Information Management

IBM

The Next Generation Data Warehouse
Smart Analytics

DB2 for z/OS Conference
Istanbul, 26-27 April 2010

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

April 2010

Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.

Agenda

- System z as the Information Management Hub
  - Business and Technology Drivers
  - Data Warehousing Components and Offerings
- IBM Smart Analytics Optimizer
  - Key Design and Operational Features
  - Supported workloads

IBM's Global CIO Study Top Priority: BI and Analytics

Leveraging analytics to gain a competitive advantage and improve business decision-making is now the top priority for CIOs

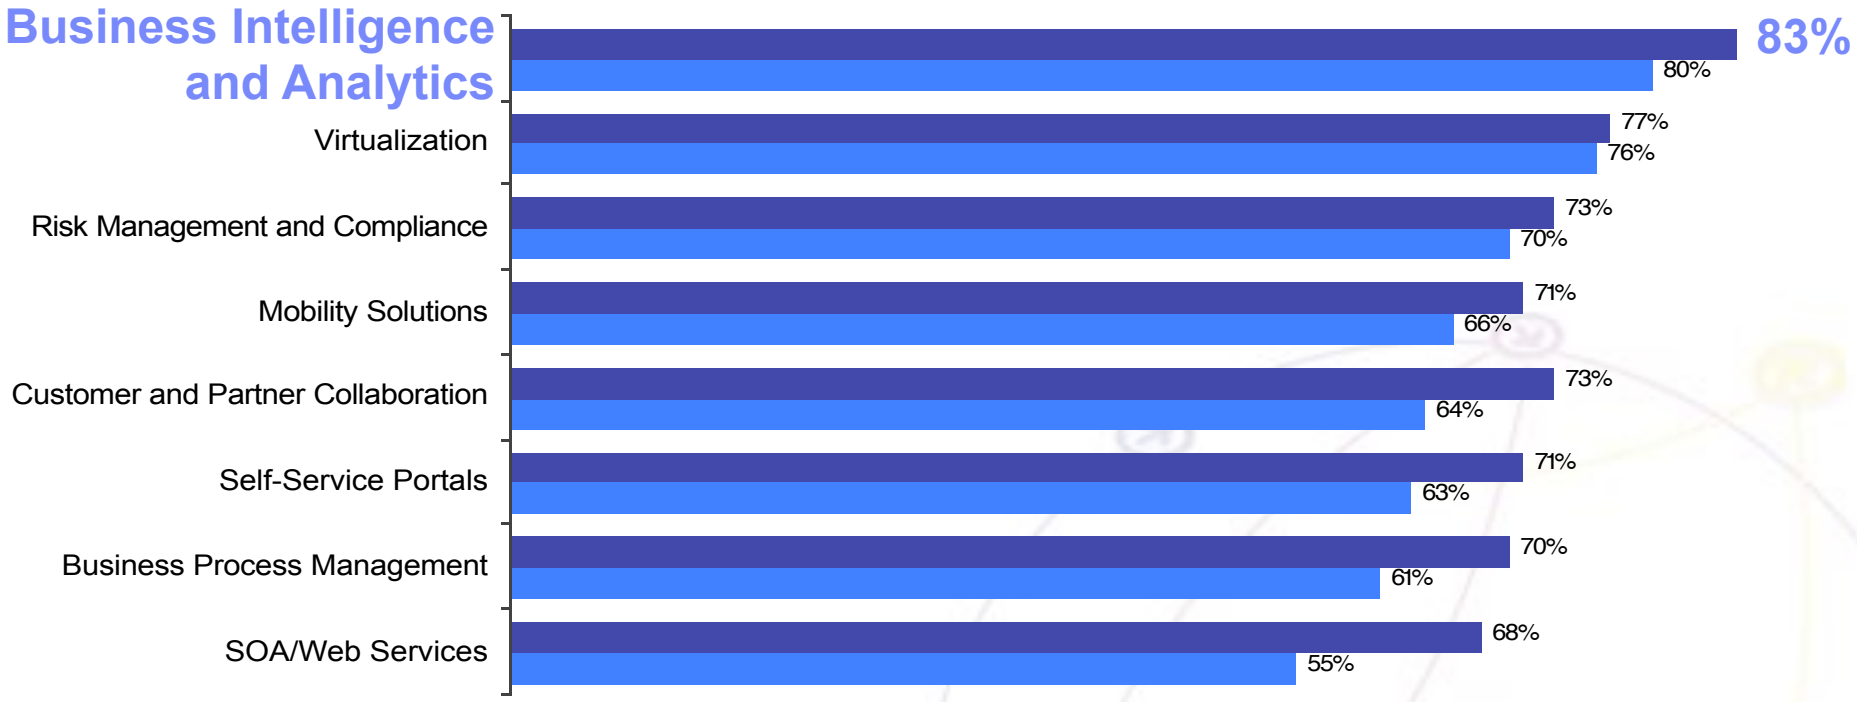

CIOs at high-growth companies pro-actively craft data into actionable information much more often than those at low-growth companies

Source: IBM Global CIO Study 2009; n = 2345

System z To Address BI/DW Business Challenges

- Most clients see a large gap between current positions and desired situations in the Business Intelligence and Data Warehousing space
- Surveys stress the importance of IM
  - By 2012, BI Platform capabilities will be embedded as a service within 75% of new business applications
  - IBM 2009 CIO Survey: 83% say BI and Analytics is the top focus area
  - Gartner 2007 Survey: Fewer than 15% of data warehouses have been designed to provide high availability, failover, disaster recovery and the remaining components of mission-critical systems.
- Changing business requirements suggests taking advantage of System z technology
  - BI/DW becoming mission critical and requires OLTP-like Quality of Service
    - Reliability, continuous availability, security, mixed workload management, …
  - Shift towards dynamic DW and operational BI
    - Combining OLTP and OLAP workloads

Information-Led Transformation for System z

More new capabilities delivered
in the past 3 years
than at any point in the mainframe's history

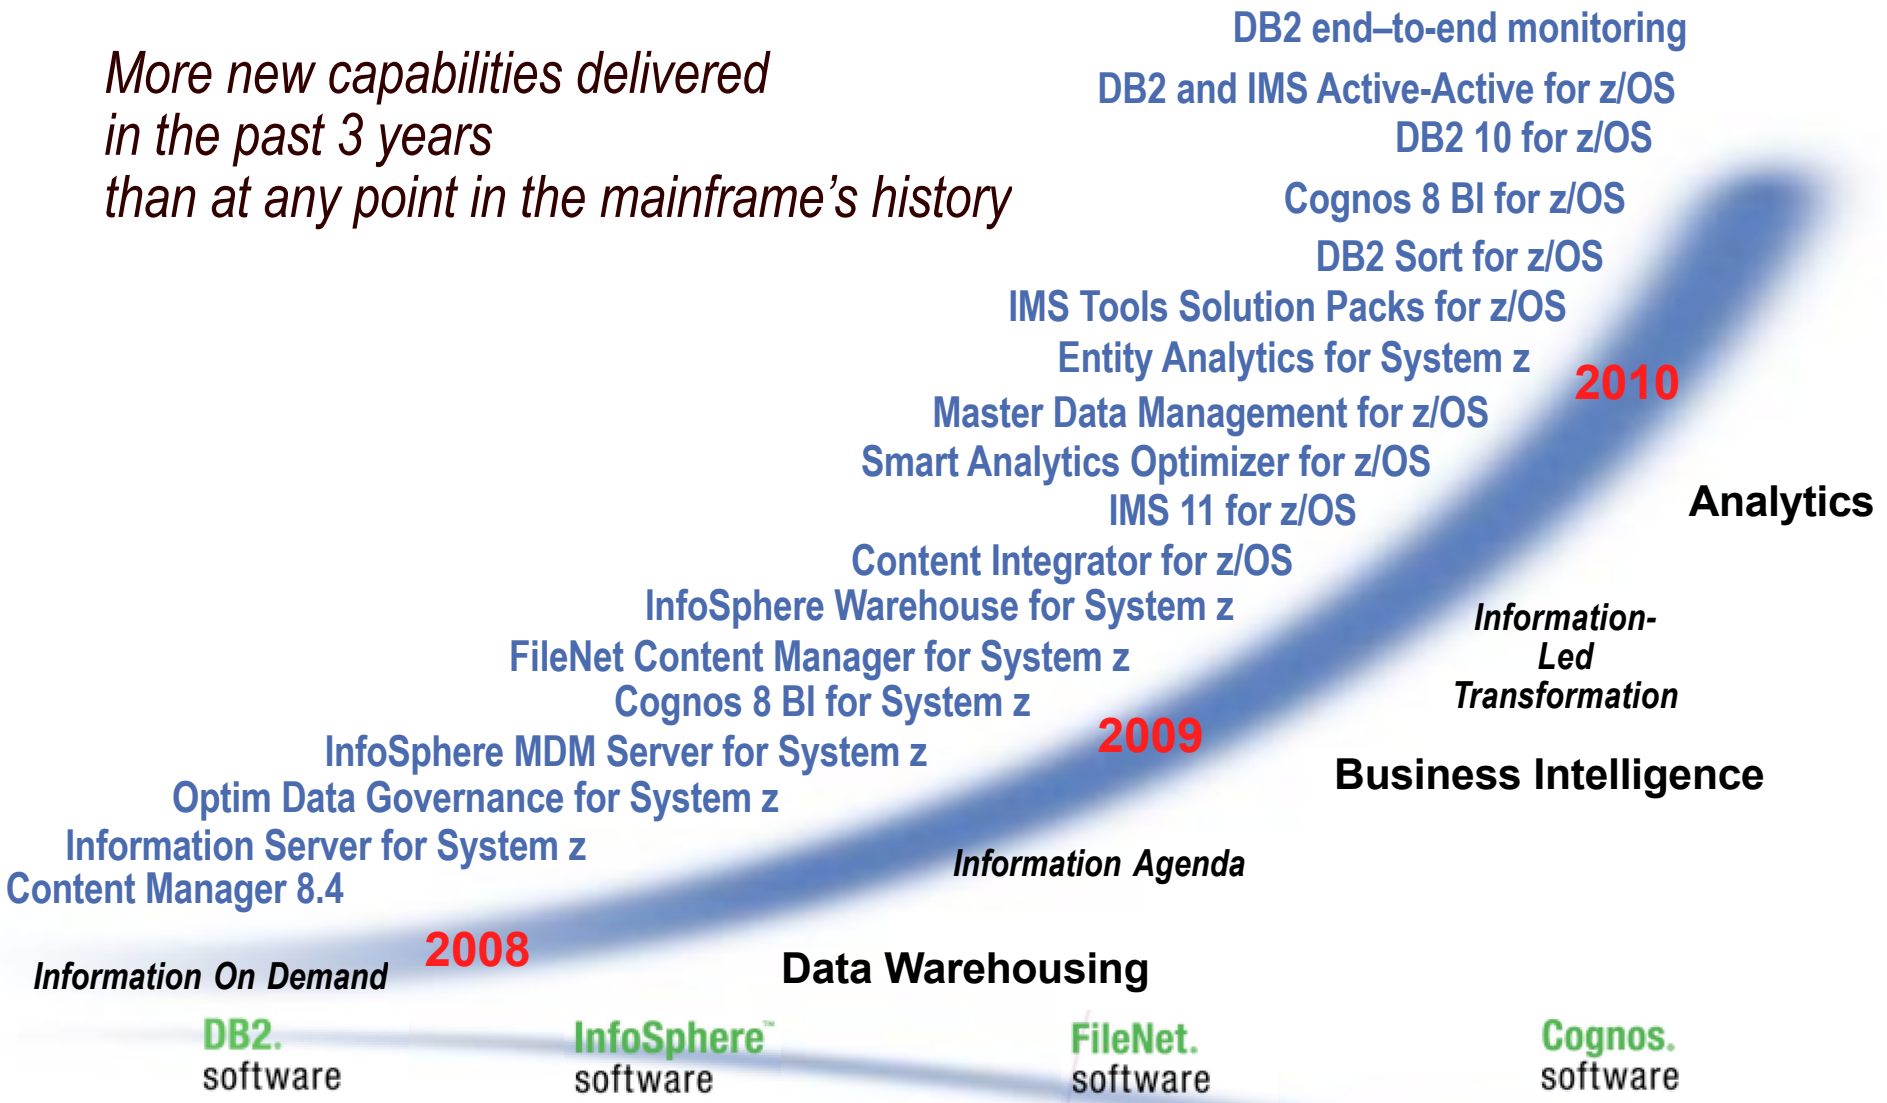

DB2 V9 –Rich, Features Filled Release: DW a Major Theme

- SHRLEVEL(REFERENCE) for REORG of LOB tablespaces
- Online RENAME COLUMN
- Online RENAME INDEX
- Online CHECK DATA and CHECK LOB
- Online REBUILD INDEX
- Online ALTER COLUMN DEFAULT
- More online REORG by eliminating BUILD2 phase
- Faster REORG by intra-REORG parallelism
- Renaming SCHEMA, VCAT, OWNER, CREATOR
- LOB Locks reduction
- Skipping locked rows option
- Tape support for BACKUP and RESTORE SYSTEM utilities
- Recovery of individual tablespaces and indexes from volume-level backups
- Enhanced STOGROUP definition
- Conditional restart enhancements
- Histogram Statistics collection and exploitation
- WS II OmniFind based text search
- DB2 Trace enhancements
- WLM-assisted Buffer Pools management
- Global query optimization
- Dynamic ANDing
- Generalizing sparse index and in-memory data caching method
- Optimization Service Center
- Autonomic reoptimization
- Logging enhancements
- LOBs network flow optimization
- Faster operations for variable-length rows
- NOT LOGGED tablespaces
- Index on expressions
- Universal Tablespaces
- Partition-by-growth tablespaces
- APPEND option at insert
- Autonomic index page split
- Different index page sizes
- Support for optimistic locking
- Faster and more automatic DB2 restart
- RLF improvements for remote application servers such as SAP
- Preserving consistency when recovering individual objects to a prior point in time
- CLONE Table: fast replacement of one table with another
- Index compression
- Index key randomization
- DECIMAL FLOAT
- BIGINT
- VARBINARY, BINARY
- TRUNCATE TABLE statement
- MERGE statement
- FETCH CONTINUE
- ORDER BY and FETCH FIRST n ROWS in sub-select and full-select
- ORDER OF extension to ORDER BY
- INTERSECT and EXCEPT Set Operations
- RANK, DENSE
- Instead of triggers
- Various scalar and built-in functions
- Cultural sort
- LOB File Reference support
- XML support in DB2 engine
- Enhancements to SQL Stored Procedures
- SELECT FROM UPDATE/DELETE/MERGE
- Enhanced CURRENT SCHEMA
- IP V6 support
- Trusted Context
- Database ROLEs
- Automatic creation of database objects
- Temporary space consolidation
- zIIP enabled SQL Procedures
- . . .

Ultimate Consolidation Opportunity

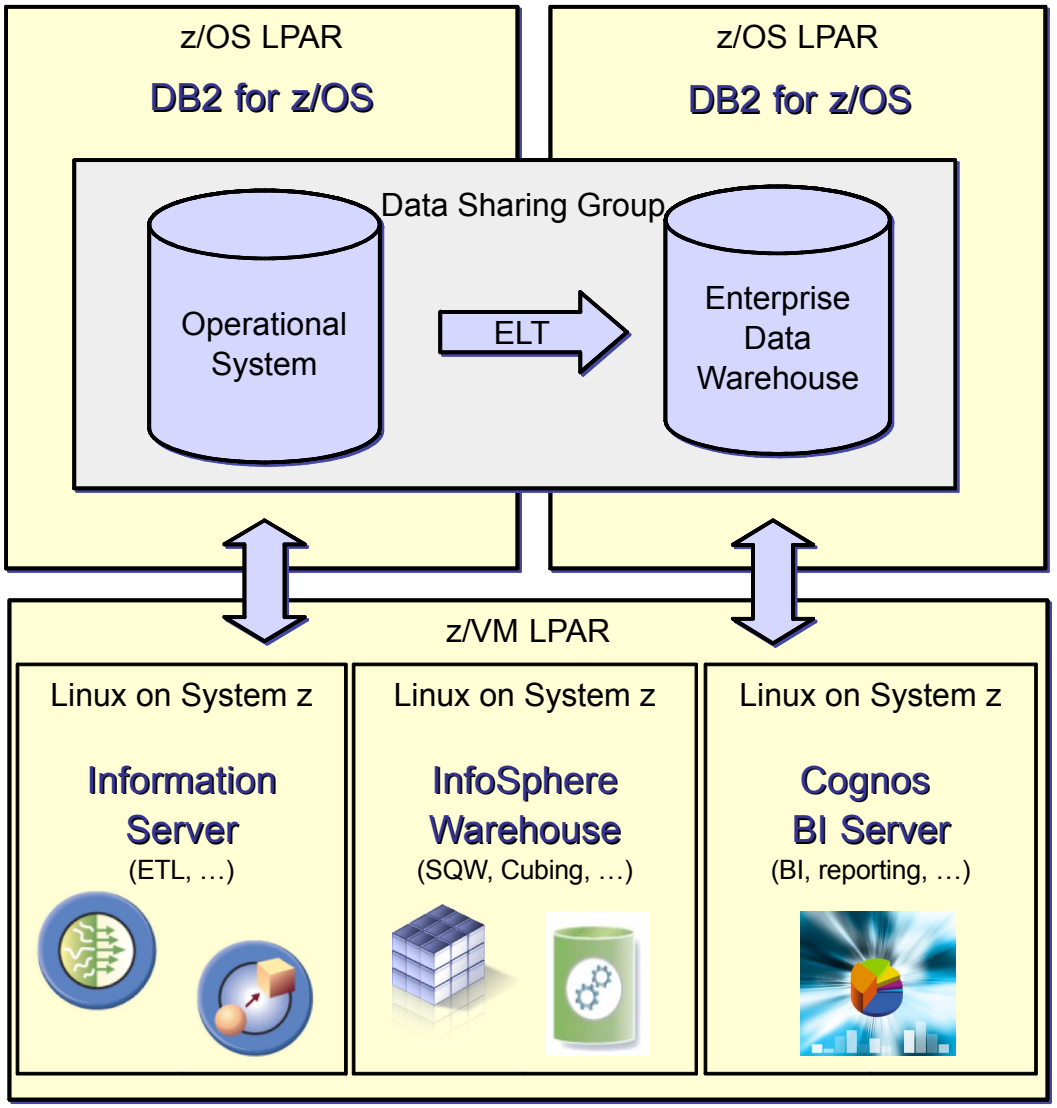

- Consolidation of mission-critical data on System z
- Leveraging existing environment, high availability, backup and governance procedures as well as skills
- Efficient data movement within a data sharing group (no network)
- Performance and TCO improvements through cubing services (data marts) and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

Leverage Data Proximity on System z

Disadvantage of "OLTP/DW Separation" approach

- Extract data / load data more expensive
- Access to OLTP data only through OLTP node (might affect operational applications work)

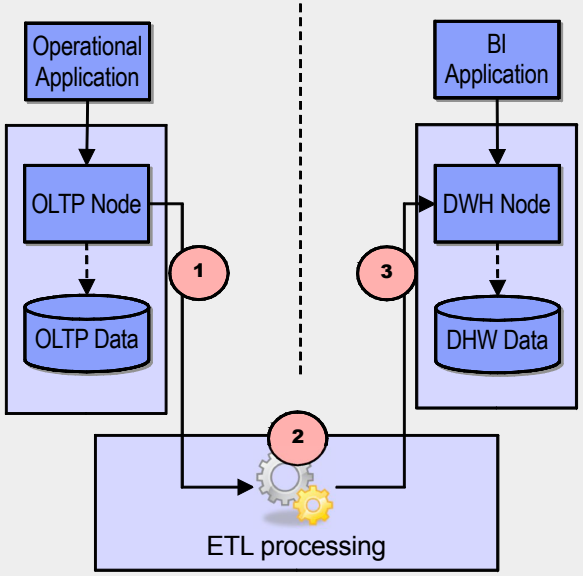

DB2 for z/OS with Data Sharing Data

- DWH node can efficiently access OLTP data
- Transformation and data movement in same data sharing group
- OLTP workload not impacted
- Effective resource management through WLM or IRD

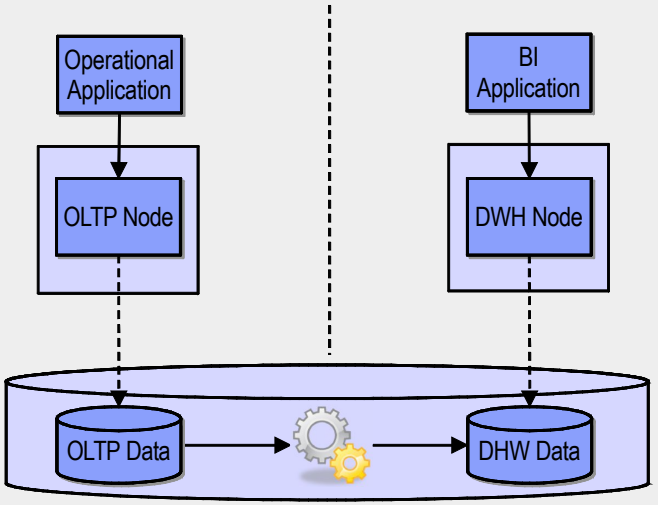

- Building a data warehouse requires access to operational data, transformations and finally loading into a data warehouse database
- With DB2 for z/OS, data warehouse tables may be accessed in the *same* data sharing group where the operational data already reside
- This setup is
  - more efficient
  - more secure (data is not leaving the platform)
  - easier to handle

Data Movement

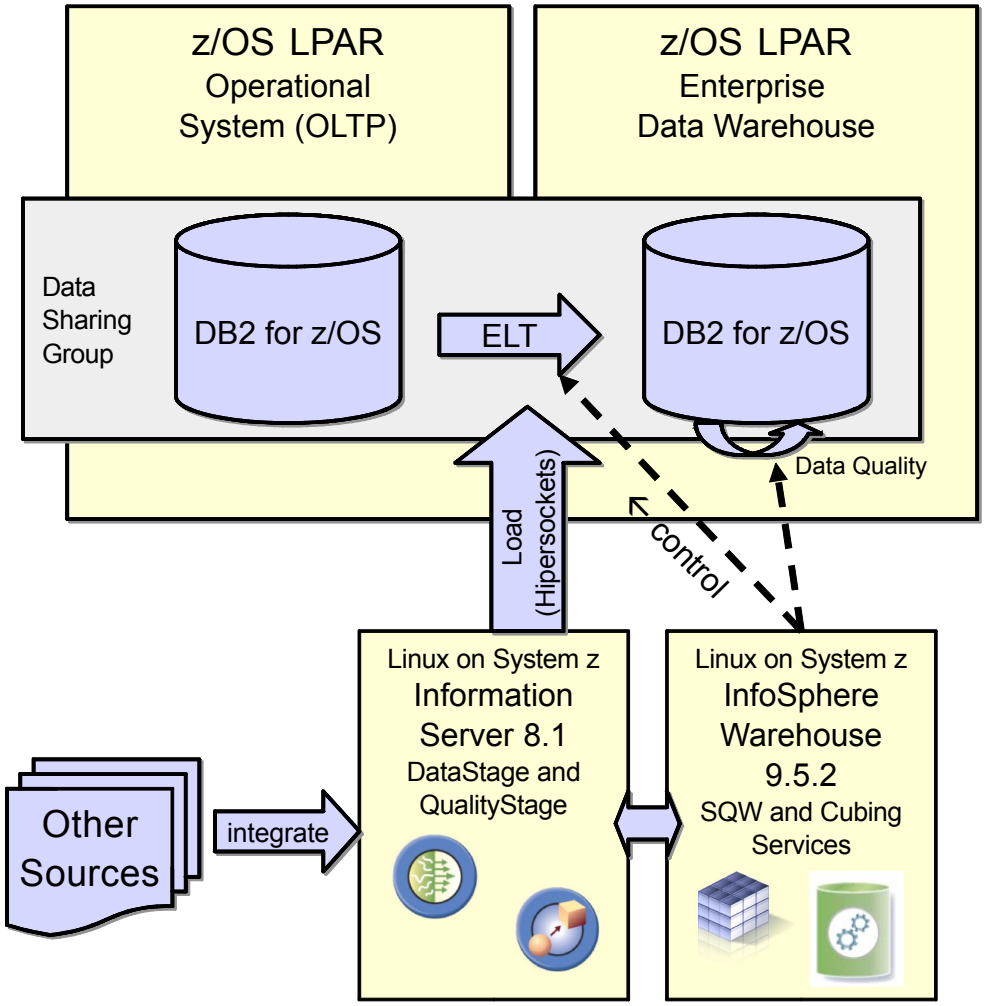

- Aggregation and data mart updates are triggered from Linux on System z and result in efficient data movement on DB2 for z/OS (no network) using combination of the SQW engine in InfoSphere Warehouse and Information Server / DataStage
- SQL triggered from Linux on System z is at least 50% **zIIP offloadable** (reduced costs)
- Information Server and InfoSphere products can integrate other relational and non-relational sources and efficiently load data into DB2 for z/OS through Hipersockets and load/pipes

Ultimate Warehouse Query Performance

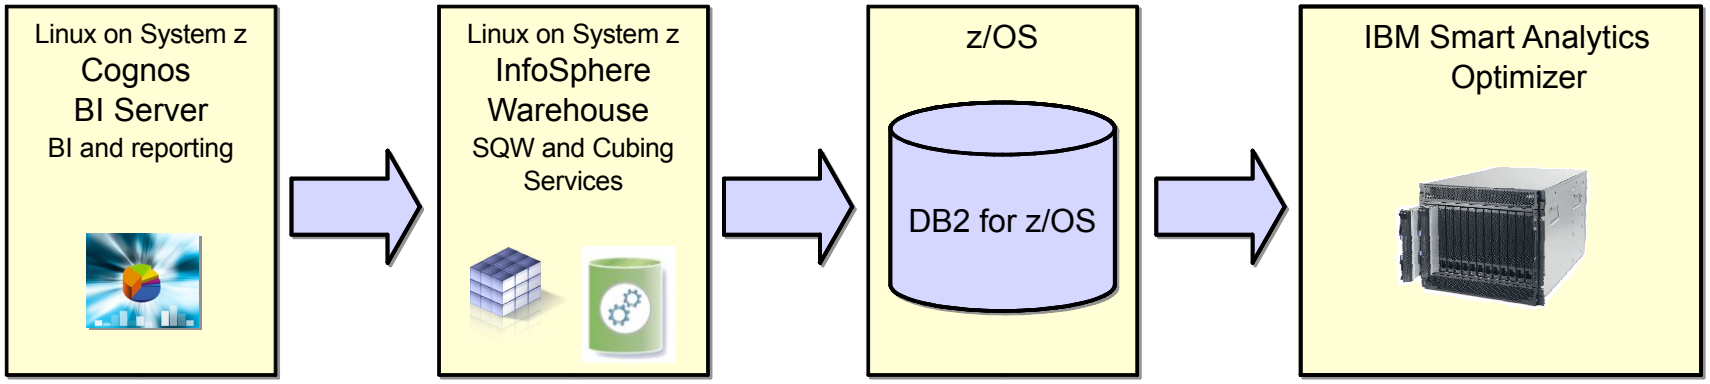

- Cognos BI reports and ad-hoc analysis result in SQL or MDX (OLAP) requests
- Optimized SQL for DB2 for z/OS and integration with InfoSphere Warehouse cubing services

- InfoSphere Warehouse cubing services provides scalable and in-memory cube model for MDX (OLAP) requests
- Additional MQTs (based on advisor) may be defined to support cube access
- Efficient access to DB2 for z/OS through Hipersockets

- Parallel query support in DB2 for z/OS
- WLM ensures that longer running queries are not jeopardizing short running queries

- Multi-core processor (Blade) and in-memory technology accelerates warehouse queries by orders of magnitude
- Scans and aggregation on highly compressed data
- Fast query response time with significantly reduced need for tuning, index creation etc.

System z & DB2: ideal data serving platform

- Data consolidation helps reduce:
  - Multiple copies
  - Cost and complexity
  - Network
  - Storage
  - Administration
  - Risk
- Leverage System z technology
  - Parallel Sysplex clustering for scalability, availability and performance
  - Data compression for TCO
  - Centralized policies
- IBM specialty engines lower TCO for data serving

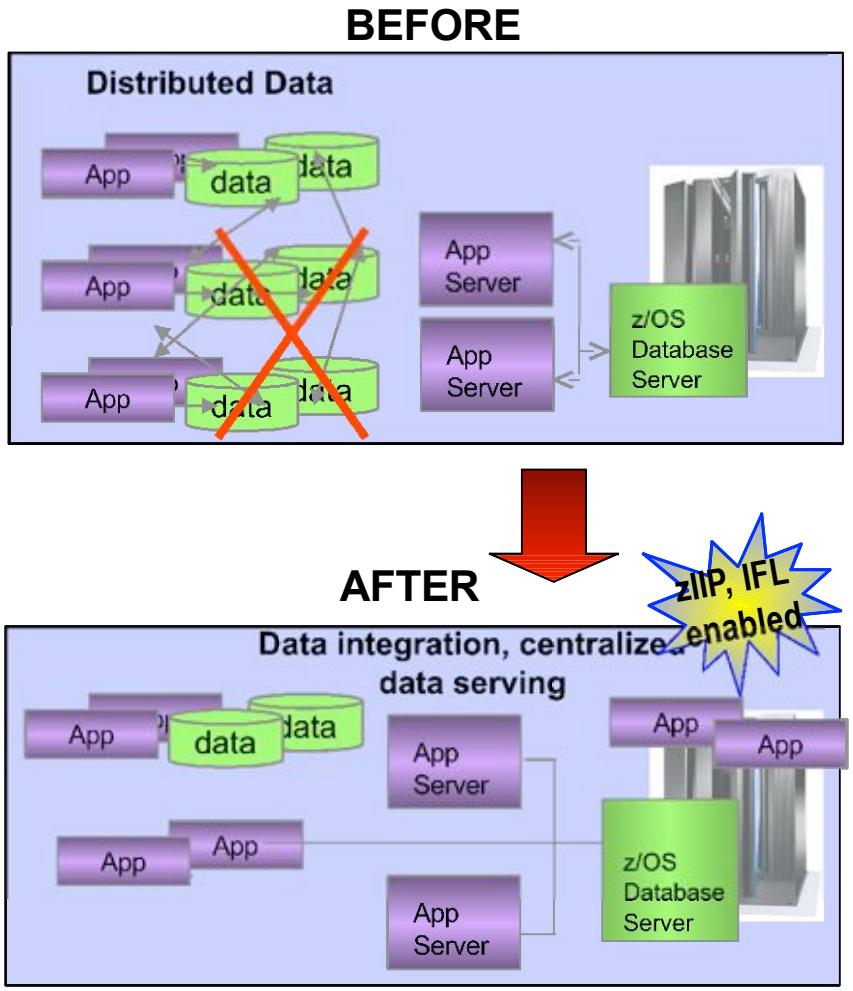

Business Challenges and Technology Trends

- Workloads increasingly ad-hoc and unpredictable
- Traditional performance tuning tools of the trade such as indexing, prebuilt aggregates and MQTs struggling to keep the pace
  - Require top DBA expertise and sophisticated tools
  - Even then not good enough
- Technology trends
  - Very large number of processor sockets and cores
  - Massive amounts of real memory
  - Specialized physical data designs: row-store vs. column-store

IBM Smart Analytics Optimizer
Adding Industry Leading Performance

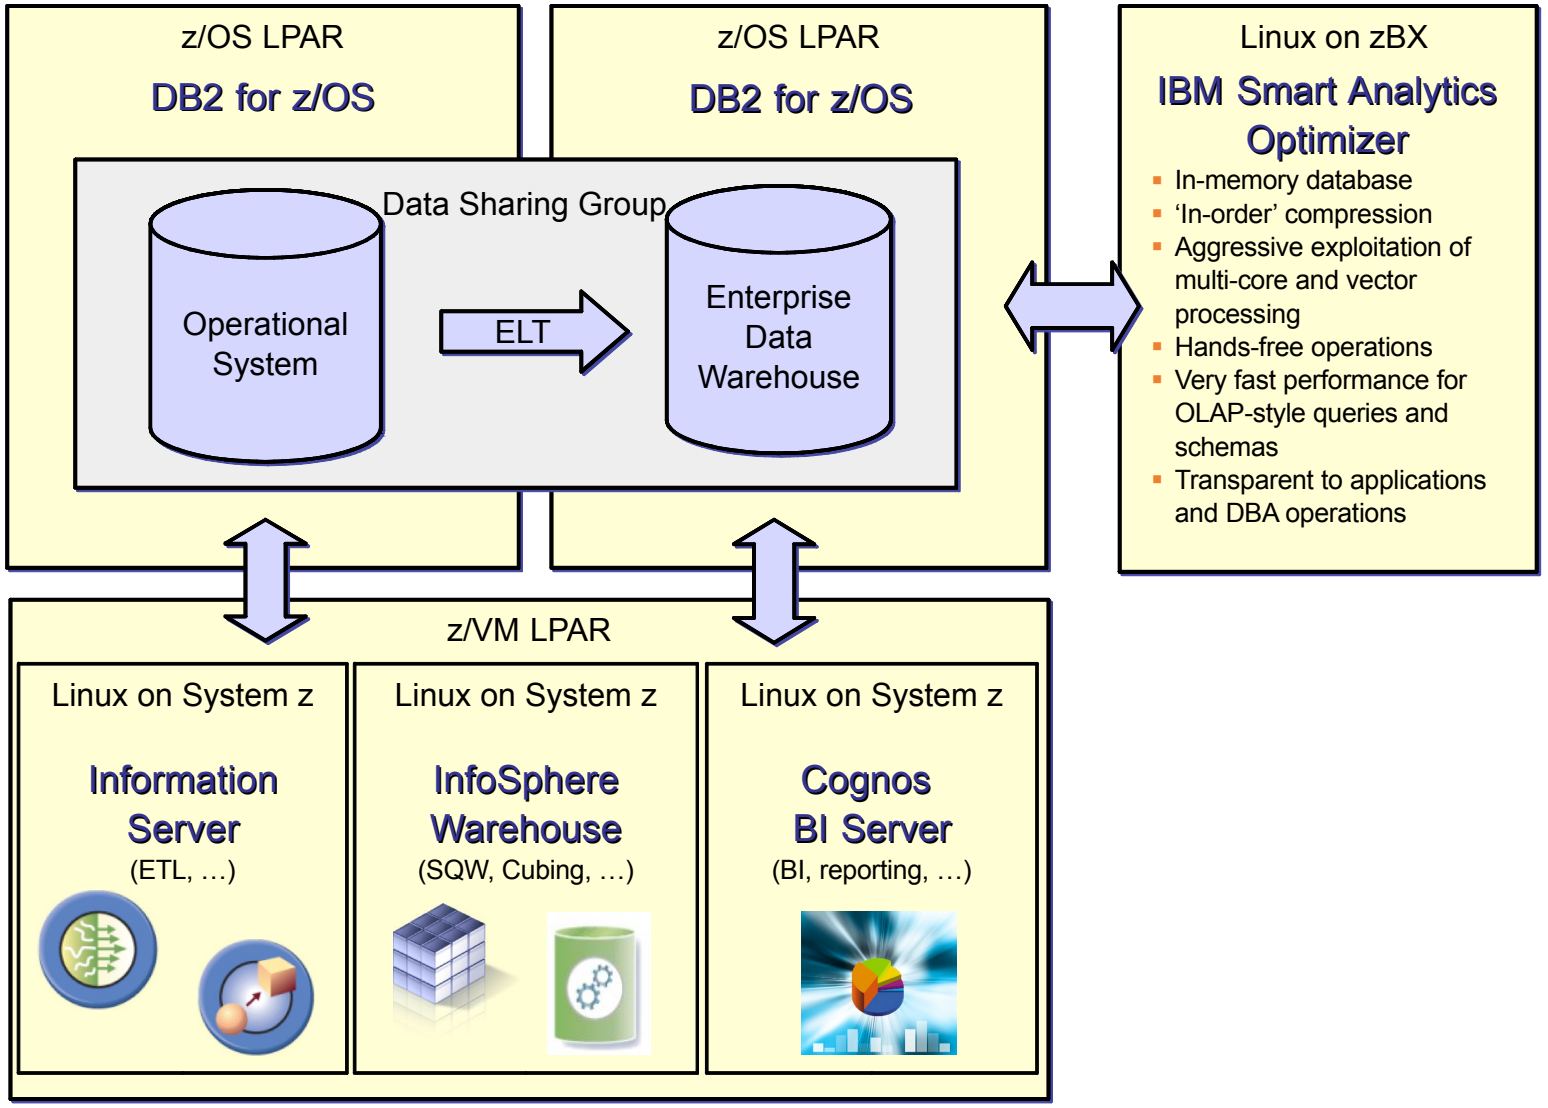

Agenda

- System z as the Information Management Hub
  - Business and Technology Drivers
  - Data Warehousing Components and Offerings
- IBM Smart Analytics Optimizer
  - Key Design and Operational Features
  - Supported workloads

Deep DB2 Integration within zHybrid Architecture

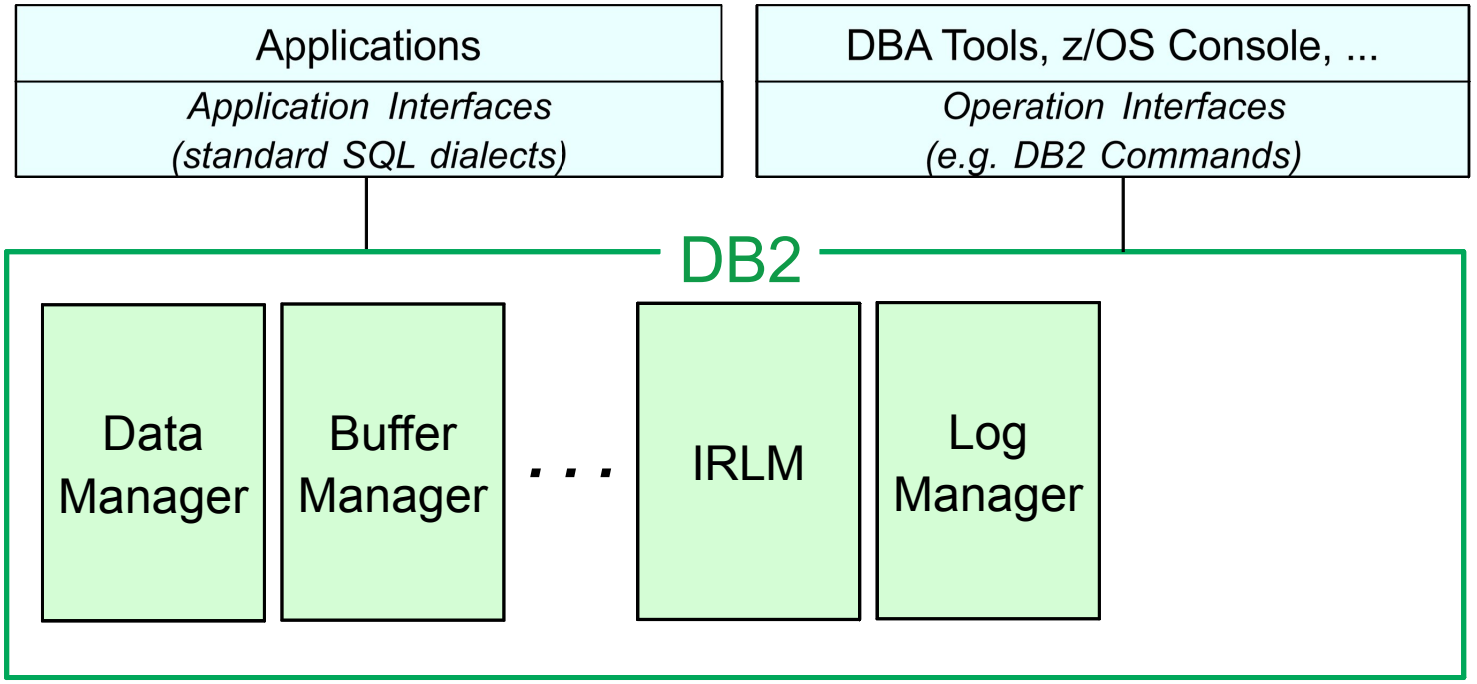

Deep DB2 Integration within zHybrid Architecture
IBM Smart Analytics Optimizer as a Virtual DB2 Component

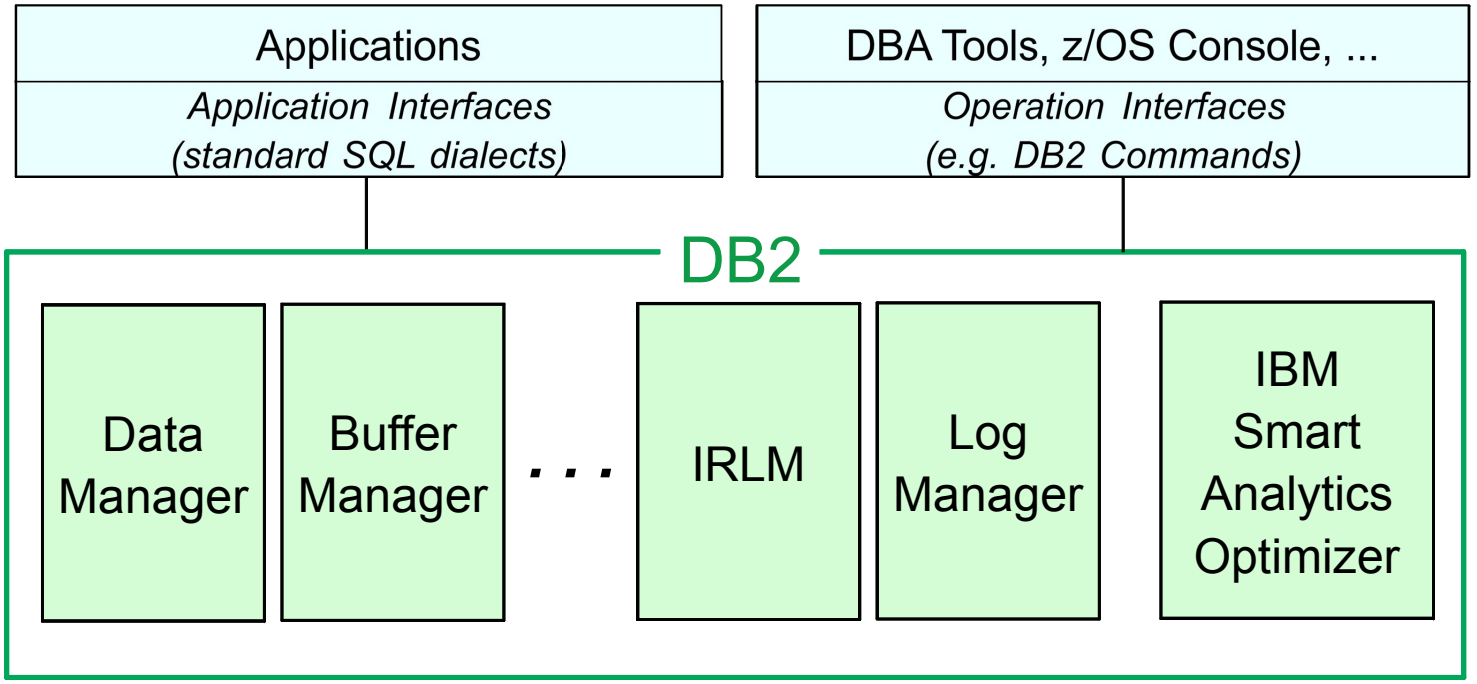

Deep DB2 Integration within zHybrid Architecture

IBM Smart Analytics Optimizer as a Virtual DB2 Component

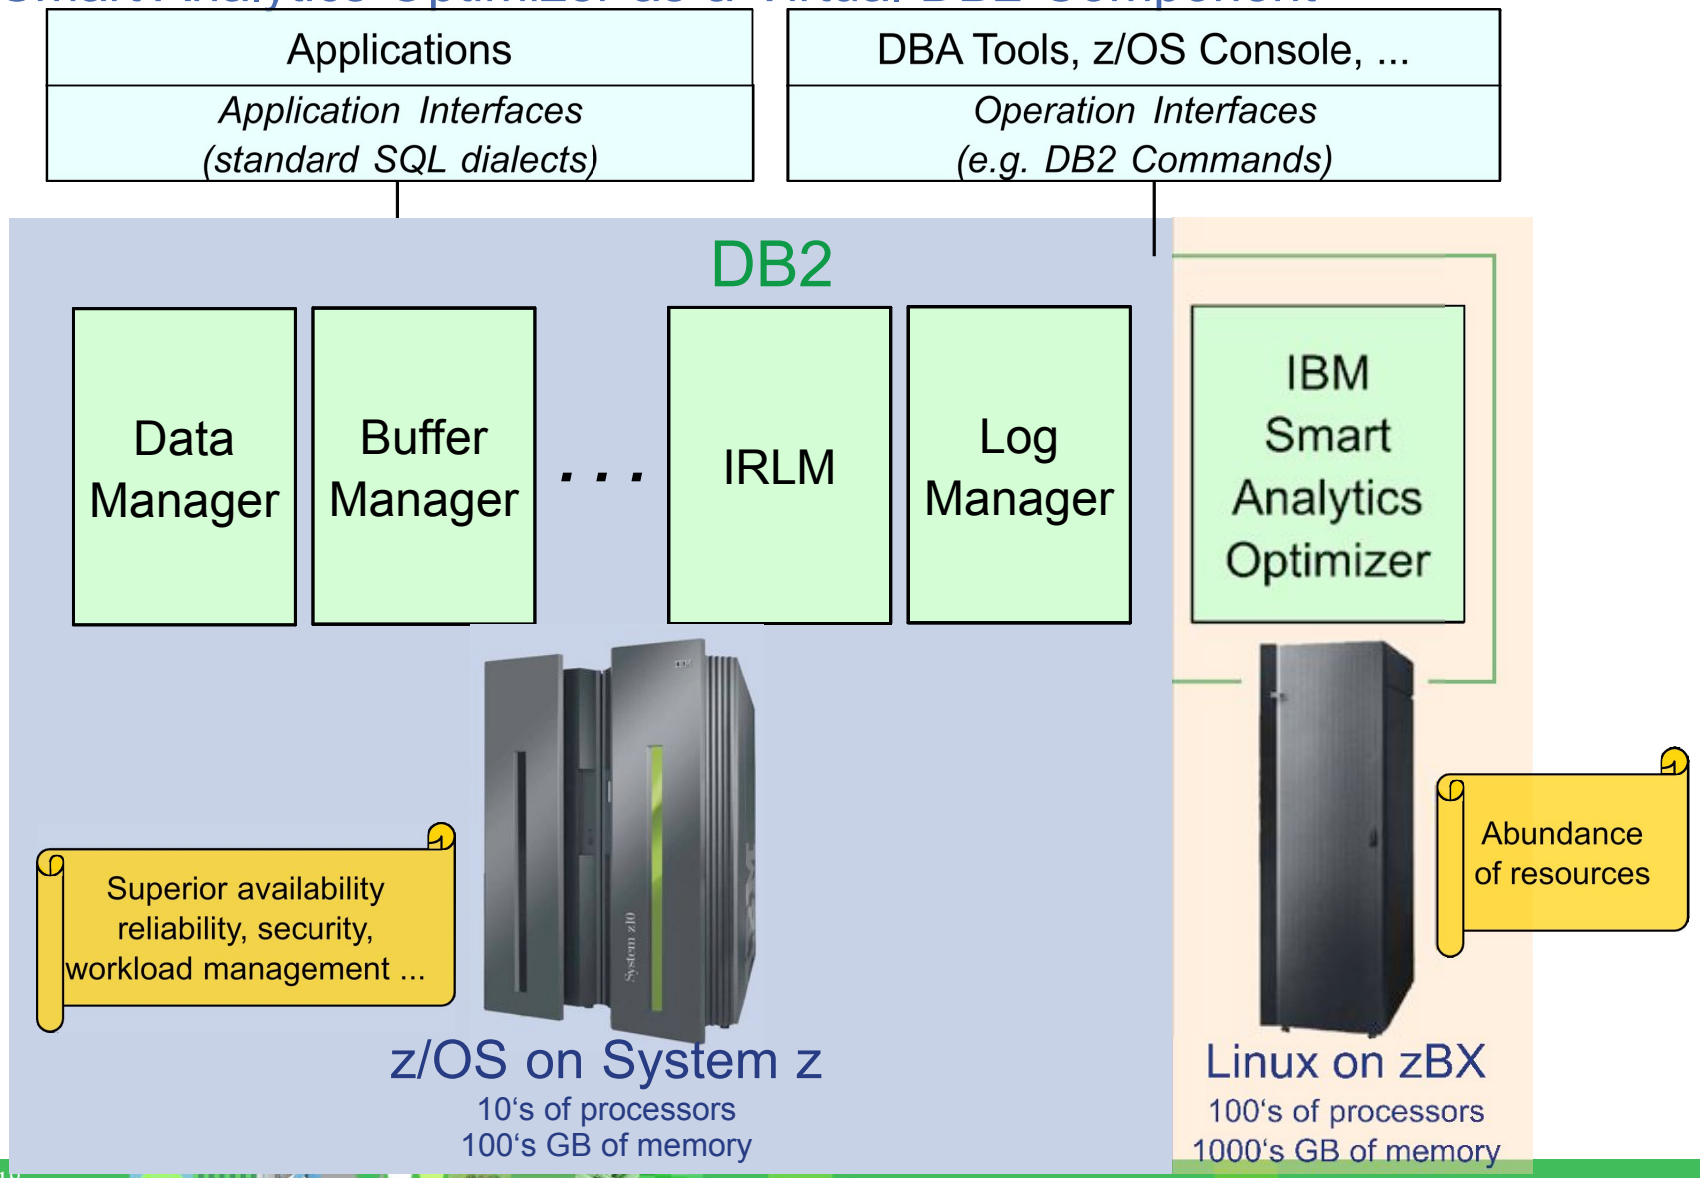

Enabling Technology – IBM Research Project BLINK

- Various Compression Techniques
  - Enables in-memory database
  - Order-preserving
  - Frequency partitioning
- Register-store: a combination of row- and column-based stores
- Multi-core friendly scans
  - Massive scale-out parallelism
  - Scans on compressed data
  - Vector processing
    Evaluation of all predicates in parallel
- Selective schema melting

OLTP vs. DW

- In a typical transactional workload, you normally fetch and use all attributes of a tuple. If you for example have a CUSTOMERS table, you wouldn‘t fetch the STREETNAME w/o also fetching the house number or ZIP code.
  - A transactional query is used to fetch few, very specific records of a relation.
- In typical Data Warehouse workloads, you tend to fetch only a small subset of each record.
  - The tables are usually very wide, having multiple measure columns.
  - Queries almost never touch all attributes of the tuples but only a small subset of the available attributes.
  - A query usually needs to evaluate/aggregate many tuples per relation.

Row-Store – Optimal Choice for OLTP

- In traditional DBMS, we use a **Row – Store** approach where each row is stored contiguously and where multiple rows are stored sequentially in I/O optimizerd data structures.
- If only few attributes are required, the complete row needs to be fetched and uncompressed.
- Lots of the data is moved and decompressed w/o even being used.

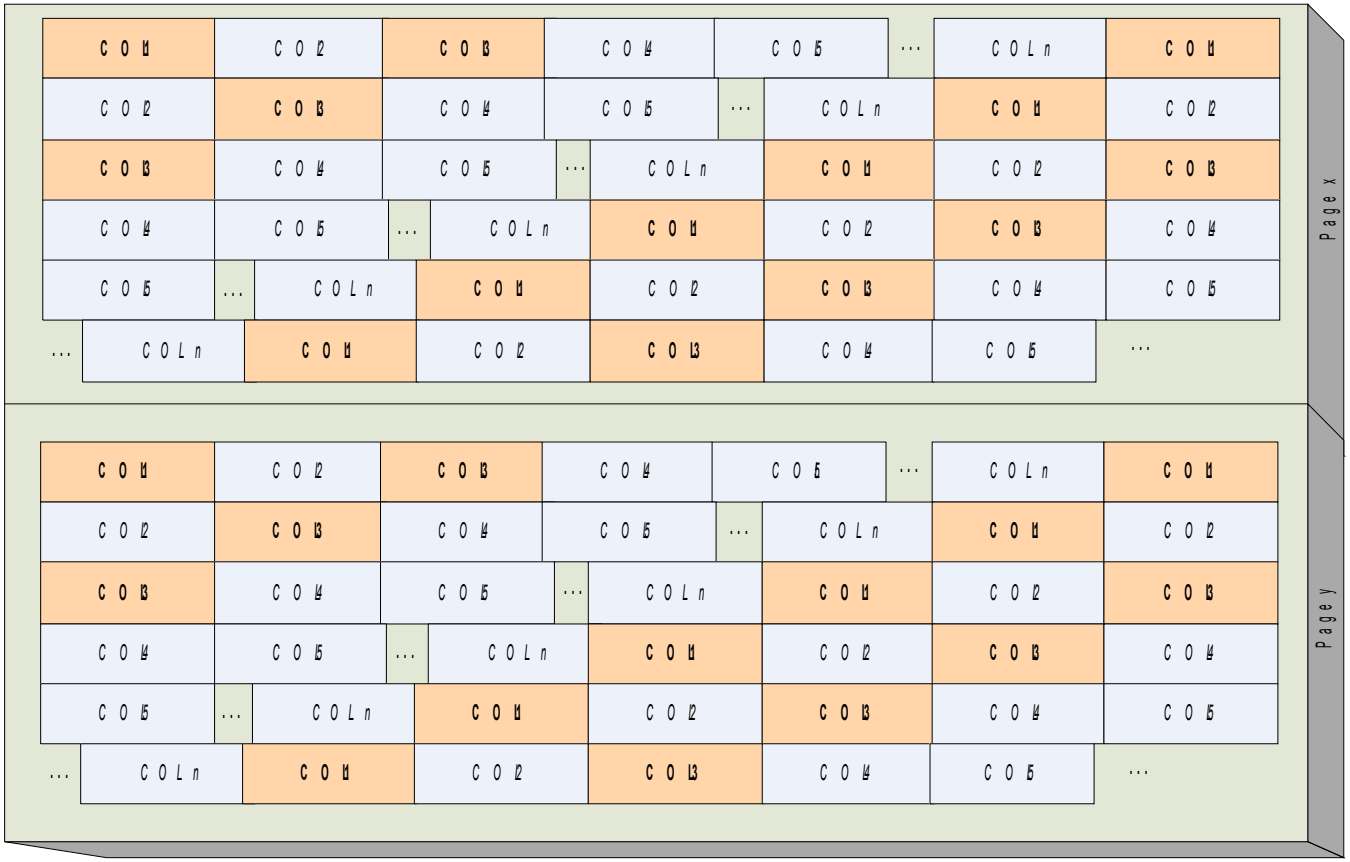

While a **Row – Store** is very efficient for transactional workloads, it is suboptimal for analytical workloads where only a subset of the attributes is needed!

Column-Store: Optimized for Certain DW Workloads

- Query Engines, which are optimized for analytical queries, sometimes use a **Column – Store** approach.
- In a **Column – Store**, the data of a specific column is stored sequentially before the data of the next column begins.
- If attributes are not required for a specific query execution, they simply can be skipped, not causing any I/O or decompression overhead.

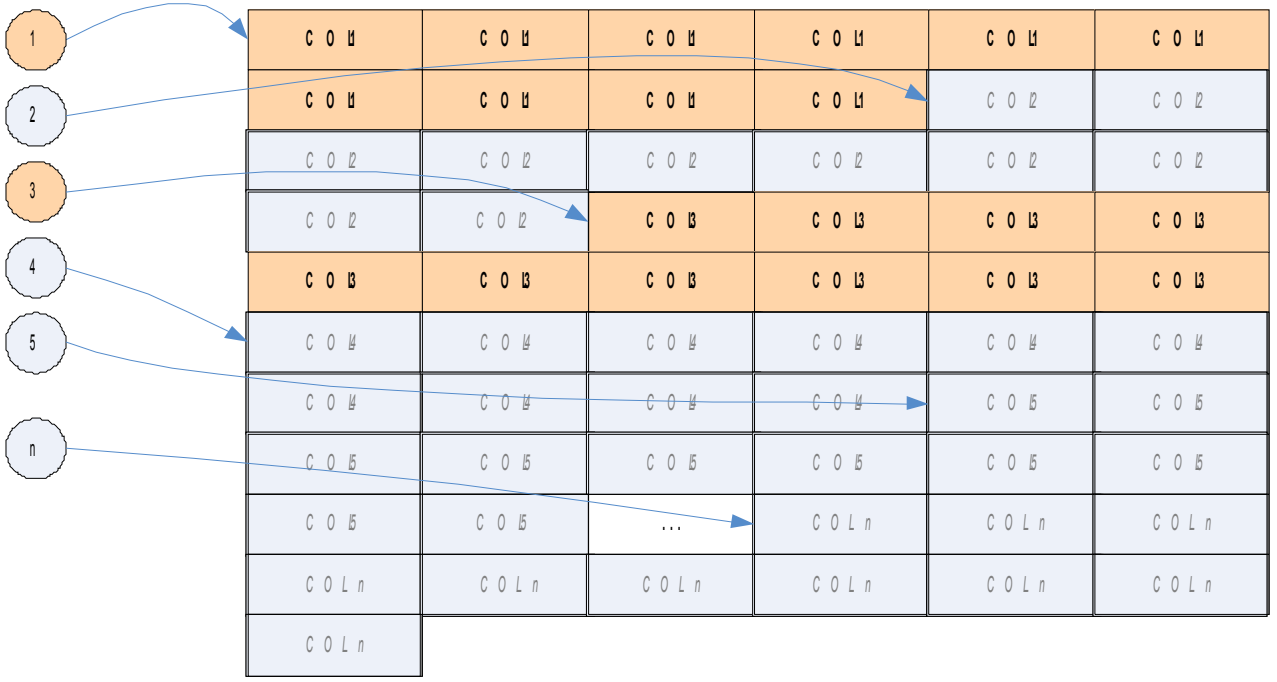

In a **Column – Store**, the data is also compressed sequentially for a column. This is an optimized approach if you plan to perform a sequential scan over your data. Random access to specific attributes in this store is not performing well.

This is normaly handled by limiting the number of tuples per column before the next column is stored. (The data is split into blocks.)

Register-Store

- Within a **Register – Store**, several columns are grouped together.
- The sum of the width of the compressed columns doesn‘t exceed a register compatible width. This could for example be 32 or 64 bit for a 64 bit system. It doesn‘t matter how many columns are placed within the register – wide data element.
- It is beneficial to place commonly used columns within the same register – wide data element. But this requires dynamic knowledge about the executed workload (runtime statistics).
- Having multiple columns within the same register – wide data element prevents ANDing of different results.

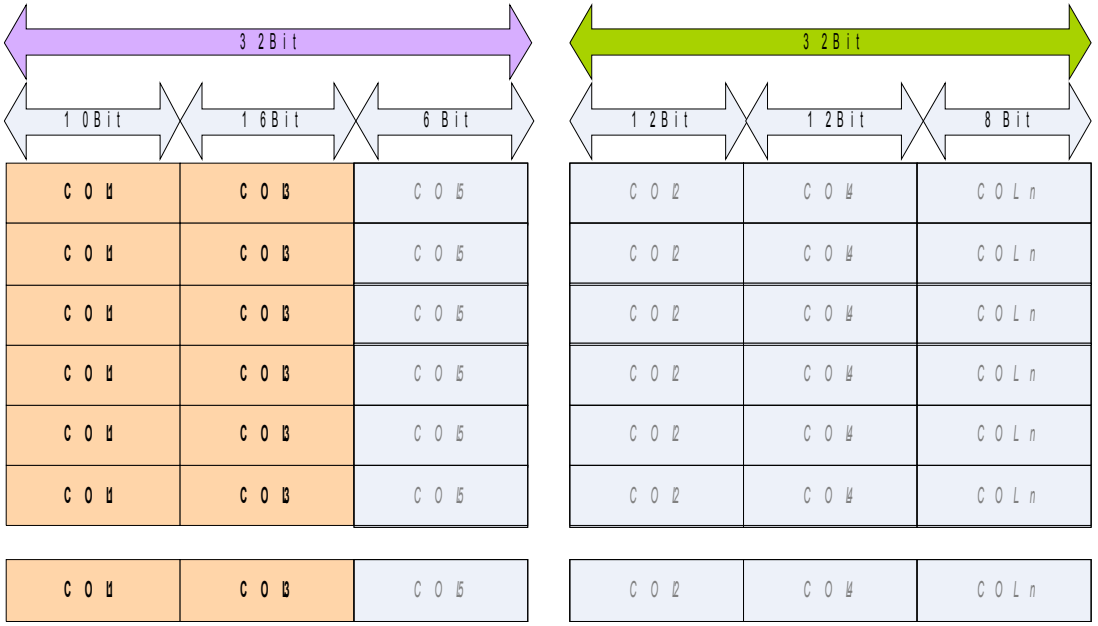

The **Register – Store** is an optimization of the Column – Store approach where we try to make the best use of existing hardware. Reshuffling small data elements at runtime into a register is time consuming and can be avoided. The **Register – Store** also delivers good vectorization capabilities.

Single Instruction Multiple Data Paradigm

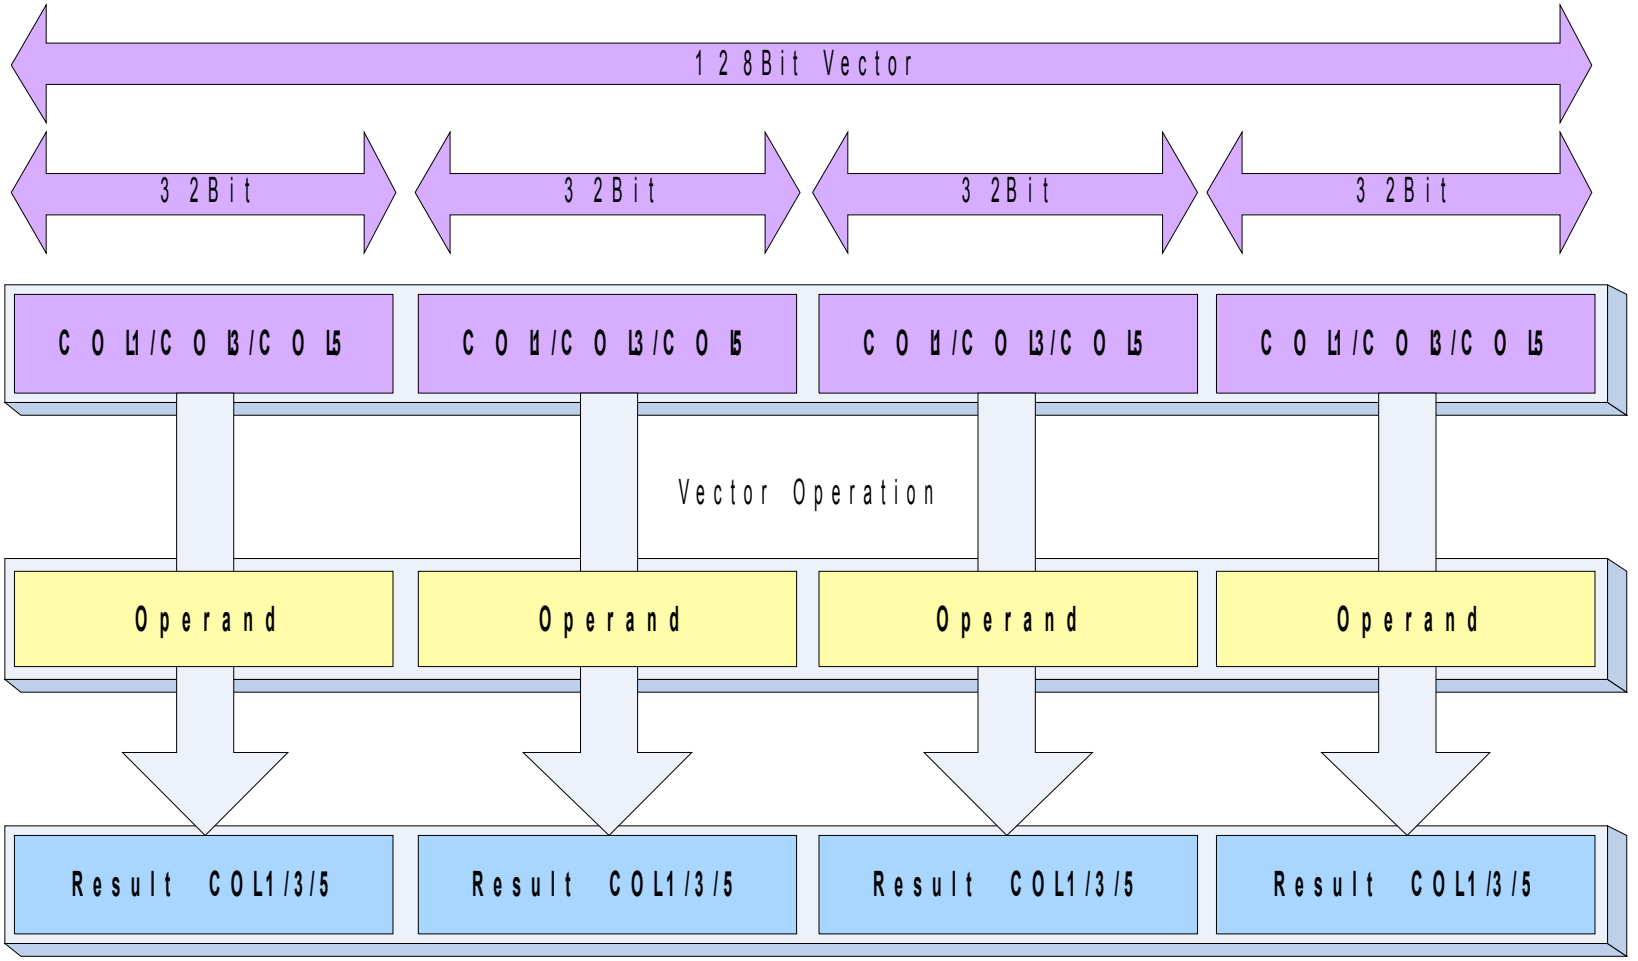

Data View

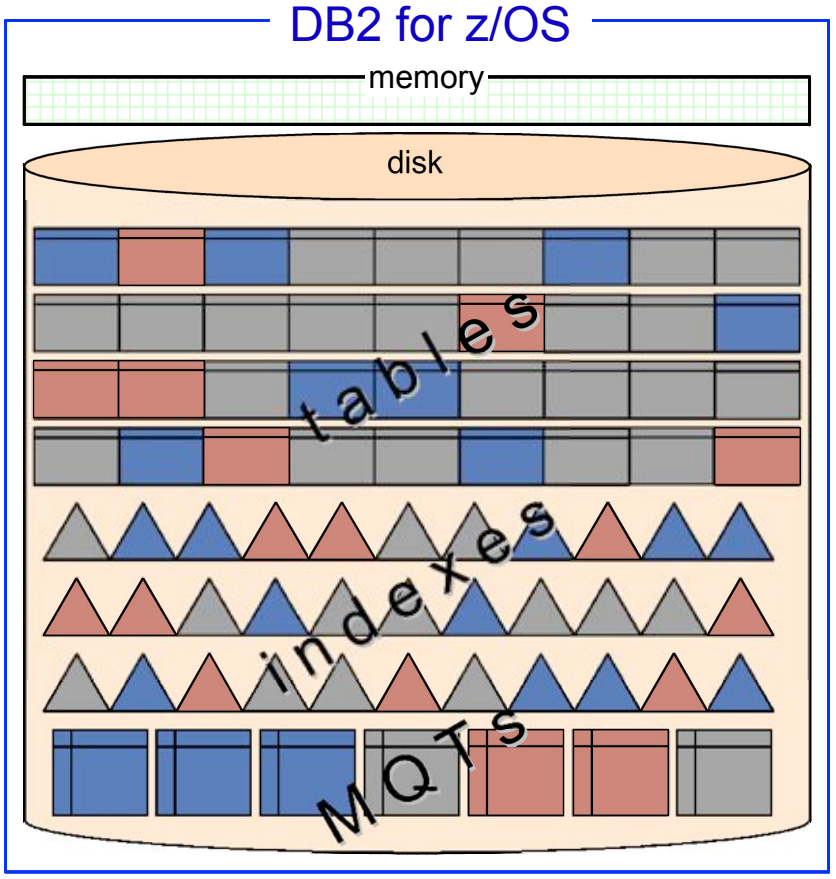

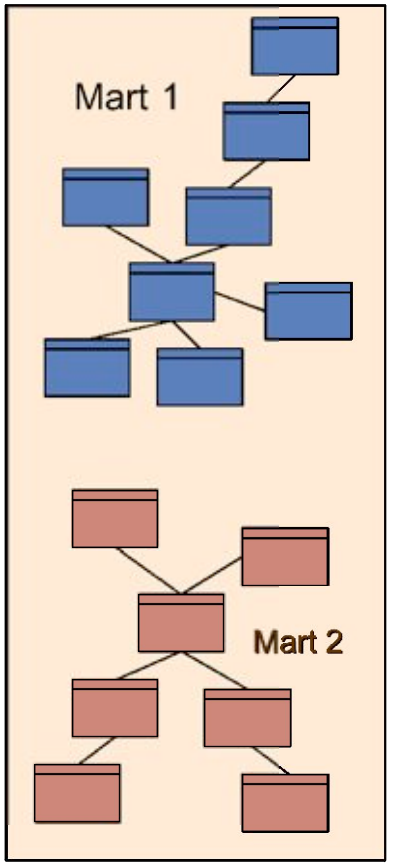

Data View: Marts – Redundant Sets of Memory Resident Tables

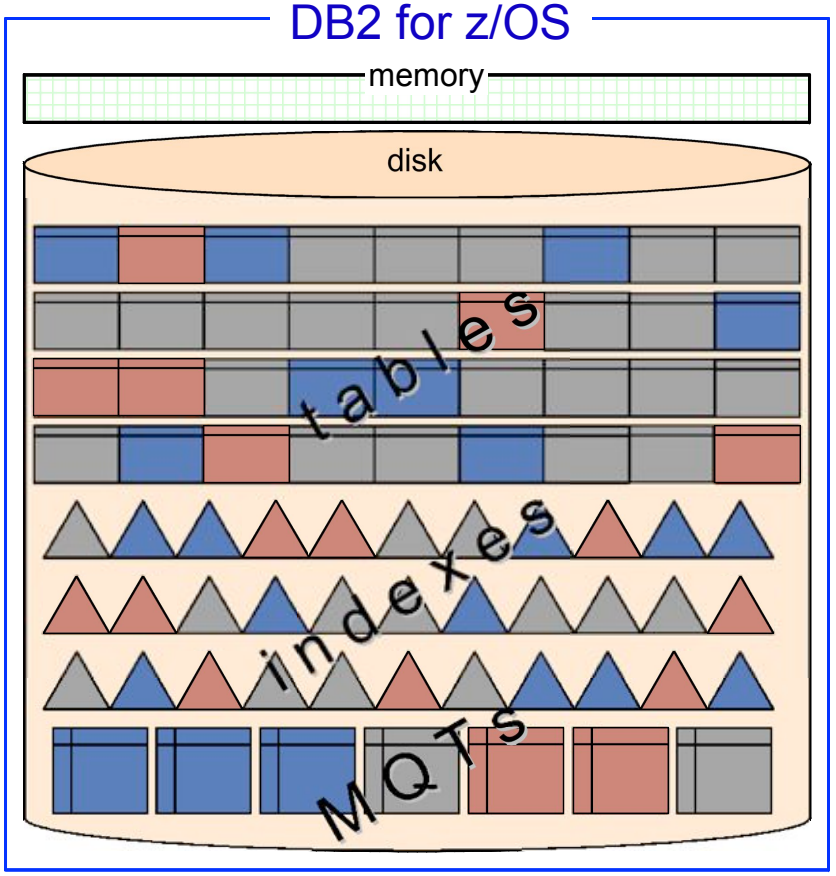

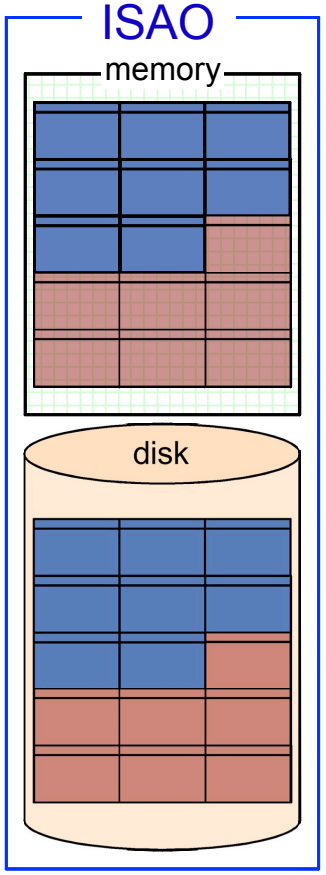

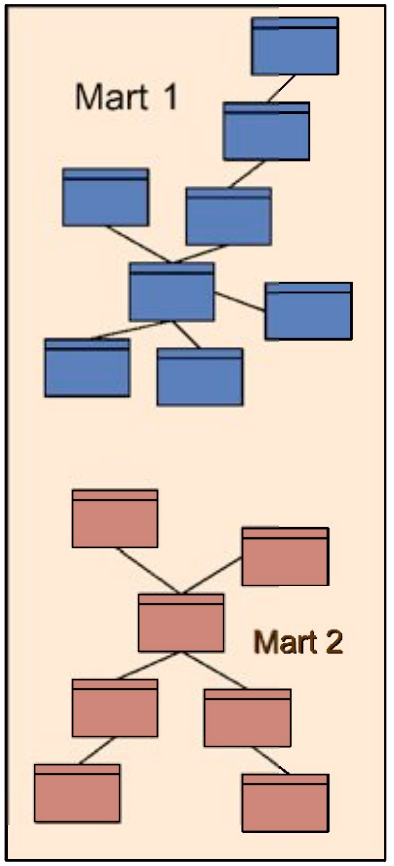

- DB2 continues to *own* and manage all data
- Access performance is influenced by traditional tuning mechanisms such as indexing, MQTs, aggregates, ...
- Typical usage: Enterprise Data Warehouse, large Data Marts

- ISAO contains fully memory resident, compressed copies of performance critical tables grouped into logically connected *marts*, bound ideally by star schema constraints
- Similar but much broader than MQTs: no column projections, no row restrictions, no row aggregations
- Typical usage: Data Marts, MQTs consolidation and replacement

ISAO Mart Definition and Deployment

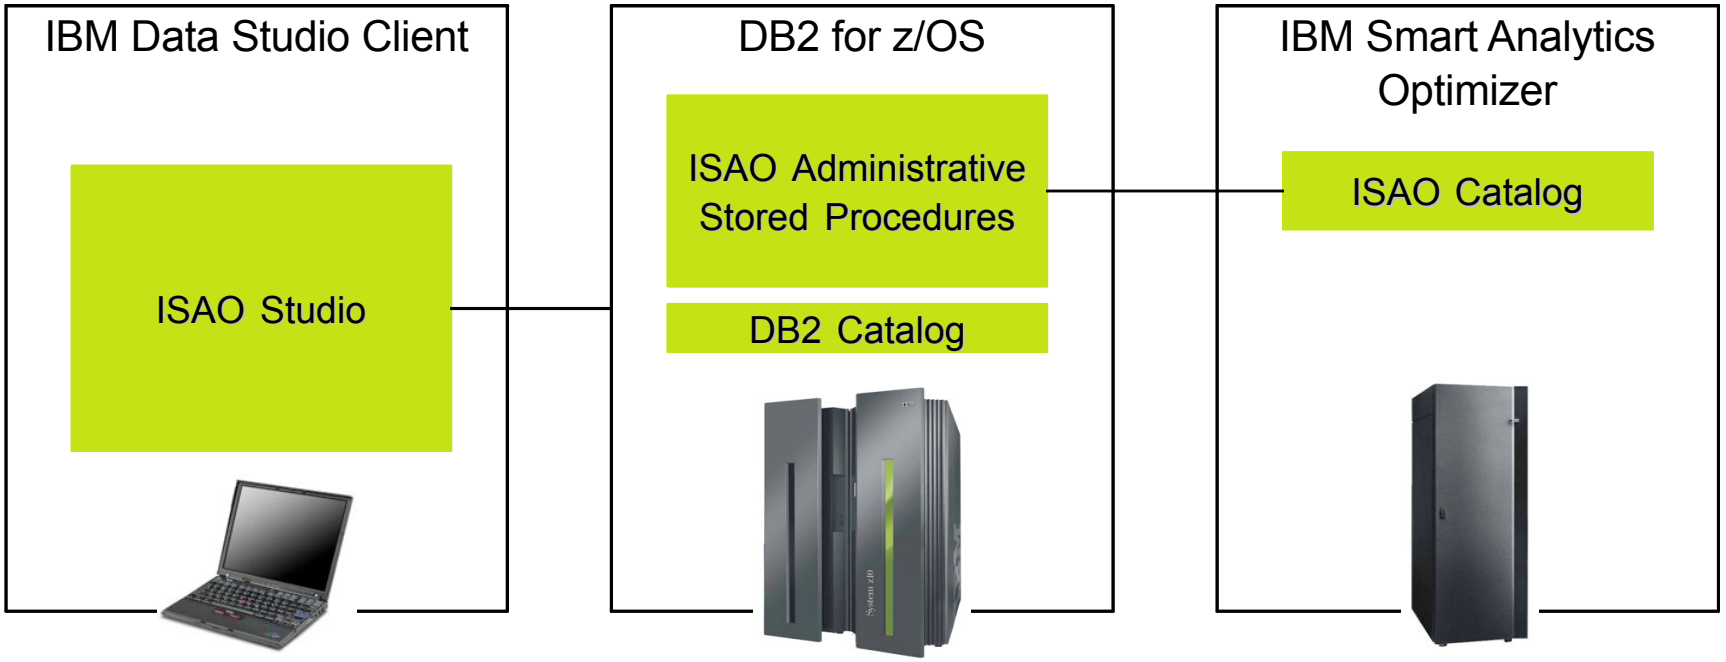

- ISAO marts need to be defined and deployed to ISAO before data is loaded and queries sent to ISAO for processing.
  - Definition: identifying tables and relations that make up marts.
  - Deployment: making marts known to DB2, i.e. storing mart meta data in the DB2 and ISAO catalog.
- ISAO Studio guides you through the process of defining and deploying marts, as well as invoking other administrative tasks.
- ISAO Stored Procedures implement and execute various administrative operations such as mart deployment, load and update, and serve as the primary administrative interface to ISAO from the outside world including ISAO Studio.

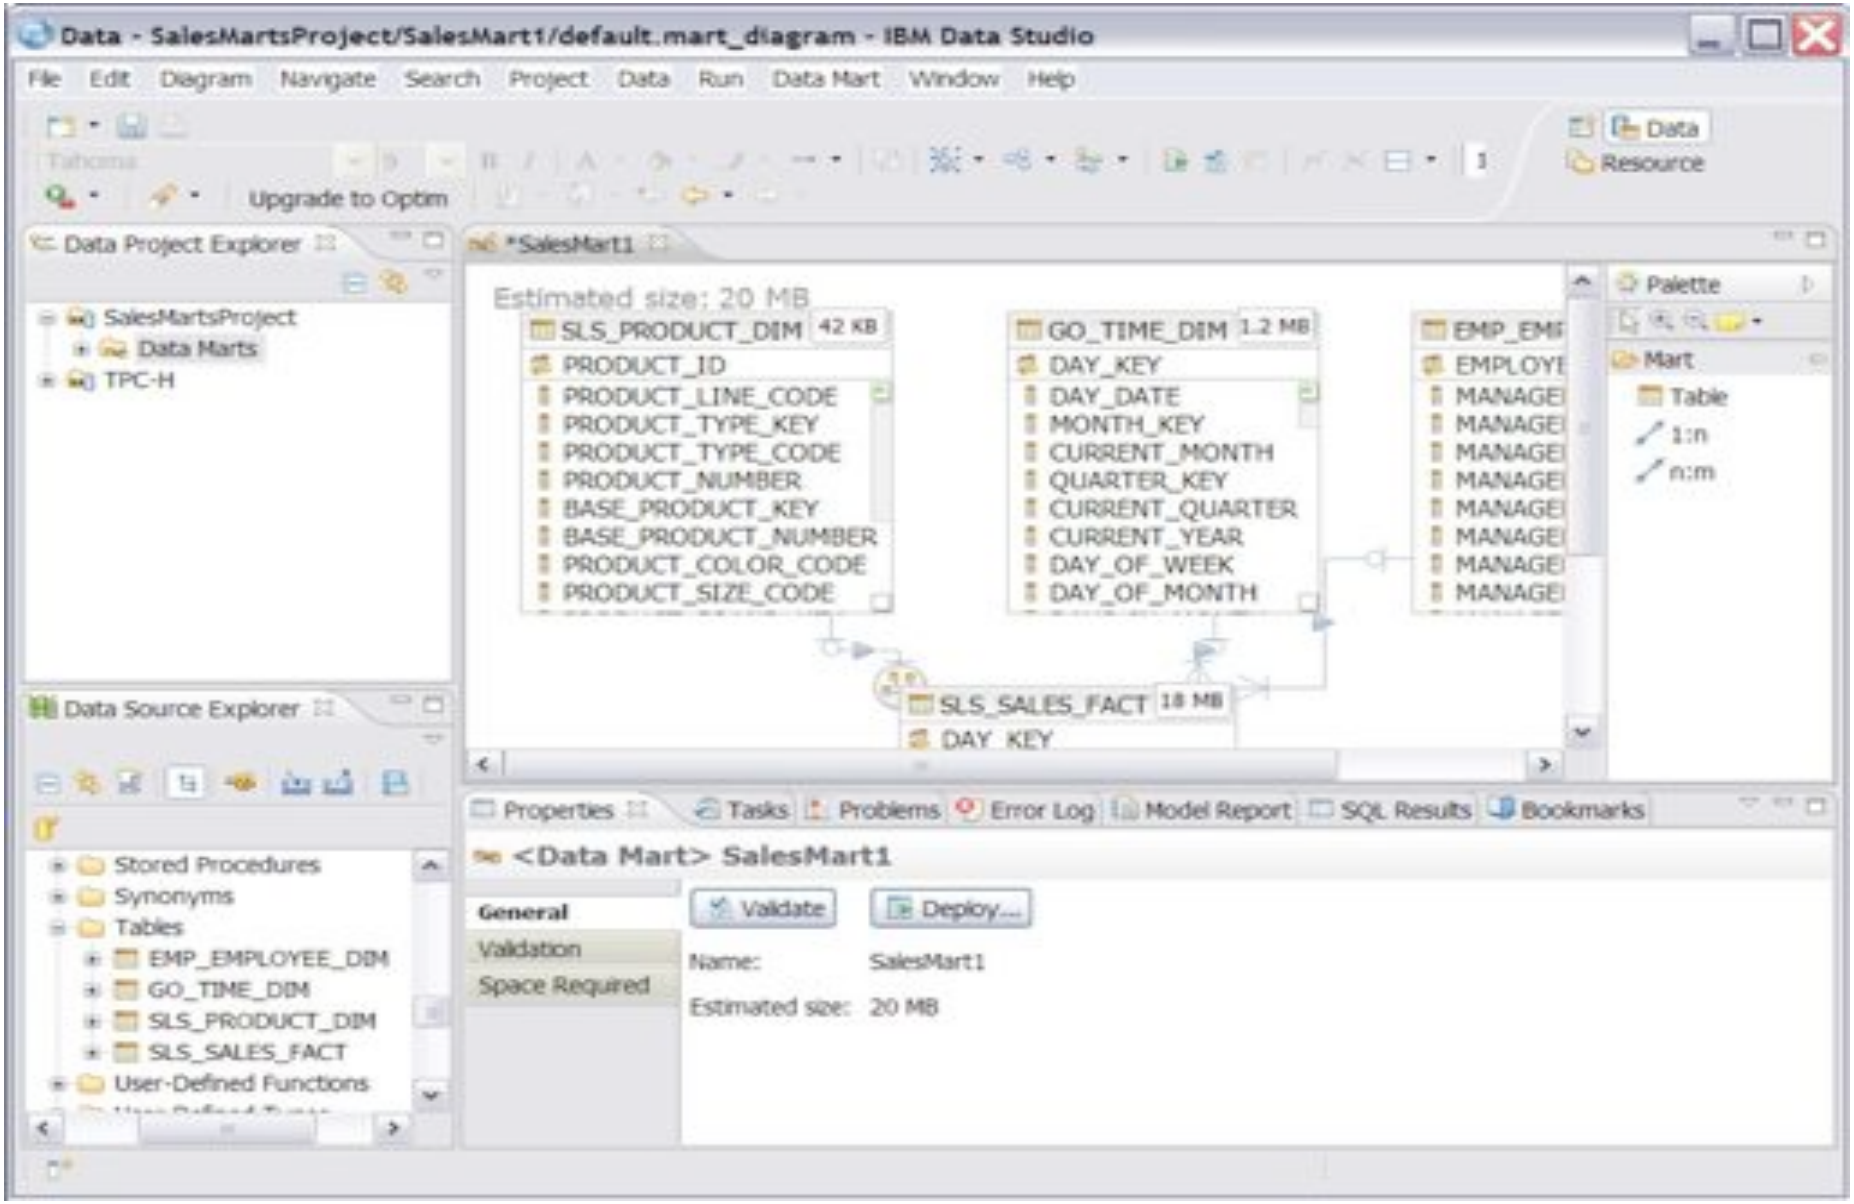

ISAO Mart Load

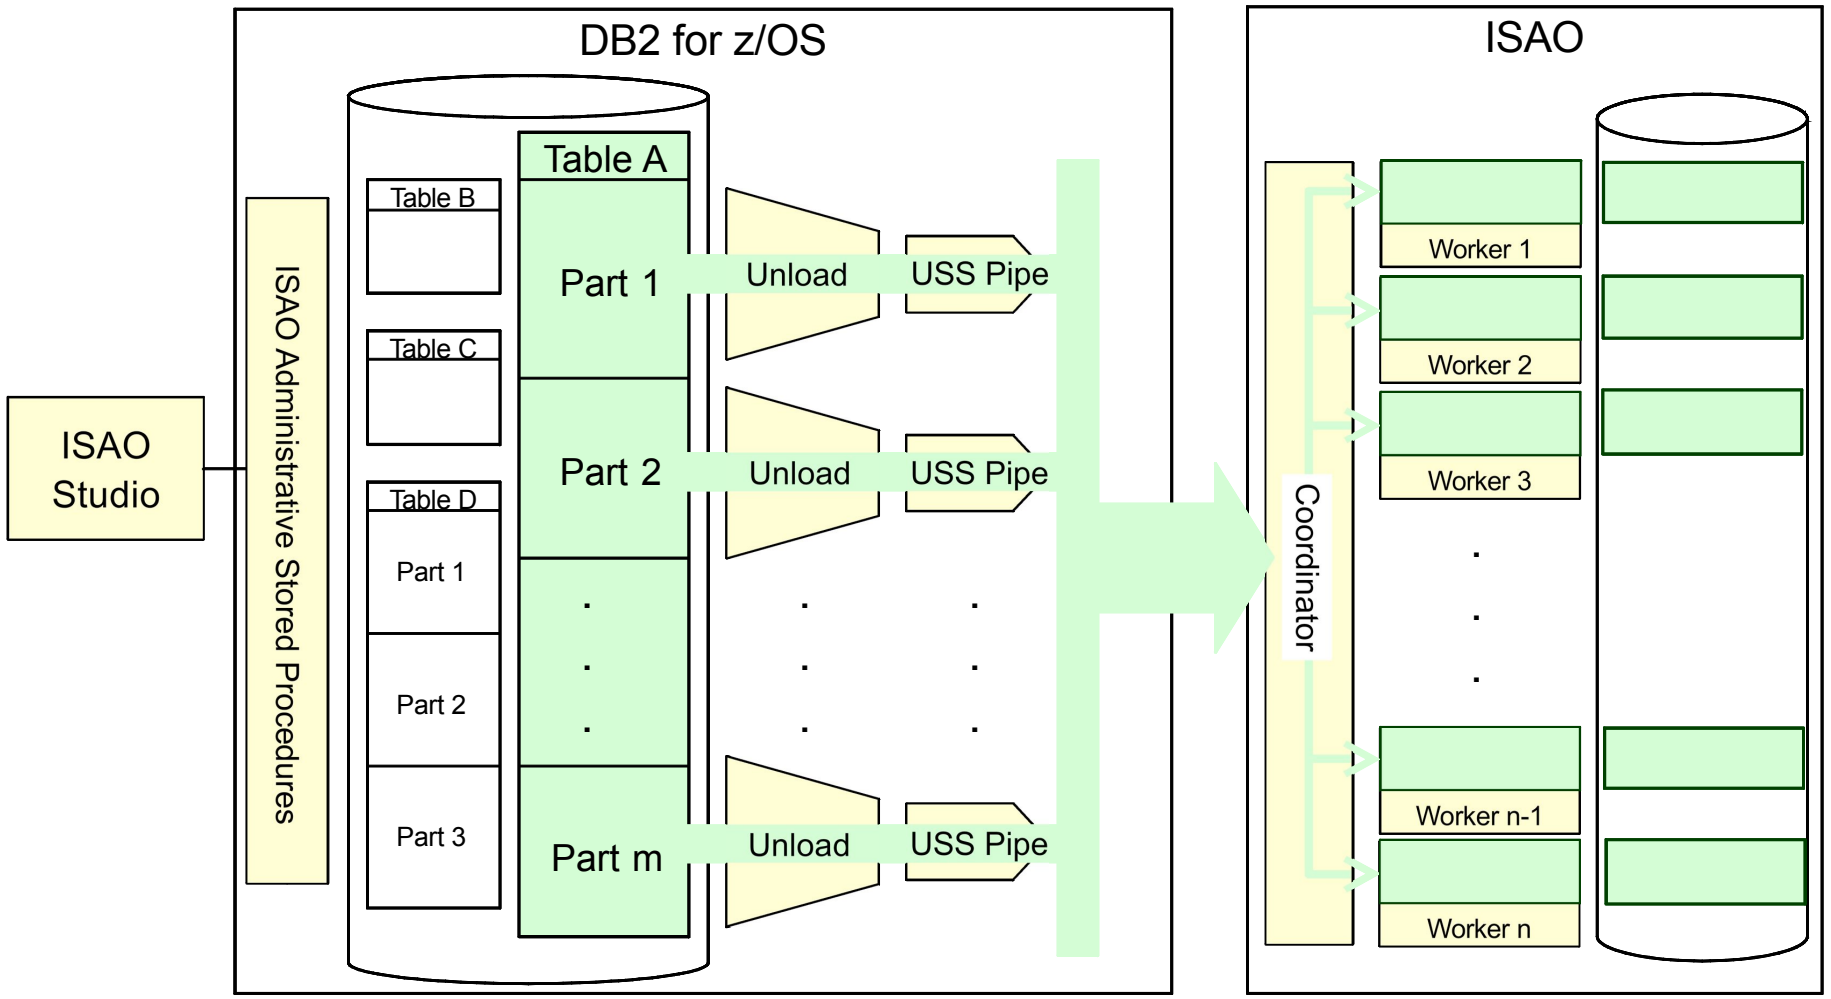

ISAO Mart Update

- Typical DW update operations:
  - LOAD RESUME and REPLACE
  - ADD and ROTATE PARTITION
  - SQL INSERT, UPDATE, DELETE
  - Delete complete partition or table
  - TRUNCATE TABLE
- ISAO will over time phase-in support for all the typical operations in this order
  1. Full table reload
  2. Updated partition reload
  3. Individual row change
- The marts update is initiated and controlled through ISAO Studio
- Queries off-loaded to ISAO before the marts are refreshed can return different result set as compared to not being off-loaded
  - In case this is not acceptable use SET CURRENT REFRESH AGE = 0

IDUG study on DW update frequency

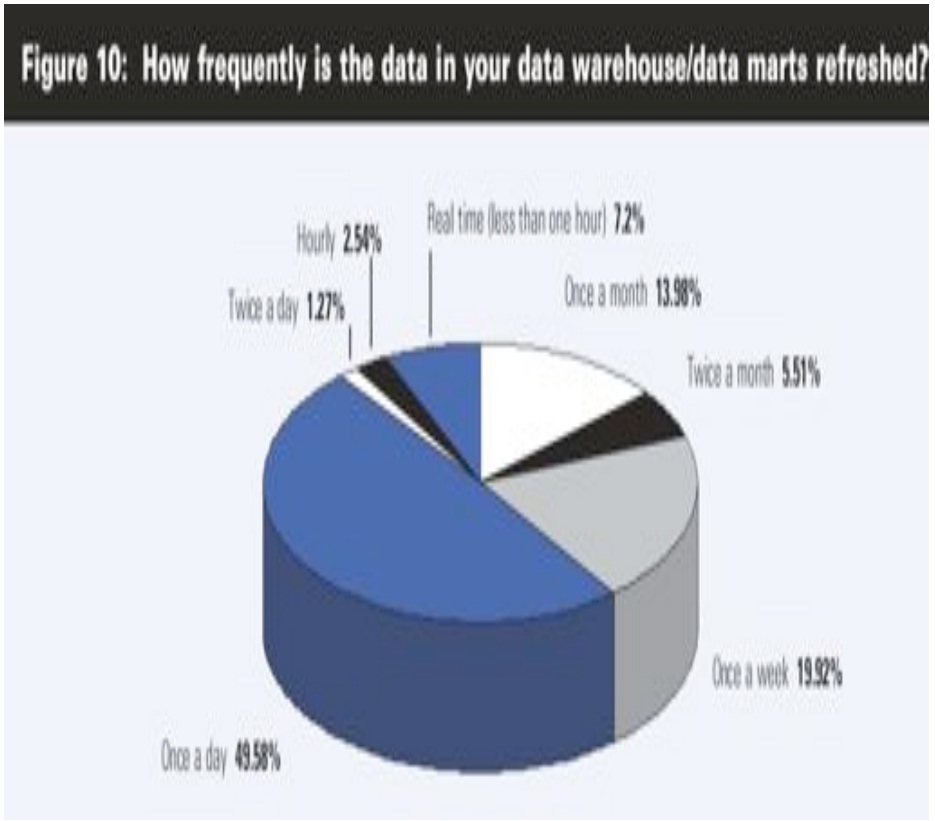

In 90% of cases DW is updated once a day or even less frequently

Query Execution Process Flow

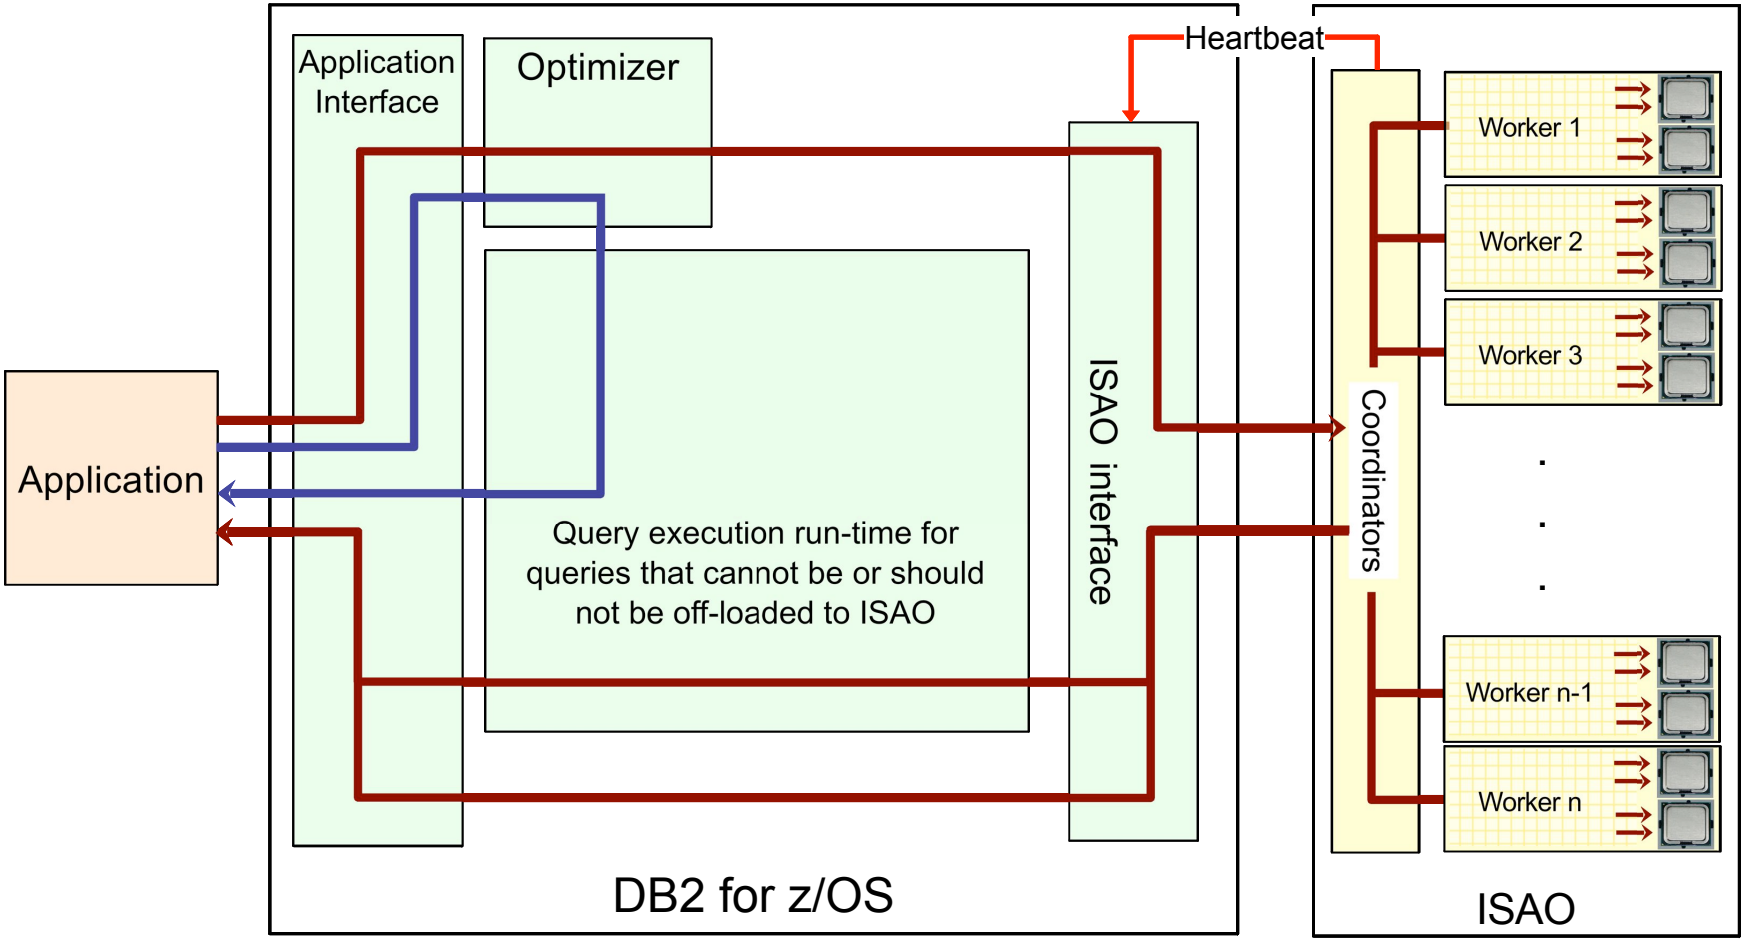

Additional DB2 Support

- Explain
  - Indicates ISAO involvement in query execution or the reason for no usage
  - New table DSN_QUERYINFO_TABLE
- Instrumentation
  - ISAO availability and performance indicators

- DB2 Commands
  - DISPLAY THREAD
  - DISPLAY ISAO
  - START ISAO
  - STOP ISAO

Possible ISAO Packaging Options

System z10

BladeCenters

ISAO Offerings	XS .5TB	S 1TB	M 2TB	L 3TB	XL 4TB
Enterprise Rack	1	1	1	2	2
BC-H Chassis, incl. double infrastructure (AMM, PDU, 10Gb Enet for data, 1Gb Enet for service, 4Gb F/C switches)	1	1	2	3	4
Blades	7	14	28	42	56
DS5020 with # of SATA disks (750GB)	16	16	16	32	32
Storage Expansion Units	0	0	0	1	1

DS5020

Testing Results

- The problem queries provided by a customer
- Expert database tunning done on all the queries
  - Q1 – Q6 even after tuning run for too long and consume lots of resources
  - Q7 improved significantly – no ISAO offload is needed
- The table shows elapsed and CPU times measured in DB2 (without ISAO)

Query	Times measured in DB2 **without** ISAO			
	Total Elapsed	CP	zIIP	Total CPU Time
Q1	**0:02:43**	0:03:52	0:02:39	**0:06:31**
Q2	**0:38:31**	0:11:52	0:36:10	**0:48:02**
Q3	**0:00:25**	0:00:04	0:00:15	**0:00:19**
Q4	**0:26:33**	0:13:43	0:20:50	**0:34:33**
Q5	**0:00:35**	0:00:09	0:00:29	**0:00:38**
Q6	**1:30:35**	5:53:30	1:29:56	**7:23:26**
Q7	**0:00:02**	0:00:02	0:00:00	**0:00:02**

Testing Results: Performance Improvement after Adding ISAO

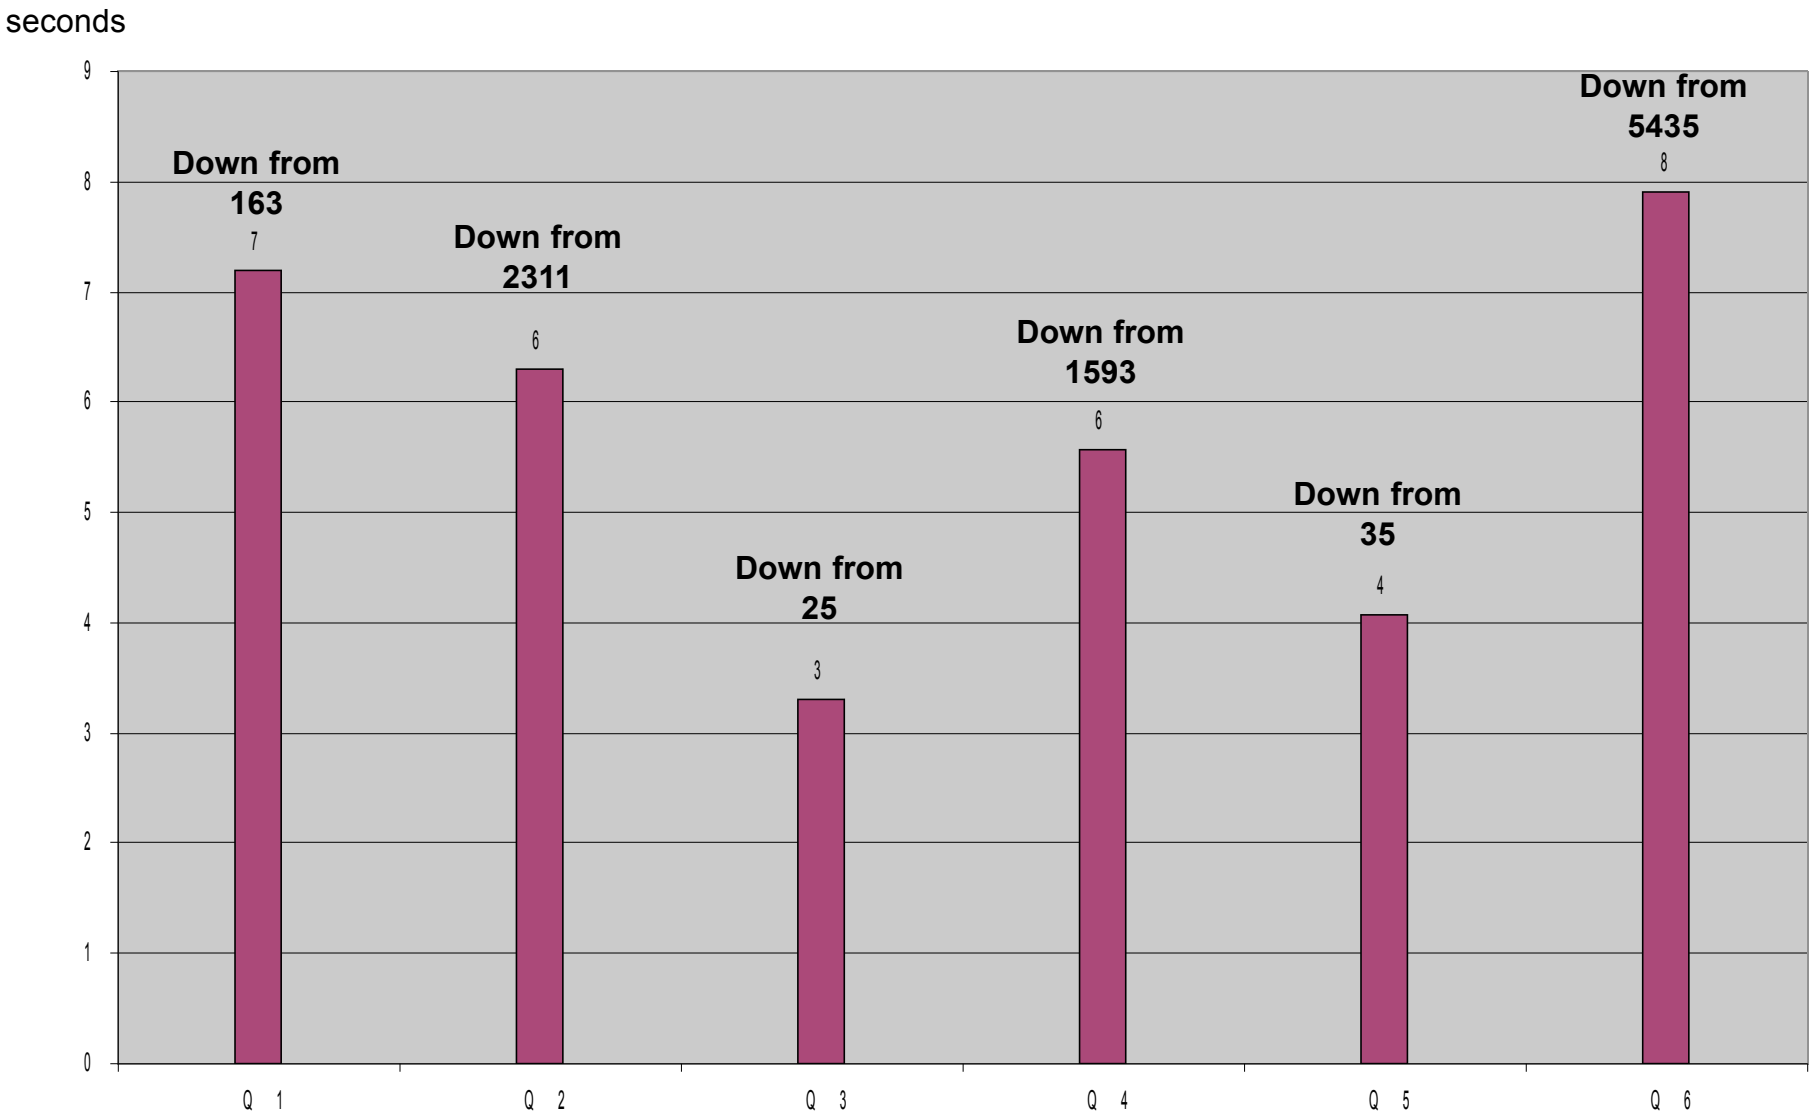

Agenda

- Business and Technology Drivers
- Key Design and Operational Features
  - IBM Smart Analytics Optimizer as a virtual DB2 component
  - ISAO engine
  - Marts
  - Query execution
- Supported workloads

What Is ISAO Ideally Suited For?

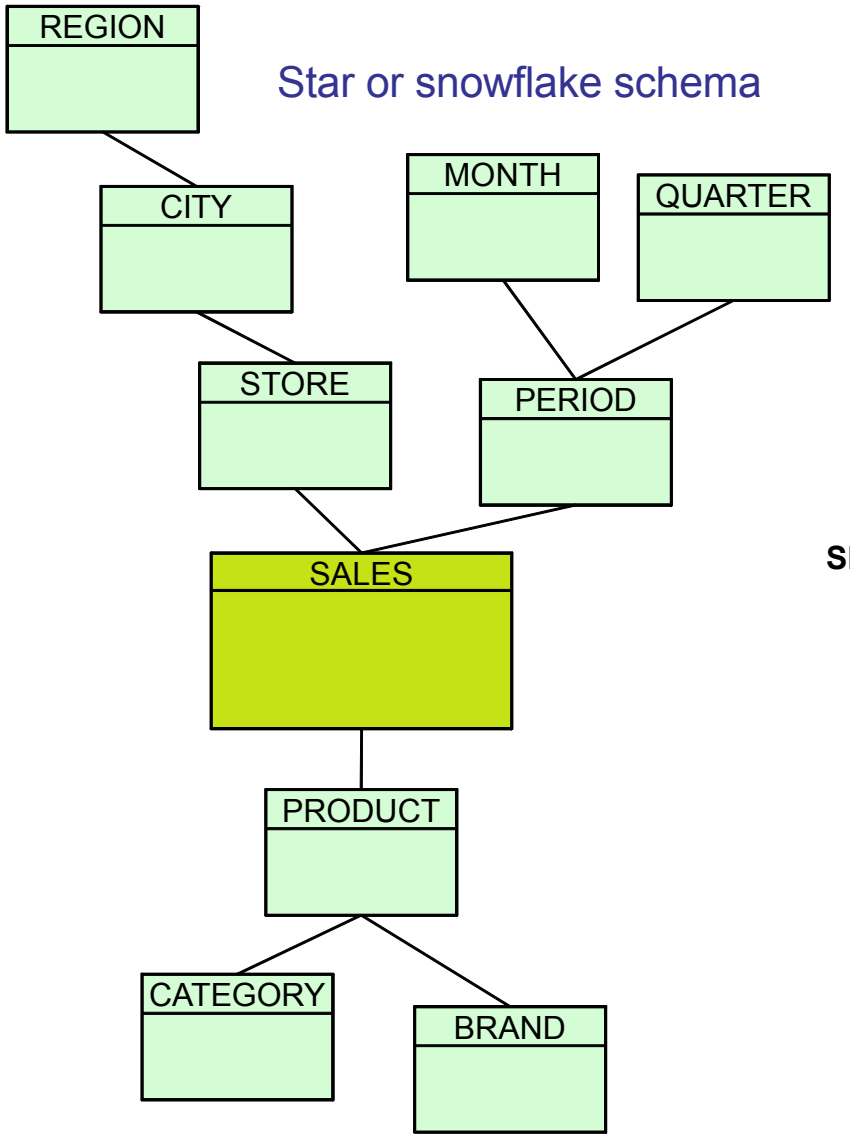

Star or snowflake schema

Complex, OLAP-style queries that typically:

- Need to scan large subset of data (unlike OLTP queries)
- Involve aggregation function such as COUNT, SUM, AVG.
- Look for trends, exceptions to assist in making actionable business decisions

```
SELECT PRODUCT_DEPARTMENT, REGION, SUM(REVENUE)
    FROM FACT_SALES F
        INNER JOIN DIM_PRODUCT P ON F.FKP = P.PK
        INNER JOIN DIM_REGION R ON F.FKR = R.PK
        LEFT OUTER JOIN DIM_TIME T ON F.FKT = T.PK
    WHERE T.YEAR = 2007
        AND P.TYPE =  'SOFTWARE'
        AND R.GEO = 'SOUTH'
    GROUP BY PRODUCT_DEPARTMENT, REGION
```

First Release Restrictions

- One query block at a time
  - If a query consists of multiple query blocks, ISAO processes them one by one
  - Outer query block that contains a subselect is not processed by ISAO (DB2 does not pass the subselect result set to ISAO)
  - Multiple query blocks can be (but do not have to be) generated by
    - Subselects in quantitative predicates (SOME, ANY, ALL)
    - EXISTS or IN predicates with subselects
    - UNION, INTERSECT, EXCEPT
- Examples:

Nested Table Expression

```
SELECT * FROM
(SELECT  C1+C2 FROM TA) TX
```

Common Table Expression

```
WITH DTOTAL (deptno, totalpay) AS
      (SELECT deptno, sum(salary+bonus)
      FROM DSN8810.EMP GROUP BY deptno)
SELECT deptno FROM DTOTAL
WHERE totalpay = (SELECT max(totalpay) FROM DTOTAL);
```

IN predicate with subquery

```
SELECT  ... FROM    ...   WHERE   ...
 AND   ( (A11.STORE_NUMBER IN
      (SELECT  C21.STORE_NUMBER
            FROM   USRT004.VL_CSG_STR   C21
            WHERE   C21.CSG_NUMBER IN (4643) ))
```

First Release Restrictions

- Limited support for very large dimension tables
  - Especially if the predicates on them are not selective
- No static SQL
- No full outer join, no right outer join
- Only equi-joins (no range join predicates)
- No queries that do not include at least one fact table
- No queries that spread across multiple marts
- Not all DB2 functions
  - No mathematical functions such as SIN, COS, TAN.
  - No user defined functins
  - No advanced string functions such as LOCATE, LEFT, OVERLAY.
  - No advanced OLAP functions such as RANK, ROLLUP, CUBE
- Not all DB2 data types such as LOBs, ROWID, XML.

Data Types Support

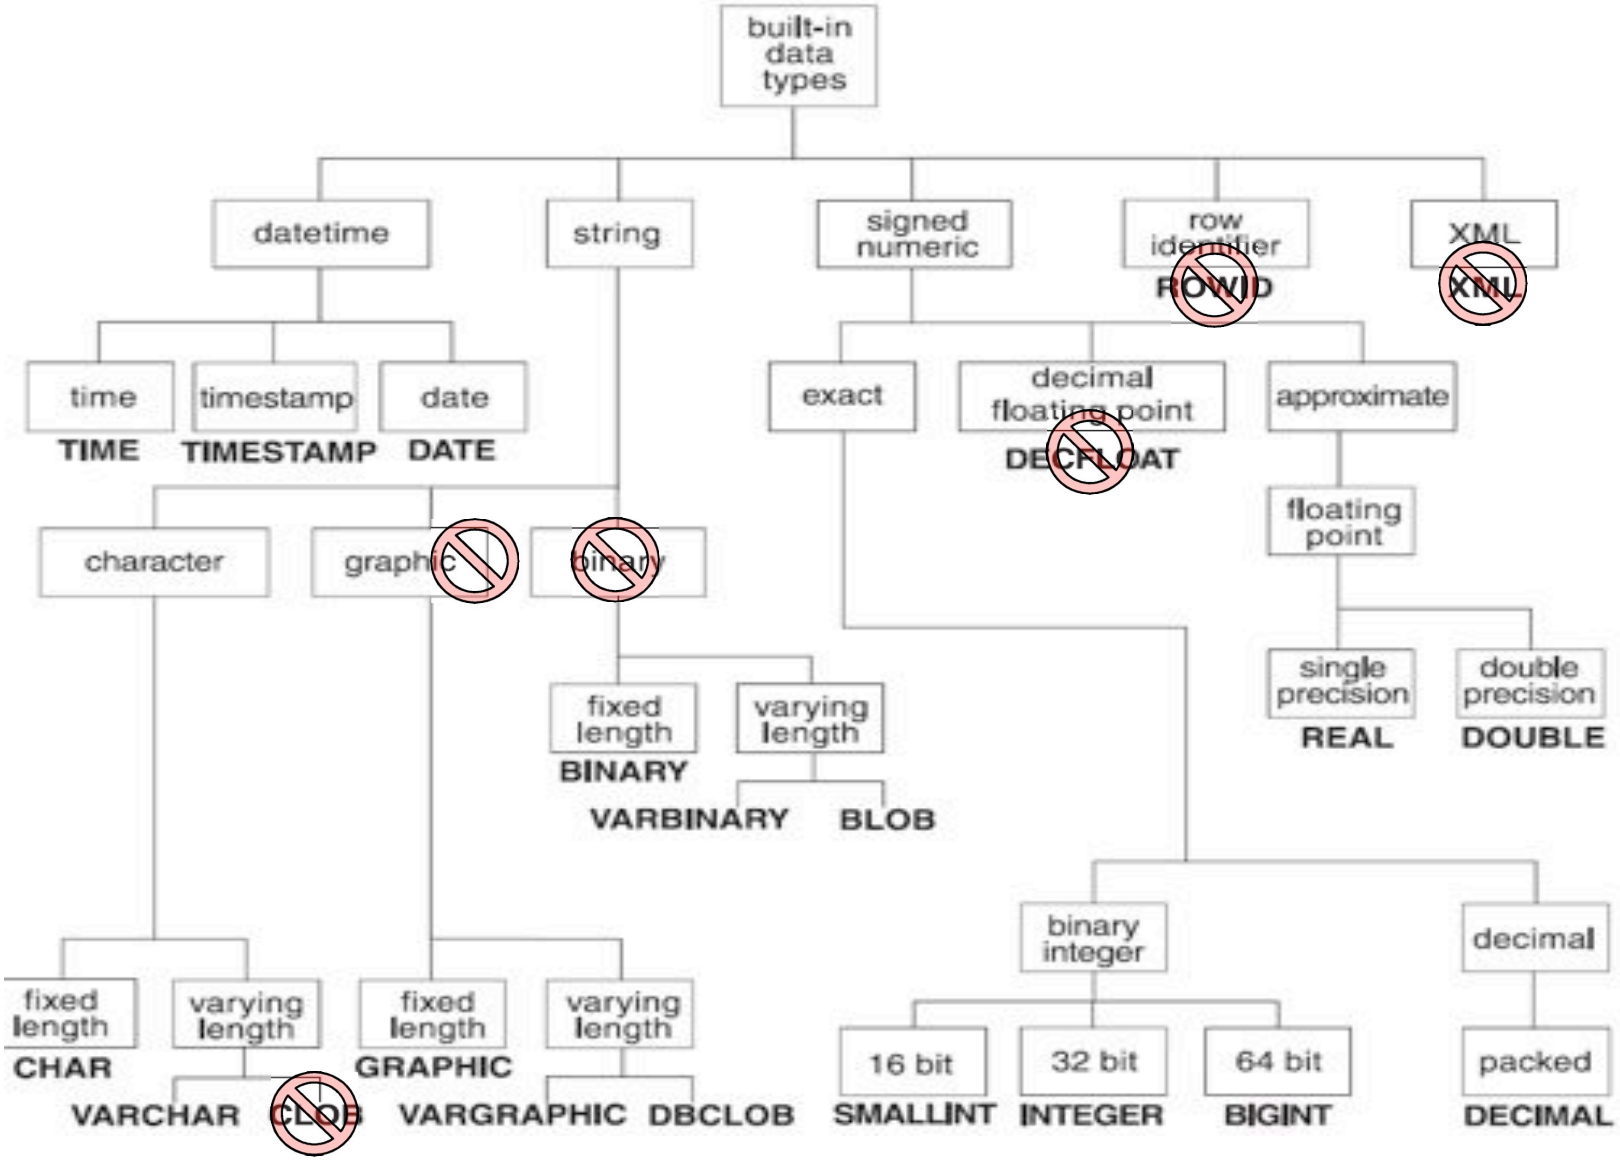

Options for Workload Analysis

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 ISAO environment. Addresses all inhibitors for offload and data mart definition questions.

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 Smart Analytics Optimizer capabilities

Report for a first assessment:

- Query offload potential
- Estimated performance gain
- Estimated CP cost saving

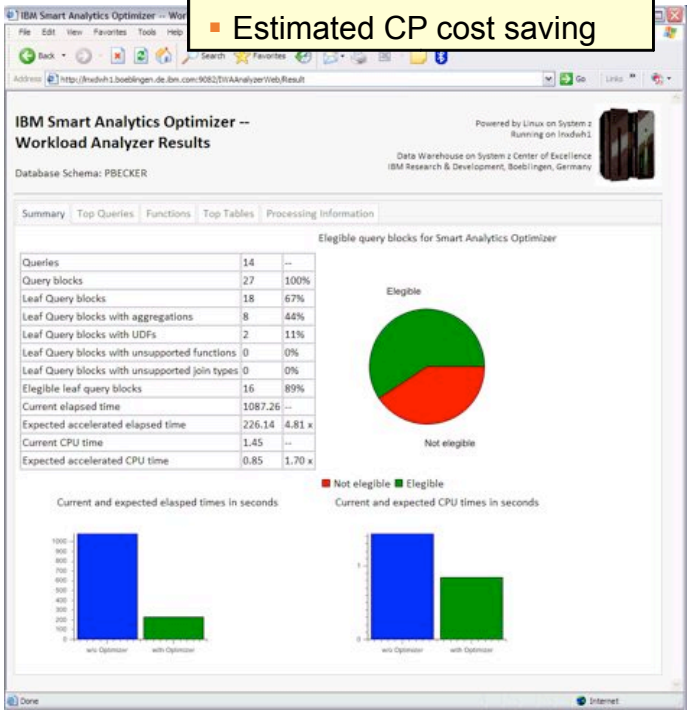

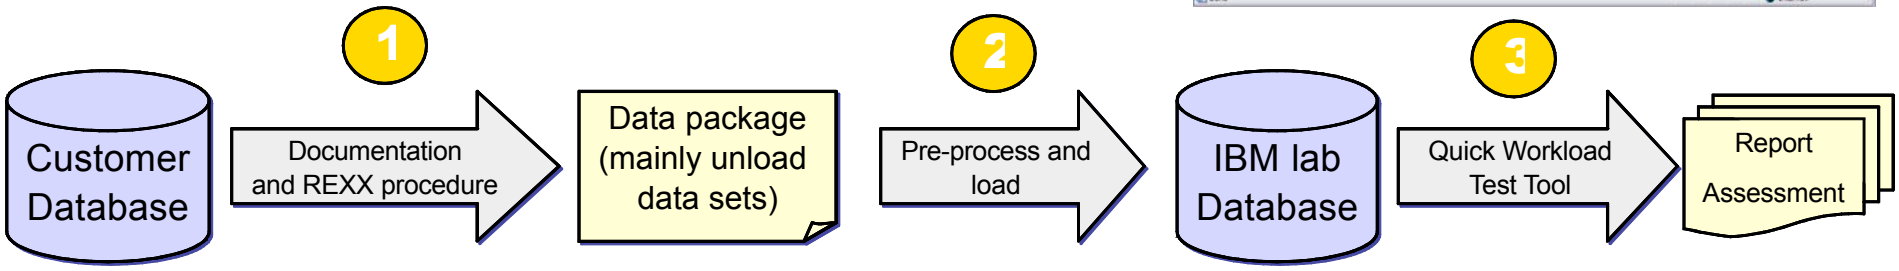

IBM Smart Analytics Optimizer

Summary of Value Proposition

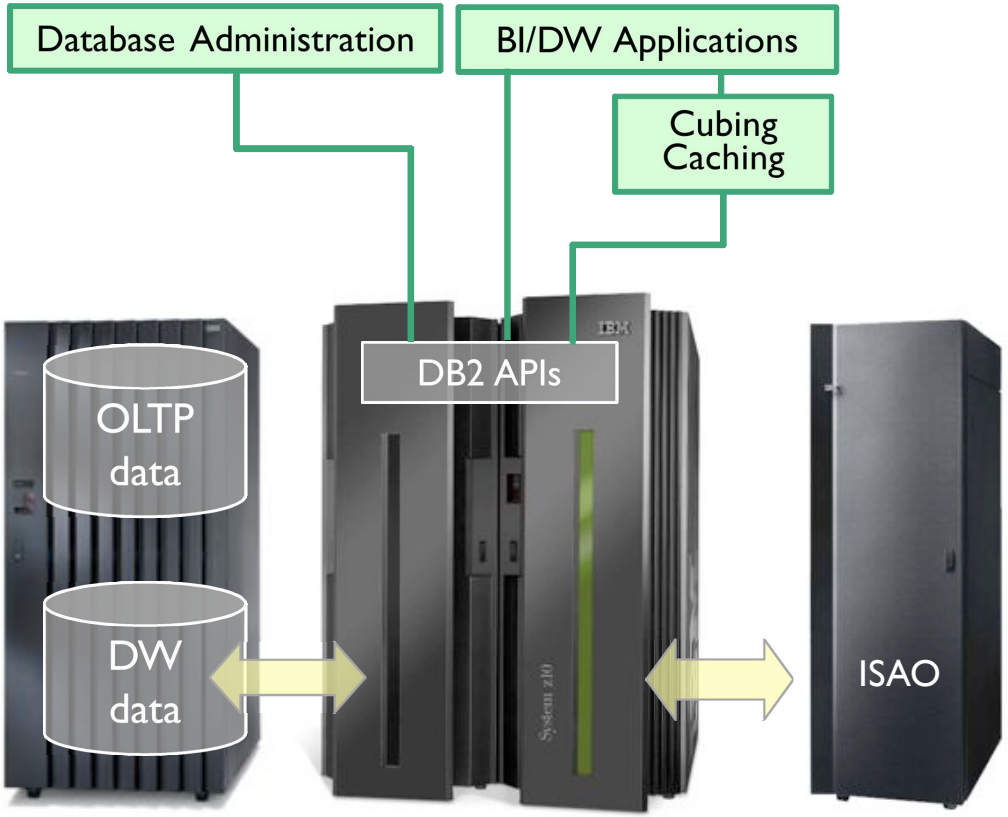

- Seamless integration of new computing paradigms into proven technology
  - Massive multi-core and vector processing
  - In-memory database
  - Hybrid row- and column-based store
  - No changes to the applications, applications continue to attach to DB2
  - Preserving traditional System z and DB2 quality of service, full fencing and protection of DB2 against possible ISAO failures
- Order of magnitude performance improvement
  - Linear scaling with the number of CPUs
- Reducing need for tedious tuning of DB2 (MQTs, aggregates, indexes, etc.)
- Appliance characteristics
  - User/reference guide assisted installation, initial configuration
  - Hands free operations
- Providing building block for Dynamic DW and Operational BI
- Augmenting System z value proposition as the overall Enterprise Data Hub
  - System z hybrid topology enables additional transparency and management integration