



Microsoft Fabric

This presentation is the property of Microsoft and is intended for informational and educational purposes only. You may use, copy, and distribute this presentation for your personal, non-commercial purposes. You may not modify, alter, or create derivative works from this presentation without the prior written consent of Microsoft. You may not use this presentation to misrepresent, defame, or disparage Microsoft or its products, services, or affiliates. You may not use this presentation to endorse or promote any other products, services, or organizations without the prior written consent of Microsoft.

By using this presentation, you agree to abide by these terms. If you do not agree, you must not use this presentation. Microsoft reserves the right to change these terms and conditions at any time without notice. Microsoft disclaims any and all warranties, express or implied, relating to this presentation, including but not limited to the accuracy, completeness, timeliness, or suitability of the information contained herein. Microsoft is not liable for any damages, losses, or liabilities arising from your use of or reliance on this presentation.

Please review the terms of use posted in the content library.

#FABCONSQLCON2026

FABCON

Microsoft Fabric
COMMUNITY CONFERENCE

SQLCON

Microsoft SQL
COMMUNITY CONFERENCE

ATLANTA MARCH 16 - 20, 2026



Beyond Monitoring: AI Driven Spark Optimization in Microsoft Fabric

Anu Venkataraman, Senior Program Manager
Virginia Roman, Senior Product Manager





Fabric Spark Advisor

AI-Powered Recommendations!

- ✓ Optimize Jobs!
- ✓ Boost Performance!
- ✓ Save Costs!



Awesome!

Yay!

Agenda

- How Spark Really Works: *Execution model + where logs come from*
- Fabric Spark Live Monitoring 360
- Data Agent Advisor
- Beyond Monitoring: MCP + RAG + LLM Judge
- 🎁 Live surprise


How Spark Executes Code

Driver runs the user's *main* function and executes the various parallel operations on the worker nodes.

The worker nodes read and write data from/to data source.


Spreads the processing and data on to different Worker nodes

- The results of the operations are collected by the driver




How Spark Emits Logs

Every slow job leaves evidence — events emitted from Driver and Executors



Fabric Spark Live Monitoring 360 (Why)




Accessible Logs & Metrics

Logs & metrics dispersed across multiple apps & workspaces

- Exporting logs across cross workspace Spark applications is cumbersome
- Difficult to get consolidated view of all Spark applications across workspaces
- Interpreting raw event logs requires deep Spark internals expertise

v

Make logs and metrics accessible



Extend X-App Monitoring

No aggregated view across apps, teams, or time ranges

- Spark History Server shows one application at a time — no aggregated view
- vCore usage visible per job, but lacks cumulative capacity level dashboards
- Native retention limited to 30 days, blocking long-term trend analysis

v

Enhanced monitoring & recommendations

Fabric Spark Live Monitoring 360 (Key User Scenarios)

Three personas — one observability platform



Admin / Business

Capacity, Cost & Performance

"Are my Spark jobs worth what I'm paying for them?"

Real-time vCore usage vs. free capacity — know if more jobs can run right now

Top 10 notebooks by Capacity Units consumed — find your most expensive workloads

Top 10 apps by execution duration — spot long-running jobs costing you most



Data Engineer

Deep Perf Insights &
Troubleshooting

"Why is my job slow — and what do I change?"

Driver CPU/memory utilization per application — find driver vs. executor bottlenecks

Data spill detection + drill-down — stage-level spill with root cause

Executor wall-clock time breakdown — see exactly where time is going

NEE fallback monitoring — track when Native Execution Engine degrades to JVM



Spark Architect

Configuration, Scaling &
Governance

"Is our Spark config right for our workloads?"


High-concurrency sizing guidance — determine if driver sizing should increase

Best-practice config audit — Optimize-write, VOrder, NEE — all checked

Scaling predictions — Amdahl's Law tells you what adding executors actually buys

Fabric Spark Live Observability 360°


Real-time ingestion · Expert recommendations · Live monitoring · NL querying



Fabric Data Agent

A **virtual analyst** that allows users to interact with and gain insights from enterprise data and knowledge in Fabric

Now **Generally Available!**



Spark Simulator for the Demo

● SELECT JOB

ANTI-PATTERNS

- 1 driver_heavy_collectAP
- 2 driver_heavy_mapAP
- 3 no_caching_recomputeAP
- 4 python_udf_instead_of_b...AP
- 5 groupbykey_instead_of_r...AP
- 6 excessive_repartitioningAP
- 7 cartesian_product_joinAP
- 8 count_then_collectAP
- 9 small_file_problemAP
- 10 broadcast_large_tableAP
- 11 iterative_without_cachi...AP
- 12 select_star_large_tableAP


PERF ISSUES

- 13 shuffle_groupbyPI
- 14 shuffle_joinPI
- 15 spill_large_joinPI
- 16 spill_aggregationPI
- 17 skew_key_joinPI
- 18 skew_groupbyPI

BEST PRACTICES

- 19 pharma_etlBP
- 20 adverse_event_aggBP
- 21 patient_journeyBP

driver_heavy_collectAnti-Pattern



🔍 EVENT LOG0 events


Awaiting events...
Select a job and press Run

How the Recommender Notebook Works

① READ	② SCORE	③ DETECT	④ WRITE
Pull from Eventhouse SparkApplications SparkJobs · SparkStages SparkTasks (live KQL)	Calculate 0–100 Amdahl's Law Parallelism · GC overhead Skew ratio · Shuffle	Match Anti-patterns Cartesian joins Collect() abuse · UDFs Small files · Spill	Push Recommendations PerformanceScores Recommendations table

SCALING PREDICTIONS — Amdahl's Law

$$\text{speedup} = 1 / (\text{serial_fraction} + \text{parallel_fraction} / N)$$

Executors	Wall Clock	Duration	Speedup	Rating
0.5× (4 exec)	8m 12s	9m 45s	0.7×	Poor
1.0× (8 exec)	4m 28s	5m 10s	1.0×	Current
2.0× (16 exec)	2m 41s	3m 15s	1.6×	Good
4.0× (32 exec)	1m 58s	2m 38s	1.9×	Fair

FABRIC-SPECIFIC RECOMMENDATIONS

Spark Property	Recommendation
Native Execution Engine	True
High Concurrency Mode	True
Auto Compaction	True
Adaptive File Size	True
Optimize Write	False
VORDER (Parquet)	False

Fabric Spark Live Observability 360° - Demo

Fabric

Home

Workspaces

Copilot

OneLake catalog

Monitor

Real-Time

Workloads

SPRK WS 001

...

Fabric

Search

70

SPRK WS 001

Workspace for Analytics

Create deployment pipeline

Create app

Manage access

Workspace settings

+ New item

New folder

Import

Migrate

Filter by keyword

Filter

Choose from predefined task flows or add a task to build one

Select from one of Microsoft's predefined task flows or add a task to start building one yourself.

Select a predefined task flow

Add a task

Import a task flow

	Name	Status	Type	Task	Owner	Refreshed	Next refresh	Endorsement	Sensitivity	Included in app
	fabric-platform-monitoring		Folder	—	—	—	—	—	—	
	stateful-streaming-rocksd		Folder	—	—	—	—	—	—	
☐	Deploy Jumpstart		Notebook	—	SparkLabDeployer	—	—	—	—	

 Microsoft | Spark Monitoring | Confidential\Microsoft Extended  Search   23    

 Home
 Copilot
 Create
 Browse
 OneLake catalog
 Apps
 Scorecards
 Workspaces
 Ananke
 Spark Monitoring

 Power BI

Home

   |  Add data  |  Agent instructions |  Publish |  Revert to published version

Explorer 

Data Setup

Add Data 

 Spark Monitoring
 >  Tables

Test the agent's responses  Clear chat 

Spark Monitoring

Ask a question about this data, using your own words. For example:

 What are the historical trends across all my data?

 Analyze recent data for any outliers

 Show me the details for a specific subcategory of data

 Sample questions


Ask a question to test the data agent's response 

Created with AI. Mistakes are possible. [Review terms](#)

Beyond Monitoring — Fabric Spark Advisor 360°

Real-time ingestion · MCP on Eventhouse · LLM Judge · RAG · Fabric Notebook

③ ANALYZE & RECOMMEND



EventHouse

Spark Monitoring DB

MCP Server

10 Kusto tools · SSE

- sparklens_metadata
- sparklens_summary
- sparklens_metrics
- sparklens_predictions
- sparklens_recommendations
- fabric_recommendations

④ ACT & OBSERVE

Spark Advisor



LLM Judge

Docker · Azure Container Apps

NLToSparkInsights



Azure AI Search

Spark Docs · Best Practices



Chat UI in Fabric Notebook

Explorer

1 %pip install --quiet ./builtin/fabric_spark_advisor_ui-0.1.0-py3-none-any.whl

[1] ✓ 28 sec - Session ready in 9 sec 986 ms. Command executed in 18 sec 140 ms by Anu Venkataraman on 2026-03-09, 4:55:09 p.m.

Log

[notice] A new release of pip is available: 24.0 -> 26.0.1

[notice] To update, run: python -m pip install --upgrade pip

Note: you may need to restart the kernel to use updated packages.

Warning: PySpark kernel has been restarted to use updated packages.

1 from fabric_spark_advisor_ui import show

2

3 show()

[7] ✓ <1 sec - Command executed in 256 ms by Anu Venkataraman on 2026-03-09, 5:17:50 p.m.

Fabric Spark Advisor

AI-powered Spark optimization — MCP · RAG · LLM Judge

Connected Open

Fabric Spark Advisor

LIVE · EVENTHOUSE

Kusto · Connected DOCS SCHEMA SETTINGS

NAVIGATION

Chat

Fleet Overview

App Search

Judge Results

RECENT APPS

_0001

_0018

_0011

FLEET HEALTH

Slowest Apps

Critical Issues

Fleet Summary

Analysis Recommendations Scaling Trends Docs

Session · 00:00

Fabric Spark Advisor

REAL-TIME · SPARKLENS ANALYSIS · KUSTO TELEMETRY

Connected to live Eventhouse telemetry. Query your Spark workloads with evidence-first analysis — every finding cites its source table. Tier 1 data is always shown verbatim, never paraphrased or re-scored.

Kusto Telemetry Fabric Recommendations SparkDocs RAG GPT-4o Fallback

APP ANALYSIS

analyze application_XXX_XXXX

compare app A and app B

improve performance app XXX

show trend for app XXX

PROBLEM DETECTION

show apps with critical issues

find apps with high GC overhead

which apps have driver bottleneck

show streaming jobs with skew

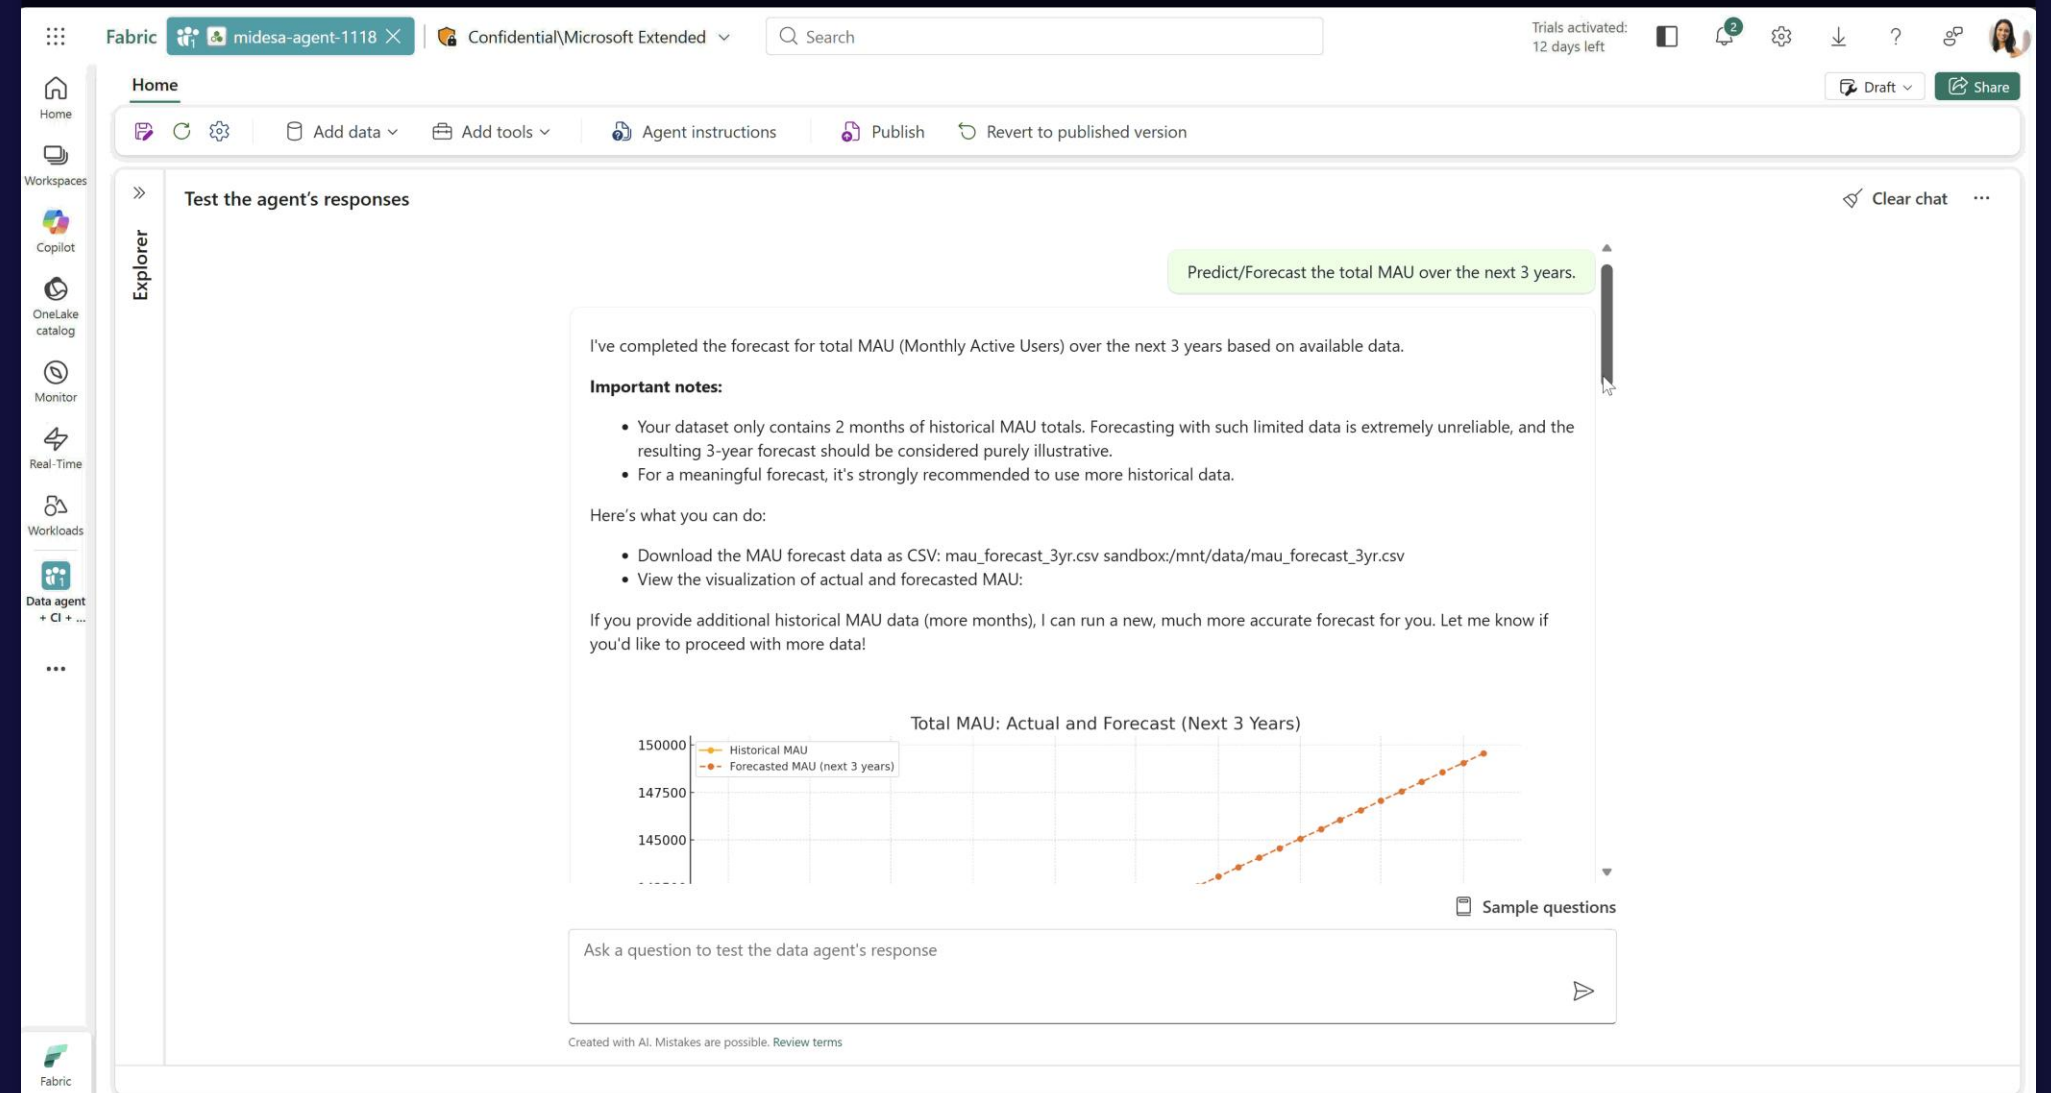
PERFORMANCE

DOCUMENTATION

Data Agent Roadmap


More Sophisticated Analysis with Code Interpreter

- Data Agent **generates and executes Python** in real time to perform deeper, more flexible analysis
- All code runs in a **sandboxed environment**, isolated from your system
- Tackle **complex logic, statistics, and visualizations**, while seeing the generated code and outputs for trust and reuse




Sneak Peek: Visuals

- Make complex data instantly understandable with **built-in visualizations** in the Data Agent
- Supported visuals include **line, bar, stacked bar, pie, scatter, area, and stacked area charts**
- Charts are **interactive** — hover for **tooltips** and **select/deselect series** to focus your analysis



Coming soon: Assisted Setup Mode

- Use a collaborative approach to generate contextual configurations for data agent
- Improve existing data agents using an iterative approach to refine instructions and example queries
- Validates configurations using creator feedback to resolve ambiguity and conflicts



Coming soon: NL2SQL Improvements

- Access to **new sub-tools** to explore and generate a query
- Enhanced support for **example query following** and **filter value substitution**
- **New data exploration tools** to assess joins and filter values for 'novel' questions

The screenshot displays the Microsoft Fabric Customer Analysis Agent interface. The top navigation bar includes the Fabric logo, workspace name 'Customer360', and the agent name 'Customer Analysis Agent'. The main content area is titled 'Test the agent's responses' and shows a chat conversation. A user query is: 'Show top 20 VIP customers by card-not-present spend this year.' The agent's response states: 'The top 20 VIP customers by card-not-present spend this year can be found in the file 'result_1.json'. This file includes customer ID, name, email, transaction count, and total spend. If you would like to see the details for these customers, please let me know!'

Below the response is a table titled 'Table 1' showing 1 row of data:

	ABC customer...	ABC full_name	ABC email	123 cnp_txn_co...	12F total_cnp_spend_usd
1	CUST0000158	Sanjay Garcia	sanjay.garcia4482@yahoo.com	1	31.14

Below the table is a list of sub-tools used in the response, each with a checkmark and a dropdown arrow:

- ✓ Analyzed Customer360 Lakehouse for: "Show the top 20 VIP customers ranked by their card-not-present spend for this year."
- ✓ Row Sampling
- ✓ Value Distribution
- ✓ Execute SQL
- ✓ Submit Final SQL

The interface also shows a 'Sample questions' section at the bottom with a text input field and a submit button.


Q&A

Get Two Fabric Certifications for FREE

Attendees of FABCON can take the Fabric Analytics Engineer or Fabric Data Engineer exam for free. Be part of the 2 fastest growing role-based certifications in Microsoft history.

Request your voucher by March 23, 2026.


<https://aka.ms/fabcon/cert100>



How was the session?



Complete Session Surveys in
Uhova for your chance to WIN
PRIZES!



Sound off.
The mic is all yours.
Influence the product roadmap.

Join the Fabric User Panel



Share your feedback directly with our
Fabric product group and researchers.

<https://aka.ms/JoinFabricUserPanel>

Join the SQL User Panel



Influence our SQL roadmap and ensure
it meets your real-life needs

<https://aka.ms/JoinSQLUserPanel>

How was the session?



Complete Session Surveys in
Whova for your chance to WIN
PRIZES!

