


# IBM Smart Analytics Optimizer



Dr. Manfred Oevers – Release Manager IBM Smart Analytics Optimizer

# Sometimes change is dramatic and changes the rules



The speed limit on a street is raised from 60 km/h to 80 km/h

*People get to their destination in  $\frac{3}{4}$  of the time*

*(Minor Impact)*

**Versus**



Multiple locations are tied together with a High Speed Train

*New opportunities, new ways of thinking about what is close vs far away* **(Major Impact)**

***"The future belongs to those who see possibilities before they become obvious."***

John Sculley  
Former CEO of Pepsi & Apple Computer

***"So many of our dreams seem impossible, then improbable, then inevitable."***

Christopher Reeve

***"The only way to predict the future is to have the power to shape it."***

Eric Hofer




# Imagine If Your Decision Makers Could...

*...predict and treat high-risk patients to proactively intervene at time of visit?*

*...identify suspects based upon initial witness information on arrival at scene?*

*...provide targeted cross sells that drive additional sales when offered at time of sale?*

*... predict correct order levels at time of purchase?*



**Physician**

**Police Detective**

**Retail Sales Manager**

**Purchasing Agent**

***...optimize every transaction, process and decision at the point of impact, based on the current situation, without requiring that everyone be an analytical expert***

# 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

# IBM Smart Analytics Optimizer –

*Capitalizing on the best of relational and the best of columnar databases*

## What is it?

The IBM Smart Analytics Optimizer is a workload optimized, appliance-like, add-on, that enables the integration of business insights into operational processes to drive winning strategies. It accelerates select queries, with unprecedented response times.



## How is it different?


- § Performance: Unprecedented response times to enable 'train of thought' analyses frequently blocked by poor query performance.
- § Integration: Connects to DB2 through deep integration providing transparency to all applications.
- § Self-managed workloads: queries are executed in the most efficient way
- § Transparency: applications connected to DB2, are entirely unaware of the Smart Analytics Optimizer
- § Simplified administration: appliance-like hands-free operations, eliminating many database tuning tasks



**Breakthrough Technology Enabling New Opportunities**

# Extreme performance for complex queries

## Extreme performance for complex queries

 Rapidly delivers information to decision makers through breakthrough technologies providing dramatic performance improvement.


- The best of row and columnar store technologies

- Highly compressed data

- Compressed data operations

- In-memory processing

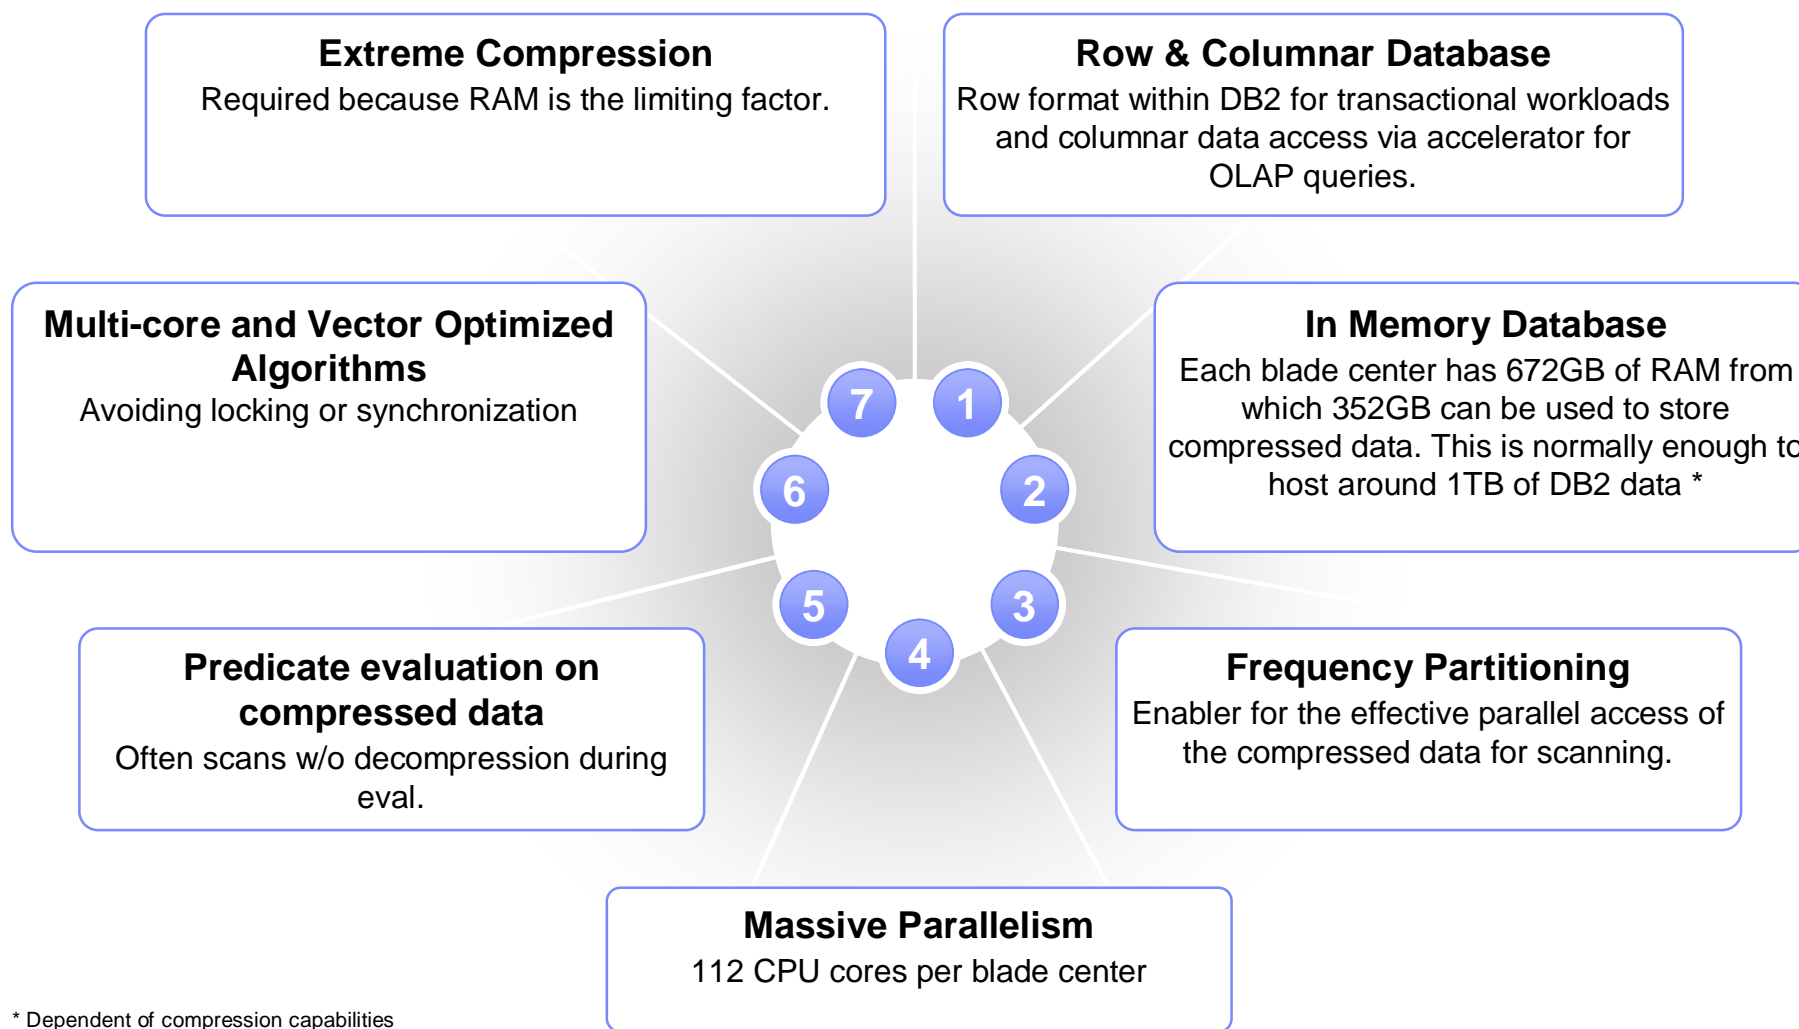
- Massively parallel architecture

 Enables decision makers to submit queries they never dared in the past, that analyze trends, predict outcomes, and produce better business results.

One Beta customer asked us to repeat a query under lab conditions because he couldn't believe the acceleration. A query execution time was reduced from 13 minutes, 42 seconds to just one second (end to end)!



# Breakthrough technologies for performance



\* Dependent of compression capabilities

# Orders of Magnitude Faster for Queries

*Beta Customer Results*

**For customers who have struggled to achieve the required performance from their complex BI queries – *the results are astounding!***

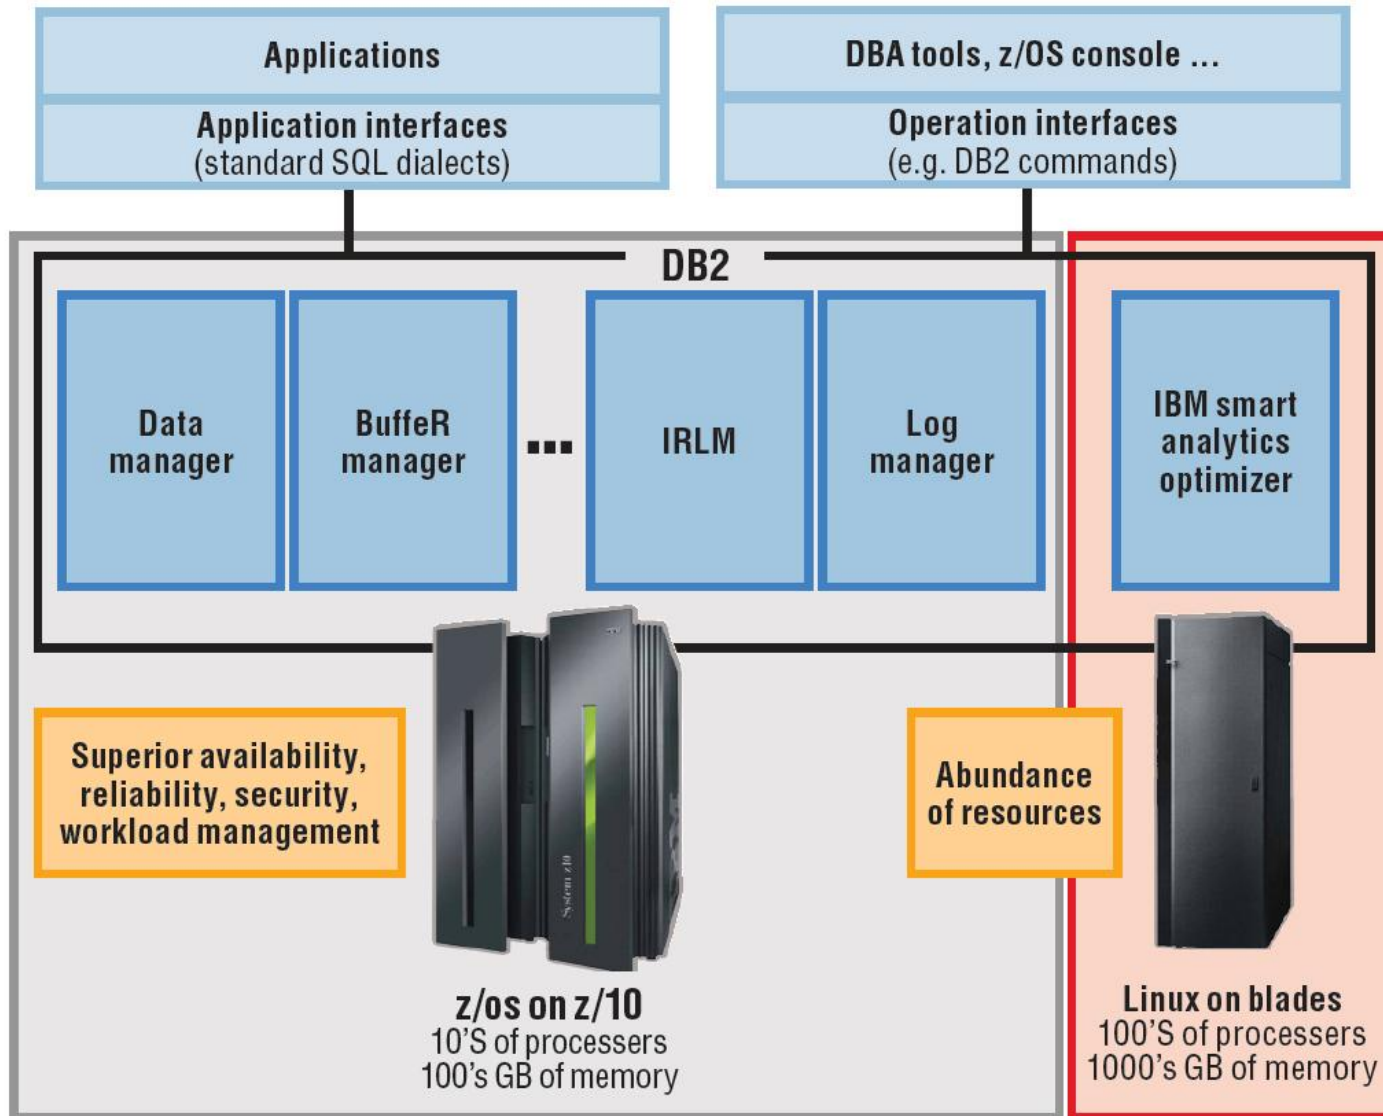


**... and its acceleration factor ::**

Runtime of queries w/o the Smart Analytics Optimizer	Factor	with Smart Analytics Optimizer
163 s	48	3s
2311 s	511	5s
25 s	12	2s
1593 s	206	8s
35 s	4	8s
5435 s	1424	4s

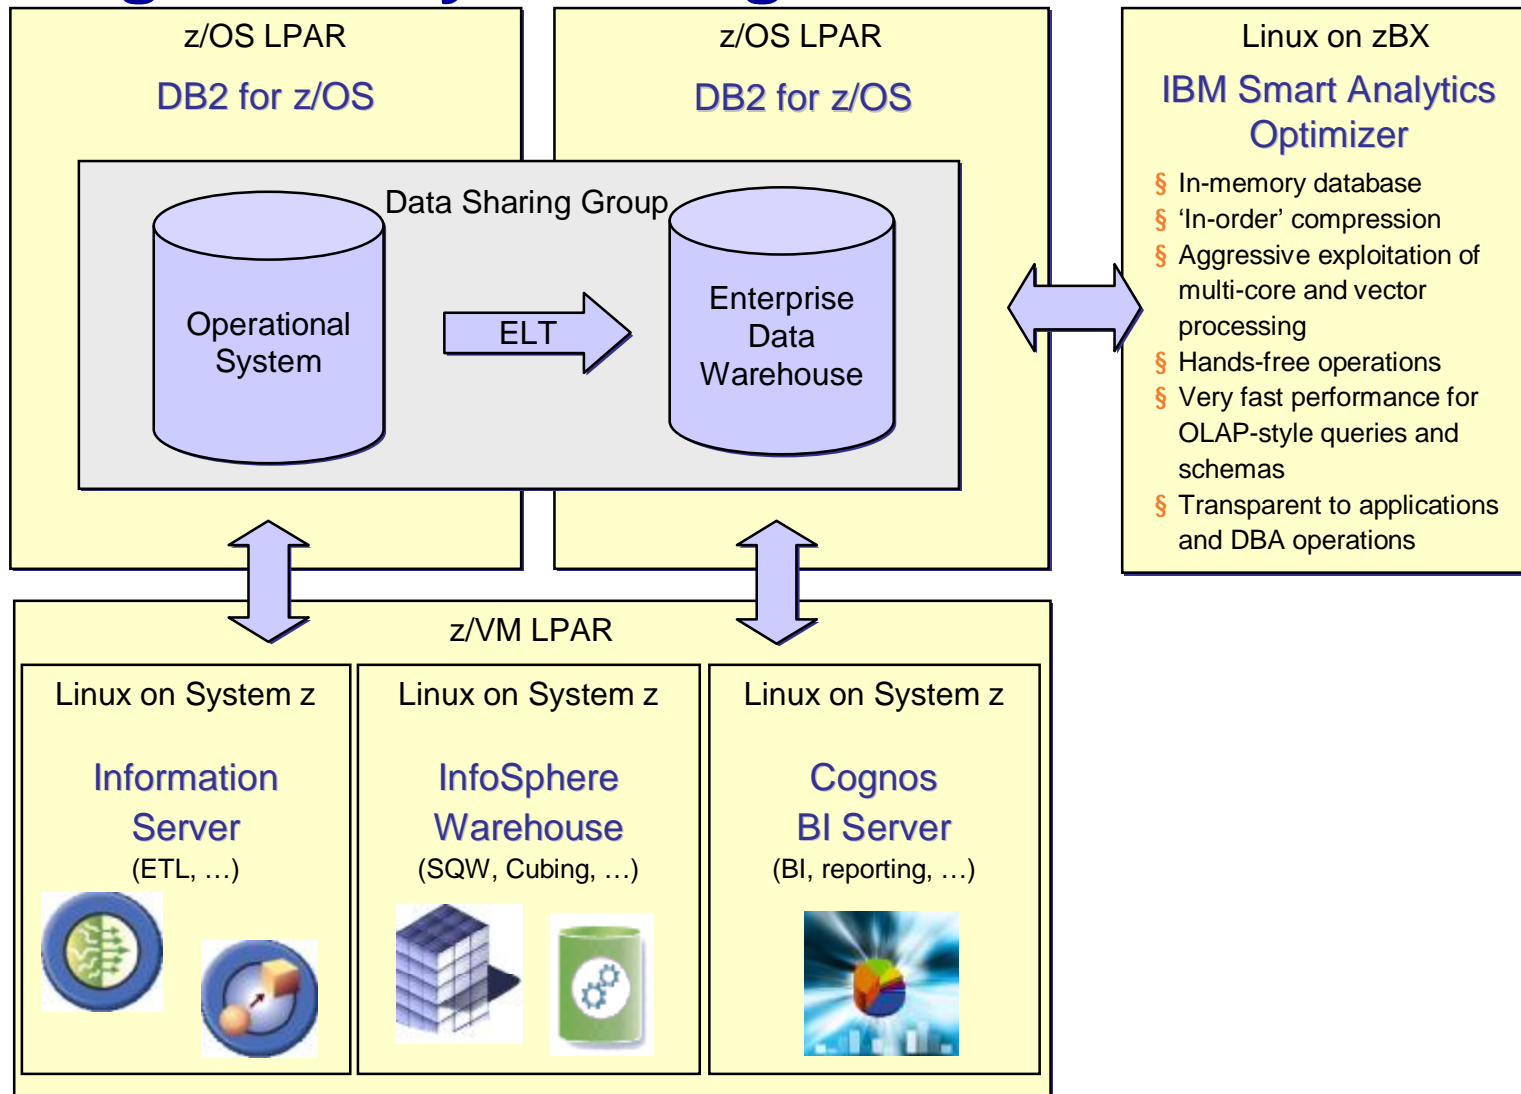


# IBM Smart Analytics Optimizer - a Virtual DB2 Component



# IBM Smart Analytics Optimizer

## Adding Industry Leading Performance

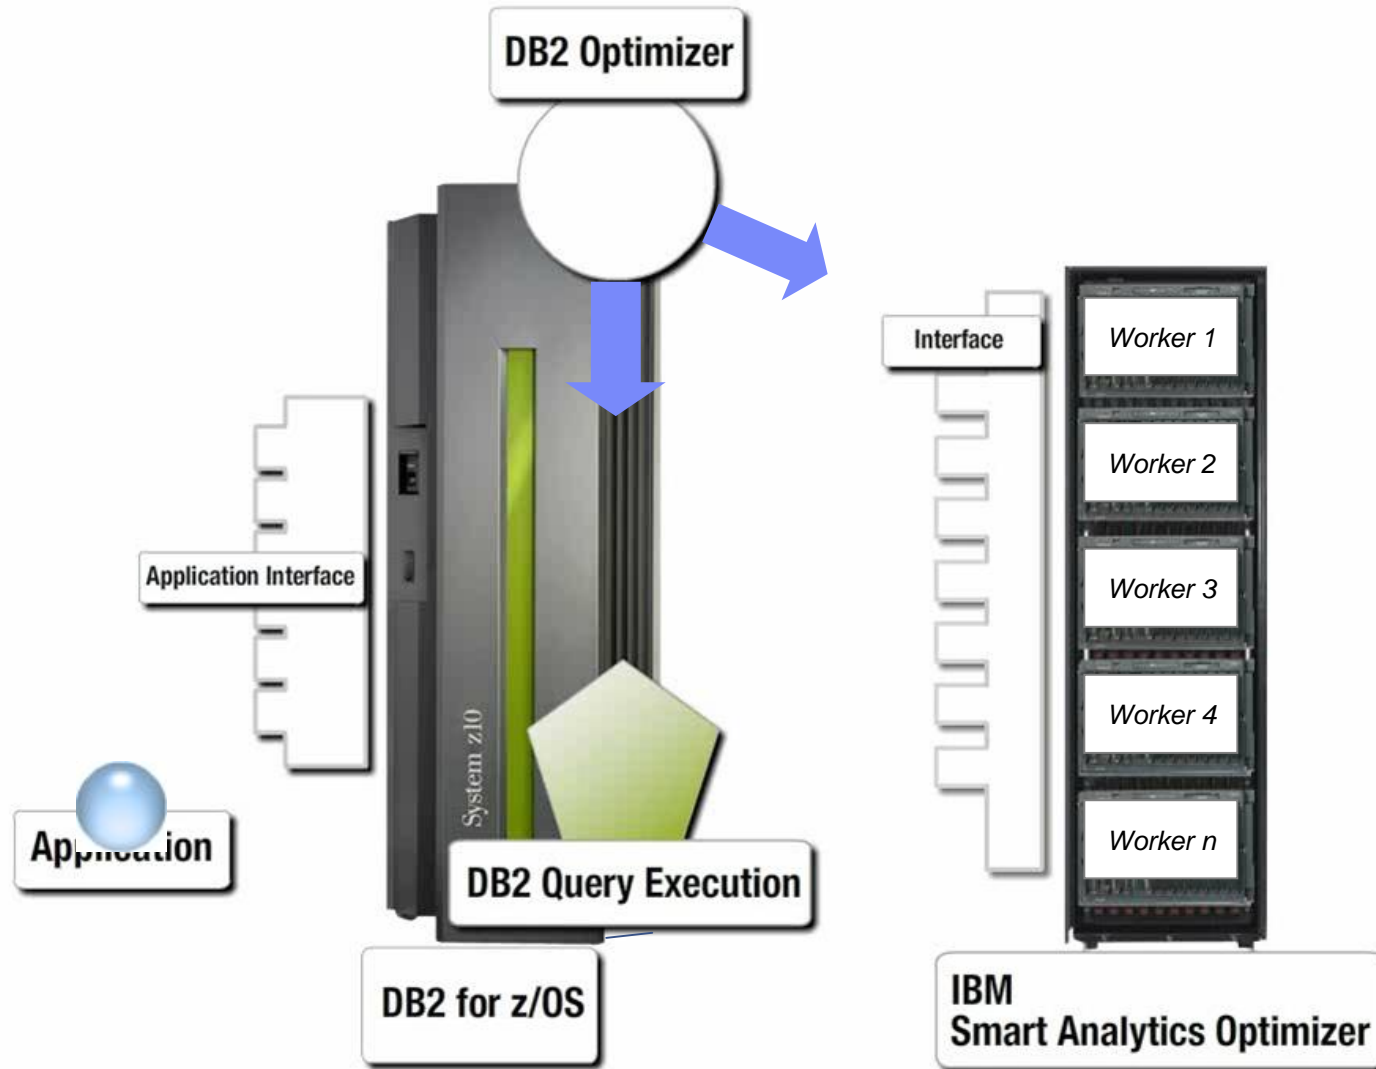


# Routing Queries to the IBM Smart Analytics Optimizer

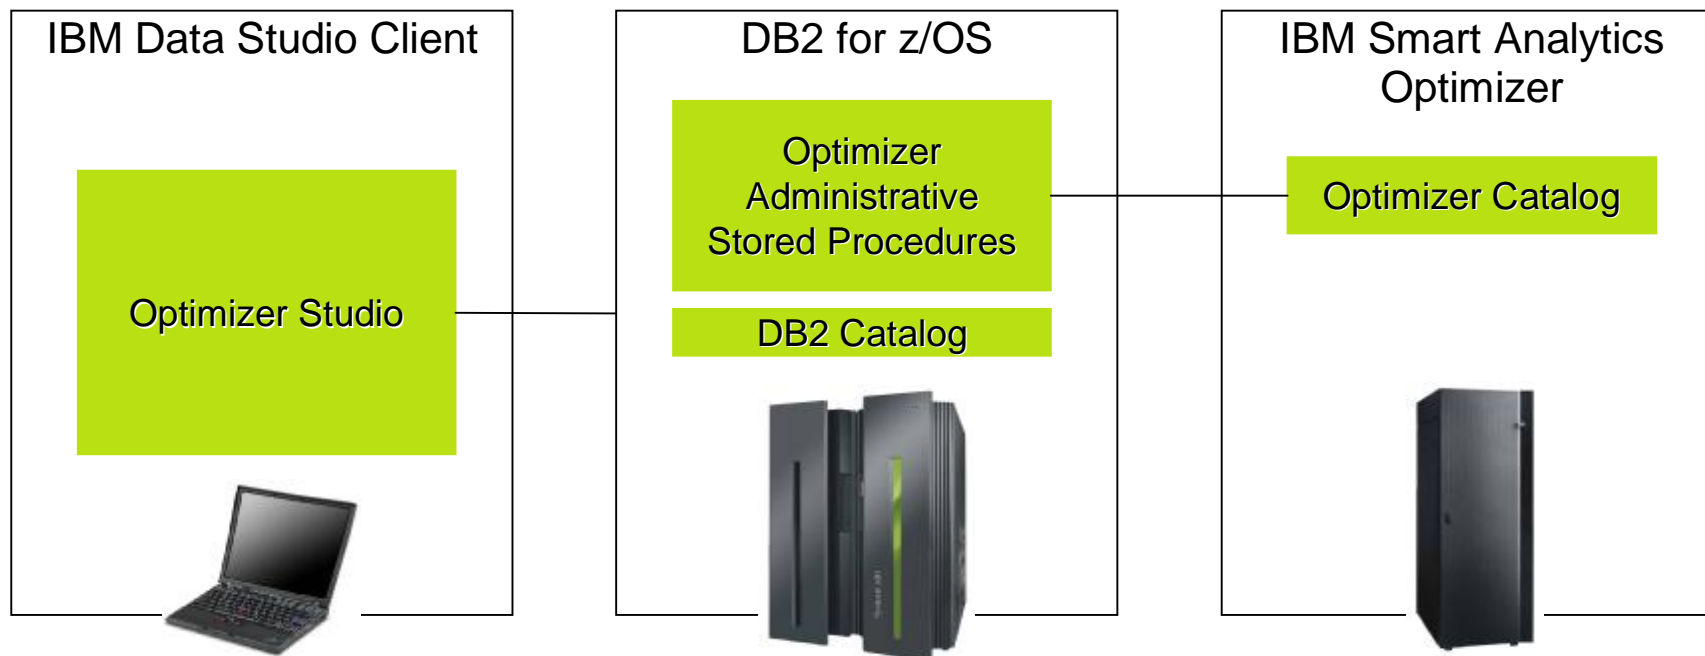
- § Routing of Queries to the Smart Analytics Optimizer is determined by the DB2 for z/OS Optimizer
- § Uses a mechanism called Accelerator Query Tables (AQTs)
- § Re-uses MQT routing functionality



# Query Execution Flow



## IBM Smart Analytics Optimizer Mart Definition and Deployment



- IBM Smart Analytics Optimizer marts need to be defined and deployed to IBM Smart Analytics Optimizer before data is loaded and queries sent to IBM 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 IBM Smart Analytics Optimizer catalog.
- IBM Smart Analytics Optimizer Studio guides you through the process of defining and deploying marts, as well as invoking other administrative tasks.
- IBM 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 IBM Smart Analytics Optimizer from the outside world including IBM 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

Tahoma 9 B I A [Font icons] | [Diagram icons] | 1

Upgrade to Optim [Upgrade icon]

**Data Project Explorer**

- SalesMartsProject
  - Data Marts
  - TPC-H

**\*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

EMP\_EMPLOYEE\_DIM

- EMPLOYEE\_KEY
- MANAGER\_KEY
- MANAGER\_CODE
- MANAGER\_CODE
- MANAGER\_CODE
- MANAGER\_CODE
- MANAGER\_CODE
- MANAGER\_CODE
- MANAGER\_CODE
- MANAGER\_CODE

SLS\_SALES\_FACT 18 MB

- DAY\_KEY

**Palette**

Mart

- Table
- 1:n
- n:m

**Data Source Explorer**

- Stored Procedures
- Synonyms
- Tables
  - EMP\_EMPLOYEE\_DIM
  - GO\_TIME\_DIM
  - SLS\_PRODUCT\_DIM
  - SLS\_SALES\_FACT
- User-Defined Functions
- User-Defined Types

Properties Tasks Problems Error Log Model Report SQL Results Bookmarks

**<Data Mart> SalesMart1**

**General** Validate Deploy...

Validation Name: SalesMart1

Space Required Estimated size: 20 MB

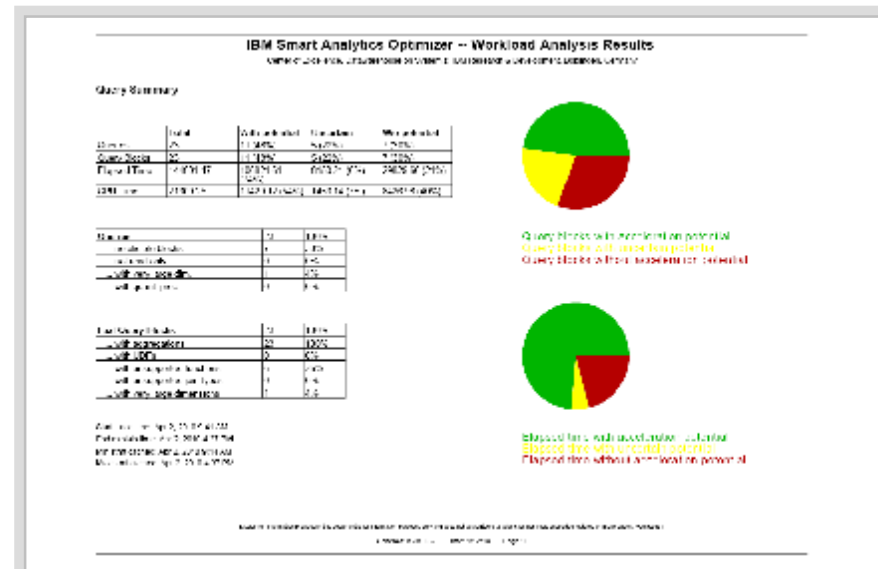
# Quick Workload Test

## 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 (Boeblingen)

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




Report for a first assessment:

- § Acceleration potential for
- § Queries
- § Estimated time
- § CP Kosten



## Preconfigured Solution Sizes for the IBM Smart Analytics Optimizer



Number of 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...

1. For existing zEnterprise 196, order **System z BladeCenter Extension**
  - Machine Type 2458 Model 002
  - Order required number of ISAOPT Blades (feature code #0610)
2. **DS5020 controller**, disks and expansion units need to be ordered separately from zBX, and housed in a separate rack.
3. Order required Value Units for **IBM Smart Analytics Optimizer Software PID 5697-AQT**



---

## *Websites for Additional Information:*

### **IBM Smart Analytics System Optimizer:**

<http://www-01.ibm.com/software/data/infosphere/smart-analytics-optimizer-z/>

### **IBM Smart Analytics System 9600:**

<http://www-01.ibm.com/software/data/infosphere/smart-analytics-system/>

### **Data Warehousing and Analytics**


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

### **Data Warehousing and Business Intelligence on System z**

<http://www-01.ibm.com/software/data/businessintelligence/systemz/>


---

# Thank you

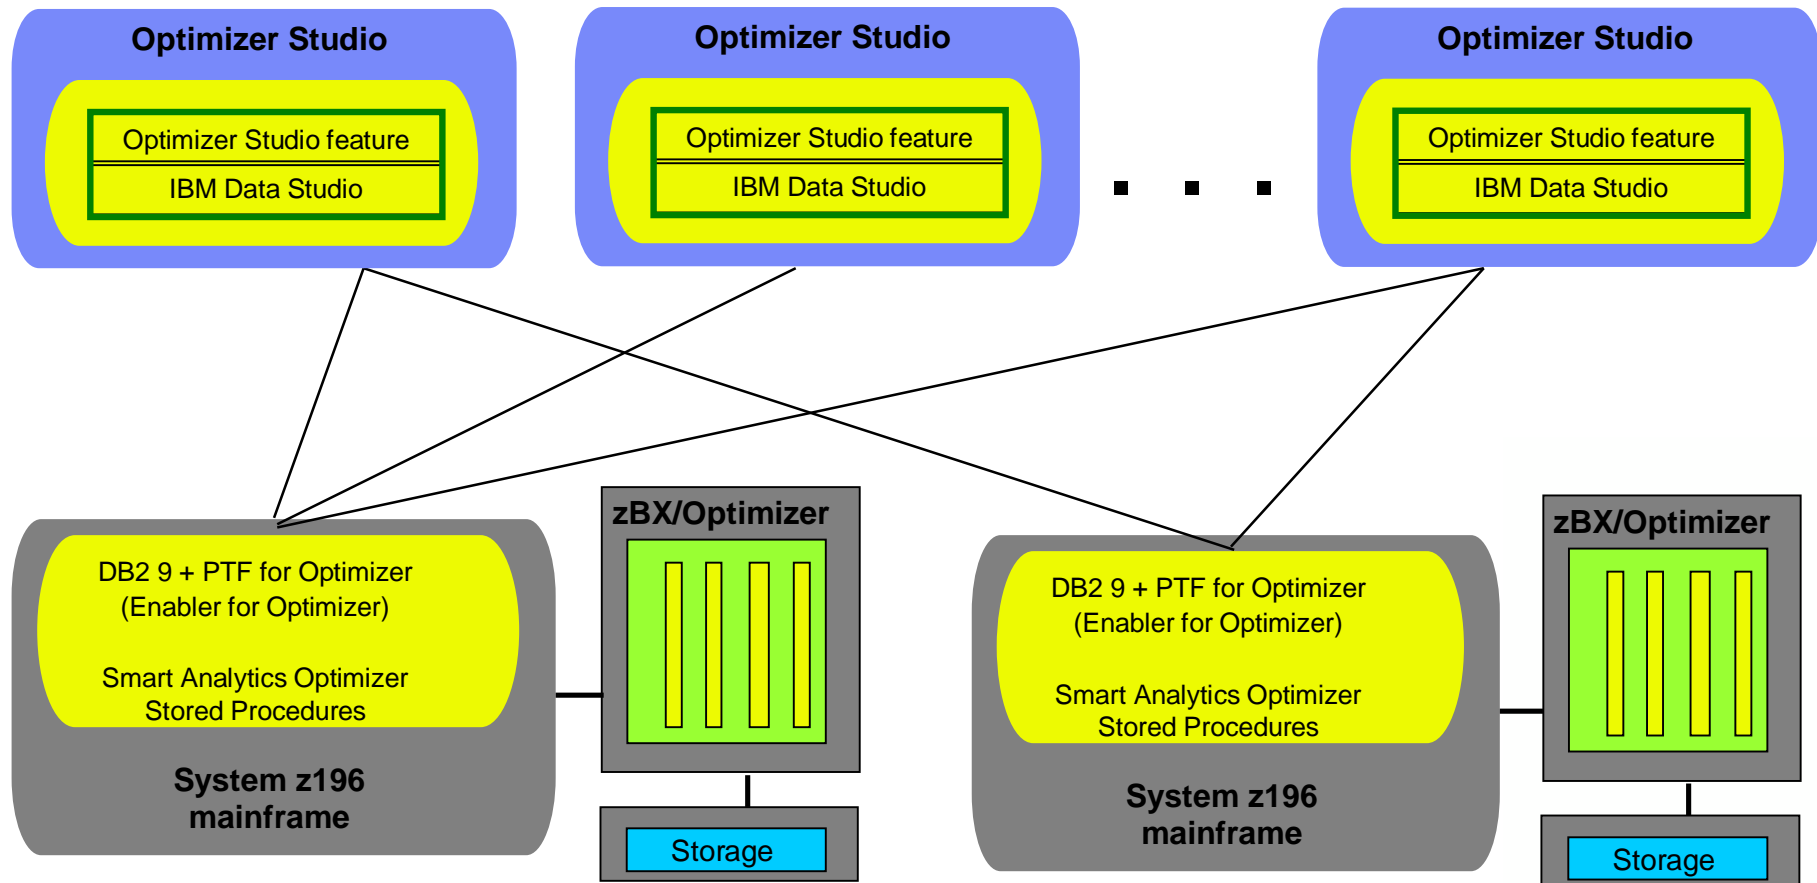


---

# Back Up

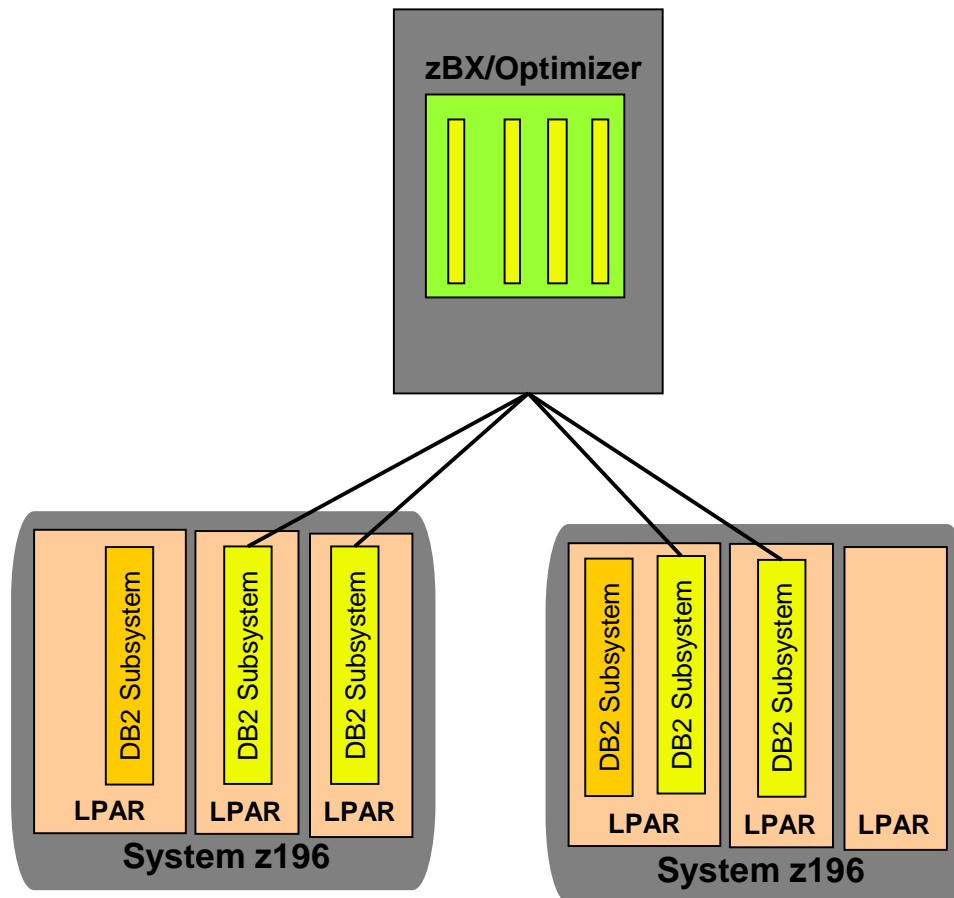


# IBM Smart Analytics Optimizer Studio configuration example



IBM Smart Analytics Optimizer Studio may access multiple DB2 subsystems and a DB2 subsystem may be accessed by multiple IBM Smart Analytics Optimizer Studio clients.

# zBX / IBM Smart Analytics Optimizer - multiple DB2 z systems configuration example



**Multiple** DB2 subsystems running on  
 § the same or different **LPARs**  
 § the same or different **CECs** (up to 8)  
 may access a shared  
 IBM Smart Analytics Optimizer/zBX  
 Via the zBX built-in **Top of Rack switch**

Connected to zBX  
 Not Connected to zBX

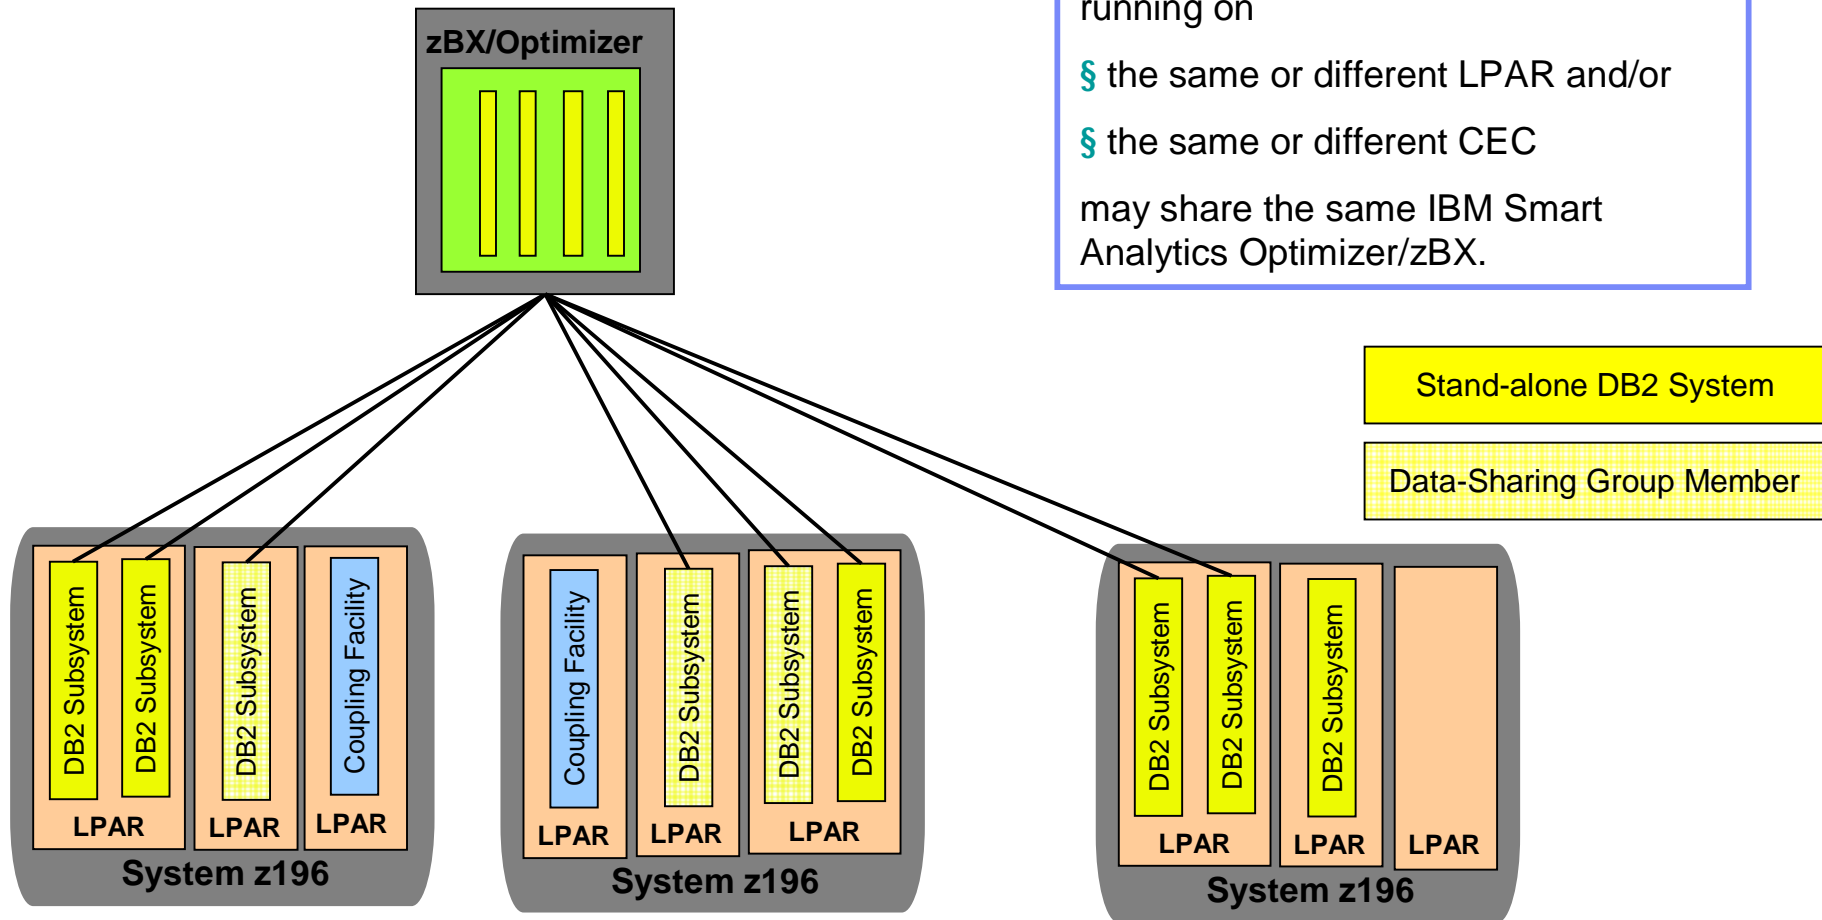
# zBX / IBM Smart Analytics Optimizer - Sysplex configuration example

All members of a **DB2 data sharing group** may access a shared IBM Smart Analytics Optimizer/zBX.


Additional standalone DB2 subsystems running on

- § the same or different LPAR and/or
- § the same or different CEC

may share the same IBM Smart Analytics Optimizer/zBX.



# Multiple zBX / IBM Smart Analytics Optimizer systems configuration example



1

Supported and Tested in GA1

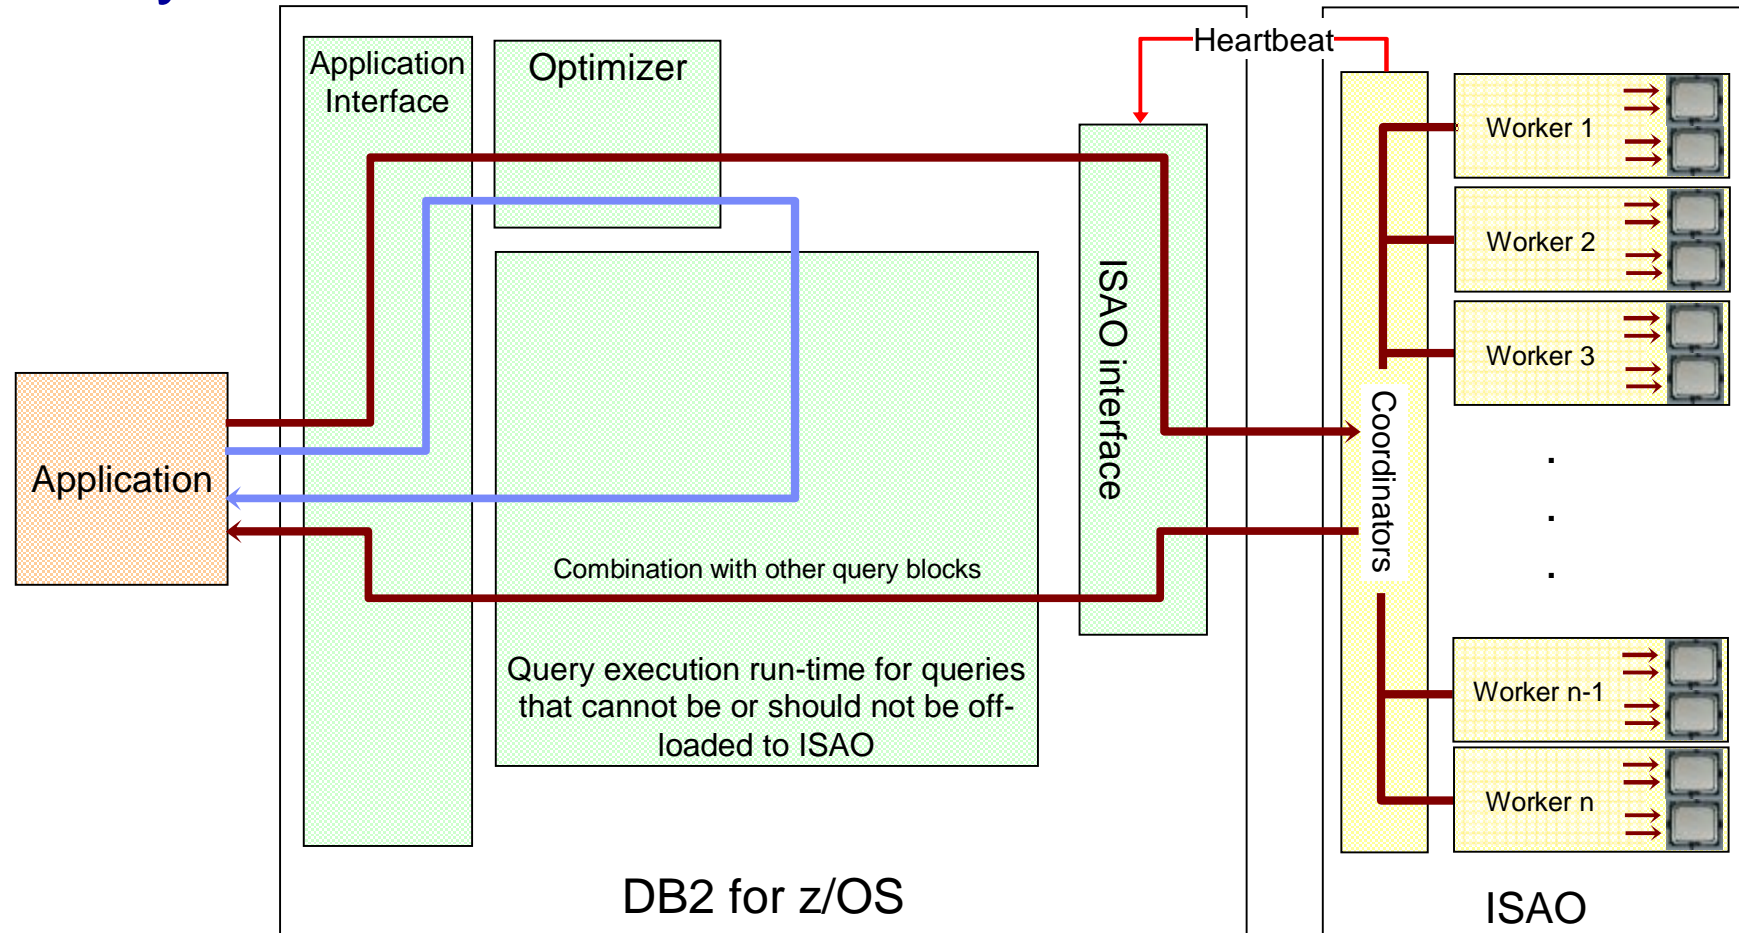
~~2~~

One DB2 subsystem must not access more than one IBM Smart Analytics Optimizer/zBX.

Only one IBM Smart Analytics Optimizer/zBX may be connected to one CEC.

- **Currently not supported**
- **Considered for future releases**

# Query Execution Process Flow



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



# IBM Smart Analytics Optimizer Installation

## Preparation tasks

- **Install IBM Smart Analytics Optimizer Studio** on a client workstation.
- **Modify DB2 for z/OS** so that IBM Smart Analytics Optimizer can be used by a DB2 subsystem on your DB2 for z/OS data server.
- **Mount** the IBM System **Storage unit DS5020 & Configure** the IBM System Storage
- Allow two **OSA-Express3** 10 GbE cards per CPC for connectivity to the zBX

## Installation of the zEnterprise BladeCenter Extension (zBX)

- **Installing the zEnterprise BladeCenter Extension.**
  - **Wiring** OSA-Express3 cards
  - **Wiring** the IBM System Storage
- **Blade entitlement.**
  - This installs the IBM Smart Analytics Optimizer Code on the zBX. The Product DVD shipped with IBM Smart Analytics Optimizer is needed for this step

## Post-installation tasks

- **Associate IBM Smart Analytics Optimizer with DB2 for z/OS** to authenticate IBM Smart Analytics Optimizer as an entitled application of DB2 for z/OS.
  - Obtaining the **pairing code** for accelerator authentication
  - Creating the **authentication token** using the Add New Accelerator wizard
- Use IBM Smart Analytics Optimizer Studio for host-based configurations and to **define data marts** that are used for **downloading relevant data** to the zEnterprise BladeCenter Extension.