

IBM Db2

Built for the world's mission
critical workloads

Aniruddha Joshi

Senior Product Manager, Data & AI

Arjun Gujral

Product Manager, Data & AI

30+

Years running the world's
mission critical workloads
with Db2

Governed,
secure



Continuously
available



Endlessly
scalable



Highly
performant



Real-time
analytics



Automated
operations



10/10

largest global banks run on Db2

- Running mission critical banking systems to support trillions of financial transactions per day
- Securing your portfolio, reference, market, and risk data from a data breach
- Detecting fraud and money laundering
- Enabling you to manage finances digitally, across any device
- Managing and analyzing trillions of assets globally



#1

Fortune 10 retailer runs on Db2

- Powering inventory and supply chain visibility across tens of thousands of stores
- Ingesting and analyzing petabytes of unstructured data from millions of customers every hour
- Personalizing omni-channel experiences for hundreds of millions of customers
- Securing customer payments globally

9/10

largest global
auto
manufacturing
companies run
Db2

- Enabling a data management foundation to support connected cars and autonomous vehicles
- Optimizing inventory and supply chain operations with real-time insights and high availability of applications
- Predicting maintenance and tuning to keep operations running safely

And we are invested in supporting the future needs of **all our customers** across the planet

76
of Fortune 100¹
companies run on Db2

100T+
Transactions per day

11.5K+
customers globally



**30+ years of
innovation**

That's our
DNA

However, **today's data** landscape has made the ability to become data-driven **infinitely complex**

Data apps & services are growing

New business opportunities are producing 15ZB of data



Web & mobile apps



Business apps



Business Intelligence tools



Dashboards



Microservices

Data engines have diversified

82% of enterprises are inhibited by data silos



Transactional databases (OLTP)



Analytics warehouses (OLAP)



Lakehouse engine



NoSQL engines



Graph engines

Data is more complex

80% of time spent on data cleaning, integration and preparation



Geospatial



Clickstream / Social media



Real-time streaming



ML models



Sensor/IoT



Graph



Video/audio



Images



Documents

Data infrastructures are more flexible

50% of data is in the public cloud, and rapidly growing



Traditional on-premises



Public cloud



Multi-cloud



Hybrid

Making it challenging to break down [data silos](#), [share data](#) responsibly, and accelerate data-driven [insights in a cost-effective way](#).

Data leaders and consumers need a simple way to manage complexity with a [single](#) database [engine](#) built for [any workload](#), [data type](#), and [skillset](#).

Db2: Achieve your strategic objectives



One database engine
for all your mission
critical workloads
across hybrid multi-
cloud environments



Mission critical apps & services



Real-time analytics & BI



Data science & machine learning

Strategic objective

Modernize global operations

Make decisions with real-time data

Adopt machine learning & AI

Pervasive challenges

Any downtime has a real
dollar impact to my
business

Applications need to
scale faster

Risk of embarrassing &
costly data breach is top
of mind

Cost of running analytics
is exceptionally high

Dashboards are too slow
and data is stale

Data is duplicated or
siload

Data quality severely
affects ML model
accuracy and adoption

Infrastructure challenges
and security risks of
moving data

How Db2 will help

Continuous availability of
applications with industry
leading SLAs

Near-infinite scale to
support any sized
workload

Security at the core with
built-in encryption & data
access controls

Control costs with a
cloud-native architecture
based on object storage
and elastic scale

Ingest data continually
across diverse sources

Access to open data
formats and secure
integration of your data
lake

Build, evaluate and train
ML models inside the
warehouse, using
curated and complete
data sets

Deploy models inside
Db2 to minimize data
movement and maximize
performance

Db2:

Cloud-first strategy

Organizationally aligned under single leadership structure fully dedicated to cloud first.

All new Db2 releases coming to cloud first, to be followed by the traditional software release. Innovation on cloud available for self-managed deployments as well.

Deploy fully managed services to IBM Cloud as the primary cloud of choice with self-managed reference architectures for others.

Engaging partnership opportunities to accelerate the delivery of fully managed DBaaS to hyperscalers.

IBM Db2 + Amazon Web Services

Partnering closely with Amazon to bring our extensive IP to AWS as fully managed services.

Db2 Warehouse led with the very first deployment on AWS – several years ago.

Other offerings available, including:

- **Db2 PureScale on AWS**
- **Db2 cloud reference architecture**
- **All offerings running on latest Db2 11.5.8**, a cloud-first deployment
- Others on the horizon including **Db2 (for OLTP workloads)**

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

Db2 Warehouse

SaaS on IBM Cloud and AWS



a fully-managed,
high performance,
elastic cloud data
warehouse

Db2 Warehouse

on IBM Cloud and AWS



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

Fully managed

by our world-class
CloudOps team, 24x7x365

***Focus on the analytics.
We'll take care of the rest.***

Our DevOps & CloudOps team is composed of engineers and experts in managing Db2 and cloud-based services.

We take care of...

- operating system updates and security patches
- new offering features and capabilities, including engine updates and console enhancements
- rolling out new hardware releases, when available
- unexpected software and hardware failure

...we're on-call 24x7x365, so you can focus on more important things.

And, if we should happen to fall short of our commitments to you, we stand behind a **99.9% SLA**.

Blazing fast Db2

Enterprise scale meets cloud

*Columnar store with in-memory processing on
compressed data sets, with option to create row-
based tables as well*

Column-organized

Tuned and optimized for analytics workloads

In-memory processing

High-speed, in-memory processing of large data sets, column-organized for fast analytics

Query compressed data

Your data is automatically compressed. When performing operations on those data sets, we can intelligently skip de-compression of certain data types, saving you time

Data skipping

Engine intelligently pulls only the data required to satisfy the query, leaving everything else on disk

Plus...

- industry leading SQL compliance
- native encryption for data in motion and at rest
- APIs for 15+ programming languages
- built-in auditing
- row/column access control
- data masking
- adaptive workload management
- time travel query
- query federation
- data lake integration with external tables
- in-database ML
- ML query optimizer
- JSON, XML, geospatial, graph
- Oracle & Netezza (PostgreSQL) compatible
- award-winning management console

Scalable

Independent scaling of storage and compute


Scale up your compute during peak demand

Scale down your compute when demand falls

Ramp up your data storage as your data needs grow

Compute & storage

[Show settings](#) ✕



Notification

Notify me when scaling is done:

☐ Email ☐ Slack

[Show API](#) ▾

Continuously available

with managed compute, storage
and cross-cloud replication



Compute

Completely automated failure detection and recovery means zero maintenance for our customers.



Storage

High-performance, cloud-managed SSD-based block storage in a redundant, highly available configuration.



Cross-region, cross-cloud replication

Continuously available, cross-region, cross-cloud active-active replication available to satisfy the toughest availability & RPO/RTO requirements.

Reliable

with snapshot-based backups

Schedule your backups

to run when it's most convenient for your business

Lightning fast restore

from snapshot if you should need it

Self-service snapshot backup

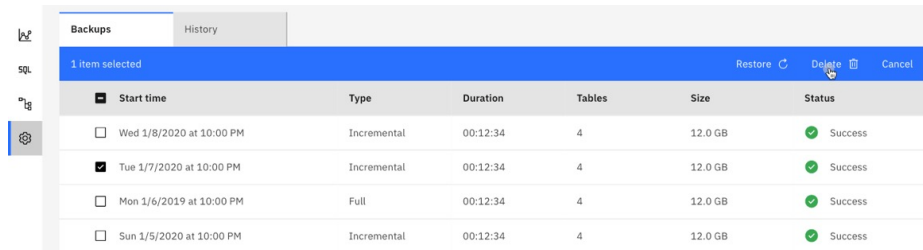
Choose the time of day for your backup to occur.

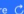
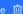




Self-service snapshot restore

Choose which backup to restore. Get your instance up and running again in minutes.

Backups for disaster recovery

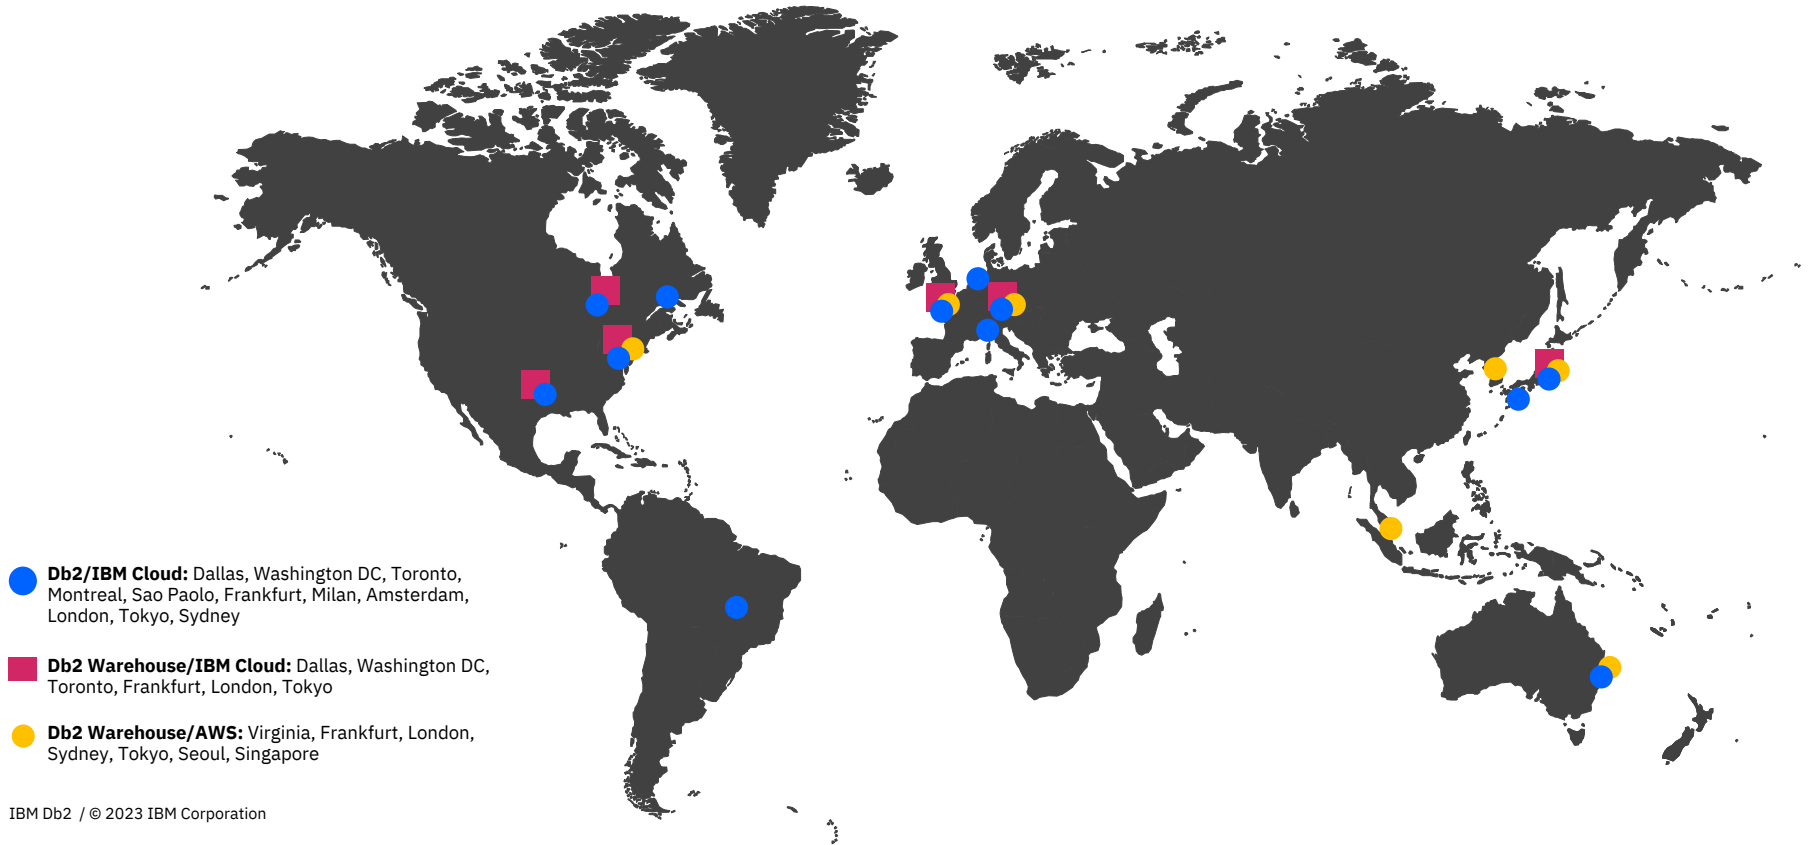
Automated backups saved to object storage, replicated across multiple availability zones and regions



Backups		History				
1 item selected						Restore  Delete  Cancel
<input checked="" type="checkbox"/>	Start time	Type	Duration	Tables	Size	Status
<input type="checkbox"/>	Wed 1/8/2020 at 10:00 PM	Incremental	00:12:34	4	12.0 GB	 Success
<input checked="" type="checkbox"/>	Tue 1/7/2020 at 10:00 PM	Incremental	00:12:34	4	12.0 GB	 Success
<input type="checkbox"/>	Mon 1/6/2019 at 10:00 PM	Full	00:12:34	4	12.0 GB	 Success
<input type="checkbox"/>	Sun 1/5/2020 at 10:00 PM	Incremental	00:12:34	4	12.0 GB	 Success

IBM Db2

SaaS Supported Regions



Client outcomes improved through IBM Db2 Warehouse SaaS



Personalize

Ability to market to leisure travelers in a much more personalized way, understanding who they are, what they need and the right way to reach them

↑ **90%** performance improvement
↑ **140** million-plus members served



Automate

Helping the company purchase media at the best value, optimize media pricing proposals, and track clients' media spend

↑ **\$80M** in unlocked billing opportunities
↓ **100s** of hours saved with AI-powered prospecting



Automate

Building robust sales analytics helps to drive customer engagement, both online and in-store

↓ **90%** reduction in reporting time on sales data
↑ Improved insights on complex data captured from all stores



Secure

Building a strong machine-learning foundation to better understand and meet the needs of the audience

↑ Ability to react quickly and optimize promotional campaigns
↑ Improved infrastructure to manage that data, data models and personnel.

A Different Approach to Cloud Data Warehousing



Analytics on massive quantities of data has always been expensive

Reality: Customers turned to cloud data warehouses as a way to drive costs down with pay-as-you-go pricing.



The modern enterprise is a global operation running 24x7x365

Reality: The new breed of cloud data warehouses are built on business model they are not always needed



Workloads have diversified and data types have become more complex

Reality: Customers are still struggling with data silos across warehouses and lakes, leaving 80%¹ of enterprise data go unused for analytics

1 Wakefield Research, The Data Activation Report, April 2022

Need: Customers are left seeking more flexibility over how they scale always-on, analytics workloads in a cost-effective way and greater access to collaborate on governed, shared data.

Db2 Warehouse

Introducing the third generation of our cloud data warehouse

1

Cut costs while improving performance

- Runs on inexpensive, cloud-native object storage (Amazon S3-based)
- ↓ storage costs over traditional block-storage
- ↑ performance with advanced caching

2

Share data responsibly

- Support for open data formats formats (Iceberg, Parquet, ORC, CSV and more)
- Single, governed, cloud-native Db2 engine over object storage with data lake integration

3

Take control over compute and storage

- Independent scale of storage and compute while keeping critical applications always-on.

The fully managed cloud data warehouse built with multi-petabyte-scale, multi-performant storage

Db2 Warehouse Gen3

Introducing the third generation of our cloud data warehouse

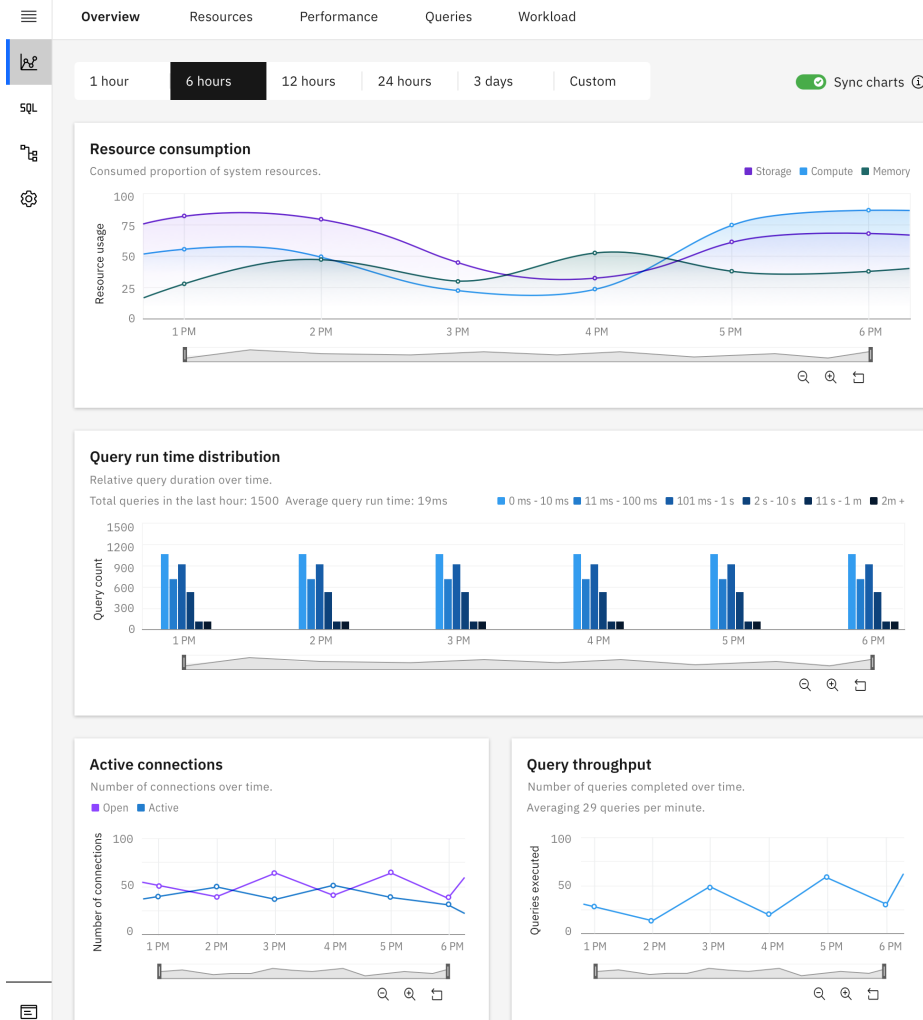
Key new features:

- **Amazon S3 support** for table storage, significantly decreasing the cost of storing data without sacrificing performance
- **Import & export multiple open data formats** (Iceberg, Parquet, ORC, CSV and more) leveraging existing compute resources dedicated to the warehouse
- **Integration with data lake** with sharing of data and metadata through use of metastores, catalogs and S3 buckets
- **Fully managed cloud data warehouse** scaling up to 2880 cores (5760 vCPUs) per cluster, multi-petabyte-scale, multi-performant storage

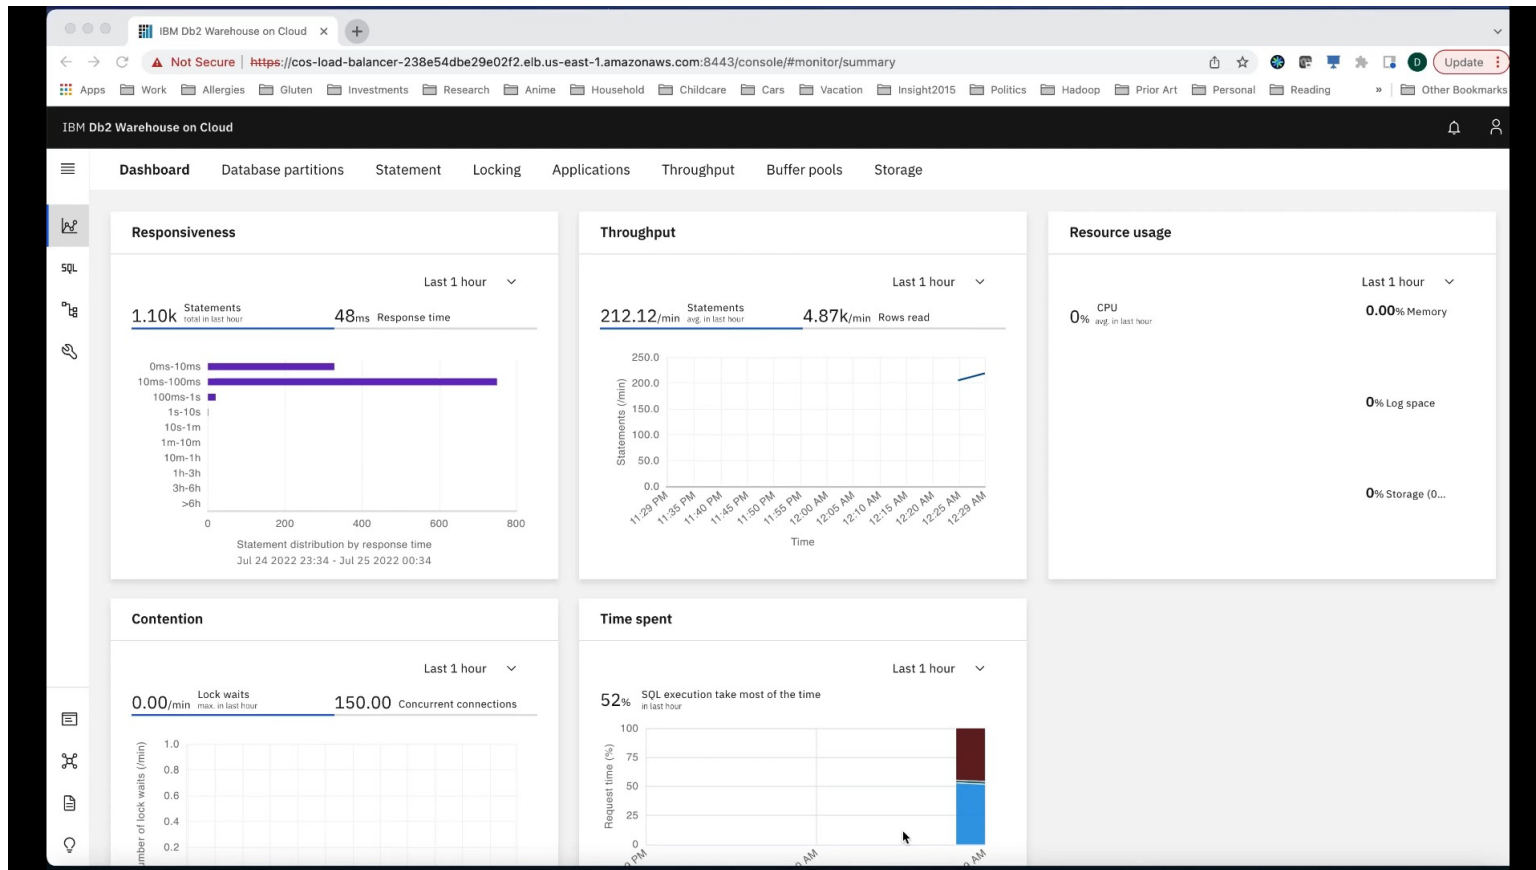
Other features:

- **Support for multiple storage options:** Amazon S3 and Block Storage
- **Storage auto-increase for Block Storage** on set threshold ensuring you never run out of storage for your workloads
- **Granular, schema-level backup/restore** to S3, restoring only the data you need
- **Cross-region snapshot backup** to Amazon S3 for disaster recovery
- **Self-service maintenance windows** for product and database engine updates
- **Integration with IBM AppID** for easy integration with Azure Active Directory
- **New APIs** for scaling, updates, backup/restore, logging

IBM Db2 / © 2023 IBM Corporation



Demo: Native Db2 tables in Amazon S3



Db2 Warehouse Gen3

Preliminary performance numbers comparing Db2 Warehouse Block vs Object Storage

2.6x

**Faster SELECT
performance¹**

**Warm local NVMe cache,
test with BDI Serial (5TB)**

- Db2W Block Storage:
2289.79s
- Db2W Object Storage:
897.08s

2.8x

**Faster INSERT
performance¹**

**Insertion of 14.5B rows into
the warehouse**

- Db2W Block Storage:
984.556s
- Db2W Object Storage:
384.674s

7.2x

**Faster UPDATE
performance¹**

**Update 10% of 14.5B rows
at random**

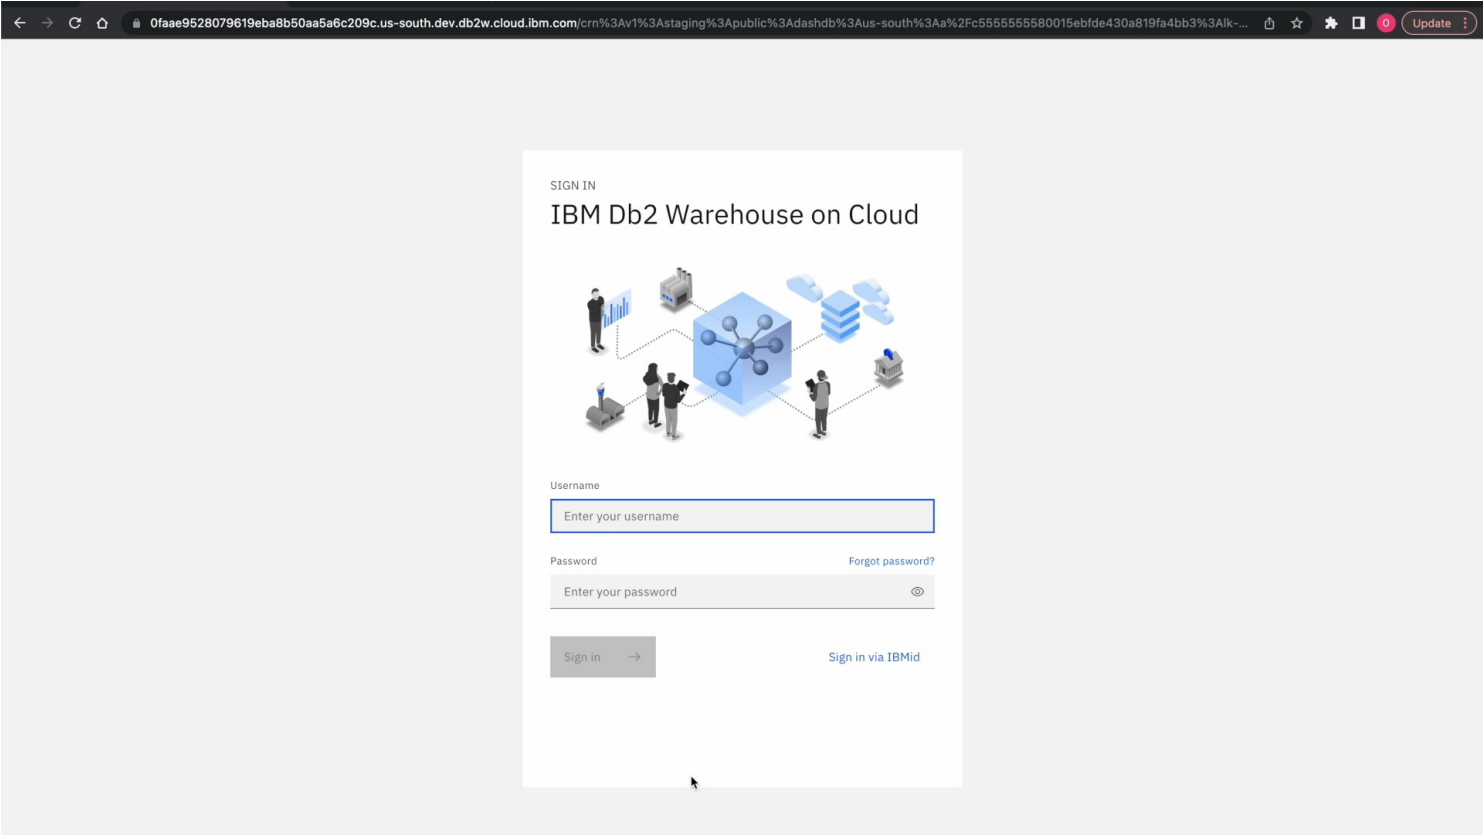
- Db2W Block Storage:
850.919s
- Db2W Object Storage:
118.342s

38x

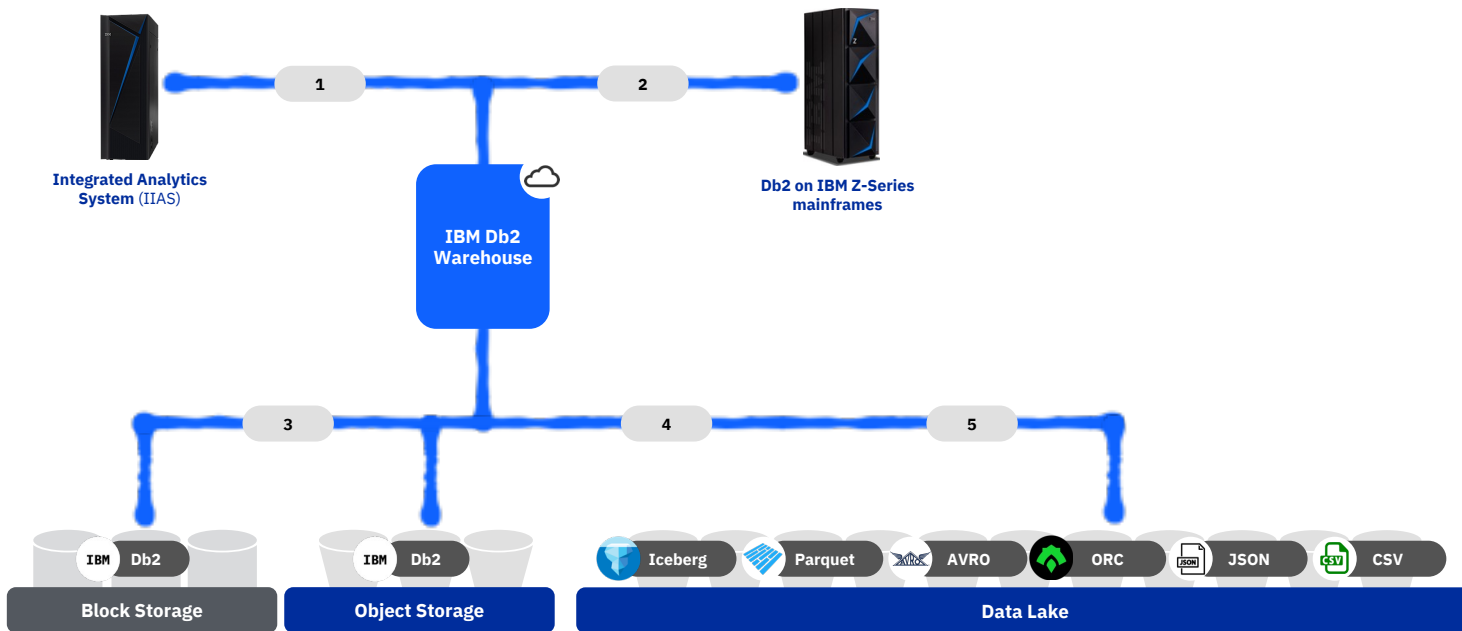
**Lower cost for object
vs block storage²**

\$23/TB/mo vs \$874/TB/mo

Demo:
Data Lake tables
using Open Data
Formats



Journey from Db2 Warehouse to Data Lake and back



- 1. Modernize the warehouse**
Upgrade from your existing on-premises analytics appliance to the cloud
- 2. Replicate data from Db2 on Z to Db2 Warehouse with live twinning**
Share data from Db2 on Z with your Db2 Warehouse & data lake
- 3. Cost-optimize cloud data warehouse storage**
Take advantage of inexpensive, highly reliable object storage
- 4. Share warehouse data with the data lake**
Join data from your warehouse with your data lake for new insights
- 5. Promote data lake data to the warehouse**
Easily promote curated data lake data to Db2 Warehouse

Db2 Warehouse Gen3
**Third generation
cloud data warehouse**

Launching in
Beta – April 17



Db2 Warehouse

Power **real-time** analytics and BI with high performance and **cost effective** cloud-native data warehousing

Benefits

- **Enable real-time insights** when you feed your data to Db2 Warehouse
- **Get the best price-performance** for your big data with elastic scale and a cloud-native architecture based on object storage
- **Collaborate on data projects** with governed data access and sharing
- **Unify your data with a single engine** that supports tight integration with your data lake over cloud object storage
- **Simplify database administration** with fully managed cloud service or on-premises container-based deployments

Use Cases

- Power real-time analytics with continual data ingest, in-memory processing, and in-database analytics
- Create, evaluate, and predict with high-performance machine learning models directly inside your data warehouse.
- Quickly build reports and dashboards by seamlessly integrating Db2 with a variety of BI tools including *IBM Cognos, Tableau, Microsoft PowerBI, Google Looker*, and more
- Securely grant access and distribute data among users and 3rd parties without data duplication.

