

IBM Db2 Update

April, 2021

Piotr Mierzejewski

Director, Db2 LUW Development



Pillars of Db2 Strategy

Build for and Powered by AI/ML

- Python UDF
- In-DB Machine Learning
- Db2 Graph – Tech Preview
- ML Optimizer – Tech Preview

Hybrid Cloud Ready and Developer Inclusive

- REST API
- Visual Studio Code Extension
- Db2 Operators
- Click2 Containerize


Enterprise Grade & Resilient

- Faster Database Activation
- Compact VARCHAR
- Pacemaker/CoroSync Automated HADR
- SSO with JWT Token

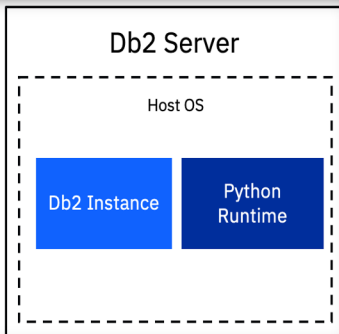
Build for and
Powered by AI/ML

Python UDF : Scoring Python Models via Db2

Watson
or
Open Source



Export the ML pipeline by
serializing *python joblib*



5x Speed up of Inferencing with IBM Db2

Model & Dataset

Logistic Regression (scikit-learn) Model

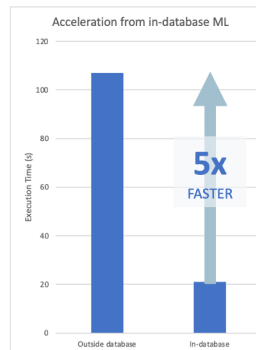
Batch size: 800k rows

Scoring on a Separate System

Data retrieved (over network) from Db2, scored, and written back to Db2: **1m47s**

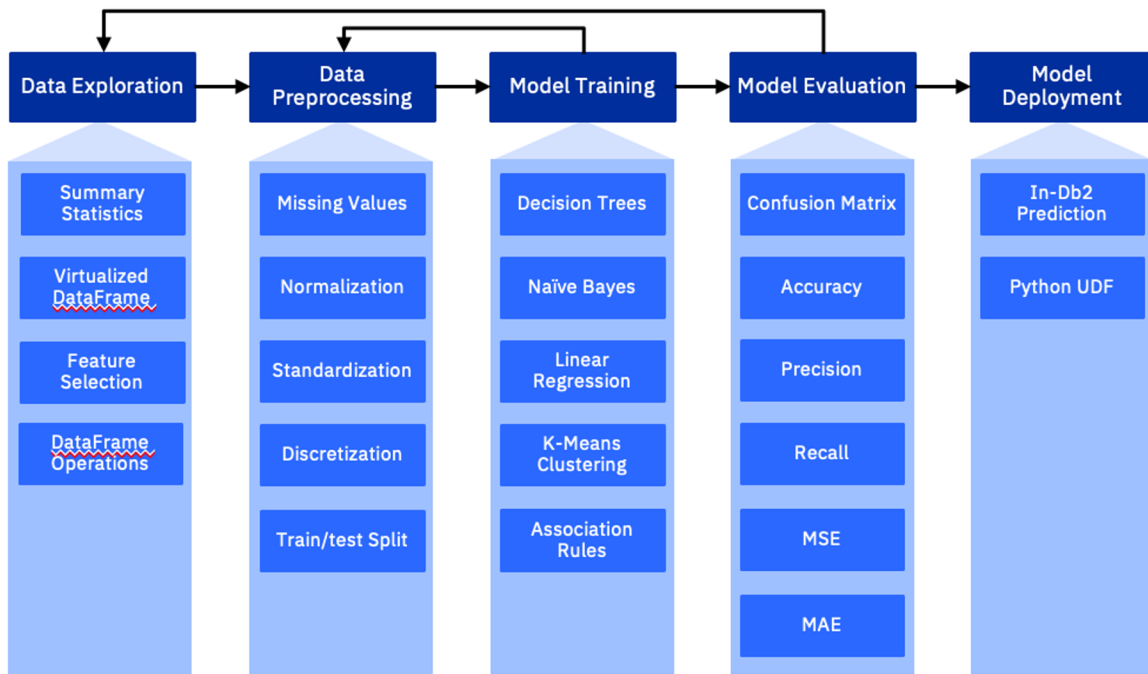
Scoring with Db2

Data retrieved, scored, and written back inside Db2: **21s**



Build for and
Powered by AI/ML

In DB Machine Learning



- Speed
- SQL Interface
- Distributed Training
- Security
- Scalability



Platforms
x86 Linux
PPC LE
zLinux (Linear Reg. unsupported)

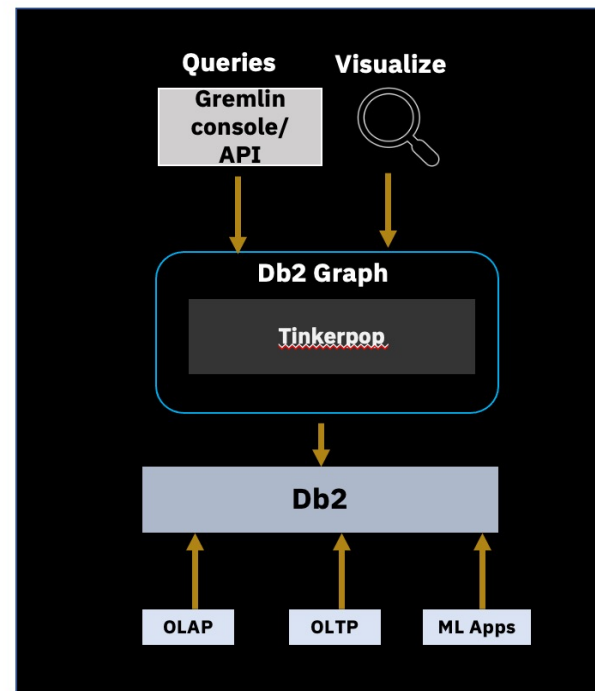
Db2 Graph - Tech Preview

Your data remains in Db2

- Continue to read and make updates to your data, without any impact on performance
- Leave your applications as they are

Simply “overlay” a virtual property graph over your Db2 data

- Uses the referential information to automatically create the Graph schema
- You can map tables or create views to map them into graph vertices and edges
- Add more relationships by editing the Graph schema, as long as the underlying data exists in Db2



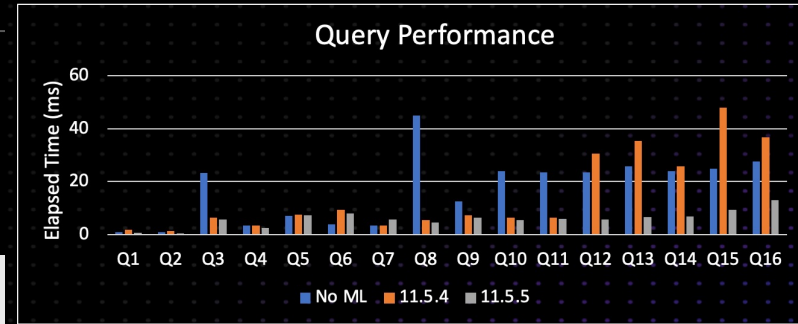
Build for and
Powered by AI/ML

ML based Query Optimization

Get intelligent query processing with ML-based optimizer that uses ML and Neural networks to learn from past experience

- Uses ML to learn from experience and **improve decisions using feedback** from query execution and statistics
- Applies ML to better **estimate cardinalities and predicate filtering**, which is key to better query performance
- Optimizes** memory and compute capacity resources and database resources


Get up to **10x** better query performance with Db2 v11.5.5



Without machine learning



With machine learning

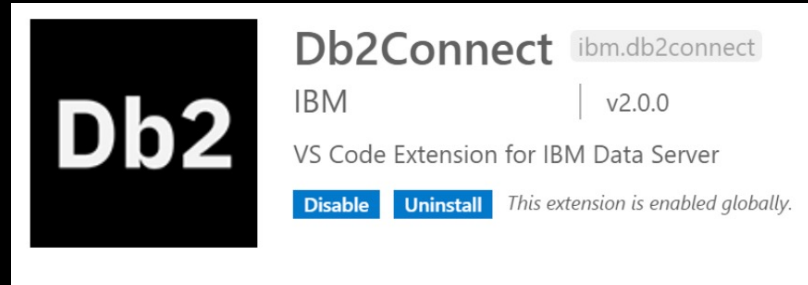


Hybrid Cloud Ready
and Developer
Inclusive

Visual Studio Code Extension & Drivers

New and Improved Db2 VS Code Extension with features including:

- Connection Manager for your connections
- Explore Database objects via UI Explorer
- Create/Open/Edit in local machine or from database
 - SQL Scripts
 - Stored Procedures
 - UDFs
 - Triggers
 - Views
 - Alias
 - Table
- Deploy Stored Procedure, UDFs, Triggers to the database.
- Export Results as CSV, JSON and XML.
- Keyboard shortcut support for different activities



Db2Connect `ibm.db2connect`
IBM | v2.0.0
VS Code Extension for IBM Data Server
[Disable](#) [Uninstall](#) *This extension is enabled globally.*

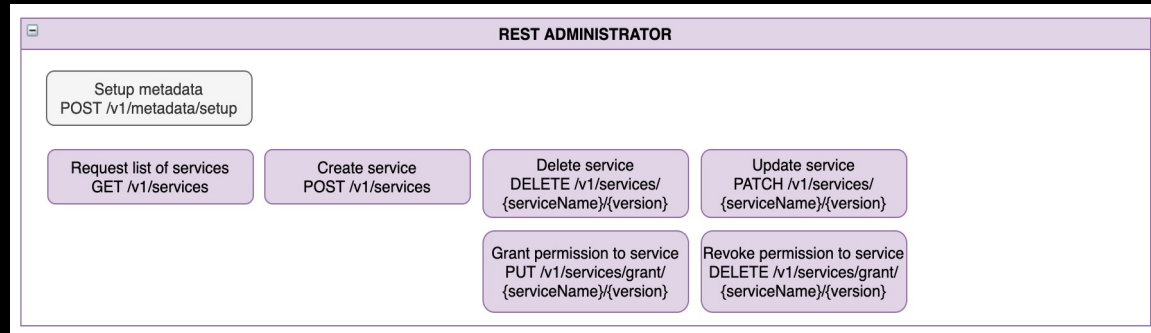
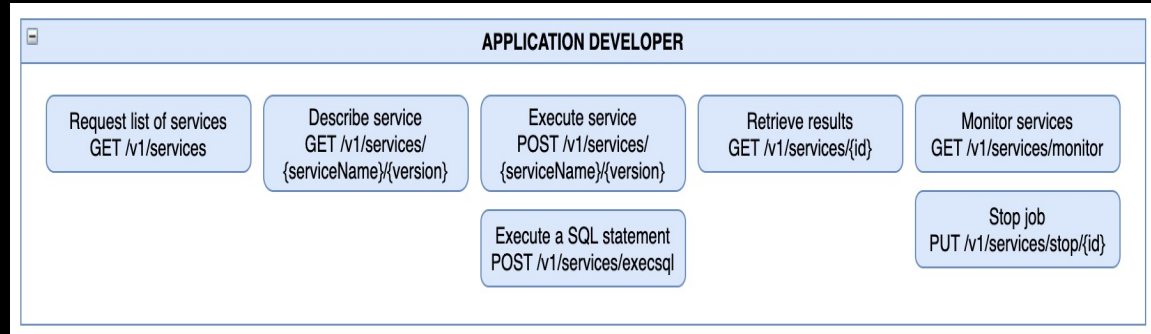
Driver	Updates with 11.5.5
.NET	.NET 3.1 .NET 4.8 .NET5 (Dec 2020)
Python	Python 3.9
Java	JDBC 4.3 and Open JDK 13 Runtime
Node.js	15.x
Ruby	2.7.x

Supported Framework	Updates with 11.5.4
Django	Django 2.2, Django 3.1 (Jan 2020)
Rails	Rails 6.0
.NET Entity Framework	6.4 (Dec 2020), .NET5 Core(Dec 2020)

Hybrid Cloud Ready
and Developer
Inclusive

Db2 REST API Service

- Provides the application developer with APIs to **create**, **discover** and **execute** their own end-points, referred to as services.
- Each **developer-defined service** is associated with a single SQL statement.
- Services can be executed synchronously or asynchronously, as a job that supports result paging.
- **Db2 REST** is deployed as a docker container supporting Db2 11.5 (and 11.1).
- Authenticated users can use the REST end-points from any REST HTTP client **without installing Db2 drivers**.
- The Db2 REST server accepts an HTTP request, processes the request body and returns results in **JSON** format.



The Db2 Cartridge

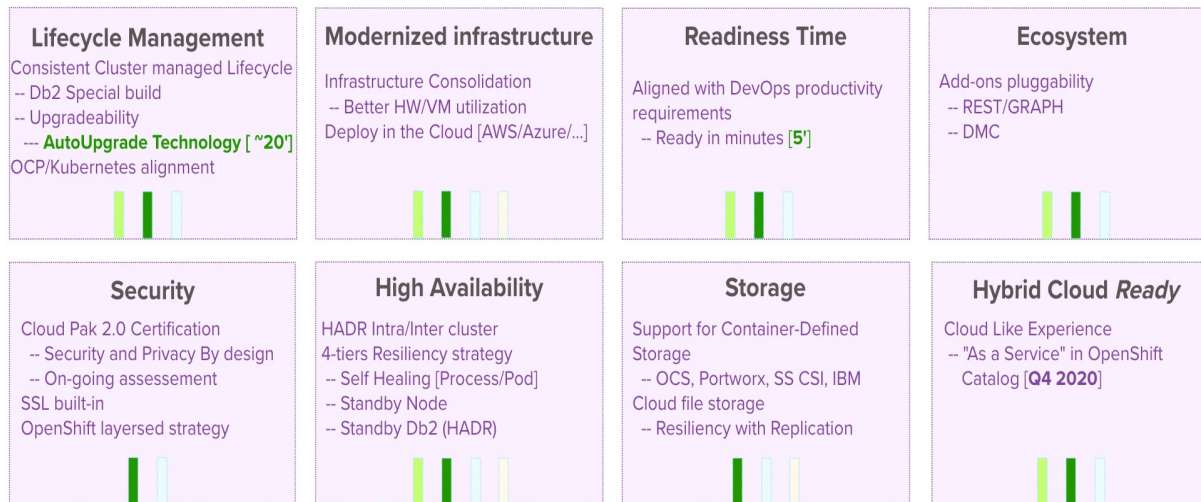
Cartridge VALUE MATRIX

Time-Saving

Cost-Saving

Modernize

Extended HA



Certification standards

- Runs on OpenShift
- Certified operators
- Fully containerized
- T1-T3 support
- Vulnerability scans

Capabilities level ⓘ

- Basic install
- Seamless upgrades
- Full lifecycle
- Deep insights
- Auto pilot

All Performance Results for **Db2U on RHOS**

Metrics	Db2U on RHOS
Install Time	[Automated & includes instance/Database Creation] Time: ~ 2' to 4' (Db2)
Upgrade Time	[Automated] Time: 20'
<u>HADR</u> Takeover time	[Automated] Time: 2-4' minute
Recovery on Db2 Failure	[Automated] Time: 3 minutes + Partial Crash Recovery time* (observed < 1 min on SMALL)
Recovery on Hardware Failure	[Automated] Time: 6 minutes + Partial Crash Recovery time*
Scaling Compute	Time: 2 minutes

- Partial Crash Recovery time*: Usually complete in minutes with worst case at 20/30 minutes

Reference architecture -

<https://www.redhat.com/en/resources/IBM-Db2-Warehouse-MPP-on-OpenShift-Container-Storage-detail>

Hybrid Cloud Ready
and Developer
Inclusive

Db2 C2C (Click2Containerize)

- It's **Containerization**

- *NOT Migration*
- No Backup Restore
- No Export / Import
- No Reconfiguration
- Data Handled Securely

The screenshot shows a web application window titled "Click to Containerize" with the IBM logo in the top right corner. The interface is divided into six main sections, each with a right-pointing arrow at the bottom right:

- Containerization Overview:** Describes the tool's purpose and provides a link for more details.
- Register Database:** Explains the registration process and lists required information like database name, host, and port.
- Database Utilities:** States that configuration is local and provides a link to view catalogued databases.
- Analyzer Settings:** Details how OpenShift and Dictionary settings can be modified and provides a link to change reference databases.
- Containerize Db2:** Describes how to generate a shell script and includes a note about requiring administrator access.
- Additional Support:** Lists contact information for Phil Downey and George Baklarz, and provides a link to the OpenShift Lab.


An "Exit" button is located at the bottom left of the interface.

Hybrid Cloud Ready
and Developer
Inclusive

Db2 C2C (Click2Containerize)

Db2 On Prem

Db2 on Linux



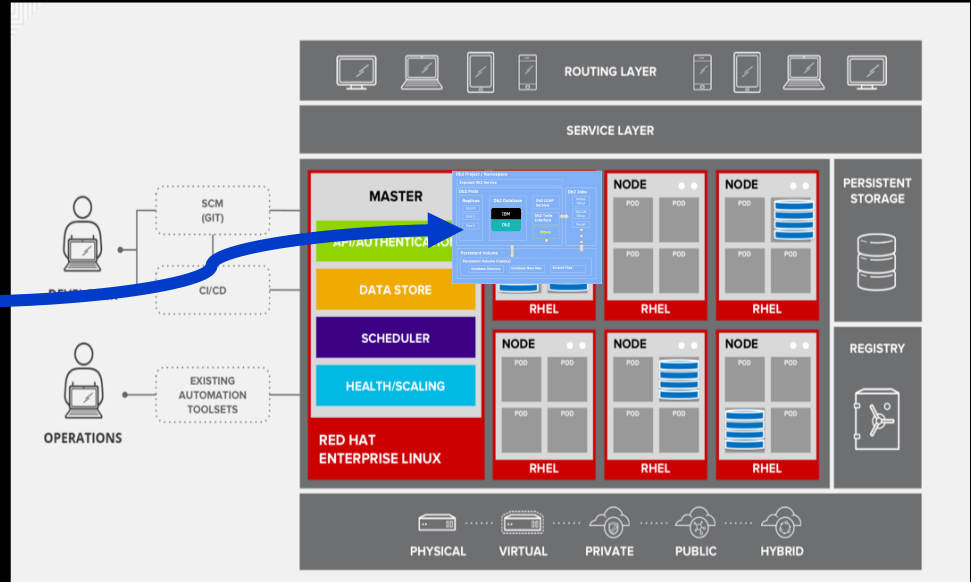
Oc Client

Internet/
Network

Tested over 15,000 km
Internet connection

- Logs
- Mirror Logs
- Storage Path
- Storage Path X
- External UDFs
- DB Config Files
- KeyStore

Tunable Parallel Data Transfer
Data Compression
Fully Auditable

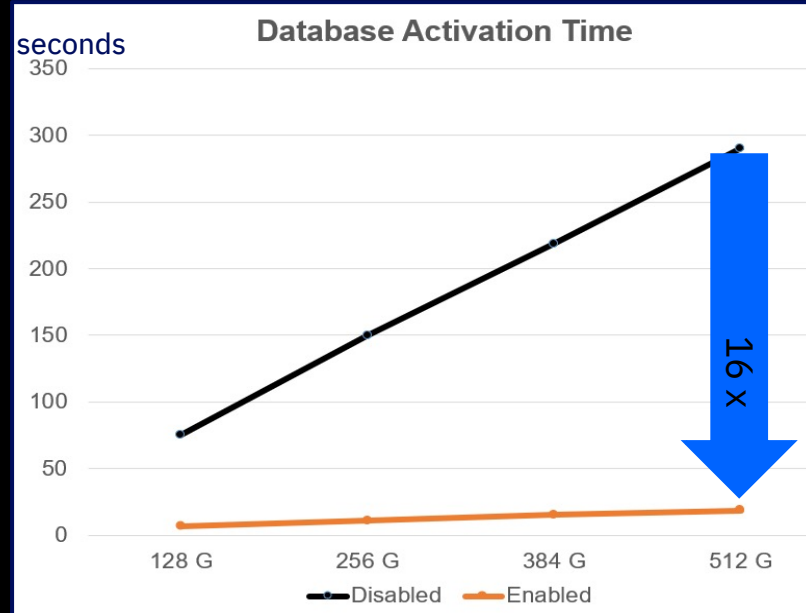


Containerized Db2 Hybrid Cloud Service

Enterprise Grade &
Resilient

Faster Database Activation

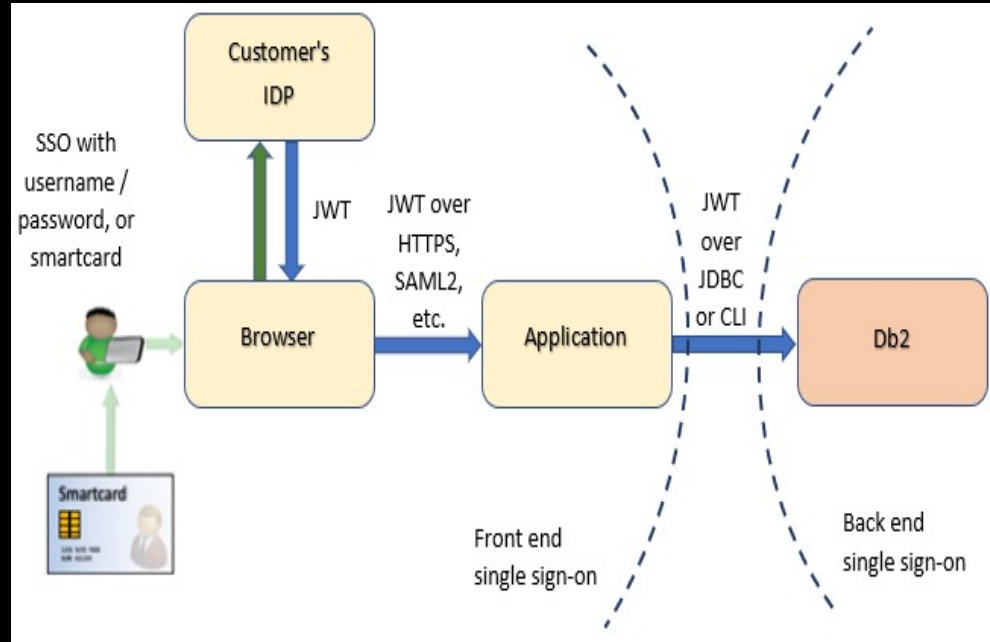
- Applicable in all – Upgrade ,
Crash Recovery scenarios
 - The buffer pools are committed
to memory using an
Asynchronous Process
 - Tests with SAP Db2 database on
AIX system
- ➔ Up to factor **16x faster**



Enterprise Grade &
Resilient

Single Sign-ON with JWT (JSON Web Token)

- **Authentication without username/password via JSON Web Token (JWT)**
- **Identity Provider (IDP) signs the token that's passed on to Db2**
- **Db2 is configured to "trust" IDPs and use their public key to verify the token**
- **If the token is valid, Db2 uses the identity within the JWT for authentication**



Enterprise Grade & Resilient

HADR with Pacemaker & Corosync



V11.5.4.0 Technical Preview

- Cluster manager-aware integrated Db2 commands
- Integrated data collection via db2support
- New cluster manager configuration utility – db2cm
- Multiple instances & databases support
- Enhanced quorum type support with QDevice
- RHEL 8.1, SLES 15 SP1 support on Intel and Linux on IBM Z
 - Validated on AWS with RHEL 8.1



V11.5.5.0

~~Technical Preview~~  **GA**

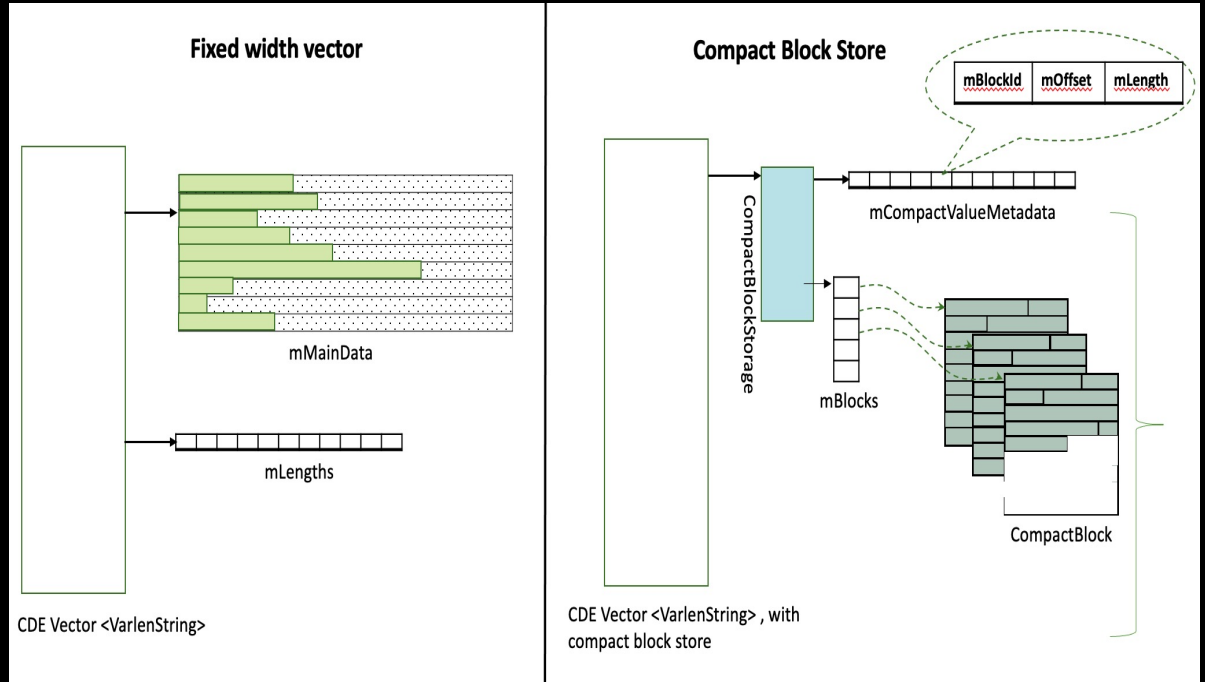
- Multiple Standby Support
- Fast redeployment via import & export support
- Two node support with fencing on AWS



Future GAs

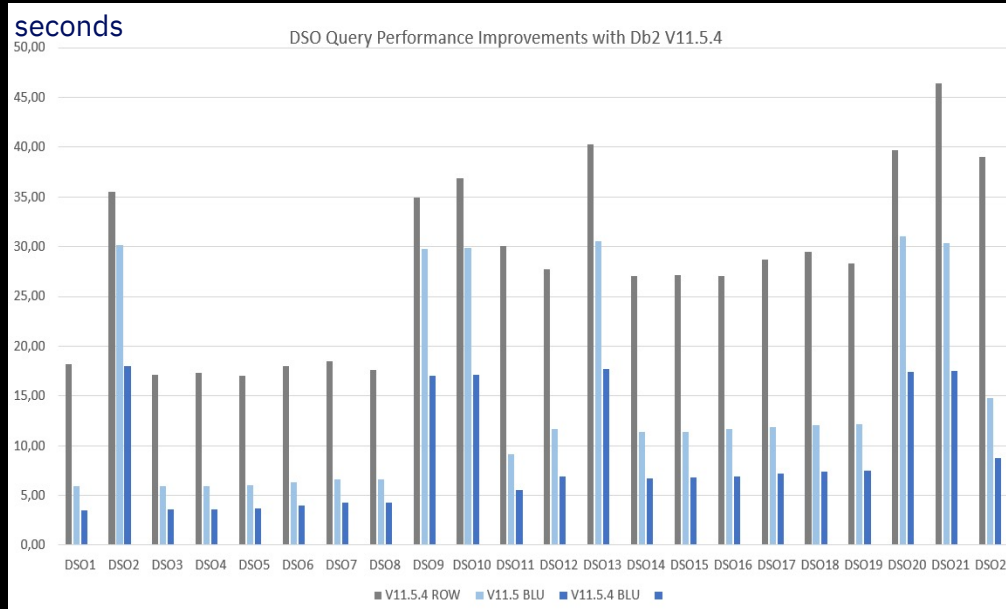
- Integrated bundling and install
- Support on POWER (AIX and PPCLE)
- Active-Passive HA Configuration
- Mount point monitoring
- DPF HA
- More Cloud validations
- Majority Quorum

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



Db2 runtime optimizations for Group By and HSJN with string data types

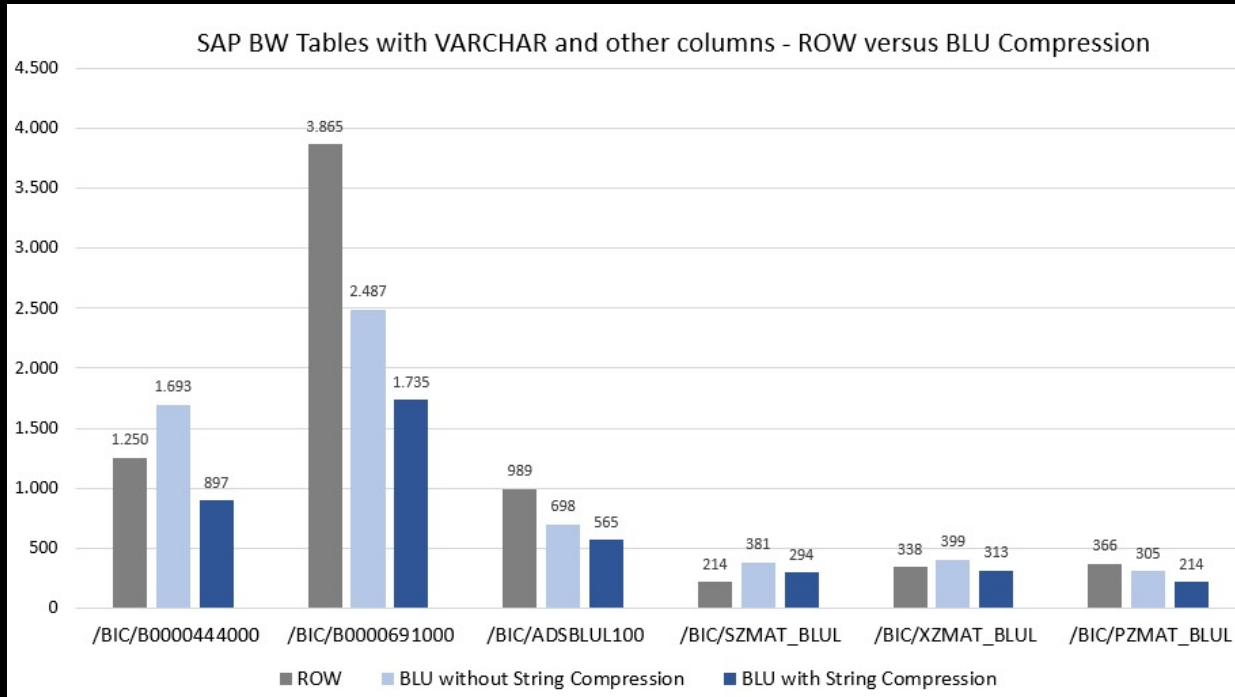
→ SAP DSO queries run between **factor 1.6 and 1.8 faster** with Db2 11.5.4 compared to 11.5 GA



Enterprise Grade &
Resilient

Columnar Page-Based String Compression Results

SAP BW tables **40%-90% better** compression then Db2 v11.5 GA



Market Awareness
and Enablement

Completed webinars series – 1,000+ signups!

Db2 Everywhere Information Webinar Series 3 Sessions (Oct 29 – Dec 8)

Db is AI
Ready

Oct 29

Db is
Developer
Inclusive

Nov 10

Db is
Resilient

Dec 8

Signup: ibm.biz/Db2isDeveloper-inclusive

Db2 11.5.5 Technical Webinar Series 8 Sessions (Nov 12 - Dec 14)

Unboxing
Db2
11.5.5

Get an overview of all features
enhancements coming in Db2 11.5.5.
Nov 12 @ 11AM EST

Perf. and
Security
Updates

Deep dive in feature enhancements for
security and performance in 11.5.5.
Nov 19 @ 11AM EST

Modern &
Mixed
Workloads

Deep dive in feature enhancements for
mixed workloads such as JSON, Graph
with demos.
Dec 1 @ 11AM EST

Db2 Cloud
Pak for Data
Updates

Deep dive in feature enhancements for
Db2 Containerization on Cloud Pak for
Data.
Dec 2 @ 11AM EST

Db2
Resiliency &
Availability

Deep dive in feature enhancements for
availability and resiliency in 11.5.5.
Dec 3 @ 11AM EST

Db2
Cloud
SaaS

Deep dive in feature enhancements for
Db2 Cloud Services
Dec 8 @ 11AM EST

Data
Management
Console

Deep dive in feature enhancements data
management console focused on DBAs.
Dec 10 @ 11AM EST

Packaging
&
Pricing

Get all the updates on Db2 pricing and
packaging coming with Db2 11.5.5
Dec 14 @ 11AM EST

Signup: ibm.biz/db2webinar

Market Awareness
and Enablement

Db2 Client Success & Recognitions 2020!

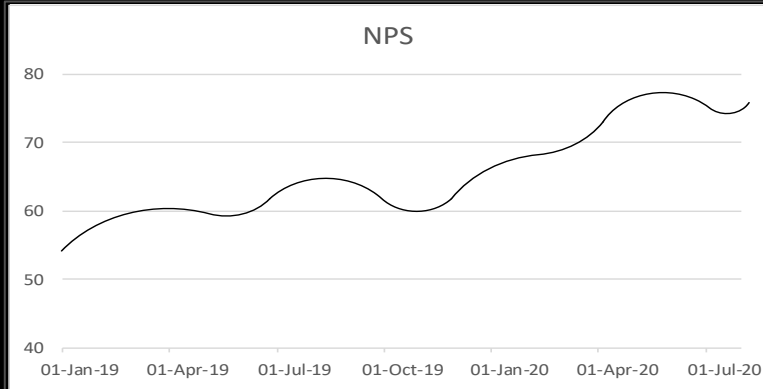


- Support NPS YTD – **73 (Highest-ever)**
- Over **80%** of surveys are **Promotors**
- **50%** of surveys are a **Perfect 10**
- Providing solutions **20% quicker**



reddot winner 2020

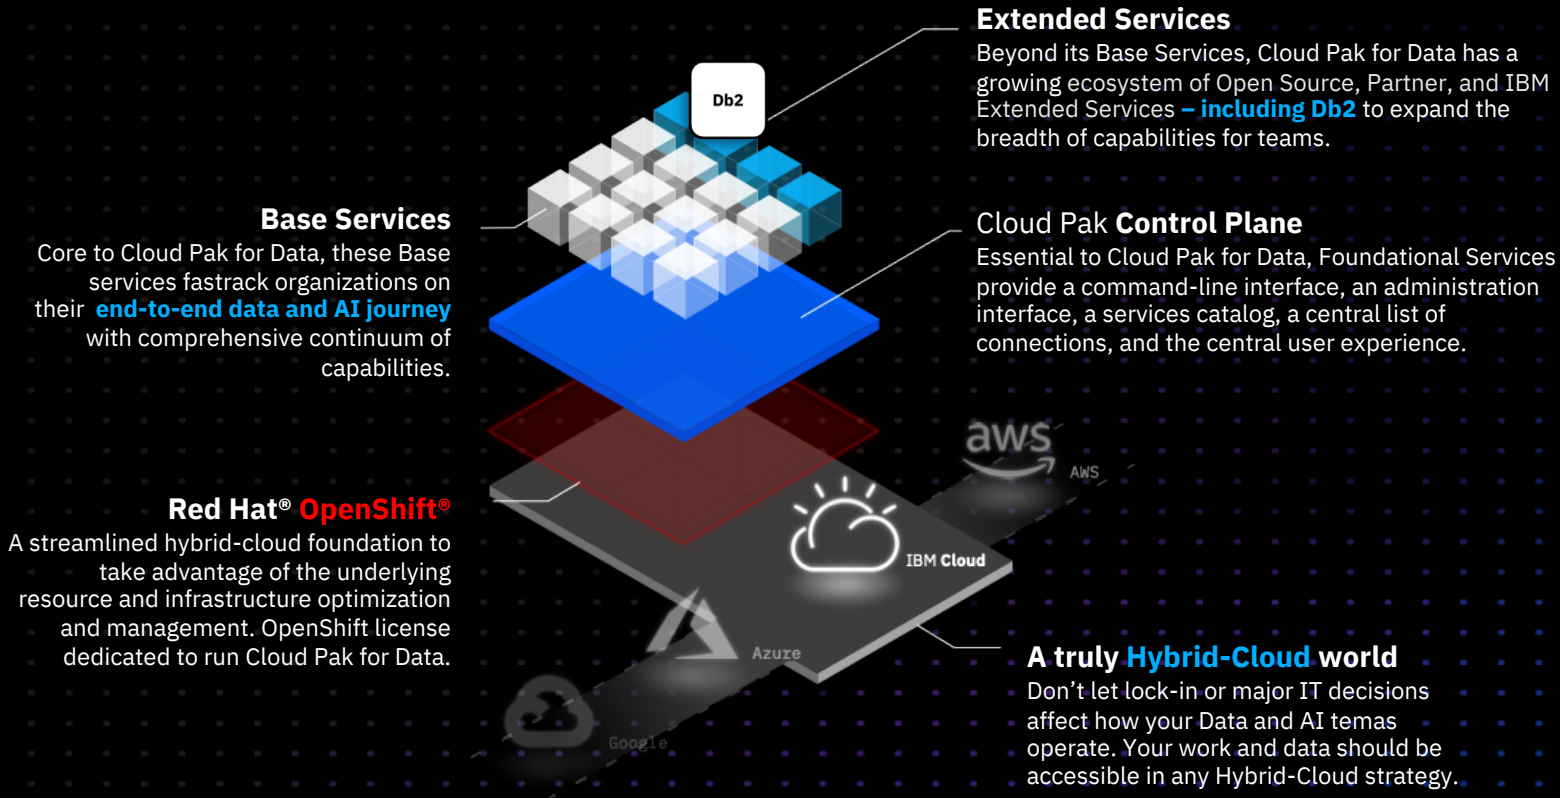
Nikkei Computer Magazine
(Japan) ranks IBM Db2 #1 in
customer satisfaction for
database software



		1	2	3	4	5	6	7	8	
		総合満足度	性能・機能	信頼性	運用性	コスト	サポート	継続意向度	有効回答数	
順位 (前位)	会社名	重視度	Importance							
全社平均 (「有効回答数」は合計)		Overall average	64.6	75.1	74.5	65.6	41.9	52.2	61.7	1072
1(1)	日本IBM (Db2(IBM・DB2), Informixなど)	IBM	71.9	79.2	81.4	73.3	51.0	64.0	66.6	56
2(2)	日本マイクロソフト (SQL Serverなど)	Microsoft	68.9	75.1	75.5	72.4	54.7	55.5	68.4	393
3(3)	日本オラクル (Oracle Databaseなど)	Oracle	59.2	76.3	74.4	58.6	22.1	46.6	53.4	432
4(4)	富士通 (Symfowareなど)	Fujitsu	59.1	56.5	67.7	55.6	54.1	56.8	54.7	40

Cloud Pak for Data + Db2

Simplifies, unifies and automates your journey to AI



IBM Db2 Resources

Information Resources:

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

Developer Resources:

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

Additional resources

White paper

[IBM Db2 on Cloud Pak for Data](#)

Blog

[Easing into your AI journey with IBM Db2 on Cloud Pak for Data](#)

Webinar

[Your journey to AI starts with the right database](#)

[Optimizing your data management infrastructure with Db2](#)

Hands-on lab

[Db2 on Cloud Pak for Data hands-on lab](#)

Video

[Introducing Db2 on Cloud Pak for Data](#)

[Machine learning SQL optimization with IBM Db2](#)

[Machine learning confidence-based query matching with IBM Db2](#)

[Graph analytics in IBM Db2](#)