



# Big Data

*Ayhan Önder*  
*Client Technical Professional, Netezza*

20.09.2012



# The World of DATA



# “Data is the new Oil”

*In its raw form, oil has little value. Once processed and refined, it helps power the world.*



*“Big Data has arrived at Seton Health Care Family, fortunately accompanied by an analytics tool that will help deal with the complexity of more than two million patient contacts a year...”*

## The New York Times

*“At the World Economic Forum last month in Davos, Switzerland, Big Data was a marquee topic. A report by the forum, “Big Data, Big Impact,” declared **data a new class of economic asset, like currency or gold.***



*“Increasingly, businesses are applying analytics to social media such as Facebook and Twitter, as well as to product review websites, to try to “understand where customers are, what makes them tick and what they want”, says Deepak Advani, who heads IBM’s predictive analytics group.”*

## THE WALL STREET JOURNAL

*“Companies are being inundated with data—from information on customer-buying habits to supply-chain efficiency. But many managers struggle to make sense of the numbers.”*



*“...now Watson is being put to work digesting millions of pages of research, incorporating the best clinical practices and monitoring the outcomes to assist physicians in treating cancer patients.”*

## Los Angeles Times


*The Oscar Senti-meter — a tool developed by the L.A. Times, IBM and the USC Annenberg Innovation Lab — analyzes opinions about the Academy Awards race shared in millions of public messages on Twitter.”*



*“Data is the new oil.”  
Clive Humby*

# The Big Data Opportunity




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



## What Our Customers Tell Us . . .

- Don't know what should be analyzed
  
- Volumes can be extremely high
  - Potentially valuable data is dormant or discarded (size/performance)
  - Too expensive to justify integrating large volumes of unstructured data
  
- Much of their data is unstructured, or in widely varying structures, which are difficult to analyze
  
- Difficult to integrate information distributed across multiple systems and the Internet
  
- Some information has a short useful lifespan
  
- Analysis needed in the context of their existing information

### Information Sources

- Point of Sale 
- Call Center 
- Customer 

Structured Information



Reporting, Trend Analysis, Decision Making



At Rest Analytics

### Users




Executive



Power User








Call Center



Customers

### Untapped Sources

- 
- 
- 
- 
- 

V<sup>3</sup>

Variety  
Structured & Unstructured

Velocity  
Speed of data  
In motion  
analytics

Volume  
Petabytes,  
Zetabytes



### Big Data Platform

- Ability to handle V<sup>3</sup>
- Enable Analytics  
In motion & At Rest  
Text, Video, Audio...
- Ease of Use for  
Developers
- Enterprise Class  
GPFS, Security,  
Governance,  
Data Integration

Ease of Use for  
End Users




Easy  
Visualization  
Tools

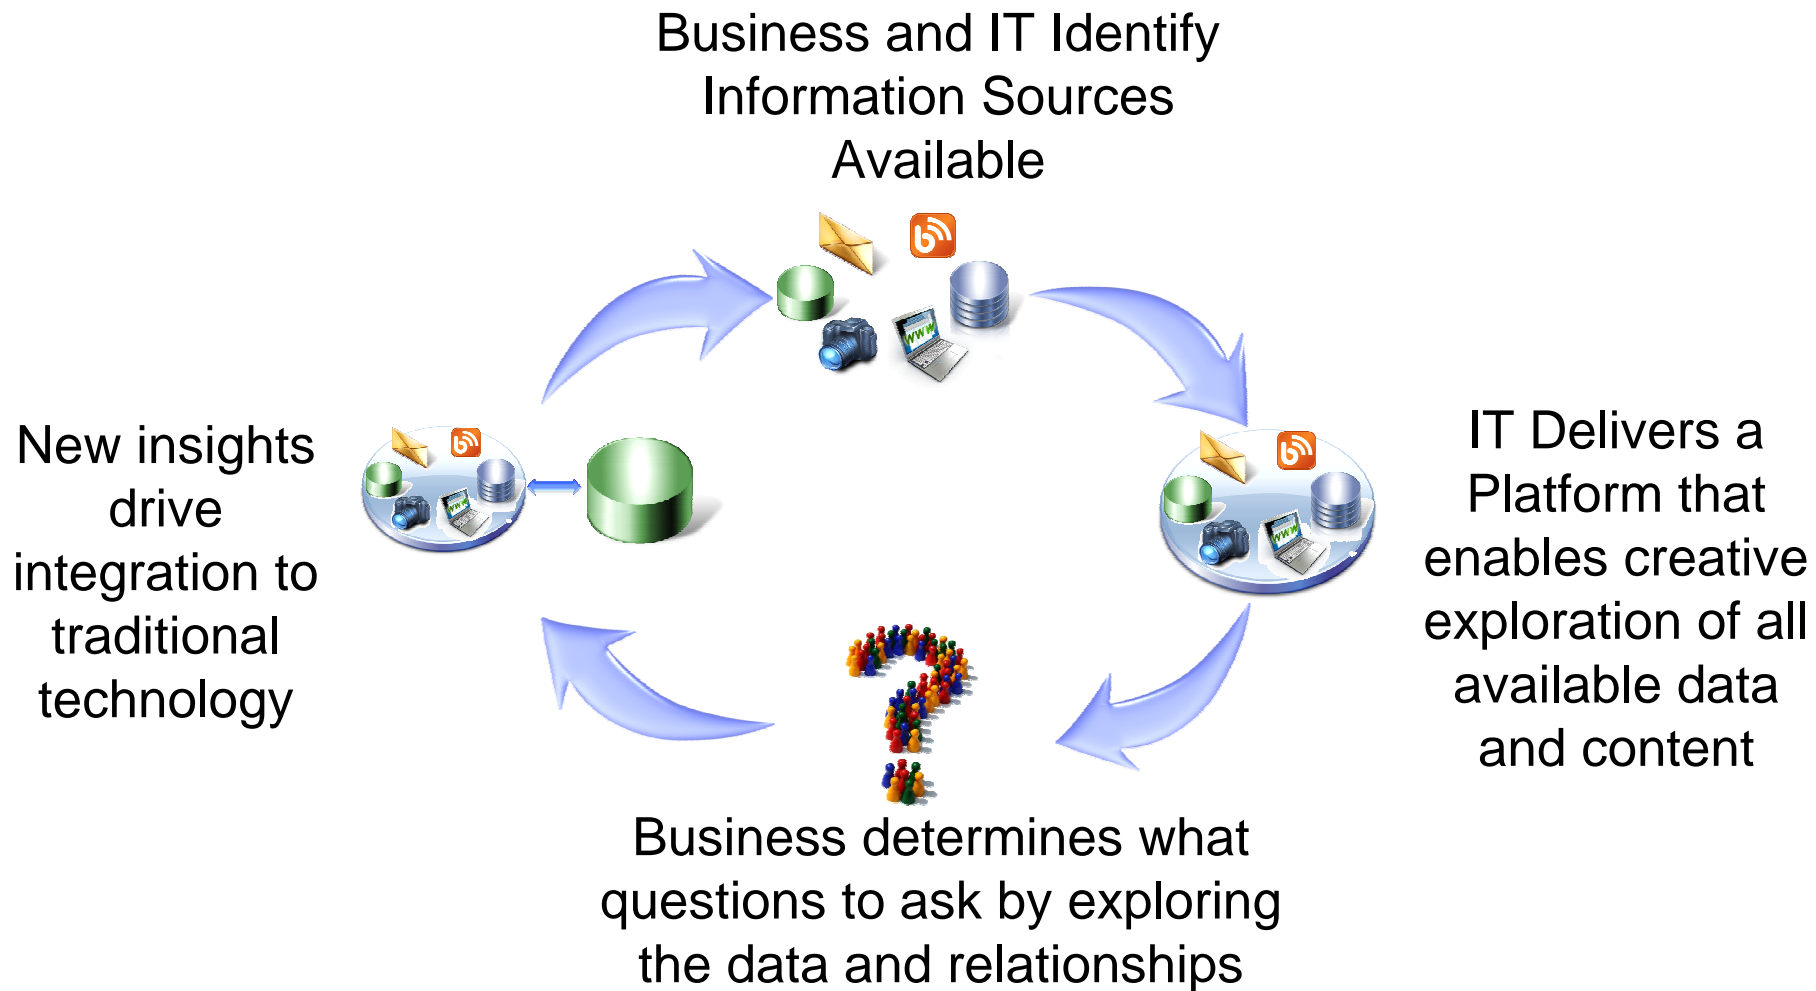
Better Business  
Outcomes



# Traditional Analytics: *Business Requirements Drive Solution Design*



# Big Data Analytics: *Information Sources Drive Creative Discovery*






# Merging the Traditional and Big Data Approaches

**Traditional Approach**  
*Structured & Repeatable Analysis*

**Big Data Approach**  
*Iterative & Exploratory Analysis*

Business Users  
 Determine what  
 question to ask




IT  
 Structures the  
 data to answer  
 that question




Monthly sales reports  
 Profitability analysis  
 Customer surveys



IT  
 Delivers a platform to  
 enable creative  
 discovery

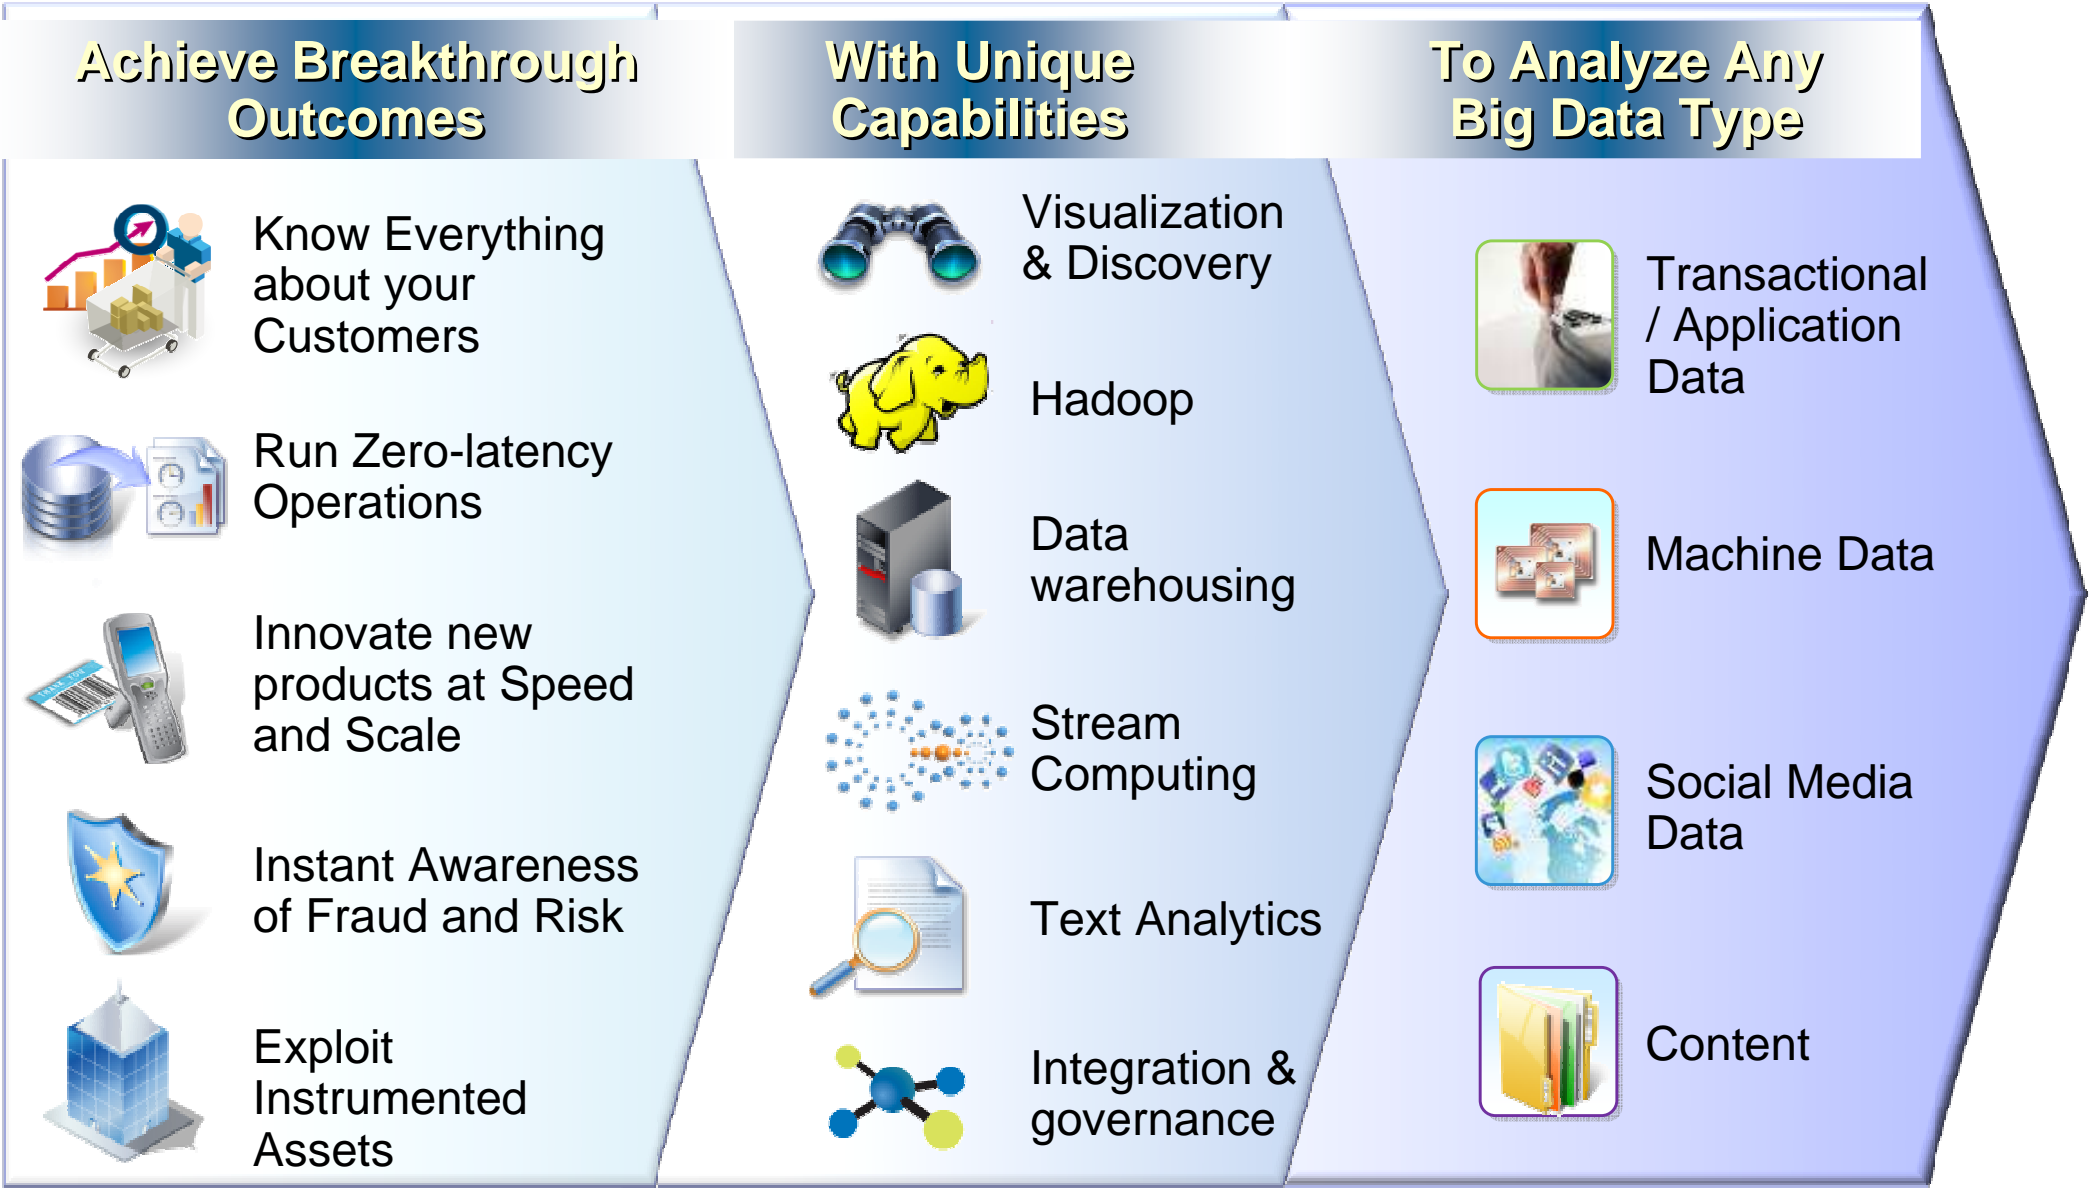


Business  
 Explores what  
 questions could be  
 asked




Brand sentiment  
 Product strategy  
 Maximum asset utilization

# Achieve Breakthrough Outcomes with Big Data capabilities




# Most Use Cases Combine Multiple Technologies



## Pre-processing

Ingest and analyze unstructured data types and convert to structured data




## Combine structured and unstructured analysis

Augment data warehouse with additional external sources, such as social media



## Combine high velocity and historical analysis

Analyze and react to data in motion; adjust models with deep historical analysis




## Reuse structured data for exploratory analysis

Experimentation and ad-hoc analysis with structured data

# Opportunities To Exploit Big Data Are Everywhere

## Analyze Information in Motion

- Smart Grid management
- Multimodal surveillance
- Real-time promotions
- Cyber security
- ICU monitoring
- Options trading
- Click-stream analysis
- CDR processing
- IT log analysis
- RFID tracking & analysis



## Analyze a Variety of Information

- Social media/sentiment analysis
- Geospatial analysis
- Brand strategy
- Scientific research
- Epidemic early warning system
- Market analysis
- Video analysis
- Audio analysis


## Discovery & Experimentation

- Sentiment analysis
- Brand strategy
- Scientific research
- Ad-hoc analysis
- Model development
- Hypothesis testing
- Transaction analysis to create insight-based product/service offerings



## Analyze Extreme Volumes of Information

- Transaction analysis to create insight-based product/service offerings
- Fraud modeling & detection
- Risk modeling & management
- Social media/sentiment analysis
- Environmental analysis




## Manage and Plan

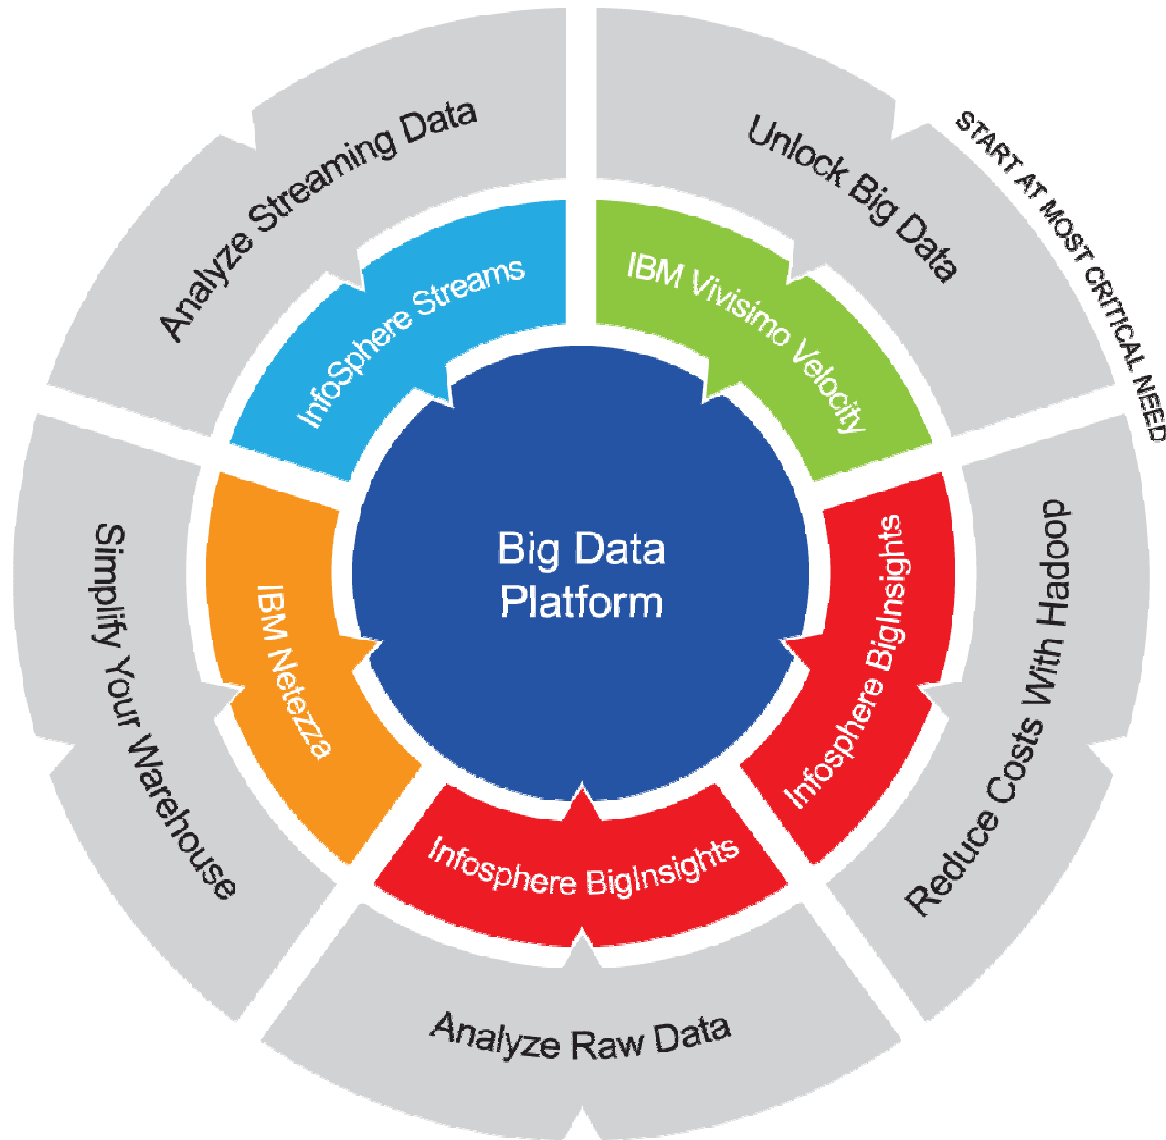
- Operational analytics – BI reporting
- Planning and forecasting analysis
- Predictive analysis

# Big Data: Integrated in Enterprise Information Architecture

- Big Data is Here to Stay
- It Must Not Be a Silo
  - It must be fully integrated for it to deliver value
- It Must be Easy to Deploy and Integrate

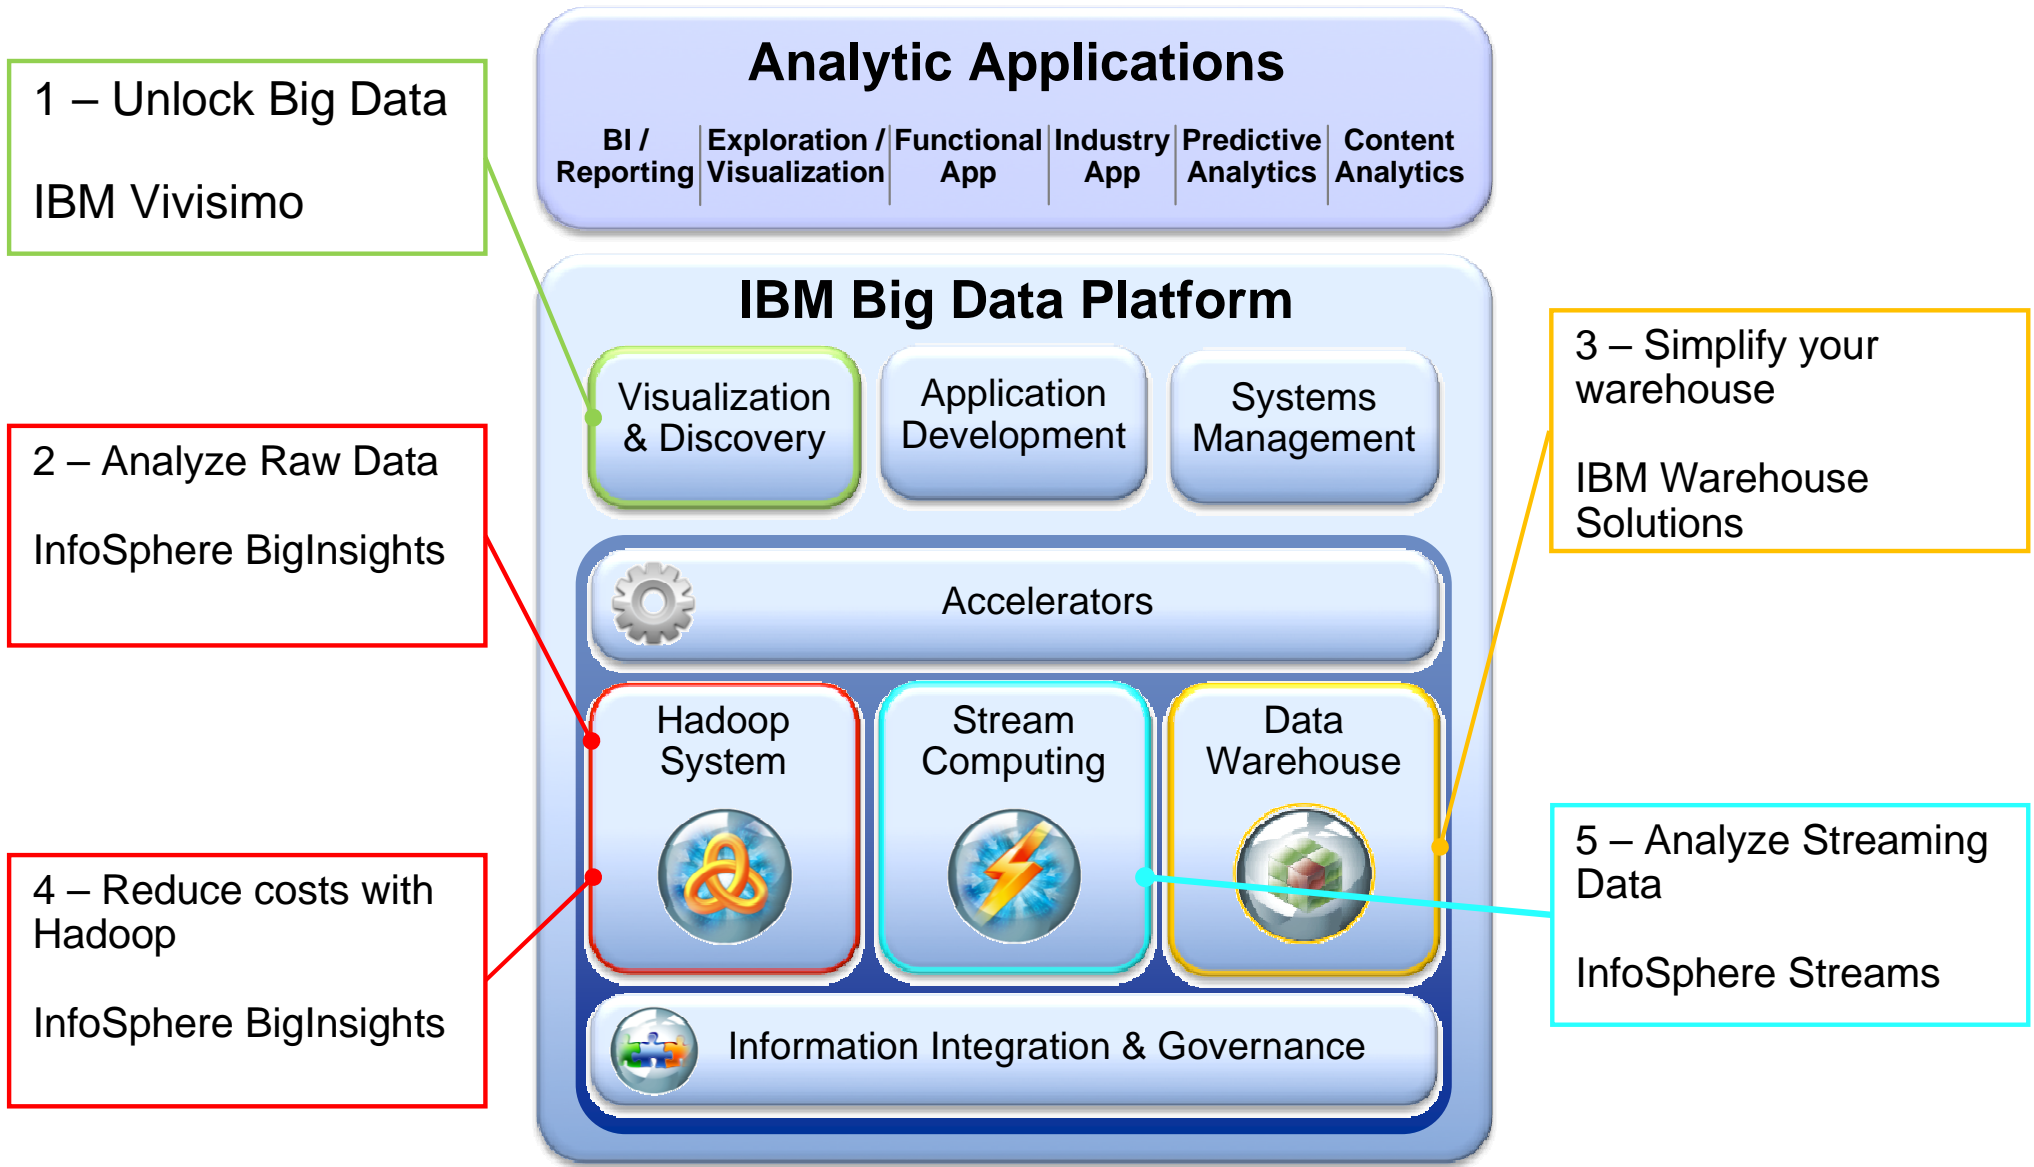


# Business-centric Big Data enables you to start with a critical business pain and expand the foundation for future requirements



- “Big data” isn’t just a technology—it’s a business strategy for capitalizing on information resources
- Getting started is crucial
- Success at each entry point is accelerated by products within the big data platform
- Build the foundation for future requirements by expanding further into the big data platform

# Entry Points are Accelerated by Products Within the Big Data Platform



days for a single query

constant tuning

“

*Nearly 70% of data warehouses experience performance-constrained issues of various types.*

”

- Gartner 2010 Magic  
Quadrant

specialized resources required

months to deploy



## Highway or off-road? Building a machine for its workload




# Transactional workloads vs. analytic workloads


Two *VERY* different requirements for storing and processing data



# Traditional Data Warehouse Approach



# Netezza – Simplicity



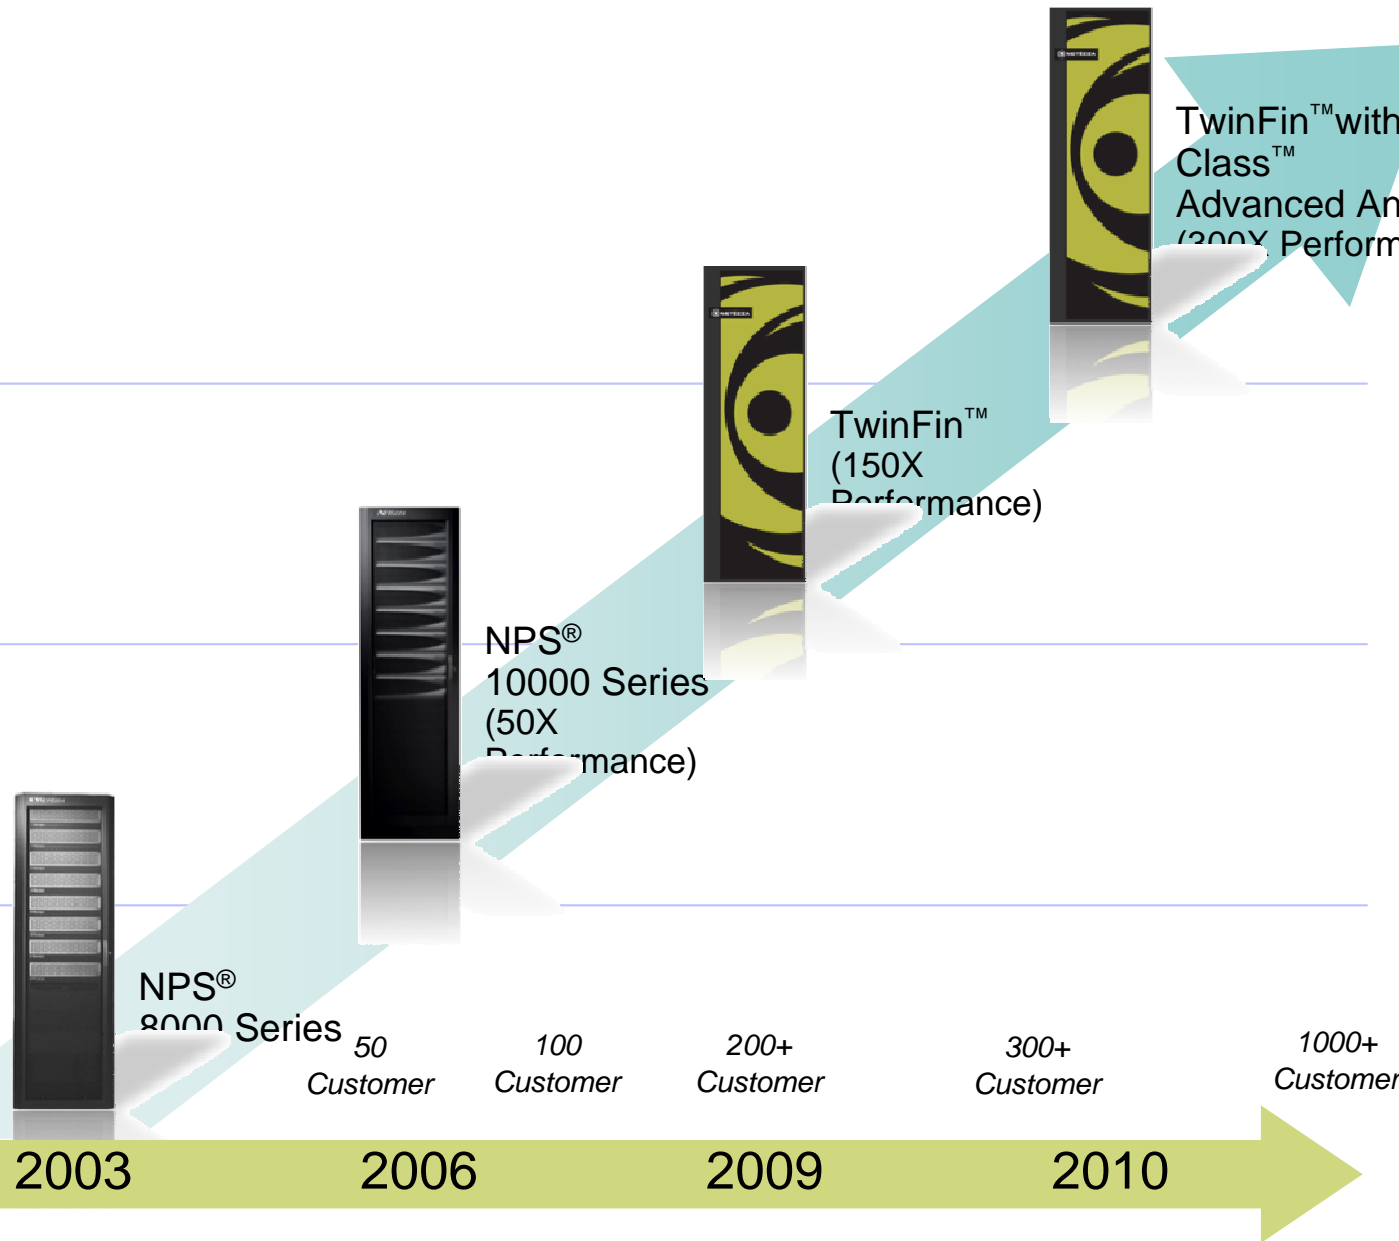
# Netezza

First Analytical Appliance in the world

First Petabyte Size Data Warehouse

First 100 TB Size Data Warehouse

Impact



“ *Netezza was part of the inspiration for Exadata ...*

*We'd like to thank them for forcing our hand and forcing us to go into the hardware business.* ”

-- Larry Ellison, CEO  
Oracle Corporation  
January 27 2010

[http://oracle.com.edgesuite.net/ivt/4000/8104/9238/12652/lobby\\_external\\_flash\\_clean\\_480x360/default.htm](http://oracle.com.edgesuite.net/ivt/4000/8104/9238/12652/lobby_external_flash_clean_480x360/default.htm)



## IBM Netezza

### The Simple Appliance for Serious 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

# Inside IBM Netezza data warehouse appliances

## Optimized Hardware + Software

Purpose-built for high performance analytics; requires no tuning

## True MPP


All processors fully utilized for maximum speed and efficiency

### Streaming Data

*Hardware-based query acceleration for blistering fast results*

### Deep Analytics

*Complex analytics executed in-database for deeper insights*





## Netezza delivers simplicity

- Up and running 6 months before being trained
- 200X faster than Oracle system
- ROI in less than 3 months



“

*Allowing the business users access to the Netezza box was what sold it.*”

-- Steve Taff,  
Executive Dir. of IT Services



# IBM Netezza Appliance family



**IBM Netezza 100**

Development &  
Test System

1 TB to 10 TB



**IBM Netezza 1000**

Data Warehouse  
High-Performance  
Analytics

1 TB to 1.5 PB



**IBM Netezza High Capacity Appliance**

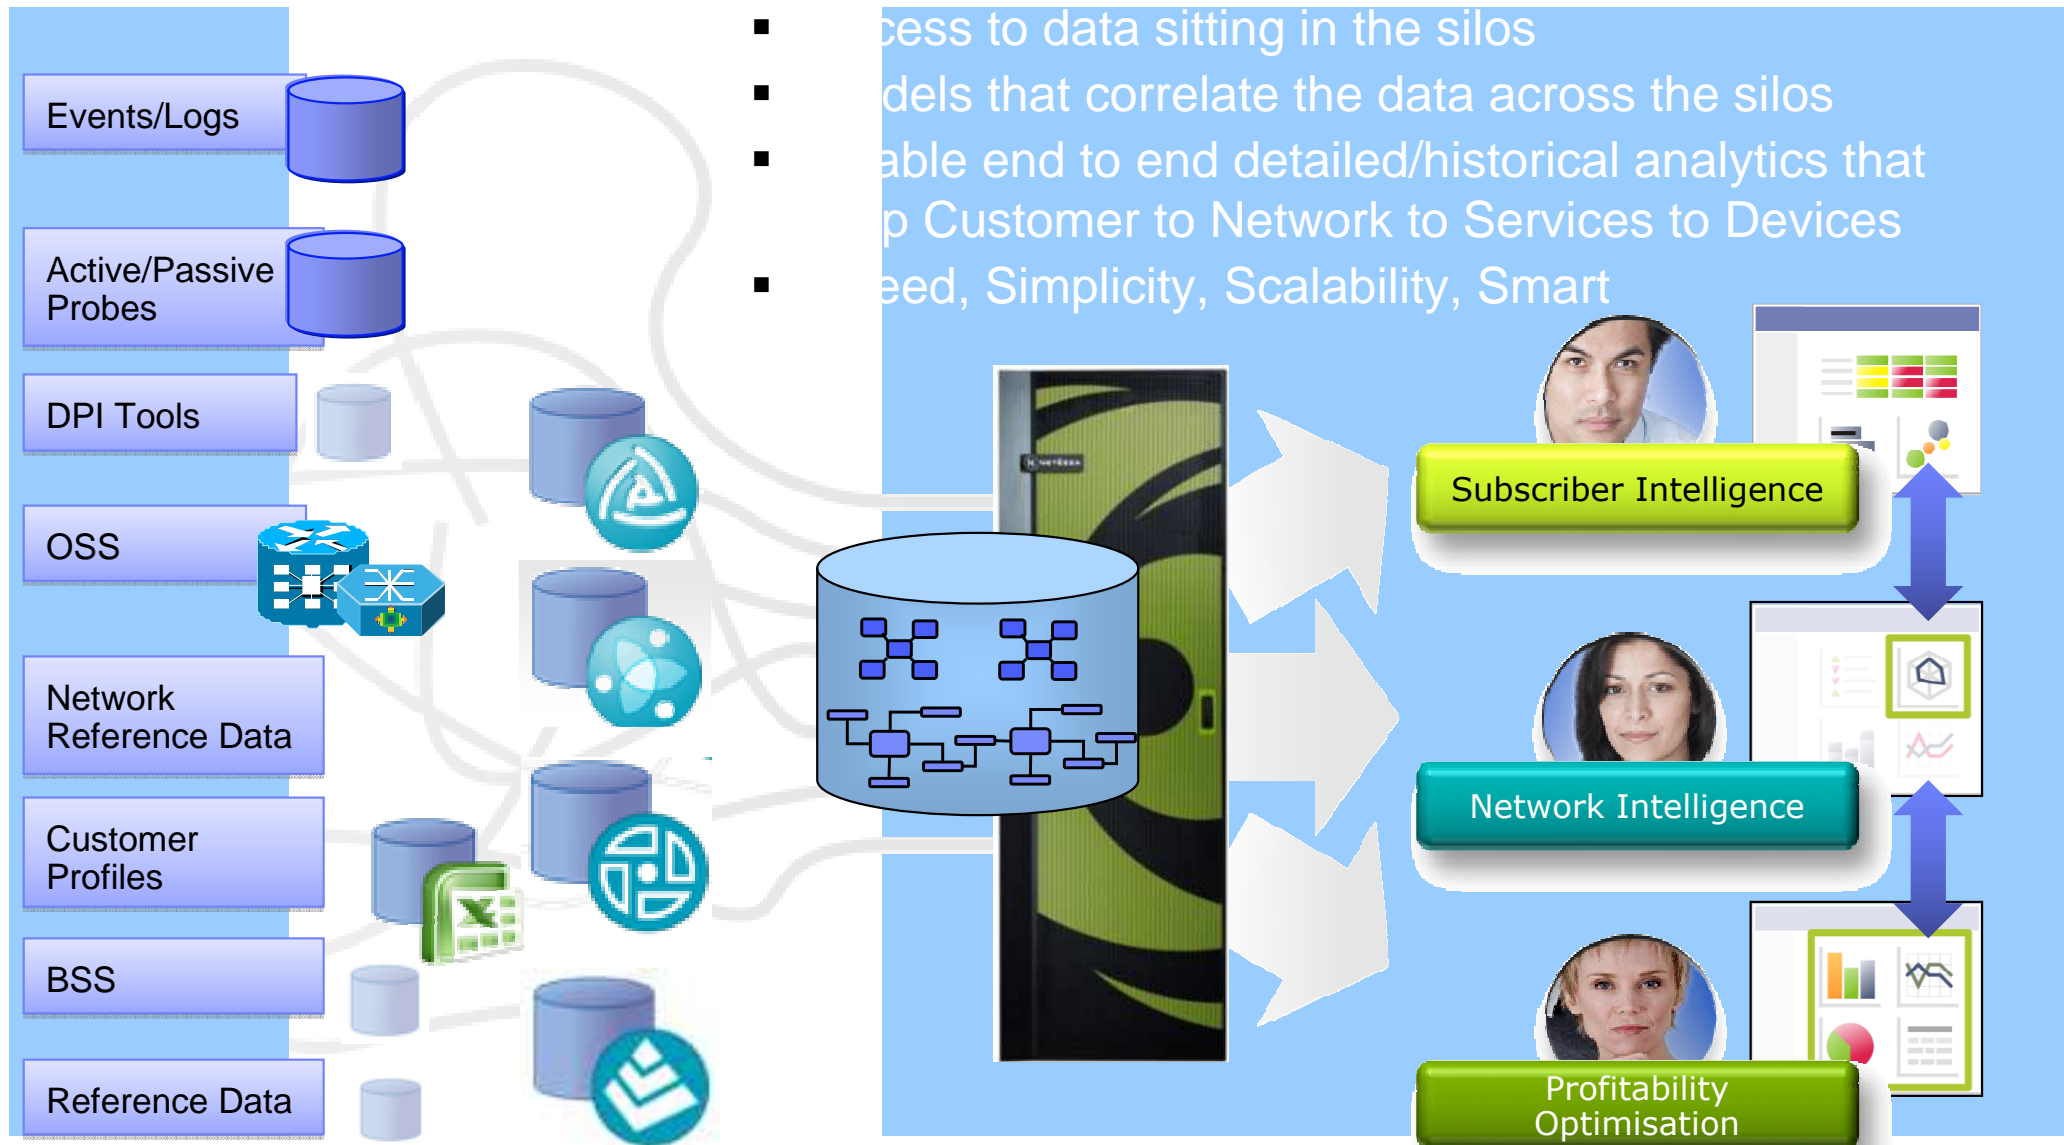
Queryable Archiving  
Backup/DR

100 TB to 10 PB



# Network Analytics with Netezza


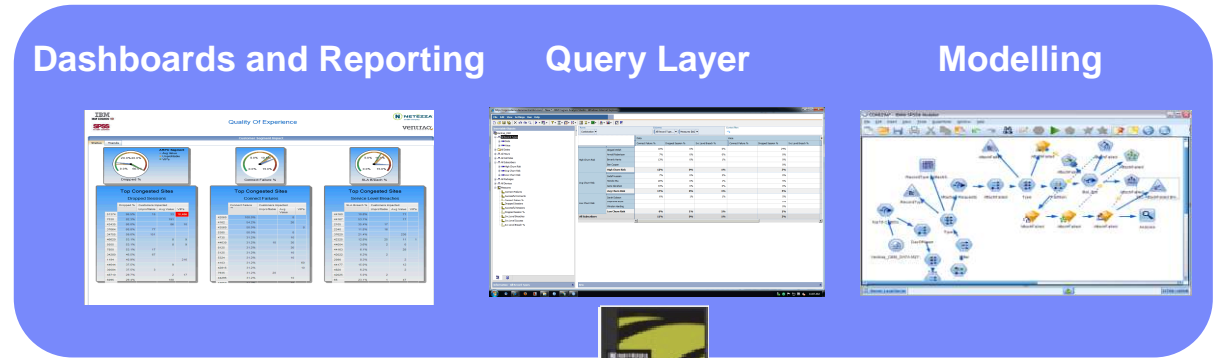
*Single view of Network, Customers, Services and Devices*



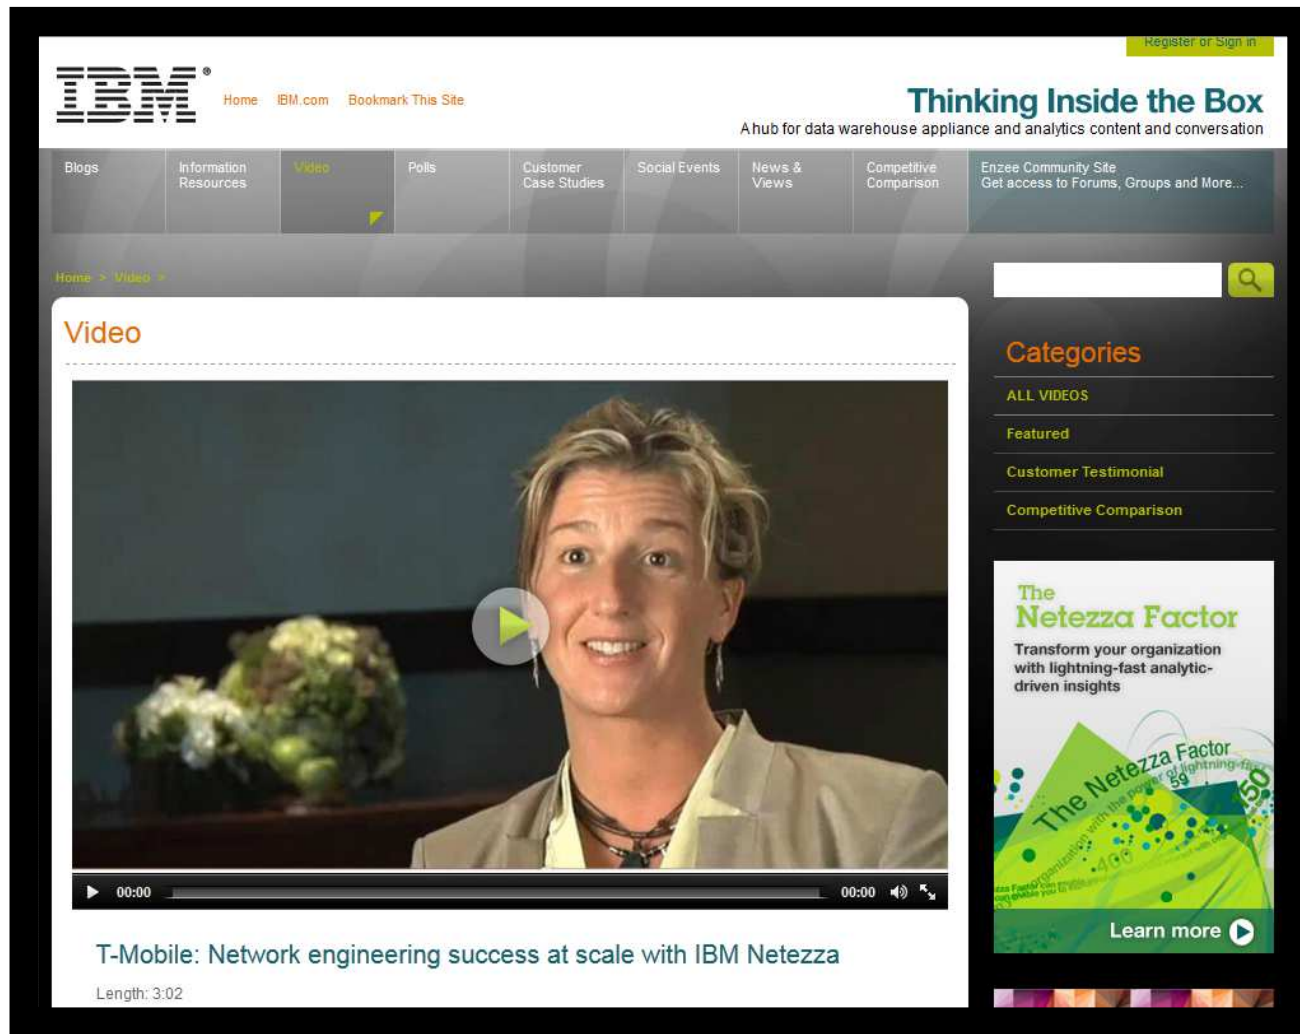
- Access to data sitting in the silos
- Models that correlate the data across the silos
- Available end to end detailed/historical analytics that span Customer to Network to Services to Devices
- Speed, Simplicity, Scalability, Smart

# Out of the box solutions for Network Analytics

- Data acquired directly from network elements or systems, OSS/BSS applications or other data stores via landing zone
- Load, analyze and explore massive volumes of detailed events
- Pre-built KPIs/KQIs, reports and dashboards to accelerate time to value
- Single source of network data for consumption across the enterprise
- Configurability of alerts for proactive management




# Results – In T-Mobile’s (USA) own words



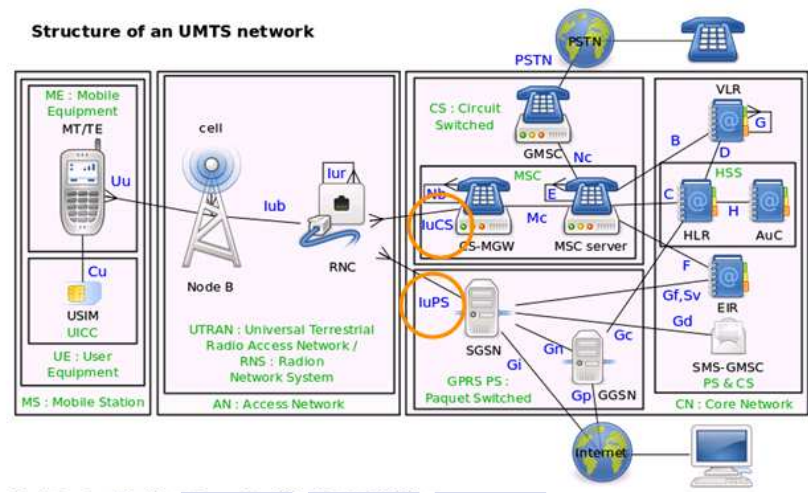
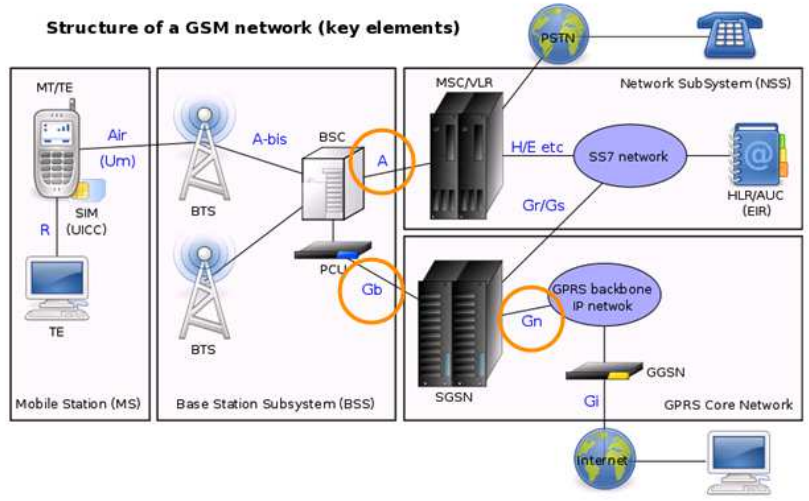
The screenshot shows the IBM 'Thinking Inside the Box' website. The main navigation bar includes 'Home', 'IBM.com', and 'Bookmark This Site'. The site title is 'Thinking Inside the Box' with the subtitle 'A hub for data warehouse appliance and analytics content and conversation'. A secondary navigation bar contains links for 'Blogs', 'Information Resources', 'Video', 'Polls', 'Customer Case Studies', 'Social Events', 'News & Views', and 'Competitive Comparison'. A search bar is located on the right side of the page. The main content area features a video player with a play button overlay. Below the video player, the title 'T-Mobile: Network engineering success at scale with IBM Netezza' is displayed, along with the video length 'Length: 3:02'. To the right of the video player, there is a 'Categories' section with links for 'ALL VIDEOS', 'Featured', 'Customer Testimonial', and 'Competitive Comparison'. Below the categories, there is a promotional graphic for 'The Netezza Factor' with the text 'Transform your organization with lightning-fast analytic-driven insights' and a 'Learn more' button.

<http://thinking.netezza.com/video/t-mobile-network-engineering-success-scale-ibm-netezza>

# T-Mobile USA



# Turkcell, Big Data PoC



Anritsu

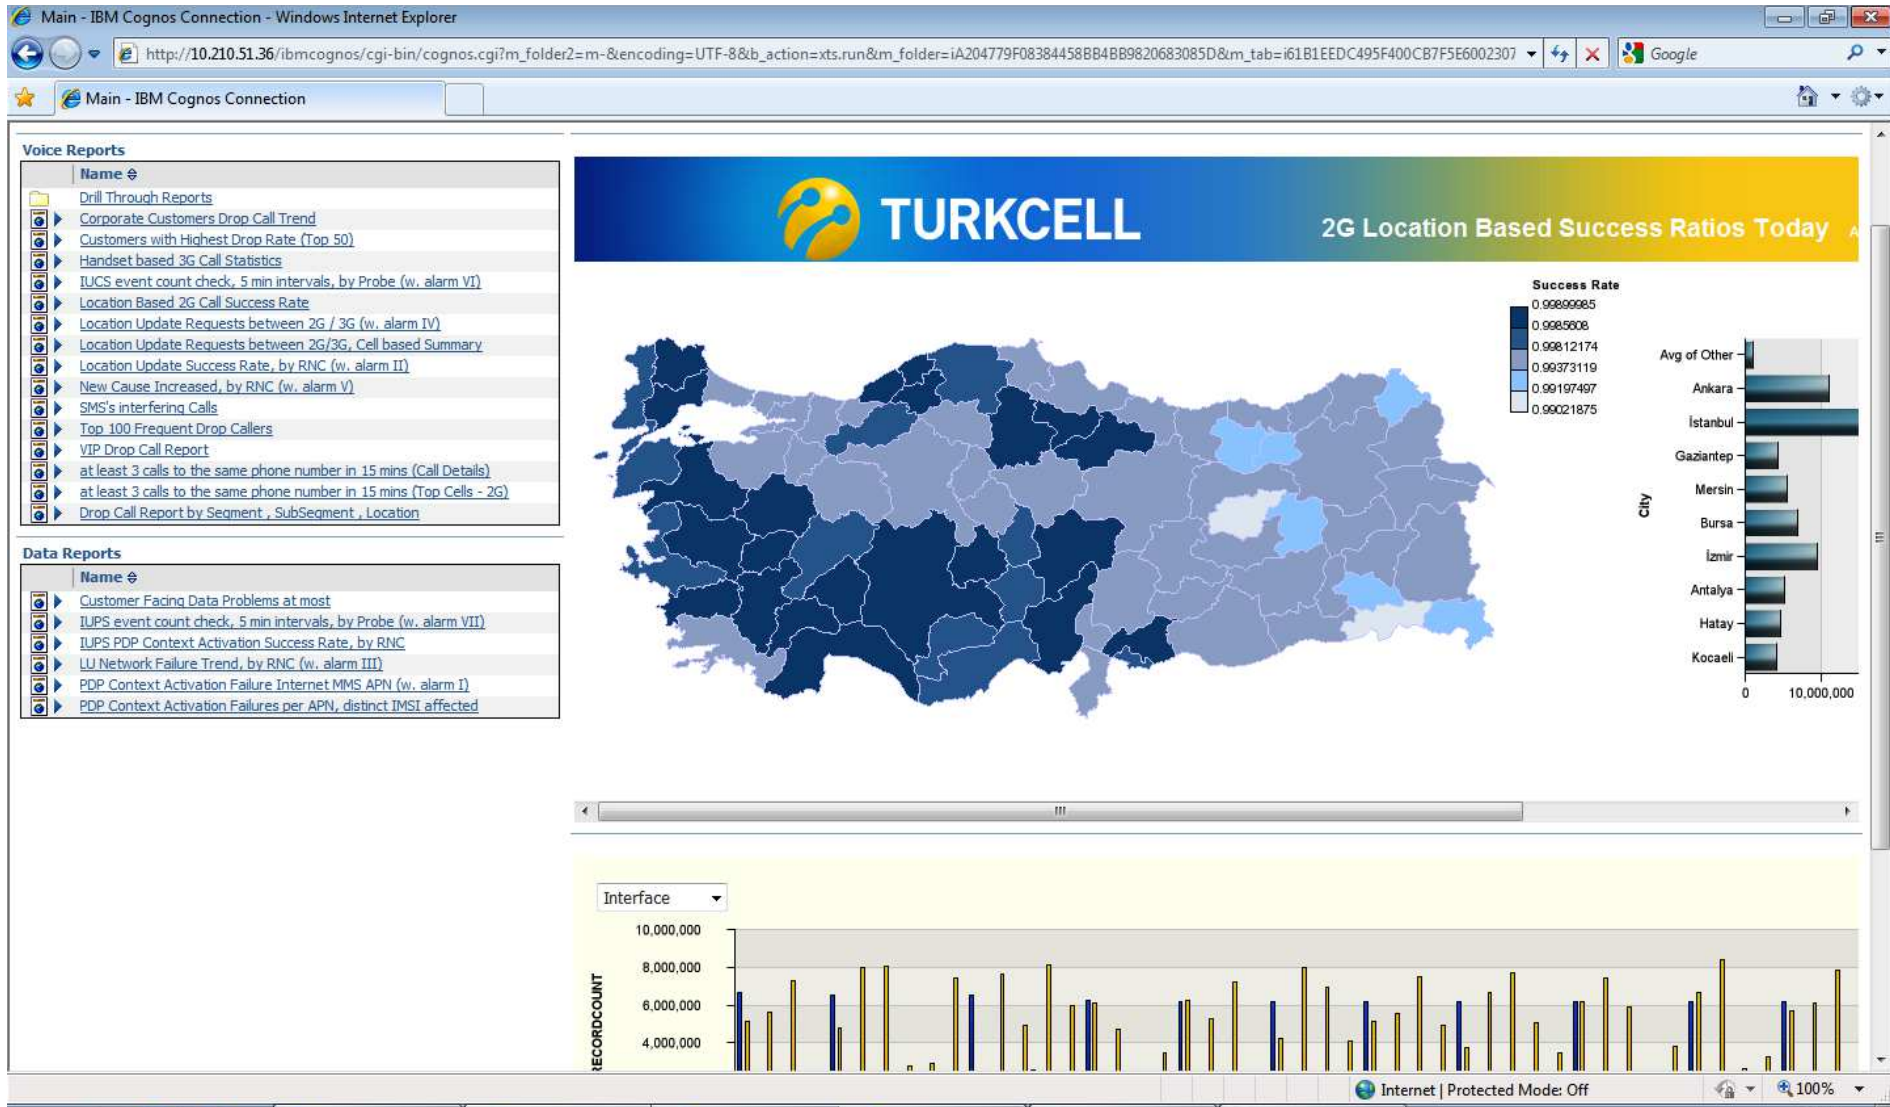
FTP →

- ADR** Single file in each 5 minutes  
Avg.1.3 billion event per day
- GB** Single file in each 5 minutes  
Avg.440 million event per day
- GN** Multiple files continuously  
Avg.5.2 billion event per day
- IuCS** Single file in each 5 minutes  
Avg.440 million event per day
- IuPS** Single file in each 5 minutes  
Avg.625 million event per day

NAT    SGSN    MAP




# Cognos, Data & Voice Reports



# Customers with Highest Drop Rate (Top 50)

Customers facing Highest Droprate (Top50) for the given date. Produced report includes 2G / 3G filtering, as well as graphs allowing drop ratio and drop count comparison for each IMSI.



## Query Performance

	<b>Observed Duration in Cognos</b> (includes report rendering / download time, while data loads are going on)
SMS' Interferring Calls	< 15 sec.
PDP Context Activation Failures per APN, distinct IMSI affected	< 5 sec.
PDP Context Activation Failure Internet MMS APN	< 10 sec.
New Cause Increased, by RNC	< 3 sec.
LU Network Failure Trend, by RNC	< 3 sec.
Location Update Success Rate, by RNC	< 4 sec.
Location Update Requests between 2G / 3G, Cell based Summary	< 4 sec.
Location Update Requests between 2G / 3G	< 4 sec.
Location Based 2G Call Success Rate	< 4 sec.
IUPS PDP Context Activation Success Rate, by RNC	< 3 sec.
IUPS event count check, 5 min intervals, by Probe	< 3 sec.
IUCS event count check, 5 min intervals, by Probe	< 3 sec.
Handset Based 3G Call Statistics	< 6 sec. (for 3 days)
Drop Call Report by Segment , SubSegment , Location	< 10 sec.
Customers with Highest Drop Rate (Top 50)	< 10 sec.
Customer Facing Data Problems at most	< 5 sec.
Corporate Customers Drop Call Trend	< 5 sec.
at least 3 calls to the same phone number in 15 mins (Top Cells - 2G)	< 10 sec.
at least 3 calls to the same phone number in 15 mins (Call Details)	< 5 sec.

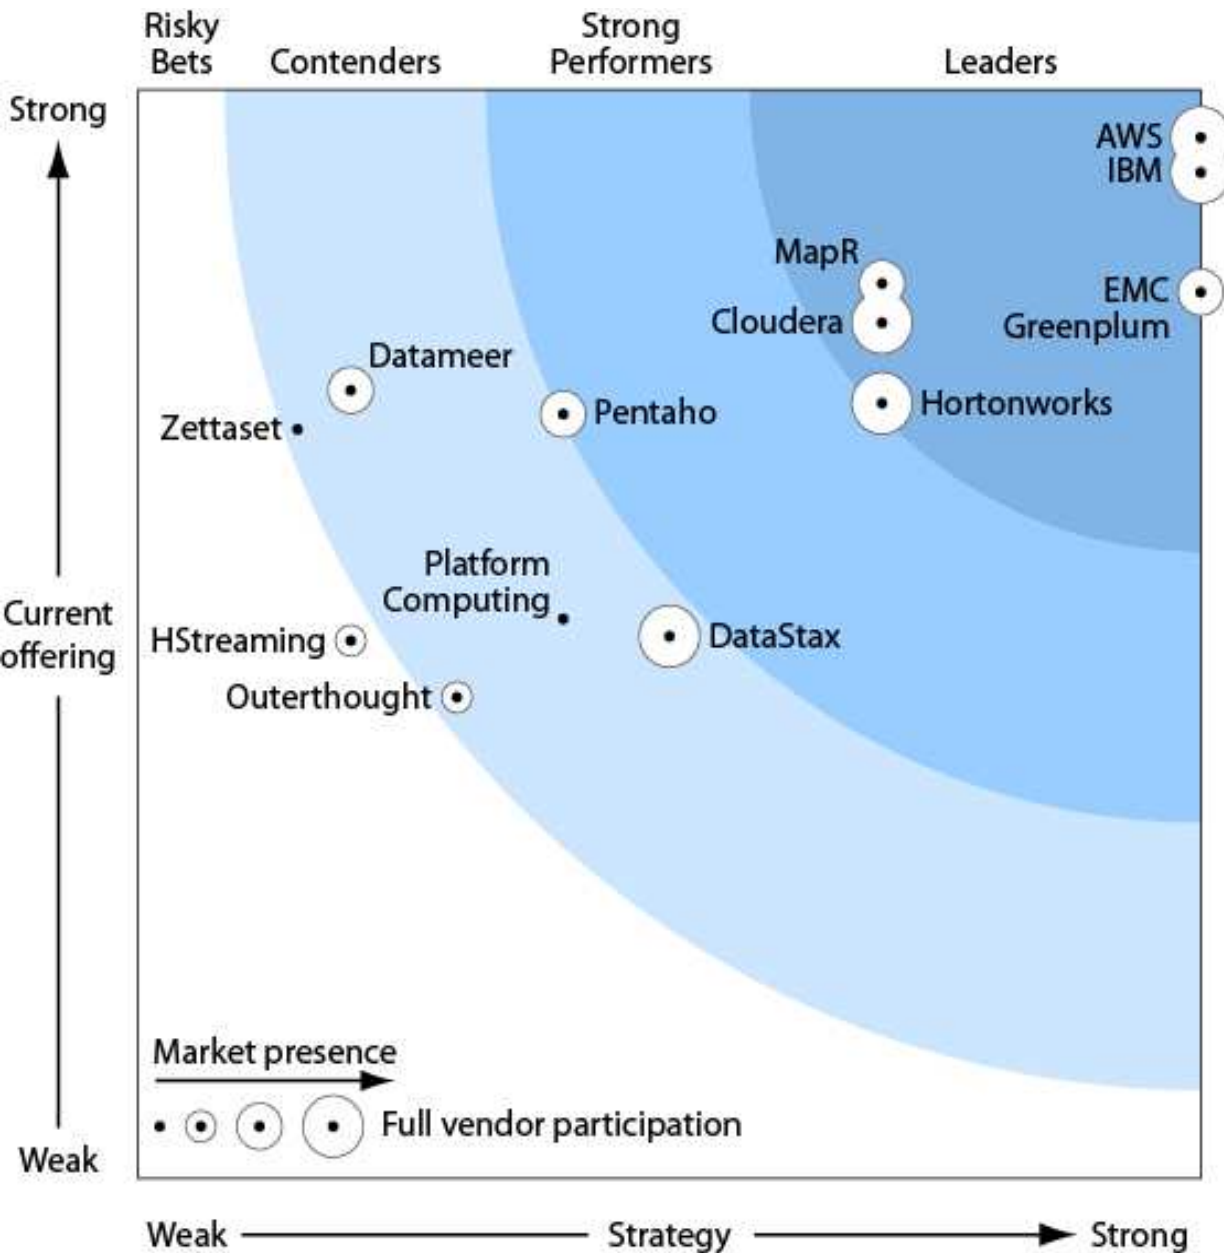
# BigInsights Summary



- BigInsights = analytical platform for persistent “Big Data”
  - Based on open source & IBM technologies
  - Managed like a start-up . . . .
  
- 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; simplifies development and maintenance*
  
- IBM advantage
  - Combination of software, hardware, services and advanced research



# First Ever Forrester Wave on Hadoop, 2012Q1



*“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


## Big Database Vendors Adopt Hadoop

### Microsoft climbs onto Hadoop bandwagon

Joint Microsoft-Hortonworks effort aims to deliver a Hadoop distribution for Windows Server and Windows Azure

By Jaikumar Vijayan

October 12, 2011 03:40 PM ET

 Add a comment

 Like 6

 +1 0



Computerworld - Microsoft is the latest of the world's top IT vendors to climb aboard the Hadoop 'big data' bandwagon.

### Oracle Jumps on Hadoop Bandwagon with Cloudera

By Jennifer LeClaire

January 10, 2012 12:10PM

Follow NewsFactor on [twitter!](#)

 SHARE 

Charles King, principal analyst at Pund-IT, said the Oracle-Cloudera partnership is a win-win. As he sees it, Oracle is coming late to the Hadoop-based Big Data party with Cloudera. He points to EMC's Greenplumb acquisition and IBM's Netezza acquisition, as well as HP's Vertica acquisition in this space as evidence.

### Teradata taps Hortonworks to improve Hadoop story

By Derrick Harris | Feb. 21, 2012, 6:33am PT | No Comments

# IBM Watson





# InfoSphere Streams - Streaming Analytics for Big Data

- 10** Built to analyze data in motion
  - Multiple concurrent input streams
  - Massive scalability
  
- 10** Process and analyze a variety of data
  - Structured, unstructured content, video, audio
  - Advanced analytic operators
  
- 10** Enables Adaptive Real-Time Analytics
  - With Data Warehousing
  - With Hadoop Systems






# Data In Motion – Improving What They Already Have

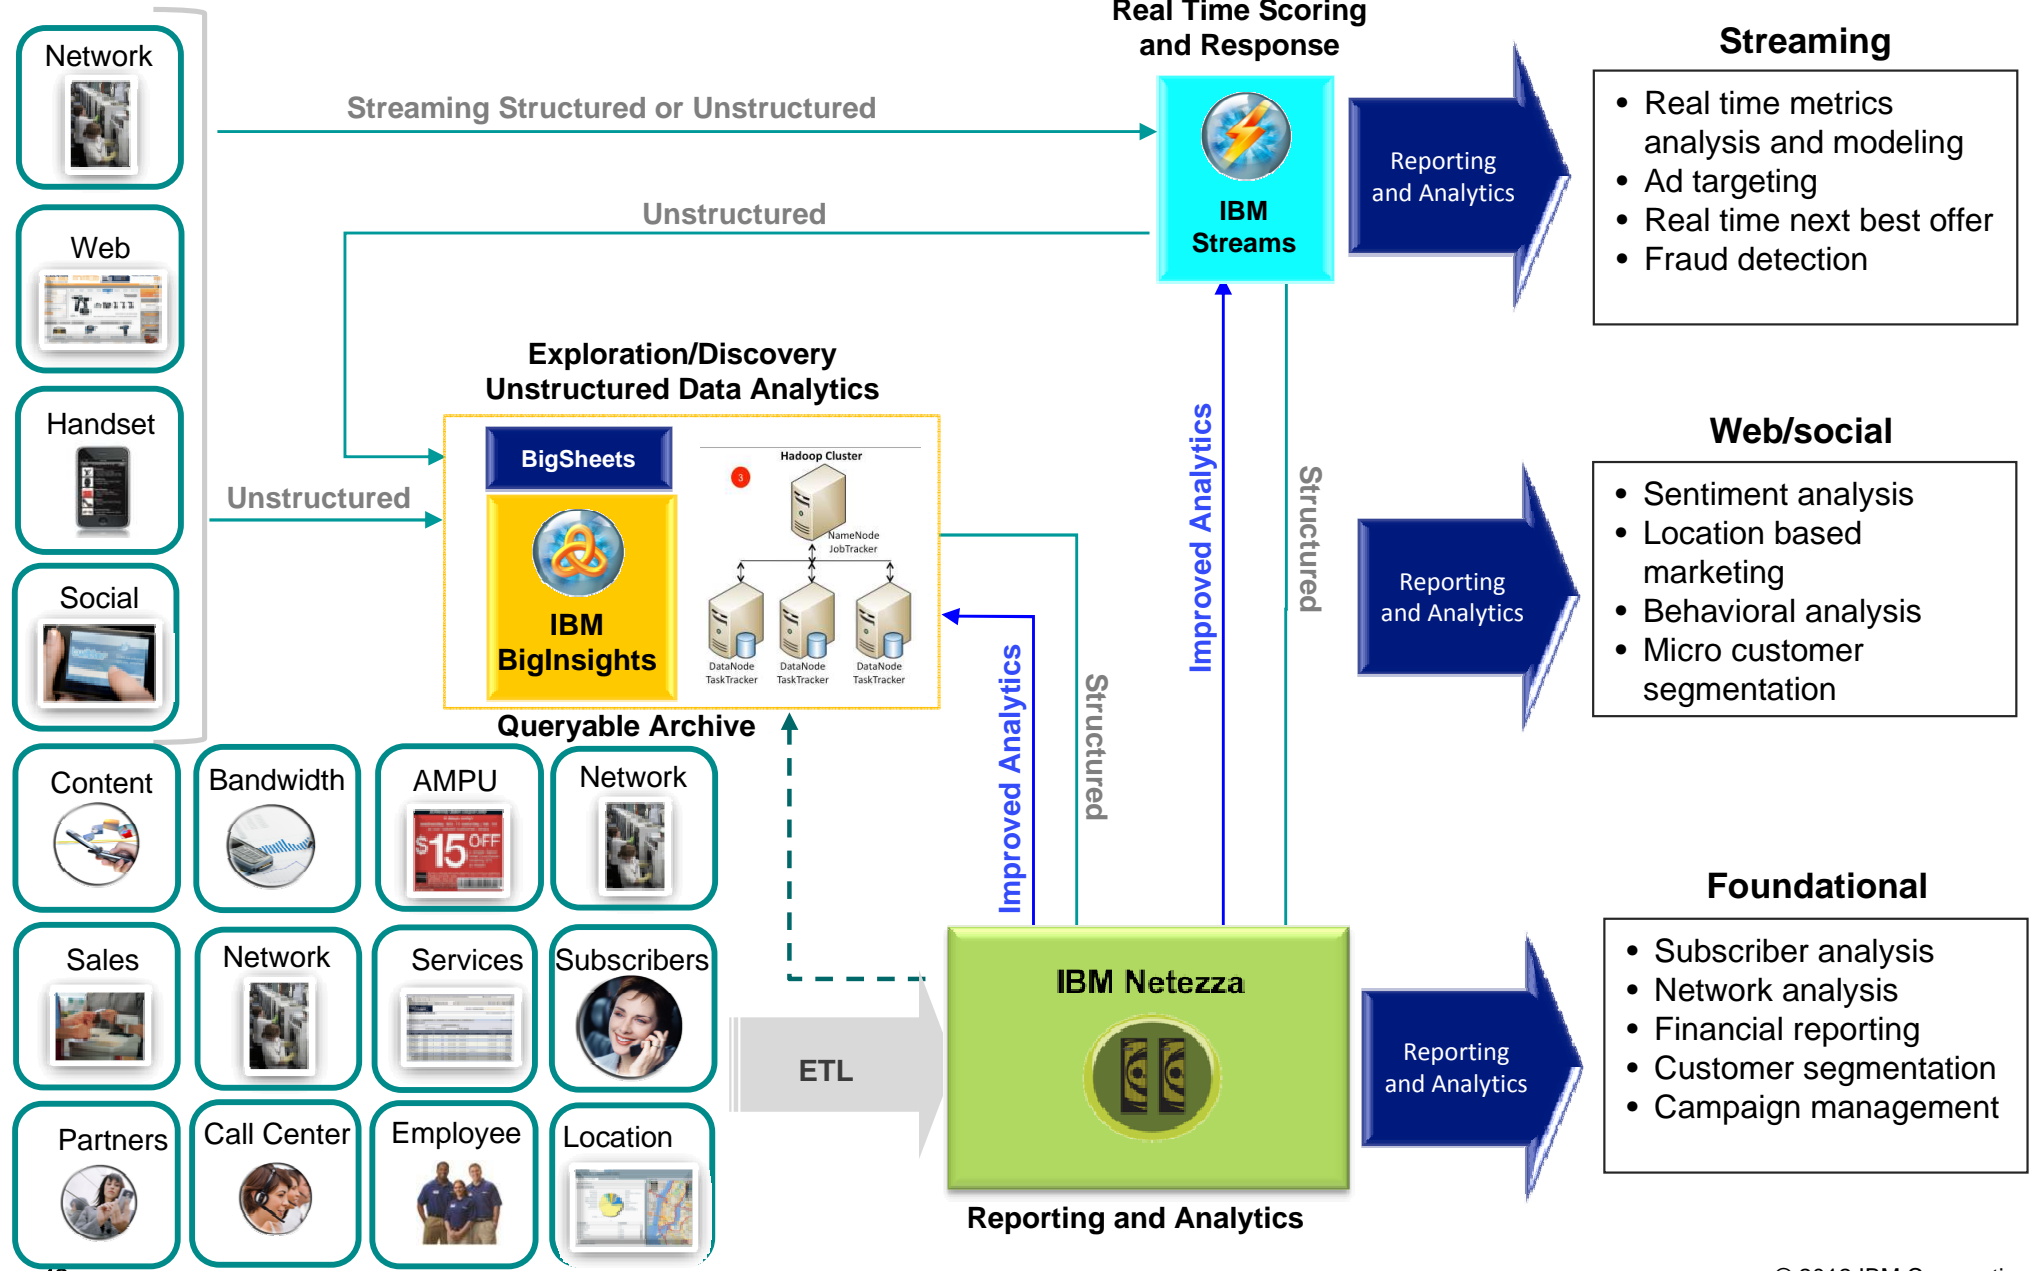
Uses InfoSphere Streams and IBM Netezza  
 Significant Benefits:




- Over 90% reduction in time to merge/load call record data
- Over 90% reduction in storage
- Increased network quality,
- improved customer satisfaction,
- reduced churn



# IBM Big Data integrates Streams, Hadoop and Netezza in an integrated platform



For more information:  
[ibm.com/bigdata](http://ibm.com/bigdata)



#ibmbigdata