# Informacja relewantna czy optymalna?
# IBM Smart Analytics Optimizer

*Grzegorz T. Kolecki*
kolecki@pl.ibm.com

by Namik Hrle

# Disclaimer



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

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

# Trademarks

# Agenda

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

# Business Challenges

- Changing business requirements
- BI/DW becoming mission critical and require 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
- Sounds like a perfect System z match

# Ultimate Consolidation Opportunity

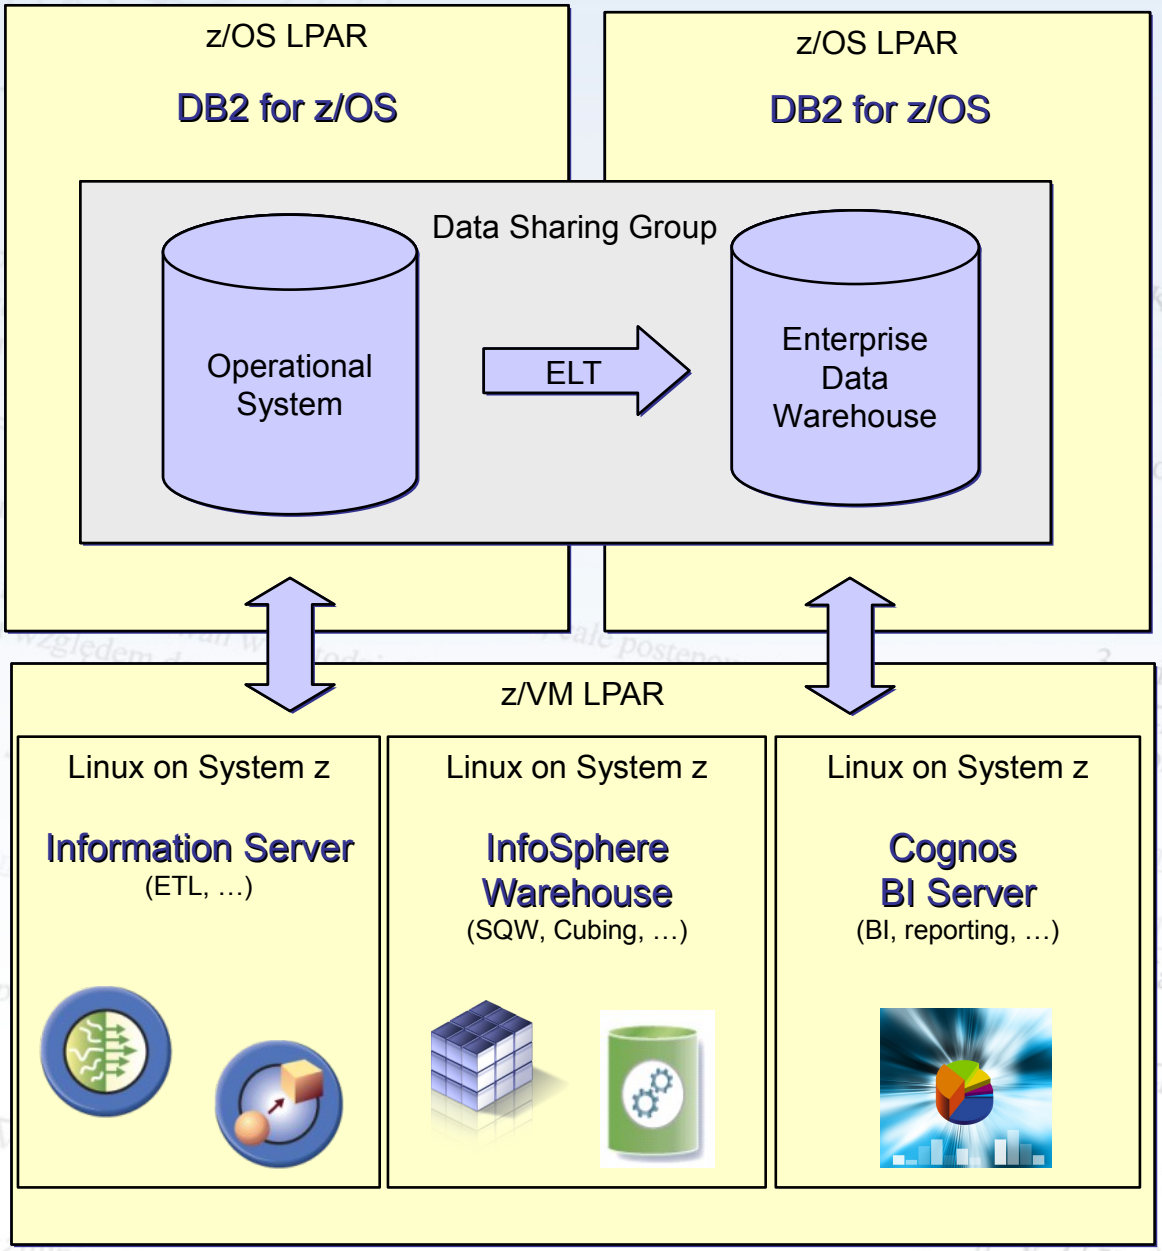


- Significant IBM investment in System z (DB2, IS, ISW, Cognos, ...) warehousing capabilities
- Consolidation of mission-critical data on System z
- Leveraging existing environment, high availability, backup and governance procedures as well as skills
- Efficient data movement within a data sharing group (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

# Technology Trends

- Abundance of computing resources
  - Very large number of processor sockets and cores
  - Massive amounts of real memory
  - Specialized physical data designs: row-store vs. Column-store

- Hybrid systems
  - zEnterprise

- Traditional performance tuning tools of the trade such as indexing, aggregates and MQTs struggling to keep the pace
  - Top DBA expertise and sophisticated tools required
  - Even then not good enough due to ad-hoc, unpredictable nature of workload

# IBM Smart Analytics Optimizer

## Adding Industry Leading Performance

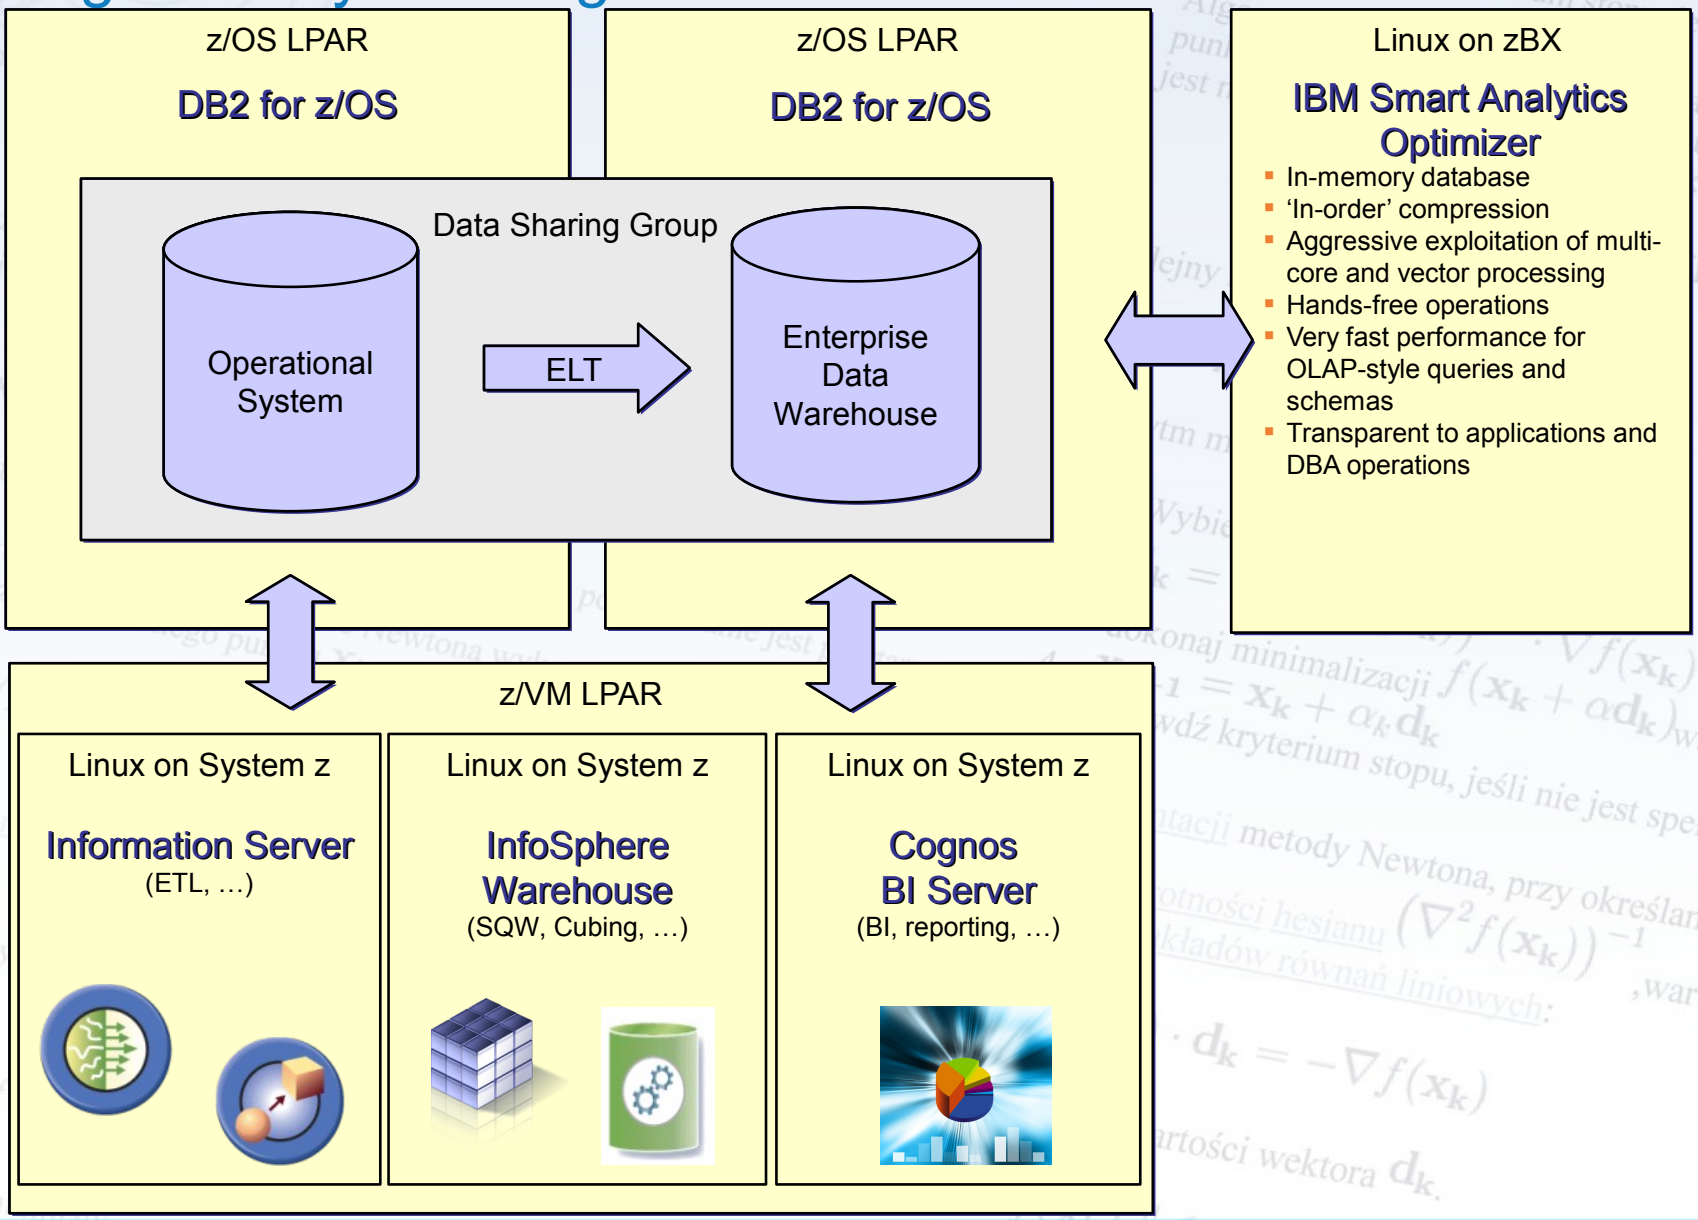

# Agenda

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

# DB2 Architecture

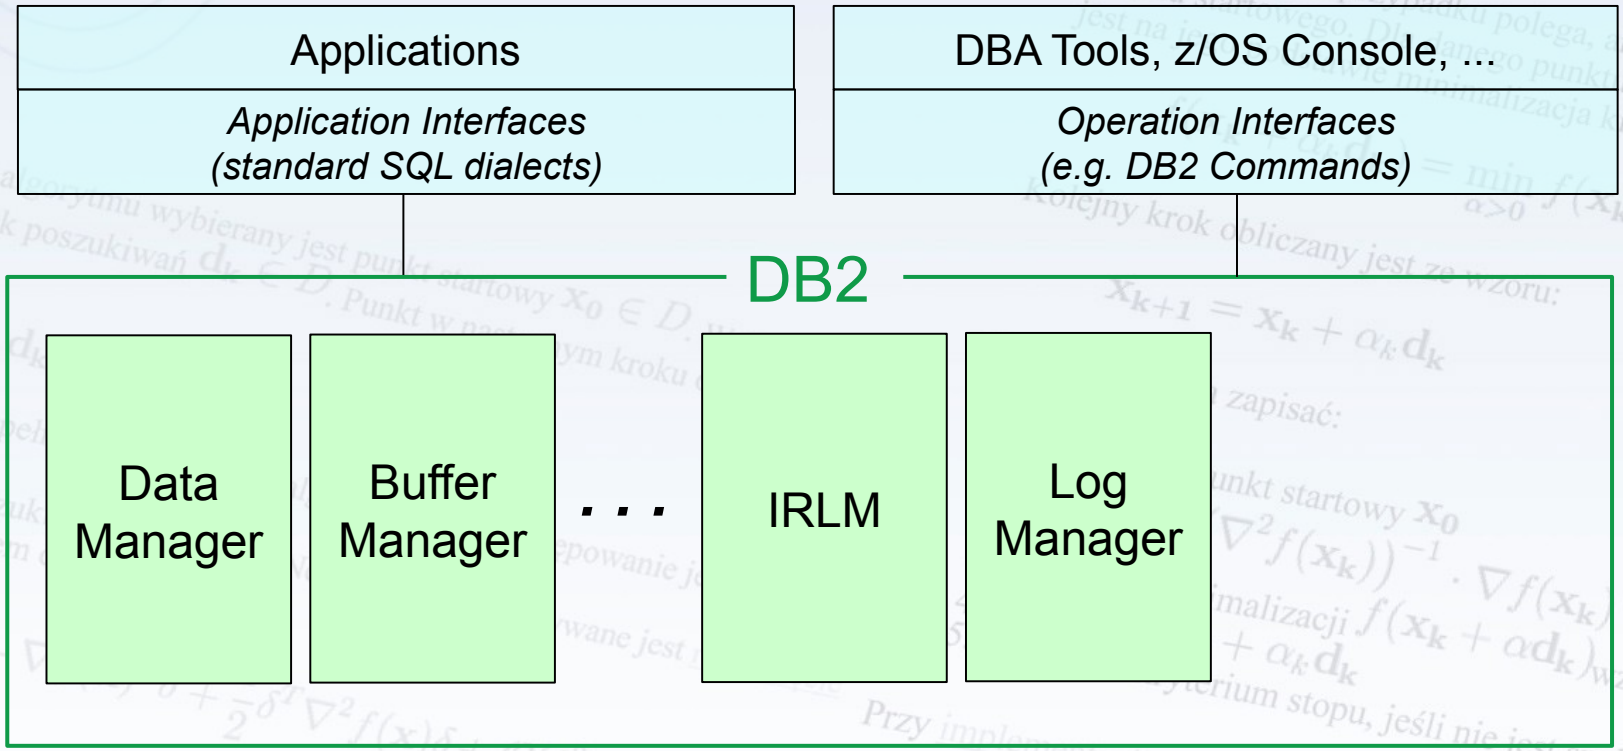

# IBM Smart Analytics Optimizer
## As a Virtual DB2 Component

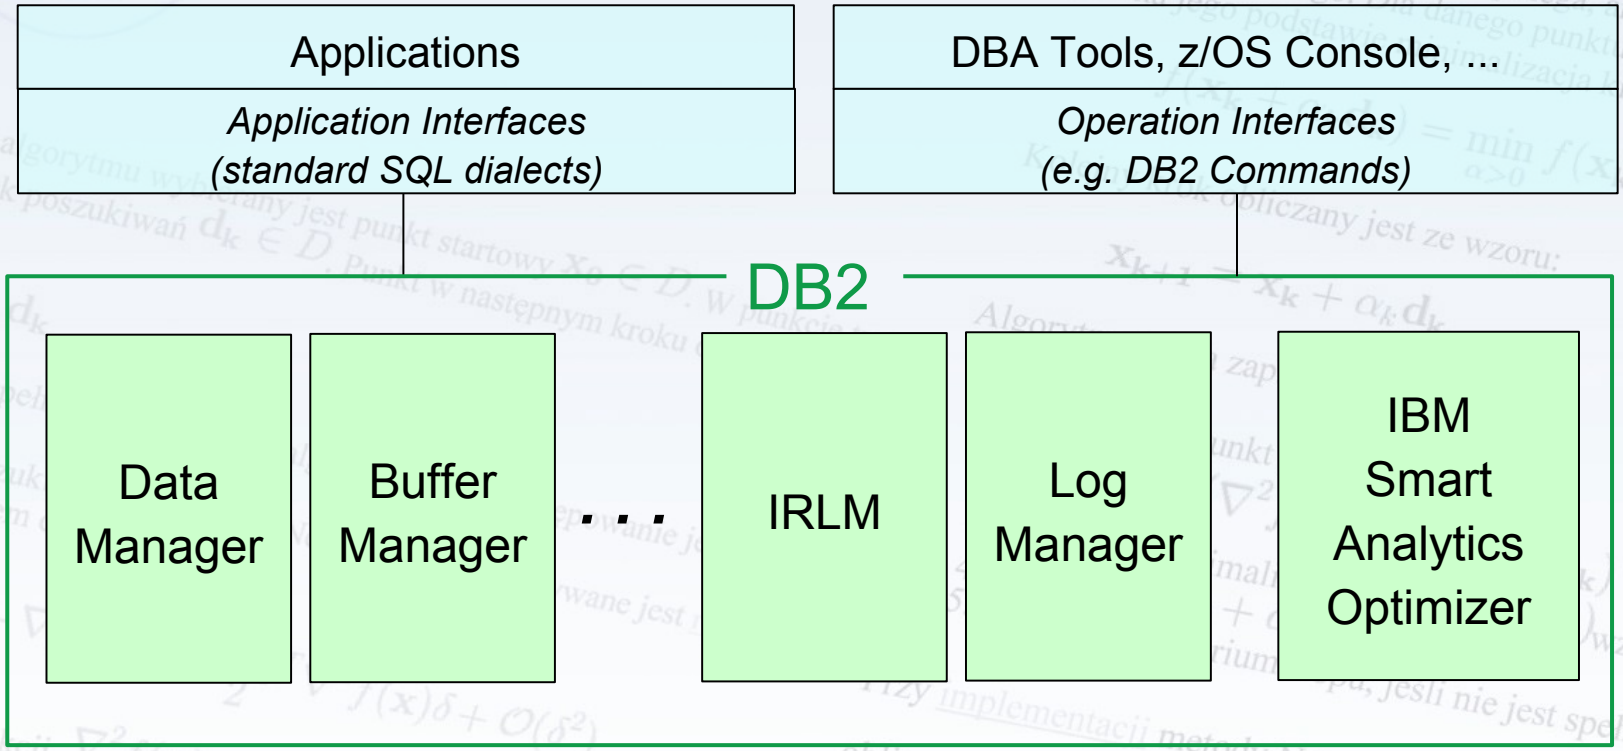

# Deep DB2 Integration within zHybrid Architecture

Applications

*Application Interfaces (standard SQL dialects)*

DBA Tools, z/OS Console, ...

*Operation Interfaces (e.g. DB2 Commands)*

DB2

Data Manager

Buffer Manager

. . .

IRLM

Log Manager

IBM Smart Analytics Optimizer

Superior availability reliability, security, workload management ...

z/OS on System z

10‘s of processors

100‘s GB of memory

Abundance of resources

Linux on zBX

100‘s of processors

1000‘s GB of memory

# 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)**

- of AIX® and Linux applications
- High performance optimizers and appliances to accelerate time to insight and reduce cost
- Dedicated high performance private network

[1] All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

# Enabling Technology
## IBM Research Project BLINK

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

# OLTP vs. DW

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

# Row-Store – Optimal Choice for OLTP

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

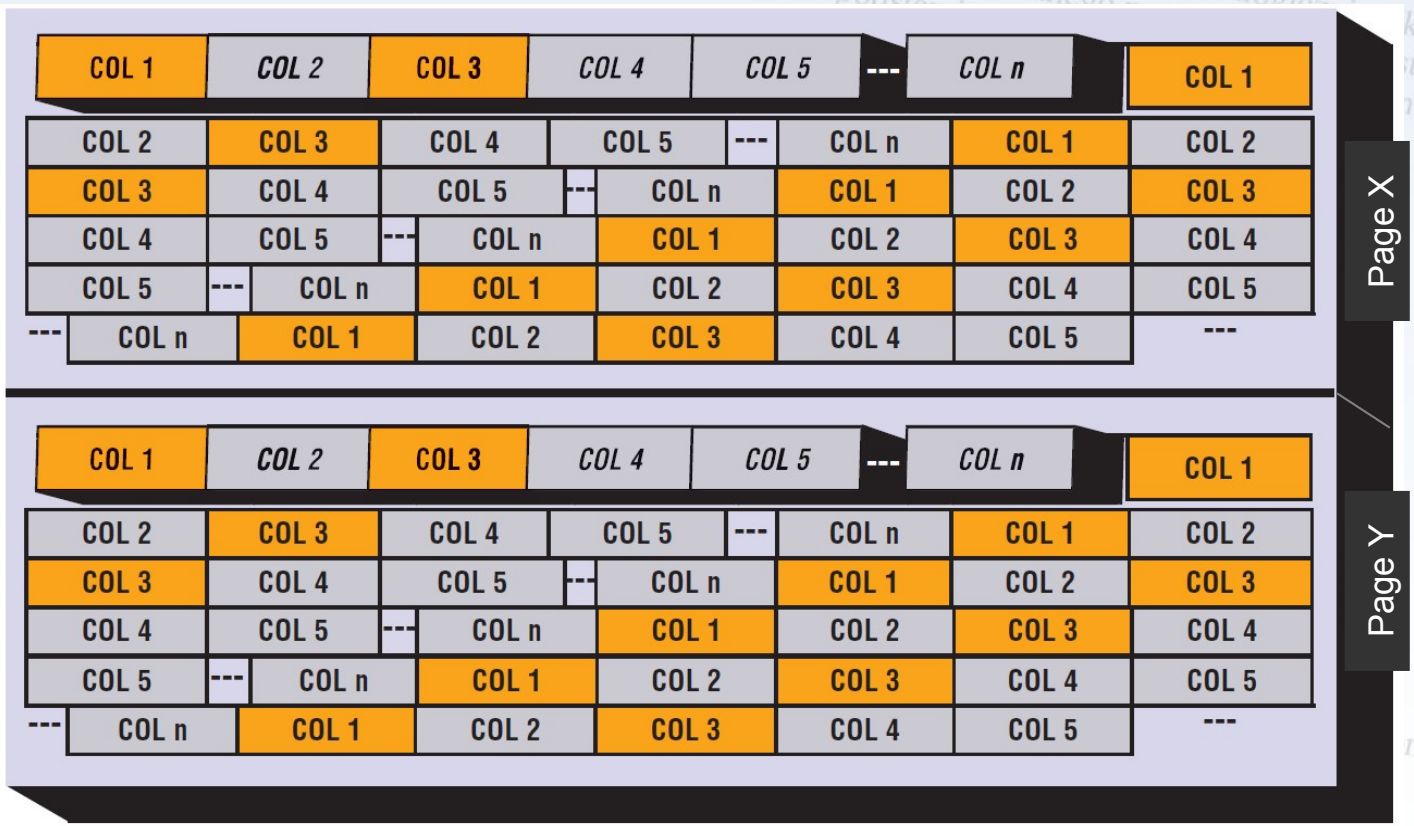


While a **Row – Store** is very efficient for transactional workloads, it is suboptimal for analytical workloads where only a subset of the attributes is needed!

# Column-Store: Optimized for DW Workloads

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

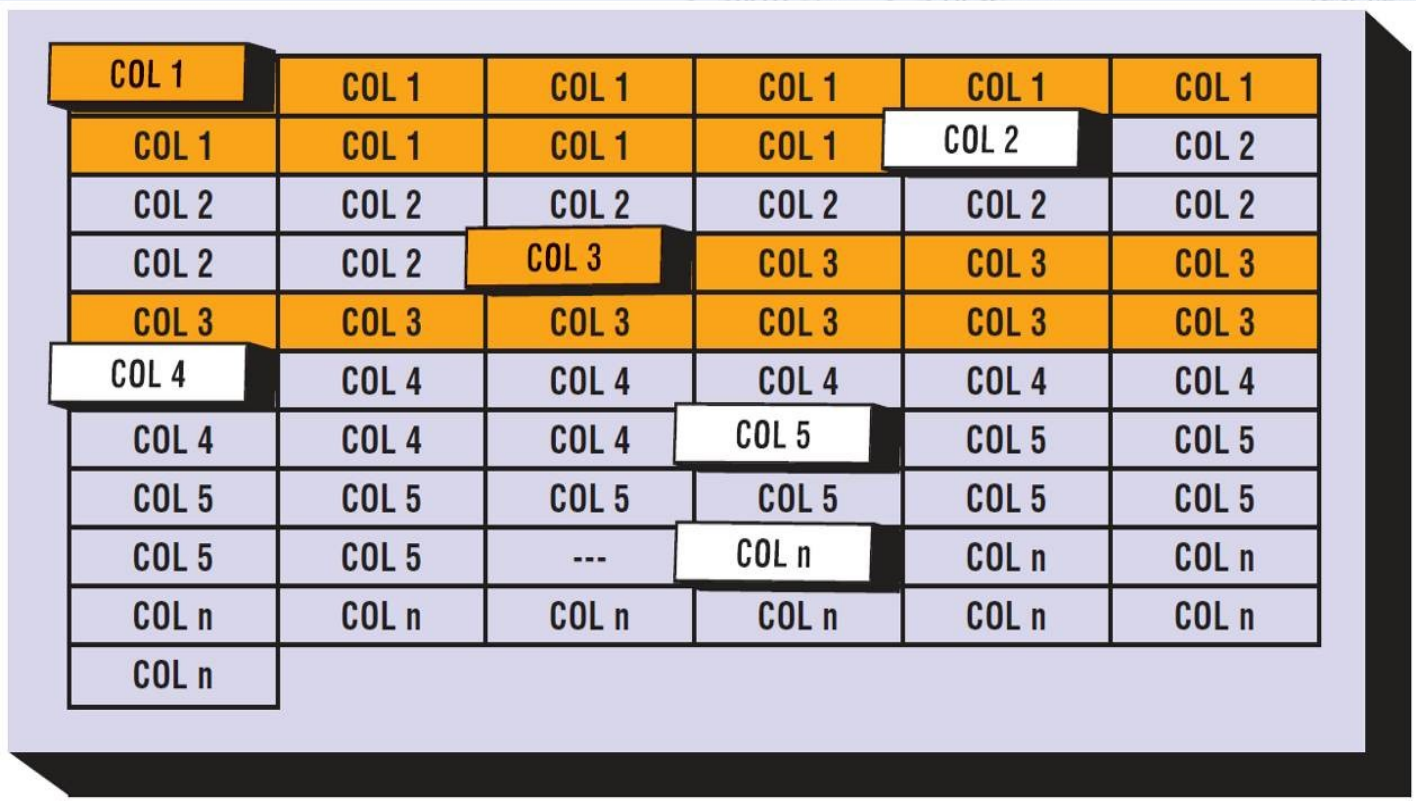


In a **Column – Store**, the data is also compressed sequentially for a column. This is an optimized approach if you plan to perform a sequential scan over your data. Random access to specific attributes in this store is not performing well.

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

# Register-Store

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

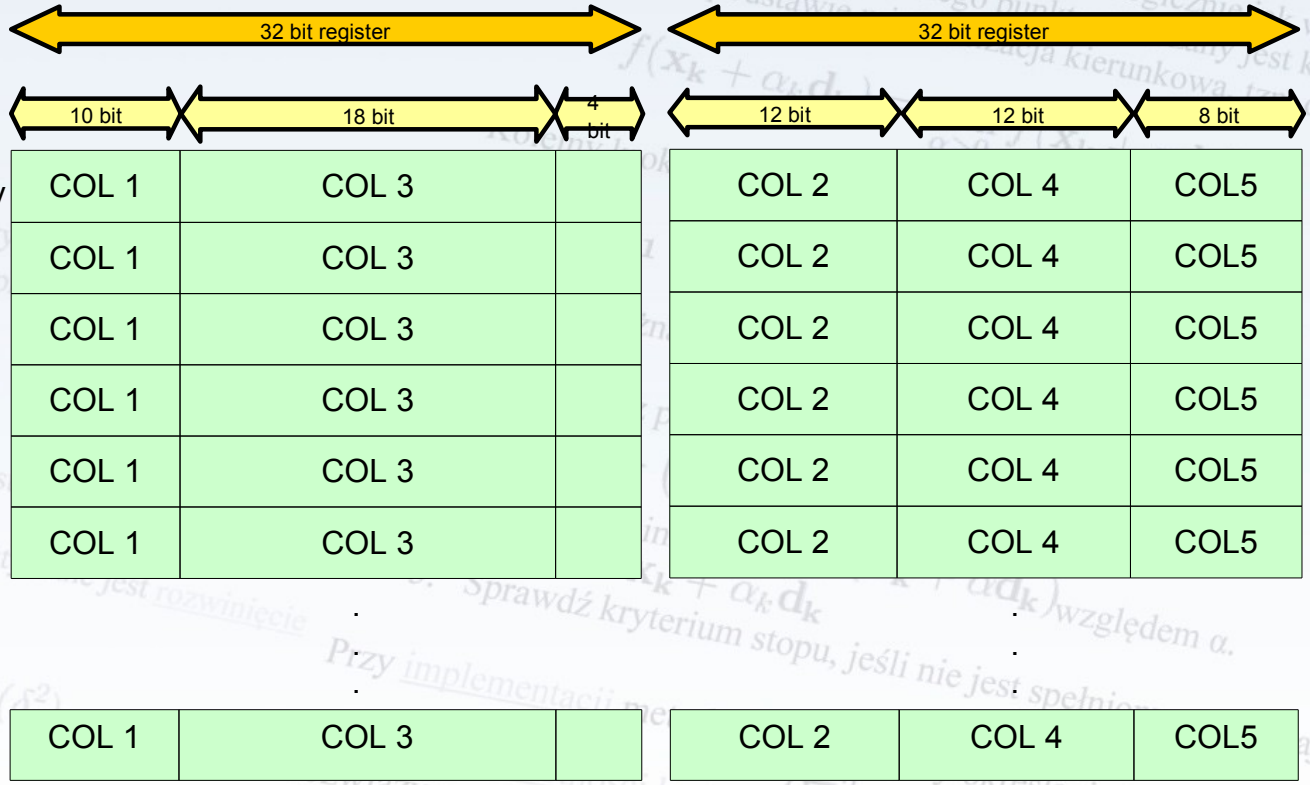


The **Register – Store** is an optimization of the Column – Store approach where we try to make the best use of existing hardware. Reshuffling small data elements at runtime into a register is time consuming and can be avoided. The **Register – Store** also delivers good vectorization capabilities.

# Single Instruction Multiple Data Paradigm

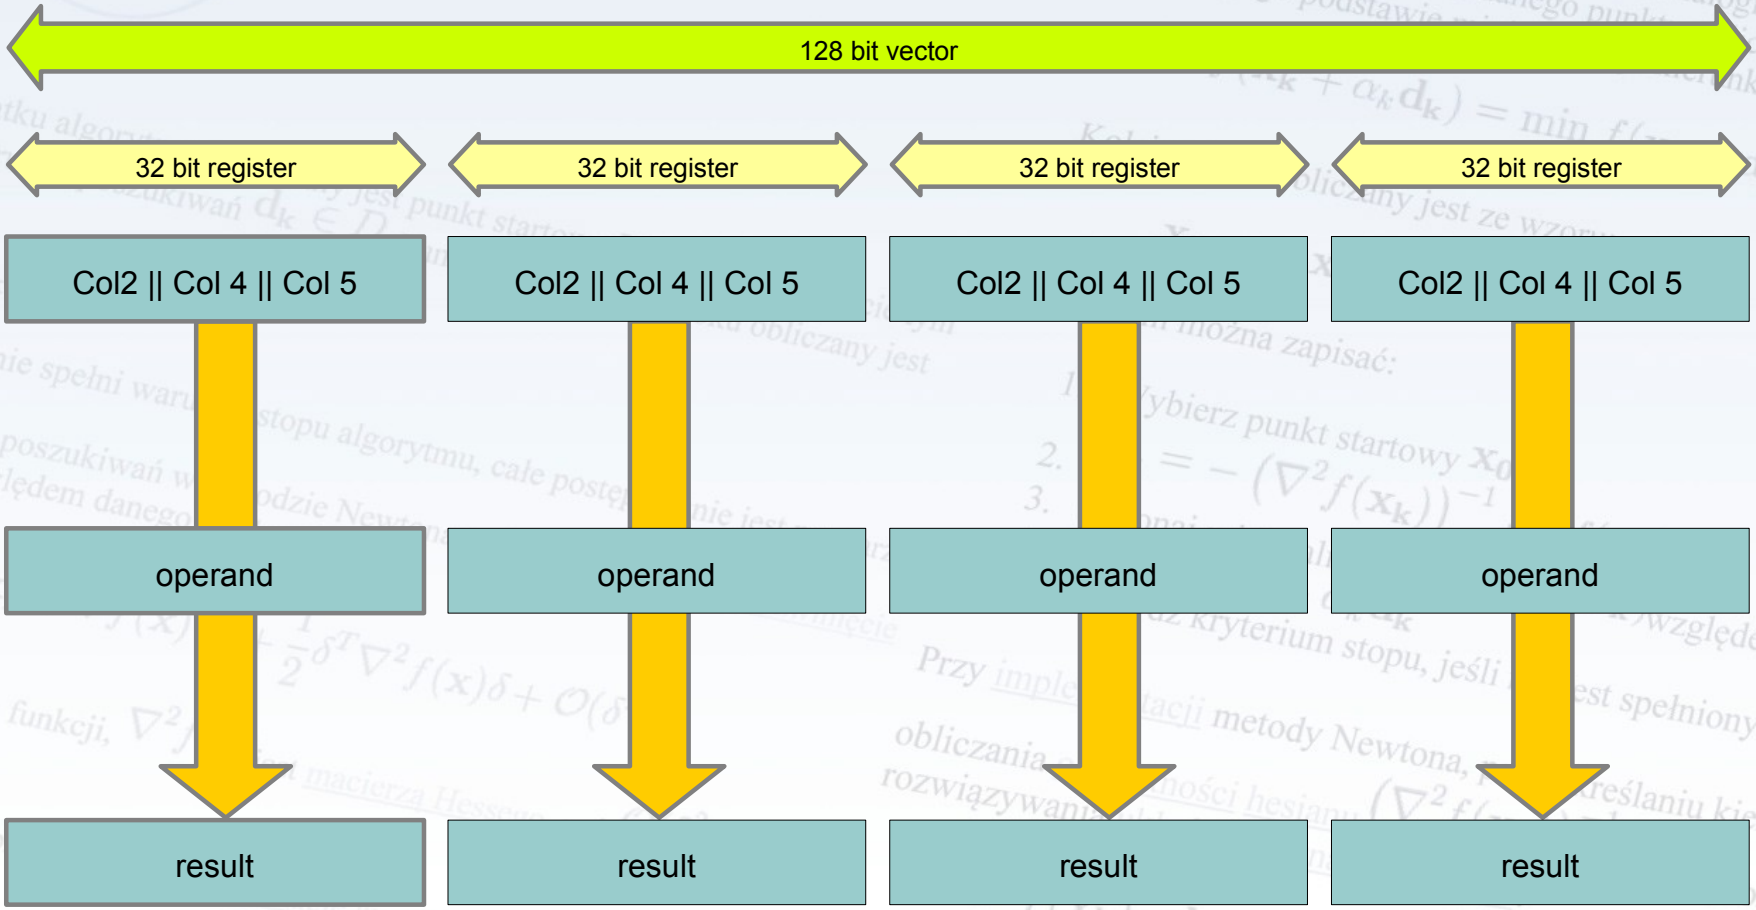

# Data View

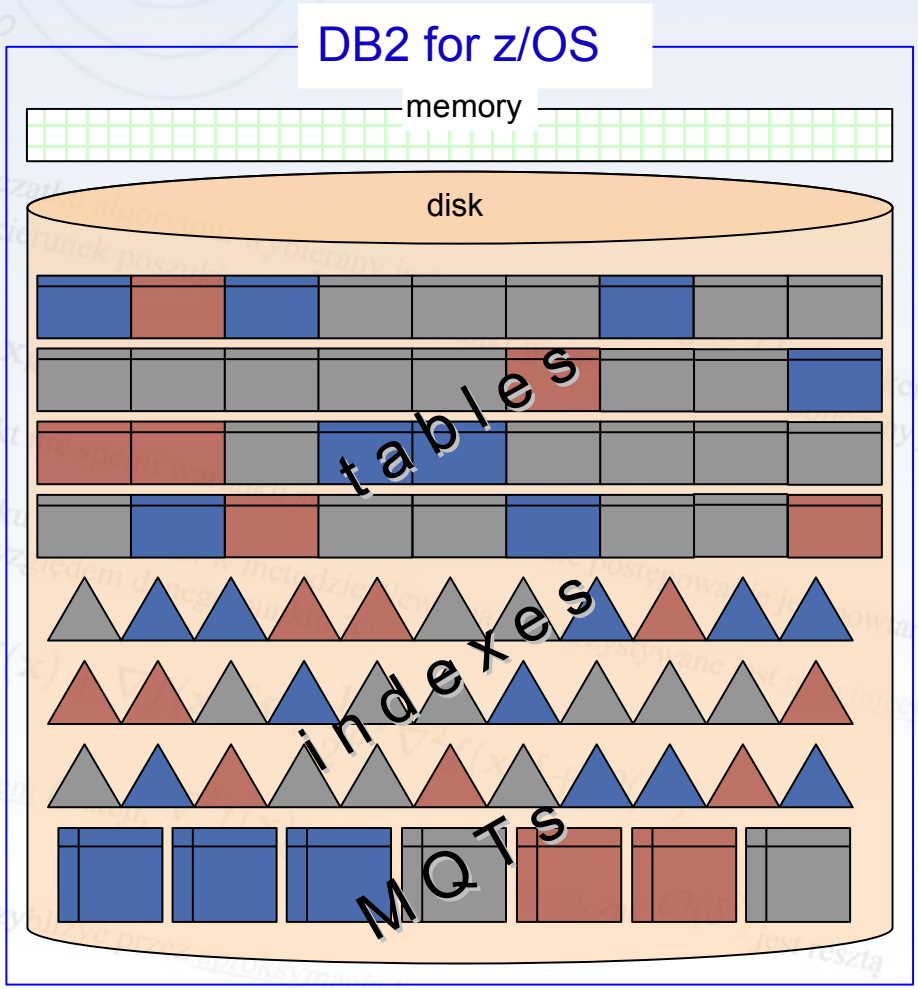


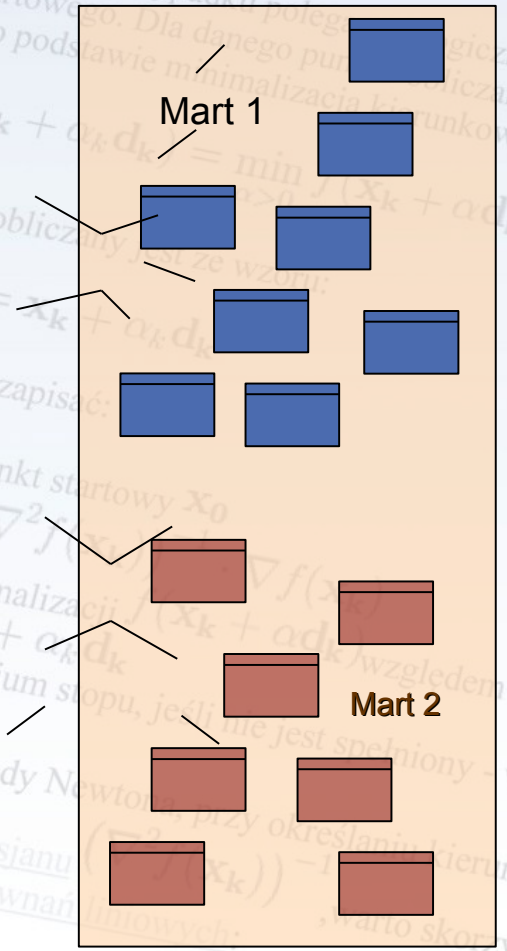

# **Marts:** Redundant Sets of Memory Resident Tables

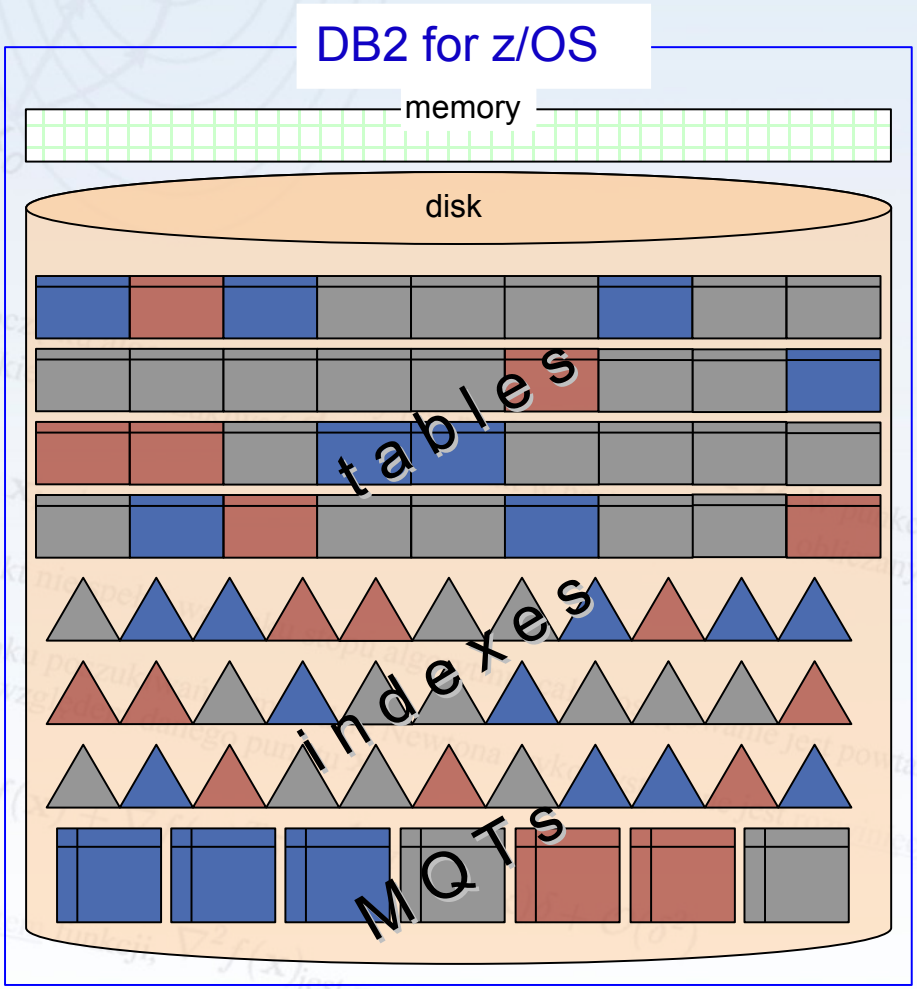


- DB2 continues to *own* and manage all data
- Access performance is influenced by traditional tuning mechanisms such as indexing, MQTs, aggregates, ...
- Typical usage: Enterprise Data Warehouse, large Data Marts

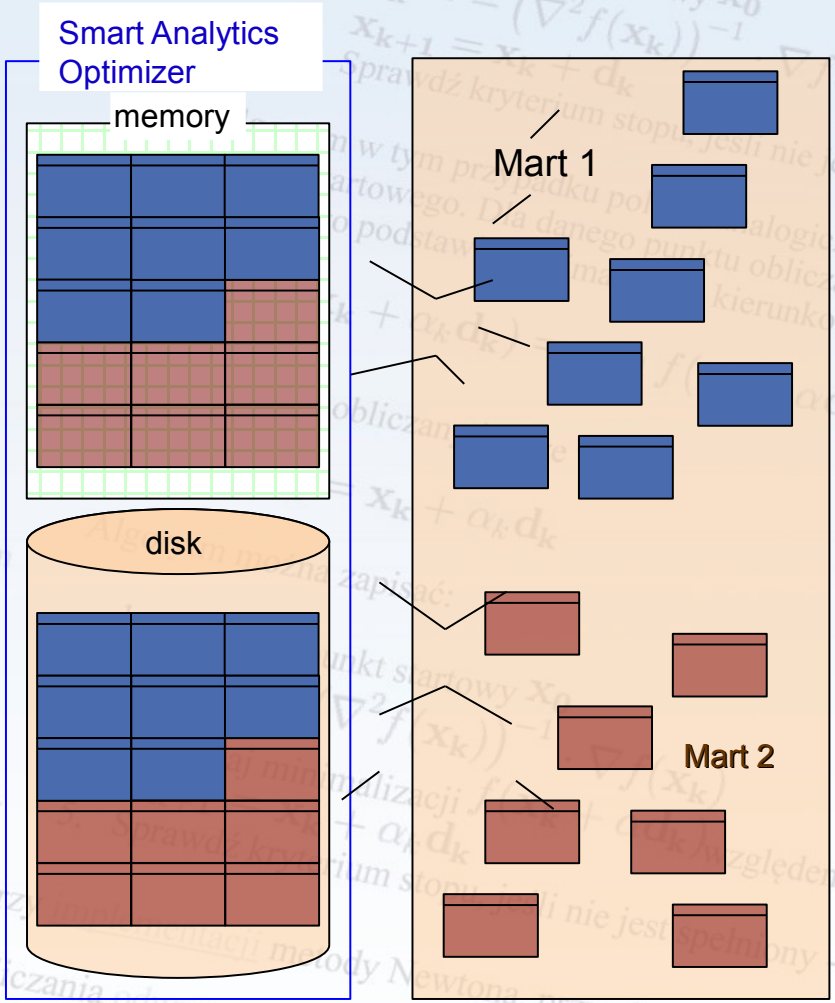


- Smart Analytics Optimizer contains fully memory resident, compressed copies of performance critical tables grouped into logically connected *marts*, bound ideally by star schema constraints
- Similar but much broader than MQTs: no column projections, no row restrictions, no row aggregations
- Typical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

# Smart Analytics Optimizer Mart Definition and Deployment

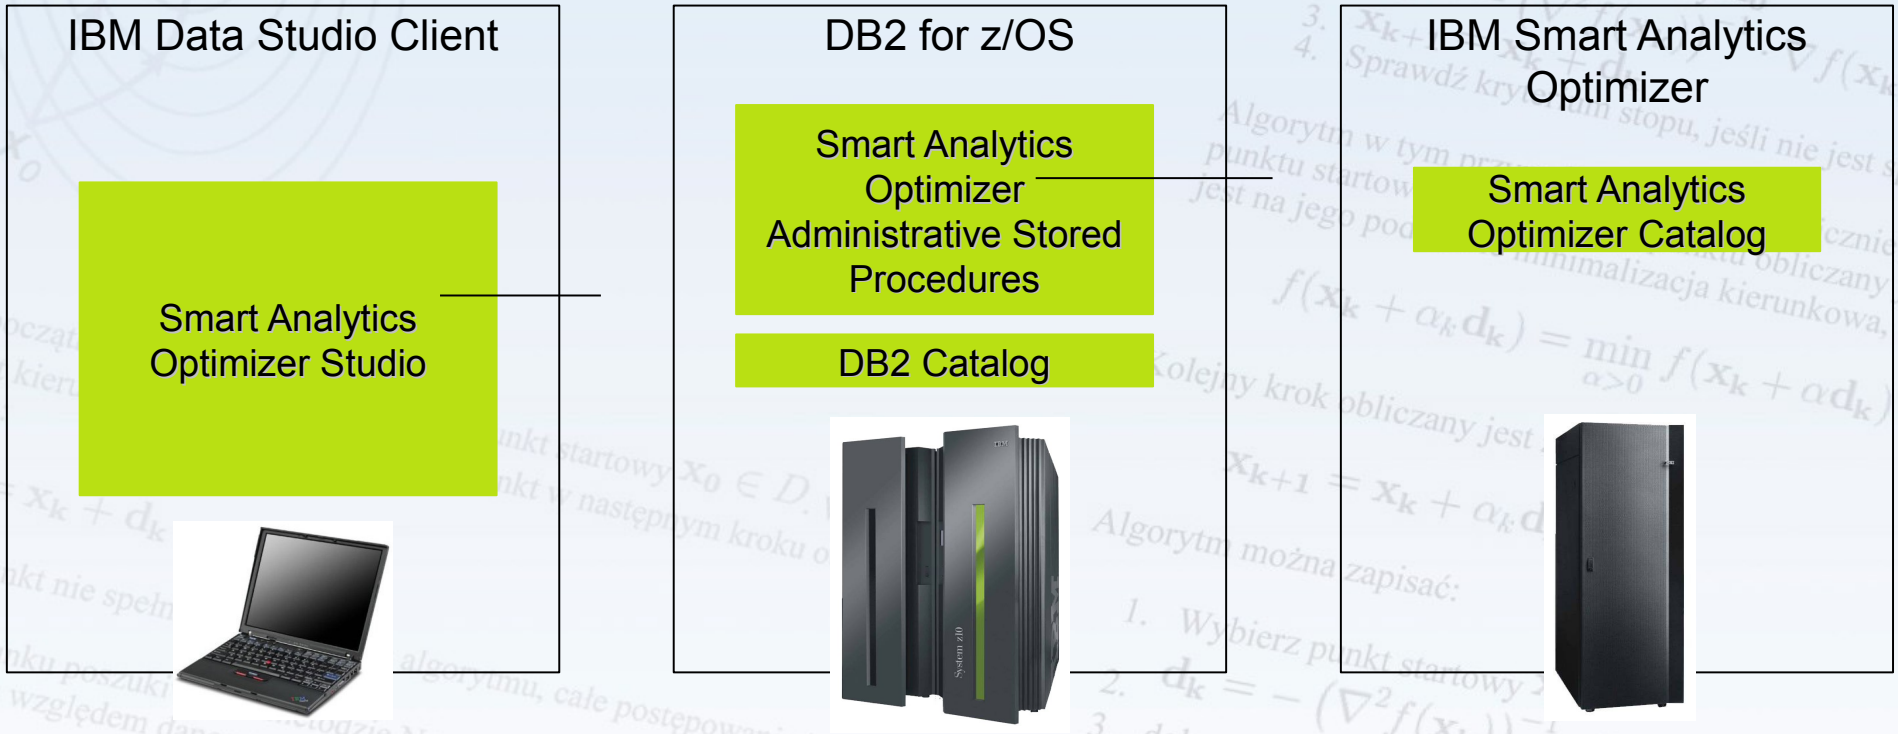


- Smart Analytics Optimizer marts need to be defined and deployed to Smart Analytics Optimizer before data is loaded and queries sent to Smart Analytics Optimizer for processing.
  - Definition: identifying tables and relations that make up marts.
  - Deployment: making marts known to DB2, i.e. storing 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.

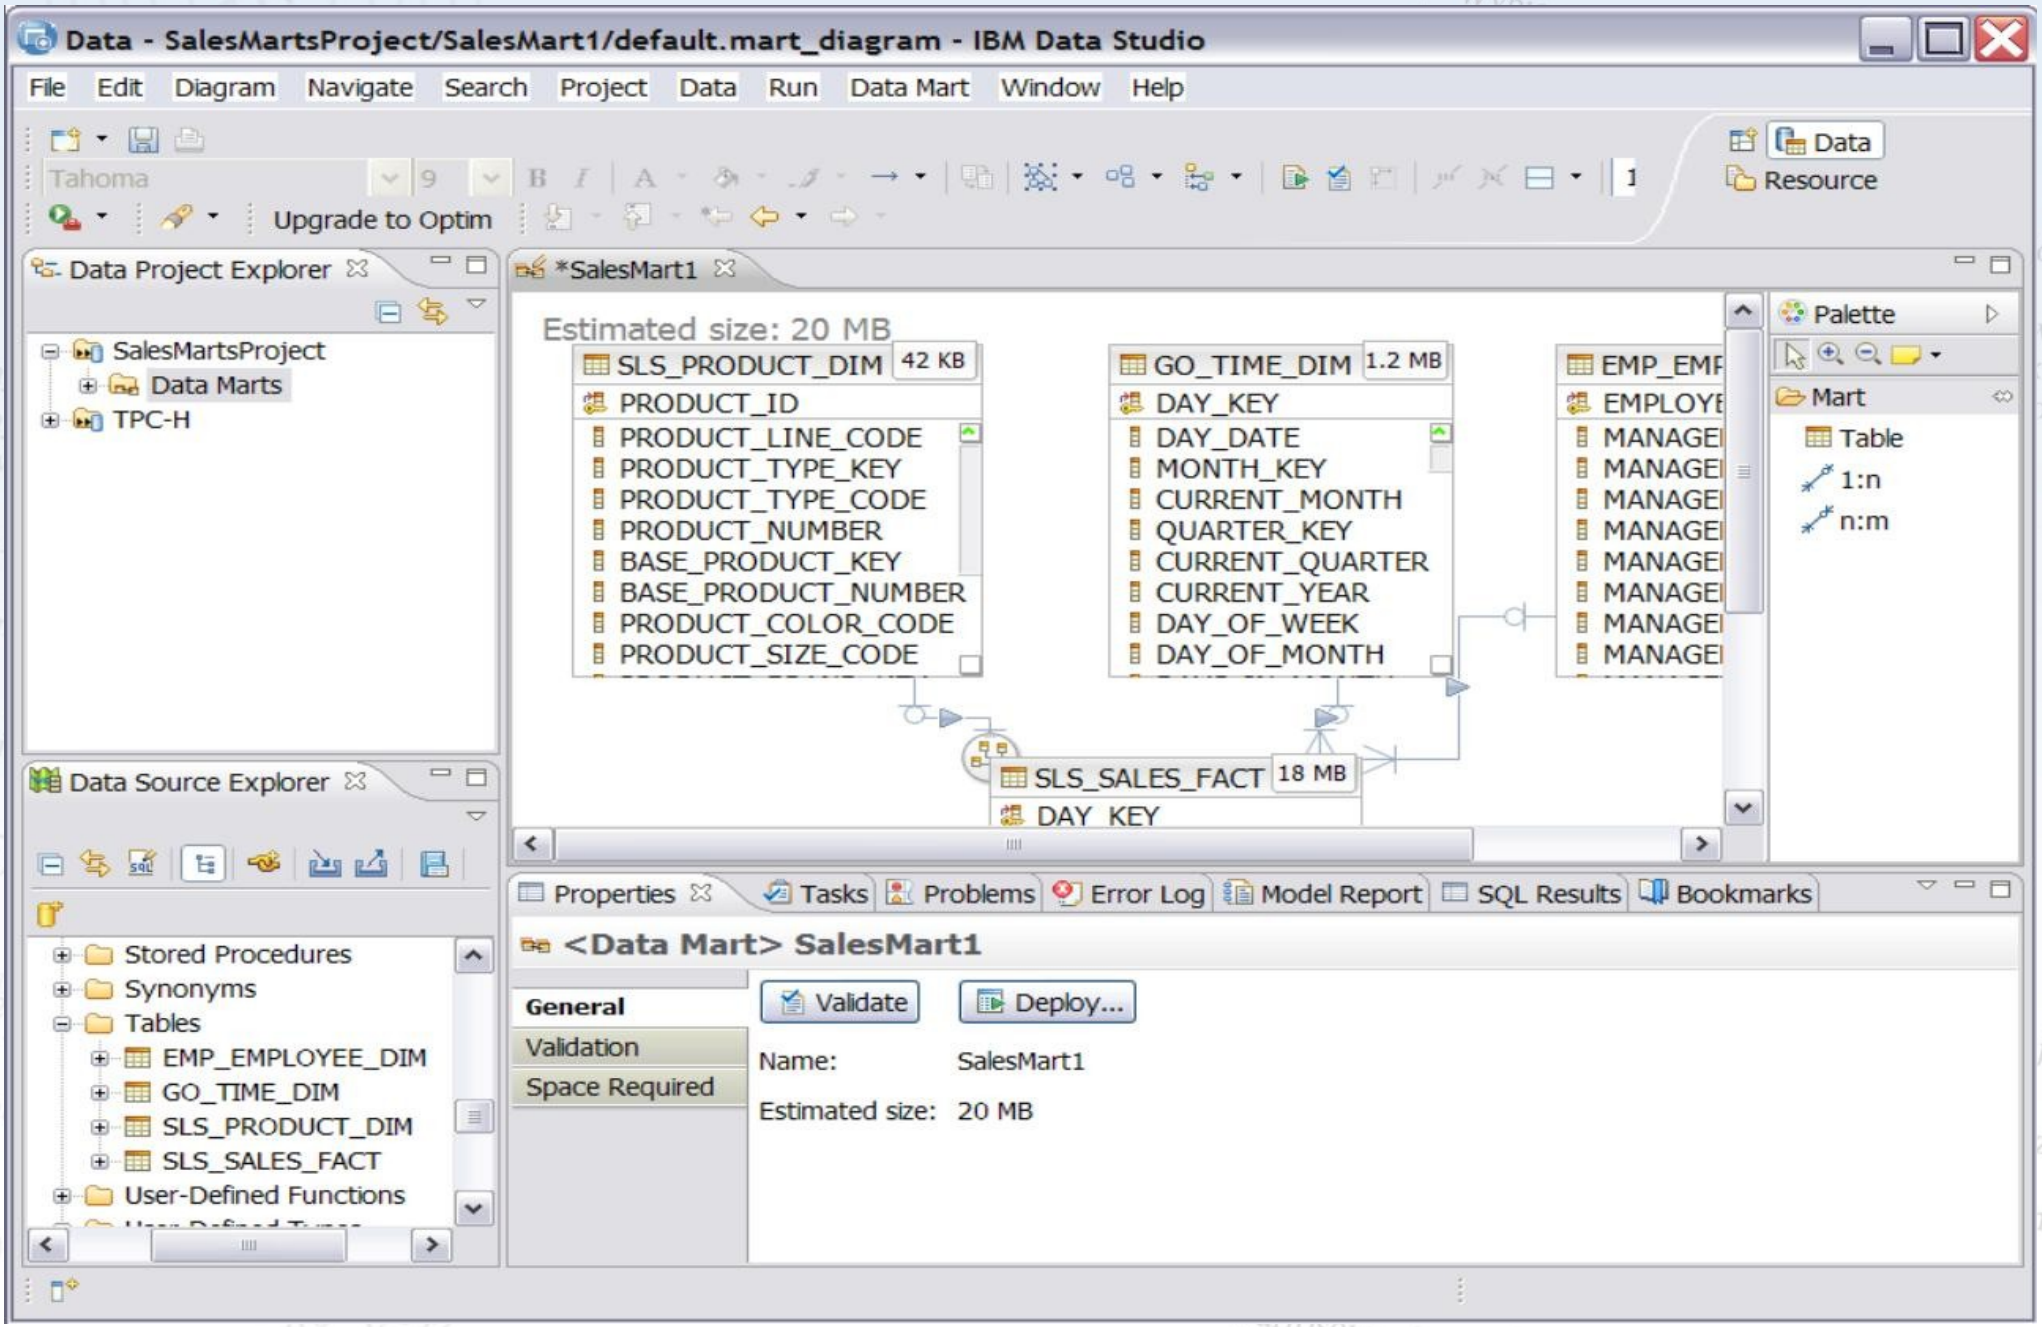
Data - SalesMartsProject/SalesMart1/default.mart_diagram - IBM Data Studio
File Edit Diagram Navigate Search Project Data Run Data Mart Window Help
Upgrade to Optim
Data
Resource
Data Project Explorer
SalesMartsProject
Data Marts
TPC-H
Data Source Explorer
Stored Procedures
Synonyms
Tables
EMP_EMPLOYEE_DIM
GO_TIME_DIM
SLS_PRODUCT_DIM
SLS_SALES_FACT
User-Defined Functions
*SalesMart1
Estimated size: 20 MB
SLS_PRODUCT_DIM 42 KB
PRODUCT_ID
PRODUCT_LINE_CODE
PRODUCT_TYPE_KEY
PRODUCT_TYPE_CODE
PRODUCT_NUMBER
BASE_PRODUCT_KEY
BASE_PRODUCT_NUMBER
PRODUCT_COLOR_CODE
PRODUCT_SIZE_CODE
GO_TIME_DIM 1.2 MB
DAY_KEY
DAY_DATE
MONTH_KEY
CURRENT_MONTH
QUARTER_KEY
CURRENT_QUARTER
CURRENT_YEAR
DAY_OF_WEEK
DAY_OF_MONTH
SLS_SALES_FACT 18 MB
DAY_KEY
Palette
Mart
Table
1:n
n:m
Properties
Tasks
Problems
Error Log
Model Report
SQL Results
Bookmarks
<Data Mart> SalesMart1
General
Validation
Space Required
Validate
Deploy...
Name: SalesMart1
Estimated size: 20 MB

# Smart Analytics Optimizer Mart Load

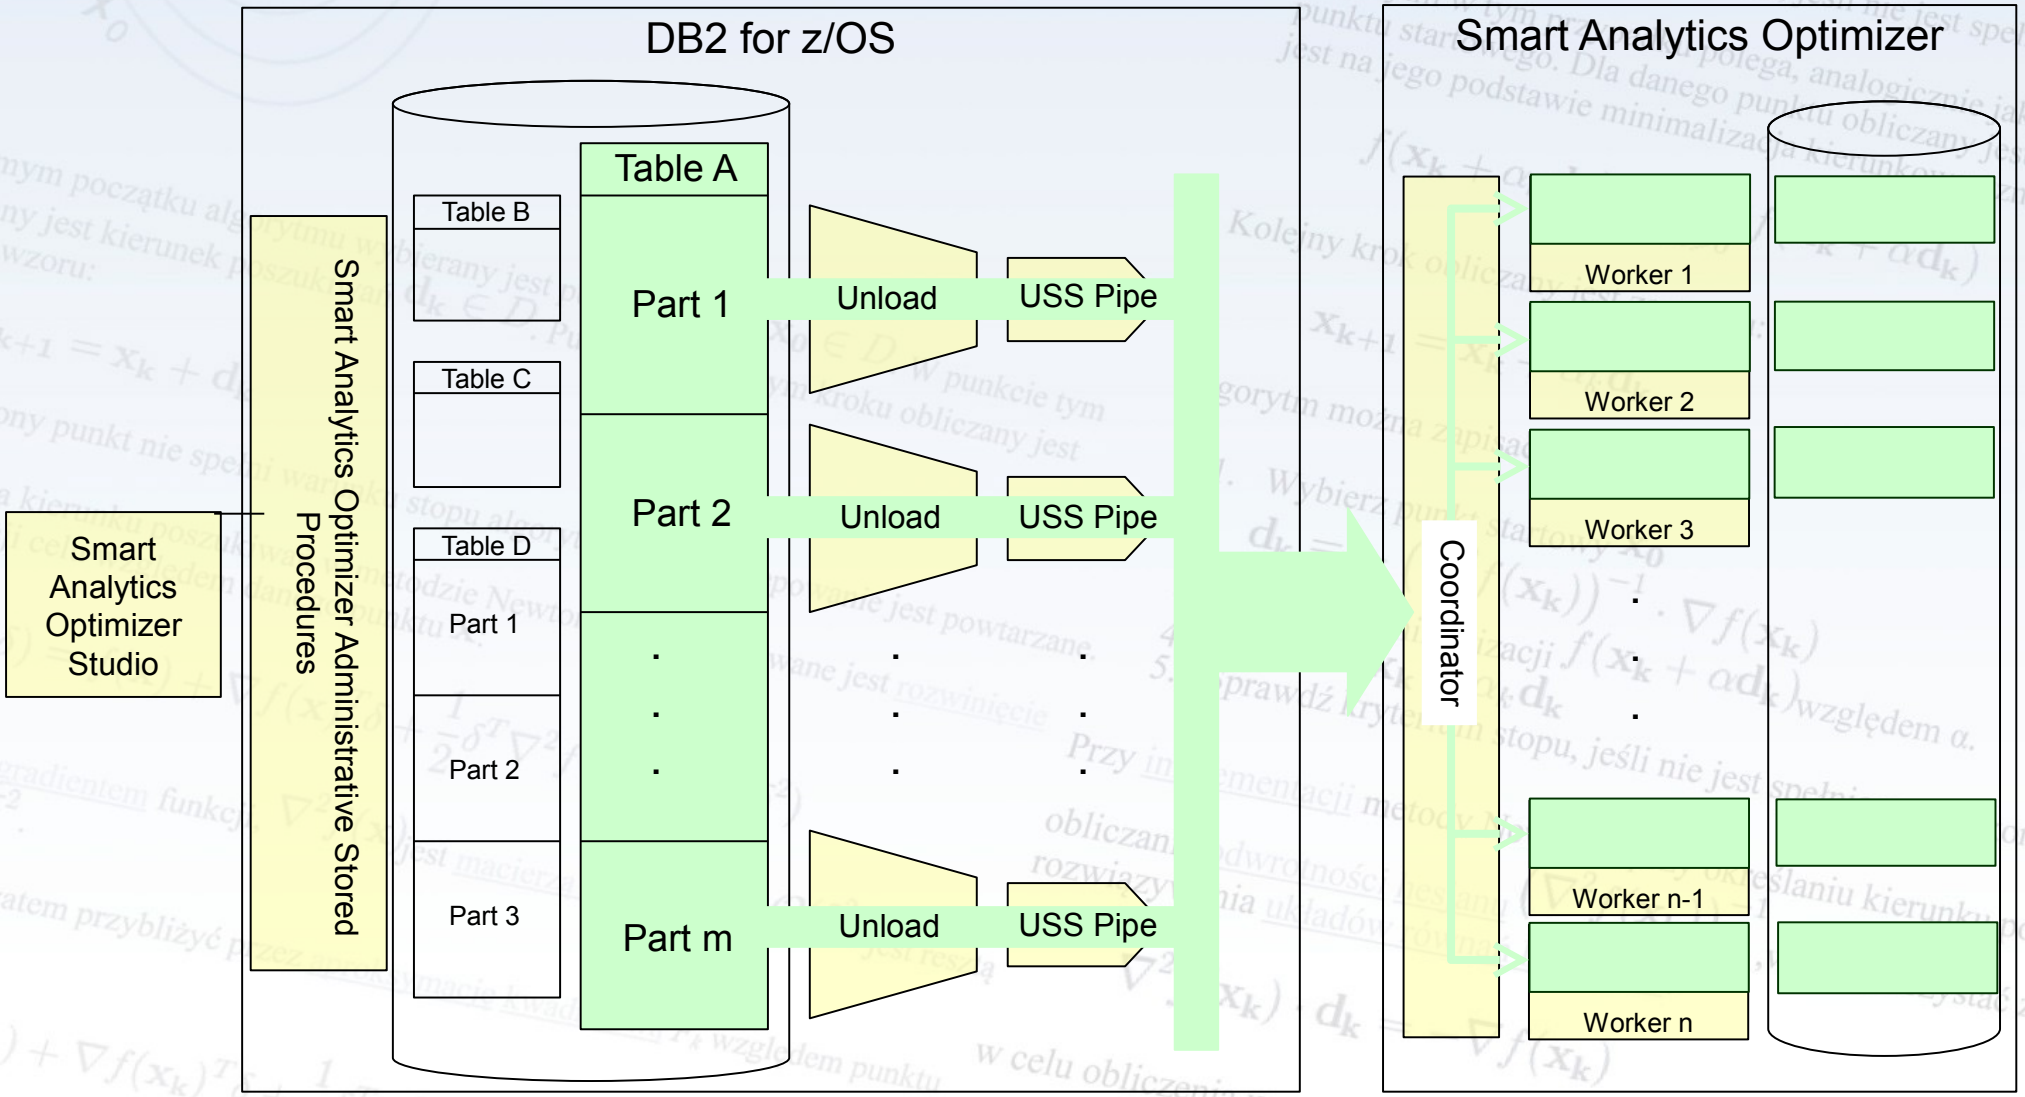

# Smart Analytics Optimizer Mart Update

Typical DW update operations:

- LOAD RESUME and REPLACE
- ADD and ROTATE PARTITION
- SQL INSERT, UPDATE, DELETE
- Delete complete partition or table
- TRUNCATE TABLE

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

IDUG study on DW update frequency

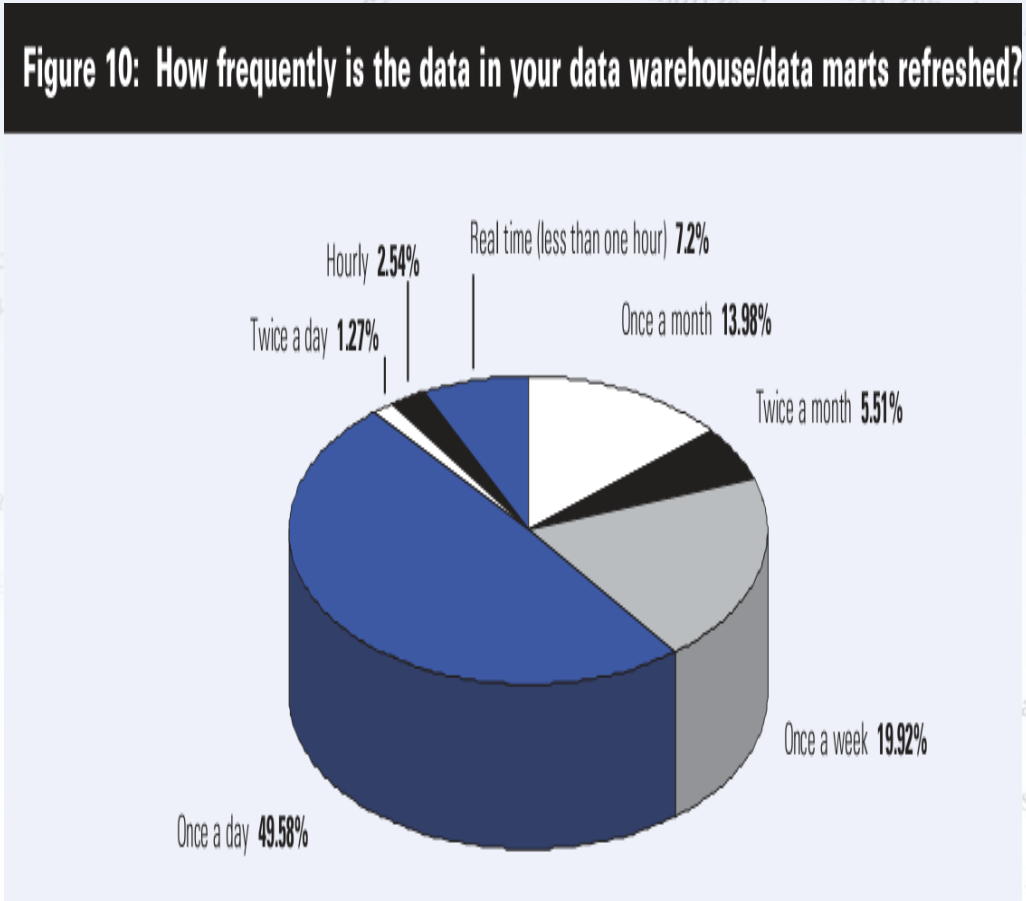


In 90% of cases DW is updated once a day or even less frequently

# Query Execution Process Flow

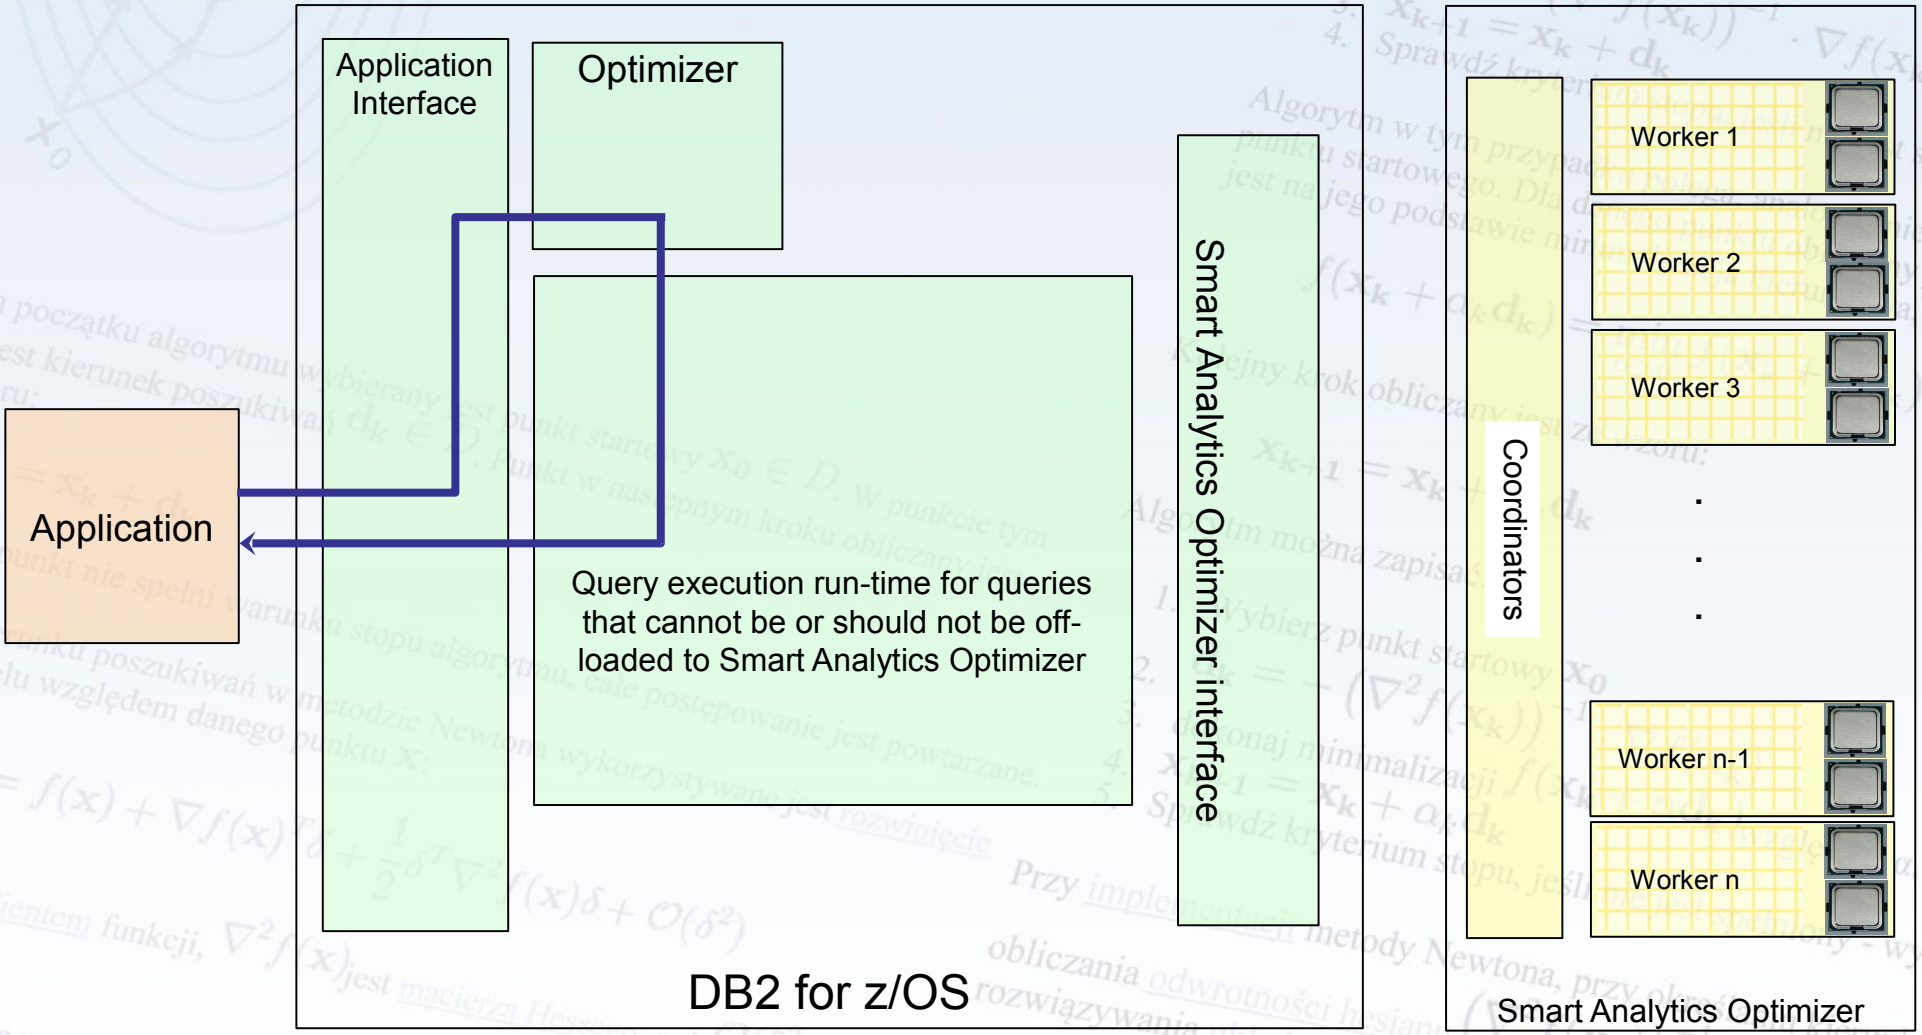

# Query Execution Process Flow

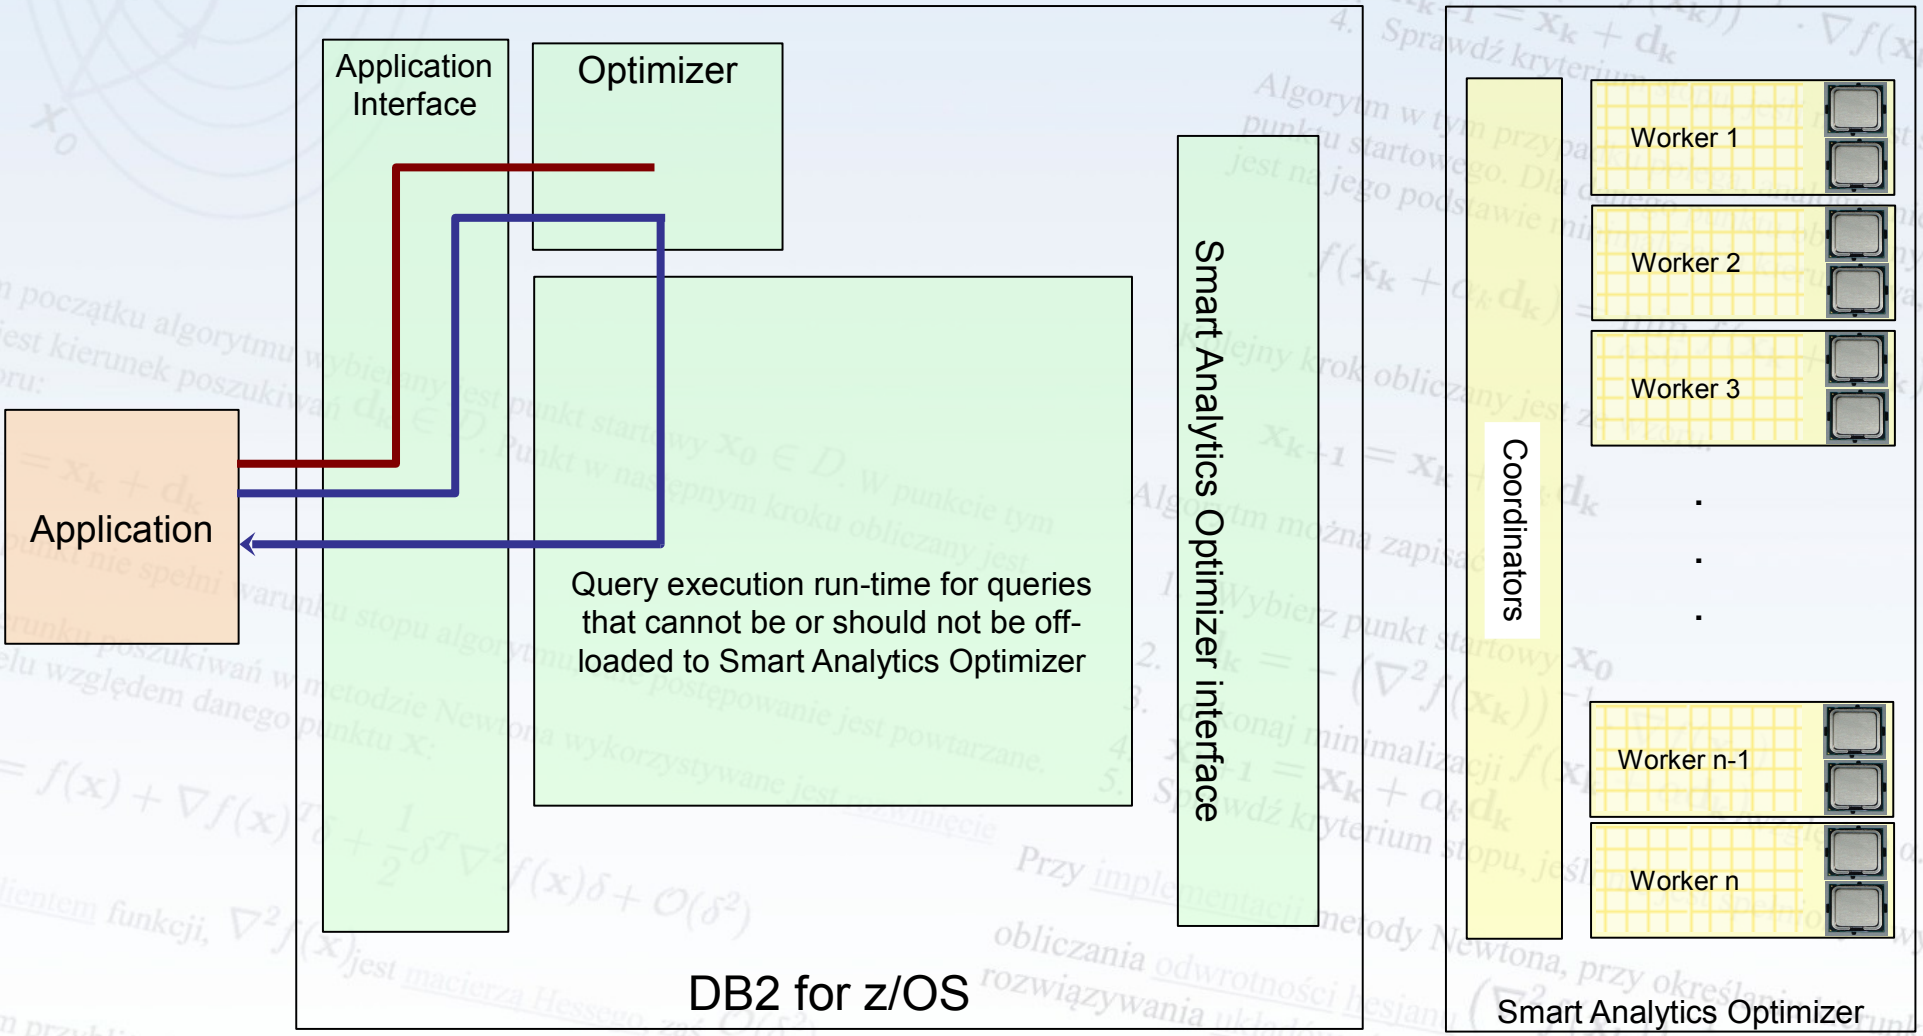


Queries executed without Smart Analytics Optimizer

Queries executed with Smart Analytics Optimizer

# Query Execution Process Flow

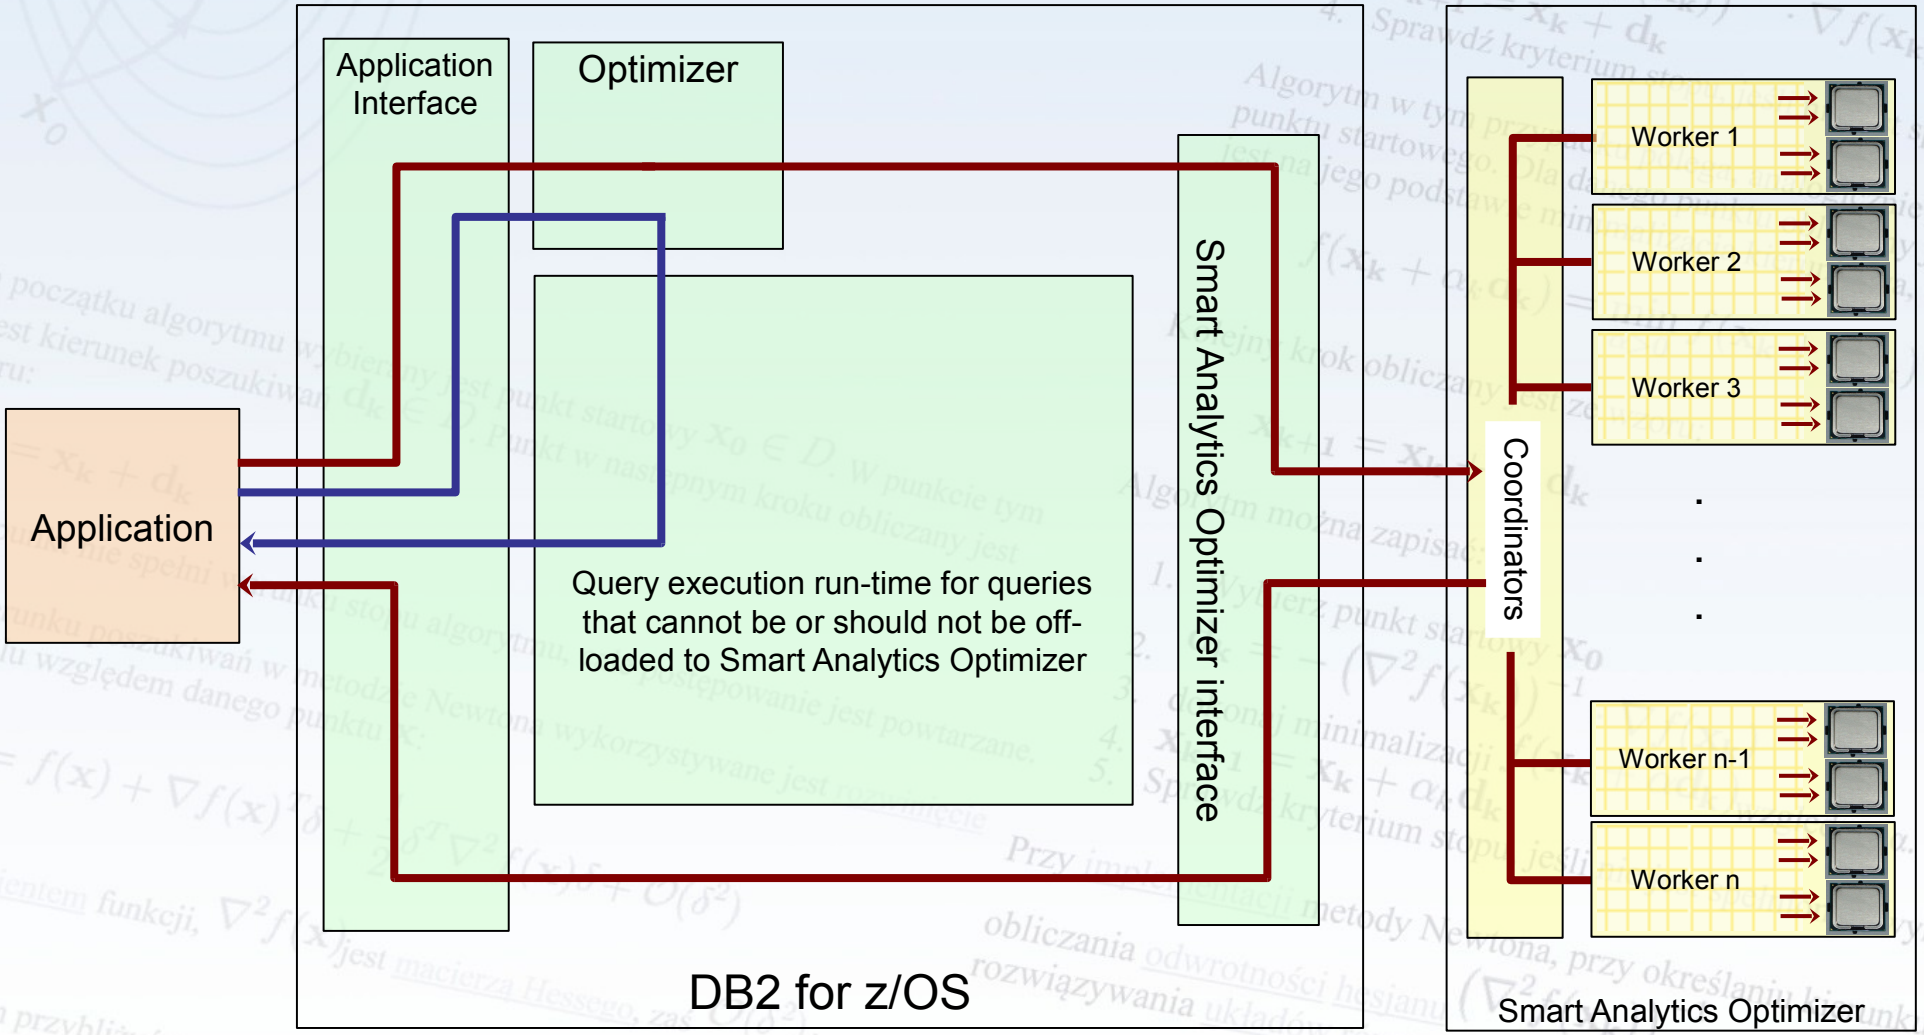

# Query Execution Process Flow

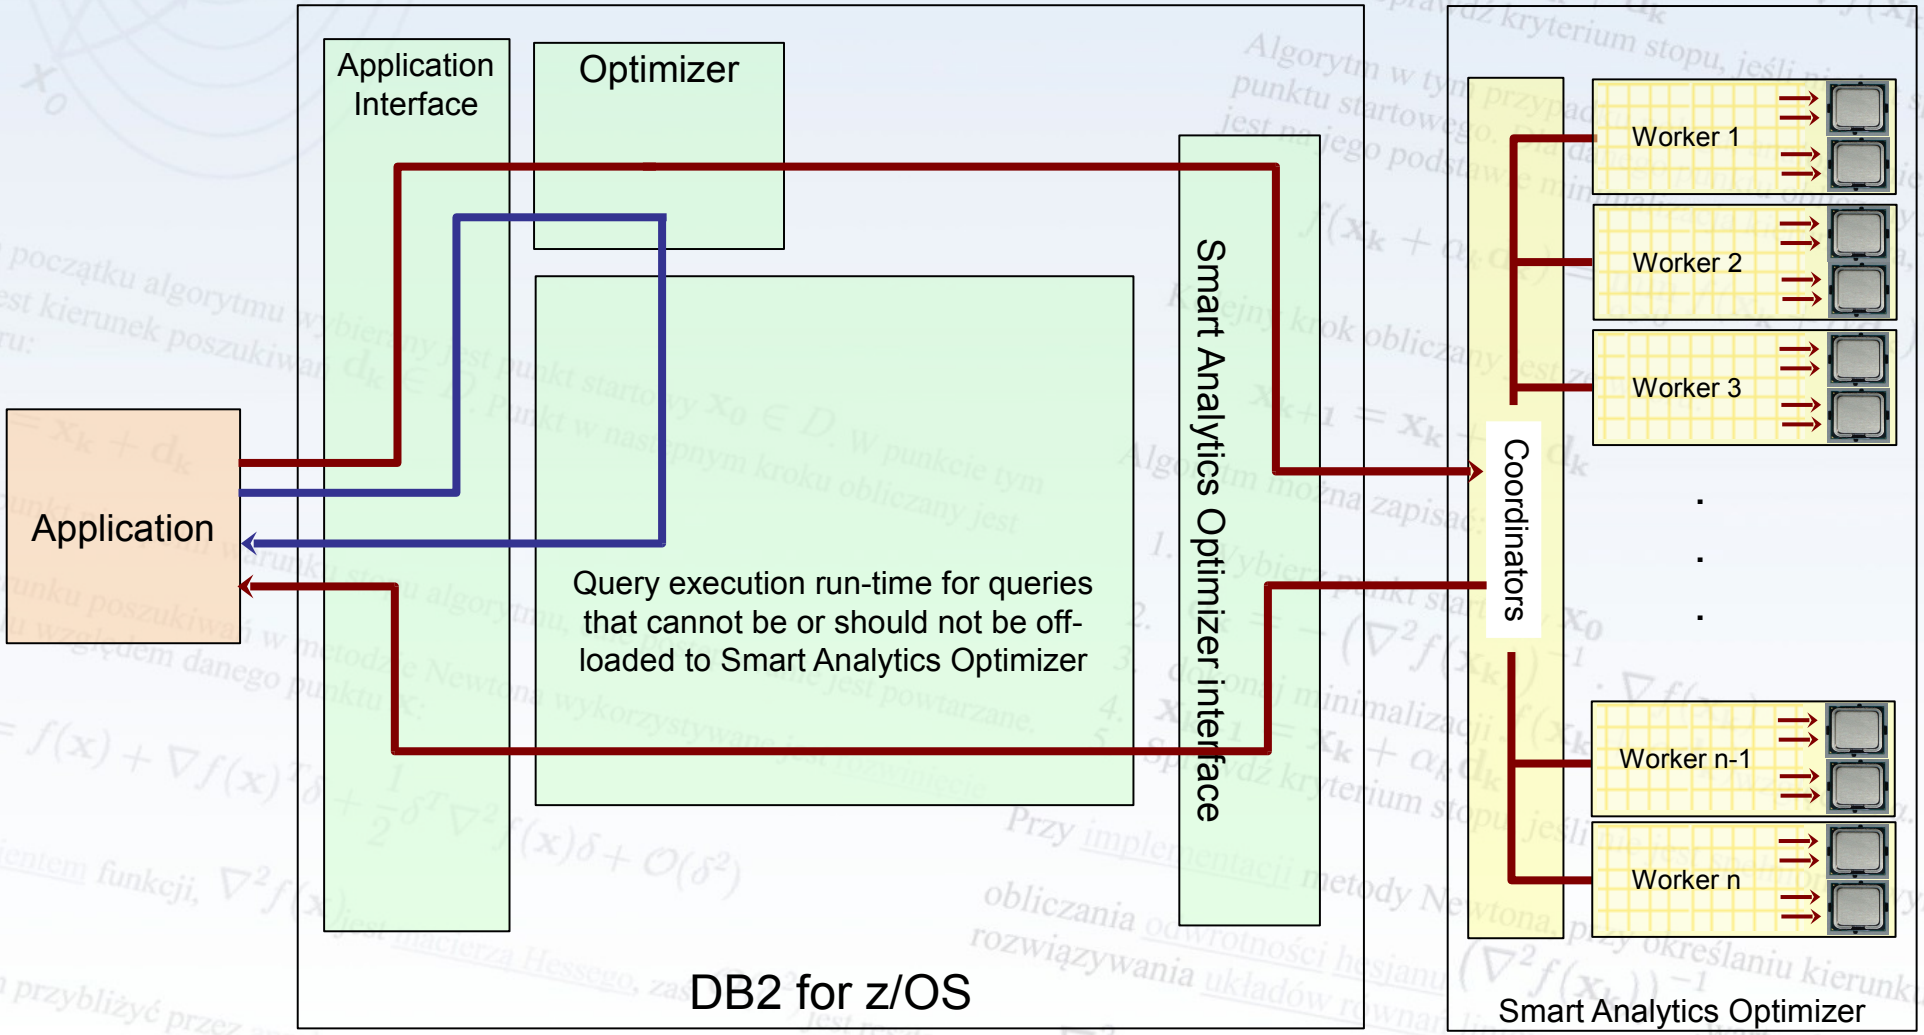

# Query Execution Process Flow

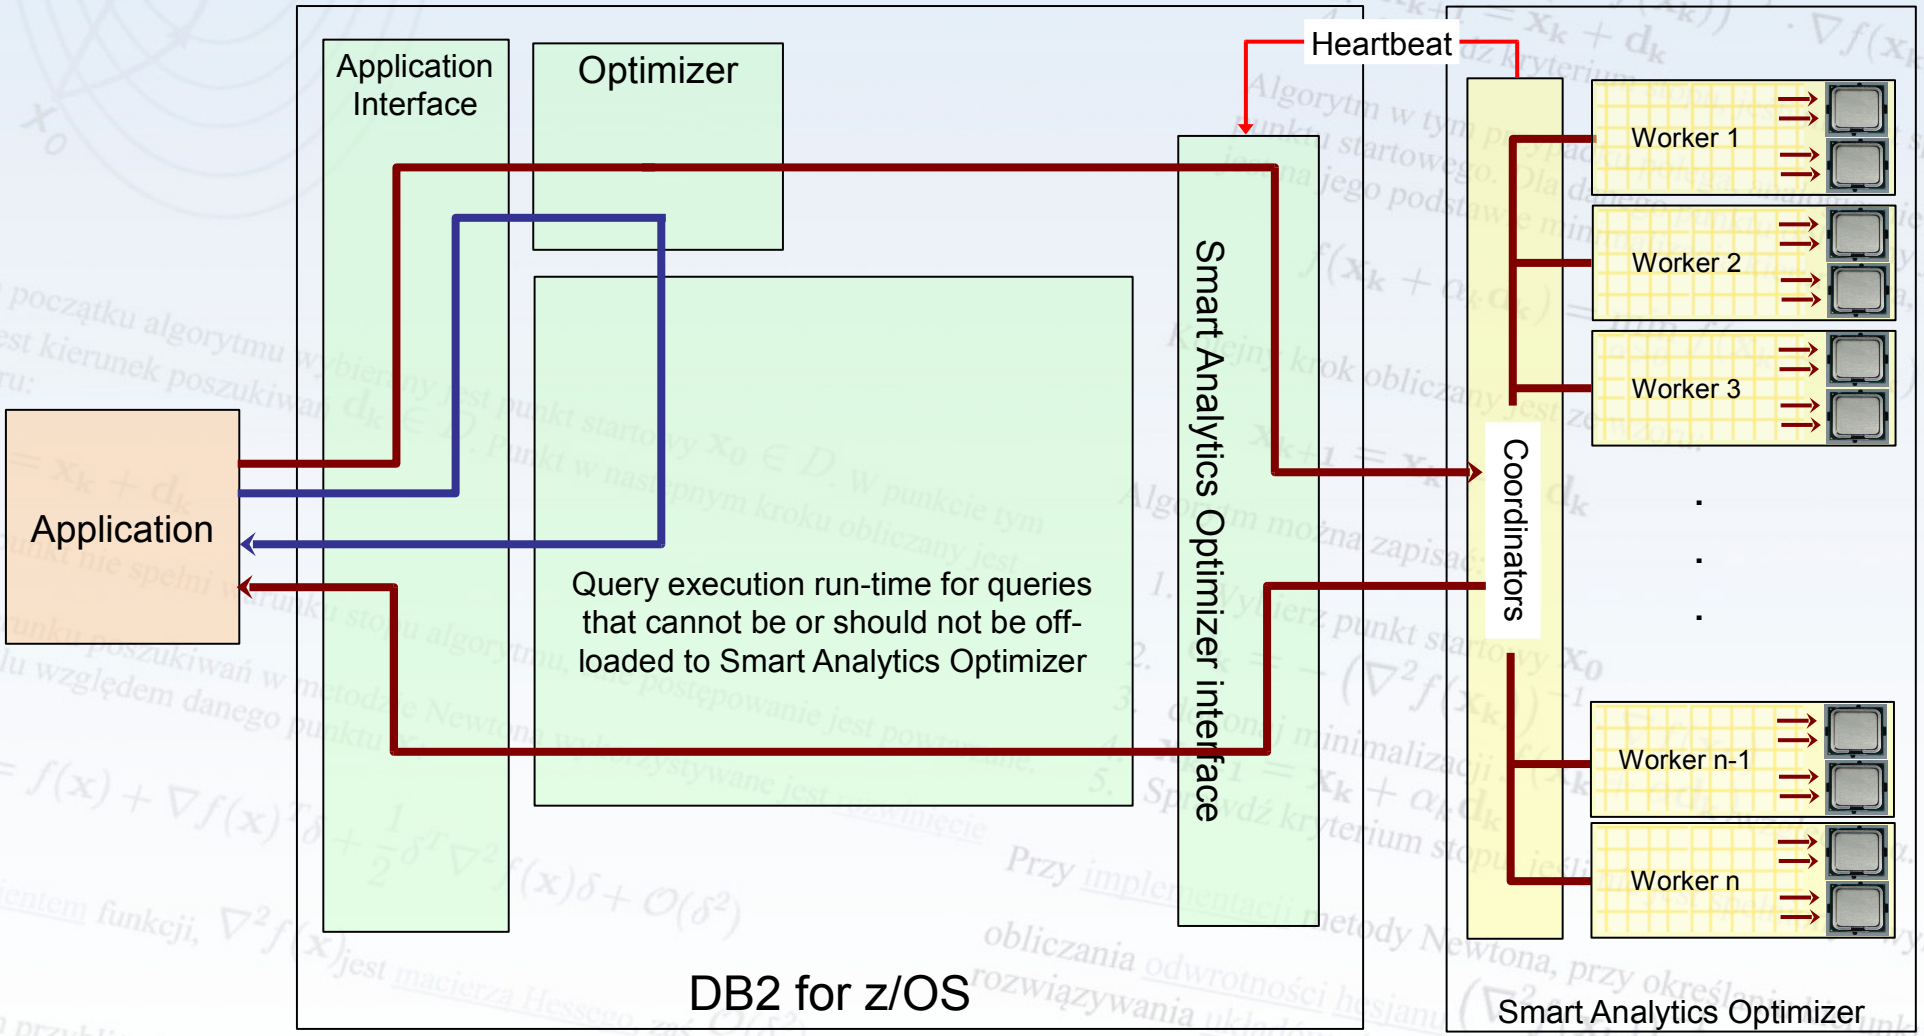


Heartbeat (Smart Analytics Optimizer availability and performance indicators)

Queries executed without Smart Analytics Optimizer

Queries executed with Smart Analytics Optimizer

# DISPLAY ACCELerator
## Resources Usage Indicators

- Queue length - average
- Queue length - high watermark
- Cluster processing capacity in *bogomips*
- Number of CPU cores
- Accumulated number of successful query requests
- Accumulated number of failed query requests
- Accumulated number of failed query requests due to Smart Analytics Optimizer being in invalid state, such as over-committed or in maintenance
- Available and in-use disk storage
- Current number of active coordinator nodes
- Average CPU utilization of coordinator nodes in last 60 seconds
- Average memory utilization on coordinator nodes since Smart Analytics Optimizer start
- Available memory for coordinator nodes
- Current number of active worker nodes
- Average CPU utilization of worker nodes in last 60 seconds
- Average memory utilization on worker nodes since Smart Analytics Optimizer start
- Available memory for worker nodes
- Average shared memory utilization on worker nodes since Smart Analytics Optimizer start
- Maximum shared memory utilization on worker nodes since Smart Analytics Optimizer start
- Available shared memory for worker nodes (controlled by a configuration parameter)
- Ratio between replicated (typically dimensions) and distributed (typically fact table) data
- Average query queue wait time in last 60 seconds
- Maximum query queue wait time since Smart Analytics Optimizer start

# Additional DB2 Support

- Explain
  - Indicates Smart Analytics Optimizer involvement in query execution or the reason for no usage
  - New table DSN_QUERYINFO_TABLE
- Instrumentation
  - Smart Analytics Optimizer availability and performance indicators
- DB2 Commands
  - DISPLAY THREAD
  - DISPLAY ACCELerator
  - START ACCELerator
  - STOP ACCELerator

# DISPLAY THREAD Support

- Limiting output to threads currently being executed by Smart Analytics Optimizer
  - New ACC filter on DISPLAY THREAD command
- Explicit indication if a thread is being executed by Smart Analytics Optimizer
  - Status field shows AC
- Showing Smart Analytics Optimizer's name and network identifiers (IP address, ports)
  - When DETAIL output is requested on DISPLAY THREAD

```
-DIS THD(*) ACC(*) DETAIL

DSNV401I # DISPLAY THREAD REPORT FOLLOWS -
DSNV402I # ACTIVE THREADS -
NAME   ST A REQ ID  AUTHID    PLAN     ASID TOKEN
BATCH AC  *   231 BI  ADMF001  DSNTEP2 0053 55
V666 ACC=BLINK1,ADDR=::FFFF:9.30.30.133..446:1076
```

# Testing Results

- The problem queries provided by a customer
- Expert database tunning done on all the queries
  - Q1 – Q6 even after tuning run for too long and consume lots of resources
  - Q7 improved significantly – no 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 Smart Analytics Optimizer

| Query | Query Elapsed Time | | | Query CPU Consumption on System z | | |
|---|---|---|---|---|---|---|
| | DB2 only | DB2 with Smart Analytics Optimizer | Speed-up | DB2 only | DB2 with Smart Analytics Optimizer | Saving |
| Q1 | 0:02:43.0 | 0:00:03.4 | 48 | 0:06:31.0 | 0.004495 | ~100% |
| Q2 | 0:38:31.0 | 0:00:04.5 | 511 | 0:48:02.0 | 0.004713 | ~100% |
| Q3 | 0:00:25.0 | 0:00:02.2 | 12 | 0:00:19.0 | 0.099702 | 99.48% |
| Q4 | 0:26:33.0 | 0:00:07.8 | 206 | 0:34:33.0 | 0.005174 | ~100% |
| Q5 | 0:00:35.0 | 0:00:08.3 | 4 | 0:00:38.0 | 0.520915 | 98.63% |
| Q6 | 1:30:35.0 | 0:00:03.8 | 1424 | 7:23:26.0 | 0.003979 | ~100% |
| Q7 | 0:00:02.0 | 0:00:02.0 | 1 | 1.361983 | 1.361983 | 0.00% |
| Total | 2:39:24.0 | 0:00:32.0 | 298 | 8:53:31.0 | 2.000961 | 99.99% |

Uniform elapsed times

Average speed-up almost 300 times

Significant CPU saving

Note: Your mileage will vary. This particular query mix is suited exceptionally well for IBM Smart Analytics Optimizer

# Agenda

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

# What Is Smart Analytics Optimizer Ideally Suited For?

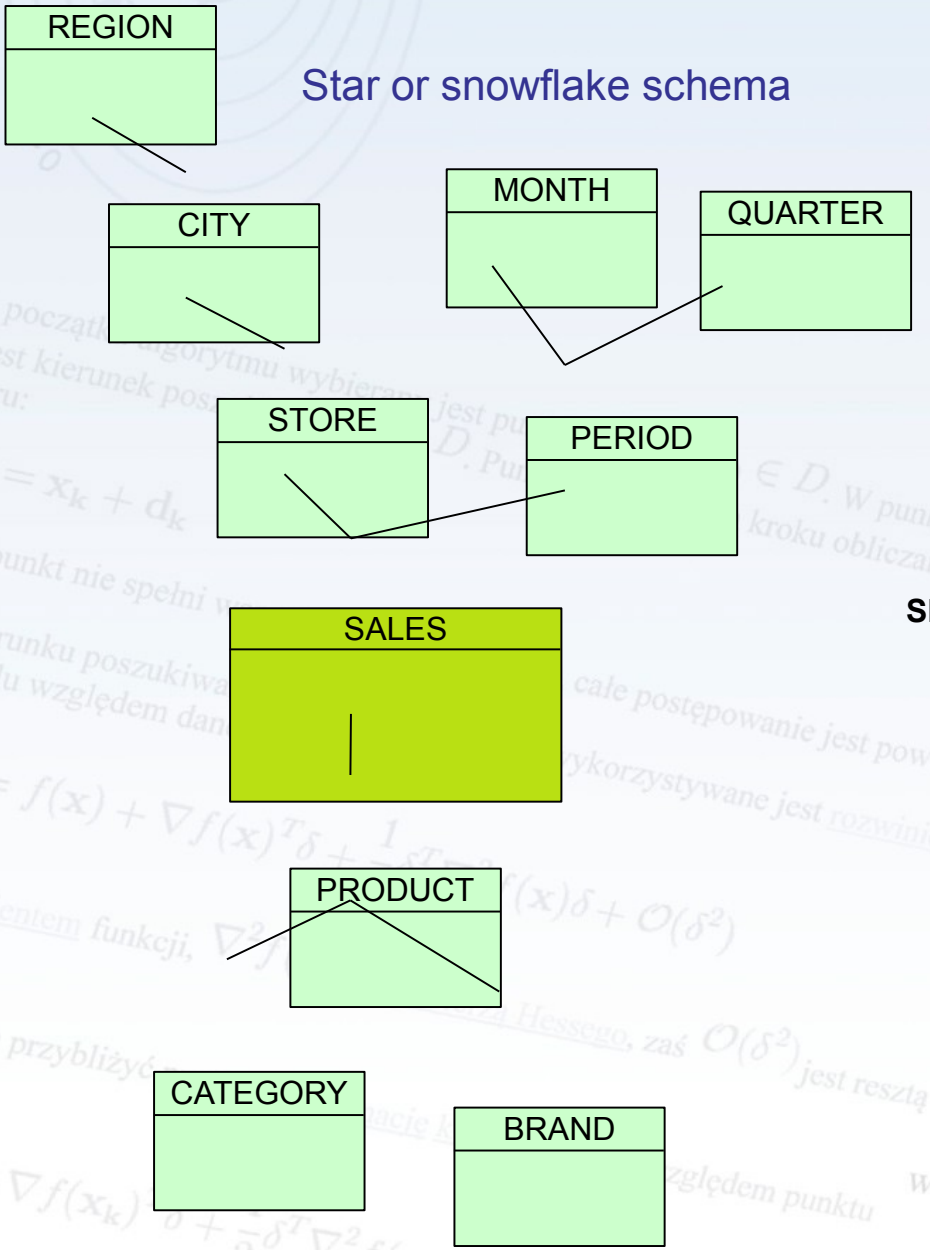


Complex, OLAP-style queries that typically:

- 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
```

# First Release Restrictions

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

Nested Table Expression

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

Common Table Expression

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

IN predicate with subquery

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

# First Release Restrictions

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

# Data Types Support

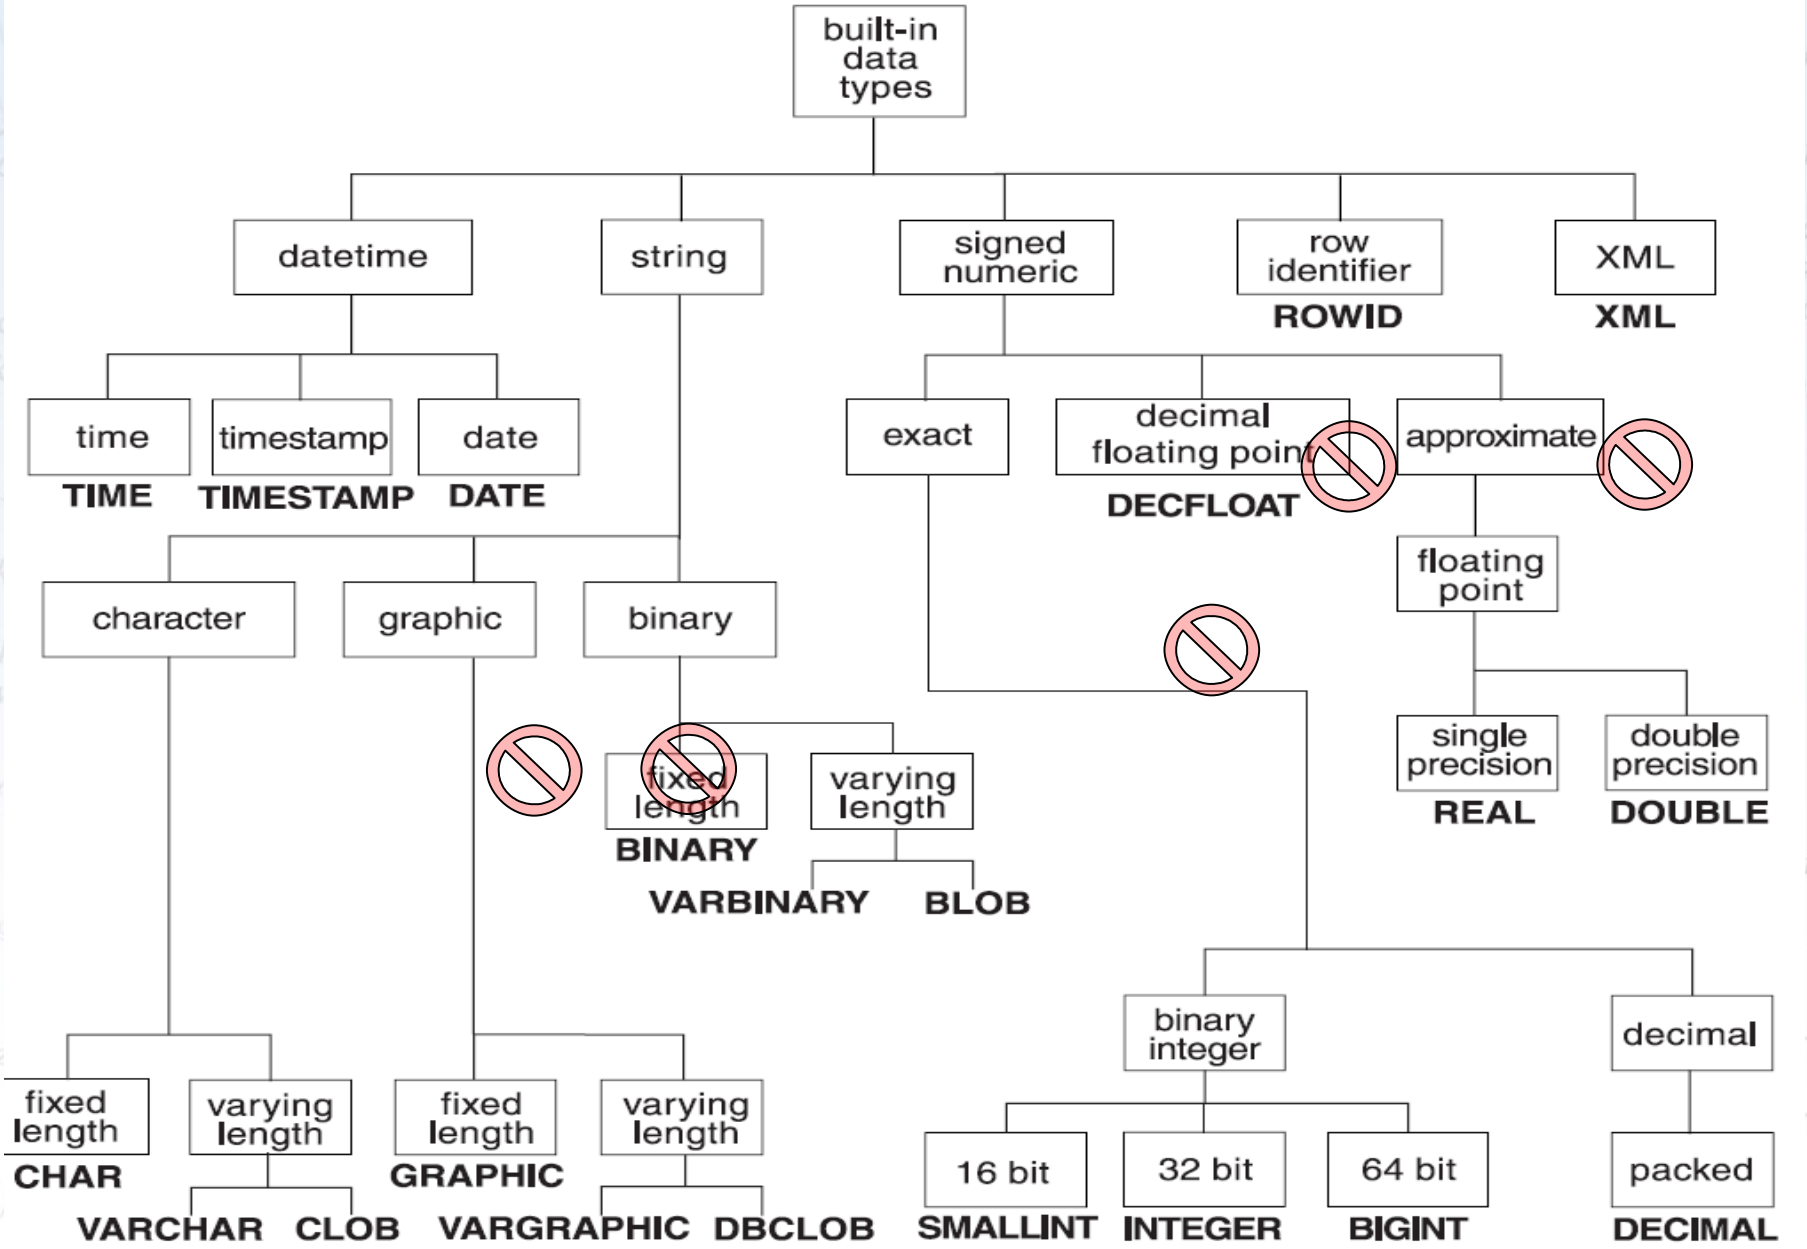

# Options for Workload Analysis

| Stage | Purpose |
|---|---|
| Questionnaire | • Initial assessment based on size, query response time, update characteristics and customer pain points |
| Quick Workload Test | • Assessment based on dynamic customer workload, runtime statistics, table sizes and SQL. |
| Detailed Online Workload Analysis | • Assessment based on data mart definition for customer data model and offload capabilities in a real Smart Analytics Optimizer environment. Addresses all inhibitors for offload and data mart definition questions. |

Contact Data Warehousing on System z Center of Excellence at dwhz@de.ibm.com

# Quick Workload Test

Report for a first assessment:

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

## 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

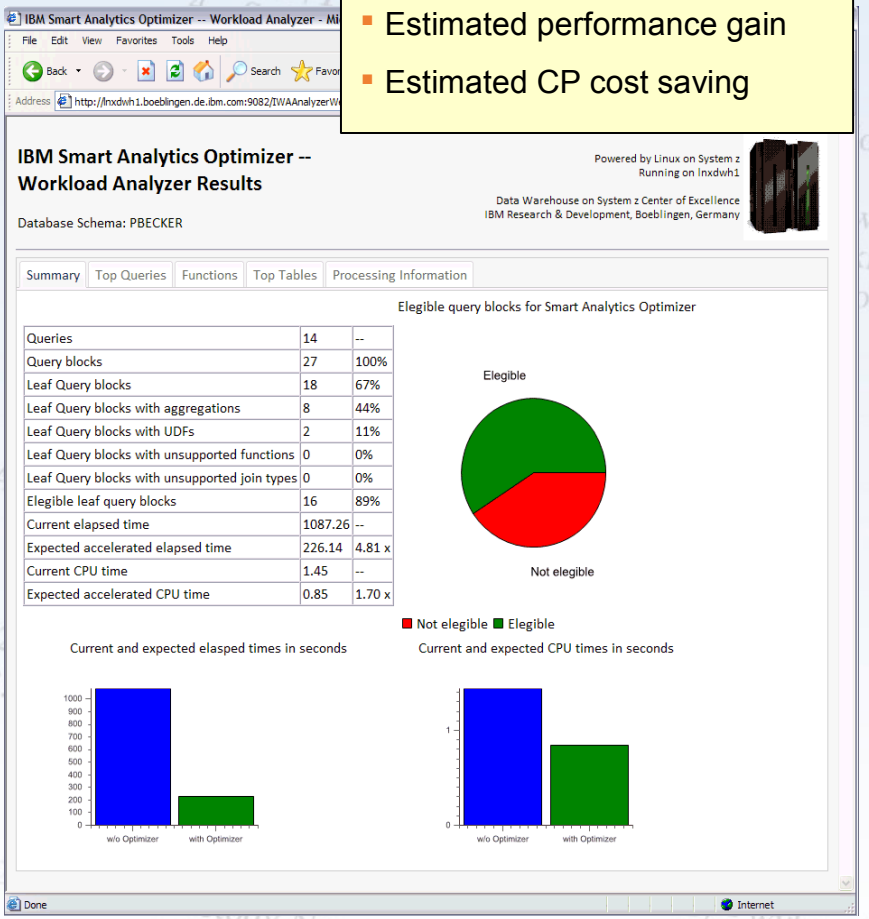


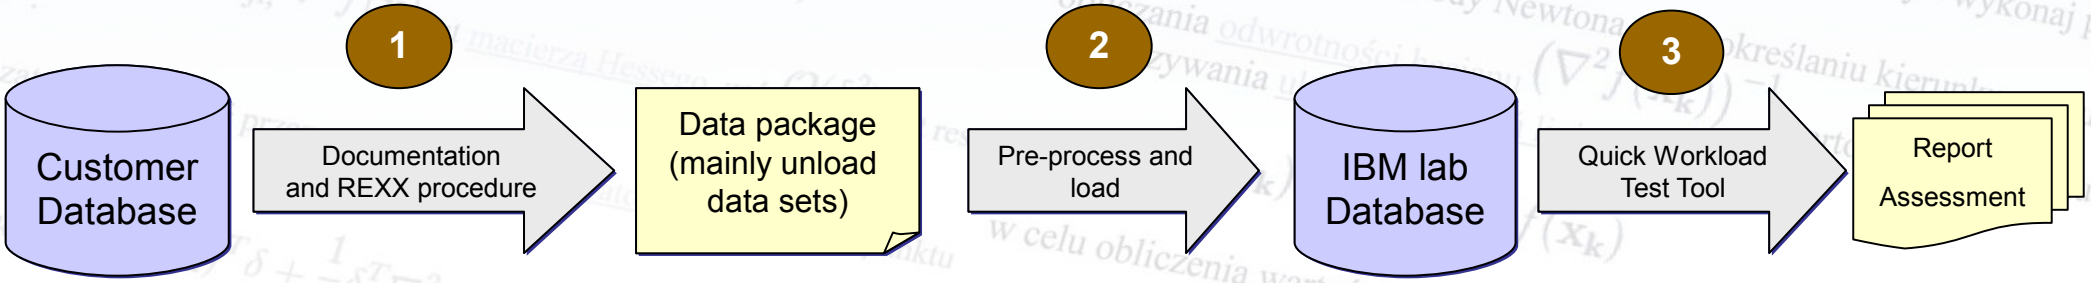

# Summary of Value Proposition

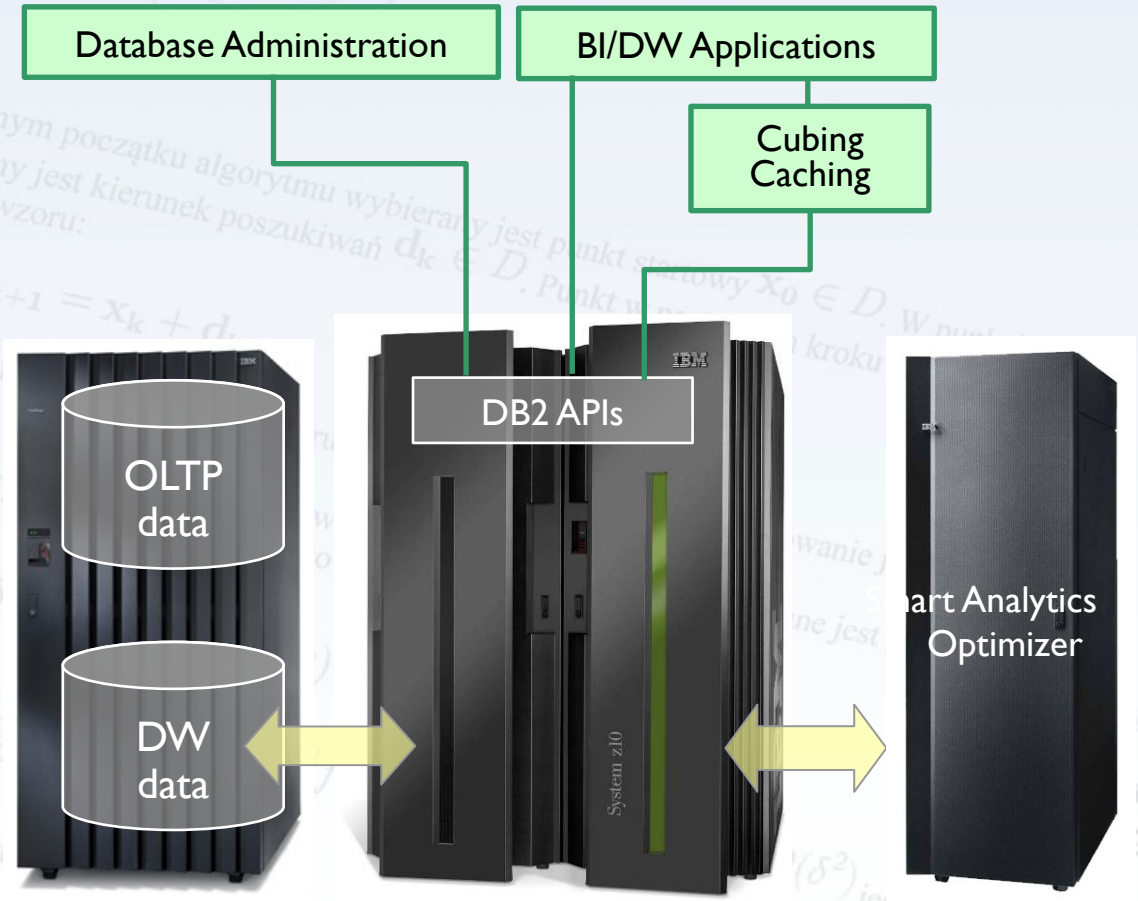


Seamless integration of new computing paradigms into proven technology
- Massive multi-core and vector processing
- In-memory database
- Hybrid row- and column-based store
- No changes to the applications, applications continue to attach to DB2
- Preserving traditional System z and DB2 quality of service, full fencing and protection of DB2 against possible Smart Analytics Optimizer 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

# IBM Smart Analytics Optimizer

## Unlocking unprecedented value from enterprise data

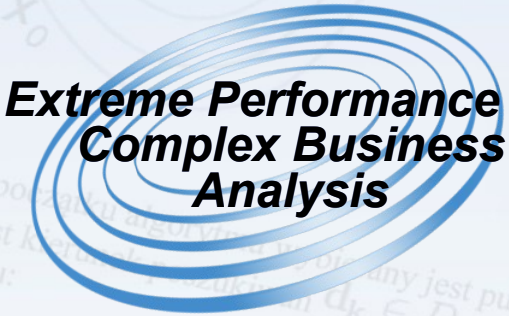

***Extreme Performance for Complex Business Analysis***

*Breakthrough technologies providing dramatic performance improvement*

***Database Performance Appliance***

*Quickly and simply deploy, hands-free operations, no query tuning, with application transparency*

***Proven operational characteristics***

*Inherits the availability, reliability, and security of System z*