

Microsoft AI landscape


Microsoft AI Ecosystem

Unofficial conceptual view of AI services and resources used to incorporate AI technology into organizations to achieve more with less.

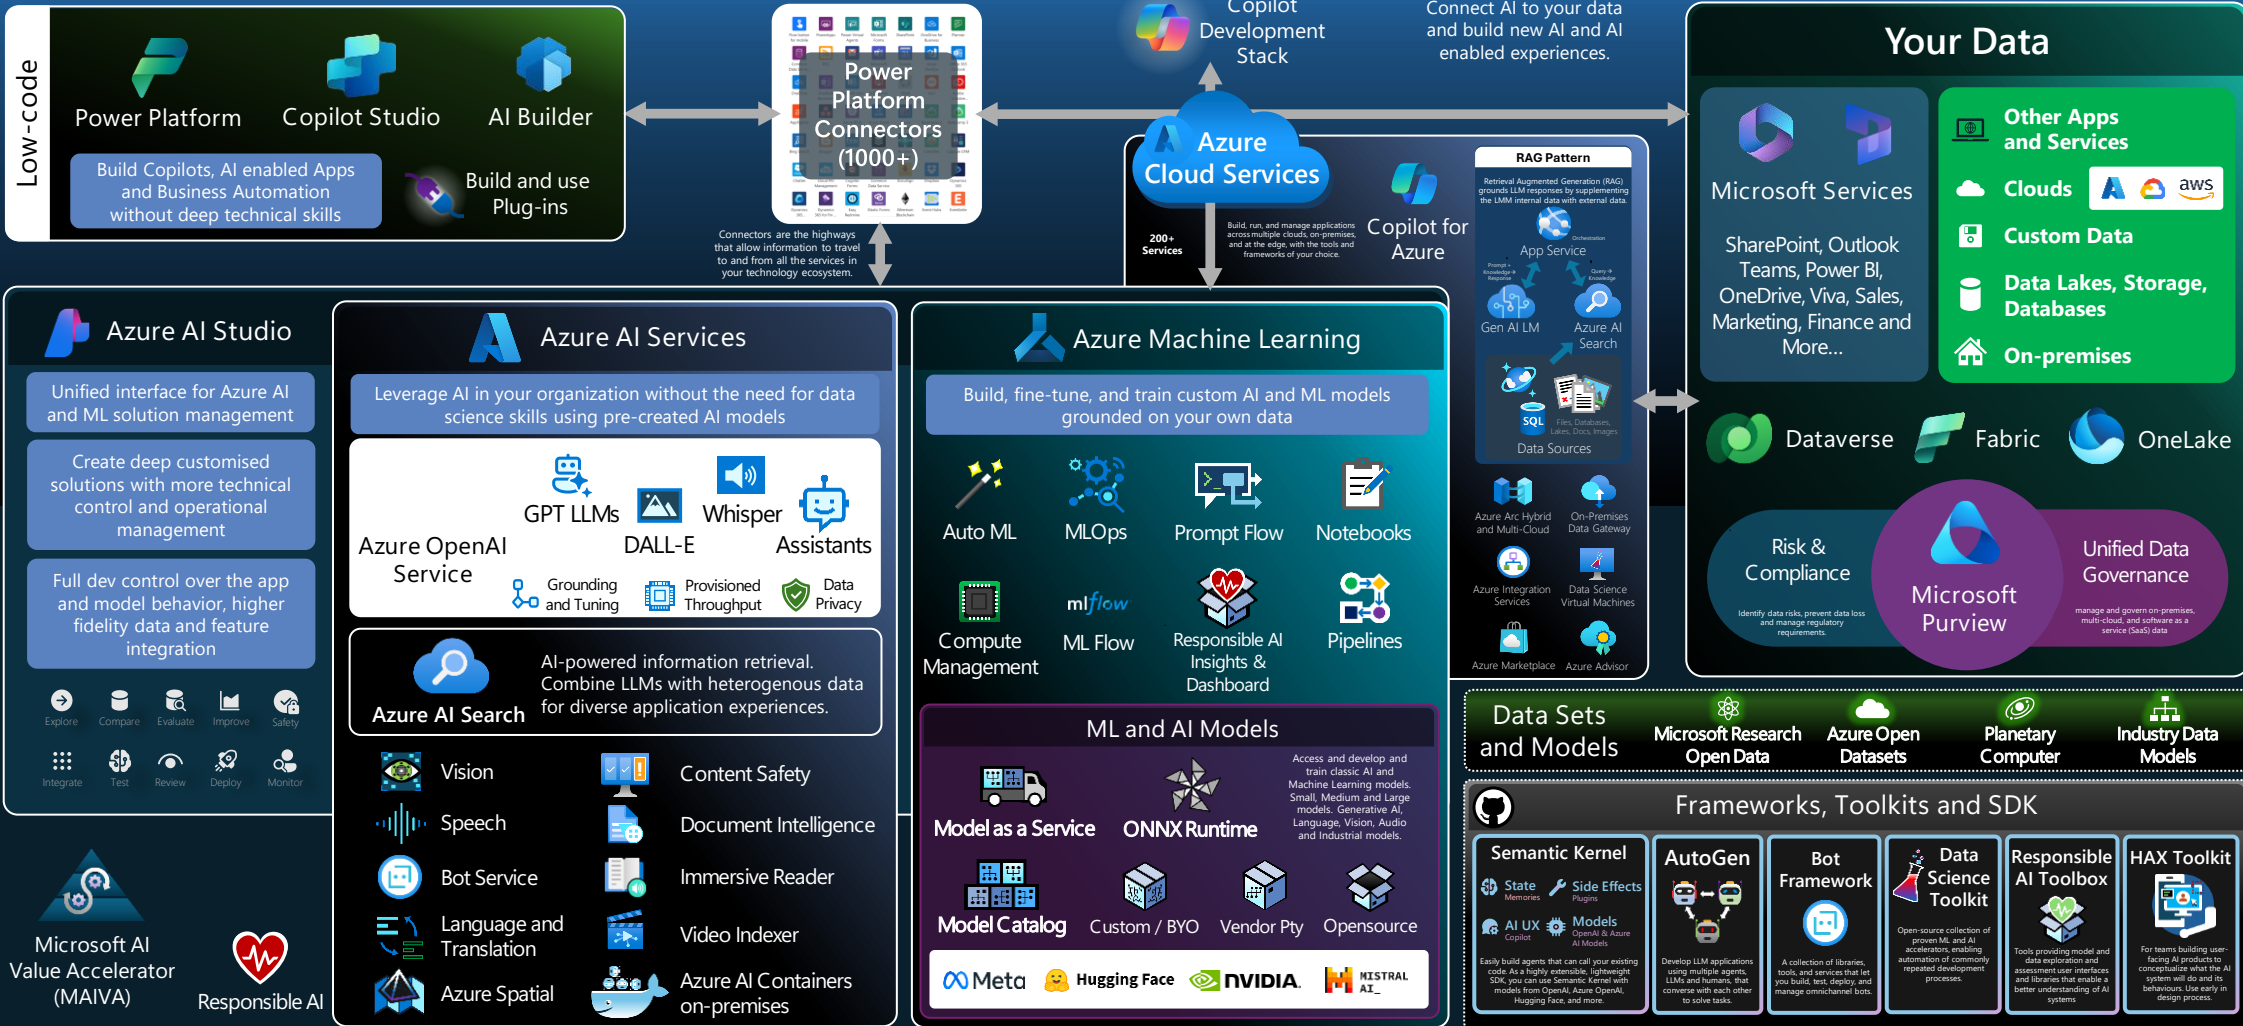
github.com/Pokeystuff/MSAIEcosystem



Extend Copilot reach using Plugins



Use



Global Infrastructure
End to end systems optimization

Datacenter



GPU & AI
Silicon



CPU



IO



Network



Circular
Centers



Sustainability

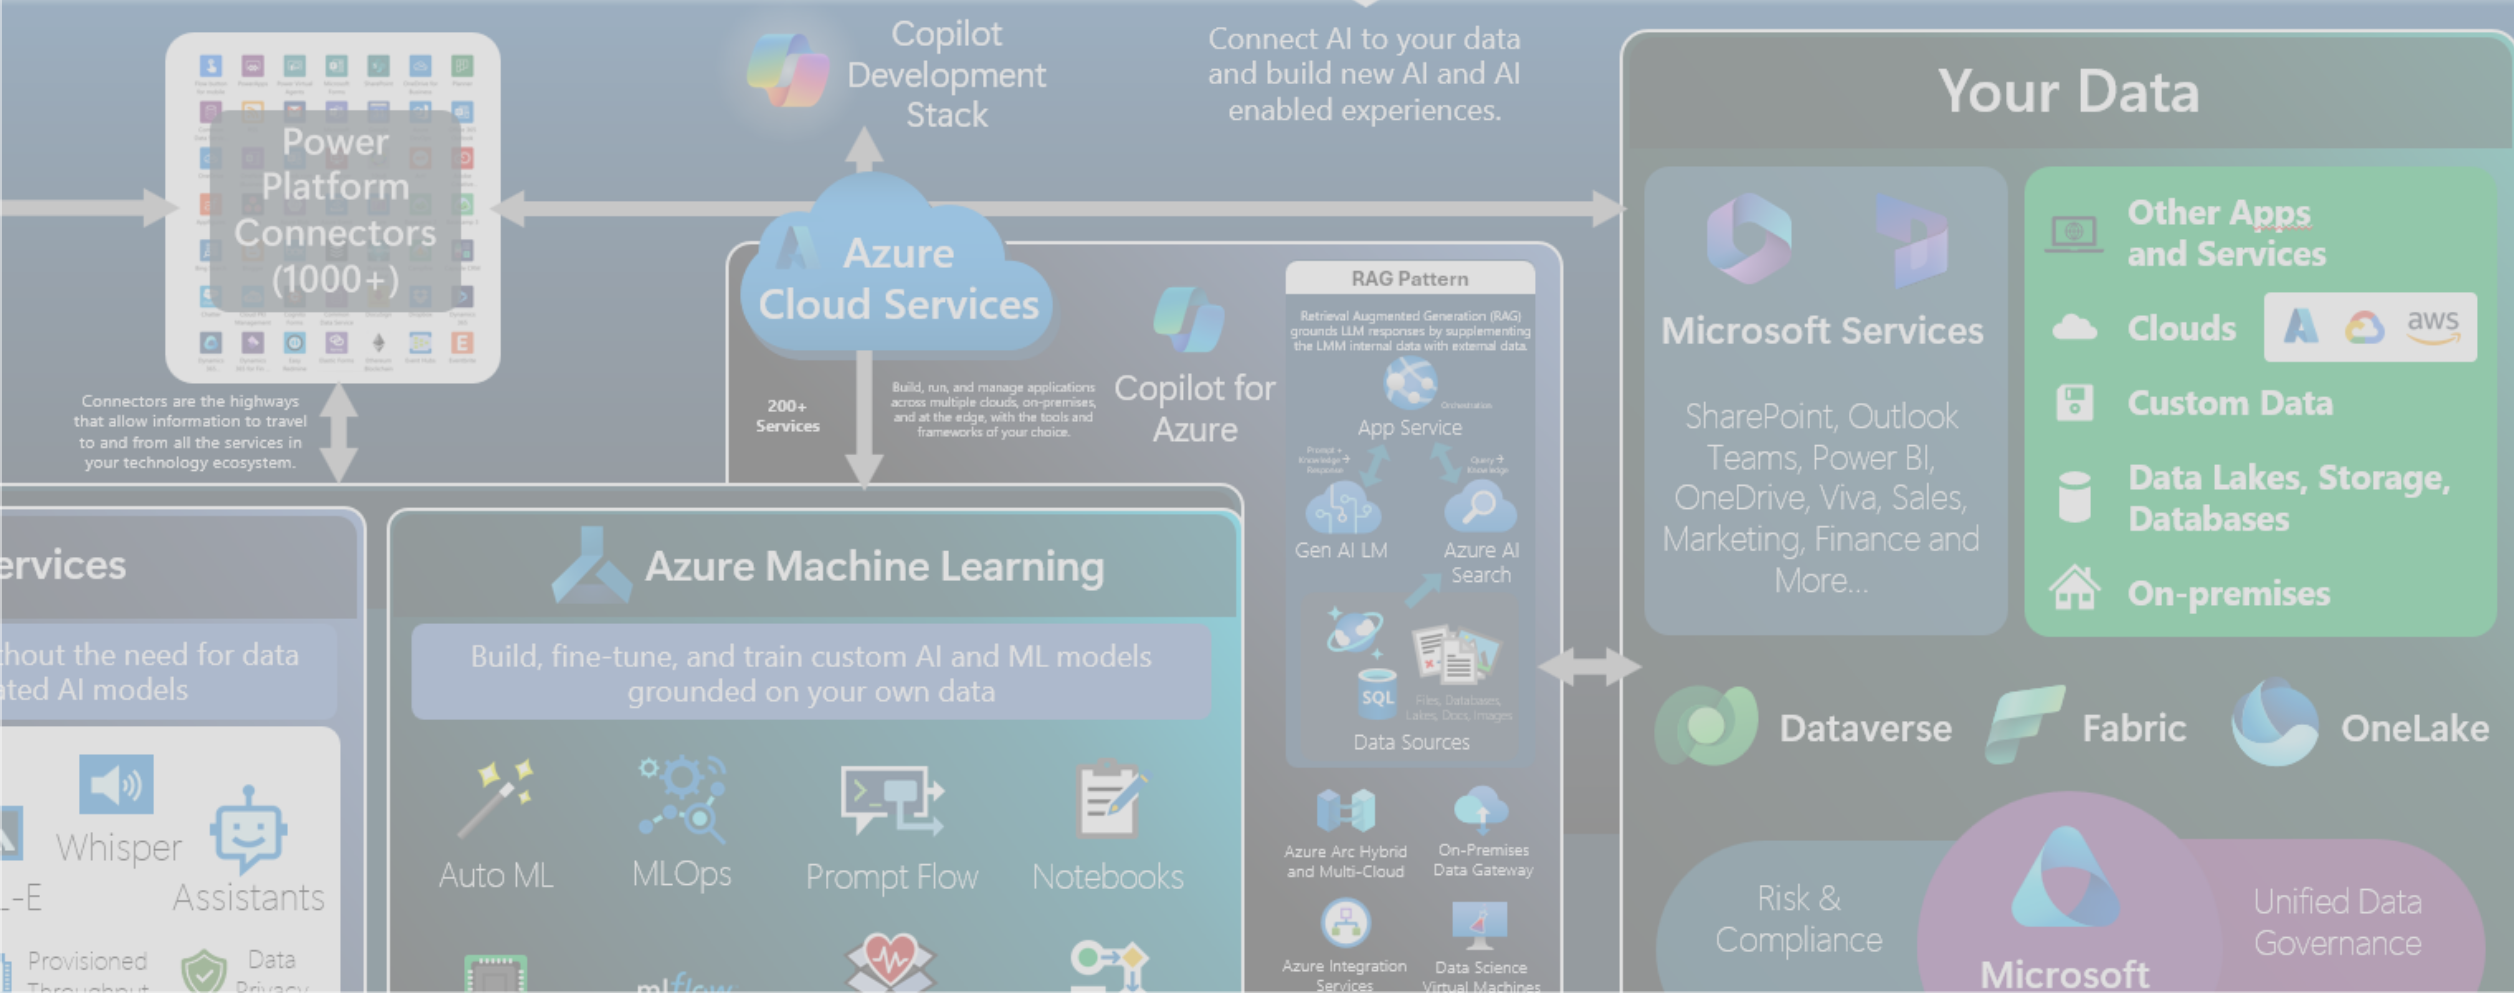


Created by Dean Corcoran

V1.51, Not an official Microsoft resource.
All trademarks are property of their respective owners



stem



Copilot

Copilot is AI-powered web chat with
commercial data protection



Chat data is not saved

Chat prompts and responses are discarded.



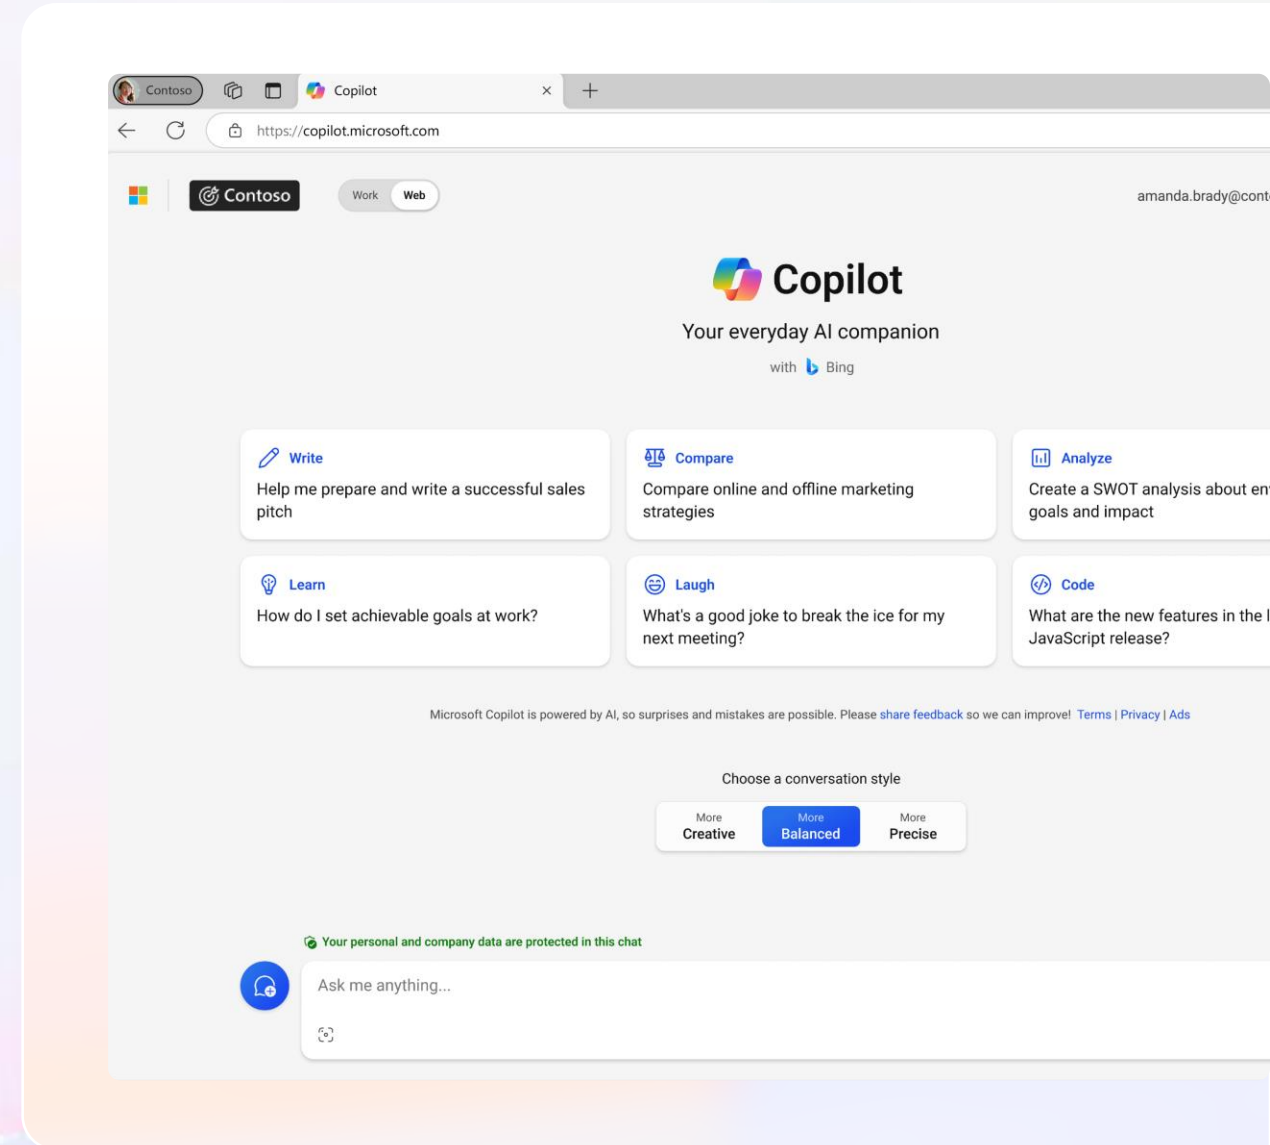
No eyes-on access

No one at Microsoft can view your data.







Data is not used to train the model

Chat data is not used to train the large language models used by AI.



Copilot.microsoft.com

What's the difference between Copilot and ChatGPT?

	 ChatGPT	 Copilot	 Copilot with CDP	 Copilot for Microsoft 365
	Personal	Personal	Work	Work
GPT Large Language Model	✓	✓	✓	✓
AI-Powered Web Search, Answers & Content Generation		✓	✓	✓
Commercial Data Protection			✓	✓
Enterprise Security, Privacy & Compliance				✓
Microsoft 365 Graph (content & context)				✓
Microsoft 365 Apps				✓
Licensing	Public & Free	Public & Free	Available with eligible Microsoft 365 license	Additional \$30/user/month

Microsoft Copilots

M365 Copilot

Bing



Better Q&A and task completion

Copilots for Web

Edge




Better interaction with web content

Word



Better reading and writing assistance

Outlook



Better e-mail management

Excel




Better data analysis

PowerPoint




Better presentations

Teams



Productive Meetings

Business Chat



Better knowledge management

Designer



Better digital creations

Windows Copilot



Better interaction with OS, apps, and files

Copilot for Creativity

Copilot for Everyday

Copilots for Productivity

Dynamics Copilot



Better sales and customer support

Copilots for Business

Fabric Copilot



Better data analytics and business intelligence

Copilots for Analytics


Security Copilot



Better threat detection, identification, and mitigation

Copilot for Security

GitHub Copilot



Better code development

Copilot for Development

Power Platform Copilot



Better creation of apps, workflows, and agents

Copilot for Low/No Code Development

File Home Insert Draw Design Layout References Mailings Review View Help

Paste

Aptos (Body) 12 A Aa A abc A

B I U X X A A

Clipboard Font

Paragraph

Normal No Spacing Heading

Paragraph Styles

Comments

Editing

Catch up

Share

Editing

Dictate

Sensitivity

Add-ins

Editor

Copilot

Draft with Copilot

Help me draft a admission recruiting proposal based on the following content: Virtual engagement: The team ideated hosting virtual campus tours, webinars with student panels, virtual information sessions, virtual reality tours, and virtual open houses. Outreach and partnerships: The team explored partnering with high schools for outreach, engaging with alumni for networking opportunities, and segmenting the audience. Promotion: The team experimented with featuring courses and highlights on social media, creating promotion materials and videos, using eye-catching visuals, and offering incentives for early applicants. Referral program: The team considered creating a referral program for current students and hosting senior recruitment talk in high schools. Open School day in Metaverse. Provide details to every items

828/2000

Generate Reference a file

Admission Recruiting Proposal

A plan to attract and engage prospective [students](#)

Virtual engagement

- Host virtual campus tours that showcase the facilities, resources, and culture of the institution.
- Organize webinars with student panels that share their experiences, insights, and tips for applicants.
- Conduct virtual information sessions that cover the admission process, requirements, deadlines, and financial aid options.
- Offer virtual reality tours that immerse the visitors in the campus environment and activities.
- Invite prospective students to virtual open houses that feature interactive sessions with faculty, staff, and current students.

Outreach and partnerships

- Partner with high schools for outreach and provide them with information and materials about the institution.
- Engage with alumni for networking and mentoring opportunities and leverage their connections and testimonials.
- Segment the audience based on their interests, preferences, and backgrounds and tailor the communication and engagement strategies accordingly.

Promotion

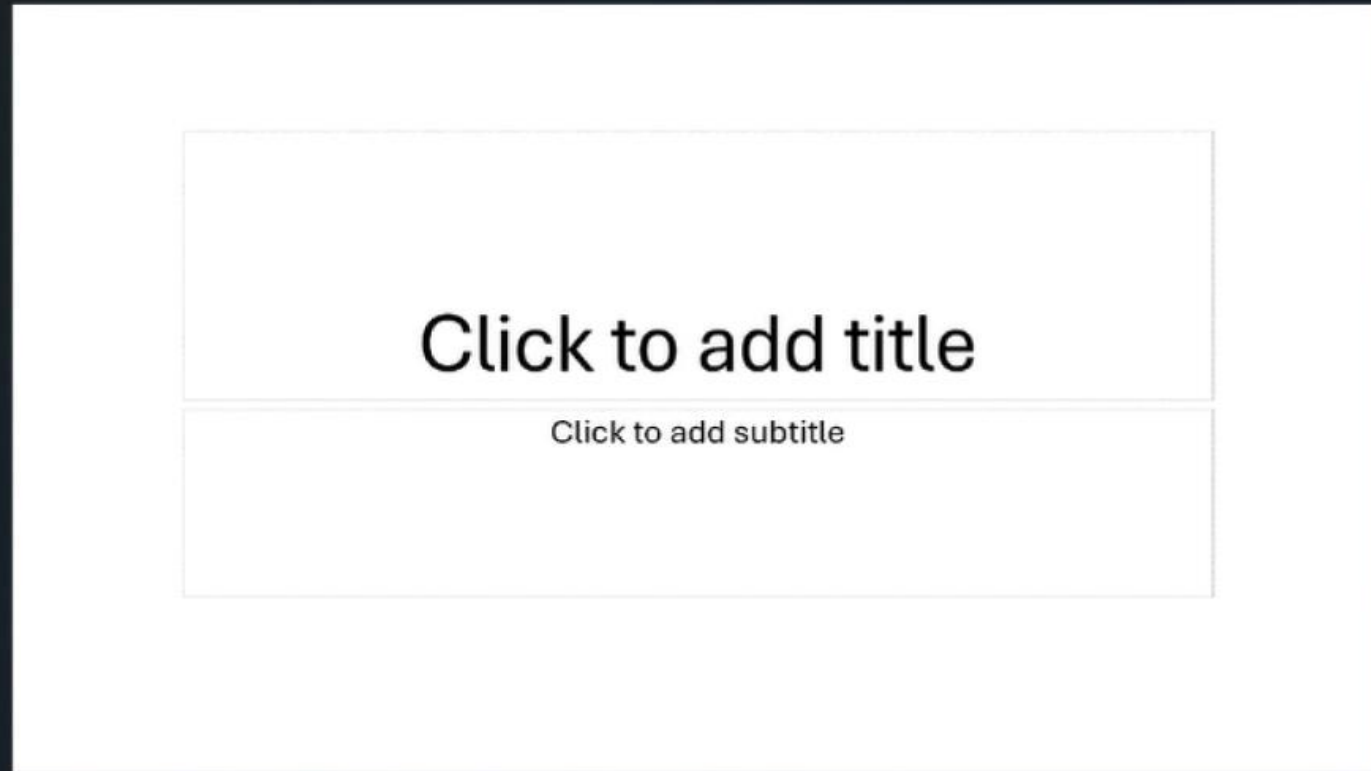

- Feature courses and highlights on social media platforms and encourage user-generated content and feedback.
- Create promotion materials and videos that highlight the benefits, values, and outcomes of studying at the institution.
- Use eye-catching visuals and catchy slogans that capture the attention and curiosity of the target audience.
- Offer incentives for early applicants such as fee waivers, scholarships, priority registration, and exclusive events.

Referral program

- Create a referral program for current students that rewards them for referring their friends, relatives, and acquaintances to apply to the institution.
- Host senior recruitment talk in high schools and invite current students to share their stories and answer questions from the seniors.

Open School day in Metaverse

- Host an open school day in the metaverse, a virtual world where people can interact with each other and the environment using avatars.
- Allow prospective students to explore the campus, attend classes, join clubs, and participate in events in the metaverse.
- Provide them with a realistic and immersive experience of studying at the institution and interacting with the community.
- Collect feedback and data from the visitors and follow up with them with personalized messages and offers.



Click to add notes

Copilot

Create presentation from file
[Admission Recruiting Proposal.docx](#)

Here's the outline for the presentation I'm generating:

- Admission Recruiting Proposal
- Executive Summary
- Virtual Engagement
- Outreach and Partnerships

AI-generated content may be incorrect

Drafting a presentation...

Stop generating

Using [this presentation](#)
Ask questions and work with this presentation
0/2000

SENSITIVITY Due to content created by Copilot, your organization automatically applied the sensitivity label: General



Agenda

- Executive Summary
- Virtual Engagement
- Outreach and Partnerships
- Promotion
- Referral Program
- Open School Day in Metaverse
- Feasibility of Strategies

Designer

Agenda

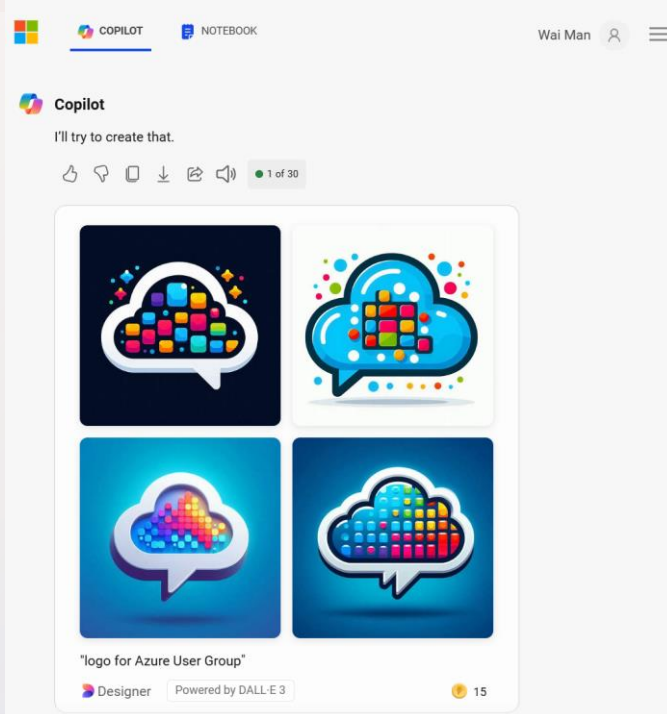
- Executive Summary
- Virtual Engagement
- Outreach and Partnerships
- Promotion
- Referral Program
- Open School Day in Metaverse
- Feasibility of Strategies

Agenda

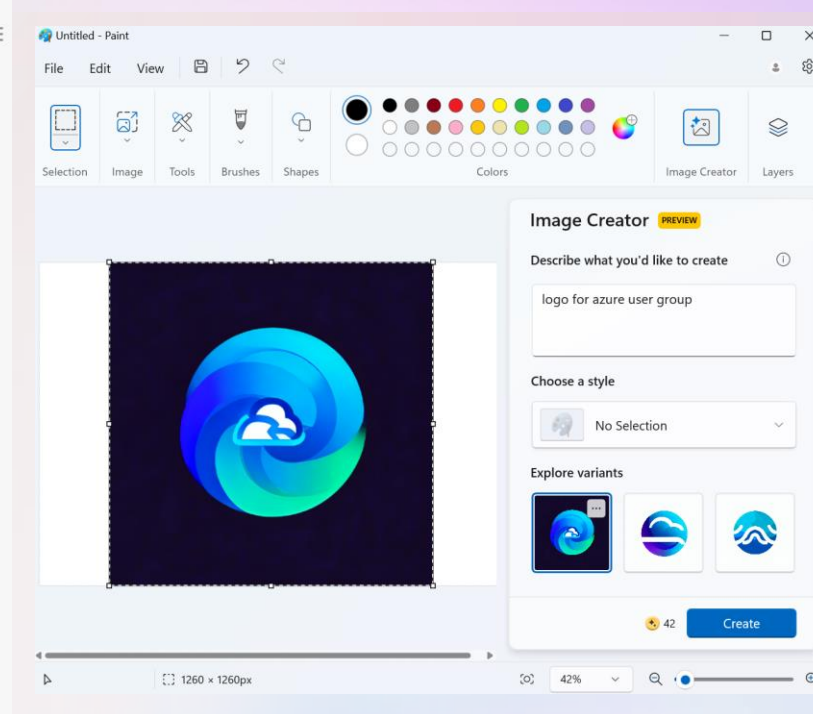
Agenda

Agenda

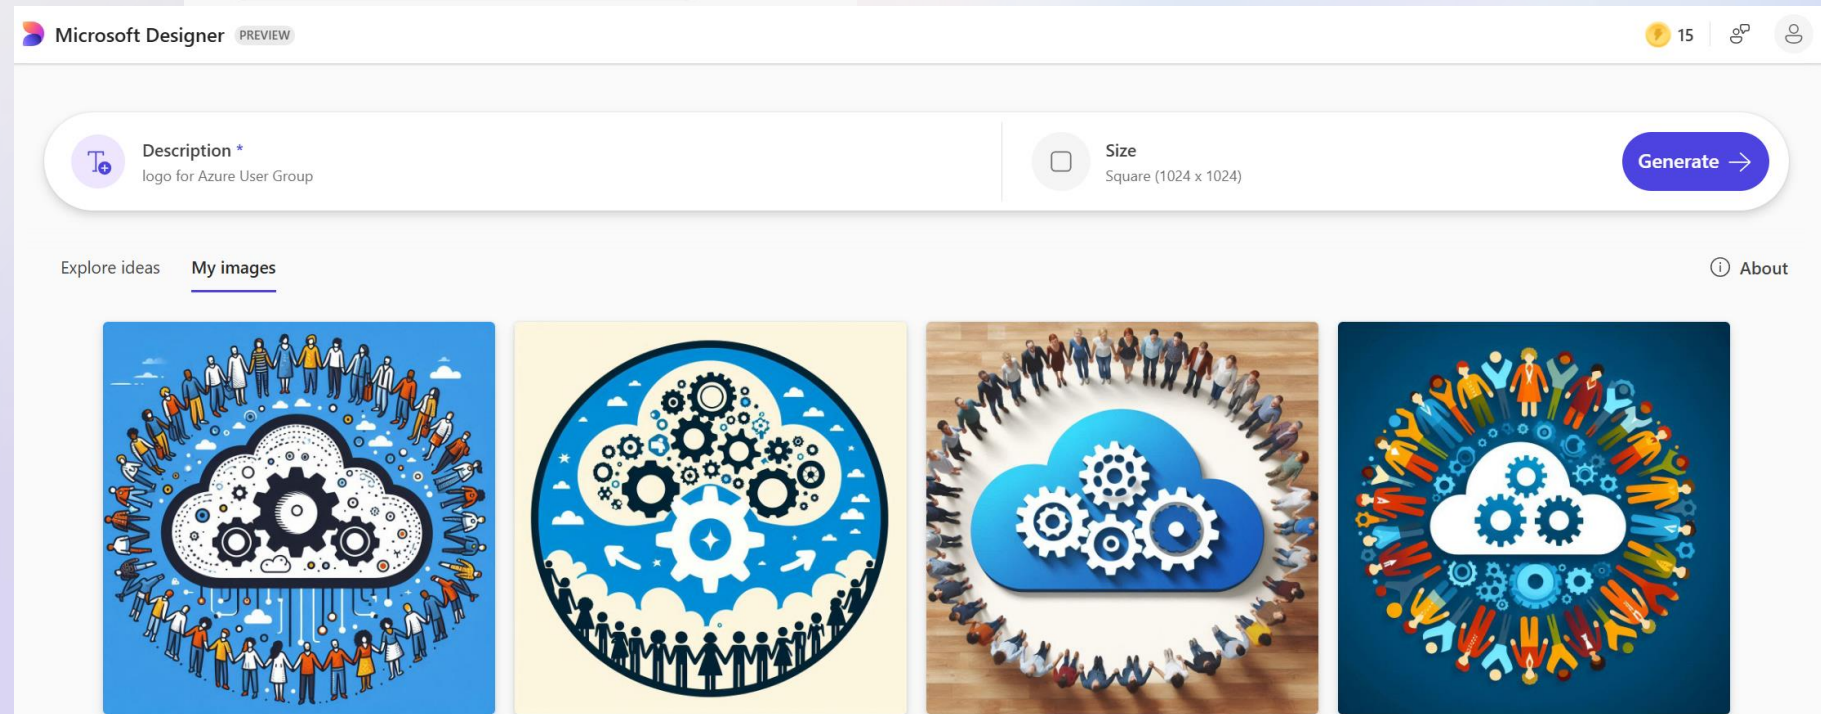
Copilot



Paint



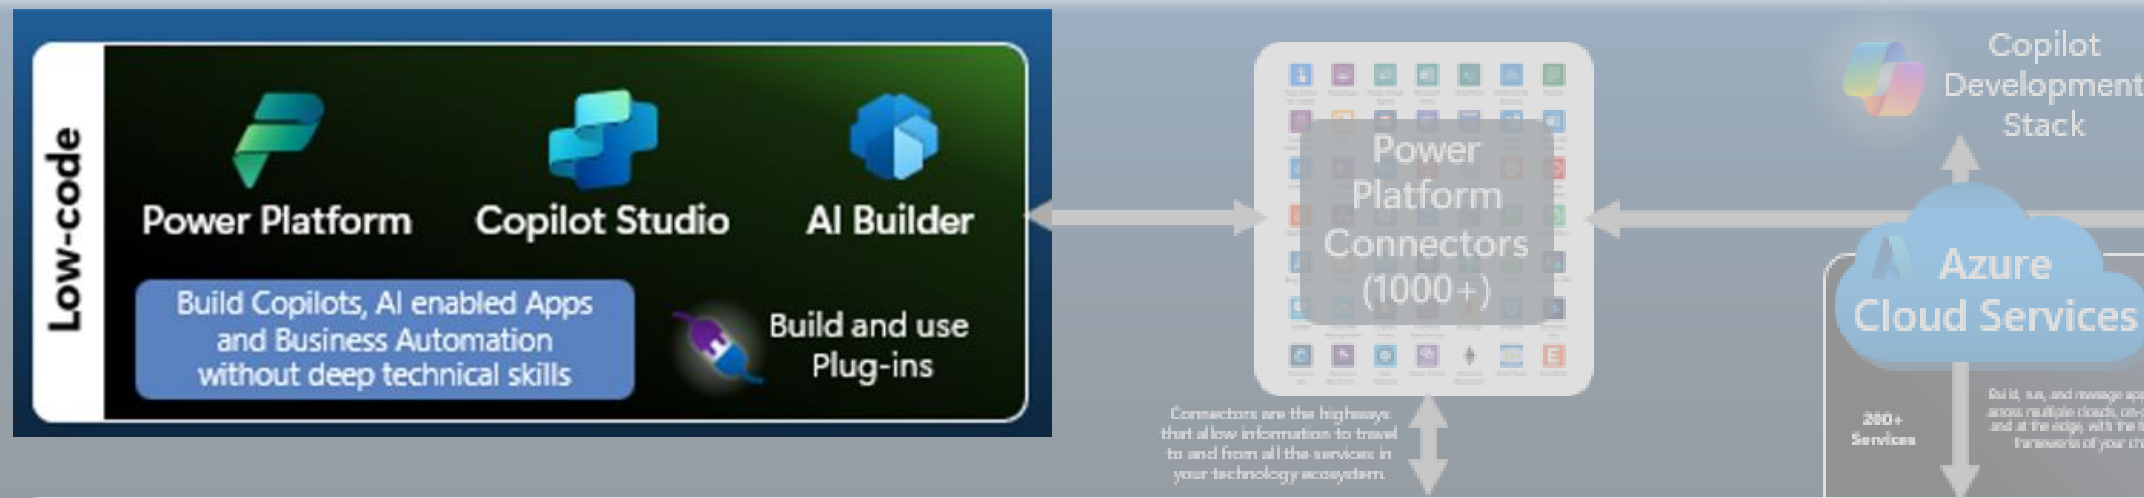
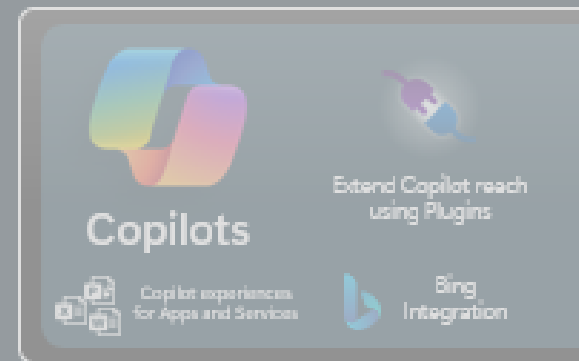
Microsoft
Designer



Microsoft AI Ecosystem

Unofficial conceptual view of AI services and resources used to incorporate AI technology into organizations to achieve more with less.

github.com/Pokeystuff/MSAIEcosystem



Azure AI Studio

Unified interface for Azure AI and ML solution management

Create deep customised solutions with more technical control and operational management

Full developer control over the app

Azure AI Services


Leverage AI in your organization without the need for data science skills using pre-created AI models

Azure OpenAI Service



Azure Machine Learning

Build, fine-tune, and train custom AI and ML grounded on your own data

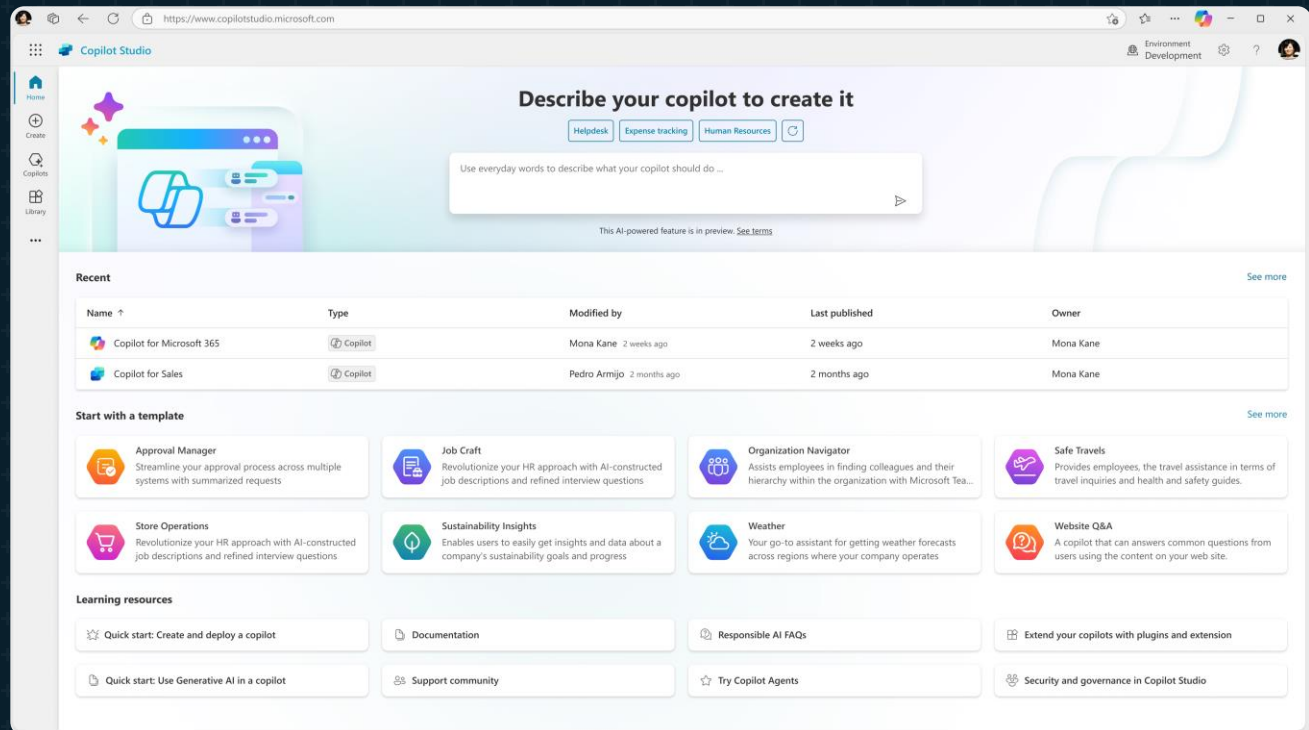


Build

Microsoft Copilot Studio

Your copilot, your way

Copilot Studio is an end-to-end conversational AI product for **building your own copilots** or **extending Microsoft Copilot** with generative AI, large language models and **your data**.



The screenshot shows the Microsoft Copilot Studio web interface. At the top, there's a navigation bar with 'Copilot Studio' and a user profile. The main heading is 'Describe your copilot to create it', followed by a text input field with the placeholder 'Use everyday words to describe what your copilot should do...'. Below this, there's a 'Recent' section with a table of existing copilots. The table has columns for Name, Type, Modified by, Last published, and Owner. Two copilots are listed: 'Copilot for Microsoft 365' and 'Copilot for Sales'. Below the table is a 'Start with a template' section with eight cards for different use cases: Approval Manager, Job Craft, Organization Navigator, Safe Travels, Store Operations, Sustainability Insights, Weather, and Website Q&A. At the bottom is a 'Learning resources' section with links to quick start guides, documentation, responsible AI FAQs, and support community.


Name	Type	Modified by	Last published	Owner
Copilot for Microsoft 365	Copilot	Mona Kane 2 weeks ago	2 weeks ago	Mona Kane
Copilot for Sales	Copilot	Pedro Armijo 2 months ago	2 months ago	Mona Kane

HR and benefits

Use everyday words to describe what your copilot should do

This AI-powered feature above is currently in preview. [See terms](#) and [supplemental terms](#)

[See more](#)

Name	Type	Last modified	Last published	Owner
 HR Copilot Analytics	Copilot	# CCI Bots 1 day ago	2 days ago	SvenErik Nielsen
 HR Copilot	Copilot	# CCI Bots 1 day ago	1 day ago	Gary Pretty

[See more](#)

Templates are currently a preview feature. See [supplemental terms](#)



Provides answers to common travel questions and related health and safety guidelines



Improve the efficiency of retail frontline workers by enabling easy access to store procedures and policies



Enables users to easily get insights and data about a company's sustainability goals and progress



Assists employees in finding colleagues and their hierarchy within the organization using Microsoft Teams



A Copilot with voice capabilities.



Your go-to assistant for getting weather forecast

Describe your copilot to create it

- Helpdesk
- Expense tracking
- HR and benefits

I want a copilot named Redi. It is an assistant for Virgin money credit card customers, helping to answer common question and help with common tasks, like changing address or checkin on the latest transactions

This AI-powered feature above is currently in preview. [See terms](#) and [supplemental terms](#)

	Type	Last modified	Last published	Owner
Analytics	Copilot	# CCI Bots 1 day ago	2 days ago	Sven
	Copilot	# CCI Bots 1 day ago	1 day ago	Gary


[Skip to configure](#)
[Create](#)


Copilot

✓ Primary language is English

[Edit language](#)

✓ An assistant for *Virgin Money* credit card customers, helping to answer common questions and assist with tasks such as changing addresses or checking the latest transactions.

Great! Your copilot Redi will be an assistant for Virgin Money credit card customers. It will help answer common questions and assist with tasks such as changing addresses or checking the latest transactions. Is there anything else you'd like to add to the description?

Just now



Low-code

Build Copilots, AI enabled Apps and Business Automation without deep technical skills



Build and use Plug-ins

Connectors are the highways that allow information to travel to and from all the services in your technology ecosystem.

Connectors (1000+)

Azure Cloud Services

200+ Services

Build, run, and manage applications across multiple clouds, on-premises, and at the edge, with the tools and frameworks of your choice.

Copilot for Azure

RAG Pattern

Retrieval Augmented Generation (RAG) grounds LLM responses by supplementing the LLM internal data with external data.



Azure AI Studio

Unified interface for Azure AI and ML solution management

Create deep customised solutions with more technical control and operational management

Full dev control over the app and model behavior, higher fidelity data and feature integration

Explore Compare Evaluate Improve Safety

Integrate Test Review Deploy Monitor



Microsoft AI Value Accelerator (MAIVA)




Responsible AI



Azure AI Services

Leverage AI in your organization without the need for data science skills using pre-created AI models

Azure OpenAI Service



Azure AI Search

AI-powered information retrieval. Combine LLMs with heterogeneous data for diverse application experiences.

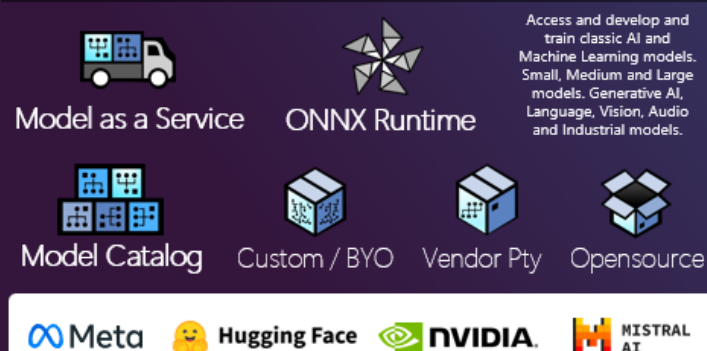


Azure Machine Learning

Build, fine-tune, and train custom AI and ML models grounded on your own data




ML and AI Models



Data Sets and Models



Semantic Kernel



Global Infrastructure
End to end systems optimization

Datacenter



GPU & AI Silicon



CPU



IO



Network



Circular Centers



Sustainability



I ML THEREFOR I AM?

The Microsoft Azure AI Portfolio

Azure AI Studio

The place to test, build and deploy AI solutions

Azure AI Services

Pre-built models, APIs and SDKs to infuse into custom apps



Azure OpenAI Service



Azure AI Search



Azure AI Speech



Azure AI Vision



Azure AI Content Safety



Azure AI Document Intelligence



Azure AI Language



Azure AI Translator

Azure Machine Learning

Advanced tools for designing and fine-tuning specialized AI models



Responsible AI Dashboard



Model Catalog



Prompt Flow



MLOps And LLMOps

Azure AI Infrastructure

State-of-art supercomputing to power AI workloads

Florence

GPT-4 and GPT-3.5-Turbo

Embeddings

Meta Llama 2

Turing

Whisper


DALL-E

Hugging Face

Microsoft and OpenAI partnership



Ensure that artificial
general intelligence (AGI)
benefits humanity



Empower every person and
organization on the planet
to achieve more

Azure OpenAI Service

Language

GPT-4,
4-Turbo
3.5-Turbo

Multi-Modal

GPT-4o
GPT-4 Vision

Fine Tuning

GPT-4
GPT-3.5-Turbo

Images

DALL·E 3

Transcription & Translation

Whisper

On Your Data GA

Azure AI Studio GA

Assistants API Preview

Pricing Innovation

New in Azure AI

General availability

Azure AI Studio

General availability

GPT-4o

General availability

Phi-3 SLMs

General availability

GPT-4 Turbo with Vision, DALL·E 3, Whisper

Generally Available (Refresh)

On your data

Preview (Refresh)

Assistants API

Preview

Fine Tuning for GPT-4

Preview

Batch API

Preview

Configurable Content Filters for Dall-E and GPT-4-Turbo

Preview

Integration with Azure Communication Services for WhatsApp

Preview

Video Translation

Azure OpenAI compliance

To date Azure OpenAI Service has been certified or assesses against 40 of the 103 compliance offerings that Microsoft measures global compliance against (Global, Regional & Industry)

23 NYCRR 500	Canada Privacy Law	DoD DISA SRG L2	NERC
Argentina PDPA	CDSA	SOC 1, 2, 3	OSPAR
CSA Star Attest	FIPS 140-2	SOX	SEC Regulation
CSA Star Cert	FedRAMP High	European BA Export	NIA
CSA Star Self-	ISO 27018:2019	(EAR)	PCI 3DS
Assessment	ISO 27701:2019	HIPAA BAA	PCI DSS Level 1
ISO 20000-1:2011	ISO 9001:2015	Germany C5	NIST CSF
ISO 22301:2012	HiTRUST	GDPR	TruSight
ISO 27001:2013	MARS-E (US)	GxP (FDA21 CFR11)	WCAG 2.0
ISO 27017:2015	DESC	HDS (France)	Singapore MTC
	UK G-Cloud	UK Cyber	TISAX
		Essentials+	

The Azure OpenAI Service certifications are described in the [Microsoft Azure Compliance Offerings document](#)



Azure AI Studio

Generally available

ai.azure.com

Azure AI Studio | Model catalog

Contoso Responsible R&D
ContosoPharmaceutical-0034

← To all projects
Current project
Contoso-0034

Project overview
Settings

Get started
Model catalog
Model benchmarks
Prompt catalog

Project playground
Chat
Completion
Assistants
Images

Tools
Code
Prompt flow
Tracing
Evaluation
Fine-tuning

Components
Data
Indexes
Deployments
Content filters

Find the right model to build your custom AI solution

Announcements

Phi-3 is now available
Microsoft's Phi-3-mini SLMs offer groundbreaking performance at a small size.
[View models](#) [Read blog](#)

Build the future of AI with Meta Llama 3
Serverless APIs for Meta-Llama-3-8B-Instruct and Meta-Llama-3-70B-Instruct models.
[View models](#) [Read blog](#)

Cohere models are here
Cohere's Enterprise AI models include Command R, Command R+, and Embed v3.
[View models](#) [Read blog](#)

All filters Publishers Deployment type Tasks Licenses

Search

Models 1,234

gpt-4-32k ✓ Text generation	llama-7b ✓ Text generation	mistralai-Mixtral-8x7B-Instruct ✓ Chat completion
databricks-dolly ✓ Text generation	Deci-DeciLM-7B ✓ Text generation	gpt-35-turbo ✓ Text generation
llama-12b ✓ Text generation	microsoft-phi-2 ✓ Text generation	llama-finetune-12b ✓ Text generation
gpt-4 ✓ Text generation	tiiuae-falcon-40b ✓ Text generation	gpt-4 ✓ Text generation
gpt-4-32k ✓ Text generation	llama-7b ✓ Text generation	mistralai-Mixtral-8x7B-Instruct ✓ Chat completion
databricks-dolly ✓ Text generation	Deci-DeciLM-7B ✓ Text generation	gpt-35-turbo ✓ Text generation
llama-12b ✓ Text generation	microsoft-phi-2 ✓ Text generation	llama-finetune-12b ✓ Text generation
gpt-4 ✓ Text generation	tiiuae-falcon-40b ✓ Text generation	gpt-4 ✓ Text generation

Image may not reflect actual user interface.

New in Azure AI Studio

Azure AI Studio general availability

Model catalog & benchmarking
generally available

Azure OpenAI Service: GPT 4 Turbo
with Vision generally available

New MaaS partners and
models in catalog in preview

HiddenLayer partnership

Assistants API

Prompt flow: Debugging
& tracing in preview

Quick-start templates integrated with Azure
Developer (AZD) CLI

Integration with AI Toolkit
for Visual Studio Code

and more!

Azure AI offers the best model selection

LLMs and SLMs

Flagship LLMs

- GPT-4
- Mistral Large
- Llama3 70b
- Command R+

Small Language Models

- Phi3
- Mistral OSS models
- Llama3 8b

Modalities, tasks and tools

Multi-modal

- GPT-4o, Phi3-vision

Image generation

- Dalle3, Stability AI

Embedding models

- Ada, Cohere
- Function calling & JSON support

Regional and domain specific

- Core42 JAIS Arabic language LLM
- Mistral Large is focused on European languages
- Nixtla TimeGEN-1 - Timeseries forecasting

Open and propriety

Premium models first on Azure: OpenAI, Mistral Large, Cohere Command R+

100s of Open models from HuggingFace

Open models from Meta, Databricks and Snowflake, Nvidia



Microsoft Phi



Azure OpenAI



Mistral AI



Meta AI



Databricks



Cohere



Hugging Face



NVIDIA



Deci AI



Nixtla



G42



Snowflake

Model catalog offers flexible deployment options



Managed compute

- Deployed to managed Azure VMs
- One click deployment that boots up an optimized container and initializes the weights.
- Pay-per-GPU billing
- 100s of open models



Serverless APIs with Azure Open AI Service and Models as a Service

- Ready to use model APIs
- Pay-per-token or pay-for-provisioned-capacity billing
- 30+ flagship models

[Azure OpenAI](#)**Playground**[Chat](#)[Completions](#)[DALL-E](#)[Assistants \(Preview\)](#)**Management**[Deployments](#)[Models](#)[Data files](#)[Quotas](#)[Content filters \(Preview\)](#)

Get started

Text generation



Assistants playground

PREVIEW

Speed up development of GPT-powered AI Assistants with prebuilt conversation state management and customization tools.

[Try it now](#)

Text generation



Bring your own data

Ground your own data on advanced AI models to create conversational copilots that aid user comprehension, task completion, and decision-making.

[Try it now](#)

Text generation



Chat playground

Design a customized AI assistant using ChatGPT. Experiment with GPT-3.5-Turbo and GPT-4 models.

[Try it now](#)

Text generation



Completions playground

Experiment with completions models for use cases such as summarization, content generation, and classification.

[Try it now](#)

Image generation




DALL-E playground

PREVIEW

Generate unique images by writing descriptions in natural language.

[Try it now](#)

Speech to text



Whisper Speech-to-Text

PREVIEW

Transcribe your audio using the Whisper model. Use text prompts to improve transcription quality.

[Create a Whisper deployment](#)

- «
- Azure OpenAI
 - Playground
 - Chat
 - Completions
 - DALL-E**
 - Assistants (Preview)
 - Management
 - Deployments
 - Models
 - Data files
 - Quotas
 - Content filters (Preview)

Azure OpenAI Studio > DALL-E playground

[Privacy & cook](#)

DALL-E playground

Deployment: dall-e-3 View code Settings

Tile Size: ☐ Medium tiles

Prompt ⓘ

logo for azure user group

Genera



logo for azure user group



logo for azure user group



logo for azure user group





Azure OpenAI

Playground

Chat

Completions

DALL-E

Assistants (Preview)

Management

Deployments

Models

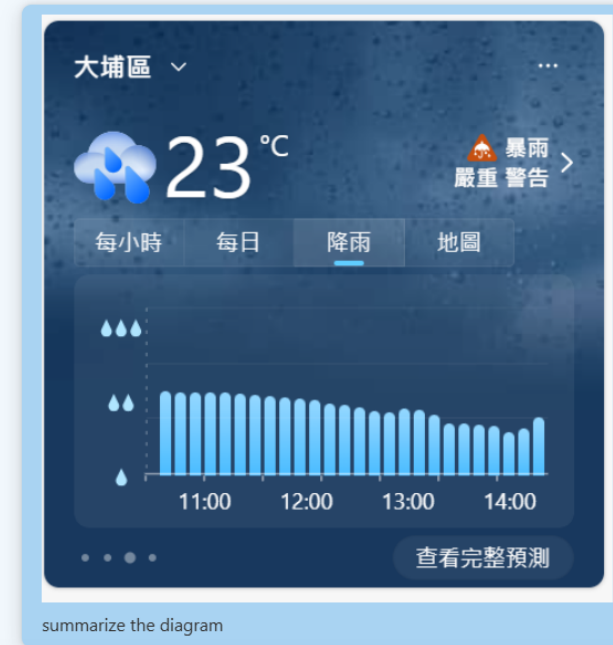
Data files

Quotas

Content filters (Preview)

Chat playground

Deploy to

[Import setup](#) [Export setup](#) [Show panels](#)[Clear chat](#) | [Playground settings](#) [View code](#) ☐ Show JSON

summarize the diagram

The image displays a weather forecast interface. It shows the current temperature as 23°C, with weather conditions indicated by a cloud and raindrops, suggesting it is raining. There is also a warning sign, possibly indicating a weather alert such as a storm. The graph at the bottom represents rainfall intensity over time, from 11:00 to 14:00, showing an increasing trend in rainfall as the time progresses. The interface includes tabs for hourly, daily, and other forecasts, though they are not in focus in this image. The text at the bottom suggests an option to view more detailed weather information.

New chat session started

Type user query here. (Shift + Enter for new line)

11/8192 tokens to be sent

Chat playground

Deploy to

Import setup

Export setup

Show panels

Setup

Prompt Add your data

Apply changes Safety system message

Use a system message template

Using templates

Use a template to get started, or just start writing your own system message below. Want some tips? [Learn more](#)

Select a template

System message

You are an AI assistant that helps people find information.

Examples

Using examples

Add examples to show the chat what responses you want. It will try to mimic any responses you add here so make sure they match the rules you laid out in

Clear chat Chat capabilities View code Show JSON

please extract the text



The text in the image is "LYNN flextime".

Type user query here. (Shift + Enter for new line)

164/128000 tokens to be sent

Configuration

Deployment Parameters

Deployment *

gpt-4o

Session settings

Past messages included

10

Current token count

Input tokens progress indicator
164/128000

Chat playground

Deploy to ▾

↑ Import setup ↓ Export setup ⚙ Show panels ▾

Setup

Prompt Add your data

📄 Apply changes 🛡 Safety system message

Use a system message template

Using templates

Use a template to get started, or just start writing your own system message below. Want some tips? [Learn more](#)

Select a template ▾

System message ⓘ

You are an AI assistant that helps people find information.

Examples ⓘ

Using examples

Add examples to show the chat what responses you want. It will try to mimic any responses you add here so make sure they match the rules you laid out in

Select all images with an Orange.



Configuration

Deployment Parameters

Deployment *

gpt-4o ▾

Session settings

Past messages included ⓘ

10

Current token count ⓘ

Input tokens progress indicator
259/128000

The images containing an orange are the ones in the first column (top and bottom), and the bottom-right corner image.

[Azure OpenAI](#)**Playground**[Chat](#)[Completions](#)[DALL-E](#)[Assistants \(Preview\)](#)**Management**[Deployments](#)[Models](#)[Data files](#)[Quotas](#)[Content filters \(Preview\)](#)

Get started

Text generation



Assistants playground

PREVIEW

Speed up development of GPT-powered AI Assistants with prebuilt conversation state management and customization tools.

[Try it now](#)

Text generation



Bring your own data

Ground your own data on advanced AI models to create conversational copilots that aid user comprehension, task completion, and decision-making.

[Try it now](#)

Text generation



Chat playground

Design a customized AI assistant using ChatGPT. Experiment with GPT-3.5-Turbo and GPT-4 models.

[Try it now](#)

Text generation



Completions playground

Experiment with completions models for use cases such as summarization, content generation, and classification.

[Try it now](#)

Image generation




DALL-E playground

PREVIEW

Generate unique images by writing descriptions in natural language.

[Try it now](#)

Speech to text



Whisper Speech-to-Text



PREVIEW

Transcribe your audio using the Whisper model. Use text prompts to improve transcription quality.



[Create a Whisper deployment](#)

 Azure OpenAI

Playground

 Chat Completions DALL-E Assistants (Preview)

Management

 Deployments Models Data files Quotas Content filters (Preview)

Add data

☒ Data source☐ Review and finish

Select or add data source

Your data source is used to ground the generated results with your data. Select an existing data source or create a new data connection with Azure Blob Storage, databases, search, URLs, or local files as the source the grounding data will be built from. [Learn more about data privacy and security in Azure AI.](#)

Select data source *

Select data source



Azure AI Search

Azure Blob Storage (preview)

Azure Cosmos DB for MongoDB vCore

Elasticsearch (preview)

URL/web address (preview)

Upload files (preview)

Next

Cancel

Using Vector Indexes in Azure Databases and/or Azure AI Search

Vector Indexes in Azure databases are preferred when:

- You have structured or semi-structured **operational data** (e.g. chat history, customer profiles, etc.) in that database
- Simplified architecture for **single source of truth** combining **vector similarity search inline with database queries**. No need to synchronize separated database and vector index
- The workload benefits from mission-critical **OLTP database characteristics** (e.g. stronger guarantees for availability and performance)



MongoDB
familiarity
NoSQL / semi-
structured data



Azure Cosmos DB
for MongoDB
(vCore-based)

PostgreSQL
familiarity
Relational data



Azure Database
for PostgreSQL
(flex)

Azure AI Search is preferred when:

- You need to index structured/unstructured data (e.g. images, docx, PDFs) from a **variety of internal and external data sources**
- Your application requires **state-of-the-art search technology** for higher search quality (e.g. hybrid full-text/vector search, fuzzy, autocomplete, semantic re-ranking, multi-language, metafiltering)
- The workload requires **multi-modal search** and/or multi-modal embeddings to perform OCR, image analysis, translation, etc.
- You're building a **Bing-like search experience** in a custom application




Azure AI Search




 Azure OpenAI

Playground

 Chat Completions DALL-E Assistants (Preview)

Management

 Deployments Models Data files Quotas Content filters (Preview)

Chat playground



Import setup

A new web app...

A new copilot in Copilot Studio (preview)

Deploy to

Setup





Prompt Add your data

Gain insights into your own data source. Your data is stored securely in your Azure subscription. [Learn more about how your data is protected.](#)

Data source: Azure AI Search
Index: chat-with-data-index

Search Resource: chat-with-data-search

Advanced settings >

 Remove data source Clear chat Playground settings View code



Show JSON



Start chatting

Test your assistant by sending queries below. Then adjust your assistant setup to improve the assistant's responses.

Type user query here. (Shift + Enter for new line)


11/8192 tokens to be sent  



Start chatting

This chatbot is configured to answer your questions



What is the benefit of 



Generally Available

GenAI gateway features in Azure API Management

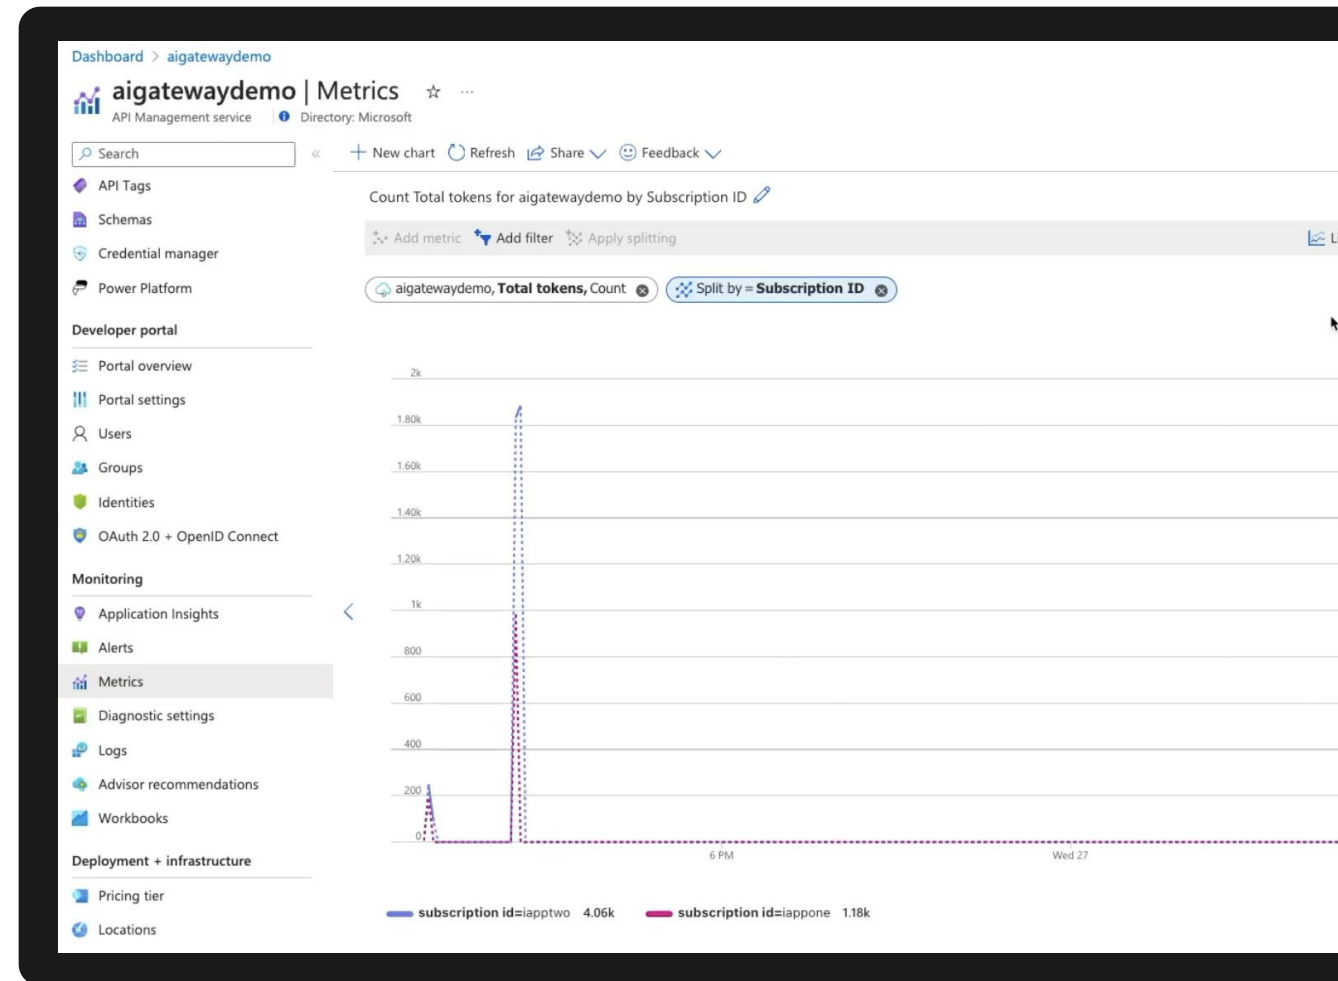
Centralize management of model and prompt endpoints

- Load balance across OpenAI endpoints, PTU/PAYG, and different models
- Track token consumption and attribution for streaming and non-streaming requests
- Improve performance and conserve tokens with semantic caching
- Secure access to models with Managed Identity

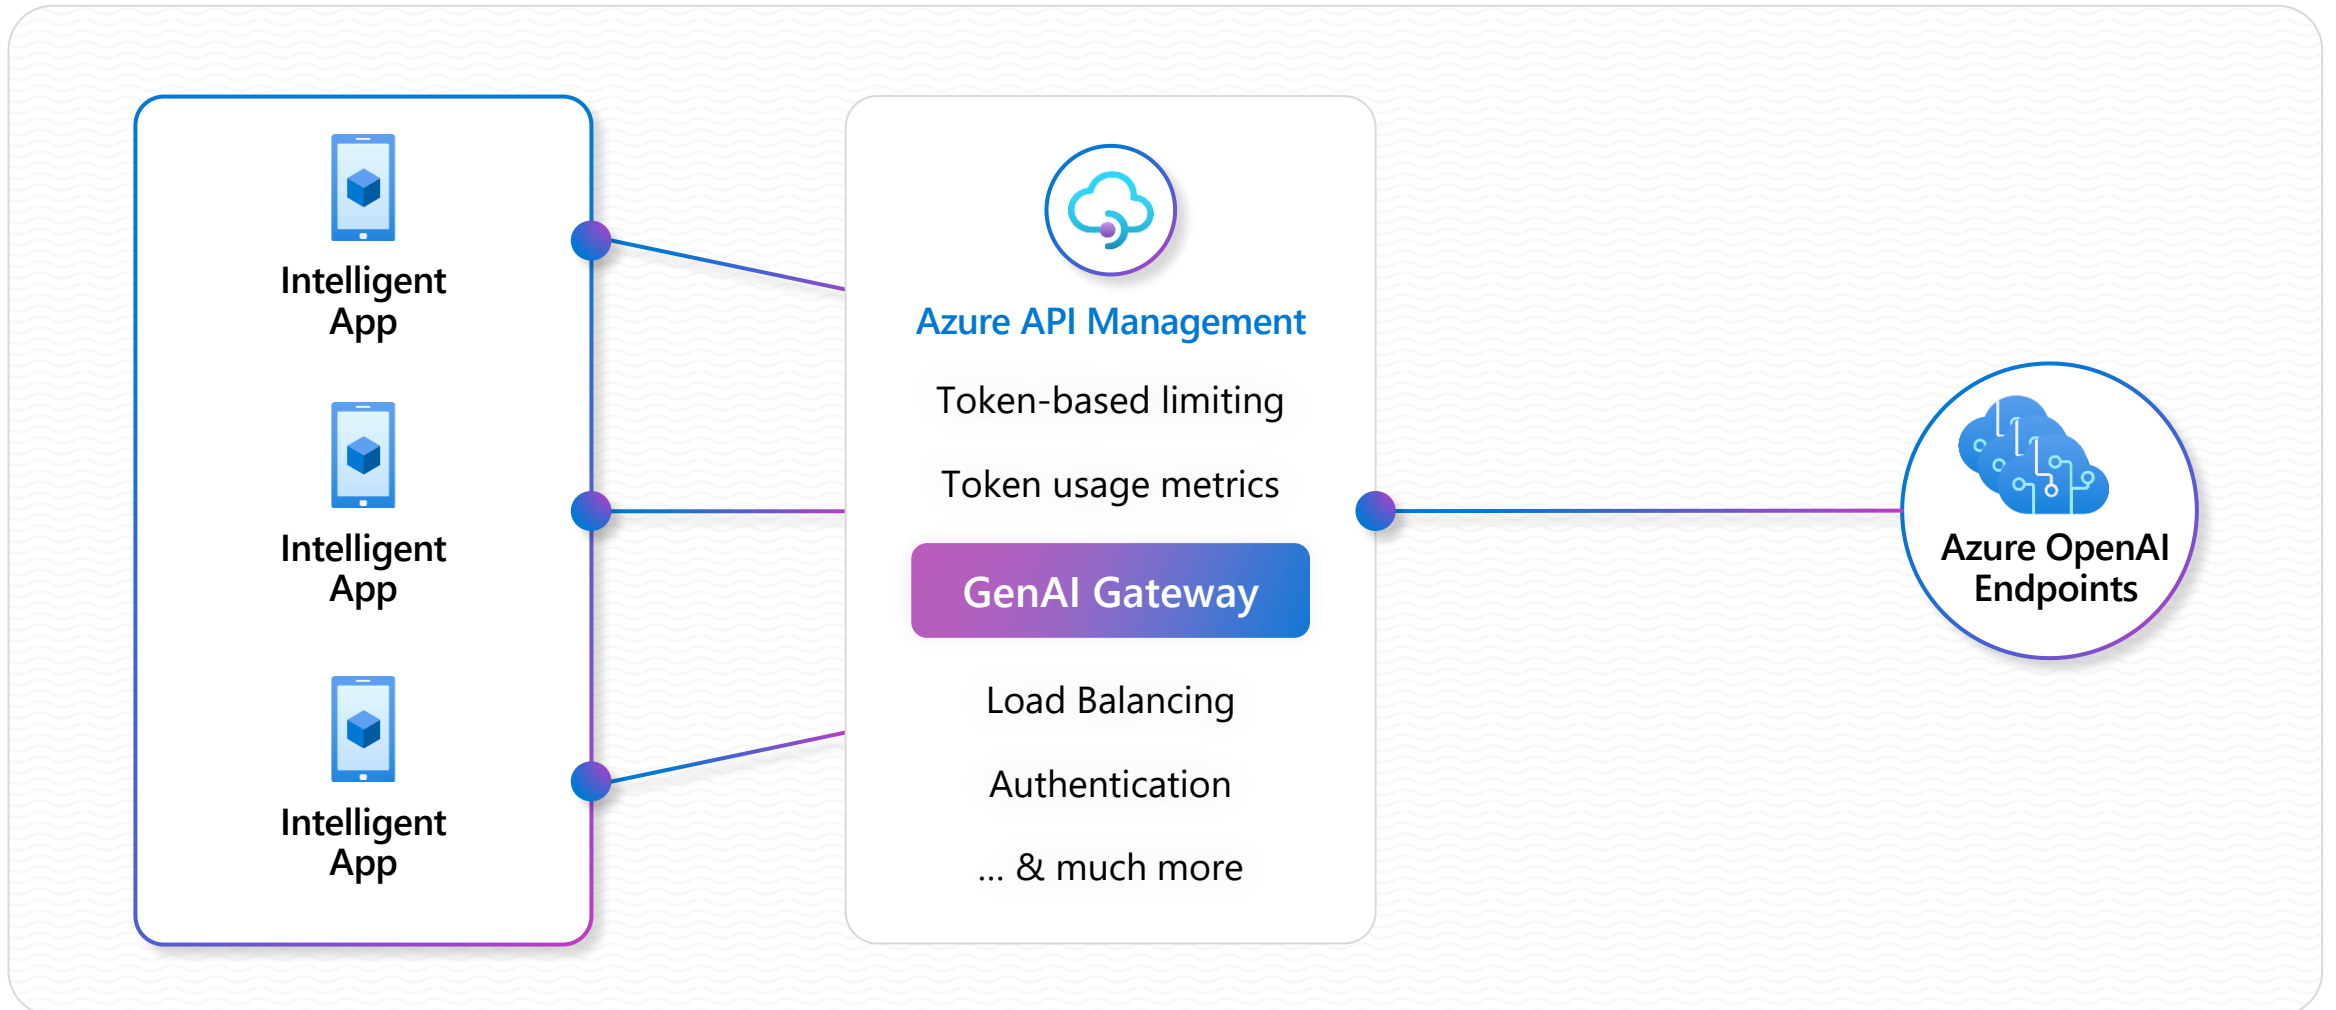
Public preview

Guide: <https://aka.ms/genai-gateway>

Landing zone accelerator: <https://aka.ms/apim-genai-lza>

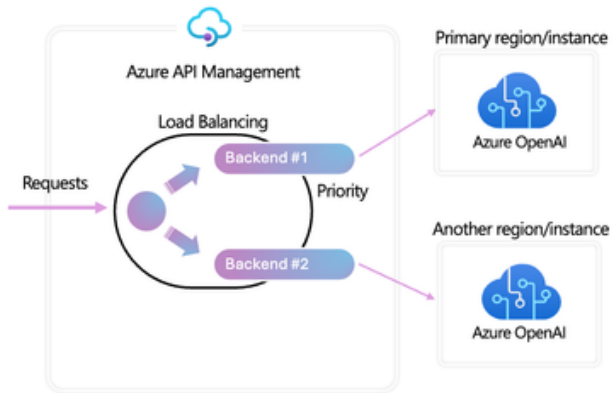


Scaling Up: Multiple Apps, Multiple OpenAI Endpoints



GenAI gateway features in Azure API Management

Load Balancer and Circuit Breaker

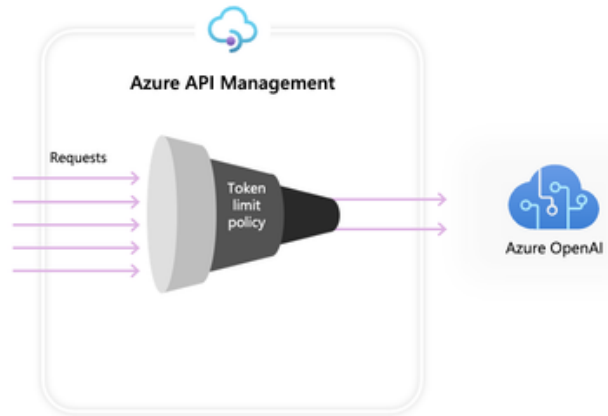


Distribute requests across PTU and Pay-as-you-go instances

Define load balancing pools to include multiple Azure OpenAI endpoints

Configure circuit breaker rules for successful failover

Azure OpenAI Token Limit policy



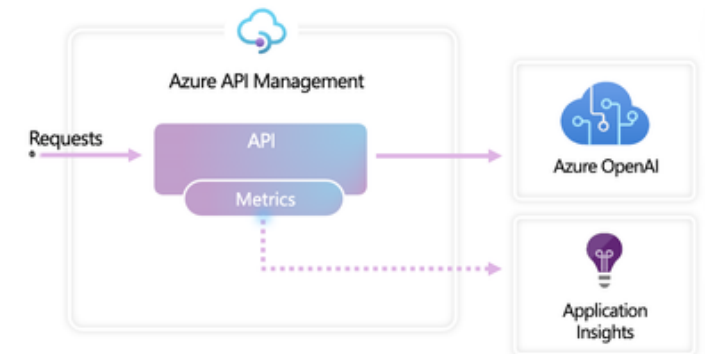
Configure tokens per minute (TPM) limits based on counter keys

Define similarity score threshold for caching

policy.xml

```
<azure-openai-token-limit
  counter-
  key="@context.Subscription.Id"
  tokens-per-minute="1000"
  estimate-prompt-tokens="false" />
```

Azure OpenAI Token Metric policy




Facilitate accurate cross-charging based on token consumption

Collect token usage data

policy.xml

```
<azure-openai-emit-token-metric
  namespace="AzureOpenAI">
  <dimension name="User ID" />
  <dimension name="Subscription ID" />
</azure-openai-emit-token-metric>
```



Generally Available

Azure API Center

Create a **centralized API inventory** for **discovery**, **consumption**, and **governance** regardless of API type, lifecycle stage, or deployment location.

Overcoming API Hurdles:

- **Fragmented API Inventory**
 - Poor API Discoverability and Reuse
 - Not all APIs are visible or managed
- **Compliance and Security Risks**
 - Inconsistent API design plus inability to enforce policies consistently
- **Reduced Efficiency and Productivity**
 - Difficulties in discovering relevant APIs across the organization

[Azure API Center: Your Comprehensive API Inventory and Governance Solution](#)

The screenshot displays the Microsoft Azure API Center interface. The top navigation bar includes the Microsoft Azure (Preview) logo, a search bar, and the Copilot icon. The main header shows the breadcrumb 'Home > api-center-microsoft-demo' and the title 'api-center-microsoft-demo | APIs' with a star icon. Below the header, there is a search bar, a 'Register an API' button, and a 'Refresh' button. The left sidebar contains a navigation menu with options: Overview, Activity log, Access control (IAM), Tags, Resource visualizer, Events, Assets (expanded), Metadata, APIs (selected), Environments, and Governance. The main content area shows a table of APIs with columns for API title, API type, and Lifecycle. The table lists several APIs: starwars api, climate api, cat facts api, uspto api, ebay api, and petstore api. Below the table, there is a section for the 'Contoso API portal' with a search bar and a list of APIs. The bottom section shows a filter by API type (REST selected) and a list of APIs: Demo Conference API, Climate, and Cat Facts.

API title	API type	Lifecycle
starwars api	REST	Production
climate api	REST	Development
cat facts api	REST	Design
uspto api	REST	Production
ebay api	REST	Production
petstore api	REST	Testing

API Center portal

Search for an API

Filter by

API type

- ☒ REST
- ☐ GraphQL
- ☐ gRPC
- ☐ SOAP
- ☐ Webhook
- ☐ WebSocket

Displaying 1 to 3 of 3 items

Sort by

API type: REST X Clear all

API | REST

Demo Conference API

API | REST

Climate

API | REST

Cat Facts

← To all hubs

Current hub
wm-test-ai-sweden

Hub overview

All projects

Get started

Model catalog

Model benchmarks

Prompt catalog

AI Services

Playgrounds

Chat

Assistants PREVIEW

Shared resources

Deployments

Connections

Compute instances

Users

Content filters PREVIEW

Integrate with generative AI



Speech analytics Preview

Transcribe audio and video recordings and generate enhanced outputs like summaries or extract valuable information such as key topics, Personal Identifiable Information (PII), sentiment and more.



Try out GPT-4 Turbo

GPT-4 Turbo can see, understand, and make inferences from visual inputs.



Prompt Shields for Generative AI Preview

Prompt Shield ensures your generative AI stays secure by detecting and blocking jailbreak attempts. Keep your community safe and trustworthy with this essential feature.



Summarize with generative AI

Get high-quality summarizations with a simple API call to simplify information at enterprise scale.

Infuse your solutions with AI capabilities



Speech

Enhance customer experiences through speech to text, text to speech, and speech translation features.

[View all Speech capabilities](#)



Language + Translator

Analyze, summarize and translate using LLM-powered natural language processing capabilities.

[View all Language + Translator capabilities](#)



Vision + Document

Discover information and insights from documents, images and video with OCR and multi-modal AI.

[View all Vision + Document capabilities](#)



Content Safety

Detect harmful, offensive, or inappropriate user-generated or AI-generated content in your app including text, image, and multi-modal APIs.

[View all Content Safety capabilities](#)

What's new

← AI Services



Current hub
wm-test-ai-sweden



Speech

Vision + Document

Language + Translator

Content Safety



Vision + Document

Give your apps the ability to read text, analyze images, process documents and detect faces with technology like optical character recognition (OCR) and machine learning.

Integrate Vision with generative AI

Try out GPT-4 Turbo

GPT-4 Turbo can see, understand, and make inferences from visual inputs.

[Open model catalogue](#)

Try out Phi-3-vision

Phi-3-vision is a cost-effective multi-modal small language model (SLM) with vision capabilities.

[Open model catalogue](#)

Try out GPT-4o

GPT-4o is OpenAI's most advanced model to date. This multimodal model handles both text and image inputs while generating text outputs.

[Open model catalogue](#)

View all other vision capabilities

Document Intelligence

Face

Image analysis

Prebuilt models for specific documents

Invoices

Extract invoice ID, customer details, vendor details, ship to, bill to, total tax, subtotal, line items and more.

[Try demo](#)

Receipts

Extract time and date of the transaction, merchant information, amounts of taxes, totals and more.

[Try demo](#)

Identity documents

Extract expiration name, dates, machine readable zone, and more from passports and US driver licenses.

[Try demo](#)

Health insurance cards

Extract information such as, medical network, member name, and deductible.

[Try demo](#)

US Tax forms

US Mortgage forms

Marriage certificates

Credit cards

Featured

Optical character recognition

Spatial analysis

Face

Image analysis



Video Retrieval and Summary

Preview

Generate a brief summary of the main points shown in video. Locate specific keywords and jump to the relevant section.

[Try it out](#)



Recognize products on shelves

Preview

Identify products on shelves, gaps in product availability, and compliance for planograms.

[Try it out](#)



Customize models with images

Preview

Create custom image classification and object detection models with images using Vision Studio and Azure ML.

[Start a project](#)



Search photos with image retrieval

Preview

Retrieve specific moments within your photo album. For example, you can search for: a wedding you attended last summer, your pet, or your favorite city.

[Try it out](#)



Add dense captions to images

Generate human-readable captions for all important objects detected in your image.

[Try it out](#)



Remove backgrounds from images

Preview

Easily remove the background and preserve foreground elements in your image.

[Try it out](#)



Add captions to images

Generate a human-readable sentence that describes the content of an image.

[Try it out](#)




Detect common objects in images

Recognize the location of objects of interest in an image and assign them a label.

[Try it out](#)

Sample image 2



Detected attributes JSON

Threshold
value 15

footwear (52.50%)

person (76.50%)

Laptop (52.30%)

seating (53.30%)

person (85.60%)

person (72.30%)

seating (67.80%)

table (61.30%)


Vision Studio > Add dense captions to images

Use one of your own files or choose from a sample below.


 Drag and drop a file here
 or
[Browse for a file](#)
 or
[Take a photo](#)



Sample image 2



Detected attributes JSON

- A woman wearing a mask sitting at a table with a laptop
- A person using a laptop
- A woman wearing a face mask
- A close-up of a green chair
- A person wearing a colorful cloth face mask
- A person using a laptop
- A woman wearing a colorful fabric face mask
- A close-up of a laptop on a table
- A woman wearing a face mask
- A woman wearing a face mask sitting at a table with a laptop

Speech Studio

Speech to text

Quickly and accurately transcribe in more than 100 languages and dialects. Enhance the accuracy of your transcriptions by creating a custom speech model that can handle domain-specific terminology, background noise, and accents. [Learn more about speech to text](#)

**Real-time speech to text**

Quickly test live transcription capabilities on your own audio without writing any code.

[Try out Real-time speech to text](#)

**Whisper Model in Azure OpenAI Service**

Quickly test live transcription capabilities on your own audio utilizing your Azure OpenAI resource and use prompts to improve the quality of the transcripts.

[Try out Whisper Model in Azure OpenAI Service](#)

**Batch speech to text**


Quickly test batch transcription capabilities to transcribe a large amount of audio in storage and receive results asynchronously using Azure Speech models or OpenAI Whisper model.

[Try out Batch speech to text](#)

**Custom Speech**

Add your own data and adapt to specific speaking styles, vocabulary, and more with a customized speech to text model.

[Start a Custom Speech project](#)

**Pronunciation Assessment with speech to text**

Get instant feedback on pronunciation accuracy and fluency by reading a script aloud.

[Try out Pronunciation Assessment](#)

**Speech Translation**

Translate speech into other languages of your choice with low latency.

[Try out Speech Translation](#)

Batch speech to text

Version 3.1

Batch speech to text enables you to transcribe a large amount of audio in storage. You can point to audio files with a shared access signature (SAS) URI and asynchronously receive transcription results.

[View documentation](#) [View sample code](#) [Use Speech CLI](#)

Platforms
[Docker container](#) [Cloud](#)

Try it out

Language of audio

Chinese (Cantonese, Traditional)

Model ⓘ

20231117

Model to compare (optional)

20240228 Whisper Large V2

[Show advanced options](#)

Use with your resource

☒ I acknowledge that this application uses the resource **wm-test-speech-sea** and will incur usage to my account. [Choose a different resource](#)

Upload audio files

Local file

Web URLs

Drag and drop files
Or
[Browse files...](#)

Batch jobs

20240504_101702 (2 models)

✓

20240504_101702 (2 models)

✓

d1.mp3

Clear all

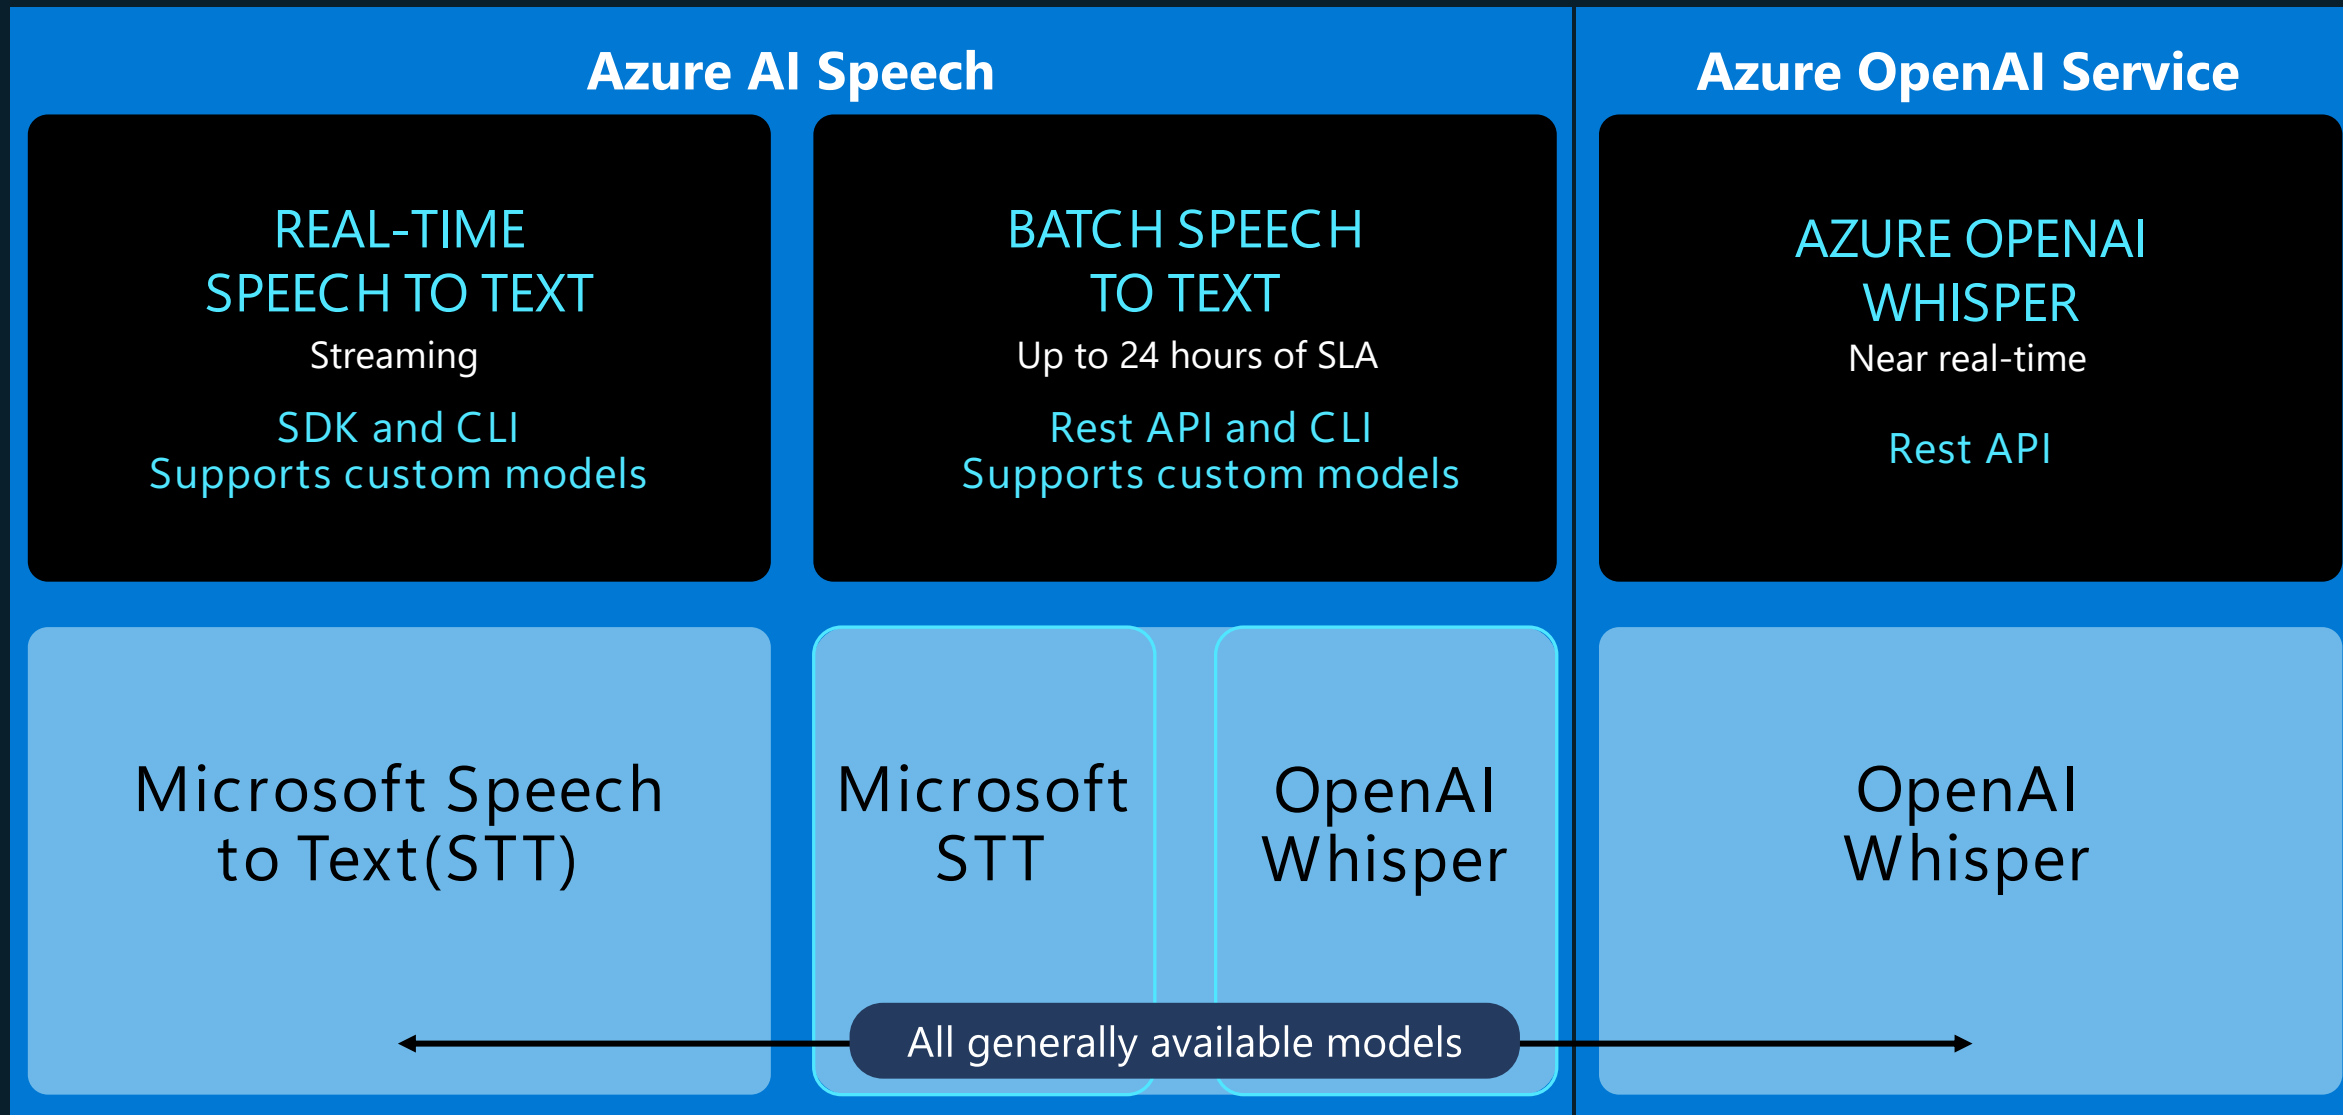
Developer resources

Power Platform connector

Audio from: <https://github.com/fcbond/hkcancor>

Azure AI has three differentiated speech to text offerings

All of them with generally available models





Get started with Document Intelligence Studio

Document Intelligence Studio is an online tool to visually explore, understand, train, and integrate features into your applications. You can easily extract key data from documents or create your own custom models with a no-code approach.

[Learn more](#)



Document analysis

Extract text, tables, structure, key-value pairs, and named entities from documents.



Read

Extract printed and handwritten text along with barcodes, formulas and font styles from images and documents.

[Try it out](#)



Layout

Extract tables, check boxes, and text from forms and documents.

[Try it out](#)



General documents

Extract key value pairs and structure like tables and selection marks from any form or document.

[Try it out](#)

Prebuilt models

Extract data from unique document types using the following prebuilt models.

← To all Vision + Document

Current hubwm-test-ai-sweden

- GPT-4 Vision

Phi-3-vision

GPT-4o

- Document Intelligence
- Invoices

Receipts

Identity documents

Health insurance cards

US Tax forms

US Mortgage forms

Marriage certificates

Credit cards

Contracts

Read
- Layout

- Face
- Detect faces in an image
- Image analysis
- Object detection

Image captioning

Dense captioning

Connected Azure AI Services

wm-test-ai-eastus-aiservices (eastus)

Connect to or create an Azure AI Services resource

Results

Drag & drop file here or [Browse for files](#)

Sample

layout-pageobject.pdf

Sample

layout-financial... rt.jpg

Sample

layout-checklist.jpg

Run analysis

Query fields

Analyze options

Layers

This is the header of the document.

This is title

1. Text

Latin refers to an ancient Italic language originating in the region of Latium in ancient Rome.

2. Page Objects

2.1 Table

Here's a sample table below, designed to be simple for easy understand and quick reference.

Name	Corp	Remark
Foo		
Bar	Microsoft	Dummy

Table 1: This is a dummy table

2.2. Figure

Figure 1: Here is a figure with text

3. Others

AI Document Intelligence is an AI service that applies advanced machine learning to extract text, key-value pairs, tables, and structures from documents automatically and accurately:

✓ Clean

✓ Precise

✓ Vague

✓ Coherent

✗ Incomprehensible

Turn documents into usable data and shift your focus to acting on information rather than compiling it. Start with prebuilt models or create custom models tailored to your documents both on premises and in the cloud with the AI Document Intelligence studio or SDK.

Learn how to accelerate your business processes by automating text extraction with AI Document Intelligence. This webinar features hands-on demos for key use cases such as document processing, knowledge mining, and industry-specific AI model customization.

This is the footer of the document.

11 Page

<

1

of 1

>

+

-

↺

↻

Content

Result

MarkdownText...

<!-- PageHeader="This is the header of the document." -->

This is title

1. Text

Latin refers to an ancient Italic language originating in the region of Latium in ancient Rome.

2. Page Objects

2.1 Table

- On this page
- Introduction

Try it out

Run the code

Run the code

To all Vision + Document

Current hub
wm-test-ai-sweden

GPT-4 Vision

Phi-3-vision

GPT-4o

Document Intelligence

Invoices

Receipts

Identity documents

Health insurance cards

US Tax forms

US Mortgage forms

Marriage certificates

Credit cards

Contracts

Read

Layout

Face

Detect faces in an image

Image analysis

Object detection

Image captioning

Dense captioning

layout-pageobject.pdf

Sample

layout-financial... rt.jpg

Sample

layout-checklist.jpg

Run the code

1. View sample code and

Select a resource to view

</> View code

2. Follow the quickStart

With your resource info

Start with SDK

3. Learn more

View documenta

Sample code

You can use the following code to start integrating your current capability and settings into your application.

https://wm-test-ai-eastus-aiservices1638064577.cognitiveservices.azure.com/

Python

```
1  # coding: utf-8
2
3  # -----
4  # Copyright (c) Microsoft Corporation. All rights reserved.
5  # Licensed under the MIT License. See License.txt in the project root for
6  # license information.
7  # -----
8
9  """
10 FILE: sample_analyze_layout.py
11
12 DESCRIPTION:
13     This sample demonstrates how to extract text, tables, figures, select
14     given through a file.
15
16 PREREQUISITES:
17     The following prerequisites are necessary to run the code. For more d
18
19     -----Python and IDE-----
20     1) Install Python 3.7 or later (https://www.python.org/), which should
```

Resource key

.....



Region

eastus



You should use environment variables or a secret management tool like Azure Key Vault to prevent accidental exposure of your key in applications. [Learn more](#)

Copy

Close

Get started with Content Safety Studio

Run moderation tests

Explore, try out, and view sample code for different types of content.



Moderate text content

Run moderation tests on text contents. Assess the test results with detected severities. Experiment with different threshold levels.

[Try it out](#)



Moderate image content

Run moderation tests on image contents. Assess the test results with detected severities. Experiment with different threshold levels.

[Try it out](#)



Moderate multimodal content

Region not supported

Run moderation tests on image and text combined contents. Assess the test results with detected severities.

Explore safety solutions for Gen-AI

Try out the latest capability for AI.



Groundedness detection

Groundedness detection detects ungroundedness generated by the large language models (LLMs).



Prompt Shields

Region not supported

Prompt Shields provides a unified API that addresses the following types of attacks:



Protected material detection

Region not supported

Use protected material detection to detect and protect third-party text material in LLM




Safety metaprompt

Use the framework of metaprompt that helps you potentially mitigate different types of harm.



Drag and drop an image here (image should be JPEG, PNG, BMP, TIFF, GIF or WEBP; max size: 4MB; max resolution: 2048x2048 pixels) or
[Browse for a file](#)

**Safe content**

An image of two children holding hands and smiling at sunset.

**Self-harm content**

An image of a girl holding a firearm with the intention of self-harm.

**AI-generated sexual content**

An image generated by AI model containing subtle adult content.

ON THIS PAGE[Introduction](#)[Try it out](#)[1. Select a sample or upload your own](#)[2. Test](#)[3. View results](#)[Next steps](#)**2. Test****Configure filters**[View code](#)**Image preview**☐ Blur image

Set the Severity thresholds for each category. Content with a severity level less than the threshold will be allowed. [Learn more about categories and threshold](#)

Category	Threshold level	↺
<input checked="" type="checkbox"/> Violence	Medium Allow Low / Block Medium and High	
<input checked="" type="checkbox"/> Self-harm	Medium Allow Low / Block Medium and High	
<input checked="" type="checkbox"/> Sexual	Medium Allow Low / Block Medium and High	
<input checked="" type="checkbox"/> ...	Medium	

Content Safety Studio > Moderate image content



 Run test

☒ Sexual

Allow Low / Block Medium and High

☒ Hate

Medium

Allow Low / Block Medium and High

ON THIS PAGE

Introduction

Try it out

1. Select a sample or upload your own

2. Test

3. View results

Next steps

3. View results

This content has been **Allowed**

All categories are accepted based on the severity thresholds in Configure filters they are currently set to.

Category and risk level detection results

The content will be annotated as Safe, Low, Medium or High.

Category	Severity level	Threshold	Judgement
Hate	Safe	Medium	Allowed
Violence	Safe	Medium	Allowed
Sexual	Safe	Medium	Allowed
Self-harm	Safe	Medium	Allowed

Customer Copyright Commitment – the Elevator pitch

Using Microsoft AI? We have your back if you get sued for IP infringement*

- If you get sued for using our enterprise copilot products, we'll defend the claim and pay the judgment or settlement.
- *There are some qualifiers:
 - You have to use our risk mitigation tools – either built-in to copilots, or described [here](#) for 'build your own' AI apps
 - Treat it as a copilot – you shouldn't use output that you know or ought to know will infringe a third party
 - Your content shouldn't cause the claim – make sure of your training data
 - We **won't** cover trademark claims for Copilot created content that you use in commerce or the course of trade.



Data Sets and Models



Microsoft Research
Open Data



Azure Open
Datasets



Planetary
Computer



Industry Data
Models



Frameworks, Toolkits and SDK

Semantic Kernel



State
Memories



Side Effects
Plugins




AI UX
Copilot



Models
OpenAI & Azure
AI Models

Easily build agents that can call your existing code. As a highly extensible, lightweight SDK, you can use Semantic Kernel with models from OpenAI, Azure OpenAI, Hugging Face, and more.

AutoGen



Develop LLM applications using multiple agents, LLMs and humans, that converse with each other to solve tasks.

Bot Framework



A collection of libraries, tools, and services that let you build, test, deploy, and manage omnichannel bots.



Data Science Toolkit

Open-source collection of proven ML and AI accelerators, enabling automation of commonly repeated development processes.

Responsible AI Toolbox



Tools providing model and data exploration and assessment user interfaces and libraries that enable a better understanding of AI systems

HAX Toolkit



For teams building user-facing AI products to conceptualize what the AI system will do and its behaviours. Use early in design process.



Get started

These issues have been labeled as a great place to start in a project.

📁 microsoft/retina

issue labeled about 11 hours ago

ideally agent's init container enables telemetry based on ConfigMap flag

[Good First Issue](#) [Help Wanted](#) [Lang/Go](#) [Type/Fix](#) [Priority/1](#) [Scope/\\$](#)

📁 Azure-Samples/azure-search-openai-demo-java

issue labeled about 19 hours ago

Triggering the AOAI content filter results in a "Type error" in the frontend

[Enhancement](#) [Good First Issue](#) [Stale](#)

📁 Azure-Samples/azure-search-openai-demo-java

issue labeled about 19 hours ago

Convert main.parameters.json to main.bicepparam

[Enhancement](#) [Good First Issue](#)

📁 Azure-Samples/azure-search-openai-demo-java

issue labeled about 19 hours ago

Convert main.parameters.json to main.bicepparam

Resources

← ↻ 🏠 https://www.ds-toolkit.com

Discover your toolkit

The data science toolkit is an open-source collection of proven ML and AI implementation accelerators. Accelerators enable the automation of commonly repeated development processes to allow data science practitioners focus on delivering complex business value and spend less time on basic setup.

I am looking for



Filter

[Clear all](#)

Use the Data Science Toolkit to

- ☐ Enable DevOps and Monitoring (5)
- ☐ Get started with a specific use-case (19)
- ☐ See examples of use-cases that could help me (10)
- ☐ Set up an environment (2)
- ☐ Start a new project (6)

Asset Type

- ☐ Accelerator (37)
- ☐ Demo (11)

Sort by: [Name](#) ▼

GET STARTED WITH A SPECIFIC USE-CASE

Anomaly Detection

Anomaly Detection using the iJungle technique to detect anomalies in very large datasets.

Created: Mar. 2, 2023 by [RC](#) [GP](#) [🌐](#)

Asset Type: Accelerator

[See more](#)

SET UP AN ENVIRONMENT

Azure Data Labs

Quickly deploy Azure Data services using Terraform and GitHub Actions.

Created: Mar. 2, 2023 by [WG](#) [🌐](#) [🌐](#)

Asset Type: Accelerator

[See more](#)

GET STARTED WITH A SPECIFIC USE-CASE

Azure Object Detection

Custom computer vision accelerator, with rapid deployment and configuration for cloud and edge scenarios.

Created: Mar. 2, 2023 by [🌐](#) [🌐](#) [🌐](#)

Asset Type: Accelerator

[See more](#)

SEE EXAMPLES OF USE-CASES THAT COULD HELP ME

Azure OpenAI App Demo

Flask Application demonstrating various use cases of Azure Open AI

GET STARTED WITH A SPECIFIC USE-CASE

Binary Classification

Scalable and efficient binary classification

SEE EXAMPLES OF USE-CASES THAT COULD HELP ME

Brainstormer

An OpenAI model empowered brainstorming assistant: fueling data



Prompty

Public preview

- A language agnostic prompt asset for creating prompts and engineering the responses
- Uses common markdown format
- Supported by popular orchestration frameworks LangChain, Semantic Kernel, and Prompt Flow:
 - Python (Prompt Flow and LangChain)
 - C# (Semantic Kernel)
 - (Future Work) works in TypeScript/JavaScript
- Understood in Azure AI Studio
- Support direct run in VS Code with extension

<https://prompty.ai/>

```
1  ---
2  name: Basic Prompt
3  description: A basic prompt that uses the GPT-3 chat API to answer questions
4  authors:
5    - sethjuarez
6    - jietong
7  model:
8    api: chat
9    configuration:
10      azure_deployment: gpt-35-turbo
11  sample:
12    firstName: Jane
13    lastName: Doe
14    question: What is the meaning of life?
15  ---
16  system:
17    You are an AI assistant who helps people find information.
18    As the assistant, you answer questions briefly, succinctly,
19    and in a personable manner using markdown and even add some
20    personal flair with appropriate emojis.
21
22  # Customer
23    You are helping {{firstName}} {{lastName}} to find answers to their questions.
24    Use their name to address them in your responses.
25
26  user:
27    {{question}}
```

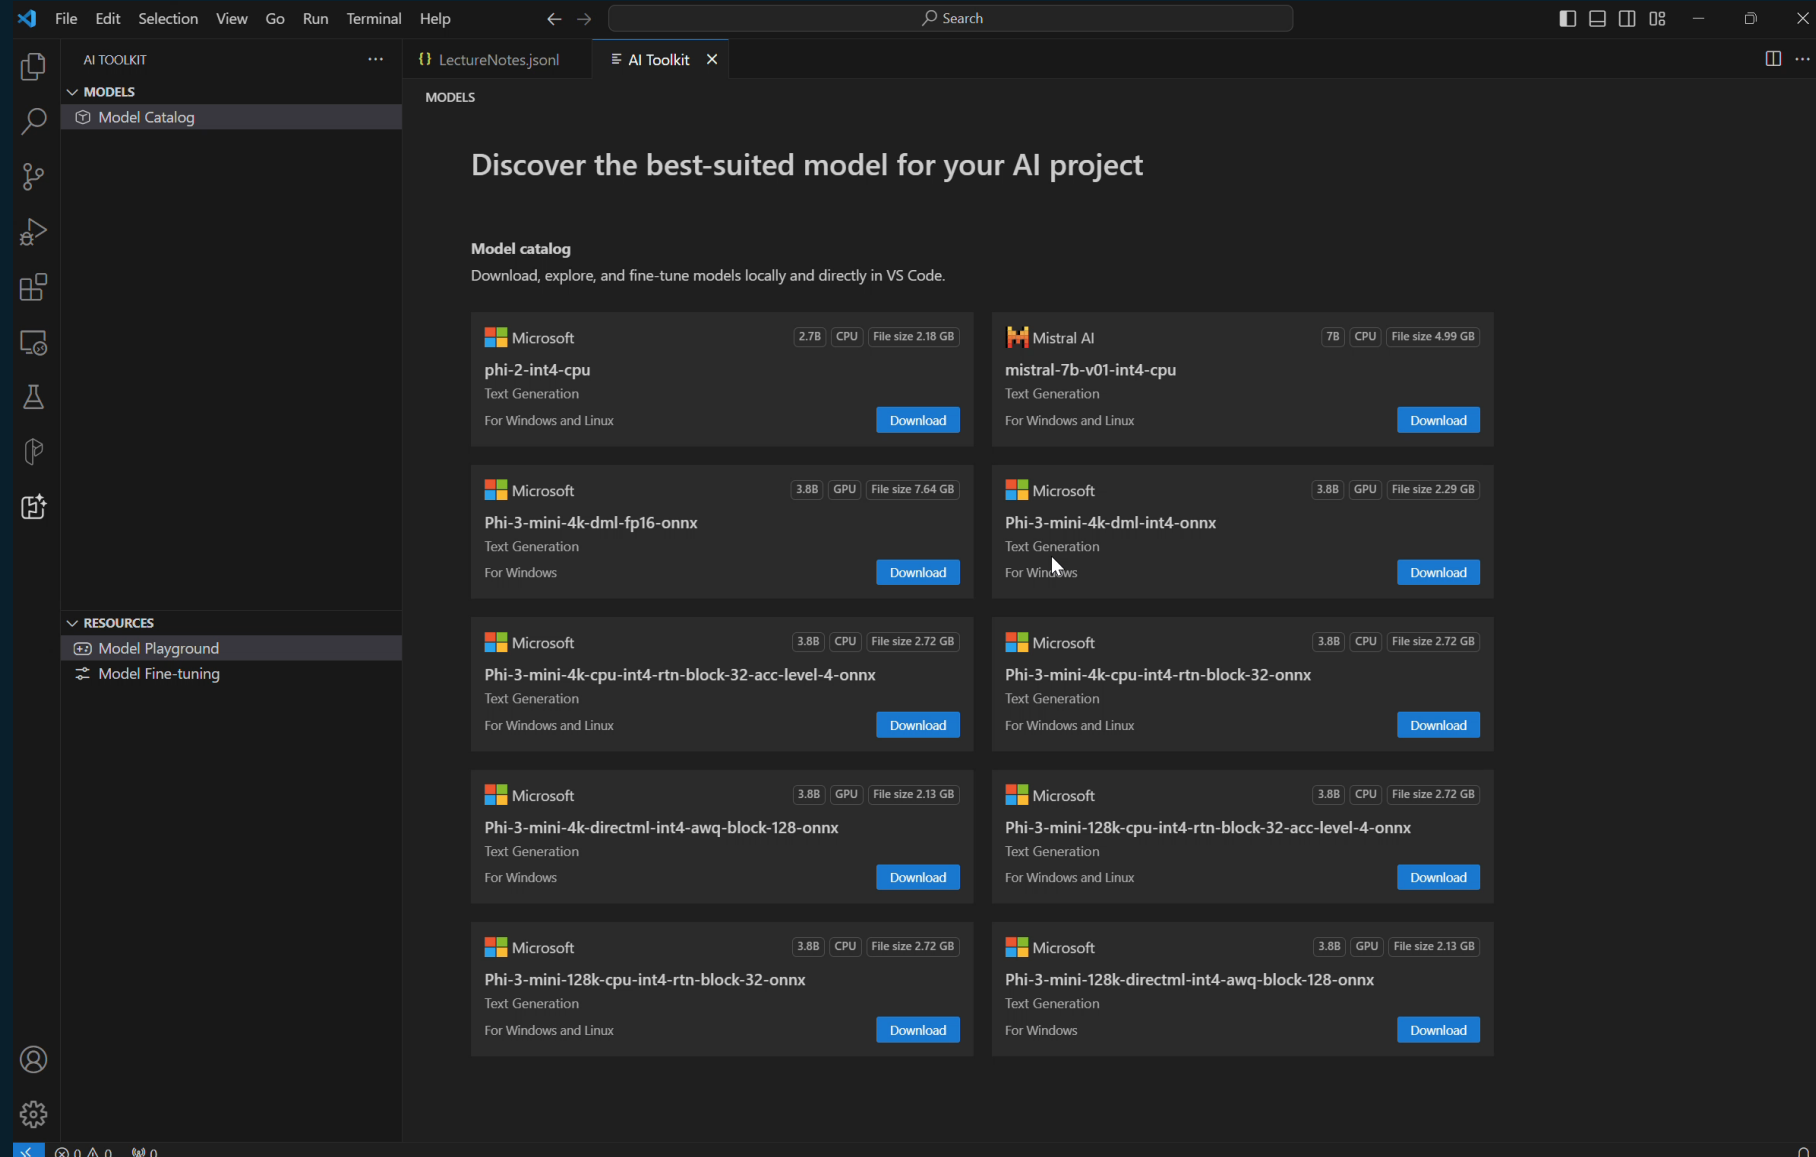
Fine Tuning with AI Toolkit for VS Code

Public preview

Available as a VS Code extension

Fine-tune locally on device or offload to the cloud

Generate ready to use fine-tuned models for your application



Empower every part of your organization with AI

Out-of-the-box AI

Just activate
Unlock productivity
Persona-based Copilots

Users Business Units SOC IT

Copilots

Copilot for M365

Copilot for Dynamics 365



Copilot in Azure



GitHub Copilot



Copilot for Security

Composable AI

Deploy in days
No code to Pro code
Democratize innovation

Makers

Copilot Studio



Copilot in Power Platform
Copilot in Fabric

AI Builder

Dataverse

Developers

Semantic Kernel SDK

Azure AI Studio

Azure OpenAI service

GitHub Copilot

Azure Application Platform

Model-as-a-Service

Azure AI Services

Pretrained Models

Build-your-own AI

Tailored models
Develop monetizable IP
Maximize data science skills

Data Scientists



Azure Machine Learning



Azure Databricks



Azure Prompt Flow

Microsoft Fabric

Model Catalog

Fine Tuning

Best in Class Foundation Models

Unified Data Platform

Azure AI Infrastructure

Time to Value

Differentiation



Thank you.