

# ISAO – IBM Smart Analytics Optimizer

## Accelerating DB2 for z/OS



James M. Wilson

Consulting IT Specialist – System z

[jamesmwi@us.ibm.com](mailto:jamesmwi@us.ibm.com)

March 8, 2010

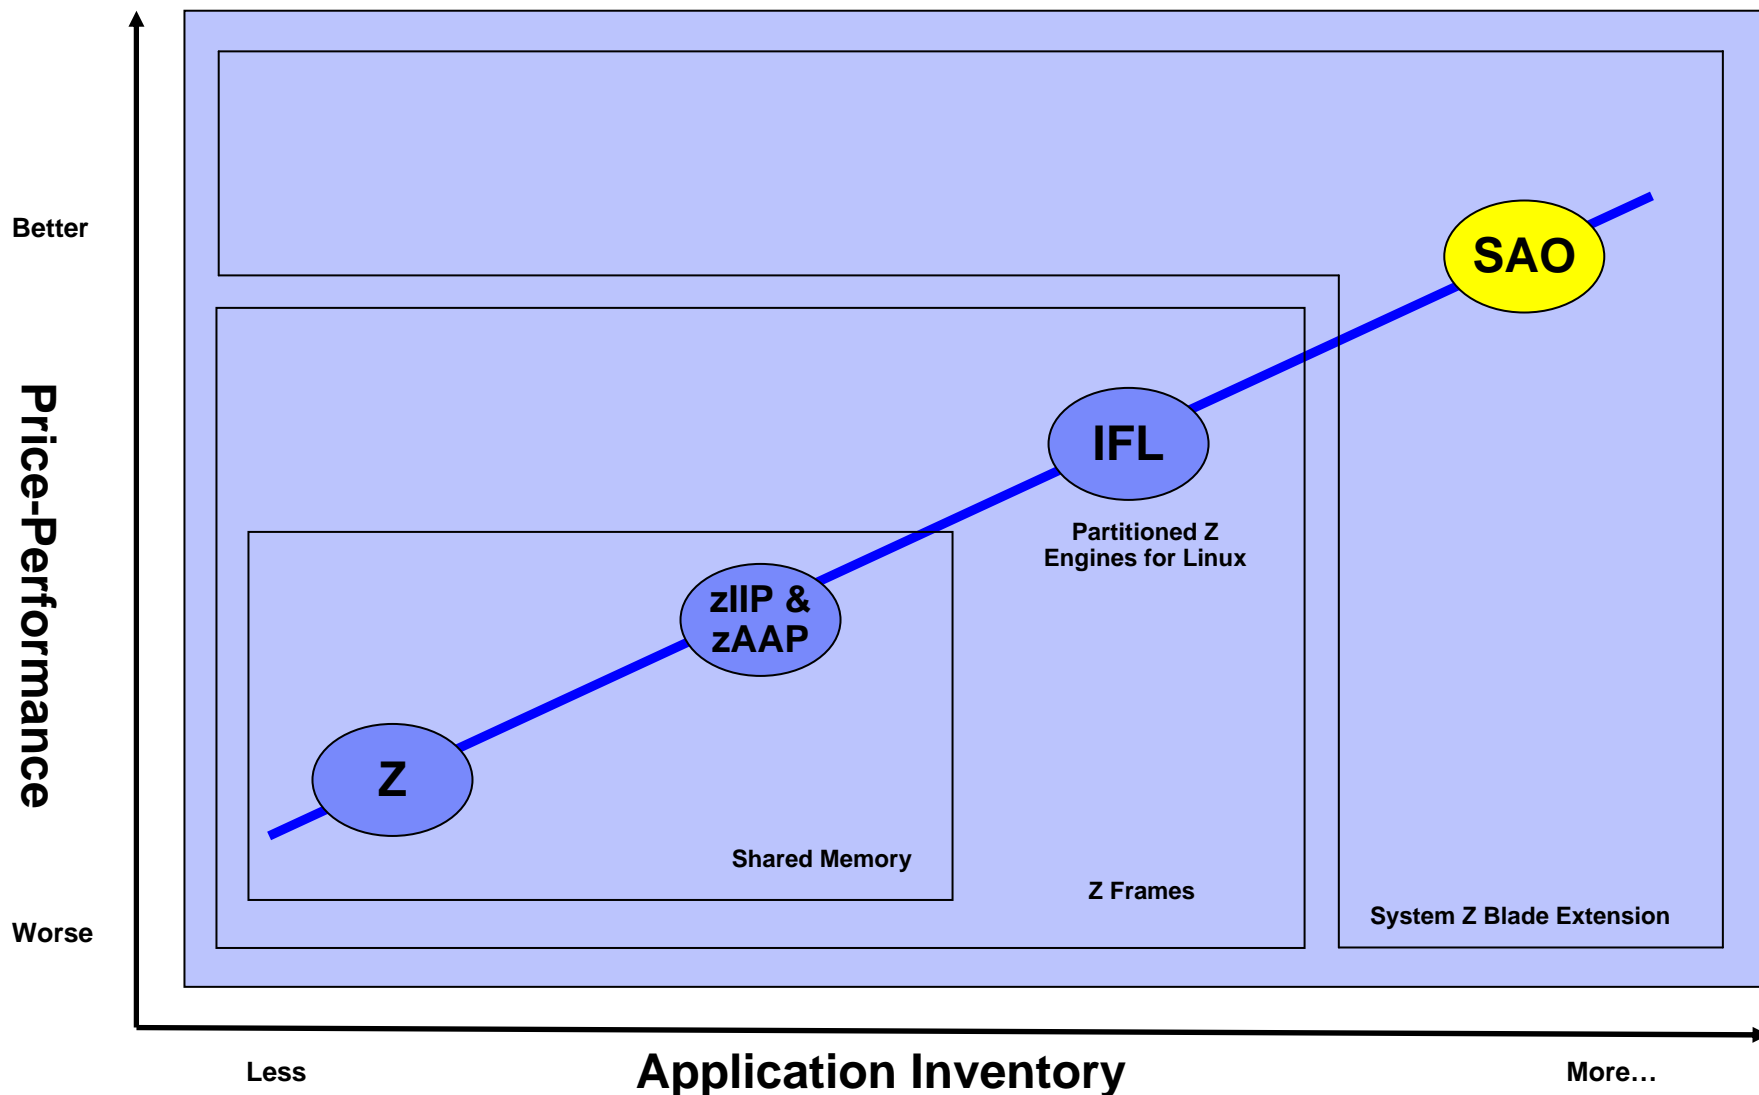
# Challenges: Performance, Scalability, TCO

- **Modern BI/DW requirements such as orders of magnitude faster query execution call for new approaches**
- **DB2 for z/OS has the first class QoS characteristics, however, it is 'only' a relational DBMS**
  - RDBMSs are the most widely used sophisticated data repositories with huge ecosystems of applications built on top of them, however, they are not the ultimate answer for everything
  - Deficiencies come from the way the data is stored and managed
    - optimizing for limited cache
    - supporting limited CPU parallelism
  - RDBMSs attempt to address performance and scalability challenges with their standard tools of the trade: indexing, prebuilt aggregates, MQTs, ...
    - Requires very sophisticated tools and top DBA expertise which significantly drives up TCO
    - Increasingly not good enough due to ad-hoc, unpredictable nature of the DW/BI requests

# ***Smart Analytics Optimizer:***

## *Accelerating DB2 for z/OS OLAP Workloads*

# SAO – Platform View from a System HW perspective (extension to the system z Specialty engines)



# IBM Smart Analytics Optimizer

## What is it?

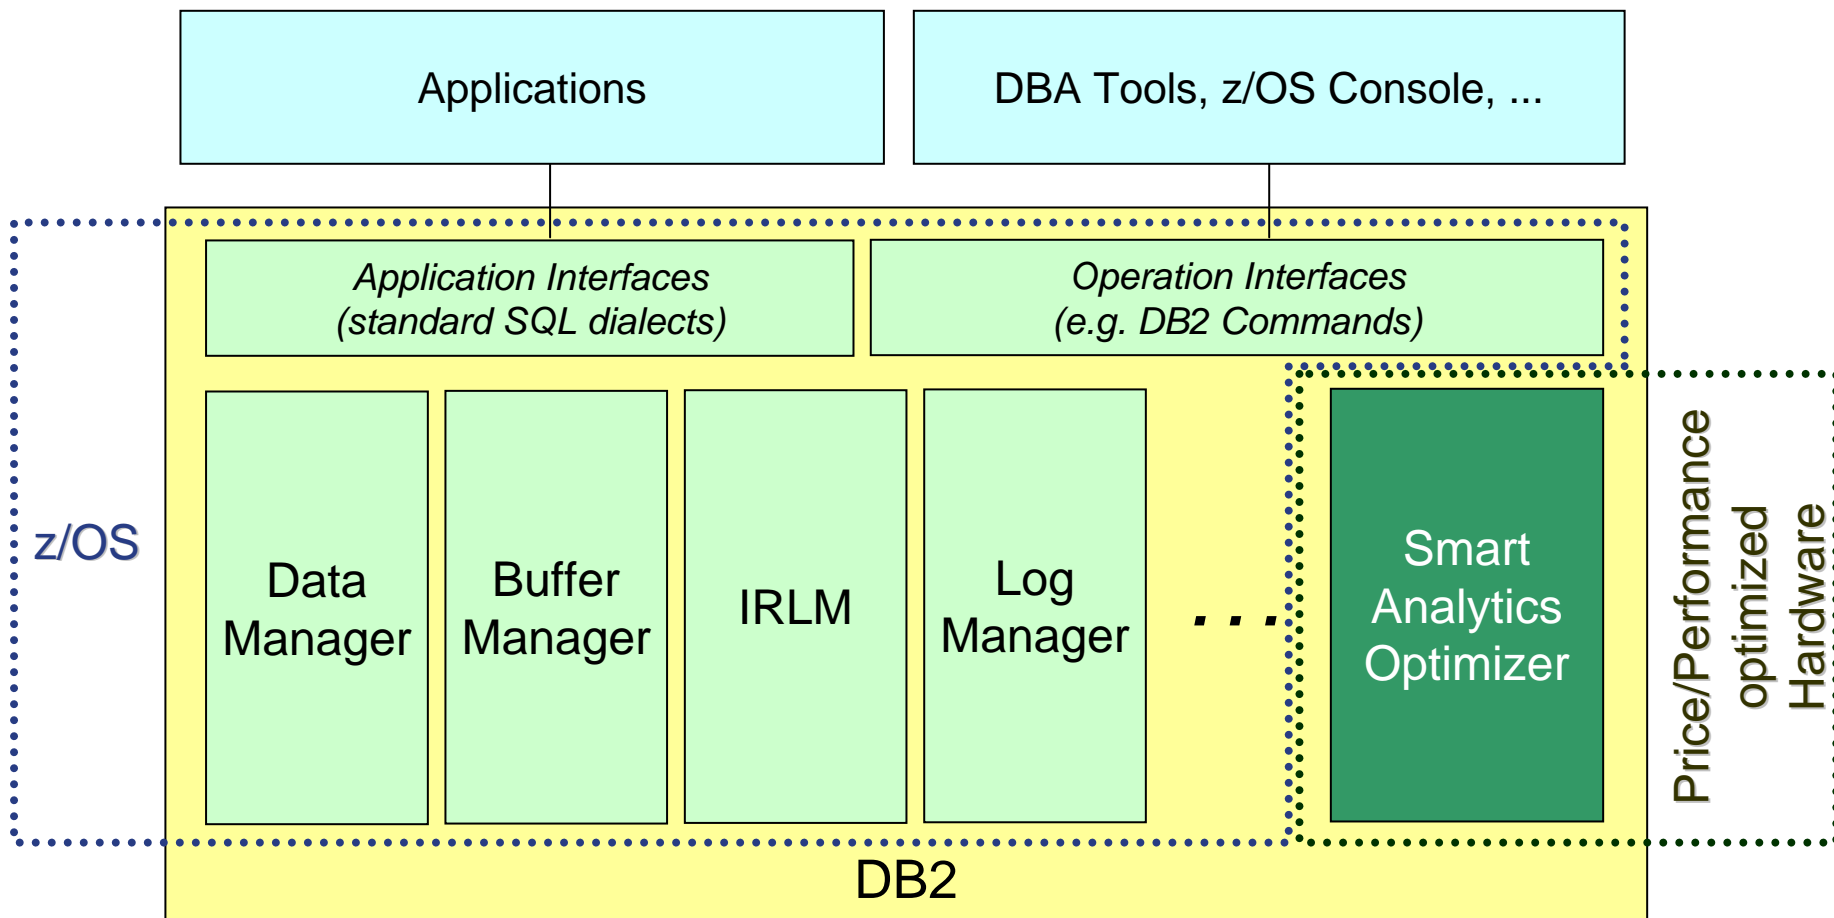
- ✓ A high performance, appliance-like add-on delivering order-of-magnitude faster, predictable, analytic query responses transparently to all users.



## How is it different

- **Performance:** through implementing leading technology trends: hybrid row/column store, predicate evaluation on compressed data, multi-core and vector optimized algorithms
- **Integration:** the data continues to be managed and secured by the most reliable database platform - DB2 for z/OS
- **Self-managed workloads:** queries are executed in most efficient way irrespective of their type (OLTP vs. OLAP)
- **Transparency:** applications connect to DB2 and are entirely unaware of ISAO presence
- **Simplified administration:** appliance form factor and hands-free everyday operations, reduced need for complex query tuning


# Smart Analytics Optimizer (SAO) – Platform View



# SAO Features

- **A special purpose, network attached blades system**
  - Offload typical DW queries from traditional database server to the accelerator
  - Based on IBM research prototype – defining new frontiers in performance and scalability
- **No changes to the applications**
  - Applications continue to attach to DB2.
  - DB2 transparently to the applications exploits the accelerator when applicable query needs to be executed
  - Full fencing and protection of DB2 against possible accelerator failures
- **Improving performance of typical DW queries by orders of magnitude**
- **Achieving linear scaling with the number of CPUs**
- **Reducing need for tedious tuning of DB2 (MQTs, indexes, etc.)**
- **Significantly improved price/performance and TCO as a combined effect of:**
  - Offloading very CPU intensive operations from System z
  - Using price/performance optimized hardware
  - Orders of magnitude performance improvement for offloaded queries
  - Reduced DBA effort for tuning offloaded queries
- **Appliance-like form-factor**
  - User/reference guide assisted installation, initial configuration
  - Hands free operations

# Adding the Accelerator without changing the environment



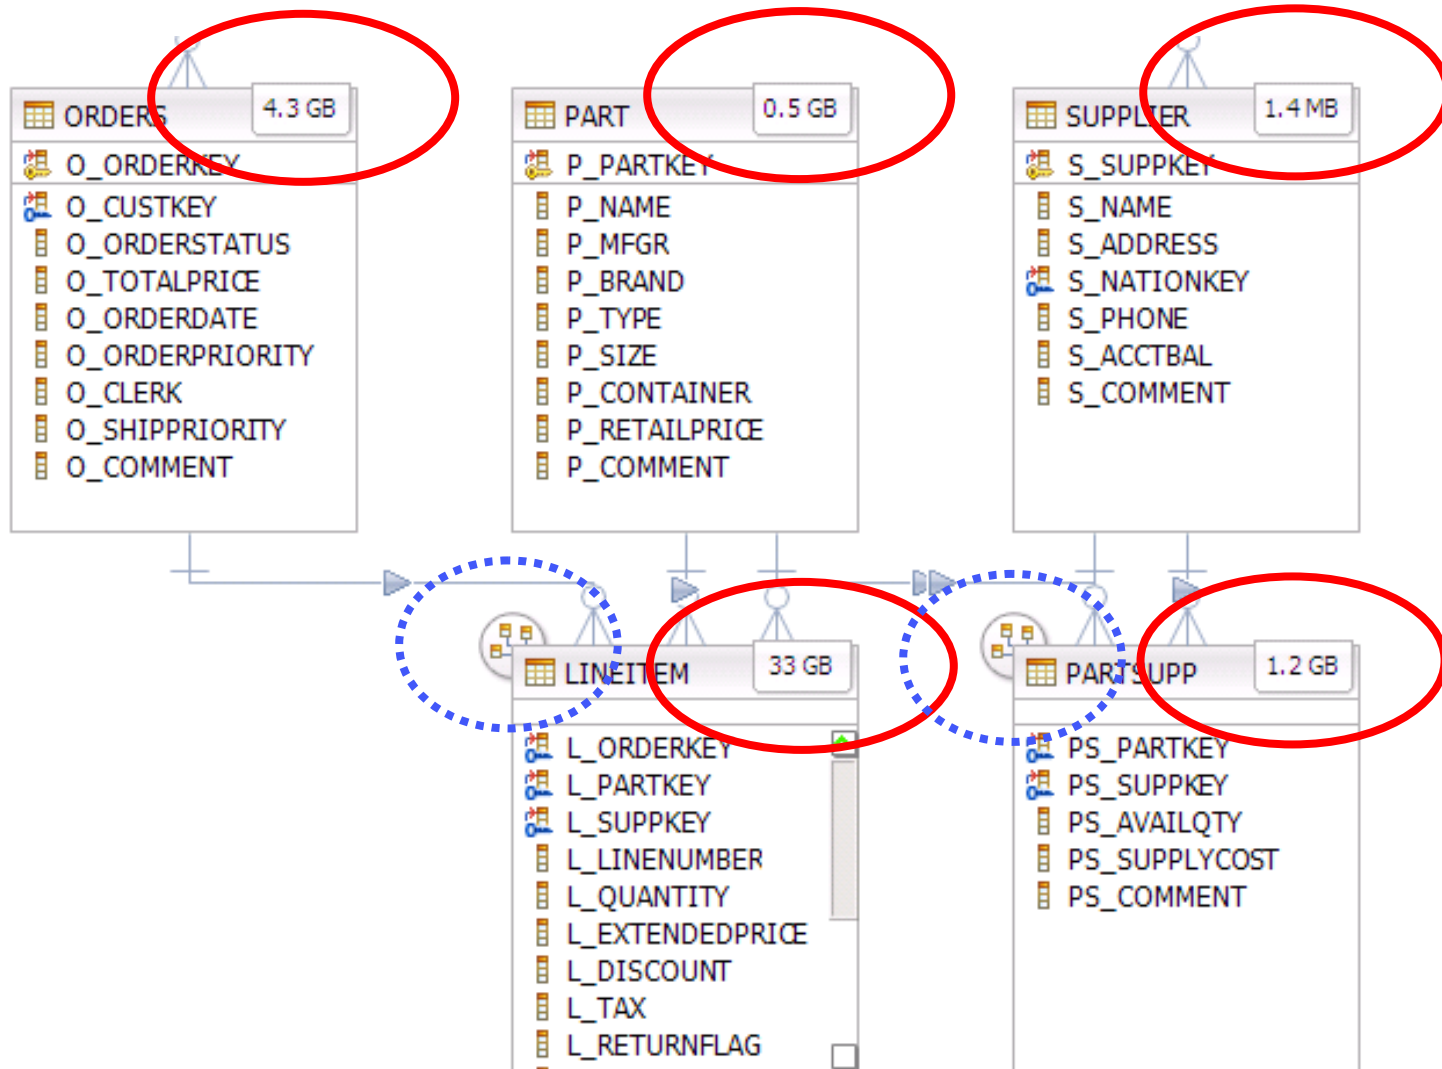


# Defining which data to accelerate

- 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 it's own catalog. The related data is read from the DB2 tables and transferred to the InfoSphere Warehouse Accelerator.
- The InfoSphere Warehouse Accelerator transforms the data into a highly compressed, scan optimized format which is kept locally (in memory) on the Accelerator



## 4b. Zoom In: Size Estimates and Fact table property



# SAO Load creates In-Memory Replica of Data

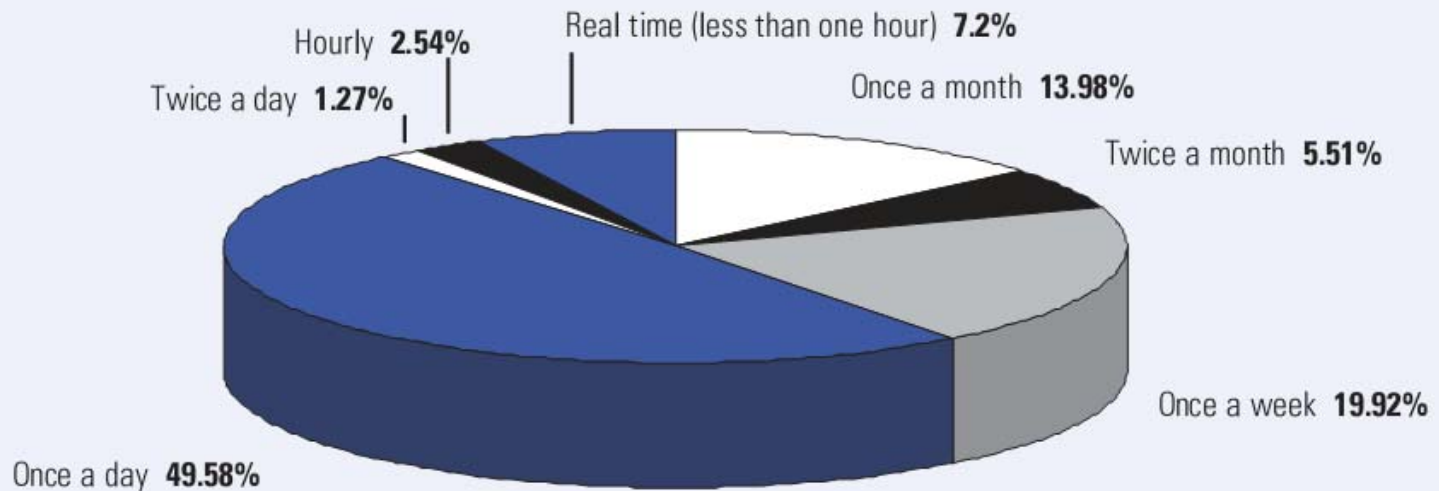
- **A highly compressed, version of all data of the MARTs is kept in the memory of the InfoSphere Warehouse Accelerator**
- **The data in memory is a snapshot of the original data which is still stored within DB2.**
- **Data changes on the original data need to be captured on the DB2 side and applied on the memory structures of the InfoSphere Warehouse Accelerator**
  - Planned latency between data change capture and applying the changes might cause different versions of the data to be queried.
  - Queries see snapshot data (for some time in the past) as in MQT approach – if a mart is enabled for acceleration queries will be eligible to be routed to the snapshot

Support for SET CURRENT REFRESH AGE

## IDUG study on update frequency – periodic update is OK

- 90 % populate their warehouse once a day or less frequently

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

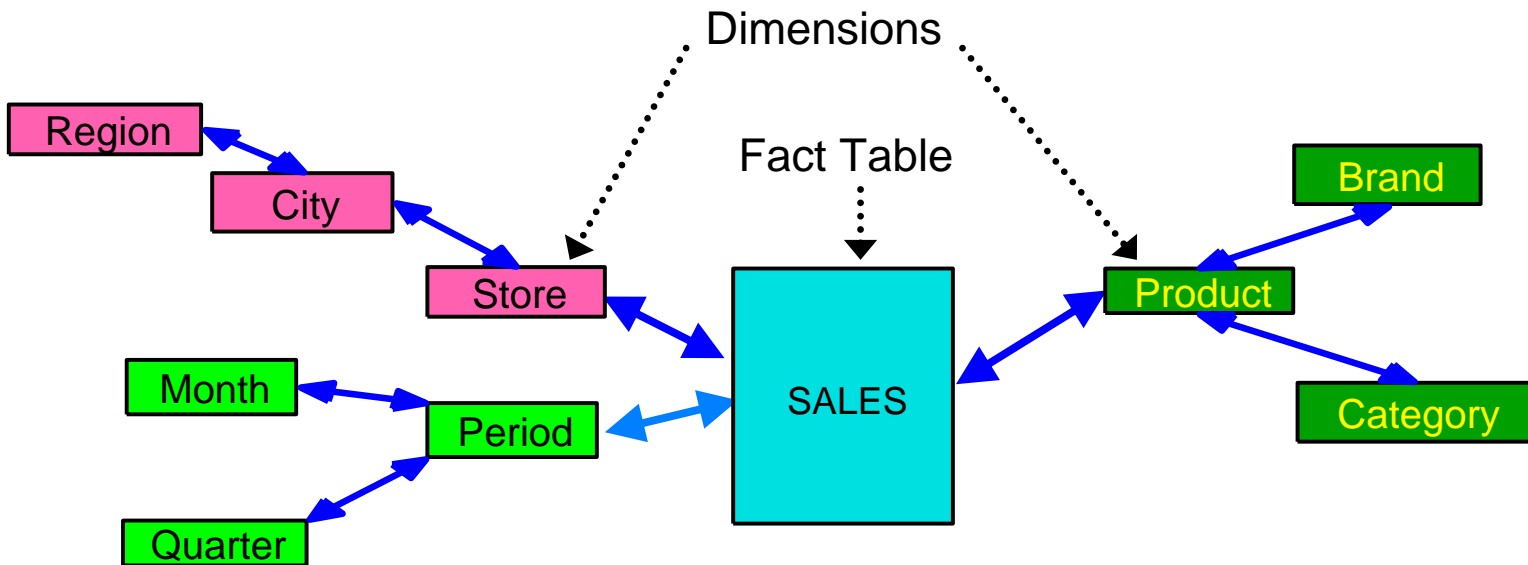


## Data Maintenance (SAO Currency)

- **In the initial release, we concentrate on DWH specific population methods:**
  - LOAD RESUME
  - Roll In/Out of Ranges/Partitions
  - Delete complete Ranges/Partitions
  - Reload a complete Range/Partition
  - Based on the DB2 UNLOAD utility
- **LOG based data change capture under evaluation**
  - Have a stored procedure which reads the LOG and delivers all SAO relevant changes to the appliance
  - Might be useful for minor updates and/or INSERT based table population
- **Other data maintenance scenarios which are not covered by partitioned LOAD or incremental updates cause a full reload of a MART.**

# Target Market: Business Intelligence (BI)

- Characterized by:
  - “Star” or “snowflake” schema:



- Complex, ad hoc queries that typically
  - Look for trends, exceptions to make actionable business decisions
  - Touch large subset of the database (unlike OLTP)
  - Involve aggregation functions (e.g., COUNT, SUM, AVG,...)
  - **The “Sweet Spot” for SAO!**

# What the Accelerator is designed for

- Fast scans over large (fact) tables
- OLAP-style queries over large fact tables in relational star schema with grouping and aggregations

```
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
```

# Matching of queries for SAO support

- **DB2 for z/OS will reuse partial MQT functionality to find out which queries are eligible for SAO offload and which are not.**
- **This implies that a subset of the MQT restrictions is inherited (at least for release 1 of SAO)**
  - Only a single query block at a time can be routed to SAO (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 SAO in the outer query (DB2 would need to pass a subselect result to SAO)




# The following queries run in DB2 and not accelerated

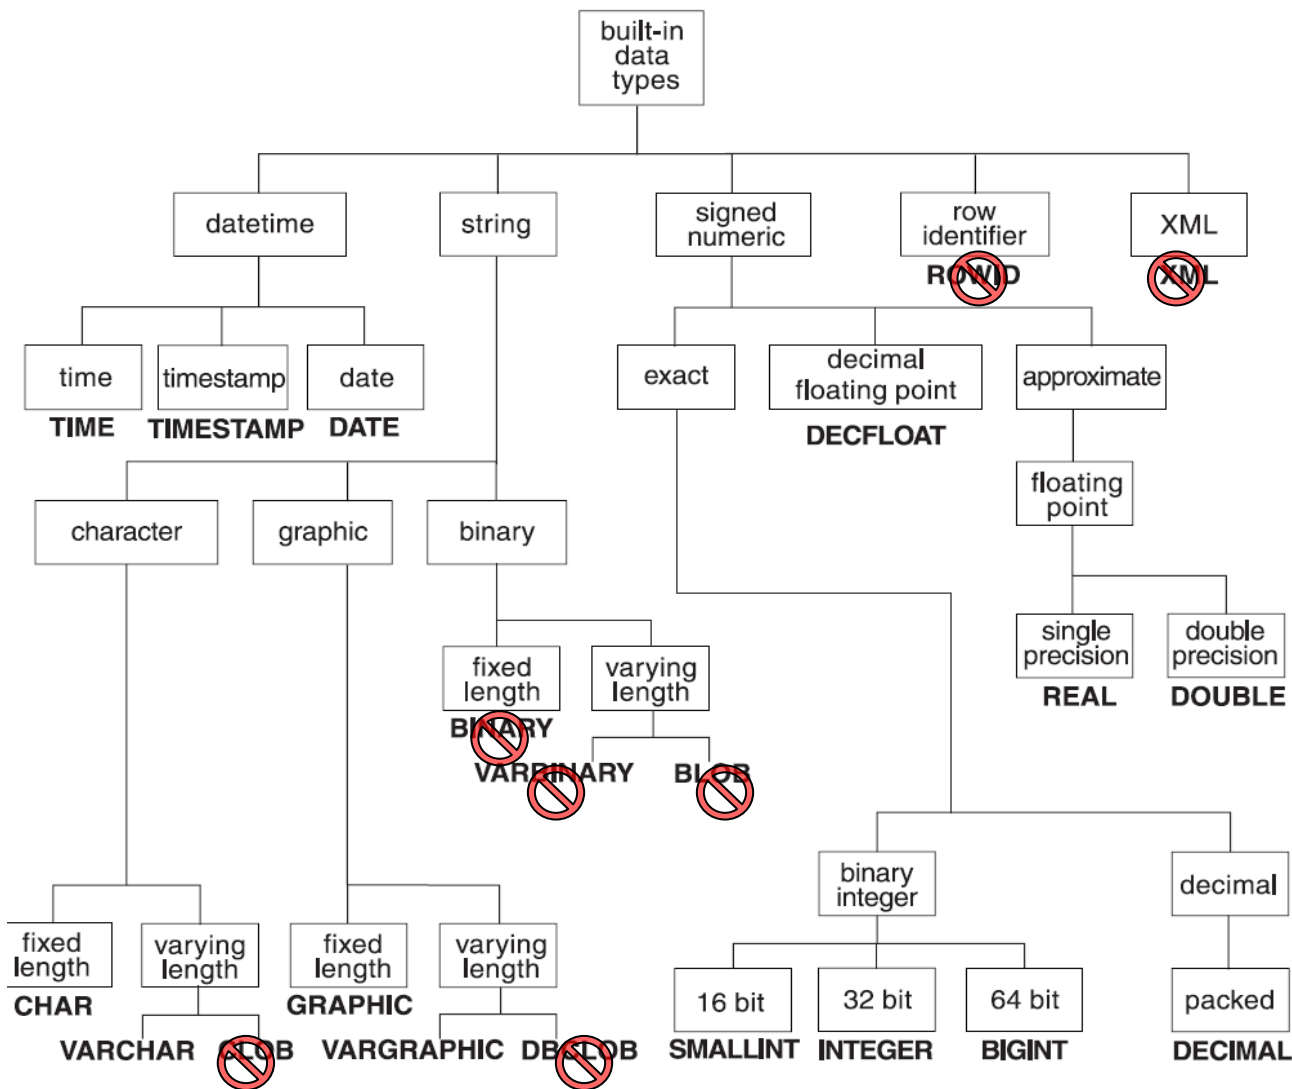
- DB2 matches one query block at a time and routes a single query block to SAO
  - ➔ Queries which consist of several Query Blocks, are not seen as whole query by the accelerator but only as single, independent blocks
- Outer query block containing inner query blocks is not routed
  - ➔ The results of subselects can not be used by SAO in the outer query (DB2 would need to pass a subselect result to SAO)
- The following SQL may generate multiple query blocks
  - **subselects and common table expressions**
    - subselects in quantitative predicates (SOME, ANY, ALL)
    - EXISTS or IN predicate with subselects
  - UNION, INTERSECT, EXCEPT
  - **UNION ALL views -> can only route inner query blocks**
- **Only inner join and Fact left outer join Dimension (no full, right outer join)**
- **Most DB2 built in functions are supported except**
  - mathematical functions like sin, cos, tan, exp, correlation
  - User defined functions
  - Advanced string functions like locate, left, like, overlay, position
  - Advanced OLAP functions like rank, dense\_rank, row\_number, rollup, cube

# 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 SAO (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**



# Support for data types



## Not supported:

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

# Summary why a query may not be routed

1. **Because it uses CURRENT REFRESH AGE = 0**
2. **Because it contains syntax that is not supported (e.g. Subselect or full outer join)**
3. **Because the accelerator or AQT is disabled**
4. **Because it references a table or column that is not in the accelerated mart (may be due to unsupported datatypes)**
5. **Because the query does not reference a fact table**
6. **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

## Added Explain Table for Queries are NOT Accelerated

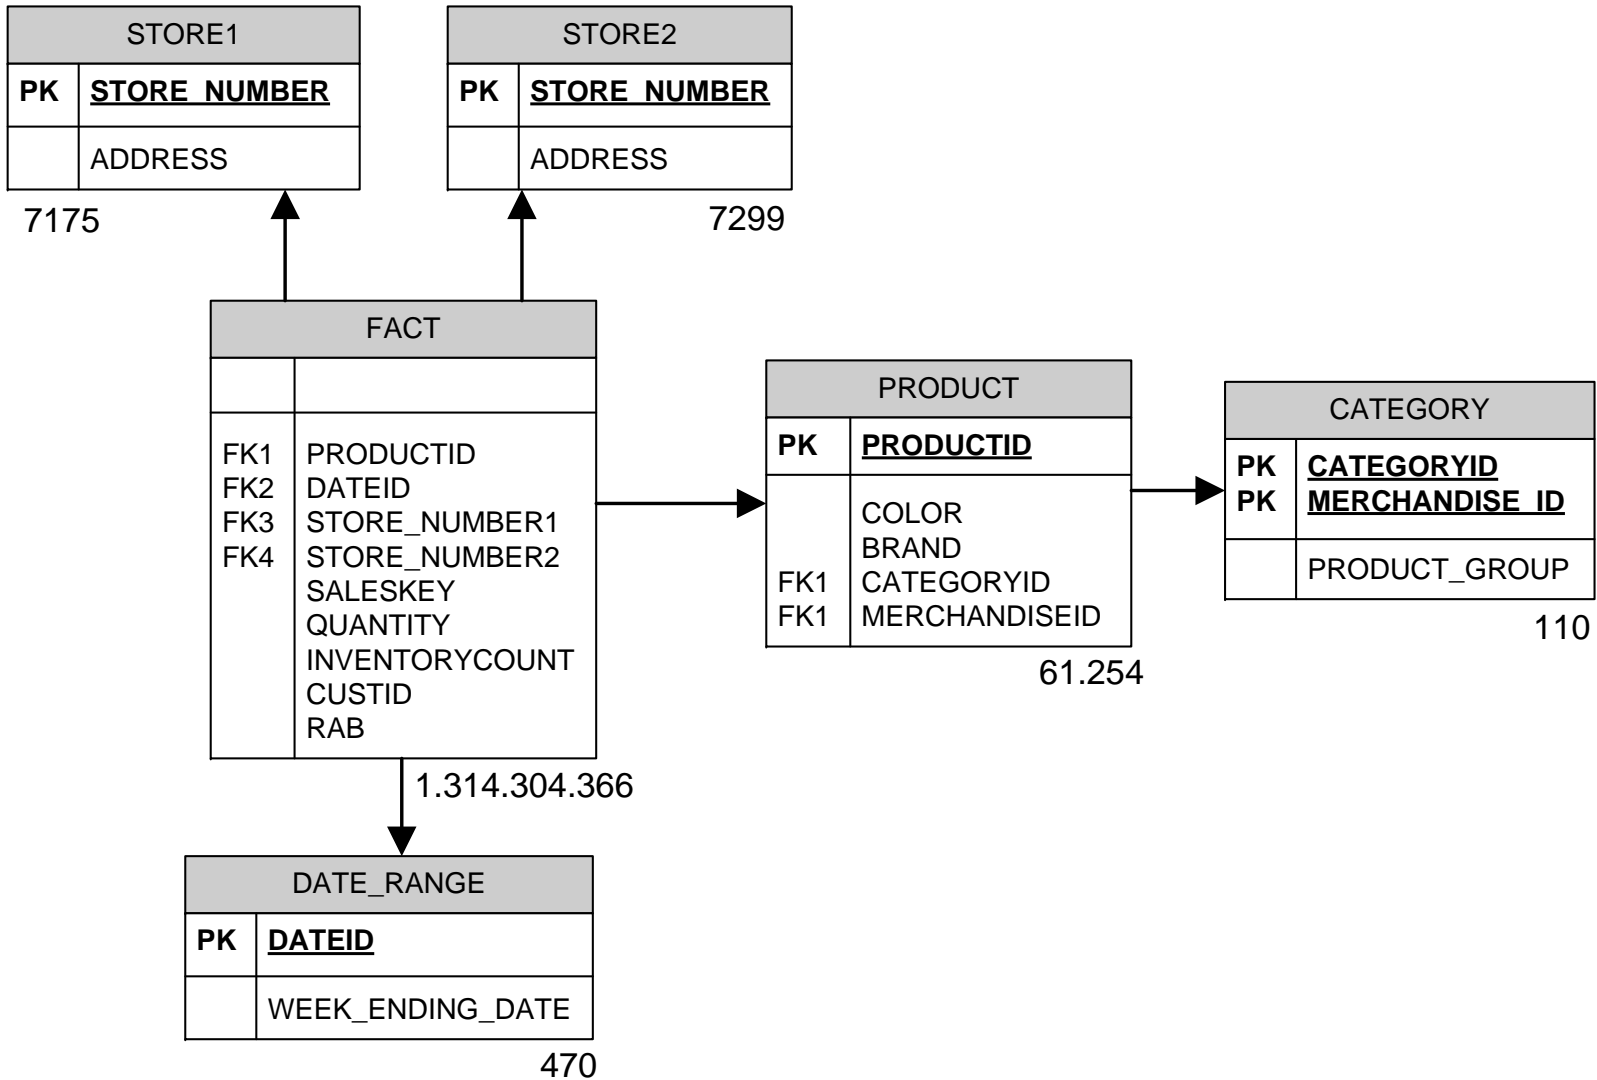
A new EXPLAIN table is added to show

- 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.
- The DDL for this new EXPLAIN table is as follows:


```
CREATE TABLE DSN_QUERYBLOCKINFO_TABLE(  
  QUERYNO INTEGER NOT NULL WITH DEFAULT,  
  QBLOCKNO SMALLINT NOT NULL WITH DEFAULT,  
  ...  
  QB_REASON SMALLINT NOT NULL WITH DEFAULT,  
  QB_INFO CLOB(2MB) NOT NULL WITH DEFAULT,  
) CCSID UNICODE;
```

- Column "QB\_INFO" contains data in XML format. This column would contain the objects that caused acceleration not to be chosen..


# The tested (sub-)schema



# Average deviations: 5.8s vs. 1520s



# Query Execution Times: DB2 with SAO - Speedup






# Summary

**Goal: Constant, short response time for BI queries in DB2 for z/OS without tuning (orders of magnitude improvement)**

- 1. In-memory database – no disk I/O**
- 2. Compression scheme that allows fixed length tuples and predicate evaluation on compressed data**
- 3. Brute-force scans for all queries instead of tuning of indexes using massive parallelisation – scales with number of blades**
- 4. Exploiting multi-core architecture and SIMD instructions of commodity hardware**
- 5. Pre-join dimension tables to fact during load (schema melting)**
- 6. Appliance attached to DB2 to which DB2 will transparently route queries without changing user application**

# What is available on System z & **How** to implement




# Some key Redbooks



- **Enterprise Data Warehousing with DB2 9 for z/OS**
  - <http://www.redbooks.ibm.com/abstracts/sg247637.html>
- **50 TB Data Warehouse Benchmark on IBM System z**
  - <http://www.redbooks.ibm.com/redbooks.nsf/RedpieceAbstracts/sg247674.html>
  - This is the draft
- **DB2 for z/OS: Data Sharing in a Nutshell**
  - <http://www.redbooks.ibm.com/abstracts/sg247322.html>
- **System Programmer's Guide To: Workload Manager**
  - <http://www.redbooks.ibm.com/abstracts/sg246472.html>
- **Workload Management for DB2 Data Warehouse, REDP-3927**
  - <http://www.redbooks.ibm.com/abstracts/redp3927.html>

# The IBM Smart Analytics System 9600



# IBM Smart Analytics System 9600

## **Includes -**

### ▪ **Hardware**

- *Appliance-Like delivery built on System z10 technology*
- *DS8000 enterprise class storage*
- *Pre-packaged in multiple scale factors to meet any requirement.*

### ▪ **Software**

- *Optimized software stack*
- *Enhance the solution with additional software add-ons*

### ▪ **Services**

- *Installed and ready to use*

### ▪ **Maintenance**

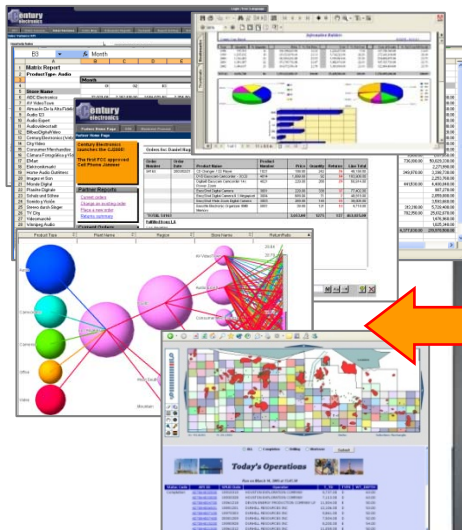
- *Up to 5 years hardware maintenance*



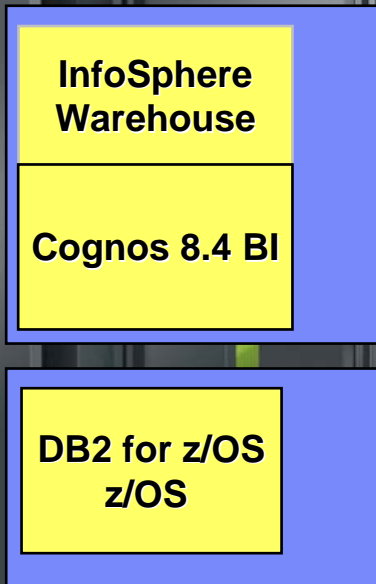
# IBM Smart Analytics System – 9600

## View of Business

*Building an end-to-end BI environment on System z*



## System z



DS8700

LPAR or stand alone

## IBM Smart Analytics System -z Foundation for growth

- Replication Server
- Federation Server/Classic Fed
- IBM Smart Analytics Optimizer

# IBM Smart Analytics System 9600 Software

Deeply Optimized by IBM Experts... Flexible Growth...

## Powerful Data Warehouse and BI Software

- DB2 for z/OS Value Unit Edition (primary) V9  
*Option for MLC*
- DB2 Utilities Suite V9
- DB2 Connect
- InfoSphere Warehouse on System z V9.5.2
- IBM Cognos 8.4 BI for Linux on System z
- z/OS Operating System Stack V1.11
- z/VM 6.1

## Optional Value Priced Add-ons

- Tivoli OMEGAMON for DB2 Performance Expert
- Tivoli Directory Server
- InfoSphere Information Server
- InfoSphere Replication Server
  - Q-Rep, CDC and Event Publisher eligible
- InfoSphere Federation Server plus Classic Federation on System z
- Tivoli ITCAM, ITUAM
- Cognos Now! For Linux on System z

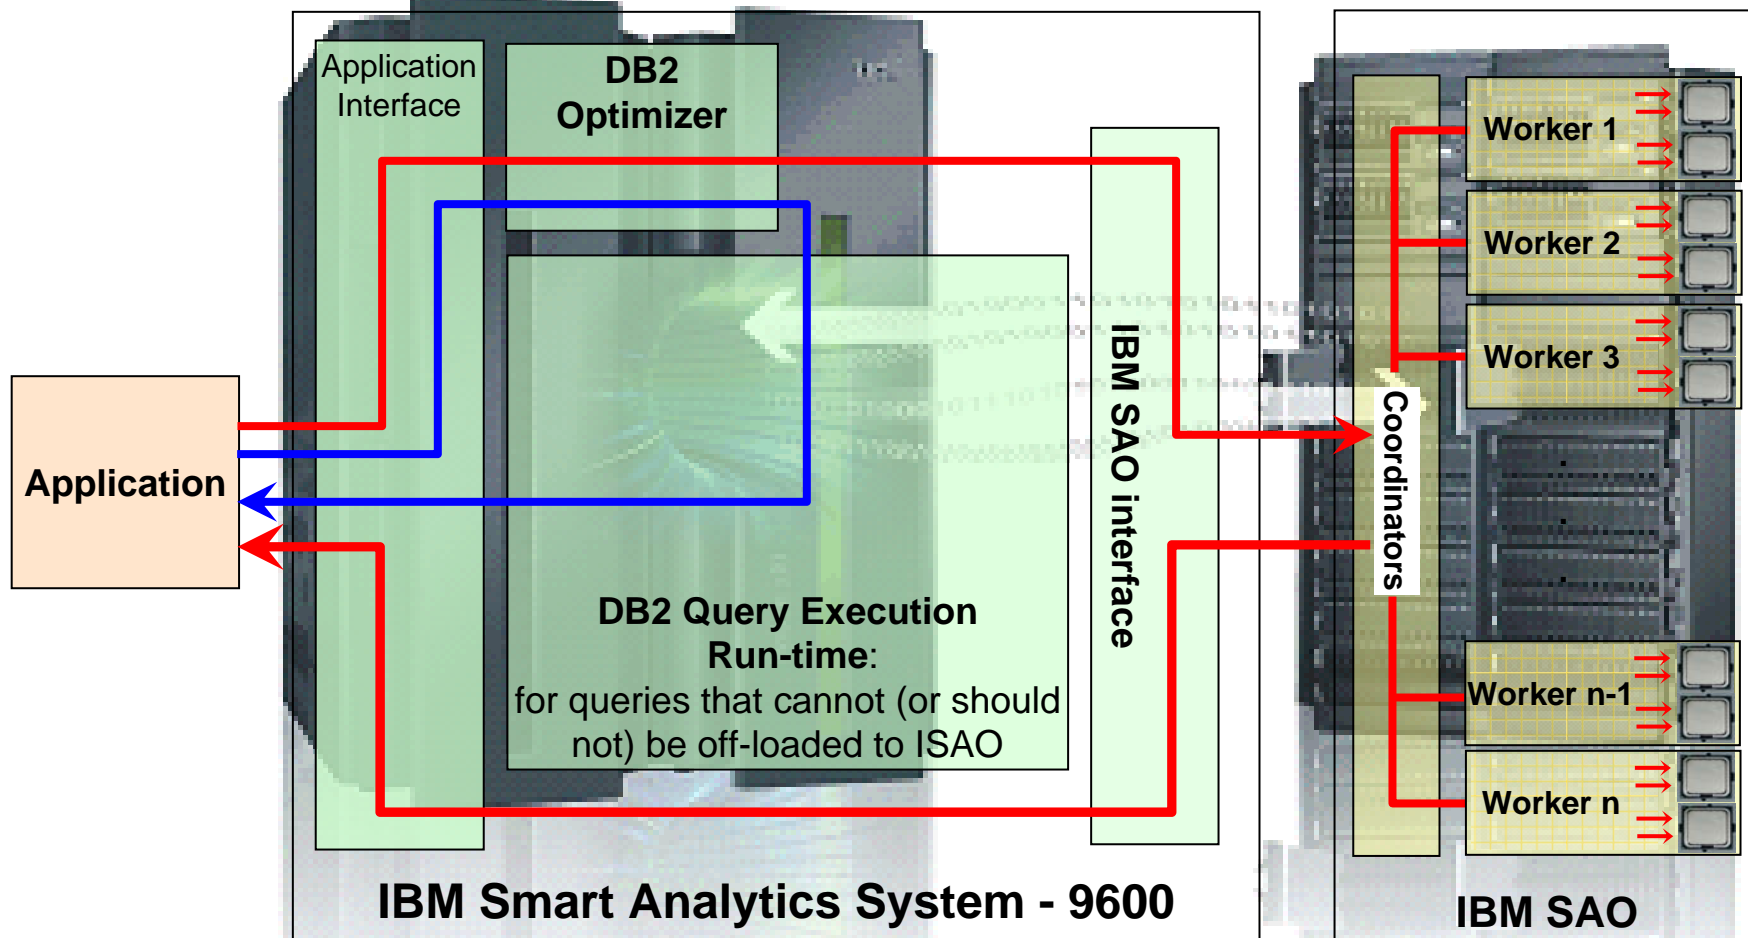


## Comparing the Smart offerings in System Z:

- **IBM Smart Analytics System:**
  - **Complete, end-to-end environment for BI workload**
  - **Processes ALL queries submitted by end-users**
  - **Software:**
    - Includes z/OS, DB2 for z/OS,
    - Linux, InfoSphere Warehouse, Cognos, DB2 Connect
  - **Supports:**
    - Data movement
    - Enduser tools (Cognos)
    - Data Storage (Data warehouse)
  - **Runs in z/OS-DB2 LPAR, Linux on Z LPAR for Tooling**
  - **Is an all purpose environment to deploy any BI workload**
- **IBM Smart Analytics Optimizer:**
  - **Self-contained, closed system, dedicated to processing select queries**
  - **MUST connect TO a DB2 for z/OS environment that is running a BI workload**
  - **Software:**
    - IBM SAO software 56997-AQT that is custom code unique to this offering, ordered as part of the system.
  - **Supports:**
    - Only executes a **SUBSET** of queries that arrive in DB2
  - **Qualifying queries:**
    - Are selected/routed to IBM SAO by DB2 for z/OS
    - Are multidimensional
    - Generally scan the FACT table loaded in memory in the IBM SAO offering



# V2 - Query Execution Process Flow



- Queries executed without IBM SAO
- Queries executed with IBM SAO

# IBM Smart Analytics System

*More Information*

**To learn more about the IBM Smart Analytics System visit:**

<http://www-01.ibm.com/software/data/infosphere/data-warehousing/>



## Collateral:

IBM Smart Analytics System 9600 Webpage:

**System z page:** <http://www.ibm.com/systems/zbi>

### Other Links:

- **Data Warehousing and BI on Z** - <http://www-2000.ibm.com/software/data/businessintelligence/systemz/>
- **Terabyte Club for System z BI customers:** <http://www-2000.ibm.com/software/data/businessintelligence/systemz/terabyte-club.html>
- **Data Governance on System z:** <http://www-2000.ibm.com/software/data/db2imstools/solutions/compliance.html>

Thank  
YOU

