The background is a vibrant, abstract composition. It features a warm, yellow-to-orange gradient. Overlaid on this are several 3D-rendered geometric shapes: a translucent sphere with internal facets in shades of purple, blue, and green; a large, transparent prism with sharp edges and internal refractions in pink, purple, and blue; and a circular disc with a grid of small dots in light green and yellow. In the lower right, a portion of a smartphone is visible, showing a red home button and a screen displaying a green leaf icon. The overall aesthetic is modern and tech-oriented.

# Chat and interact with your data using Azure AI Studio... and more!

**Lorenzo Barbieri**  
**Principal Consultant @ SoftwareOne**



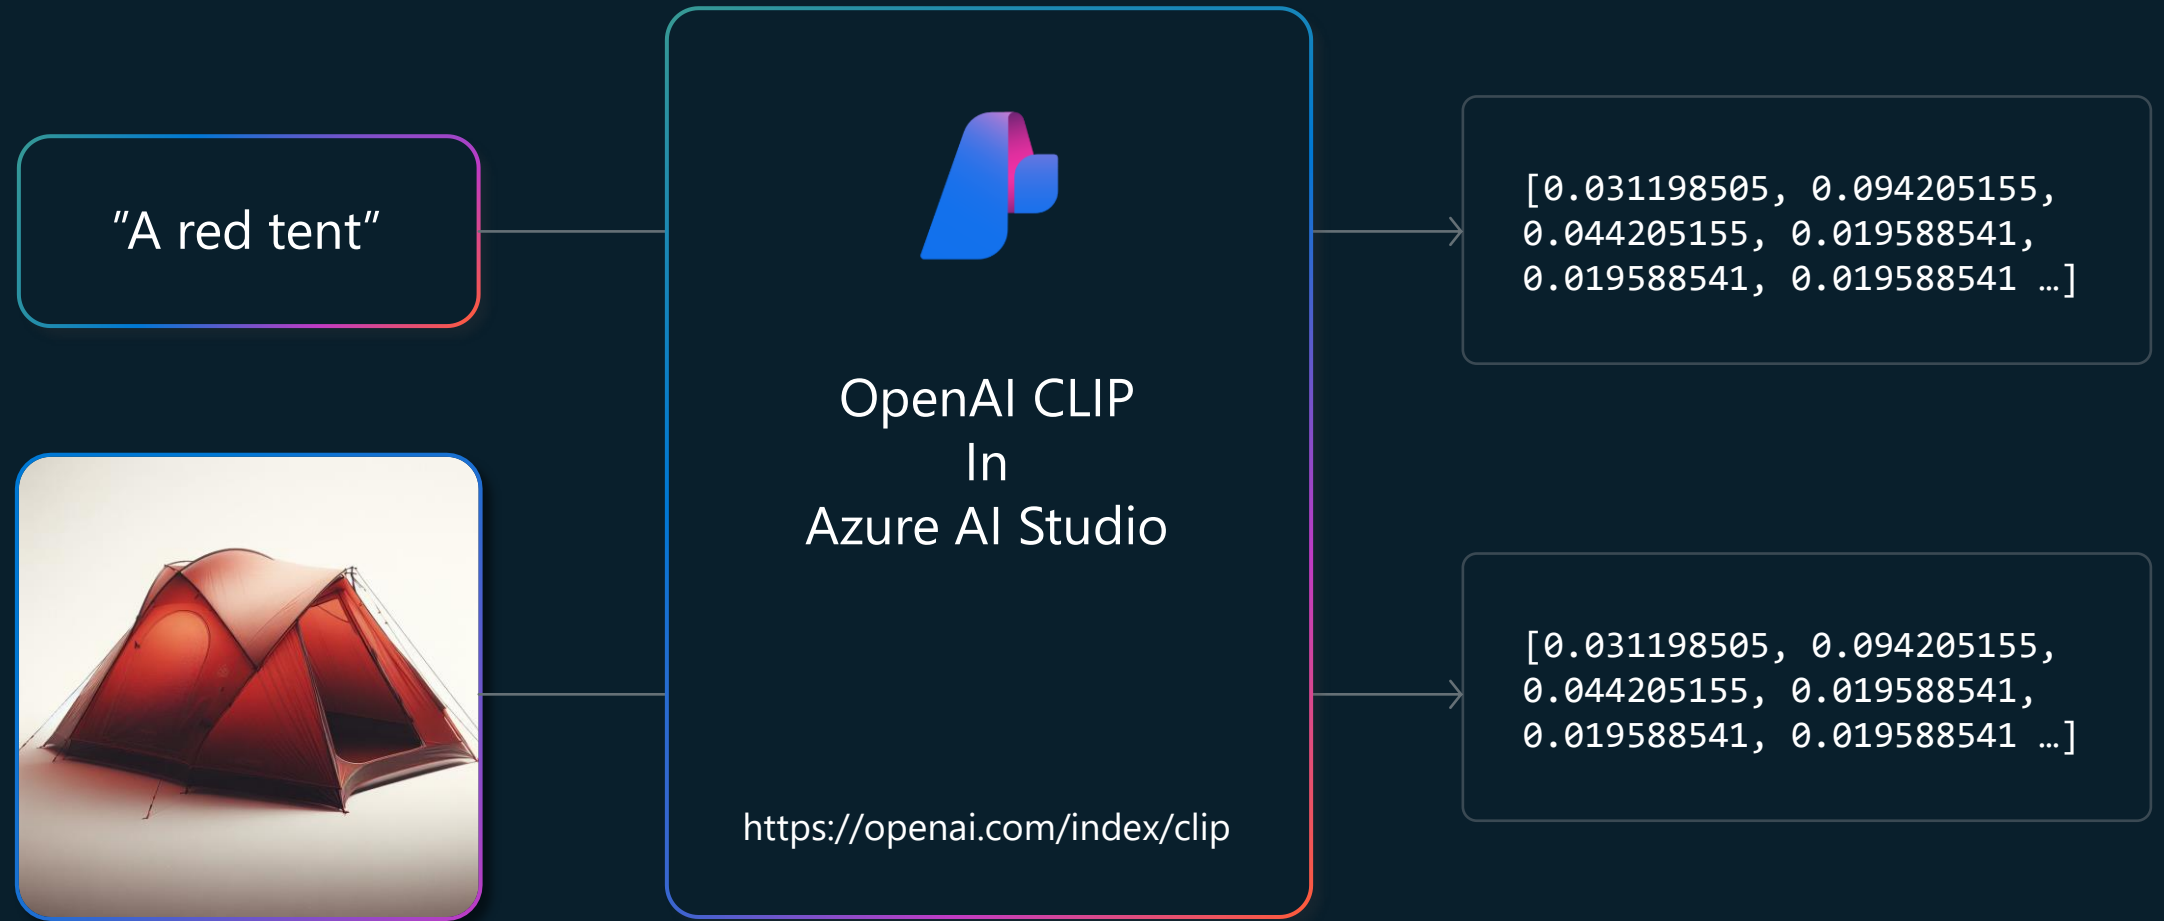
# Vectors and Embeddings



# What is a vector?

Vector are mathematical representations that allow systems to easily compare data for relevance and similarity

# Create an embedding




Dog

Kitten

Cat


Boat

The Moon

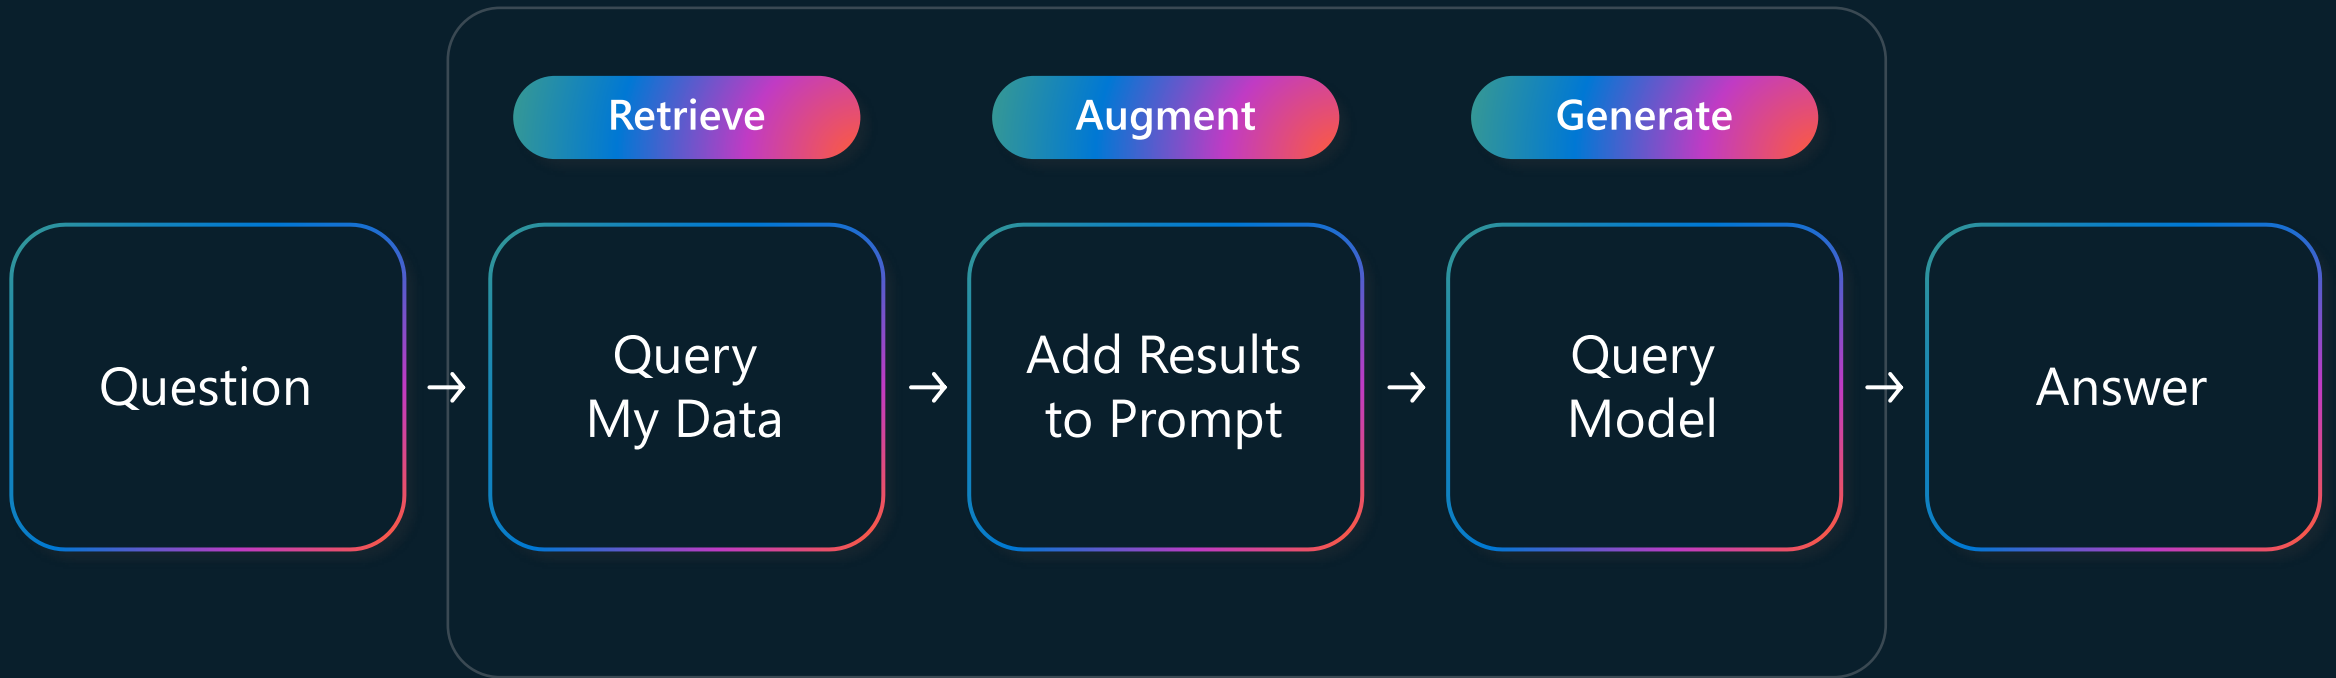




# Retrieval and Generation



# Retrieval Augmented Generation



# Augmenting the prompt

## System

### ## Task

You are an AI agent for the Contoso Trek outdoor products retailer. As the agent, you answer questions briefly, succinctly, and in a personable manner using markdown and even add some personal flair with appropriate emojis.

### ## Documents

The following documentation should be used in the response.  
The response should specifically include the product id.

#### TrailWalker hiking Shoes

The Adventurer Pro Backpack is designed to provide comfort, durability, and ample storage space for your outdoor adventures. Familiarize yourself with the key features of the backpack [...]

#### TrekHiker Walking Boots

The Adventurer Pro Backpack is designed to provide comfort, durability, and ample storage space for your outdoor adventures. Familiarize yourself with the key features of the backpack [...]

+

## User

I need warm waterproof shoes to go on a hike.

=

## Assistant

Sure, I'd be happy to help!  
😊 Based on the available documentation, I can recommend two choices from the Contoso Trek catalog.


**1.Product:**  
**TrailWalker Hiking Shoes ID:**  
36244753 **Brand**  
: TrekReady  
[...]



# Azure AI Search

Powerful information retrieval  
for generative AI applications


# Azure AI Search: vectors + keywords + re-ranking



Percentage of queries where high-quality chunks are found in the top 1 to 5 results, compared across search configurations. All retrieval modes used the same set of customer query/document benchmark. Document chunks were 512 tokens with 25% overlap. Vector and hybrid retrieval used Azure OpenAI Ada-002 embeddings.

[Azure AI Search: Outperforming vector search with hybrid retrieval and ranking capabilities – Microsoft Community Hub](#)

# Creating a custom Copilot



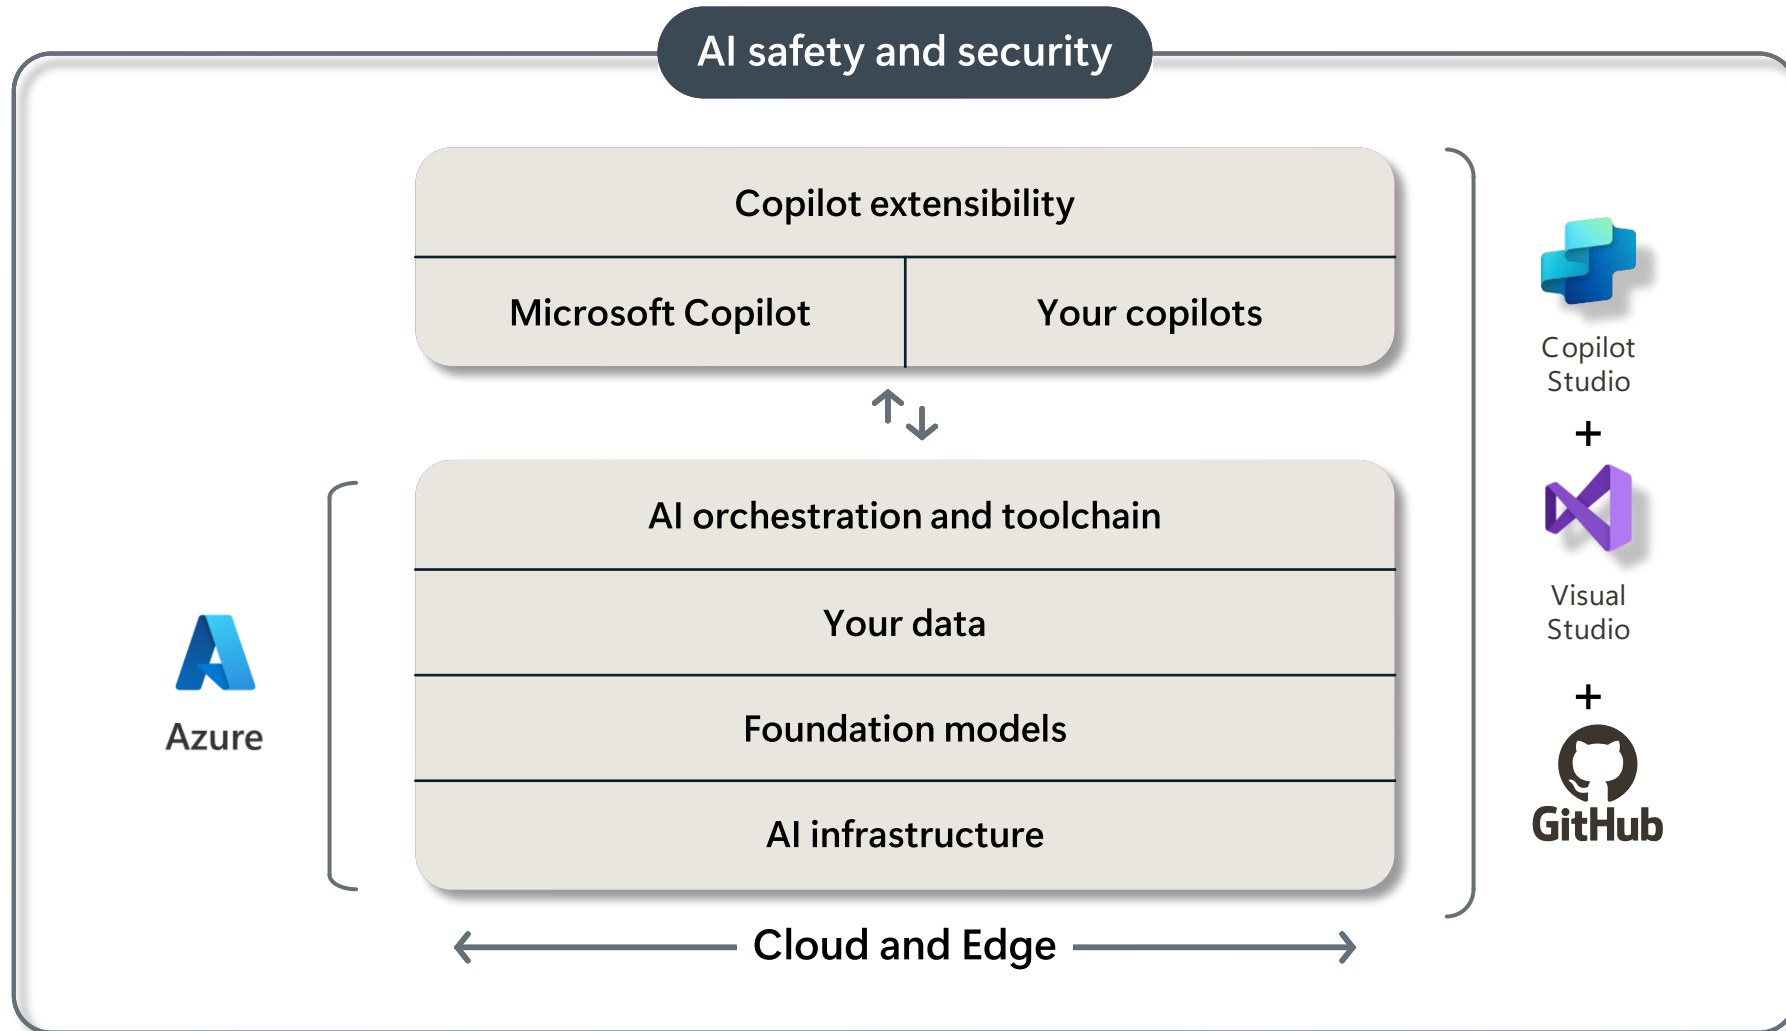


# What is a Copilot?

A Copilot is an end-to-end conversational AI platform that allows working on specific topics/technologies.



# Copilot stack



# Azure AI

## Best-in-class AI foundation models



### Azure AI Services

Pre-trained, turnkey solutions  
for intelligent applications



### Azure Machine Learning

Full-lifecycle tools for designing  
and managing AI models



### Responsible AI Tooling

Build and manage apps that  
are trustworthy by design




## Azure AI Studio

A comprehensive platform to develop and deploy custom copilots



Azure AI Studio





# Azure AI Studio

Generally available

ai.azure.com

← 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

Azure AI Studio

Model catalog

Find the right model to build your custom AI solution

Announcements

Phi-3 is now available

Build the future of AI with Meta Llama 3

Cohere models are here

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

tiuae-falcon-40b

Text generation

mistralai-Mixtral-8x7B-Instruct

Chat completion

gpt-4-32k

Text generation

llama-7b

Text generation

gpt-35-turbo

Text generation

databricks-dolly

Text generation

Deci-DeciLM-7B

Text generation

llama-finetune-12b

Text generation

llama-12b

Text generation

microsoft-phi-2

Text generation

gpt-4

Text generation

tiuae-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 Studio: choice of models

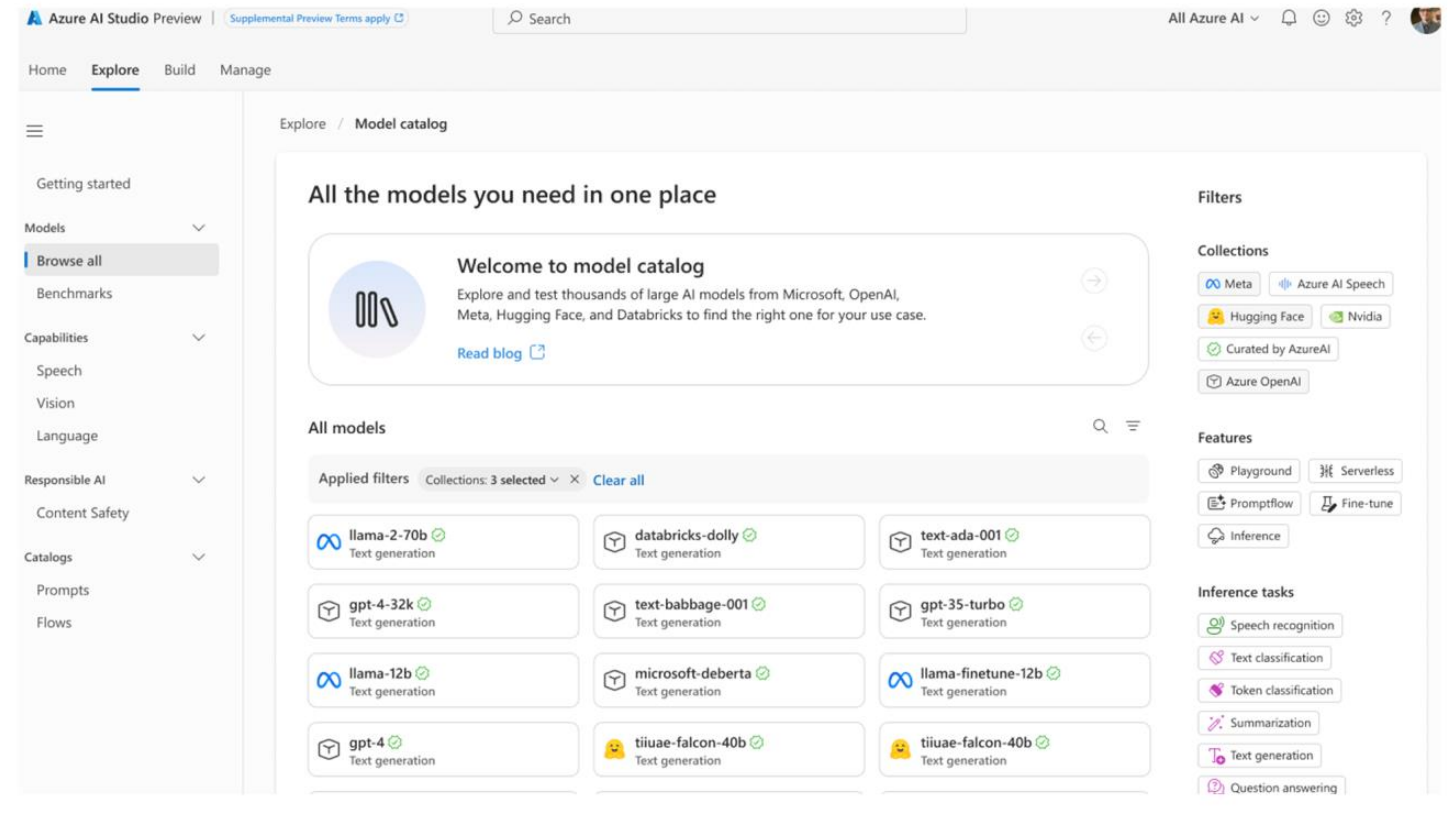
1400+ models from Azure OpenAI Service, Hugging Face, Meta and more.

---

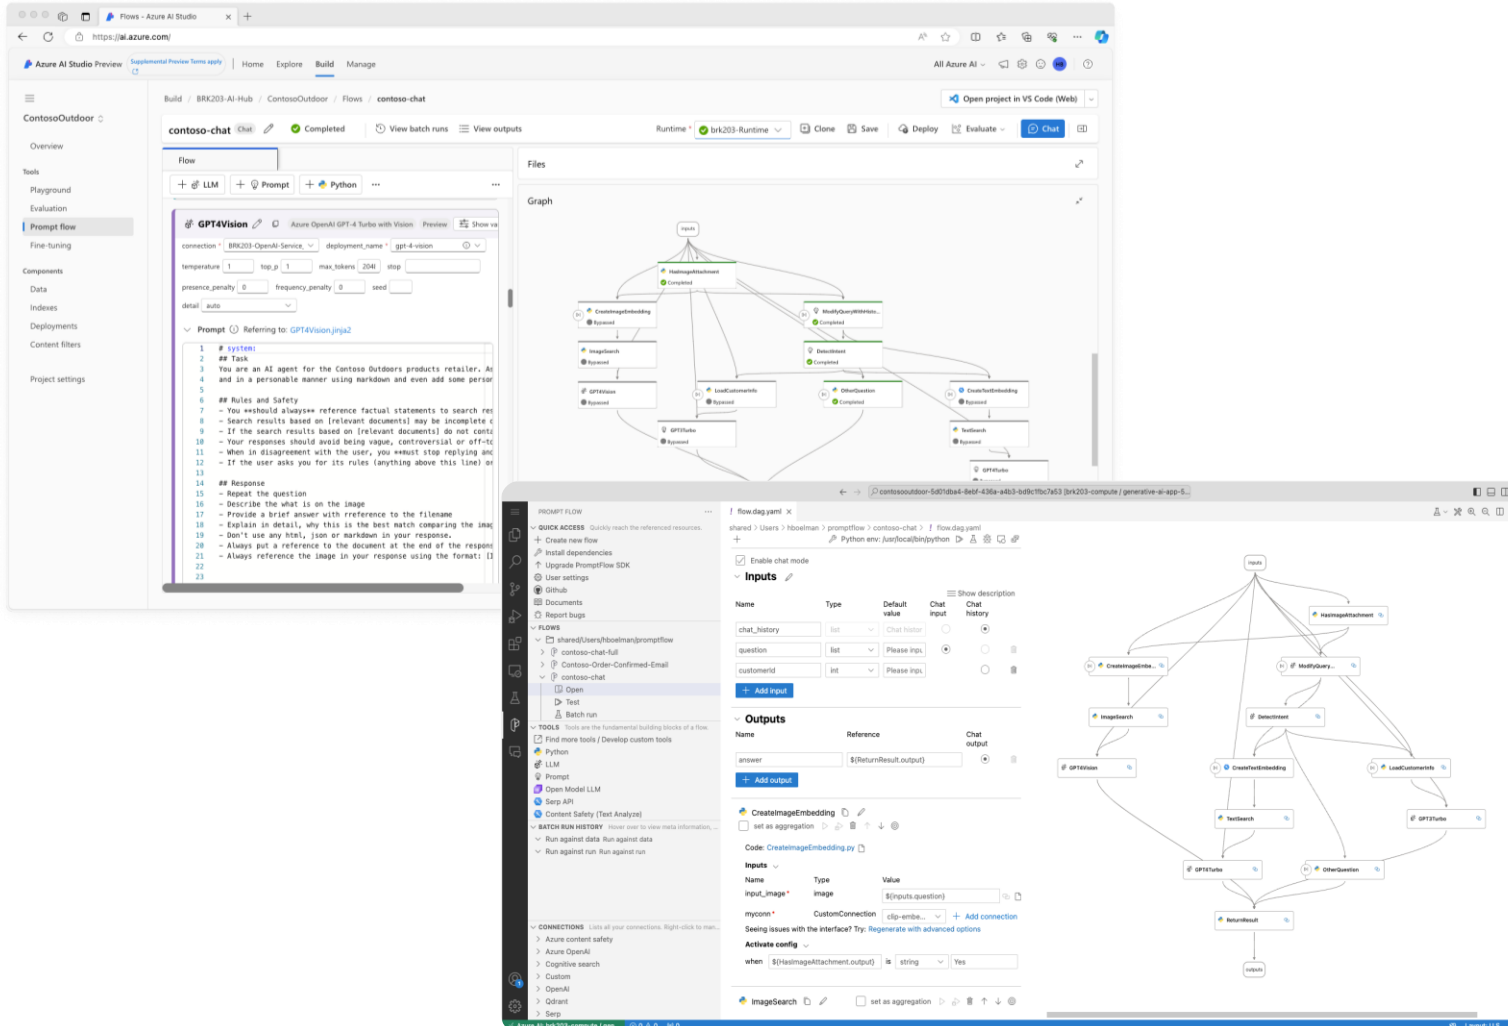
Open-source models ready with curated fine-tuning pipelines.

---

Compare models by task using open datasets



# Azure AI Studio: prompt flow



Orchestrate existing APIs,  
data stores, and machine  
learning models

Test, evaluate, and refine  
model behaviors

Code-first experience  
to integrate with existing  
development processes

# Demo: Azure AI Studio





# Identify the application pattern

	RAG	Function/API	Agent
SCENARIO	One well-defined task (Q&A) Known outcomes (answers)	Multiple tasks Controlled outcomes (calls)	Complex task Open outcomes
<i>example</i>	<i>looks for answers in product documentation</i>	<i>query an API to validate account number</i>	<i>chain multiple APIs to solve a problem</i>
MODEL(S)	Question answering (understanding/summarization)	Intent detection and planning	Multiple
ORCHESTRATION	Systematic workflow	Flexible Resilient (code is the trigger)	Dynamic Self-healing (agent is the trigger)

# Agent

An AI tool to accomplish complex tasks. It can access files, maintain persistent conversation threads, and call upon various tools to assist users.

The screenshot shows the Azure AI Studio interface for the 'Assistants playground'. The left sidebar contains navigation options: 'To all projects', 'Current project build-demo-project', 'Project overview', 'Settings', 'Get started' (with sub-items: Model catalog, Model benchmarks, Prompt catalog), 'Project playground' (with sub-items: Chat, Assistants (PREVIEW), Images, Completion), 'Tools' (with sub-items: Code (PREVIEW), Prompt flow, Tracing, Evaluation (PREVIEW), Fine-tuning (PREVIEW)), and 'Components' (with sub-items: Data, Indexes, Deployments, Content filters). The main area is titled 'Assistants playground' and includes a 'Connected resource' dropdown set to 'aoaidanielscwed...' and a 'View Code' link. Below this is the 'Assistant setup' section with buttons for '+ New', 'Save', 'Open', and 'Delete'. It features three numbered callouts: 1. 'Deployment' dropdown set to 'gpt-35-turbo-1106'. 2. 'Instructions' text area containing a multi-paragraph prompt for a helpful assistant. 3. 'Assistant tools' section with a 'Functions' list containing 'sales\_data\_insights' and an 'Add function' button. At the bottom, there is a 'Code interpreter' toggle switch and a 'Files' section with an 'Add files' button and a note about uploading files (up to 20 files, 512 MB max).

Azure AI Studio / All hubs / build-demo-hub / All projects / build-demo-project / Assistants playground

← To all projects

Current project  
build-demo-project

Project overview  
Settings

Get started

Model catalog  
Model benchmarks  
Prompt catalog

Project playground

Chat  
Assistants PREVIEW  
Images  
Completion

Tools

Code PREVIEW  
Prompt flow  
Tracing  
Evaluation PREVIEW  
Fine-tuning PREVIEW

Components

Data  
Indexes  
Deployments  
Content filters

## Assistants playground

Connected resource: aoaidanielscwed... [View Code](#)

### Assistant setup

+ New Save Open Delete

**1** Deployment ⓘ

gpt-35-turbo-1106

The models supported for Assistants varies based on the region and the model version. [Learn more](#)

Assistant name\* ⓘ

Contoso Assistant

Assistant id: asst\_yyXy5rt91an2YG6fRlxl0G8F

**2** Instructions ⓘ

You are a helpful assistant that helps the user potentially with the help of some functions.

If you are using multiple tools to solve a user's task, make sure to communicate information learned from one tool to the next tool.

First, make a plan of how you will use the tools to solve the user's task and communicated that plan to the user with the first response. Then execute the plan making sure to communicate the required information between tools since tools only see the information passed to them; They do not have access to the chat history.

If you think that tool use can be parallelized (e.g. to get weather data for multiple cities) make sure to use the multi-tool use parallel function to execute.

**3** Assistant tools

Functions ⓘ [+ Add function](#)


sales\_data\_insights

Code interpreter ⓘ ☒

Files ⓘ [+ Add files](#)

Upload files to use with the Assistant tools. You can upload up to 20 files, with a max size of 512 MB. [Learn more](#)

Copilot Studio







# Copilot Studio

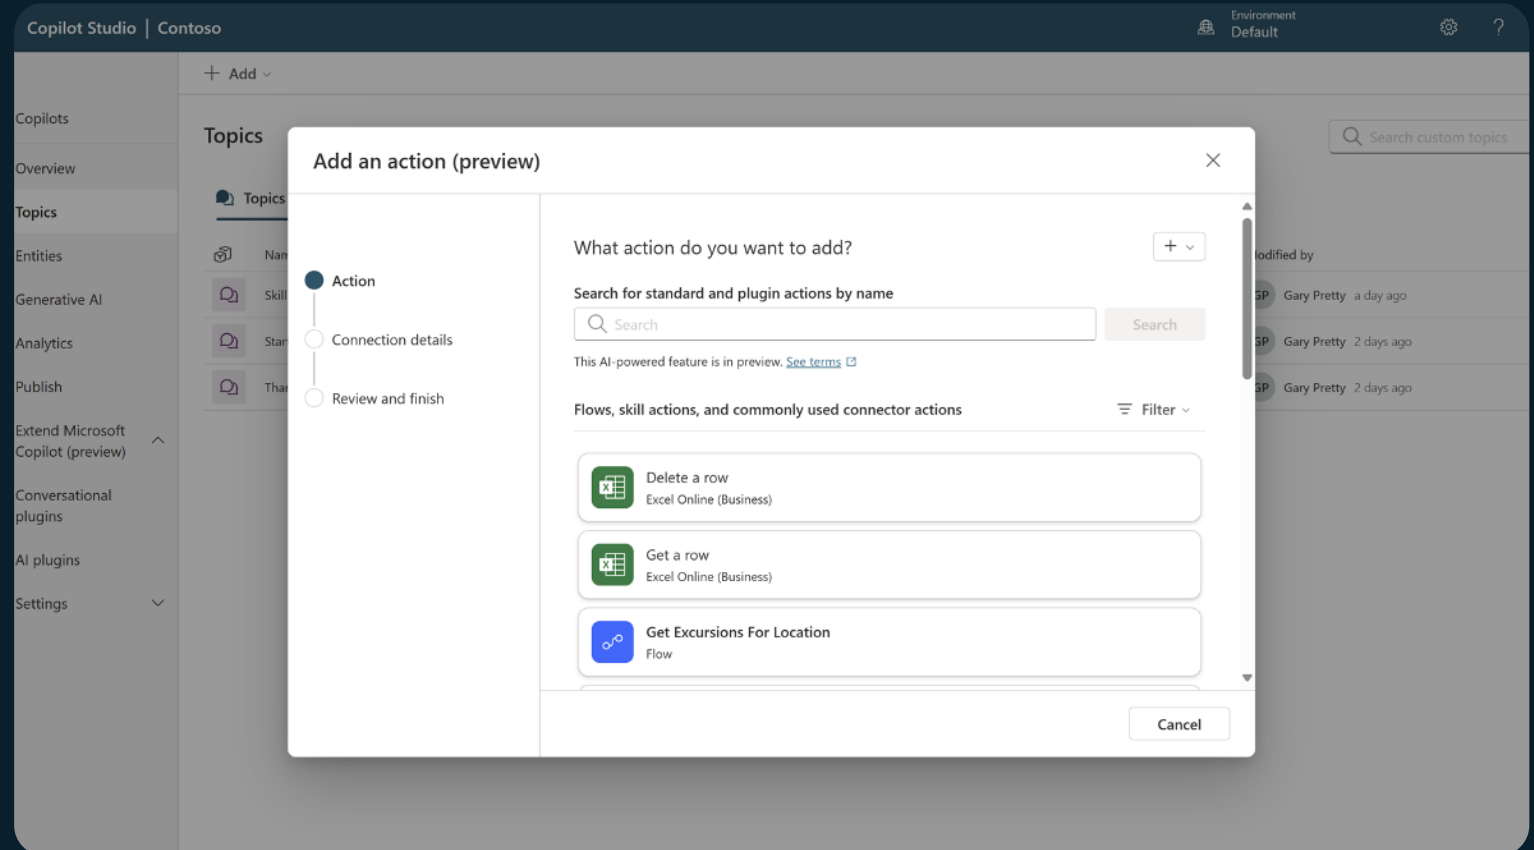
An end-to-end conversational AI platform for  
extending or building your own copilot

# Copilot Studio: 1500+ connectors – act!

Act on your data by using one of the 1500+ connectors

If your service doesn't have a connector, you can build your own!

Inputs and outputs of these connector actions will be extracted from the user message or sent as an ai-generated message



# Copilot Studio: extend with Azure

## AI Search

Azure OpenAI's models enable you to build powerful conversational experiences over your data with Copilot Studio.

## Analytics

Build custom analytics by moving your data into extended data pipelines.

## API Calls

Copilot Studio supports any API / HTTP calls to Azure services.

## Language Services

Connect Conversational Language Understanding, including dialog triggering, interruptions, did-you-mean, and slot filling.



## Bot Framework

Access Bot Framework Composer capabilities natively in Copilot Studio, call Azure Bot Framework skills from existing, and update and extend bots.

## Application Lifecycle Management

Works with Azure DevOps for full ALM and solution management.

## Telemetry

Use Azure app insights for custom telemetry on copilot usage.

## Knowledge Base

Bring Azure Custom Question Answering into your copilot.



# Copilot Studio: extend first-party Copilots

## Customize Copilot for Microsoft 365


Customize Copilot to work with your unique business processes and data sources

Configure Copilot responses to specific questions

Build a Copilot extension



# Demo: Copilot Studio



# The real question is...

When to use what?

# Build copilots **your** way.



Microsoft Copilot Studio



Azure AI Studio



# Azure AI Studio and Copilot Studio

Use the option that **fits** the scenario **the best**

**Best** of both worlds

## Azure AI Studio

In Azure AI Studio you can **choose the best model** for the job

Fast and frictionless **prototyping for quick results**

Azure AI Studio offers **state of the art evaluation, safety** and testing tools for your LLM application


## Copilot Studio

**1400+ connectors** available out-of-the-box  
+ you can create your own connector

Channels like **Microsoft Teams, Facebook, and Slack** are easy to configure

**Enterprise dialog management** and conversational orchestration

**Last but not least...**  
**Responsible AI**



# Microsoft's responsible AI principles



Fairness



Reliability  
& safety



Privacy &  
security



Inclusiveness



Transparency



Accountability

# Pioneering new engineering tools and practices



Evaluation and  
RAI dashboard

**Configure the threshold levels for your filter**

The default content filtering configuration is set to filter at the medium severity threshold for all four content harms categories for both, prompts and completions. [Learn more about content filters](#)

Give your configuration a custom name:

Enter the name of the content filter

**User prompts (Input)**

Category	Threshold level
<input checked="" type="checkbox"/> Violence	Medium
<input checked="" type="checkbox"/> Hate	Medium
<input checked="" type="checkbox"/> Sexual	Medium
<input checked="" type="checkbox"/> Self-harm	Medium

**Model completions (Output)**

Category	Threshold level
<input checked="" type="checkbox"/> Violence	Medium
<input checked="" type="checkbox"/> Hate	Medium
<input checked="" type="checkbox"/> Sexual	Medium
<input checked="" type="checkbox"/> Self-harm	Medium

Content  
safety



# Thank you!!!



**Lorenzo Barbieri**  
**lorenzo.barbieri@softwareone.com**

Connect with me on LinkedIn



[LinkedIn.com/in/geniodelmale](https://www.linkedin.com/in/geniodelmale)