IDUG

2026

Sydney | March 16 - 18

AU Db2 TECH CONFERENCE

SQL Data Insights: Past, Present, and Future
A Deep Dive into New Capabilities

Jae Lee, IBM

Session Code: A05

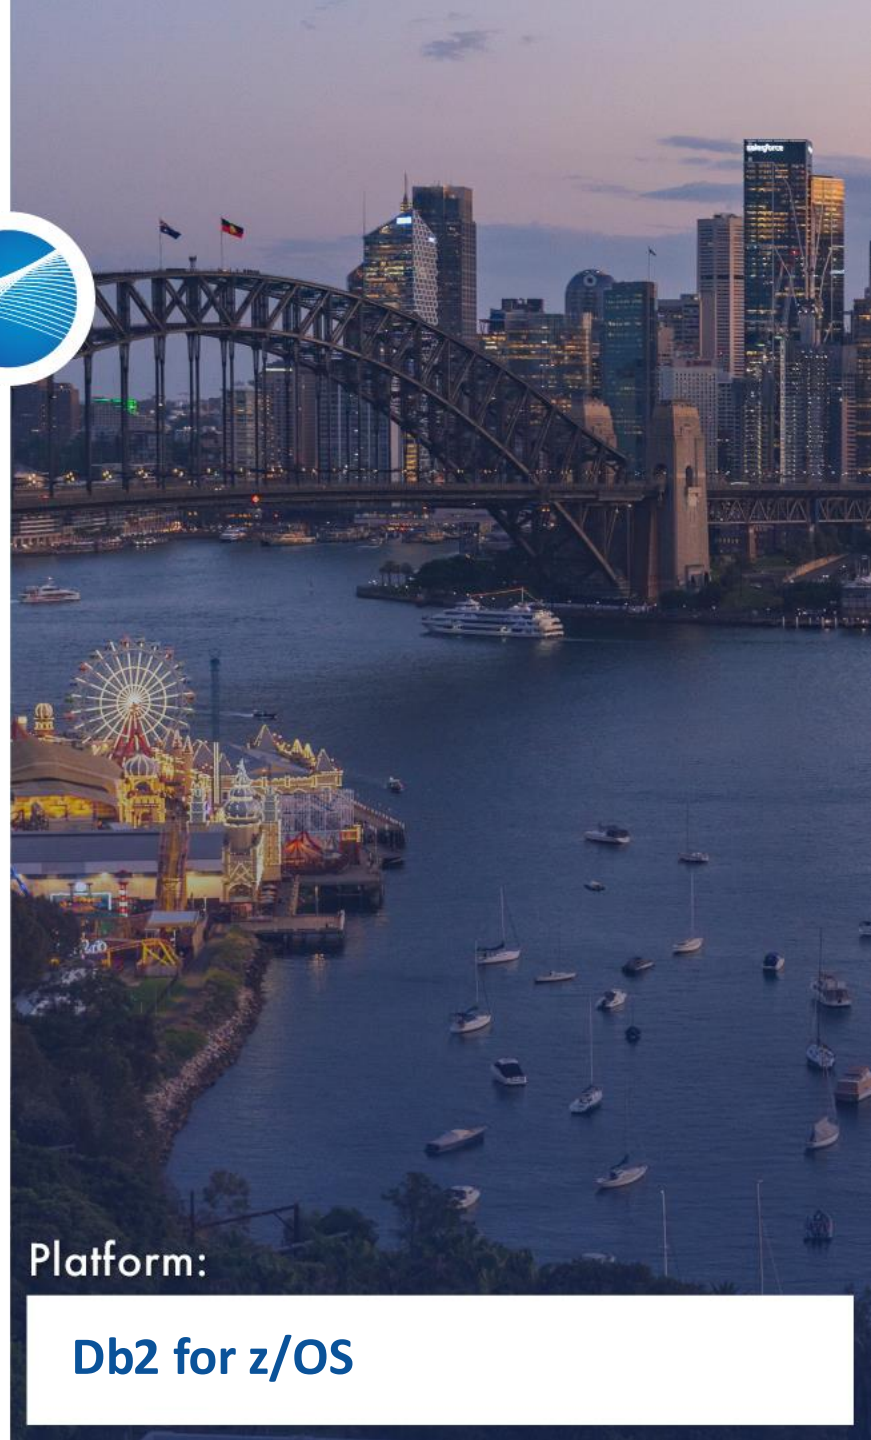

Platform:

Db2 for z/OS

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

- Introduction to SQL Data Insights
- Major Enhancements
- SQL Data Insights vNext
- Use-case on Insurance Company

What is SQL Data Insights?

- AI-powered downloadable feature of Db2 13 to let you exploit your business data through semantic AI functions
- Builds an unsupervised, Large Database Model against the Db2 data, to answer business questions
  - No data move necessary
- Integrates AI within any Db2 application via SQL
  - Semantic AI-built in functions
  - No data scientist required
- Leverage IBM Z AI functions (zAIU on Telum I and II) and zIIP enabled

What Questions Can you ask with SQL Data Insights?

Traditional SQL

- Find customers with the same attributes
- Which transactions match fraud rules
- List products in the same category
- Identify transactions outside predefined threshold

Exact match conditions that may require complex SQL

SQL Data Insights

- Find customers similar to this customer
- Which transactions look similar to a fraudulent transaction
- Discover products similar to this product
- Which transactions are outliers

SQL Data Insights: Similarity search using AI models

Large Language Models (LLMs)

generated using text data

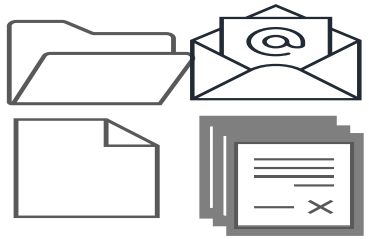

Unstructured data (emails, blogs, tweets, PDFs…)

Chunks

LLM

–OpenAI's GPT-4
–Google's Gemini, PaLM, BERT
–Meta's LlaMA
–IBM's Granite

Generate embedding per chunk

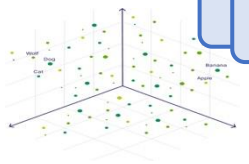

Sentence Embeddings

What snack is like Oreo?

Large Database Models (LDMs)

generated using multi-modal relational data

Key

CODE	PRODUCT_NAME	BRANDS	ENERGY_100G	FAT_100G	SUGARS_100G	PROTEIN_100G	SODIUM_100G
16087	Organic Salted Nut Mix	Grizzlies	2540.0	57.14	3.57	17.86	0.482
16124	Organic Muesli	Daddy's Muesli	1833.0	18.75	15.62	14.06	0.055
16872	Zen Party Mix	Sunridge	2230.0	36.67	3.33	16.67	0.633
18012	Cinnamon Nut Granola	Grizzlies	1824.0	18.18	21.82	14.55	0.009
18050	Organic Hazelnuts	Grizzlies	2632.0	60.71	3.57	14.29	0.004
18227	Organic Oat Groats	Pcc	1096.0	5.95	2.38	16.67	0.01

Numeric

Categorial

Structured relational data

Textification and clustering

Code_16087 Product_name_Organic Salted Nut Mix
Brands_Grizzlies Energy_100G_c7 Fat_100G_C4 …

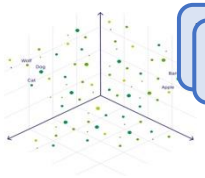

Database Embeddings

Generate embedding per unique values for each column and unique row

AI_SIMILARITY('44000007492', CODE)

Traditional Way to Train Data in Mainframe

Access Data sources

Db2
IMS
VSAM

Prepare Data

Copy Data
Format Data
Cleanse Data
Reduce Data
Label Data
…

Weeks of preparation

Data Engineer DBA

Model training

Choosing a model
Learning
Evaluating, Testing, Tuning
…

Weeks/Months of training

Data Scientist

Deployment

Access transactional data
Operationalize on production
HA/DR
Performance
…

Weeks of deployment

Application Developer

SQL DI Way

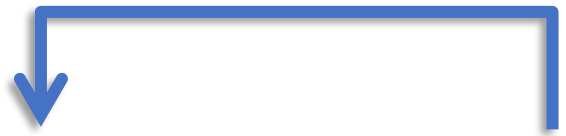

Access Data sources

Db2
IMS
VSAM

Prep to Deployment

Hours

1. Choose table & relevant columns

2. Kickoff training

3. Run semantic queries on a model

4 Ways to ask Db2

AI_SIMILARITY	• Find similar entities	FL 500
AI_SEMANTIC_CLUSTER	• Find entities that belong to a group of entities	FL 500
AI_ANALOGY	• Analyze relationship between a pair of entities	FL 500
AI_COMMONALITY	• Find the most common entities, or outliers	FL 504

AI Semantic Queries (1|2)

- AI_SIMILARITY

```
SELECT AI_SIMILARITY(CUSTOMERID, '3668-QPYBK'), CHURN.*
FROM CHURN
ORDER BY 1 DESC
FETCH FIRST 5 ROWS ONLY
```

ORDER BY 1 ASC makes DISSIMILARITY

- AI_SEMANTIC_CLSUTER

```
SELECT AI_SEMANTIC_CLUSTER(CUSTOMERID,
                           '0280_XJGEX', '6467_CHFZW', '0093_XWZFY'), CHURN.*
FROM CHURN
ORDER BY 1 DESC
FETCH FIRST 5 ROWS ONLY
```

AI Semantic Queries (2|2)

- AI_ANALOGY

```
SELECT AI_ANALOGY( 'Month-to-month' USING MODEL COLUMN CONTRACT,
                   'Electronic check' USING MODEL COLUMN PAYMENTMETHOD,
                   'One year',
                   PAYMENTMETHOD), I.*
FROM CHURN
WHERE CONTRACT = 'One year'
ORDER BY 1 DESC
FETCH FIRST 5 ROWS ONLY
```

- AI_COMMONALITY

```
SELECT AI_COMMONALITY(CUSTOMERID), CHURN.*
FROM CHURN
ORDER BY 1 DESC
FETCH FIRST 5 ROWS ONLY
```

ORDER BY 1 ASC
makes the least common - Outliers

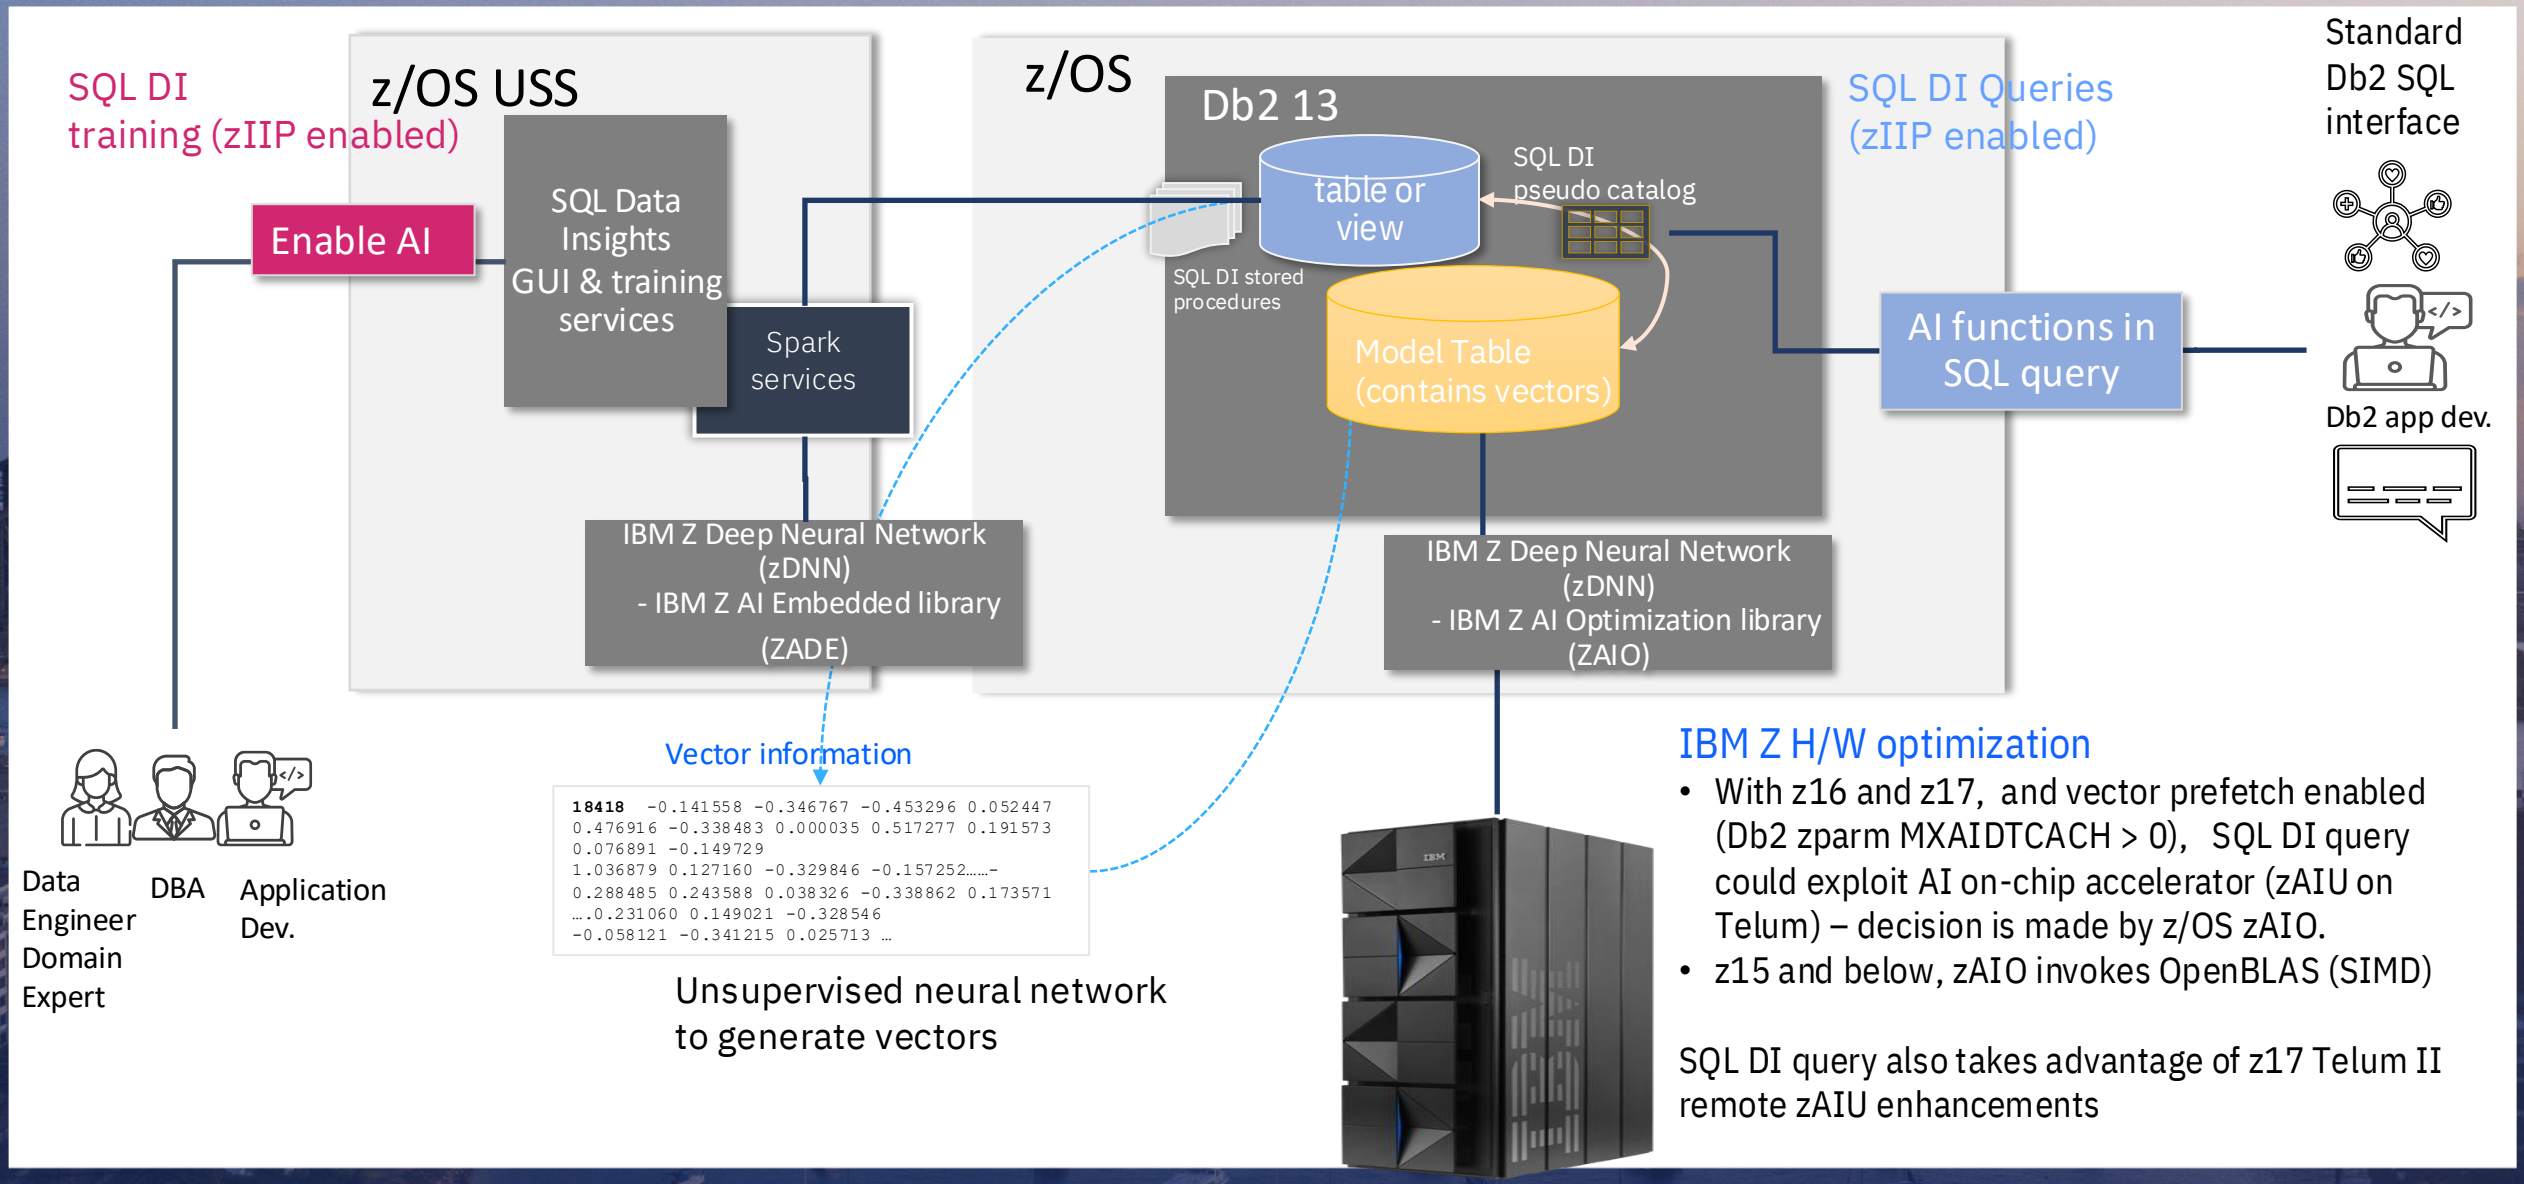

IBM Z H/W optimization

- With z16 and z17, and vector prefetch enabled (Db2 zparm MXAIDTCACH > 0), SQL DI query could exploit AI on-chip accelerator (zAIU on Telum) – decision is made by z/OS zAIO.
- z15 and below, zAIO invokes OpenBLAS (SIMD)

SQL DI query also takes advantage of z17 Telum II remote zAIU enhancements

Evolution of SQL DI

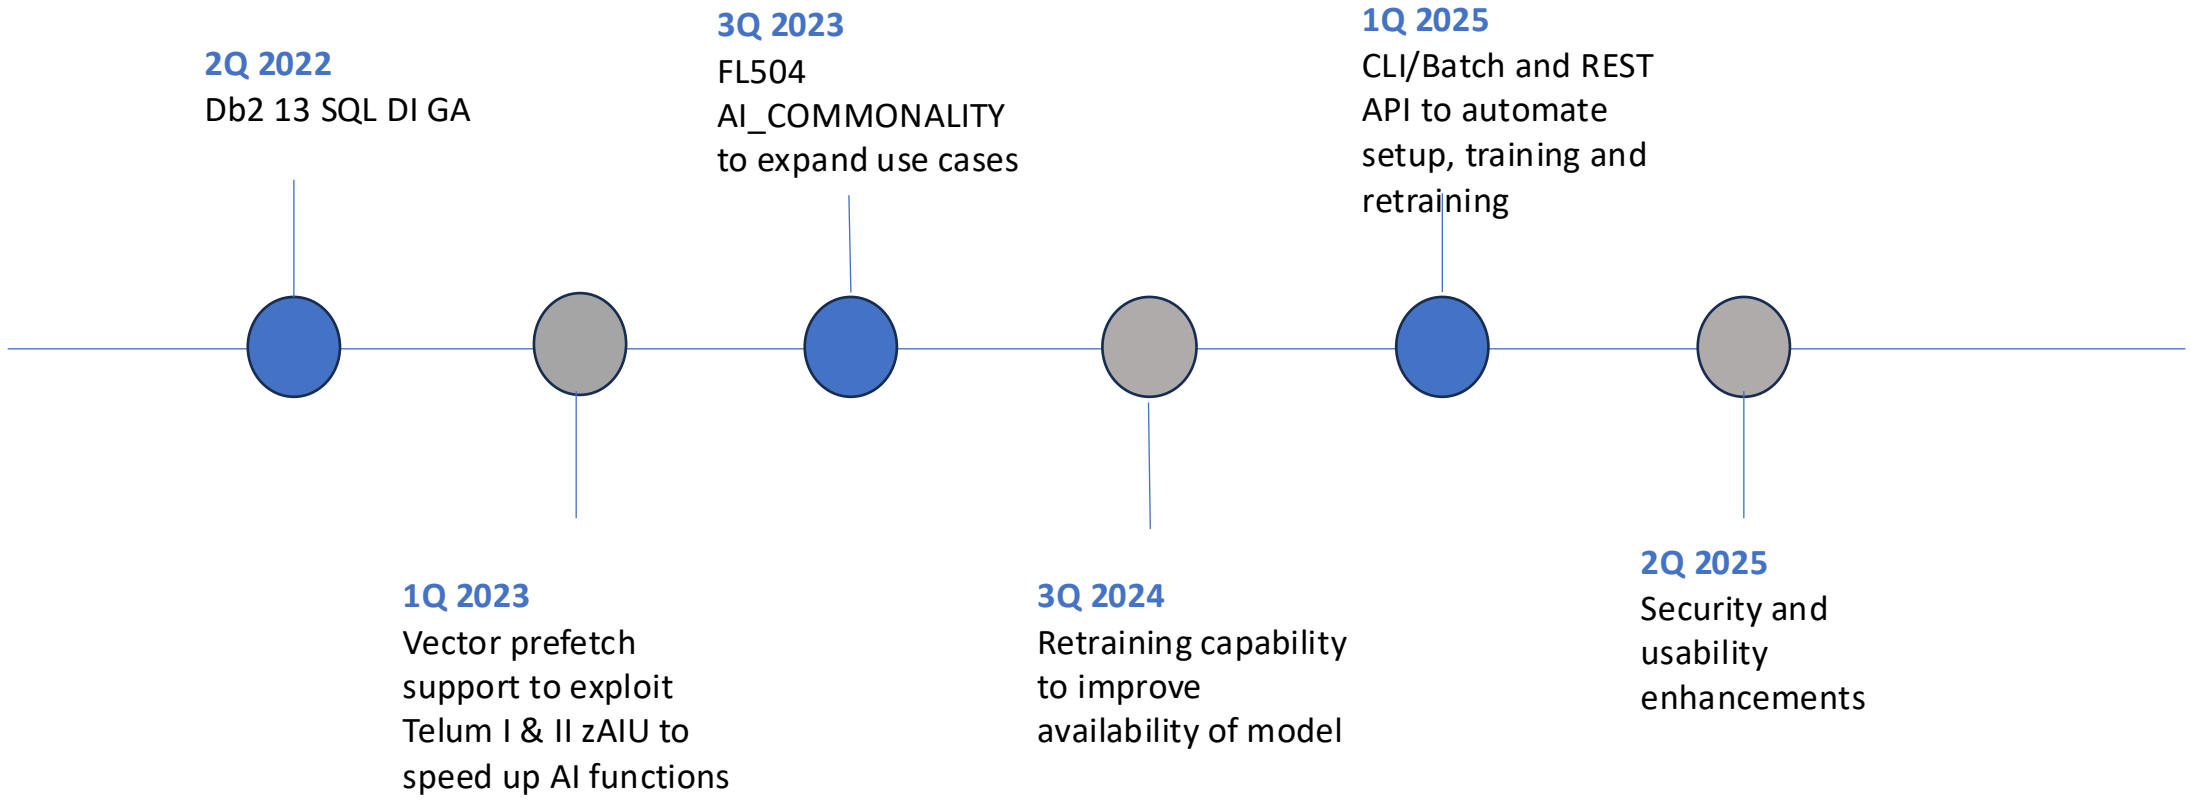

Enhancements in 2025 (1|3)

- REST API and Command Line Interface (CLI)
  - REST API can be called by any applications and CI/CD pipelines, enable automation
  - CLI can be invoked from JCL, makes is able to schedule batch jobs for model training and retraining
  - Tighten security by using Auth Token & PASSTICKET to access SQL DI & Db2

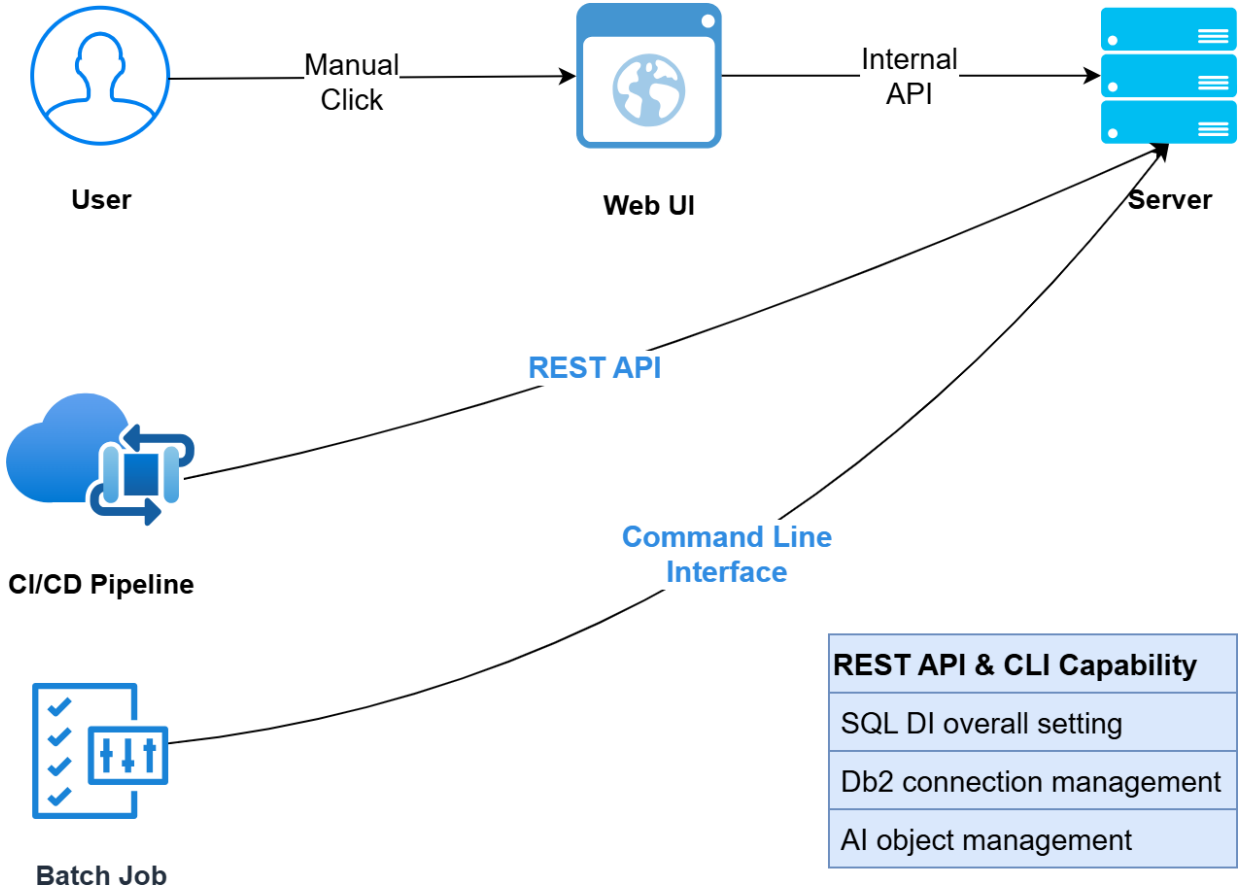

REST API & CLI Capability
SQL DI overall setting
Db2 connection management
AI object management

Enhancements in 2025 (2|3)

- Sample REST APIs and CLI commands

POST	/sqldi/v2/connections	Add a connection
PATCH	/sqldi/v2/connections/{connId}	Edit a connection
DELETE	/sqldi/v2/connections/{connId}	Remove a connection
GET	/sqldi/v2/connections/{connId}/objects	Get AI objects for a connection
POST	/sqldi/v2/connections/{connId}/objects	Create and enable an AI object for AI query
DELETE	/sqldi/v2/connections/{connId}/objects/{object}	Remove an AI object
GET	/sqldi/v2/connections/{connId}/objects/{object}/config	Get column configuration for an AI object
POST	/sqldi/v2/connections/{connId}/objects/{object}/retraining	Start model retraining
DELETE	/sqldi/v2/connections/{connId}/objects/{object}/retraining	Remove a retrained model
POST	/sqldi/v2/connections/{connId}/objects/{object}/retraining/deployment	Deploy a retrained mod

- To create an AI object and enable it for AI query, issue the following command with the `-f` option:

```
./sqldi.sh create_object -f sqldi_object_model_input
```

- To view the enablement status and other details of an object, issue the following command with the `-f` option:

```
./sqldi.sh list_object -f sqldi_object_model_input -status_rc
```

`-status_rc` is optional. When specified, the response of the command includes the AI query enablement status of the object.

- To list all objects associated with a specific connection, issue the following command with the `-f` option:

```
./sqldi.sh list_all_objects -f sqldi_object_model_input
```

- To remove an AI object, issue the following command with the `-f` option:

```
./sqldi.sh remove_object -f sqldi_object_model_input
```

- To view the column configuration of a specific object, issue the following command with the `-f` option:

```
./sqldi.sh list_column_config -f sqldi_object_model_input -o <column_config_file>
```

Enhancements in 2025 (3|3)

- Support of Secondary Authorization ID
  - Manage user authorization using secondary authorization ID reduces security management effort

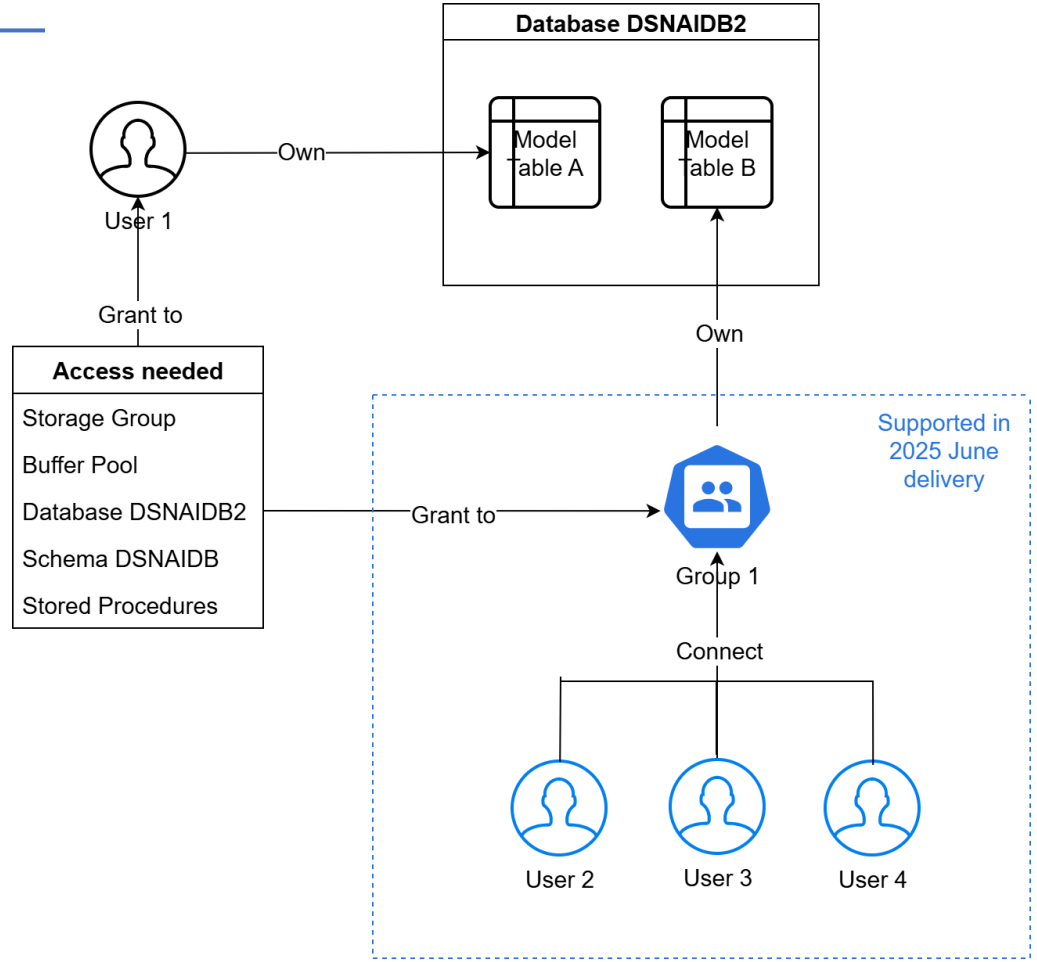

SQL DI as of Today

- Slow adoption... but catching up
  - Many clients have been consumed migrating to Db2 13
  - Not enough attention to the free feature
  - Not easy to find the use cases
  - Swiss Mobiliar have been a vocal advocate
  - FI shows their recent success on REST APIs and CLI feature
- Missing key features like
  - Support of unstructured data – 60% of clients have more than 1 TB of unstructured text in Db2 (logs, descriptions, comment, JSON, etc.)
  - Model retraining that takes less resources on production system

SQL DI Vnext Overview

Announcement: March 17
GA: March 20th

- Major new features
  - Utilizes IBM Db2 Analytics Accelerator (IDAA) to provide AI capability perform AI model training using Integrated Facility for Linux (IFL) processors
  - Supports unstructured data – comments, descriptions, notes – stored in Db2 character columns
  - Supports incremental retraining – saves time and resources
- The evolution into an advanced priced zStack offering, designed to unlock far deeper intelligence from Db2 for z/OS and IBM Z Data.

Packaging & Entitlement

- A standalone product
- Included Components
  - SQL DI (FMID HDBDD18)
  - IDAA 8.1 (FMIDs HAQTX10, JAQTX1D) - shopZ
  - IDAA SSC Installation image – FixCentral
  - License key
- Impact on Existing and New SQL DI and IDAA Customers

	Existing SQL DI Customer	New SQL DI Customer
Owns IDAA	• Upgrade to SQL DI 1.1.4 • Upgrade to IDAA 8.1.1	• Install SQL DI 1.1.4 • Upgrade to IDAA 8.1.1
Do not own IDAA	• Upgrade to SQL DI 1.1.4 • Install IDAA 8.1.1 (limited to SQL DI entitlement)	• New Installation of SQL DI 1.1.4 • Install IDAA 8.1.1 (limited to SQL DI entitlement)

Support Unstructured Data

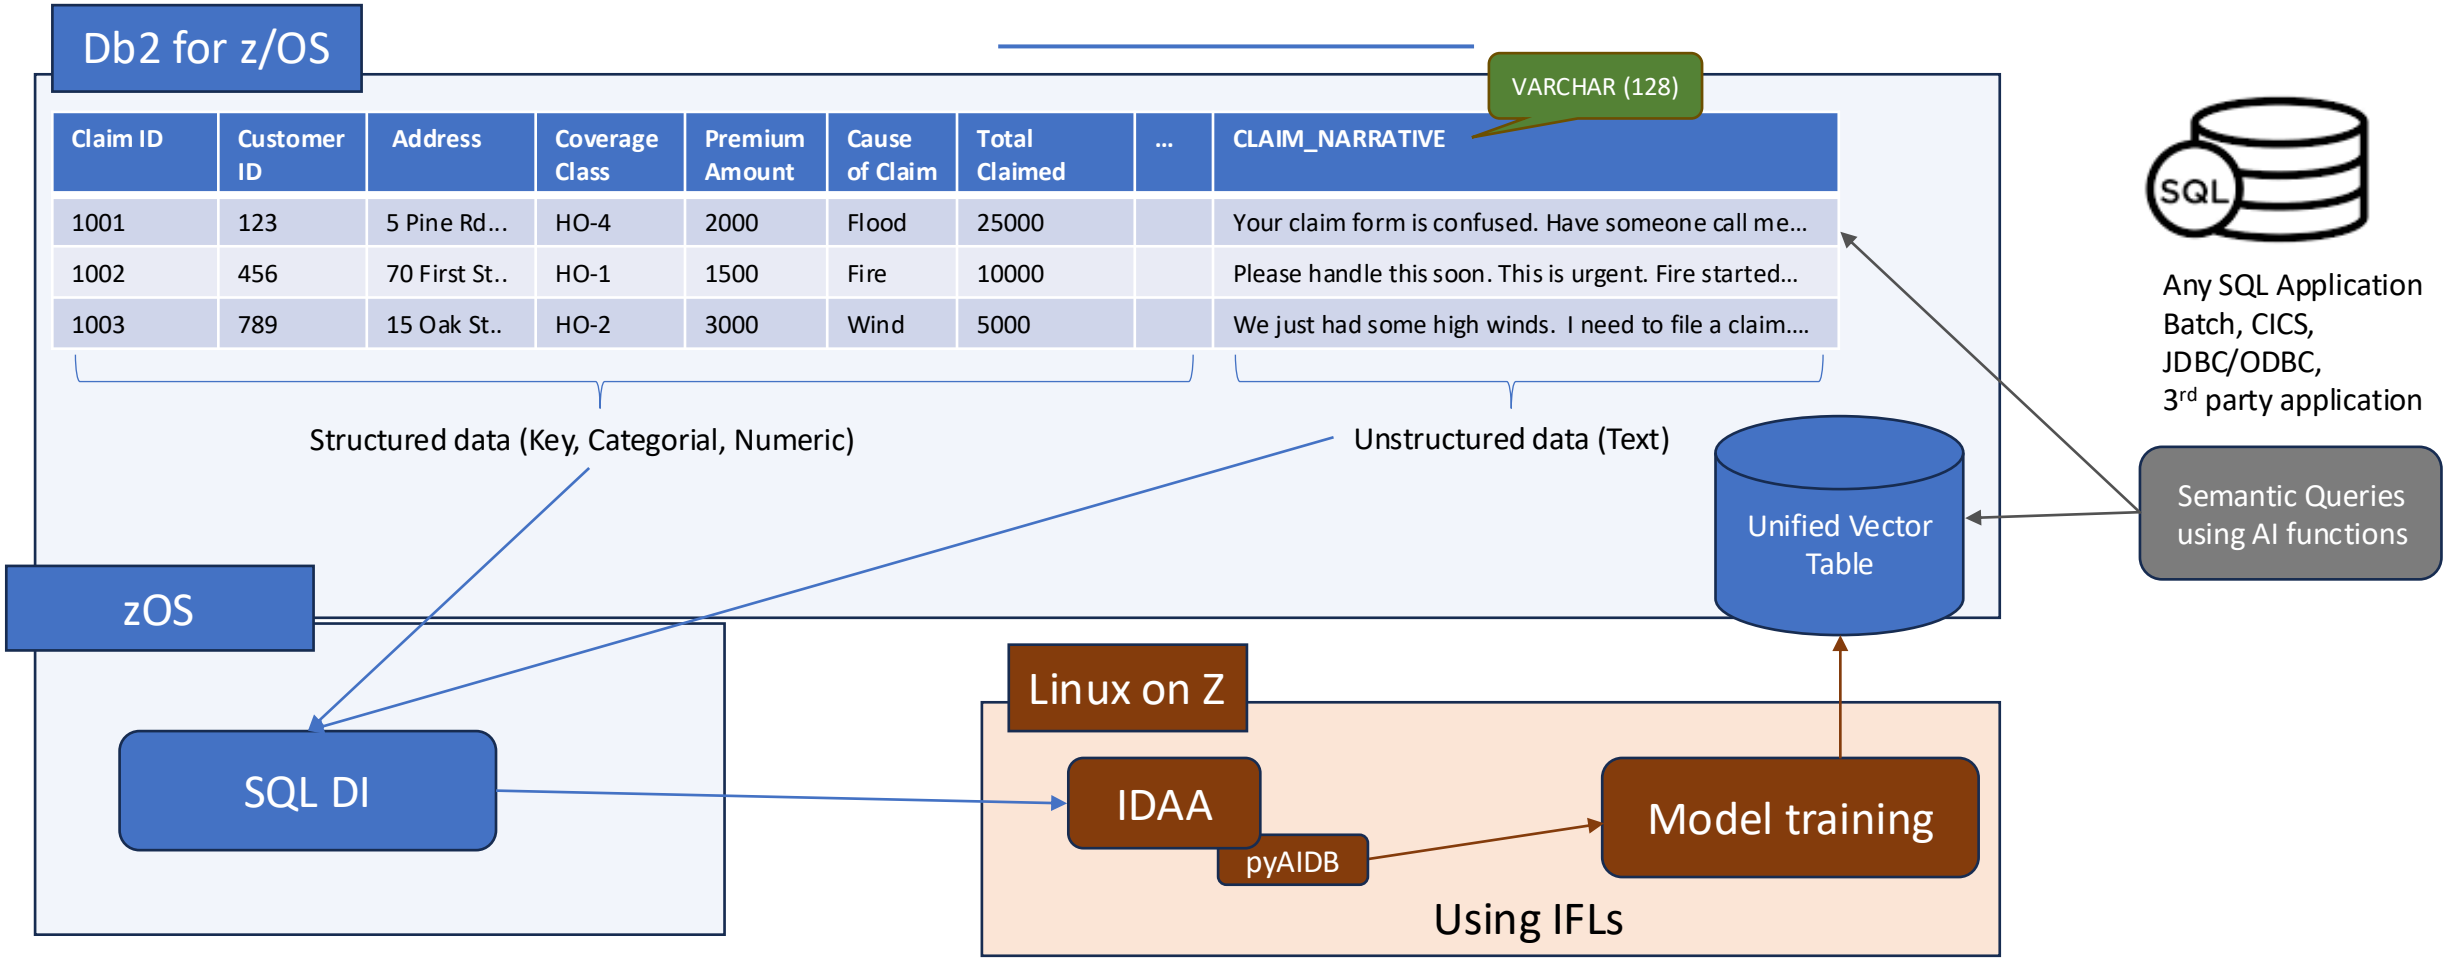

Incremental Retraining

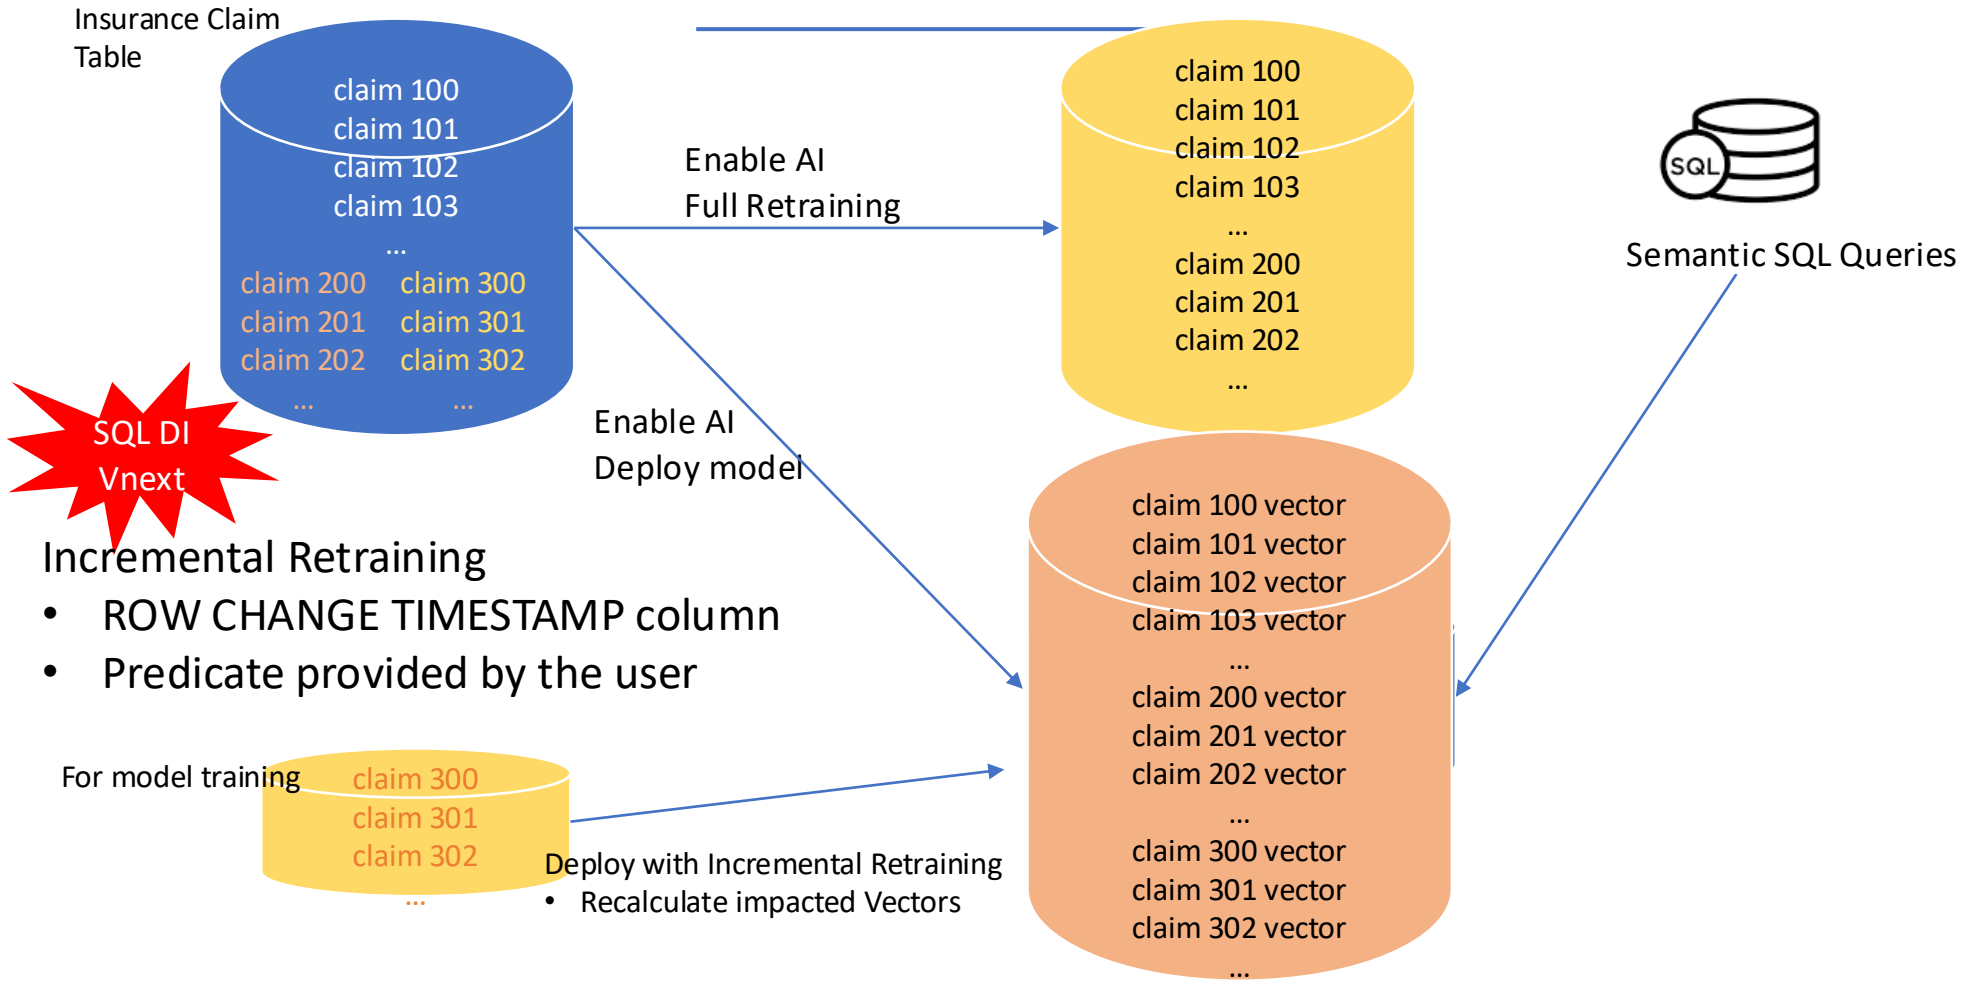

Retraining Guidelines: Full vs. Incremental

Scneario	Recommended Action
New rows Inserted only	Incremental retraining
Majority of data updated	Full retraining
Majority of data deleted	Full retraining
Incremental data > 75 % original	Full retraining
Hard to identify changed data	Full retraining
AI queries mainly use AI Key columns	Incremental retraining
AI queries mainly use non-AI Key columns (numeric or Categorical)	Full retraining

Insurance Use Case Overview

- 3 Insurance related tables
  - Insurance_Policy – 191 columns with policy coverage details
  - Insurance_Applicants – 149 columns with customer profile information
  - Insurance_Claims – 541 columns with narrative on claim used as Text
- Goal is to
  - Gain deeper insight into insurance claims
  - Identify policies that may be underpriced or undercharged
  - Identify policies and customers suitable for upselling additional coverage
- Objects used for training
  - Train three models based on individual tables
  - Create a view that INNER JOIN Policy and Applicants table and LEFT JOIN with Claims table (on policy_id column)

Connections

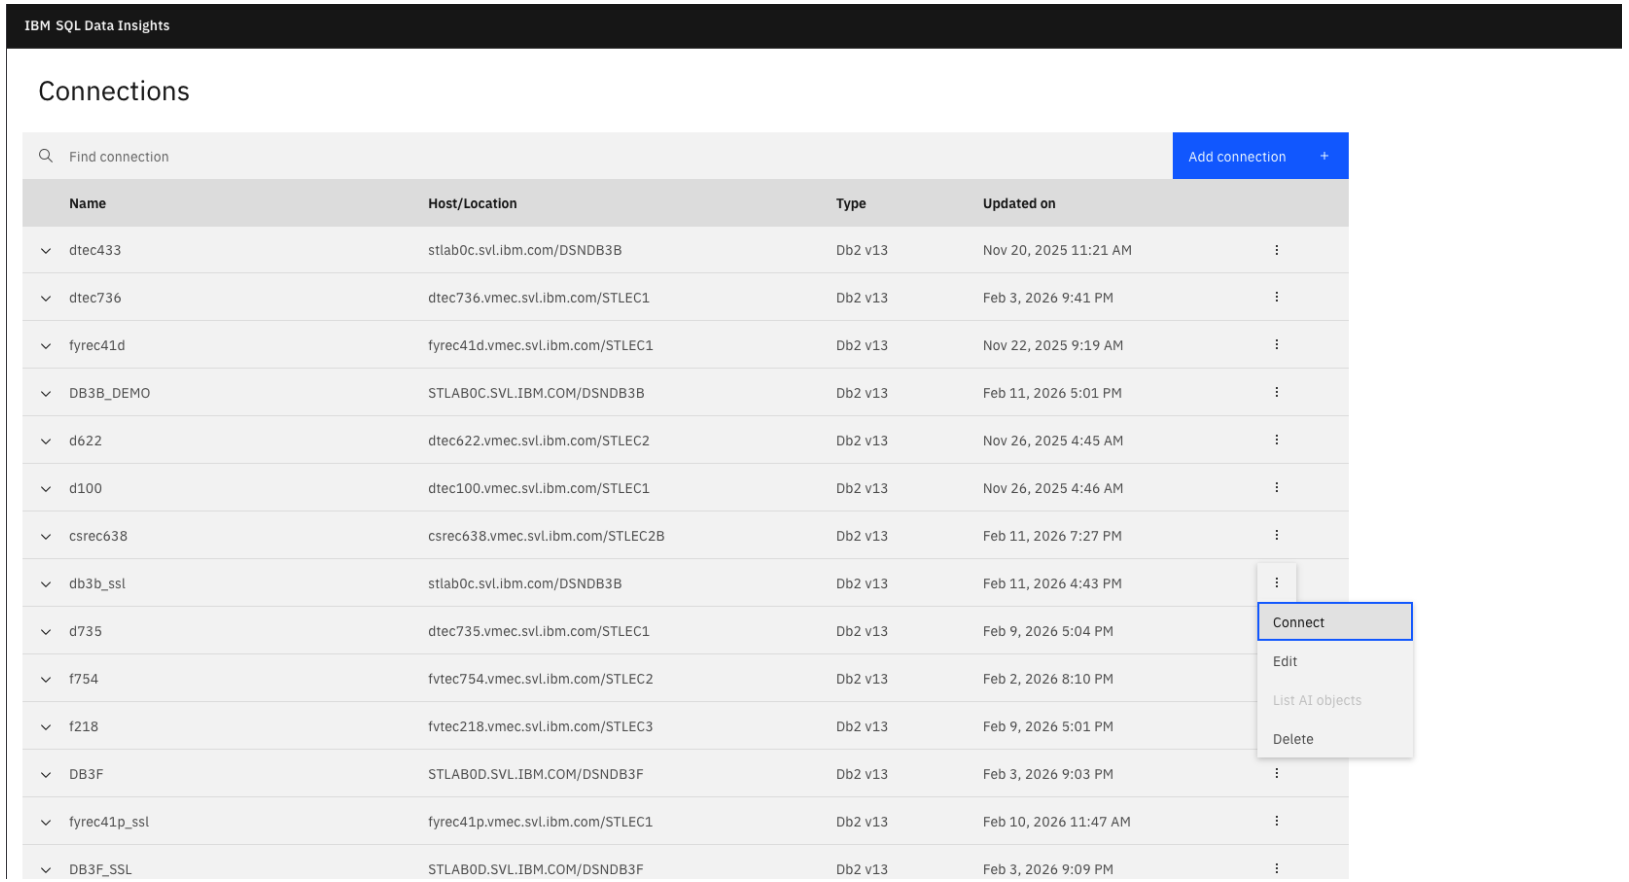

Enable AI (1|9)

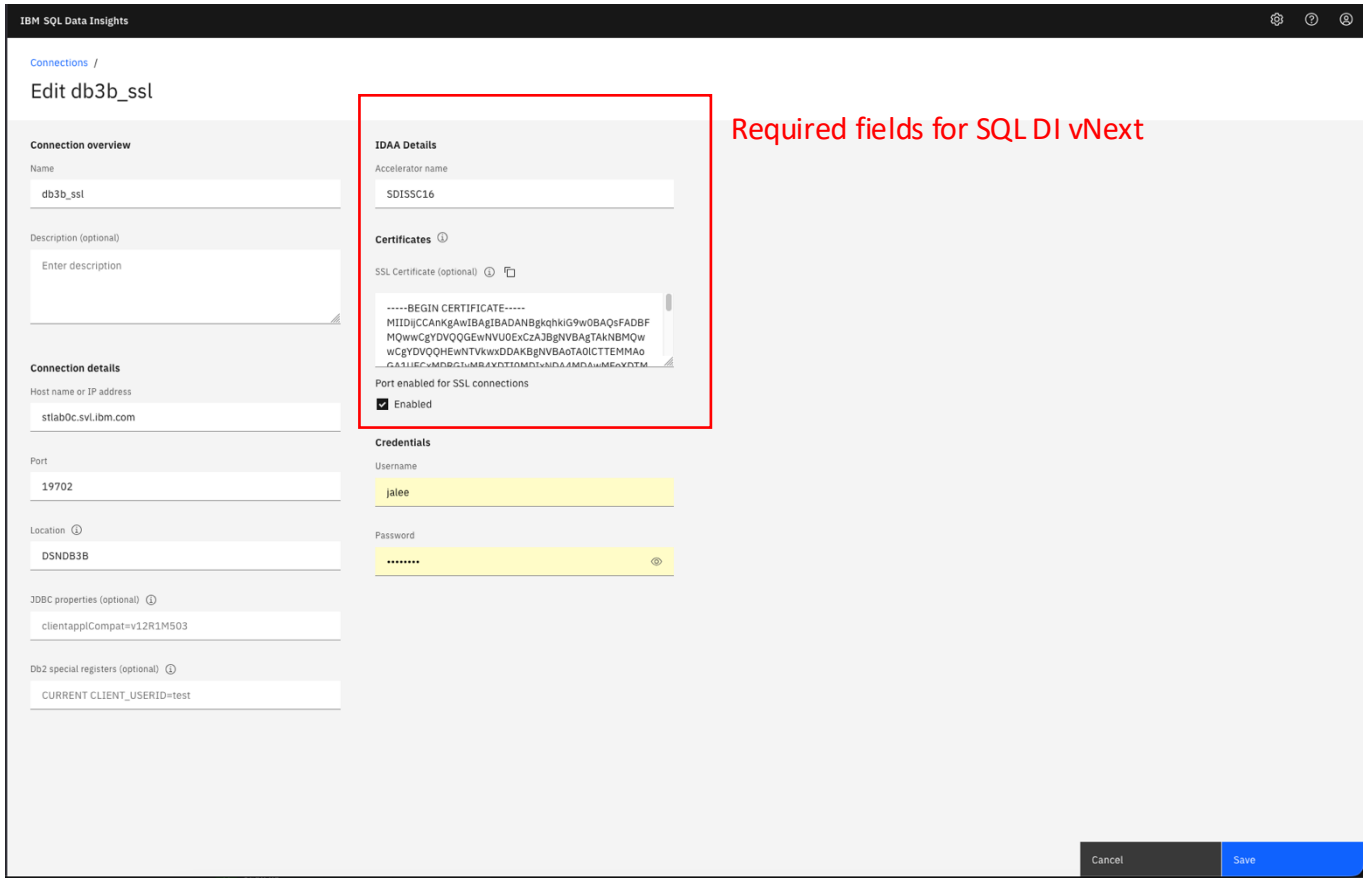

Required fields for SQL DI vNext

Enable AI (2|9)

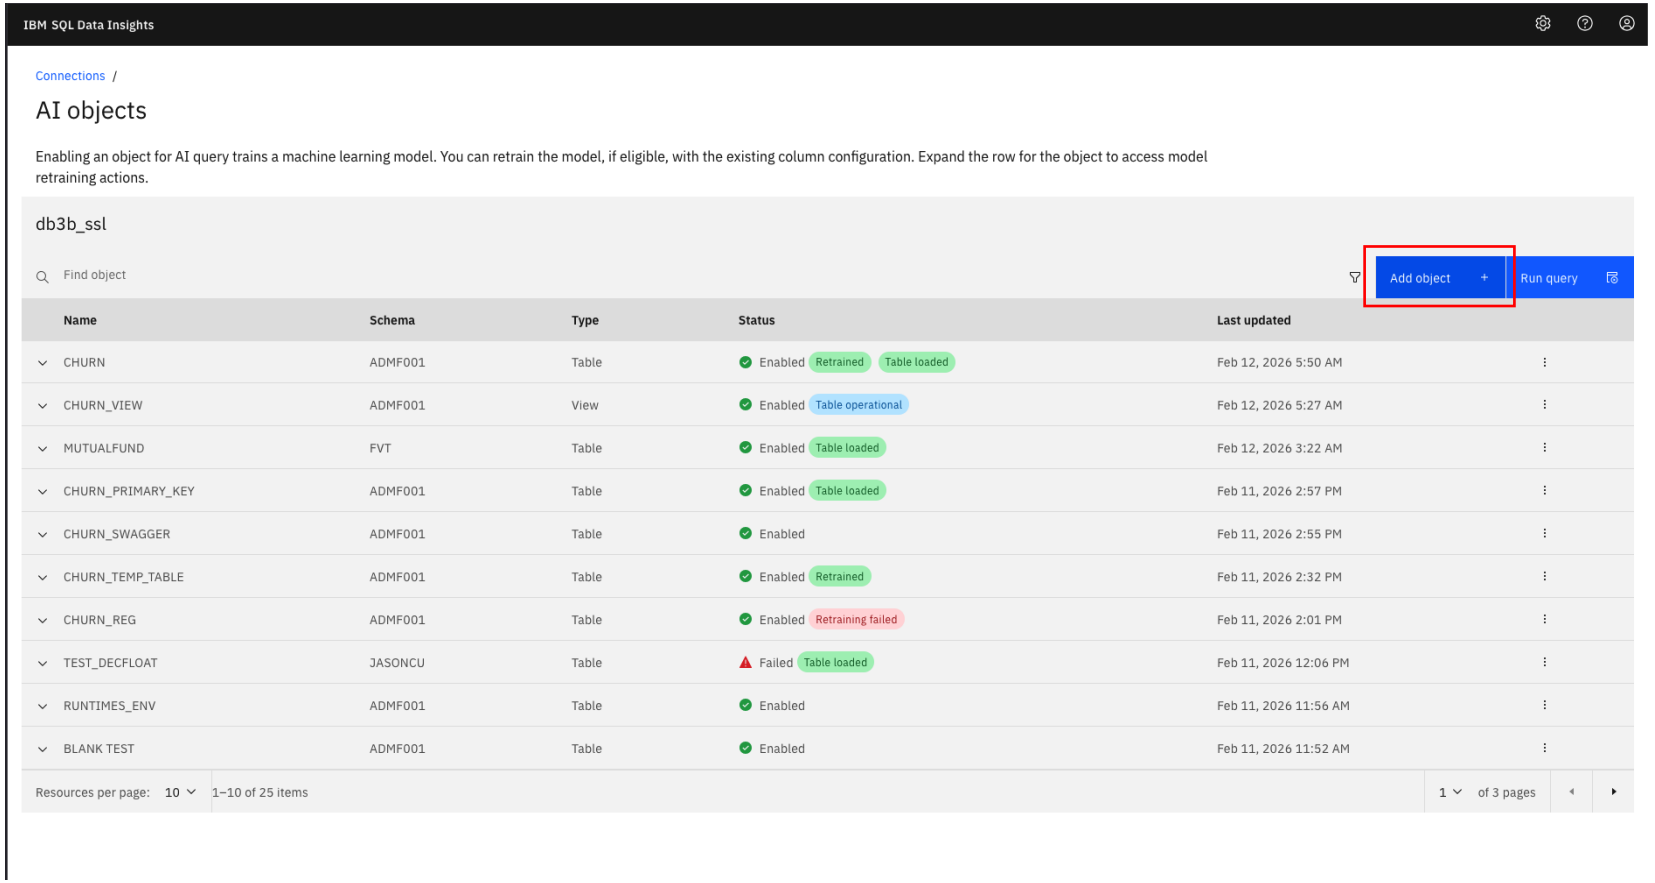
IBM SQL Data Insights
Connections /
AI objects
Enabling an object for AI query trains a machine learning model. You can retrain the model, if eligible, with the existing column configuration. Expand the row for the object to access model retraining actions.
db3b_ssl
Find object
Add object
Run query
Name
Schema
Type
Status
Last updated
CHURN ADMF001 Table Enabled Retrained Table loaded Feb 12, 2026 5:50 AM
CHURN_VIEW ADMF001 View Enabled Table operational Feb 12, 2026 5:27 AM
MUTUALFUND FVT Table Enabled Table loaded Feb 12, 2026 3:22 AM
CHURN_PRIMARY_KEY ADMF001 Table Enabled Table loaded Feb 11, 2026 2:57 PM
CHURN_SWAGGER ADMF001 Table Enabled Feb 11, 2026 2:55 PM
CHURN_TEMP_TABLE ADMF001 Table Enabled Retrained Feb 11, 2026 2:32 PM
CHURN_REG ADMF001 Table Enabled Retraining failed Feb 11, 2026 2:01 PM
TEST_DECFLOAT JASONCU Table Failed Table loaded Feb 11, 2026 12:06 PM
RUNTIMES_ENV ADMF001 Table Enabled Feb 11, 2026 11:56 AM
BLANK TEST ADMF001 Table Enabled Feb 11, 2026 11:52 AM
Resources per page: 10
1–10 of 25 items
1 of 3 pages

Enable AI (3|9)

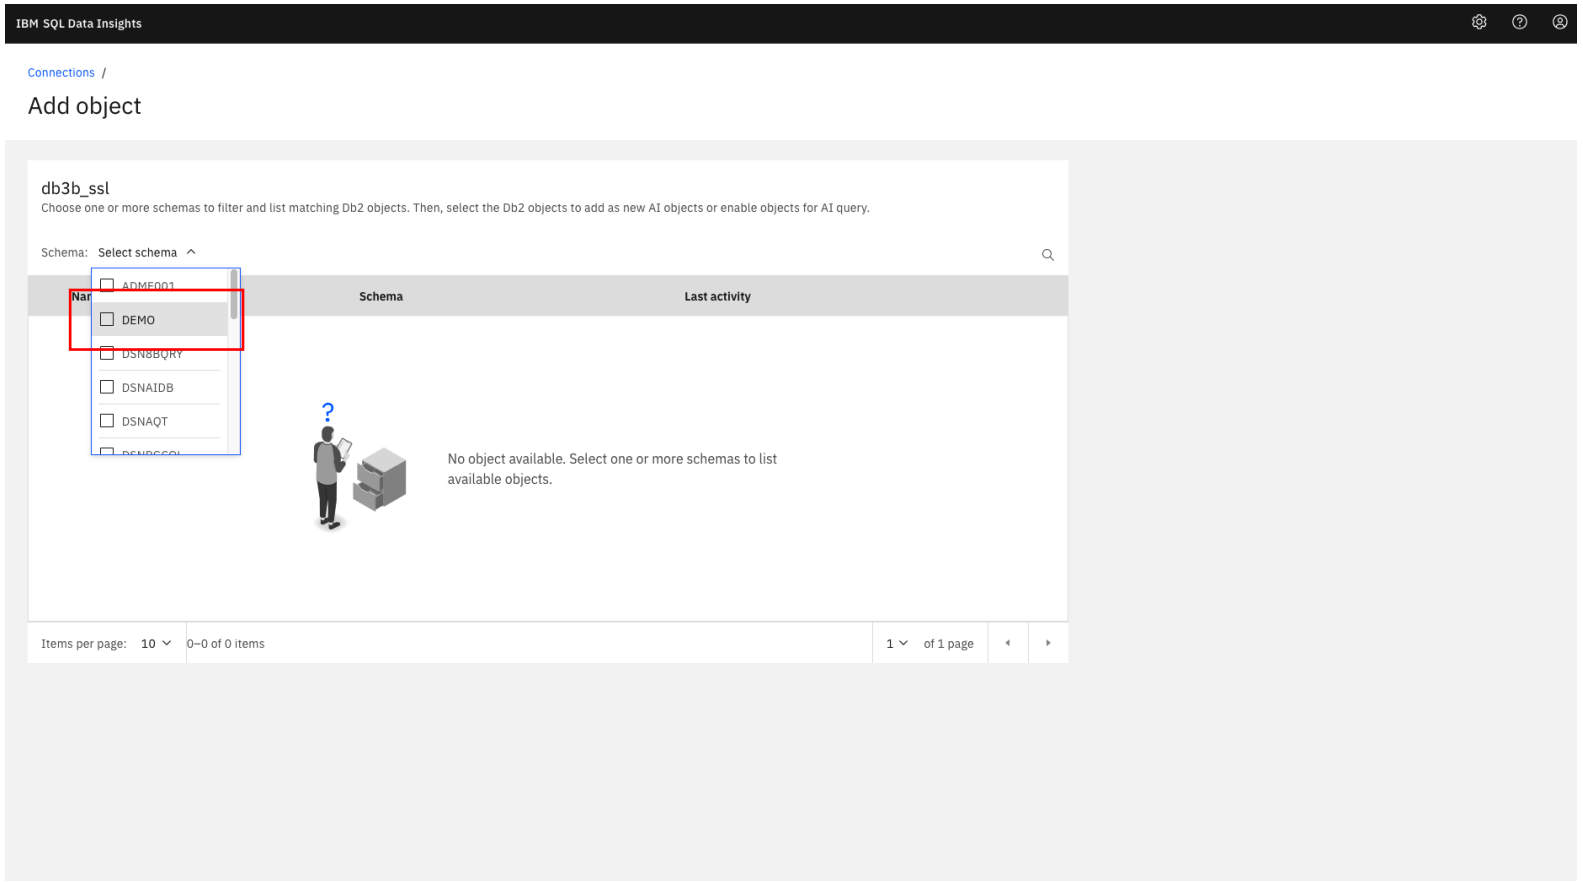
IBM SQL Data Insights
Connections /
Add object
db3b_ssl
Choose one or more schemas to filter and list matching Db2 objects. Then, select the Db2 objects to add as new AI objects or enable objects for AI query.
Schema: Select schema
ADMF001
DEMO
DSN8BQRY
DSNAIDB
DSNAQT
Schema
Last activity
No object available. Select one or more schemas to list available objects.
Items per page: 10
0–0 of 0 items
1 of 1 page

Enable AI (4|9)

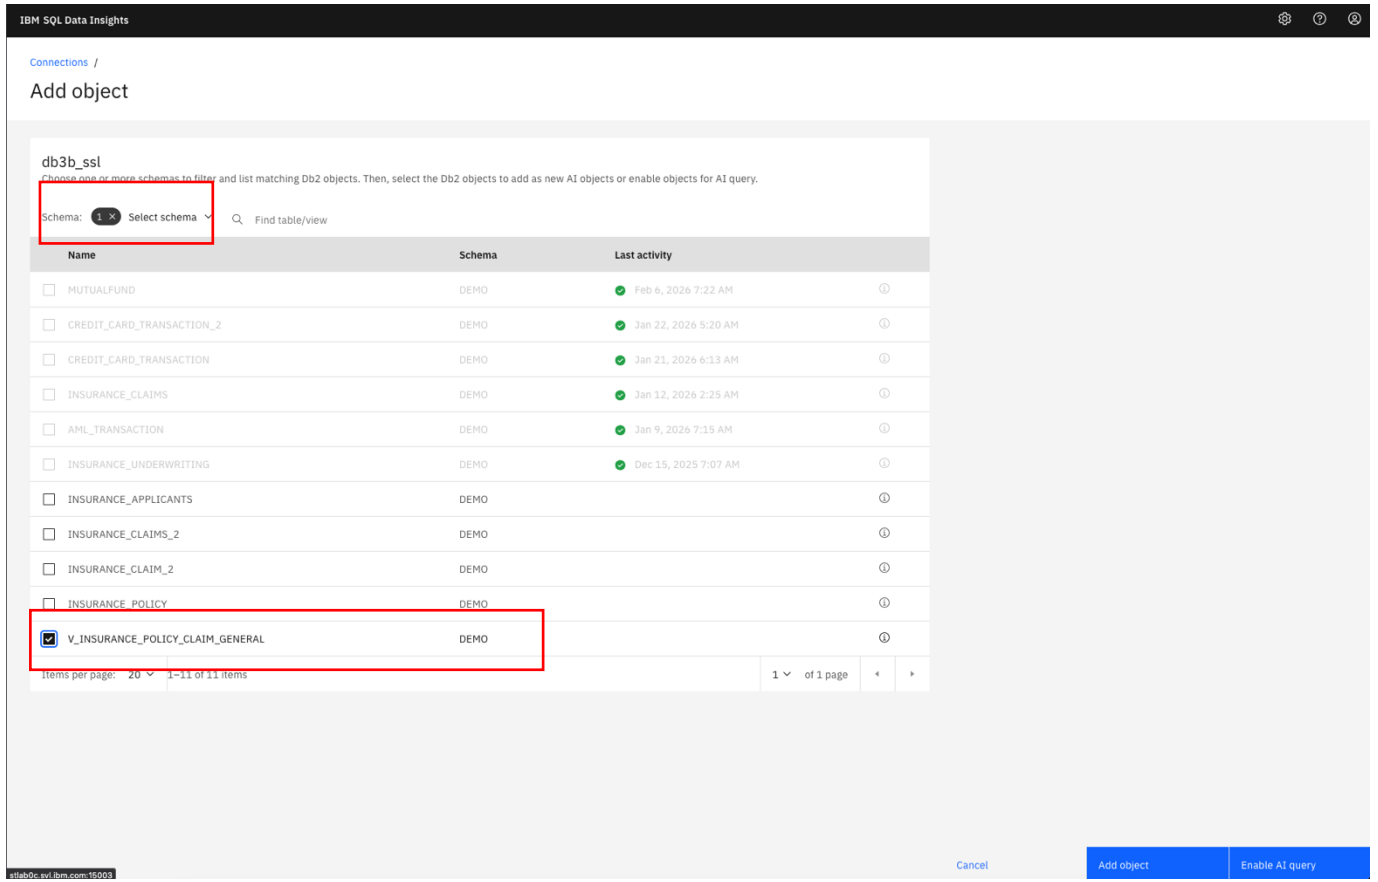
IBM SQL Data Insights
Connections /
Add object
db3b_ssl
Choose one or more schemas to filter and list matching Db2 objects. Then, select the Db2 objects to add as new AI objects or enable objects for AI query.
Schema: 1 Select schema
Find table/view
Name
Schema
Last activity
MUTUALFUND DEMO Feb 6, 2026 7:22 AM
CREDIT_CARD_TRANSACTION_2 DEMO Jan 22, 2026 5:20 AM
CREDIT_CARD_TRANSACTION DEMO Jan 21, 2026 6:13 AM
INSURANCE_CLAIMS DEMO Jan 12, 2026 2:25 AM
AML_TRANSACTION DEMO Jan 9, 2026 7:15 AM
INSURANCE_UNDERWRITING DEMO Dec 15, 2025 7:07 AM
INSURANCE_APPLICANTS DEMO
INSURANCE_CLAIMS_2 DEMO
INSURANCE_CLAIM_2 DEMO
INSURANCE_POLICY DEMO
V_INSURANCE_POLICY_CLAIM_GENERAL DEMO
Items per page: 20
1–11 of 11 items
1 of 1 page
Cancel
Add object
Enable AI query

Enable AI (5|9)

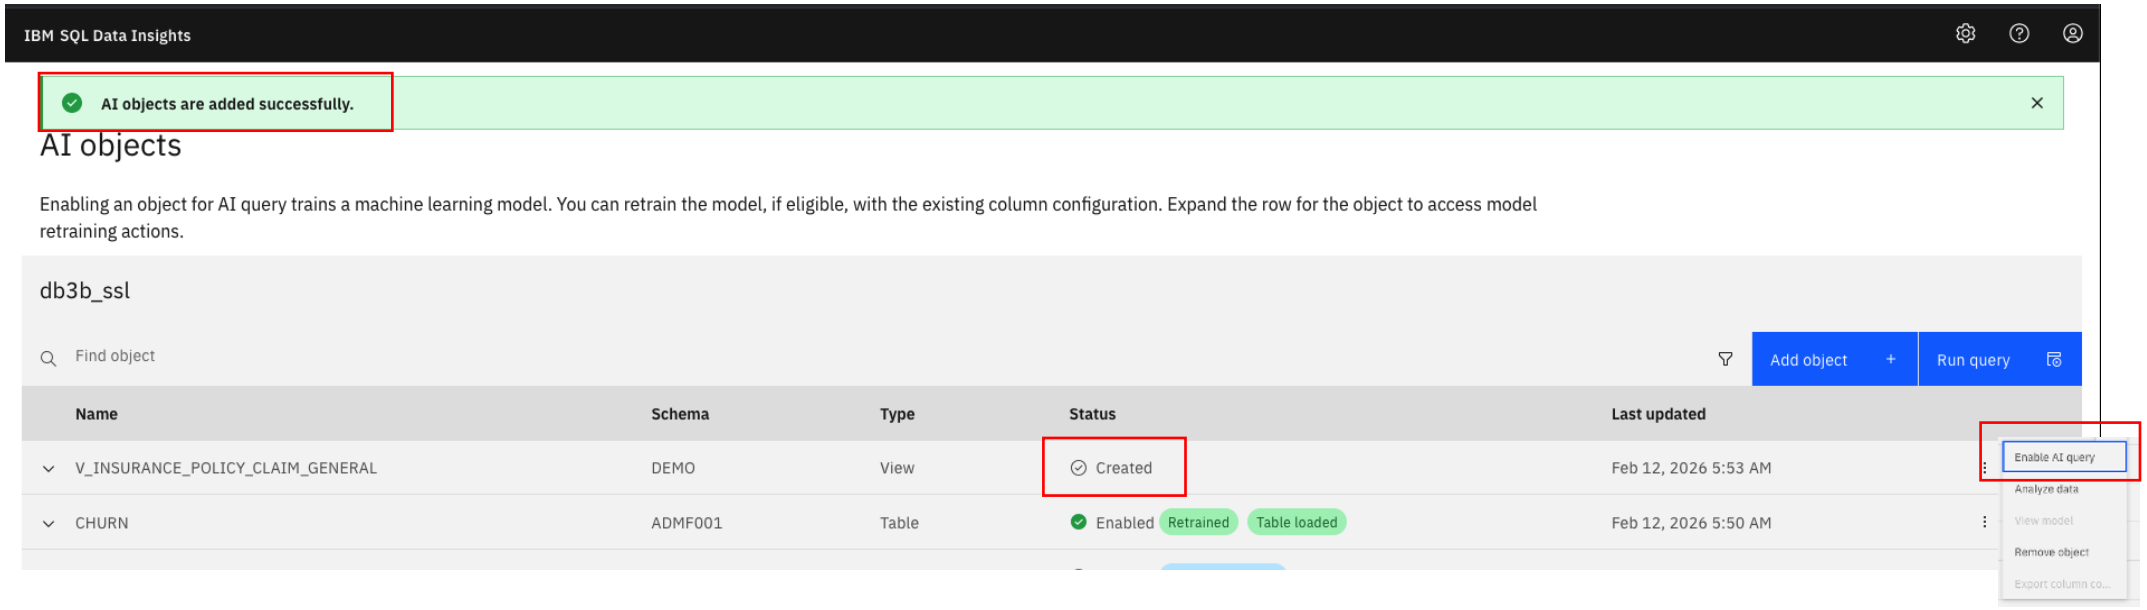
IBM SQL Data Insights
AI objects are added successfully.
AI objects
Enabling an object for AI query trains a machine learning model. You can retrain the model, if eligible, with the existing column configuration. Expand the row for the object to access model retraining actions.
db3b_ssl
Find object
Add object
Run query
Name
Schema
Type
Status
Last updated
V_INSURANCE_POLICY_CLAIM_GENERAL
DEMO
View
Created
Feb 12, 2026 5:53 AM
CHURN
ADMF001
Table
Enabled
Retrained
Table loaded
Feb 12, 2026 5:50 AM
Enable AI query
Analyze data
View model
Remove object
Export column co...

Enable AI (6|9)

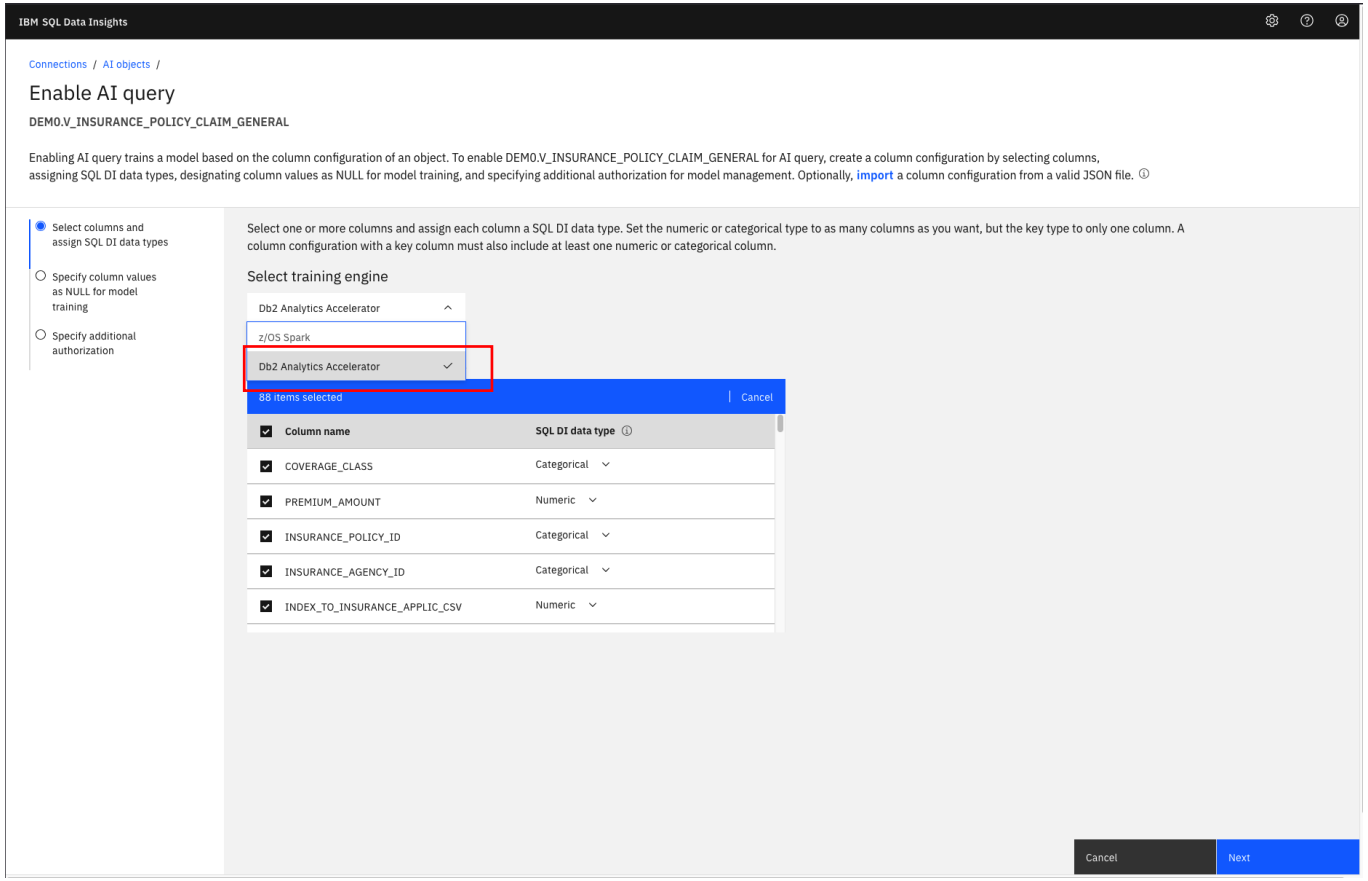
IBM SQL Data Insights
Connections / AI objects /
Enable AI query
DEM0.V_INSURANCE_POLICY_CLAIM_GENERAL
Enabling AI query trains a model based on the column configuration of an object. To enable DEM0.V_INSURANCE_POLICY_CLAIM_GENERAL for AI query, create a column configuration by selecting columns, assigning SQL DI data types, designating column values as NULL for model training, and specifying additional authorization for model management. Optionally, import a column configuration from a valid JSON file.
Select columns and assign SQL DI data types
Specify column values as NULL for model training
Specify additional authorization
Select one or more columns and assign each column a SQL DI data type. Set the numeric or categorical type to as many columns as you want, but the key type to only one column. A column configuration with a key column must also include at least one numeric or categorical column.
Select training engine
Db2 Analytics Accelerator
z/OS Spark
Db2 Analytics Accelerator
88 items selected
Cancel
Column name
SQL DI data type
COVERAGE_CLASS
Categorical
PREMIUM_AMOUNT
Numeric
INSURANCE_POLICY_ID
Categorical
INSURANCE_AGENCY_ID
Categorical
INDEX_TO_INSURANCE_APPLIC_CSV
Numeric
Cancel
Next

Enable AI (7|9)

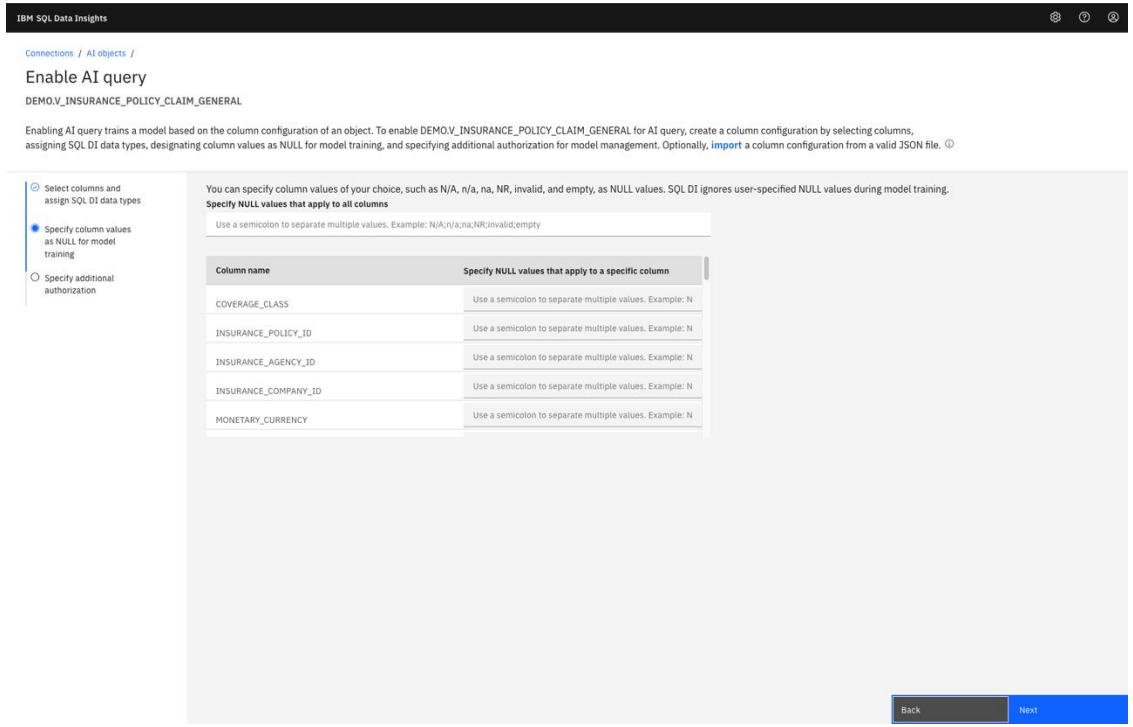
IBM SQL Data Insights
Connections / AI objects /
Enable AI query
DEMO.V_INSURANCE_POLICY_CLAIM_GENERAL
Enabling AI query trains a model based on the column configuration of an object. To enable DEMO.V_INSURANCE_POLICY_CLAIM_GENERAL for AI query, create a column configuration by selecting columns, assigning SQL DI data types, designating column values as NULL for model training, and specifying additional authorization for model management. Optionally, import a column configuration from a valid JSON file.
Select columns and assign SQL DI data types
Specify column values as NULL for model training
Specify additional authorization
You can specify column values of your choice, such as N/A, n/a, na, NR, invalid, and empty, as NULL values. SQL DI ignores user-specified NULL values during model training.
Specify NULL values that apply to all columns
Use a semicolon to separate multiple values. Example: N/A;n/a;na;NR;invalid;empty
Column name
Specify NULL values that apply to a specific column
COVERAGE_CLASS
INSURANCE_POLICY_ID
INSURANCE_AGENCY_ID
INSURANCE_COMPANY_ID
MONETARY_CURRENCY
Use a semicolon to separate multiple values. Example: N
Back
Next

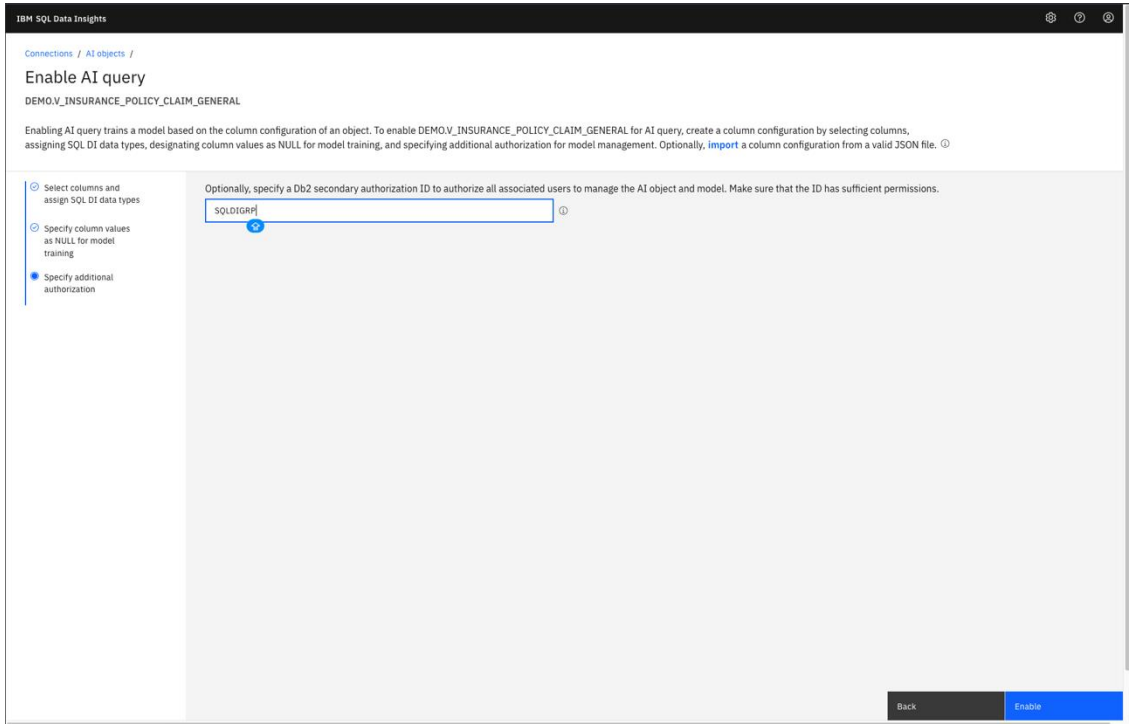
IBM SQL Data Insights
Connections / AI objects /
Enable AI query
DEMO.V_INSURANCE_POLICY_CLAIM_GENERAL
Enabling AI query trains a model based on the column configuration of an object. To enable DEMO.V_INSURANCE_POLICY_CLAIM_GENERAL for AI query, create a column configuration by selecting columns, assigning SQL DI data types, designating column values as NULL for model training, and specifying additional authorization for model management. Optionally, import a column configuration from a valid JSON file.
Select columns and assign SQL DI data types
Specify column values as NULL for model training
Specify additional authorization
Optionally, specify a Db2 secondary authorization ID to authorize all associated users to manage the AI object and model. Make sure that the ID has sufficient permissions.
SQLDIGRP
Back
Enable

Enable AI (8|9)

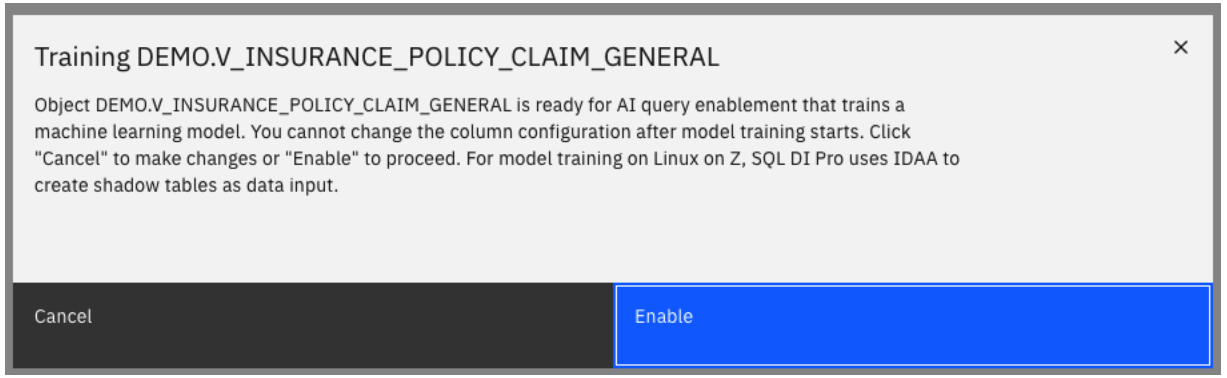
Training DEMO.V_INSURANCE_POLICY_CLAIM_GENERAL
Object DEMO.V_INSURANCE_POLICY_CLAIM_GENERAL is ready for AI query enablement that trains a machine learning model. You cannot change the column configuration after model training starts. Click "Cancel" to make changes or "Enable" to proceed. For model training on Linux on Z, SQL DI Pro uses IDAA to create shadow tables as data input.
Cancel
Enable

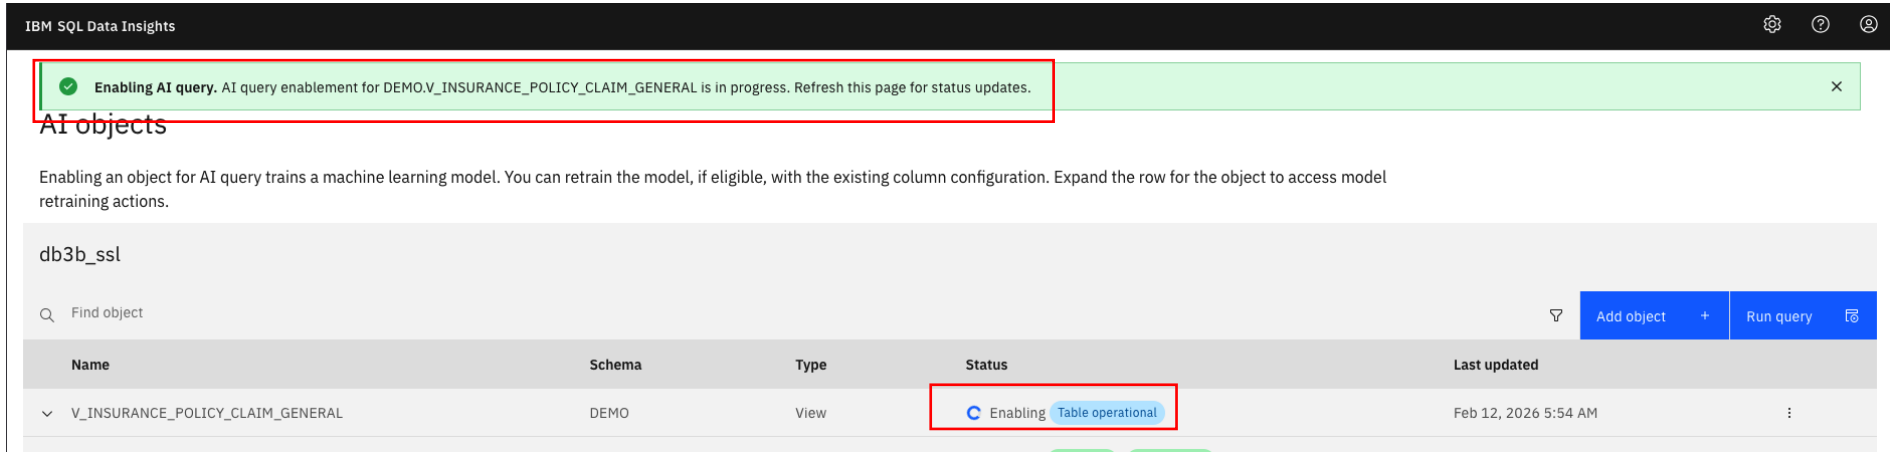
IBM SQL Data Insights
Enabling AI query. AI query enablement for DEMO.V_INSURANCE_POLICY_CLAIM_GENERAL is in progress. Refresh this page for status updates.
AI objects
Enabling an object for AI query trains a machine learning model. You can retrain the model, if eligible, with the existing column configuration. Expand the row for the object to access model retraining actions.
db3b_ssl
Find object
Add object
Run query
Name
Schema
Type
Status
Last updated
V_INSURANCE_POLICY_CLAIM_GENERAL
DEMO
View
Enabling
Table operational
Feb 12, 2026 5:54 AM

Enable AI (9|9)

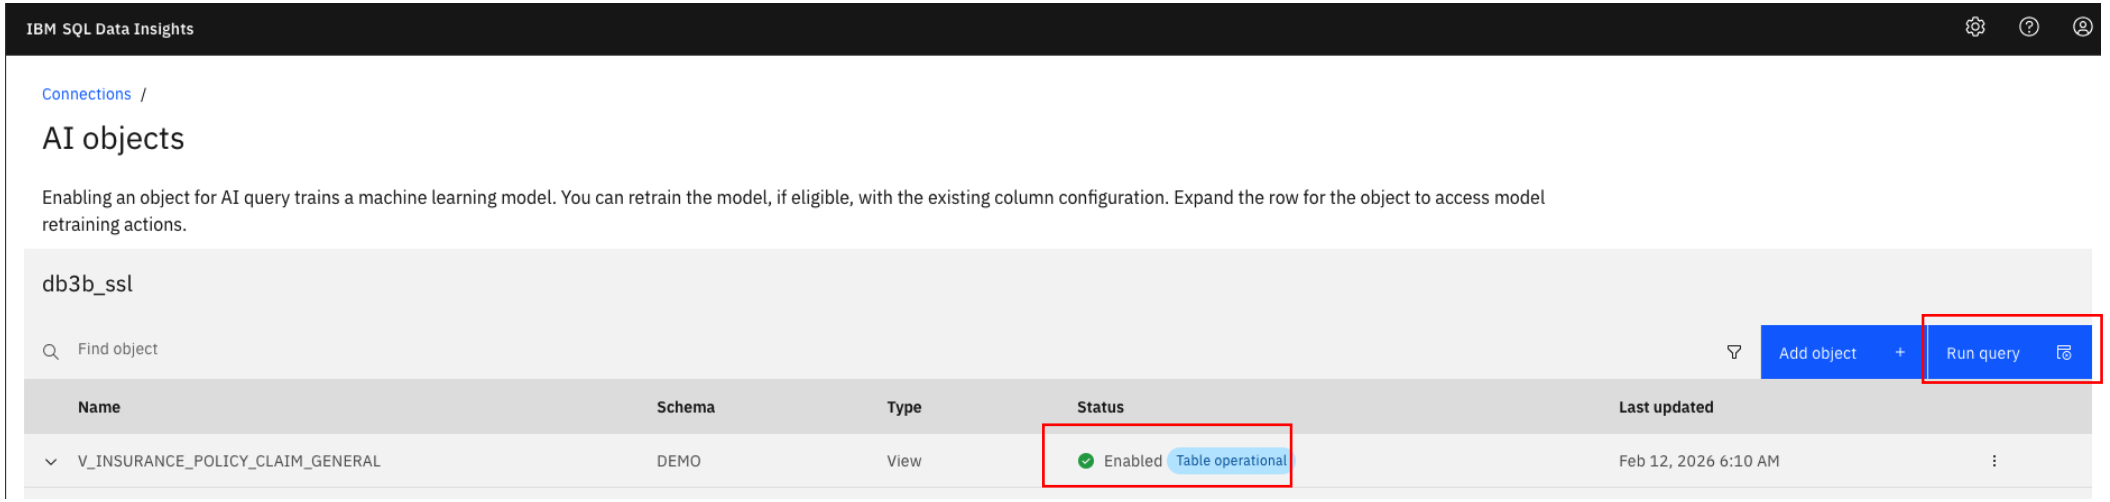
IBM SQL Data Insights
Connections /
AI objects
Enabling an object for AI query trains a machine learning model. You can retrain the model, if eligible, with the existing column configuration. Expand the row for the object to access model retraining actions.
db3b_ssl
Find object
Add object
Run query
Name
Schema
Type
Status
Last updated
V_INSURANCE_POLICY_CLAIM_GENERAL
DEMO
View
Enabled
Table operational
Feb 12, 2026 6:10 AM

Incremental Retraining (1|3)

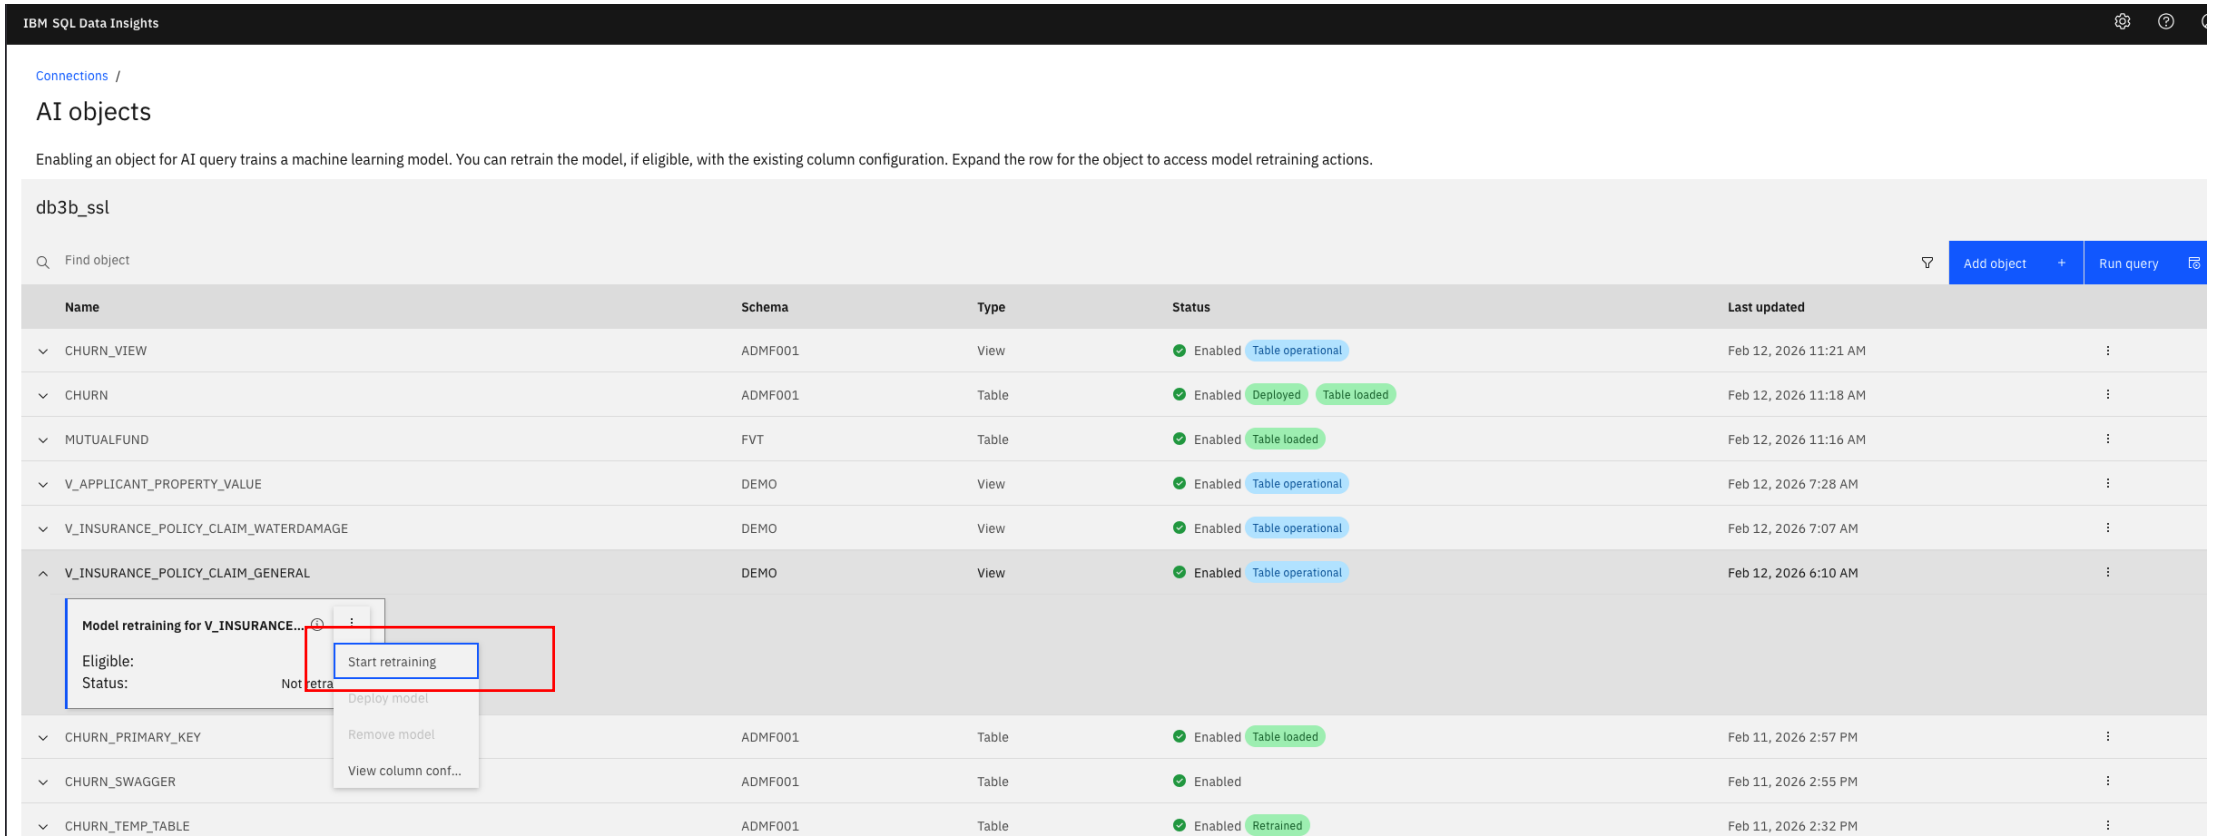

Incremental Retraining (2|3)

IBM SQL Data Insights

Connections / AI objects /

Retrain model

DEMO.V_INSURANCE_POLICY_CLAIM_GENERAL

Preview of data filtered by the specified predicate

DEMO.V_INSURANCE_POLICY_CLAIM_GENERAL

Total rows: 1701 (with 100 rows displayed)

NCY_ID	INSURANCE_COMPANY_ID	HEATING_TYPE_OF_UPDATE	HAIL_DAMAGE_COVERAGE_LIM	START_DATE	YEAR_BUILT	FIRE_DAMAGE
	1412378E0	No Updates	44000	2025-09-03	2024	2000
	141237400	No Updates	144000	2025-09-23	2015	5000
	141237400	No Updates	144000	2025-09-23	2015	5000
	141239140	No Updates	1024000	2025-09-30	2017	5000
	141239FE0	No Updates	98000	2025-09-04	2002	2000
	1412378E0	No Updates	34000	2025-09-22	1996	5000
	141237400	Full	41000	2025-09-03	1965	5000
	141239140	No Updates	30000	2025-09-12	2021	250
	141239140	Full	96000	2025-09-02	1963	5000
	141237400	No Updates	153000	2025-09-04	2024	250
	141239B00	No Updates	116000	2025-09-23	2016	250

Incremental Retraining (3|3)

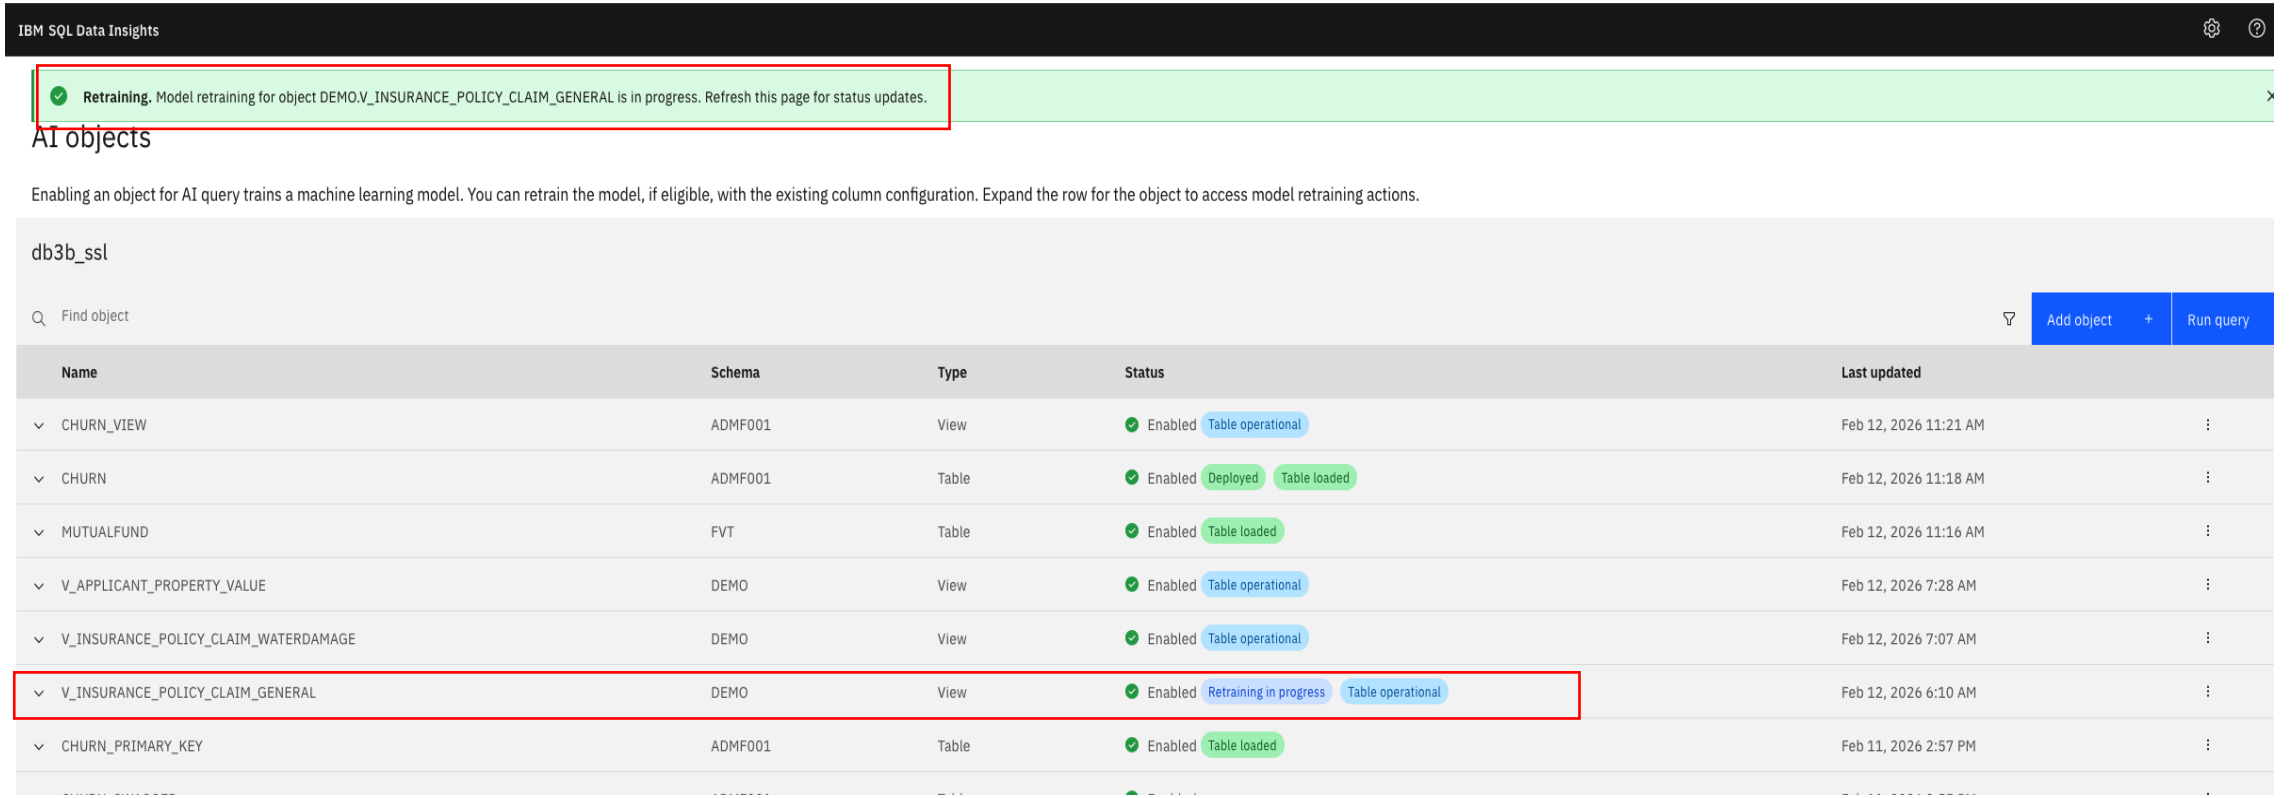

AI Query for Insurance Use Cases (1|2)

- Find claims SIMILAR to a suspicious claim

```
SELECT AI_SIMILARITY (I.CLAIM_ID, 207550) AS SCORE, I.*
FROM INSURANCE_CLAIMS I
ORDER BY SCORE ASC
FETCH FIRST 10 ROWS ONLY;
```

- Identify the LEAST COMMON claims - outliers

```
SELECT AI_COMMONALITY(I.CLAIM_ID) AS SCORE, I.*
FROM INSURANCE_CLAIMS I
ORDER BY SCORE DESC;
```

AI Query for Insurance Use Cases (2|3)

- Find policies belong to a group of policies known underwriting pattern of concern.

```
SELECT AI_SEMANTIC_CLUSTER (POLICY_ID,
                            'A08670943', 'T43877806', 'C30381936') AS SCORE
FROM INSURANCE_POLICY P
WHERE POLICY_ID NOT IN ('A08670943', 'T43877806', 'C30381936')
ORDER BY SCORE DESC;
```

- Identify policies in state of California whose underwriting risk profile is analogous to a known risky customer in New York.

```
SELECT AI_ANALOGY('NEW YORK' USING MODEL COLUMN POLICY_STATE,
                  'Q06-42-2341' USING MODEL COLUMN POLICY_ID,
                  'CALIFORNIA',
                   POLICY_ID) AS SCORE, P.*
FROM INSURANCE_POLICY P
WHERE DRIVER_STATE = 'CALIFORNIA'
ORDER BY SCORE DESC;
```

AI Query for Insurance Use Cases (3|3)

- Looking for a coverage gap
  - Identify claims on water damage that do not have coverage on flood
  - Run either AI_SIMILARITY or AI_SEMANTIC_CLUSTER to find the similar claims

```
SELECT AI_SEMANTIC_CLUSTER (CLAIM_ID,
                                    '207550', '509853', '717312') AS SCORE
FROM V_INSURANCE_POLICY_APPLICANT_CLAIM
WHERE POLICY_ID NOT IN ('A08670943', 'T43877806', 'C30381936')
ORDER BY SCORE DESC;
```

Potential Use Cases (1|2)

- Finance (Consumer Banking, Investment Advisors)
  - Find customer with similar transactions
  - Wealth Management
  - Identify stocks with similar behavior as determined by their numerical attributes (structured data)
  - Mutual Fund analysis incorporated with strategy description, manager's commentary that suites your goal (unstructured data)
- Fraud Detection
  - Anti money laundering
  - Credit card fraud detection

Potential Use Cases (2|2)

- Insurance
  - Claim analysis – Identify similar/dissimilar claims
  - Underwriting – Risk assessment for the new policy
  - Upselling – Identify opportunities for revenue generation
- IT Analytics
  - Identify tickets with similar behavior/patterns
  - Identify customers with similar ticket patterns
- Customer Analytics
  - Find similar customers based on buying patterns
  - Customer Churn Analytics

Industry Packs

- Demonstrate how structured and unstructured data can be unified under Db2 with AI-drive insights.
- Each Industry Pack will include:
  - Sample table DDL with data
  - Prebuilt model
  - Ready-to-run Jupyter notebooks
  - Semantic SQL examples using AI functions
- Industries from
  - Finance (Credit Card Transactions, Bank Transactions)
  - Insurance (Claims, Underwriting, Up-Selling)
  - Cross-Industry (Churn Prediction, Entity Resolution)
  - And more are planned to be included

Industry Pack

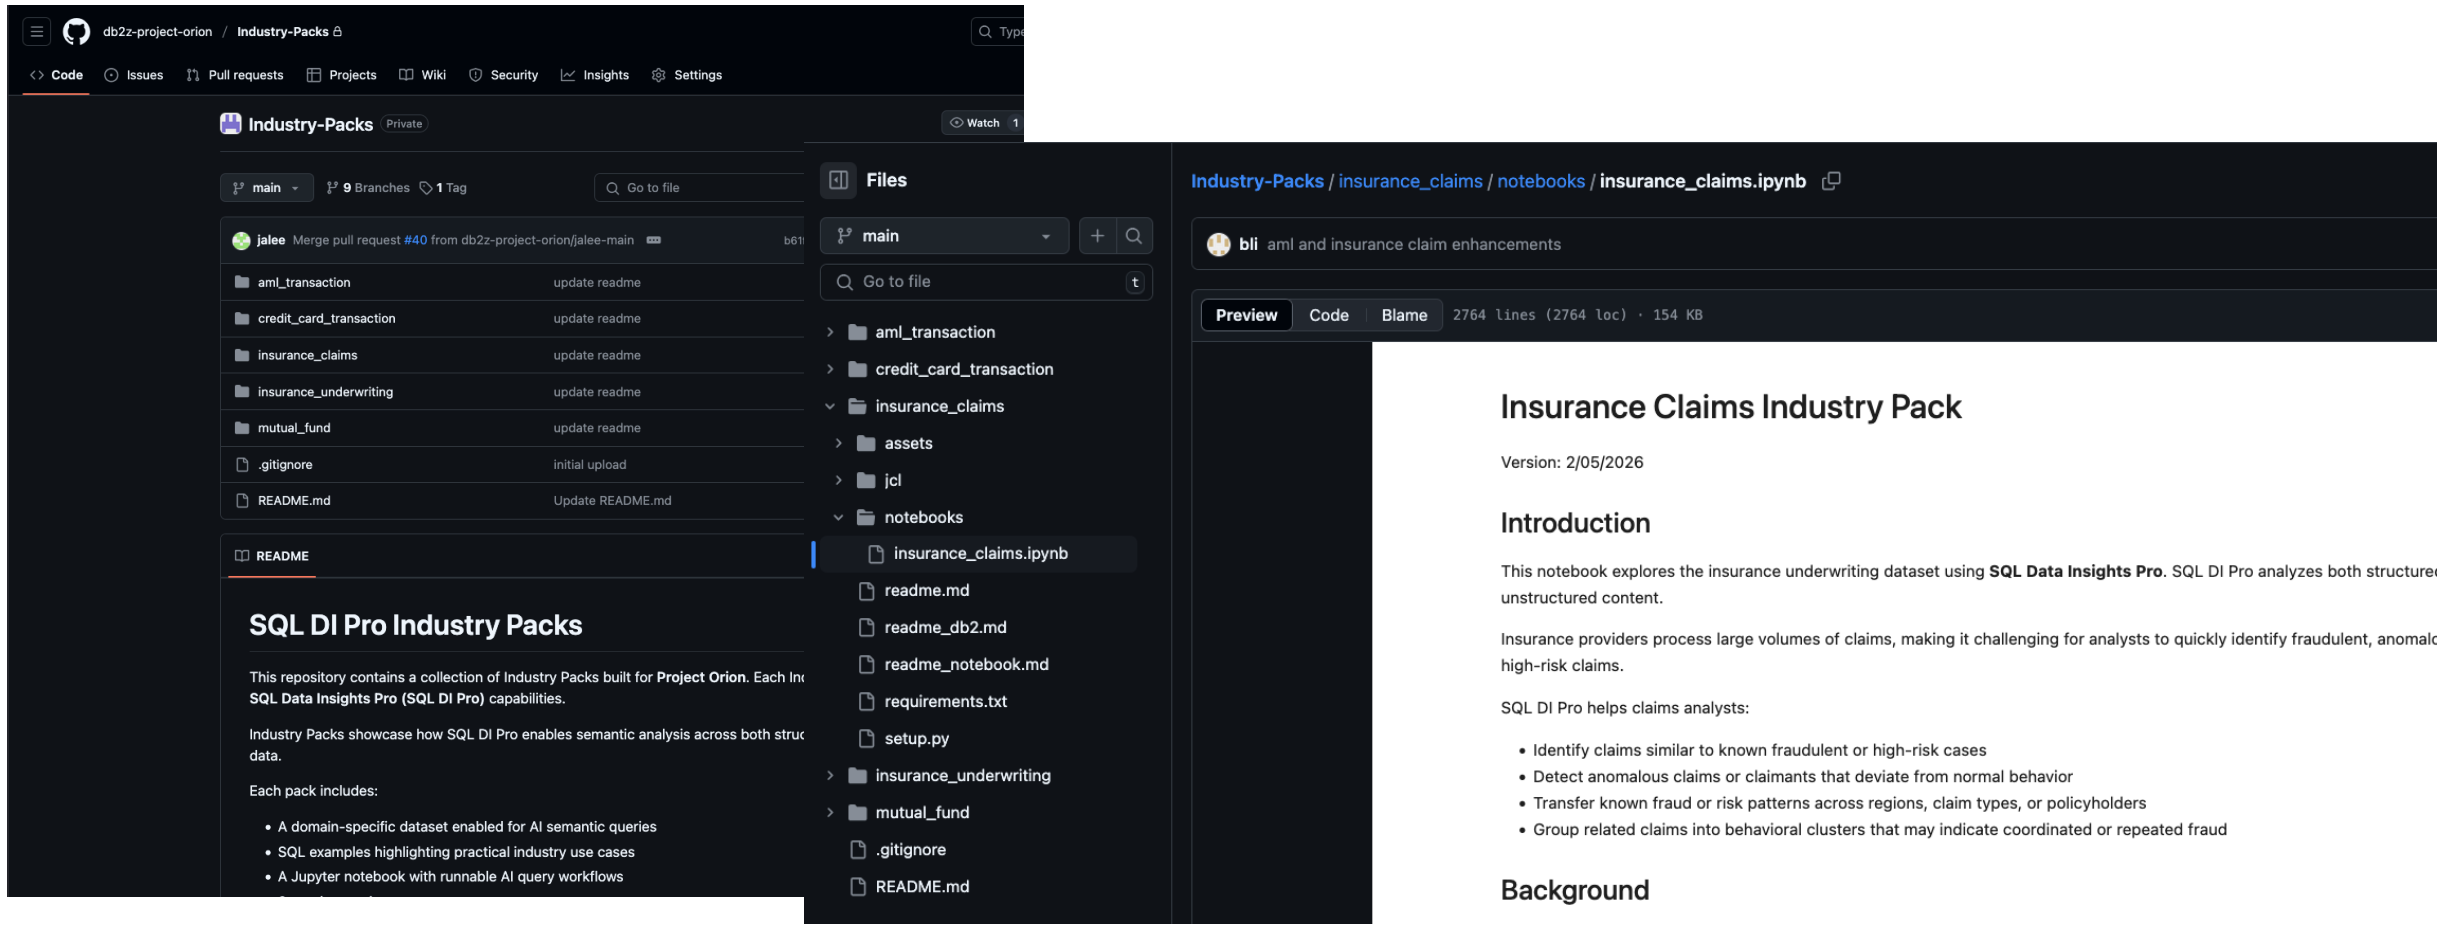

Summary: SQL DI in One-Pager

The problem

Teams have valuable structured and unstructured business data in their Db2 for z/OS tables. To get deeper insight, they often copy data to other systems or do manual analysis. That adds time, cost, and compliance risk.

Main Features

- ✓ Search by meaning, not just exact matches
- ✓ Find records similar to a case, claim, transaction, or customer profile
- ✓ Detect outliers and unusual behavior earlier in the data lifecycle
- ✓ Analyze structured data together with text, including notes, descriptions, and logs
- ✓ Incremental Retraining to include new data while saving resources

The solution

Instantly uncover patterns across your Db2 for z/OS data with advanced AI that is optimized and easy to use. SQL Data Insights discovers relationships and similarities that traditional means can't – so that teams can quickly find "cases like this", spot unusual behavior early, and make faster decisions without moving data off platform.

Persona

- ✓ Business analysts
- ✓ Application Developer
- ✓ Db2 architects / DBAs
- ✓ IT leaders

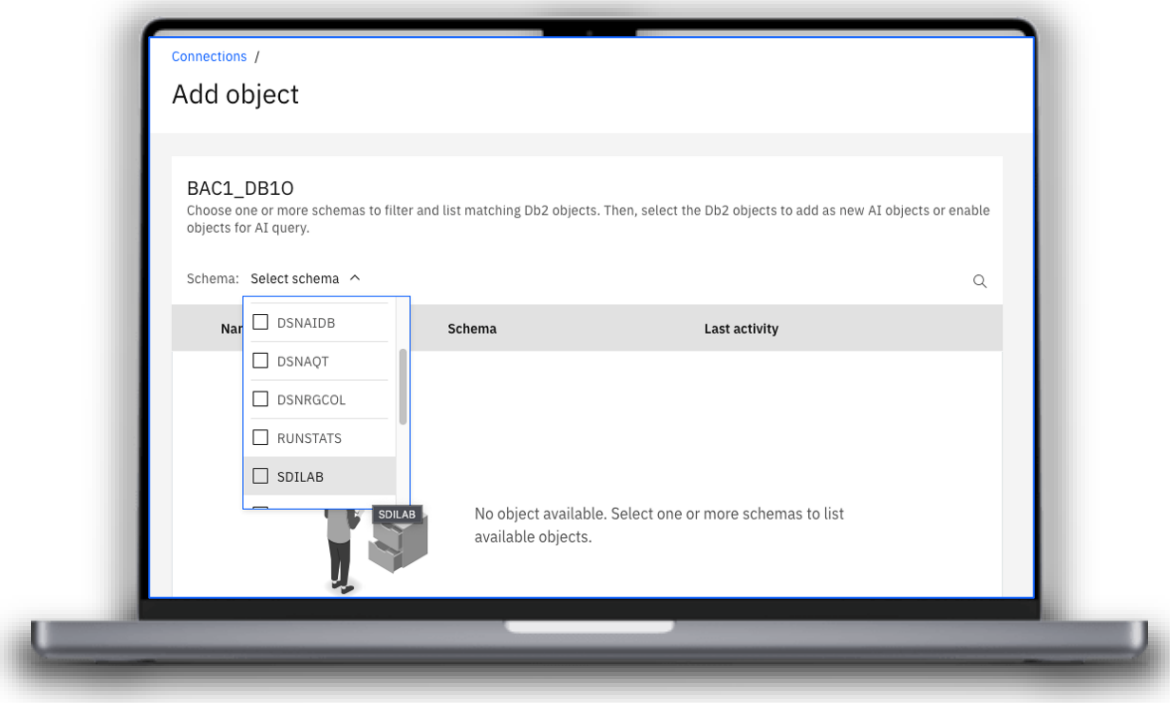

Hands-on Lab on SQL DI

C14A & C14B

Hands-on Lab: Performing AI queries with Db2 for z/OS using SQL Data Insight

Time: Wednesday, March 18 2pm – 5:30pm

Location: Cross Platform/App Dev Room

Please bring you own laptop.

IDUG

2026

Sydney | March 16 - 18

AU Db2 TECH CONFERENCE

Session Title

Jae Lee, *IBM*

Contact: jalee@us.ibm.com

Session Code: A05

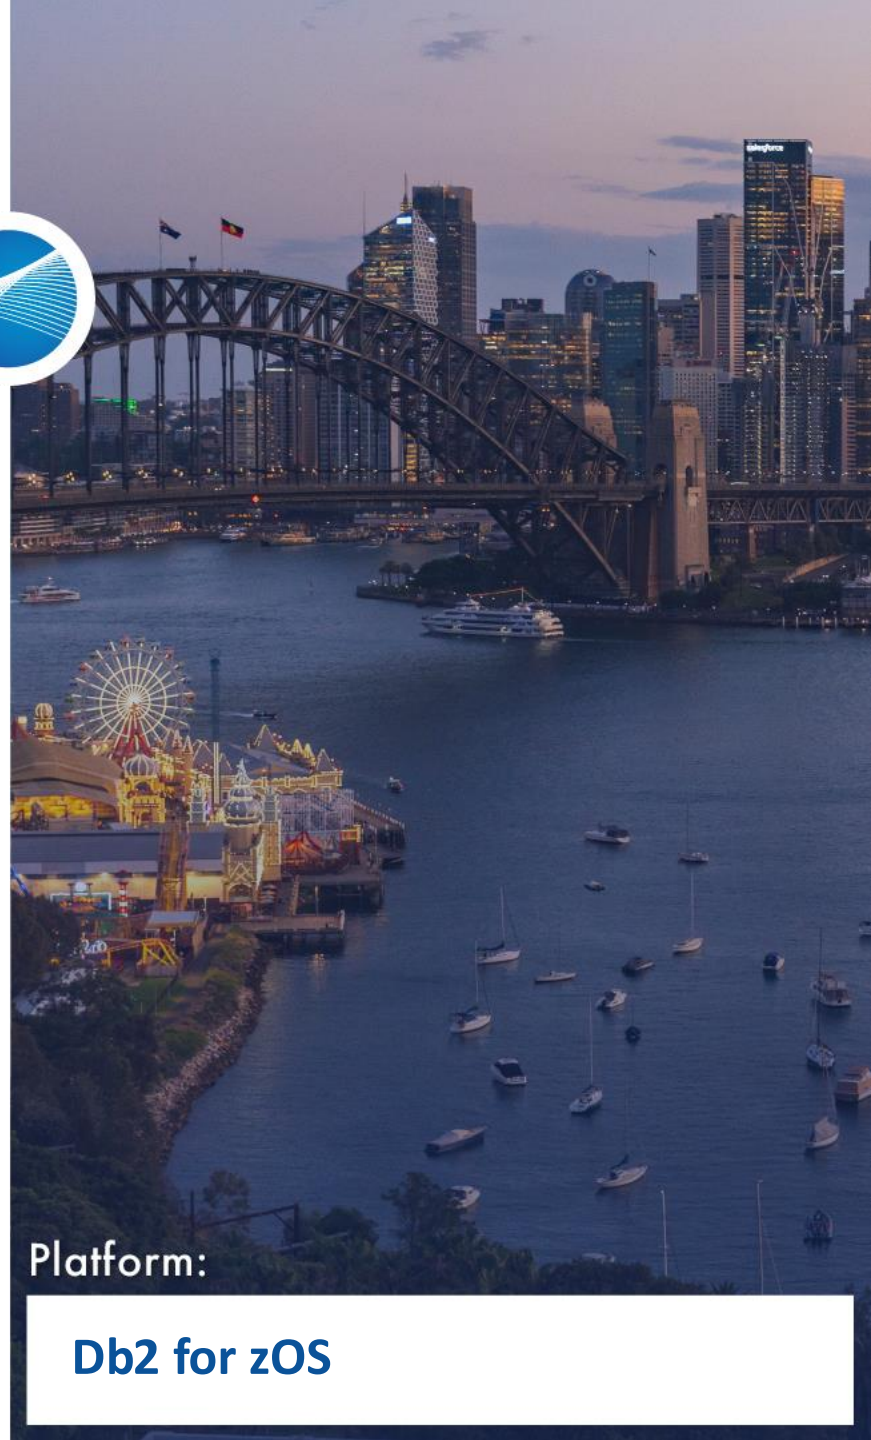

Platform:

Db2 for zOS


IDUG
2026 Australia Db2 Tech Conference