

The Analytics Transformation on the Mainframe

Alan Meyer
IBM

March 15, 2012



Agenda

- Introduction: Challenges that organizations are facing
- Business Analytics - driving better business outcomes
- Customer Examples
- Data Warehousing – the foundation for business analytics
- Accelerators – delivering immediate analytics solutions
- Supporting Software

Drowning in a sea of data with no information



Volume of Digital Data



Variety of Information



Velocity of Decision Making

1 in 3

Business leaders **frequently make decisions** based on information they don't trust, or don't have

1 in 2

Business leaders say **they don't have access to the information they need** to do their jobs

Imagine the possibilities



...predict and treat high-risk patients to proactively intervene at time of visit?

...identify suspects based upon initial witness information on arrival at scene?

...provide targeted cross sells that drive additional sales when offered at time of sale?

... predict correct order levels at time of purchase?



Physician



Police Detective



Retail Sales Manager



Purchasing Agent

...optimize every transaction, process and decision at the point of impact, based on the current situation, without requiring that everyone be an analytical expert

Price by opportunity

...matching room prices to availability and customer type

- **Maximize inventory yield**
- **Best price, best yield for both Marriott and the customer in less than one second**



Analytics-driven Organizations Can...

Increase agility

...rapidly respond with information

- ***Immediate, accurate information***
- ***No searching for data***
- ***Timeliness of analysis***



Analytics-driven Organizations Can...



Detect patterns

...rapidly map symptoms to cures

- ***Clinical research studies***
- ***Identify best practices***
- ***Reduced operating costs***



FLORIDA
HOSPITAL

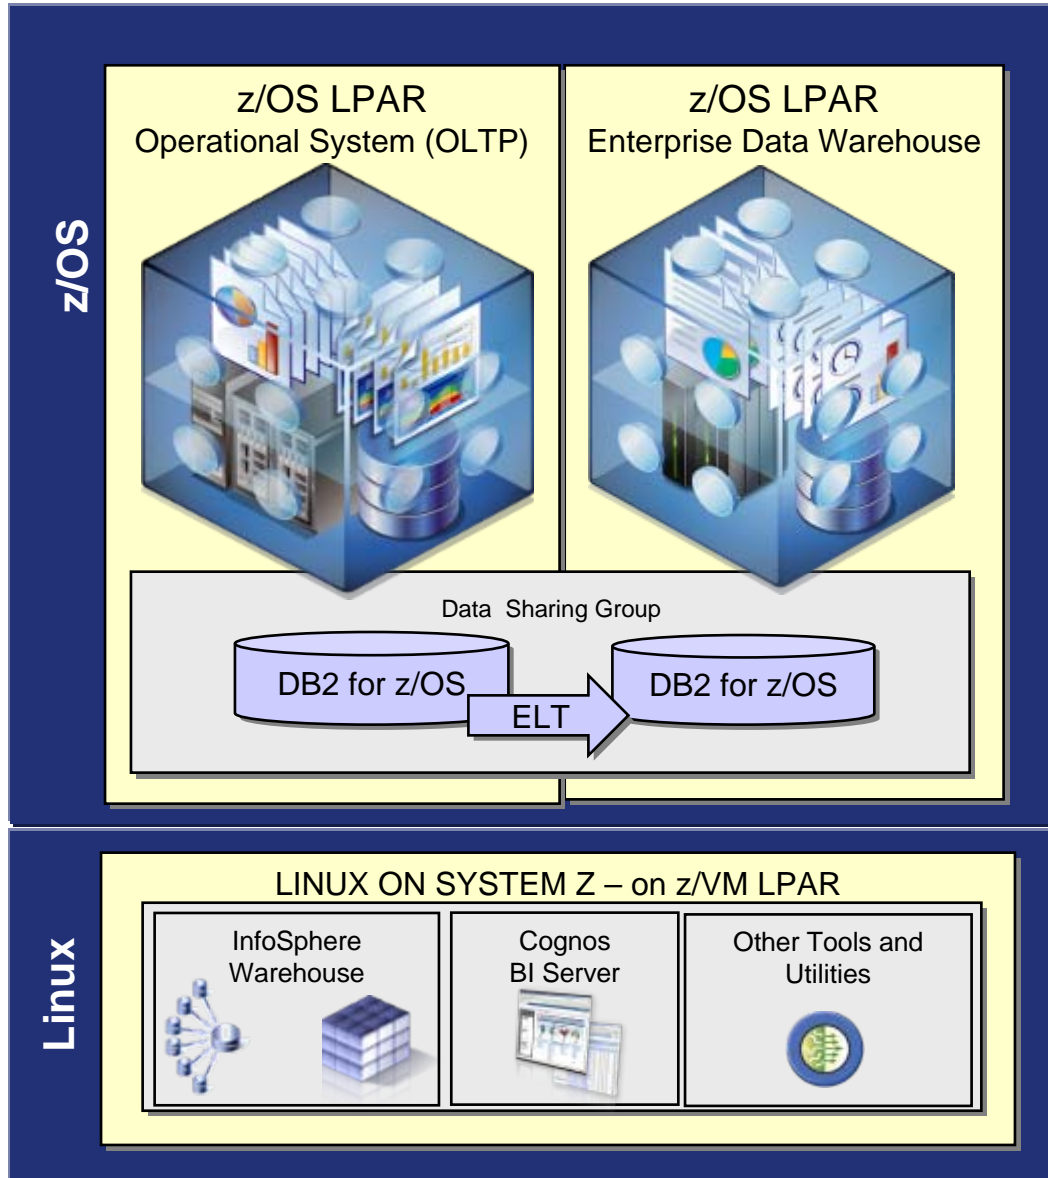
The skill to heal. The spirit to care.

The market is moving to the strengths of System z

- **Enterprises are modernizing and extending functionally**
 - Better decisions from the right information
 - Informed decisions at the point of contact
 - Consistency of information across organizations
- **Which is driving operational characteristics requirements**
 - Cost of downtime is escalating
 - The impact of unauthorized intrusion and publishing of private information is overwhelming
 - Stringent Service Level Agreements
- **Newer applications demand lower latency of the data**
 - Businesses want the most up-to-date information they can get
 - Yesterday's information was good yesterday
- **All while focusing on reducing costs/ consolidating**
 - Lower costs through reduced complexity
 - Simplified environment with easier administration
 - Lower SW costs
 - Reduced costs through elimination of redundant servers and resources
 - Reduced footprint, environmental, and administrative costs



A data warehouse solution on a System z foundation




- Minimization of data movement between operational system and data warehouse
- Lowers data latency for time sensitive decisions
- Enables consolidation and simplification of data warehouse and data marts
- Leverages existing high availability, backup, disaster recovery, and security environments
- Greater scalability of multidimensional analysis through cubing services (data marts) and DB2 enhancements
- Complex transformations and data quality driven from Linux on System z with Data Studio

Data mart consolidation

Consolidation into a single footprint

Utilize virtualization to optimize the use of resources while reducing costs and gaining new agility



- Lower software costs
- Lower hardware costs
- Lower administration costs
- Lower environmental costs
- Greater flexibility
- Less complexity
- Fewer points for security intrusions

Three ways to consolidate on System z

Choose a platform to consolidate upon



DB2 for z/OS

zEnterprise



Robust Mixed Workload

DB2 for Linux

zEnterprise



Integrated System

DB2 for Linux/AIX

zBX



Managed by System z

Functionality

The right mix of functionality and investment

Investment

Standardized Business Intelligence tools

Delivering information when, where, and how each user needs it



■ Delivers information where, when and how it is needed

- Self-service reporting and analysis
- Individualized by user
- Automated delivery of information in context
- Author once, consume anywhere

■ Full range of BI capabilities

- Query, reporting, analysis, dashboarding, realtime monitoring

■ Purpose-built SOA platform

that fits client environments and scales easily

Query Management Facility (QMF) 10

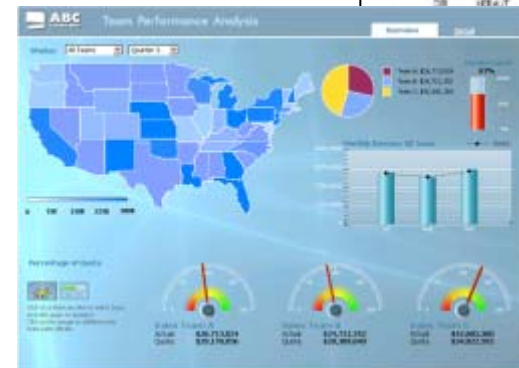
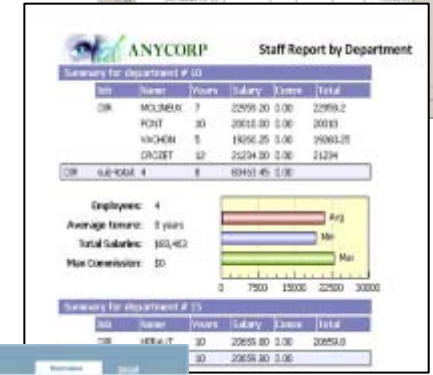
Meeting the challenges of today's Business Analytics requirements



Today's functionality with support for yesterday's applications

- 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

NAME	YEARS	SALARY	EXPN	TOTAL
SMITHS	5	2000.00	0.00	\$10,000.00
SMITHS	5	2000.00	0.00	\$10,000.00
SMITHS	5	2000.00	0.00	\$10,000.00
SMITHS	5	2000.00	0.00	\$10,000.00
SMITHS	5	2000.00	0.00	\$10,000.00
SMITHS	5	2000.00	0.00	\$10,000.00
SMITHS	5	2000.00	0.00	\$10,000.00
SMITHS	5	2000.00	0.00	\$10,000.00
SMITHS	5	2000.00	0.00	\$10,000.00
SMITHS	5	2000.00	0.00	\$10,000.00



SPSS Modeler with InfoSphere Warehouse



Full breadth of predictive analytics

Data collection, statistics, data mining, predictive modeling, deployment services...

Putting prediction in hands of the business

Decision Management

Driving better business outcomes

- Attract and retain profitable customers
- Detect and prevent fraud
- Improve resource allocation

IBM Smart 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.

- ***Secure, Available Business Analytics***
- ***Simplified administration***
- ***Proven Operational Characteristics***
- ***High Value Operational BI***

Making every decision on facts, at the point of impact

IBM Smart Analytics System 9700

High Value Data Warehousing – Standard Configuration



**System z
Z196**

Cognos 10.1

Cognos
software

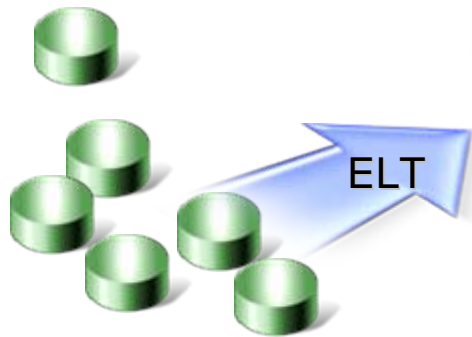


DS8800 Storage

**InfoSphere
Warehouse 9.7.3**



**SPSS
Modeler 14.2**



Operational Source Systems
Structured/ Unstructured Data



**DB2 for z/OS V10 VUE
(option for MLC)**

DB2 Utilities Suite
Image Copy, LOAD, UNLOAD,
REORG, etc

**Implementation
Services**



IBM Smart Analytics System 9710

Unprecedented Value in Deploying New Business Analytics



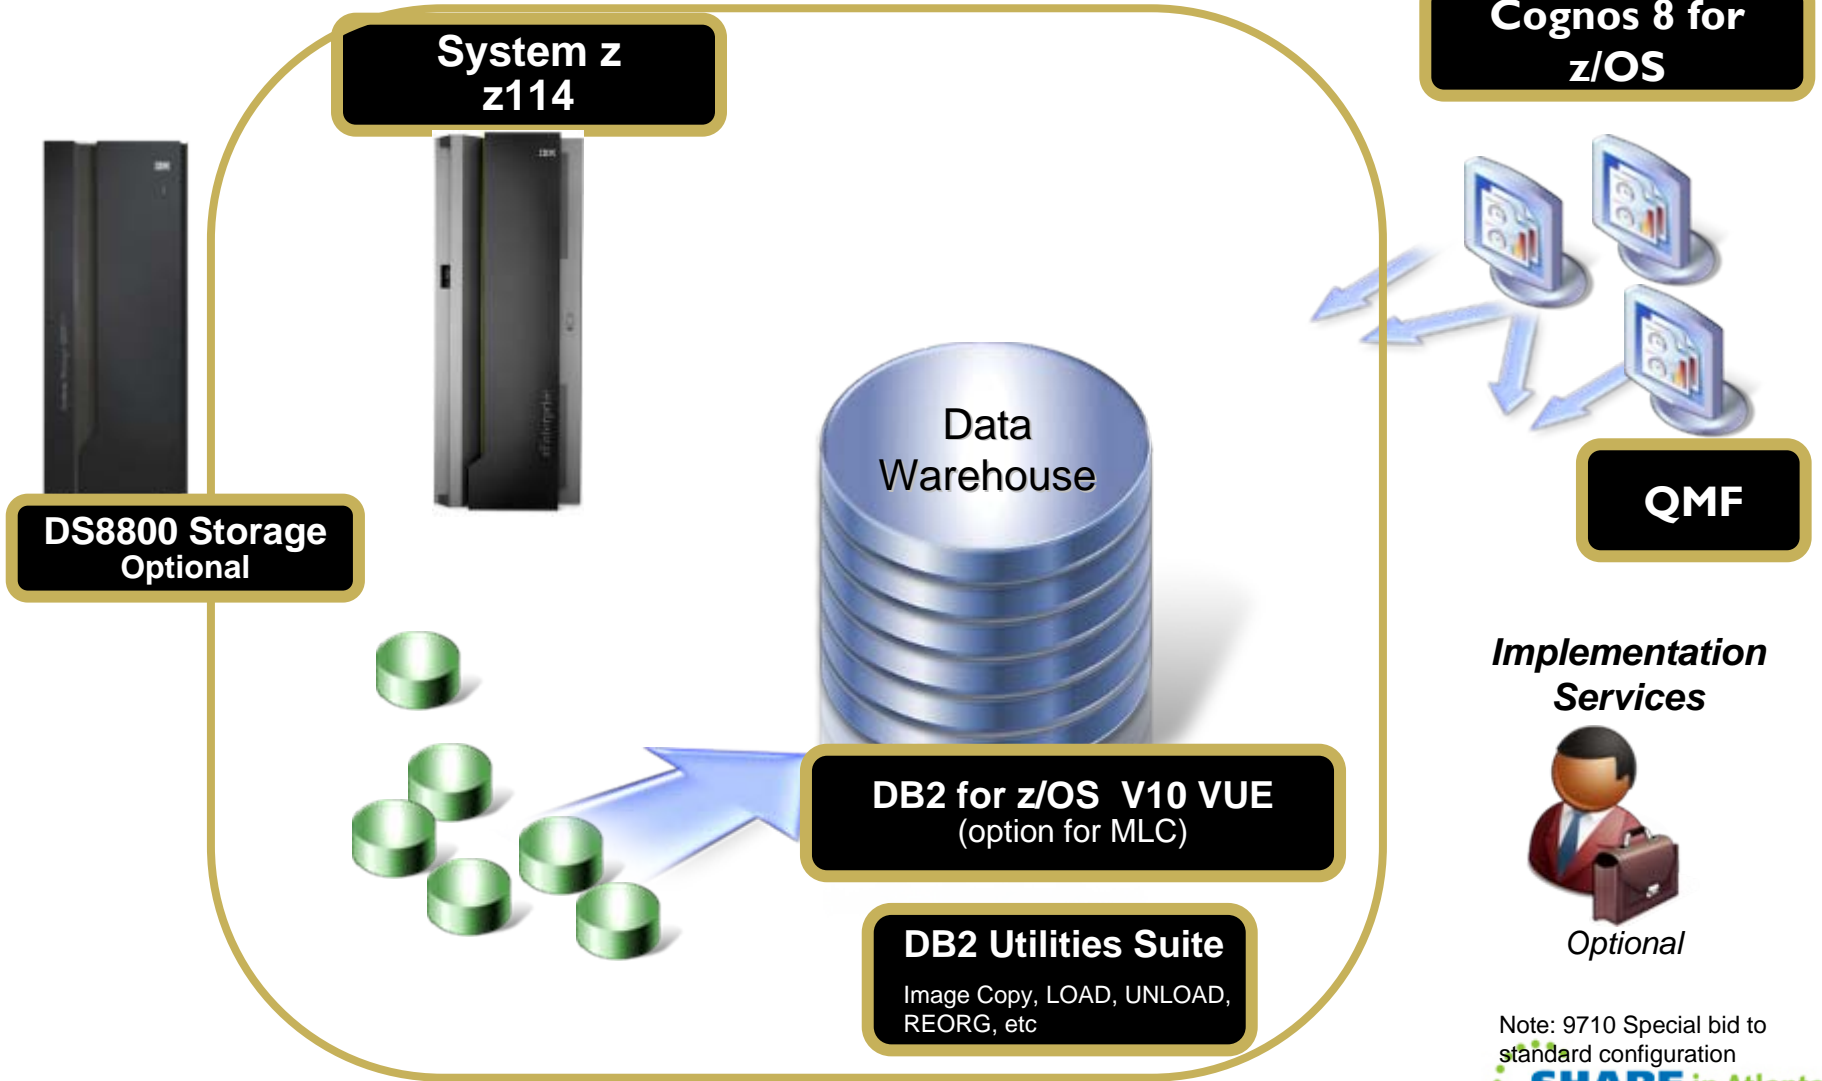
Introducing the IBM Smart Analytics System 9710 based upon the new zEnterprise 114 platform, delivering the quality of service of System z at an entry-level price.

- ***Small Footprint – Big Value***
- ***High Value Operational BI***
- ***Secure, Available Business Analytics***
- ***Foundation Offering***

Operational Businesses Analytics at an entry-level price

IBM Smart Analytics System 9710

Unprecedented Value Data Warehousing



Foundation Offering

Note: 9710 Special bid to standard configuration



Simplicity, Flexibility, Choice

IBM Data Warehouse & Analytics Solutions



IBM Netezza



True Appliance

IBM Smart Analytics System



Flexible Integrated System

IBM InfoSphere Software



Custom Solution

Warehouse Accelerators

Information Management Portfolio

(Information Server, MDM, Streams, etc)

Simplicity

The right mix of simplicity and flexibility

Flexibility

DB2 Analytics Accelerator

Accelerating decisions to the speed of business



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



Get more insight from your data

- Fast, predictable response times for “right-time” analysis
- Accelerate analytic query response times
- Improve price/performance for analytic workloads
- Minimize the need to create data marts for performance
- Highly secure environment for sensitive data analysis
- Transparent to the application

Time is money

“The more data we have, the longer our analysis takes!”

- Waiting for key queries to complete
- Waiting for fact based information to drive key decisions
- Analysts performing query analysis and rewrites
- Adding indices, MQTs in an attempt to speed up queries



Large Insurance Company

Adding value by Accelerating the Delivery of Business Reporting



				DB2 Only		DB2 with IDAA		Times Faster
Query	Total Rows Reviewed	Total Qualifying Rows	Total Rows Returned	Hours	Sec(s)	Hours	Sec(s)	
Query 1	591,941,065	2,813,571	853,320	2:39	9,540	0.0	5	1,908
Query 2	591,941,065	2,813,571	585,780	2:16	8,220	0.0	5	1,644
Query 3	813,343,052	8,260,214	274	1:16	4,560	0.0	6	760
Query 4	283,105,125	2,813,571	601,197	1:08	4,080	0.0	5	816
Query 5	591,941,089	3,422,765	508	0:57	4,080	0.0	70	58
Query 6	813,343,052	4,290,648	165	0:53	3,180	0.0	6	530
Query 7	591,941,065	361,521	58,236	0:51	3,120	0.0	4	780
Query 8	813,343,052	3,425,292	724	0:44	2,640	0.0	2	1,320
Query 9	813,343,052	4,130,107	137	0:42	2,520	0.1	193	13

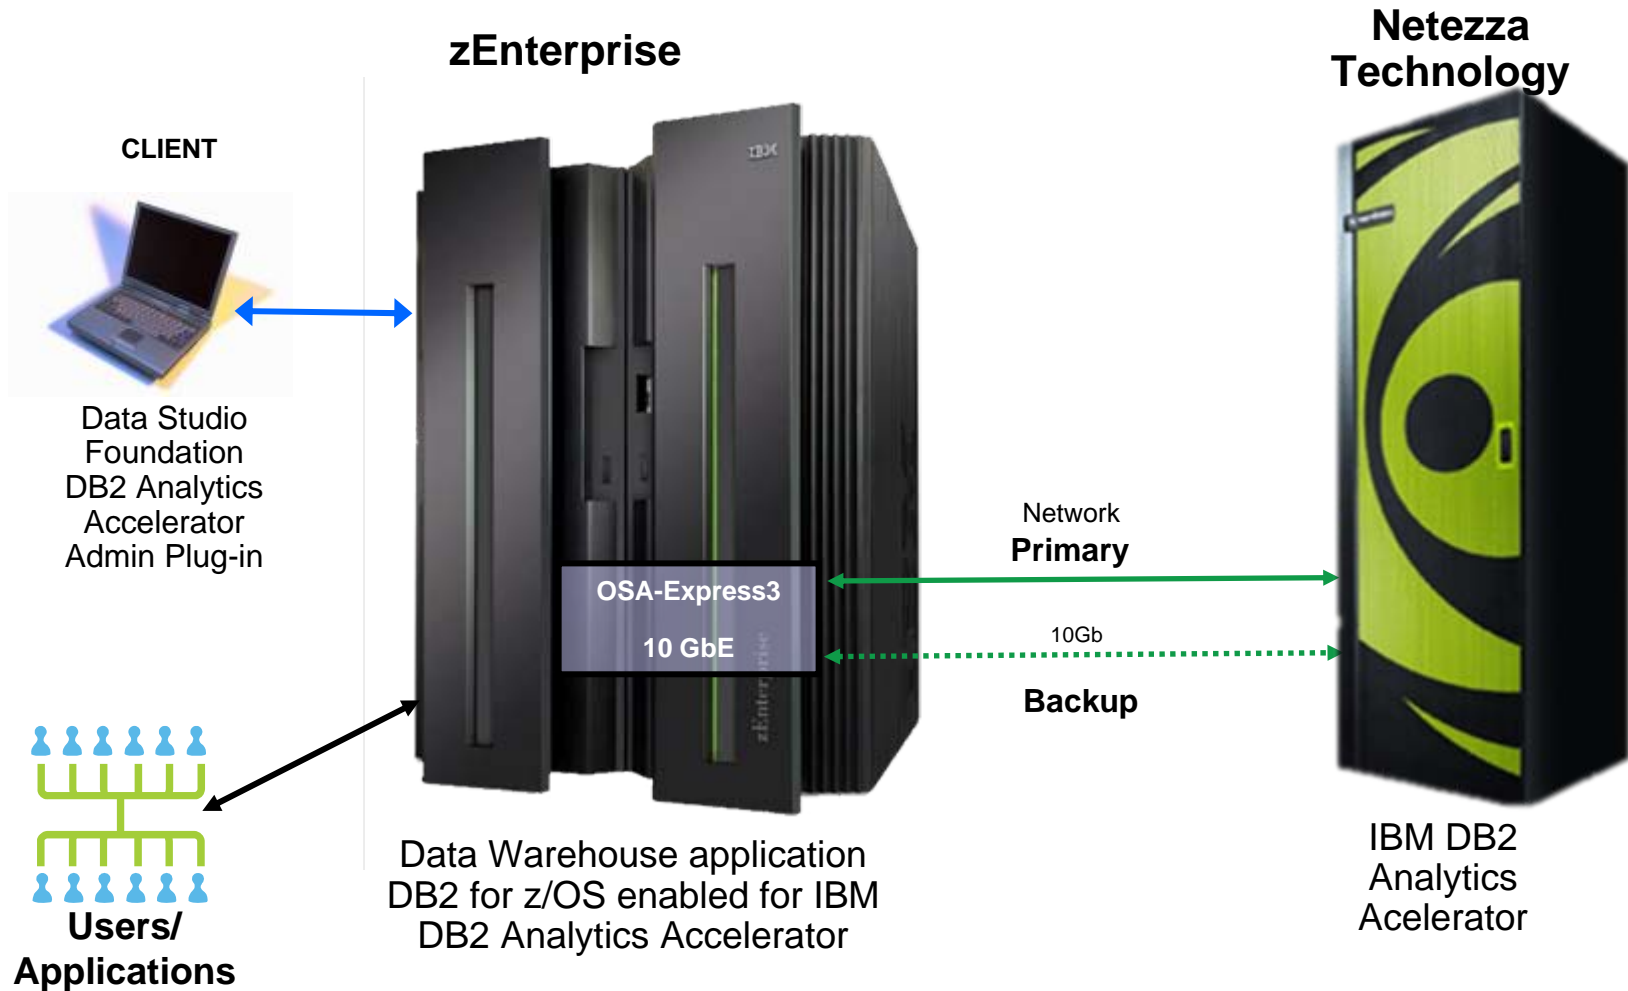
With Accelerated Time to Value

- IBM DB2 Analytics Accelerator (Netezza 1000-12)**
 Production ready - 1 person, 2 days
- Table Acceleration Setup in 2 Hours**
 - DB2 “Add Accelerator”
 - Choose a Table for “Acceleration”
 - Load the Table (DB2 Loads Data to the Accelerator)
 - Knowledge Transfer
 - Query Comparisons
- Initial Load Performance**
 400 GB Loaded in 29 Minutes
 570 Million Rows (Actual: Loaded 800 GB to 1.3 TB per hour)
- Extreme Query Acceleration - 1908x faster**
 2 Hours 39 minutes to 5 Seconds
- CPU Utilization Reduction**
 Up to 35%


Customer Quote: “we had this up and running in days with queries that ran over 1000 times faster”



IBM DB2 Analytics Accelerator V2 Product Components






Query Execution Process Flow

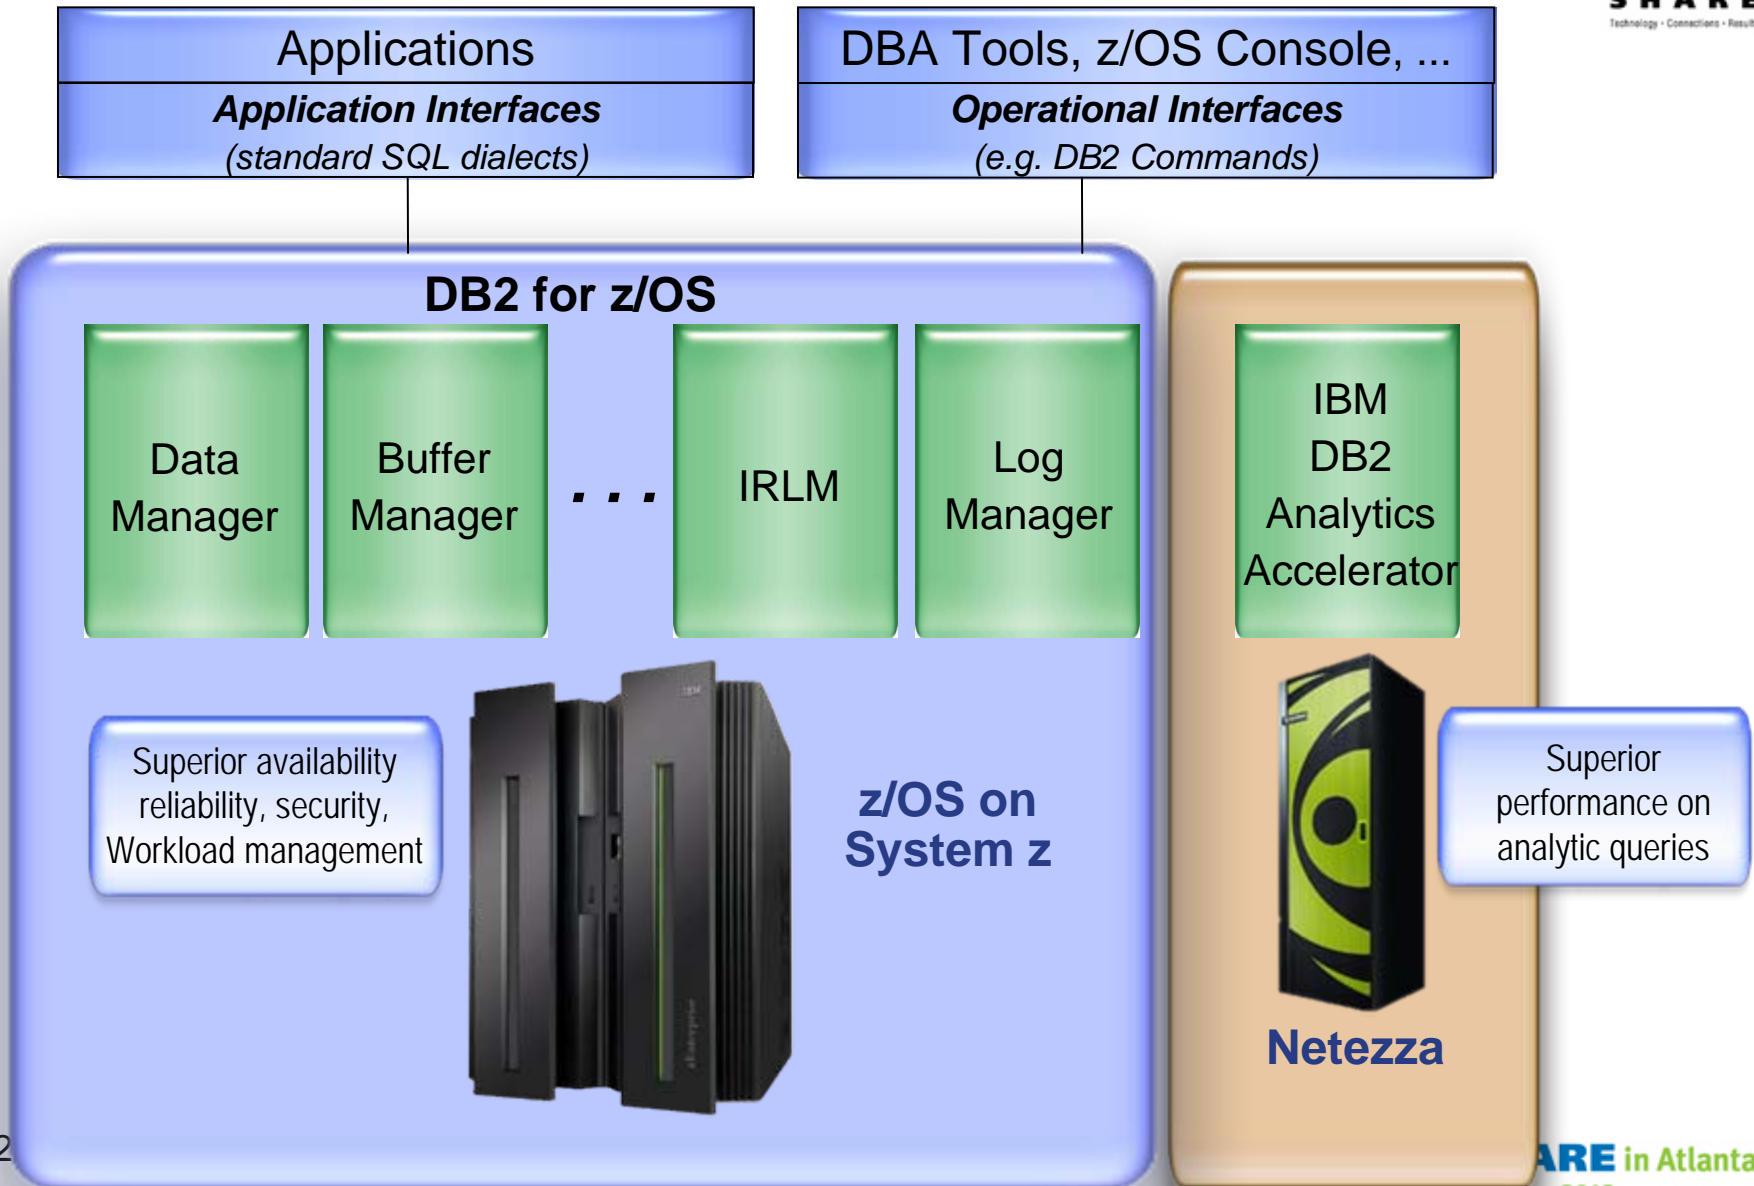


DB2 for z/OS

IDAA

-  Queries executed without IDAA
-  Queries executed with IDAA
-  Heartbeat (IDAA availability and performance indicators)

Deep DB2 Integration within zEnterprise



Netezza 1000 Appliance Scalability



	1000-3	1000-6	1000-12	1000-24	1000-36	1000-48	1000-72	1000-96	1000-120
Cabinets	1/4	1/2	1	2	3	4	6	8	10
Processing Units	24	48	96	192	288	384	576	768	960
Capacity (TB)	8	16	32	64	96	128	192	256	320
Effective Capacity (TB)*	32	64	128	256	384	512	768	1024	1280

Current IDAA Platforms →

Predictable, Linear Scalability throughout entire family

Capacity = User Data space
 Effective Capacity = User Data Space with compression

*: 4X compression assumed

DB2 Analytics Accelerator V2

Powered by Netezza 1000 Appliance



Disk Enclosures

SMP Hosts

Snippet Blades™
(S-Blades, SPUs)



Slice of User Data

Swap and Mirror partitions

High speed data streaming

High compression rate

EXP3000 JBOD Enclosures

12 x 3.5" 1TB, 7200RPM, SAS (3Gb/s)

max 116MB/s (200-500MB/s compressed data)

e.g. TF12:

8 enclosures → 96 HDDs

32TB uncompressed user data (→ 128TB)

IDAA Server

SQL Compiler, Query Plan, Optimize

Administration

2 front/end hosts, IBM 3650M3

clustered active-passive

2 Nehalem-EP Quad-core 2.4GHz per host

Processor &

streaming DB logic

High-performance database

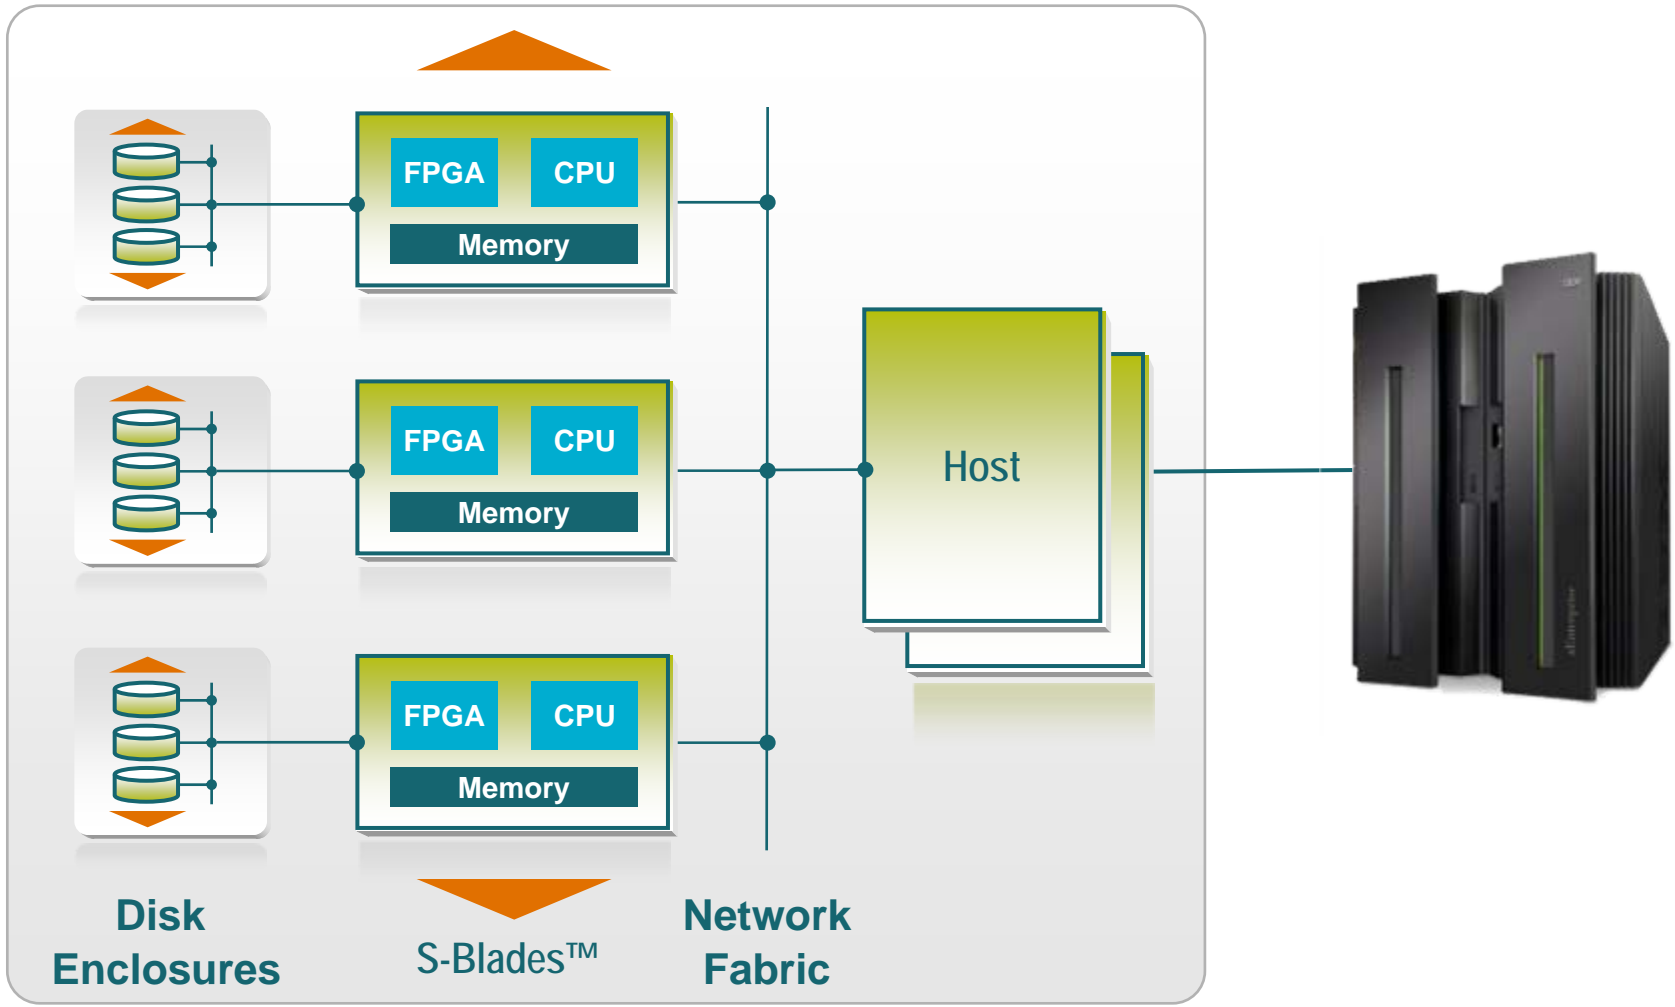
engine streaming joins,

aggregations, sorts, etc.

e.g. TF12: 12 back/end SPUs

(more details on following charts)


The Appliance Connected to a System z



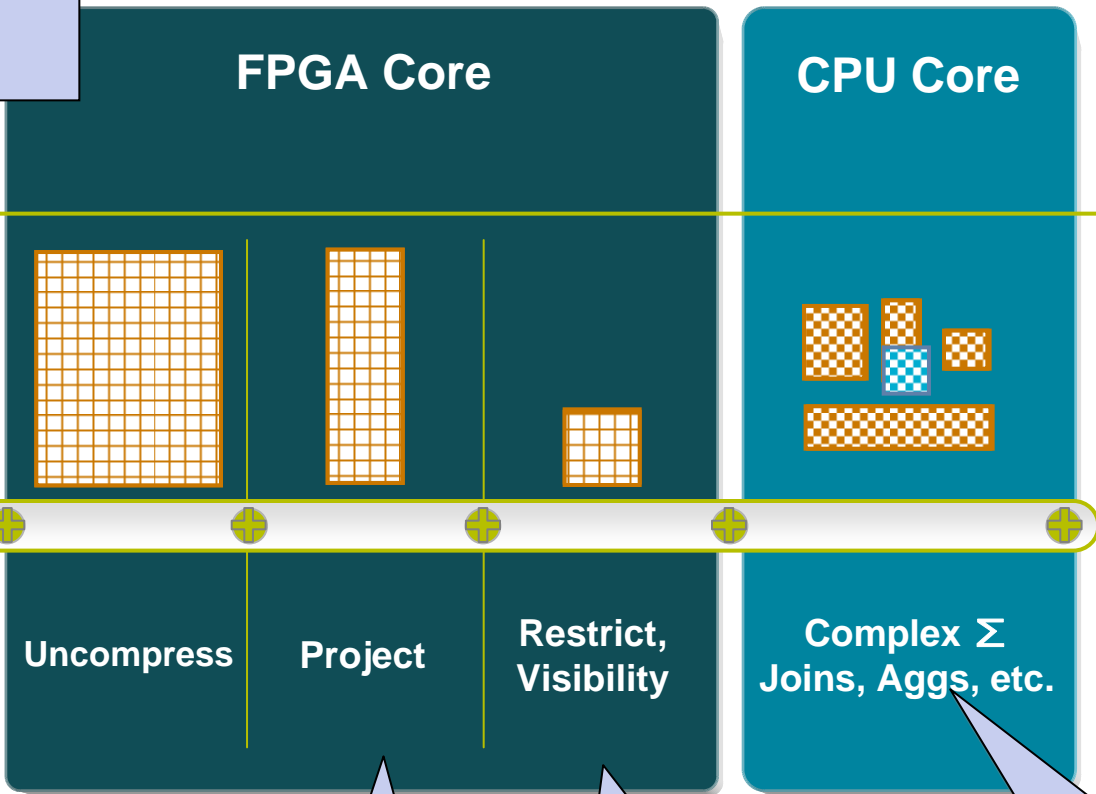
Netezza Appliance

The Key to the Speed

```
select DISTRICT,
       PRODUCTGRP,
       sum(NRX)
from   MTHLY_RX_TERR_DATA
where  MONTH = '20091201'
and    MARKET = 509123
and    SPECIALTY = 'GASTRO'
```



Slice of table
MTHLY_RX_TERR_DATA
(compressed)



```
select DISTRICT,
       PRODUCTGRP,
       sum(NRX)
```

```
where MONTH = '20091201'
and    MARKET = 509123
and    SPECIALTY = 'GASTRO'
```

sum(NRX)

Why Both?

Marrying the best of each

IBM System z



Mixed Workload System

IBM Netezza



Focused Appliance

Capitalizing on the strengths of both platforms while driving to the most cost effective, centralized solution - destroying the myth that transaction and decision systems had to be on separate platforms

Very diverse workload

Very focused workload

Tailored to your needs

A Hybrid Solution



IBM System z with IBM DB2 Analytics Accelerator

Mixed Workload System

- Mixed workload system z with operational transaction systems, data warehouse, operational data store, and consolidated data marts.
- Unmatched availability, security and recoverability
- Natural extension to System z to enable pervasive analytics across the organization.
- Speed and ease of deployment and administration

IBM Netezza

Focused Appliance

- Appliance with a streamlined database and HW acceleration for performance critical functionality
- Price/performance leader
- Speed and ease of deployment and administration
- Optimized performance for deep analytics, multifaceted, reporting and complex queries

Flexibility

The right mix of simplicity and flexibility

Simplicity

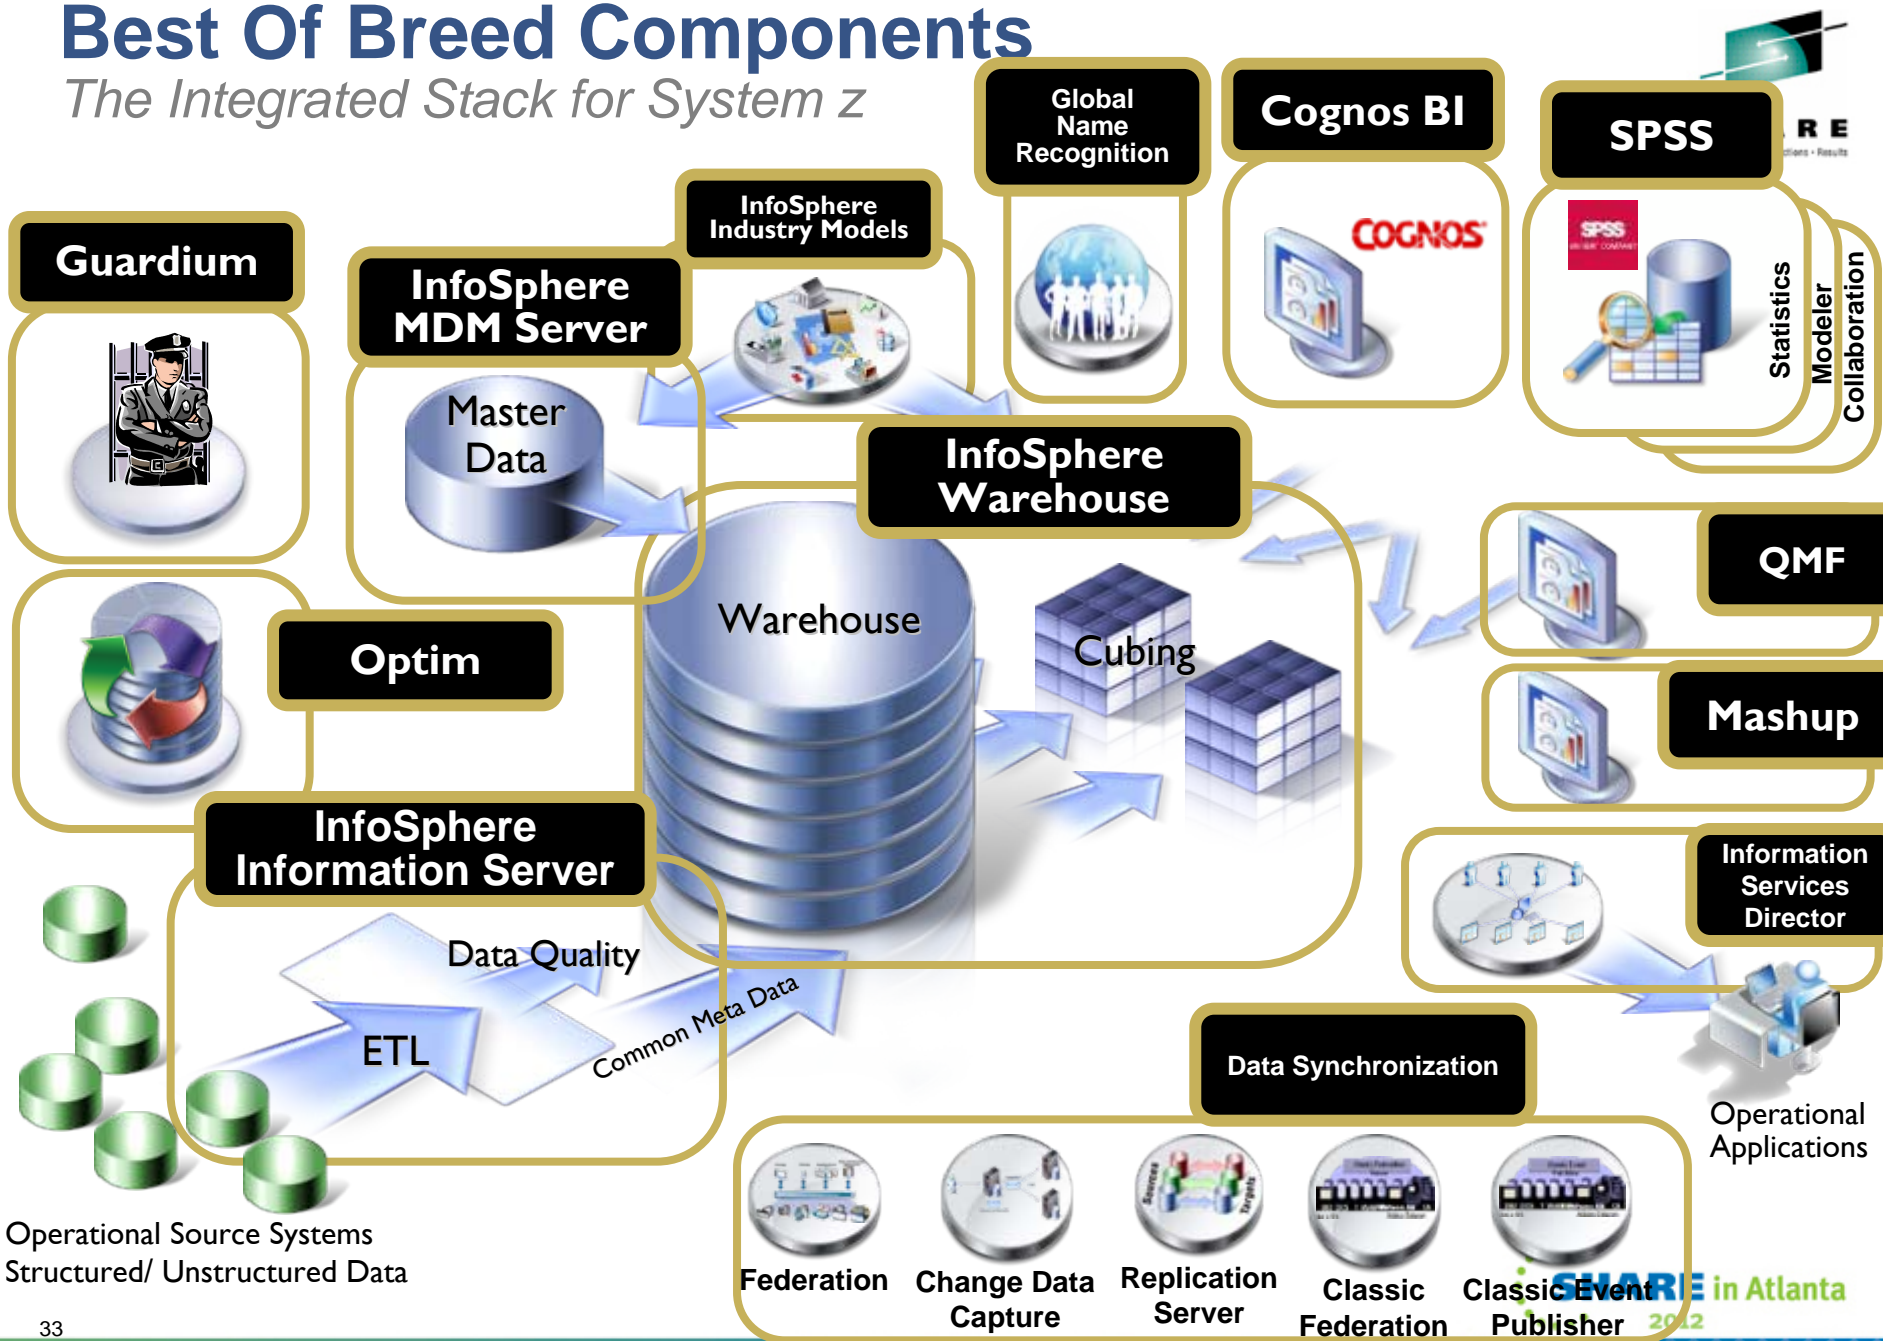
OLTP vs. Analytics – Examples

OLTP - “Transactional”	Transactional Analytics: (Operational BA)	Deep Analytics
Withdrawal from a bank account using an ATM	Approve request to increase credit line based on credit history and customer profile	Regular reporting to central bank – sum of transactions by account
Buying a book at Amazon.com	Propose additional books based on similar purchases by other customers	Which books were best-sellers in Europe over the last 2 months?
Check-In for a flight at the airport	Offer an upgrade based on frequent flyer history of all passengers and available seats	Marketing campaign to sell more tickets in off-peak times
Hand-over manufactured printers to an oversea-carrier	Optimize shipping by selecting cheapest and most reliable carrier on demand	Trend of printers sold in emerging countries versus established markets.

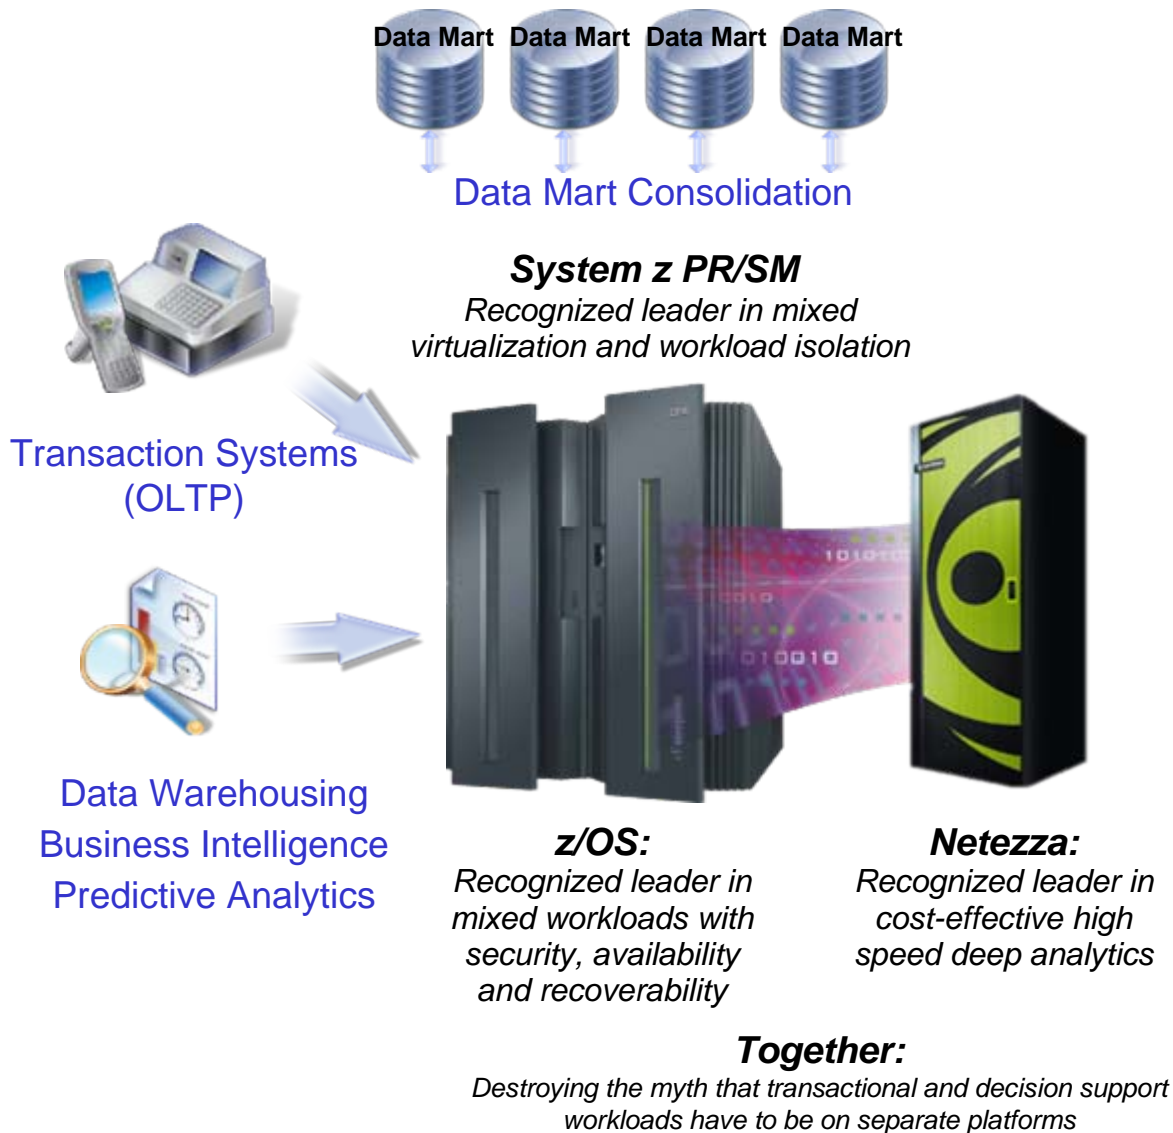


Best Of Breed Components

The Integrated Stack for System z



The Ultimate Consolidation Platform



Bringing it all together

- *Better Business Response*
- *Reduced Costs*
- *More Available*
- *More Secure*
- *Reduced Data Movement*
- *Reduced Data Latency*
- *Reduced Complexity*
- *Reduced Resources*

Today's IBM System z is ...



- The world's most trusted transaction processing and data server for business critical applications
- The world's most cost-efficient platform for data center consolidation and virtualization
- The world's most dependable and scalable hardware and middleware platform for new business applications
- A thoroughly modern application environment for traditional and Cloud delivery models



The zEnterprise 196 is the world's fastest and most scalable enterprise system. (50 BIPS)

Based on 5.2GHz core processor speed





SHARE
Technology • Connections • Results

Thank You