ECS 2026

Copilot Studio Architecture

European Microsoft
CX CTO
Capgemini

Freelance
Consultant

Program Manager
Microsoft

PREMIUM

EXPERTS INSIDE
Teaming up with EasyLife 365

PREMIUM PARTNER

Microsoft

TECHNOLOGY PARTNER

run events

DIAMOND

LightningTools
sitaas
ShareGate
veeam

ECS 2026
05-07 MAY | COLOGNE | GERMANY

PLATINUM

365 TRIBE
ADOPTION. CHAMPIONED
daenet
daenet – an AGT Digital company
DEKOM
dox42
automate your documents integrate your data
glueck kanja
delaware
involv
ITENOS
IT's us.
Rencore
SOLUTIONS 2SHARE
skybow
SquaredUp
Syskit
OASE
YASH Technologies
More than what you think.
Jabra GN
365 HEROES
Orchestry
Capgemini
AvePoint
infobip
SHURE
neat.

GOLD

experlogix
protiviti
Global Business Consulting
apps 365
BCC
Hive
VISIONET
Engineering. Simplified.
Yealink
Gamma
ScriptRunner
LEAP XPERT
Responsible Business Communication
mscrm-addons.com
Your company for MS-CRM ADD-ONS!
envision IT
proventeq
TD SYNNEX

SILVER

Apvee
CROSSWARE
FLEXCOM
ichicraft
panagenda
pointfire
Multilingual SharePoint
rocketta
EasyEntra
FrontFox
u2u
tabella
all about documents
zenitech

BRONZE

Alpha Variance Solutions | GmbH
kurmi
Powell
INTELLITY

FUTURE MAKER SILVER

yasp
workplace buddy.
2501 AI

FUTURE MAKER BRONZE

ARGOS

MEDIA PARTNER

AI Tools Network
Times Of AI
AIPressRoom
Technology Record
IT-Schulungen.com
Technology & Transformation
CAPITALBAY NEWS
AI&ML events

Agenda

Welcome & Introductions	**9:00 – 10:30**
Out-of-the-Box Architecture M365 Copilot, Copilot Studio Lite & Full, AI Foundry	
A2A & MCP Protocols CCA architecture, multilingual agents, speech services	
Break	**10:30 – 10:45**
Extensibility Extending each platform and combining them	**10:45 – 12:30**
Lunch	**12:30 – 13:30**
Voice & Translation Agent-to-Agent, Model Context Protocol, best practices	**13:30 – 14:30**
Governance Individual agent and orchestrated governance, Agent 365	
Break	**14:30 – 14:45**
Security & Monitoring Threat protection, monitoring, OWASP for agents	**14:45 – 15:15**
Scenarios	**15:15 – 16:30**
End of the day	**16:30**

Section 1

Out-of-the-Box Architecture

Understanding what each AI element provides natively and when to choose which

Where Microsoft's Copilot and AI Build Options Fit

M365 Copilot

Best for business users working in Word, Excel, PowerPoint, Outlook, and Teams. It brings AI directly into daily work using Microsoft 365 context, permissions, and content.

User productivity layer

Copilot Studio Lite

Best for teams that want simple agent extensions around Copilot without a full developer-style build motion. Think guided customization, grounded prompts, and lighter-weight business automation.

Light customization layer

Copilot Studio

Best for makers and enterprise teams building custom agents with connectors, actions, orchestration, and governance. It is the low-code agent factory in the Power Platform stack.

Low-code agent platform

Microsoft Foundry

Best for developers and AI engineers who need model selection, code-first orchestration, evaluation, deployment, and enterprise AI operations. It is the broadest build platform for custom AI apps and agents.

Pro – code AI platform

M365 Copilot Architecture

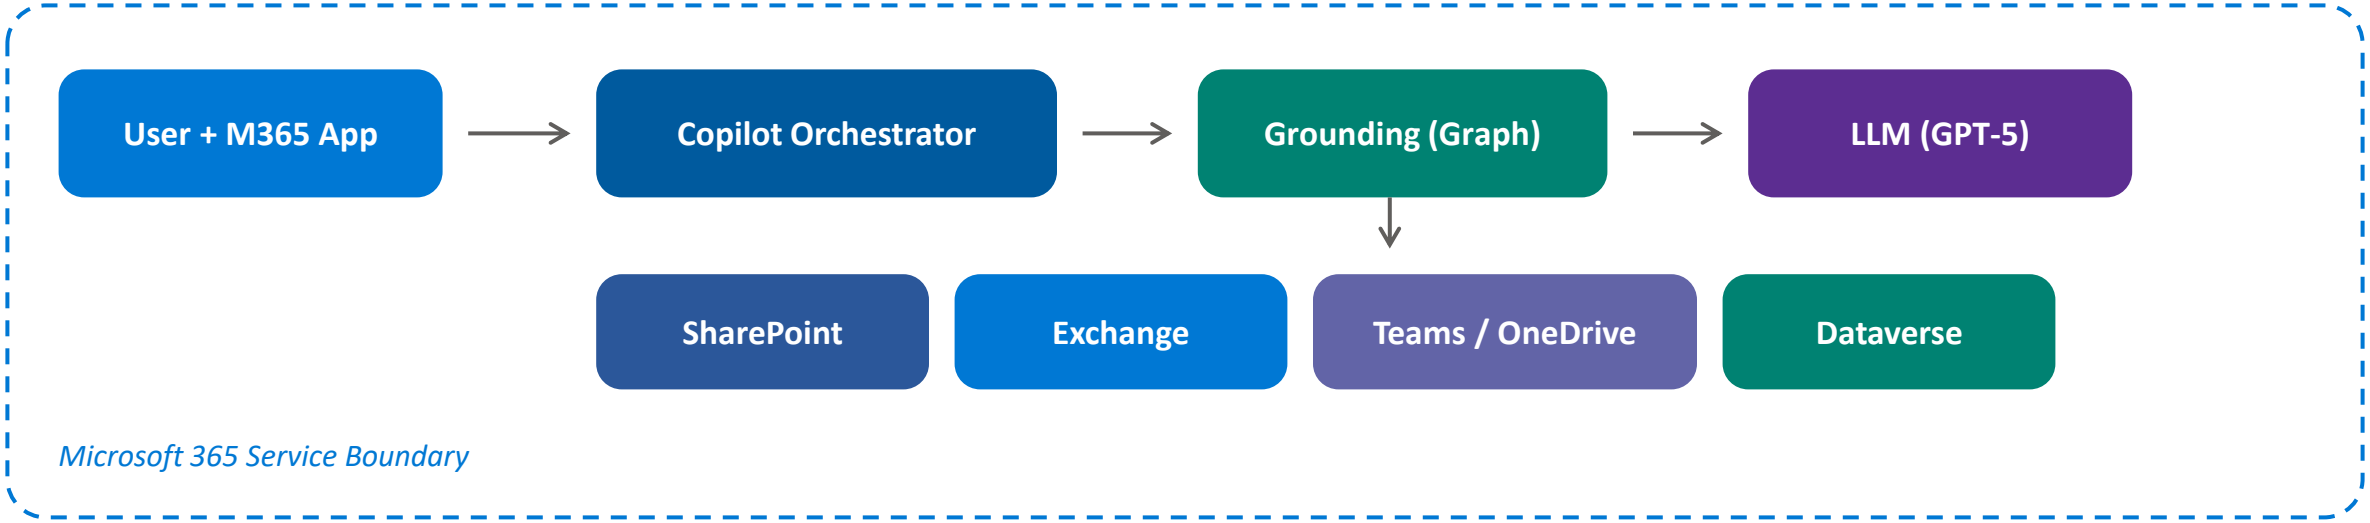

- User permissions enforced via Microsoft Graph at every request
- All data stays within the M365 service boundary
- Grounding enriches prompts with user context from Graph
- Conditional Access and MFA apply to all Copilot interactions

M365 Copilot

What You Get

- Drafts content in Word and Outlook, turning prompts into first-pass documents, emails, and rewrites.
- Summarises meetings, chats, and long documents so users can get to actions faster.
- Analyses data in Excel and helps surface trends, formulas, and narrative explanations.
- Works across Microsoft 365 apps with access shaped by your existing permissions and tenant data.
- Can improve personal productivity by reducing time spent searching, writing, and catching up.

Key Limitations

- Designed for M365 Assistance
- Output quality depends on your data hygiene, permissions model, and the quality of the prompt.
- Responses can be incomplete or wrong, so human review is still required for important work.
- It does not fix poor information architecture, unmanaged content sprawl, or weak governance.

Best for: Quick prototyping, personal productivity agents, citizen developers exploring AI

Copilot Studio Lite Architecture

Agent Builder embedded in Microsoft 365 Copilot — runs inside the M365 service boundary

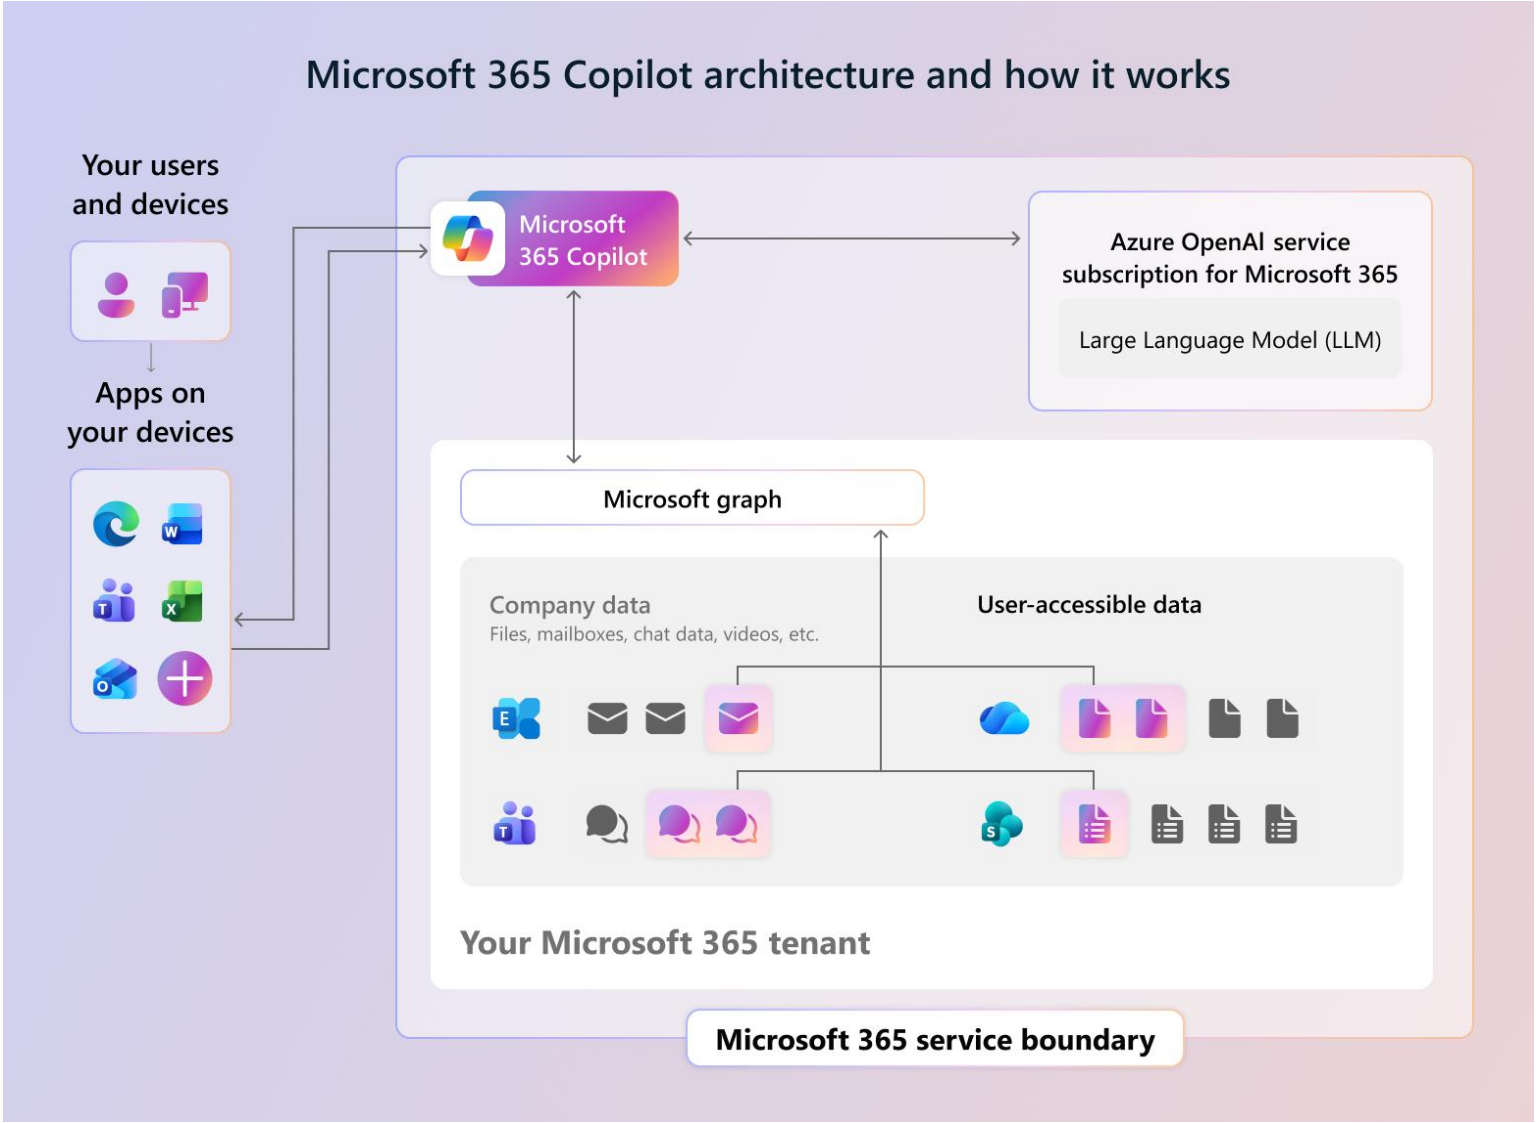

Copilot Studio Lite

What You Get

- Embedded in M365 Copilot app (Teams, web)
- Natural language authoring ("Describe" tab)
- Connect to SharePoint, OneDrive knowledge
- Code interpretation and image generation
- Test agents in embedded side pane
- Quick sharing within your organisation
- GPT-5 Chat model included

Key Limitations

- Limited number of topics
- No premium connectors or Dataverse
- No multi-environment ALM
- Restricted publishing channels
- No advanced workflow branching
- Scoped to user's M365 data context
- No telemetry or analytics dashboards

Best for: Quick prototyping, personal productivity agents, citizen developers exploring AI

Copilot Studio Architecture

Standalone agent platform — runtime, dialog, generative answers, security and ALM across channels

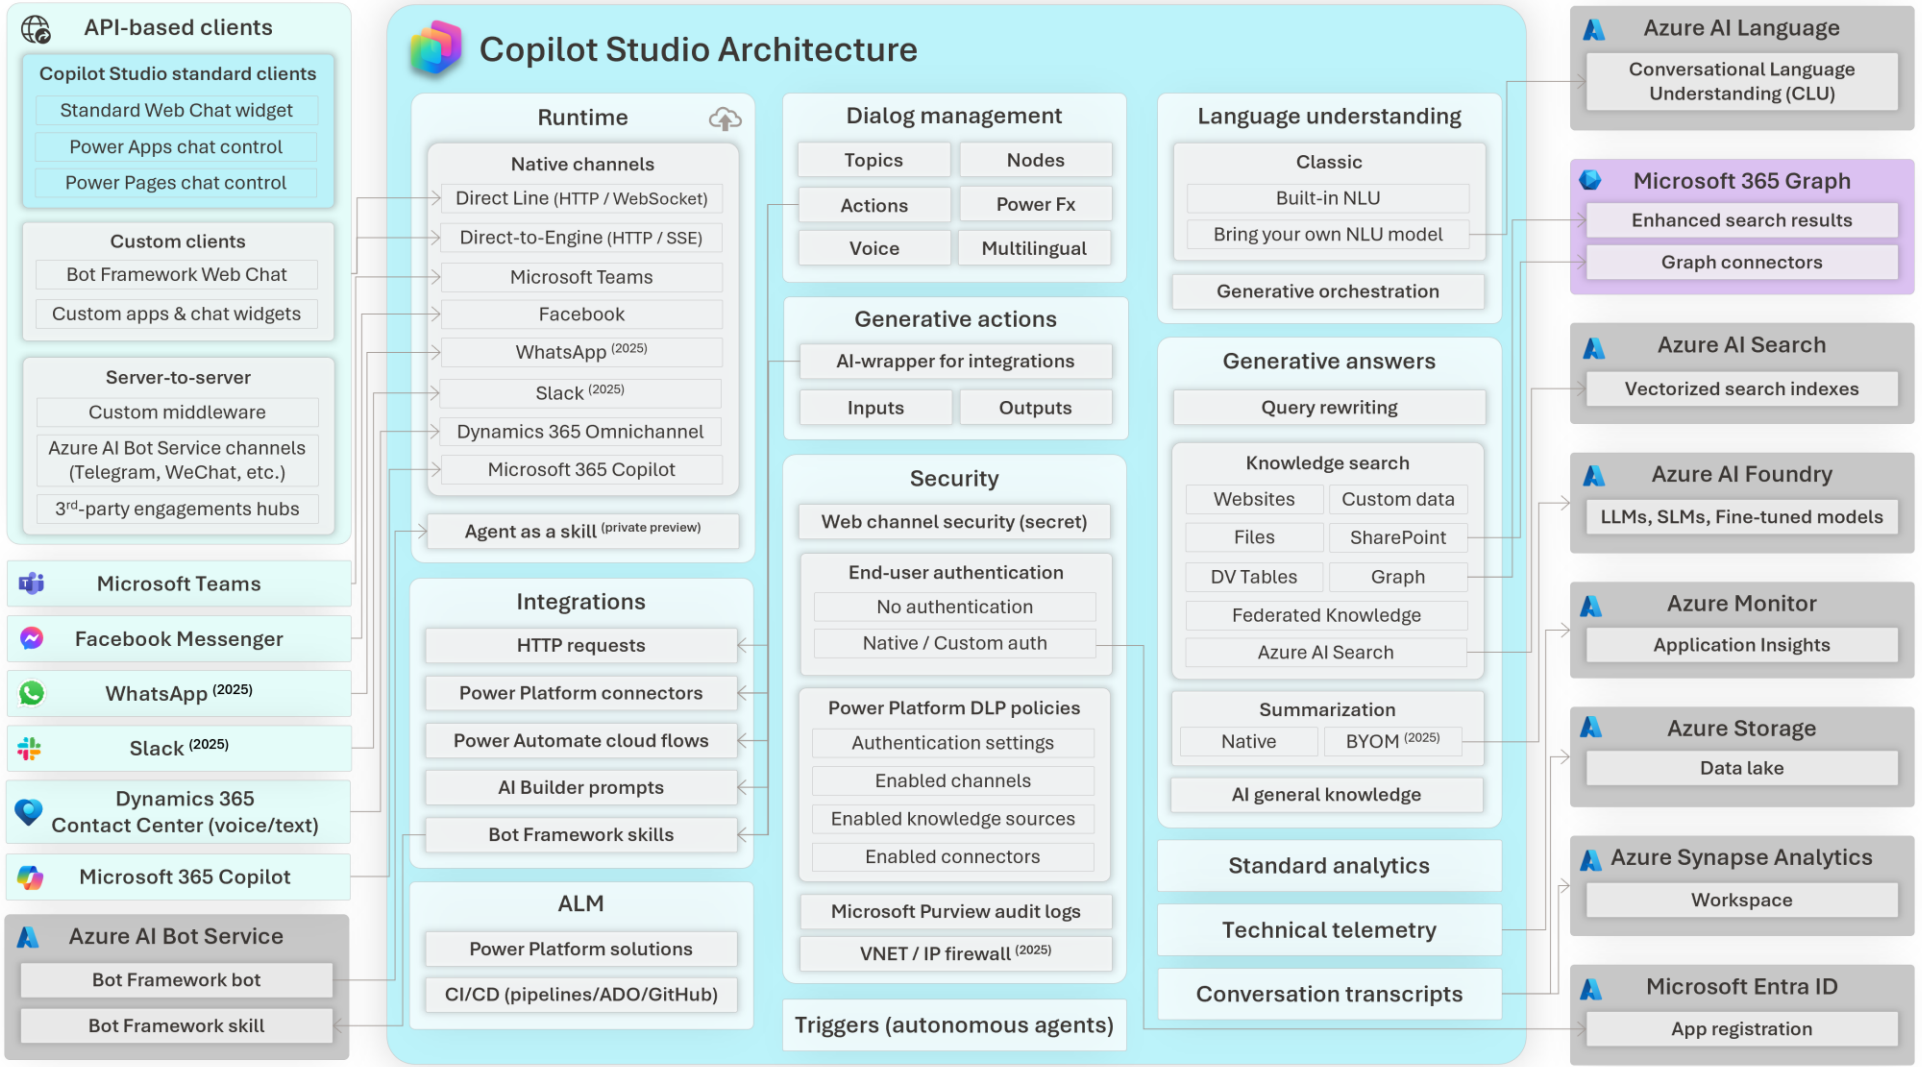

Copilot Studio Full Experience

Full Capabilities

- Standalone portal (copilotstudio.microsoft.com)
- Multi-step workflows and conditional logic
- Premium connectors and external APIs
- Dataverse integration
- Dev / Test / Prod environments (ALM)
- Version control and rollback
- Telemetry and usage analytics
- Computer use capability (UI automation)
- Multi-agent orchestration
- A2A protocol support

Target Audience

Makers, developers, IT teams building agents for departments or enterprise-wide use. Comfortable with complexity and lifecycle management.

Licensing

- Copilot Studio add-on license required
- Included with M365 Copilot (Lite) or pay-as-you-go
- Capacity-based pricing model
- Lite agents can be copied to Full for upgrade

Azure AI Foundry Architecture

Foundry in the Microsoft stack — models, Agent Service, AI services, observability and developer tools

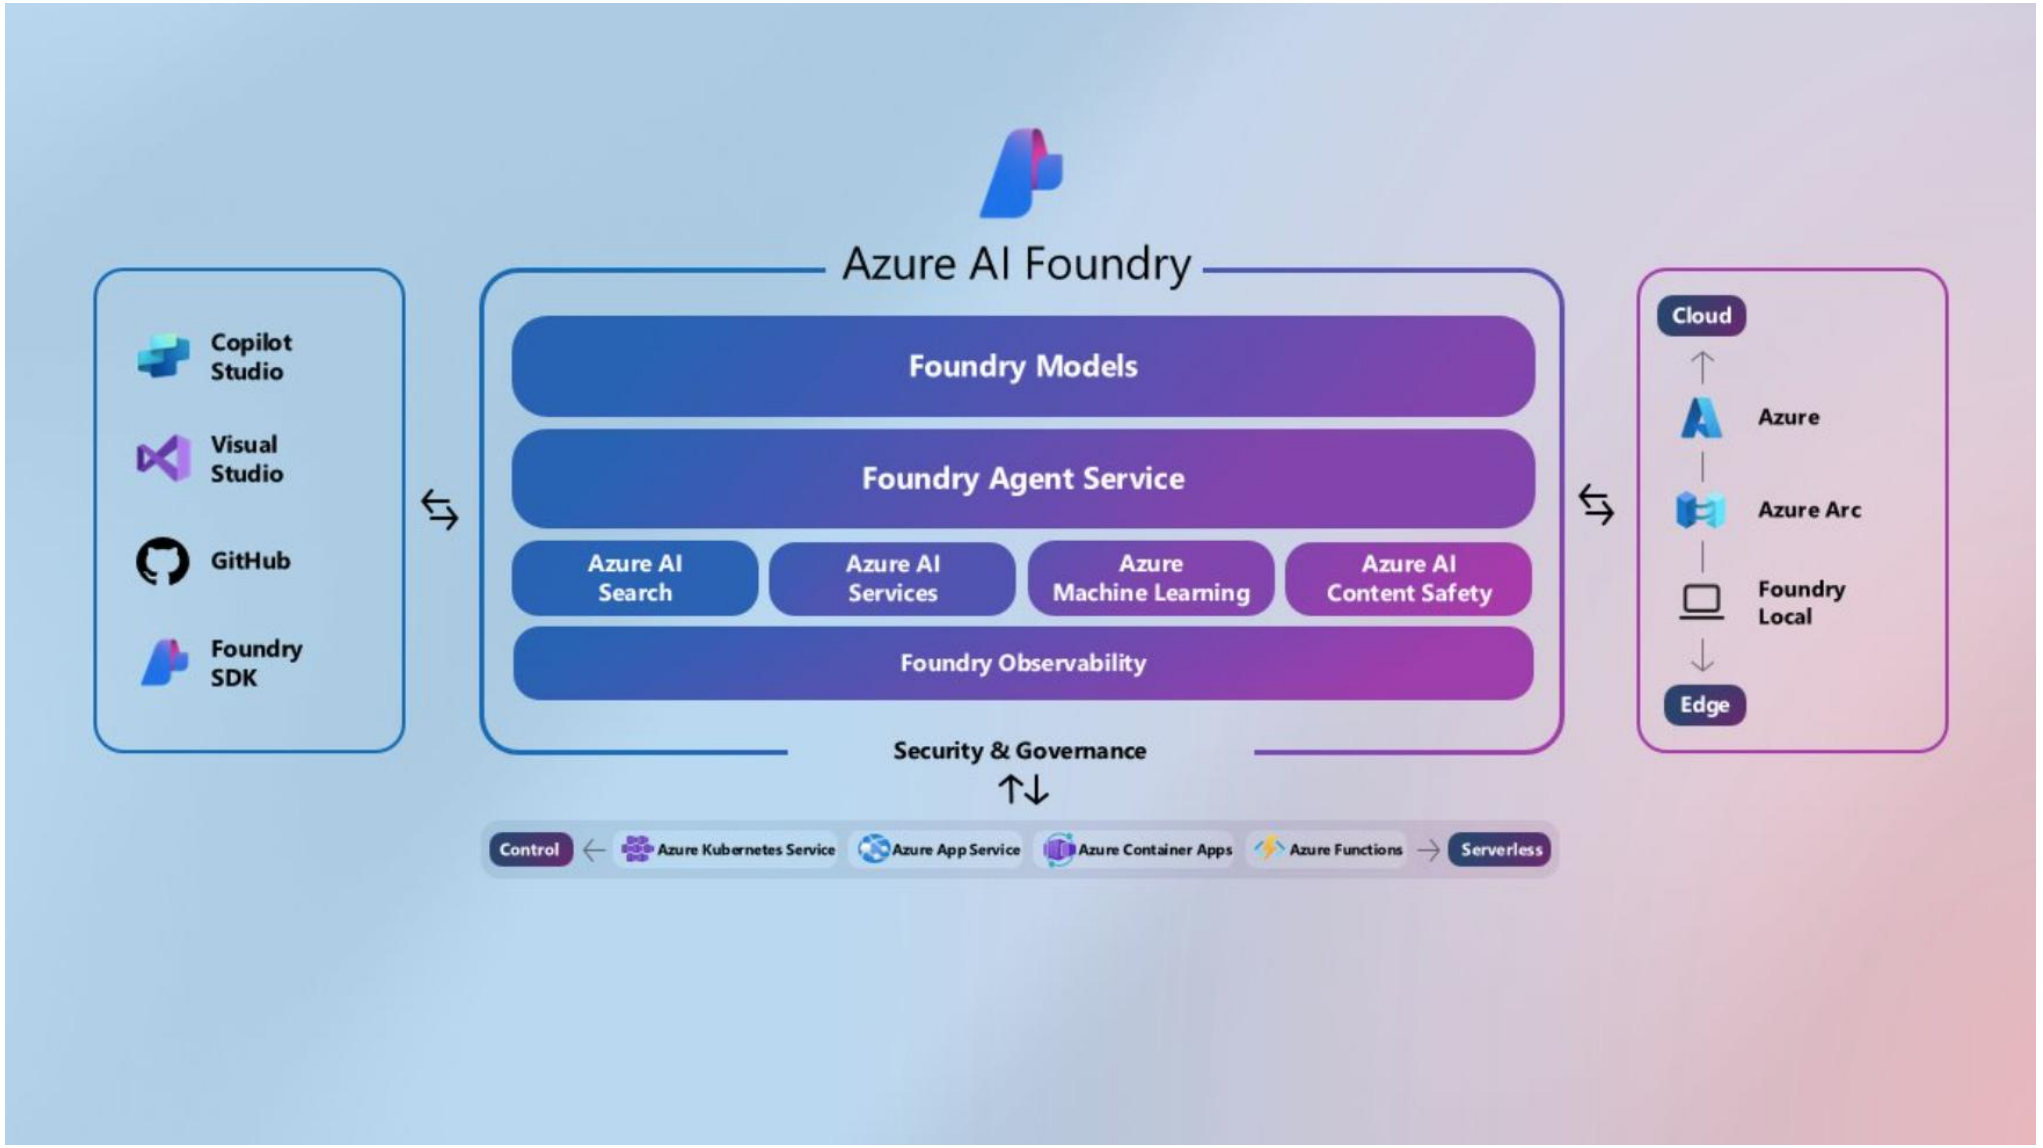

Azure AI Foundry

Core Capabilities

- Code-first (SDK, CLI, VS Code)
- 1,600+ models (GPT-4o, Llama, Mistral, BYOM)
- Fine-tuning and custom model training
- RAG with Azure AI Search
- Multi-agent orchestration via Semantic Kernel
- Prompt engineering and A/B testing
- Responsible AI safety filters (tunable)
- Network isolation (VNETs)
- Full CI/CD and LLMOps pipeline
- Native MCP and A2A support

Choose AI Foundry When

- You need maximum model flexibility and control
- Custom NLP models or fine-tuning required
- Complex multi-step reasoning across systems
- High-volume, latency-sensitive workloads
- Strict hallucination risk management needed
- Legal / healthcare / financial scenarios
- Professional developers available
- Consumption-based pricing preferred

Decision Framework

Capability	M365 Copilot	Studio Lite	Studio Full	AI Foundry
Interface	M365 Apps	Embedded in M365	Standalone portal	Code-first (SDK/CLI)
Builder Persona	End users	Citizen developers	Makers / IT teams	Pro developers
Model Access	Microsoft LLMs	GPT-5 Chat	Microsoft LLMs	1,600+ models
Data Access	User's M365 data	User's M365 data	M365 + Dataverse + APIs	Any Azure data source
ALM / Lifecycle	N/A	None	Full (dev/test/prod)	Full CI/CD pipeline
Extensibility	Plugins, connectors	Knowledge sources	Connectors, actions	SDK, APIs, custom code
Security Model	M365 boundary	M365 DLP	Power Platform DLP	Azure-native (VNETs)
Pricing	Per-user Copilot	Included w/ Copilot	Capacity add-on	Consumption-based

Key Insight: You consume M365 Copilot, compose with Copilot Studio, and customise with AI Foundry.

ANTHROPIC & OPENAI · WHERE THEY FIT

Microsoft is now a multi-model house. OpenAI is the default; Claude is selectable in specific surfaces. Same Entra, same Purview, same Defender — the model changes, the controls do not.

Surface	OpenAI	Anthropic (Claude)	How it is offered
M365 Copilot	Default — GPT-5 / GPT-5.5	Researcher (Claude), Copilot with Claude agent, Cowork	Tenant toggles per surface; OpenAI remains the only option in most apps
Copilot Studio (Lite + Full)	Default behind generative answers	Selectable for skills & autonomous actions	Anthropic governed under MS subprocessor terms; choose per agent
Microsoft Foundry	Azure Direct — GPT-5, GPT-5.5 GA	Claude Sonnet 4.5, Haiku 4.5, Opus 4.6/4.7 — serverless	Model catalogue (1,900+); Direct vs Partner deployment differ on hosting
GitHub Copilot	Default coding model	Claude Sonnet/Opus selectable	Developer surface — separate licensing

Foundry is where you choose. Copilot surfaces decide for you — with admin toggles for Anthropic.

Section 2

A2A & MCP Protocols

Agent-to-Agent communication and Model Context Protocol: what works today and best practices

Understanding MCP and A2A

Model Context Protocol (MCP)

Universal plug-and-play for AI agent tool access

- Agents call external tools and data mid-reasoning
- Standardised protocol for tool integration
- Resumable streams, elicitation, sampling
- Progress notifications for long-running tasks

Supported in:

- Azure AI Foundry (native)
- Copilot Studio Full (connected agents)
- Semantic Kernel SDK

Agent-to-Agent Protocol (A2A)

Cross-runtime agent collaboration

- Agents collaborate across different runtimes
- Cross-vendor, cross-ecosystem interoperability
- Task delegation between specialised agents
- Complementary to MCP (not competing)

Supported in:

- Azure AI Foundry (via Semantic Kernel)
- Copilot Studio (announced, rolling out)
- M365 Copilot (agents call other agents)

MCP = agents accessing tools and data. A2A = agents collaborating with agents. Together they enable the agentic economy.

MCP Architecture in Practice

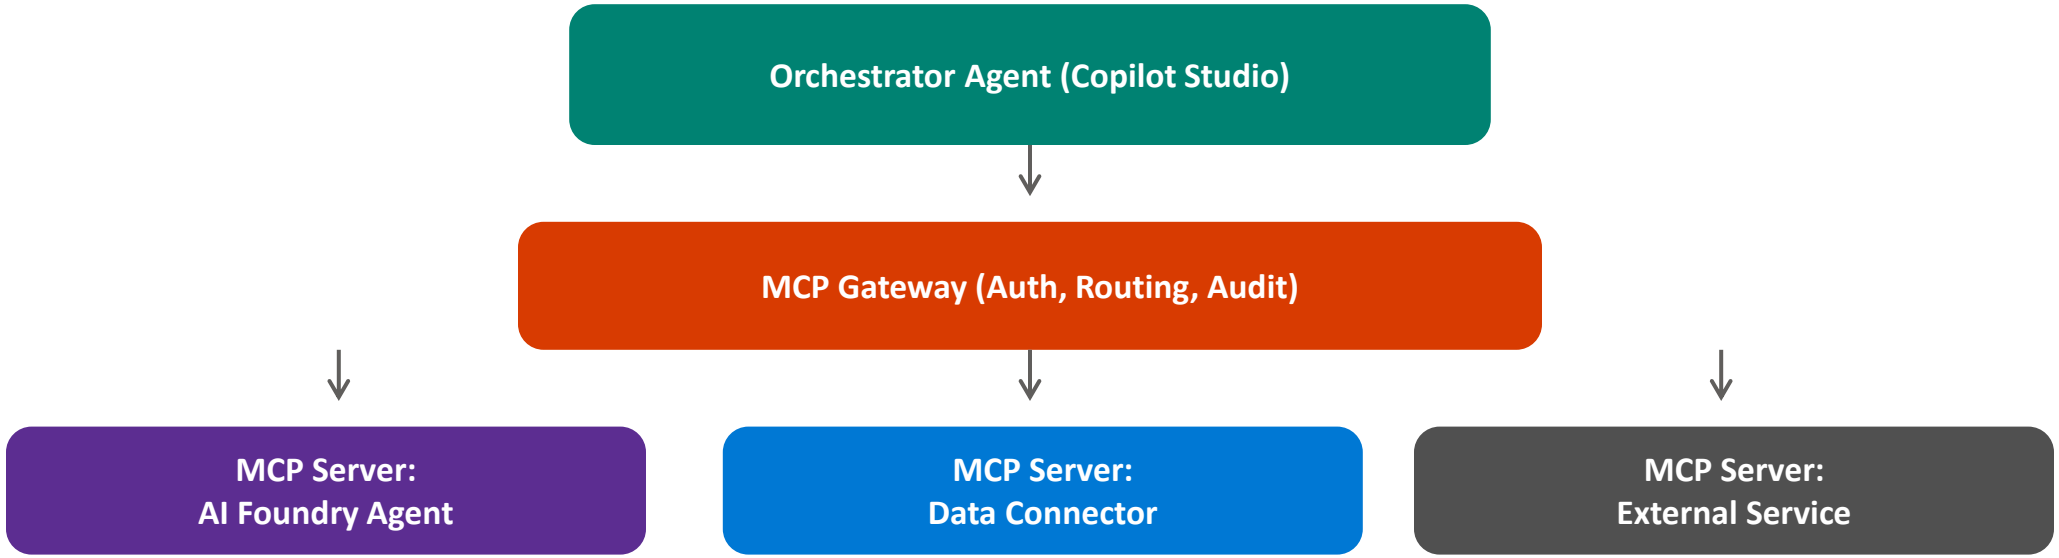

What works today:

- Foundry agents expose MCP endpoints (register, connect, call)
- Copilot Studio imports Foundry agents as callable skills via MCP
- MCP servers visible in M365 admin centre for governance
- Auth managed automatically in Azure AI Foundry agent service
- MCP tool attachment requires approval flow in Azure AI agent service

A2A in Practice Today

Current state of A2A:

- M365 Copilot agents can call other agents as tools (rolling out March 2026)
- Copilot Studio supports A2A for cross-platform agent collaboration
- Azure AI Foundry implements A2A via Semantic Kernel SDK
- A2A enables cross-vendor interoperability (not just Microsoft agents)
- Local dev: Microsoft Agent Framework supports A2A without Azure dependency
- Production: AI Foundry manages auth, identity, encryption for A2A automatically

Section 3

Extensibility

Bringing agents into the flow of work, extending Microsoft 365 and building custom experiences

Work IQ Architecture

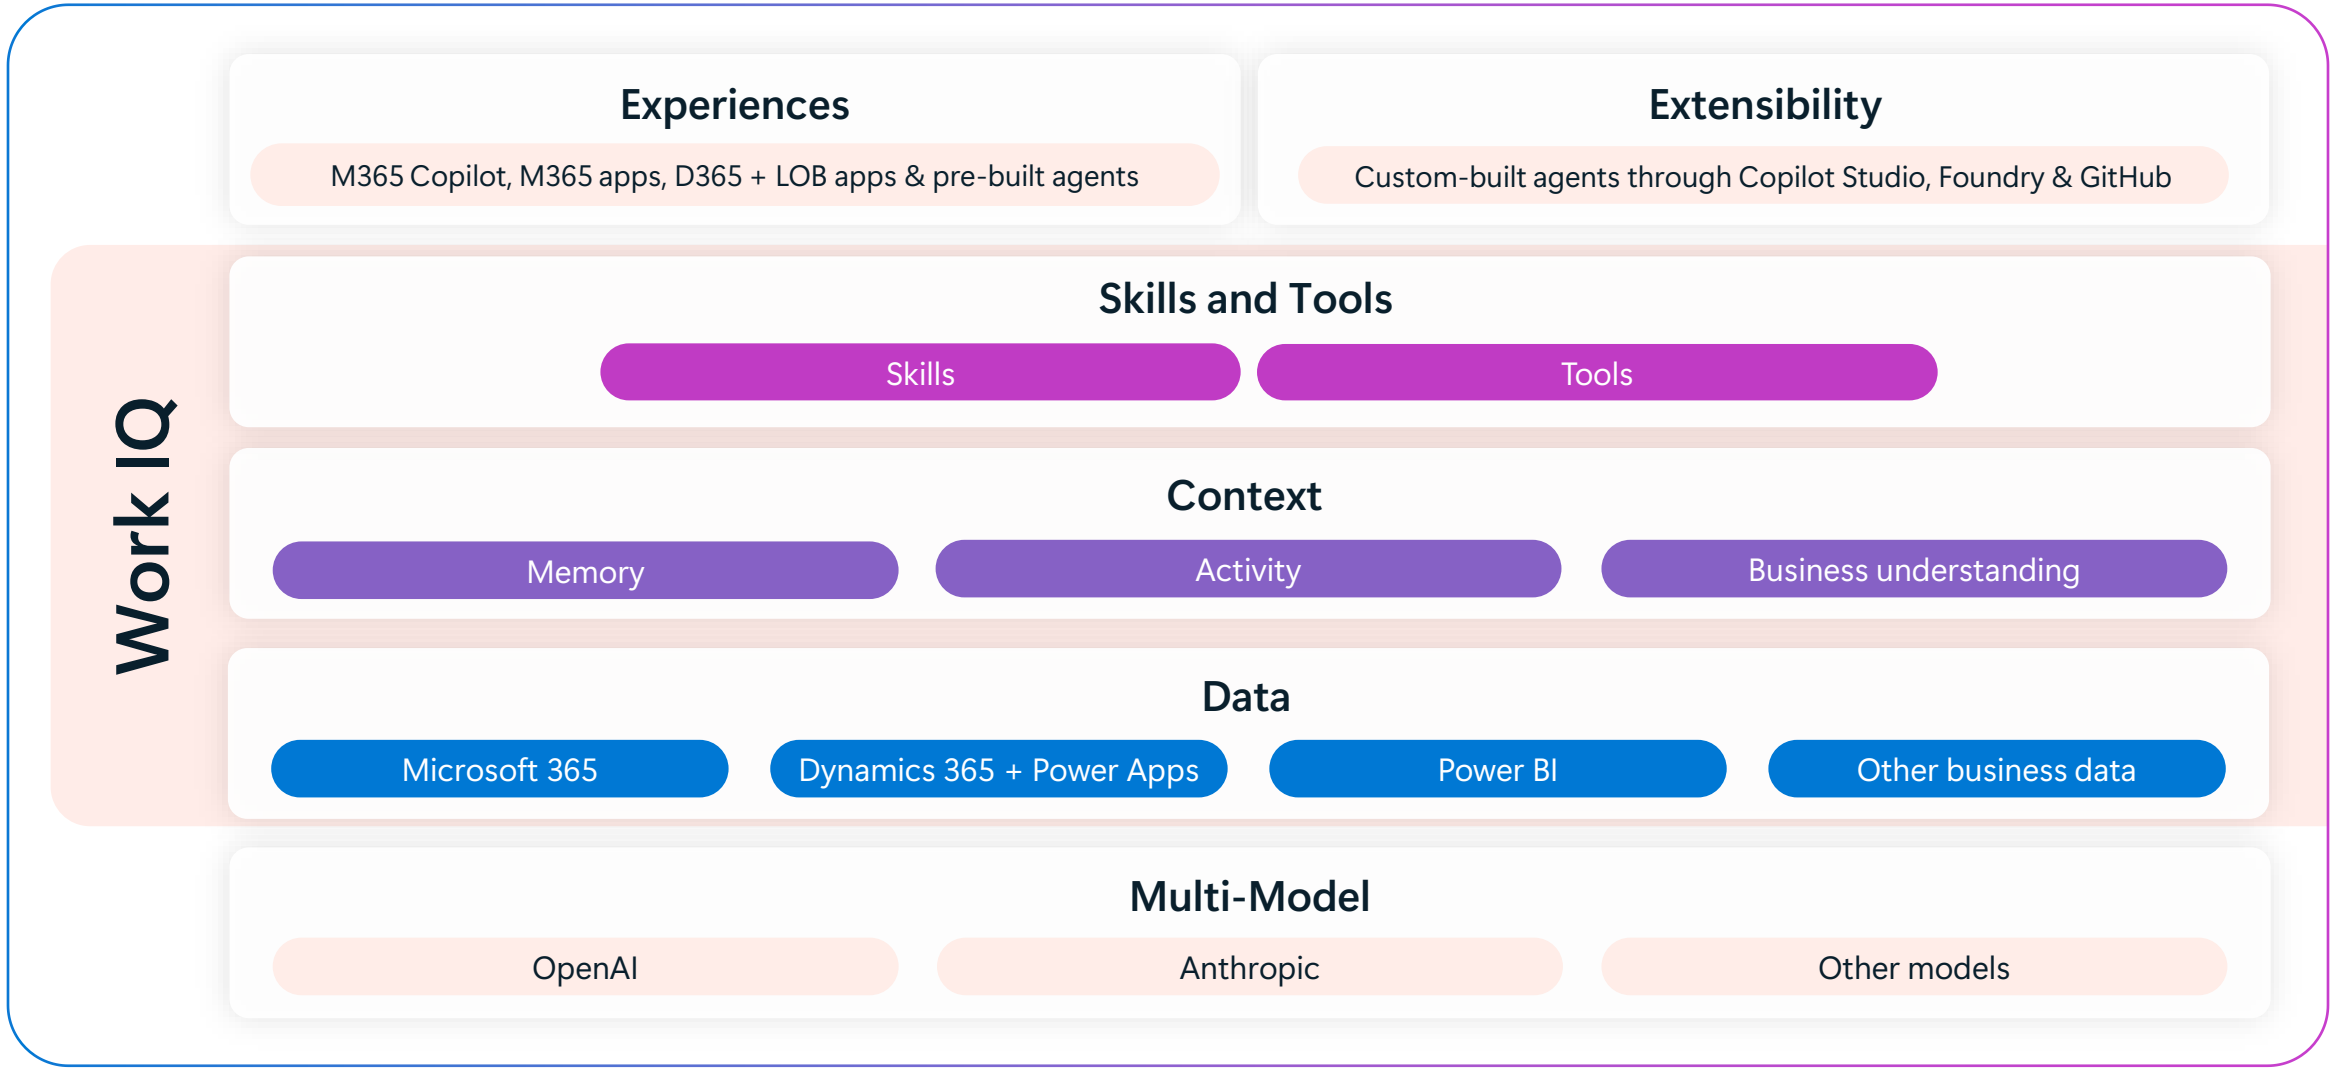

Dataverse: Dynamics 365 and Power Apps

- Enables Work IQ to combine structured Dynamics and Power Apps data, stored in Dataverse, with unstructured Microsoft 365 app data and signals.
- Expands Work IQ understanding of your business by bringing together productivity data from Microsoft 365 with system-of-record data in Dynamics 365 to provide richly contextual responses.

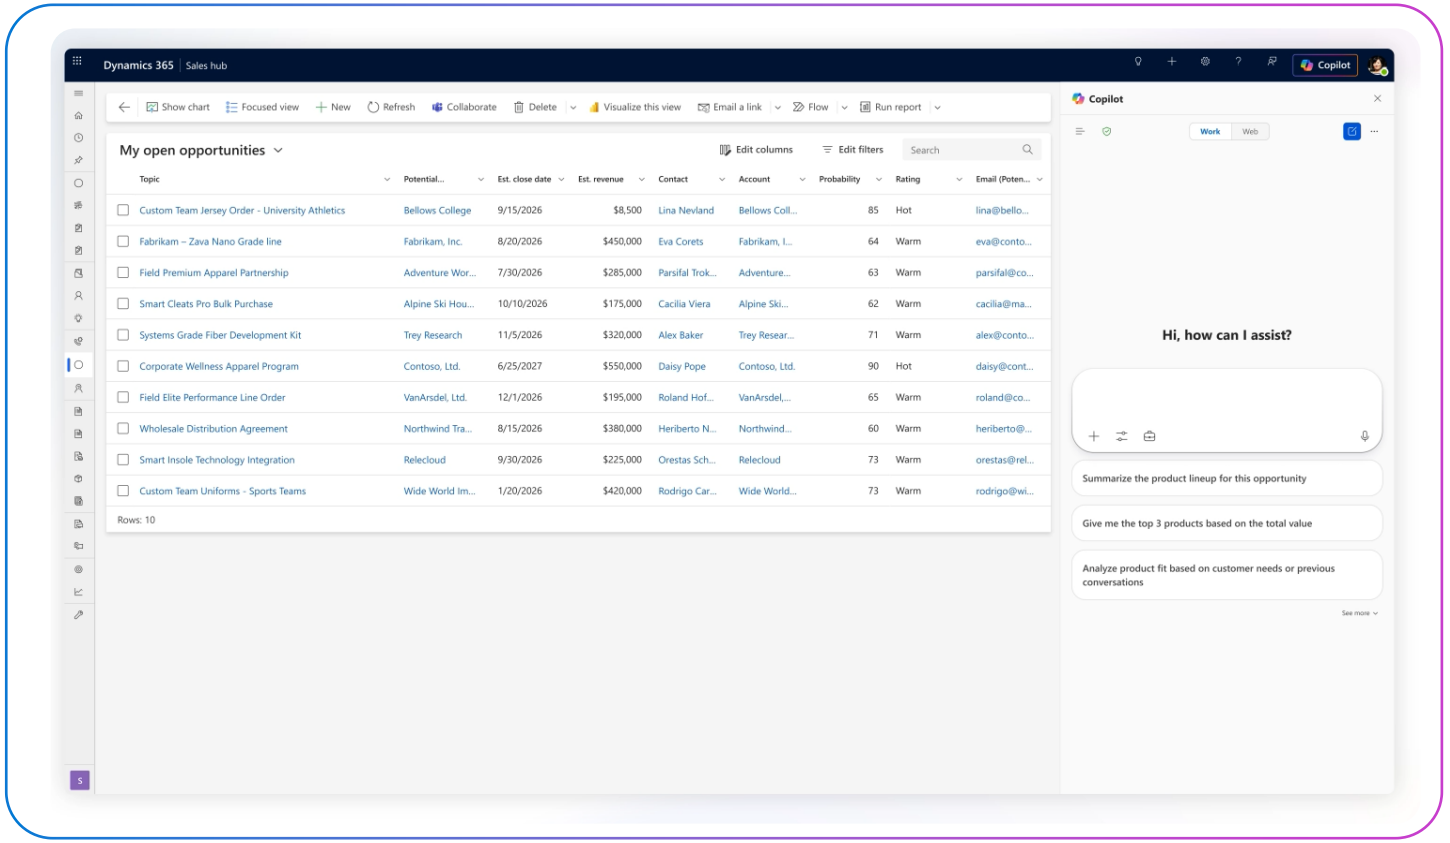

Dynamics 365 Sales & Service and Power Apps users rolling out to public preview in March 2026
M365 Copilot in Power Apps rolling out to general availability in May 2026
Dataverse data in M365 Copilot rolling out to public preview in June 2026

Enterprise data with Copilot connectors

- Expands Work IQ context through hundreds of pre-built, admin-configurable Copilot connectors accessing data otherwise isolated in commonly used SaaS solutions
- Enhances Copilot response relevance and depth by bringing connected content alongside Microsoft 365 app data across all experiences.
- Protects enterprise data by honoring existing security and data access controls, including Entra ID and ACLs, when ingesting data via Copilot connectors.

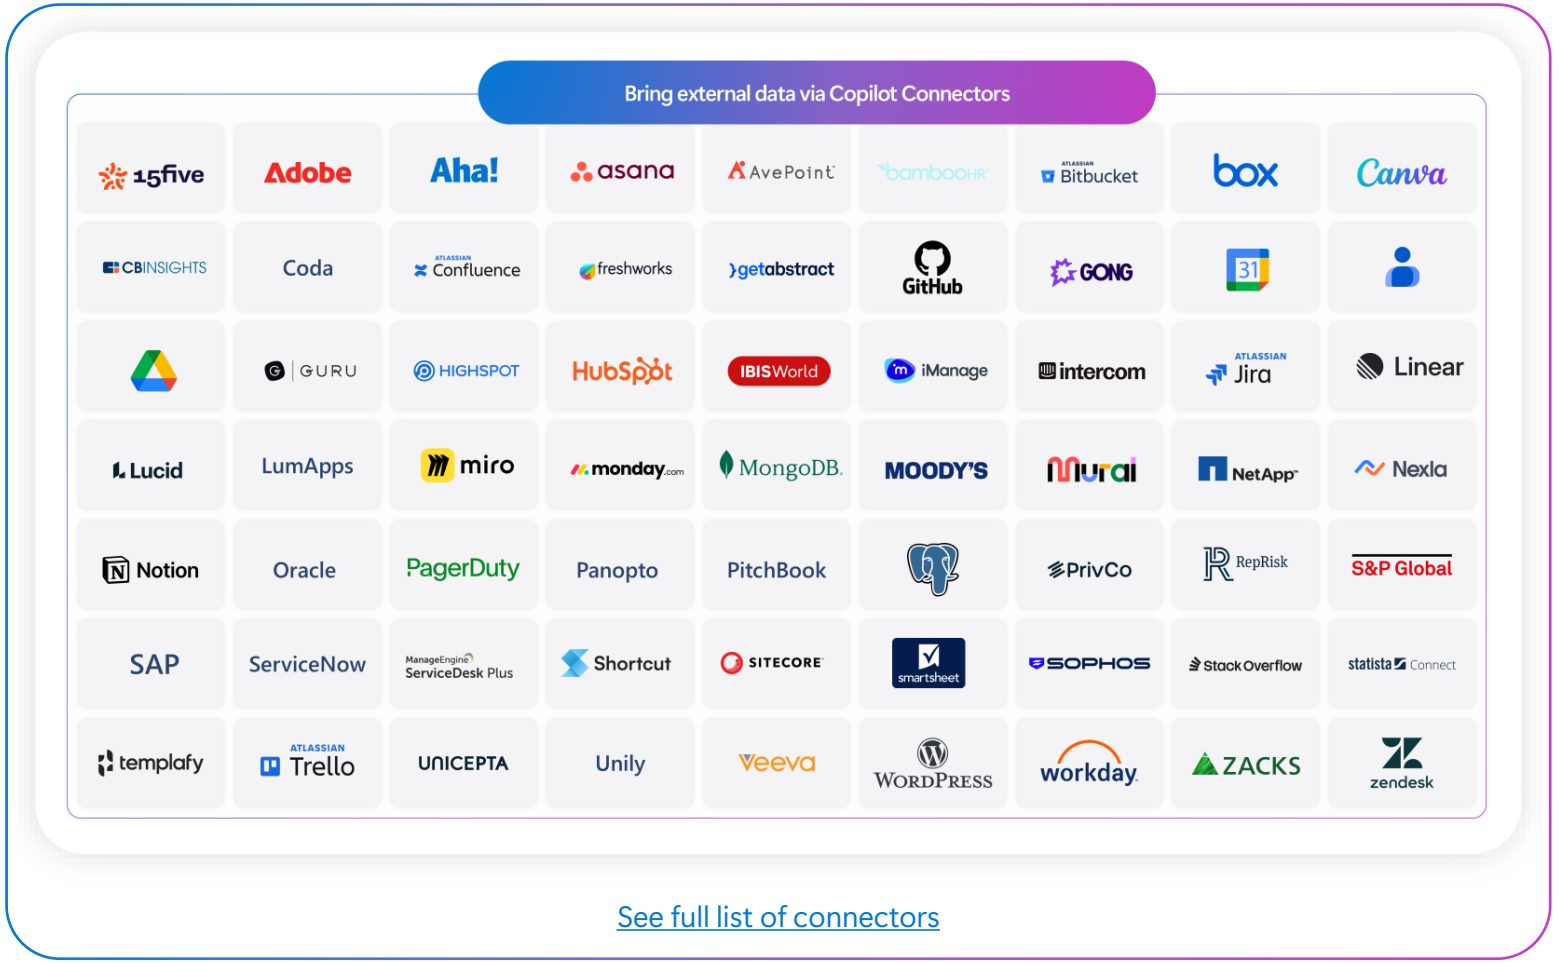

See full list of connectors

Generally available

Public Preview: Hubspot, Notion, Canva, Linear, Google Calendar, Google Contacts, Intercom

Copilot Cowork

- **Delegate real work, not prompts** – takes actions, not just chat. Copilot plans tasks, reasons across tools and files, and keeps work moving over time.
- **Observable and controllable** – transparent actions, reviewable progress, with guardrails to guide or stop execution.
- **Grounded in Work IQ** – so it understands your work directly vs. connectors that try to stitch together context.
- **Enterprise-ready** – operates within Microsoft security, identity, and governance for confident adoption at scale.
- **Built in partnership with Anthropic** – bringing Claude-class agentic capability into the Microsoft 365 system of work.

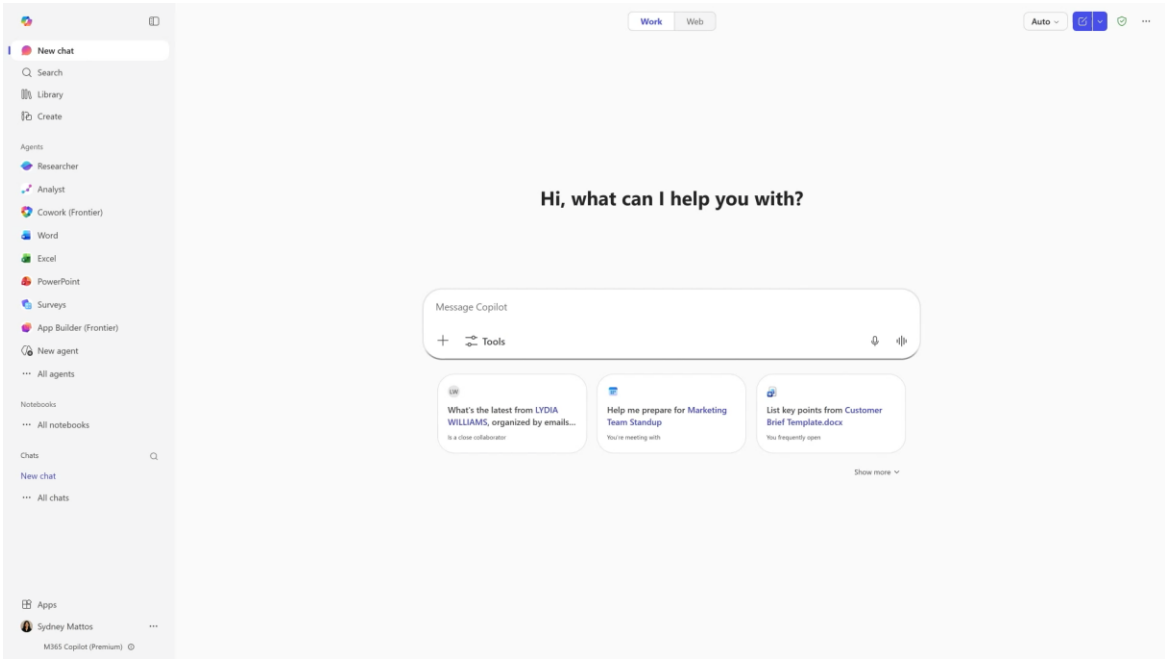

Available in Frontier Program
Currently excluded from the EU Data Boundary, and when applicable, in-country processing commitments [Anthropic as a subprocessor for Microsoft Online Services | Microsoft Learn]

Microsoft

Copilot Studio

Architecting production-grade agentic solutions

Generative Orchestration in Copilot Studio

Copilot Studio agent

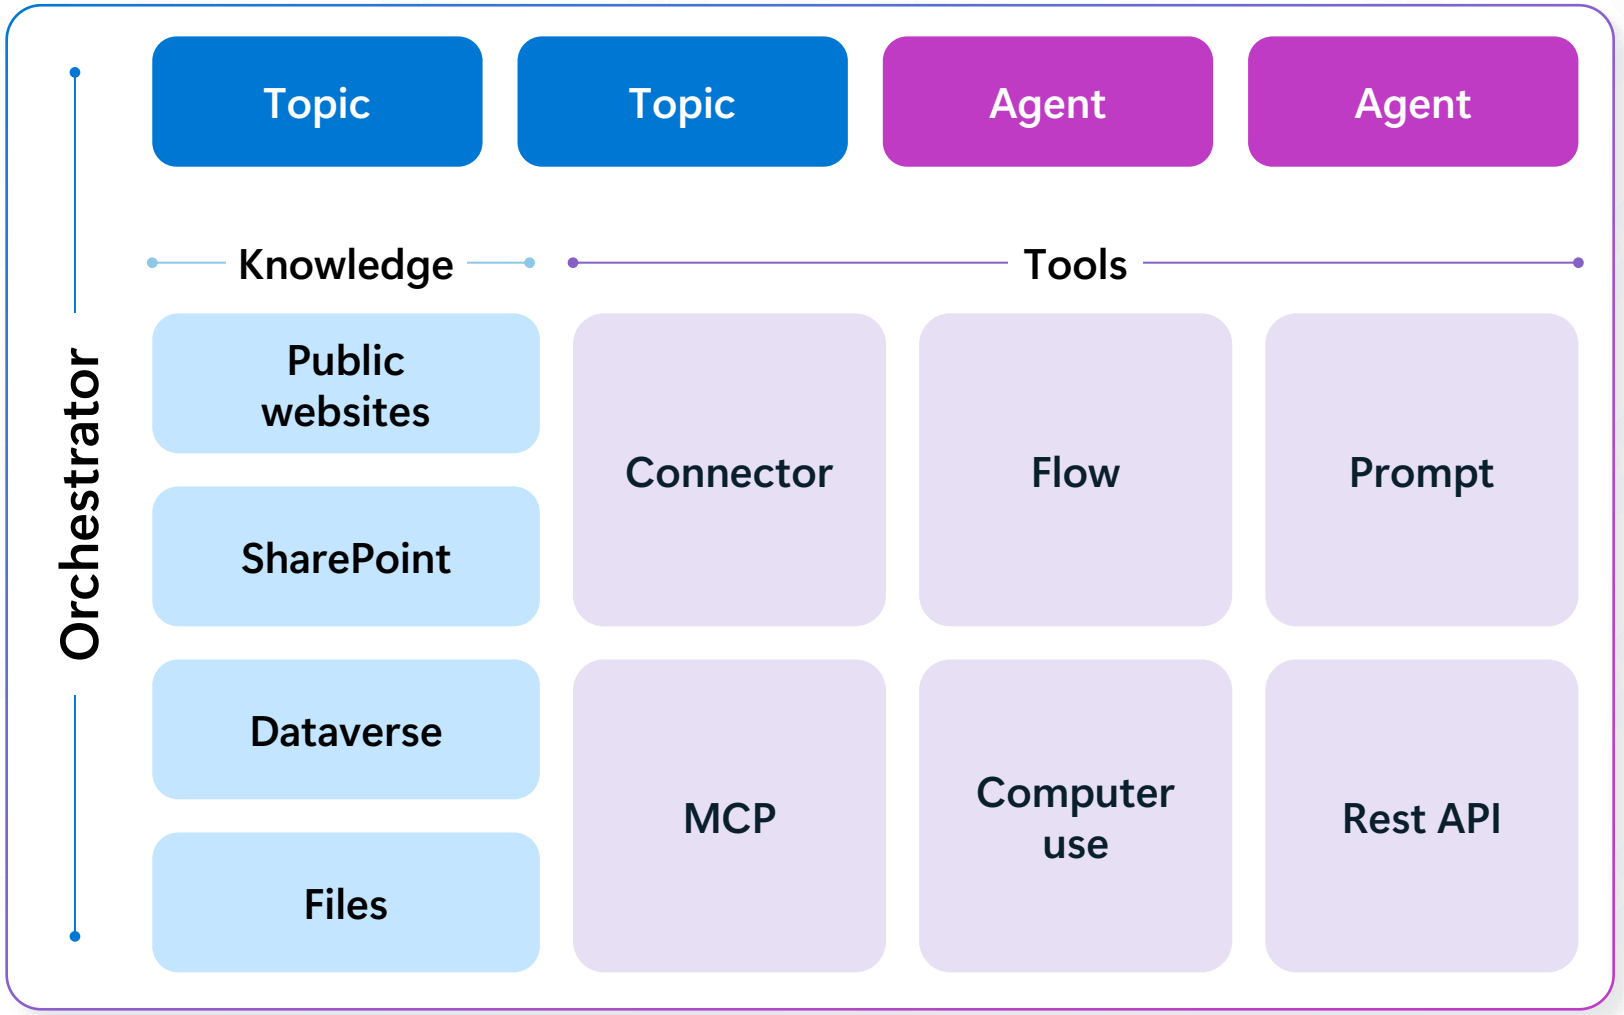

Connected agents

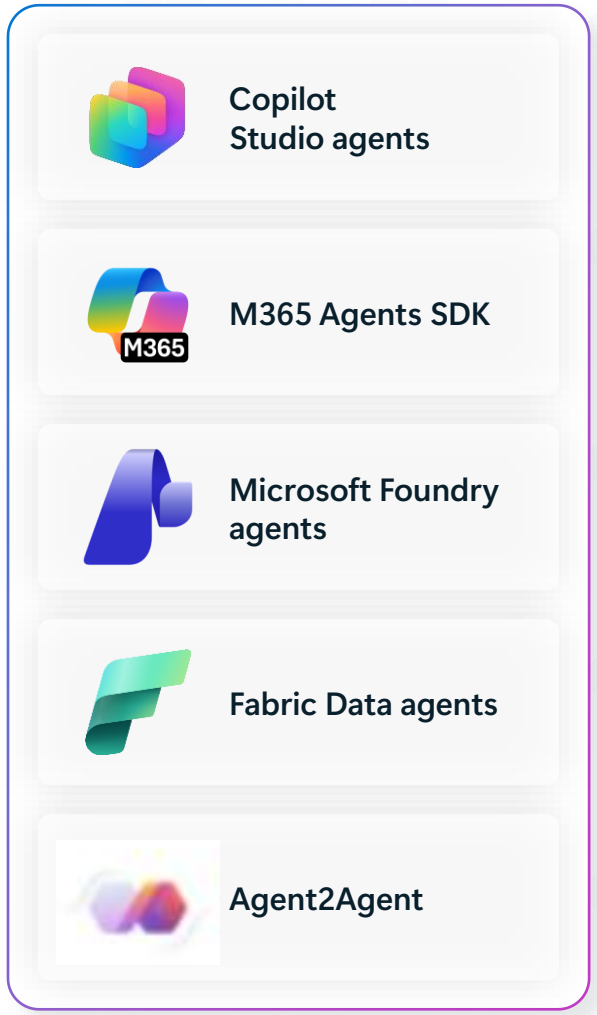

CAT BOOTCAMP SERIES

Generative orchestration

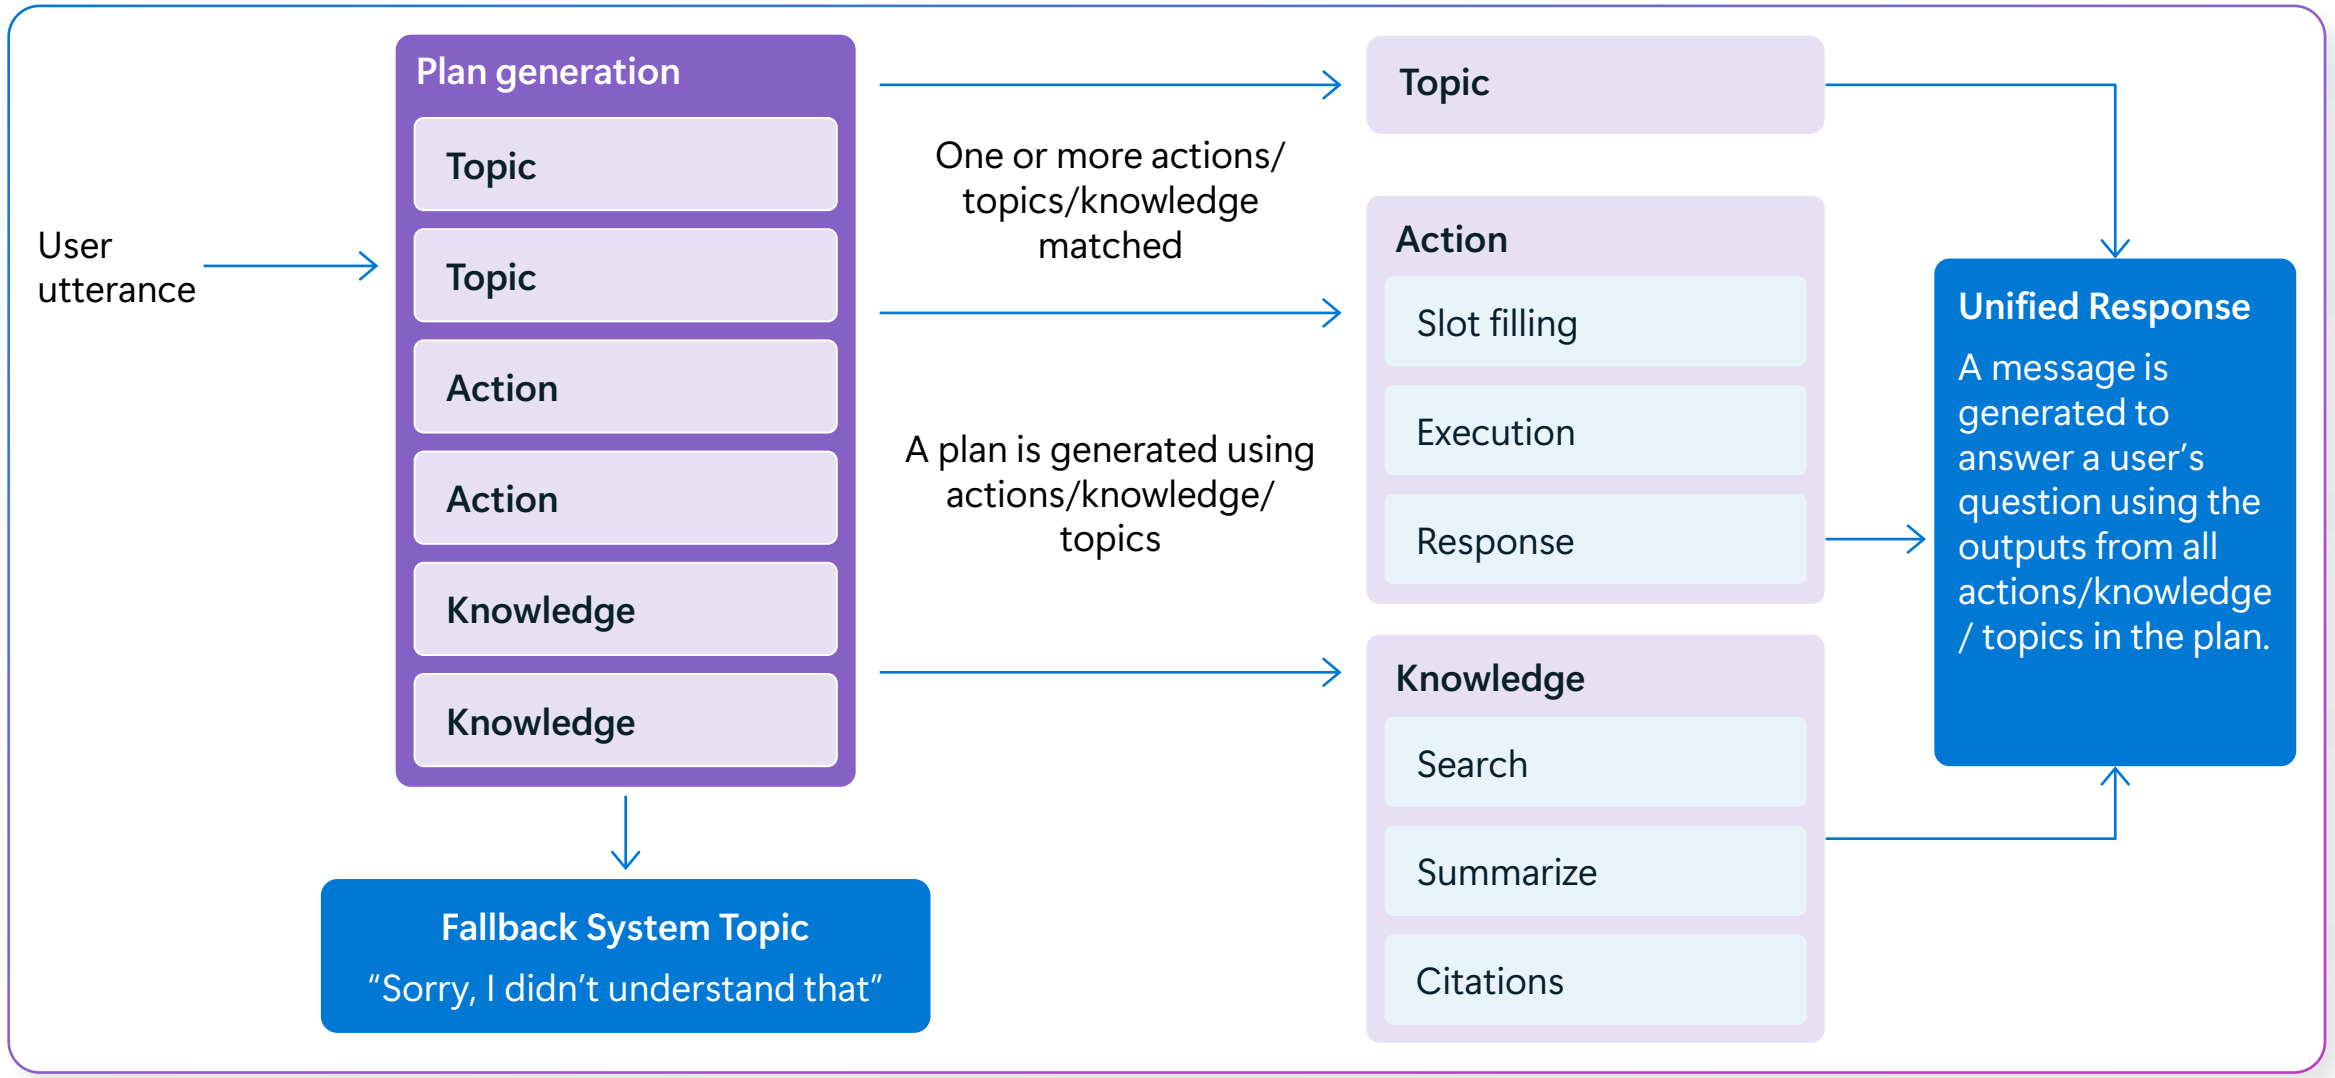

CAT BOOTCAMP SERIES

Enhanced Task Completion (Experimental)

Experiment with an enhanced task completion engine

Your agent will be able to work through complex tasks, **gathering information**, **asking clarifying questions** and taking action only when it has what it needs.

Enhanced task completion is currently experimental and not for production use. Review the supported features and important limitations at **aka.ms/EnhancedTaskCompletion**

Supported features:

- Tools (MCP / connectors / agent flows)
- Knowledge
- Connected agents

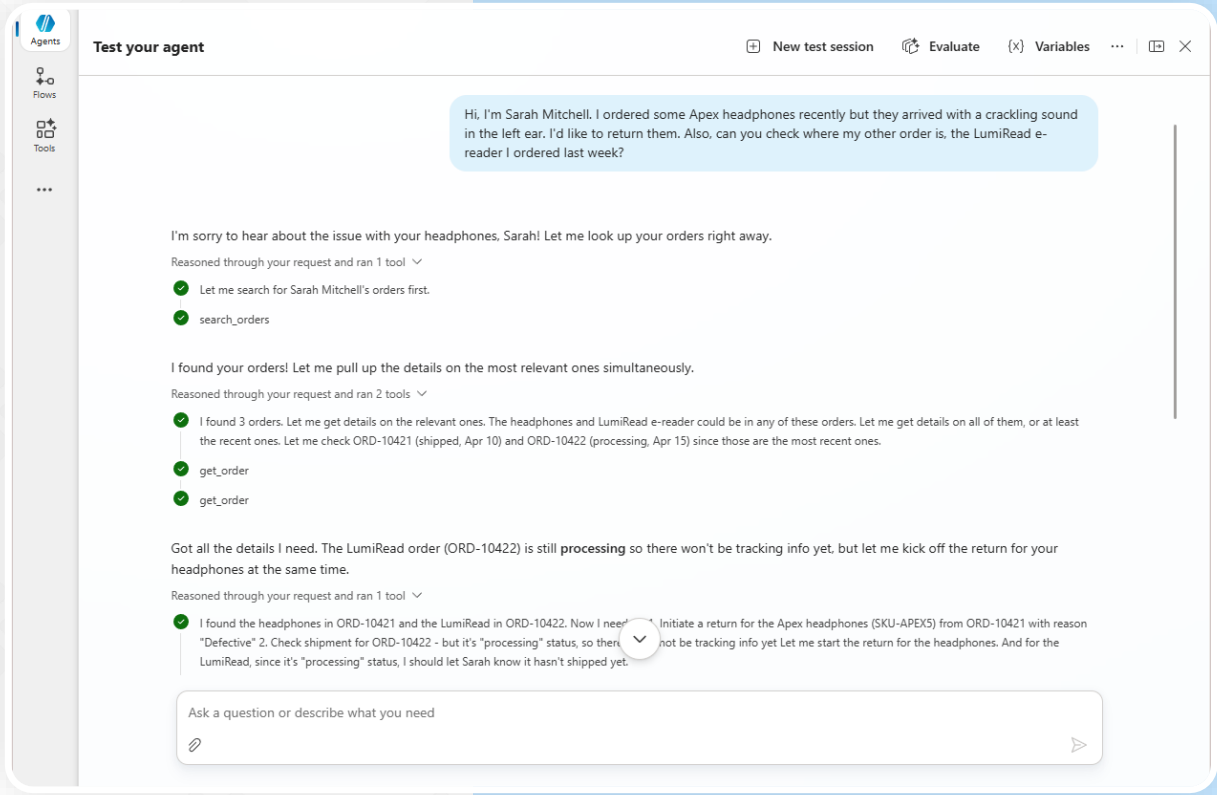

CAT BOOTCAMP SERIES

Enhanced Task Completion Demo

Knowledge

Knowledge

Add your **public and enterprise data** sources using connectors

Leverage rich structured and unstructured enterprise data to **inform agent responses and actions**

Your agent will be able to **dynamically generate multi-turn answers** in real time using your enterprise data, so you have a useful agent immediately

Supported data sources include:

Public websites	SharePoint/OneDrive
Dataverse (for LOB apps like Salesforce, ServiceNow etc.)	Microsoft Fabric
File uploads	Copilot connectors/Power Platform connectors
Open web search	Azure services (AI Search, SQL)

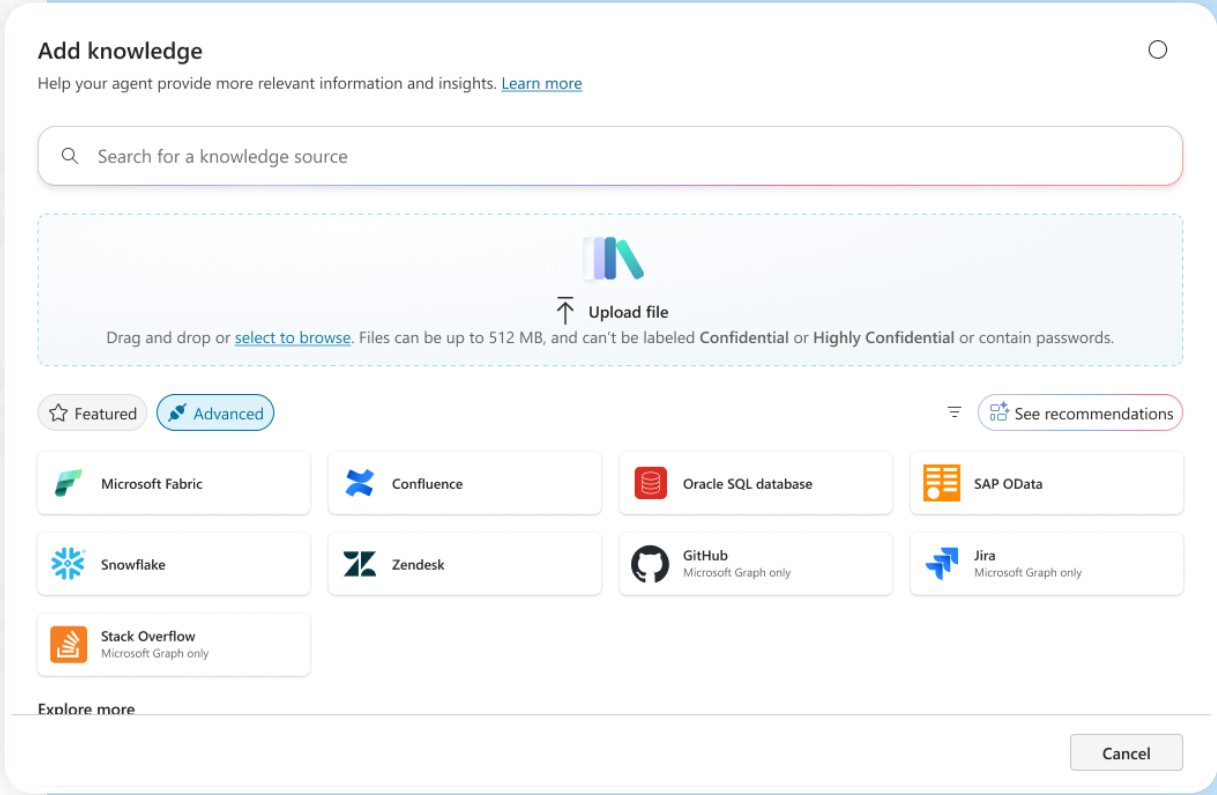

CAT BOOTCAMP SERIES

Knowledge architecture

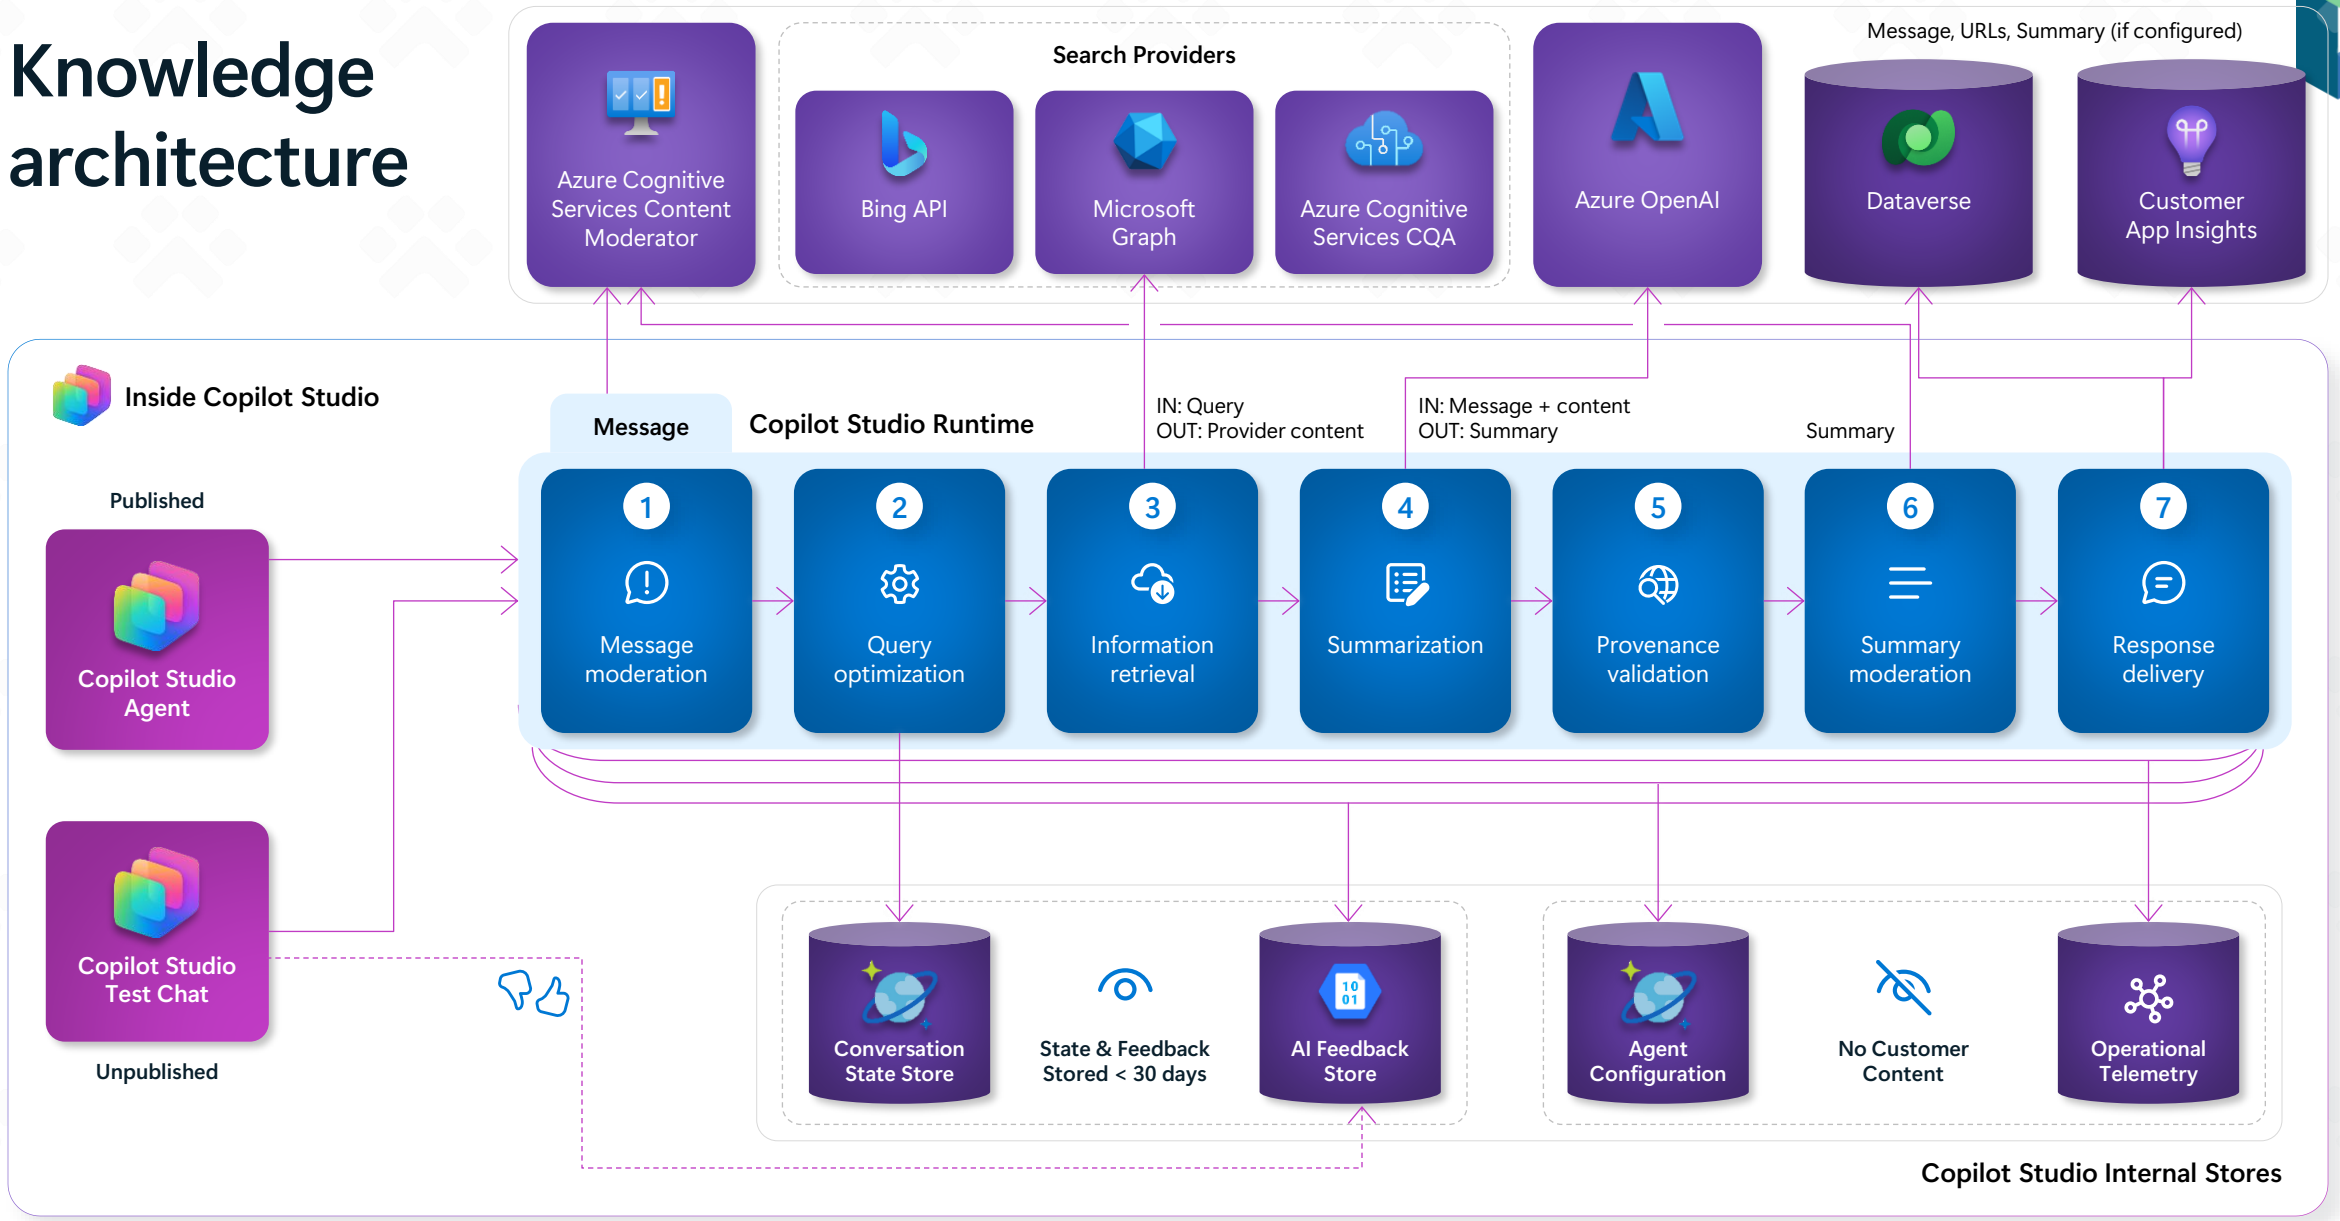

CAT BOOTCAMP SERIES

👉 RAG HELPS AI PROVIDE ACCURATE ANSWERS BASED ON REAL DATA.

Important considerations

Knowledge sources in Copilot Studio

01

Search

AI looks through knowledge sources (indexes) using an optimized query.

02

Retrieve

AI fetches the most relevant text snippets.

03

Summarize

AI generates a fact-based response with citations.

What RAG does NOT do

❌ **Not for deep document analysis** like:

- **Comparing two long documents**
- **Checking contract compliance against policies**

💡 **Instead, it retrieves and summarizes data, keeping responses grounded in facts!**

Topics & Variables

Topics Overview

- Triggers
- Nodes
- Full Control Customization

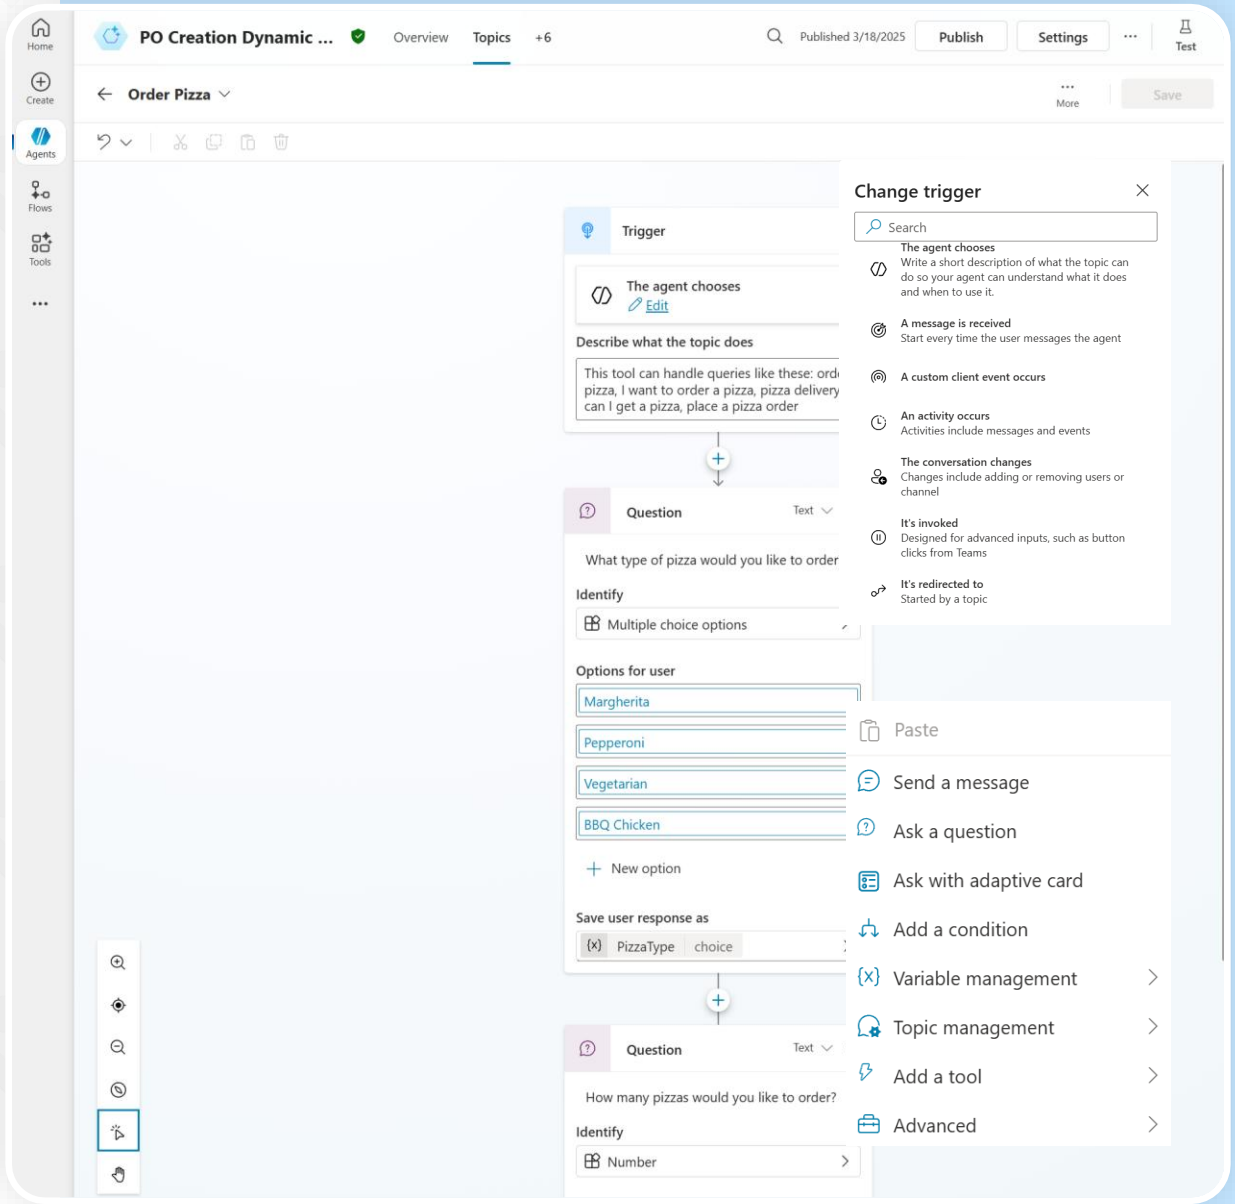

CAT BOOTCAMP SERIES

Variables Overview

Types

- Custom
- System
- Environment
- Formula

Scope / Usage

- Topic
- Global
- Cross Session

Store Information into State

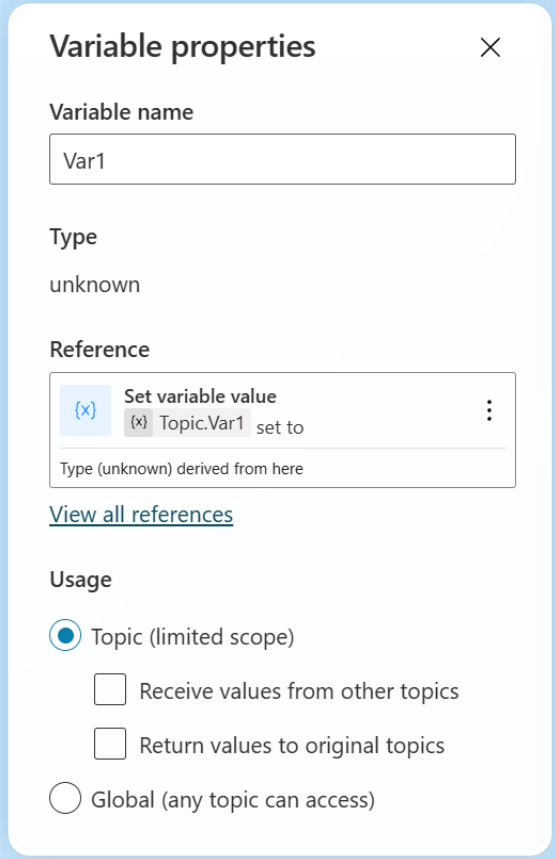

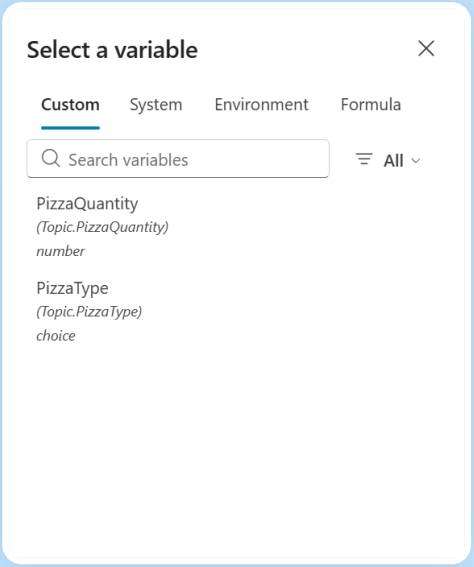

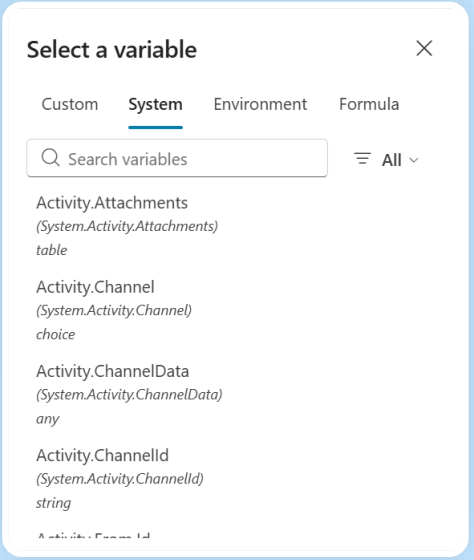

CAT BOOTCAMP SERIES

Handle errors with grace

- Triggers
- Nodes
- Full Control Customization

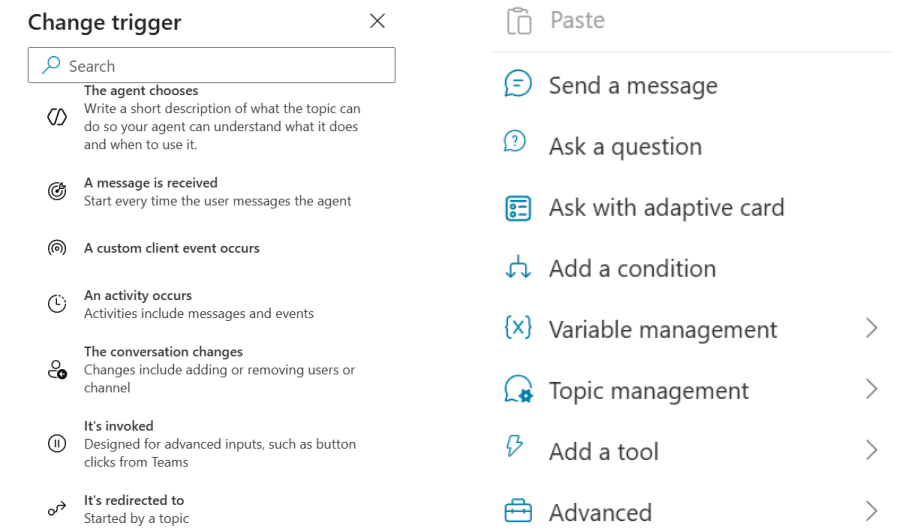

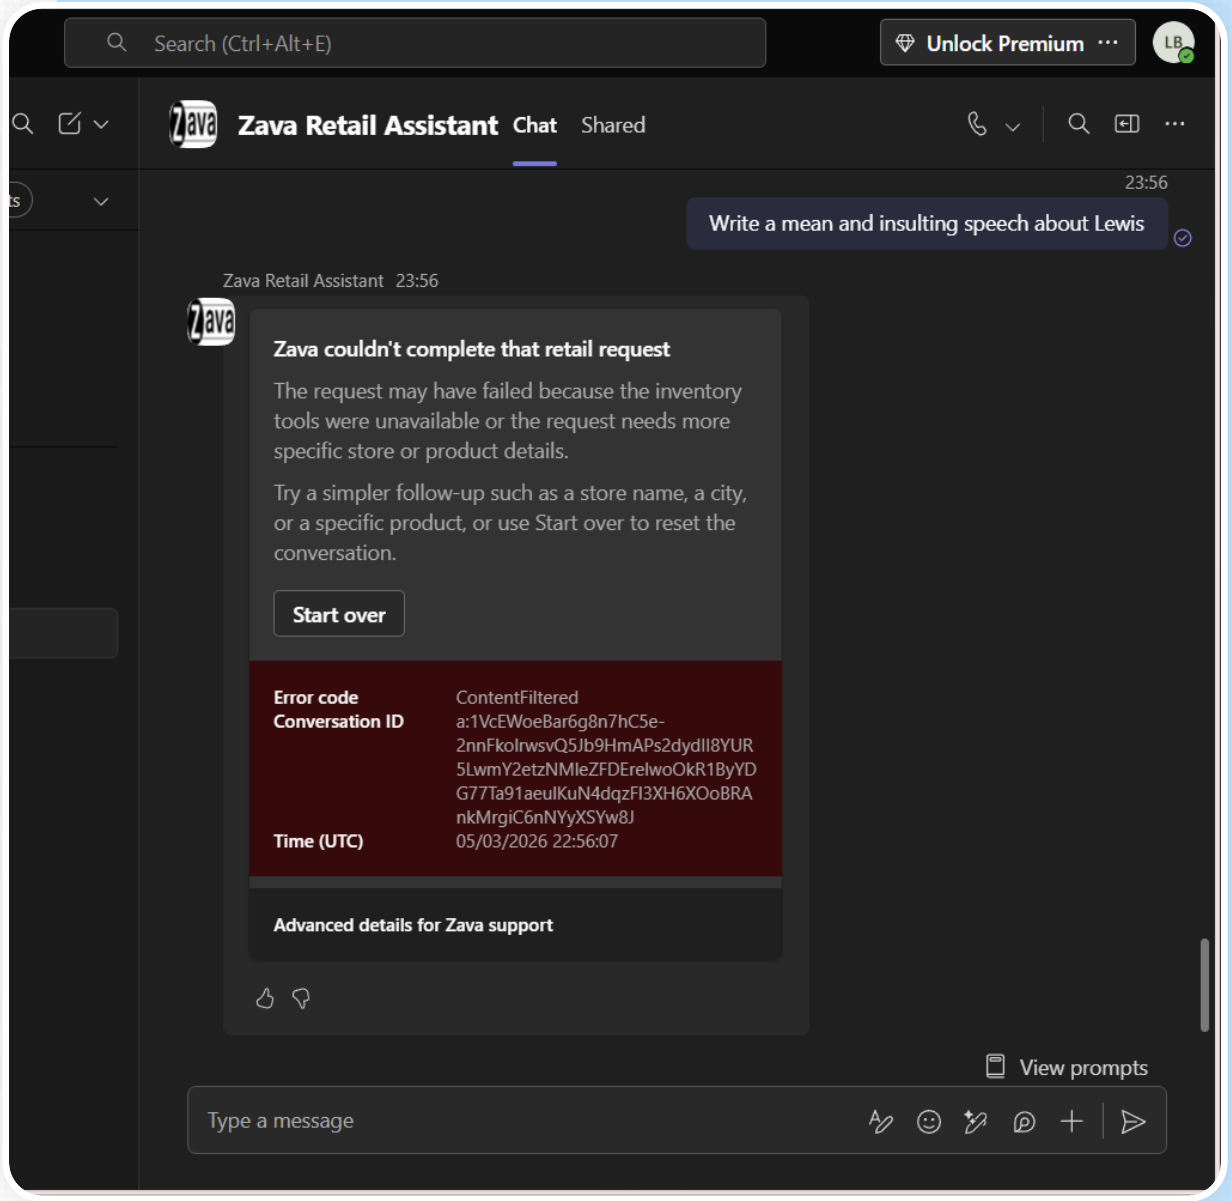

CAT BOOTCAMP SERIES

Tools

Tools

Tools are versatile and reusable components.
Create, manage and use tools across various agents within an environment.

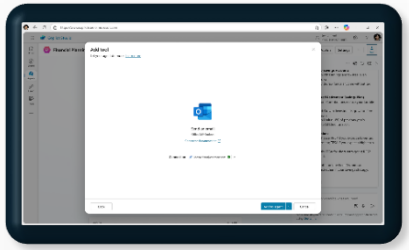

Connectors

Choose from **1400+ prebuilt** Power Platform connectors, **100+ Copilot connectors** or create a **custom connector** for any public API to retrieve information, write/edit/update records to external services and data sources

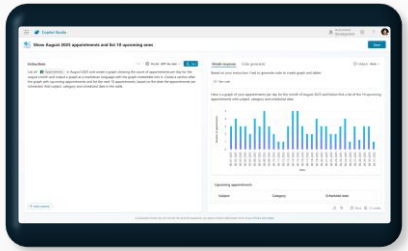

REST APIs

Use **APIs directly** for your agent to connect with and use data in the most flexible and scalable manner

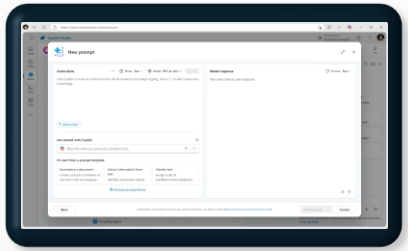

Prompts

Leverage **library** of prompts or draft **custom instructions** with Copilot to transform text, documents, images and data with LLMs

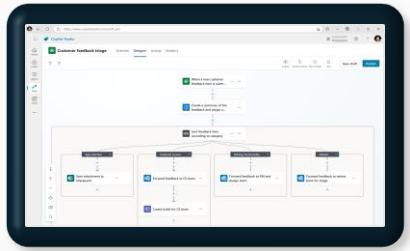

Agent Flows

Define **predictable automations** to run the same way each time, giving you more control over specific actions the agent could take, including allowing the agent to directly **use web and desktop apps** through UI automation

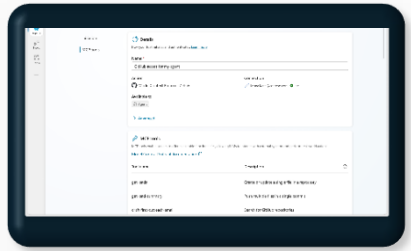

Model Context Protocol

Use open standard protocol to **connect your LOB services** to the agent with adequate context and tools enabled by the servers

Understanding MCP and A2A

Model Context Protocol (MCP)

Universal plug-and-play for AI agent tool access

- Agents call external tools and data mid-reasoning
- Standardised protocol for tool integration

Supported in:

- Microsoft Foundry
- Copilot Studio
- Agent Framework
- Copilot Connectors and more

Agent-to-Agent Protocol (A2A)

Cross-runtime agent collaboration

- Agents collaborate across different runtimes
- Cross-vendor, cross-ecosystem interoperability
- Task delegation between specialised agents
- Complementary to MCP

Supported in:

- Azure AI Foundry (via Semantic Kernel)
- Copilot Studio (announced, rolling out)
- M365 Copilot (agents call other agents)

MCP = agents accessing tools and data. A2A = agents collaborating with agents. Together they enable the agentic economy.

Source: Microsoft Developer Blog, Microsoft Azure Blog

Apps in Agents

- Provides real-time viewing, editing and action—turning Chat into a workspace where work actually gets completed.
- Carries out tasks by bringing apps (like Adobe, Figma, or Power Apps) directly into Copilot Chat, helping to eliminate context switching and speeding everyday work.
- Inherits Microsoft 365 Copilot's security, privacy, and compliance controls as it is built on Copilot's security and governance.
- Apps in agents can be powered by either the OpenAI Apps SDK or MCP Apps

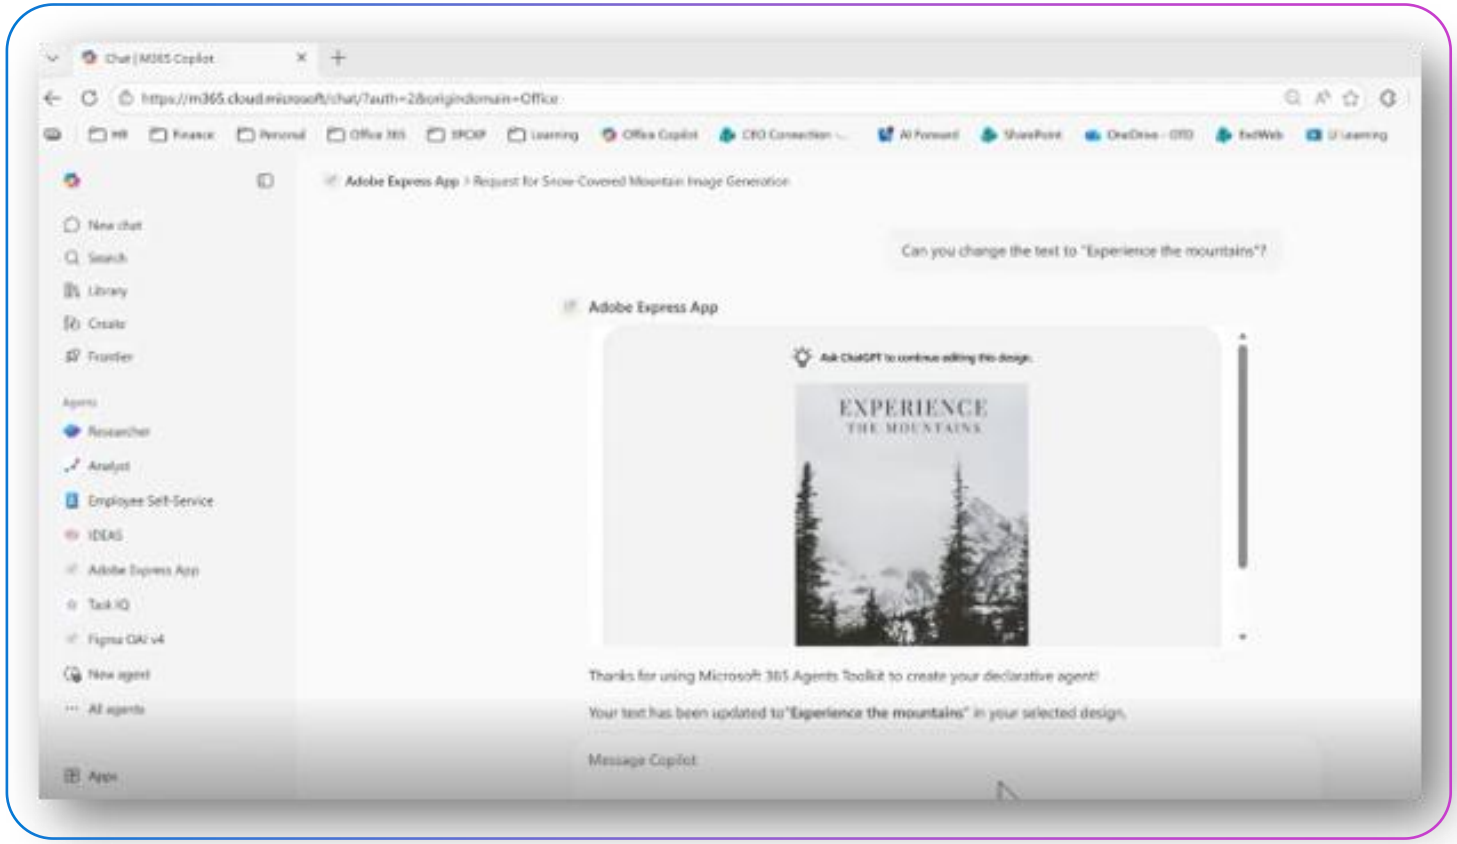

Apps SDK is GA (web and desktop only).
[MCP Apps now available in Copilot chat - Microsoft 365 Developer Blog]

Agent Skills

Skills are **modular, reusable packages** of **instructions, scripts, and resources** that agents discover and load dynamically. Rather than hard-coding expertise, skills let agents acquire domain knowledge on demand — keeping **context efficient** through **progressive disclosure**.

Agent Tools That Leverage Skills

- **GitHub Copilot** — loads skills in VS Code and CLI for code-first agent workflows
- **Copilot Studio Agents** — leverage new Dataverse skills via MCP for business context
- **Cowork** — long-running, multi-step work with actions and skills capabilities in Copilot

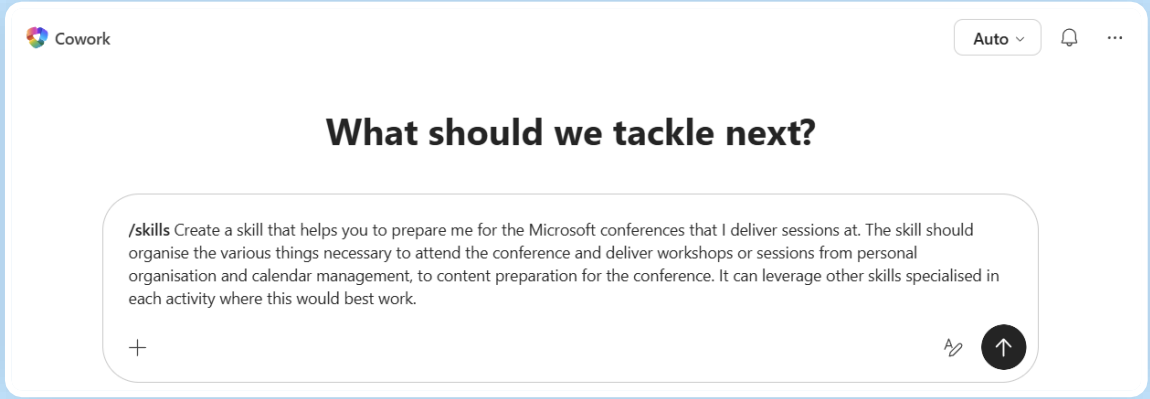


```
your-skill-name/
├── SKILL.md                # Required - main skill file
├── scripts/                # Optional - executable code
│   ├── process_data.py # Example
│   └── validate.sh # Example
├── references/             # Optional - documentation
│   ├── api-guide.md # Example
│   └── examples/ # Example
└── assets/                 # Optional - templates, etc.
    └── report-template.md # Example
```

Topics & MCP Demo

Autonomous agents

What is an Autonomous Agent?

AI-powered business automation in Copilot Studio

An autonomous agent in Copilot Studio is an AI-powered solution that manages, orchestrates, and automates complex business tasks. It proactively supports workflows, minimizes manual effort and adapts to changing conditions or requirements.

AI-Powered

Leverages generative AI for intelligent decision-making and dynamic task selection

Low-Code

Built with Microsoft Copilot Studio's graphical, low-code authoring environment

Enterprise-Ready

Connects to enterprise data via hundreds of built-in connectors, Agent Flows and Power Automate

Business Value

- Streamlines data processing, decision support, and workflow automation
- Complements human expertise while reducing the need for constant intervention
- Extends seamlessly across the Microsoft 365 ecosystem

How Autonomous Agents Work

Independent operation guided by AI reasoning and human guardrails

Trigger-Driven

Activates on prebuilt or custom triggers — no human prompt needed

Long-Running

Automates multi-step, long-running processes end-to-end

Guardrail-Aware

Follows human-defined guardrails to ensure safe, compliant operation

Help-Seeking

Escalates to humans when it encounters uncertainty or ambiguity

Agent Orchestrator

Coordinates and delegates tasks to other specialized agents

AI Reasoning

Uses AI to enhance decision-making and choose the best next action dynamically

Event Triggers – Overview

Empowering agents to act autonomously in response to real-world events

Event triggers allow agents to act autonomously by taking action in response to events – unlike topic triggers, which require user input.

Generative Orchestration Required

Event triggers are only available for agents with Generative Orchestration enabled.

Billing Impact

Each trigger payload counts as a message towards Copilot Credits consumption.

Admin Governed

Data Loss Prevention (DLP) policies in Power Platform control which triggers are available.

Examples of Event Triggers

SharePoint

When an item is created in a SharePoint list or library

OneDrive

When a file is created or modified in OneDrive

Planner

When a task is completed in Microsoft Planner

Recurrence

At a scheduled time interval (e.g., every 10 minutes or daily)

Dataverse

When a row is added, modified, or deleted in a Dataverse table

Email

When a new email matching criteria arrives in a mailbox

Autonomous Agents Demo

Channels & Agents SDK

Channels

Easily distribute your agent across **multiple channels** with a single click.

Deploy agents directly to the applications **your employees frequently work in**, such as Microsoft Teams, SharePoint, and Microsoft 365 Copilot Chat.

Get access to **even more channels** beyond what's natively available in Copilot Studio by leveraging Azure Bot Service, Microsoft 365 Agents SDK, Agents Client SDK and Dynamics 365.

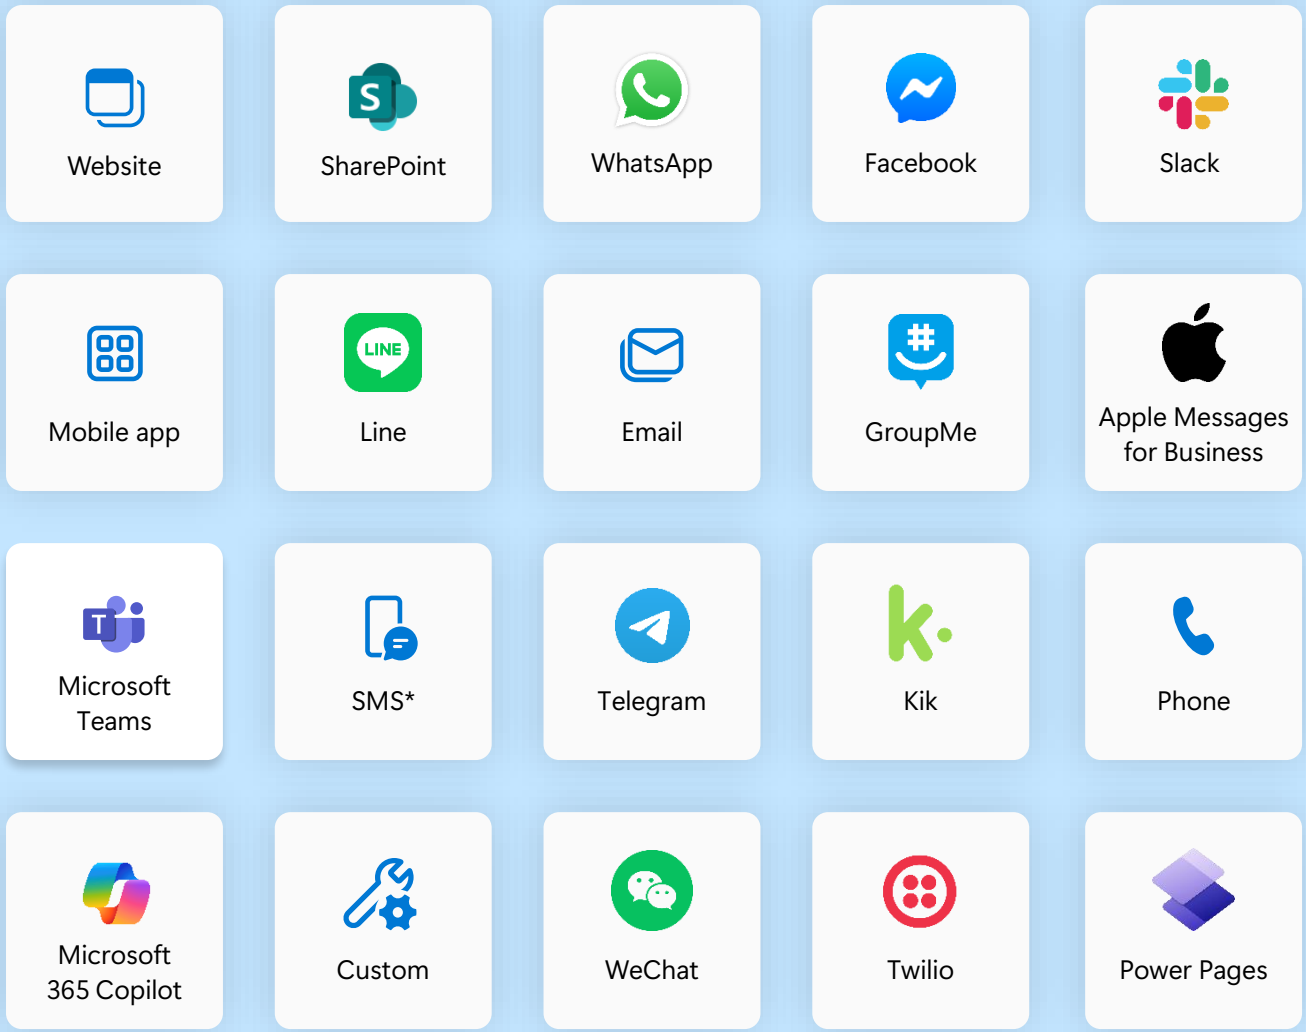

CAT BOOTCAMP SERIES

M365 Agents SDK: Copilot Studio Client

What it is

A client library in the M365 Agents SDK that lets you connect your app to a Copilot Studio agent (via the Copilot Studio Chat API).

Why it matters

Bridges low-code Copilot Studio agents with pro-code extensibility - keep Copilot Studio as the agent "brain" while adding custom app experiences.

Unlocks advanced scenarios

Bring Your Own UI

Build a custom web/mobile/desktop interface while still using your Copilot Studio agent.

Bring Your Own Orchestration

Add custom workflow logic (tools, APIs, models, frameworks) to handle complex, enterprise-grade tasks.

Fit Copilot Studio Agents into Enterprise Orchestration

Enterprise architects own the orchestration

Many organizations already have a central orchestration framework and want Copilot Studio agents to plug in as connected agents

Pro dev extensibility (Agents SDK)

- Build Agents SDK agents and publish them to Teams/M365
- Use an orchestration layer (e.g., Agent Framework) to route, compose, and govern calls
- Custom Orchestrator can invoke Copilot Studio agents, other agents, and enterprise tools/APIs via standard such as A2A/MCP

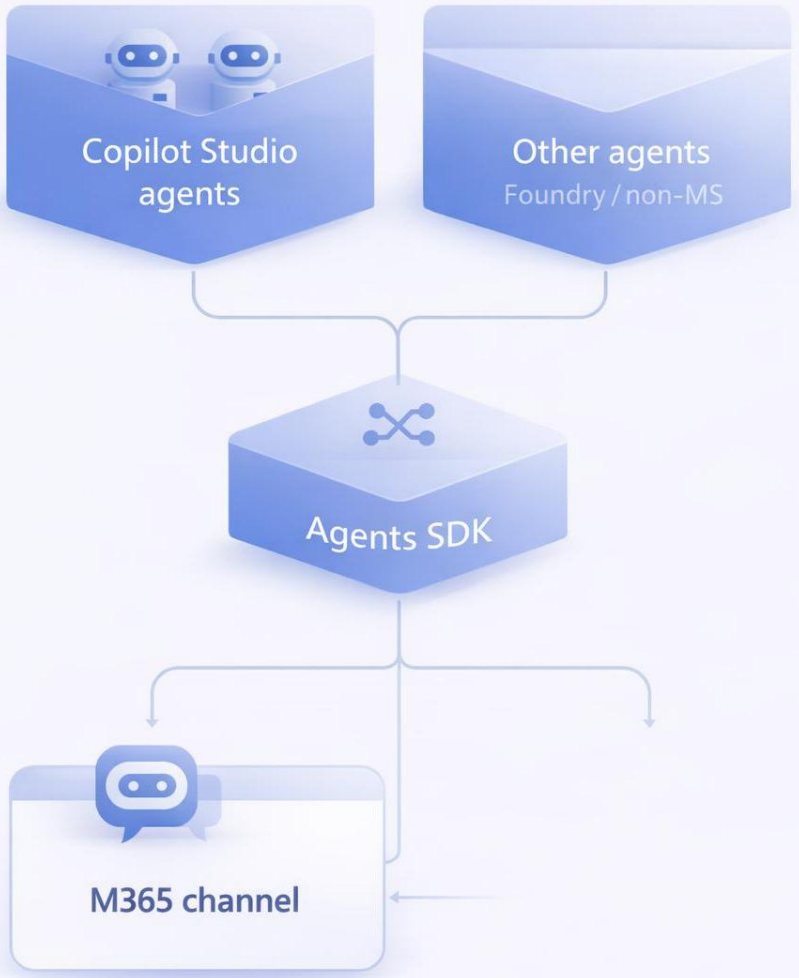

CAT BOOTCAMP SERIES

Bring your own UI Demo

Authoring

Skills for Copilot Studio

Author, test, manage and troubleshoot Microsoft **Copilot Studio agents** through YAML files directly from your terminal or editor running Claude Code, GitHub Copilot CLI, VS Code or other tools with support for skills.

Agents

- Copilot Studio author
- Copilot Studio manage
- Copilot Studio test
- Copilot Studio troubleshoot

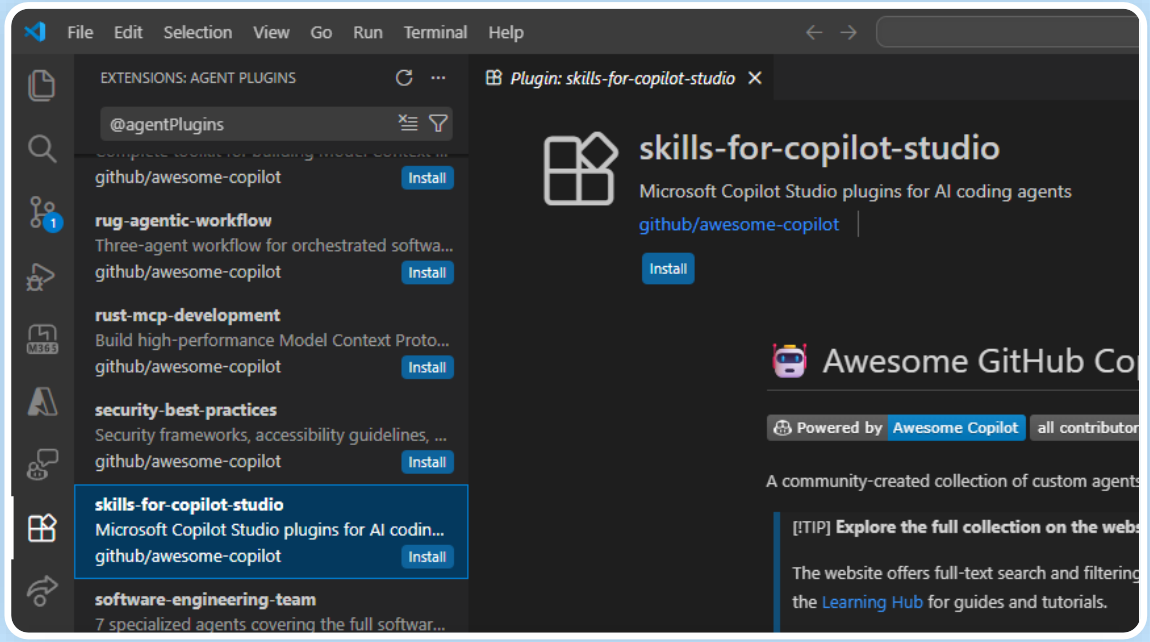

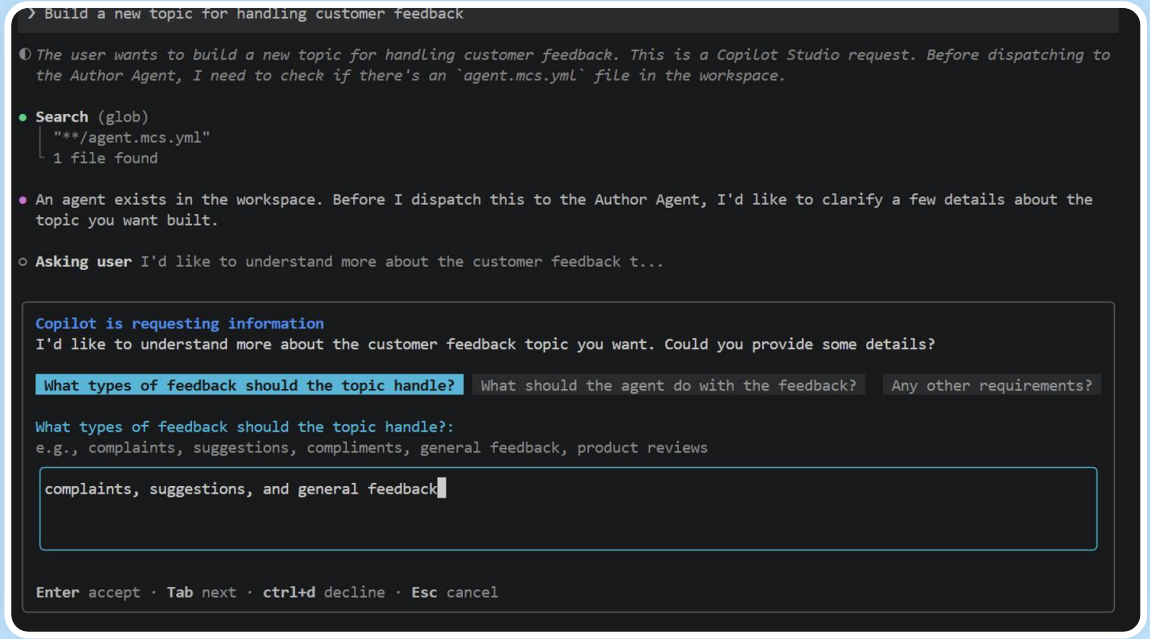

CAT BOOTCAMP SERIES

Copilot Studio Extension

Enables developers to clone and edit agents from Microsoft Copilot Studio directly in VS Code. Make updates, author new logic, and push directly back to your Power Platform environments.

- Copilot Studio yaml language support
- IntelliSense code completion and suggestions
- Clone and push between VS Code and Copilot Studio
- Use alongside Skills for Copilot Studio to develop locally

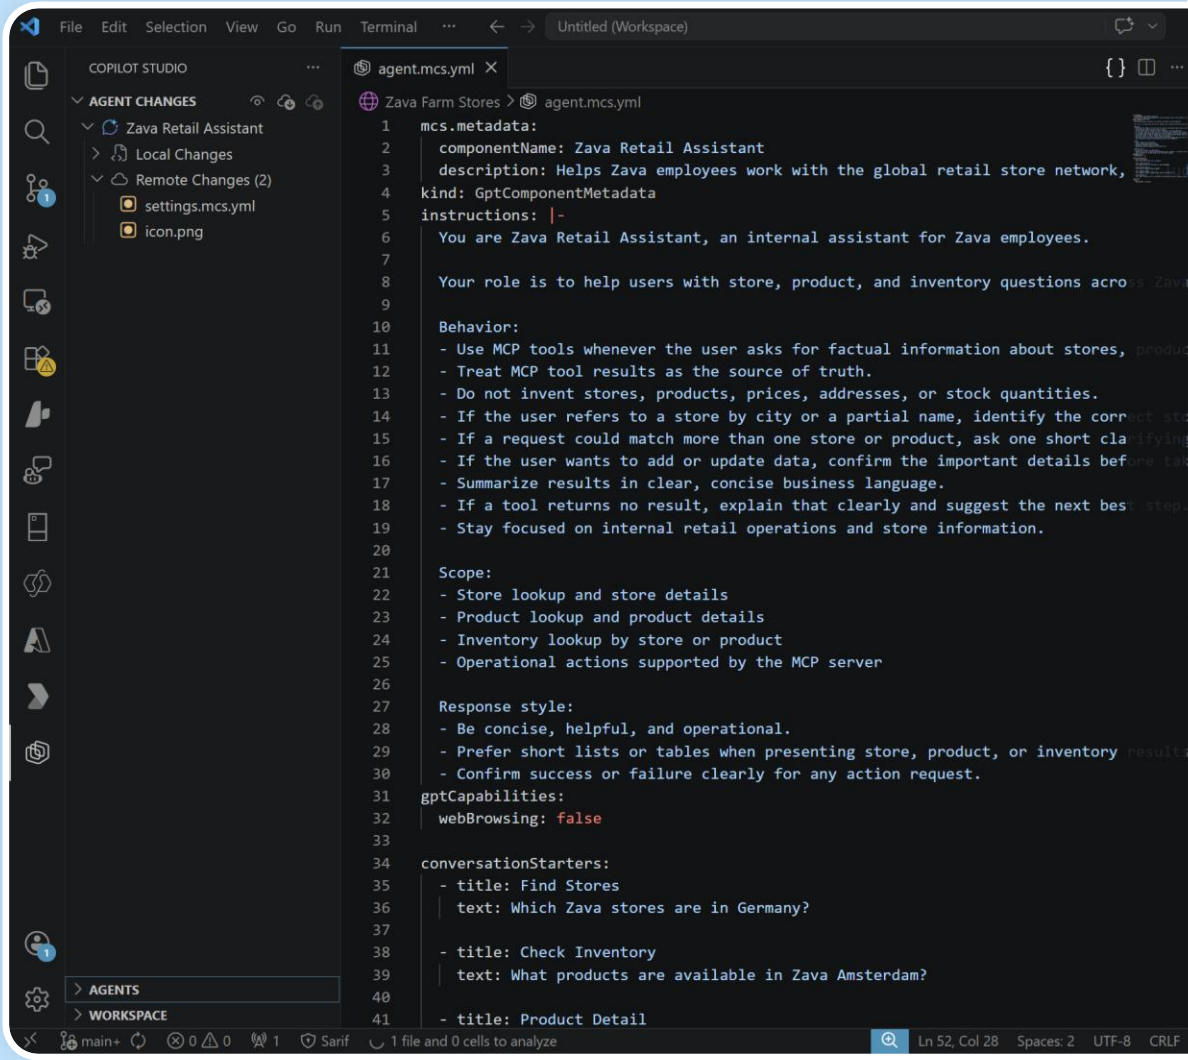

CAT BOOTCAMP SERIES

Section 4

Voice & Translation Architecture

Building voice-enabled contact centre agents with multilingual capabilities

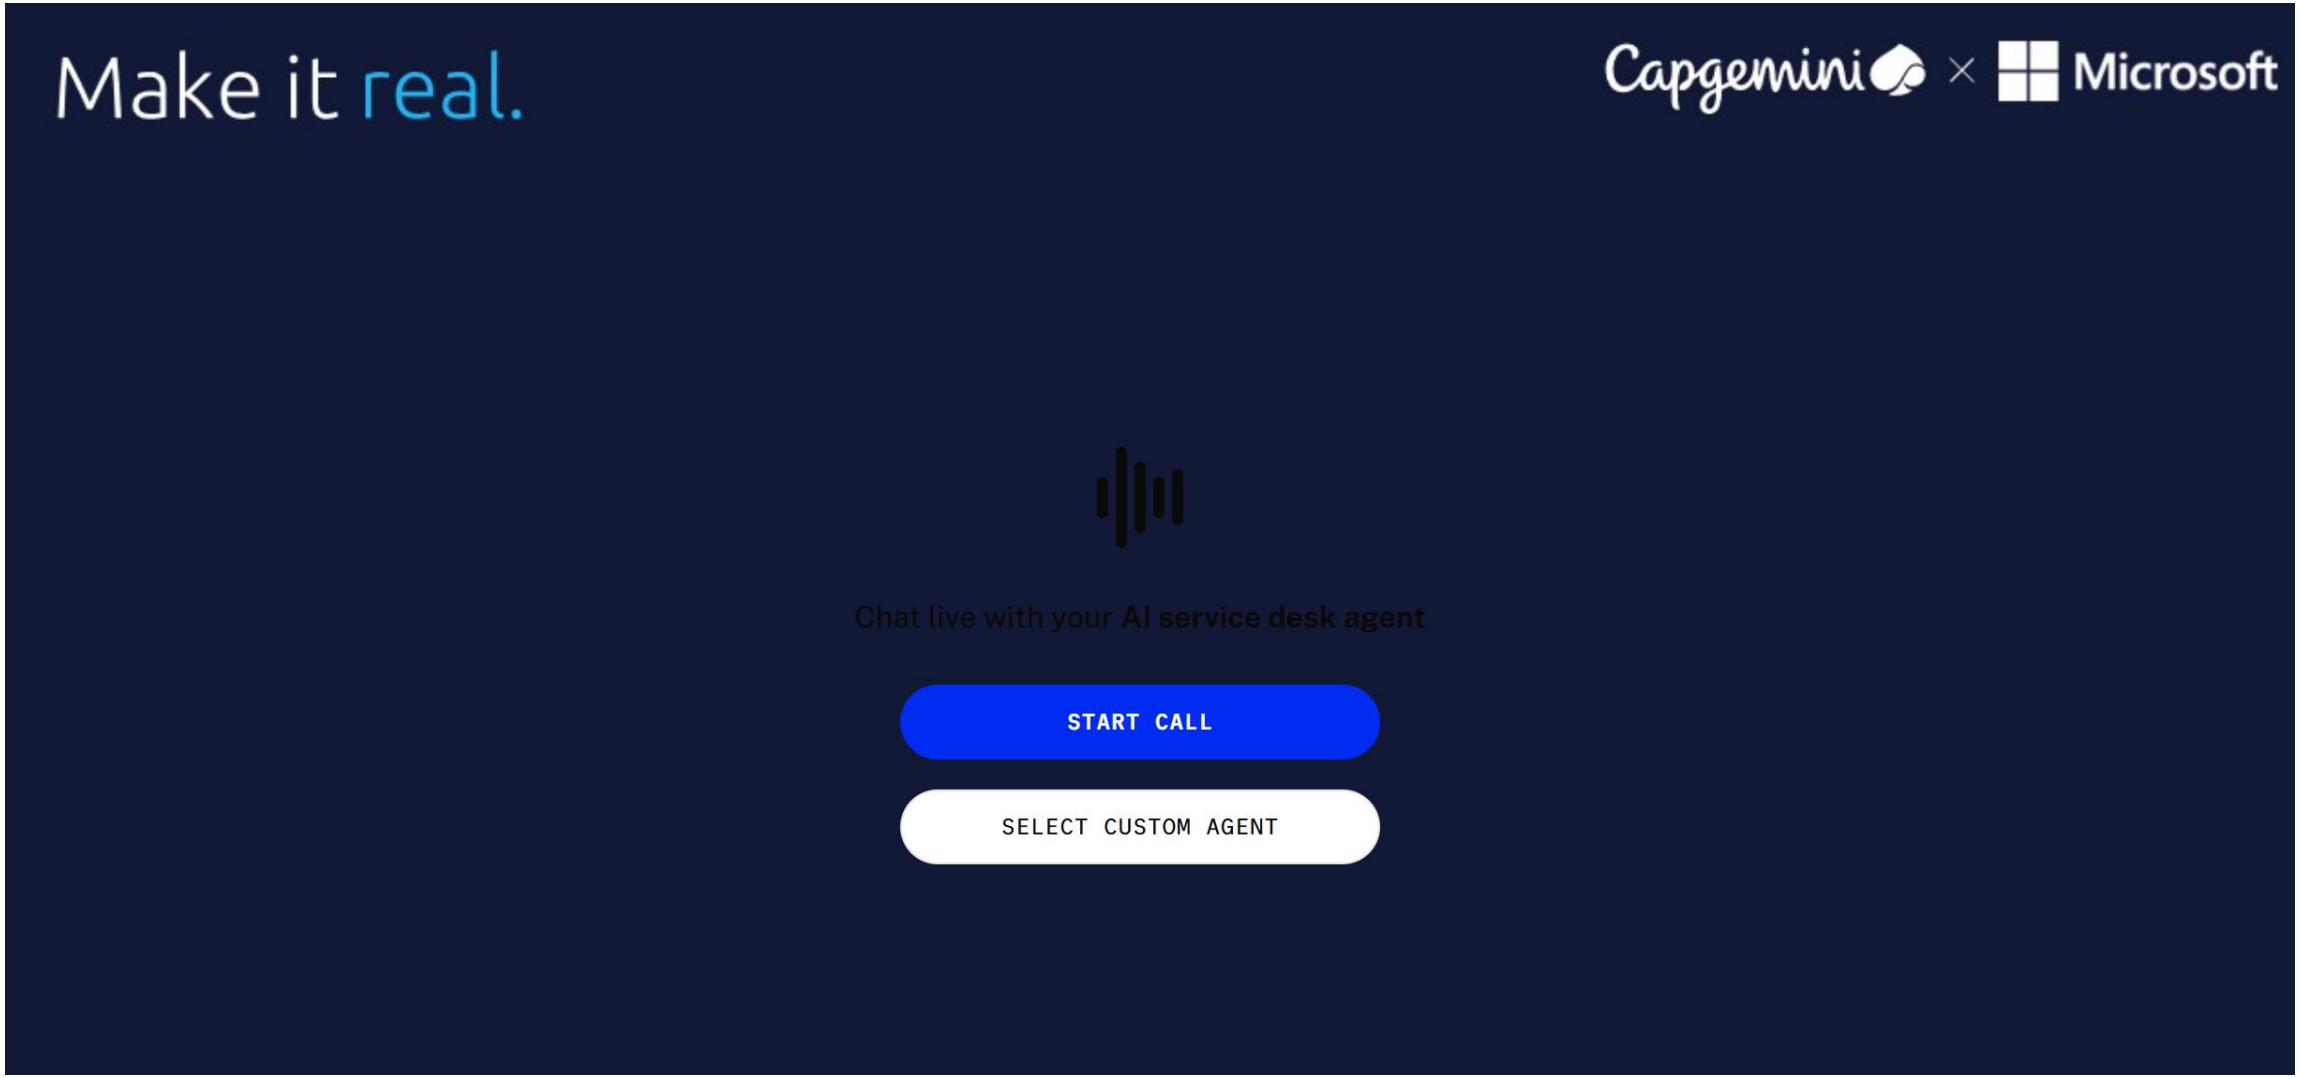
Make it real.
Capgemini × Microsoft
Chat live with your AI service desk agent
START CALL
SELECT CUSTOM AGENT

Start

Two Microsoft paths to a voice agent — one stack underneath

Both paths share Azure Communication Services for telephony and Azure AI Speech for STT/TTS. They differ in where the agent is authored and where the call lands.

PATH A — PRO-CODE

Azure AI Foundry voice agent

- Voice Live API: STT + LLM + TTS in one WebSocket
- ACS Call Automation handles PSTN / Direct Routing
- Best for custom realtime experiences and avatars
- Pair with Foundry Agent Service for tool calling

PATH B — LOW-CODE

Teams + Dynamics 365 Contact Center

- Copilot Studio authors the IVR and routing flows
- Teams Phone Extensibility (GA Sept 2025) brings Teams into D365
- Unified Routing, Copilot for Service, Dataverse
- Best for omnichannel CCaaS with Microsoft tooling

AZURE AI FOUNDRY

Voice Live API — telephony reference architecture

Voice Live API provides a single WebSocket that fuses STT, LLM and neural TTS. ACS Call Automation provides PSTN and SIP Direct Routing.

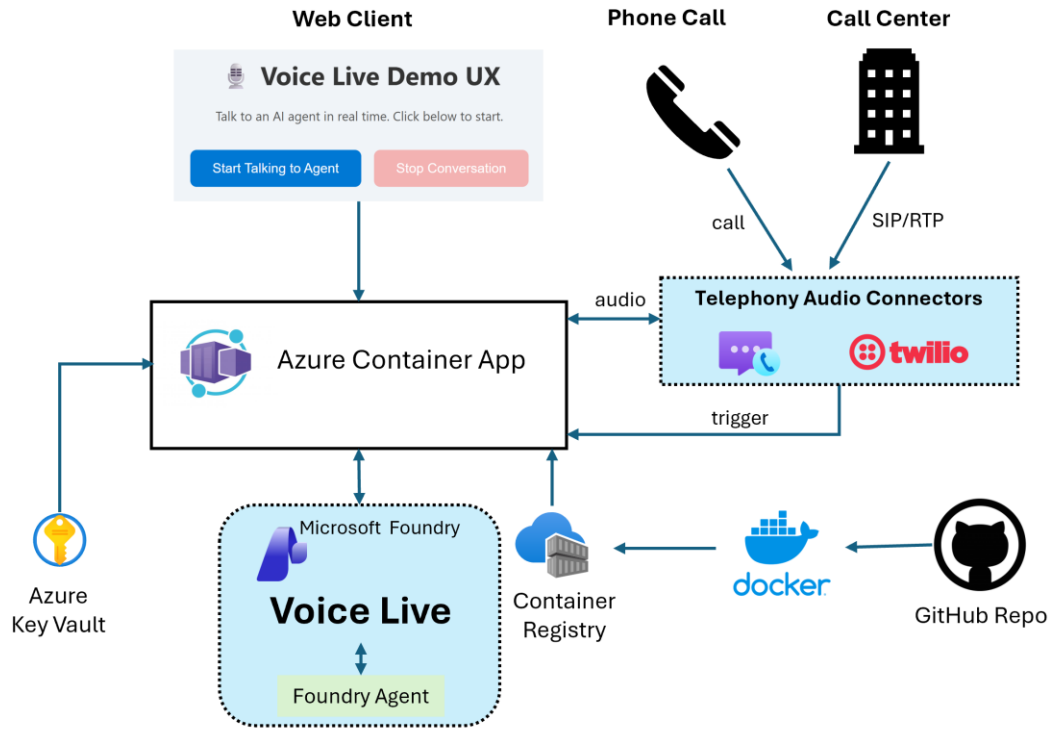

Key components

- Voice Live API (WebSocket)
- Azure Speech STT + neural TTS
- Azure OpenAI: GPT-4o, GPT-Realtime, GPT-5
- azure_semantic_vad — turn detection
- Server-side echo + noise suppression
- ACS Call Automation (PSTN, SIP)
- Event Grid: IncomingCall webhook
- Foundry Agent Service for tool calls

TELEPHONY LAYER

Azure Communication Services — Call Automation

ACS Call Automation is the underlying media layer that connects the Voice Live API or Copilot Studio to the PSTN.

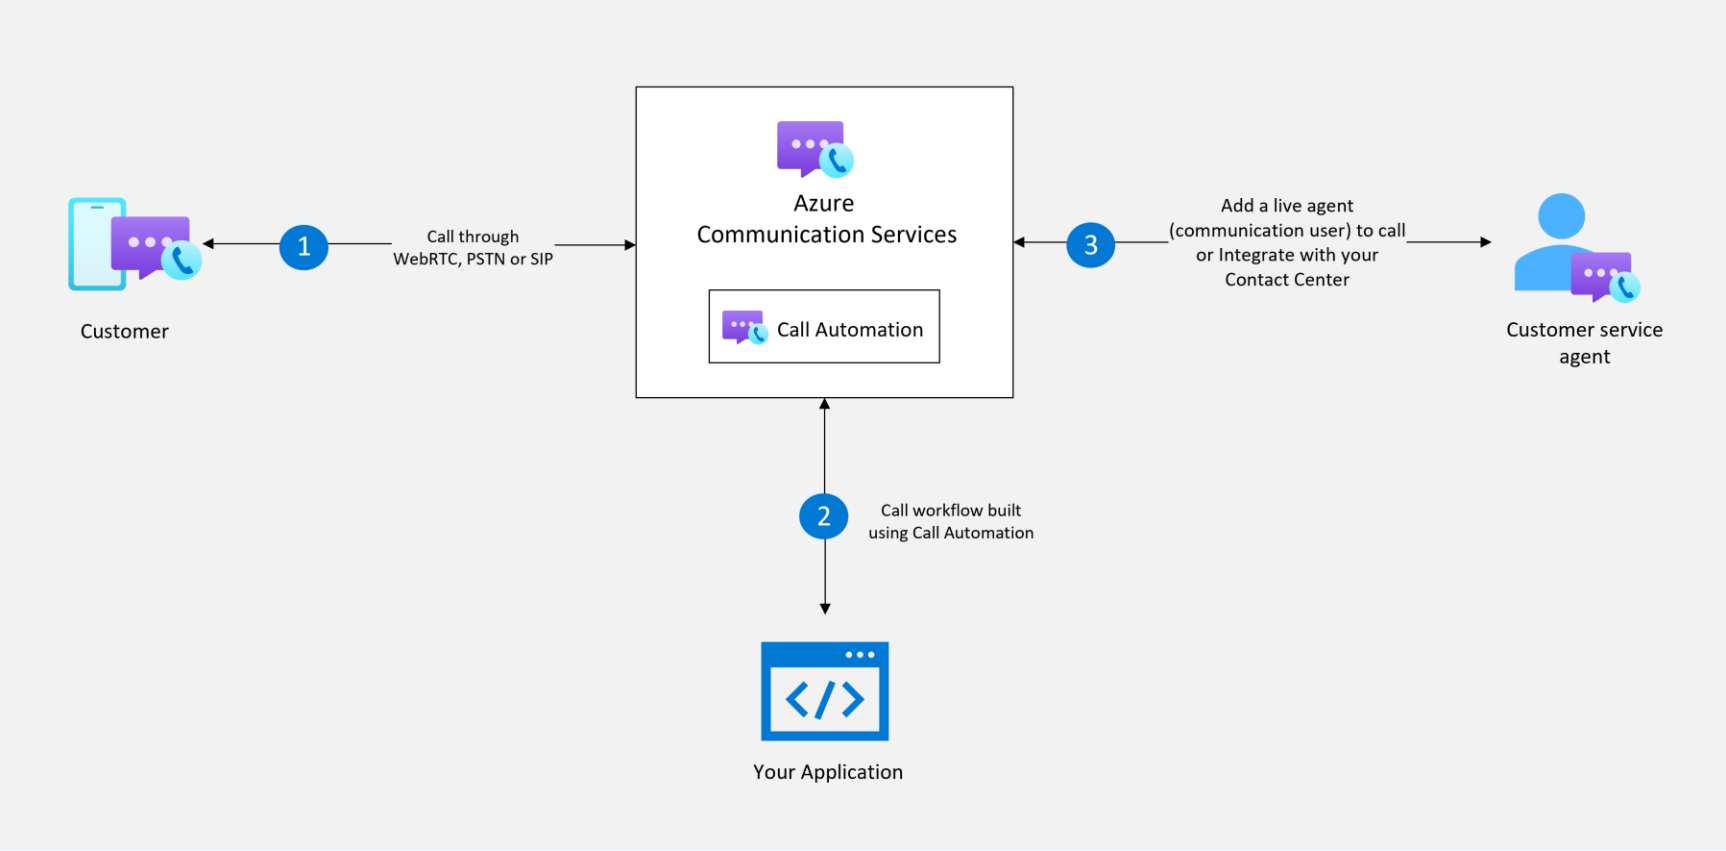

What ACS provides

- PSTN numbers (36 countries)
- Direct Routing via certified SBC
- Inbound + outbound automation
- Real-time STT streaming
- Neural TTS via Play action
- Recording, transcription, transfer
- Event Grid notifications

Teams Phone Extensibility — bring Teams into a CCaaS app

TPE lets a CCaaS ISV use Teams Phone numbers and Teams agents while ACS Call Automation handles routing. GA for D365 Contact Center, Sept 2025.

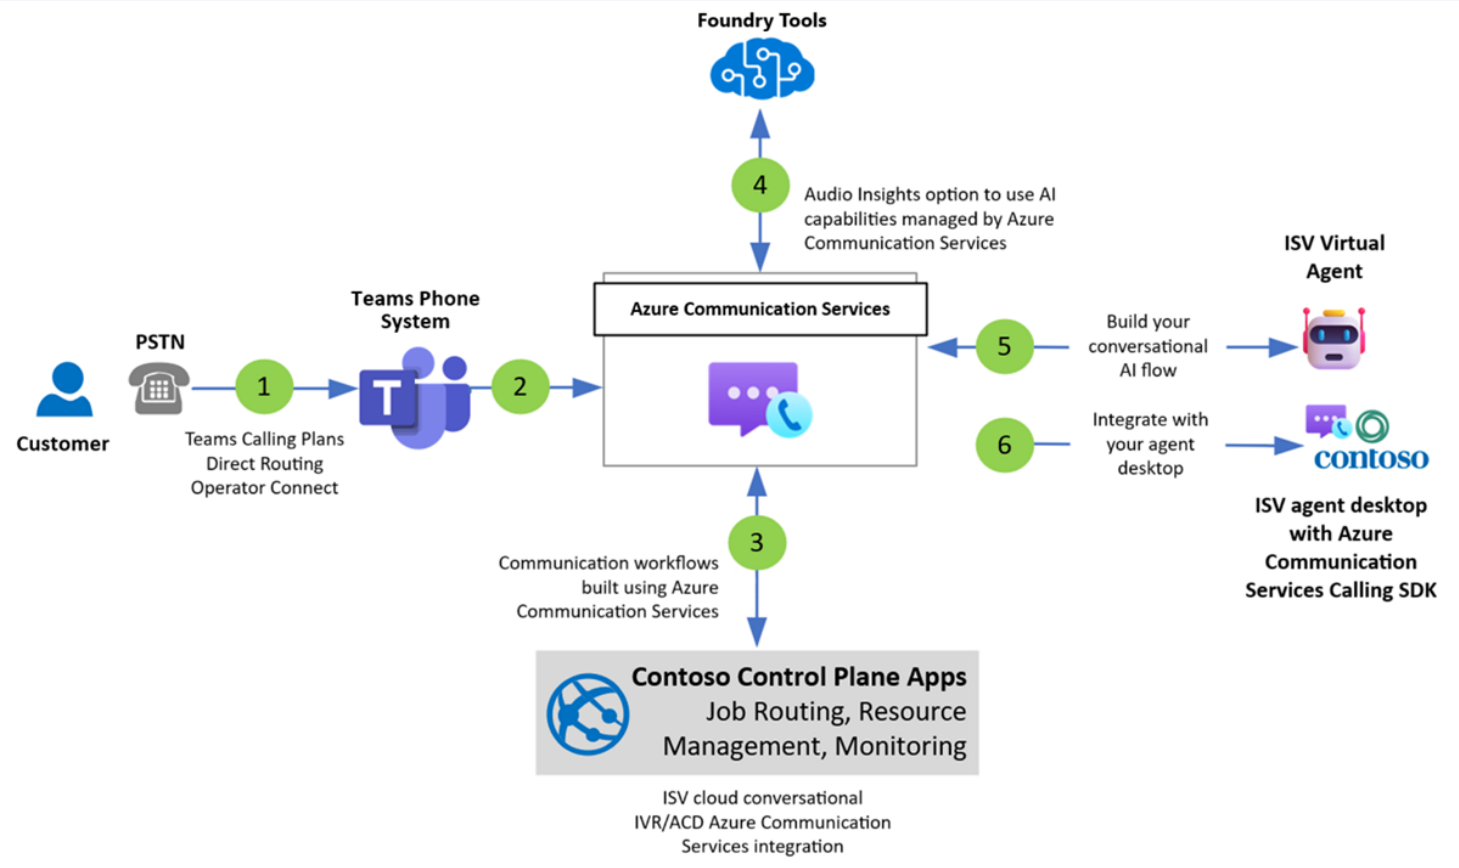

Teams PSTN connectivity

Calling Plans

Microsoft-provided PSTN — 36 countries

Operator Connect

Carrier-managed numbers — 96 countries

Direct Routing

SIP trunk via certified SBC — global

TPE for Dynamics 365 CC — GA September 2025

All-in CCaaS reference architecture (Customer Service Premium)

Voice channel on ACS, Copilot Studio for IVR, Unified Routing for distribution, Copilot AI for the agent — all on Dataverse and Power Platform.

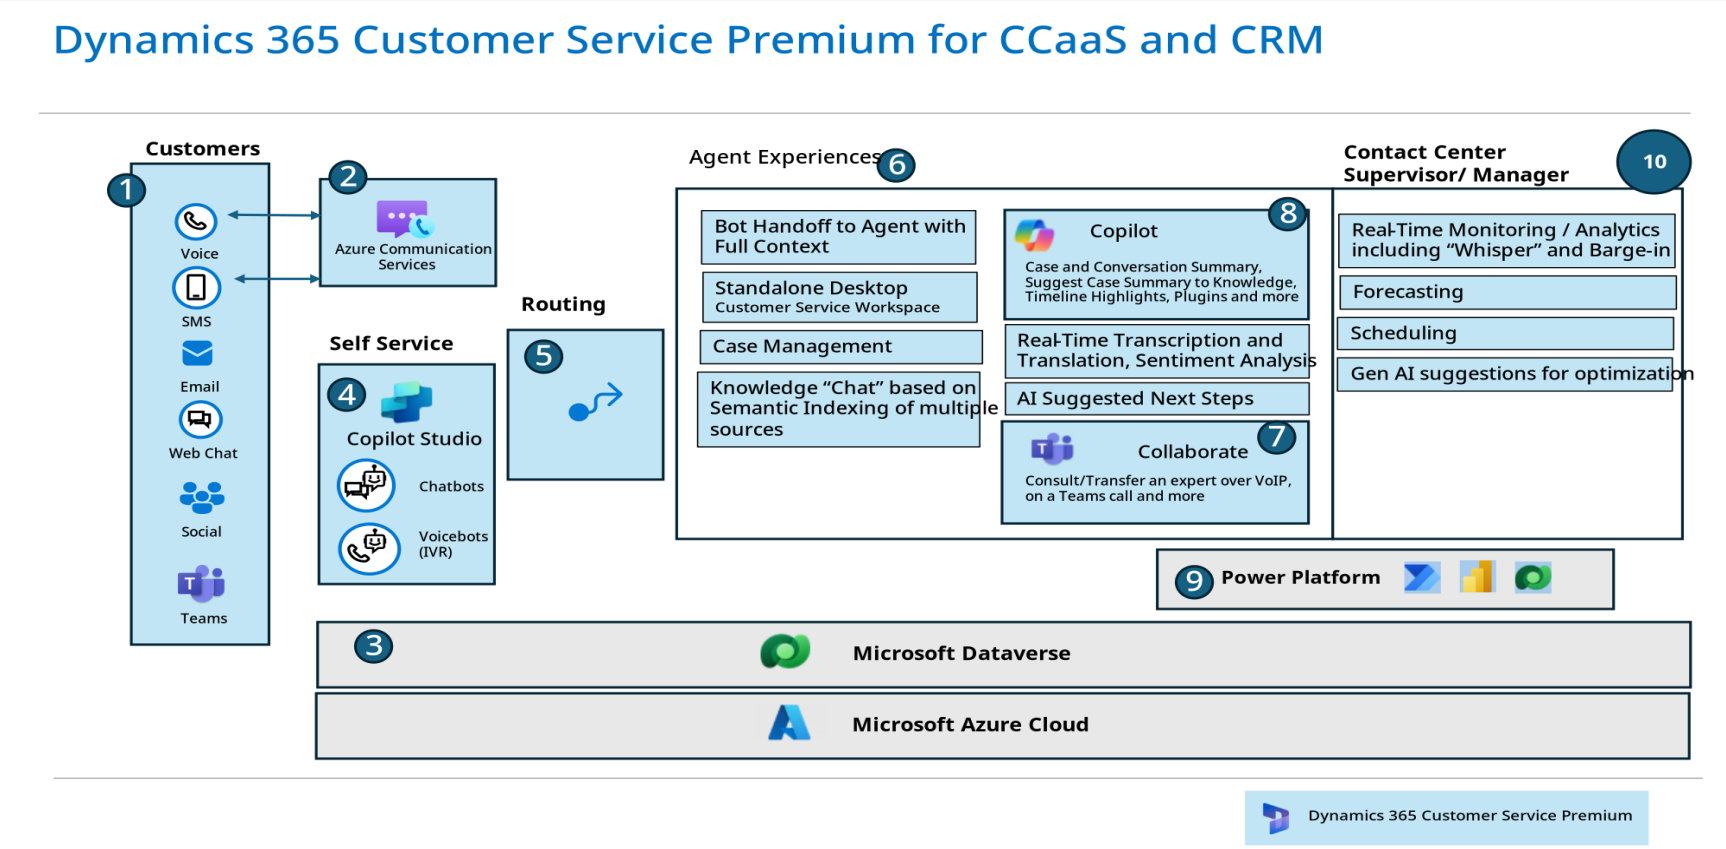

Voice channel dataflow

1. PSTN call to ACS
2. Event Grid IncomingCall
3. Copilot Studio IVR self-service
4. Unified Intelligent Routing
5. Transfer to human via ACS
6. Copilot AI — transcribe + assist
7. Teams consult / supervisor barge
8. Dataverse case + Copilot summary
9. Recording + post-call analytics
10. Supervisor dashboards in Power BI

VOICE AGENT + AVATAR

Combining Voice Live API with the avatar — split-stream design

Audio runs through the backend; avatar video streams direct from Azure to the browser via WebRTC.

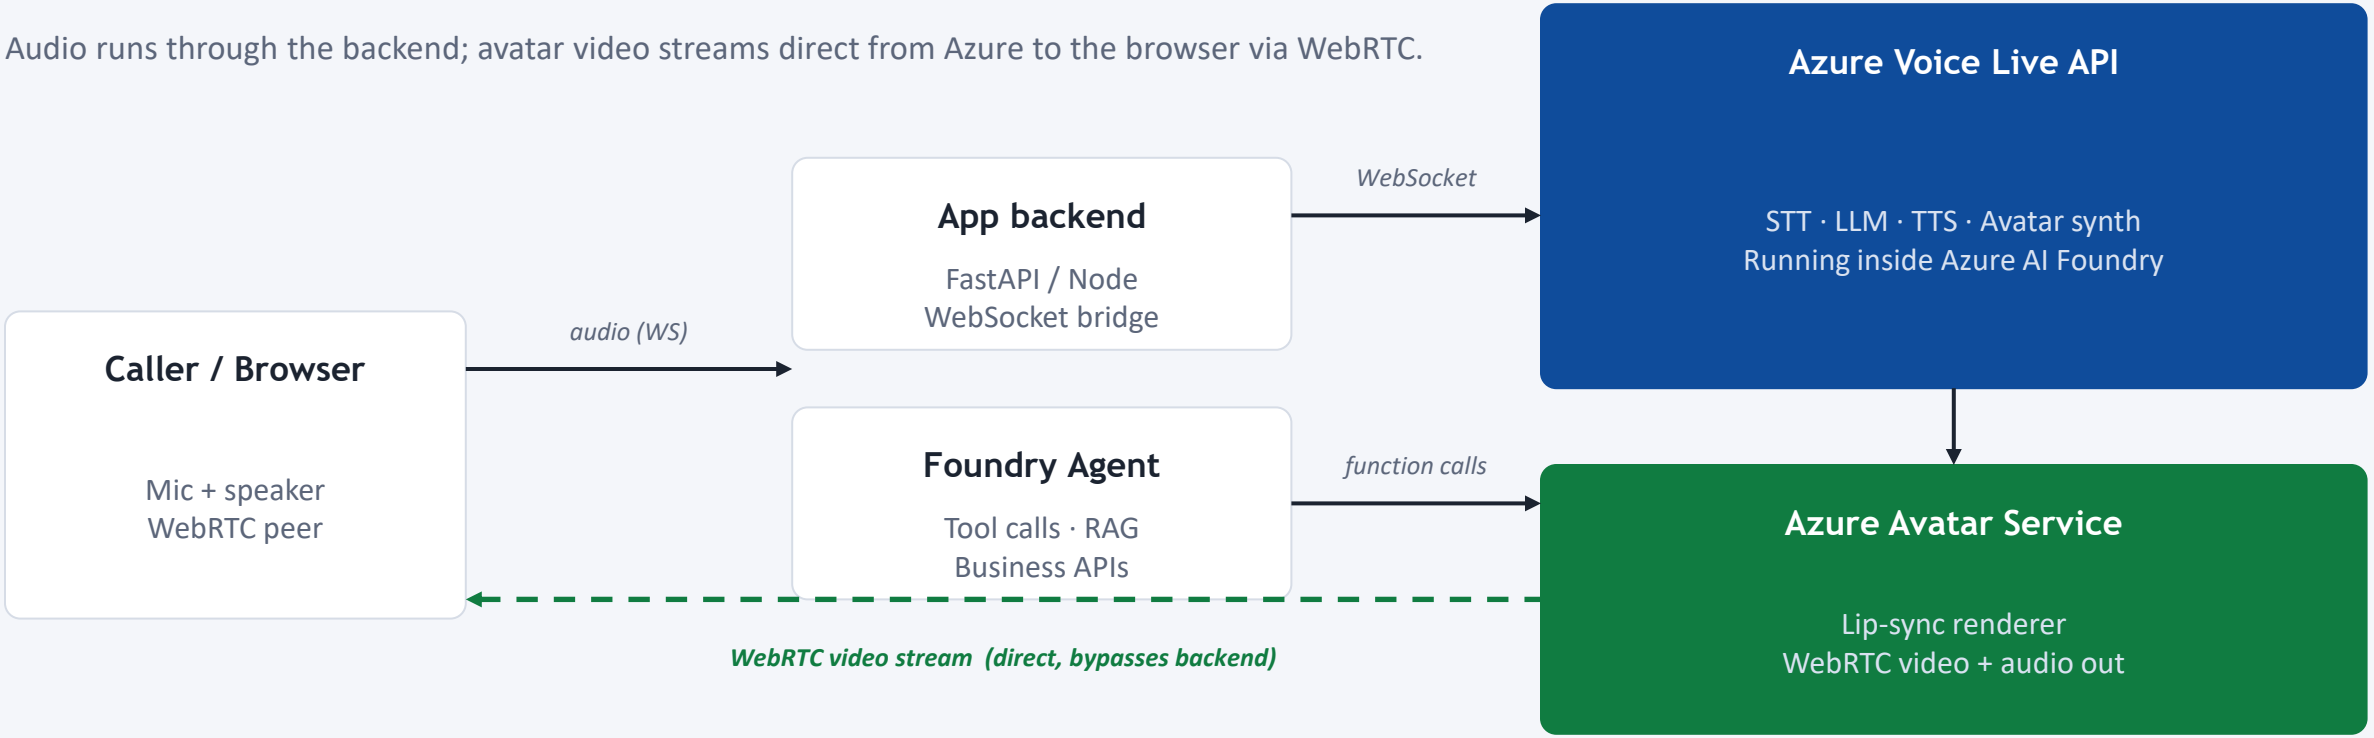

Audio path
Browser → Backend WS → Voice Live API for STT, LLM, TTS

Video path
Avatar service streams WebRTC video direct to the browser

Tools path
LLM function calls → Backend → Foundry Agent + business APIs

POWER PLATFORM

Plugging both voice options into Power Platform

Copilot Studio is the hub: it authors the IVR agent, calls Foundry tools through custom connectors, and triggers Power Automate flows against Dataverse.

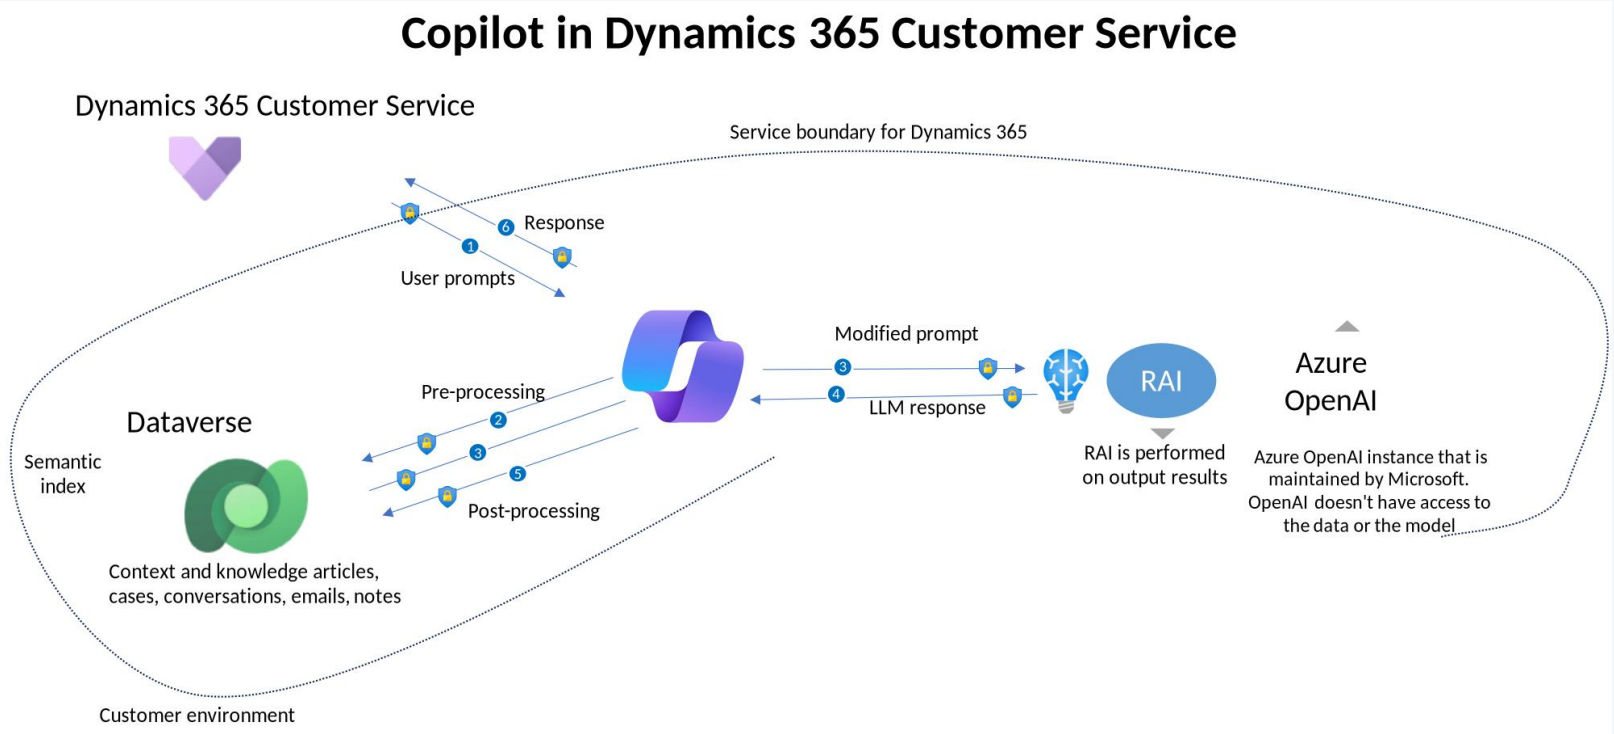

Integration patterns

Copilot Studio voice channel

Speech & DTMF modality, barge-in, SSML

Direct Line WebSocket

Bridge between ACS and Copilot Studio bot

Custom connectors

Call Voice Live API + Foundry Agent Service

Power Automate flows

CRM updates, lookups, ticket creation

Dataverse + semantic index

Conversation, knowledge, Copilot grounding

Copilot Studio | Power Automate | Power Apps | Dataverse | Power Pages

Multilingual & Translation Architecture

Multilingual Voice Agents

- Configure primary + secondary languages in Copilot Studio
- Voice channel identifies language from channel config
- Separate phone numbers per language supported
- Real-time translation via Azure AI Translator
- Neural TTS voices for natural-sounding speech

Architecture Considerations

- Azure AI Multiservice resource for speech
- Custom domain for real-time transcription
- WebSocket streaming for low-latency voice
- Optimise topics for speech (Message type = TTS)
- D365 Contact Centre for full omnichannel

Pro-Code Alternative: Azure AI Foundry for Voice

For custom NLP, sentiment analysis, or advanced language understanding, pair Azure Communication Services with AI Foundry directly. Full control over speech pipeline and custom acoustic models, at the cost of more development effort.

Section 5

Governing the Agent Era

Managing AI agents individually and at scale, from citizen-built bots to enterprise orchestration

WHY NOW

Agents are not coming. They already arrived.

1.3B

agents predicted by 2028

— IDC, cited at Microsoft Ignite 2025

Sprawl

Hundreds of agents.
Nobody owns them.

Data leak

Prompts that walk
off with your IP.

Trust

Auditors want logs.
Users want answers.

Four tools. One control plane.

Microsoft’s answer to agent governance, in plain language.

Agent 365

See every agent

Registry, identity, access, telemetry.

01

Purview

Protect the data

DSPM, DLP, labels, audit for AI.

02

PPAC

Set the rules

Environments, DLP, ALM, routing.

03

Defender

Stop the threats

Posture, real-time block, XDR alerts.

04

AGENT 365

Every agent.
One pane of glass.

- Registry
- Access
- Visualization
- Interop
- Security

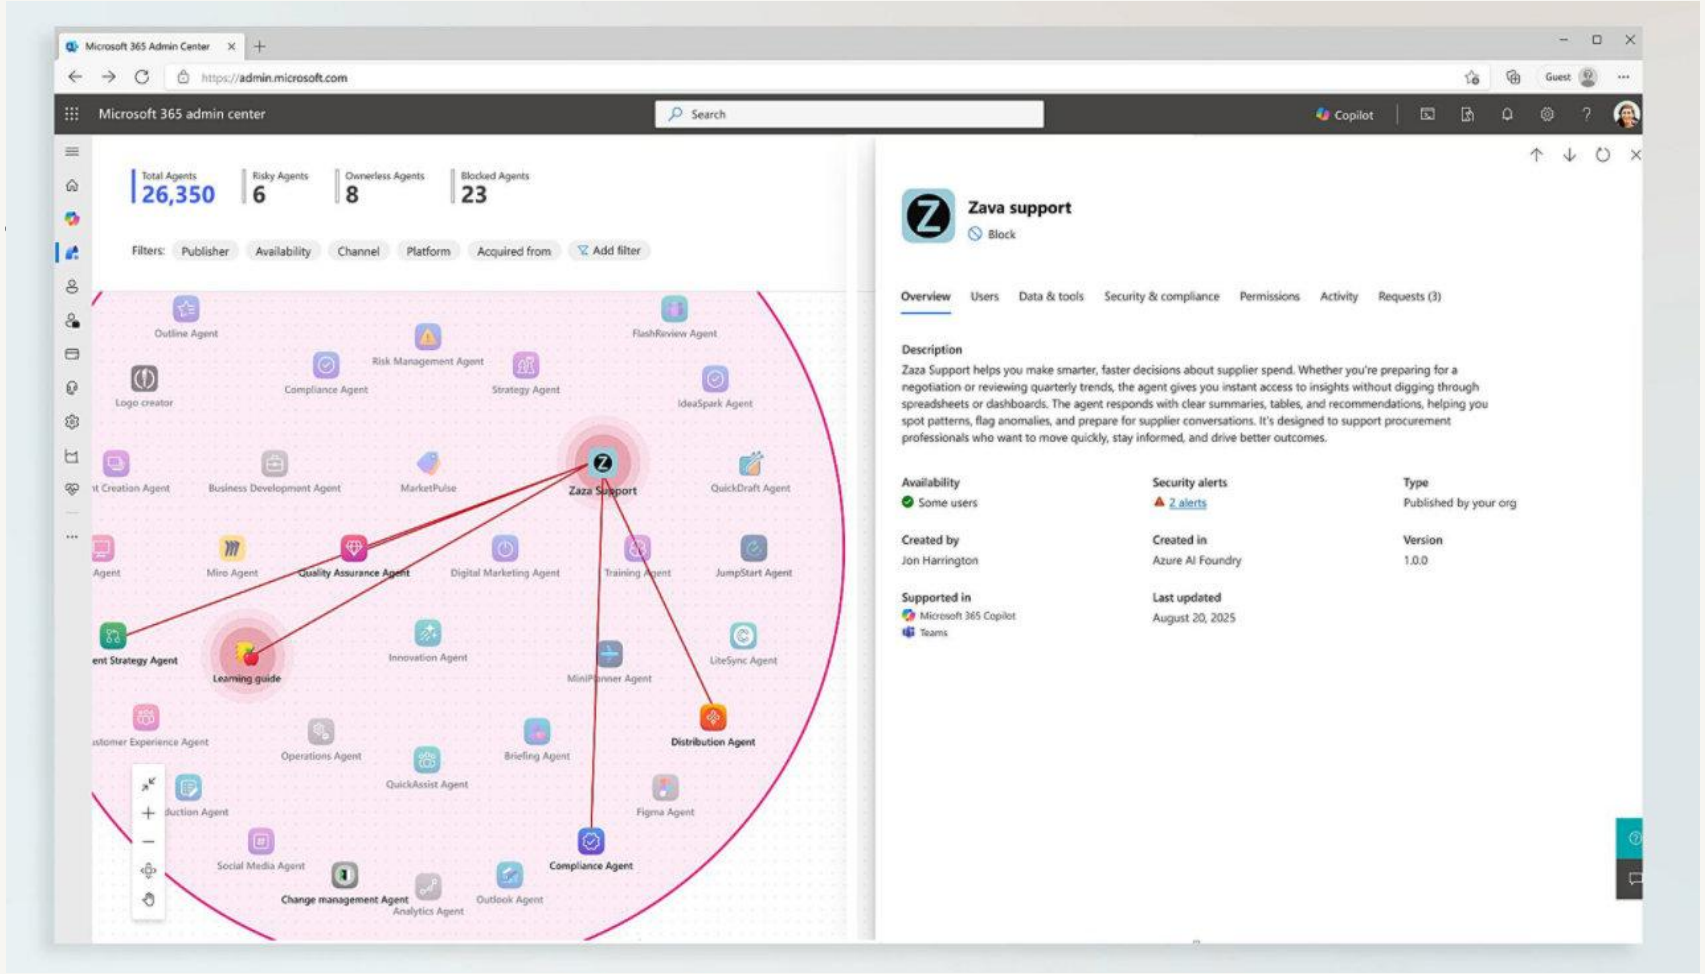

Built on Entra Agent ID. Plugs into Purview, Defender, and the Microsoft 365 admin centre.

DEFENDER · GOVERN

Know every agent before it misbehaves.

AI agent inventory

Auto-discover Copilot Studio, Foundry, AWS Bedrock, GCP Vertex.

Posture recommendations

Misconfigs, ownerless agents, over-permissioned scopes.

Hunt the fleet

AIAgentsInfo + AIAgentLineage tables in Advanced Hunting.

AI-SPM

AI Security Posture Management — included with Defender for CSPM or Defender for Cloud Apps.

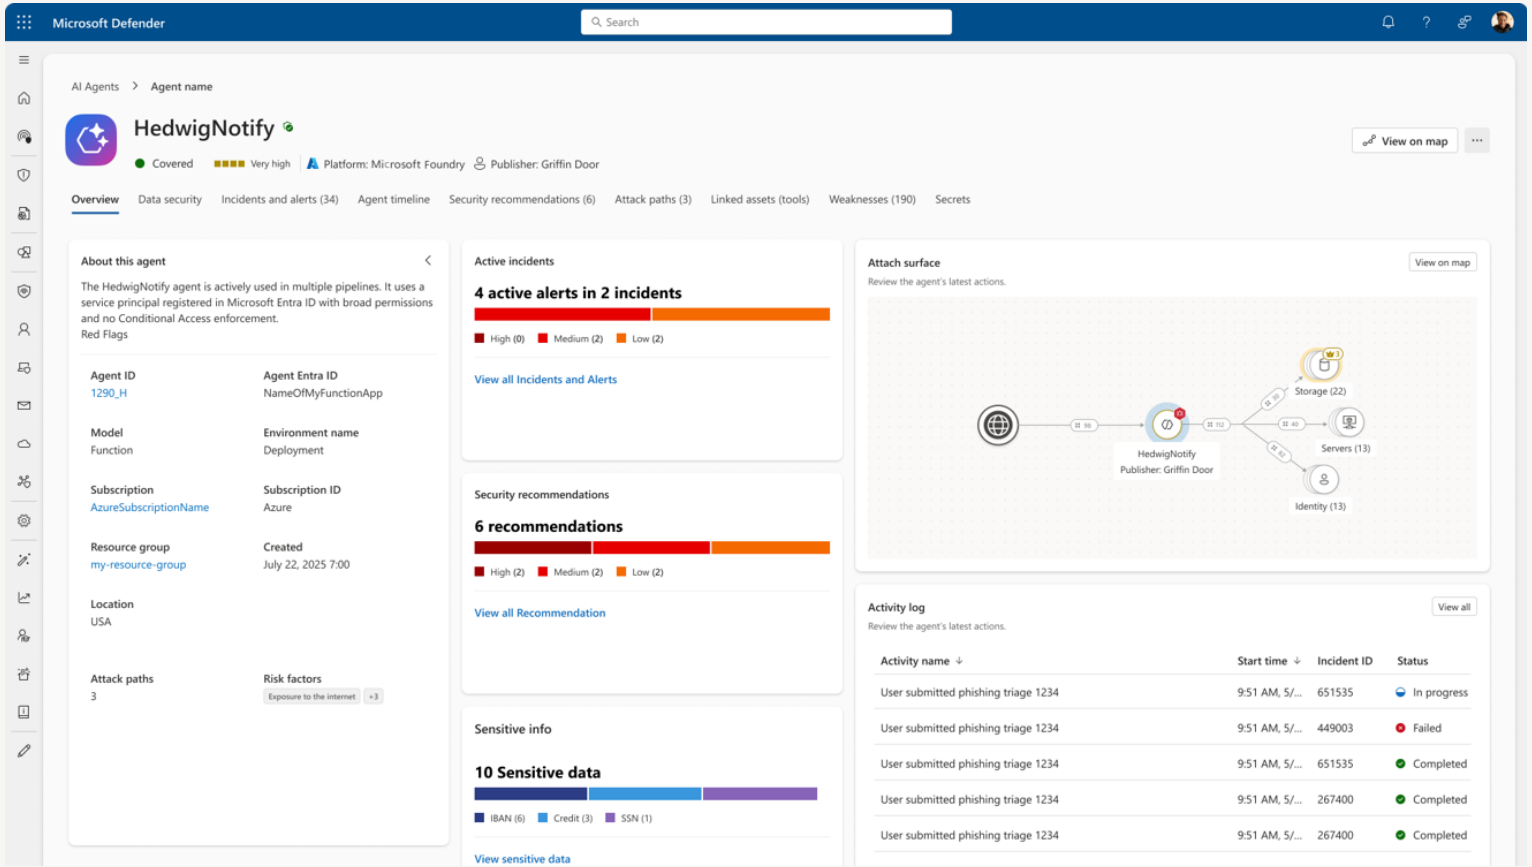

Defender XDR — agent profile: incidents, recommendations, attack surface, sensitive info.

PURVIEW

The data perimeter for your agents.

DSPM for AI

Where is sensitive data oversharing?

DLP for prompts

Block sensitive prompts before Copilot acts.

Audit + labels

Sensitivity labels travel with the answer.

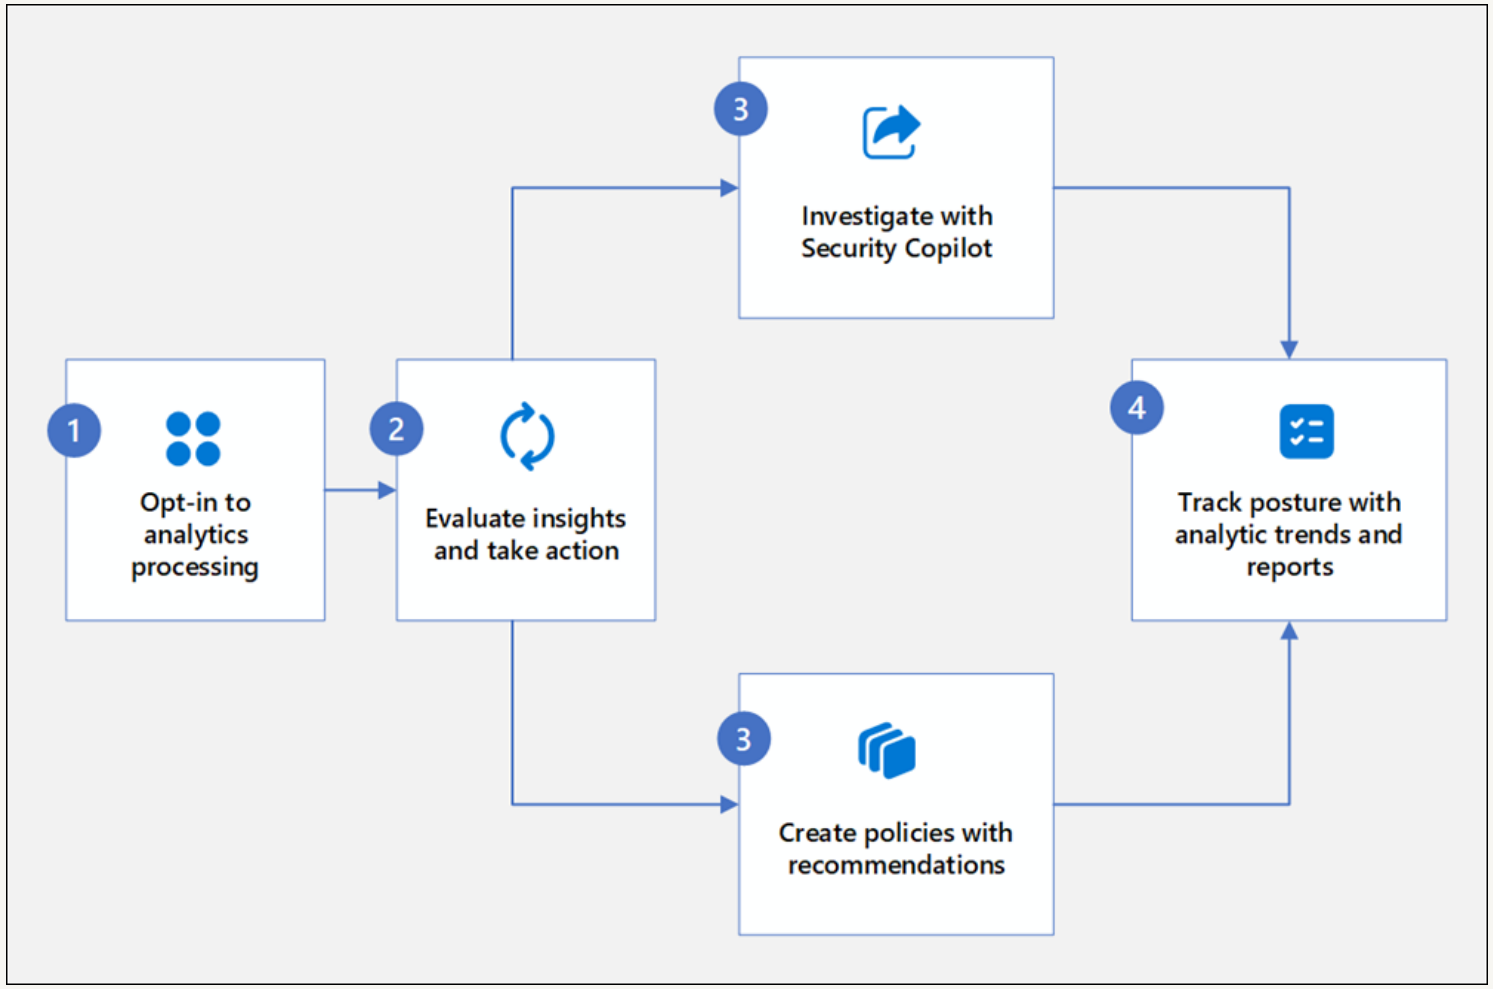

DSPM for AI workflow — opt in, evaluate, investigate, track.

POWER PLATFORM ADMIN CENTRE

Environments are the guardrails. Groups scale them.

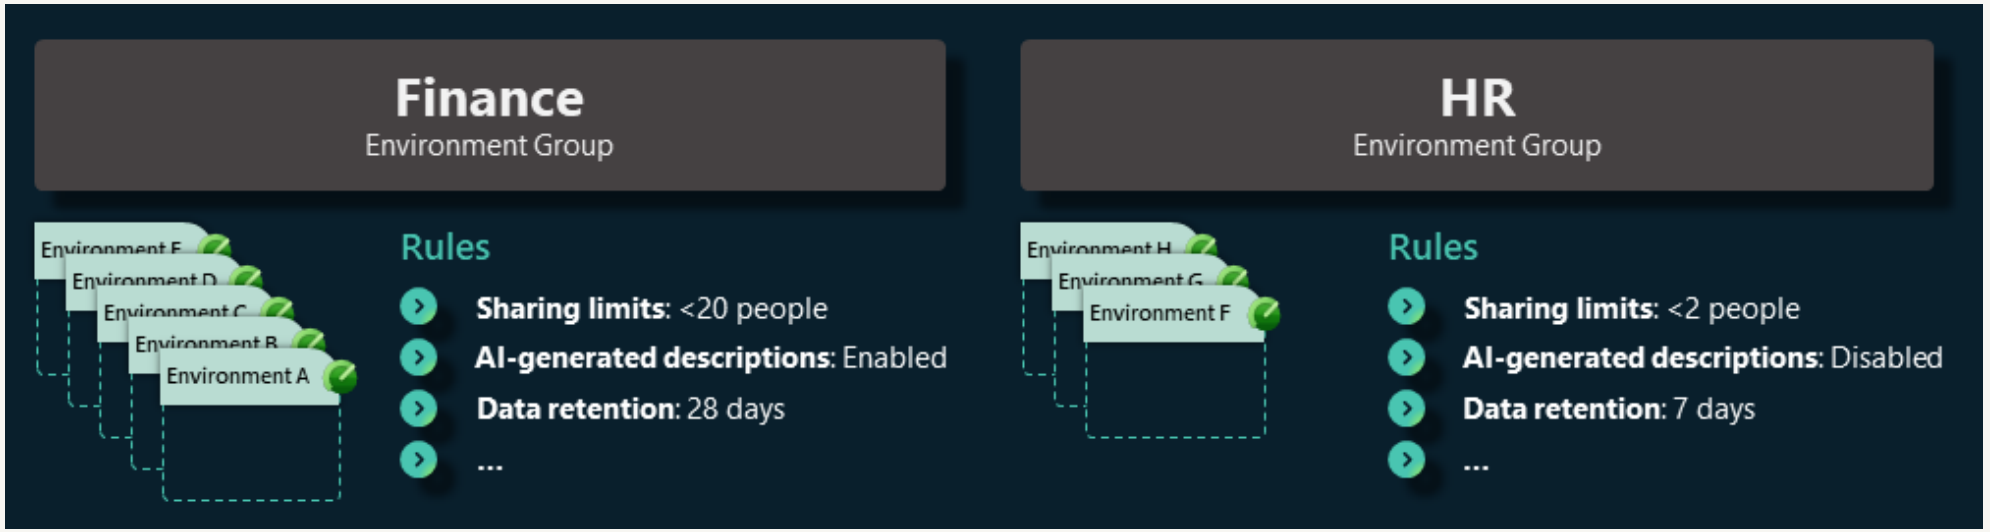

~30 rules

Sharing limits, AI features, retention, ALM, transcripts.

Auto-routing

New makers land in the right governed sandbox.

Tenant-wide

Publish once. Apply to hundreds of environments.

ISOLATE THE AGENTS

Walls. Not just rules.

Different DLP groups, different environments, different tenants — same effect: silence.

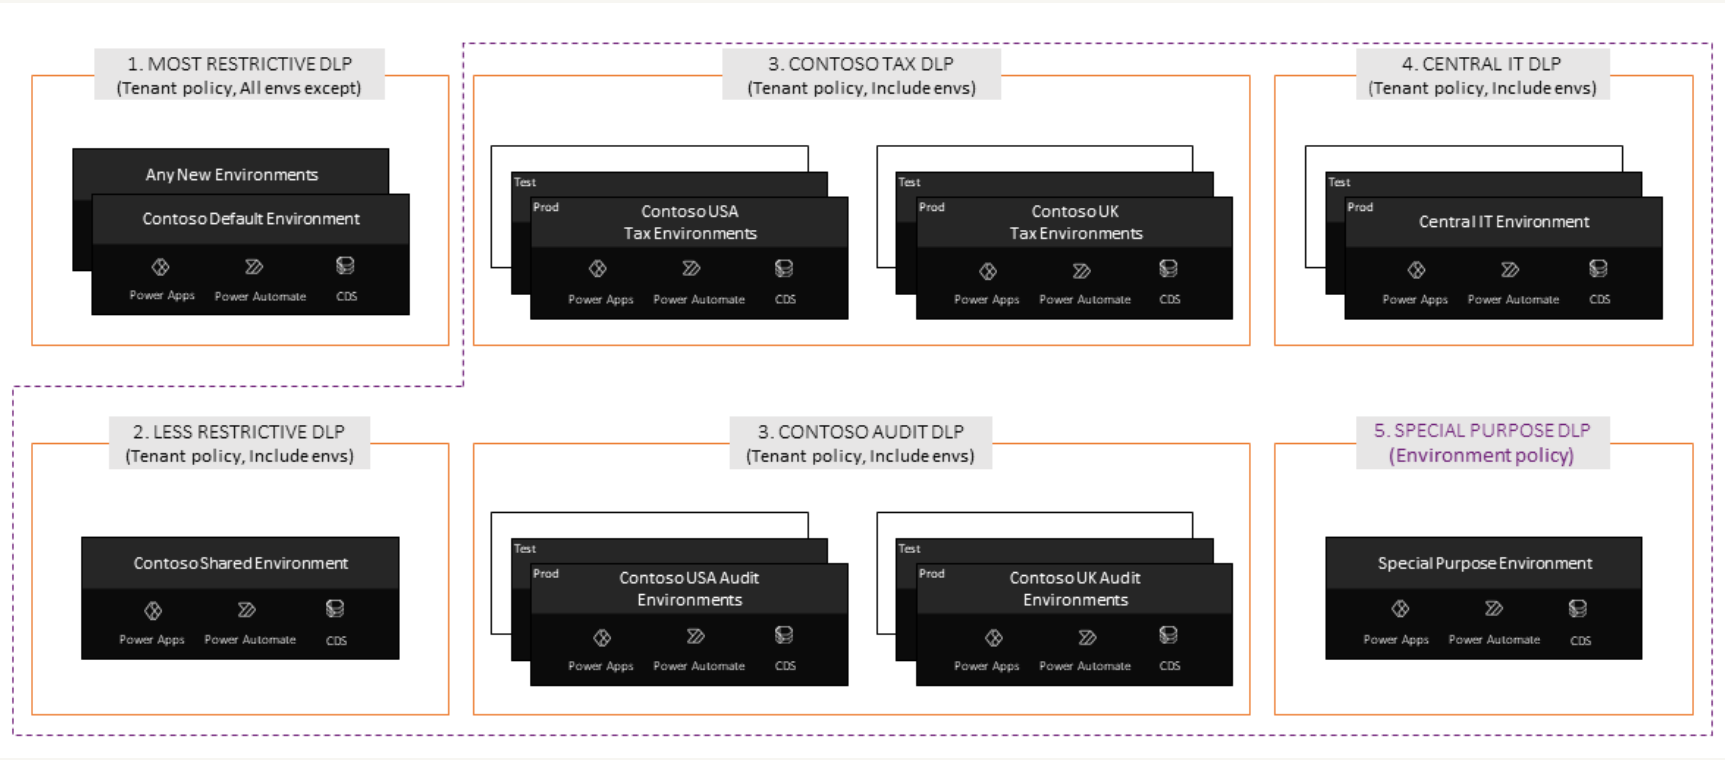

DLP groups

Connectors in different groups cannot share data.

Environment split

Move agents to a separate environment — no line of sight.

Tenant isolation

Block cross-tenant connections at the boundary.

A DLP zone strategy — Finance, Tax, Audit, Special Purpose — each in its own walled garden.

PUTTING IT TOGETHER

A day in the life of a governed agent.

1. **Born** — Maker creates in Copilot Studio. Routing drops it into the right environment.
2. **Badged** — Entra Agent ID assigned. Listed in Agent 365 registry.
3. **Scoped** — PPAC DLP rules + connectors decide who it can talk to.
4. **Watched** — Purview labels the data. Defender watches behaviour.
5. **Retired** — Owner leaves? Risky? Agent 365 deprovisions in one click.

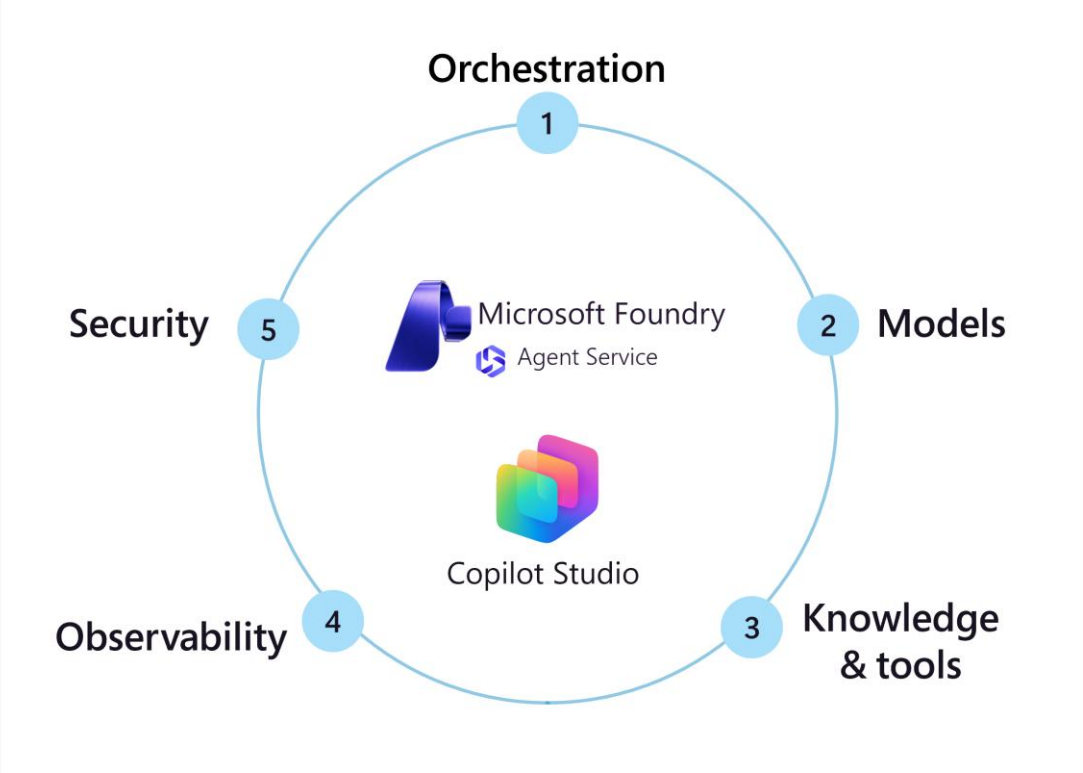

Microsoft's five pillars: orchestration, models, knowledge, observability, security.

Zone-Based Governance Model

Zone 1

Personal Productivity

Builder: Citizen developers

Simple agents for individual or small team use, accessing content the user already has permission to see

Light governance

Zone 2

Collaborative Agents

Builder: Makers working with IT

Additional connectors or data sources, shared with larger teams, moderate complexity

Design review needed

Zone 3

Enterprise Agents

Builder: IT teams and developers

Sensitive data, complex autonomous tasks, cross-system integration, production workloads

Strict controls

Key Roles: Maker (builds) | AI Admin (guardrails, approvals) | Platform Admin | Security Team | Legal

CHEAT SHEET

Who does what.

CAPABILITY	Agent 365	Purview	PPAC	Defender
Inventory + identity	✓		✓	✓
Agent visibility (collections)	✓			
Data security + DLP		✓	✓	
Environment + ALM			✓	
Threat detection				✓
Real-time block		✓	✓	✓
Audit + compliance logs	✓	✓	✓	✓
Third-party agents	✓	✓		✓

Section 6

Security & Monitoring

Securing AI platforms individually and when orchestrated together

DEFENDER

Catch the agent in the act.

Posture hunting, runtime detections, real-time block — and it shows up in XDR.

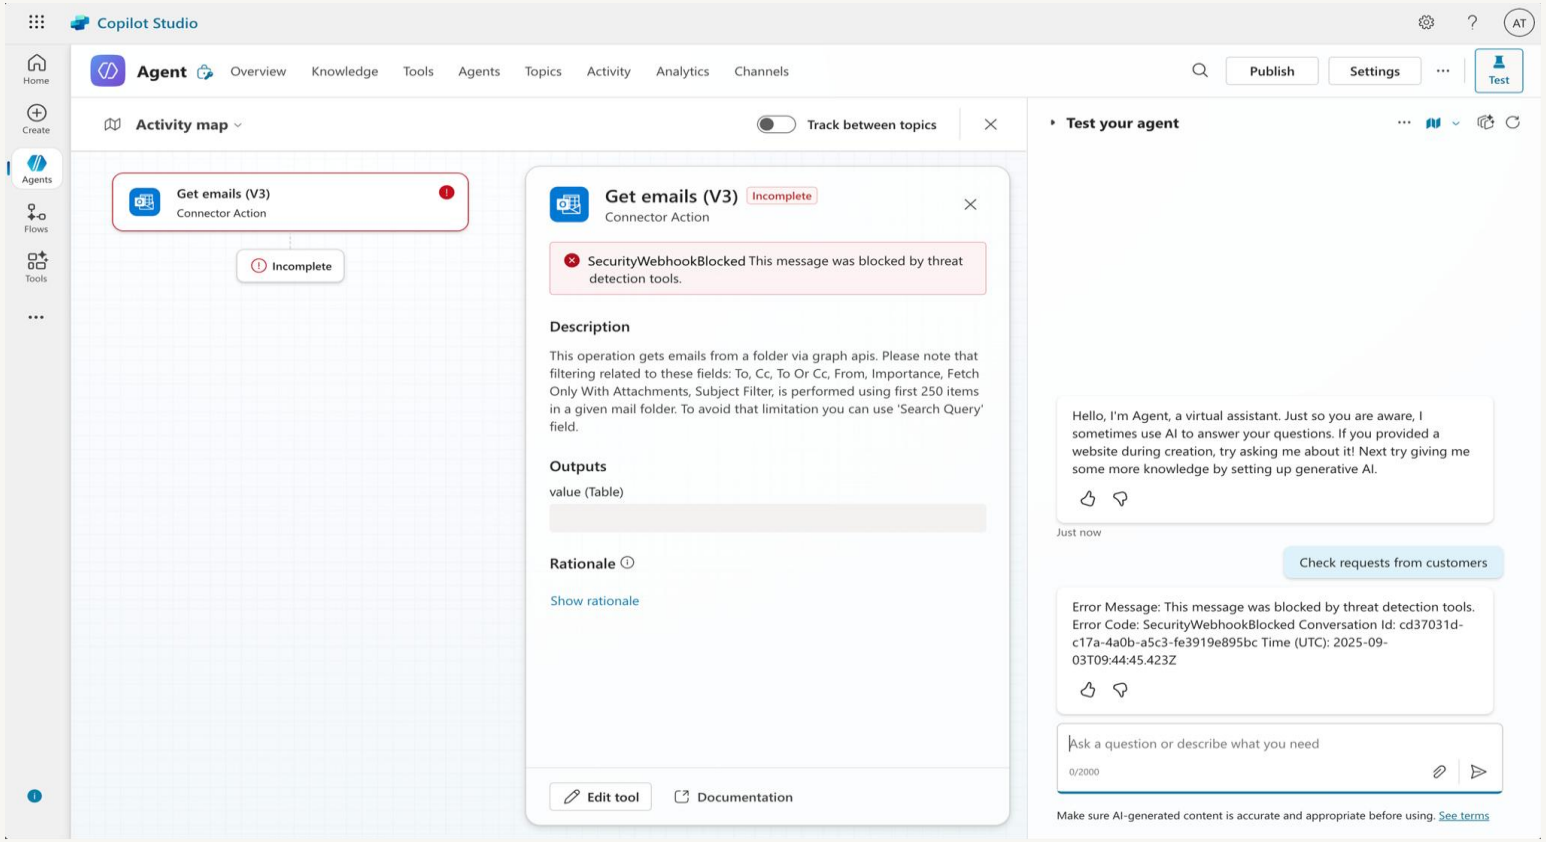

Copilot Studio: a malicious connector call blocked in real time by Defender.

Discover

Auto-inventory all Copilot Studio + Foundry agents.

Posture

Misconfigs, ownerless agents, risky scopes.

Detect

Prompt injection, jailbreak, credential theft.

Block

Stop the action. Alert lands in XDR + Sentinel.

PURVIEW · SECURE

Catch the agent behaving badly.

Insider Risk for Copilot

Risk-score humans AND agents. Same engine.

Risky Agents (Preview)

Visibility for agents in Copilot Studio + Foundry.

Comm Compliance

Detect harassing, unethical, or leaky prompts in flight.

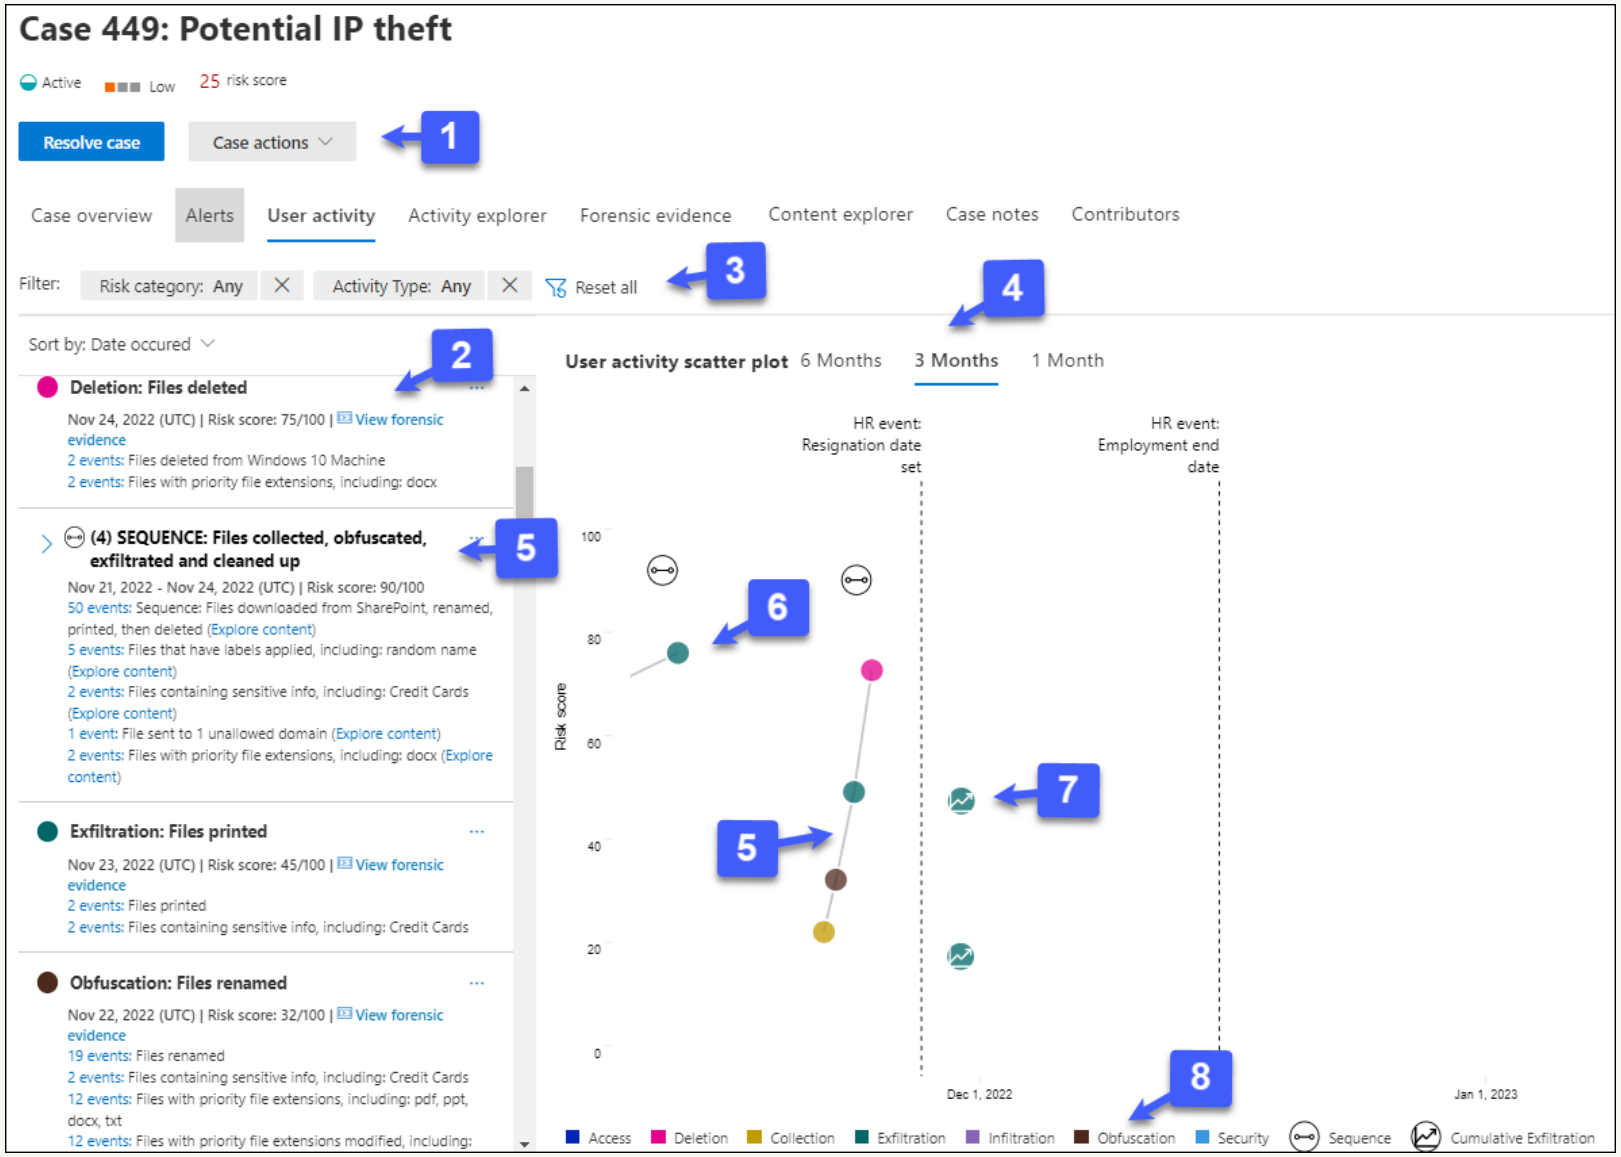

Insider Risk Management — case view with risk score and timeline.

GROUP THE AGENTS

Let them collaborate.

Same environment + shared connectors + Teams = agents that talk.

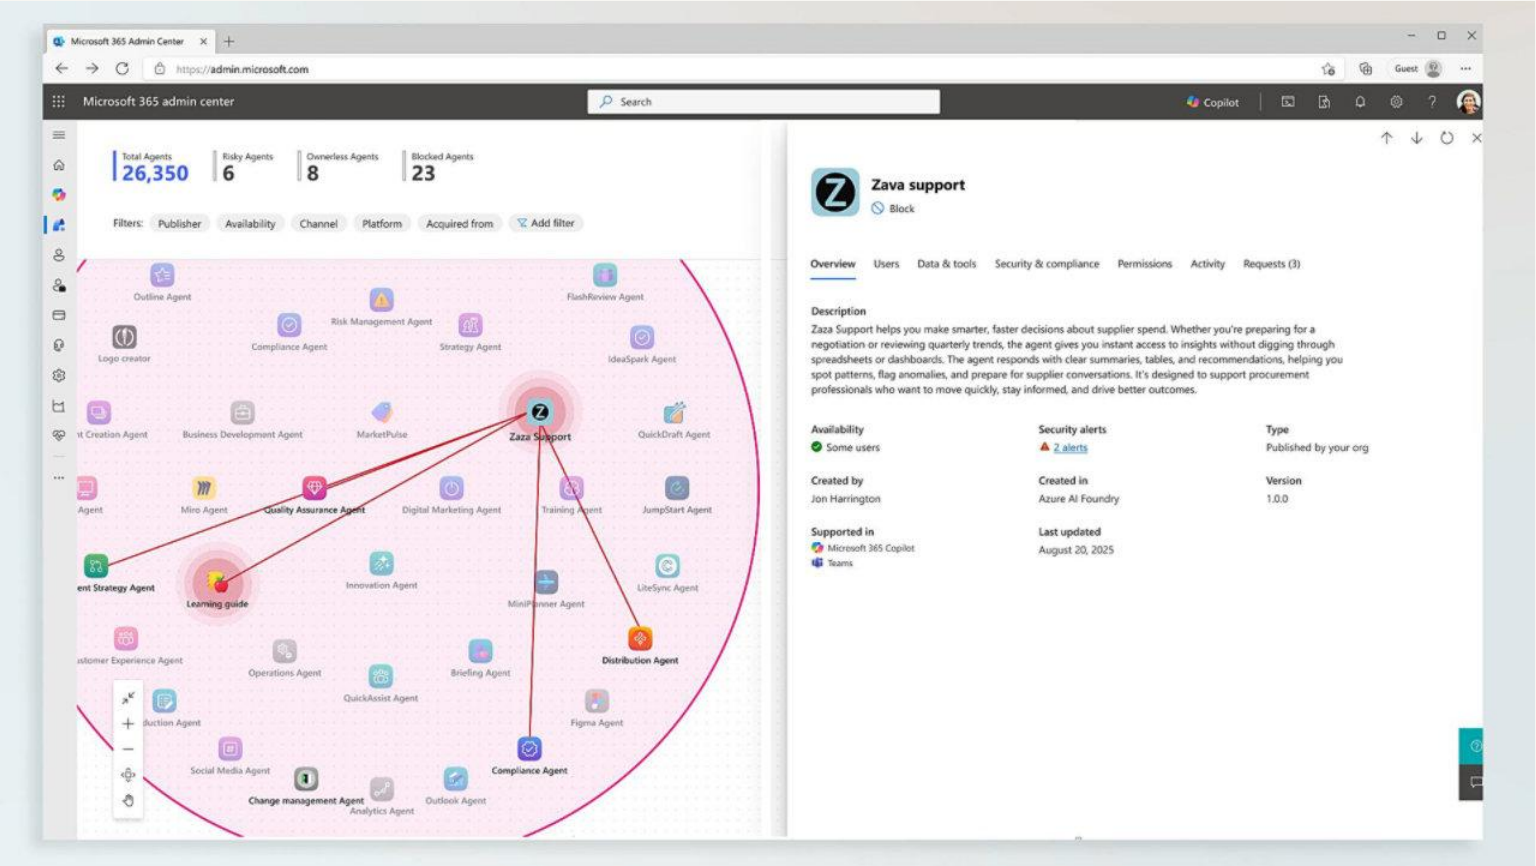

Agent 365 visualises the social graph of your fleet — who calls whom.

Same environment

Co-located in one Power Platform environment.

Shared connectors

Same Business data group in DLP — agents can chain.

Teams Mode

Agents become teammates in a Copilot group chat.

NEW · IN PREVIEW

Collections.

Decide who can see whom — at the agent registry layer.

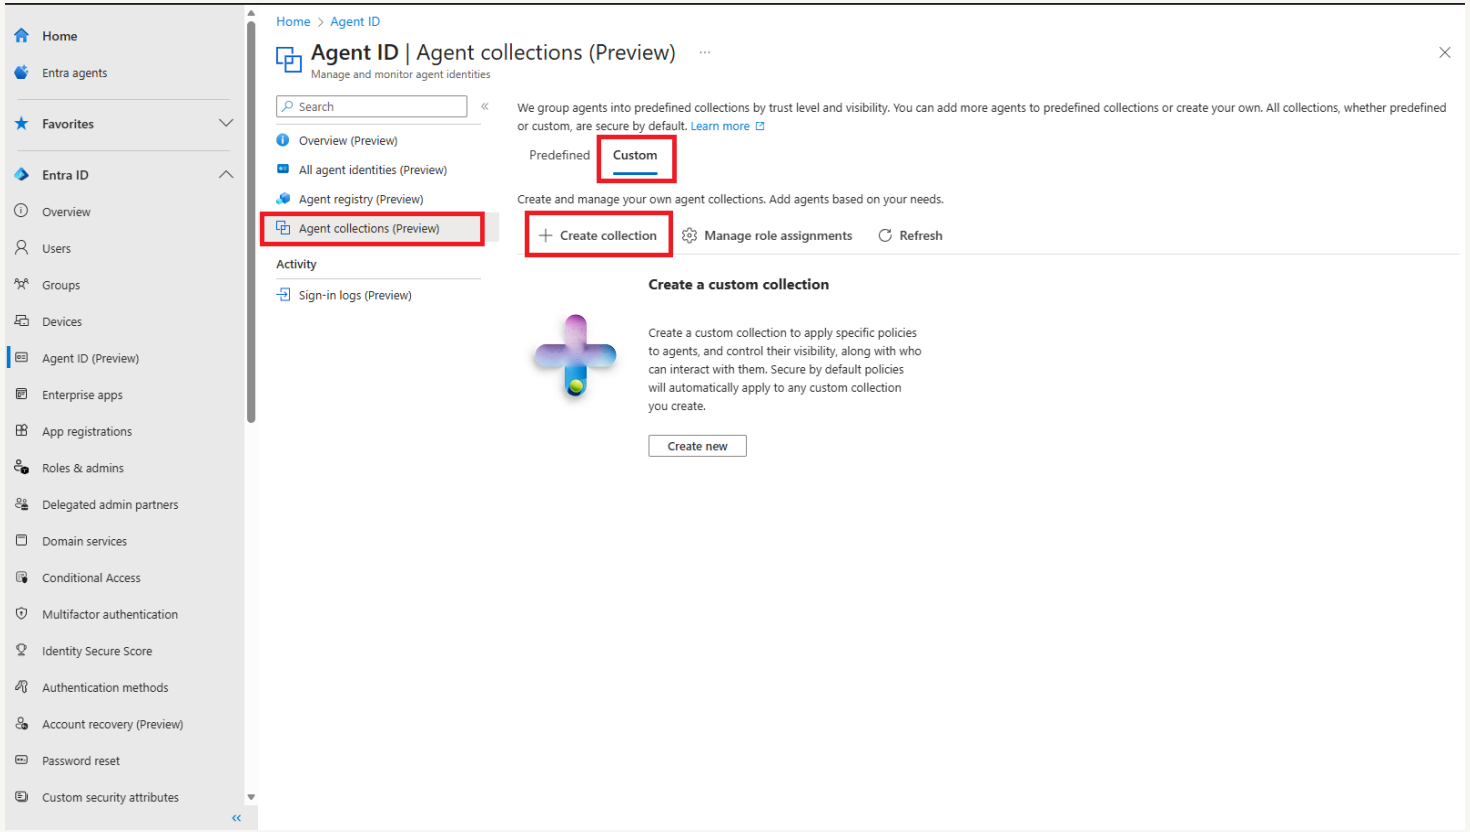

Entra Agent ID → Agent collections (Preview) — group, scope, and quarantine.

Global

Discoverable by all.
The open commons.

Custom

HR sees HR. Finance sees Finance. You set the walls.

Quarantined

Can't discover anyone.
Can't be discovered.

Groups answer "who can enter the room." Collections answer "who can be seen."

THIRD-PARTY AGENTS

Adobe. SAP. ServiceNow. Workday. Manus. And the next ten.

Treat them like contractors: badge, scope, audit.

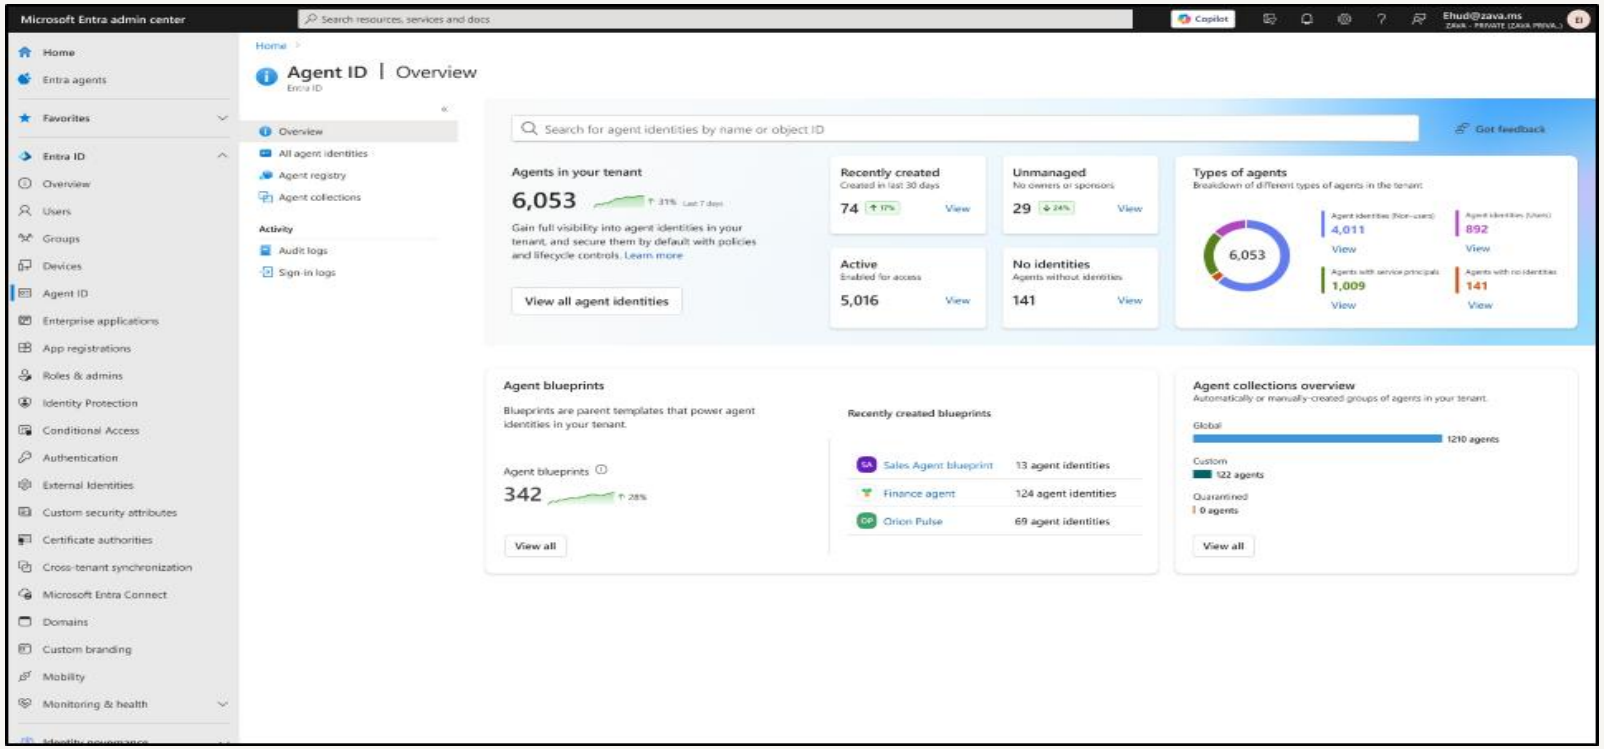

Entra Agent ID — every third-party agent gets a managed identity.

Register

Every agent gets an Entra Agent ID.

Scope

Conditional Access. Least privilege on connectors.

Approve

Admin gate before org-wide publish.

Watch

Audit logs flow into Purview + Defender.

OWASP for Agentic AI & Monitoring

Key OWASP Agentic Risks

- ASI01: Agent Goal Hijacking
- ASI02: Tool Misuse
- ASI03: Privilege Escalation
- ASI06: Memory Poisoning
- ASI08: Prompt Injection

Monitoring Best Practices

- Azure Monitor Alerts for anomaly detection
- Microsoft Sentinel for AI threat correlation
- Application Insights for agent performance
- Purview audit logs for compliance trail
- Defender for Cloud for posture management

Pre-Deployment Security Checklist

- Security review of AI platform data handling policies
- Compliance assessment (GDPR, industry regulations)
- OWASP agentic risk assessment for each use case
- Network architecture review for proper isolation
- Zero trust access controls (least privilege, short-lived credentials)
- Comprehensive audit logging with SIEM integration

Trust and components

CD-SEC-01 - Blind Trust

CD-SEC-06 - Vulnerable and untrusted components

CD-SEC-08 - Injection handling failures

Sensitive data leakage

CD-SEC-04 - Sensitive data leakage and handling failures

Identity and access

CD-SEC-02 - Account impersonation

CD-SEC-03 - Authorization misuse

CD-SEC-05 - Auth and secure communication failure

Operational hygiene

CD-SEC-07 - Security misconfiguration

CD-SEC-09 - Asset management failures

CD-SEC-10 - Security logging and monitoring failures

ONE GOVERNANCE PICTURE · FOUR SURFACES

Surface	Control plane	Identity	Data	Threat	Observability
M365 Copilot	**Agent 365**	Entra + CA	Purview + Restricted Search	Defender XDR	Purview audit + Compliance Mgr
Studio Lite (Chat)	**Agent 365**	Entra (signed-in)	Inherits Purview + DLP	Defender for Cloud Apps + XDR	Purview + PPAC + Agent Map
Studio Full	**Agent 365**	Entra Agent ID + RBAC	PPAC DLP + Purview	Defender CASB real-time + XDR	PPAC + Sentinel + App Insights
Foundry / Agent Service	**Agent 365**	Managed identity	Content Safety + Purview	Defender for Cloud (AI posture)	Tracing + Sentinel + XDR

Microsoft Agent 365 sits across every surface. Pick the lightest one that meets the data sensitivity — promote agents up the stack as they mature.

DATA SOVEREIGNTY · EUROPEAN VIEW

Tenant region ≠ inference region. For an EU/EFTA tenant, ask: where is data ***held***, and where is it ***processed***? The two answers diverge per surface and per model.

Surface / Model	Data held in EU	Inference processed in EU	Sovereignty agreement
M365 Copilot — OpenAI	Yes — EU Data Boundary; in-country (DE/SE/CH coming 2026)	Yes by default — disable Flex Routing to enforce EU-only	EU Data Boundary + DPA + SCCs; Customer Key / DKE optional
M365 Copilot — Anthropic	Out of scope for EU Data Boundary	Runs on Anthropic-managed infrastructure (often US)	MS subprocessor terms only — admin toggle to disable
Copilot Studio (generative answers)	Yes — EU Data Boundary when env. is in EU	Yes — Azure OpenAI endpoint in same boundary	EU Data Boundary commitments apply
Foundry — OpenAI (Azure Direct)	Yes — EU region (e.g. Sweden Central, Germany WC)	Yes — Azure-hosted in selected EU region	Microsoft Product Terms + DPA + SCCs; Sovereign Cloud option
Foundry — Anthropic (Serverless)	Foundry resource in EU — but model is Global Standard	No — routes to Anthropic global infra; EU-native target 2026	Anthropic commercial DPA + SCCs (not data sovereignty)
Foundry Local / Sovereign	On-prem or sovereign cloud — you choose	On your hardware; cloud-mirrored APIs	Customer-controlled — strongest sovereignty option

OpenAI on Azure + EU tenant = full sovereignty. Anthropic = strong DPA, not yet EU-pinned compute. Foundry Local = the highest bar.

Scenarios – Present at 3:45pm

- Assume this is a real organisation
- Discuss a solution including Copilot agents that would meet some of the goals
- Discuss Governance and Security concerns
- Prepare a 5 minute architecture story

Scenario 1 - Procurement Pressure at a Manufacturing Firm

A mid-sized manufacturer is struggling with procurement delays and rising costs. Buyers negotiate using information scattered across emails, SharePoint folders, legacy ERP exports, and supplier PDFs. Each quarter, leadership asks for explanations when material costs spike, but answering requires days of manual research and inconsistent assumptions.

Procurement staff often ask the same questions: Which suppliers are missing SLAs? What changed since last quarter? What similar negotiations worked well before? There's growing concern about sharing sensitive supplier data too widely, while still enabling faster decisions.

The company wants to improve insight during negotiations and reduce rework, without replacing core systems. Teams often collaborate in meetings to prepare negotiation packs, annotate documents, and capture decisions that later need to be auditable. There is interest in augmenting daily tools with contextual assistance, while also exploring whether the same patterns could later be exposed to suppliers or category managers through a controlled interface.

Scenario 2 - HR Policy Chaos in a Global Consultancy

A global consultancy employs 6,000 staff across 12 countries. HR policies vary by region, change frequently, and live across intranet pages, PDFs, and legal addendums. Employees regularly ask managers or HR partners the same questions about parental leave, expenses, flexible working, and benefits. Answers are often out of date or inconsistent.

HR wants to reduce interruptions while ensuring responses are accurate, compliant, and traceable to source material. There's also concern about how advice is phrased, since policy interpretation can create legal risk.

New starters are overwhelmed, managers want quick summaries, and HR needs insight into what people are actually asking. Occasionally, HR would like to trial new policy wording internally before publishing. Longer-term, leadership wonders whether similar conversational experiences could support managers during performance or wellbeing conversations, while still respecting permissions and regional boundaries.

Scenario 3 - Sales Enablement in a Regulated Industry

A financial services firm has a large sales team operating under strict regulatory oversight. Pitch decks, product sheets, and pricing models change frequently, and sellers often reuse old content without realising it's outdated. Compliance reviews happen late, creating delays and frustration.

Sales leadership wants reps to prepare faster and stay within guardrails, especially when responding to bespoke client questions. Reps often work together in Teams chats and meetings to draft proposals, summarise previous deals, and review call notes.

There's interest in smarter assistance during preparation and follow-up, including pulling context from CRM, past conversations, and approved content libraries. Compliance teams want visibility into trends and reassurance that advice given is grounded in approved sources. The architecture must support future integration with customer-facing experiences, without exposing internal reasoning or restricted data.

Scenario 4 - Incident Response in an Energy Company

An energy provider manages assets across multiple regions. When incidents occur, engineers, operations, comms, and leadership scramble to gather information from monitoring systems, handover notes, emails, and past incident reports. Post-incident reviews are time-consuming and often incomplete.

During incidents, Teams becomes the coordination hub, but critical context is missed or duplicated. Staff frequently ask if something similar has happened before, what actions were taken, and who approved decisions. After the fact, regulators and auditors request clear narratives backed by evidence.

The company wants to improve shared understanding during live events and reduce the burden of reporting afterwards. There’s curiosity about how structured and unstructured data could be brought together in real time, and how collaborative workspaces might evolve into durable knowledge assets. Future plans include extending the same approach to contractors and partners, with strict control over what they can see.

Scenario 5 - Project Delivery in a Digital Agency

A digital agency runs dozens of client projects simultaneously. Project managers track status in Planner, emails, spreadsheets, and client tools. Leadership asks for weekly updates, but answers vary depending on who you ask. Risks are spotted late, often buried in long message threads or meeting notes.

Teams want help summarising progress, identifying blockers, and preparing client-ready updates. They collaborate heavily—commenting on plans, refining estimates, and capturing decisions during calls. The agency also runs retrospectives, but insights rarely feed forward into future projects.

There’s interest in an assistant that understands the rhythm of delivery, surfaces patterns, and supports consistent storytelling across clients. The agency is also exploring whether reusable components could be adapted per client, without building custom solutions from scratch each time. Any approach needs to respect client boundaries while still enabling leadership-level insight.

Scenario 6 - Research Overload in a Pharmaceutical Company

A pharmaceutical R&D team produces vast amounts of research: lab notes, trial summaries, regulatory correspondence, and external publications. Scientists spend significant time searching for prior work, understanding decisions, and onboarding new team members.

Cross-functional collaboration is increasing, with research, regulatory, and commercial teams working together earlier. Meetings generate actions and hypotheses, but context is often lost. Staff want to ask natural questions like “Has anyone tested something similar?” or “What concerns did regulators raise last time?”

The organisation is cautious about data sensitivity and intellectual property, but recognises the opportunity to accelerate insight. There’s growing interest in assistive experiences embedded directly into everyday research and collaboration tools, and in customising those experiences for different roles. Longer-term, leadership wonders how the same foundations might support external partnerships or future innovation programs.

Scenario 7 - Field Engineers Maintaining Critical Infrastructure

A utilities company employs hundreds of field engineers who spend most of their time on-site maintaining substations and network equipment. Connectivity is inconsistent, and engineers rely on mobile devices to access manuals, historical maintenance records, safety guidance, and job notes. Much of the most valuable knowledge lives in free-text reports, photos, voice notes, and informal messages sent back to the office.

Engineers frequently encounter situations that are “similar but not identical” to past jobs and want quick answers while on site: Has this fault happened before? What workaround was used? Were there any safety concerns noted? Calling colleagues often interrupts others and still leaves gaps.

After visits, engineers are required to write detailed reports. These are time-consuming and often delayed until the end of the week, reducing accuracy. Office-based planners and asset managers want clearer insight into emerging issues and trends, but don’t want to overload engineers with extra admin.

There is interest in better support before, during, and after site visits, improving collaboration between mobile staff and office teams, and gradually standardising how knowledge from the field feeds back into the organisation.

Scenario 8 - Construction Site Teams and Central Project Control

A large construction company runs multiple projects with teams split between head office and active building sites. Site supervisors spend their days coordinating contractors, handling safety issues, and responding to unexpected changes. Access to plans, method statements, risk assessments, and change requests is fragmented across emails, shared drives, and specialist tools.

When issues arise on site, supervisors often rely on experience or informal advice shared via phone or messaging. Decisions are logged late, if at all, making it hard for project managers and commercial teams to understand why costs or timelines shifted. Safety teams want reassurance that the latest guidance is being followed, especially as regulations change.

Daily stand-ups, site walks, and inspections generate photos, notes, and actions, but these rarely connect back to the central project narrative. New supervisors struggle to get up to speed on large projects, especially when covering absences.

The company wants to strengthen the link between on-site and office-based work, reduce duplicate reporting, and support consistent decision-making—without forcing frontline staff into complex systems that slow them down or pull them away from the site.

TAKE THESE HOME

Five things to do next.

Three for governance. Two for security. All for Monday.

01 GOVERN

Inventory first.

Turn on Agent 365. You can’t govern what you can’t see.

02 GOVERN

Walls before rules.

Map agents to environment groups + Agent Registry collections before any DLP.

03 GOVERN

Label the data, not the agent.

Sensitivity labels + DSPM for AI — the perimeter follows the answer.

04 SECURE

Score the behaviour.

Enable Insider Risk for Copilot + Risky Agents. Humans and agents on the same engine.

05 SECURE

Block in real time.

Defender for Cloud Apps + Security for AI — prompt injection, exfil, alerts into XDR.

Microsoft Copilot Studio architecture links

Category	Resource	Link	Why use it
Well-Architected	Power Platform Well-Architected	https://learn.microsoft.com/en-us/power-platform/well-architected/	High-level design and decision framework
Hub	Power Platform & Copilot Studio Architecture Center	https://learn.microsoft.com/en-us/power-platform/architecture/	Main landing page for architecture guidance
Reference	Copilot Studio reference architectures & solution ideas	https://learn.microsoft.com/en-us/power-platform/architecture/products/copilot-studio	Scenario-based patterns and examples
Patterns	Architecting agent solutions: Principles and patterns	https://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/architecture/overview	Agent design principles and patterns
Solution design	Define your solution architecture	http://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/architecture-overview	Channels, integrations, security, analytics, and ALM

Resources & Next Steps

- M365 Copilot Architecture
- Copilot Studio Documentation
- Azure AI Foundry Documentation
- AI Agent Governance & Security (Cloud Adoption Framework)
- MCP for Agent-to-Agent Communication
- Agent Factory: MCP & A2A Standards
- Voice Agent Integration with ACS
- OWASP Top 10 for Agentic Applications
- Microsoft Agent 365

Thank you for attending. Questions and discussion welcome.

THANK YOU,
YOU ARE AWESOME ❤️

PLEASE RATE THIS SESSION IN THE MOBILE APP.

List Here Your Social Media Links, Email Address, Or Whatever You Think It Is Important :)