



Unlock the Power of Data:
Your Journey with Generative AI

**INTRODUCING
YOUR AUTODATA COPILOT**





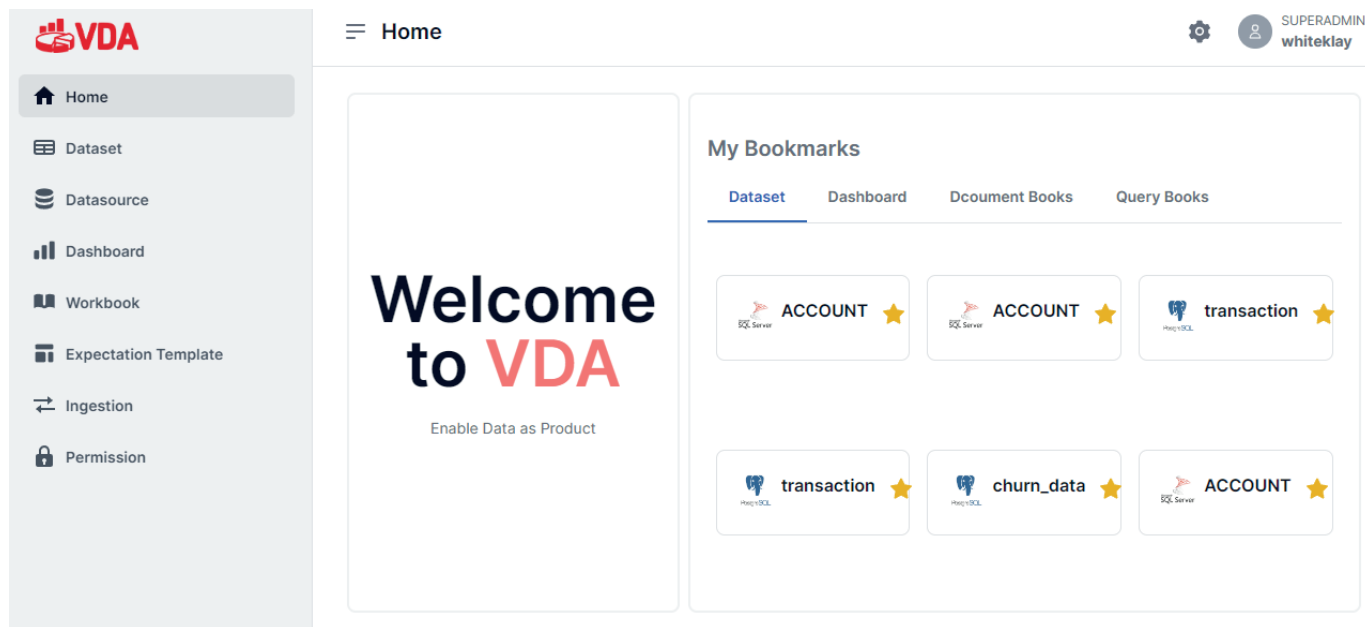
DISCOVER THE FUTURE OF DATA

Still Struggling with Data Engineering?

Have you ever wondered what stories your data could tell? What if you could explore it like a strategic roadmap, seeking answers to questions that drive your business forward?

In today's world, data is everywhere, and making sense of it can be a challenge. What if you'd an engineering assistant?

VDA is your virtual data assistant. It's here to help you make sense of your information.



Are you ready to transform the way you look at data? Step into a world where questions are the keys to insights, and data is your trusted companion.

DATASHEET VDA



VDA: YOUR AUTODATA COPILOT

Imagine if you could effortlessly understand your data, ask it questions, visualise and uncover valuable insights. Enter VDA, your Virtual Data Assistant— a game-changer designed to transform the way we work with data. It's akin to having a knowledgeable companion by your side, simplifying the process and enabling better decisions.

Your Data, Your Questions

Instead of shifting through data on your own, VDA lets you ask questions:

What are the spending habits of our top banking customers?

What's the average balance in our customers' bank accounts?

What patterns emerge from our customers' recent transactions?

VDA brings you closer to natural language interaction with your data.

Conversations with Data

With the power of natural language, VDA helps you interact with your datasets and derive insights out of it in a matter of seconds. Learn about customer behaviour, market trends, opportunities with few clicks.

The Power of Insights

With VDA, you'll discover valuable insights that will shape your decisions. It's all about making informed choices based on data.

The screenshot shows the VDA interface with a sidebar on the left containing navigation items: Home, Dataset, Datasource, Dashboard, Workbook, Expectation Template, Ingestion, and Permission. The main area is titled "Add KPI to 1_Customer 360" and features a "Statement" input field with the text "top 10 customers with high balance". Below the input, a note reads "KPIs can be anything like - Top 10 customers, Gender wise saving amounts". The interface displays the generated SQL query and a table of results.

```
SQL
SELECT
customers.customer_id as
customer_id,
customers.first_name as
first_name,
customers.last_name as
last_name, customers.age as
age, customers.gender as
gender, accounts.balance as
balance FROM customers
```

customer_id	first_name	last_name	age	gender	balance
329	Jill	Phillips	68	Male	99951
6	Crystal	Sanchez	52	Male	99915
455	Alan	Hodge	61	Female	99829
393	Anna	Smith	67	Female	99732
408	Wayne	Berry	19	Male	99606
273	Bridget	Smith	30	Male	99572
231	Kimberly	Wright	40	Female	99555
40	Michelle	Cooper	48	Male	99510
272	Hector	Carpenter	36	Female	99498
349	Angela	Prince	23	Male	99250

CREATING A DATA STORY



Source



Build



Analyze



Visualize

Source

As part of VDA, the Virtual Data System, our goal is to provide you with a suite of automated tools. These tools efficiently gather information from diverse sources, structured or unstructured. Using AI and ML techniques, we create a robust data catalog and employ generative AI to generate SQL statements for various use cases. Additionally, our platform offers powerful search capabilities, accommodating both specific and broad search requirements across your datasets.

Build

The VDA's auto data catalog feature offers a method for building data products through the use of logical data models, allowing you to define your data models. Further, it provides you a way where you can ingest these data models or you can ingest any kind of datasets using our build ingestion framework.

So it is metadata on top of metadata which can help you to do data classification where you can provide certain kind of information. It will help you to do data governance, understanding how this data is getting populated.

Analyze

The initial step in our analysis process involves conducting a data search, data review, and data preview. Following this, the subsequent phase caters to a business analyst's specific work requirements, often involving fieldwork.

As part of our Virtual Data System, we've introduced a series of workbooks designed to streamline various tasks for our business analysts. These workbooks serve as comprehensive toolsets, providing essential features that simplify tasks that might otherwise be time-consuming or require comprehensive documentation.

Visualize

After sourcing and ingesting the data, developing a logical data model, and conducting various analyses, the next crucial step is data visualization.

In this regard, we have introduced an 'Auto Dashboarding' feature. Using natural language, users can effortlessly build dashboards. This feature simplifies the process of creating informative dashboards quickly and efficiently. This version maintains a formal tone and focuses on the key information, eliminating grammatical errors and unnecessary conversational elements.



KEY FEATURES & BENEFITS

1

Data Cataloging:

Connectors to diverse data sources allow the VDA to perform efficient data cataloging. It automatically extracts metadata from these sources, providing analysts with a comprehensive view of available data assets.

The ability to understand data context enables analysts to make informed decisions and facilitates effective collaboration across the organization.

2

Assistive Data Modeling:

VDA (Visual Data Assister) is a powerful tool that simplifies data modeling by offering an intuitive UI window to effortlessly scan multiple source data and tables. Collaborate seamlessly with your team to create a tailored target model that aligns with your specific use case. With natural language capabilities, you can define transformation logic, ensuring precision in your model. VDA then takes care of the heavy lifting by automatically generating SQL code to implement your logical model, making data modeling a breeze.

3

NLP-Powered Dashboard:

VDA's Conversational Dashboard feature empowers users to interact with their data using natural language. Simply input KPIs and queries in plain language, and watch as VDA generates informative textual reports or dynamic dashboards in response. It's a seamless and intuitive way to transform data into actionable insights.

4

Ingestion:

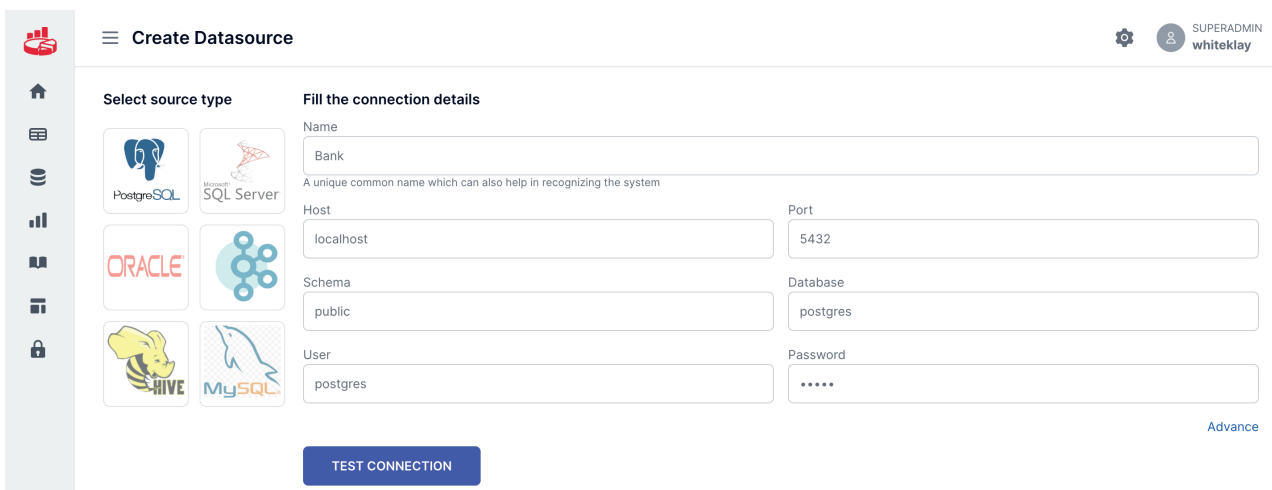
VDA provides an easy way for the business analyst to create ingestion pipelines of your choices.

The concept of data ingestion is traditionally associated with the domain of data engineering, whereas business analysts often rely on proficient data engineers or alternative toolsets when engaging in data ingestion processes.

DATA SOURCES

A Datasource is an entity within the Virtual Data Assistant that serves as a container for a collection of metadata. Metadata refers to information about datasets, such as data source location, schema, data types, and other relevant properties. Essentially, a Datasource is like a virtual folder that groups related datasets together, making it easier for users to manage and access data efficiently.

DATASOURCE CREATION USING CONNECTORS



The screenshot shows the 'Create Datasource' page. On the left is a navigation sidebar with icons for home, list, database, charts, and settings. The main area is titled 'Create Datasource' and features a user profile for 'SUPERADMIN whiteklay'. It is divided into two sections: 'Select source type' and 'Fill the connection details'.

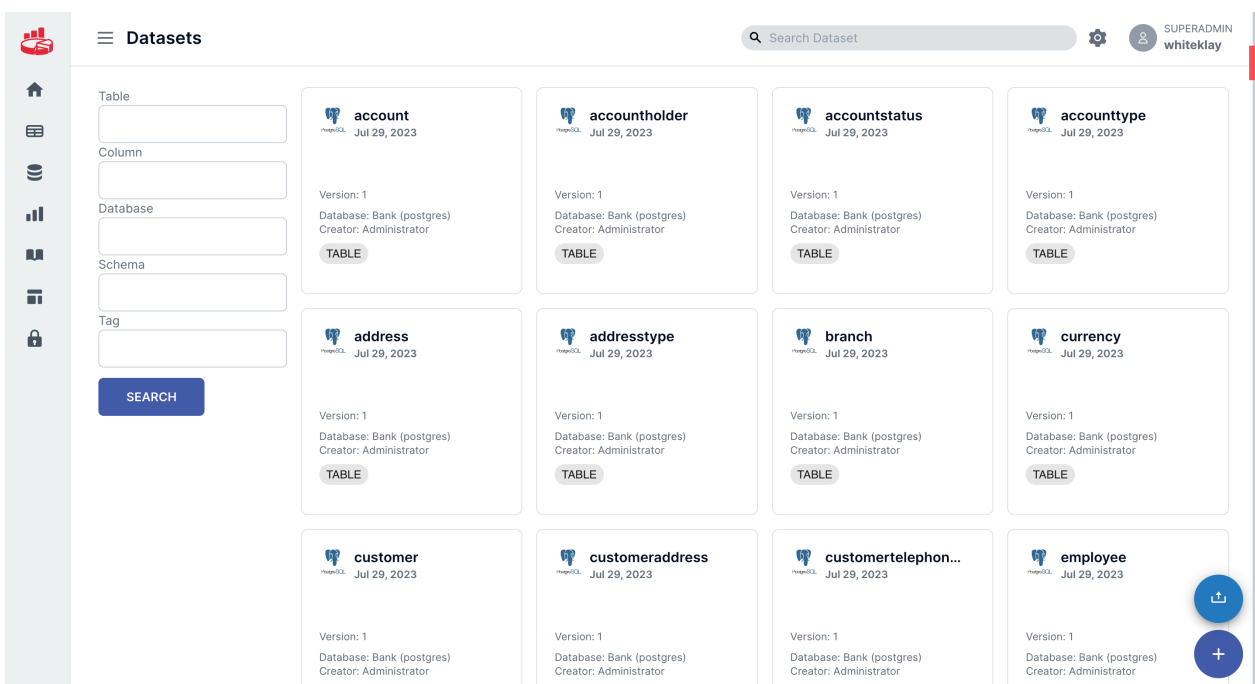
Select source type: Includes icons for PostgreSQL, Microsoft SQL Server, ORACLE, and HIVE. Below these are icons for MySQL and another database connector.

Fill the connection details: A form with the following fields:

- Name:** Bank (with a note: "A unique common name which can also help in recognizing the system")
- Host:** localhost
- Port:** 5432
- Schema:** public
- Database:** postgres
- User:** postgres
- Password:** masked with dots

At the bottom right of the form is an 'Advance' link, and at the bottom center is a blue 'TEST CONNECTION' button.

SEARCHING DATASETS WITHIN A DATASOURCE CREATION USING CONNECTORS



The screenshot shows the 'Datasets' page. It features a search bar at the top right with the text 'Search Dataset'. On the left is a navigation sidebar with icons for home, list, database, charts, and settings. The main area displays a grid of dataset cards.

Search Filters (Left Side):

- Table:
- Column:
- Database:
- Schema:
- Tag:

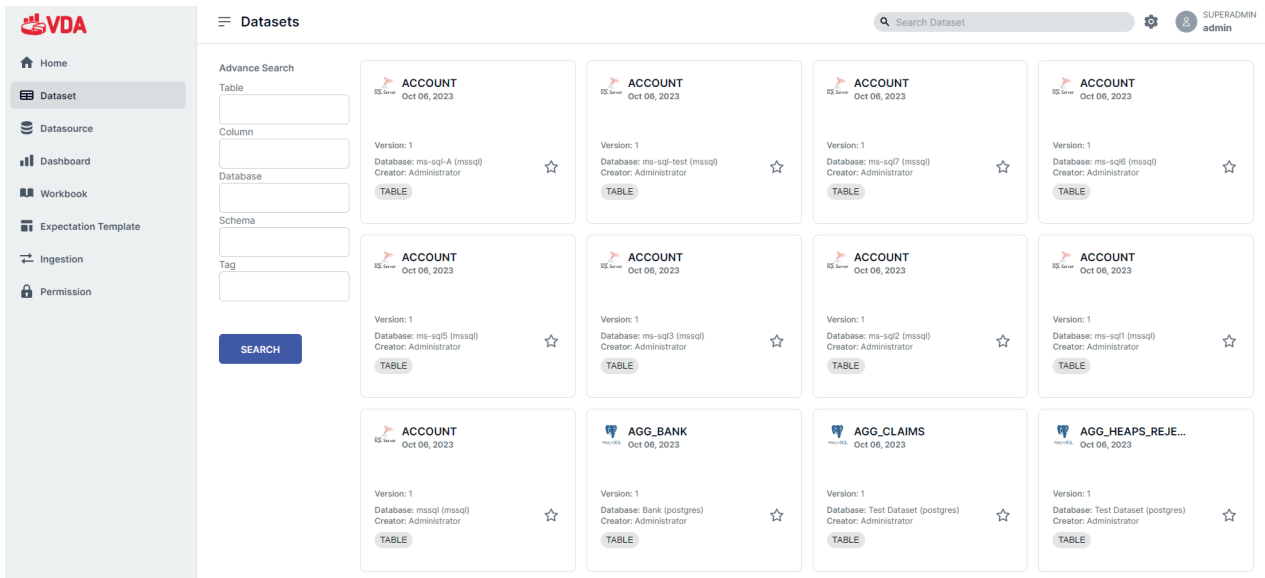
Dataset Cards (Grid): Each card shows a dataset name, version, database, and creator.

- account:** Version: 1, Database: Bank (postgres), Creator: Administrator
- accountholder:** Version: 1, Database: Bank (postgres), Creator: Administrator
- accountstatus:** Version: 1, Database: Bank (postgres), Creator: Administrator
- accounttype:** Version: 1, Database: Bank (postgres), Creator: Administrator
- address:** Version: 1, Database: Bank (postgres), Creator: Administrator
- addressstype:** Version: 1, Database: Bank (postgres), Creator: Administrator
- branch:** Version: 1, Database: Bank (postgres), Creator: Administrator
- currency:** Version: 1, Database: Bank (postgres), Creator: Administrator
- customer:** Version: 1, Database: Bank (postgres), Creator: Administrator
- customeraddress:** Version: 1, Database: Bank (postgres), Creator: Administrator
- customertelephon...:** Version: 1, Database: Bank (postgres), Creator: Administrator
- employee:** Version: 1, Database: Bank (postgres), Creator: Administrator

At the bottom right of the grid are two circular buttons: a blue one with a download icon and a blue one with a plus sign.

DATASETS

Datasets are the lifeblood of data analysis, and Virtual Data Assistant (VDA) empowers users with a feature-rich environment to harness the full potential of data. Datasets within the VDA serve as powerful entities that showcase data as a product, providing valuable insights and enabling data-driven decision-making. Let's delve into the key features of Datasets:



1. Metadata, Overview, and Tags

Datasets are accompanied by comprehensive metadata, offering essential information about the data they contain. This metadata includes a list of columns, descriptions for each column, and data types, providing users with a clear understanding of the dataset's structure and content. Additionally, an Overview section offers a concise summary of the dataset's key characteristics and significance. To further enhance organization and accessibility, users can add tags to datasets. Tags serve as labels or keywords, making it easier to search and categorize datasets effectively.

2. Preview - See Live Data

Datasets come to life with the "Preview" feature, offering users a glimpse of the actual data contained within. This live data preview allows analysts to quickly assess the dataset's contents before diving into in-depth analyses. It also aids in identifying any potential data issues or anomalies that may require attention.

3. Stories - Fostering Collaborative Discussions

The "Stories" section provides a collaborative platform for users to engage in discussions about datasets. It serves as a dynamic forum where data analysts, stakeholders, and team members can share insights, ask questions, and exchange valuable perspectives related to the dataset. Additionally, users can upload important documents, transforming the Stories section into a central hub for discussions and knowledge-sharing.

4. Transformations - Empowering Custom Datasets

For custom datasets, the VDA offers the "Transformations" feature. This powerful tool enables users to define data transformations using templates, facilitating data modeling and customization. Analysts can upload transformation logic and document the process, streamlining the preparation and refinement of custom datasets.

5. Publish and Generate SQL - Effortless Code Generation

Users can publish versions for transformations and also utilize the generate SQL option to generate code for the defined transformations. The "Generate SQL" feature leverages the power of OpenAI to automate code generation based on the logic defined in transformations. With just a few clicks, users can effortlessly convert their transformation logic into SQL code.



DASHBOARDS

Data visualization is an integral part of data analysis, as it allows users to derive insights and communicate findings effectively. However, creating visually appealing and informative dashboards traditionally involved time-consuming manual processes. With our groundbreaking feature of Natural Language to Dashboard Creation, data analysts can effortlessly build dynamic dashboards using Key Performance Indicators (KPIs) expressed in plain language, all powered by Superset, an open-source data visualization tool.

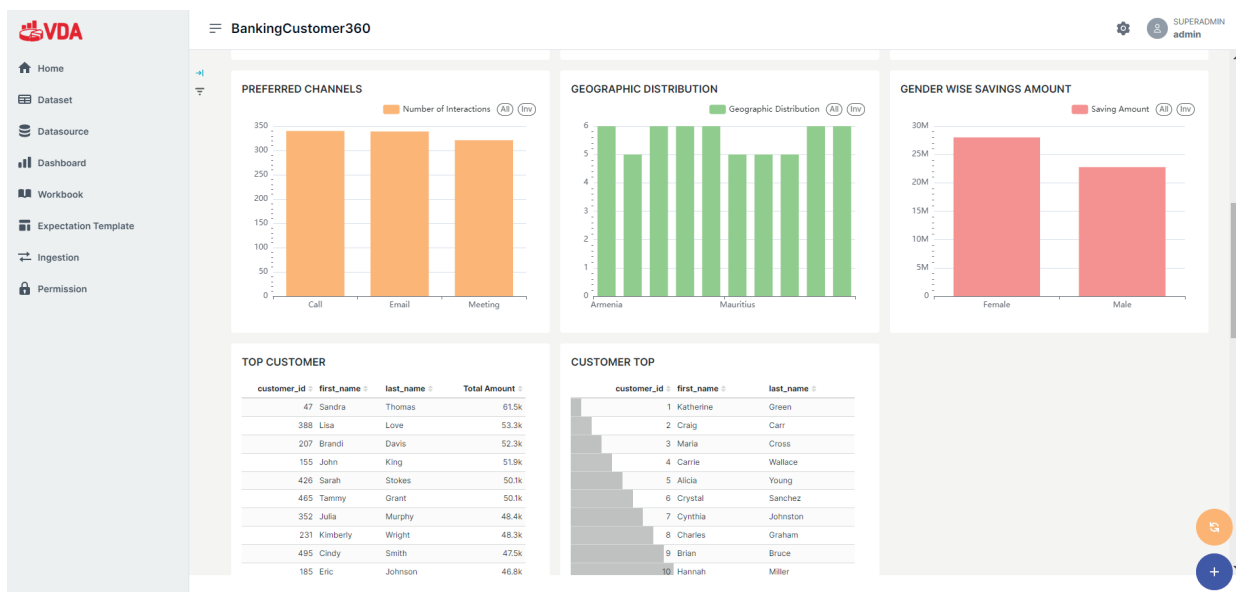
Below are the steps involved in the dashboard creation process :-

1. Select Datasource and Dashboard Description
2. Add KPIs Using Natural Language
3. Using the Natural Language Dashboard Creation feature, the user adds Key Performance Indicators (KPIs) to the dashboard by expressing their requirements in plain language.
4. The system generates SQL queries based on the natural language inputs and displays a preview of the resulting data in real-time on the right-hand side.
5. The user has the option to review the generated SQL queries and validate the results presented in the preview section.
6. If needed, the user can manually modify the SQL queries to fine-tune the data retrieval or apply custom logic for more advanced data processing.
7. To further enhance the dashboard's visual appeal and user experience, the user can provide design instructions, such as specifying the chart types, color schemes, layout, and other visual elements.
8. An API connects to Superset, the open-source data visualization tool, and adds the specified charts to the dashboard based on the SQL queries and design preferences.

Further enhance Charts in Superset

Once the charts are added to the dashboard, users have the flexibility to explore and edit the charts further within Superset.

With Superset's rich features, users can perform fine-grained customization, apply additional data filters, and create interactive dashboards tailored to their unique analytical needs.



WORKBOOK

Workbooks are dynamic playgrounds designed to empower data analysts, providing them with versatile tools to streamline and enhance their data analytics process. Our feature-rich Workbooks offer a trio of powerful options: Document Book, Query Book, and Expectation Book. Each workbook caters to different aspects of data analysis, offering an intuitive and efficient user experience.



DOCUMENT BOOK



QUERY BOOK

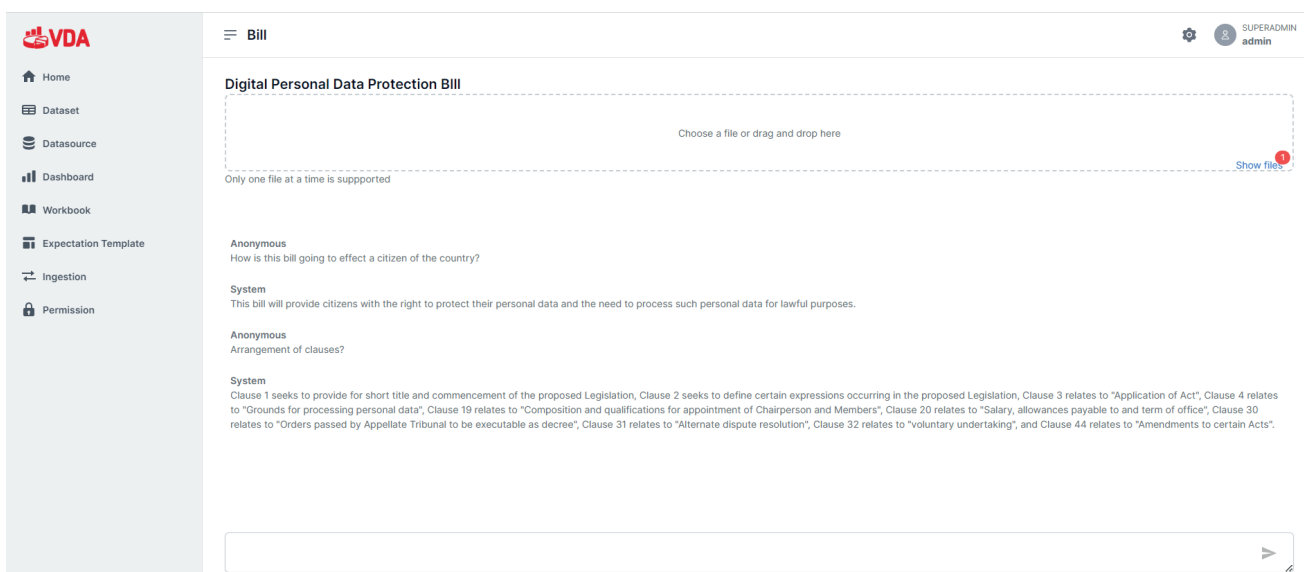


EXPECTATION BOOK

Document Book:

The Document Book revolutionizes the way data analysts work with unstructured data. By harnessing the power of Natural Language Processing (NLP), analysts can effortlessly upload and query documents using simple, everyday language. This interactive capability transforms the analysis of textual data into a seamless and insightful experience, significantly reducing the effort required to derive valuable insights from documents.

By leveraging the Document Workbook, analysts can now gain valuable insights from textual documents without having to sift through extensive pages of documentation manually. The feature's ability to scope workbooks, ask contextually relevant questions, and provide precise answers in real-time transforms document analysis into a seamless and efficient process, driving informed decision-making and enhanced collaboration among teams. Embrace the Document Workbook and revolutionize the way you interact with textual data, making knowledge discovery and analysis faster, more accurate, and more impactful.



The screenshot shows the VDA Document Book interface. On the left is a navigation sidebar with icons for Home, Dataset, Datasource, Dashboard, Workbook, Expectation Template, Ingestion, and Permission. The main content area is titled "Bill" and shows a document titled "Digital Personal Data Protection Bill". At the top right of the main area, there is a settings gear icon and a user profile icon labeled "SUPERADMIN admin". Below the title, there is a large dashed box for file upload with the text "Choose a file or drag and drop here" and a "Show files" button with a red notification dot. A note below the upload area states "Only one file at a time is supported". The document content is displayed in a scrollable view with the following text:

Anonymous
How is this bill going to effect a citizen of the country?

System
This bill will provide citizens with the right to protect their personal data and the need to process such personal data for lawful purposes.

Anonymous
Arrangement of clauses?

System
Clause 1 seeks to provide for short title and commencement of the proposed Legislation, Clause 2 seeks to define certain expressions occurring in the proposed Legislation, Clause 3 relates to "Application of Act", Clause 4 relates to "Grounds for processing personal data", Clause 19 relates to "Composition and qualifications for appointment of Chairperson and Members", Clause 20 relates to "Salary, allowances payable to and term of office", Clause 30 relates to "Orders passed by Appellate Tribunal to be executable as decree", Clause 31 relates to "Alternate dispute resolution", Clause 32 relates to "voluntary undertaking", and Clause 44 relates to "Amendments to certain Acts".



Query Book:

The Query Book empowers data analysts to interact with data through natural language inputs, seamlessly converting plain language queries into SQL. With Text to SQL capabilities, analysts can efficiently and accurately write complex SQL queries without the need for extensive coding knowledge. This feature enhances productivity and enables quick access to the precise data needed for analysis. Query Book is designed to streamline the process of SQL query generation, documentation, and validation. It empowers data analysts with an intuitive environment to create, store, and preview queries conveniently, eliminating the need to switch between different platforms for exploration and report generation.

Using ABC Bank as datasource [CHANGE](#)

Statement: top 10 customers from transactions

KPI Statement like Top 10 customers, Best selling items

DETAILS

Description: This query retrieves the top 10 customers based on the number of accounts they have.

Title: Top 10 Customers from Transactions

SQL: SELECT customer.customerid, customer.firstname, customer.lastname FROM customer ORDER BY customer.numberofaccounts DESC LIMIT 10

Buttons: SAVE, Preview

query-test

Top 10 Customers from Transaction

ABC Bank

This query retrieves the top 10 customers based on the number of accounts they have.

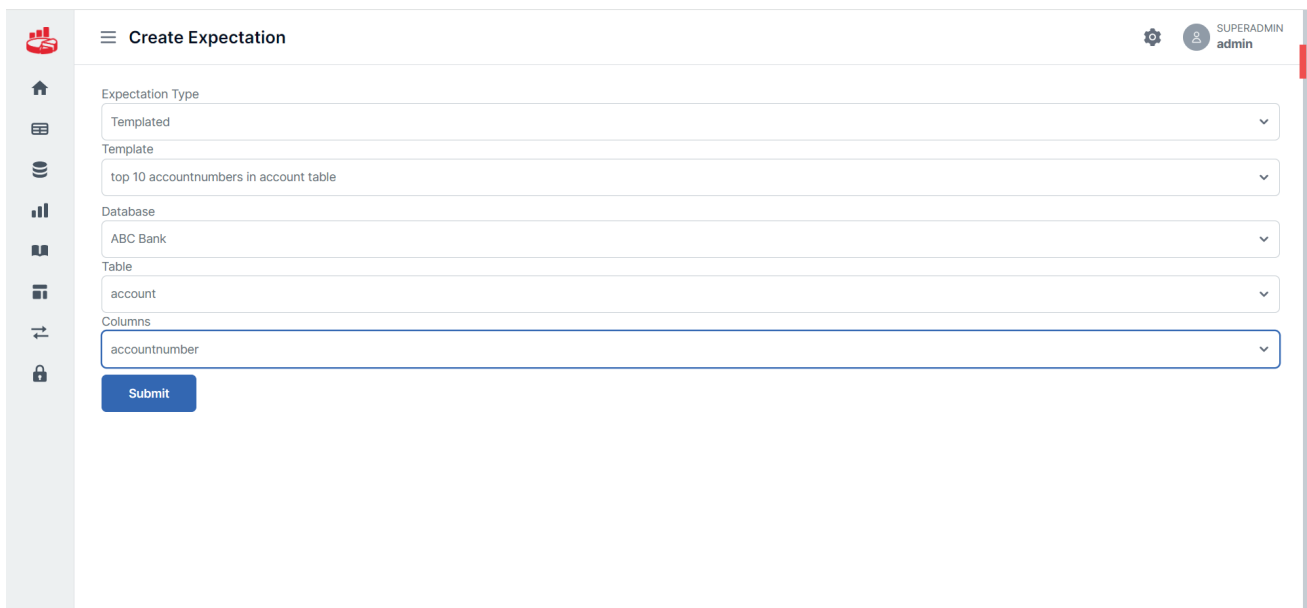
Preview

customerid	firstname	lastname
42	Rip	Ongin
51	Prentice	MacGahy
29	Jim	McAreevey
5	Kermy	Dummer
48	Alyce	Karpenya
20	Giraud	Hauxley
12	Chevy	Secretan
28	Kingsly	Lehmann
10	Margarete	Hawkey
54	Barron	Jillions

Expectation Book:

The Expectation Book elevates data quality assurance by allowing analysts to define and execute data quality tasks using expectations. By setting expectations for data, analysts can validate and monitor data accuracy, completeness, and consistency, ensuring reliable and trustworthy analytical outcomes. This feature fosters data confidence and empowers analysts to make data-driven decisions with certainty.

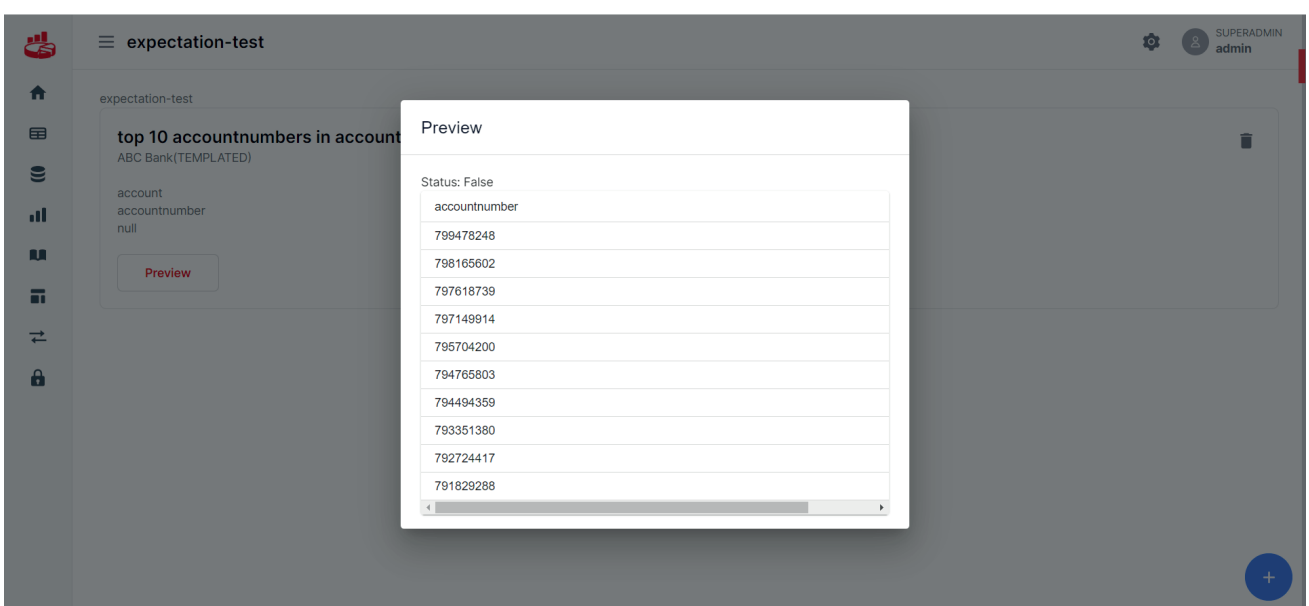
Expectation Books are designed to facilitate data quality assurance and validation. It empowers data analysts to define and run data quality tasks using expectations, ensuring that data meets predefined criteria and adheres to specific requirements. By setting and validating expectations, analysts can instill confidence in the data, leading to more reliable and trustworthy analytical outcomes.



The screenshot shows the 'Create Expectation' form. The form includes the following fields:

- Expectation Type:** Templated
- Template:** top 10 accountnumbers in account table
- Database:** ABC Bank
- Table:** account
- Columns:** accountnumber

A 'Submit' button is located at the bottom of the form.



The screenshot shows the 'expectation-test' page. The main content area displays the following information:

- expectation-test**
- top 10 accountnumbers in account**
- ABC Bank (TEMPLATED)
- account
- accountnumber
- null

A 'Preview' button is visible. A modal window titled 'Preview' is open, showing the following data:

accountnumber
799478248
798165602
797618739
797149914
795704200
794765803
794494359
793351380
792724417
791829288



INGESTION

The concept of "ingestion" primarily falls within the realm of data engineering. In contrast, business analysts often find themselves relying on proficient data engineers or resorting to alternative toolsets when dealing with data ingestion.

Ingestion revolves around three types: full refresh, incremental and historical, and you just fill in this particular form and it will perform the ingestion wait.

Create Ingestion Pipeline

Name: CRM pipeline
A unique common name which can also help in recognizing the system

Source: BankingTest

Customer_fact Incremental Primary: L_test Incremental: Select...	accounts Full-Refresh	churn_data Load type
customers Load type	geographic_data Load type	interactions Load type
marketing_campaigns Load type	test-dataset Load type	transactions Load type

SUBMIT

The concept of "ingestion" primarily falls within the realm of data engineering. In contrast, business analysts often find themselves relying on proficient data engineers or resorting to alternative toolsets when dealing with data ingestion.

Ingestion encompasses three primary types: full refresh, incremental, and historical. By completing a specific form, the ingestion process is initiated.

For instance, when considering incremental ingestion, questions such as determining the primary key and incremental key come into play. The metadata, previously collected using our sourcing tool, provides valuable information about which columns should serve as incremental and primary key columns. This approach significantly reduces the time required to move from data search to insights, often achieved within hours, thus minimizing dependency.

From a developer's perspective, this approach reduces the reliance of business analysts on developers for ingestion-related tasks and various coding activities, including SQL-based processes. This, in turn, accelerates the process of building a data narrative.



**Start your data-driven future today with VDA by your side.
Uncover the possibilities, make better decisions, and
unlock the full potential of your data.**

www.microui.in
sales@microui.in