Db2 NIGHT SHOW:

LATEST FROM THE LAB: WHAT'S NEXT FOR Db2 V11.5

Keri Romanufa– STSM and Chief Architect for Db2

SPECIAL GUEST

Adam Storm – Distinguished Engineer, Data & AI

And a few words on IBM, Db2 and the world.

AGENDA

Special Guest: Adam Storm

The Monologue: The new v11.5 numbering scheme

Interview: Db2 v11.5

Top 10: Deep Dive

Musical Performance

**and the occasional "commercial break and poll" by host Martin Hubel

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 :

•**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.

KERI'S SUMMARY:

I talk mostly about the future 1H 2020 roadmap in this presentation.

Things can and will change.

Nothing in here implies a commitment for something to be available in *any* time frame or even ever.

Things can and will change.

THE MONOLOGUE

Db2 V11.5: WHAT'S UP WITH THE NUMBERING?

11.5 follows a continuous delivery model (as did 11.1)

New function will increment Mod level (as did 11.1)

but it will also reset the FP level to 0 (*new*)

Fix-only levels will increment the FP level (as did 11.1)

This is addresses the issue that 11.1 "ifixes" were not easily distinguishable from the base mod pack nor each other.

Db2 family has 2 install options:

1) containerized
2) multi-platform install

CONTAINERIZED INSTALLS:

RHOS based container.

Supports single node and DPF (MPP); does not support pureScale.

Ships ~every 2 months.

Currently always have new function*, so are modpacks.

- v11.5.1.0 & v11.5.2.0 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).

Db2 V11.5 MULTI-PLATFORM INSTALLS:

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

Supports all Db2 configurations: single node, DPF, and pureScale

Modpacks ship at non-fixed intervals, generally 6-12 months apart.

Between modpacks, there may be fixpacks. 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!)

THE Db2 11.5 LIFECYCLE

Multi-platform install path -- modpacks and fixpacks

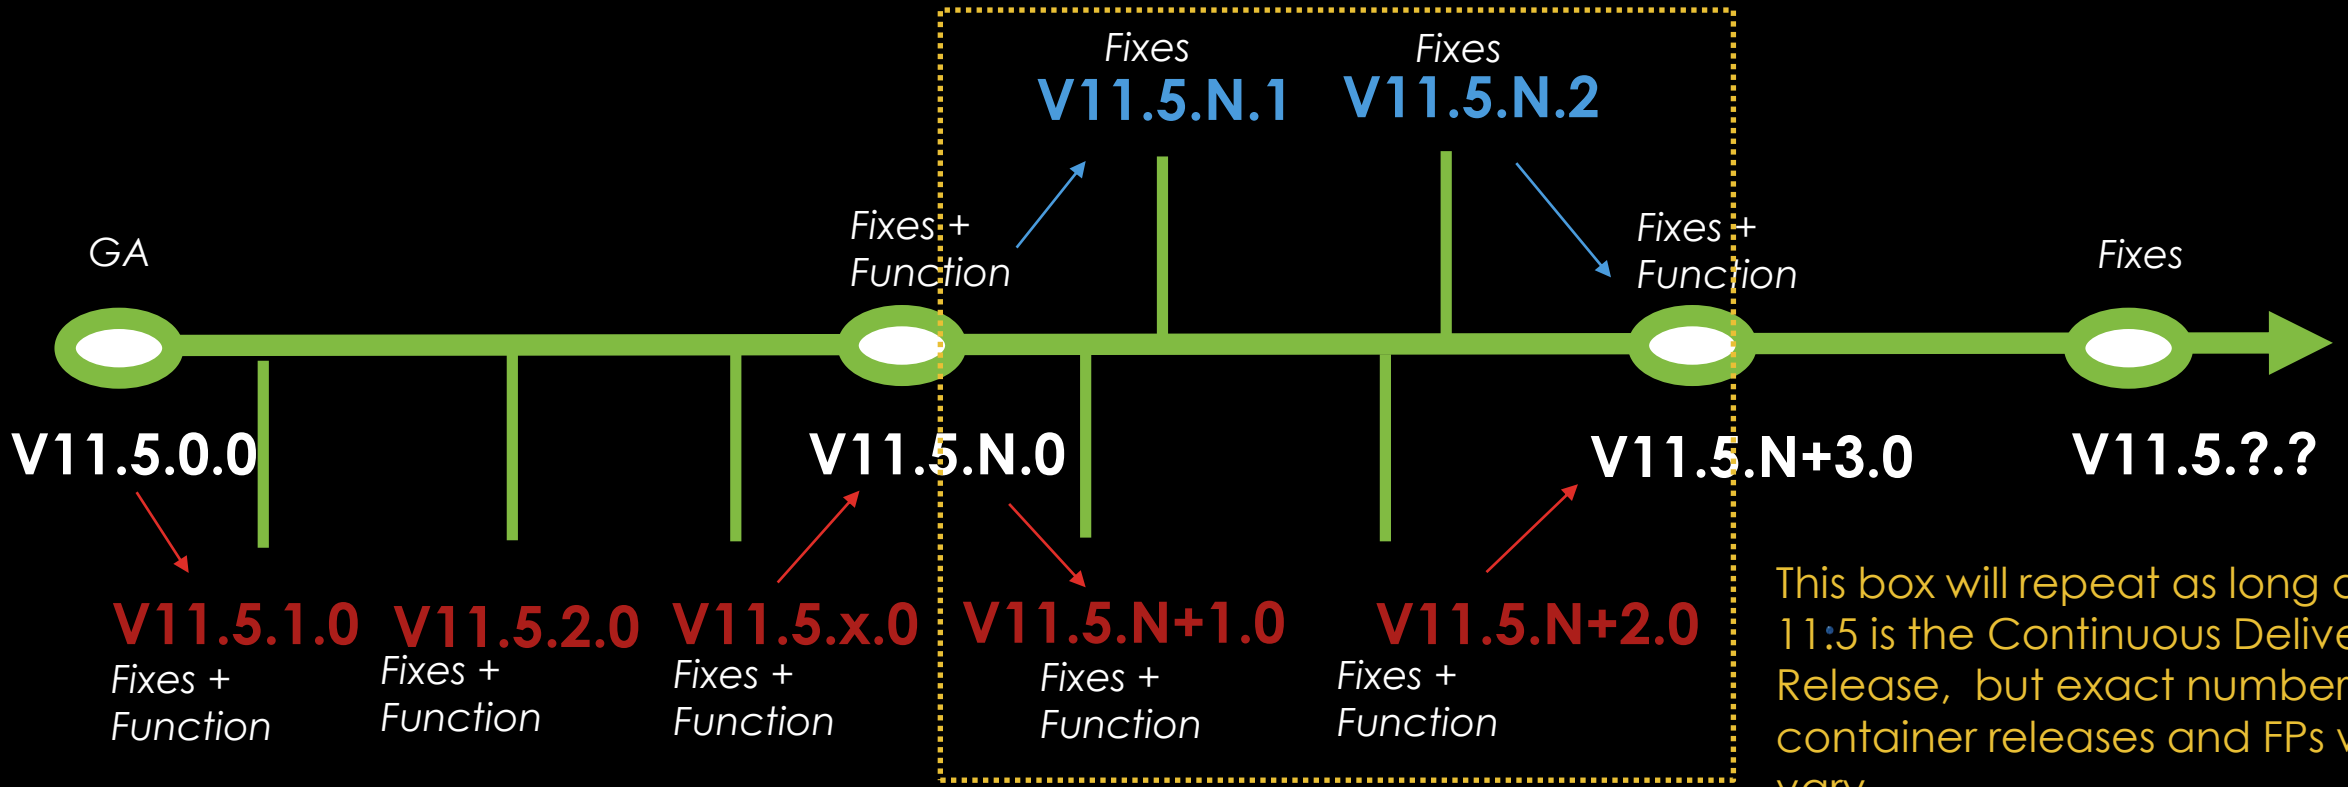

This box will repeat as long as 11:5 is the Continuous Delivery Release, but exact number of container releases and FPs will vary.

Containerized Install Path -- always* modpacks!

INTERVIEW WITH Db2 V11.5

So… Db2.. Tell me a bit about yourself.? What did you do last summer? What about this year; can you tell me what your plans are for the 1st half of 2020?

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***

Cloud

Db2 Warehouse on Cloud

—

Fully-managed, cloud data warehouse

Cloud

Db2 *on Cloud*

—

Fully-managed, cloud transactional data store

Cloud

Db2 *Hosted*

—

Not managed - we install Db2 and hand the keys over to you

Cloud

Hosted Analytics

—

Big SQL and Data Science Experience

Db2 **SQL Engine**

Db2

& Db2 Warehouse

—

Transactional or analytics SQL database deployed on commodity hardware

Db2 Cartridge *& Db2 Warehouse On CP4D*

--

Transactional or analytics SQL database deployed on commodity hardware

Integrated Analytics System

—

Dedicated analytics appliance

Db2 Event Store

—

Event-driven data processing and real-time analytics

Db2 Big SQL

—

Open source Hadoop with Cloudera

Db2 11.5 GA Highlights

Cognitive and Emerging

- ML Optimizer Tech Preview
- Federation support for Block Chain Tech Preview

Hybrid & Multi-cloud

Docker container

Application Development

- Net core 2.2 support
- GO language Driver
- IDE for Visual Studio Code – a db2 extension
- Augmented Data Explored (ADE) beta

BLU & MPP Enhancements

- Next Gen Bulk Insert for columnar tables
  - Vectorized Insert (and Update)
  - Reduced UNDO logging
- Update+Delete performance enhancements
- Automatic Dictionary Creation (ADC)
- Automatic REORG RECOMPRESS (compression of uncompressed data used to build dictionary
- Vectorized ADC
- Support for LOB data type in columnar tables
- Numerous improvements to columnar query performance.

Db2 11.5 GA Highlights

Availability

- Advanced Log Space Management Tech Preview
- Parallel Logging when using Mirrored Logs

pureScale

- pureScale Improved Cluster-wide free space management
- pureScale host-based firewall support
- Automate setup for public Ethernet monitoring
- pureScale increase ports, XI connection and worker limits
- pureScale OLIC and extent reclaim on by default
- Currently Committed across pureScale members
- pureScale 2x faster LOAD w/ Range Partitioned Tables

Core Engine

- External Table Support
- 4k Sector Support
- Auto Column Group Stats (CGS)
- Numerous additional monitor elements
- DROP/CREATE TABLE [IF EXISTS]
- Create Table As (CTAS) enhancements
- DBMS_APPLICATION_INFO and UTL_RAW package support (Oracle Compat)
- PLSQL Enhancements
- WLM:
  - Simplified thresholds
  - Cascaded drop service class

Db2 11.5 1H Roadmap Highlights

Graph database

Graph database + Relational database in one store.

Overlay a graph on top of Db2 data, and query Db2 via Gremlin language.

Hybrid & Multi-cloud

Db2 container for RH OpenShift

Application Development

- Drivers and NoSQL wire listener updates
- Db2 REST APIs
- Python UDFs
- Programming language updates:
  - .Net 2.2,
  - .Net drivers for MS Azure,
  - Node JS v12,
  - Ruby/Rails latest,
  - Django 2.0

BLU & MPP Enhancements

- Automatic INDOUBT Resolution for DPF
- Adaptive workload management for WORKLOAD=ANALYTICS
  - Session priorities
  - CPU Controls
- Columnar Trickle Insert
- Columnar Page based VARCHAR compression
- Delayed synopsis table population
- Query Perf (early aggregation, early distinct, full outer join, join residual predicate support, NULL=NULL)
- Memory Stability (e.g COMPACT VARCHAR, Full Outer Join)
- Truncate Table -- rollback support
- COUNT DISTINCT OLAP
- RID Scaler Function for Columnar/DPF
- Alter Table add column support for LOB

Roadmap: https://ibm-analytics-roadmaps.mybluemix.net/

Db2 11.5 1H Roadmap Highlights

Availability

- Automated HADR with Pacemaker for RHEL- Tech Preview
- pureScale support for Power9 native
- Advanced Log Space Management
- Faster pureScale online modpack& fixpack update
- Ability to block reorg pending operations through registry variable
- New column in MON_GET_HADR HADR_LAST_TAKEOVER_TIME

Core Engine

- Optimizer Version Control
- Faster Database Activation
- Skipped locked Data For Queries
- Nested WITH support
- ALIAS support in WHERE clause
- Better Cleanup of temporary files used for LOBs during a Remote Load from Cursor
- Faster Index Splitting at non-leaf levels under high contention
- UTL_RAW, DBMS_LOCK, DBMS_STAT package support (Oracle Compat)

Security

- SSO with JWT Token
- Authentication Caching
- Security-Enhanced (SE) Linux support on RHEL 7 & 8
- Allow SSL_SVR_LABEL to be changed online

Cognitive and Emerging

- ML Optimizer - Tech Preview 2
- In Database Analytic functions
- Spatial Analytics
- Federation:
  - Parallelism
  - Block Chain – GA
  - Numerous JDBC connectors

TOP 10 DEEP DIVE

Wow, so many things to choose from. How'd you get to a list of "Top 10",

Skipped the ones Piotr will cover in May such as RESTAPI, GRAPH, ML Opt, In DB analytics, Python UDF< containers, etc.

Skipped tech preview items, still ran out of room!

SECURITY -- AUTHENTICATION CACHE

- Intended to relieve performance impact due to bottleneck on authentication backend
  - Workloads with extremely short duration connections which 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

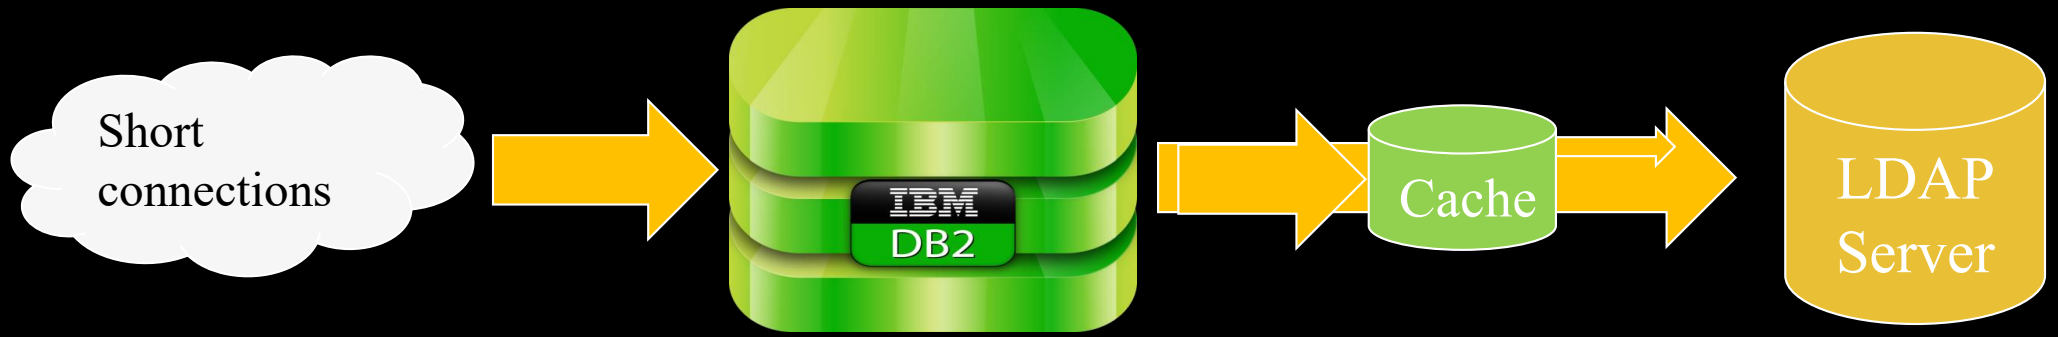

SECURITY -- AUTHENTICATION CACHE

Controlled by 2 registry variable:

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

STORAGE -- BLU COMPRESSION IMPROVEMENTS

2 *new* page-level compression methods for string data:

- 'repeating pattern'-based compression (LZ4 based)
  - Work's 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, coordinates, $ amounts, etc

Stored compressed on disk and in buffer pool

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 modpacks of Db2 v11.5

LAB RESULTS:

BLU COMPRESSION IMPROVEMENTS

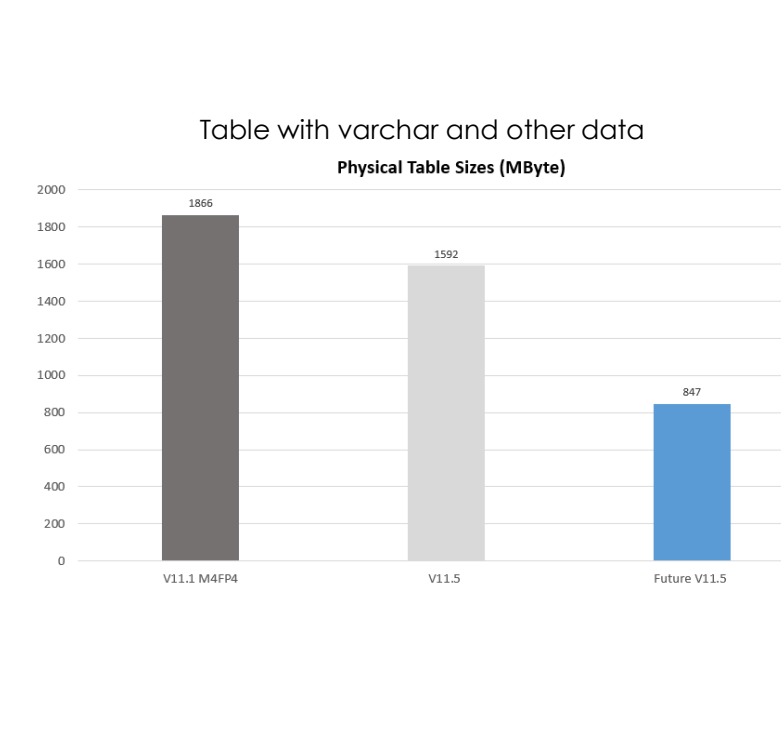

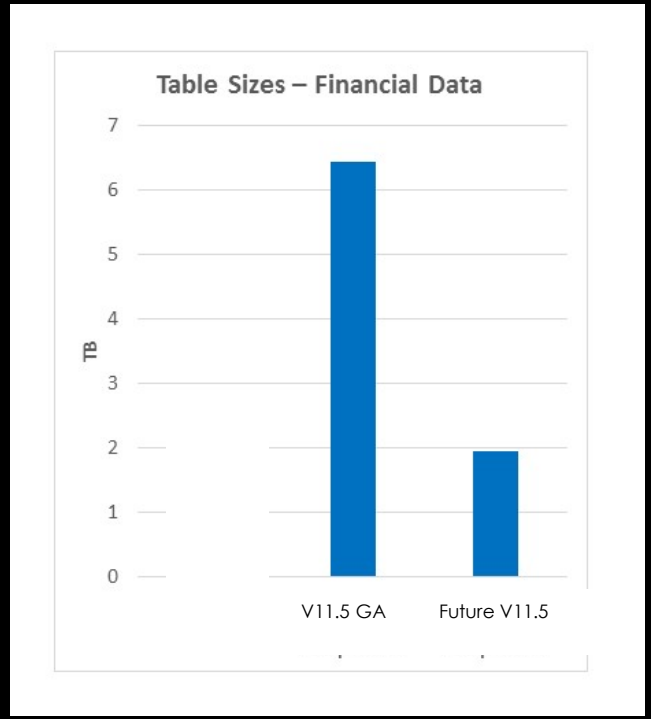

Internal Lab Results based on customer-like data.

SIMPLICITY-- RID SCALAR FUNCTION FOR BLU & DPF

Already exists for ROW tables. Now supported for COLUMNAR and w/ DPF when DBPARITIONNUM is used.

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: 0001020B000000000000000158

First you map the lock name to a TSNID:

select varchar(value,30) as TSNID from table(mon_format_lock_name('0001020B000000000000000158'))
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…
```

SIMPLICITY-- OPTIMIZER VERSION CONTROL

Goal:

Allow a DBA to 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 regression(s) may be experienced after an upgrade or update. Could also be used pro-actively. (Versus today, where a large number of both documented and undocumented reg vars would need to be set.)

It is not recommended to be run with long term.

SIMPLICITY-- OPTIMIZER VERSION CONTROL

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 then override it's setting.

SIMPLICITY-- BLOCK OPERATIONS THAT RESULT IN REORG PENDING

Goal: prevent a user accidentally running ALTER TABLE statements that put 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 NULL
- ALTER COLUMN SET DATE TYPE except in the following situations:
  - Increase VARCHAR/VARGRAPHIC length
  - Decrease a non-indexed VARCHAR/VARGRAPHIC length, without truncating trailing blanks

SIMPLICITY--BLOCK OPERATIONS THAT RESULT IN REORG PENDING

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 tablescan-only access).

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

PERFORMANCE-- FASTER BUFFER POOL STARTUP

ON by DEFAULT

Speeds BP startup by up to 120x.

The larger the BP, the larger the speed-Up.

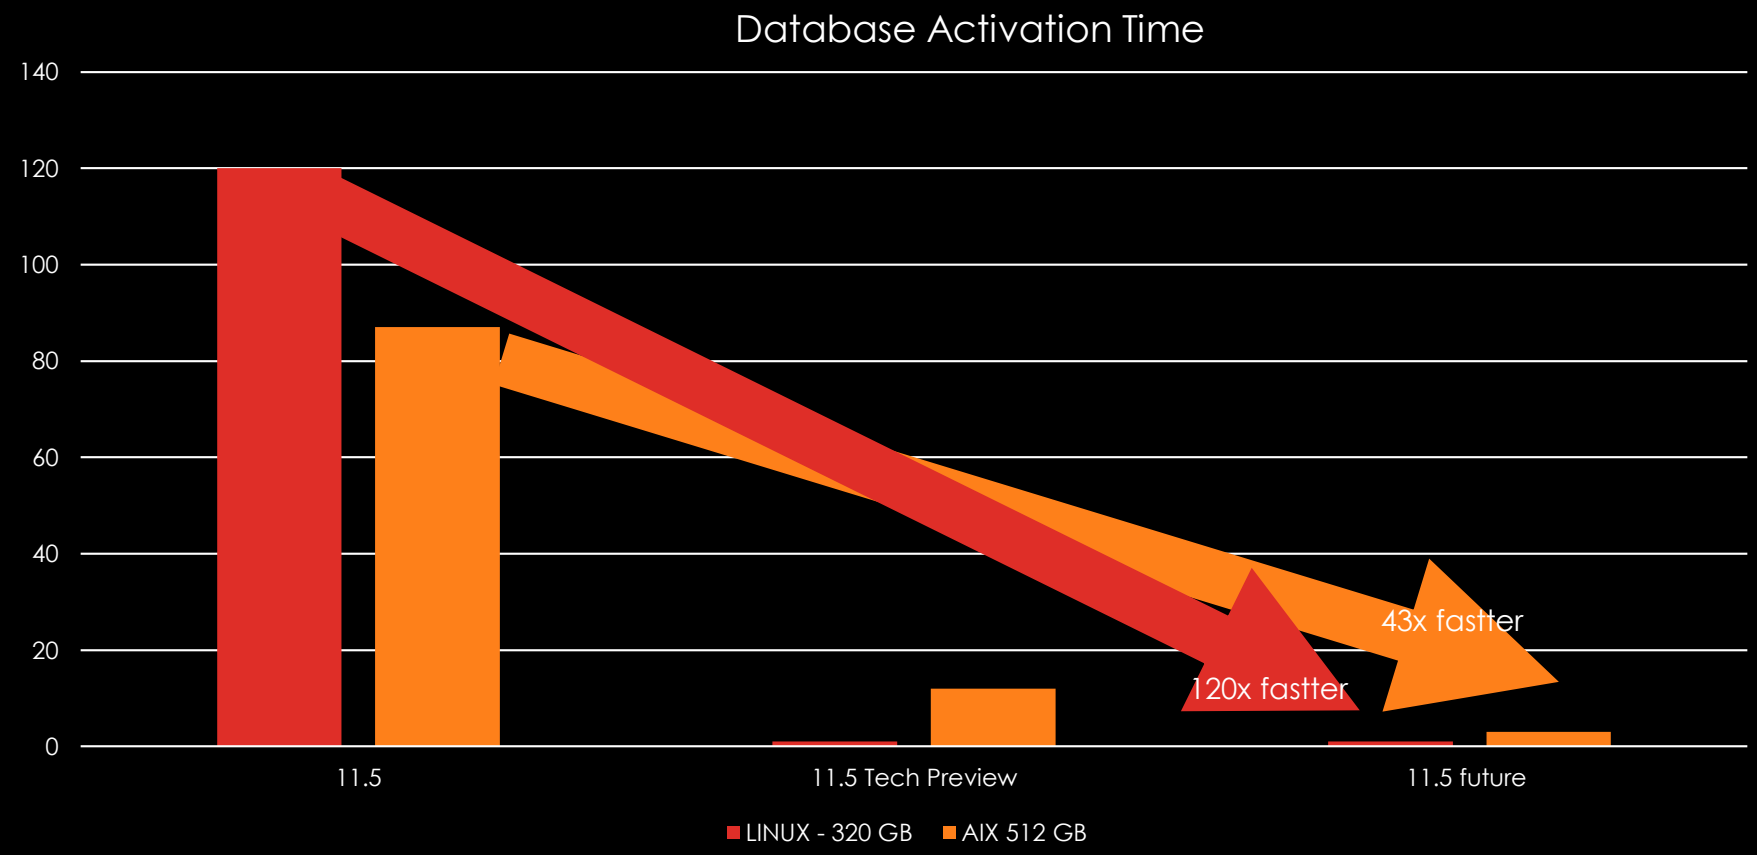

PERFORMANCE -- COLUMNAR TRICKLE INSERT

Goals:

Speed up “trickle” inserts and
reduce memory footprint.

Decrease size of small tables.

New registry variable:

DB2_COL_INSERT_GROUPS

Default OFF

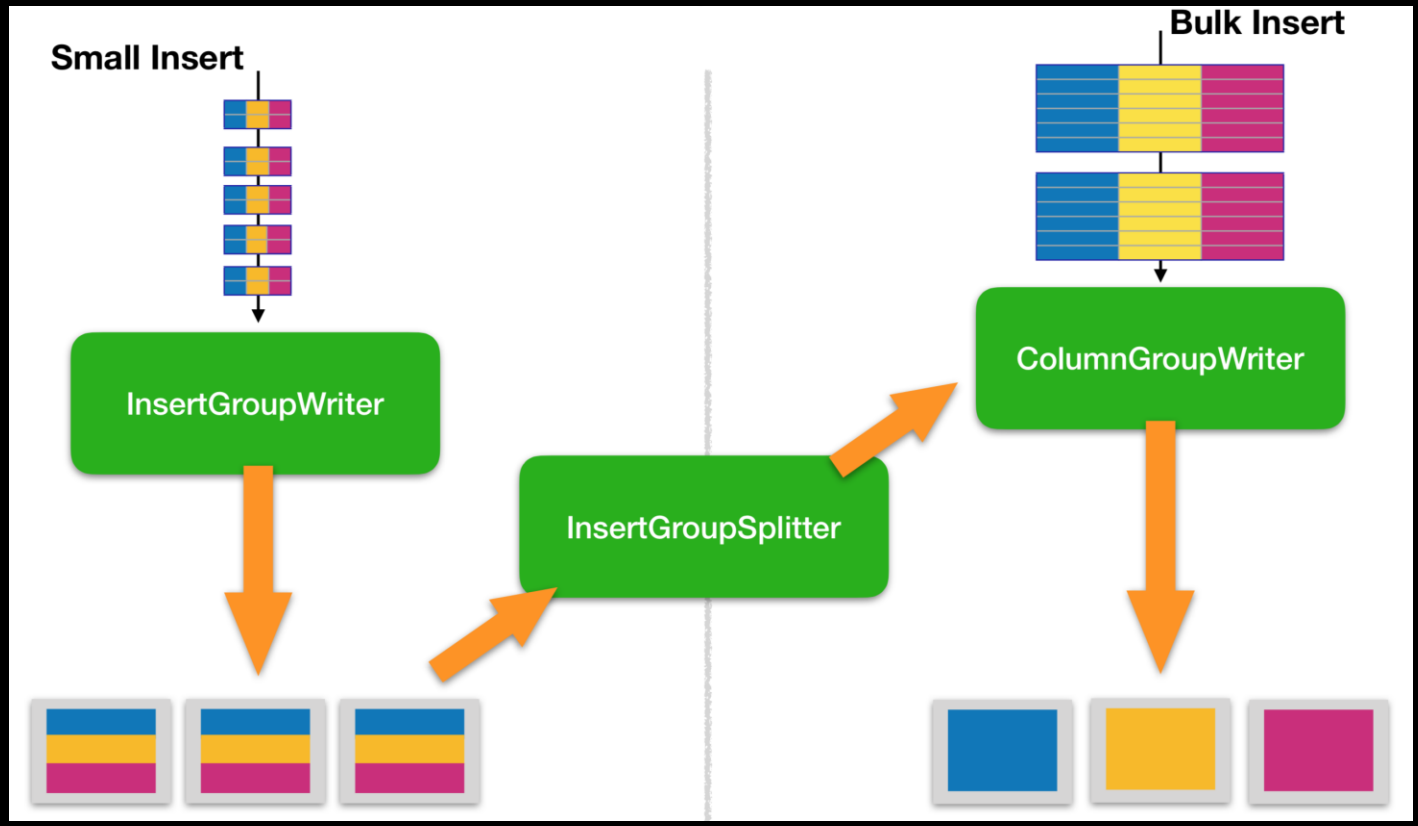

Once set to ON, any tables using this new insert mechanism will not be readable by previous modpacks of Db2 v11.5

Insert Groups

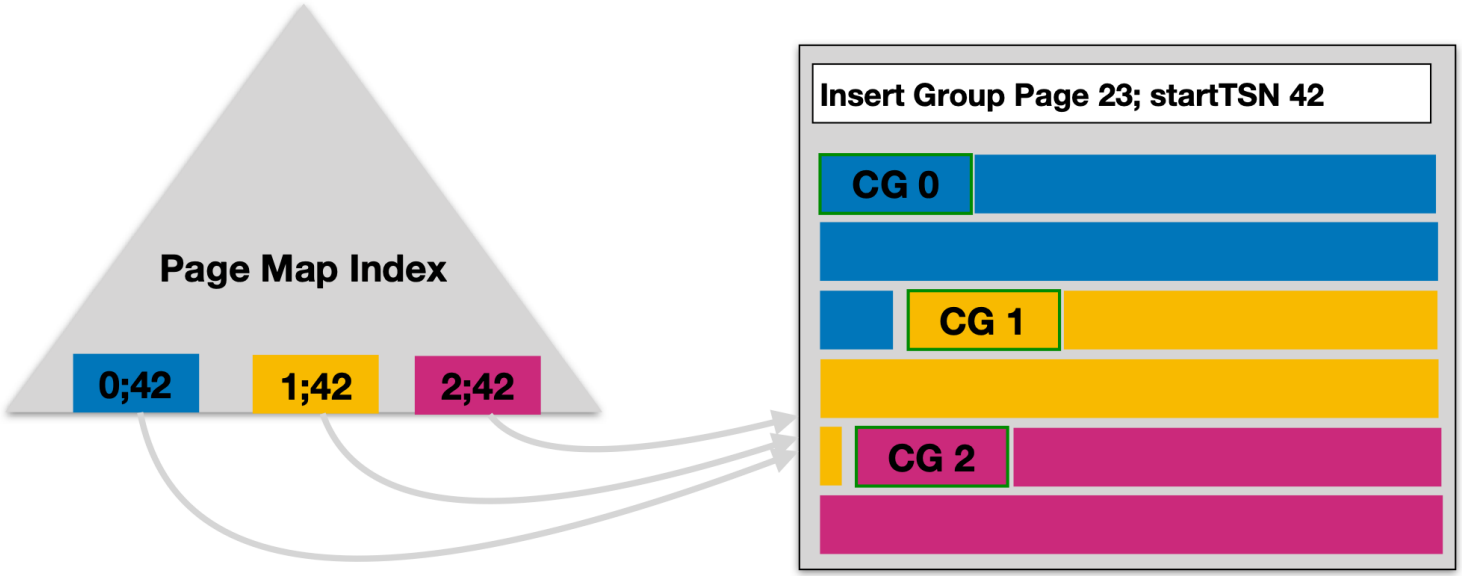

- Unencoded data of multiple column groups is stored in an **Insert Group**.
- Data is in **columnar format;** still one PMI entry per column group.
  - **Seamless integration** into existing infrastructure.

Small inserts are fast, readers scan vectors of columnar data.

Insert Groups

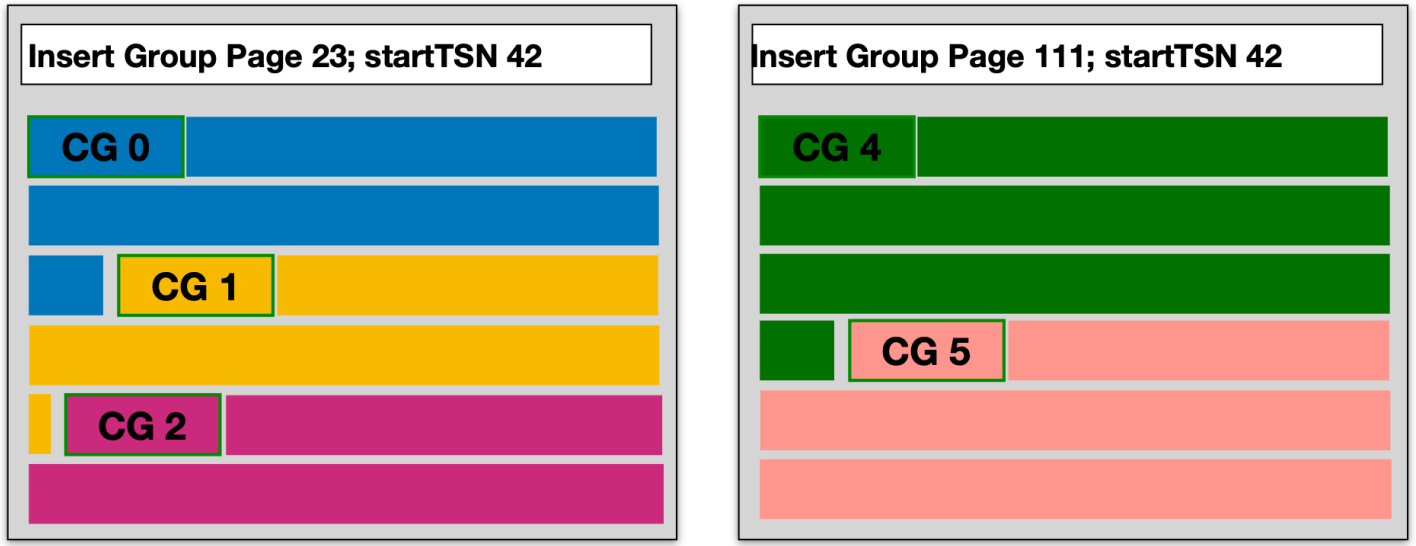

- Flexible assignments of column groups to insert groups.
  - Fixed-length vs. variable-length, large varchar columns.
  - Co-existence with other approaches, e.g., text compression.

Split Insert Group

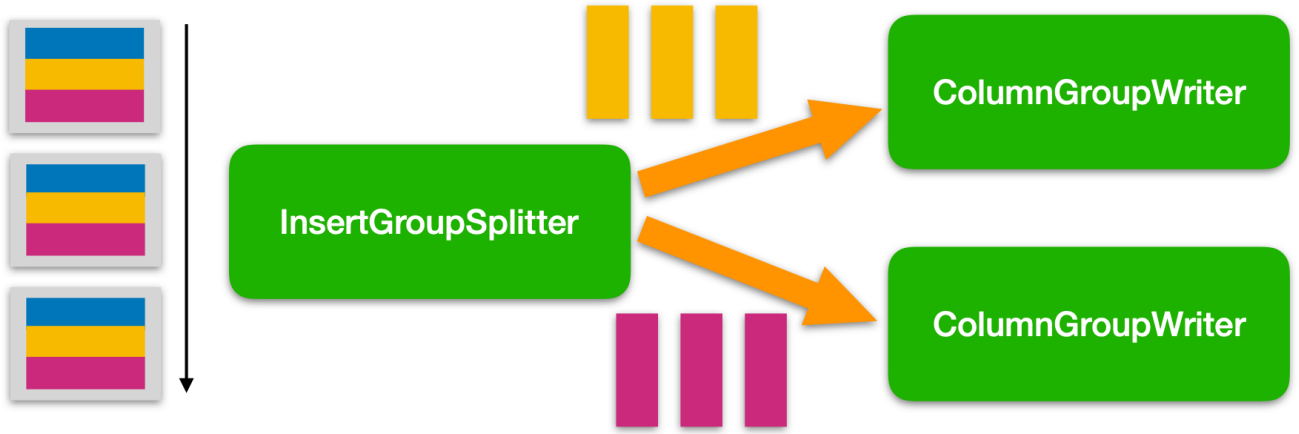

- **Synchronous** within a transaction (same approach as used for page compression).
- Tries to split when it thinks that a column group page can be filled.
  - Very **few pages** are **in Insert Group format**.
- Each column group is processed independently.

SIMPLICITY-- ADAPTIVE WLM

- Objective is to deliver automatic workload management within Db2 that ensures system stability and responsiveness with zero tuning
  - Don't overcommit the system yet ensure it's well utilized
  - Schedule jobs appropriately to ensure fairness and appropriate responsiveness

- Basic premise:
  - Incoming work guided to different "lanes" based on expected resource consumption and duration
    - Each "lane" gets a defined resource allotment to maximize predictability
    - Resources considered: Sort memory & CPU
  - Scheduling of work in each "lane" based on dynamic view of resource availability
    - Admission of new work is on the fit of the work to the available resources in that class and latency
    - Includes historical feedback based on past executions

THE INTRODUCTION OF ADAPTIVE WLM

- Currently being used by default in Db2 Warehouse family
  - New default workload configuration
- Will introduce as new, optional behaviour for default configuration in a future Db2 11.5 mod pack
  - Need to validate further with more row-centric workloads

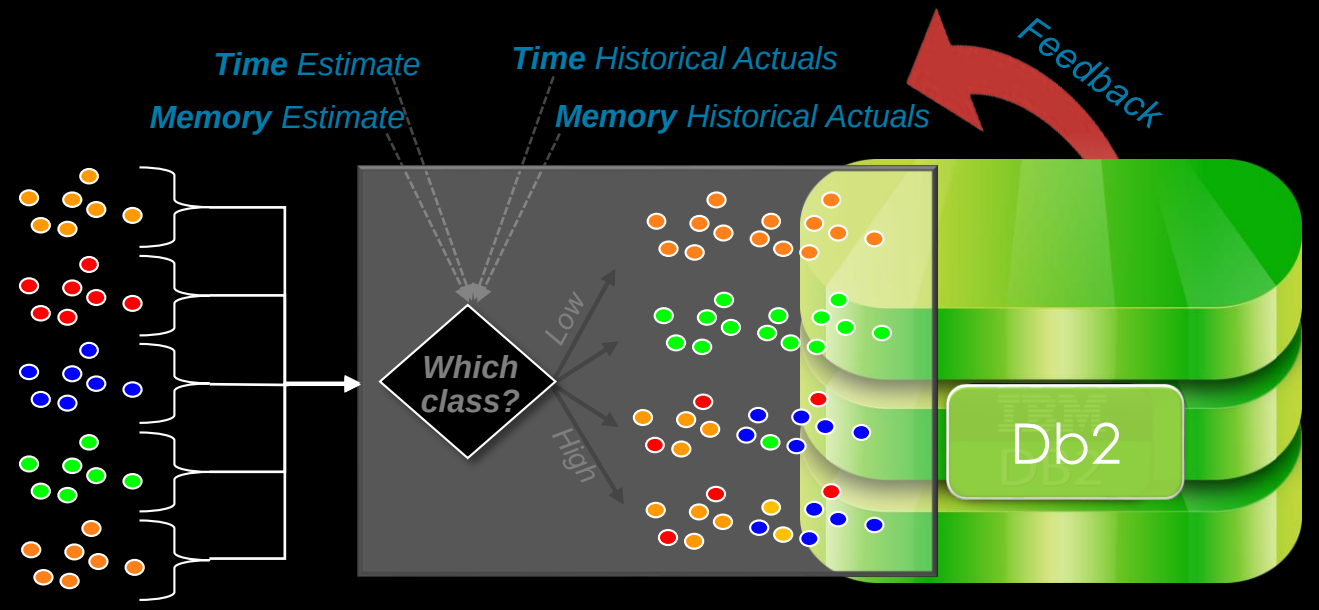

BRINGING AWLM TO Db2 V11.5 …

- AWLM will have a staged delivery to Db2
  - Stage 1: AWLM available for use by analytics (BLU) customers
    - Available to customers who set DB2_WORKLOAD = ANALYTICS
  - Stage 2: AWLM available for use by all Db2 customers
- Moving to AWLM will require an outage to allow the implementation of a new default system configuration
  - A new "opt-in" procedure will be provided

SIMPLICITY-- ADVANCED LOG SPACE MANAGEMENT

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

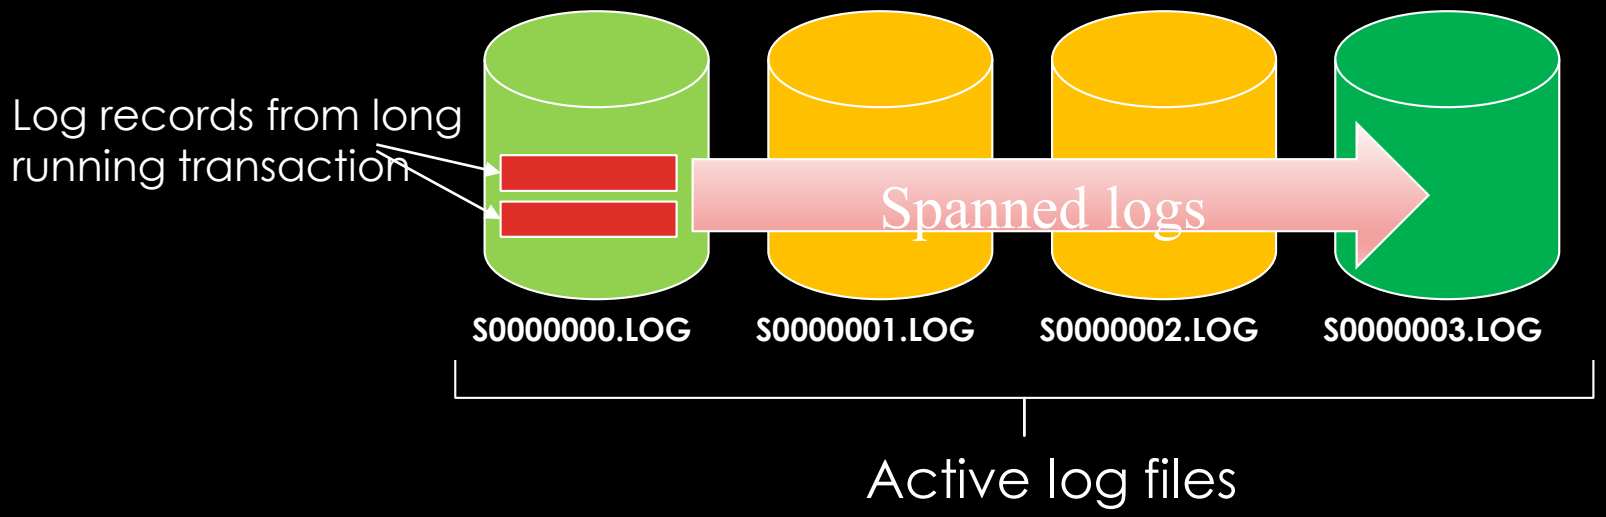

SIMPLICITY -- ALSM

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

SPATIAL ANALYTICS

- 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
- Currently Tech Preview in Db2 11.5 GA

SPATIAL ANALYTICS

- 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

SPATIAL PROCESSING IN Db2

	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

Operating System Support Matrix For DB2 11.5 Server non-pureScale

Distro	OS	Db2 11.5 GA	New Additions	To be Discontinued
AIX	AIX	- AIX 7.1TL5SP3 - AIX 7.2TL3SP2 - P7, P8, P9		- AIX 7.1* - Power 7
RHEL	- X86-64 - Linux390 - PPCLE	- RHEL 7.5	- RHEL 8.1 w/ p9 native	
SLES	- X86-64 - Linux390 - PPCLE	- SLES 12 SP3	- SLES 15 SP1 (no Automated HA)	
Ubuntu	- X86-64 - Linux390 - PPCLE	- Ubuntu 16.04 LTS		- Linux390 (32bit)
CentOS	- X86-64	- RHEL 7.5		
Windows	Windows	- Windows 2012R2 - Windows 2016 - Windows 10	- Windows 2019	
MAC OS	** DSDRIVER ONLY	- 10.10 Yosemite	- 10.15 Catalina	- 10.10 Yosemite

List is minimum level. For non--pureScale all minor levels after are supported by default (ie. 7.5 means also 7.6 support but not 8.x support)
For PPCLE:

- RHEL 7.5 (and above) is support on Power 9 but only in p8 compat mode
- Db2 does not support use of alternate kernels

Operating System Support Matrix For DB2 11.5 Server pureScale

Distro	OS	Db2 11.5	New Additions	To be Discontinued
AIX	AIX	- AIX 7.1TL5SP3 - AIX 7.2TL3SP2 - P7, P8, P9	- AIX 7.2 TL4SP1 - P9 Native on AIX 7.2TL4SP1	- AIX 7.1* - Power 7
RHEL	- X86-64 - PPCLE	- RHEL 7.5	- RHEL 8.1 - RHEL 7.6	
SLES	- X86-64 - PPCLE	- SLES 12 SP3 - SLES 12 SP4		

pureScale support is on specific Major and Minor level
RHEL 7.5 (and above) support on Power 9 native not supported, Db2 does not support use of alternate kernels

MUSICAL PERFORMANCE

Happy Belated April Fool's Day!

Really, it's better off for all of us this way… unless you want Alexa to play something.

WANT TO KNOW MORE DETAIL ?

Things can (and do) change. So to keep up-to-date:

- We revisit development priorities frequently (e.g. every quarter) in response to customer and market demand/feedback
- Public roadmap:
  http://ibm.biz/AnalyticsRoadmaps
- We want to hear from you!
  - http://ibm.biz/IBMAnalyticsIdeasPortal
- Piotr Mierzejewski is doing another Db2 Night Show May 1st

CREDITS

Everyone on the Db2 team, and all the Db2 family teams that contribute back to the core Db2 Common SQL Engine.

But especially the folks who I "borrowed" slides from to make this presentation!

THANKS!

And to Martin Hubel for running the Db2 Night Show!

Q&A