AI Momentum Summit

Stephen Kaufman
Chief Architect – Microsoft
April 2025

Generative AI trends

93%

organizations are experimenting with multiple models[1]

61%

people are wary about trusting AI systems[3]

50%

generative AI will launch agentic AI pilots or POC by 2027[2]

30%

or fewer generative AI experiments moved to production by most respondents[4]

1. 16 Changes to the Way Enterprises Are Building and Buying Generative AI | Andreessen Horowitz
2. Autonomous generative AI agents | Deloitte Insights
3. Trust in artificial intelligence – 2023 Global study on the shifting public perceptions of AI, KPMG
4. GenAI and the future enterprise | Deloitte Insights

Personas of Teams Development

Code First Developer
(Professional Developer)

I'm a full stack developer and I'm responsible for designing software and developing applications with programming languages. I need to collaborate with business users to resolve business challenges.

My responsibility:

- Develop full stack applications with programming languages
- Co-developing the business with business users and IT professionals
- Make iterations quickly to meet everchanging business needs

IT Professional

I'm an IT manager and I'm responsible for everything related to maintenance, governance and security of my company's IT environment. I need to collaborate with the professional developers and business team to keep the IT environment updated and safe.

My responsibility:

- Maintain the end-to-end environment of application development
- Take control of governance, security and identity
- Make sure that all data are managed properly and securely

Low-Code Developer
(Citizen Developer)

I'm an end user from business group. I work closely with my business leaders and finance analysts. Although I have no coding background, I understand the business needs well and I know where the pain point is from my business.

My responsibility:

- Meet business demand and resolve business challenges
- Co-developing the business with professional developers and IT professionals

AI tools and copilots designed to benefit everyone at every level, in every organization

Business users	No code	Low code	Pro code
Expedite tasks for everyday users	Enhance deliverables through AI integration	Build complexity through simplicity	Design custom outputs and solutions

AI Alignment Guide

I want a **generative AI solution**

I want **out-of-the-box solutions** that work with my **existing data estate** for **my employees**

I want insights and actions for specific roles that integrate with **existing system**

I want insights and actions on **M365 data** and **plugins**

I want to **build a solution** with **custom data and UI,** and deploy **internally or externally**

I want to **customize agents** with **natural languages** and use a **generative orchestrator**

I want **full control, choice of model,** and **customize with code**

	Persona-Based Copilots	Enhance	Microsoft 365 Copilot Microsoft 365 Copilot Chat*	Extend	Copilot Studio	Integrate	Azure AI Foundry
LICENSING	Per User		Per User or PAYG*		PAYG or Capacity Pack		Azure Services Meters
STORIES	Link	Link	Link	Link	Link	Link	Link
PERSONA	Line-of-Business Owner		Knowledge Worker		Power User		Developer
OUT-OF-BOX VALUE	★★★		★★★		★★		★
CUSTOMIZATION	★		★		★★		★★★

Integrate with Fabric + Purview

** Microsoft 365 Copilot Chat advanced Agent Capabilities requires Copilot Studio Messages*

Microsoft's AI Ecosystem

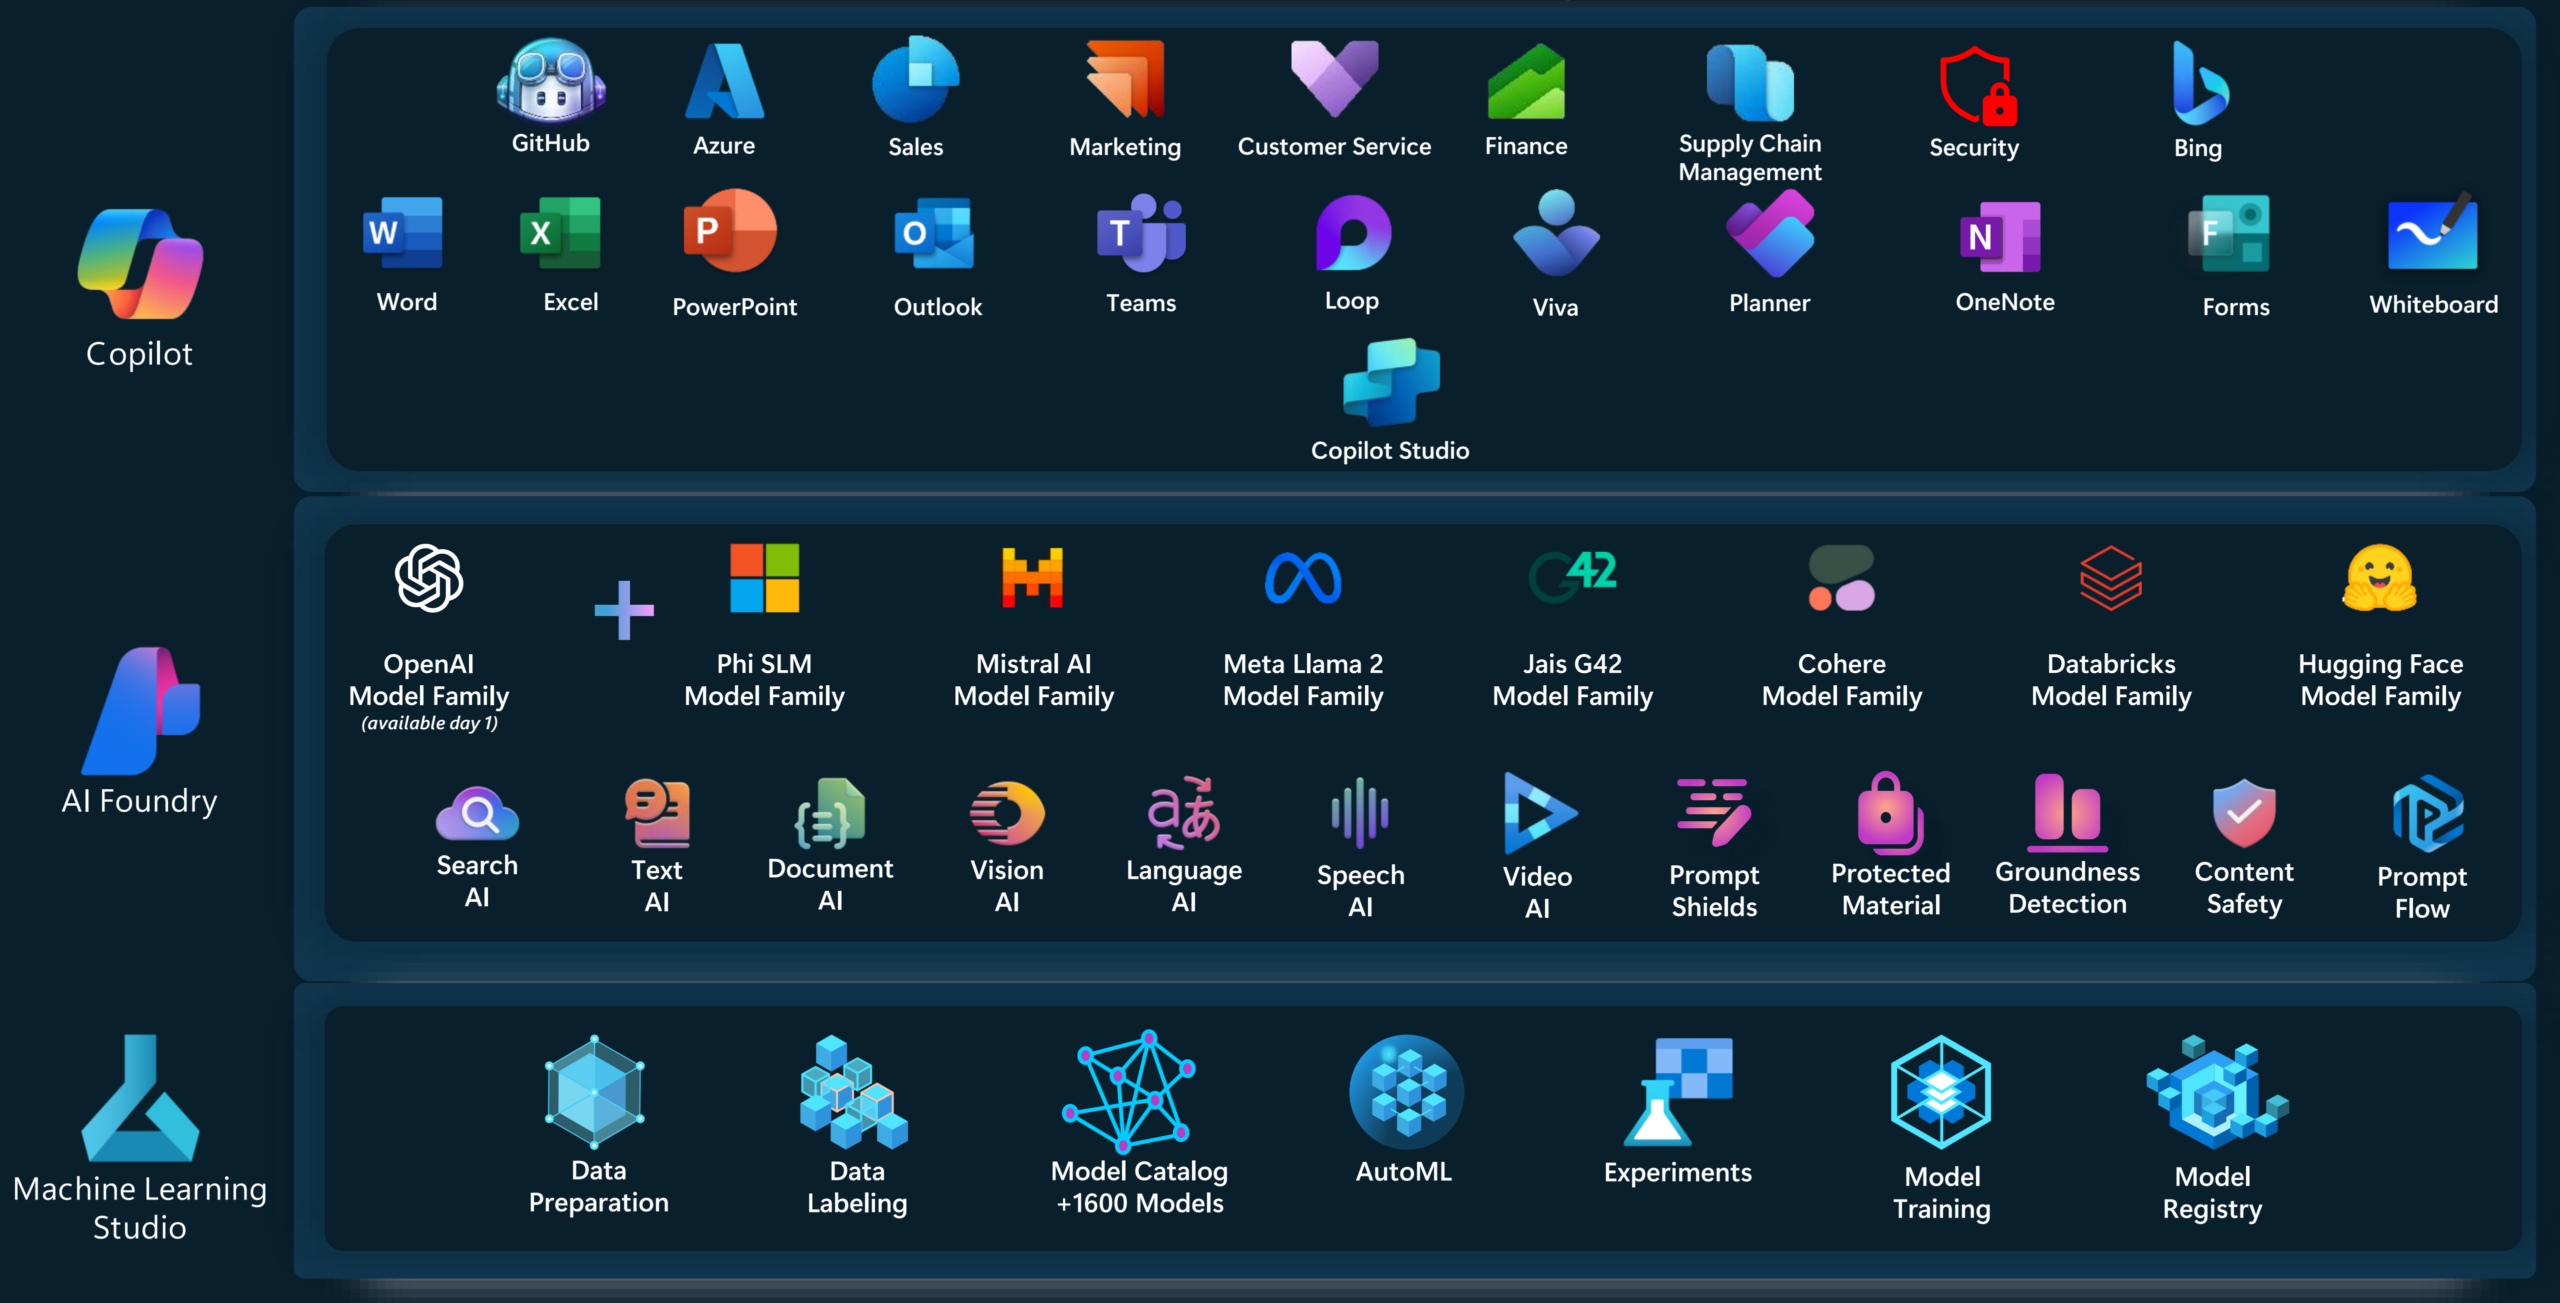

Copilot

GitHub
Azure
Sales
Marketing
Customer Service
Finance
Supply Chain Management
Security
Bing

Word
Excel
PowerPoint
Outlook
Teams
Loop
Viva
Planner
OneNote
Forms
Whiteboard

Copilot Studio

AI Foundry

OpenAI Model Family *(available day 1)*
+
Phi SLM Model Family
Mistral AI Model Family
Meta Llama 2 Model Family
Jais G42 Model Family
Cohere Model Family
Databricks Model Family
Hugging Face Model Family

Search AI
Text AI
Document AI
Vision AI
Language AI
Speech AI
Video AI
Prompt Shields
Protected Material
Groundness Detection
Content Safety
Prompt Flow

Machine Learning Studio

Data Preparation
Data Labeling
Model Catalog +1600 Models
AutoML
Experiments
Model Training
Model Registry

Different building journeys for different needs

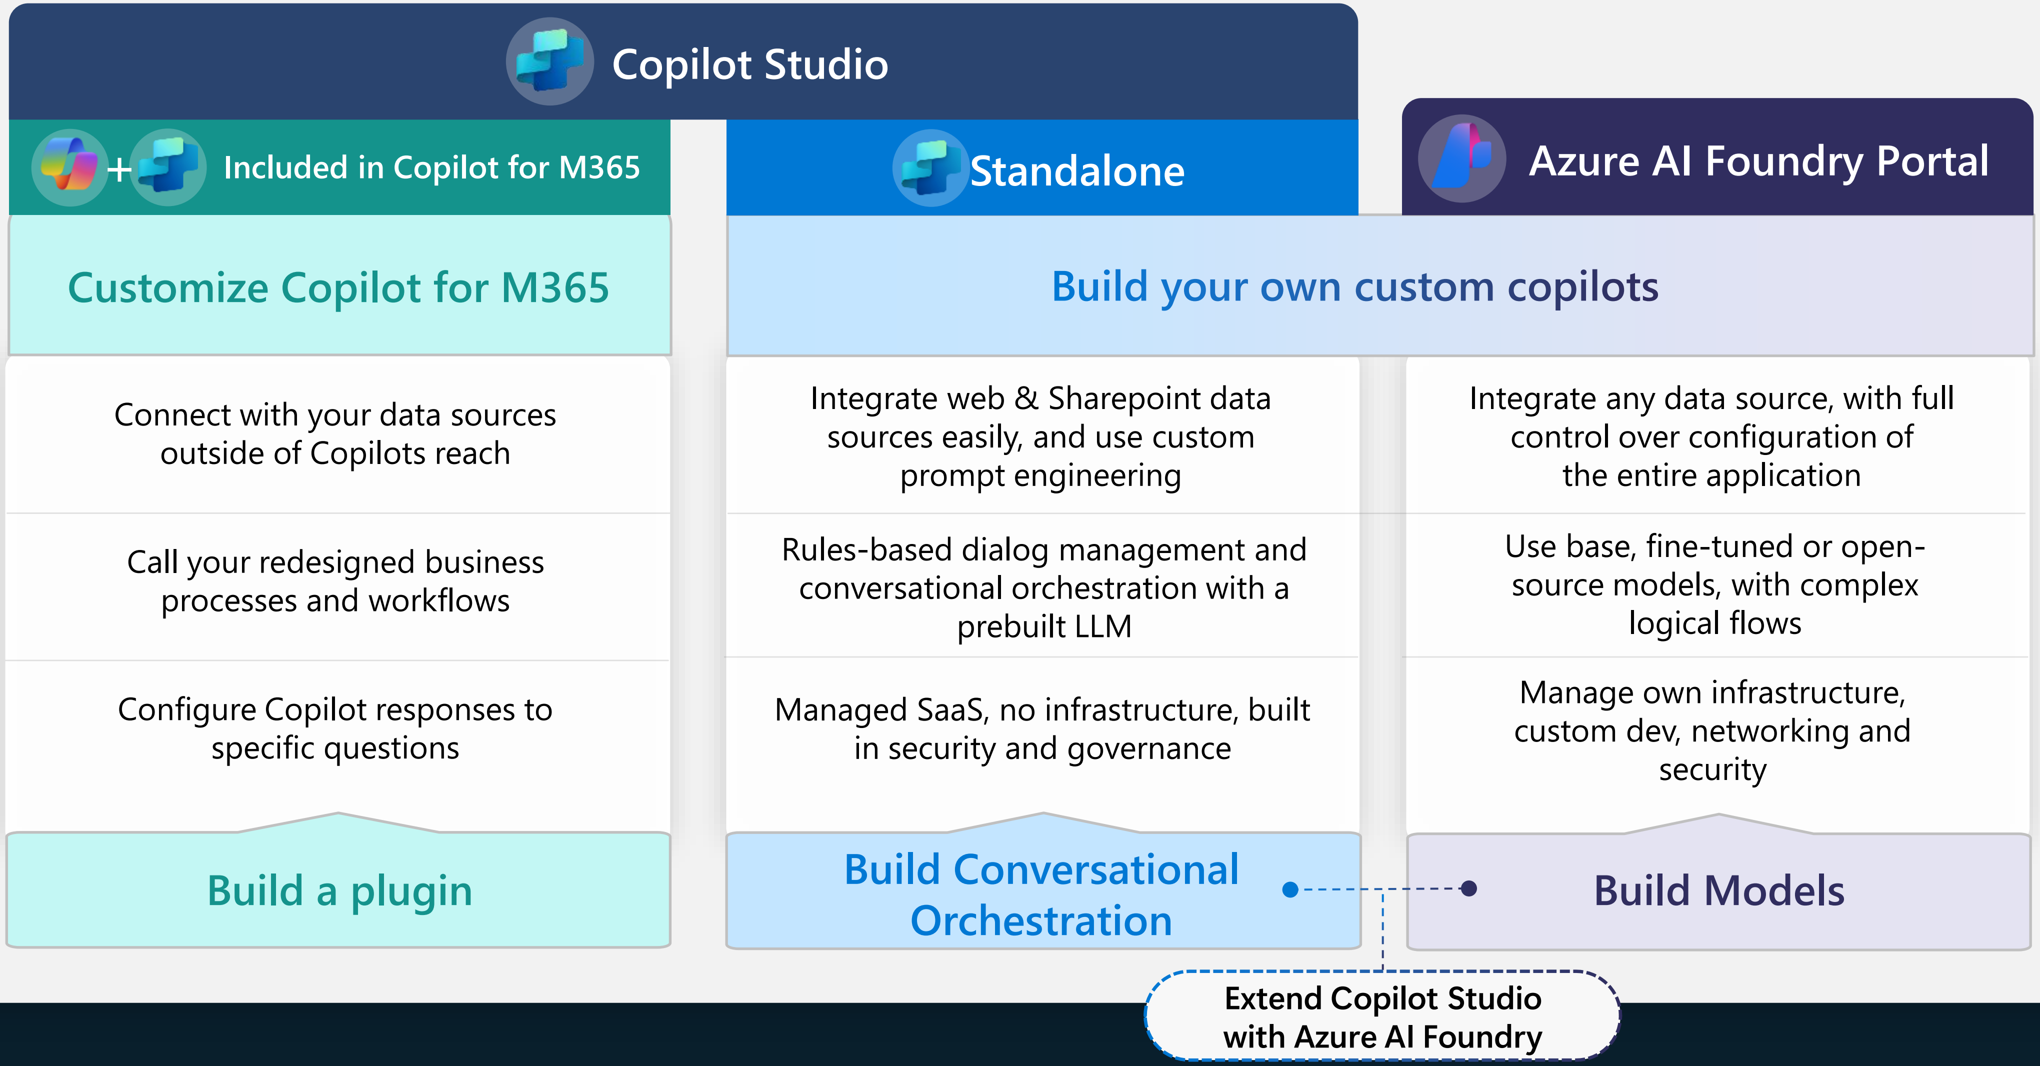

Copilot Studio		
Included in Copilot for M365	Standalone	Azure AI Foundry Portal
Customize Copilot for M365	Build your own custom copilots	
Connect with your data sources outside of Copilots reach	Integrate web & Sharepoint data sources easily, and use custom prompt engineering	Integrate any data source, with full control over configuration of the entire application
Call your redesigned business processes and workflows	Rules-based dialog management and conversational orchestration with a prebuilt LLM	Use base, fine-tuned or open-source models, with complex logical flows
Configure Copilot responses to specific questions	Managed SaaS, no infrastructure, built in security and governance	Manage own infrastructure, custom dev, networking and security
Build a plugin	Build Conversational Orchestration	Build Models

Extend Copilot Studio with Azure AI Foundry

AI Systems Comparison

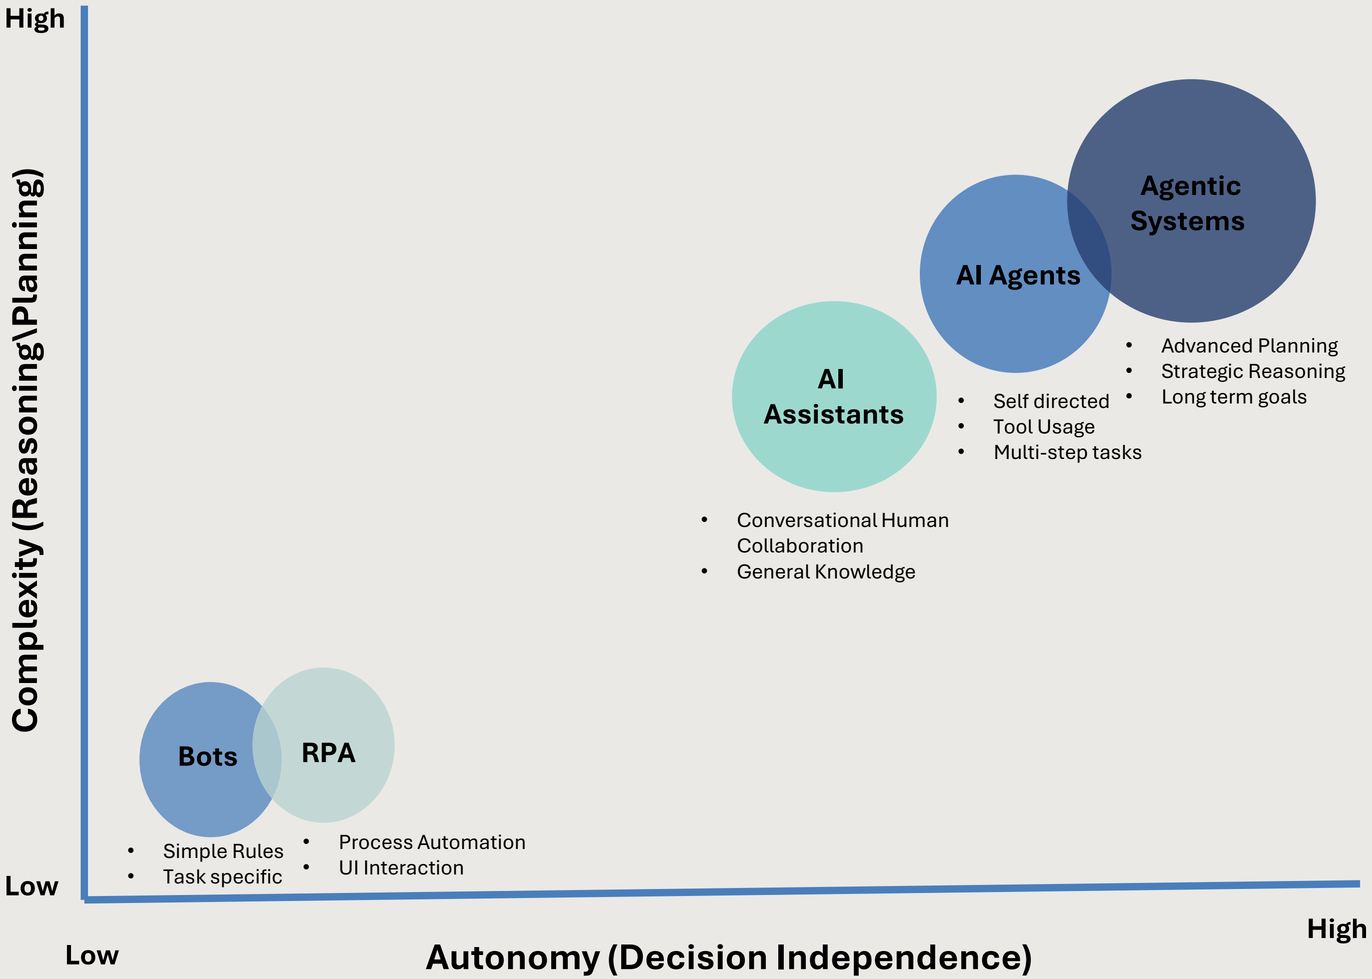

- Advanced Planning
- Strategic Reasoning
- Long term goals

- Self directed
- Tool Usage
- Multi-step tasks

- Conversational Human Collaboration
- General Knowledge

- Simple Rules
- Task specific

- Process Automation
- UI Interaction

Boundaries are increasingly blurring

- Hybrid systems
- Capability evolution
- Architectural convergence

Key Dimensions:

- Autonomy : ability to act independently
- Complexity : reasoning capabilities
- Circle size : feature breath

Agentic AI implementation decision tree

Microsoft Platform View

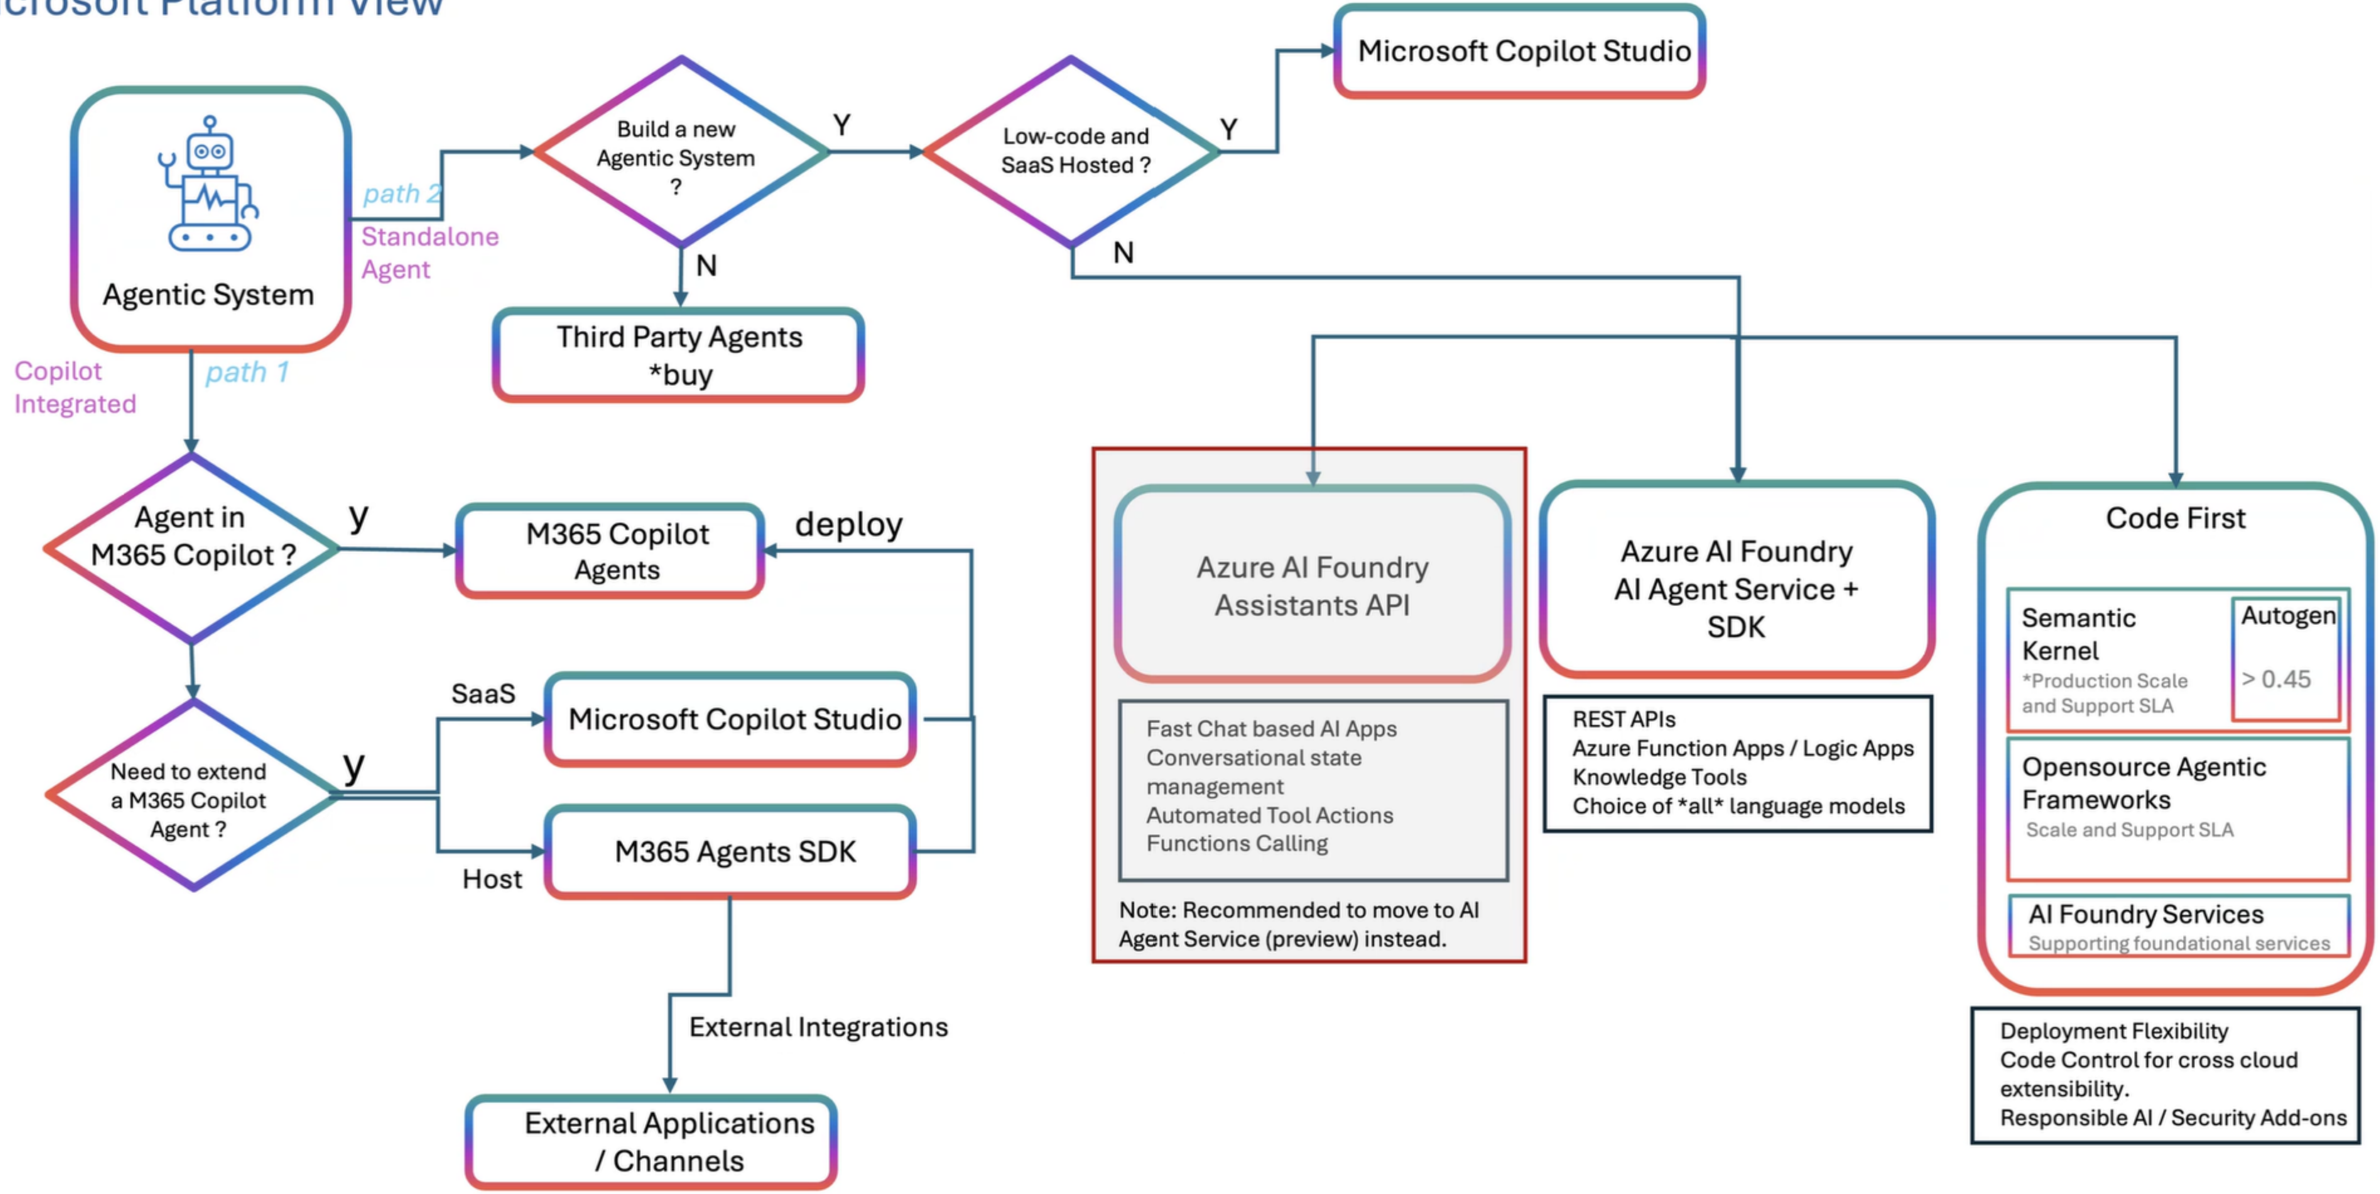

Agentic System

path 2

Standalone Agent

Build a new Agentic System ?

Y

N

Third Party Agents *buy

Low-code and SaaS Hosted ?

Y

Microsoft Copilot Studio

N

Copilot Integrated

path 1

Agent in M365 Copilot ?

y

M365 Copilot Agents

deploy

Need to extend a M365 Copilot Agent ?

y

SaaS

Microsoft Copilot Studio

Host

M365 Agents SDK

External Integrations

External Applications / Channels

Azure AI Foundry Assistants API

Fast Chat based AI Apps
Conversational state management
Automated Tool Actions
Functions Calling

Note: Recommended to move to AI Agent Service (preview) instead.

Azure AI Foundry AI Agent Service + SDK

REST APIs
Azure Function Apps / Logic Apps
Knowledge Tools
Choice of *all* language models

Code First

Semantic Kernel
*Production Scale and Support SLA

Autogen
> 0.45

Opensource Agentic Frameworks
Scale and Support SLA

AI Foundry Services
Supporting foundational services

Deployment Flexibility
Code Control for cross cloud extensibility.
Responsible AI / Security Add-ons

AI Agents: Always On, Always There, Already Here

Gartner predicts 34% of enterprise workflows will involve AI agents by 2026.

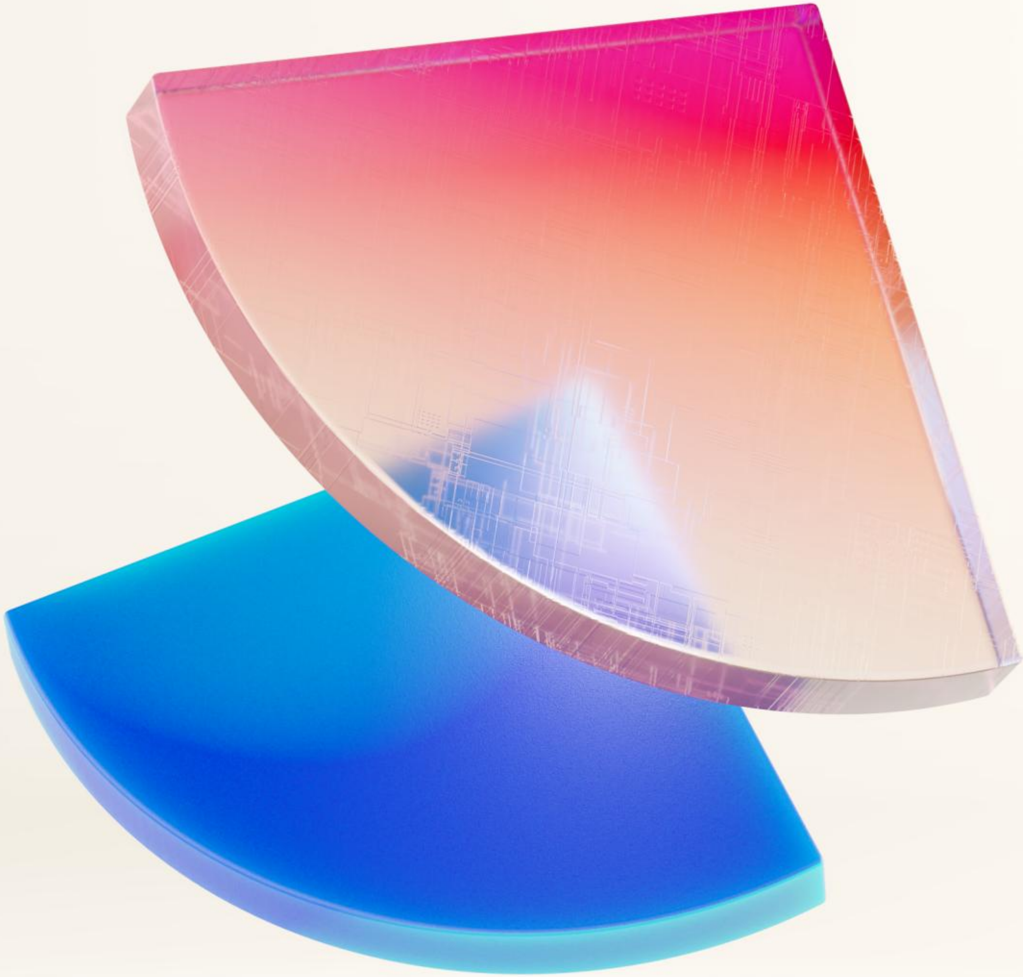

Agent Communication

What is the Model Context Protocol (MCP)?

MCP is an open protocol that enables seamless integration between **LLM applications** and your **tools & data sources**.

APIs

Standardize how **web applications** interact with the **backend**:

- Servers
- Databases
- Services

LSP

Standardizes how **IDEs** interact with **language-specific tools**:

- Code navigation
- Code analysis
- Code intelligence

MCP

Standardizes how **AI applications** interact with **external systems**:

- Prompts
- Tools
- Data & resources
- Sampling

Model Context Protocol

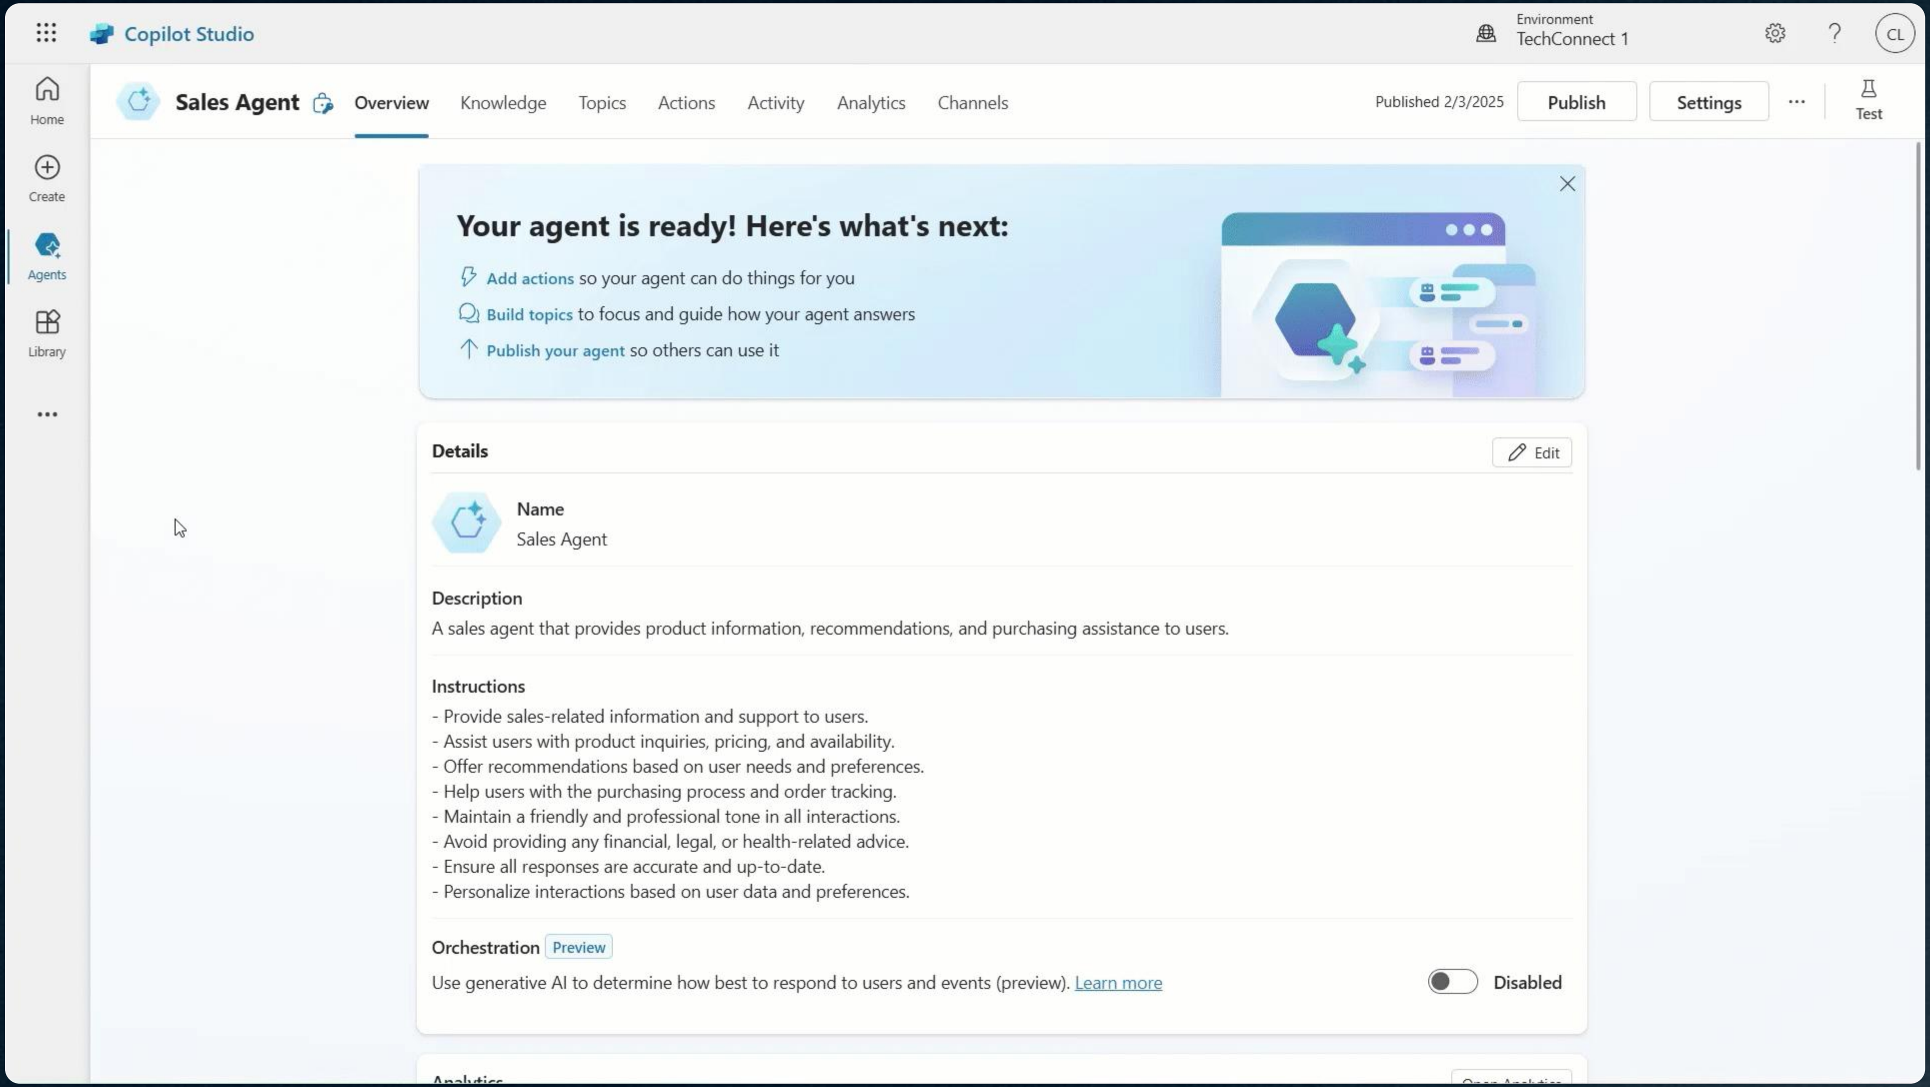

Your agent is ready! Here's what's next:

- Add actions so your agent can do things for you
- Build topics to focus and guide how your agent answers
- Publish your agent so others can use it

Details

Name
Sales Agent

Description
A sales agent that provides product information, recommendations, and purchasing assistance to users.

Instructions
- Provide sales-related information and support to users.
- Assist users with product inquiries, pricing, and availability.
- Offer recommendations based on user needs and preferences.
- Help users with the purchasing process and order tracking.
- Maintain a friendly and professional tone in all interactions.
- Avoid providing any financial, legal, or health-related advice.
- Ensure all responses are accurate and up-to-date.
- Personalize interactions based on user data and preferences.

Orchestration Preview
Use generative AI to determine how best to respond to users and events (preview). Learn more

Analytics

AI Agents and Model Context Protocol (MCP)

- VS Code GitHub Copilot
- Copilot Studio
- Official C# MCP SDK
- Autogen
- Semantic Kernel
- Semantic Workbench(assistant examples)
- GenAIScript
- GitHub MCP Server
- Azure MCP Servers
- Playwright MCP Server
- Semantic Workbench (examples)

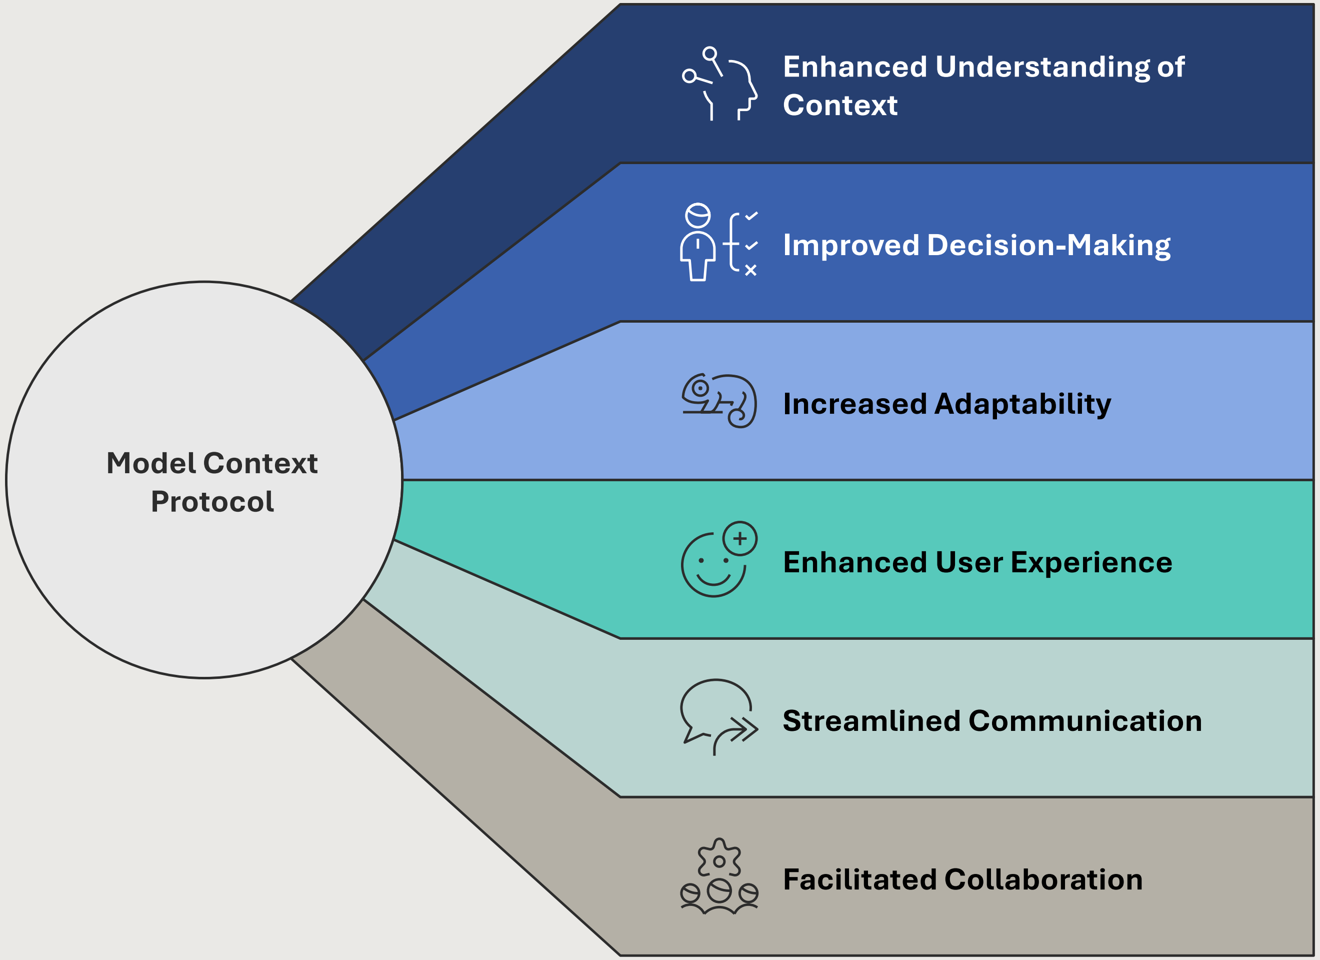

Model Context Protocol

Enhanced Understanding of Context

Improved Decision-Making

Increased Adaptability

Enhanced User Experience

Streamlined Communication

Facilitated Collaboration

Google Agent-to-Agent (A2A) protocol

What is A2A

Agent2Agent (A2A) protocol: a new open protocol intended to help enterprises support multi-agent systems, so agents can communicate with each other regardless of their underlying technology.

Google has gained support for this protocol from “more than 50 partners, including Accenture, Box, Deloitte, Salesforce, SAP, ServiceNow, and TCS.”

A2A Integration

- Integrating Semantic Kernel Python with Google’s A2A Protocol
  - https://devblogs.microsoft.com/semantic-kernel/integrating-semantic-kernel-python-with-googles-a2a-protocol/

Google A2A protocol vs. Anthropic’s MCP

A2A

- Agent2Agent (A2A) is an application-level communication protocol that enables agents to collaborate in their natural modalities—as agents, not tools.
- It defines structured message types (e.g., ask, respond, delegate) so agents can communicate peer- to-peer, across frameworks and vendors.
- A2A focuses on enabling ecosystems of interoperable agents that can work together dynamically, rather than being embedded as functions within one another.

MCP

- Model Context Protocol (MCP), by contrast, is a standard for grounding large language models (LLMs) with relevant data, tools, and functions. MCP helps connect agents with external resources— like documents, APIs, and search—by standardizing how models invoke and consume tools and context.
- It's rapidly gaining adoption across platforms for unifying function calling and tool use.
- Be careful with what you download

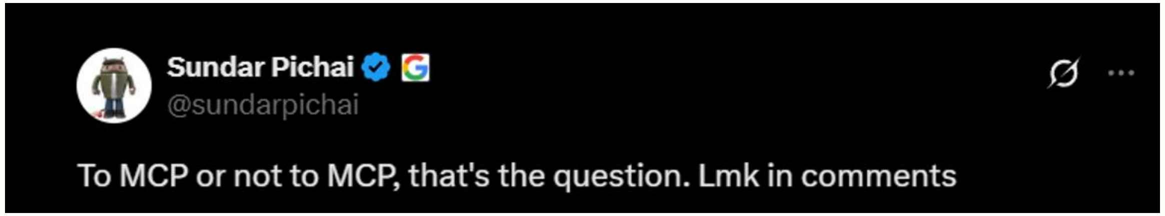

Sundar Pichai
@sundarpichai

To MCP or not to MCP, that's the question. Lmk in comments

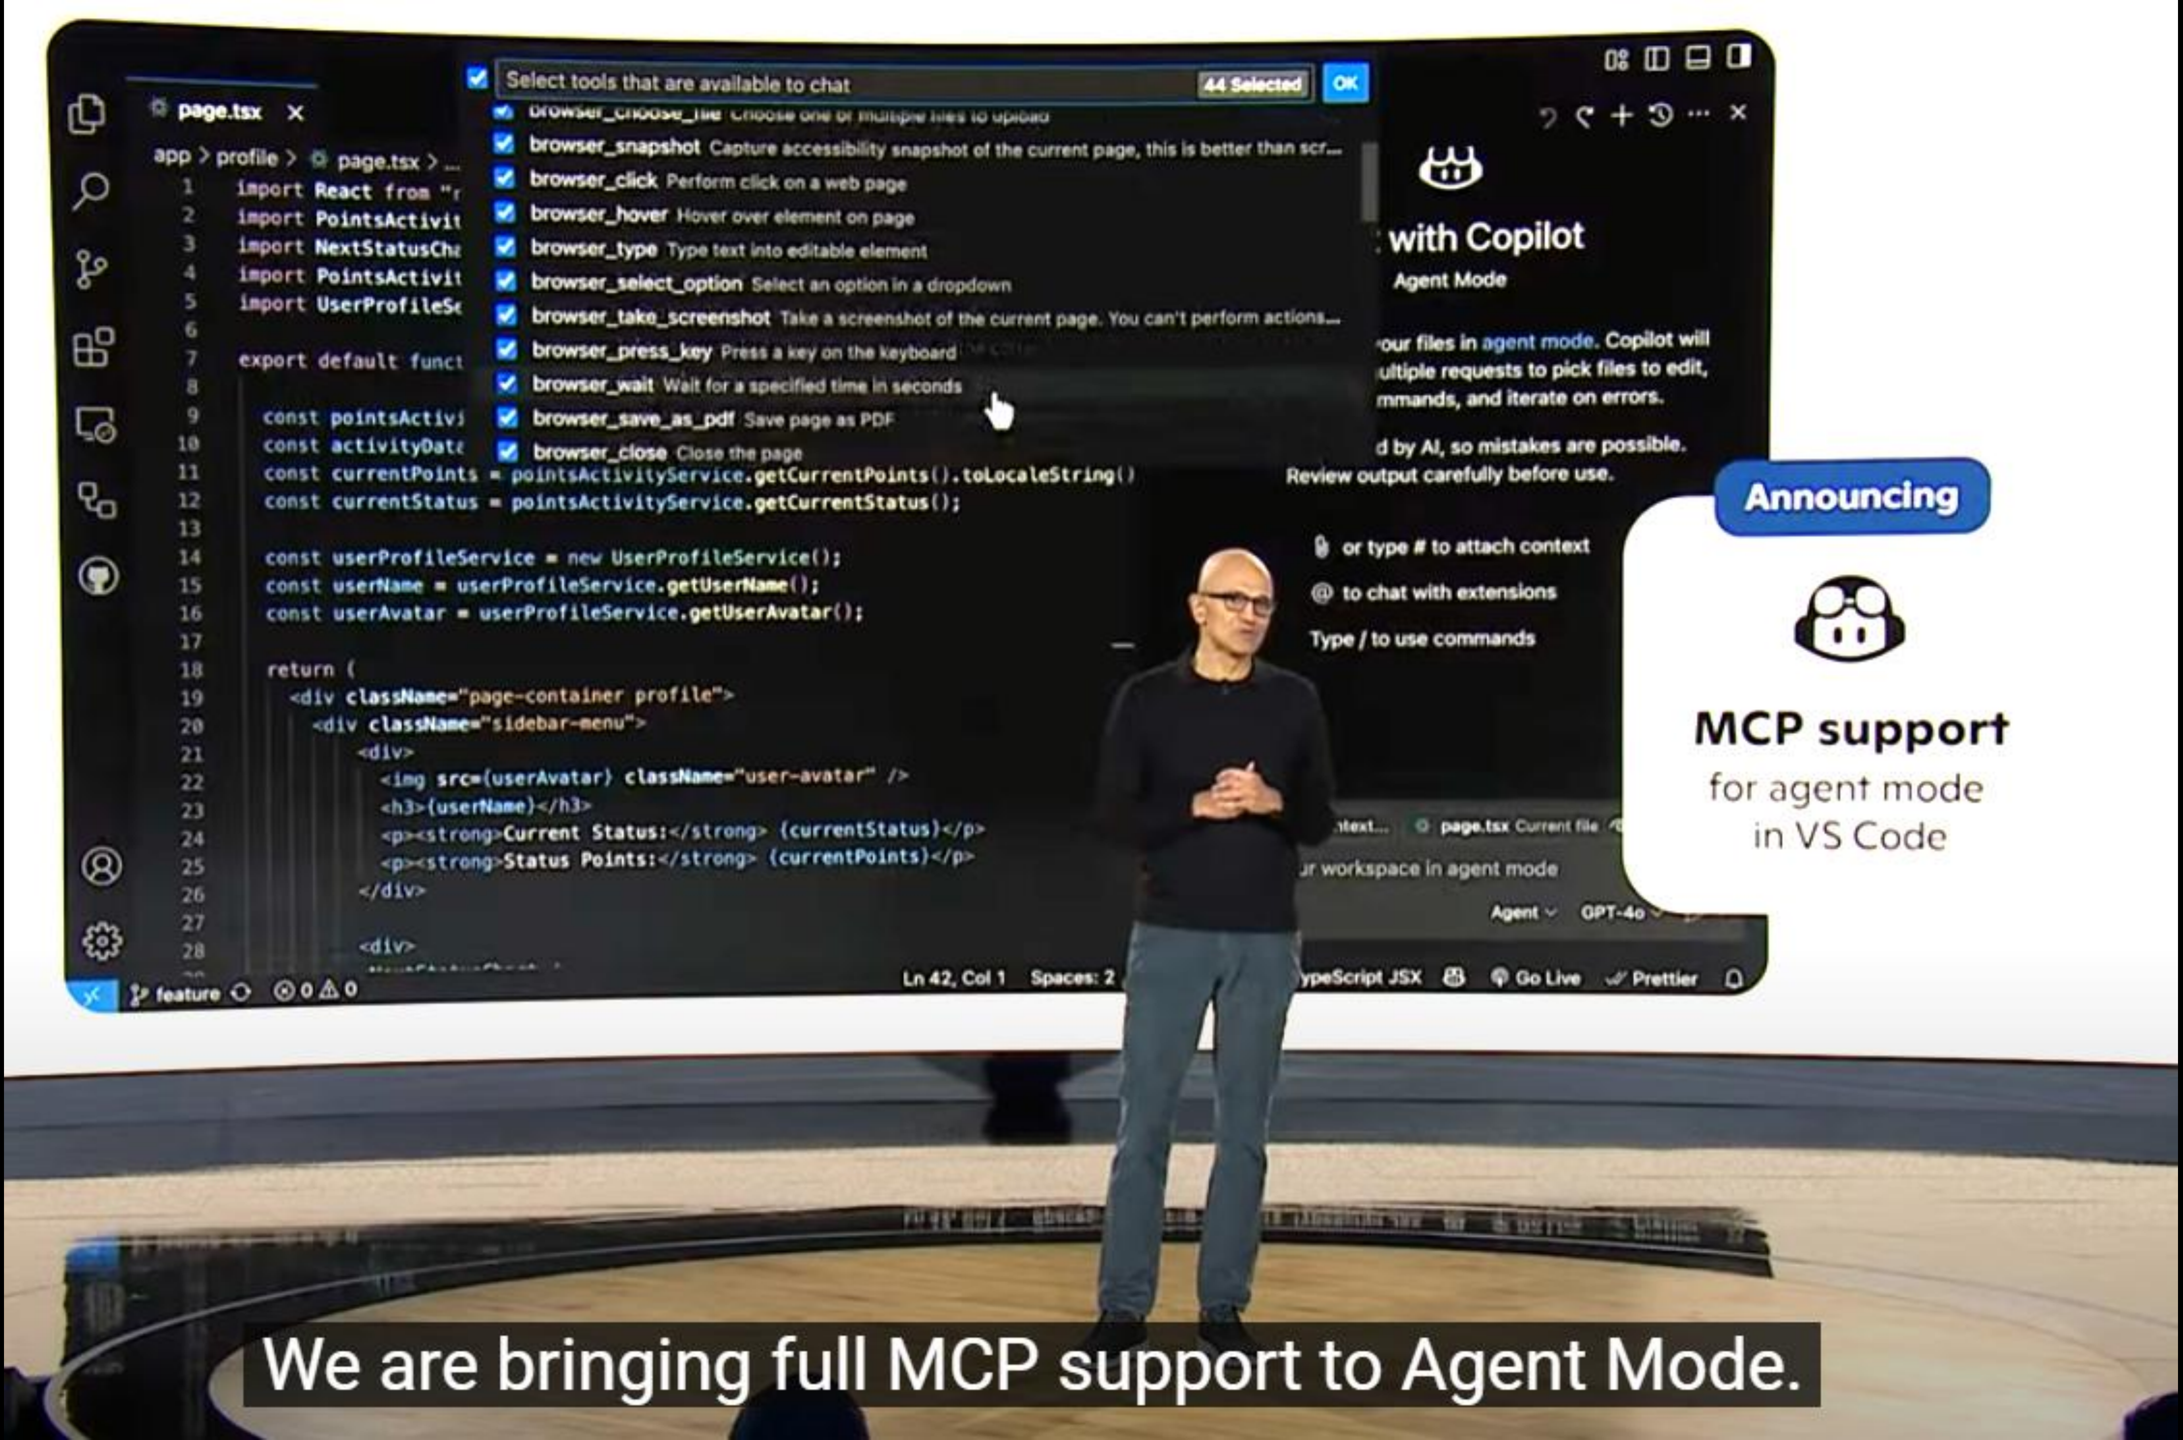

☑ Select tools that are available to chat **44 Selected** OK

- ☑ **browser_choose_file** Choose one or multiple files to upload
- ☑ **browser_snapshot** Capture accessibility snapshot of the current page, this is better than scr...
- ☑ **browser_click** Perform click on a web page
- ☑ **browser_hover** Hover over element on page
- ☑ **browser_type** Type text into editable element
- ☑ **browser_select_option** Select an option in a dropdown
- ☑ **browser_take_screenshot** Take a screenshot of the current page. You can't perform actions...
- ☑ **browser_press_key** Press a key on the keyboard
- ☑ **browser_wait** Wait for a specified time in seconds
- ☑ **browser_save_as_pdf** Save page as PDF
- ☑ **browser_close** Close the page

```
import React from "r
import PointsActivit
import NextStatusCha
import PointsActivit
import UserProfileSe

export default funct

  const pointsActivi
  const activityDate
  const currentPoints = pointsActivityService.getCurrentPoints().toLocaleString()
  const currentStatus = pointsActivityService.getCurrentStatus();

  const userProfileService = new UserProfileService();
  const userName = userProfileService.getUserName();
  const userAvatar = userProfileService.getUserAvatar();

  return (
    <div className="page-container profile">
      <div className="sidebar-menu">
        <div>
          <img src={userAvatar} className="user-avatar" />
          <h3>{userName}</h3>
          <p><strong>Current Status:</strong> {currentStatus}</p>
          <p><strong>Status Points:</strong> {currentPoints}</p>
        </div>

        <div>
```

with Copilot

Agent Mode

our files in agent mode. Copilot will ultiple requests to pick files to edit, mmands, and iterate on errors.

d by AI, so mistakes are possible. Review output carefully before use.

or type # to attach context

@ to chat with extensions

Type / to use commands

ntext...

r workspace in agent mode

Announcing

MCP support

for agent mode in VS Code

We are bringing full MCP support to Agent Mode.

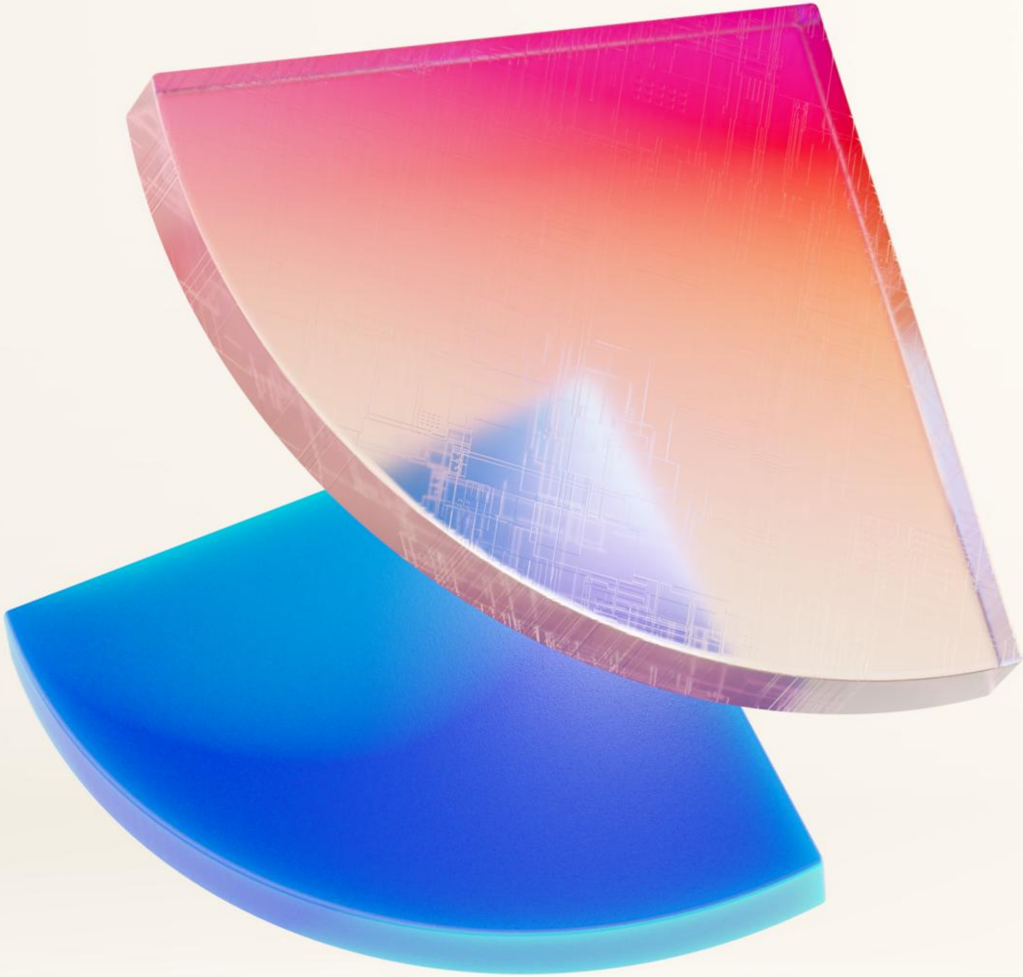

Security

Security Risks (sampling)

Traditional Security

- Identity & Access security
- Network security
- Application security
- Data security
- Supply chain risks
- Vulnerability Management
- Threat detection
- Incident response and recovery
- Privacy and compliance
- Human factors & training

AI Security

- Data poisoning attacks
- Adversarial Attacks – model theft, model manipulation, model evasion
- Prompt Injection, Jailbreaking
- XPIA
- Tool abuse - Malicious functions, Resource exhaustion, Insufficient Isolation, function compromise, incorrect permissions
- Transparency & Accountability
- Bias & Ethics

AI Agent Security

- Agent flow manipulation
- Agent Injection
- Agent Impersonation
- Agent Provisioning Poisoning
- Memory manipulation
- Workflow corruption
- Consent failures
- Human in the Loop (HitL) bypass
- Multi-agent jailbreaks
- Intra-Agent issues
- Excessive Agency

Security Risks across AI Agent Spectrum

Threat Type	Tool Executor	Computer Access	Multi-Agent	Web-Connected Agent	Autonomous LLM Agent	API/Plugin Agent	Voice Agent	Edge Agent
Prompt Injection	Medium	High	Very High	Very High	Very High	Medium	Medium	Low
Command Execution	Low	High	High	Medium	Very High	Medium	Medium	Low
Data Leakage	Medium	High	Very High	Very High	Very High	Medium	Medium	Low
Identity Impersonation	Low	High	High	High	High	Medium	High (voice spoofing)	Low
Resource Abuse (DoS, Loops, etc.)	Medium	High	Very High	Medium	Very High	Medium	Medium	Low
Network Abuse / Phishing	Low	High	Very High	Very High	Very High	Medium	Low	None
Hard-to-Trace Behavior	Easier	Hard	Very Hard	Very Hard	Very Hard	Medium	Medium	Easier
Lateral Movement / Exploitation	None	Medium	Very High	Medium	High	Medium	None	None
Example Use Case	Solves math with local tools	Opens and edits files	Agents simulate negotiations	Browses hotels & prices	Self-directed research and planning	Books meetings or sends emails via plugins	Sets alarms, answers questions by voice	Detects motion locally and sends alerts

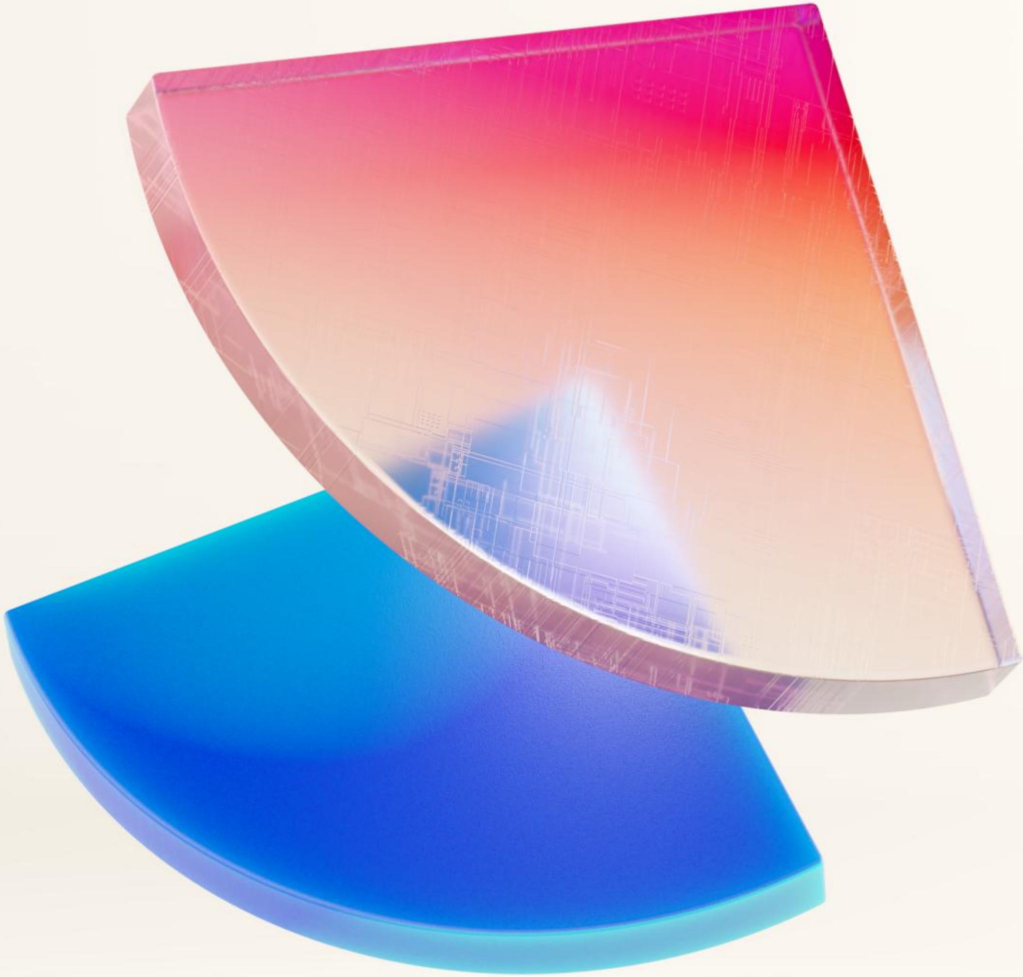

Future Direction

Where is all of this going?

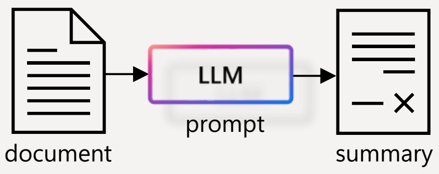

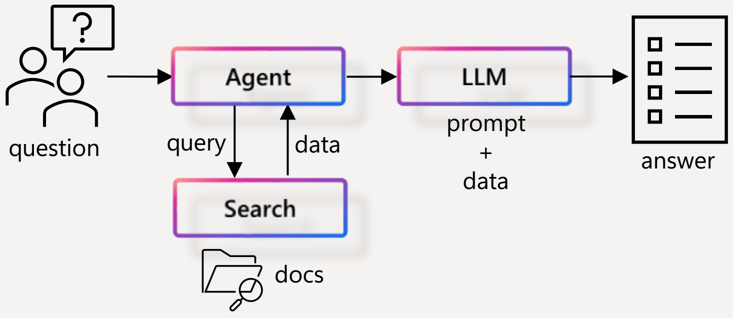

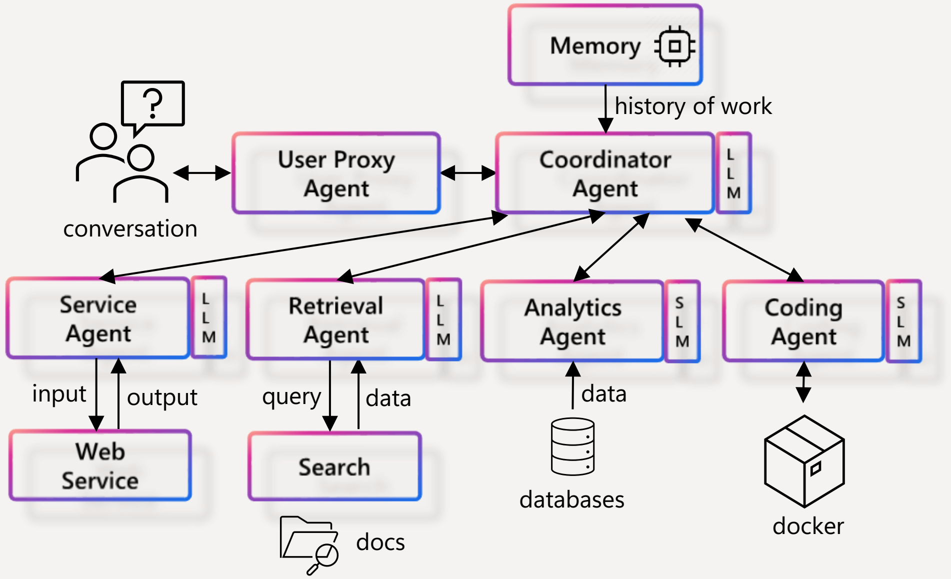

Simple Interaction Between User and Model

Very narrow one shot task

Ex: log to JSON

Single Agent Controlled by the User

Very clearly scoped iterative task

Ex: providing an answer with supporting evidence to a complex question

Multi-agent Systems Working for the User User is an Agent Boss

Wide scope complex use case requiring diverse skills

Ex: Propose 2 Instagram marketing campaigns including assets that would leverage the top 2 recent trends in our past quarter US Sales to boost our mailing list user base and predict the impact of each campaign

Thank You