

IBM Infrastructure for Smarter Analytics

Smarter**Analytics** Live

Quando la conoscenza trasforma il business

Milano, 5 giugno 2012 (ore: 9.00-17.30)
Sede Gruppo 24 ORE, Via Monte Rosa 91



Analytics has evolved from business initiative to business imperative

Analytically sophisticated companies outperform their competition

2.2x

more likely to
outperform
industry peers

260%

more likely to be
top performers

Analytics is expanding from enterprise data to big data

Volume

12 terabytes
of Tweets create daily

Velocity

5 million
trade events per second

Variety

100's video feeds
from surveillance cameras

350 billion
meter readings per annum

500 million
call detail records per day

80% data growth
are images, video, documents...

Organizations drive transformation by starting with **one of these four high-value initiatives**

Examples:

1

Grow, retain and satisfy customers



- Churn management
- Social media sentiment analysis
- Propensity to buy/Next best action

2

Increase operational efficiency



- Predictive maintenance
- Supply chain optimization
- Claims optimization

3

Transform financial processes



- Rolling plan, forecast and budget
- Financial close process automation
- Real-time dashboards

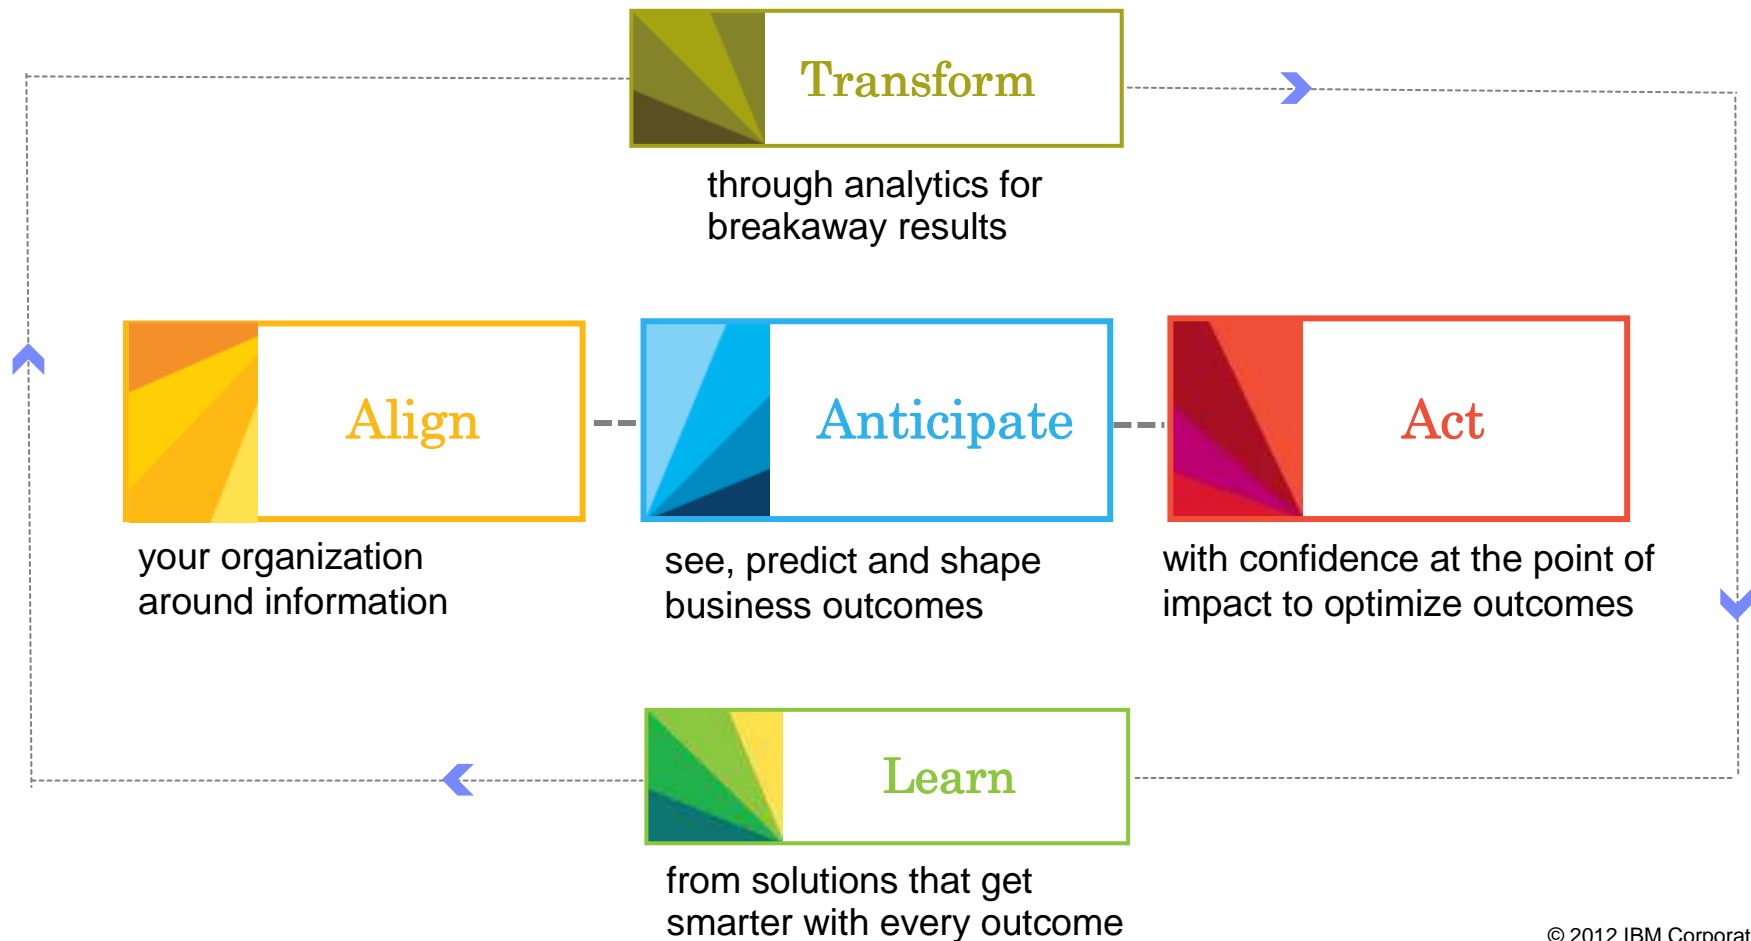
4

Manage risk, fraud & regulatory compliance



- Operational and financial risk visibility
- Policy and compliance simplification
- Real-time Fraud identification

IBM Smarter Analytics is a **holistic approach** that turns information into insight and insight into business outcomes.



The Smarter Analytics approach embodies several key **IT Infrastructure tenets**

Align
 Deploy an information and big data strategy that flows from your business architecture.

Anticipate
 Leveraging and Integrating Business Analytics to deliver actionable insights

Act
 Embed analytics into your processes and empower a culture of data-driven decision making



Creating a **scalable**, trusted information and systems foundation that improves IT economics and optimizes analytic workload performance using all available data and information.

Optimizing high performance parallel technologies to support complex decision making, spotting trends and anomalies, predicting business outcomes.


Deploying analytics throughout the organization, it's customers and suppliers using **resilient** architectures either on premise or in the cloud.

The trend in the system landscape

Scale In

Analytics and Big Data Workloads, MultiCore Chip, integrated Systems

Process-Centric era



- Business processes automation
- Focused on bottom-line improvement through SG&A reduction
- Typically long business cycle
- Terabytes of largely structured data

Scale Out

Web Application

Information-Centric era

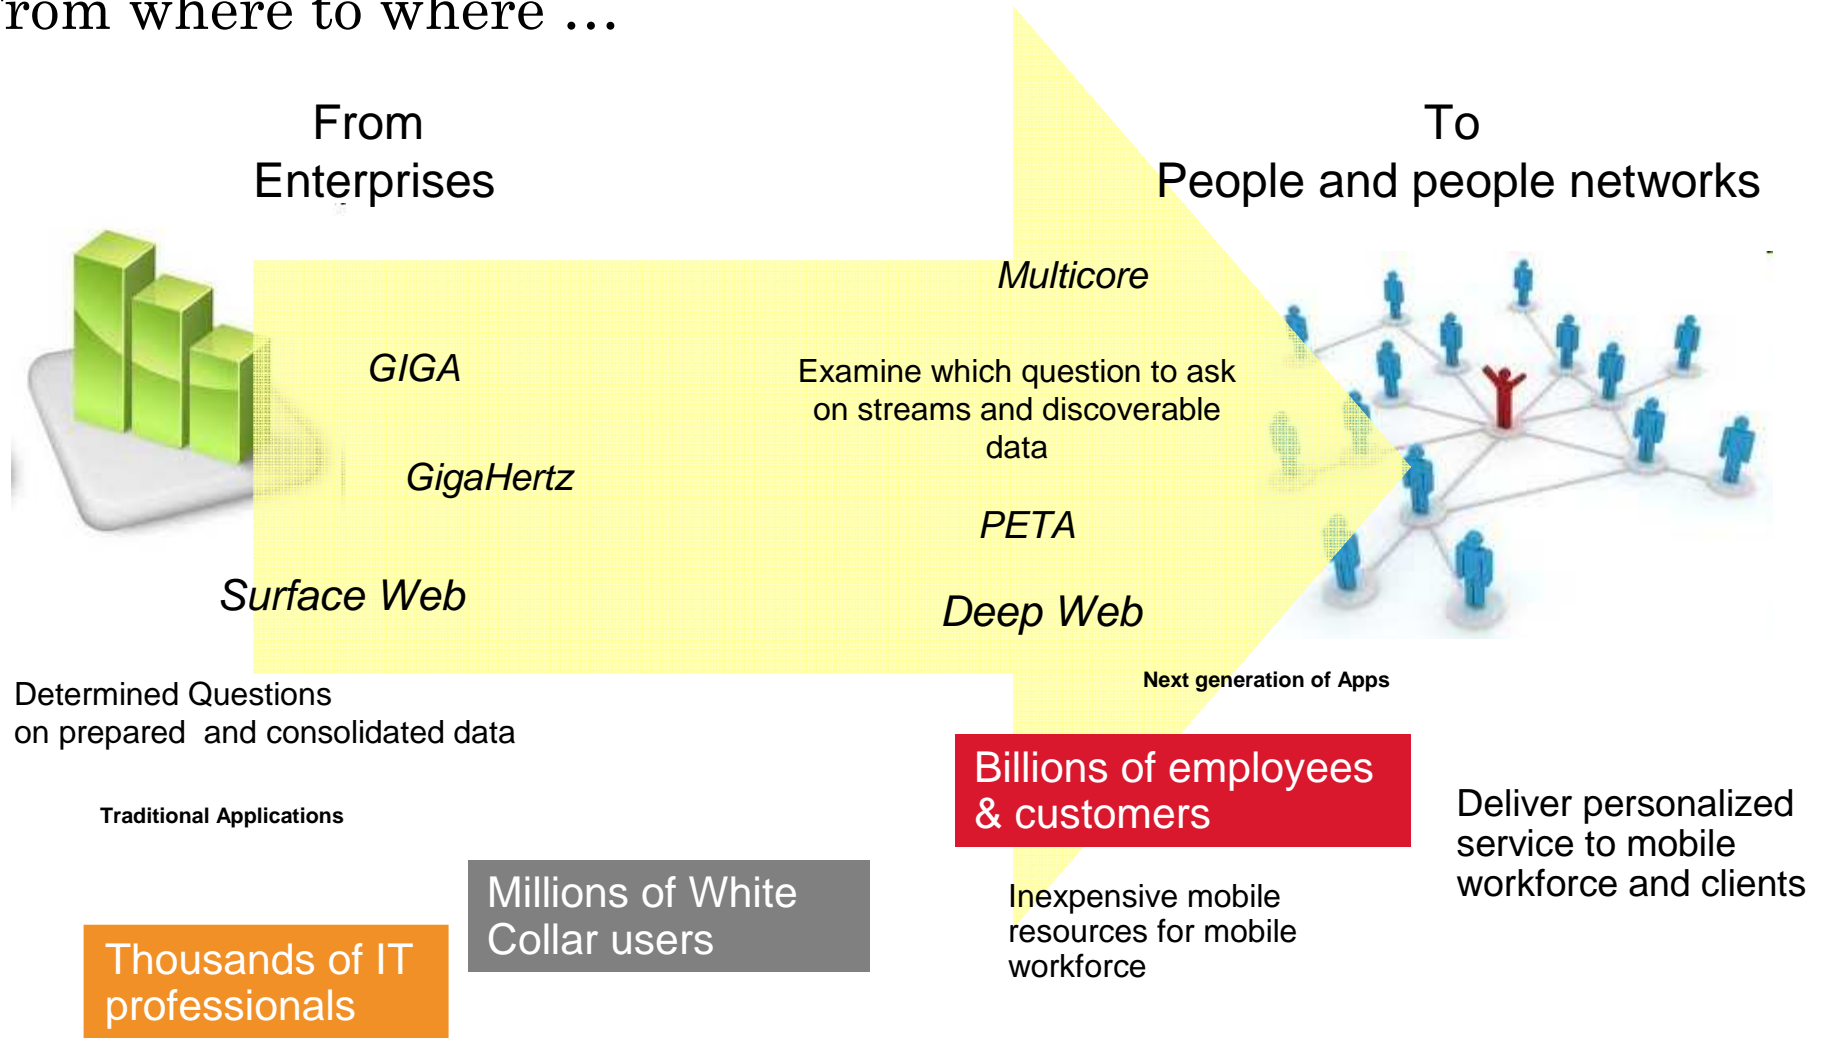


- Real-time pattern based action
- Focused on top-line growth through revenue acquisition
- Reactive for shorter product cycles
- Zettabytes of largely unstructured data

Scale Up

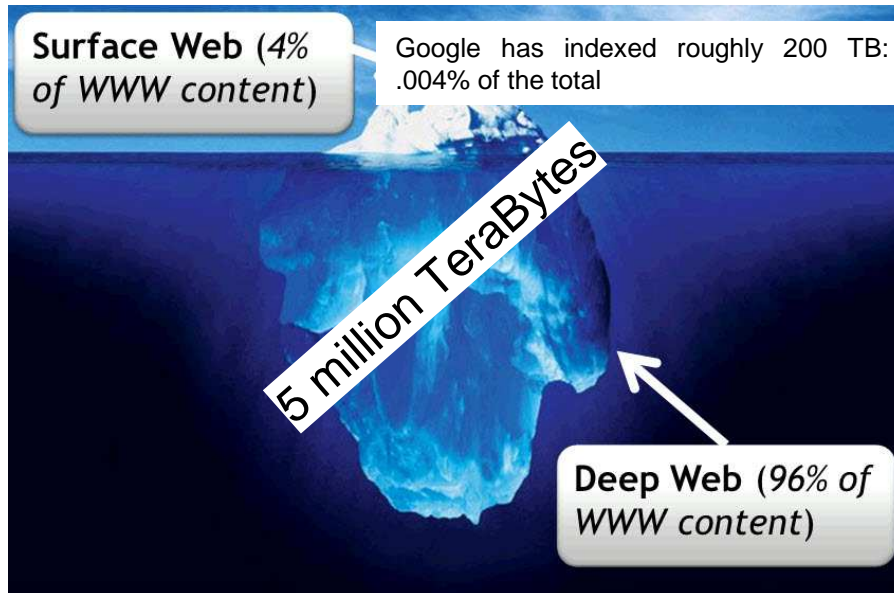
Large OLTP Database and ERP.

From where to where ...

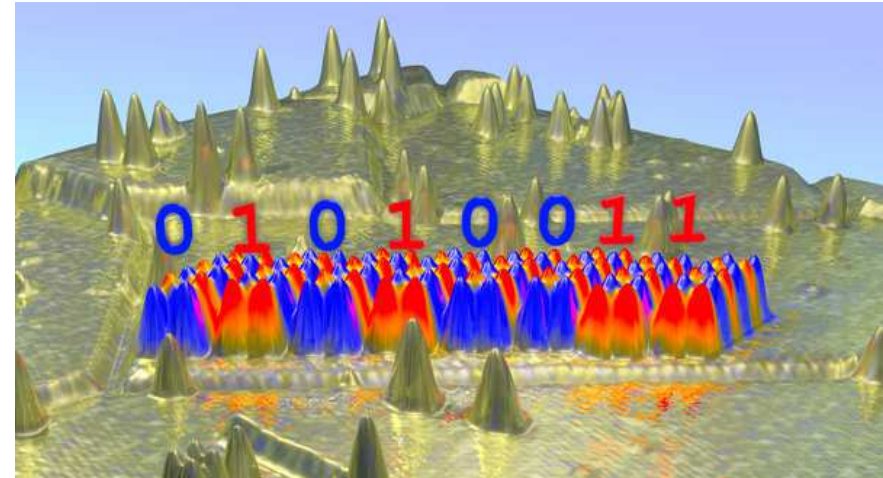


New advances in technology enable new questions

- New technology capabilities
- New Source of Data



InfoSphere BigNights
InfoSphere Streams

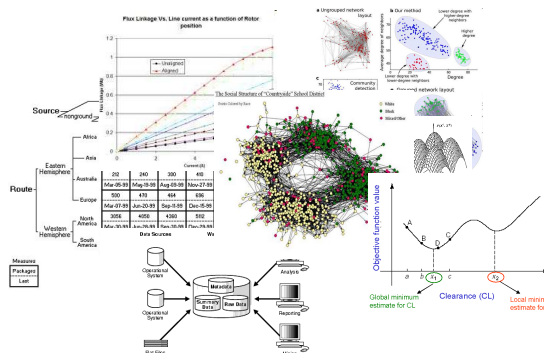


January 12, 2012 SAN JOSE, Calif. — Researchers at I.B.M. have stored and retrieved digital 1s and 0s from an array of just 12 atoms, pushing the boundaries of the magnetic storage of information to the edge of what is possible. Until now, the most advanced magnetic storage systems have needed about one million atoms to store a digital 1 or 0.

The group at I.B.M.'s Almaden Research Center here, led by Andreas Heinrich, has now created the smallest possible unit of magnetic storage by painstakingly arranging two rows of six iron atoms on a surface of copper nitride.

Scalable: Different types of analytics require a **scalable** IT infrastructure

Different types of analytics (OLTP, Data warehouse, Streaming Data, OLAP, Operational Analytics, Ad-hoc reporting, time series, deep analytics)...



...need to access data differently and require compute *and* storage resources that are distinctly different and often highly scalable.

Optimized: Enabling analytics anytime and anywhere requires an **optimized** system “tuned for the task”

Turning information into insight requires information as it happens and analysis as needed....



...thru flexible systems designed to access the latest information regardless of type or location; allocating the right resource at the right time; providing new resources as needed for optimal analysis on-demand.

Resilient: Empowering a culture of data-driven decision making requires a **resilient** IT infrastructure

As a business imperative, Smarter Analytics is about pushing analytics to the edge of the organization and beyond....to employees, customers, and suppliers...



...thru mission critical, reliable systems and design, a responsive infrastructure that can handle large numbers of users, secure systems that work seamlessly.

...in 2011 IBM showed to the world an example on how an IT Infrastructure can be designed to achieve superior outcomes...

A System Designed For Answers:

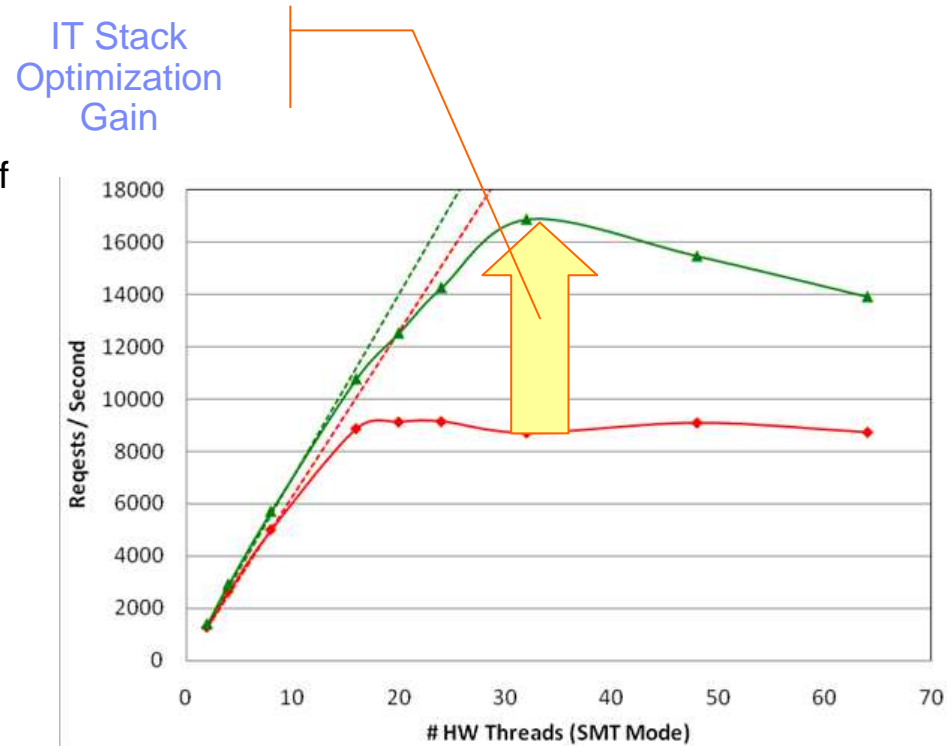
- Designed to answer the questions of today and tomorrow using deep QA Architecture
- Addresses natural language questions with accuracy and confidence in seconds not hours or days
- Processes massive array of information including unstructured “Big Data”
- Built on a cluster of commercially available Power 7 servers, optimized for complex analytic workloads and used by thousands today.



Workload Optimized Systems (WOS)

Workload Optimized Systems are

- An approach to maximize the overall potential of each component of an IT solution related to a specific client's business workload
- Easy to manage and to fit in the current IT environment
- Drive Innovation and new business capabilities
- An IBM 360° approach
 - IBM Systems,
 - IBM Software
 - IBM Research
 - Workload specialists involvement
 - Best Practices



Clients require a range of workload optimized systems

**Client-built
with optimized
components**



Need flexibility to deploy multiple workloads of different types—e.g., data management, messaging, web facing etc.

**Integrated,
Optimized Systems**



Requires moderate flexibility to tune small number of workloads—e.g., information management and analytics


Appliances



Flexibility not required—need high performance at low cost for a specific workload

IBM Systems for Smarter Analytics

IBM Systems is the market leader in support of analytical workloads



IDC Server Workload study 2011

Workloads:

- Data Analysis/Data Mining
- Data Warehousing/Data Mart
- Scientific/Engineering/Industrial R&D



What the Analysts are saying

"...at the Smarter Analytics Leadership Summit, it was clear that the **company's hardware, software, and services** people are all working closely together to build integrated business analytics and optimization solutions tuned to perform extremely well — and designed around achieving the outcomes required by IBM customers (on an industry-by-industry basis)." Clabby Analytics

"When we commenced this exercise we expected to find that there were some areas in which IBM excelled and others in which Oracle did so. We have been surprised to find that that is not the case and that the **IBM Smart Analytics System out-competes Oracle Exadata in almost every area we have examined**" Philip Howard, Bloor Research

"**IBM has the deepest Hadoop platform and application portfolio.** IBM ... has its own Hadoop distribution; an extensive professional services force working on Hadoop projects; extensive R&D programs developing Hadoop technologies ... and software, appliance, and cloud offerings." Source: The Forrester Wave™: Enterprise Hadoop Solutions















Three-year costs for Smart Analytics System 7700 are **43 and 40 percent less than those for Oracle and Teradata systems** respectively. Source: Cost/Benefit Case for Enterprise Warehouse Solutions, International Technology Group

IBM and Analytics at a glance:

- More than **\$14B** in Acquisitions Since 2005
- More than **10,000** Technical Professionals
- More than **7,700** Dedicated Consultants
- **Largest** Math Department in Private Industry
- More than **27,000** Business Partner Certifications

Industry leading portfolio for Smarter Analytics...

Manages the data explosion to spot trends, predict outcomes, and take meaningful actions


Align	Anticipate	Act
 <p>Netezza (x, Storage)</p>  <p>Smart Analytics System (x, Power, z, Storage)</p>  <p>Custom Data Warehouse (x, Power, z, Storage)</p>  <p>DS8000 DS5000 DS3000</p>  <p>XIV SONAS StorWize V7000 SVC</p>	 <p>DB2 Analytics Accelerator (z, x, Storage)</p>  <p>SAP HANA Appliance (x, Storage)</p>  <p>Optimized for IBM systems (x, Power, z, Storage)</p>  <p>Systems for InfoSphere BigInsights (x, Power, Storage)</p>  <p>Optimized for IBM systems (x, Power, z, Storage)</p>	 <p>IBM Watson (Power, Storage)</p>  <p>Optimized for IBM systems (x, Power, z, Storage)</p>  <p>Systems for InfoSphere Streams (x, Power, Storage)</p>  <p>Systems for ISV BAO (x, Power, z, Storage)</p>

- Cost effective Big Data storage
- Integrate and govern data
- Single, trusted information source

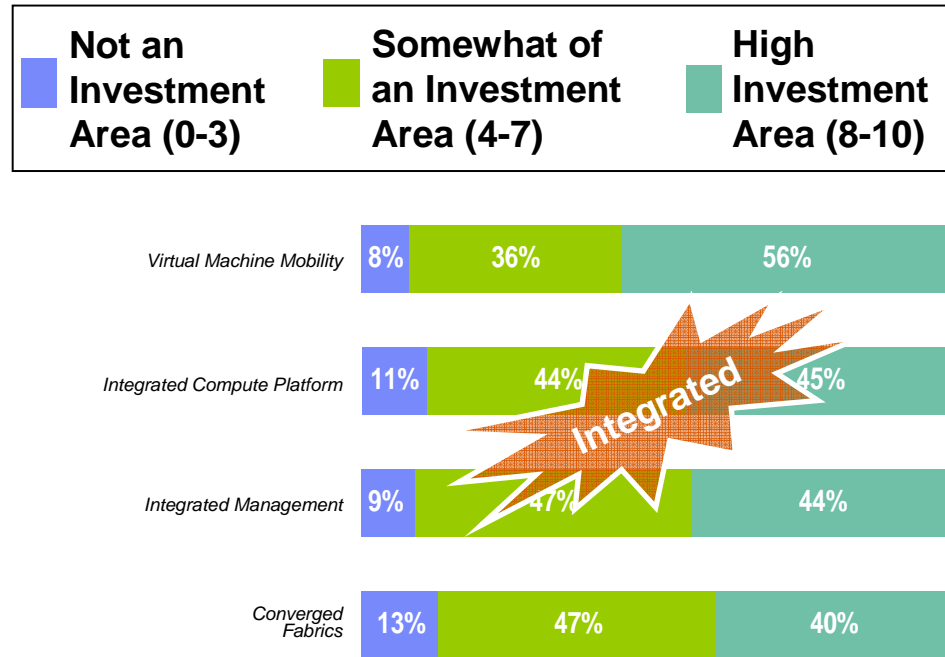
- Address scale and complexity
- Model trends and scenarios
- Predict business outcomes

- Embed analytics across enterprise
- Accelerate real-time decision making
- Systems that reason and learn

Customers want larger integrated systems



Priority Investment Trends in next 18 months



Clients are seeking solutions to the complexities associated with inefficient networking, server sprawl and manual virtualization management.

Integrated system pre-packages server, storage, network, virtualization and management and provides an automated, converged & virtualized solution with fast time to value & simple maintenance

The Time has come for a new breed of Expert Integrated Systems
Announcing the first two members of the IBM PureSystems family

PureFlex

New

Infrastructure System:
*Expert at sensing and
anticipating resource
needs to optimize your
infrastructure*



PureApplication

New

Platform System:
*Expert at optimally
deploying and running
applications for rapid
time-to-value*



Built-in expertise ■ Integration by design ■ Simplified experience

IBM PureSystems... “integration by design”



Optimizes the complete solution stack

- All hardware and software components factory integrated and optimized
- Single point of unified lifecycle management
- Integrated monitoring & maintenance
- Integrated and elastic application and data runtimes
- Application patterns allocate system resources for optimal performance, security and reliability
- Fully virtualized and built for cloud
- Storage tuned to data needs

- Integrated Security

IBM's Systems and Storage provide an ideal foundation for Workload Optimized Systems and Appliances

Power Systems

- IBM Smart Analytics Systems (7700/7710)
 - **11 percent** less TCA
 - **43 percent** less TCO than Oracle Exadata*
 - **16 percent** less TCA
 - **40 percent** less TCO than Teradata*
- Optimization of Cognos & SPSS analytics on AIX
- Power i BI Bundle
 - DB2 WebQuery for IBM i
 - Analytics packaged solution for midmarket
- IBM Watson



System x

- IBM Smart Analytics Systems (5600/5710)
 - **80 percent** reduction in data storage space
- IBM SAP Hana Appliance
 - Scales to the largest SAP BW installations
 - Real time analytic for SAP operational data
 - "Accelerates analytics as much as 10,000x..." **
- IBM System x and Informix Warehouse Accelerator
- IBM System x and Cognos TM1
- IBM Netezza



System z

- IBM Smart Analytics Systems (9700/9710)
 - *Secure, available business analytics*
 - **70 percent less TCO** compared to distributed platforms
 - Up to 52% lower security admin costs
 - *Each processor uses less energy than a 40 Watt light bulb*
- IBM DB2 Analytics Accelerator (IDAA)
 - Combining the best of System z and Netezza
 - Queries run up to 1000x faster (hrs to seconds)
 - *Operational in 3 days, ROI in 4 months*



Storage

- IBM Real-Time Compression Appliance
 - **Shrink active data up to 80%**
 - Keeps more data on-line for improved analytics and decision making
- IBM ProtecTIER
 - Optimized backups & restores
- IBM Easy Tier
 - **Self-tuning storage** by moving data to the right place
- IBM Disk Systems
 - XIV, DS8000, Storwize V7000, DS3000, SVC
- IBM NAS Systems
 - SONAS, N Series, Storwize V7000U




* Based on relative performance for complex mixed workload environments
https://www14.software.ibm.com/webapp/iwm/web/signup.do?source=sw-infomgt&S_PKG=dwanalyst6

** Bill McDermott, Co-CEO SAP, Keynote Address at FKOM 2012

Optimizing for Different Analytic Workloads



Business Analytics



Deep Analytics



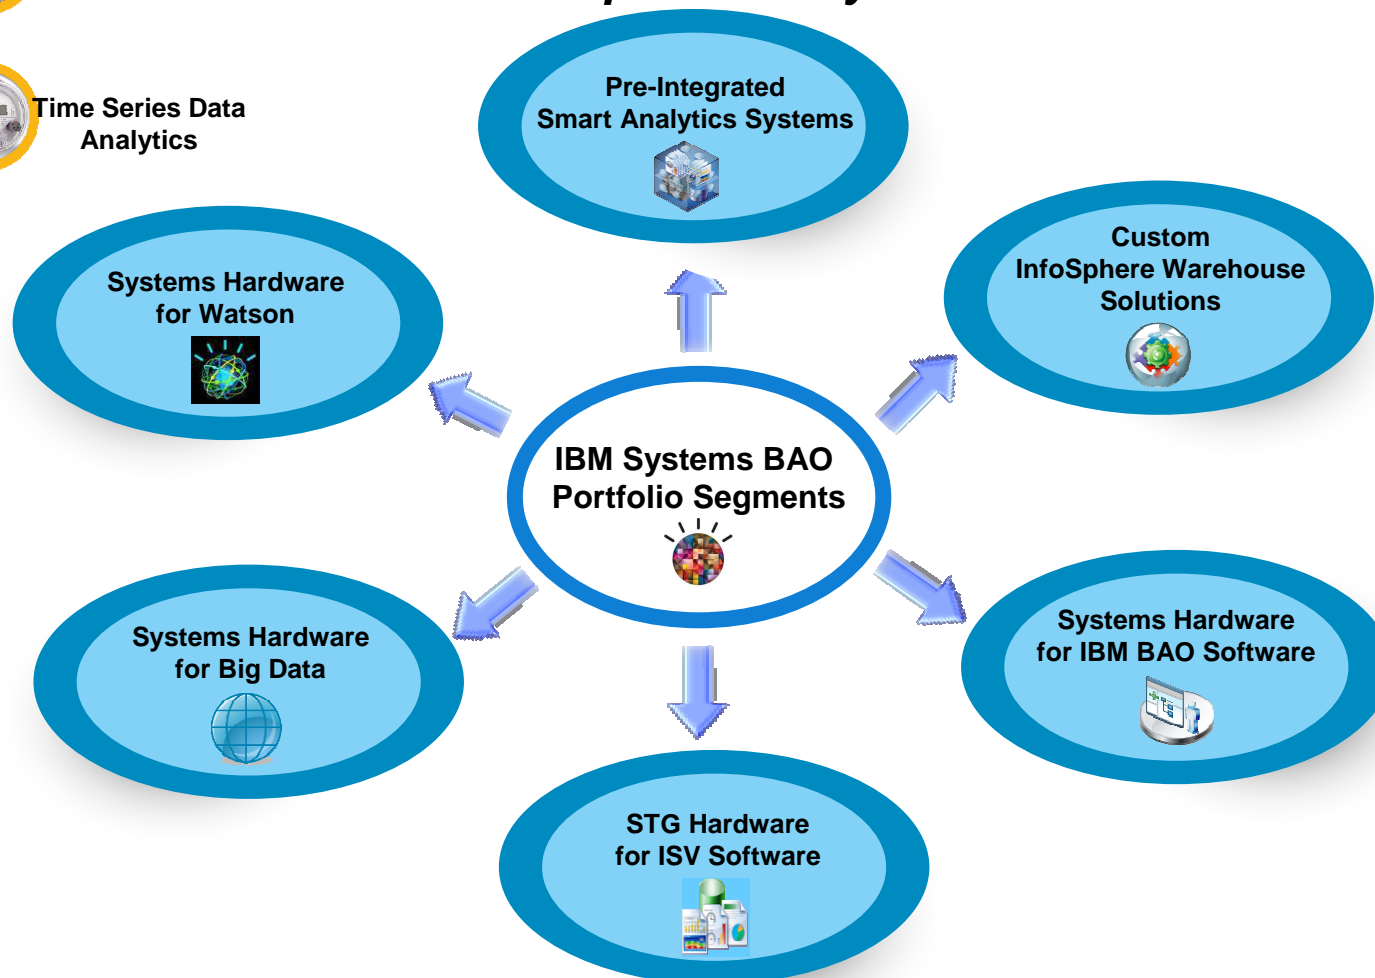
Operational Analytics



Time Series Data Analytics

IBM Systems for Smarter Analytics

Workload Optimized Systems



IBM System Storage for Big Data

Align System Storage Designed to Support Big Data



Self-tuning storage systems
Stop storing so much
Move data to the right place
Store more with what's on the floor



Managing unstructured data growth
Big File Systems with Big Files
Global Name Space
Data when you need it, where you need it



Data Protection and Retention
Perfect balance of performance and cost
Continuous data availability
Smart Archive



XIV SONAS SVC
 DS8000 DS3000 StorWize V7000



IBM Power Systems for SPSS & Cognos

Optimized for maximum performance



Your complete solution for Business Analytics

Cognos.
software

Cognos BI optimized for maximum performance on POWER7

•**40% better** performance with Cognos Business Intelligence V10.1.1 on POWER7/AIX 7.1, over Windows 2008 on x86

SPSS optimized for maximum performance on POWER7

•**22% better** performance for real-time scoring with SPSS Collaboration and Deployment Services V4.2 on POWER7/AIX 7.1, over Windows 2008 on x86

•**38 times better** performance for real-time scoring with IBM SPSS Collaboration and Deployment Services V4.2 optimized for POWER7 over default POWER7 environment configuration settings.



SPSS
AN IBM COMPANY

Sources: Best Practices and Advantages for Cognos BI on POWER7 (IBM, BP); Best Practices for SPSS on POWER7 (IBM, BP); Power Systems Analytics with Cognos and SPSS (IBM, BP)

IBM DB2 Analytics Accelerator for System z

Best of both worlds – for the next generation Analytics workload

Anticipate

Extreme performance for complex business analytics

Workload-optimized appliance add-on for System z based on Netezza

- Deeply integrated with DB2 for z/OS
- Completely transparent to applications and end users

Fast, predictable response times for “right time” analysis

Improves price/performance for analytics workloads

Minimize the need to create data marts for performance

Highly secure environment for sensitive data analysis



Unprecedented Speed, Reliability and Security

Exploits hardware accelerators

Response times to enable “train of thought” analyses

IBM Systems solution for SAP HANA™

Simple, Seamless, & Scalable

Anticipate

Your complete solution for Business Analytics

Real time analytics for SAP operational data

Certified eX5 workload optimized models


preconfigured with software preload

Accelerates analytics on SAP

operational data for ERP data marts and Business Warehouse

Scales out to support the largest environments

Services to deploy, manage, & maintain



Speed

Accelerates analytics as much as 10,000x

Solution

The fastest growing product in SAP history



IBM Netezza Appliance

The Simple Data Warehouse Appliance for Serious Analytics

Anticipate

Industry's fastest TTV and lowest TCO

Purpose-built analytics engine


- Runs analytic computations directly in the appliance, thereby accelerating analytic queries and shortening query times

Integrated database, server & storage

- Maximizes performance by running parallelized, in-database advanced analytics algorithms

Low total cost of ownership

- Reduces the cost and expands the available disaster recovery options, serving as a consolidated hot-standby platform during an outage



Transforming information into business insight

Speed

10-100x faster than traditional systems

Simplicity

Minimal administration and tuning

Scalability

Peta-scale user data capacity

Smart

High-performance advanced analytics

IBM Watson: A Workload Optimized System

integrated optimized system for business analytics



Watson will transform how technology is applied

A System Designed For Answers:

- Designed to answer the questions of today and tomorrow using deep QA Architecture
- Addresses natural language questions with accuracy and confidence in seconds not hours or days
- Processes massive array of information including unstructured “Big Data”
- Built on a cluster of commercially available Power 7 servers, optimized for complex analytic workloads and used by thousands today.



Speed



Operates at up to 80 Teraflops

Scalability

Scales out with and searches vast amounts of unstructured information with UIMA & Hadoop open source components



Watson: A Workload Optimized System

- 90 x IBM Power 750 servers
- 2,880 POWER7 cores
- POWER7 3.55GHz chip
- 500GB per sec on-chip bandwidth
- 10Gb Ethernet network
- 15 Terabytes of memory
- 20 Terabytes of disk, clustered
- Operates at up to 80 Teraflops
- Runs IBM DeepQA software stack
- Scales out with and searches vast amounts of unstructured information with UIMA & Hadoop open source components
- SUSE Linux provides a cost-effective open platform which is performance-optimized to exploit POWER 7 systems
- 10 racks include servers, networking, shared disk system, cluster controllers



Summary and Thank You

Smarter**Analytics** Live

Quando la conoscenza trasforma il business

Milano, 5 giugno 2012 (ore: 9.00-17.30)
Sede Gruppo 24 ORE, Via Monte Rosa 91

