

IBM TechXchange

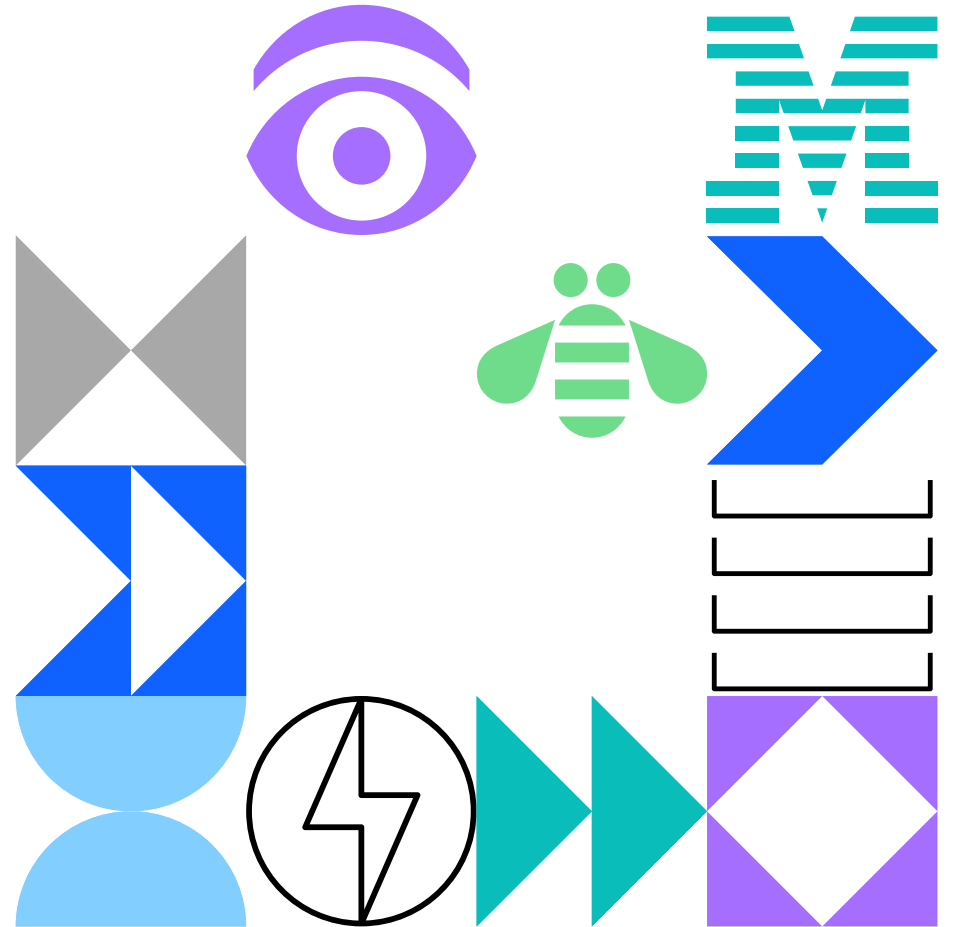
Announcing the IBM Database Assistant for Db2 Powered by watsonx

Session 3333

David Kalmuk
IBM, Distinguished Engineer

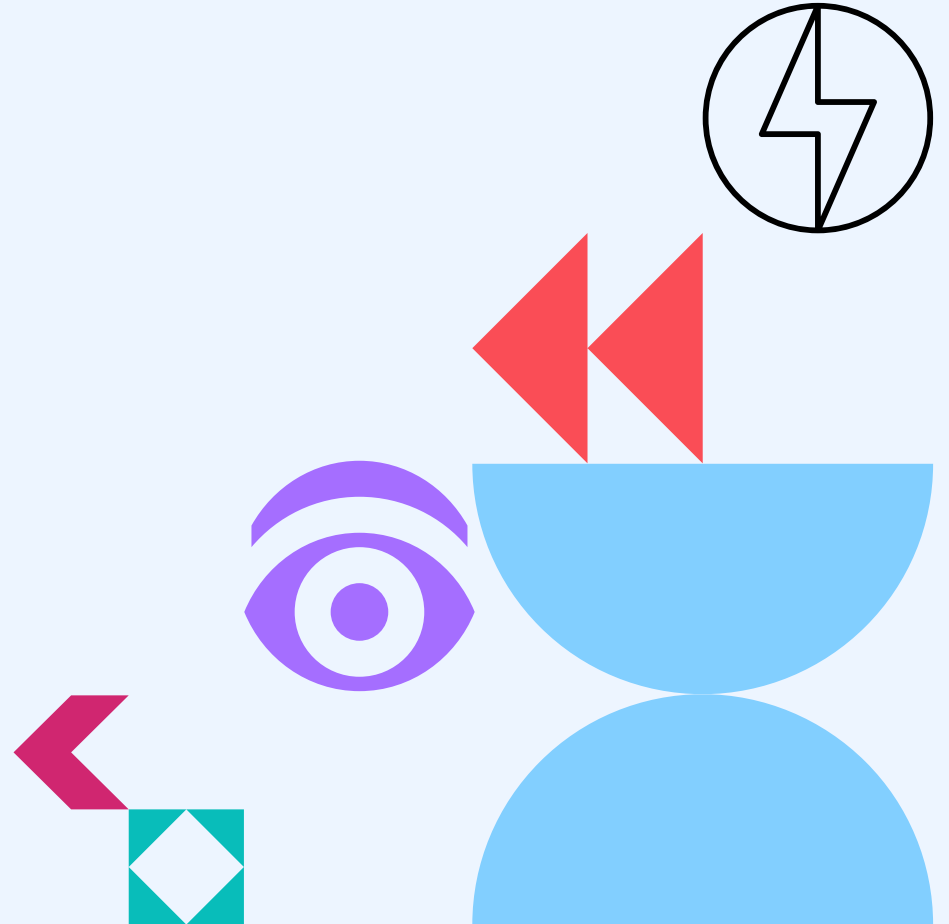
Krishna Guntuka
IBM, Advisory Software Engineer

October 21-24, 2024
Mandalay Bay Convention Center
Las Vegas, Nevada



Agenda

- 01 Introduction
- 02 Usage
- 03 Demo
- 04 Looking Forward
- 05 Q&A

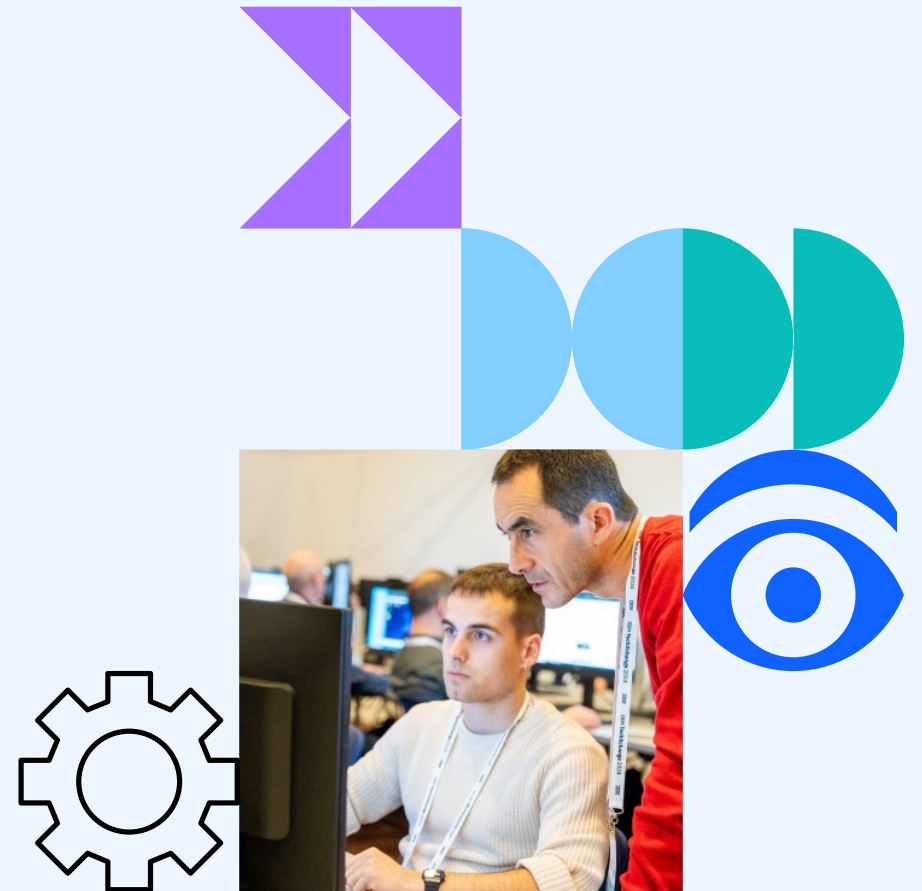


IBM TechXchange



Introduction

October 21-24, 2024
Mandalay Bay Convention Center
Las Vegas, Nevada









Sergio,
Database Administrator

Overview

Sergio is responsible for the overall health of his company's databases. He monitors the performance of the databases, ensures the databases are secure, and plans for the system's daily usage capacity. When problems arise, Sergio is on-call 24/7.

Responsibilities

-  Database performance monitoring + tuning
-  Database security
-  Storage and capacity planning
-  Installation, upgrade, configuration, migration

Team makeup

- Junior DBAs
- Application developers
- Data engineers
- Project managers

Goals

- Ensure that databases are running smoothly and securely (keep the lights on)
- Minimize interruptions by others
- Stay ahead of problems instead of reacting when problems arise



Sergio,
Database Administrator

Pain Points

Finding the right information / documentation

"I search for documentation daily, sometimes hourly. IBM documentation can feel like boiling the ocean. I use Google."

Training junior DBAs

"Staffing is an issue. Workloads are increasing. A value add would be a way to handle increased workloads/mixed workloads and not have to increase staffing."

Identifying and resolving the root problem

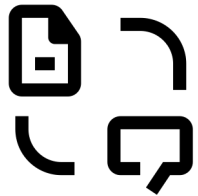
"I have a lot of information, but I don't know what is relevant to my current issue."

Optimizing + tuning the database

"Optimizing performance is complex and requires expertise. The current Tuning UX in DMC is complicated and not ideal."

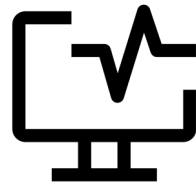
Introducing Database Assistant

powered by watsonx



Db2 Expert

Get answers to your Db2 questions, faster



Monitoring Metrics

Quickly access key Db2 metrics using natural language queries



Simplified Troubleshooting

Get recommendations for troubleshooting common database issues



Coming Soon!

Advanced Analysis

Accurately identify root cause of performance issues, bottlenecks, deadlocks

Solve challenges finding the right information through a

Conversational Knowledge Base

Finding the answer to database questions can be challenging:

Finding the right resource. Sometimes if the content doesn't change between versions, it won't be reproduced for the newer version and it causes confusion - user thinks it's outdated when it actually isn't.

Common resources

Senior DBAs (\$\$)
Message boards
Github
Stack Overflow
Blogs
IBM Redbooks and White papers
Support information
Others who have had similar issues
(our key consultants have created a community amongst themselves and often refer to each other for difficult problems)

"If I need something from IBM I don't go to the Knowledge Center, I use Google."

"I search for documentation daily, sometimes hourly. IBM documentation can feel like boiling the ocean. I use Google."

Our users have expressed a lot of concern around training the next generation of Db2 DBAs, and how easily they can access expert knowledge

Simplify the process of *navigating through multiplicity of database tasks* through

Conversational Database Skills

AI assisted navigation of basic database tasks such as:

- *Viewing database summary information*
- *Listing tables + schemas + indexes*
- *Checking active resource usage (CPU, IO, Memory)*
- *Checking storage utilization*
- *Viewing active sessions*
- *Viewing active queries*
- *Analyzing where time is being spent*
- *Analyzing lock waits*
- *Analyzing Top N queries + connections*



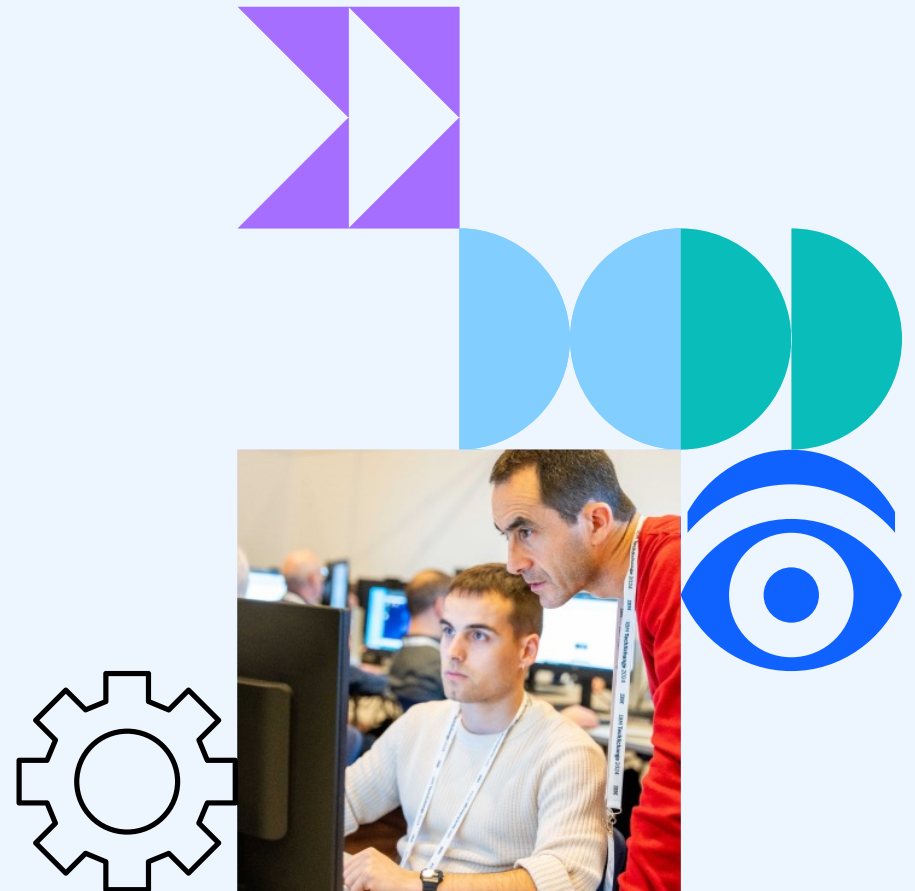
For a non-expert user, tasks normally involves cross referencing public documentation with unguided adhoc navigation of the available console panels and following a multi-step process to locate the required information, diagnose, and then resolve an issue.

IBM TechXchange



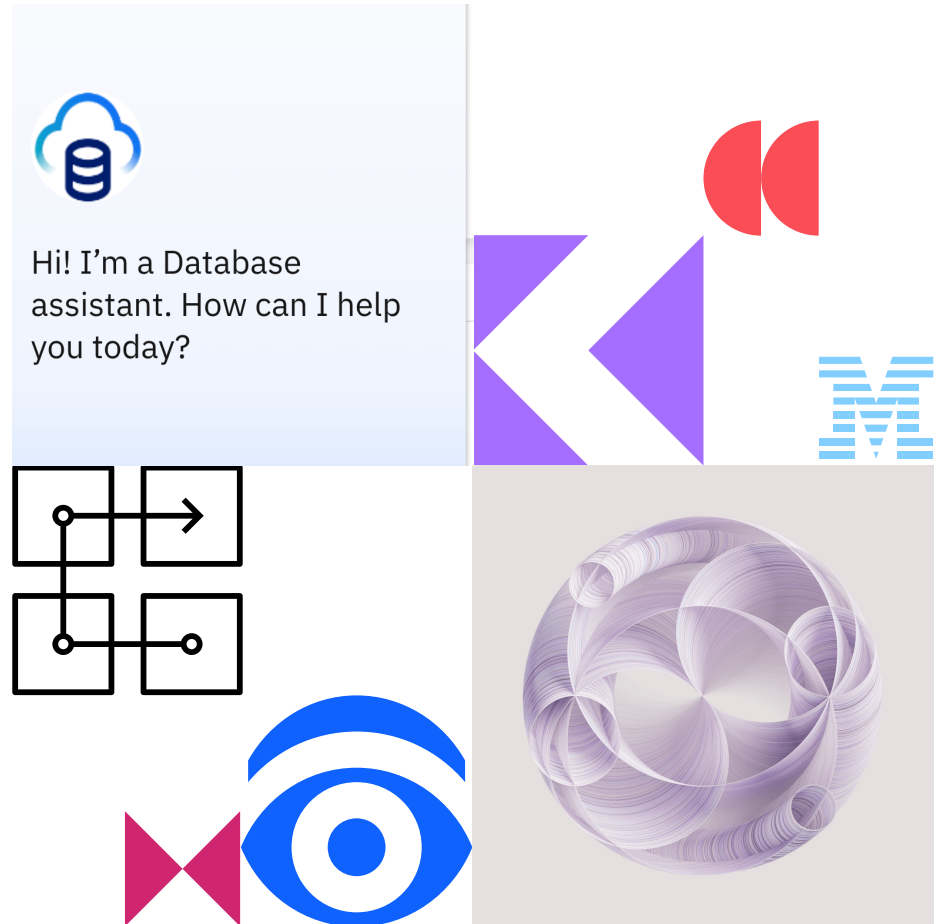
Assistant Usage + Demo

October 21-24, 2024
Mandalay Bay Convention Center
Las Vegas, Nevada

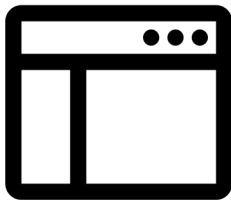


Benefits

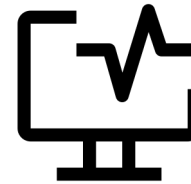
- Easy access to targeted grounded answers for technical questions.
- Reduce context switching and switching between different tools while diagnosing and fixing issues.
- AI guided tasks and troubleshooting to streamline the DBA's job.
- The Assistant is developed using a RAG based AI system to minimize hallucinations by retrieving information from trusted sources.



Where does it operate?

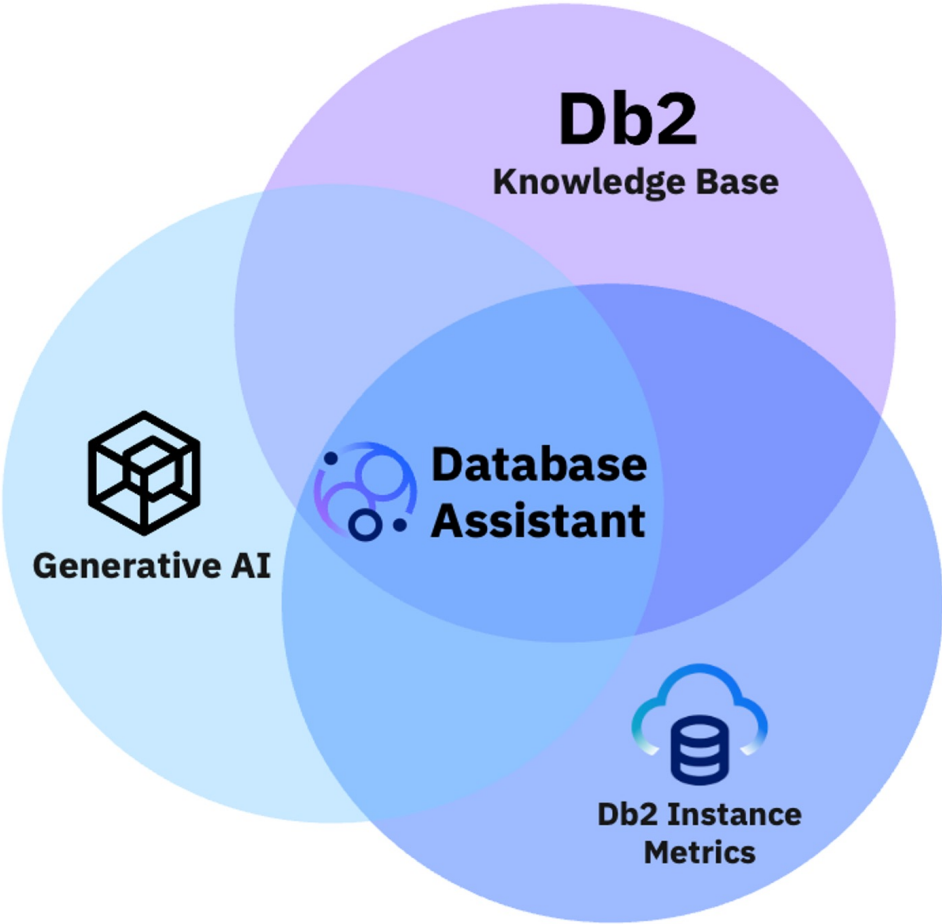


Database Assistant is built directly into your **Database Management Console (DMC)**.

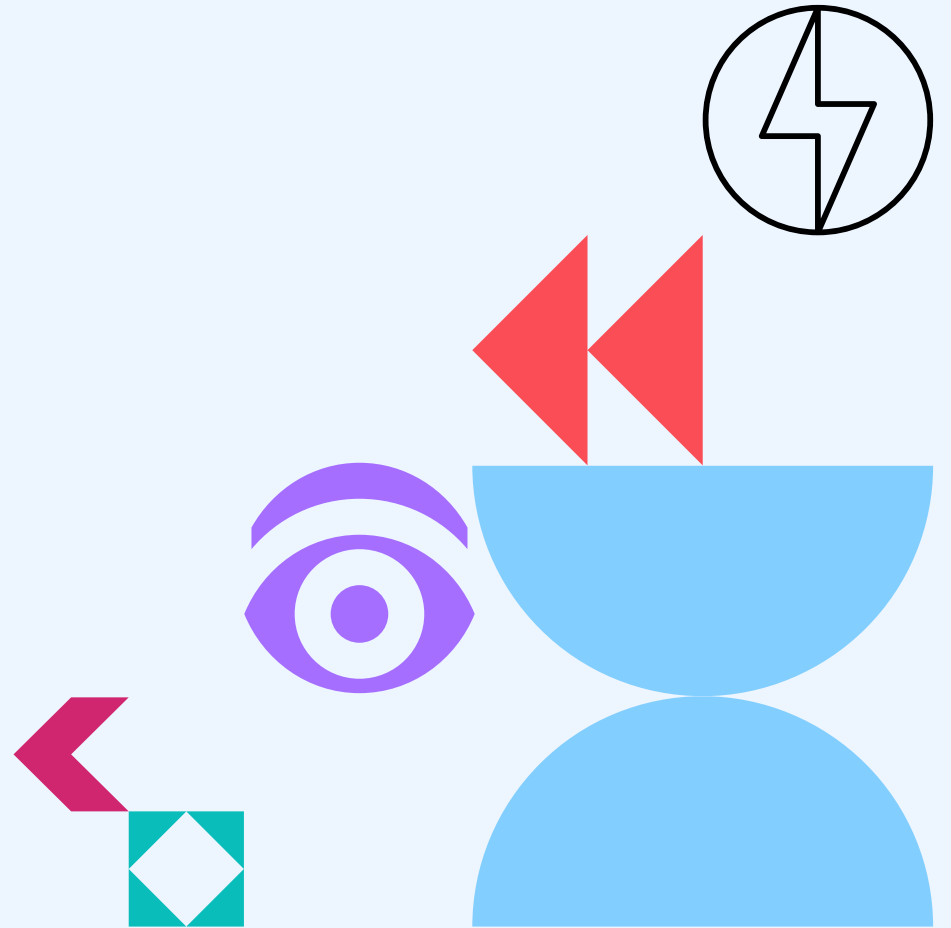


Database Assistant provides **real-time metrics** of your database instance.

Components



Demo





Scenario: Db2 Expert

Database Assistant AI

You 12:24 PM

How do I create a range partitioned table

watsonx 12:25 PM

To create a range partitioned table, you can use the CREATE TABLE statement with the PARTITION BY RANGE clause, as shown in the Db2 documentation. The basic syntax is:

```
CREATE TABLE table_name (  
    column_name data_type null_attr  
) IN table_space_list  
PARTITION BY RANGE (column_expres  
(STARTING FROM constant ENDING
```

For example, the following statement creates a table where rows with a value of 1 to 20 are in PART0, rows with a value of 21 to 40 are in PART1, and so on:

```
CREATE TABLE foo(a INT)  
PARTITION BY RANGE (a)  
(STARTING FROM (1) ENDING AT (1
```

This will create 5 data partitions: PART0, PART1, PART2, PART3, and PART4. ▾

Type something...

Finding the answer to database questions can be challenging.

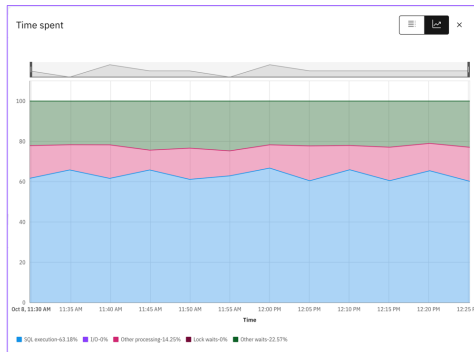
Quick Answers: Type your database question into the chatbot and get immediate, expert-level advice.

Tailored Solutions: Responses are tailored to your specific Db2 version.

Trustworthy Sources: Database Assistant provides source documents to ensure you have reliable, accurate information.



Scenario: **Monitoring Metrics**



Database Assistant

You 12:26 PM

Show me the summary of time spent on the database

watsonx 12:26 PM

Here is how time is being spent on your system from October 8, 2024 11:26:58 AM to October 8, 2024 12:25:58 PM .

Time spent
See information about how time was spent on your system

[View details](#)

Type something...

Need to check on key Db2 metrics without digging through dashboards?

Instant Access: Check key database metrics instantly using natural language queries.

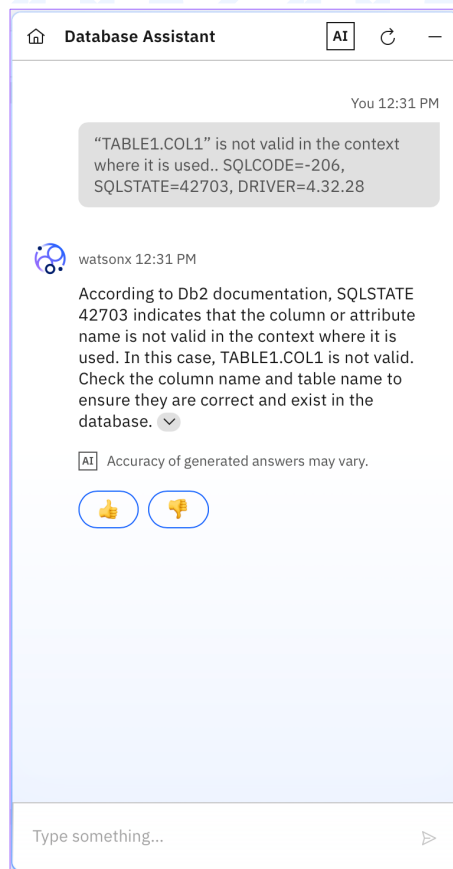
Filter by Time: Refine results by specifying a time period.

No Dashboards: Avoid complex dashboards and menus—get the exact data when you need it.

Proactive Management: Gain real-time insights to resolve issues faster and optimize performance.



Scenario: Troubleshooting



Facing an issue? Describe your problem to Database Assistant

Immediate Recommendations: Describe your issue, and Database Assistant will guide you through the solution.

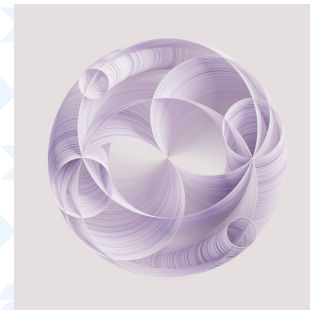
Error Identification: Instantly identify the root cause of errors, minimizing trial-and-error.

Fewer Disruptions: Resolve issues quickly to avoid downtime and operational disruptions.

Stay Focused: More time for strategic improvements, less time firefighting technical problems.

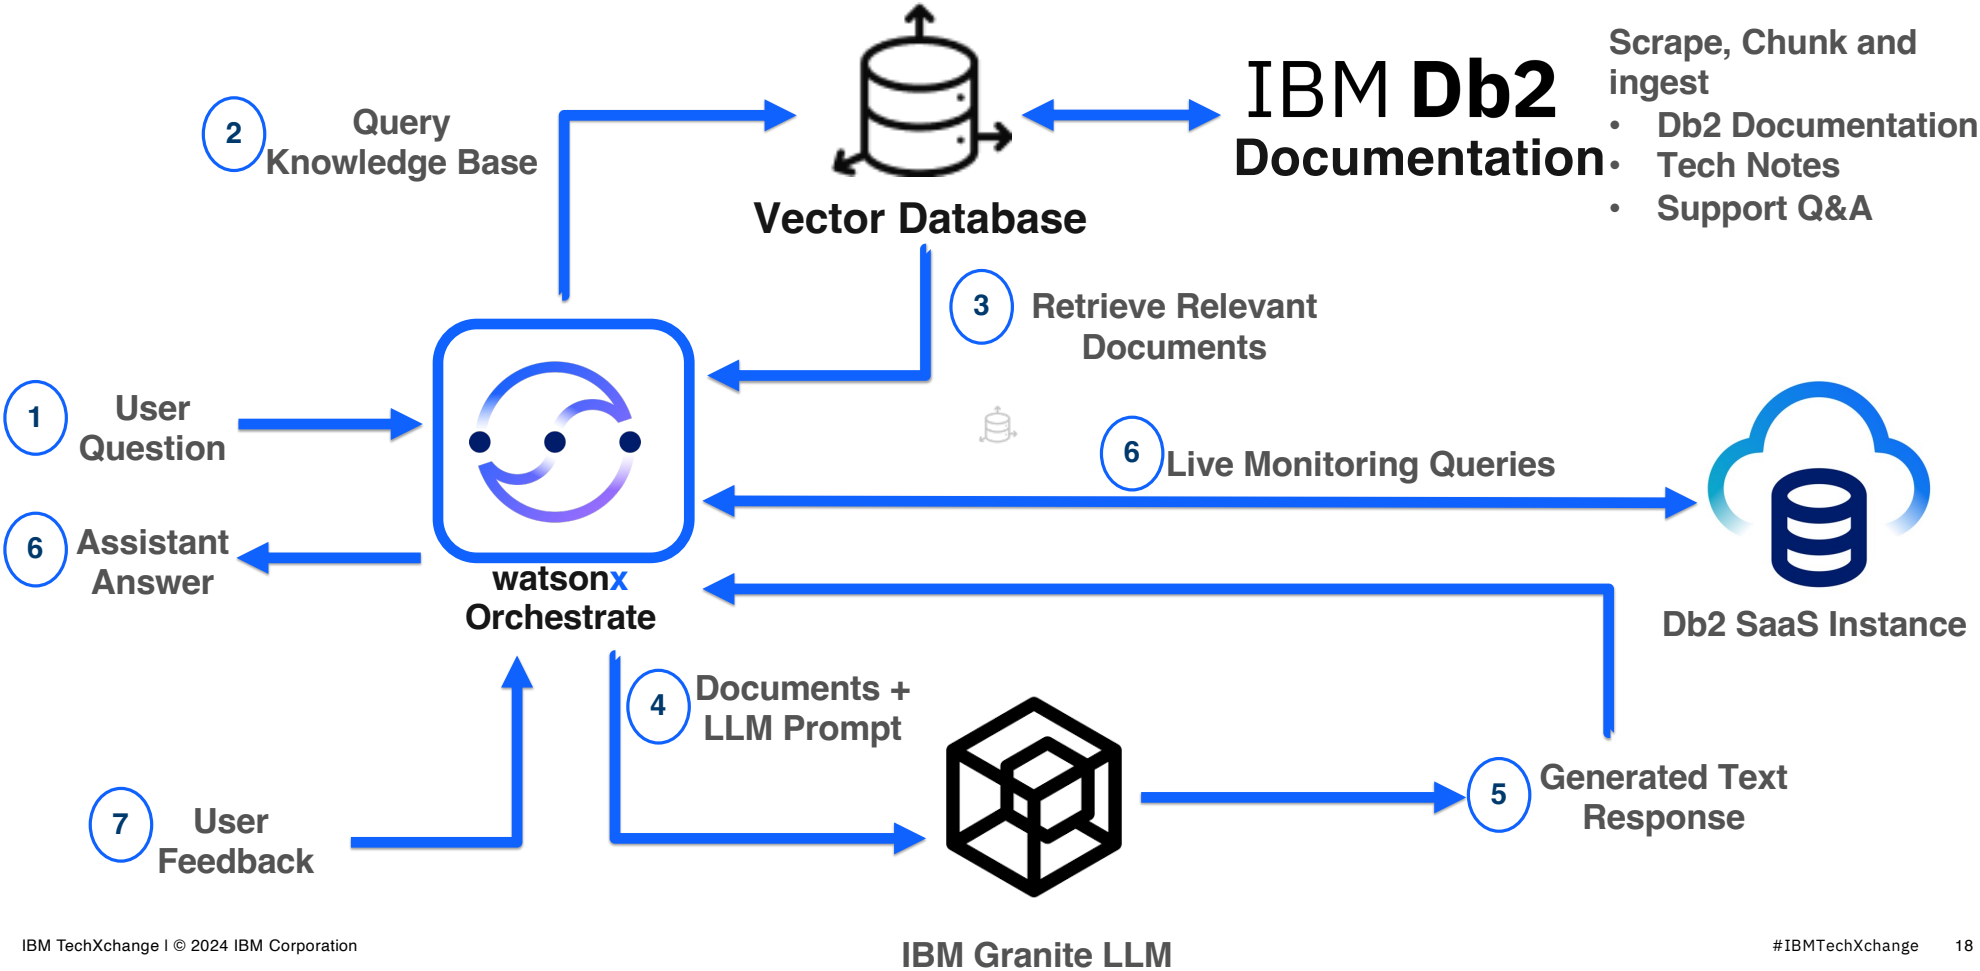
Technologies

- IBM Watsonx Orchestrate
- Vector Database
- IBM LLM – Granite 3.0



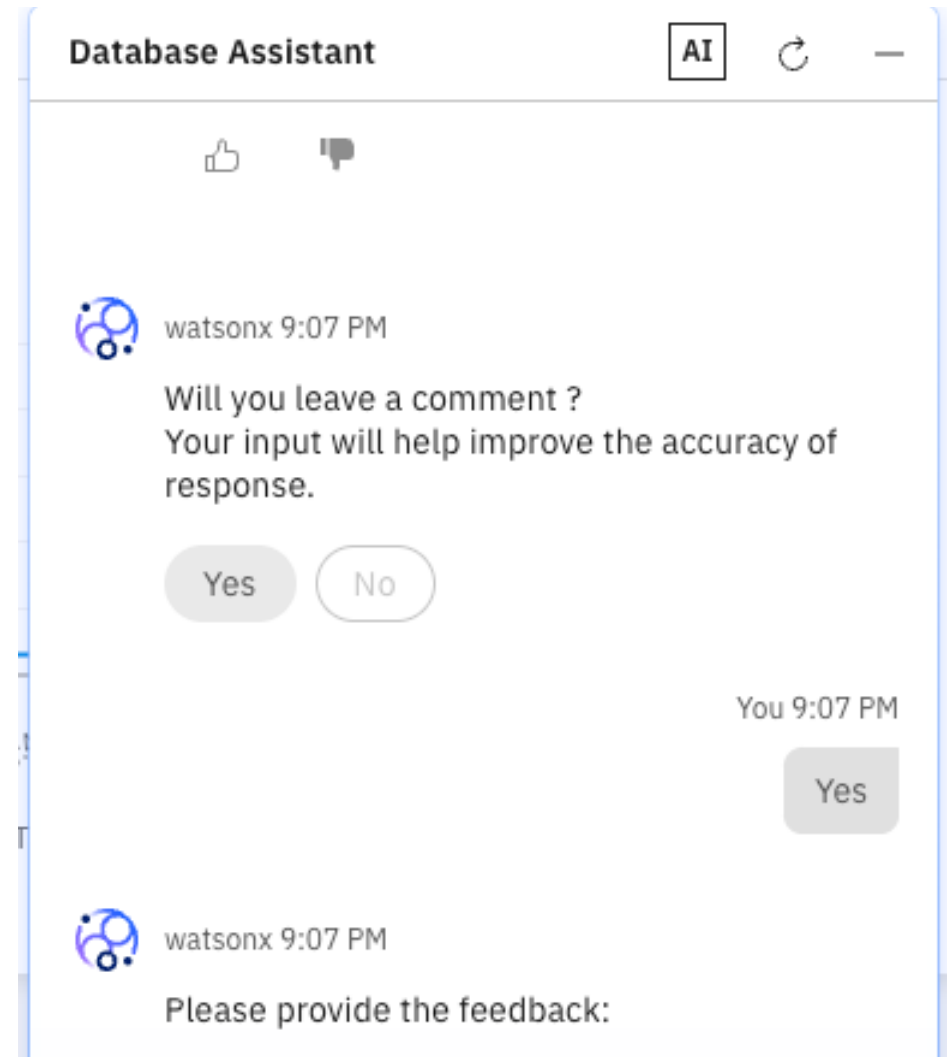
watsonx™ Orchestrate

Technical Architecture



Learning from Continuous Feedback

- One of the key differentiators for the assistant is continuous feedback and agile improvement.
- Users can provide active feedback on the quality of responses which is then factored into improvements to answer quality.
- Accuracy and quality of answers will continuously improve over time



IBM TechXchange



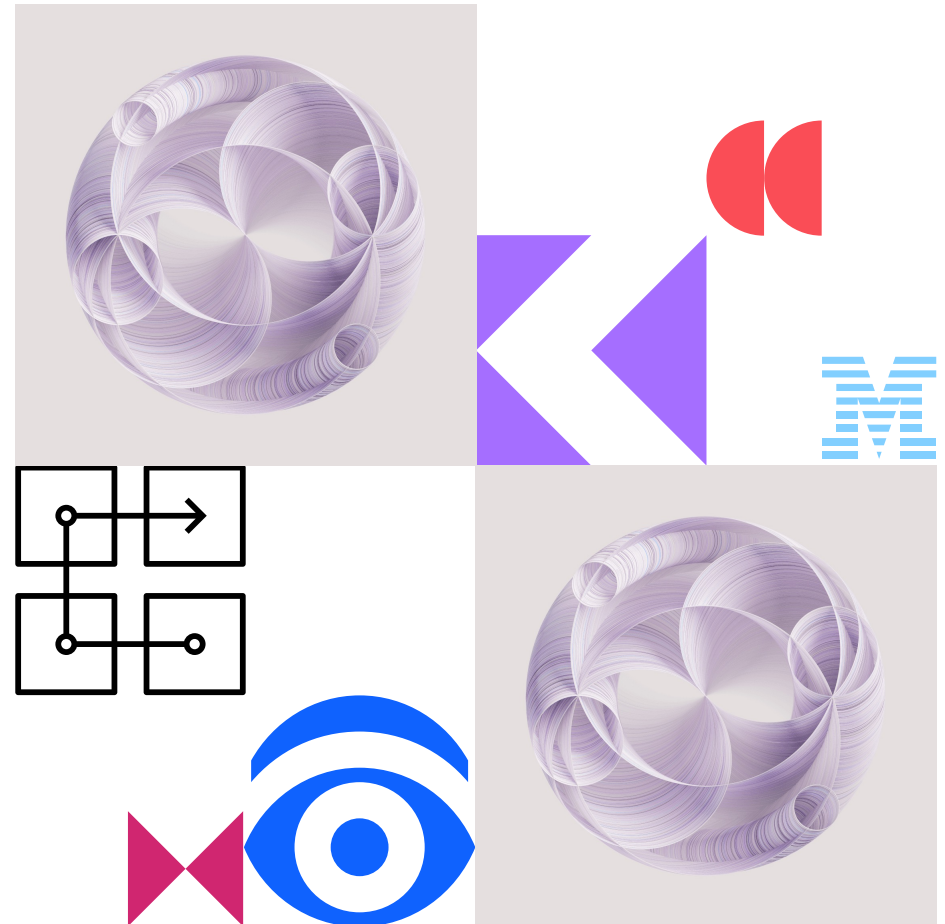
Looking Forward

October 21-24, 2024
Mandalay Bay Convention Center
Las Vegas, Nevada



The Power of Generative AI

- AI and Generative AI is a dynamic and rapidly evolving field.
- New models, new approaches, new techniques.
- Active research in the industry and within IBM
- We intend to rapidly evolve and expand the capabilities of the Database Assistant
- Both in terms of filling out existing capabilities and leveraging new research and state of the art....

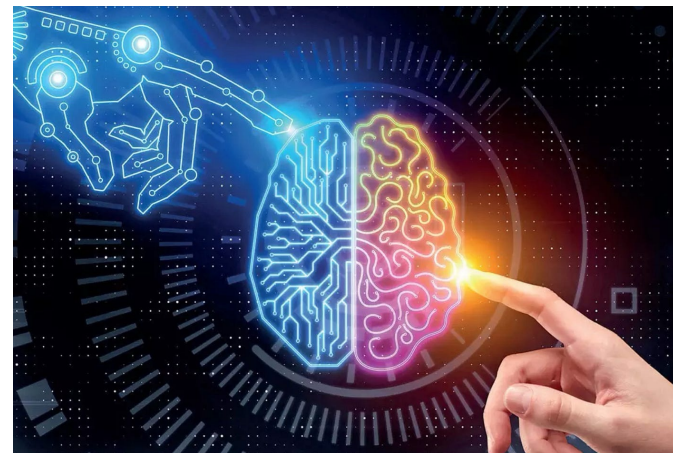


Simplify the process of *navigating through multiplicity of database tasks* through

Advanced Conversational Database Skills

AI assisted advanced tasks including:

- *Query tuning*
- *Index recommendations*
- *MQT recommendations*
- *Partitioning recommendations*
- *Clustering recommendations*
- *Configuration recommendations*
- *Resolving locking issues*
- *Resolving blocking transactions*
- *Resolving blocking queries*
- *Resolving compression + table skew issues*
- *Resolving IO performance issues*
- *Resolving WLM queueing issues*



For a non-expert user, tasks normally involves cross referencing public documentation with unguided adhoc navigation of the available console panels and following a multi-step process to locate the required information, diagnose, and then resolve an issue.

Simplify the process of *navigating complex database administration tasks* through

Conversational Telemetry Analysis and Forecasting

Combination of conversational input and standardized monitoring metrics opens up options for wider analytics and forecasting leveraging natural language innovations being developed within IBM:

- Look at trends in resource consumption to project future capacity needs
- Identify workload patterns to find optimization opportunities
- Identify outliers in behaviour to proactively identify problems

It combines reasoning AI that understands your database metrics and builds models automatically, so you know what happened, why it happened, what will happen, and what you can do. Now everyone in your organization can harness an analytics brain to think, reason, discover insights, and take actions on your database—all at the speed-of-thought.

Simplify the process of *navigating complex database management tasks* through

AI Assisted Problem Solving

One of our longer-term goals is to enable AI assisted troubleshooting by leveraging LLM driven reasoning incorporating multi-modal inputs:

User Problem / Question →

→ *Applicable Documentation + Support Information (RAG Search)* →

→ *Applicable Live System Metrics (LLM selected skill invocation)* →

→ *Solution generation (LLM driven solution extrapolation + summarization from knowledge based + metrics)* →

→ ***AI Generated Troubleshooting Guidance For My System***

Leverage AI that synthesizes knowledge base + live metrics to provide intelligent guidance and accelerate DBA tasks

Q&A



David Kalmuk

IBM, Distinguished Engineer

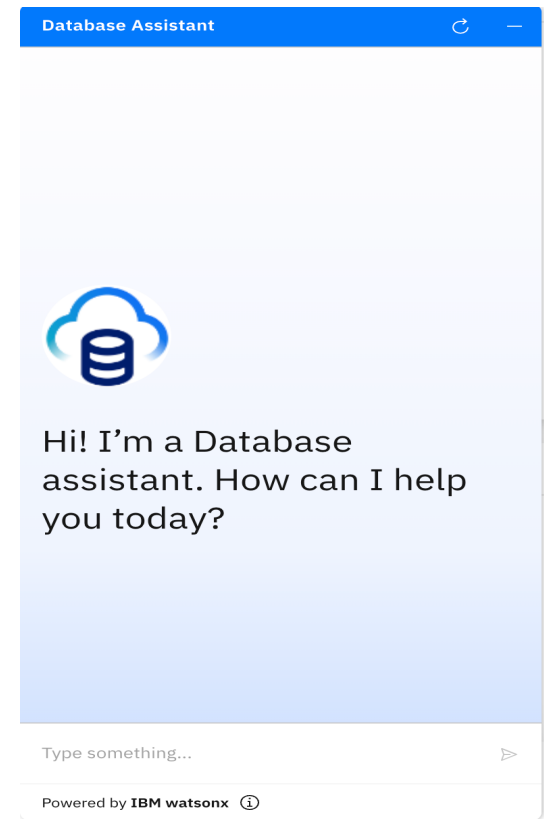
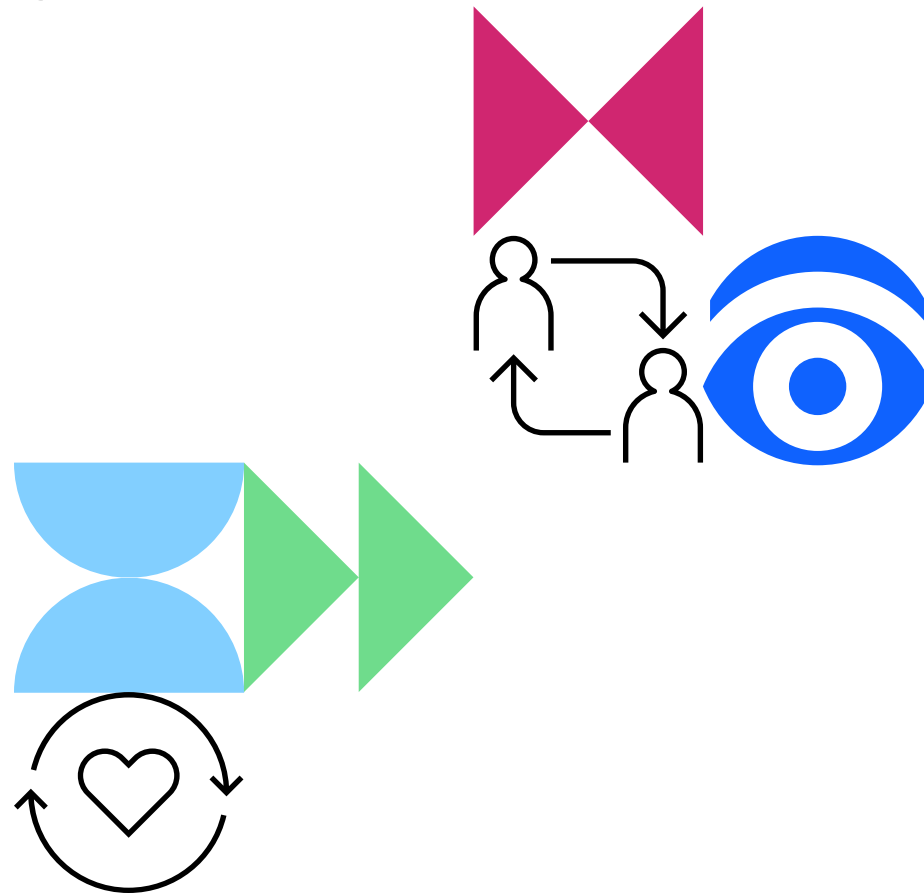
Contact: dckalmuk@ca.ibm.com

Krishna Guntuka

IBM, Advisory Software Engineer

Contact: Krishna.Guntuka@ibm.com

Thank You



IBM