THINK BIG

James Lynch
VP & GM Financial Services
Big Data

- Smarter Analytics
- Big Data
- Netezza
- Smart Consolidation

Imagine a world where:

Your organization can run its analytic processes ***wherever*** its data lies: at rest or in motion…structured or unstructured.

regardless of the shape, size, or the speed at which it flows through your organization.

Studies show that those organizations with more sophisticated analytics capabilities outperform their competition

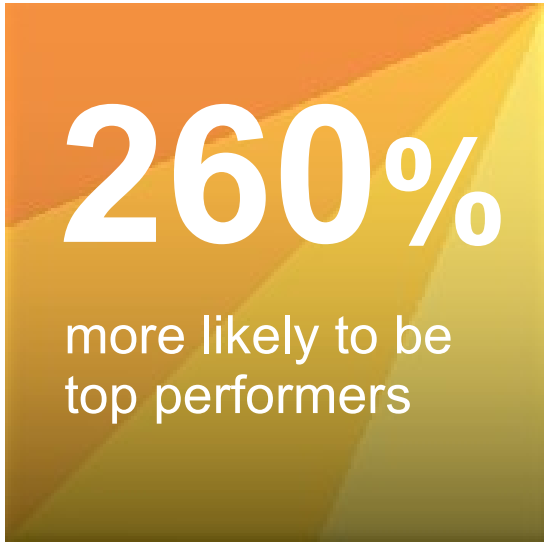

Source: The New Intelligent Enterprise, a joint *MIT Sloan Management Review* and IBM Institute of Business Value analytics research partnership.

Better use of information is required for differentiation

4 in 5 business leaders see **information** as a source of **competitive advantage**

3 in 4 business leaders say more **predictive** information would drive **better decisions**

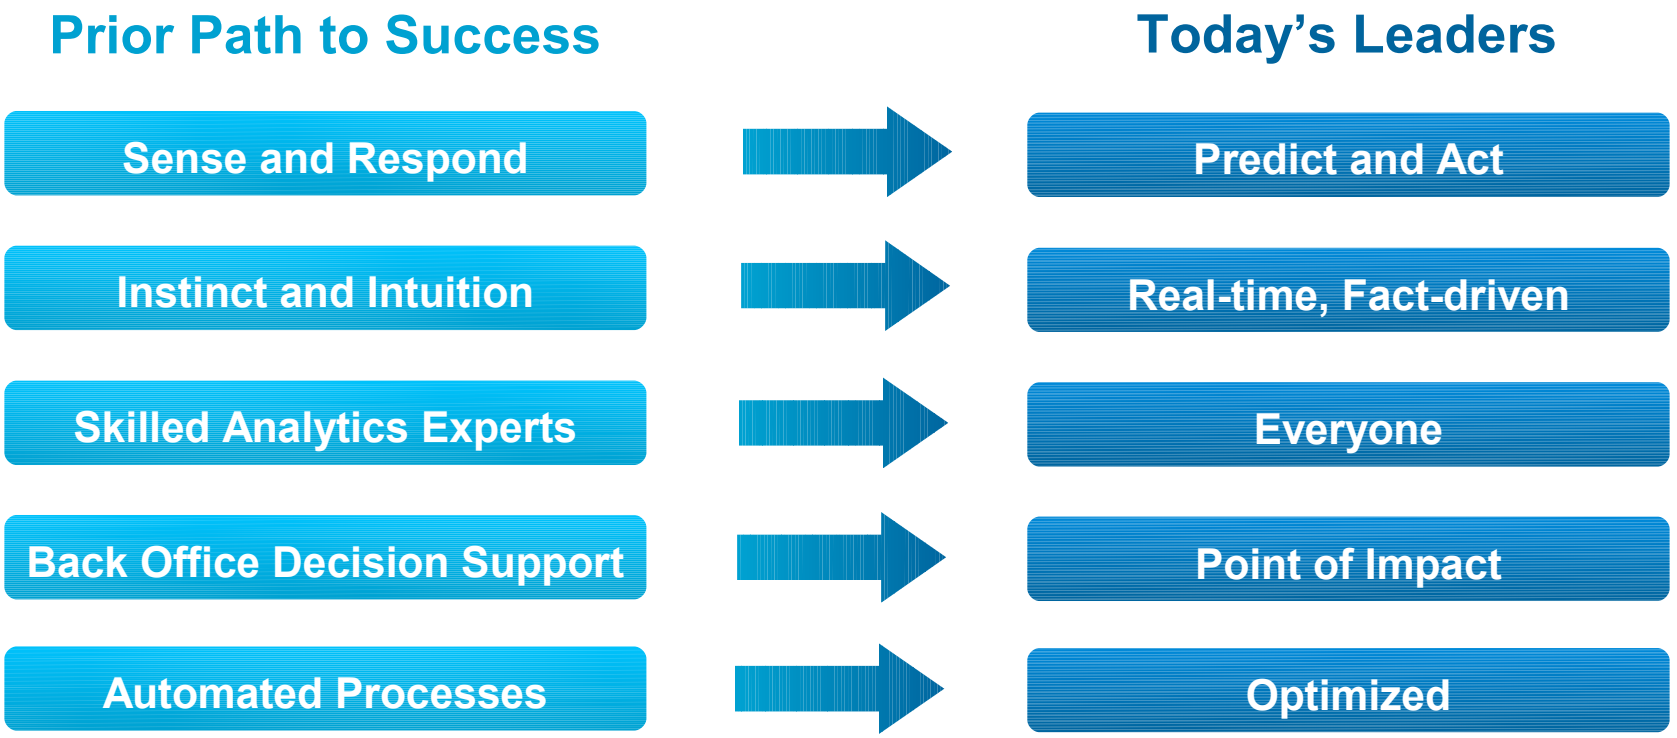

What is “BIG DATA”?

Large volumes

All kinds of data – structured and unstructured

Valuable insight, but difficult to extract

Often extremely time sensitive

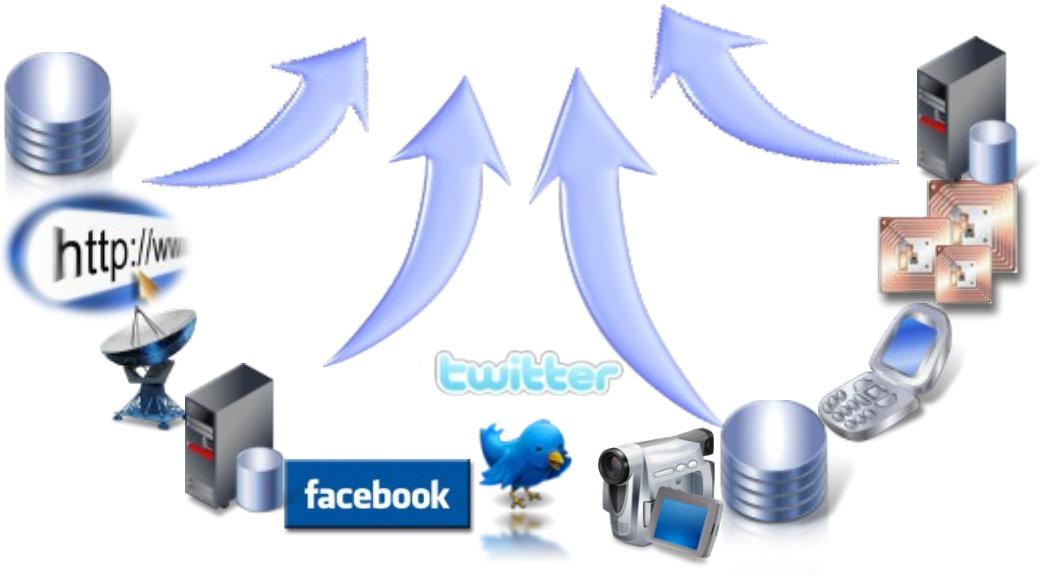

What does a Big Data platform do?

*Extract insights from a large **volume** of data, including a wide **variety** of types, with high **velocity***

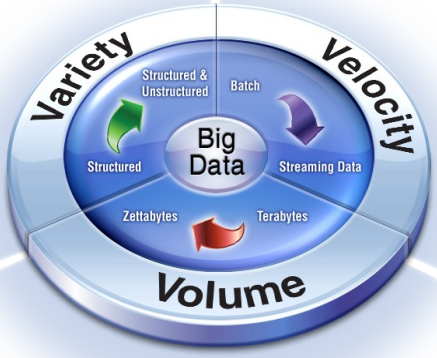

Analyze a Variety of Information

Novel analytics on a broad set of mixed information that could not be analyzed before

Multiple relational & non-relational data types and schemas

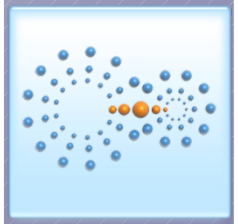

Analyze Information in Motion

Streaming data analysis

Large volume data bursts & ad-hoc analysis

Analyze Extreme Volumes of Information

Cost-efficiently process and analyze petabytes of information

Manage & analyze high volumes of structured, relational data

Discover & Experiment

Ad-hoc analytics, data discovery & experimentation

Manage & Plan

Enforce data structure, integrity and control to ensure consistency for repeatable queries

IBM's Big Data Platform

More than Hadoop

InfoSphere BigInsights

Hadoop-based low latency analytics for variety and volume

Hadoop

InfoSphere Information Server

High volume data integration and transformation

Information Integration

Stream Computing

InfoSphere Streams

Low Latency Analytics for streaming data

MPP Data Warehouse

IBM InfoSphere Warehouse

Large volume structured data analytics

IBM Netezza High Capacity Appliance

Queryable Archive Structured Data

IBM Netezza 1000

BI+Ad Hoc Analytics Structured Data

IBM Smart Analytics System

Operational Analytics on Structured Data

IBM Informix Timeseries

Time-structured analytics

The Forrester Wave™: Enterprise Hadoop Solutions, Q1 2012

Forrester states that, "IBM has the deepest Hadoop platform and application portfolio."

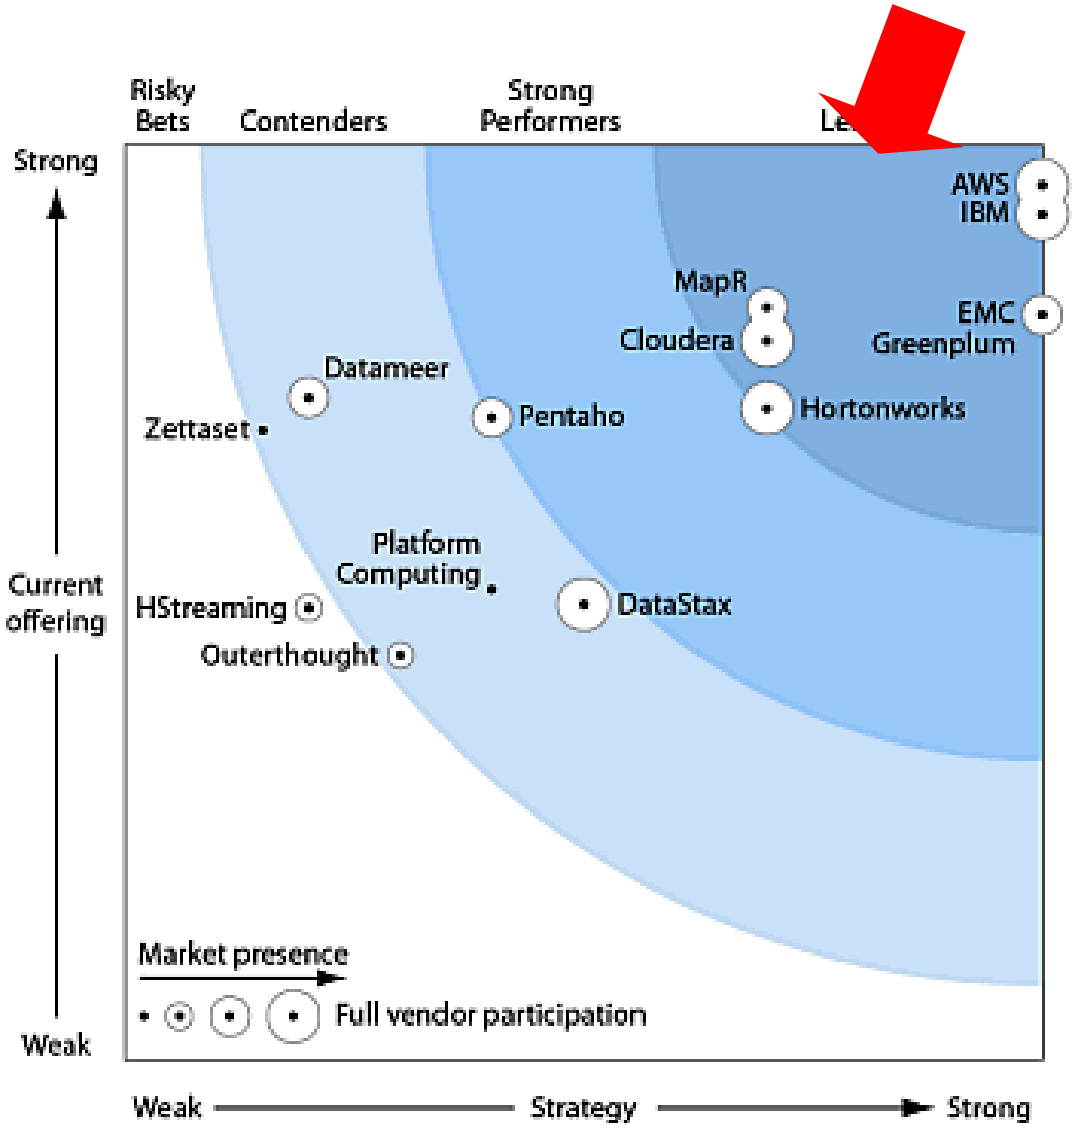

This translates into a few key focus areas…

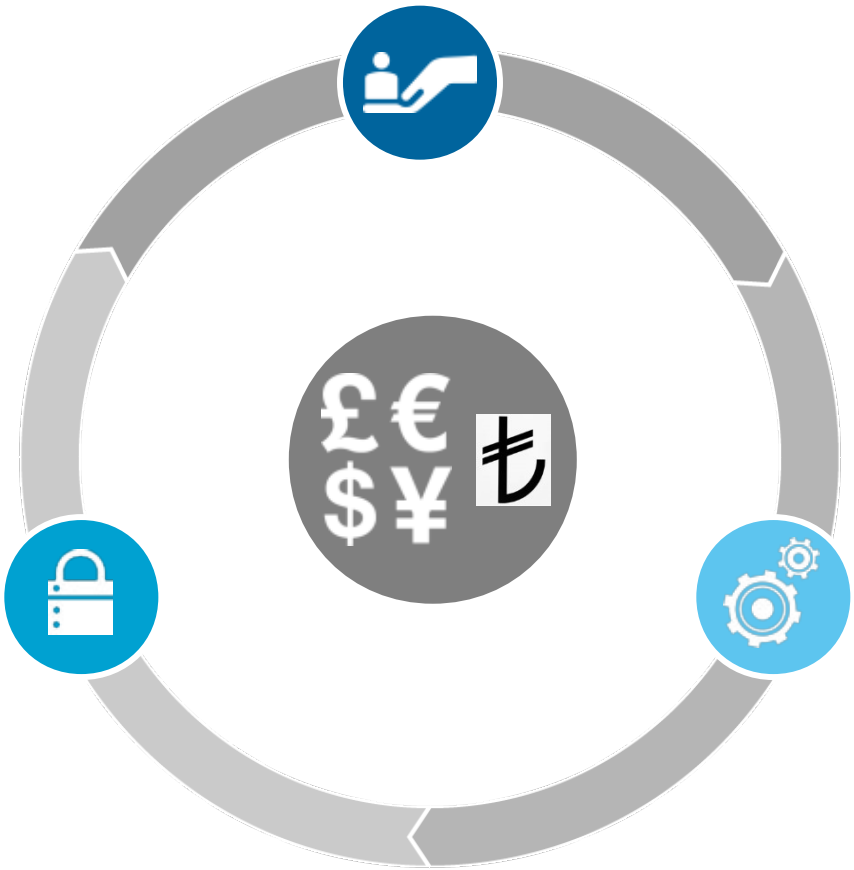

Providing a Path to Big Data…

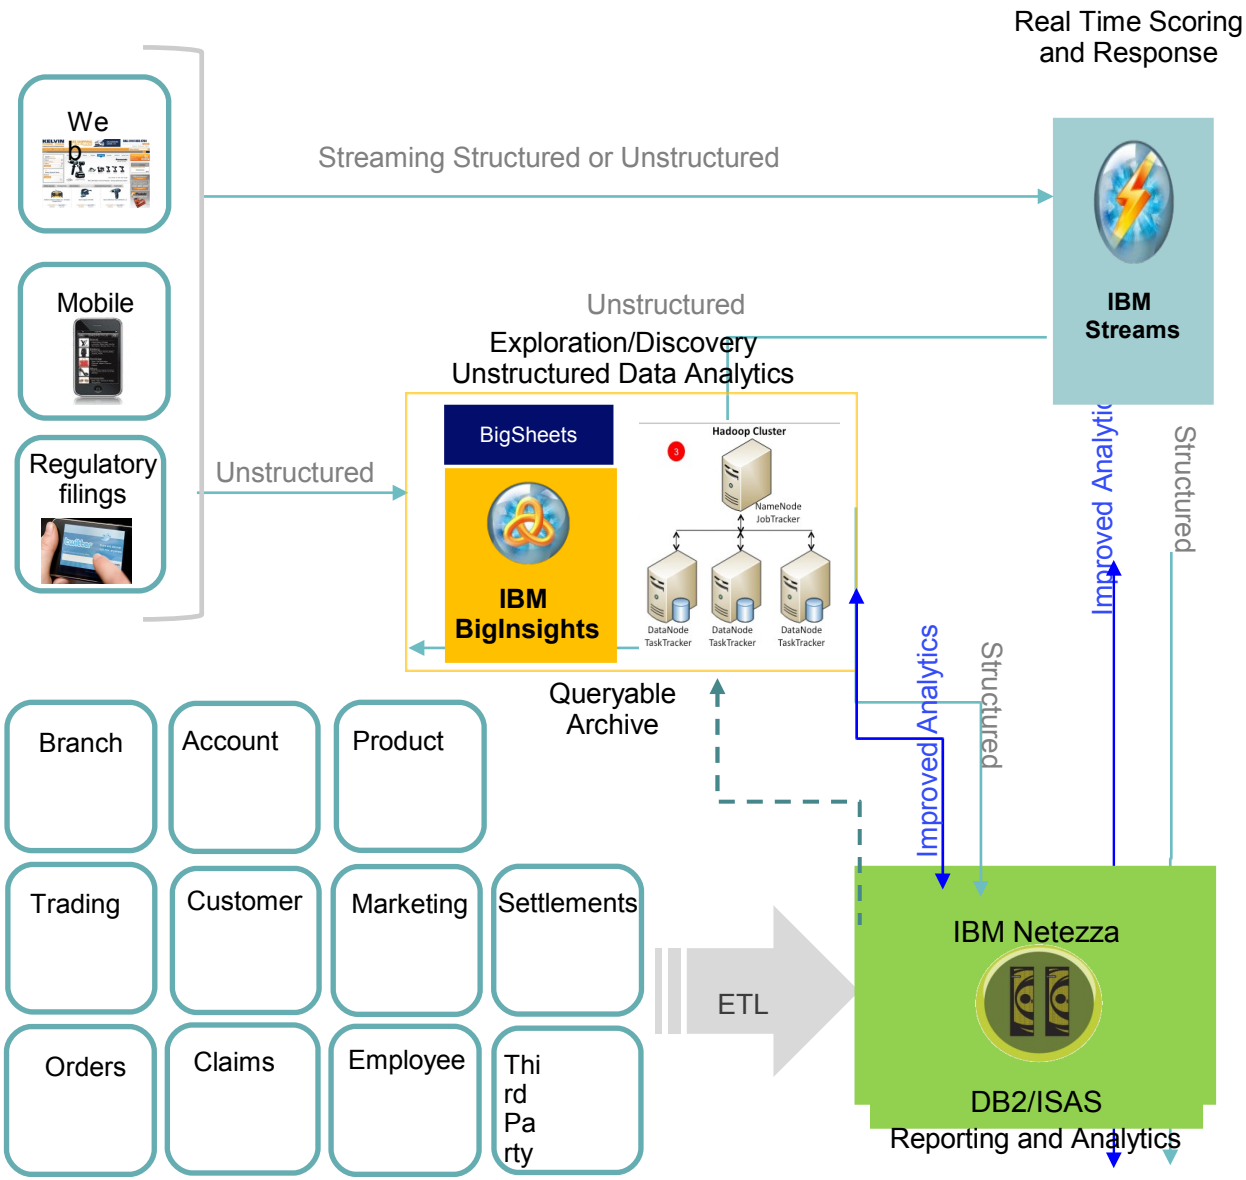

Streaming

- Real time fraud management
- Intra-day risk/liquidity analysis
- Real time trading analysis
- Real time portfolio analysis

Web/Social/Text/Voice

- Sentiment analysis
- Social CRM
- Call center analysis
- Enhanced risk Metrics
- Influencer network
- Employee sentiment

Foundational

- Capital adequacy
- Risk
- Fraud
- Retention
- Segmentation
- Campaign Management
- Surveillance
- Compliance

Big Data brings **Intensified** opportunities and challenges…

Fluid changes to data sets and analytic requirements, no steady state…

Sound familiar?

This **magnifies** the importance of Netezza's architectural benefits

IBM Netezza Value Proposition

The Simple Appliance for Serious Analytics

- Leader in data warehouse & analytics appliances
- 600+ customers, 1500+ systems
- Added momentum since IBM acquisition

- **Speed**: 10-100x faster than traditional systems
- **Simplicity**: Minimal administration and tuning
- **Scalability**: Peta-scale user data capacity
- **Smart**: High-performance advanced analytics

SmarterAnalytics

IBM

Digital Media

ACXIOM
datalogix
epsilon
Kelley Blue Book THE TRUSTED RESOURCE kbb.com
MERKLE A Database Marketing Agency
TRIDENTMARKETING
appnexus
edmunds.com
exelate
MediaMath
MicroAd
m!c MEDIA INNOVATION GROUP
[x+1]

Financial Services

AARP
GUY CARPENTER
ICE
insure the box
NYSE Euronext

Health & Life Sciences

MASSACHUSETTS
Capital BlueCross
HARVARD MEDICAL SCHOOL
株式会社日本医療データセンター Japan Medical Data Center
University at Buffalo The State University of New York
UWHealth

Retail / Consumer Products

APERITY
CATALINA MARKETING
ESSELUNGA
BARNES&NOBLE BOOKSELLERS
carter's babies and kids
NOOS
SAPPORO
Yum!

Telecom

AIRWAVE
colt smarter / faster / further
MACH YOUR MOBILE VISION. REALIZED.
TEOCO
Virgin media
THE Carphone Warehouse
iBasis a KPN affiliate
QUALCOMM
T-Mobile
XO Communications

Other

Con-way FREIGHT
FOXWOODS RESORT + CASINO
NIC Nippon Information and Communication
OLYMPUS Your Vision, Our Future

Netezza – Simplicity over Complexity

Relief

NO dbspace/tablespace sizing and configuration

NO redo/physical log sizing and configuration

NO journaling/logical log sizing and configuration

NO page/block sizing and configuration for tables

NO extent sizing and configuration for tables

NO temp space allocation and monitoring

NO RAID level decisions for dbspaces

NO logical volume creations of files

NO integration of OS kernel recommendations

NO maintenance of OS recommended patch levels

NO JAD sessions to configure host/network/storage

Simple partitioning strategies: HASH or ROUND ROBIN

Results

Data Marts On Demand

"The Netezza system requires 90% fewer DBA's than the Oracle system" – *Top five Investment Management firm*

"Data loads on the incumbent system took six hrs. on average; on Netezza they take less than 30 minutes" – *Large consumer credit provider*

"Time to deliver new data sources into production reduced by 30-40% with Netezza" – *Top five bank*

"4 to 1 reduction in resource requirements" – *Top five bank*

"A 17 step process has been reduced to 6 steps with Netezza" – *Top securities regulator*

"Netezza eliminated 30-40% of DW development costs attributed to tuning, enabling 75% reduction in support costs"

Netezza Approach

Accelerated POC

Client Terms **Client Data** **Client Site**

Exemplifies NETEZZA's Simplicity, Power and Ease

TOP 5 Global Bank IBM Netezza Implementation

Over 2 Dozen Applications Across 13 LOB's thx

Deployments

2011
- Corporate Treasury
- Quantitative Risk
- ZINC (Front Office Risk + MBS)
- Market Risk
- eSMART (Mater Data Mgt Hub) - **InfoSphere, DB2 and Netezza solutions**

2010
- Associate Monitoring
- CashPro (customer portal)
- CRM (Profit Mart)
- Merchant Services
- Enterprise Credit Risk

2009
- CSAR (client sales reporting)
- Treasury and Payments (GPS MIS)
- Card Services
- Global Wealth – Distro AUM

- Teradata / Oracle / Sybase replacements
- BI = Cognos, MicroStragegy, OBIEE
- ETL = DataStage, Informatica, Talend

Improvement Metrics

- No indexes, partitions, optimizer hints or system parameters
- Queries reduced from hours or minutes to minutes or seconds; impossible queries enabled
- Large volumes of data load in minutes

Results

- **3-5x reduction** in raw disk space
- Reduced DBA man hours; resources perform higher value activities
- Near real time response on years of data
- **thousands more queries** can run per day
- Mixed workloads with 100+ concurrent users run 20x faster

Strategic Initiatives

Big Data	Big Insights, Streams and Netezza appliances to support analysis of increasing volume, variety and velocity of data
Smart Consolidation	EDW augmentation, deploying Netezza appliances to provide more cost effective and better performing analytic applications

Tactical Solutions

IDAA	Netezza appliance that attaches directly to System z enabling clients to query and analyze mainframe data without utilizing System z resources
SAS Optimization	Netezza appliance to accelerate performance of SAS analytics and reduce costs

IDAA Next Generation Accelerator

Netezza appliance that attaches directly to System z enabling clients to query and analyze mainframe data without utilizing System z resources

Challenges

- Reporting on or analyzing System z data restricted due to operational applications
- Accessing System z data drives MIPs
- Operational reporting requires set up and management of separate data marts

Value

- Enable new reporting & analytics without impacting System z performance or MIPs
- Reduce MIPs cost
  - Avoiding 400 MIPs = cost of IDAA Netezza 1000-3
  - Avoiding 800 MIPs = cost of IDAA Netezza 1000-6
  - Avoiding 1600 MIPs = cost of IDAA Netezza 1000-12
- Reduce data mart infrastructure costs and data integration complexity

Client Examples

Query	Total Rows Reviewed	Total Rows Returned	DB2 Only Hours	DB2 Only Sec(s)	DB2 with IDAA Hours	DB2 with IDAA Sec(s)	Times Faster
Query 1	2,813,571	853,320	2:39	9,540	0.0	5	1,908
Query 2	2,813,571	585,780	2:16	8,220	0.0	5	1,644
Query 3	8,260,214	274	1:16	4,560	0.0	6	760
Query 4	2,813,571	601,197	1:08	4,080	0.0	5	816
Query 5	3,422,765	508	0:57	4,080	0.0	70	58
Query 6	4,290,648	165	0:53	3,180	0.0	6	530
Query 7	361,521	58,236	0:51	3,120	0.0	4	780
Query 8	3,425.29	724	0:44	2,640	0.0	2	1,320
Query 9	4,130,107	137	0:42	2,520	0.1	193	13

Queries run faster

- Save CPU resources
- People time
- Business opportunities

Actual customer results, October 2011

Aetna™

DB2 Analytics Accelerator: "we had this up and running in days with queries that ran over 1000 times faster"

Swiss Re

DB2 Analytics Accelerator: "we expect ROI in less than 4 months"

Accelerating decisions to the speed of business

SAS Optimization

Netezza appliance to accelerate performance of SAS analytics and reduce costs

Challenges

- Long running SAS jobs (>1 hour)
- Large, unmanageable SAN stores for SAS data sets
- Pulling data across network to SAS server to store and process data

Value

- Reduce time for analytics
- Enable additional, value-add analytics
- Reduce SAN storage costs
- Reduce cost of SAS servers and licenses
- Reduce need for SAS Grid (triple price of normal SAS license)

Client Examples

Client	Results
Bank of America	SAS Quantitative Statistics Department saw a reduction of 20-hour SAS jobs down to 20 minutes. Through user education, they decided to implement IBM/Netezza as the back end data store for their SAS ETL and Statistical analysis
citibank	Took Basel II process from 40 hours to ½ hour
ZURICH	Took 62 hour monthly close process down to 2½ hours

- Processing is much faster due to removing the need to pull massive amounts of data to the SAS application server, using IBM/Netezza's In-Database functionality
- IBM Netezza running with SAS has a proven record of accomplishment demonstrating significant relief against these challenges

Why IBM

- ✓More than $14B in acquisitions Since 2005
- ✓More than 10,000 technical professionals
- ✓More than 7,500 dedicated consultants
- ✓Largest math department in private industry
- ✓More than 27,000 Business Partner certifications

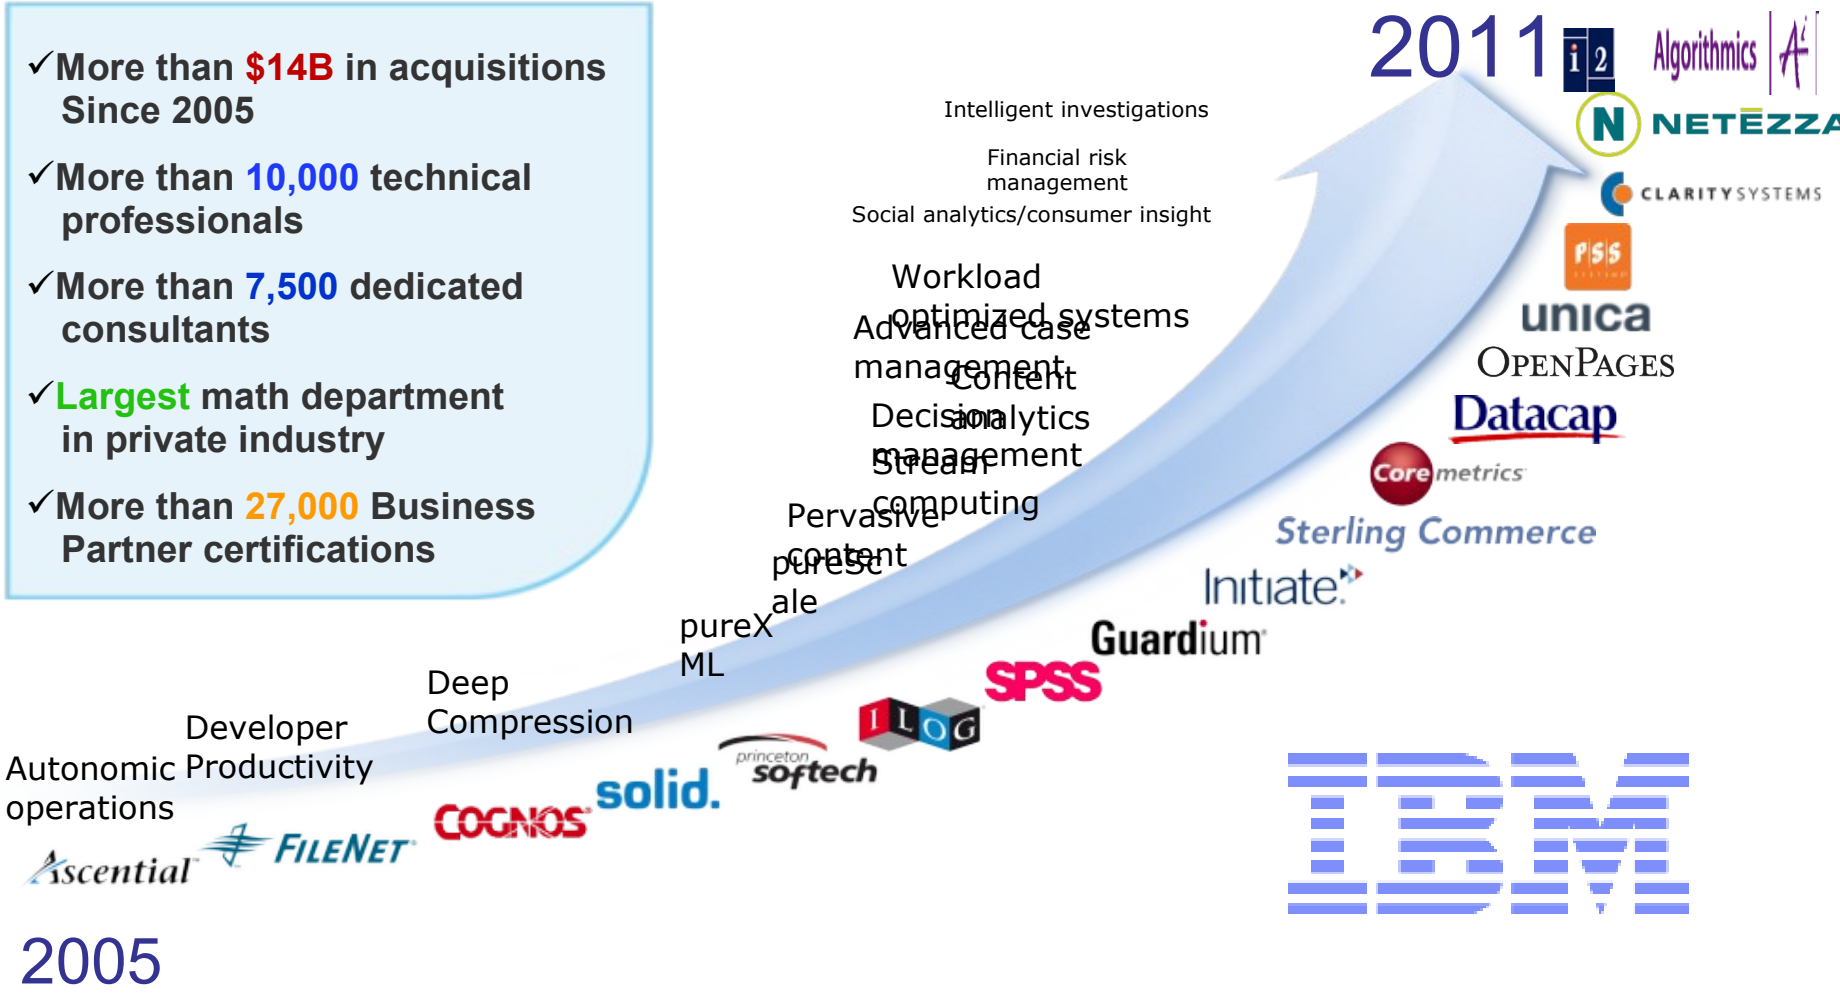

THINK BIG

THANK YOU!