


Wolfgang Nimfuehr– Dipl. Ing.

Big Data Enterprise Class

Information Agenda Executive Consultant

Big Data Tiger Team

IBM Software Group Europe



The Information Explosion in Data and Real World Events

44x


as much Data and Content
Over Coming Decade

2009
800,000 petabytes



Velocity
Variety
Volume

2020
35 zettabytes



80%

Of world's data
is unstructured



SI decimal prefixes		Binary usage
Name (Symbol)	Value	
kilobyte (kB)	10 ³	2 ¹⁰
megabyte (MB)	10 ⁶	2 ²⁰
gigabyte (GB)	10 ⁹	2 ³⁰
terabyte (TB)	10 ¹²	2 ⁴⁰
petabyte (PB)	10 ¹⁵	2 ⁵⁰
exabyte (EB)	10 ¹⁸	2 ⁶⁰
zettabyte (ZB)	10 ²¹	2 ⁷⁰
yottabyte (YB)	10 ²⁴	2 ⁸⁰

1 in 3 Business leaders frequently make decisions based on information they don't trust, or don't have

1 in 2 Business leaders say they don't have access to the information they need to do their jobs

83% of CIOs cited "Business intelligence and analytics" as part of their visionary plans to enhance competitiveness

60% of CEOs need to do a better job capturing and understanding information rapidly in order to make swift business decisions

Organizations Need Deeper Insights

Challenge

Study a Large Volume and Variety of Data to Find New Insights



-  *Multi-channel customer sentiment and experience analysis*
-  *Support medical diagnostics
Detect life-threatening conditions*
-  *Predict weather patterns to plan optimal wind turbine usage, and optimize capital expenditure on asset placement*
-  *Make risk decisions and frauds detection based on real-time transactional data*
-  *Identify criminals and threats from disparate video, audio, and data feeds*

Leveraging Big Data Analytics

How do you address the challenges presented by empowered market participants generating mountains of data?

People and devices are talking more than ever ...



1 Can you capture data generated by these interactions?

2 Can you do it in real-time?

3 Can you turn that data into insights to predict customer / competitive / market behavior?


Source:

¹ – Barrera, Clod and Wojtowecz. "Cloud Leads Five Storage Trends for 2011." CIO. January 27, 2011 .

² – <http://www.internetworldstats.com/stats.htm>.

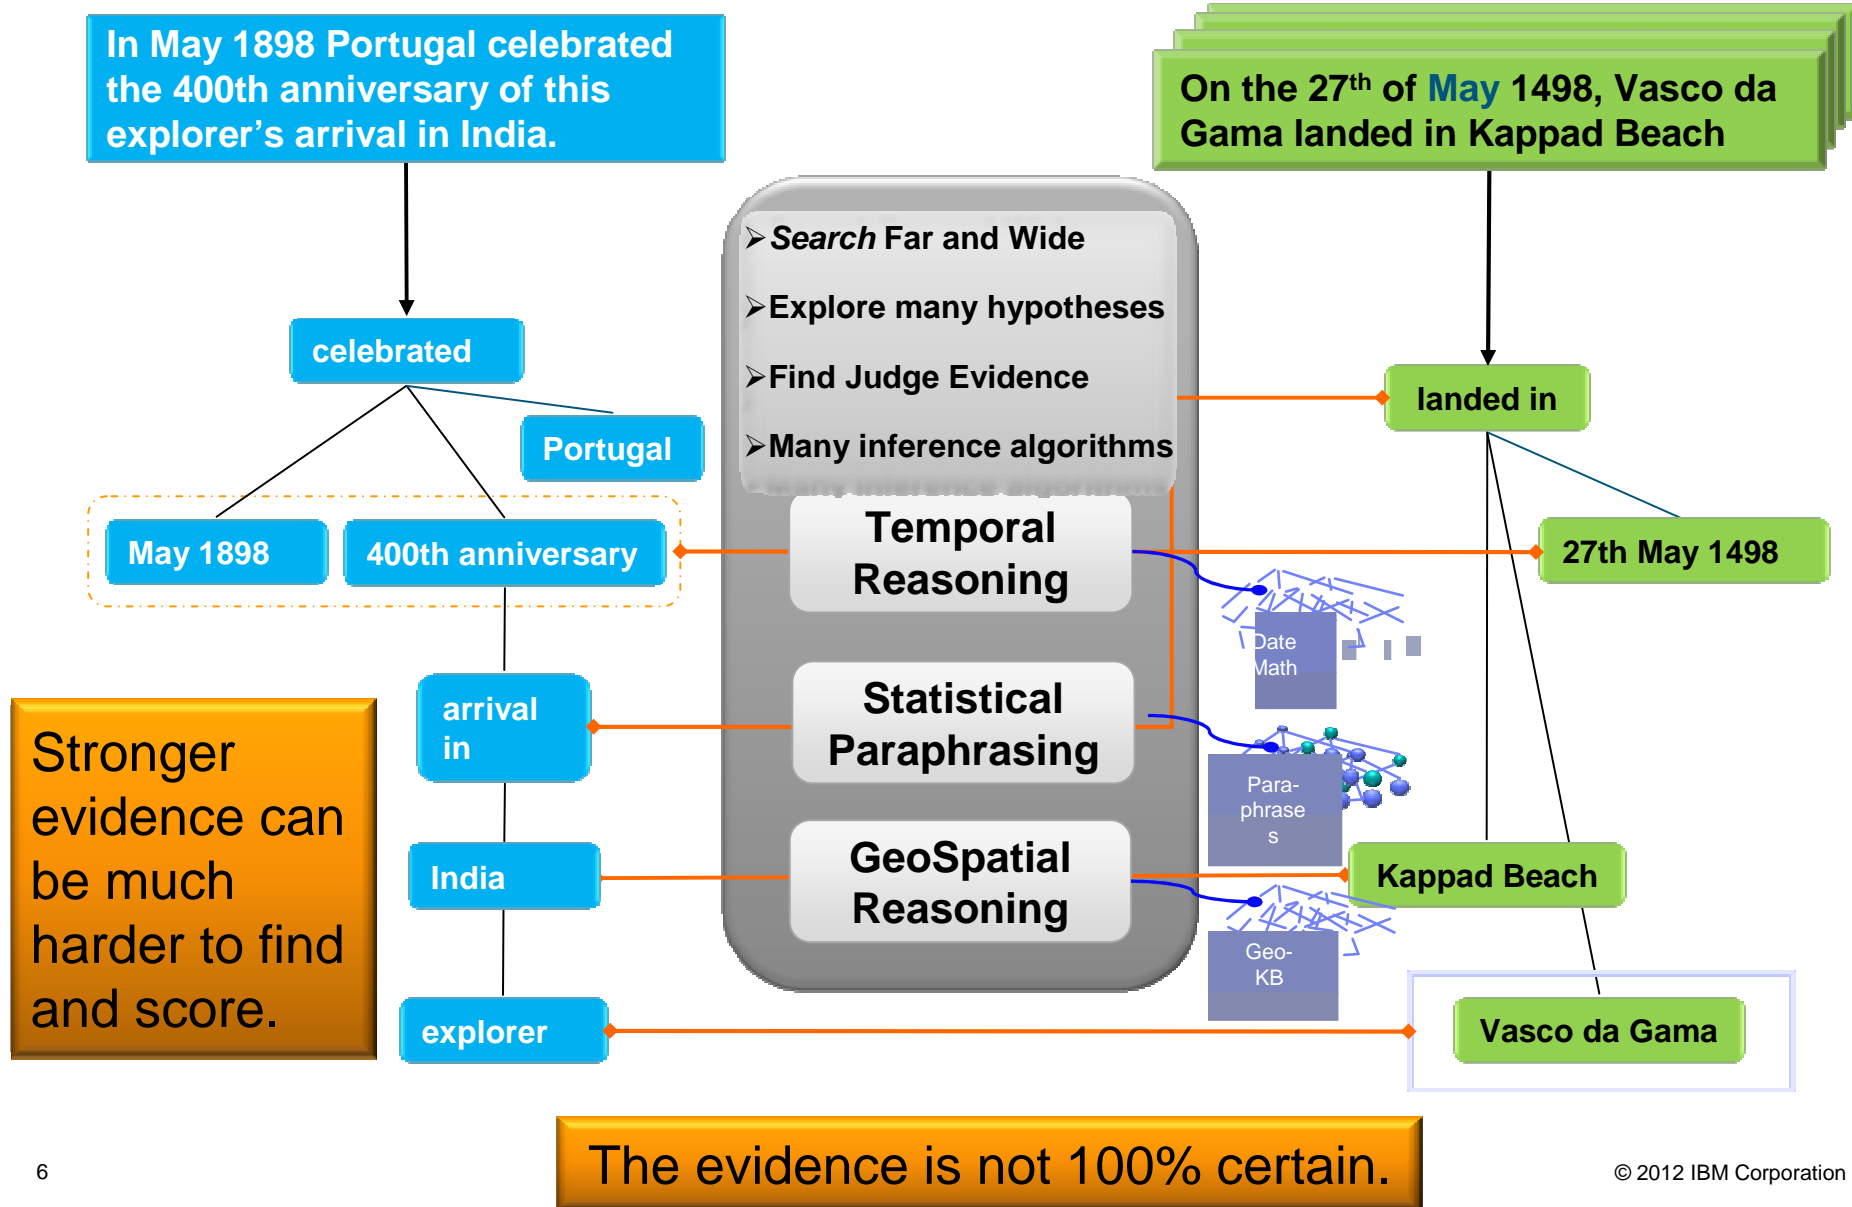
³ – <http://www.abiresearch.com/press/3584-More+than+Seven+Trillion+SMS+Messages+Will+Be+Sent+in+2011>

On Feb 16 2011 the IBM Watson system won Jeopardy!

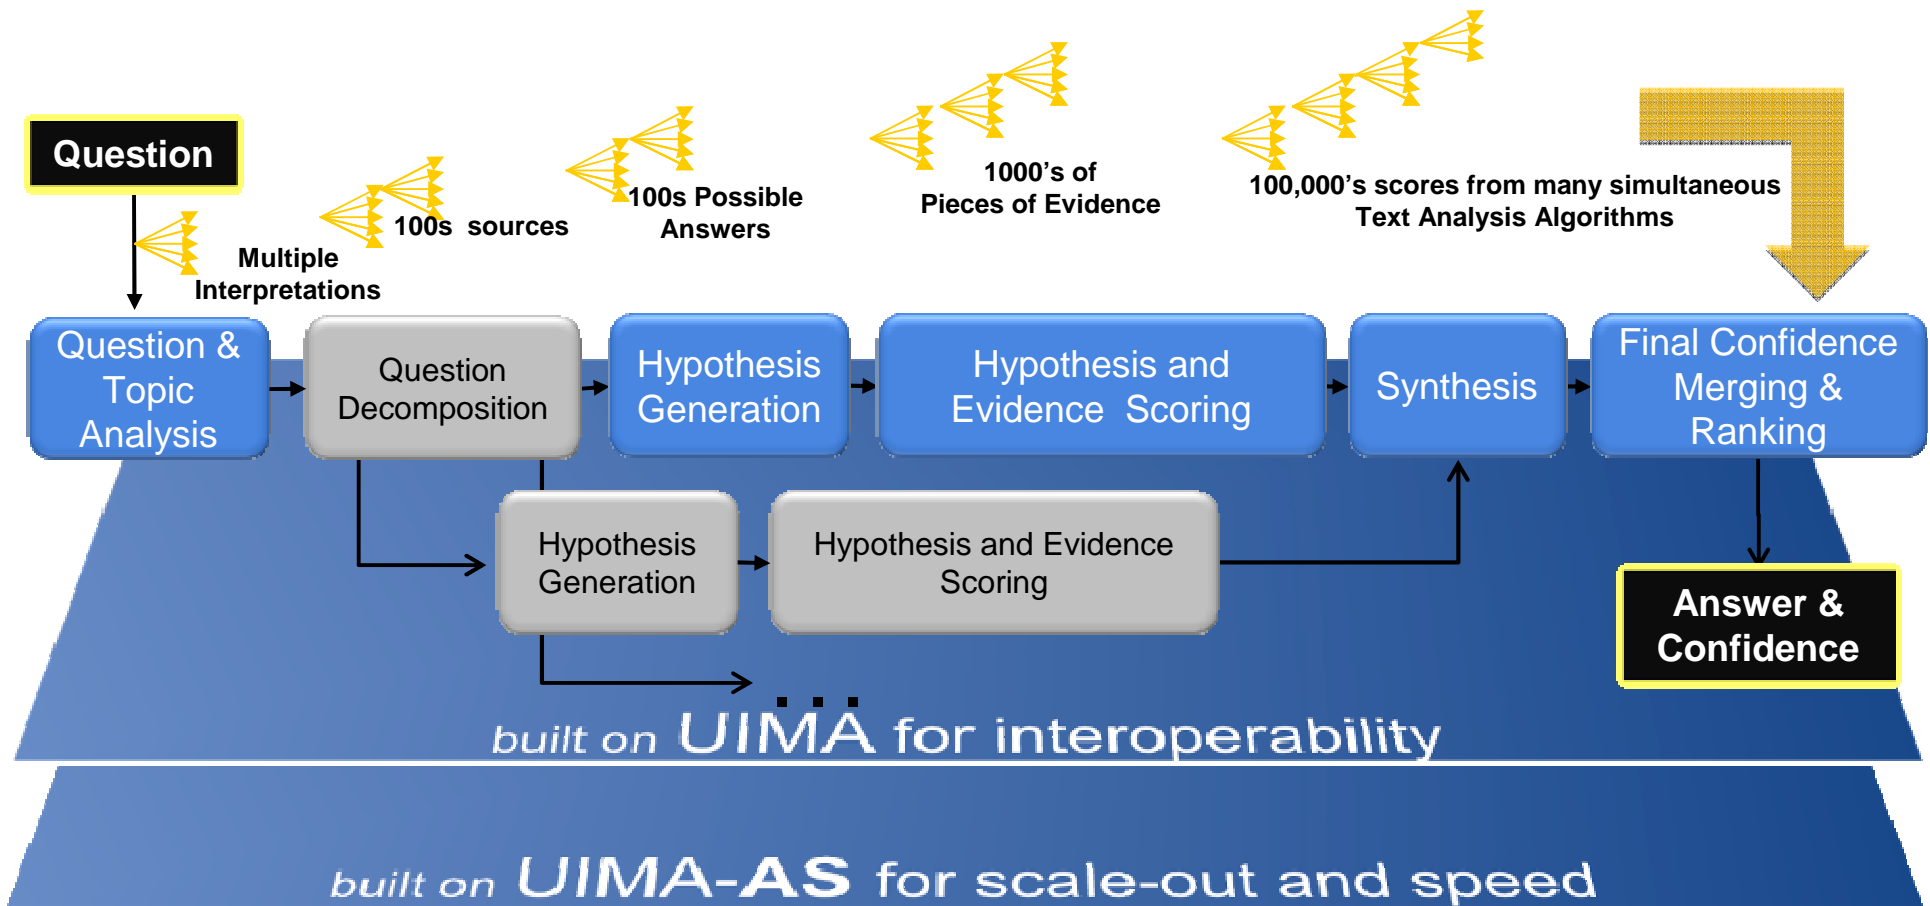


Can we design a computing system that rivals a human's ability to answer questions posed in natural language, interpreting meaning and context and retrieving, analyzing and understanding vast amounts of information in real-time?

Different Types of Evidence: Deeper Evidence



Deep QA Massively Parallel Probabilistic Evidence-Based Architecture



Big Data use cases across all industries

Financial Services

- Fraud detection
- Risk management
- 360° View of the Customer



Utilities

- Weather impact analysis on power generation
- Transmission monitoring
- Smart grid management

Transportation

- Weather and traffic impact on logistics and fuel consumption



IT

- Transition log analysis for multiple transactional systems
- Cybersecurity

Health & Life Sciences

- Epidemic early warning system
- ICU monitoring
- Remote healthcare monitoring



Retail

- 360° View of the Customer
- Click-stream analysis
- Real-time promotions

Telecommunications

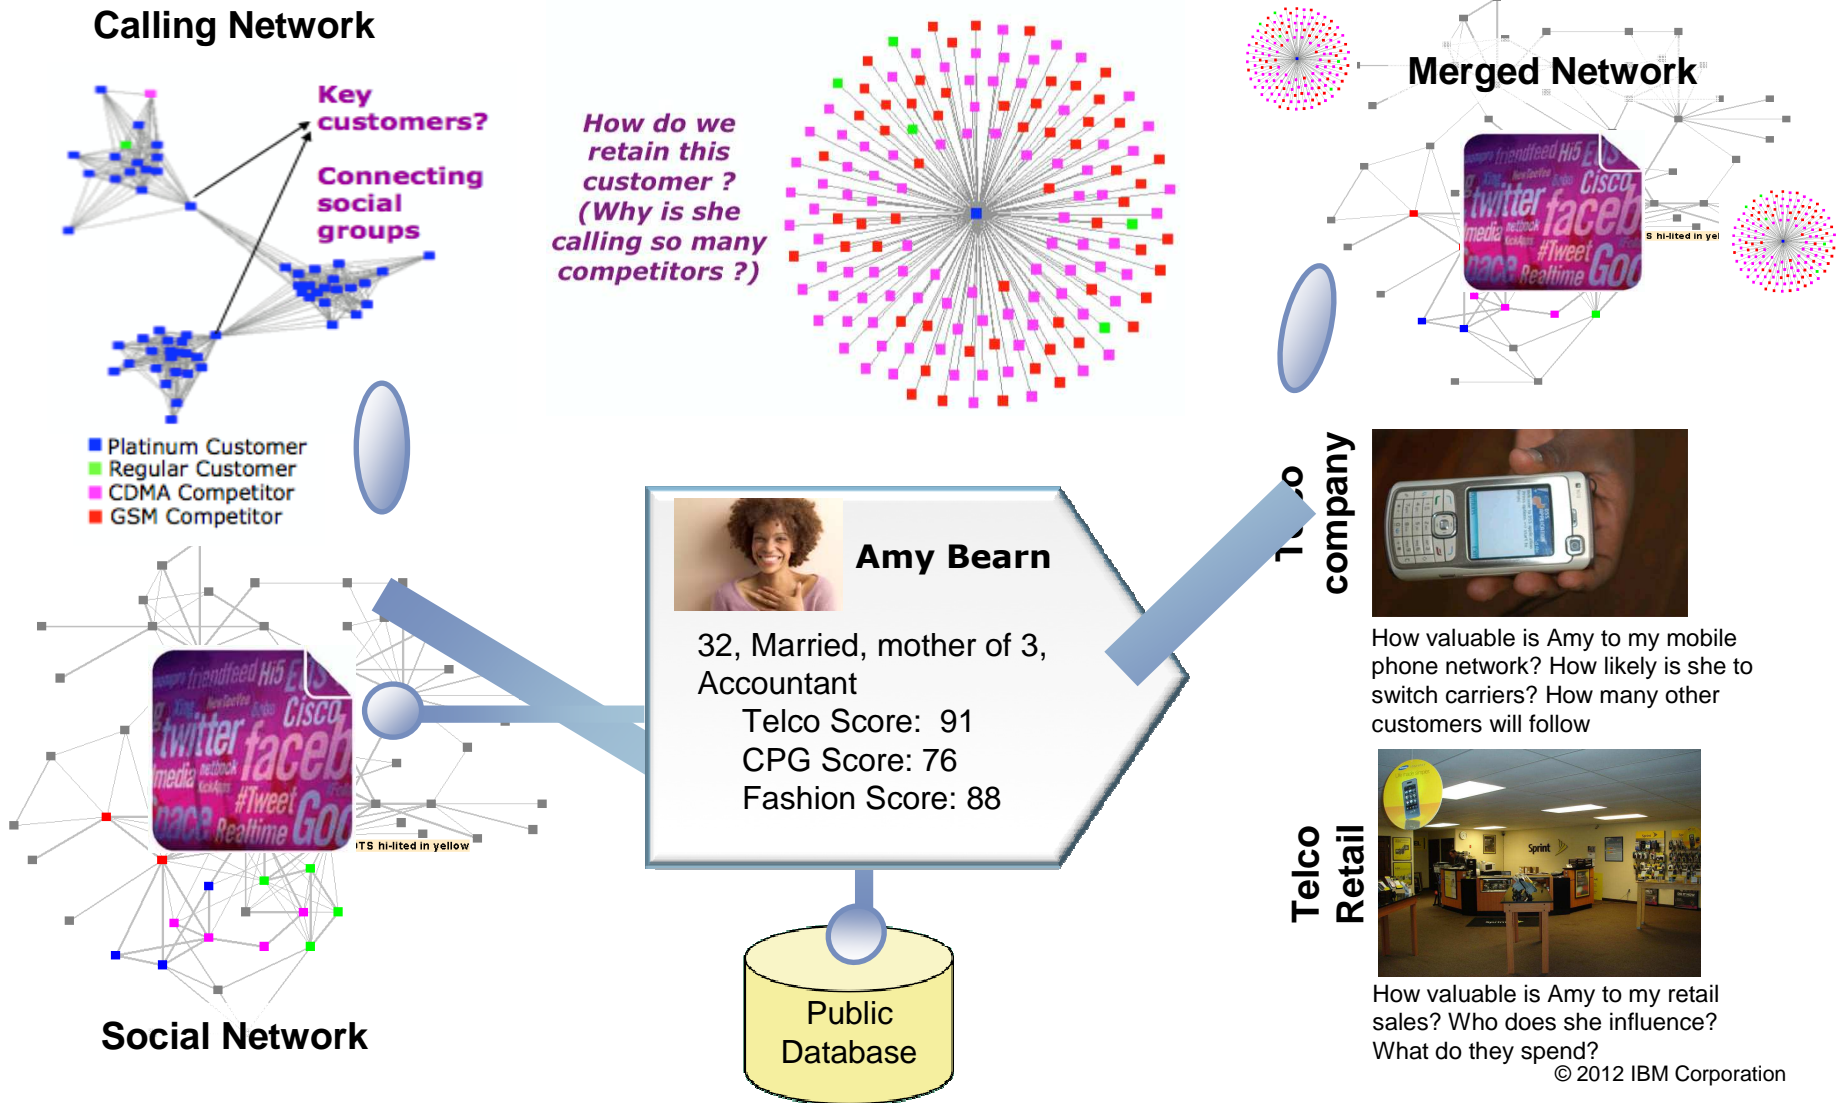
- CDR processing
- Churn prediction
- Geomapping / marketing
- Network monitoring



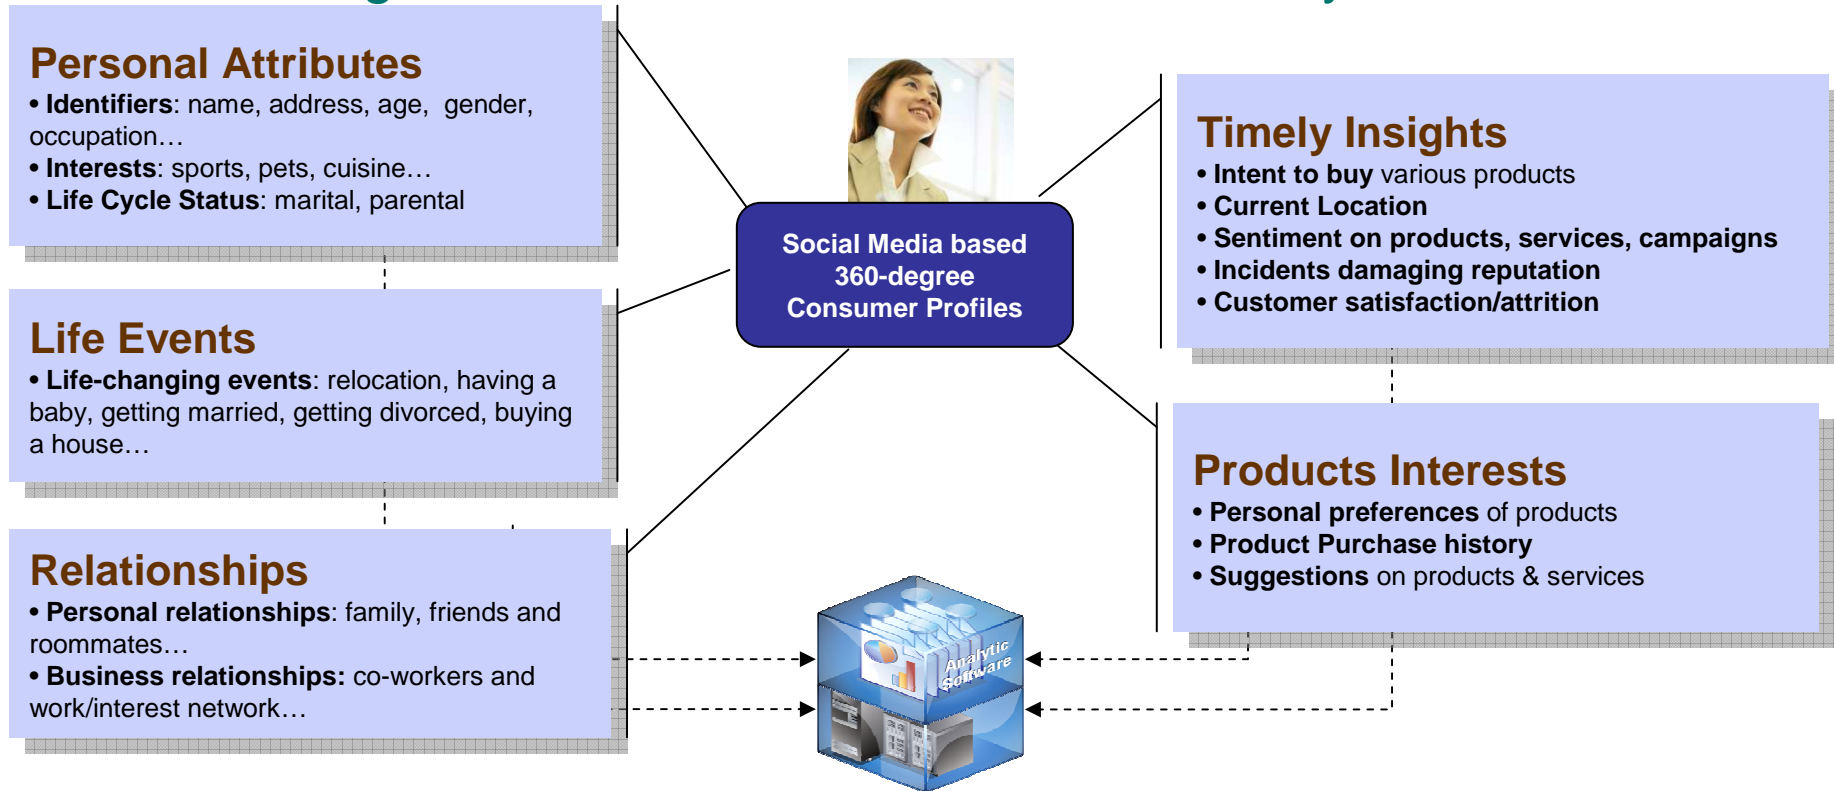
Law Enforcement

- Real-time multimodal surveillance
- Situational awareness
- Cyber security detection

Big Data Pattern Monetizing Relationships - not just Transactions



Use Case: Big Data 360° Profile / Sentiment Analysis



Monetizable intent to buy products

- I need a new digital camera for my food pictures, any recommendations around 300?
- What should I buy?? A mini laptop with Windows 7 OR a Apple MacBook!??!

Location announcements

I'm at Starbucks Parque Tezontle <http://4sq.com/fYReSj>

Life Events

- College: Off to Stanford for my MBA! Bbye chicago!
- Looks like we'll be moving to New Orleans sooner than I thought.

Intent to buy a house

I'm thinking about buying a home in Buckingham Estates per a recommendation. Anyone have advice on that area? #atx #austinrealestate #austin

Use Case: Institutional Risk Application

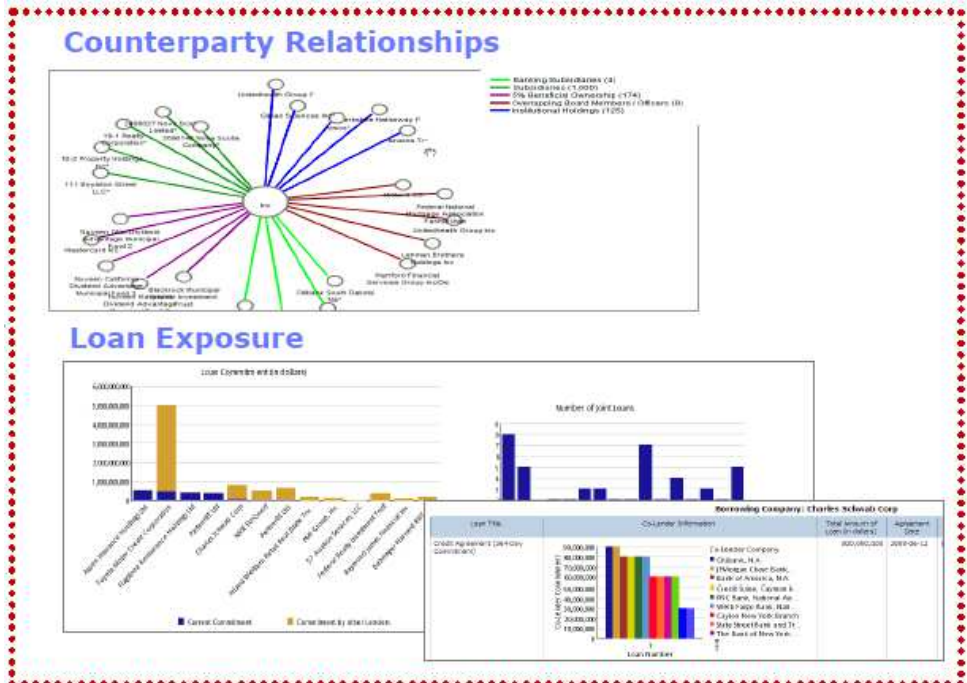
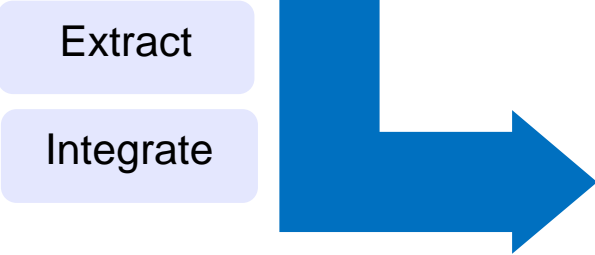
Comprehensive view of publicly traded companies and related people based on regulatory filings

Annual Report

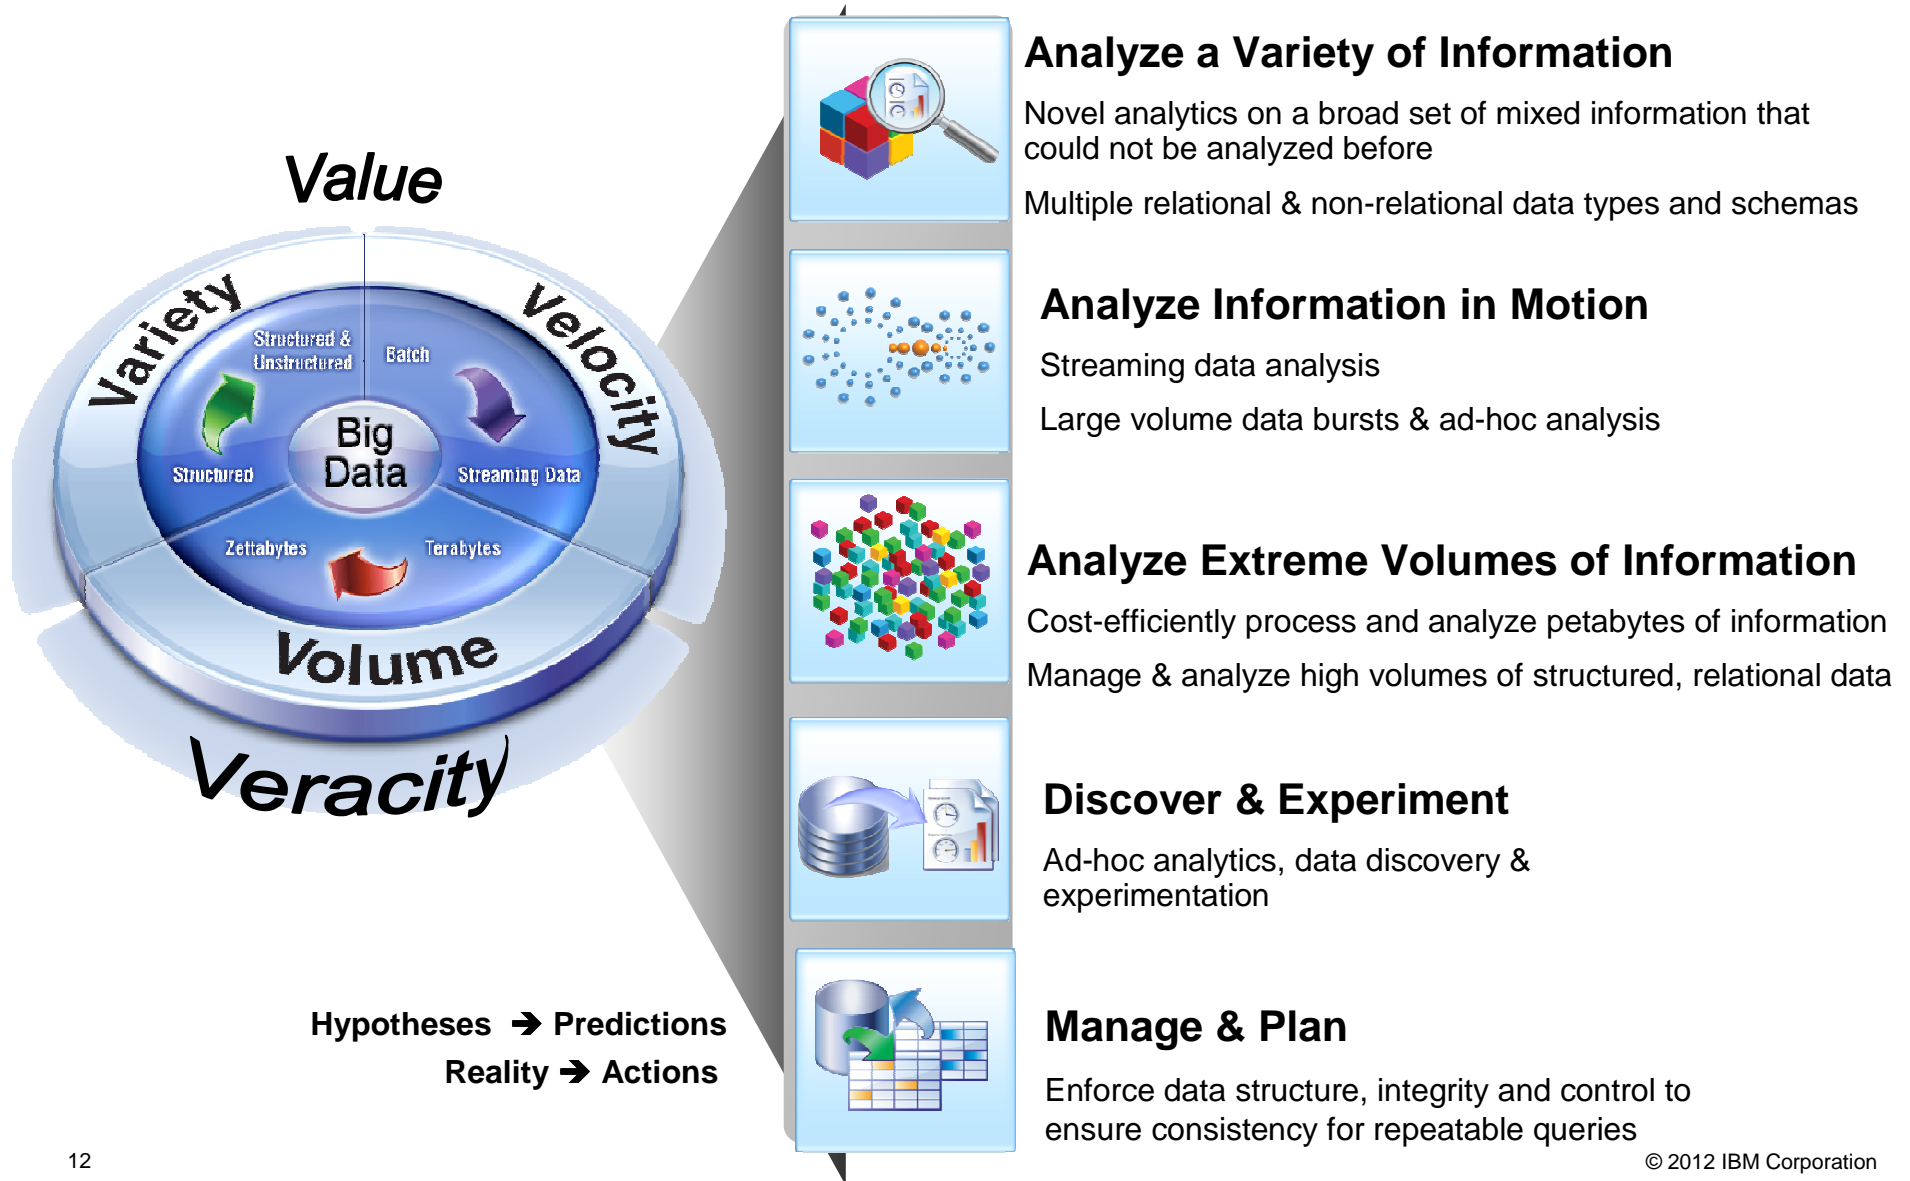
Loan Agreement

Proxy Statement

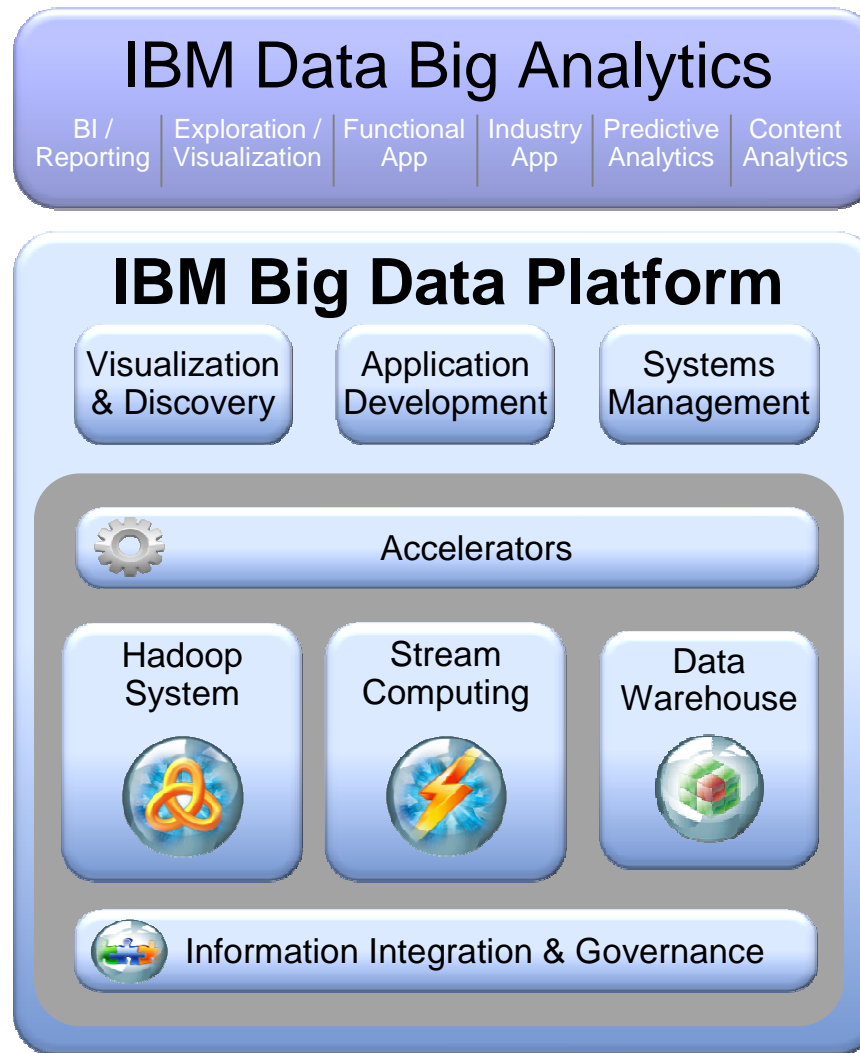
Insider Transaction



Requirements for a Big Data Solution Platform



Big Data Analytics harnesses the value of the new mix of information



IBM Vivissimo Velocity

Unique Federated Discovery and Navigation Technology

Advanced On-the-Fly Analytics

- State-of-the-art real-time text and meta-data analytics

Accuracy

- More relevant results due to position-based indexing

Security

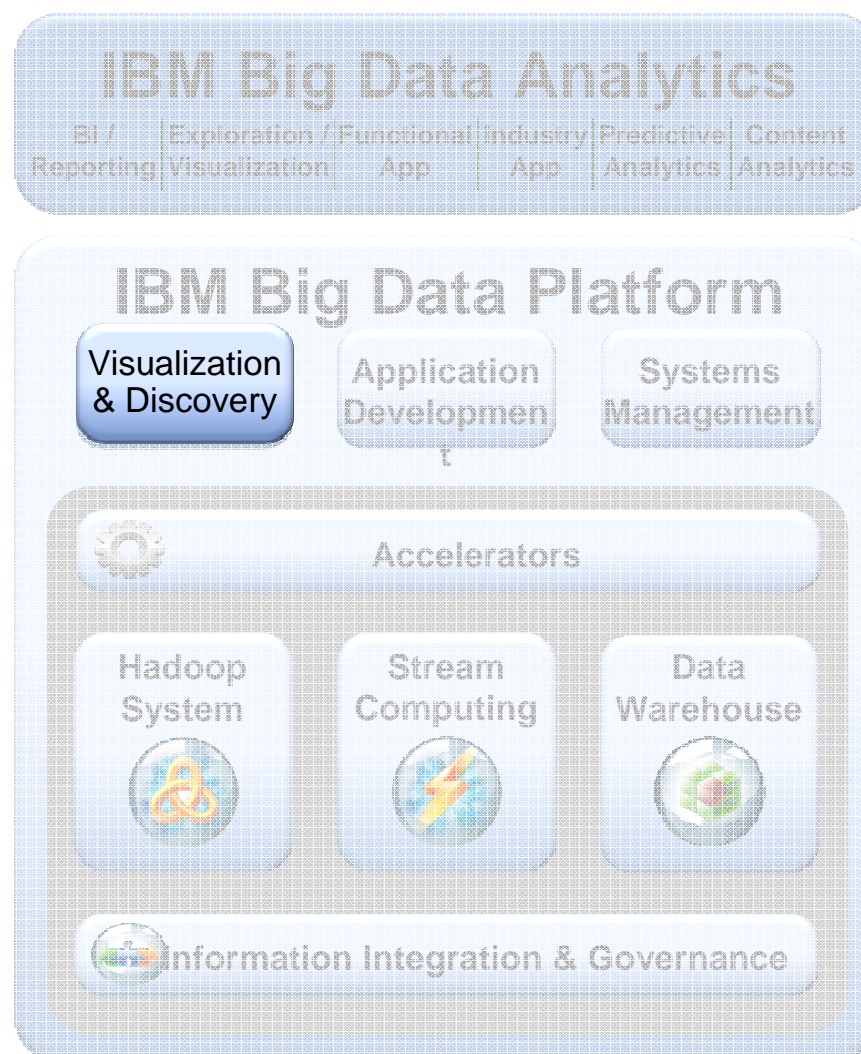
- Respects the security rights of underlying systems

Scalability

- Scales to trillions of records

Integration

- Rich connectors to a variety of data sources



Vivissimo Velocity Navigating Enterprise Class Information across different sources

The screenshot displays the Vivissimo Velocity search interface. At the top, there are navigation tabs for various sources: All, Lotus Notes, SharePoint, Documentum, Employees, Trade Publications, and Desktop. A search bar contains the query 'collaboration'. The interface is annotated with several callouts:

- Dynamic categorization:** Points to the 'Topic Clusters' sidebar on the left, which lists categories like 'Portals, Content & Collaboration (16)', 'Marketing (28)', and 'Social (16)'. Below this is a 'Refinements' sidebar with filters for 'Mobile (55)', 'Windows (51)', and 'Review (32)'. At the bottom left, a 'Folder' sidebar lists 'Mobile (8)', 'Sales Education (8)', and 'Analyst (5)'.
- Access across many sources:** Points to the top navigation tabs.
- Leveraging structured and unstructured content:** Points to the 'Employees' profile for Stacy Monarke, showing her photo, name, department (Pittsburgh, PA), title (Director of Product Management), extension (971), and email (S.Monarke@vivissimo.com). It also highlights 'Top Tags' such as 'collaboration (10)', 'competitors (7)', and 'social tagging (4)'.
- Highly relevant, personalized results:** Points to the search results, including an email titled 'Confirmation Package for the Gartner Portals, Content & Collaboration Summit' and a document titled 'Portals, Collaboration & Content Conference brochure deadline'.
- Enhanced by social collaboration:** Points to the 'My Tags' section, which includes 'Tradeshaw', 'Gartner', and 'PCC', and the 'In Marketing' section with tags like 'Analyst' and 'Mobile'.
- Refinements based on structured information:** Points to the 'Graphical Refinements' on the right, which includes a bar chart for 'Date' (Selected: 6/1/2005 - 2/25/2009) and another for 'Average Rating' (Selected: 1 - 4.5). Below this is a 'Folders' sidebar with a tree view of folders like 'Vivissimo', 'All Results', 'BI', 'Competitor', and 'Sales Education'.

InfoSphere BigInsights

Analytical platform for Big Data at-rest

Based on open source & IBM technologies

Distinguishing characteristics

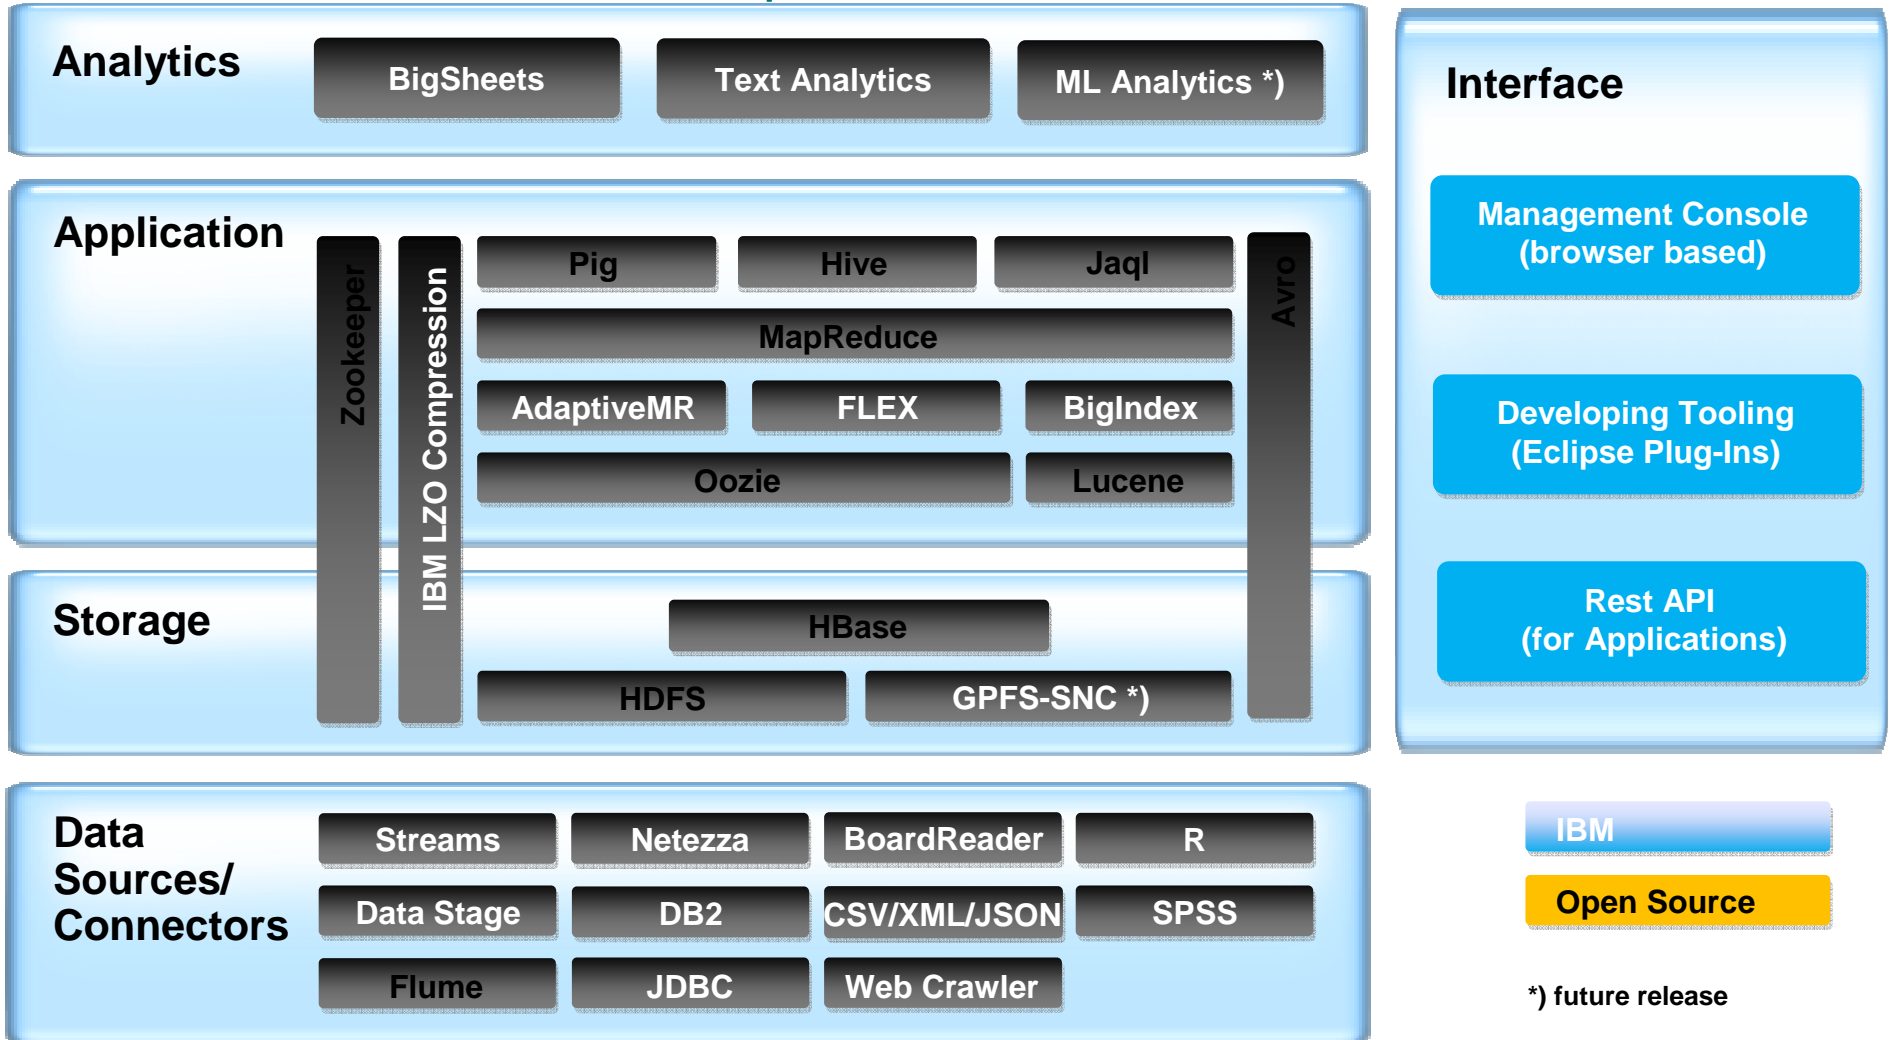
- Built-in analytics enhances business knowledge
- Enterprise software integration complements and extends existing capabilities
- Production-ready platform with tooling for analysts, developers, and administrators speeds time-to-value and simplifies development/maintenance

IBM advantage

- Combination of software, hardware, services and advanced research



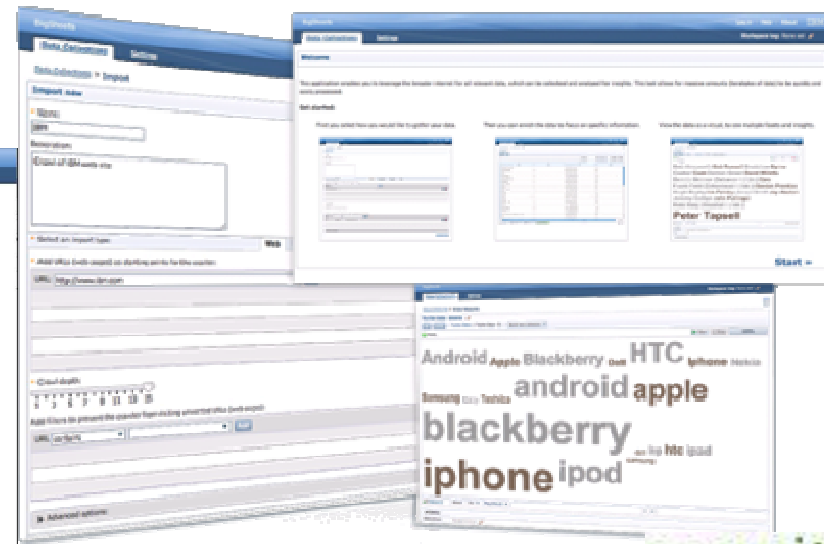
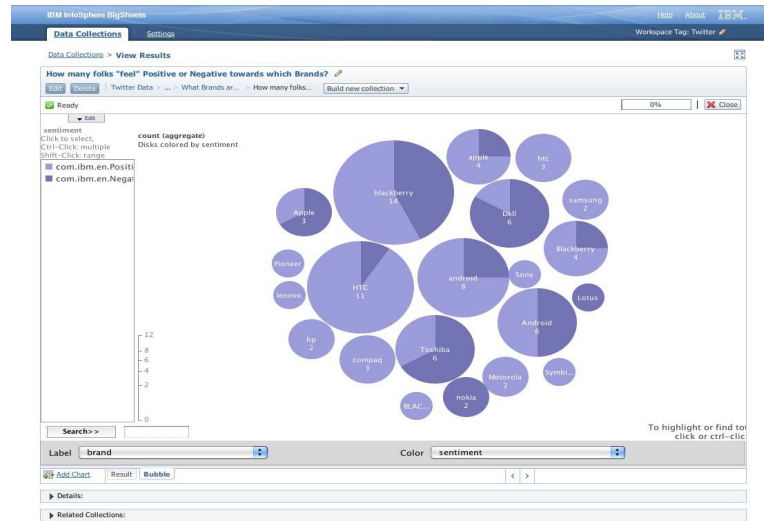
InfoSphere BigInsights Embrace and Extend Hadoop



BigSheets (Many Eyes)

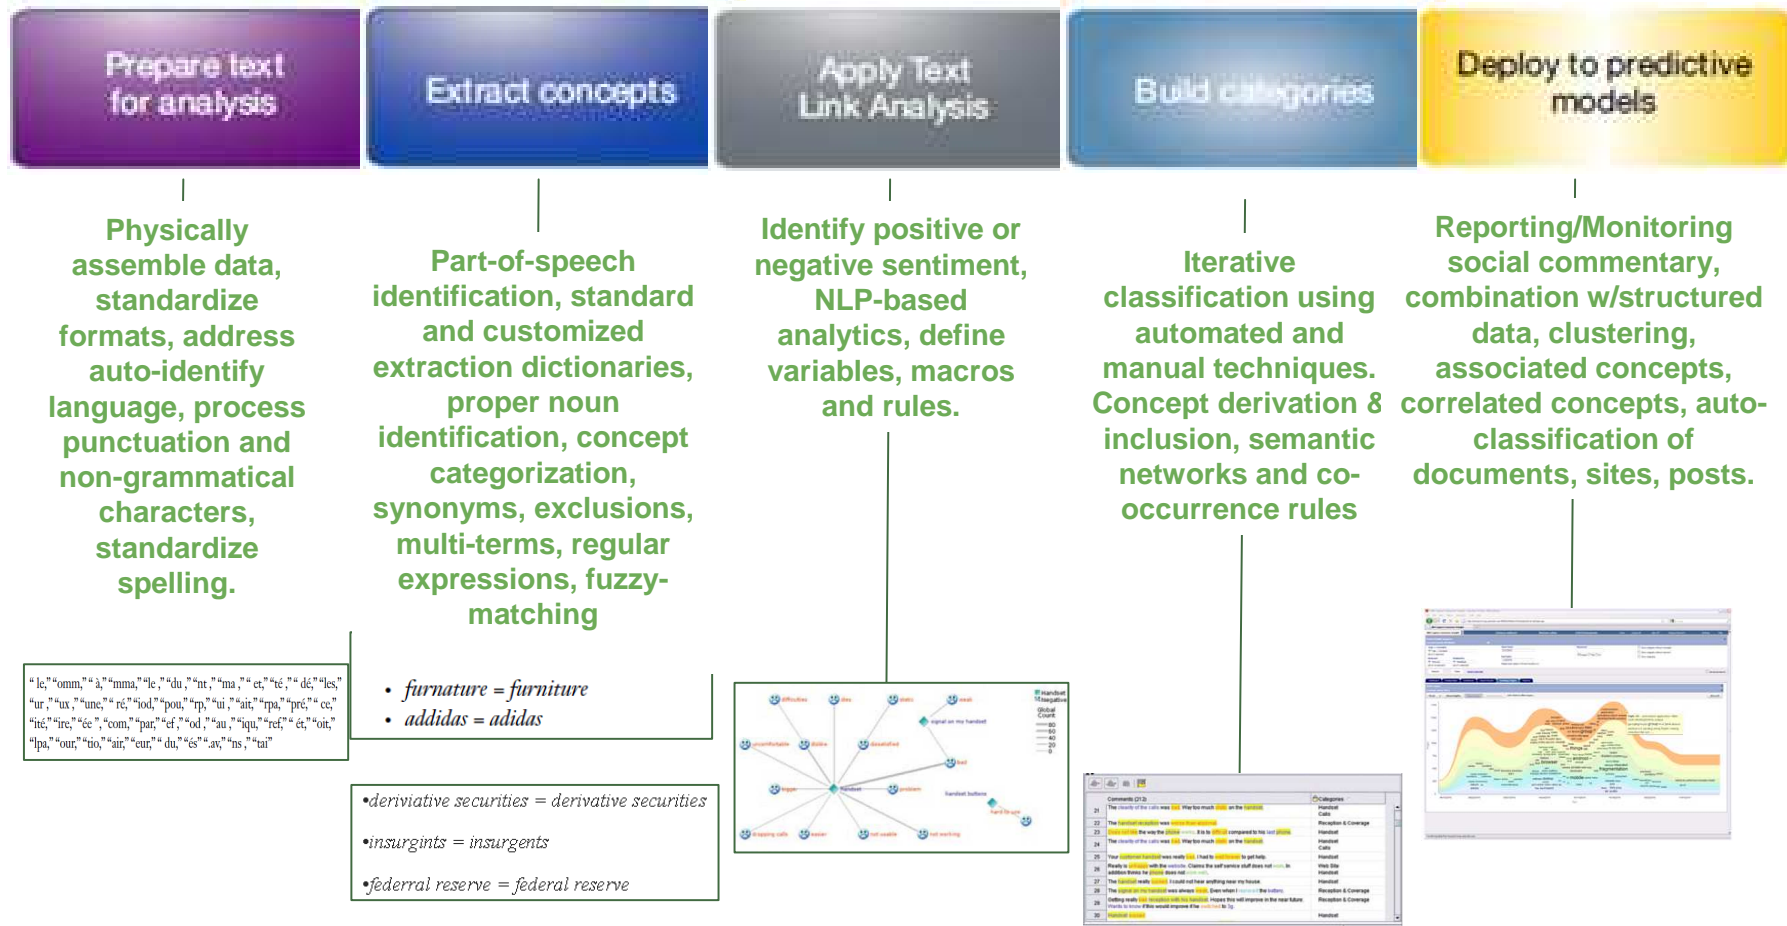
A visual tool for data manipulation and prototyping

- Ad-hoc analytics for LOB user
- Analyze a variety of data - unstructured and structured
- Spreadsheet metaphor for exploring/ visualizing data
- Browser-based



Text Analytics

Turns disparate words into measurable insights



Pre-configured text annotators ready for distributed processing on Big Data

Support for native languages including double-byte



Reduced turbine placement
identification from weeks to
hours

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

Capabilities Utilized

Hadoop System

Results

- 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

InfoSphere Streams

Analytical platform for Big Data in-motion

Built to analyze data in motion

- Multiple concurrent input streams
- Massive scalability

Process and analyze a variety of data

- Structured, unstructured content, video, audio
- Advanced analytic operators



InfoSphere Streams Massively Scalable Stream Analytics

Linear Scalability

- Clustered deployments – unlimited scalability

Automated Deployment

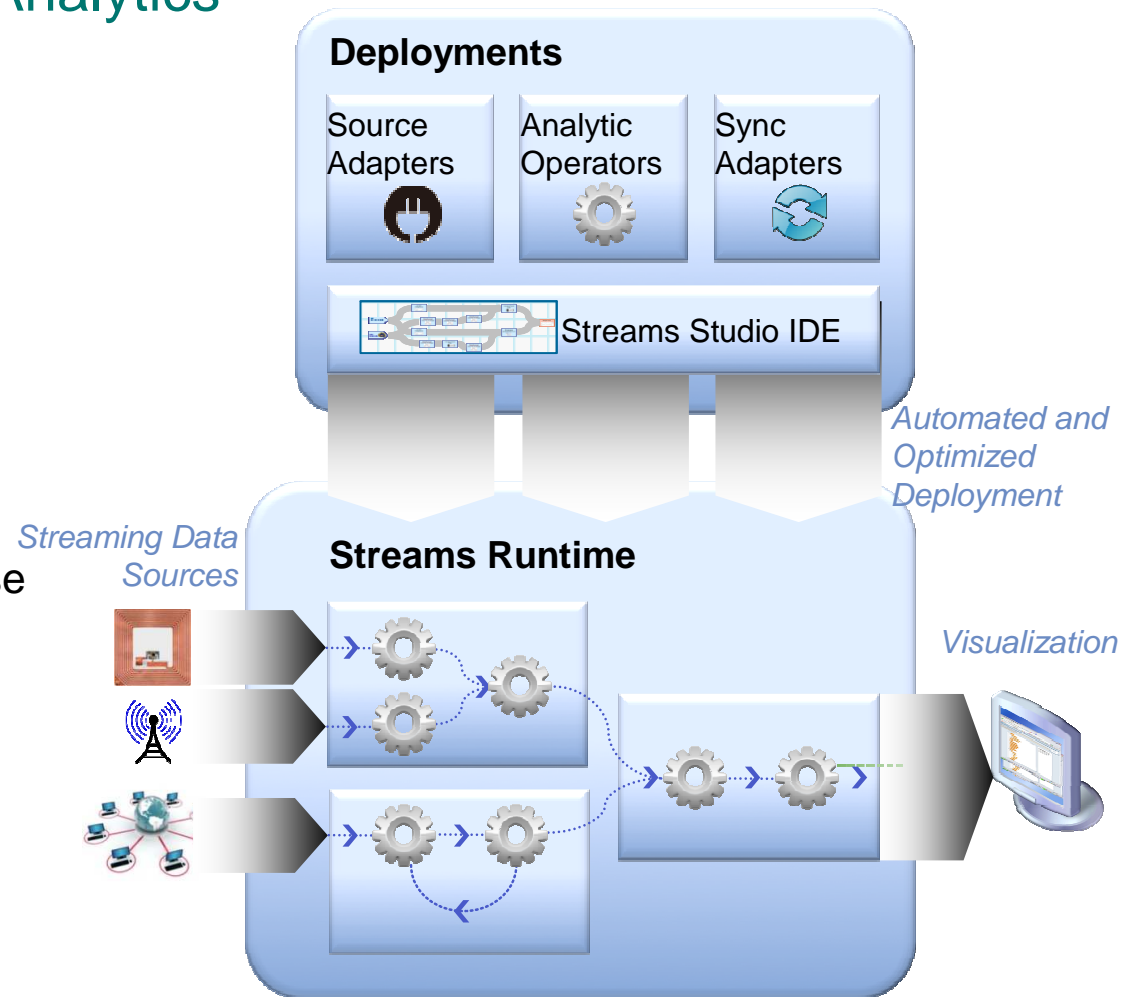
- Automatically optimize operator deployment across clusters

Performance Optimization

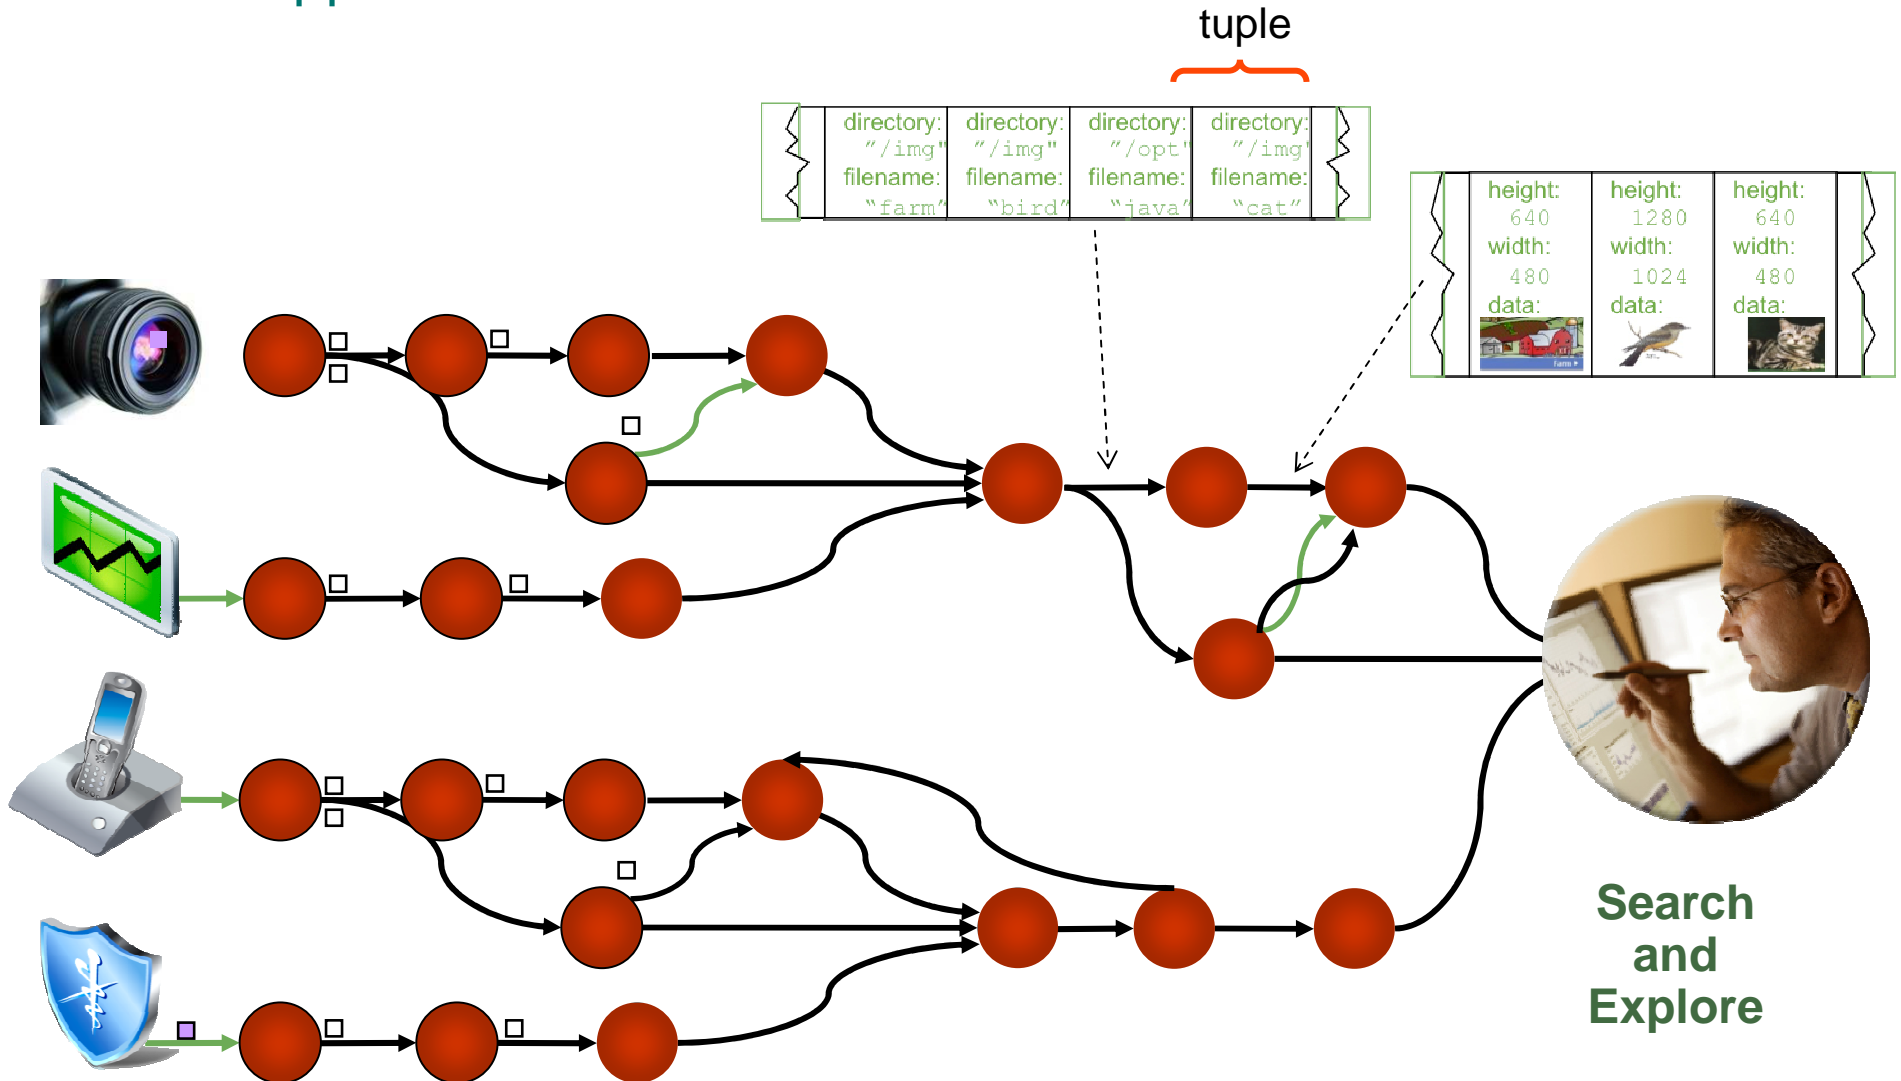
- JVM Sharing – minimize memory use
- Fuse operators on same cluster
- Telco client – 25 Million messages per second

Analytics on Streaming Data


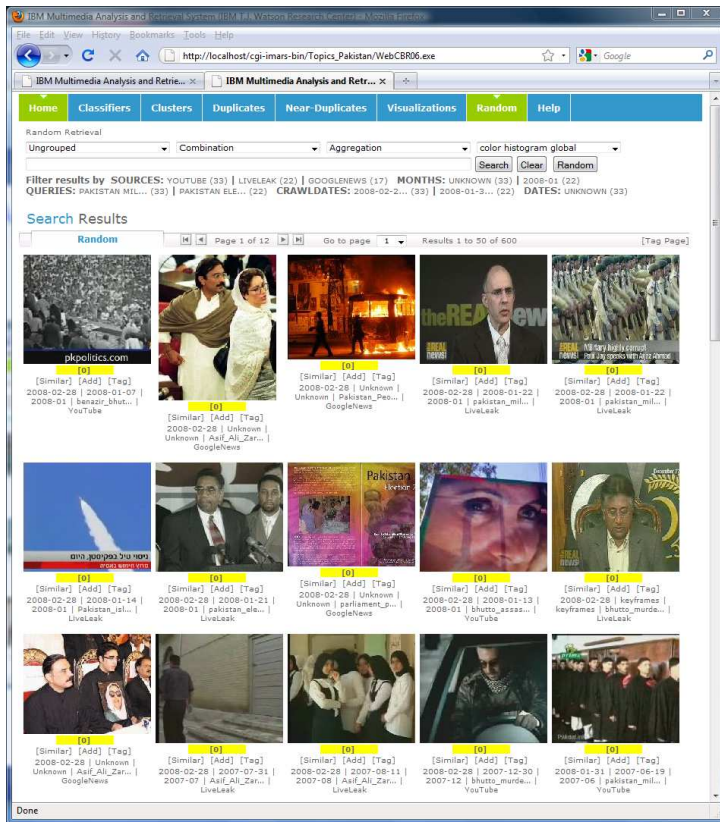
- Analytic accelerators for a variety of data types
- Optimized for real-time performance



Streams approach illustrated



Use Case Automatic Recognition of Scenes of Interest



Air Attack



Explosion



Demonstration



Individual



Wreckage



Military

Automatic scene recognition, clustering, duplicate and near-duplicate detection from image and video content



“Helps detect life threatening conditions up to 24 hours sooner”

University of Ontario Institute of Technology (UOIT) Detects Neonatal Patient Symptoms Sooner

Capabilities Utilized

Stream Computing

- Performing real-time analytics using physiological data from neonatal babies
- Continuously correlates data from medical monitors to detect subtle changes and alert hospital staff sooner
- Early warning gives caregivers the ability to proactively deal with complications

Results

- Helps detect life threatening conditions up to 24 hours sooner
- Lower morbidity and improved patient care





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

Real-time visualization and visibility into the arrival times of their 1000 buses on 150 routes

TerraEchos Turns to IBM Big Data for Low Latency Surveillance Data Analysis

Capabilities Utilized

Stream Computing

- Deployed security surveillance system to detect, classify, locate, and track potential threats at highly sensitive national lab
- Stream computing collects and analyzes acoustic data from fiber-optic sensor arrays
- Analyzed acoustic data fed into TerraEchos intelligence platform for threat detection, classification, prediction & communication

Results

- Enables Terraechos solution to analyze and classify streaming acoustic data in real-time
- Provides lab & security staff with holistic view of potential threats & non-issues
- Enables a faster and more intelligent response to any threat

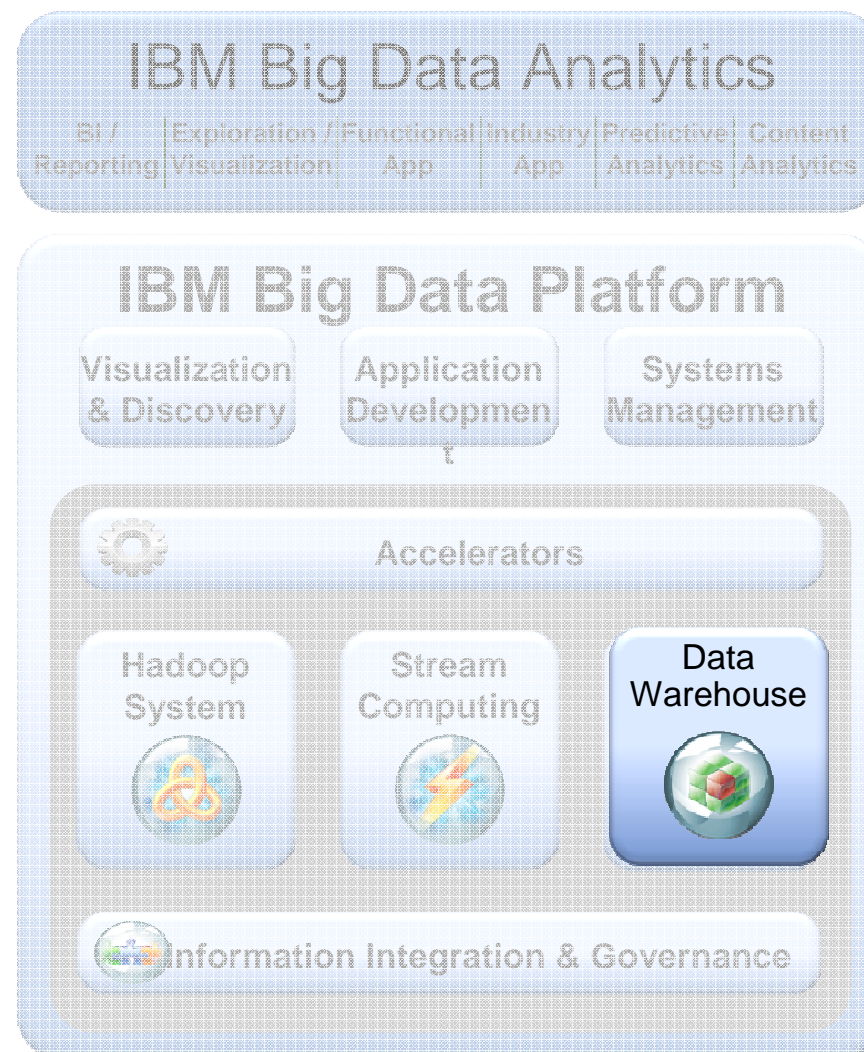
Identifies and classifies potential security threats miles away




IBM Netezza The True Data Warehouse Appliance

**Simpler, Faster, More accessible
Enterprise Class Analytics**

- Purpose-built analytics engine
- Integrated database, server and storage
- Standard interfaces
- Low total cost of ownership
- Speed: 10-100x faster than traditional system
- Simplicity: Minimal administration and tuning
- Scalability: Peta-scale user data capacity
- Smart: High-performance advanced analytics





Barnes & Noble helps
suppliers track sales and
inventory in real time

Capabilities Utilized:

Data Warehouse Appliance

Need

- Publishers absorb the losses on returns when they print too many books, and have stock-outs when they print too few
- To provide publishers with the ability to get real-time insight into sales and view inventory trends over time

Benefits

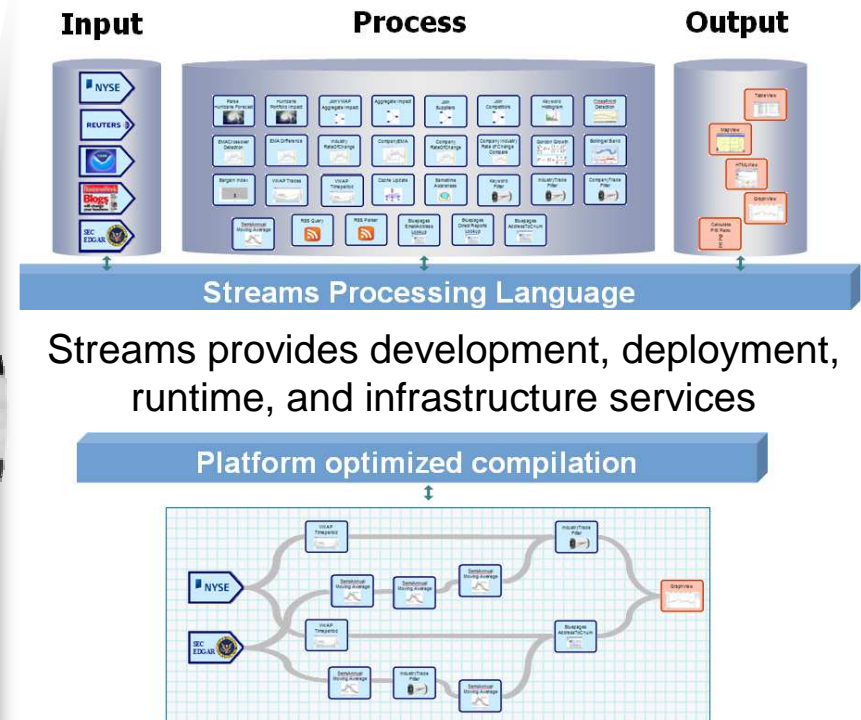
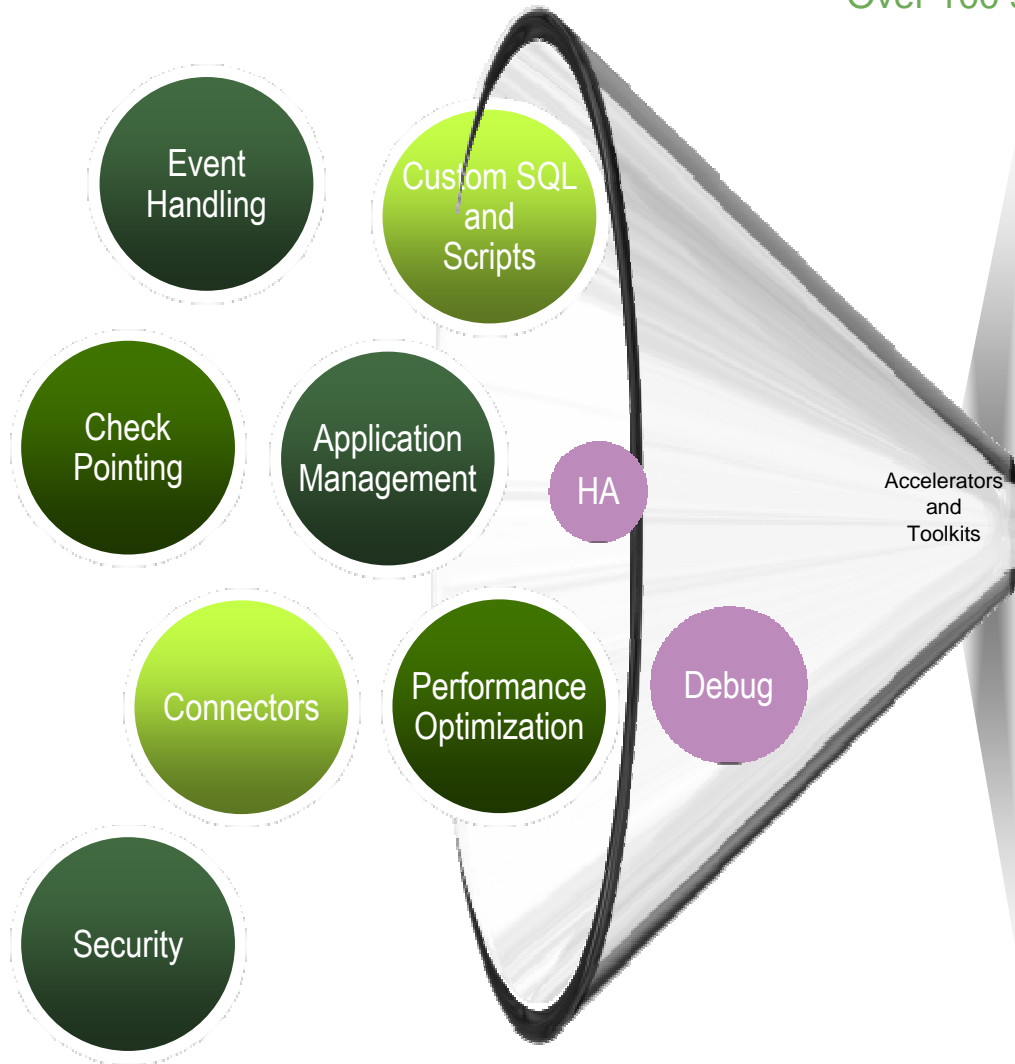
- Decreased time to run queries from weeks to seconds and enabled an 80% reduction in time to run compared to its previous system
- Reduced inventory levels and inventory carrying costs

80% reduction in run time

BARNES & NOBLE
BOOKSELLERS

Without a Big Data Platform You Code...

Over 100 sample applications and toolkits with industry focused toolkits with 300+ functions and operators



Streams provides development, deployment, runtime, and infrastructure services

“TerraEchos developers can deliver applications 45% faster due to the agility of Streams Processing Language...”
– Alex Philip, CEO and President, TerraEchos

IBM is Committed to Innovation



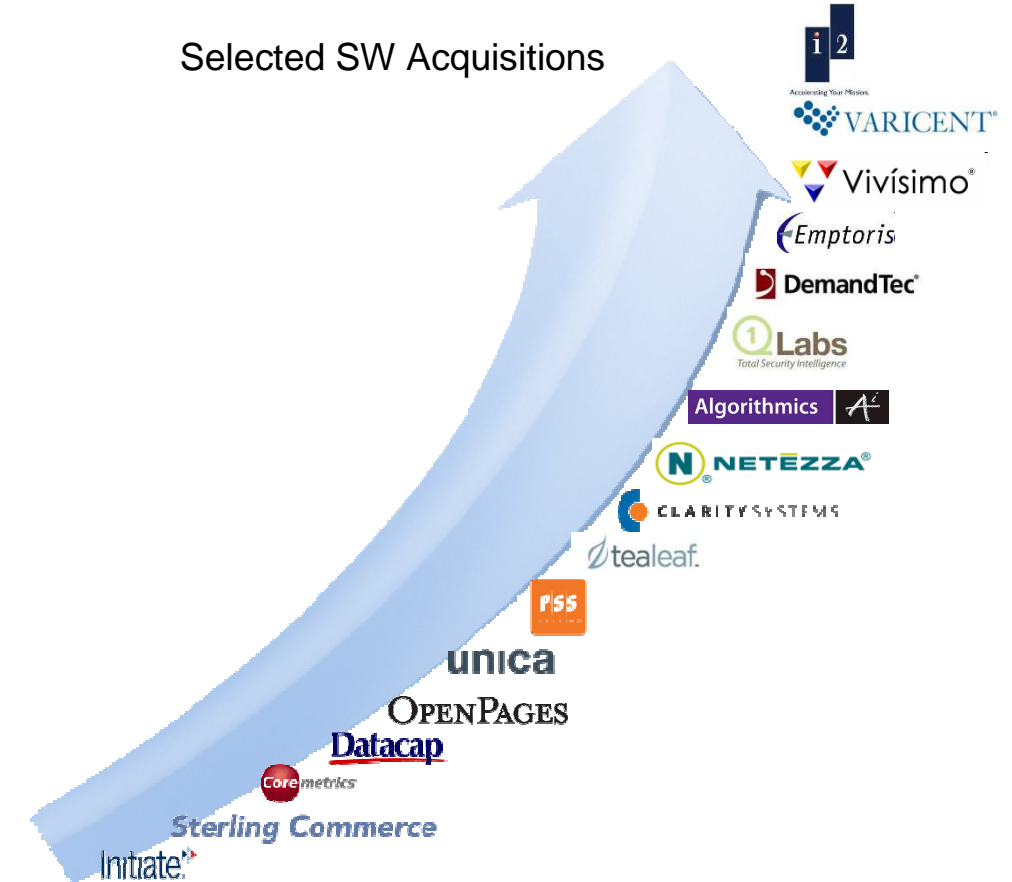
IBM Research

- Almaden
- Austin
- Melbourne
- Sao Paulo
- Beijing
- Haifa
- Delhi
- Ireland
- Yamato
- Watson
- Zurich

- **\$16B+** in acquisitions since 2005
- **10,000+** technical professionals
- **~8000** dedicated consultants
- **27,000+** business partner certifications
- **8** Analytics Solutions Centers
- **100** analytics-based research assets; almost **300** researchers

Selected SW Acquisitions

2012



“Watson is going to revolutionize many, many industries and it will fundamentally change the way we interact with computers & machines.”

John Kelly, SVP & Head of IBM Research

2005



* TeaLeaf, Varicent Vivismo pending acquisition close

More Information ...

bigdatauniversity.com/

ibm.com/software/data/infosphere/biginsights
ibm.com/software/data/infosphere/streams



Understanding Big Data
Analytics for Enterprise Class Hadoop and Streaming Data

- Learn how IBM hardens Hadoop for enterprise-class scalability and reliability
- Gain insight into IBM's unique in-motion and at-rest Big Data analytics platform
- Learn tips and tricks for Big Data use cases and solutions
- Get a quick Hadoop primer

CHRIS EATON DIRK DEROIS
 TOM DEUTSCH GEORGE LAPIS
 PAUL ZIKOPOULOS

ibm.com/software/data/bigdata/

youtube.com/user/ibmbigdata

Questions & Answers



Thank You

Legal Disclaimer

© IBM Corporation 2012. All Rights Reserved.

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

