

Db2 V11.5 – Ready for the Workloads and Deployments of Tomorrow

Les King

lking@ca.ibm.com

September 2023

Data Server Day – Stockholm, Sweden

Agenda

- Investment Drivers
- Modern Deployments
- Modern Workloads
- Additional Information
- Questions

IBM Db2

Portfolio of database solutions

Built to run the world's mission critical workloads

Cloud/SaaS

Db2

SaaS

—
Relational database delivered as a service



Cloud/SaaS

Db2 Warehouse

SaaS

—
Cloud data warehouse delivered as a service



Software

Db2

—
Relational database built to run the world's mission critical workloads



Software

Db2 Warehouse

—
High-performance data warehouse for deep analytics and machine learning



Software


Db2 BigSQL

—
SQL-on-Hadoop engine, delivering MPP and advanced data query



The AI Ladder

A prescriptive approach to accelerating the journey to AI



AI


INFUSE – Operationalize AI with trust and transparency

ANALYZE – Scale insights with AI everywhere

ORGANIZE – Create a trusted analytics foundation

COLLECT – Make data simple and accessible

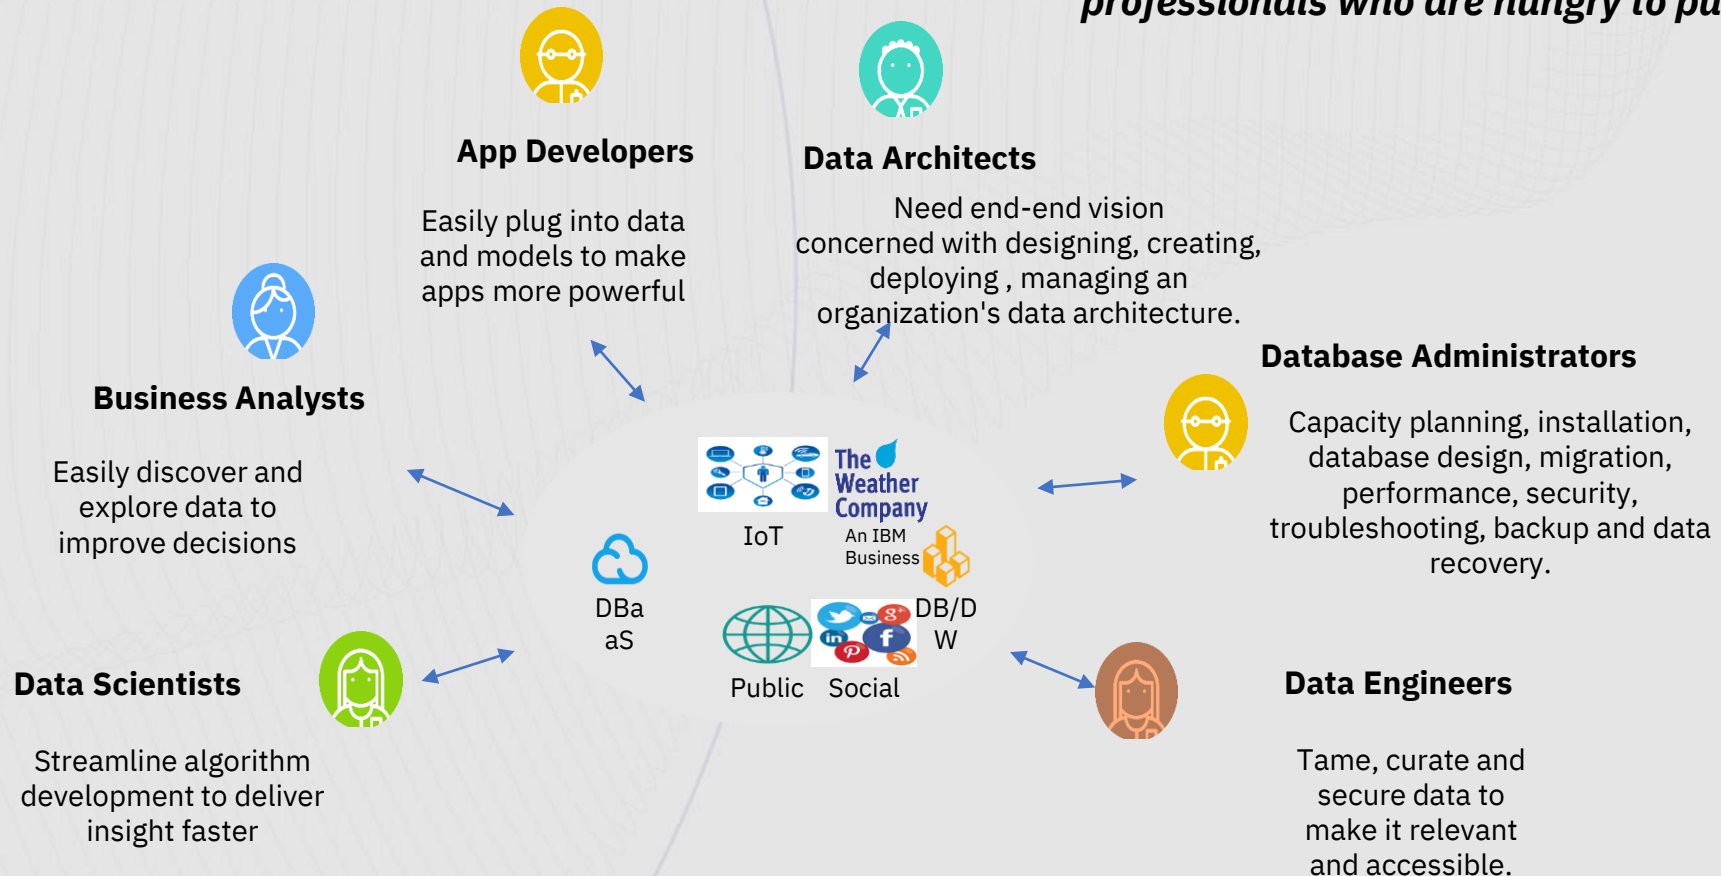
Data of every type, regardless
of where it lives



MODERNIZE
your data estate
for an AI and
multicloud world

Evolving the Needs of All Data Professionals

As data maturity increases, so does the number of data professionals who are hungry to put data to work



Db2 – A Modern Database System

Containerized Deployment Options

The ability to deploy Db2 leveraging containerized platforms including RHOS and non-RHOS platforms

Fully Managed DBaaS

The ability to deploy Db2 as a fully managed service for public, multi and hybrid cloud environments

Cornerstone of COLLECT in Journey to AI

Db2 is the cornerstone for COLLECT in Cloud Pak for Data to cover end-to-end needs to leverage AI

Self Managed Db2 Deployments

The ability to deploy proven reference architectures of Db2 – for any workload – on all key cloud providers

Data Virtualization

Db2 contains a data virtualization component which allows Db2 to be a doorway to all of your business data

In-Db2 Machine Learning

Allows data scientists and developers to bring machine learning local to the data stored within Db2

Multi-Model – NoSQL and NewSQL Data Store

Db2 is a multi-model data store supporting native relational, JSON, BSON, Graph, Spatial, Text and XML

Mixed Workloads

Db2 can handle any combination of workloads including real-time data ingestion, multi-model and mixed.

Two Dimensions of Modernization

- **Deployment Form**
- **Workload Dynamics**

Db2 – Ready for Modern Deployments

Containerized Deployment Options

The ability to deploy Db2 leveraging containerized platforms including RHOS and non-RHOS platforms

- RHOS
- ROSA (Amazon RHOS)
- ARO (Azure RHOS)
- GCP RHOS
- EKS, AKS, GCE (K8S)
- Operators
- SMP + HADR + DPF
- Modernization Tooling

Fully Managed DBaaS

The ability to deploy Db2 as a fully managed service for public, multi and hybrid cloud environments

- Db2 on Cloud (IBM)
- Db2 Warehouse on Cloud (IBM)
- Db2 Warehouse on Cloud (AWS)
- Roadmap: Azure, GCP, RDS, SQL Database

Cornerstone of COLLECT in Journey to AI

Db2 is the cornerstone for COLLECT in Cloud Pak for Data to cover end-to-end needs to leverage AI

- Integrated into CP4D
- CP4D can access Db2 anywhere

Self Managed Db2 Deployments

The ability to deploy proven reference architectures of Db2 – for any workload – on all key cloud providers

- Any topology
- Reference Architectures for DPF
- Reference Architectures for pureScale
- AWS, Azure, GCP

Db2U – Now & Future


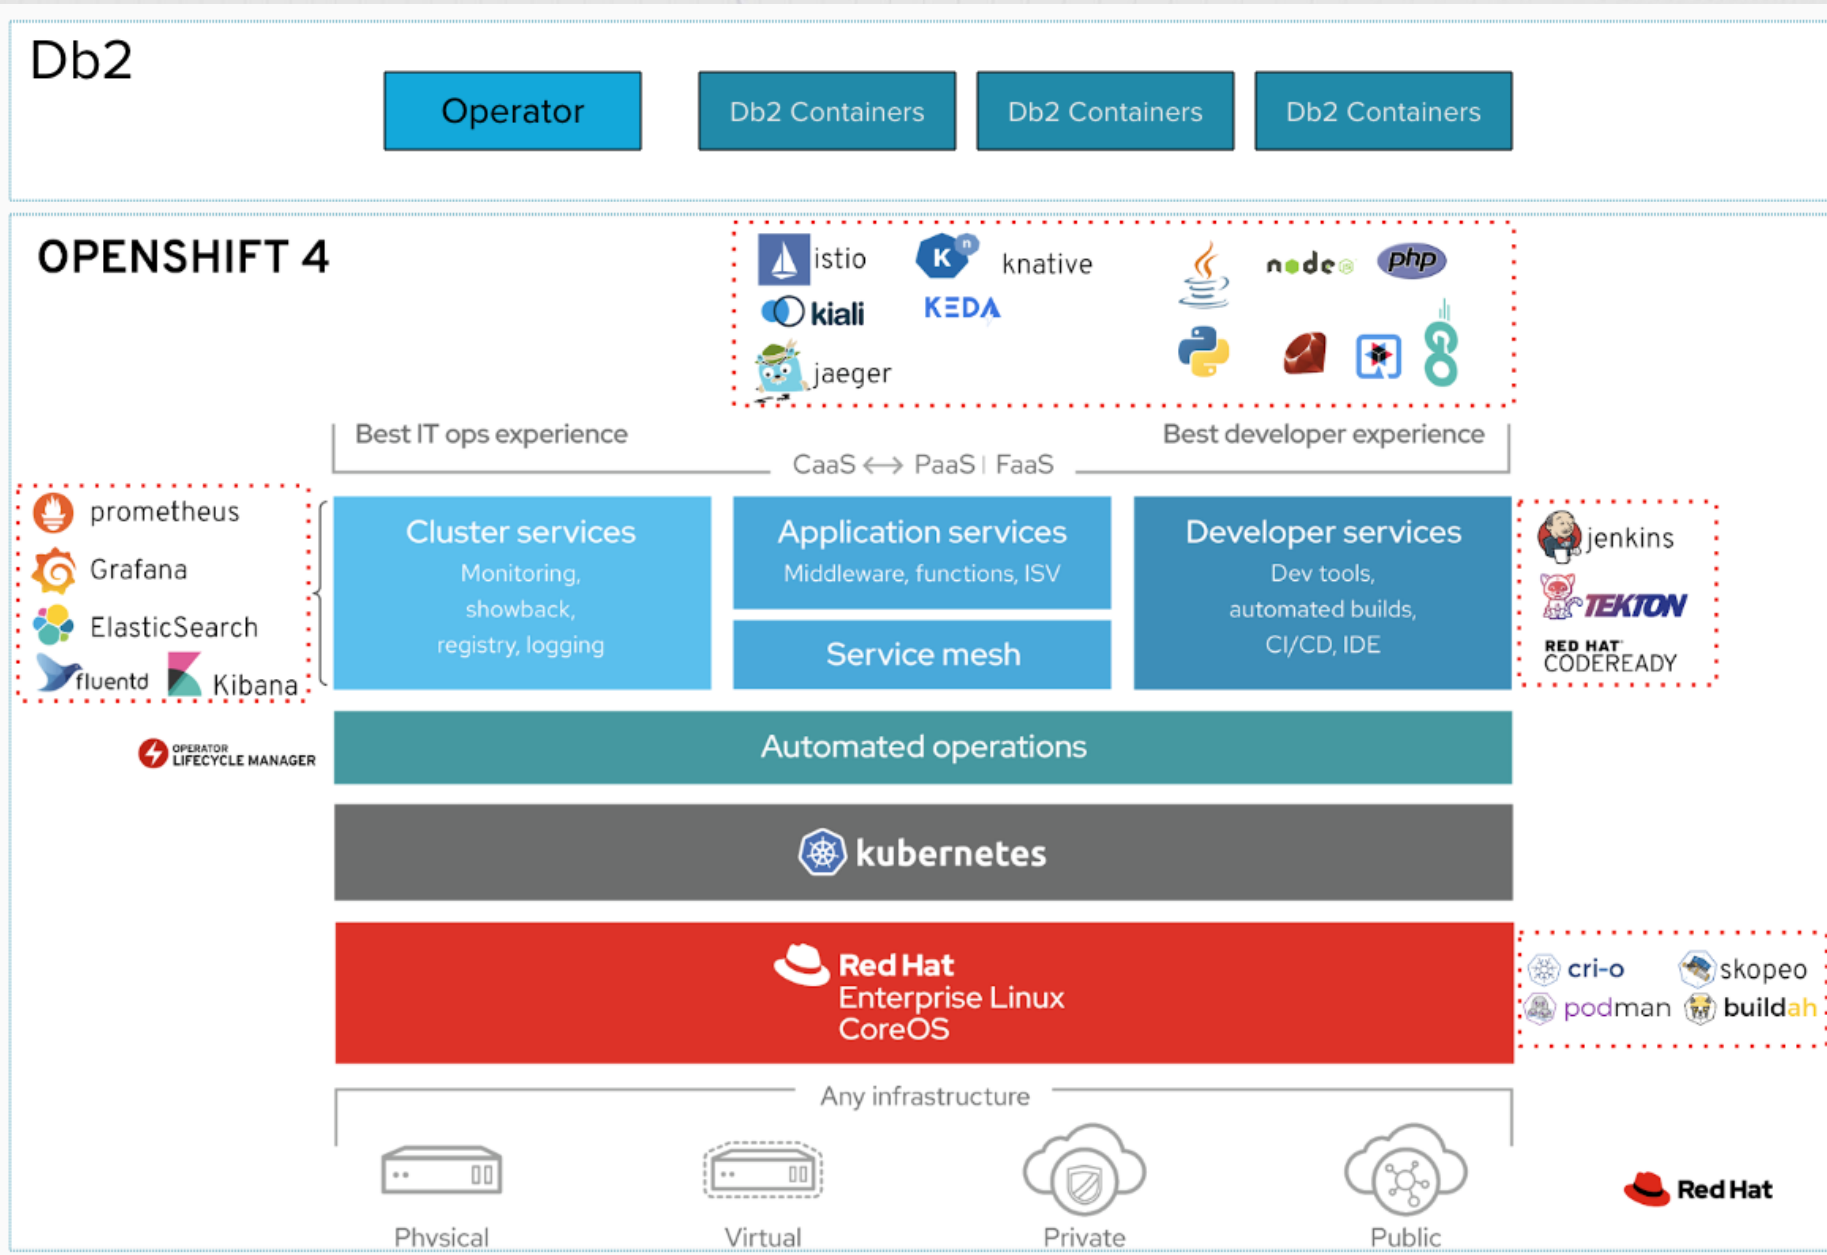
What is it ?

- ✓ Db2 Universal Container on Kubernetes
- ✓ SaaS Offering on IBM Cloud
- ✓ OLTP and OLAP workloads
- ✓ Elastic – scale up or down
- ✓ Cloud-Native

Now Available (11.5.8 !)

- ✓ Deploy on Amazon EKS
- ✓ Deploy on Azure AKS
- ✓ Deploy on Google Cloud GKE
- ✓ RHOS on AWS & Azure - ROSA and ARO
- ✓ BAR Integration with Velero

Db2 & OpenShift - Containerized Deployment



Db2U Operator – Performance Evaluation



All Performance Results for **Db2 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>


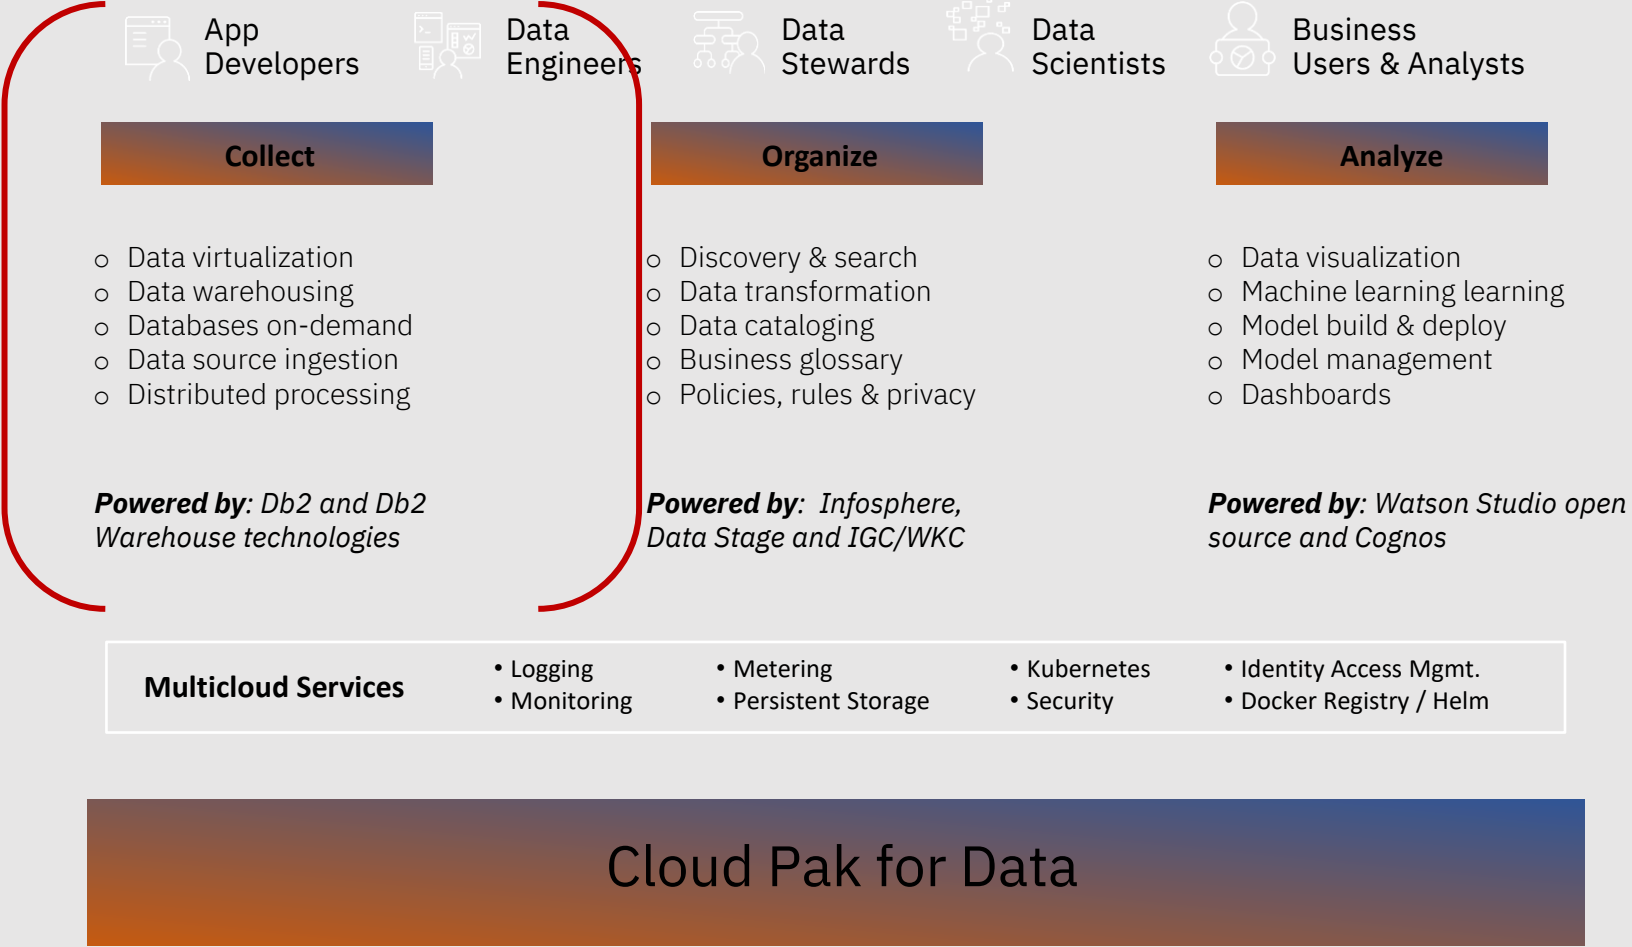
Db2 V11.5 – Cornerstone of “Collect” in Cloud Pak for Data

Foundational “out of the box” multi-cloud Data & AI services

The Ladder to AI



InfuseAI

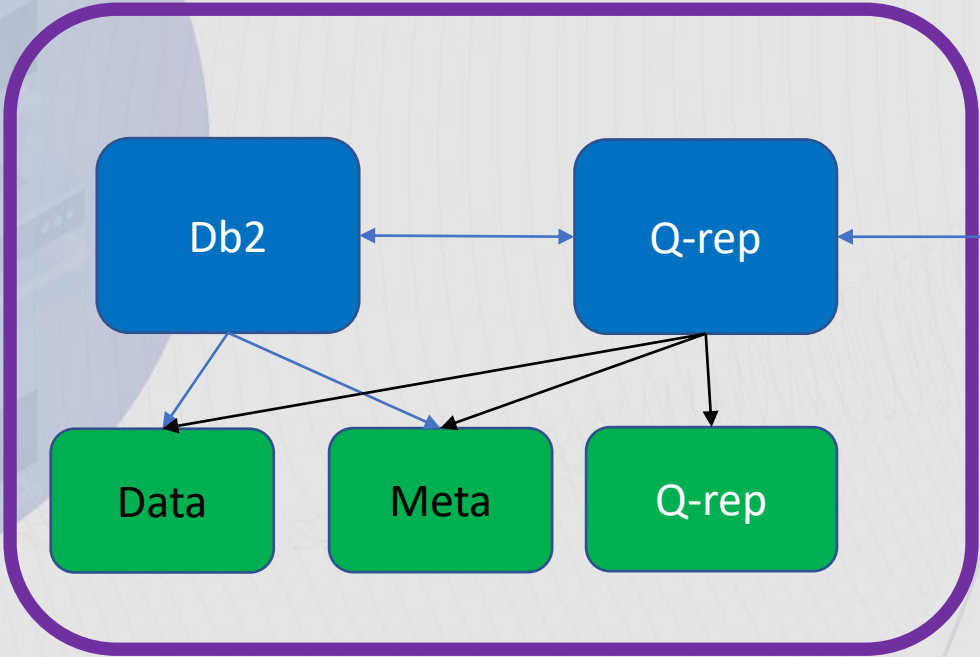


Db2U – What's New in V11.5.6 and CPD 4.0

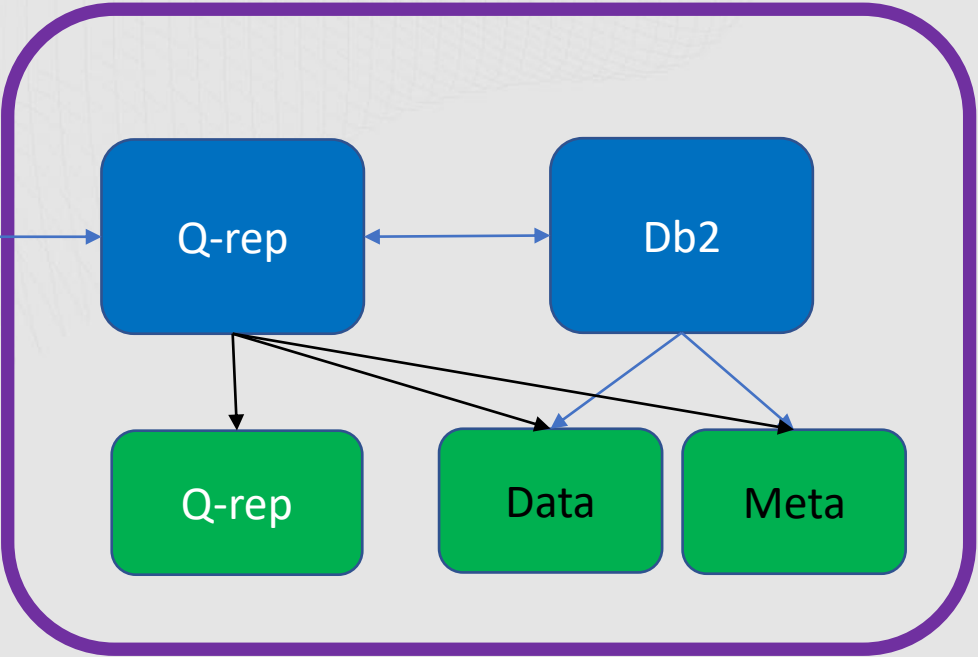
- Expanded Adoption Locations
 - Watson Knowledge Catalog, Open Pages, IBM Operator Catalog (CP4D)
- Deployment
 - Operator
 - Multiple Namespaces (Multiple Databases) – They share resources of container
 - API
 - LPV – Local Persistent Volume
- Security
 - User Group (ie: LDAP Groups)
 - Custom Privileges (ie: UserDefined)
 - IAM Integration (CPP Identity Provider)
- CPD 4.0 is an Operator-driven installation
 - New Operators for our CPD integration

Db2U – What's New in V11.5.7 and CPD 4.0

Instance 1



Instance 2




Modernization Tooling

Accelerated Pathway for Customers to Containerization & Hybrid Cloud

Db2 On Prem

Db2 on Linux




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

Oc Client

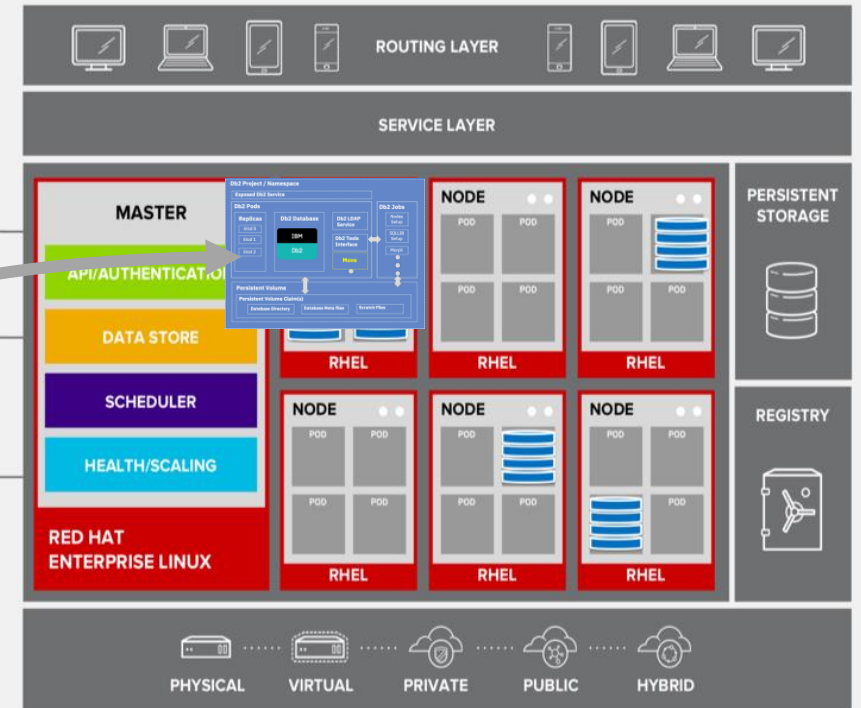
Internet/
Network

**Tested over 15,000 km
Internet connection**

**Tunable Parallel Data Transfer
Data Compression
Fully Auditable**



Db2 for OpenShift / Cloud Pak for Data



Containerized Db2 Hybrid Cloud Service

Db2 Cloud Services




Fully managed / SaaS

Leave the IT work to our CloudOps team 24x7x365



On-Demand Scalable

Scale your environment as your business grows



Highly available and Reliable

Maintain business continuity with peace of mind



Performance and Secure

Leverage blazing fast speed and Security enterprise ready Db2 Engine



All-in Pricing

Control your cloud spending with no-hidden charges.

<Developer>

Rich Developer Ecosystem

Allow developers and business users to self-serve their needs.

Db2 Cloud Services – Today & Tomorrow

Db2 Topology	Fully-Managed (SaaS)		Self-Managed (Reference Architecture)	
Db2 SMP	IBM AWS	Azure	IBM AWS	Azure GCP
Db2 SMP w/HADR	IBM AWS	Azure	IBM AWS	Azure GCP
Db2 DPF	IBM AWS	Azure	IBM AWS	Azure GCP
Db2 DPF w/Q-Rep (DR)	IBM AWS	Azure	IBM AWS	Azure GCP
Db2 pureScale <i>Phase 1=TCPIP + TSA</i> <i>Phase 2=RDMA + pacemaker</i>	TBD		IBM-1; IBM-2* AWS-1; AWS-2	Azure-1; Azure-2 GCP-1; GCP-2*
Db2 pureScale w/HADR (DR) <i>Phase 1=TCPIP + TSA</i> <i>Phase 2=RDMA + pacemaker</i>	TBD		IBM-1; IBM-2* AWS-1; AWS-2	Azure-1; Azure-2 GCP-1; GCP-2*

Available Today

Future

Db2 Warehouse (SaaS) – Gen 3 (11.5.9)




Fully managed / SaaS

Focus on the analytics, we'll take care of the rest



Blazing-fast

Columnar-organized, memory-optimized data warehouse




Scalable & elastic

Independently scale and manage compute & storage



Continuously available

Managed compute, highly available storage, cross-cloud replication



Reliable

Double protection with disaster recovery & self-service backup/restore



Secure

Encrypted at the storage level by the vendor and customer can protect data using their Key Protect keys.

Deploy as fully-managed DBaaS on **IBM Cloud** and **Amazon Web Services**

Db2 – Ready for Modern Workloads

Data Virtualization

Db2 contains a data virtualization component which allows Db2 to be a doorway to all of your business data

- Relational Sources
- Cloud Sources
- Open Source Sources
- NoSQL Sources
- Native Clients
- ODBC, JDBC, REST, NoSQL
- Pushdown Performance
- In-memory MQT

In-Db2 Machine Learning

Allows data scientists and developers to bring machine learning local to the data stored within Db2

- Data Exploration
- Model Training
- Model Evaluation
- Model Deployment
- Data Preprocessing
- Inferencing
- Error Detection
- Support for many models

Multi-Model – NoSQL and NewSQL Data Store

Db2 is a multi-model data store supporting native relational, JSON, BSON, Graph, Spatial, Text and XML

- Graph
- XML/JSON/BSON
- Spatial
- Text
- Gremlin
- XQuery/Mongo/FLWOR
- ESRI
- ACID Properties

Mixed Workloads

Db2 can handle any combination of workloads including real-time data ingestion, multi-model and mixed.

- ML Optimizer
- ML Memory Management
- CDI (Trickle-feed)
- Access multi models
- Access Remote sources
- HTAP
- OLTP + OA + Reporting
- OLAP (All Combinations)

Workloads – Where Db2 Plays

	OLTP	Mixed Workloads		Traditional Enterprise Data Warehouse (EDW)			
Workload	OLTP	Operational	Extreme Analytics	Operational Data Store	Operational Analytics	Analytics	Data Marts
Use	Transactions	Transactions with operational analytics	All types of analytic workloads and federated sources	Simple queries	Single-record look-up / IUD	Deep analytics – bulk scan	Computationally heavy / mining
Data Types	Traditional structured	Traditional structured + events + JSON + Graph + ...	Traditional structured + events + JSON + Graph + ...	Traditional structured	Traditional structured	Traditional structured	Traditional structured
Performance	1000s TPS	1000s TPS + 10s QPS	1000 QPS, many complex	100s QPS	10s-100s of S/IUD per second	100s QPS – 1000 QPS	Lower volume long running
Inserts	1000/s	1000/s	Continual Data Ingest	snapshots Seconds delay	Continual Data Ingest (Trickle feed) – OR - Batch - hourly/n-times daily/daily/etc		
Schema	Normalized	Normalized	De-normalized	Normalized	De-normalized	De-normalized	De-normalized
Applications	SQL	SQL + NoSQL + NewSQL	SQL + NoSQL + NewSQL	SQL	SQL	SQL	SQL

Workloads – New Workloads

	OLTP	Mixed Workloads		Traditional Enterprise Data Warehouse (EDW)			
Workload	OLTP	Operational	Extreme Analytics	Operational Data Store	Operational Analytics	Analytics	Data Marts
Use	Transactions	Transactions with operational analytics	All types of analytic workloads and federated sources	Simple queries	Single-record look-up / IUD	Deep analytics – bulk scan	Computationally heavy / mining
Data Types	Traditional structured	Traditional structured + events + JSON + Graph + ...	Traditional structured + events + JSON + Graph + ...	Traditional structured	Traditional structured	Traditional structured	Traditional structured
Performance	1000s TPS	1000s TPS + 10s QPS	1000 QPS, many complex	100s QPS	10s-100s of S/IUD per second	100s QPS – 1000 QPS	Lower volume long running
Inserts	1000/s	1000/s	Continual Data Ingest	snapshots Seconds delay	Continual Data Ingest (Trickle feed) – OR - Batch - hourly/n-times daily/daily/etc		
Schema	Normalized	Normalized	De-normalized	Normalized	De-normalized	De-normalized	De-normalized
Applications	SQL	SQL + NoSQL + NewSQL	SQL + NoSQL + NewSQL	SQL	SQL	SQL	SQL

Db2 – Handling Modern Workloads

Powered by AI



Confidence-based query results
leveraging ML-SQL



Up to 10x better query performance
powered by an ML-Optimizer



No data movement & single view
on all data
delivered by Data Virtualization



Auto resource optimization
delivered by Adaptive Workload Management

Built for AI



Faster data exploration
by using In-Db2 Machine Learning



Build AI based applications
with Python, R, GO , JSON and Jupyter notebooks



Model Complex Relationships
by using Db2's Multi-Model Capabilities



Blockchain Ready
using Db2 Blockchain Connector

Db2's Machine Learning Optimizer

ML based optimizer improves query performance using predictive estimations to improve query access plans.

The Db2 Machine Learning Optimizer provides an additional level of intelligent optimization to deliver query execution strategies that improve on traditional cost-based query optimization

Basic workload cost optimizers can suggest query execution strategies, but they aren't sensitive to recent changes in the database, and they can't learn from experience

The Db2 Machine Learning Optimizer, by contrast, incorporates feedback from actual query performance to recommend execution strategies that may deliver improved results

Up to **10x** fast query performance!

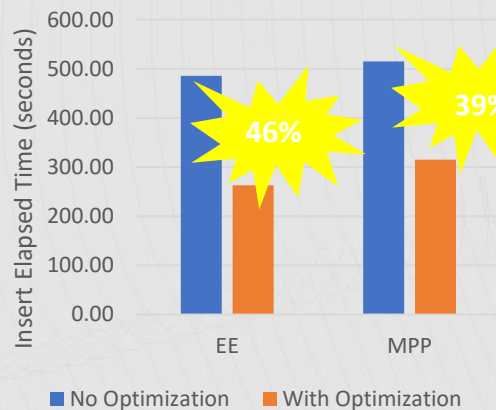


Synopsis Table Maintenance (11.1)


- Buffer tuples in memory and write them to synopsis table every 1024 rows
 - Up to 46% improvement in insert performance!
 - Up to 87% reduction in synopsis table size!
 - Up to 60% reduction in elapsed time with queries with range predicates!

OLTP
+ OLAP
= HTAP


Elapsed Time of INSERT
With and Without Synopsis
Optimization



Size of Synopsis Table
With and Without
Synopsis Optimization



Query Elapsed Time
With and Without Synopsis
Optimization


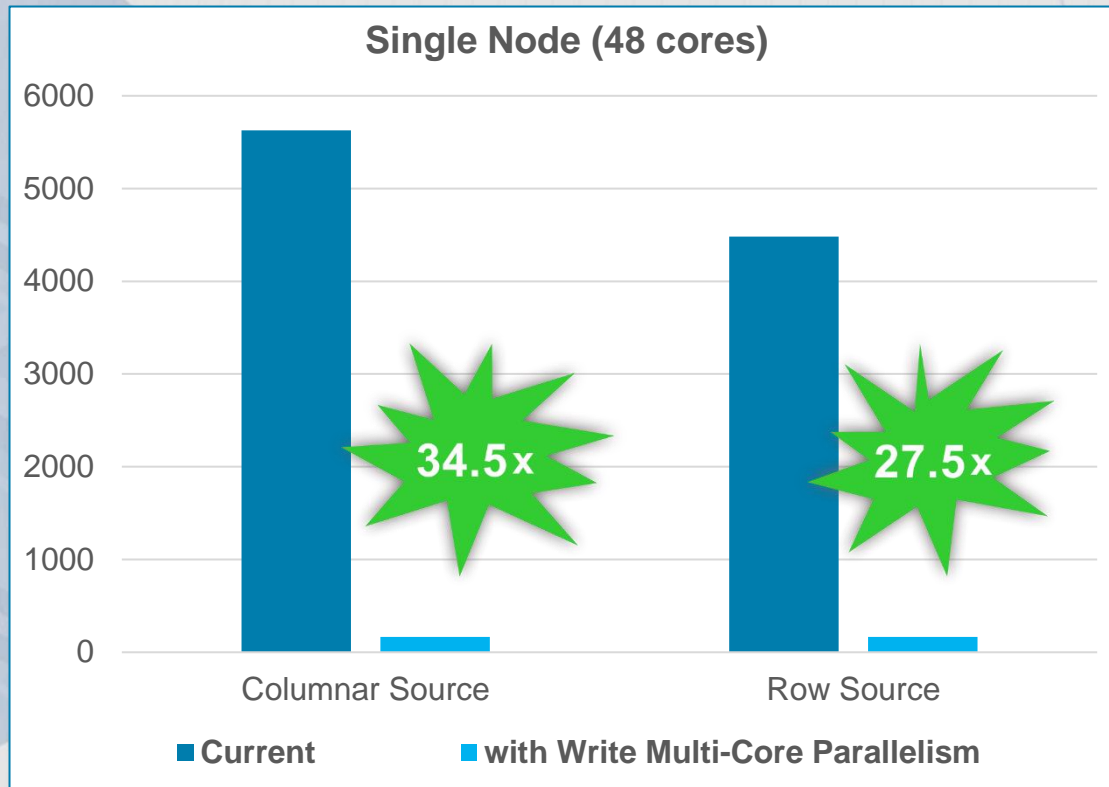


- IBM p760 / POWER7+ 32 cores / 1TB RAM
- Table with 50 columns.

- Single user
- Insert test 100K rows with commit count = 1
- Select query with 10 pairs of range predicates

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Multi-core parallelism (11.1)



- 100 Million Rows
- 47 GB

- DPF tests used 12 logical DB partitions, & co-located tables
- Same 48 core server used for single-node and DPF tests


Columnar – Index Support (11.1)

- DB2 11.1. currently implicitly creates unique indexes to support PRIMARY or UNIQUE KEY constraints
 - A unique index can be used to access data if at most 1 row qualifies and every key in the index must have an equality predicate OR the FETCH FIRST 1 ROW ONLY clause is specified
- As of Db2 11.1.3.3, Db2 now supports CREATE INDEX for column-organized tables:
 - Unique and non-unique indexes
 - DROP and ALTER also supported
 - ALTER INDEX only supports COMPRESS [YES | NO]
- These new indexes can be used to access column-organized data using most of the same methods as row-organized data and without restrictions on the number of qualifying rows


Trickle Feed INSERT in Db2 (11.5.6)

BENEFITS to both INSERTs and UPDATES
Performance on par with BULK INSERTs


Trickle Feed Insert



Column Group Buffer




Column Split



Column Vectors

Column Insert Buffer Writer



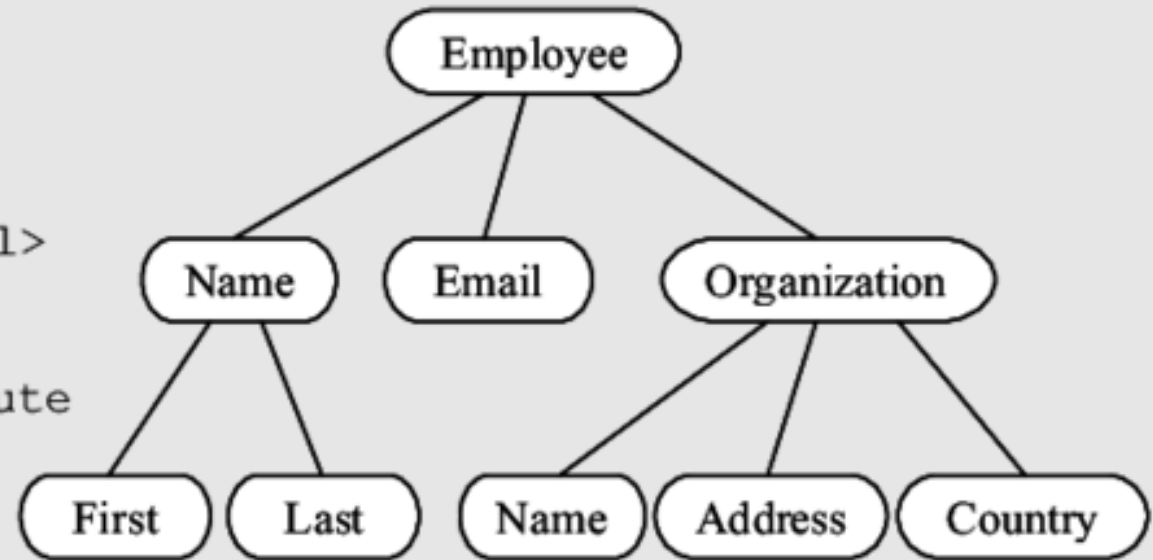
Logging is done for fuller Insert Buffers
Less dirty pages to process in the bufferpool
Reduced Storage Space and Log Space

What does Db2 mean by Multi-Model support ?

- **Storage** – natively storing the data to ensure no loss of data – no force fitting into row/column structure (ie: shredding)
- **Performance** – ability to index the data in a meaningful way to provide tier 1 performance for both ingestion and queries
- **Integration** – ability to query data in each model of data – within the same query
- **SQL Support** – ability to work with the model of data using SQL
- **NoSQL Support** – ability to work with the model of data using a natural query language for that model of data
- **NewSQL Support** – support for transactional awareness (ACID properties) when using that model of data
- **Enterprise Requirements** – leverage Db2's availability, security, recoverability, etc for that model of data
- **Output** – ability to decide between traditional relational result set of model specific output

XML (pureXML in 10.5)

```
<Employee>
  <Name>
    <First>Lassi</First>
    <Last>Lehto</Last>
  </Name>
  <Email>Lassi.Lehto@fgi.fi</Email>
  <Organization>
    <Name>
      Finnish Geodetic Institute
    </Name>
    <Address>
      PO Box 15,
      FIN-02431 Masala
    </Address>
    <Country CountryCode="358">Finland</Country>
  </Organization>
</Employee>
```



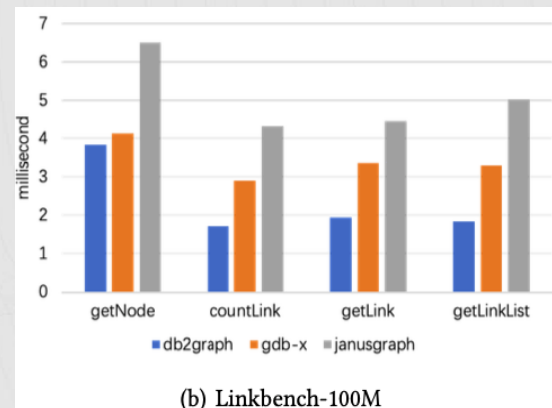

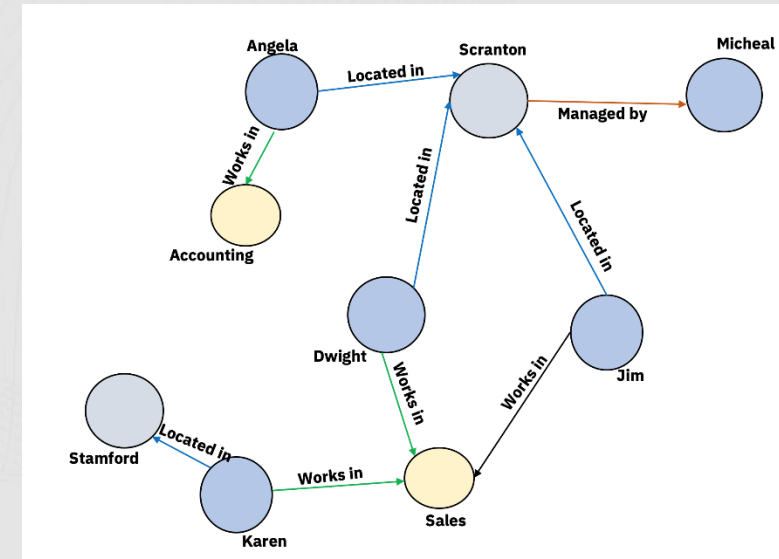
JSON & BSON Support in Db2 (11.5.5)

- Natively store JSON or BSON data in Db2
 - Get top tier performance with INDEX support
 - Under the covers we leverage:
 - VARCHAR / VARBINARY for smaller documents
 - CLOBS / BLOBS for larger documents
 - Full ACID and enterprise level availability, scalability, performance, security and recoverability available
- SQL Support
 - Proprietary UDFs with many JSON/BSON functions
 - SQL:2016 Standards (report) – 70% complete
 - NoSQL Support
 - Mongo API wire listener
 - Data Ingestion
 - LOAD, IMPORT, INSERT, etc
 - CVC Enhancements

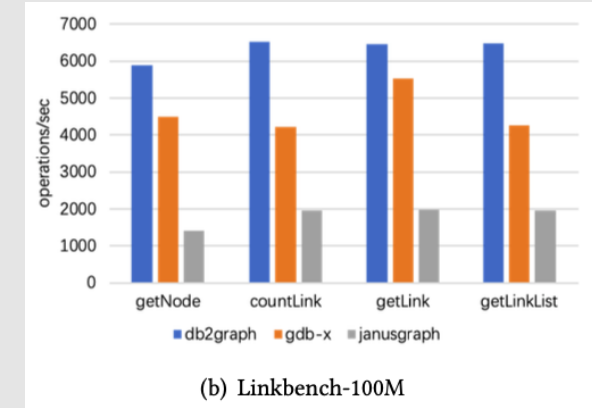
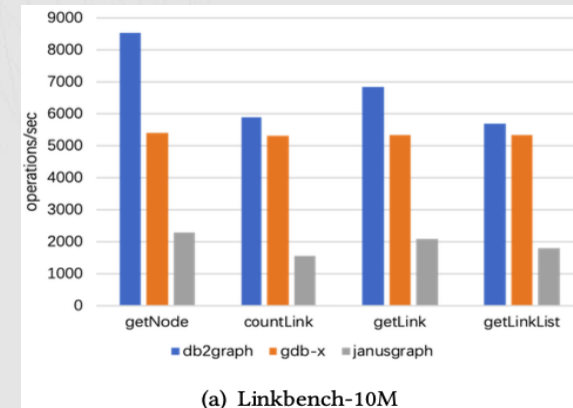
Db2 Graph (11.5.6)

- **An enterprise grade in-database Graph solution**

- **Store and query** data and relationships
- **Traverse data** across depths and find hidden relationships between data points
- Perform **Graph analytics and traditional SQL** on the same storage engine
- Use in conjunction with transactional data and other applications, in **real-time**

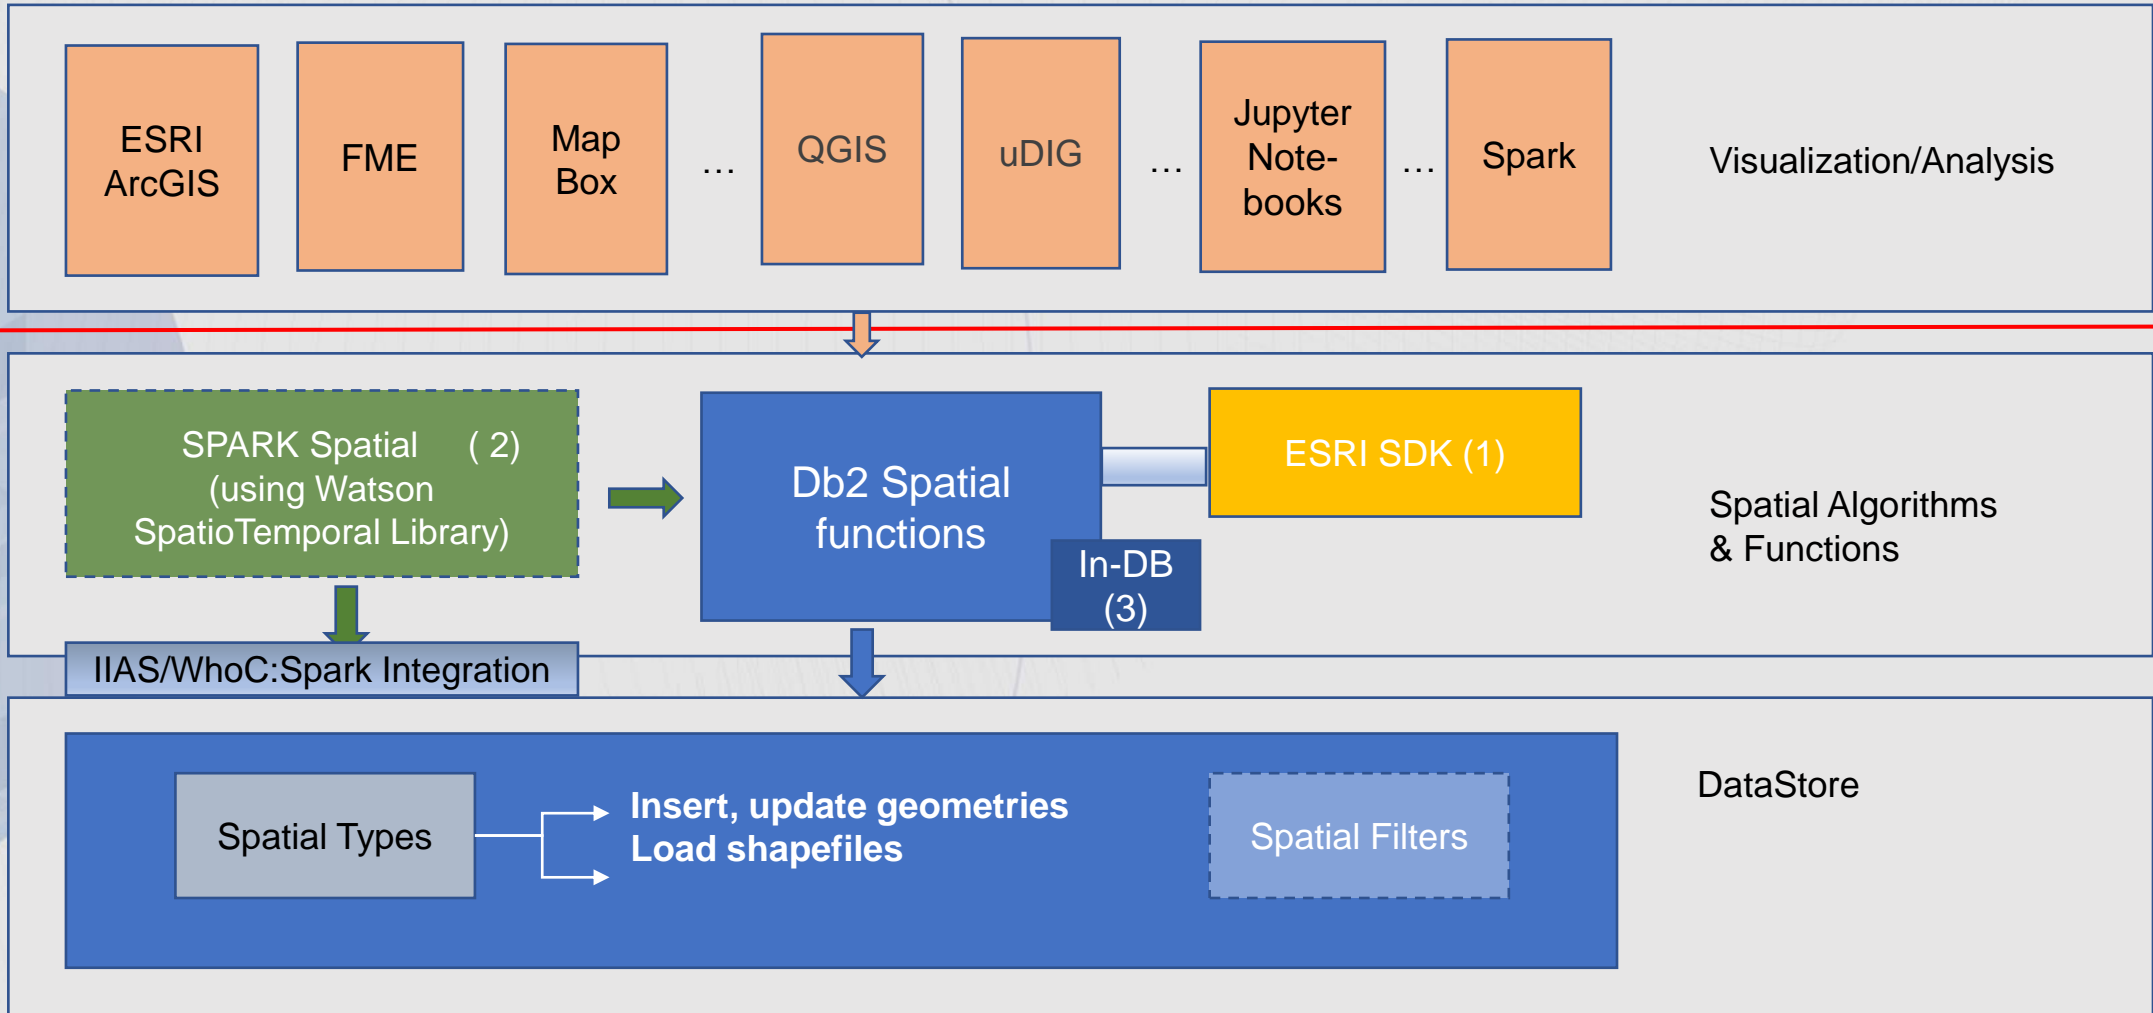


Latency – lower is better



Throughput – higher is better

Spatial Analytics (11.5.7)



DB2 Text Search Features

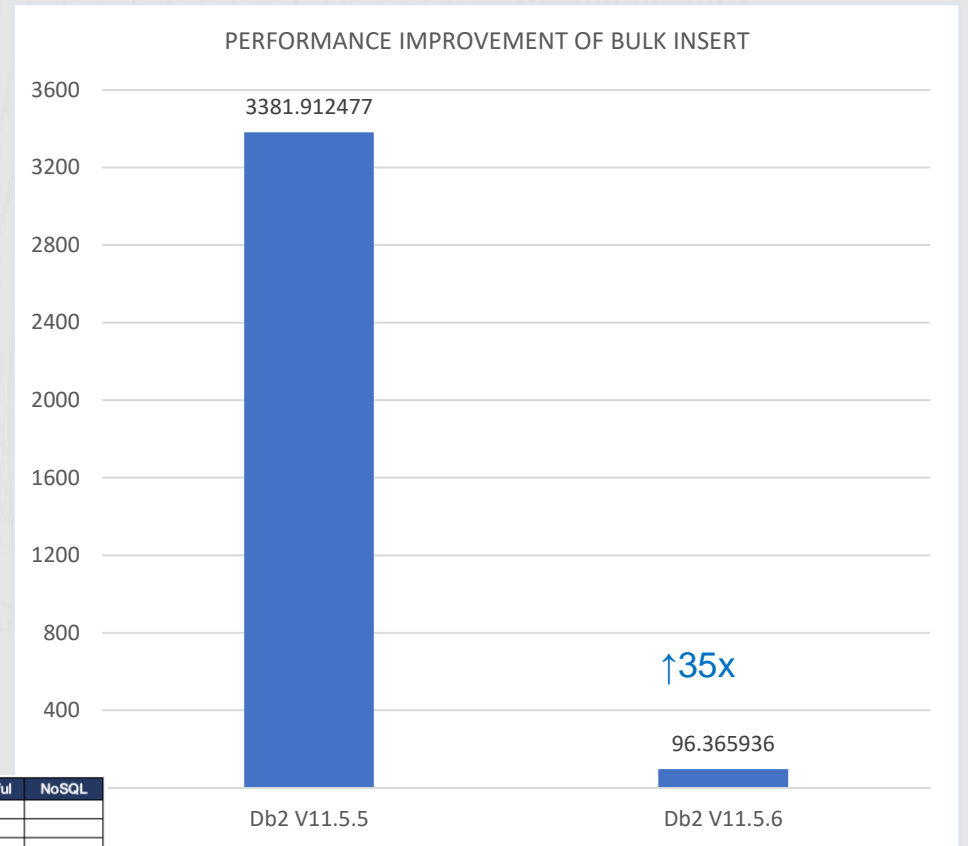
- **Advanced search technology**
 - SQL functions CONTAINS and SCORE
 - SQL, SQL/XML and XQUERY support including XPath-like syntax to search XML docs
 - Built-in SQL functionality combined with the DB2 optimizer
 - Linguistic processing for 20+ languages/locales
 - Synonym dictionary support
- **Document Indexing**
 - Native XML support
 - Multiple document formats, including rich text (via Accessories Suite)
 - Incremental and asynchronous index updates
- **Integration with Db2**
 - Integrated install; No additional license required
 - Stored procedures for administration
 - Scheduler for index updates
 - Index type TEXT

Data Virtualization – Federation (11.5.8)

- Connectivity – Spark JDBC Connectivity Support
- Functionality – Column Length Variation for Code Page Conversion
- Functionality – Nickname Hidden Column Support
- Performance – Federation DRDA Bulk Insert for Db2 Family Data Sources

Category	Data Source	Native	ODBC	JDBC	RESTful	NoSQL
Relational	Db2 LUW	Yes		Yes		
	Db2 for IBM i	Yes				
	Oracle	Yes	Yes	Yes		
	MS SQL Server	Yes	Yes	Yes		
	Informix	Yes				
	Sybase	Yes				
Warehouse / Appliance	IIAS	Yes		Yes		
	Netezza		Yes	Yes		
	Teradata	Yes		Yes		
	SAP HANA		Yes	Yes		
	Greenplum		Yes	Yes		
Open Source	MySQL Community		Yes	Yes		
	MySQL Enterprise		Yes	Yes		
	PostgreSQL		Yes	Yes		
	MariaDB		Yes	Yes		
	Derby			Yes		
Hadoop	IBM Db2 BigSQL	Yes		Yes		
	Hive		Yes	Yes		
	Spark		Yes	Yes		
	Impala		Yes			
Files	Delimited	Yes				
	Excel	Yes	Yes			
	XML	Yes				
	JSON					Yes
	CSV	Yes				
Mainframe	Db2 for z/OS	Yes		Yes		
	IBM DVM for z/OS			Yes		

Category	Data Source	Native	ODBC	JDBC	RESTful	NoSQL	
Message Queue	MQSeries	Yes					
	Db2 Warehouse	Yes		Yes			
Cloud	MS Azure SQL		Yes				
	Oracle Cloud		Yes				
	Amazon AWS Redshift			Yes			
	Google BigQuery			Yes			
	Amazon AWS S3			Yes			
	Salesforce			Yes			
	Snowflake		Yes	Yes			
	NoSQL	Hyperledger Fabric					Yes
		MongoDB					Yes
CouchDB						Yes	
Hbase HDFS						Yes	
Cassandra						Planning	
Redis						Planning	
Jira					Yes		
Aha!					Yes		
GitHub					Yes		
HubSpot					Yes		
TeamCity					Yes		
api.spacexdata.com					Yes		
earthquake.usgs.gov					Yes		
Google Calendar API					Yes		
groupkt.com					Yes		
Yelp					Yes		



Supported Before v10.5
Supported In v11.1
Supported In v11.5 GA
Supported in v11.5.4
Supported in v11.5.5
Supported in v11.5.6
Working / Planning


In-Db2 Machine Learning (11.5.6)

Train, Tune, Cleanse, Explore, Evaluate, Manage, Error Detection, Inferencing


Integrated Python and R Library for exploring and manipulating data

Accelerated and Distributed Machine Learning Algorithms in Db2


Random Forest



Decision Tree




Logistic Regression



The IBM Data Lakehouse

The IBM data lakehouse brings together the advantages of data warehouses and data lakes within a new managed cloud service and self-managed on any Cloud or on-premise.



A Low-Cost and extensible Query Engine

Presto is an **open-source, fast** and reliable **SQL engine** for Data Analytics and data lake houses.



A proven and reliable metadata repository

The hive **metastore** is the de facto standard in open-source data lake metadata management




Stores data in Object Store buckets in the Iceberg open data format to **facilitate data access and sharing** across applications




Open, Flexible, and Modular

Lakehouse is designed to enable customers to standardize their data formats and metadata with unified data governance




Evolve your big data platform

The simplest path to upgrade from traditional big data platforms, either as a side car or moving data to cloud object storage



Limitless scalability and elasticity

Explore, shape, and analyze data at any scale by separating storage and compute

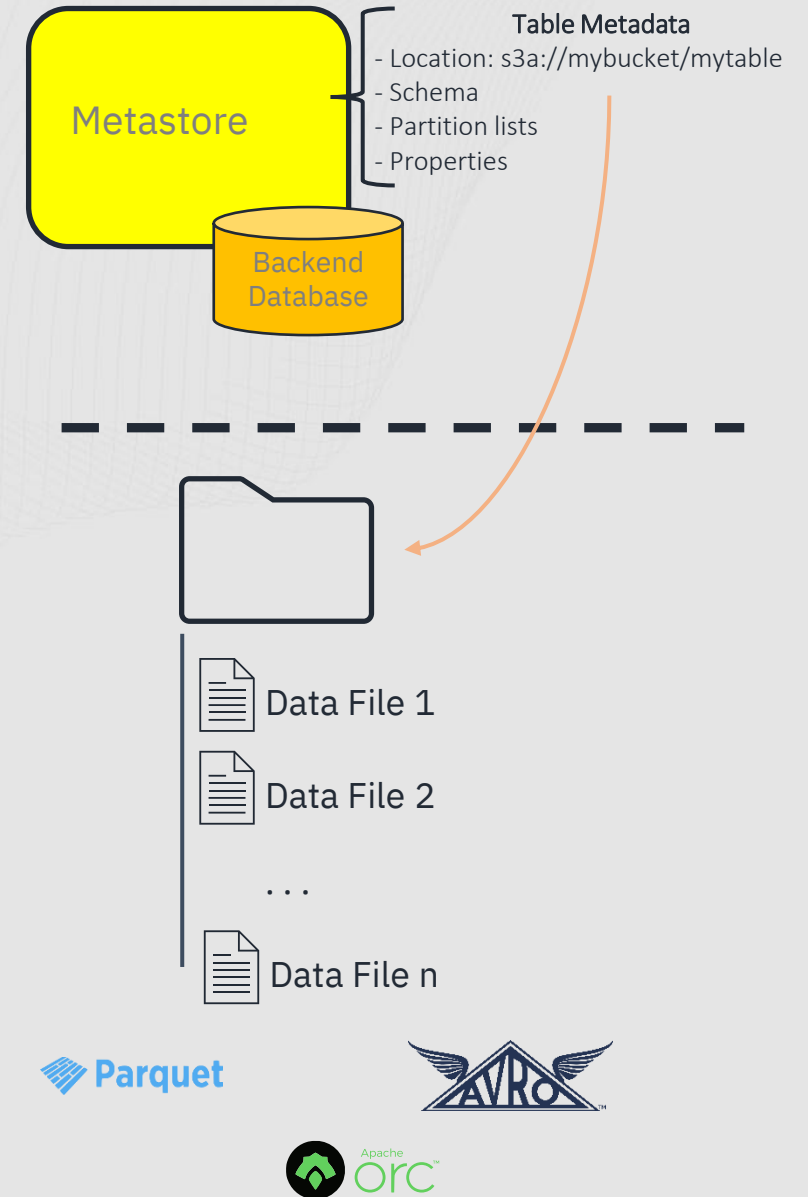


Integrates readily with Db2

Warehouse and Netezza to support the **right data engine for the right workloads at the right cost**


Datalake Tables (11.5.9)

- A Data Lake “Table” is a collection of files serialized using an **Open Data File** (ODF) format (CSV, ORC, Avro, Parquet ...) stored on remote storage (HDFS, S3, COS, ...)
- The **metadata** of the table is stored in a Metastore server
 - Location
 - Schema
 - Partition lists
- An engine querying the table must query the metadata first and can proceed to read the data from remote storage
- Benefits
 - Interoperability of open data formats
 - Ease of use



Object Storage support for Table Storage (11.5.9)

- **Db2WoC 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 store data across a mix of block storage and object storage, based on business or technical requirements
 - Db2WoC uses different mechanisms in order to facilitate reads and writes to object storage
- **No applications and workload changes necessary in order to use this feature**
 - Db2WoC handles all the necessary interfacing to object storage, thus existing applications and warehouse workloads do not have to be changed in order to make use of this
 - Specific Db2 tablespace available backed by S3 for customer use
 - Insert, Update, Delete data as needed 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



More Information - In-Db2 ML

Demos:

[Build a Customer Segmentation Model with Db2 \(K-Means Clustering\)](#)

[Build a Classification Model with Db2 \(Decision Tree\)](#)

[Build a Regression Model with Db2 \(Linear Regression\)](#)

[Integrate a Db2-native model with a Cognos Dashboard](#)

[Deploying a ML Model Trained on Cloud Pak for Data to Db2](#)

Hands-On:

[Tutorials and Jupyter Notebooks](#)

[Pre-configured Hands-on Environment](#)

Documentation:

[Db2 11.5 Knowledge Center](#)

More Information - In-Db2 ML – Hand-on-Training

- Gentle Introduction to Machine Learning Concepts
- Hands-on: build and deploy ML models using Db2's built-in ML Stored Procedures
- Hands-on: build and integrate open-source Python models with IBM Db2

The Workshop is delivered via requestable VMs available for both IBM employees and customers.

Contact: Shaikh Quader

Email Address: shaikhq@ca.ibm.com

More Information – Spatial Analytics

Documentation:

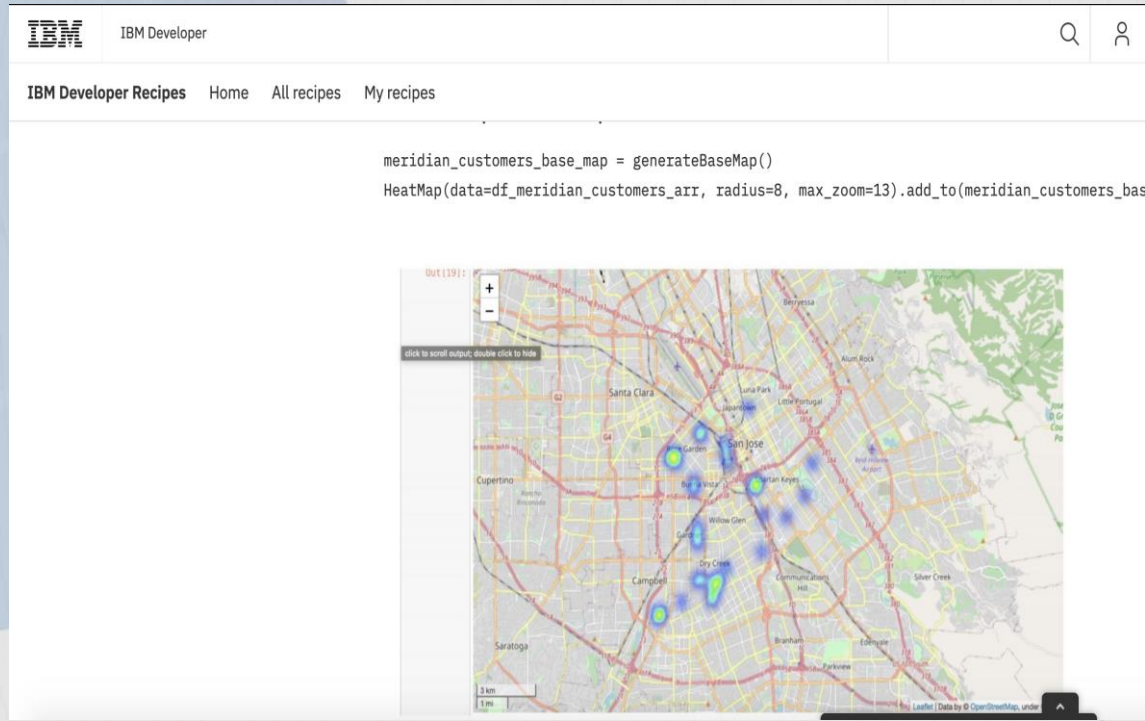
<https://www.ibm.com/docs/en/db2/11.5?topic=spatial-data>

Banking location usage insight scenario

<https://developer.ibm.com/recipes/tutorials/ibm-db2-spatial-analytics-bank/>

Runkeeper Scenario

<https://ibm.app.box.com/s/kb7baqbchmfptyf903q9fyqvpe3w093n>

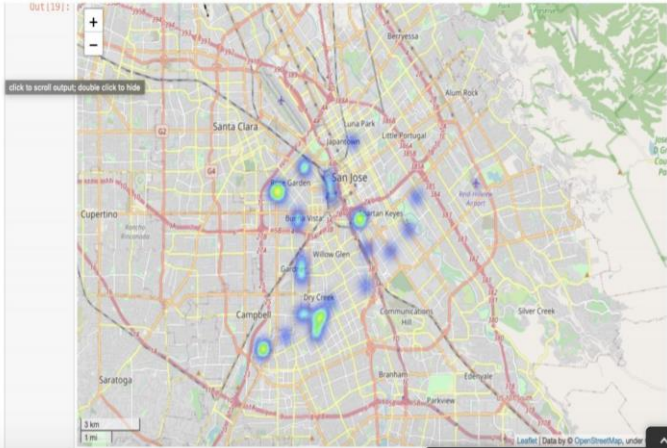


IBM Developer

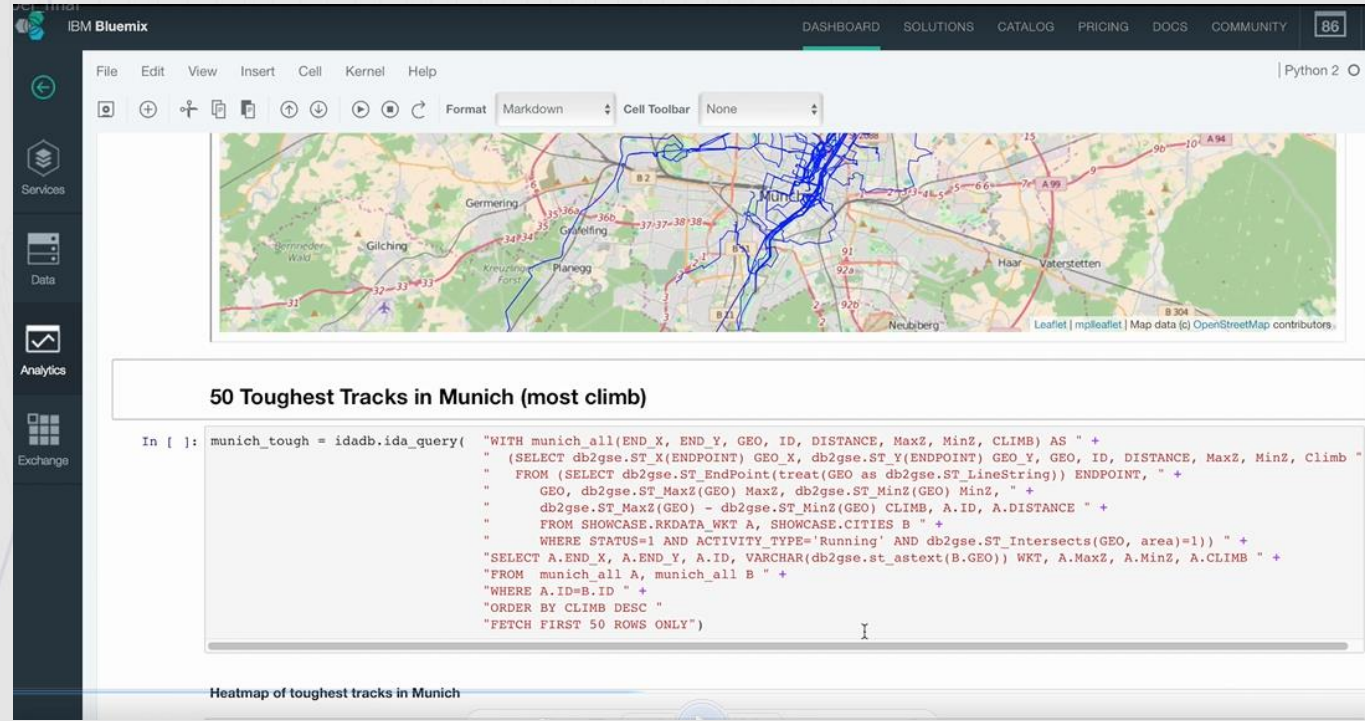
IBM Developer Recipes Home All recipes My recipes

```
meridian_customers_base_map = generateBaseMap()  
HeatMap(data=df_meridian_customers_arr, radius=8, max_zoom=13).add_to(meridian_customers_base_map)
```

Out [19]:



click to scroll output; double click to hide




IBM Bluemix

DASHBOARD SOLUTIONS CATALOG PRICING DOCS COMMUNITY 86

File Edit View Insert Cell Kernel Help Python 2.0

Format Markdown Cell Toolbar None



Leaflet | mapbox, OpenStreetMap contributors

50 Toughest Tracks in Munich (most climb)

```
In [ ]: munich_tough = idadb.ida_query( "WITH munich_all(END_X, END_Y, GEO, ID, DISTANCE, MaxZ, MinZ, CLIMB) AS " +  
" (SELECT db2gse.ST_X(ENDPOINT) GEO_X, db2gse.ST_Y(ENDPOINT) GEO_Y, GEO, ID, DISTANCE, MaxZ, MinZ, Climb " +  
" FROM (SELECT db2gse.ST_EndPoint(treat(GEO as db2gse.ST_LineString)) ENDPOINT, " +  
" GEO, db2gse.ST_MaxZ(GEO) MaxZ, db2gse.ST_MinZ(GEO) MinZ, " +  
" db2gse.ST_MaxZ(GEO) - db2gse.ST_MinZ(GEO) CLIMB, A.ID, A.DISTANCE " +  
" FROM SHOWCASE.RKDATA WKT A, SHOWCASE.CITIES B " +  
" WHERE STATUS=1 AND ACTIVITY_TYPE='Running' AND db2gse.ST_Intersects(GEO, area)=1) " +  
"SELECT A.END_X, A.END_Y, A.ID, VARCHAR(db2gse.st_astext(B.GEO)) WKT, A.MaxZ, A.MinZ, A.CLIMB " +  
"FROM munich_all A, munich_all B " +  
"WHERE A.ID=B.ID " +  
"ORDER BY CLIMB DESC " +  
"FETCH FIRST 50 ROWS ONLY")
```

Heatmap of toughest tracks in Munich

More Information

Db2 JSON, BSON Support
ibm.biz/db2json

Db2 Graph Database and Graph Query Demos:
<https://youtu.be/C5vmcYKEN-U>
https://youtu.be/5_5UMeGWHV8

In-Db2 Machine Learning Demos:
[Build a Customer Segmentation Model with Db2](#)
[Build a Classification Model with Db2 \(Decision Tree\)](#)
[Build a Regression Model with Db2 \(Linear Regression\)](#)
[Integrate a Db2-native model with a Cognos Dashboard](#)

Hands-On:
[Tutorials and Jupyter Notebooks](#)

Documentation:
[Db2 11.5 Knowledge Center](#)

More Information – Db2 Resources

Information Resources:

- Db2 Roadmap - <http://ibm.biz/AnalyticsRoadmaps>
- Db2 RFE (Idea) Portal - <http://ibm.biz/submitdb2idea>
- Db2 Recorded Educational Webinars- <http://ibm.biz/db2webinar>
- Subscribe to Db2 technical newsletter - <http://ibm.biz/db2nlsignup>
- Connect with the Db2 online community - <http://ibm.biz/db2tribe>

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>

Thank You

Speaker: Les King

Company: IBM

Email Address: lking@ca.ibm.com