

IBM DB2 Analytics Accelerator

Optimizing Your Key Business Decisions

Udo Hertz, Dir. Information Management Development Boeblingen



- Business and Technology Drivers

- Data Warehousing Solution on zEnterprise
 - IDAA Key Design and Operational Features
 - Powered by Netezza

- IDAA Use Cases and Next Steps



Knowing what happened is no longer adequate.

Business leaders need to know

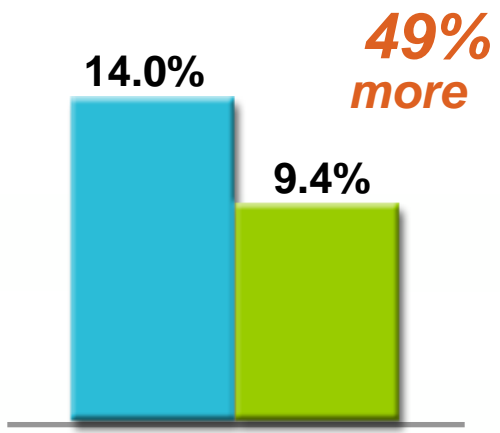
what is happening now,

what is likely to happen next and


what actions they should take.

.... Consistently Outperform


Revenue Growth
5 Year CAGR (2004-2008)



Profit Growth
5 Year CAGR (2004-2008)



Return on Invested Capital
5 Year Average (2004-2008)



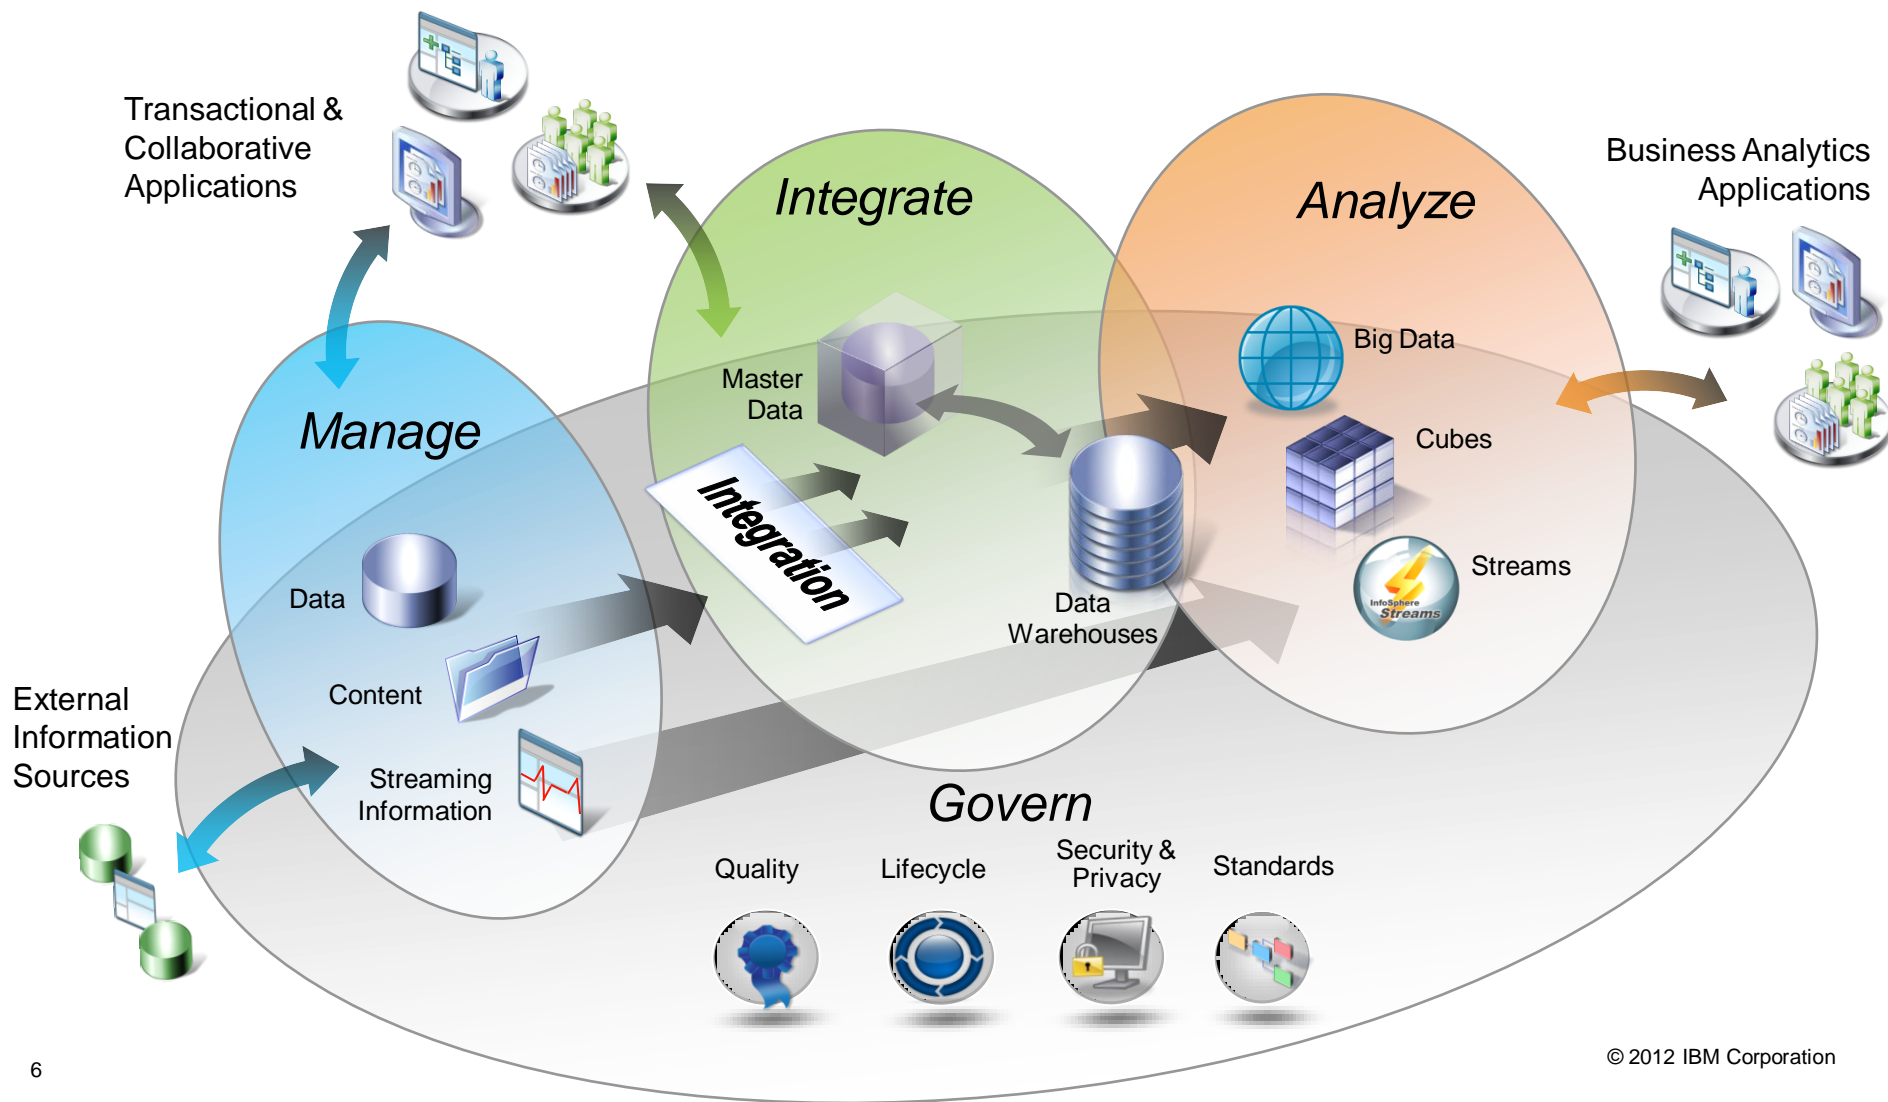
- Finance organizations with business insight
- All other enterprises

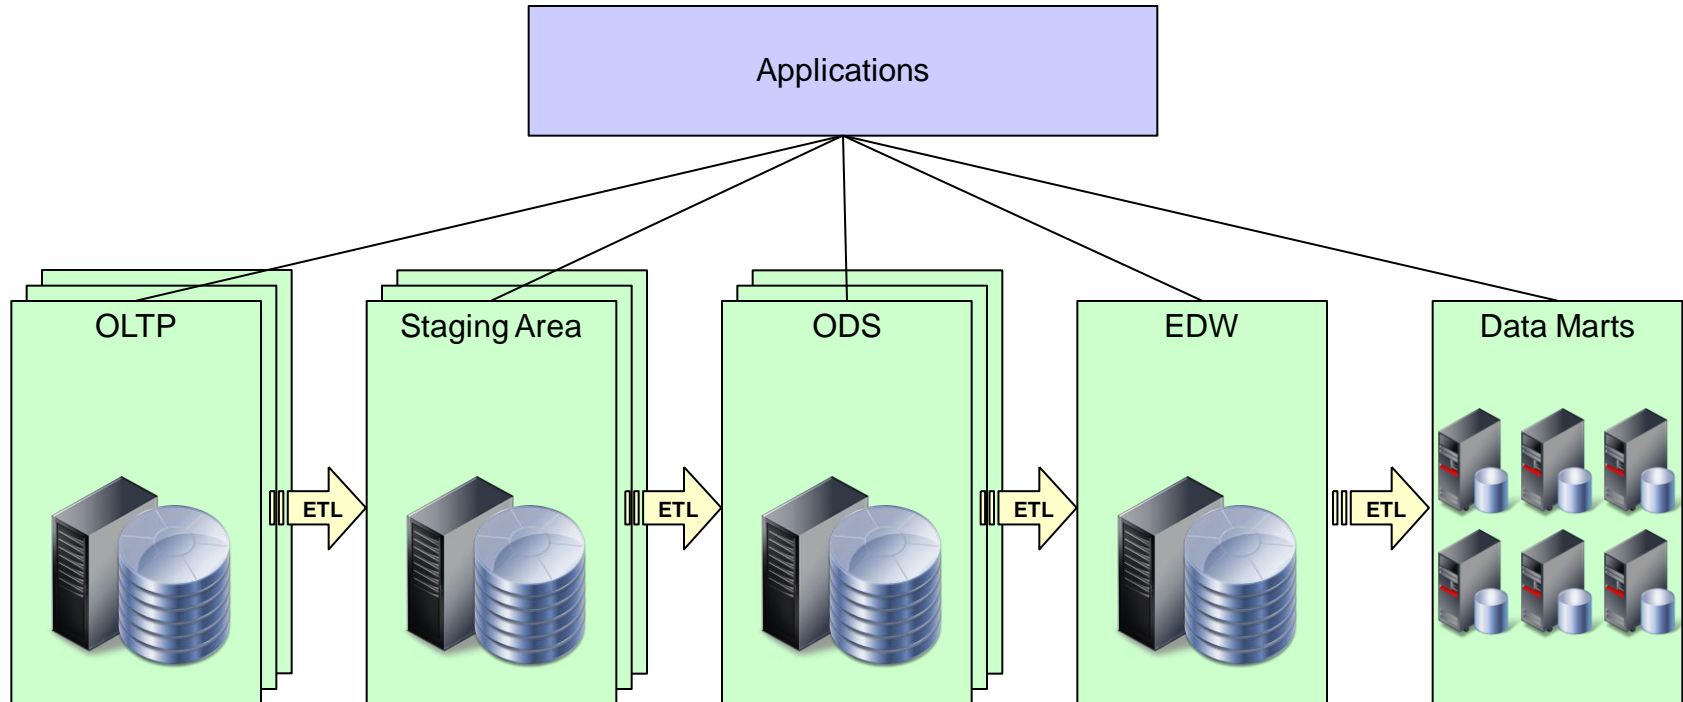
Revenue Growth: N = 580; Profit: N = 435; ROIC: N = 606
Source: IBM Global Business Services, The Global CFO Study 2010

- 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



Delivering trusted information for smarter business decisions across your entire information supply chain



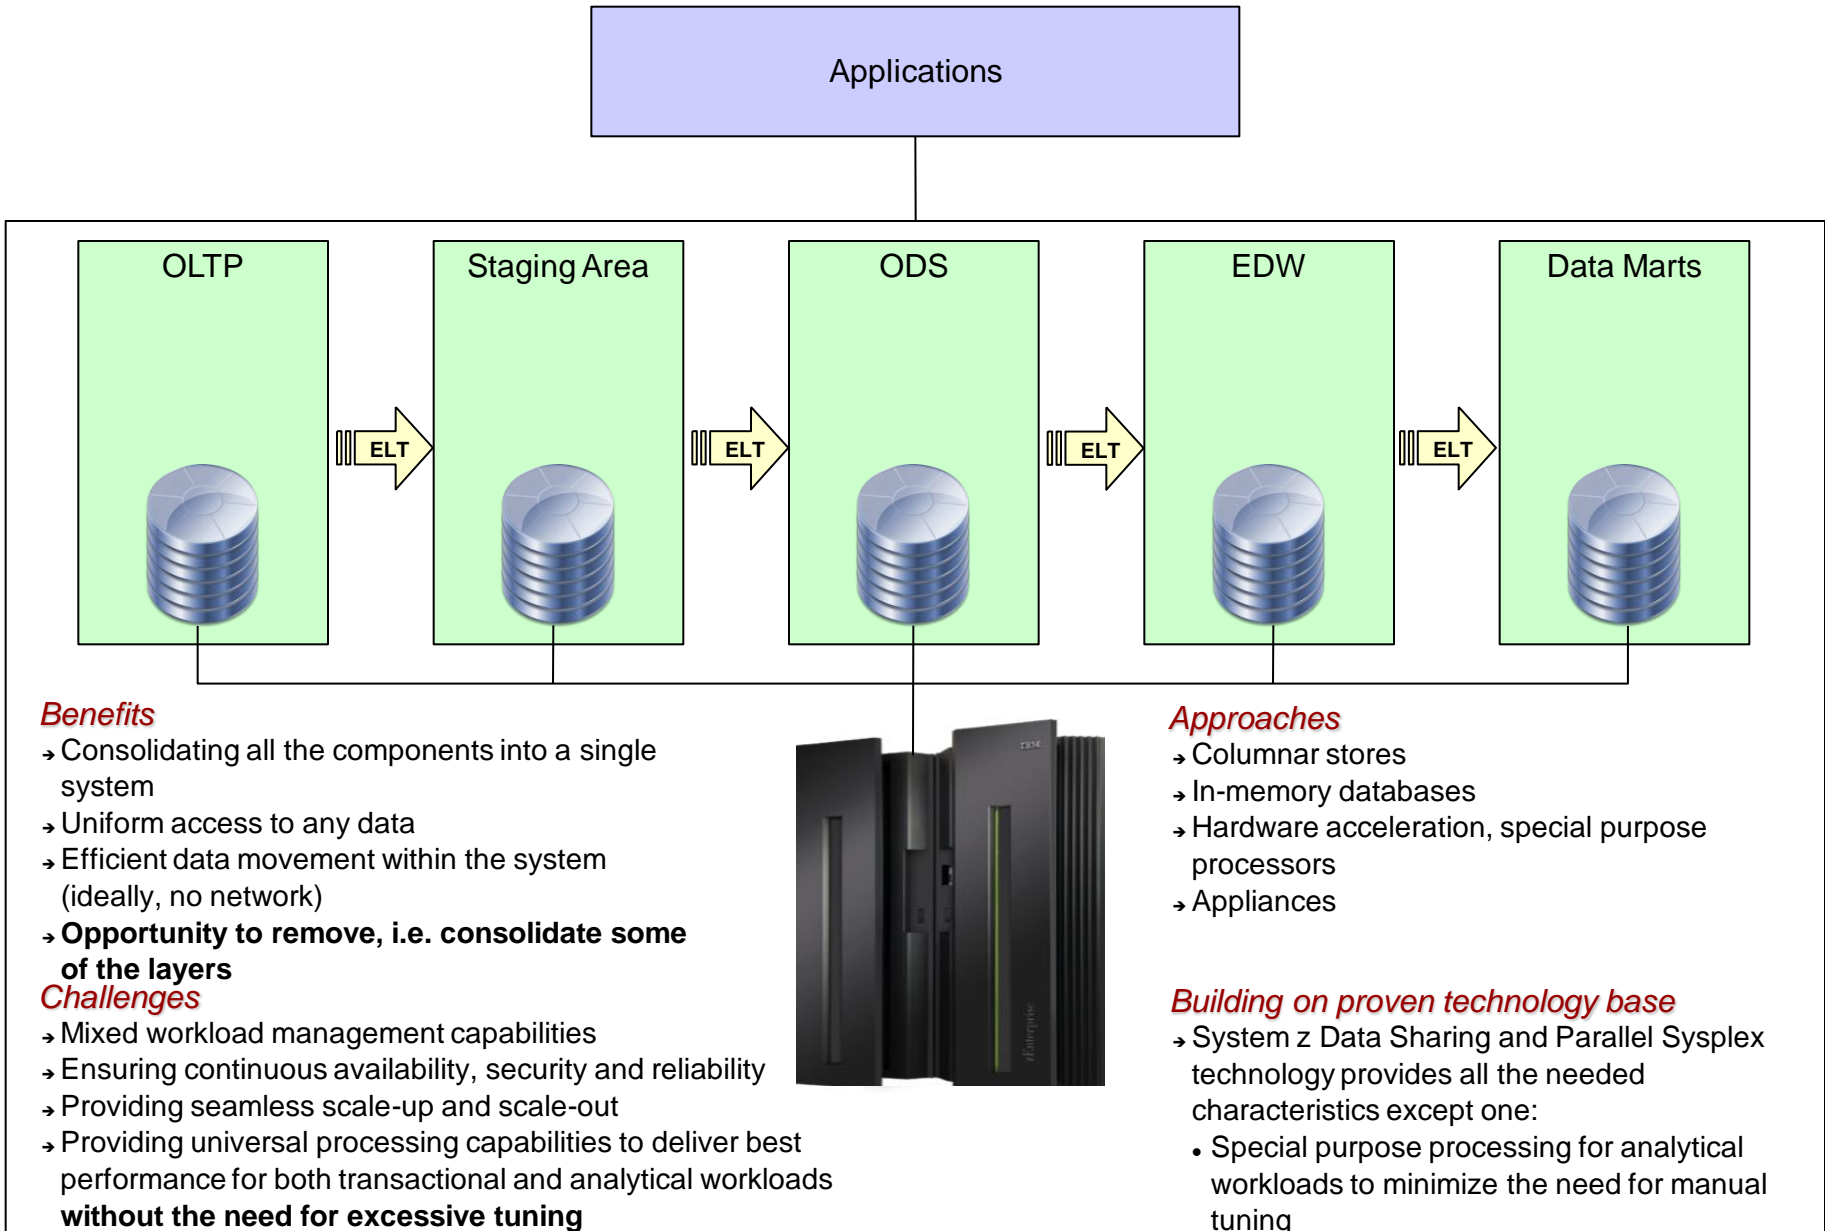


Historical reasons:

- Different access patterns
 - impact on performance
- EDW as the data integration hub
 - again, impact on performance
- Different life-cycle characteristics
 - and again, impact on performance
- Different Service Level Agreements (SLA)
 - Lack of broadly available workload management capabilities
 - Choice of lower cost-of-acquisition offerings

Negative ramifications:

- Complexity
 - both in systems management and in applications
- Difficulties in supporting real time analytics
- Inability to match ever more demanding SLA requirements
- High total cost of ownership

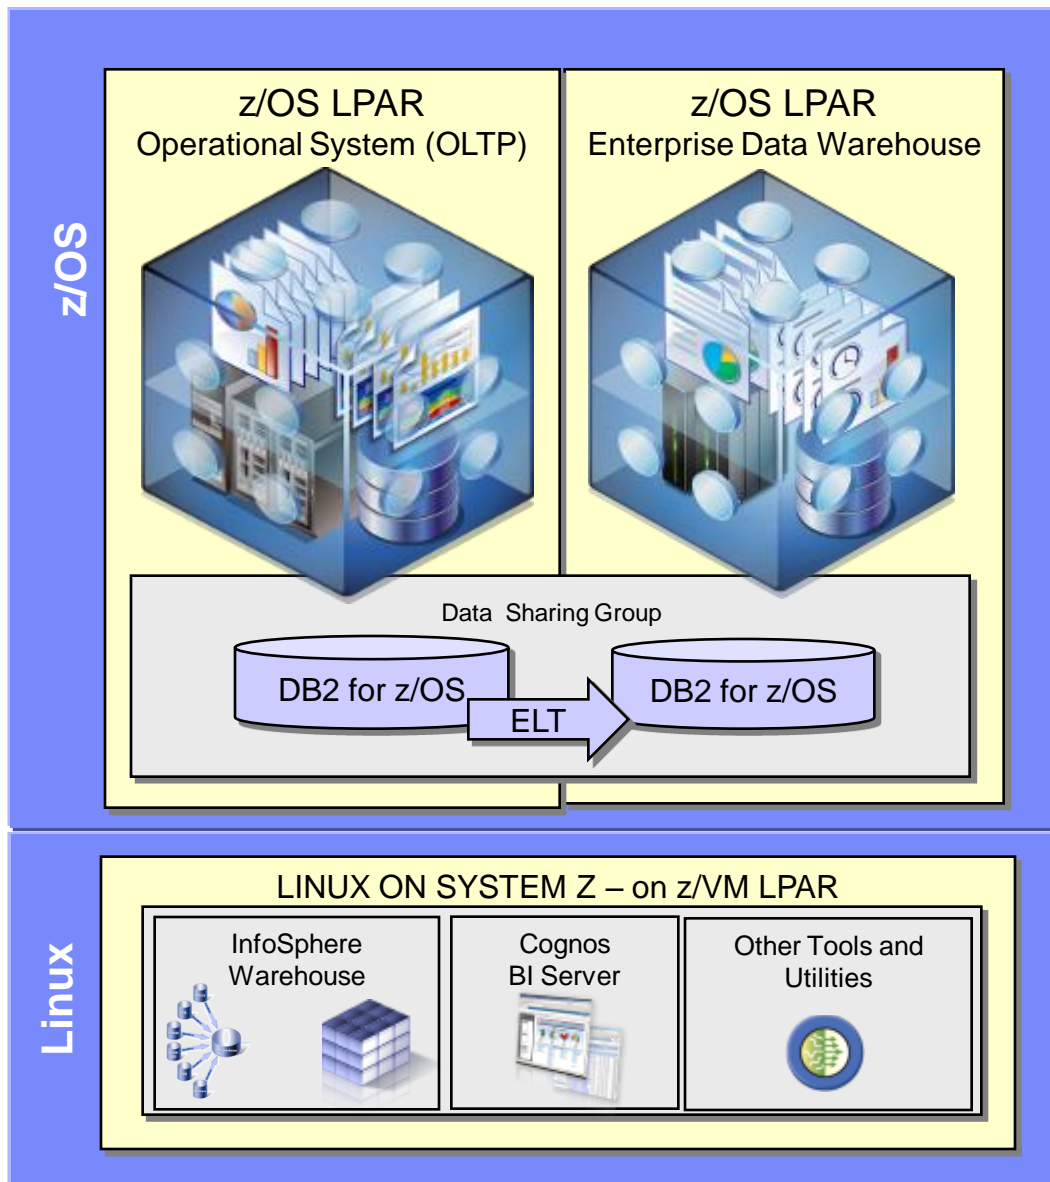


- Business and Technology Drivers

- Data Warehousing Solution on zEnterprise
 - IDAA Key Design and Operational Features
 - Powered by Netezza

- IDAA Use Cases and Next Steps

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

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

... delivering information when, where, and how each user needs it



- **Delivers information where, when and how it is needed**

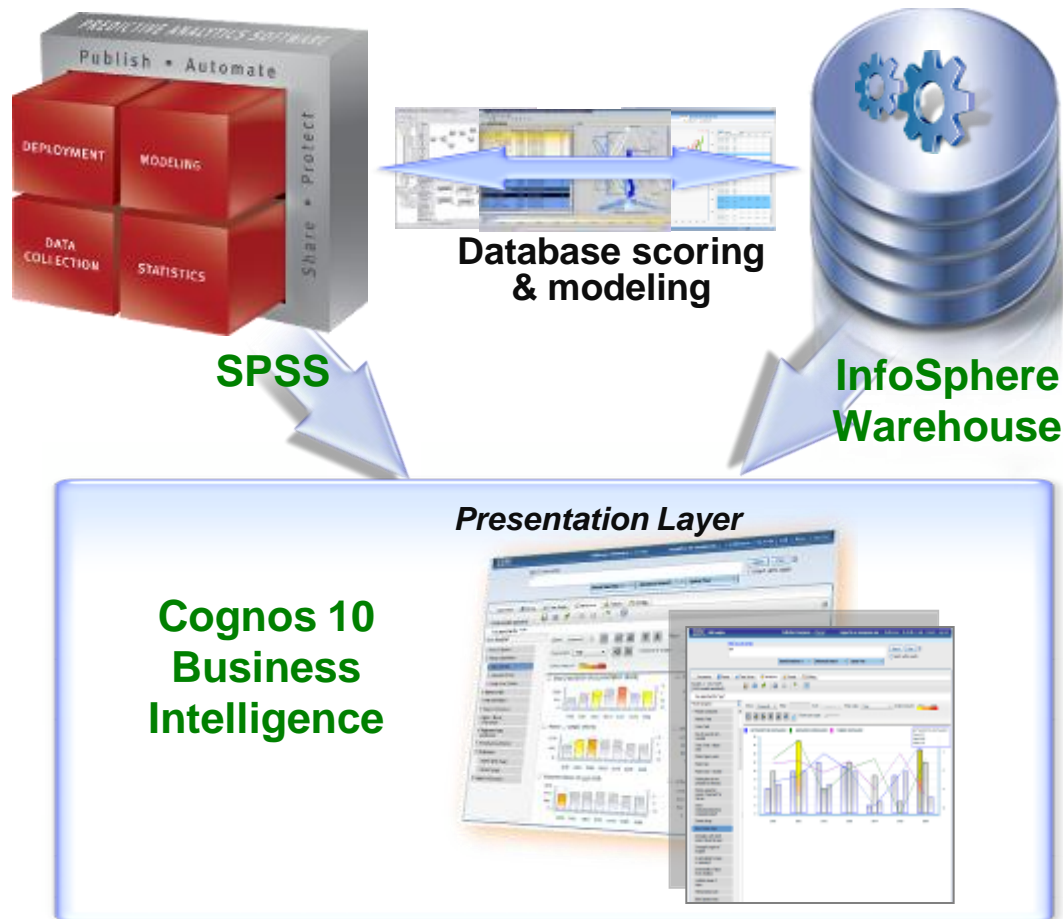
- Self-service reporting and analysis
- Individualized by user
- Automated delivery of information in context
- Author once, consume anywhere

- **Full range of BI capabilities**

- Query, reporting, analysis, dashboarding, realtime monitoring

- **Purpose-built SOA platform**

that fits client environments and scales easily



Full breadth of predictive analytics

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

Putting prediction in hands of the business

Decision Management

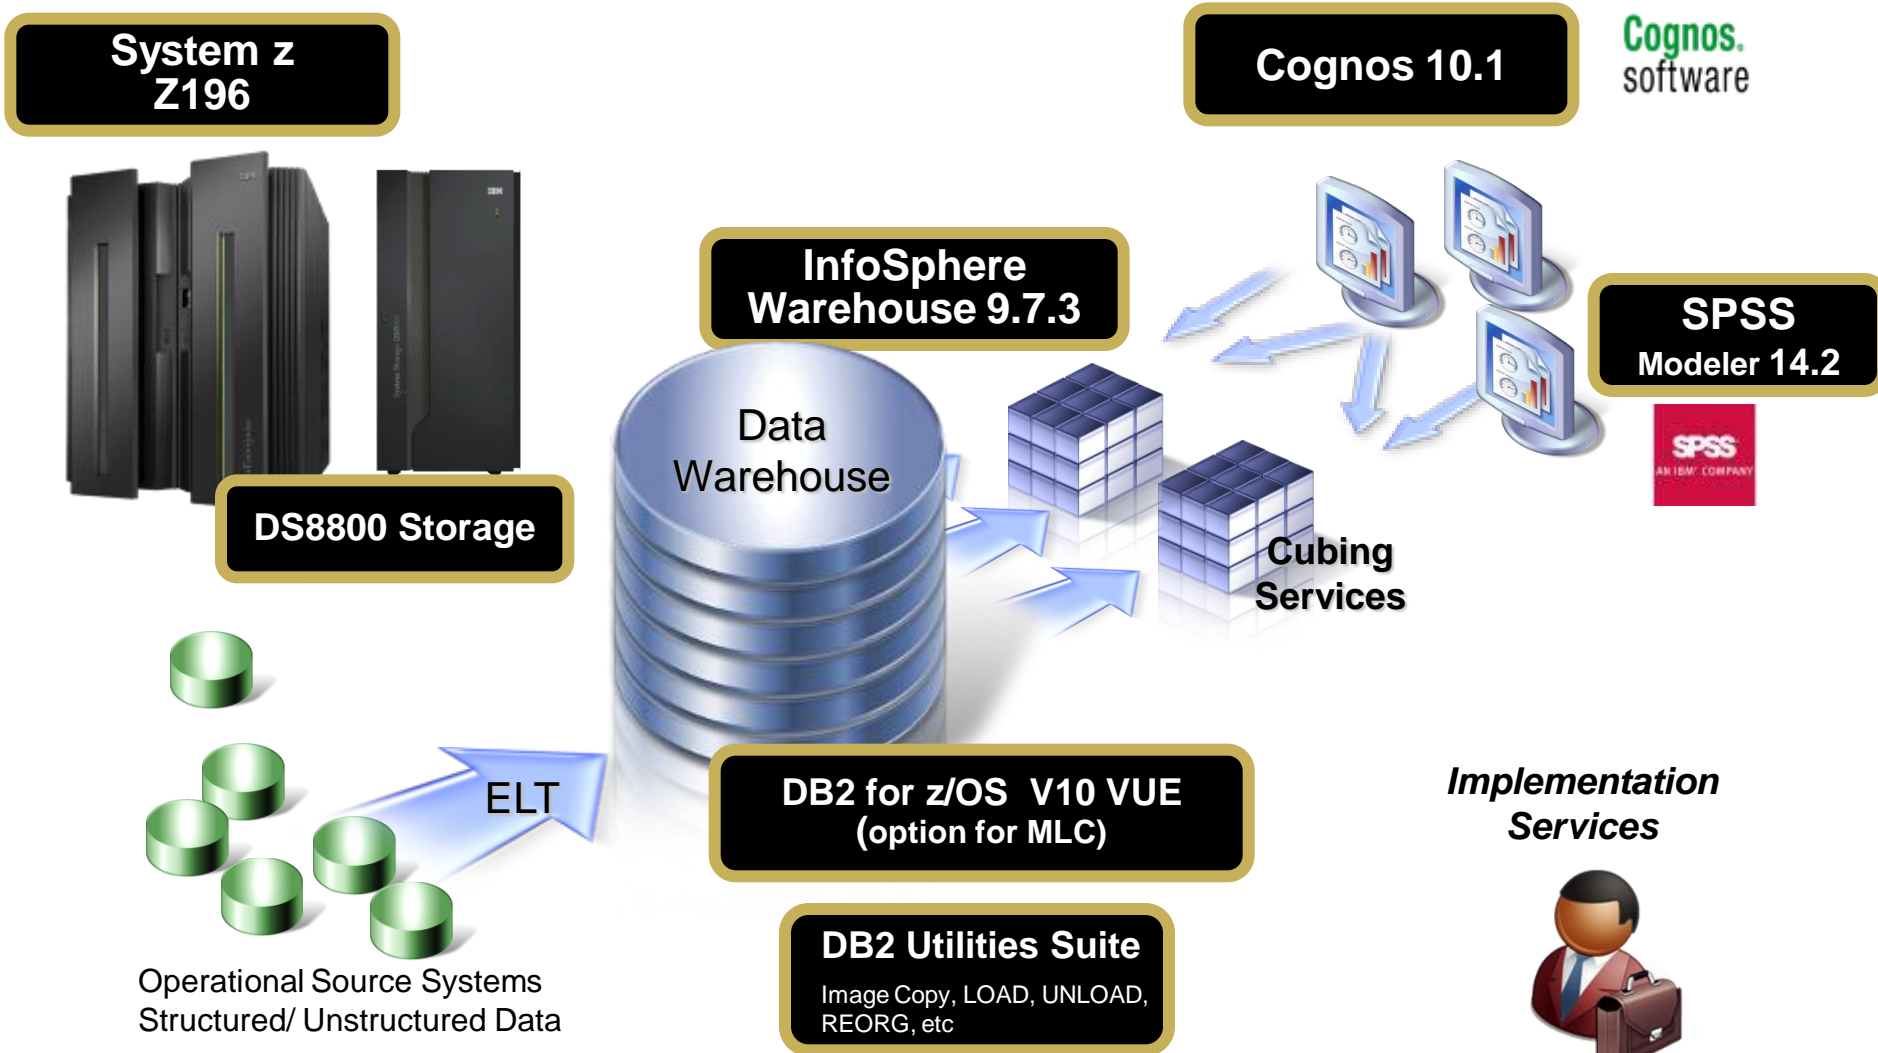
Driving better business outcomes

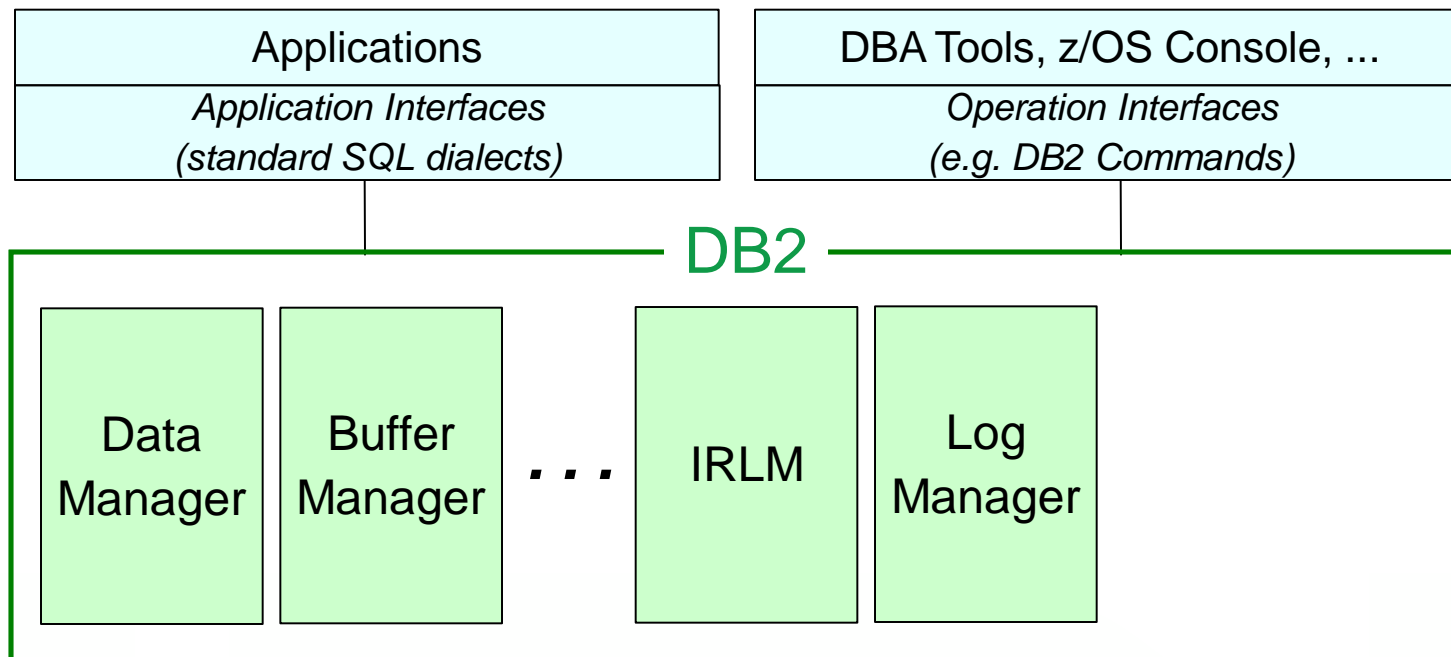
- Attract and retain profitable customers
- Detect and prevent fraud
- Improve resource allocation

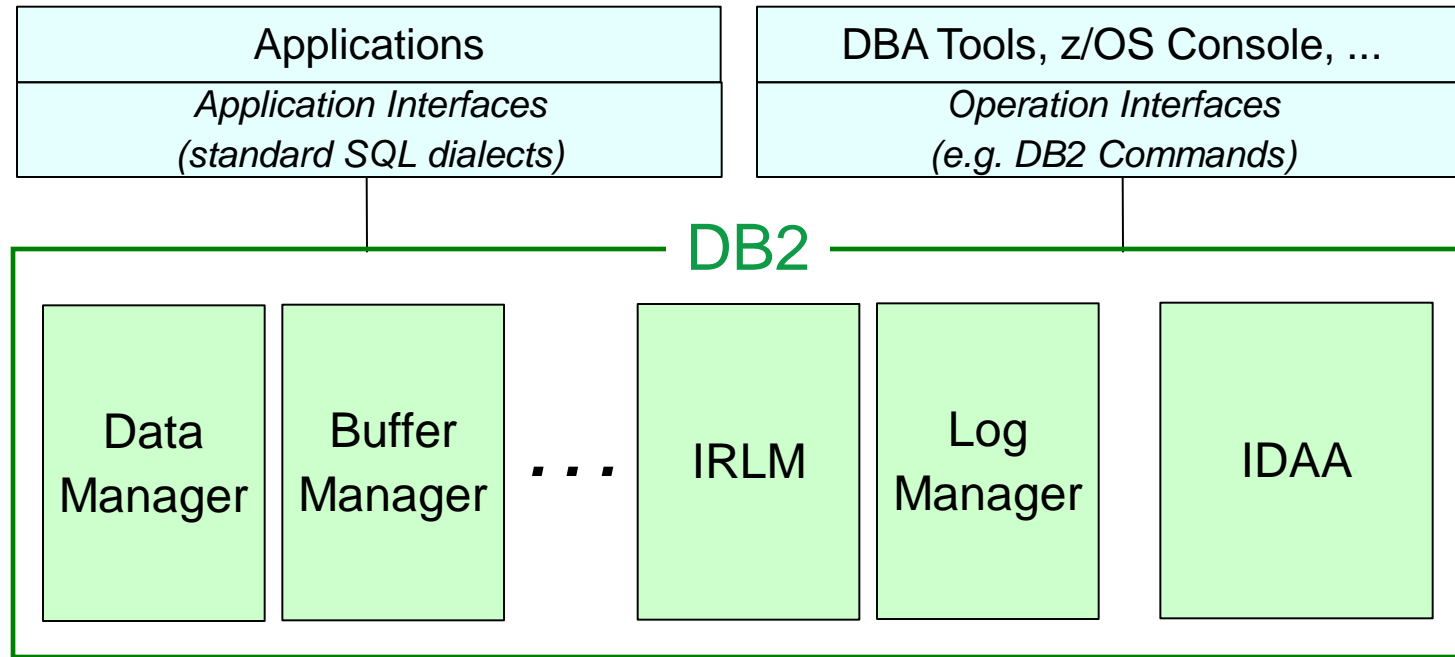
IBM Smart Analytics System 9700



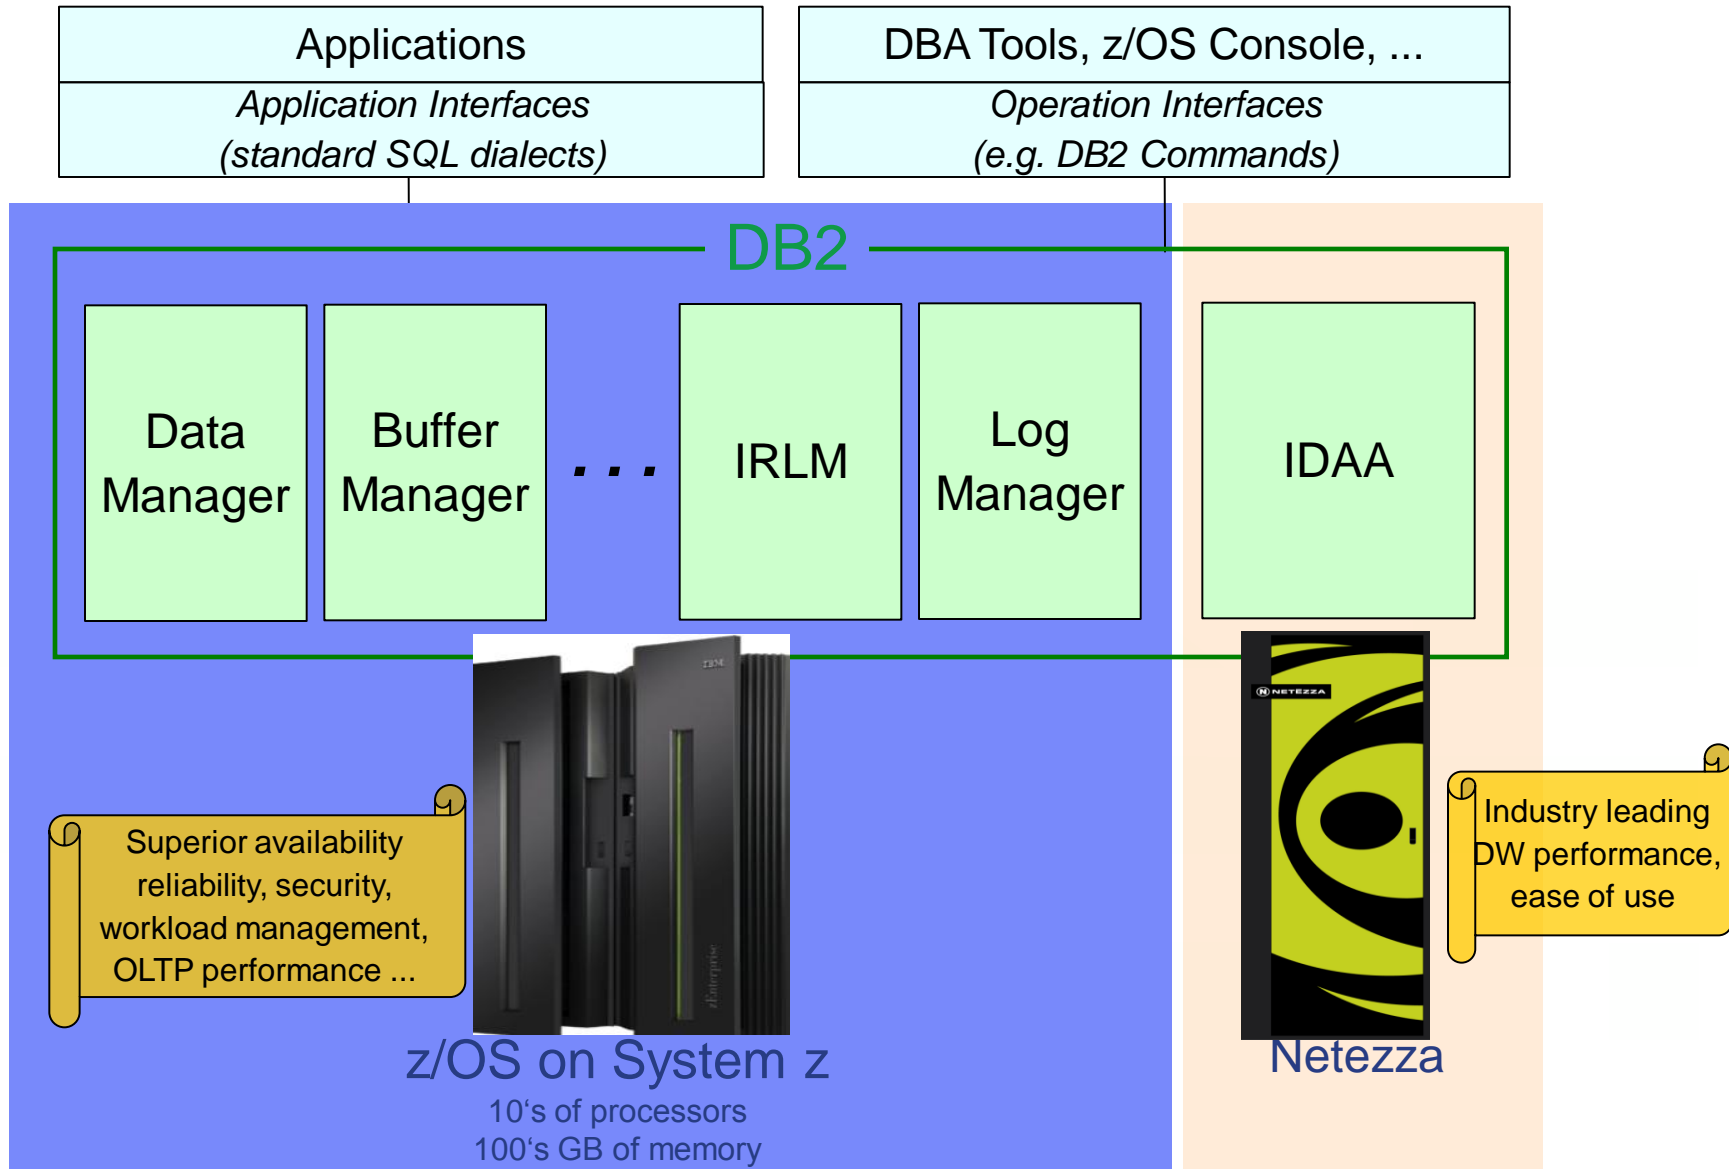
High Value Data Warehousing – Standard Configuration



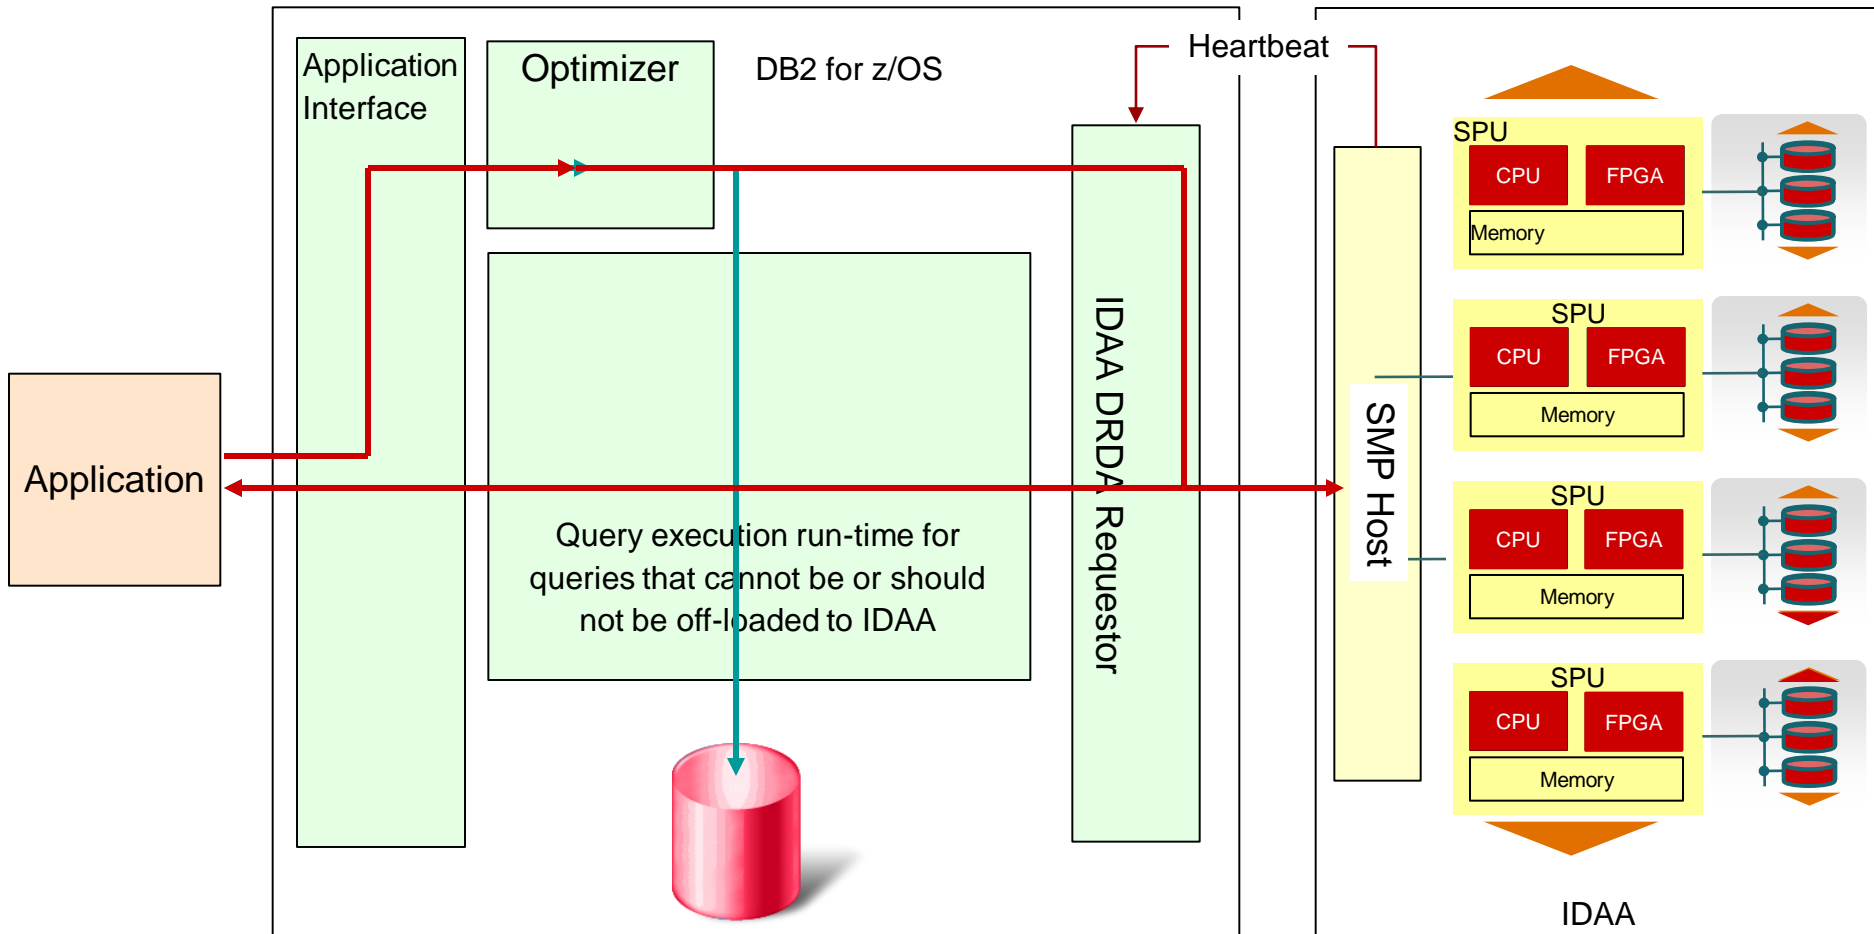




Deep DB2 Integration



Query Execution Process Flow



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

- 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%



			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

Queries run faster

- Save CPU resources
- People time
- Business opportunities

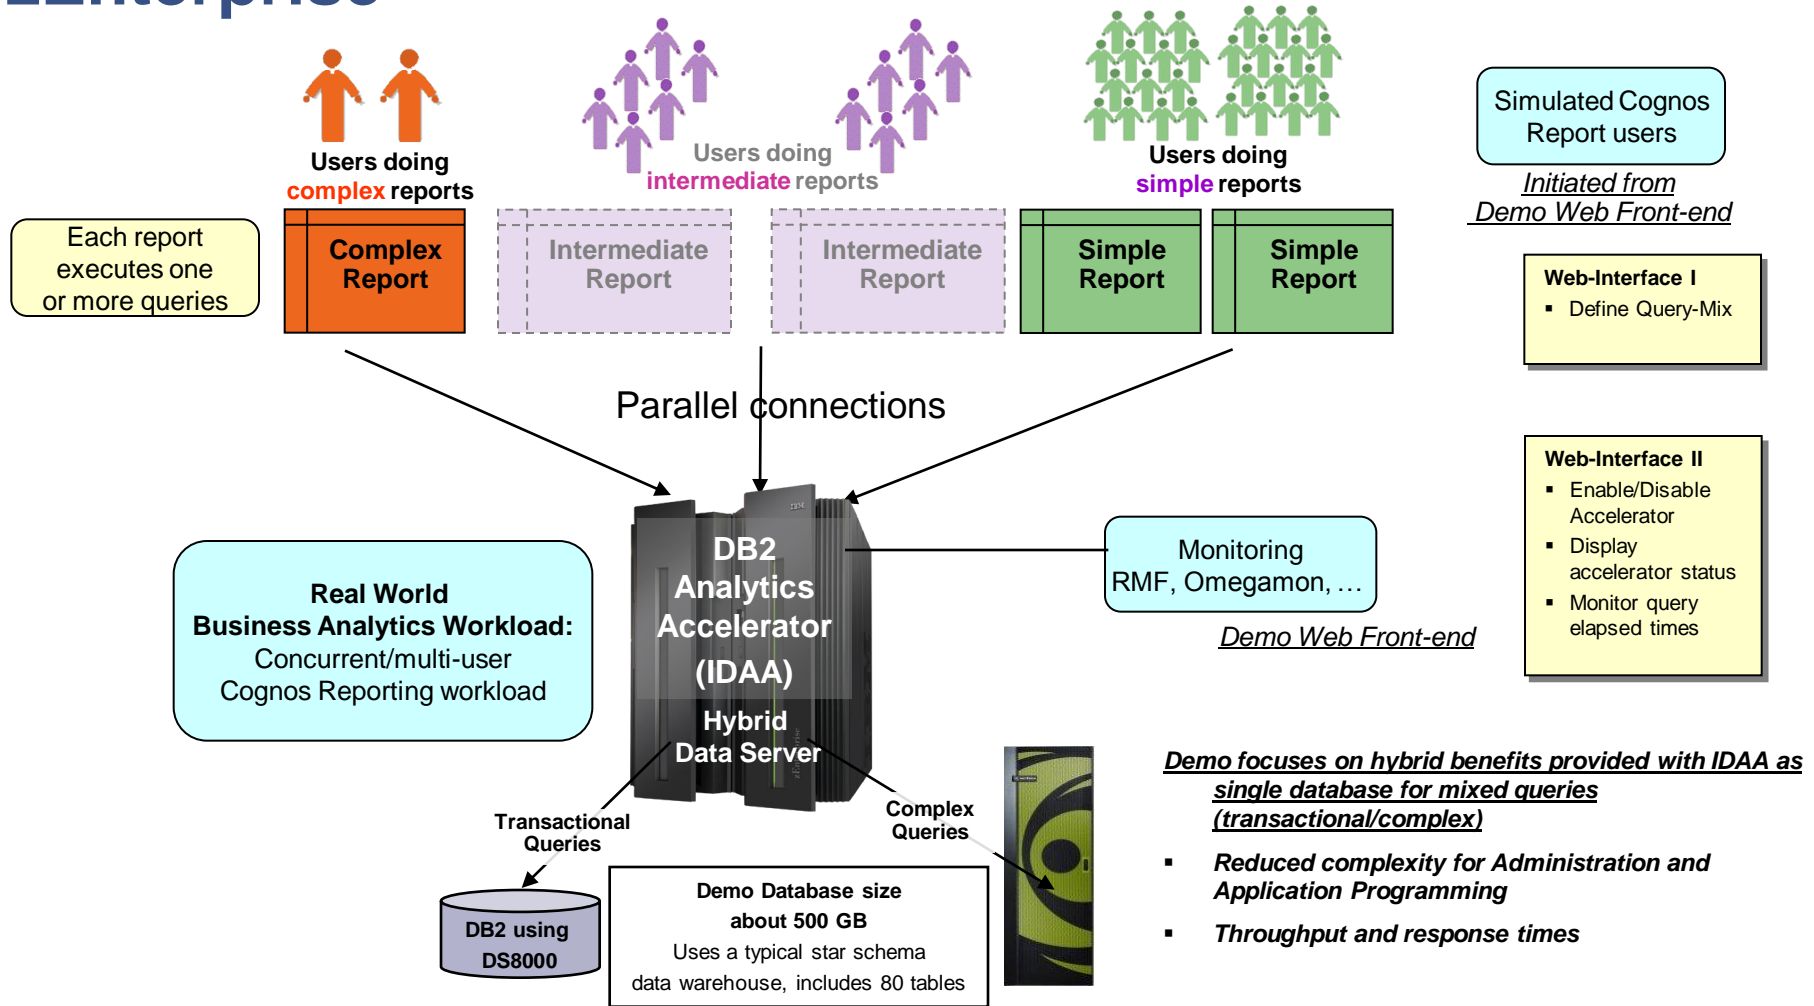
Actual customer results, October 2011

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

Live Demo: IDAA – BI Day with mixed workload on zEnterprise




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

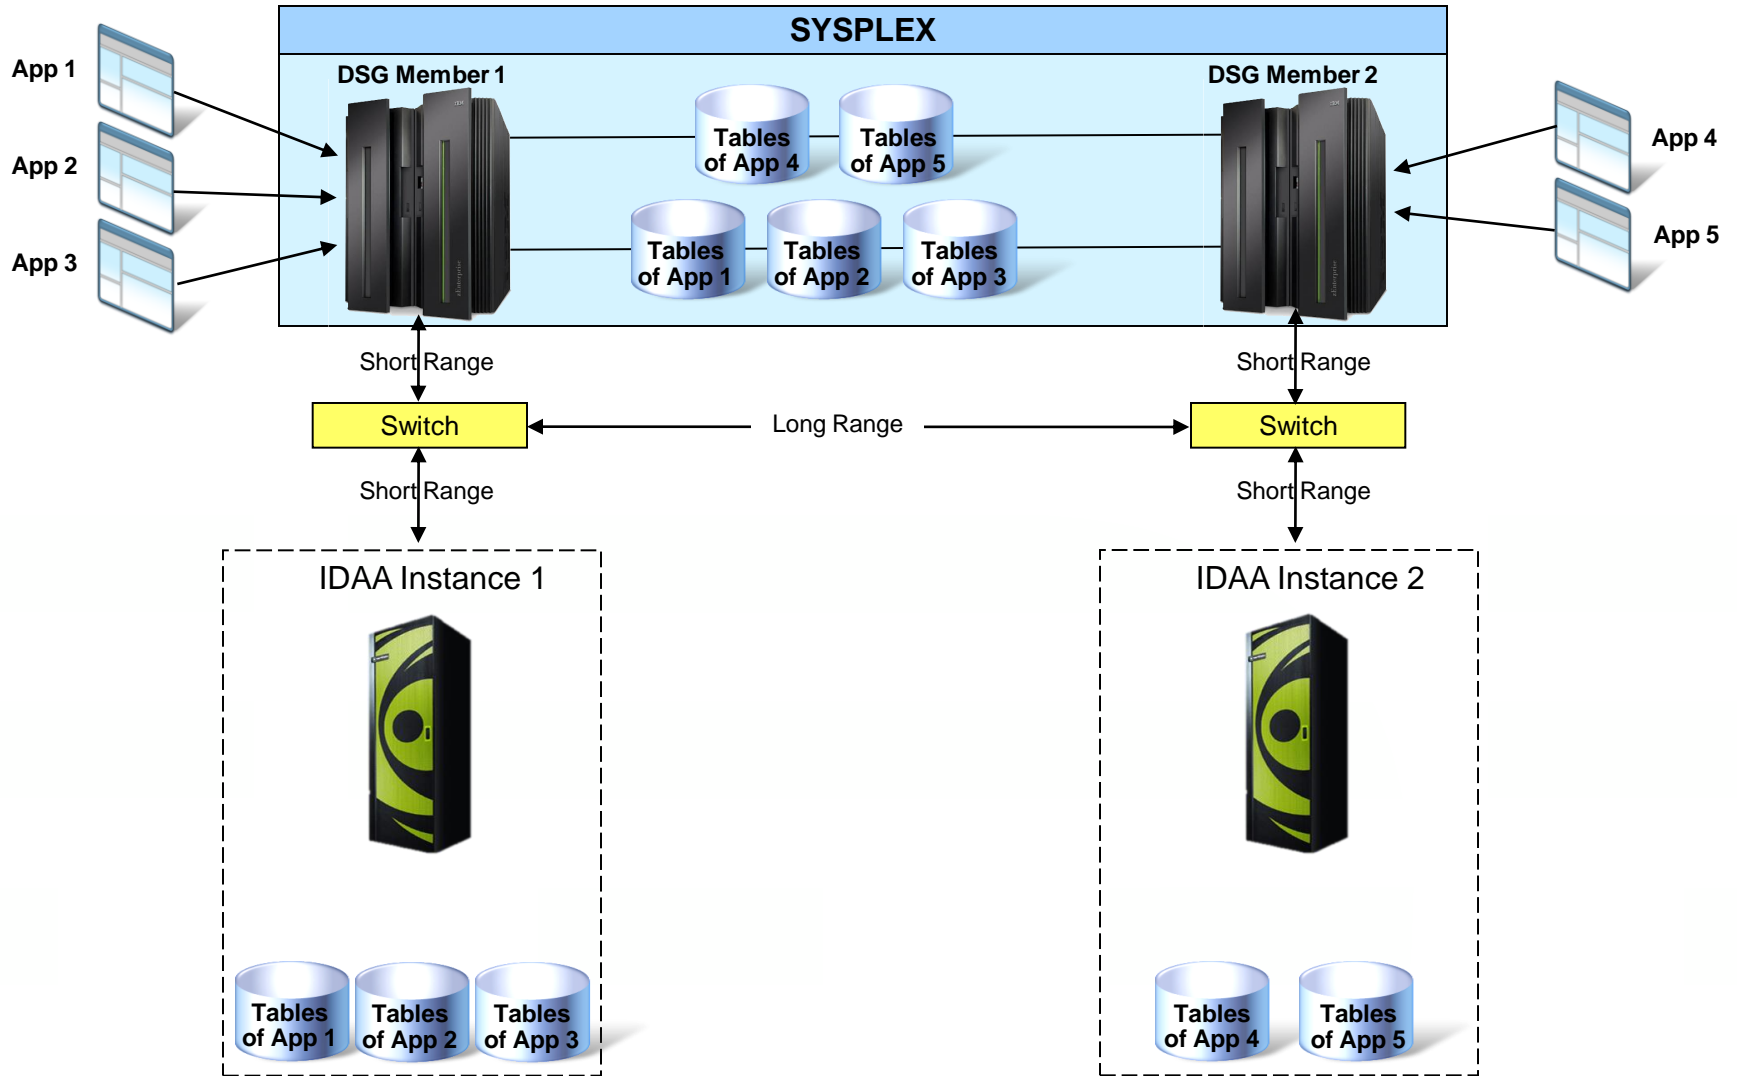


Better utilization of IDAA resources
Scalability
High availability

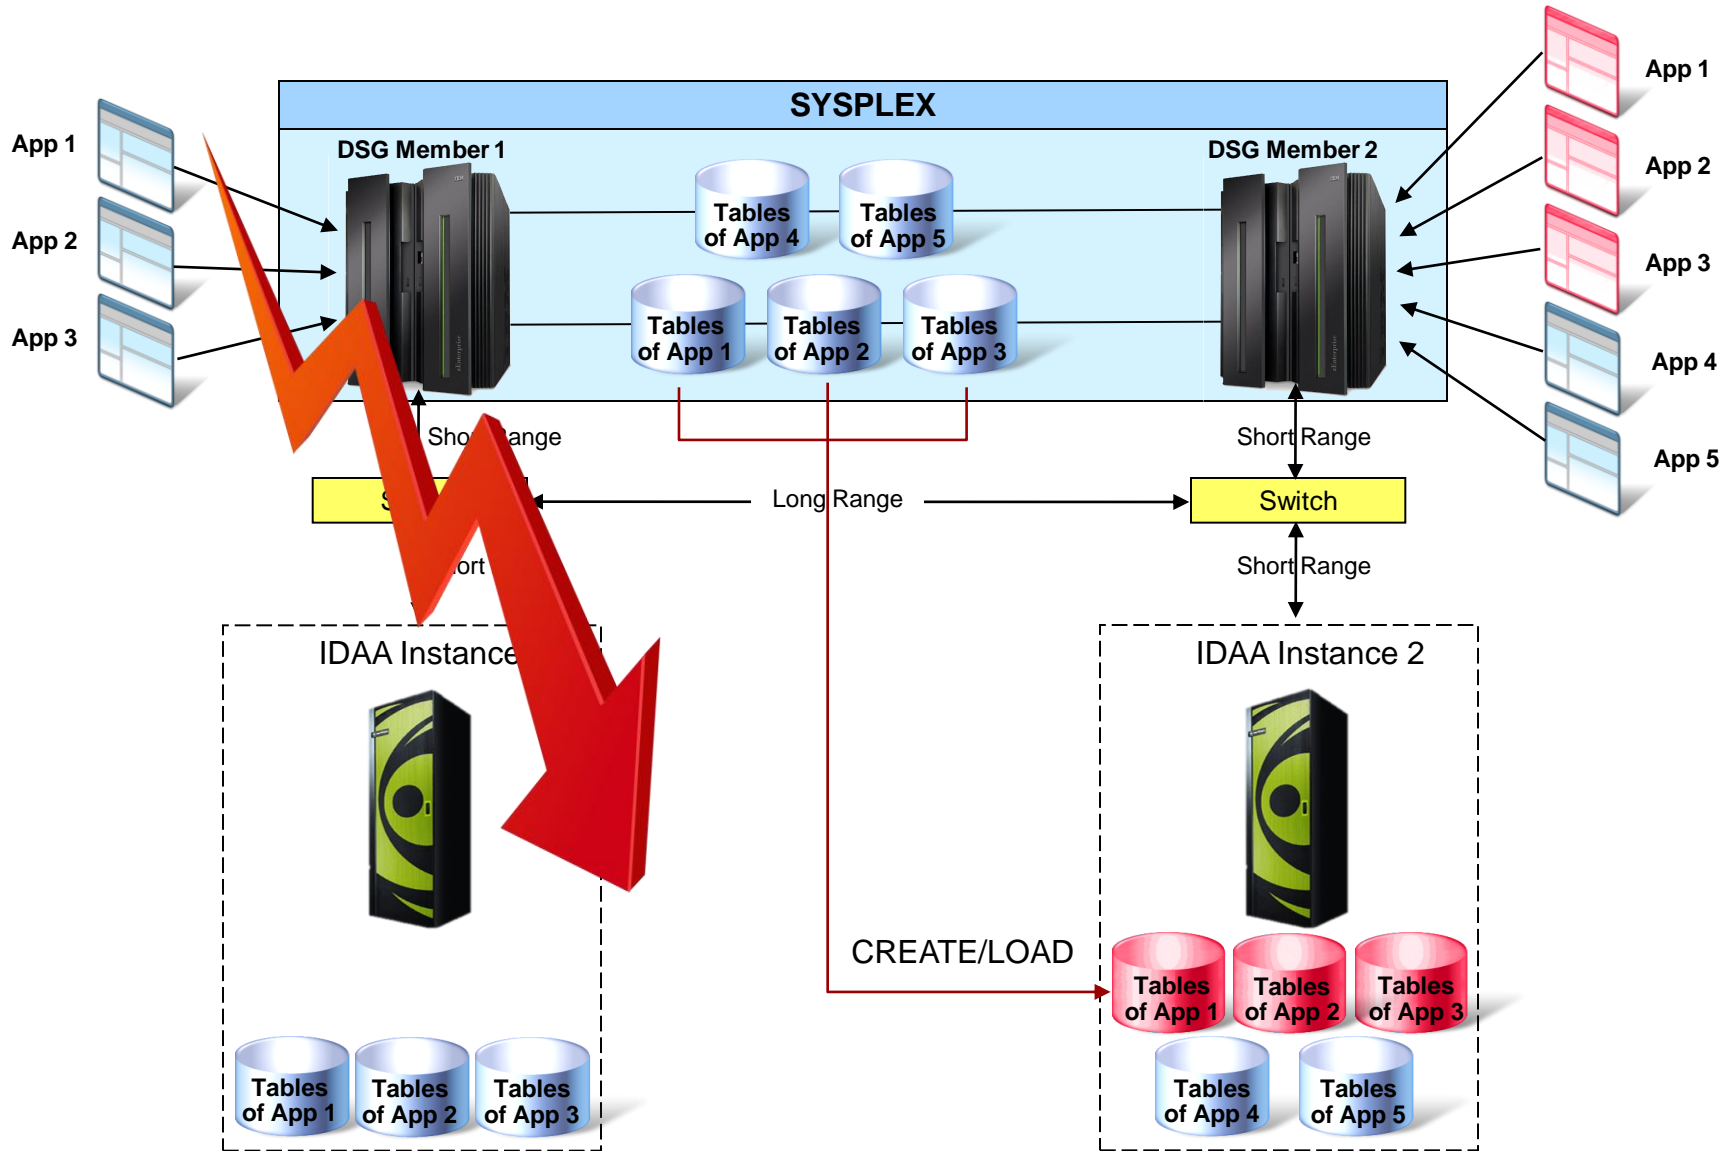
Full flexibility for DB2 systems:

- 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

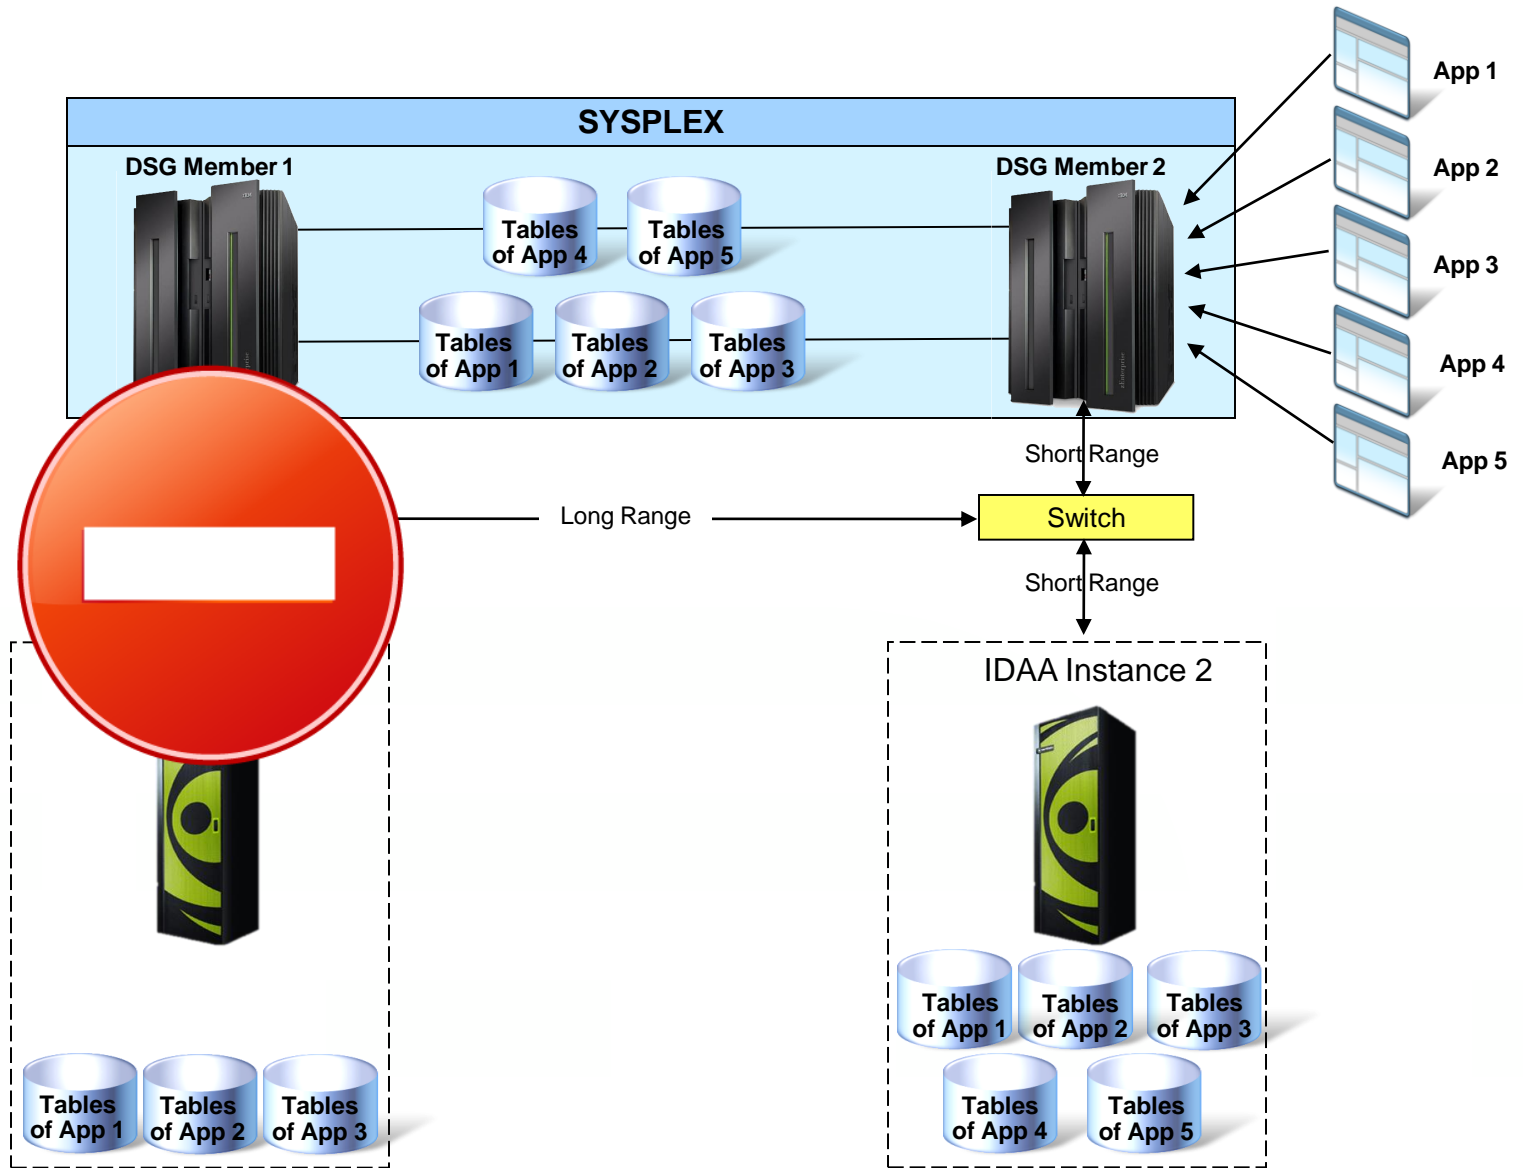
IDAA Disaster Recovery Considerations (1 of 3)



IDAA Disaster Recovery Considerations (2 of 3)



IDAA Disaster Recovery Considerations (3 of 3)



- Business and Technology Drivers

- Data Warehousing Solution on zEnterprise
 - IDAA Key Design and Operational Features
 - Powered by Netezza

- IDAA Use Cases and Next Steps


IDAA V2 Powered by IBM 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. in model 1000-12:
8 enclosures → 96 HDDs
32TB uncompressed user data (→ 128TB)

IDAA Server
SQL Compiler, Query Plan, Optimize
Administration
2 front/end hosts, IBM 3650M3 or 3850X5
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. in model 1000-12: 12 back/end SPUs
(more details on following charts)

The Netezza S-Blade™

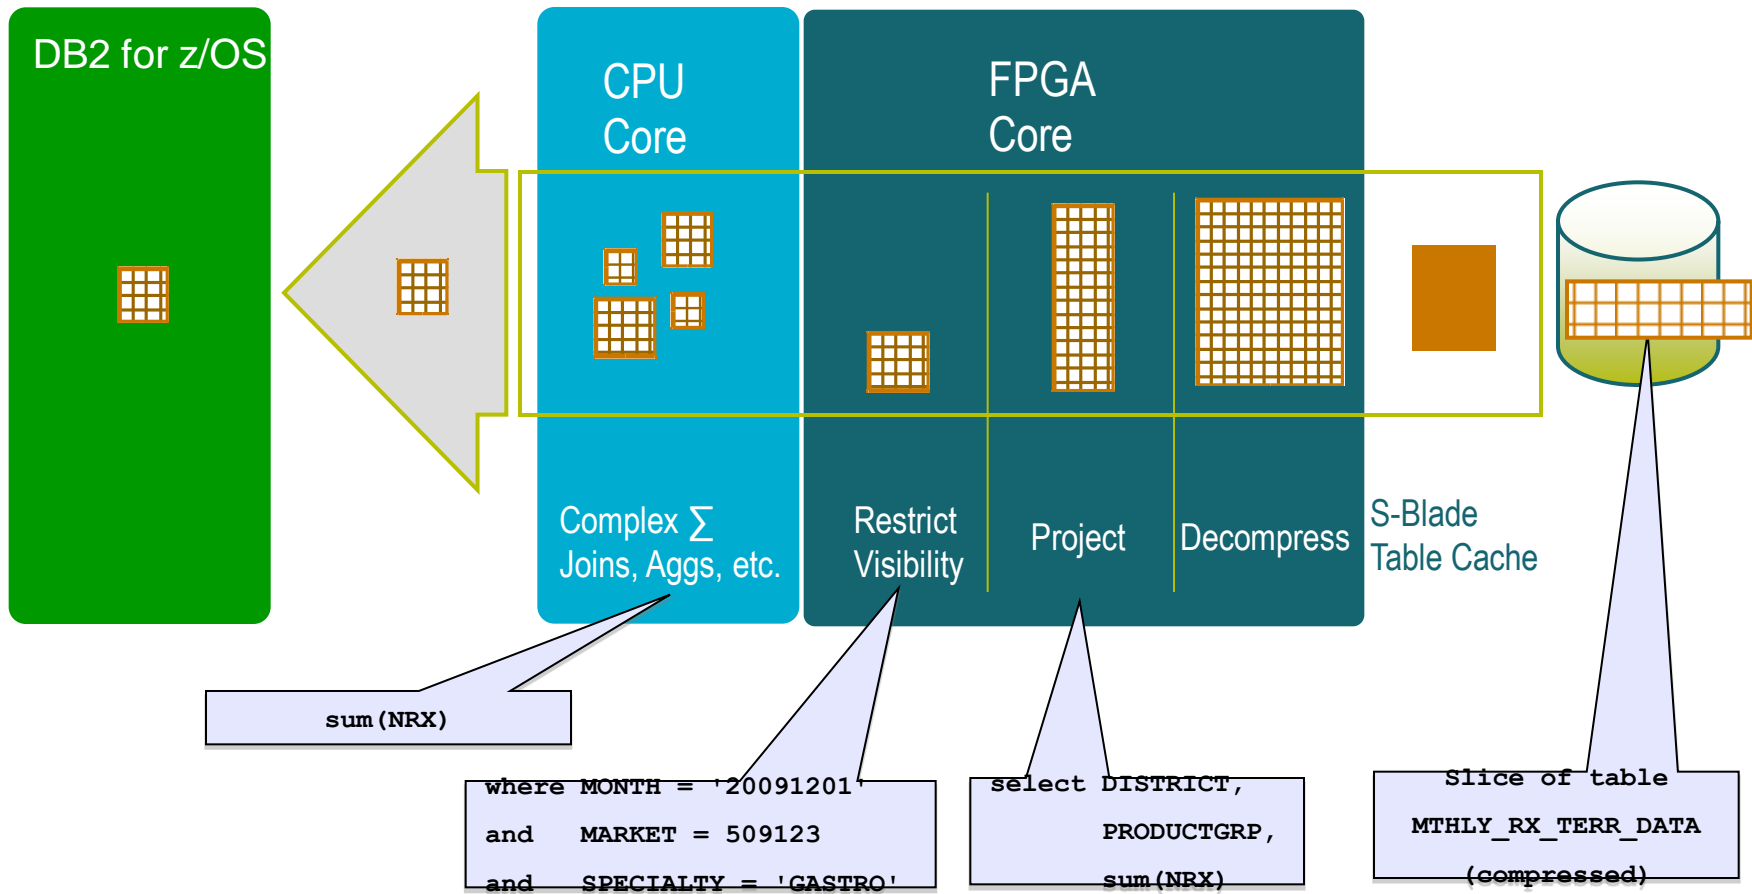


Applying Data Stream Processing to DB2 Queries




... the key to the speed

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




IBM Netezza 1000 Appliance Models




Model	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

Capacity = User data space
 Effective Capacity = User data space with compression (4x compression assumed)

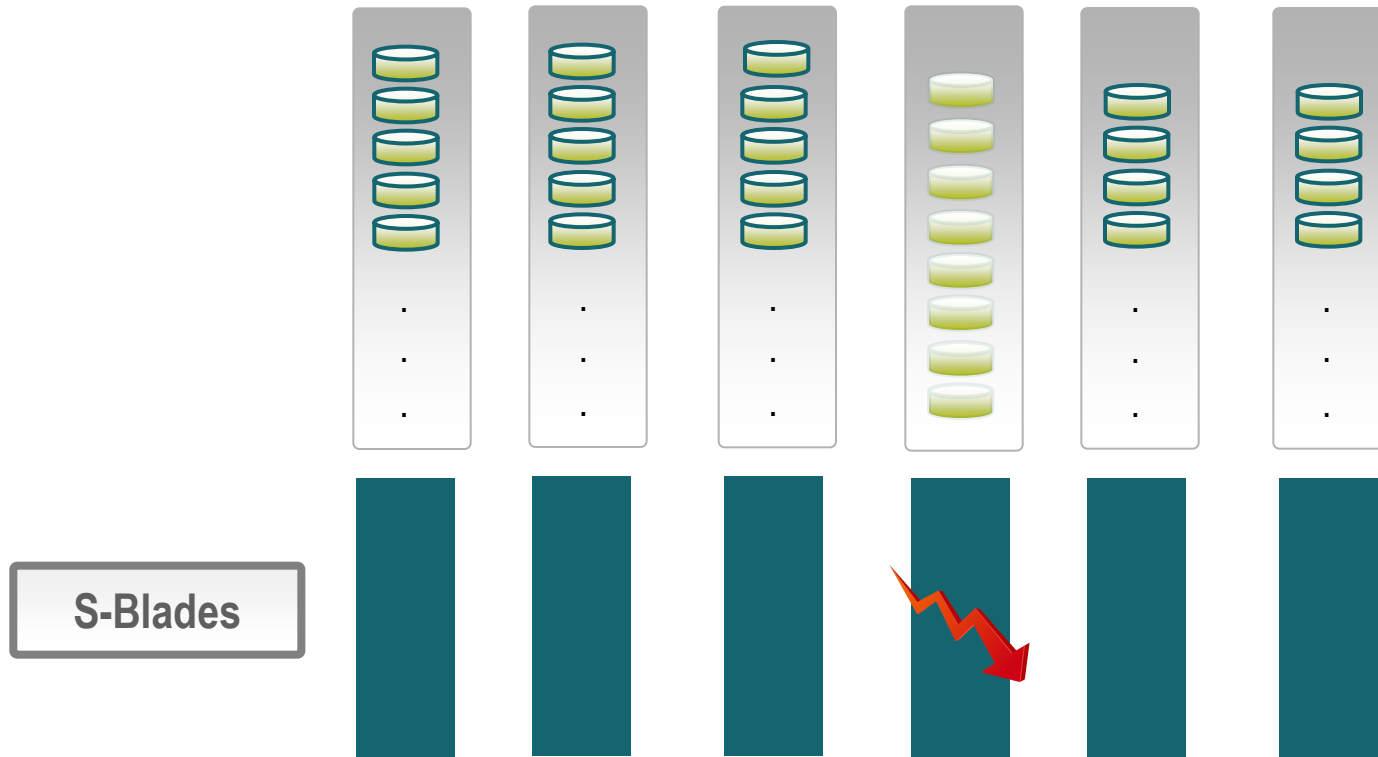


- All user data and temp space mirrored
- Disk failures transparent to queries and transactions
- Failed drives automatically regenerated
- Bad sectors automatically rewritten or relocated

Shielding Against S-Blade™ Failures



- S-Blade failure is automatically detected



- 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

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



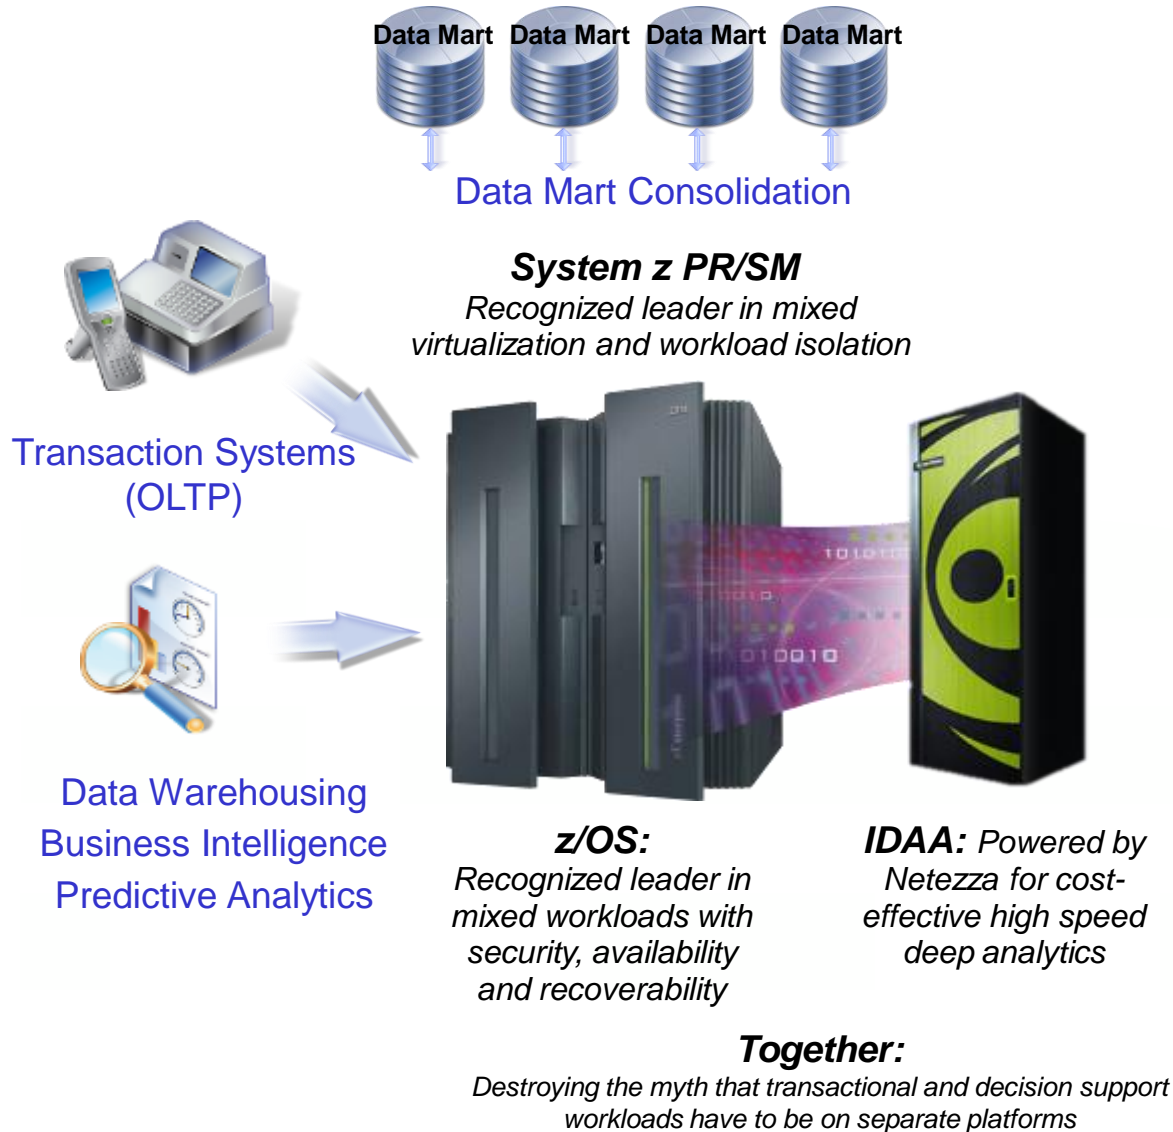
Get more insight from your data

- 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

- Business and Technology Drivers

- Data Warehousing Solution on zEnterprise
 - IDAA Key Design and Operational Features
 - Powered by Netezza

- IDAA Use Cases and Next Steps




Bringing it all together

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

zEnterprise with IDAA Use Cases – Potential Entry Points



- On average, **70% of the data** that feeds data warehousing and business analytics solutions **originates on the System z platform**



How is that data being analyzed today?

IDAA use case

Analyzed on distributed EDW

Consolidation of EDW

Analyzed on a sprawl of Data Marts

Consolidation of Data Marts

Analyzed on distributed EDW

Offload data and analytics from the EDW to optimized DB2 z with IDAA

Not being analyzed yet (new workload)

More Insights

- new BI applications
- new operational BI

With long running queries (>1min) in DWH or Operational Data Stores on DB2 for z/OS.

Modernization

- Partner with IBM to assess best IDAA entry point for your needs:
 - jointly assess your Data Warehouse & BI architecture and current challenges
 - define IDAA architecture
 - identify distributed workload for IDAA
 - define use case for new project
 - conduct Workload Assessment of long running queries on DB2 for z/OS
 - define PoC and next steps

- Contact the Data Warehousing on System z Center of Excellence at dwhz@de.ibm.com

Thank You

The words 'Thank You' are rendered in a large, bold, 3D sans-serif font. The letters are a metallic grey color and are set against a white background. Below the text, there is a clear, slightly blurred reflection of the letters on a reflective surface, creating a sense of depth and three-dimensionality.