IBM Smart Analytics Optimizer

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

IBM*	System z*
IBM Logo*	System z10
DB2*	Tivoli*
Dynamic Infrastructure*	z10
GDPS*	z10 BC
HyperSwap	z/OS*
InfoSphere	z/VM*
Parallel Sysplex*	z/VSE
RACF*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
INFINIBAND, InfiniBand Trade Association and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

- Overview over the IBM Smart Analytics Optimizer
  - Which challenges are addressed?
  - What components are used?
- Different Data Stores for different Workloads
  - What is a Row-, Column- or Register Store?
  - Why is this relevant for query performance?
- Define what should be optimized
  - How to define what data / workload to accelerate?
  - How is data transferred and distributed?

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 ISAO
- ***Simplified administration:*** appliance-like hands-free operations, eliminating many database tuning tasks

Breakthrough Technology Enabling New Opportunities

Extreme performance for complex queries

Game Changing Performance

- ✓ **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)!

IBM

Breakthrough technologies for performance

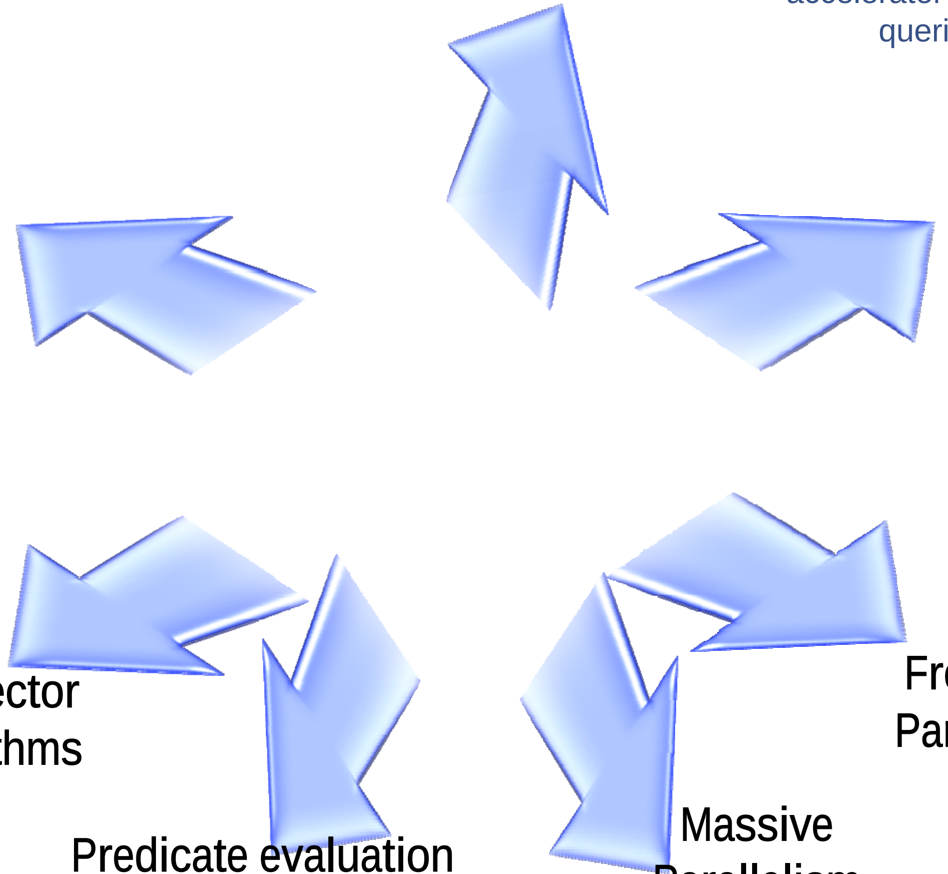

Row & Columnar Database

Row format within DB2 for transactional workloads and columnar data access via accelerator for OLAP queries.

In Memory Database

Each blade center has 672GB of RAM from which 352GB can be used to store compressed data. This is normally enough to host around 1TB of DB2 data (dep. on compression capabilities)

Frequency Partitioning

Enabler for the effective parallel access of the compressed data for scanning.

Massive Parallelism

112 CPU cores per blade center

Predicate evaluation on compressed data

Often scans w/o decompression during eval.

Multi-core and Vector Optimized Algorithms

Avoiding locking or synchronization

Extreme Compression

Required because RAM is the limiting factor.

IBM Smart Analytics Optimizer – Beta Customer

Orders of magnitude faster for queries within Sweetspot

The ISAO Sweetspot:

- Complex OLAP-style queries
- Scan large subset of data
- Look for trends – period on period analysis
- Assist in making actionable business decisions

... and its acceleration factor : :

Runtime of queries w/o ISAO	Factor	with ISAO
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

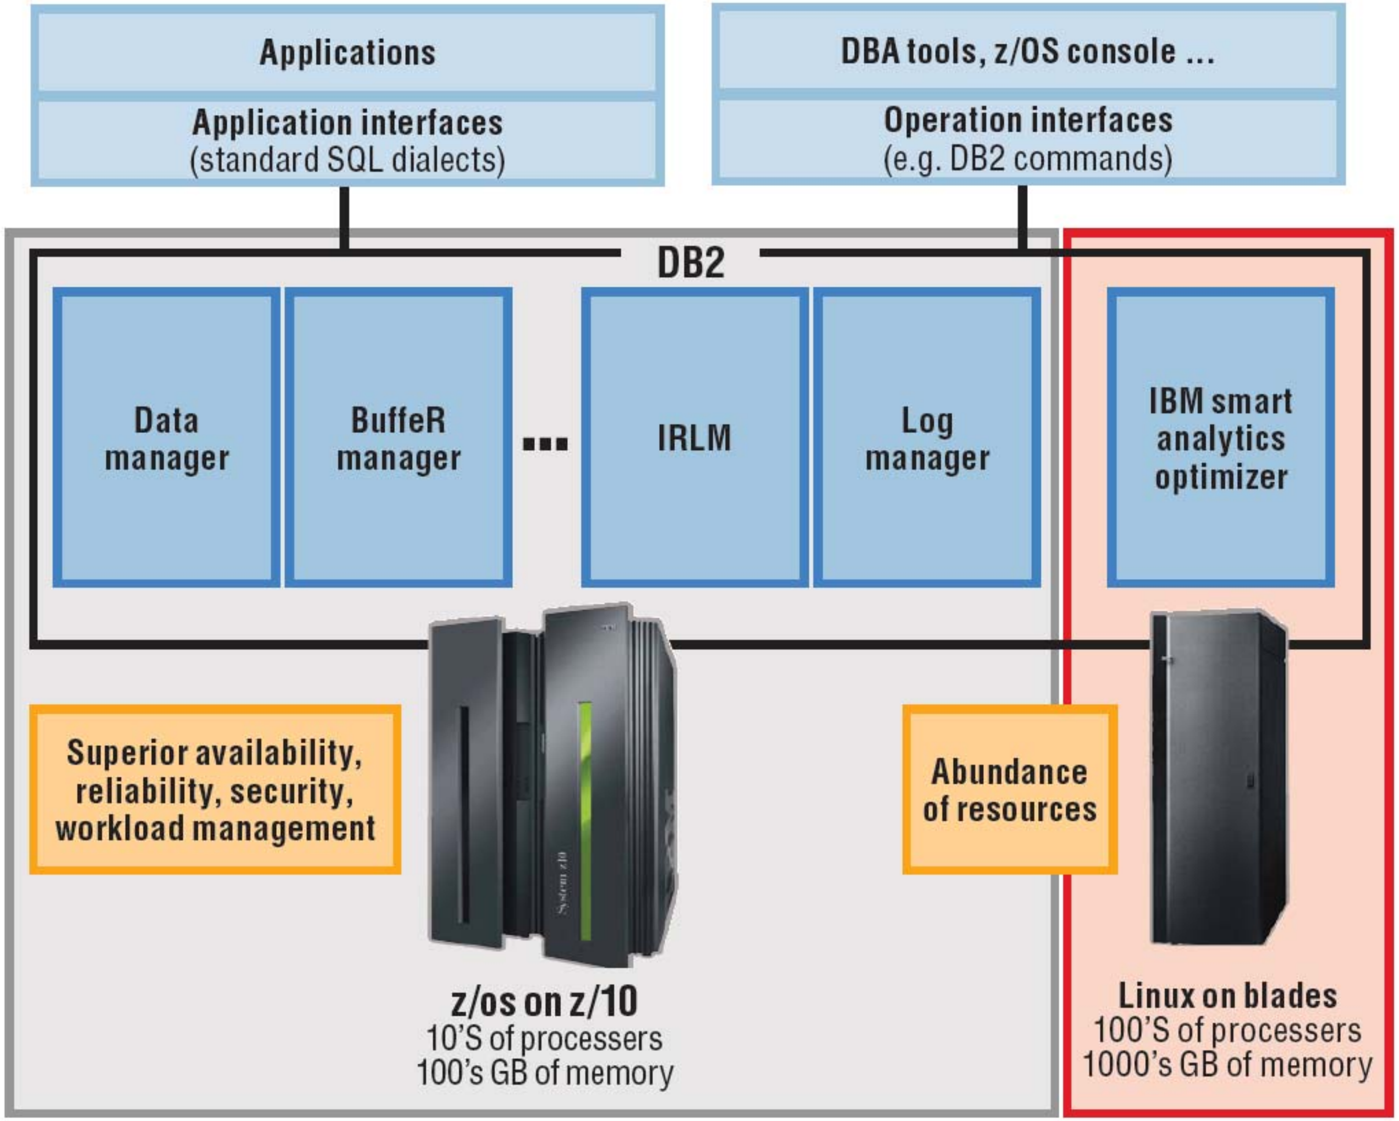

Query Execution Flow

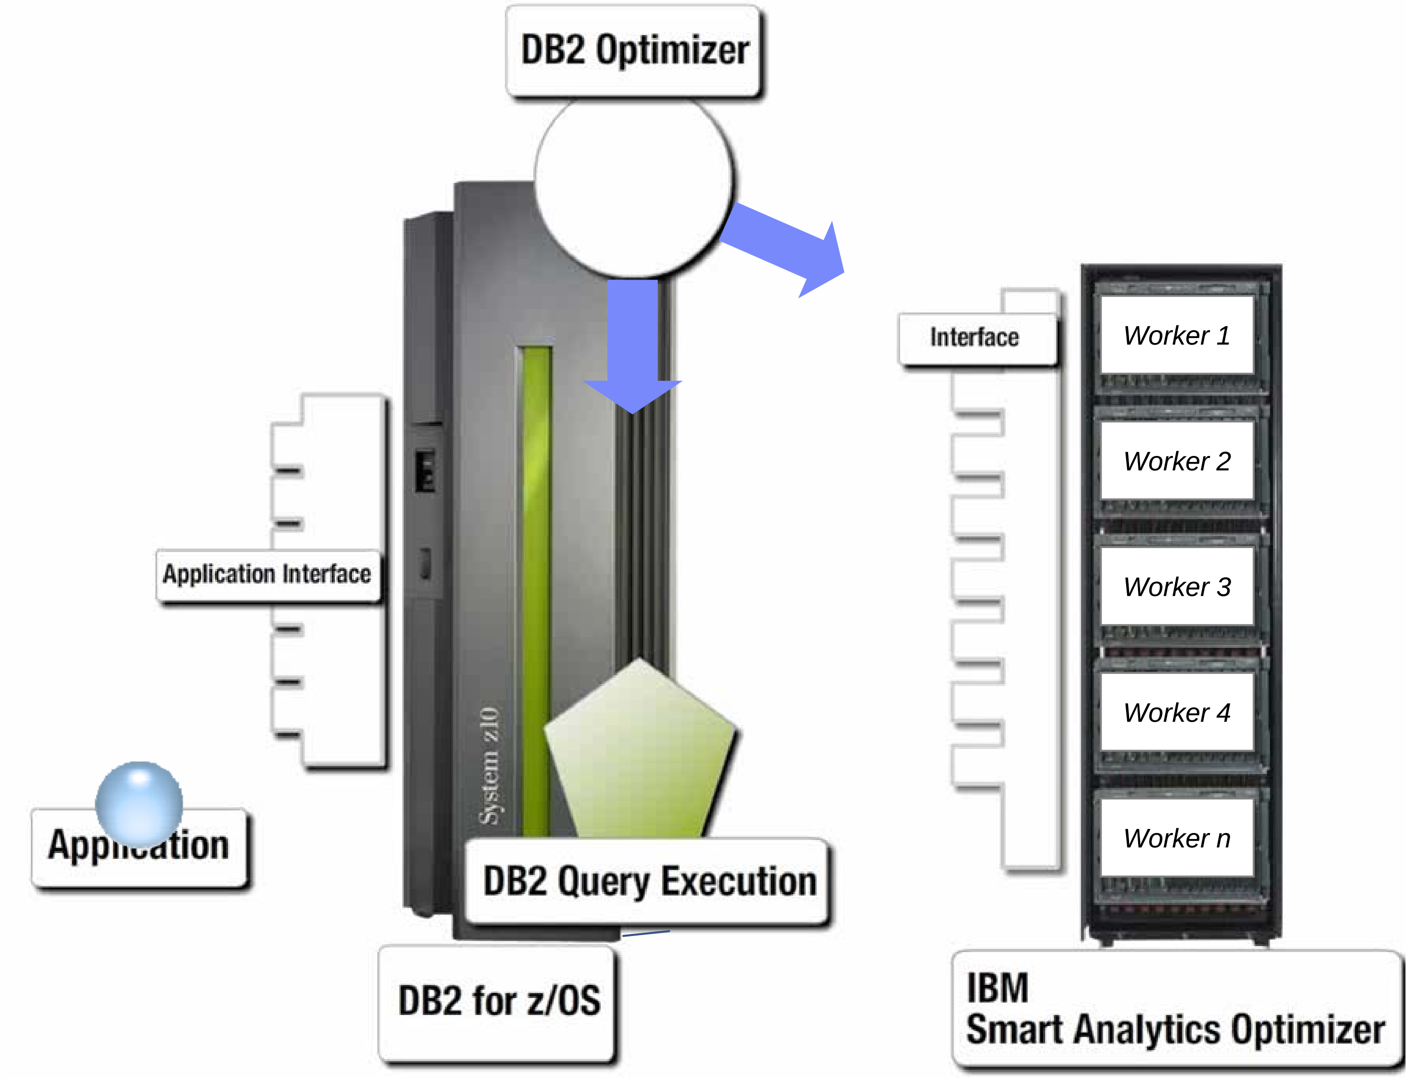

IBM Smart Analytics Optimizer - Characteristics

- 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
  - Fencing and protection of DB2 against possible accelerator failures
- Order of magnitude performance improvement
- Reducing need for tedious tuning of DB2 (MQTs, indexes, etc.)
- Appliance-like form-factor
  - Hands free operations
- Significantly improved price/performance and TCO as a combined effect of:
  - Offloading very CPU intensive operations
  - Orders of magnitude performance improvement for offloaded queries
  - Reduced DBA effort for tuning offloaded queries
- Hybrid technology
  - Enabling Dynamic DW and Operational BI
  - Preserving traditional DB2 quality of service
  - Having transactional and analytical workload being handled by DB2

Agenda

- Overview over the IBM Smart Analytics Optimizer
  - Which challenges are addressed?
  - What components are used?
- Different Data Stores for different Workloads
  - What is a Row-, Column- or Register Store?
  - Why is this relevant for query performance?
- Define what should be optimized
  - How to define what data / workload to accelerate?
  - How is data transferred and distributed?

Row oriented data store

Each row stored sequentially

- Optimized for record I/O
- Fetch and decompress entire row, every time
- Result –
  - Very efficient for transactional workloads
  - Not always efficient for analytical workloads

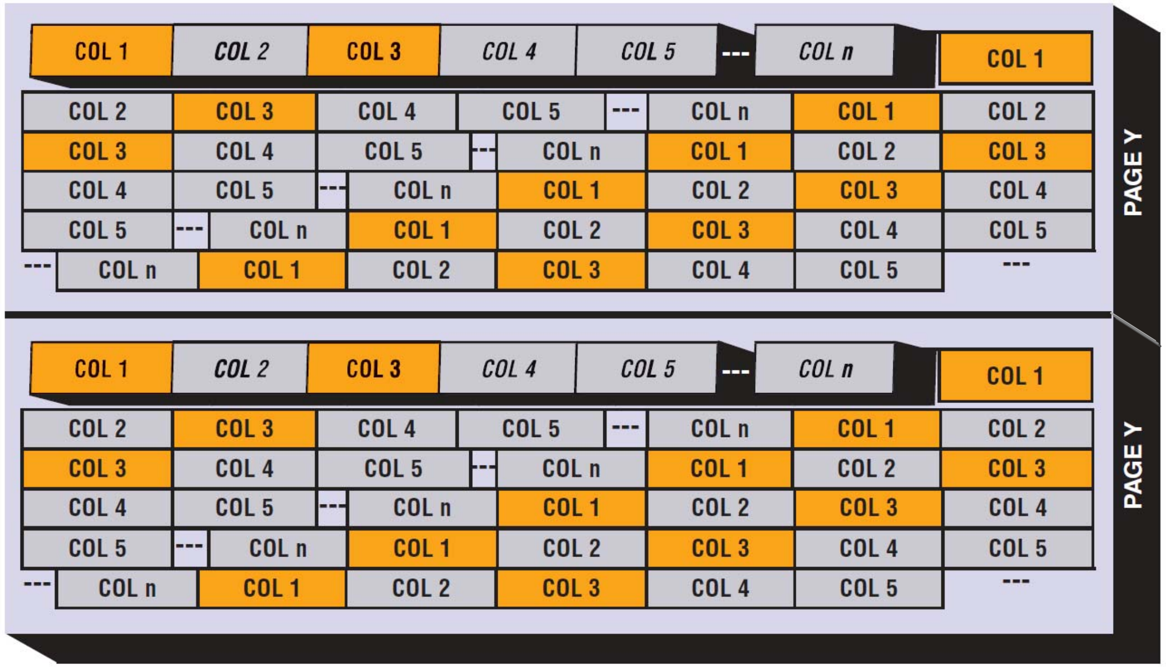

If only few columns are required the complete row is still fetched and uncompressed

Columnar data store

Data is stored sequentially by column

- Data is compressed sequentially for column:
  - •Aids sequential scan
  - •Slows random access

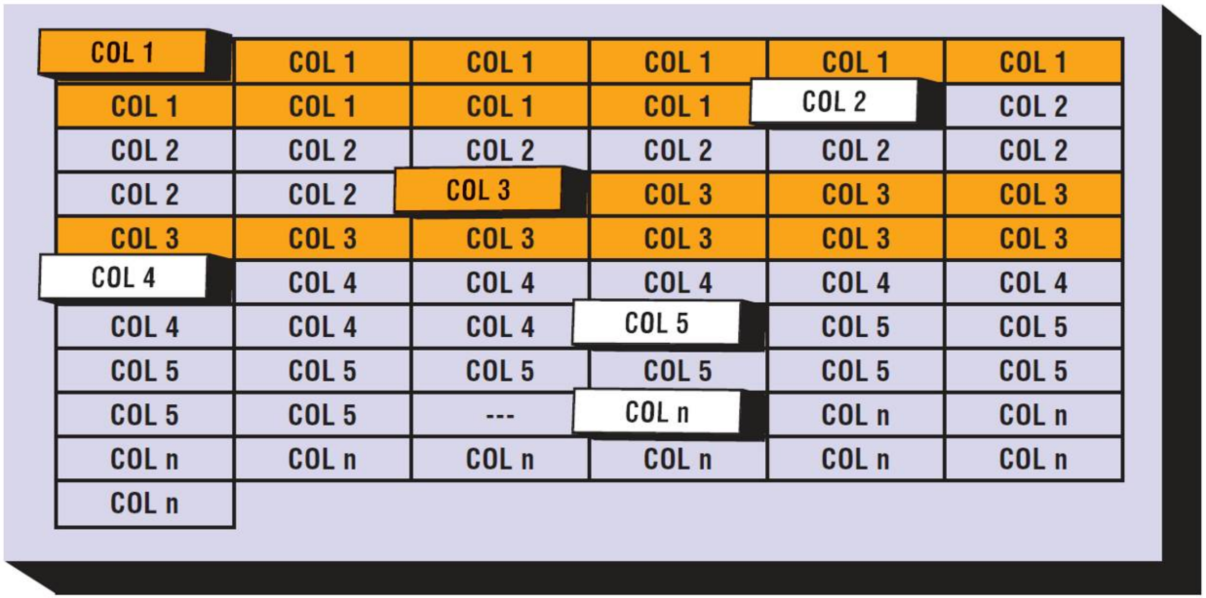

If attributes are not required for a specific query execution, they are skipped completely.

Columnar data store not optimal for all queries

Individual record joins require many extra steps

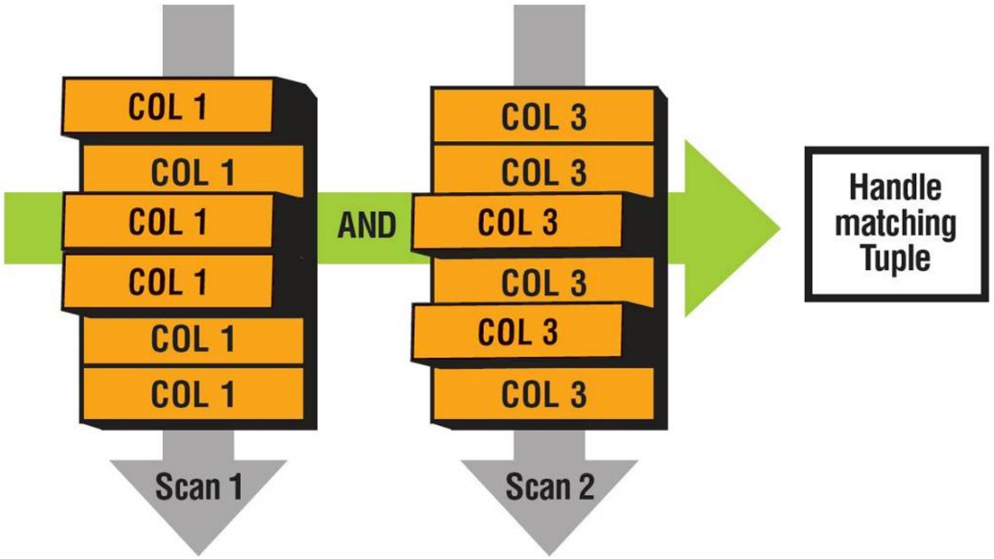

Multi-record joins difficult –

- Predicates processed separately
- Results from each column needs to be ANDed to determine a match
  - Significant additional processing

Random Record Fetch

Random record access is not performed well on pure columns stores.

Data is processed in compressed format

- 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 utilizes the full capabilities of 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.

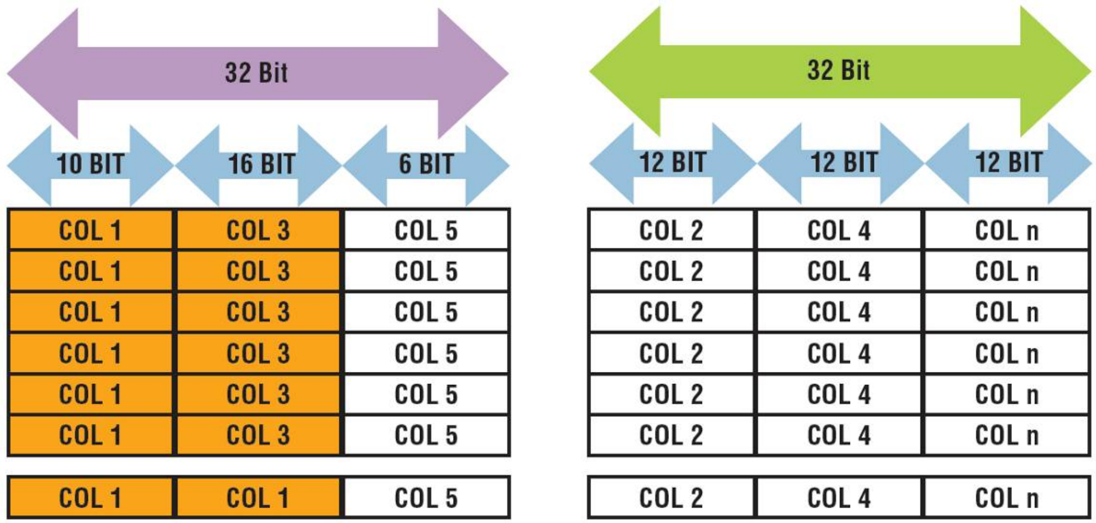

Predicate evaluation is done against compressed data!

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

Agenda

- Overview over the IBM Smart Analytics Optimizer
  - Which challenges are addressed?
  - What components are used?
- Different Data Stores for different Workloads
  - What is a Row-, Column- or Register Store?
  - Why is this relevant for query performance?
- Define what should be optimized
  - How to define what data / workload to accelerate?
  - How is data transferred and distributed?

The “Sweet – Spot” Schema

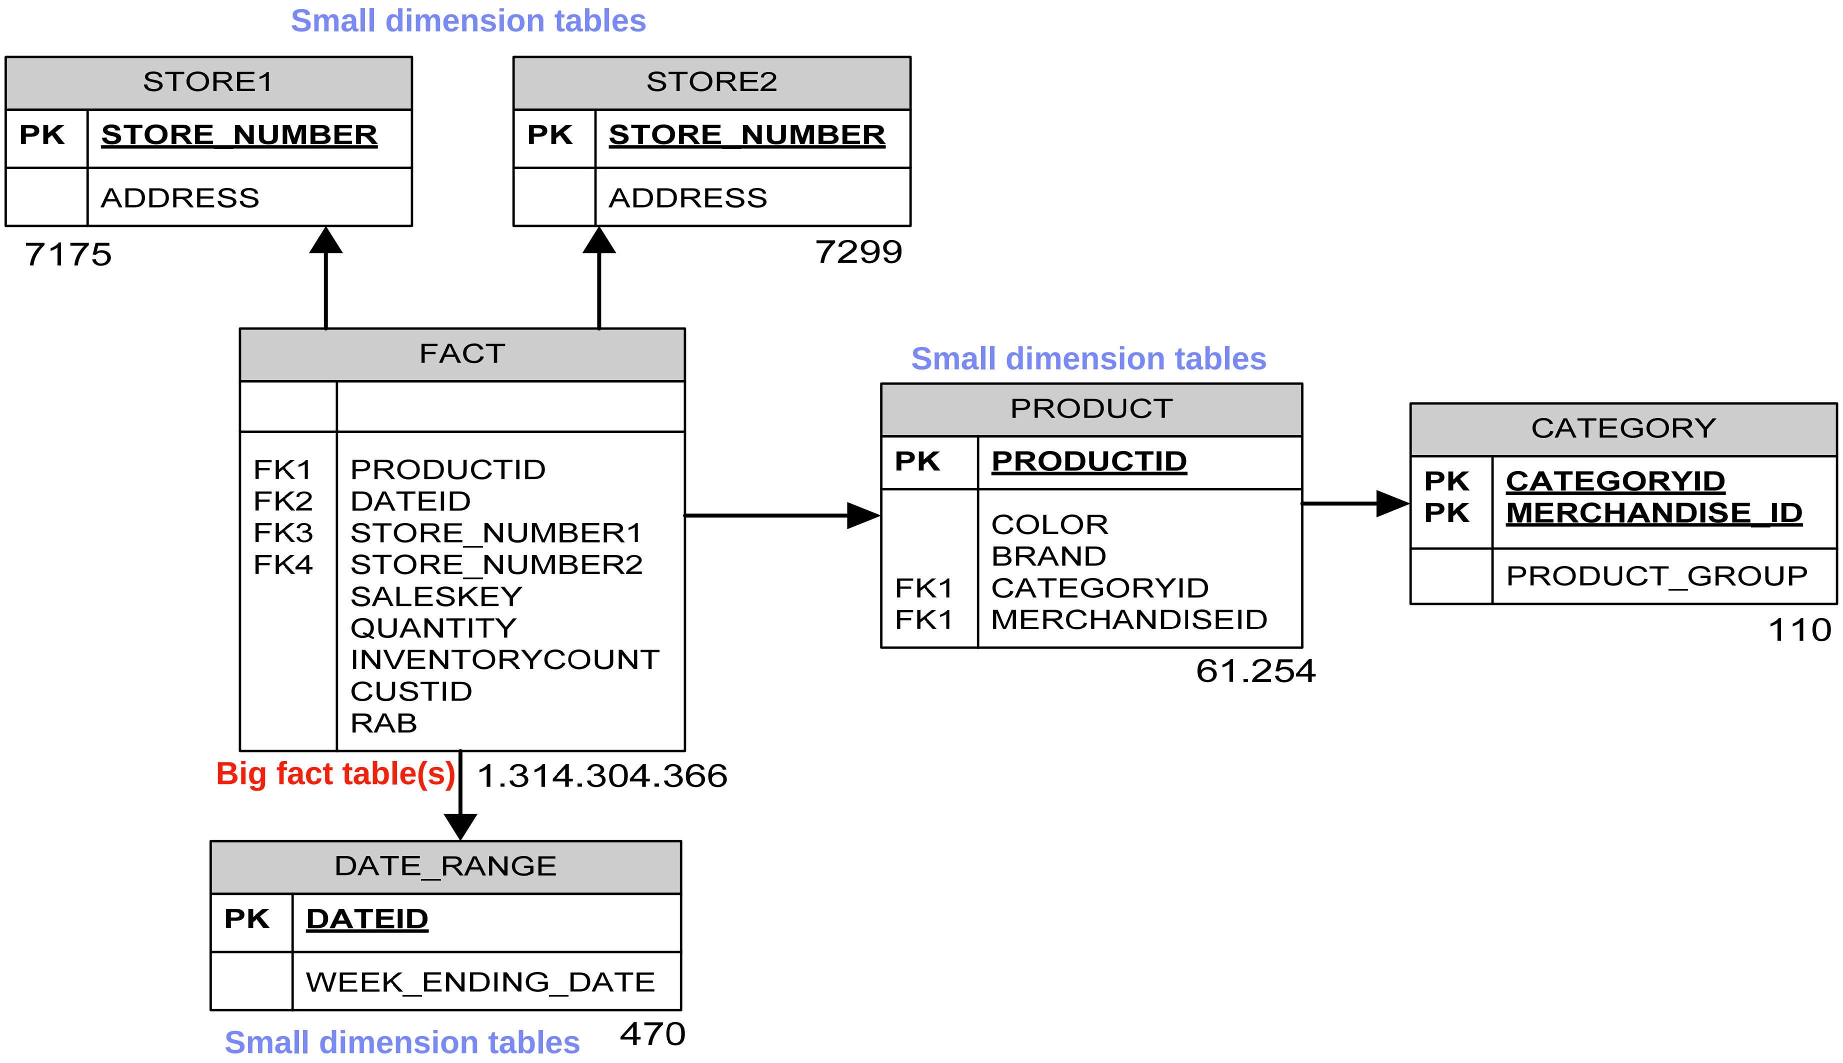

Join Strategy

- 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)

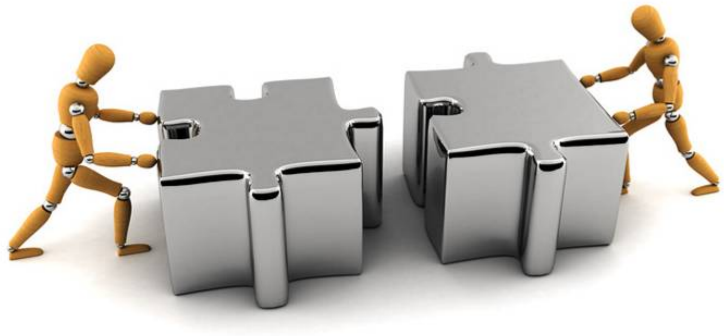

Defining which data to accelerate

- 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 Optimizer transforms the data into a compressed, scan optimized format which is kept locally in memory.

What the Smart IBM Analytics Optimizer 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
  AND P.TYPE = ´CAPIT207´
GROUP BY PRODUCT_DEPARTMENT, REGION
```

Analytics Critical for Driving Competitive Advantage

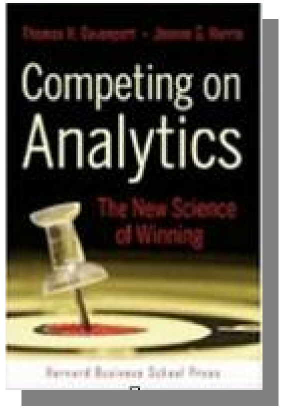

"At a time when companies in many industries offer similar products and use comparable technology, high-performance business processes are among the last remaining points of differentiation."

Tom Davenport, "Competing on Analytics"

BI/Analytics #1 investment to improve competitiveness

IBM Global CIO Study 2009

Ten Most Important Visionary Plan Elements

Interviewed CIOs could select as many as they wanted

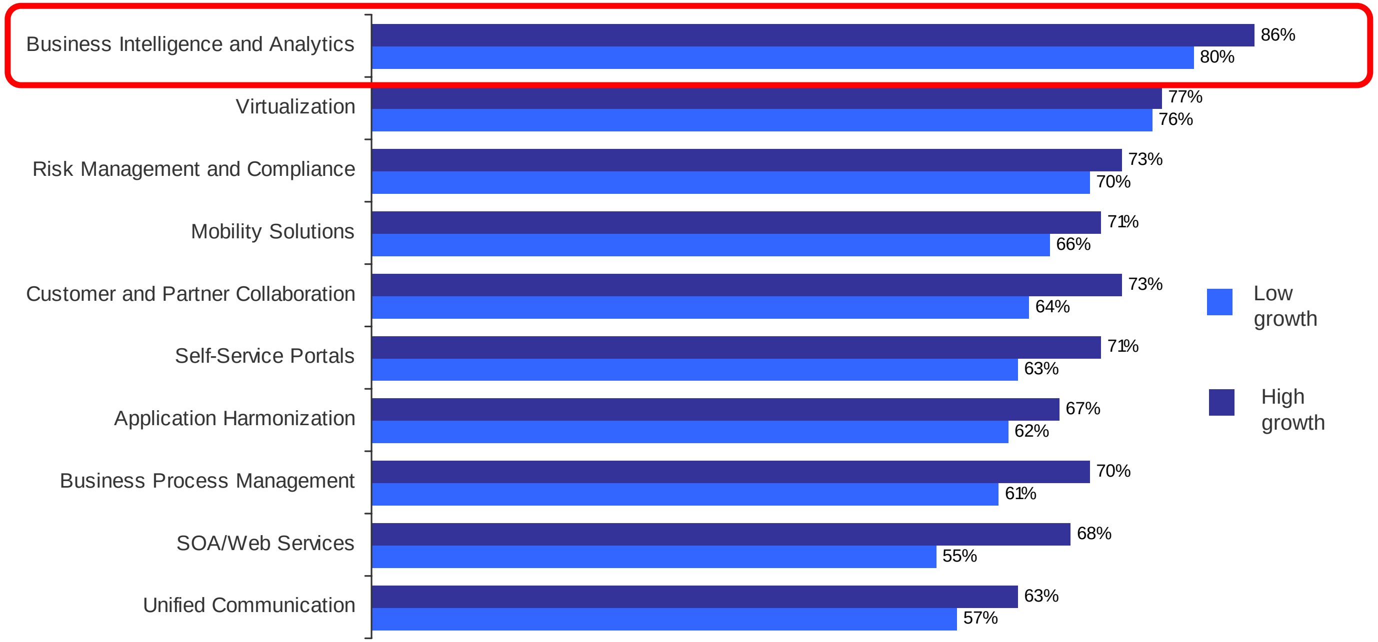

Source: IBM Global CIO Study 2009; n = 2345

Thank you