

*Smarter Decision Making  
Leverage Big Data to Gain New Actionable Insights*

George Lapis



# The World is Changing and Becoming More...




INSTRUMENTED



INTERCONNECTED



INTELLIGENT



The resulting explosion of information creates a need for a new kind of intelligence

*...to help build a Smarter Planet*

# The Social Layer in an Instrumented Interconnected World

12+ TBs  
of tweet data  
every day




? TBs of  
data every day



25+ TBs of  
log data every  
day



30 billion RFID  
tags today  
(1.3B in 2005)



4.6  
billion  
camera  
phones  
world wide



100s of  
millions  
of GPS  
enabled  
devices  
sold  
annually



76 million smart  
meters in 2009...  
200M by 2014





http://www.

2+  
billion  
people on  
the Web  
by end  
2011

**2011**  
1.8 ZB

**2009**  
1 ZB  
1 ZB=1T

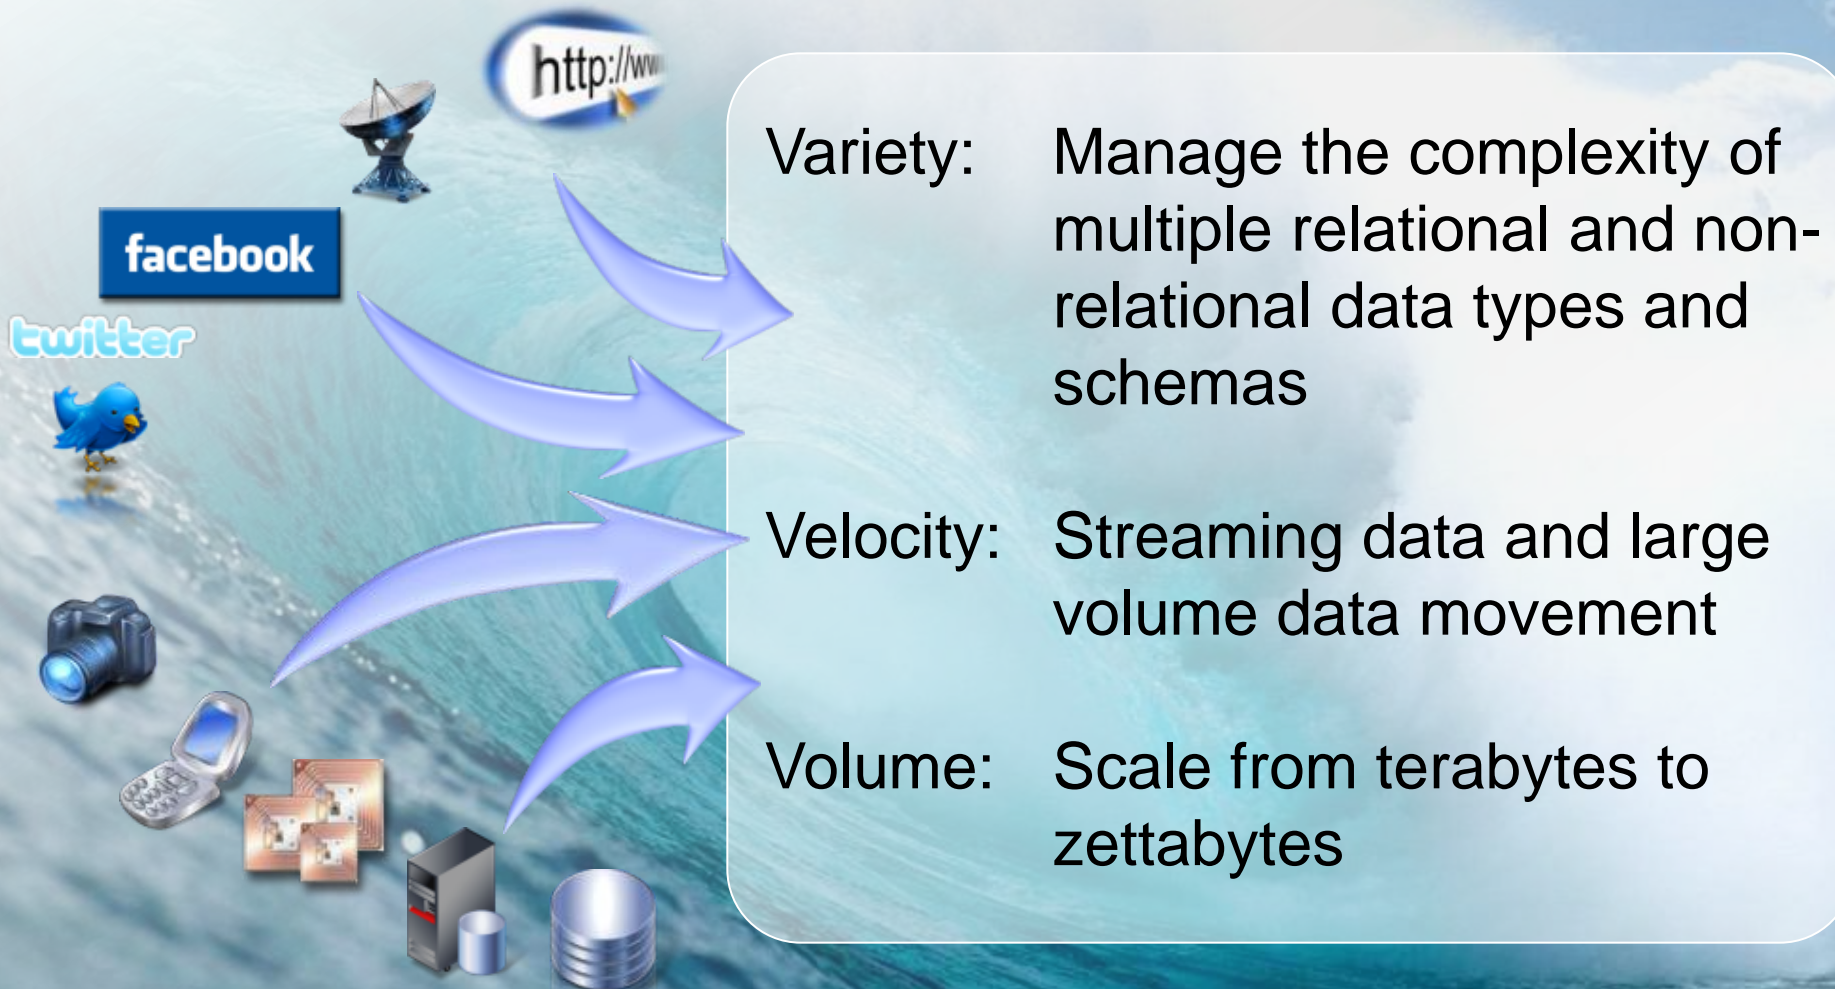
**2010**  
**2020**






# The Big Data Opportunity

*Extracting insight from an immense volume, variety and velocity of data, in context, beyond what was previously possible.*




# Benchmark of Global Big Data Activities (Oct 2012)

## Realizing a competitive advantage




Nearly two out of three respondents reports realizing a **competitive advantage** from information and analytics

## Big data activities



Three out of four organizations have **big data activities** underway; and one in four are either in **pilot or production**

## Big data objectives



Improving **customer experience** by better understanding behaviors drives almost half of all active big data efforts followed by **Operational Optimization**

[www.ibm.com/2012bigdatastudy](http://www.ibm.com/2012bigdatastudy)

IBM Institute for Business Value and the University of Oxford Saïd Business School

# The new era of analytics delivers value across the enterprise

## Customer Service Representatives

...offer personalized price promotions to different customer segments in real-time

## Network Operations

...identify network bottlenecks in real-time for faster resolution



External Data

## Executive Leaders

...get real-time reports and analysis based on data inside as well as outside the enterprise (web, social media etc.)

## Business Analysts

... analyze social media buzz for the new services/offerings to gauge initial success and any course correction needed

## Finance

...analyze all Call Detail Records (CDRs) to identify and reduce revenue leakage due to unbilled / underbilled CDRs

## Marketing

... analyze subscriber usage pattern in real-time and combine that with the profile for delivering promotional or retention offers

## Business Development

... find and deliver new mechanisms to monetize network traffic and partner with upstream content providers






***Big Data enabled doctors from University of Ontario to apply neonatal infant monitoring to predict infection in ICU 24 hours in advance***



**IBM Data Baby  
youtube.com**





# Vestas optimizes capital investments based on **2.5 Petabytes** of information.

- Model the weather to optimize placement of turbines, maximizing power generation and longevity.
- Reduce time required to identify placement of turbine from weeks to hours.
- Incorporate 2.5 PB of structured and semi-structured information flows. Data volume expected to grow to 6 PB.

**Vestas**



# Cisco turns to IBM big data for intelligent infrastructure management

- Optimize building energy consumption with centralized monitoring
- Automate preventive and corrective maintenance

## Capabilities Utilized:

- Streaming Analytics
- Hadoop System
- Business Intelligence

## Applications:

- Log Analytics
- Energy Bill Forecasting
- Energy consumption optimization
- Detection of anomalous usage
- Presence-aware energy mgt.
- Policy enforcement







# Dublin City Centre Increases Bus Transportation Performance

## Capabilities Utilized:

### ***Stream Computing***

- Public transportation awareness solution improves on-time performance and provides real-time bus arrival info to riders
- Continuously analyzes bus location data to infer traffic conditions and predict arrivals
- Collects, processes, and visualizes location data of all bus vehicles
- Automatically generates transportation routes and stop locations

## Results:

- Monitoring 600 buses across 150 routes
- Analyzing 50 bus locations per second
- Anticipated to Increase bus ridership





# Asian telco reduces billing costs and improves customer satisfaction.

Capabilities:

Stream Computing  
Analytic Accelerators


Real-time mediation and analysis of  
**6B CDRs per day**

Data processing time reduced from  
**12 hrs to 1 sec**

**Hardware cost reduced to 1/8<sup>th</sup>**

Proactively address issues  
(e.g. dropped calls) impacting customer satisfaction.









## Data Set

- 1.1B tweets
- 5.7M blog and forum posts
- 3.5M relevant messages
- 97K referencing Product\_A
- 18K referencing Product\_B

## Information extracted

- Buzz and sentiment
- Gender, Location and Occupation
- Fans
- Intent to in purchase
- Specific attributes of products



# Enhancing Fraud Detection for Banks and Credit Card Companies

## Scenario

- Build up-to-date models from transactional to feed real-time risk-scoring systems for fraud detection

## Requirement

- Analyze volumes of data with response times that are not possible today
- Apply analytic models to individual client, not just client segment.



# Build Faster Real-time Trading Systems

## Scenario

- Identify and execute trades
- Process over 5M events per second with average latency of 150 microseconds

## Requirement

- Consuming, analyzing and acting on market data while maintaining sub-millisecond response time under extreme data loads
- Incorporate content feeds, news text, audio, video, to establish greater context for better decisions



# Transaction Analysis for Banking Industry

## Scenario

- Analyze transaction issues from federated systems and applications to provide up-to-date account status with less turnaround time

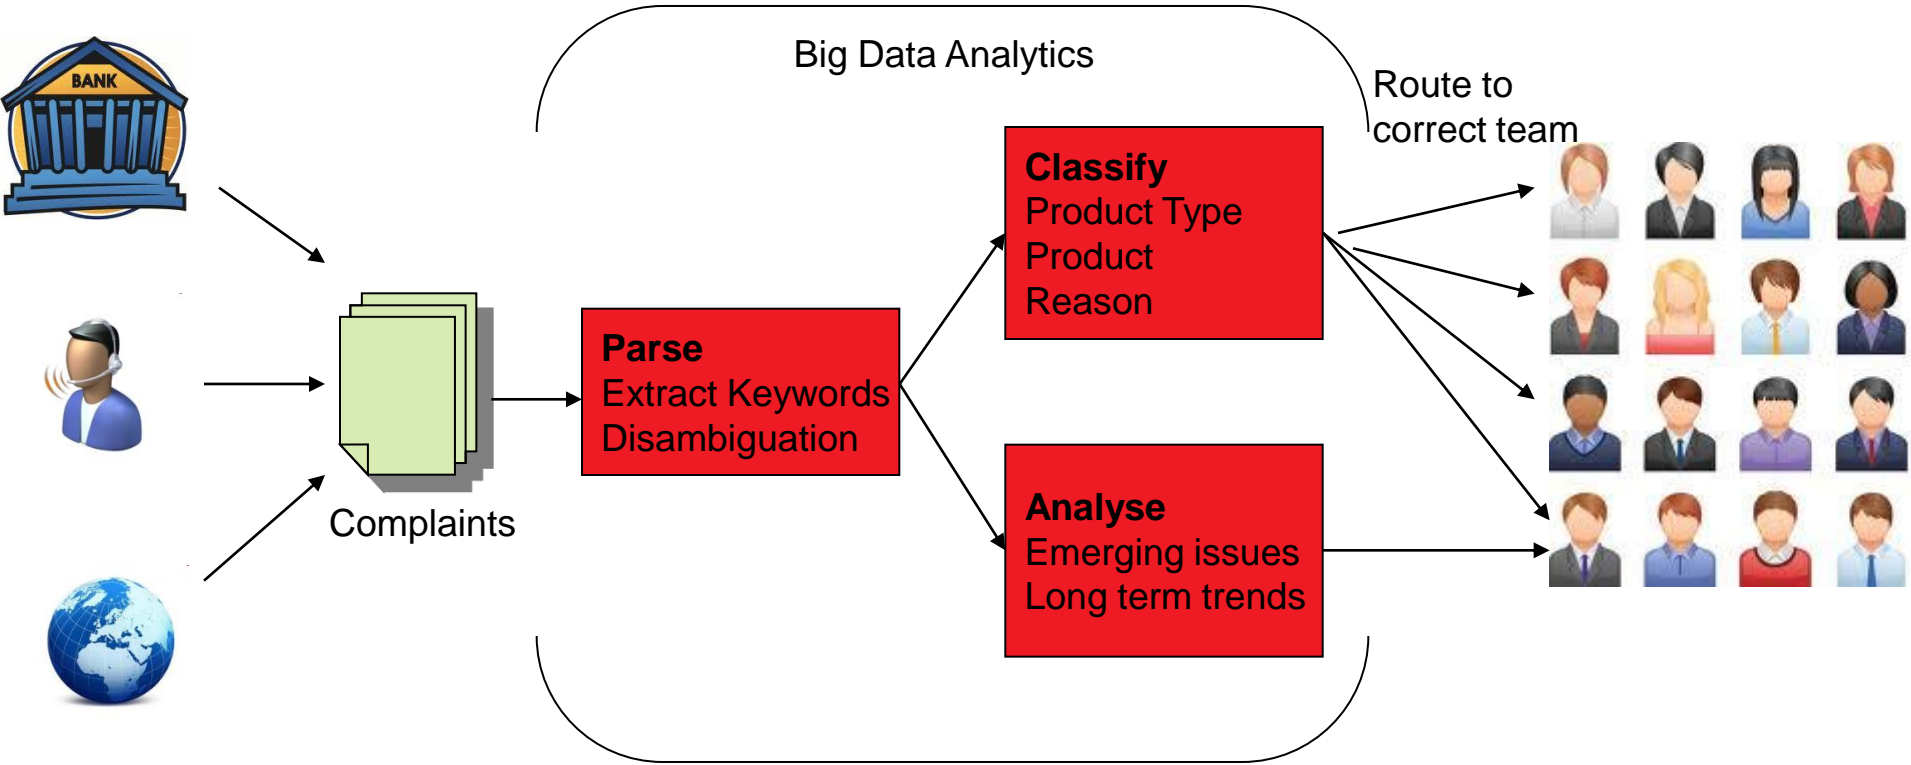
## Requirement

- Collect, aggregate, and analyze log data from various application systems
- Handle logs in different formats and correlating errors across applications
- Reduce response time to less than 2 minutes

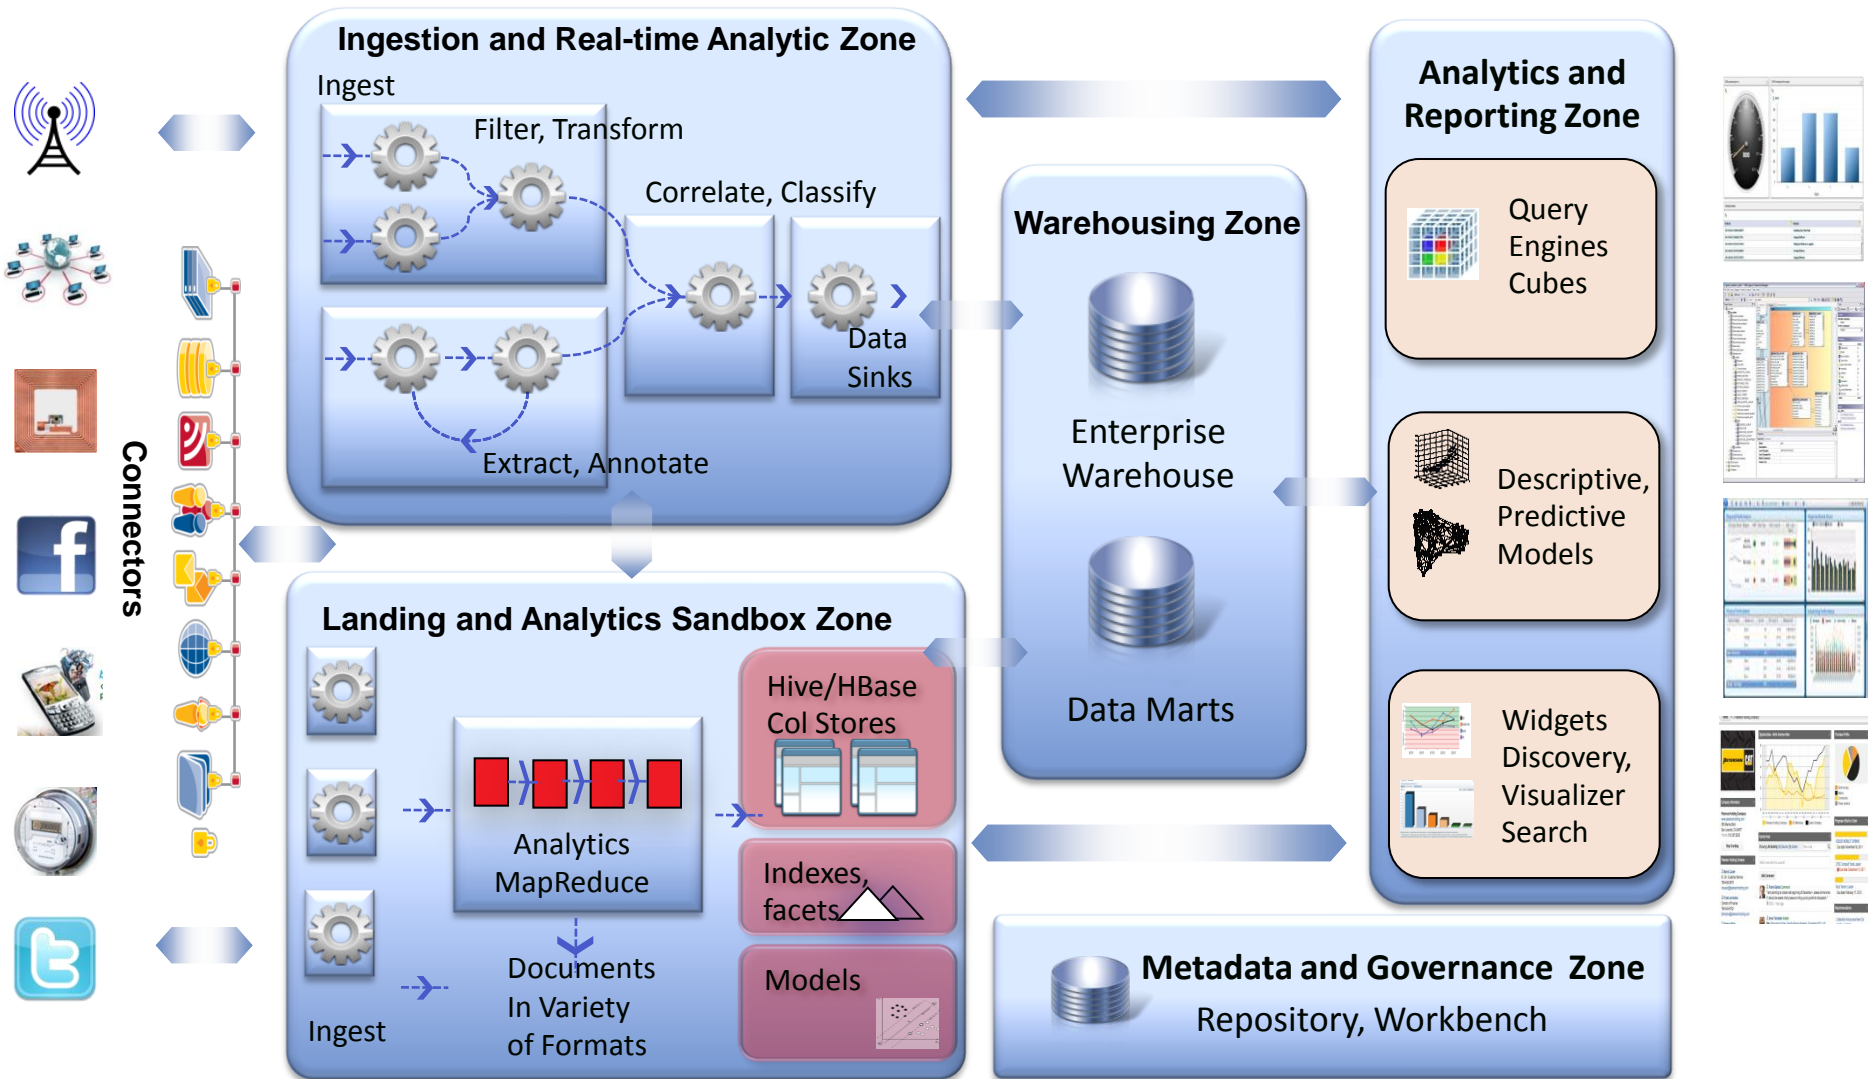




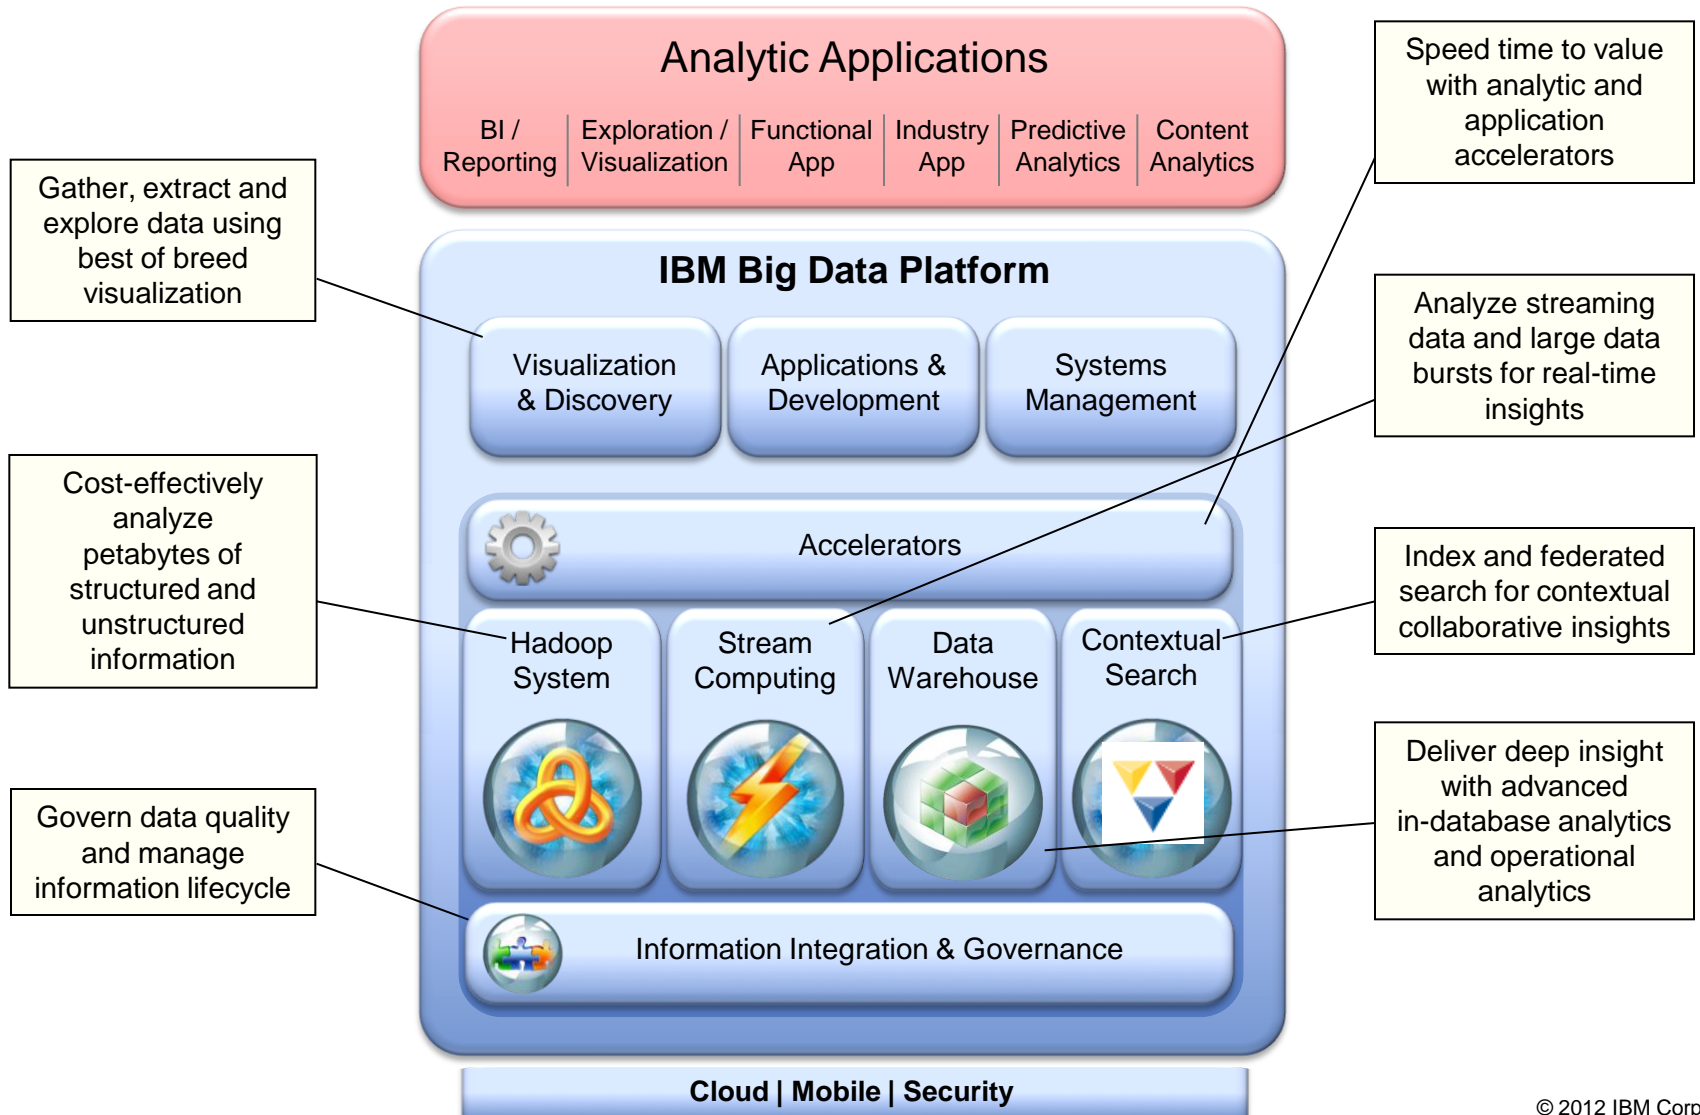
# CRM Automatic Complaints Classification



# Emerging Pattern of Big Data Implementation



# Big Data Platform and Application Framework





# Big Data Platform – Hadoop Distribution-Agnostic



## Platform Capabilities

- Built-in analytics
  - Text analytics engine, annotators, Eclipse tooling
  - Interface to project R (statistical platform)
- Deep integration with enterprise software stack
- Analytical tool for analysts
- Ready-made business process accelerators
- Integrated installation of supported open source and other components
- Web Console for admin and application access
- Platform enrichment: additional security, performance features, . . .
- World-class support
- Full open source compatibility

## Business benefits

- Quicker time-to-value due to IBM technology and support
- Reduced operational risk
- Enhanced business knowledge with flexible analytical platform
- Leverages and complements existing software

Open source components







# Massively Scalable Stream Analytics

## Linear Scalability

- Clustered deployments – unlimited scalability

## Automated Deployment

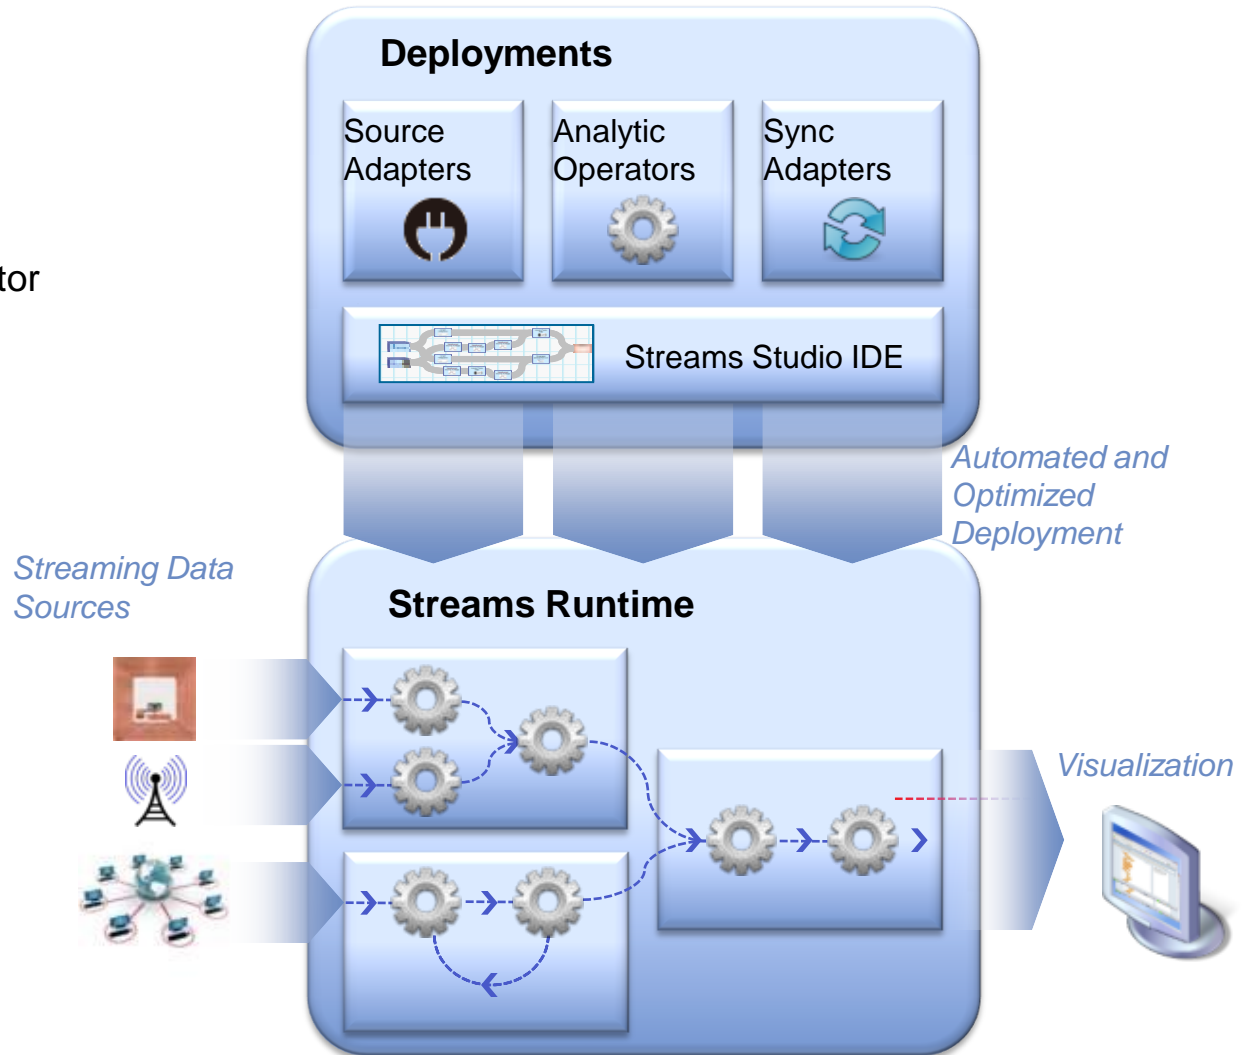
- Automatically optimize operator deployment across clusters

## Performance Optimization

- JVM Sharing – minimize memory use
- Fuse operators on same cluster
- Telco client – 25 Million messages per second

## Analytics on Streaming Data

- Analytic accelerators for a variety of data types
- Optimized for real-time performance



# Big Data Accelerators Improve Time to Value



## Telecommunications

CDR streaming analytics  
Deep Network Analytics



## Retail Customer Intelligence

Customer Behavior and Lifetime Value Analysis



## Finance

Streaming options trading  
Insurance and banking DW models




## Media and Entertainment

Sentiment Analytics, Intent to purchase, micro segmentation,



## Public transportation

Real-time monitoring and routing optimization



## Oil and Gas

Streaming statistical analysis, new energy source discovery




Over 100 sample applications



User Defined Toolkits

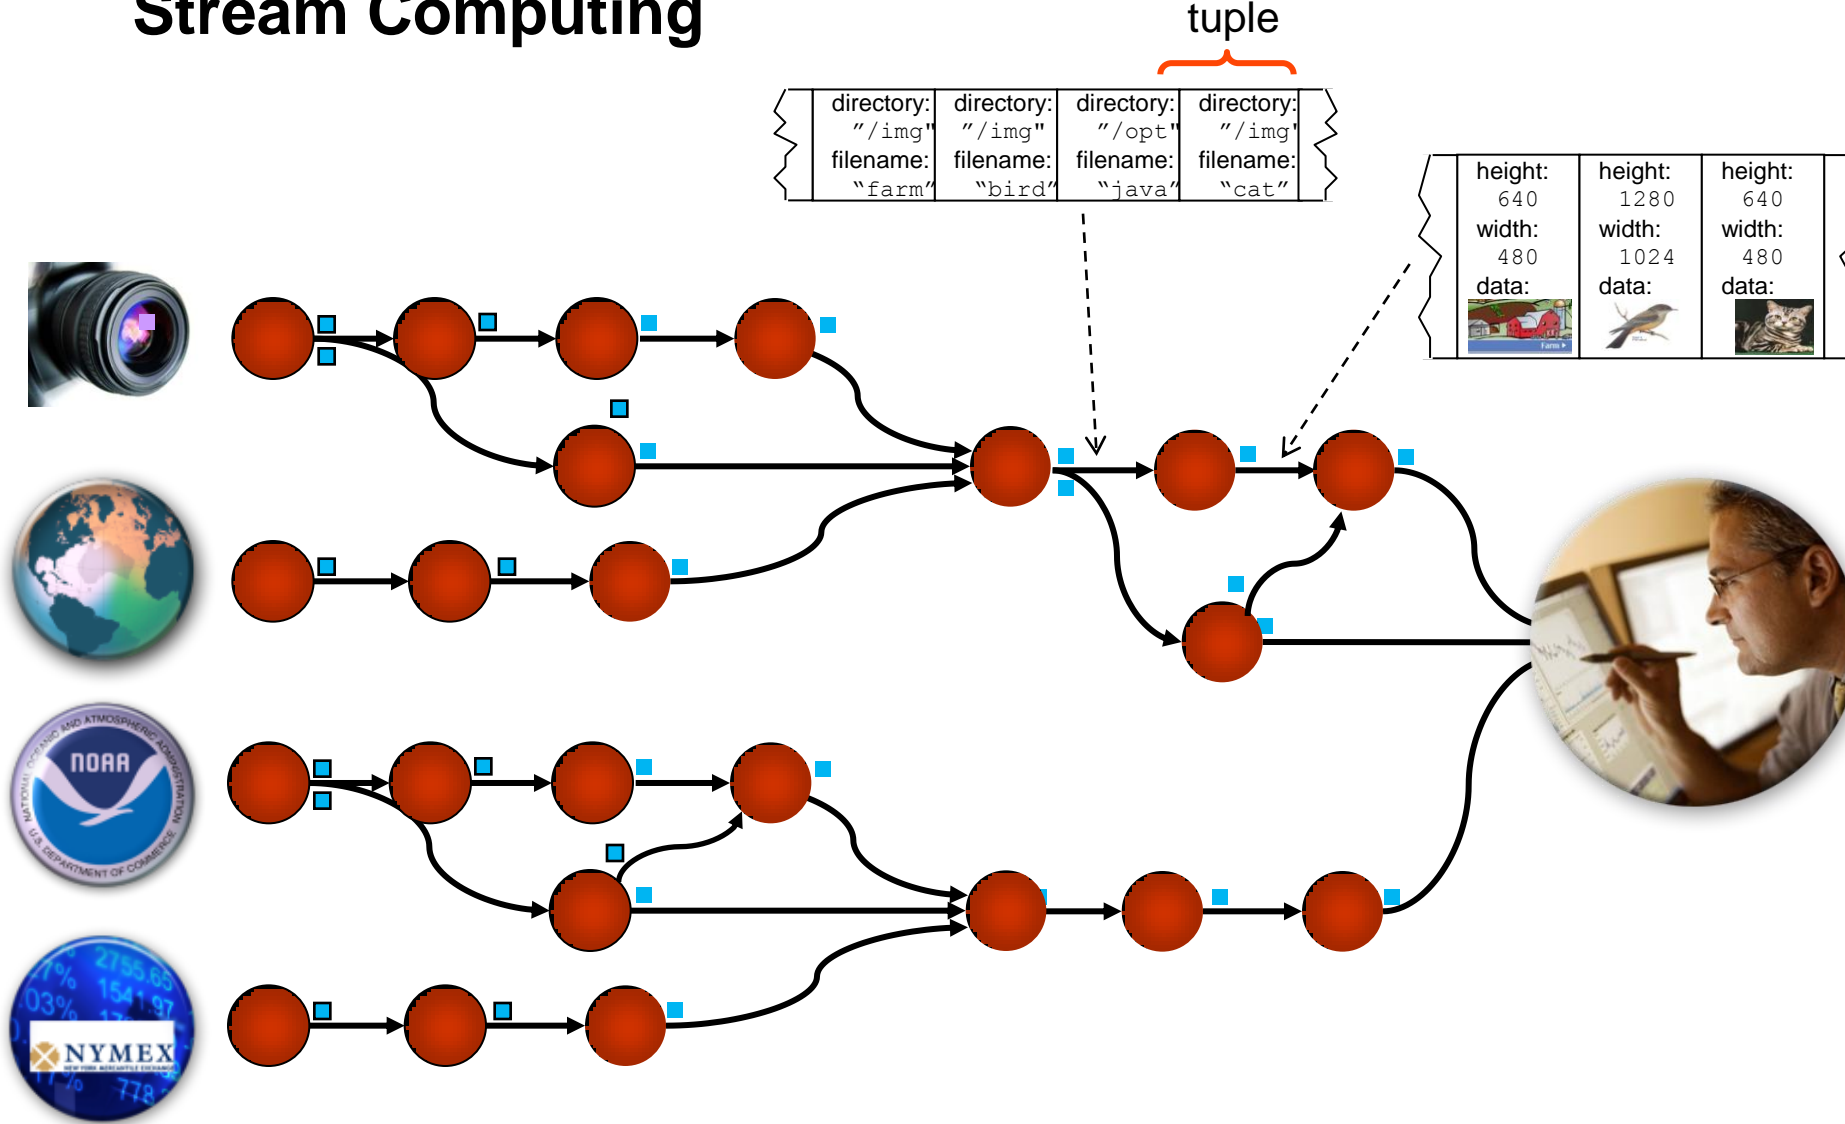


Standard Toolkits



Industry Data Models  
*Banking, Insurance, Telco, Healthcare, Retail...*


# Stream Computing



# Real-time Stream Analytics

→ continuous ingestion  
 → continuous analysis

infrastructure provides services for  
 scheduling analytics across h/w nodes  
 establishing streaming connectivity



achieve scale

by partitioning applications into components

by distributing across stream-connected hardware nodes

where appropriate,  
 elements can be “fused” together  
 for lower communication latencies





## IBM InfoSphere BigInsights

### Visualization & Discovery

BigSheets

Dashboard & Visualization

### Applications & Development

Apps

Text Analytics

MapReduce

Workflow

Pig & Jaql

BigSQL & Hive

### Systems Management

Admin Console

Monitoring

### Integration

JDBC

Netezza

DB2

Streams

DataStage

Flume

Sqoop

### Advanced Analytic Engines

Machine Learning (ML)

Text Processing Engine & Extractor Library (AQL+HIL)

BigSQL

R

Mahout

### Workload Optimization

Integrated Installer

Enhanced Security

Splittable Text Compression

Adaptive MapReduce

ZooKeeper

Oozie

Jaql

Flexible Scheduler

Lucene

Pig

Hive

Index

### Runtime

MapReduce

### Data Store

HBase

Column Store

### File System

HDFS

GPFS (EAP)

### Management

Security

Audit & History

Lineage

HCatalog

# New classes of applications for end-users



Real Time Analytics

Internet Scale Analytics

In-Database Analytics

Enterprise Data Connectors

Federated Discovery

Navigation and Visualization

Application Framework

**OPTIMUM INVESTMENTS**

Logged in as Frank Gelato | Help

**Followed High Net Wealth Clients**

- Isabella Jones
- Thomas Jackson
- Theresa Mayer
- Michael Kleinfelder

**Tracked Products**

- 529 Plan
- 401K
- Money Market IRA
- Fixed Income & Bonds

**Financial Blogs**

CNBC: Warren Buffett: 'Disruptive' Debt Limit Debates Are 'Waste of...  
Everything Warren Buffet

CNBC: CNBC Transcript: Warren Buffett on Russian Roulette, Tax...  
Everything Warren Buffet

BLOOMBERG: Munger Treats 'Hard-Core Addicts' as Wesco Stock Exits...  
Everything Warren Buffet

**Investment News**

Texas Gains New Billion Dollar Bank Business Wire - 35 minutes ago

Ally Financial Reports Preliminary Second Quarter 2011 Financial Results  
PR Newswire - 54 minutes ago

Pinnacle West Reports Second-Quarter Results  
TheStreet.com - 55 minutes ago

**Action Needed**

Sentiment	Customer	Type	Format	Time	Product
Negative	Isabella Jones	Support	Tweet	05.29.12 10:30 am EST	Mutual funds
Neutral	Thomas Jackson	Support	Support ticket - email	05.29.12 10:25 am EST	401K
Positive	Isabella Jones	Sales	Tweet	05.29.12 10:05 am EST	529
Negative	Theresa Mayer	Satisfaction	Blog	05.29.12 09:30 am EST	401K

**Activity Feed**

Isabella Jones - @tzyJones  
Wow is the stock market really this bad, my monthly mutual fund statement looks terrible! Time to call my financial advisor.  
Twitter - Minutes ago

Theresa Mayer - tmayer@gmail.com  
Great advice on retirement planning! Thinking about increasing my 401K contribution but are there other retirement plans I should be looking at given I hope to retire in 10-15 yrs?  
<http://mextavenue.org/blog/why-women-need-embrace-retirement-planning>  
Blog - Minutes ago

Thomas Jackson  
How can I change my 401K contribution using your online system?  
Remedy - Minutes ago

**Sentiment**

**Sentiment By Age**

# IBM's Big Data Business Partner Ecosystem

**230+**  
Big Data Business Partners

**KARMAPHERE**  
**Datameer** Powerfully Simple™  
**DIGITAL REASONING SYSTEMS**  
**Kitenga** Reinvigorating Information  
**MATTIVO** INSIGHT THAT MATTERS  
**JASPERSOFT**  
**StackIQ**  
**CLICKFOX**  
**REVOLUTION ANALYTICS**  
**QUEST SOFTWARE**  
**DATASIFT**  
**HADAPT**

**boardreader™**  
Connecting Communities through Search  
**CISCO**  
**EasyAsk.** Search That Means Business  
**lucid** IMAGINATION  
**infogroup**  
**tresata** make • data • work  
**TerraEchos**  
**syncsort**  
**ALOGIX** SYSTEMS GROUP  
**FUZZY LOGIX**  
**Platform Computing**

**Atigeo™**  
Compassionate Technology  
**CADOW SOFTWARE**  
**GOGRID**  
**TOSHIBA**  
Leading Innovation >>>  
**Sharpe Engineering Inc.**  
**BROCADE**  
**eZly**  
**SEMA sistemas**  
**ARISTA**  
**Zettaset**

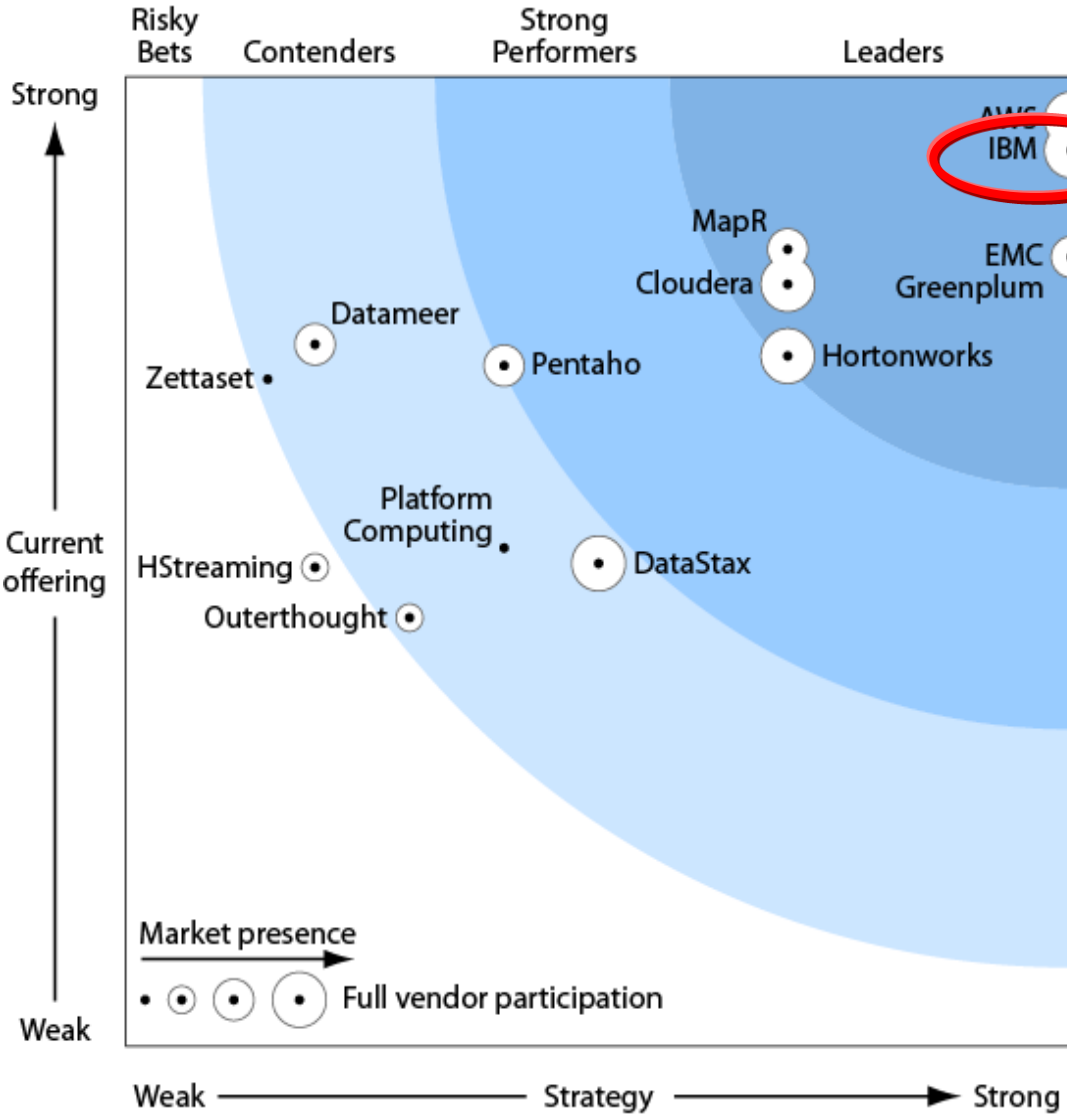
**Capgemini** CONSULTING. TECHNOLOGY. OUTSOURCING.  
**Infosys**  
**lexa** Make your data  
**DEL**  
**EIGHTHOUSE** COMPUTER SERVICES  
**Perficient**  
**ESTUATE**  
**btrg**  
**ONE POINT SOLUTIONS**  
**TAMGROUP**  
**WaveStrong**

**SYSTEMS** Refresh Solutions, Inc.  
**Cognizant**  
**WIPRO** Applying Thought  
**MEPS** REAL TIME, INC.  
**NP** NewVantage Partners



February 2012 “The Forrester Wave™: Enterprise Hadoop Solutions, Q1 2012”

## Forrester Wave™: Enterprise Hadoop Solutions, Q1 '12



**“IBM has the deepest Hadoop platform and application portfolio. IBM, an established EDW vendor, has its own Hadoop distribution; an extensive professional services force working on Hadoop projects; extensive R&D programs developing Hadoop technologies; connections to Hadoop from its EDW.”**

–The Forrester Wave™: Enterprise Hadoop Solutions, 1Q12



THINK

BIG

BIG