IBM Cloud and Cognitive Software

Unboxing Db2 v11.5 to v11.5.4

October 8, 2020

Walter Alvey
Data Storage, IBM Analytics

© 2020 IBM Corporation

Preface

- These slides are based heavily on two other presentations:
  - "Unboxing Db2 11.5 and a Look Into the Future" given by Mike Springgay at IDUG in Rotterdam in October 2019.
  - "Spotlight on Nebula (Db2 11.5.4)" given by Keri Romanufa at Virtual IDUG in Summer 2020.
  - Some slides were updated by me. (Blame any errors on me!)
- We can cover only a fraction of the new features.
- Apologies for the mixture of master slides.
  - Using any one of the master slides consistently messes up the formatting of least some of the other slides!

Please note :

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.
- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

NOTICE AND DISCLAIMERS :

•© 2020 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.

•U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

•Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed "as is" without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.

•IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply."

•Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

•Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

•References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

•Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

•It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5 and later
- Db2 information resources and resources for developers

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5 and later
- Db2 information resources and resources for developers

Our family of **Hybrid Data Management** solutions built on the **Db2** common SQL engine

Write your SQL once deploy against any form factor ***run anywhere***

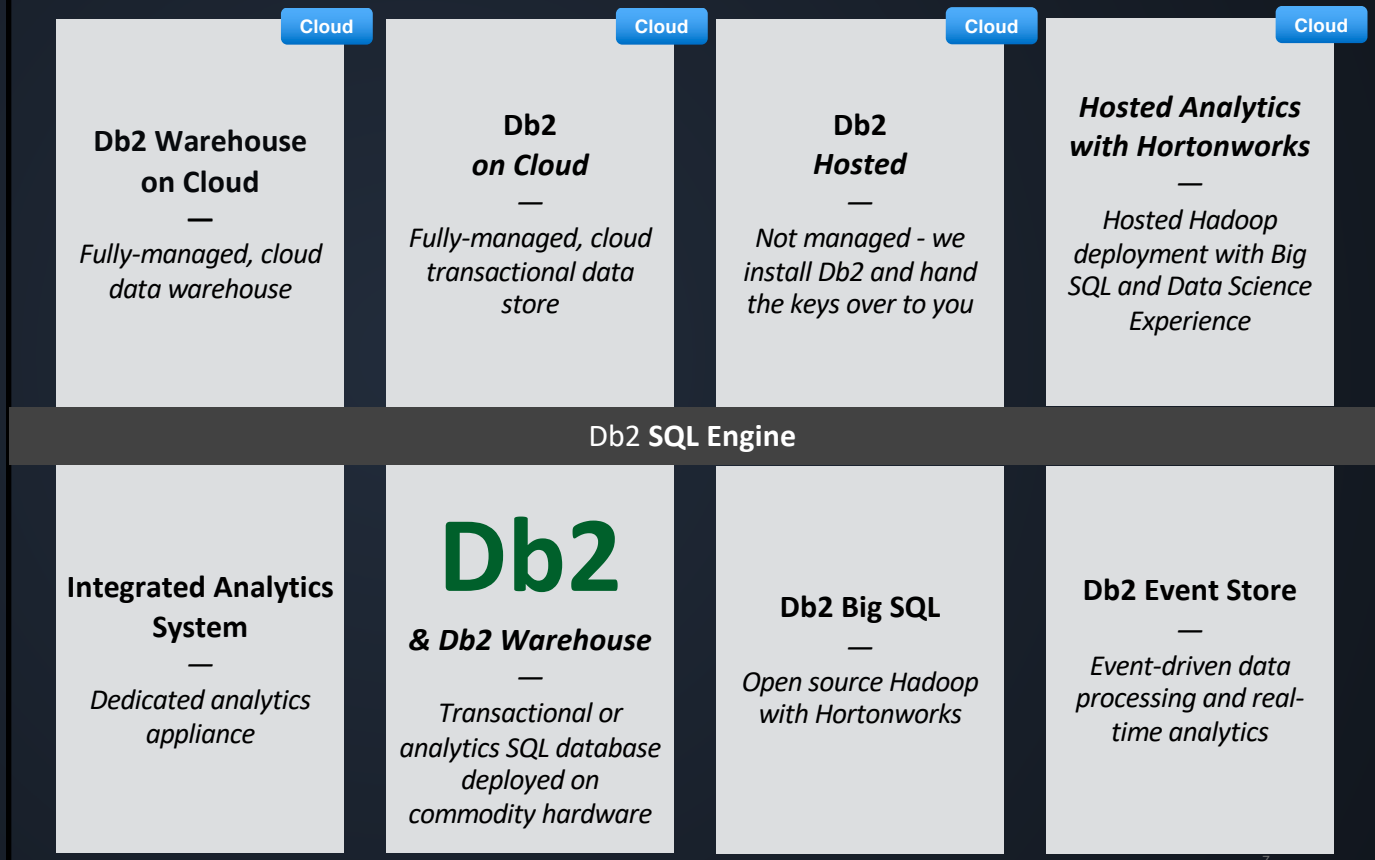

IBM ©2018 IBM Corporation

Convergence

Move all Db2 family members to the same version of the Db2 Common SQL Engine with continuous delivery model

- Reduced latency in feature availability
- Wider availability of existing Db2 Common SQL Engine capabilities across Db2 family
- Consistent experience across family members
  - Can more easily move applications from one form factor to another

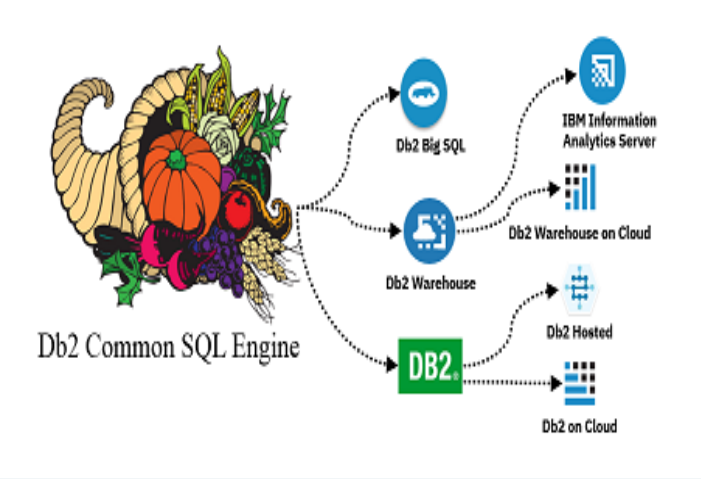

Individual family members pick up updates at different intervals

- E.g. Db2 offers modification updates every 6-12 months while IIAS (appliance), Cloud, and Docker-based environments can move to a new level every 2 months
- Monthly Early Access Program (EAP) updates

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5 and later
- Db2 information resources and resources for developers

Db2 11.5: What's Up With the Numbering?

Db2 11.5 follows a continuous delivery model (same as 11.1)

If a delivery vehicle has:

- new function -- will increment Mod level (same as 11.1)
  - But it will also reset the FP level to 0 (*NEW*)

  → 11.5.M.0

- ONLY fixes -- will increment just the FP level (same as 11.1)
  - This addresses the issue where 11.1 "ifixes" were not easily distinguishable from the base mod pack nor other ifixes off the same mod pack.

    → 11.5.M.1 and higher

The Db2 11.5 lifecycle

BLUE - Multi-platform install path -- mod packs and fix packs

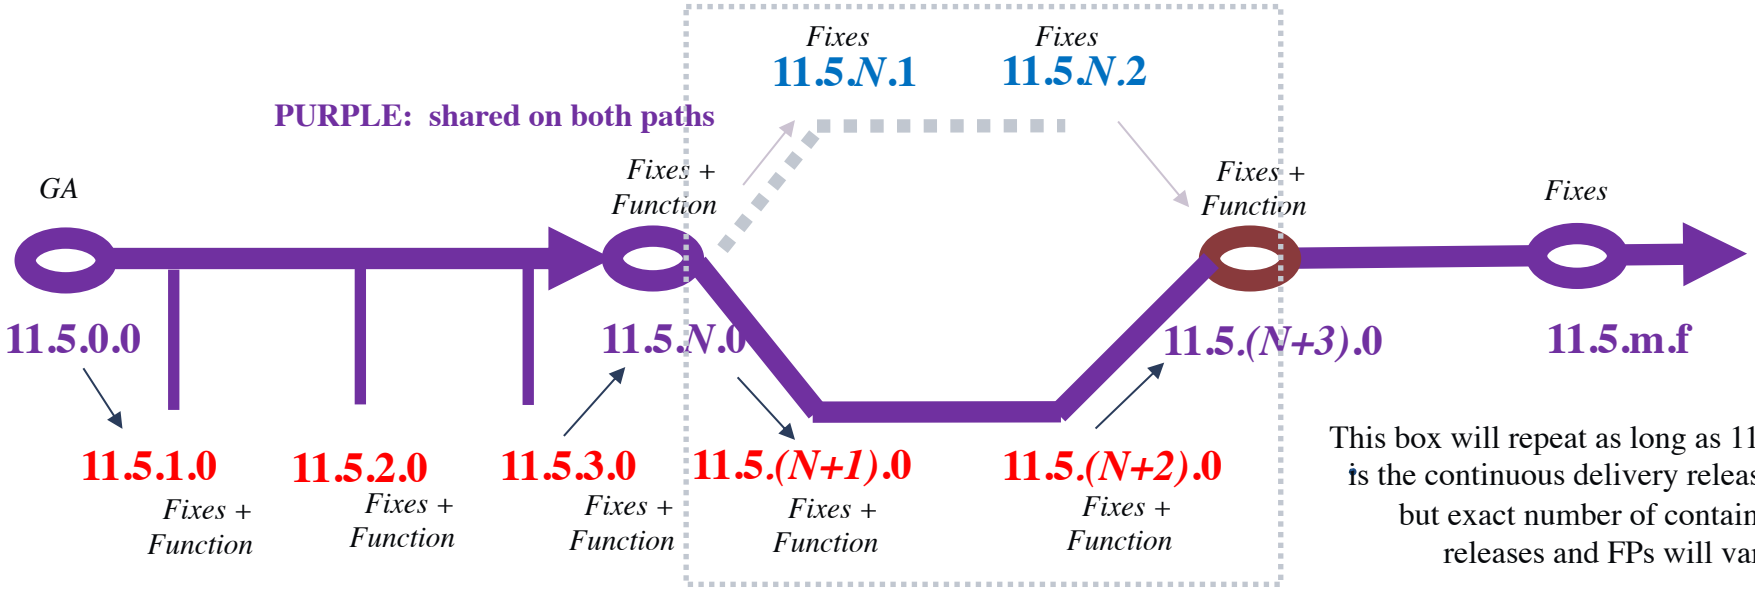

This box will repeat as long as 11.5 is the continuous delivery release, but exact number of container releases and FPs will vary.

RED - Containerized Install Path -- always* mod packs!

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5 and later
- Db2 information resources and resources for developers

Categories of enhancements

- Availability
- BLU
- Compression and storage
- Federation
- Logging
- Manageability
- Monitoring
- Performance
- pureScale
- SQL compatibility
- Security
- Spatial
- Workload management

Categories of enhancements

- The next 18 slides are a bulleted list of all the enhancements, grouped by release, then by category.
- Those that we will cover in this presentation are in red.
- Most are in 11.5.4.

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5 and later
- Db2 information resources and resources for developers

List of what's new in V11.5

- **Manageability enhancements**
  - External table support
    - With the CREATE EXTERNAL TABLE statement, you can store data outside a Db2 database while retaining the ability to query that data.
  - LOB support in columnar tables
    - Use data types CLOB, BLOB, and DBCLOB in columnar tables.
  - Support for compiled SQL Procedural Language (PL) scalar functions in DPF (MPP) environments.
- **Monitoring enhancements**
  - New monitoring metrics for determining failure rate of SQL statements
  - New monitor interfaces for workloads at the service superclass level
  - Improved table functions provide access to information without having to perform joins
  - Improved table functions provide direct access to WLM statistics
  - New monitor element for concurrently executing activities
  - New monitor elements report thread and memory usage for workload and service class objects

List of what's new in V11.5 *(continued)*

- **Performance enhancements**
  - Support for storage devices that use a 4KB sector size
  - SQL insert and update into columnar tables
    - Optimized to support large data sets efficiently, including creation of encoding dictionary.
  - Improved encoding dictionary support for SQL insert and update
    - Enhanced to provide encoding in some cases where encoding was not previously possible including for Declared Global Temporary Tables (DGTTs), Not Logged Initially (NLI) tables, and tables created using Create Table As Select (CTAS) statements.
  - Support for deferred memory commit (tech preview in v11.5, fully supported in v11.5.4)
    - Improves the performance of large memory allocations, such as those that occur during database activation and some STMM tuning operations.
  - Automatic collection of column group statistics
    - Updated to discover and collect column group statistics on base tables.
- **Compression and space reclamation enhancements**
  - Auto recompress for columnar tables
    - Uses the new dictionary to recompress rows at start of the table that were compressed using the older dictionary
  - Improved reclamation of empty extents in columnar tables

List of what's new in V11.5 *(continued)*

- **SQL compatibility enhancements**
  - DROP TABLE *IF EXISTS* and CREATE TABLE *IF NOT EXISTS*
  - New module DBMS_APPLICATION_INFO
    - Procedures that set and return customer client information
  - New module UTL_RAW
    - Routines for manipulating BINARY and VARBINARY data.
- **Workload management enhancements**
  - Simplified declaration of THRESHOLD objects
    - The ACTIVITIES and ENFORCEMENT clauses on the CREATE THRESHOLD statement are now optional, leading to simpler declaration of THRESHOLD objects.
  - Cascaded drop of WLM service classes
    - Dropping a disabled service class with the DROP statement now drops all dependent service subclasses, thresholds, and work action sets.
- **Data server client and driver enhancements**
  - Call level interface (CLI) driver enhancements
  - IBM Data Server Provider for .NET enhancements
  - IBM Data Server Driver for JDBC and SQLJ enhancements

List of what's new in V11.5 *(continued)*

- **pure Scale enhancements**
  - Improved re-use of free pages across Db2 members
    - For tables that have the default APPEND OFF option set, re-use of free space on data pages has been improved. Particularly in the case where the majority of DELETE operations are performed on a distinct subset of members in the cluster.
  - Simplified geographically dispersed Db2 pure Scale cluster (GDPC) configuration
    - Automate the configuration of public network equivalency resources for the CF and the condition response pair for the public ethernet IP subnet.
  - Improved CF recovery time
    - On all Db2 pure Scale configurations, CF recovery is faster in public Ethernet failure scenarios.
  - Cross Invalidation Connections in RDMA Networks
    - Multiple cross invalidation (XI) connections in RDMA networks can improve TPS performance in GDPC clusters.
  - IBM Spectrum Scale changes to Db2 pure Scale instances
    - The default cipherList security mode is changed to AUTHONLY from EMPTY. Some IBM Spectrum Scale configuration parameters have changed.
    - The configuration automatic GPL compilation feature is enabled.
  - New registry variable supports enhanced security
    - Host-based firewall enablement is now an option for single-node, DPF (MPP), and pure Scale Db2 configurations, through the use of the new DB2_FIREWALL_PORT_RANGE registry variable.
  - Improved castout performance

List of what's new in V11.5

- **Spatial data analysis enhancements** (tech preview in v11.5, GA in v11.5.4)
  - Perform spatial analysis on spatial data in column-organized and row-organized tables.
  - Use filtering based on geometry properties to speed up query performance.
- **Federation enhancements** (tech preview only)
  - Access to Hyperledger Fabric Blockchain data.

List of what's new in V11.5.1

- **Availability enhancements**
  - Improvements for how node failures are handled
    - This reduces the window of disruption to the end user during node failure recovery on DPF (MPP) configurations.
- **Manageability enhancements**
  - New buffered scan option available in the EXPLAIN_ARGUMENT table
    - A new column, BUFFSCAN, provides the option of enabling a buffered scan. Scan results can identify which access method was used on a row, in the event of a performance issue.
  - Ability to lift or enforce locks placed on tables by INSERT operations
    - A new registry variable DB2_INSERT_RELEASE_LOCK_ON_DUP has been added, which allows users to enforce or lift locks on tables caused by INSERT operations.
  - External storage enhancements to the CREATE EXTERNAL TABLE statement
    - The CREATE EXTERNAL TABLE statement now supports data storage on Microsoft Azure Blob Storage, with the ability to query the data.
- **Performance enhancements**
  - Enhanced query performance against column organized tables
    - Query performance against column organized tables is enhanced by improvements to memory utilization. This feature is most utilized by complex queries on column organized tables, and in particular by queries involving wide value types that have a large percentage of the column values shorter than the defined schema width.

List of what's new in V11.5.1 *(continued)*

- **SQL compatibility enhancements**
  - Netezza expressions can now be referenced in a WHERE clause by an alias
    - A new Netezza Platform Software (NPS) compatibility feature allows expressions to be referenced in the WHERE clause by its alias in a select list.
  - New module added to the SYSIBMADM schema
    - A new module, DBMS_LOCK, has been added to the SYSIBMADM schema.
    - The DBMS_LOCK package makes Oracle lock management services available to PL/SQL developers.

List of what's new in V11.5.2

- **Monitoring enhancements**
  - New monitor element to identify a failover event
  - New monitor element to identify time of execution on a coordinator member
    - A new monitor element, coord_time_created, has been created to help analyze the behavior of an activity, in conjunction with other activity history elements.
  - New activity event monitor elements for joining and pruning activity event monitor tables
- **Manageability enhancements**
  - The SSL_SVR_LABEL database manager parameter can now be configured dynamically.
    - The SSL_SVR_LABEL database manager configuration parameter can now be updated dynamically. There is no longer a need to restart the instance to change this parameter.
    - Db2 is now able to change the SSL server certificate used for incoming connections while the instance is running.
- **Performance enhancements**
  - Improved performance of Hash Joins (HSJN) that have a single equality join predicate and one, or more, non-equality joins predicates when accessing columnar tables
    - This improvement enables Hash Joins joining columnar tables or streams to apply additional non-equality join predicates during the main Hash Join for optimized performance.
  - Early aggregation / early distinct for column-organized queries
    - Additional strategies have been introduced for column-organized query planning. These strategies will help improve query performance by aggregating and removing duplicates earlier as the query is run.
    - These strategies are counterparts to existing strategies used in row-organized query planning.

List of what's new in V11.5.2 *(continued)*

- **SQL compatibility enhancements**
  - WITH clauses can now be nested
    - The query body of a common-table expression (that is, a WITH clause) can now contain additional common-table expressions.
  - A TRUNCATE TABLE statement no longer requires an IMMEDIATE clause
    - For column-organized tables, the IMMEDIATE clause is now optional for a TRUNCATE statement.
    - When the IMMEDIATE clause is not specified, the operation can be stopped at any point in the transaction's scope before it completes.
    - The truncated table is then immediately available for use within the same unit of work.
    - For a TRUNCATE statement that is issued without the IMMEDIATE clause, you can issue a ROLLBACK statement to undo the TRUNCATE operation, even if another data-changing operation was issued after the original TRUNCATE statement.
    - This will undo everything, including the truncate operation.
    - After this is done, you can reclaim storage manually by running the REORG RECLAIM operation, or you can wait for the health monitor trigger to reclaim storage automatically (improved in v11.5.4 to use background async task).
  - LFINSTRING specifies how to interpret unescaped line-feed characters in string data in an external table
    - When you create an external table that uses a text file format, a new option called LFINSTRING lets you specify how unescaped line-feed (sometimes called LF or newline) characters in string data are to be interpreted for that table.
    - This option is not supported for unload operations, and applies only to line-feed characters, not to carriage-return line-feed (CRLF) characters.

List of what's new in V11.5.3

- **Performance enhancements**
  - Introduction of a cache for User ID and Password based authentication and group lookup
    - A new cache for User ID and Password based authentication plugins has been introduced to relieve pressure on backend authentication mechanisms.
    - The cache stores information about successful authentications and compares the information from new, incoming authentication requests against the cached entries to see if a valid match is found. If the match is found, then the new authentication request is considered successful and subsequent Db2 post-authentication processing begins.

List of what's new in V11.5.4

- **Compression and storage enhancements**
  - Improved compression for string data types
    - For string data types, Db2 now has the ability to use page-based compression algorithms in addition to the existing column dictionary-based compression algorithms. Especially during bulk insert or update operations involving columns that contain many unique values, the compression may be greatly improved for character, graphic, and binary data types.
    - There are two page-based compression algorithms and Db2 will automatically choose the algorithm that provides the best compression.
      - The first algorithm targets data that has longer repeated patterns, such as text and URLs.
      - The second algorithm targets data that has 16 or fewer distinct bytes and only short repeating patterns. Such data typically includes numbers, dates, times, and timestamps when stored as strings.
    - This feature is disabled by default because it is fallback incompatible, and it must be enabled explicitly. For details, see the DB2_COL_STRING_COMPRESSION registry variable.
  - Reduced Synopsis Table Storage for Small Tables
    - Synopsis tables for small column-organized tables can have excessive overhead due to Db2's partitioning and storage allocation. The unused allocated storage for synopsis tables can be excessive in comparison to the base tables.
    - The excessive storage consumption by synopsis tables can be avoided without performance penalty by deferring the creation of synopsis tuples until storage consumption overhead can be minimized.
    - This feature is disabled by default because it is fallback incompatible, and it must be enabled explicitly. For details, see the DB2_COL_SYNOPSIS_SETTINGS registry variable.

List of what's new in V11.5.4 *(continued)*

- **Compression and storage enhancements** *(continued)*
  - Fallback-incompatible features are the features that, once enabled, will cause the database to be incompatible with any prior release.
    - Because of the fallback impact, those features are now disabled by default to give users more control.
    - Users must make a conscious decision that they will not fallback and explicitly turn on the new features using external registry variables.
    - After upgrading to the latest release, falling back to a prior release remains a supported operation until the explicit enabling of any fallback-incompatible features.
    - Once any fallback-incompatible feature is enabled, fallback compatibility is no longer guaranteed. Falling back to an earlier release will likely require restoring the database from a backup from that earlier release.
- **Database logging enhancements**
  - Advanced Log Space Management (ALSM)
    - Advanced Log Space Management (ALSM) helps to minimize application failures caused by log full errors (SQL0964N) as the result of a long running transaction holding back the active log.
    - This is particularly useful if a long running transaction does not generate much log data in comparison to other concurrently running transactions.
- **Data movement enhancements**
  - Improved efficiency in removal of temporary storage data
    - When performing a LOAD operation from a CURSOR file type, where the cursor was defined using the DATABASE keyword option (also known as the REMOTEFETCH load operation), LOB data which is fetched from the cursor is more efficiently and expediently purged from temporary storage, thereby preventing unnecessary temporary tablespace full conditions.

List of what's new in V11.5.4 *(continued)*

- **SQL compatibility enhancements**
  - TRUNCATE TABLE without IMMEDIATE clause supports asynchronous reclaim at the end of the transaction
    - Storage can be automatically reclaimed through a background asynchronous task.
  - Lock avoidance for catalog tables for external user queries only.
    - In previous versions of Db2, when CC (Currently Committed) was enabled, then DB2_SKIPINSERTED, DB2_EVALUNCOMMITTED, and DB2_SKIPDELETED were not supported for user-initiated catalog table scans.
    - This restriction is lifted by the option to change the locking behavior, thus improving the concurrency for user-initiated catalog scans.
    - The isolation levels that are supported depend on the optimization.
    - You can set the DB2COMPOPT registry variable with the LOCKAVOID_EXT_CATSCANS option to enable catalog scans on external queries. This registry variable setting does not impact the behavior of internal queries on the Db2 catalog tables.
  - New/enhanced built-in scalar functions
    - ASCII_STR, NCHR, TO_MULTI_BYTE, UNICODE_STR, TO_SINGLE_BYTE
  - New Netezza TIMESTAMP string support
    - The Netezza timestamp format (MM-DD-YYYY HH24:MM:SS) is recognized in Db2.
  - WITH clause in nested-table-reference and derived table usage
    - The query body of a common-table expression (WITH clause) can now contain additional nested-table-reference and derived table usage, except for subqueries in predicates. For these subqueries, a WITH clause in nested-table-references and derived tables is not possible.

List of what's new in V11.5.4 *(continued)*

- **SQL compatibility enhancements** *(continued)*
  - NULL ordering
    - In Db2, NULL values are considered higher than any other values. By enabling NULL ordering, NULLS are considered as the smallest values in sorting.
    - You can enable this new option by setting the DB2_REVERSE_NULL_ORDER registry variable to DB2_REVERSE_NULL_ORDER=TRUE.
    - By default, the DB2_REVERSE_NULL_ORDER registry variable is set to FALSE.
  - External Table with COMPRESS GZIP option does not need data object with .gz extension
    - When you use the COMPRESS GZIP option, you can now choose to specify the value with or without the .gz extension for the DATAOBJECT or FILE_NAME option.
  - Changed behavior of DECIMAL scalar function empty string in NPS mode
    - In NPS mode, casting an empty string to DECIMAL now returns 0.
  - DAYS_BETWEEN, WEEK_BETWEEN, MONTHS_BETWEEN, HOURS_BETWEEN, MINUTES_BETWEEN, SECONDS_BETWEEN scalar functions behavior change in NPS mode
    - In NPS compatibility mode, the DAYS_BETWEEN, WEEK_BETWEEN, MONTHS_BETWEEN, HOURS_BETWEEN, MINUTES_BETWEEN, SECONDS_BETWEEN scalar functions always return a positive number.
  - New SKIP LOCKED DATA clause for row-organized tables
    - The SKIP LOCKED DATA clause specifies that rows locked by other transactions are skipped. This occurs when incompatible locks that would block the progress of the statement are held on the rows.
  - RID function available for columnar tables.

List of what's new in V11.5.4 *(continued)*

- **High availability enhancements** (tech preview only)
  - Support for Pacemaker as the cluster manager for automated failover to HADR standby
    - You can now deploy automated failover to HADR standby solution for on-premise and cloud instances with Pacemaker as the cluster manager.

List of what's new in V11.5.4 *(continued)*

- **Manageability enhancements**
  - Ability to block ALTER TABLE actions that would put the table into reorg pending state
    - Users can disallow ALTER TABLE operations that place a table into a reorg pending state, preventing the unintentional loss of full access to a table. For more information, refer to DB2_BLOCK_REORG_PENDING.
  - Alternative handling of table space error during database recovery, including HADR standby
    - Users can specify not to use the default handling when database recovery encounters a table space in an invalid or error state.
    - The default handling is to continue the recovery on other valid table spaces and leave the invalid table space to be recovered in a future recovery operation. This default behavior is preferable when the affected table space is only a small portion of the database and most applications can function with the valid table spaces.
    - Users can specify a different behavior by setting the new DB2_FAIL_RECOVERY_ON_TABLESPACE_ERROR registry variable to ROLLFORWARD.
  - Ability to specify the query optimizer version
    - The DB2_OPTIMIZER_VERSION registry variable can be used to have the query optimizer run at a previous Db2 version in order to disable query optimizations introduced between that version and the current Db2 version, to avoid performance regressions due to inappropriate access plan changes.

List of what's new in V11.5.4 *(continued)*

- **Performance enhancements**
  - Faster database activation
    - Deferred memory commit is now fully supported in production environments which leads to faster database activation.
    - Introduced as a technical preview in Version 11.5 GA, this feature improves database activation time by initializing the bufferpool and locklist memory asynchronously in the background.
    - This allows database activation to be complete sooner and for application connections to be connected quicker.
  - Index management improvements
    - Index management performance enhancements have been made in high index contention workload scenarios through improvements to index page split locking mechanisms.
  - New values for database configuration parameter stmt_conc
    - Setting stmt_conc to COMMENTS (and COMM_LIT) ensures that SQL statements that are otherwise identical but end with a different simple comment text (beginning with '--') share the same entry in the package cache thus saving space in the package cache and avoiding unneeded compilations.

List of what's new in V11.5.4 *(continued)*

- **Security enhancements**
  - Consume JWT (JSON Web Token) for SSO
    - Log in to Db2 without providing a userid and password using JWT (JSON Web Tokens).
- **Workload management enhancements**
  - Adaptive Workload Manager
    - Analytics databases can now enable and use the Adaptive Workload Manager, which will provide increased flexibility, stability and simpler configuration and tuning.
    - In order to enable the Adaptive Workload Manager, the WLM_ENABLE_ADMISSION_CTRL stored procedure must be run and the database needs to be restarted.
- **Problem determination enhancements**
  - Db2 Historical Monitoring (db2histmon)
    - To aid with problem determination and improve the availability of diagnostic and monitoring data at the first occurrence of a problematic event, the new Db2 Historical Monitoring (db2histmon) scripts and framework are now available.
    - These scripts deploy a broad set of monitoring and diagnostic data collection, at varying intervals, and manage the archival and purging of this historical data. Data collection sets can be customized as needed and can be deployed on all modern versions of Db2.
    - Download the db2histmon scripts: https://github.com/IBM/db2histmon
- **Data server client and driver enhancements**
  - .NET provider enhancements
  - CLI driver enhancements
  - IBM Data Server Driver for JDBC and SQLJ enhancements

List of what's new in V11.5.4 *(continued)*

- **Federation server enhancements**
  - Optimizes federation JDBC wrapper for connecting to various (total of 21) data sources
  - Support for AutoREST Connector
    - Enables Autonomous REST Connector which is a JDBC connector for accessing to various JSON-based data sources through RESTful API
  - Support for JWT (JSON Web Token) for SSO
    - Logs into a remote data source without providing a userid and password. Experience single sign-on (SSO) with Db2 applications.
  - Inter-partition parallelism with parallel fetching
    - Exploits all the applicable partitions to involve in the execution of one Federated SQL statement to accelerate the execution speed which can gain better performance improvement.
  - Support bulk insert for Oracle through Oracle Net8 wrapper
    - Inserts a set of records into Oracle in one insert operation, and gain performance benefit against one record per one insert operation.
  - Federation FMP enhancement - Buffer size optimization
    - Improves federation nickname query performance by optimizing the block fetching buffer size.
  - Pushdown capability enhancement for MongoDB through NoSQL wrapper
    - Filters will be pushed down to MongoDB. This will reduce the data transfer.
  - Federation pushdown capability enhancement on FFNR for various (total of 12) data sources
    - This will reduce the data transfer.

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5 and later
- Db2 information resources and resources for developers

Performance - Columnar Insert/Update/Delete

- Improved performance of insert, update, and delete in columnar tables
  - Exploit all aspects of modern processors including single-instruction-multiple-data (SIMD) and specialized instructions available on some CPUs.
  - Improvements to data structures to maximize cache-line efficiency.
  - Improved algorithms and data structures for bulk insert operations.
    - Improvements to row-by-row (trickle) inserts is in progress!
  - Split bulk inserts (eg, INSERT INTO … (SELECT FROM …)) into multiple threads running in parallel.
  - Many small improvements such as reduce number of copies and moves of data.

- Critical to maximize ETL/ELT batch performance.

Performance - Vectorized Insert/Update

Pre-v11.5 approach

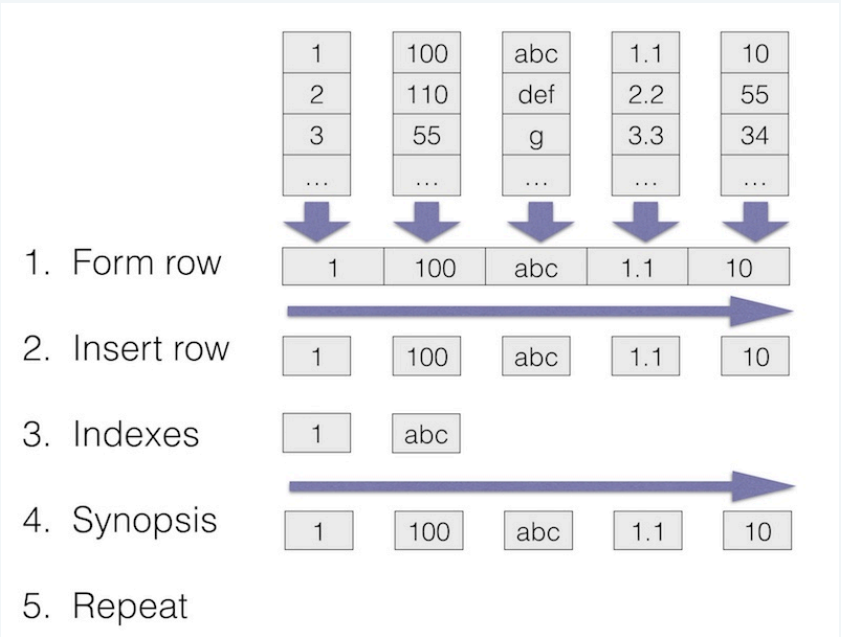

v11.5 approach

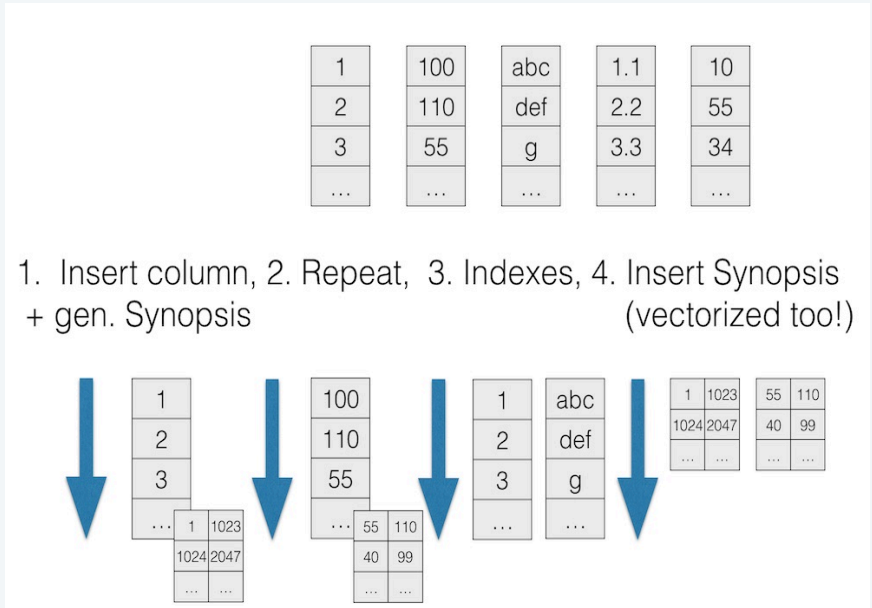

Performance - ETL Performance Example

- Data ingest rate
  - 1 TB/hour before enhancements
  - ***Now ~5 TB/hour (IIAS)***
  - Depends on data, machine, workload, etc
  - Matches or exceeds Netezza in many cases
- >10 TB data
- Table remains online

Parallel Insert Degree/Time/Speedup

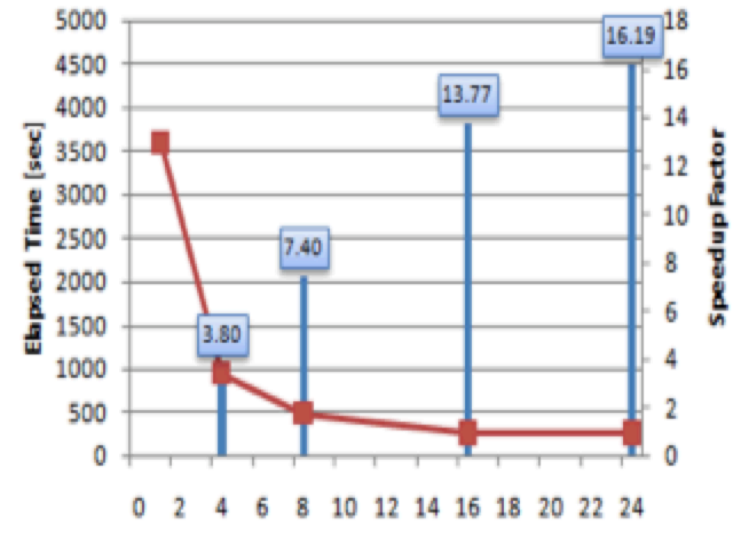

Performance – Reduce the number and size log records

- Reduce number of UNDO log records
  - Pages freed as part of rollback don't need to be undone.
  - Can show up to 50% savings in log space and improved performance of rollback.
  - Always on.
- Reduce size of XOR log records
  - Improve the algorithms that determine changed data so that delta log records are smaller.
- Reduced REDO logging
  - After some insert work has occurred, rather than write log records and then flush the page after the commit, we can do the flush at commit time and skip the log write.
  - Can show up 95% reduction in log space.
  - Rollforward to *pointInTime* or END OF LOGS syntax will be blocked.
  - For that reason, this is off by default in the on-prem form factors.

Storage Reduction – Reorg recompress for columnar tables

- When Automatic Dictionary Creation (ADC) or Vectorized-ADC is used, some portion of the data will be inserted before the dictionary is created
  - REORG RECOMPRESS is focused on these uncompressed first portions of the data.
  - NOTE: Amount of data inserted before creating a dictionary has been increased.

- REORG RECOMPRESS simply applies the compression dictionary to the first portion of the data that was uncompressed initially.
- Done automatically by a new background task (the new compression daemon) once a new dictionary is available and no writer is using the old dictionary.
- Background task will be enhanced in the future!

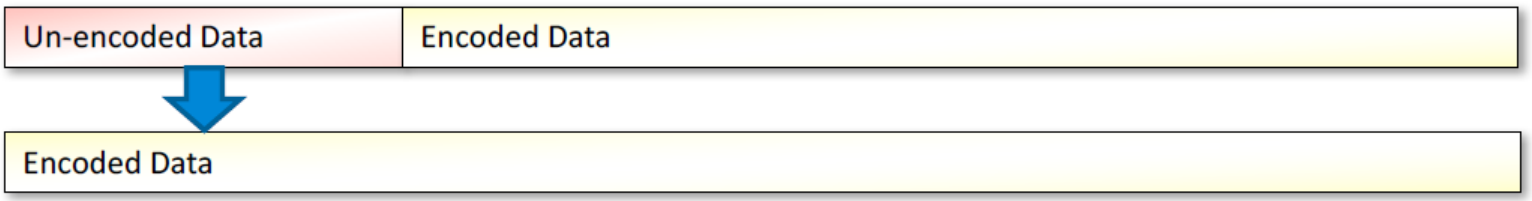

Storage -- BLU Compression IMPROVEMENTS 1/2

2 NEW page-level compression methods for string data:

- 'repeating pattern'-based compression (LZ4 based)
  - Works well for geospatial data, URLs, comment fields, etc.
  - Compression rate typically 2-4x but depends on frequency & length of patterns
- nibble based compression
  - 2x compression for strings with <=16 unique characters
  - Works well for phone numbers, dates, times, timestamps, etc

- Db2 will choose the best compression method.
- Compressed data is in the buffer pool and on-disk.
- Both of these **are off by default** and require a registry variable to enable. Once set, any tables using these new compression mechanisms will not be readable by previous mod packs of Db2 11.5.

Storage -- BLU Compression IMPROVEMENTS 2/2

Internal Lab Results based on customer-like data.

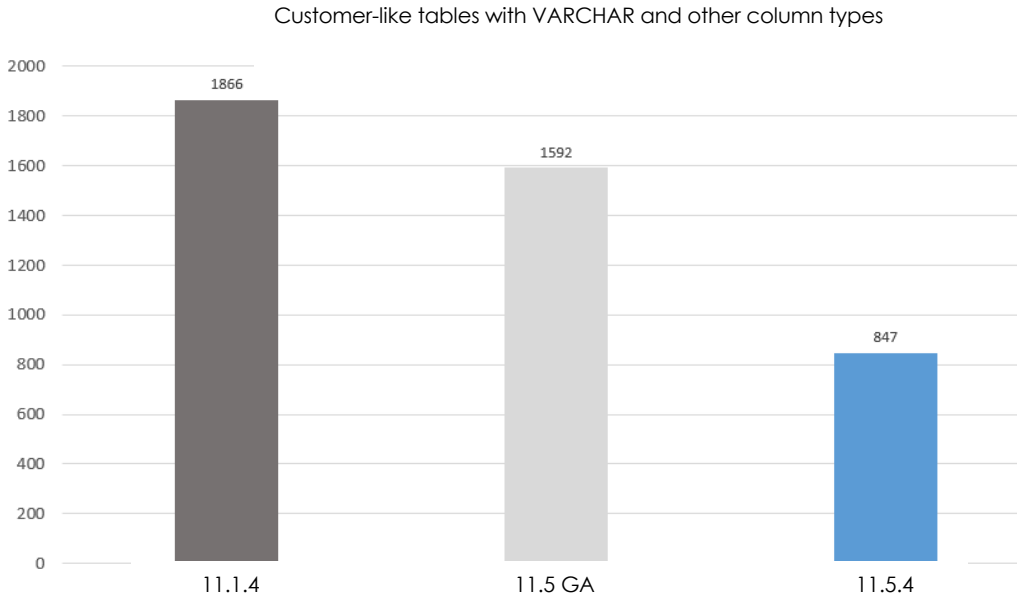

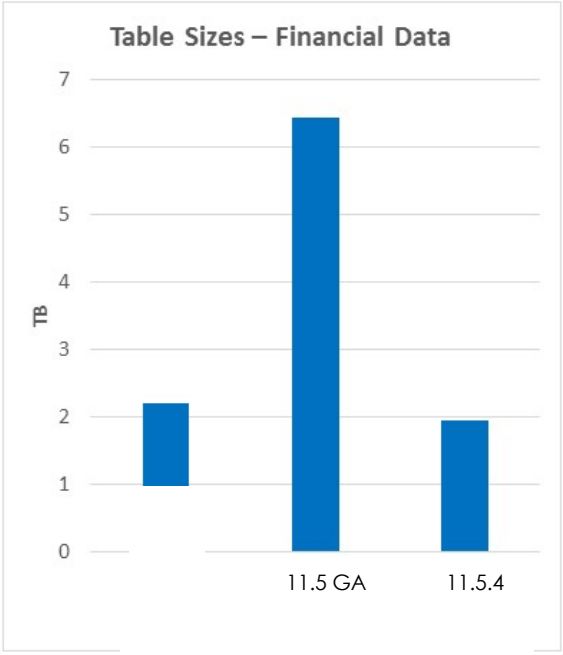

Synopsis tuple deferral

- Decreases the storage used by small tables.
- What is a synopsis table?
  - A synopsis table is a column-organized table that is automatically created and maintained by the system to store metadata for an associated user-defined column-organized table.
  - Each row of the synopsis table stores the minimum and maximum values for each column across a range of rows in the base table.
  - Used by queries to skip over data that is of no interest to a query during evaluation of certain type of predicates.
- Problem:
  - Synopsis tables for small column-organized tables can have excessive storage overhead due to Db2's partitioning and storage allocation.
- Solution:
  - The excessive storage consumption by synopsis tables can be avoided without performance penalty by deferring the creation of synopsis tuples until storage consumption overhead can be minimized.
- This feature is disabled by default because it is not fallback compatible, and it must be enabled explicitly. For details, see the DB2_COL_SYNOPSIS_SETTINGS registry variable.

Synopsis tables

Synopsis table

C1MIN	C1MAX	C2MIN	C2MAX	TSNMIN	TSNMAX
1	4	'A'	'E'	0	4
4	8	'F'	'J'	5	9

Base table

C1	C2
1	'A'
3	'B'
2	'C'
4	'D'
3	'E'
5	'F'
4	'G'
6	'H'
7	'I'
8	'J'

Synopsis tuple deferral

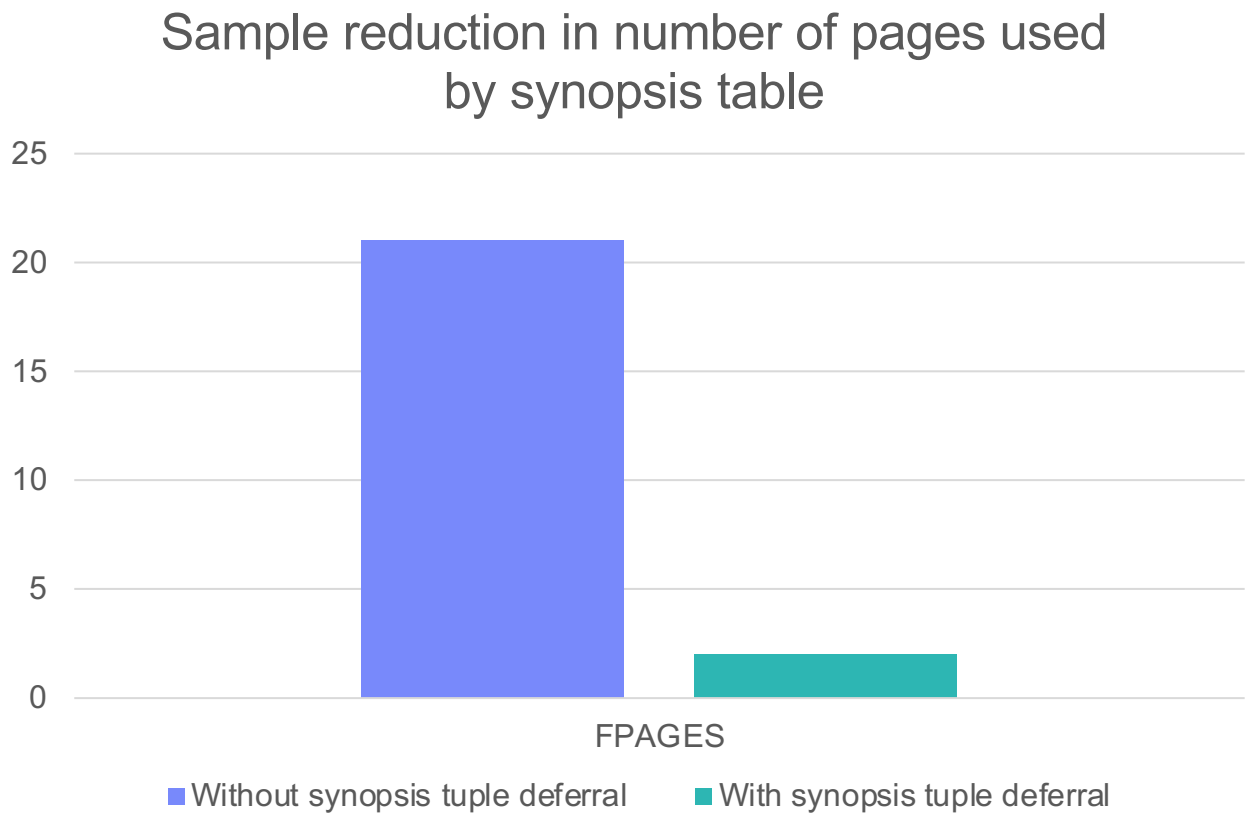

Base table with 1 integer column and 9000 rows

Gap Closing - Large Objects in Columnar Tables

Datatypes supported CLOB, BLOB, DBCLOB, NCLOB

Columnar database engine consumers can now avoid storing Large objects in row store tables.

Descriptor and in-lined data that is small enough is stored on the columnar data page.
Larger objects are stored using the same buddy space as row tables.

Simplicity -- RID (row ID) Scalar Function for BLU & DPF

Already exists for ROW tables. Now supported for COLUMNAR and for DPF

For DPF, DBPARITIONNUM should used with it.

Can be used to map lock information or any TSN to a row:

Example, mon_get_locks shows a row lock on table SAPD01.BBNLTEST named: 0001020B00000000000000158

First you map the lock name to a TSNID:

```
select varchar(value,30) as TSNID from table(mon_format_lock_name('0001020B00000000000000158'))
as t where name = 'TSNID'

TSNID

------------------------------

1
```

Next, you take the TSNID it returned, and select the matching row data:

```
select * from SAPD01.BBNLTEST where RID() = 1

…row will be returned…
```

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5 and later
- Db2 information resources and resources for developers

Security - pureScale Support for host-based firewalls 1/2

- In a multi-node environment, Db2 will need to run commands on remote nodes for certain operations.
- Host-based firewalls are an option for single-node, DPF, and now pureScale configurations.
- db2set DB2_FIREWALL_PORT_RANGE=<startPort>-<endPort>
  - Minimum of 2000 ports need to be specified.
  - The port range chosen must be within the range of non-privileged ports (1024 - 65535).
- Ensures that this cross-node communication will be performed using ports in the range specified.
- Used for all Db2 activities, from before install, throughout deployment, and during production.

Security - pure Scale Support for host-based firewalls 2/2

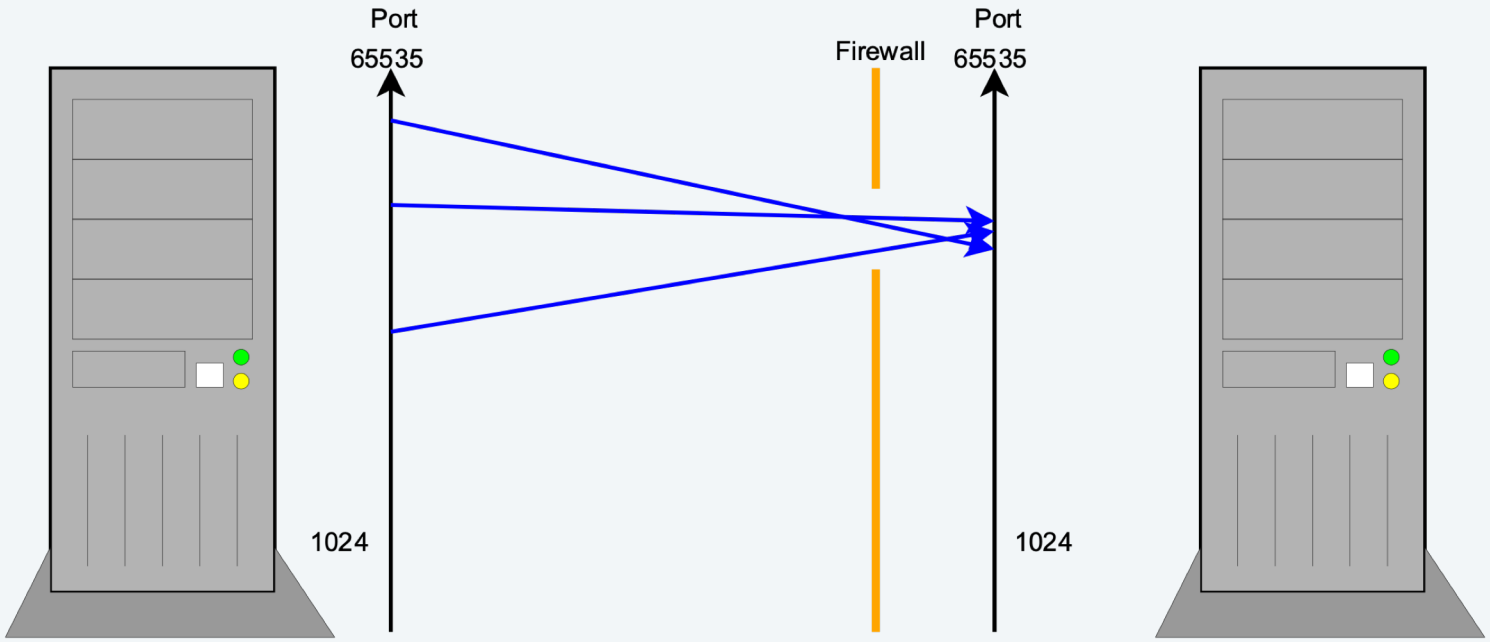

Simplicity - pS cluster wide free space management

- For tables that have the default APPEND OFF option set, re-use of free space on data pages has been improved.
  - APPEND OFF (the default) is specified on ALTER TABLE.
  - It means inserted data is placed in available (free) space.
- Reuse free space especially when the majority of DELETE operations are performed on a distinct subset of members in the cluster.
- Free space indicator has been moved to the cluster cache facility (CF)
  - Now all the members can check for new free space efficiently and reliably.
- This is the new default behavior starting in Db2 v11.5 GA.
- The following registry variable can be set to revert back to old behavior of free space searching:

  db2set DB2_SD_DISABLE_GLOBAL_FSCR_SEARCH=TRUE

Gap Closing - Db2 pure Scale Cross-member Currently Committed

Using cursor stability (CS) isolation, under **currently committed** semantics only committed data is returned to readers, but …

- Readers do not wait for writers to release locks.
- Instead readers return data that was committed before the write started.

Before this enhancement:

- In pureScale environments, the reader avoids a lock wait when:
  - The row being INSERTed by the other application resides on any member.
  - The row being UPDATEd or DELETEd by the other application *resides on the same member.*
- In other words … if the row being updated or deleted (locked) resides on a different member, the reader will still wait.

With this enhancement:

- Now the reader will go to the other member to return the currently committed data from the other member's logs.

Performance- 2x faster pureScale LOAD of range-partitioned tables

Before LOAD starts, it requires that target table pages first get flushed/purged from the buffer pool.

Prior to v11.5, the flush/purges were done individually for each partition, requiring multiple remote procedures calls (RPCs) to each member.

Starting with v11.5, the flushes/purges are grouped by partitions, indexes, and LOB columns.

Simplicity - Automatic config of public network monitoring

Manual commands for GDPC	V11.1	V11.5
After typical install	33	0
After every maintenance operation	6	0

Automatically create public ethernet condition response pair for GDPC (during setup, repair and enter/exit maintenance)

GDPC = Geographically Dispersed pureScale Cluster

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5 and later
- Db2 information resources and resources for developers

Simplicity-- Advanced Log Space Management 1/2

As part of a focus on simplifying log management, the first delivery will result from a focus on avoiding "log full" scenarios caused by long running, low volume transactions hold up log space

E.g. transactions that span multiple log files with little content in the intermediate files

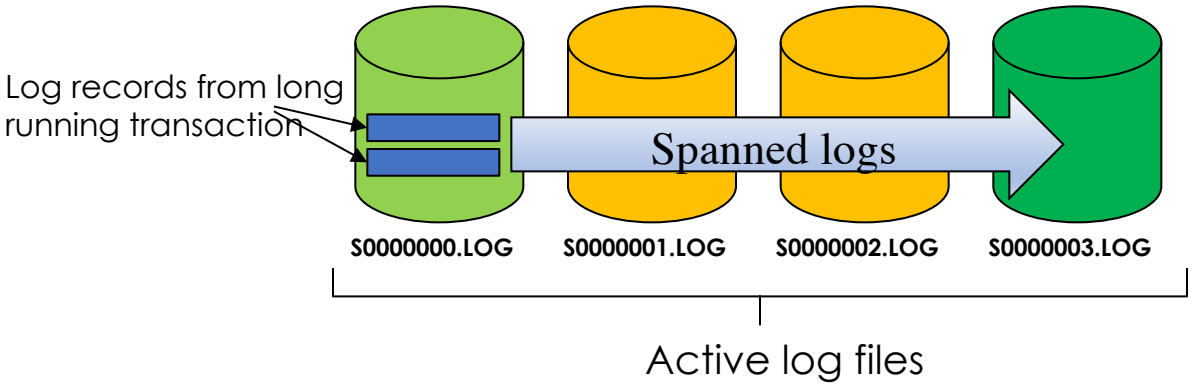

Simplicity-- Advanced Log Space Management 2/2

Solution is to extract log records for long running active transactions to a separate file and allow intermediate log files to be closed, archived, and reused

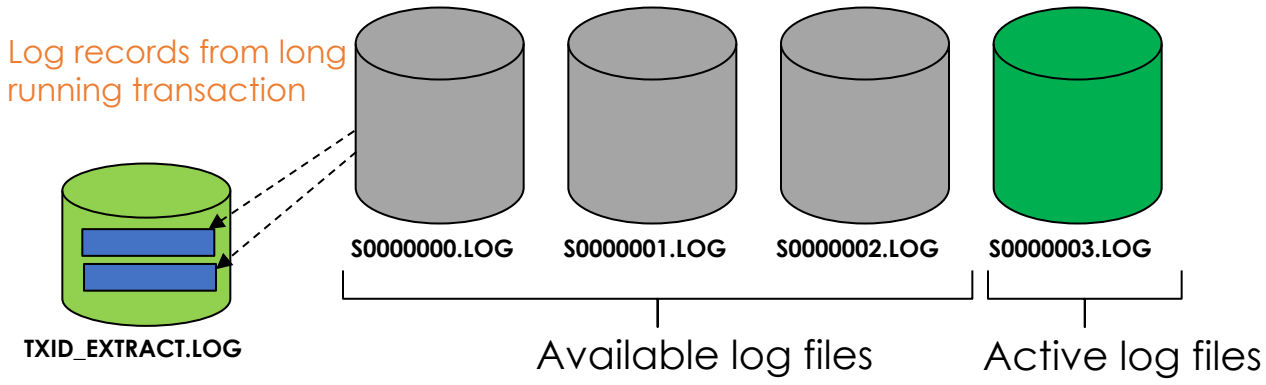

Adaptive Workload Management

- The adaptive workload manager automatically controls the admission of queries for execution to ensure that concurrently executing queries do not overwhelm the available resources.
- Adaptive WLM's objective is to
  - Deliver effective, automatic workload management by default with zero tuning.
  - Make sure that the system is well-utilized but does not become overcommitted.
  - Schedule jobs to ensure fairness and appropriate responsiveness.
- This new technology is being used today in:
  - Db2 Warehouse
  - Db2 Warehouse on Cloud
  - IBM Integrated Analytics System (IIAS)
  - Plan to add supported in future for more row-centric workloads after further validation and tuning.

Adaptive Workload Management

- The recipe is a mixture of **admission control** and **dynamic resource management** modified by **intelligent job scheduling** with a dash of **historical feedback**
- Admission control is based on whether the estimated query resource "footprint" will fit within the remaining system resources.
  - Resources considered: sort memory, CPU, number of threads.
  - The query is allowed to run if it fits, otherwise it waits.
- Incoming work guided to different "lanes" based on expected resource consumption and duration.
  - Each "lane" gets a defined resource allotment to maximize predictability.
  - Scheduling of work in each "lane" is based on dynamic view of resource availability.
- Includes historical feedback based on past executions.

Under the Hood: Managing Mixed Workloads for Predictable Performance

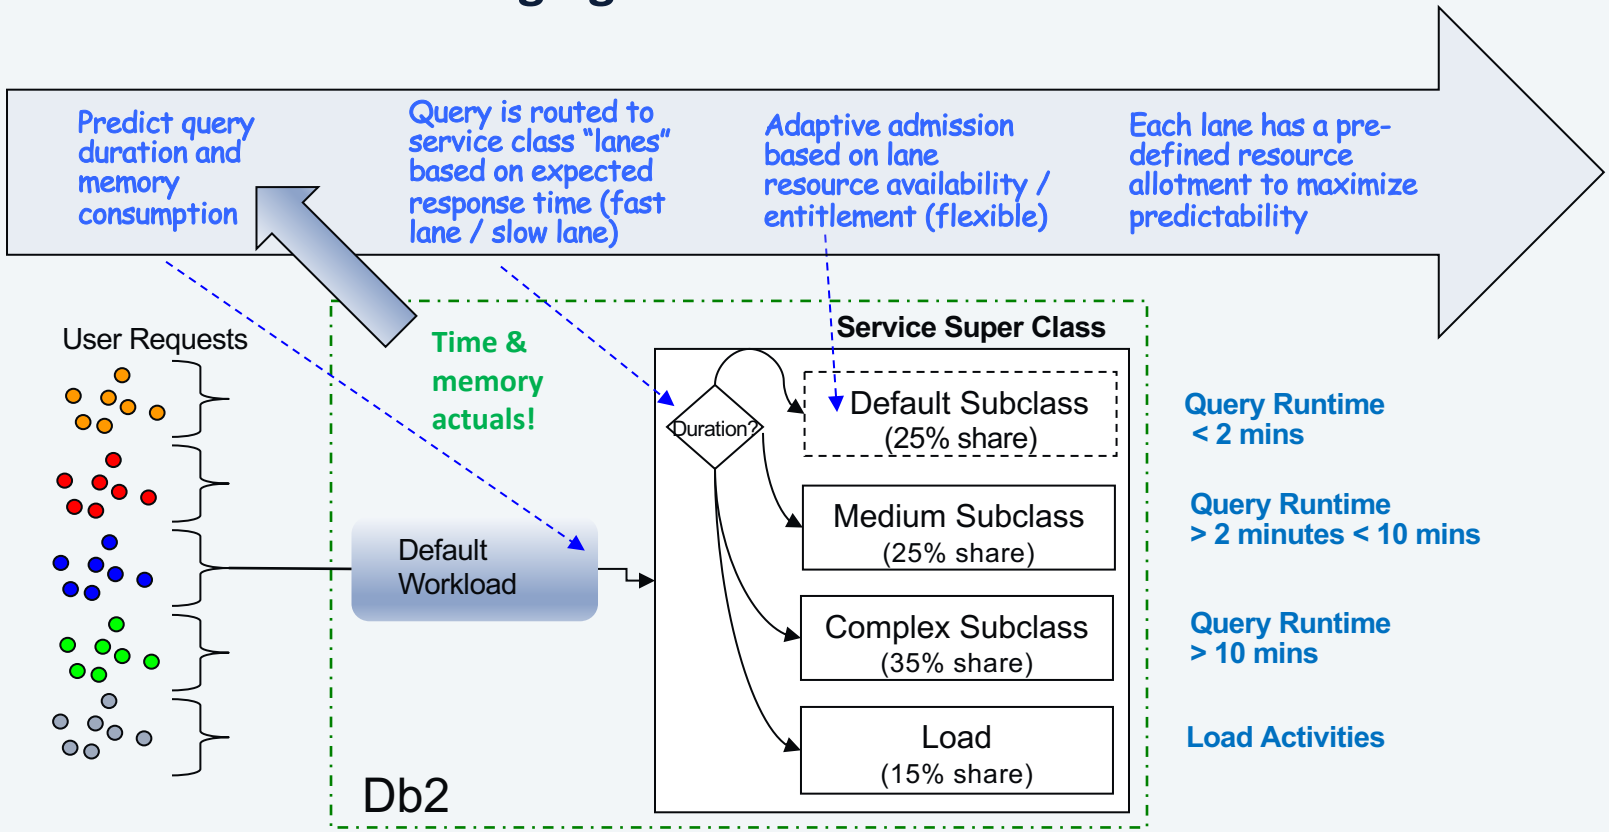

Under the Hood: Latency Oriented Job Scheduling

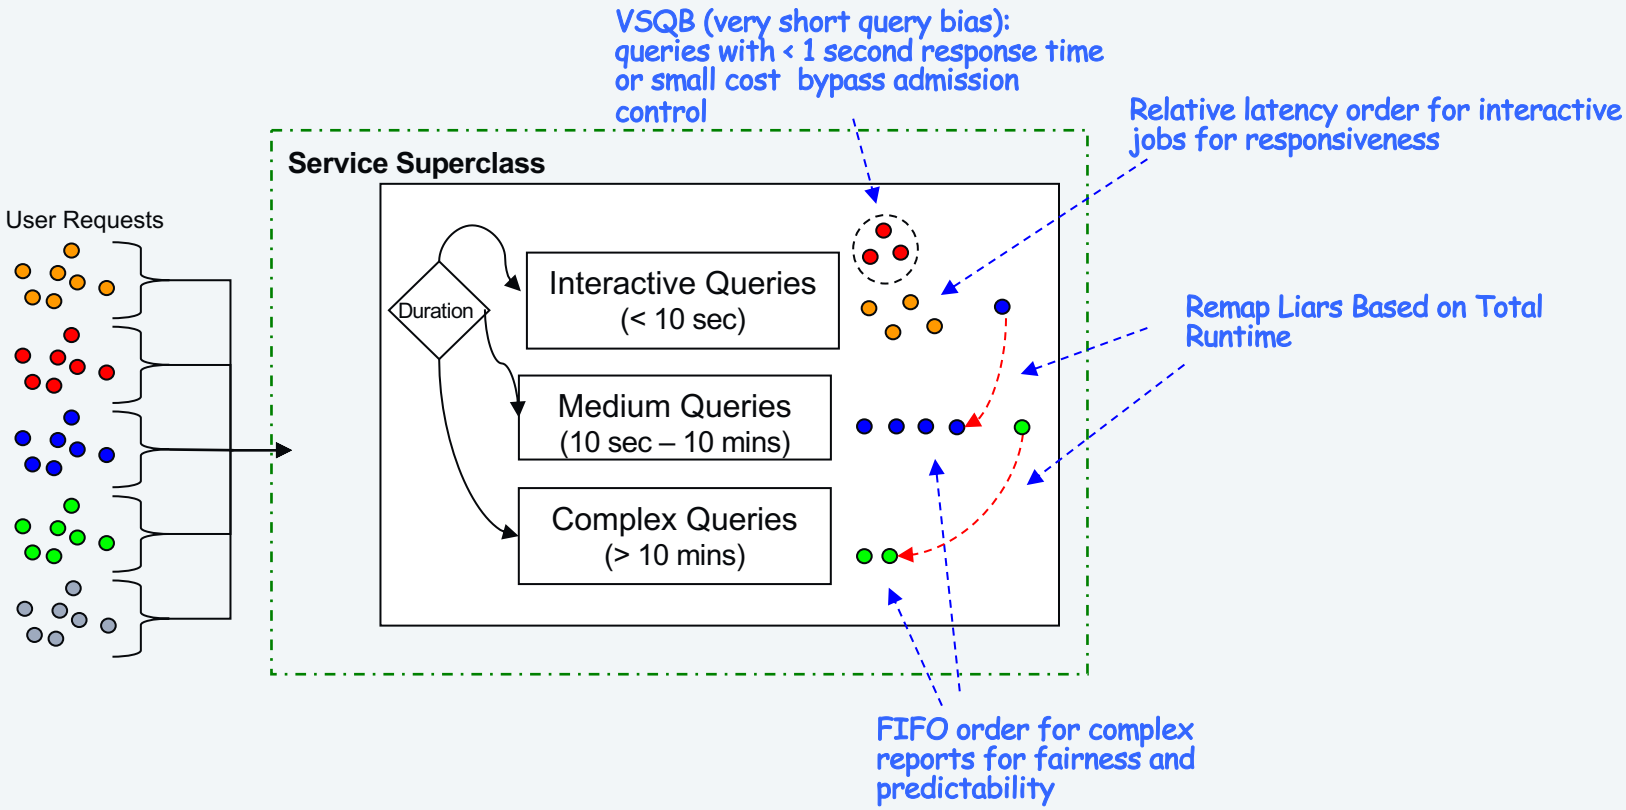

Support for storage devices using 4K sector size

This support can be enabled by setting the DB2_4K_DEVICE_SUPPORT registry variable to ON (default is OFF)

Allows Db2 databases to be created on devices that allow only 4K IOs.

Restrictions and limitations:

- The use of database managed (DMS) raw containers is not supported.
- Backup and load copy files will be slightly larger.
- There may be a small performance penalty accessing:
  - Small non-inlined LOBs
  - Backup or load copy files created prior to the enablement of 4K device support.

SQL Compatibility - External Tables

A simple mechanism that lets you access external "files" within an SQL statement just like a table

- Similar to federating on flat files (but allow write access).
- Can also be used to load from or unload to external files.
- Can be used to define a permanent external table or directly within a SQL statement.

Currently supports CSV/delimited files and local/remote sources

Example:

```
create external table ext_orders(order_num INT, order_dt TIMESTAMP)
                                USING(dataobject('/tmp/order.tbl') DELIMITER '|');

insert into orders (select * from ext_orders);
```

Example:

```
insert into orders (select * from external '/tmp/orders.txt' using(REMOUTESOURCE
GZIP delimiter ','));
```

SQL Compatibility - DROP/CREATE TABLE [IF EXISTS]

DROP TABLE IF EXISTS myTable

- IF EXISTS clause can be added to suppress SQL0204N errors
- 0 success returned if table exist or not

CREATE TABLE IF NOT EXISTS myTable (C1 int)

- IF NOT EXISTS clause add to suppress SQL0601N error
  - Warning “SQL4136W Table or view "SPRINGGA"."MYTABLE" already exists.” returned if table does exist
- Otherwise no impact if table exists
- Table created as normal if not exists

Security -- Authentication cache 1/3

- The authentication cache stores information about successful authentications and compares the information from new, incoming authentication requests against the cached entries to see if a valid match is found.
- Intended to relieve performance impact due to bottleneck on authentication backend
  - Workloads with short duration connections that occur repeatedly using the same, limited set of authorization IDs
- Applies to all password authentication plugins supported by Db2.
- Cache is configured based on maximum number of users to be cached and duration of time for cache (default is 3 minutes)
  - Credentials are not cached.

Security -- Authentication cache 2/3

Without a cache – every connection needs to call out to LDAP Server

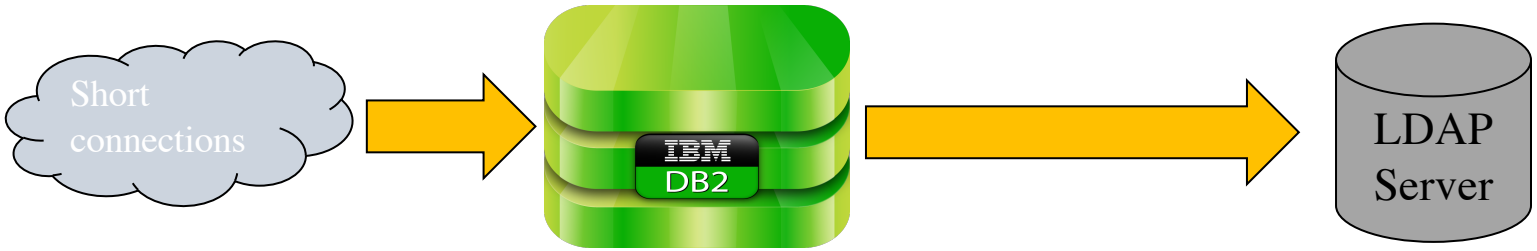

Security -- Authentication cache 2/3

With a cache – only a small subset need to, the rest are serviced by the cache.

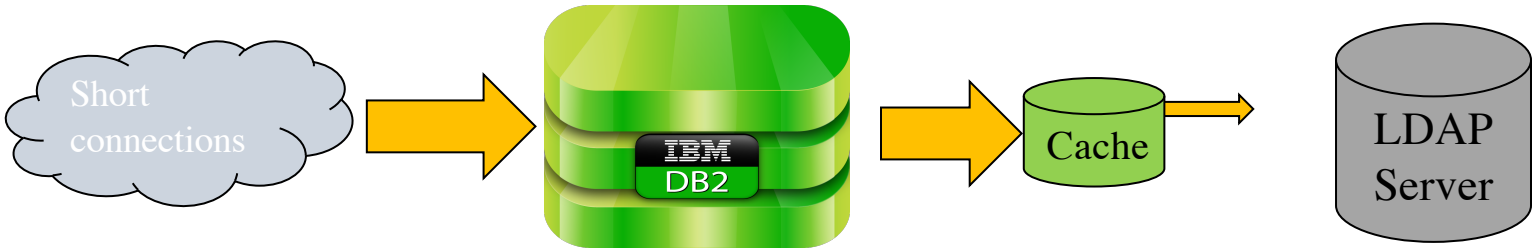

The cache exists at each database member that receives CONNECT requests

The cache contents are independent of the contents of the cache at any other database member.

Security -- Authentication cache 3/3

Controlled by 2 database configuration parameters:

AUTHN_CACHE_USERS –

- Number of entries to be kept in the Db2 Authentication cache.
- Default 0 (== OFF)

AUTHN_CACHE_DURATION –

- Time in minutes for which an entry is considered valid and available for reuse in the Db2 Authentication cache.
- Default 3min

New SQL statement, to allow SECADM and DBADM to clear the cache:

- FLUSH AUTHENTICATION CACHE

GET_MON_DB also updated to be able to monitor the efficiency of the cache

SIMPLICITY-- OPTIMIZER VERSION CONTROL 1/2

- The Db2 optimizer uses some heuristics.
- Occasionally changes to the optimizer made in a fix pack will cause an existing query to not perform as well as it under the old optimizer.
- Previously, in order to revert the behavior of the optimizer, a large number of both documented and undocumented reg vars had to be set, often to cryptic values.
- Now a DBA can quickly revert the optimizer behavior in terms of:
  - Query rewrites/transformations
  - Access plan generation

  to match a previous version/release.
- Single control that can quickly be used in emergency situations where a regression may be experienced after an upgrade or update. Can also be used pro-actively.
- It is not recommended to be run with long-term.

SIMPLICITY-- OPTIMIZER VERSION CONTROL 2/2

New registry variable:

DB2_OPTIMIZER_VERSION

Default = current version

Can be set to any 4 part release name since 10.5.0.0

E.g 11.5.0.0 to match 11.5 GA

Does not prevent new plans due to change in statistics or other non-optimizer code changes.

This is applied first, followed by any other SQL compiler registry variables which can then override its underlying settings.

SIMPLICITY-- BLOCK OPERATIONS THAT RESULT IN REORG PENDING 1/2

- Some ALTER TABLE statements can put a table into reorg pending state.
- In this state, many types of queries cannot be run.
  - You must perform a table reorganization before the table becomes available

Goal:

- Prevent a user accidentally running an ALTER TABLE statements that puts a table into reorg pending and affect subsequent access to table (until the reorg is run).

List of operations that place table into reorg pending:

- DROP COLUMN
- ALTER COLUMN SET NOT NULL
- ALTER COLUMN DROP NOT NULL
- ALTER COLUMN SET DATA TYPE except in the following situations:
  - Increase VARCHAR/VARGRAPHIC length
  - Decrease a non-indexed VARCHAR/VARGRAPHIC length, without truncating trailing blank

SIMPLICITY-- BLOCK OPERATIONS THAT RESULT IN REORG PENDING 2/2

New registry variable:

DB2_BLOCK_REORG_PENDING

Default OFF

When set to ON, this prevents “ALTER TABLE” operations that would put the table into reorg pending state (and limit activity to read-only table scan-only access).

ALTER table statements that would result in reorg pending state will fail with SQL0270N reason code=129

Availability -- FASTER Database STARTUP

ON by DEFAULT

Speeds database startup by
up to 120x.

The larger the buffer pools and lock list the larger the speed-up.

Implemented by deferred memory commit, which improves the performance of large memory allocations, such as those that occur during database activation.

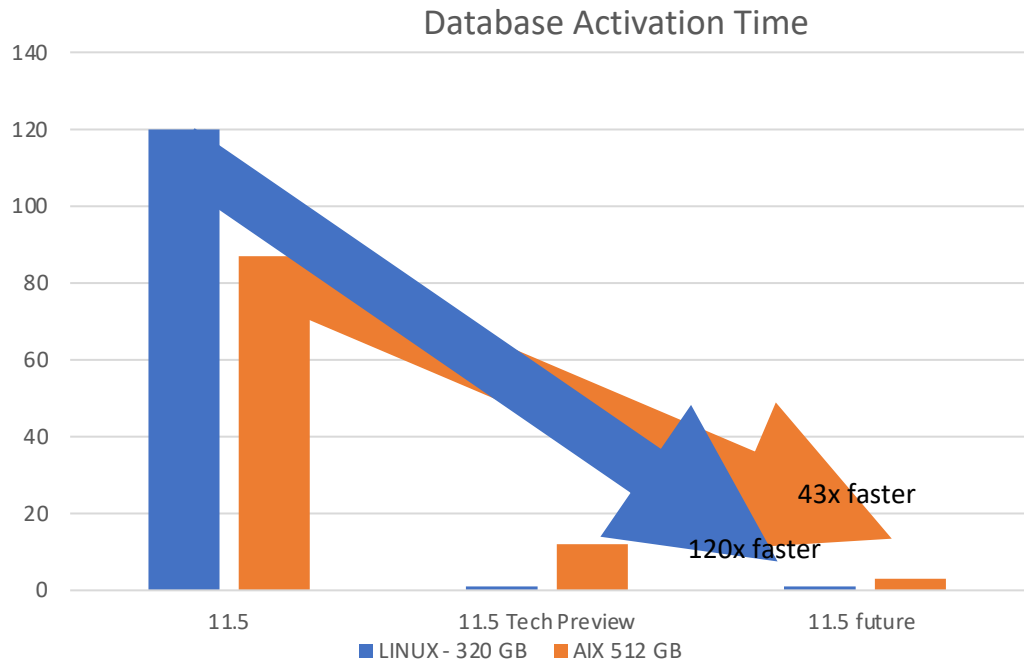

Modern Dev - Spatial Analytics 1/3

- Spatial support for analytic applications using BLU technology
  - Supports both row and column-organized tables
- Similar functionality as existing Db2 Spatial Extenders
- Some external and architectural differences
- Tech Preview in Db2 11.5 GA, GA in 11.5.4

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5
- Db2 information resources and resources for developers

pureScale support!

pure Scale

- pure Scale support for Power9 native
- Faster pure Scale online modpack & fixpack update (concurrent)
- Automatic Cluster Validation (periodic)
- Lightspeed Remote Direct Memory Access (RDMA) ping

Performance – Columnar Trickle Insert

Goals:

Speed up “trickle” inserts and

reduce memory footprint.

Decrease size of very small tables.

The idea is to get data in as fast as possible (in unencoded row format), then go back to split it up and encode it.

This will be off by default when it ships, user will have to turn on as any table created or inserted into after it is enabled can no longer be accessed by previous level (no fallback support).

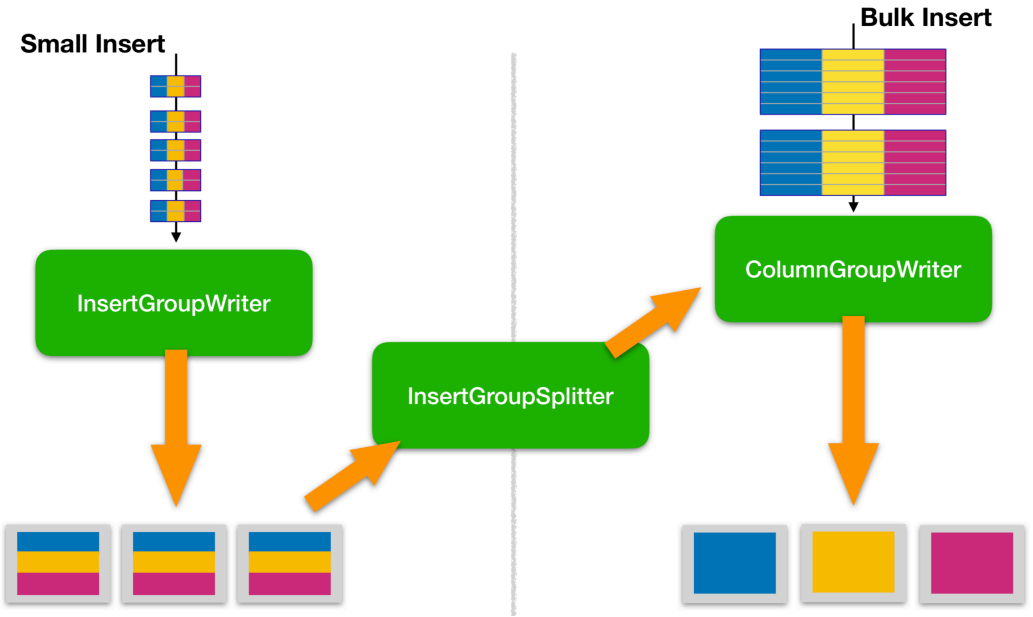

Compatibility – Multi-tenancy

User-defined tenants

- In a **multi-tenant database** each tenant's data is isolated and remains invisible to other tenants.
- A DBADM can create a tenant to set up an independent catalog namespace within a Db2 database.
- All connections made to a database are initially associated with the SYSTEM tenant.
- A SET TENANT statement must be issued to associate a connection with a user-defined tenant .

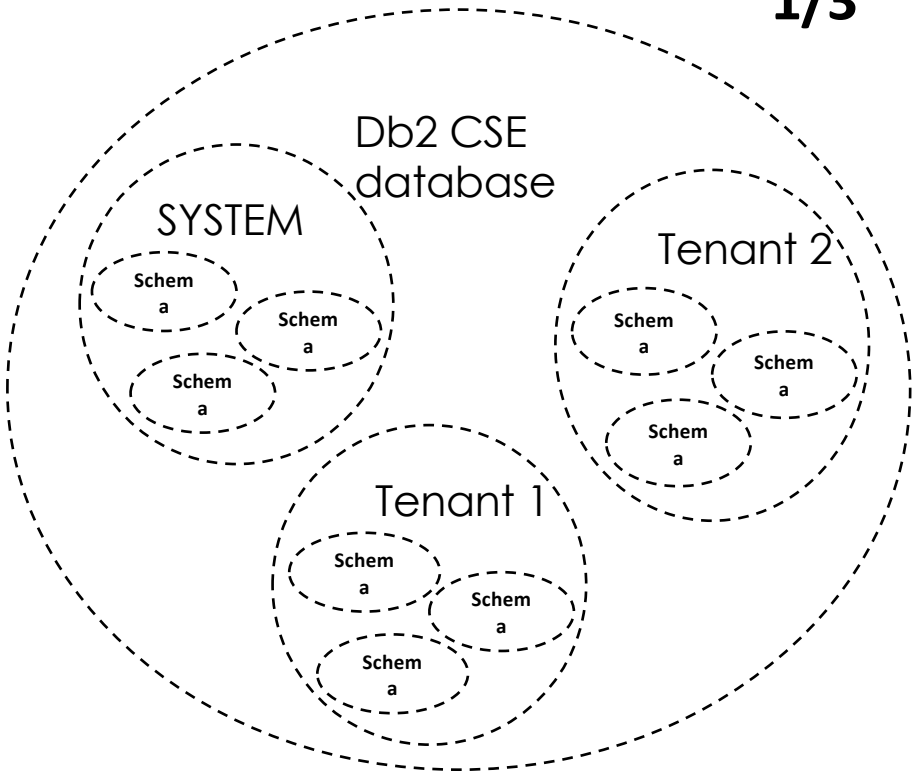

Compatibility – Multi-tenancy 2/3

A simple example (setup)

```
CREATE DATABASE TEST

CONNECT TO TEST

CREATE TENANT WORLD1
GRANT USAGE ON TENANT WORLD1 TO USER1

CREATE TENANT WORLD2
GRANT USAGE ON TENANT WORLD2 TO USER2
```

Compatibility – Multi-tenancy 3/3

A simple example

```
CONNECT TO TEST USING USER1

SET CURRENT TENANT = WORLD1

CREATE TABLE MINE.T1 (C1 INT)

INSERT INTO MINE.T1 VALUES (1)

SELECT * FROM MINE.T1
    -> 1
```

```
CONNECT TO TEST USING USER2

SET CATALOG WORLD2

CREATE TABLE MINE.T1 (C1 CHAR(1))

INSERT INTO MINE.T1 VALUES ('A')

SELECT * FROM MINE.T1
    -> 'A'
```

AI - Db2, the only database w/ ML Optimizer in addition to Cost-based Optimizer

Machine Learning Optimizer that improves from experience to make query optimization simple, reliable and stable

- Number of rows flowing through the various operations impacts performance for most common issues and is calculated via cardinality estimates.
- Basic tuning to improve cardinality estimation is high impact
- **Phase 1 – Cardinality Estimation**
  - Initial phase – Support equality & range local operators with no expressions
  - Future phases – Cardinality support for expressions, predicates & BLU, support for join enumerations.

Agenda

- Meet the family of hybrid data management solutions built on Db2
- Release numbering
- Categories of enhancements
- A (long) list of enhancements in v11.5 through v11.5.4
- Closer look at some BLU enhancements
- Closer look at some pureScale enhancements
- Closer look at other selected enhancements
- A peek at v11.5.5 and later
- Db2 information resources and resources for developers

IBM Db2 Resources

Information Resources:

- Db2 Roadmap - http://ibm.biz/AnalyticsRoadmaps
- Db2 RFE (Idea) Portal - http://ibm.biz/submitdb2idea
- Get latest Info- Subscribe to Db2 technical newsletter - http://ibm.biz/db2nlsignup
- Connect with the Db2 online community - http://ibm.biz/db2tribe
- Stay up to date with all the new demos on Youtube - http://ibm.biz/db2-youtube
- Interested in getting the latest beta version - http://ibm.biz/getdb2beta

Developer Resources:

- Db2 Developer Page to get started - http://ibm.biz/db2developer
- For Experienced Db2 developers, get your fav Db2 code sample on github - http://ibm.biz/db2github
- Want to try Machine Learning with Db2, check out - http://ibm.biz/learndb2
- Want details on Db2 Python Driver - http://ibm.biz/db2-drivers-python
- Want Details on Db2 PHP Driver - http://ibm.biz/db2-drivers-php
- Want Details on Db2 Node.js Driver - http://ibm.biz/db2-drivers-node
- Download the free Db2 python e-book - http://ibm.biz/db2pythonbook

IBM Cloud and Cognitive Software

Extra slides

October 8, 2020

Walter Alvey
Data Storage, IBM Analytics

© 2020 IBM Corporation

Compact Varchar Phase 2

Impact

- Improved memory efficiency for wide VARCHARs in CDE Group By and Join queries
- Reduce memory consumption, spill I/O and OOMs primary impact focus
- Performance improvements
- Increase in concurrency within Group By and Join operator

Results

- Overall workload elapse time, memory footprint and spilling greatly improved
  - Performance: Up to 2.9X overall workload, 17.6X individual query
  - Memory reduction: Up to 1.1X overall workload, 2.5X individual query
  - Spilling reduction: Up to 5.6X overall workload, >1200X individual query

Compact Varchar Performance Metrics

Metric	DS Augmented for Group By	DS Augmented for Join	CVC2 Targeted for Group BY	CVC2 Targeted for Join
Number of queries	99	42	360	42
Number improved (5% or more)	11	14	171	30
Number regressed (5% or more)	5	2	7	4
Performance				
Workload: overall workload improvement	1.8X	3X	1.2X	1.3X
Single query: most improved	5.1X	4.7X	11.7X	7.7X
Memory Reductions				
Workload: overall workload memory usage	1.0X	1.1X	1.1X	1.1X
Single query: largest memory reduction	2.0X	1.3X	1.9X	2.2X
Spilling Reductions				
Workload: overall reduced spilling	1.5X	1.6X	5.6X	1.4X
Single query: best spilling improvement	342X	4.6X	>1200X	3.9X

*AI – ADE tool

Get to the insights faster – Use Natural Language Query (NLQ) for data analysis

Conversational analytics is the future…

- Data scientists often face the challenge of exploring data sets that are not familiar to them
  - Time consuming, Not sure what to look for
  - Need understanding of what is in the data set
- Need an intuitive data exploration tool
  - Alexa, Google Home, etc. prove trend
  - Natural extension of conversation to data
  - Provide multiple views of the data with minimal touch
- **Augmented Data Explorer** available to download - https://www.ibm.com/us-en/marketplace/db2-augmented-data-explorer
  - Seamlessly plugs and plays with Db2

*Developer - Programming IDE

- Bridging the gap between developers and Db2 application development
- Install Visual Studio Code and install Db2 Extension (link: ibm.co/2wvxcGs)
- Connect with Db2 and test your application

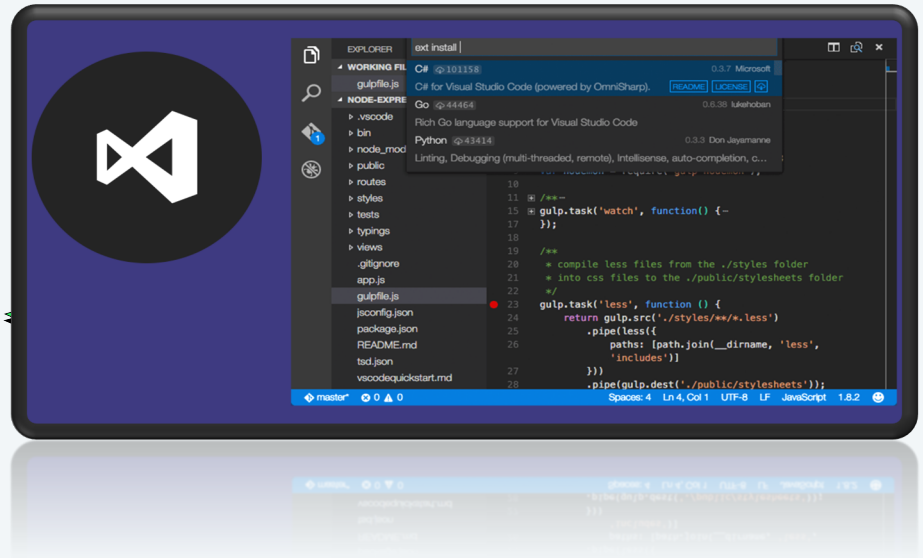

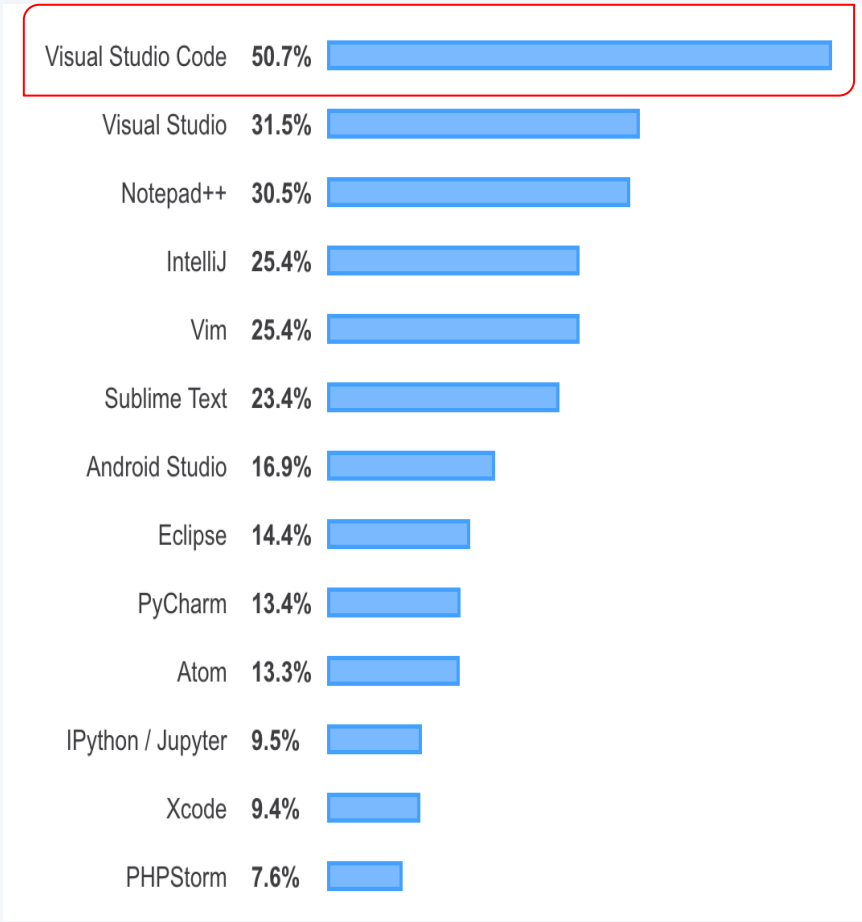

*Developer - GoLang for Db2

- Now available @ https://github.com/ibmdb/go_ibm_db

How to run sample program

example1.go:-

```
package main

import (
    _ "github.com/ibmdb/go_ibm_db"
    "database/sql"
    "fmt"
)

func main(){
    con:="HOSTNAME=host;DATABASE=name;PORT=number;UID=username;PWD=password"
 db, err:=sql.Open("go_ibm_db", con)
```

- https://github.com/ibmdb/ - contains 8 open source drivers (Go, Python, Jupyter Notebook, Sequelize, PHP, Java (Spring), Node.js, Ruby)

Modern Dev - Spatial Analytics 2/3

- Dedicated data type to hold shapes
  - LOB-based datatype SYSIBM.ST_GEOMETRY allows storing large geometries
  - Dedicated subtypes for points, linestrings, polygons, etc
- Pre-loaded spatial catalog data with SQL procedures for customization
  - E.g. add custom coordinate systems
- SQL functions based on the SQL/MM and OGC standards:
  - Construct and maintain/modify shapes
  - Determine relations between shapes
  - Get properties
    - `SELECT ST_DISTANCE(ST_CENTROID(GEOM1), GEOM2) FROM TAB1, TAB2`
- Same feature/interfaces for column and row-organized tables
- Enabled via SYSINSTALLOBJECTS procedure

Modern Dev - Spatial Analytics 3/3

	Spatial Extender	Spatial Analytics
Processing Method	In-Database	In-Database
Data Organization	Row-Store	Column-Store Row-Store
Index / Filter Type	Spatial Grid	N/A *
Spatial Joins	Yes	Yes
Function Type	Planar (with few exceptions)	Planar (with few exceptions)
Support for custom Coordinate Systems	Yes	Yes
Support for Spatial Reference Systems	Yes, default = 0, undefined	Yes, default = 4326, WGS84
Maximum Shape Size (compressed)	4 MB	4 MB

WGS = World Geodetic System = standard for use in cartography, geodesy, satellite navigation
WGS84 = latest revision of the standard
4326 = spatial reference system ID

Db2 11.5: How Can I Install Db2?

Db2 family has 2 deployment (install) options:

1) containerized

2) multi-platform install

Multi-platform Installs:

Supports all Db2 platforms: Linux (Intel,Power,Z), AIX and Windows

Supports all Db2 configurations: single node, DPF, and pure Scale

Mod packs generally ship 6-12 months apart.

Between mod packs, there may be “fix-only” fix packs. These are similar to the concept of 11.1 ifixes but 2 differences:

1. Will contain all APARs not just a subset
2. Will bump up the FP field of the version. (Easy to recognize!)

Containerized Installs:

RedHat OpenShift (RHOS) based container.

Supports single node and DPF (MPP); does not yet support pure Scale.

Ships ~every 2-3 months (subject to change).

Currently always has new function*, so are mod packs.

- 11.5.1.0 , 11.5.2.0, 11.5.3.0 are already available
- these levels are also used to sync the Db2 Family

* This could change once there is a new continuous development release (no planned date).