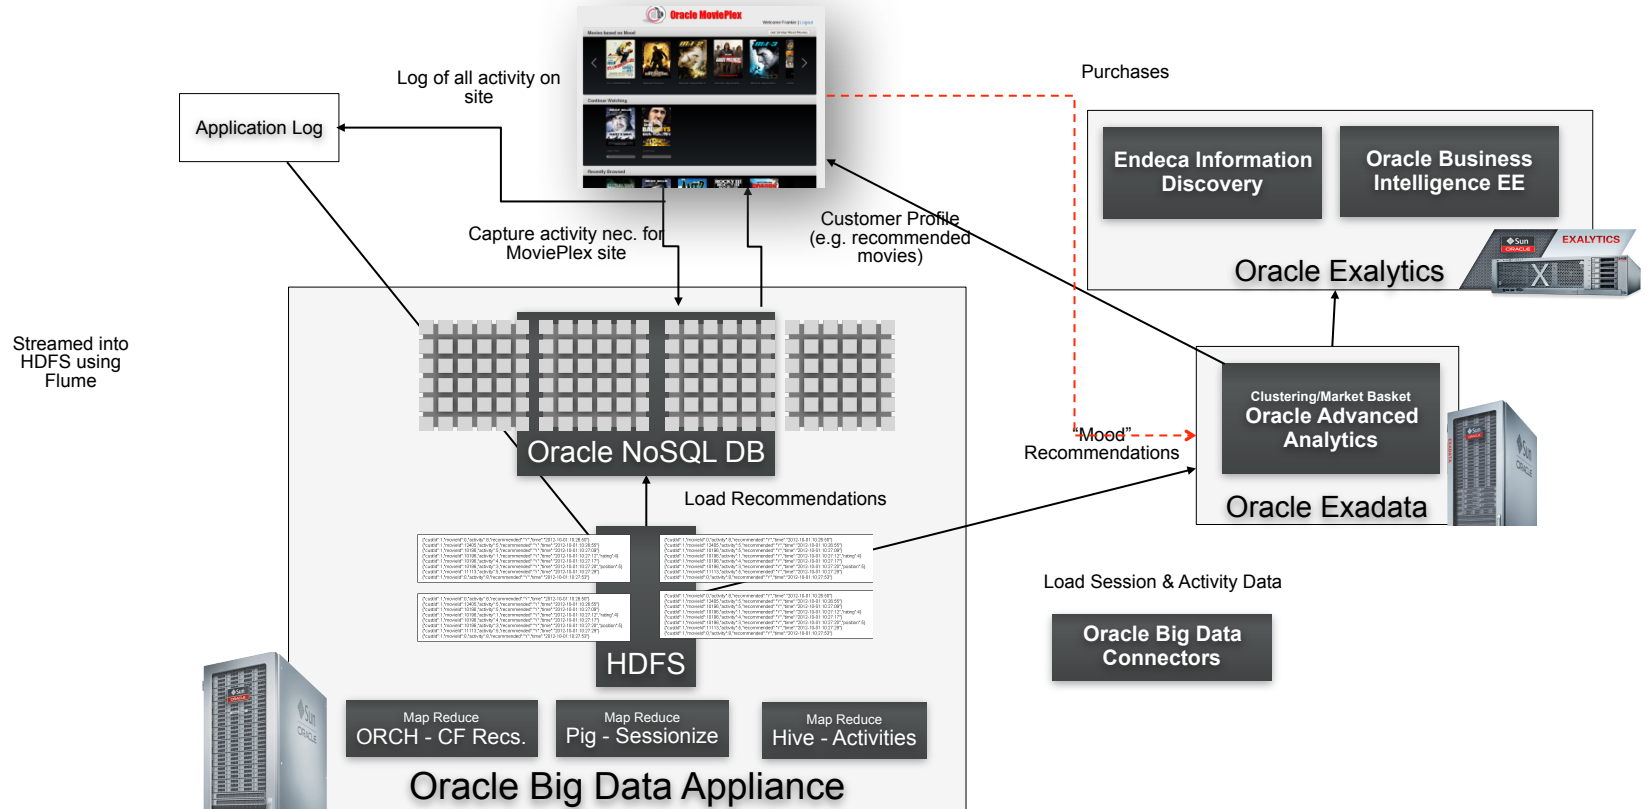
MoviePlex Architecture



MoviePlex Architecture



MoviePlex Architecture



Hardware and Software



Engineered to Work Together

ORACLE®