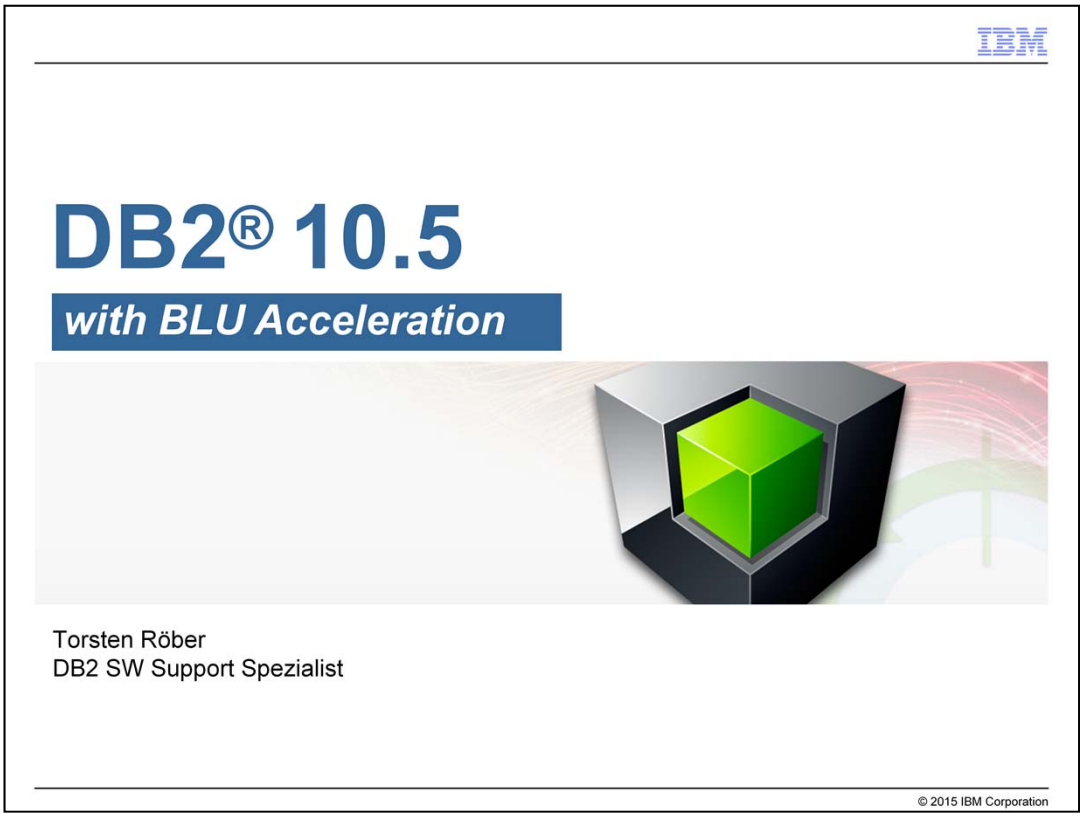

This presentation provides an overview of DB2 10.5 with BLU Acceleration including the new DB2 Cancun Release Shadow Tables and also details on DB2 pureScale updates in DB2 10.5 and DB2 Cancun Release 10.5.0.4.

IBM® DB2 Version 10.5 for Linux, UNIX, and Windows offers accelerated analytic processing by introducing a new processing paradigm and data format within the DB2 database product. Advantages include

- Significant reductions in time-to-value
- Increased consumability
- Minimal DBA design requirements
- Reduced query tuning and debugging efforts

DB2 Cancun Release 10.5.0.4 provides additional significant enhances in the DB2 pureScale space as well as adding the Shadow Tables feature allowing both OLTP and OLAP queries to run with superior performance characteristics against a single database!

Industry-leading compression, large performance gains for analytic queries, and large reductions in performance variation round out the benefits of deploying this technology.

Contributors to this presentation (last name order)

-George Baklarz

-Chris Eaton

-Bob Harbus

-Kelly Schlamb

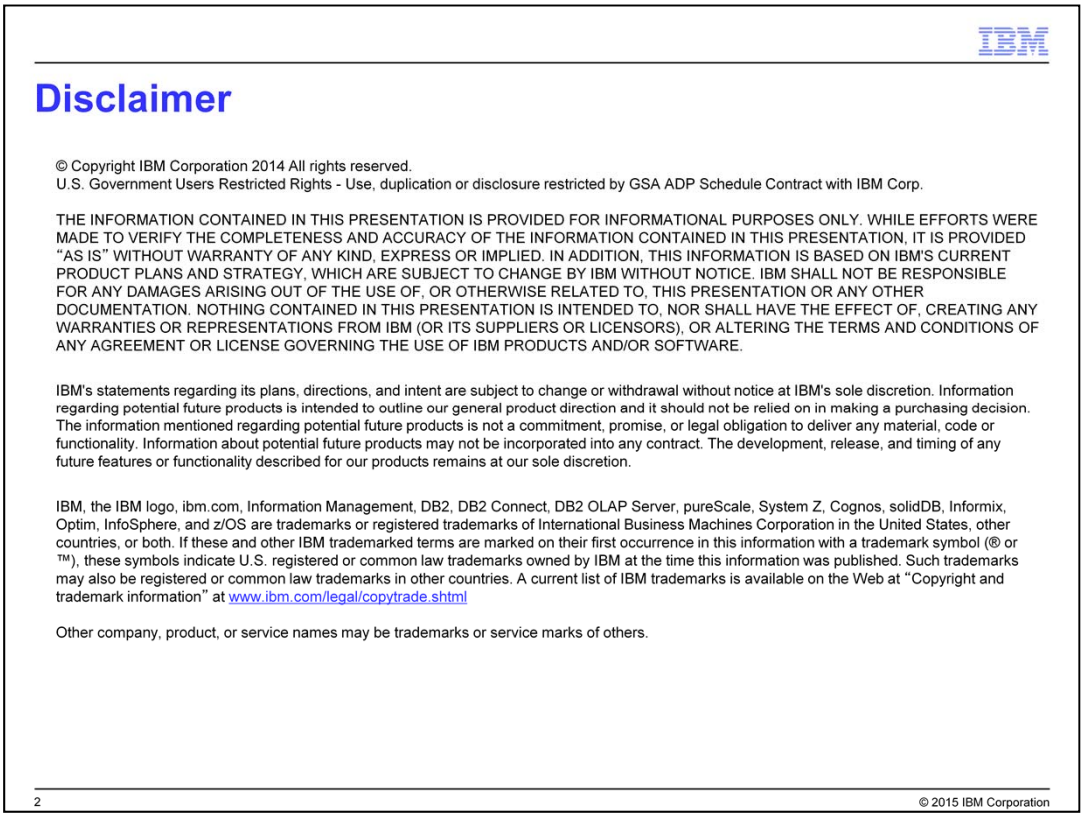

Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

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.

IBM, the IBM logo, ibm.com, Information Management, DB2, DB2 Connect, DB2 OLAP Server, pureScale, System Z, Cognos, solidDB, Informix, Optim, InfoSphere, and z/OS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Other company, product, or service names may be trademarks or service marks of others.

This chart must be presented and is a legal requirement.

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.

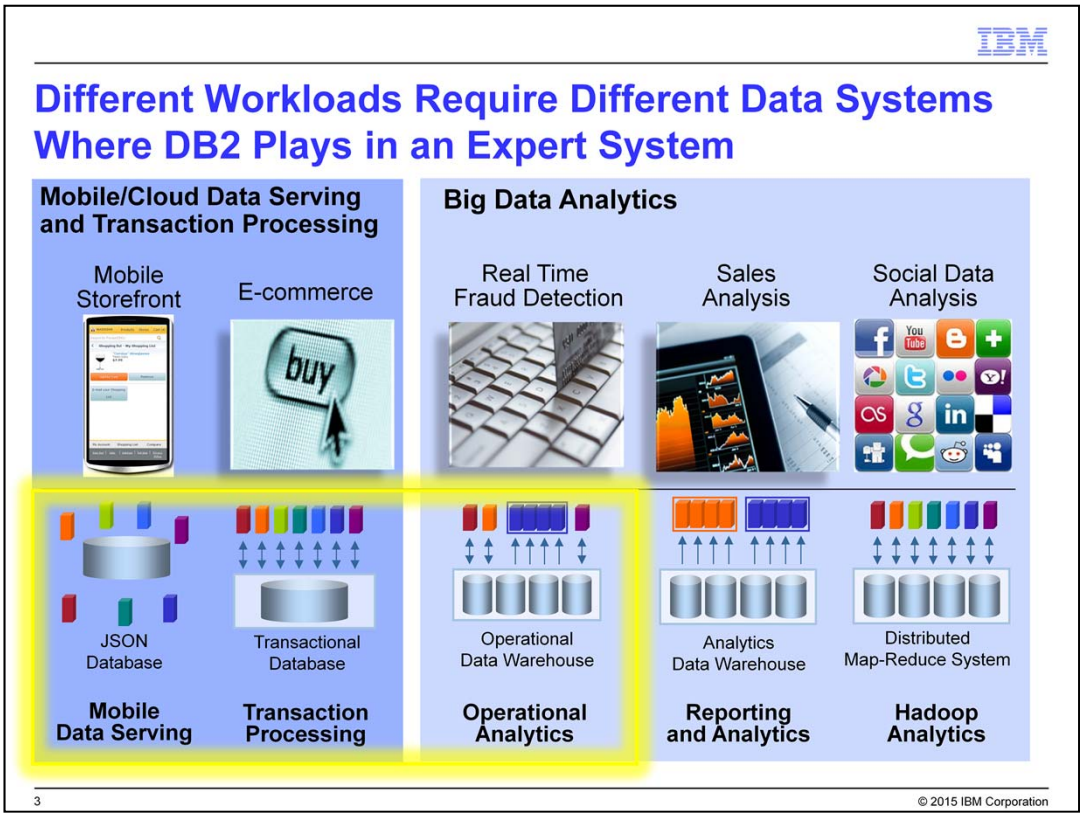

The idea of designing and building things for a specific purpose is not new. IBM believes that designing and building systems optimized for specific workloads is the best approach. We understand the breadth and complexity of different workloads facing enterprises and we know that a single system approach can't always satisfy diverse application requirements. However, we continue to add features to the DB2 product line to allow it to participate as part of many of these workloads.

This first charts shows the role that DB2 plays in Expert Systems. DB2 10.5 is the core database for both the PureData for Transactions and the PureData for Operational Analytic appliances.

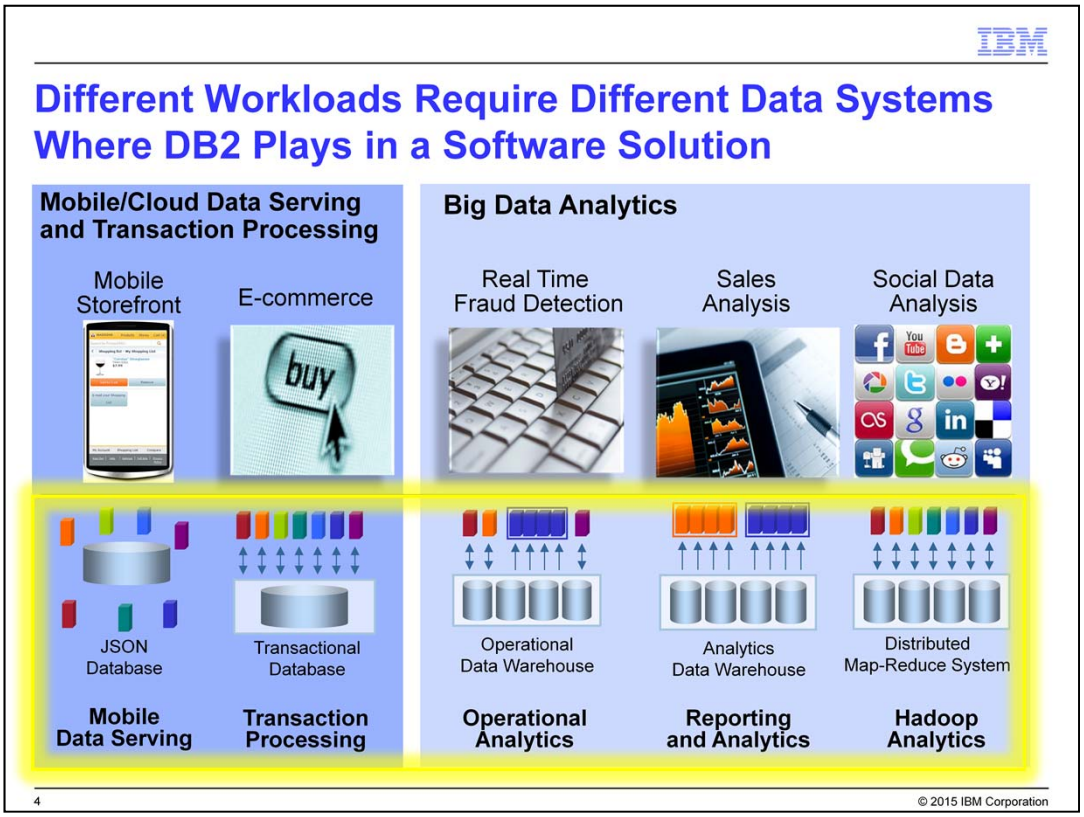

In addition to acting as the primary database for a variety of Appliances, DB2 can also be part of a software solution stack. DB2 10.5 introduces a number of features to help customers in all of these areas:

- Transaction Processing with new Disaster Availability and Recovery features included in the DB2 pureScale feature
- Operational Analytics with a new Column-Organized database engine within DB2
- Ability to have both OLTP and OLAP queries against a single database without concerns about performance – with Shadow Tables in DB2 Cancun Release 10.5.0.4

Customers will find that DB2 can be used for a larger variety of workloads and environments than ever before.

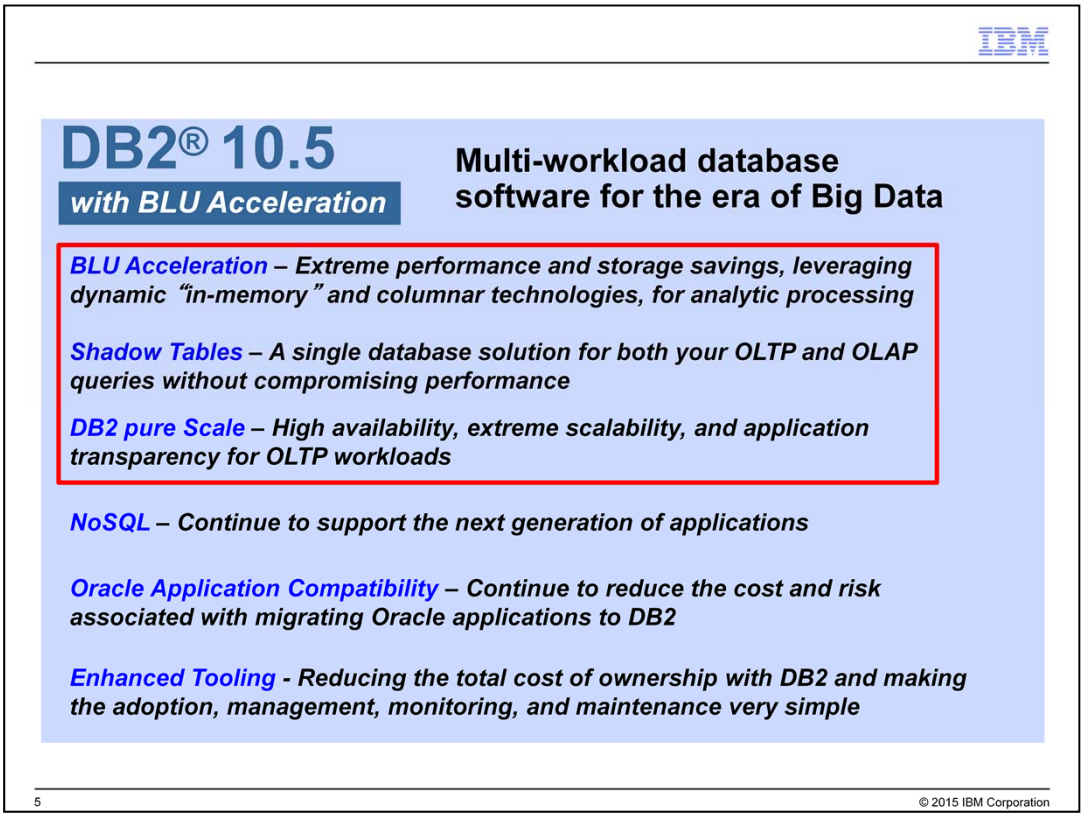

DB2 10.5 introduces a number of significant enhancements that will appeal to a large variety of customers. This presentation is going to cover the following topics

BLU Acceleration

The DB2 Version 10.5 release introduces compressed column-organized tables for DB2 databases, and includes broad support for data mart (analytic) workloads having complex queries that are commonly characterized by multi-table joins, grouping and aggregation, and table scans over a star schema.

DB2 Cancun Release 10.5.0.4 Shadow Tables

DB2 Cancun Release 10.5.0.4 provides the new Shadow Tables feature functionality. This allows you to have a single database for both your OLTP and analytic OLAP queries, with no performance penalty. Let DB2 do all the heavy lifting for you!

DB2 pureScale

The DB2 pureScale Feature provides extreme capacity and application transparency which exceed even the strictest industry standard. Continued improvements in continuous availability and high availability, the DB2 pureScale Feature tolerates both planned maintenance and component failure with ease.

NoSQL Support

DB2 supports the Resource Description Framework (RDF) that you can use as a standard data interchange framework for modeling information. Applications can store and query RDF data in DB2® Enterprise Server Edition for Linux, UNIX, and Windows (DB2 Enterprise Server Edition) databases. As part of the launch of DB2 10.5, a technical preview of JSON support will be made available. This technology preview is not part of the code being shipped at General Availability.

Oracle Application Compatibility

If you work with relational database products other than DB2 products,V10.5 builds on existing functionality, interfaces, and compatibility features to provide additional enhancements that make DB2 products more familiar to you. These enhancements reduce the time and complexity of enabling applications that are written for other relational database products to run quickly in a DB2 environment.

Enhanced Tooling

Many of the existing tools used within DB2 will be updated to work with the DB2 pureScale and DB2 BLU

Acceleration technology.

DB2 10.5 with BLU Acceleration

We now take you through the details on the internals of DB2 with BLU Acceleration.

What is DB2 with BLU Acceleration?

- **Ground breaking 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

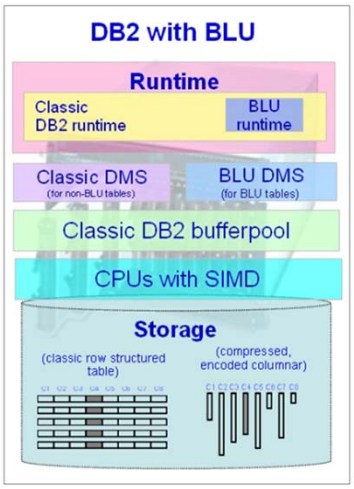

7

This slide describes at a high level what DB2 with BLU Acceleration is. What is the key business value of implementing BLU Acceleration?

This is a new technology that has been developed by IBM and integrated directly into the DB2 engine. BLU Acceleration is a new storage engine along with integrated runtime (directly into the core DB2 engine) to support the storage and analysis of column organized tables. The BLU Acceleration processing is parallel to the regular, row-based table processing found in the DB2 engine. This is not a bolt-on technology nor is it a separate analytic engine that sits outside of DB2. Much like when IBM added XML data as a first class object within the database along with all the storage and processing enhancements that came with XML, now IBM has added column organized tables directly into the storage and processing engine of DB2.

Simply put, this is a column-organized table store in DB2. Along with this store are many benefits including significantly improved performance, massive storage savings and ease of implementation and ease of management.

This feature allows us to deliver on these performance and storage innovations while also optimizing the use of main-memory, improving I/O

efficiency and exploiting CPU instructions and characteristics to enhance the value derived from your database investments.

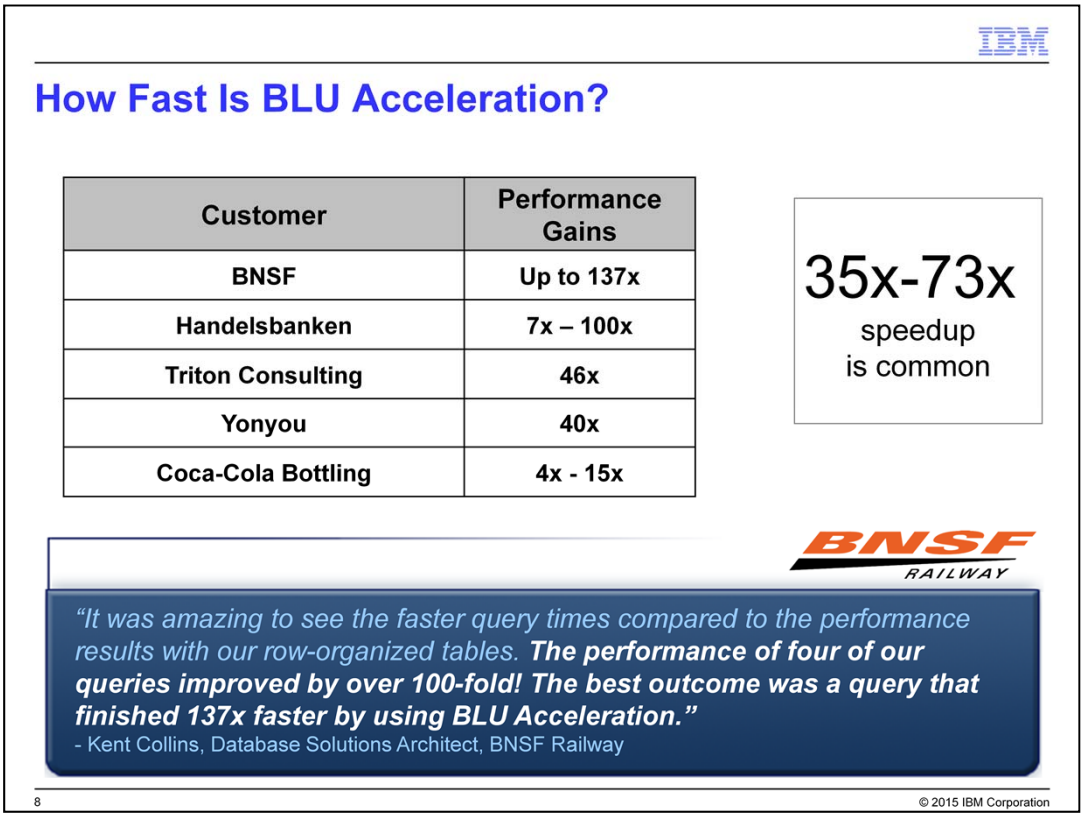

This slide shows the performance improvements some customers have experienced using DB2 with BLU Acceleration.

While individual customer results do vary, the average generally falls in the range of 35-73x improvement for typical customer analytic workloads!

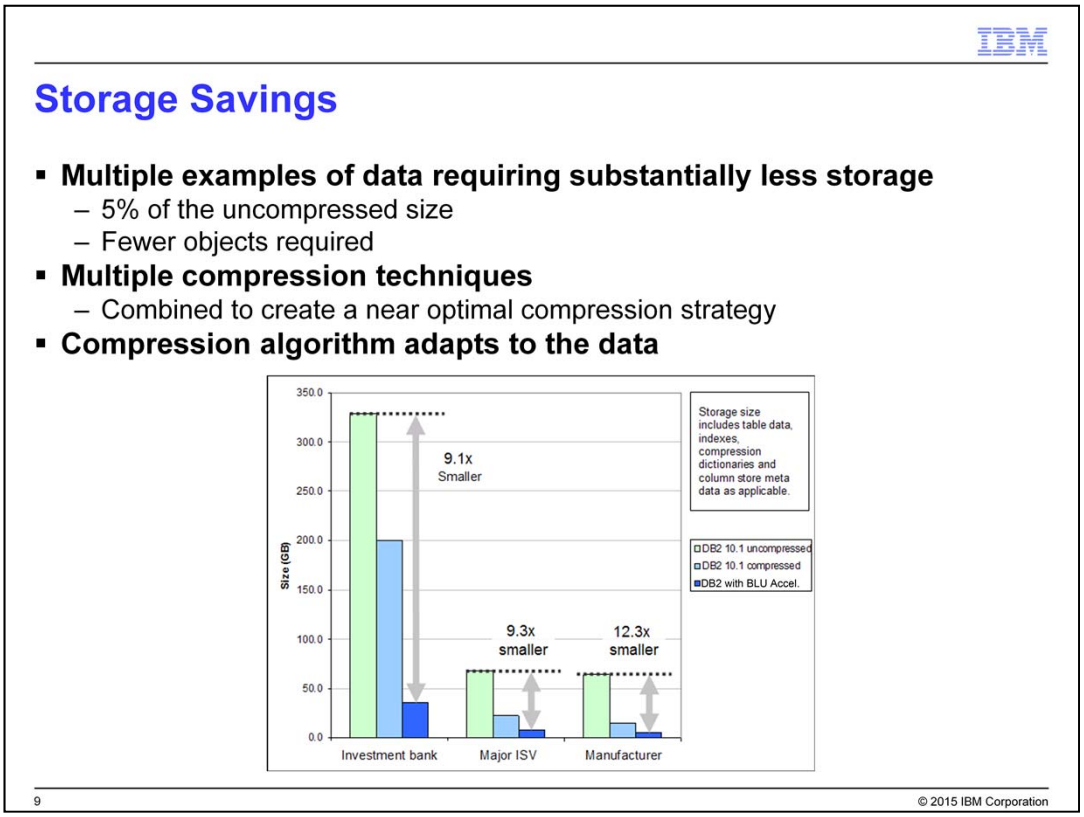

A key benefit of the DB2 with BLU Acceleration technology is the storage savings. By virtue of the columnar storage method DB2 for LUW is often able to provide significant storage savings even over DB2 10.1 adaptive compression (which currently leads our competitors by a wide margin in storage savings). This provides savings on disk for the base tables and also provides savings in backup image size and time for backup.

Compression, however, is not the only storage benefit. DB2 with BLU Acceleration does not require auxiliary performance structures like indexes, materialized views/materialized query tables (MQTs), multi dimensional clustered tables (MDCs), etc, to gain superior performance. Customers will gain significant storage savings by not requiring these objects and also save time in not having to maintain them.

Lower cost of operational analytics

- Reduces the number of objects that need to be deployed to achieve the required level of performance saves time and effort, as well as being able to deploy and start using the DB2 with BLU Acceleration technology faster. In addition effort required for ongoing tuning is reduced if not eliminated
- Reduces the cost and time of performance tuning for analytical workloads. Since DB2 with BLU Acceleration has superior performance right out-the-box, additional indexes and tuning is not required
- Does not require scale out technology to achieve required performance. Enables better usage of existing hardware technology or leverage single

server deployments to deliver the performance of larger clusters from competitors

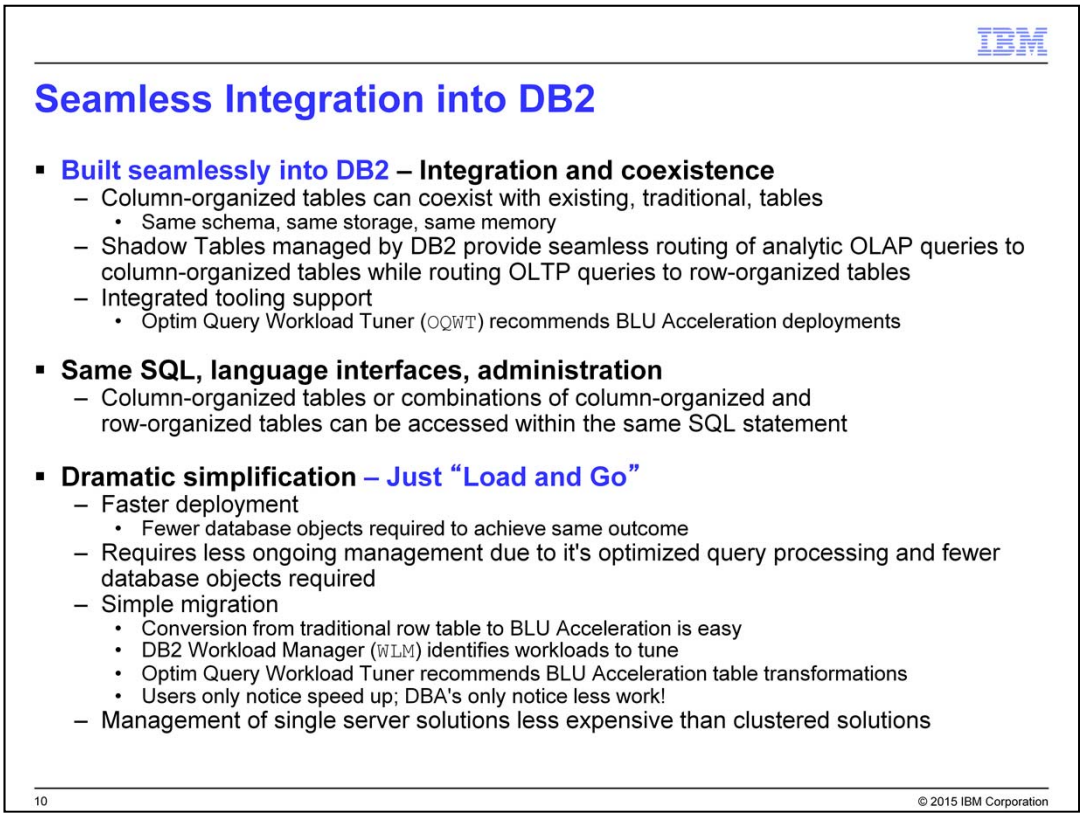

Built into DB2 with consistent SQL. etc.

- Row-based table storage and column-based table storage both coexist in the same database and SQL can access them both at the same time
- Consistent SQL, language interfaces, administration; Reuses DB2 process model, storage, utilities
- Integration Tooling Support
  - Integrated DBA support through Data Studio
  - Integrated workload monitoring and workload query tuning support with Optim Performance Manager (OPM) and OPTIM Query Workload Tuner
  - Integrated design with IBM InfoSphere Data Architect (IDA) and InfoSphere Warehouse (ISW) Design Studio
  - Integrated data movement support with ISW SQW (SQL Warehousing Tool)
  - Cognos Automated Summary Table creation using BLU Acceleration with Cognos ROLAP (Relational online analytical processing)

Simplification

- Reduces the cost and time of performance tuning for analytical workloads
- Every effort has been made to simplify the adoption process for column-organized tables. The DDL changes for creating these table are simple. Moreover, database-level overrides (such as the **dft_table_org** database configuration parameter) further simplify the adoption process
- The **db2convert** conversion utility enables customers to convert existing tables in row orientation to column orientation. Note that the table conversion process does temporarily require space for both the source and target table. Because there is no online process to convert a column-organized table back to a row-organized table, the best practice is to perform a backup before converting tables to column orientation
- Consumability, time to value, Simplification – Load and then run queries!

- No indexes
- No reorganization (it's automated)
- No runstats (it's automated)
- No MDC or MQT
- No hash partitioning, statistical views, or optimizer hints

Essentially many of the steps that are typically required (especially in competitor solutions) need more design, management and tuning. These are tasks you no longer need to worry about with DB2 with BLU Acceleration.

BLU Acceleration is super fast and super easy. You just create BLU tables, load your data and start running your queries. It is that easy.

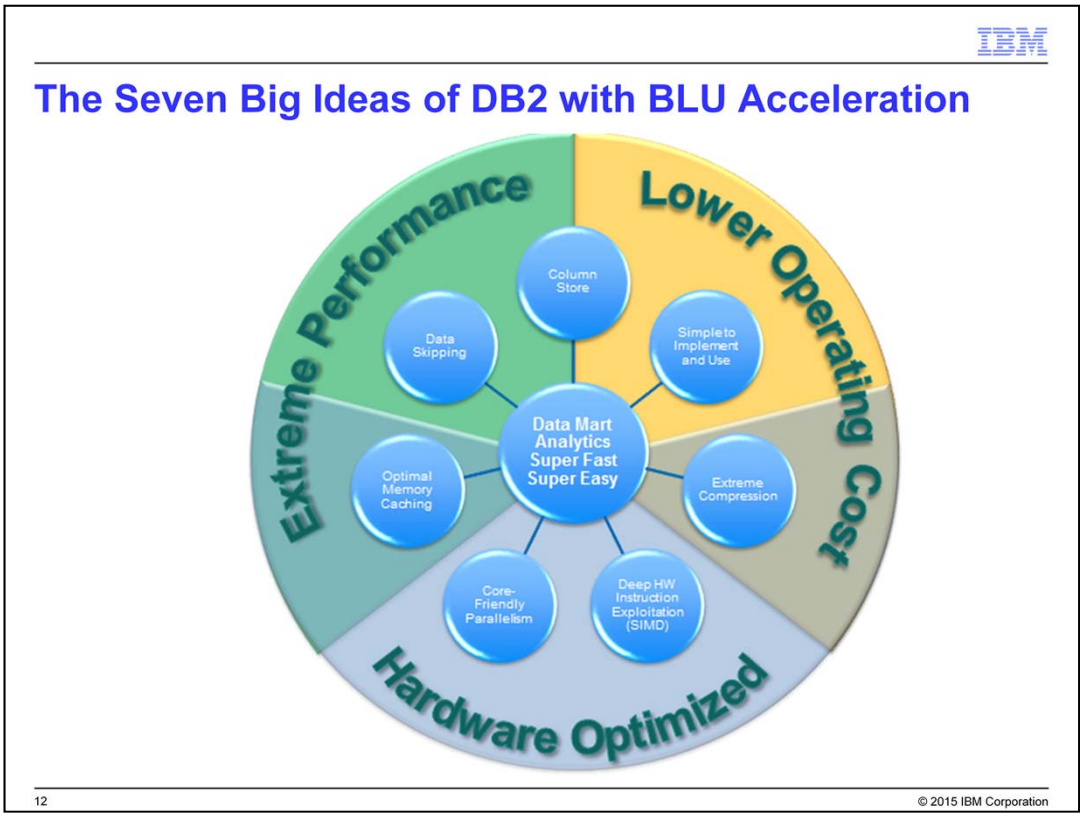

BLU Acceleration is based on 7 innovations that have been added to DB2 which we call "Big Ideas". Each of these big ideas is a technology capability that provides business value and en-mass make up what we refer to as **BLU Acceleration**.

Lower Operation Costs

•Column Store – by storing table data in column organized format we not only save significantly on storage costs but we also improve I/O and memory efficiency. This lowers operating costs

•Simple to Implement and Use – as has already been mentioned, with BLU Acceleration you just create the column organized table and then load and go. That's it – no tuning, no indexes, etc. This lowers administration and development costs significantly

•Extreme Compression – by using compression and sophisticated encoding algorithms, DB2 can save significantly on storage costs including power, cooling, and management of that storage

Hardware Optimized

•Extreme Compression – in addition to the lower costs, the compression algorithms used exploit processor characteristics to improve performance. The compression we use works with a register friendly encoding technique to improve processor efficiency

•Deep Hardware Instruction Exploitation – we will discuss this in more detail on a future slide, but with SIMD (see below) processing we are multiplying the performance of the processor by having instructions work on multiple data elements simultaneously

•Core Friendly Parallelism – Access plans on column organized tables will leverage all of the cores on the server simultaneously to delver better analytic query performance

•Optimal Memory Caching – With row organized tables, a full table scan ends up putting data into the bufferpool that is often not required. For column organized tables, if there are columns that are involved in joins or other predicates in many queries then we can pack the bufferpool full of those columns while keeping other columns out of memory if they are not regularly used. This improves performance and optimized the memory available

Extreme Performance

•Optimal Memory Caching – as stated above, this not only helps to optimize hardware but also improves overall workload performance

•Data Skipping – by keeping track of which pages of data contain which column values, we can avoid a lot of I/O and query processing by simply skipping data we already know would not qualify for the query

•Column Store – in addition to lowering costs, by selecting only columns that are part of a query we can increase performance of queries by an order of magnitude in some cases

SIMD: (**S**ingle **I**nstruction stream **M**ultiple **D**ata stream) A computer that performs one operation on multiple sets of data. It is typically used to add or multiply eight or more sets of numbers at the same time for multimedia encoding and rendering as well as scientific applications. Hardware registers are loaded with numbers, and the mathematical operation is performed on all registers simultaneously.

I/O: Input/**O**utput

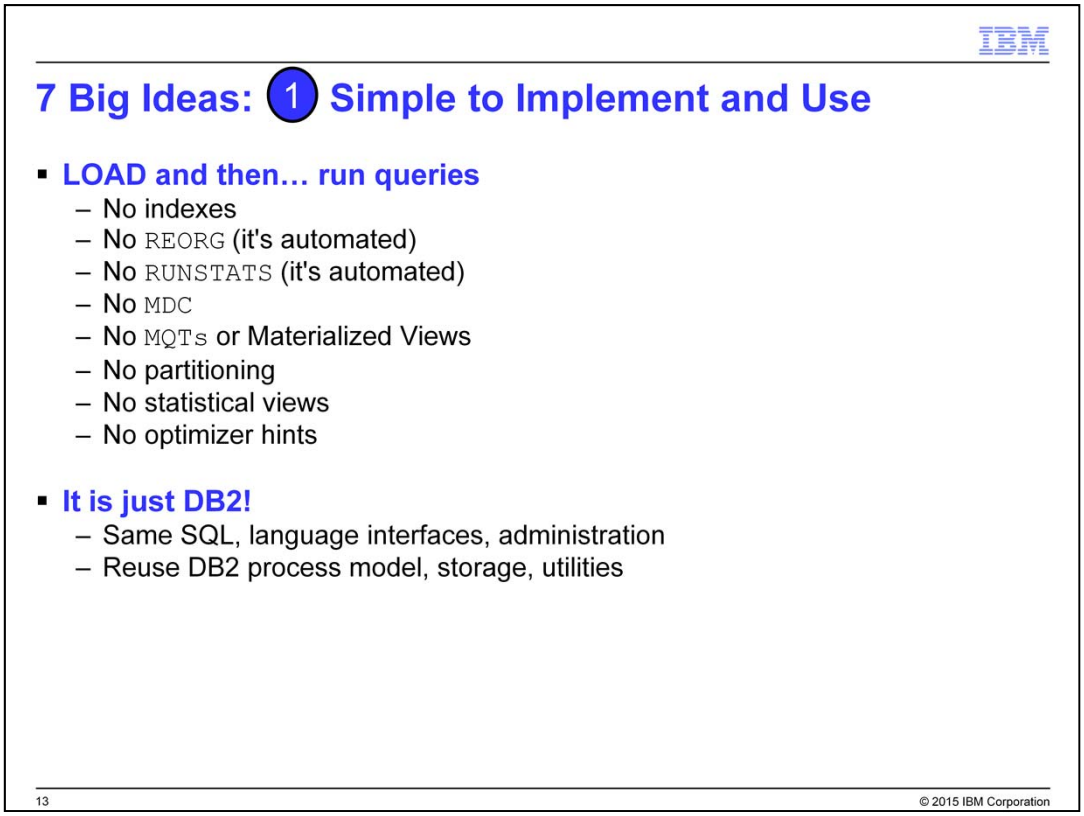

The simplicity of DB2 with BLU Acceleration is one of the key value propositions. The fact that it is really simple to use is key – remember it is all part of the DB2 kernel.

Because of the simplicity and the built in space reclamation and statistics gathering, and also the fact that other structures such as indexes, MDC tables, MQTs, etc are not needed really to add to the value.

From a customer perspective it is LOAD and GO! You can start running your queries immediately after data load and start getting the performance gains of DB2 with BLU Acceleration immediately.

REORG: DB2 Table/Index Reorganization

RUNSTATS: DB2 Table/Index Statistics gathering

SQL: Structured Query Language

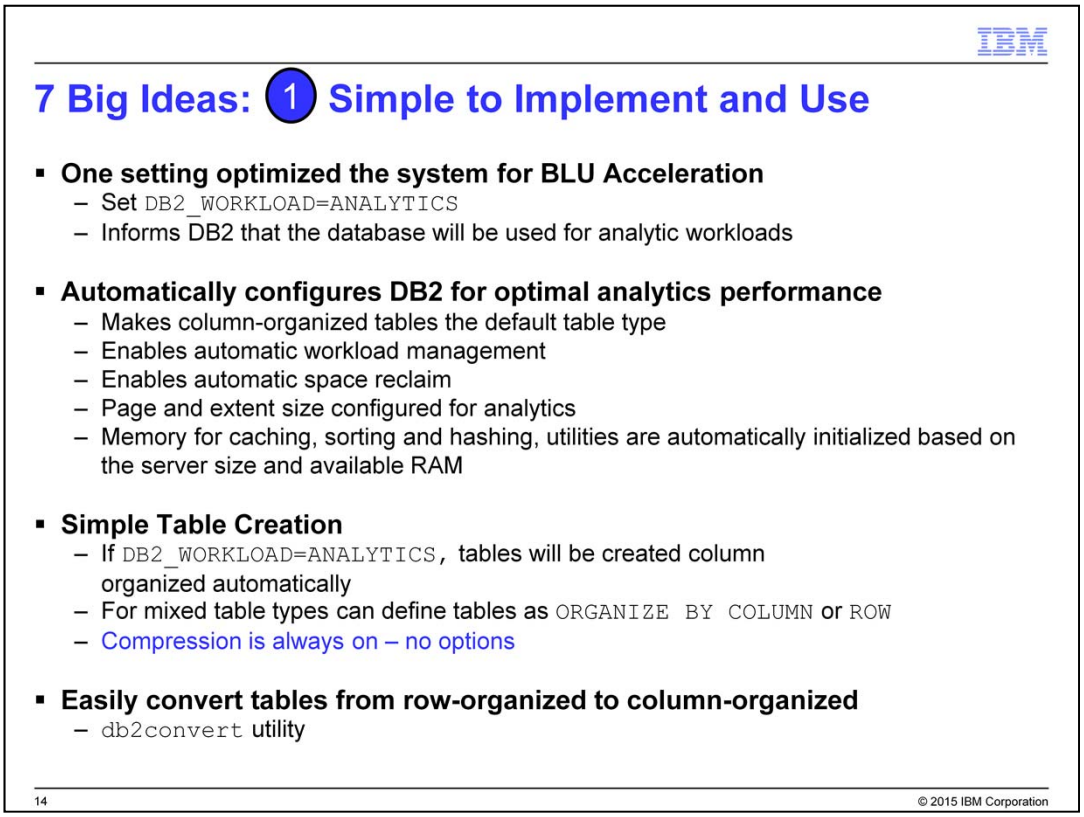

It's really valuable to set DB2_WORKLOAD=ANALYTICS before creating a new DB2 database with BLU Acceleration. When you create the database, it will detect this setting and configure itself for analytic processing, with the changes that follow. One switch, great value. And it's personalized to your server – based on your CPU type, memory, and disks.

If for some reason you can't or don't want to set this, don't worry. Everything that it automates can be enabled with a few commands in just a few minutes.

DB2_WORKLOAD is a DB2 for LUW Registry variable.

RAM: Random Access Memory

CPU: Central Processing Unit

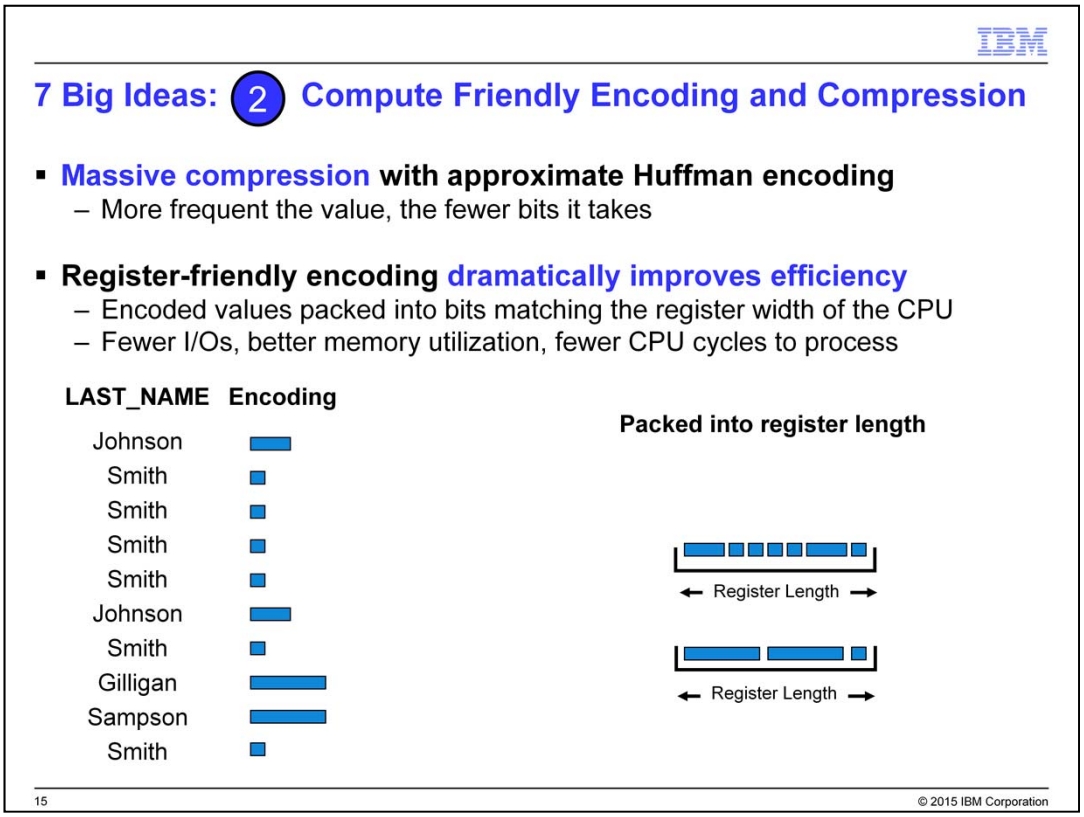

Key to DB2 with BLU Acceleration is how we encode the data on disk – this is important not only because we get significant compression from it, but also because we are using a different encoding scheme which allows us to scan data while it is encoded. This is all part of the DB2 with BLU Acceleration scanning engine. This ability to scan (and do predicate evaluation for instance) while encoded, allows us to delay materialization as long a possible. This means we can do late materialization and this allows us to do as much work without uncompressing the data. This saves processing effort and I/O throughout. The bottom line is we do as much as we can without decompression.

With respect to the Huffman encoding… you can think of it this way .. Something that appears many many times should be able to be compressed more than other things which do not appear as often. For instance the letter 'e' may appear many times so you can encode it with a single bit (1) thereby getting very good compression for this item which shows up many times. However for say the letter 'q' which may not appear as often you, maybe only once, you encode this with seven bits (7). The end result is those items appearing more often get higher level of compression. The letter e example was only an example of the concept. FYI: In BLU we do not encode at a per letter basis. The encoding is always an entire column value. And we combine this with prefix encoding and offset coding, to give us a more complete compression story.

Once encoded and comparable and analyzable – we pack it as tightly as possible in a collection of bits to the register width of the CPU of the server. Note that the above representation is only an example to illustrate the point and is not intended to represent exactly how the code for column organized tables works.

In the end we will see less I/O, better utilization of the memory, and more effective use of the CPU. Storage will also be significantly less. Of course we cannot forget the much better performance.

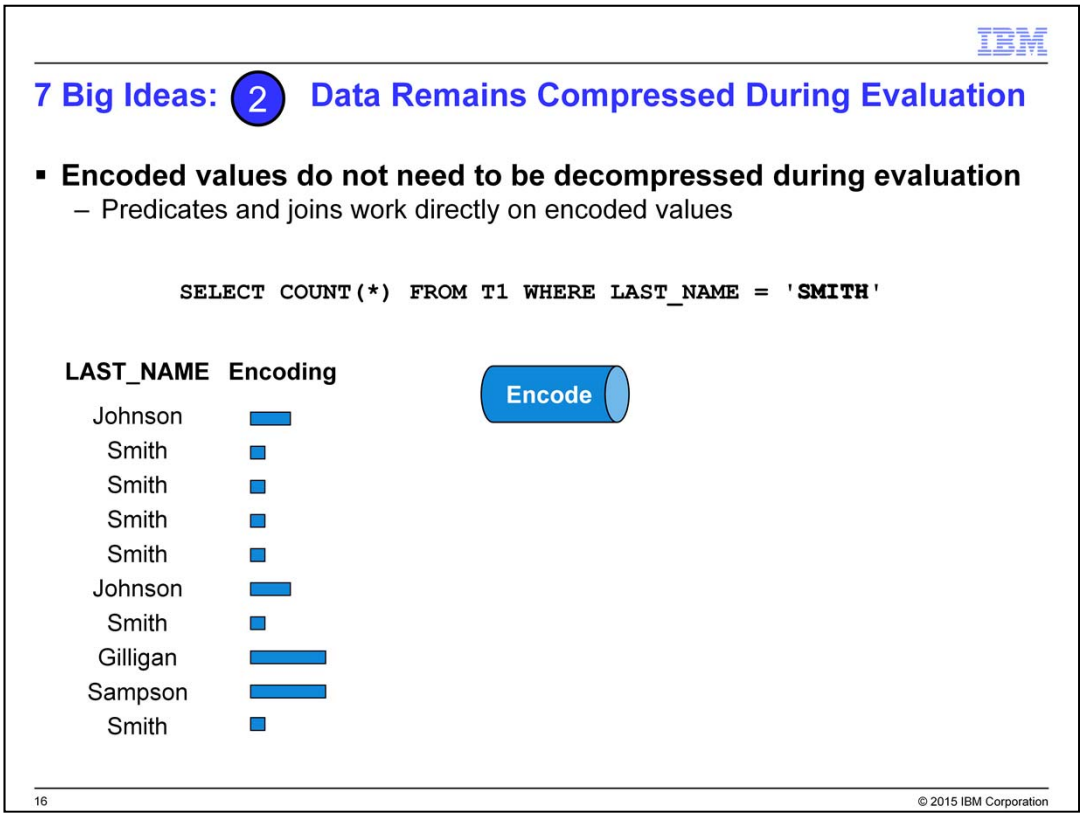

The second key to compression benefits is that we do not decode/uncompress the data until we have to in order to return a result to the user. For example in the case of the above query, rather than uncompressing the values of the LAST_NAME column and then comparing them to the predicate “= SMITH” with BLU Acceleration we actually encode/compress the value “SMITH” and then compare it directly to the encoded values in the column that are stored in the data pages.

Likewise if we are joining to tables on a given column, we can compare the encoded values rather than having to uncompress the data to do the comparison.

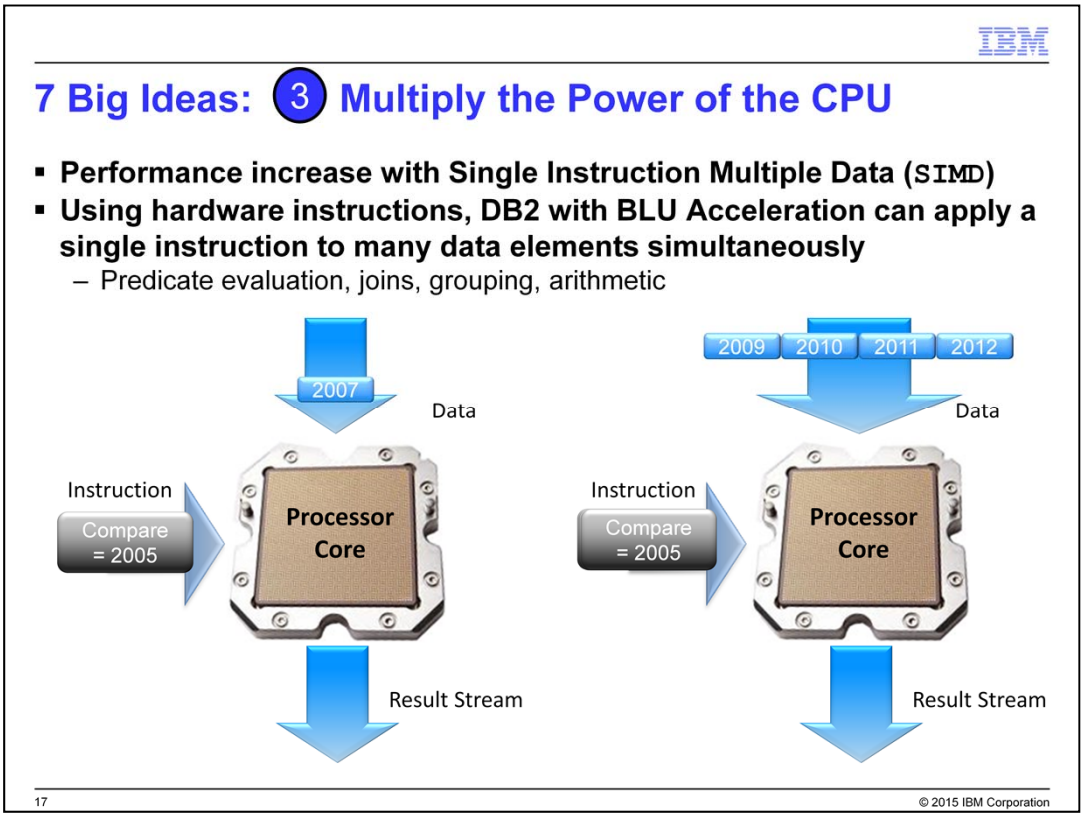

The 3rd big idea for BLU is the use of special hardware instructions to work on multiple data elements with a single instruction. This is known as Single Instruction Multiple Data or SIMD. There are special instructions available on various hardware platform to accomplish this. We have special BLU Acceleration code to do this. And we put as much data (128 bits) into a SIMD register. So now with 1 instruction I can multiply the power of the CPU – I can get results of all those data values in the register. This means we can use the power of the CPUs now to do not only scanning, but joins, etc.

In the above example we show 4 column values being acted on at one time but this is only for illustration purposes. It is quite possible to have >4 data elements being processed by a single instruction with this technology.

SIMD: (**S**ingle **I**nstruction stream **M**ultiple **D**ata stream) A computer that performs one operation on multiple sets of data. It is typically used to add or multiply eight or more sets of numbers at the same time for multimedia encoding and rendering as well as scientific applications. Hardware registers are loaded with numbers, and the mathematical operation is performed on all registers simultaneously.

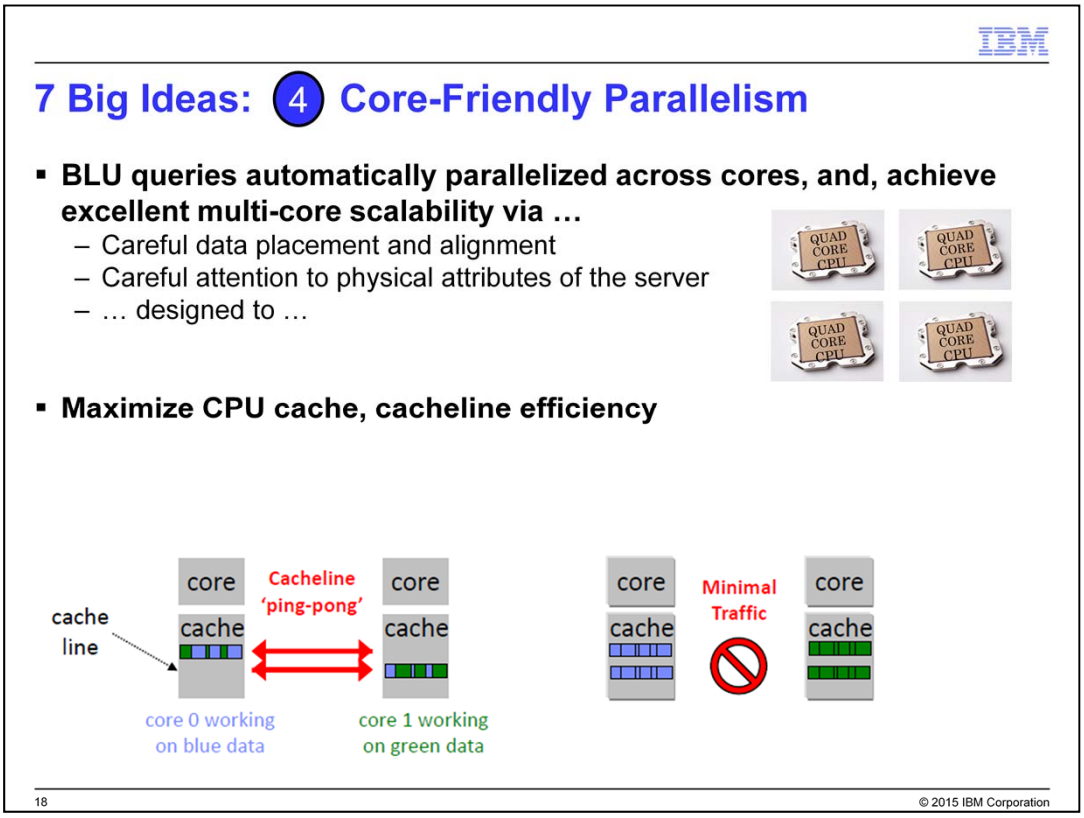

With DB2 with BLU Acceleration we pay very close attention to multi-core parallelism. DB2 with BLU Acceleration is designed from the ground up to take advantage of the cores you have and to always drive multi-core parallelism for the queries you have. This is all done in shared memory – this is not DPF parallelism.

In DB2 with BLU Acceleration we really look into memory latency and memory access. At this point in time we are almost up-to 16-way parallelism and we are pushing for 32-core and 64-core for future.

DPF: Data Partitioning Feature of DB2. Provides a shared-nothing parallel database environment for DB2/InfoSphere Warehouse

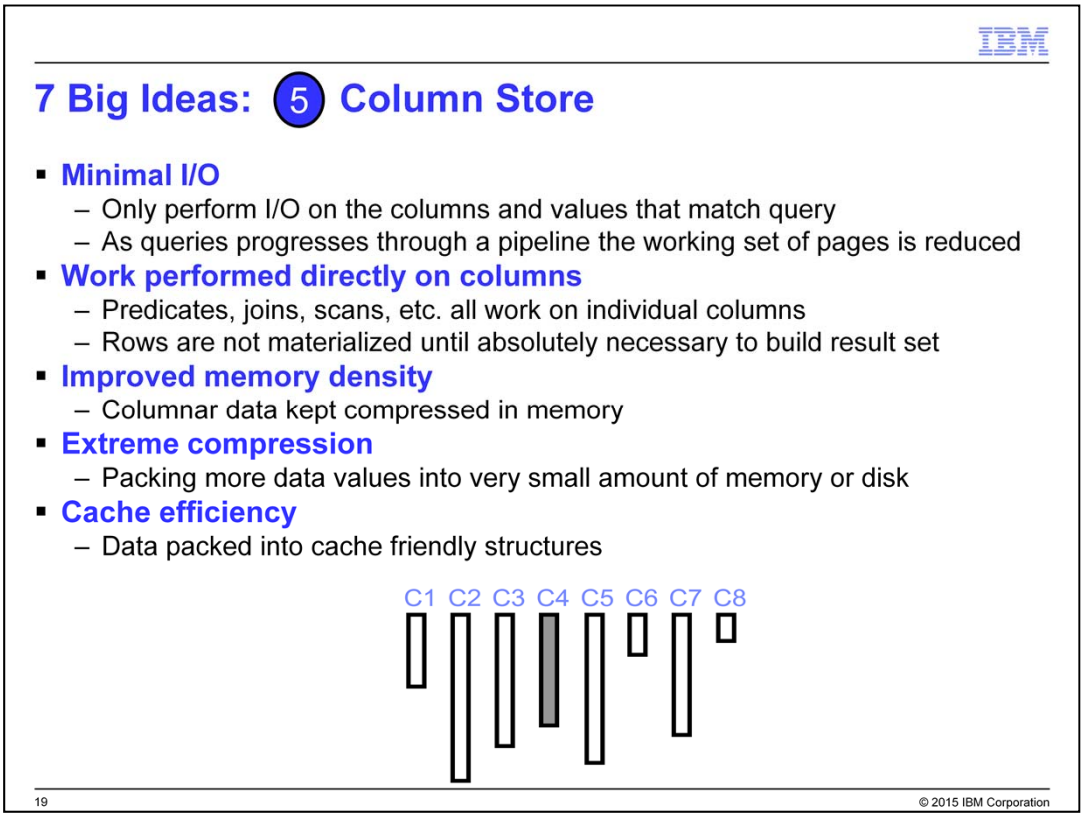

Using a column store and the encoding we spoke of earlier we gain big advantages in DB2 – a store and results that our major competitors do not have and cannot achieve! DB2 can!

By using a column store and the encoding we are able to get an additional level of compression for DB2 and also very importantly we are able to minimize the I/O that we need to do to satisfy scans and queries. Recall also that we are able to do evaluation against the compressed data (no decompression needed – we use late materialization) so this also reduces the effort we need to do.

Our I/O is much better, only reading what you need. Often you can see up to 95% of the I/O go away as most analytic workloads only access a small amount of data.

DB2 works directly on the column data and as the query progresses, fewer and fewer rows will qualify and so the working set of data pages accessed in

the next column worked on is reduced. In addition, the rows are not “stitched” back together until absolutely necessary. In most cases this will be when the answer set is returned to the end user or in some cases when we join a column organized table with a row organized table.

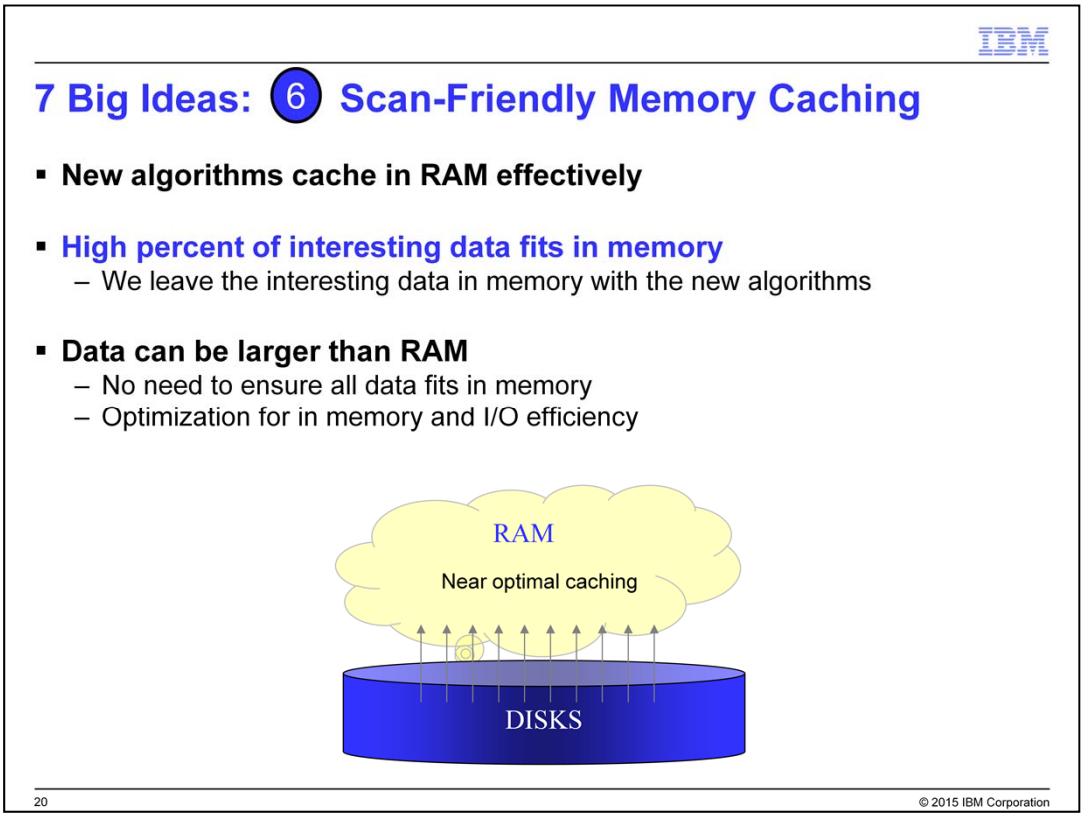

Generally when doing table access, the full table will not fit in memory; meaning more I/O.

Recall DB2 with BLU Acceleration uses excellent (10x) compression and columnar access.

We discussed with customers their ratio of memory to disk and most indicated they have a 15-50% ratio of memory to disk. What does this mean for customers using DB2 with BLU Acceleration? There are very good opportunities that the customers will be able to fit a high percentage of their active data in memory (possibly from 80-100%). So we have a better way to use that memory in DB2 with BLU Acceleration!

We have a new al algorithm changing victimization in BP – now getting almost 100% benefit from BP – almost able to fit 80% of active data now in BP. This is not as effective with row store due to row store requiring the full row and not just active columns. Algorithm runs side by side with row algorithm.

BP: DB2 Bufferpool – used for storing DB2 data, index, and temporary data pages in memory during DB2 runtime execution.

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

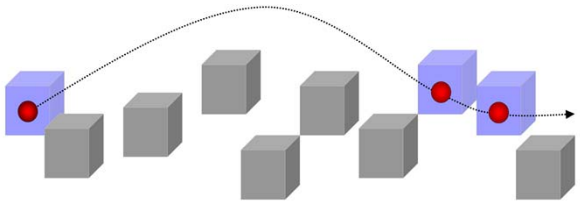

21

- Similar to Netezza zone maps
- This is IO and memory savings – even if pages are in memory we can skip over them.
- This is a range of records– every 2000 records we have the min max information – automatically maintained during load, Insert/Update/Delete (IUD)
    - Stored in a separate table called a synopsis table

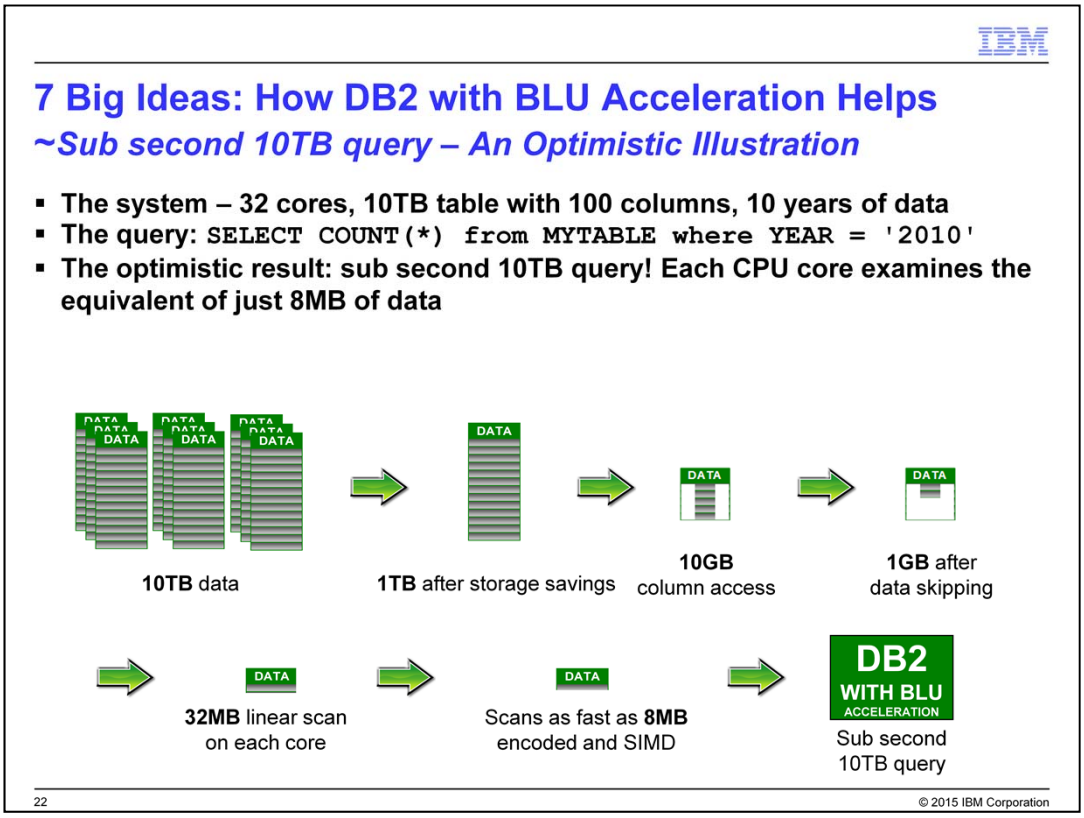

Before BLU Acceleration – no way a sub second query without an index would be possible… how can we possibly achieve this with DB2 with BLU Acceleration?

- Compress 10x
- Query only accesses 1 column -> so 1/100 column – 1% - 10GB of 1TB
- So using data skipping we can skip over 9 years and only look at 1 year – now 1GB of data
- Now divide across 32 cores for scan – almost linearly – each core 32 MB of data (Note this is a very optimistic end but something we are working towards)
- Scan will happen faster on encoded – say 4x faster than traditional – as fast as 8MB of data on traditional system. No issue doing this in a sub second! YAY!

GB: Gigabyte - A unit of information equal to one billion (10**9) or, strictly, 2**30 bytes.

TB: Terabyte - A unit of information equal to one million million (10**12) or

strictly, 2**40 bytes.

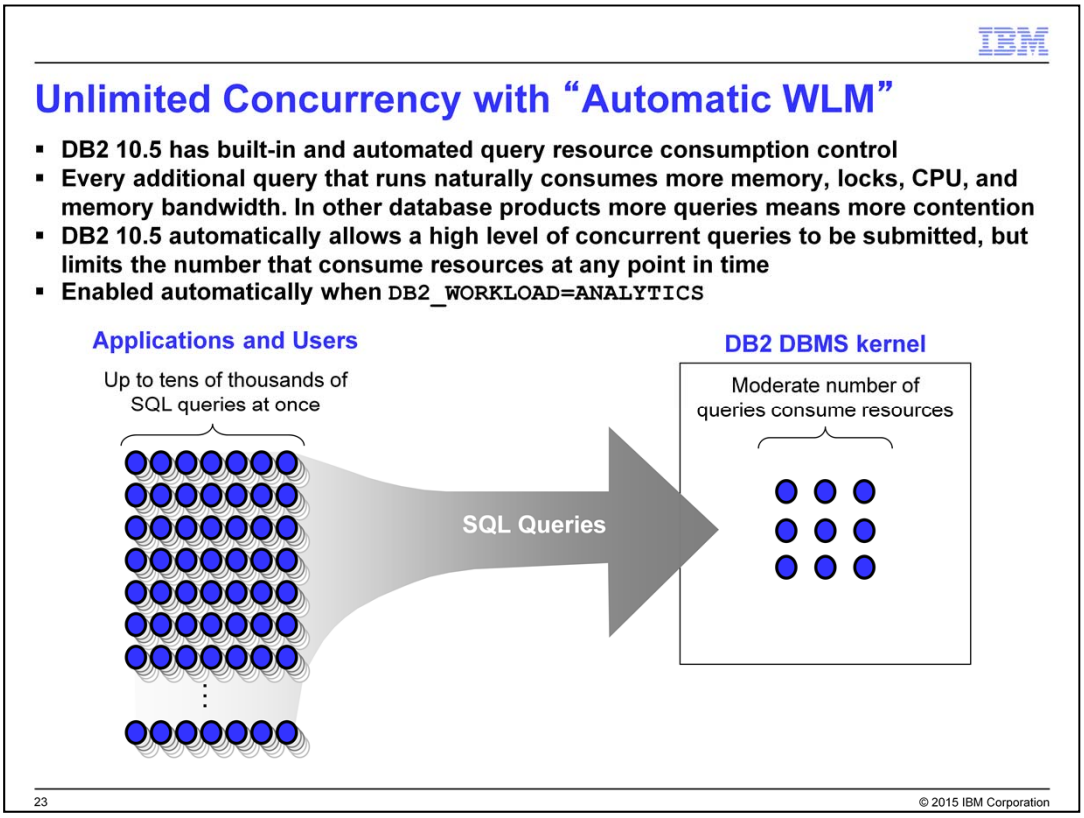

DB2 10.5 has built-in and automated query resource consumption control

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

That means more memory and CPU for each query that's running. In the end, the entire workload benefits!

This is a bit like people leaving a crowded theatre. If 1000 people try to exit at the same in a mad rush it goes slowly. But if people line up courteously and exit 4 by 4, everyone gets through more efficiently.

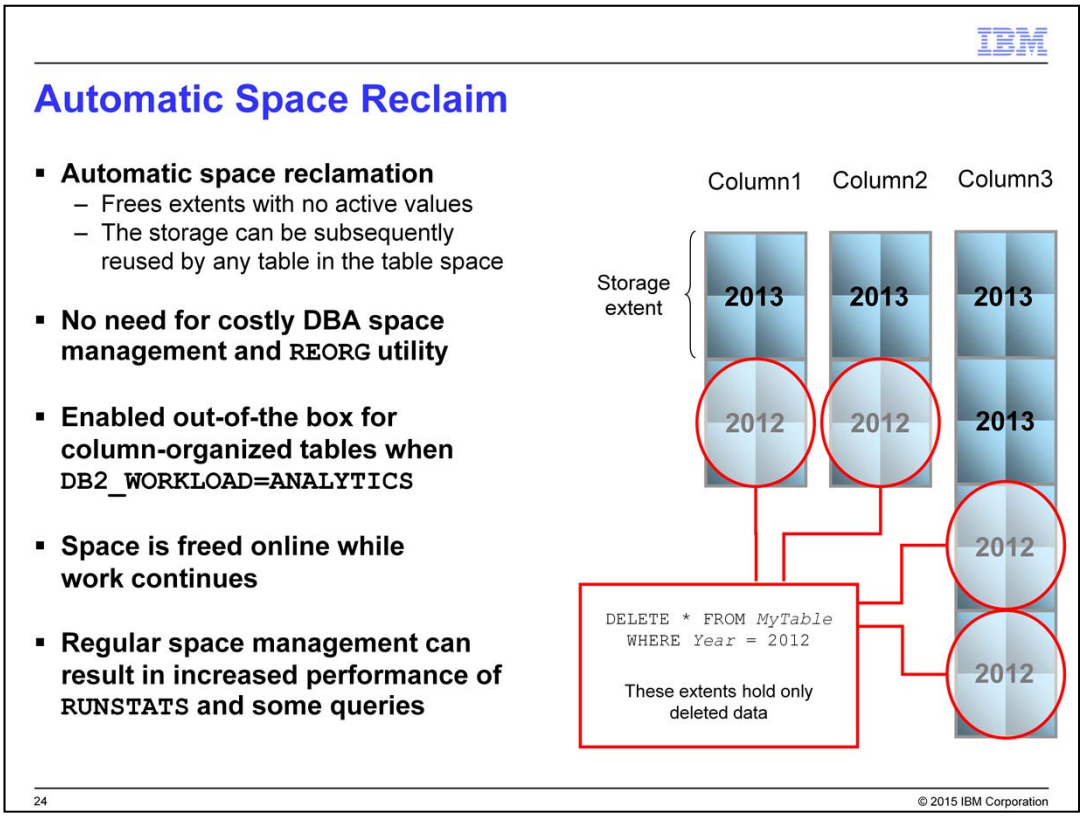

Like other features with BLU Acceleration we wanted to make storage management a breeze. We have enabled automatic space reclaim when DB2_WORKLOAD=ANALYTICS. Storage extents that hold only entirely deleted values are detected and freed by a background process. It's a silent and automatic garbage collection.

In this example we see a table with 8 storage extents holding 3 columns of data.

A SQL statements deletes (rolls out) all of the data associated with a time period, in this case 2012. The automatic storage reclaim logic detects that 4 extents are now holding deleted values, and returns them back to database storage where they can be reused in the future by this or other tables requiring storage.

One more burden the DBA need not worry about!

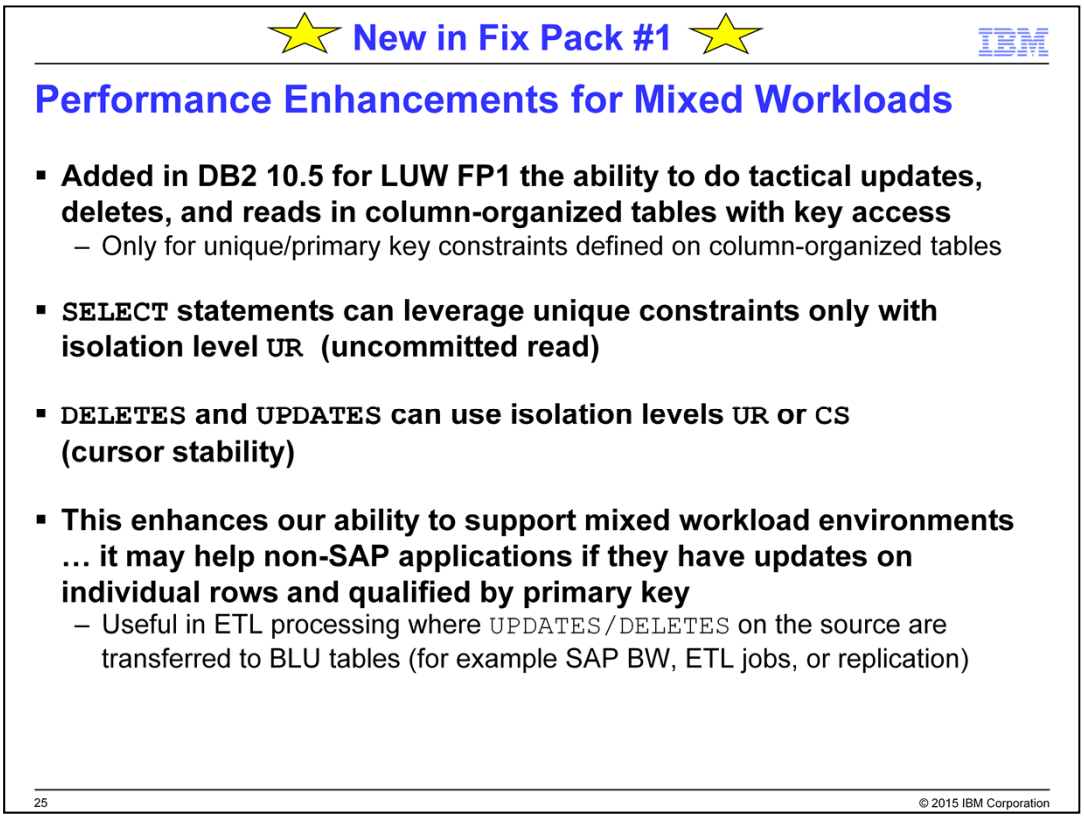

A performance enhancement in Fix Pack #1 for DB2 (with BLU Acceleration) provides the ability to do tactical updates, deletes, and reads in column-organized tables with key access.

This enhances the ability of BLU Acceleration to support mixed workload environments.

Some additional quotes from our customers on DB2 10.5.

DB2 Cancun Release 10.5.0.4
BLU Acceleration Enhancements

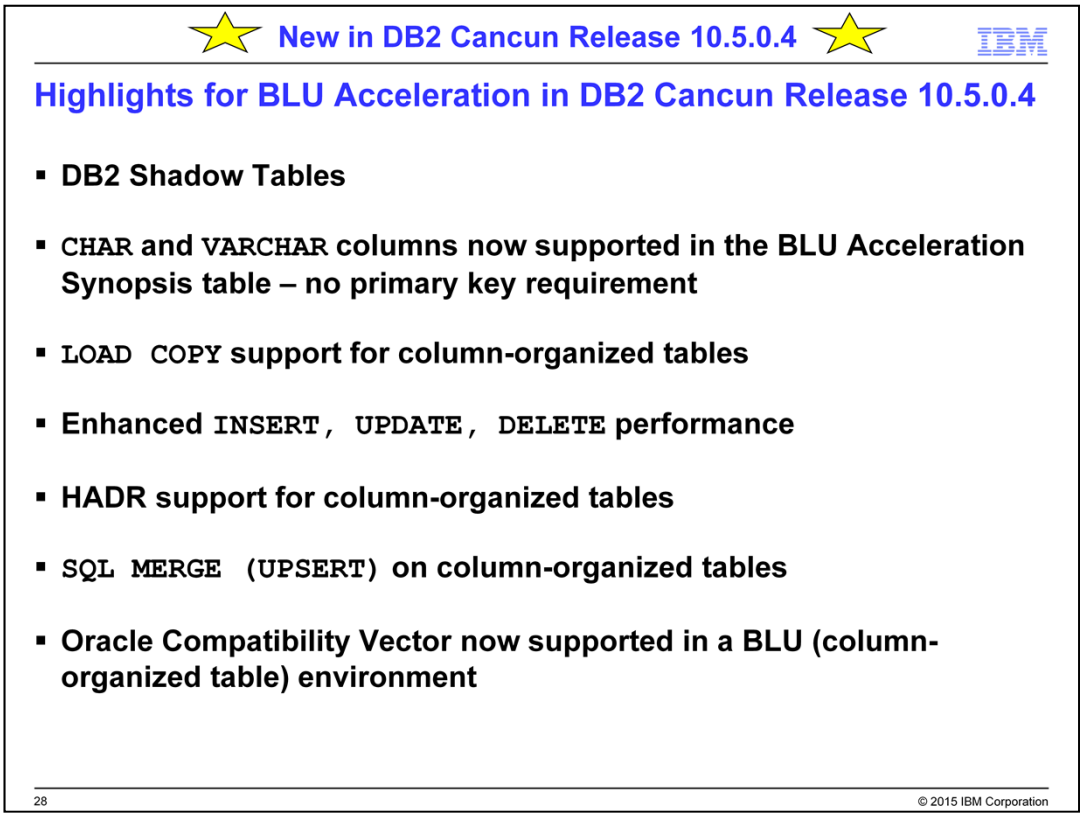

This slide highlights some of the key items that were added in DB2 Cancun Release 10.5.0.4 (FP4). Some of these items will be discussed in more detail in the following slides.

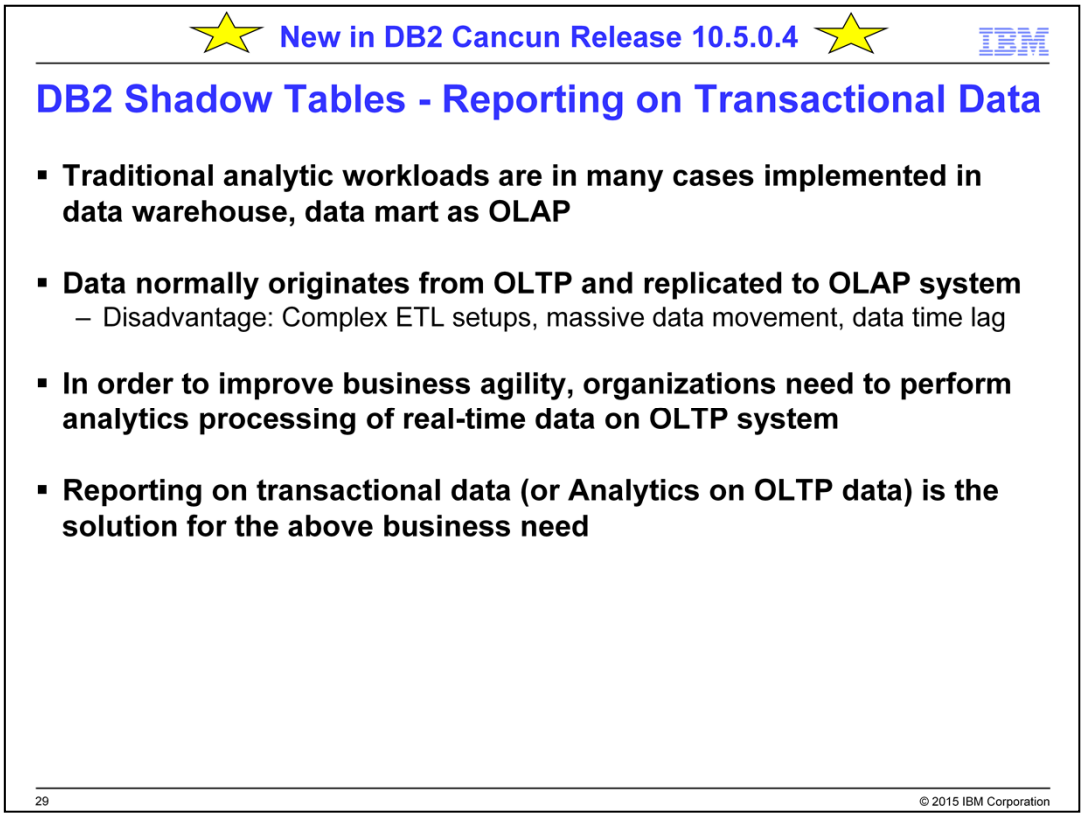

Analytics on OLTP data environments are those where you will have both OLTP and OLAP SQL execution. In some cases the environment is such that you can get excellent performance for one, but not so good for the other. So often customers are asked to make a tradeoff in execution of their SQL – one type over the other.

DB2 Shadow Tables in an OLTP environment are meant to provide a single solution which you can get good execution of your SQL queries regardless of the type.

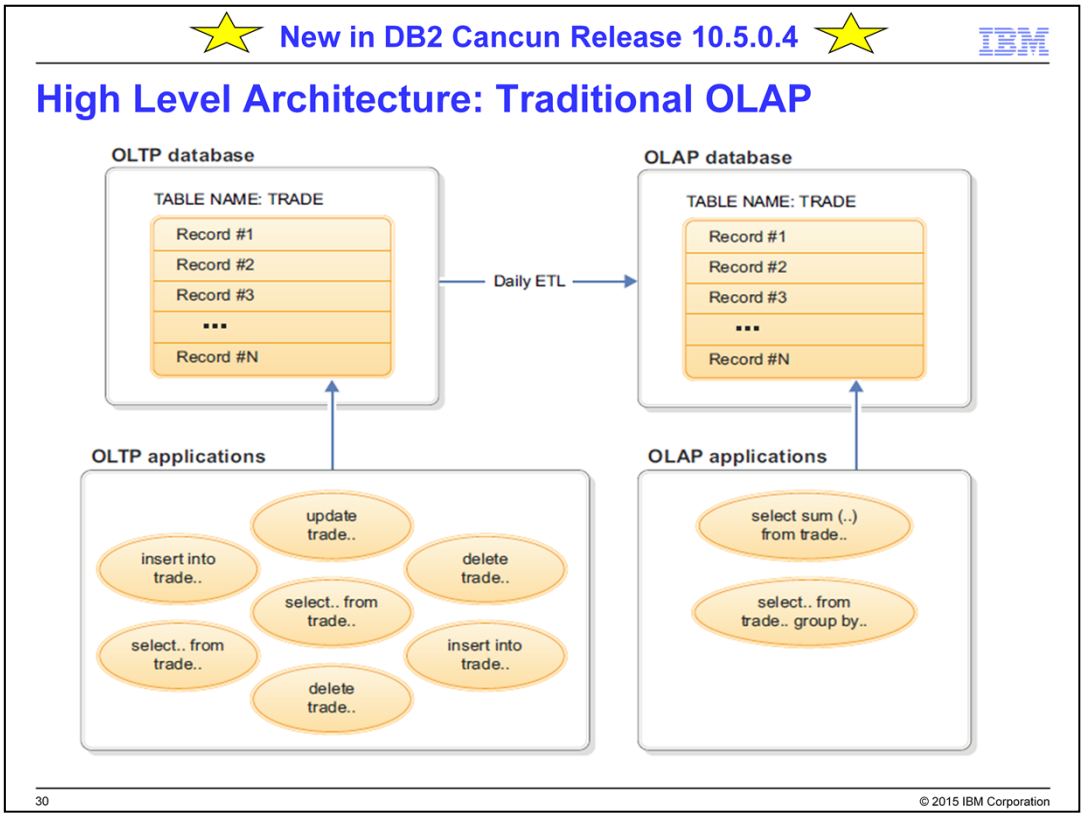

Here is an illustration of the traditional OLTP vs. OLAP setup.

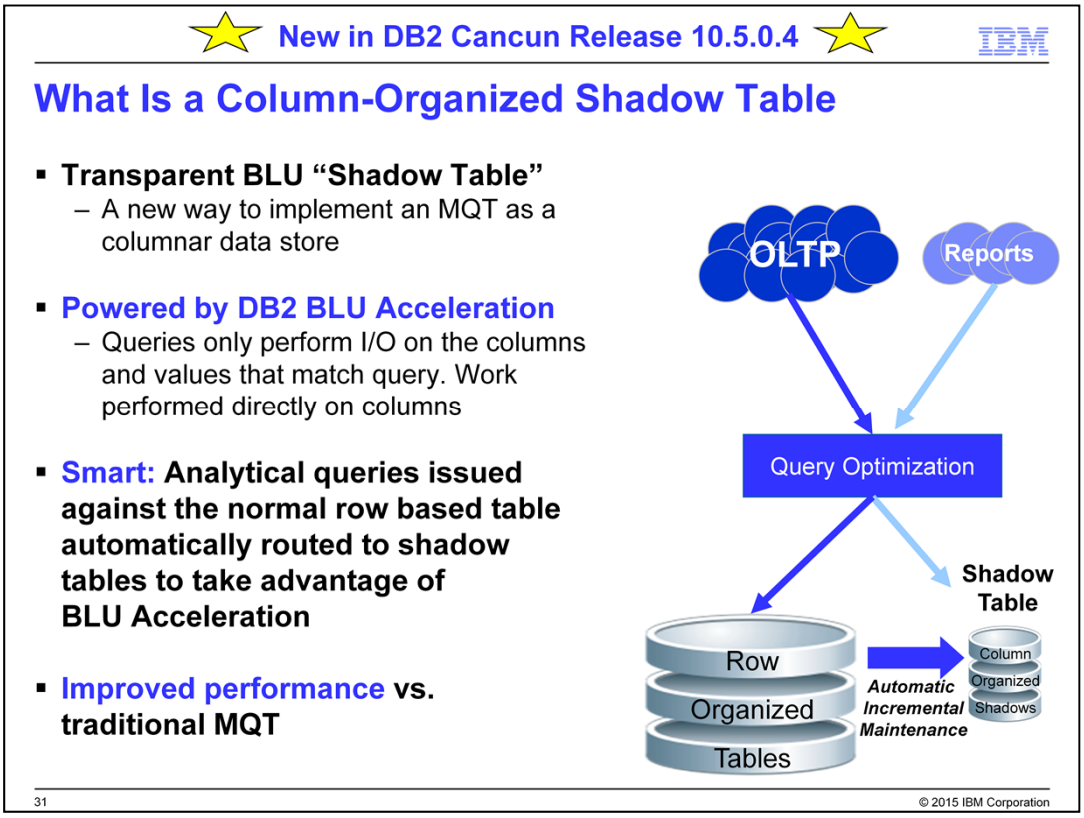

"Shadow Tables" allows you to have a column organized table on top of the row organized base table. The column organized version is either an exact duplicate of the row based table or a subset of the columns in the row table. Think of it like a column organized MQT on top of the base row organized table. When a query comes in to the system, the optimizer decides if the query is routed to the row based version of the table or the column based version of the table (just as it does automatic routing to MQTs today).

This is different than what Oracle is focused on in that these shadow tables are real physical materialized objects (not just in memory) but they have all the in-memory benefits that BLU tables have. This allows for running reports directly on your OLTP system.

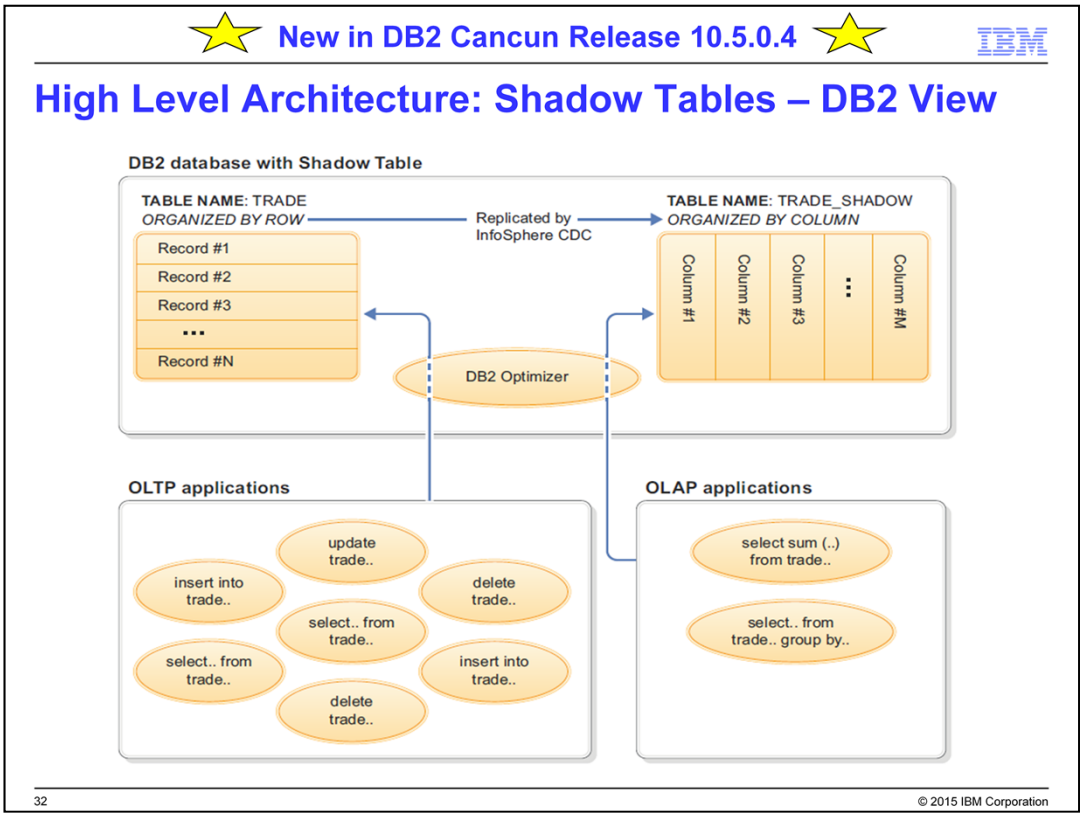

Here we see a different view of the CDC/DB2 Shadow Table architecture – solely from the DB2 side of things.

Here DB2 does all the hard work for you. If the SQL statement is OLTP related then DB2 will route the SQL statement to the row-organized tables to give the industry leading performance of the DB2 engine on row-organized tables. If the query is more OLAP centered and could benefit from column-organized BLU table execution, then DB2 will route the query to the DB2 Shadow Tables.

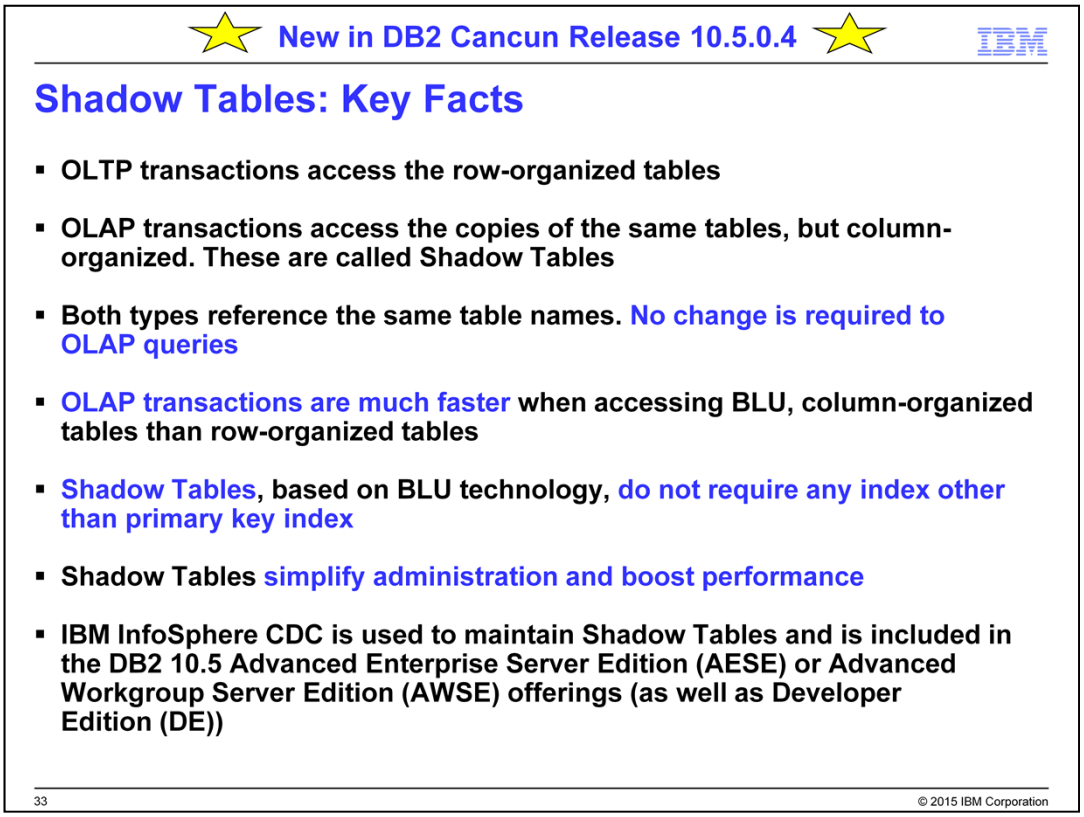

Here we highlight some of the key facts behind DB2 Shadow Tables and why they can assist in the overall performance of the Analytics on OLTP data environment.

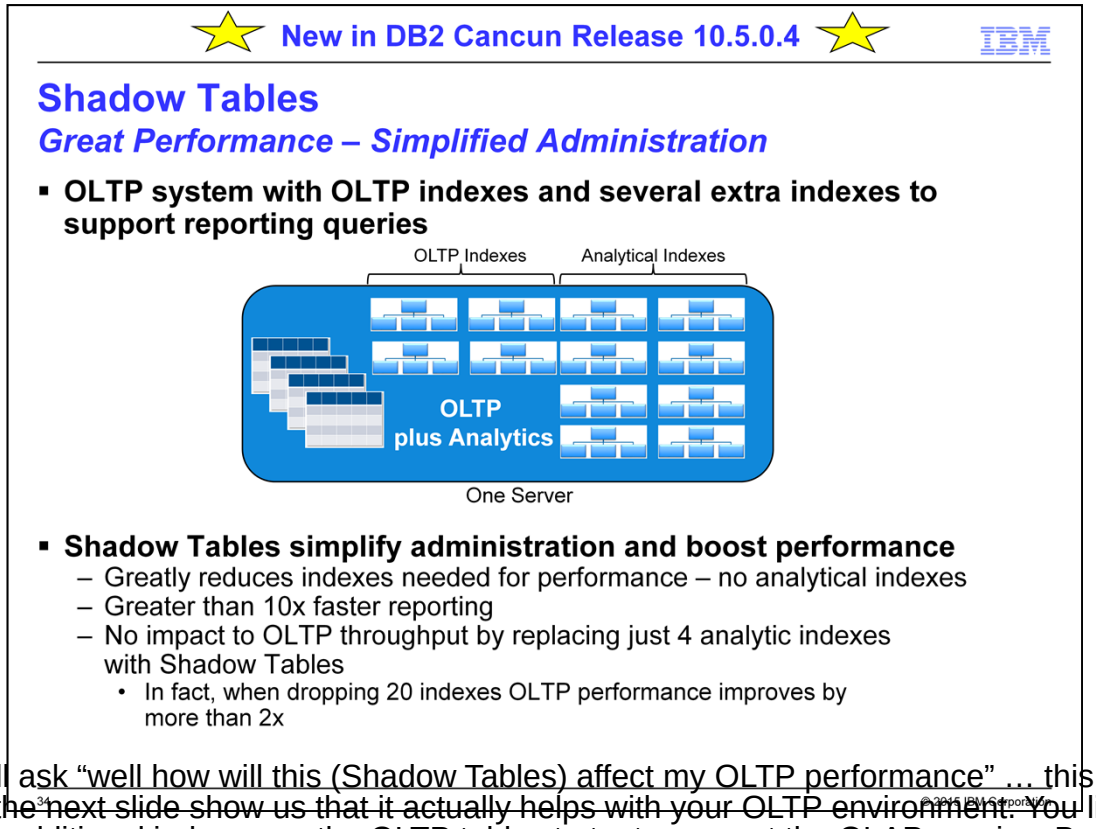

Many will ask "well how will this (Shadow Tables) affect my OLTP performance" … this slide and the next slide show us that it actually helps with your OLTP environment. You likely have additional indexes on the OLTP tables to try to support the OLAP queries. By adding Shadow Tables to your environment you can get rid of these indexes! If you can replace (remove) at least 4 of these indexes you will not see any performance issue (see next slide). In fact the more indexes you remove the more performance gain you will see in your OLTP environment side of things!

Shadow Tables (Example)

"Shadow Tables" allows you to have a column organized table on top of the row organized base table. The column organized version is either an exact duplicate of the row based table or a subset of the columns in the row table. Think of it like a column organized MQT on top of the base row organized table. When a query comes in to the system, the optimizer decides if the query is routed to the row based version of the table or the column based version of the table (just as it does automatic routing to MQTs today).

This is different than what Oracle is focused on in that these shadow tables are real physical materialized objects (not just in memory) but they have all the in-memory benefits that BLU tables have. This allows for running reports directly on your OLTP system.

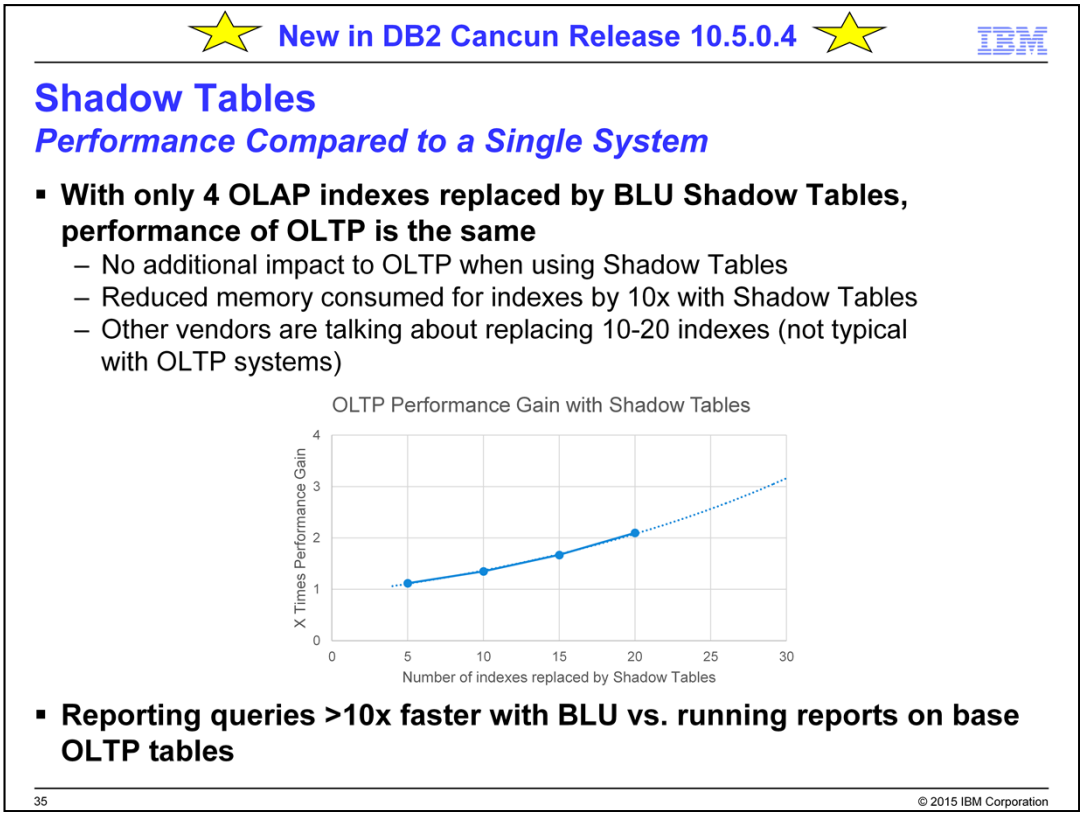

Here we have some performance tests that were run where the 4 OLAP indexes were replace with Shadow Tables (which do not need indexes remember!) and how the performance got even better the more OLAP indexes we got rid of. So Shadow Tables help greatly not only with your OLAP environment, but also with your OLTP environment! Win Win and WIN!

Shadow Table – Performance

Note that these are lab results and do not reflect customer experiences at this time.

DB2 with BLU Acceleration Customer Quotes

DataProxy LLC

"Our initial results using ***DB2 with BLU Acceleration*** *were impressive; massive analytic* ***query speed-up, up to 273x improvement.*** *With the* ***Shadow tables that are part of the DB2 Cancun Release,*** *we can now leverage that* ***great analytic performance in our transactional ODS*** *and let DB2 do all the work for us.* ***One database, one system, no extra effort on our part!"*** – Ruel Gonzalez - Information Services

"In the telecom business, reporting on data as it is happening is critical to our clients success. We need to be able to ***run complex reports directly on our transactional data*** *to provide a competitive advantage for our business.* ***Shadow Tables****, in the DB2 Cancun Release,* ***gives us that ability****. They are simple to create and very easy to manage and our end users are now* ***reporting directly on our OLTP data without having to add in a separate BI system****."* – Paul Peters, Lead Database Administrator, VSN Systemen BV

DB2 with BLU Acceleration Customer Quotes (cont.)

*"When we first started using BLU Acceleration in DB2 10.5 we were very pleased with the **10x compression** and the **10x performance improvement**. We are now using the DB2 Cancun Release, specifically **Shadow Tables**, to allow our clients to **run reports directly on top of transactional tables**. The **results are delighting my end users** and we **don't see any impact to our transactional performance**."*

– Paul Peters, Lead Database Administrator, VSN Systemen BV

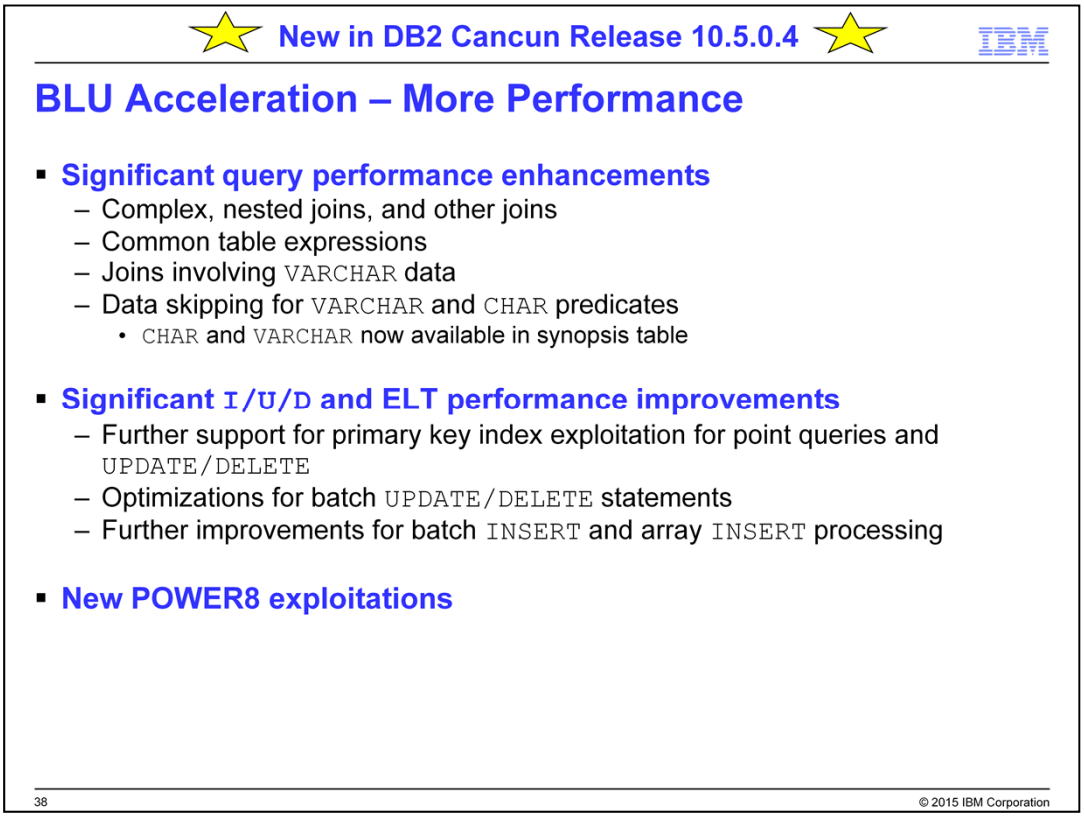

On the performance enhancements side there is more work being pushed down into the encoded space. This means that we can do more processing on data while it is still compressed and in columnar rather than row format. Support for common table expressions on BLU tables as well as adding some VARCHAR columns to the synopsis table so they can benefit from data skipping.

There are also enhancements in memory management where by have improved the algorithms to make better use of available memory.

We also have improvements to update and delete performance by leveraging the primary key index for single row lookups. This greatly improves the ingest performance and the ELT performance of BLU tables.

There are also enhancements in POWER8 that DB2 BLU takes advantage of for SIMD and other processing capabilities.

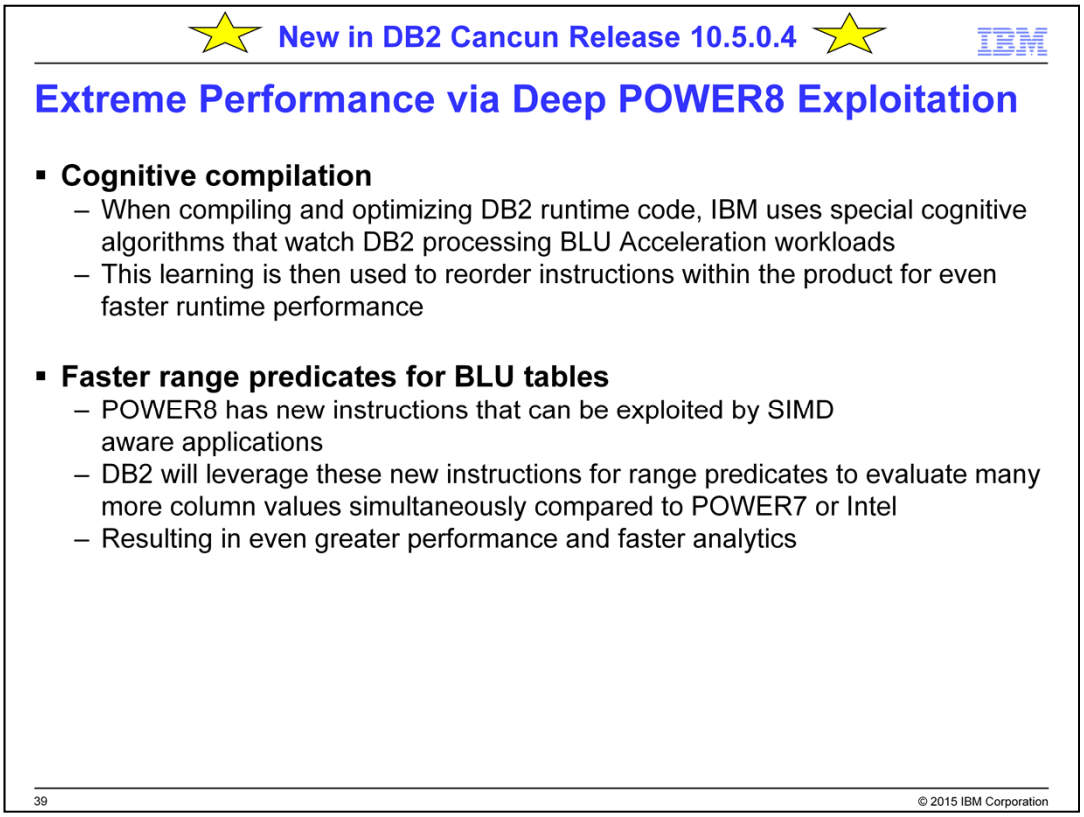

In DB2 there is a step run after the product is compiled to look at how the instructions are executed at runtime. This uses an IBM patented learning algorithm that profiles the running code to optimize the binary libraries for improved performance. For example does an IF statement usually take the first action or the ELSE section of code? By using cognitive profiling, the optimizer learns how DB2 executes it's binary instructions (and in DB2 Cancun Release specifically how it executes BLU instructions). The code is reordered to improve performance by up to 20% for certain workloads.

Finally in Power8 there are new 128bit instructions that allow BLU to process even more data in a single register and therefore improve the SIMD processing which is already an advantage of BLU acceleration. The result is faster query processing especially for range predicates with BLU Acceleration.

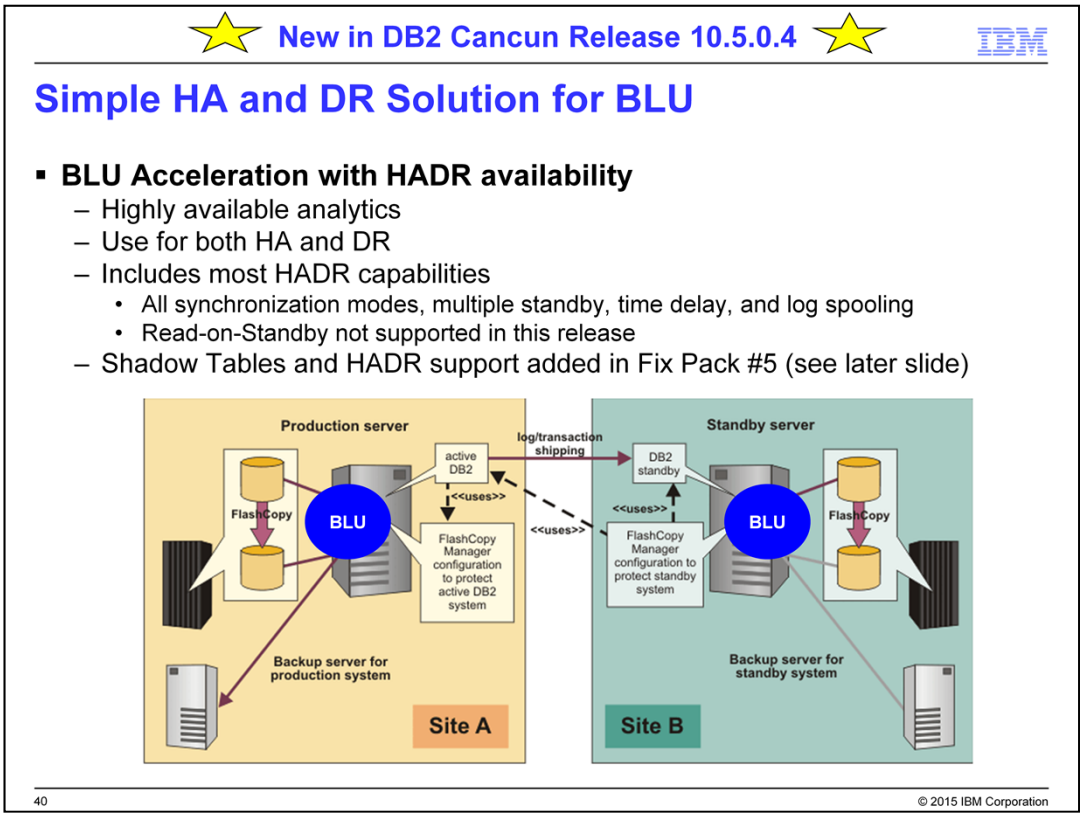

In DB2 Cancun Release 10.5.0.4 is the ability to use HADR on databases that have column organized tables in them. All sync modes will be supported as will multiple standbys along with time delay and log spooling.

There will be an ability to do HA and DR for column-organized shadow tables – details are still forthcoming there.

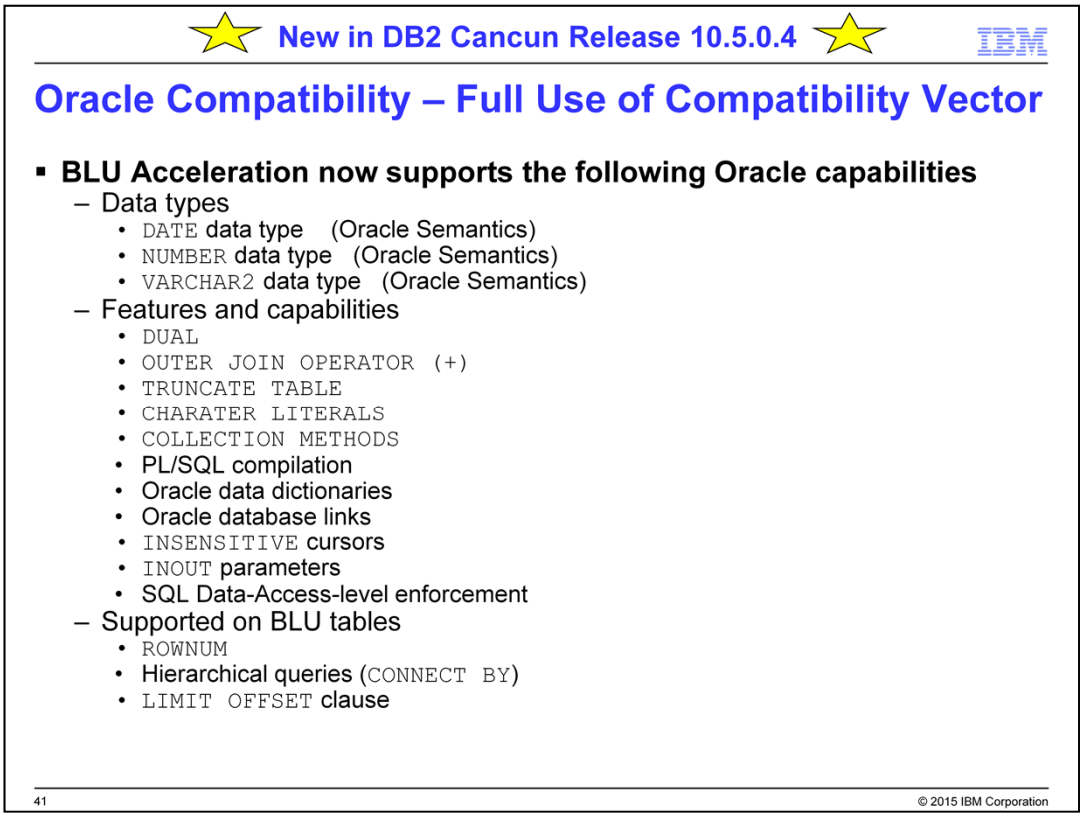

In DB2 Cancun Release 10.5.0.4 we lift the Oracle Compatibility Vector restriction on BLU Acceleration column-organized tables. Note however that some items, while supported now, are not pushed down into the columnar execution space (the last three items on the slide).

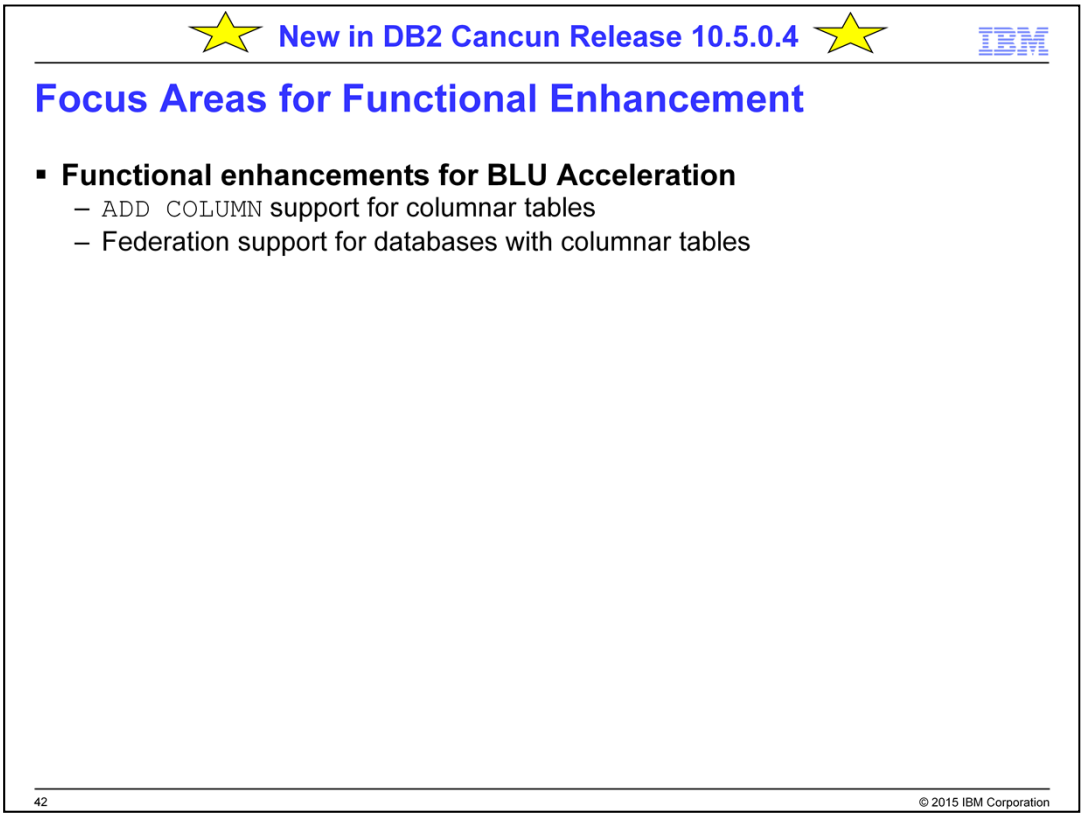

Finally, some additional restrictions are lifted for columnar tables.

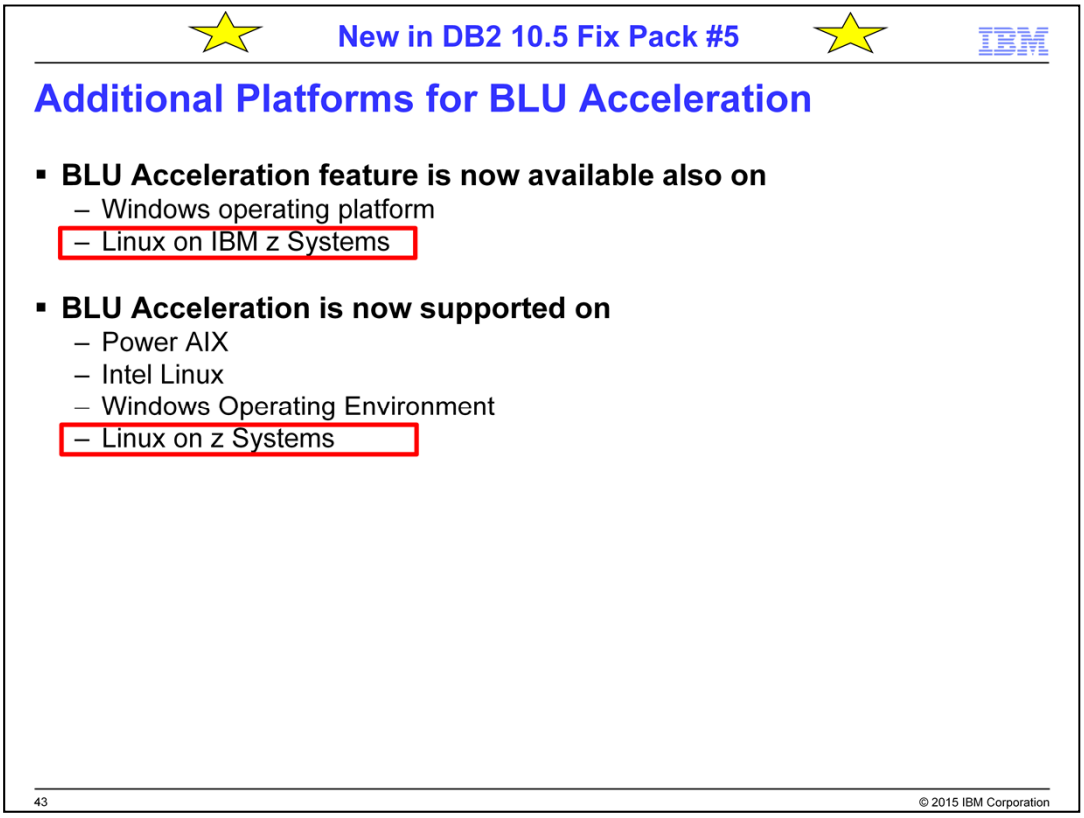

BLU is now available on both Windows and zLinux (Linux on zSeries). This in addition to previous platforms of Power AIX and Intel Linux.

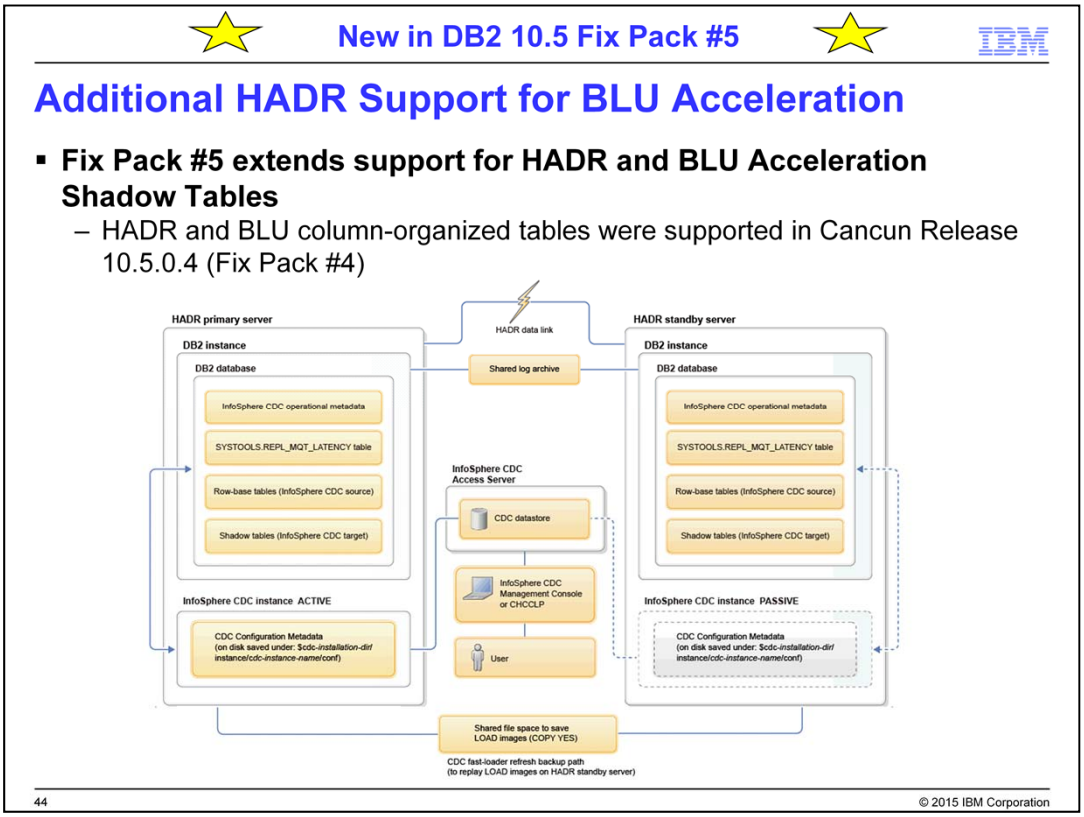

HADR for column-organized tables was added in Cancun Release 10.5.0.4 (Fix Pack #4). This Fix Pack #5 adds official support for HADR and Shadow Tables.

Some comments on this support

-- After the fail-over Shadow Tables are available

-- CDC is configured to be Active-Passive ... Passive on the Standby node

-- CDC will need to be started on the *new* primary after any fail-over event

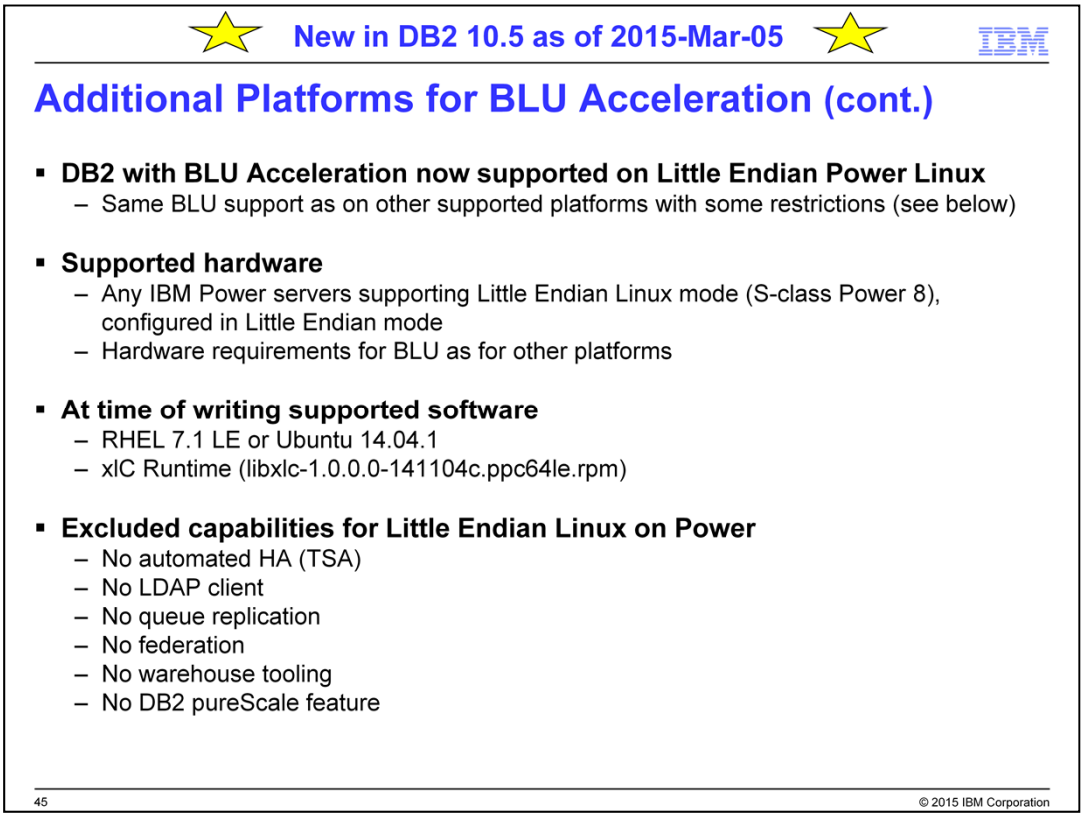

As of March 5, 2015, DB2 LUW and DB2 with BLU Acceleration are supported on Little Endian (LE) Linux on Power. This slide provides some details and restrictions of this support.

DB2 10.5 for LUW (non BLU Acceleration)

In DB2 Cancun Release 10.5.0.4 we added some things to DB2 LUW in general that are not specifically related to BLU Acceleration – this section will cover these items. Additional detail for some of the items will be provided

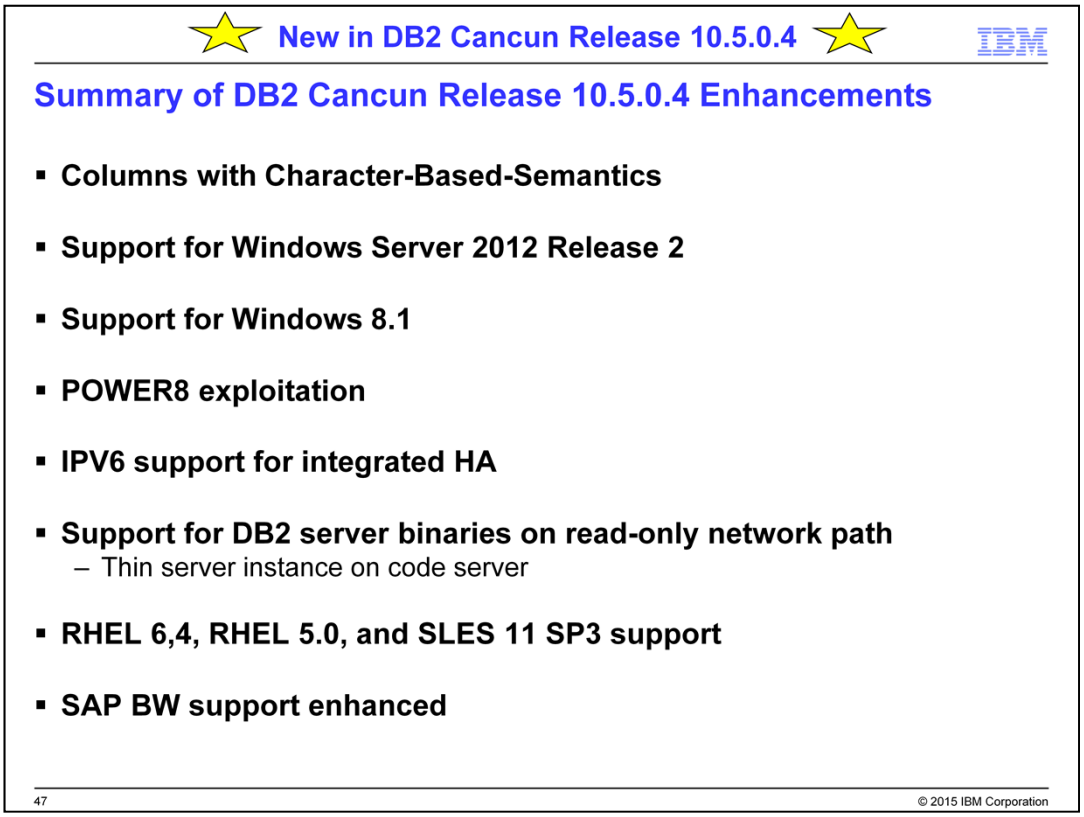

This slide lists some of the enhancements in the non-BLU Acceleration space in DB2 Cancun Release 10.5.0.4. Not all of these items are discussed in detail in later slides.

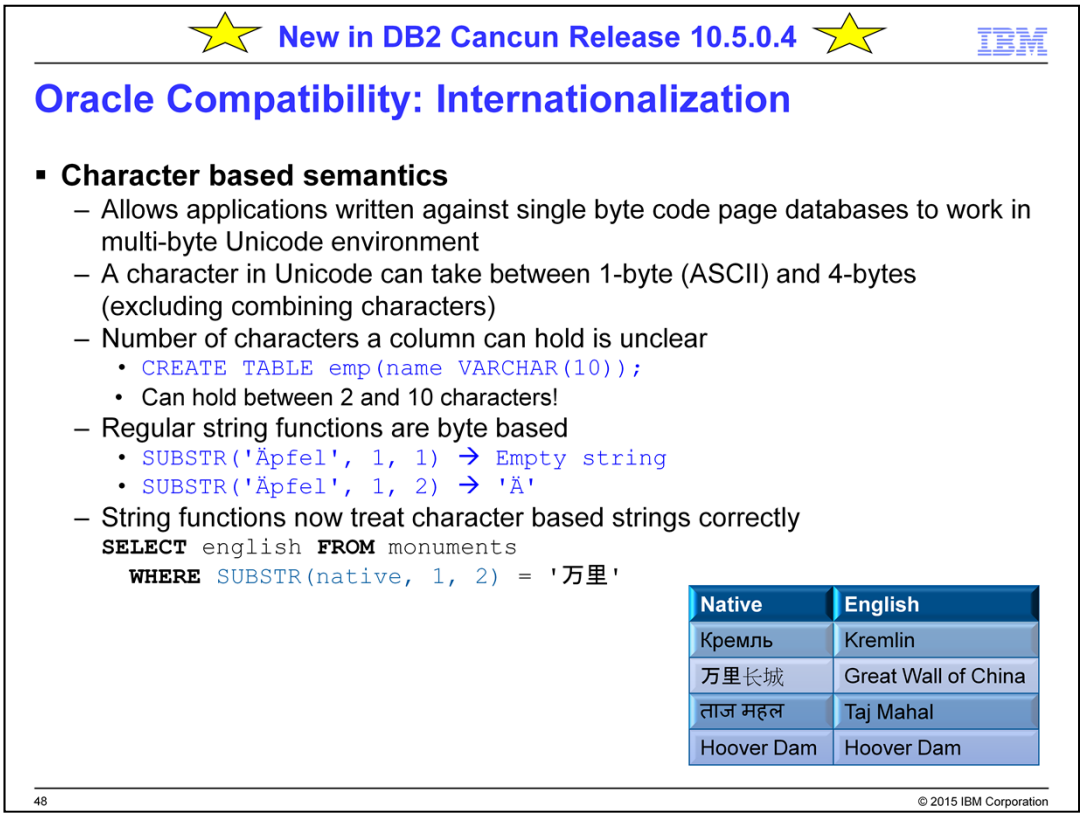

Character Based Semantics

•Applications written in Single-Byte code pages will not return the proper result when dealing with Unicode (multi-byte) languages. The reason for this is that many of the functions that applications use are based on byte-boundaries (SUBSTR) but Unicode characters can be from 1 to 4 bytes in length. A column that is 10 bytes wide could contain between 2 (4 byte Unicode) to 10 (1 byte) characters. The function SUBSTR(string,1,1) should return the first valid character from a string rather than the first byte.

•This new feature in DB2 will change the behaviour of the functions and comparison operators so that they look at character boundaries rather than byte boundaries – without having to change the existing application code.

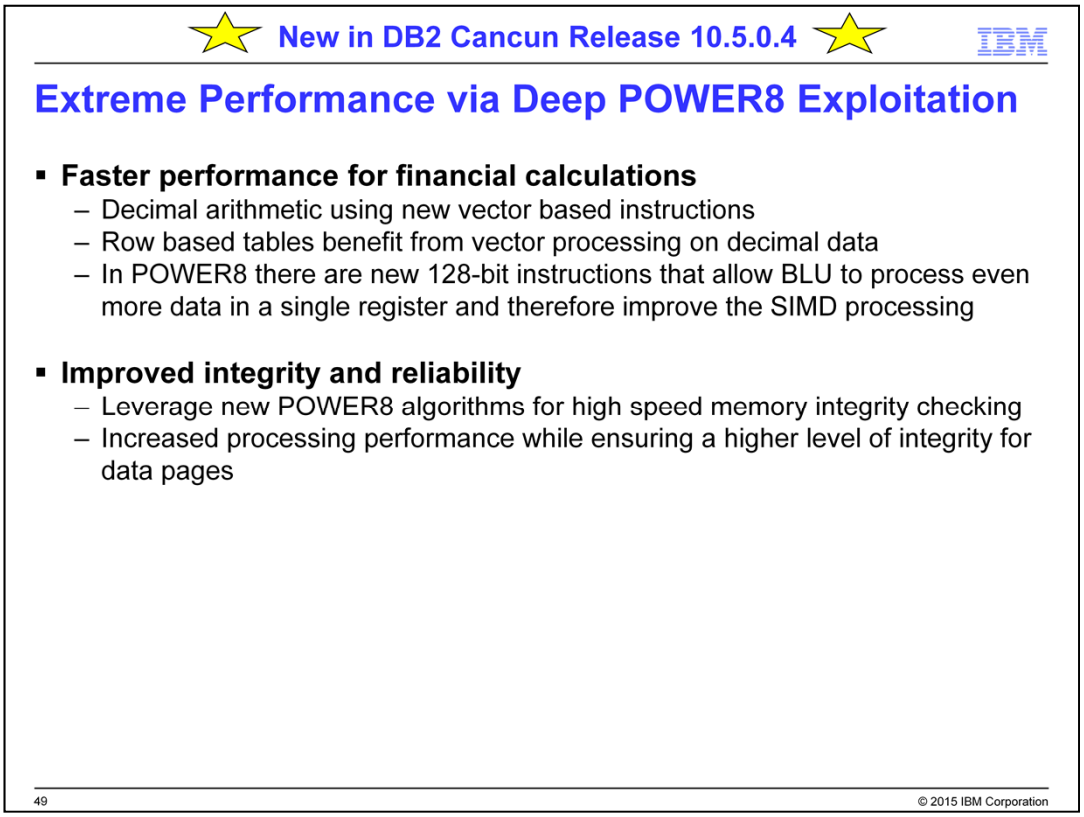

All of the exploitation of Power8 is very deep technical enhancements that are not easily described (low level algorithmic enhancements, leveraging Power8 instructions, latch code, etc). This slide describes the results of these enhancements which in general is increased performance. There are several exploitations that result in specific advantages for DB2 on Power8 over competitors.

The first is the use of vector processing instructions for row based tables that improve decimal arithmetic performance. Since decimal data is typically used for financial elements in a database (price, tax, bid, ask, etc) we should consider this enhancement in business terms as faster performance for financial calculations.

The second advantage is the use of high speed Power8 algorithms that allow for very fast data page checking. This allows for improved data page checking in memory and the use of integrity checking in more areas of the code for increased reliability. Operations like RUNSTATS, REORG, queries that access many data pages and heavy insert workloads will benefit from this high speed integrity checking.

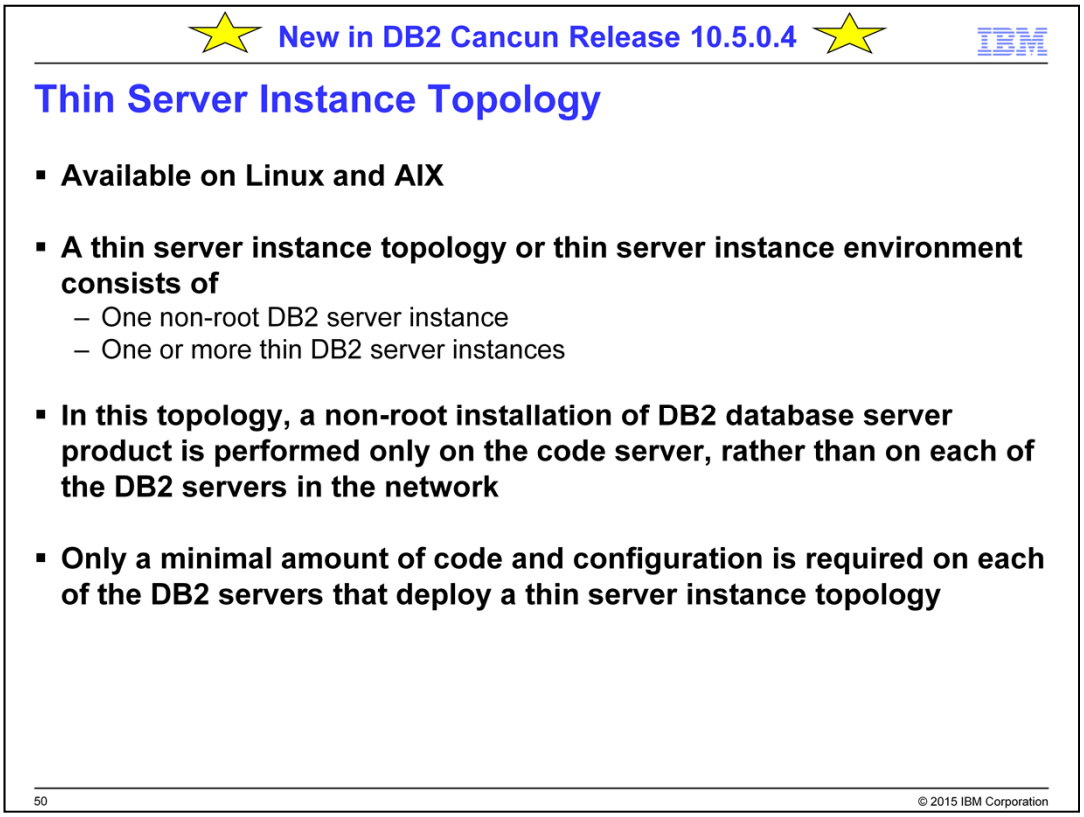

A thin server instance topology or thin server instance environment consists of one non-root DB2® server instance and one or more thin DB2 server instances. In this topology, a non-root installation of DB2 database server product is performed only on the code server, rather than on each of the DB2 servers in the network. Only a minimal amount of code and configuration is required on each of the DB2 servers that deploy a thin server instance topology.

You can create a thin server instance either locally on a code server or on a remote server.

In the first case, perform a non-root installation of DB2 database server product on a code server and share the DB2 installation path as a read-only copy to other non-root users in the network through Network File System (NFS). Multiple non-root users from the same system (code server) can access the read-only copy of the shared DB2 installation path and create a thin server instance locally on the code server.

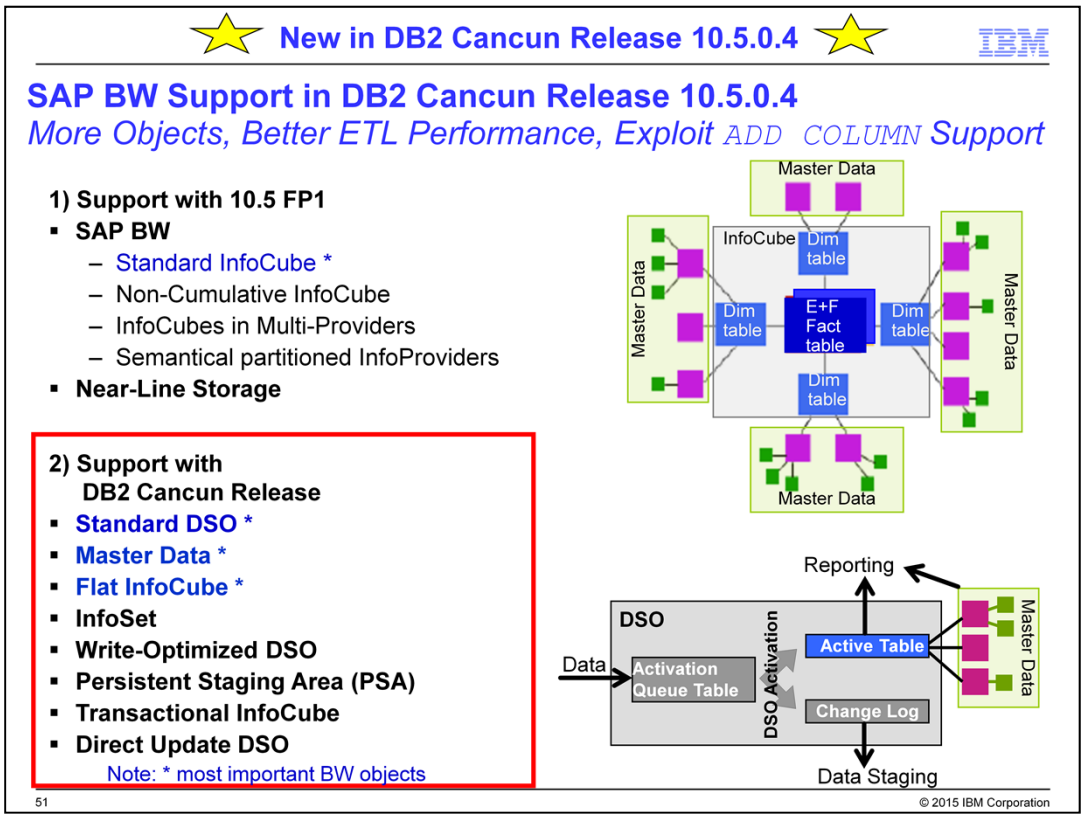

DB2 Cancun Release 10.5.0.4 adds more SAP BW support and function for BLU Acceleration. More important BW objects are now supported as column-organized

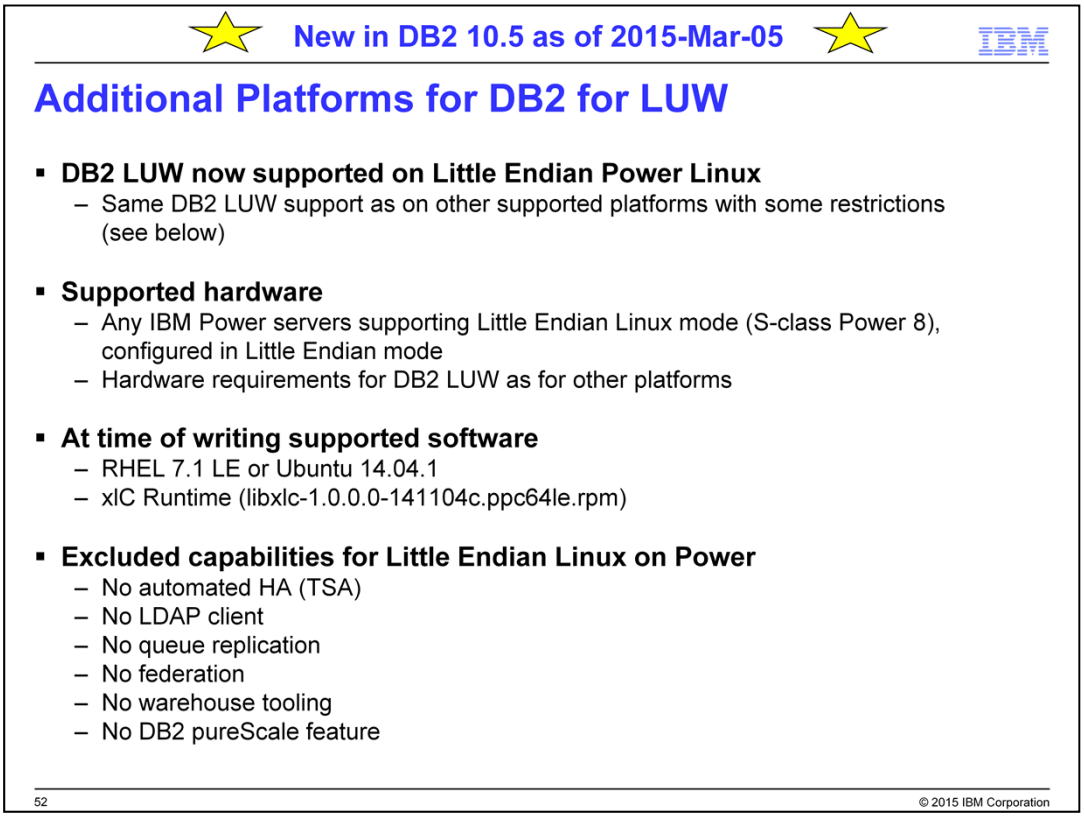

As of March 5, 2015, DB2 LUW and DB2 with BLU Acceleration are supported on Little Endian (LE) Linux on Power. This slide provides some details and restrictions of this support.

Increasingly, businesses desire or are mandated to encrypt sensitive data to meet organizational or regulatory requirements. The DB2 Encryption Offering (also knows as DB2 Native Encryption) assists organizations to meet those requirements by providing, natively within the DB2 engine itself, encryption capabilities that encrypt data at rest for the entire database, including backup images, log files, etc. DB2 Native Encryption meets the requirements of NIST SP 800-131 compliant cryptographic algorithms and utilizes FIPS 140-2 certified cryptographic libraries. It is available as part of the Advanced Editions (AESE and AWSE) and DB2 Express-C. It is also available as a purchasable offering on DB2 Express Server Edition, DB2 Workgroup Server Edition, and DB2 Enterprise Server Edition.

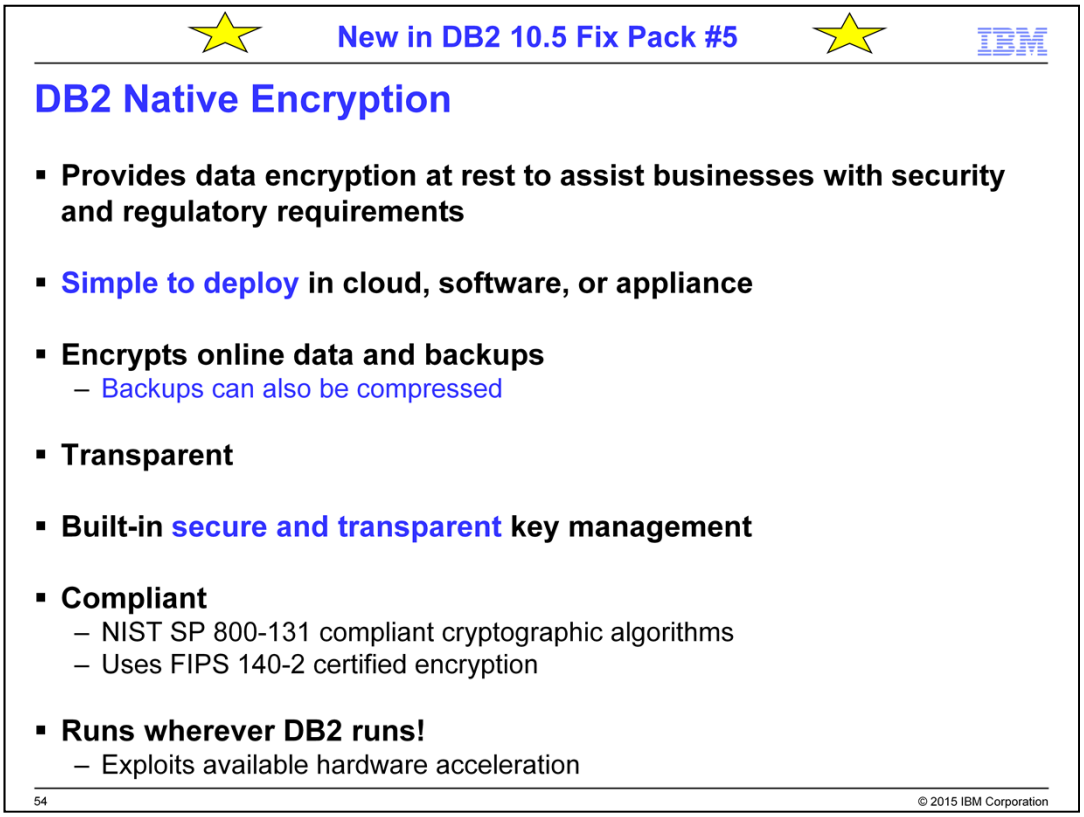

1. **DB2 Native Encryption is transparent to your applications and schemas**
   DB2 Native Encryption encrypts your data as it is written to disk. It is implemented within the DB2 kernel itself. This means that encryption is totally transparent to your applications and database schemas
2. **Key Management is secure and transparent**
   DB2 Native Encryption uses a standard two-tier model for key management. The Data Encryption Key (DEK) constitutes the first tier. The DEK is the actual key used to perform data encryption. The DEK is then encrypted with a second key and stored within the database (or backup image). The second key is called the Master Key (MK) and constitutes the second tier. In the industry, this model is referred to as envelope encryption. The MK is stored outside the database in a Public-Key Cryptography Standards (PKCS#12) compliant keystore
3. **DB2 Native Encryption encrypts both your online data and your backup images**
   To encrypt your online data, you need to create your database with the new ENCRYPT option of the CREATE DATABASE command. By default, your database encryption uses Advanced Encryption Standard (AES) in Cipher-Block Chaining (CBC) mode with a 256 bits key. But other encryption algorithms and key sizes are available. Every database has its own unique Data Encryption Key (DEK). The encryption for backup images is independent of online database encryption
4. **DB2 Native Encryption employs certified and compliant cryptography, and exploits hardware acceleration for cryptographic operations.**
   Certification and compliance are critical when it comes to encryption solutions. DB2 Native Encryption uses FIPS 140-2 certified cryptographic modules. Additionally, only cryptographic algorithms that are compliant with NIST SP 800 - 131 are employed by DB2 Native Encryption. Similarly, performance is critical for database workloads. DB2 Native Encryption is capable of exploiting recent innovations in processor technology such as the

Intel AES-NI. This exploitation is automatically detected and transparently exploited by DB2 Native Encryption

5. **DB2 Native Encryption supports encrypting your existing DB2 databases**
   It is possible to convert an existing unencrypted database into an encrypted database.

DB2 Native Encryption (cont.)

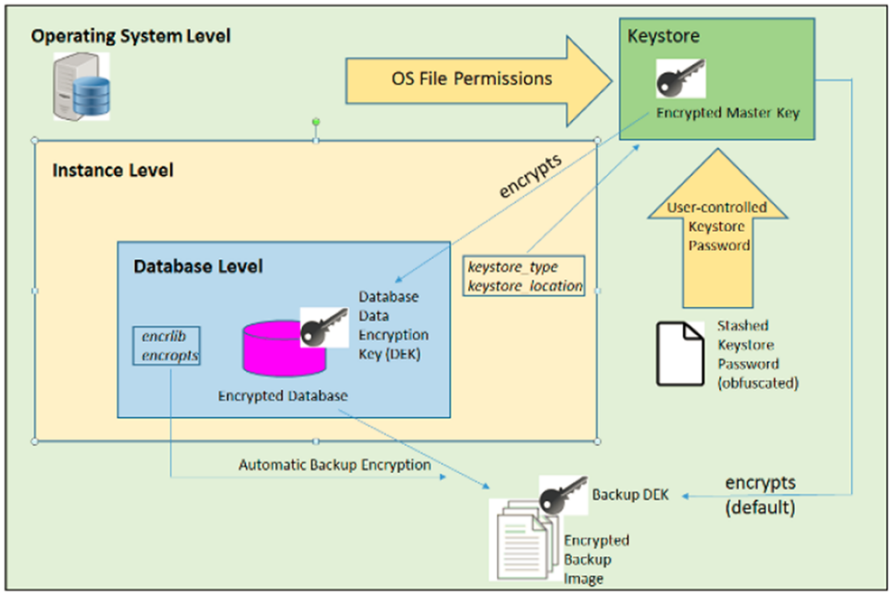

This diagram gives you an overall pictorial view of the DB2 Native Encryption environment and its key components

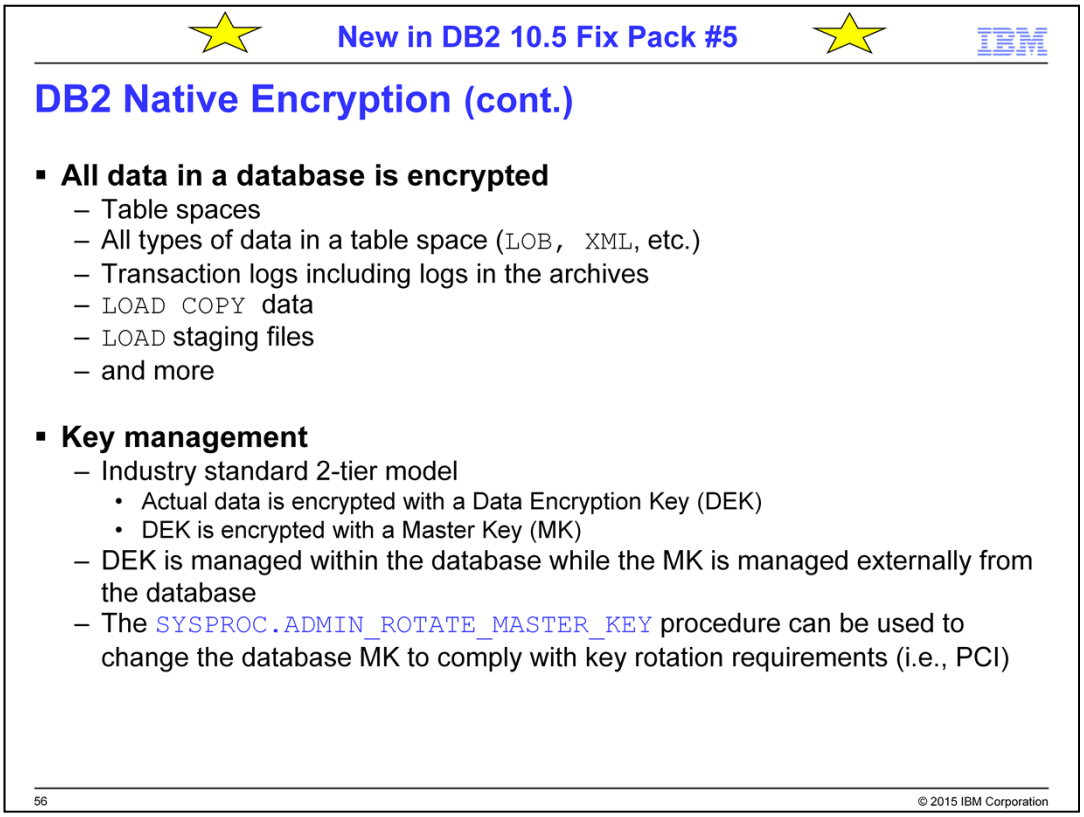

With DB2 Native Encryption, all parts of the database are protected, from user data, to system data, to logs, to LOAD COPY data and more.

DB2 Native Encryption uses a standard two-tier model for key management. The Data Encryption Key (DEK) constitutes the first tier. The DEK is the actual key used to perform data encryption. The DEK is then encrypted with a second key and stored within the database (or backup image). The second key is called the Master Key (MK) and constitutes the second tier. In the industry, this model is referred to as envelope encryption. The MK is stored outside the database in a Public-Key Cryptography Standards (PKCS#12) compliant keystore. There are two security protection measures for your keystore. The first is file permissions. You need to make sure that only the DB2 instance owner has read/write access to the keystore. The second is encryption of the actual content of the keystore. You need to make sure you create your keystore with the password option. The content of the keystore (i.e., your master keys) is encrypted using a symmetric key derived from that password using a hashing algorithm. Without the password, the content of the keystore cannot be decrypted. DB2 Native Encryption also allows you to rotate your database MK to comply with your corporate security policies. You rotate your database MK by calling the new ADMIN_ROTATE_MASTER_KEY

procedure. The procedure decrypts your database DEK with the old MK and then re-encrypts it with the new MK.

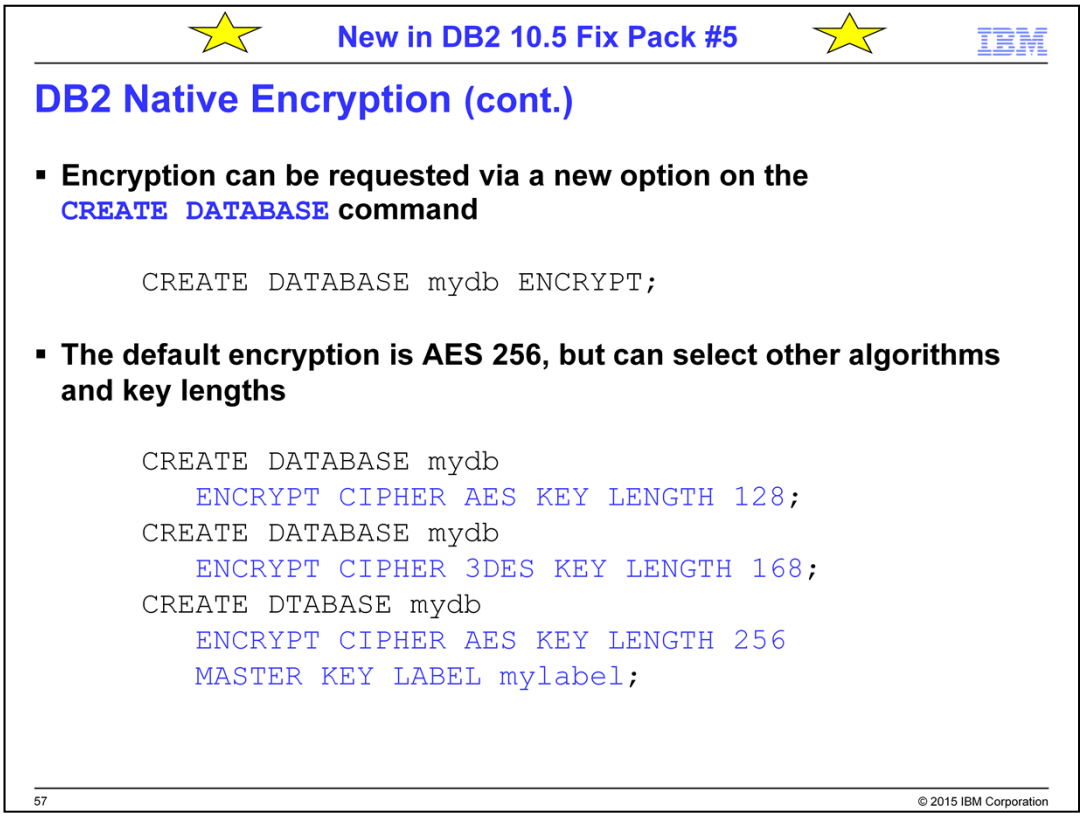

Using DB2 Native Encryption is very simple.

To encrypt your online data, you need to create your database with the new ENCRYPT option of the CREATE DATABASE command. By default, your database encryption uses Advanced Encryption Standard (AES) in Cipher-Block Chaining (CBC) mode with a 256 bits key. But other encryption algorithms and key sizes are available. As for online data, every backup image has its own unique DEK. By default, the backup image DEK is encrypted with the database MK although a different MK can be used.

DB2 Native Encryption supports encrypting your existing DB2 databases

It is possible to convert an existing unencrypted database into an encrypted database. The approach is as follows. First, you take a backup of your existing database using the BACKUP DATABASE command. Then, you restore that backup image into a new database using the RESTORE DATABASE command. When invoking the RESTORE DATABASE command, you specify the new ENCRYPT option. This new option mirrors exactly the ENCRYPT option of the CREATE DATABASE command. That is, the default is that your new database will be encrypted using AES 256. But you can choose different algorithms and key sizes if so desired.

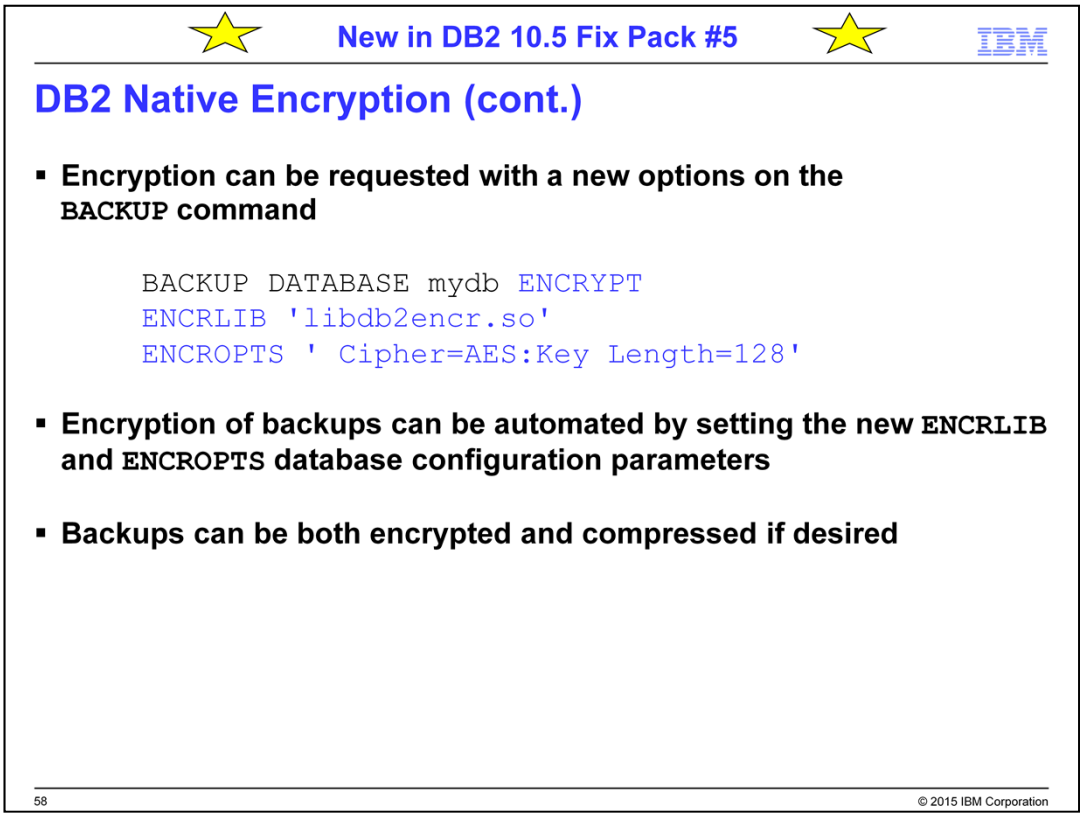

The encryption for backup images is independent of online database encryption. That is, you can choose to encrypt your backup images even if your online database is not encrypted. You can request an encrypted backup image by explicitly specifying the ENCRYPT option of the BACKUP DATABASE command. Alternatively, you can enforce and automate backup images encryption by configuring the new ENCRLIB and ENCROPTS database configuration parameters. For encrypted databases, these two parameters are automatically configured by DB2. This means that when your database is encrypted, the default is that your backup images are automatically encrypted. Also, by default a backup image is encrypted with AES 256, but a different algorithm and key size can be chosen.

DB2 10.5 pureScale

We now take you through some of the enhancements of DB2 pureScale.

DB2 10.5 pureScale Enhancements
Enhanced availability, optimized for OLTP Workloads

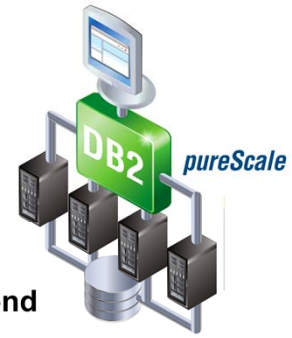

- **DB2 pureScale**
  - Robust infrastructure for OLTP workloads
  - Provides improved availability, performance, and scalability
  - Transparent scalability beyond 100 nodes[1]
  - Leverages z/OS cluster technology
- **NEW pureScale enhancements in GA and beyond**
  - Online member add
  - No planned downtime required for pureScale maintenance updates[2]
  - HADR designed to failover in seconds[3]
  - TCP/IP "sockets" based interconnect
  - Virtualized pureScale deployments using VMware and KVM

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.

60

In additional to BLU Acceleration, DB2 10.5 include enhancements to DB2 pureScale. pureScale is built-in database cluster technology that provide highly scalable and available data services for online transaction processing workloads.

-This is a shared disk cluster architecture

-Runs on both Unix (AIX) and Linux (x86)

-Leverages DB2 for z/OS cluster management technology – which is well known as the best in the industry

-Provides high availability built into the architecture

-Provides very high level of scalability also built into the architecture (we have scaled to >100 members)

-Provides application transparency (application does not need to be cluster aware, as with Oracle)

-With DB2 10.5, we have these new pureScale enhancements -

- HADR support to distances over 1000 KMs
- disaster recovery options
- backup and restore between pureScale and non- pureScale environments
- rolling online fix pack updates

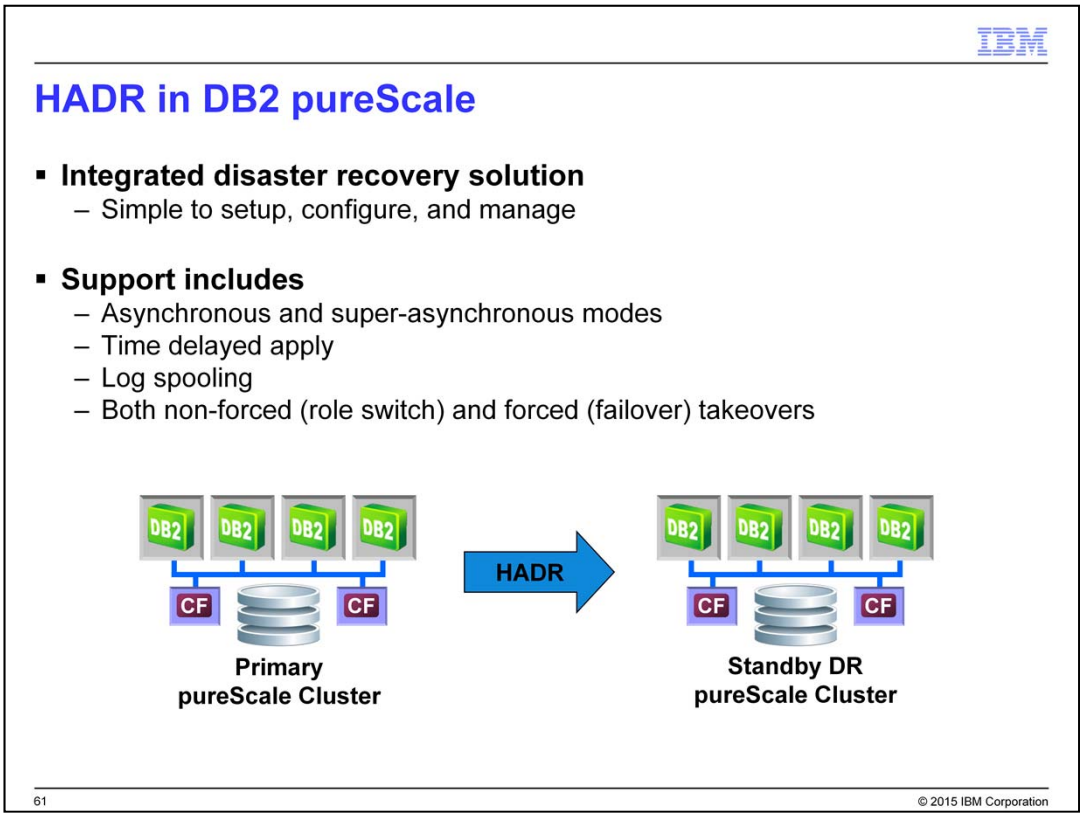

Previously, HADR was not supported in DB2 pureScale environments. However, DB2 pureScale itself is already inherently highly available and so it was just the DR (disaster recovery) piece of HADR that was missing. HADR support for pureScale was added in DB2 10.5. HADR in DB2 pureScale supports a single standby, time delayed apply on the standby, and log spooling.

You can also do forced or non-forced takeovers on the standby system. The TAKEOVER HADR command, which you can issue on the standby database only, changes the standby database to a primary database. When you do not specify the BY FORCE option, the primary and standby databases switch roles. When you do specify the BY FORCE option, the standby database unilaterally switches to become the primary database. In this case, the standby database attempts to stop transaction processing on the old primary database. However, there is no guarantee that transaction processing will stop. You are supposed to use the BY FORCE option to force a takeover operation for failover conditions only.

This HADR support initially includes asynchronous and super asynchronous modes (HADR modes are explained here: http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0011724.html). With the asynchronous mode, there is a guarantee of no data loss during non-forced takeovers.

An asynchronous DR (disaster recovery) solution such as this is perfectly suited to the needs of many of our customers. With HADR, our pureScale customers now have a DR infrastructure that is very simple to setup and manage. Of course, there is a class of customers that do have a business requirement for zero data loss or an RPO (recovery point objective) of zero with respect to DR site failover. These customers do require a synchronous DR solution and there are other options available to do that today with pureScale, such as storage-based replication and Geographically Dispersed pureScale Clusters (GDPC).

That said, does your customer really need synchronous DR? If you ask any customer what their RPO needs are, the immediate answer is often "zero data loss". However, once you've had a conversation with them and their requirements are truly understood, this is frequently found not to be the case. Also, constraints imposed on the solution often result in synchronous DR being dismissed as impractical. For instance, the latency introduced based on the distance to the DR site may have too much of an impact on the performance of the primary site. For this reason, use of synchronous HADR today (nearsync or sync) is limited to within 100 km and it is typically much less than that in practice. This makes it a "metro" type of DR solution which only protects customers from a limited class of disasters. Another constraint might be the cost of providing the necessary bandwidth and service levels for the pipe between the primary and disaster sites. This is needed for synchronous DR because delays in the data flow have an impact on the performance and responsiveness of the primary site. For all of these

reasons, customers often choose an asynchronous DR solution and pureScale's HADR capabilities are a perfect fit.

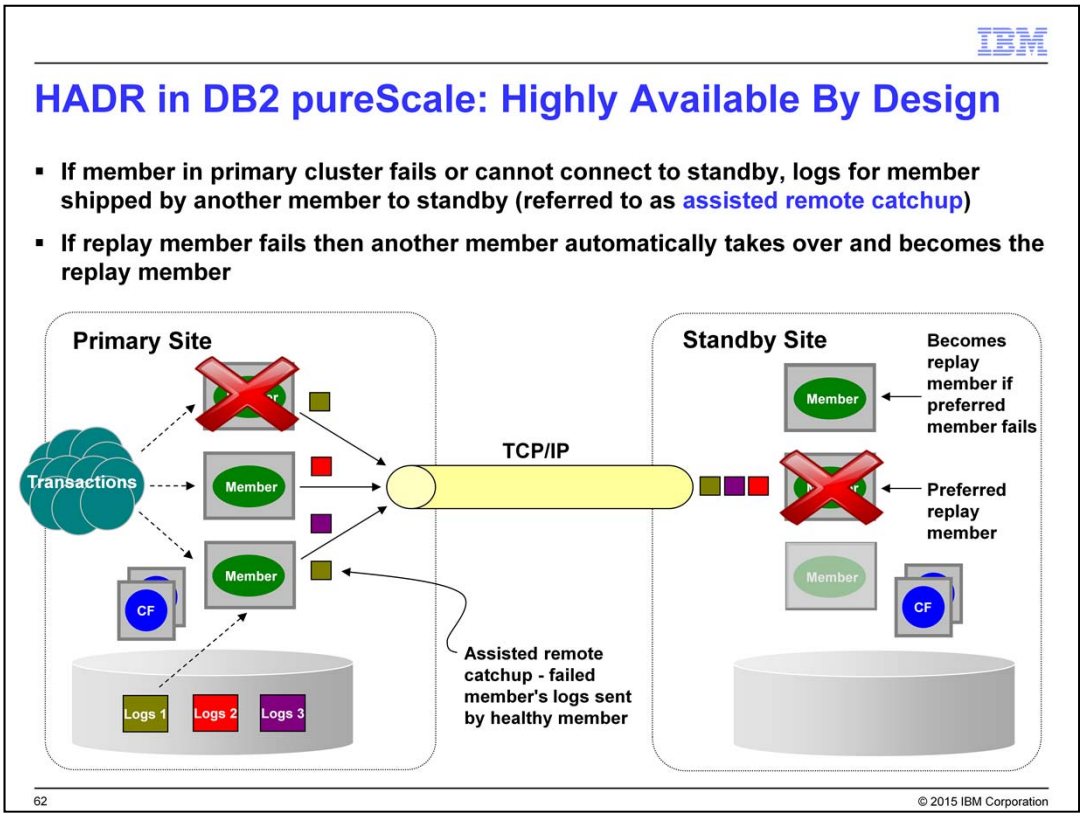

Because pureScale is intended to be a highly available environment, we've made HADR highly available as well with respect to its toleration of member failures on the primary and standby systems.

For instance, if a member in the primary cluster fails or cannot connect to the replay member on the standby then that member's logs are shipped to the replay member by one of the other primary members. This is referred to as assisted remote catchup.

Also, if the replay member on the standby goes down normally or abnormally, DB2 will automatically migrate the log replay duties to another member. As long as there is one online member at the standby, replay will continue. Once replay moves away from the preferred replay member, it won't automatically move back there if the instance gets started again. Also, there is no user control on selection among non-preferred members.

This slide shows an example of how two 3-member pureScale clusters are setup in an HADR configuration. When the animation starts, the logs from each of the three members are sent via TCP/IP to the secondary site (and the single replay member applies the changes; the database is not activated on the other standby members). Subsequently, the first member fails. If there are logs that need to be sent from that first member then this is done using another member (the third member in this particular example). Because the logs are stored on the cluster file system, members can read the log files for other members as necessary. This is referred to as assisted remote catchup. Then, if the replay member on the standby fails, another member will take over replay duties.

Rolling Database Fix Pack Updates

- **Transparently install pureScale fix packs in an online rolling fashion**
- **No outage experienced by applications**
- **Single `installFixPack` command run on each member/CF**
  - Quiesces member
    - Existing transactions allowed to finish (configurable timeout, default is 2 minutes)
    - New transactions sent to other members
  - Installs binaries
  - Updates instance
    - Member still behaves as if running on previous fix pack level
  - Unquiesces member

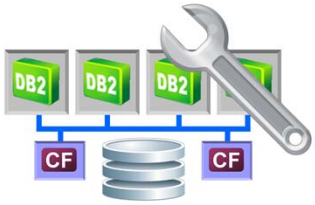

- **Final `installFixPack` command to complete and commit updates**
  - Instance now running at new fix pack level

63

Previously, when applying DB2 pureScale fix packs it was necessary to completely stop the entire cluster so that the fix pack could be applied to all of the members and CFs. This is no longer the case in DB2 10.5 and you can perform this kind of maintenance without bringing the cluster down.

It is termed a "rolling update" because you can perform the maintenance on one host at a time, rolling through each of the members and CFs. During the application of the fix pack to a particular host, only one member or CF is offline but at least one other member and CF should still be online and so the cluster itself is still online and able to perform work on behalf of applications.

The newly enhanced installFixPack command is used to do the update of the binaries on an individual host by host basis. The installFixPack command also has new options for committing the changes and for doing a pre-commit check first. The process of a rolling update is covered on the next slide.

Note that this does not imply "rolling upgrade" support. In other words, we will not necessarily support being able to do a rolling version/release upgrade from the DB2 10.5 release to whatever major release follows it.

Also note that you won't find much in the DB2 10.5 GA Information Center about rolling fix pack update support. The reason for this is that you can't really take advantage of it yet. Until FP1 is released, you can't actually apply a fix pack in a rolling fashion. If a customer happens to require a special build on top of GA, the special build will be installable online. To cover that case, we will be including the documentation directly within the special build in a PDF file.

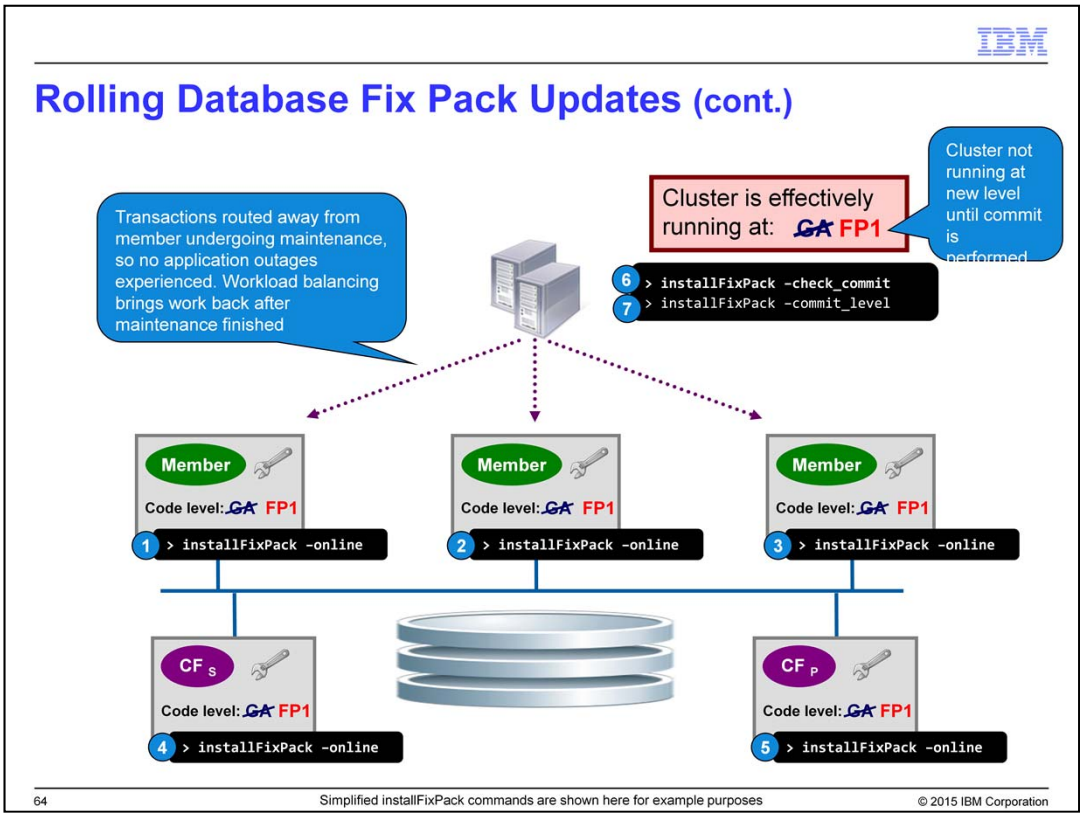

This slide shows an example of a rolling fix pack update. Initially, the instance starts at the GA code level and we will go through the process of updating the instance to the fix pack 1 (FP1) code level. The commands on the slide are just used to show the steps involved and don't include all of the options that you'd have to specify in practice.

It is suggested that you start with the members prior to doing the CFs. The reason being that if there is a problem then you are likely to see it with the members first and this way you can identify any serious issues before going through the process of updating the primary CF (more on this below).

You will start with one member and issue the installFixPack command with the appropriate options (e.g. *media-dir*/installFixPack -p *FP-install-path* -I *instance-name* -online -l *log-file-name*). This will quiesce the transactions running against it, allowing them to finish executing (but not allowing new ones to start), stop the member, install and update the binaries, and start the member back up. New transactions are not sent to the member undergoing maintenance, but once the maintenance process is complete it will start allowing them again. Workload balancing should start sending new work it's way. At this point, the installed and updated code is FP1 but the member is not allowed to use any new features that might exist in FP1. It's still running in GA mode.

This process would then need to be repeated for each of the remaining members, one at a time.

From here, you will move on to the CFs. It is suggested that you start with the secondary CF, allowing the primary to stay as primary for now (this way you just have a single primary CF->standby CF failover). You issue the installFixPack command with the appropriate options (e.g. *media-dir*/installFixPack -p *FP-install-path* -I *instance-name* -online -l *log-file-name*). This will stop the secondary CF, get it updated to the FP1 code level, and start it up again. You can then move on to do the primary CF next. You should first ensure, though, that the secondary has reached peer state before doing this (since the secondary will go into a catchup mode once it's restarted).

After all of this is done, all of the members and CFs are running with the new code level. However, they are still running in GA mode. To make the switch over to FP1 and have the instance actually running at this level, it requires committing the changes. First, though, you want to make sure that everything has in fact been updated successfully. This check can be done by executing installFixPack with a special check option (e.g. *media-dir*/installFixPack -check_commit -p *FP-install-path* -I *instance-name*). If everything is good you can then commit the changes using another option of installFixPack (e.g. *media-dir*/installFixPack -commit_level -p *FP-install-path* -I *instance-name* -l *log-file-name*). All of the commands, including checking and committing the changes is fully online. There is no cluster outage required at all.

Use the new ENV_GET_INSTANCE_CODE_LEVELS table function to obtain information about an online fix pack update in progress such as the state, the architecture level, and the code level of instances, members, and CFs. You can also use the db2pd -ruStatus command to get this kind of information. Consider running it after each step to ensure that things are in the state you expect them to be.

Canceling an online fix pack update on a member or cluster caching facility (CF) requires that you reinstall the fix pack software with the same code level as the current effective code level.

In an HADR environment, it is suggested that you update the standby cluster first in a rolling fashion, followed by the primary cluster.

However, you do not commit the changes to either cluster until both have been completely updated with the new binaries. One both clusters have been updated, you commit the changes first on the standby cluster, followed by the primary cluster.

If an older version of the binaries are no longer needed on the machines in the cluster then they can be removed using the db2_deinstall command, running it from the installation path of the previously installed DB2 copy.

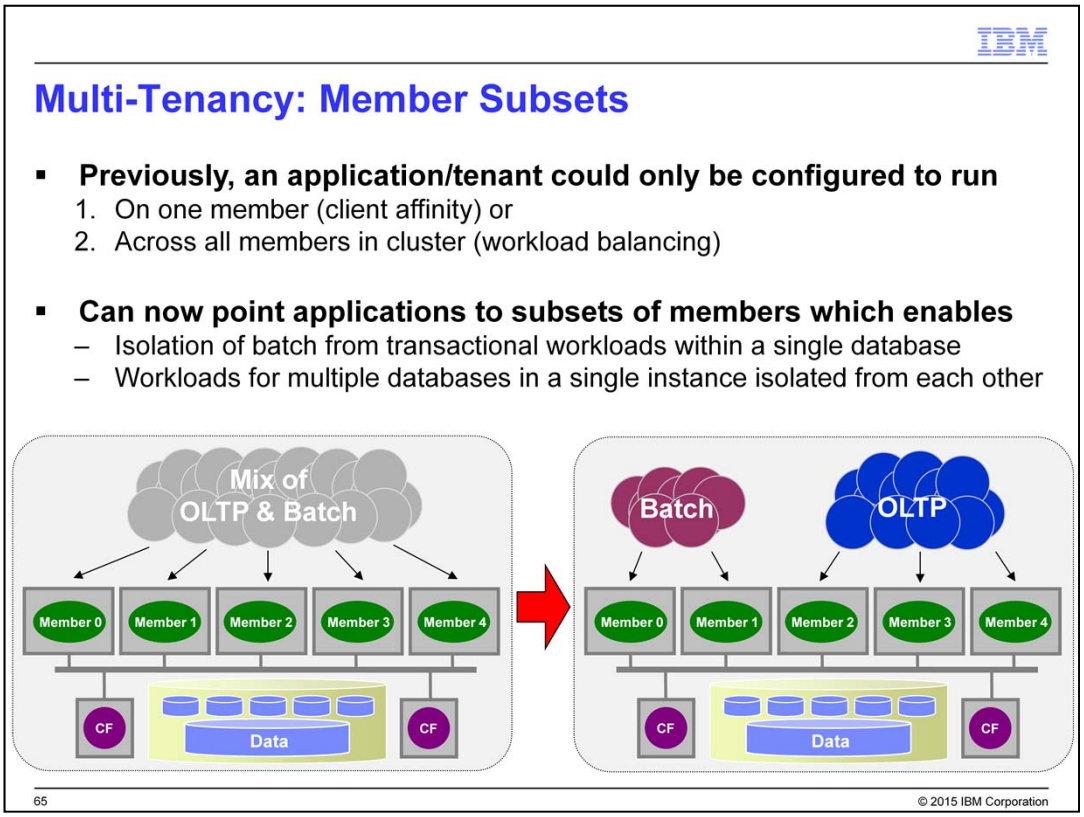

In previous versions of DB2 pureScale, you could choose to have transactions from a particular application running either against exactly one member (using client affinity) or across all of the members in the cluster (using workload balancing). Different applications could be configured with different approaches (i.e. some could use workload balancing and others could use client affinity).

You can now define member subsets such that you can do workload balancing within just a subset of the members in the cluster. There are different reasons that you might want to do this, including:

- You have different types of workloads being executed within a single database – e.g. batch and transaction processing – and you want to keep them running on separate members. In the past, you could direct the batch to a single member but the transaction applications that needed workload balancing could still run on that batch member as well (although member/server load would come into consideration when choosing where to run the transactions on).
- Similarly, some users have wanted to be able to use a specific member for administrative tasks – like backups – and use the remaining members for the real workload.
- Multiple databases are being consolidated into a single DB2 pureScale instance, where the workloads are distinct requiring their own specific database configuration. You could do this before with client affinity but that meant a single member. Now, you can have multiple members designated for a given database and its applications.

Member subsets can be defined as an inclusive subset (the default) or an exclusive subset. If it is an inclusive subset then members which are not included in the member subset are included in the server list for HA purposes only when no member of the subset is online. If the member subset is defined as an exclusive subset then members which are not included in the member subset are excluded from the server list. In this case, if all members of the subset are down then recovery will not be performed for them on any of the remaining healthy members outside of the subset.

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

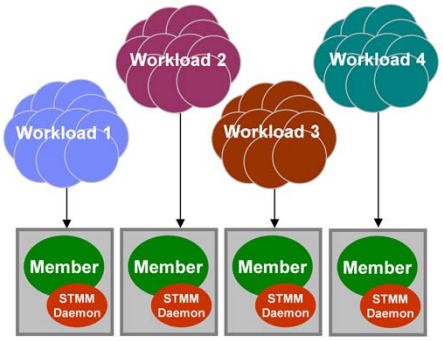

66

Previously, the behavior of self-tuning memory management (STMM) in DB2 pureScale was that a single member was chosen to be the STMM tuning member. This could either be explicitly chosen by the user or it could be assigned randomly by DB2. The STMM tuner would only be aware of the work being done on its own member and when it chose to make tuning changes, these changes would be distributed to the other members and be applied there. To be specific, this only took place for those members that had the configuration parameter SELF_TUNING_MEM set to ON.

Now, DB2 pureScale allows per-member STMM tuning. Each member can have its own tuner and can make and apply decisions locally based on the workload being run on it. This supports workload consolidation and affinitized workloads much better where those workloads are different across the members.

Additionally, the INSTANCE_MEMORY database manager

configuration parameter can be set on a per-member basis. This allows for better support of heterogeneous environments where different member hosts have different amounts of memory.

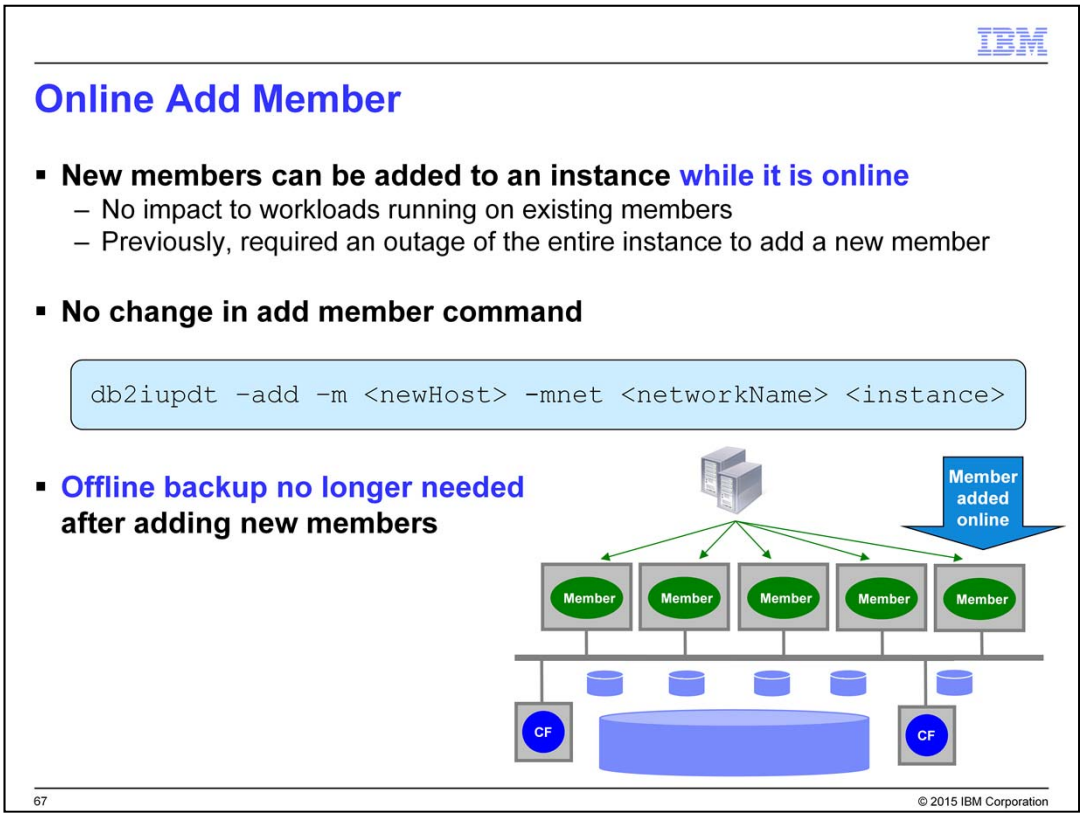

Previously, if you wanted to add one or more members to a DB2 pureScale instance, you had to stop the instance. The resulted in a cluster-wide outage. Now, you can add members in an online fashion, without impacting the work being done across the existing members.

Also, after adding new members, DB2 forced you to take a full offline database backup for every database in the instance prior to starting up those databases again. This could take a significant amount of time depending on the size of the database(s). The reason for this is that the addition of a member means the introduction of a new log stream for that member and DB2 was not able to rollforward through that addition. However, that is no longer the case and you are no longer forced to take the backup now.

db2iupdt is the command that is still used to add members, as was the case in the past.

The drop member operation is still an offline operation, requiring that the cluster be stopped (and an offline backup be taken before starting up a database again). This should be an infrequent type of operation. While some customers are interested in temporarily adding capacity, most are interested in adding capacity permanently. For those that want to add it temporarily, the member in question can always be stopped (and have it's host's resources be used for other purposes).

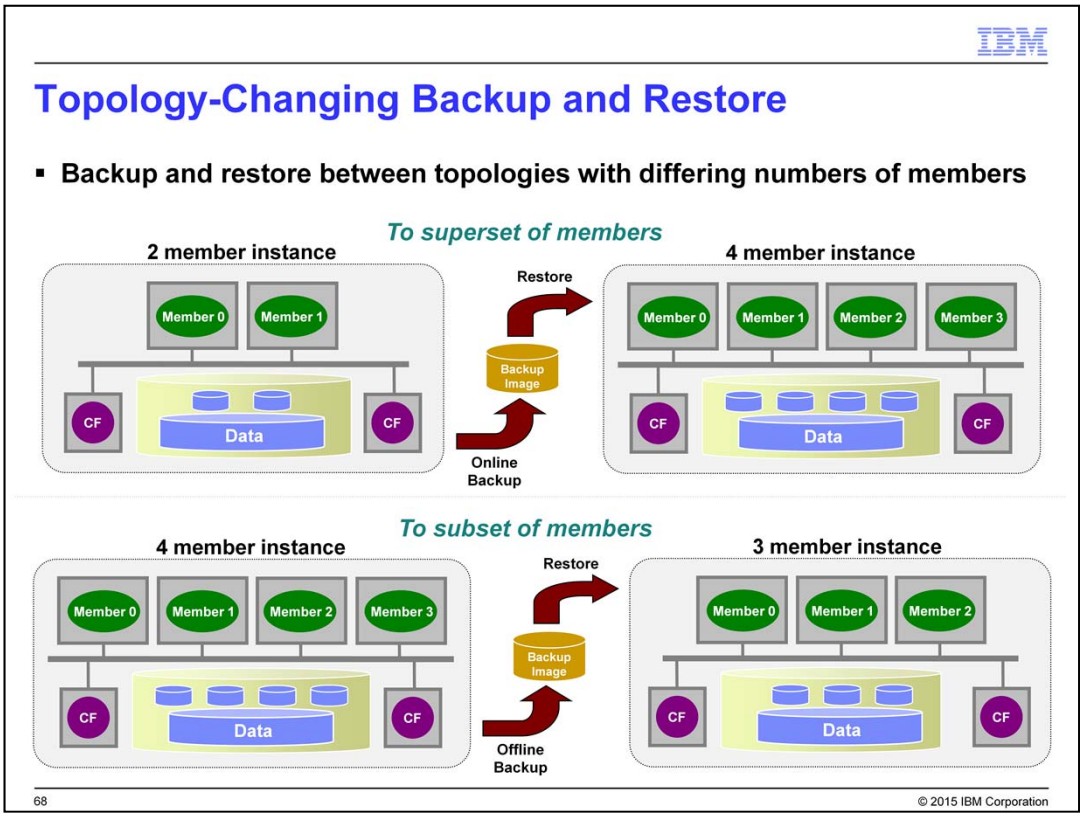

Previously, when doing a restore of a DB2 pureScale database image you had to restore it to an instance that had a matching topology – in terms of the numbers of members and the member numbers (for instance, the backup of database from instance with members 0, 1, 2, and 3 restored to a DB2 pureScale instance with members 0, 1, 2, and 3). Otherwise, the restore would not be allowed.

In DB2 10.5, you can now restore a pureScale database backup to a different number of members. Also, you can restore a non-DB2 pureScale backup image to a DB2 pureScale instance (and vice-versa). The next slide covers the latter.

There are some restrictions that apply in terms of the types of backups that are allowed and whether rollforward is allowed after the restore. If you are restoring a backup image that was generated with a topology that is a subset of the target instance topology (for instance: (0, 1, 2) to (0, 1, 2, 3) or (1, 2, 4) to (1, 2, 3, 4, 5)) then the source backup image can either be from an offline or online backup. Also, you can restore either at the database level or the table space level. With this scenario you can also perform a rollforward afterwards. This is what is shown in the first example here ("To superset of members").

If you are restoring to a subset of the members, as in the second example here ("To subset of members"), the source backup image must be from an offline backup and you are not allowed to do a rollforward following the restore. Also, before the database can be activated and used, a full offline database backup must be taken. This is done to provide the user with a new recovery starting point if they ever need to restore the database again and will likely need to perform a rollforward. The restore performed in this subset example can only be done at the database level. Table space level restore requires a roll forward to be done after it, which is not allowed here.

Having to take a full offline database backup could be an issue for some users, especially if they have a very large database. If the user is willing to skip this – at the risk of having unrecoverable data in the case of a serious failure prior to their next scheduled backup – then it is possible to use the db2dart command to turn off the "backup pending" state for the database. The full command is: db2dart <dbName> /CHST /WHAT DBBP OFF

All of this applies to snapshot backups as well (keeping in mind that snapshot backups are typically online, but can be taken offline).

On a semi-related note, you can restore a DB2 9.8 pureScale database into DB2 10.1 or DB2 10.5 (and have the database automatically upgraded as part of the process).

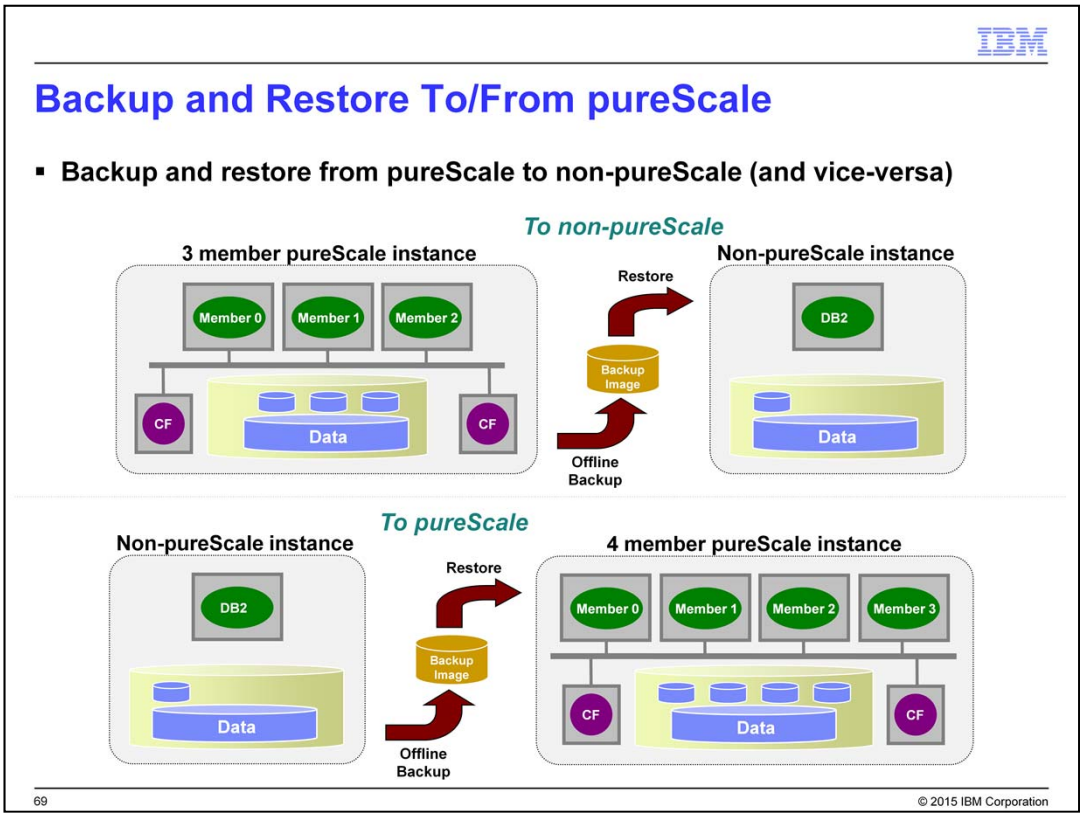

You can also restore a non-DB2 pureScale backup image to a DB2 pureScale instance (and vice-versa). In the case of restoring from non-DB2 pureScale to DB2 pureScale, per the DB2 pureScale prerequisites the database must be using automatic storage for all of the table spaces (restore will fail otherwise). The target DB2 pureScale storage must be on GPFS but it does not matter what kind of file system was being used on the original non-pureScale source system.

This top example on this slide shows a situation where we are moving a database from a pureScale instance to a non-pureScale instance. The bottom example shows the reverse of this.

In both of these cases, the source backup image must be from an offline backup and you are not allowed to do a rollforward following the restore. Also, before the database can be activated and used, a full offline database backup must be taken. This is done to provide the user with a new recovery starting point if they ever need to restore the database again and will likely need to perform a rollforward.

Having to take a full offline database backup could be an issue for some users, especially if they have a very large database. If the user is willing to skip this – at the risk of having unrecoverable data in the case of a serious failure prior to their next scheduled backup – then it is possible to use the db2dart command to turn off the "backup pending" state for the database. The full command is: db2dart <dbName> /CHST /WHAT DBBP OFF

All of this applies to snapshot backups as well (keeping in mind that snapshot backups are typically online, but can be taken offline).

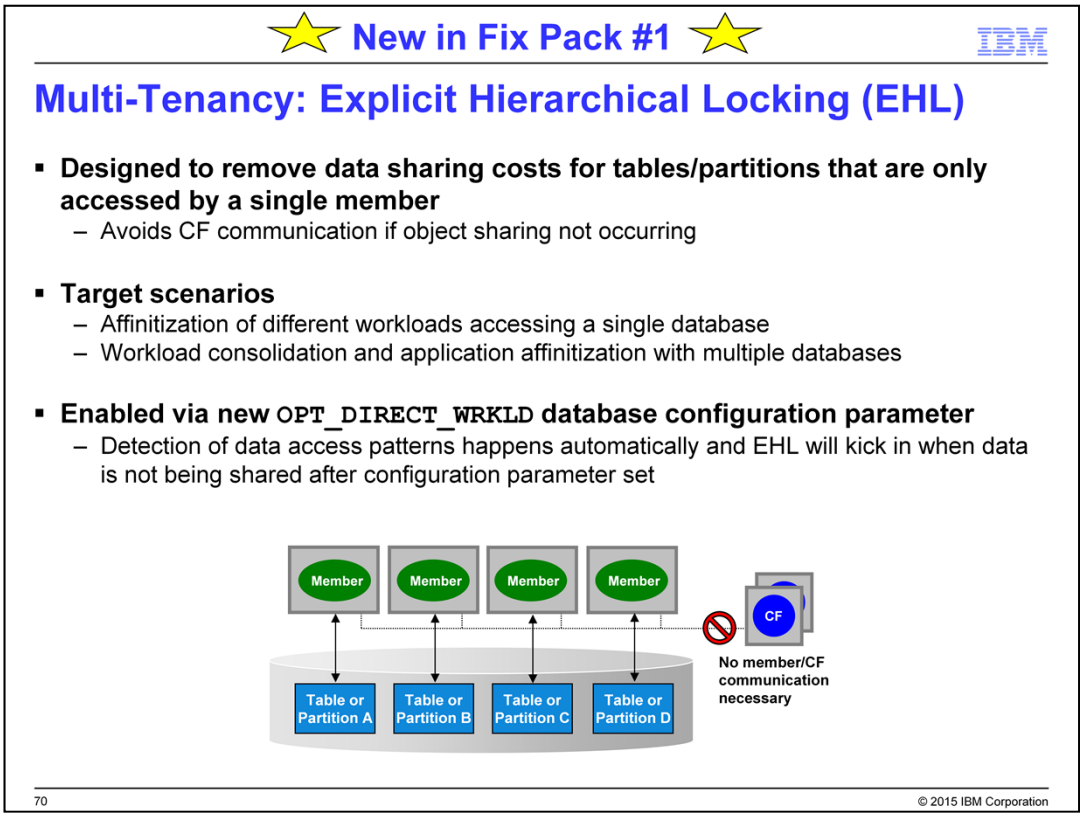

Explicit Hierarchical Locking (EHL) is a strategy that is employed to remove the data sharing costs associated with transactions running in DB2 pureScale if the table (or table partition) being accessed by a member is not being accessed by other members. When dealing with these "not shared" (or "directed access") tables, the member does not need to do the typical communication it does with the CF, such as asking the CF for row locks and flushing pages to the CF at commit time.

We won't get into low level details, but the name and strategy come from our ability to exploit the implicit locking hierarchy that exists between table locks, row locks, and page locks. If we hold a table lock at the member level (regardless of what individual transactions are doing within the member) then there's no need to communicate with the CF to get row locks. Of course, that means that concurrency is reduced across the members, but we're talking about situations where we believe there should be no other members wanting concurrent access to that particular piece of data. And if there are then we fall out of this mode and go back to regular data sharing mode for the table or partition in question.

Explicit Hierarchical Locking is critical for several workload types which can take advantage of this optimization. In particular the following main uses cases exist:

- One member configurations, or batch window workloads which typically only occur using a single member.
- Grid deployments, where each application is affinitized to a single member and where most of its data access is data associated only with that particular application.

This behavior can be enabled or disabled at the database level through the new OPT_DIRECT_WRKLD database configuration parameter, which can be set to YES or NO.

The default for new databases is YES. The parameter can be updated dynamically and it cannot be configured for non-DB2 pureScale instances. The parameter is required for rare cases where EHL might be hurting the performance for a workload, or if customers are not willing to live with the HA (longer recovery time) repercussions (which will be described on a subsequent slide).

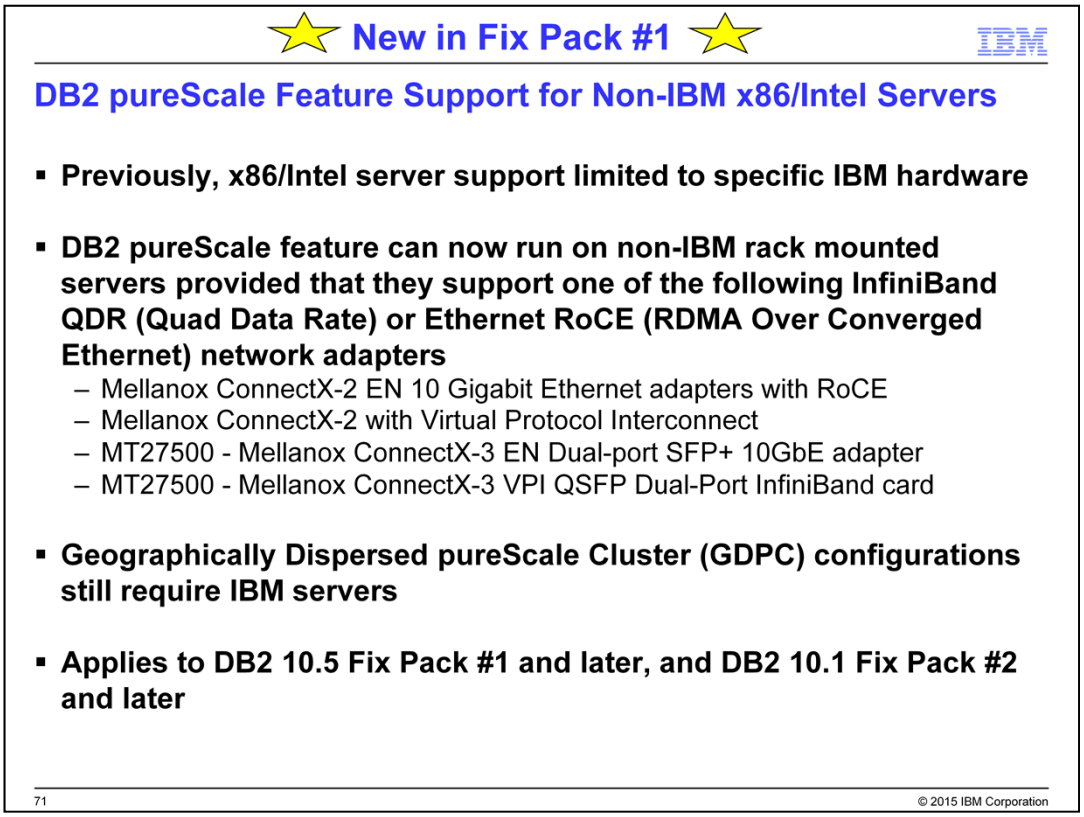

Prior to DB2 10.5 Fix Pack 1, the DB2 pureScale feature was only supported on a limited number of IBM System x x86/intel-based servers. Starting with DB2 10.5 Fix Pack 1, the DB2 pureScale feature can now be deployed on non-IBM rack mounted servers (including Dell and HP hardware). This also applies to DB2 10.1 Fix Pack 2 and later.

The use of a high speed interconnect (10 GbE RoCE or QDR InfiniBand) is still required and there are specific network adapter cards that are supported (see the ones listed on the slide). Therefore, for a server to support the DB2 pureScale feature, it must support one of these adapters.

Given the widely varying nature of such systems, IBM cannot practically guarantee to have tested on all possible systems or variations of systems. In the event of problem reports for which IBM deems reproduction necessary, IBM reserves the right to attempt problem reproduction on a system that may not match the system on which the problem was reported.

This new support for non-IBM hardware does not apply to Geographically Dispersed pureScale Cluster configurations (GDPC). GDPC configurations are only supported with IBM servers (those that were previously explicitly listed in the documentation prior to DB2 10.5 Fix Pack #1).

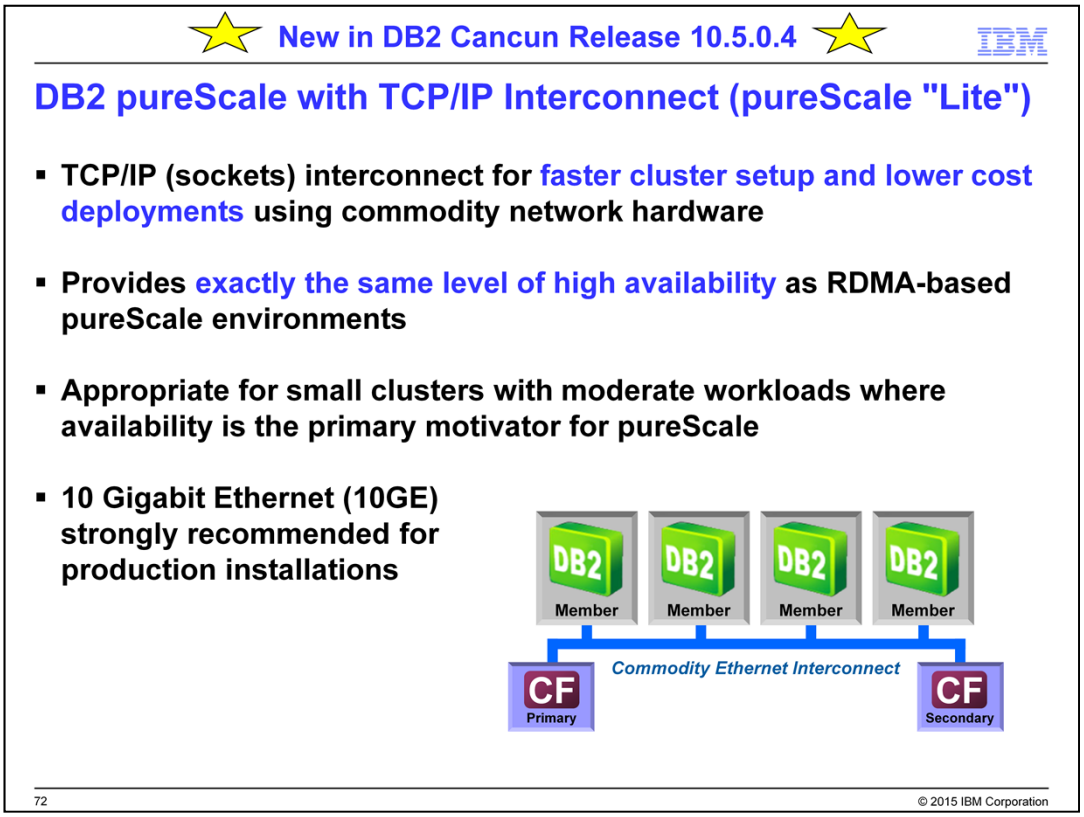

Previously, a DB2 pureScale environment supported two network configurations for low latency, high-speed interconnect communication between DB2 members and cluster caching facilities (CF): remote direct memory access (RDMA) protocol over an InfiniBand (IB) network, and, RDMA protocol over Converged Ethernet (RoCE) network (i.e. RDMA over 10 Gigabit Ethernet).

Starting in DB2 Cancun Release, you can run pureScale on a TCP/IP network, without requiring special RDMA capable adapters. This is sometimes referred to as running in "sockets" mode.

Running your DB2 pureScale environment on a TCP/IP network has many advantages including faster cluster setup and the ability to test the DB2 pureScale technology without special hardware requirements. This configuration provides very good performance, and is aimed at workloads with modest data sharing needs. For the most demanding write-intensive data sharing workloads, an RDMA-based network offers the best performance and is still the recommended interconnect choice. The words "modest" and "moderate" have been used here to describe appropriate workloads, which are rather subjective. As we get more experience with customer workloads in this environment, we'll be able to provide better guidance around this.

No additional hardware, firmware, or software is required to install pureScale with the TCP/IP sockets interconnect option. The only requirement is to have a network that is accessible by all the hosts. All hosts in the instance must use the same type of interconnect, must be able to access each other, and must be on the same subnet. You set up your TCP/IP network as you normally would, set up all hosts on the same subnet, and test the host name resolution and connectivity.

It is a best practice that your TCP/IP network is 10 Gigabit Ethernet (10 GE) or higher. However, if your workload has only modest network usage requirements (e.g. for a development or QA system), you can use 1 Gigabit Ethernet (1 GE). By default, DB2 will block the use of a network slower than 10 GE unless the DB2_SD_ALLOW_SLOW_NETWORK registry variable is enabled. Typical performance impact of 10GE TCP/IP is 30% relative to RDMA-based interconnect (based on initial internal testing – your mileage may vary).

A new database manager configuration parameter has been added to reflect the transport method being used by the cluster. It will be set to TCP or RDMA as appropriate when the instance is created, but it can be changed

afterwards if the interconnect type needs to change (either from TCP to RDMA, or RDMA to TCP.

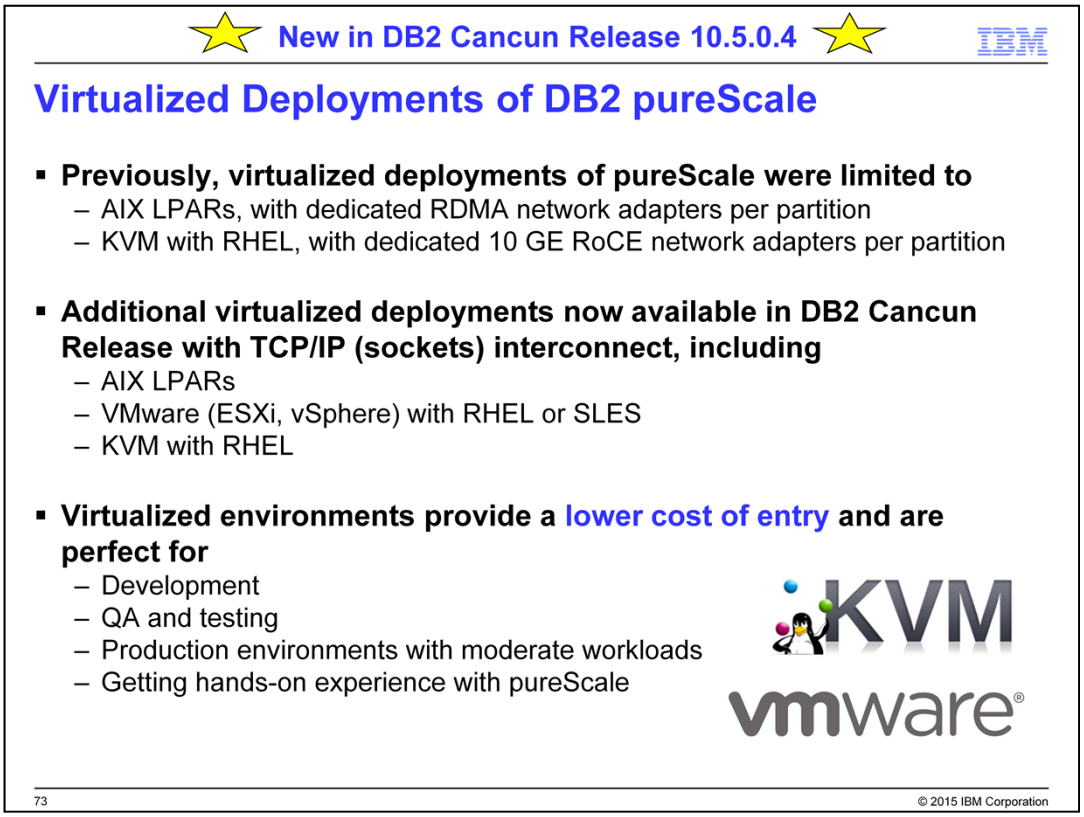

Previously, virtualization for pureScale environments was limited to logical partitions (LPARs) on the AIX operating system on Power hardware and to KVM (Kernel Virtual Machine) with Red Hat Linux on Intel x86 hardware. However, in both cases there was still a requirement to have dedicated RDMA network interconnect adapters per partition. For AIX this was InfiniBand or 10GE RoCE. For KVM with RHEL this was 10GE RoCE.

Starting with DB2 Cancun Release, a pureScale environment can be run in a set of virtual machines using sockets (TCP/IP) as the underlying interconnect. This includes AIX LPARs on Power and VMware or KVM for Linux operating systems (RHEL or SLES). For VMware, specifically this means the use of VMware ESXi or vSphere. See the Information Center for details on specific versions/levels.

KVM = Kernel Virtual Machine. KVM is a virtualization infrastructure that's part of the Linux kernel.

Additional GDPC Configurations

- **A "stretch" or geographically-dispersed pureScale cluster (GDPC) spans two sites at distances of tens of kilometers**
  - Provides active/active access to one or more shared databases across the cluster
  - Enables a level of DR support suitable for many types of disasters (e.g. fire, localized power outage)

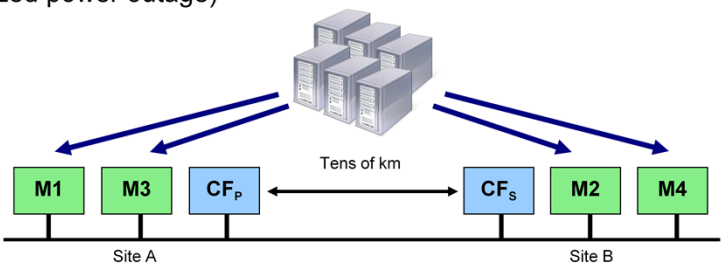

- **GDPC now supports the following configurations**
  - AIX with InfiniBand (IB) network
  - AIX with 10 Gigabit Ethernet (10GE) RoCE network (new)
  - RHEL with 10 Gigabit Ethernet (10GE) RoCE network
  - SUSE with 10 Gigabit Ethernet (10GE) RoCE network (new)

Previously, Geographically-Dispersed pureScale Clusters (GDPC) were limited to AIX with InfiniBand interconnect and RHEL Linux with a 10GE RoCE interconnect. Starting with DB2 Cancun Release, additional configurations are now available. This includes AIX with 10GE RoCE and SUSE with 10GE RoCE.

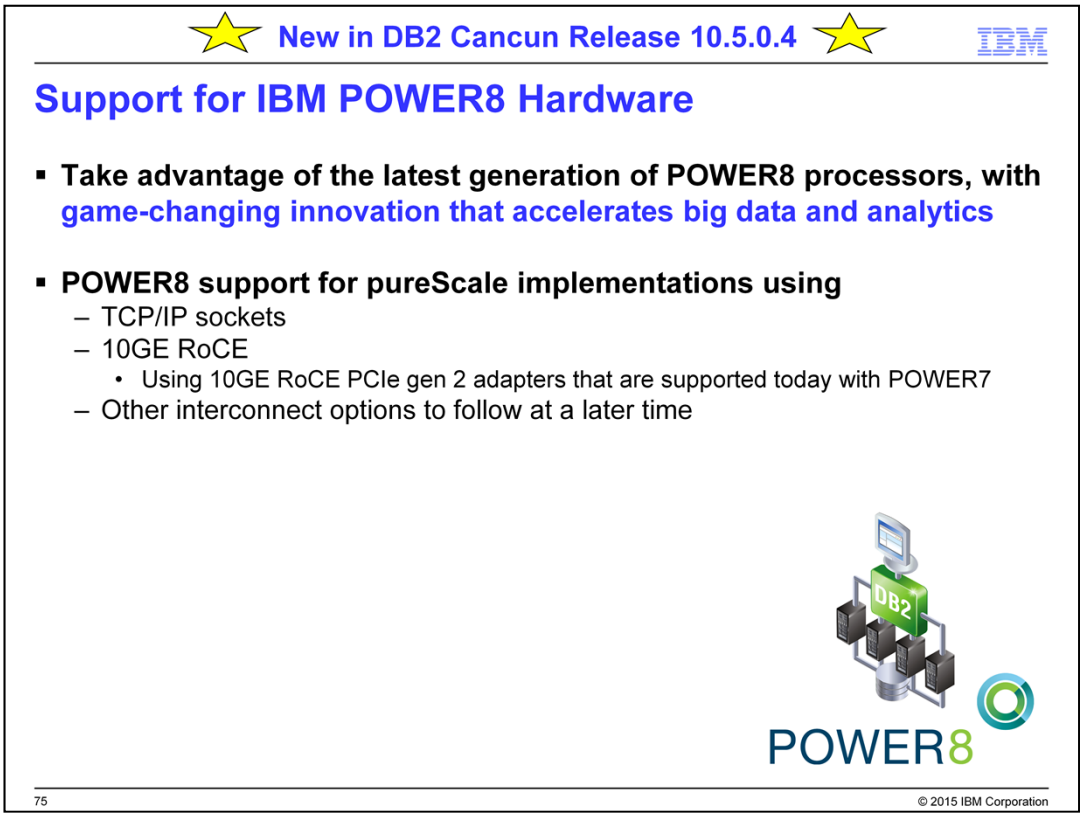

In DB2 Cancun Release, POWER8 hardware will be supported (with AIX-based implementations; DB2 pureScale does not support pLinux).

Specifically, support will start with 10GE RoCE using the same cards that are supported for pureScale today with POWER7. Also supported will be the non-RDMA TCP/IP configuration previously discussed.

DB2 Cancun Release will not initially support the InfiniBand interconnect. This should come post-DB2 Cancun Release.

In-place (Online) Table Reorganization

- **Online table reorganization now fully supported in pureScale**
  - Reclaim free space
  - Eliminate overflows
  - Re-establish clustering
- **Example**

```
REORG TABLE <tableName> INPLACE ALLOW READ ACCESS
```

In-place table reorganization, also referred to as "online table reorg" is supported for pureScale in DB2 Cancun Release.

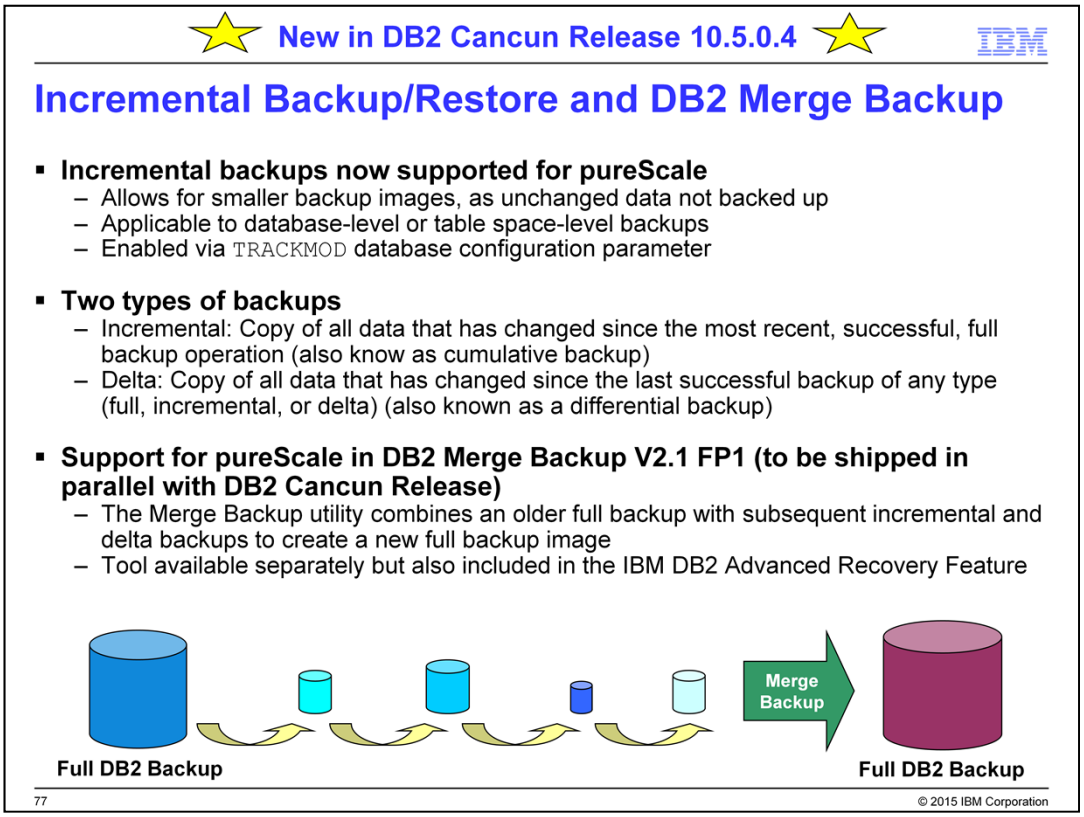

Incremental backup and restore operations are permitted starting in DB2 Cancun Release. Prior to this, only full database or full table space level backups could be performed. An incremental backup is a backup image that contains only pages that have been updated since the previous backup was taken (see the different incremental types shown on the slide for specifics). In addition to the updated data pages, each incremental backup image also contains all of the initial database metadata (such as database configuration, table space definitions, database history, and so on) that is normally stored in full backup images. Incremental backups can be performed at both the database level and at the table space level. Two types of incremental backups are supported, as described on the slide.

With the ability to now do incremental backups in pureScale, the DB2 Merge Backup (MBK) tool has also been enhanced to function in a pureScale environment. This tool allows you to merge an older full backup image with subsequent incremental and delta backup images into a new full backup image. This "flattening" of the images into one means less disk space is required to hold the backup image and it also means less recovery objects to manage. The DB2 Merge Backup tool is available individually and it is also included in the IBM DB2 Advanced Recovery Feature. It can be used with any edition of DB2 except for DB2 Express-C (however, pureScale itself is only available through the advanced editions and the developer edition). Specifically, the pureScale support starts with DB2 Merge Backup v2.1 FP1, which was released in parallel with DB2 Cancun Release.

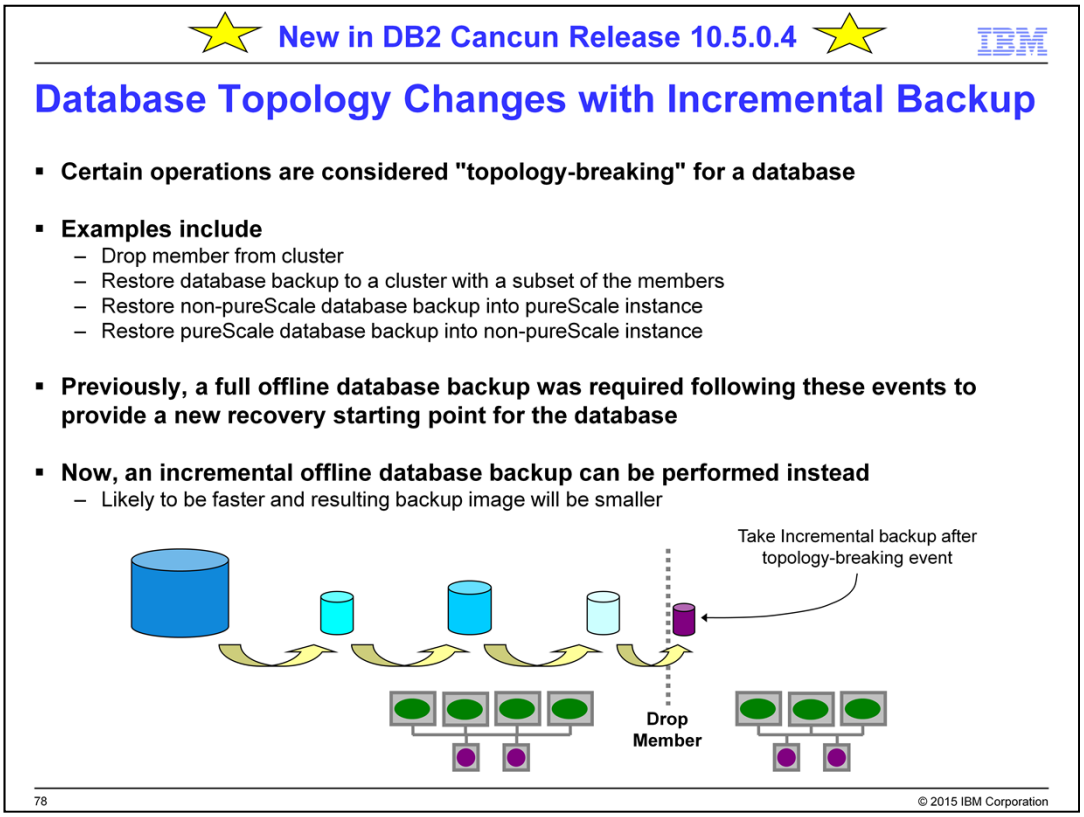

DB2 10.5 GA introduced the ability to add new members to a pureScale instance and database while keeping the cluster online. However, dropping a member required that the instance (and database(s)) be offline. This is considered a "topology-breaking" event and topology-breaking events prior to DB2 Cancun Release required that a new full offline database backup be taken. Having a full offline database backup provides a new recovery starting point for the database such that a restore and rollforward can be performed in the future if ever required. There are other topology-breaking events that had this backup requirement as well. For instance, restoring a backup to a cluster that is not the same member topology, or is not a superset of the original members associated with the database. Also, restoring a non-pureScale backup into a pureScale instance and vice-versa. Adding new members to a cluster or restoring a backup image to a superset of the members are *not* considered topology-breaking events and so they don't have this backup requirement.

In DB2 Cancun Release, if incremental backup is enabled, an incremental offline database backup (incremental or delta) can be used instead of a full offline database backup. This can speed up the backup time and the resulting backup image should be much smaller (depending on the type of backup and how many changes have been made since the last one).

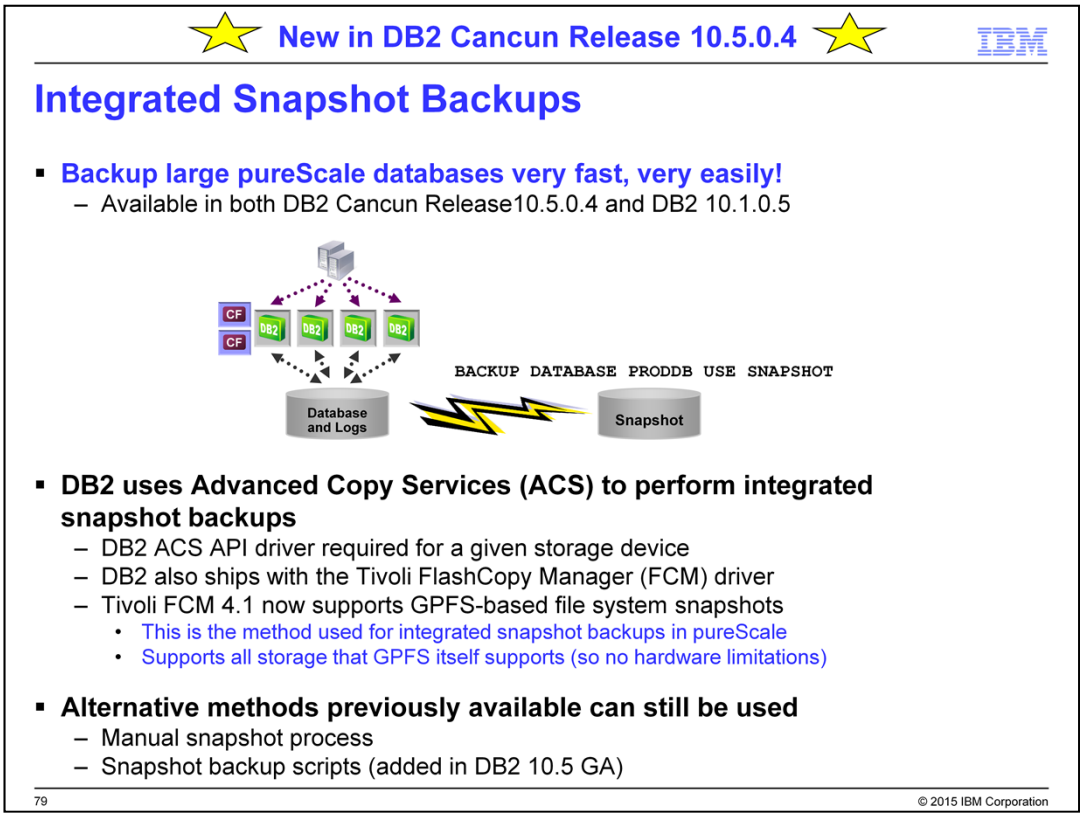

DB2 Cancun Release 10.5.0.4 (as well as DB2 10.1 FP5) introduces the ability to use integrated snapshot backups (e.g. BACKUP DATABASE <dbName> … USE SNAPSHOT) in pureScale. Previously, you could not use this functionality. However, you could perform all of the necessary the steps manually (using SET WRITE SUSPEND, storage-specific snapshot commands, and SET WRITE RESUME) or you could use the snapshot script feature which was added in DB2 10.5 GA (e.g. BACKUP DATABASE <dbName> … USE SNAPSHOT SCRIPT '<scriptName>'). Integrated snapshots, though, utilize DB2 Advanced Copy Services (ACS) which has an API that allows for libraries to be written for specific storage hardware. With the snapshot script support in DB2 10.5, scripts can be written that implement the ACS API, which means that you can use integrated snapshots, even for storage devices that do not provide a vendor library.

The terms snapshot and flashcopy are often used interchangeably. In the context of Tivoli FlashCopy Manager specifically (described in more detail on the next slide), the term *FlashCopy* is used for IBM System Storage DS8000, IBM System Storage SAN Volume Controller, and IBM Storwize V7000 storage devices. A FlashCopy creates a point-in-time copy in which the target volume represents an exact copy of the data on a source volume at the time the FlashCopy was started. For IBM XIV Storage System, IBM System Storage N series, NetApp, and file systems such as GPFS, the term *snapshot* is used. A snapshot represents a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The source volumes and snapshots are located on the same storage system. Similarly, a file system snapshot represents a point-in-time copy of a file system or file set within a file system. The space that is required for the snapshot is allocated automatically within the same storage system or file system and can increase over time.

DB2 Advanced Copy Services (ACS) (shipped with DB2) enables you to use the fast copying technology of a storage device to perform the data copying part of backup and restore operations. A backup operation that uses DB2 ACS is called a snapshot backup. To perform snapshot backup and restore operations, you need a DB2 ACS API driver for your storage device. DB2 ships with Tivoli FlashCopy Manager (FCM) which supports various storage vendors. For a list of supported storage hardware for this integrated driver, refer to the Tivoli documentation here: http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/topic/com.ibm.itsm.fcm.unx.doc/c_fcmu_ovr_overview.html. Also see the following documentation for more information: http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/topic/com.ibm.itsm.fcm.unx.doc/t_protect_fcmul.html.

In DB2 Cancun Release 10.5.0.4 (as well as DB2 10.1 FP5), Tivoli FlashCopy Manager and the GPFS version that gets shipped with DB2 fully supports integrated snapshot backups in pureScale environments. This means that you can perform integrated snapshot backups like you can with non-pureScale configurations. The one difference is that the snapshots are GPFS file system-based for pureScale (or more specifically, they can be GPFS file system or file set-based). Because they are based on GPFS snapshot technology, there is no dependency on the underlying storage hardware. As long as GPFS supports the hardware (and we're going to assume that this is the case for somebody using GPFS with pureScale), this technology will work with it.

A requirement for using GPFS snapshots with a pureScale database is that the transaction log path must be in a different file set or different file system than where the database data (database path and table space data) reside.

When a snapshot is performed, data isn't actually being copied for the purposes of creating a completely separate physical copy of it. Instead, pointers to the data blocks that make up the files in the file system are physically copied (to a location within the same file system), but not the data itself. Following the snapshot process, you still just have one copy of the actual data. Once a change needs to be made to the database, though, a data block that gets modified (e.g. a row is updated on a data page and that page is written to disk) is copied within the file system such that you now have the old copy – which stays associated with the snapshot image – and the new copy, which is associated with the current version of the database. This way, the snapshot still represents what the database looked like when the snapshot operation was performed. This is referred to as a "copy-on-write" methodology or algorithm. If you have needs to copy that snapshot image off to another location – for the purposes of having it available for DR purposes, then there are ways of doing that. The specific steps involved are beyond the scope of this presentation.

With respect to speeds, backups are relatively fast because of this copy-on-write methodology. Some internal performance work done within the lab showed that a 200GB database could be backed up in about a minute. However, the restore time will take longer, with the actual time depending on the quantity of data changed since the snapshot was taken, since all of the data blocks that were changed will need to be physically copied back into place.

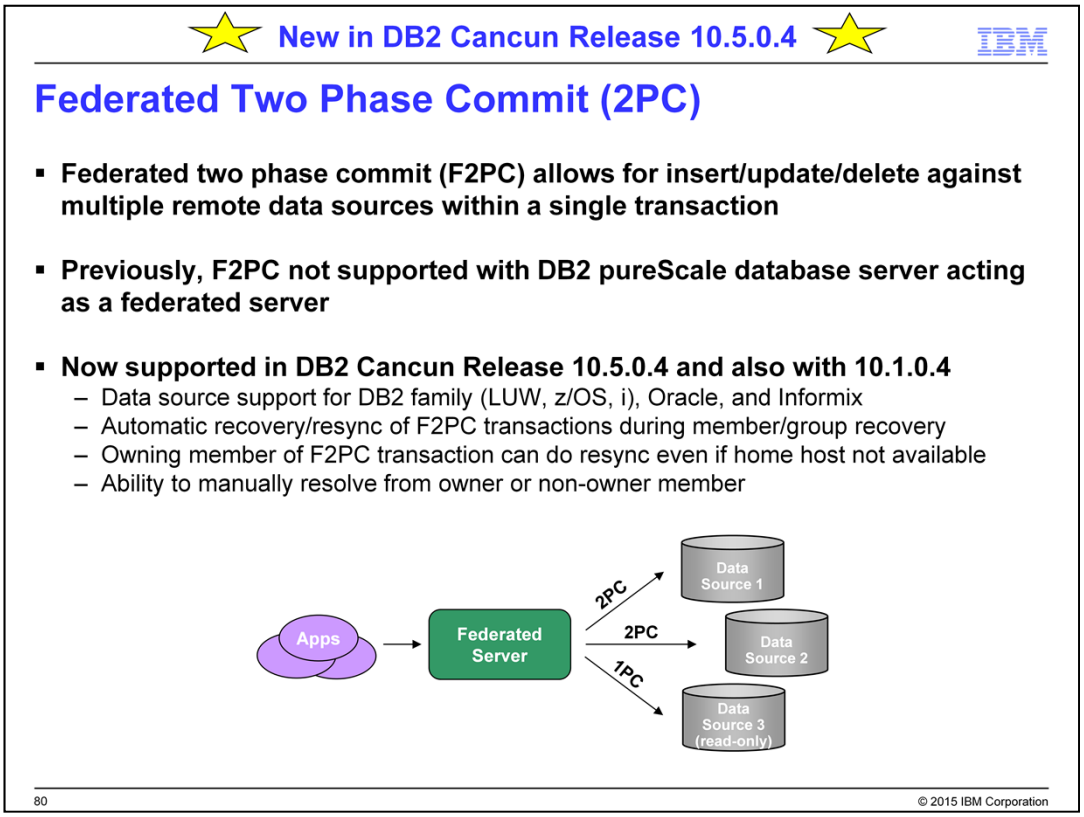

Federated Two Phase Commit (F2PC) is a scenario in which the Federation Server coordinates distributed transactions among multiple remote data sources. In such a scenario, the user executes insert/update/delete statements against multiple nicknames (or federated three-part-names) that belong to different remote data servers. Federation Server acts as a Transaction Manager (TM) and ensures data ACID among the impacted data servers. The F2PC feature was released in DB2 9.1 for single node instances but it has not been permitted in multi-member configurations (DPF and pureScale).

Prior to DB2 Cancun Release, if you tried to define a federated data source in a pureScale environment using the CREATE SERVER statement (or a nickname using the CREATE NICKNAME statement) then you could not define it with the DB2_TWO_PHASE_COMMIT server option (it would fail with an SQL1881N error).

F2PC is supported for pureScale starting with DB2 Cancun Release. In this case, the DB2_TWO_PHASE_COMMIT option must be turned on for all servers involved in a F2PC transaction where updates will be made.

Specifically, the following things can be done in a pureScale environment now:

•Perform distributed updates against multiple (different) data source servers. Data consistency is guaranteed. Data sources can be members of the DB2 family (DB2 for LUW, DB2 for z/OS, DB2 for i), Oracle, or Informix.

•Automatic recovery/resync of F2PC transaction during member/group recovery, if a failure of the member/group occurred in the middle of a F2PC transaction being executed.

•Owning member of a F2PC transaction can do resync even if the member's home host (where it normally runs when the member and host are healthy) is unavailable and the member is residing on another host in restart light mode. Owning member means that member on which the application was connected and the transaction was run.

•Ability to run the LIST INDOUBT TRANSACTIONS command from owner or non-owner member and drive F2PC transaction outcome (commit or rollback) manually. The need for this happens often after a member/group failure that results in a F2PC transaction being in the indoubt state, but has not yet been resync'd (e.g. remote data source may no longer be available).

This capability has also been added to DB2 10.1 FP4, but there are restrictions that do not exist with DB2 Cancun 10.5.0.4. For instance, only DB2 family members (DB2 for LUW, z/OS, i) are supported as a remote data source.

DB2 Spatial Extender

- You can store, manage, and analyze spatial data in a DB2 pureScale environment
- No difference between pureScale and non-pureScale in terms of setting up Spatial Extender and creating a project that uses spatial data

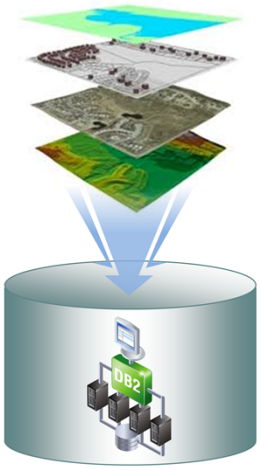

Starting with DB2 Cancun Release, you can now use DB2 Spatial Extender to generate, analyze, store, and manage spatial information about geographic features in DB2 pureScale environments. To start using Spatial Extender, you must set up Spatial Extender and create a project that uses spatial data. The same procedure to do this is followed in pureScale instances as it is done in non-pureScale instances.

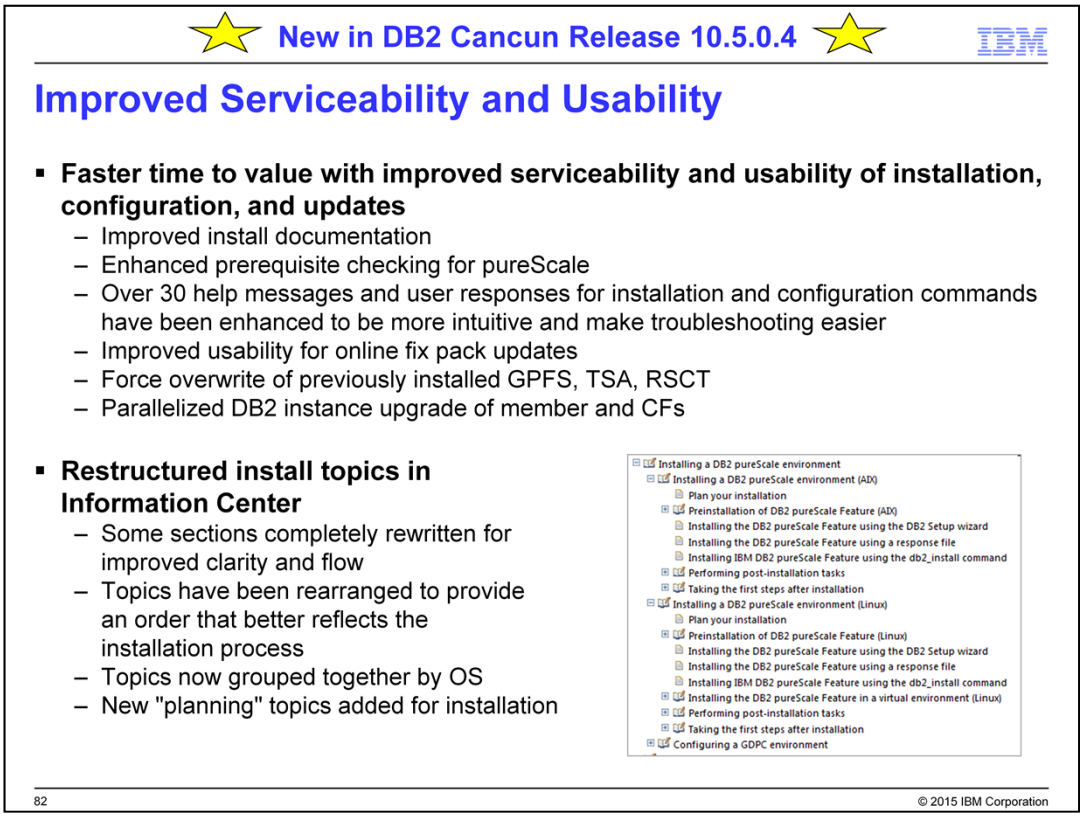

Many improvements have been made to pureScale with respect to improved serviceability and usability in the DB2 Cancun Release (DB2 10.5 FP4), particularly around installation, configuration, and upgrade.

For example, the Information Center has undergone a significant restructuring with respect to the install topic. Topics have been better organized to reflect the actual steps involved in installing pureScale (keeping things specific to AIX grouped together, and things specific to Linux group together), some sections have been completely rewritten, and some new sections have been added.

Another example is that the prerequisite checking for pureScale has been enhanced. The db2prereqcheck tool can be used to validate that the environment is in a state that pureScale requires.

Over 30 help messages and user responses have been enhanced to be more clear and make the troubleshooting process easier should a problem be encountered with the installation or configuration of a pureScale environment.

Various improvements have been made around the usability and robustness of online (rolling) fixpack updates. The process has been enhanced with additional validation and improved error reporting as shown on the slide.

When applying a fix pack in a DB2 pureScale environment, if GPFS, Tivoli SA MP (TSA MP), or RSCT was already manually installed or updated manually (maybe due to a hot fix having to be applied outside of a DB2 fix pack), installFixPack will fail with an error stating that the install cannot proceed. This is done to let the user know that it is not possible to move back to this older level. If the user wants to force the install to overwrite this already installed version then this is possible now using the new "-f" (force) option.

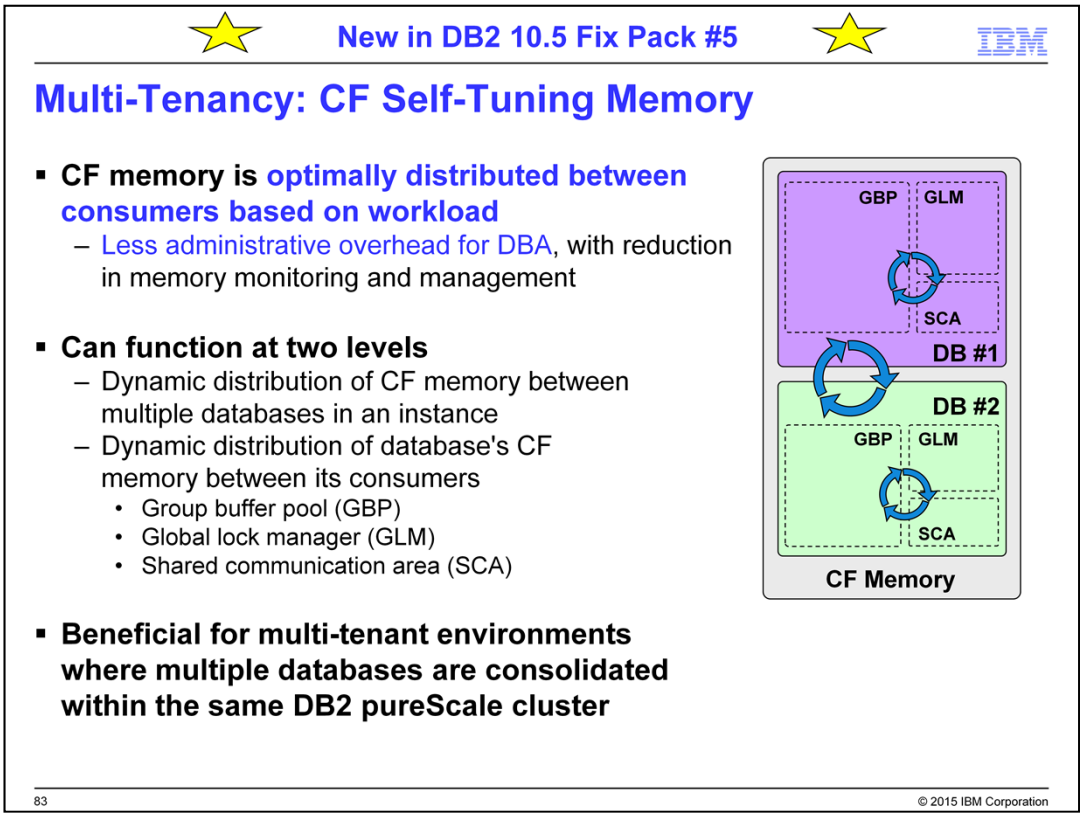

Starting in DB2 10.5 FP5, self-tuning memory management is supported for the cluster caching facility (CF) memory. If enabled, the CF memory will be automatically and continuously reconfigured based on the databases that are active and the workloads being run.

CF self-tuning memory is enabled by setting the registry variable DB2_DATABASE_CF_MEMORY to AUTO. When enabled, CF self-tuning memory avoids out of memory conditions by tuning the database memory parameters: CF_DB_MEM_SZ, CF_GBP_SZ, CF_LOCK_SZ, and CF_SCA_SZ. CF memory will be distributed amongst the consumers that are setup to be automatic. It is perfectly valid to have a database's CF memory be static (e.g. CF_DB_MEM_SZ set to a fixed value) but the individual consumers within that database (GBP, GLM, and SCA) be automatically tuned within that fixed amount of memory.

In a multiple database environment, CF memory is configured automatically based on workload and available memory. Databases that are already active automatically give up CF memory for newly activated databases until a workload-based distribution of CF memory is reached. In addition, when a database is added, there is no downtime to reconfigure CF memory allocation. For each active database, memory will be allocated between the individual memory consumers (GBP, GLM, and SCA) of each database.

The diagram on the slide represents CF memory where there are two active databases – DB #1 and DB #2. CF memory can be dynamically distributed between those two databases, and within each of them, it can be further distributed between the three consumers (GBP, LOCK, SCA).

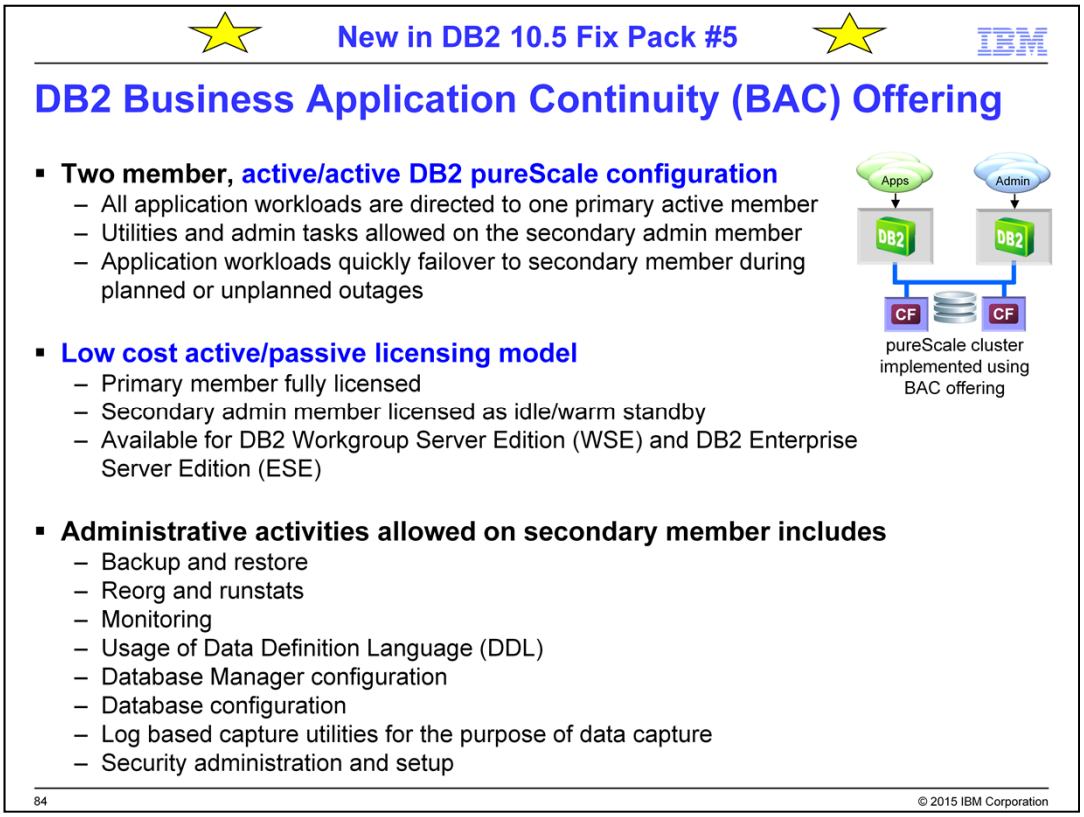

Prior to the introduction of the DB2 Business Application Continuity (BAC) offering, the DB2 pureScale feature was only available as part of DB2 Advanced Workgroup Server Edition (AWSE) and DB2 Advanced Enterprise Server Edition (AESE) (as well as the DB2 Developer Edition). In some cases, this was cost prohibitive to those customers who wanted the high availability that pureScale offers, but didn't really need the scalability of it nor the other features available only in the advanced editions. This is where the new Business Application Continuity (BAC) offering comes in.

With the BAC offering, a customer can setup a two member pureScale cluster where their application workloads are permitted to run on one of the members only. From this perspective, the licensing is considered active/passive because full licensing is only required for the one member. However, we are also allowing administrative work to take place on the secondary member and so this secondary member is still considered active, or hot.

As mentioned, only one of the members in a pureScale cluster using the BAC offering is allowed to run application workloads. However, the other member – the secondary admin member – is allowed to run various DB2 utilities and administrative activities. This slide lists those things that are permitted to run on the secondary member. You can see that it is a fairly extensive list of things, which means that a customer's hardware investment in having this active second node can be fully justified. It isn't a passive node that is just waiting to take over in case of a failure with the primary member. Note that the primary member that is running the application workload can also run these administrative tasks as well, but it will typically make more sense to keep them off the primary member and running on the secondary member instead.

This offering is only available on DB2 10.5 FP5 and later. It is not available for DB2 10.1, nor any DB2 10.5 fix pack prior to FP5.

NOTE: When talking to customers, we have to be very careful about the terminology used and how this solution is represented. A common question that gets asked in RFPs and things like that is "does this solution provide active/active clustering?" and the answer to this is "yes, of course!" Both members are active with useful work being permitted. However, the difference between this solution and a regular pureScale cluster is that there are restrictions around what work is permitted where – it's application work on one member, and administrative work

on the other. Now, the licensing follows an active/passive model because real user application work is only permitted on one member, but the full resources associated with both members of the cluster are there to be used. Hence it is an active/active or hot/hot solution. Therefore, this presentation tries not to use the term active/passive when talking about this offering (unless it's in the context of the licensing). Instead, it typically differentiates the two different types of members with the terms primary active member (or preferred member) vs. the secondary admin member (or just secondary member or admin member). Unfortunately, the actual license terms tends to use the terms passive and standby, but the product management team would prefer this not be the case when talking to people about it.

Other DB2 LUW Enhancements

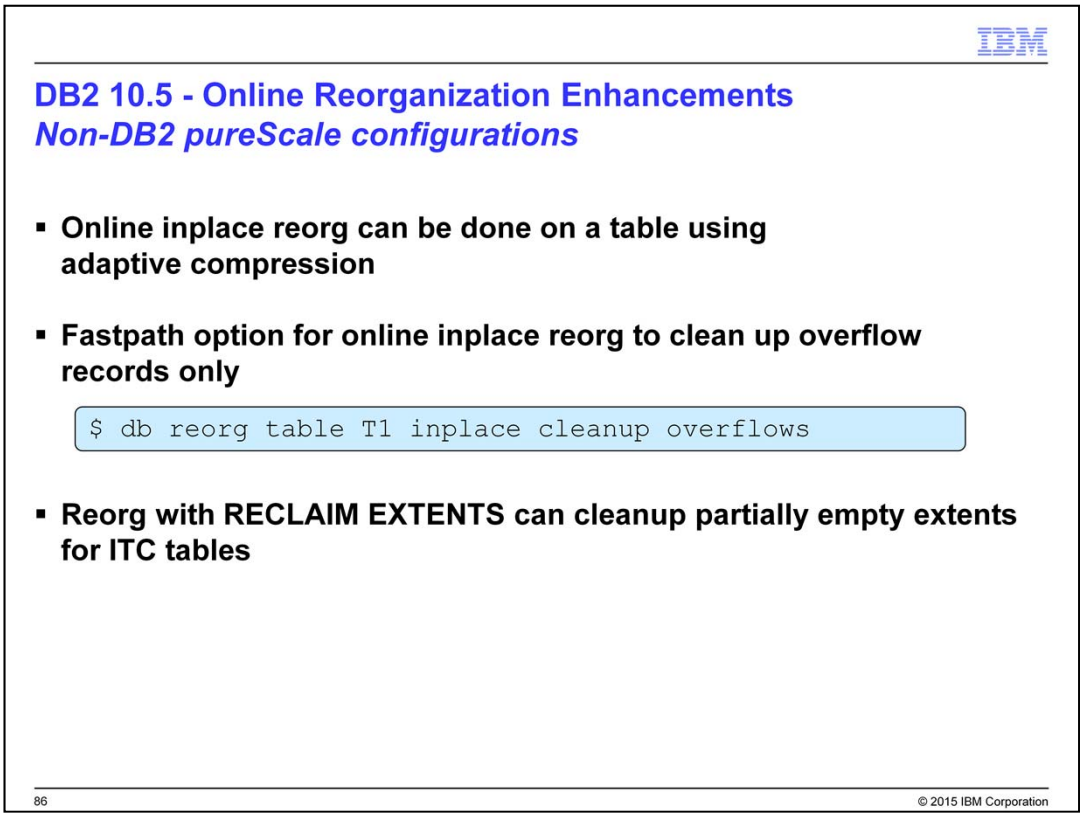

DB2 10.5 adds some additional scope of function for online reorganization.

Note: IBM Internal Use Only : at eGA there will not be Online inplace reorg in pureScale. This will come at a later date.

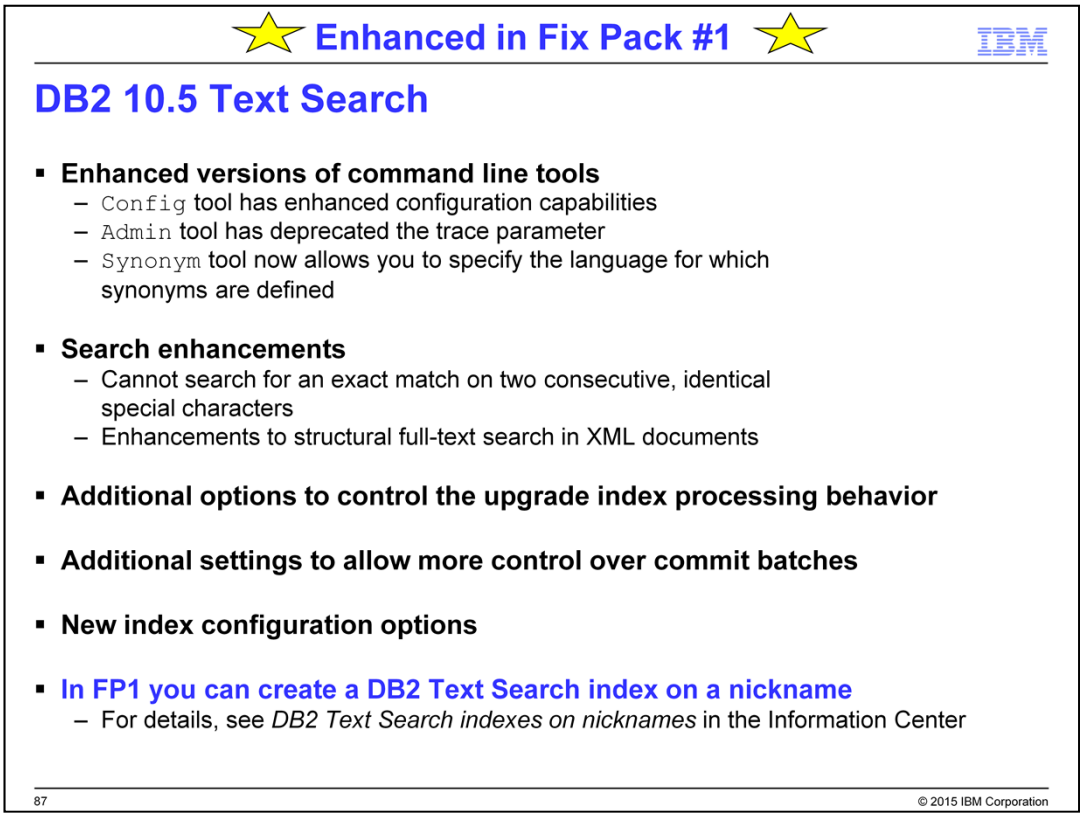

This relates to new features and enhancements in DB2 10.5 Text Search. Full details can be found in this technote: http://www-01.ibm.com/support/docview.wss?uid=swg21643180 Text below with additional links

DB2 Text Search enables an IBM DB2 for Linux, UNIX, and Windows database user to create applications with full text-search capabilities by embedding full text-search clauses in SQL and XQuery statements. Its new generation text search component from IBM.
The functionality of Text Search is enhanced/extended in V10.5.

Command line tools:

DB2 Text Search 10.5 provides enhanced version of command line tools and the details are explained below:

Config Tool
Enhanced configuration capabilities for example, ability to define system level defaults which would apply to all index collections, ability to change all the configuration parameters via the configTool etc are added newly.

The configuration changes for V10.5 set by the user using the configTool will be updated to sqllib/db2tss/config/configuration.xml instead of the old file sqllib/db2tss/config/config.xml.

For more information on each configuration parameter, see the information center.

Admin Tool
The trace parameter is deprecated in this release, use -logLevel instead to enable/disable the trace.

For more information on this tool, see the information center.

Synonym tool
Now, we can also specify the language for which synonyms are defined using -language <2-character language code> parameter.

For more information on this tool, see the information center.

Search enhancements:

Special characters

You cannot search for an exact match on two consecutive, identical special characters. Queries of this type return documents that contain only one of the special characters.

For more information on special characters, see the information center.

Structural full-text search in XML documents

The following search features are supported by XML search:
Boolean operators (basic search)
exact match
fuzzy search
proximity search
stop words
synonyms
wildcard characters

For more information and examples on xml search, see the information center.

Trouble shooting Text Search:

Logging and tracing for the DB2 Text Search server

There are changes in the steps to collect the data required to trouble shoot text search related problems, for more details please refer to the link:

For more information on logging and tracing, see the information center.

Monitoring queues for DB2 Text Search index updates

You can gather more monitoring information now with this release and this can be used to tune the DB2 Text Search server configuration, if you experience indexing performance issues.

For more information on monitoring queues, see the information center.

Control update index processing:

To control the update behavior several options are available now, from specifying a commit-batch to secure completed work, to parameters that determine the number of maximum parallel update processes, maximum collection numbers, etc. With this release, the administrators can control the time-window that an index update may use and thus avoid load on the system during peak time.

The following capabilities are enabled with this release:
Ability to specify the commit cycle in hours to simplify the specification of commit-batches
Provide an option to gracefully end the index update and continue processing the next time the update starts by enabling the use of a maximum time window for updates to prevent that index processing continues to run during peak workload times
Simplify schedule handling by combining create and initial update
Allow control of which data to process in the execution of the initial update
Enable an option to specify a custom logtype to control whether triggers are created and how the auxiliary logtable is defined

Commit batches

Text index update processing already provides a feature to specify the commitsize via the updateautocommit argument. To provide further control, additional settings are available now to determine whether the commitsize should be treated as rows or hours, and how many batches to process.
Each commitcyle comes with a significant overhead, and should only be used for very long-running index updates.

For more information on commit batches, see the information center

New index configuration options

A new option 'INITIALMODE' is introduced in this release, with this option, you can run the initial update immediately when the index is created. Also, you can defer it to the first update operation or skip it altogether.
Another option LOGTYPE' is introduced to set a custom logtype so that the text index administrator can decide whether triggers should be added to populate the primary log table.

For more information on these new index configuration options, see the information center

Set command locks:

It is possible to set the command lock manually, for example, to restrict access when an administrative operation is applied on the collection level.
For more information on command locks, see the information center.

Some of the other capabilities include the following:

Enhanced error handling, specifically query syntax errors

Enhanced indexing and optimization changes which would speed up indexing and optimization, and reduce impact of indexing/optimization on search.
Stronger support for multilingual collections (provide more accurate results across languages)
Support for embedded documents, archive files, such as ZIP, RAR and TAR files by extracting, concatenating, and indexing their content.

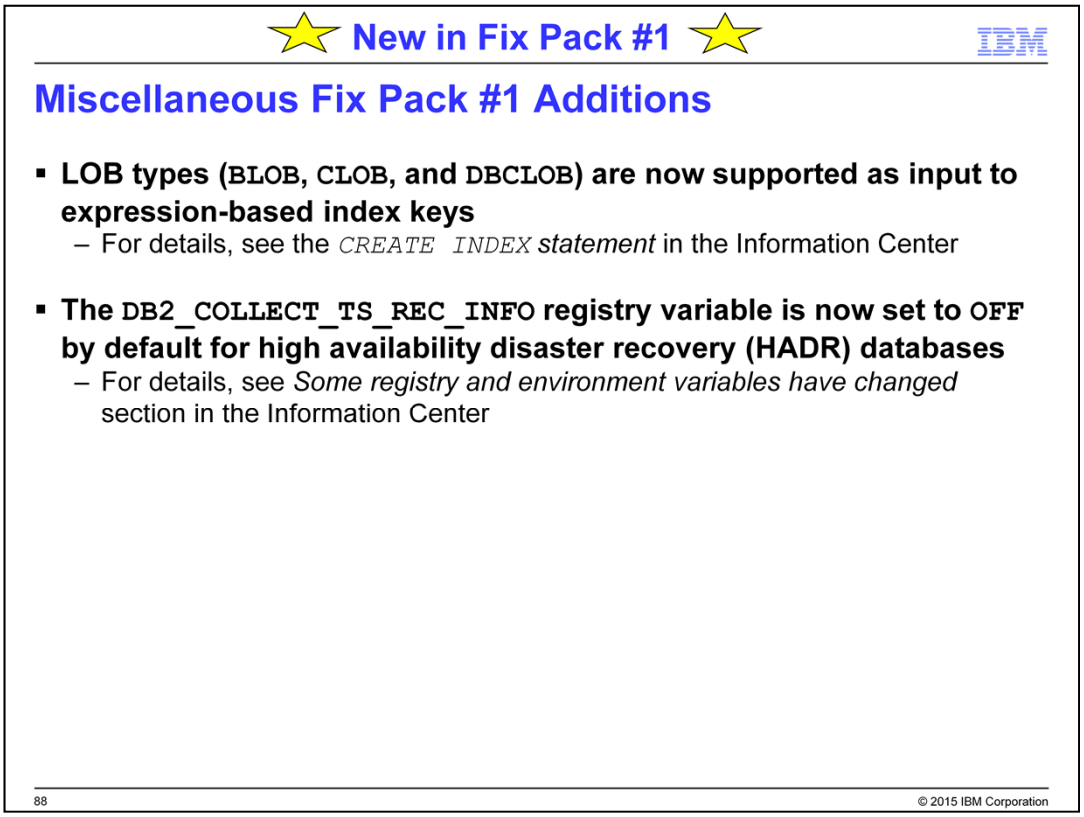

A number of other changes have occurred in DB2 10.5 LUW FP1. This slide lists them at a high level.

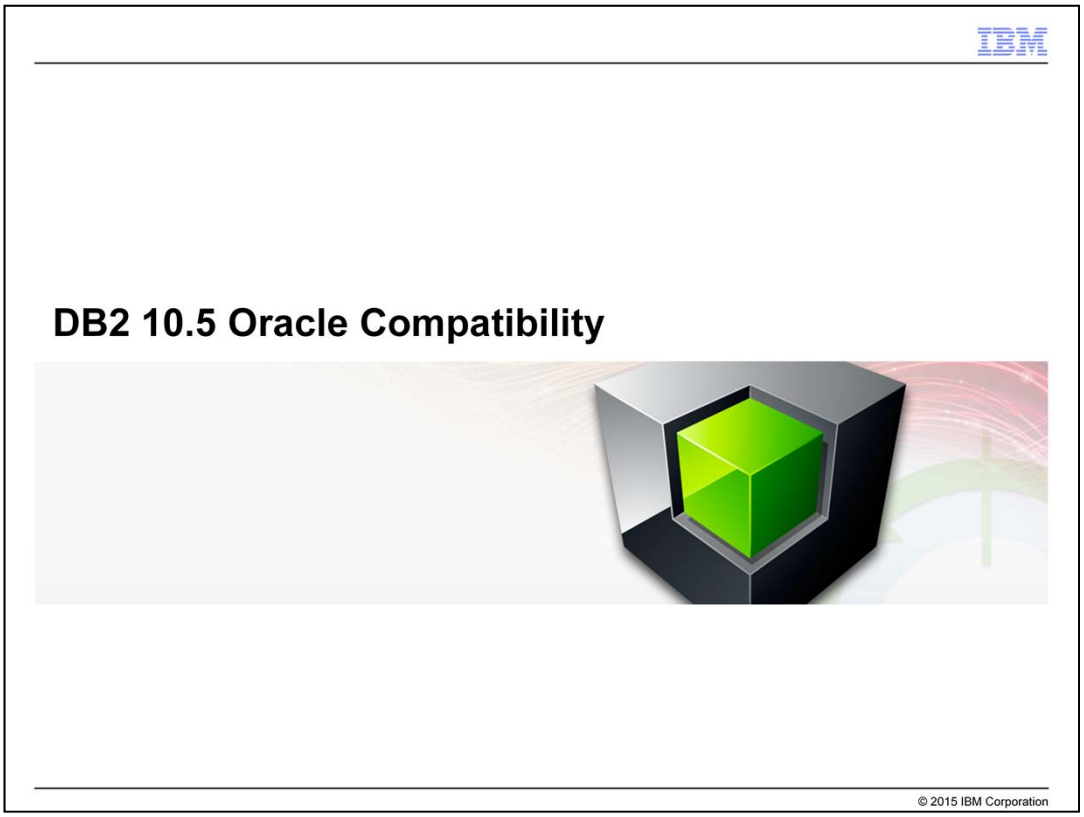

DB2 has included Oracle compatibility features for many releases. DB2 10.5 continues to add significant compatibility features that make it easier to port Oracle applications to DB2.

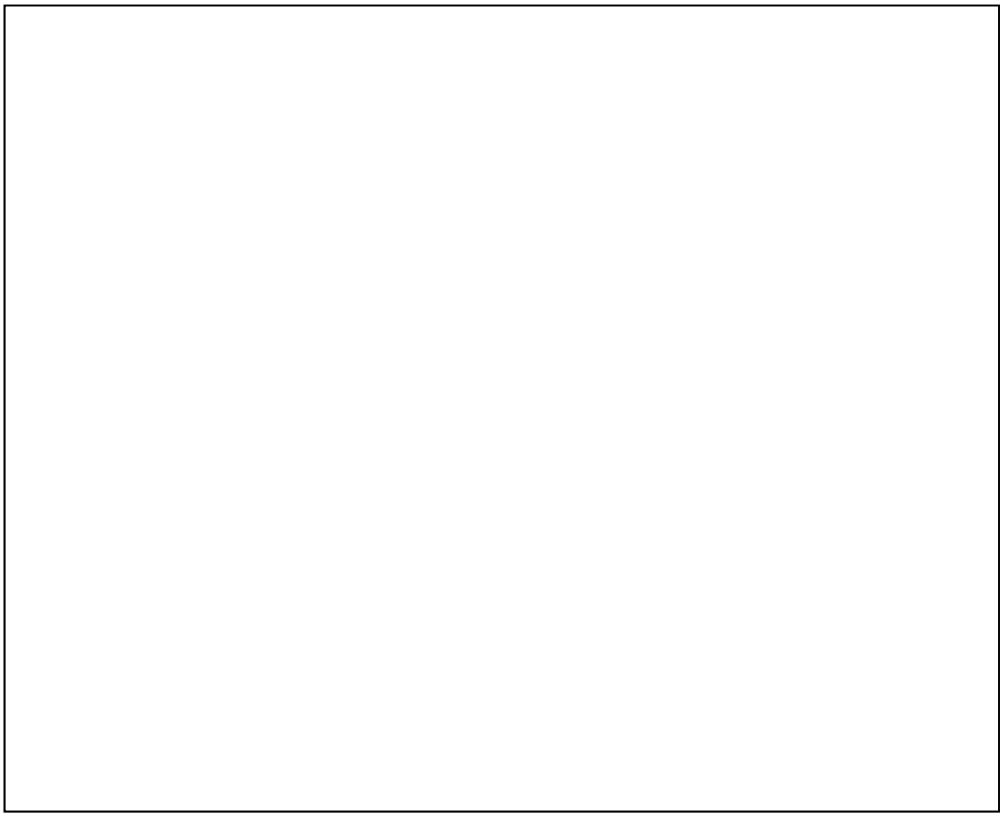

DB2 Oracle compatibility goes deeper than just supporting SQL syntax. This chart lists all of the features that DB2 has implemented to make the transition from Oracle to DB2 as easy as possible. Compatibility also means "native" implementation. There is no need to convert, translate, or change most SQL since the DB2 engine will understand the Oracle syntax natively.

This is a list of high level features and “Oracle buzzwords”.
OCI = Oracle's CLI
Oracle Forms “4GL” order entry language
SQL*PLus = CLP
built-in package library = All sorts of module extensions such as send mail, semaphoring etc... DB2 support 11 of the most common.
Concurrency control = “readers don't block writers and writers don't block readers”

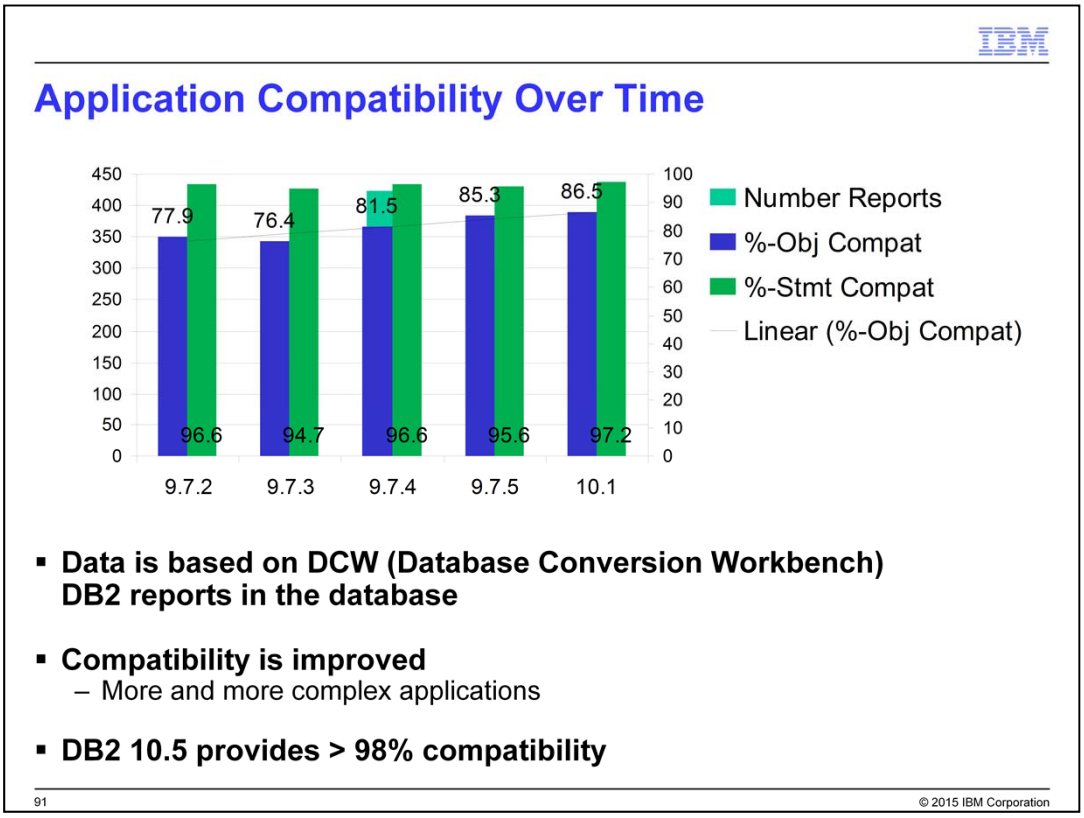

A number of tools are available that will help gauge how much of an Oracle application will port to DB2 without charge. This chart represents the application compatibility across a wide variety of customer workloads. Each DB2 fix pack has included new compatibility features that has helped reduce the amount of SQL modifications required.

The chart shows a number of trends

- Compatibility continues to improve over releases and fix packs
- Customers are beginning to port more complex applications as they gain confidence in the compatibility
- The DCW (Database Conversion Workbench) tool (used to gauge compatibility) is helping drive more compatibility features into DB2

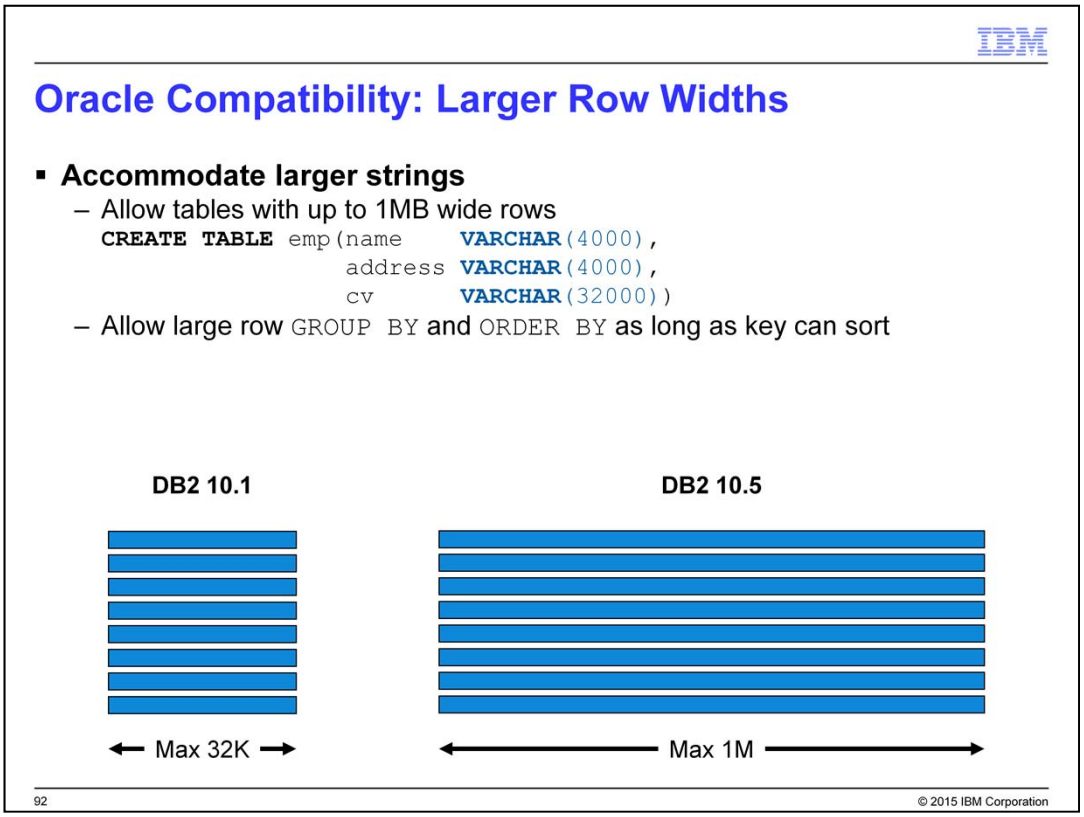

Larger Strings

•Currently a row in DB2 must fit on one page (4/8/16/32K) – excluding LOBs and XML objects. Many applications define columns to have a maximum size (like 4K) which are rarely reached. However, these column sizes are taken into account when creating a table and errors occur due to the total size of a row exceeding 32K.

•DB2 10.5 lifts the row size restriction so that rows can be "larger" than an individual page. Any VARCHAR columns that cause the row to exceed the page size will be "spilled" onto a LOB page transparently.

•Answer sets that exceed 32K is size (due to joins, GROUP BY, or ORDER BY) can also be handled with this new feature.

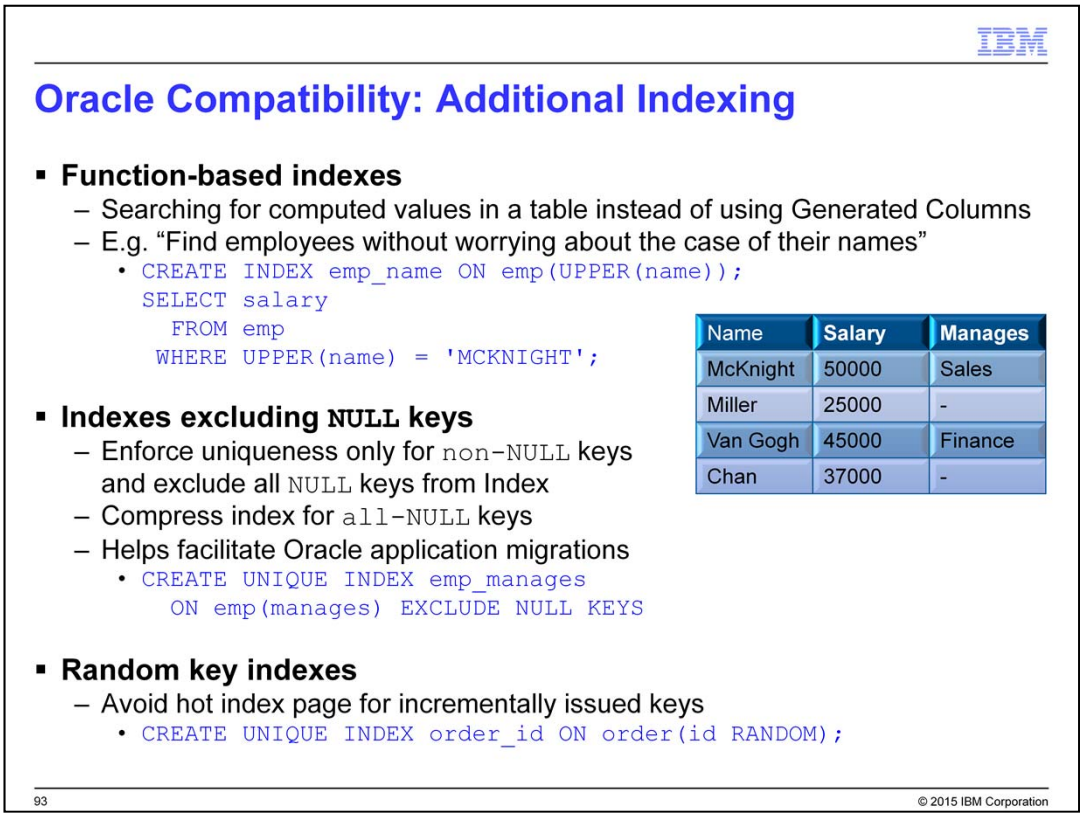

All of the items found on this page are applicable to DB2 workloads in addition to Oracle workloads.

Function-based Indexes

•Function-based indexes allow the definition of an index to be based on a calculation (i.e., UPPERcase)

•Prior to DB2 10.5, a GENERATED COLUMN would be required in the actual table to hold the calculation

•The GENERATED COLUMN would then be used as the source of the INDEX

•Using GENERATED COLUMNs required extra space in the table and caused issues when loading data into the table

Indexes Excluding Nulls

•A UNIQUE index can only contain one NULL value (since another NULL would not meet the UNIQUE requirement)

•Many applications require a column to be UNIQUE, but may not know the value at the time of row creation

•This type of index will support multiple NULL values and only require uniqueness when they become NOT NULL

Random Key Indexes

•Many applications generate "new key" values (i.e., ORDER NUMBER, EMPLOYEE NUMBER)

•When multiple concurrent users are generating sequential values, the same INDEX page becomes a "hotspot" due to multiple update requests

•Adding the RANDOM qualifier to the INDEX will add a hidden HASH value that will split sequential keys onto separate pages and remove much of this contention

•DB2 pureScale can use this feature to minimize contention between members updating the same index

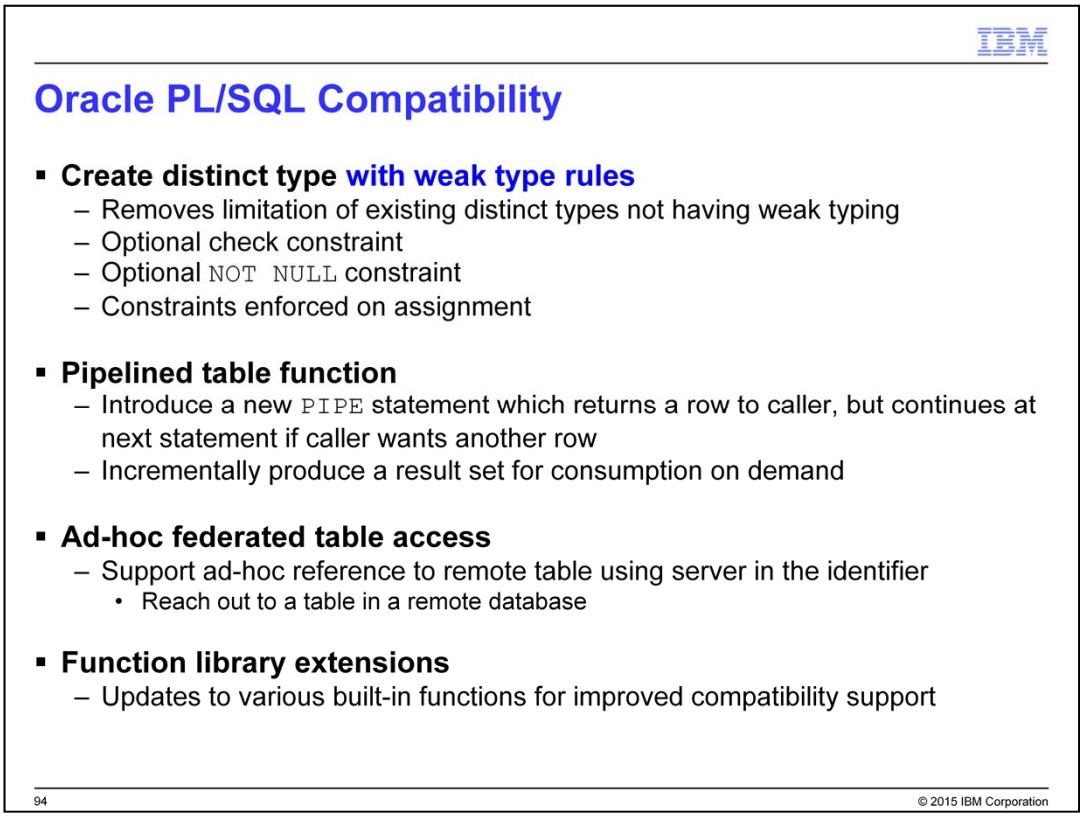

Weak Distinct Types

•Previous versions of DB2 allowed weak typing on base types ("1" = 1 is valid)

•Any subtypes (MONEY based on DECIMAL(31,2)) that a developer created was always strongly typed and didn't allow this weak typing

•In DB2 10.5 weak assignment with distinct types is allowed with an additional feature that allows for constraint checking as part of the definition

Pipelined Table Functions

•A Table function is used to return a result set (that looks like a table) back to the calling SQL statement or application

•Current DB2 table functions must return all of the data before returning to the statement that called it

•In some cases not all of the data may not be required, but there was no means to stop the table function

•Pipelining allows the table function to return a row at a time and stop when no more rows are required by the calling application

Ad-Hoc Federated Table Access

•DB2 allows federated access (to other DB2, Informix, Oracle and other data sources) via the use of the NICKNAME keyword

•NICKNAMEs needed to be defined prior to running a SQL statement against the other data source

•Three-part table names allow the SQL to use the SERVER.OWNER.TABLE syntax to access a table rather than using NICKNAMEs

Function Library Extensions

•A number of functions in the DB2 library have been updated (MOD, RTRIM, LTRIM, SUBSTR4)

•RTRIM/LTRIM – now allows for second parameter to specify a trim-character

•MOD (Modulo) – now supports DOUBLE and DECFLOAT

•SUBSTR4 – equivalent to SUBSTR but a negative start number is used to begin the substring from the end of

the string

JSON stands for JavaScript Object Notation; this is something that's become very, very popular with the emergence of JavaScript as a pretty common programming language.

As of DB2 10.5 for LUW Fix Pack #1, the JSON storage in DB2 for LUW is officially supported and is no longer a technology preview.

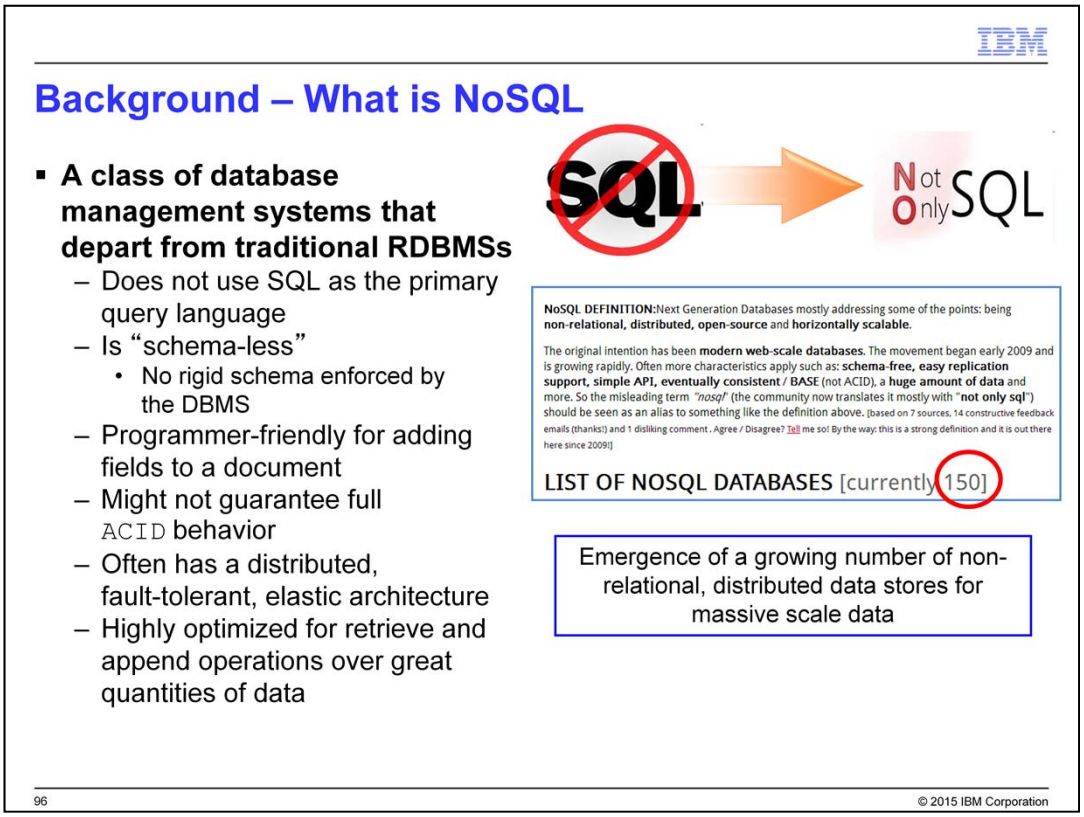

NoSQL = "No use of SQL" or better: "Not *only* SQL"

NoSQL denotes a class of database systems that depart from traditional RDBMSs in one or multiple ways:

- Data format/data model
- Query language, APIs
- Data consistency
- etc.

Goals: performance, scalability, simplicity, schema flexibility – for specific uses case and access patterns

ACID is the acronym used to describe the four properties of an enterprise-level transaction

•ATOMICITY: a transaction should be done or undone completely. In the event of a failure, all operations and procedures should be undone, and all data should rollback to its previous state

•CONSISTENCY: a transaction should transform a system from one consistent state to another consistent state

•ISOLATION: each transaction should happen independently of other transactions occurring at the same time

•DURABILITY: Completed transactions should remain permanent, even during system failure

Details on NoSQL Databases and the current list of NoSQL databases and count: http://nosql-database.org/

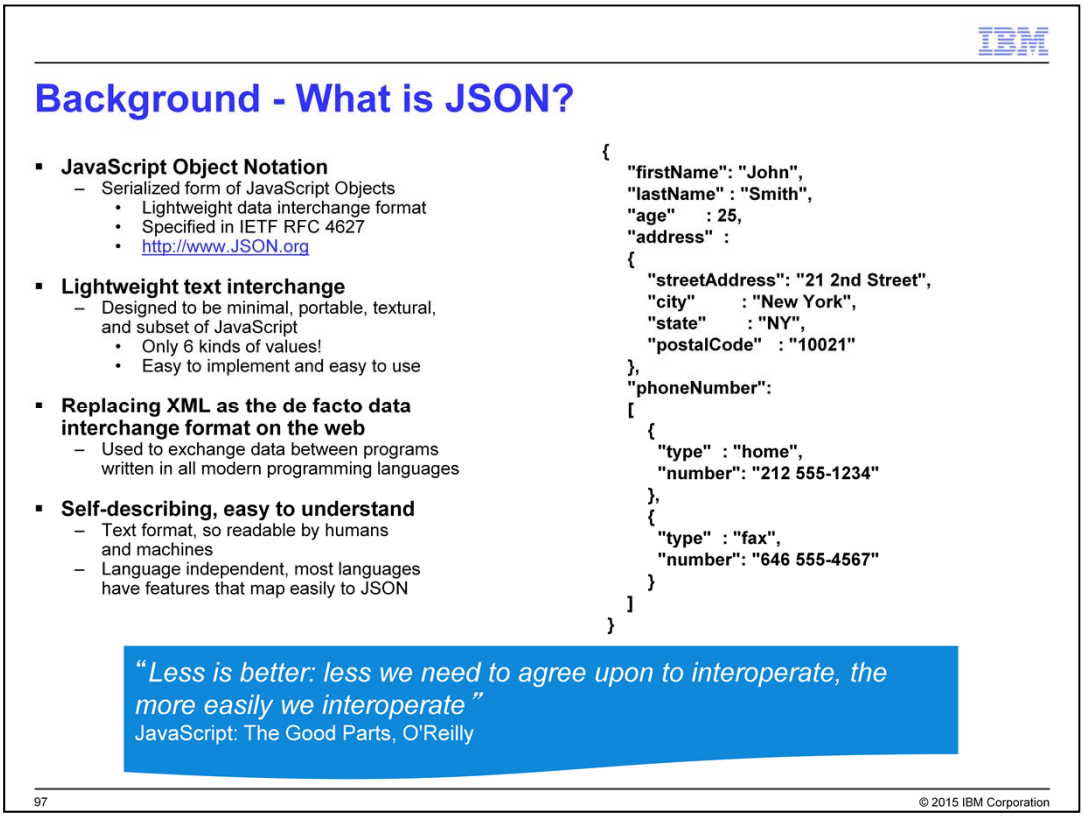

JavaScript Object Notation, is a text-based open standard designed for human-readable data interchange. It is derived from the JavaScript scripting language for representing simple data structures and associative arrays, called objects. Despite its relationship to JavaScript, it is language-independent, with parsers available for many languages. (ref: http://en.wikipedia.org/wiki/JSON)

Characteristics of a JSON Object

- Strings: wrapped in double quotes. "\" is used for escapement and "/" can be escaped so JSON can be embedded in HTML <script> tags. (e.g. </ in HTML only starts </script>, but JSON allows <\/, to produce the same result and not confuse HTML.
- Numbers: Integer, Real and Scientific point, like JavaScript numbers, Leading zero is not allowed.
- Objects: Unordered set of (comma separated) name/value pairs, where a name can be any string, a value can be any JSON value, including arrays and objects. Objects can be nested to any depth, but in reality are kept relatively flat.
- Arrays: Ordered sequence of values, where a value can be any JSON value, including arrays and objects.
- Boolean

- True
- False
- Null

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

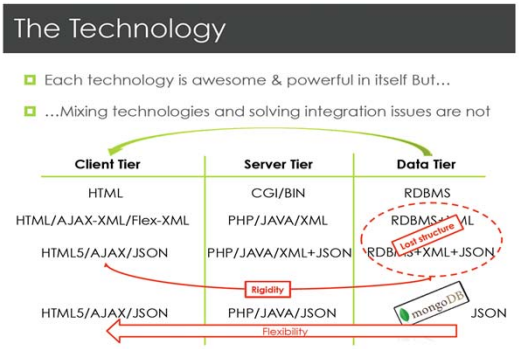

98 © 2015 IBM Corporation

The market is not just suggesting, in fact it's telling us that basically the client tier is moving to HTML5, Ajax, and JSON. As more and more people find out about server-side JSON support in databases, and take notice their client tier supports JSON, they see it's a lot easier to move things back and forth without losing information and you can become much more agile and flexible. Now, this was the same claim that we made about XML databases that all the RDBMS vendors went down that path (albeit with different sophistication levels). Ironically, by the time XML database support got into a state where it was ready for prime time, application developers had already kind of moved off of XML and they were moving on to JSON.

Why did they move to JSON – these fickle developers? And there were a number of things that caused them to go there. One reason is that JSON is a lot more simplistic than XML. Now if you look at the feature and functions of XML, I'd conclude that it's more powerful, but what the application development community found is that JSON was good enough for the things they want to do in the application tier and it mad programming tasks easier. That lead developer to take the

attitude that they didn't really care that XML was more powerful, they found JSON easier to work with and decided to use it.

This application shift in attitude led to the emergence of databases that use JSON to store data, such as MongoDB. The fact that these JSON inspired databases makes them kind of interesting to this class of developers that are kind of gravitating to JSON for their object models, their UI, and so on.

Image slide reference: “How MongoDB adds Flexibility and Agility to the Innovation Cycle” by Olivier Poupeney, DreamFace Interative CEO/CoFounder

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

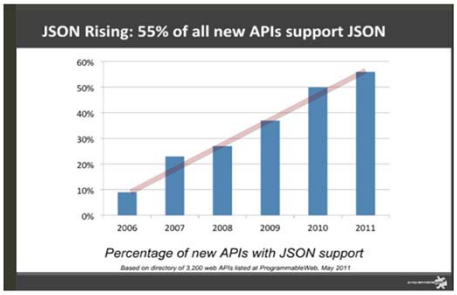

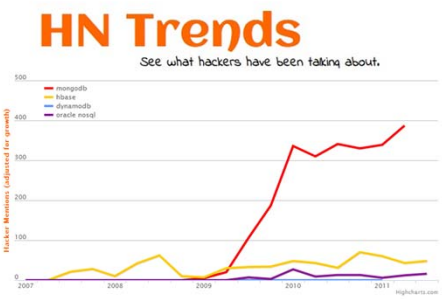

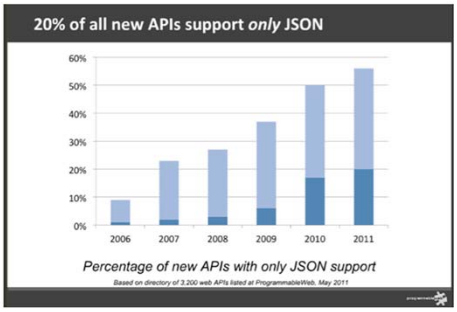

99

Here's some information that came from some of the application development conferences recently. It's really showcasing the rise of JSON and the 'good bye' to XML from a programmers point of view (POV). JSON stands for JavaScript Object Notation; this is something that's become very, very popular with the emergence of JavaScript as a pretty common programming language.

There are various reasons why XML is falling out of favour with developers: XML is more verbose, XML documents are typically a lot larger than the same document in JSON (for example. <Mytag>value</Mytag> versus Mytag:value), there is lots of development work going on in JavaScript, and JSON has a very tight affinity with JavaScript, and most of the new development tools and APIs are supporting JSON, and many of them no longer support XML

+CLICK+

And you can see in the upper right hand corner you can see all the APIs that do support JSON – if you went back five years ago hardly anything supported JSON.

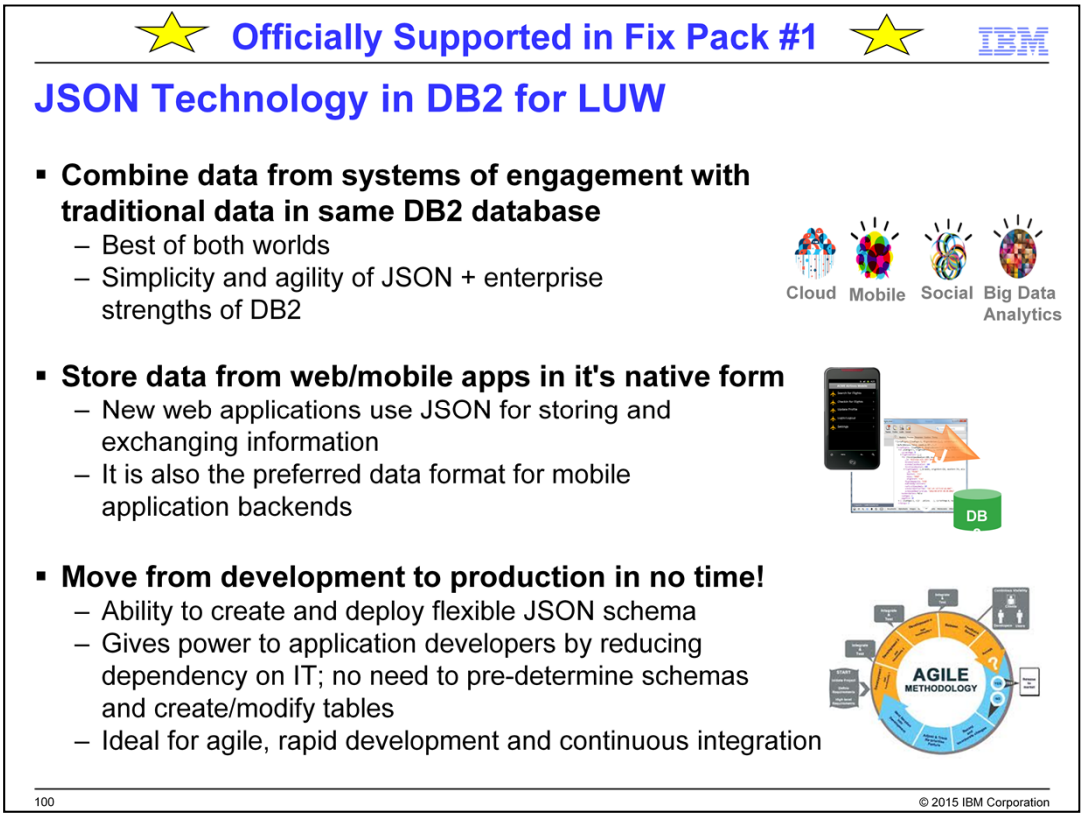

- A JSON Technology Preview was released during the same time as DB2 10.5. It was not part of the DB2 10.5 release.
- As of DB2 for LUW Fix Pack #1, it is now considered part of the DB2 10.5 release and is officially supported.

- JSON support in DB2 will allow you to store/query/update JSON documents in DB2 for LUW using a Java API that is similar to the MongoDB APIs. Plans are also underway to support the same from a node.js environment in a DB2 release.
- Internally, at a high level, it is storing the content into a BLOB column in a DB2 table and there is one built-in function that will allow indexing of the JSON fields in the BLOB. There will also be other JSON user-defined functions (UDFs) the JSON API will use to achieve querying or updating the back-end.
- For an introduction to MongoDB please see this url: http://www.mongodb.org/about/introduction/

- For more information on node.js, please see this url as an example (or google 'what is node.js'): http://radar.oreilly.com/2011/07/what-is-node.html

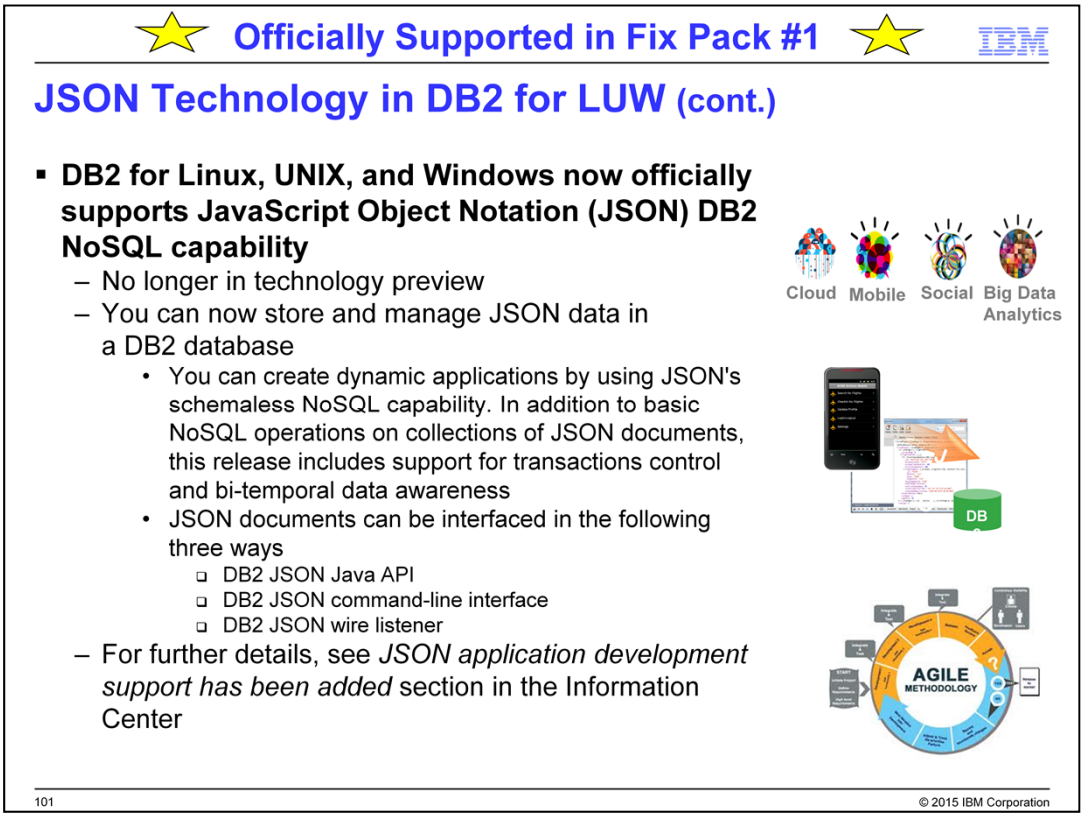

- A JSON Technology Preview was released during the same time as DB2 10.5. It was not part of the DB2 10.5 release.
- As of DB2 for LUW Fix Pack #1, it is now considered part of the DB2 10.5 release and is officially supported.

- JSON support in DB2 will allow you to store/query/update JSON documents in DB2 for LUW using a Java API that is similar to the MongoDB APIs. Plans are also underway to support the same from a node.js environment in a DB2 release.
- Internally, at a high level, it is storing the content into a BLOB column in a DB2 table and there is one built-in function that will allow indexing of the JSON fields in the BLOB. There will also be other JSON user-defined functions (UDFs) the JSON API will use to achieve querying or updating the back-end.
- For an introduction to MongoDB please see this url: http://www.mongodb.org/about/introduction/
- For more information on node.js, please see this url as an example (or google 'what is node.js'): http://radar.oreilly.com/2011/07/what-is-node.html

DB2 10.5 Packaging Simplification

This next section describes some of the packaging changes that will be part of the new DB2 10.5 release.

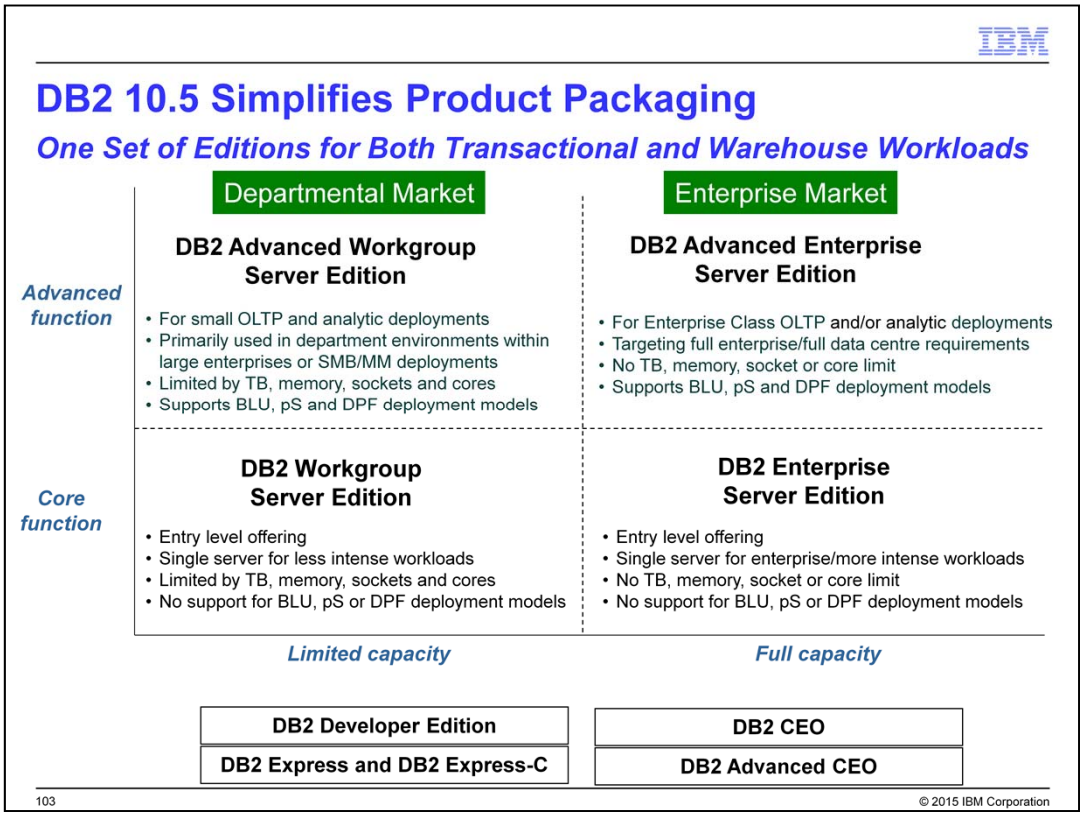

- Down from 13 to 6 editions (not including CEO offerings)
- Down to 1 new optional feature – Advanced Recovery Feature
  - High Performance Unload
  - Merge Backup
  - Recovery Expert
- Grandfather clause rules will be in effect
  - No customer will lose entitlement for currently active licenses
- Ability to purchase stand-alone products will remain for some time
  - Tools, Storage Optimization Feature, DB2 pureScale, etc
  - Required for customers who are on older versions of DB2

DB2 10.5 Business Value Offerings

<<Speaker Name Here>>
<<Speaker Title Here>>
<<For questions about this presentation contact Speaker Name speaker@us.ibm.com>

This presentation will take you through the 4 new Business Value Offerings that were added as of 10.5 FP5.

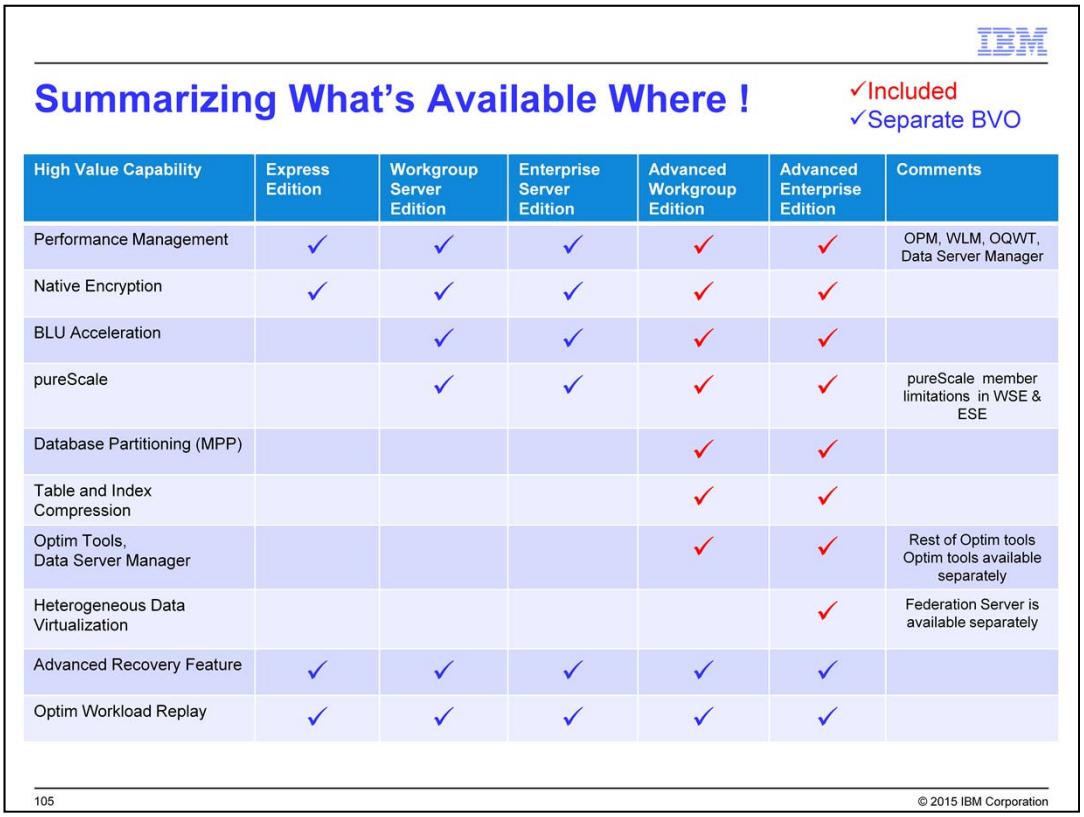

Summarizing What's Available Where !

✓Included
✓Separate BVO

High Value Capability	Express Edition	Workgroup Server Edition	Enterprise Server Edition	Advanced Workgroup Edition	Advanced Enterprise Edition	Comments
Performance Management	✓	✓	✓	✓	✓	OPM, WLM, OQWT, Data Server Manager
Native Encryption	✓	✓	✓	✓	✓	
BLU Acceleration		✓	✓	✓	✓	
pureScale		✓	✓	✓	✓	pureScale member limitations in WSE & ESE
Database Partitioning (MPP)				✓	✓	
Table and Index Compression				✓	✓	
Optim Tools, Data Server Manager				✓	✓	Rest of Optim tools Optim tools available separately
Heterogeneous Data Virtualization					✓	Federation Server is available separately
Advanced Recovery Feature	✓	✓	✓	✓	✓	
Optim Workload Replay	✓	✓	✓	✓	✓	

105

This slide is a one page summary to help complete the picture of what the new packaging looks like. What you basically see here is that we have a variety of capabilities that are purchasable on top of some of our editions.

A subset of the high value capabilities from the Advanced Edition are available through 4 separate BVOs two of which are available on Express and all four of which are available on Workgroup and Enterprise.

The blue checkmarks imply that the capability on the left is purchasable (either via BVO or separate feature) on top of the base edition. For example, DB2 Native Encryption capability can be purchased on top of Express, Workgroup and Enterprise Edition via the DB2 Encryption offering.

The red checkmarks imply that the capability on the left is included in that edition.

As you can see, some of the high value capabilities are still only available by trading up to an Advanced Edition. For example, Database Partitioning, table and index compression and the full suite of Optim tools are only available in the Advanced Editions.

Please note that the heterogeneous data virtualization or heterogeneous federation is the one difference between Advanced Workgroup and Advanced Enterprise.

Please also make note of course that many of the Optim tools are sold separately as well as federation server is sold separately.

The Advanced Recovery Feature and Optim Workload Replay are available as separately purchasable on all editions.

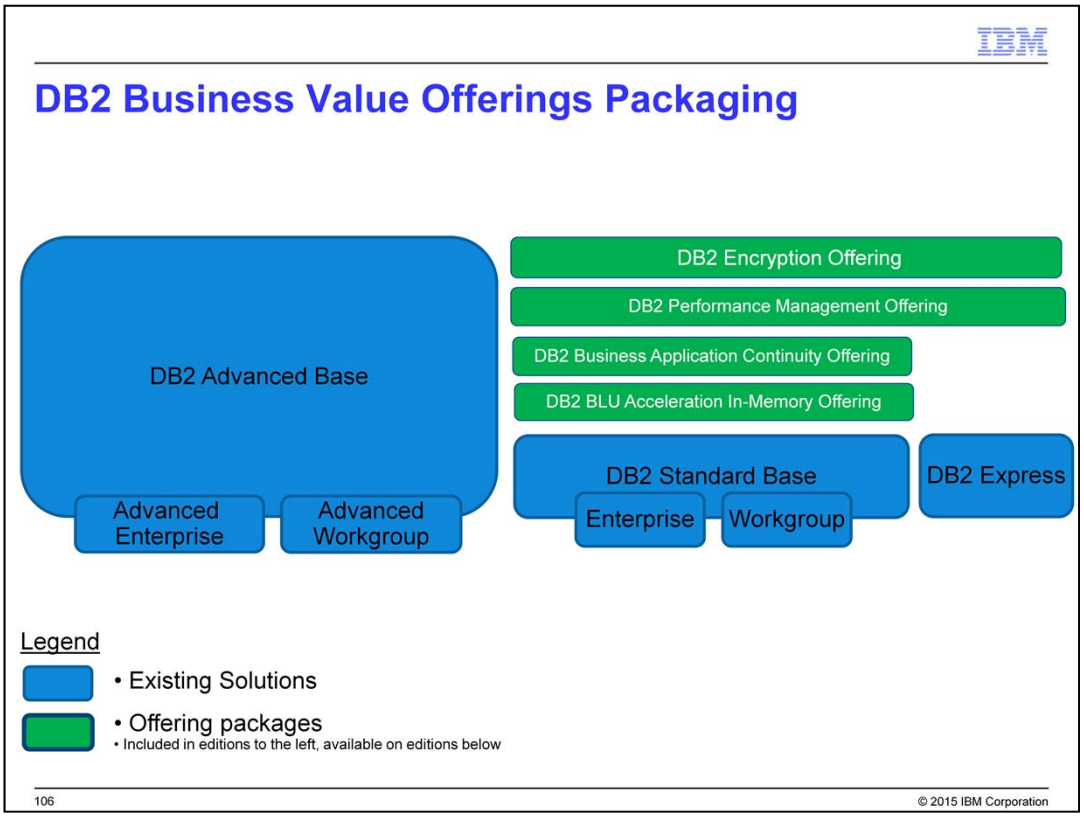

So let's take a look at the 4 new Business Value Offerings in a more pictorial way here on Page 14. So in the Advanced base content all of the value that we're talking about here is included. So BLU as you know is included in the Advanced Editions. pureScale is included. OPM is included.

Workload Manager is included and as I mentioned earlier the new DB2 Native Encryption is included in the Advanced Editions. But now what you have is the ability to sell a couple of these on top of DB2 Express.

This is the extension of the trade up play so there's a conversation that can continue if a client does not have a budget for a full trade up to Advanced Edition and all four of these are included on -- as separately purchasable offerings on top of either the Workgroup Server Edition or the Enterprise Server Edition.

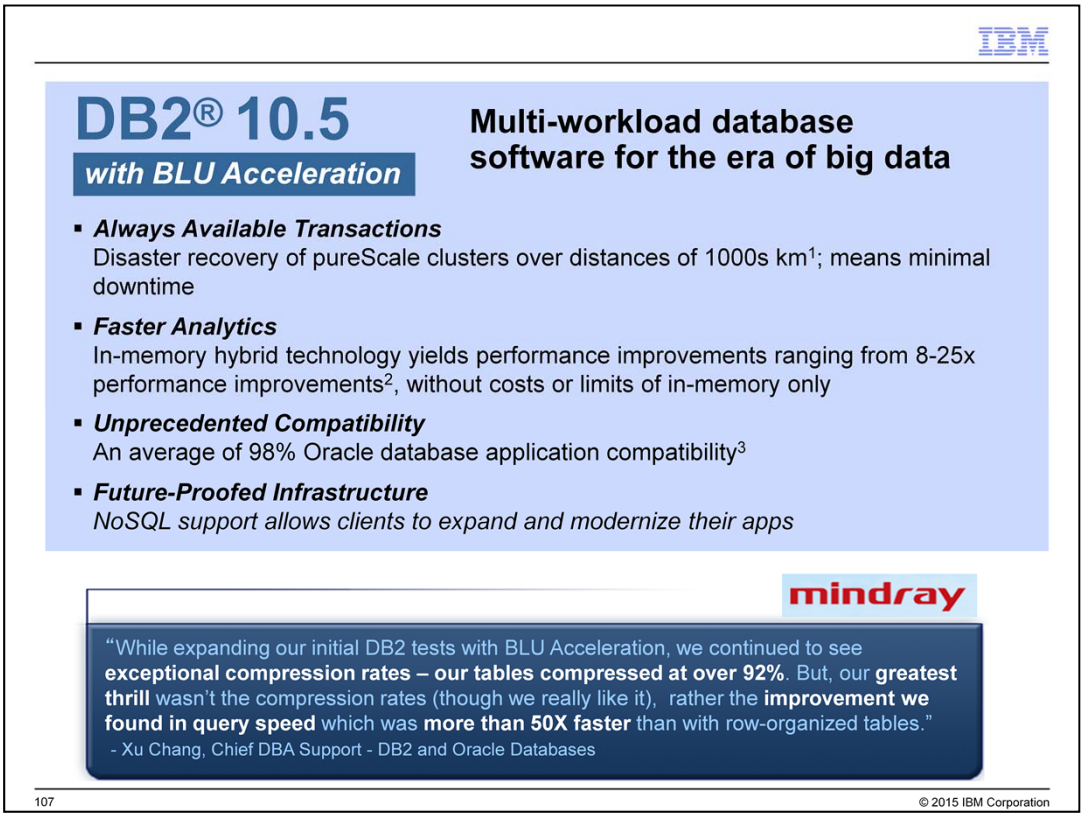

1. Based on IBM design for normal operation under typical workload. Individual results will vary depending on individual workloads, configurations and conditions, network availability and bandwidth.
2. Based on internal IBM testing of sample analytic workloads (not including transactional or OLAP workloads) comparing queries accessing row-based tables on DB2 10.1 vs. columnar tables on DB2 10.5. Performance improvement figures are cumulative of all queries in the workload. Individual results will vary depending on individual workloads, configurations and conditions.
3. Compared prices exclude applicable taxes, and are subject to change without notice. IBM: assumes 120 Processor Value Units. Oracle: assumes 1.0 processor multiplier. Both including Y1 maintenance/support.
6. Based on internal tests and reported client experience from 28 Sep 2011 to 07 Mar 2012.

DB2 10.5 gives customers significant benefits in the areas of:

- Always Available Transactions
  Continuous availability with disaster recovery up to 1000s km; means minimal downtime
- Faster Analytics
  In-memory hybrid yields 25x* faster analytics without costs or limits of in-memory only
- Unprecedented Compatibility
  > 99% Oracle Database application compatibility
- Future-Proofed Infrastructure
  NoSQL support allows clients to expand and modernize their apps

Let’s take a moment for Q&A.

Thank You!