IBM Technical **Summit**

Stay ahead.

Gopi Varadarajulu

Seven Big Ideas or Habits for High Speed Analytics with DB2 BLU

Legal Notice

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including 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 results similar to those stated here.

© 2013 IBM Corporation

Agenda

- **Overview of DB2 with BLU Acceleration**
- **Use cases**
- **BLU Acceleration technology internals**
- **Early customer experiences**

© 2013 IBM Corporation

What is DB2 with BLU Acceleration?

- **Large order of magnitude benefits**
  - Performance
  - Storage savings
  - Time to value

- **New technology in DB2 for analytic queries**
  - CPU-optimized unique runtime handling
  - Unique encoding for speed and compression
  - Unique memory management
  - Columnar storage, vector processing
  - Built directly into the DB2 kernel

- **Revolution or evolution**
  - BLU tables coexists with traditional row tables - in same schema, storage, and memory
  - Query any combination of row or BLU tables
  - Easy conversion of tables to BLU tables
    - Change everything, or change incrementally

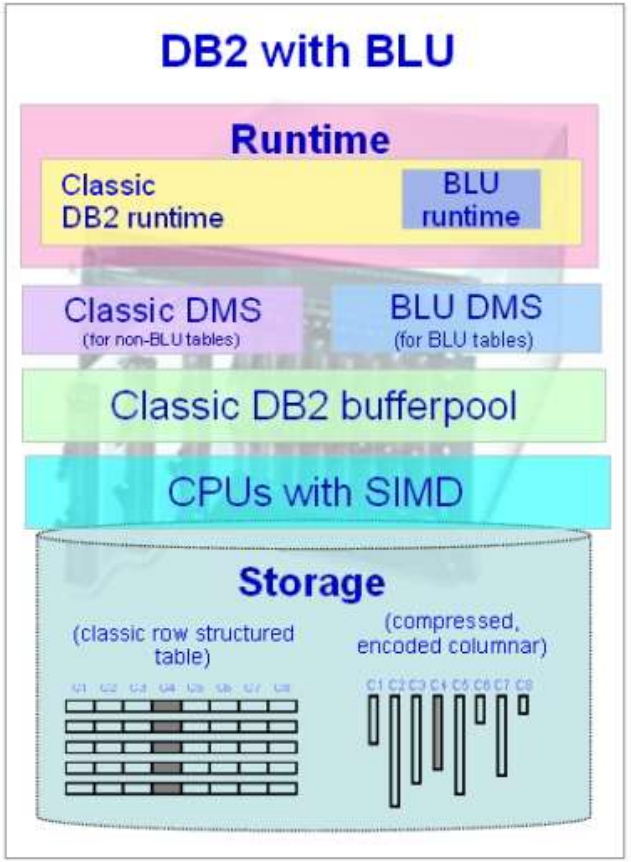

© 2013 IBM Corporation

How fast is it?
Results from the DB2 10.5 Beta

Customer	Speedup over DB2 10.1
Large Financial Services Company	46.8x
Global ISV Mart Workload	37.4x
Analytics Reporting Vendor	13.0x
Global Retailer	6.1x
Large European Bank	5.6x

10x-25x
improvement is common

BNSF RAILWAY

"It was amazing to see the faster query times compared to the performance results with our row-organized tables. ***The performance of four of our queries improved by over 100-fold! The best outcome was a query that finished 137x faster by using BLU Acceleration."***
- Kent Collins, Database Solutions Architect, BNSF Railway

© 2013 IBM Corporation

Storage Savings

- **Multiple examples of data requiring substantially less storage**
  - 95% smaller than uncompressed data size
  - Fewer objects required – no storage required for indexes, aggregates, etc
- **Multiple compression techniques**
  - Processing takes place on compressed data
- **Compression algorithm adapts to the data**

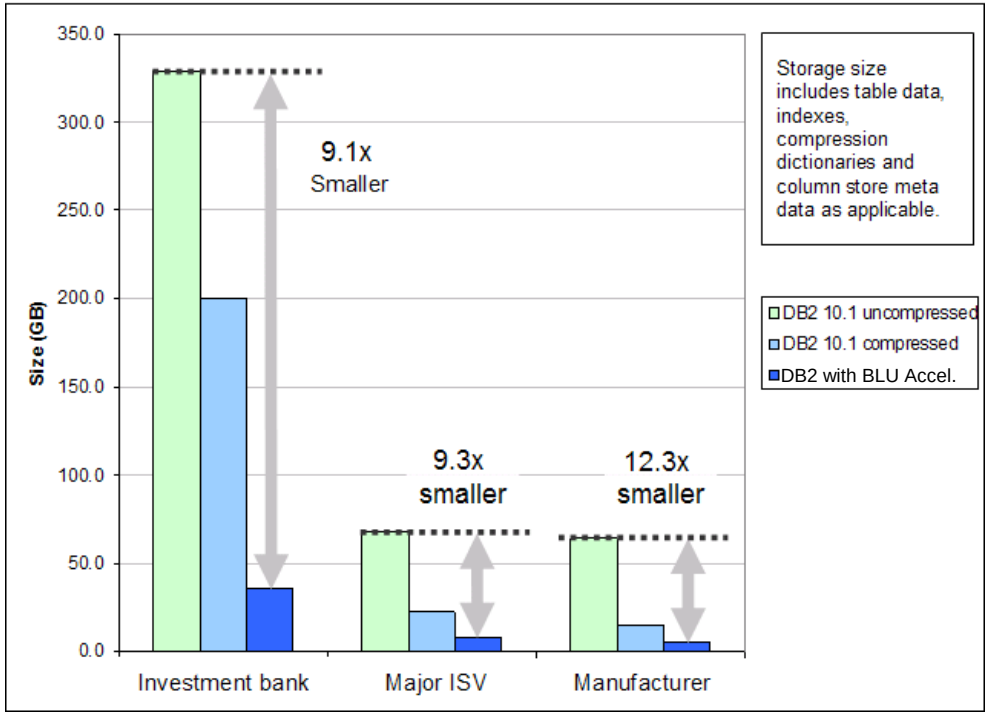

© 2013 IBM Corporation

Seamless Integration into DB2

- **Built seamlessly into DB2 – integration and coexistence**
  - Column-organized tables can coexist with existing, traditional, tables
    - Same schema, same storage, same memory
  - Integrated tooling support
    - Optim Query Workload Tuner recommends BLU Acceleration deployments
- **Same SQL, language interfaces, administration**
  - Column-organized tables or combinations of column-organized and row-organized tables can be accessed within the same SQL statement
- **Dramatic simplification – Just "Load and Go"**
  - Faster deployment
    - Fewer database objects required to achieve same outcome
  - Requires less ongoing management due to it's optimized query processing and fewer database objects required
  - Simple migration
    - Conversion from traditional row table to BLU Acceleration is easy
    - DB2 Workload Manager identifies workloads to tune
    - Optim Query Workload Tuner recommends BLU Acceleration table transformations
    - Users only notice speed up; DBA's only notice less work!
  - Management of single server solutions less expensive than clustered solutions

© 2013 IBM Corporation

Analytic Database Management Complexity

DATABASE

MICROSOFT
SYBASE
TERADATA
ORACLE

Database Design and Tuning

- Decide on partition strategies
- Select Compression Strategy
- Create Table
- Load data
- Create Auxiliary Performance Structures
  - Materialized views
  - Create indexes
    - B+ indexes
    - Bitmap indexes
- Tune memory
- Tune I/O
- Add Optimizer hints
- Statistics collection

Repeat

© 2013 IBM Corporation

Simple to Deploy and Operate

- **Operations**
  - Simply **Load and Go**
  - Installation to business value in ~2 days
  - Ease of evaluation and performs as advertised
- **BI developers and DBAs – faster delivery**
  - No configuration or physical modeling
  - No indexes or tuning – out of the box performance
  - Data Architects/DBA focus on business value, not physical design
- **ETL developers**
  - No aggregate tables needed – simpler ETL logic
  - Faster load and transformation times
- **Business analysts**
  - Train of thought analysis – 5x to 100x faster
  - True ad-hoc queries – no tuning, no indexes
  - Ask complex queries against large datasets

© 2013 IBM Corporation

BLU Acceleration Use Cases

© 2013 IBM Corporation

Analytics Data Mart

From Transactional Database

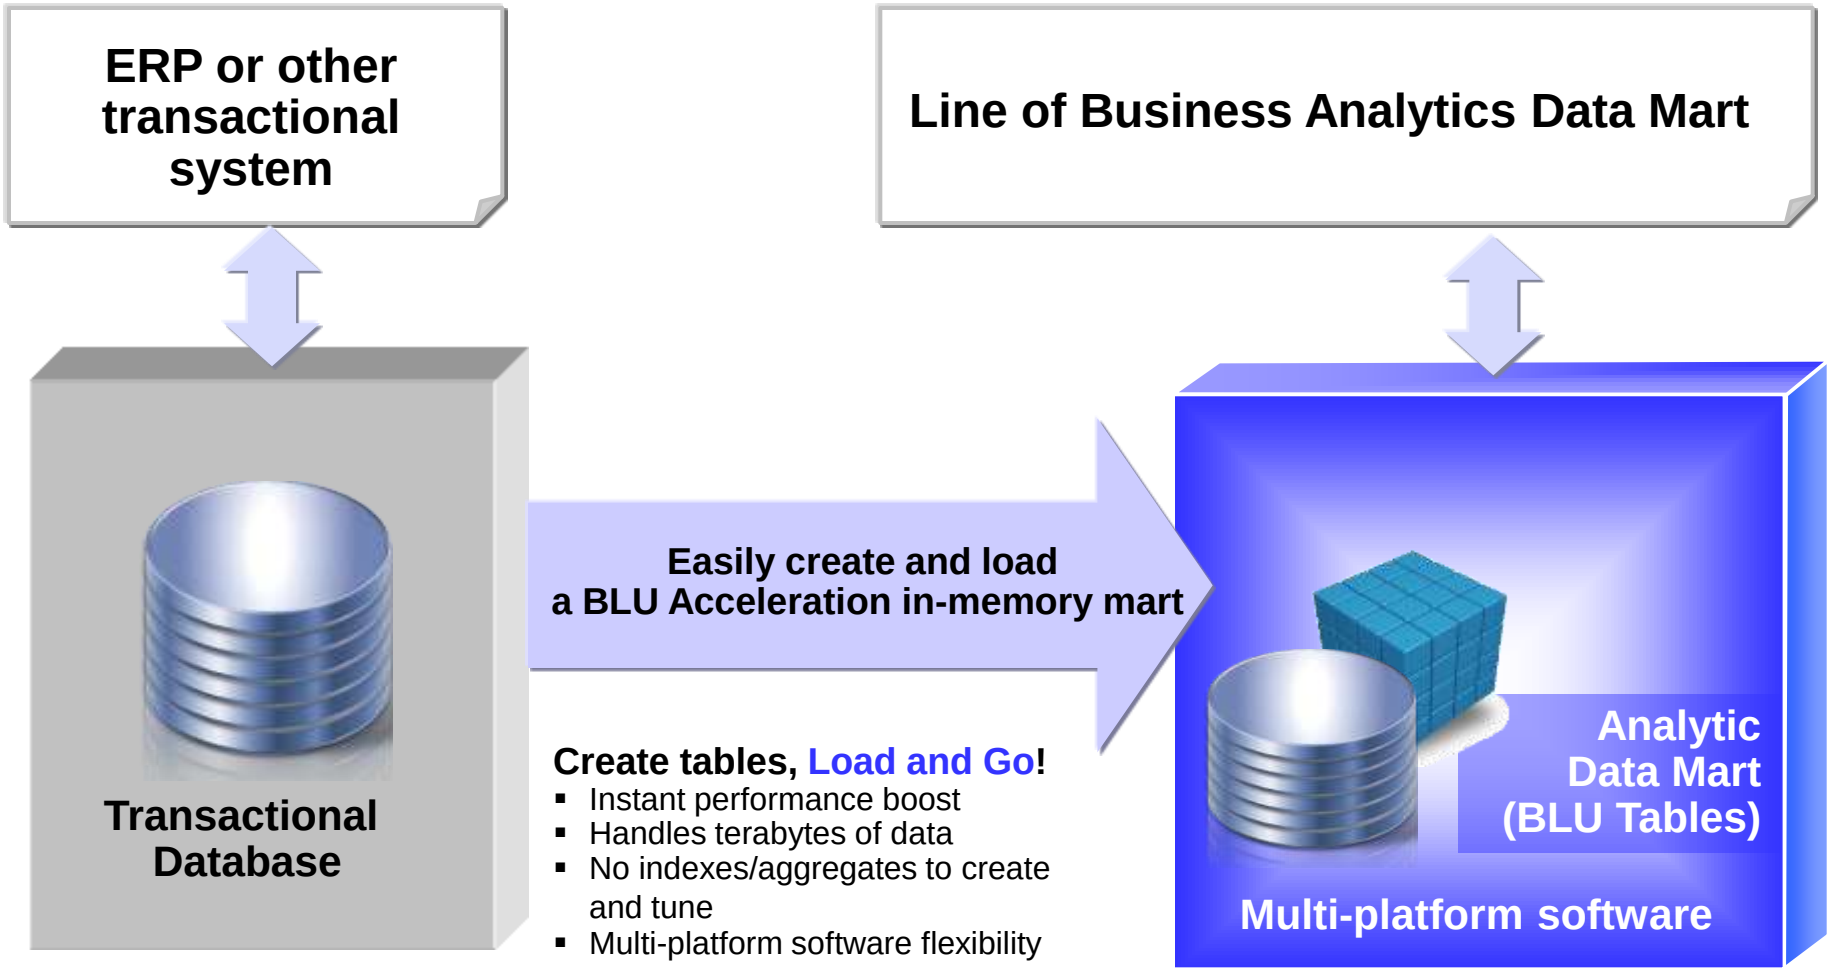

© 2013 IBM Corporation

Use Case – Enterprise Data Warehouse Offload
Data Mart Acceleration

EDW Application

OLAP Application

Cognos BI
with BLU Acceleration

Poor performing Oracle or Teradata warehouse

Easily create and load a BLU Acceleration in-memory mart

Create tables, Load and Go!
- Instant performance boost
- Handles terabytes of data
- No indexes/aggregates to create and tune
- Multi-platform software flexibility

Cognos BI
with BLU Acceleration

Analytic Data Mart (BLU Tables)

Multi-platform software

© 2013 IBM Corporation

DB2 with BLU Acceleration Technology Internals

© 2013 IBM Corporation

BLU Acceleration is Deeply Integrated With the DB2 Kernel

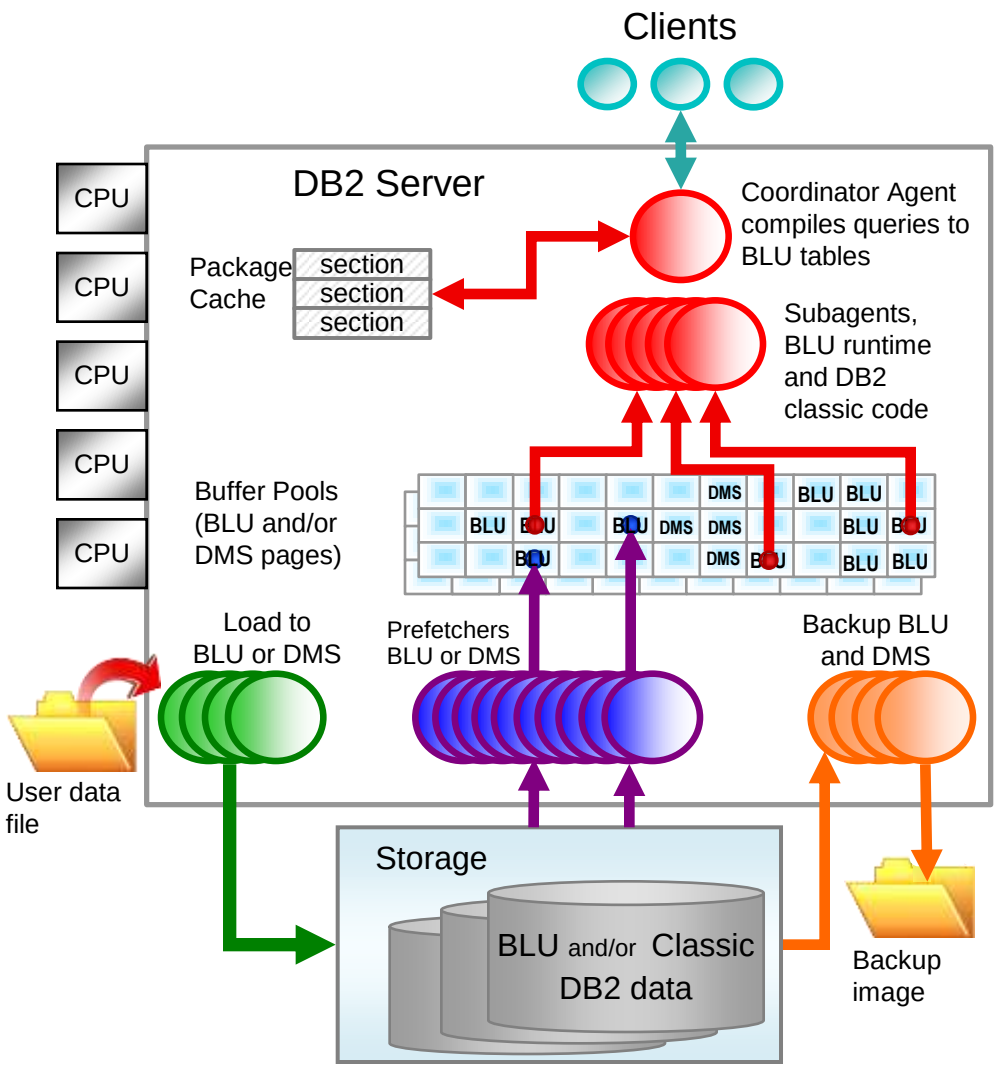

- **Client/Server**
  - BLU Acceleration uses DB2 client server infrastructure. Complete transparency to the application
- **Compiler**
  - BLU Acceleration uses the DB2 compiler to accept SQL, parse, perform semantic checking, and package creation
- **Process model – BLU Acceleration uses**
  - DB2 subagents
  - Prefetchers
  - TCB and Packed Descriptor for metadata
- **Memory**
  - BLU Acceleration uses DB2 bufferpool for storage allocation and caching
  - BLU Acceleration uses DB2 sort heap and package cache
  - OSS memory allocation for private work areas
- **Storage**
  - BLU Acceleration uses normal DB2 table spaces for storage allocations
  - Page sizes: 4K-32K
- **Utilities**
  - `LOAD, BACKUP, RESTORE, EXPORT, SNAPSHOT, db2top, db2pd, etc.`

© 2013 IBM Corporation

The Seven Big Ideas of DB2 with BLU Acceleration

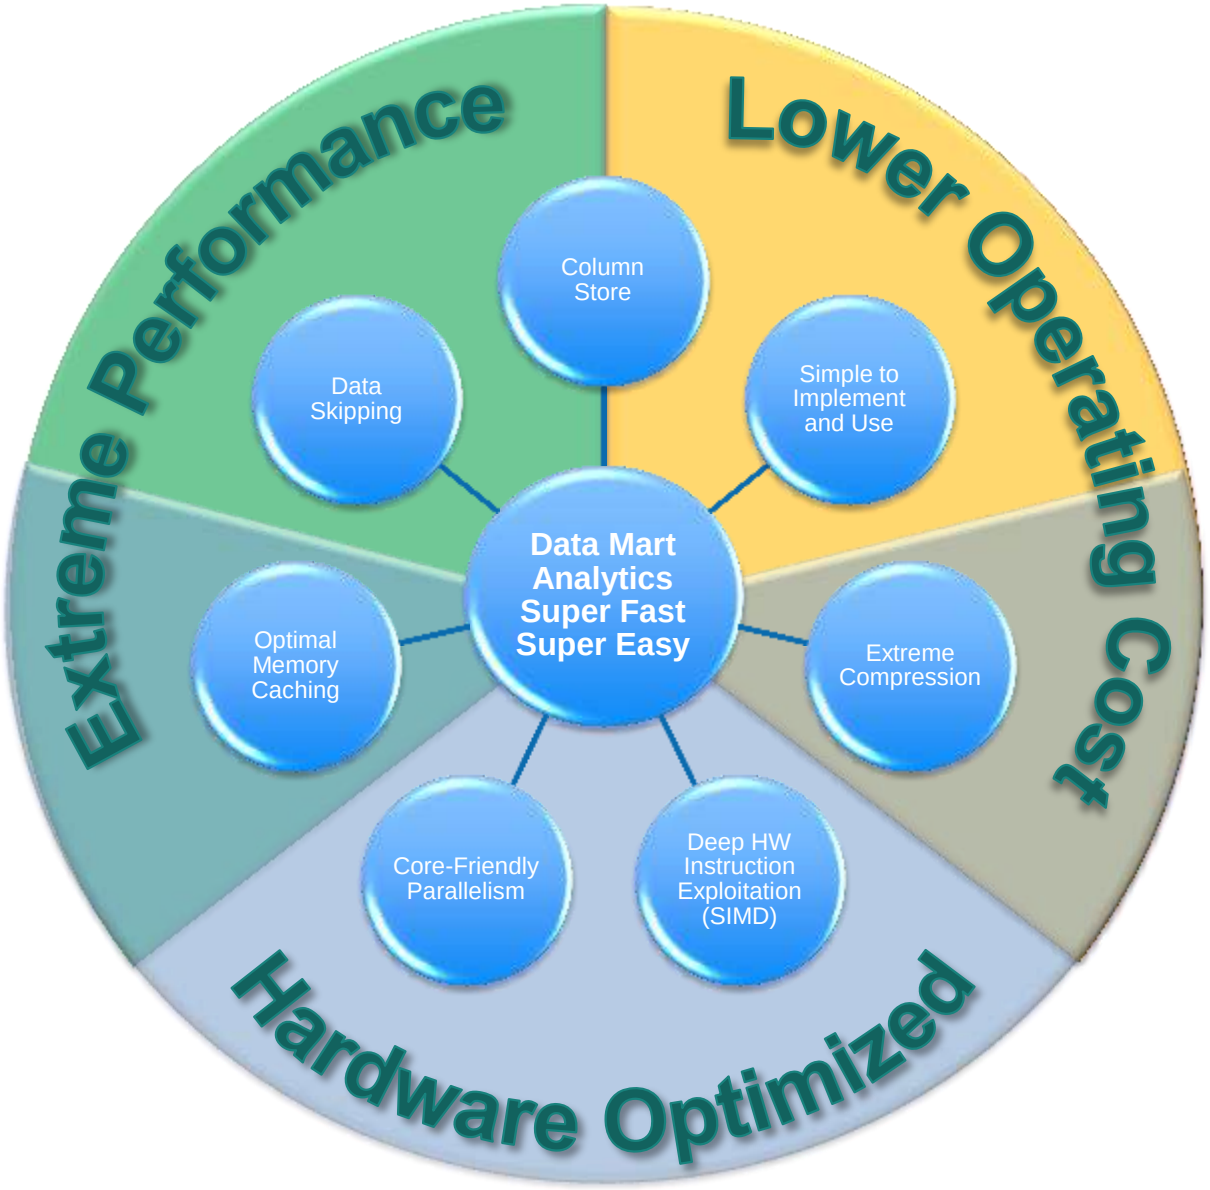

© 2013 IBM Corporation

7 Big Ideas: 1 Simple to Implement and Use

- **LOAD and then… run queries**
  - No indexes
  - No `REORG` (it's automated)
  - No `RUNSTATS` (it's automated)
  - No MDC or MQTs or Materialized Views
  - No partitioning
  - No statistical views
  - No optimizer hints

- **It is just DB2!**
  - Same SQL, language interfaces, administration
  - Reuse DB2 process model, storage, utilities

"The BLU Acceleration technology has some obvious benefits: It makes our analytical **queries run 4-15x faster** and **decreases the size of our tables by a factor of 10x**. But it's when I think about **all the things I don't have to do with BLU**, it made me appreciate the technology even more: **no tuning, no partitioning, no indexes, no aggregates**."
-Andrew Juarez, Lead SAP Basis and DBA

© 2013 IBM Corporation

7 Big Ideas: 1 Simple to Implement and Use

- **One setting optimized the system for BLU Acceleration**
  - Set `DB2_WORKLOAD=ANALYTICS`
  - Informs DB2 that the database will be used for analytic workloads

- **Automatically configures DB2 for optimal analytics performance**
  - Makes column-organized tables the default table type
  - Enables automatic workload management
  - Enables automatic space reclaim
  - Page and extent size configured for analytics
  - Memory for caching, sorting and hashing, utilities are automatically initialized based on the server size and available RAM

- **Simple Table Creation**
  - If `DB2_WORKLOAD=ANALYTICS,` tables will be created column organized automatically
  - For mixed table types can define tables as `ORGANIZE BY COLUMN` or `ROW`
  - Compression is always on – No options

- **Easily convert tables from row-organized to column-organized**
  - `db2convert` utility

© 2013 IBM Corporation

7 Big Ideas: 2 Compute Friendly Encoding and Compression

- **Massive compression with approximate Huffman encoding**
  - More frequent the value, the fewer bits it takes
- **Register-friendly encoding dramatically improves efficiency**
  - Encoded values packed into bits matching the register width of the CPU
  - Fewer I/Os, better memory utilization, fewer CPU cycles to process

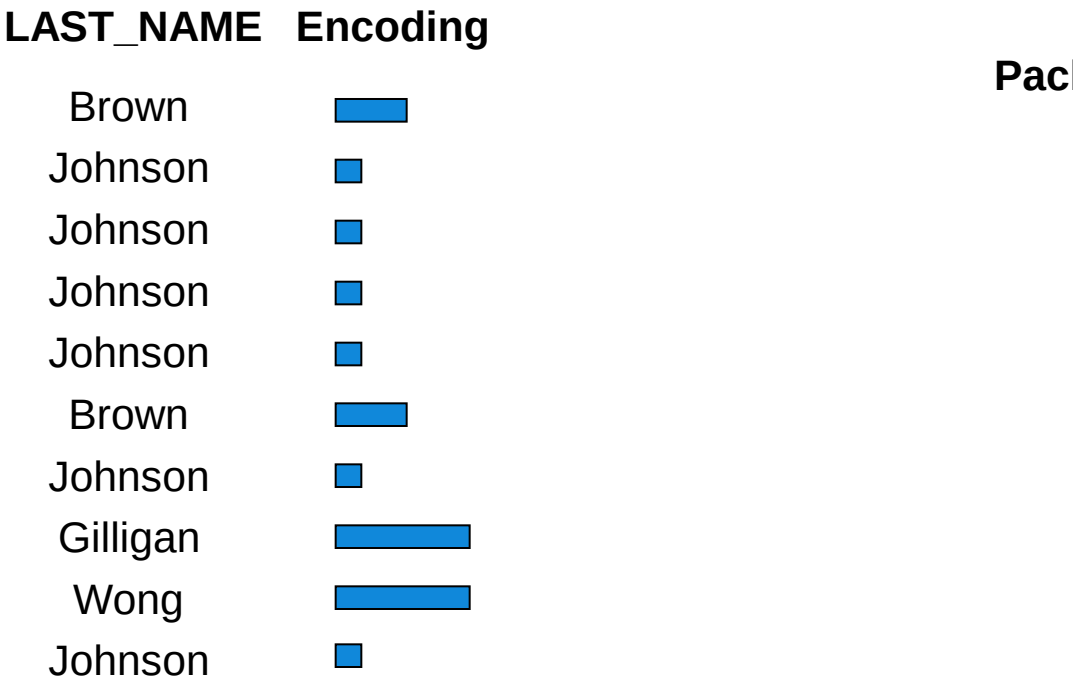

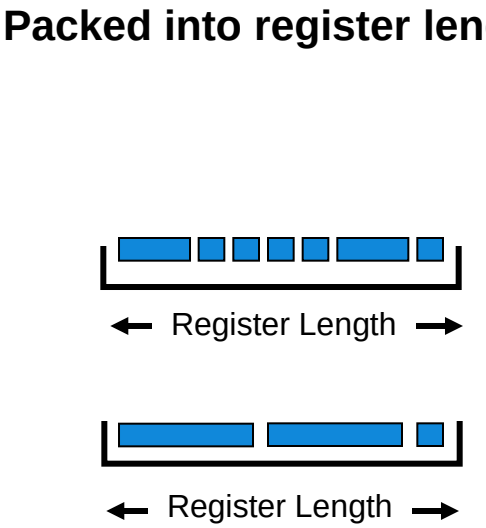

© 2013 IBM Corporation

7 Big Ideas: 2 Data Remains Compressed During Evaluation

- **Encoded values do not need to be decompressed during evaluation**
  - Predicates (=, <, >, >=, <=, Between, etc), joins, aggregations and more work directly on encoded values

```
SELECT COUNT(*) FROM T1 WHERE LAST_NAME = 'Johnson'
```

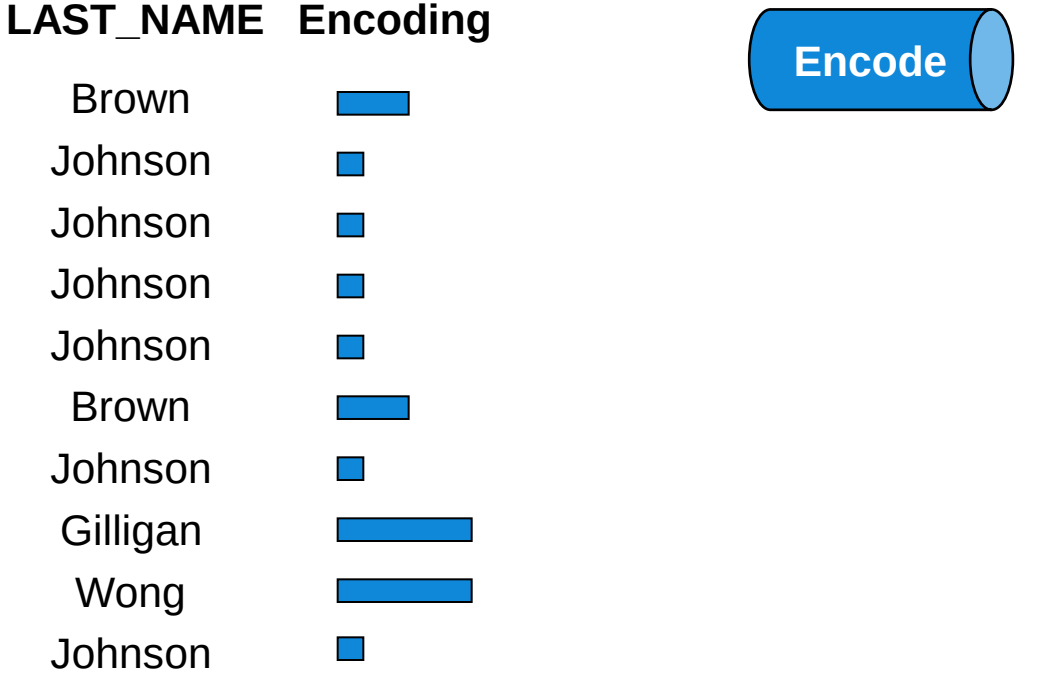

© 2013 IBM Corporation

7 Big Ideas: 3 Multiply the Power of the CPU

- **Performance increase with Single Instruction Multiple Data (SIMD)**
- **Using hardware instructions, DB2 with BLU Acceleration can apply a single instruction to many data elements simultaneously**
  - Predicate evaluation, joins, grouping, arithmetic

intel

*"Intel is excited to see a **25x improvement in query processing performance** using DB2 10.5 with BLU acceleration over DB2 10.1. To achieve these amazing gains, IBM has taken advantage of the Advanced Vector Extensions (AVX) instruction set on Intel® Xeon® processor E5-based systems."*

- Pauline Nist, GM, Enterprise Software Alliances, Datacenter & Connected Systems Group

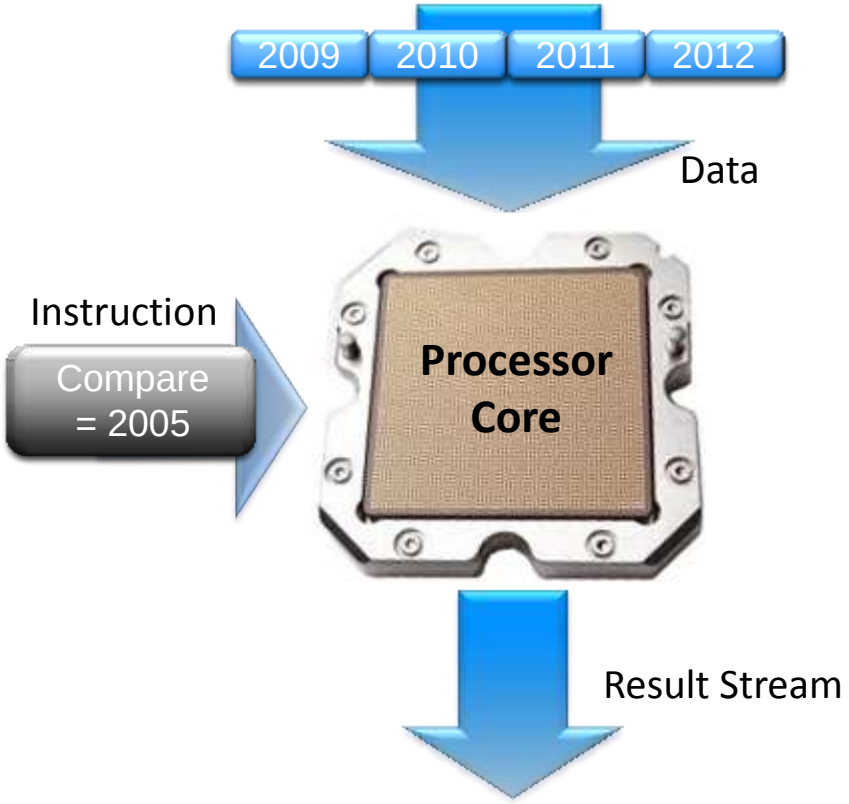

© 2013 IBM Corporation

7 Big Ideas: 4 Core-Friendly Parallelism

- **Careful attention to physical attributes of the server**
  - Queries on BLU Acceleration tables automatically parallelized
- **Maximizes CPU cache, cacheline efficiency**

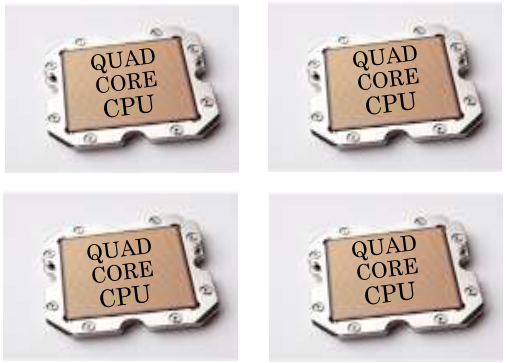

BNSF RAILWAY

"During our testing, we couldn't help but notice that ***DB2 10.5 with BLU Acceleration is excellent at utilizing our hardware resources****. The core-friendly parallelism that IBM talks about was clearly evident and* ***I didn't even have to partition the data across multiple servers****."*

- Kent Collins, Database Solutions Architect, BNSF Railway

© 2013 IBM Corporation

7 Big Ideas: 5 Column Store

- **Minimal I/O**
  - Only perform I/O on the columns and values that match query
  - As queries progresses through a pipeline the working set of pages is reduced
- **Work performed directly on columns**
  - Predicates, joins, scans, etc. all work on individual columns
  - Rows are not materialized until absolutely necessary to build result set
- **Improved memory density**
  - Columnar data kept compressed in memory
- **Extreme compression**
  - Packing more data values into very small amount of memory or disk
- **Cache efficiency**
  - Data packed into cache friendly structures

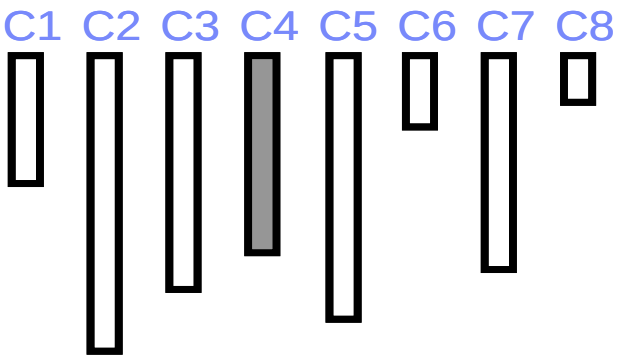

© 2013 IBM Corporation

7 Big Ideas: 6 Scan-Friendly Memory Caching

- **New algorithms cache in RAM effectively**
- **High percent of interesting data fits in memory**
  - We leave the interesting data in memory with the new algorithms
- **Data can be larger than RAM**
  - No need to ensure all data fits in memory
  - Optimization for in memory and I/O efficiency

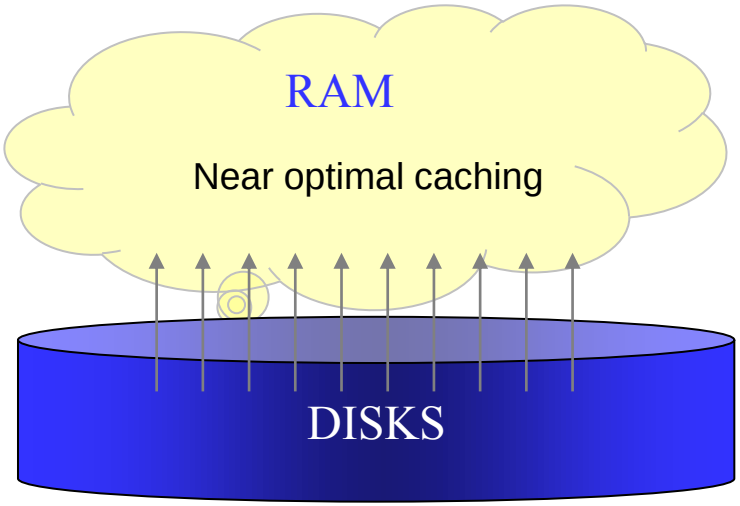

© 2013 IBM Corporation

7 Big Ideas: 7 Data skipping

- **Automatic detection of large sections of data that do not qualify for a query and can be ignored**
- **Order of magnitude savings in all of I/O, RAM, and CPU**
- **No DBA action to define or use – truly invisible**
  - Persistent storage of min and max values for sections of data values

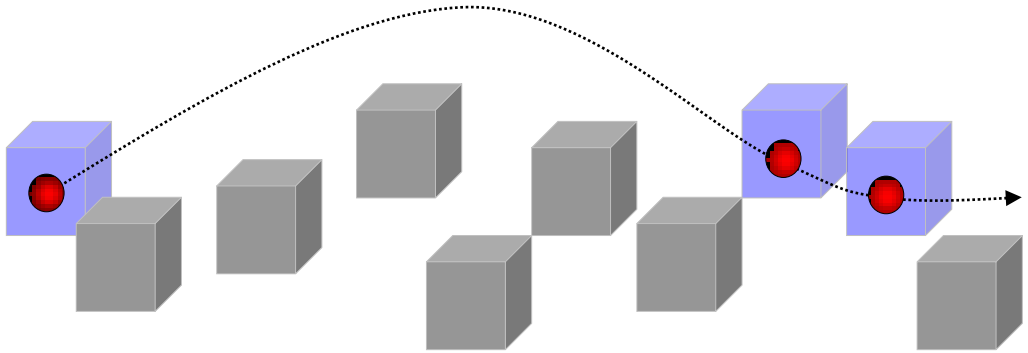

© 2013 IBM Corporation

Optimize the Entire Hardware Stack

In-Memory Optimized

- **Memory latency optimized for**
  - Scans
  - Joins
  - Aggregation
- **More useful data in memory**
  - Data stays compressed
  - Scan friendly caching
- **Less to put in memory**
  - Columnar access
  - Late materialization
  - Data skipping

CPU Optimized

- **CPU acceleration**
  - SIMD processing for
    - Scans
    - Joins
    - Grouping
    - Arithmetic
- **Keeping the CPUs busy**
  - Core friendly parallelism
- **Less CPU processing**
  - Operate on compressed data
  - Late materialization
  - Data skipping

I/O Optimized

- **Less to read**
  - Columnar I/O
  - Data skipping
  - Late materialization
- **Read less often**
  - Scan friendly caching
- **Efficient I/O**
  - Specialized columnar prefetching algorithm

© 2013 IBM Corporation

7 Big Ideas: How DB2 with BLU Acceleration Helps

~Sub second 10TB query – An Optimistic Illustration

- **The system – 32 cores, 10TB table with 100 columns, 10 years of data**
- **The query: `SELECT COUNT(*) from MYTABLE where YEAR = '2010'`**
- **The optimistic result: sub second 10TB query! Each CPU core examines the equivalent of just 8MB of data**

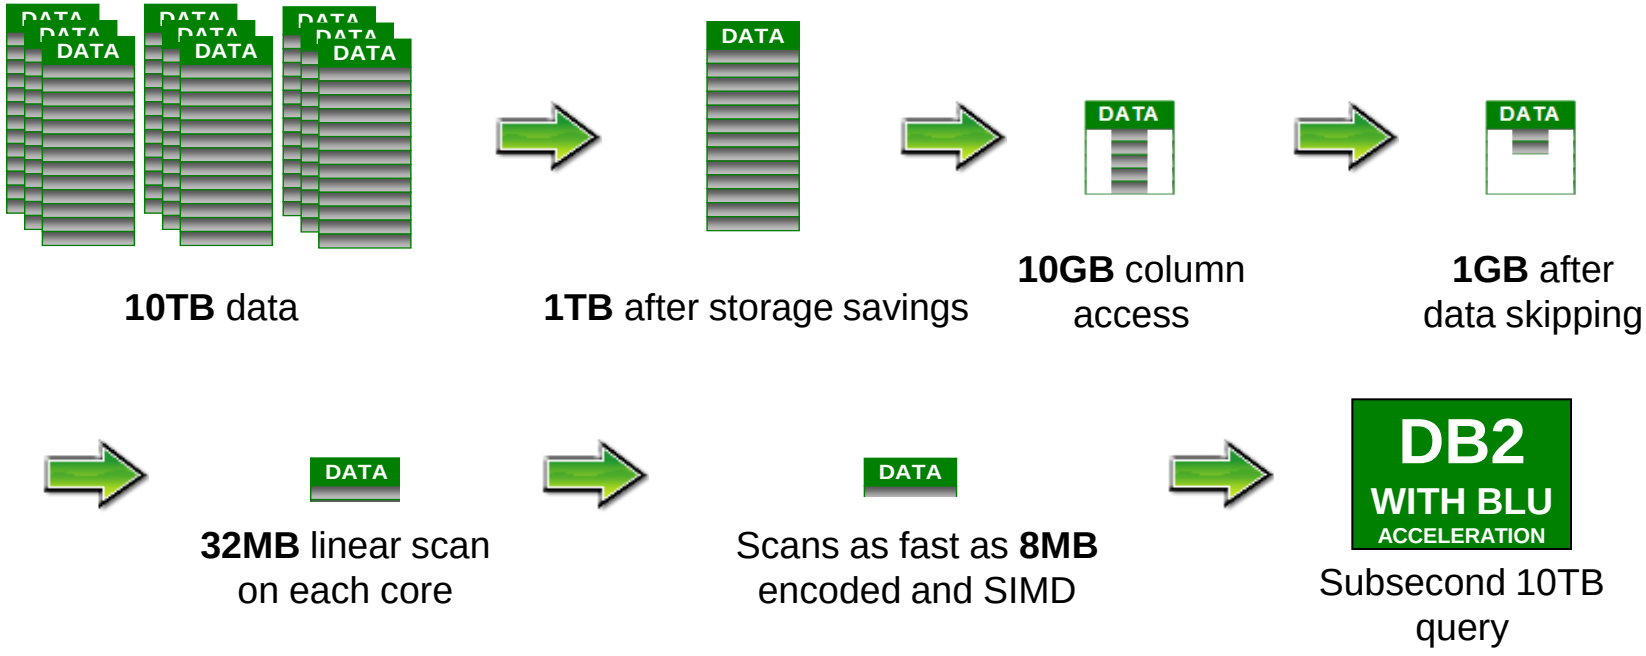

10TB data

1TB after storage savings

10GB column access

1GB after data skipping

32MB linear scan on each core

Scans as fast as **8MB** encoded and SIMD

Subsecond 10TB query

© 2013 IBM Corporation

Unlimited Concurrency with "Automatic WLM"

- **DB2 10.5 has built-in and automated query resource consumption control**
- **Every additional query that runs naturally consumes more memory, locks, CPU, and memory bandwidth. In other database products more queries means more contention**
- **DB2 10.5 automatically allows a high level of concurrent queries to be submitted, but limits the number that consume resources at any point in time**
- **Enabled automatically when `DB2_WORKLOAD=ANALYTICS`**

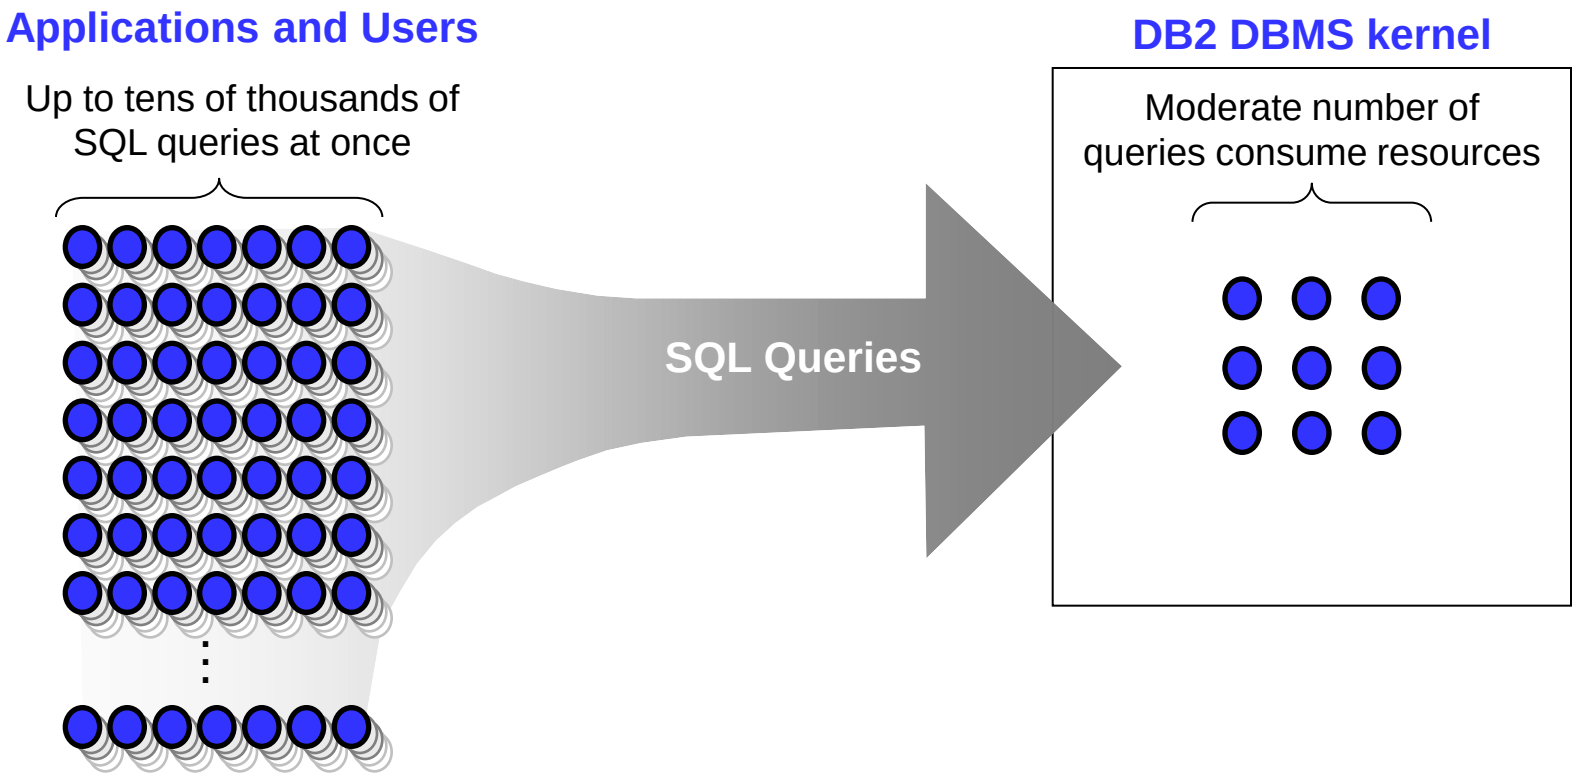

© 2013 IBM Corporation

Informational Uniqueness

- **DB2 10.5 introduces informational uniqueness constraints**
  - Enforced uniqueness remains the default
  - Informational (i.e., `NOT ENFORCED`) constraints do not enforce uniqueness
  - Valuable when data is coming from a trusted source

- **Benefits**
  - Less storage required! No index is created to enforce the constraint.
  - No runtime overhead to maintain unique indexes during `LOAD, INSERT, UPDATE, or DELETE`
  - The uniqueness definition informs the query compiler of unique data, enabling opportunities for superior query execution plans

```
CREATE TABLE t1 (c1 INTEGER NOT NULL, c2 INTEGER,
                                 PRIMARY KEY (c1) NOT ENFORCED);

ALTER TABLE t1 ADD CONSTRAINT unique1 UNIQUE (c2) NOT ENFORCED;
```


© 2013 IBM Corporation

Mixing Row and Columnar Tables

- **DB2 10.5 supports mixing row and columnar tables seamlessly**
  - In the same tablespace and bufferpools
  - In the same query
- **Best query performance for analytic queries usually occurs with all tables columnar**
- **Mixing row and columnar can be necessary**
  - Point queries (highly selective access) favor row-organized tables with index access
  - Small, frequent, write operations favor row-organized tables

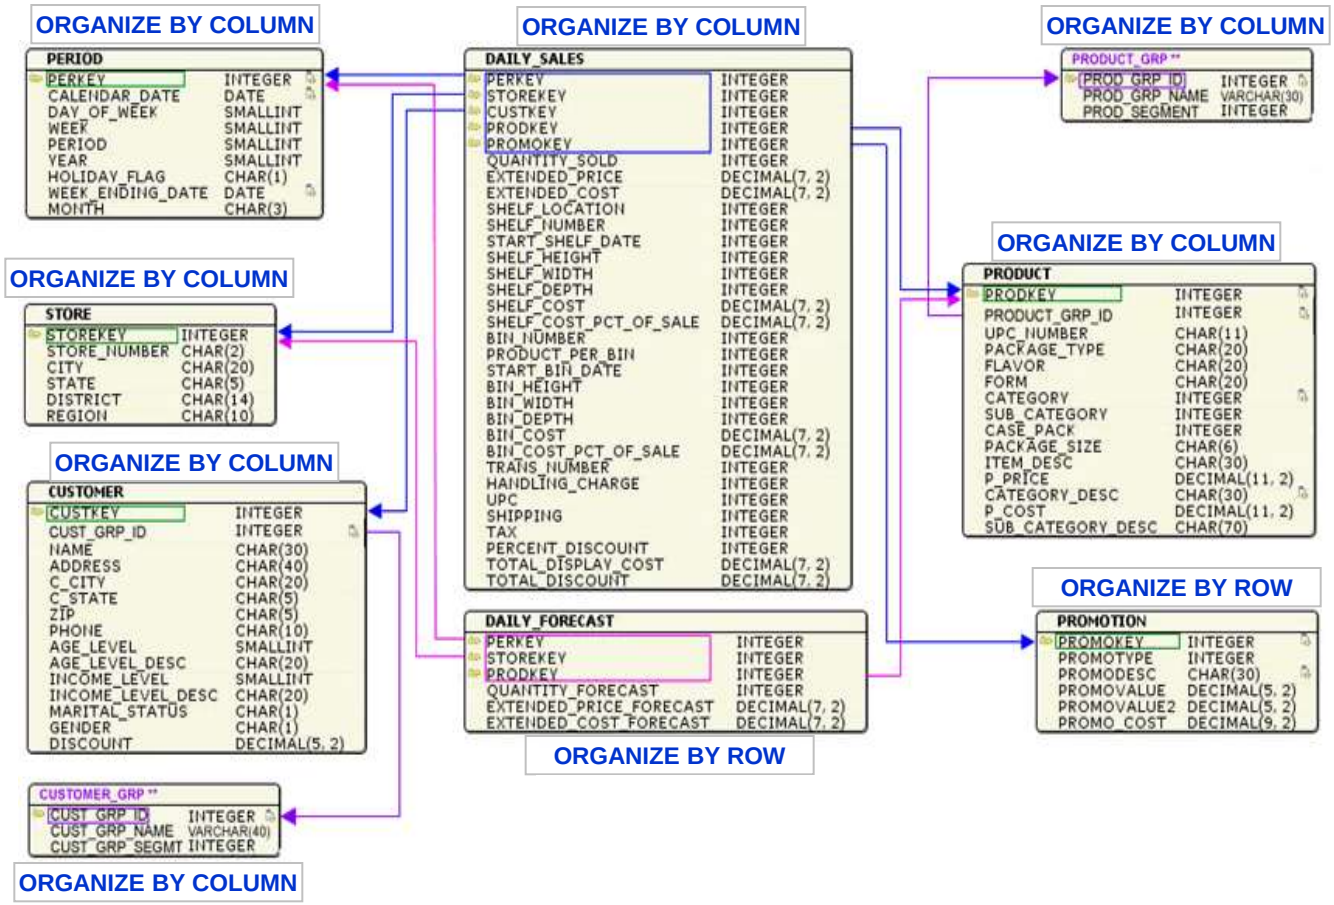

© 2013 IBM Corporation

Automatic Space Reclaim

- **Automatic space reclamation**
  - Frees extents with no active values
  - The storage can be subsequently reused by any table in the table space
- **No need for costly DBA space management and `REORG` utility**
- **Enabled out-of-the box for column-organized tables when `DB2_WORKLOAD=ANALYTICS`**
- **Space is freed online while work continues**
- **Regular space management can result in increased performance of `RUNSTATS` and some queries**

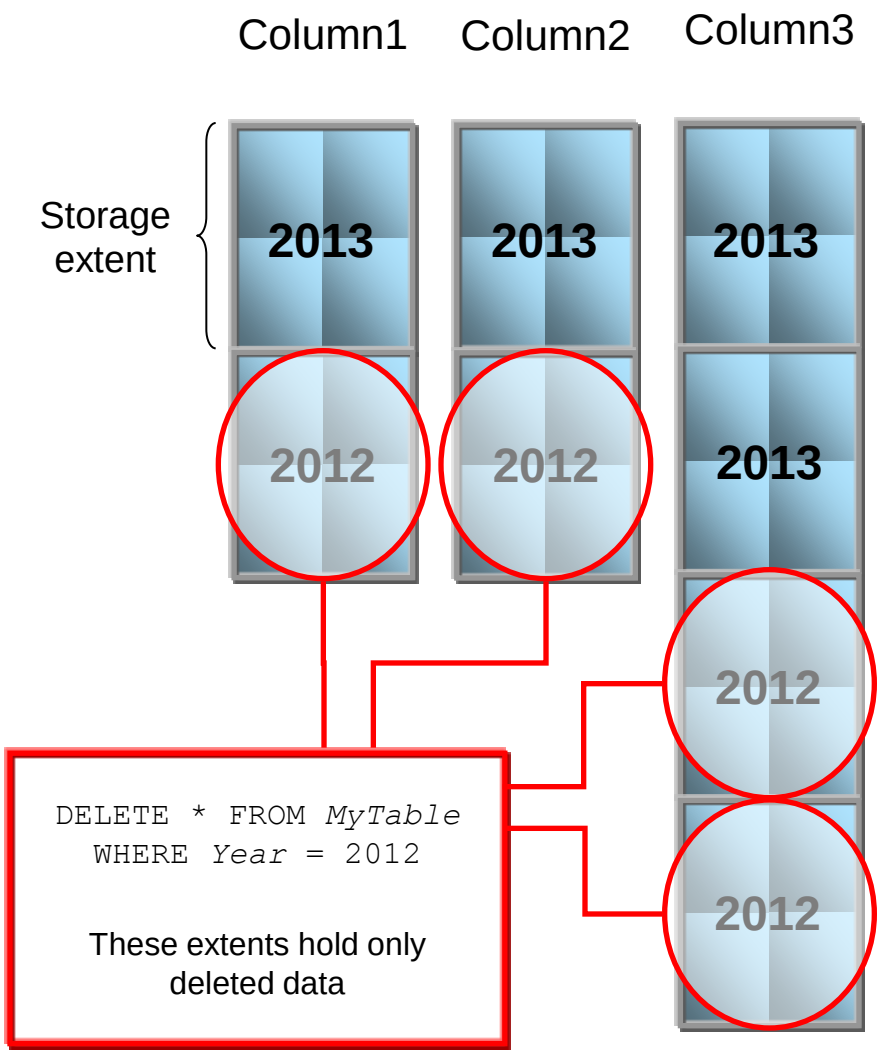

© 2013 IBM Corporation

Early Customer Examples

© 2013 IBM Corporation

BNSF Railway – Faster Queries, Less Management

"When we compared the performance of column-organized tables in DB2 to our traditional row-organized tables, we found that, on average, ***our analytic queries were running 74x faster when using BLU Acceleration."***

- Kent Collins, Database Solutions Architect, BNSF Railway

"What was really impressive is the fact that we could get ***significantly better performance with DB2 10.5 using BLU Acceleration without having to create indexes or aggregates on any of the tables****. That is going to* ***save us a lot of time*** *when designing and tuning our workloads."*

- Kent Collins, Database Solutions Architect, BNSF Railway

© 2013 IBM Corporation

CCBC – Significantly Less Storage, Better Performance

*"**10x.** That's how much smaller our tables are with BLU Acceleration. Moreover, **I don't have to create indexes or aggregates, or partition the data**, among other things. When I take that into account in our **mixed table-type** environment, that **number becomes 10-25x**."*

-Andrew Juarez, Lead SAP Basis and DBA

*"When I converted one of our schemas into DB2 10.5 with **BLU Acceleration tables, the analytical query set ran 4-15x faster**."*

- Andrew Juarez, Lead SAP Basis and DBA

© 2013 IBM Corporation

Value of DB2 with BLU Acceleration?

BLU Acceleration

Next Generation Database for Analytics

- **Extreme performance out-of-the-box**
- **Massive storage savings**
  - No indexes required
- **Lower cost of operational analytics**

Seamlessly Integrated

- **Built seamlessly into DB2**
- **Consistent SQL, interfaces, administration**
- **Dramatic simplification**
  - Less to design
  - Less to tune
  - **Just Load and Go**

Hardware Optimized

- **In memory optimized**
  - Compressed in memory
- **Modern CPU Exploitation**
- **I/O optimized**
  - Only read columns of interest

© 2013 IBM Corporation

What Does BLU Acceleration Give Us?

- **Order of magnitude improvements**
  - Consumability
  - Speed
  - Storage savings
- **Breakthrough technology**
  - Combines and extends the very best technologies
  - Over 25 patents filed and pending
  - Leveraging years of IBM R&D spanning 10 laboratories in 7 countries worldwide
- **Typical experience**
  - Simple to implement and use
  - 10x-20x performance gains
  - 5x-20x storage savings vs. uncompressed data with indexes

Request a beta nomination form
db2beta@ca.ibm.com

© 2013 IBM Corporation