Microsoft | Microsoft Fabric


#FABCONSQLCON2026
FABCON
Microsoft Fabric
COMMUNITY CONFERENCE
SQLCON
Microsoft SQL
COMMUNITY CONFERENCE
ATLANTA MARCH 16 - 20, 2026

Can we add AI to this?

AI Functions

One-line transformations

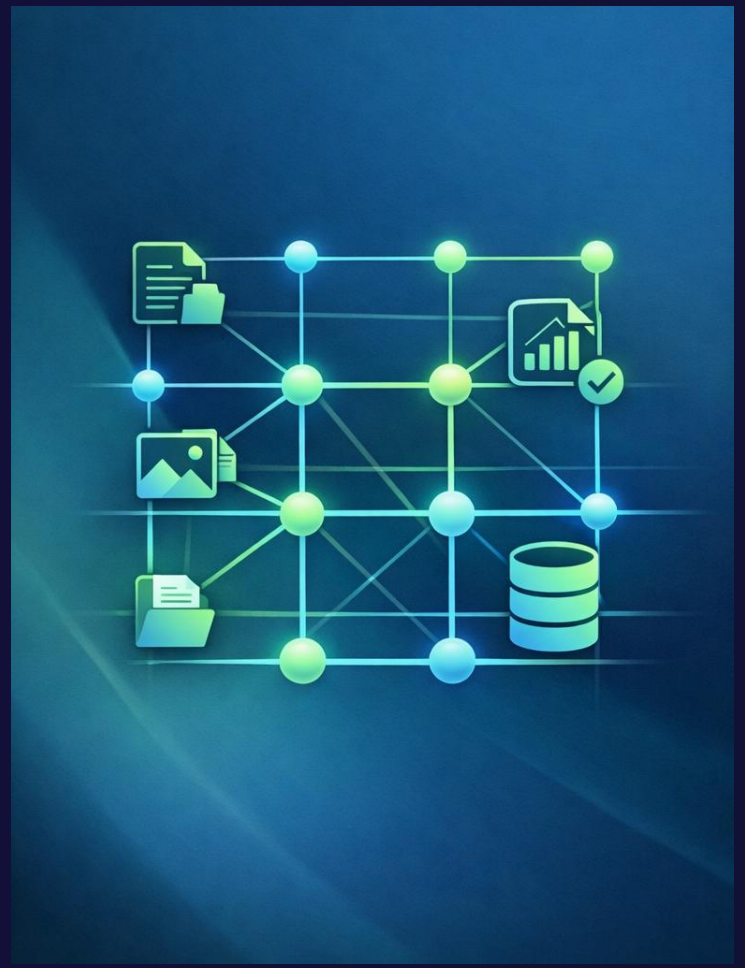

Powering Data with AI Functions in Fabric

Virginia Roman

Product Manager
AI Products
Microsoft Fabric

Rana Singh

Engineering Manager
AI Powered Data Engineering
Microsoft Fabric

Agenda

Overview
- AI Functions overview (pandas & PySpark)
- Business value & common use cases

Integrations across Fabric
- Dataflow Gen2
- Data Warehouse (SQL)
- Data Wrangler

Multimodal input support
- Images, PDFs, text files

Experience enhancements
- Improved cost transparency
- Eval & starter notebooks

Real world use case
- How a Fortune 500 saved 1,000+ hrs. with AI Functions

What are AI Functions?

1-Liners that perform high level intelligence tasks
by applying LLMs to your rows of data

text
It was the best of times, ….
Call me Ishmael….
All happy families are alike; …

```
df.ai.summarize()
```

summary
The novel explores the struggles…
A sailor recounts his voyage….
The novel delves into the complexities of love…

Specialized AI Function	Description
`ai.analyze_sentiment`	Identify emotions and intent
`ai.classify`	Classification into categories
`ai.extract`	Entity or substring extraction from data
`ai.fix_grammar`	Improve coherency, grammar, and spellings
`ai.similarity`	Compute cosine-similarity between text embeddings
`ai.summarize`	Customizable summary
`ai.translate`	Translation into any major language
General AI Function	**Description**
`ai.embed`	Generate vector embeddings for text
`ai.generate_response`	Do anything with custom prompt + structured output

Key Benefits

- ⚡ **Single line of code**
- ⚙️ **No API setup** — no keys, endpoints, or configuration
- 🔀 **High concurrency**: process multiple rows in parallel
- 🔒 **Row level isolation**: Context between rows is not mixed
- 🌐 **Flexible model choice**: not locked into one model

Use Cases

- 🧹 **Data cleanup** at scale
- 🔎 **Extracting intelligent insights** from large datasets
- 📄 **Complex** contract **analysis**
- 📊 Build **intelligent dashboards** from predictable and structured AI functions output
- 🧪 AI functions to **evaluate AI functions**: LLM as a judge

Impact

- 🚀 “We transformed a **1,000+ hour manual** and error-prone **process to only 1 week** and with 100% consistency”
- ⏳ “Before, this would take us **2 days**. Now at most it takes **10 minutes**”
- 🎯 “**ai.extract** makes **parsing really easy** and **allows our data analysts to own it** end to end”

AI Functions Demo: Text Transformation with Pandas

Generally Available

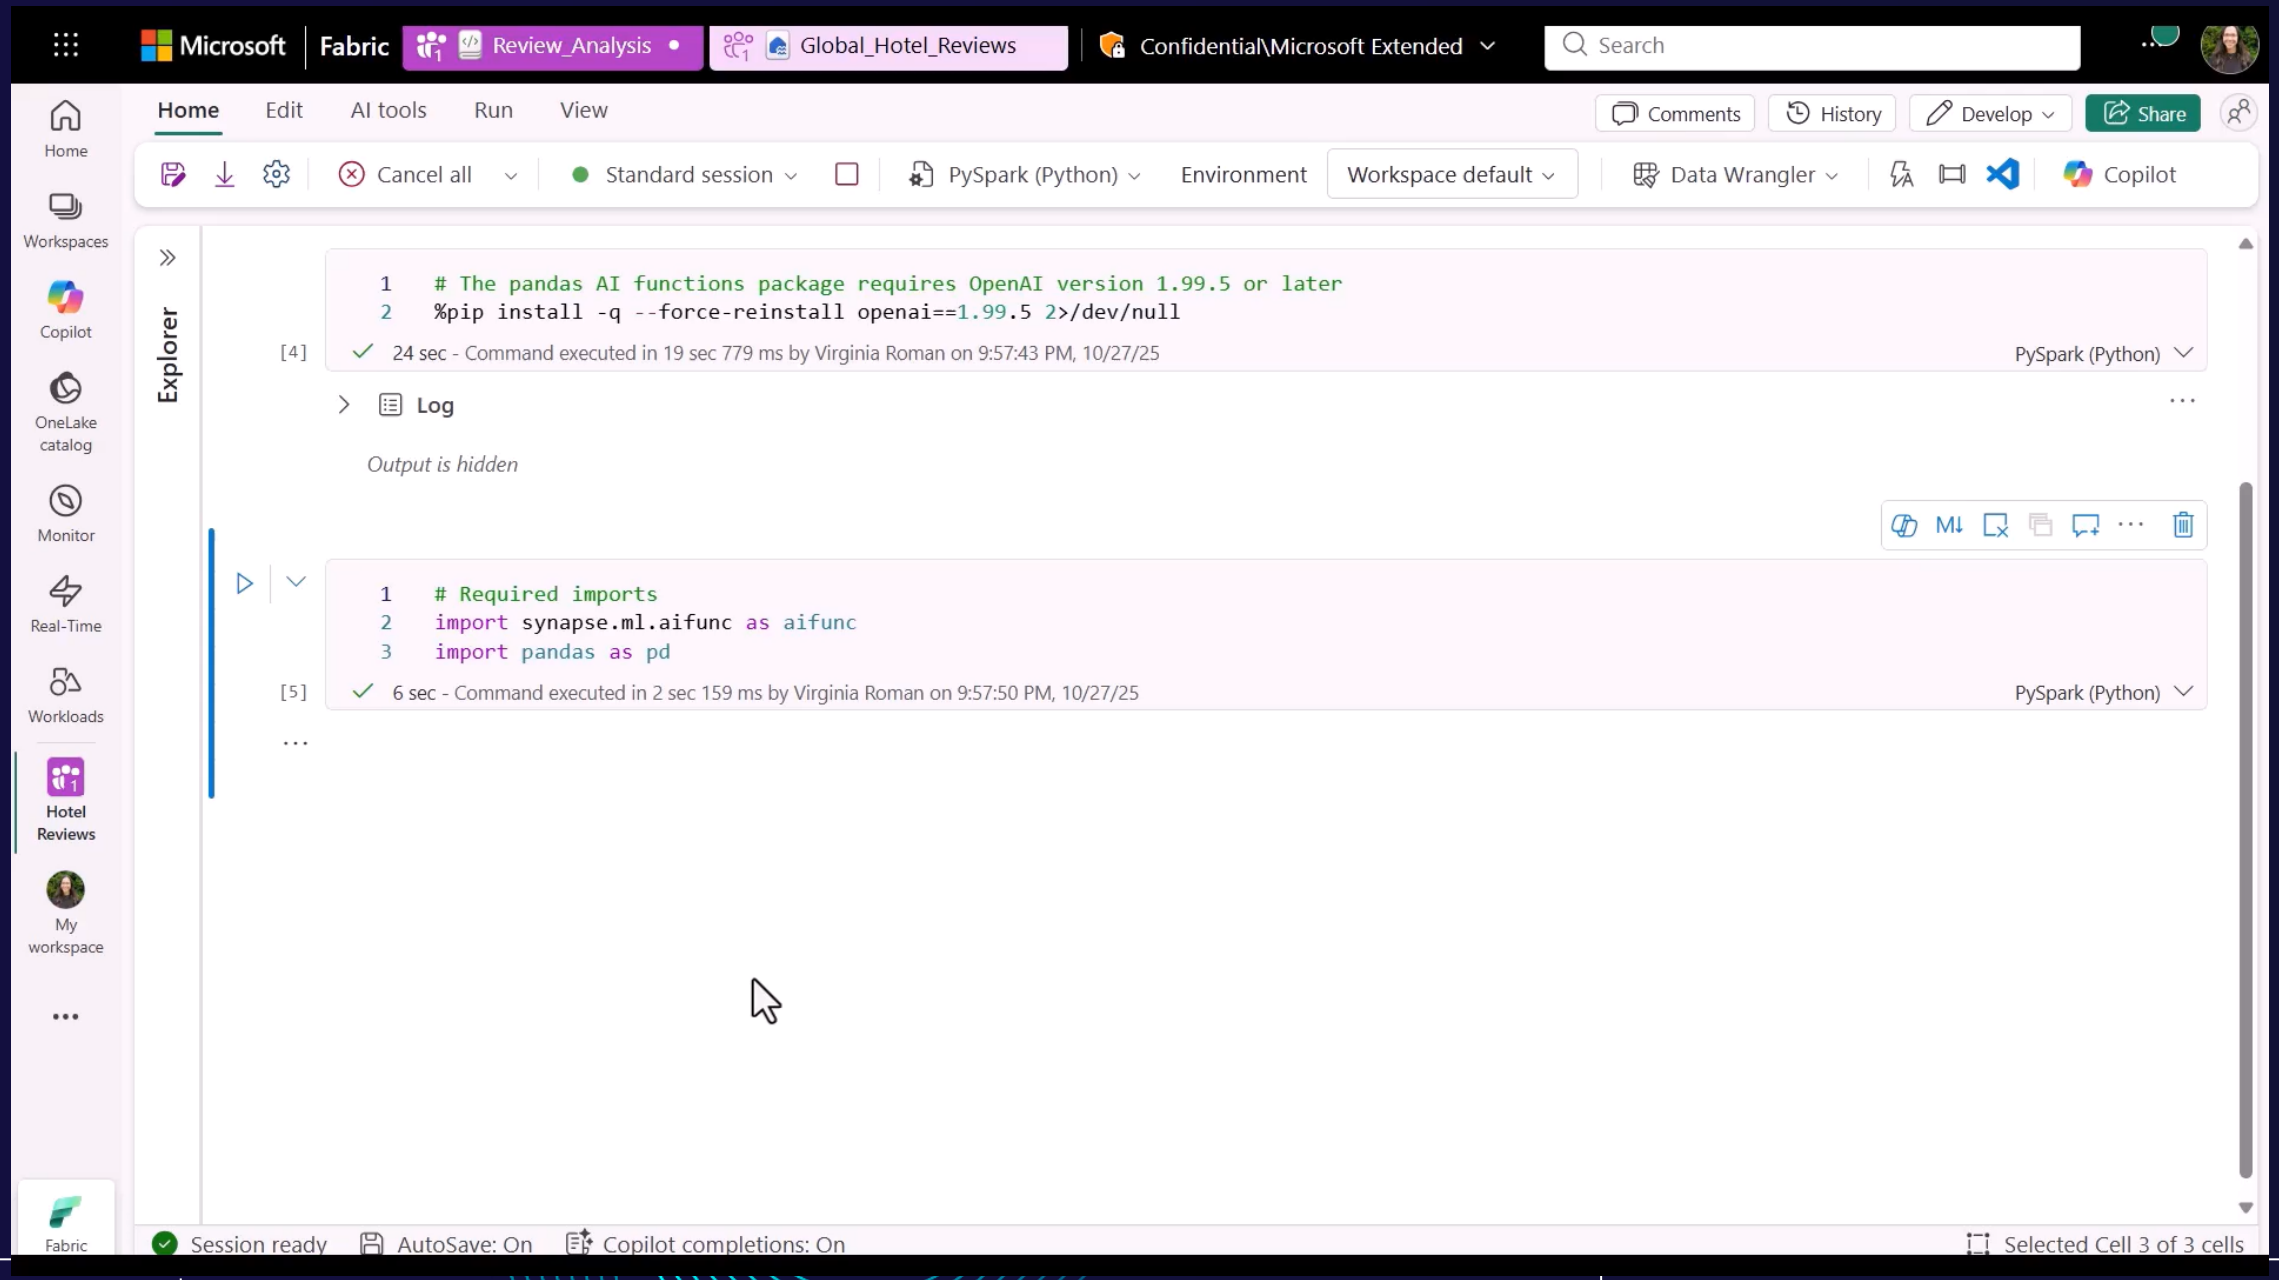

AI Functions Demo: Notebook UI

Generally Available

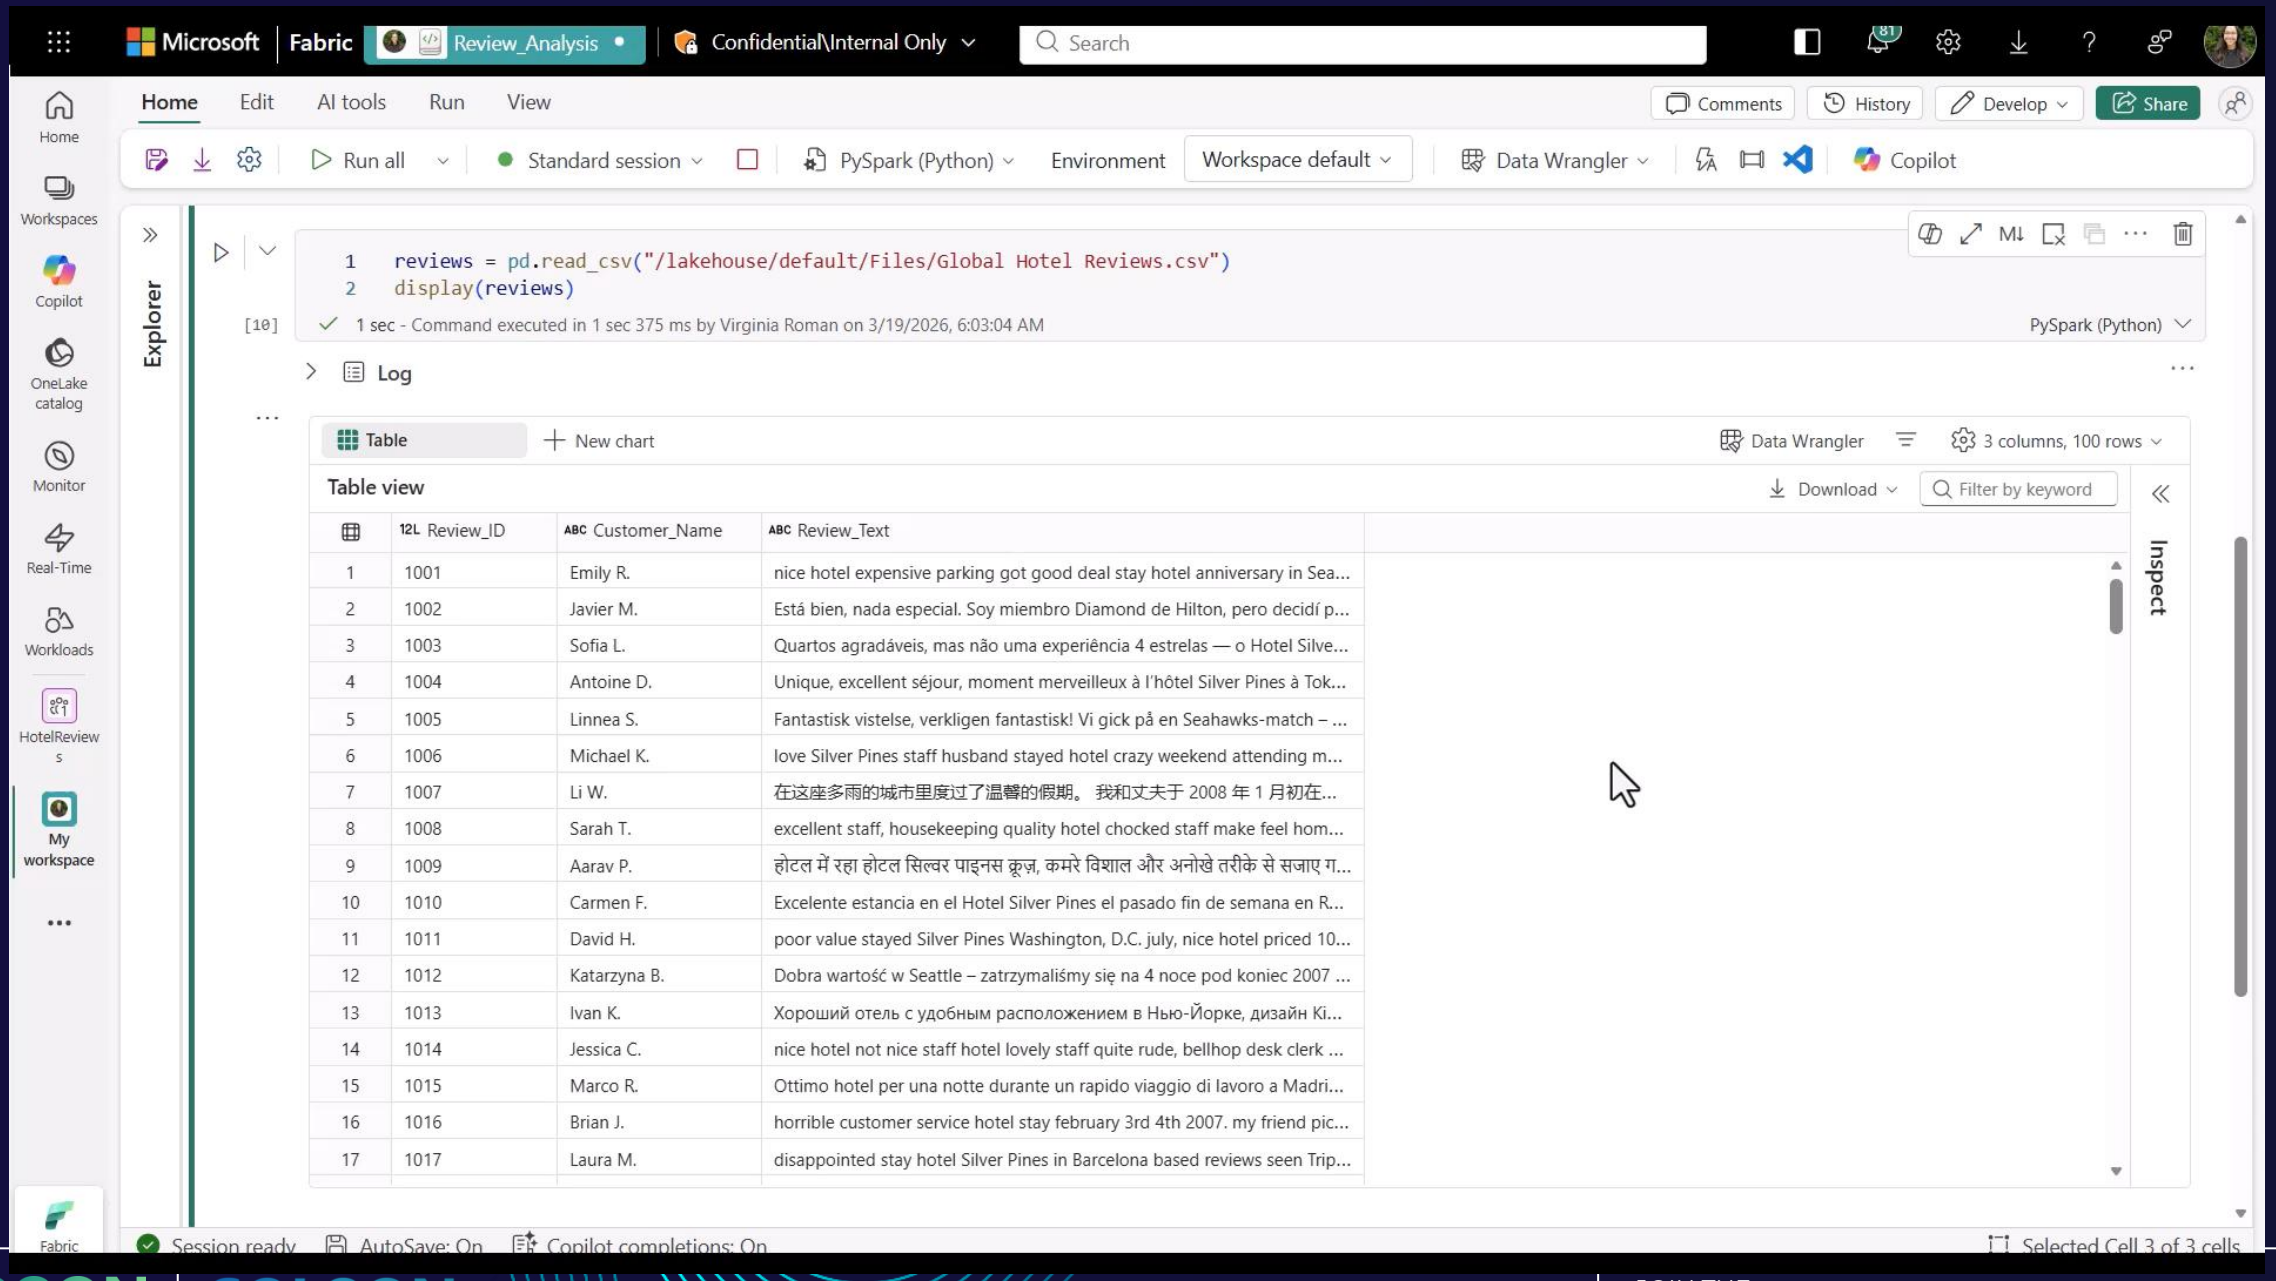

	Review_ID	Customer_Name	Review_Text
1	1001	Emily R.	nice hotel expensive parking got good deal stay hotel anniversary in Sea...
2	1002	Javier M.	Está bien, nada especial. Soy miembro Diamond de Hilton, pero decidí p...
3	1003	Sofia L.	Quartos agradáveis, mas não uma experiência 4 estrelas — o Hotel Silve...
4	1004	Antoine D.	Unique, excellent séjour, moment merveilleux à l'hôtel Silver Pines à Tok...
5	1005	Linnea S.	Fantastisk vistelse, verkligen fantastisk! Vi gick på en Seahawks-match – ...
6	1006	Michael K.	love Silver Pines staff husband stayed hotel crazy weekend attending m...
7	1007	Li W.	在这座多雨的城市里度过了温馨的假期。 我和丈夫于 2008 年 1 月初在...
8	1008	Sarah T.	excellent staff, housekeeping quality hotel chocked staff make feel hom...
9	1009	Aarav P.	होटल में रहा होटल सिल्वर पाइनस क्रूज़, कमरे विशाल और अनोखे तरीके से सजाए ग...
10	1010	Carmen F.	Excelente estancia en el Hotel Silver Pines el pasado fin de semana en R...
11	1011	David H.	poor value stayed Silver Pines Washington, D.C. july, nice hotel priced 10...
12	1012	Katarzyna B.	Dobra wartość w Seattle – zatrzymaliśmy się na 4 noce pod koniec 2007 ...
13	1013	Ivan K.	Хороший отель с удобным расположением в Нью-Йорке, дизайн Ki...
14	1014	Jessica C.	nice hotel not nice staff hotel lovely staff quite rude, bellhop desk clerk ...
15	1015	Marco R.	Ottimo hotel per una notte durante un rapido viaggio di lavoro a Madri...
16	1016	Brian J.	horrible customer service hotel stay february 3rd 4th 2007. my friend pic...
17	1017	Laura M.	disappointed stay hotel Silver Pines in Barcelona based reviews seen Trip...

AI Functions in Data Wrangler

Public Preview

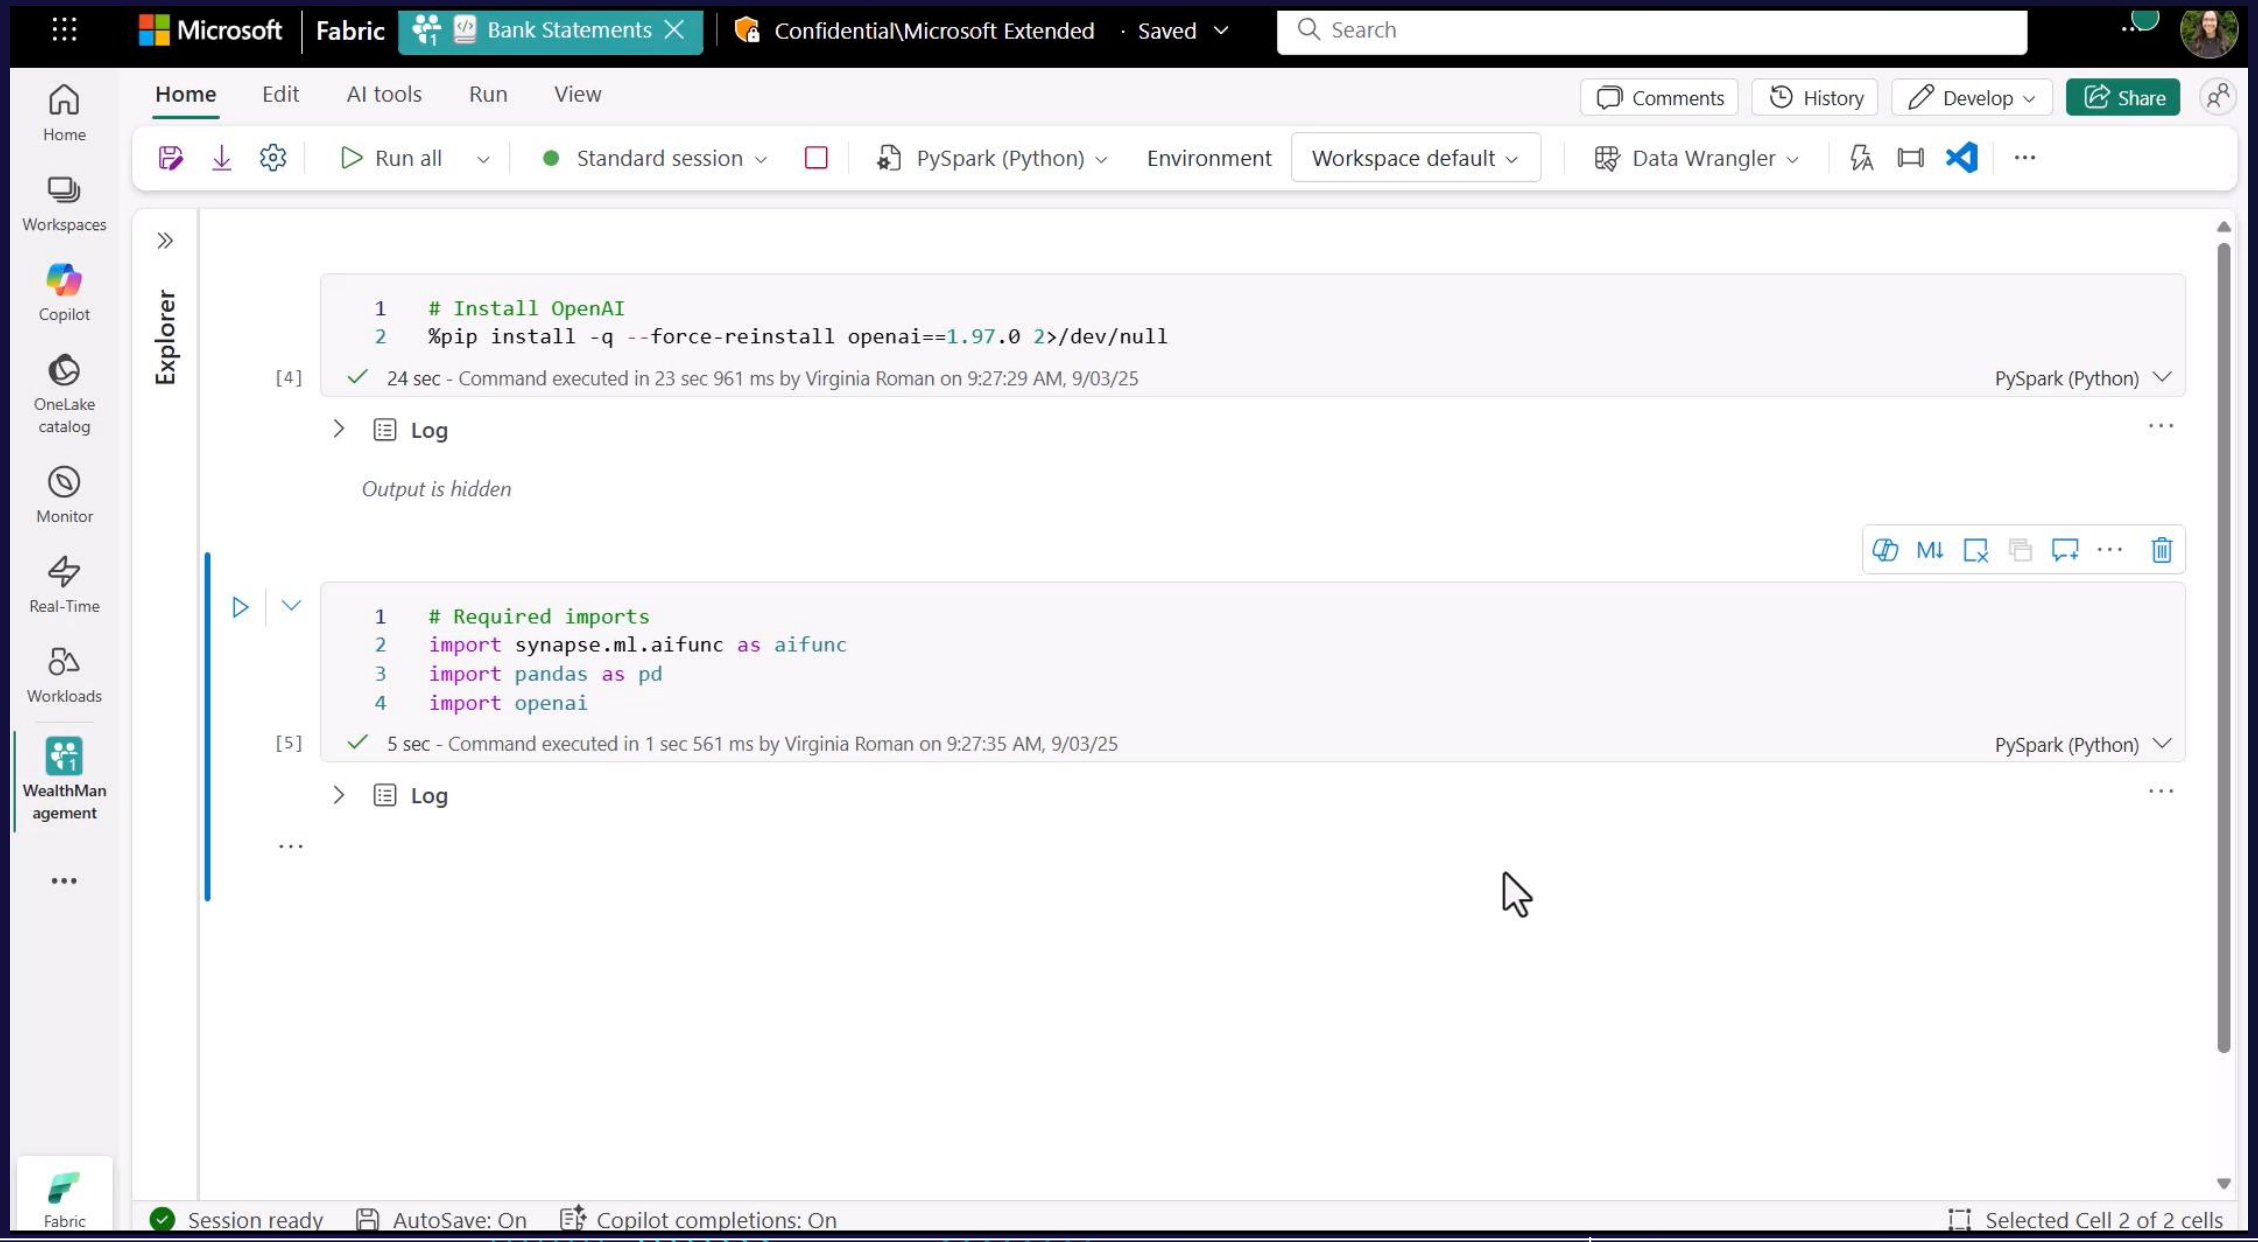

Broader Fabric Integration

AI Functions in Data Warehouse SQL

Public Preview

AI Functions in Dataflow Gen2

Generally Available

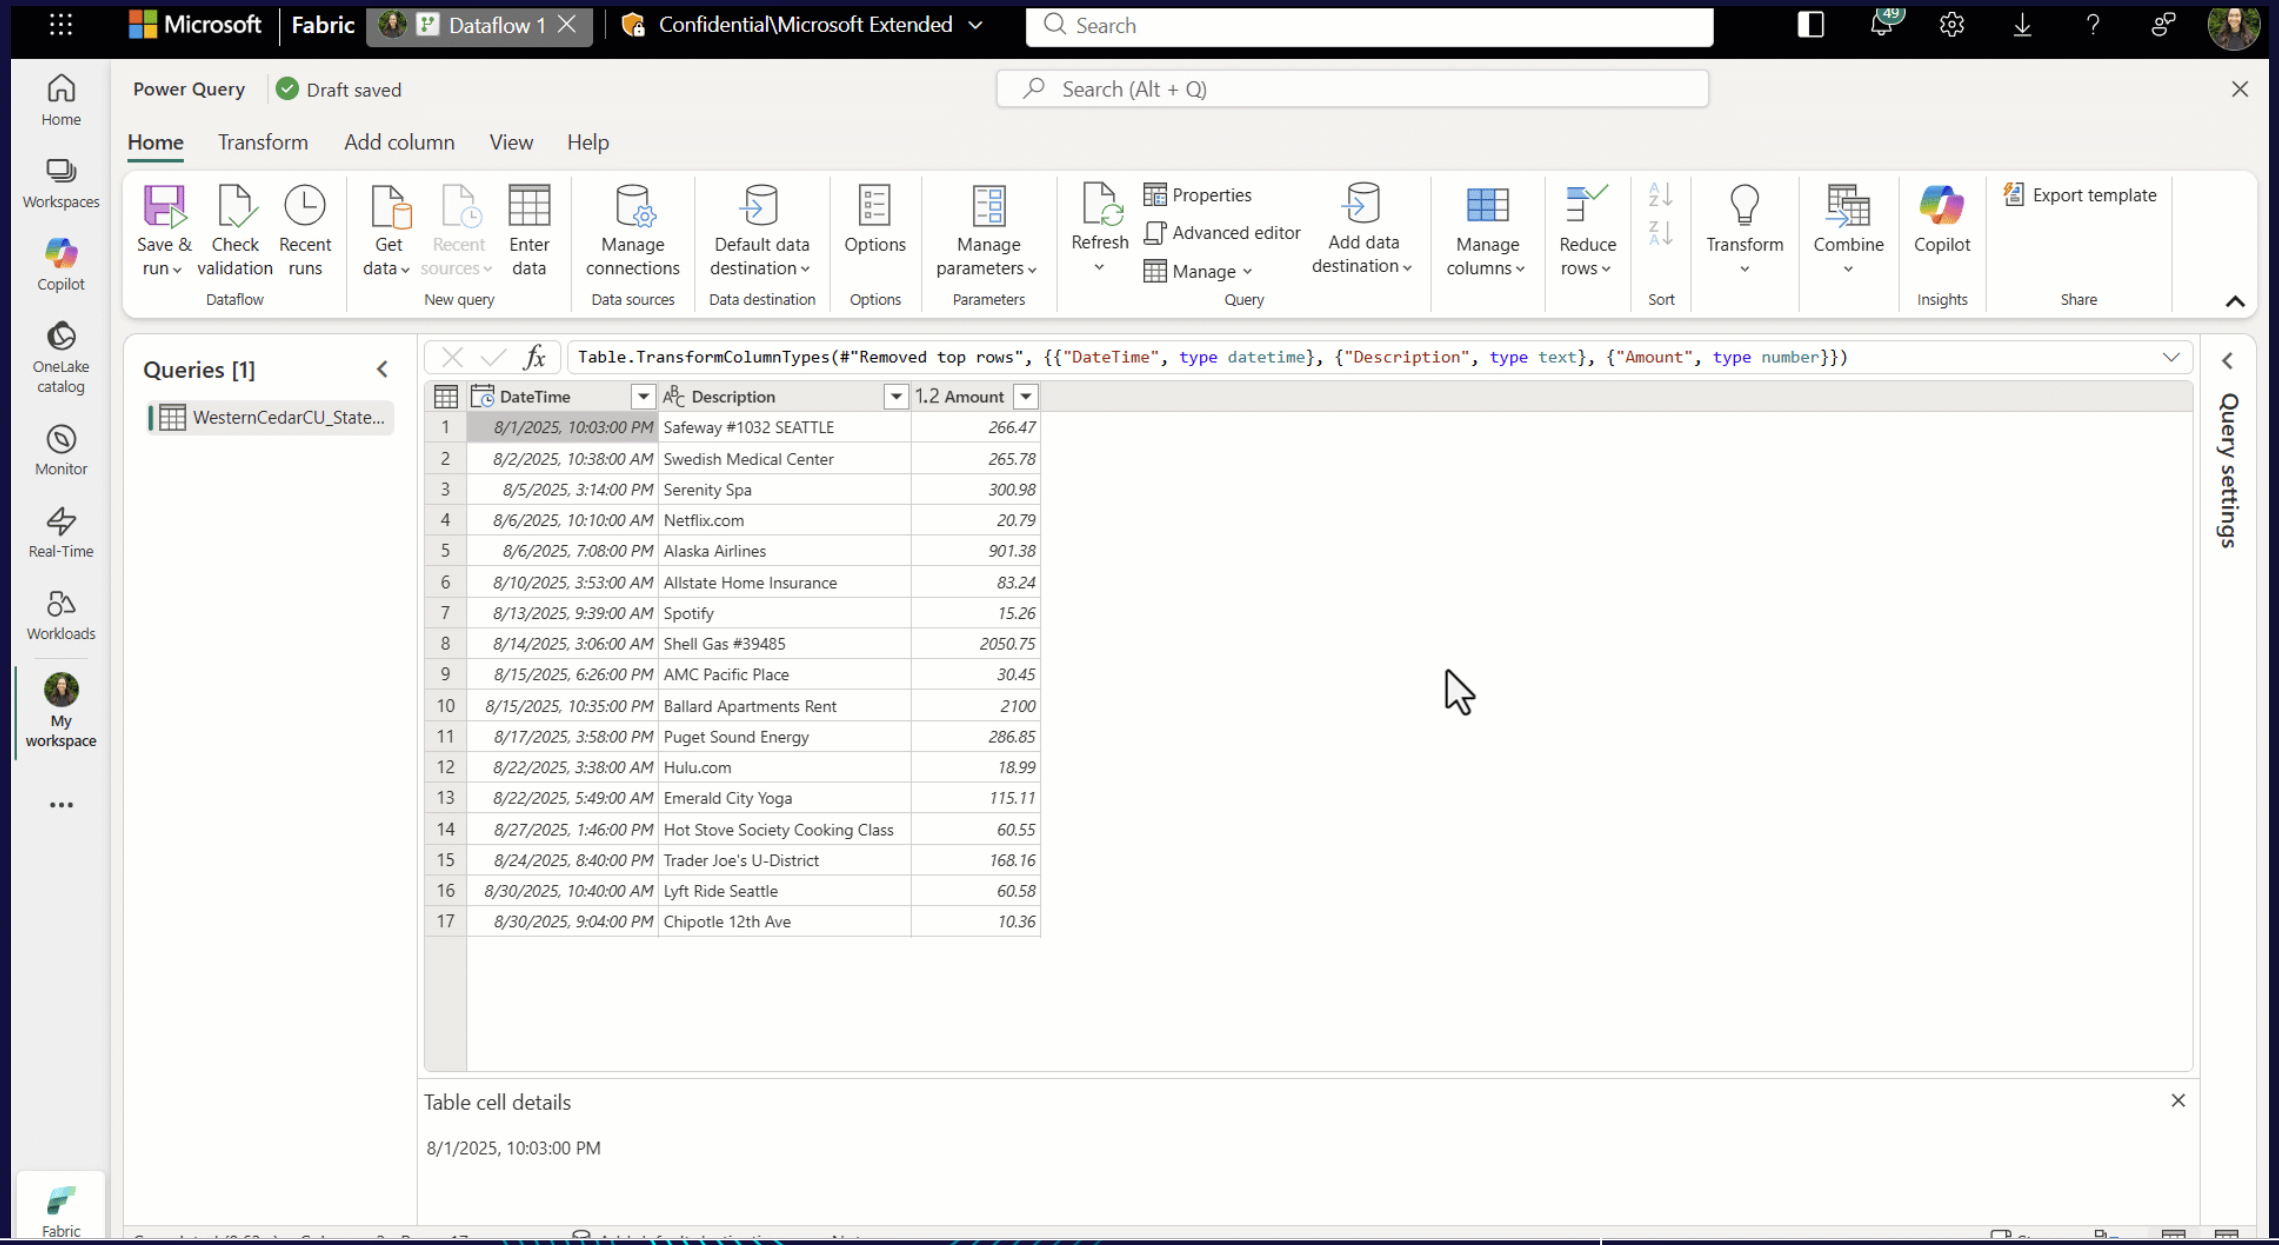

	DateTime	Description	1.2 Amount
1	8/1/2025, 10:03:00 PM	Safeway #1032 SEATTLE	266.47
2	8/2/2025, 10:38:00 AM	Swedish Medical Center	265.78
3	8/5/2025, 3:14:00 PM	Serenity Spa	300.98
4	8/6/2025, 10:10:00 AM	Netflix.com	20.79
5	8/6/2025, 7:08:00 PM	Alaska Airlines	901.38
6	8/10/2025, 3:53:00 AM	Allstate Home Insurance	83.24
7	8/13/2025, 9:39:00 AM	Spotify	15.26
8	8/14/2025, 3:06:00 AM	Shell Gas #39485	2050.75
9	8/15/2025, 6:26:00 PM	AMC Pacific Place	30.45
10	8/15/2025, 10:35:00 PM	Ballard Apartments Rent	2100
11	8/17/2025, 3:58:00 PM	Puget Sound Energy	286.85
12	8/22/2025, 3:38:00 AM	Hulu.com	18.99
13	8/22/2025, 5:49:00 AM	Emerald City Yoga	115.11
14	8/27/2025, 1:46:00 PM	Hot Stove Society Cooking Class	60.55
15	8/24/2025, 8:40:00 PM	Trader Joe's U-District	168.16
16	8/30/2025, 10:40:00 AM	Lyft Ride Seattle	60.58
17	8/30/2025, 9:04:00 PM	Chipotle 12th Ave	10.36

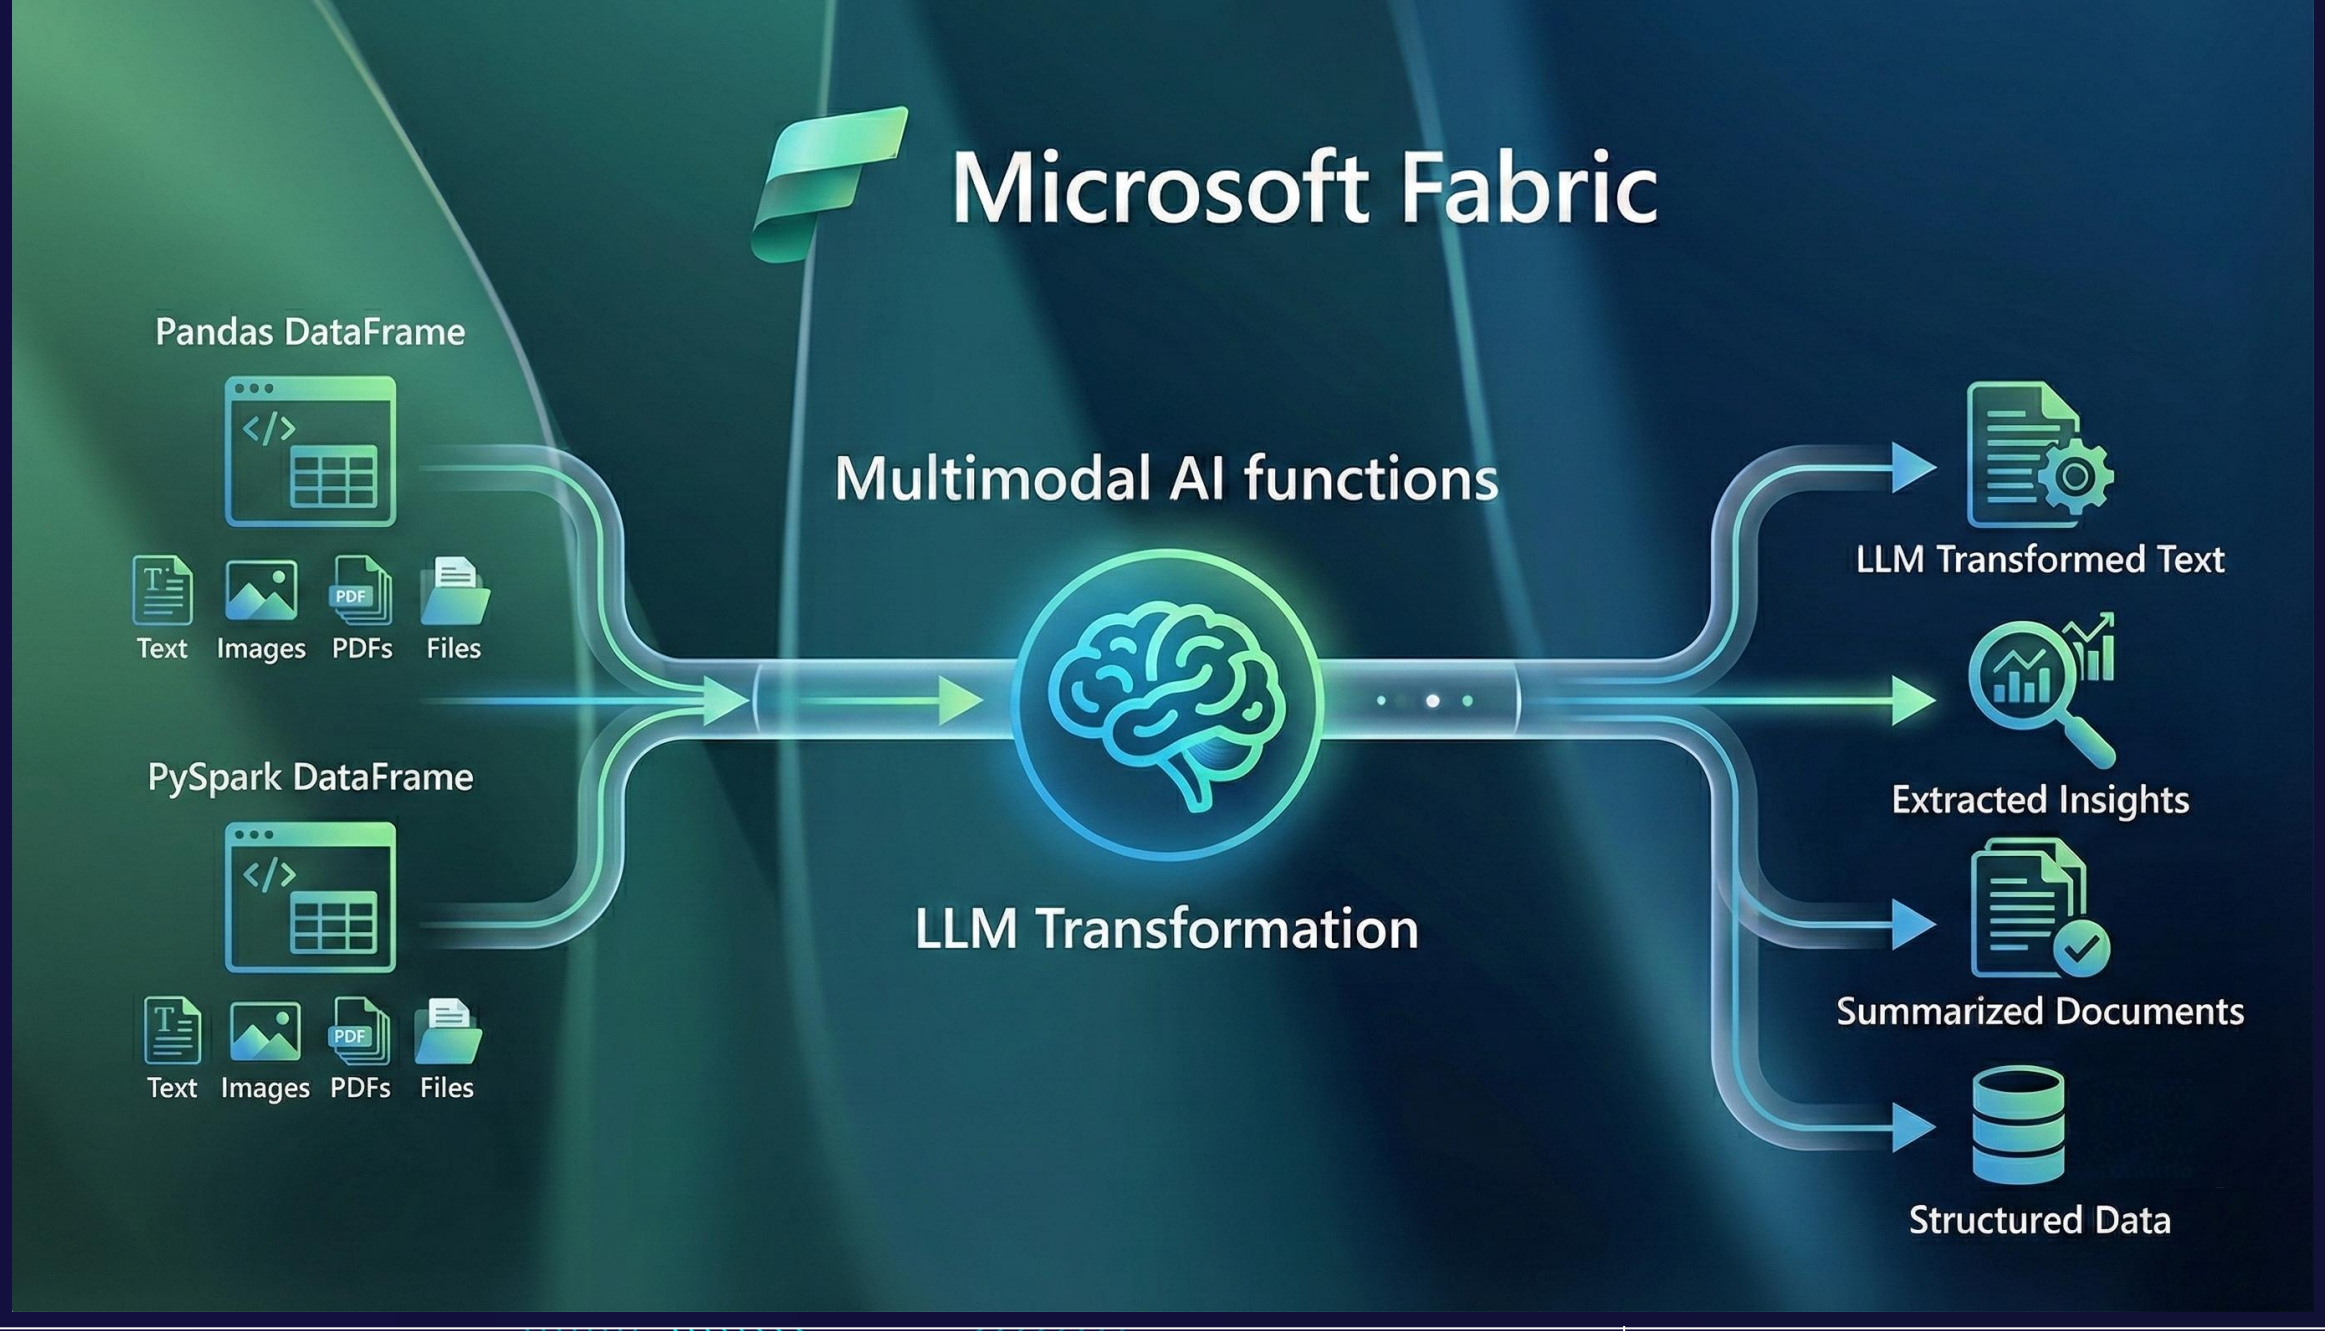
Microsoft Fabric
Pandas DataFrame
Text
Images
PDFs
Files
PySpark DataFrame
Text
Images
PDFs
Files
Multimodal AI functions
LLM Transformation
LLM Transformed Text
Extracted Insights
Summarized Documents
Structured Data

AI Functions Demo: Multimodal Files Transformation

Public Preview

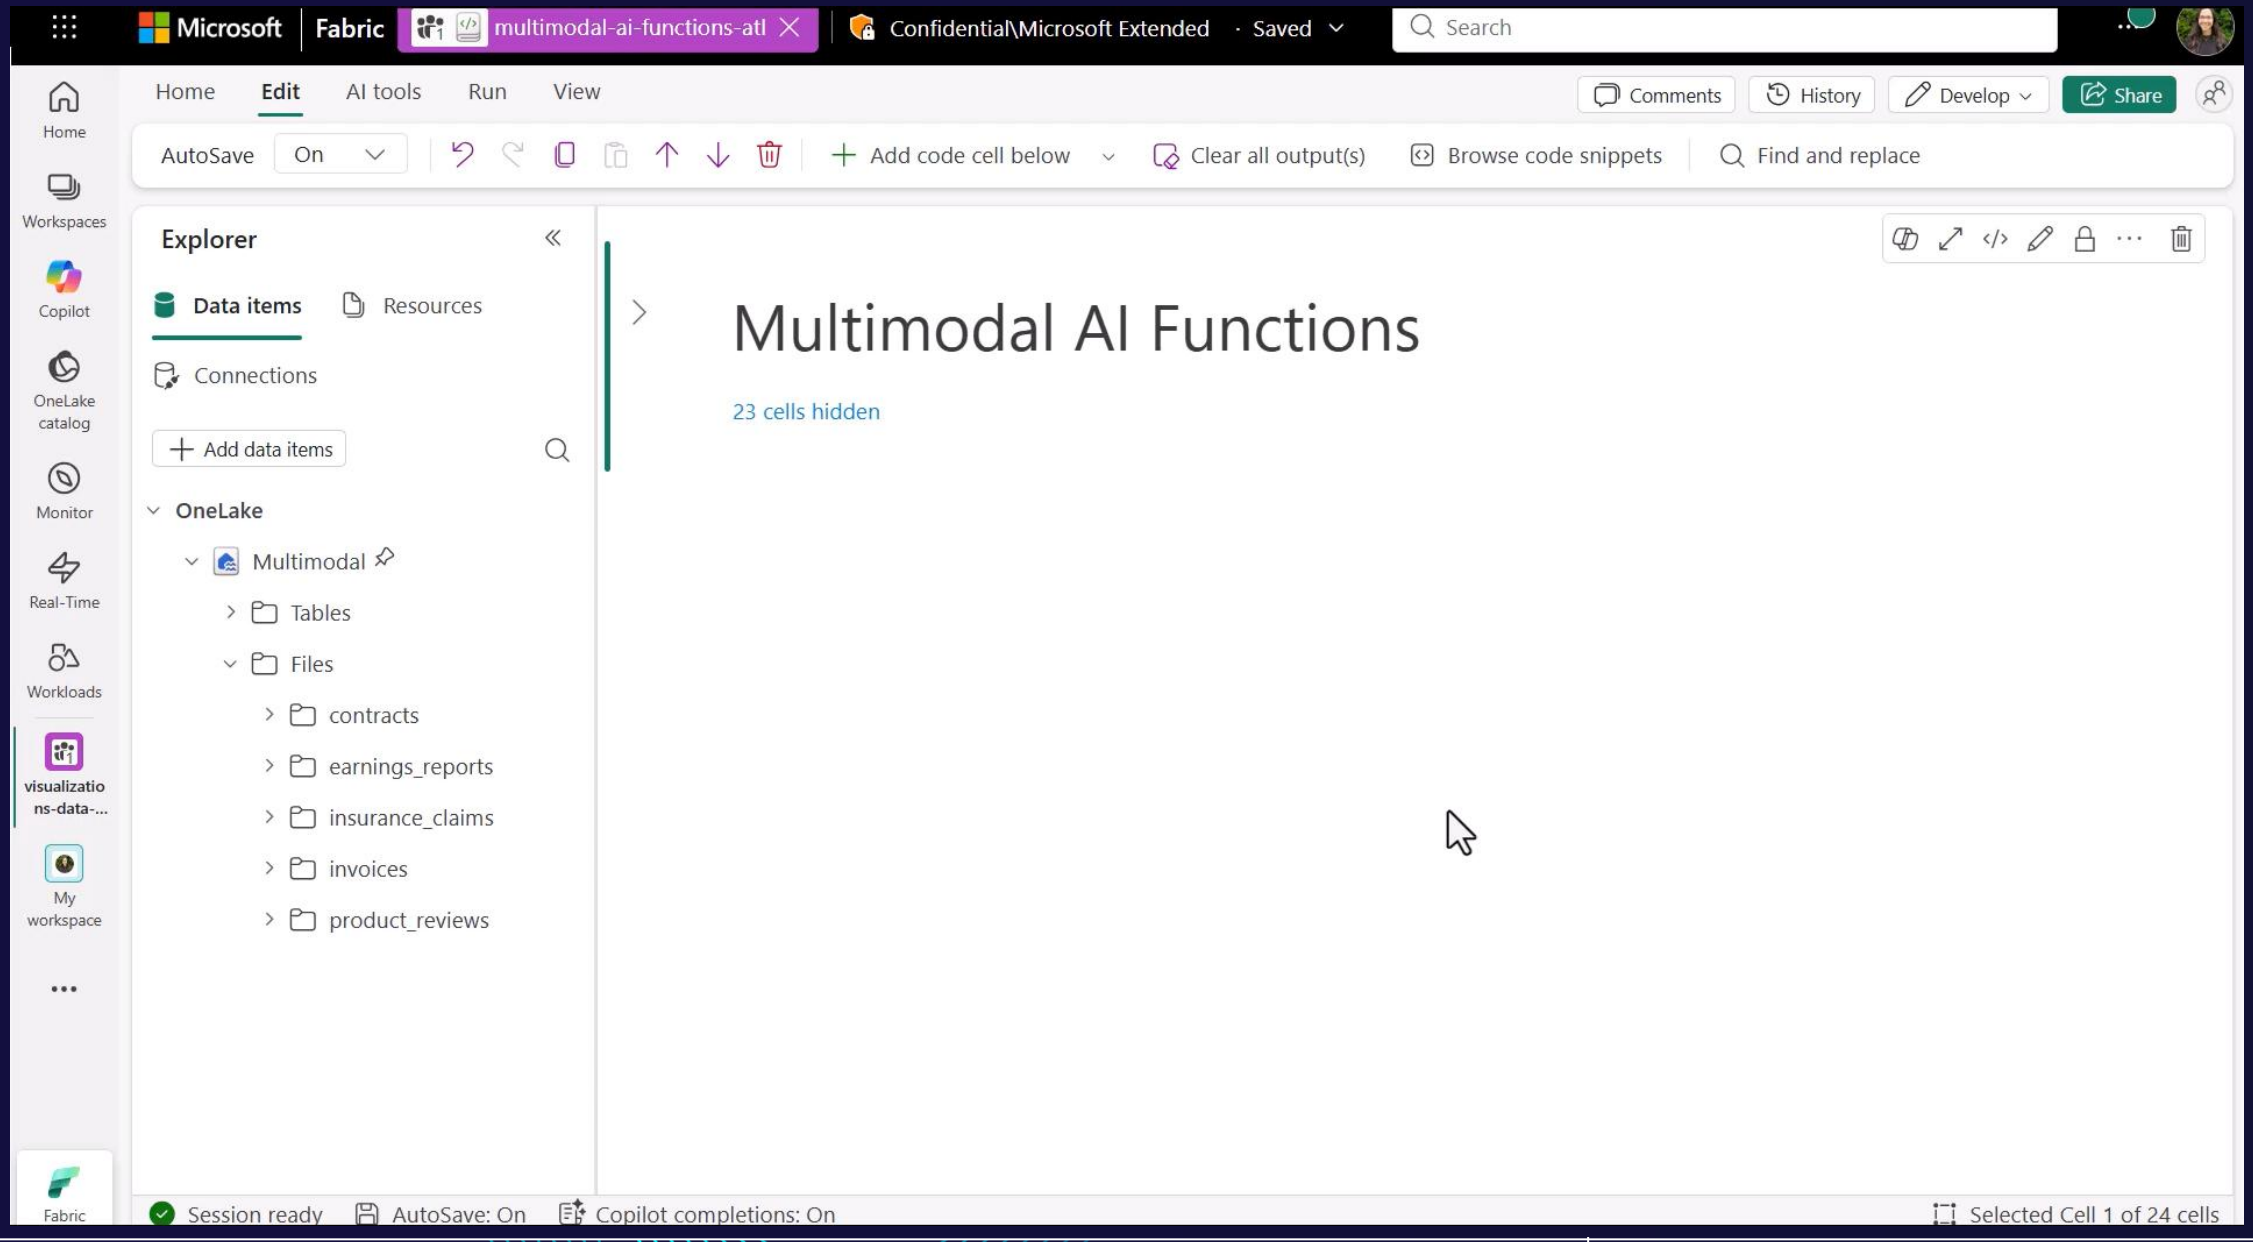

Experience Enhancements

Knobs & Levers

Parameter	Description	Default
concurrency	Number of parallel LLM calls	200
progress_bar_mode	Details in pandas TQDM progress bar	basic
Model Parameters	**Description**	
embedding_deployment_name	Generate vector embeddings for text	text-embedding-ada-002
model_deployment_name	Do anything with custom prompt + structured output	gpt-4.1-mini
reasoning_effort	Reasoning effort to be used	None
temperature	Randomness of output	1.0
top_p	Restruct model token probability space	None
verbosity	Controls output length	None

Configure LLM

- Change Models: Bigger LLMs generally perform better in diverse scenarios

```
aifunc.default_conf.model_deployment_name = "gpt-5"
```

Microsoft Fabric

Fabric LLMs for AI Functions

- gpt-4.1-mini (default)
- gpt-4.1
- gpt-5

Upcoming Fabric LLMs options

- gpt-5-mini (new default end of April 2026)
- gpt-5.1 and beyond

- BYO: Bring Your LLM deployments compatible with `responses` or `chat_completions` API

Microsoft Foundry

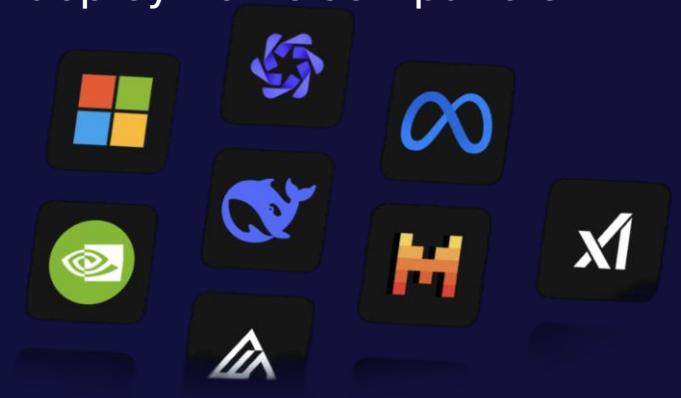

Cost Transparency

- Real Time Token Usage and CU Prediction in pandas tqdm Progress Bar:

  `aifunc.default_conf.progress_bar_mode = “stats”`

Cost Transparency

- Real Time Token Usage and CU Prediction in pandas tqdm Progress Bar:

  `aifunc.default_conf.progress_bar_mode = "stats"`

- Post AI Function detailed stats with **`ai.stats`**

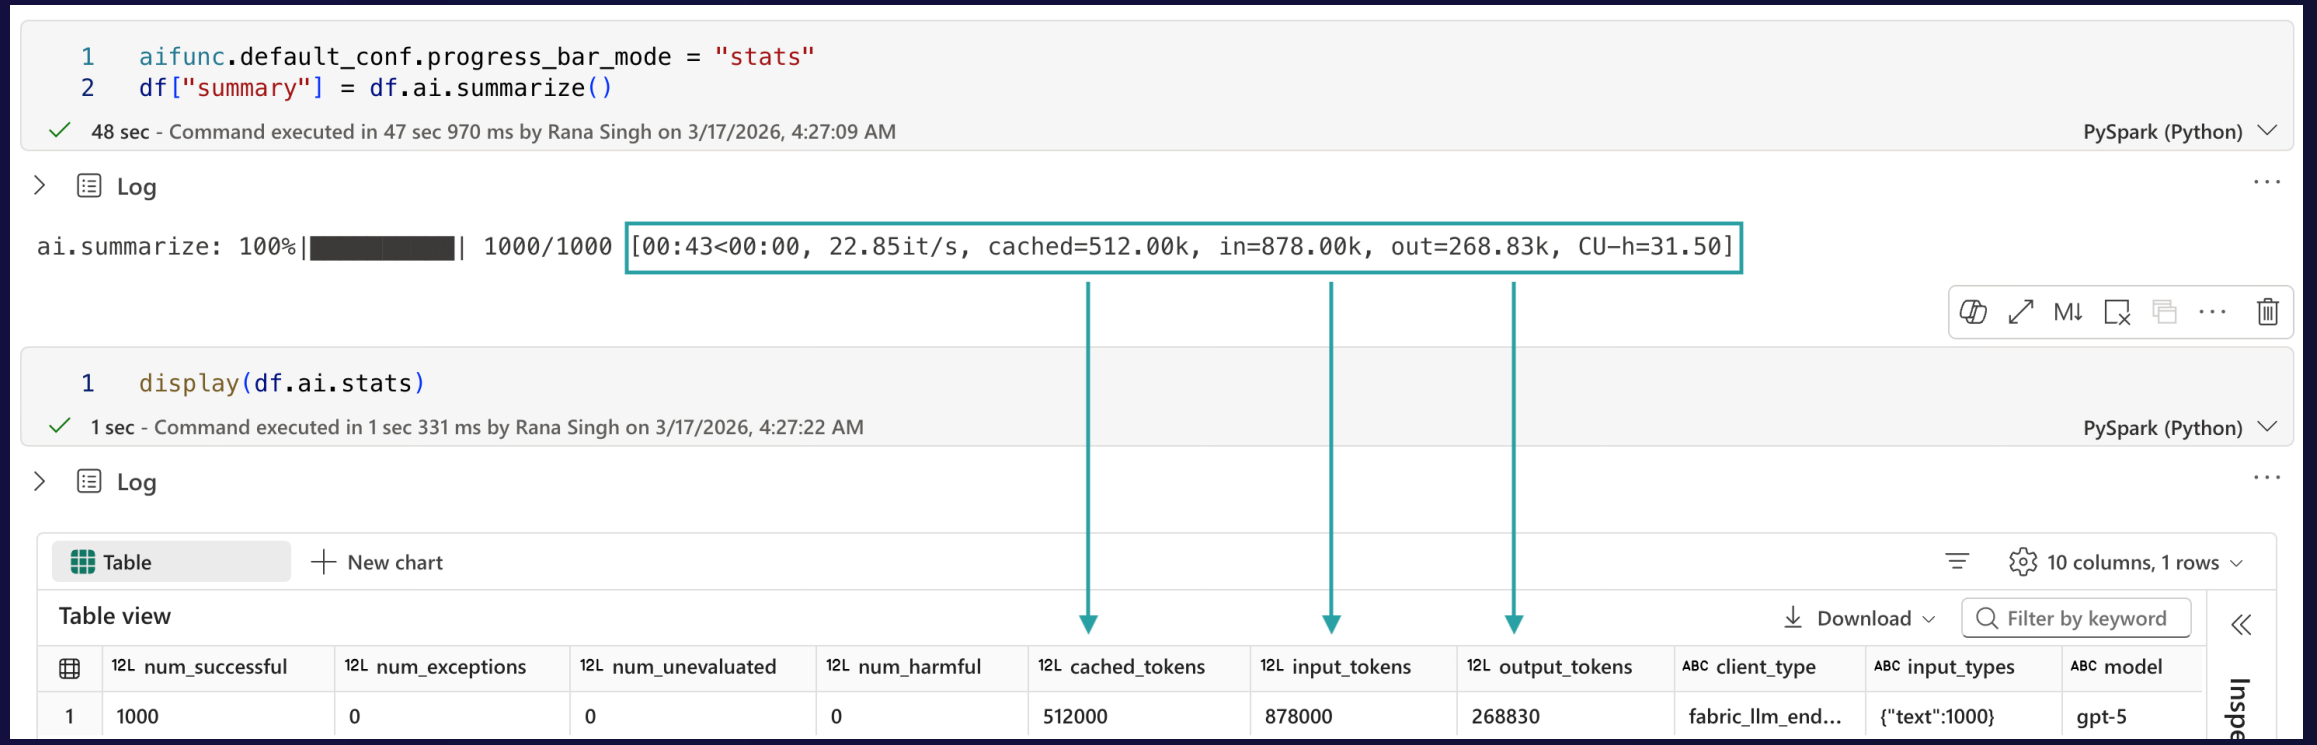

- Coming soon: tiktoken based input token and cost calculator

Cost Transparency: Progress Bar Modes

Visual Comparison

Set `aifunc.default_conf.progress_bar_mode` to one of the following:

`"basic"` (default) — Clean progress tracking

```
ai.extract: 100%|████████| 1000/1000 [00:10<00:00, 22.1it/s]
```

`"stats"` — Full metrics with token counts and CU prediction estimates

```
ai.extract:  67%|██████▒▒▒▒| 670/1000 [00:06<00:04, 22.1it/s, cached=150.03k->250.50k, in=800.82k->1.21M, out=68.65K->101.06k, CU-h=3.60->5.33]
                                                                ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
                                                                        token stats + CU prediction estimates

ai.extract: 100%|████████| 1000/1000 [00:10<00:00, 22.1it/s, cached=305.80k, in=1.21M, out=101.06k, CU-h=5.17]
                                                              ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
                                                                    final values (no arrows)
```

`"disable"` — No output

```
(nothing displayed)
```

Cost Transparency: New Operation Name

Before

- AI Functions usage was bundled into other operations such as Spark or Data Warehouse CU
- Hard to isolate or track AI-specific costs

Change

- New dedicated "`AI Functions`" operation name
- Usage tracked separately in billing in Microsoft Fabric Capacity Metrics app

Impact

- Clear visibility into AI Functions usage
- Easier cost management & reporting
- Understand usage patterns of AI Functions

Eval Notebooks

Dogfood AI Functions in LLM-as-judge workflows to evaluate performance of specialized AI Functions on your data + tips to make AI Functions work better for your use case

aka.ms/fabric-aifunctions-eval-notebooks

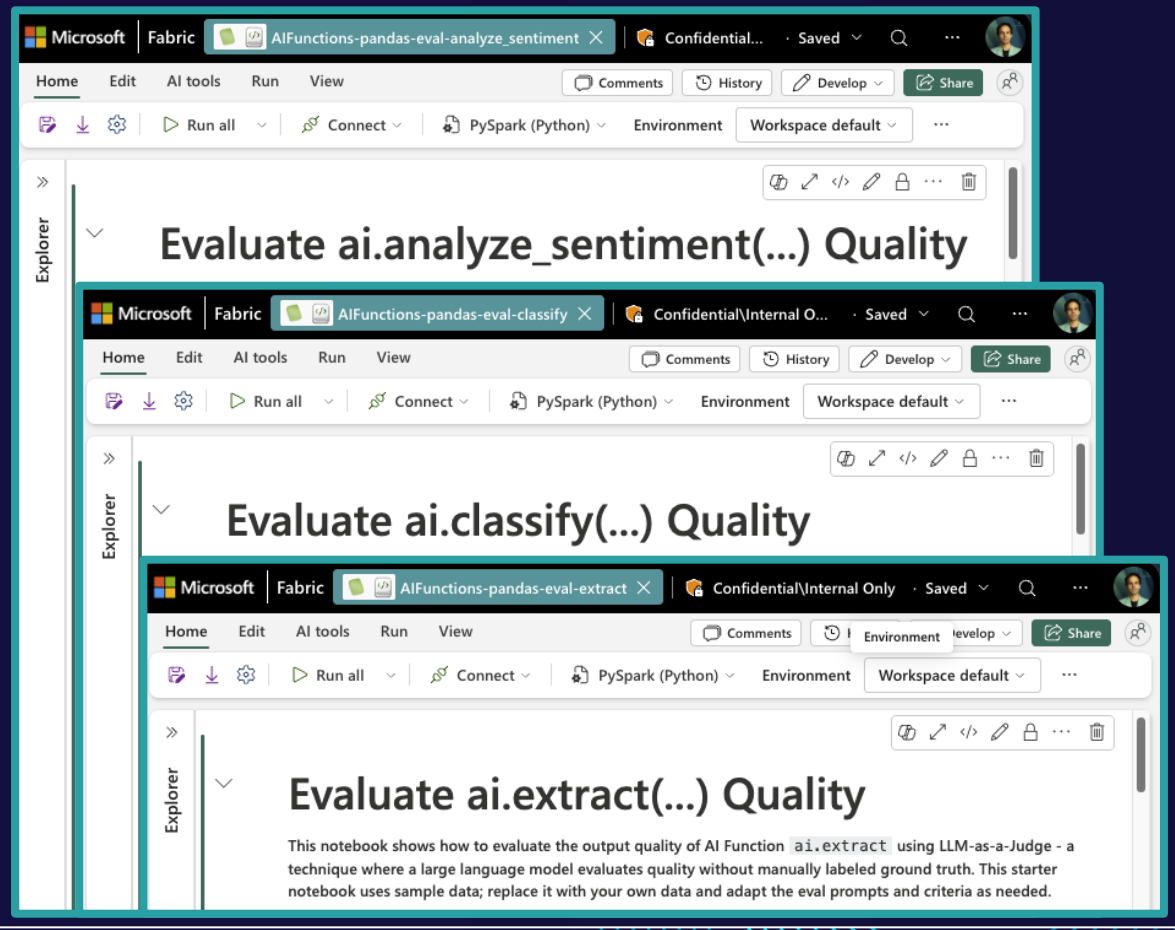

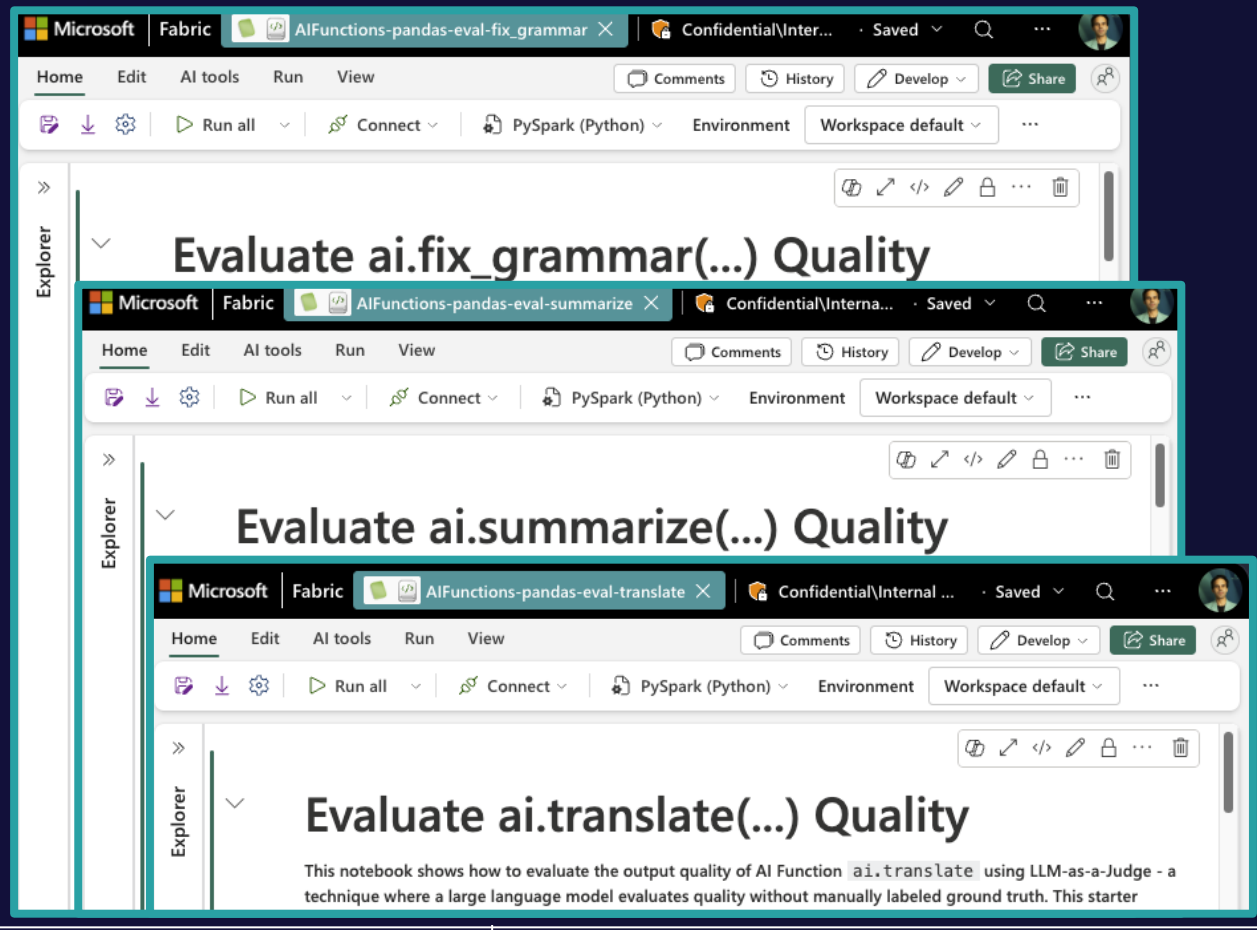

Starter Notebooks

Ready to play AI Functions powered data analytics workflows in pandas or PySpark

aka.ms/fabric-aifunctions-starter-notebooks

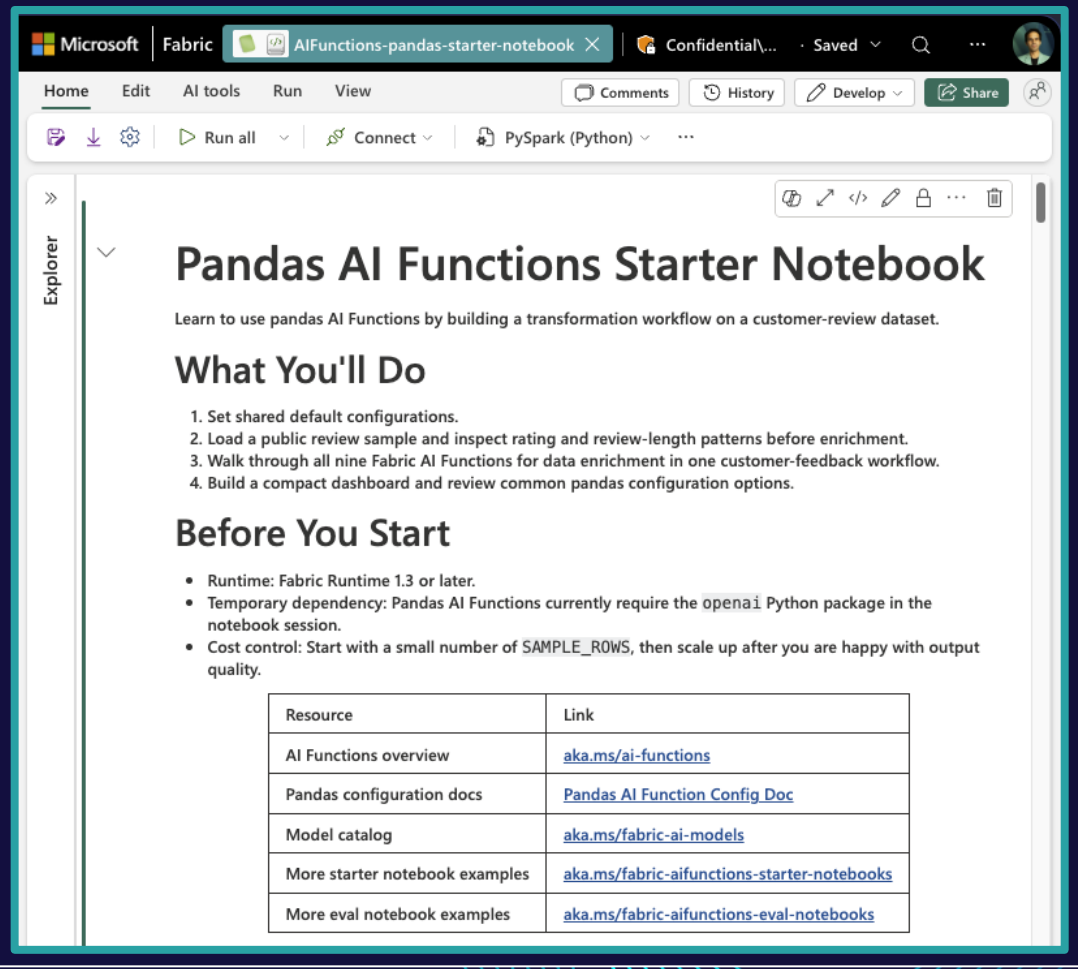

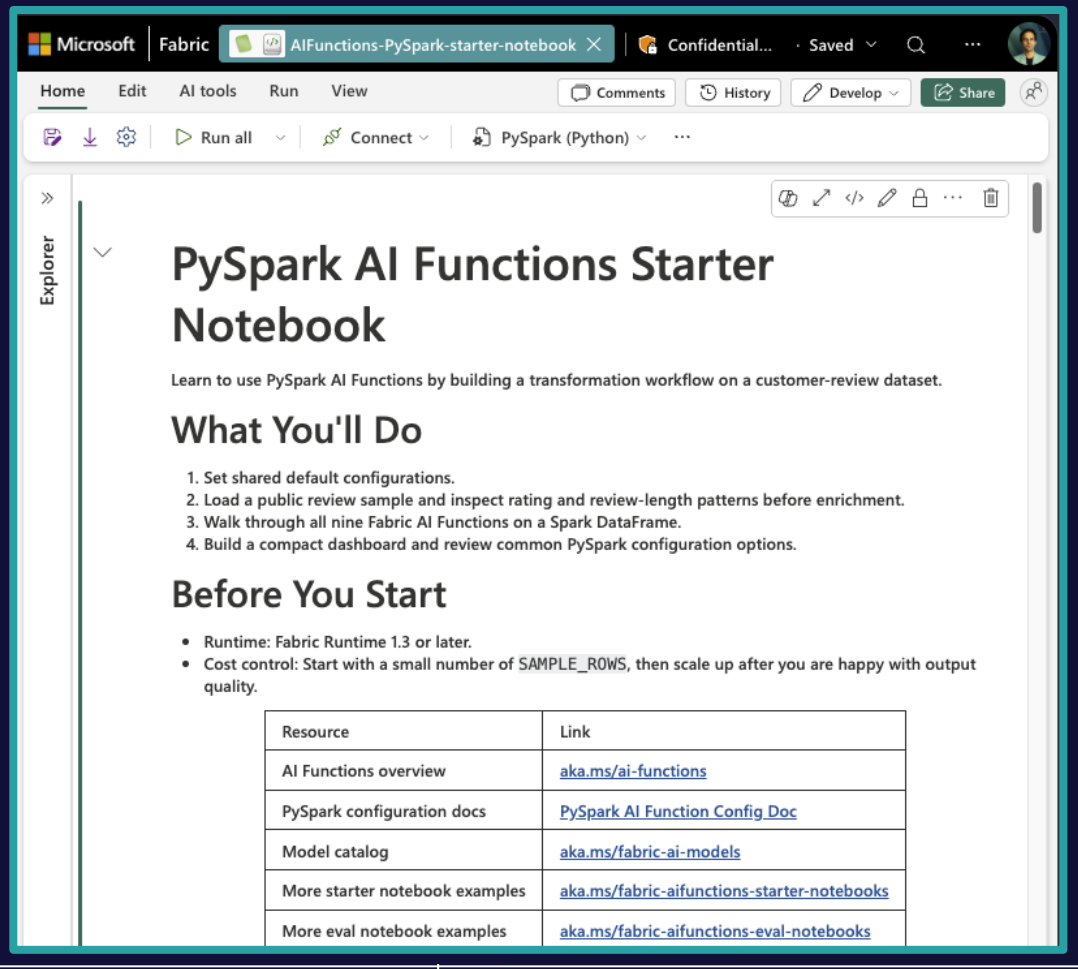

How a Fortune 500 Saved 1,000+ hrs. with AI Functions

How a Fortune 500 Automated Reporting with AI Functions

Problem

- 📊 **1,000+ active initiatives** across the portfolio
- ⏱️ Monthly status reporting was **manual and time-consuming**
- 🔎 Leaders lacked clear, consistent, up-to-date **business value visibility**
- 🚦 Traditional reports focused on **status, not outcomes**

AI Functions Usage

- 🤖 Used **ai.generate_response()** on JIRA portfolio data
- 🧠 **Analyzed text fields** across projects within each initiative
- 📝 Generated **executive-ready value summaries** automatically
- 🎯 Tuned prompts to reflect **company-specific language and criteria**

Impact

- ✅ **~75%** of monthly status reporting **automated**
- ⏳ **Thousands of hours** saved across initiative managers
- 📈 **Consistent, scalable** visibility into portfolio value
- 💡 Leaders can prioritize based on **business impact**, not manual updates

Resources

- Documentation: aka.ms/ai-functions
- AI Functions Starter Notebooks: aka.ms/fabric-aifunctions-starter-notebooks
- AI Functions Eval Notebooks: aka.ms/fabric-aifunctions-eval-notebooks
- Suggest ideas on the Fabric Ideas forum: ideas.fabric.microsoft.com

Q&A

How was the session?

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

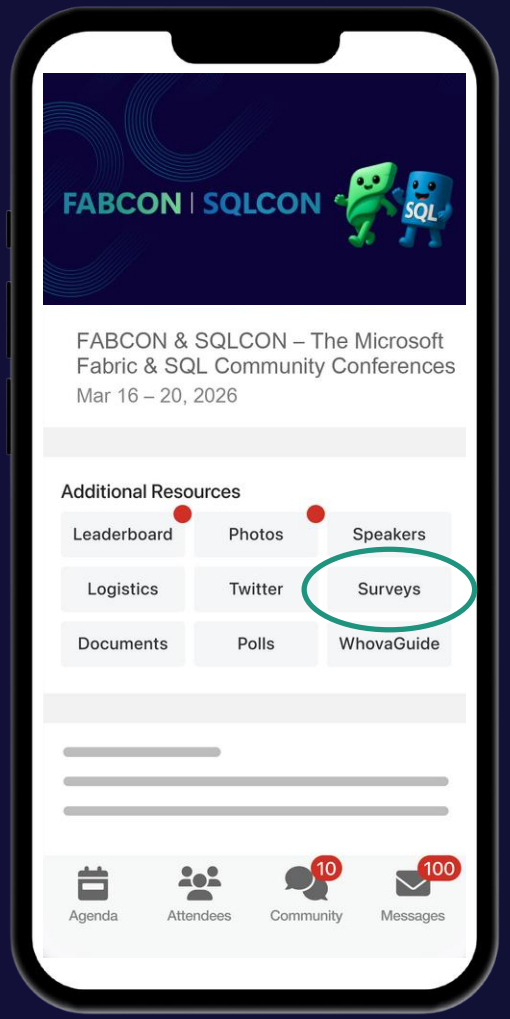

Sound off.
The mic is all yours.
Influence the product roadmap.

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

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