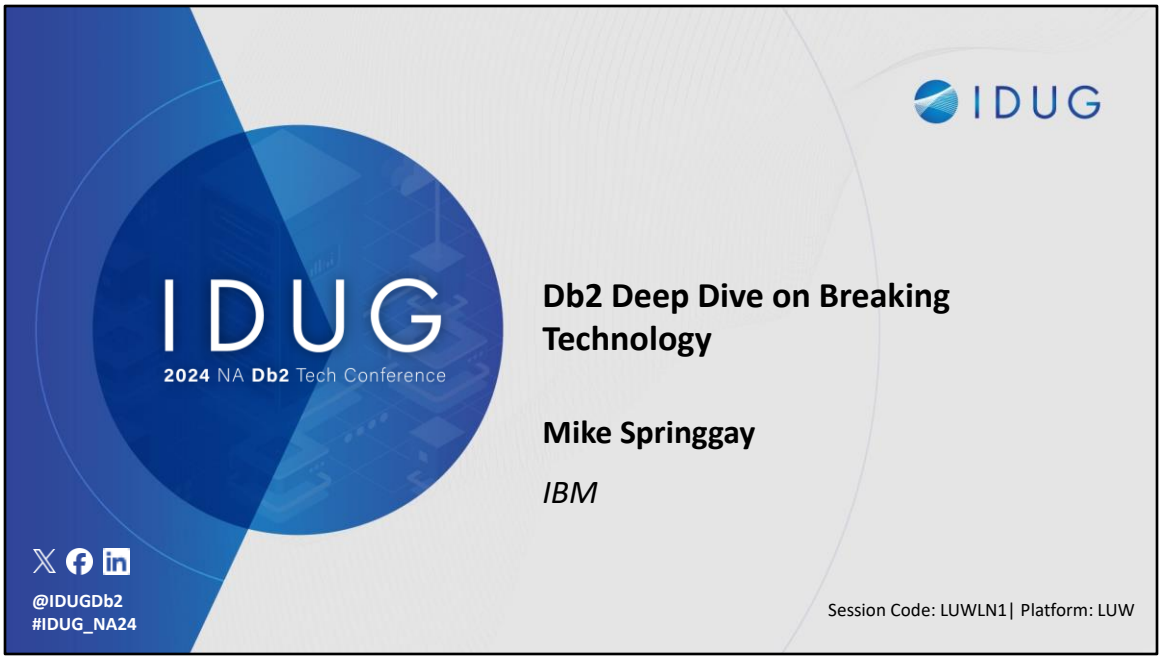
IDUG
2024 NA Db2 Tech Conference
IDUG
Db2 Deep Dive on Breaking Technology
Mike Springgay
IBM
@IDUGDb2
#IDUG_NA24
Session Code: LUWLN1| Platform: LUW

Notices and disclaimers

- Certain comments made in this presentation may be characterized as forward looking under the Private Securities Litigation Reform Act of 1995.
- Forward-looking statements are based on the company's current assumptions regarding future business and financial performance. Those statements by their nature address matters that are uncertain to different degrees and involve a number of factors that could cause actual results to differ materially. Additional information concerning these factors is contained in the Company's filings with the SEC.
- Copies are available from the SEC, from the IBM website, or from IBM Investor Relations.
- Any forward-looking statement made during this presentation speaks only as of the date on which it is made. The company assumes no obligation to update or revise any forward-looking statements except as required by law; these charts and the associated remarks and comments are integrally related and are intended to be presented and understood together.

Agenda

- 2023 Deliverables
- 11.5.9 Overview
- vNext: A sneak peek

Agenda

- 2023 Deliverables
- 11.5.9 Overview
- vNext: A sneak peek

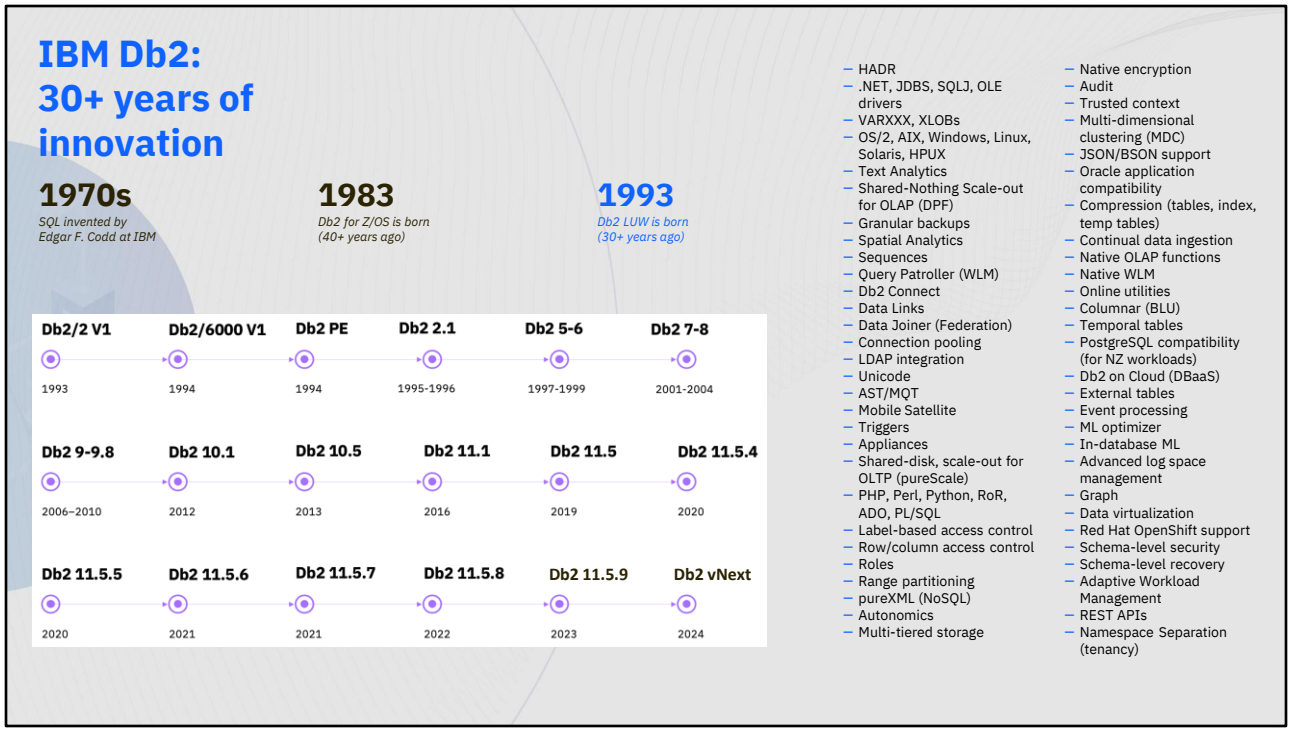

Db2 is a database that needs no introduction, but sometimes people can forget the rich history of innovation that underpins this technology that is built for the worlds mission critical workloads – you're going to hear that repeated a lot this week because its true, and because we are going back to basics with our messaging, highlighting our strengths and focusing on where we win.

We work with customers everyday to get them on the latest version of Db2 so they can take advantage of our latest innovations, but we also still support 6 different versions across 12 different operating systems and platforms and a few customers that have been running the same version of Db2 for over 10-years…. A real testament to the maturity and reliability of the technology.

Today's Db2 is much much more than a relational database for applications…. Db2 supports virtually every type of workload out there. Its hands-down the the worlds best database for SAP. Db2 pureScale supports 99.999% availability – that's 5 minutes of downtime per year. Advanced analytics, graph and machine learning workloads can be executed in the database without having to copy the data, improving performance and security and reducing cost – the proverbial better, faster, cheaper – and more secure.

Db2 continues to gain market share against competitors in the self-managed software segment of the market which includes both on-premises private clouds and public cloud IaaS – and when it comes to SaaS, as you'll hear today and throughout the week, we are laser focused on delivering a best-in-class managed service of Db2 on our customer's cloud of choice.

Db2 2023

Reference architecture for self-service deploy

Fully documented, self-service deployment of Db2u container on cloud managed Kubernetes/OpenShift & on-premises

Db2 Warehouse Gen 3 on AWS and 11.5.9

Fully managed cloud data warehouse featuring Db2 tables on S3, support for open data formats and watsonx.data integration

Features will also be available as software via 11.5.9

Db2 RDS as a managed service on AWS

Db2 for OLTP and mixed workloads available as a fully managed service on AWS

UX overhaul for management console

Continued investment to improve the cloud user experience for developers and DBAs

IBM Db2 +
Amazon Web Services

Partnering closely with Amazon to bring ohur Db2 offerings to AWS

Other offerings available, including:

- **Db2 Warehouse on Cloud (managed service)**
- **Db2 RDS for OLTP workloads (managed service)**
- **Db2 pureScale on AWS**
- **Db2 Container reference architecture**

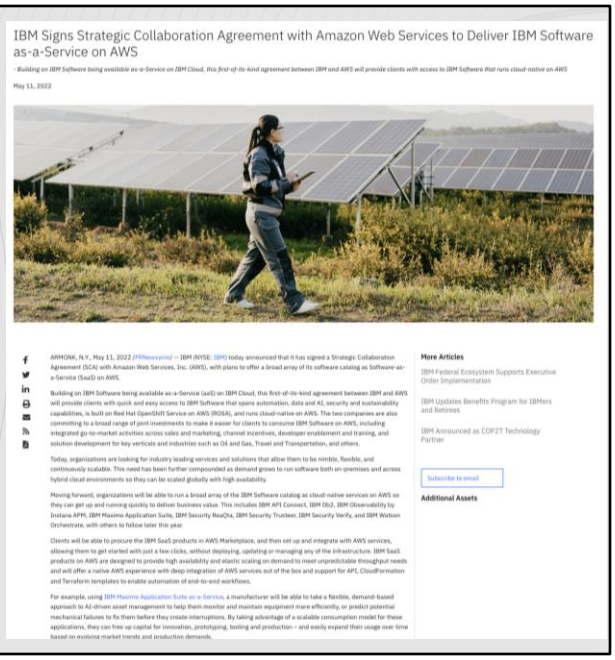

IBM Signs Strategic Collaboration Agreement with Amazon Web Services to Deliver IBM Software as-a-Service on AWS

- Building on IBM Software being available as-a-Service on IBM Cloud, this first-of-its-kind agreement between IBM and AWS will provide clients with access to IBM Software that runs cloud-native on AWS

May 11, 2022

ARMONK, N.Y., May 11, 2022 /PRNewswire/ -- IBM (NYSE: IBM) today announced that it has signed a Strategic Collaboration Agreement (SCA) with Amazon Web Services, Inc. (AWS), with plans to offer a broad array of its software catalog as Software-as-a-Service (SaaS) on AWS.

Building on IBM Software being available as-a-Service (aaS) on IBM Cloud, this first-of-its-kind agreement between IBM and AWS will provide clients with quick and easy access to IBM Software that spans automation, data and AI, security and sustainability capabilities, is built on Red Hat OpenShift Service on AWS (ROSA), and runs cloud-native on AWS. The two companies are also committing to a broad range of joint investments to make it easier for clients to consume IBM Software on AWS, including integrated go-to-market activities across sales and marketing, channel incentives, developer enablement and training, and solution development for key verticals and industries such as Oil and Gas, Travel and Transportation, and others.

Today, organizations are looking for industry leading services and solutions that allow them to be nimble, flexible, and continuously scalable. This need has been further compounded as demand grows to run software both on-premises and across hybrid cloud environments so they can be scaled globally with high availability.

Moving forward, organizations will be able to run a broad array of the IBM Software catalog as cloud-native services on AWS so they can get up and running quickly to deliver business value. This includes IBM API Connect, IBM Db2, IBM Observability by Instana APM, IBM Maximo Application Suite, IBM Security ReaQta, IBM Security Trusteer, IBM Security Verify, and IBM Watson Orchestrate, with others to follow later this year.

Clients will be able to procure the IBM SaaS products in AWS Marketplace, and then set up and integrate with AWS services, allowing them to get started with just a few clicks, without deploying, updating or managing any of the infrastructure. IBM SaaS products on AWS are designed to provide high availability and elastic scaling on demand to meet unpredictable throughput needs and will offer a native AWS experience with deep integration of AWS services out of the box and support for API, CloudFormation and Terraform templates to enable automation of end-to-end workflows.

For example, using IBM Maximo Application Suite as-a-Service, a manufacturer will be able to take a flexible, demand-based approach to AI-driven asset management to help them monitor and maintain equipment more efficiently, or predict potential mechanical failures to fix them before they create interruptions. By taking advantage of a scalable consumption model for these applications, they can free up capital for innovation, prototyping, tooling and production -- and easily expand their usage over time based on evolving market trends and production demands.

More Articles

IBM Federal Ecosystem Supports Executive Order Implementation

IBM Updates Benefits Program for IBMers and Retirees

IBM Announced as COP27 Technology Partner

Subscribe to email

Additional Assets

Amazon RDS for Db2

- **Fully managed service** that allows you to spend more time building, less time managing

Push-button scaling with a few clicks or API call to cut costs

- **Automated backups, snapshots, and failover** to support durability of business-critical workloads

- **Availability and reliability** with high availability and automated multi-AZ data replication

- **Isolation & Security** including encryption in motion and at rest, network isolation, and permissions

Easy migration from on-premises to AWS RDS for Db2 with options for one-time and continuous data movement

- **Native IBM and AWS integrations** to scale analytics and AI with
- Db2 Warehouse, watsonx.data,
- and Amazon S3

Db2 pureScale on AWS Marketplace

Available since July 2022

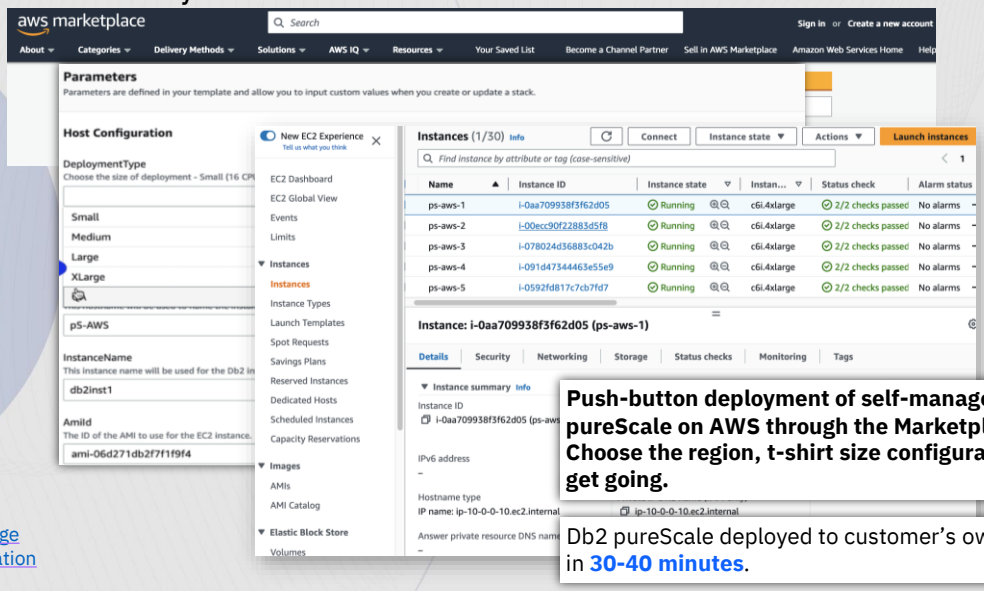

Push-button deployment of self-managed pureScale on AWS through the Marketplace. Choose the region, t-shirt size configuration and get going.

Db2 pureScale deployed to customer's own account in **30-40 minutes**.

- Product Page
- Documentation

BYOL – Bring Your Own Licence

Support

IBM Db2 pureScale
Support level is based on the software license agreement

AWS Infrastructure
AWS Support is a one-on-one support channel that is staffed 24x7x365 with experienced support engineers. AWS Support offers four support plans: Basic, Developer, Business, and Enterprise. The Basic plan is free of charge and offers support for account and billing questions and service limit increases. The other plans offer an unlimited number of technical support cases with pay-by-the-month pricing and no long term contracts, providing the level of support that meets your needs. Learn more

2023 Additions:
New 1 click to provision a second DR cluster!

New template to add member(s)
New template to add disk(s) to file system

Business Transformation with Db2 Warehouse – On-premises P10 and X86 Architecture Solutions

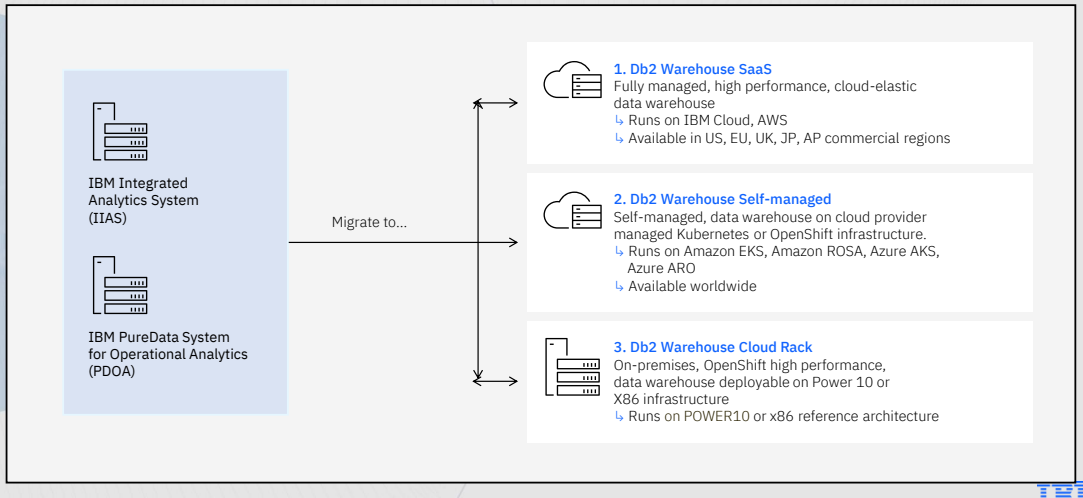

Agenda

- 2023 Deliverables
- 11.5.9 Overview
- vNext : A sneak peek

Db2 11.5.9

Core Engine

- RHEL 9.2 support
  - RHEL 7 will no longer support integrated TSA (for HA or pS)
  - Ubuntu 22.4 support (no TSA/PCMK support)
- Remote Storage (S3) available on additional Linux platforms (PPCLE and zLinux)
- Advanced Data masking support (Linux only until vNext)
- Audit Logs to Remote Storage (S3)
- Restricted TCPIP Listener Mode
- Ability to protect any History File entry from pruning
- Online Synopsis Table Rebuild
- Federation Enhancements
  - Support for AWS Aurora PostgreSQL and PostgreSQL 12-15
  - Watsonx.data support via JDBC
- ECMTS, Java, Gskit, TSA, GPFS and Pacemaker stacks updates

pureScale

- RHEL 9.2 & 8.8
- SLES 15 FP5,
- AIX 7.3 TL1 support

- Z16 support
- CX-6 card support
- RoCE virtualization support on AIX 7.3

- CF reduction in latch contention with high read rates

Db2 11.5.9

Containerization: Hybrid & Multi-cloud

- Reference Architectures for Db2 Warehouse on Power, x86 and Cloud
- Support for tablespace storage to be on S3 (aka Remote Table spaces) – Db2 Warehouse only
- Support for Open Data Formats: Iceberg, Parquet, AVRO, ORC - Db2 Warehouse only
  - Including native watsonx.data integration

- Operator driven Db2 native backup and restore
- Db2 W Scaling - vertical by adjusting resources per pod
- Db2 W Scaling - horizonal by scaling out number of pods
- HADR role-aware Kubernetes service for Db2U. Route clients to primary.

* Delivered (1Q 2023) in latest containers

Object storage support for table storage

- Db2WoC Gen3 now supports Amazon S3 object storage for database table storage, where customer data resides within the database
- Customer saves cost by using object storage instead of block storage
  - Allows customers to choose to store data on block storage or object storage, based on business or technical requirements
  - Db2WoC uses different mechanisms to facilitate reads and writes to object storage
  - Enables a consumption-based model for the storage, with all the benefits of automatic and unlimited storage scaling
- No applications and workload changes necessary to use this feature
  - Db2WoC handles all the necessary interfacing to object storage, thus existing applications & warehouse workloads do not have to be changed to make use of this
  - Specific Db2 tablespace available backed by S3 for customer use
  - Insert, Update, Delete data into and out of tables within object storage
  - Move and copy data to and from column-organized tables residing in block storage and object storage
  - Query data seamlessly no matter where it resides (in block or object storage), in isolation or in combination with each other

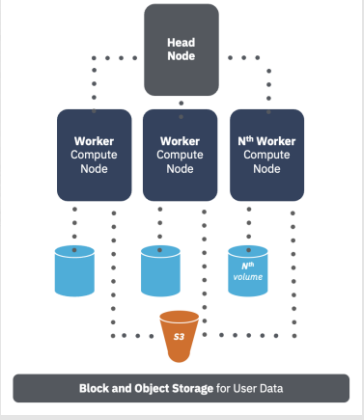


```
TBSP_NAME          MEMBER TBSP_TYPE   CACHING_TIER
------------------ ------ ---------- ------------
...
OBJSTORESPACE1          0 DMS        ENABLED
OBJSTORESPACEUTMP       0 DMS        ENABLED
```

DATALAKE tables support

- Db2WoC Gen3 now also supports Open Data Formats as DATALAKE tables, allowing for seamless access to other data within the enterprise for integrated workloads.
- Leverage existing compute resources dedicated to the warehouse
  - Facilitate data use to and from the Db2 Warehouse to quickly access a variety of enterprise data
  - Leverage the high performance Db2 engine for queries against enterprise data
- Db2WoC provides interfaces for customers to leverage their enterprise data residing in object storage as DATALAKE tables
  - Supports both regular and Iceberg DATALAKE table types, based on existing data formats or for business/technical requirements such as ACID compliance
  - Browse, explore, and query enterprise data in both Db2 and DATALAKE formats, using either the web-based UI, or through SQL
  - Access data in place within DATALAKE tables, joining as necessary with Db2 based data for queries
  - Access data within DATALAKE tables and import into Db2 formatted tables
  - Create new DATALAKE tables in S3 and export from Db2 formatted tables

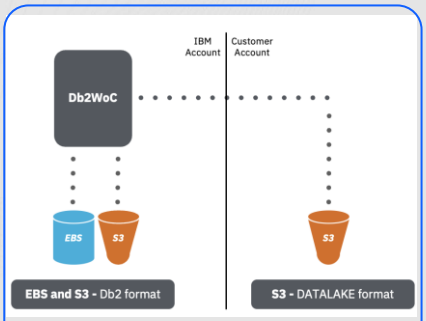


```
CREATE DATALAKE TABLE my_datalake_table
  ( id INT,
    name VARCHAR(8)
  )
  STORED AS PARQUET
  STORED BY ICEBERG
  TBLPROPERTIES ('external.table.purge'='true')
  LOCATION 'DB2REMOTE://mybucket-alias//
                        my_datalake_table';
```

Restricted TCP Listener Mode (1/2)

Purpose

- Allow secured application connections using connect type 2

Problem

- TCP listener is not running (SVCENAME, DB2COMM)
- Applications use connect type 2 using secure port only
- Sync Point Manager requires unsecured TCPIP (DRDA protocol)
- Distributed transaction processing is not possible and db2diag.log message is logged

Restricted TCP Listener Mode (2/2)

Solution

- TCPIP listener can be started in restricted mode, dropping any connection request other than:
  - resynchronization (SPM)
  - DRDA ping

How to Setup

- Configure SVCENAME in DBM CFG
- DB2COMM registry variable does not contain value TCPIP

Logging

- One EVENT db2diag.log message in instance start (if configured)
- One INFO db2diag.log if clients attempt to connect and are rejected logging client IP

Online Columnar Synopsis Table Rebuild (1/2)

Purpose of rebuilding synopsis table:

- Reduce synopsis table excessive storage usage due to sparsity in data pages
- Recovery from issues with synopsis table if encountered

Restriction:

- When rebuilding a synopsis table, the corresponding base table is in EXCLUSIVE mode, not available for read or write. (The database is ONLINE, but table not accessible until complete.)

Online Columnar Synopsis Table Rebuild (2/2)

Rebuild synopsis table syntax:

```
REORG SYNOPSIS FOR TABLE  <base table> REBUILD
```

Example:

```
REORG SYNOPSIS FOR TABLE customer REBUILD
```

The REORG SYNOPSIS command needs to be the first statement in a UOW. If not, SQL0428N will return.

If the REORG SYNOPSIS command does not complete successfully (SQL2223N), the synopsis table will remain in an unusable state.

- The synopsis table will not be included in any query plan
- Querying the synopsis table will receive an error (SQL0668N reason code 12).

Data Masking (1/3)

Purpose: Provide enhanced data privacy capabilities

Adds built-in masking function that support:

- Redaction
- Partial Redaction
- Substitution
- Obfuscation

Can be combined with Row and Column Access Control to provide advanced data privacy

Based on the IBM Research Next Generation Data Masking Engine (Magen) Library
For advanced obfuscation, the following list of subproperties are supported.
Preserve Format
Irreversible masking
Repeatable Consistency
Random Consistency
Input validation
For the advanced options and obfuscation, the following data classes are supported.
Person name / First name / Last name / Email Address
Gender Honorific
US Street Name, US State Capital Name, US zip code, US phone number, US street name, City (English), Country Name,
US Social Security Number, All the 50 US state driving license number formats, Commercial and Government Entity Code
Employment Status
Ethnicity, Eye Color, Hair Color
Hobby/Leisure Activity
Legal Marital/Civil Status, Name Suffix, Relationship, Religion

Data Masking (2/3)

```
DATA_MASK(expression, mask-type, mask-parameters, mask-format, seed)
```

Expression	Expression that returns value to be masked
Mask-type	An expression that specifies masking operation to be perform
Mask-parameters	String that provides additional parameters to influence masking operation
Mask-format	What format to apply when Obfuscation format preserving / format preserving fabrication
Seed	A string value to use as seed to hashing function

Data Masking (3/3)

Mask-Type	Masking Name	Data-Type Support	Description
0	REDACT	CHAR , VARCHAR*	redact string via mask-parameters
1	REDACT PRESERVE LENGTH	CHAR , VARCHAR*	redact character via mask-parameter
2	SUBSTITUE	All support data types	Strings SHA-256 hash & base64 encode. Others within type range
3	PARTIAL REDACT	CHAR , VARCHAR*	Pattern via mask-parameters
4	FORMAT PRESERVING	CHAR , VARCHAR, DATE TIMESTAMP*	matches format via mask-format
5	DATE AGING	DATE and TIMESTAMP*	mask-parameters how date is aged
7	IDENTIFIER	CHAR , VARCHAR*	Alphabets and digits masked other characters remain as-is
9	FORMAT PRESERVING PARTIAL	CHAR , VARCHAR*	Only for Email currently
10	FORMAT PRESERVING FABRICATION	CHAR , VARCHAR, DATE TIMESTAMP*	obfuscation without input validation
11	NUMERIC SHIFT	Numeric Types	mask-parameters integer percentage to shift input value by

** All other data-types redacted to default corresponding value*

All other data-types redacted to default corresponding value

SUBSTITUE - 43 bytes length required to avoid collisions on hash and encoding of string values. No seed default value is used

FORMAT PRESERVING, FORMAT PRESERVING FABRICATION, IDENTIFIER, NUMERIC SHIFT - If seed is empty or null, random seed is used.

FORMAT PRESERVING PARTIAL and PARTIAL READACT – limited support

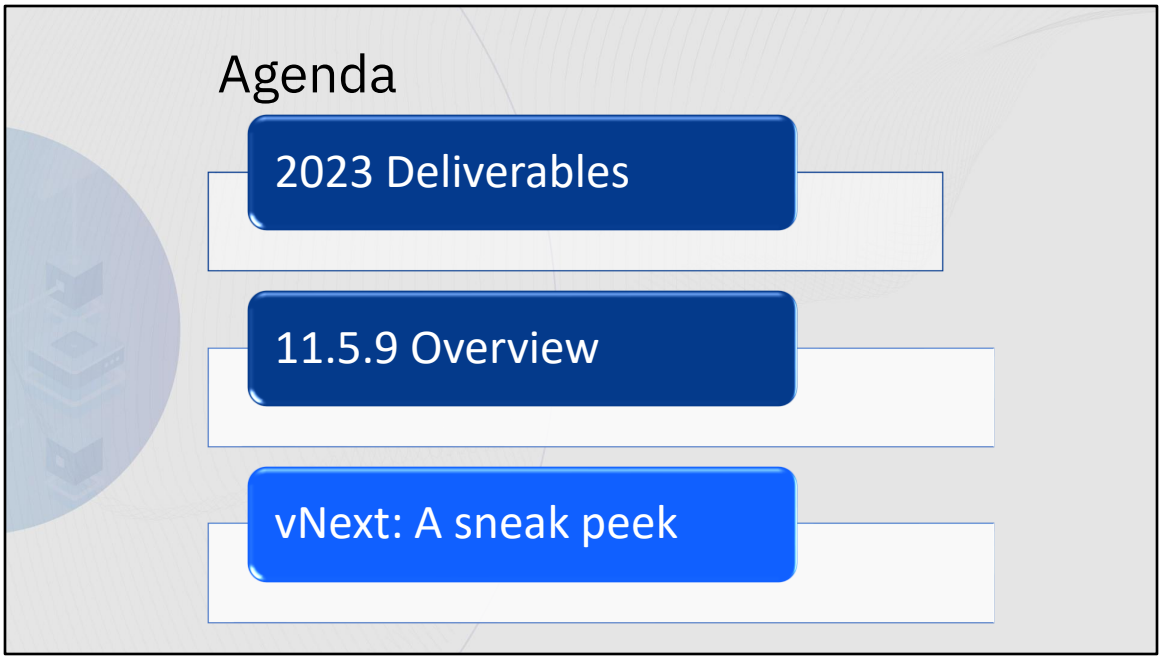
Agenda
2023 Deliverables
11.5.9 Overview
vNext: A sneak peek

Db2 Four big bets for 2024

Continued Investment in Db2 on Amazon RDS

Roadmap evolution including new licensing options and other enhancements that make it easy to modernize

Db2 Warehouse Gen3 on IBM Cloud

Fully managed cloud data warehouse featuring Db2 tables on Cloud Object Storage, support for open data formats and watsonx integration

Db2 infused with Generative AI

We're adding Gen AI capabilities into Db2. Stay tuned.

UX overhaul for management console

Continued investment to improve the user experience for devs and DBAs

Db2 12 Candidates

Planned for 2024, Db2 12 will bring significant enhancements to Db2 pureScale, name space separation, generative AI-powered insights, a new AI optimizer and hundreds of other enhancements.

Db2 pureScale improvements

Replacement of TSA with Pacemaker technology for cluster management, leading to significantly faster failure recovery times

Name space separation with TENANT construct

Create a logical separation between one or more database schemas, easily isolating differing sets of tables from each other

AI–powered query optimizer

Allows Db2 to continuously learn from customer's queries and achieve up to 3x query performance improvement over prior version

Db2 infused with Generative AI

We're adding Gen AI capabilities to Db2. Stay tuned.

- **Improvements to backup performance** by initiating multiple threads to process a single table space
- **Mac M1/M2 driver support** for developers on macOS using Apple Silicon chip
- **Db2 pureScale HADR support for enterprise-grade end-to-end SSL** encryption
- **Online index reorg for Db2 pureScale** allowing index reorg while table remains online/available
- **ADMIN_MOVE_TABLE** performance enhancements
- **Security enhancements** with AUDIT exceptions, Trusted Context and data masking
- Continuing investment in **cloud object storage performance**
- **Schema evolution with DROP and RENAME** support for online schema updates to columnar tables
- **UPDATE and JOIN** performance enhancements for columnar tables
- **Logical backup/restore** experience improvements
- **Recovery time improvements** in the unlikely event of crash
- **Federation enhancements** with support for Snowflake, Oracle 23c and performance improvements

IBM

pureScale: Continue investment to make it the best OLTP solution for the cloud, with the highest performance, availability and recovery times. In addition to paceMaker,, bringing the highest performance equivalent to on-premises environment. Continued investment in reducing downtime, with work we are doing for online index reorg

Looking ahead we intend to invest in AWS EFA support to support higher write wokloads (i.e RDMA equivalency) and increased scale operations with online-drop member.

Db2 Namespace Separation

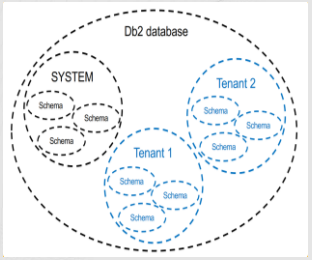

- What is Namespace Separation (Tenant)?
  - A single physical database supporting multiple “Tenants”
  - All tenants share the same infrastructure and resources
  - Each tenant has a private, isolated perspective for their own objects
- What problem does having Tenants solve?
  - Netezza/Postgres compatibility
  - Cost-savings through consolidation
  - Reduction of fixed overhead costs associated with individual databases
  - Centralization/simplification of database operations
  - Sharing unique environments without collisions
  - Significant cost-savings for our customers with many small development systems

The default SYSTEM tenant

- The initial set of catalogs established when the database is created
  - This is referred to as the SYSTEM tenant
  - Cannot be removed
  - Contains catalog information for shared resources and Db2 defined objects

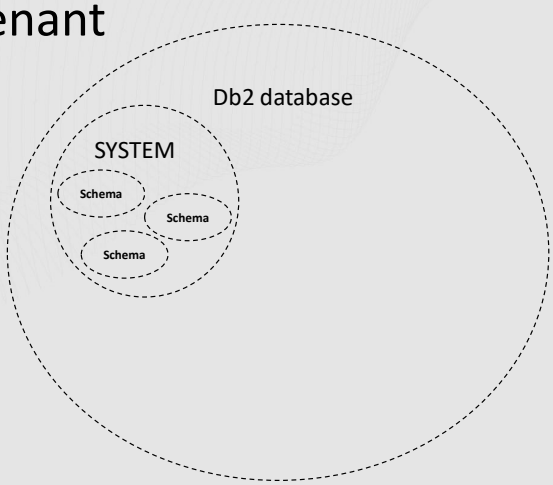

User-defined tenants

- 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

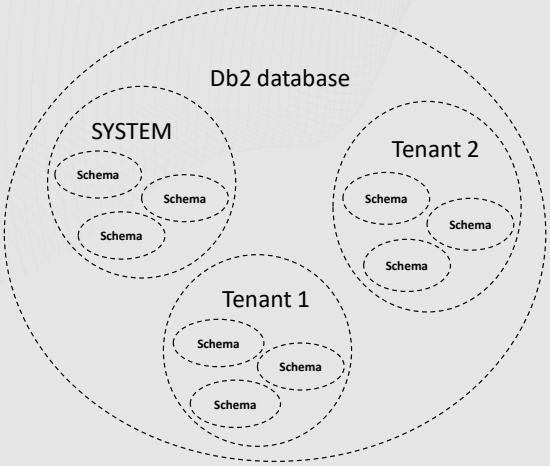

Scope of Namespace delivery

- New CREATE/DROP TENANT DDL
- New SET TENANT statement and CURRENT TENANT special register
- New GRANT/REVOKE USAGE privilege to control tenant access
- Introduction of TENANT as WLM workload definition attribute
- Introduction of tenant awareness to key (not all) interfaces
  - Introduce tenant information on key monitoring interfaces
  - Introduce tenant input option on key tools

Authorization in the world of tenants

- Scope of existing Db2 authorities and privileges not affected
  - Database level authorities apply across all tenants
- Authorization records for individual objects are recorded in the catalogs for the tenant in which the object is defined
  - Db2 provided objects are defined in the SYSTEM tenant
- Database role membership definitions are shared across all tenants
  - Authorization granted to a role are recorded in the tenant where the target object is defined

Db2 Namespace Authorizations

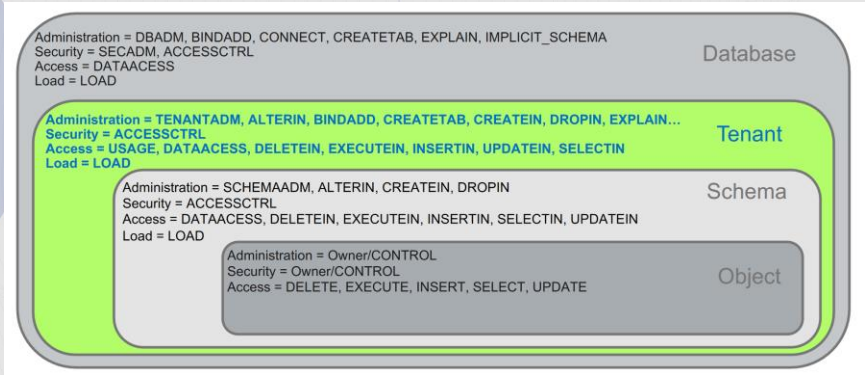

A simple “Namespace” example (setup)

```
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
```

A simple “Namespace” example

CONNECT TO TEST USING USER1

SET CURRENT TENANT = WORLD1

VALUES (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

VALUES (CURRENT TENANT)
- WORLD2

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

INSERT INTO MINE.T1 VALUES (‘A’)

SELECT * FROM MINE.T1
- ‘A’

Security Enhancements

Trusted Context – TRUST PROCEDURE

- A stored procedure can be used as an attribute of trust for a TRUSTED CONTEXT
  - You can examine more than just a single address/encryption level, such as system global variables – for example for a range of IP addresses, time of day
  - We are adding an application token, sent during connect, that can provide application-level details/identification
    - Callback in application to obtain at client, and available to TRUST PROCEDURE

Audit Exception

- Ability to turn off auditing within a Trusted Context
  - To be combined with previous solution
- AUDIT ADD EXCEPTION FOR TRUSTED CONTEXT <tcname>

Table space and Buffer pool authority enhancements

- Our model is SYSADM/SYSCTRL should control instance wide / machine resources that are outside the database, memory, disk etc.
  - But can delegate to DBADM, Database users
- DBADM will now be able to create Automatic Storage Tablespaces inside storage groups that the SYSADM has created
- DBADM will now be able to create Bufferpools, up to database_memory limit that SYSADM can set
- We are looking to also create new TBSPACEADM authority to manage tablespaces within a storage group, and USE privilege on bufferpools, so these can be pushed down to application administrators, schemaadms, tenantadms

V12: Misc

- Use trusted context role for package execution/static SQL
  - Aha! 407
- Change to KMIP SSL label now an online operation
  - Renewing client side authentication certificate
- JWT from authentication exposed as Global Variable for use with RCAC
- Improve update/upgrade with RCAC dependency on system objects (catalog views, functions etc.)
  - Aha! 1004
- TLS + SERVER can connect in place of SERVER_ENCRYPT authentication

GRANT DBADM statement changes

- Change

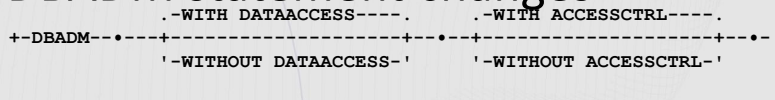

- TO

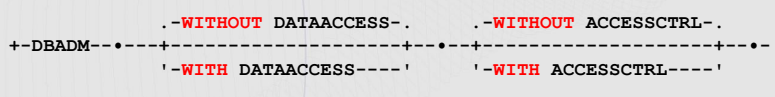

- Silently granting high level privileges is a security risk
- No changes to existing grants, this only affects the GRANT statement
- To revert behaviour:
  - **db2set DB2_ALTERNATE_AUTHZ_BEHAVIOUR=DBADM_ADDTL_AUTHS**
  - Automatically set when DB2_WORKLOAD=SAP
- The following is still valid accepted (preferred) syntax:
  - GRANT DBADM, DATAACCESS, ACCESSCTRL ON DATABASE

Authorities/privileges no longer granted to PUBLIC during database creation (non-restrictive)

- CONNECT privilege
  - Any user that can authenticate can connect
    - Might mean anyone in LDAP, not just intended Db2 users
- Following authorities could allow a malicious user to create database objects that use up instance resources
  - IMPLICT_SCHEMA
  - CREATETAB
  - BINDADD
  - USE on USERSPACE1
  - CREATEIN on SQLJ and NULLID schemas
- No changes to existing database
- To revert behaviour:
  - db2set DB2_ALTERNATE_AUTHZ_BEHAVIOUR= PUBLIC_DBCREATE

TLS / SSL Changes

- TLS 1.0 and 1.1 removed from Db2 server
  - Also removing support for SHA-1 and 3DES cipher suites
- TLS 1.3 – on by default at server
  - Restrictions on certificates when using TLS 1.3
    - No SHA1 or SHA224 signatures
    - No RSA keysize < 2048
- Hostname Validation On By Default
  - Validation by client that hostname for connection matches what is in certificate
  - **Requires certificate to be created properly**
    - May cause connection failures if proper certs not configured
  - Many places that TLS is used and hostname validation turned on:
    - Clients (CLI/ODBC/JDBC), including XA
    - HADR (between primary and standby)
    - Connecting to KMIP Key Manager (Native Encryption)

41

TLSV1 is very insecure and rapidly being removed by industry
TLSV12 supported since 9.7 fp9, 10.1 fp 4, 10.5 fp3
z/OS V2R1 -now out of support
Clients older than above will not be able to connect with TLS
Clients still support TLS 1.0 and 1.1 for downlevel compatibility

SHA-1 cipher suites and 3DES cipher suites for TLS -- only issue if explicitly configured in SSL_CIPHERSPECS dbm cfg - default is to pick strongest cipher.

Hostname Validation:
Existing feature added in 11.5.6, being turned on by default
Hostname validation for various topologies extensively documented :
https://www.ibm.com/docs/en/db2/11.5?topic=transit-hostname-validation

3DES with native encryption (1|2)

- 3DES is no longer considered secure
- You will still be able to decrypt any 3DES data, but will need to use AES when encrypting new databases/backups
- Recall database and backups can have independent algorithms
  - By default, ENCROPTS db cfg has backups use the same algorithm
- Following are still supported
  - ACTIVATE DATABASE when it is using 3DES
  - Restoring from a backup encrypted with 3DES
  - Restoring on top of an existing database that is using 3DES
  - Migrating a database from v11 to v12 that is using 3DES
- AES is the replacement algorithm that should be used
  - Master Keys are always AES, it's only the DEK of the database that can be specified by the user – existing master keys are still valid/secure
  - AES should perform better than 3DES

3DES with native encryption (2|2)

- Create a new database using 3DES is no longer supported
- Restoring into a new database using 3DES is no longer supported
  - In this scenario RESTORE command would already be specifying
    - RESTORE … ENCRYPT CIPHER 3DES
  - Reading from the backup is OK, it's the new database that will require AES
- Backup a database where backup is encrypted with 3DES is no longer supported
  - This is the default for databases encrypted using 3DES
  - ENCROPTS (db cfg or backup cmd) needs to be updated to specify AES
- To migrate a database to AES (not required)
  - Backup, drop, and restore as new database using AES

Authentication Changes

- DATA_ENCRYPT authentication discontinued
  - Only supports DES
  - Replaced by TLS/SSL as strategic network encryption mechanism
- CLIENT Authentication
- FED_NOAUTH DBM CFG
- Change password plugin
  - Default OS plugin provides equivalency. Legacy from earlier implementations

Columnar Enhancements

Columnar Improvements

- Enhanced Insert and Compression enabled by default
  - Page based string compression
  - Trickle Insert performance improvements
  - Deferred Synopsis creation for small tables
- Schema Evolution: DROP and RENAME columns
  - ADMIN_MOVE_TABLE – online columnar table moves
  - V12 will see schema evolution evolve throughout the mod packs
- Compact Varchar Extended to hash Join

ColumnarHigh Cardinality String Data Types

States VARCHAR
001 = Colorado
001 = Colorado
010 = Kentucky
001 = Colorado
011 = Illinois

States have high frequency so are encoded with dictionary

Product Description VARCHAR
Blue dress with unicorns girls size 6X
Red dress with hearts girls size 6X
Red dress with bears girls size 6X
Blue uniform shirt boys size 5
Red uniform shirt boys size 5

Free flowing text stored unencoded

- Frequency-based compression not effective for high cardinality string datasets so percentage of values encoded < 10%
- String data dominates storage cost

Columnar: Page-based String Compression

- Page-Based String Compression Type 2 very effective for compressing high cardinality hex, date, timestamp, numeric data stored in string data types
- Must have <= 16 distinct characters per compressed data page
- Sample data
  - HEX: 59721B038CB9AC5A5DD09055529A64CA4D000000
  - Postal-codes: 95133-7670
  - Telephone Numbers: 1-408-775-6978
  - Date:12/21/2003 (or other date formats)
  - Time:11:32:02 (or other time formats)
  - CDR (Call Detail Record for telecommunications):00650068D34B41799911903603
- Db2 automatically determines which compression scheme to use

Trickle Feed Insert Enhancements (1|2)

- Used only when small number of rows are being inserted (aka data trickling in).
- Inserted rows are split to one or more “insert groups” – still columnar format just inserting more columns per page.
- Number of insert groups depends on types of columns, average length, etc. But generally, will be much less than total number of columns.
- Data going into these insert groups are not compressed.

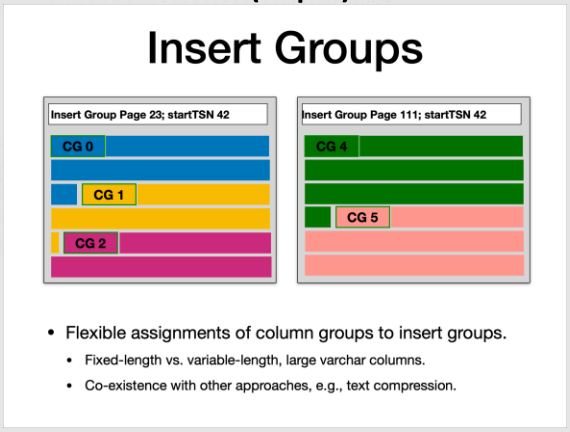

Trickle Feed Insert Enhancements (2|2)

- The insert group pages are almost always temporary -- a window of the most recent 'trickle' inserts. Exception is small tables.
- Insert group pages/rows are automatically moved (aka split) to the single column per page format.
- The split is done synchronously. Triggered as soon as it's predicted that full column group pages can be created.

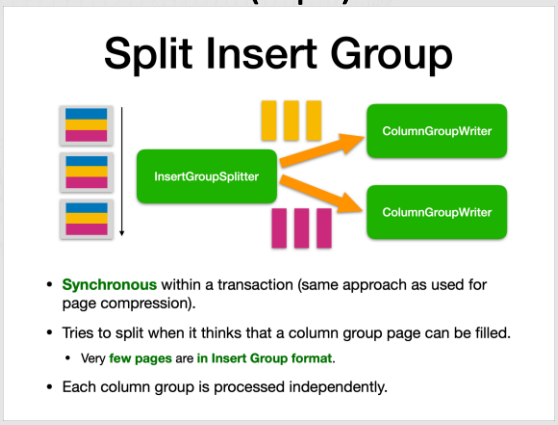

Columnar: Compact Varachar (CVC) Recap

- Improved memory stability and performance of the columnar engine
  - Specifically for workloads involving queries with VARCHAR columns
  - Dramatically reduce Out of Memory (OOM) / -955C errors
  - Improve individual query and overall workload performance
- Reduce need to modify schema to better size varchar columns
- Phased approach:
  - Vector and Work Units ✓
  - Group by and Join ✓
  - Aggregation and OLAP ✓
  - M to N Joins Coming in vNext ✓

Columnar: Compact Varchar Hash Join (1/2)

- CDE Hash Join improvements:
  - Used for Equality Joins in CDE
  - Reduce memory requirement for large “Payloads”
    - i.e. Inner/Dimension tables columns
  - Improve performance for joins with large “Payloads”
  - HSJoin converted to use new “Compact Block Store”
  - Variable length inners and wide inners are stored in the CBS
- CVC Extended to cover all variable length data types
  - VARGRAPHIC, VARBINARY, LOB descriptors

Columnar: Compact Varchar Hash Join (2/2)

- Variable length “Payloads”:
  - “Payloads” are smaller because they are stored compacted
  - Performance improved overall by 2.4x - 2.5x across a wide range of Sort Heap sizes (600K, 300K, 150K, 75K pages)
  - Individual queries saw performance improvements of 5x – 60x!!
  - Overall memory requirements reduced by 75% - 80% (4x -5x “less”)
  - Many queries no longer spill, leading to the huge improvements
- Fixed Length Payloads
  - “Payloads” are the same length, but handled as a “blob” by Hash Join
  - Performance improved overall by 34%– 59% across a wide range of Sort Heap sizes (600K, 300K, 150K, 75K pages)
  - Individual queries saw performance improvements of 3.7x – 4.5x!!
  - Overall memory requirements reduced by 20% - 50% (26% - 2x “less”)

53

Michael m3fp3

Availability Improvements

Online Index Reorg for pureScale (1/2)

- Problem:
  - bWant to rebuild index to achieve optimal index structure and performance without taking an outage or scheduling a maintenance window
  - pureScale index reorg rebuild prevents any access to base table for duration of index rebuild

Online Index Reorg for pureScale (2/2)

- vNext brings parity with non-pureScale Environments
  - index reorg rebuild with read and/or write access permitted on the base table
- Index Reorg Rebuild can be started on any member and the reorg runs on the member it is issued
- Only one instance of reorg rebuild can occur on a table's indexes in the cluster
- Multiple instances of rebuild can happen concurrently against different tables in various modes on a mix of members
- A reorg rebuild "on data partition" locks the partition where as all other reorg rebuilds lock the table

Online Index Reorg for pureScale restrictions:

- Existing non-pureScale index rebuild restrictions and:
  - Circular logging not supported
  - Online add member not supported during index reorg (write access only)
  - Pending add member blocks index reorg (write access only)
  - Extent Movement/Table Space Reclaim paused for duration of rebuild
  - Supports tables with up to 200 defined indexes

HADR Read Access on standby during database upgrade of primary

- HADR now allows read access on standby during primary upgrade
  - New upgrade procedure: upgrade primary database first while keeping standby database at v11.5 to provide read only access
  - After primary database has completed the upgrade, move applications back to the primary then upgrade the standby database
  - Trade-off: read access during upgrade of primary vs no HA protection until standby completes upgrade and reaches PEER state.
  - If multiple standbys, the following order to avoids gap
    - Upgrade principle standby, then primary, while keeping aux standby for read-access
    - Upgrade aux standby after upgrade of primary is complete

Database upgrade can cause a long outage to applications.
Trade-off read access vs HADR protection after upgrade completes.
Upgrade of standby is after upgrade of primary, so during the standby upgrade, unplanned outage of primary will not be able to failover to standby, until standby completes upgrade and reaches PEER state

Backup: Intra-tablespace parallelism (1/2)

Single node with 1FC Card:

- 2.7X faster in extreme skewed case, best scenario
- Flat for balanced case
- Improvement limited by headroom of IO bandwidth on the system

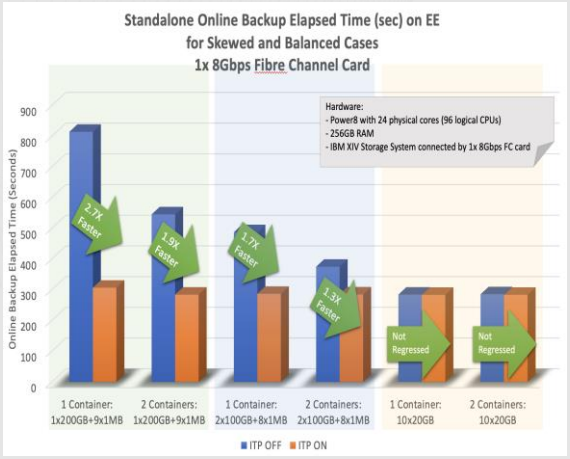

Backup: Intra-tablespace parallelism (2/2)

Single node with 2 FC Cards:

- 3.8X faster in extreme skewed case, best scenario.

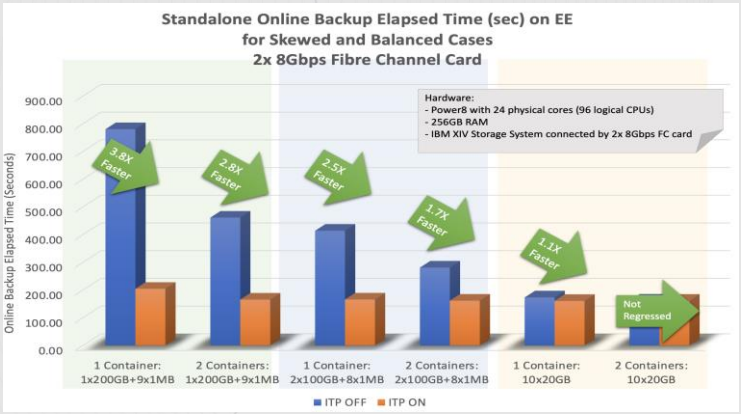

Things to be aware of ….

- minimum OS levels under consideration:
  - RHEL 9.4, UBI9, SUSE 15SP6, Ubuntu 22.04LTS,
  - AIX 7.3 TL2
  - Windows 11, Server 2022
  - Mac (driver only): Sonoma (M1/M2/M3) 64bit
- Linux pacemaker will be the only integrated cluster manager
  - including pureScale
  - AIX remains with TSA
- ICU only latest version will be carried forward
  - mod packs may update ICU versions – ICU 74 targeted default

ICU Older Library Deprecation

When collation is specified, on-disk items are affected:

Indexes on Character data columns
MDC (Multi-Dimensional Clustering) on character data columns
PDF partitioning on Character data columns
MQTs with Character Data columns
Range Partitioning on Character Data columns

The refresh would lead to invalidation and a rebuild of indexes and MTQs
Tables that have some form of partitioning would require tooling to assist or a repartitioning to get data “in the right place”

Mike Springgay is currently the overall Db2 Architect. Prior to that he was responsible for architecting extensions to Db2's Warehouse capabilities, SQL compatibility features, routine infrastructure and client server connectivity.