

# Big Data and Analytics on z Systems



## 5 key takeaways

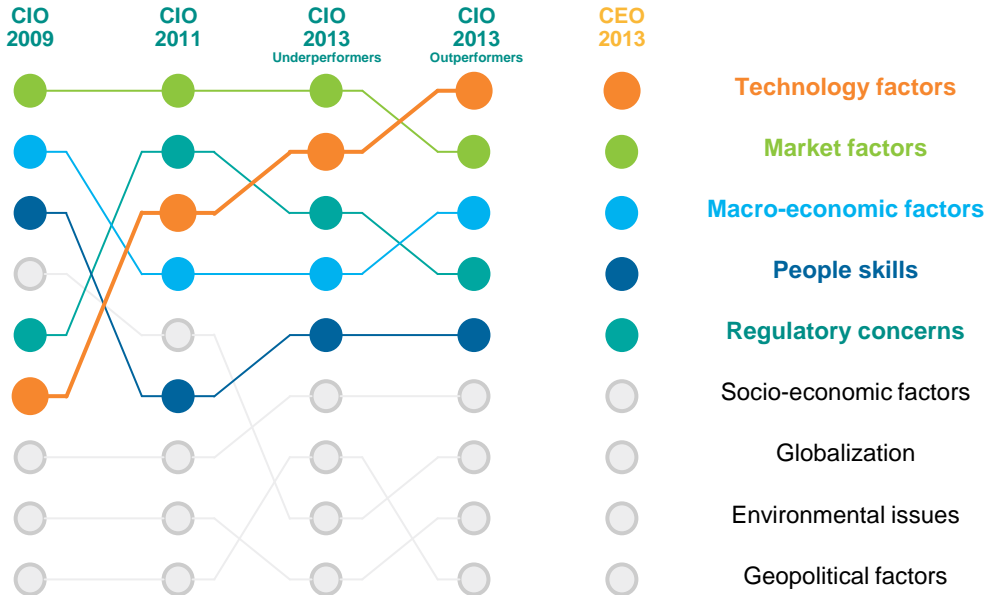
---



- **Many organizations are trying to deliver instantaneous, on-demand customer service** with IT systems designed to provide after-the-fact intelligence
- **Achieving insight with every transaction demands** a holistic implementation of an integrated data lifecycle with business-critical systems
- **z Systems has the vision, strategy and technology** to fuse transactions and analytics by eliminating the latency and complexity pitfalls that develop with a distributed approach
- **z Systems "operational analytics" builds advanced decision management support** on this integrated data platform injecting intelligence into operations without sacrificing performance
- **Truly transformational business opportunities** require truly transformational infrastructure - and that infrastructure is z Systems

- **Why analytics**
- Current limitations of analytics
- Integrated transaction and analytic processing
- Case studies
- Where to go for more information

# Today's leaders recognize the importance of technology




# Analytics separates underperformers from outperformers

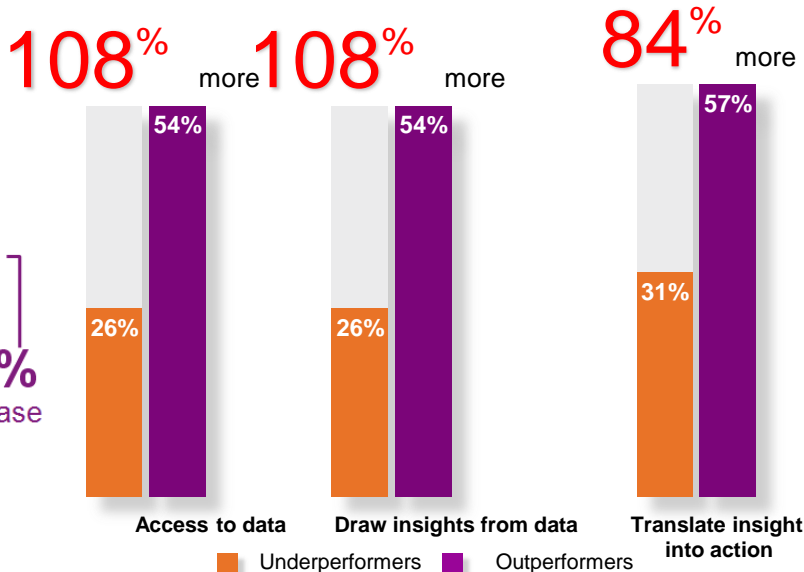


Outperformers strongly differentiate their organizations in three key areas

Respondents who say analytics creates a competitive advantage




Sources: IBM Institute for Business Value



Source: 2012 IBM CEO study, "How good is your organization at driving value from data?"

# The Big Data starting point

Transactional sources are the dominant data types analyzed



N =465 (multiple responses allowed)

Source: Gartner (September 2013)

Gartner research note "Survey Analysis - Big Data Adoption in 2013 Shows Substance Behind the Hype" Sept 12 2013


Analyst(s): Lisa Kart, Nick Heudecker, Frank Buytendijk

# Analytics has evolved to a business imperative



More organization are using analytics to create a competitive advantage...

Respondents who believe analytics creates a competitive advantage



Source: The New Intelligent Enterprise, a joint MIT Sloan Management Review and IBM Institute of Business Value analytics research partnership. Copyright © Massachusetts Institute of Technology 2011

...and leaders are outperforming their competitors in key financial measures

**1.6x** Revenue growth

**2.0x** EBITDA growth

**2.5x** Stock price appreciation

Source: Outperforming in a data-rich, hyper-connected world, IBM Center for Applied Insights study conducted in cooperation with the Economist Intelligence Unit and the IBM Institute of Business Value, 2012

# Evolving customer needs are driving new business IT models



- From...**
- Applications
  - Static Infrastructure
  - Programmed Systems
  - Structured Data At Rest
  - Stable Well-Defined Workloads
  - Standard Devices
  - Proprietary Standards
  - Corporate-owned IT



- To...**
- Composable Services**
  - Dynamic** Services, defined by Software
  - Cognitive** Systems
  - Unstructured **Data in Motion**
  - Unpredictable **Workloads**
  - A Variety of **Devices**
  - Open** Innovation
  - Infrastructure **As-a-Service**

These shifts in the industry are informing IBM's strategy

Data is becoming the world's new natural resource

500 million DVDs worth of data is generated daily

1 trillion connected objects and devices by 2015

80% of the world's data is unstructured

Cloud is transforming IT and business processes into digital services

85% of new software is being built for cloud

25% of the world's applications will be available in the cloud by 2016

72% of developers say cloud-based services or APIs are central to their app designs

Social, mobile and access to data are changing the way we are understood and engaged

80% of individuals are willing to trade their information for a personalized offering


84% of millennials say social and user-generated content has an influence on what they buy

5 minutes: response time users expect once they have contacted a company via social media






# Data is becoming the new natural resource



**1 trillion**

connected objects and devices on the planet generating data by 2015



**2.5 billion**

gigabytes of data generated every day

**80%**


of the world's data is unstructured. Audio, Video, Sensor data, Social media. All represent new areas to mine for insights.

Today, every discussion about changes in technology, business and society must begin with data. In its exponentially increasing volume, velocity and variety, data is becoming a new natural resource. It promises to be for the 21<sup>st</sup> century what steam power was for the 18<sup>th</sup>, electricity for the 19<sup>th</sup> and hydrocarbons for the 20<sup>th</sup>

# CIOs in outperforming enterprises are focusing particularly heavily on developing the resources to acquire deeper customer insights




## Combining internal and external data for better insights



**25%** | **49%**  
Underperformers | Outperformers

## Customer analytics drive big data initiatives



**27%** | **41%**  
Underperformers | Outperformers


Source: IBM Institute for Business Value, Global C-Suite Study, 2013

- Why analytics
- **Current limitations of analytics**
- Integrated transaction and analytic processing
- Case studies
- Where to go for more information

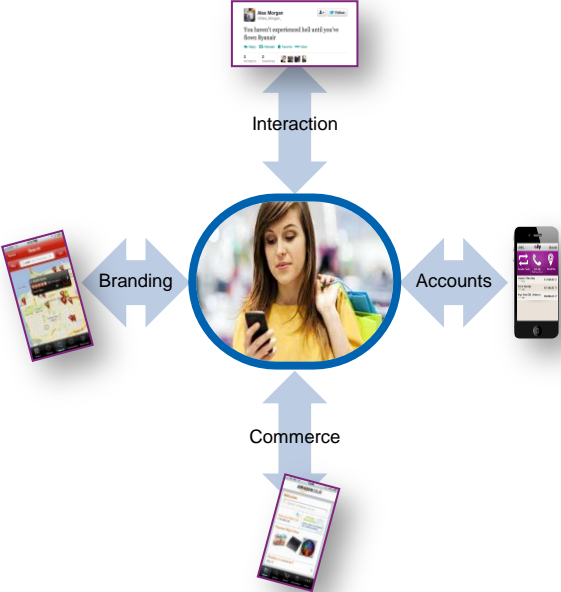
# The business relationship has changed forever



**Then:** *“I have an offer – let me find a customer I can sell to”*



**Now:** *“I have a customer – what do they need most?”*



**Customer experience is the competitive advantage for top-line growth**

# Leaders must leverage data to outperform the competition



Drive **top-line growth** via transformational new services

80%

of marketers send **same content** to all subscribers

6%

of businesses “**extremely satisfied**” with ability to use customer data for decisions

+7.6%

annual increase in **customer lifetime value** for firms that use engagement analytics

Dramatically “**maximize yield**” on investments in standard electronic business practices

\$226M

estimated **fraud loss** to healthcare

16%

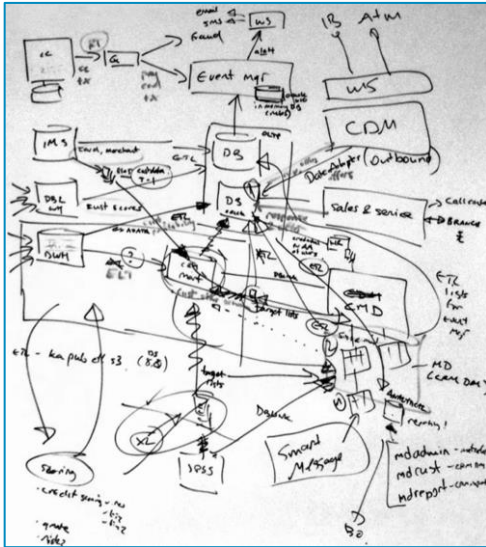
government **tax revenues lost** to non-compliance

\$100M

typical banking regulation **non-compliance fine**

***IT must be exploited as a business strategy***

# Challenges with traditional analytics processing



## Significant complexity

Data is move from operational databases to separated data warehouses/data marts to support analytics

## Analytics latency

Transactional data is not readily or easily available for analytics when created

## Lack of synchronization

Data is not easily aggregated and users are not assured they have access to “fresh” data

## Data duplication

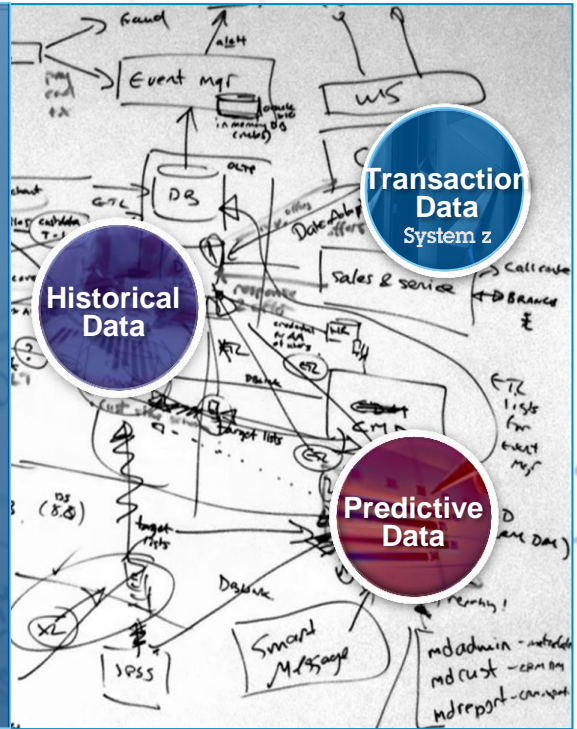
Multiple copies of the same data is proliferated throughout the organization

## Excessive costs

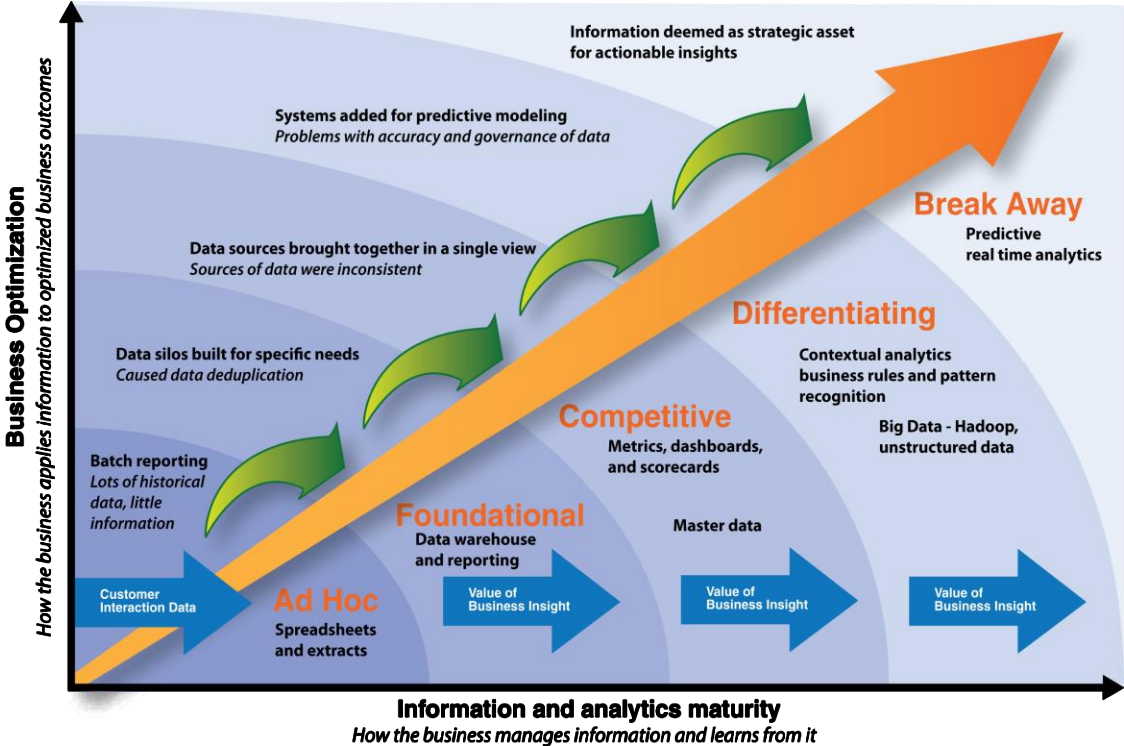
An IT infrastructure that was not designed nor can support real-time analytics

Business has fundamentally changed – but IT remains aligned to the old way of doing business

*There can no longer be multiple systems; there must be a **fully-integrated, end-to-end system** that executes intelligent business processes*




# Analytic Roadmap





# Transactions and analytics processed together

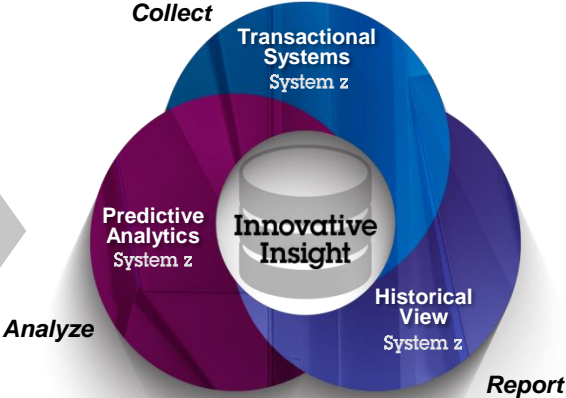

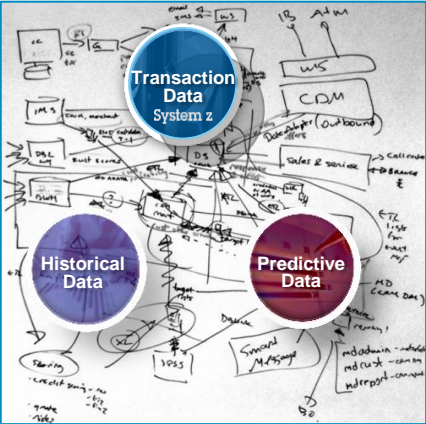


- Purchase made
- Resources consumed
- Bill paid
- Claim submitted
- Information updated
- Call center contacted

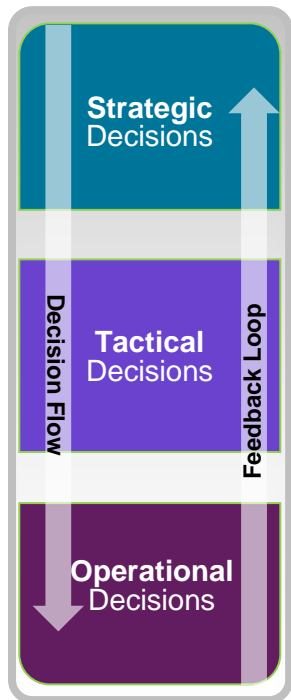
- What happened?
- How many, how often, where?
- What actions are needed?
- What will happen if?
- What will produce the best outcome?

***Analytics as part of the flow of business; insights on every transaction***

# Enabling innovative insight



# Enabling an integrated Decision Management System



- Guide overall direction of the enterprise
- *Should we expand overseas?*
- Time to insight: **highly variable**

- Manage and control operations
- *What upsell is appropriate?...can I push products in real-time as customers navigate the store?*
- Time to insight: **short to mid-term**

- Handle every customer interaction
- *Is this payment request fraudulent?*
- Time to insight: **immediate**

Accurate warehouses and accelerated reporting

Analytics embedded in operational systems

# Clients find value in running analytical and transactional on one platform



## Extract, Transform and Load (ETL)

### Multiple copies of data

Redundancy, inconsistency, complexity and cost

### Significant compute power

Increases cost

### Transaction and analytics isolation

Increases time to insight

1TB ETL per day, Initial copy plus three derivatives costs  
**>\$8 million** over 4 years



Operational data



Analytical data

Source: CPO internal study. Assume dist. send and load is same cost as receive and load.. Also, assume 2 switches and 2 T3 WAN connections.

Clients report using 30% - 40% of z Systems cycles to move data

# z Systems is uniquely positioned



zBC12



zEC12



- Creating personalized client experiences
- Speeding delivery of new products and services
- Integrates business operations within a single system
- A robust infrastructure

# The most valuable insights occur when the analysis executes where the data originates



## **zEnterprise** – a single solution to:

- Provide a single source of data for multiple analytics applications
- Access, Combine & Manage a relevant mix of information
- Deliver timely and secure access for more accurate answers

72%

of responders plan to analyze transactional data from enterprise applications using Big Data technologies

80%

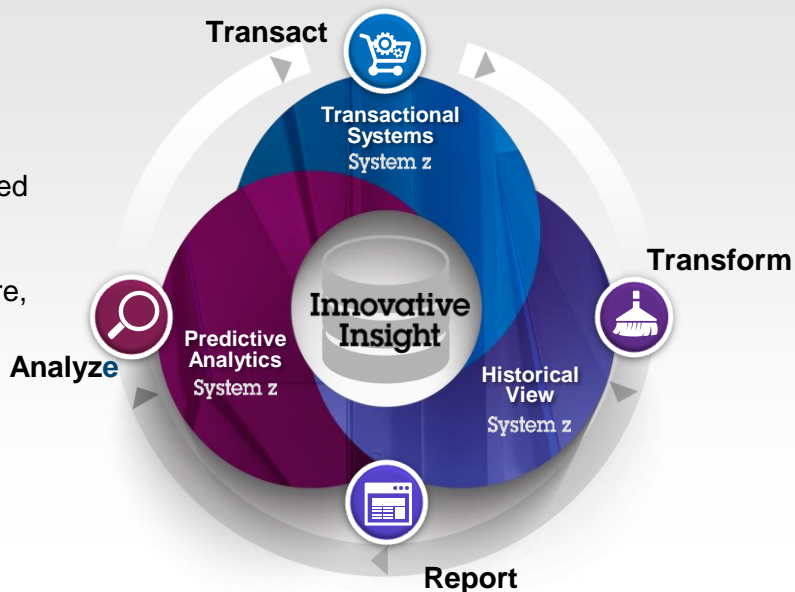
of world's corporate data resides or originates on mainframes

# The z Systems strategy

*Integrating operations & analytics in one streamlined lifecycle*

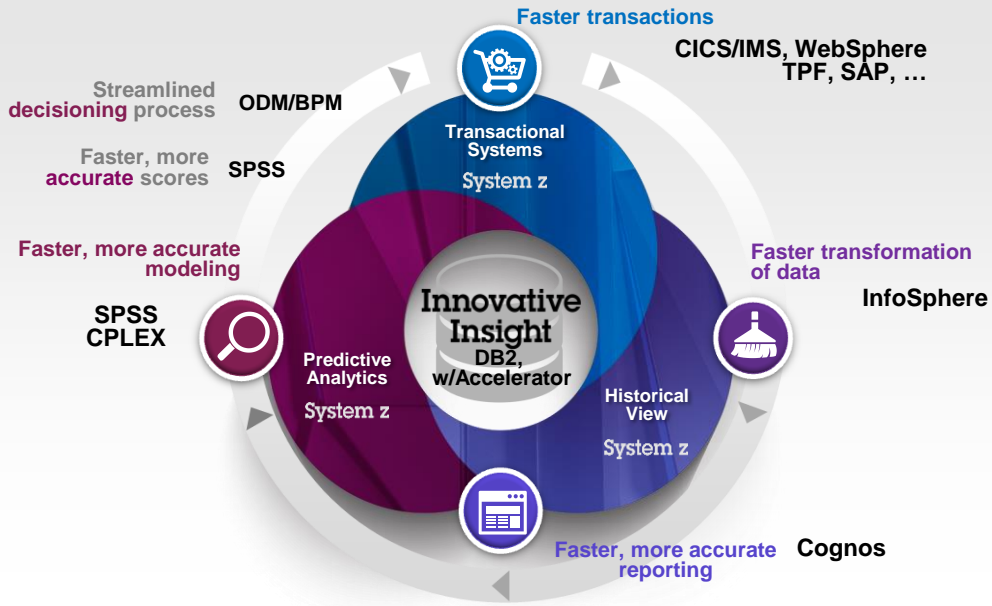


- Better business response
- Reduced data movement, reduced complexity, reduced configuration resources
- More accurate, more secure, more available



# The z Systems strategy

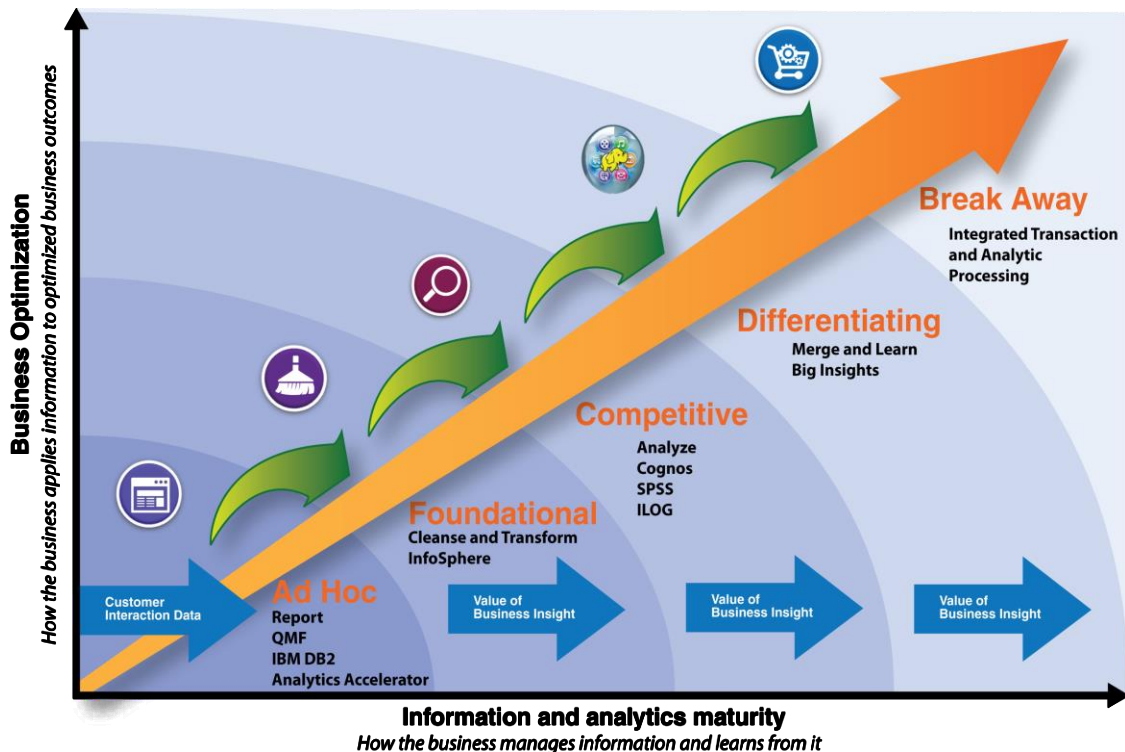
*Integrating operations & analytics in one streamlined lifecycle*





# z Systems Point of View

Building a foundation to grow with business needs



# z Systems: Faster transactions

*Around an integrated view of business-critical data*



- Designed for systems with mission critical SLAs
- Optimized for high performance and scalability
- Designed for continuous operation and disaster recovery
- Dynamic, goal-oriented workload management
- Unique z Systems hardware architecture capabilities
- Fully utilize all existing capacity for lower cost



CICS/IMS, WebSphere, TPF, SAP

# Accelerate data warehousing on z Systems

## DB2 Analytics Accelerator improves the speed of business decisions



*Blending z Systems and Netezza technologies to deliver unparalleled, mixed workload performance for complex analytic business needs.*



- Enables integration of analytic workloads and transaction processing on zEnterprise
- Ensures cost-effective, timely, accurate and secure insight in near real-time
- Benefit from the proven z Systems qualities of service Accelerates complex queries, up to 2000x faster
- Improves access to historical data and lowers the cost of storing, managing and processing historical data
- Minimizes latency
- Reduces zEnterprise capacity requirements
- Improves security and reduces risk
- Complements existing investments

# z Systems: Faster transformation of data

*Around an integrated view of business-critical data*



- Enables many OLAP queries against standard OLTP/ODS data schema
- In-Database transformation for reduced data movement
- High performance and linear scalability
- Multi-tier distributed ETL configurations minimized or eliminated
- Data currency in warehouse optimized via incremental update
- Accelerator data load from DB2 managed & highly optimized
- Accelerator Loader provides integration of non-DB2 data sources



InfoSphere, DB2 Accelerator

# z Systems: Faster reporting, OLAP and ad-hoc query

*Around an integrated view of business-critical data*



- Real-time DB2 query optimization determines best fit architecture query execution
- Single point of entry for all requests
- Application transparency, no change to SQL
- Complex tuning tasks eliminated
- Underlying complexity hidden from users, IT staff
- Analytics run up to 2000x faster
- Enables more exploration and creativity



Cognos, DB2 Accelerator

# Remember the challenge we spoke about earlier



Drive **top-line growth** via transformational new services

80%

of marketers send **same content** to all subscribers

6%

of businesses “**extremely satisfied**” with ability to use customer data for decisions

+7.6%

annual increase in **customer lifetime value** for firms that use engagement analytics

Dramatically “**maximize yield**” on investments in standard electronic business practices

\$226M

estimated **fraud loss** to healthcare

16%

government **tax revenues lost** to non-compliance

\$100M

typical banking regulation **non-compliance fine**

***IT must be exploited as a business strategy***

# z Systems: Faster, more accurate modeling

*Around an integrated view of business-critical data*



- SPSS predictive analytics modeling is done where historical data is – in DB2 Accelerator
- Move from rules-only approach to incorporating predictive models
- In-database transformation enables data preparation and statistical computation optimization
- Eliminates/reduces data movement and network bandwidth constraints
- Exploits massive parallelism for performance
- Enables frequent model refresh, facilitating more accurate models and better business decisions
- Improves both data governance and security



SPSS, CPLEX,  
DB2 Accelerator

# z Systems: Faster, integrated transactional scoring

*Around an integrated view of business-critical data*



- Actionable insight on every transaction
- SPSS GUI deployment of scoring models as DB2 UDFs
- Analytic scoring done within the scope of an OLTP transaction with negligible impact to SLAs
- Scores executed 10-100x faster than making distributed system calls to remote scoring engines
- Will exploit future generation z Systems hardware
- Real time decisions at the point of interaction with the customer
- Analytics in the flow of business to stop fraud, increase customer loyalty, increase revenue and reduce risk.



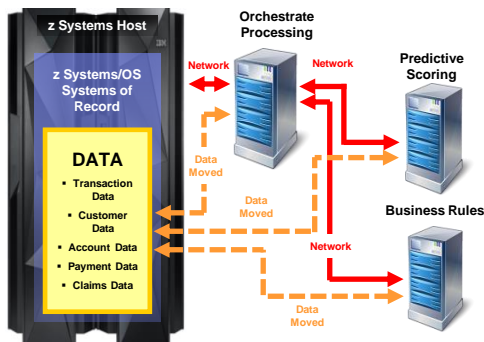
SPSS, CPLEX,  
DB2 Accelerator



# Analytics scenarios for customers with z Systems data



## Move the data to the analytics



- Performance of critical transactions **may not meet SLAs** due to data movement
- Customer needs to **create security infrastructures** across multiple servers
- Customer needs to **create audit infrastructure** across multiple servers to ensure governance
- Customer needs to **create availability and DR function** for multi-server transaction flows and in-transit data

## Move the analytics to the data



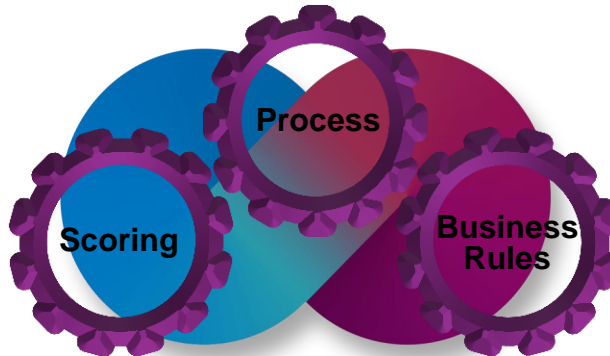
- Unparalleled, proven performance execution for models and rules, with **NO data movement**
- **Leverage existing best of class security** with z Systems infrastructure
- **Leverage existing transaction level auditing** and logging for governance
- **Leverage existing, tested HA / DR capabilities** already configured with z Systems

# Integrate operational analytics into OLTP with 3 core capabilities



## Fast, efficient process orchestration

- Encapsulate interaction with existing OLTP
- Composable based on organizational needs
- Efficiently prepare inputs for and invoke predictive models and rules



## Integrate Advanced Analytics

- Predictive insight on each transaction
- Determine likelihood of fraud, likelihood of opportunity to enhance customer value....
- Co-locate with data for huge performance scale and efficiency

*Faster, more accurate scores*

## Automate real-time decisions

- Apply organization specific thresholds
- Matrix results from one or more predictive scores
- Introduce line of business specific parameters in decision process

*Streamlined decisioning process*

# Why Operationalize In-Transaction Analytics on z Systems?

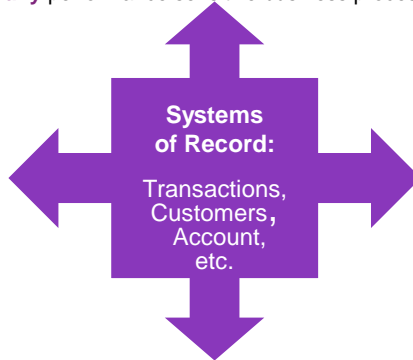


## Performance

- For clients with z Systems data, gain business benefit of advanced analytic capabilities **without moving data** – if data is moved, clients may not be able to meet business critical SLAs
- Analytics has access to more current data for yielding more **accurate** score results
- Provides value for **any** performance sensitive business process – real-time, near real-time, batch

## Availability

- Leverage **existing z Systems HA / DR** infrastructures and extend to fraud detection capabilities
- **Flexibility** to determine optimal commit scopes for analytic operations based on business recovery needs



## Governance / Security

- **Avoid data proliferation**, preserve tighter data governance and auditability
- Preserve the **security envelope** of transactions during analytic operations

## z Optimizations

- **Extend** underlying z accounting, metrics, monitoring , workload management infrastructures to analytic operations
- Extend the use of unique z **optimizations** to analytic processes, Hardware Crypto Acceleration, zIIPs, ...

# What's keeping mainframe clients from using Hadoop for mainframe data?



They are worried about data governance as the data moves off of z.

Data is considered secure as long as it is on z.

How do you secure sensitive data once it has left z?

The ingestion of data from z into external Hadoop system is turning into a bottleneck.



## *Lead to key requirements*

z Systems needs to be in "control" of the data, and existing security policies must be applied

There needs to be high speed / optimized connectors between traditional z/OS LPARs and the Hadoop cluster


# Two z Systems options for using Hadoop




## IBM InfoSphere BigInsights + Veristorm vStorm Connect

On-platform analysis of data that does not fit with relational tools

Safely combine z/OS and external data for analysis (e.g. improving claims response time)



- VSAM
- QSAM
- SMF, RMF
- System logs
- Operator logs
- Application logs
  - etc.



## IBM InfoSphere BigInsights

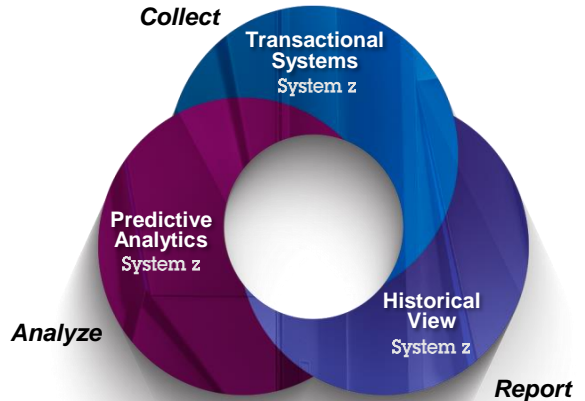
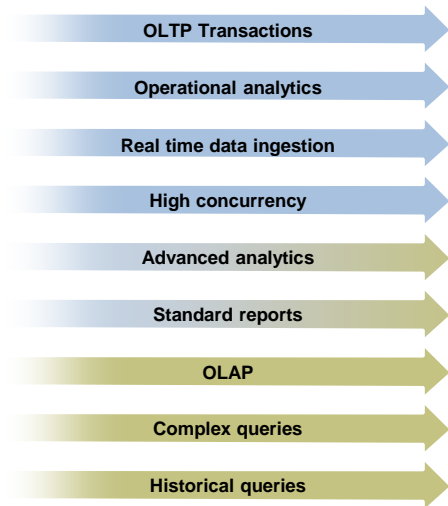


DB2 z/OS can integrate insights from big data sources to augment analysis (e.g. improve accuracy of fraud detection with "nuggets" from social media)



# z Systems Transaction & Analytic Processing:

*Everything is online – analytics in the right place*



*Best of class Data Life Cycle Management for:  
Fighting fraud, preventing financial crimes, generating customer insights, ...*

- Why analytics
- Current limitations of analytics
- Integrated transaction and analytic processing
- **Case studies**
- Where to go for more information

**Business Challenge:**

How to maximize profitability as its business grows

**Technical Challenge:**

Running its growing transaction processing and analytics workloads side by side without increasing compute requirements

**Solution:**

Deploy IBM DB2 Analytics Accelerator for z/OS bringing together transactional processing and analytics workloads in a cost-effective solution

achieved its  
objective of  
**providing  
access to the  
most timely,  
accurate data**  
to improve  
customer  
satisfaction

“Queries that used to take five hours to complete are now processed in just 20 seconds in the optimized mainframe environment—and we can run them any time, day or night, with no interruption to our production systems on the mainframe.”

*Thomas Baumann, IT Performance Architect at Swiss Mobiliar*



**PETROL**

**dramatically decreases  
analytics query times**

**IBM**

**Business Challenge:**

How to improve customer service and satisfaction in order to drive greater revenue

**Technical Challenge:**

Existing analytic processes were unable to manage the analysis of historic and transaction data from Petrol's retail stores, service stations and home oil/gas businesses

**Solution:**

Implemented IBM DB2 Analytics Accelerator to support high performance queries and IBM SPSS to make real time, point of sale product recommendations

**Increased  
retail sales  
revenue**  
through point-of-  
sale improvement,  
suggest-sell insight

*IBM provides us with tools that align with smarter commerce, enabling us to deliver the right message to the right person at the right time, to understand product affinities and intelligently drive the sale all in a customer centric way”*



**BANCA CARIGE**

Cassa di Risparmio di Genova e Imperia



**leverages new approach to real-time analytics to boost productivity by 400%.**

### **Business Challenge:**

How to maximize value from big data in order to improve product development and customer relationships

### **Technical Challenge:**

Unable to quickly extract actionable insights from big data and identify market opportunities in order to adapt or expand its offering to meet customer demand

### **Solution:**

Created a secure analytics platform, to extract true business value from their big data for better business decisions about everything from product development to special offers to promotions

**1000+ users**

simultaneously get high-speed analytics on

**real-time data**

**Time cut from months to weeks**

to deliver the insight needed to develop and release new marketing campaigns

*“IBM DB2 Analytics Accelerator enables us to support the additional workloads that come with business growth without activating more cores on the mainframe.”*

- Why analytics
- Current limitations of analytics
- Integrated transaction and analytic processing
- Case studies
- Where to go for more information

# Growing your IBM skills a new model for training



## Meet the authorized IBM Global Training Providers in the Enterprise Solution Showcase

- Access to training in more cities local to you, where and when you need it, and in the format you want
  - Use [IBM Training Search](#) to locate training classes near to you
- Demanding a high standard of quality / see the paths to success
  - Learn about the [New IBM Training Model](#) and see how IBM is driving quality
  - Check [Training Paths and Certifications](#) to find the course that is right for you
- [Academic Initiative](#) works with colleges and universities to introduce real-world technology into the classroom, giving students the hands-on experience valued by employers in today's marketplace
- [www.ibm.com/training](http://www.ibm.com/training) is the main IBM training page for accessing our comprehensive portfolio of skills and career accelerators that are designed to meet all your training needs.



Global Skills Initiative

Avnet Academy   
Global Training Provider



Global Knowledge.

INGRAM  
MICRO<sup>®</sup>

LearnQuest

**Thank You**

