DB2 Analytics Accelerator – Optimizing your Key Business Decisions

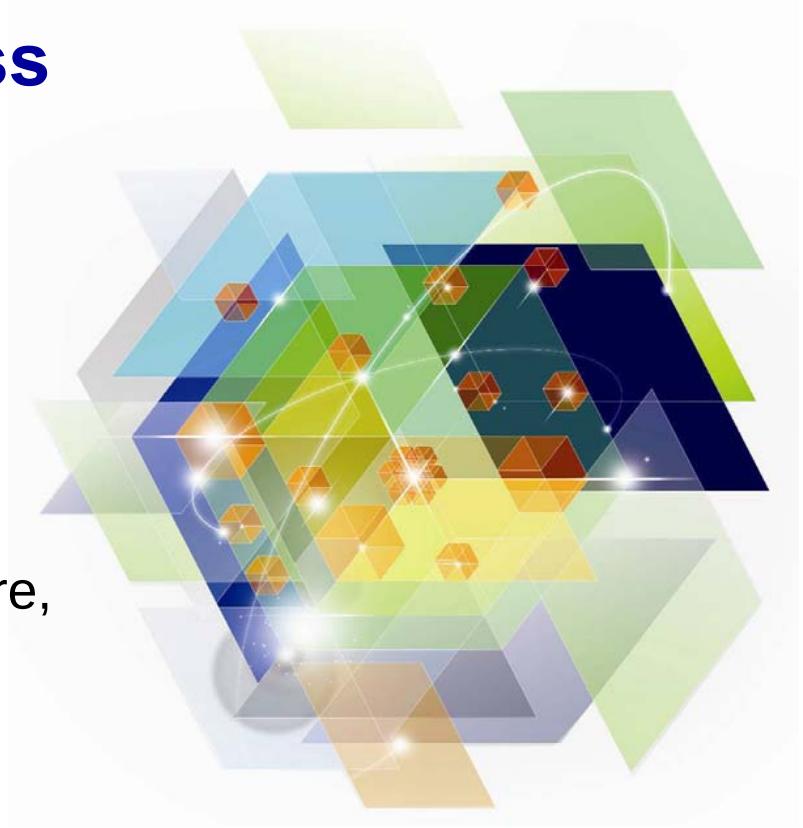

Dan Wardman, Vice President,
Information Management, Mainframe Software,
IBM Software Group

March 20,2012

Knowing what happened is no longer adequate.

Business leaders need to know

what is happening now,

what is likely to happen next and

what actions they should take.

The market is moving to the strengths of System z

- **Enterprises are expanding the role of analytics**
  - 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 must be met
- **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

Operational information to drive the business

Using information to make fact-based decisions

- **See:** *Observe trends, identify events, and detect changes within the business and external activities that affect the business.*
- **Respond:** *Take actions using current information in the context of the issue, at the time of decision*
- **Evaluate:** *Categorize, classify, and understand the information of the business, predicting future outcomes and behaviors.*
- **Transform:** *Optimize business processes to be compliant, that reduce costs, increase sales, and enable better customer service*

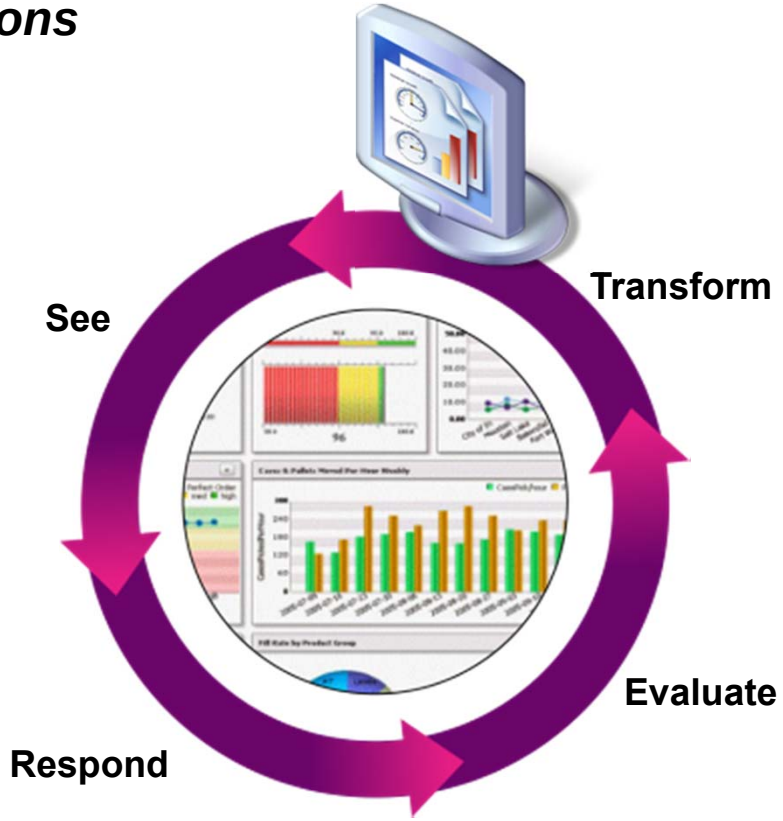

Foundation for Business Analytics and Optimization

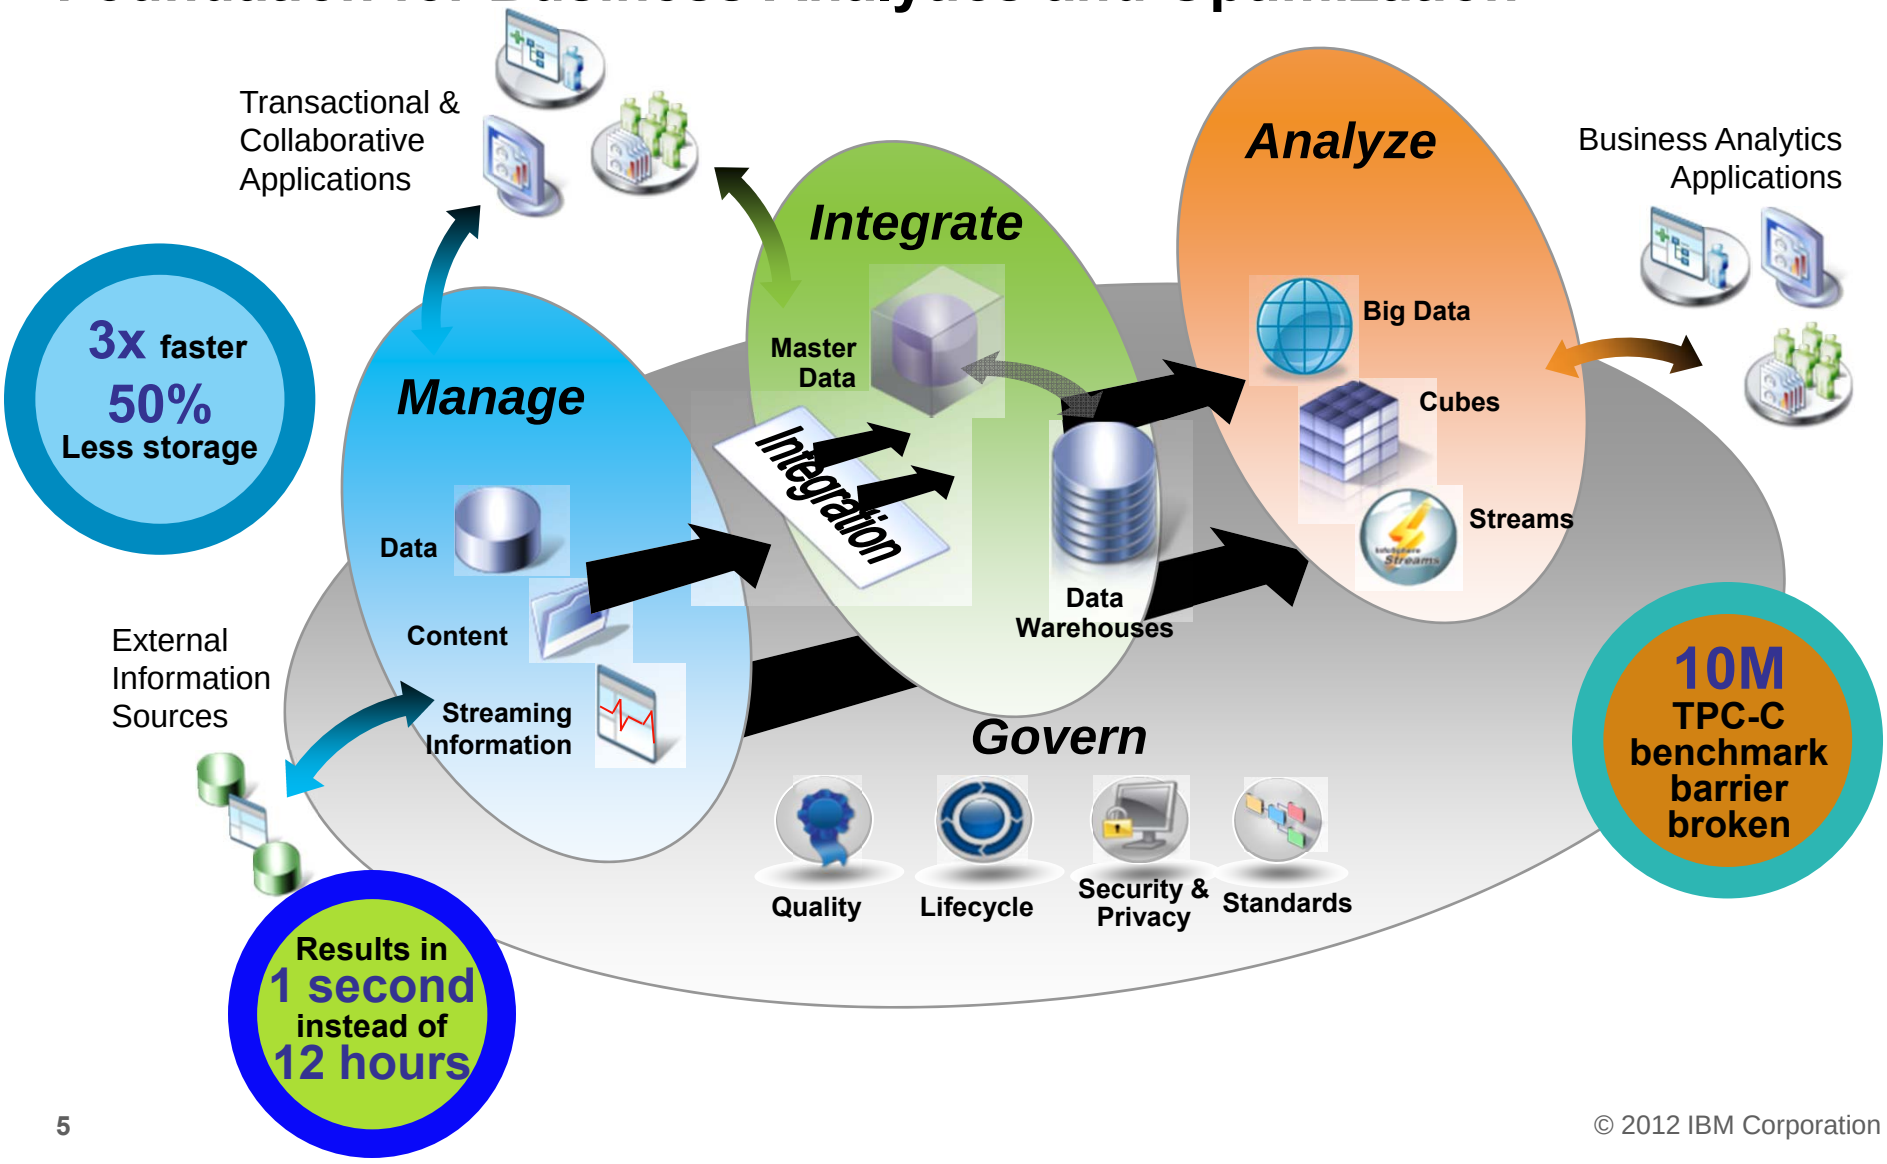

Data Warehouses have become isolated and outdated

Any data warehouse which has not had a significant architectural or infrastructure update since 2008 should undergo a significant evaluation and most likely requires an immediate retrofit for technology and design advances

Source: Gartner, The State of Data Warehousing in 2011, Jan 2011

- Much of the information to drive the business is known but not available to the decision makers
- Information in the DW is limited to a small number of people in the organization
- Little to no interactivity with other systems
- Not built with the same criteria as the operational systems
- Difficult to manage and maintain multiple servers and copies of the data
- Minimal control over who is accessing the data

Only 8.2% of the employees of a typical organization regularly use BI applications

Source: Gartner ,The BI Survey 8,, 03/09

A data warehouse solution on a System z foundation

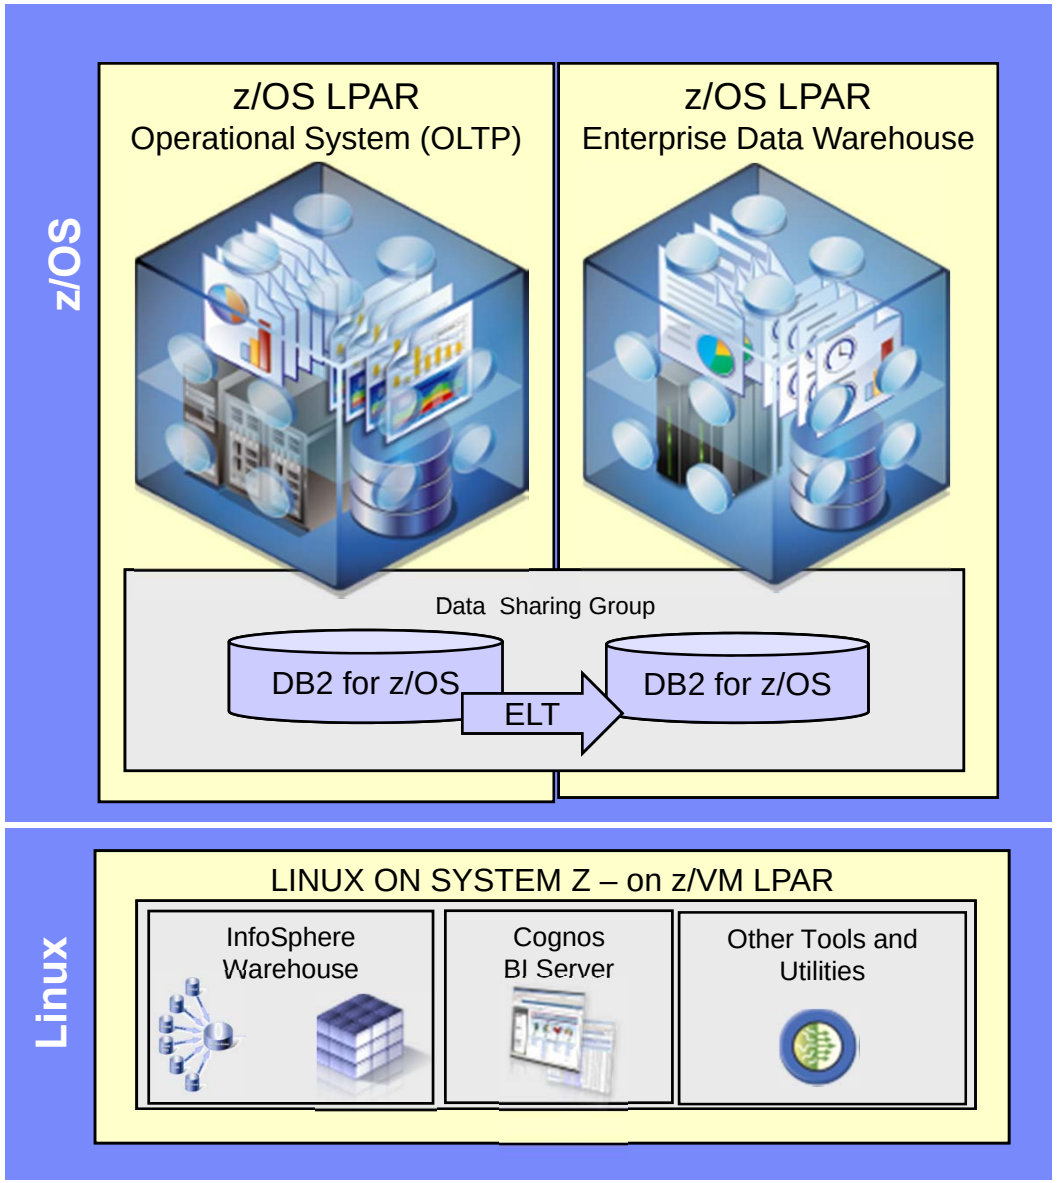

- Minimizes data movement between operational system and data warehouse
- Lowers data latency for time sensitive decisions
- Enables consolidation and simplification of data warehouse and data marts
- Leverages existing high availability, backup, disaster recovery, and security environments
- Provides greater scalability of multidimensional analysis through cubing services (data marts) and DB2 enhancements

DB2 10 for z/OS

CPU reductions for transactions, queries, and batch

- Out-of-the-box CPU reductions of 5-10% for traditional workloads
- Up to additional 10% CPU savings using new functions or avoiding constraints
- Out-of-the box CPU reductions of up to 20% for new workloads

Scales with less complexity and cost

- 5-10x more concurrent users – up to 20,000 per subsystem
- Significant scale-up capabilities in addition to existing scale-out support
- Consolidate to fewer LPARs and subsystems

Improved operational efficiencies and lower administration cost

- Automatic diagnostics, tuning, and compression

Even better performance

- Elapsed time improvement for small LOBS and Complex Queries

64 bit Evolution
Virtual Storage Relief

Temporal Data

Integrated XML Support

Query Processing Enhancements

Business Security & Compliance

Better Productivity

Standardized Cognos 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

The New 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

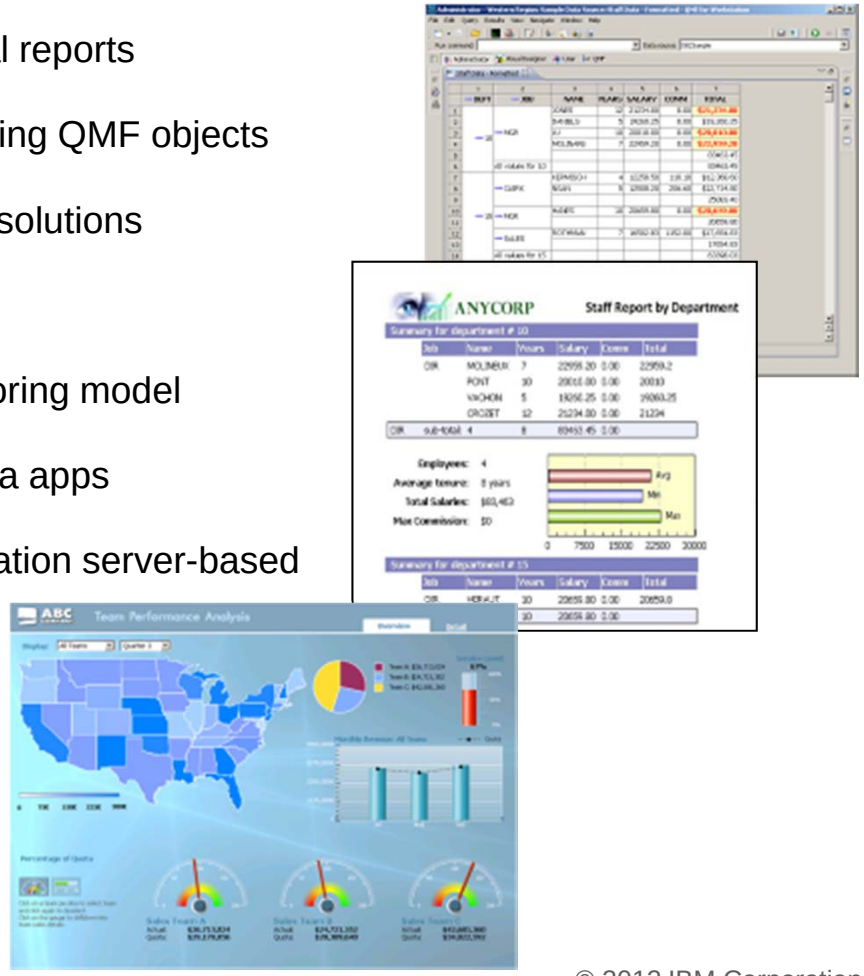

SPSS Modeler with InfoSphere Warehouse

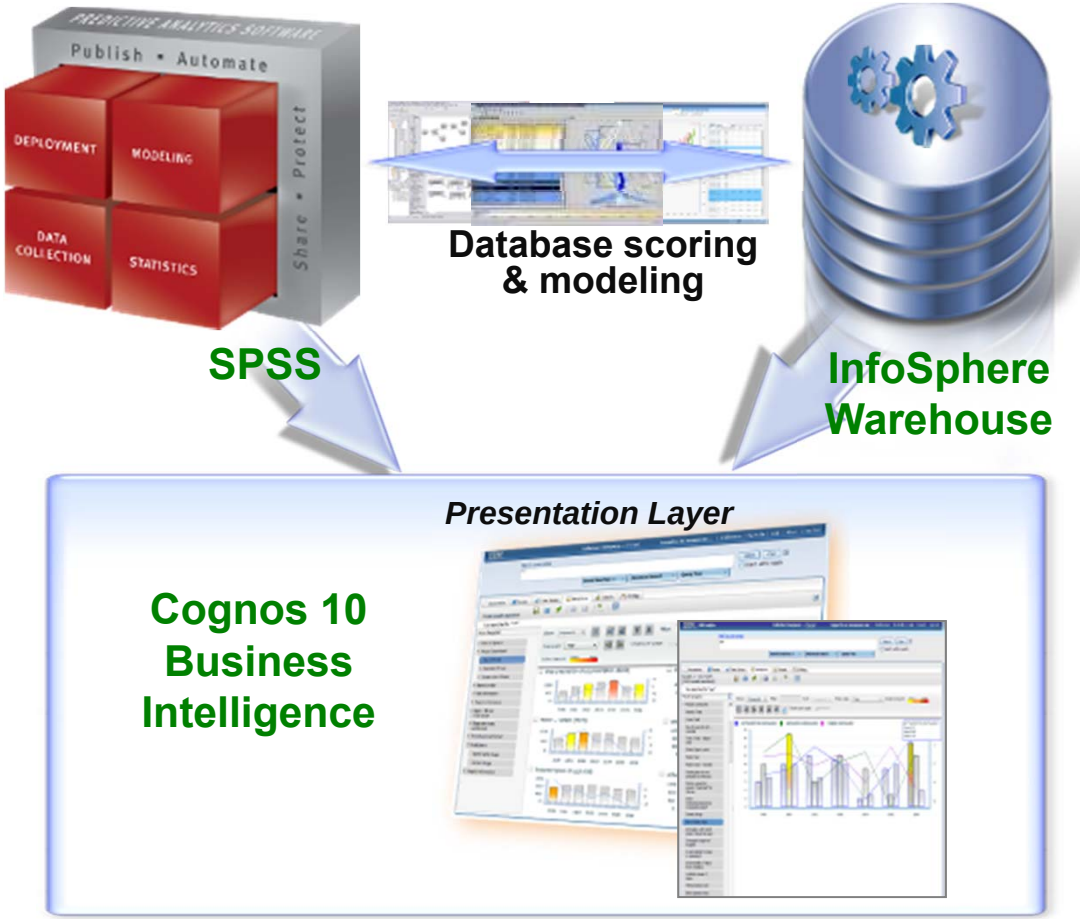

Full breadth of predictive analytics

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

DS8800 Storage

Cognos 10.1

Cognos. software

InfoSphere Warehouse 9.7.3

Data Warehouse

SPSS Modeler 14.2

SPSS AN IBM COMPANY

Cubing Services

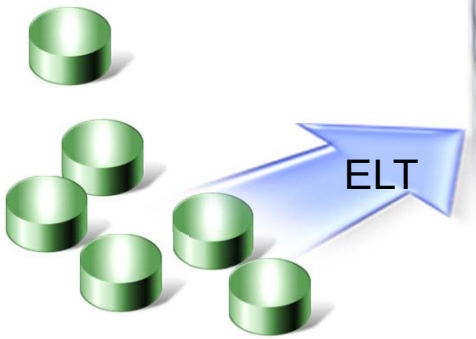

ELT

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

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

Fast Time to Value

- IBM DB2 Analytics Accelerator (Netezza 1000-12)
  - ➔ Production ready - 1 person, 2 days

- Table Acceleration Setup … **2 Hours**
  - DB2 “Add Accelerator”
  - Choose a Table for “Acceleration”
  - Load the Table (DB2 copy to Netezza)
  - Knowledge Transfer
  - Query Comparisons

- Initial Load Performance …
  - ➔400 GB “Loaded” in 29 Min
    570 million rows (Loads of 800GB to 1.3TB/Hr)
- Actual Query Acceleration … **1908x faster**
  - ➔2 Hours 39 Minutes to 5 Seconds
- CPU Utilization Reduction
  - ➔35% to ~0%

Actual customer results, October 2011

Performance & Savings

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

Queries run faster

- Save CPU resources
- People time
- Business opportunities

Actual customer results, October 2011

DB2 Analytics Accelerator: "we had this up and running in days with queries that ran over 1000 times faster"

DB2 Analytics Accelerator: "we expect ROI in less than 4 months"

http://www.livestream.com/ibmsoftware/video?clipId=pla_1bc6db16-ac1a-48c2-b50d-2ad13c6ba7ec

Advance to 31 minute mark for DB2 Analytics Accelerator section of keynote

Accelerating decisions to the speed of business

Performance & Savings

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

Queries run faster

- Save CPU resources
- People time
- Business opportunities

Actual customer results, October 2011

DB2 Analytics Accelerator: "we had this up and running in days with queries that ran over 1000 times faster"

DB2 Analytics Accelerator: "we expect ROI in less than 4 months"

Accelerating decisions to the speed of business

Deep DB2 Integration within zEnterprise

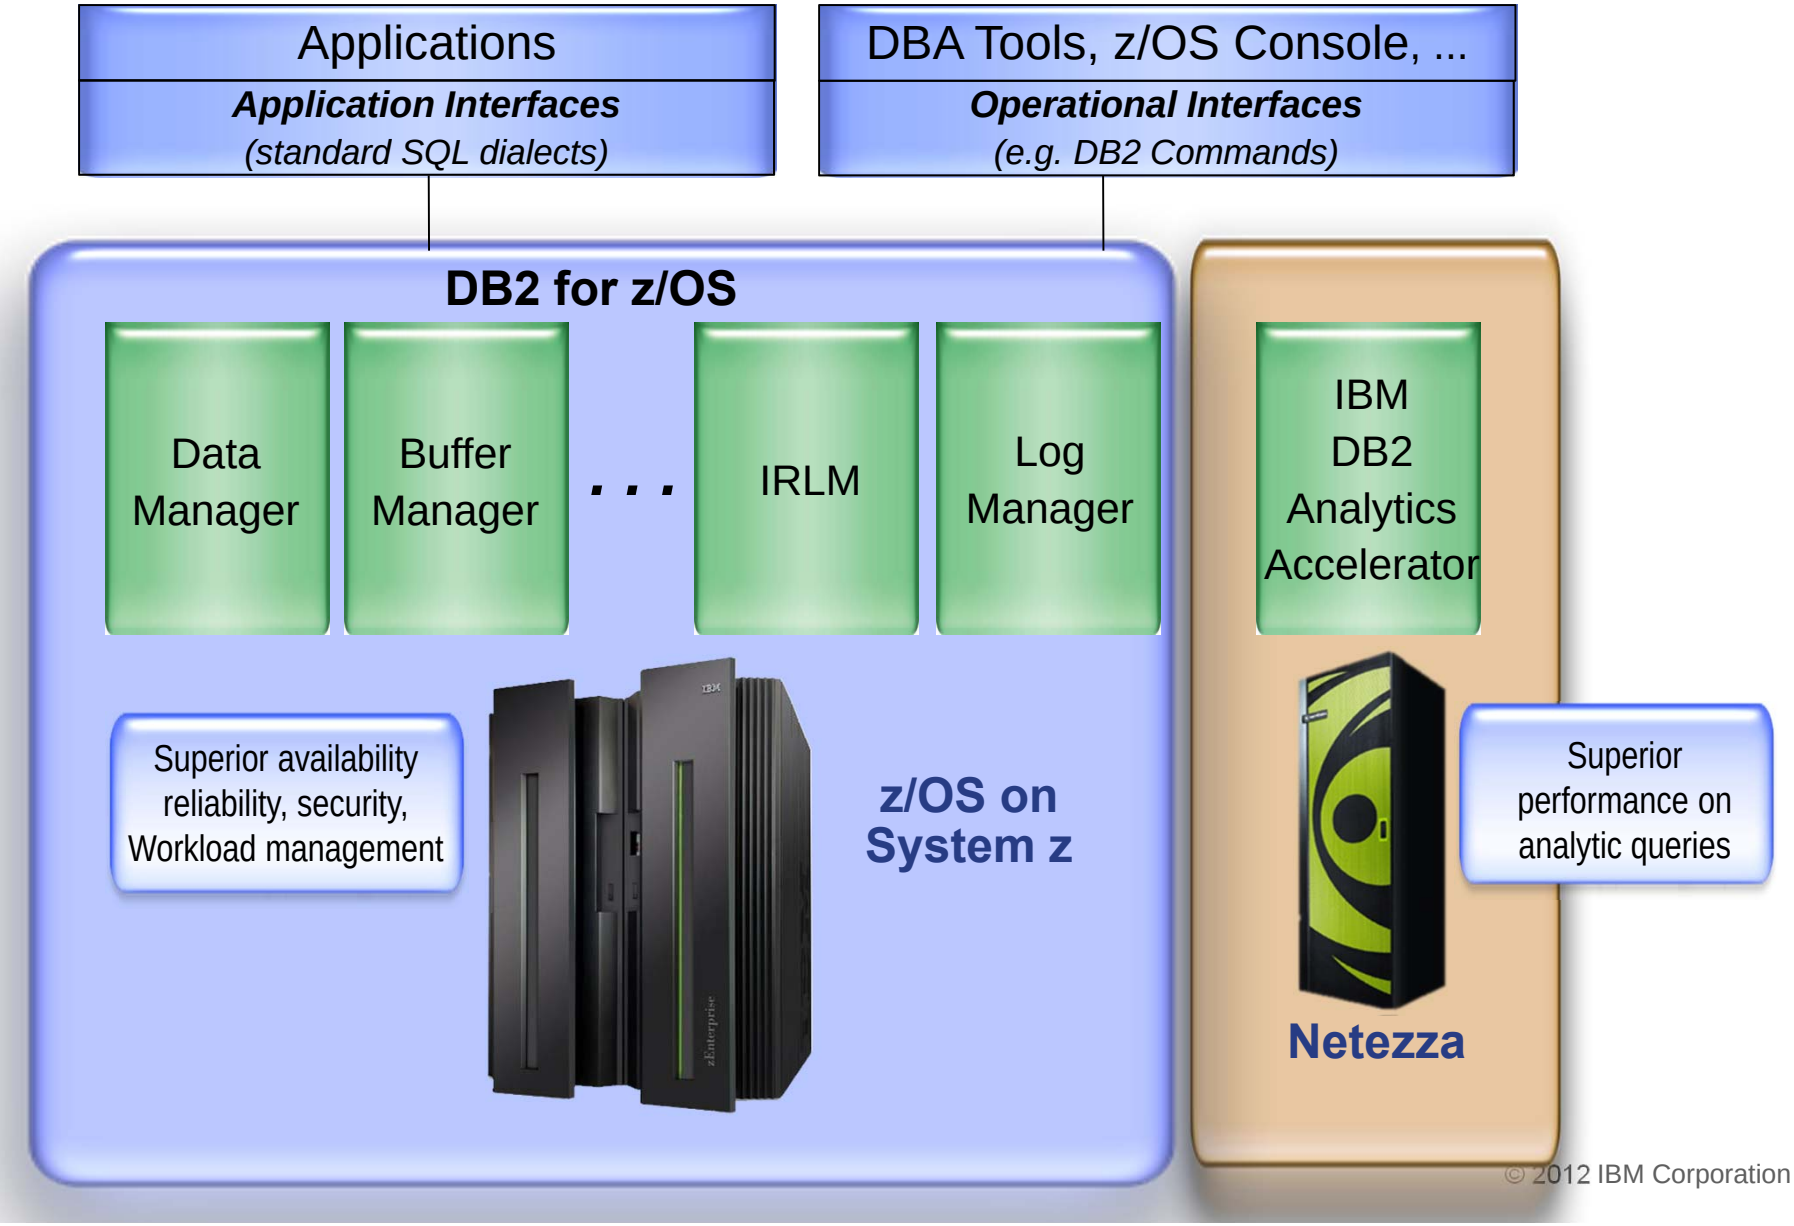

IBM DB2 Analytics Accelerator V2 Product Components

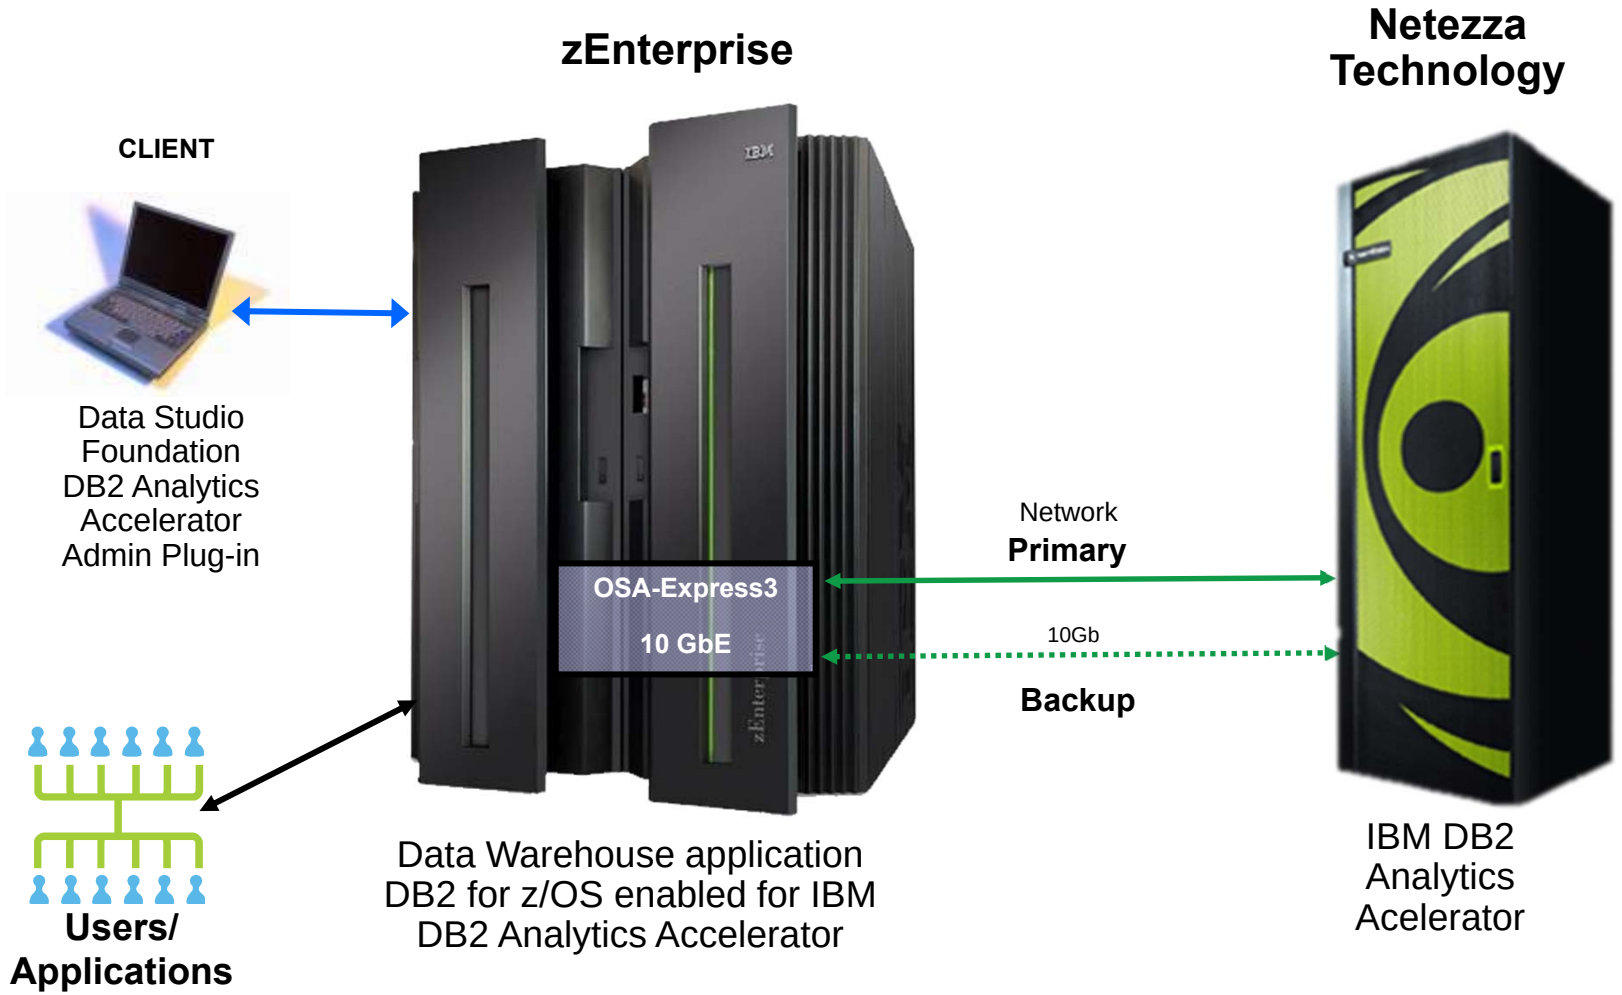

Analytics Accelerator Table Definition and Deployment

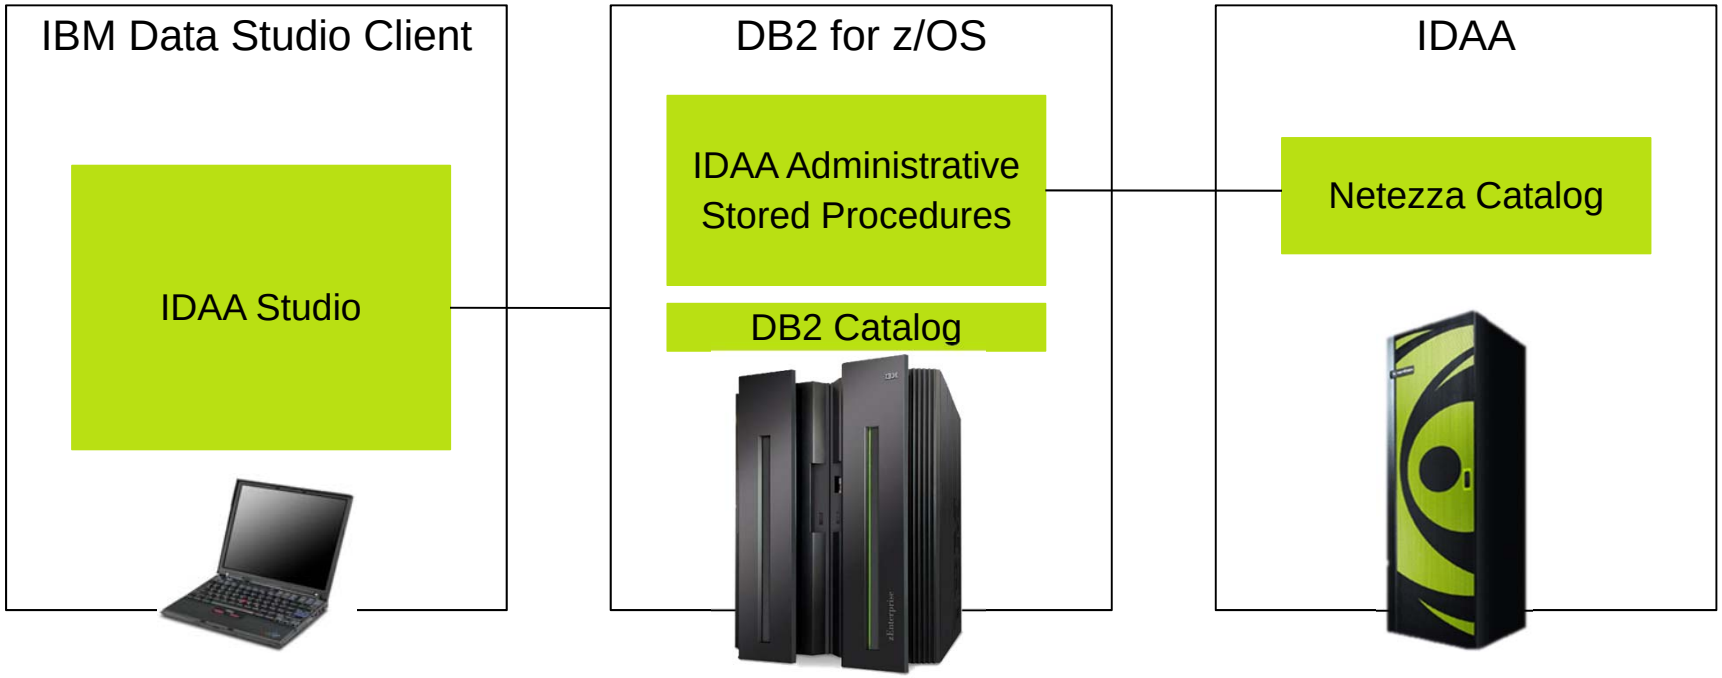

- The tables need to be defined and deployed to IDAA before data is loaded and queries sent to it for processing.
  - Definition: identifying tables for which queries need to be accelerated
  - Deployment: making tables known to DB2, i.e. storing table meta data in the DB2 and Netezza catalog.
- IBM DB2 Analytics Accelerator Studio guides you through the process of defining and deploying tables, as well as invoking other administrative tasks.
- IBM DB2 Analytics Accelerator Stored Procedures implement and execute various administrative operations such as table deployment, load and update, and serve as the primary administrative interface to IDAA from the outside world including IDAA Studio.

Query Execution Process Flow

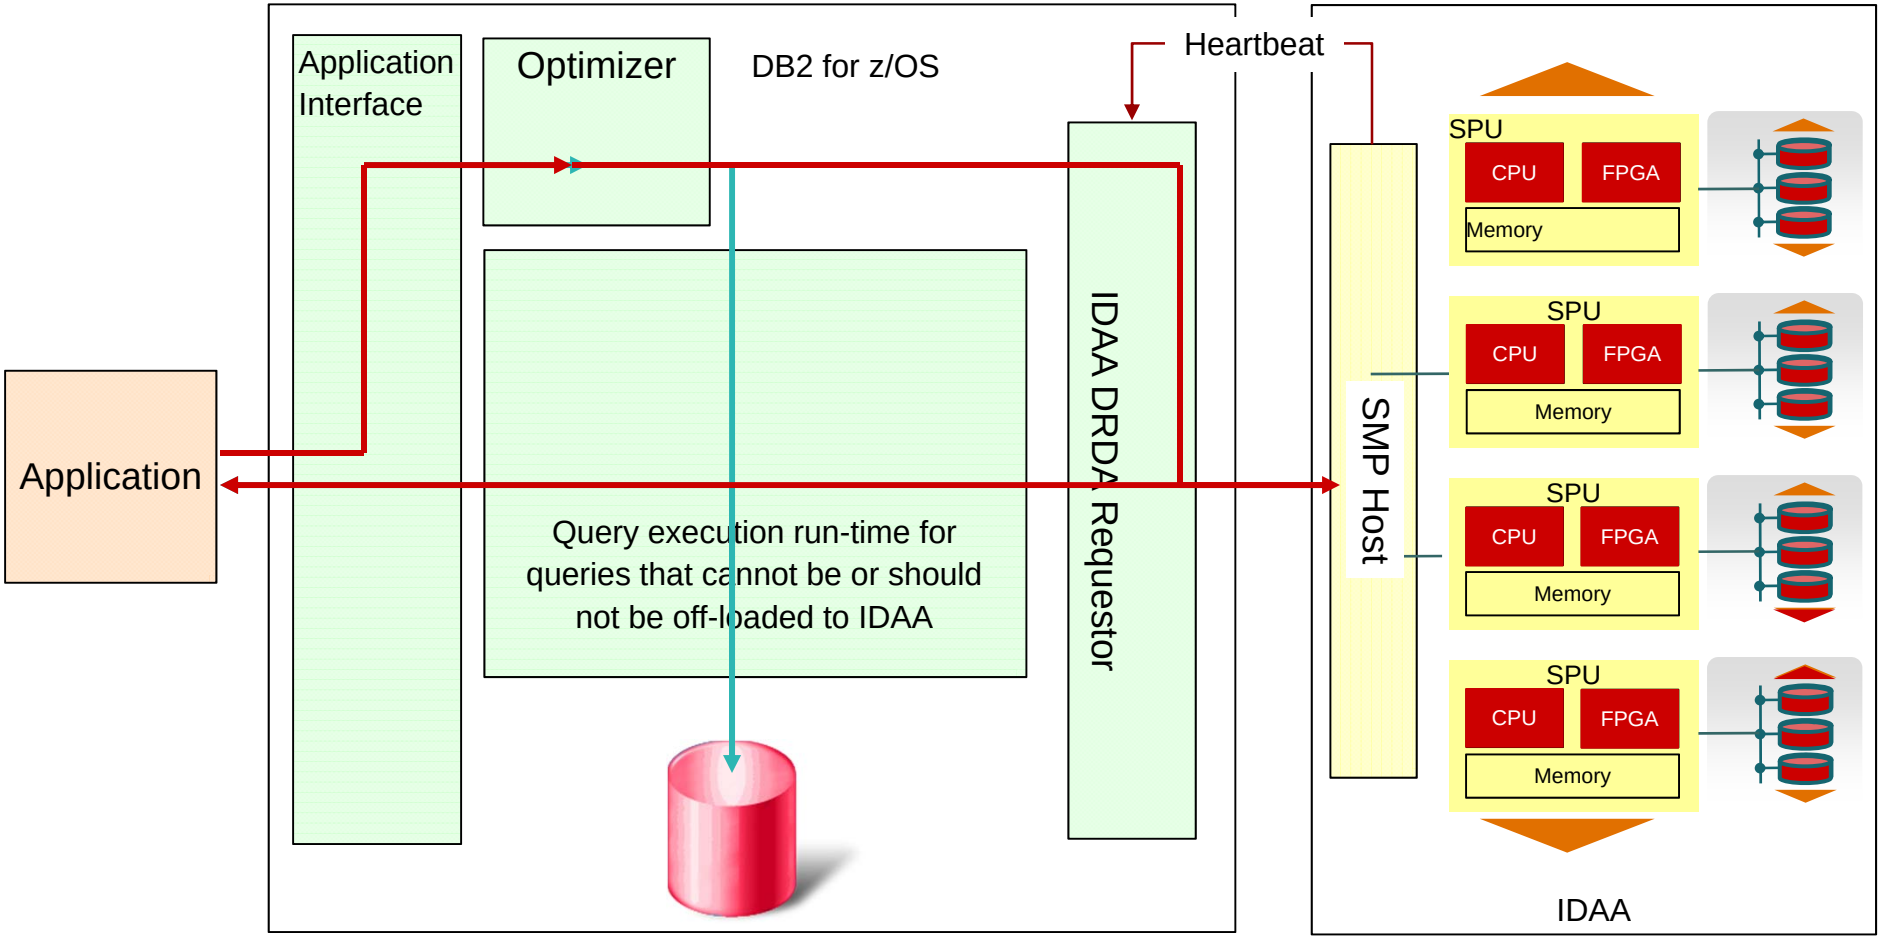

Queries executed without IDAA

Queries executed with IDAA

Heartbeat (IDAA availability and performance indicators)

DB2 Analytics Accelerator V2

Powered by Netezza 1000 Appliance

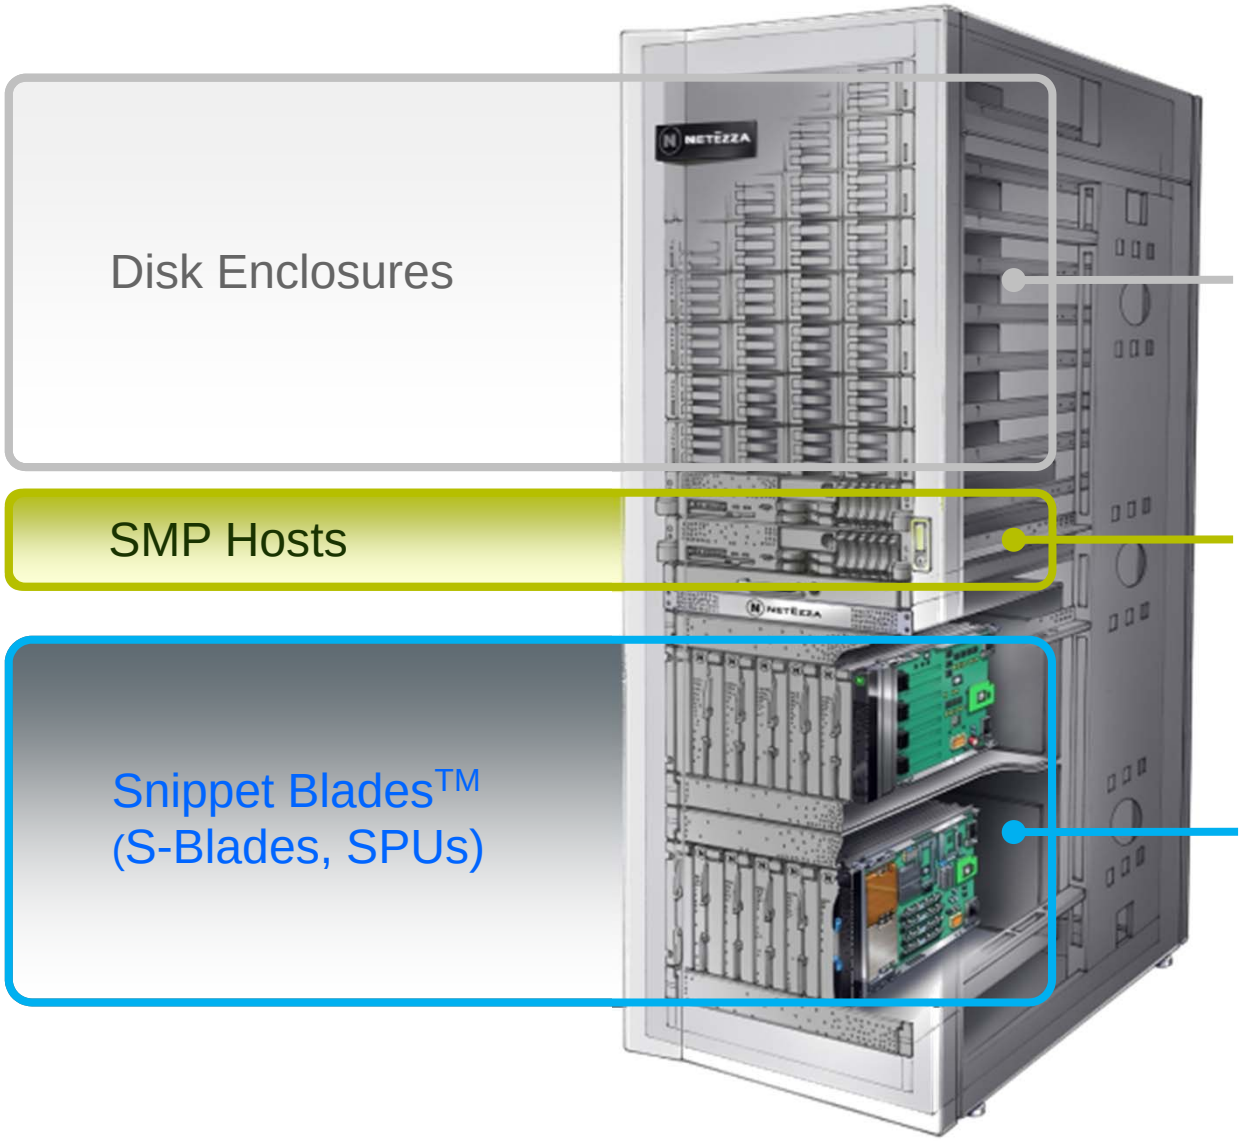

Slice of User Data
Swap and Mirror partitions
High speed data streaming
High compression rate
EXP3000 JBOD Enclosures
12 x 3.5" 1TB, 7200RPM, SAS (3Gb/s)
max 116MB/s (200-500MB/s compressed data)
e.g. TF12:
8 enclosures → 96 HDDs
32TB uncompressed user data (→ 128TB)

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

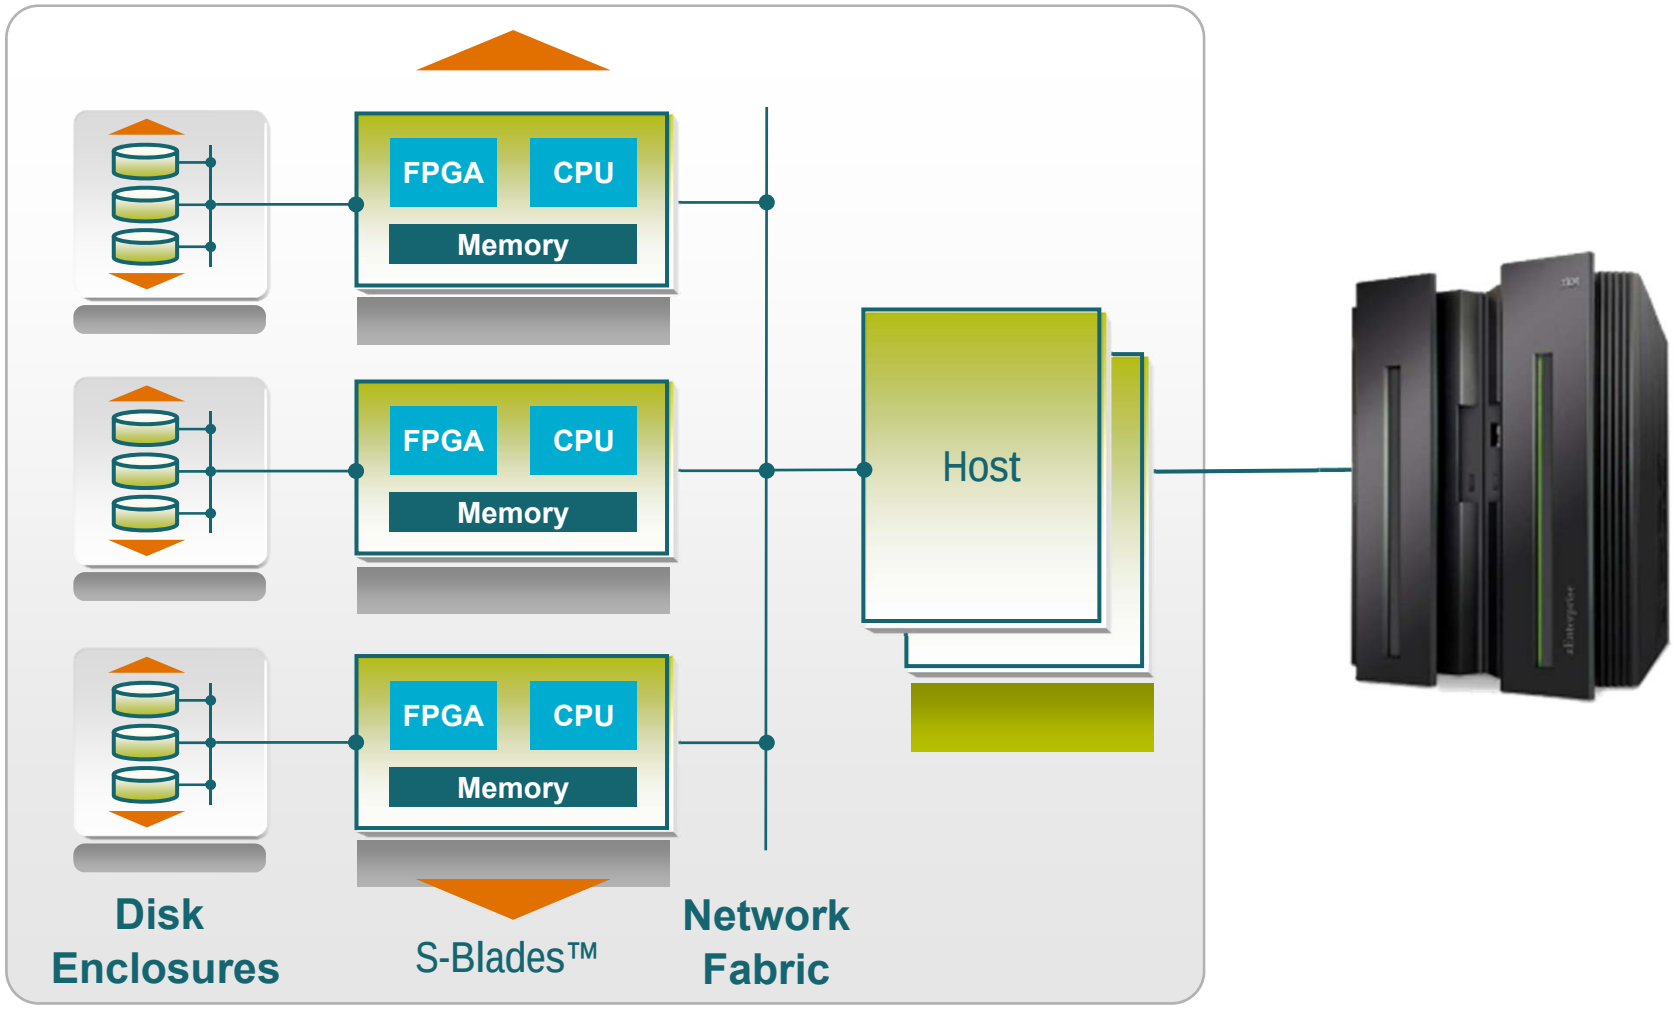

Netezza Appliance

The Key to the Speed

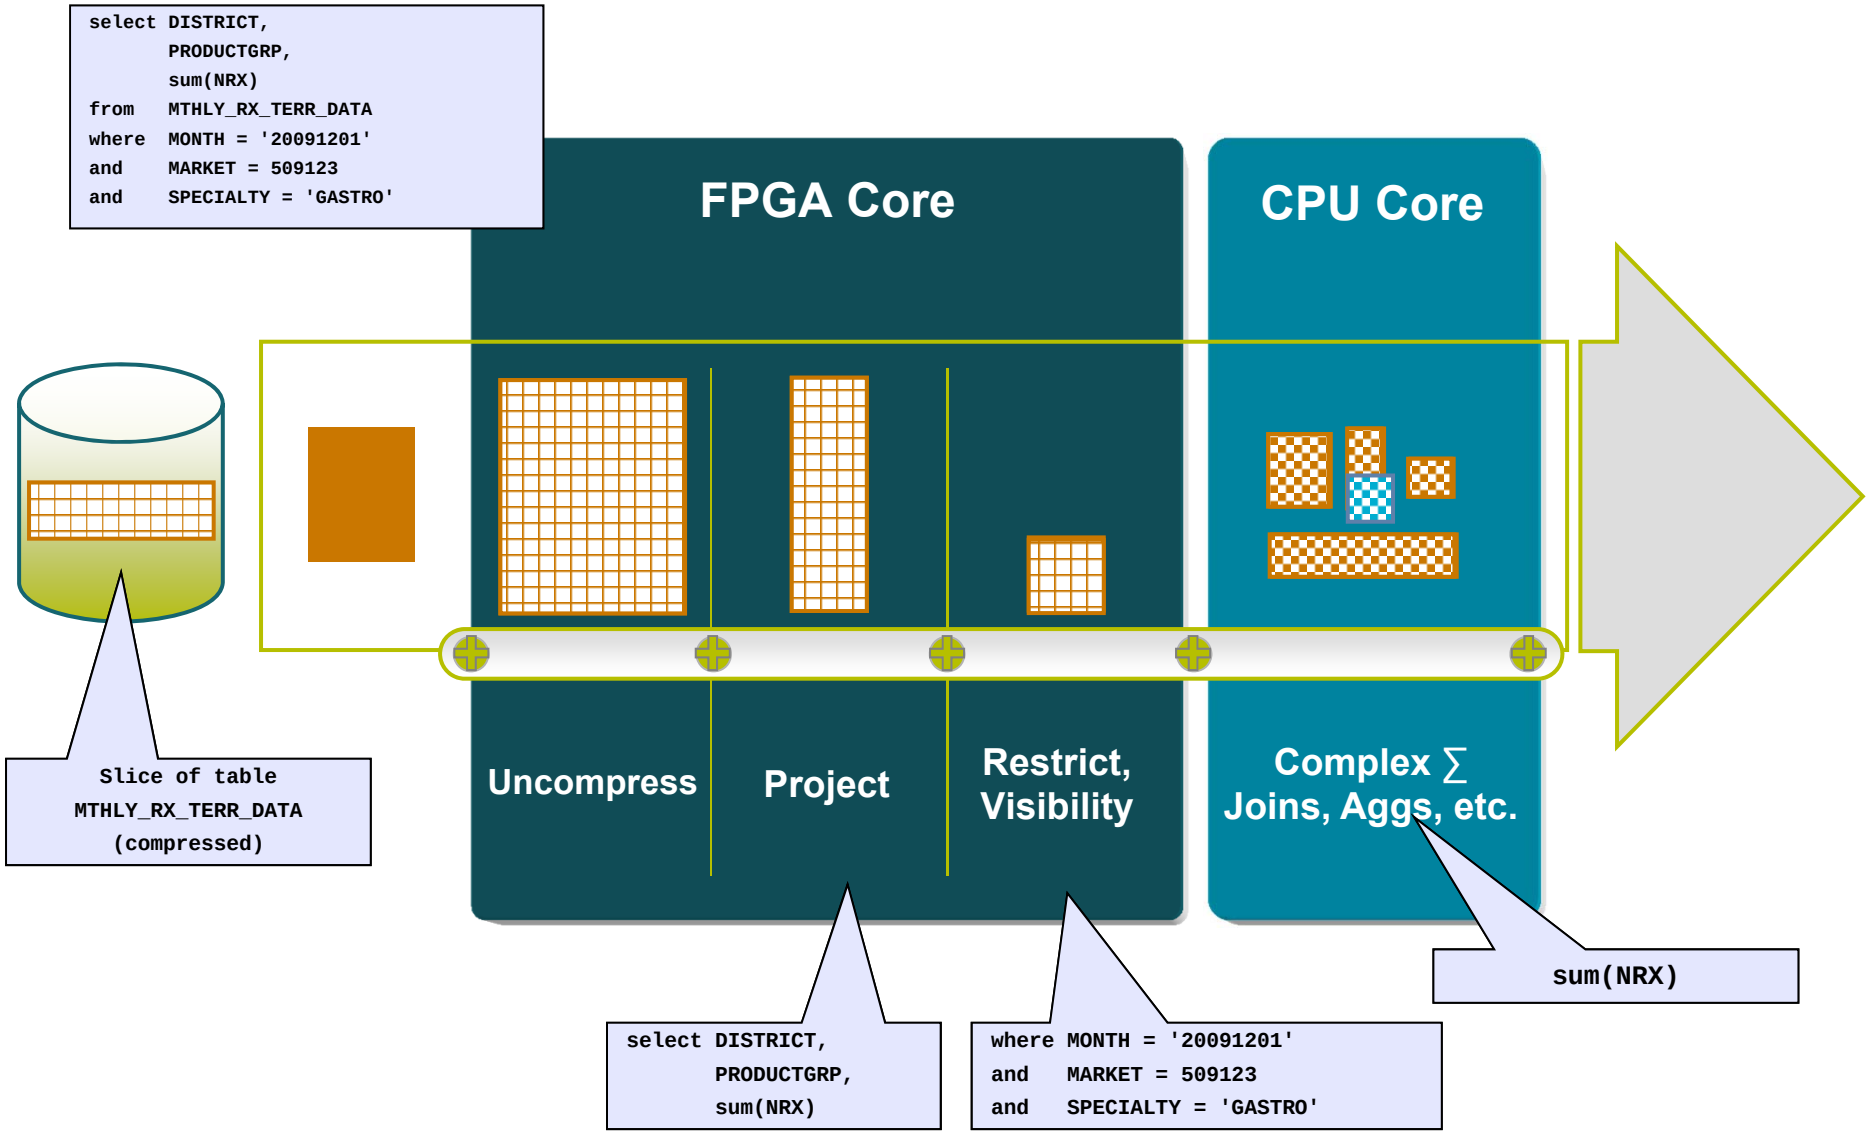

Netezza 1000 Appliance Scalability

	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

Future →

Predictable, Linear Scalability throughout entire family

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

*: 4X compression assumed

Accelerator Data Load

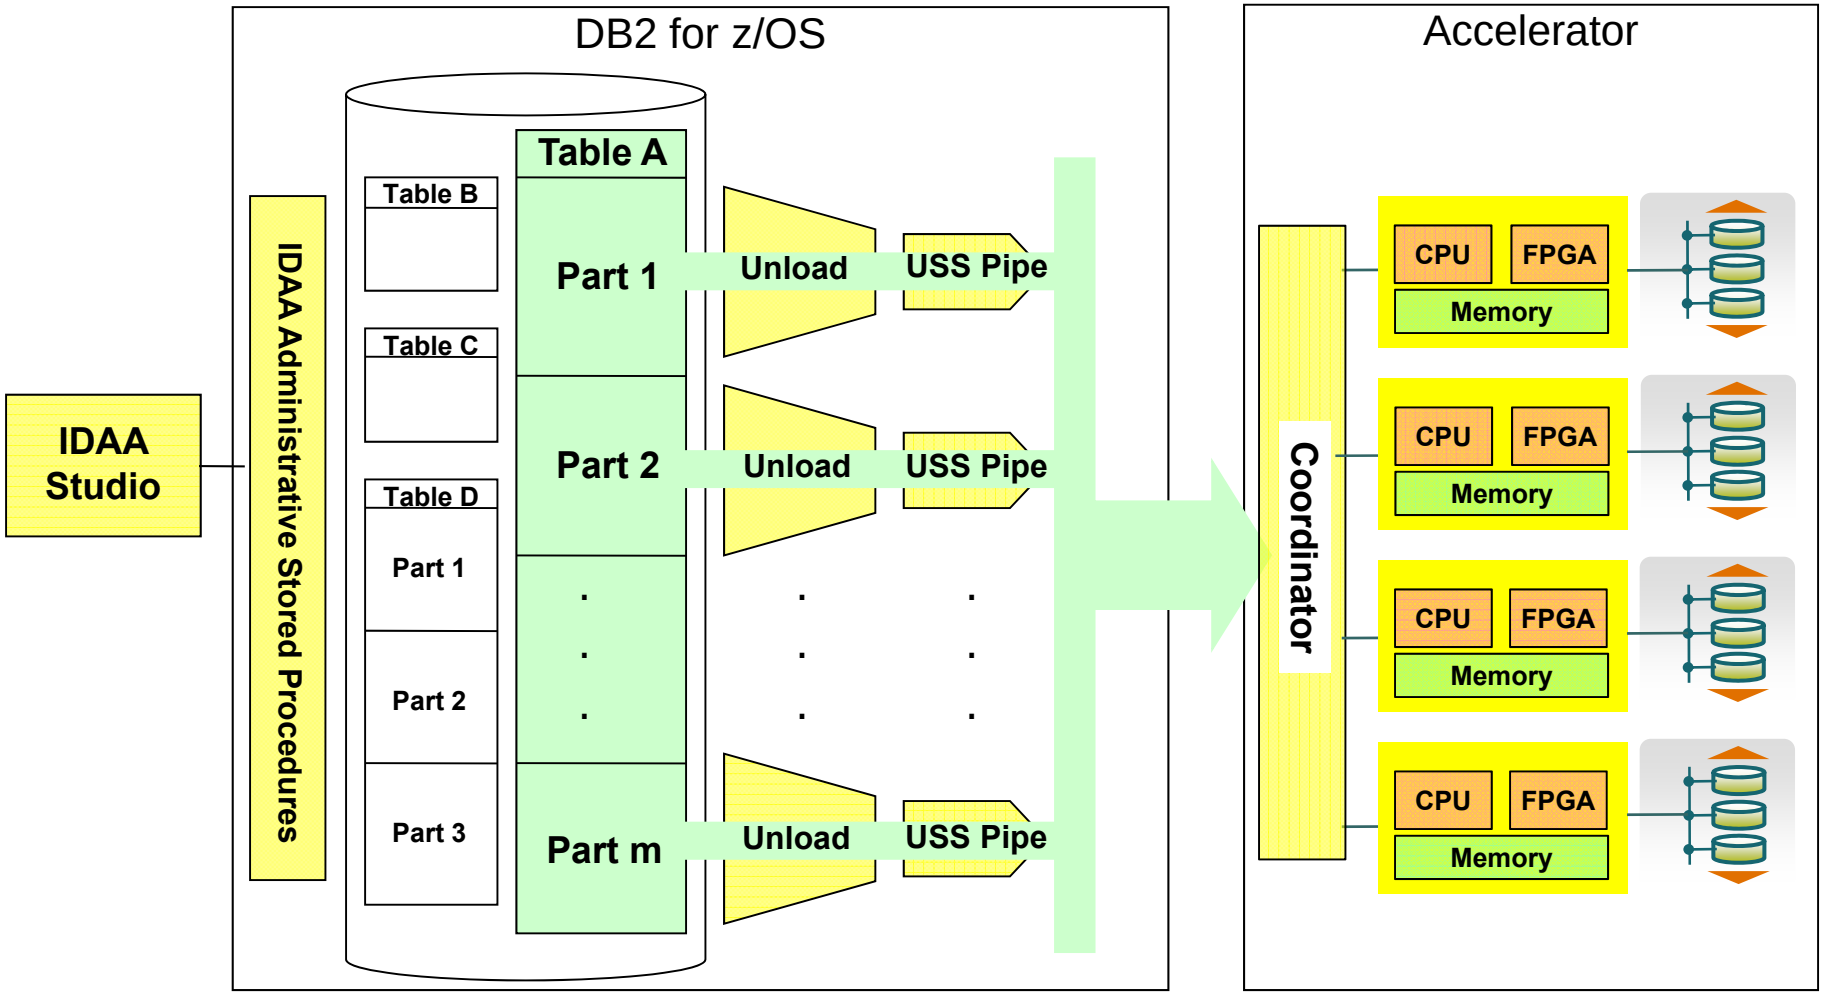

- 1 TB / h – can vary, depending on CPU resources, table partitioning, …
- Update on table partition level, concurrent queries allowed
- Trickle-feed update under discussion

Connectivity Options

Multiple DB2 systems can connect to a single IDAA

A single DB2 system can connect to multiple IDAAs

Multiple DB2 systems can connect to multiple IDAAs

Better utilization of IDAA 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

Shielding Against Disk Failures

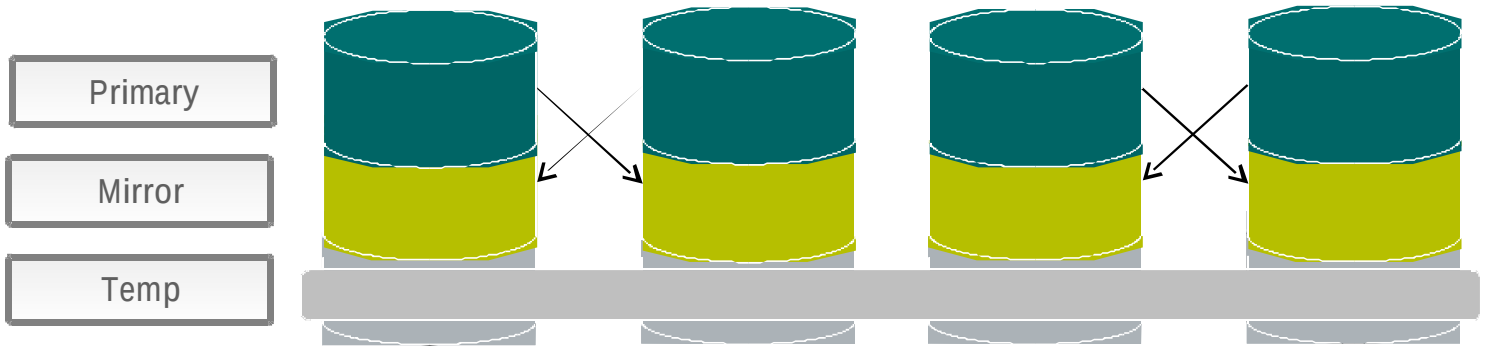

- All user data and temp space mirrored
- Disk failures transparent to queries and transactions
- Failed drives automatically regenerated
- Bad sectors automatically rewritten or relocated

Shielding Against S-BladeTM Failures

- S-Blade failure is automatically detected

Shielding Against S-BladeTM Failures

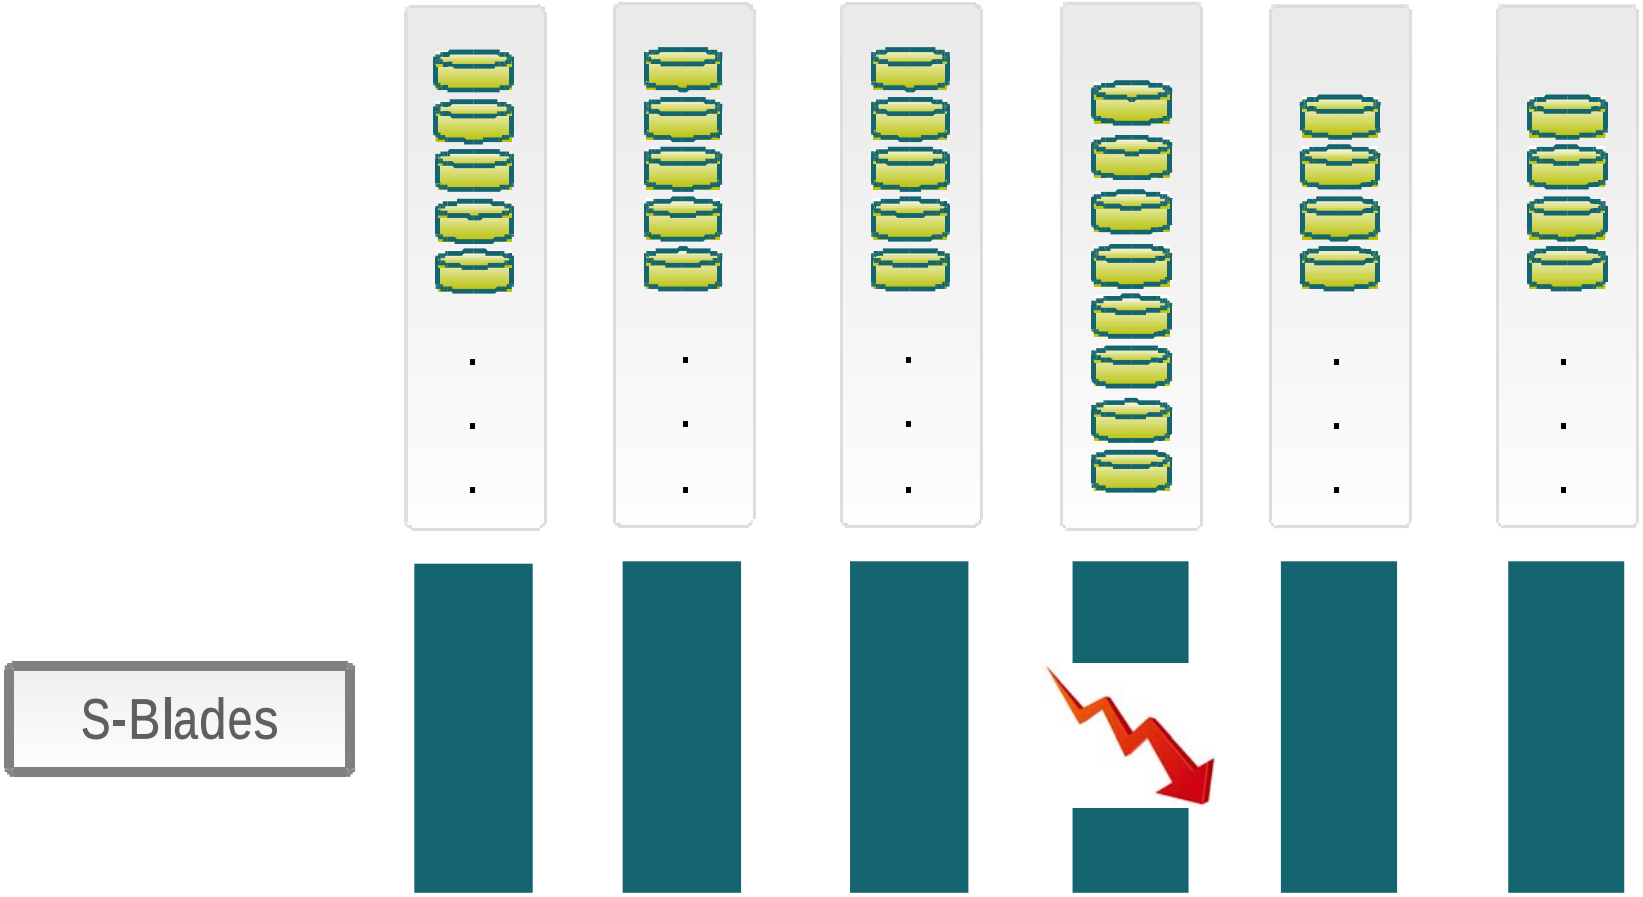

- Drives automatically reassigned to active S-Blades within a chassis
- Read-only queries (that have not returned data yet) automatically restarted
- Transactions and loads interrupted
- Loads automatically restarted from last successful checkpoint

Why Both?

Marrying the best of both worlds

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	IBM Netezza
Mixed Workload System	**Focused Appliance**
• 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	• 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

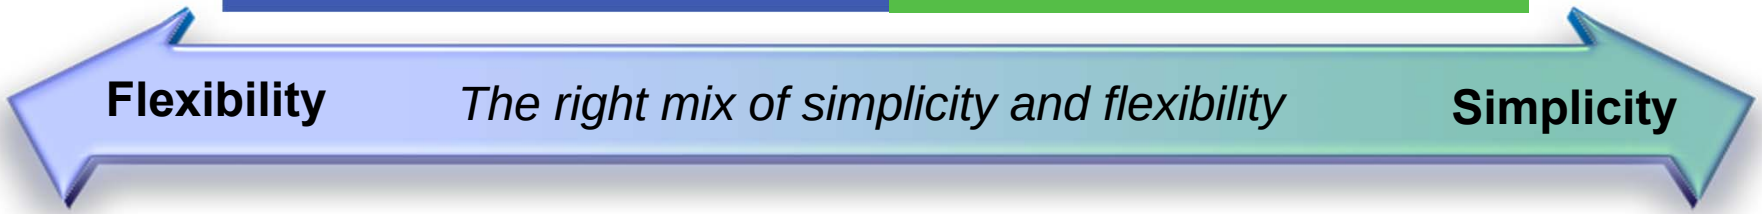

What is the value?

- Quickly delivers analytics to operational applications
- High speed analytics where the data is generated
- Enables train-of-thought analysis with high speed complex queries
- Substantially reduces operational costs by removing the need for complex query tuning
- Creates a highly secure environment for highly sensitive analysis (EAL5)
- Speeds batch reporting cycle to meet stricter SLAs
- Enables decision makers to perform business analysis they never dared in the past
- Enables query acceleration across multiple applications and systems
- Capitalizes on DB2 skills and certification removing the need to learn or convert to another SQL environment

"Back of the Envelope" ROI

Consider the MIPs of your z196 / z114
Consider software MLC reduction: z/OS, CICS, DB2…
Consider hardware value of MIPs redeployed

Examples of 6 month ROI:

- Avoiding 400 MIPs is roughly the cost of an IDAA Netezza 1000-3
- Avoiding 800 MIPs is roughly the cost of an IDAA Netezza 1000-6
- Avoiding 1600 MIPs is roughly the cost of an IDAA Netezza 1000-12

Next step: Quick Workload Test

- **Customer**
  - Collecting information from dynamic statement cache, supported by step-by-step instruction and REXX script (small effort for customer)
  - Uploading compressed file (up to some MB) to IBM FTP server
- **IBM / Center of Excellence**
  - Importing data into local database
  - Quick analysis based on known DB2 Analytics Accelerator capabilities

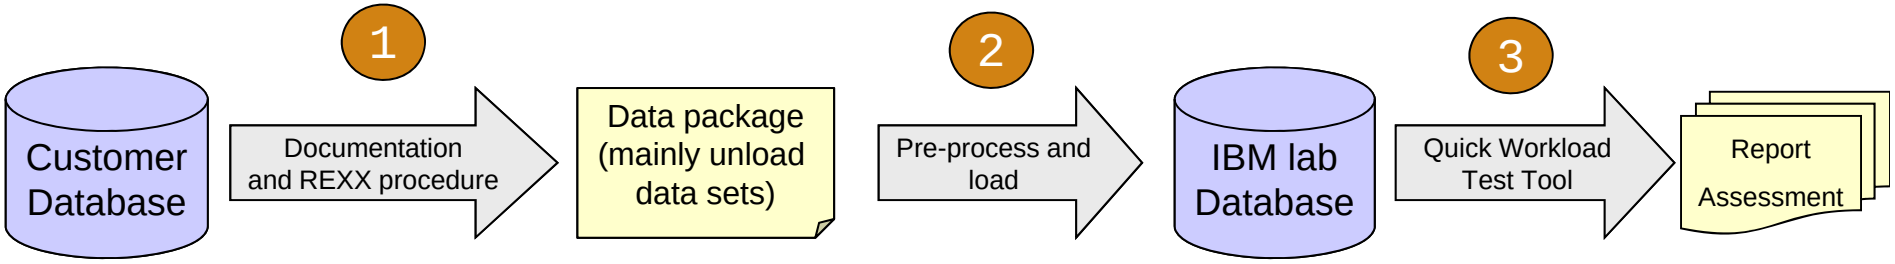

IDAA 2012 Roadmap

IBM DB2 Analytics Accelerator Vnext - Highlights

Further strengthen DB2 for z/OS competitive position to host mission critical operational BI workloads

- New capability: Online Storage Server
  - Exploit the storage capacity of the Netezza appliance for a true “multi-temperature” data solution controlled by DB2 for z/OS.
  - Less Cost for DataMarts on System z through reduced data duplication.
- Significantly enhanced Data Currency
  - Advanced asynchronous data propagation with change data capture technology
- Improved Performance
  - Faster data load process with reduced CPU requirements
- Enhanced Query Support
- Platform Support
  - More IBM Netezza systems

Beta program starting in July 2012

Prerequisites: zEnterprise (z196, z114) - DB2 10

Today's Data Life Cycle "Architecture"

- *Async and Distributed*

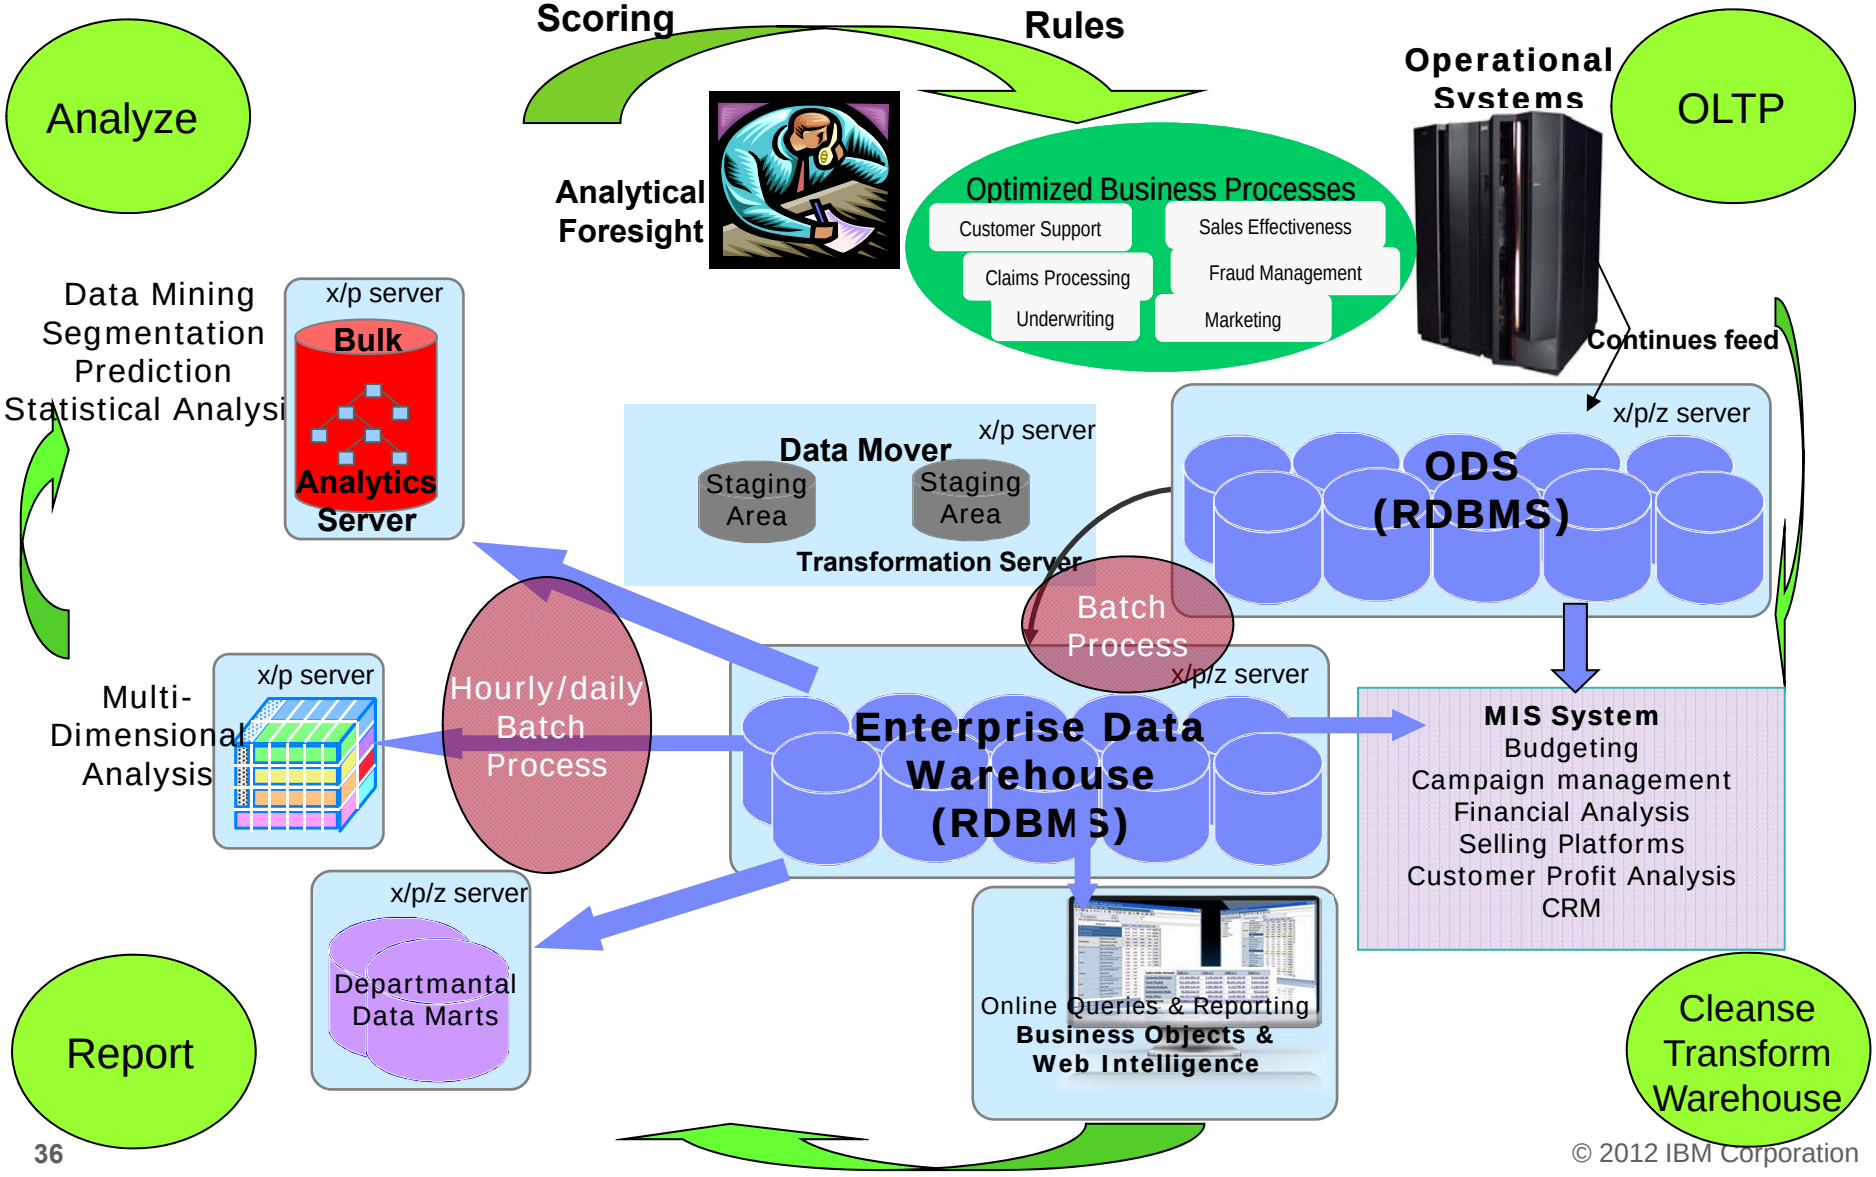

Fit The Solution To The Workload – Access to all Data

Business Analytics / ODS	OLTP
■Complex queries against a data warehouse ■Read only ■Star schema design often used to speed up queries ■Different modes of operation	▪Transactions against operational data ▪Reads and writes ▪Normalized schema design to eliminate redundancy ▪Multiple user throughput operation

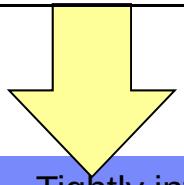

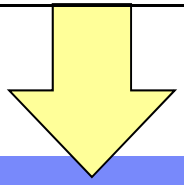

Tightly integrated solution optimized for workloads

Simplification

Single entry point to the true data

High fidelity data

Performance

IDAA (IBM DB2 Analytics Accelerator) DB2 z/OS ODS members	DB2 z/OS OLTP members

37

Business Analytics Life Cycle Architecture on System z

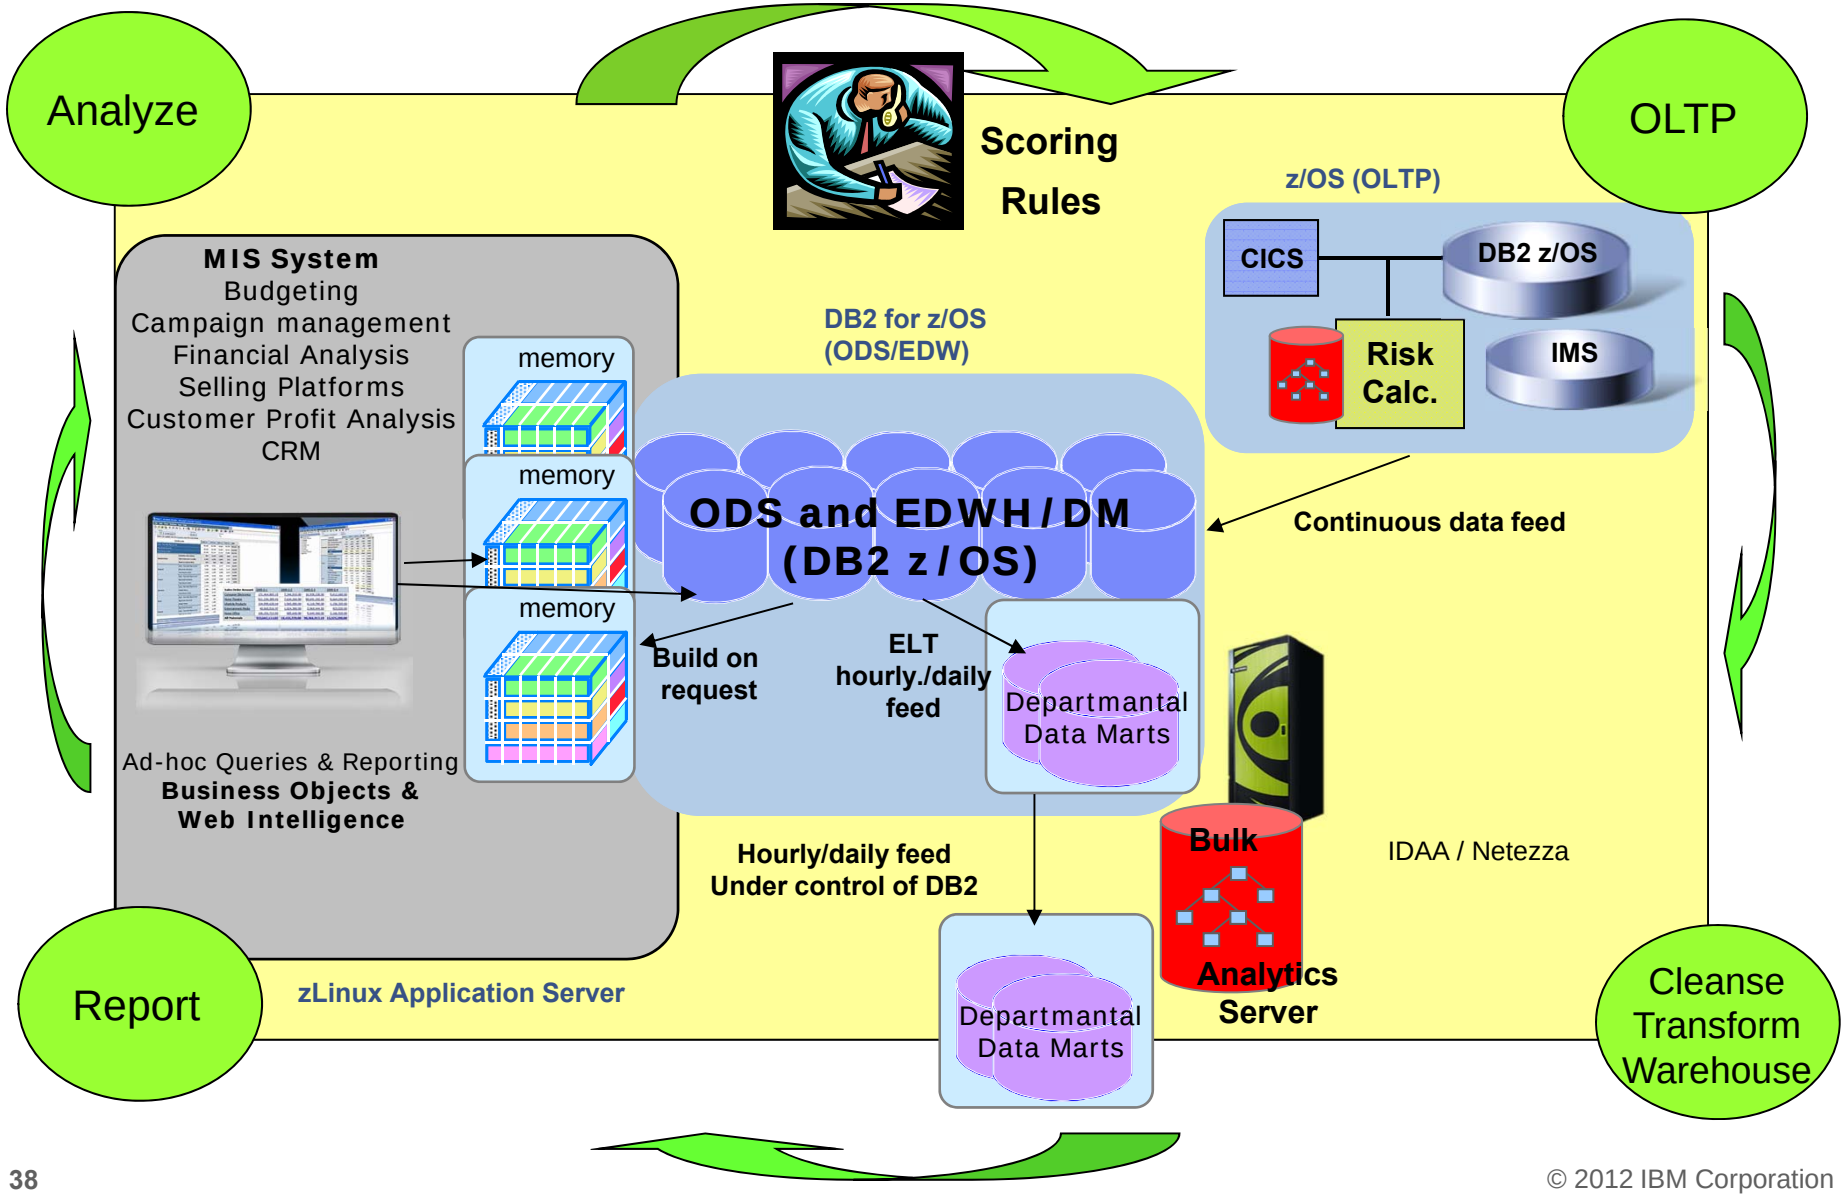

The Ultimate Consolidation Platform

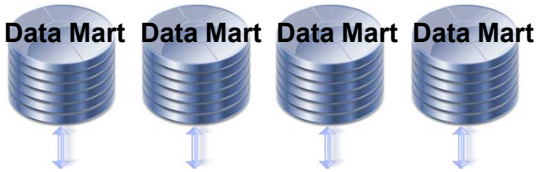

Data Mart Consolidation

Transaction Systems (OLTP)

Data Warehousing
Business Intelligence
Predictive Analytics

System z PR/SM
Recognized leader in mixed virtualization and workload isolation

z/OS:
Recognized leader in mixed workloads with security, availability and recoverability

IDAA: *Powered by Netezza for cost-effective high speed deep analytics*

Together:
Destroying the myth that transactional and decision support workloads have to be on separate platforms

Bringing it all together

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

In Summary, IBM Continues To Invest:

Your Partner in Business Analytics and Data Warehousing Solutions

- More than $14B in Acquisitions Since 2005
- More than 10,000 Technical Professionals
- More than 7,500 Dedicated Consultants
- Largest Math Department in Private Industry
- More than 27,000 Business Partner Certifications