IBM Db2

Built for the world’s mission critical workloads

Mike Springgay (springga@ca.ibm.com)
Db2 Architect

March 28, 2024

Notices and disclaimers

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.

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.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM.

Not all offerings are available in every country in which IBM operates.

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

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at: www.ibm.com/legal/copytrade.shtml.

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 Preview
- CSB & New Support Models

Agenda

- 2023 Deliverables
- 11.5.9 Overview
- vNext Preview
- CSB & New Support Model

IBM Db2: 30+ years of innovation

1970s

SQL invented by Edgar F. Codd at IBM

1983

Db2 for Z/OS is born (40+ years ago)

1993

Db2 LUW is born (30+ years ago)

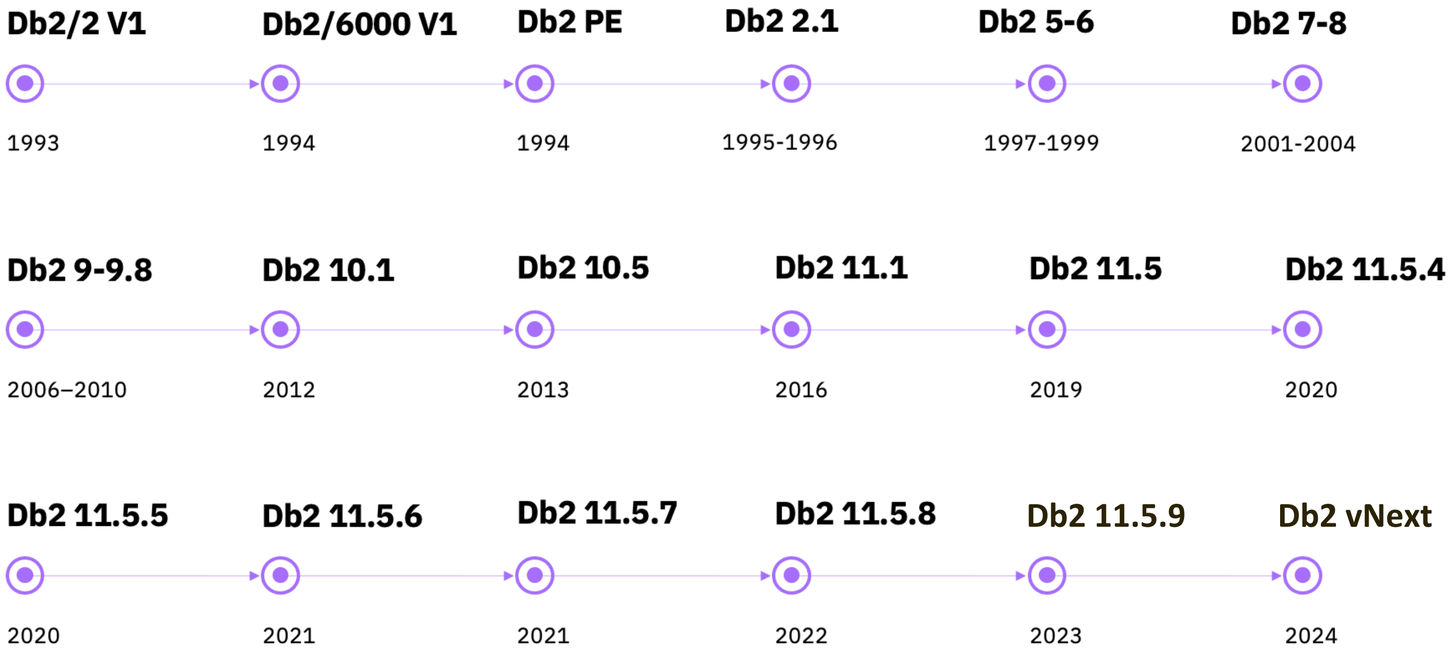

- HADR
- .NET, JDBS, SQLJ, OLE drivers
- VARXXX, XLOBs
- OS/2, AIX, Windows, Linux, Solaris, HPUX
- Text Analytics
- Shared-Nothing Scale-out for OLAP (DPF)
- Granular backups
- Spatial Analytics
- Sequences
- Query Patroller (WLM)
- Db2 Connect
- Data Links
- Data Joiner (Federation)
- Connection pooling
- LDAP integration
- Unicode
- AST/MQT
- Mobile Satellite
- Triggers
- Appliances
- Shared-disk, scale-out for OLTP (pureScale)
- PHP, Perl, Python, RoR, ADO, PL/SQL
- Label-based access control
- Row/column access control
- Roles
- Range partitioning
- pureXML (NoSQL)
- Autonomics
- Multi-tiered storage
- Native encryption
- Audit
- Trusted context
- Multi-dimensional clustering (MDC)
- JSON/BSON support
- Oracle application compatibility
- Compression (tables, index, temp tables)
- Continual data ingestion
- Native OLAP functions
- Native WLM
- Online utilities
- Columnar (BLU)
- Temporal tables
- PostgreSQL compatibility (for NZ workloads)
- Db2 on Cloud (DBaaS)
- External tables
- Event processing
- ML optimizer
- In-database ML
- Advanced log space management
- Graph
- Data virtualization
- Red Hat OpenShift support
- Schema-level security
- Schema-level recovery
- Adaptive Workload Management
- REST APIs
- Multi-tenancy

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 our Db2 offerings to AWS

Other offerings available, including:

- **Db2 RDS for OLTP workloads (managed service)**
- **Db2 Warehouse on Cloud (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

Introducing Amazon RDS for Db2

RUN FULLY-MANAGED IBM Db2 DATABASES ON AWS

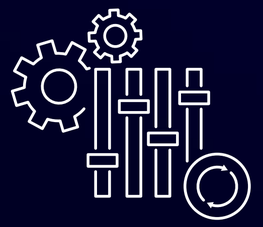

Increase efficiency

Focus on innovation

Reduce costs

Automates undifferentiated Db2 tasks, such as provisioning, backups, patching, and monitoring

Easily migrate existing IBM Db2 databases

Launch Db2 Version 11.5 databases in minutes and enable high availability with RDS Multi-AZ deployment

Bring your own IBM software licenses, supporting Standard and Advanced Editions

Supports transactional, mixed and analytics workloads, including Oracle compatibility

Overview of Amazon RDS for Db2

Easy to administer

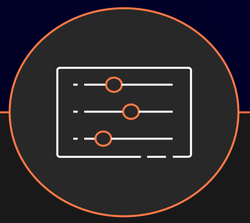

- Create database with few clicks in few mins
- No infrastructure provisioning, software installation, or patching
- Built-in monitoring

Performant and scalable

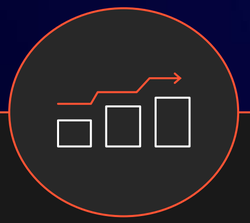

- Power your database with push-button compute scalability
- Auto scale your storage

Available and durable

- Achieve high availability with Amazon RDS Multi-AZ deployments
- Automated backup, snapshots, and failover

Secure and compliant

- Protect data with encryption at rest and in transit
- Achieve compliance with key industry compliance programs

Based on IDC Study-39% lower database operation costs

RDS for Db2 2024 Releases since launch

03/20/2024
AWS License Manager
AWS License Manager now allows you to track IBM Db2 licenses on Amazon Relational Database Service (RDS)

03/11/2024
Support for M6i and R6i in additional AWS Regions
Amazon RDS for Db2 expands support for M6i and R6i in additional AWS Regions

03/11/2024
Support for X2iedn instances in additional AWS Regions
Amazon RDS for Db2 expands support for X2iedn instances in additional regions

03/06/2024
Support for io2 Block Express for sub-millisecond latency and 99.99% durability
Amazon RDS now supports io2 Block Express for consistent sub-millisecond latency and 99.999% durability

02/15/2024
Support for Audit Logging
Amazon RDS for Db2 now supports audit logging

01/29/2024
Support for EBCDIC Collation Sequence
Amazon RDS for Db2 now supports EBCDIC collation sequence

01/19/2024
Support for Cross-Region Automated Backups
Amazon RDS for Db2 now supports Cross-Region Automated Backups

01/19/2024
Amazon RDS for Db2 now supports local time zones
Local time zone support for Amazon RDS for Db2

As of April’24

Amazon RDS for Db2 resources

- AWS News Blog
- Service page
- Technical Documentation for RDS Db2

Contact

Karthik Gopalakrishnan Senior Product Manager
carthik@amazon.com

Db2 pureScale on AWS Marketplace

Available since July 2022

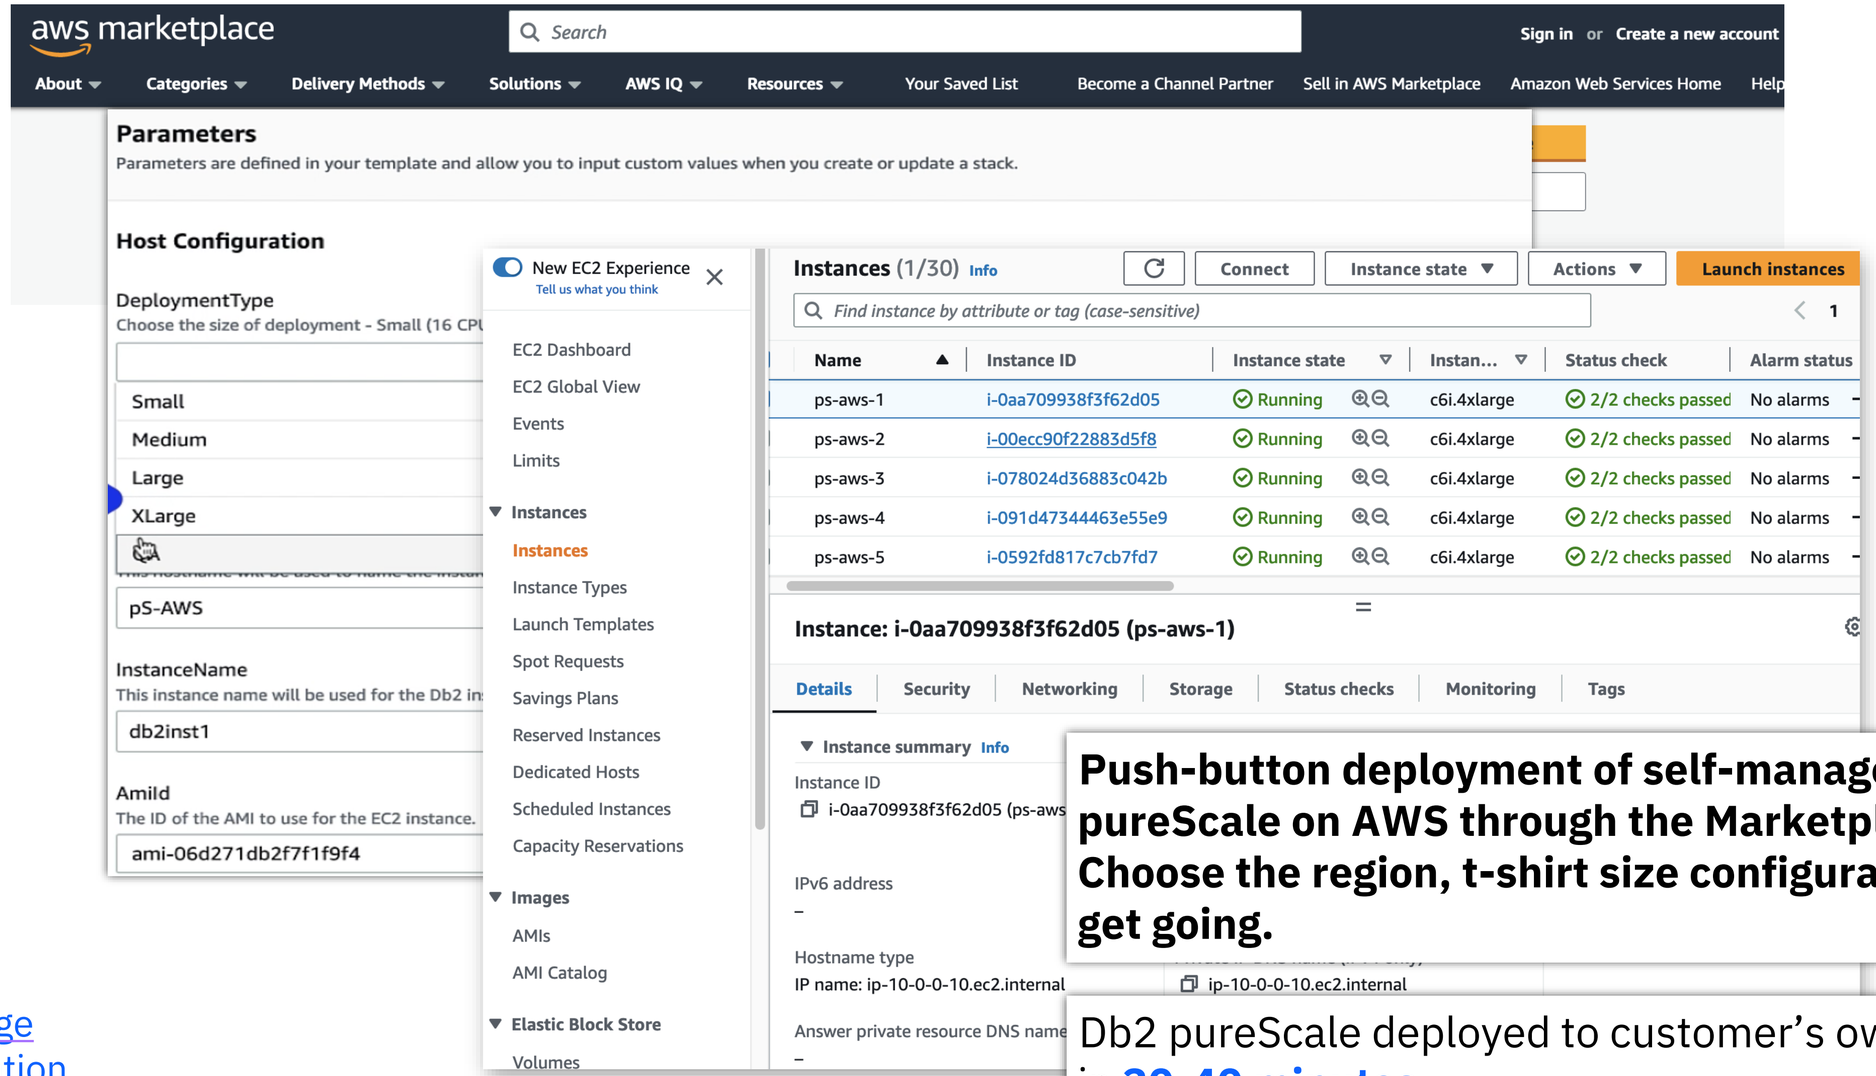

Push-button deployment of self-managed pureScale on AWS through the Marketplace. Choose the region, t-shirt size configuration and get going.

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

- Product Page
- Documentation

IBM Db2 Warehouse on Cloud Gen3

Introducing our 3rd generation of cloud data warehouse (Gen3) built for your always-on workloads

1 Skinny cost, big performance

Highly compressed **data running natively on highly-resilient cloud object storage**

↓ **storage costs** over traditional block storage

↑ **performance** with advanced caching

2 Responsible data sharing

Built on decades of **leadership in data privacy, protection and security**

Support for **open table and data formats** including Iceberg, Parquet, ORC, CSV for maximum collaboration

3 Scale analytics across the enterprise

Seamless integration and common metadata service with the new watsonx.data service

Singular view of your entire analytics estate

Shift workload and **run fit-for-purpose engines** that best meet your needs

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

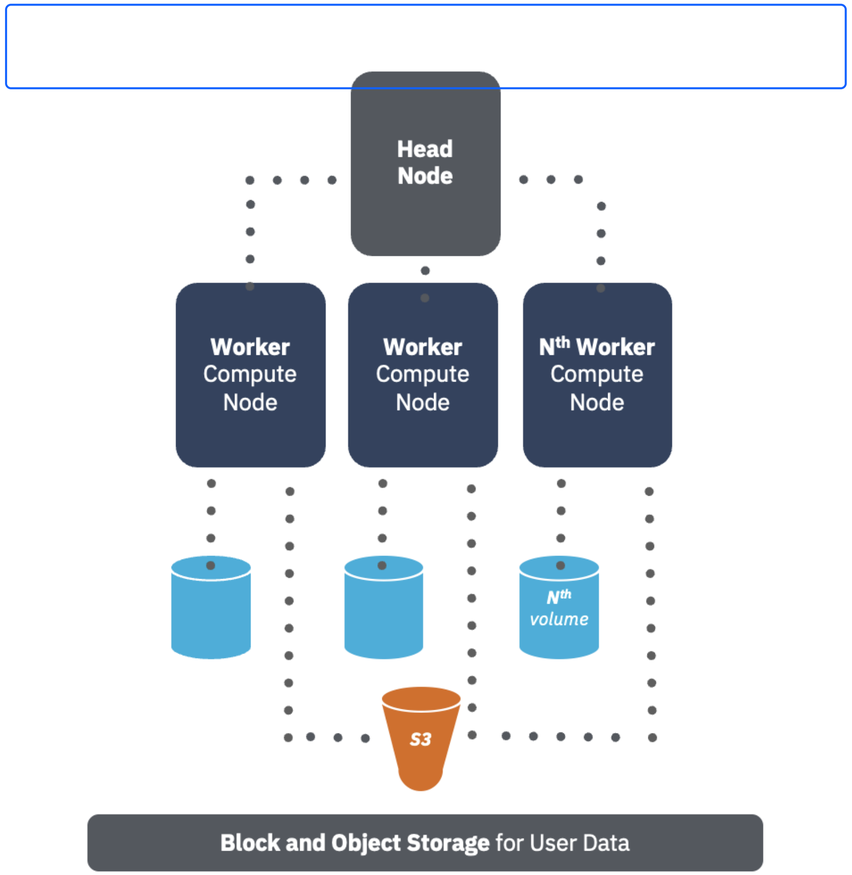


```
TBSP_NAME          MEMBER TBSP_TYPE  CACHING_TIER
------------------ ------ ---------- ------------
...
OBJSTORESPACE1          0 DMS        ENABLED
OBJSTORESPACEUTMP       0 DMS        ENABLED
```

Db2 Warehouse on Cloud Gen3

34x

Less expensive to host Db2 data on object vs block storage[1]

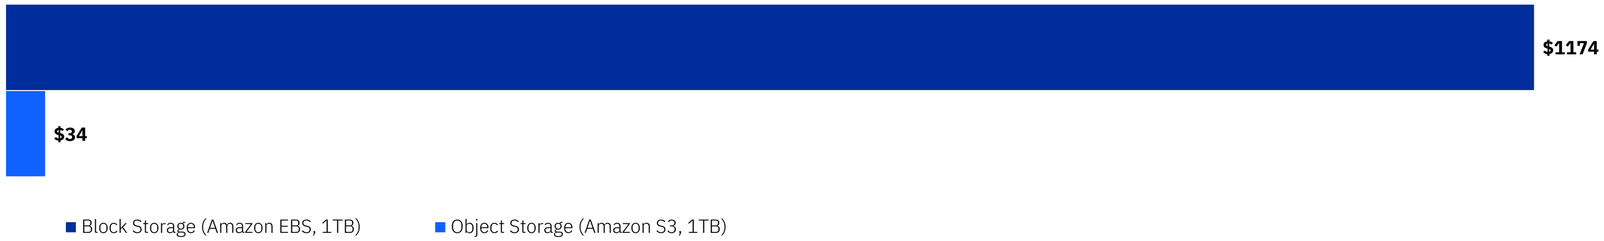

Block Storage (Amazon EBS) vs Object Storage (Amazon S3)
Cost reflects Amazon's list price for block storage (various tiers & IOPS levels) required to host an incremental 1TB of Db2 data

[1] Block vs Object Storage comparison depicts difference between published prices for Amazon EBS 1TB of io1 at 6 IOPS/GB (and additional tiers to support Db2 data) vs Amazon S3. This metric is not an indicator of future storage pricing for Db2 Warehouse Gen 3.

Db2 Warehouse on Cloud Gen 3
Performance

4x

Faster query performance

When Gen3 is compared against the prior generation

Note: Lower number is better

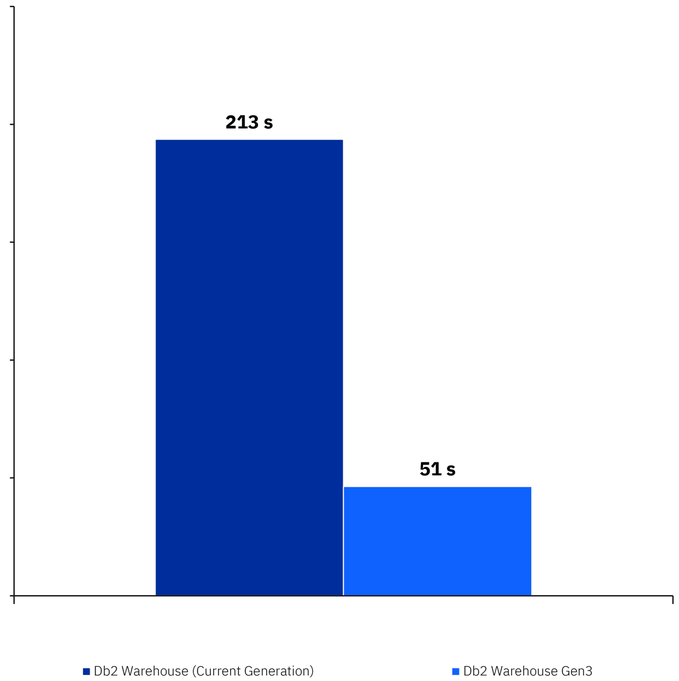

IBM Big Data Insight (BDI) Benchmark
simulates real-world deep analytics, reporting, and dashboard queries

10TB Db2 data warehouse
residing either on block storage (current generation) or object storage (Gen3)

16 concurrent users
running a variety of ML, reporting, and dashboard queries

Cold cache start
for both the in-memory buffer pools or the NVMe cache

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

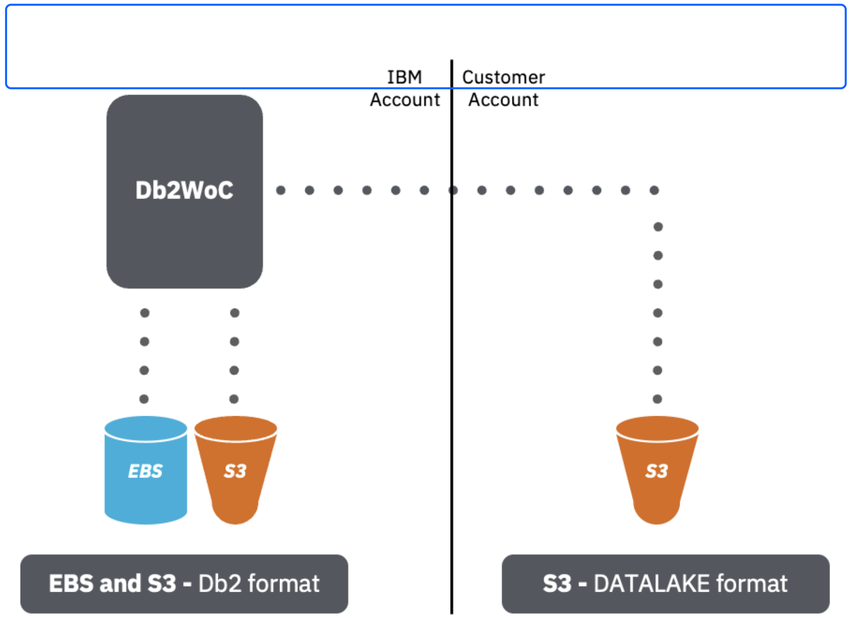


```
CREATE DATALAKE TABLE my_datalake_table
  ( id INT,
    name VARCHAR(8)
   )
   STORED AS PARQUET
   STORED BY ICEBERG
   TBLPROPERTIES ('external.table.purge'='true')
   LOCATION 'DB2REMOTE://mybucket-alias//
                         my_datalake_table';
```


Db2 Warehouse Architecture Solutions

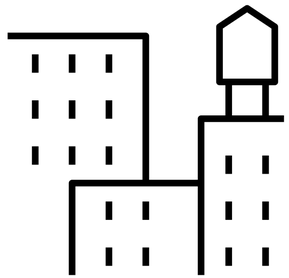

On-premises offerings that provide a smooth upgrade path from PDOA/IIAS (Sailfish) appliances[1].

- Maximizes simplicity, minimizes time to value, and ensures fully supported environments
- *Optimized for performance* and scalability
- **Prescriptive** configuration of IBM hardware, Db2 Warehouse software and Technology Expert Lab services
- IBM built and lab tested
- Includes *Gold-glove Installation* and *Solution Support* Services with a single point of entry.

[1]As opposed to Appliances where Db2 Warehouse licenses are an integrated part of the solution, in a Reference Architecture Db2 License can be separated from the Solution and moved to another infrastructure.

[1]Clients can also procure the prescribed hardware, software, and Expert Lab services a-la-carte to create a roll-your-own (RYO) or Do-It-Yourself (DIY) self-managed on-premises deployment of a Reference Architecture. However, for RYO/DIY self-managed deployment of a Reference Architecture Solution, IBM will provide only component level installation and support services and **Gold-glove Installation** and ***Solution Support Services*** **will not be available to RYO/DIY Reference Architecture Solutions.**

Db2 Warehouse Power 10 Architecture

On-Premise HW Stack

Servers:

- OpenShift Management: 3x Power 1022s
- Workers: From 3 to 80 - IBM Power 1022

Storage:

- IBM FlashSystem 7300 for Tier-1
- IBM FlashSystem 5200 for Tier-2

High Speed Networking

- IBM 100 GB Ethernet Switching
- IBM Storage Networking SAN

Management:

- IBM HMC

Warehouse Software Stack

- RHEL OpenShift
- Db2 Warehouse on OpenShift
- IBM PowerVC
- IBM Spectrum Scale Data Management Edition
- PowerVM Enterprise

IBM Services

Pre-sales:

- Size an appropriate configuration for customer's existing workload.
- Collaboration with business partners.

Install:

- Build the infrastructure hardware stack and pre-load the OCP/SS software stack in IBM MFG
- Gold-glove installation of the infrastructure hardware/software stack at customer's premises
- Remote installation/customization of Db2 Warehouse Software on the infrastructure at customer's premises
- Migration Assessment session with the customer

- **Experience:** A balanced resilient architecture for maximum reliability and performance
- **Flexibility:** The data warehouse starts small, but grows as you do
- **Efficiency:** E-config enables the full stack pre-selected and configured as a "solution edition" set-up
- **Storage:** IBM FlashSystem adds simplicity and consolidation, capable of delivering more work with fewer drives and lower cost

PDOA/IIAS vs. P10 Cloud Rack for Db2 Warehouse

IBM Integrated Analytics System

- Power 8 2U Servers
- FS900 Flash Modules
- 4x10 Gb bonded network
- 16 Gb HBAs
- Db2 Warehouse

P10 Cloud Rack for Db2 Warehouse

- Power 10 2U Servers
- FS7300 Flash Modules
- 100 Gb bonded network
- 32 Gb HBAs
- Modernized Db2 Warehouse on OpenShift

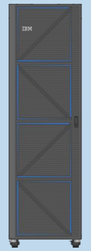

IBM Integrated Analytic System

- 8 Data MLNs per Server
- 3 Power8 Cores per MLN
- 1.6TB Usable Tier 1 space per MLN
- 51GB Memory Per MLN

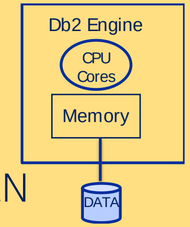

P10 Cloud Rack for Db2 Warehouse

- 8 Data MLNs per Server
- 4 Power10 Cores Per MLN
- 3TB Usable Tier 1 space per MLN
- 100GB Memory Per MLN

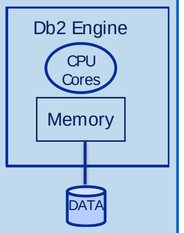

Fusion HCI System

Next generation hyper-converged infrastructure delivers the simplicity of OpenShift on virtual machines with the benefits of OpenShift on bare-metal

OPENSHIFT

IBM Storage
Fusion

Maintenance
Monitoring
Automation

Hardware	Operating System	Fusion data services	Lifecycle management

Accelerate container projects

- Fast turnkey deployment of bare-metal OpenShift on x86 compute
- Fast NVMe storage
- Container data services for persistent volumes, backup / restore, Metro DR, and Regional DR
- Engineered for resiliency and performance
- Validated with Db2 Warehouse

Cost Avoidance

- Reduce OpenShift subscription costs by up to 90%
- Eliminate hardware procurement, system design, test, and debug
- Self-contained networking simplifies deployment and improves performance
- Automated lifecycle services reduces administrative costs

PDOA/IIAS vs. IBM Storage Fusion HCI (x86) for Db2 Warehouse

IBM Integrated Analytics System

- Power 8 2U Servers
- FS900 Flash Modules
- 4x10 Gb bonded network
- 16 Gb HBAs
- Db2 Warehouse

Fusion for Db2 Warehouse

- 1U, 64 Core Lenovo Servers
- 7.68TB NVMe PCIe Gen4 drives
- 25 GbE bonded network for client connectivity
- 100 GbE network for backend
- Db2 Warehouse OpenShift Container

IBM Integrated Analytic System

- 8 Data MLNs per Server
- 3 Power8 Cores per MLN
- 1.6TB Usable Tier 1 space per MLN
- 51GB Memory Per MLN

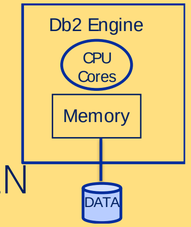

Fusion for Db2 Warehouse

- 8 Data MLNs per Server
- 7 x86 Cores Per MLN
- 7.0 TB Usable Tier 1 space per MLN
- 100 GB Memory Per MLN

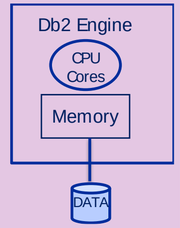

Business Transformation with Db2 Warehouse – On-premises P10 and X86 Reference Architecture Solutions

IBM Power Systems

x86

IBM Integrated Analytics System (IIAS)

IBM PureData System for Operational Analytics (PDOA)

Migrate to...

1. Db2 Warehouse SaaS

Fully managed, high performance, cloud-elastic data warehouse

↳ Runs on IBM Cloud, AWS

↳ Available in US, EU, UK, JP, AP commercial regions

2. Db2 Warehouse Self-managed

Self-managed, data warehouse on cloud provider managed Kubernetes or OpenShift infrastructure.

↳ Runs on Amazon EKS, Amazon ROSA, Azure AKS, Azure ARO

↳ Available worldwide

3. Db2 Warehouse Cloud Rack

On-premises, OpenShift high performance, data warehouse deployable on Power 10 or X86 infrastructure

↳ Runs on POWER10 or x86 reference architecture

Agenda

- 2023 Deliverables
- 11.5.9 Overview
- vNext Preview
- CSB & New Support Modelh

Db2 11.5.9

Hybrid & Multi-cloud

Containerization:

- 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.
- Support for Native Object Store – Db2 Warehouse only
- Support for ODF – Db2 Warehouse MPP only

Application Development

- Federation Enhancements
  - Support for AWS Aurora PostgreSQL
  - Support for PostgreSQL 12-15
  - Addition pushdown of:
    - timestamps to MongoDB
    - strings including trailing blanks comparison
    - numeric functions
    - date/time functions
    - CSE
  - Watsonx.data support via JDBC

- Client Enhancements:
  - Uplift of levels

Columnar Enhancements & Db2 Warehouse

- Support for tablespace storage to be on S3 (aka Remote Table spaces)
- Support for Open Data Formats: Iceberg, Parquet, AVRO, ORC
- Reference Architectures for Db2Wh on Power and x86
- Online Synopsis Table Rebuild

* Delivered (1Q 2023) in latest containers

Db2 11.5.9

Availability

- Ability to protect any History File entry from pruning
- Improved Replay Performance
- Db2 HADR with Pacemaker on Google Cloud Platform validated with SAP Workload

pureScale

- RHEL 8.8 support (might be post GA), SLES 15 FP5 support (might be post GA), 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

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)
- Advanced Data masking support
- Audit Logs to Object Store
- Restricted TCPIP Listener Mode
- ECMTS, Java, Gskit, TSA, GPFS and Pacemaker stacks updates

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

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*

Agenda

- 2023 Deliverables
- 11.5.9 Overview
- vNext Preview
- CSB & New Support Model

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

Agenda

- 2023 Deliverables
- 11.5.9 Overview
- vNext Preview
- CSB & New Support Model

Db2 11.5: CSB vs. past “Special build”

Similarities:

- both are how customers get APARs fixes “asap”
- both are identified by a special build ID (versus change in mod or FP level)

Differences (→ Why is the CSB much better?).

- C stands for Cumulative : The CSB is all APARs (including security) fixed since the last mod pack.
- Any/all customers pick up the same stream, the previous special build model resulted in “1-of” builds for customers
- Security SB is just a CSB w/ the latest security fixes-- they are part of the cumulative stream.
- Being a continuous stream, it gets continuous testing including pipeline, functional, performance and stress; previous SB had minimal testing (decided per fix).

When a customer is installing a new Db2 instance or upgrading to version 11.5, it is recommended to choose the most recent CSB stream, which will automatically include the latest updates.

Note that we retain the capability to perform individual, one-off updates (referred to as unofficial SB). However, these are primarily employed for debugging purposes or for partial fixes that are not intended or prepared for widespread use.

Db2 11.5: Going Forward

Rely on CSB stream, no fix packs.

Fix packs simply package many APARs together.

👍 Good:

- as a periodic encouragement to move up; “look here’s a bunch of great fixes”
- to give a target level/date to help support/justify a plan to update (to management, other teams, etc.)

⚠ But:

- often dates don’t align with when customers want to do a new install, or upgrade or even pre-planned updates.
- it is stale the moment it ships, let alone by the time the next one comes out

  → e.g., mod 11.5.8 released it’s first CSB with ~50 additional APAR fixes the same day
- same issue tends to be hit multiple times till next FP out (which can be months)

Db2 11.5: Going Forward

Choosing the most recent CSB addresses the majority of the problems.

- Selecting a particular CSB level at regular intervals (currently within approximately a 3-month timeframe)
- Giving preference to security-focused levels that are documented as "recommended" levels.
- Encourage customers to transition to these recommended levels, much like they would for fix packs.
- Customers would still have the option to choose the latest level when they decide to implement it.

The rollout for this is still in progress! We have one more mod pack on it's way!

KIs requires updates to include CSB level they are fixed in.

- Ensure KIs are published and listed to match CSBs available * DONE *
- Ensure the existence of a document, like the security vulnerabilities page, to direct users to the latest and 'selected' CSB levels. * DONE *

Certain fixpack-based documentation requires updates.

** see https://www.ibm.com/support/pages/node/6856211*


Q
Questions
&
Answers
A

Thank You

Mike Springgay

STSM, Db2 Architect

springga@ca.ibm.com