Big Data and Real-Time Analytics on System z

Steve Mink

Information Management System z Strategy and Product Management, IBM Software Group

Agenda

- Defining analytics and the impact of new paradigms
- Traditional approaches to analytics
- Modern hybrid transactional / analytics environments
  - Gartner's perspective
  - DB2 Analytics Accelerator for z/OS
  - Cognos BI, SPSS, Cognos TM1, CMA
- Leveraging all of your data assets
  - InfoSphere BigInsights, Veristorm
- Wrap up with Q&A

What is Analytics?

- *Analytics derive insight from data*
- *Organizations use analytics to help optimize business performance*
- *Analytics are only as good as the underlying data foundation*

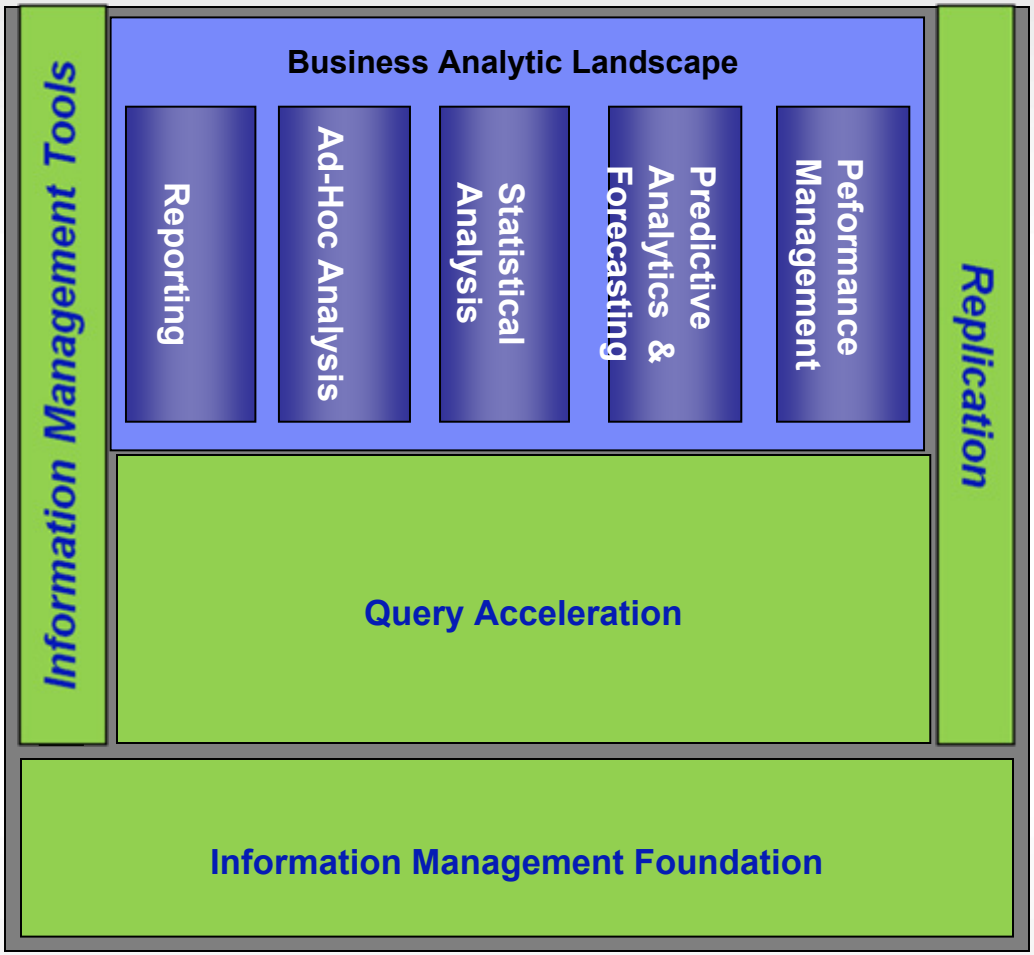

What is Business Critical Analytics?

- An analytics application that is tightly integrated with transaction systems and critical to the optimal running of a business
- Make decisions and deliver business insight based on real time or near real time data
- Failure of these analytics applications for any length of time can result in lost business
- Typically support a large concurrent user population with high volume of requests

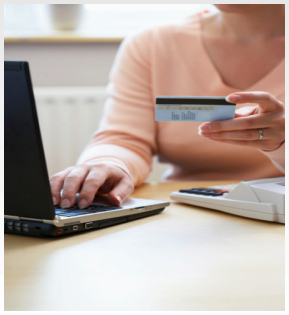

Preventing Fraud

Reducing Customer Churn

Cross-selling, up-selling customers

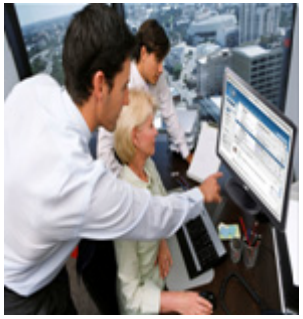

Realtime Operational Reporting

These applications require high degree of reliability, availability, scalability and low data latency

More users across the organization want access to analytics

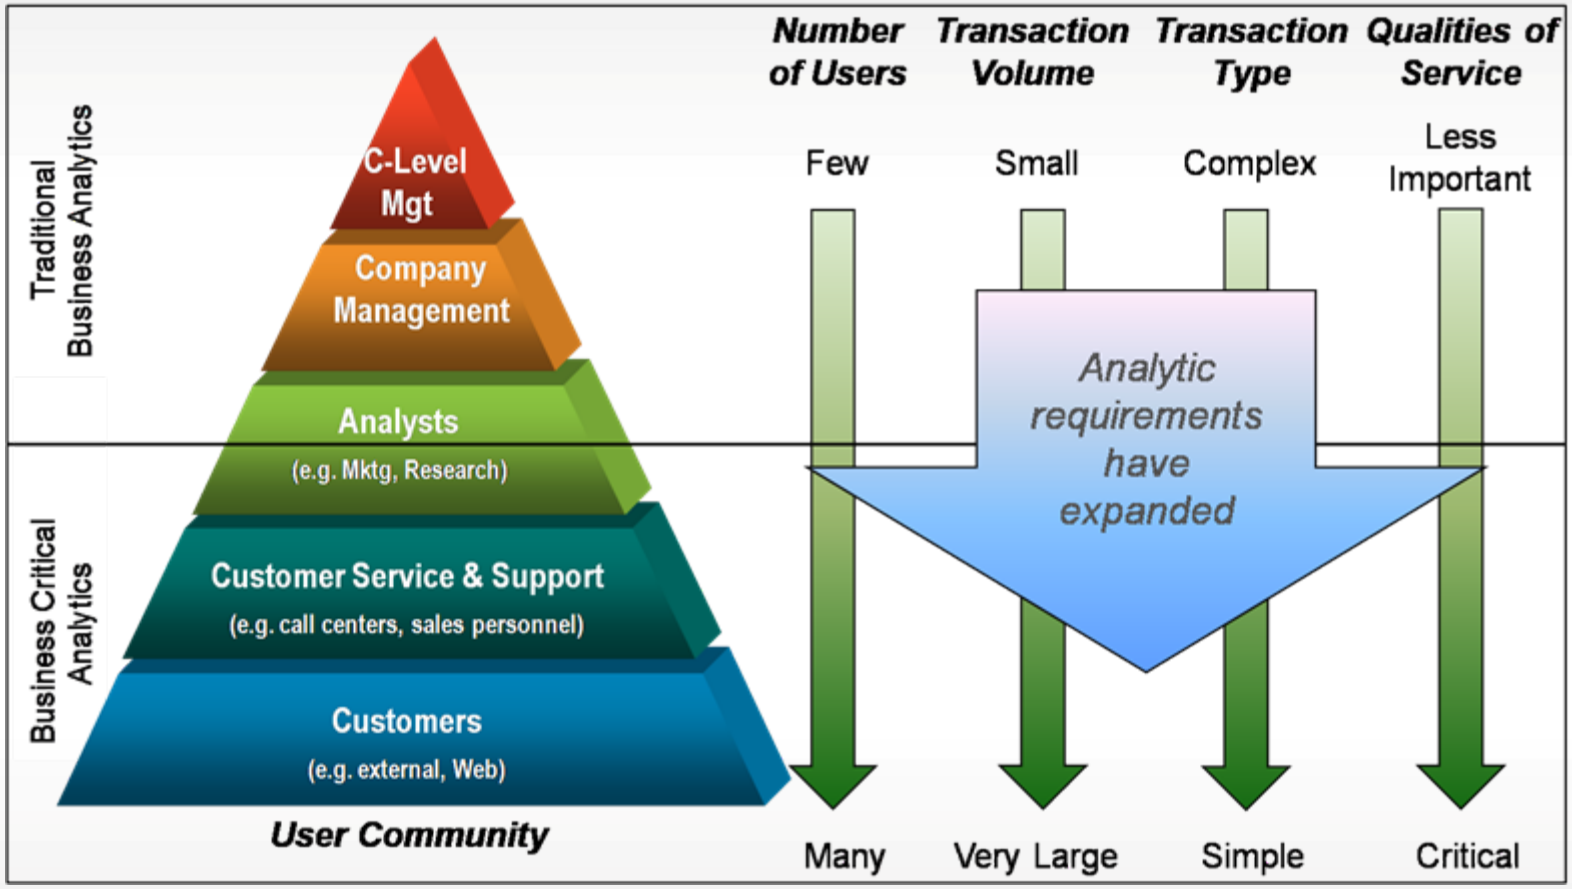

New paradigms impacting Big Data & Analytics

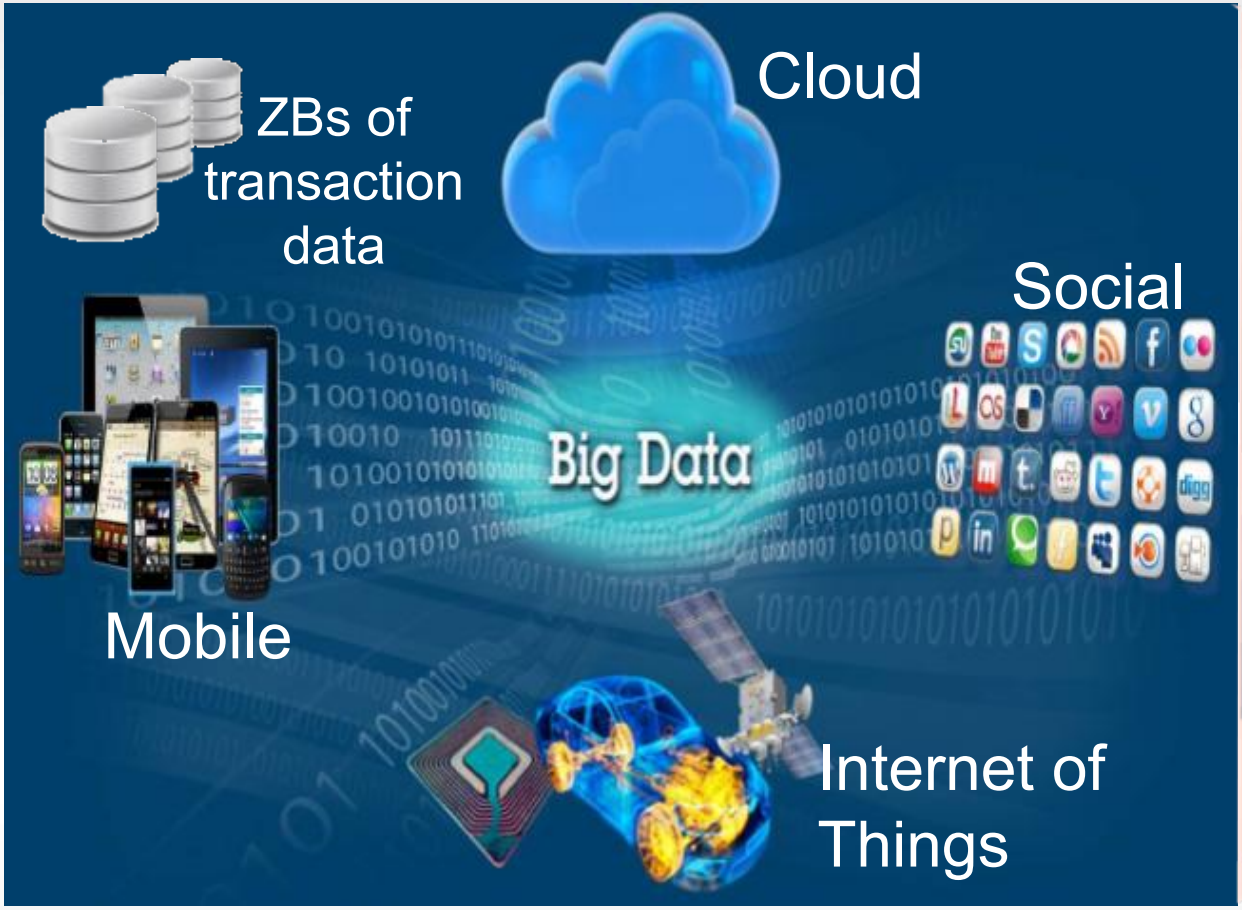

The Big Data Starting Point

Types of Data Analysed

Transactional sources are the dominant data types analyzed in big data initiatives

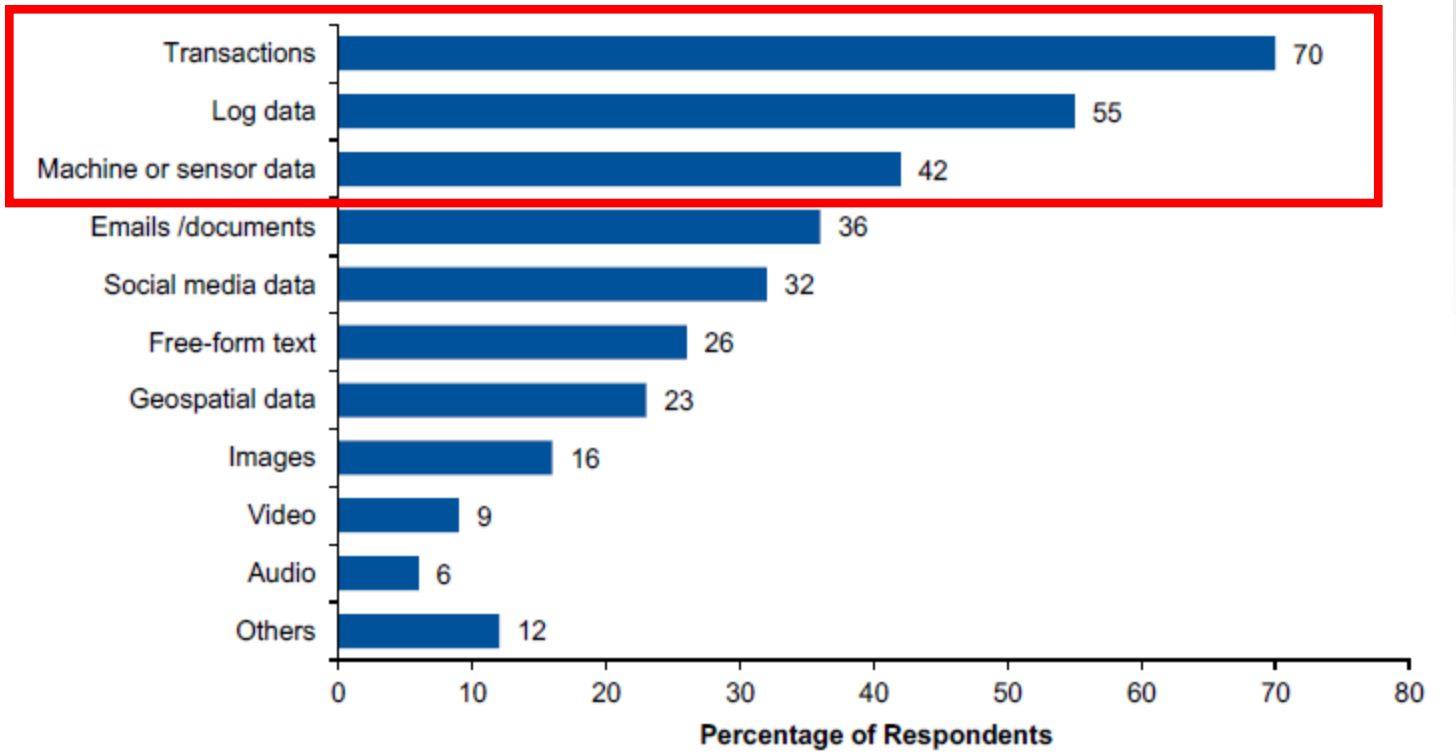

N =465 (multiple responses allowed)

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

The reality of Analytics

1. Organizations are using analytics to outperform their competition
2. More users across the organization want access to business critical analytics
3. Business critical analytics demands low latency, high qualities of service and performance
4. Spreading analytic components across multiple departments can increase data latency, cost, complexity and governance risk
5. Bringing analytic components to where data originates improves data governance, while minimizing data latency, cost and complexity

Traditional Approach to Analytic Systems

Operational Applications

Transaction Processing

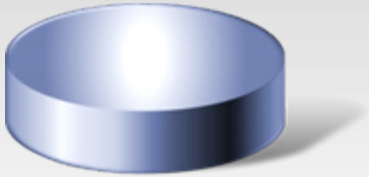

Shared Everything DB

High volume business transactions and batch processing running concurrently

Latency?
Security?
Data Governance?
Complexity?

Analytic Applications

Data Store, Business Intelligence, Predictive Analytics

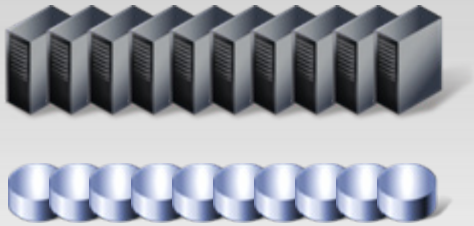

Shared Nothing DB

Low volume complex queries and batch reporting

Typically results in . . .

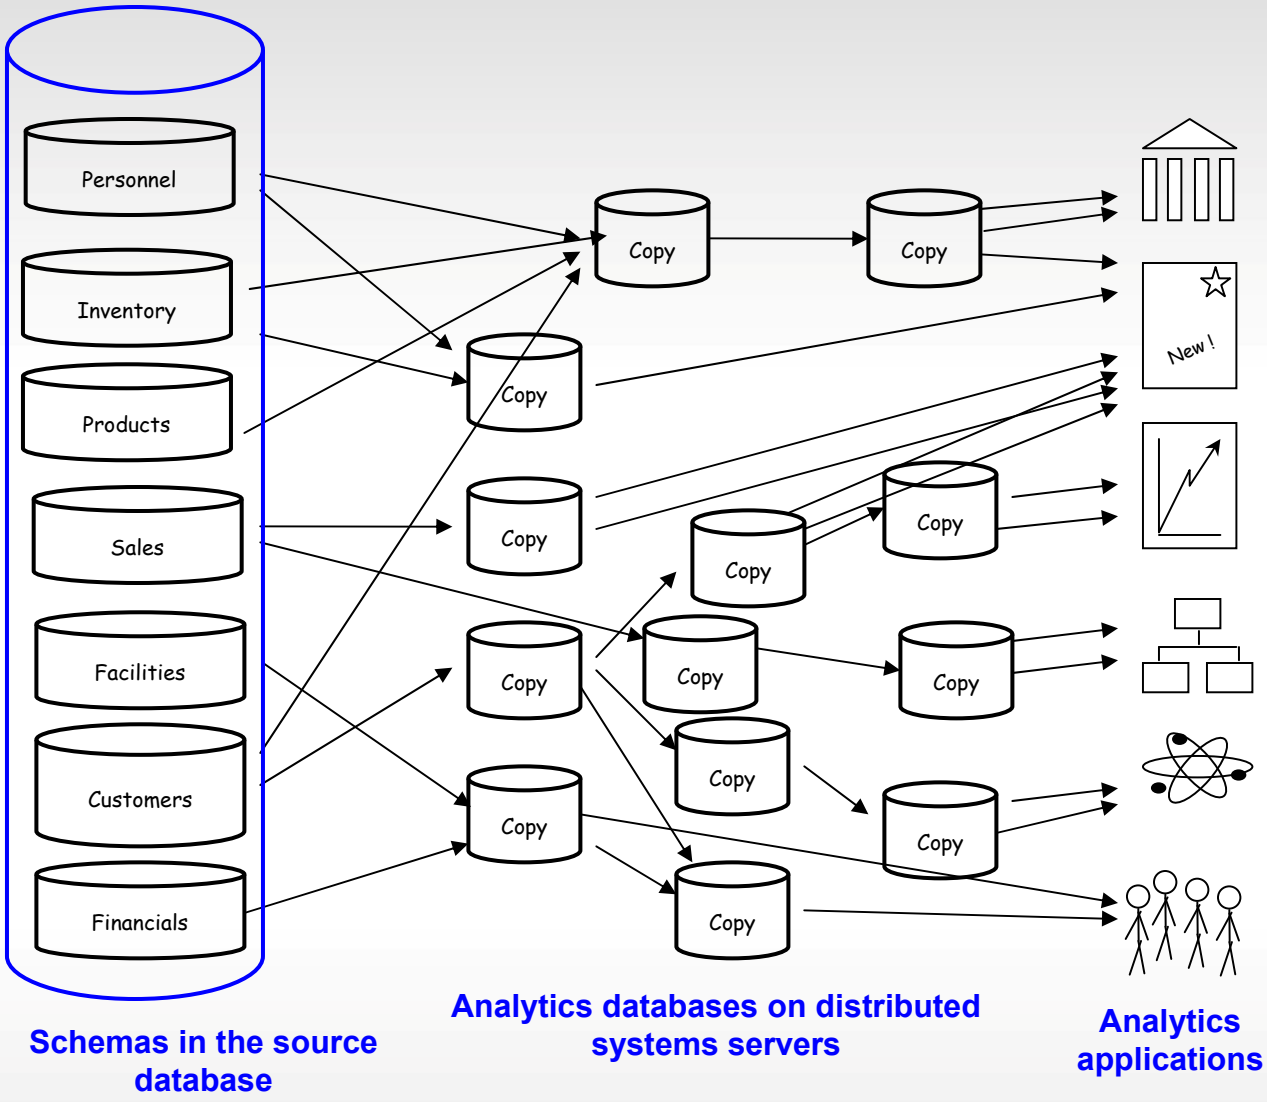

Problems:

- Data latency: time between transaction and insight
- Expensive, resource-intensive data replication processes
- Greater risk of data security breaches
- Data governance issues: copies of data can become inconsistent – do users trust the data?
- Data currency challenges: copies of data can become out-of-date – users demand timely data
- Proliferation of data silos impedes integration, reduces value derived from data assets

The Hybrid Transactional/Analytics Environment on System z

Combined Workloads

Transactional Processing, Traditional Analytics & Business Critical Analytics

Hybrid DB

Reduced Latency. Greater Security. Improved Data Governance. Reduced Complexity.

High volume business transactions and batch reporting running concurrently with complex queries

The Analyst Community Has Taken Notice!

Gartner®

- "By eliminating analytic latency and data synchronization issues, hybrid transaction/analytical processing will enable IT leaders to simplify their information management infrastructure"

- "This architecture will drive the most innovation in real-time analytics over the next 10 years via greater situation awareness and improved business agility"

Gartner Research Note G00259033: Gartner 01-2014 Hybrid Transaction Analytical Processing Will Foster Opportunities

DB2 Analytics Accelerator

Enabling a Hybrid Transactional / Analytics Environment on System z

- **What is it?**
  - A high performance appliance that integrates Netezza technology with zEnterprise technology, to deliver dramatically faster business analytics
- **What does it do?**
  - Accelerates complex queries, up to 2000x faster
  - Improves access to 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

PETROL

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 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"

Swiss Mobiliar
Insurance & Pensions

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 insight 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

BANCA CARIGE
Cassa di Risparmio di Genova e Imperia

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."

Query Execution Process Flow

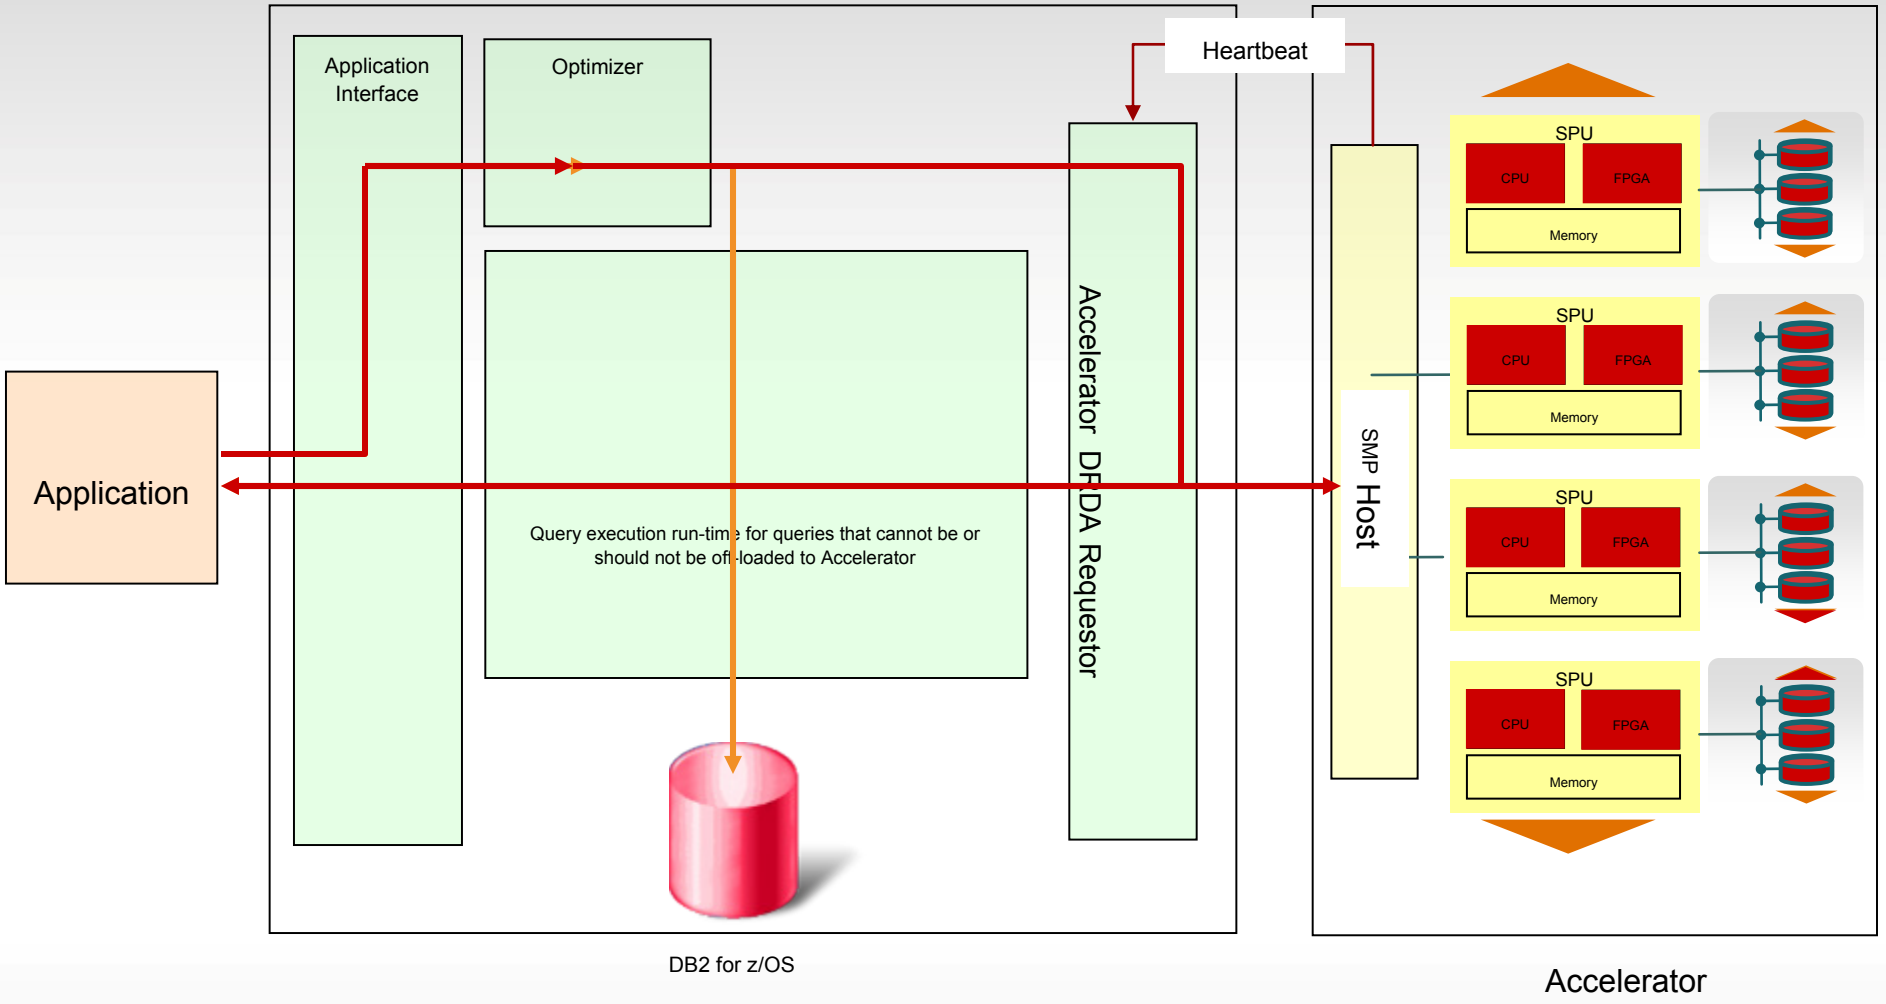

Queries executed without Accelerator

Queries executed with Accelerator

Heartbeat (availability and performance indicators)

DB2 Analytics Accelerator Version 4.1

Delivering critical insight at the speed of business

Accelerate a broader spectrum of queries

-Static SQL, multi-row FETCH, and multiple encodings on the same Accelerator

Improve Enterprise Robustness & Scalability

-Enhanced workload balancing

-Improved incremental update performance

Enhance High Performance Storage Saver

-Improved ease of use through built-in restore, better access control for archived partitions and protection of moved partitions

Supports new DB2 and PureData Technology

-DB2 11

-PureData Analytics N2001 hardware

DB2 Analytics Accelerator Roadmap

- Enhancing current capabilities
- Enabling more query acceleration
- Increasing transparency
- Supporting new use cases
- Utilizing PDA technology advantages

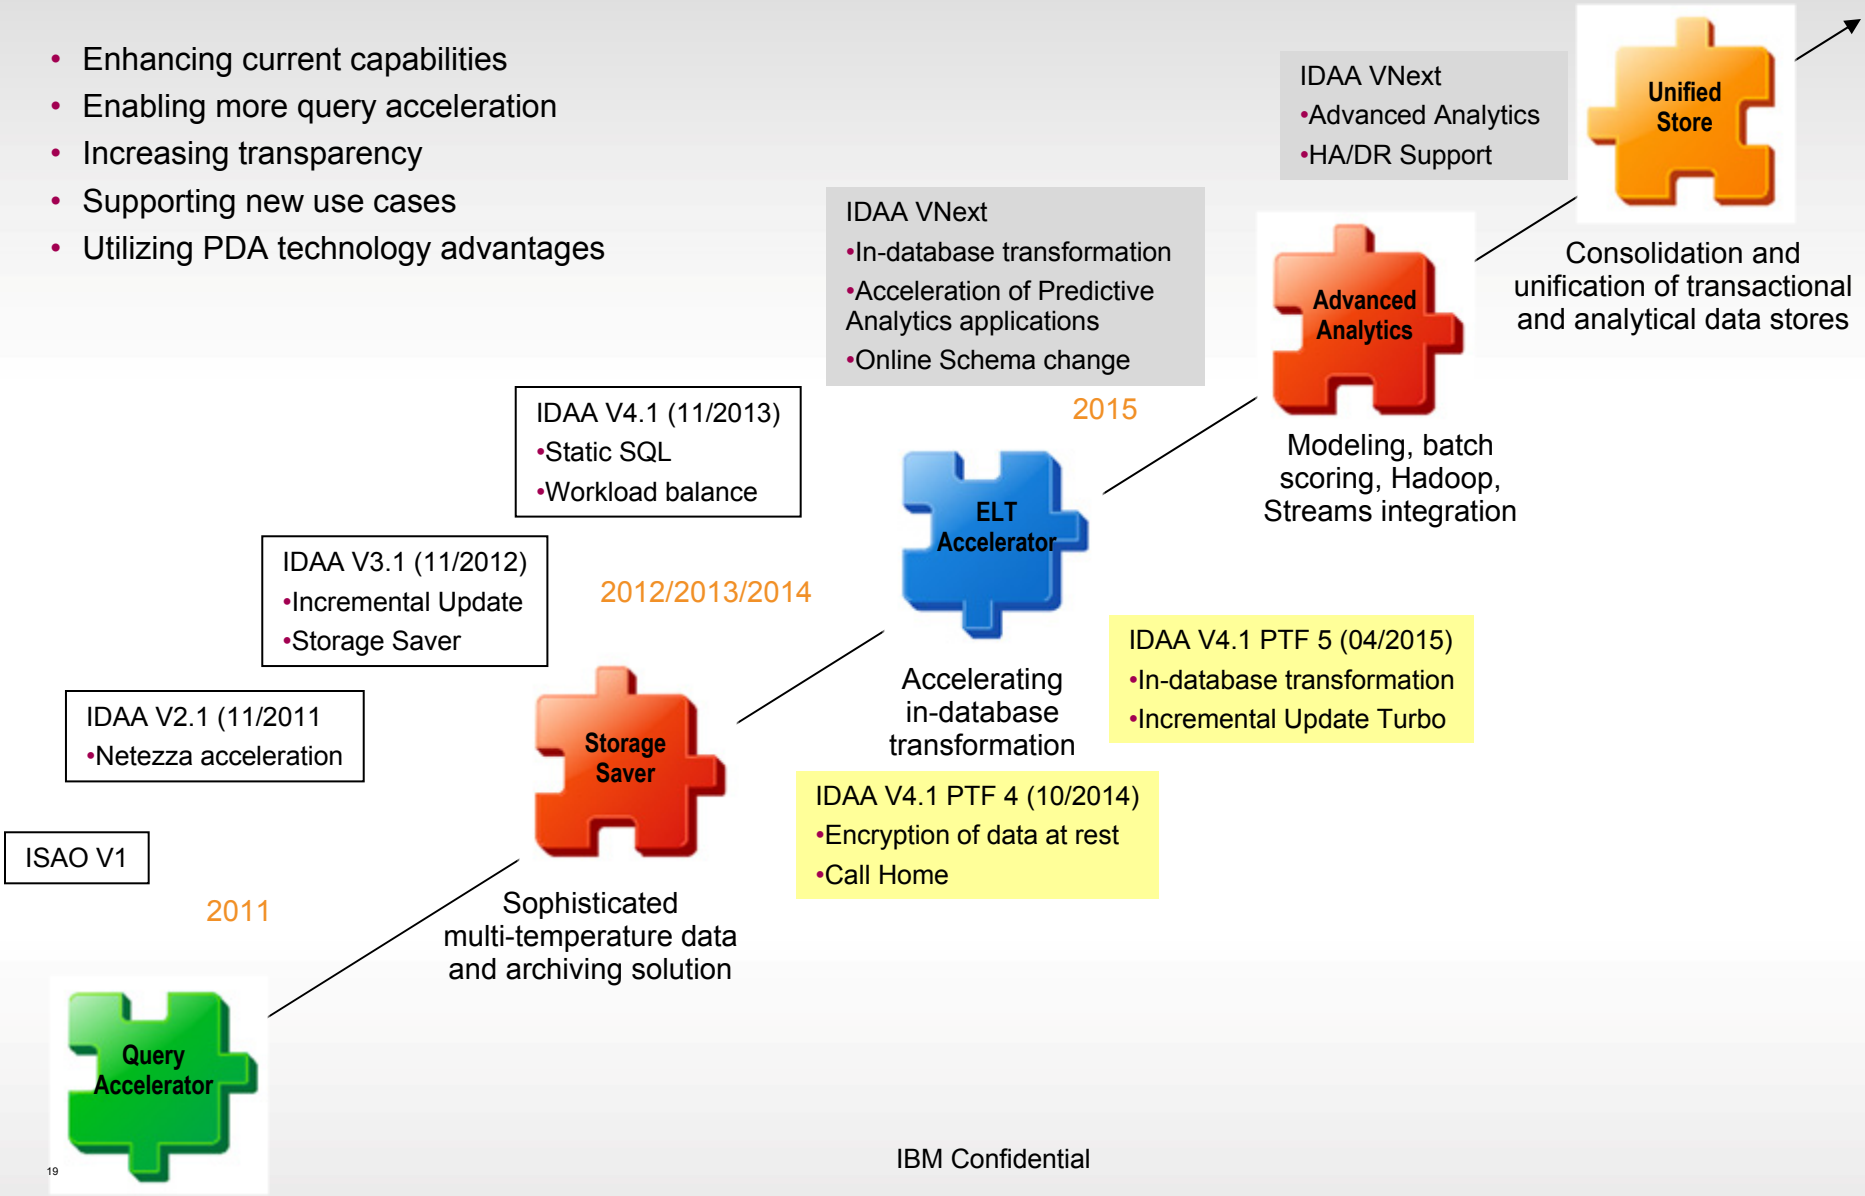

IBM Confidential

Business Analytics Solutions on zEnterprise

Business analytics capabilities	Business outcomes/benefits
Cognos Business Intelligence (BI) 	▪ Understand your current & potential state ▪ Monitor results & fine-tune your business ▪ Inform strategy with a view into the future
SPSS Predictive Analytics (PA) 	▪ Predict customer segment & category affinity ▪ Market Basket Analysis to identify the next best offer for the customer ▪ Overlay browsing history onto purchase history to profile customers
TM1 Performance Management (PM) 	▪ Reporting, analysis, operational & financial planning and consolidation ▪ Product profitability solutions across customers, business lines & channels ▪ Sales Performance Management to improve pay-for-performance programs & efficiency in incentive compensation processes

In-database scoring for both historical data and real-time transactional data

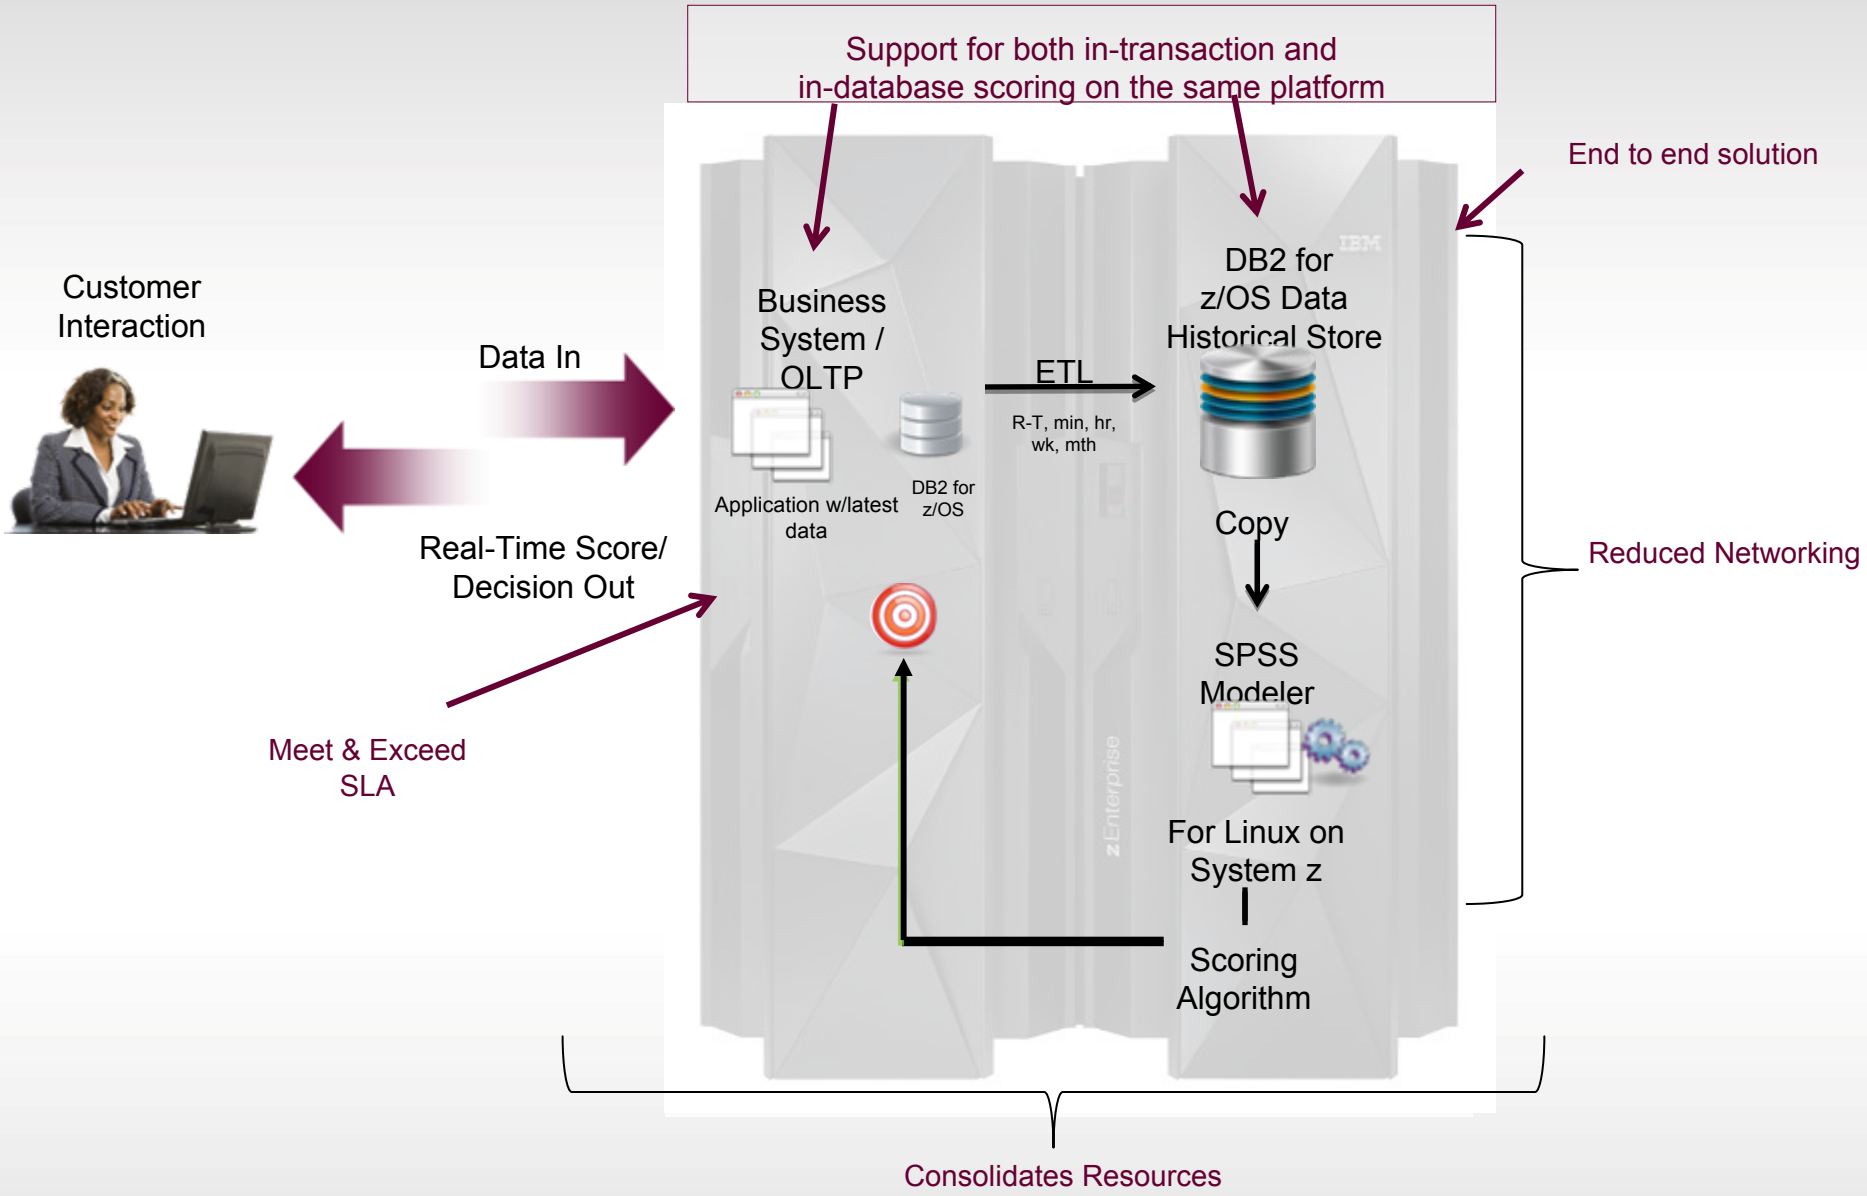

System z: A Hybrid Transaction and Analytics Processing (HTAP) platform

A single workload-optimized system for accelerating decisions to the speed of business

Everything is online – analytics in the right place

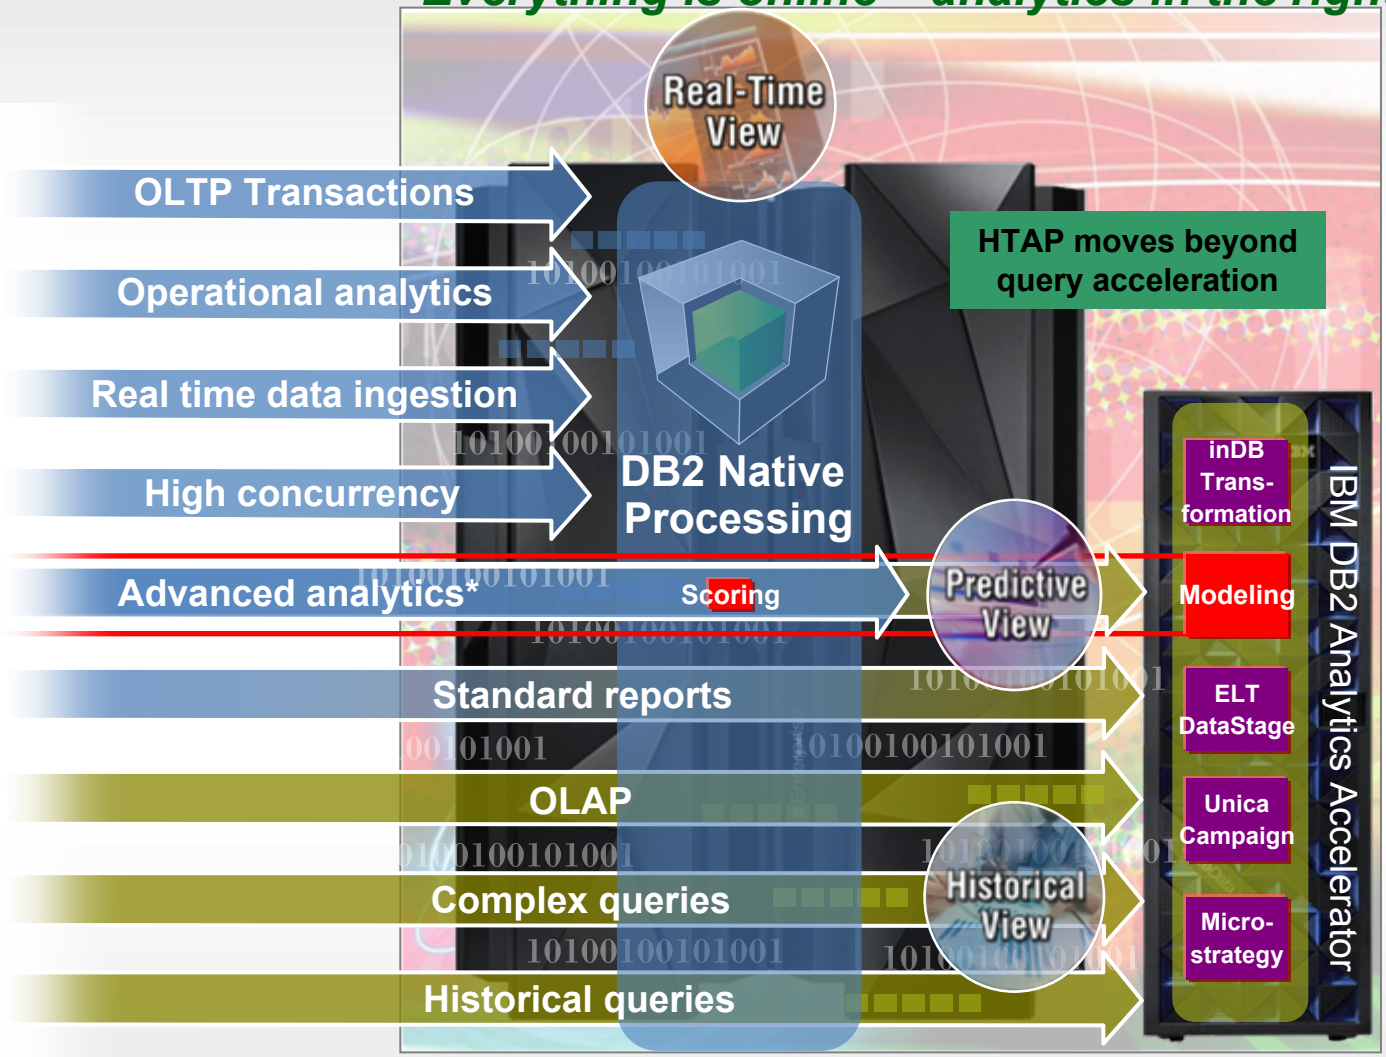

More timely insights from data

- High-speed analytics easily integrated into operational applications
- Historical views are quickly analyzed for more train-of-thought analysis
- Decision makers can perform business analysis they never dared in the past
- Secured environment for highly sensitive data
- Speed operation reporting cycle to meet stricter service level agreements
- Configuration simplification

Operational benefits

- Single point of entry
- Reduced data movement
- High fidelity data
- Dynamic routing for most efficient fit for purpose execution architecture
- Single environment for security, logging, back-up, and recovery
- Competitive price/performance

Imagine the possibility of leveraging all of your data assets

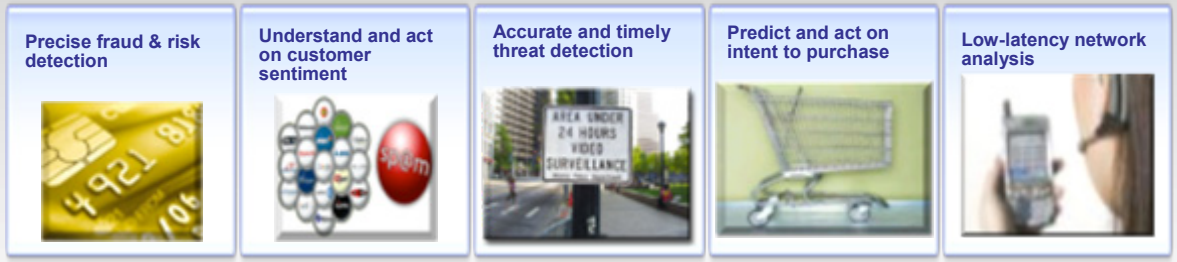

Traditional Approach

Structured, analytical, logical

Data
Rich, historical, private, structured
Customers, history, transactions

The "Circle of Trust"

Data warehouse & business analytics moving closer to this data

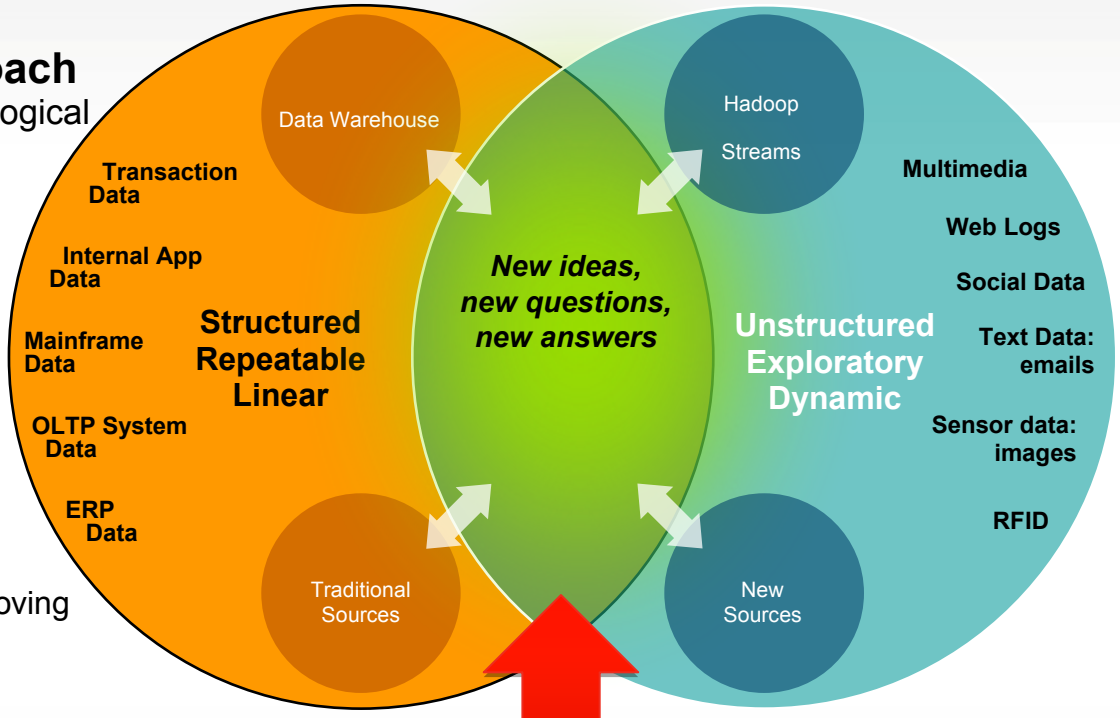

New Approach

Creative, holistic thought, ir

Data: Intimate, unstructured.
Social, mobile, GPS, web, photos, video, email, logs

The real benefit is derived from integration of new data sources with traditional corporate data

- How can you query across both realms?
- How can you preserve security and lower TCO?
- How can you avoid costs and risks of offloading?

InfoSphere BigInsights for Linux on System z - 2.1.2

GA'd Aug 5 2014

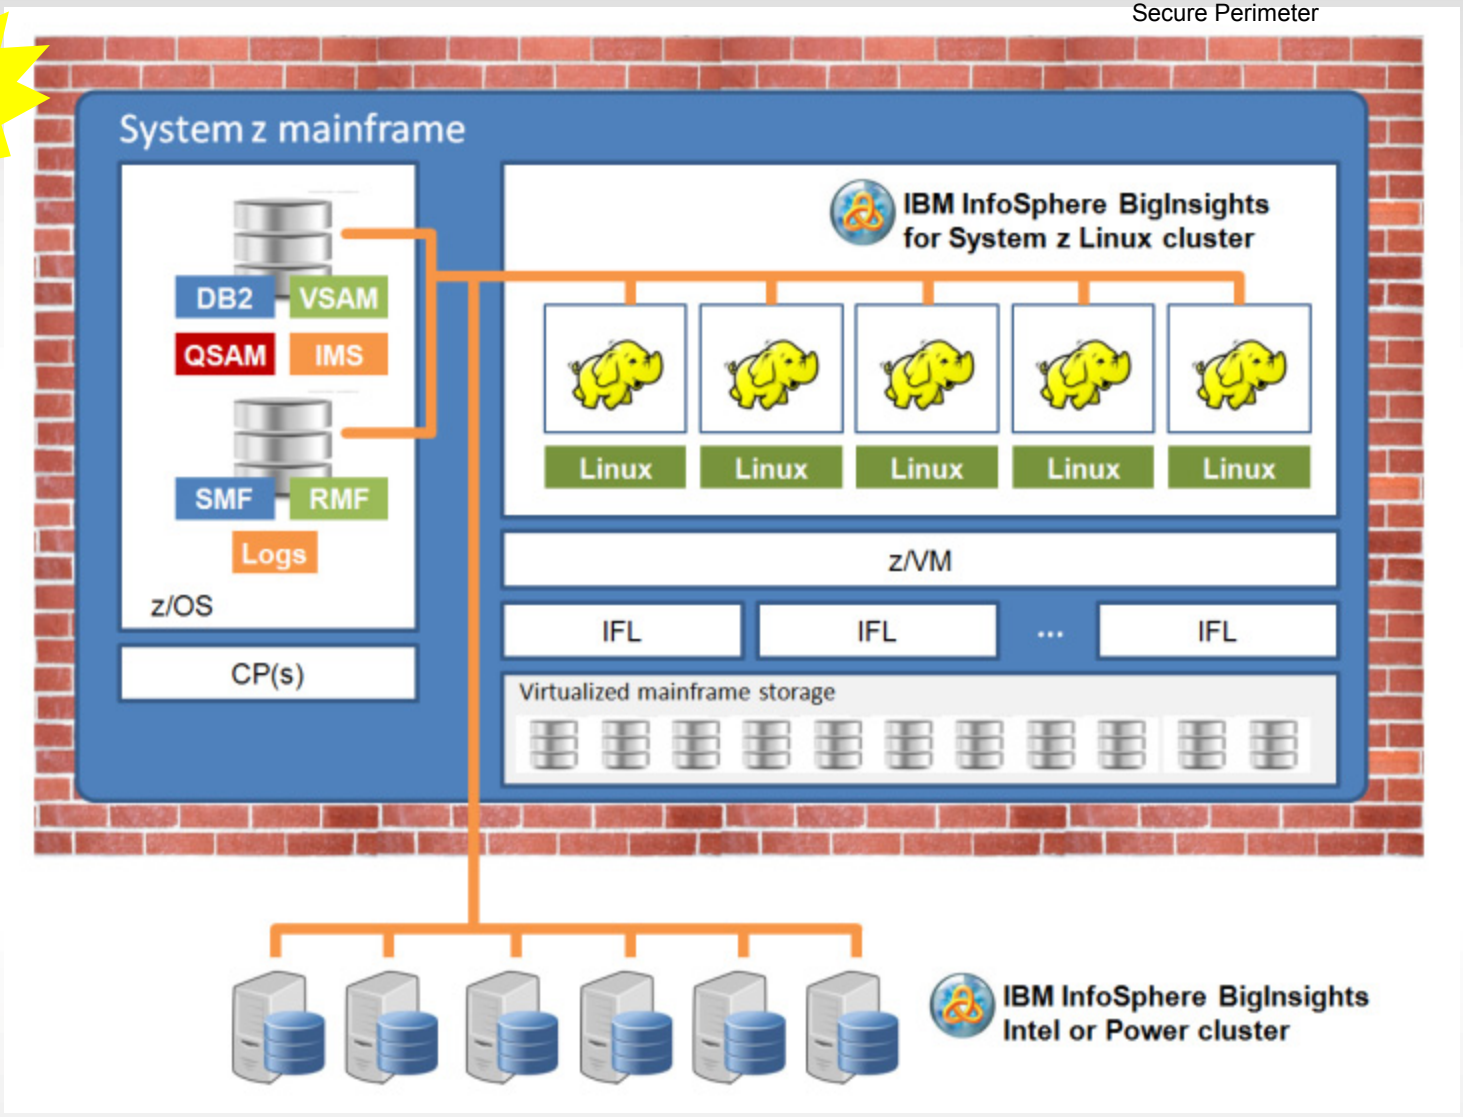

InfoSphere BigInsights 2.1.2 for Linux on System z

What's included?

- Enterprise Edition only offered on System z Linux
- Available from Fix Central August 5th, 2014
- GA release based on **BigInsights 2.1.2**
- Support RHEL 6.4
- Apache MapReduce and HDFS only (no GPFS FPO or Adaptive MapReduce)
- Pricing identical to System x and Power – per node/VM licensing, perpetual
- 5-6 nodes recommended per IFL
- While all components shown are entitled, some may need to be run on other platforms in this initial release

Adaptive MapReduce ad GPFS FPO are not supported in this release of BigInsights Enterprise Edition on System z.
Other components including Streams and Data Explorer have components that will need to run on other platforms

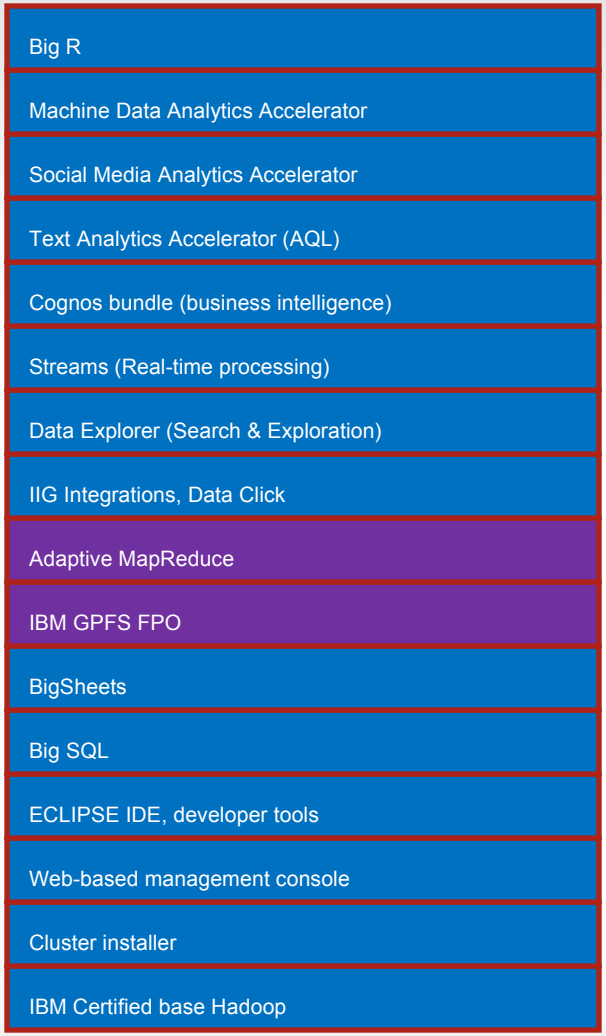

Veristorm – vStorm Connect

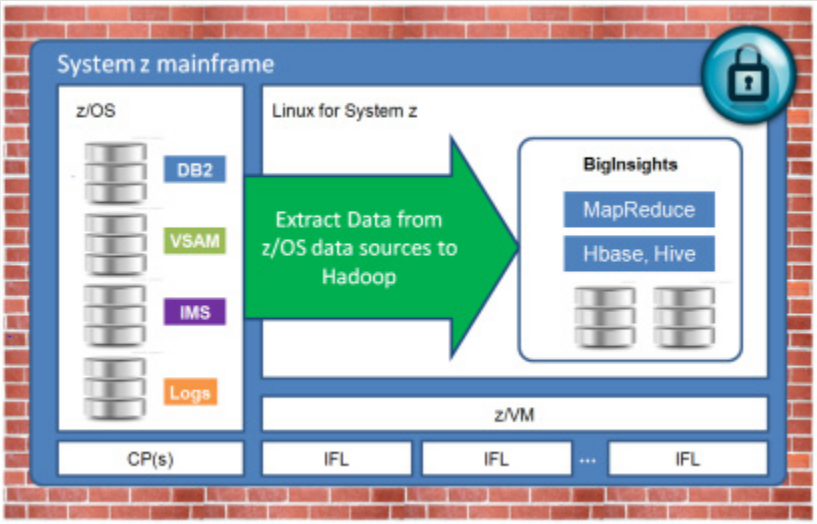

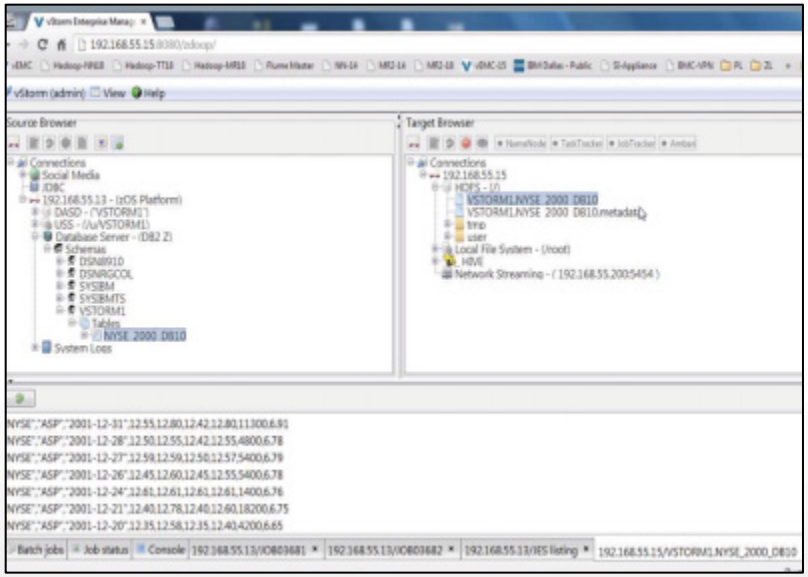

Easy to use ingestion engine

- Native data collectors accessed via graphical interface
- Light-weight; no programming required
- Multiple z/OS data sources
- Conversions handled automatically
- Streaming technology leverages USS (no z/OS engines) with no DASD required for staging

A secure pipe for data

- RACF integration – no need for separate or special credentials
- Data streamed over secure channel using hardware crypto
- Combining with BigInsights for Linux on System z means data never leaves the box

Mainframe efficiencies

IBM System z Transactional/Analytic Environment

Delivering real-time insight for proven competitive advantage

Helping IT Drive Innovation

- **Improved bottom line results** driven from real-time business critical insight
- **Highly competitive decision making** derived from extremely timely, accurate analytic information
- **Rapid IT execution** against new & changing business demands though a proven, integrated platform
- **Minimized risk** with superior security and data governance technologies
- **Added value realized out of current investments** for reduce total cost of ownership

End to end solution for real-time enterprise-wide insight

- Transactional Data/Applications
- Data Warehousing
- Rapid query acceleration
- Business Intelligence
- Predictive Analytics (including real-time in-transaction scoring)
- Performance Management

Business critical analytics at the point of customer impact

- Timely, accurate, secure data
- Proven availability, scalability, performance

Modern, flexible platform

- Rapid deployment & expansion
- Start from your top requirement
- Grow without re-architecting
- Reduce cost and complexity

Integration with Big Data

- Volume, Velocity, Varity & Veracity
- Simplified integration of hadoop, streaming & transactional data

Thank You

zEnterprise Analytics System 9700 and 9710

Flexible zEnterprise Deployment Models

A cost-competitive, integrated combination of hardware, software and services to deliver business reporting and business critical analytics

- **Solution Priced** - for deployment as an additional logical partition (LPAR) on an existing system or as a new system
- **Preselected** – to deliver a comprehensive, yet flexible end-to-end solution
- **Pretested** – to meet business reporting and critical analytic demands

zEnterprise Analytics System 9700

zEnterprise EC12

zEnterprise Analytics System 9710

zEnterprise BC12