


# Smart Analytics Optimizer for DB2 for z/OS: How Does it Work?

Willie Favero  
IBM Data Warehouse on System z Swat Team (DB2 SME)

Thursday, March 3, 2011: 4:30 PM-5:30 PM  
Session # 8786



**SHARE**

Technology • Connections • Results

# Notices

This information was developed for products and services offered in the U.S.A.

Note to U.S. Government Users Restricted Rights — Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.


**SHARE**  
in Anaheim  
2011



# Trademarks

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

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



# Topics

- The IBM Smart Analytics Optimizer

- The basics
- Beta customer experience
- Under the hood
- Setting it up


- Is it right for me?





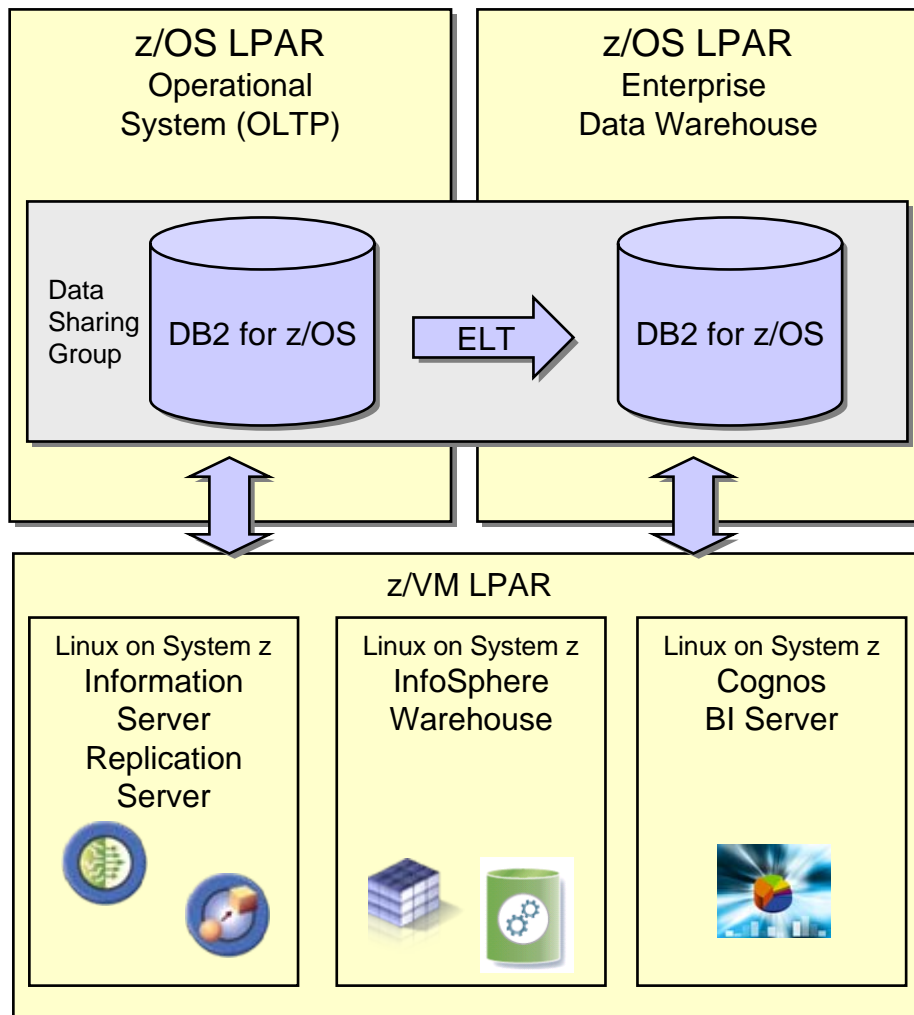
# Business Challenges and Technology Trends

- 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 x86 blades<sup>1</sup> for tens of thousands of AIX® and Linux applications
- High performance optimizers and appliances to accelerate time to insight and reduce cost



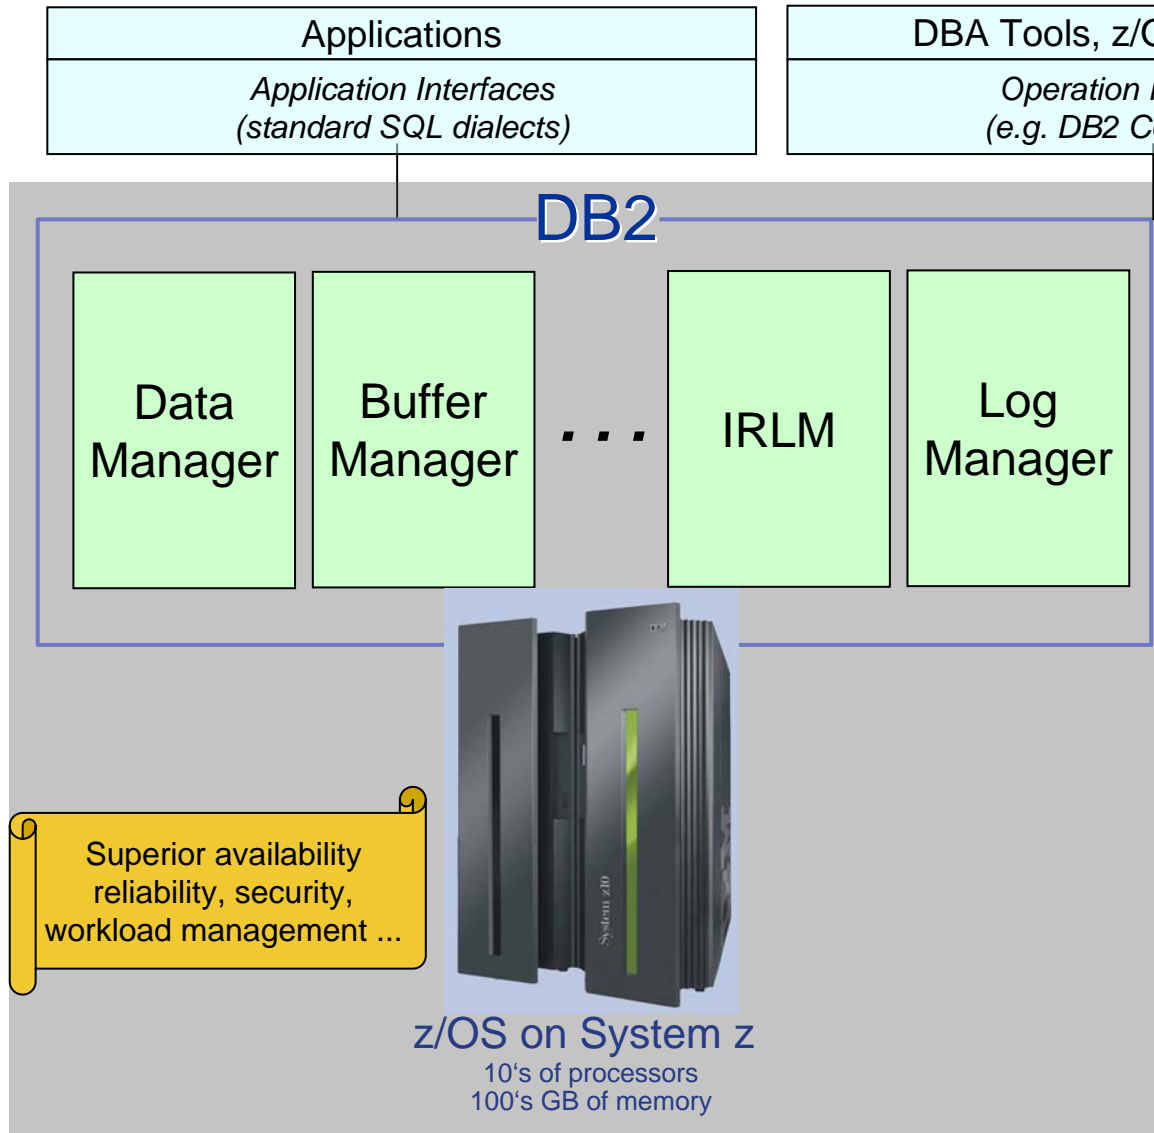
<sup>1</sup> All statements regarding IBM future direction and intent are subject to change or withdrawal without notice and represents goals and objectives only.

▪ Dedicated high performance private network




# IBM Smart Analytics Optimizer – a Virtual DB2 Component

**SHARE**  
Technology • Connections • Results

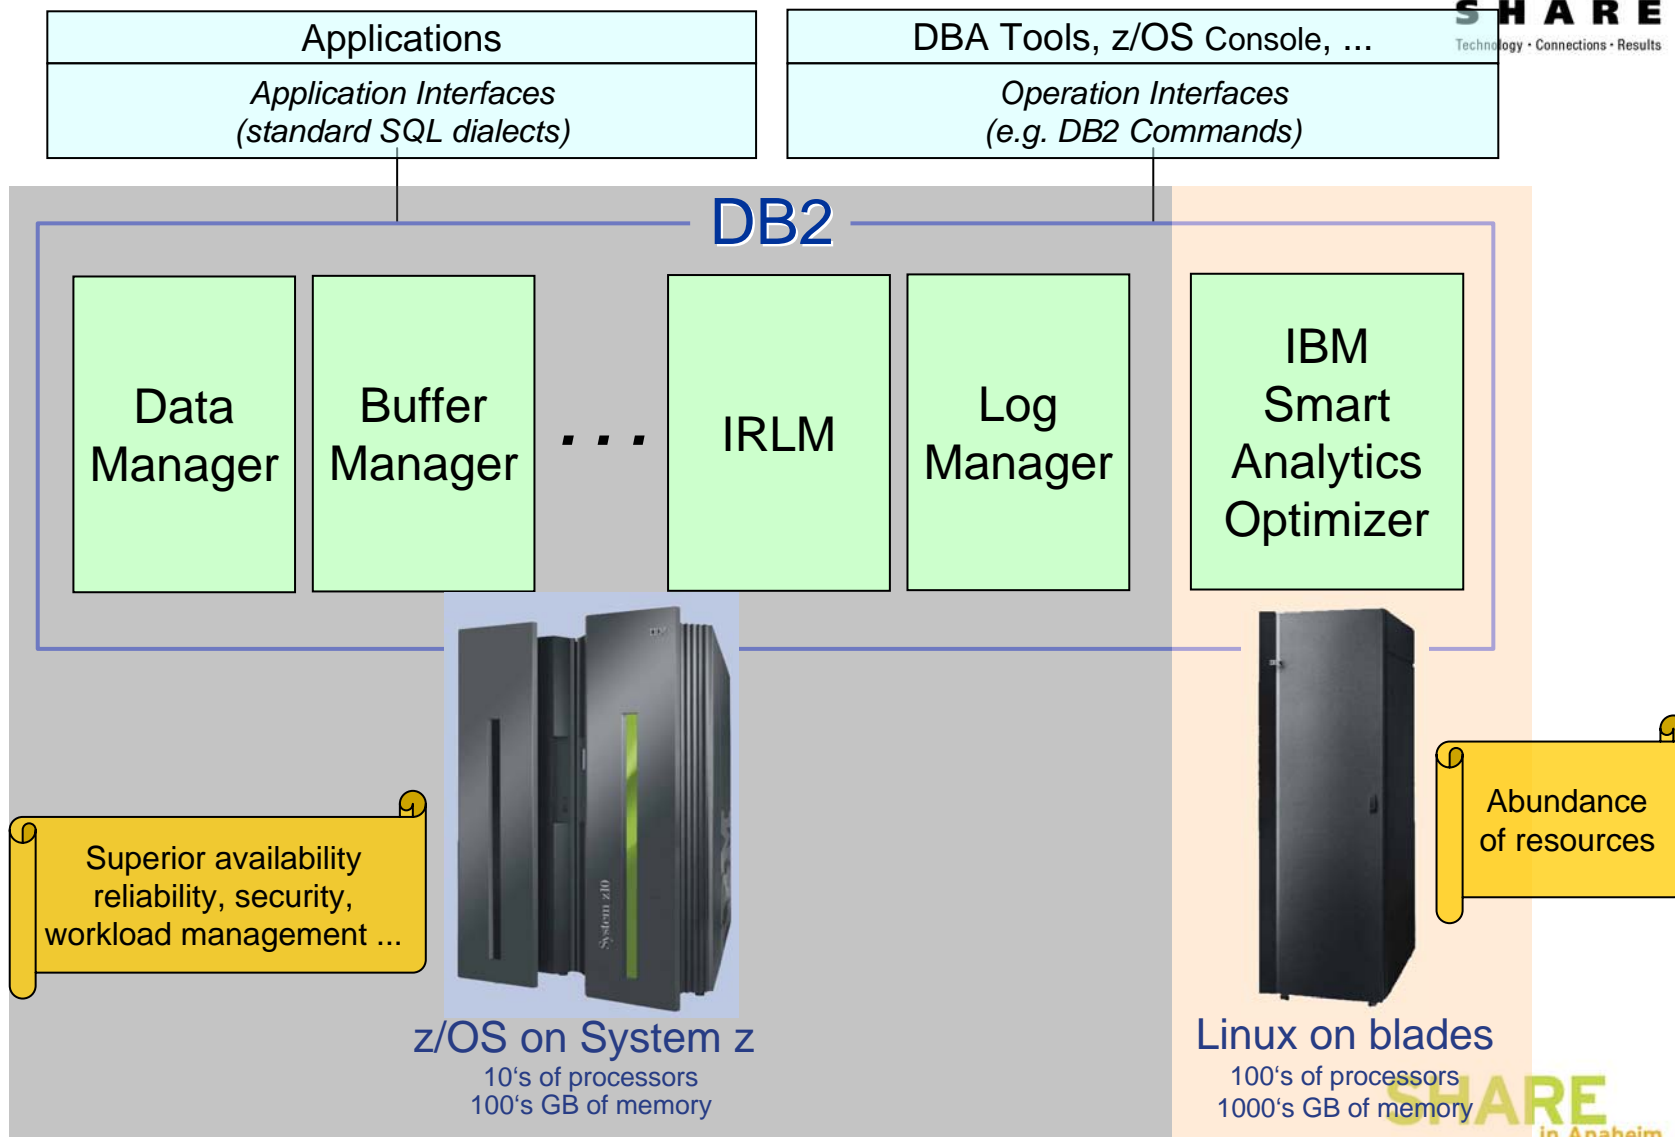






**SHARE**  
Technology · Connections · Results


# IBM Smart Analytics Optimizer – a Virtual DB2 Component





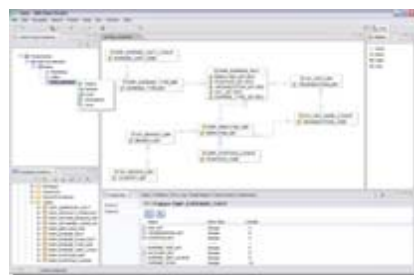
# Characteristics of the IBM Smart Analytics Optimizer

- 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 accelerated queries
  - Reduced DBA effort for tuning accelerated queries
- Hybrid technology
  - Enabling Dynamic DW and Operational BI
  - Preserving traditional DB2 quality of service
  - Having transactional and analytical workload being handled by DB2



# 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.

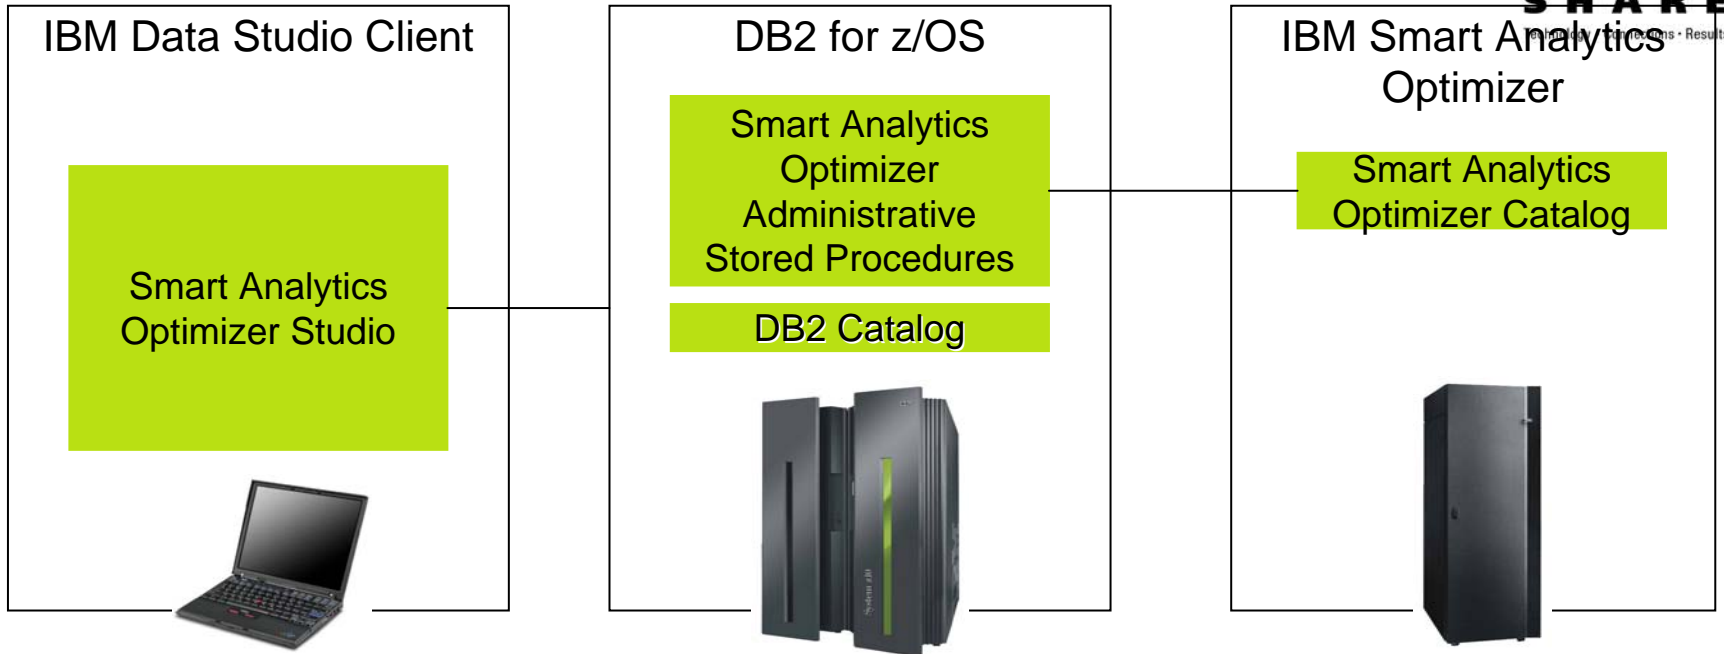


Define




Transform

# 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

SHARE

The screenshot shows the Eclipse IDE interface for a Data Mart project. The main window displays a diagram of the TPC-H Mart structure. The tables and their relationships are as follows:

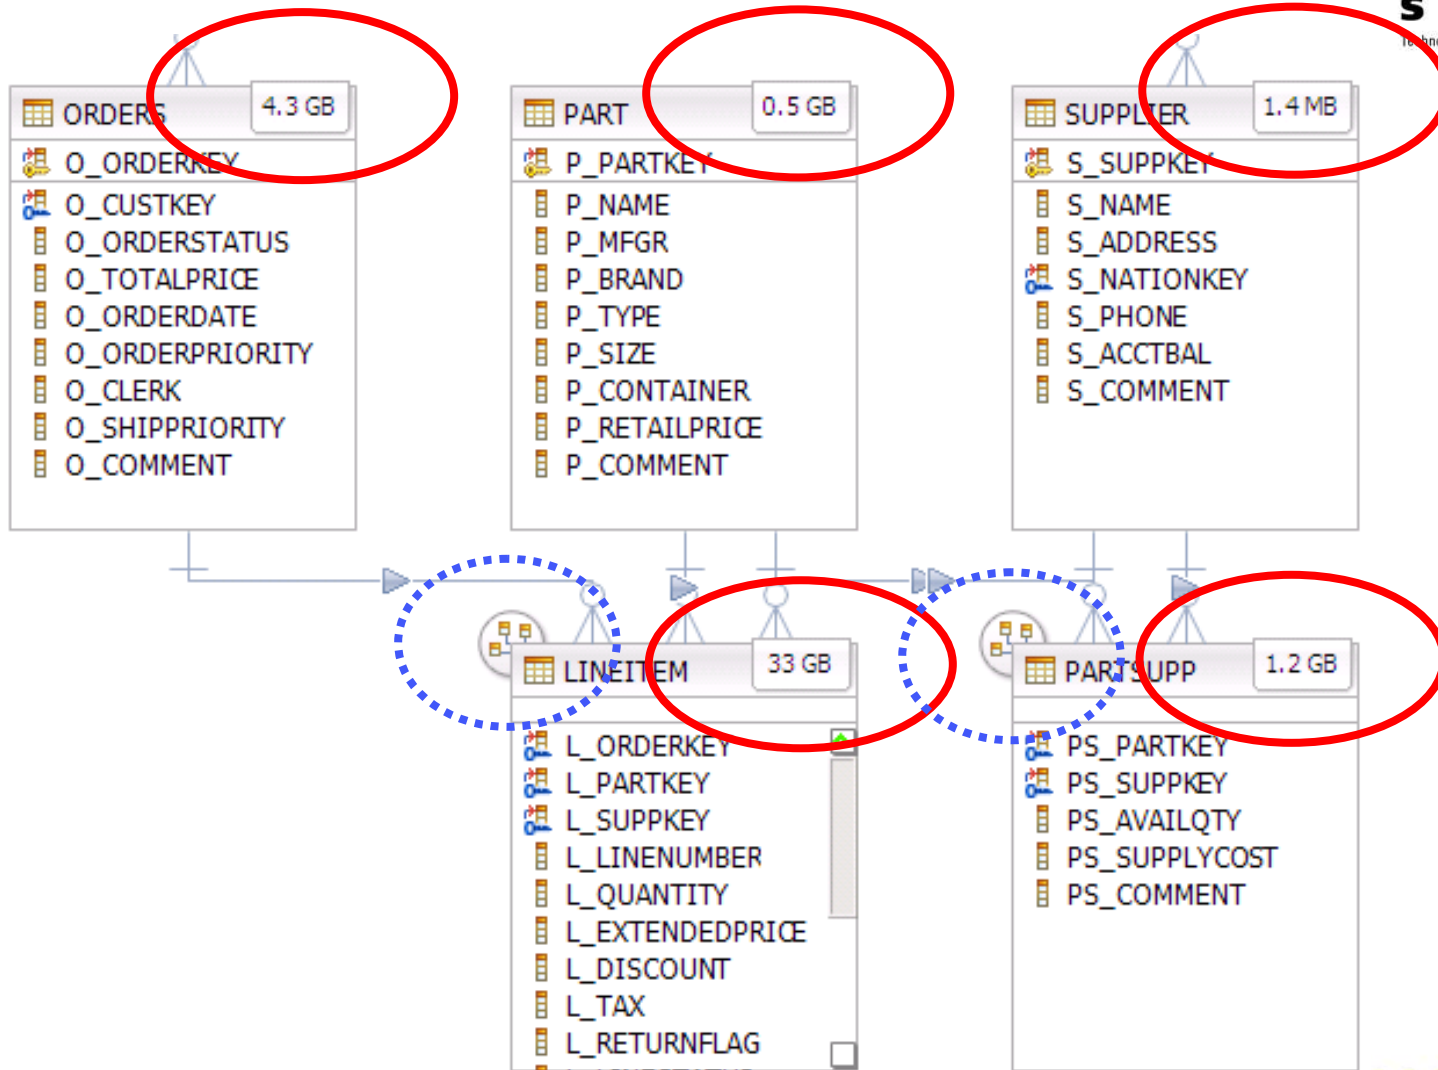
- ORDERS** (4.3 GB) is connected to **PART** (0.5 GB) and **SUPPLIER** (1.4 MB).
- PART** is connected to **LINEITEM** (33 GB).
- SUPPLIER** is connected to **PARTSUPP** (1.2 GB).
- LINEITEM** is connected to **PARTSUPP**.


The Properties window at the bottom shows the following details for the Mart:

Mart TPC-H Mart	
<b>General</b>	Deploy... Validate
<b>Validation</b>	Name: TPC-H Mart
<b>Space Required</b>	Estimated Size: 39 GB

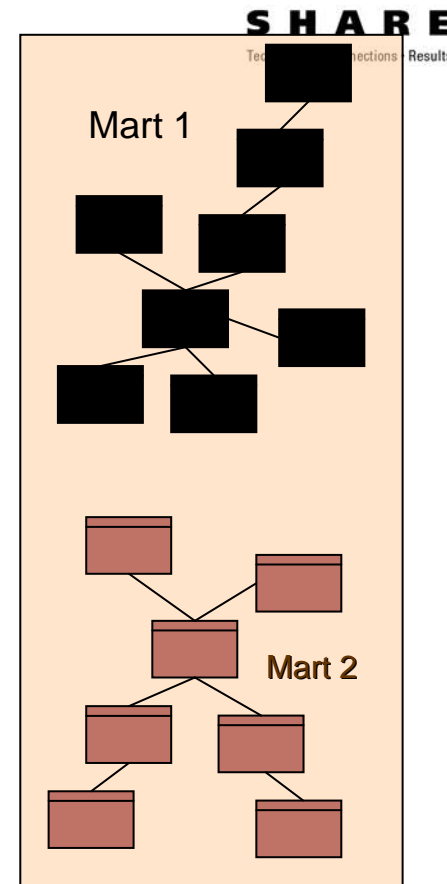

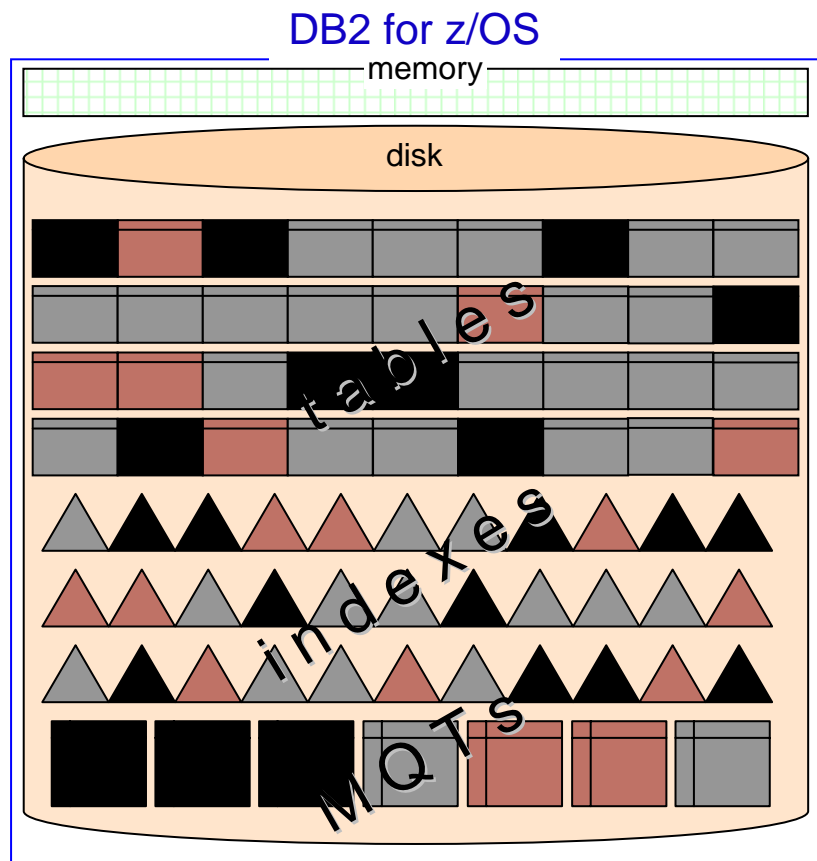


# Zoom In: Size Estimates and Fact table property






# Data View: Marts – Redundant Sets of Memory Resident Tables



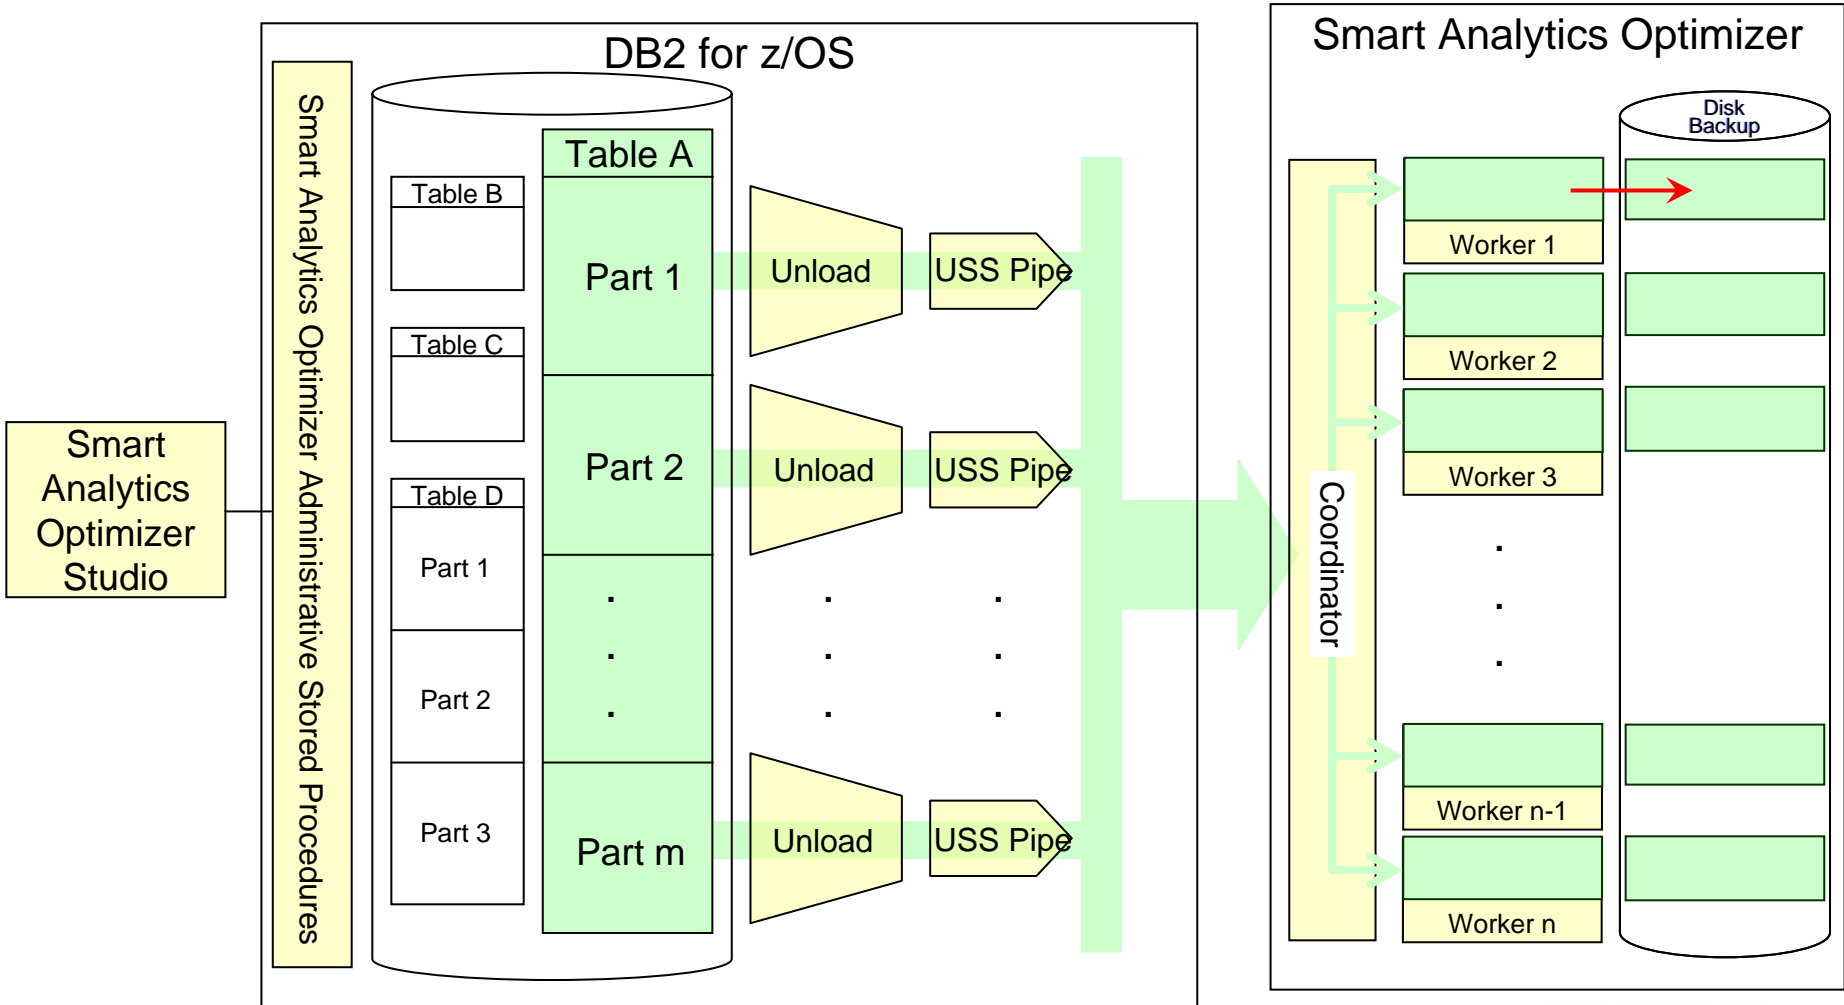
- DB2 continues to *own* and manage all data
- Access performance is influenced by traditional tuning mechanisms such as indexing, MQTs, aggregates, ...
- Typical usage: Enterprise Data Warehouse, large Data Marts

- Smart Analytics Optimizer contains fully memory resident, compressed *copies* of performance critical tables grouped into logically connected *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

# Smart Analytics Optimizer Mart Load



**SHARE**  
Technology · Connections · Results



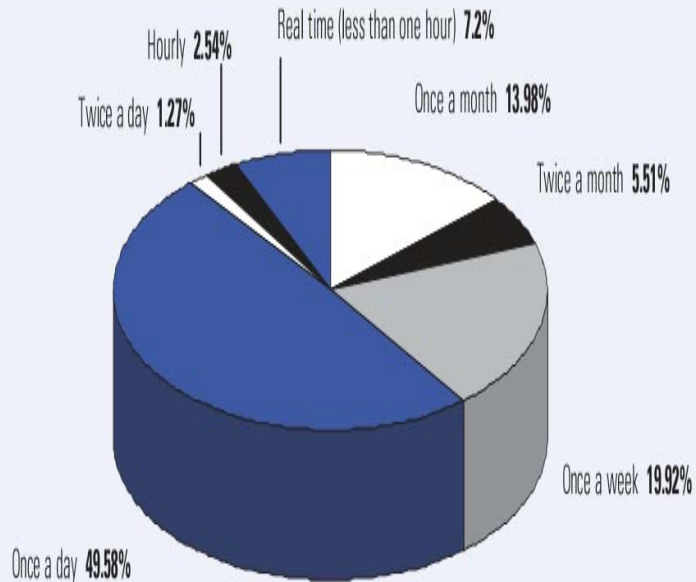
**SHARE**  
in Anaheim  
2011



# Smart Analytics Optimizer Mart Update


## 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

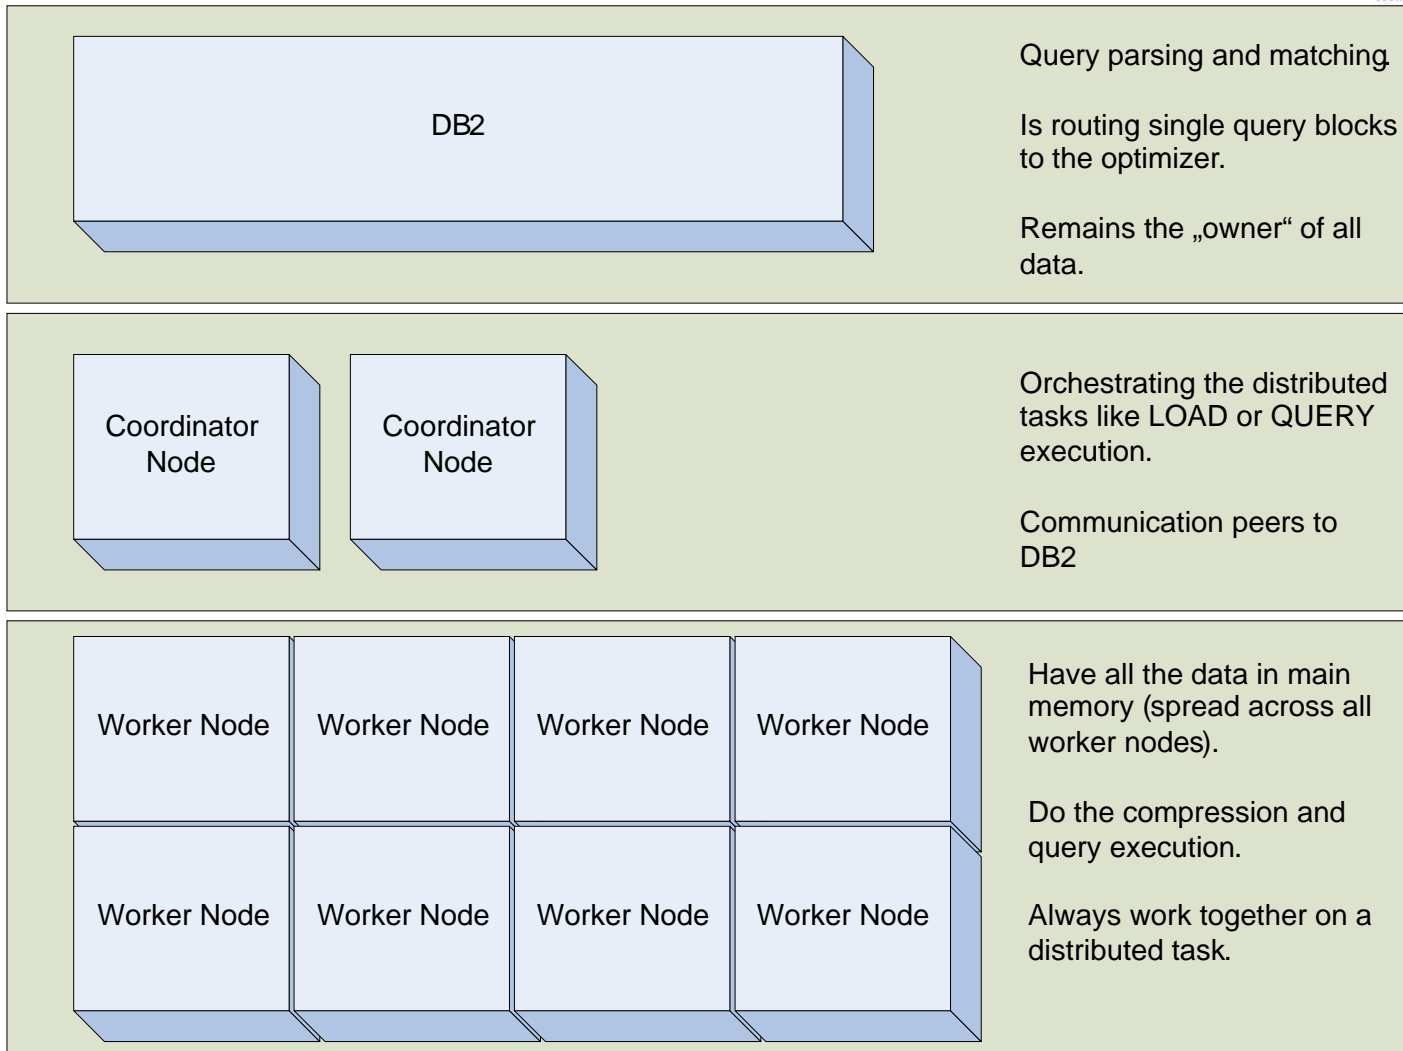
- Typical DW update operations:
  - LOAD RESUME and REPLACE
  - ADD and ROTATE PARTITION
  - SQL INSERT, UPDATE, DELETE
  - Delete complete partition or table
  - TRUNCATE TABLE
  
- Smart Analytics Optimizer will over time phase-in support for all the typical operations in this order
  1. Full table 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




*Pause for*  
*Questions*



# IBM Smart Analytics Optimizer Node Overview

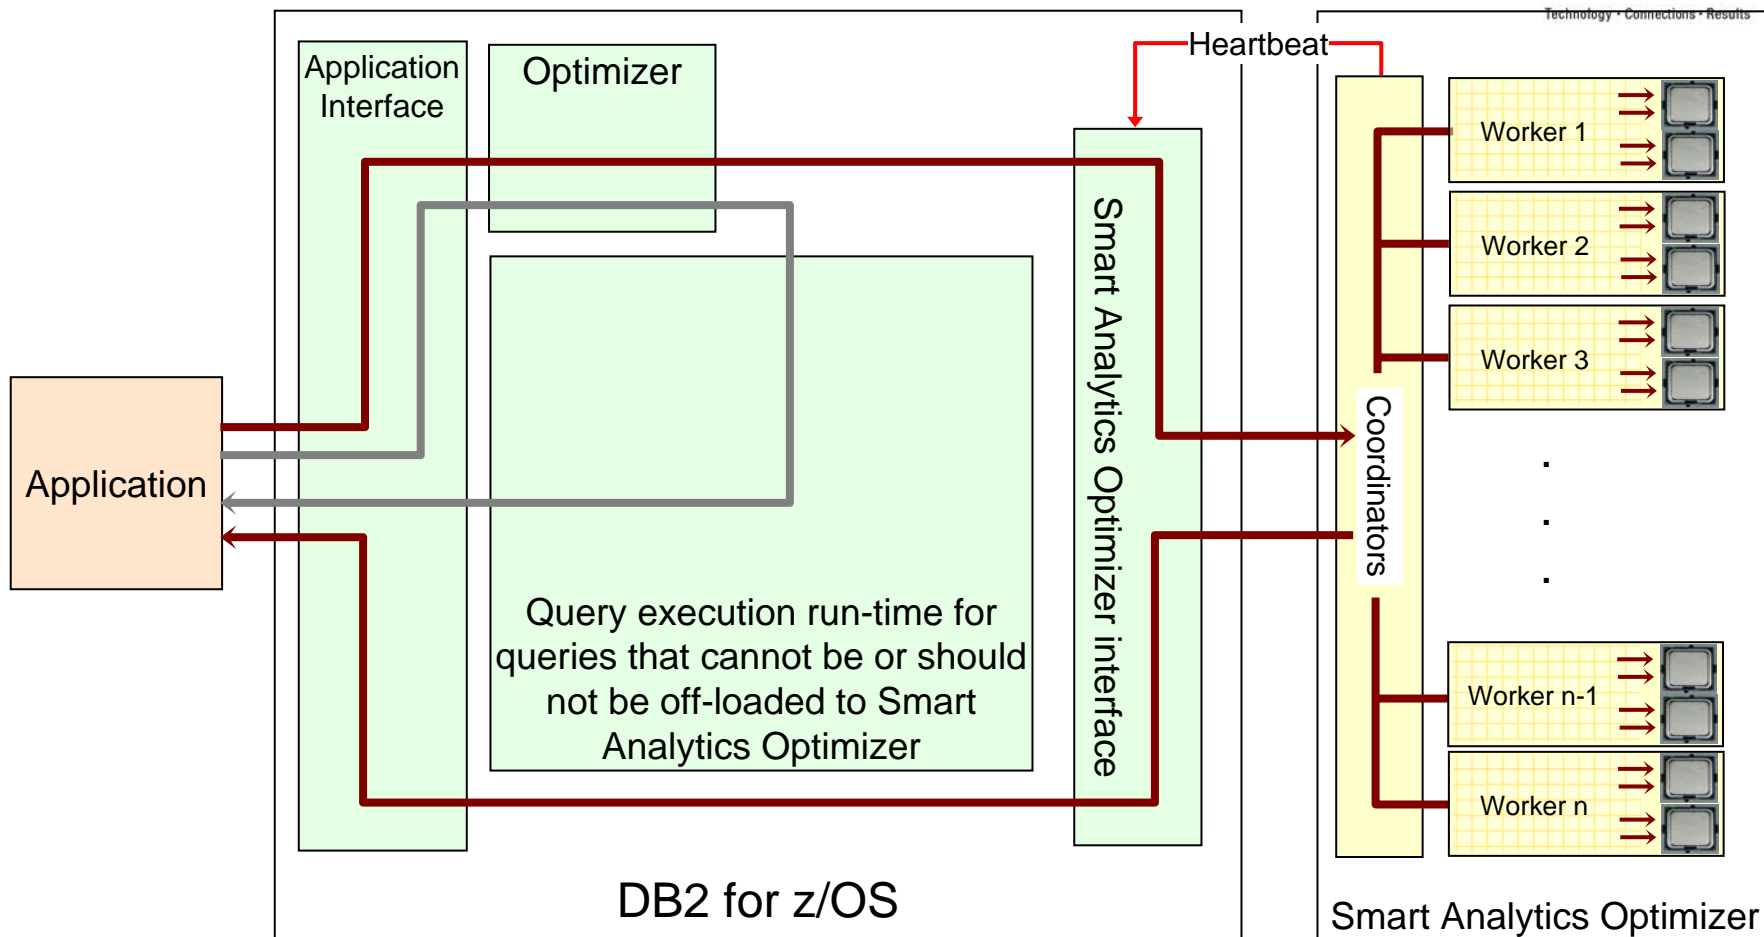




**SHARE**


Technology • Connections • Results

# Query Execution Process Flow



- Heartbeat (Smart Analytics Optimizer availability and performance indicators)
- Queries executed **without** Smart Analytics Optimizer
- Queries executed **with** Smart Analytics Optimizer





**SHARE**  
Technology · Connections · Results


# 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

<Table> DSN\_QUERYINFO\_TABLE

General				
Columns	Name	Primary Key	Data Type	Length
Relationships	QUERYNO	<input type="checkbox"/>	INTEGER	
Documentation	QBLOCKNO	<input type="checkbox"/>	SMALLINT	
	QINAME1	<input type="checkbox"/>	VARCHAR	128
	QINAME2	<input type="checkbox"/>	VARCHAR	128
	APPLNAME	<input type="checkbox"/>	VARCHAR	24
	PROGNAME	<input type="checkbox"/>	VARCHAR	128
	VERSION	<input type="checkbox"/>	VARCHAR	122
	COLLID	<input type="checkbox"/>	VARCHAR	128
	GROUP_MEMBER	<input type="checkbox"/>	VARCHAR	24
	SECTNOI	<input type="checkbox"/>	INTEGER	
	SEQNO	<input type="checkbox"/>	INTEGER	
	EXPLAIN_TIME	<input type="checkbox"/>	TIMESTAMP	
	TYPE	<input type="checkbox"/>	CHAR	8
	REASON_CODE	<input type="checkbox"/>	SMALLINT	
	QI_DATA	<input type="checkbox"/>	CLOB	2097152
	SERVICE_INFO	<input type="checkbox"/>	BLOB	2097152
	QB_INFO_ROWID	<input type="checkbox"/>	ROWID	

**SHARE**  
in Anaheim  
2011



# Additional DB2 Support - Commands


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

```
DSNX810I -D2PAA DSNX8CMD STOP ACCEL FOLLOWS -  
DSNX860I -D2PAA DSNX8STO STOP ACCELERATOR SUCCESSFUL FOR IWA01  
DSNX9022I -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  
DSNX9022I -D2PAA DSNX8CMD '-START ACCEL' NORMAL COMPLETION
```


# Display Accelerator Command



```

DSNX810I  -DB91 DSNX8CMD DISPLAY ACCEL FOLLOWS -
DSNX830I  -DB91 DSNX8CDA
ACCELERATOR          MEMB  STATUS  REQUESTS  QUED  MAXQ  FAIL
-----
BACGEN              DB91  STARTED          0     0     0     0
PORT=0 IPNAME=VIRTUAL
DEMOACC1           DB91  STARTED        365     0     1    28
PORT=1400 IPNAME=897A45D3
IPADDR=192.168.139.100 HEALTHY
DETAIL STATISTICS
  LEVEL   = AQT01010
  STATUS  = FULLY OPERATIONAL
  AVERAGE QUEUE WAIT          =          0
  MAXIMUM QUEUE WAIT          =          0
  PROCESSING CAPACITY         = 1314278
  TOTAL NUMBER OF PROCESSORS  =         224
  AVERAGE CPU UTILIZATION ON COORDINATOR NODES =         .35%
  NUMBER OF ACTIVE COORDINATOR NODES          =          3
  PHYSICAL MEMORY AVAILABLE ON COORDINATOR NODES = 47622 MB
  PHYSICAL MEMORY AVERAGE USAGE ON COORDINATOR NODES = 653 MB
  AVERAGE CPU UTILIZATION ON WORKER NODES     =         .24%
  NUMBER OF ACTIVE WORKER NODES                =         11
  PHYSICAL MEMORY AVAILABLE ON WORKER NODES    = 40395 MB
  PHYSICAL MEMORY AVERAGE USAGE ON WORKER NODES = 7879 MB
  SHARED MEMORY DATA AVAILABLE ON WORKER NODES = 25423 MB
  SHARED MEMORY DATA AVERAGE USAGE ON WORKER NODES = 6728 MB
  MAXIMUM SHARED MEMORY DATA IN USE ON WORKER NODES = 6741 MB
DISPLAY ACCEL REPORT COMPLETE
DSNX9022I  -DB91 DSNX8CMD '-DISPLAY ACCEL' NORMAL COMPLETION

```



# Display Thread Command

```

DSNV401I  -DB91 DISPLAY THREAD REPORT FOLLOWS -
DSNV402I  -DB91 ACTIVE THREADS -
NAME      ST A   REQ ID              AUTHID   PLAN      ASID  TOKEN
BATCH     AC *     7 PBECKERD          PBECKER  DSNTIB91 0035   737
V666 ACC=DEMOACC1,ADDR=192.168.139.100..1400:9991
V441-ACCOUNTING=56000999:00000000:47630000
V436-PGM=*.DSNTIAUL, SEC=2, STMT=3644
TSO       T     307 FNEUMAN          FNEUMAN  ADB       0080   417
V441-ACCOUNTING=3300
TSO       T     7 PBECKER          PBECKER  ADB       0070   751
V441-ACCOUNTING=3300
TSO       T *    3 PBECKER          PBECKER  ADB       0070   756
V441-ACCOUNTING=3300
DISPLAY ACTIVE REPORT COMPLETE
DSN9022I  -DB91 DSNVDT '-DIS THREAD' NORMAL COMPLETION
***
  
```





# Preconfigured Solution Sizes for the IBM Smart Analytics Optimizer




Blades	7	14	28	42	56
<b>Handles up to xx of Raw DB2 Table Data</b>	0.5TB	1TB	2TB	3TB	4TB
Enterprise Rack	1	1	1	2	2
BladeCenter H Chassis	1	1	2	3	4
DS5020 with # of SATA disks (1 TB)	16	16	16	32	32
Storage Expansion Units	0	0	0	1	1

- Includes redundant hardware infrastructure for availability including PDUs, TOR 10GbE switches for data, TOR 1GbE switches for management, FC switches, etc...

- **For zEnterprise 196, order Machine Type 2458 Model 002 with feature #0610**
- **DS5020 disk needs to be ordered separately from zBX, and housed in its own rack.**





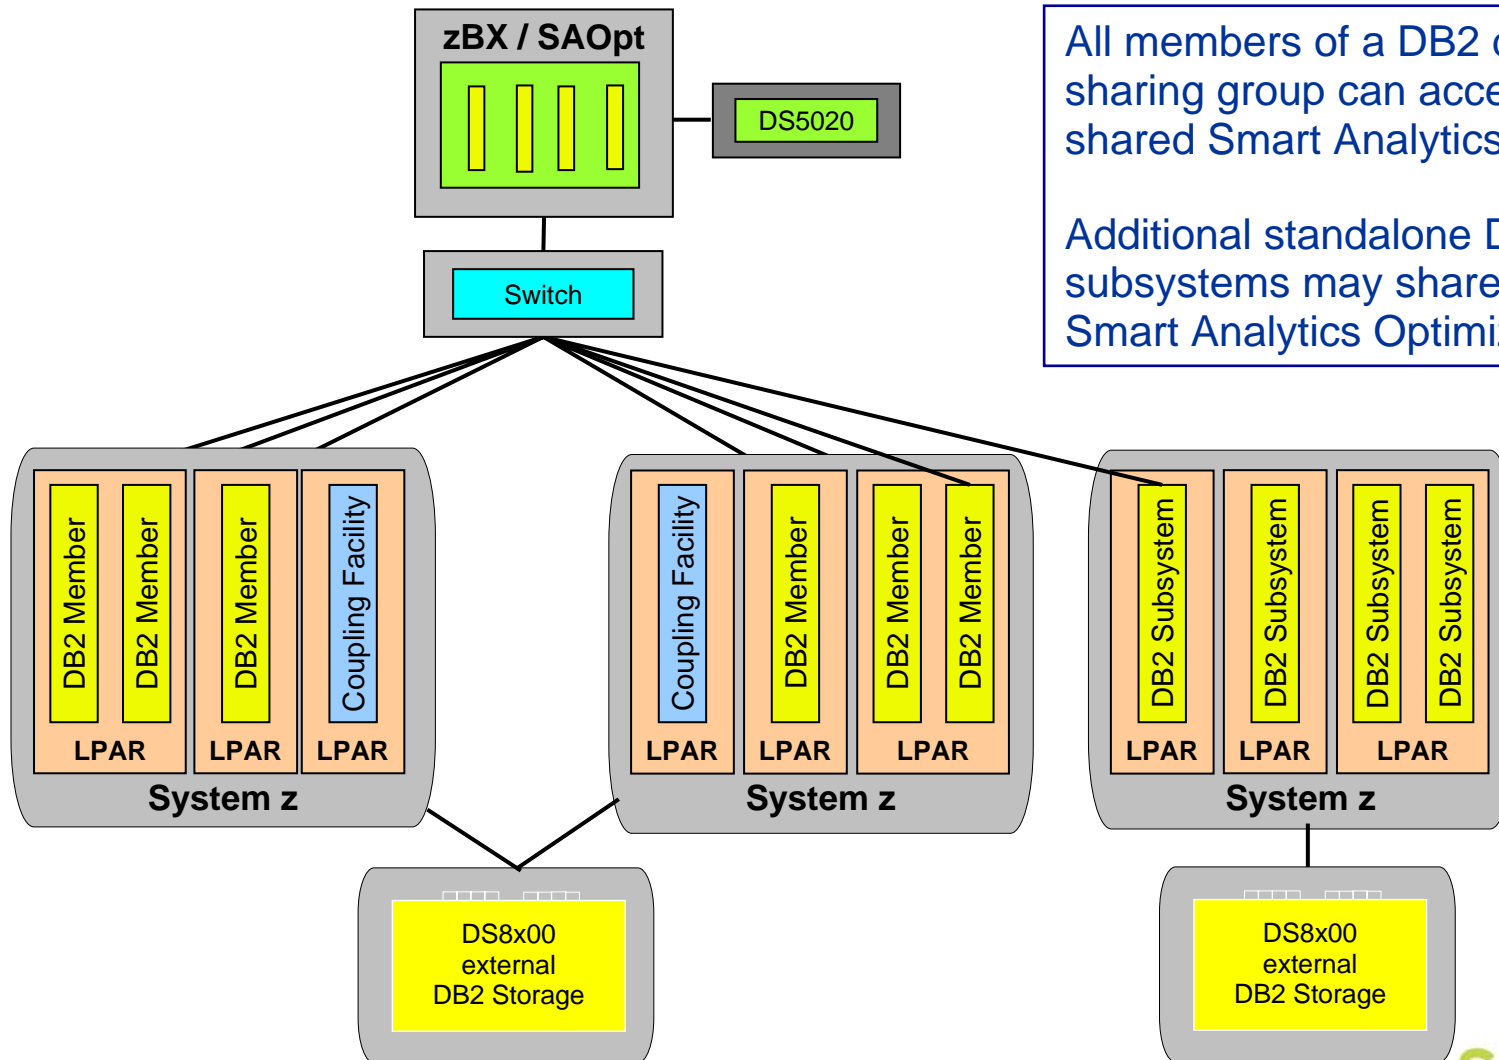
# zBX with Smart Analytics Optimizer Sysplex configuration


**SHARE**

Technology • Connections • Results

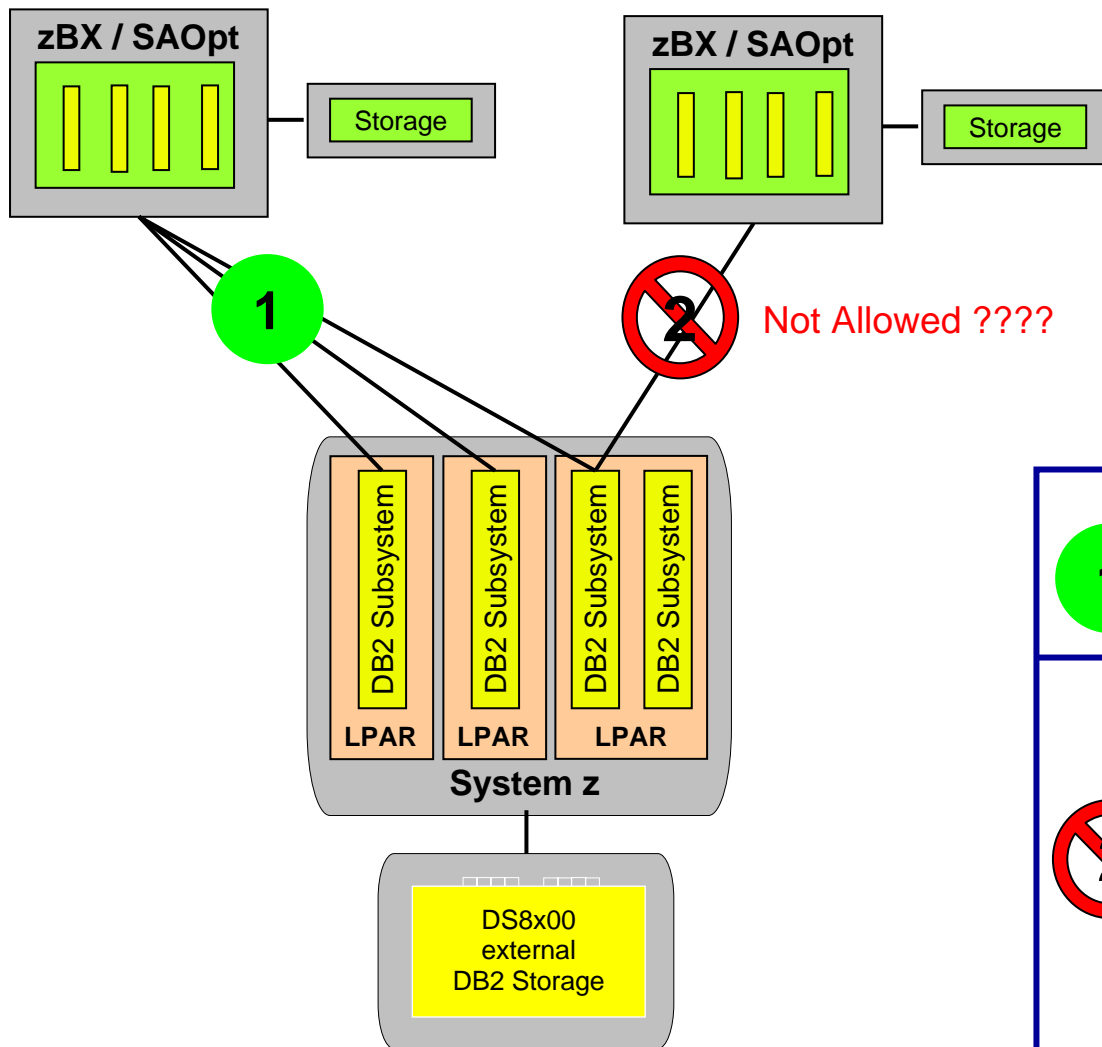
All members of a DB2 data sharing group can access a shared Smart Analytics Optimizer.

Additional standalone DB2 subsystems may share the same Smart Analytics Optimizer.






# Multiple Smart Analytics Optimizer Configuration




	OK
	One DB2 subsystem must not access more than one Smart Analytics Optimizer / zBX. Only one Smart Analytics Optimizer / zBX may be connected to one CEC.



# Testing Results – Native on DB2

- 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 <b>without</b> Smart Analytics Optimizer			
	Total Elapsed	CP	zIIP	Total CPU Time
<b>Q1</b>	<b>0:02:43</b>	0:03:52	0:02:39	<b>0:06:31</b>
<b>Q2</b>	<b>0:38:31</b>	<b>0:11:52</b>	<b>0:36:10</b>	<b>0:48:02</b>
<b>Q3</b>	<b>0:00:25</b>	0:00:04	0:00:15	<b>0:00:19</b>
<b>Q4</b>	<b>0:26:33</b>	<b>0:13:43</b>	<b>0:20:50</b>	<b>0:34:33</b>
<b>Q5</b>	<b>0:00:35</b>	0:00:09	0:00:29	<b>0:00:38</b>
<b>Q6</b>	<b>1:30:35</b>	<b>5:53:30</b>	<b>1:29:56</b>	<b>7:23:26</b>
<b>Q7</b>	<b>0:00:02</b>	0:00:02	0:00:00	<b>0:00:02</b>




# Testing Results – Smart Analytics Optimizer

Performance Improvement after Adding IBM Smart Analytics Optimizer

**SHARE**  
Technology • Connections • Results

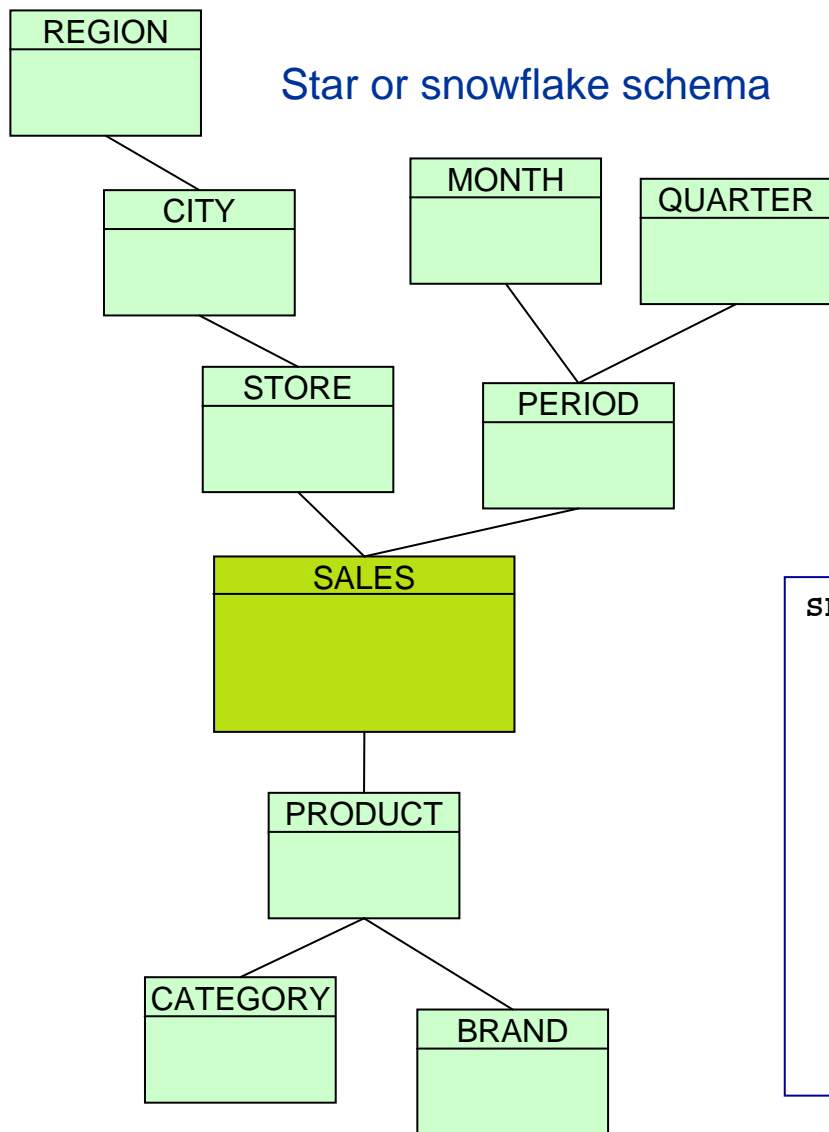
Query	Query Elapsed Time		
	DB2 only	DB2 with Smart Analytics Optimizer	Speed-up
Q1	0:02:43.0	0:00:03.4	48
Q2	0:38:31.0	0:00:04.5	511
Q3	0:00:25.0	0:00:02.2	12
Q4	0:26:33.0	0:00:07.8	206
Q5	0:00:35.0	0:00:08.3	4
Q6	1:30:35.0	0:00:03.8	1424
Q7	0:00:02.0	0:00:02.0	1
Total	2:39:24.0	0:00:32.0	298

in Anaheim  
2011



**SHARE**  
Technology · Connections · Results

# What Is Smart Analytics Optimizer 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

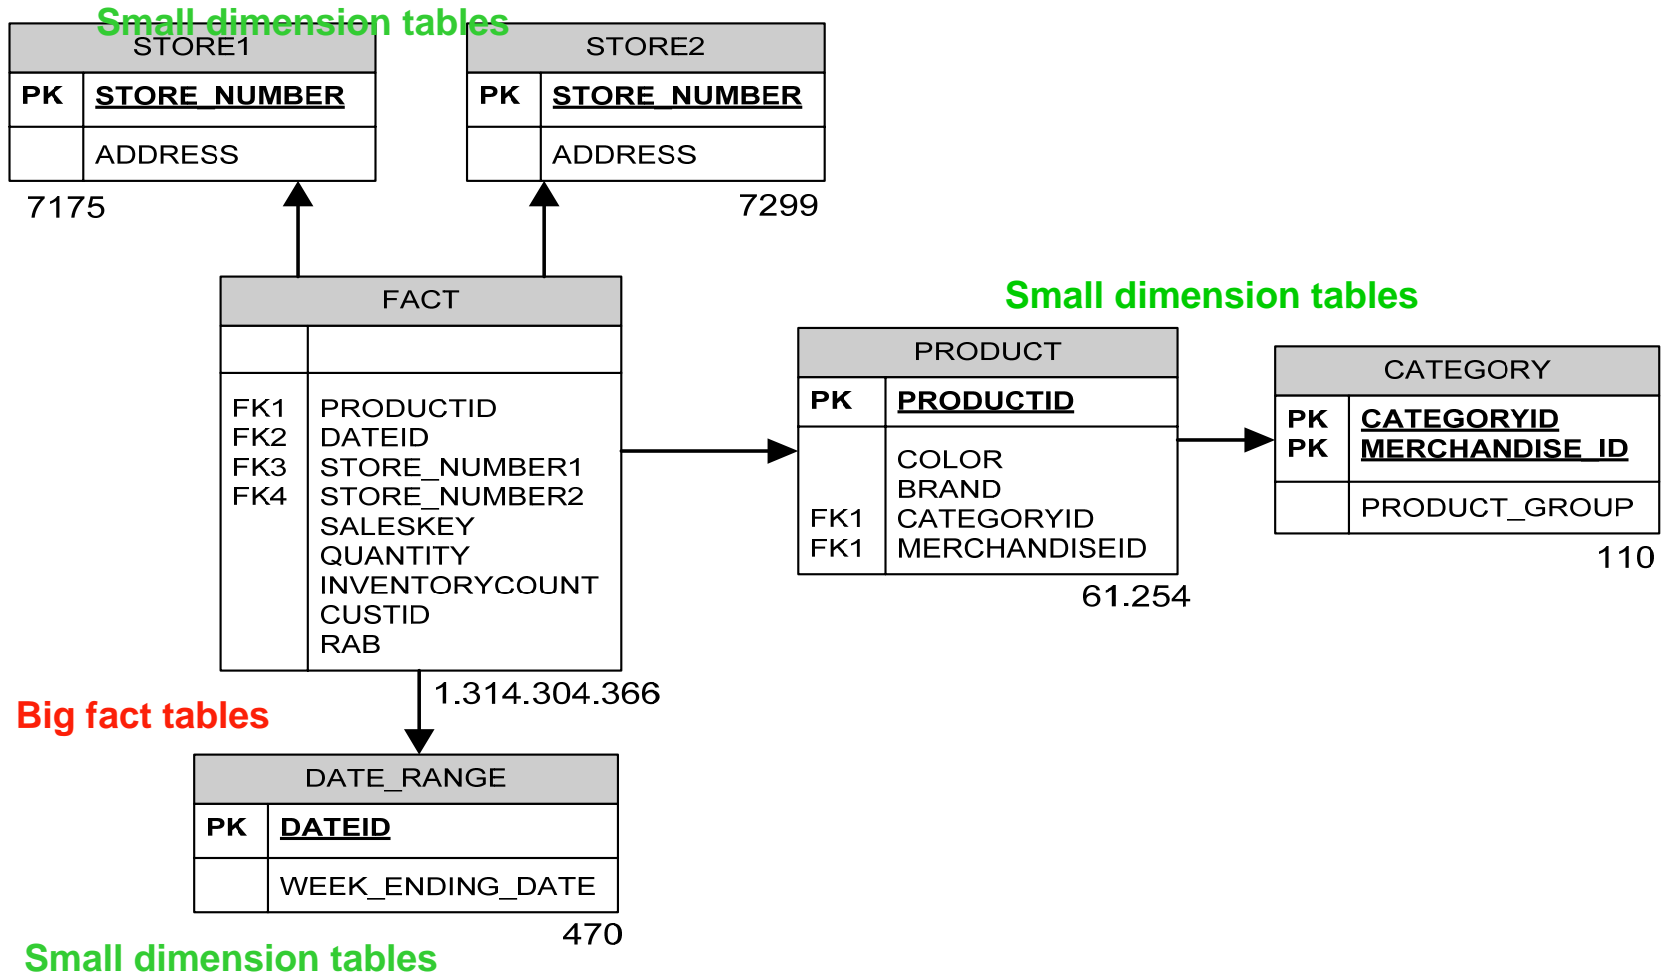
```

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
      AND P.TYPE = 'SOFTWARE'
      AND R.GEO = 'SOUTH'
GROUP BY PRODUCT_DEPARTMENT, REGION
    
```

**SHARE**  
in Anaheim  
2011



# 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 )



# Potentially multiple query blocks

## Derived table (nested table expression)

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


The intermediate results (data) may **ONLY** flow from the accelerator to DB2 and not the other way

## 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






**SHARE**  
Technology • Connections • Results

# First Release Restrictions

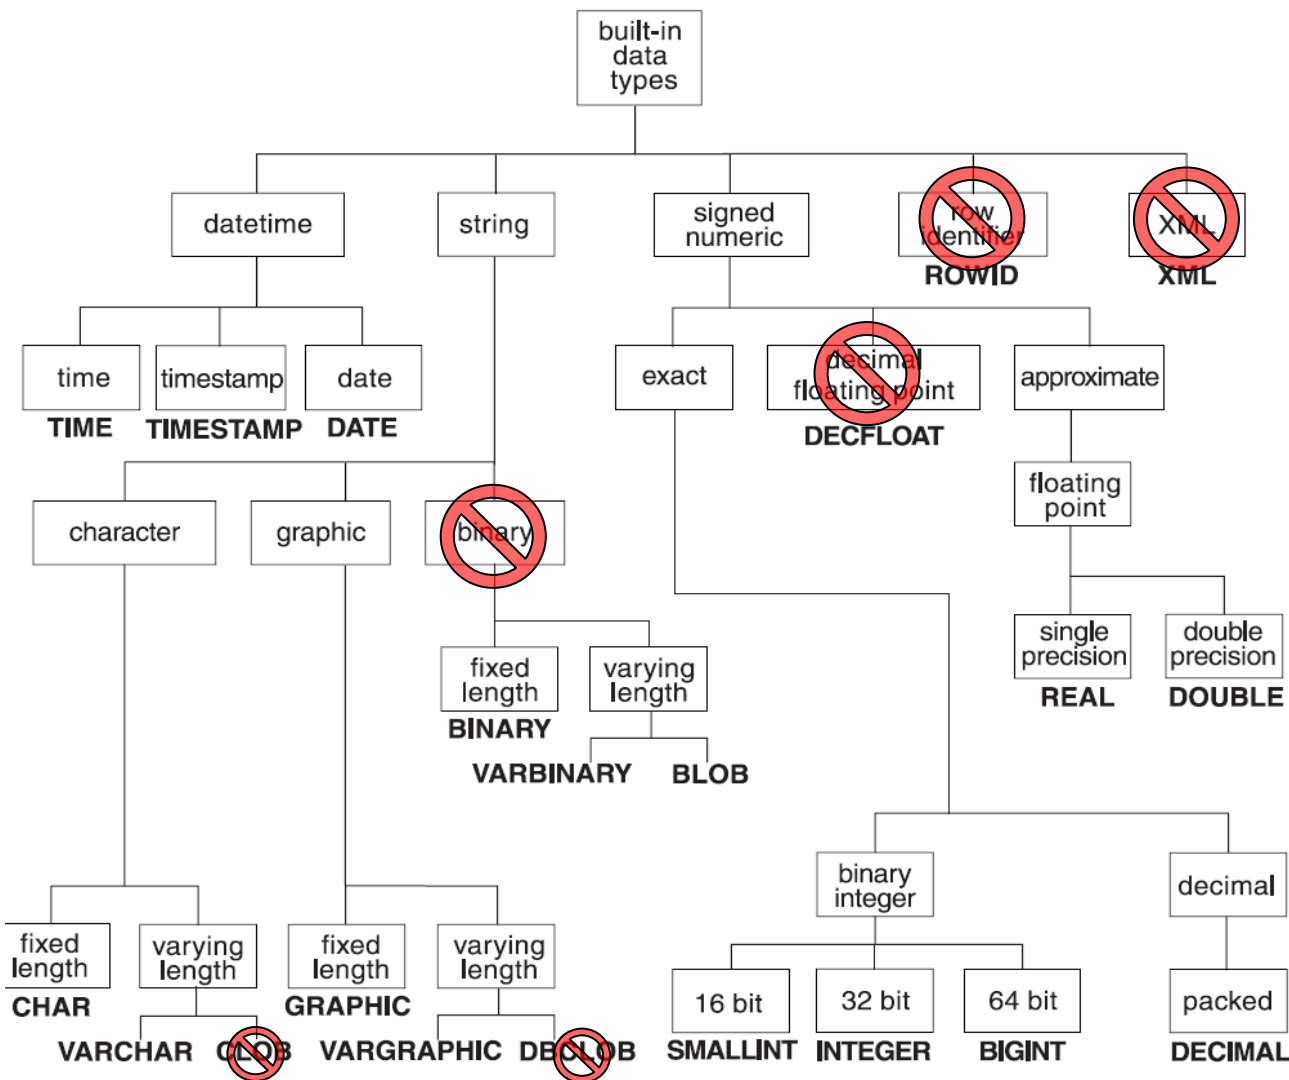
- Limited support for very large dimension tables
  - Especially if the predicates on them are not selective
- No static SQL
- No full outer join, no right outer join
- Only equi-joins (no range join predicates)
- No queries that do not include at least one fact table
- No queries that spread across multiple marts
- Not all DB2 functions
  - No mathematical functions such as SIN, COS, TAN.
  - No user defined 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

**SHARE**  
in Anaheim  
2011




*Pause*  
*simply for effect*

# Data Types Support



## Not supported:

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




**SHARE**  
Technology • Connections • Results

# Prerequisites


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

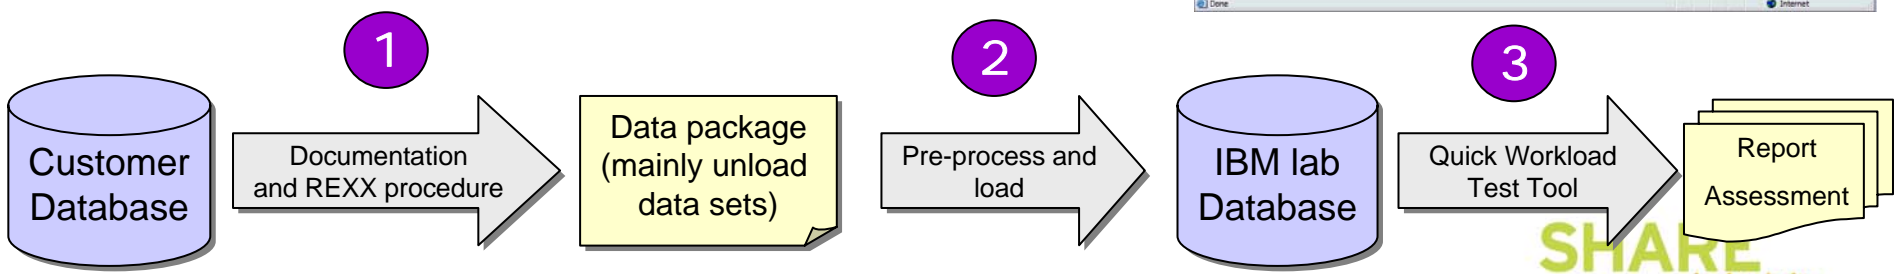
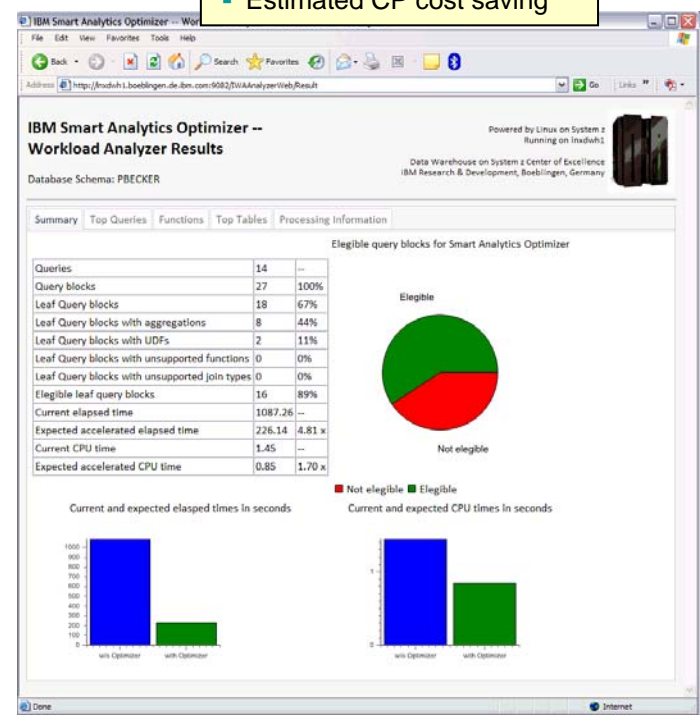



# Quick Workload Test

- 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






**SHARE**  
Technology • Connections • Results

# 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](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](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



**SHARE**  
Technology • Connections • Results

# 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





**SHARE**  
Technology • Connections • Results

# 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



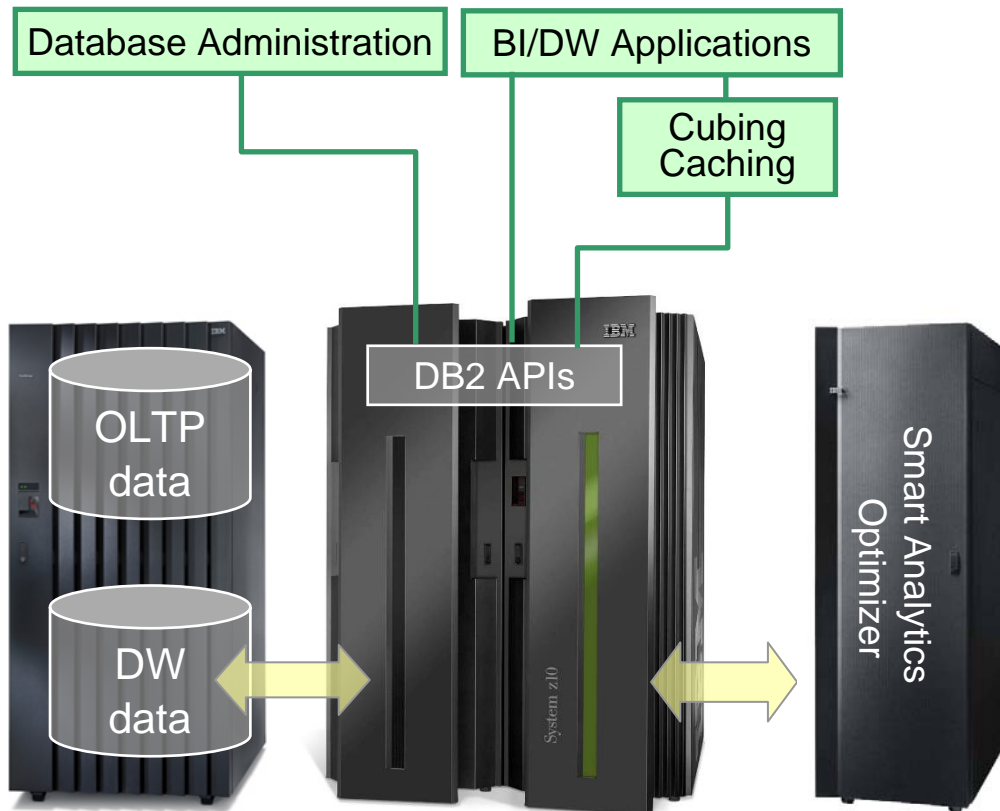


**SHARE**

Technology · Connections · Results

# 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



# 50TB System z Data Warehouse Benchmark



- Poughkeepsie Lab / Silicon Valley Lab joint effort
- Deliver proof points of System z scalability in Business Intelligence environment
- Establish capability of System z to scale to larger volumes
- Develop best practices of managing large data warehouses
- Drive unique value of System z
  - Workload Manager: capabilities to manage mixed workloads
  - Operational BI: large volumes of users, smaller queries
  - zIIP: lower cost BI solution
  - Data compression: minimal overhead with hardware compression




[SG24-7674](#)

**SHARE**  
in Anaheim  
2011

# Co-locating Transactional and Data Warehouse Workloads


- Poughkeepsie Lab / Silicon Valley Lab joint effort
- Managing analytic and traditional workload best practices
- Using workload manager (WLM) to manage mixed workloads
- OLAP vs data warehouse workloads
- Data sharing vs non data sharing considerations
- Single vs multiple LPAR configurations
- Implementation considerations
- Benchmark results




[SG24-7726](#)

**SHARE**  
in Anaheim  
2011

# Other Significant System z Data Warehouse Redbooks



[SG24-7637](#)  
[Enterprise Data Warehousing with DB2 9 for z/OS](#)




[SG24-7813](#)  
[InfoSphere Warehouse: A Robust Infrastructure for Business Intelligence](#)




[REDP-4606](#)  
[Using IBM System z As the Foundation for Your Information Management Architecture](#)





# More Significant System z Data Warehouse Redbooks

**SHARE**  
Technology • Connections • Results



**IBM**

C. Nin Lai

## Workload Management for DB2 Data Warehouse

**Abstract**

This IBM® Redpaper addresses the workload management issues encountered in a DB2 warehouse environment. It delivers a set of guidelines for overcoming these obstacles by utilizing the z/OS® Workload Manager. It also provides tips and techniques of properly implementing a Workload Manager policy.

**Management of large queries**


A key objective of workload management is to prevent "loner" queries from monopolizing system resources. All issue here are not the long queries created accidentally by users, such as forgetting to code a join predicate. Many often the long queries reflect legitimate business queries. Certainly they need to be executed, however, unless they are of a critical nature and answers are required as soon as possible, they shouldn't pose a significant impact to the short queries by consuming most of the system resources.

One technique is to use a predictive governor to screen all the queries prior to execution. Queries that are believed as long would be scheduled for execution during off-peak hours. But this approach proves to be labor intensive and drains valuable DBA's time. Besides, the governor can't catch all the long-running queries. Occasionally some slip through and take away system resources from the short queries.

Another solution calls for using a runtime governor to monitor resource usage during query execution. Queries that consume more than a threshold of CPU cycles will be cancelled automatically. This isn't a good solution either, because a high threshold would waste many CPU cycles prior to query cancellation. On the other hand, a low threshold would trigger many false alarms, cancelling many short or medium queries unnecessarily.

© Copyright IBM Corp. 2014. All rights reserved. [ibm.com/redbooks](http://ibm.com/redbooks) 1

[REDP-3927](#)  
[Workload Management for DB2 Data](#)




## Best Practices for SAP Business Information Warehouse on DB2 UDB for z/OS V8

Benefits of SAP Business Warehouse on z/OS

Best Practices

Troubleshooting




IBMs Data  
Dave Bellon  
Nick Dylek  
Andrew Harman  
Theodor Kuebler  
Ting Ly  
Mike Maritz  
Bala Prabhakar  
Michael Theissen

**Redbooks**

[ibm.com/redbooks](http://ibm.com/redbooks)

[SG24-6489](#)  
[Best Practices for SAP BI using DB2 9 for z/OS](#)

Draft Document for Review August 24, 2010 9:22 pm



**IBM**

IBM White  
Parvez Kaniad  
Fernando Neuge

## IBM zEnterprise BladeCenter Extension Model 001

IBM® has extended the role of the System z10® family of servers by incorporating a piece of infrastructure based on the IBM BladeCenter®. It is called the IBM zEnterprise™.

The zBX brings computing capacity of systems in blade form factor to the System z10 servers. The configuration and management of the zBX is provided by z10 in a seamless, integrated way.

Combined, the System z10 and zBX create a workload optimized technology that supports an integrated hardware and software solution, named the IBM Smart Analytics Optimizer.

In this Redpaper™ we introduce the zBX Model 001 and detail its components. We also explain the basic concepts of the IBM Smart Analytics Optimizer solution and the building blocks for the zBX.

The information in this paper can be used for planning purposes and to help define the configurations that best fit your requirements.


This Redpaper discusses the following topics:

- "zBX concept" on page 2
- "zBX hardware description" on page 3
- "Smart Analytics Optimizer overview" on page 4
- "IBM Smart Analytics Optimizer software" on page 10
- "zBX capacity for zBX" on page 18
- "zBX environment" on page 19
- "Installation overview" on page 25

© Copyright IBM Corp. 2010. All rights reserved. [ibm.com/redbooks](http://ibm.com/redbooks) 1

[REDP-4668](#)  
[IBM zEnterprise BladeCenter Extension](#)


**SHARE**  
in Anaheim  
2011



**SHARE**  
Technology • Connections • Results

# More information on zEnterprise

- IBM zEnterprise Announcement Landing Page: [ibm.com/systems/zenterprise196](http://ibm.com/systems/zenterprise196)
- IBM zEnterprise HW Landing Page: [ibm.com/systems/zenterprise196](http://ibm.com/systems/zenterprise196)
- IBM zEnterprise Events Landing Page: [ibm.com/systems/breakthrough](http://ibm.com/systems/breakthrough)
- IBM Software: [ibm.com/software/os/systemz/announcements](http://ibm.com/software/os/systemz/announcements)
- IBM System Storage: [ibm.com/systems/storage/product/z.html](http://ibm.com/systems/storage/product/z.html)
- IBM Global Financing: [ibm.com/financing/us/lifecycle/acquire/zenterprise/](http://ibm.com/financing/us/lifecycle/acquire/zenterprise/)
- Global Technology Services:
  - vanity: [ibm.com/services/zenterprise](http://ibm.com/services/zenterprise)



**SHARE**

15 • Results

धन्यवाद

Hindi

多謝

Traditional Chinese

ขอบคุณ

Thai

Спасибо

Russian

Gracias

Spanish

Thank You

English

شكراً

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

**SHARE**  
in Anaheim  
2011



**SHARE**  
Technology • Connections • Results

# Willie Favero

## Senior Certified Consulting IT Software Specialist

Dynamic Warehousing on System z Swat Team

IBM Silicon Valley Laboratory

IBM Academic Initiative Ambassador for System z

IBM Certified Database Administrator - DB2 Universal Database V8.1 for z/OS

IBM Certified Database Administrator – DB2 9 for z/OS

IBM Certified System Administrator – DB2 9 for z/OS

IBM Certified Database Administrator – DB2 10 for z/OS

IBM zChampion

[wfavero@attglobal.net](mailto:wfavero@attglobal.net)