Session 5

Why IBM zEnterprise Analytics, Why Now?

Data Warehousing, Business Intelligence & Predictive Analytics

Presenters name, title

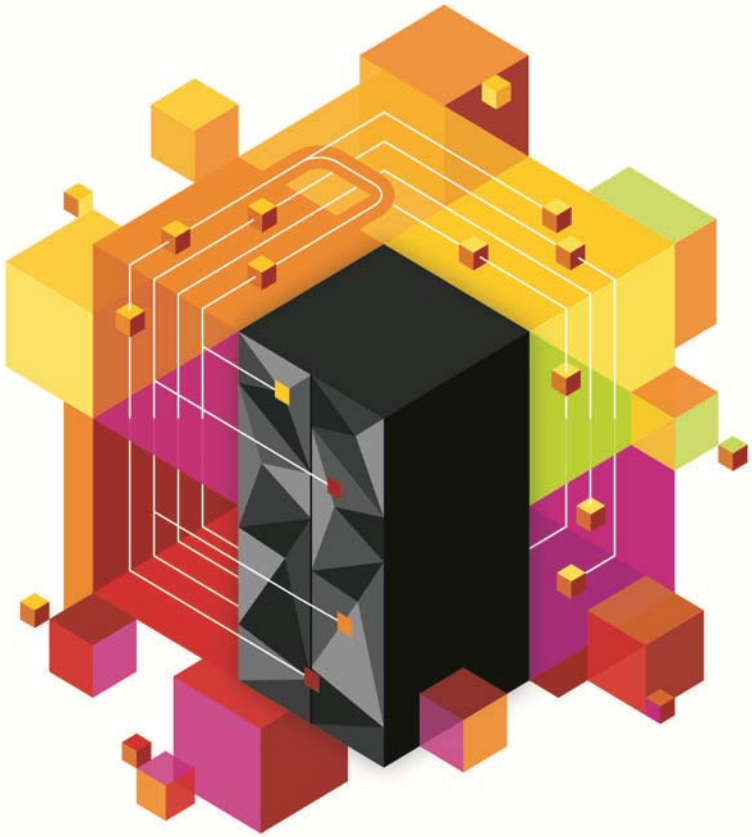

Analytics on zEnterprise delivers…

Ave. **87% savings in CPU, 96% savings in servers** over 5 yrs. for BI deployments

BI servers that run at **90%+ capacity** without impacting SLA

30%- 45% performance improvement with zEC12 over z196

Meet SLAs & score **3000-5000+ transactions in Real-time**

Run complex queries up to **2000x Faster**

80% less capacity for Data Warehousing

95% savings in host disk space for historical data

BI system **admin savings alone, pays for the HW investment** in 5 years

Organizations are using information to differentiate

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**

Prior Path to Success		Today's Leaders
Sense and Respond	→	Predict and Act
Instinct and Intuition	→	Real-time, Fact-driven
Skilled Analytics Experts	→	Everyone
Back Office Decision Support	→	Point of Impact
Automated Processes	→	Optimized

Source: CIO Survey

Traditional Analytics Market

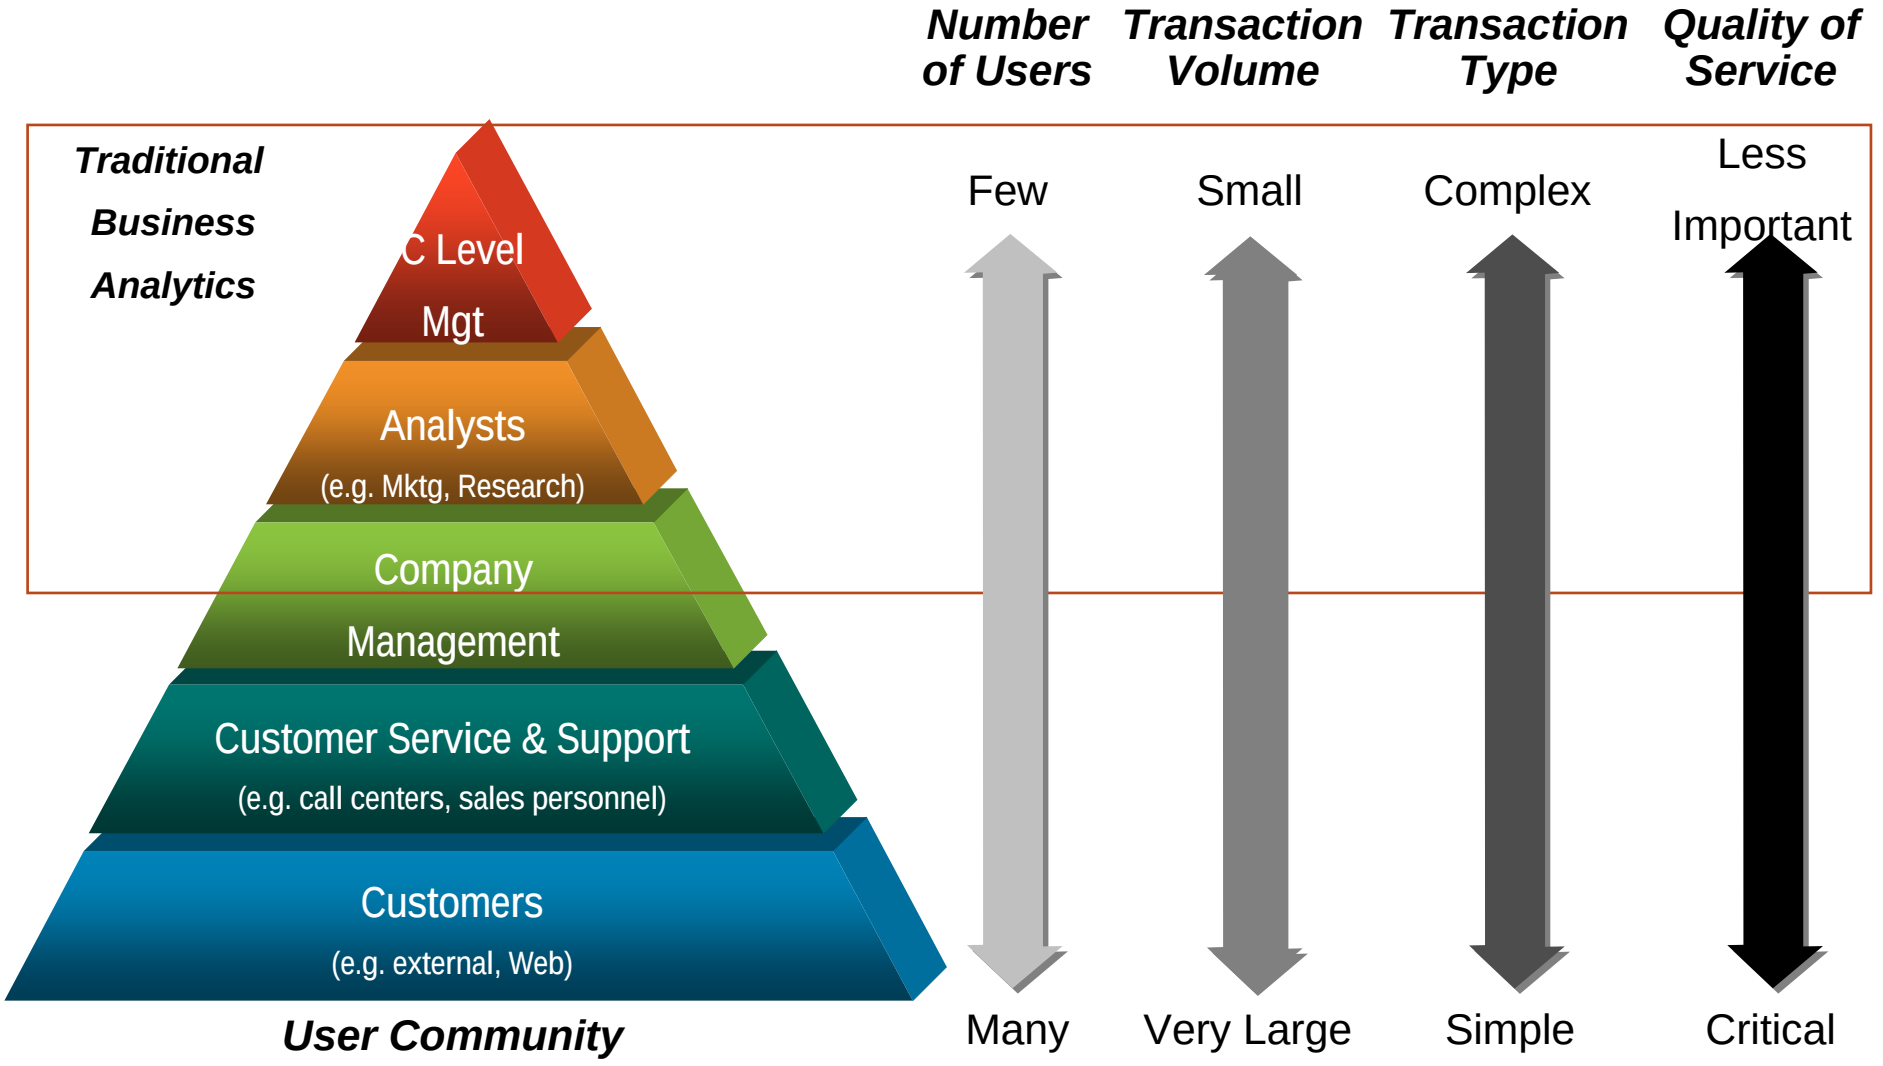

Analytic requirements have expanded

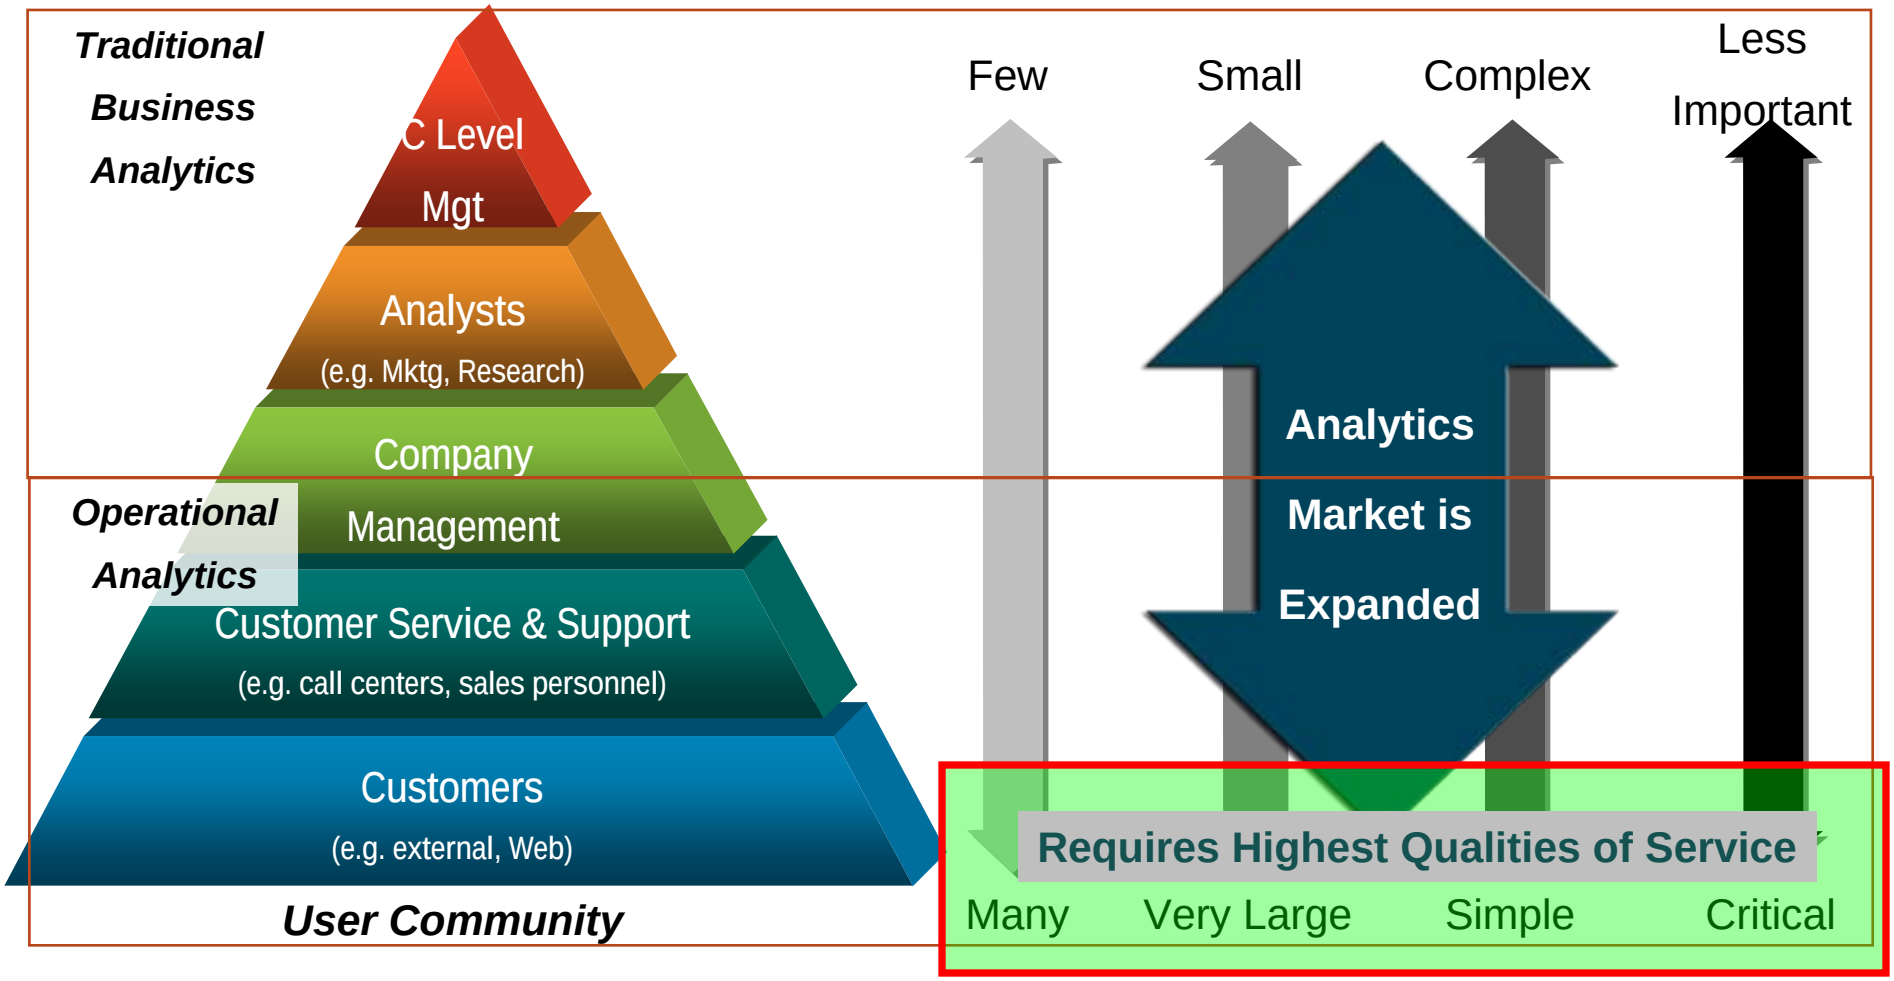

IBM zEnterprise

Creating business insight where the data originates

- Delivers continuous availability for OLTP and analytics on a single, integrated platform
  - As little as 5 minutes unplanned downtime per year
- A large % of the data that is accessed for analytics resides or originates on IBM zEnterprise
  - 66% of business transactions for U.S. retail banks
  - 80% of world's corporate data
- Businesses that run on zEnterprise
  - 25 of the top 25 worldwide banks
  - 23 of the top 25 U.S. retailers
  - 9 of the top 10 global life/health insurance providers
  - 64% of Fortune 500
  - 45% of Fortune 1000
  - 71% of Fortune Global 500
- The ability to virtualize more than1000 business applications on a single machine
- 1,300+ ISVs run zEnterprise today, with more than 275 of these selling over 800 applications on Linux

Traditional Approach to workload optimized Systems

Operational Applications	Analytic Applications
Transaction Processing	Data warehousing
	
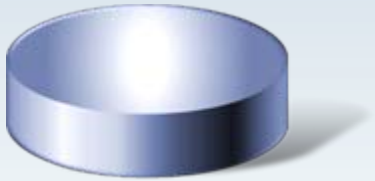	
Shared Everything DB	Shared Nothing DB
High volume business transactions and batch reporting running concurrently	Low volume complex queries

The zEnterprise Hybrid Solution

Mixed Workloads for Next Generation Business Analytics

Mixed Workload Applications

Operational Business Analytics

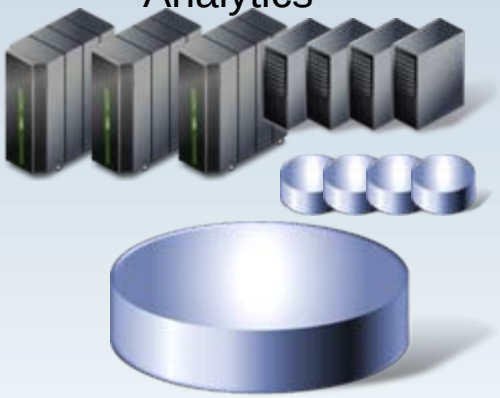

Hybrid DB

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

IBM zEnterprise

an enterprise information hub providing an end-to-end, integrated foundation for modern analytics

Business Analytics
- Business Intelligence
- Predictive Analytics

Data Warehousing
- Data Warehouses
- Operational Data Stores
- Data Marts

Business System / OLTP

Insight Out

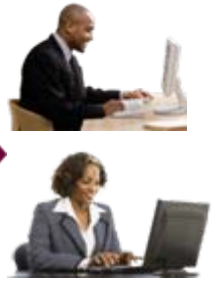

Customer Interaction

Data In

Real-Time Score/ Decision Out

DB2 10 for z/OS

A highly tuned database for OLTP & Analytics

- CPU reductions for transactions, queries, and batch
- Scales with less complexity and cost
- Improved operational efficiencies and lower administration cost
- Even better performance

➢64 bit Evolution Virtual Storage Relief

➢Temporal Data

➢Integrated XML Support

➢Query Processing Enhancements

➢Business Security & Compliance

➢Better Productivity

DB2 Query Management Facility (QMF) 10

- Executive dashboards & significantly enhanced visual reports
- New QMF content remains fully compatible with existing QMF objects
- Rapid development and deployment enterprise-wide solutions
- Lightweight installation and administration
- Minimal learning curve - zero coding, drag-drop authoring model
- Embeddable BI – can be integrated into web and Java apps
- Database-based licensing model – not user or application server-based
- 150 new BI and analytic functions

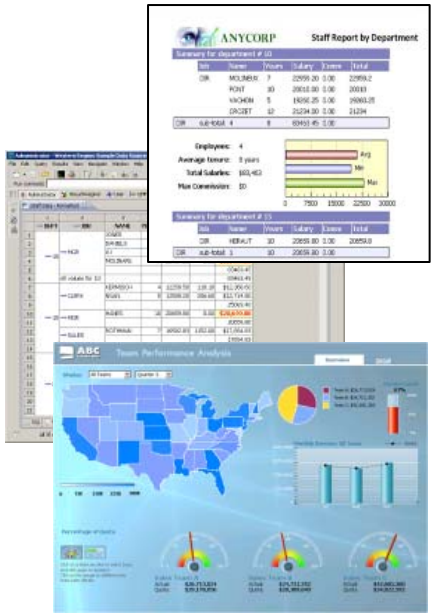

DB2 Analytics Accelerator for z/OS

Blending zEnterprise and Netezza technologies

A high performance analytics accelerator appliance for IBM zEnterprise, delivering dramatically faster complex business analysis transparently to all users.

Fast

Complex queries run up to 2000x faster while retaining single record lookup speed

Cost Saving

Eliminate costly query tuning while offloading complex query processing

Appliance

No applications to change, just plug it in, load the data, and gain the value

DB2 Analytics Accelerator V3

Further extending the features

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

More insight from your data

- Unprecedented response times for “right-time” analysis
- Complex queries in seconds rather than hours
- Transparent to the application
- Inherits all System z DB2 attributes
- No need to create or maintain indices
- Eliminate query tuning
- Fast deployment and time-to-value

IBM DB2 Analytics Accelerator Product Components

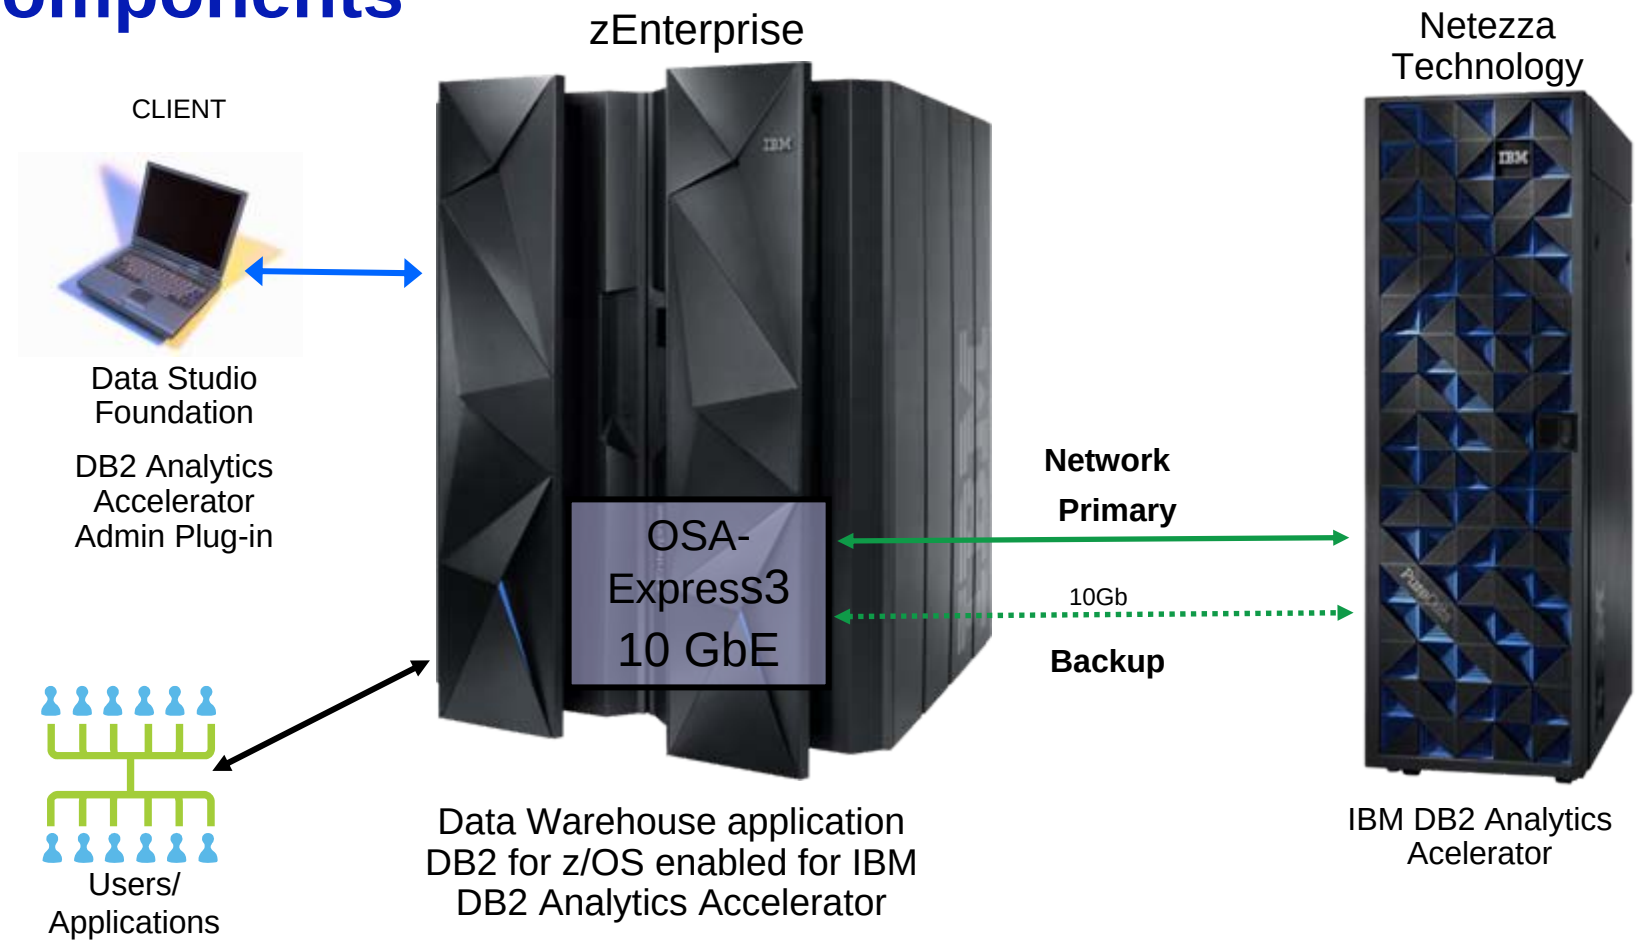

Note: There are several connection options using switches to increase redundancy

Deep DB2 Integration within zEnterprise

Applications

Application Interfaces
(standard SQL dialects)

DBA Tools, z/OS Console, ...

Operational Interfaces
(e.g. DB2 Commands)

DB2 for z/OS

Data Manager

Buffer Manager

. . .

IRLM

Log Manager

IBM DB2 Analytics Accelerator

Superior availability reliability, security, Workload management

z/OS on System z

Superior performance on analytic queries

Query Execution Process Flow

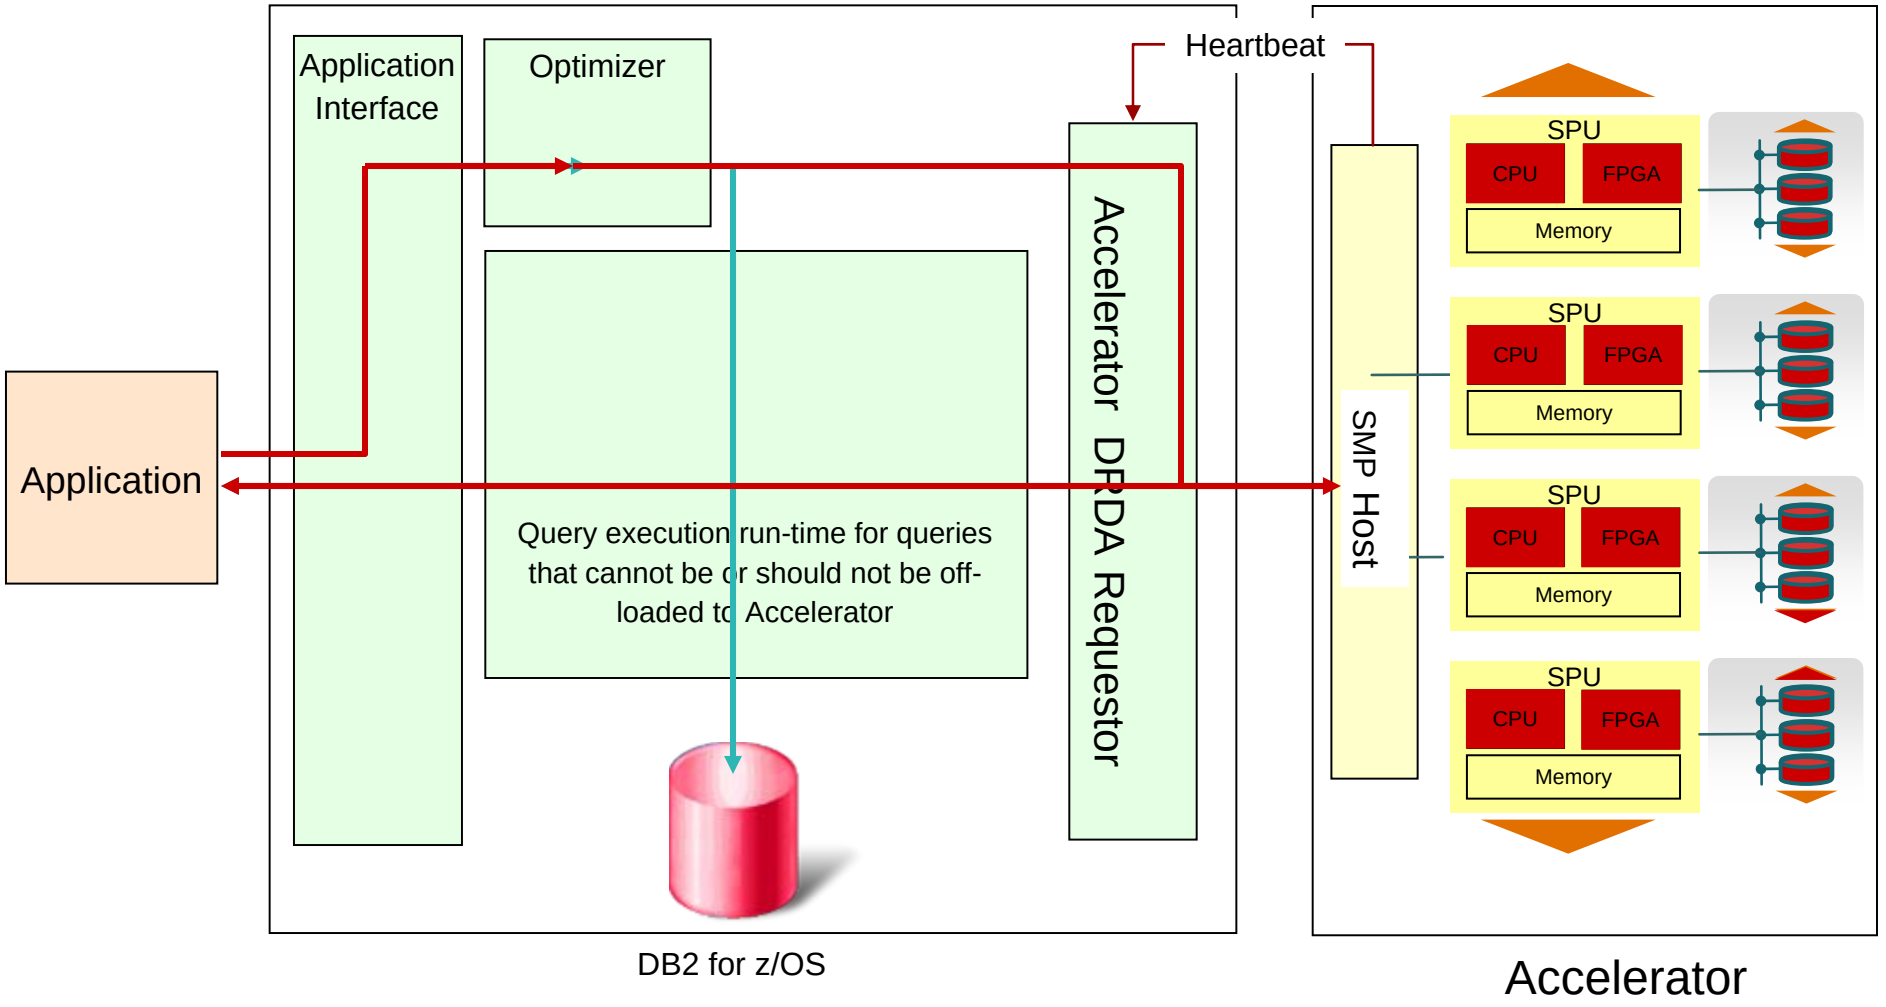

What's new in DB2 Analytics Accelerator V3

- ✓ *Lowering the cost of historical data*
- ✓ *Better decisions through lower latency of data*
- ✓ *Dramatic improvement in scale and growth opportunities*
- ✓ *Lowering the cost of analytic computing*

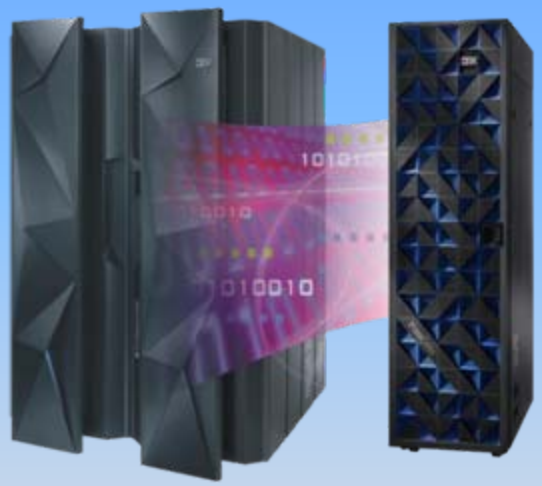

- **High Performance Storage Saver**
- Significantly reduces the cost for storage resources
- Option to store data only once: in the accelerator
- **Incremental Update**
- Data changes are propagated to the accelerator as they happen
- Uses change data capture technology
- Extends the accelerator use to reporting on operational data
- **New optimization**
- Tables or partitions refresh much faster and less resources intensive
- Optimized unloading data from DB2
- **High Capacity**
  - Capacity has been extended to 1.28 PB for a single Accelerator
- **New functions**
- More queries eligible for acceleration

High Performance Storage Saver (NEW!)

Reduces the cost of high speed storage

- You can choose the disk location for historical data to reduce host data warehouse storage usage by over 95%
  - When data no longer requires updating, reclaim the DB2 storage
  - Tables can be resident on:
    1. DB2 Only
    2. DB2 and Accelerator
    3. Accelerator Only
  - Special Registers control behavior
    - CURRENT QUERY ACCELERATION
    - CURRENT GET_ACCEL_ARCHIVE
  - Managed by zParms
  - Enhanced Heuristics

Ability to store historic data on the Accelerator only

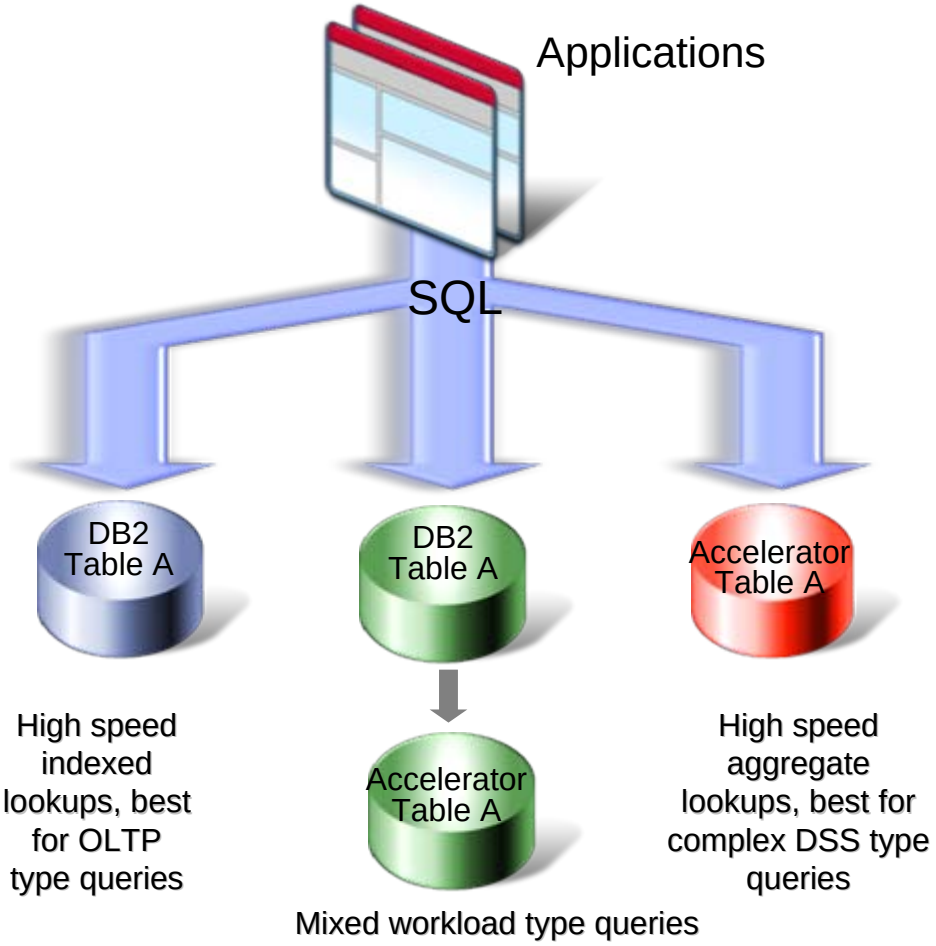

Save Over 95% of Host Disk Space for Historical Data

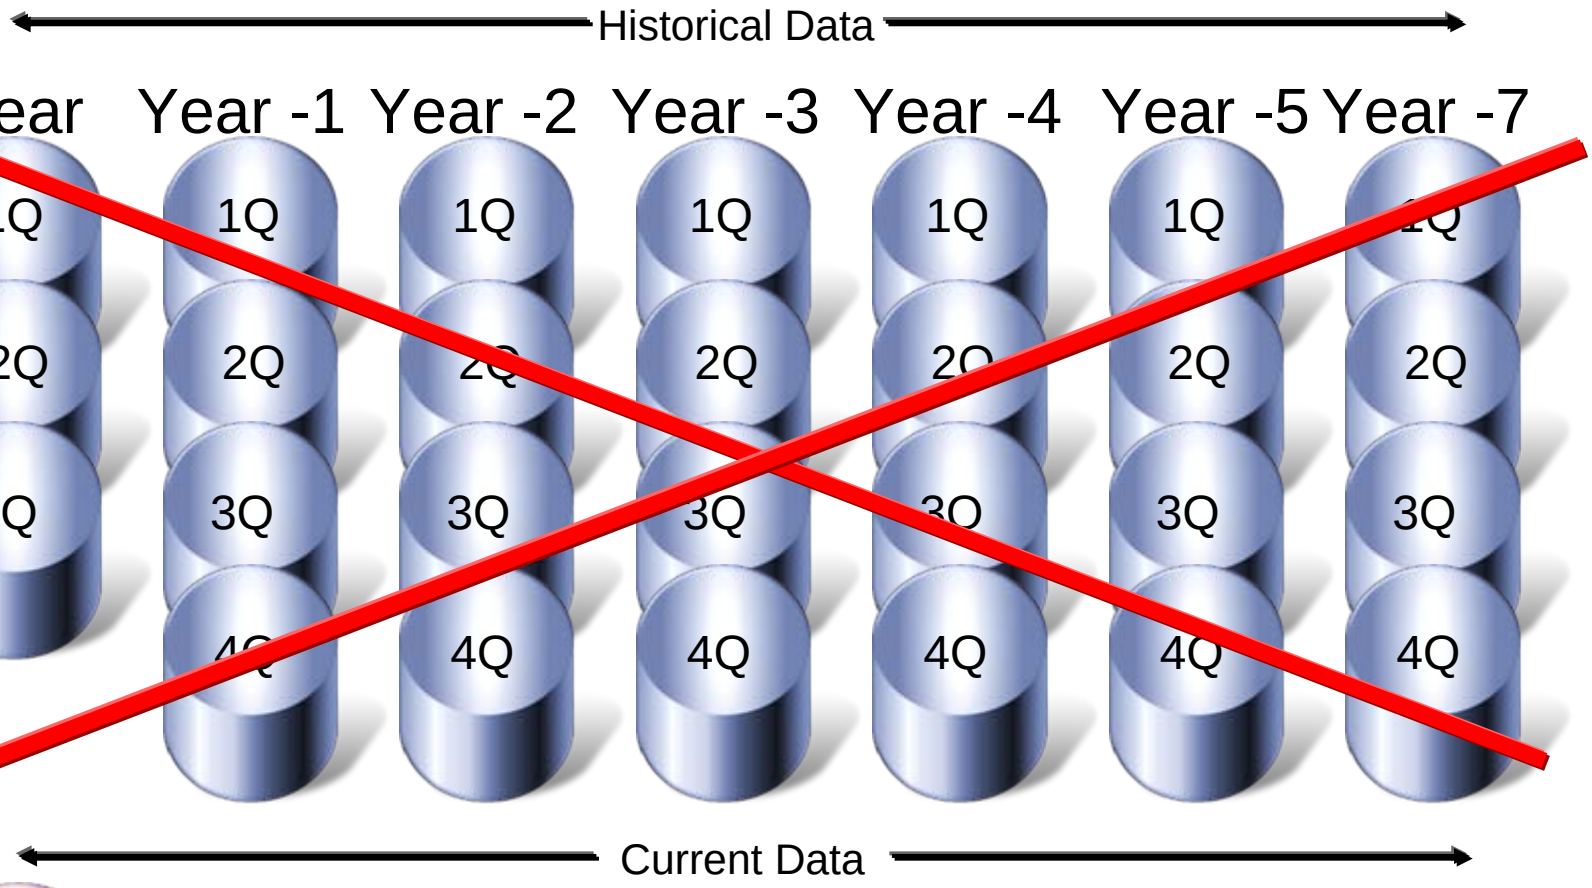

4Q

One Quarter = 3.57% of 7 years of data

One Month = 1.12% of 7 years of data

One month = 2.78% of 3 years of data

Incremental Update

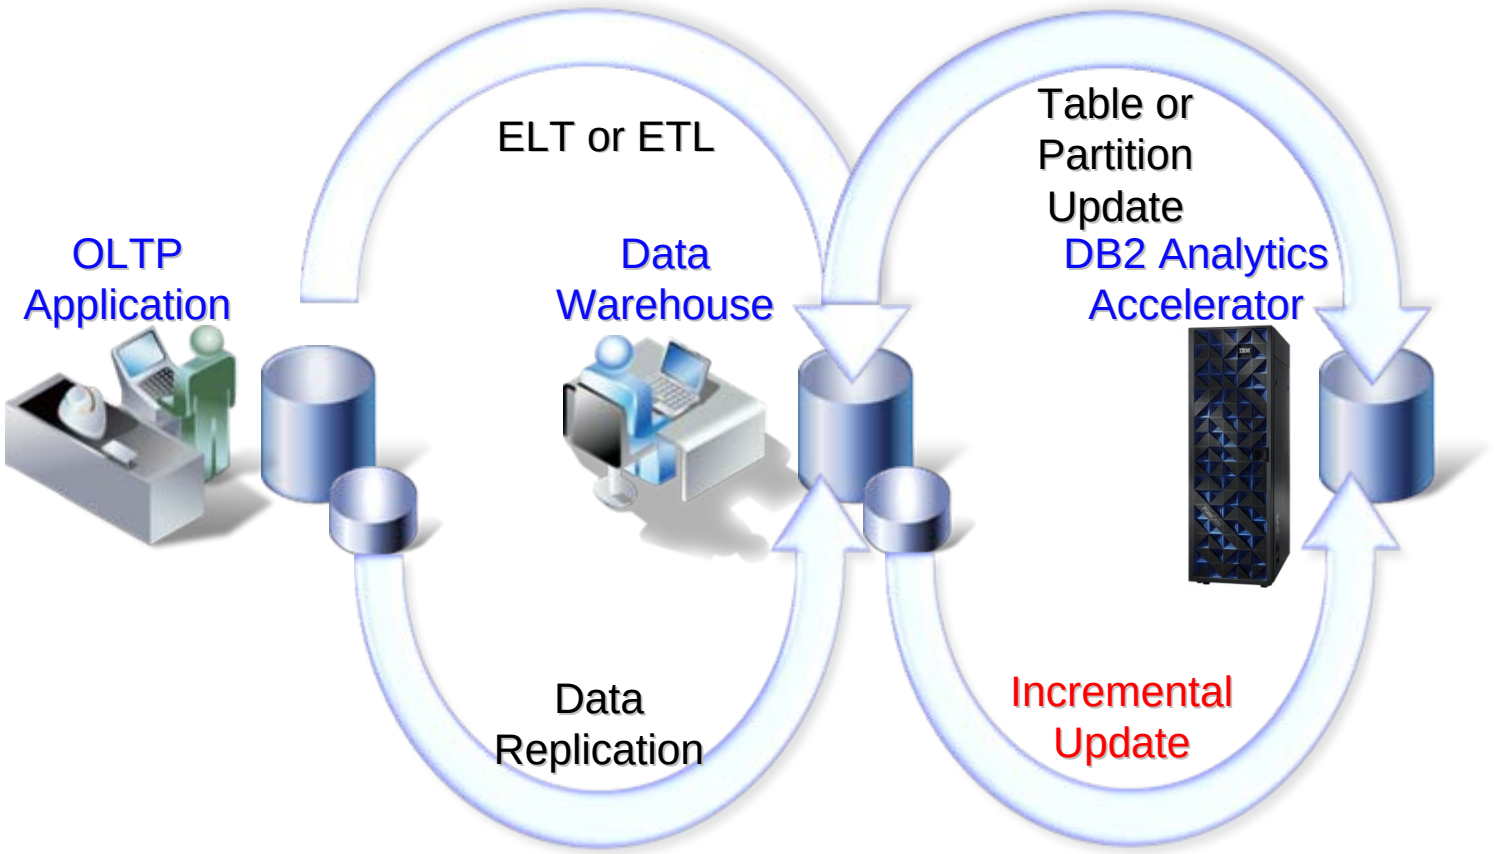

Synchronizing data to lower data latency from days to minutes/seconds

Connectivity Options

Multiple DB2 systems can connect to a single Accelerator

A single DB2 system can connect to multiple Accelerators

Multiple DB2 systems can connect to multiple Accelerators

The same table can be stored in the multiple Accelerators (except High Performance Storage Saver tables)

Better utilization of Accelerator resources

Scalability

High availability

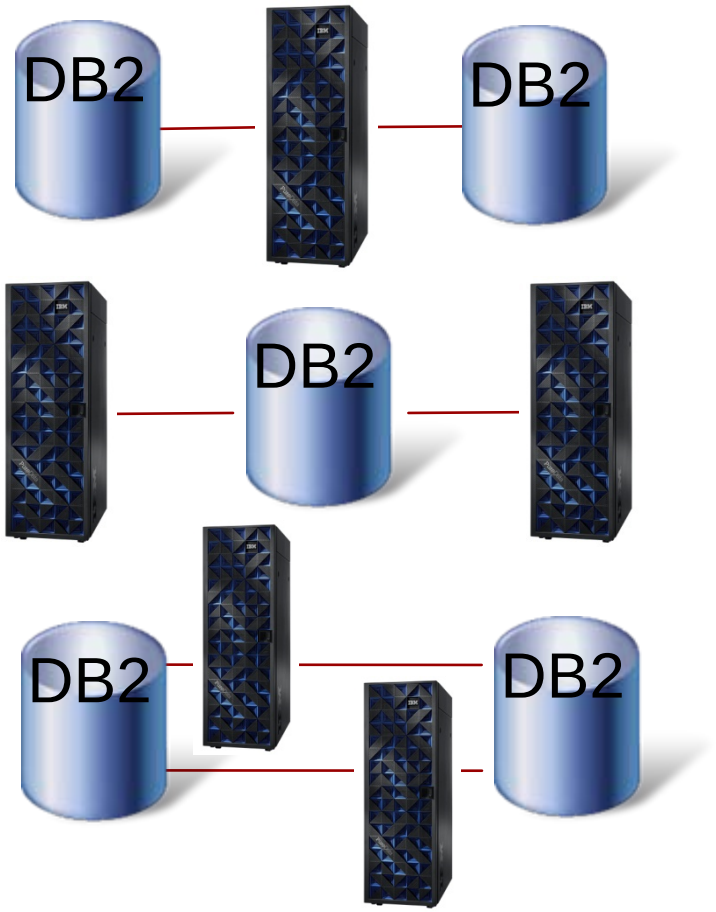

Full flexibility for DB2 systems:

- residing in the same LPAR
- residing in different LPARs
- residing in different CECs
- being independent (non-data sharing)
- belonging to the same data sharing group
- belonging to different data sharing groups

21

Introducing the PureData System for Analytics N2001

The fastest performance of Netezza technology to date!

Accelerate Performance of Analytic Queries

- **3X faster performance**[1] *for Big Data analytics*
- **128 GB/sec effective scan rate per rack**[2] *to tackle Big Data faster*

Increase Efficiency of your Data Center

- **50% greater data capacity per rack**[3] *helps optimize data center efficiency*
  - **More capacity and less power per rack** *than both Oracle and Teradata*

Simplicity and Ease of Administration

- **Improved system management and resilience** *to spend less time managing and more time delivering value*

[1] Based on a comparison of the IBM PureData System for Analytics N2001 to the IBM PureData System for Analytics N1001. The performance speed refers to the query times on both macro-analytic and mixed workload tests as conducted in IBM engineering lab benchmarks. The N2001 query times were an average of 3x faster than those of the N1001. Individual results may vary.

[2] 128 GB/sec scan rate assuming an average of 4x compression across the system. Individual results may vary.

[3] Capacity of IBM PureData System for Analytics N2001 compared to previous generation IBM PureData System for Analytics N1001.

What's New?

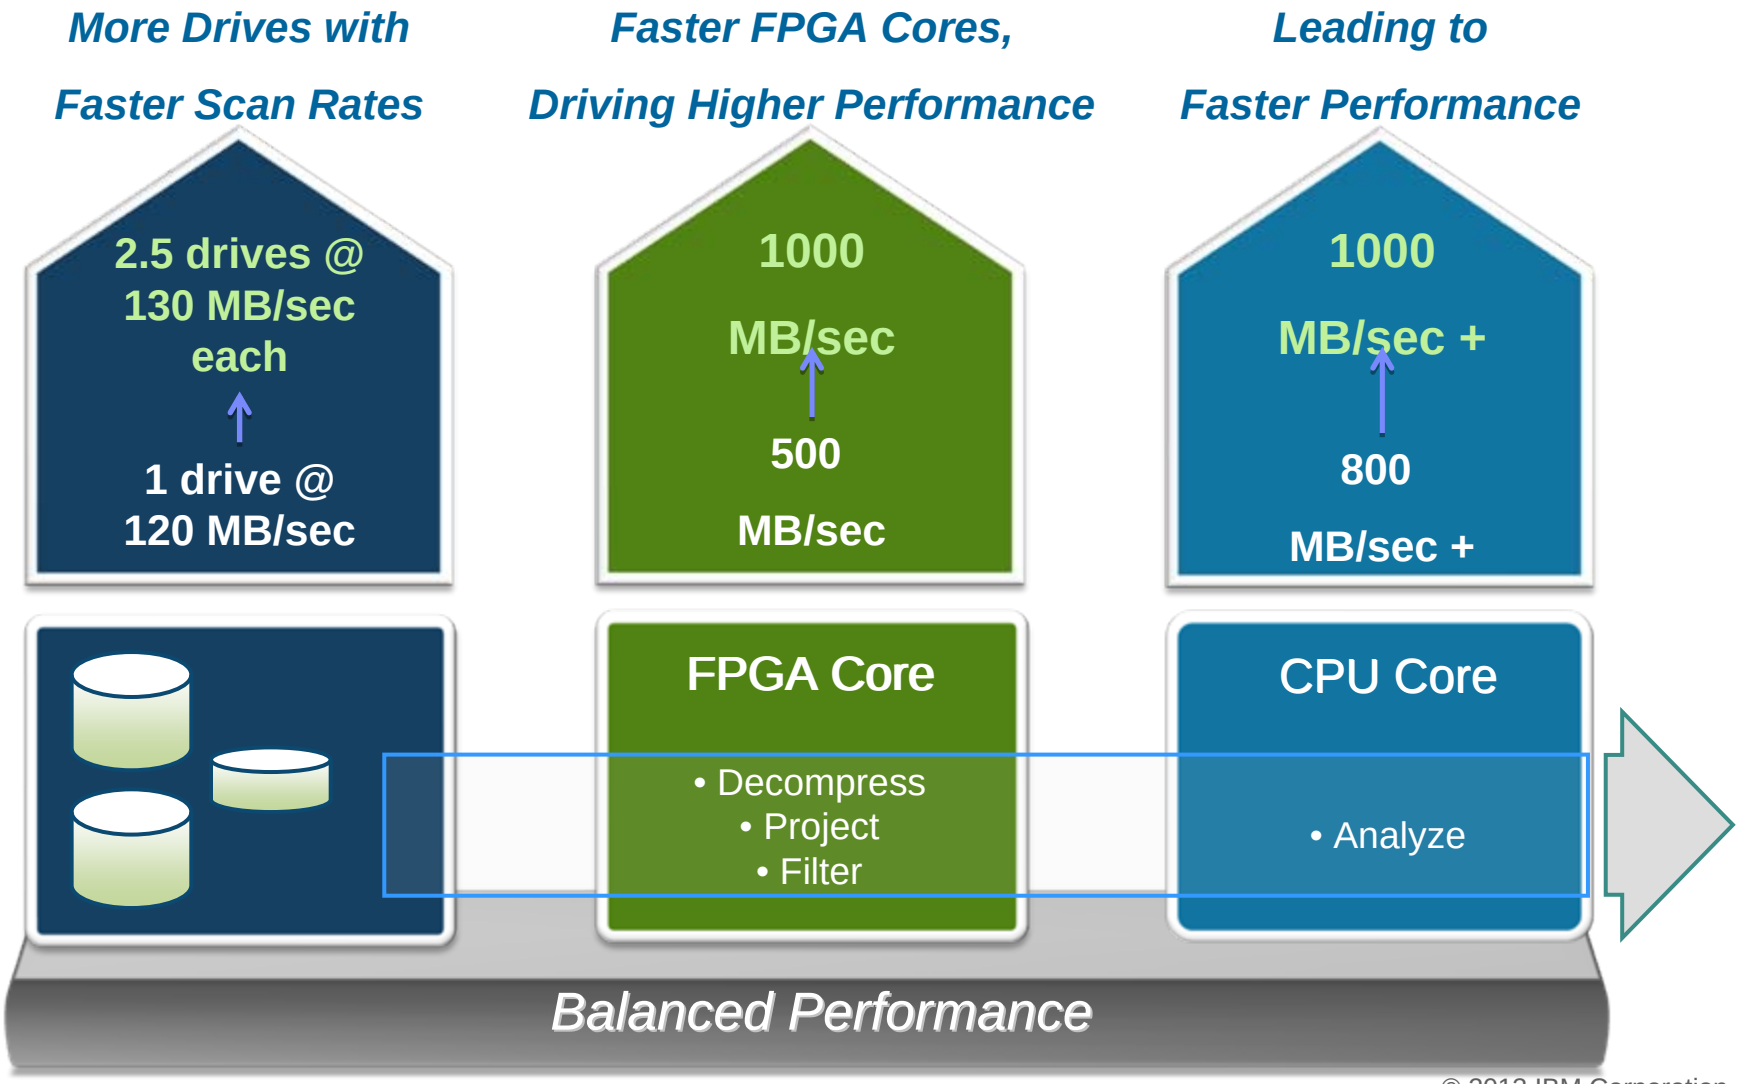

Key to the Speed

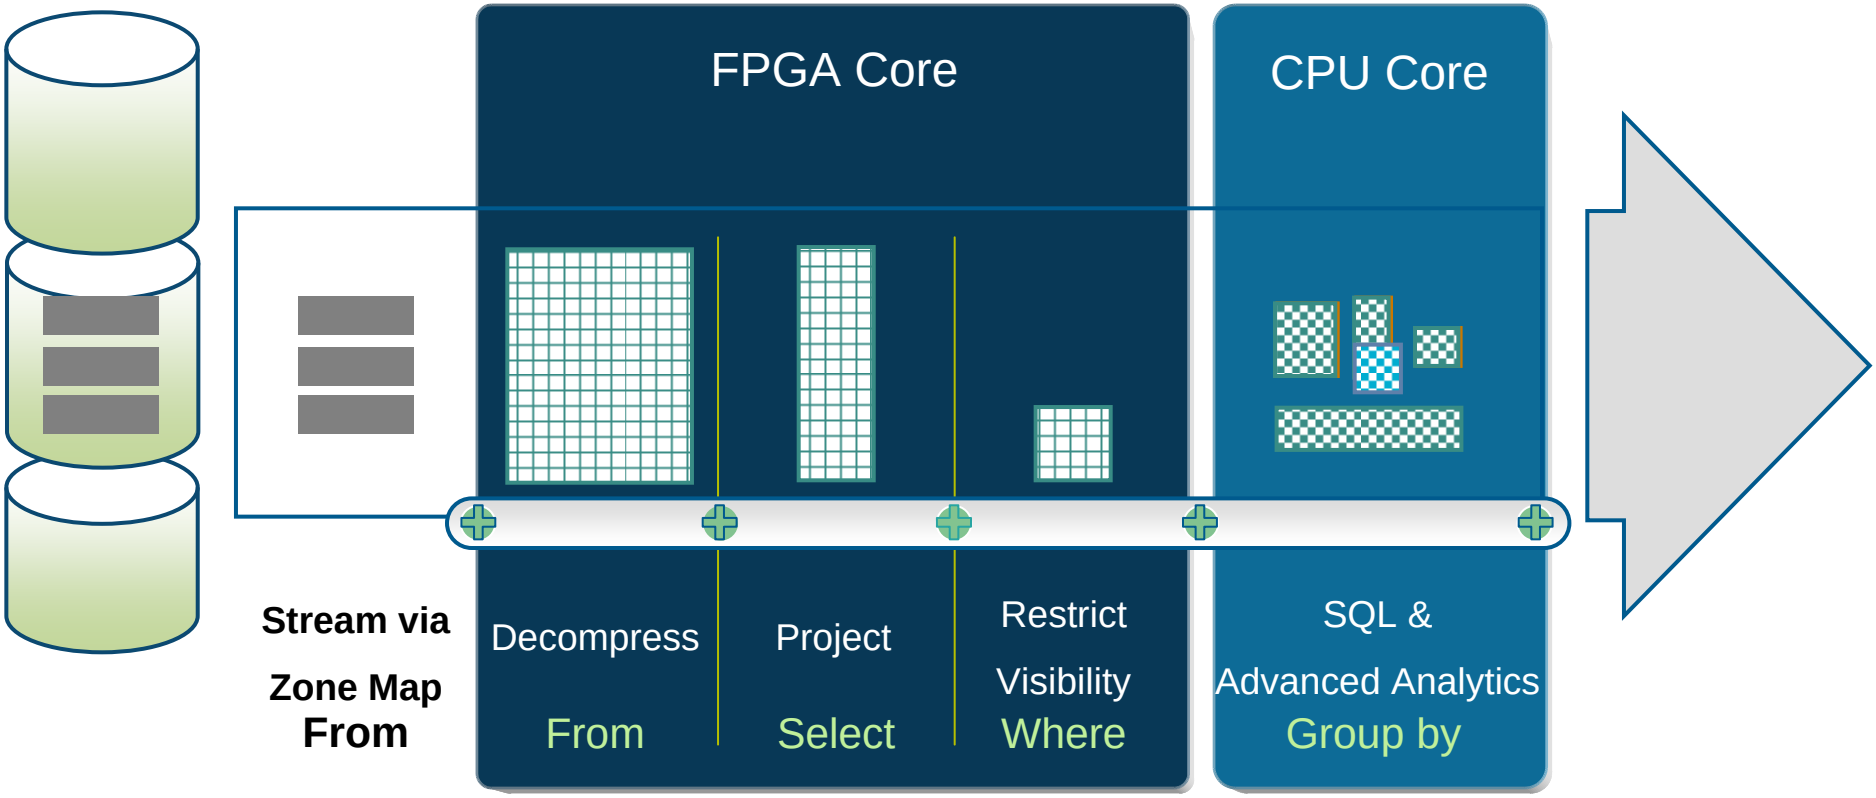

Select State, Age, Gender, count(*) From BLUDB.Customer Where BirthDate > '01/01/1960' And State in ('FL', 'GA', 'SC', 'NC') Group by State, Age, Gender Order by State, Age, Gender

Performance & Savings

270 of the Mixed Workload Queries

Executes in DB2 returning results in seconds or sub-seconds

30 of the Mixed Workload Queries took minutes to hours

			DB2 Only		DB2 with IDAA		Times Faster
Query	Total Rows Reviewed	Total Rows Returned	Hours	Sec(s)	Hours	Sec(s)	
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

Successfully accelerated the problem queries without affecting the rest

IBM zEnterprise Analytics System 9700

Mixed Workloads for Next Generation Business Analytics

The next generation of System z analytics; an integrated solution of hardware, software and services that enables customers to rapidly deploy cost effective game changing

analytics across their business.

Preselected

All the necessary components are identified and integrated into an end-to-end solution

Pretested

Over 20 different customer typical configurations are presized and tested

Solution Priced

Aggressively priced for a cost-effective add-on or new deployment for customers with critical data operations

IBM zEnterprise Analytics System 9700

Flexibility in Critical Data Decision Systems

Organized for simplicity & functionality

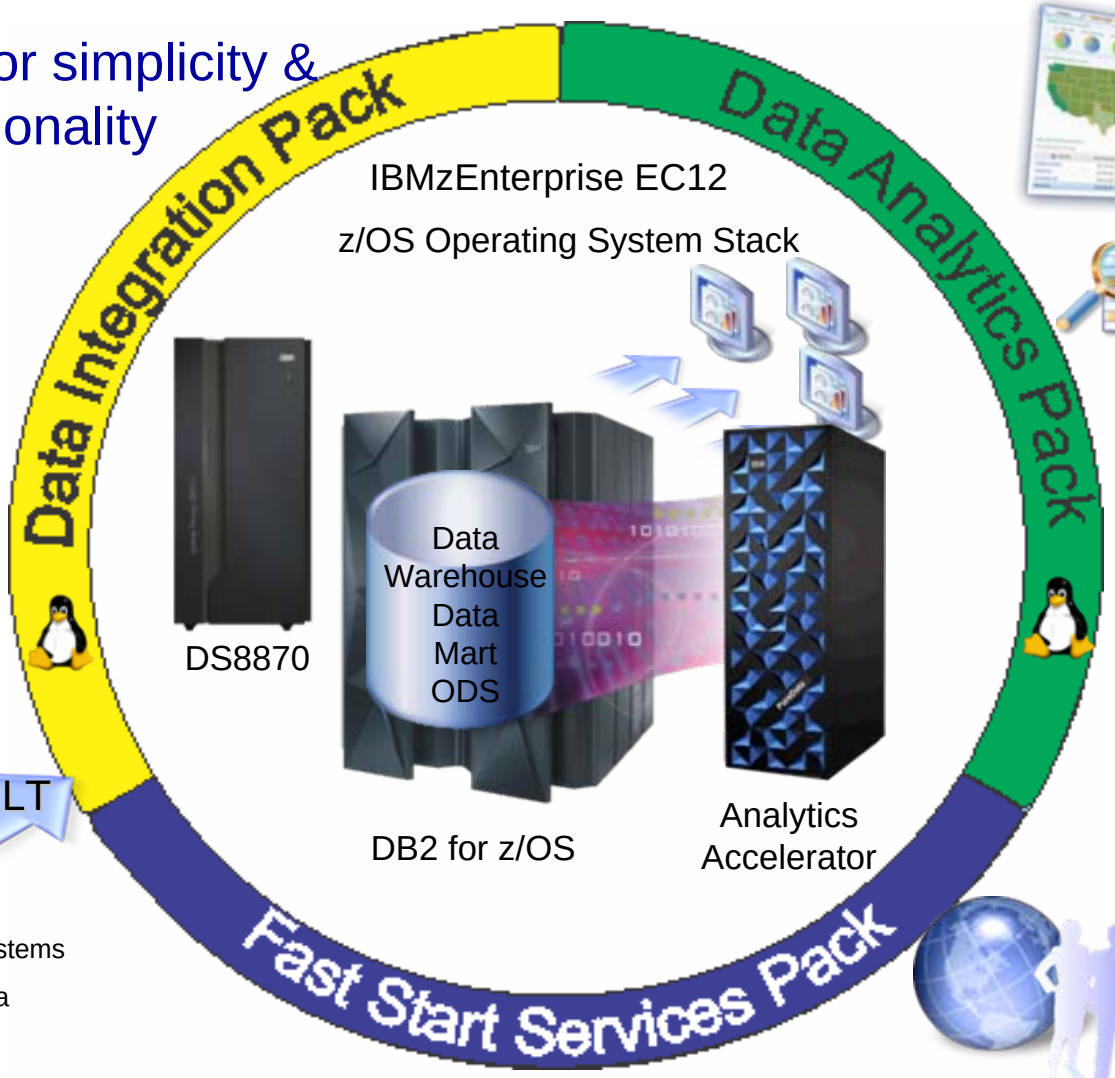

IBM zEnterprise Analytics System 9710

Cost Effective Critical Data Decision Systems

Organized for simplicity & functionality

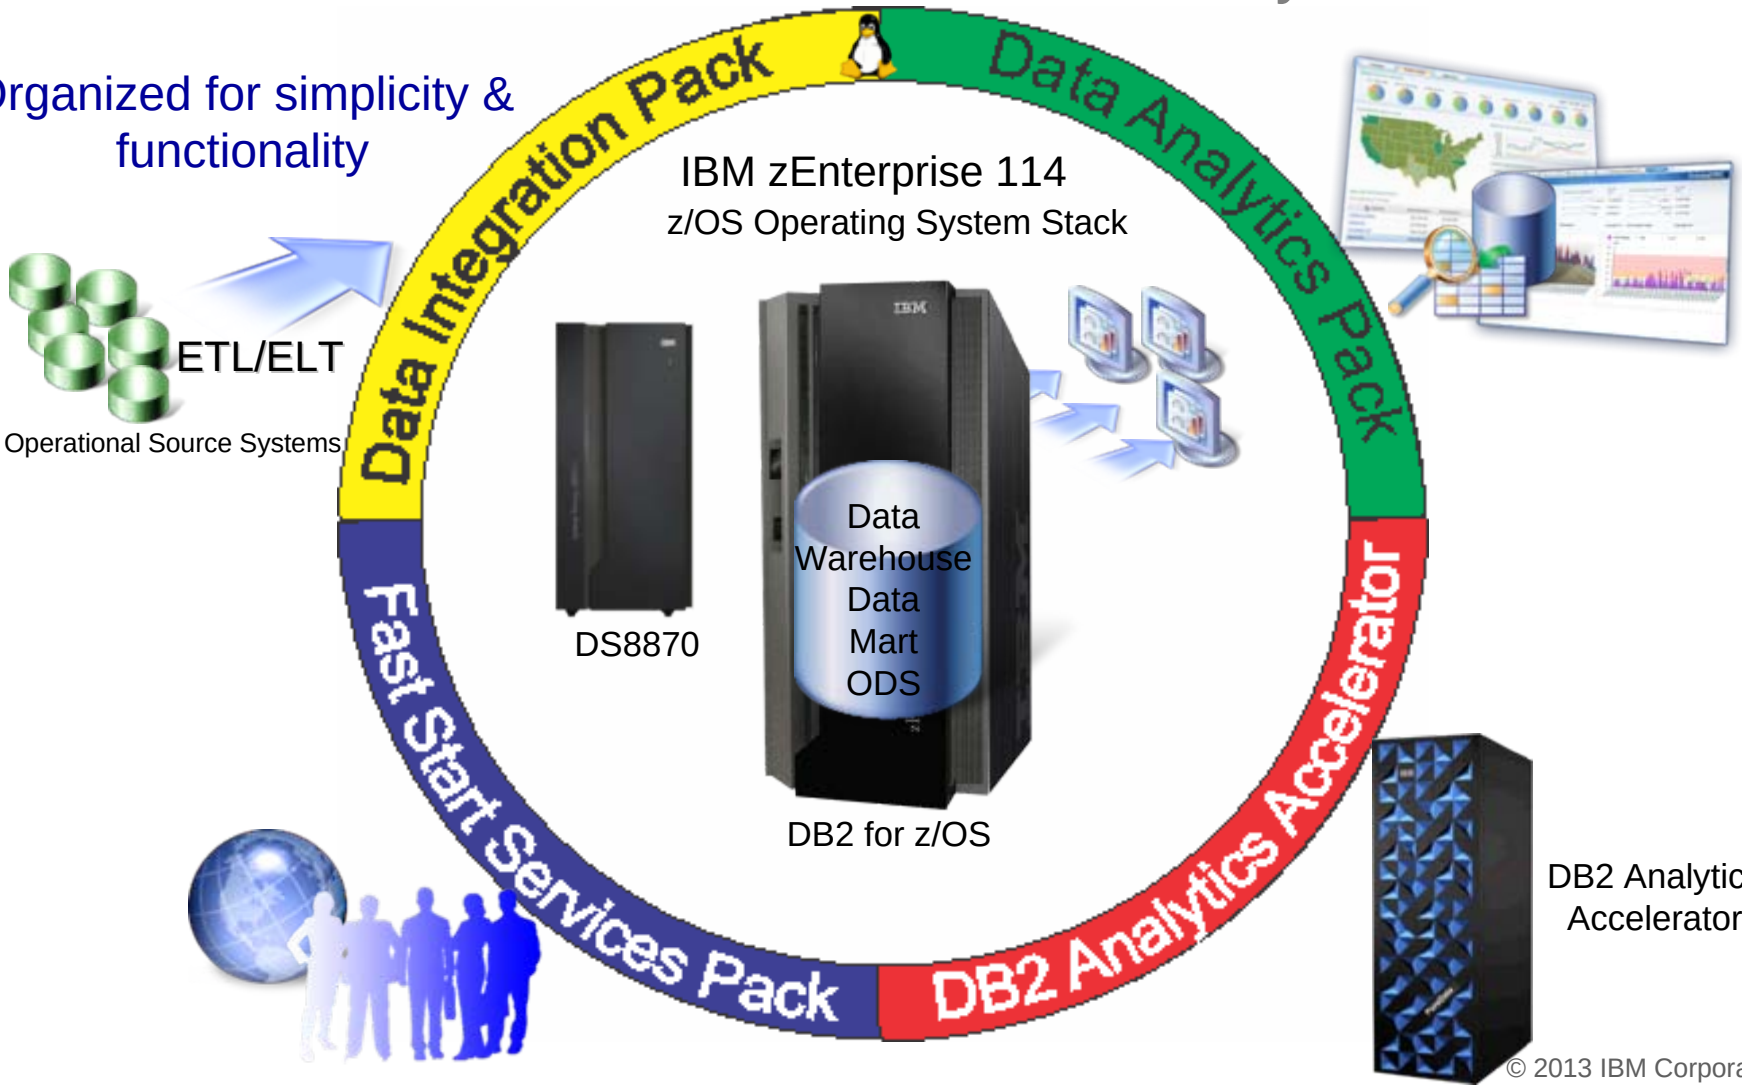

Cognos BI v10.2 for Linux and z/OS on zEnterprise

Improved User Experience

Enable users to focus on what's important

Improved Performance & Scale

Meeting the demands of growing data volumes

* Not supported on z/OS

Predictive Analytics for Linux on zEnterprise

IBM SPSS Statistics for Linux on System z

- Apply math to decision making and research for commercial, government, and academic users

IBM SPSS Modeler for Linux on System z

NEW

- Data mining tool used for generating hypotheses and scoring
- Text analysis for unstructured data to model consumer behavior
- **In-Transaction Scoring with DB2 z/OS: Embeds the Scoring Algorithm Directly within the Transactional Application**

SPSS Decision Management for Linux on System z

- Employs both predictive models and business rules to automatically generate recommended actions

SPSS Collaboration and Deployment Services for Linux on System z

- Provides role-based models and security for in scoring, job scheduling, repository services, and integration

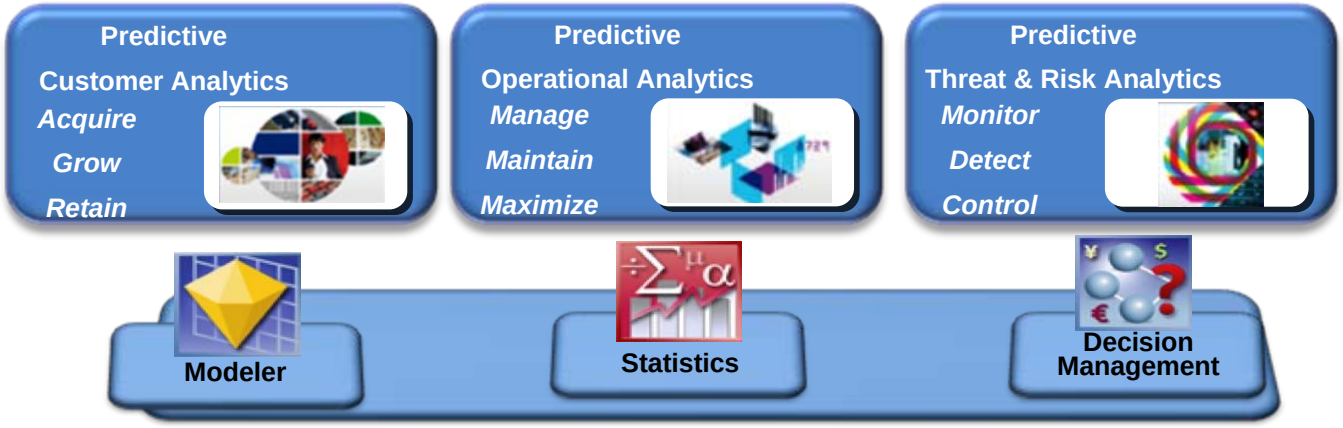

Real time scoring on the zEnterprise platform…

- Helps **improve the success rate of up sell / cross sell opportunities, fraud detection customer service** by using the most current transactional data to gain a more accurate view of the “next best action” to take.
- **Increases the speed and accuracy of scoring in real time** by imbedding the scoring algorithm in DB2 and running it directly within the transactional application resulting in:
  - **Reduced costs and complexity associated with network traffic** by bringing the analytics to the data on zEnterprise
  - Ability to **meet and exceed SLAs** associated with delivery of real time analytics by offering the **same service level as the OLTP applications**

Sample Real-time Scoring Use Cases

By Industry	
Banking	Monitor credit card usage in real-time to proactively detect/prevent
Insurance	Score claims in real-time to immediately identify fraudulent claims and identify up-sell opportunities
Government	Combine the details of a current crimes in progress with lessons learnt from past crimes to determine the safest course of action for the officer
Retail	Combine today's purchase details, with current market information and historical purchase patters to determine the best upsell opportunity when they are ready to spend money.
Telco	Combine todays complaint with the current account status and previous behavior to determine best upsell

The “how” of Predictive Scoring

Identify predictive models/patterns found in historical data

Use those predictive models with variables to score transactions & identify the best possible future outcomes

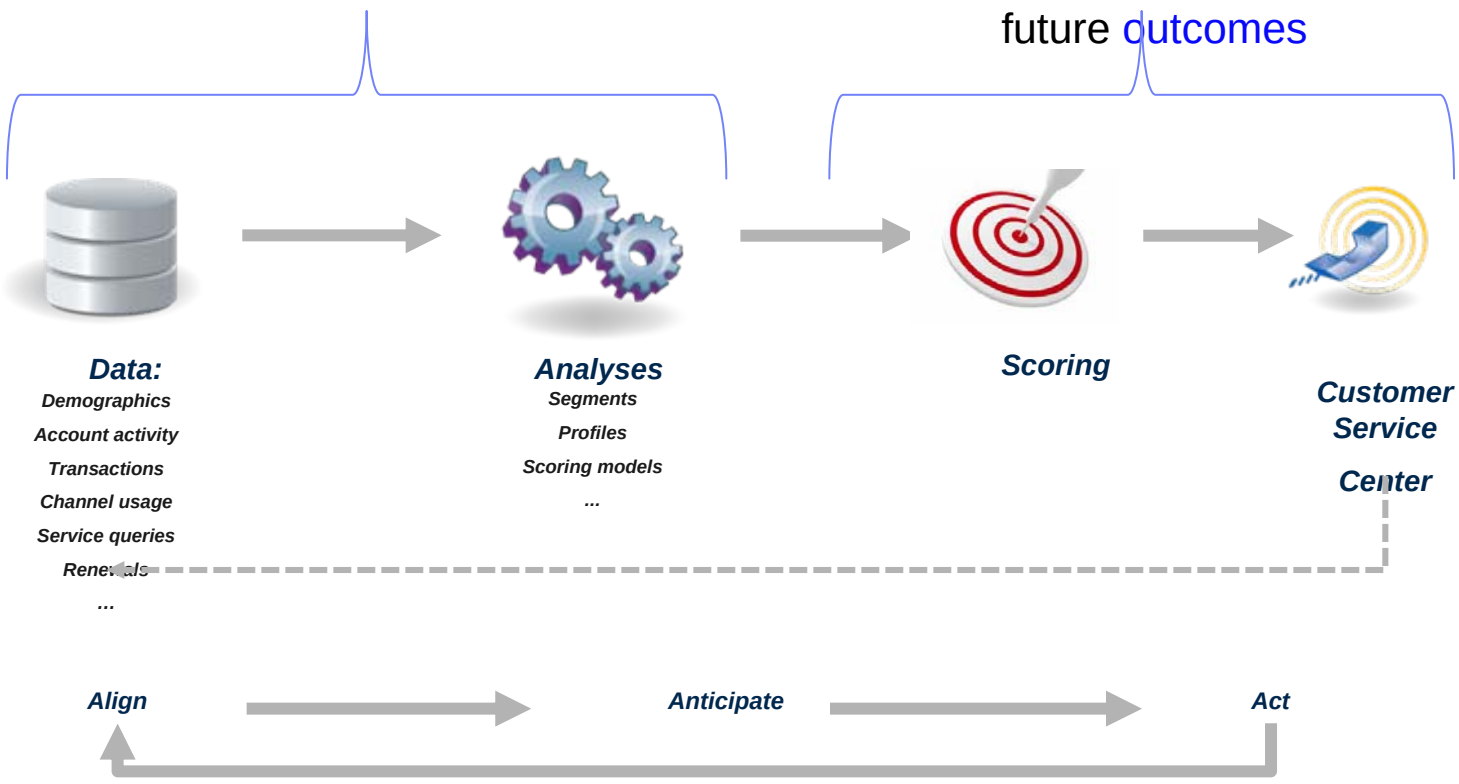

Scoring Options with an OLTP Application

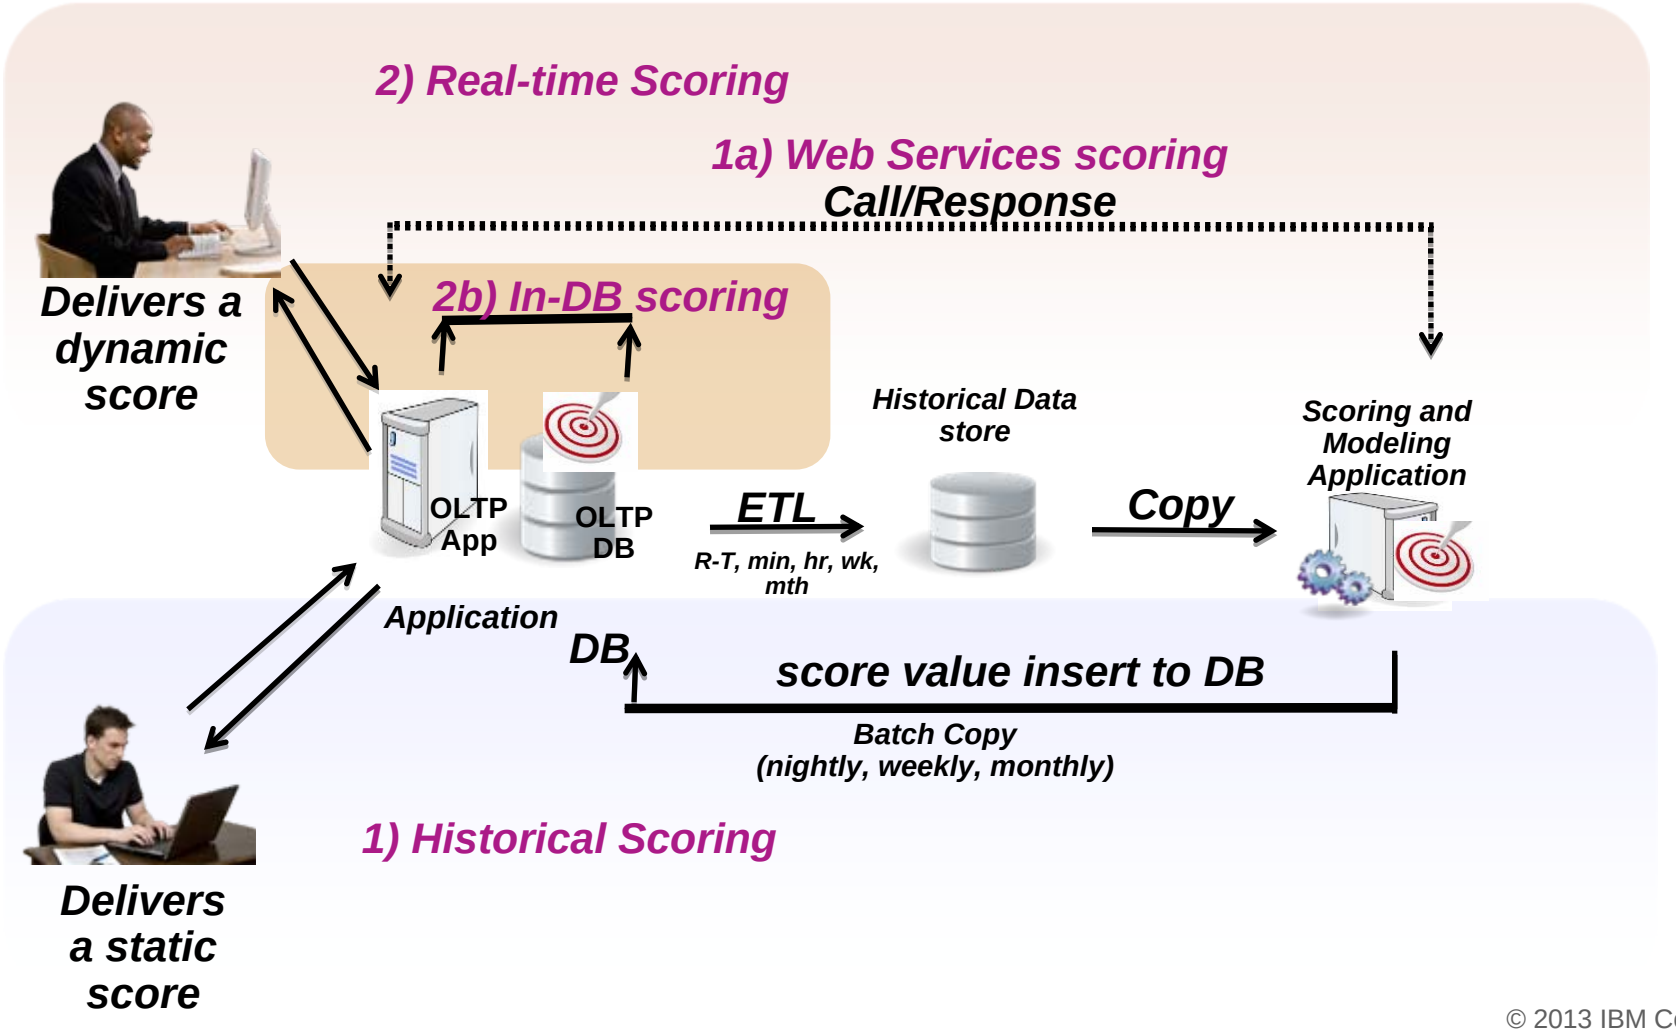

Integrating real time scoring for DB2 for z/OS into an OLTP application

1. Create a stream to build a model.
2. Execute the model building node to produce a model apply node.
3. Evaluate the model against a separate partition of the historical data to test the model
4. Publish the model to the DB2 for z/OS
5. Add the scoring call to application

To rebuild the model as business policies change, repeat steps 1..4 above

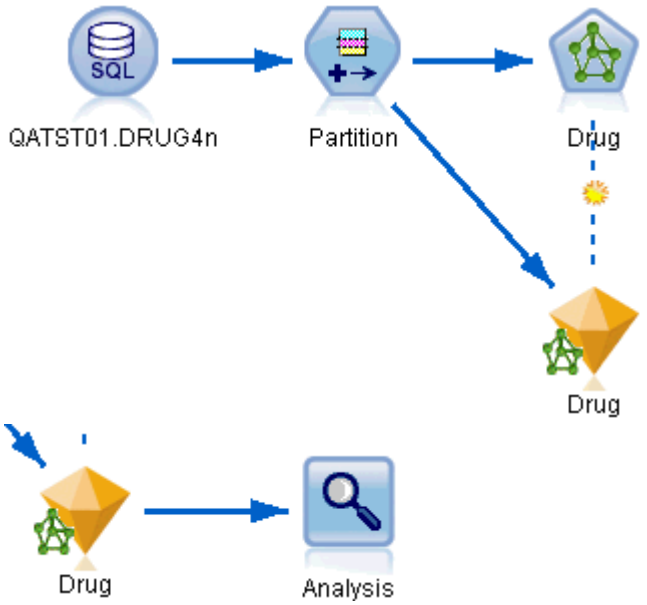

Modeler 15 Real-time Scoring with DB2 for z/OS

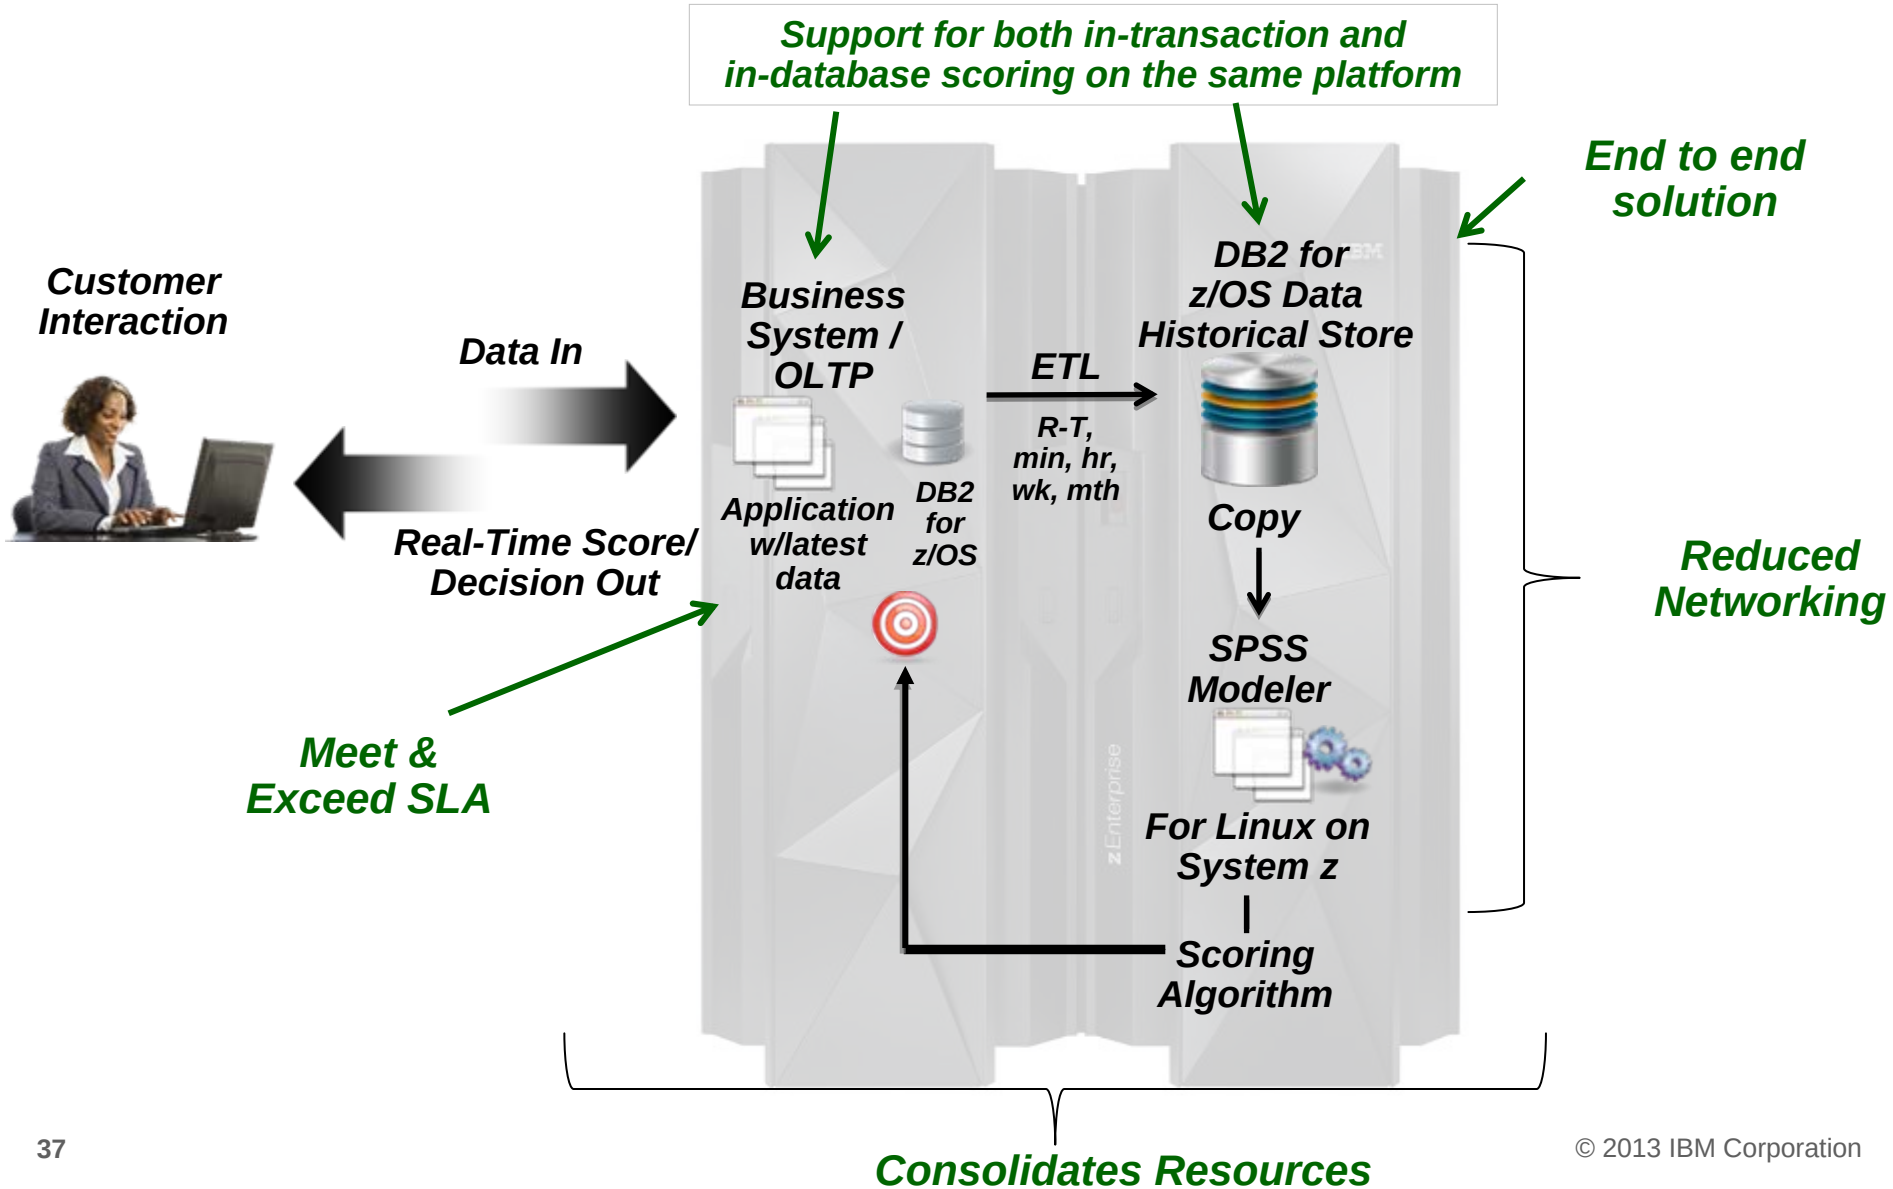

An enterprise information hub on a single, integrated, secure platform

Best OLTP/ Transactional Solutions

Industry leader in mission critical transactional systems

Best In Analytics

Industry recognized leader in BI, PA & DW solutions

Best In Flexibility

Start with your most critical business issue & quickly realize biz value with the flexibility to expand & grow across the enterprise

Transaction Processing (OLTP)

Business Analytics

Predictive Analytics

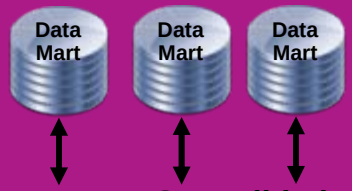

Data Mart Consolidation

zEnterprise

Recognized leader in workload management with proven security, availability, recoverability

DB2 Analytics Accelerator for z/OS *powered by Netezza technology*

Recognized leader in cost-effective high speed deep analytics

Unprecedented mixed workload flexibility & virtualization providing the most options for cost effective consolidation

Learn More

- *Visit the IBM Mainframe Business Analytics & Data Warehousing Website*
  http://www.ibm.com/software/os/systemz/badw/

- *Join the IBM Analytics Networking Community*
  http://db2forzos.ning.com/group/datawarehousebusinessintelligenceonsystemz

IBM

Thank You