Introducing the IBM DB2 Analytics Accelerator

Session IDW-1040

Alan Meyer
Sr Marketing Manager, Data Warehousing
almeyer@us.ibm.com

Namik Hrle
IBM Distinguished Engineer
hrle@de.ibm.com

IBM Software
Information On Demand 2011

Disclaimer

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Agenda

- Overview of the Announcement
- Business and Technology Drivers
- Key Design and Operational Features
- Powered by Netezza
- Supported Workloads and Workload Applicability Assessment

DB2 Analytics Accelerator for DB2 z/OS V2.1

Capitalizing on the best of both worlds – System z and Netezza

What is it?

The IBM DB2 Analytics Accelerator is a workload optimized, appliance add-on that enables the integration of business insights into operational processes to drive winning strategies. It accelerates queries, with unprecedented response times.

How is it different?

- **Performance**: Unprecedented response times to enable 'train of thought' analyses frequently blocked by poor query performance.
- **Integration:** Connects directly to DB2 through deep integration transparently to applications.
- **Self-managed workloads:** queries are executed in the most efficient way
- **Transparency:** applications connected to DB2 typically with no modification
- **Simplified administration:** appliance hands-free operations, eliminating many database tuning tasks

Marrying the best of System z and Netezza technologies

How is it different from the Smart Analytics Optimizer?

The IBM DB2 Analytics Accelerator V2.1 is technology change from V1.1, moving from a memory based deployment to a disk based, hardware accelerated solution. This change should continue to be transparent to the application.

Data volumes expanded to meet any need

- V1 - 4TB max
- V2 -128TB+

Improved query execution

- Context switching
- Full query execution

Lower cost

- Lower cost in data storage
- More workload offloaded from System z

Expanded query types

- Many V1 restrictions removed

Announced October 12th
will GA in November 24

IDAA V2.1 Preserves V1 Key Value Propositions

- DB2 continues to own data (both OLTP and DW)
  - Access to data (authorization, privileges, …)
  - Data consistency and integrity (backup, recovery, …)
  - **Enables extending System z QoS characteristics to BI/DW data as well**
- Applications access data (both OLTP and DW) only through DB2
  - DB2 controls whether to execute query in DB2 mainline or route to IDAA
  - DB2 returns results directly to the calling application
  - **Enables mixed workloads and selection of optimal access path (within DB2 mainline or IDAA) depending on access pattern**
- IDAA continues to be implemented as DB2 internal component
  - DB2 provides key IDAA status and performance indicators as well as typical administration tasks by standard DB2 interfaces and means
  - No direct access (log-on) to IDAA
  - **Enables operational cost reduction through skills, tools and processes consolidation**
- Investment protection
  - **The full value of ISAO V1 will be applied to IDAA V2.1**

Why Both?

Marrying the best of each

IBM Netezza

Focused Appliance

IBM System z

Mixed Workload System

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 focused workload ⟷ **Very diverse workload**

Tailored to your needs

A Hybrid Solution

IBM Netezza	IBM System z with DB2 Analytics Accelerator
Focused Appliance	**Mixed Workload System**
• Appliance with a streamlined database and HW acceleration for key warehouse functionality • Price/performance leader • Speed and ease of deployment and administration • Optimized performance for a specific workload range	• 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

Combining the best transaction system with the best analytics system

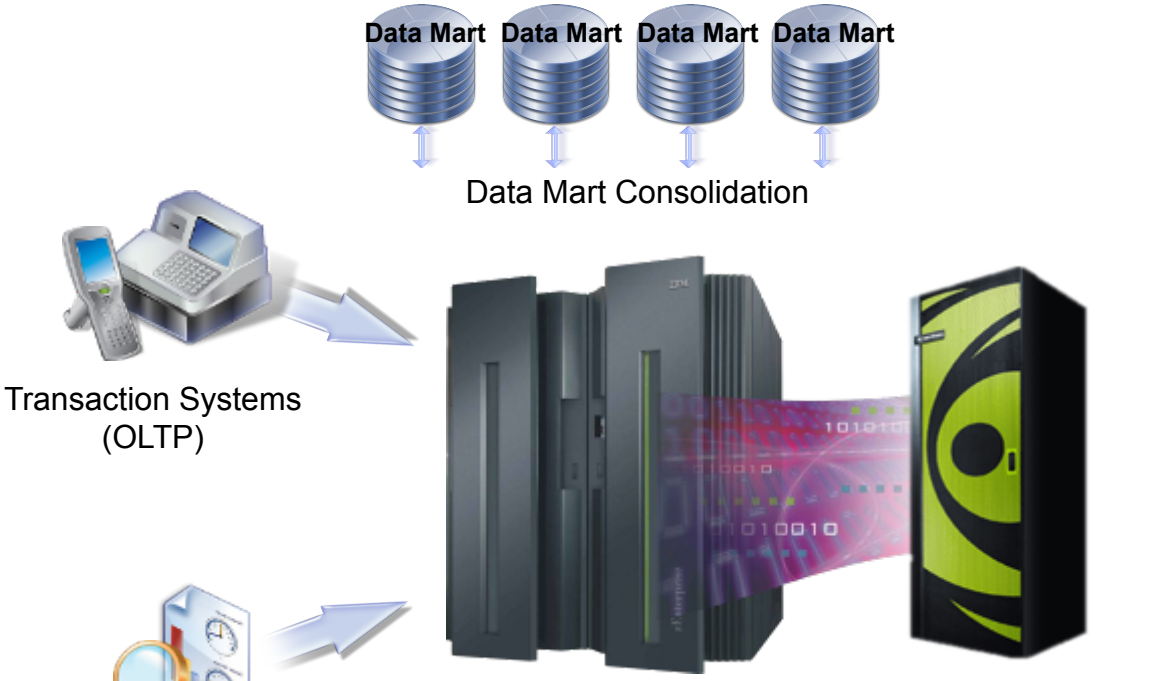

z/OS:
Recognized leader in mixed workloads with security, availability and recoverability for OLTP

Netezza:
Recognized leader in cost-effective high speed deep analytics

Best in OLTP

Industry recognized leader in mission critical transaction systems

Best in Data Warehouse

Proven appliance leader in high speed analytic systems

Best in Consolidation

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

Agenda

- Overview of the Announcement
- Business and Technology Drivers
- Key Design and Operational Features
- Powered by Netezza
- Supported Workloads and Workload Applicability Assessment

Business Challenges and Technology Trends

- **Changing business requirements**
  - BI/DW becoming mission critical and requires OLTP-like QoS
    - reliability, continuous availability, security, mixed workload management, …
    - orders of magnitude faster execution of complex, ad hoc queries
    - predictable query performance
  - Shift towards dynamic DW and operational BI
    - Combining OLTP and OLAP workloads
- **Traditional performance tuning tools of the trade such as indexing, pre-built aggregates and MQTs struggling to keep the pace**
  - Require top DBA expertise and sophisticated tools
  - Even then not good enough due to ad-hoc, unpredictable nature of the workload
- **Technology trends**
  - Very large number of processor sockets and cores
  - Specialized query processing engines
  - Appliances

Ultimate Consolidation Opportunity

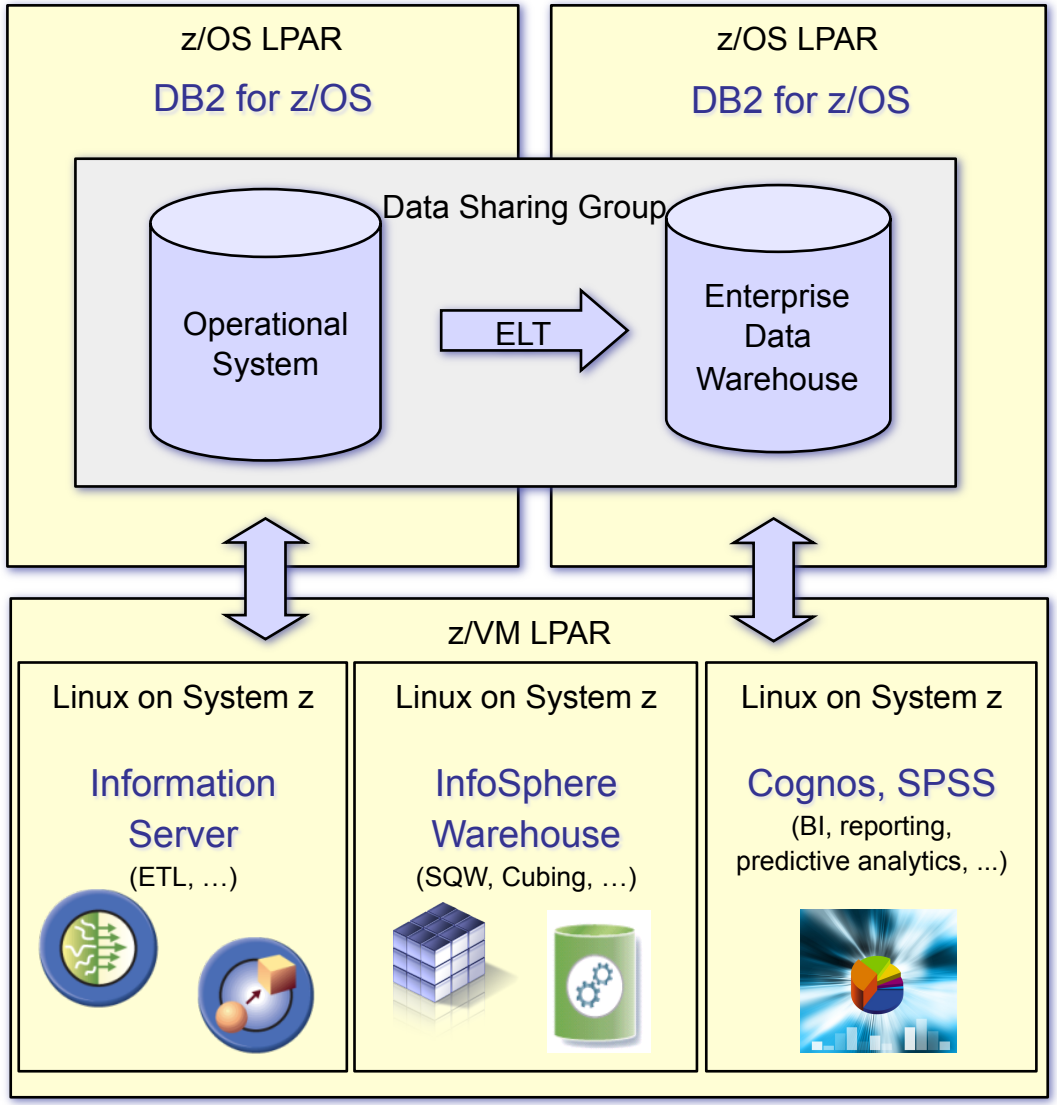

- Consolidation of mission-critical data on System z
- Leveraging existing environment, high availability, backup and governance procedures as well as skills
- Efficient data movement within a data sharing group (no network)
- Performance and TCO improvements through cubing services (data marts) and DB2 enhancements
- Complex transformations and data quality are driven with Information Server
- Cognos and SPSS as the industry leading business intelligence and predictive analytics tools

IBM DB2 Analytics Accelerator

Adding Industry Leading Performance

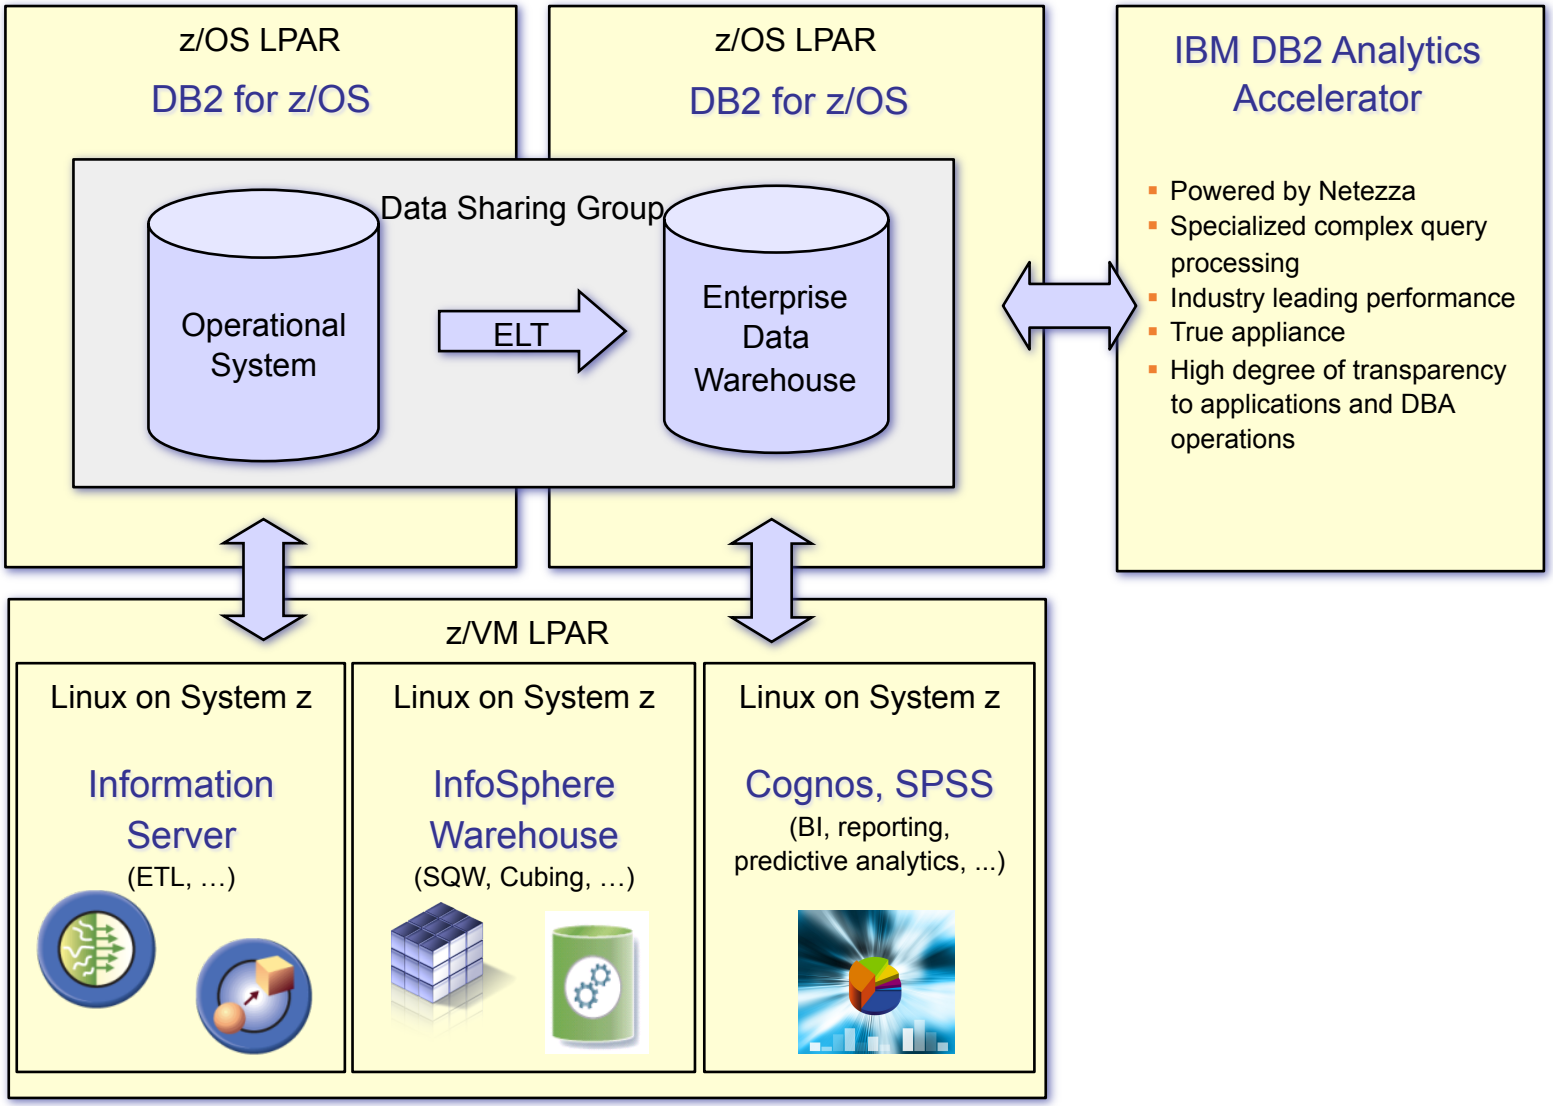

Agenda

- Overview of the Announcement
- Business and Technology Drivers
- Key Design and Operational Features
- Powered by Netezza
- Supported Workloads and Workload Applicability Assessment

DB2 Components

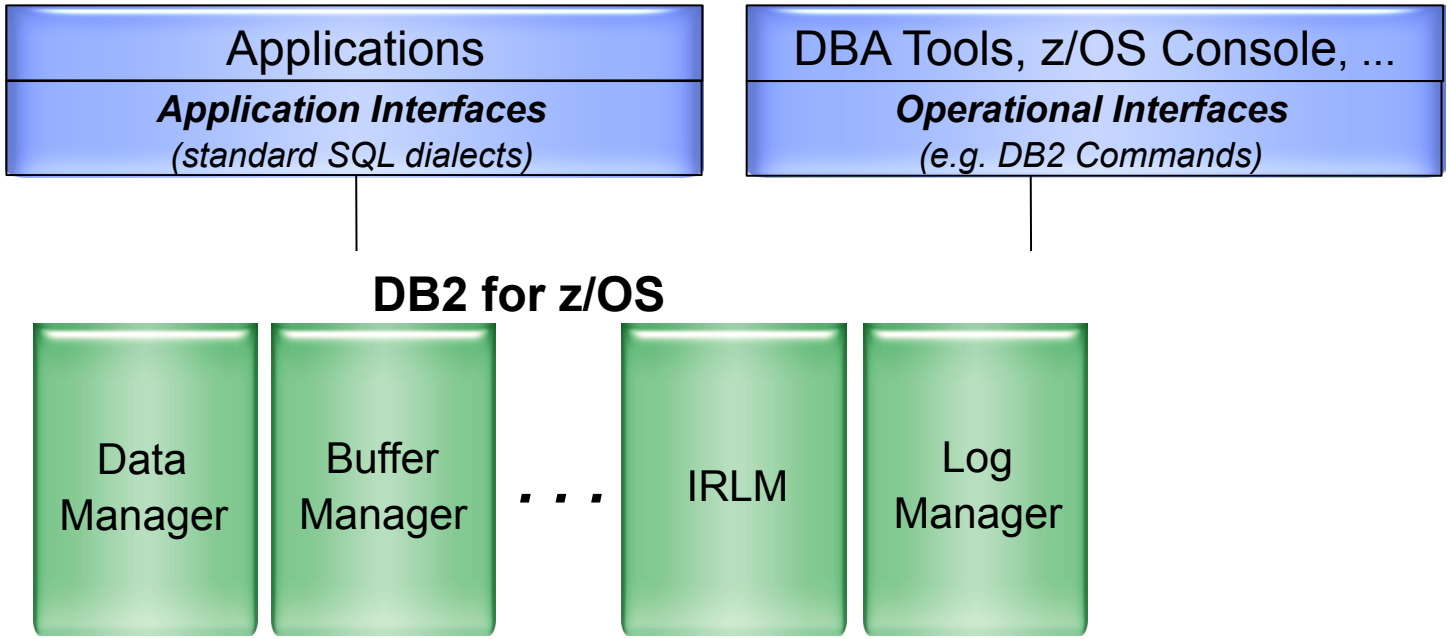

IBM DB2 Analytics Accelerator as a Virtual DB2 Component

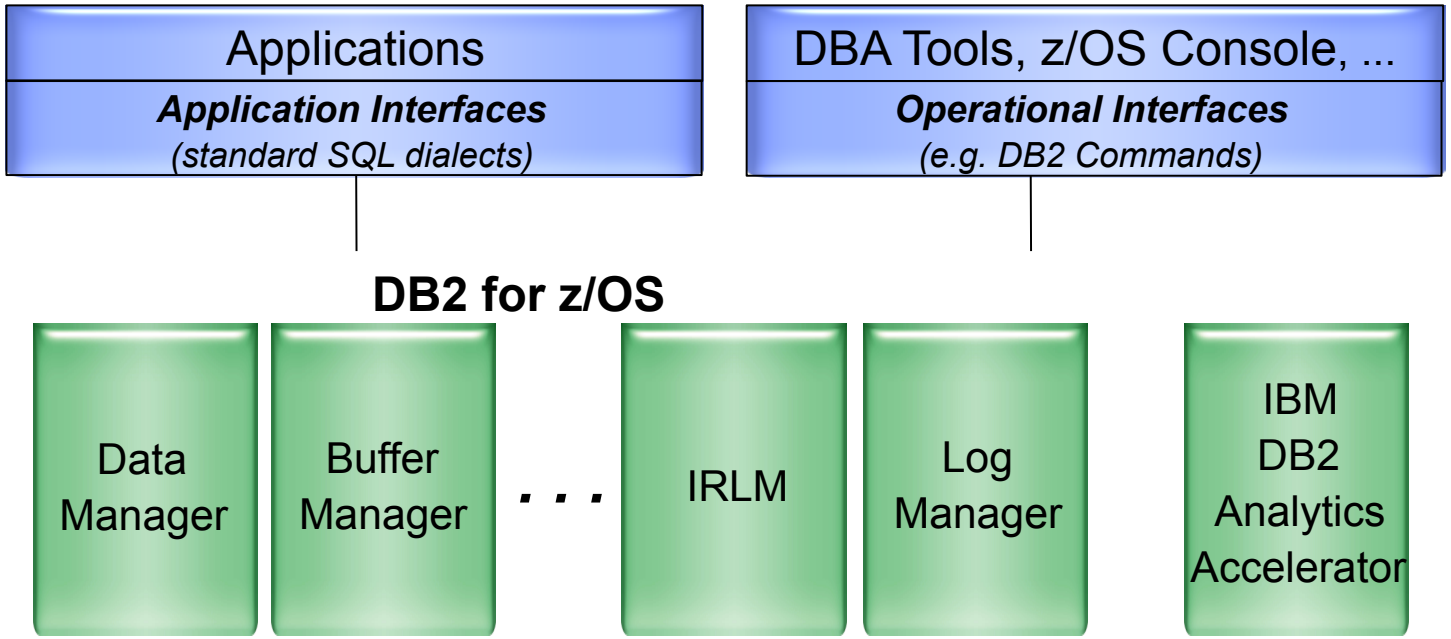

Deep DB2 Integration

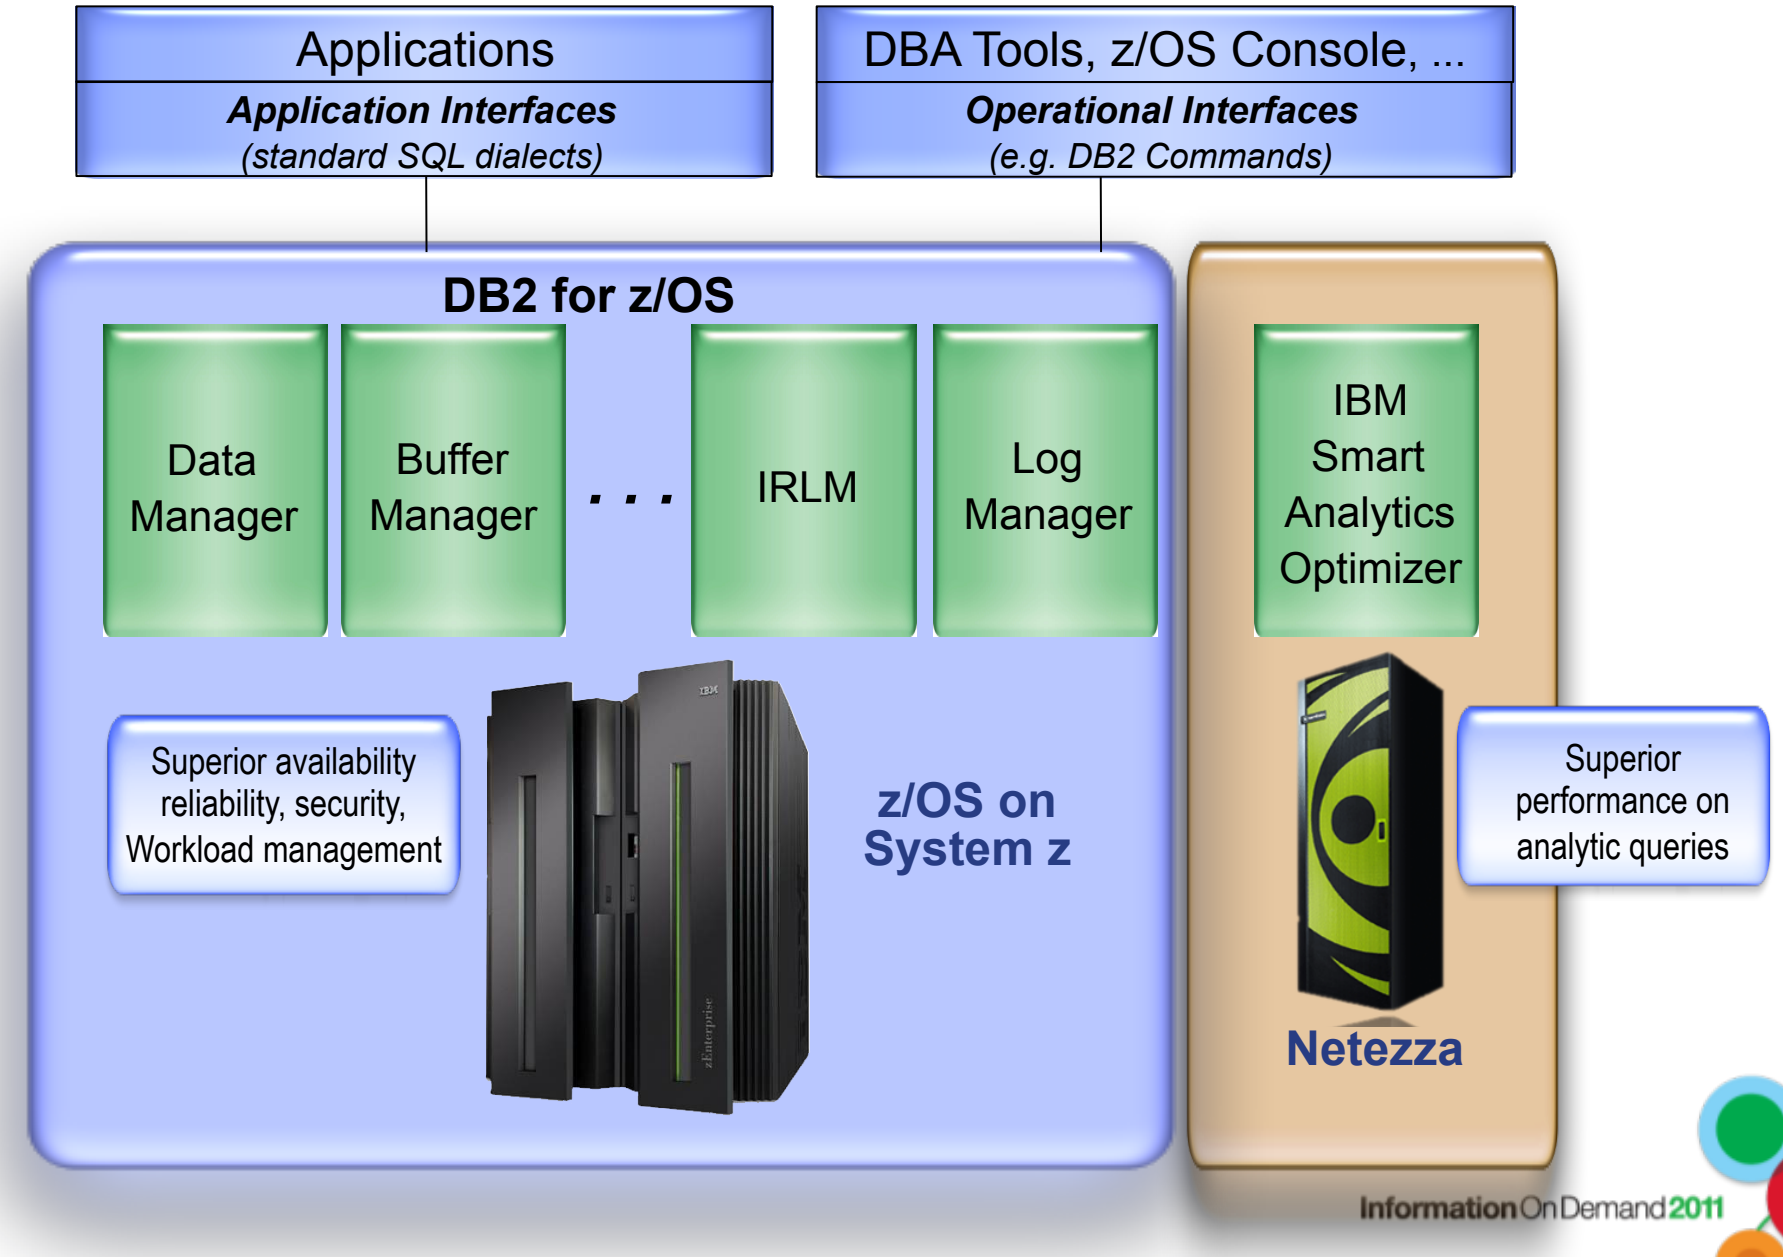

Query Execution Process Flow

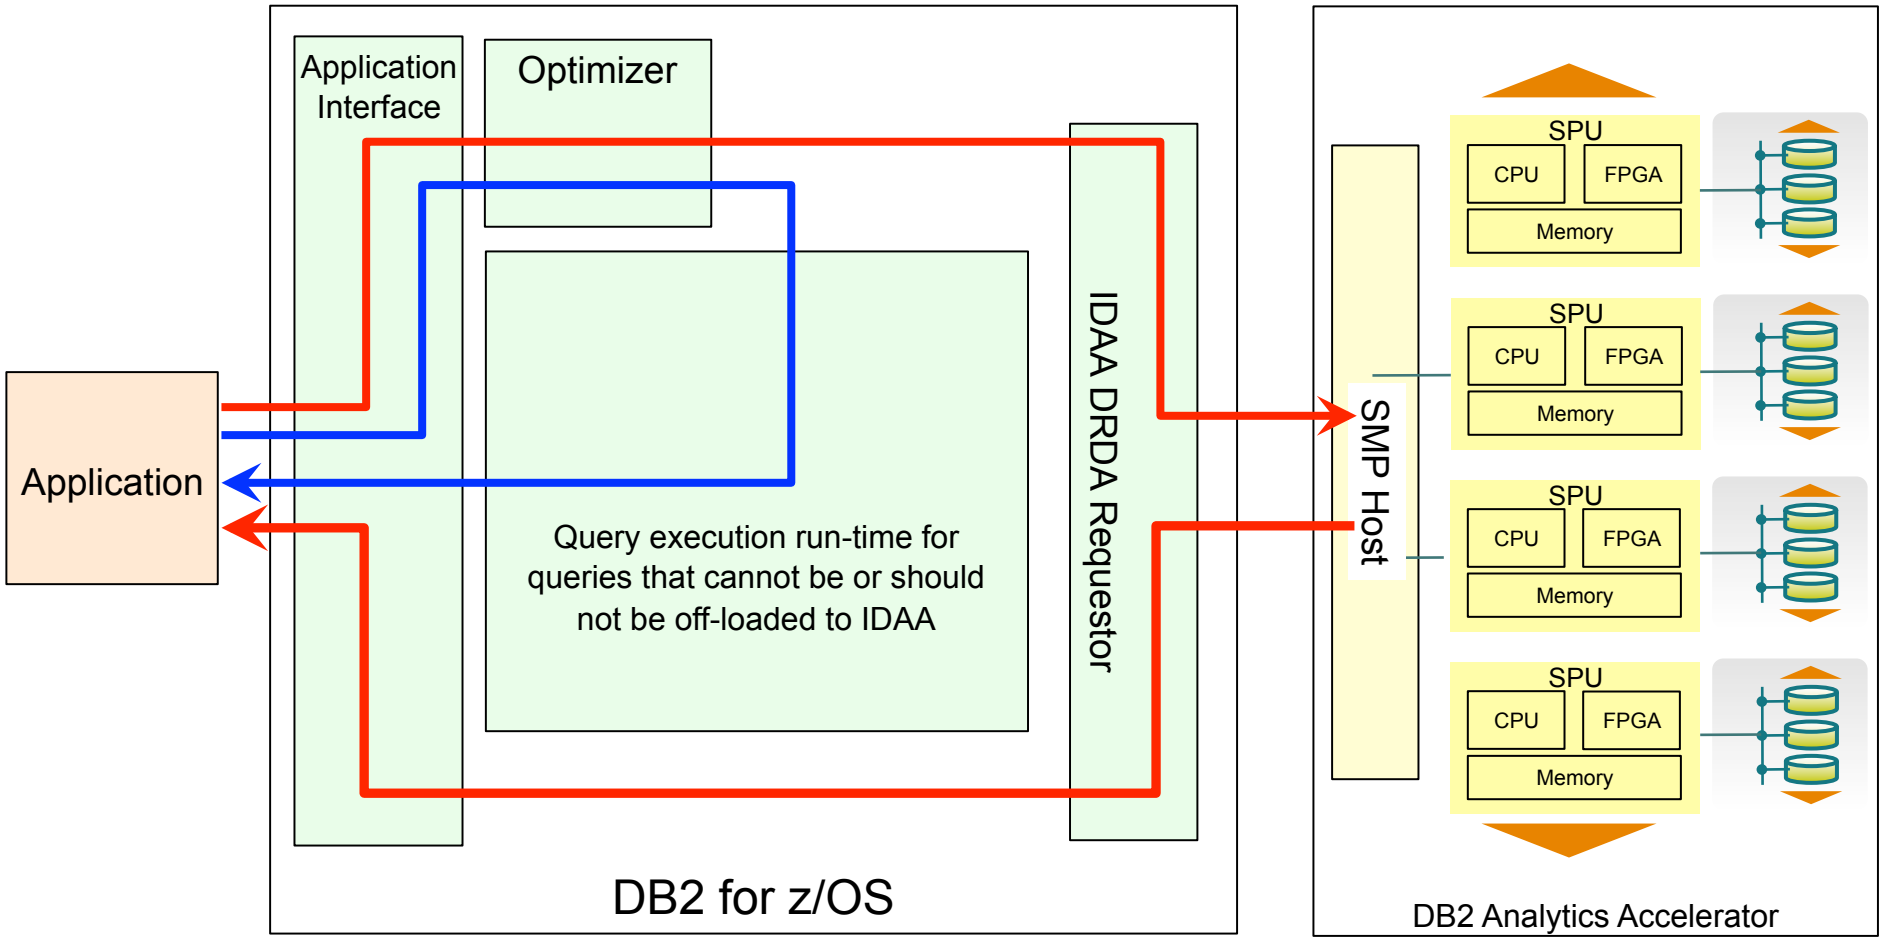

Heartbeat (DB2 Analytics Accelerator availability and performance indicators)

Queries executed without DB2 Analytics Accelerator

Queries executed with DB2 Analytics Accelerator

DB2 Analytics Accelerator Content Maintenance

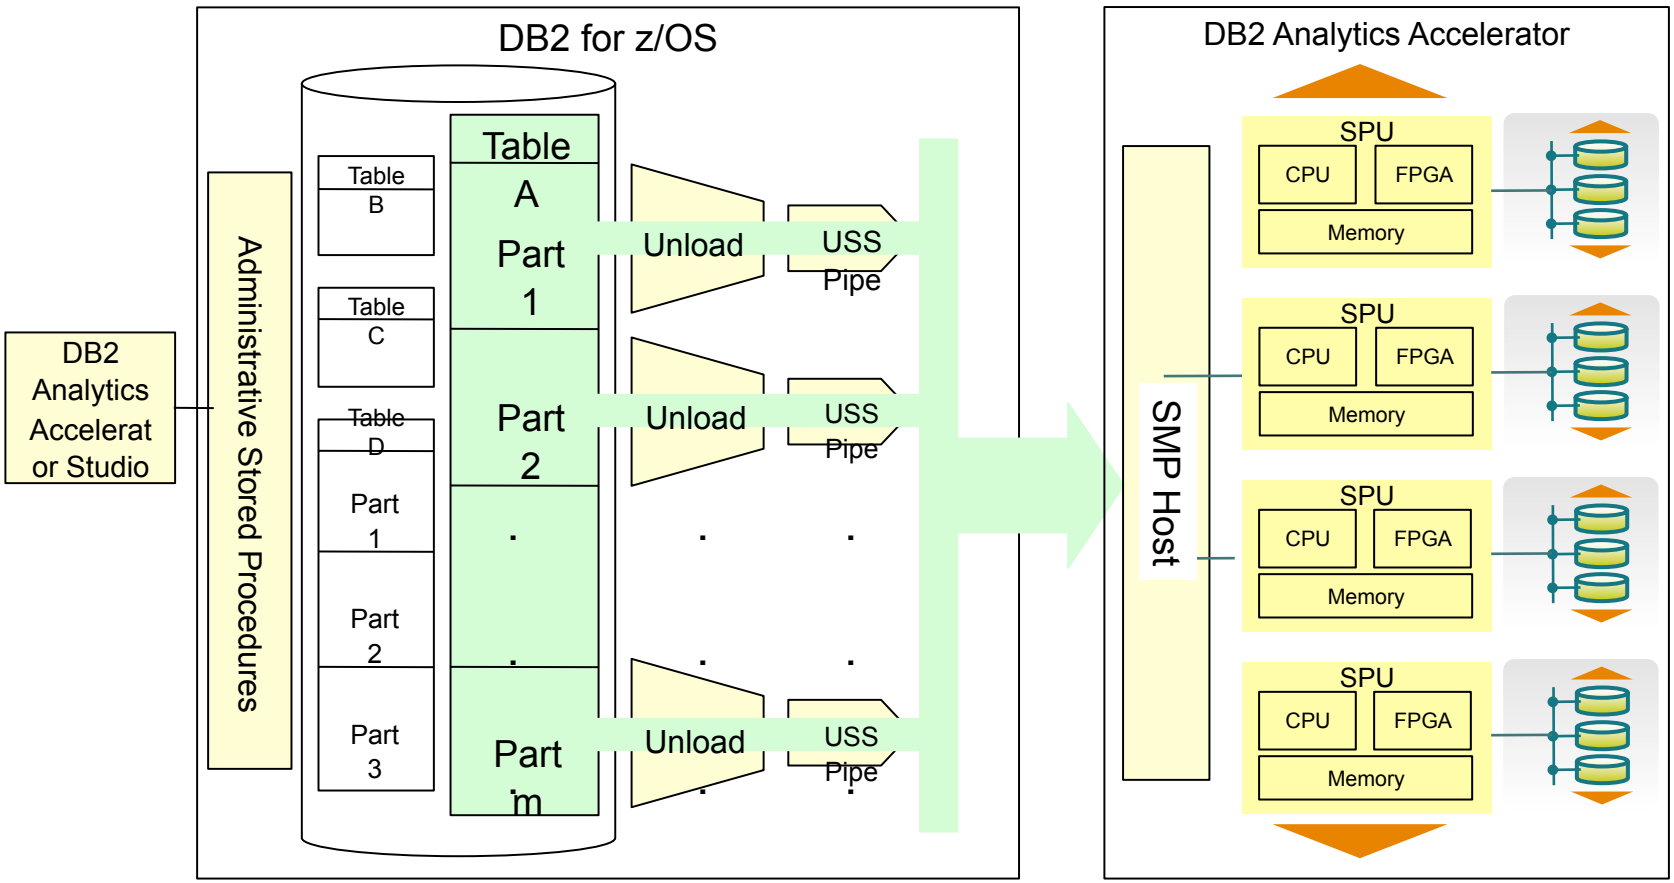

- Partitions belonging to the same table can be loaded in parallel
  - → User-defined degree of parallelism
- Updates are done on a per-table or per-partition level

IDAA Table Definition and Deployment

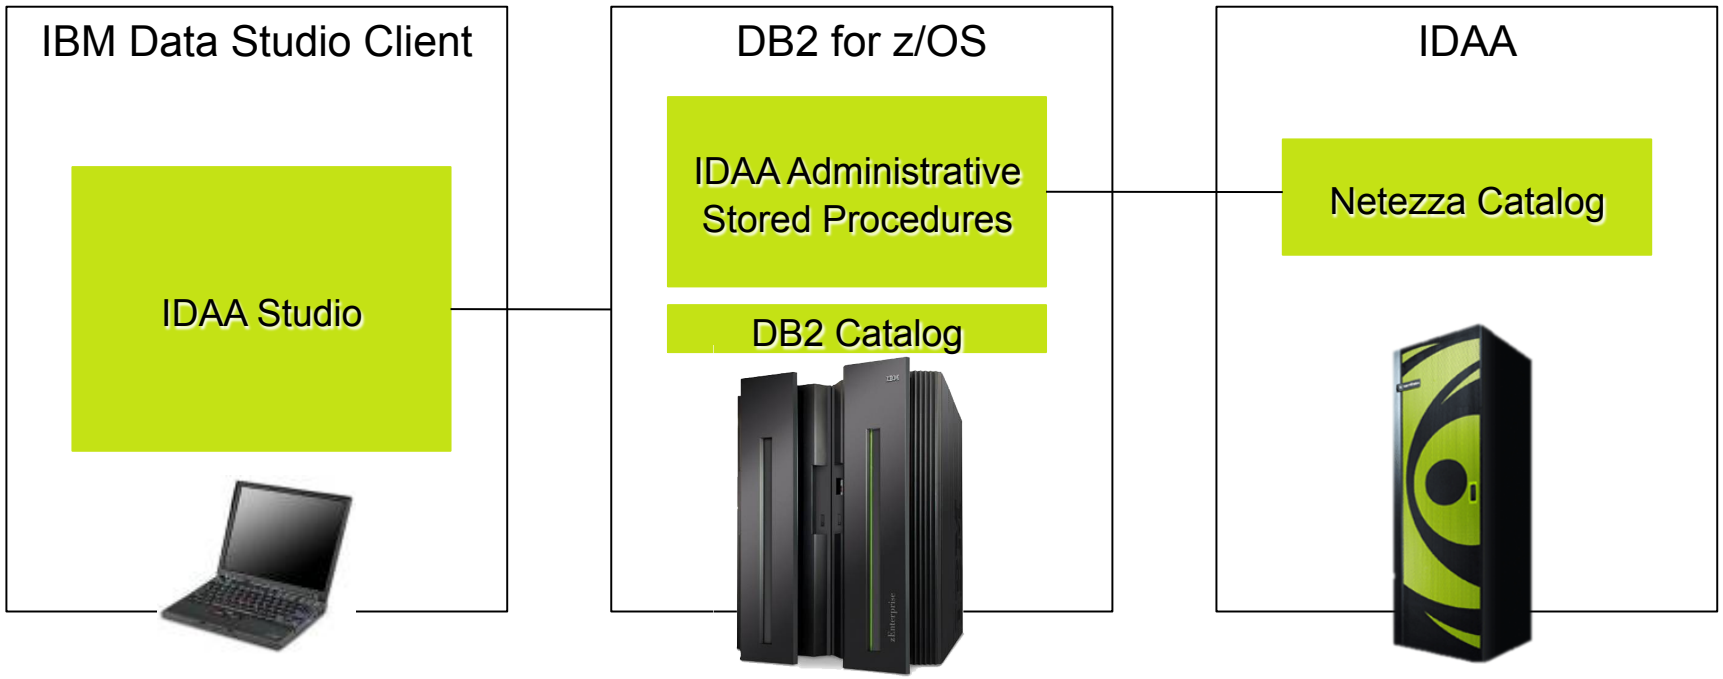

- The tables need to be defined and deployed to IDAA before data is loaded and queries sent to it for processing.
  - Definition: identifying tables for which queries need to be accelerated
  - Deployment: making tables known to DB2, i.e. storing table meta data in the DB2 and Netezza catalog.
- IDAA Studio guides you through the process of defining and deploying tables, as well as invoking other administrative tasks.
- IDAA 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.

Additional DB2 Support

- Explain
  - Indicates DB2 Analytics Accellerator involvement in query execution or the reason for no usage
  - New table DSN_QUERYINFO_TABLE
- Instrumentation
  - DB2 Analytics Accellerator availability indicators
  - Accounting, Statistics and Record Trace performance indicators
  - Statement-level reporting through IDAA Studio
- DB2 Commands
  - DISPLAY THREAD
  - DISPLAY ACCELerator
  - START ACCELerator
  - STOP ACCELerator

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

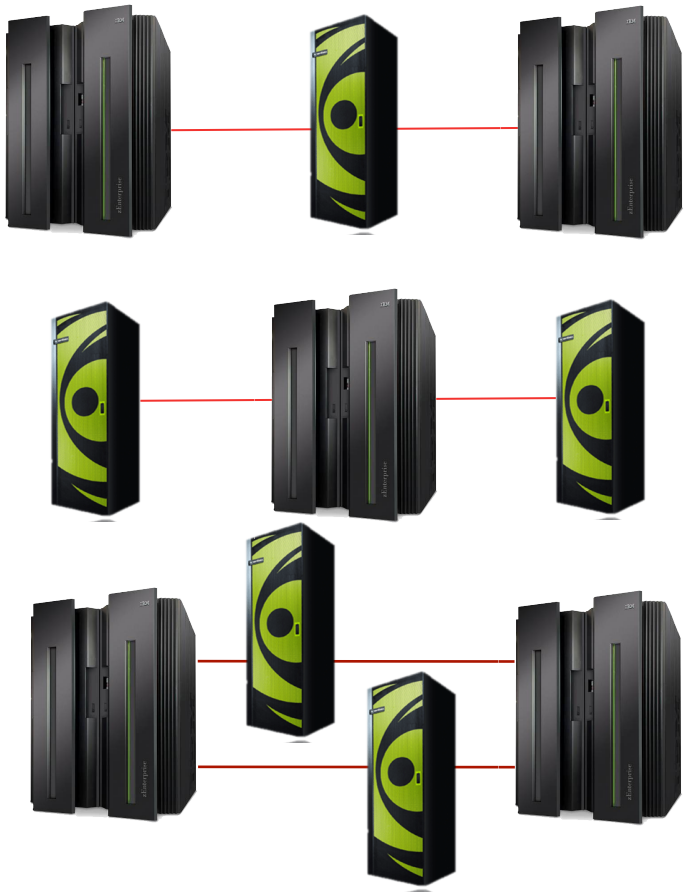

Better utilization of IDAA resources
Scalability
High availability

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

Disaster Recovery Considerations (1 of 3)

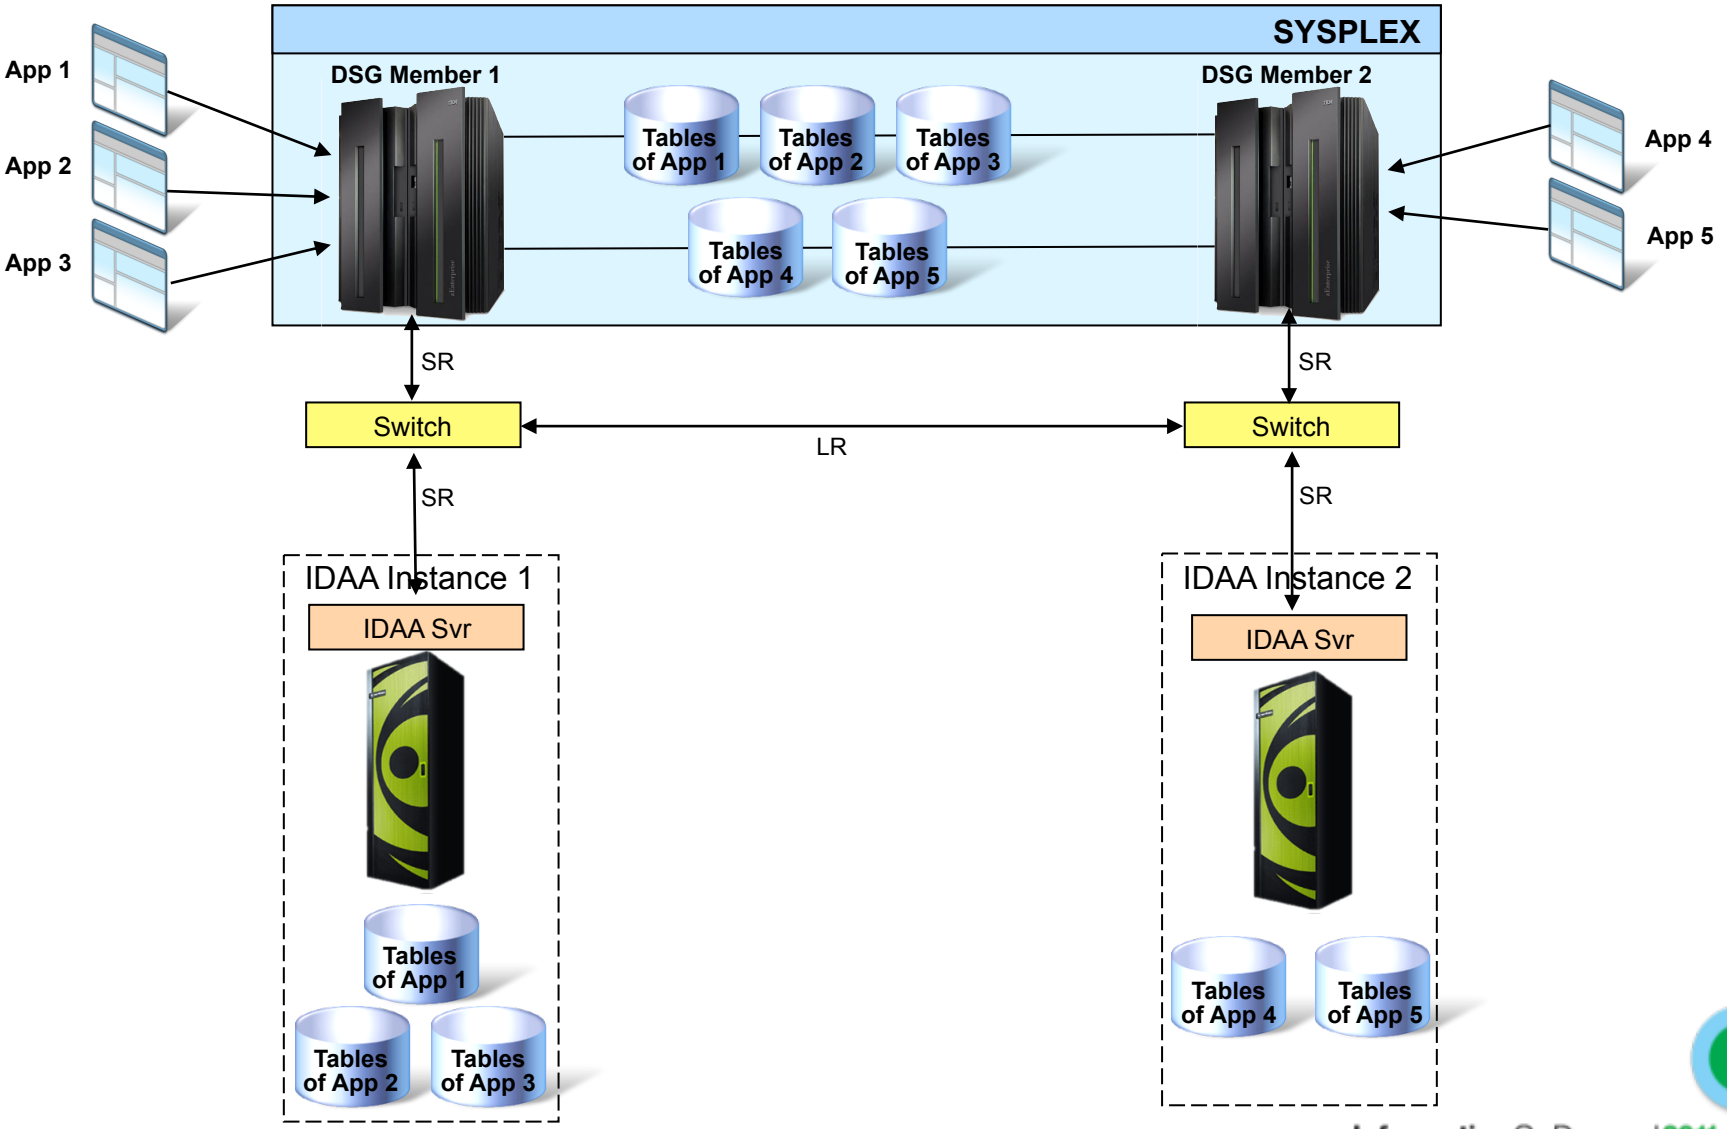

Disaster Recovery Considerations (2 of 3)

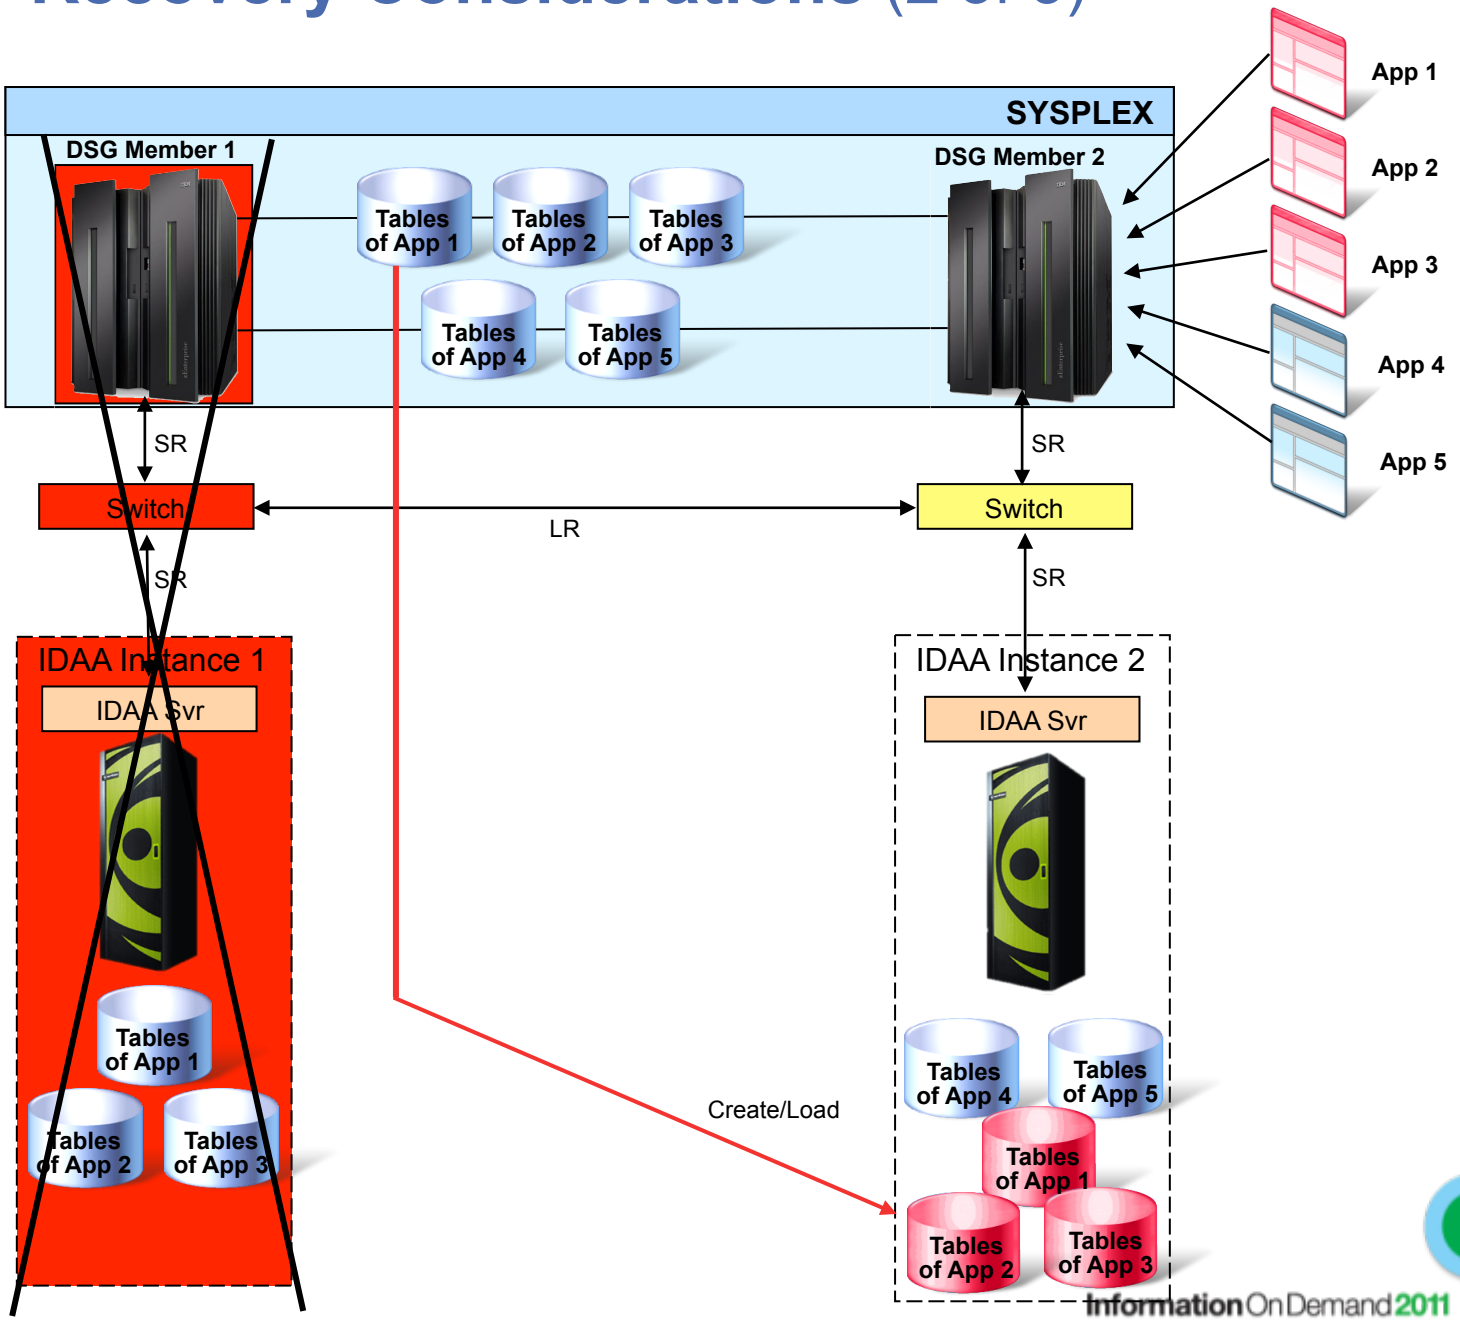

Disaster Recovery Considerations (2 of 3)

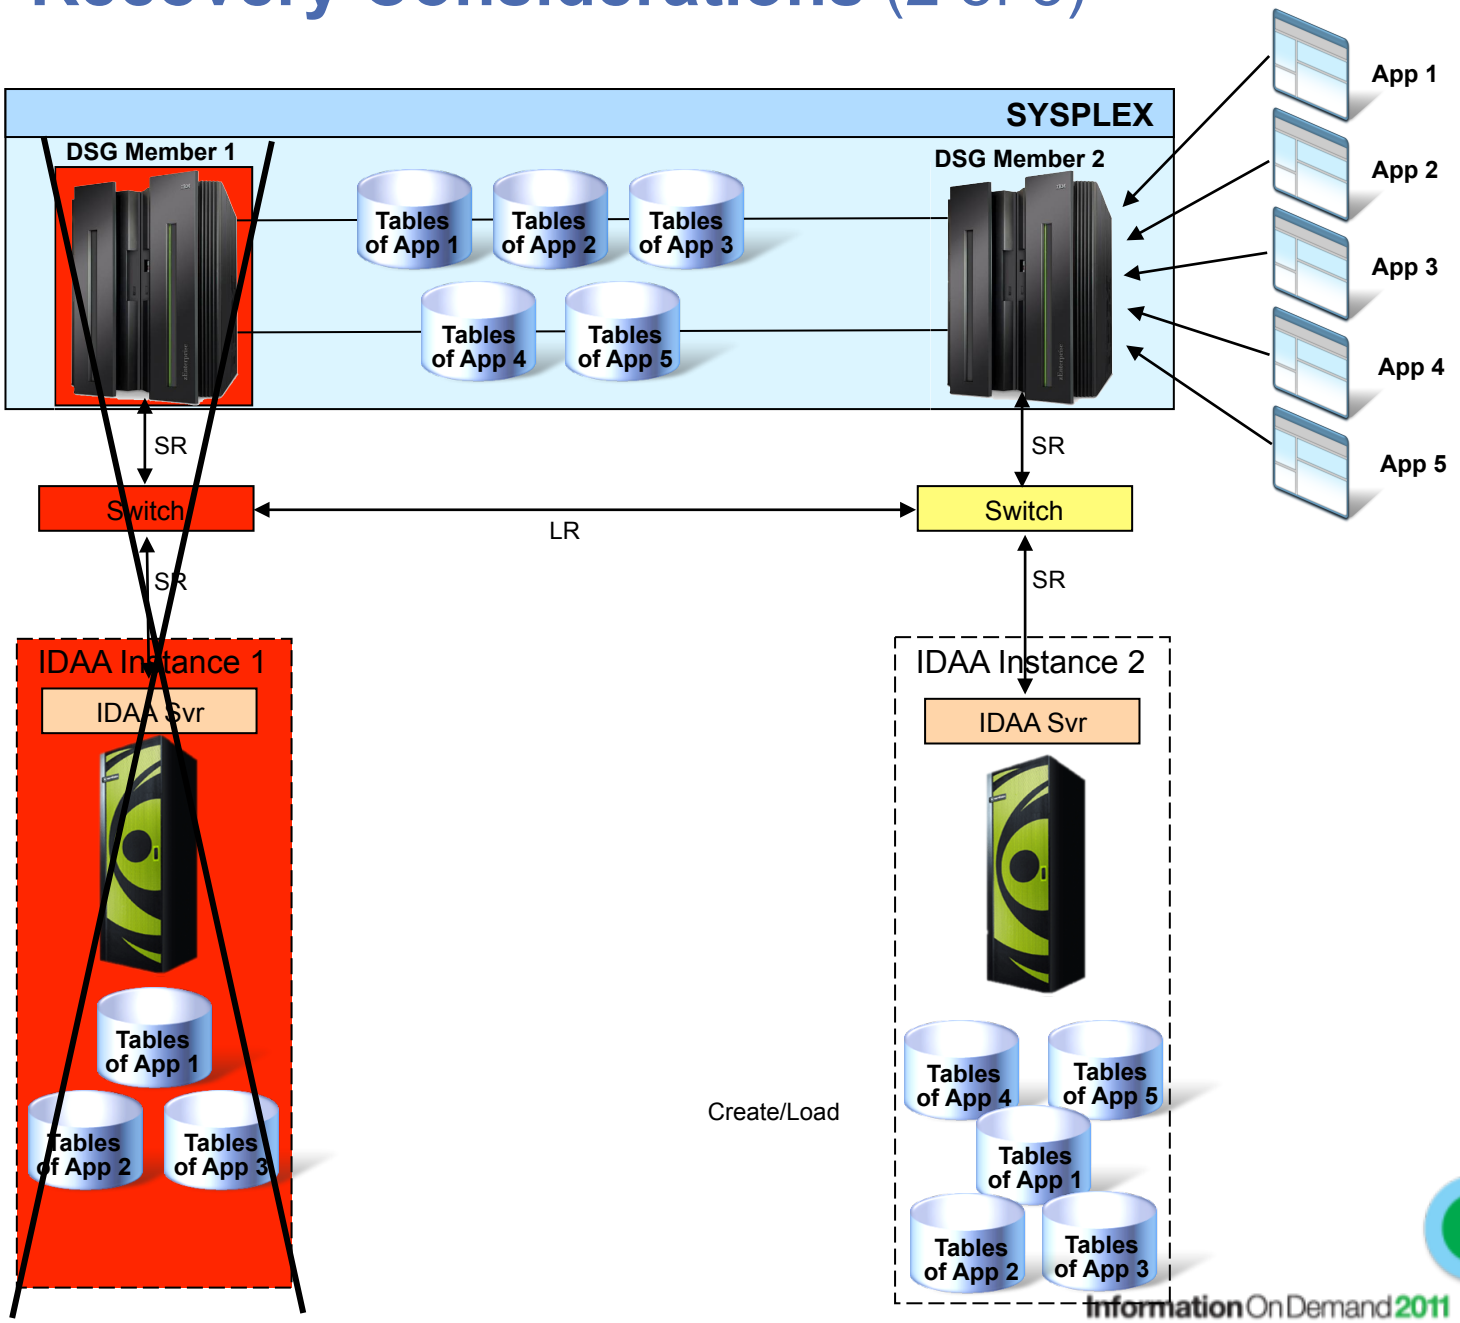

Agenda

- Overview of the Announcement
- Business and Technology Drivers
- Key Design and Operational Features
- Powered by Netezza
- Supported Workloads and Workload Applicability Assessment

Powered by Netezza 1000 Appliance

Disk Enclosures

SMP Hosts

Snippet Blades™
(S-Blades, SPUs)

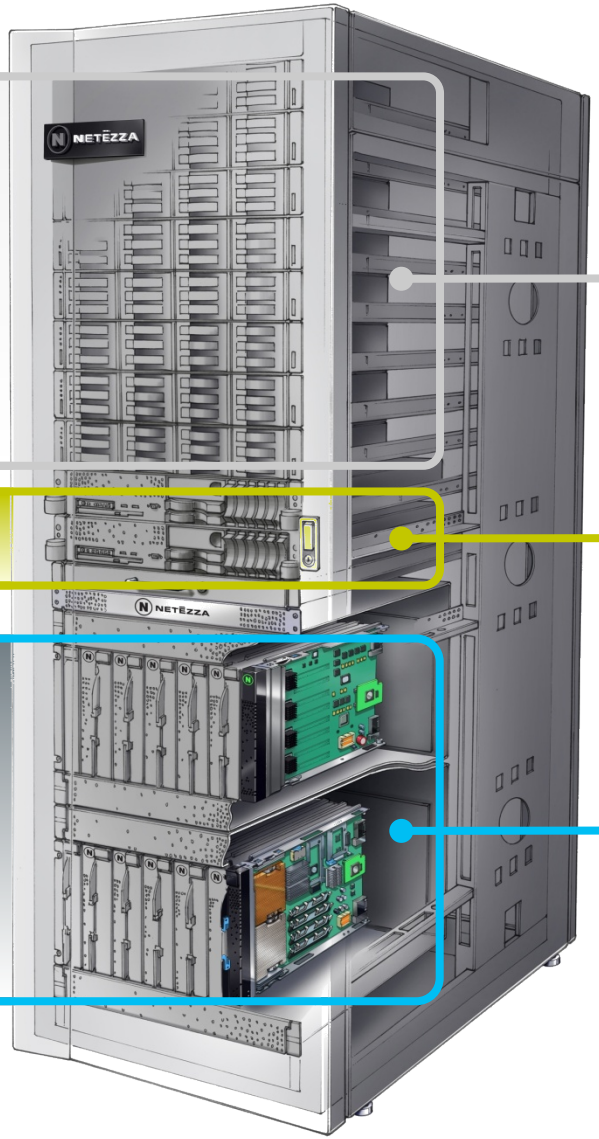

Slice of User Data
Swap and Mirror partitions
High speed data streaming
High compression rate
EXP3000 JBOD Enclosures
12 x 3.5" 1TB, 7200RPM, SAS (3Gb/s)
max 116MB/s (200-500MB/s compressed data)
e.g. TF12:
8 enclosures → 96 HDDs
32TB uncompressed user data (→ 128TB)

IDAA Server
SQL Compiler, Query Plan, Optimize
Administration
2 front/end hosts, IBM 3650M3 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. TF12: 12 back/end SPUs
(more details on following charts)

The Netezza S-Blade™

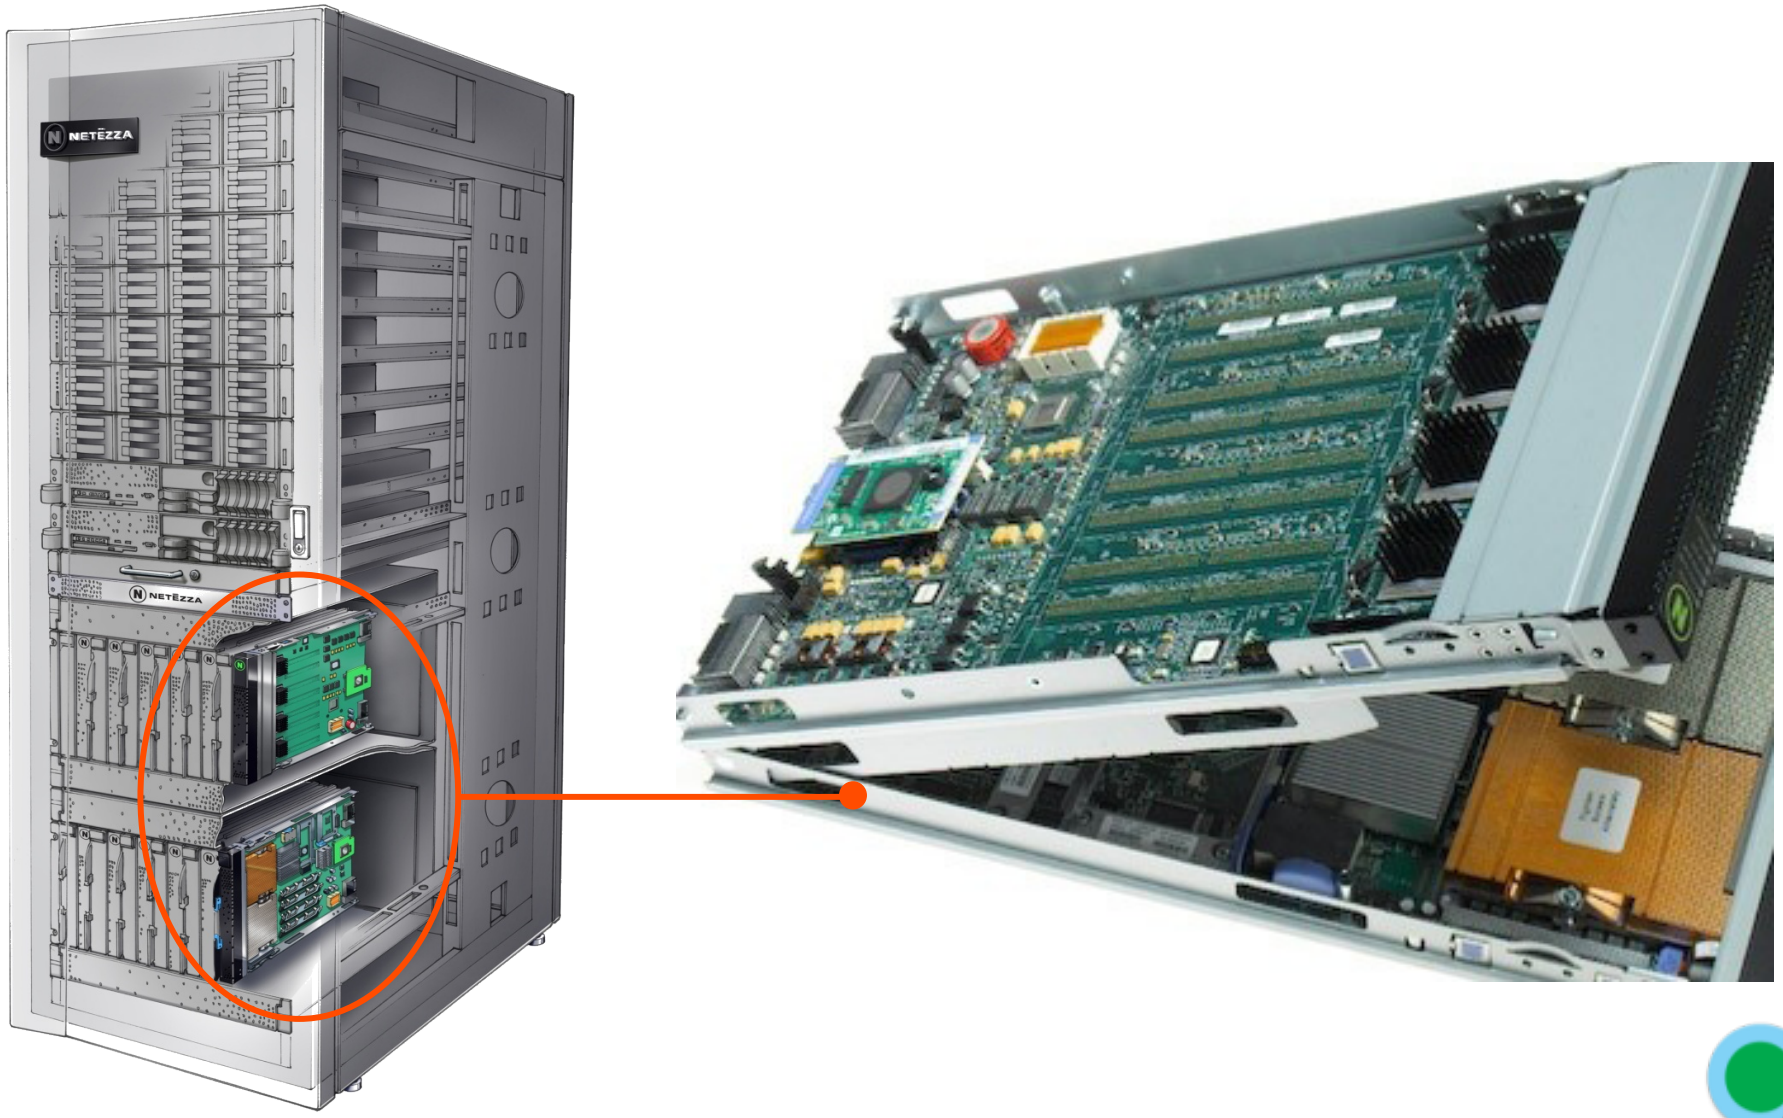

Bringing Netezza AMPP™ Architecture to DB2

AMPP = Asymmetric Massively Parallel Processing

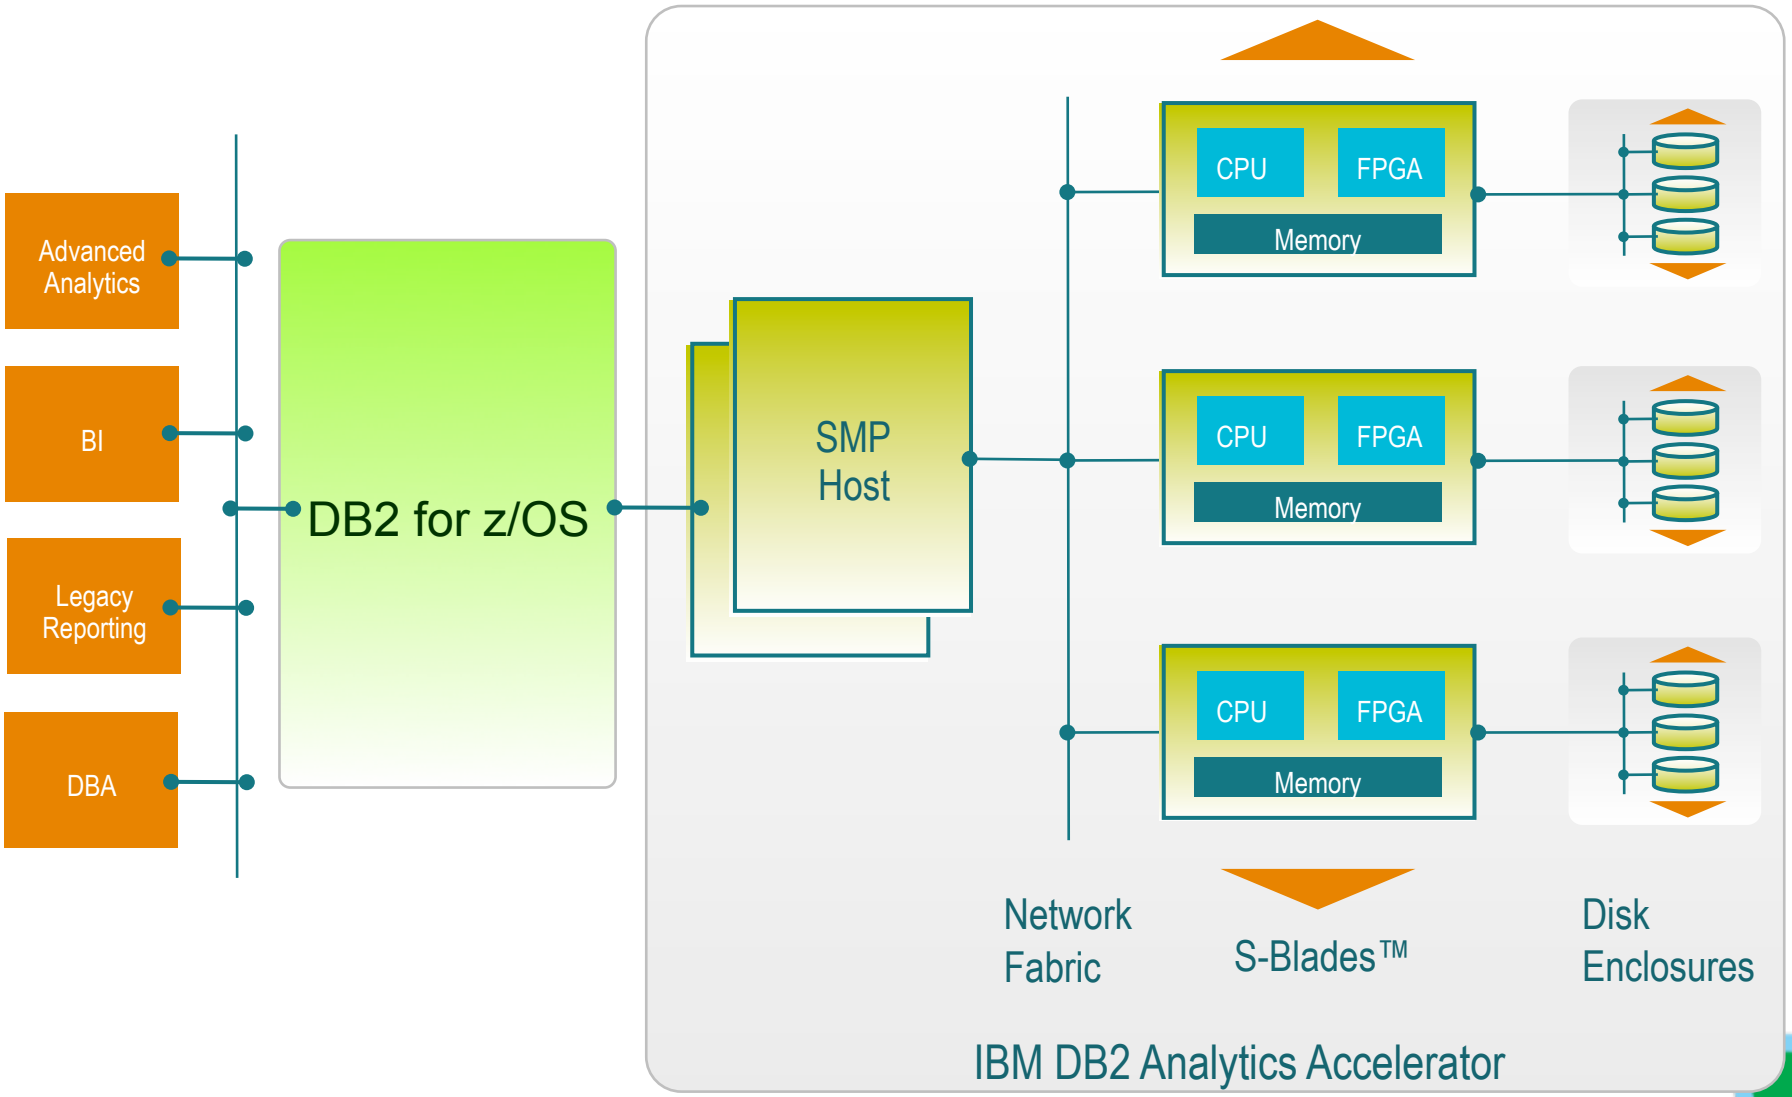

The Key to the Speed

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

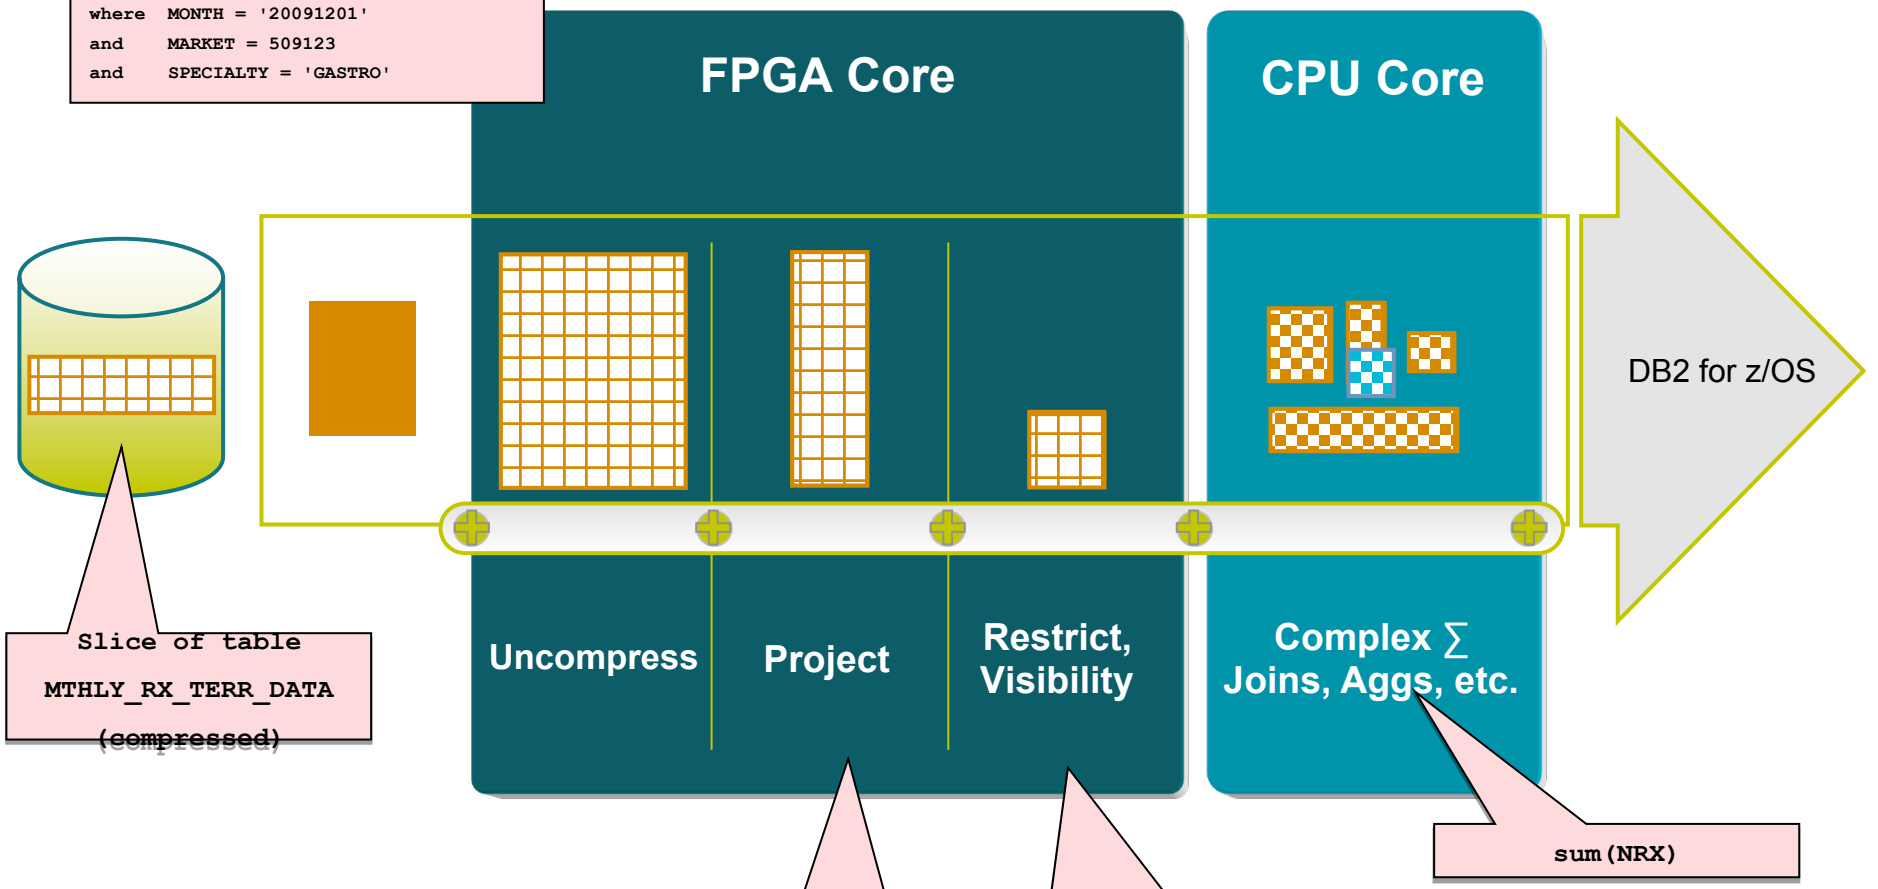

Shielding Against Disk Failures

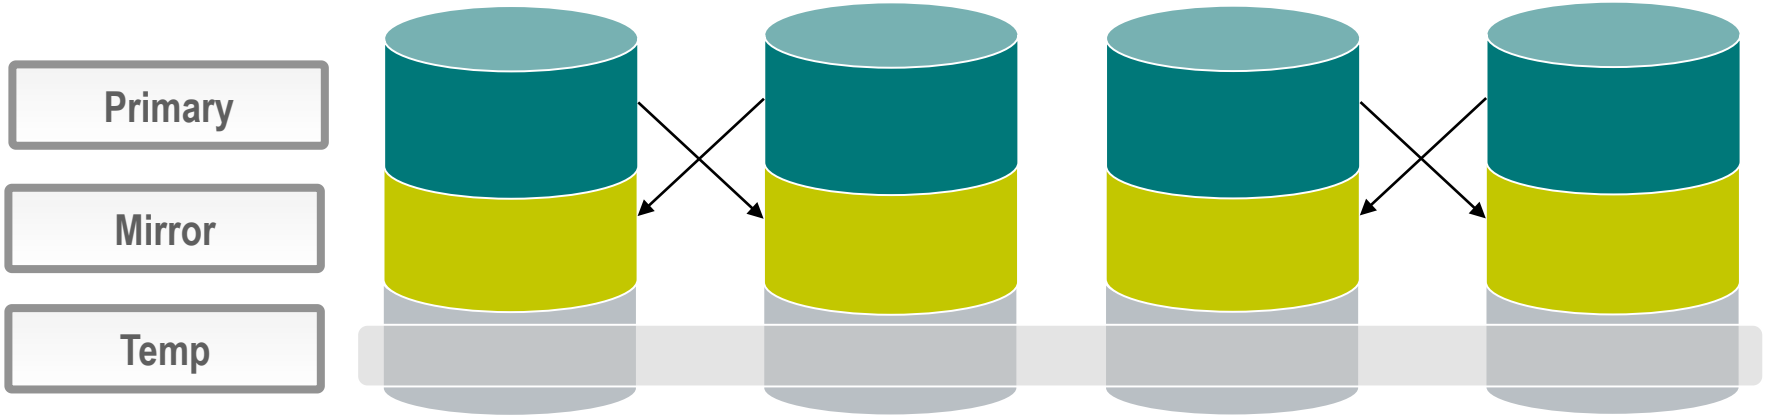

- All user data and temp space mirrored
- Disk failures transparent to queries and transactions
- Failed drives automatically regenerated
- Bad sectors automatically rewritten or relocated

Shielding Against S-Blade™ Failures

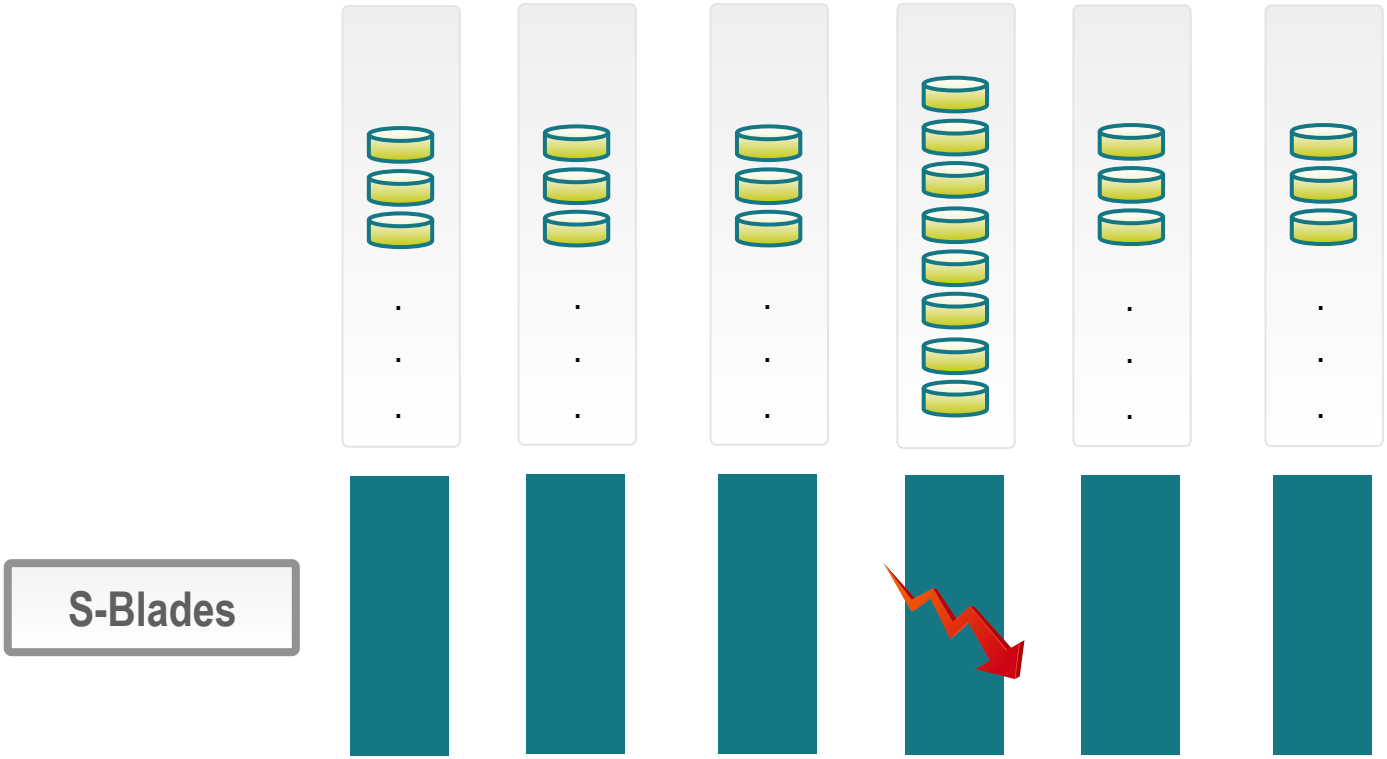

- S-Blade failure is automatically detected

Shielding Against S-Blade™ Failures

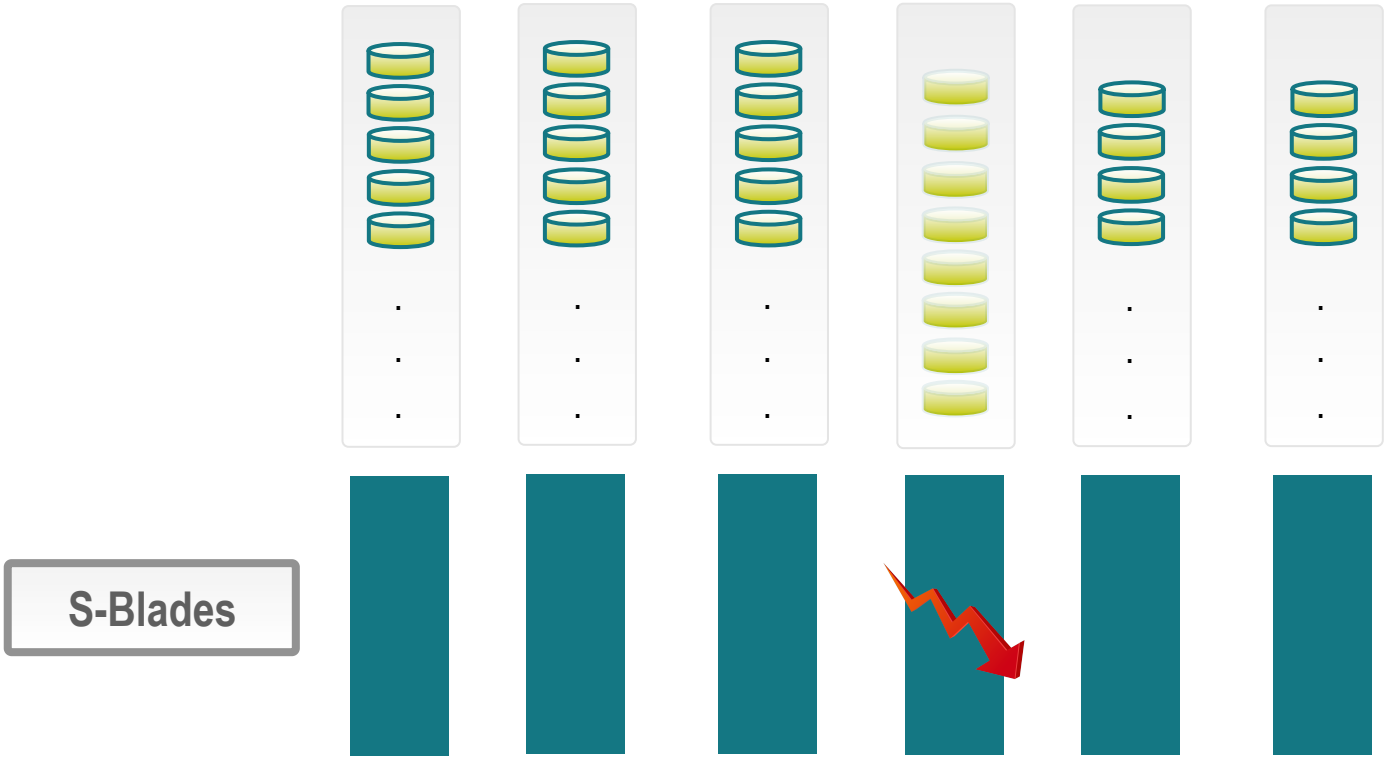

- Drives automatically reassigned to active S-Blades within a chassis
- Read-only queries (that have not returned data yet) automatically restarted
- Transactions and loads interrupted
- Loads automatically restarted from last successful checkpoint

Agenda

- Overview of the Announcement
- Business and Technology Drivers
- Key Design and Operational Features
- Powered by Netezza
- Supported Workloads and Workload Applicability Assessment

Query off-load applicability

- IDAA V2.1 is based on Netezza which supports rich set of SQL and data types
  - BI tools such as Cognos has been running on Netezza for years and will run on IDAA as well
- Due to very large number of query types and SQL functions not all of them could have been contained in the release. Here are the key restrictions:
  - No static SQL
  - Not all DB2 functions, such as HEX, POSITION, SIN, ...
  - No user defined functions
  - No correlated table expressions or recursive correlated table expressions
  - No correlated subquery in the SELECT list
  - Not UTF-16 and MIXED/DBCS EBCDIC
  - No multiple encoding schemes in the same statement
  - Not all DB2 special registers: CURRENT PATH, SERVER, SQLID, SCHEMA, APPLICATION ENCODING SCHEME
  - Not all DB2 data types such as LOBs, ROWID, XML.
- None of these restrictions is seen as a design problem.

Options for Workload Analysis

Stage	Purpose
Questionnaire	• Initial assessment based on size, query response time, update characteristics and customer pain points
Quick Workload Test	• Assessment based on dynamic customer workload, runtime statistics, table sizes and SQL.
Detailed Online Workload Analysis	• Assessment based on data mart definition for customer data model and offload capabilities in a real DB2 Analytics Accelerator environment. Addresses all inhibitors for offload and data mart definition questions.

Resources:

- Instruction for the process at https://w3.tap.ibm.com/w3ki08/display/isao/Process
  - Includes capturing description and jobs/scripts to support the procedure
- Contact the Data Warehousing on System z Center of Excellence at dwhz@de.ibm.com

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

Report for a first assessment:
- Acceleration potential for
  - Queries
  - Estimated time
  - CP cost

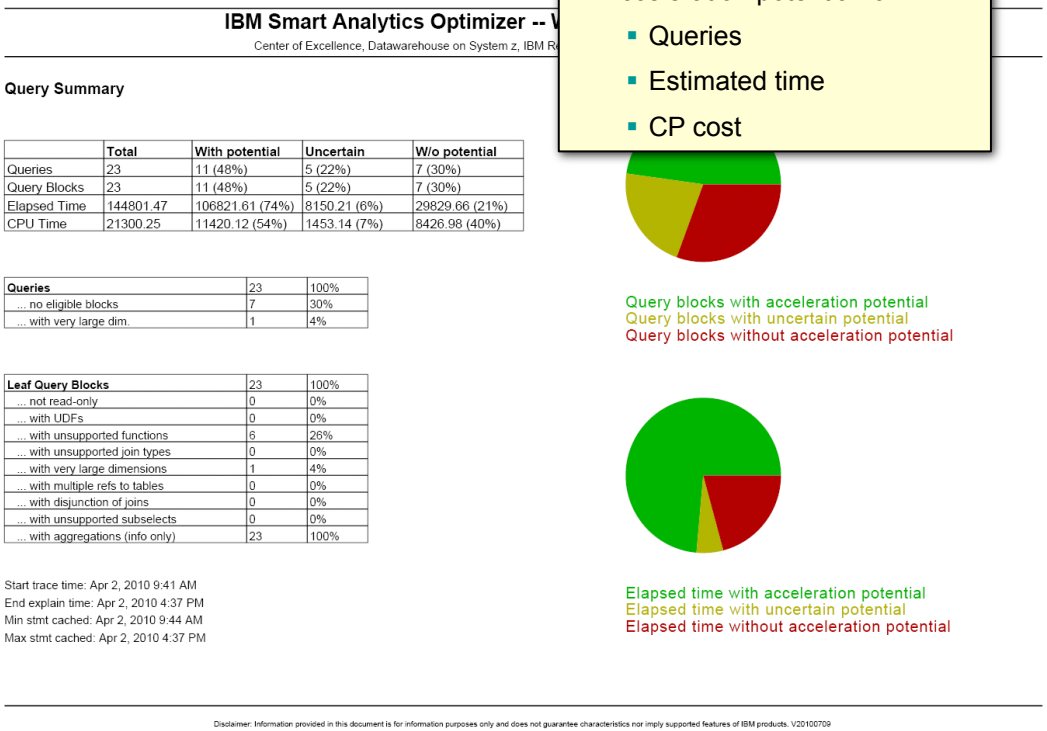

IBM Smart Analytics Optimizer -- V

Center of Excellence, Datawarehouse on System z, IBM R

Query Summary

	Total	With potential	Uncertain	W/o potential
Queries	23	11 (48%)	5 (22%)	7 (30%)
Query Blocks	23	11 (48%)	5 (22%)	7 (30%)
Elapsed Time	144801.47	106821.61 (74%)	8150.21 (6%)	29829.66 (21%)
CPU Time	21300.25	11420.12 (54%)	1453.14 (7%)	8426.98 (40%)

Queries	23	100%
... no eligible blocks	7	30%
... with very large dim.	1	4%

Leaf Query Blocks	23	100%
... not read-only	0	0%
... with UDFs	0	0%
... with unsupported functions	6	26%
... with unsupported join types	0	0%
... with very large dimensions	1	4%
... with multiple refs to tables	0	0%
... with disjunction of joins	0	0%
... with unsupported subselects	0	0%
... with aggregations (info only)	23	100%

Start trace time: Apr 2, 2010 9:41 AM
End explain time: Apr 2, 2010 4:37 PM
Min stmt cached: Apr 2, 2010 9:44 AM
Max stmt cached: Apr 2, 2010 4:37 PM

Query blocks with acceleration potential
Query blocks with uncertain potential
Query blocks without acceleration potential

Elapsed time with acceleration potential
Elapsed time with uncertain potential
Elapsed time without acceleration potential

Disclaimer: Information provided in this document is for information purposes only and does not guarantee characteristics nor imply supported features of IBM products. V20100709

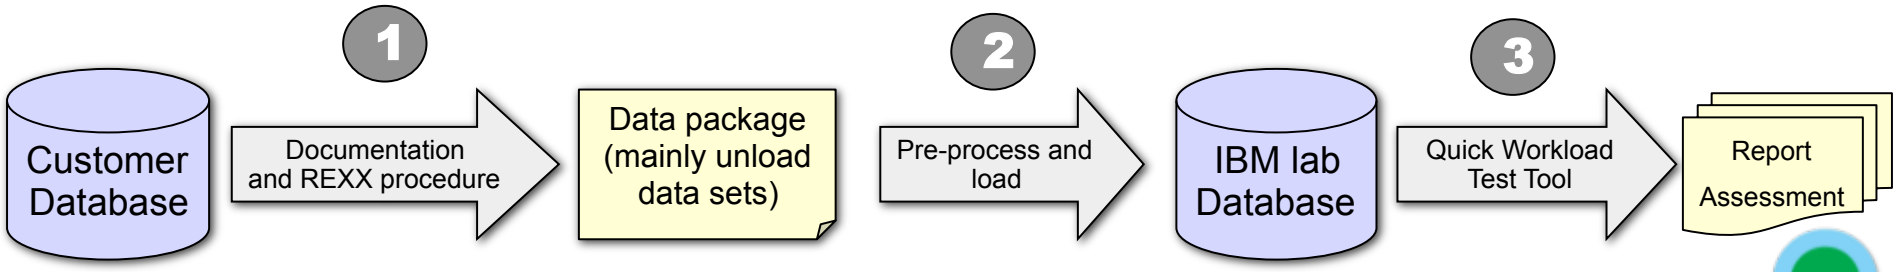

Workload Assessment Report

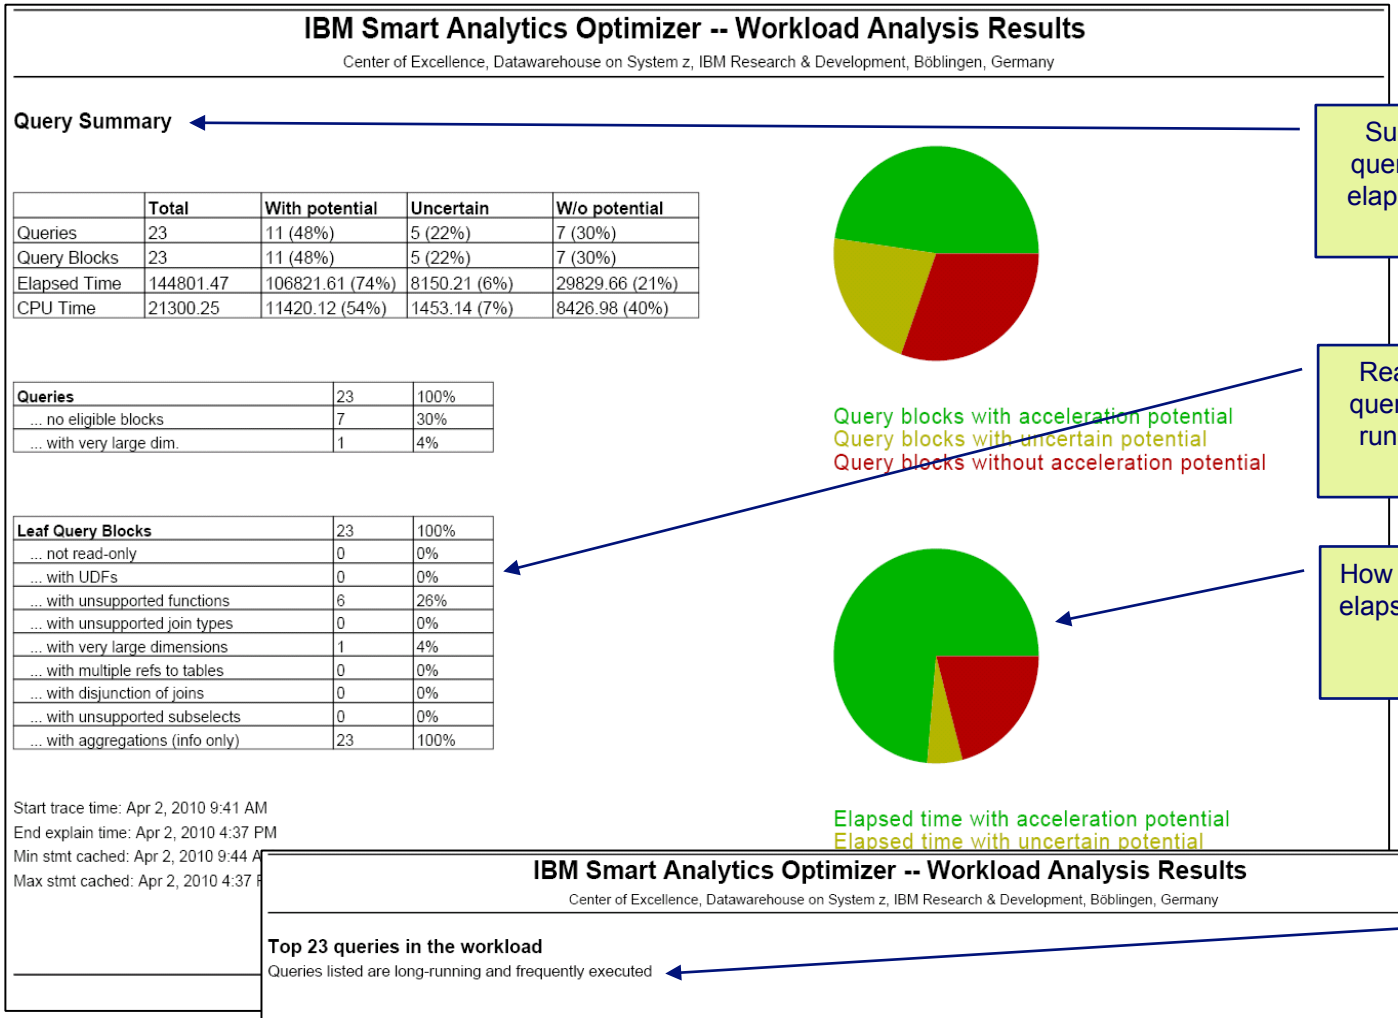

IBM Smart Analytics Optimizer -- Workload Analysis Results

Center of Excellence, Datawarehouse on System z, IBM Research & Development, Böblingen, Germany

Query Summary

	Total	With potential	Uncertain	W/o potential
Queries	23	11 (48%)	5 (22%)	7 (30%)
Query Blocks	23	11 (48%)	5 (22%)	7 (30%)
Elapsed Time	144801.47	106821.61 (74%)	8150.21 (6%)	29829.66 (21%)
CPU Time	21300.25	11420.12 (54%)	1453.14 (7%)	8426.98 (40%)

Queries	23	100%
... no eligible blocks	7	30%
... with very large dim.	1	4%

Leaf Query Blocks	23	100%
... not read-only	0	0%
... with UDFs	0	0%
... with unsupported functions	6	26%
... with unsupported join types	0	0%
... with very large dimensions	1	4%
... with multiple refs to tables	0	0%
... with disjunction of joins	0	0%
... with unsupported subselects	0	0%
... with aggregations (info only)	23	100%

Start trace time: Apr 2, 2010 9:41 AM
End explain time: Apr 2, 2010 4:37 PM
Min stmt cached: Apr 2, 2010 9:44 A
Max stmt cached: Apr 2, 2010 4:37 P

Query blocks with acceleration potential
Query blocks with uncertain potential
Query blocks without acceleration potential

Elapsed time with acceleration potential
Elapsed time with uncertain potential

IBM Smart Analytics Optimizer -- Workload Analysis Results

Center of Excellence, Datawarehouse on System z, IBM Research & Development, Böblingen, Germany

Top 23 queries in the workload

Queries listed are long-running and frequently executed

Query	Elapsed Time	Executed	Blocks	Eligible		Elapsed Time with pot./uncert.	Table Sizes	Statement
8	83117.95	1	1		1	100% / 0%	283.33 GB	(Primauth: SCHRAY) SELECT YEAR(L_SHIPDATE) ,MONTH(L_SHIPDATE) ,SUM(L_QUANTITY*L_EXTENDEDPRICE) FROM SYSADM1.LINEITEM GROUP BY YEAR(L_SHIPDATE) ,MONTH(L_SHIPDATE)
5	16595.42	1	1		1	100% / 0%	39.08 GB	(Primauth: SCHRAY) SELECT YEAR(O_ORDERDATE) ,MONTH(O_ORDERDATE) ,SUM(O_TOTALPRICE) FROM SYSADM1.ORDERS GROUP BY YEAR(O_ORDERDATE) ,MONTH(O_ORDERDATE)
7	10592.58	1	1		0	0% / 0%	24.32 GB	(Primauth: SCHRAY) SELECT S_NATIONKEY ,SUM(FLOAT(PS_SUPPLYCOST * PS_AVAILQTY)) AS SUM_SUPPLYCOST FROM SYSADM1.SUPPLIER ,SYSADM1.PARTSUPP WHERE S_SUPPKEY = PS_SUPPKEY GROUP BY S_NATIONKEY
12	8915.12	1	1		0	0% / 0%	46.64 GB	(Primauth: SCHRAY) SELECT C_MKTSEGMENT ,COUNT(*) AS #ORDERS ,SUM(FLOAT(O_TOTALPRICE)) AS SUM_TOTALPRICE ,AVG(FLOAT(O_TOTALPRICE)) AS AVG_TOTALPRICE FROM SYSADM1.ORDERS ,SYSADM1.CUSTOMER WHERE C_CUSTKEY = O_CUSTKEY GROUP BY C_MKTSEGMENT

Summary based on queries, query blocks, elapsed time and CPU time

Reasons why certain query blocks might not run on DB2 Analytics Accelerator

How much of the current elapsed time may run on DB2 Analytics Accelerator

Detailed query-level assessment of the workload

SQL statement per query

Elapsed time per query


Thank You