IBM Smart Analytics Optimizer

for DB2 for z/OS

Willie Favero
Senior Certified Consulting IT Software Specialist
Data Warehouse for System z Swat Team - DB2 SME

Trademarks

- This presentation contains trade-marked IBM products and technologies. Refer to the following Web site:

  http://www.ibm.com/legal/copytrade.shtml

You should also stop in to…

Session 7800

IBM Smart Analytics Optimizer Overview and Beta Experiences

- Bill Schray
- Thursday, 3:00 PM
- Room 303

Topics

- **The IBM Smart Analytics Optimizer**
  - The basics
  - Beta customer experience
  - Under the hood
  - Setting it up
- **Is it right for me?**

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

The IBM Data Warehouse Solution on System z

Ultimate Consolidation Opportunity

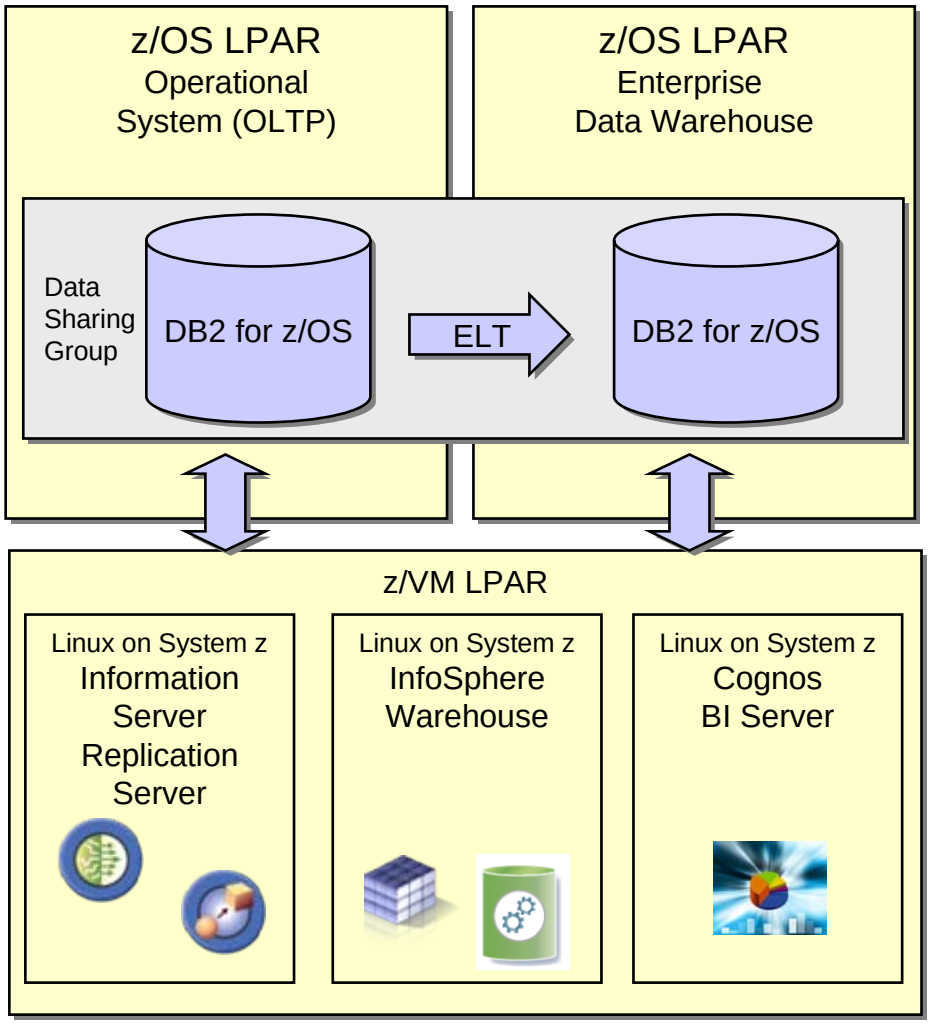

- Consolidation of mission-critical data on System z
- Ability to leverage existing environment, high availability, backup and governance procedures as well as skills
- Efficient data movement within a data sharing group (no network)
- Performance and TCO improvements through cubing services (data marts)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

IBM zEnterprise System – Best in Class Systems and Software Technologies

A system of systems that unifies IT for predictable service delivery

Unified management for a smarter system:
zEnterprise Unified Resource Manager

- Part of the IBM System Director family, provides platform, hardware and workload management
- Unifies management of resources, extending IBM System z® qualities of service across the infrastructure

The world's fastest and most scalable system:
IBM zEnterprise™ 196 (z196)

- Ideal for large scale data and transaction serving and mission critical applications
- Most efficient platform for Large-scale Linux® consolidation
- Leveraging a large portfolio of z/OS® and Linux on System z applications
- Capable of massive scale up, over 50 Billion Instructions per Second (BIPS)

Scale out to a trillion instructions per second:
IBM zEnterprise BladeCenter® Extension (zBX)

- Selected IBM POWER7™ blades and IBM System x® Blades[1] for tens of thousands of AIX® and Linux applications
- High performance optimizers and appliances to accelerate time to insight and reduce cost
- Dedicated high performance private network

IBM Smart Analytics Optimizer – a Virtual DB2 Component

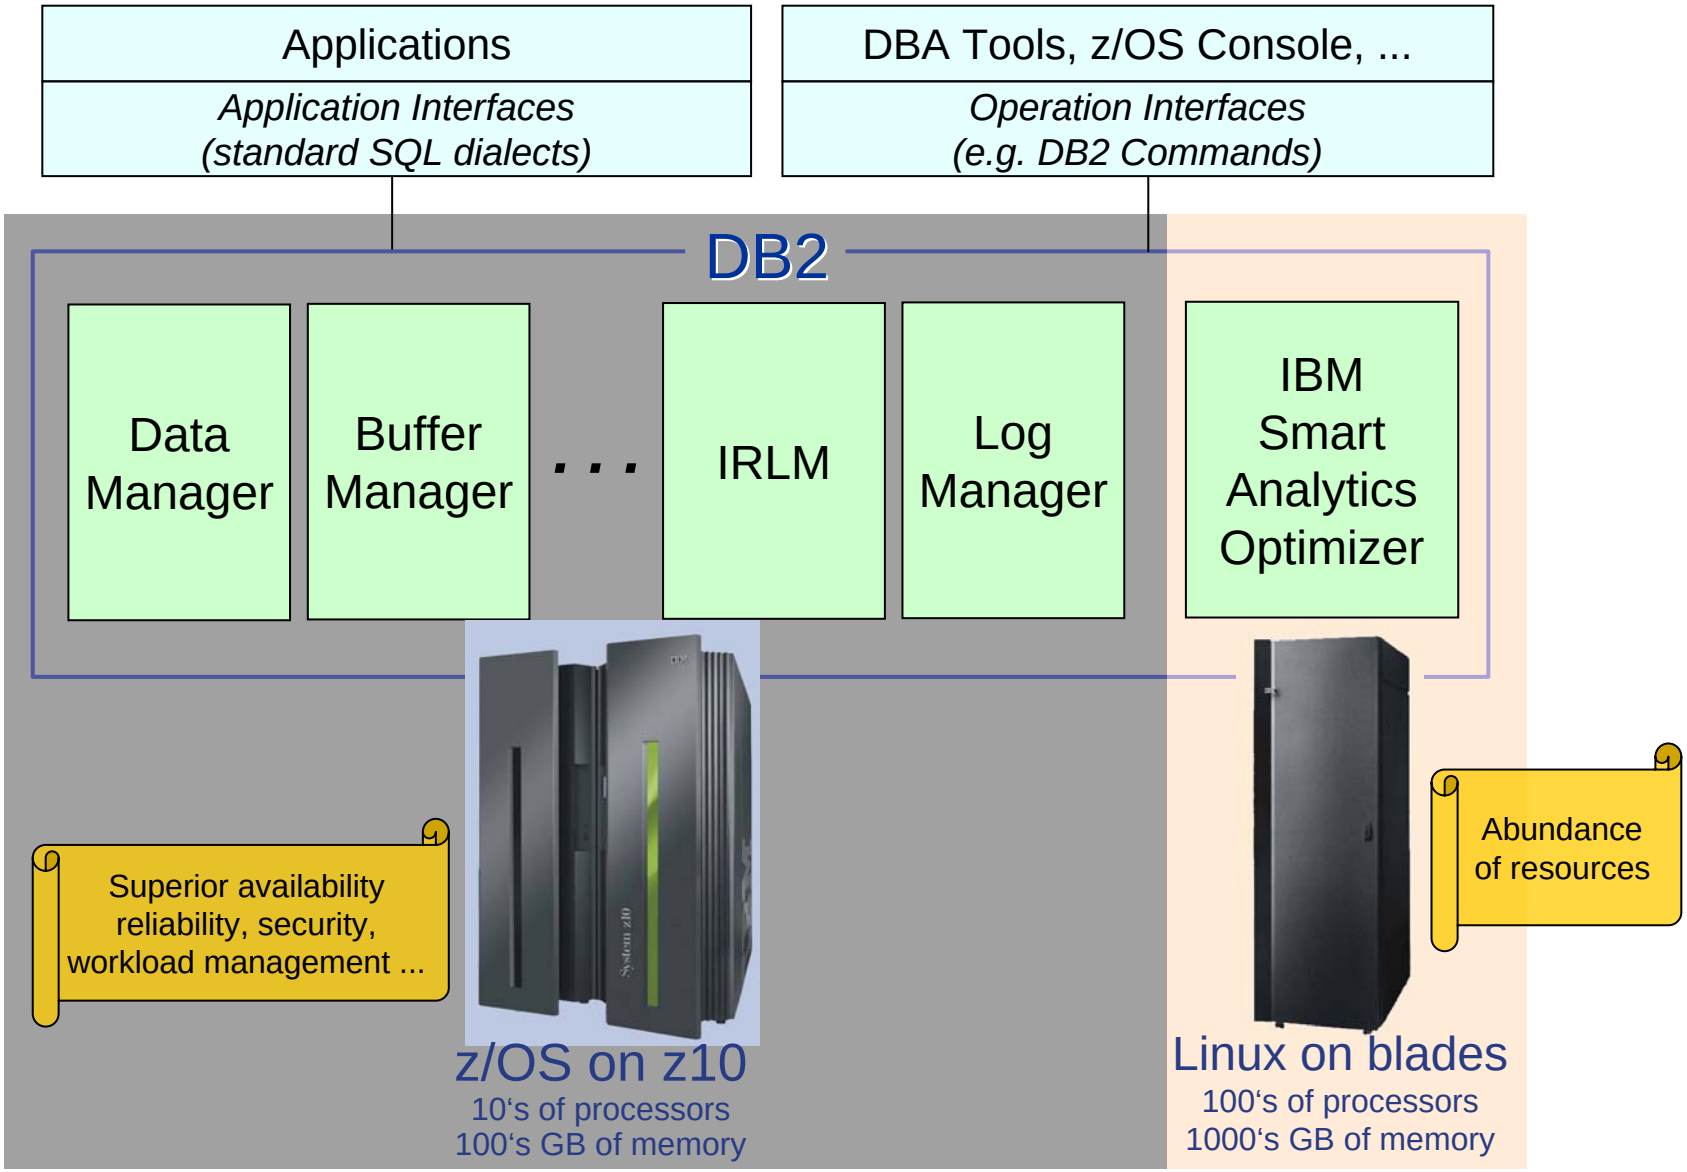

Characteristics of the IBM Smart Analytics Optimizer

- A special purpose, network attached blades system
- No changes to the applications
  - Applications continue to attach to DB2
  - When applicable query needs to be executed DB2 exploits the accelerator transparently to the applications
  - Full fencing and protection of DB2 against possible accelerator failures
- Order of magnitude performance improvement
- Reducing need for tedious tuning of DB2 (MQTs, indexes, etc.)
- Appliance-like form-factor
  - Hands free operations
- Significantly improved price/performance and TCO as a combined effect of:
  - Offloading very CPU intensive operations
  - Orders of magnitude performance improvement for offloaded queries
  - Reduced DBA effort for tuning offloaded queries
- Hybrid technology
  - Enabling Dynamic DW and Operational BI
  - Preserving traditional DB2 quality of service
  - Having transactional and analytical workload being handled by DB2

Data View: Marts – Redundant Sets of Memory Resident Tables

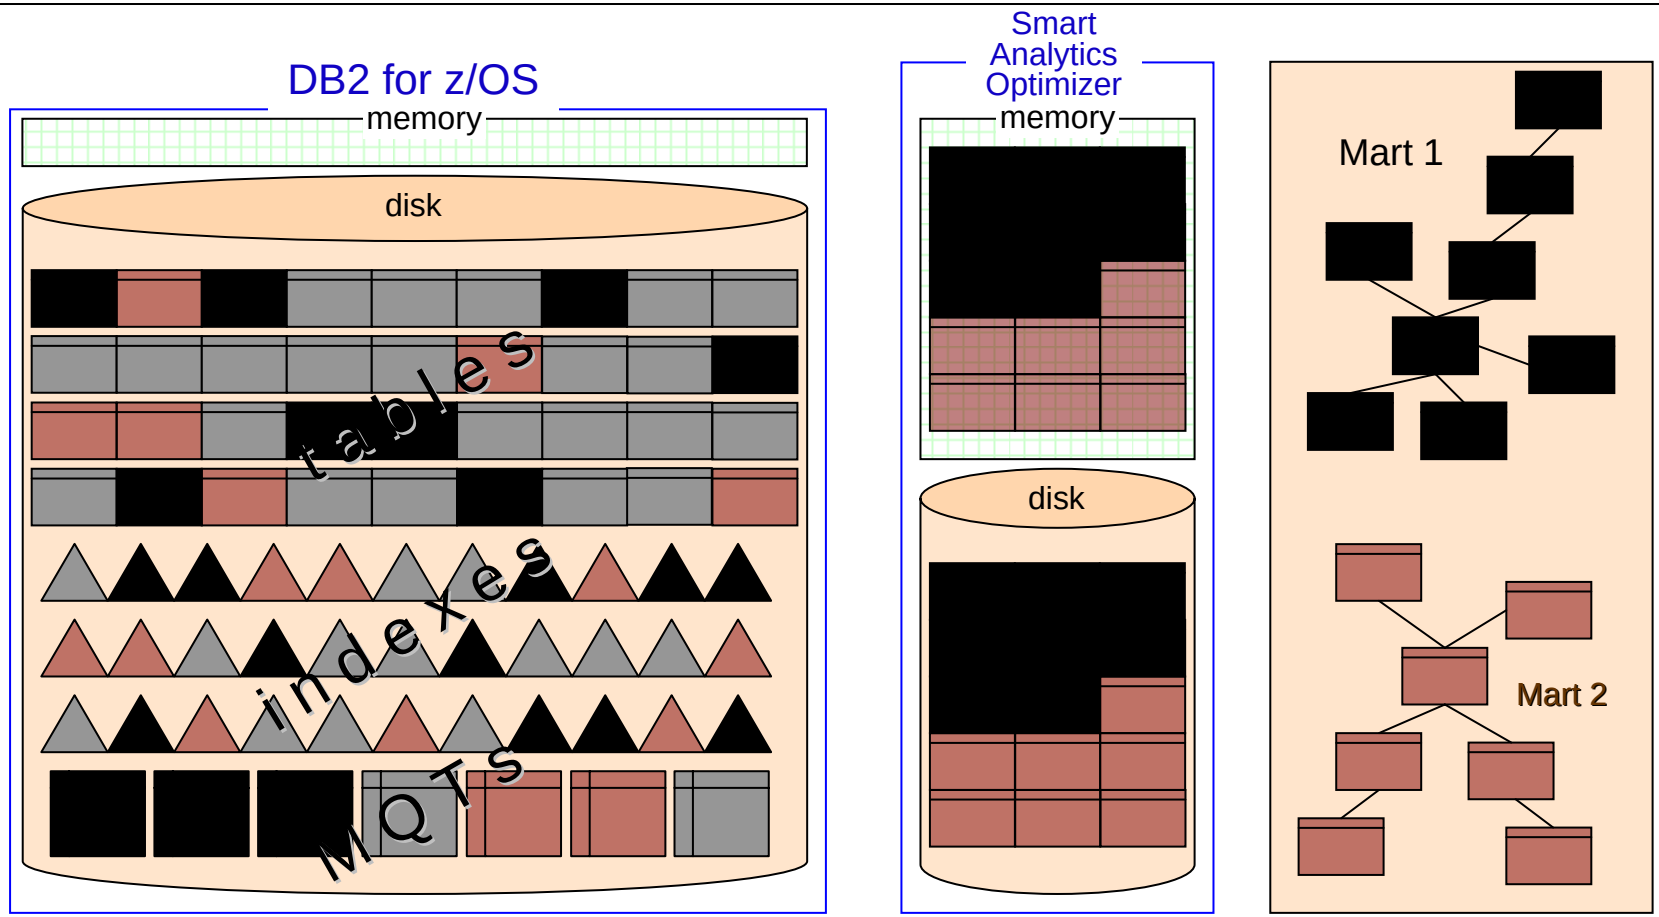

- DB2 continues to *own* and manage all data
- Access performance is influenced by traditional tuning mechanisms such as indexing, MQTs, aggregates, ...
- Typical usage: Enterprise Data Warehouse, large Data Marts

- Smart Analytics Optimizer 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

Smart Analytics Optimizer: Mart Definition and Deployment

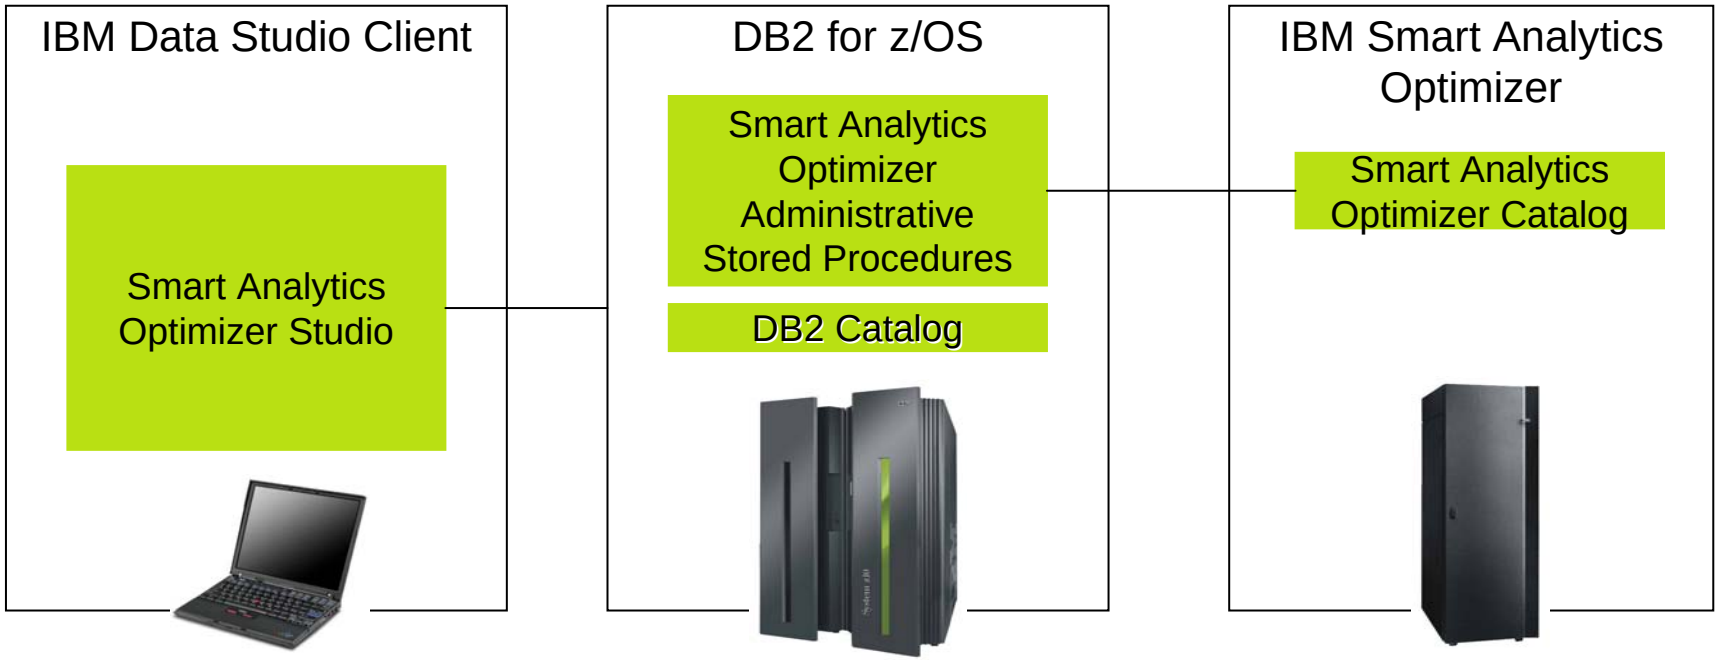

- Smart Analytics Optimizer marts need to be defined and deployed to Smart Analytics Optimizer before data is loaded and queries sent to Smart Analytics Optimizer for processing.
  - Definition: identifying tables and relations that make up marts.
  - Deployment: making marts known to DB2, i.e. storing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

Browse the Mart With Table Relations

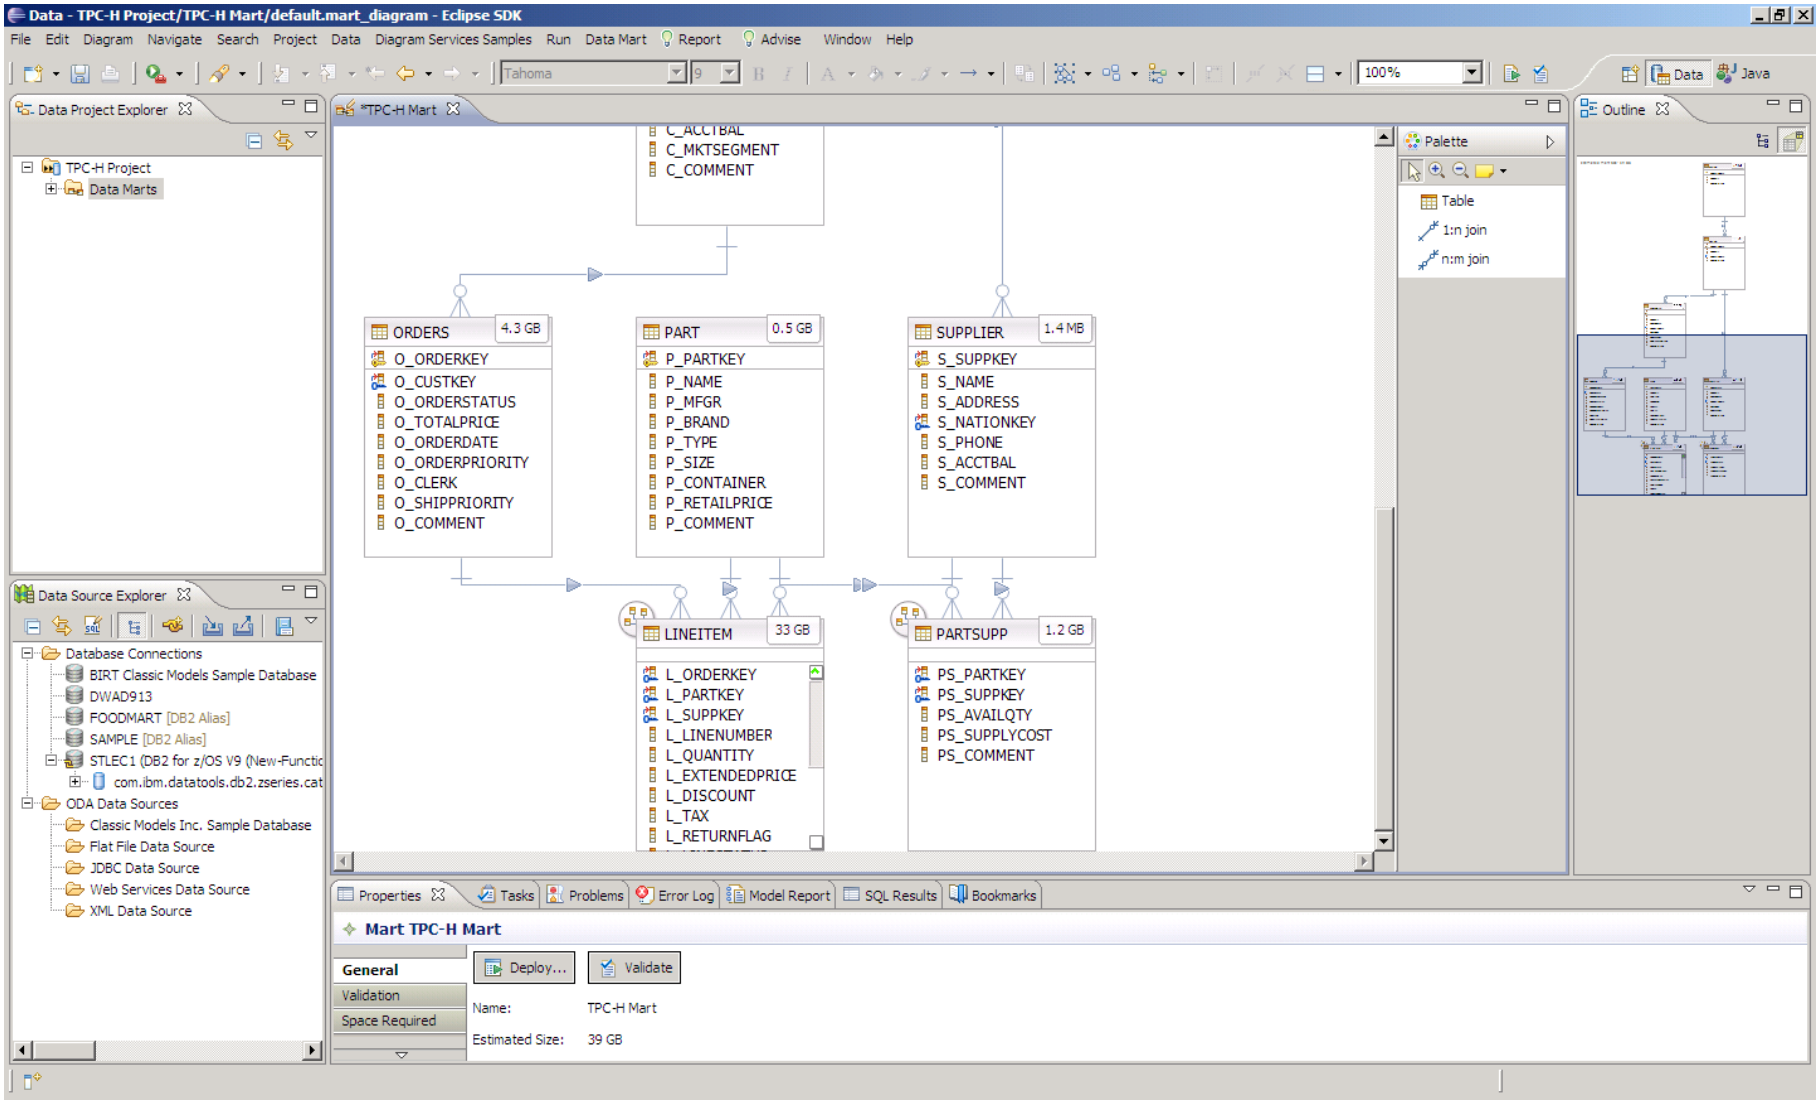

Zoom In: Size Estimates and Fact table property

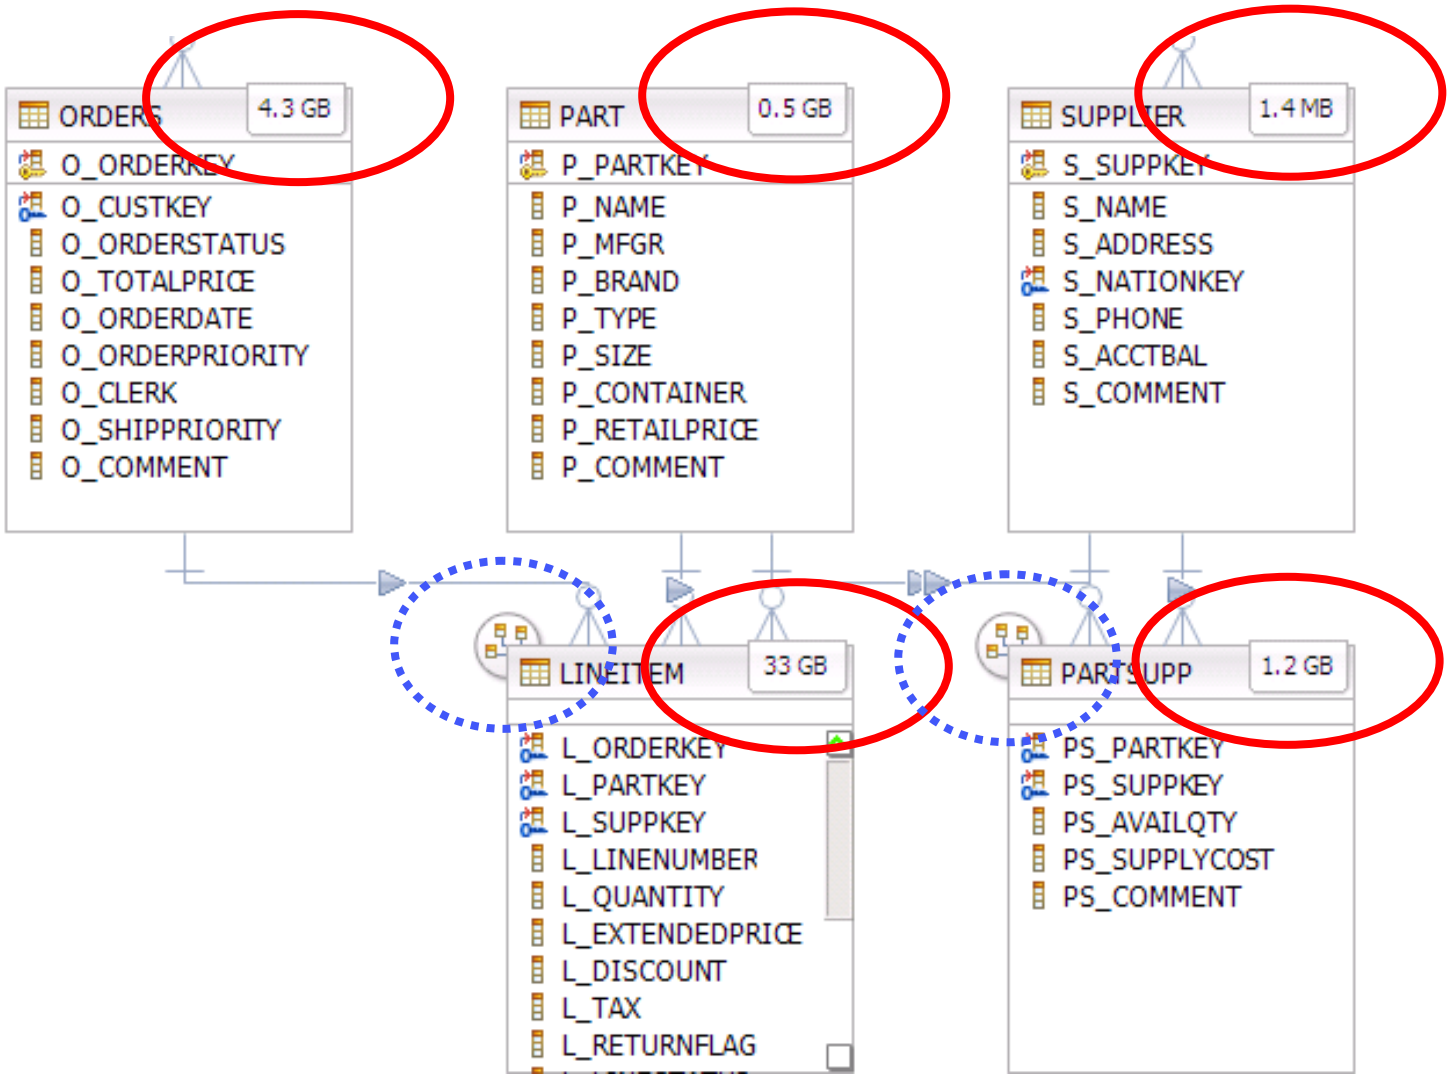

Smart Analytics Optimizer Mart Load

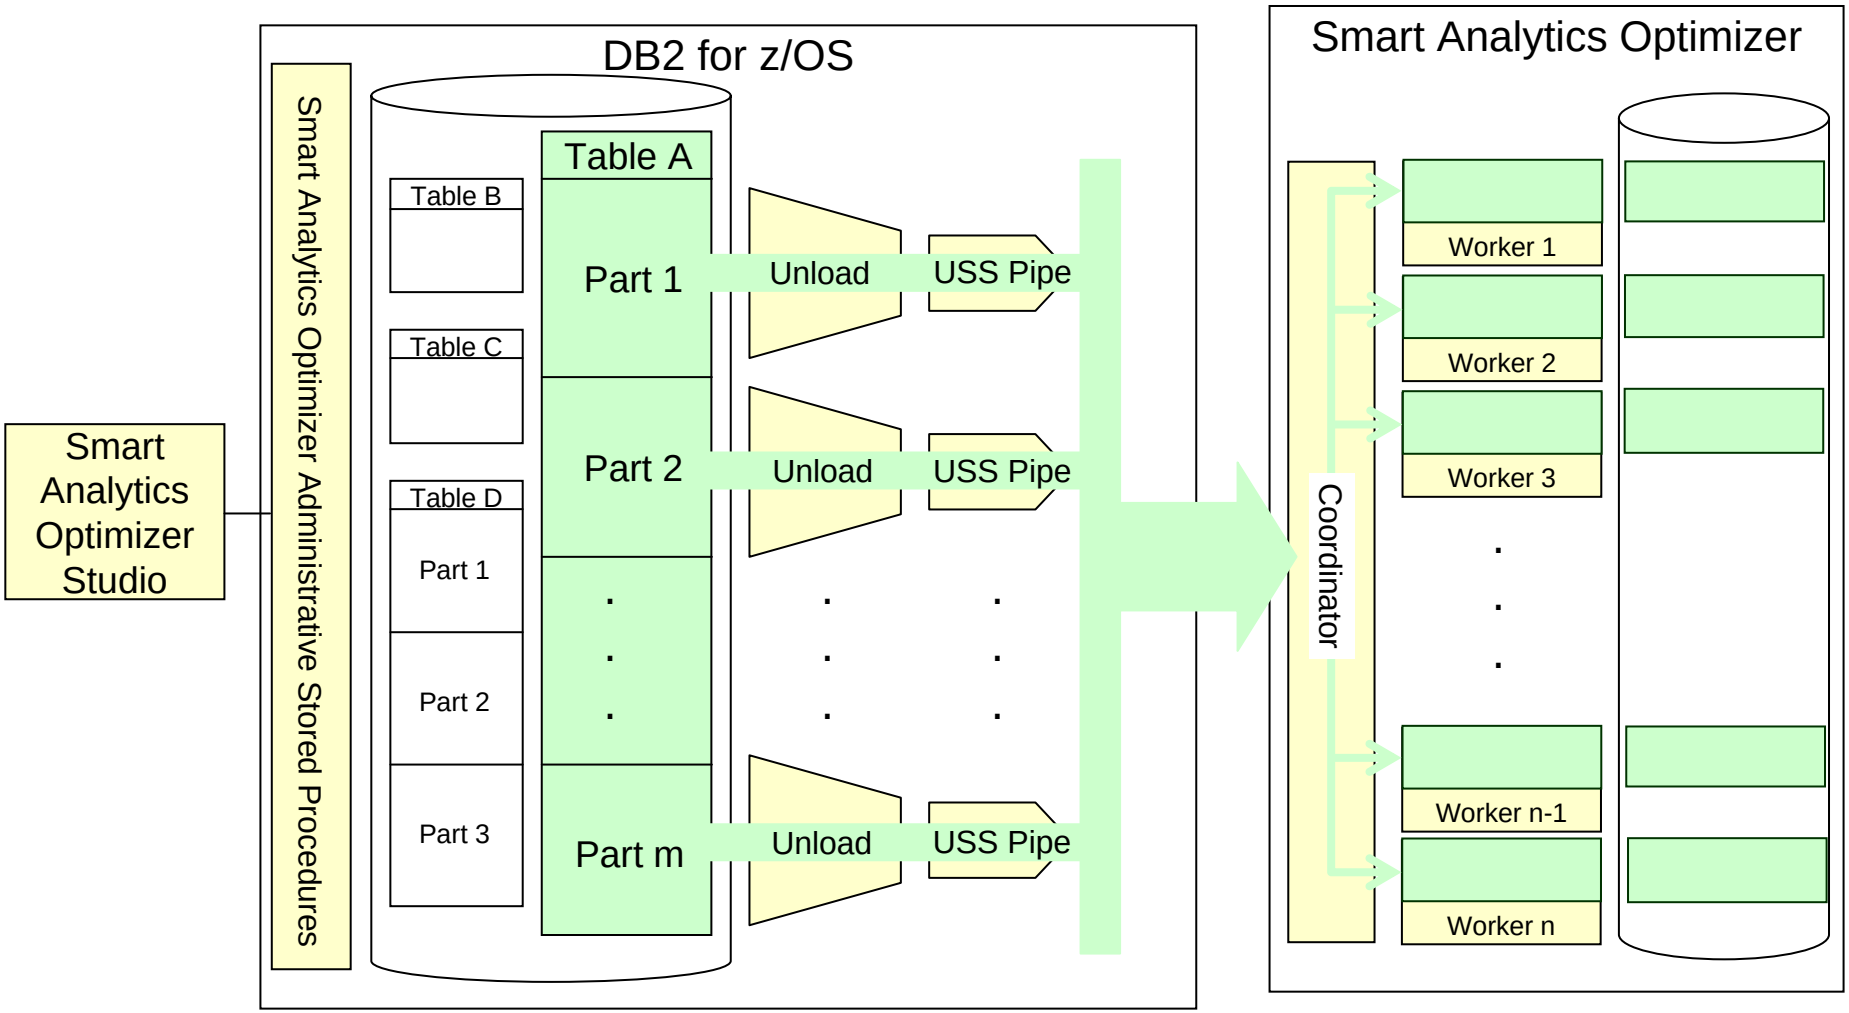

Defining, which data to accelerate

- A MART is a logical collection of tables which are related to each other. For example all tables of a single star schema would belong to the same MART.
- The administrator uses a rich client interface to define the tables which belong to a MART together with the information about their relationships.
- DB2 for z/OS creates definitions for these MARTs in the own catalog. The related data is read from the DB2 tables and transferred to the Smart Analytics Optimizer.
- The accelerator transforms the data into a highly compressed, scan optimized format which is kept locally (in memory) on the Accelerator.

Query Execution Process Flow

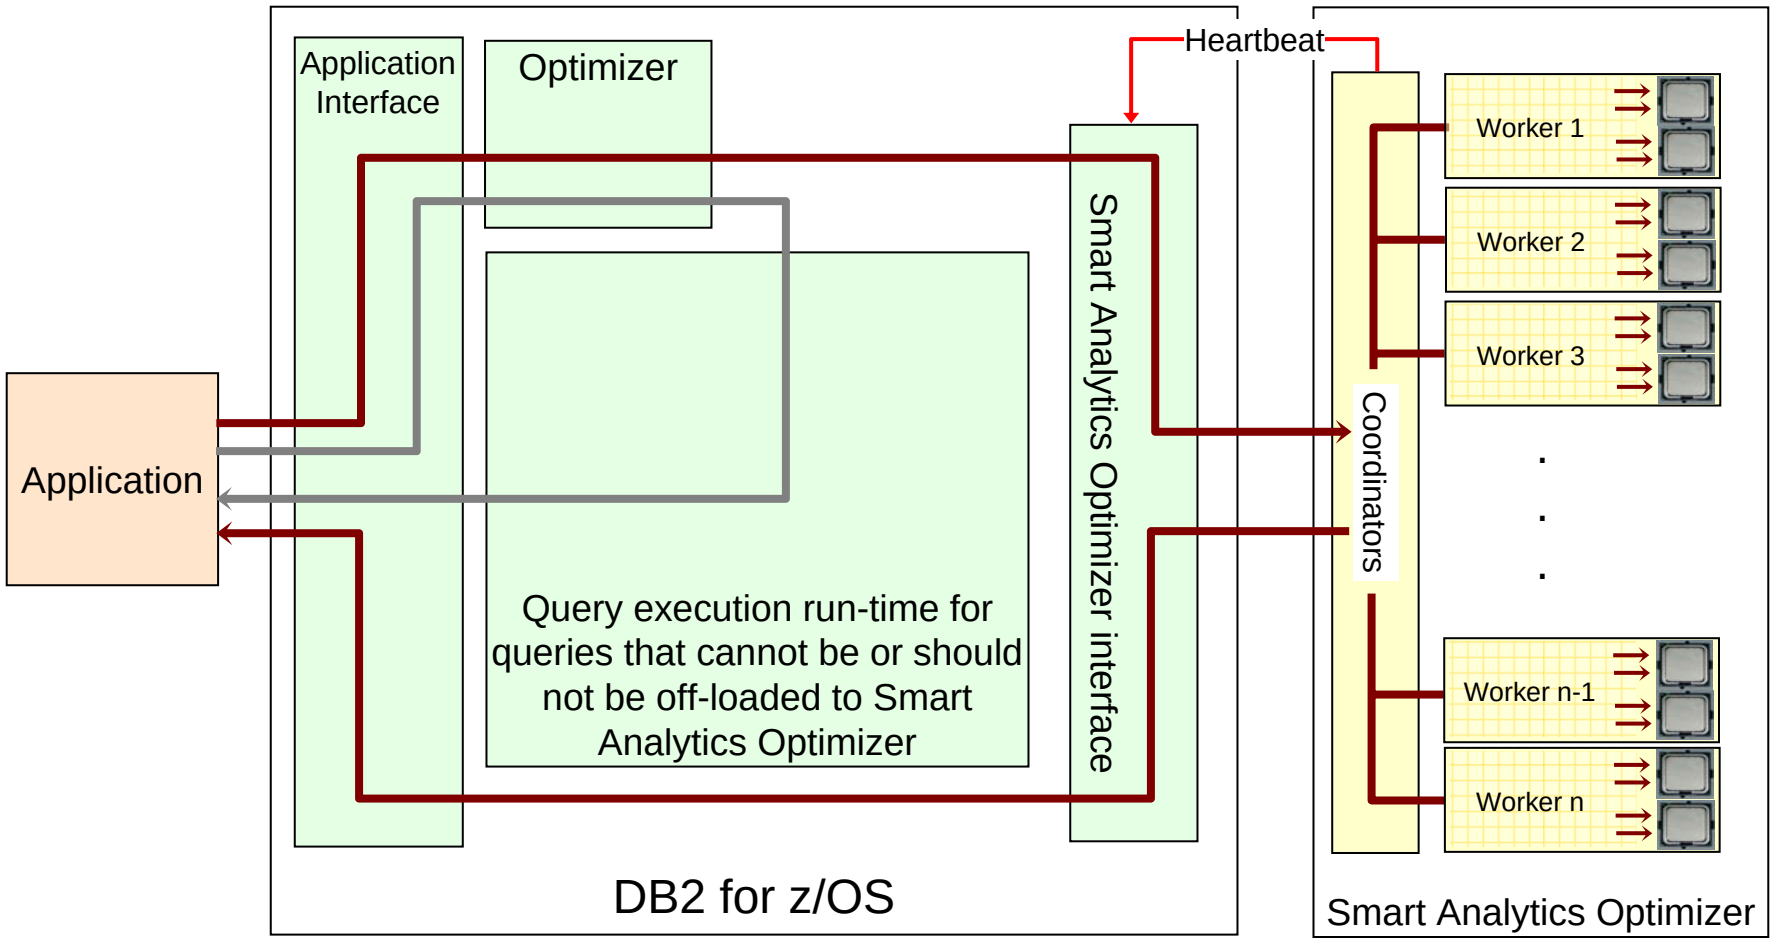

Heartbeat (Smart Analytics Optimizer availability and performance indicators)

Queries executed **without** Smart Analytics Optimizer

Queries executed **with** Smart Analytics Optimizer

IBM Smart Analytics Optimizer Overview

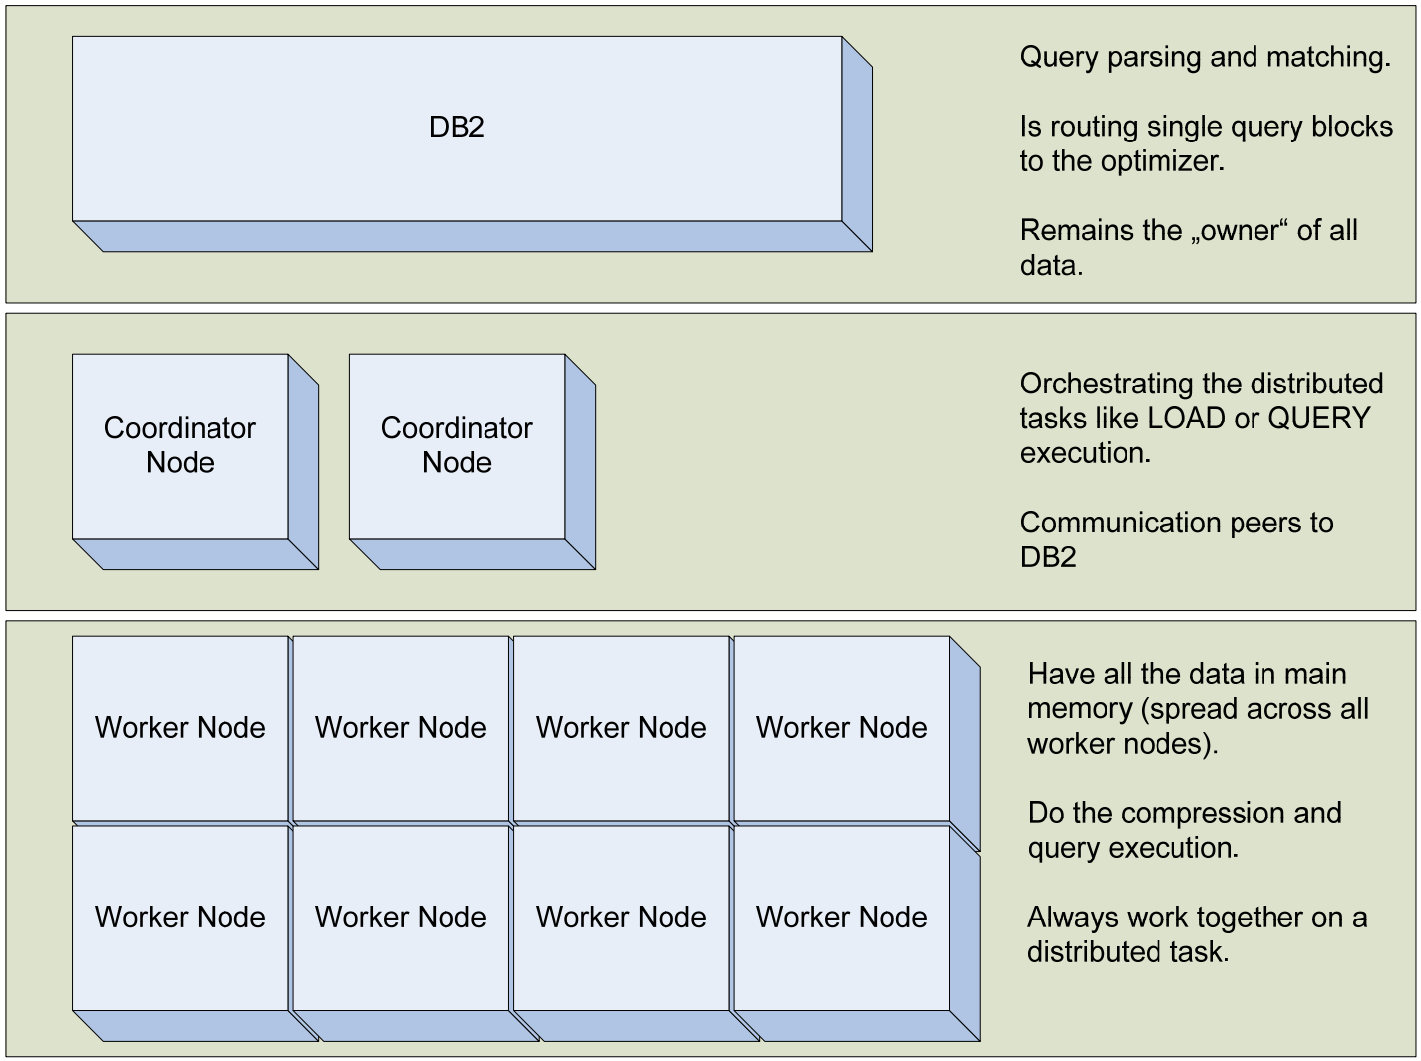

Smart Analytics Optimizer Mart Update

IDUG study on DW update frequency

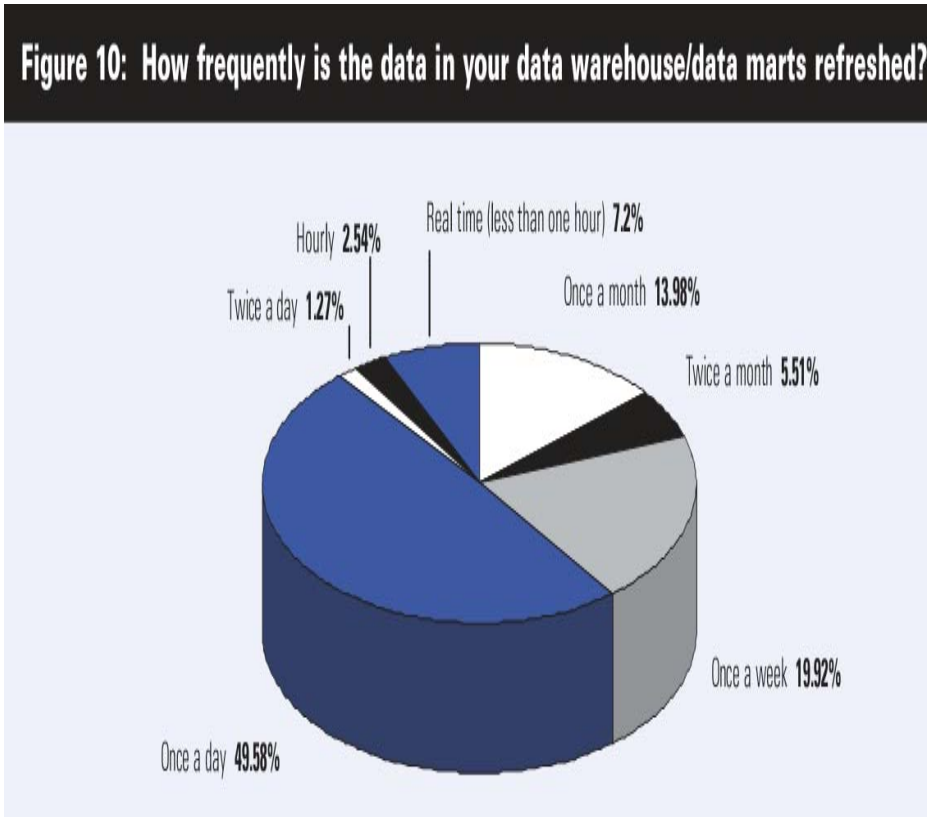

In 90% of cases DW is updated once a day or even less frequently

- Typical DW update operations:
  - LOAD RESUME and REPLACE
  - ADD and ROTATE PARTITION
  - SQL INSERT, UPDATE, DELETE
  - Delete complete partition or table
  - TRUNCATE TABLE
- Smart Analytics Optimizer 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 Smart Analytics Optimizer Studio
- Queries off-loaded to Smart Analytics Optimizer 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

Additional DB2 Support - EXPLAIN

- **DSN_QUERYINFO_TABLE**
- Whether or not a query block is eligible for automatic query rewrite, and if not eligible show the reason why it's not eligible.
- If eligible for automatic query rewrite, which materialized/accelerated query tables were considered, and for each one that wasn't chosen the reason why it was not chosen.
  - REASON_CODE – why AQT matches and why not
  - QI_DATA – part of SQL statement that causes failing match

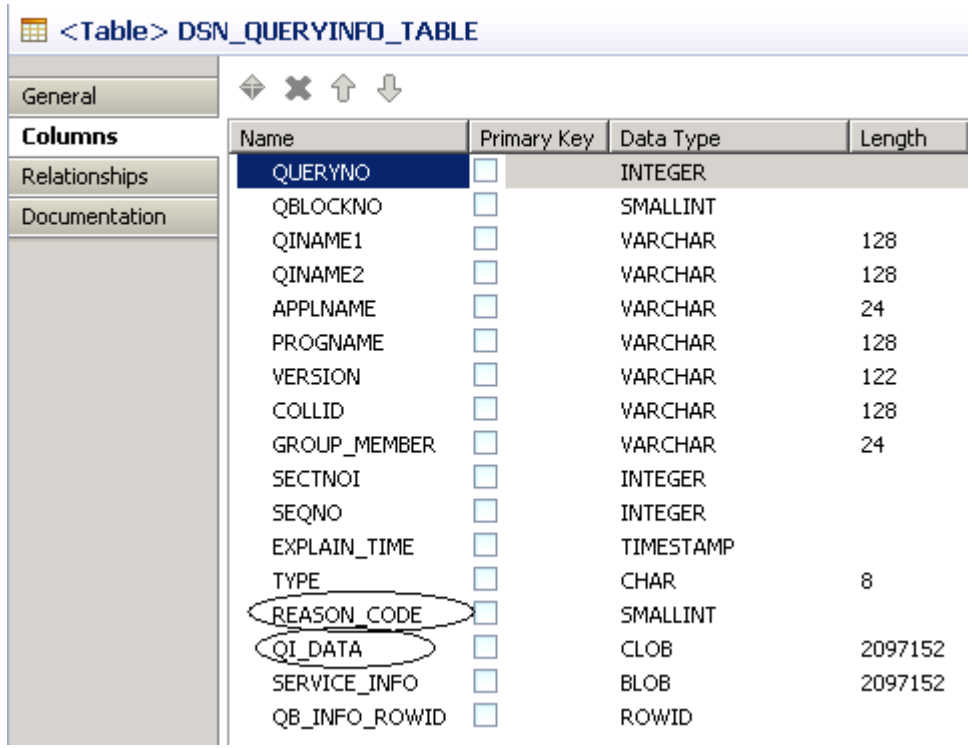

Additional DB2 Support - Commands

Stop Accel-Command "-STO ACCEL(IWA001)"

```
DSNX810I  -D2PAA DSNX8CMD STOP ACCEL FOLLOWS -
DSNX860I  -D2PAA DSNX8STO STOP ACCELERATOR SUCCESSFUL FOR IWA01
DSN9022I  -D2PAA DSNX8CMD '-STOP ACCEL' NORMAL COMPLETION
```

Start Accel-Command "-STA ACCEL(IWA001)"

```
DSNX810I  -D2PAA DSNX8CMD START ACCEL FOLLOWS -
DSNX820I  -D2PAA DSNX8STA START ACCELERATOR SUCCESSFUL FOR IWA01
DSN9022I  -D2PAA DSNX8CMD '-START ACCEL' NORMAL COMPLETION
```

Additional DB2 Support - DISPLAY

Display Accel-Command "-DIS ACCEL(IWA001) DETAIL(*) LIST(*)"

```
DSNX810I  -D2PAA DSNX8CMD DISPLAY ACCEL FOLLOWS -
DSNX830I  -D2PAA DSNX8CDA
ACCELERATOR                      MEMB  STATUS  REQUESTS QUED MAXQ FAIL
-------------------------------- ----  -------- -------- ---- ---- ----
IWA01                            DPAA STARTED        10    0    1    0
PORT=1400 IPNAME=00010601
IPADDR=192.168.150.21 HEALTHY
IPADDR=192.168.150.22 HEALTHY
IPADDR=192.168.150.23 HEALTHY
CUMULATIVE STATISTICS SINCE 14:36:13 OCT  7, 2009
  STATUS = FULLY OPERATIONAL
  AVERAGE QUEUE WAIT                                    =          4
  MAXIMUM QUEUE WAIT                                    =          1
  PROCESSING CAPACITY                                   =    1314278
  TOTAL NUMBER OF PROCESSORS                            =        224
  AVERAGE CPU UTILIZATION ON COORDINATOR NODES          =          2
  NUMBER OF ACTIVE COORDINATOR NODES                    =          3
  PHYSICAL MEMORY AVAILABLE ON COORDINATOR NODES        =      45235 MB
  PHYSICAL MEMORY AVERAGE USAGE ON COORDINATOR NODES    =       3039 MB
  AVERAGE CPU UTILIZATION ON WORKER NODES               =          0
  NUMBER OF ACTIVE WORKER NODES                         =         11
  PHYSICAL MEMORY AVAILABLE ON WORKER NODES             =      45697 MB
  PHYSICAL MEMORY AVERAGE USAGE ON WORKER NODES         =       2577 MB
  SHARED MEMORY DATA AVAILABLE ON WORKER NODES          =       3160 MB
  SHARED MEMORY DATA AVERAGE USAGE ON WORKER NODES      =      33047 MB
  MAXIMUM SHARED MEMORY DATA IN USE ON WORKER NODES     =      33373 MB
DISPLAY ACCEL REPORT COMPLETE
DSN9022I  -D2PAA DSNX8CMD '-DISPLAY ACCEL' NORMAL COMPLETION
```

Orderable Smart Analytics Optimizer Options

DS5020

Smart Analytics Optimizer Offerings	XS .5TB	S 1TB	M 2TB	L 3TB	XL 4TB
Enterprise Rack	1	1	1	2	2
BC-H Chassis, incl. redundant infrastructure (AMM, PDU, TOR 10Gb Ethernet switch for data, TOR 1Gb Ethernet for management, 4Gb F/C switches)	1	1	2	3	4
Blades	7	14	28	42	56
DS5020 with # of SATA disks (1 TB)	16	16	16	32	32
Storage Expansion Units	0	0	0	1	1

- **For z10, order Machine Type 2458 Model 001**
- **For zEnterprise 196, order Machine Type 2458 Model 002**
- **Model 001 is MES upgradable to Model 002, when z10 is upgraded to zEnterprise 196**
- **DSS5020 disk needs to be ordered separately from zBX, and housed in its own rack.**

Testing Results

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

Query	Times measured in DB2 **without** Smart Analytics Optimizer			
	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 IBM Smart Analytics Optimizer

Query	Query Elapsed Time		
	DB2 only	DB2 with Smart Analytics Optimizer	Speed-up
Q1	0:02:43.0	0:00:03.4	48
Q2	0:38:31.0	0:00:04.5	511
Q3	0:00:25.0	0:00:02.2	12
Q4	0:26:33.0	0:00:07.8	206
Q5	0:00:35.0	0:00:08.3	4
Q6	1:30:35.0	0:00:03.8	1424
Q7	0:00:02.0	0:00:02.0	1
Total	2:39:24.0	0:00:32.0	298

What Is Smart Analytics Optimzer Ideally Suited For?

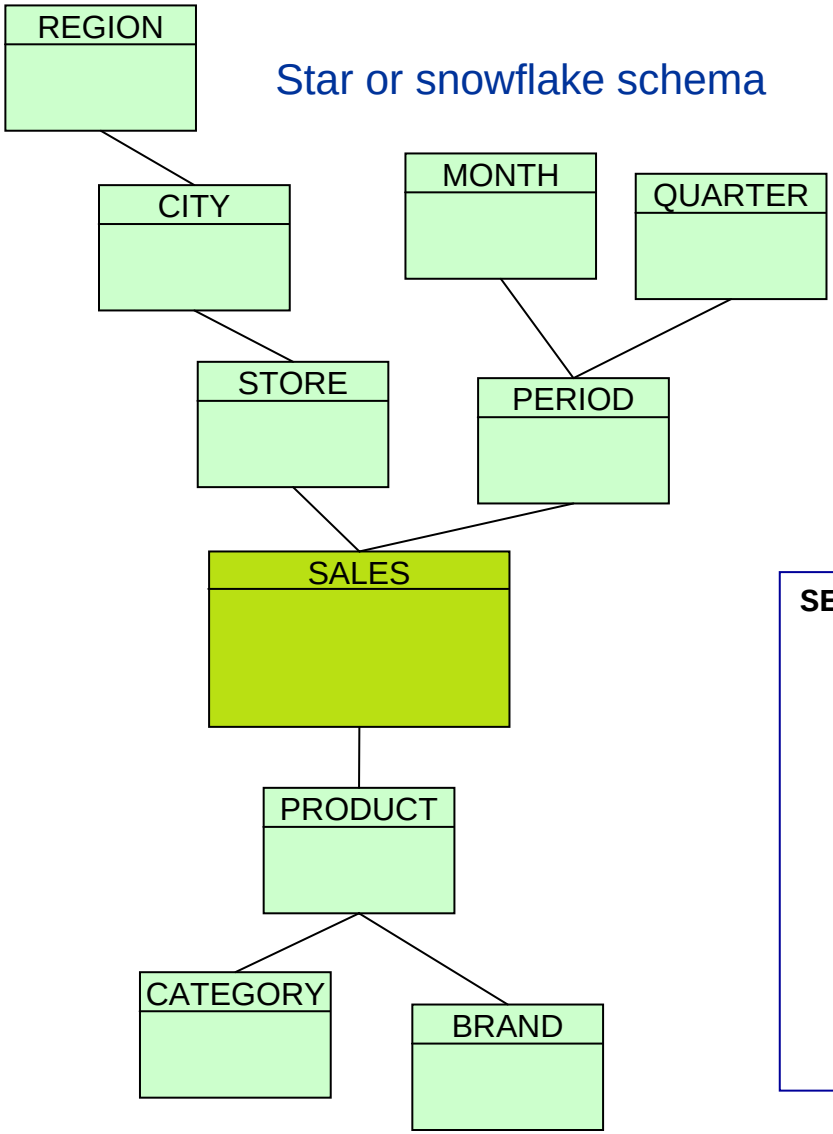

Star or snowflake schema

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
```

The “Sweet – Spot” Schema

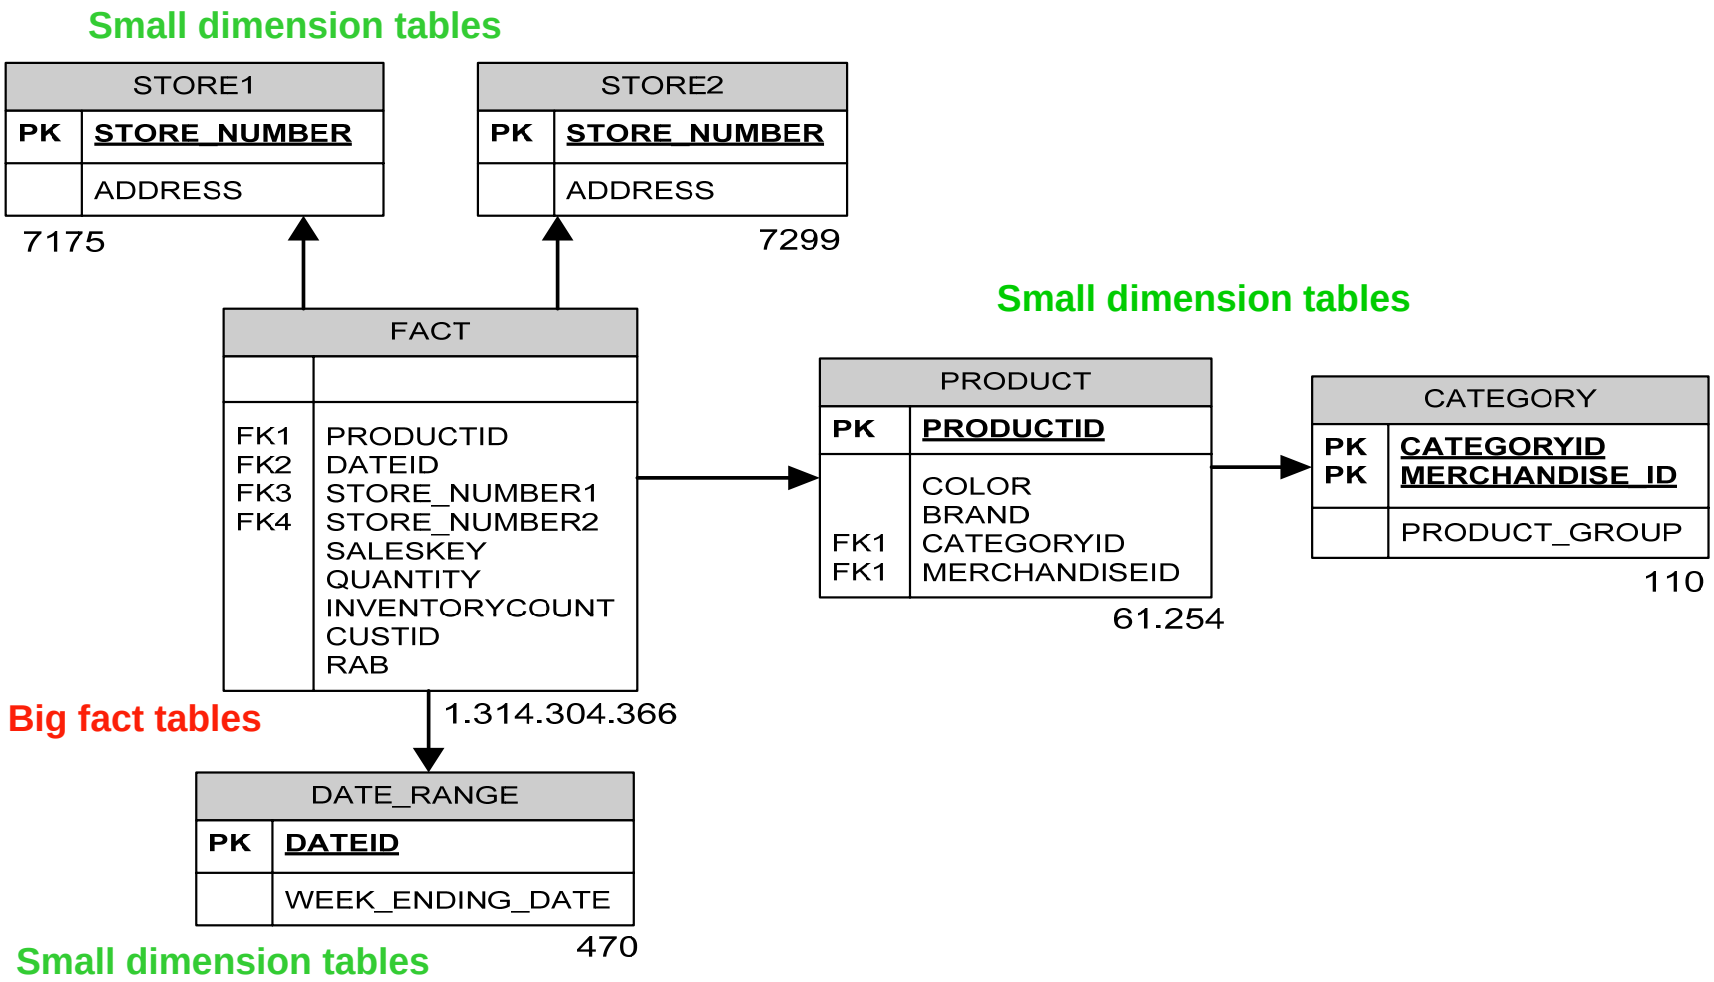

Join Strategy within the Cluster

- The Fact table is split into multiple parts and distributed evenly across the Worker nodes within the cluster.
  - Bigger Fact tables “just“ require enough Worker nodes to contain the compressed data in memory.
- The Join Strategy between Dimension Tables and the Fact table data is always a collocated join.
  - This means that all dimension tables are fully replicated to each of the worker nodes.
  - Space requirements for dimension tables therefore needs to be multiplied with cluster size (amount of Worker Nodes)

Matching of queries for Smart Analytics Optimizer

- DB2 for z/OS reuses partial MQT functionality to find out which queries are eligible for Smart Analytics Optimizer offload and which are not.
- This implies that a subset of the MQT restrictions is inherited
  - Only a single query block at a time can be routed to the Optimizer (Queries which consist of several Query Blocks, are not seen as whole query by the accelerator but only as single, independent blocks)
  - The results of subqueries can not be used by the Optimizer in the outer query (DB2 would need to pass a subselect result to Smart Analytics Optimizer)

Examples with potentially multiple query blocks

Derived table (nested table expression)

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

Derived table (common table expression=CTE)

```
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) ))
```

Supported schemas

- A MART consists of a set of tables together with their referential constraints.
  - Fact tables are considered to be the tables which have the highest join depth.
- Only Queries, including at least the fact table, can be routed to Smart Analytics Optimizer (Queries which are only scanning the dimensions have to be handled by DB2)
- Multiple fact tables are allowed within the same MART definition but:
  - Queries can not handle table across MART boundaries

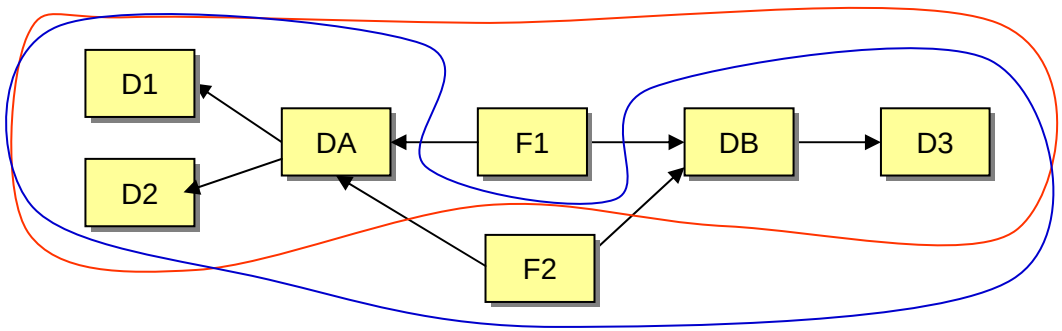

Summary why a query may not be routed

- Because it uses static SQL
- Because it uses CURRENT REFRESH AGE = 0
- Because it contains syntax that is not supported (e.g. Subselect or full outer join)
- Because the accelerator or MART are disabled
- Because it references a table or column that is not in the accelerated mart (may be due to unsupported datatypes)
- Because the query does not reference a fact table
- Because the optimizer decides DB2 for zOS can do better (DB2 has a cost-based threshold)
  - E.g. Query with selective predicate on indexed column is executed in DB2

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

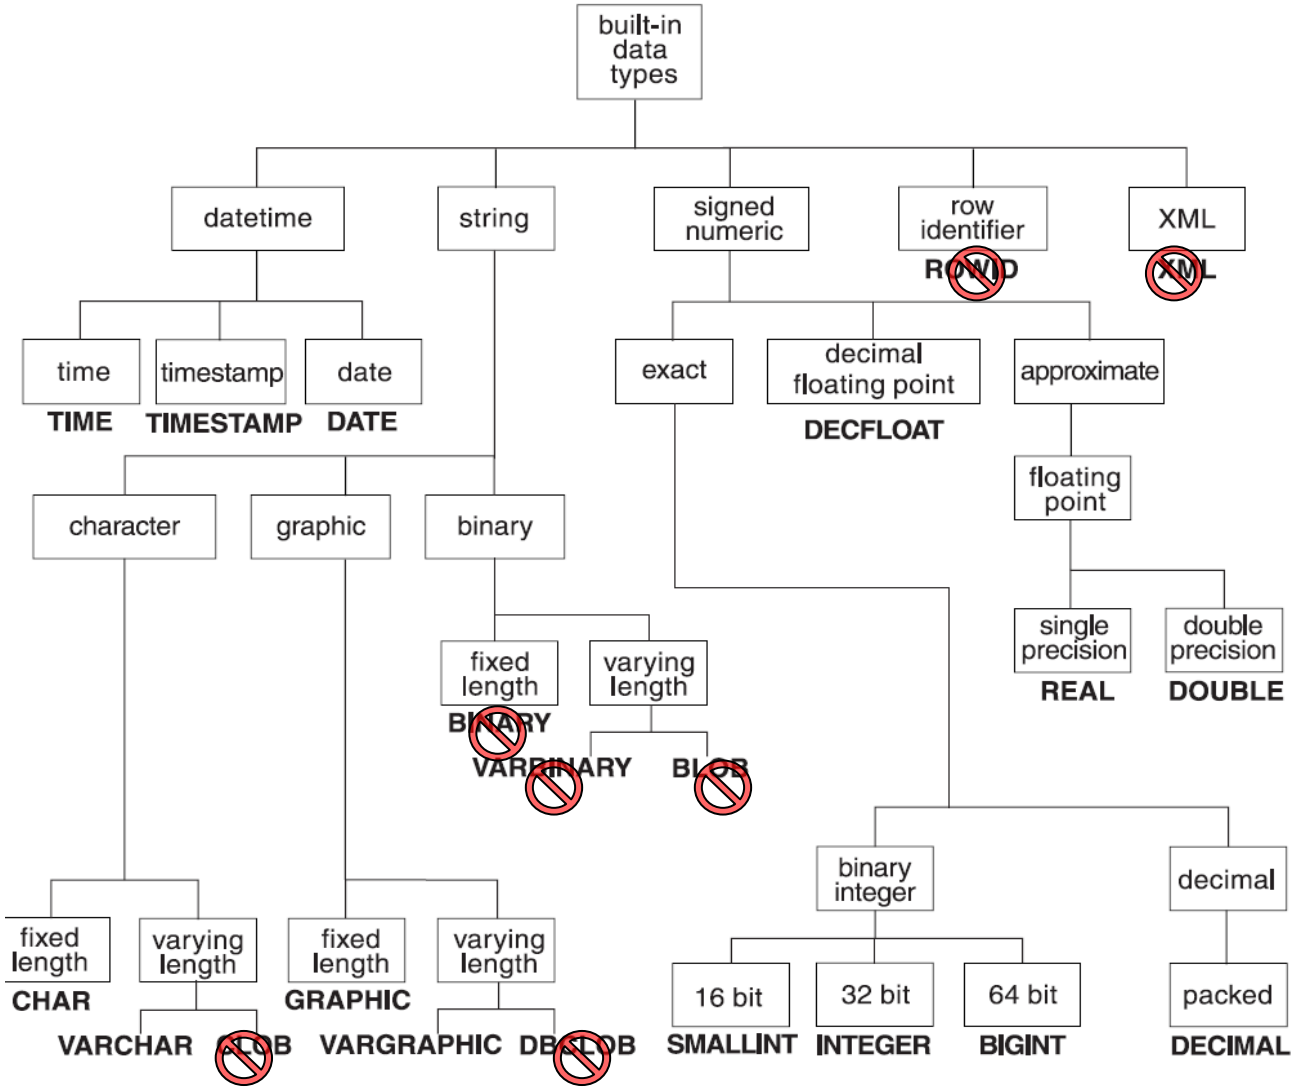

Not supported:

- Any kind of LOB
- ROWID
- XML
- Binary data

Prerequisits

- DB2 9 for z/OS – With maintenance
- z/OS 1.10 - With maintenance
- USS pipe support for TEMPLATE utility - PTF UK43948 for APAR PK70269
- XML Toolkit for z/OS V1.10.0, Program Number 5655-J51, FMID HXML190
- Data Studio V2.2 (IDE, no-charge download)
- IBM Smart Analytics Optimizer Studio
- System z BladeCenter Extension (zBX)
  - Z10 w/zBX and z196 w/zBX
- OSA Express3 10 GbE connection
- PC Windows XP or higher, 5 GB free disk, 2 GB RAM

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 Smart Analytics Optimizer 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

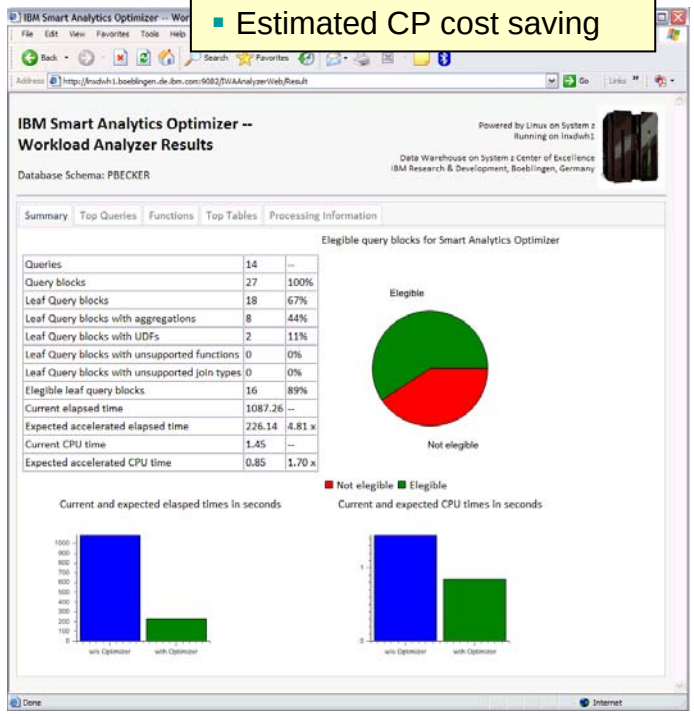

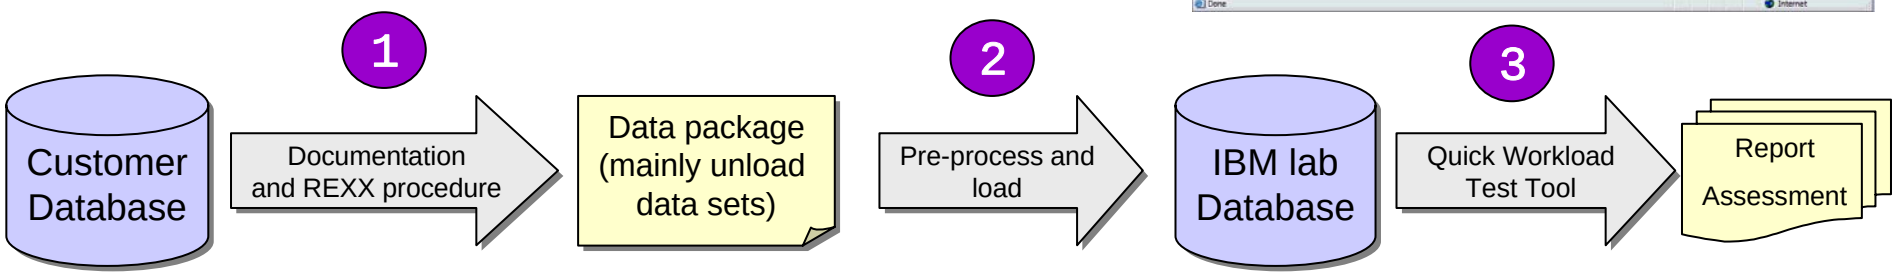

Activate Dynamic Statement Cache

- Dynamic Statement cache must be active
  - Install panel
    - CACHE DYNAMIC SQL field

  or

  - DSNZPARM
    - DSN6SPRM CACHEDYN

Start Trace for Cache (1 of 2)

- Non-Data-Sharing

```
-START TRACE(MON) CLASS(1) IFCID(316,317,318) DEST(SMF)
```

- Data-Sharing

```
-START TRACE(MON) CLASS(1) IFCID(316,317,318) DEST(SMF)
 SCOPE(GROUP)
```

- Verify

```
-DIS TRACE(*)
```

Start Trace for Cache (2 of 2)

- **IFCID 316**
  - First 60 bytes of SQL statement plus identifying information and statistics
    - http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.xe_db2.doc/ko2rrd20228.htm
- **IFCID 317**
  - Used in addition to IFCID 316 to obtain the full SQL statement text
    - http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.xe_db2.doc/ko2rrd20228.htm
- **IFCID 318**
  - Acts as a switch for IFCID 316 to collect all available information
    - Stop and Start for new interval

Create Tables Used by EXPLAIN

- Create tables used by EXPLAIN
  - Optimization Service Center (OSC)

    or
  - Optim Query Tuner (was Optimization Expert)

    or
  - DSNTIJOS job in SDSNSAMP

    or
  - A copy of DSNTIJOS is also included in the Assessment ZIP file

Get Dynamic Statement Cache

- Execute
  - **`EXPLAIN STMTCACHE ALL`**
  - Can use SPUFI, DSNTEP2, or anything that allows this SQL statement to be executed
- Ensure SQLID (SCHEMA) is set correctly for the EXPLAIN tables
- Verify/validate correct information was gathered

Run REXX Exec

- Provided REXX exec will EXPLAIN all SQL dumped from the statement cache using
  - owner.PLAN_TABLE
  - owner.DSN_STATEMNT_TABLE
  - owner.DSN_FUNCTION_TABLE
  - owner.DSN_PREDICAT_TABLE
  - owner.DSN_STRUCT_TABLE
  - owner.DSN_QUERY_TABLE
    - with CLOB column for statement text in XML format (V9)
  - owner.DSN_DETCOST_TABLE

Unload Everything and Send

- Using supplied job, unload EXPLAIN tables and selected DB2 Catalog tables
  - Read documentation carefully and pay attention to warnings
- Prepare unloads using ADRDSSU and TRSMAIN
- FTP to testcase at boulder

Clean Up

- Customer may want to cleanup all of the tables/data sets used during this process.
  - Sample cleanup jobs are provided

Customer Analysis – 1 of 2

IBM Smart Analytics Optimizer -- Workload Analysis Results

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

Query Summary

Queries	2729	
Queries potentially accelerated	2713	99%
Query Blocks	2733	100%
Leaf Query blocks	2729	100%
... with aggregations	13	0%
... with UDFs	0	0%
... with unsupported functions	1	0%
... with unsupported join types	15	1%
Eligible leaf query blocks	2713	99%
Current elapsed time	8110.00	--
Eligible leaf query blocks elapsed time	5796.66	71%
Current CPU time	1555.71	--

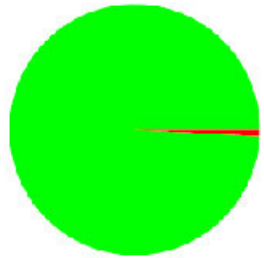

Query blocks with acceleration potential
Query blocks without acceleration potential

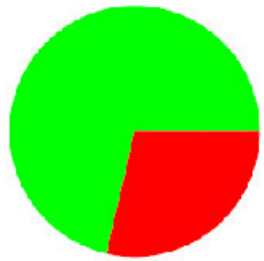

Elapsed time with acceleration potential
Elapsed time without acceleration potential

Customer Analysis – 2 of 2

IBM Smart Analytics Optimizer -- Workload Analysis Results

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

Query	Elapsed Time	Executed	Blocks	Eligible		Eligible Time	Table Sizes	Size Join Ind	Statement
149256	38.02	1	1		1	100.00%	118.50 MB		(Primauth: T. SELECT M D (T1.
138238	37.22	1	1		0	0.00%	12.93 GB		(Primauth: SELECT
87602	37.33	24	1		1	100.00%	267.25 MB		(Primauth SE XT' FROM ORDER BY
148204	36.77	1	1		1	100.00%	4.05 GB		(Primauth SELECT (T R_MO_NBR' FOR FETCH ONLY
146105	33.67	2	1		0	0.00%	12.93 GB		(Primauth: SELECT R C ORDER BY T1. FOR FETCH ONLY
140675	33.32	2	1		1	100.00%	60.00 MB		(Primauth SE BY Only
149330	32.46	1	1		1	100.00%	677.11 MB		(Primauth: SELECT F T (ORDER BY T1.
139154	29.40	1	1		1	100.00%	7.62 GB		(Primauth: SELECT D

IBM Smart Analytics Optimizer - Summary

Summary of Value Proposition

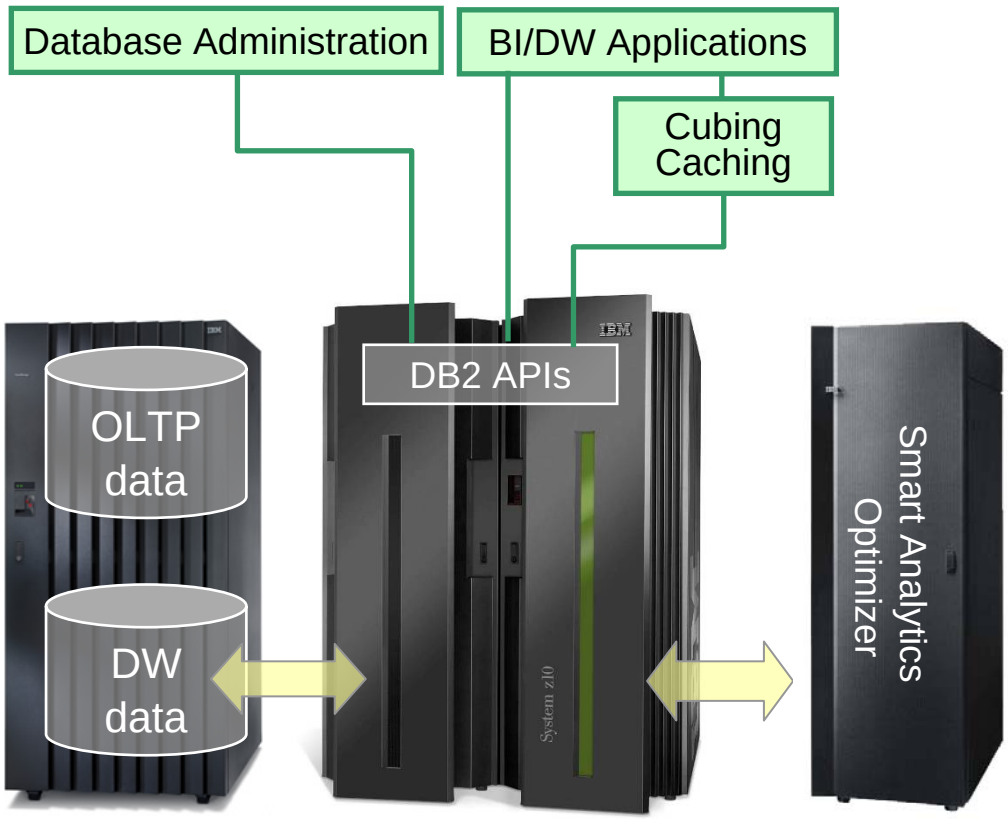

- Seamless integration of new computing paradigms into proven technology
  - Massive multi-core and vector processing
  - In-memory database
  - No changes to the applications; applications continue to attach to DB2
    - When applicable query needs execution, DB2 transparently exploits Smart Analytics Optimizer
  - Preserving traditional System z and DB2 Quality of Service (QoS), full fencing and protection of DB2 against possible Smart Analytics Optimizer failures
  - Hands free operation
- Augmenting System z value proposition as the overall Enterprise Data Hub
  - System z hybrid topology enables additional transparency and management integration
- Order of magnitude performance improvement
  - Linear scaling with the number of CPUs
- Reducing need for tedious tuning of DB2 (MQTs, aggregates, indexes, etc.)
- Providing building block for Dynamic Data Warehousing and Operational BI

- Using IBM System z as the Foundation for Your Information Management Architecture, REDP-4606
  - Focuses on Smart Analytics Optimizer
    - http://www.redbooks.ibm.com/redpieces/abstracts/redp4606.html
- IBM zEnterprise BladeCenter Extension Model 001, REDP-4668
    - http://www.redbooks.ibm.com/redpieces/abstracts/redp4668.html
- 50 TB Data Warehouse Benchmark on IBM System z, SG24-7674
    - http://www.redbooks.ibm.com/abstracts/sg247674.html
- Housing Transactional & Data Warehouse Workloads on System z,
  - SG24-7726-00
    - http://www.redbooks.ibm.com/redpieces/abstracts/sg247726.html
- Enterprise Data Warehousing with DB2 9 for z/OS, SG24-7637
    - http://www.redbooks.ibm.com/abstracts/sg247637.html
- Workload Management for DB2 Data Warehouse, REDP-3927
    - http:// www.redbooks.ibm.com/abstracts/redp3927.html

धन्यवाद

Hindi

多謝

Traditional Chinese

ขอบคุณ

Thai

Спасибо

Russian

Gracias

Spanish

Thank You

English

شكرا

Arabic

Merci

French

Obrigado

Brazilian Portuguese

Bedankt

Nederlands

多谢

Simplified Chinese

Danke

German

நன்றி

Tamil

ありがとうございました

Japanese

감사합니다