





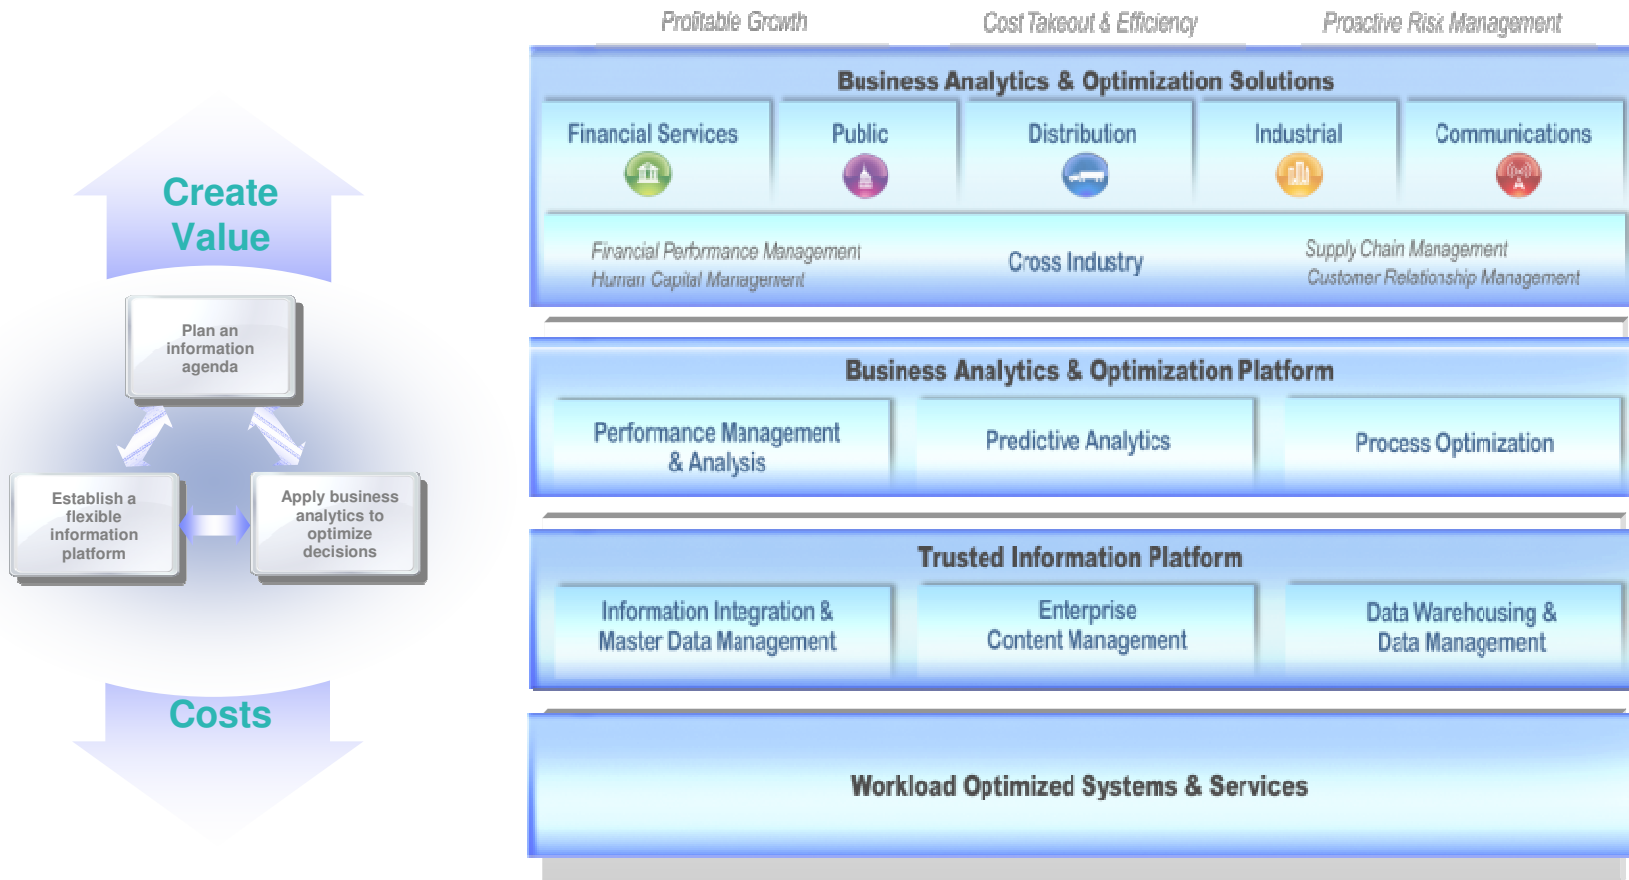
Introducing IBM Smart Analytics Optimizer

Dan Wardman
VP, Information Management Mainframe Software





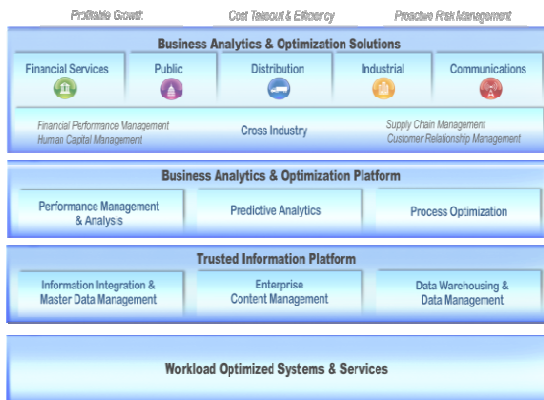
Getting started with an information-led transformation





Information-Led Transformation for System z

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



and much more!

DB2 X for z/OS

Smart Analytics Optimizer

IMS 11

2010

Analytics

Content Integrator for z/OS

InfoSphere Warehouse for System z

FileNet Content Manager for System z

Cognos 8 BI for System z

Open Access DW Tools

Business Intelligence

Data Studio pureQuery for System z

2009

Information Agenda

InfoSphere MDM Server for System z

Optim Data Governance for System z

App Dev Tools

Information Server for System z

Content Manager 8.4

Data Archiving Test Data Mgmt

2008

Data Warehousing

IOD Launch 2006

pureXML

SOA

ETL & Data Quality

DB2 software

InfoSphere software

Cognos software

FileNet software



The Resurgence of Data Warehousing and Business Intelligence on System z

→ **IBM has invested hundreds of millions of dollars to bring new state of the art capabilities and solutions to System z in support of customers' warehouse and BI requirements**

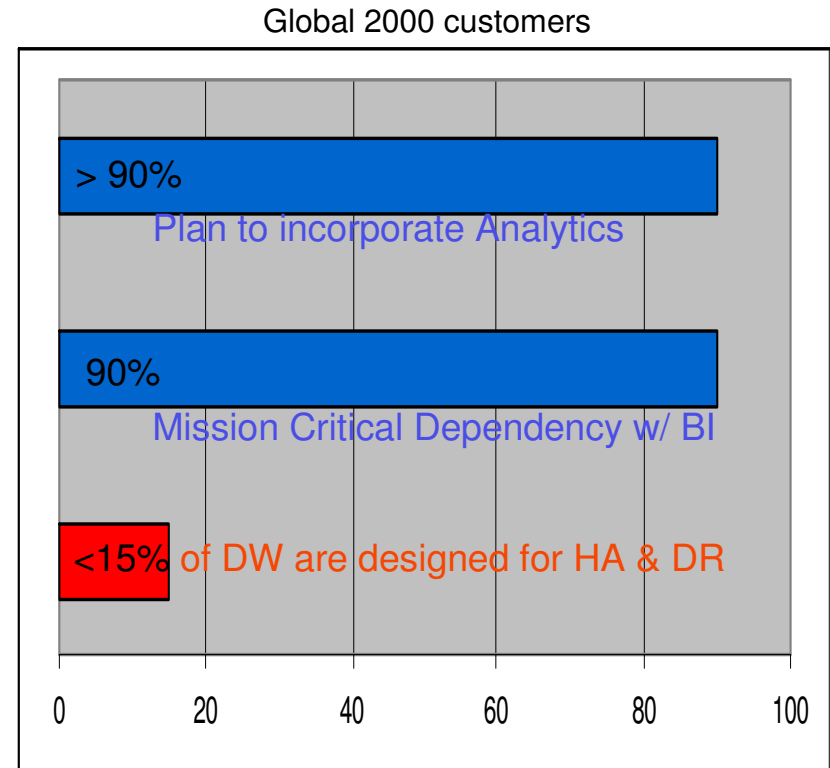
- A 2007 study by IDUG found that nearly 50% of IDUG respondents are already using DB2 for z/OS for data warehousing. 78% indicated a desire for more capabilities in warehousing, query and reporting.
- *Analysts Agree! IBM's Data Warehousing & BI breadth on System z is a game changer in the market.*
Donald Feinberg, Gartner





Mission Critical Workloads Require Highest QoS

- **More than 90%** of Global 2000 companies plan to incorporate analytics into multiple operational applications that access the data warehouse by 2010, but fewer than 15% of data warehouses have been designed to provide high availability, failover, disaster recovery and the remaining components of mission-critical systems.
- By the end of 2009, **90% of Global 2000** companies will have implemented some type of mission-critical dependency between the warehouse and at least one revenue supporting or cost-controlling operational application — up from less than 25% in 2007.
- **Fewer than 15% of data warehouses** in 2007 have been designed to provide high availability, failover, disaster recovery and the remaining components of mission-critical systems.

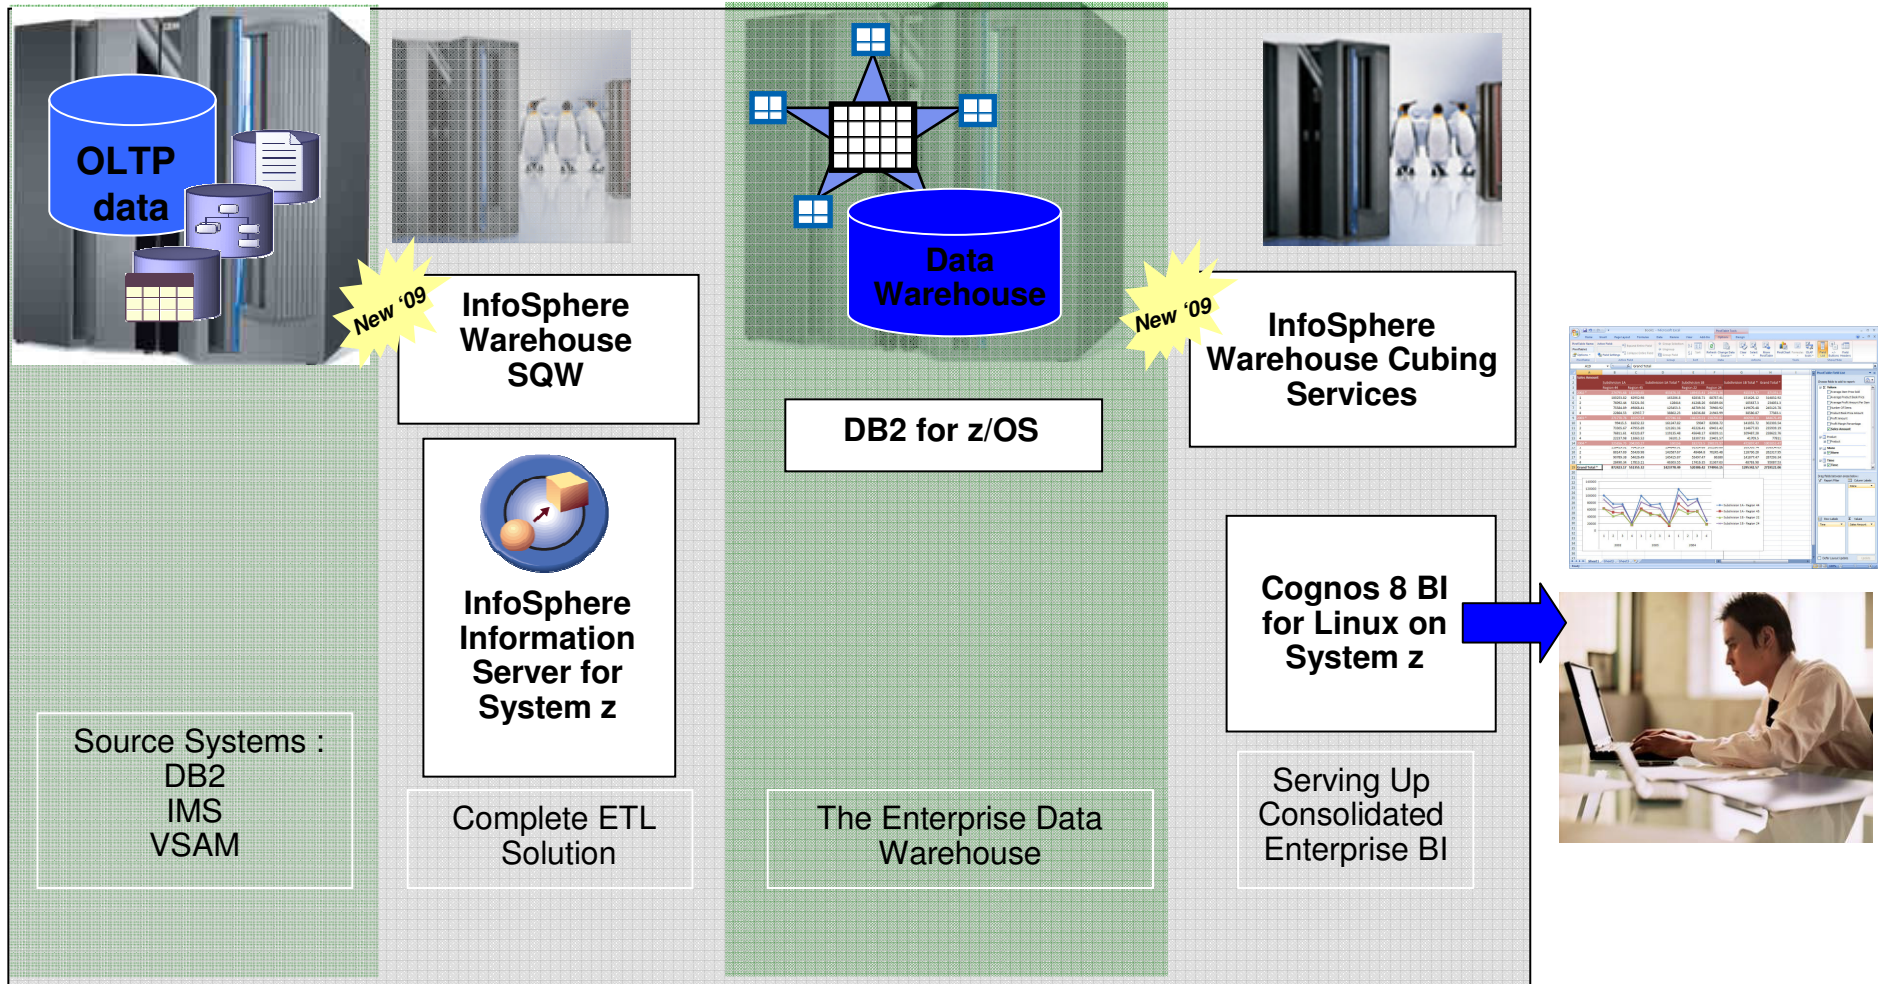


- Sounds like a good match for System z value proposition
 - Proven reliability and continuous availability capabilities
 - Exploiting synergistic effects of proximity to the operational data

¹ Operational Analytics and the Emerging Mission-Critical Data Warehouse, 14 May 2007



The Data Warehouse and BI Solution on System z



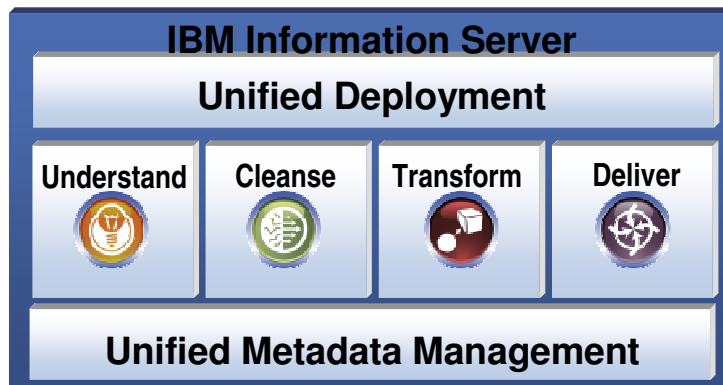
Combining the Reliability and Availability of DB2 for z/OS with Cost Effective Applications running on Linux for System z



InfoSphere Information Server for System z

Accelerating the delivery of trusted information

Profile, cleanse, and transform information from heterogeneous data sources to drive greater business insight



- Significant cost savings on System z
- Scalable to any volume and processing requirements
- Fully integrated, auditable data quality
- Metadata-driven integration for increased productivity



InfoSphere MDM Server for System z

The first multi-domain, multi-function MDM product in the market

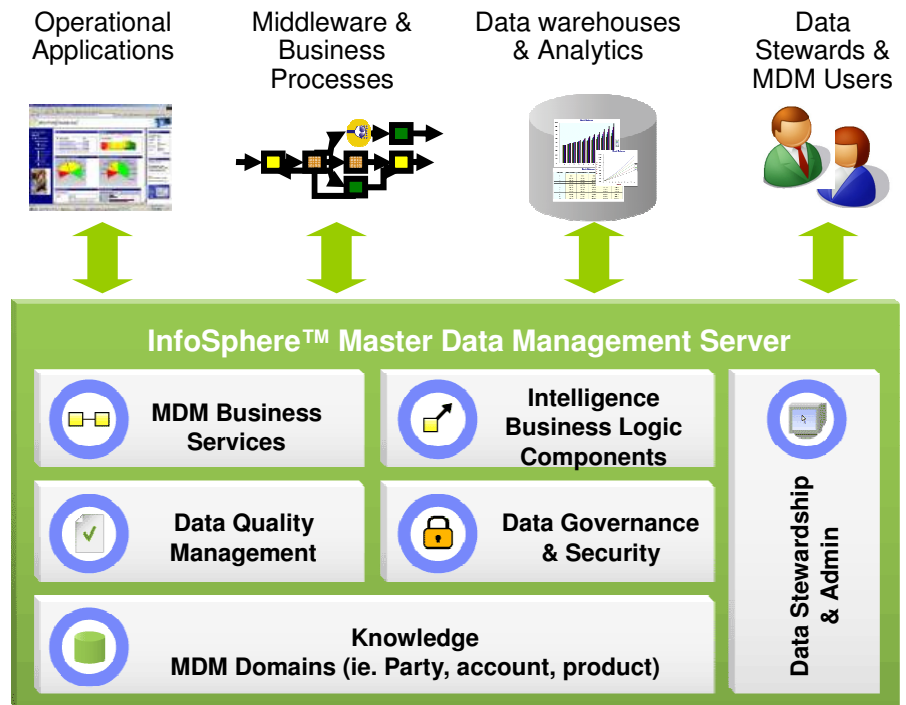
■ Packaged to address all types of MDM implementations

- From small “registry” projects to strategic “transaction-hub” deployments
- Allows clients to grow as required by implementing existing functionality
- Significantly lowers client risk and time/cost to implement

■ Enables a SOA Library

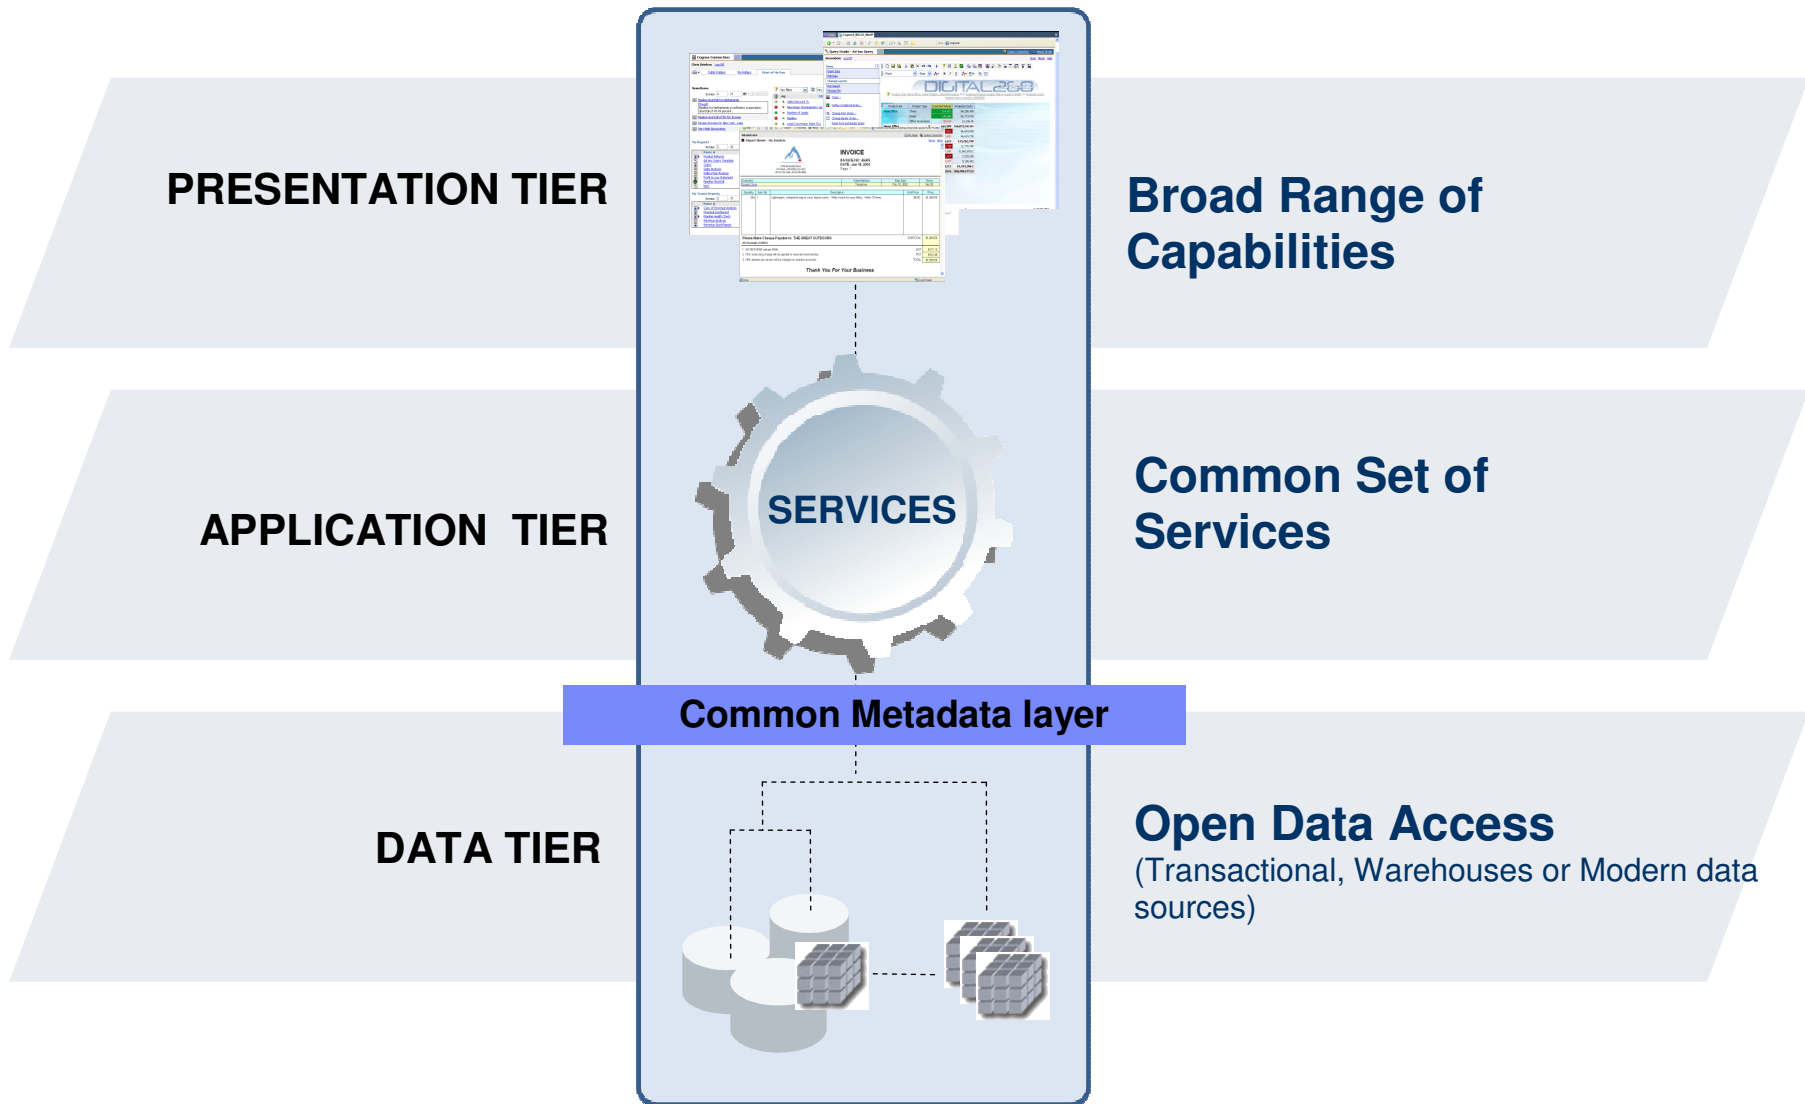
- 800 pre-packaged business services
- Significant out of the box functionality
- Reduces total cost of ownership

■ Provides leading performance & scalability





The IBM Cognos 8 platform





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

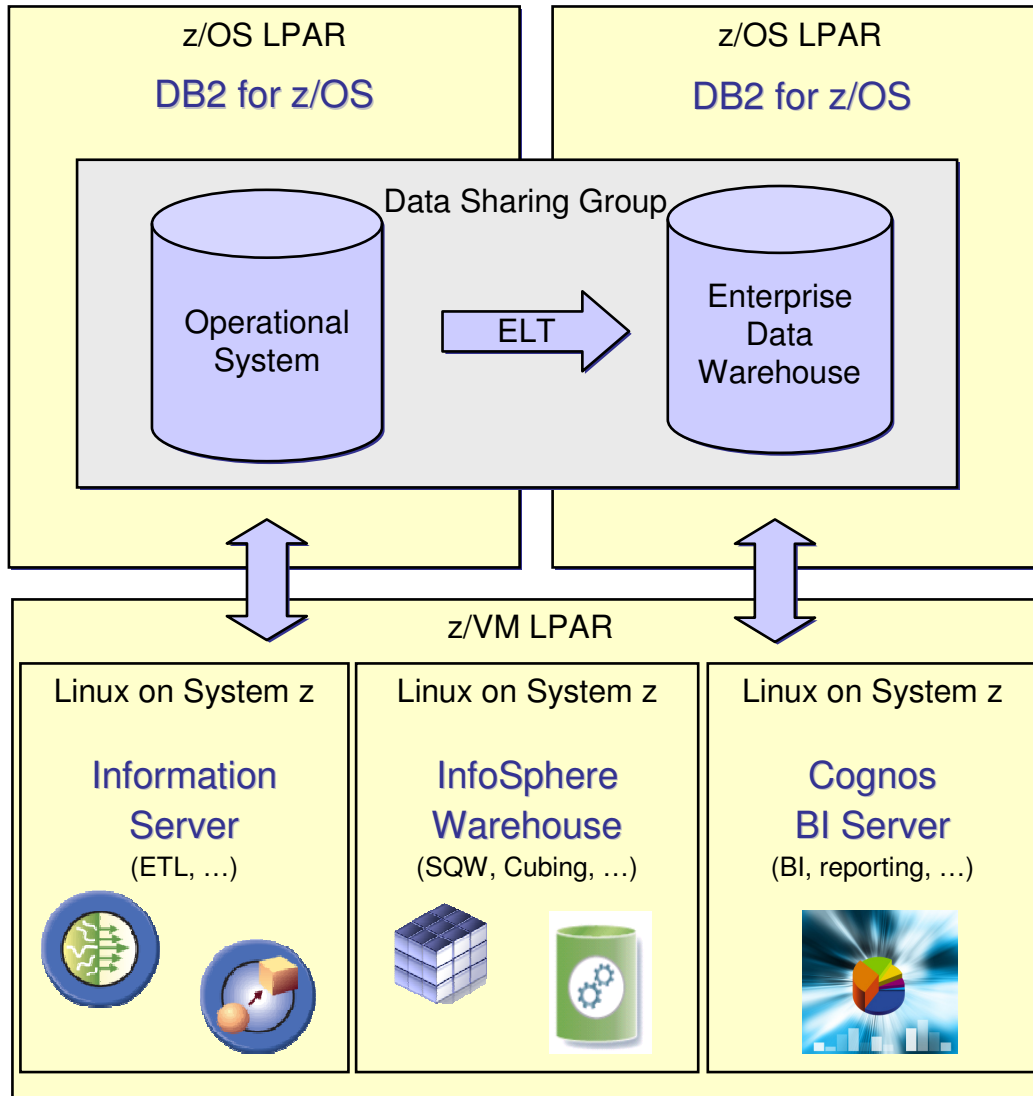


Business Challenges and Technology Trends

- 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, prebuilt 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
 - Massive amounts of real memory
 - Specialized physical data designs: row-store vs. column-store



Ultimate Consolidation Opportunity

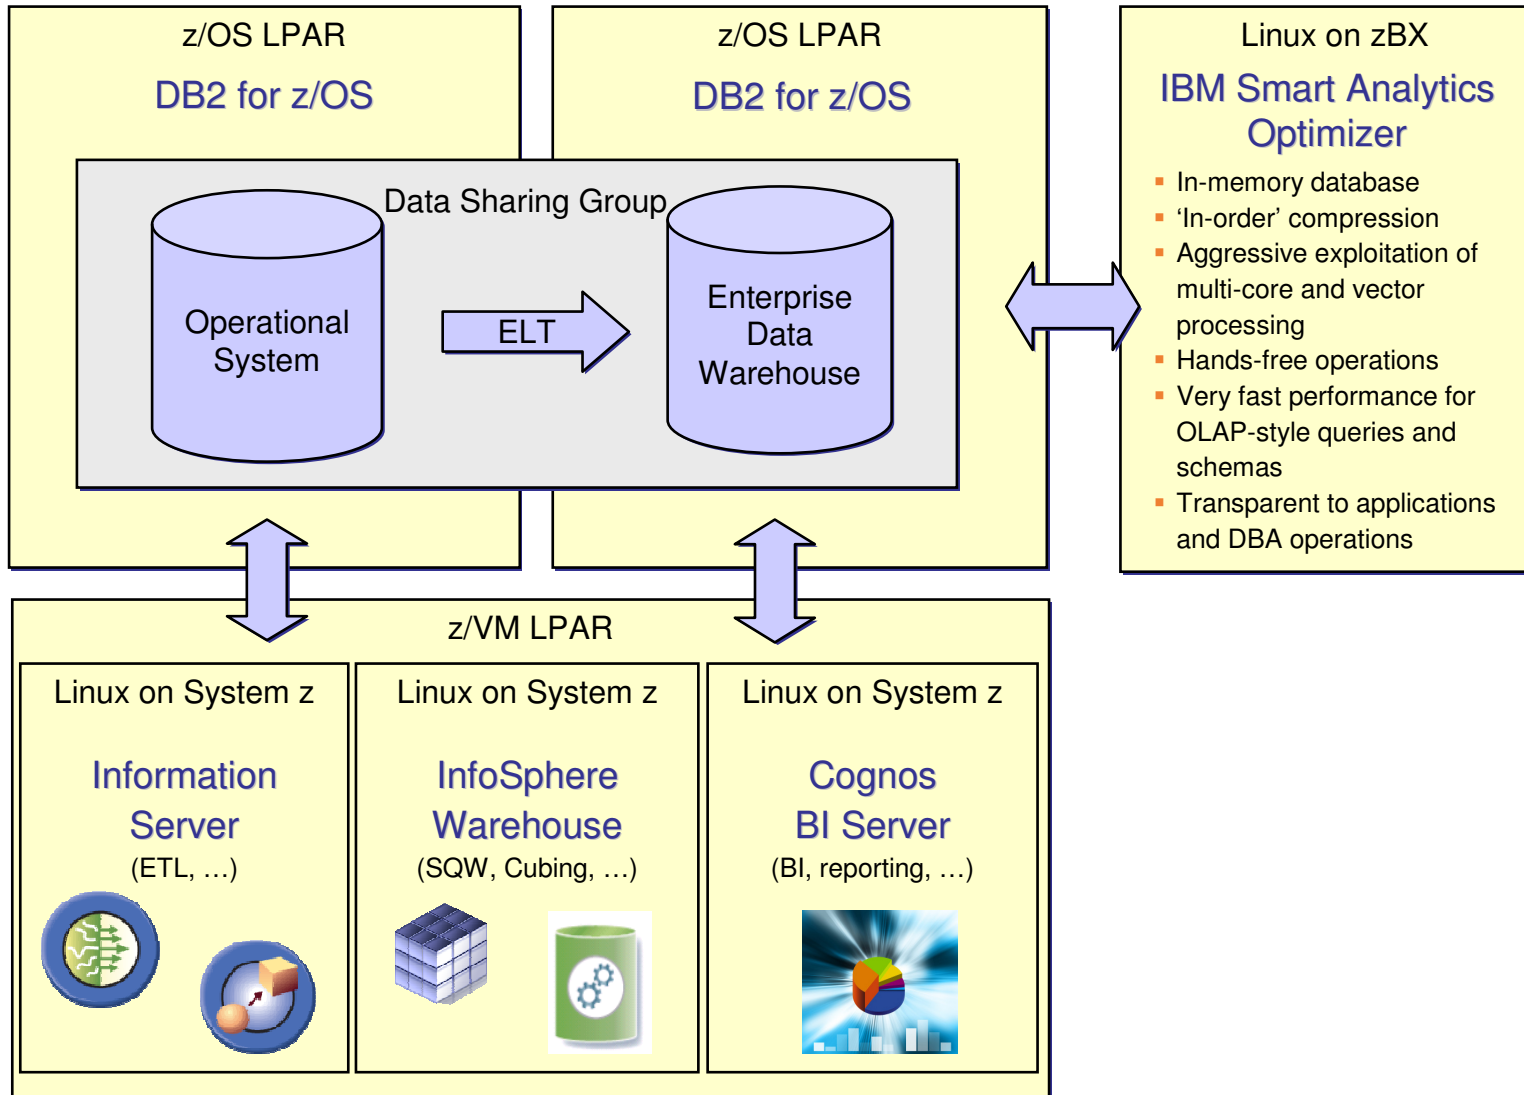


- Consolidation of mission-critical data on System z
- Leveraging existing environment, high availability, backup and governance procedures as well as skills
- Efficient data movement within a data sharing group (no network)
- Performance and TCO improvements through cubing services (data marts) and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



IBM Smart Analytics Optimizer

Adding Industry Leading Performance



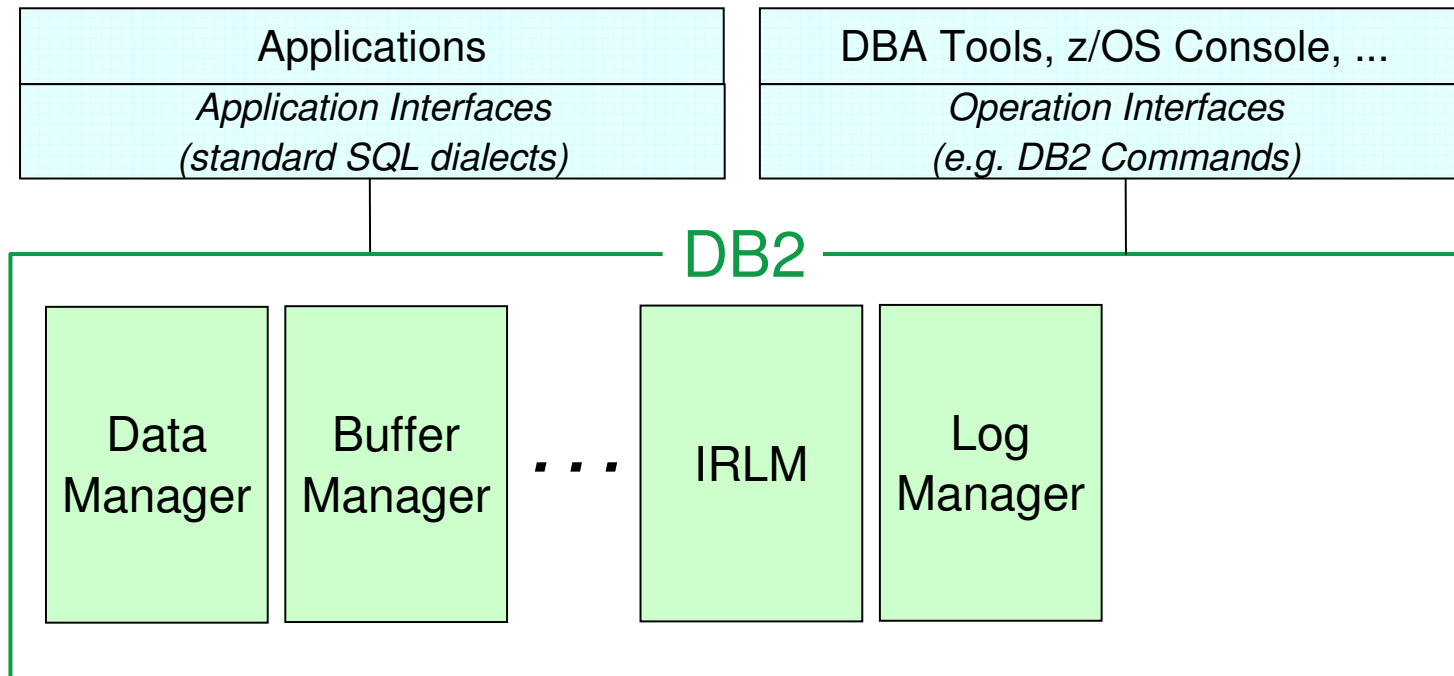


Agenda

- 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



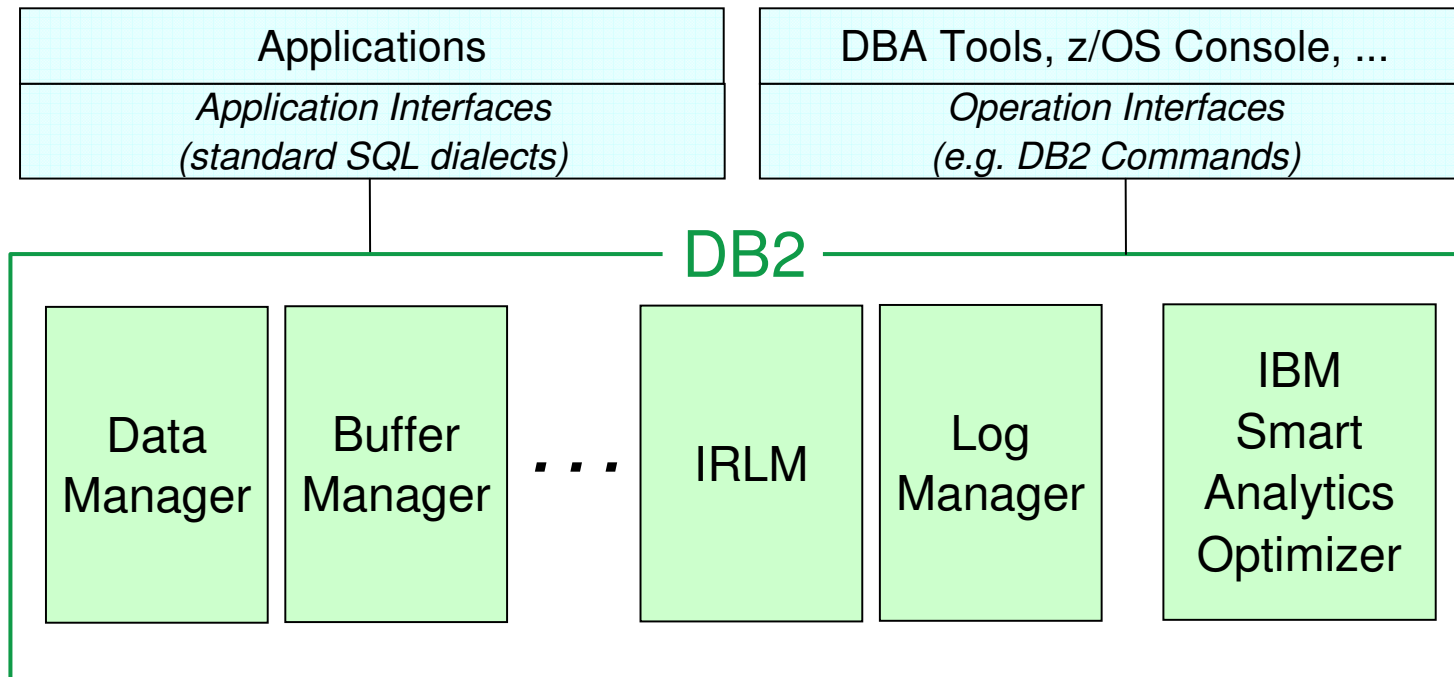
Deep DB2 Integration within zHybrid Architecture





Deep DB2 Integration within zHybrid Architecture

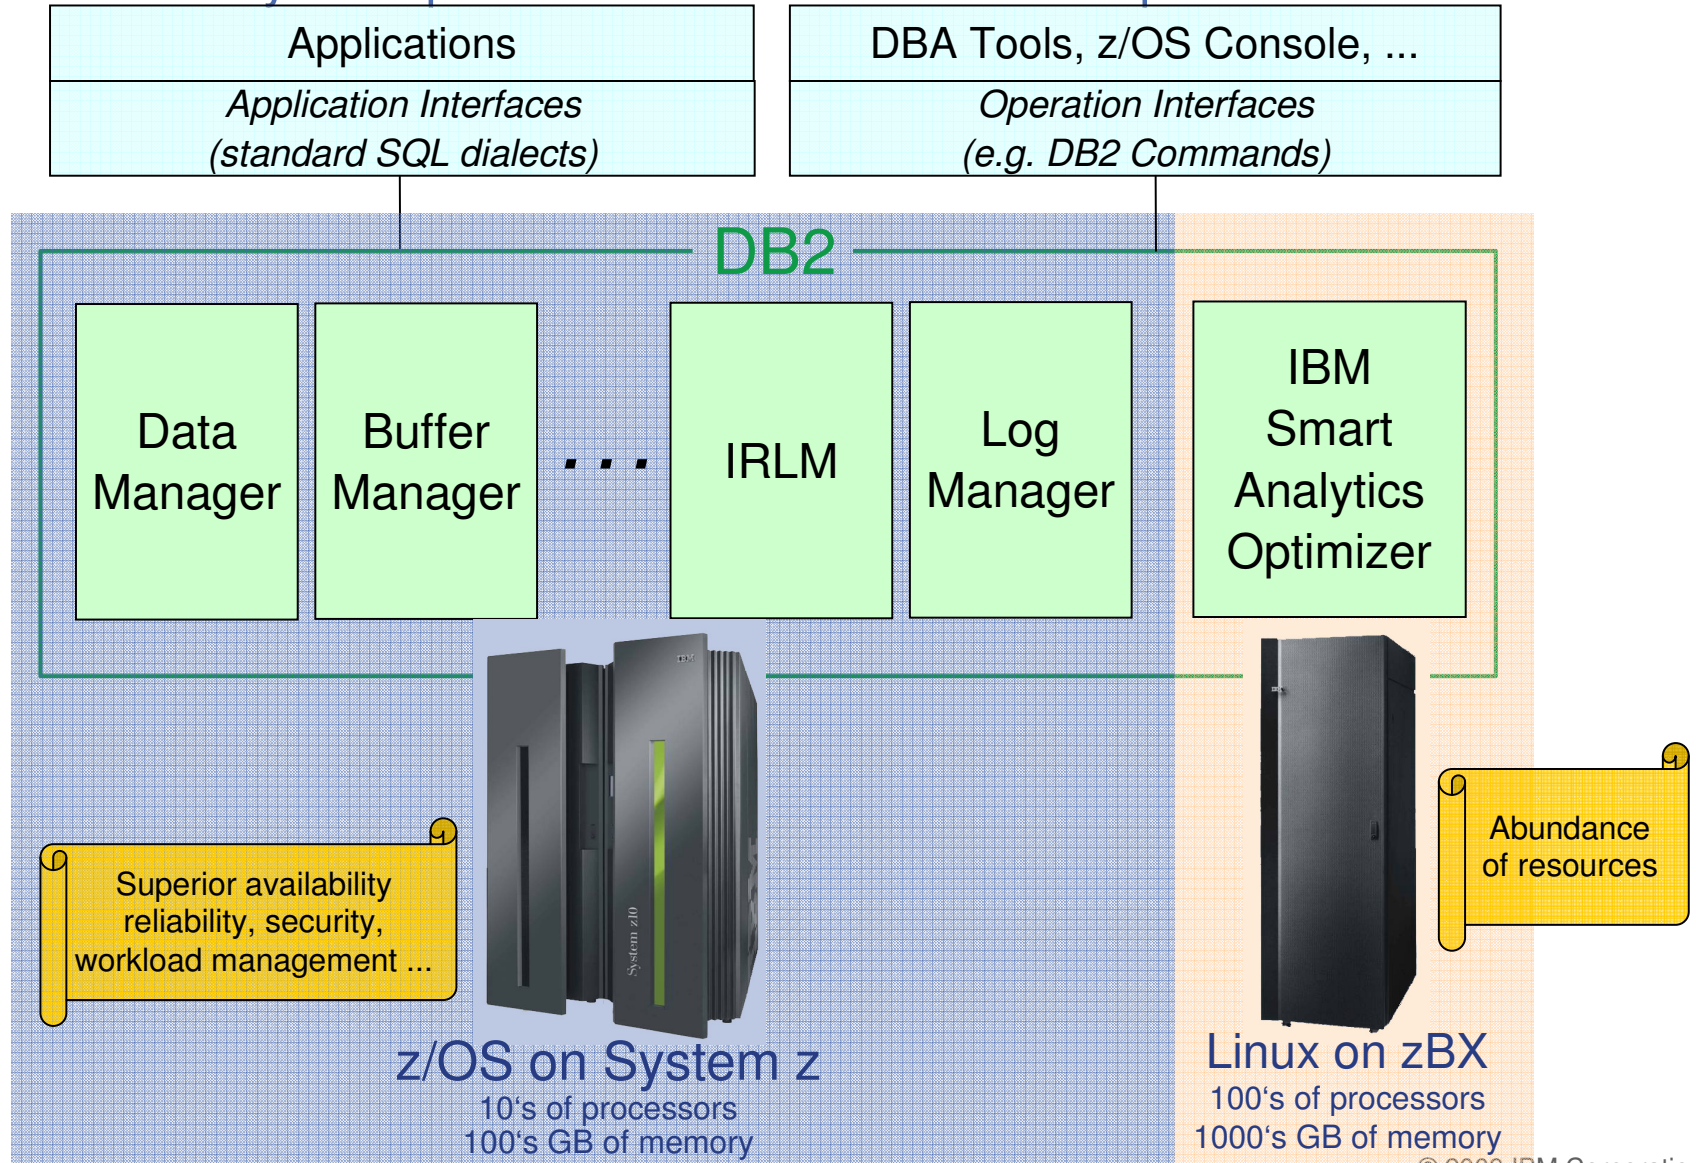
IBM Smart Analytics Optimizer as a Virtual DB2 Component





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



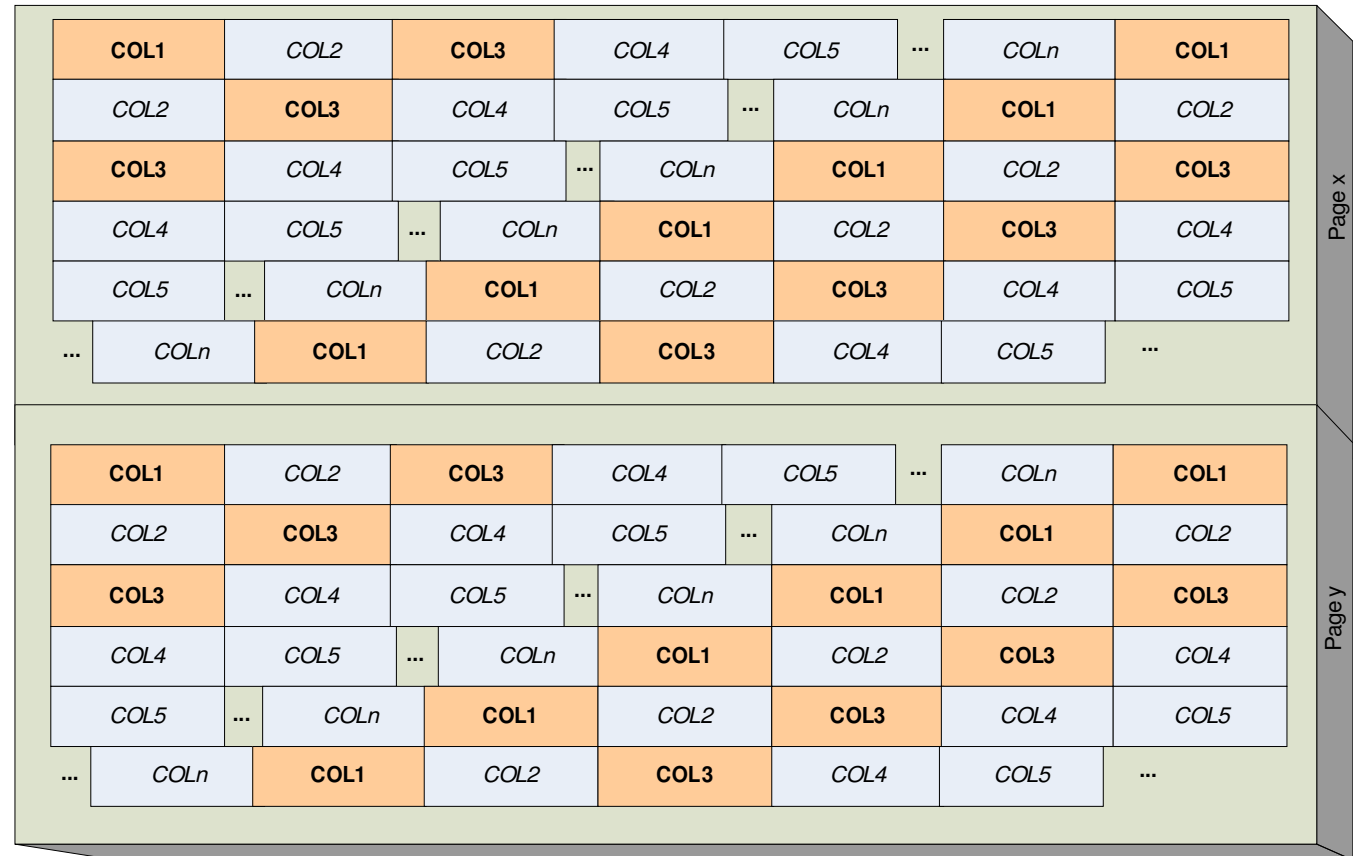
OLTP vs. DW

- 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 optimized 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.

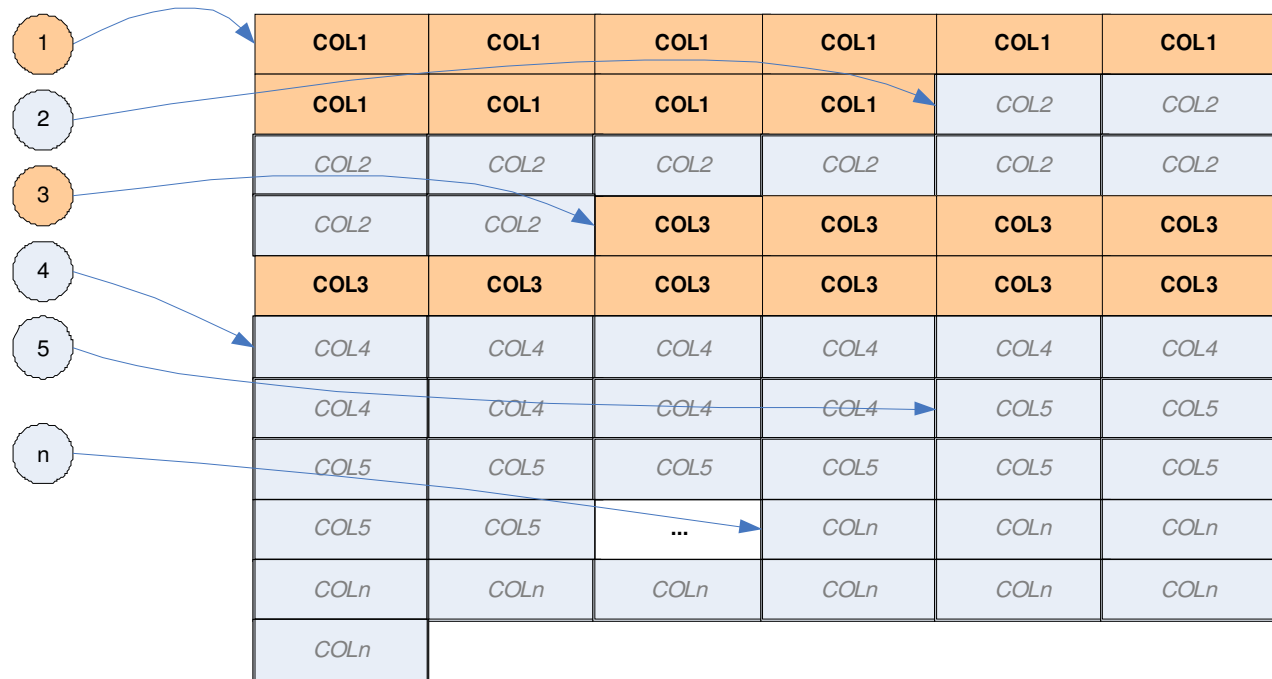


While a **Row – Store** is very efficient for transactional workloads, it is suboptimal for analytical workloads where only a subset of the attributes is needed!



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.

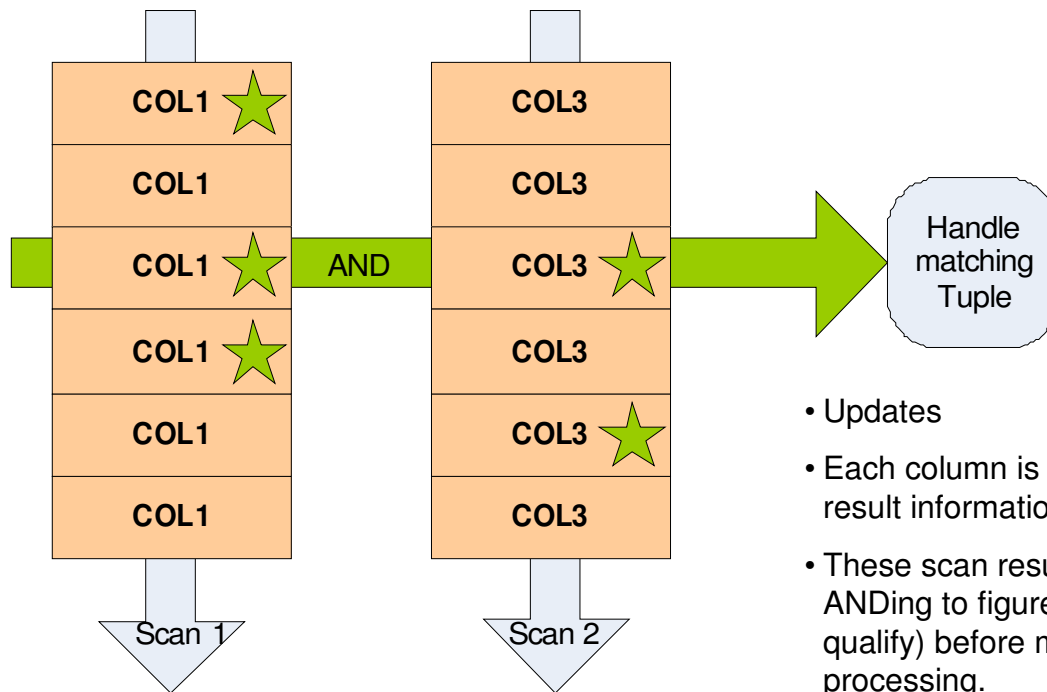


In a **Column – Store**, the data is also compressed sequentially for a column. This is an optimized approach if you plan to perform a sequential scan over your data. Random access to specific attributes in this store is not performing well.

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



Column-Store Deficiencies

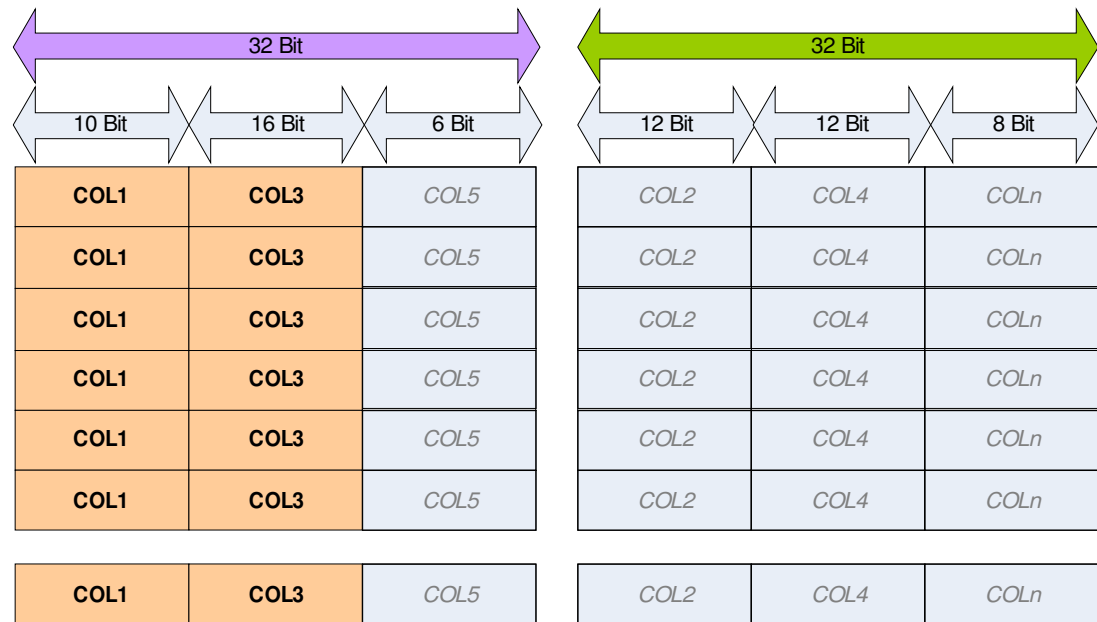


- Updates
- Each column is scanned independently, resulting in multiple result information where predicates did or did not match.
- These scan results from each column need to be combined by ANDing to figure out if all predicates did match (tuple did qualify) before measure columns can be accessed for processing.
- This ANDing is a significant processing effort which increases with tuple count and amount of columns.
- The access to measure columns for processing (i.e. aggregation) is a „random access“ which is not performing well on pure columns stores.
- The width of a compressed column is often not matching processor architectures.



Register-Store

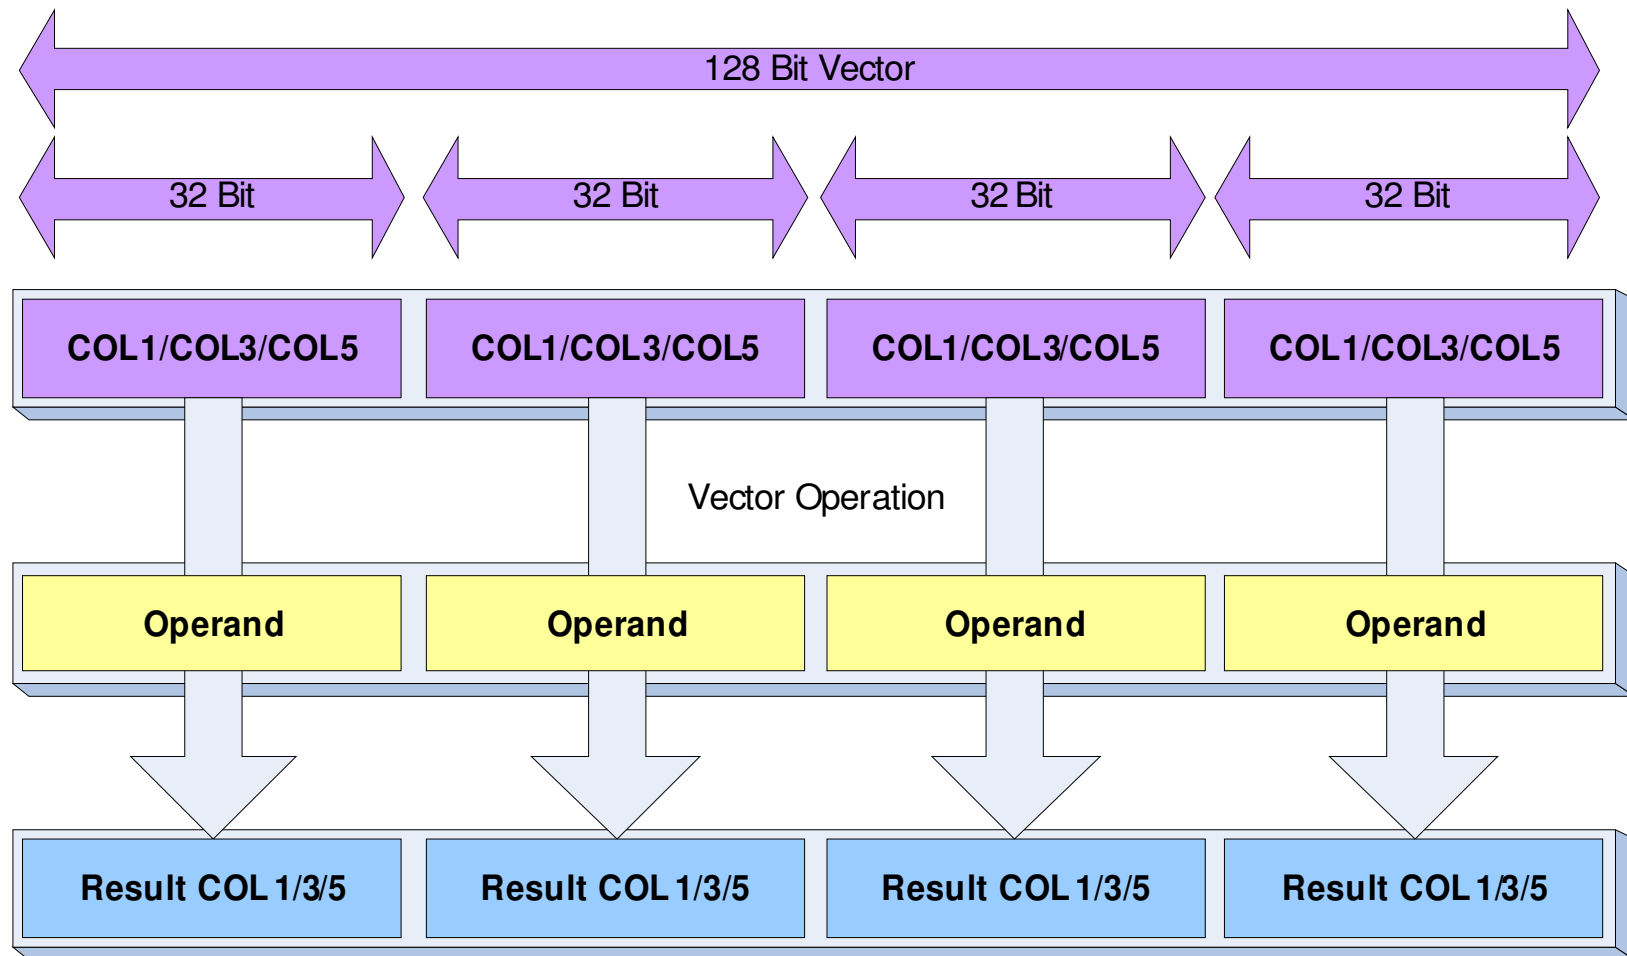
- Within a **Register – Store**, several columns are grouped together.
- The sum of the width of the compressed columns doesn't exceed a register compatible width. This could for example be 32 or 64 bit for a 64 bit system. It doesn't matter how many columns are placed within the register – wide data 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.

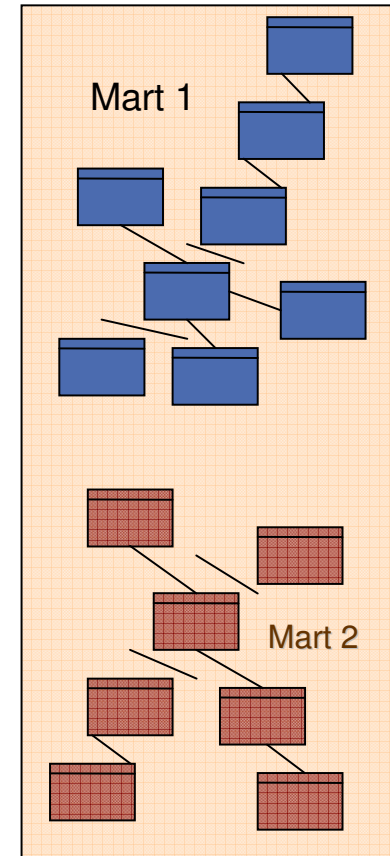
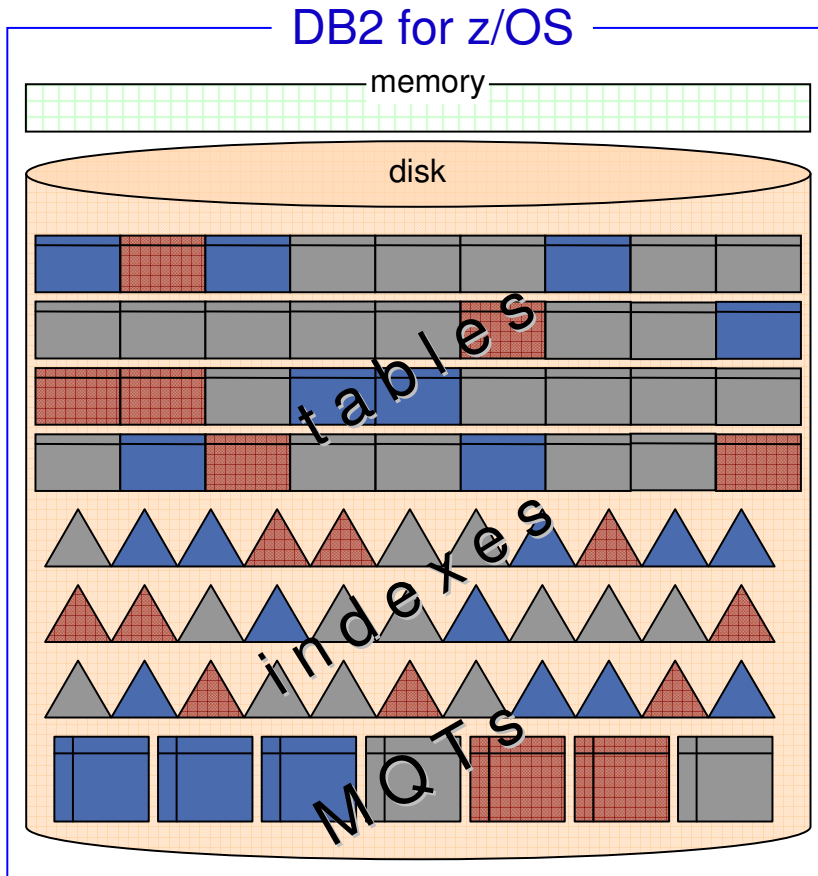


Single Instruction Multiple Data Paradigm



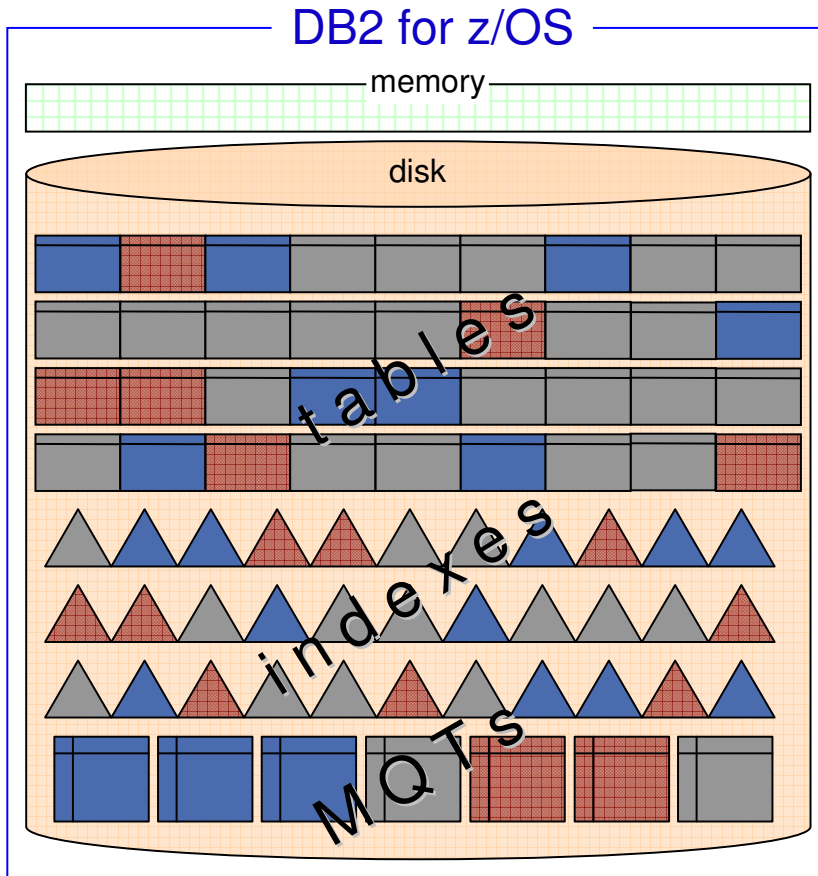


Data View

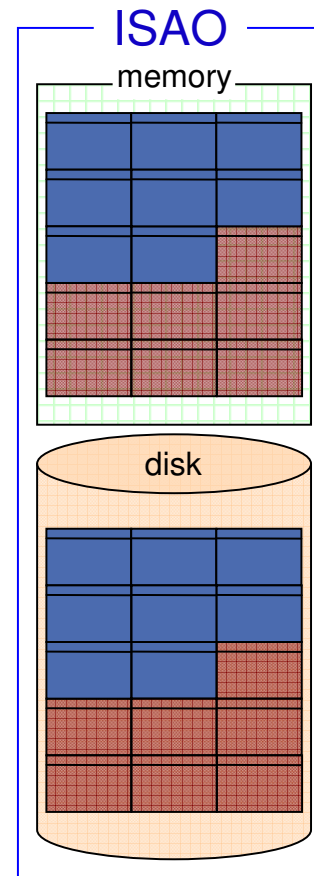




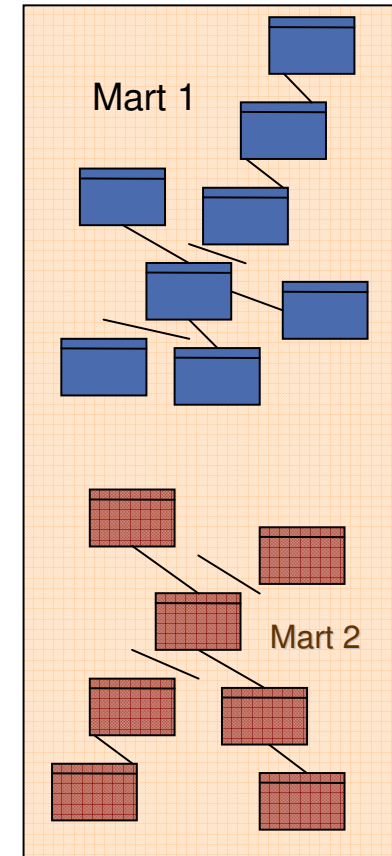
Data View: Marts – Redundant Sets of Memory Resident Tables



- 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

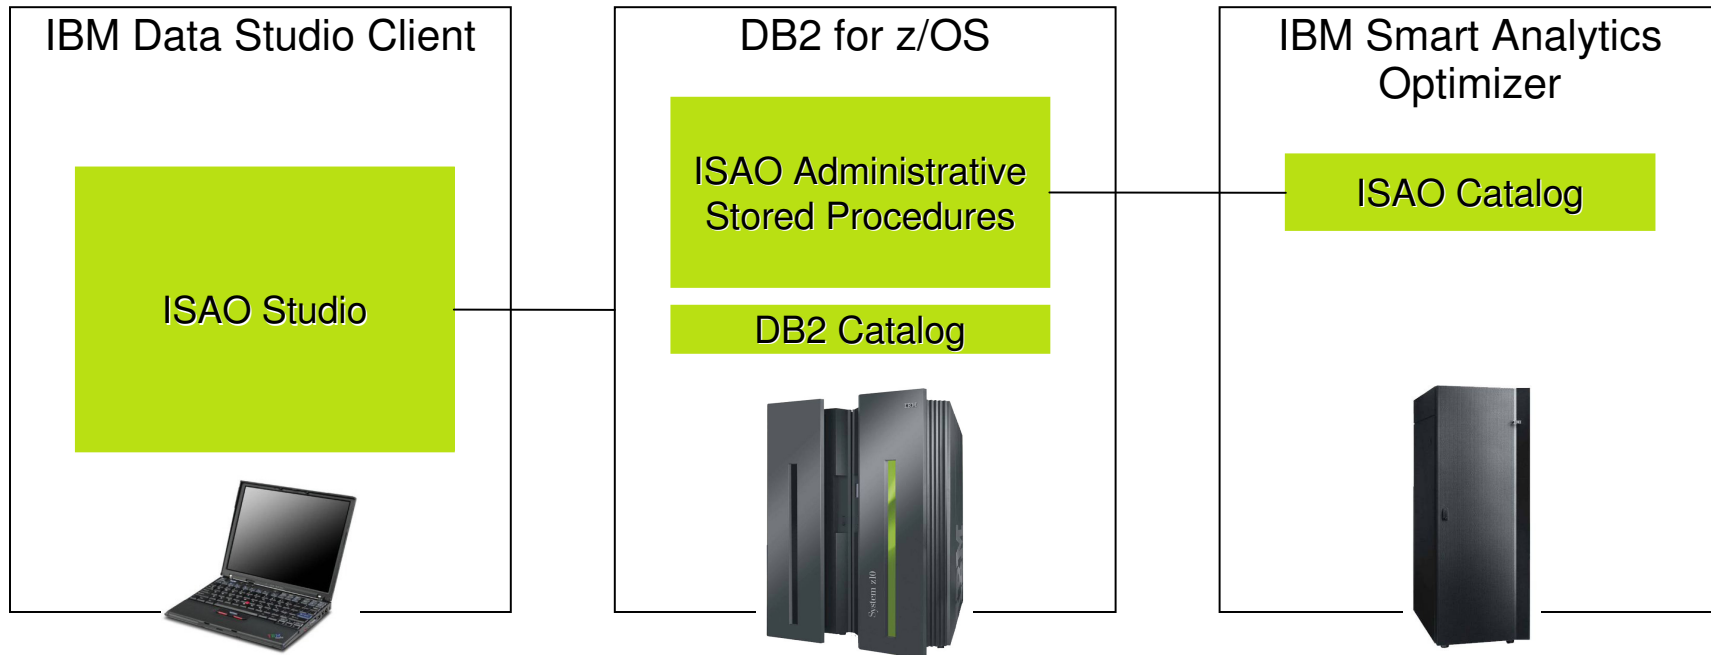


- ISAO contains fully memory resident, compressed *copies* of performance critical tables grouped into logically connected *mart*s, bound ideally by star schema constraints
- Similar but much broader than MQTs: no column projections, no row restrictions, no row aggregations
- Typical usage: Data Marts,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





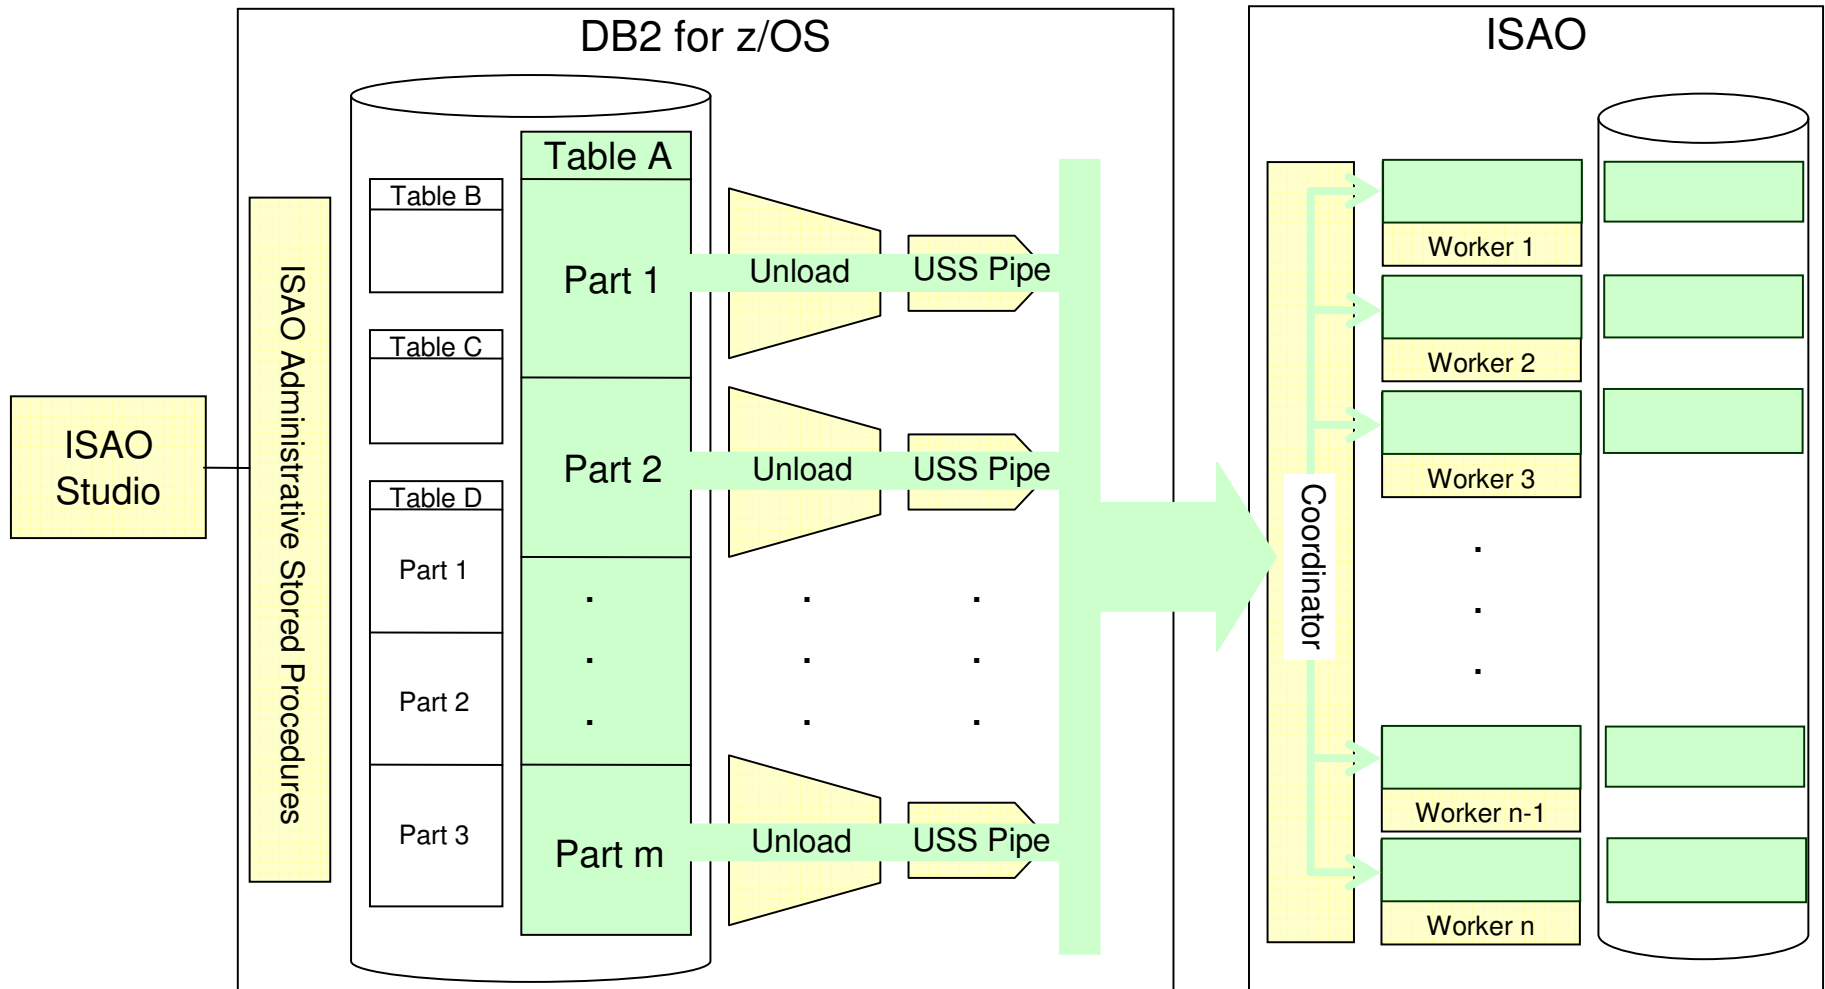
ISAO Mart Definition and Deployment



- 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

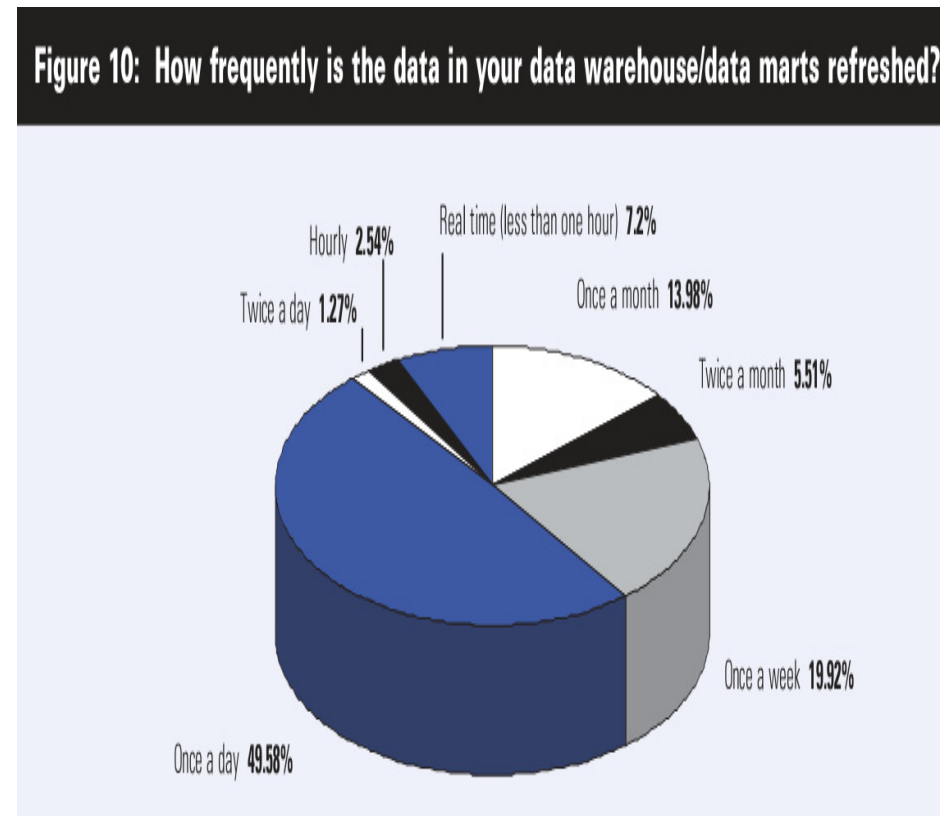
- ISAO will over time phase-in support for all the typical operations in this order
 1. Full table reload
 2. Updated partition reload
 3. Individual row change

- 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

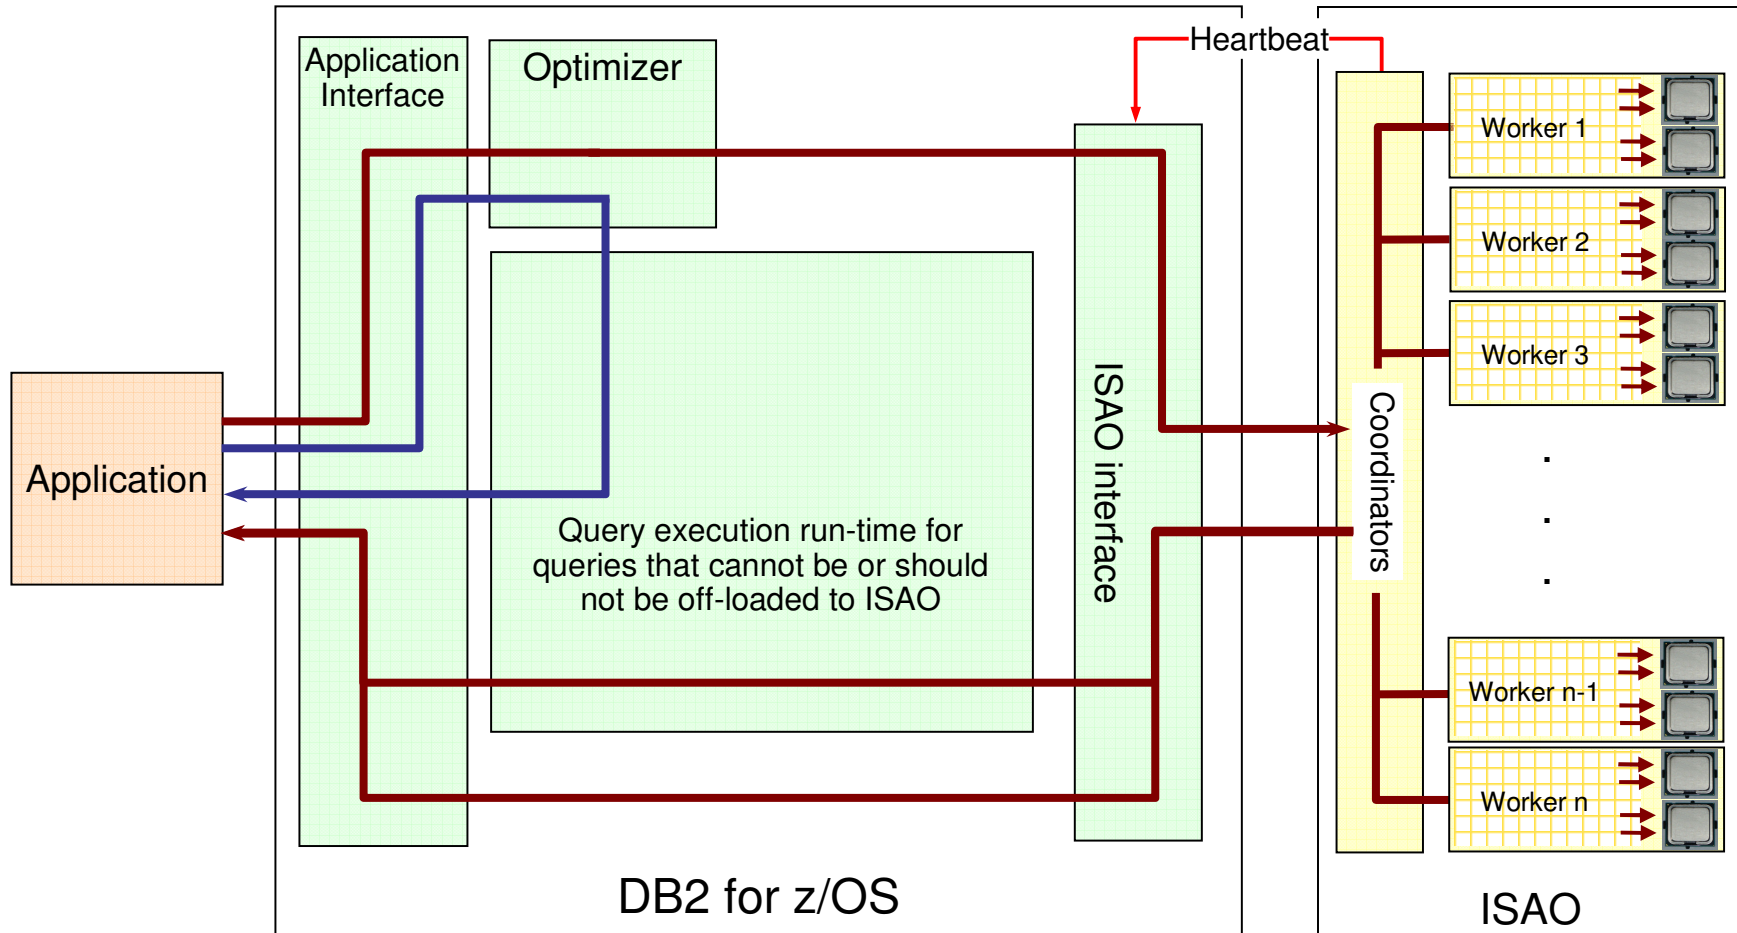
Figure 10: How frequently is the data in your data warehouse/data marts refreshed?



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



Query Execution Process Flow



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



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

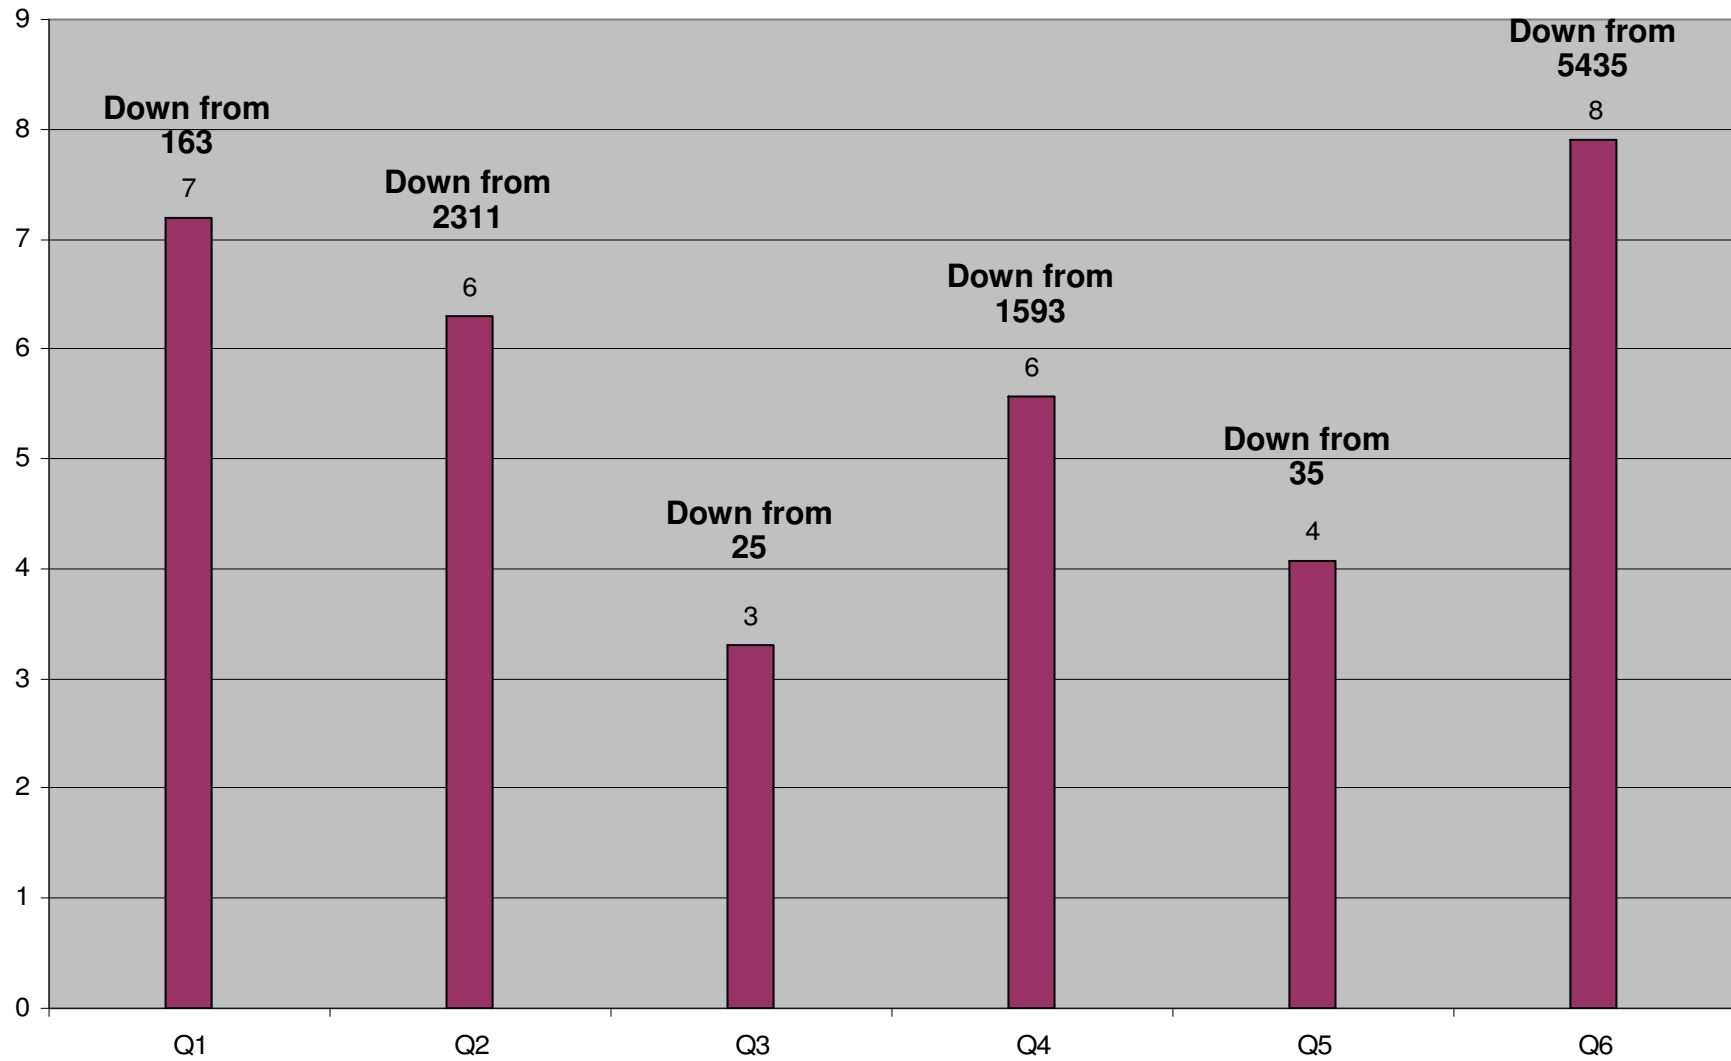
- The problem queries provided by a customer
- Expert database tuning done on all the queries
 - Q1 – Q6 even after tuning run for too long and consume lots of resources
 - Q7 improved significantly – no ISAO offload is needed
- The table shows elapsed and CPU times measured in DB2 (without ISAO)

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

seconds



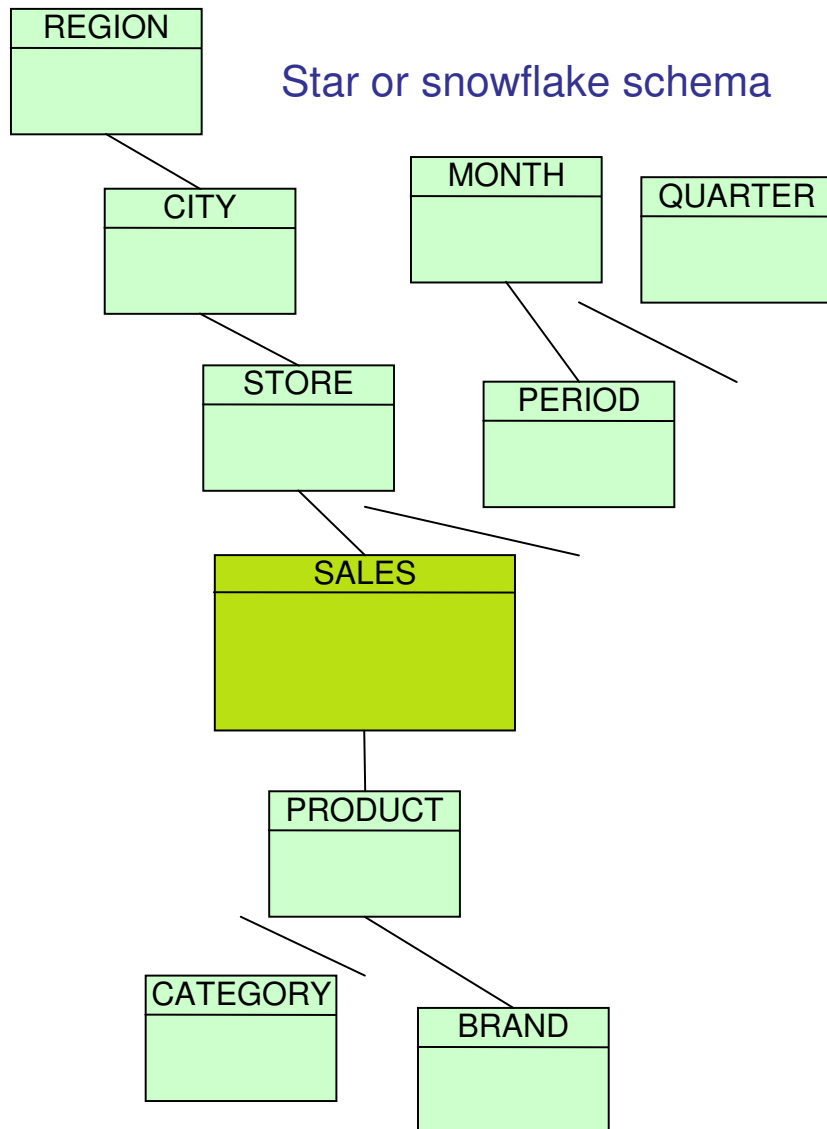


Agenda

- 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?



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
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:

```
SELECT * FROM Nested Table Expression  
(SELECT C1+C2 FROM TA) TX
```

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

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

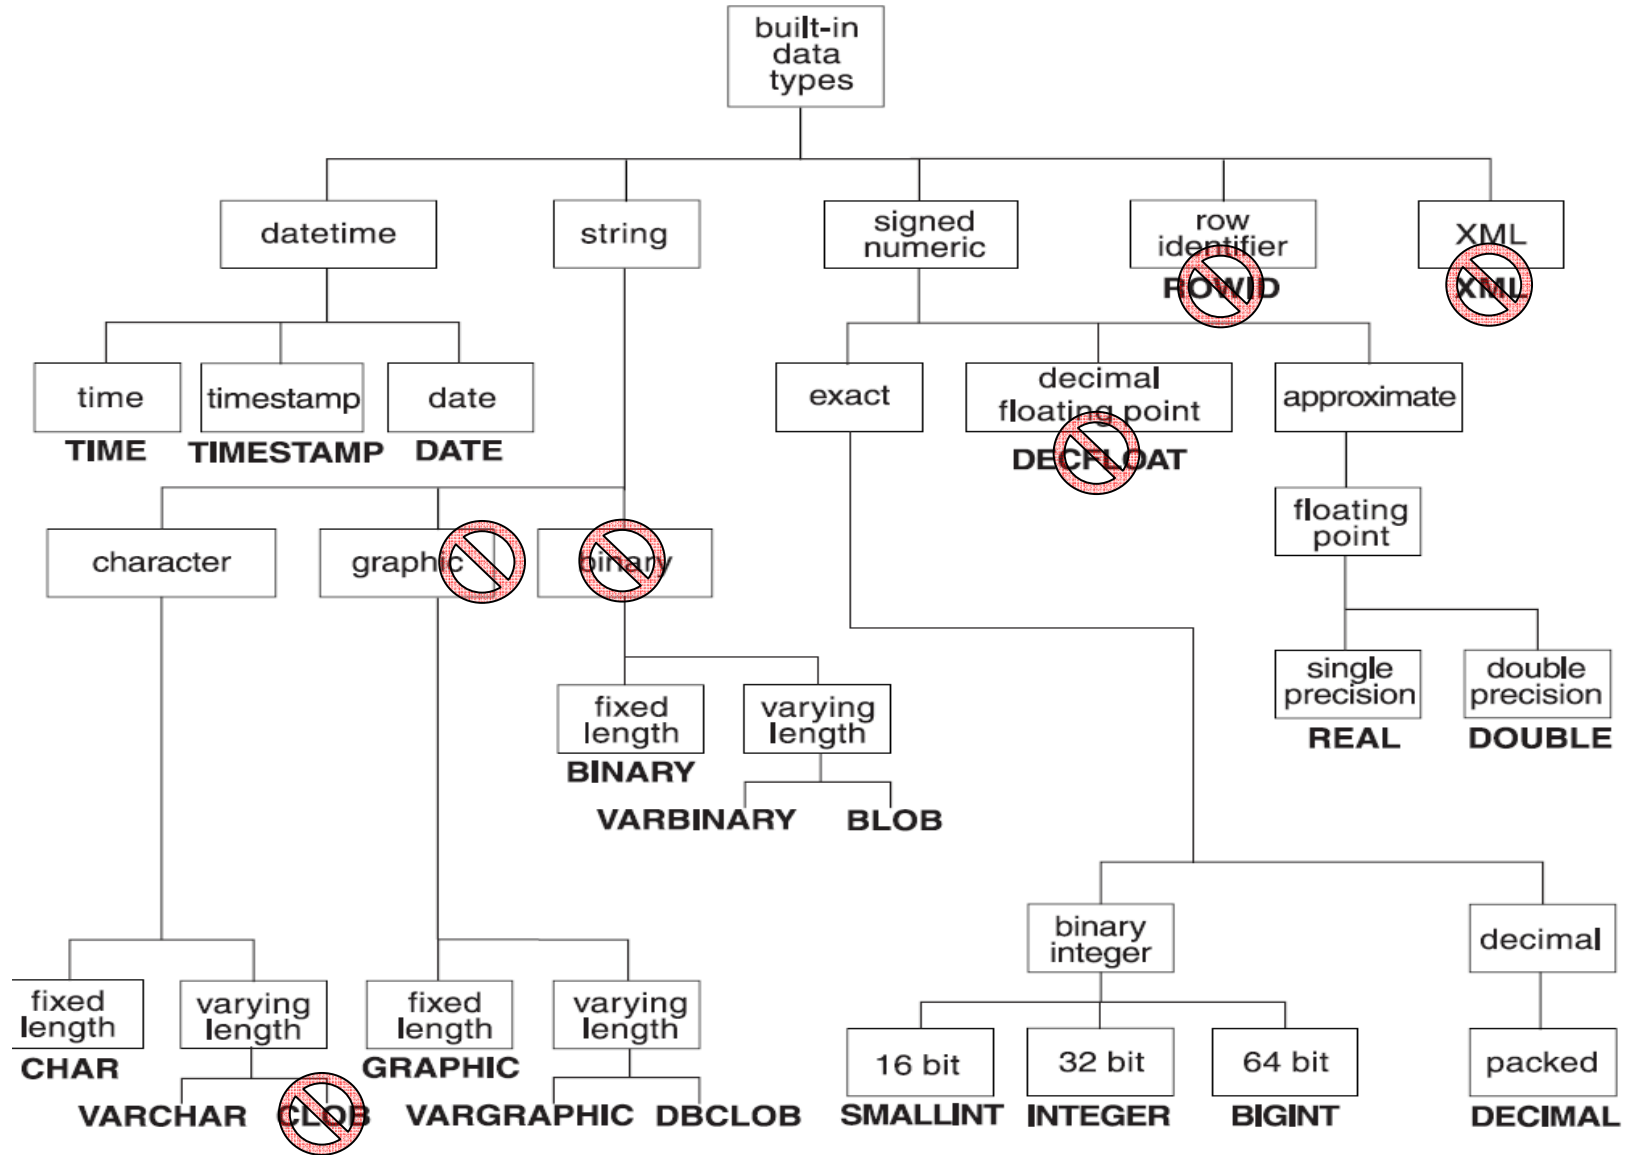


First Release Restrictions

- 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 functions
 - 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	<ul style="list-style-type: none">▪ Initial assessment based on size, query response time, update characteristics and customer pain points
Quick Workload Test	<ul style="list-style-type: none">▪ Assessment based on dynamic customer workload, runtime statistics, table sizes and SQL.
Detailed Online Workload Analysis	<ul style="list-style-type: none">▪ 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

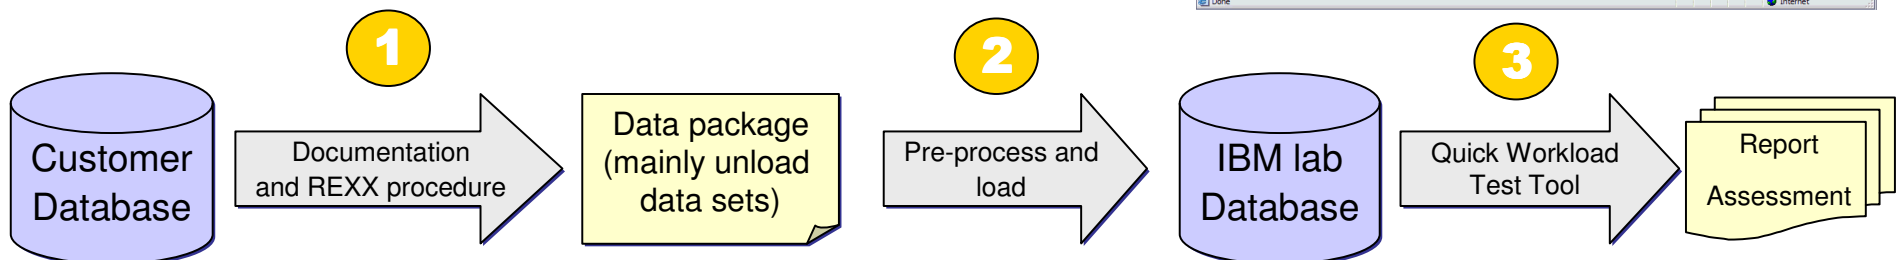
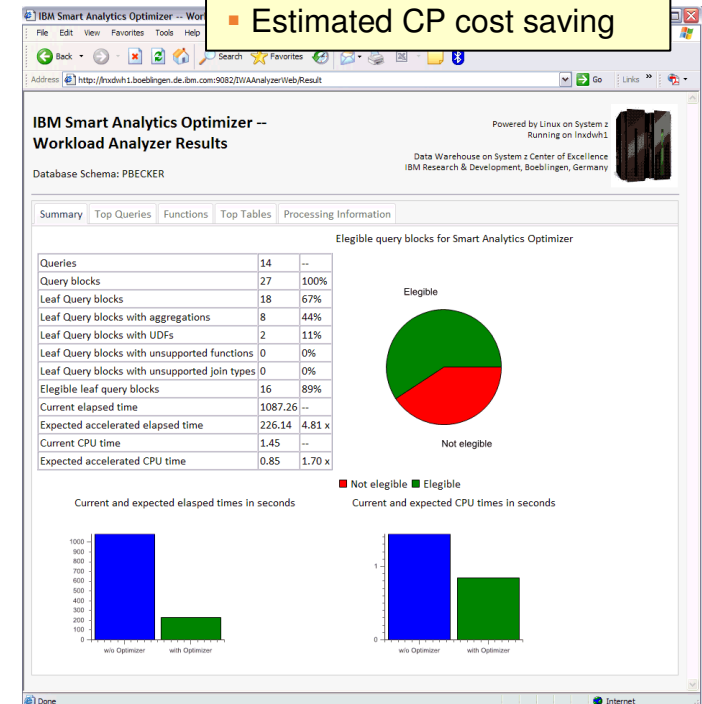
- Collecting information from dynamic statement cache, supported by step-by-step instruction and REXX script (small effort for customer)
- Uploading compressed file (up to some MB) to IBM FTP server

→ IBM / Center of Excellence

- Importing data into local database
- Quick analysis based on known Smart Analytics Optimizer capabilities

Report for a first assessment:

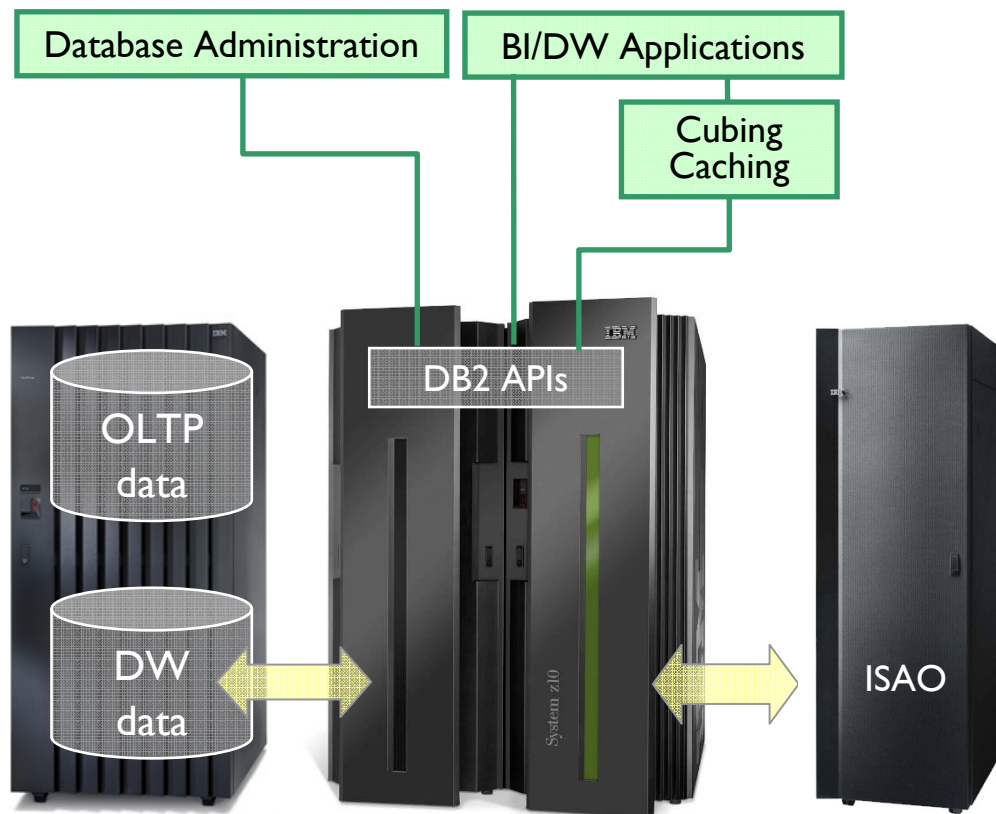
- Query offload potential
- Estimated performance gain
- Estimated CP cost saving





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



THANK
YOU