DB2® 10.5
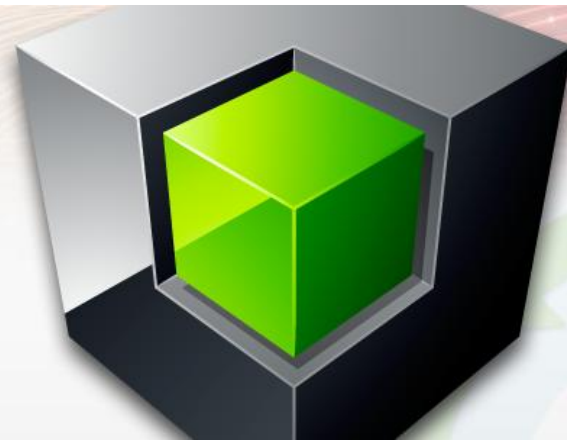
with BLU Acceleration

New

Craig Richardson
Technical Sales Specialist - Core Database
Mid-Atlantic Region

Disclaimer

Different Workloads Require Different Data Systems Where DB2 Plays in an Expert System

Mobile/Cloud Data Serving and Transaction Processing

Big Data Analytics

Mobile Storefront	E-commerce	Real Time Fraud Detection	Sales Analysis	Social Data Analysis

JSON Database	Transactional Database	Operational Data Warehouse	Analytics Data Warehouse	Distributed Map-Reduce System
Mobile Data Serving	**Transaction Processing**	**Operational Analytics**	**Reporting and Analytics**	**Hadoop Analytics**
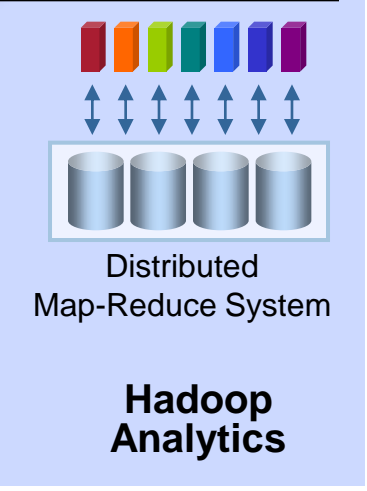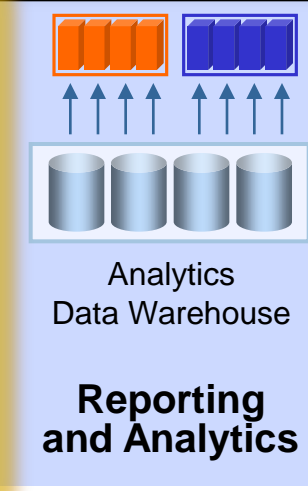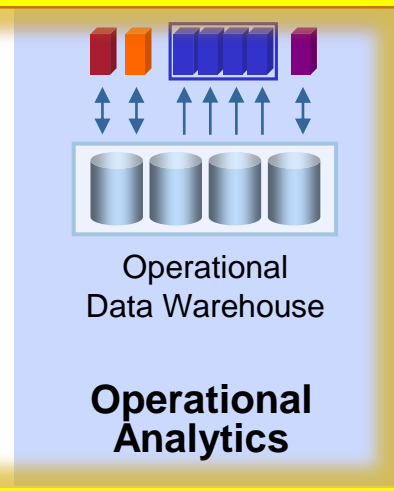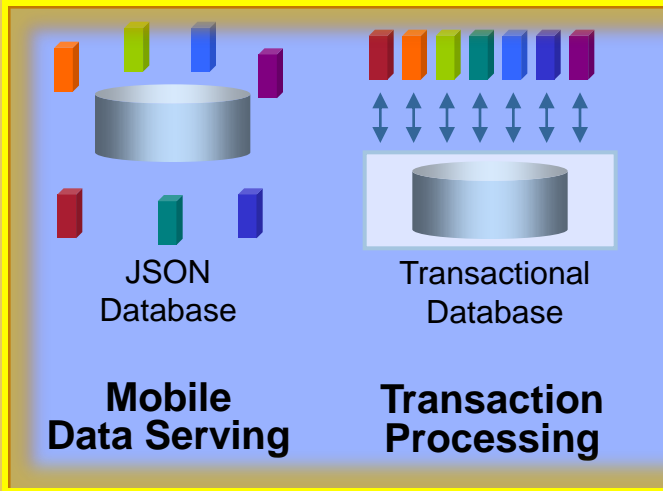
Different Workloads Require Different Data Systems Where DB2 Plays in a Software Solution

Mobile/Cloud Data Serving and Transaction Processing

Big Data Analytics

| Mobile Storefront | E-commerce | Real Time Fraud Detection | Sales Analysis | Social Data Analysis |

JSON Database · Transactional Database · Operational Data Warehouse · Analytics Data Warehouse · Distributed Map-Reduce System
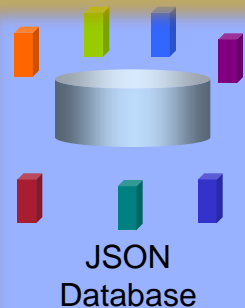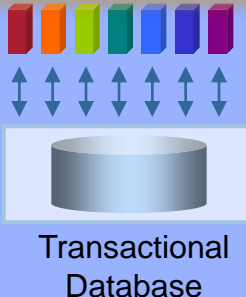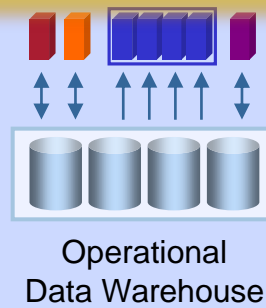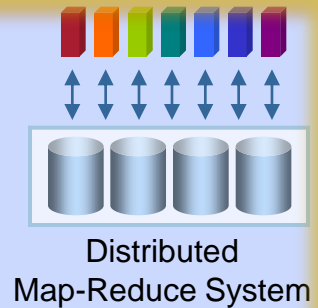
Mobile Data Serving **Transaction Processing** **Operational Analytics** **Reporting and Analytics** **Hadoop Analytics**

DB2® 10.5
with BLU Acceleration

New

Multi-workload database software for the era of Big Data

BLU Acceleration – Extreme performance and storage savings, leveraging dynamic "in-memory" and columnar technologies, for analytic processing

DB2 pure Scale – High availability, extreme scalability, and application transparency for OLTP workloads

NoSQL – Continue to support the next generation of applications

Oracle Application Compatibility – Continue to reduce the cost and risk associated with migrating Oracle applications to DB2

Enhanced Tooling - Reducing the total cost of ownership with DB2 and making the adoption, management, monitoring, and maintenance very simple

DB2 10.5 with BLU Acceleration

What is DB2 with BLU Acceleration?

- **New technology** for analytic queries in DB2 LUW
  - DB2 column-organized tables add columnar capabilities to DB2 databases
    - Table data is stored column organized rather than row organized
    - Using a vector processing engine
    - Using this table format with star schema data marts provides significant improvements to storage, query performance, ease of use, and time-to-value
  - New unique runtime technology which leverages the CPU architecture and is built directly into the DB2 kernel
  - New unique encoding for speed and compression
    - This new capability is both main-memory optimized, CPU optimized, and I/O optimized

DB2 with BLU

Runtime

| Classic DB2 runtime | BLU runtime |

| Classic DMS (for non-BLU tables) | BLU DMS (for BLU tables) |

Classic DB2 bufferpool

CPUs with SIMD

Storage

(classic row structured table)

(compressed, encoded columnar)
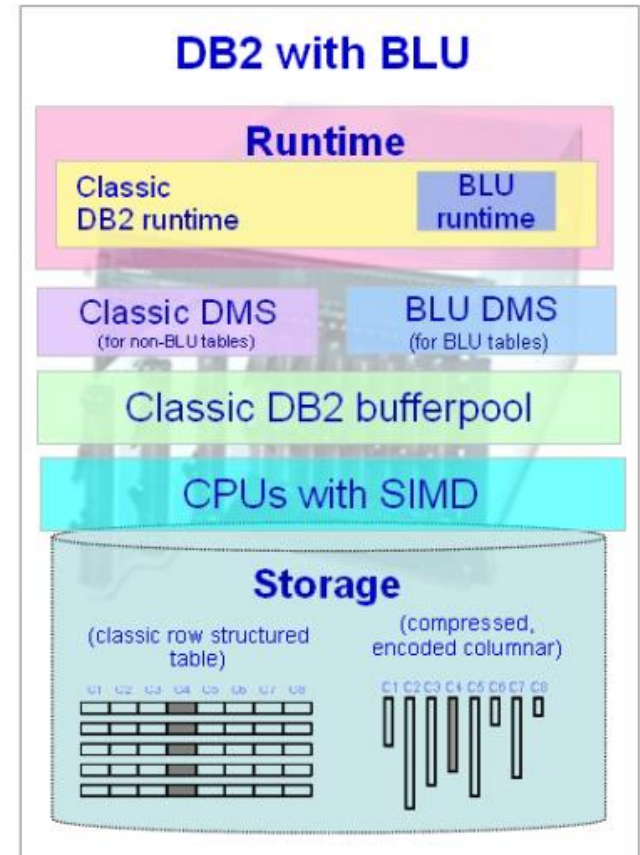
How Fast Is BLU Acceleration?

Customer	Performance Gains
BNSF	Up to 137x
Handelsbanken	7x – 100x
Triton Consulting	46x
Yonyou	40x
Coca-Cola Bottling	4x - 15x

10x-25x

speedup
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

Storage Savings

- **Multiple examples of data requiring substantially less storage**
  - 5% of the uncompressed size
  - Fewer objects required
- **Multiple compression techniques**
  - Combined to create a near optimal compression strategy
- **Compression algorithm adapts to the data**
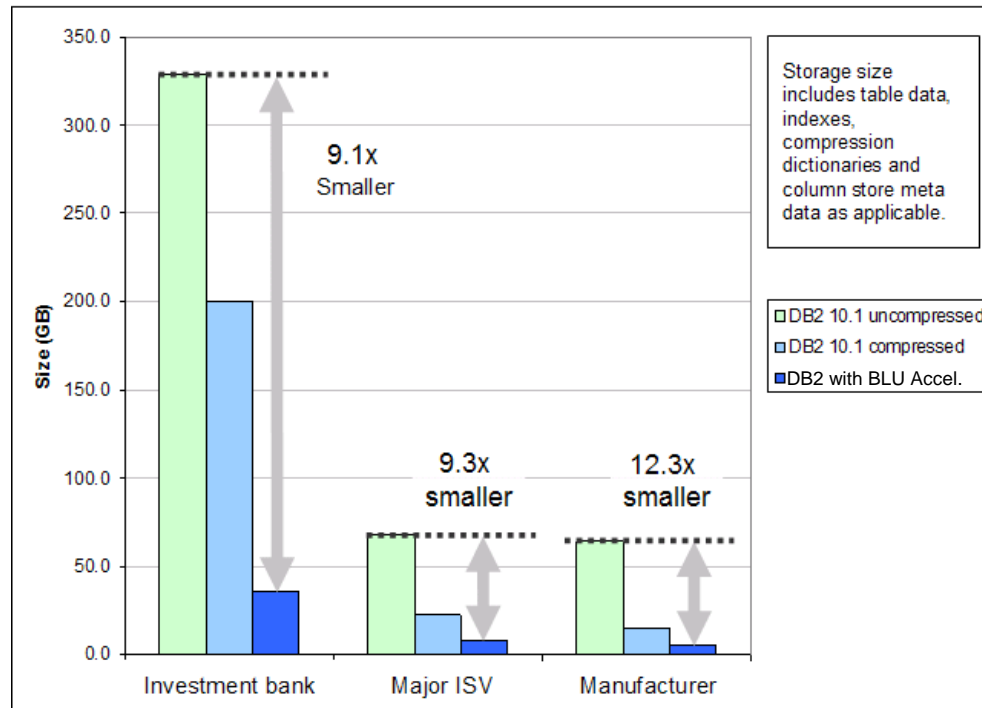
Seamless Integration into DB2

- **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 (`OQWT`) 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 (`WLM`) 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

Super Fast, Super Easy – Create, Load, and Go!

Database Design and Tuning

1. Decide on partition strategies
2. Select Compression Strategy
3. Create Table
4. Load data
5. Create Auxiliary Performance Structures
   - Materialized views
   - Create indexes
     - B+ indexes
     - Bitmap indexes
6. Tune memory
7. Tune I/O
8. Add Optimizer hints
9. Statistics collection

Repeat
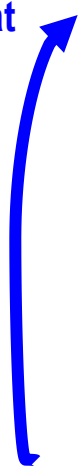
DB2 with BLU Acceleration

1. Create Table
2. Load data

The Seven Big Ideas of DB2 with BLU Acceleration
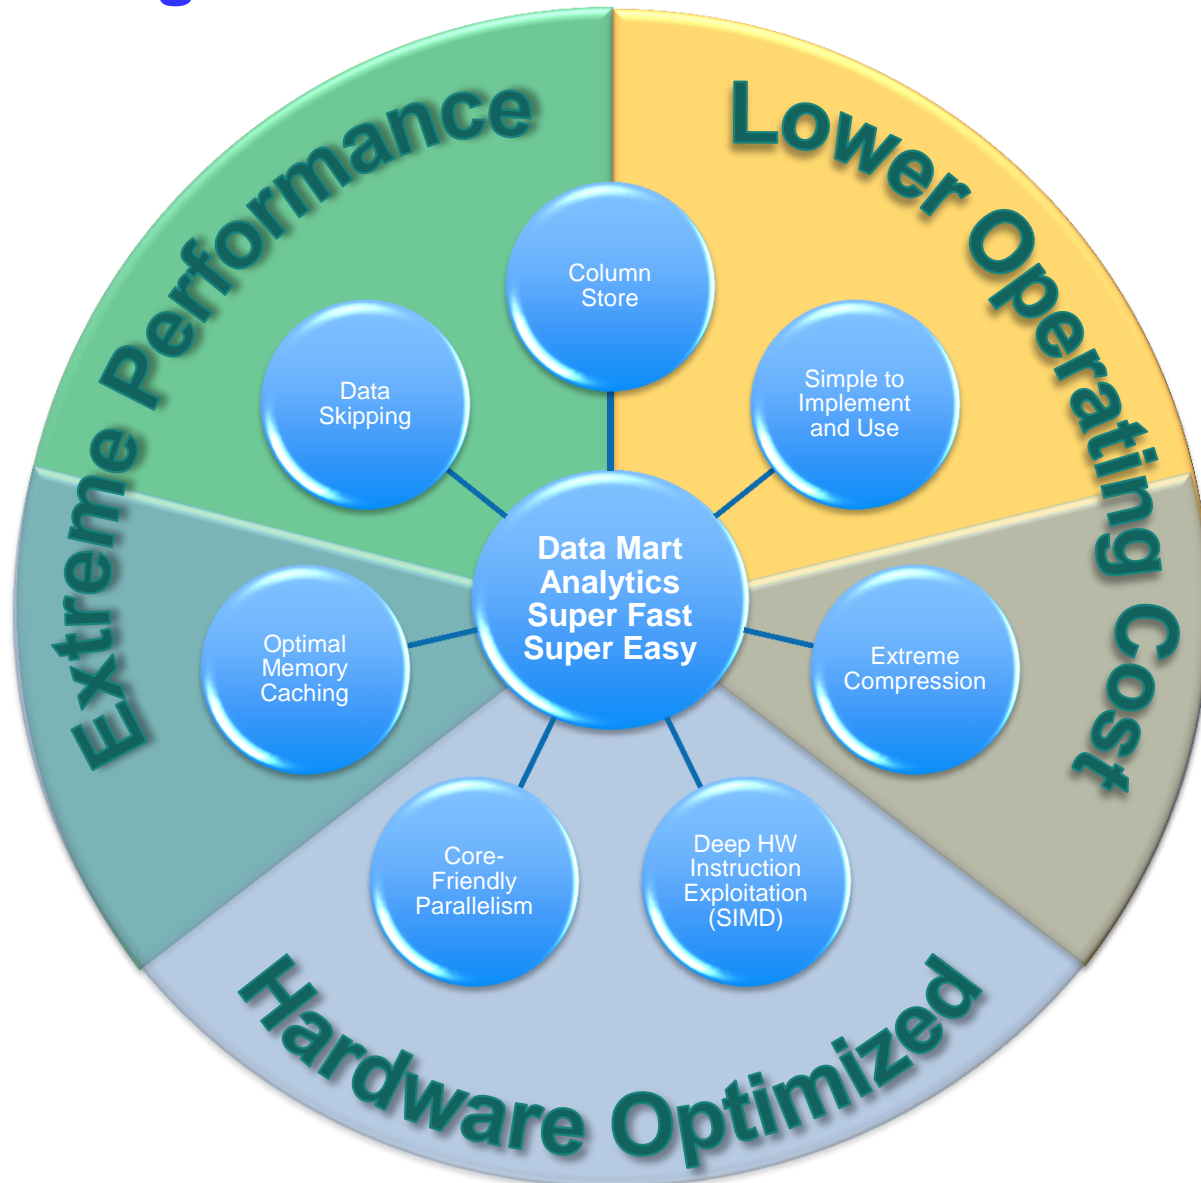
7 Big Ideas: ① Simple to Implement and Use

- **LOAD and then… run queries**
  - No indexes
  - No `REORG` (it's automated)
  - No `RUNSTATS` (it's automated)
  - No `MDC`
  - No `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

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

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

LAST_NAME Encoding

Packed into register length

LAST_NAME
Johnson
Smith
Smith
Smith
Smith
Johnson
Smith
Gilligan
Sampson
Smith

← Register Length →

← Register Length →
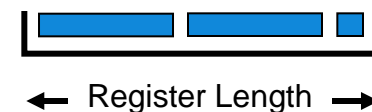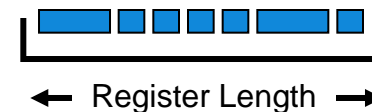
7 Big Ideas: ② Data Remains Compressed During Evaluation

- **Encoded values do not need to be decompressed during evaluation**
  - Predicates and joins work directly on encoded values

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

LAST_NAME   **Encoding**

Johnson

Smith

Smith

Smith

Smith

Johnson

Smith

Gilligan

Sampson

Smith

Encode

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

2007

Data

Instruction

Compare = 2005

Processor Core

Result Stream

2009　2010　2011　2012

Data

Instruction

Compare = 2005

Processor Core

Result Stream
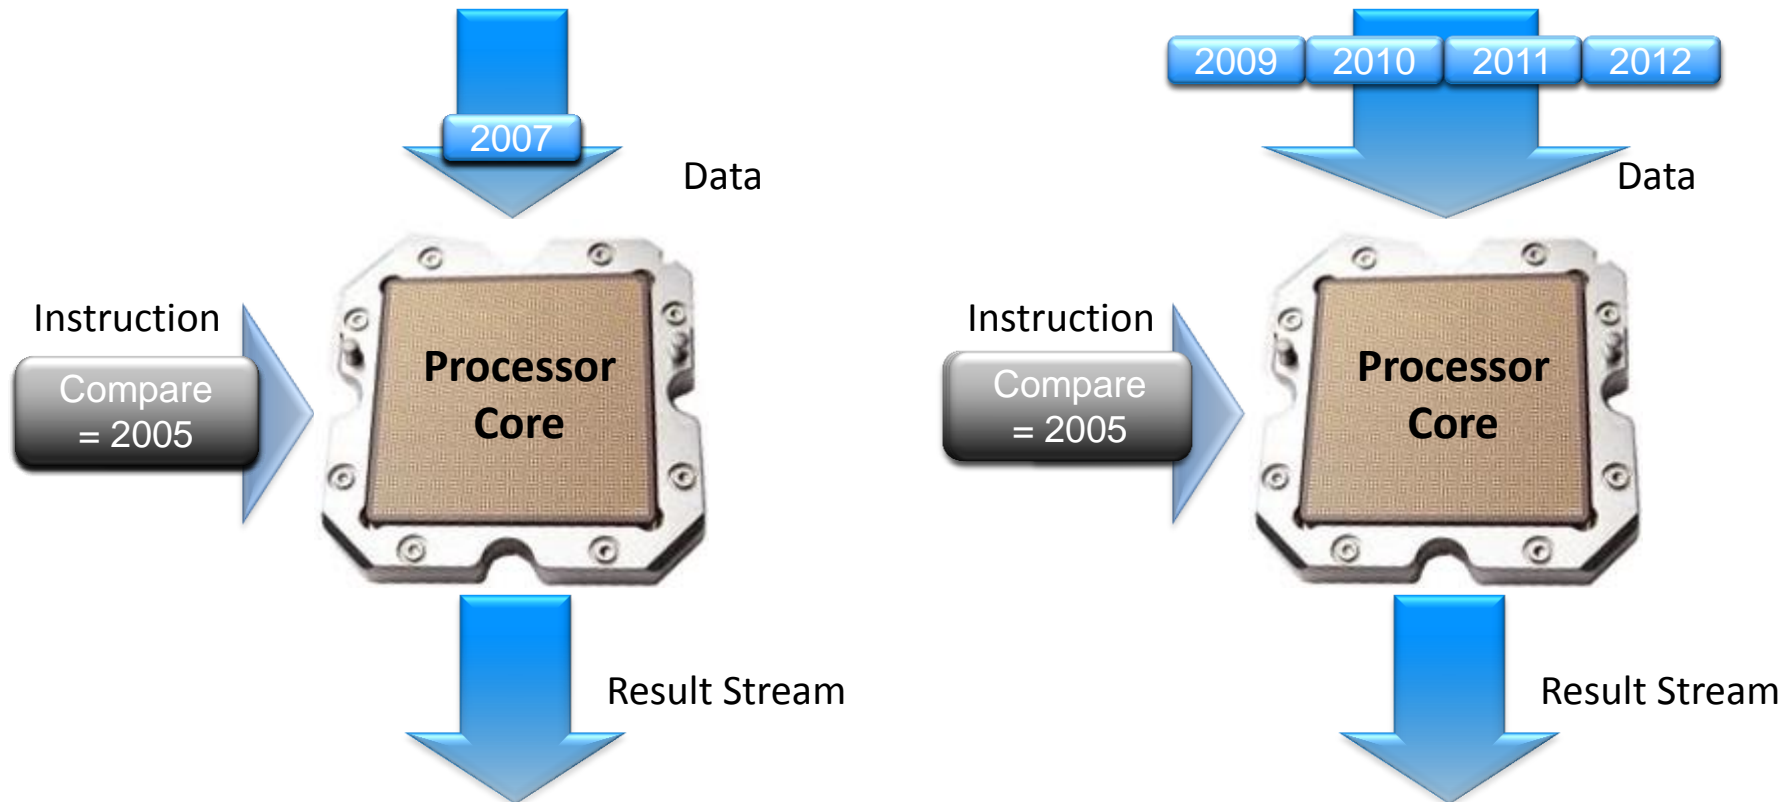
7 Big Ideas: ④ Core-Friendly Parallelism

- **BLU queries automatically parallelized across cores, and, achieve excellent multi-core scalability via …**
  - Careful data placement and alignment
  - Careful attention to physical attributes of the server
  - … designed to …
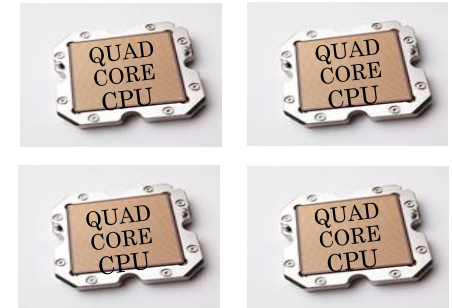
- **Maximize CPU cache, cacheline efficiency**

cache line

Cacheline 'ping-pong'

core 0 working on blue data

core 1 working on green data

Minimal Traffic
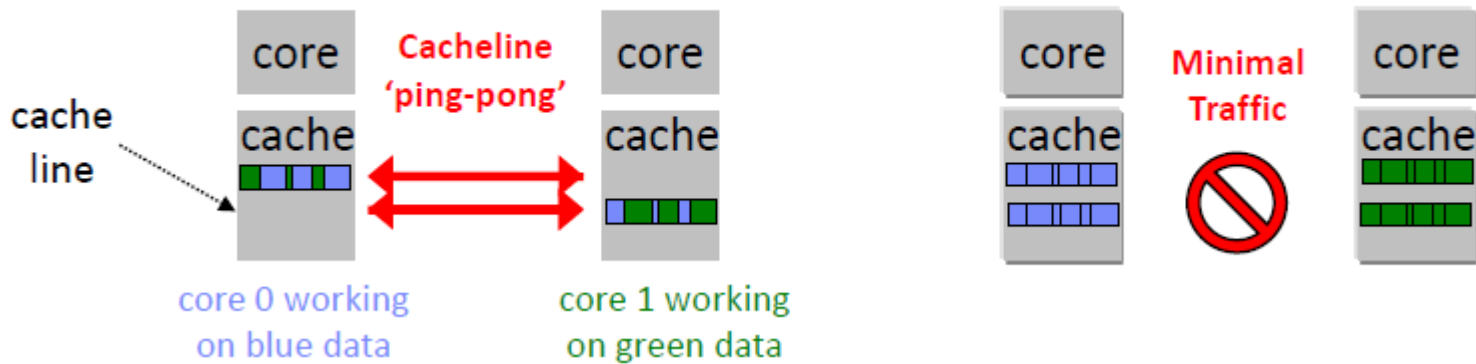
7 Big Ideas: ⑤ Column Store

- **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

C1 C2 C3 C4 C5 C6 C7 C8
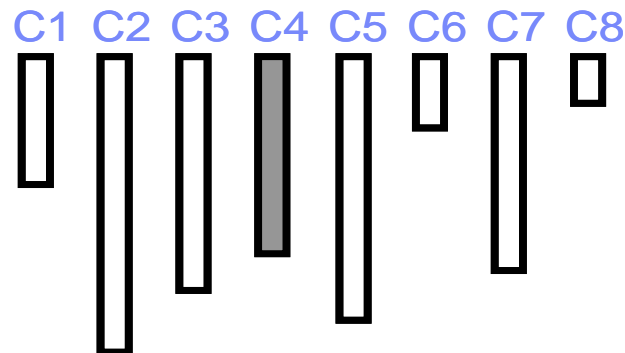
7 Big Ideas: ⑥ Scan-Friendly Memory Caching

- **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

RAM

Near optimal caching

DISKS
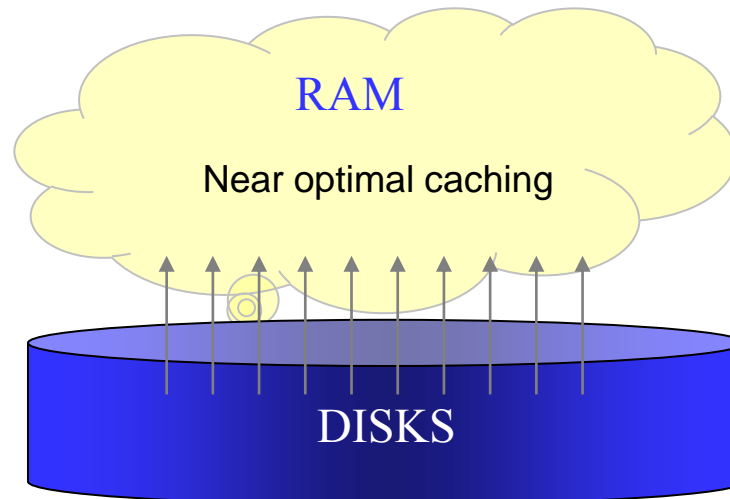
7 Big Ideas: (7) Data Skipping

- **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
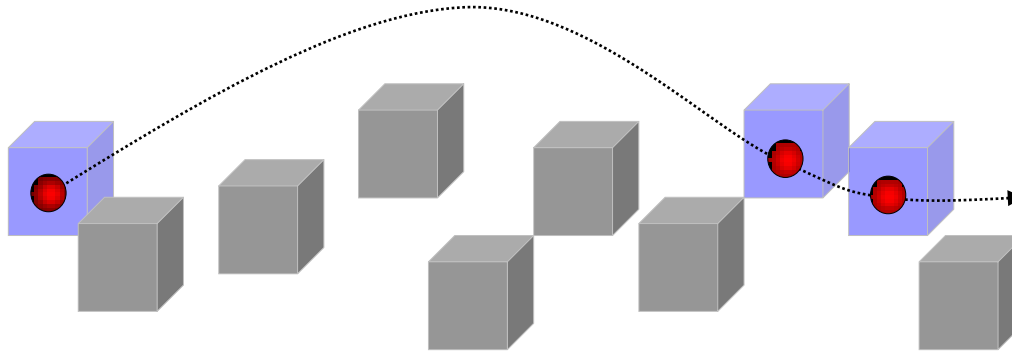
7 Big Ideas: How DB2 with BLU Acceleration Helps
~*Sub second 10TB query – An Optimistic Illustration*

- **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

DB2 WITH BLU ACCELERATION
Sub second 10TB query
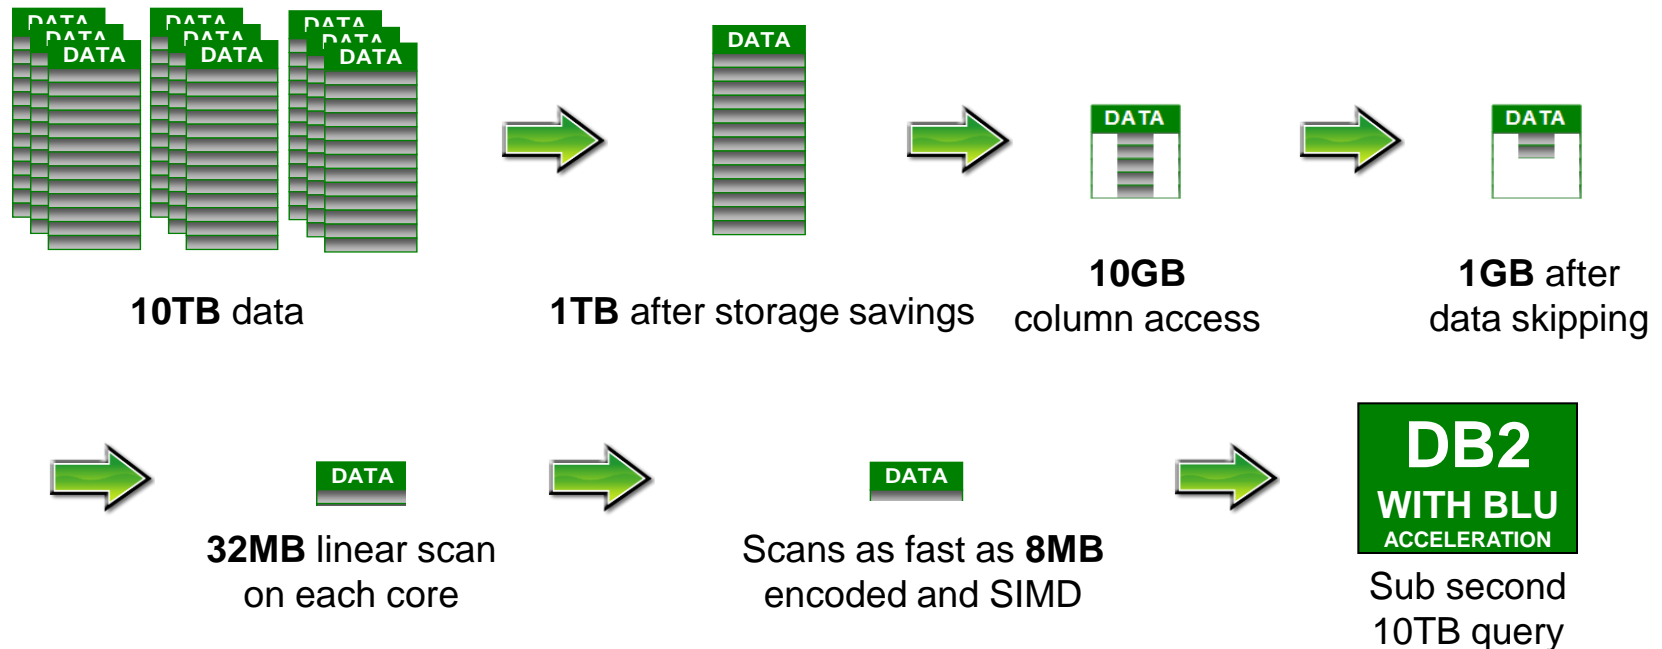
Unlimited Concurrency with "Automatic WLM"

- **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`**

Applications and Users

Up to tens of thousands of SQL queries at once
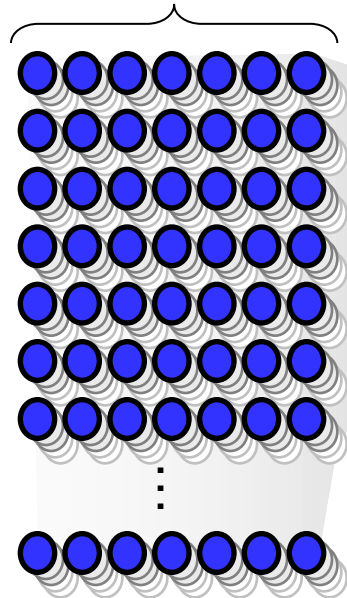
DB2 DBMS kernel

Moderate number of queries consume resources

SQL Queries
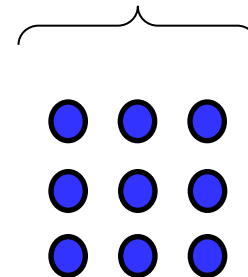
Automatic Space Reclaim

- **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**

Column1   Column2   Column3

Storage extent

2013   2013   2013

2012   2012   2013

2012

```
DELETE * FROM MyTable
    WHERE Year = 2012
```

These extents hold only deleted data

2012
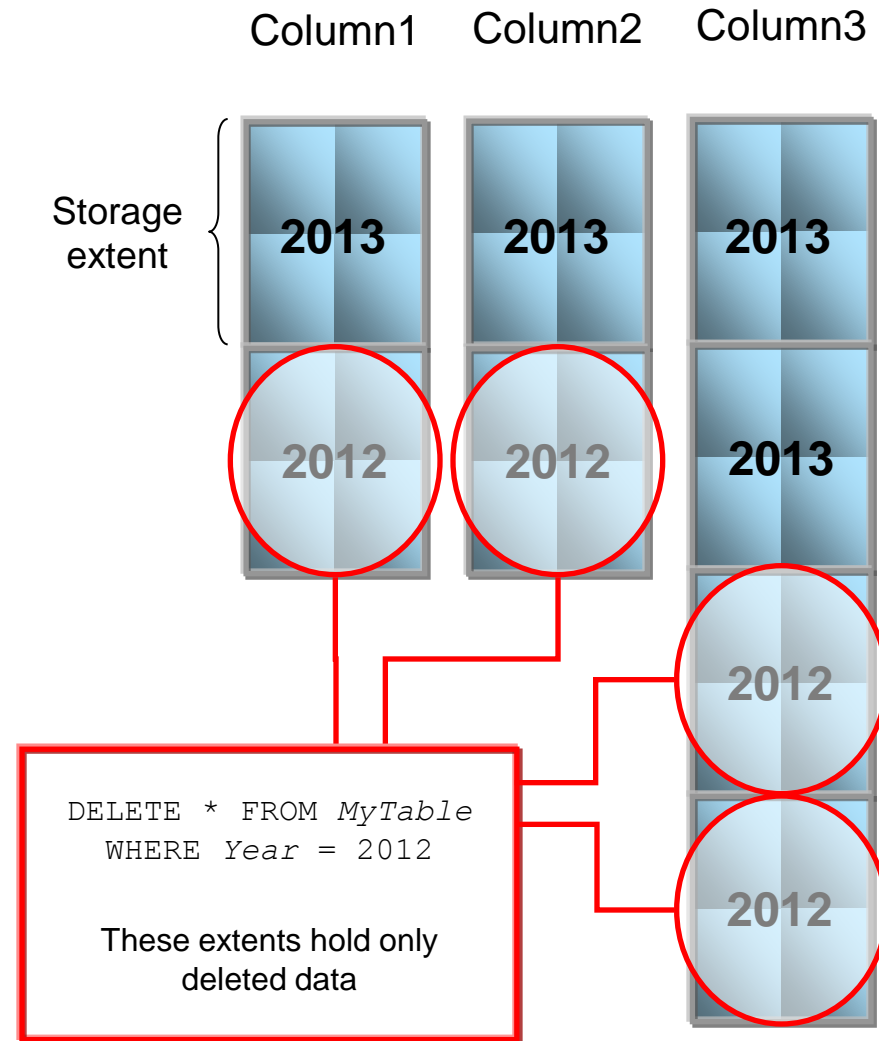
Performance Enhancements for Mixed Workloads

- **Added in DB2 10.5 for LUW FP1 the ability to do tactical updates, deletes, and reads in column-organized tables with key access**
  - Only for unique/primary key constraints defined on column-organized tables

- **`SELECT` statements can leverage unique constraints only with isolation level `UR` (uncommitted read)**

- **`DELETES` and `UPDATES` can use isolation levels `UR` or `CS` (cursor stability)**

- **This enhances our ability to support mixed workload environments … it may help non-SAP applications if they have updates on individual rows and qualified by primary key**
  - Useful in ETL processing where `UPDATES/DELETES` on the source are transferred to BLU tables (for example SAP BW, ETL jobs, or replication)

DB2 with BLU Acceleration Early Customer Quotes

BNSF *RAILWAY*

"Using DB2 10.5 with BLU Acceleration, our storage consumption went down by about 10x compared to our storage requirements for uncompressed tables and indexes. In fact, I was surprised to find a 3x increase in storage savings compared to the great compression that we already observed with Adaptive Compression on the DB2 10.5 server.*"*
- Kent Collins, Database Solutions Architect, BNSF Railway

Coca-Cola
BOTTLING CO.
CONSOLIDATED

"One of the things I really like about BLU Acceleration is that it enables me to put **column-organized tables beside row-organized tables in the same database.** *In our mixed environment, we realized an amazing* **10-25x reduction in the storage** requirements for the database when taking into account the compression ratios, along with all the things **I no longer need to worry about: indexes, aggregates, and so on.***"*
-Andrew Juarez, Lead SAP Basis and DBA

DB2 10.5 pureScale

DB2 10.5 pureScale Enhancements
Enhanced availability, optimized for OLTP Workloads

- **DB2 pureScale**
  - Robust infrastructure for OLTP workloads
  - Provides improved availability, performance, and scalability
  - Transparent scalability beyond 100 nodes[1]
  - Leverages z/OS cluster technology

- **NEW pureScale enhancements**
  - Online member add
  - No planned downtime required for pureScale maintenance updates[2]
  - HADR designed to failover in seconds[3]
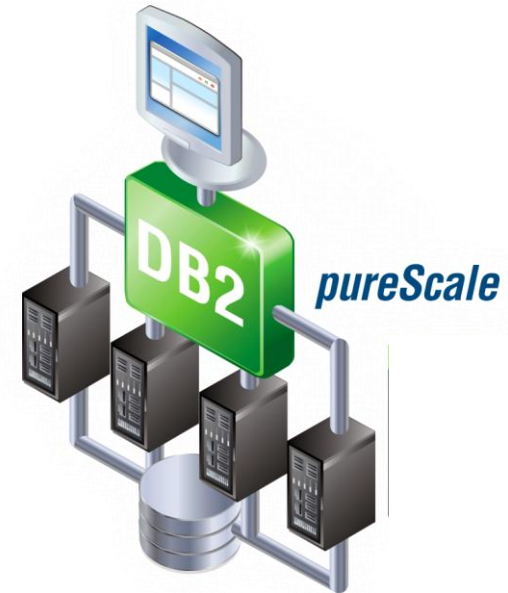
1. Available with DB2 Advanced Enterprise Server Edition.
2. Based on IBM design for normal operation with rolling maintenance updates of DB2 server software on a pureScale cluster. Individual results will vary depending on individual workloads, configurations and conditions, network availability and bandwidth.
3. Based on IBM design for normal operation under typical workload using HADR and pureScale clusters. Individual results will vary depending on individual workloads, configurations, and conditions, network availability and bandwidth.

HADR in DB2 pureScale

- **Integrated disaster recovery solution**
  - Simple to setup, configure, and manage

- **Support includes**
  - `ASYNC` and `SUPERASYNC` modes
    - `SYNC`/`NEARSYNC` under development
  - Time delayed apply
  - Log spooling
  - Both non-forced (role switch) and forced (failover) takeovers

Primary Cluster   HADR   **Standby DR Cluster**
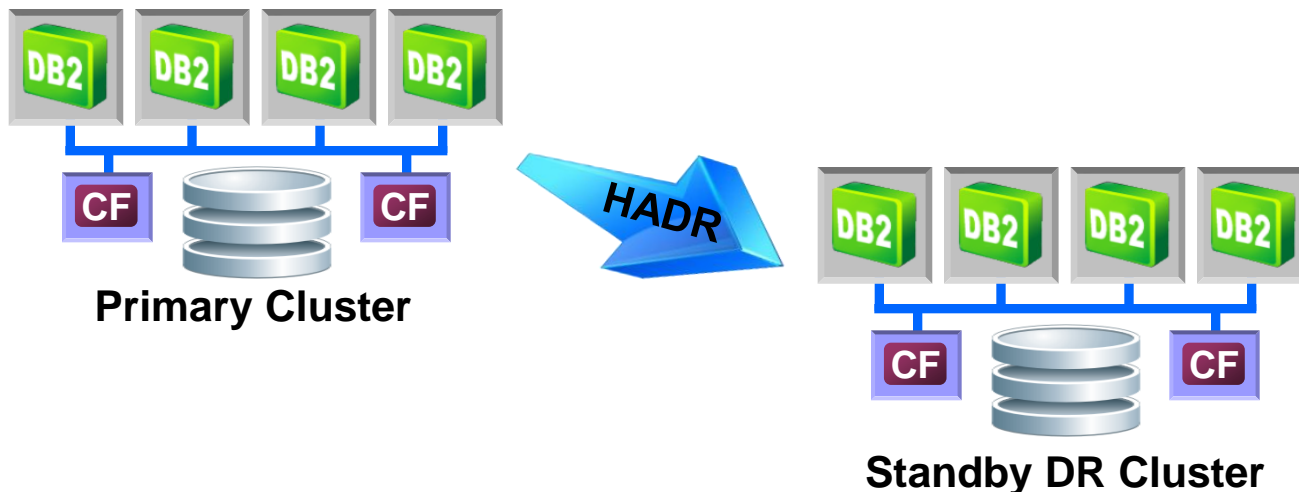
HADR in DB2 pureScale: Example

Primary site

Transactions

Member

Member

CF

Logs 1   Logs 2   Logs 3

TCP/IP

Failed member's logs

Standby site

Replay member

Member

Member

CF
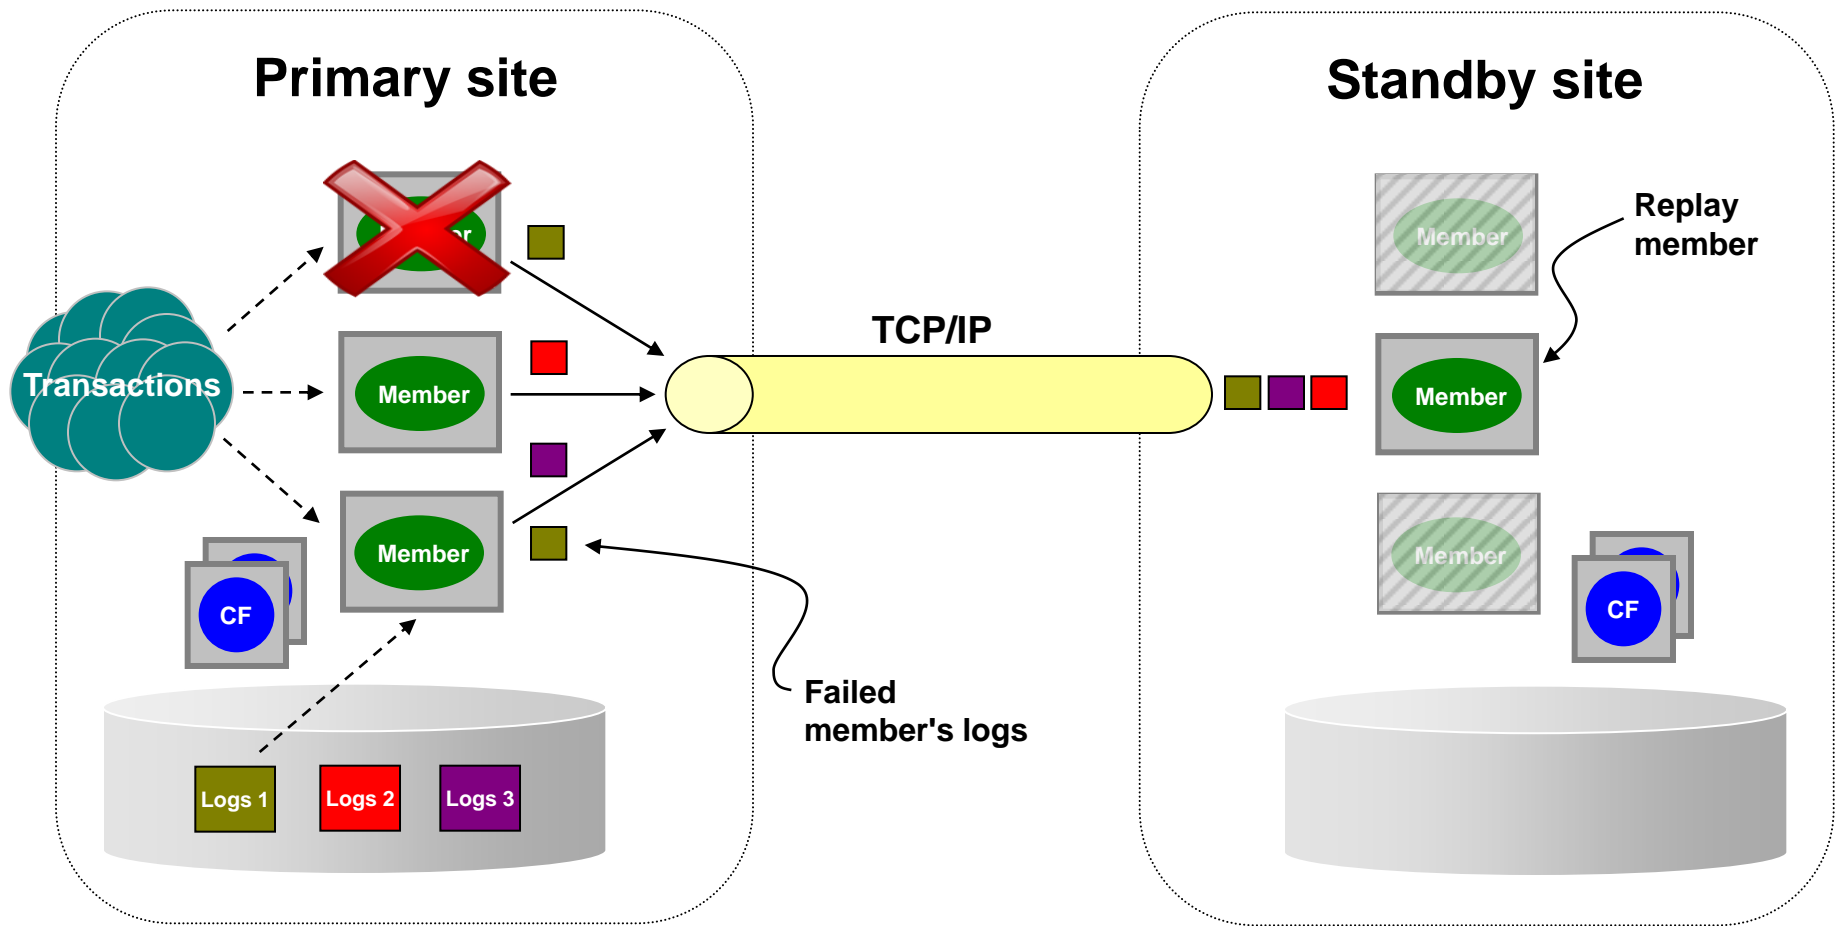
Rolling Fix Pack Updates

- **DB2 pureScale fix packs can be applied in an online rolling fashion**
  - Transparently install DB2 pureScale fix packs with no outage

- **New options for `db2iupdt` to do to online update, do a pre-commit check, and to subsequently commit the changes**

- **Includes updates of CFs and members**

Rolling Fix Pack Updates – Example

C
C

Code level: ~~GA~~ FP1
Code level: ~~GA~~ FP1

```
1. db2stop member1
   quiesce

2. db2iupdt
   -online_update
   member1 inst1

3. db2start member1
```

```
4. db2stop member2
   quiesce

5. db2iupdt
   -online_update
   member2 inst1

6. db2start member2
```

Member 1
Member 2

```
7. db2iupdt -commit_new_level inst1
```

Two member cluster (each at GA level) with clients (C) connecting into each member

1. **Member 1 is quiesced – clients all move to Member 2**

2. **DB2 binaries updated on Member 1**

3. **Member 1 started again and a portion of the clients get rerouted to Member 1 to balance the workload**

4. **Member 2 is quiesced – clients all move to Member 1**

5. **DB2 binaries updated on Member 2**

6. **Member 2 started again and a portion of the clients get rerouted to member 2 to balance the workload**

At this point, code is at FP1 level, but can't use any new FP1 features; can test stability and roll down to GA level if necessary

7. **Updates are committed**

The instance is now completely running at FP1 and new features can be used; cannot roll down to GA any longer.

Multi-Tenancy: Member Subsets

- **Previously, an application/tenant could only be configured to run**
  1. On one member (client affinity) or
  2. Across all members in cluster (workload balancing)

- **Can now point applications to subsets of members which enables**
  - Isolation of batch from transactional workloads
  - Multiple databases in a single instance to be isolated from each other
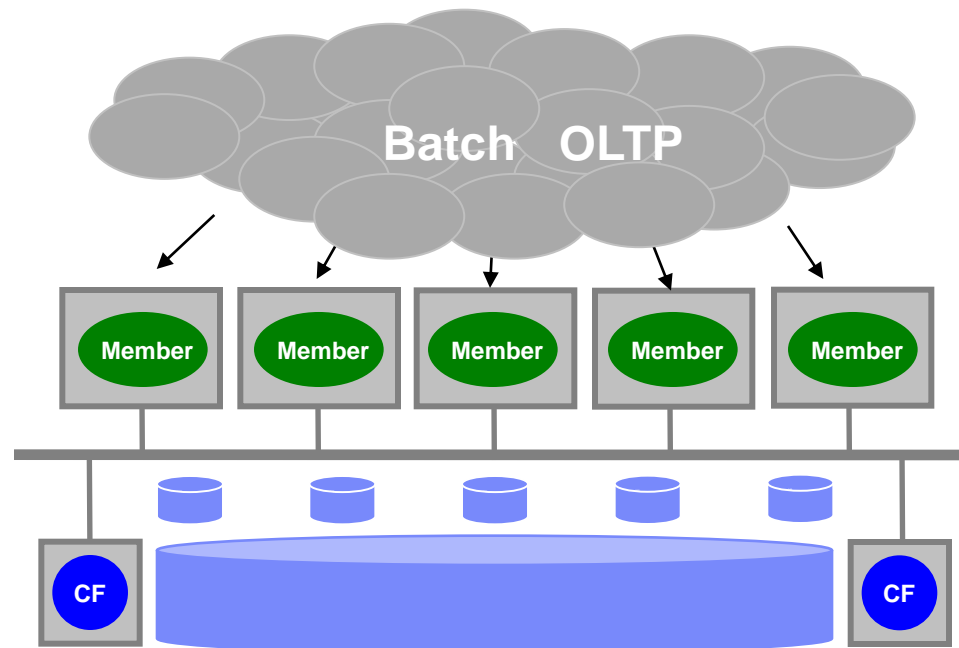
Multi-Tenancy: Self-Tuning Memory Management (STMM)

- **Prior DB2 pureScale STMM design**
  - Single tuning member makes local tuning decisions based on workload running on that member
    - Other member becomes tuning member in case of member failure
  - Broadcasts tuning decisions to other members
  - Works well in single homogeneous workload scenarios

- **DB2 pureScale now allows per-member STMM tuning**
  - Workload consolidation
  - Multi-tenancy
  - Batch workloads
  - Affinitized workloads
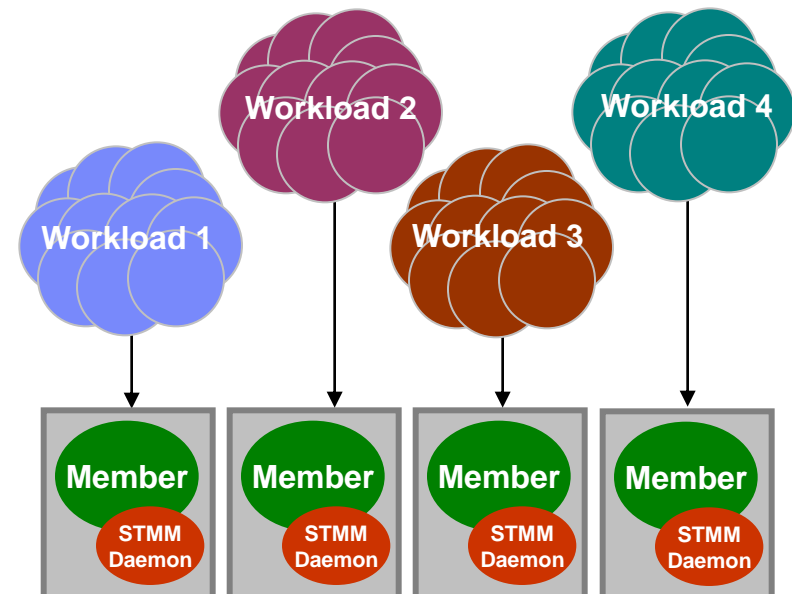
Online Add Member

- **New members can be added to an instance while it is online**
  - No impact to workloads running on existing members
  - Previously, required an outage of the entire instance to add a new member

- **No change in add member command**

```
db2iupdt –add –m <newHost> -mnet <networkName> <instance>
```

- **Offline backup no longer needed** after adding new members

Member added online
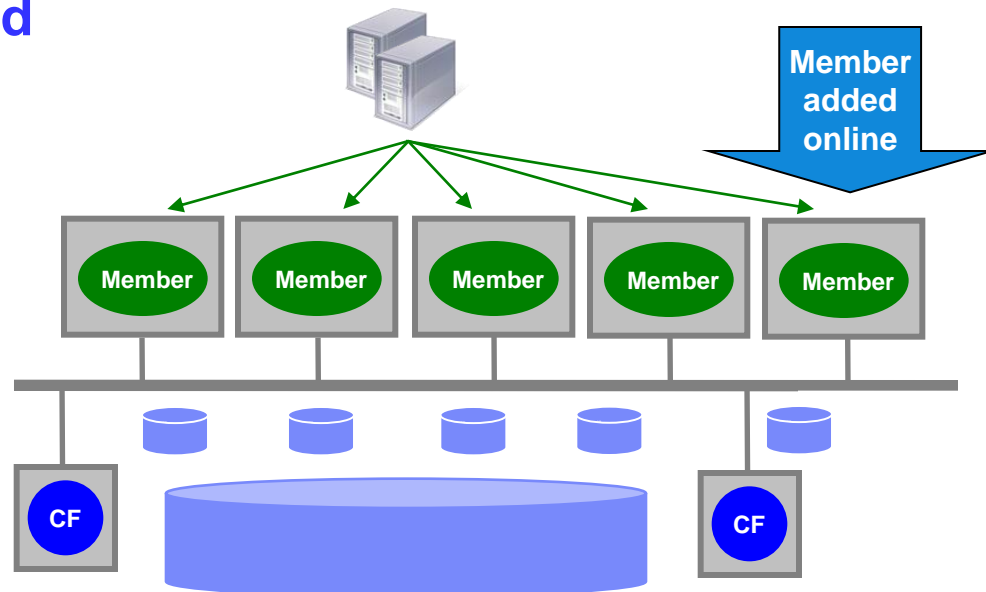
Topology-Changing Backup and Restore

- **Backup and restore between topologies with differing numbers of members**

- **Backup and restore from DB2 pureScale to non-DB2 pureScale (and vice-versa)**
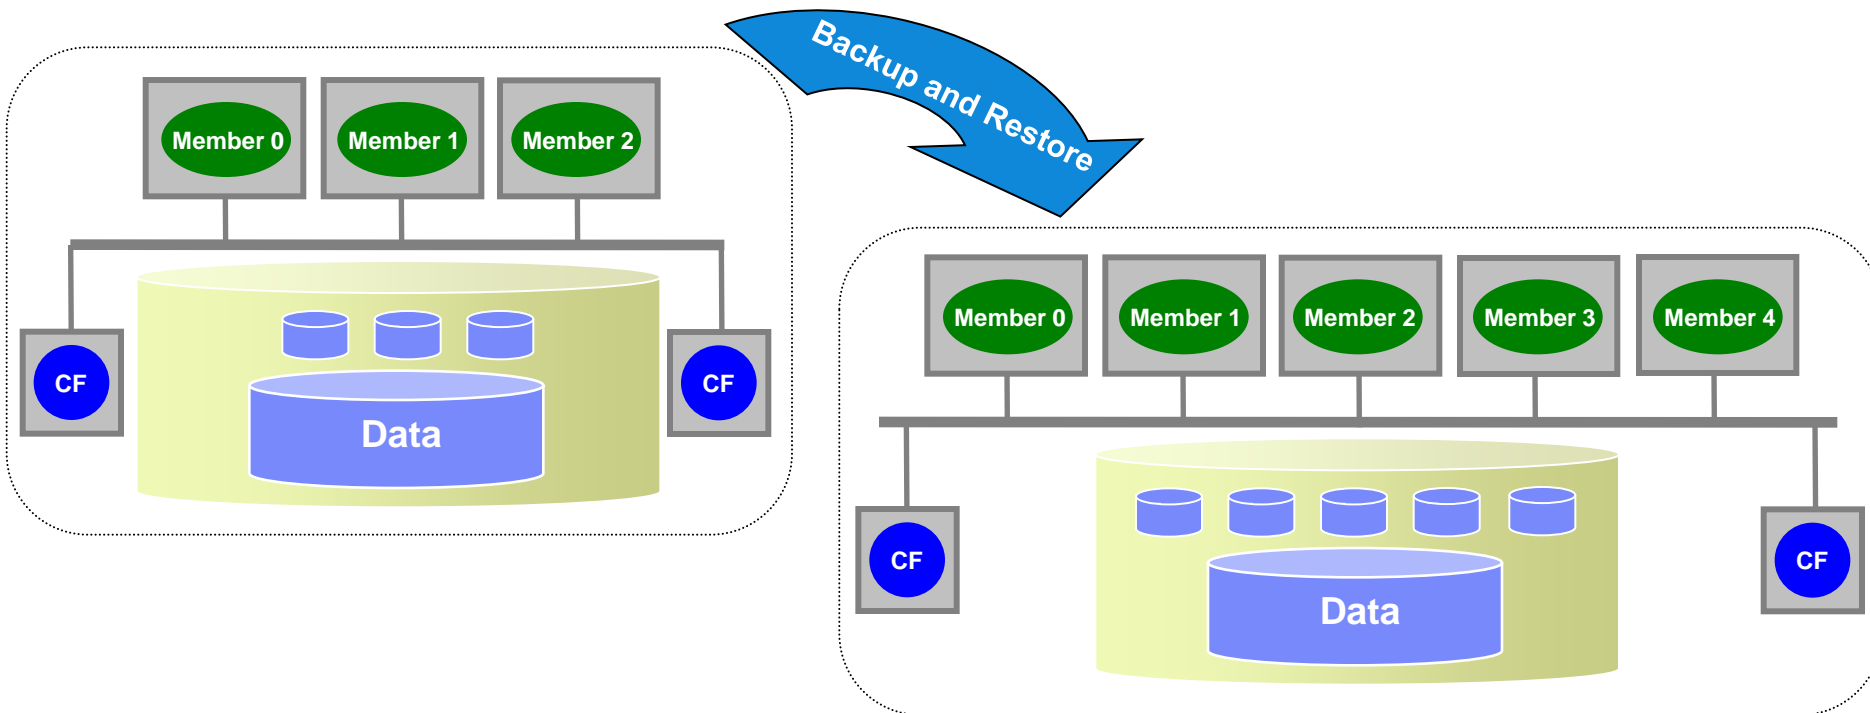
Multi-Tenancy: Explicit Hierarchical Locking (EHL)

- **Designed to remove data sharing costs for tables/partitions that are only accessed by a single member**
  - Avoids CF communication if object sharing not occurring

- **Target scenarios**
  - Workload affinitization
  - Workload consolidation and application affinitization

- **Enabled via new `OPT_DIRECT_WRKLD` database configuration parameter**
  - Detection of data access patterns happens automatically and EHL will kick in when data is not being shared after configuration parameter set

No member/CF communication necessary
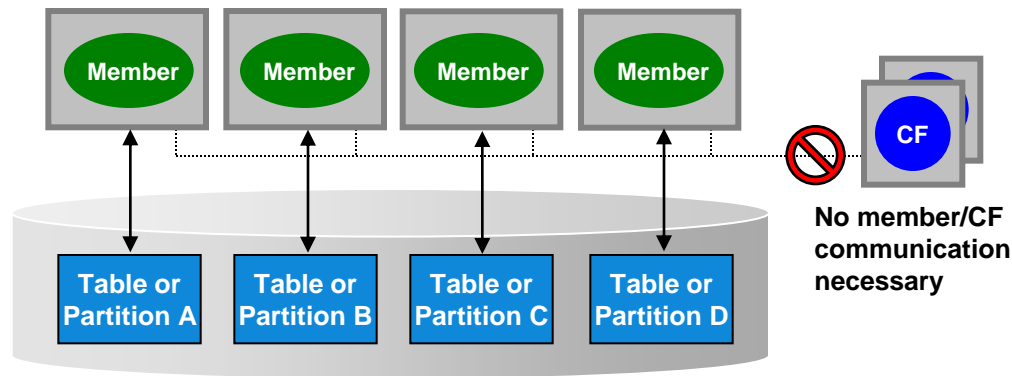
DB2 pureScale Feature Support for Non-IBM x86/Intel Servers

- **Previously, x86/Intel server support limited to specific IBM hardware**

- **DB2 pureScale feature can now run on non-IBM rack mounted servers provided that they support one of the following InfiniBand QDR (Quad Data Rate) or Ethernet RoCE (RDMA Over Converged Ethernet) network adapters**
  - Mellanox ConnectX-2 EN 10 Gigabit Ethernet adapters with RoCE
  - Mellanox ConnectX-2 with Virtual Protocol Interconnect
  - MT27500 - Mellanox ConnectX-3 EN Dual-port SFP+ 10GbE adapter
  - MT27500 - Mellanox ConnectX-3 VPI QSFP Dual-Port InfiniBand card

- **Geographically Dispersed pureScale Cluster (GDPC) configurations still require IBM servers**

- **Applies to DB2 10.5 Fix Pack #1 and later, and DB2 10.1 Fix Pack #2 and later**

DB2 10.5 - Online Reorganization Enhancements
Non-DB2 pureScale configurations

- **Online inplace reorg can be done on a table using adaptive compression**

- **Fastpath option for online inplace reorg to clean up overflow records only**

```
$ db reorg table T1 inplace cleanup overflows
```

- **Reorg with RECLAIM EXTENTS can cleanup partially empty extents for ITC tables**

DB2 10.5 Text Search

- **Enhanced versions of command line tools**
  - `Config` tool has enhanced configuration capabilities
  - `Admin` tool has deprecated the trace parameter
  - `Synonym` tool now allows you to specify the language for which synonyms are defined

- **Search enhancements**
  - Cannot search for an exact match on two consecutive, identical special characters
  - Enhancements to structural full-text search in XML documents

- **Additional options to control the upgrade index processing behavior**

- **Additional settings to allow more control over commit batches**

- **New index configuration options**

- **In FP1 you can create a DB2 Text Search index on a nickname**
  - For details, see *DB2 Text Search indexes on nicknames* in the Information Center

Miscellaneous Fix Pack #1 Additions

- **LOB types (`BLOB`, `CLOB`, and `DBCLOB`) are now supported as input to expression-based index keys**
  - For details, see the *`CREATE INDEX` statement* in the Information Center

- **The `DB2_COLLECT_TS_REC_INFO` registry variable is now set to `OFF` by default for high availability disaster recovery (HADR) databases**
  - For details, see *Some registry and environment variables have changed* section in the Information Center

DB2 10.5 Oracle Compatibility

Oracle Compatibility Built into DB2
Lower Transition Cost and Less Risk

Concurrency Control	→	**Native support**
Oracle SQL dialect	→	**Native support**
PL/SQL	→	**Native support**
PL/SQL Packages	→	**Native support**
Built-in package library	→	**Native support**
Oracle JDBC extensions	→	**Native support**
OCI	→	**Native support**
Oracle Forms	→	**Through partners**
SQL*Plus Scripts	→	**Native support**
RAC	→	**DB2 pureScale**

Changes are the exception. Not the rule.

Application Compatibility Over Time
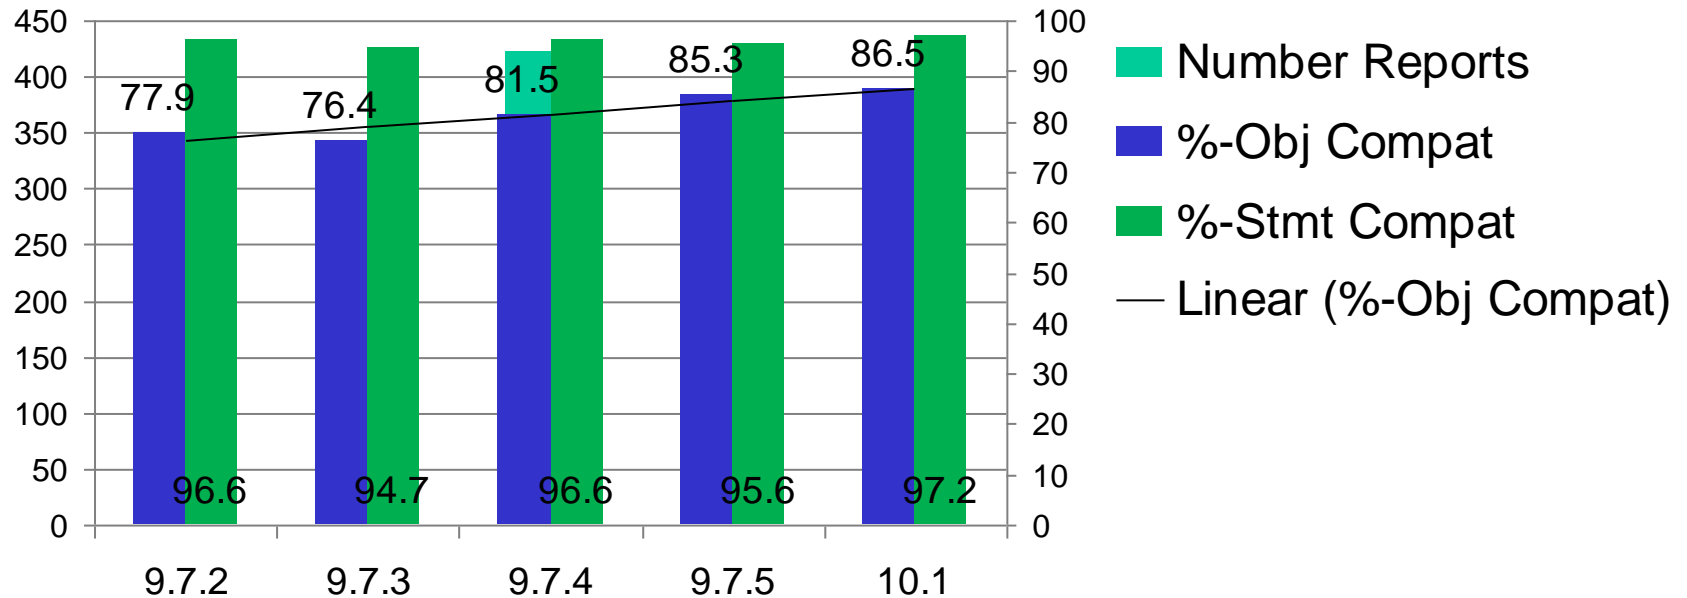
- **Data is based on DCW (Database Conversion Workbench) DB2 reports in the database**

- **Compatibility is improved**
  - More and more complex applications

- **DB2 10.5 provides > 98% compatibility**

Oracle Compatibility: Larger Row Widths

- **Accommodate larger strings**
  - Allow tables with up to 1MB wide rows
    ```
    CREATE TABLE emp(name     VARCHAR(4000),
                     address  VARCHAR(4000),
                     cv       VARCHAR(32000))
    ```
  - Allow large row GROUP BY and ORDER BY as long as key can sort

DB2 10.1
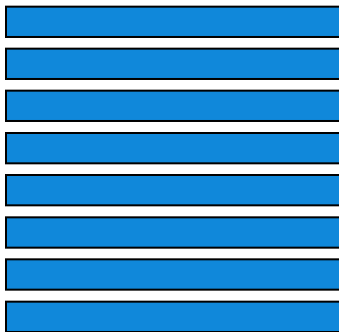
DB2 10.5

◄— Max 32K —►

◄——————————— Max 1M ———————————►

Oracle Compatibility: Additional Indexing

- **Function-based indexes**
  - Searching for computed values in a table instead of using Generated Columns
  - E.g. "Find employees without worrying about the case of their names"
    - ```
      CREATE INDEX emp_name ON emp(UPPER(name));
      SELECT salary
        FROM emp
       WHERE UPPER(name) = 'MCKNIGHT';
      ```

Name	Salary	Manages
McKnight	50000	Sales
Miller	25000	-
Van Gogh	45000	Finance
Chan	37000	-

- **Indexes excluding NULL keys**
  - Enforce uniqueness only for non-NULL keys and exclude all NULL keys from Index
  - Compress index for all-NULL keys
  - Helps facilitate Oracle application migrations
    - ```
      CREATE UNIQUE INDEX emp_manages
          ON emp(manages) EXCLUDE NULL KEYS
      ```

- **Random key indexes**
  - Avoid hot index page for incrementally issued keys
    - ```
      CREATE UNIQUE INDEX order_id ON order(id RANDOM);
      ```

Oracle PL/SQL Compatibility

- **Create distinct type with weak type rules**
  - Removes limitation of existing distinct types not having weak typing
  - Optional check constraint
  - Optional `NOT NULL` constraint
  - Constraints enforced on assignment

- **Pipelined table function**
  - Introduce a new `PIPE` statement which returns a row to caller, but continues at next statement if caller wants another row
  - Incrementally produce a result set for consumption on demand

- **Ad-hoc federated table access**
  - Support ad-hoc reference to remote table using server in the identifier
    - Reach out to a table in a remote database

- **Function library extensions**
  - Updates to various built-in functions for improved compatibility support

JSON Technology

Background – What is NoSQL

- **A class of database management systems that depart from traditional RDBMSs**
  - Does not use SQL as the primary query language
  - Is "schema-less"
    - No rigid schema enforced by the DBMS
  - Programmer-friendly for adding fields to a document
  - Might not guarantee full `ACID` behavior
  - Often has a distributed, fault-tolerant, elastic architecture
  - Highly optimized for retrieve and append operations over great quantities of data
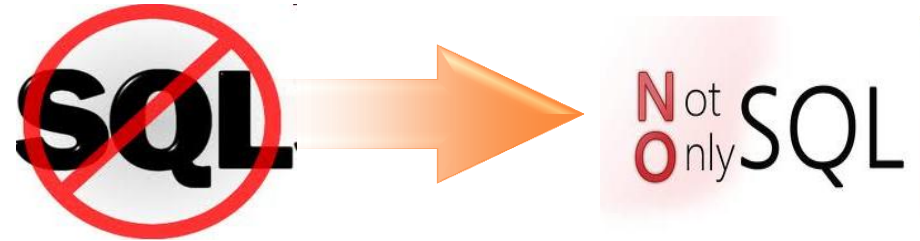
NoSQL DEFINITION: Next Generation Databases mostly addressing some of the points: being **non-relational, distributed, open-source** and **horizontally scalable**.

The original intention has been **modern web-scale databases**. The movement began early 2009 and is growing rapidly. Often more characteristics apply such as: **schema-free, easy replication support, simple API, eventually consistent / BASE (not ACID)**, a **huge amount of data** and more. So the misleading term *"nosql"* (the community now translates it mostly with **"not only sql"**) should be seen as an alias to something like the definition above. [based on 7 sources, 14 constructive feedback emails (thanks!) and 1 disliking comment . Agree / Disagree? Tell me so! By the way: this is a strong definition and it is out there here since 2009!]

LIST OF NOSQL DATABASES [currently 150]

Emergence of a growing number of non-relational, distributed data stores for massive scale data

The JSON-XML Shift

- **Developers find it easier to move data back and forth without losing information in JSON vs. XML**
  - XML is more powerful and more sophisticated than JSON
  - But JSON found to be 'good enough" → It makes programming tasks easier

- **By the time RDBMS world got very sophisticated with XML, developers had chosen JSON**
  - Application shift lead to emergence of database that store data in JSON (i.e., MongoDB)
  - JSON on the server side is appealing for developers using JSON on the client tier side

json

The Technology

- ☐ Each technology is awesome & powerful in itself But…
- ☐ …Mixing technologies and solving integration issues are not

Client Tier	Server Tier	Data Tier
HTML	CGI/BIN	RDBMS
HTML/AJAX-XML/Flex-XML	PHP/JAVA/XML	RDBMS+XML
HTML5/AJAX/JSON	PHP/JAVA/XML+JSON	RDBMS+XML+JSON
HTML5/AJAX/JSON	PHP/JAVA/JSON	JSON

Lost structure

Rigidity

mongoDB

Flexibility
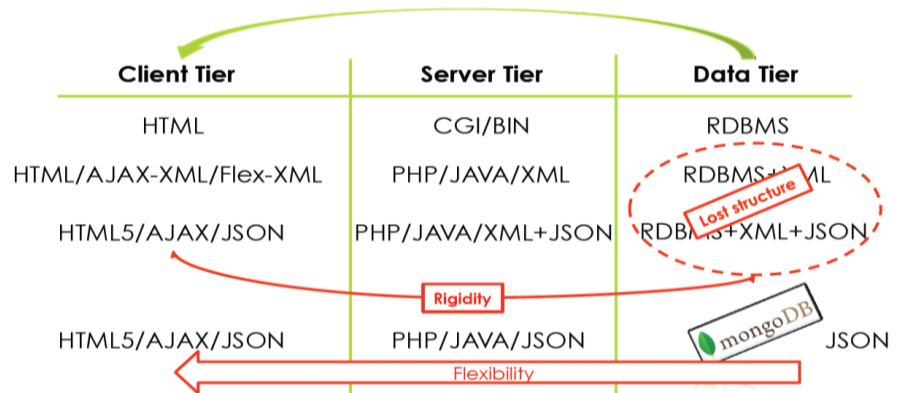
Open APIs State of the Market

- **JSON is the new cool**
  - XML declining: 5 years ago hardly any JSON

- **Why? JSON is**
  - Less verbose and smaller docs size
  - <Mytag>value</Mytag> vs. Mytag:value
  - Tightly integrated with JavaScript which has a lot of focus
  - Most new development tools support JSON and not XML

JSON Rising: 55% of all new APIs support JSON

Percentage of new APIs with JSON support
Based on directory of 3,200 web APIs listed at ProgrammableWeb, May 2011

HN Trends
See what hackers have been talking about.

- mongodb
- hbase
- dynamodb
- oracle nosql

Hacker Mentions (adjusted for growth)

Highcharts.com

20% of all new APIs support *only* JSON

Percentage of new APIs with only JSON support
Based on directory of 3,200 web APIs listed at ProgrammableWeb, May 2011

JSON Technology in DB2 for LUW

- **Combine data from systems of engagement with traditional data in same DB2 database**
  - Best of both worlds
  - Simplicity and agility of JSON + enterprise strengths of DB2

Cloud Mobile Social Big Data Analytics

- **Store data from web/mobile apps in it's native form**
  - New web applications use JSON for storing and exchanging information
  - It is also the preferred data format for mobile application backends

DB2

- **Move from development to production in no time!**
  - Ability to create and deploy flexible JSON schema
  - Gives power to application developers by reducing dependency on IT; no need to pre-determine schemas and create/modify tables
  - Ideal for agile, rapid development and continuous integration

AGILE METHODOLOGY
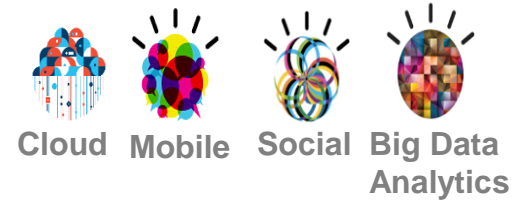
JSON Technology in DB2 for LUW (cont.)

- **DB2 for Linux, UNIX, and Windows now officially supports JavaScript Object Notation (JSON) DB2 NoSQL capability**
  - No longer in technology preview
  - You can now store and manage JSON data in a DB2 database
    - You can create dynamic applications by using JSON's schemaless NoSQL capability. In addition to basic NoSQL operations on collections of JSON documents, this release includes support for transactions control and bi-temporal data awareness
    - JSON documents can be interfaced in the following three ways
      - ❑ DB2 JSON Java API
      - ❑ DB2 JSON command-line interface
      - ❑ DB2 JSON wire listener
  - For further details, see *JSON application development support has been added* section in the Information Center

Cloud   **Mobile**   **Social**   **Big Data Analytics**

AGILE METHODOLOGY
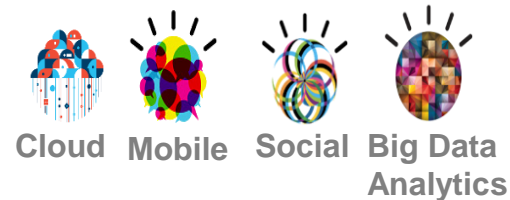
DB2 10.5 Packaging Simplification

DB2 10.5 Simplifies Product Packaging

One Set of Editions for Both Transactional and Warehouse Workloads

Departmental Market	**Enterprise Market**
DB2 Advanced Workgroup Server Edition	**DB2 Advanced Enterprise Server Edition**

Advanced function

DB2 Advanced Workgroup Server Edition

- For small OLTP and analytic deployments
- Primarily used in department environments within large enterprises or SMB/MM deployments
- Limited by TB, memory, sockets and cores
- Supports BLU, pS and DPF deployment models

DB2 Advanced Enterprise Server Edition

- For Enterprise Class OLTP and/or analytic deployments
- Targeting full enterprise/full data centre requirements
- No TB, memory, socket or core limit
- Supports BLU, pS and DPF deployment models

Core function

DB2 Workgroup Server Edition

- Entry level offering
- Single server for less intense workloads
- Limited by TB, memory, sockets and cores
- No support for BLU, pS or DPF deployment models

DB2 Enterprise Server Edition

- Entry level offering
- Single server for enterprise/more intense workloads
- No TB, memory, socket or core limit
- No support for BLU, pS or DPF deployment models

Limited capacity   *Full capacity*

DB2 Developer Edition	**DB2 CEO**
DB2 Express and DB2 Express-C	**DB2 Advanced CEO**

DB2® 10.5

with BLU Acceleration

Multi-workload database software for the era of big data

- ***Always Available Transactions***
  Disaster recovery of pureScale clusters over distances of 1000s km[1]; means minimal downtime

- ***Faster Analytics***
  In-memory hybrid technology yields performance improvements ranging from 8-25x performance improvements[2], without costs or limits of in-memory only

- ***Unprecedented Compatibility***
  An average of 98% Oracle database application compatibility[3]

- ***Future-Proofed Infrastructure***
  NoSQL support allows clients to expand and modernize their apps

DOWJONES

*"Before we made a final decision we benchmarked some of the key database management systems. That includes Oracle, SQL Server and DB2. **We ended up choosing DB2 for several reasons. One was reliability, second was performance and perhaps the most important factor was ease of use***"*
– Bashir Khan, Director of Data Management and Business Intelligence

Thank You!