

What's New in DB2 10.1 for Linux, UNIX, and Windows



Kelly Schlamb
WW IM Technical Sales Acceleration, DB2 pureScale Specialist
kschlamb@ca.ibm.com

Agenda

- **Lower operational costs**
 - Adaptive compression
 - Multi-temperature data management
 - Workload Management enhancements (Reliability)
 - Range (table) partitioning enhancements
 - Improved SQL query performance and index management
 - Real-time data warehousing

- **Ease of development**
 - Temporal tables (time travel query)
 - Row and Column Access Control
 - SQL compatibility enhancements (Oracle compatibility)
 - NoSQL (RDF / triple-store) graph store

- **Reliability**
 - DB2 pureScale feature and enhancements
 - DB2 HADR feature enhancements

- **DB2 Advanced Edition enhancements**

DB2 10.1 Adaptive Compression

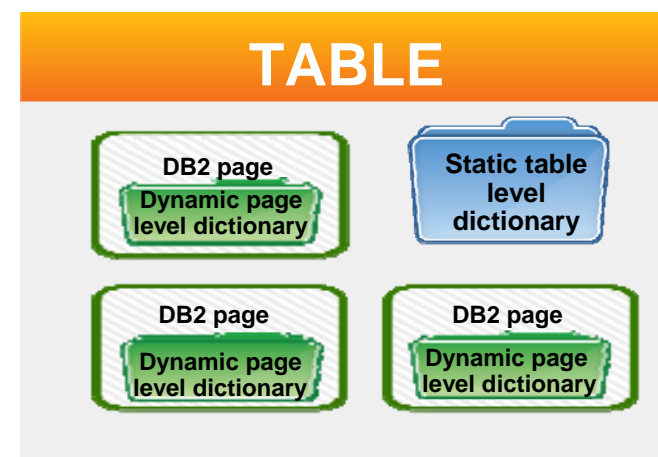
Low Operational
Costs

- Adaptive compression is an **advanced row compression** technique that uses two levels of compression dictionaries (table-level and page-level) to **improve compression ratios**, particularly as data changes

- **How it will help you**
 - **Lower costs**
 - Postpone upcoming storage purchases
 - Lower ongoing storage needs
 - **Better compression rates = increased storage savings**
 - Easier administration with reduced need for table reorganization
 - Compression rates remain very high, **even as data grows and changes**
 - Table reorganization not required to maintain high compression rates
 - **Higher performance**
 - Faster queries for I/O-bound environments
 - Faster backups

Adaptive Compression

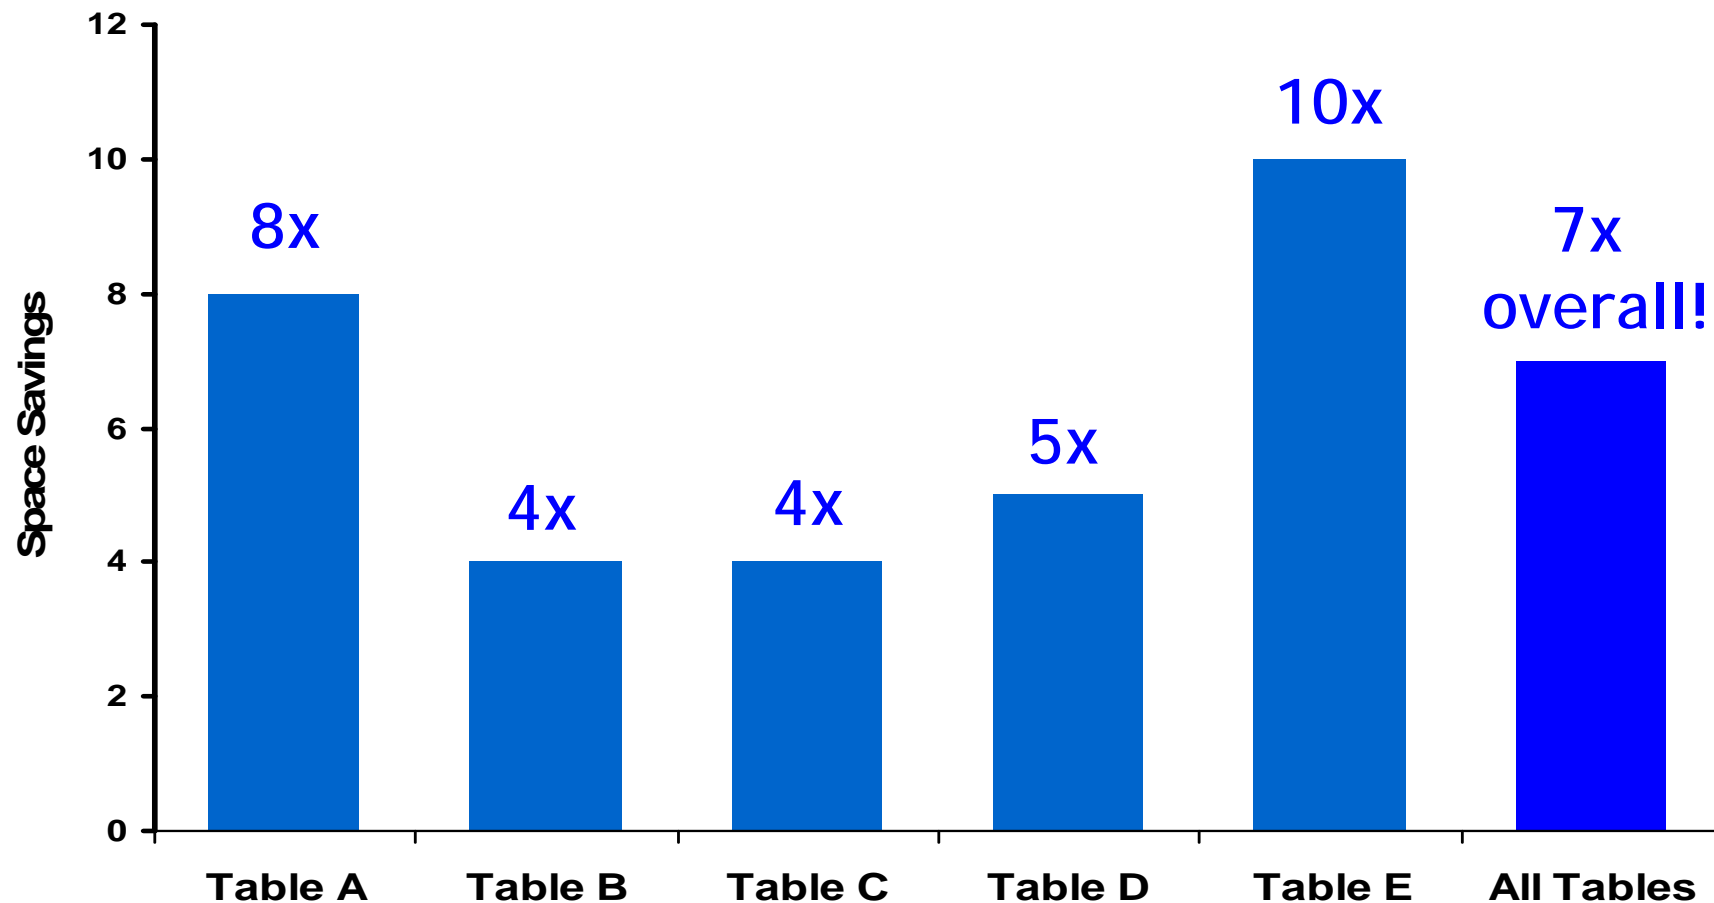
- **Adaptive Compression is an enhancement to the Classic Row Compression found in DB2 9.7**
 - Compress rows by using a combination of two types of dictionaries
 - Global static table-level dictionary
 - Local page -level dictionaries
- **Benefits**
 - Page level dictionaries **adapt to data** skew over a period of time
 - **No REORGs required** to maintain high compression ratio as data changes
 - **Less disk space** for data and logs
 - 2x storage savings for tables over Classic Row Compression
 - 5x-8x overall table compression
 - **Reduced I/O**
 - Fewer pages to process



- **Better compression ratios than Classic Row Compression**
- **Over time reduces need of a REORG to find locally recurring patterns**



Adaptive Compression Provides Up to 10x Savings



Low Operational Costs

Results based on customer testing on customer data during the DB2 10 Early Access Program testing.

Archive Log Compression Added in DB2 10

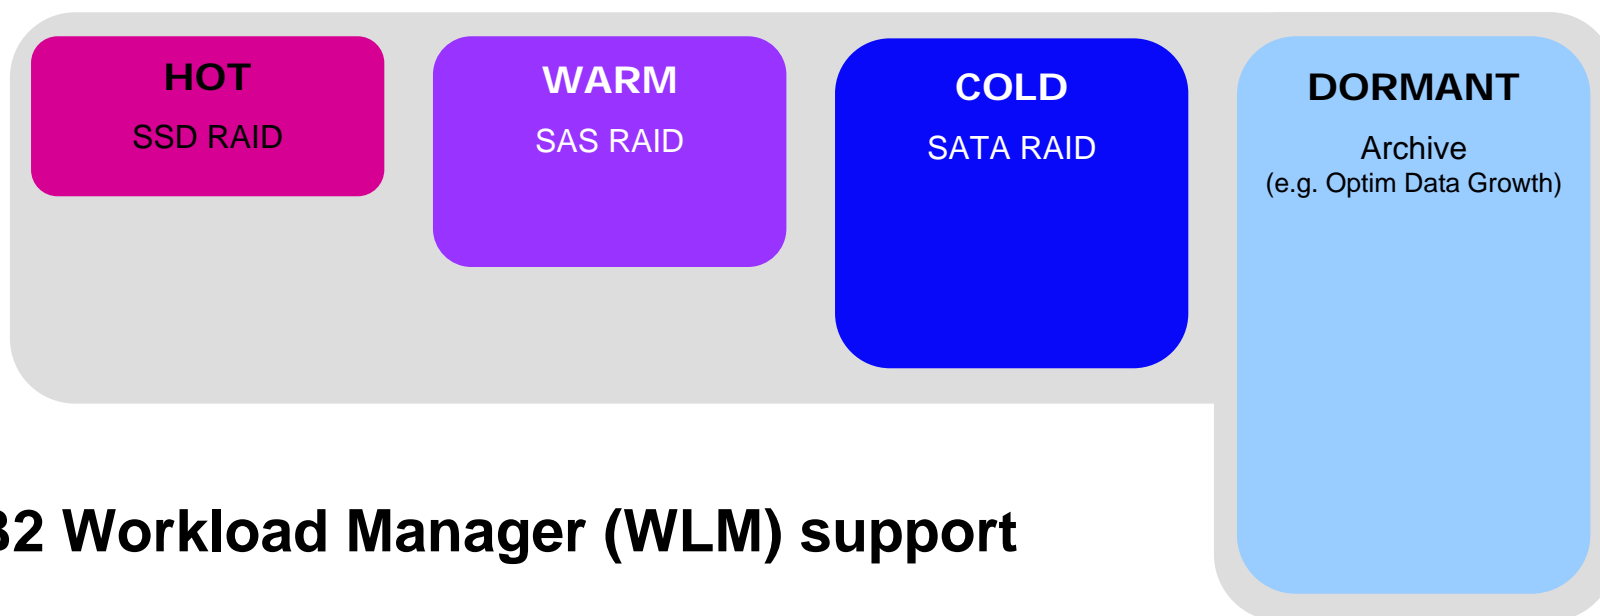
Low
Operational
Costs

- **New archive log compression reduces log archive storage**
- **Simply turn it on and DB2 does the compression for you**
 - `logarchcompr1` database configuration parameter set to ON
- **Large SAP using DB2 for LUW customer generates 60GB of log per day and they keep 8 weeks of archives**
 - Storing 3.3TB of archived log files
 - Compression of 4x results in storage of only 825GB for 8 weeks

Multi-Temperature Data Management

Increase Ability to Meet SLAs. Postpone Hardware Upgrades

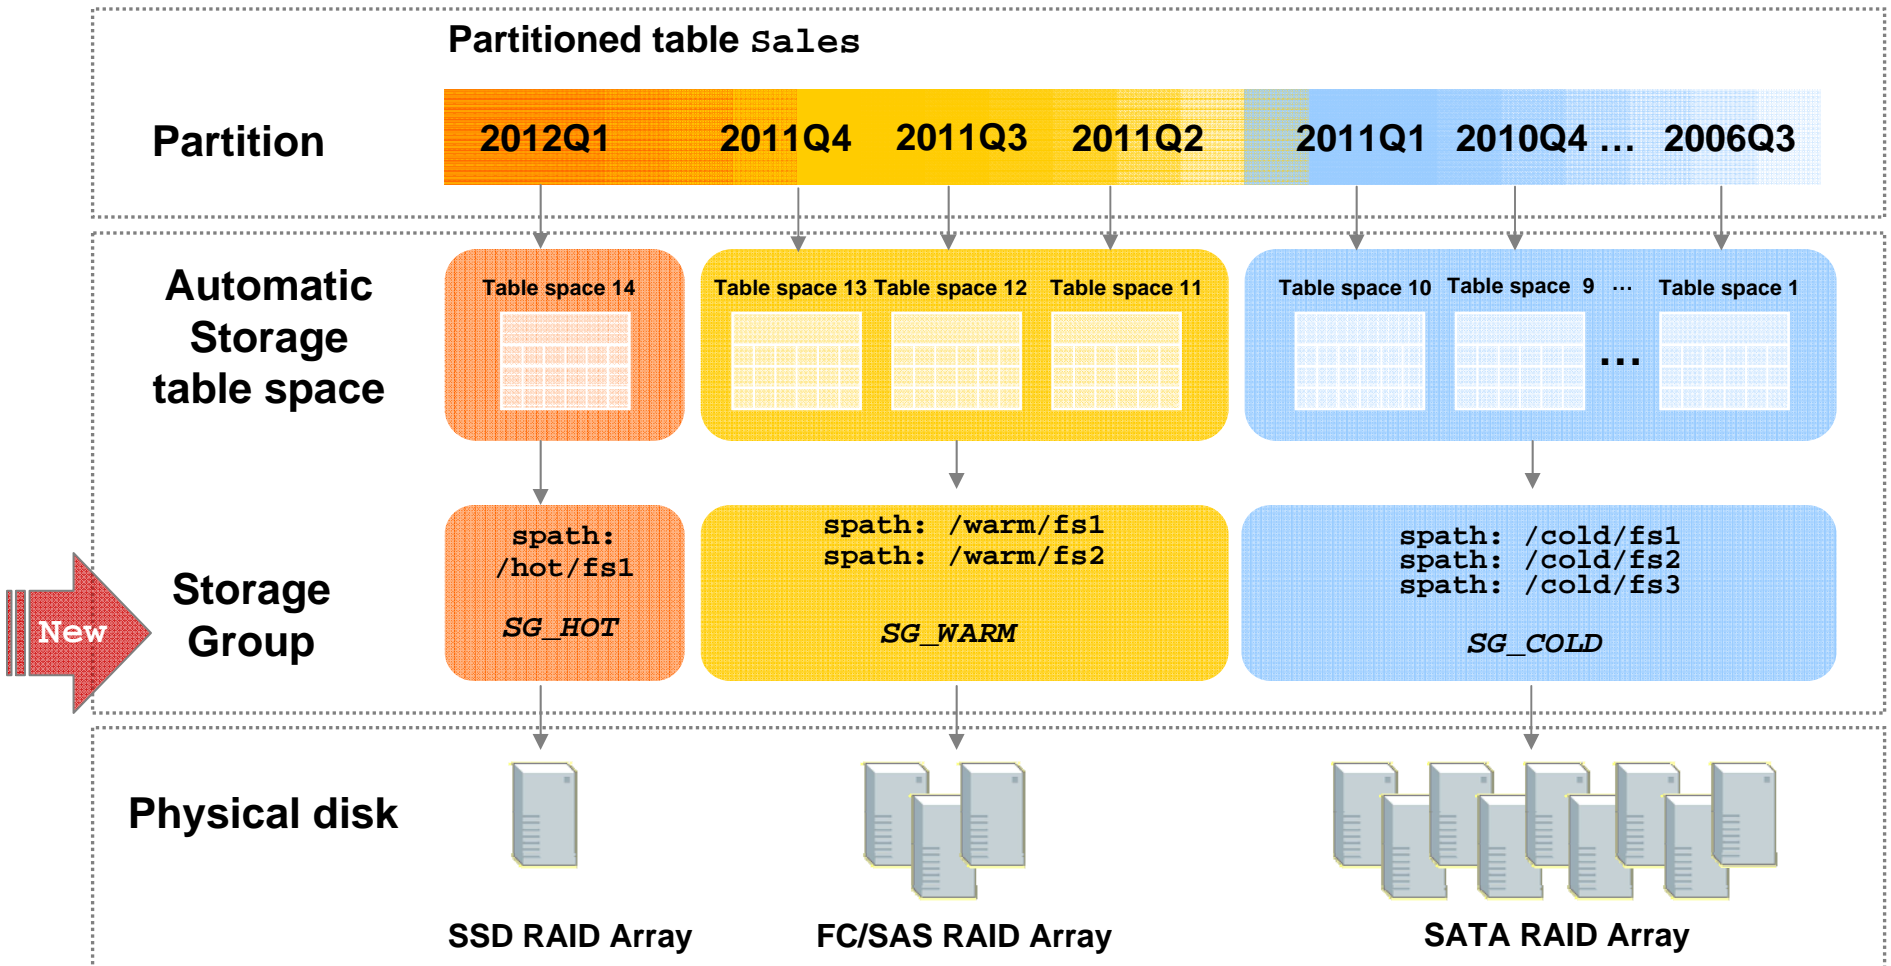
- **Storage pools for different tiers of storage**
 - For range partitions, **policy-based automated movement** of data



- **DB2 Workload Manager (WLM) support**

Introducing DB2 10.1 for LUW Storage Groups

Low Operational Costs



Multi-Temperature Data Management

Integration With the DB2 Workload Manager

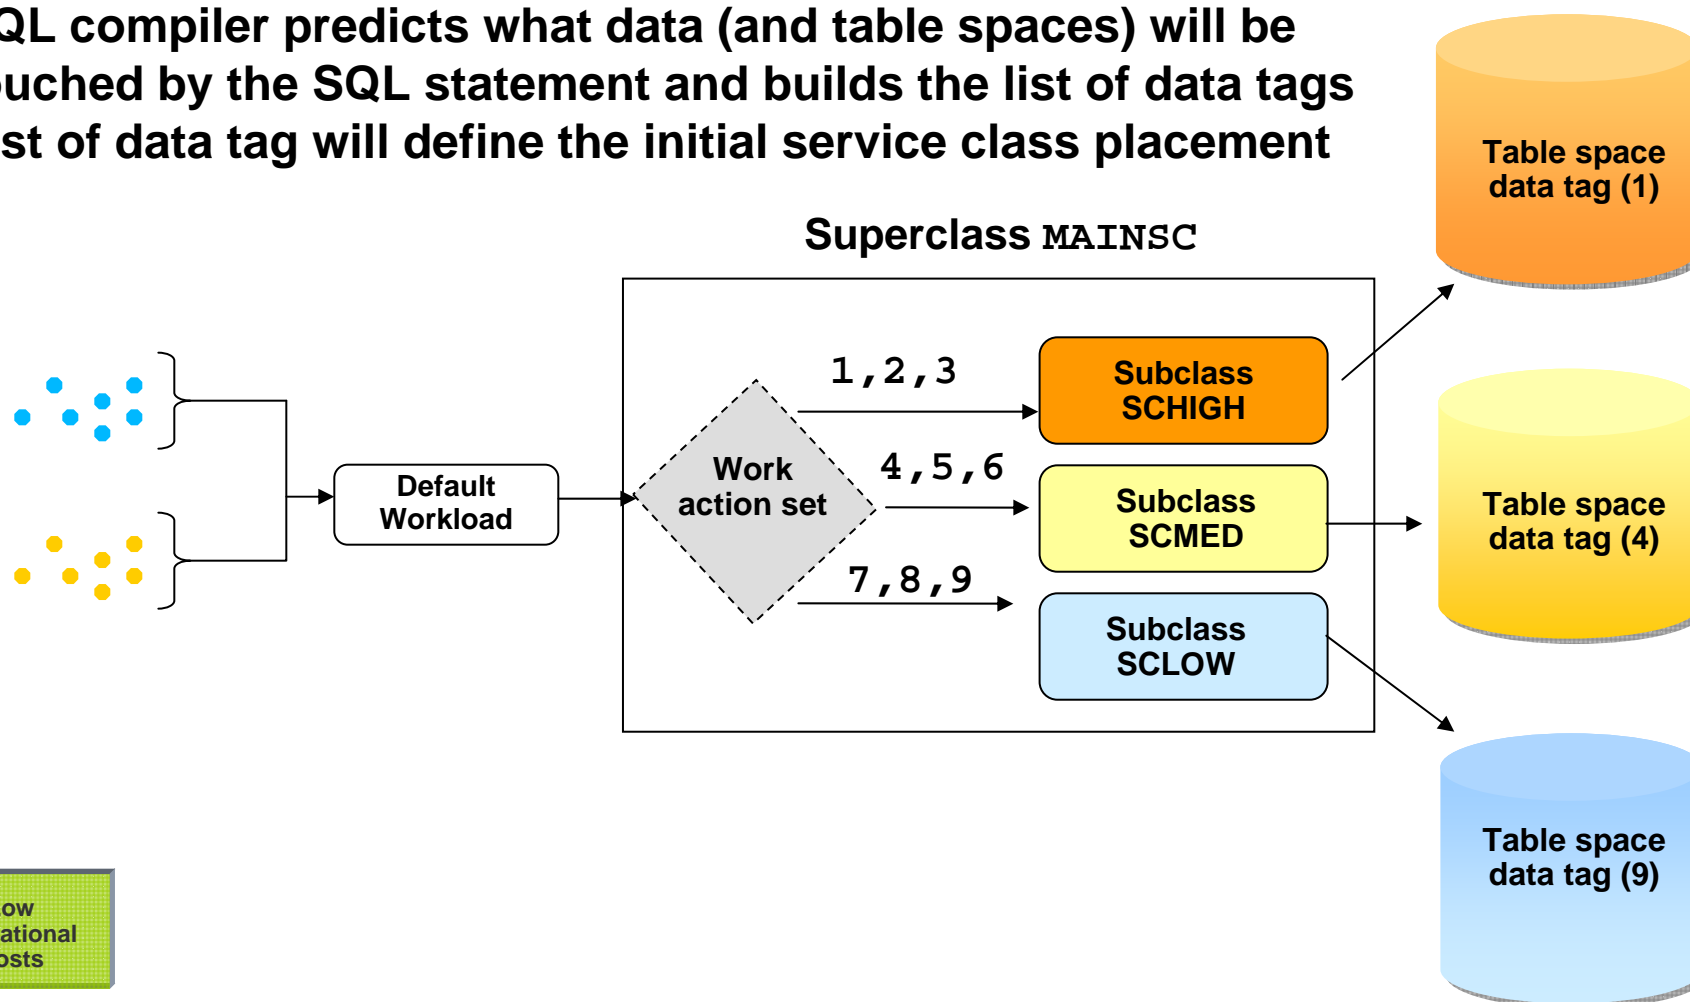
Low Operational
Costs

- **DB2 WLM work class and threshold DDL have been extended to support the new data tag attribute**
- **Work class sets are predictive based on compilation information**
 - Sometimes there isn't enough information at compile time (for example, queries with parameter markers) to predict which table spaces will be touched. For this reason...
 - **Activities can be mapped to a service subclass** based on the data expected to be touched before the activity starts to run
- **Data tag thresholds are reactive and use information that is available at runtime**
 - **Activities can be remapped** to a service subclass based on the data touched by the activity as it is running

Workload Management - WLM Environment At a Glance

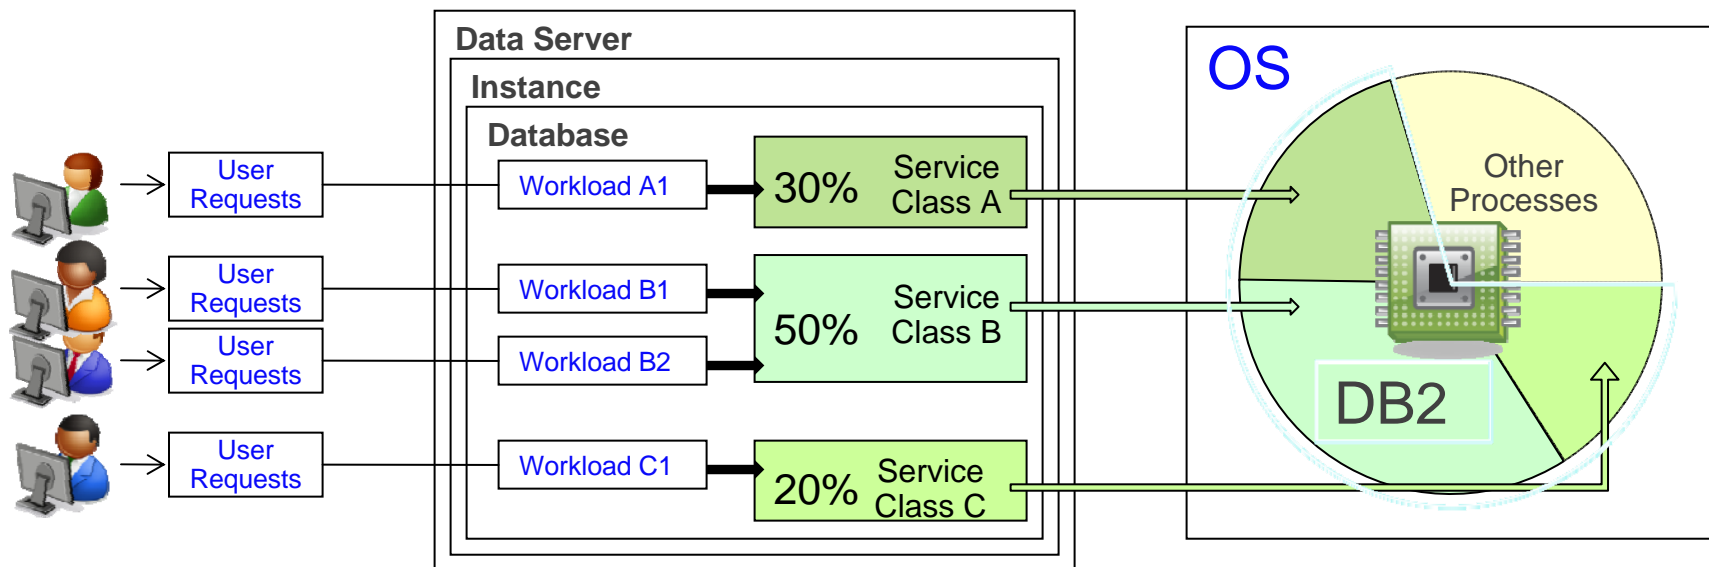
Predictive

- Map statement to service class (**Predictive**)
- SQL compiler predicts what data (and table spaces) will be touched by the SQL statement and builds the list of data tags
- List of data tag will define the initial service class placement



WLM Dispatcher

- **DB2 9.7 allowed integration with OS WLM to provide additional control over resources, in particular, the amount of CPU resources allocated to each service class**
 - Only supported on specific platforms
 - DBAs would require extra knowledge and privileges to set up the OS WLM
- **WLM Dispatcher is a CPU scheduling technology in the DB2 engine at instance level**
 - Allows for fine allocation of CPU amongst DB2 work executing in user and maintenance service classes
 - It works in parallel with OS WLM mechanisms
 - Supported for all DB2 LUW platforms
 - Provides additional flexibility with use of hard and soft CPU shares, and CPU limits

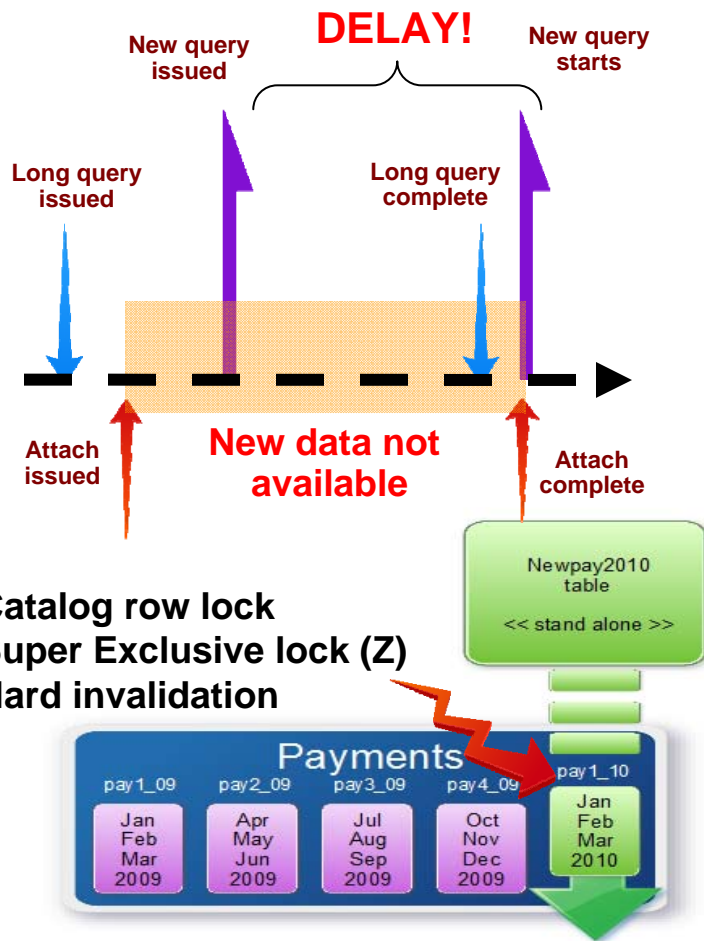


Improved ADD/ATTACH Behavior

Low Operational Costs

DB2 9.7


- Command waits for existing queries to complete and then delays new queries



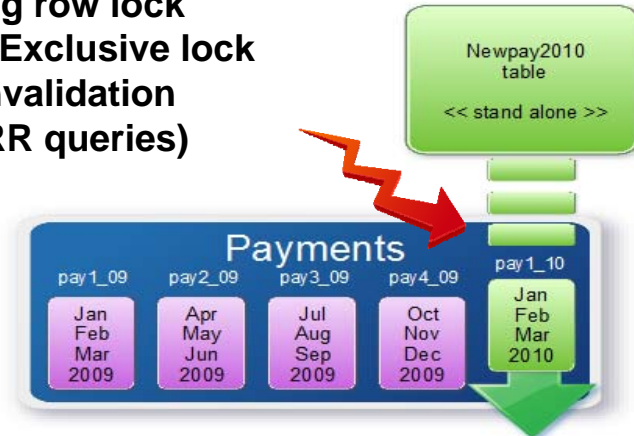
- Catalog row lock
- Super Exclusive lock (Z)
- Hard invalidation

DB2 10.1 for LUW

- ATTACH or ADD completes immediately, allowing new partition to be accessible sooner!



- Catalog row lock
- Intent Exclusive lock
- Soft invalidation (non-RR queries)



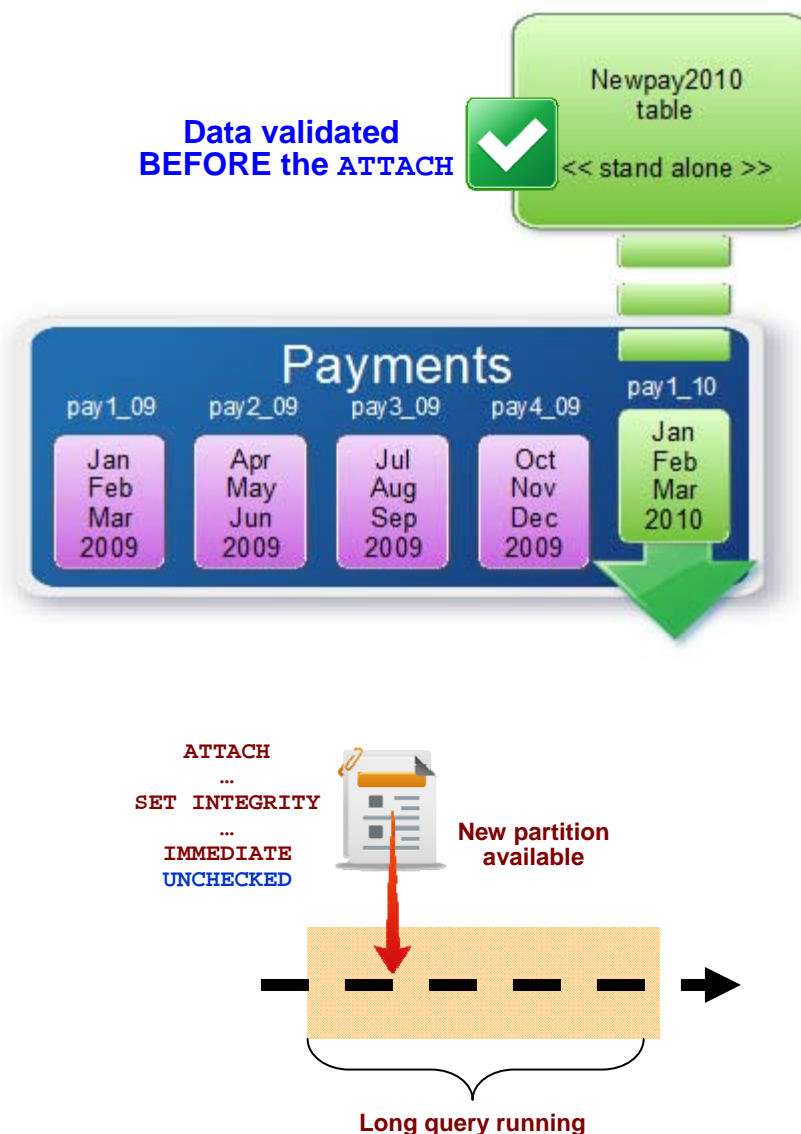
SET INTEGRITY Behavior Comparison - DB2 10.1

■ SET INTEGRITY ... IMMEDIATE UNCHECKED is now available!

- No checking is done except an invalid partitioned index is detected and rebuilt
- Application must validate data to ensure correct query results

■ Restrictions

- If any non-partitioned indexes exist, then internally converted to ALL IMMEDIATE CHECKED



Low
Operational
Costs

Up to 3x Faster Query Performance

Increase Ability to Meet SLAs. Postpone Hardware Upgrades

- **Multi-core parallelism enhancements**
- **Performance improvements for**
 - Queries over star schemas
 - Queries with joins and sorts
 - Queries with aggregation
 - Hash joins
- **Higher performance**
 - Up to 35% faster out-of-the-box performance
 - Up to 3x faster when using new features
- **Lower costs**
 - Postpone hardware upgrades



*"DB2 10.1 performance improvements helped us in achieving nearly 5X faster query response times."
-- Bin Ma, Division Manager, Beijing Join-Cheer Software Co. Ltd*

Index Management Redefined

Increase Ability to Meet SLAs. Lower Administration Costs

- **Jump scan**
 - Reducing the number of indexes needed
- **Smart index prefetching**
 - Faster index access and fewer index reorganizations
- **Smart data prefetching**
 - Faster index scans and fewer index reorganizations
- **Predicate evaluation avoidance**
 - Faster index scans
- **Higher performance**
 - Faster index performance
- **Lower costs**
 - Fewer indexes to maintain
 - Dramatic reduction in index reorganization

“Jump Scan optimizes buffer usage by 75 to 80%, resulting in very good improvement in overall performance and saving the CPU cycles.”

-- Shanmukhaiah D, Cognizant Technology Solutions.

Improved Performance for Queries with Composite Indexes

- For workloads with many ad-hoc queries, it is often difficult to optimize a database for high performance for all queries
 - Index gaps are common in queries against tables with composite indexes

Composite Index on
columns A,B,C

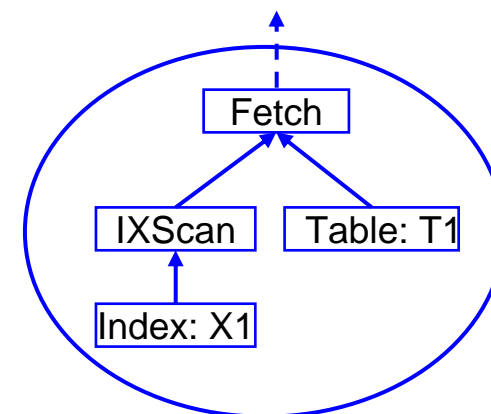
```
SELECT * FROM t WHERE a=5 AND c=10
```

Query contains an index gap on
column B in the composite index

- DB2 10.1 enables you to
 - Eliminate need for additional indexes using Jump Scan
 - Eliminate need for re-orgs related to index scan and ISCAN/FETCH performance
 - Eliminate extra CPU cost by skipping duplicate keys when possible
- Jump scans provide a solution to the index gap problem
 - Ad-hoc queries may utilize Jump Scan without additional index definitions if general-purpose composite indexes are available

Smart Index/Data Pre-fetching


- **Currently, when data is loaded into or updated in an indexed table, index leaf pages become non-sequential**
 - Unclustered pages reduce asynchronous I/O and increase synchronous I/O
 - Performance for such queries becomes sub-optimal
 - An index re-org is an expensive operation to fix this
- **In DB2 10.1, **smart Index pre-fetch** minimizes negative I/O performance impact**
 - Introduces read-ahead pre-fetching for index scans
 - Lets the optimizer select the optimal pre-fetching method
 - When sequential detection is not sufficient, dynamically switch to read-ahead pre-fetch during query execution
- **Smart data pre-fetch**
 - Similar approach to smart index pre-fetching
 - Used with Index Scan + Data Fetch [ISCAN-FETCH]
 - Uses sequential detection and switches to read ahead pre-fetching when needed



Smart index pre-fetching against disorganized index

Combines sequential detection with read-ahead pre-fetching

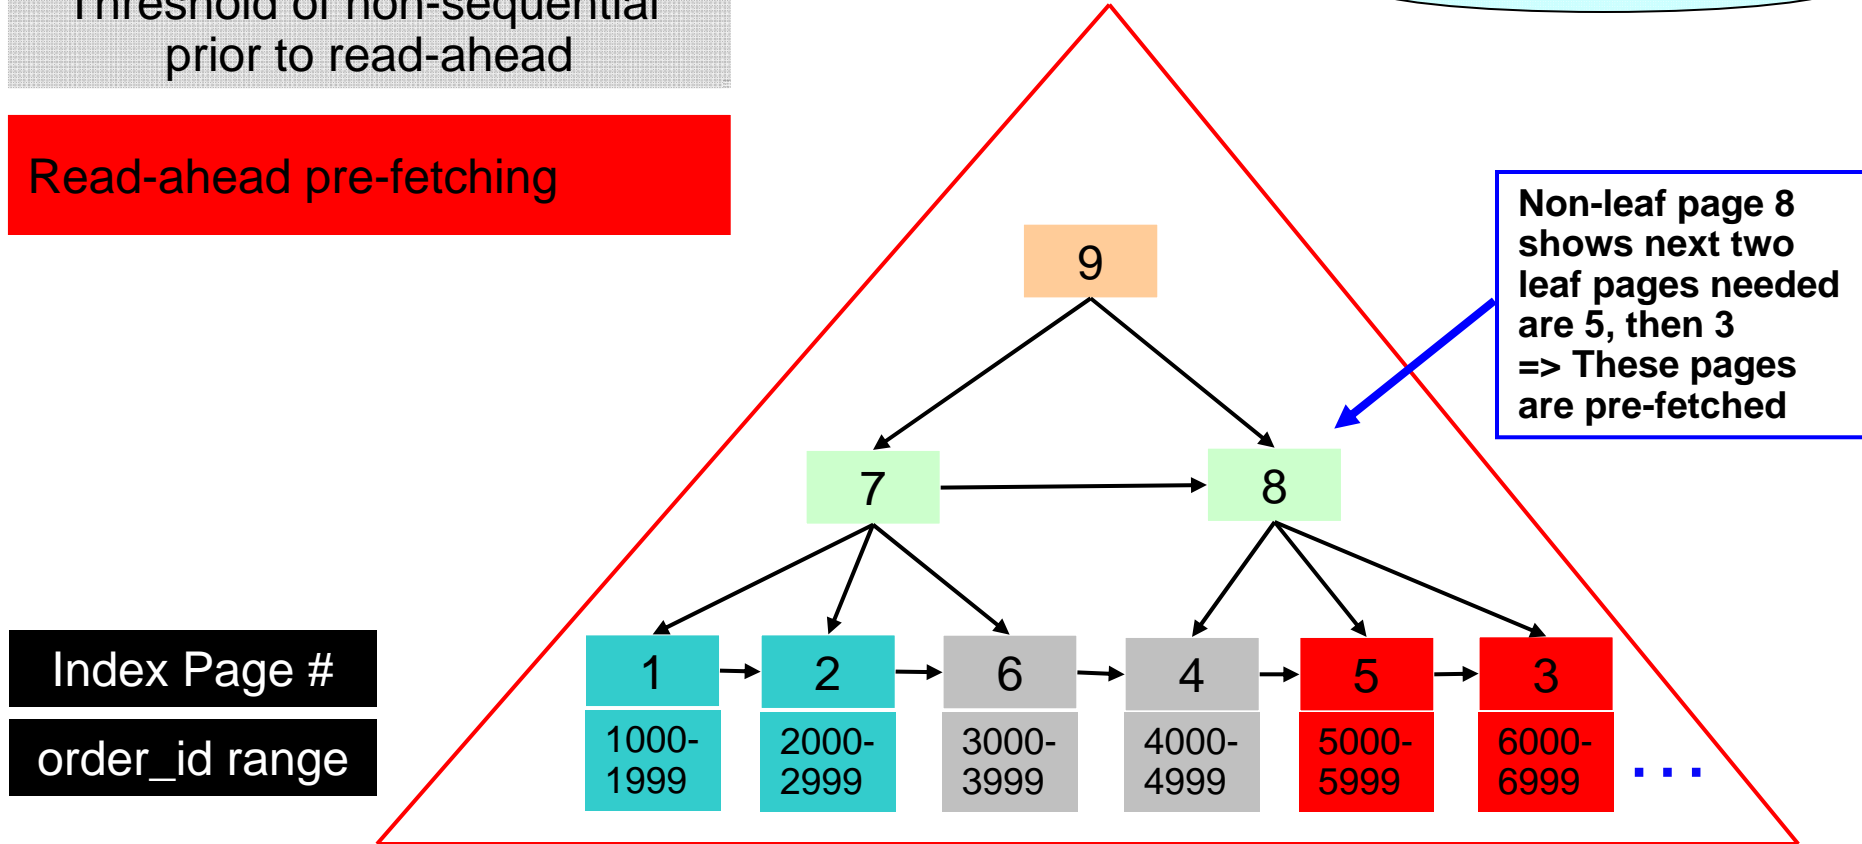
Index scan needs leaf pages 1,2,6,4,5,3



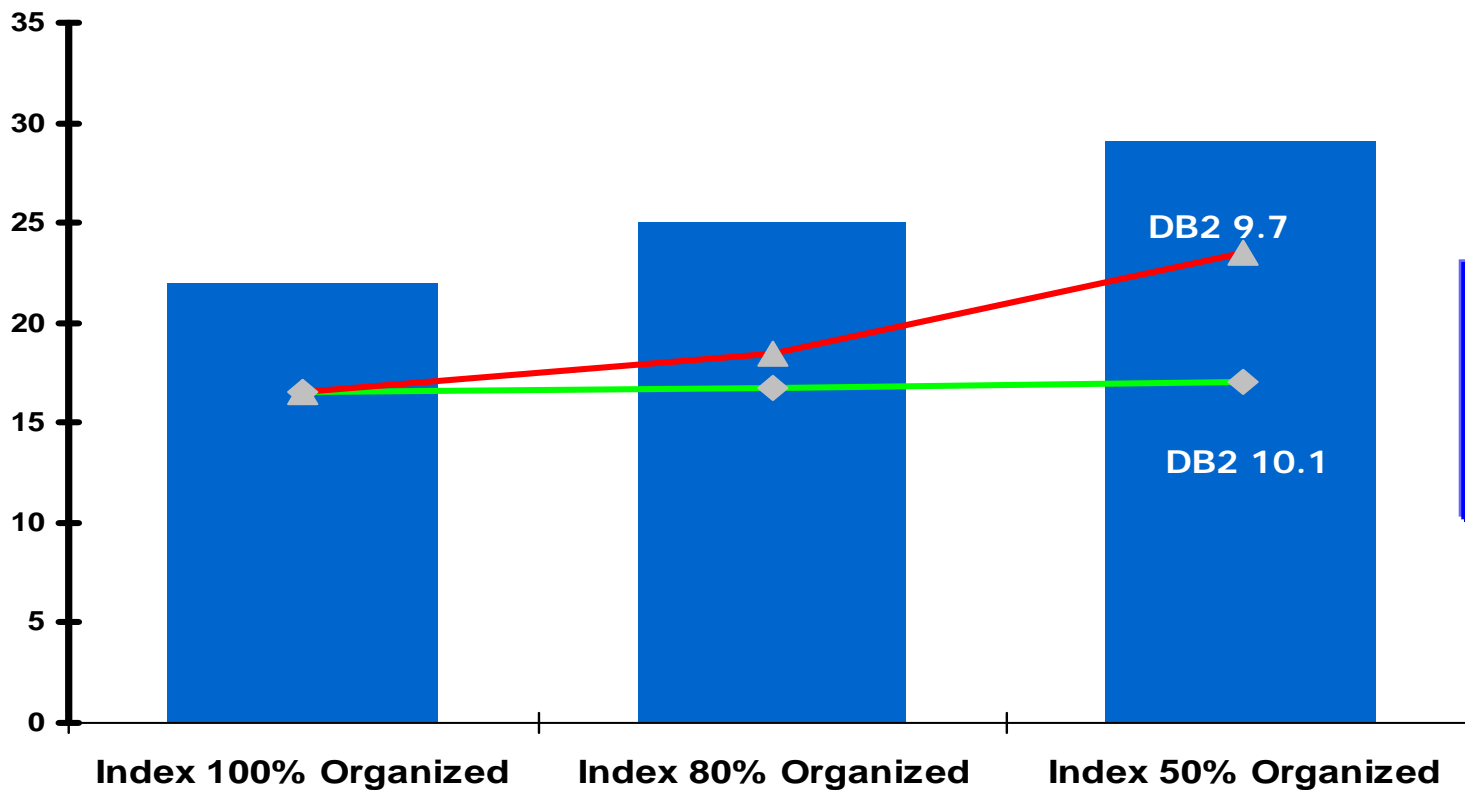
Sequential detection pre-fetching

Threshold of non-sequential
prior to read-ahead


Read-ahead pre-fetching



Smart Index Prefetching Means Easier Management



**DB2 10.1
28%
faster!**



Low Operational Costs

Results based on IBM internal testing

Real-Time Data Warehousing: INGEST Utility

Low Operational Costs

■ Continuous feed of data

- New, **high-speed, client-side utility** that streams data from files or pipes into a target table
 - **Fast!** Uses an optimized multi-threaded design to process data in parallel and to ingest data into multiple database partitions in parallel
 - **Available!** Operates at the row level using row-level locking. Allows normal read/write access against the target table while it is running
 - **Continuous!** Not only can it run against pre-processed data directly from files, INGEST can run in a continuous manner by reading a stream from a pipe
 - **Robust!** Designed to be tolerant of unexpected failures. Can be restarted from last commit point
 - **Flexible and Functional!** Supports a number of input formats and has rich data manipulation capabilities

LIS.TEC
ISV Distributor

*“You can now **continuously feed data** into your data warehouse **at a high rate** even whilst you are **running queries against the tables** in our data warehouse. InfoSphere Warehouse 10 represents a greatly strengthened offering for the data warehouse market.”*

—Ivo Grodtke, LIS.TEC GmbH

Time Travel Query

Easily Analyze Historical Trends and Predict Future Demand

- Query data as it was at any point in the past, or as it will be at some point in the future
- Provide temporal logic and analysis
- Can be based on system time, application-based time and queried using AS OF queries
- **How with it help you?**
 - Higher performance
 - Native support for fast performance
 - Lower costs
 - Up to 45x fewer lines of code than Java
 - Eliminate need to maintain and update custom temporal implementations
 - Easy to administer (simply turn on for any table)



*"The use of standardized SQL syntax for temporal operations and the integration deep into the database engine, make **DB2 a leader in second generation bitemporal data management - Bitemp 2.0!**"*

—Craig Baumunk, Principal at BitemporalData.com

Time Travel Query – Temporal Tables

- **System-period temporal table**
 - DB2 tracks the **period** when a row is valid (beginning when it is inserted, and ending when it is either updated or deleted)
 - All currently valid rows are stored in one table
 - When a row is no longer currently valid, it is automatically moved to an associated history table that is transparent to users and applications
 - DB2 will automatically query the history table and return applicable rows based on the SQL executed against the base table

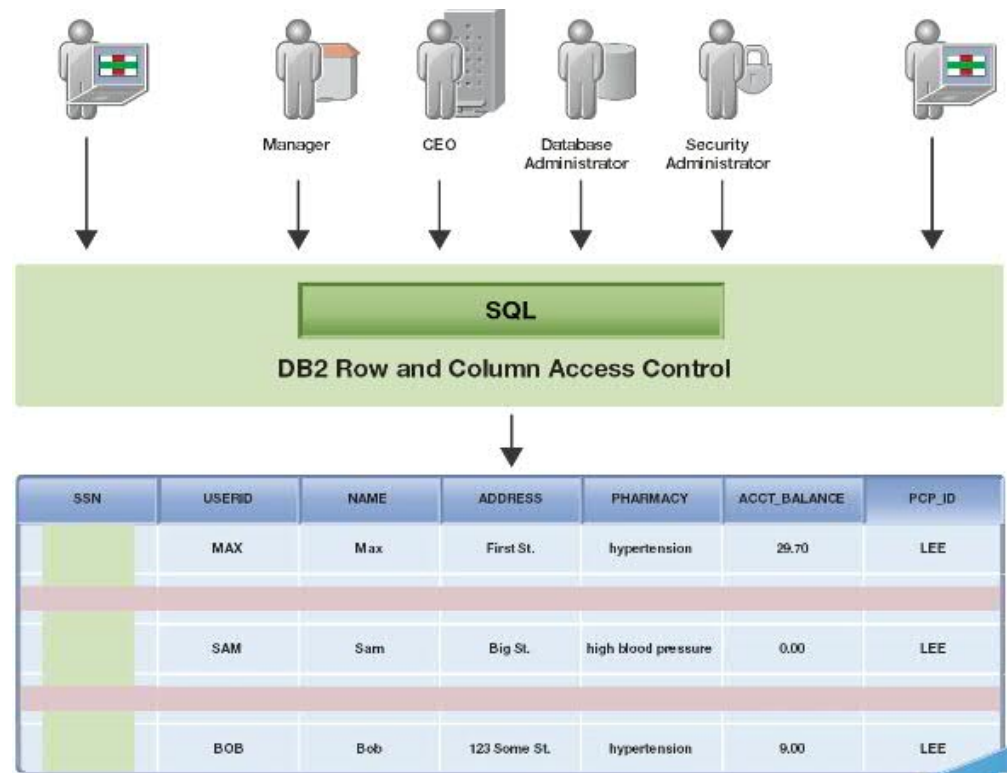
- **Application-period temporal table**
 - The **user or application updates** the beginning and end of the period in which the information is valid
 - All data is kept in a single table and DB2 will automatically split rows based on SQL activity against the existing rows

- **Bi-temporal table**
 - Combines characteristics of both types of temporal tables

Row and Column Access Control

Ease of Development

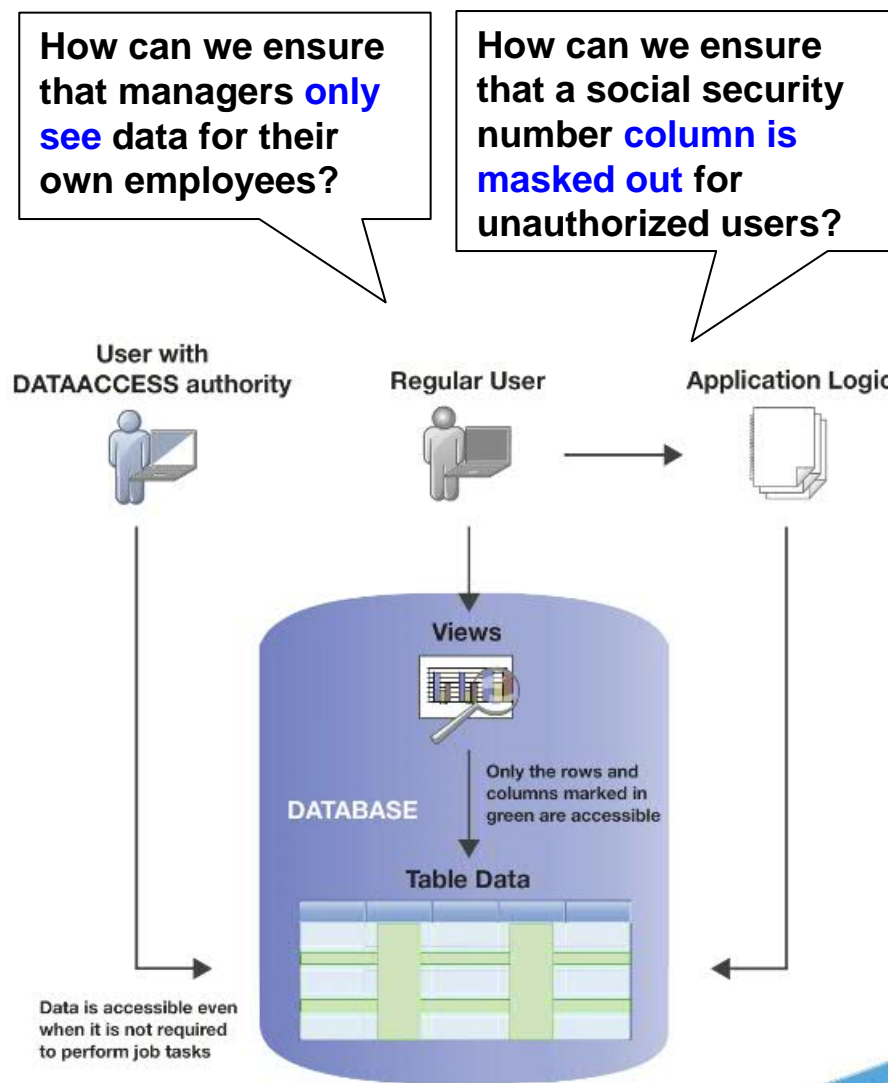
- **Additional layer of data security** introduced in DB2 10.1 for LUW
- **Complementary** to table level authorization
- **Allows access only to subset of data** useful for job task
- **Controls access to a table at the row, column, or both**
- **Two sets of rules**
 - Permissions for rows
 - Masks for columns



Why Use Row and Column Access Control?

Ease of Development

- Currently, data access is restricted via views or application logic
- Users with direct access to databases can bypass these layers
 - Example: Users with `DATAACCESS` authority can still view all data
- DB2 10.1 provides a **new way to control data access at row/column level**
 - Set up rich security policies
 - Prevents administrators with `DATAACCESS` authority from accessing all data in a database
 - No dependency on application logic
 - Facilitates table level multi-tenancy



DB2 10 PL/SQL Compatibility

Ease of Development

Average PL/SQL Compatibility Moves Above 98%

*"The total cost of ownership with DB2 running on IBM systems is almost **half the cost of Oracle Database on Sun systems.**"*
--- Reliance Life Insurance

*"We switched from Oracle Database to IBM DB2 and **cut our costs in half, while improving performance and reliability of business applications.**"*
Sandro Reátegui Banco de Crédito del Peru

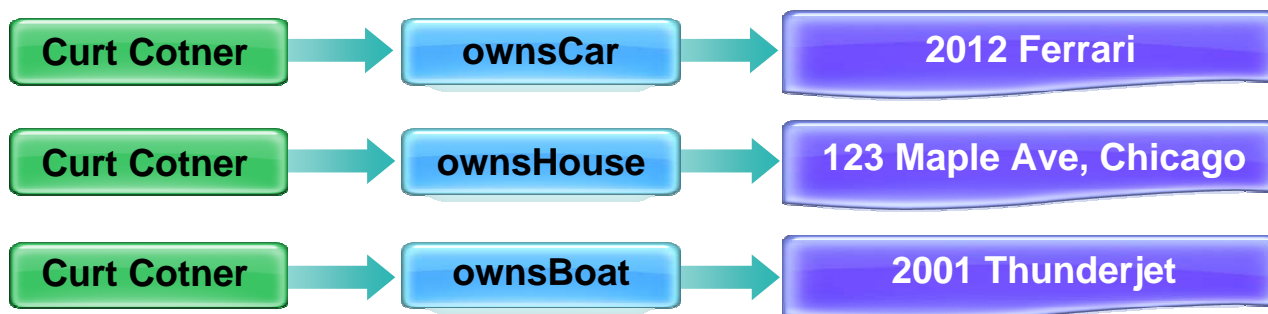
- Moved from Oracle Database to IBM DB2
 - Used "compatibility features"
 - **3-30x faster query performance**
 - **200% improvement in data availability**
- JSC Rietumu Banka

9.7.1	SUB STRB	Increase compatibility
9.7.1	UDF Parameters: INOUT	Increase compatibility
9.7.1	FORALL/BULK COLLECT	Increase compatibility
9.7.1	Improve BOOLEAN	Increase compatibility
9.7.1	Conditional Compilation	Enhancement
9.7.1	Basic DPF Support	Broaden coverage
9.7.1	OCI Support	Broaden coverage
9.7.2	UDF Parameters: DEFAULT	Increase compatibility
9.7.2	Obfuscation	Enhancement
9.7.2	NCHAR, NVARCHAR, NCLOB	Increase compatibility
9.7.3	NUMBER Performance	Performance
9.7.3	Runtime "purity level" Enforcement	Increase compatibility
9.7.3	RATIO_TO_REPORT Function	Increase compatibility
9.7.3	RAISE_APPLICATION_ERROR	Increase compatibility
9.7.3	Small LOB Compare	Increase compatibility
9.7.4	Multi-action Trigger & Update Before Trigger	Increase compatibility
9.7.4	Autonomous Tx Improvements	Increase compatibility
9.7.4	LIKE Improvements, LISTAGG	Increase compatibility
9.7.4	ROW & ARRAY of ROW JDBC Support	Increase compatibility
9.7.5	Pro*C Support	Increase compatibility
9.7.5	Nested Complex Objects	Increase compatibility
10	Local Procedure Definitions	Increase compatibility
10	Local Type Definitions	Increase compatibility
10	PL/SQL Performance	Performance

DB2 NoSQL Graph Store

Rapid Application Development

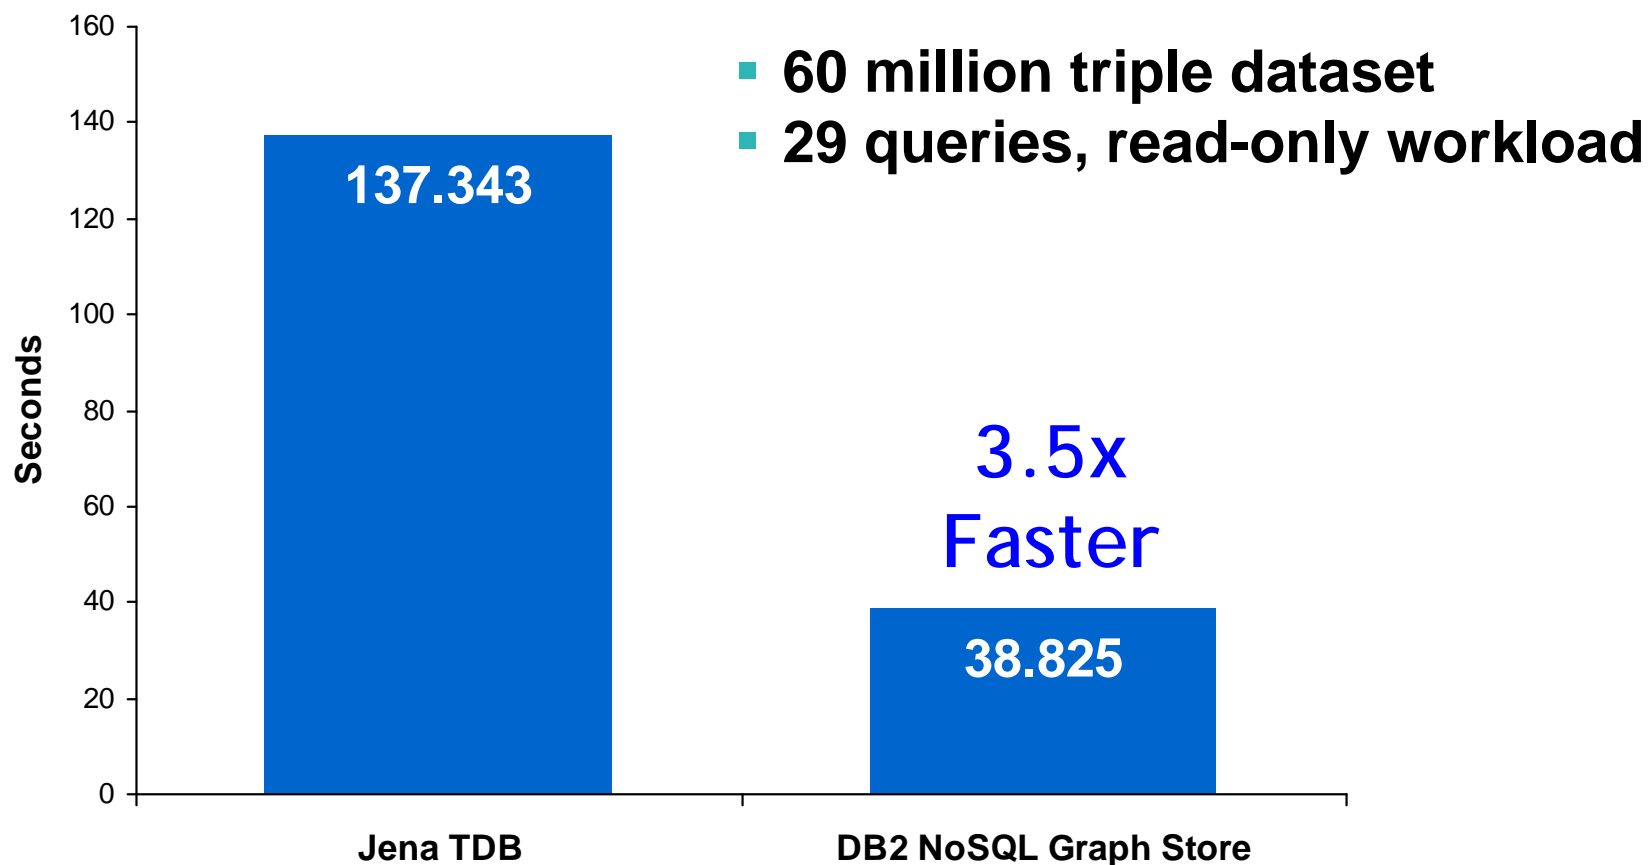
- NoSQL->Graph Store->RDF
- **Optimized method** to store graph triples in DB2
<subject> <predicate> <object>



- **Supports SPARQL 1.0 query language**
- **Higher performance**
 - Almost 4x faster for leading open source semantic Web framework
- **Lower costs**
 - Rapid development with schema-less approach
 - Easy adaption as needs evolve
 - Simpler data management for triples

DB2 NoSQL Graph Store (cont.)

Accelerates Leading Open Source Semantic Web Framework by up to 3.5x




Based on internal benchmark tests of Rational Jazz graph store usage, comparing DB2 10 Graph Store with Jena TDB version 0.8.10.

Ease of
Development

DB2 pureScale Feature

Reliability


- The DB2 pureScale feature is a clustering technology that reduces the risk and cost of business growth
- It allows you to scale your system with near-linear efficiency and predictable query performance, and without modifying applications or changing how data is distributed



DB2 10.1 pureScale Enhancements

Increase Ability to Meet SLAs. Easily Add or Remove Capacity


- **Further improving IBM's shared-disk cluster capability**
- **New features and function in DB2 10.1**
 - Installation enhancements
 - Workload management for DB2 pureScale
 - Range partitioning support
 - Additional backup and restore options
 - Support for 10-gigabit Ethernet
 - Support for multiple InfiniBand adapters and switches




HADR Supports Multiple Standby Servers

Increase Ability to Meet SLAs. Disaster Recovery

- **HADR now supports more than one stand-by server**
 - If primary server fails, principle standby takes over
 - If principle standby then fails, can switch to auxiliary standby
 - Auxiliary standby can provide complete offsite availability, while maintaining speed of local standby
- **Time delay apply available for the standby**



HADR Multiple Standbys (cont.)



- **All standbys are directly connected to primary**
 - No daisy chaining/cascading of standbys
- **Reads on standby supported on all standbys**
- **Takeover (forced and non-forced) supported from any standby**
 - After takeover, configuration parameters on new primary's standbys will be changed automatically so they point to the new primary

DB2 10.1 Advanced Editions Offer a Broad Solution Platform *Higher Value at a Very Competitive Bundled Price*

- **Based on DB2 ESE but includes additional key features and products**
 - DB2 10.1 Enterprise Server Edition, which includes
 - Advanced Security [[Row and Column Access Control](#), Label Based Access Control]
 - [Time Travel Query](#)
 - [Multi-Temperature Data Management](#)
 - DB2 Advanced Enterprise Server Edition additionally includes
 - Storage Optimization Feature, including [Adaptive Compression](#)
 - [Continuous Data Ingest](#)
 - DB2 Workload Management, for monitoring and controlling workloads
 - Homogeneous Replication Feature for DB2 (restricted use [between 3 DB2 LUW systems](#))
 - InfoSphere Federation Server (restricted use between Oracle and DB2 LUW)
 - Tooling
 - IBM Data Studio
 - [InfoSphere Data Architect \(10 authorized users\) *](#)
 - [InfoSphere Optim Configuration Manager *](#)
 - [InfoSphere Optim Performance Manager Extended Edition \(add \[Extended Insight *\]\(#\)\)](#)
 - [InfoSphere Optim Query Workload Tuner *](#)
 - [InfoSphere Optim pureQuery Runtime *](#)

This color denotes new functions and products available in DB2 AESE 10.1 for LUW

* Optim tool currently not included in earlier versions

Low Operational
Costs

In Closing - Building On the Pillars of DB2



Low Operational Costs

- Adaptive compression
- Multi-Temperature Data Management
- Faster query response
- Improved index mgmt
- Real-time data warehousing

Ease of Development

- Temporal capabilities
- Row and Column Access Control
- SQL compatibility enhancements
- NoSQL graph store

Reliability

- DB2 pureScale enhancements
- Workload management enhancements
- HADR support extended to multiple standby servers