Microsoft

Multi-Modal Real-time Agents

Anahita Afshari, Partner Solution Architect
Juan Llovet de Casso, Partner Solution Architect

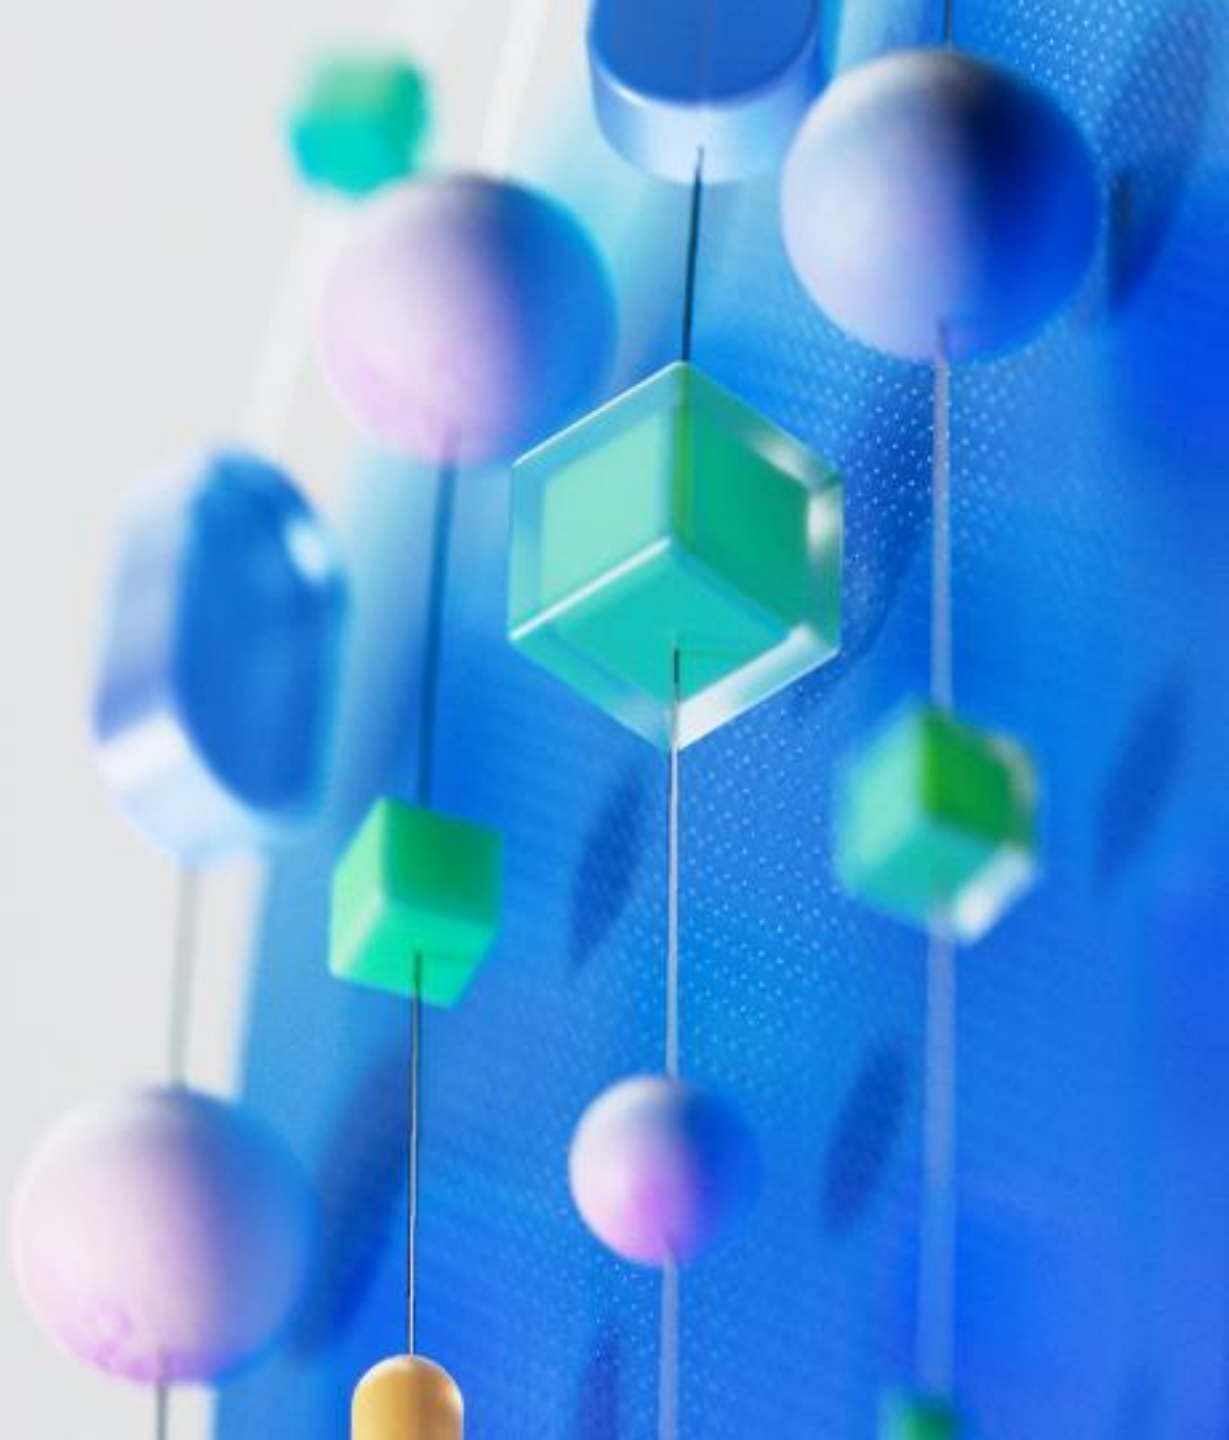

Agenda

- 1.Agents with Microsoft
- ·Frontier
- ·Foundry
- ·Architecture patterns
- 2.Fabric
- 3.Demo

Becoming a Frontier Partner

Success Framework

Enrich
employee experiences

Reinvent
customer engagement

Reshape
business processes

Bend the curve
on innovation

Augment with AI
Across all roles and business functions

Transform with AI
Software development and delivery capabilities

Revolutionize with AI
Monetization models, AI and Agent-based solutions

Disrupt with AI
With products and services that shape your customer's future

What is an agent?

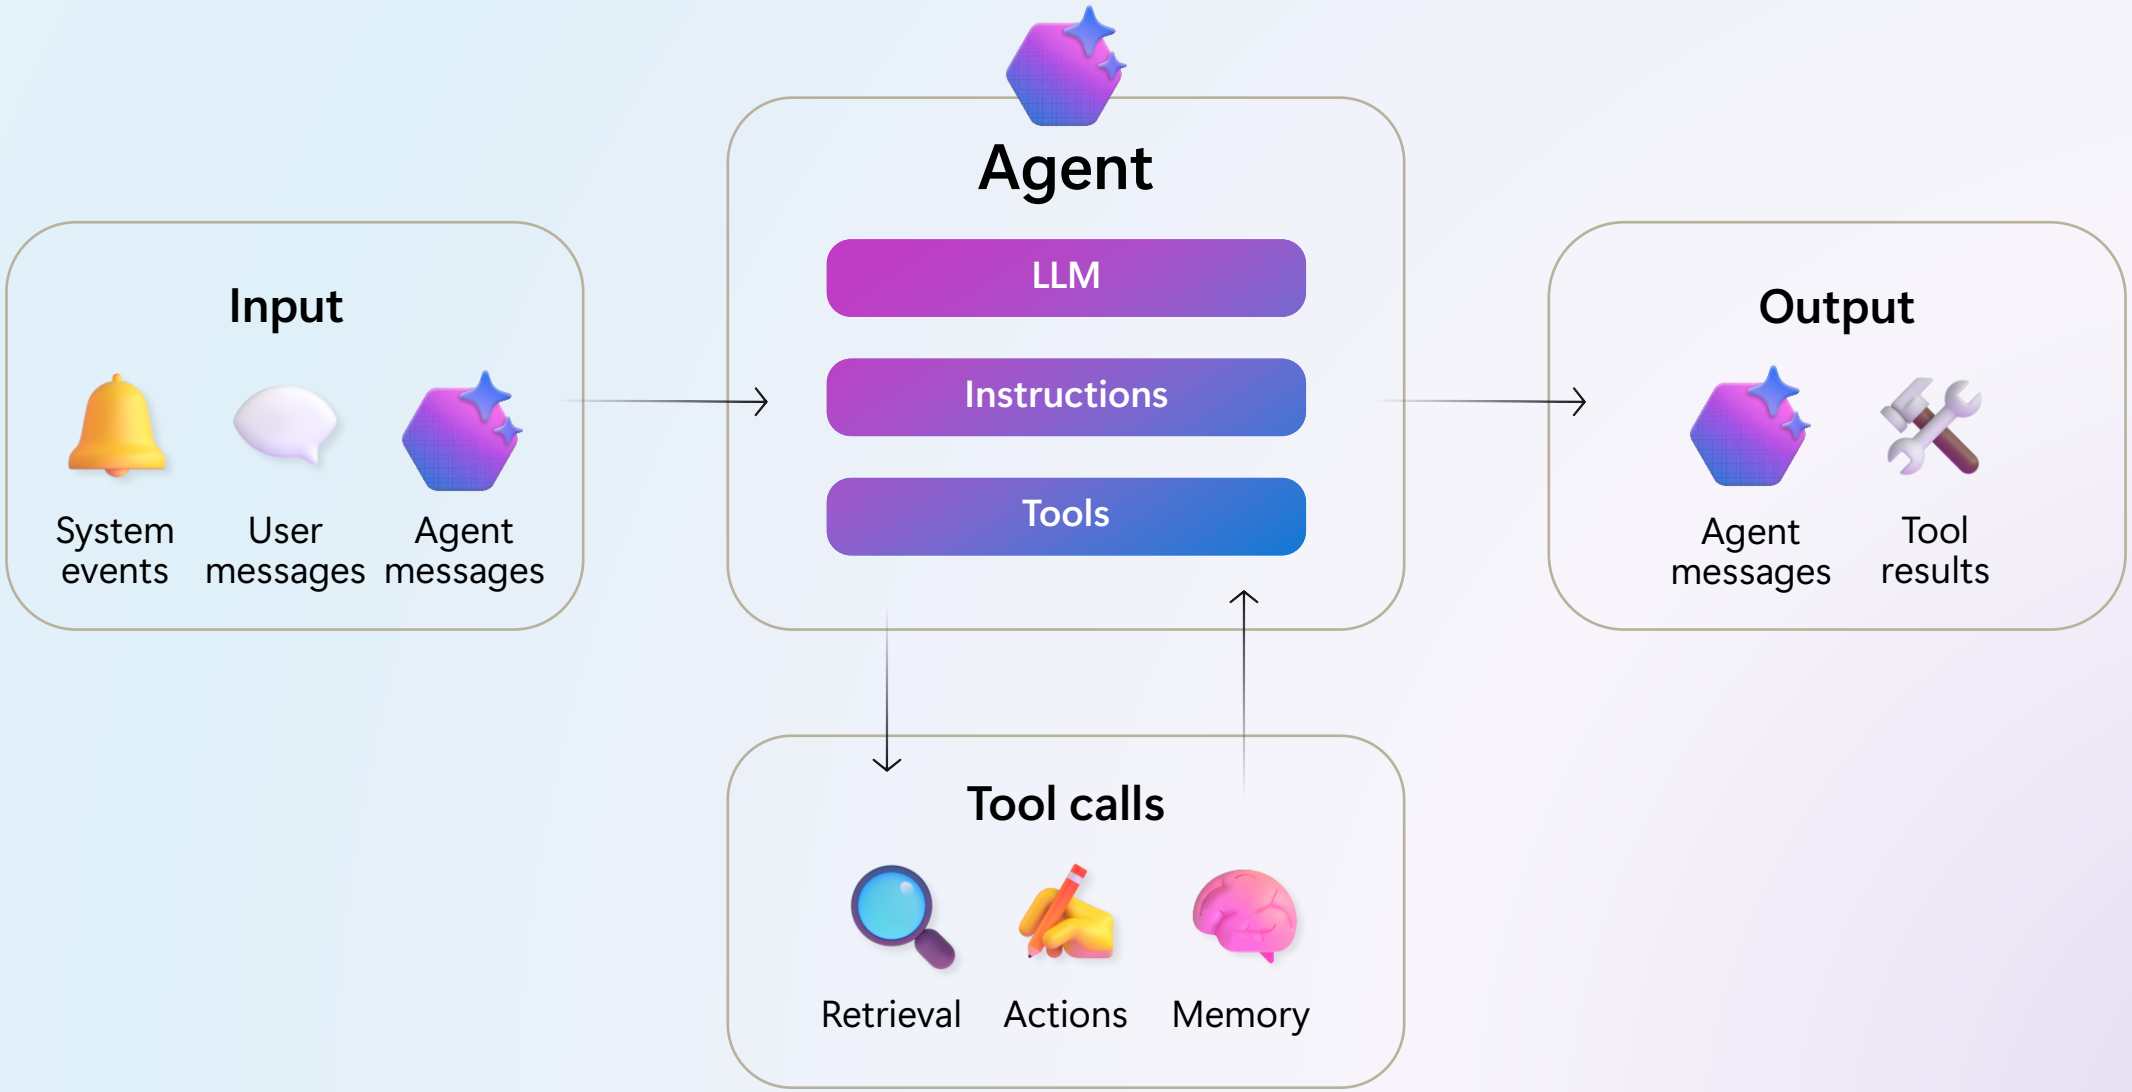

Source: What Are AI Agents, and How Do They Help Businesses?

Agentic AI changes how business operate

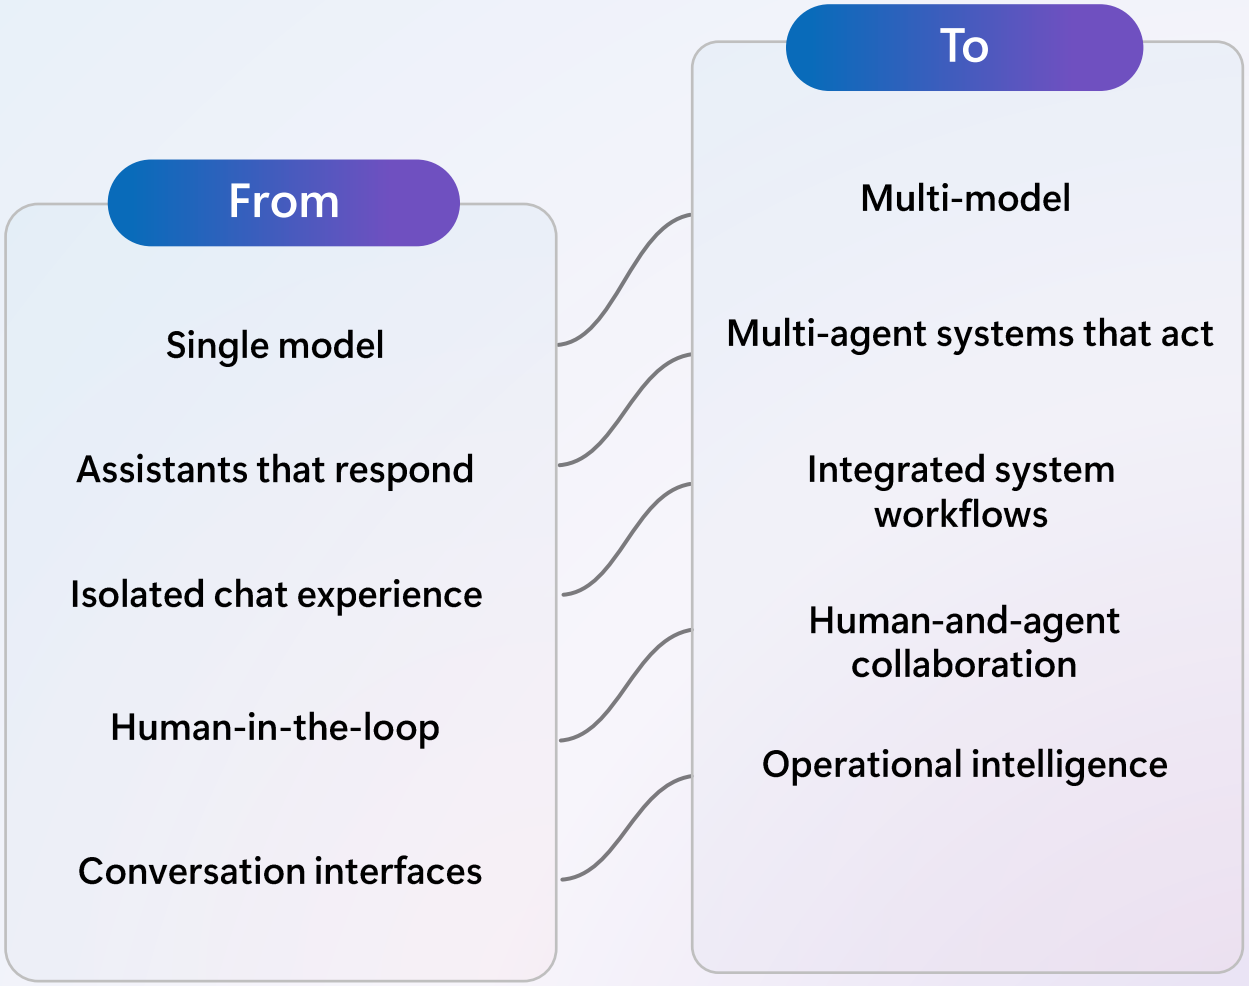

Organizations that master agents can automate multi-step workflows, integrate deeply with systems, reduce cost and cycle time, and deliver new business model

GitHub Copilot's Benefits Extend Beyond Developers...

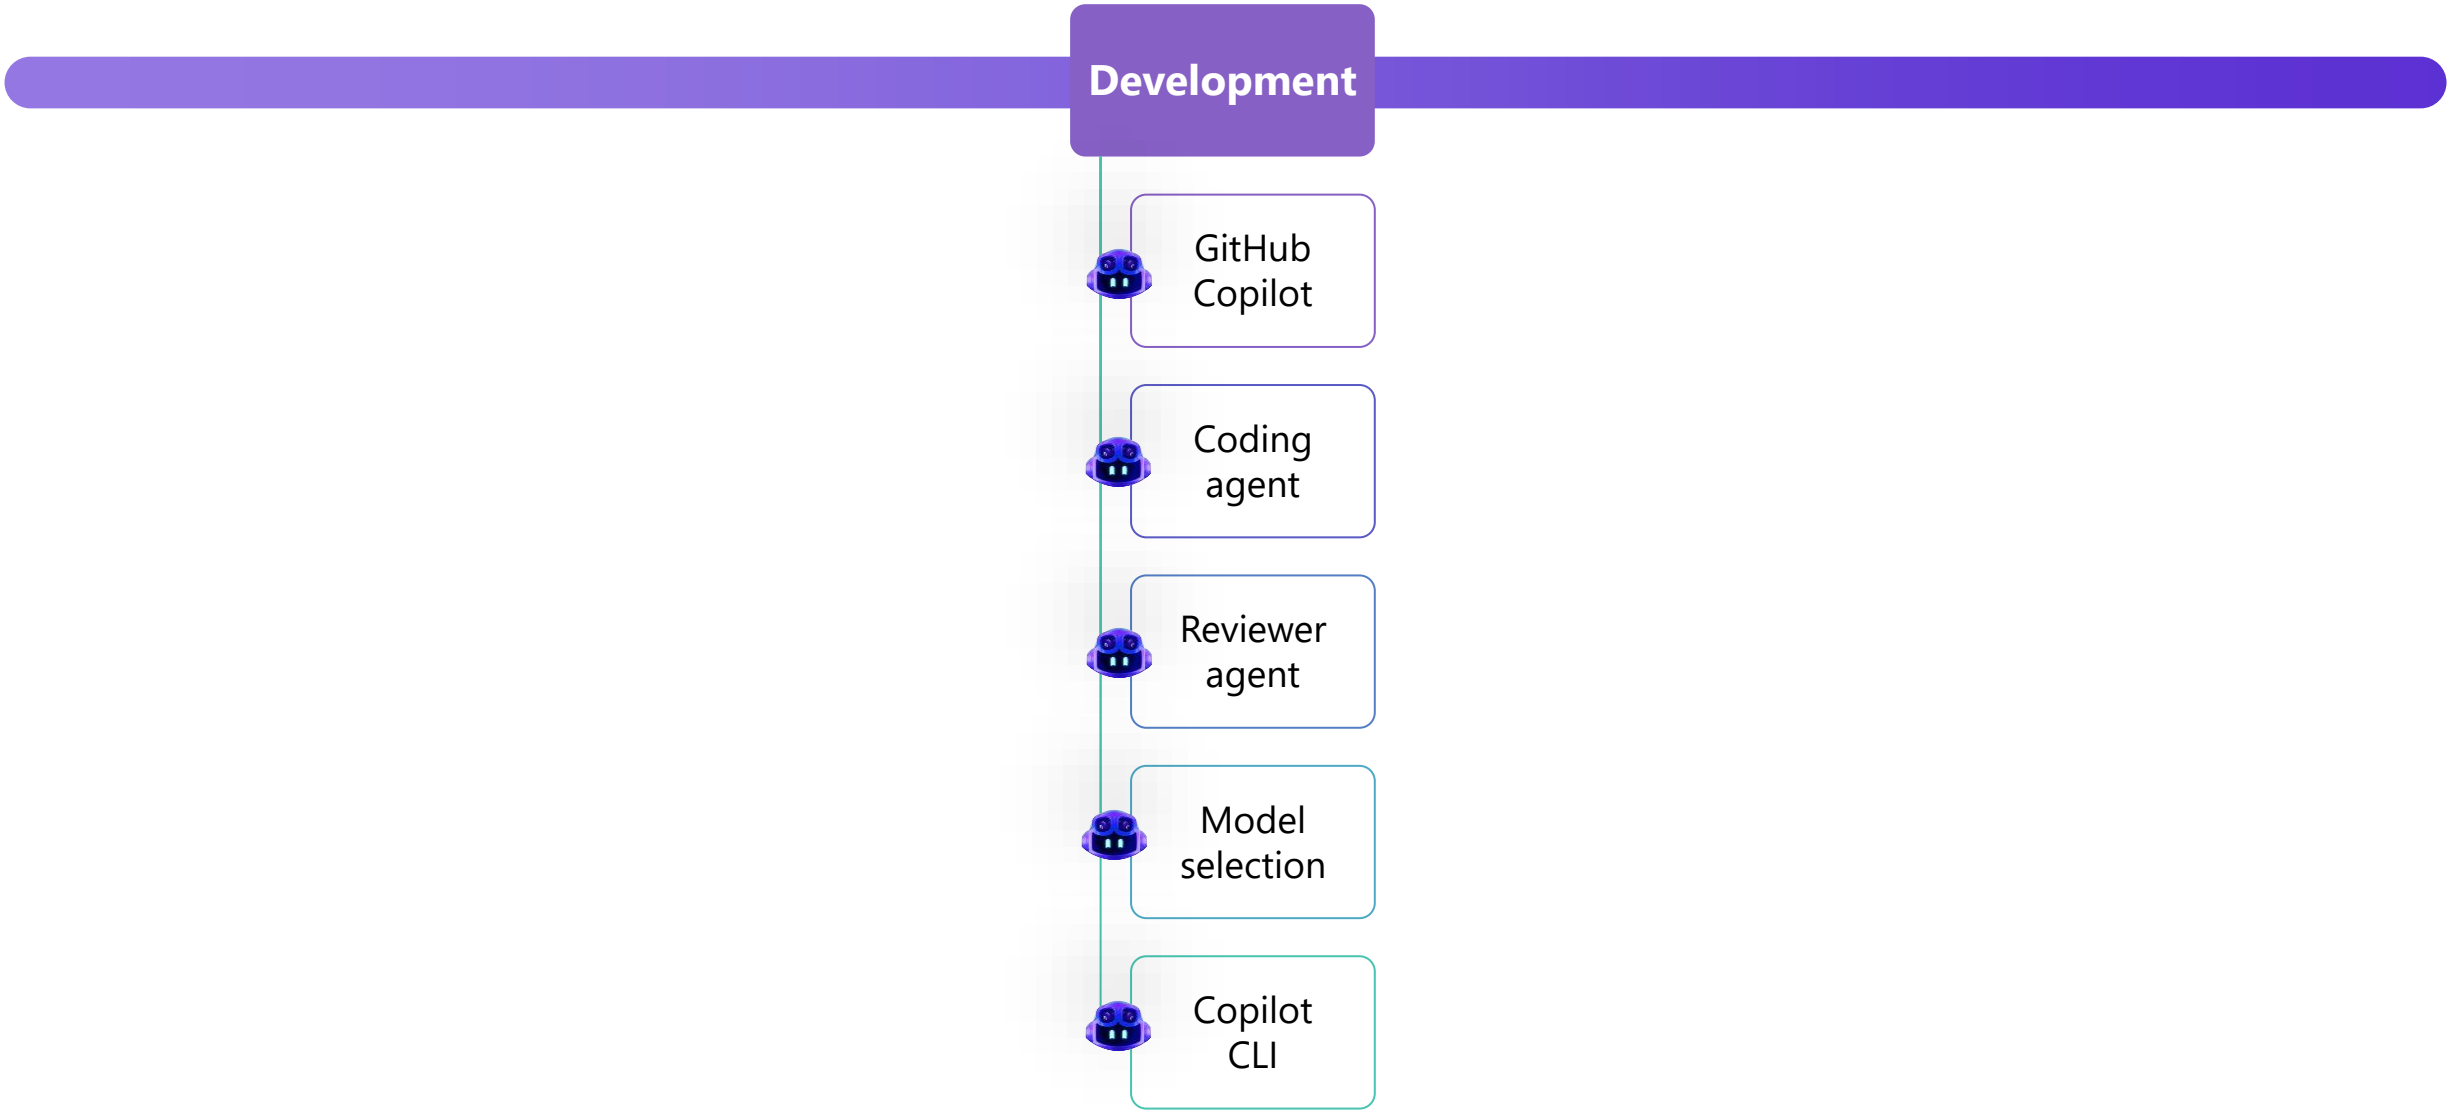

Copilots & Agents Across the Solution Delivery Lifecycle…

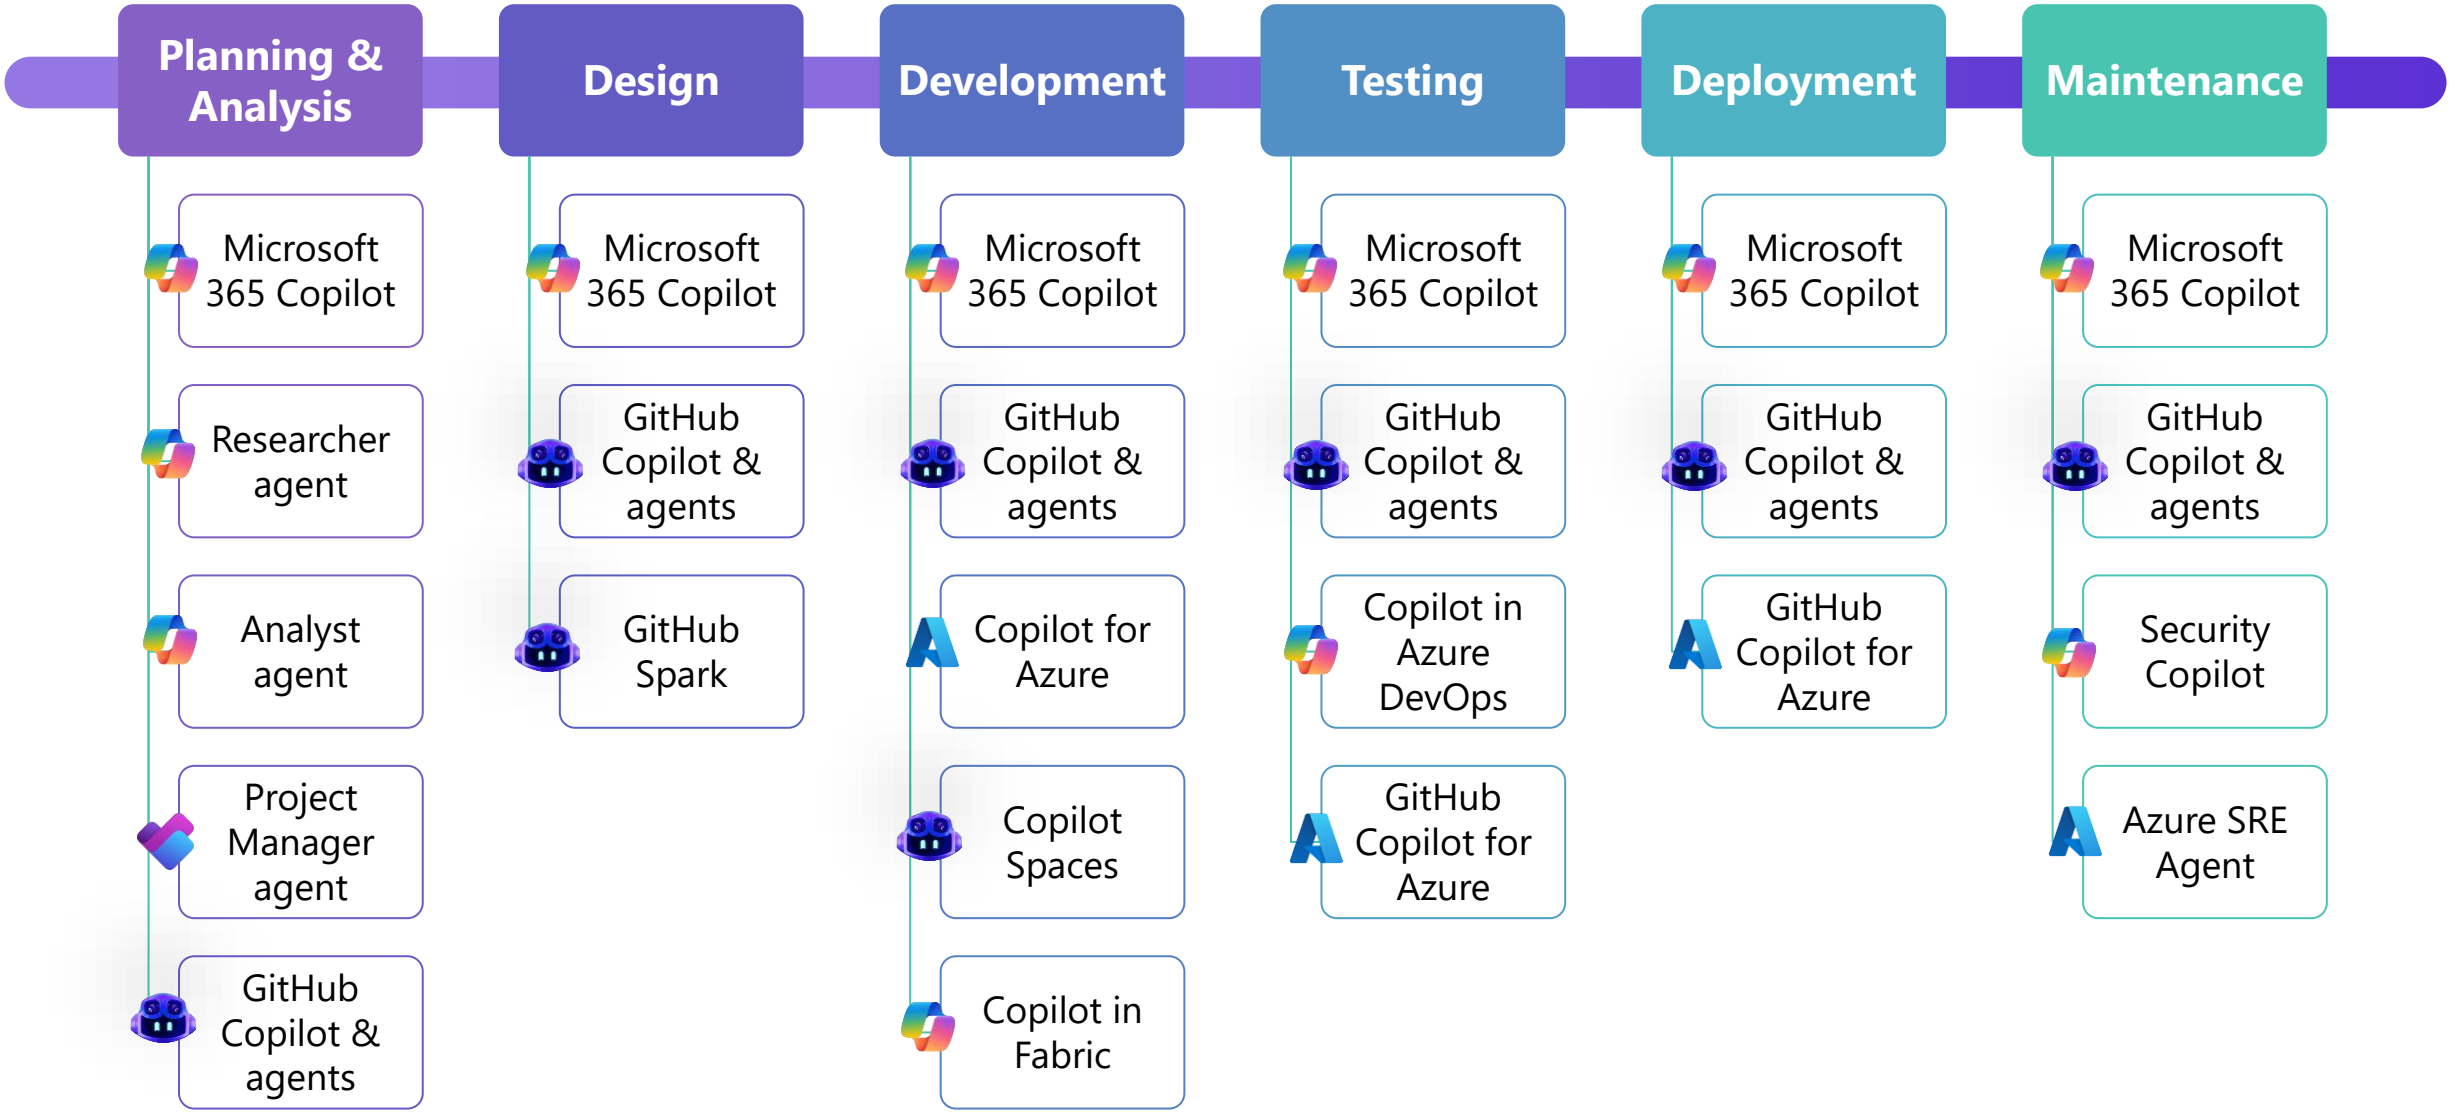

Increasing Productivity For All Roles...

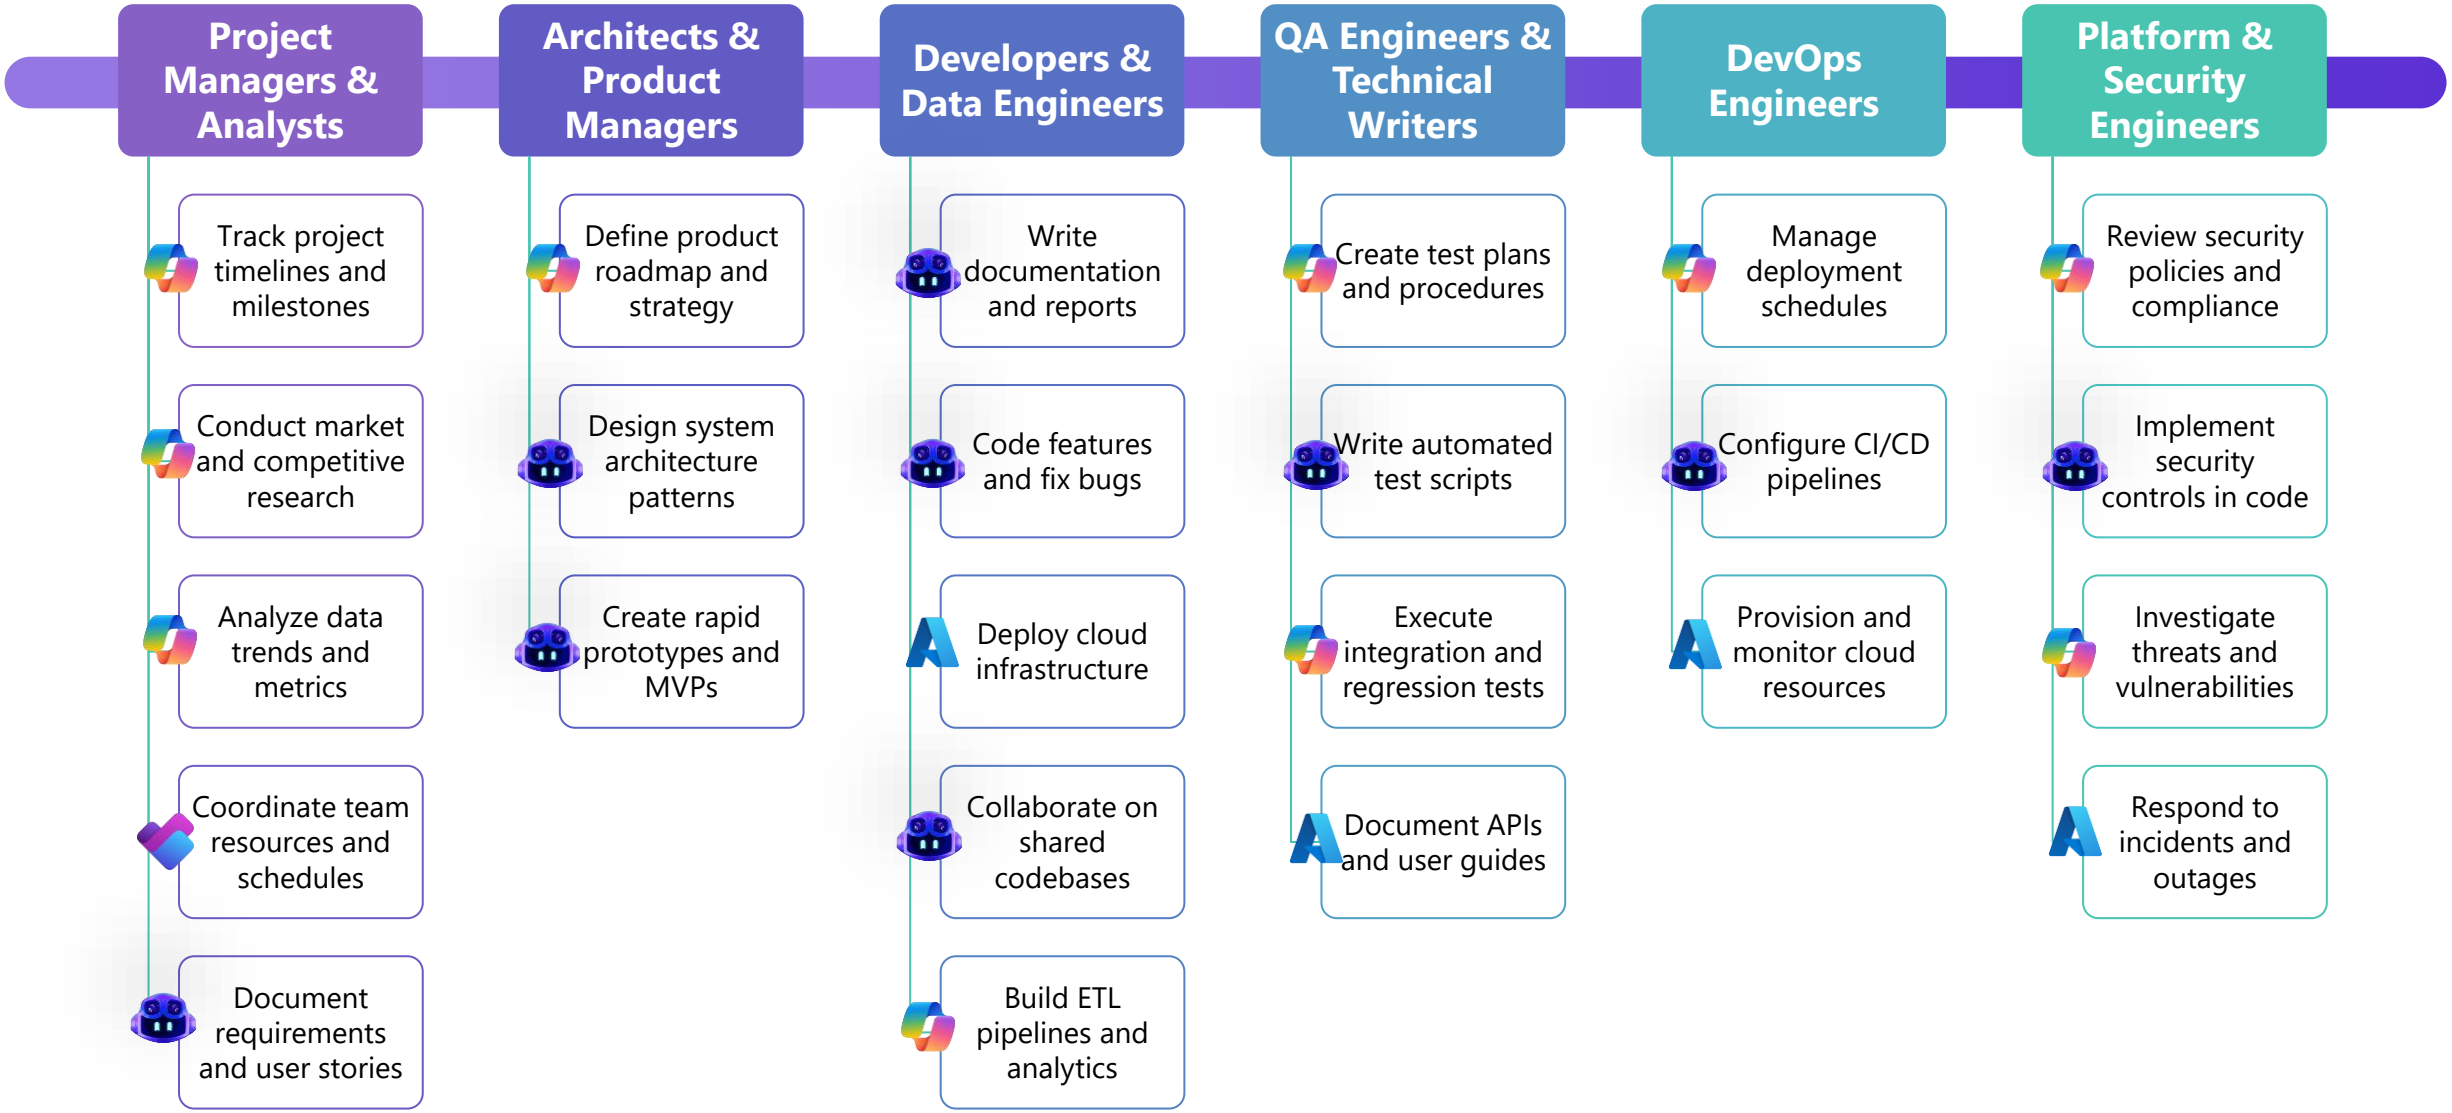

Across All Business Functions in the Organization

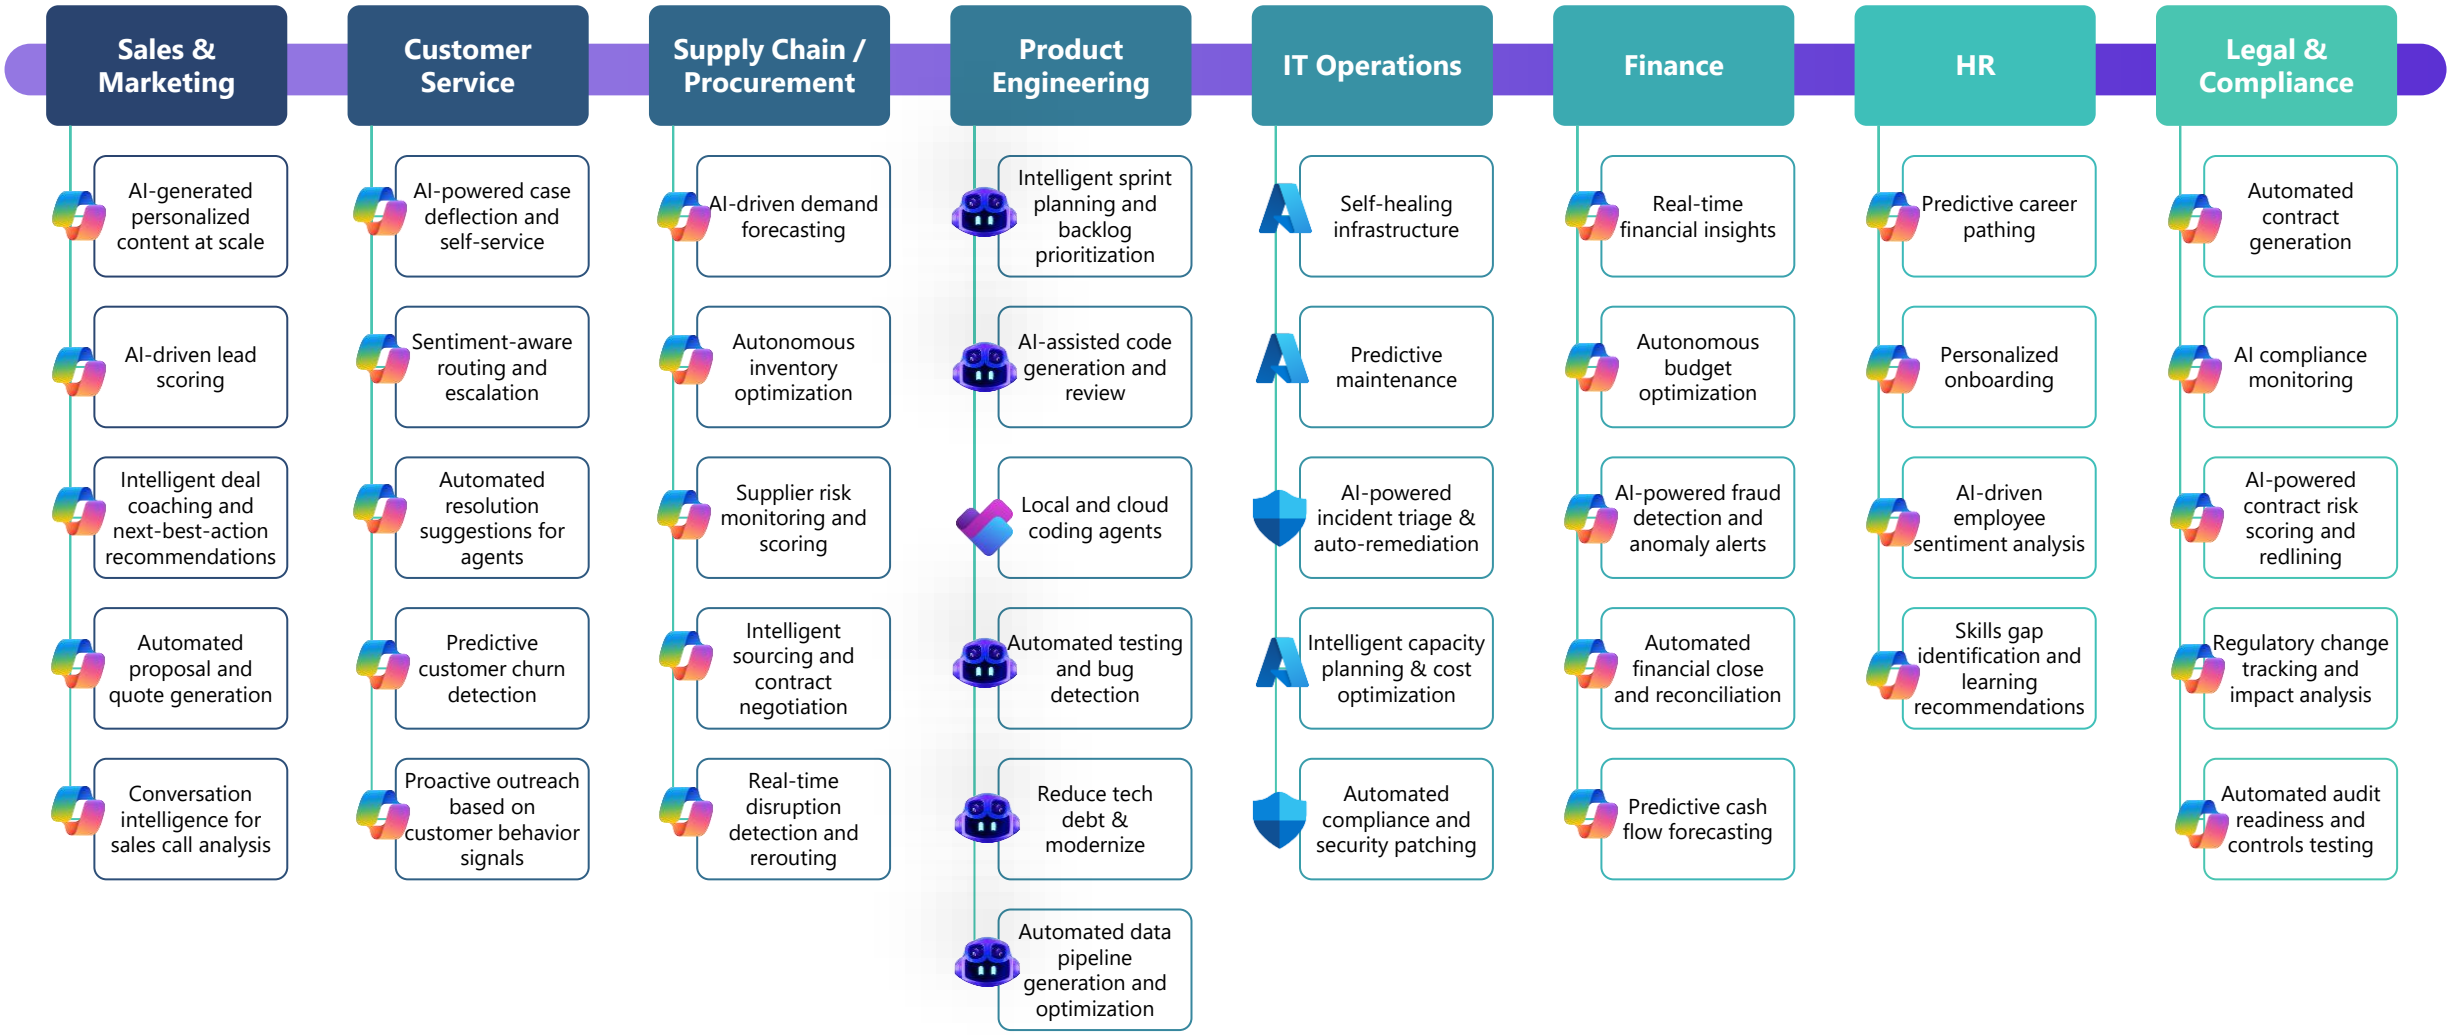

Agentic components

Reasoning
aka planning

Create a JSON describing the user issues

Transcript

I'll look at the logs to see what happened

Check logs

✕ Machine is past due for maintenance

I've captured the underlying issue in the details

```
{
  "name" : "John Clar
  "issue": "Washing m
```

Go beyond transcriptions and slot filling by allowing agents to deduce the *why*

Acting
aka tool calling

Fill out and submit the workorder

Populate field: name

Populate field: issue

✓ Name field populated

✓ Issue field populated

Submit form

The form has been successfully

Agents can take a simple request and chain multiple tool calls to complete the ask

Learning
aka memory

Assign technician

Update work order

✕ Alex is already booked at this same time

\+ Saved memory: Alex is OOF next week

Let me try again with an available technician.

Update work order

✓ Successfully assigned technician

Only make a mistake once. With memory, agents can recall prior experiences

Agentic Architectures

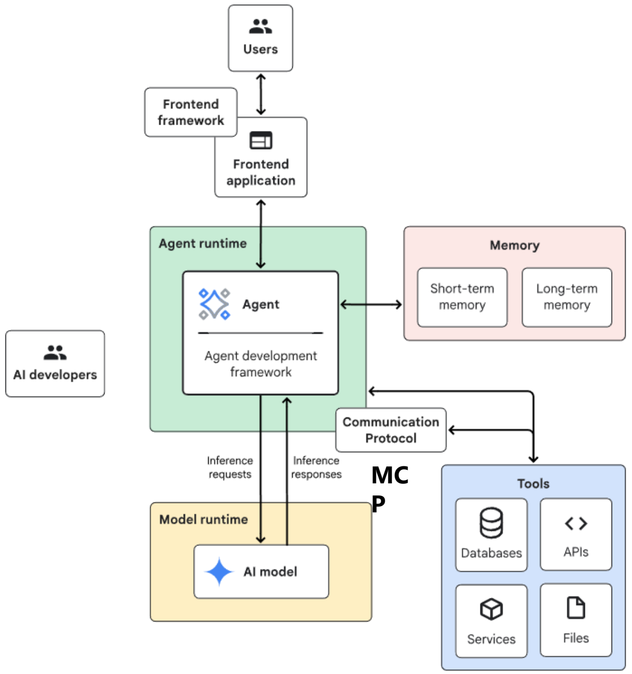

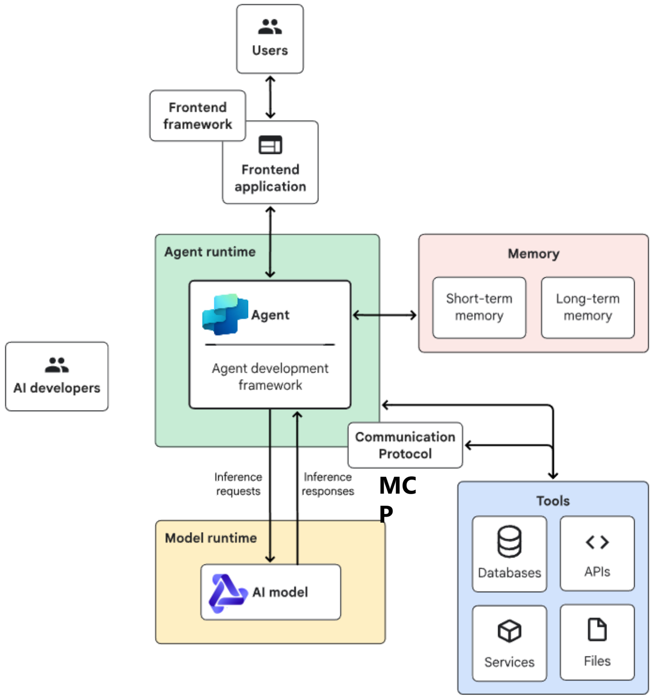

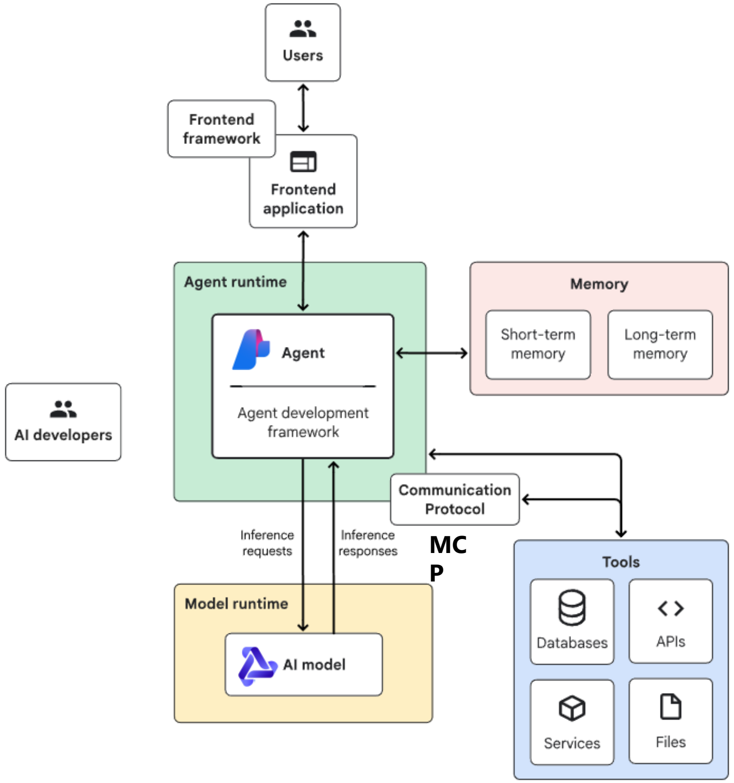

Agents are becoming more capable

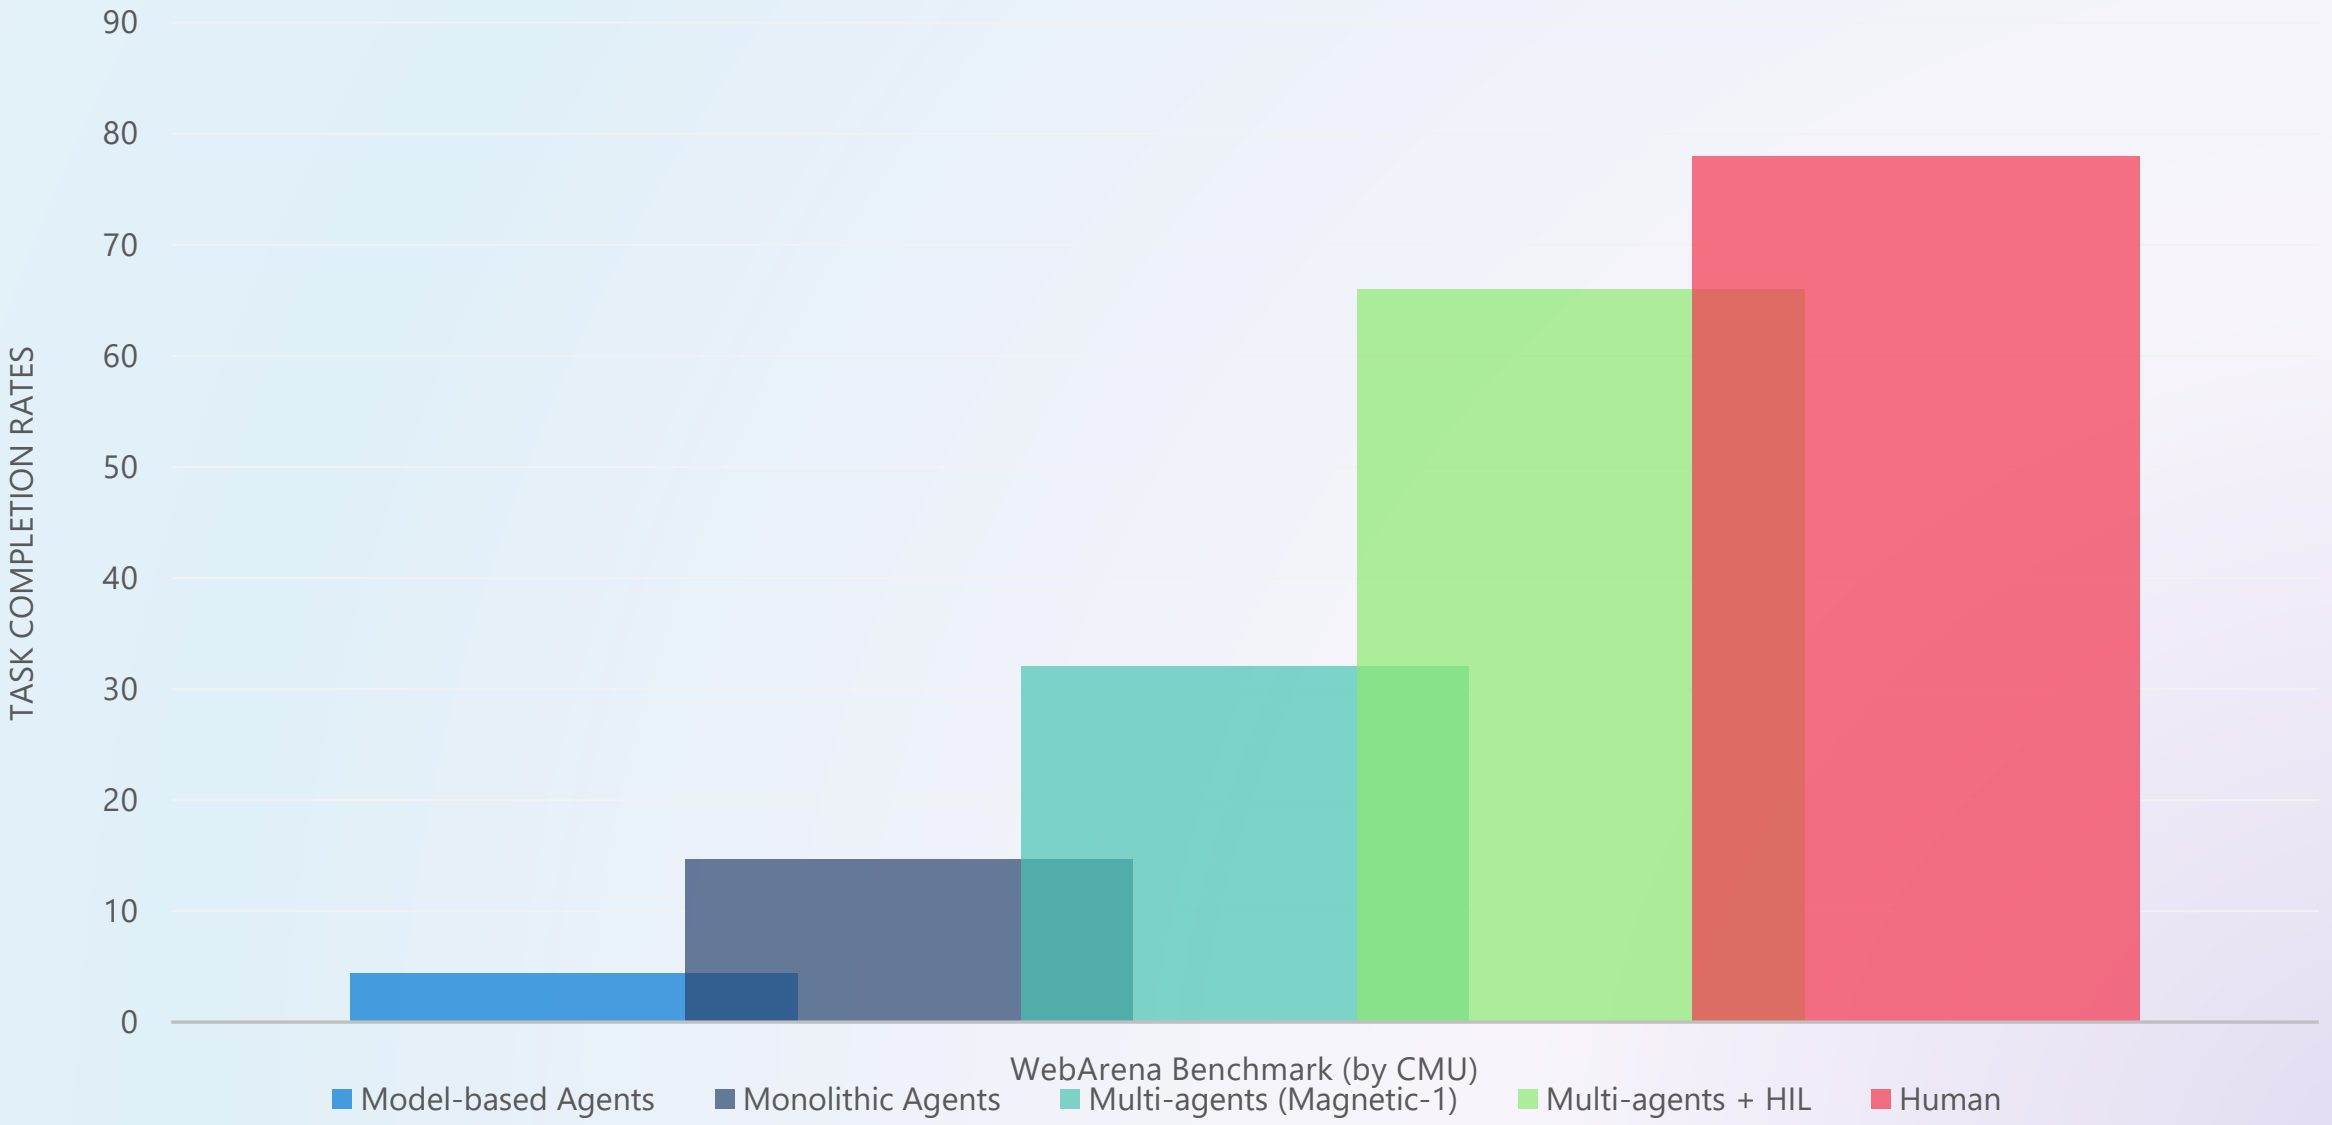

Microsoft Foundry

Streamline development with native IDE experiences

GitHub · Visual Studio · Visual Studio Code · Copilot Studio

Models · Agent Service · IQ · Tools · Machine Learning

Build context-aware and action-oriented agents with 1,400+ pre-built connections and MCP tools

Adobe · ATLASSIAN · SAP · servicenow · UiPath · Microsoft Bing · Microsoft Fabric · Microsoft OneLake

Control Plane

Leverage a complete signals management layer with Microsoft Security integrations

Microsoft Agent 365 · Microsoft Defender · Microsoft Purview · Microsoft Entra

Security, compliance, and governance

Foundry Agent Service

Build, connect, and scale intelligent agents – open, integrated, and enterprise-ready

Open by Design	Connected Intelligence Everywhere	Enterprise Trust & Reach

AI.Azure.com

Foundry Agent Service

Open by design
Build with any framework or protocol

Connected intelligence
Empower agents with memory, knowledge, and over 1,400 MCP-enabled connectors

Enterprise-grade security and reach
Deploy confidently and one-click publishing to Microsoft 365

Build AI apps and agents that scale

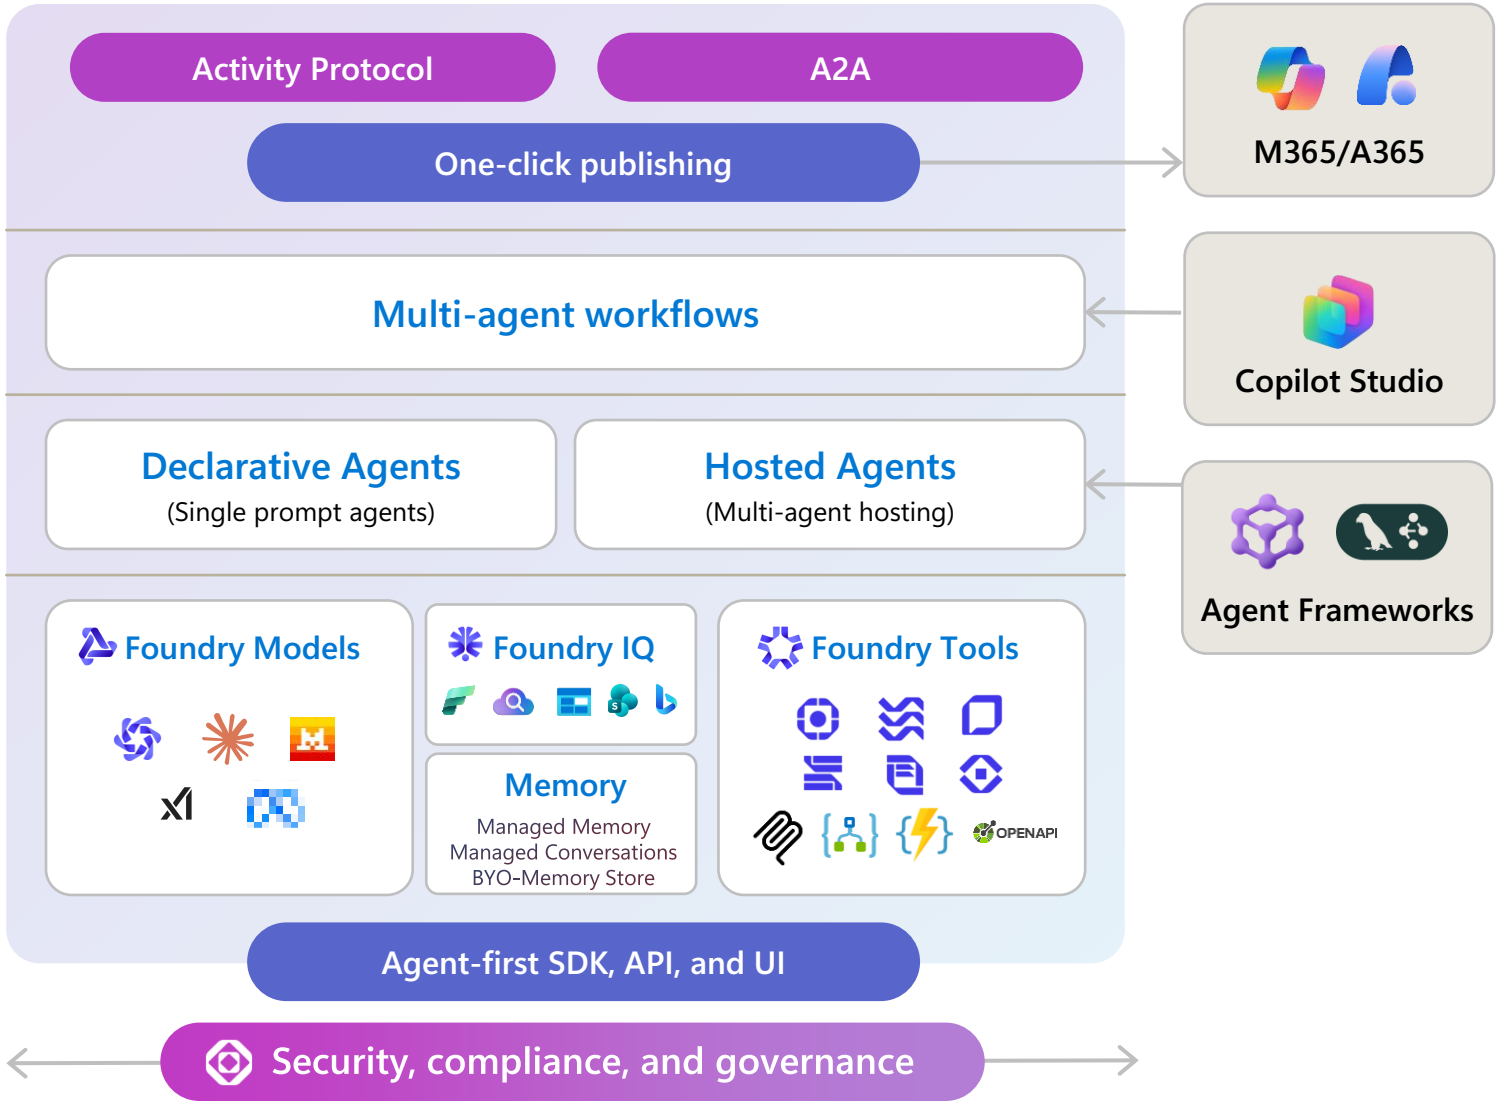

Microsoft Foundry

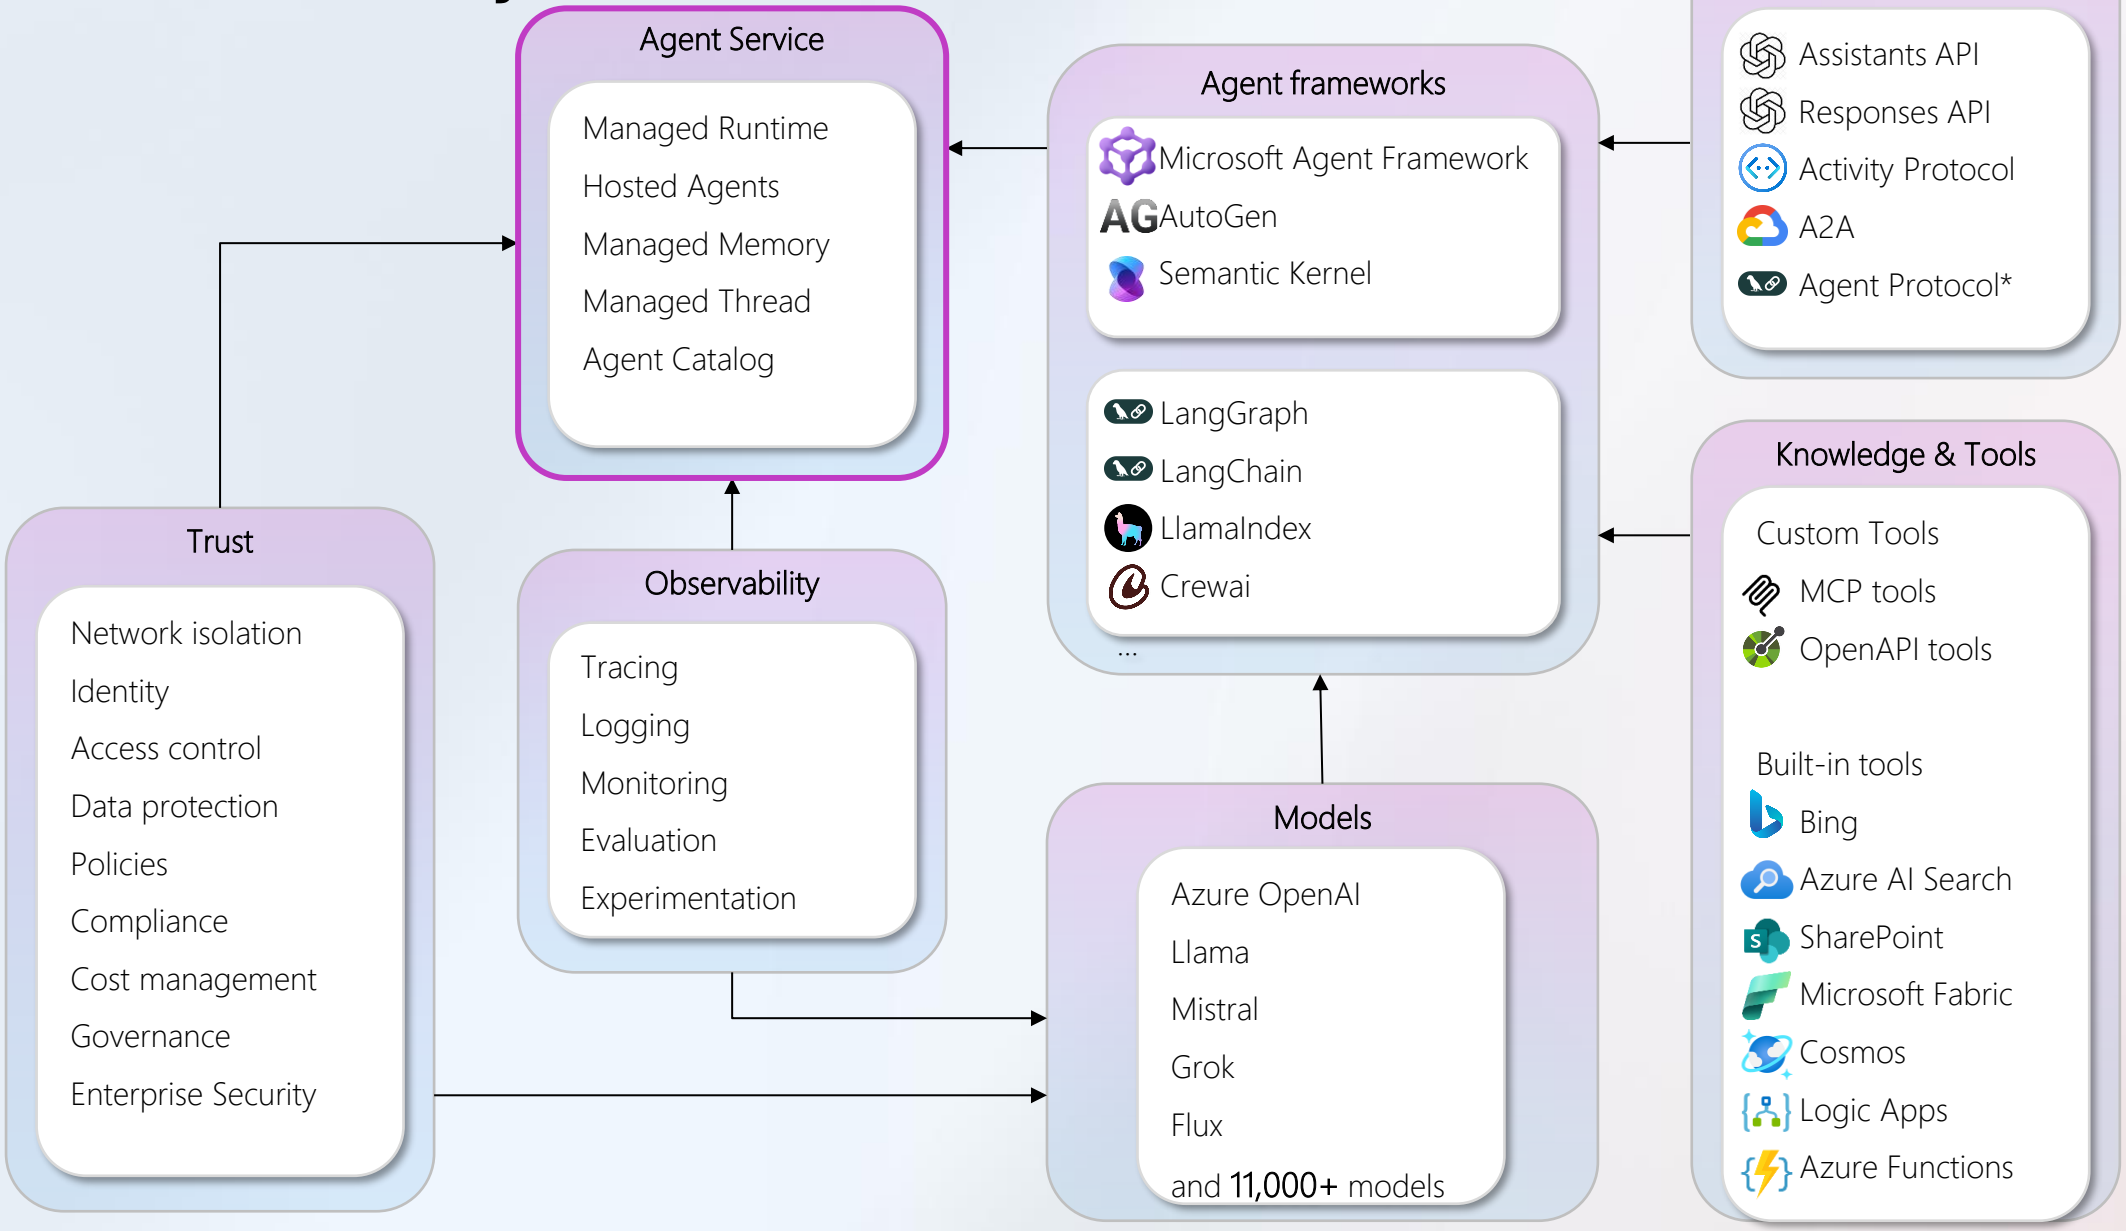

Building on top of industry-leading agentic APIs

OpenAI Responses API

Core API primitive for building agentic applications

- OpenAI Models
- OpenAI Tools

Azure OpenAI Responses API

Enterprise-grade Responses API with Azure's security and compliance

Enterprise-guarantee with 99.9% SLA
+

- Azure OpenAI Models
- Azure OpenAI Tools

Foundry Agent Service

Unified agent platform built on industry-leading APIs

- Support for other Foundry Models and Tools
- Managed memory
- Multi-agent workflows
- Hosted agents
- Built-in evaluation and observability
- One-click publishing to M365

++

Enterprise-guarantee with 99.9% SLA + BYO-resources in your tenant

- Azure OpenAI Models
- Azure OpenAI Tools

Microsoft Agent Factory

Available Now

Accelerate AI innovation and adoption across your organization

Scale AI with one plan

Build agents using Copilot Studio and Microsoft Foundry with a single pre-paid plan

Access deep expertise

Rapidly co-build solutions with Forward Deployed Engineers (FDE) and partners

Grow org-wide AI skills

Get customized, expert-led AI training for every role in your organization

The AI App and Agent Factory

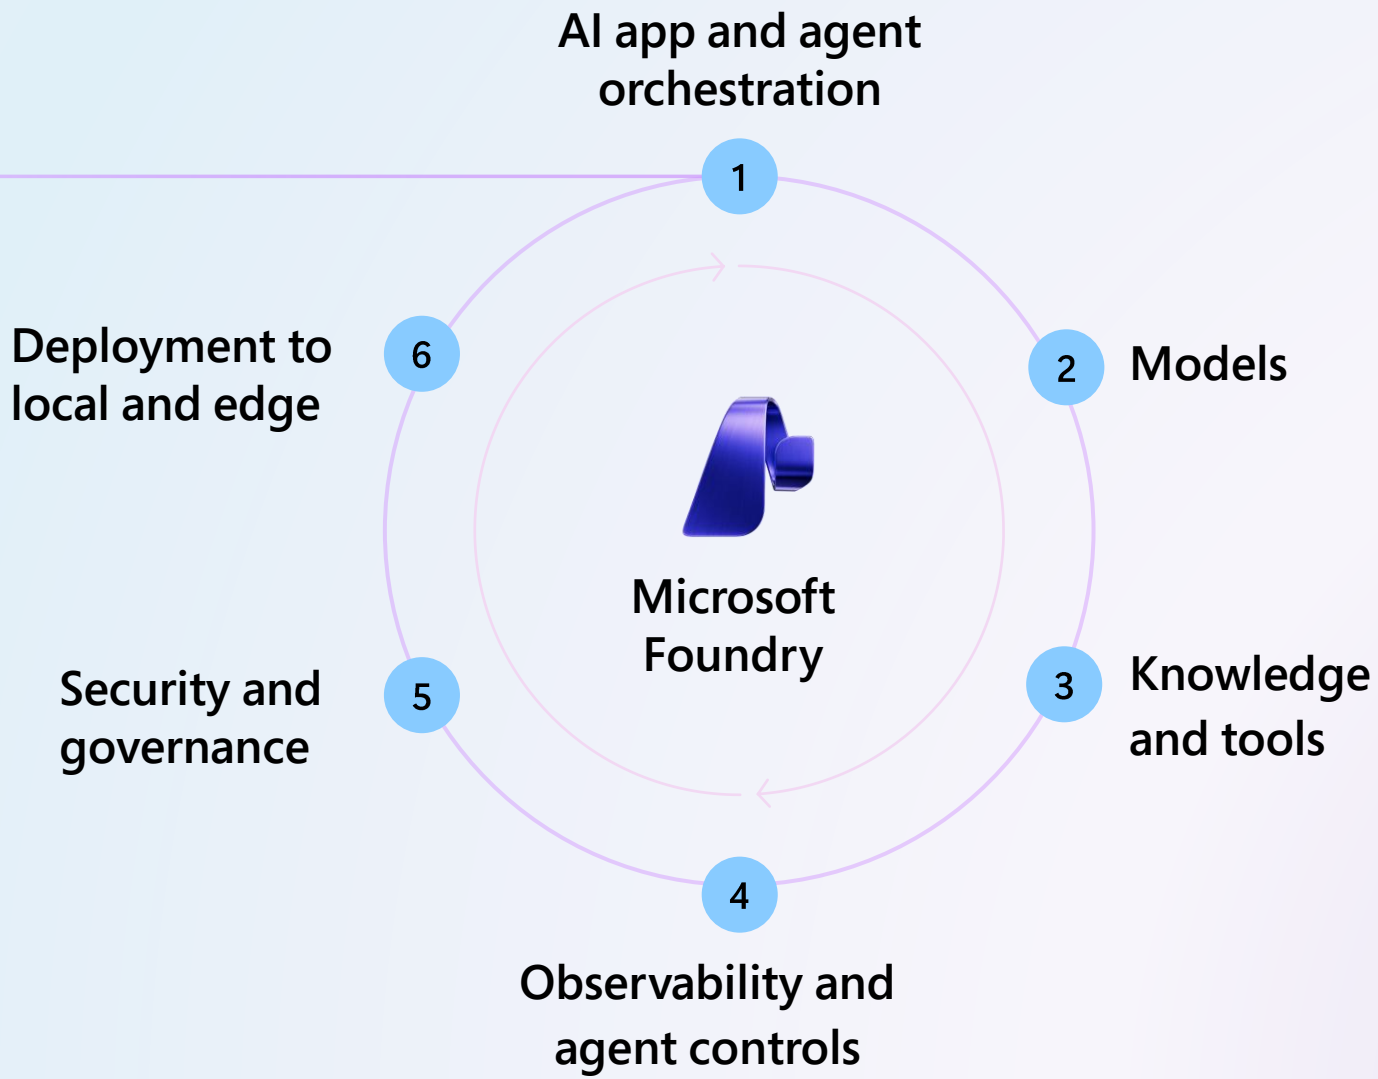

The AI App and Agent Factory

AI app and agent orchestration

1

- Declarative agents
- Hosted agents
- Workflow agents
- Built-in memory
- Agent controls
- Channel integration
- Open frameworks and protocols

Models

2

- Azure OpenAI
- Anthropic Claude
- Meta Llama
- DeepSeek
- Mistral AI
- xAI Grok
- Black Forest Labs
- Microsoft Phi

Knowledge and Tools

3

- SharePoint
- Bing
- OneLake
- Cosmos DB
- PostregreSQL
- Blob Storage
- Azure Functions
- MCP server

+ 1,400+ connectors to biz applications

Observability and controls

4

- Tracing/logging
- Evaluations
- Monitoring
- Quality, safety, security
- Cost & Throughput
- Compliance

Security and governance

5

- VNET
- BYO-resources
- AI Gateway
- Entra Agent ID
- AI red teaming
- Access controls
- Data classification
- Policies
- Compliance

Deploy anywhere

5

- Foundry Local
- On-prem
- Multi-cloud
- Microsoft 365 Copilot
- Teams
- Agent 365

Evolving into Multi-Agent Systems

Agent Orchestration

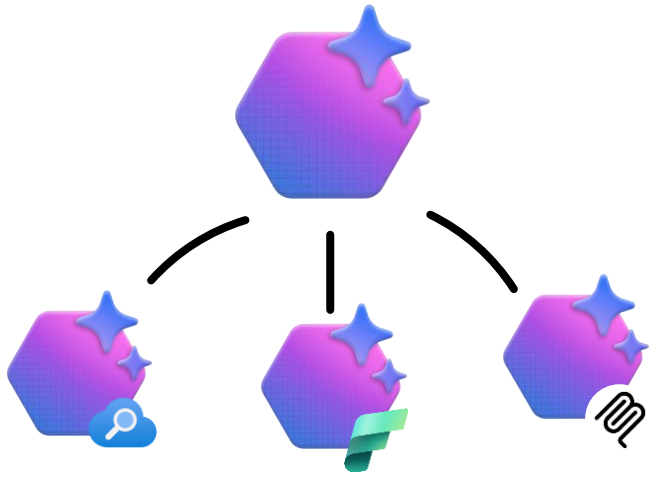

LLMs dynamically decide the control flow and direct their own processes and tool usage, maintaining control over how they accomplish tasks.

Workflow Orchestration

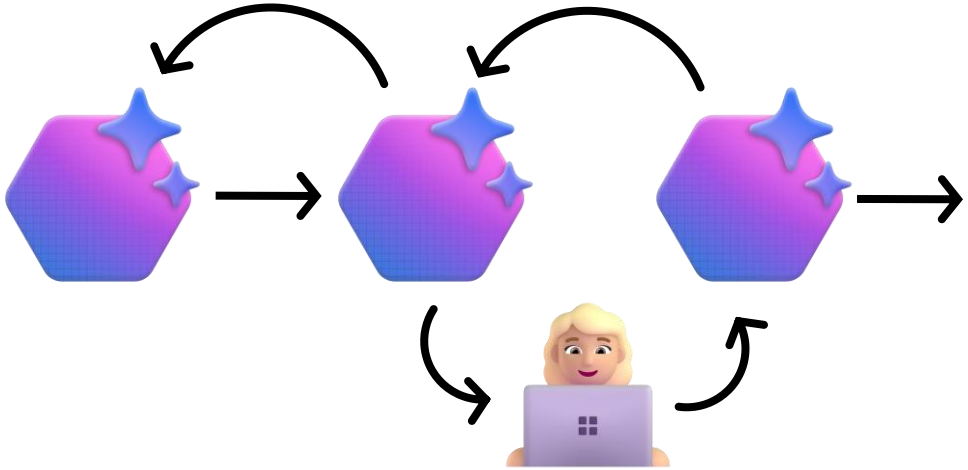

You decide the control flow orchestrating LLMs, Agents, tools and more via predefined code paths.

Typical Multi-Agent Architectures

Sequential

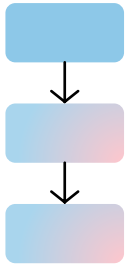

Concurrent

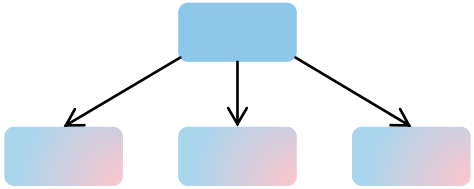

Handoff

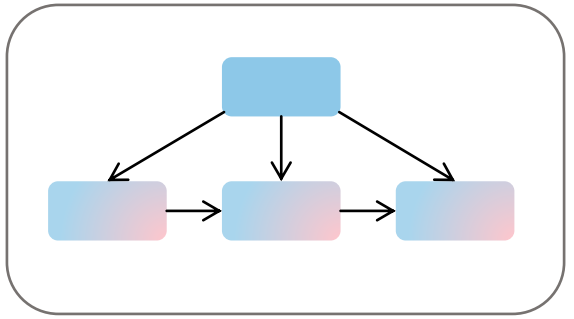

Group Chat

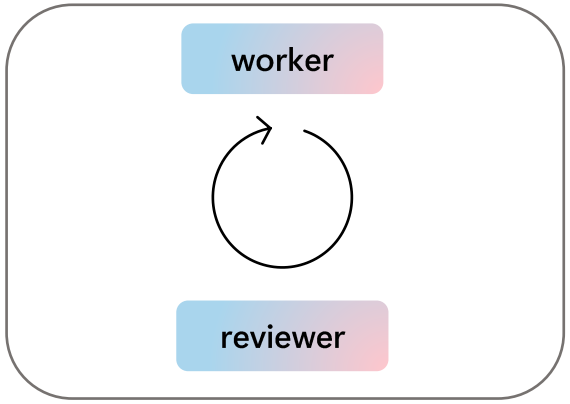

Magentic

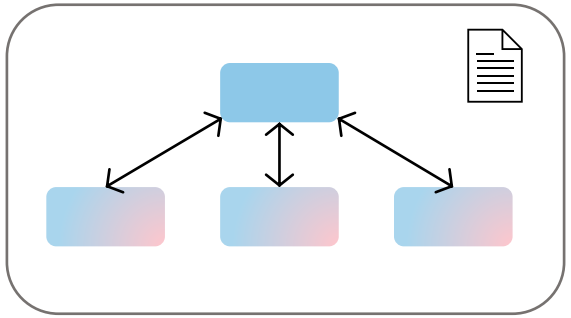

Workflow Process

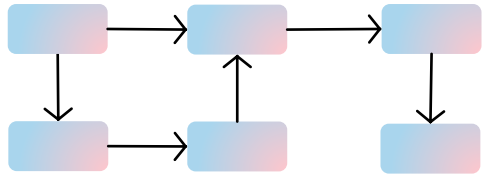

Multi-Agent Orchestration Patterns

Sequential

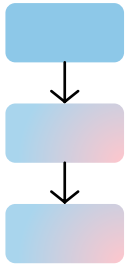

Concurrent

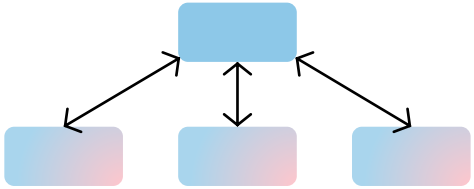

Handoff

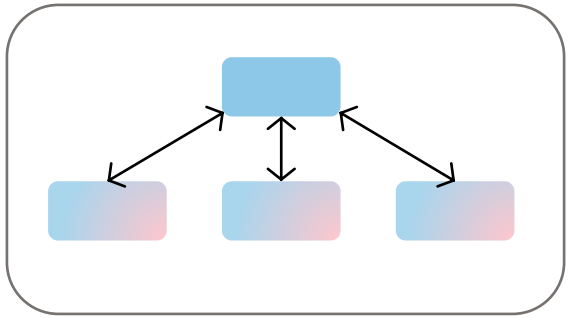

Group Chat

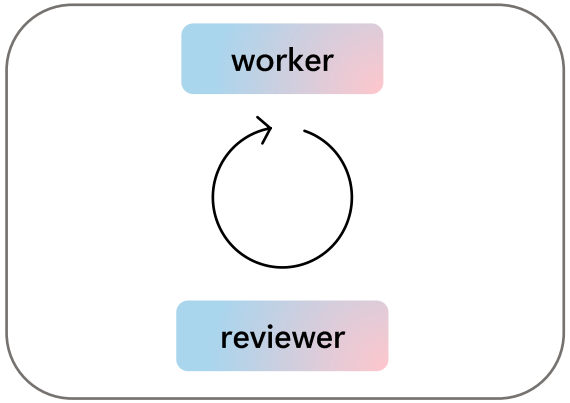

Magentic

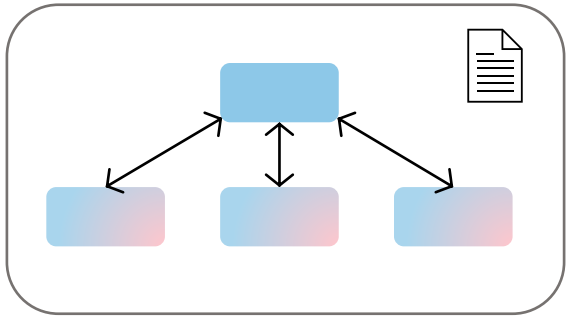

Workflow Process

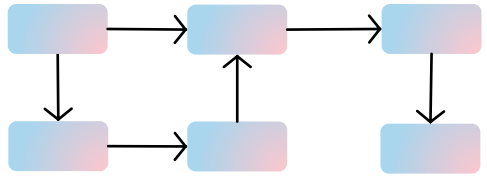

Sequential Pattern

- The sequential orchestration pattern chains AI agents in a predefined, linear order. Each agent processes the output from the previous agent in the sequence, which creates a pipeline of specialized transformations.
- The sequential orchestration pattern solves problems that require step-by-step processing, where each stage builds on the previous stage.

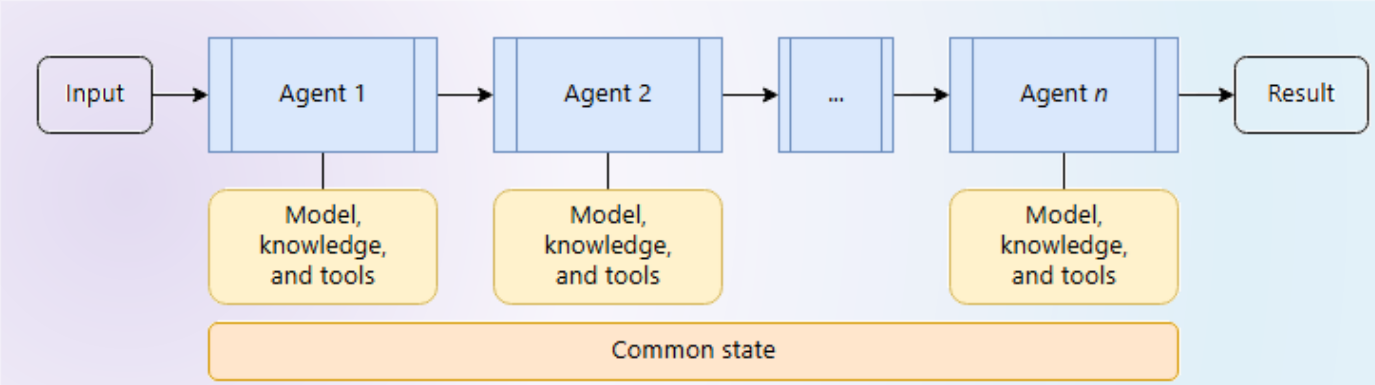

Concurrent Pattern

- The concurrent orchestration pattern runs multiple AI agents simultaneously on the same task. This approach allows each agent to provide independent analysis or processing from its unique perspective or specialization.

- This pattern addresses scenarios where you need diverse insights or approaches to the same problem.

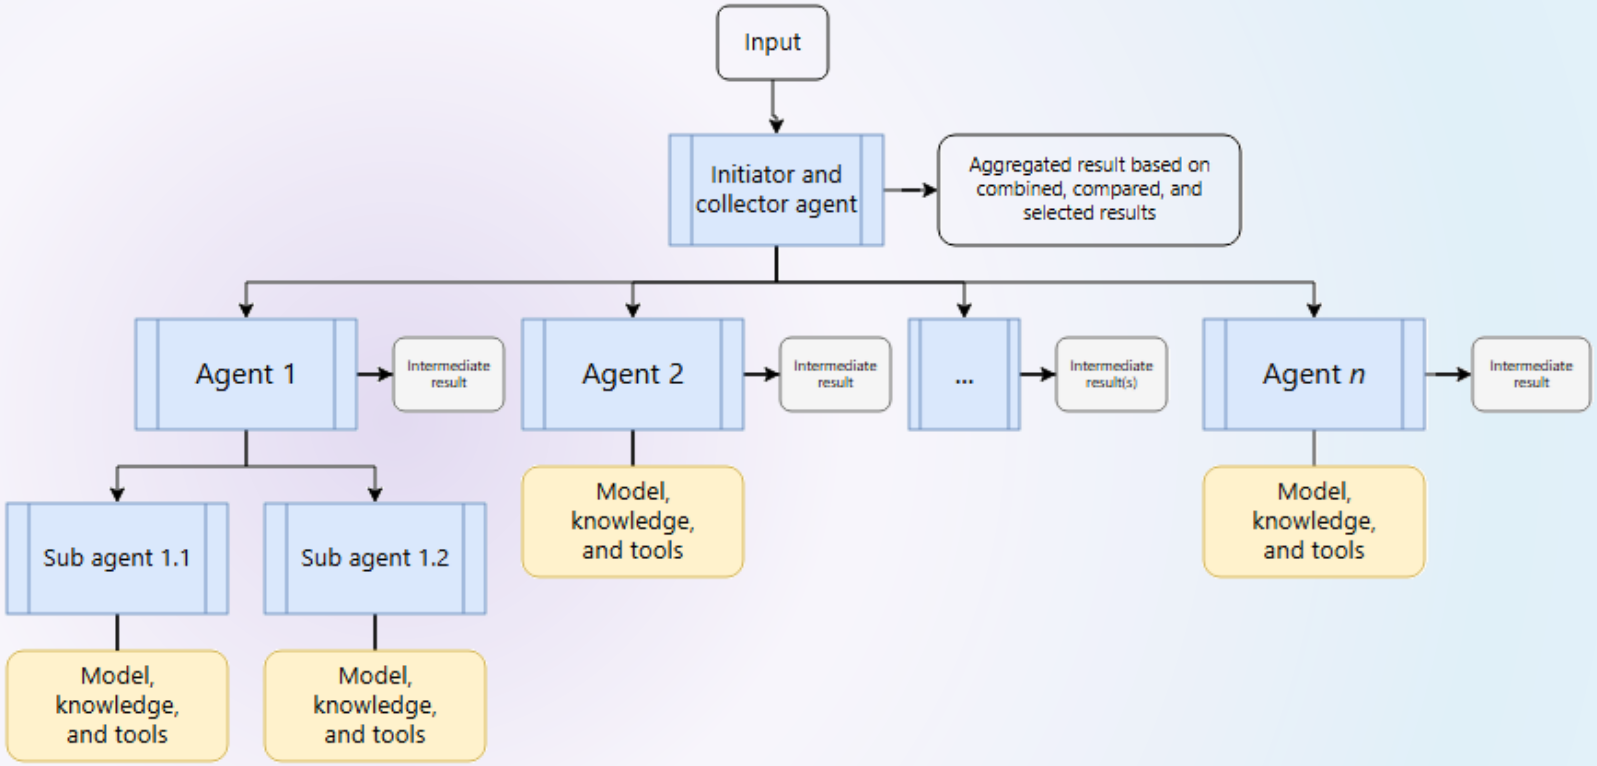

Group Chat Pattern

- The group chat orchestration pattern enables multiple agents to solve problems, make decisions, or validate work by participating in a shared conversation thread where they collaborate through discussion.

- This pattern addresses scenarios that are best accomplished through group discussion to reach decisions.

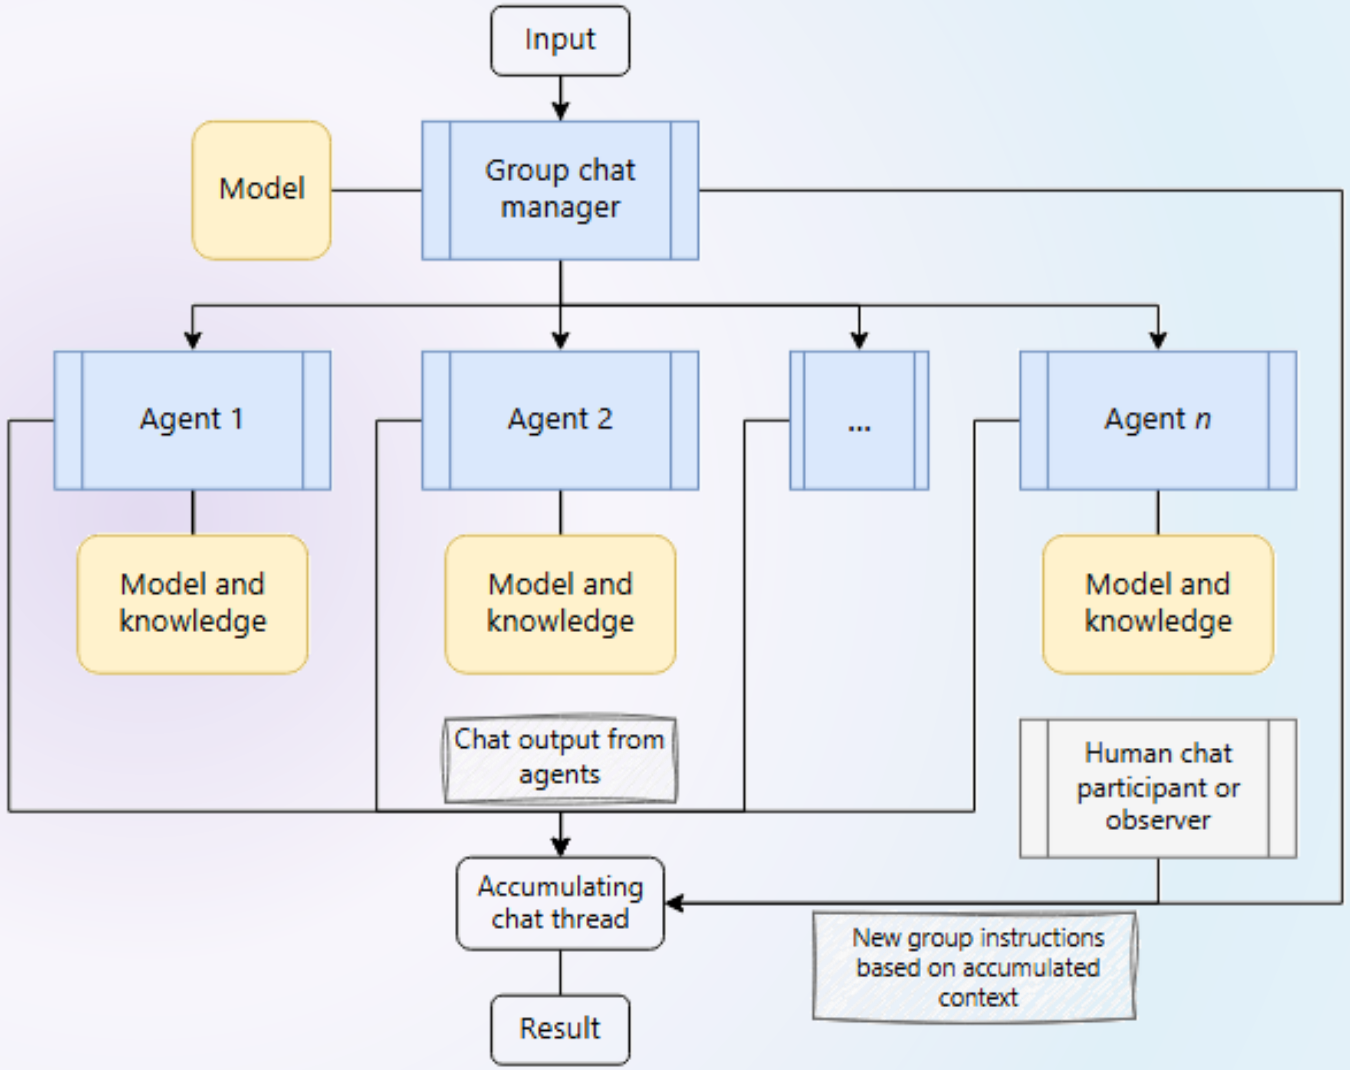

Handoff Pattern

- The handoff orchestration pattern enables dynamic delegation of tasks between specialized agents. Each agent can assess the task at hand and decide whether to handle it directly or transfer it to a more appropriate agent based on the context and requirements.

- This pattern addresses scenarios where the optimal agent for a task isn't known upfront or where the task requirements become clear only during processing.

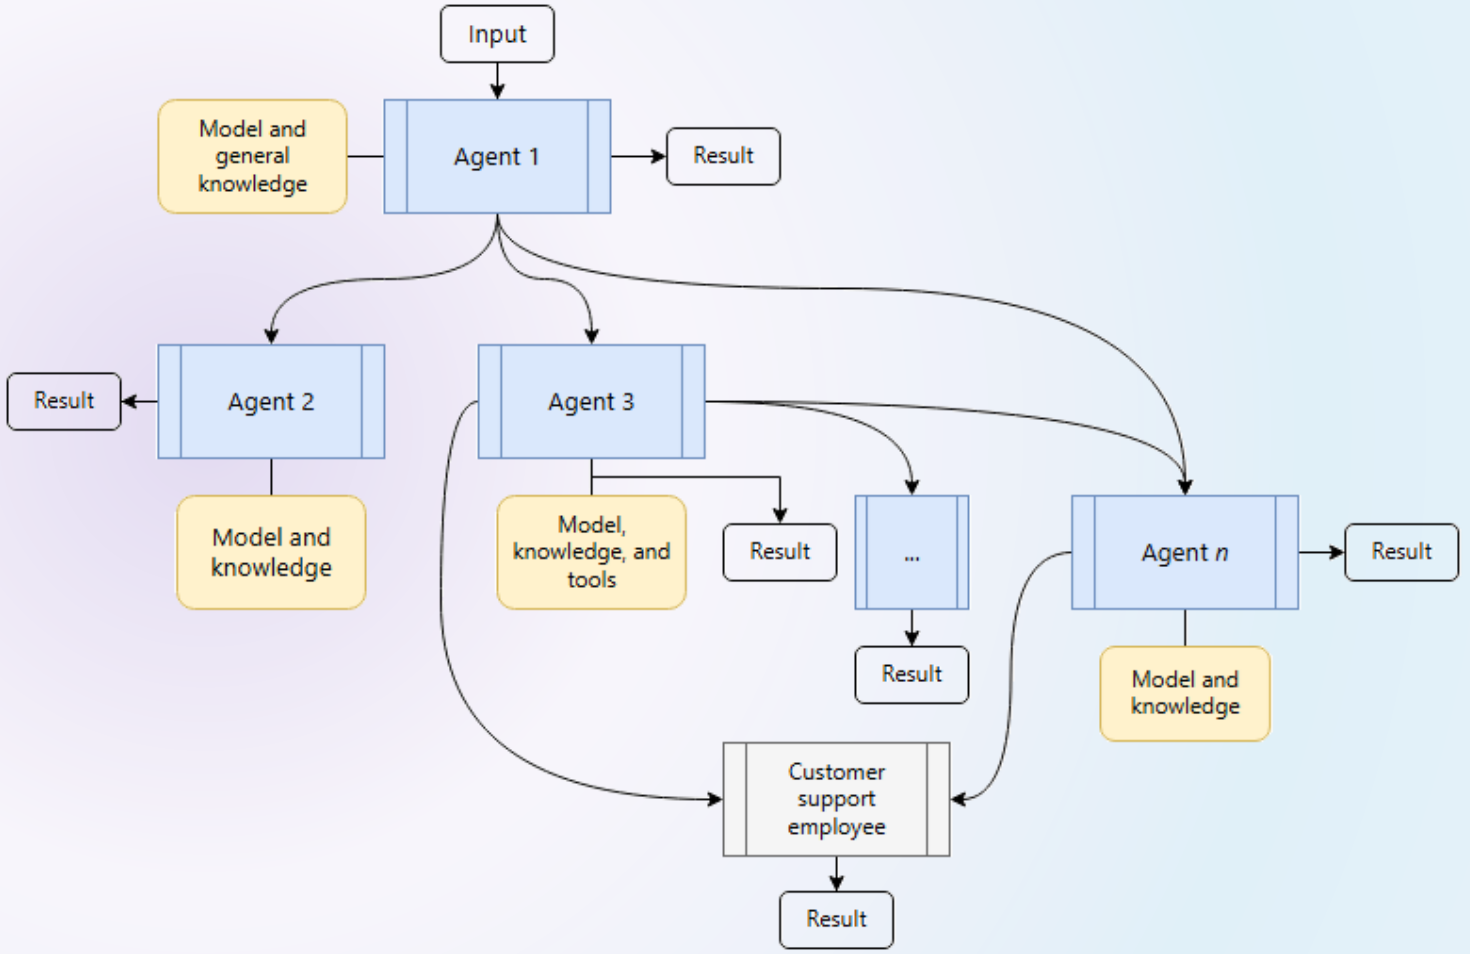

Magnetic Pattern

- The magentic orchestration pattern is designed for open-ended and complex problems that don't have a predetermined plan of approach.
- The manager agent communicates directly with specialized agents to gather information as it builds and refines the task ledger

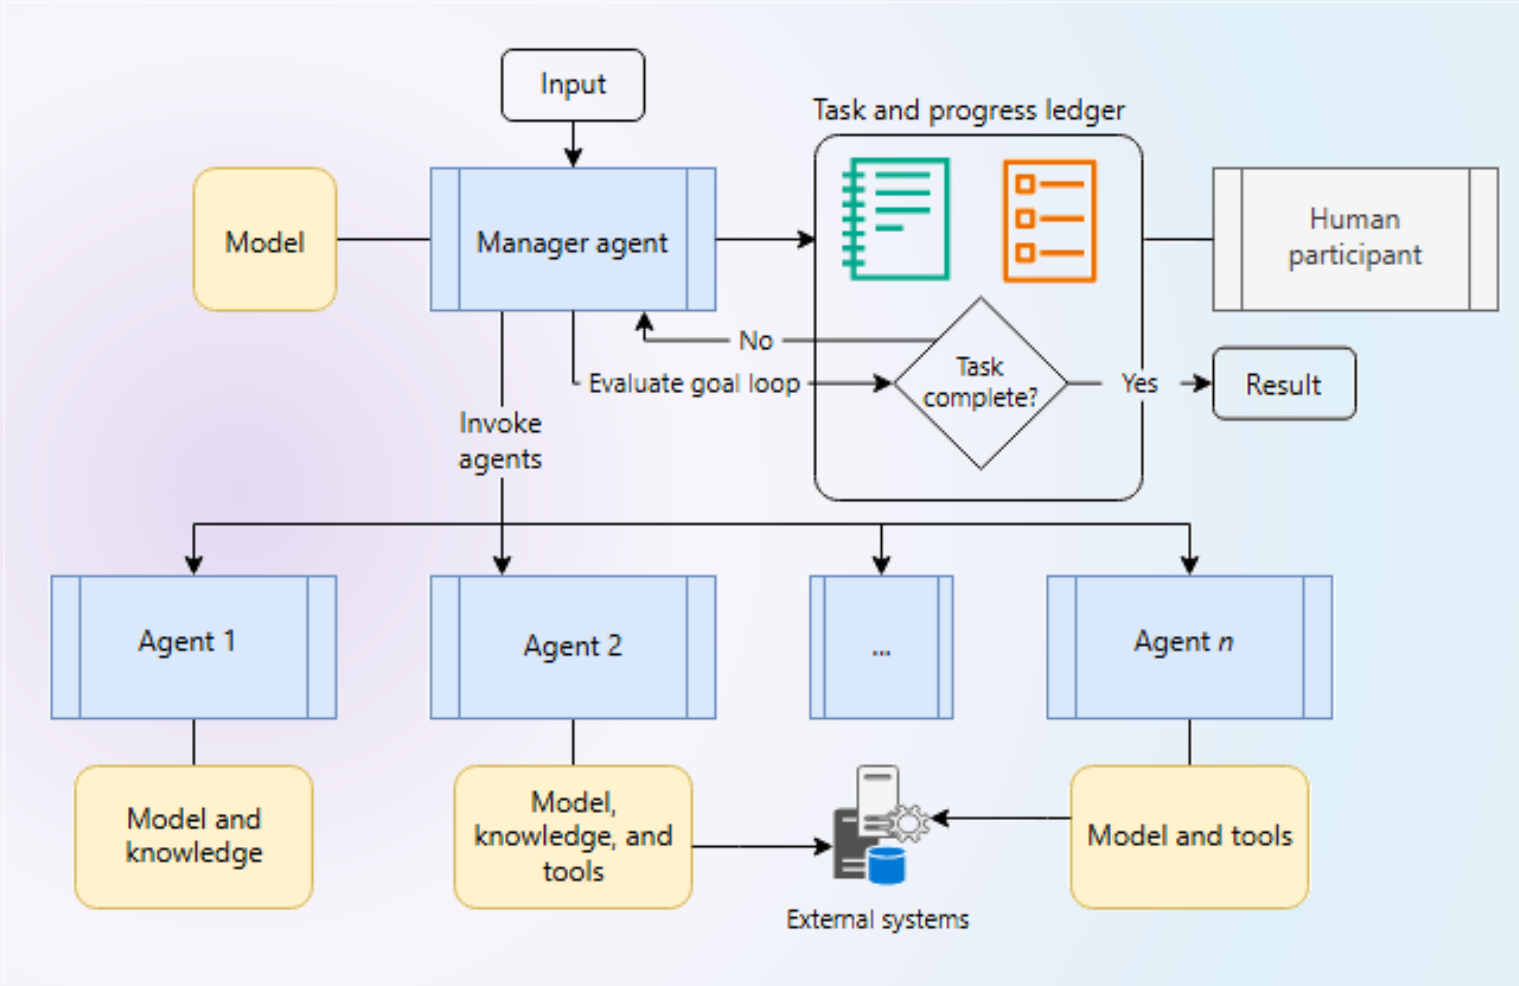

Multi-agent workflows

- Empower developers to build **multi-step, cross-agent logic** with **explicit transitions**, state control, and structured data sharing
- Offer **server-managed persistence**, enabling agents to:
  - Resume after a crash or scale event
  - Run for **minutes, hours, or days**
  - Track history, retries, and progress
- Support **rule-based or LLM-assisted transitions**, offering **traceable, testable, and governable flows**

Connected

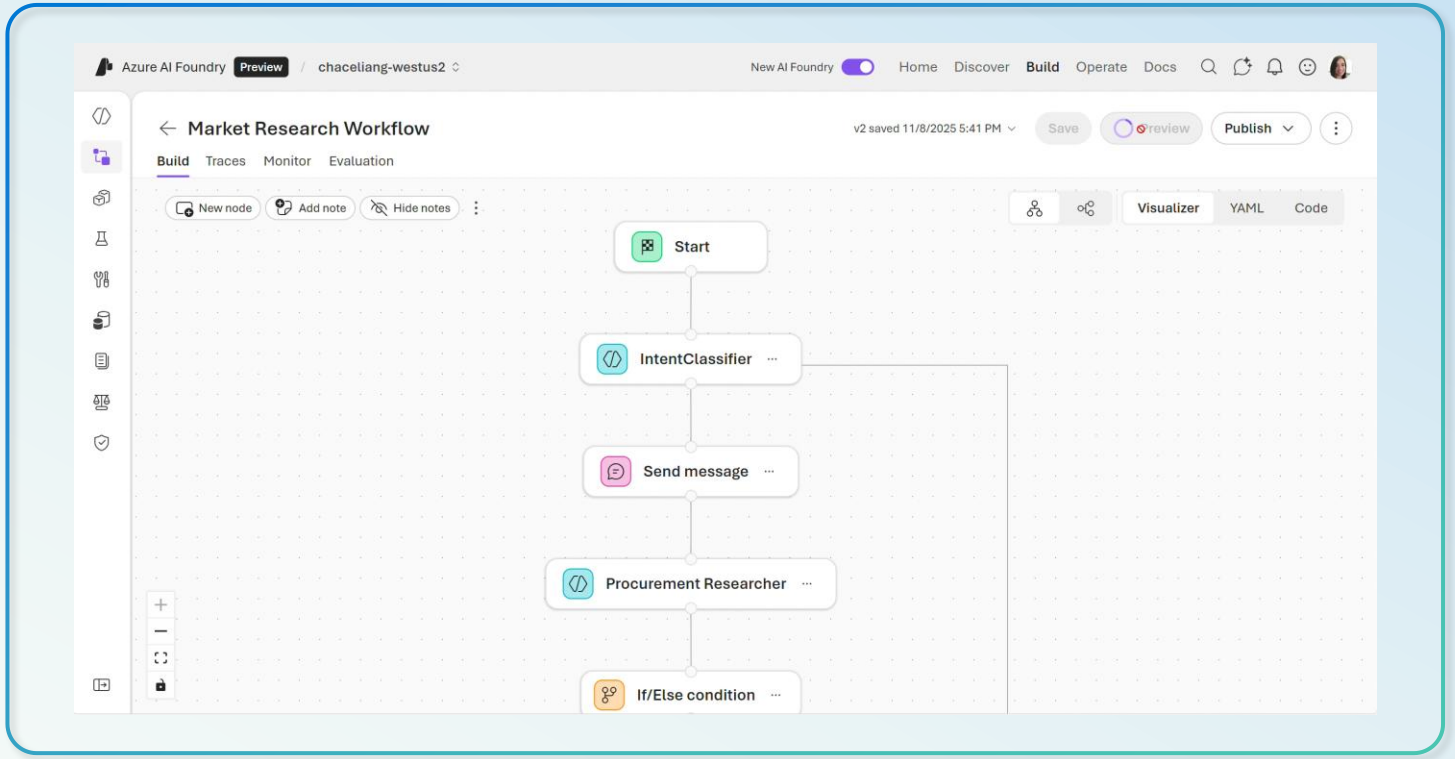

Workflow

- Express multi-step reasoning and tool use as structured workflows
- Combine LLM reasoning + deterministic logic + API calls
- Support for looping, branching, fan-in, fan-out

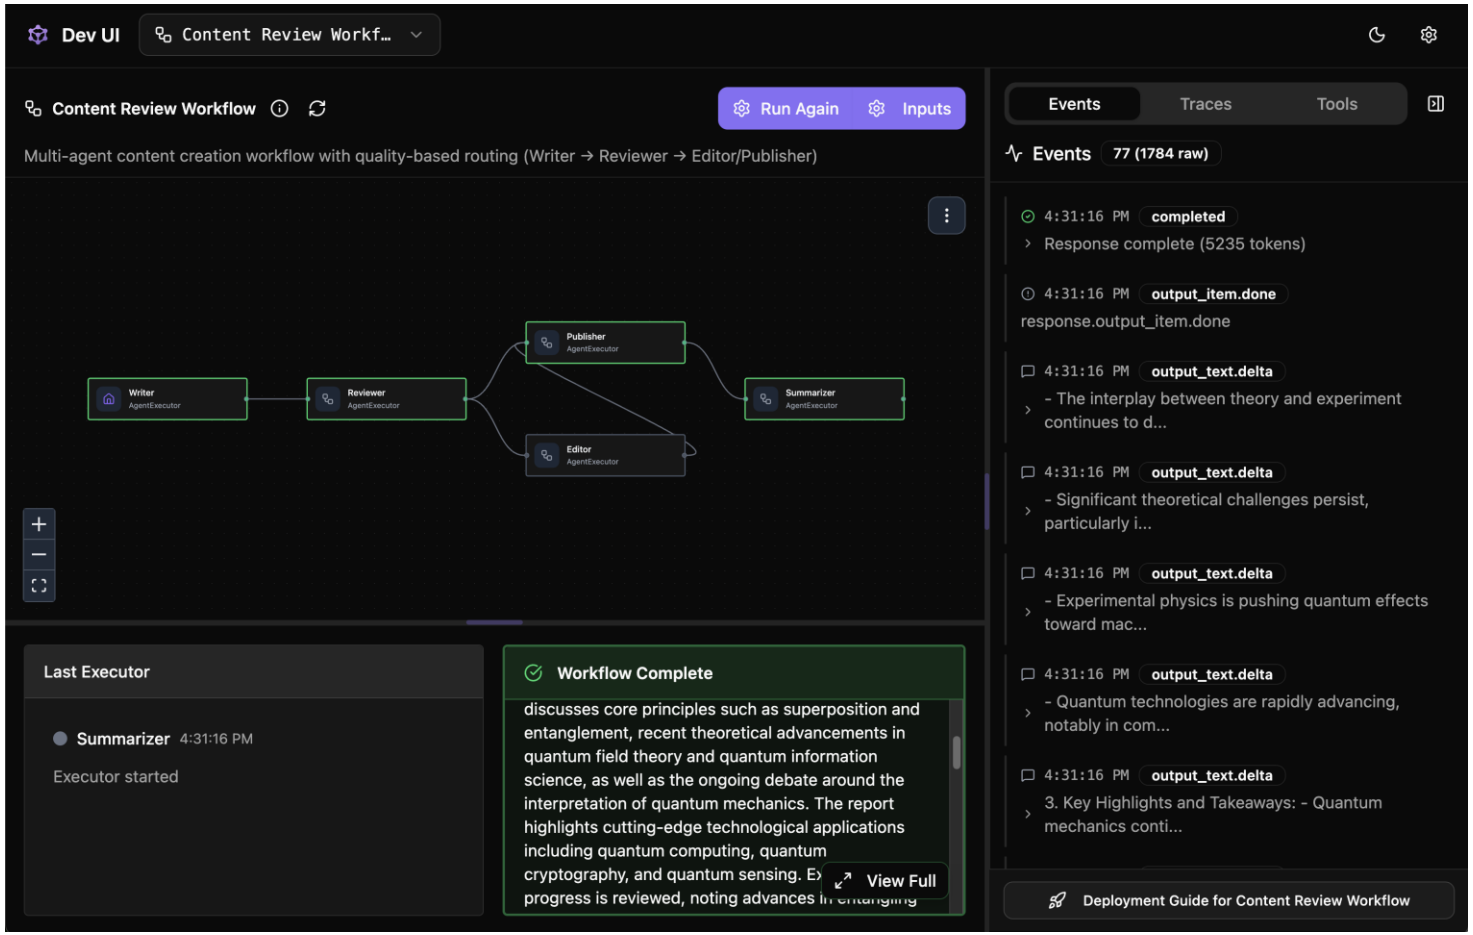

Powering Agents with

Fabric
Real-Time Intelligence

Microsoft Fabric

The unified data platform for AI transformation

Data Factory

Analytics

Databases

Real-Time Intelligence

IQ

Power BI

Fabric Platform

AI in Fabric

OneLake

Governance

Components of Fabric's Real-Time Intelligence

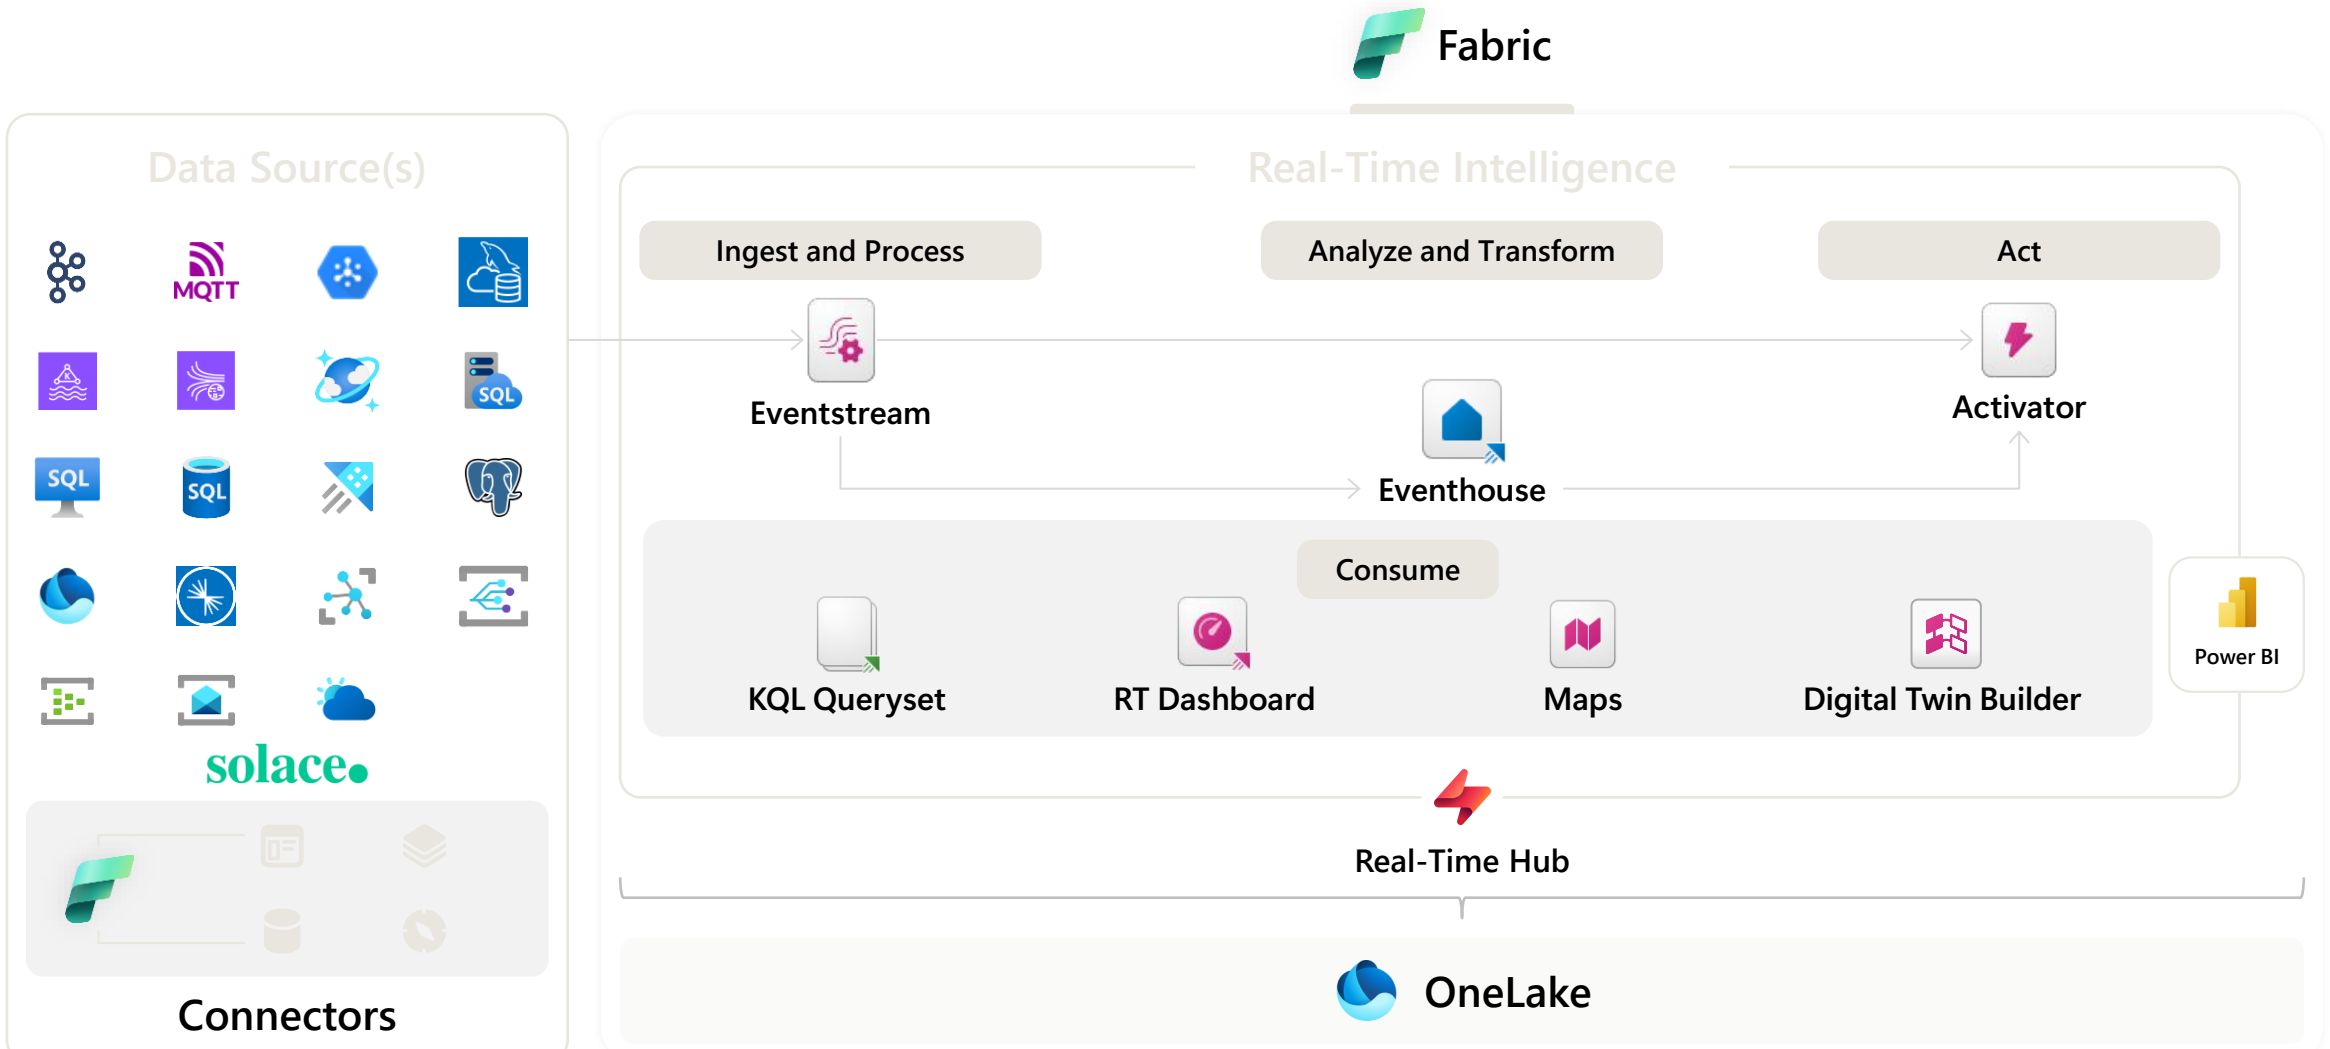

Data agents deliver impactful insights

Fabric data agents leverage business context and understanding from IQ to drive more accurate responses without needing extensive prompting.

Key Capabilities:

- Ground natural language queries in the shared definitions of business entities, attributes, relationships, and metrics
- Formulate precise queries across tabular, graph, geospatial, and digital twin models
- Define custom business semantics and grounding unique to your organization
- Support multiple data sources (lakehouse and warehouse tables, mirrored DB and shortcut data, Semantic models, Eventhouse KQL DB)
- Data agents can be consumed inside and outside of Fabric, including Foundry Agent Service and Copilot Studio

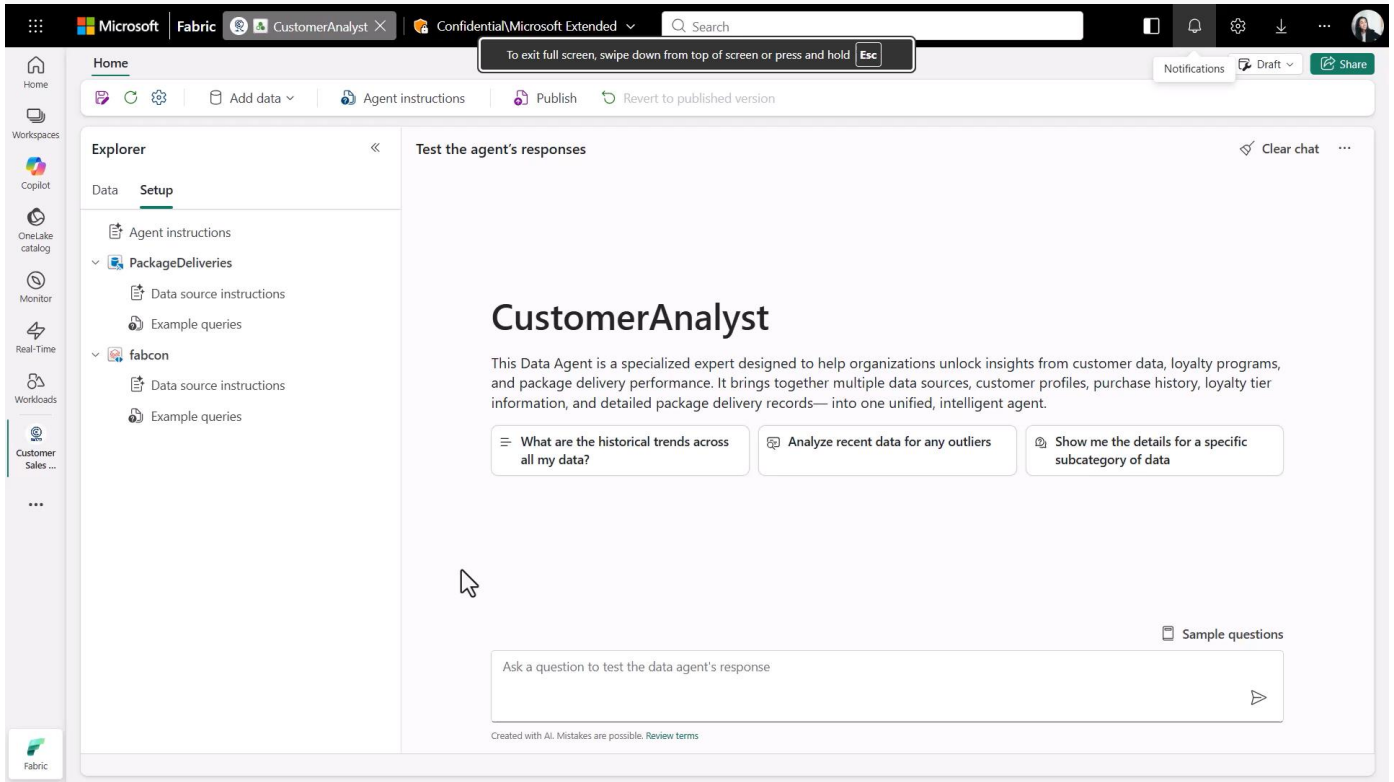

An end-to-end Real-Time Intelligence experience

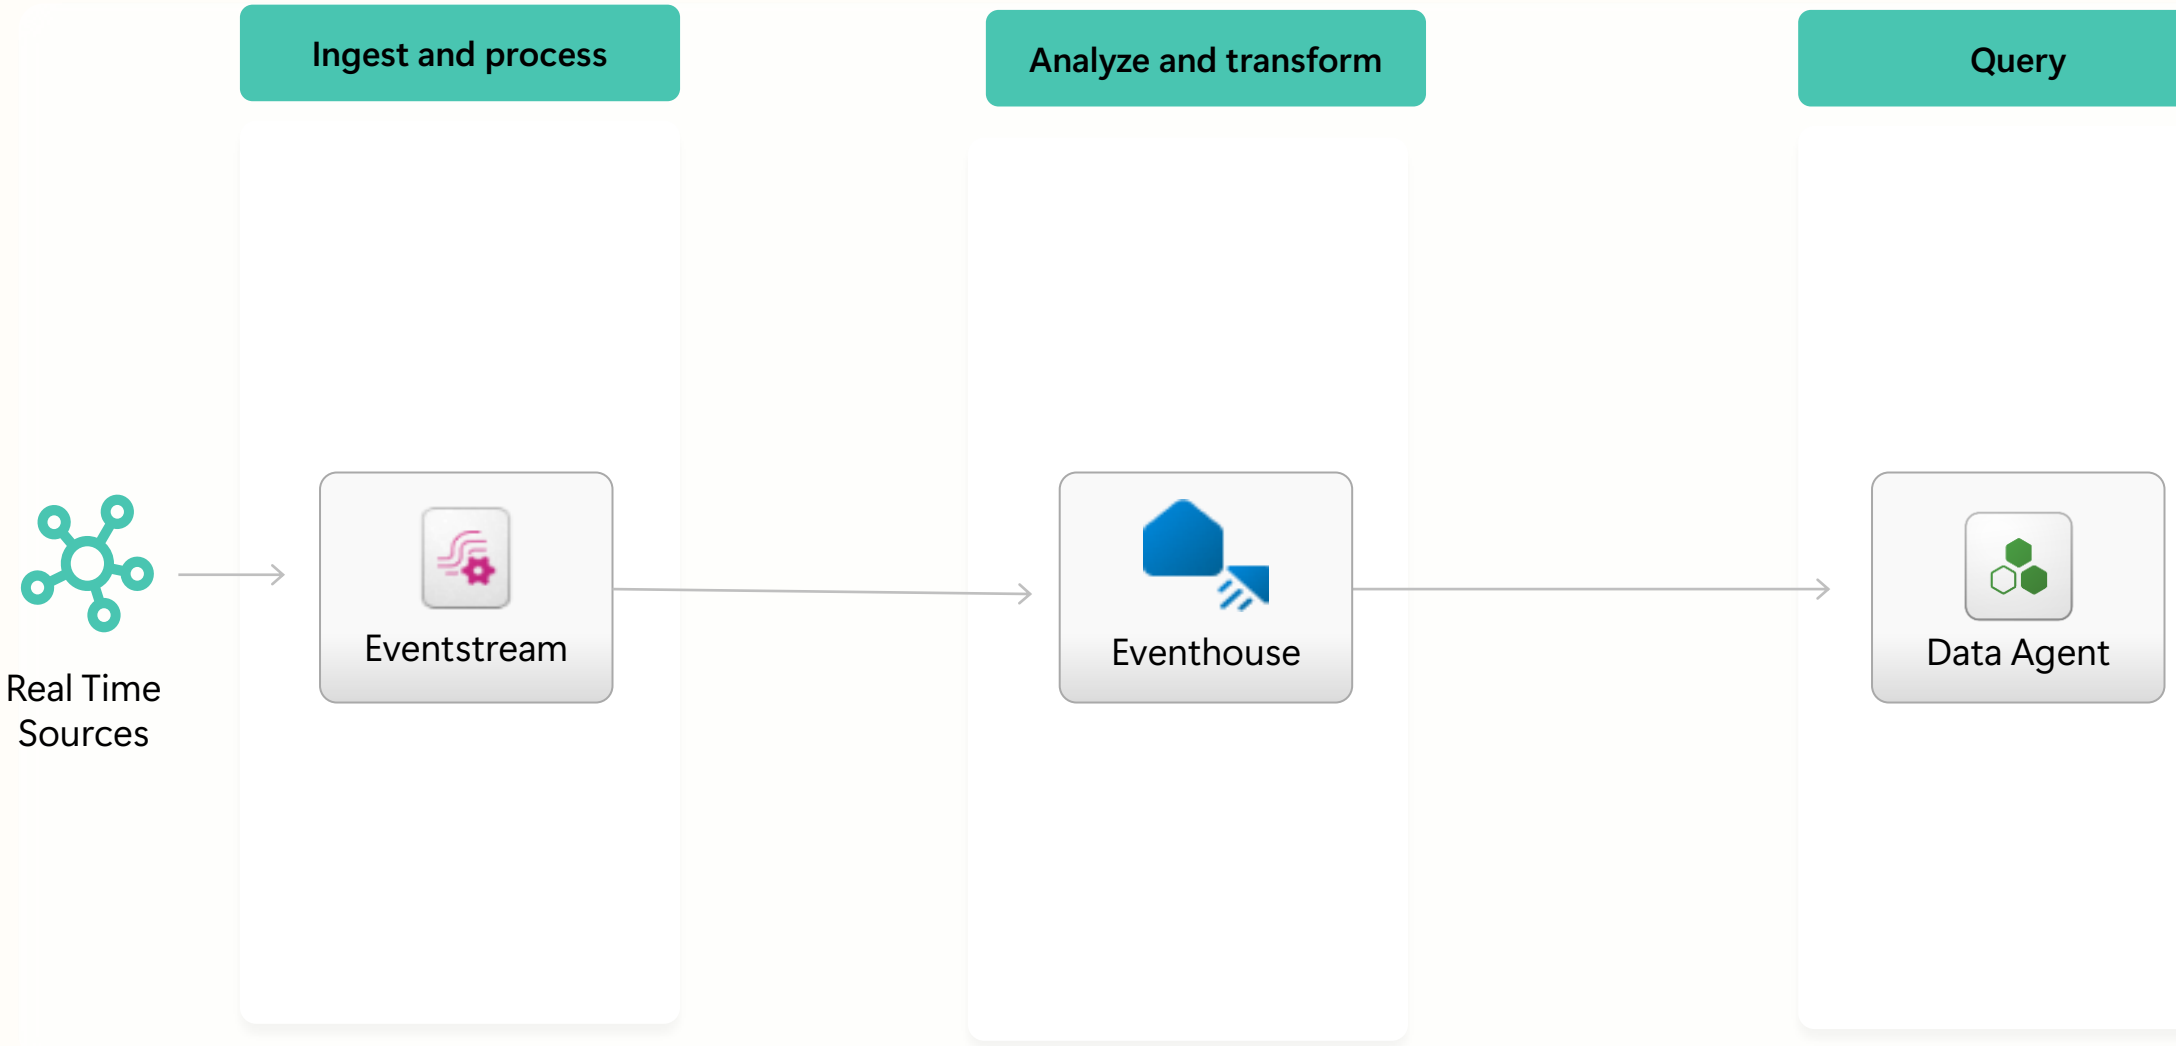

Fabric data agent integration with Microsoft Foundry

Integrate Foundry agents with Fabric data agents to unlock powerful data analysis capabilities. Fabric data agents transform enterprise data into conversational Q&A systems, allowing users to interact with data through chat and uncover actionable insights

Key Capabilities:

- Agentic workflows connect with secure enterprise data for grounding using Microsoft knowledge sources, such as data agents
- Define custom business semantics and grounding unique to your org
- Enterprise action tools enable your agent to automate tasks across apps and APIs

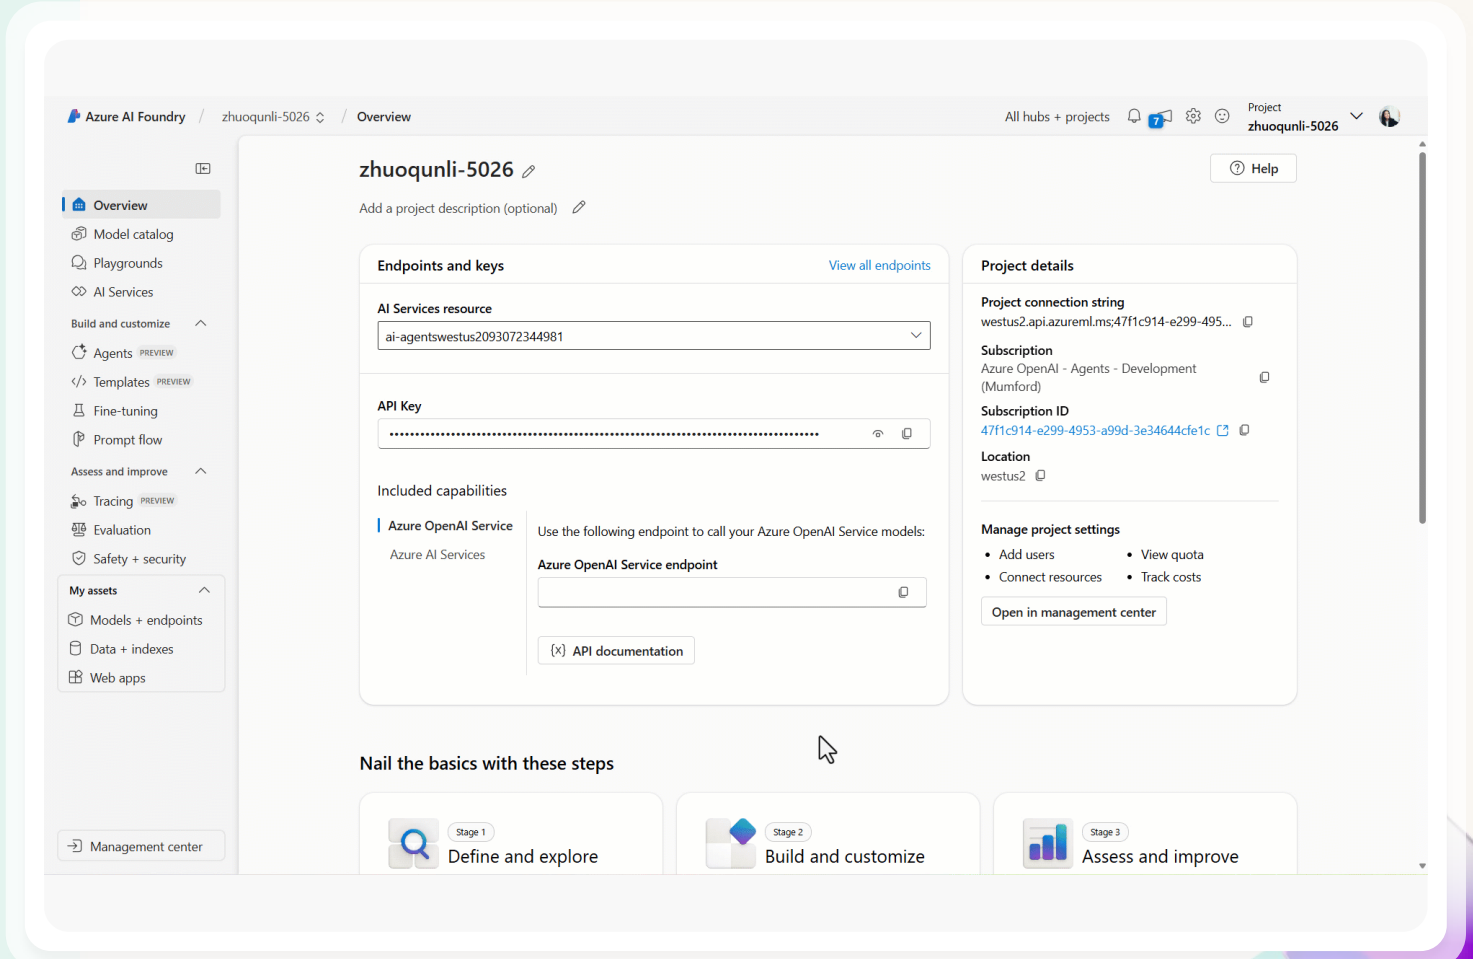

Why innovate with
AI apps and agents on Azure

Empower

Empower developers to easily build AI apps and agents

Enable developers with **integrated AI-native tools**

Reduce developer ramp-up and unlock AI-driven innovation across teams

Accelerate

Build, optimize, and govern AI innovation at scale

Orchestrate models, agents, and AI apps through a **unified platform**

Get end-to-end visibility and centralized resource management across departments

Drive

Scale AI apps and agents seamlessly and efficiently

Streamline deployment and scaling of AI services with enterprise-grade databases and app services

Optimize resource utilization, run AI apps with high performance, and align AI scaling with ROI goals

Govern

Secure, govern, and ensure the safety of AI

Mitigate risk and ensure security and governance with powerful tools

Empower responsible AI practices with built-in security, governance, and observability

Creating powerful, trustworthy, secure AI apps and agents

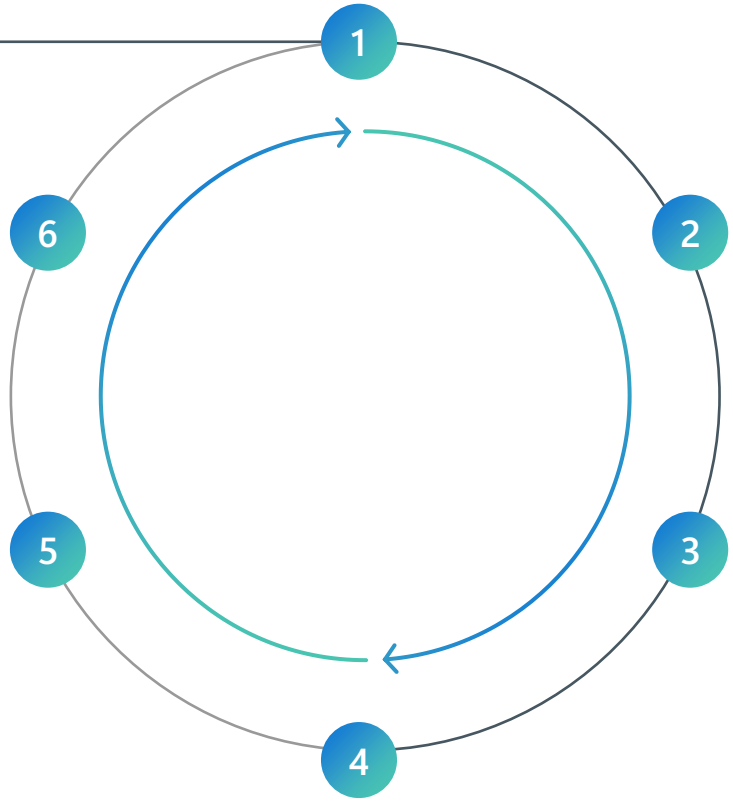

Demo Architecture

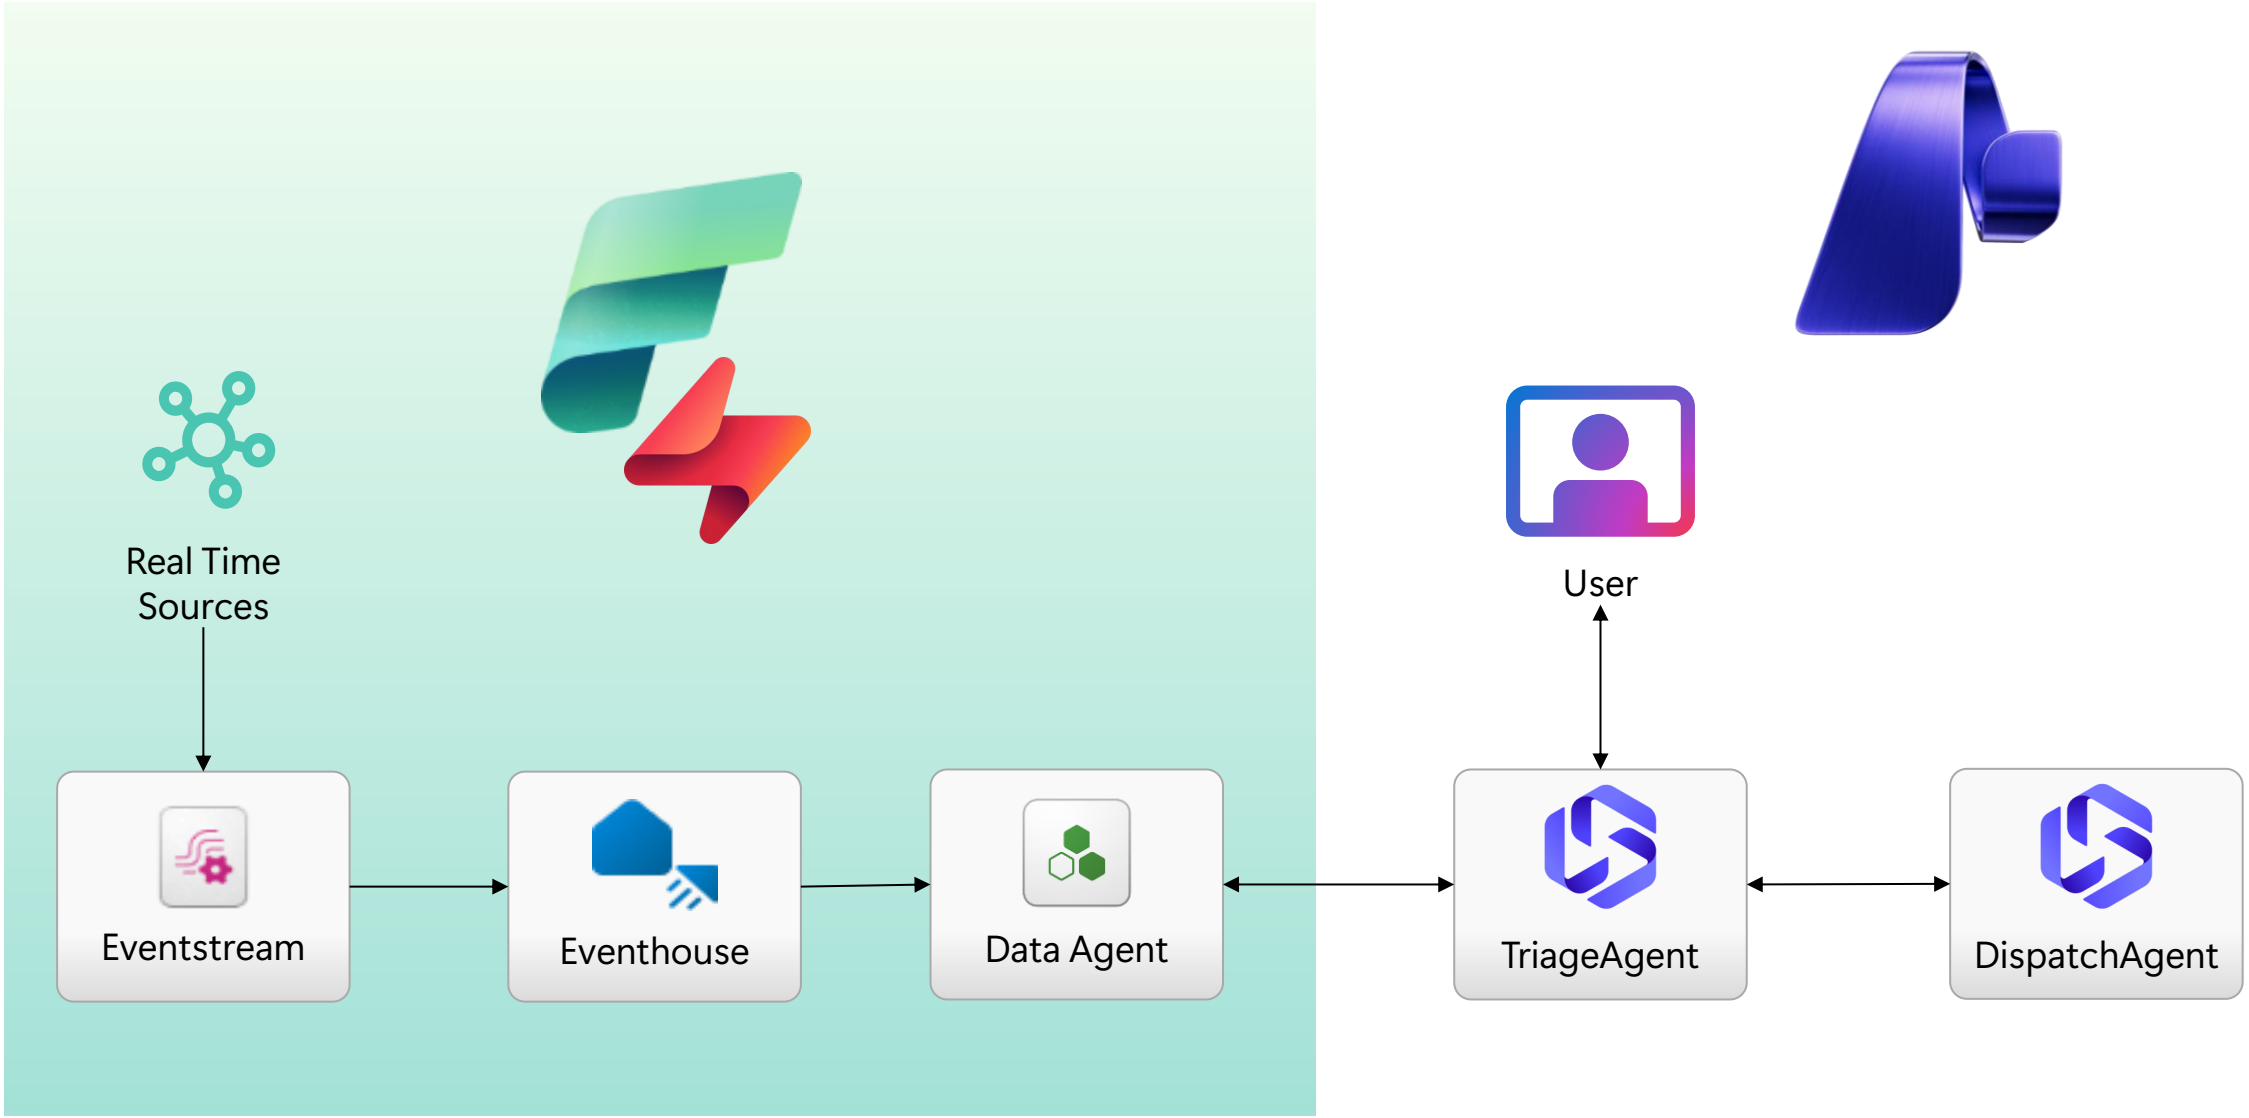

Microsoft Azure

Demo

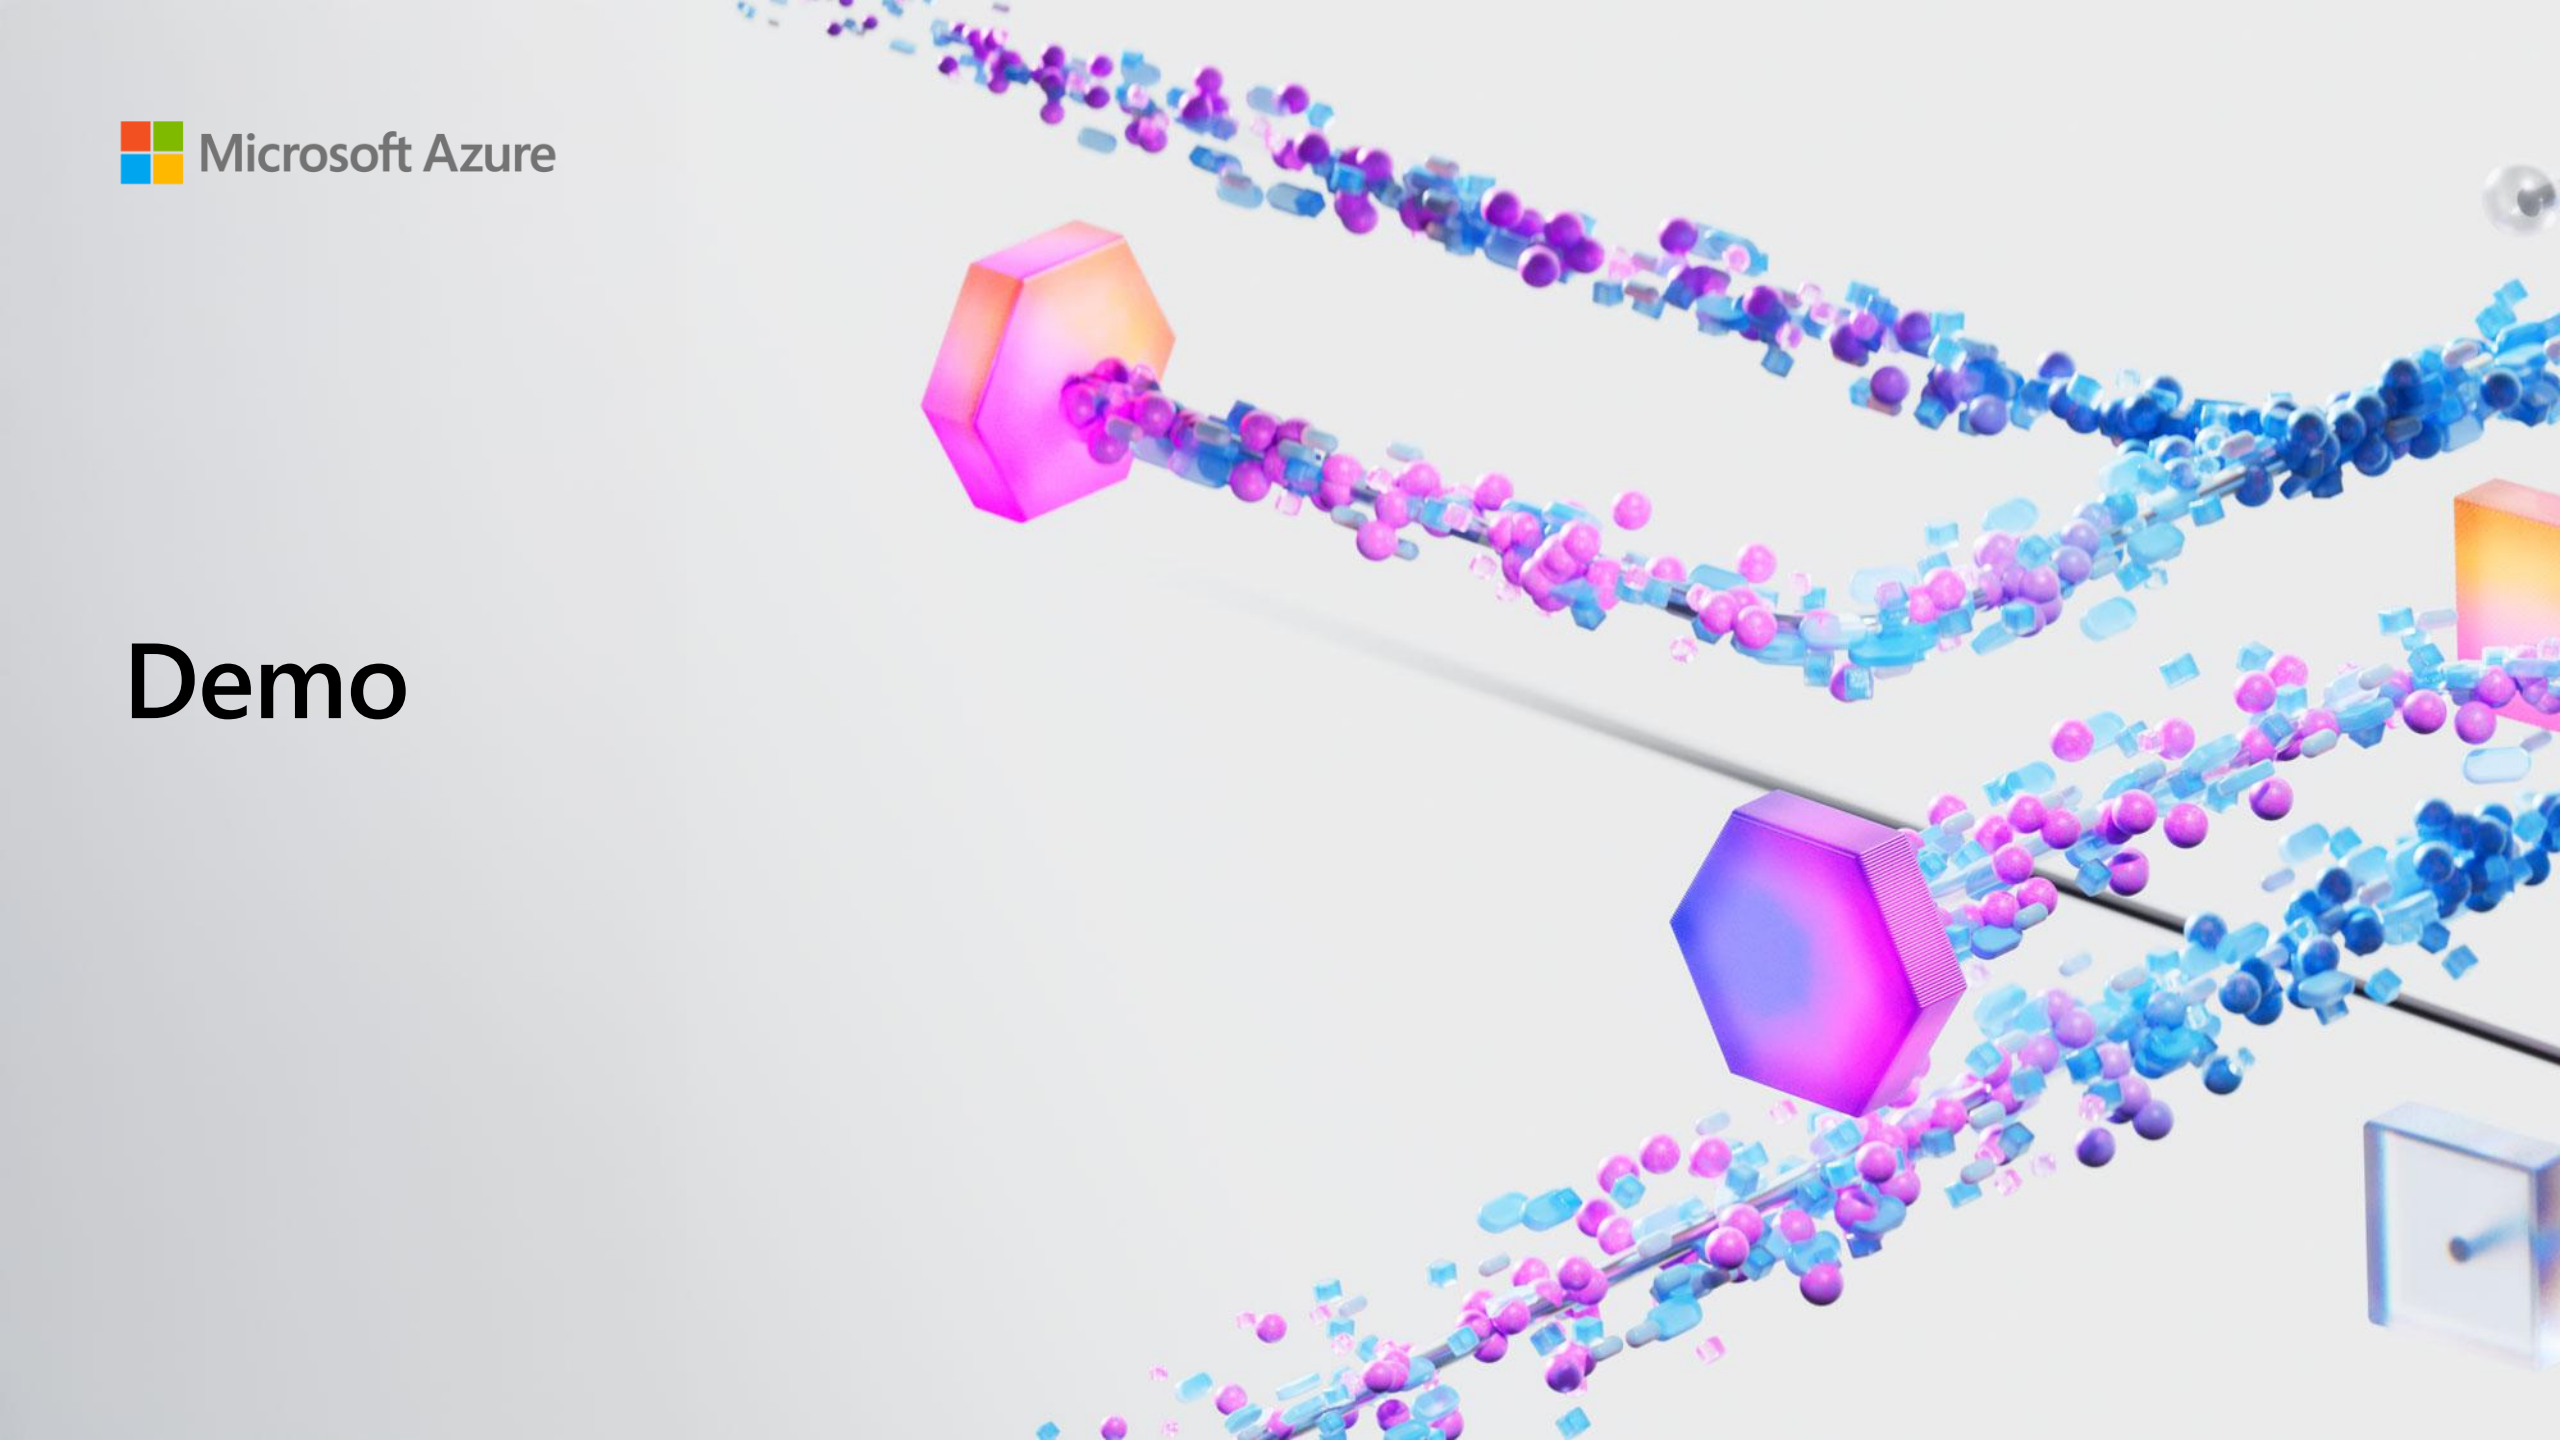

Microsoft

Thank you

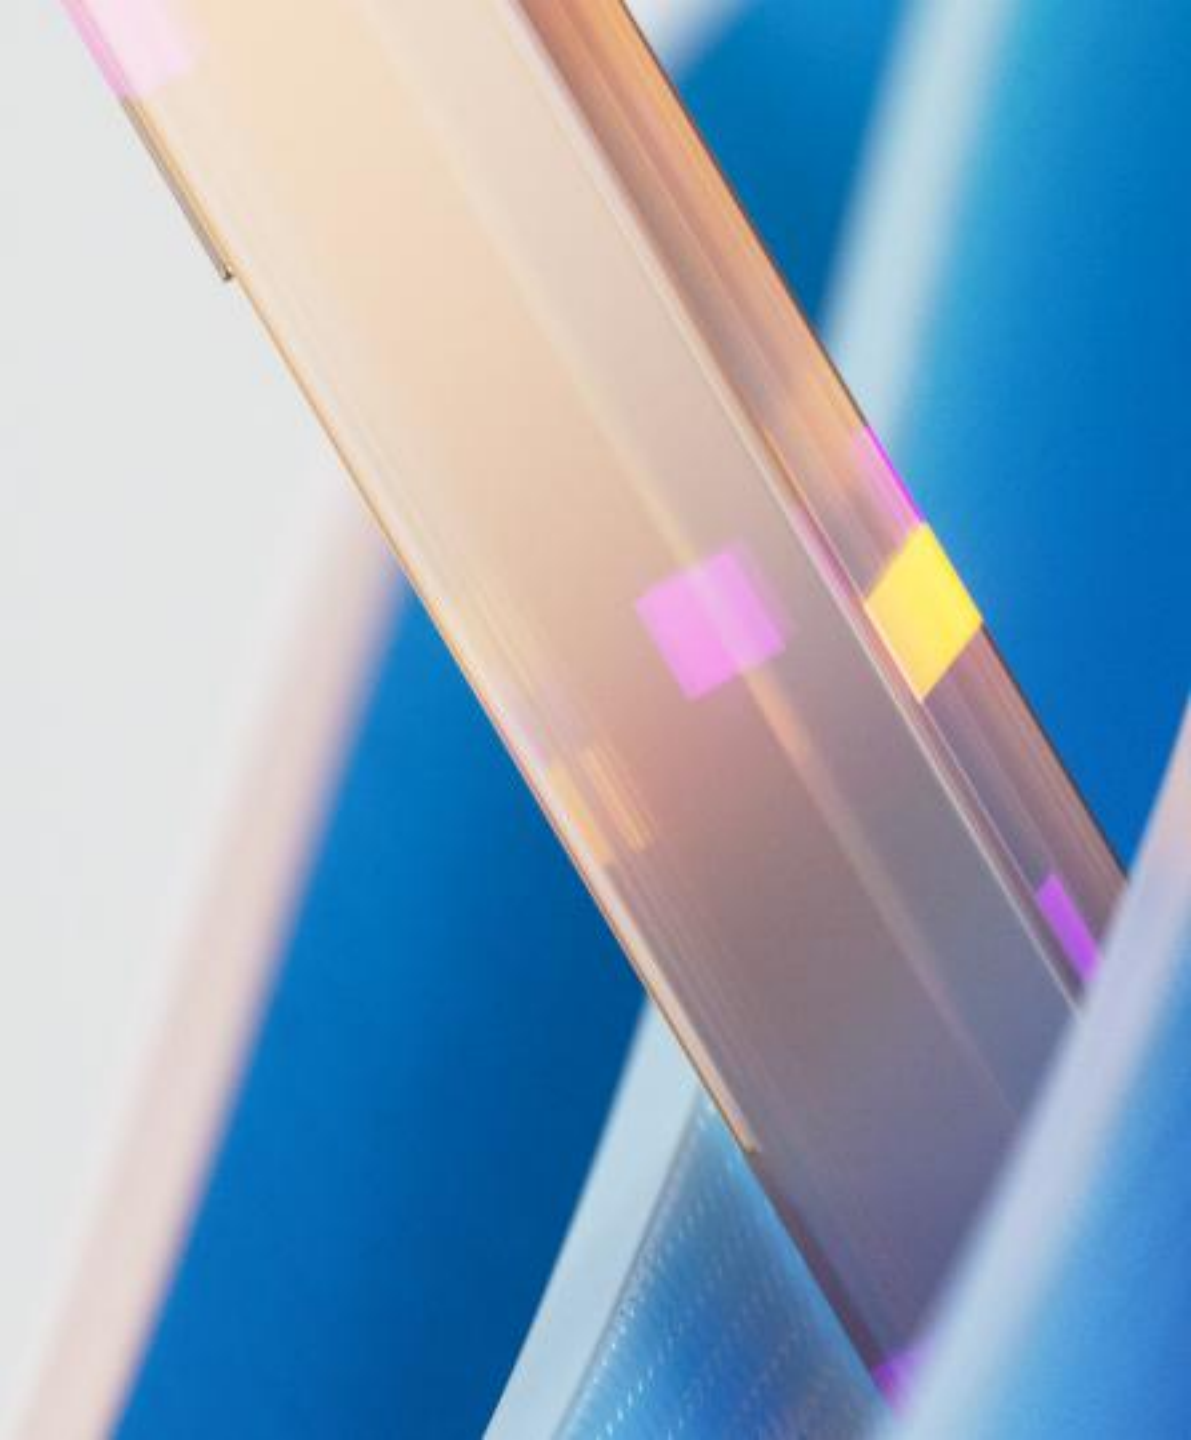

Screenshot slide

- Text here

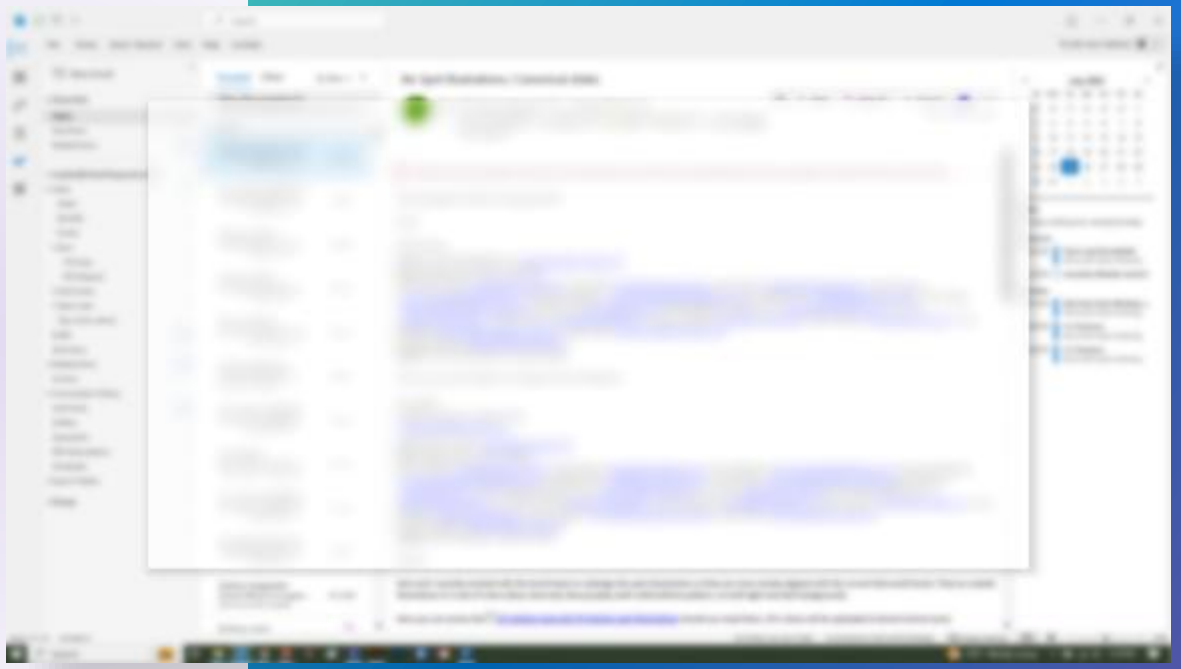

Photo slide

- Text here

Photo slide

- Text here

Microsoft

Thank you

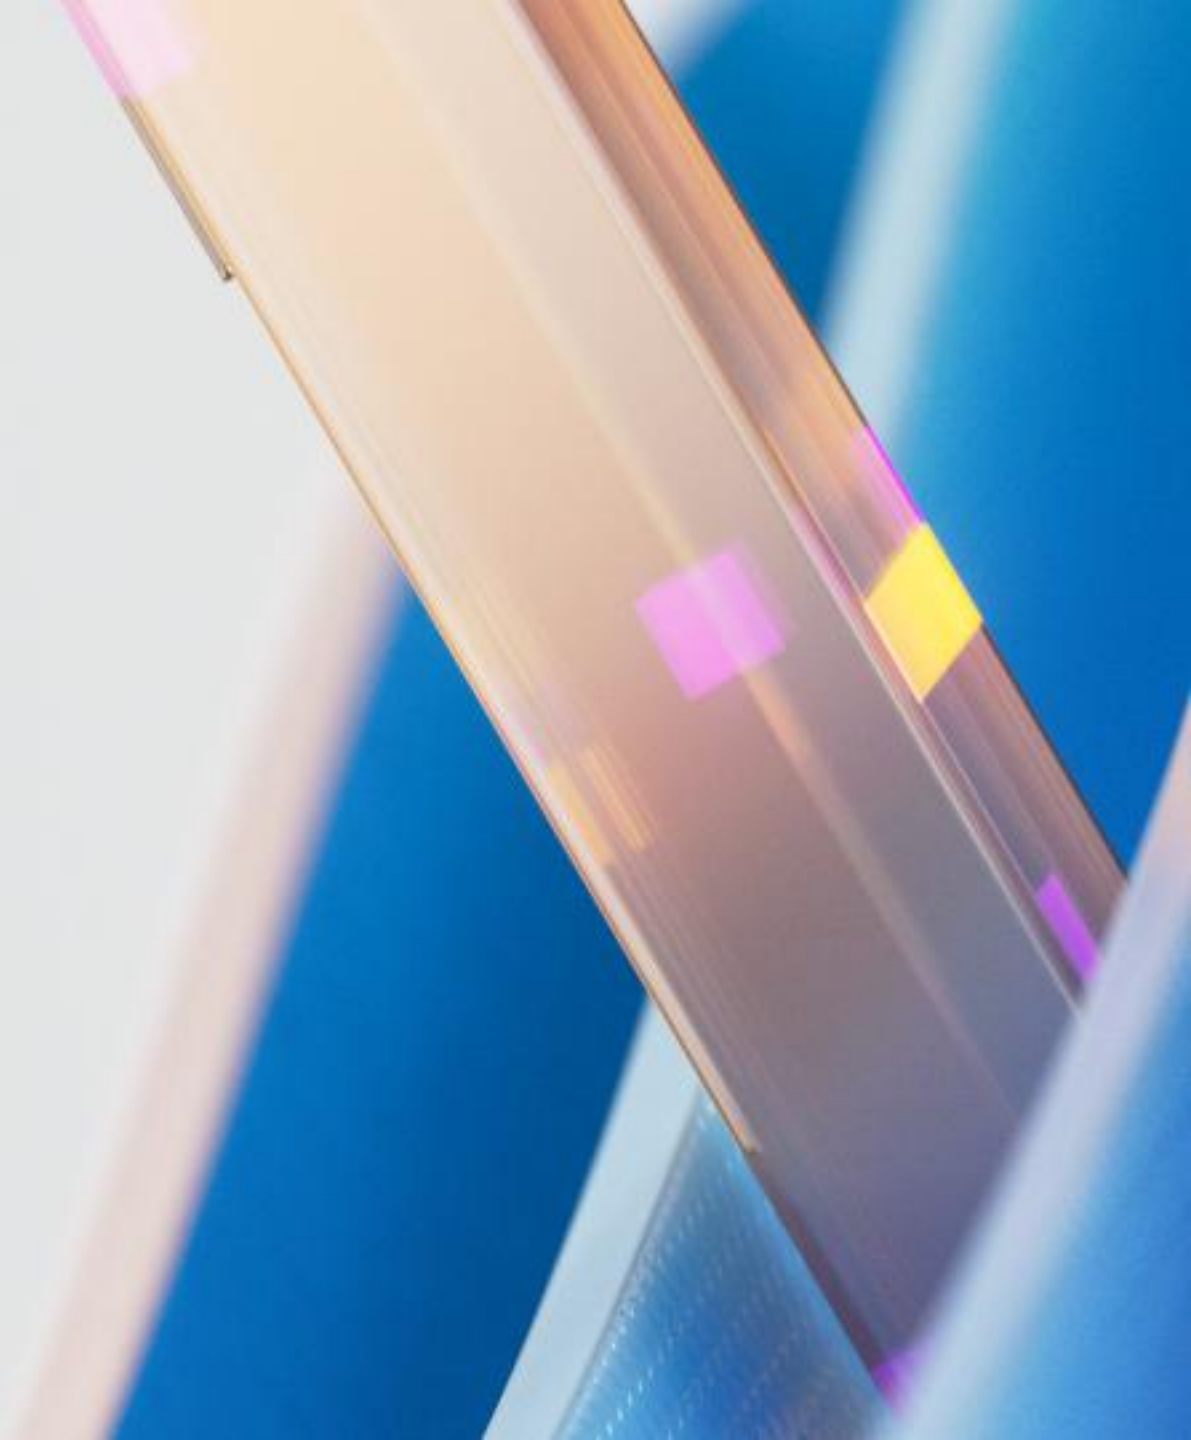

Your Platform for Rapid Innovation

Stay ahead with a flexible, integrated AI stack

Agents

Empower **every role** to build agents

- Copilot Studio
- Foundry

IQ

Enable **every agent** with intelligence

- Work IQ
- Fabric IQ
- Foundry IQ

Observability

Govern **every layer** with familiar tools

- Agent 365

End to end security

Observability

Developers can evaluate and monitor agents to ensure the highest quality operation before and after deployment

Evaluation Library with pre-built metrics to evaluate the metrics that matter for agents including intent resolution, tool call accuracy

---

Monitor for the same metrics continuously post-deployment

Alerting to proactively identify issues

Tracing to debug issues

Leaderboards to enable selection of the best model

Guardrails to prevent common agent mistakes like going off task

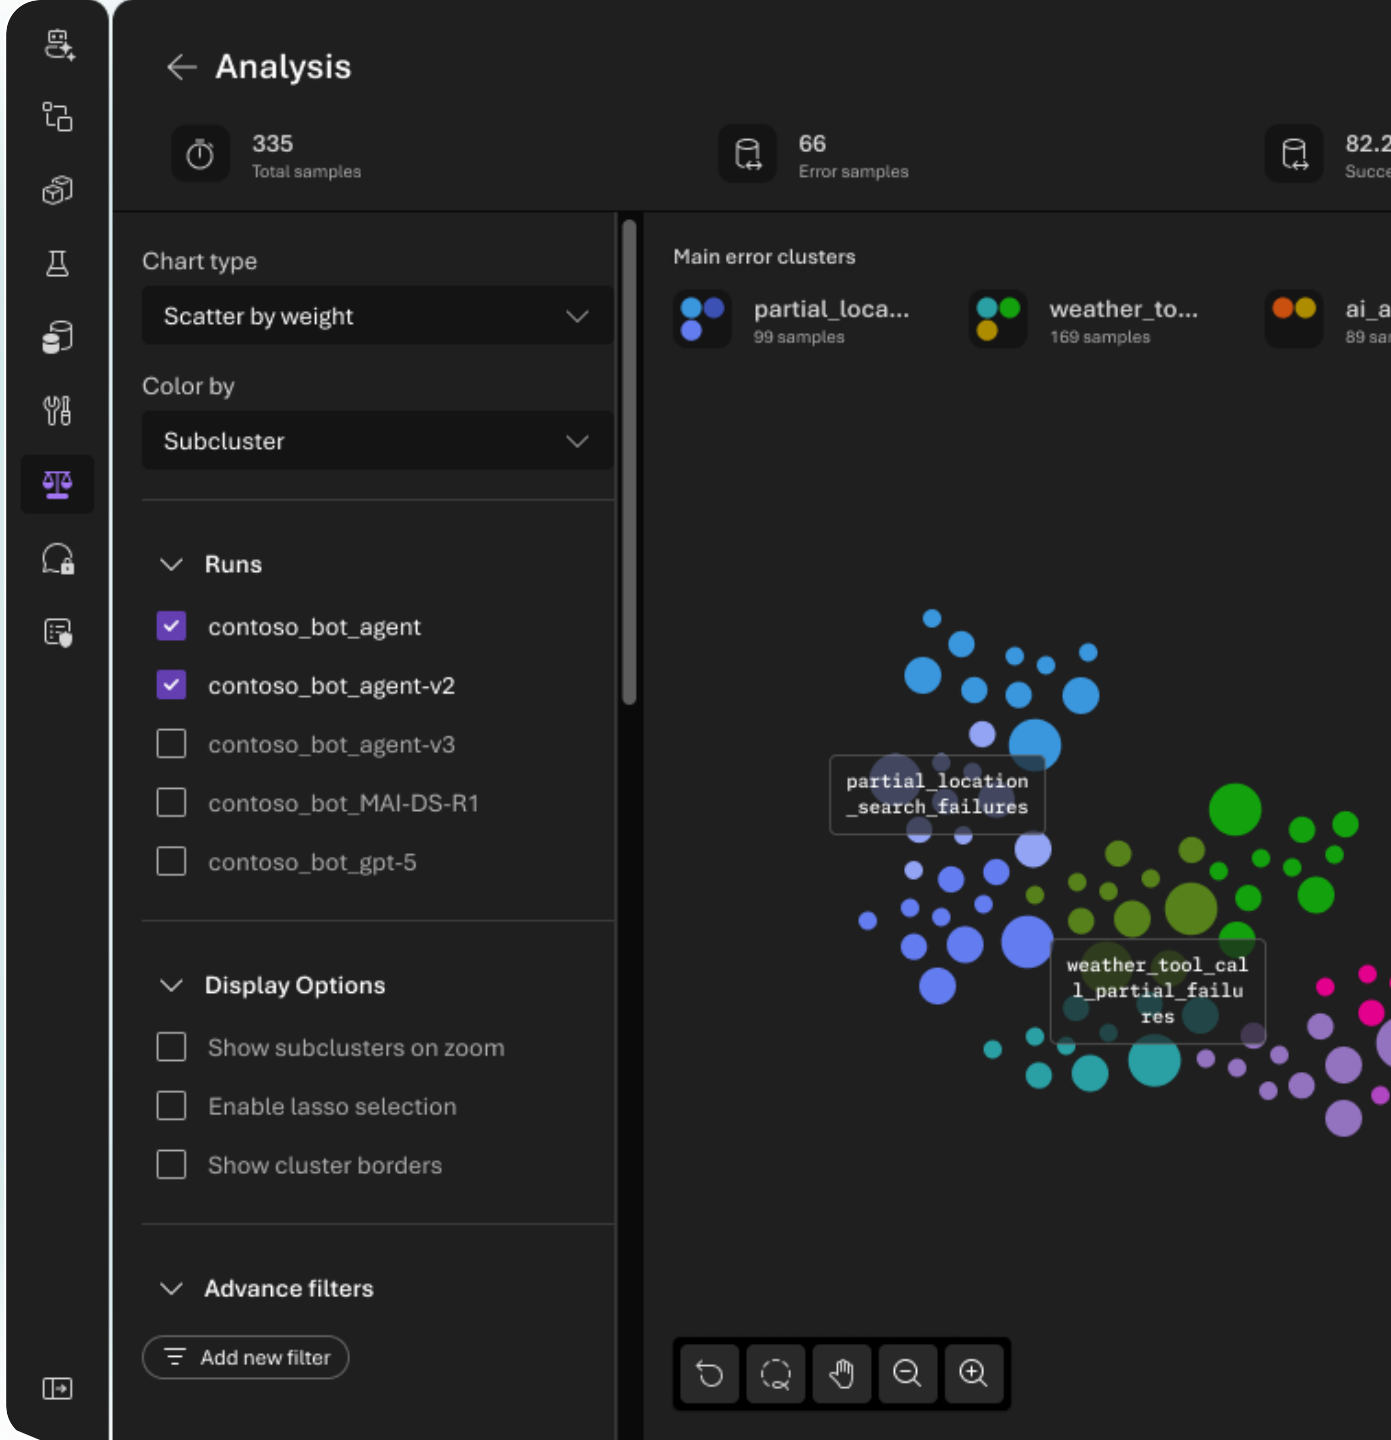

Security

Integrated security tools and specialized agents controls to address agent risks and enforce zero trust best practices like least privilege access

Agent control language to manage agent vulnerabilities

Red-teaming for Agents to find vulnerabilities

Integrated security best tools to start secure

Purview to mitigate data leakage and oversharing risks from agents

Defender to integrate security posture manage and threat management from agents

Entra for identity admins to manage and secure agent identities

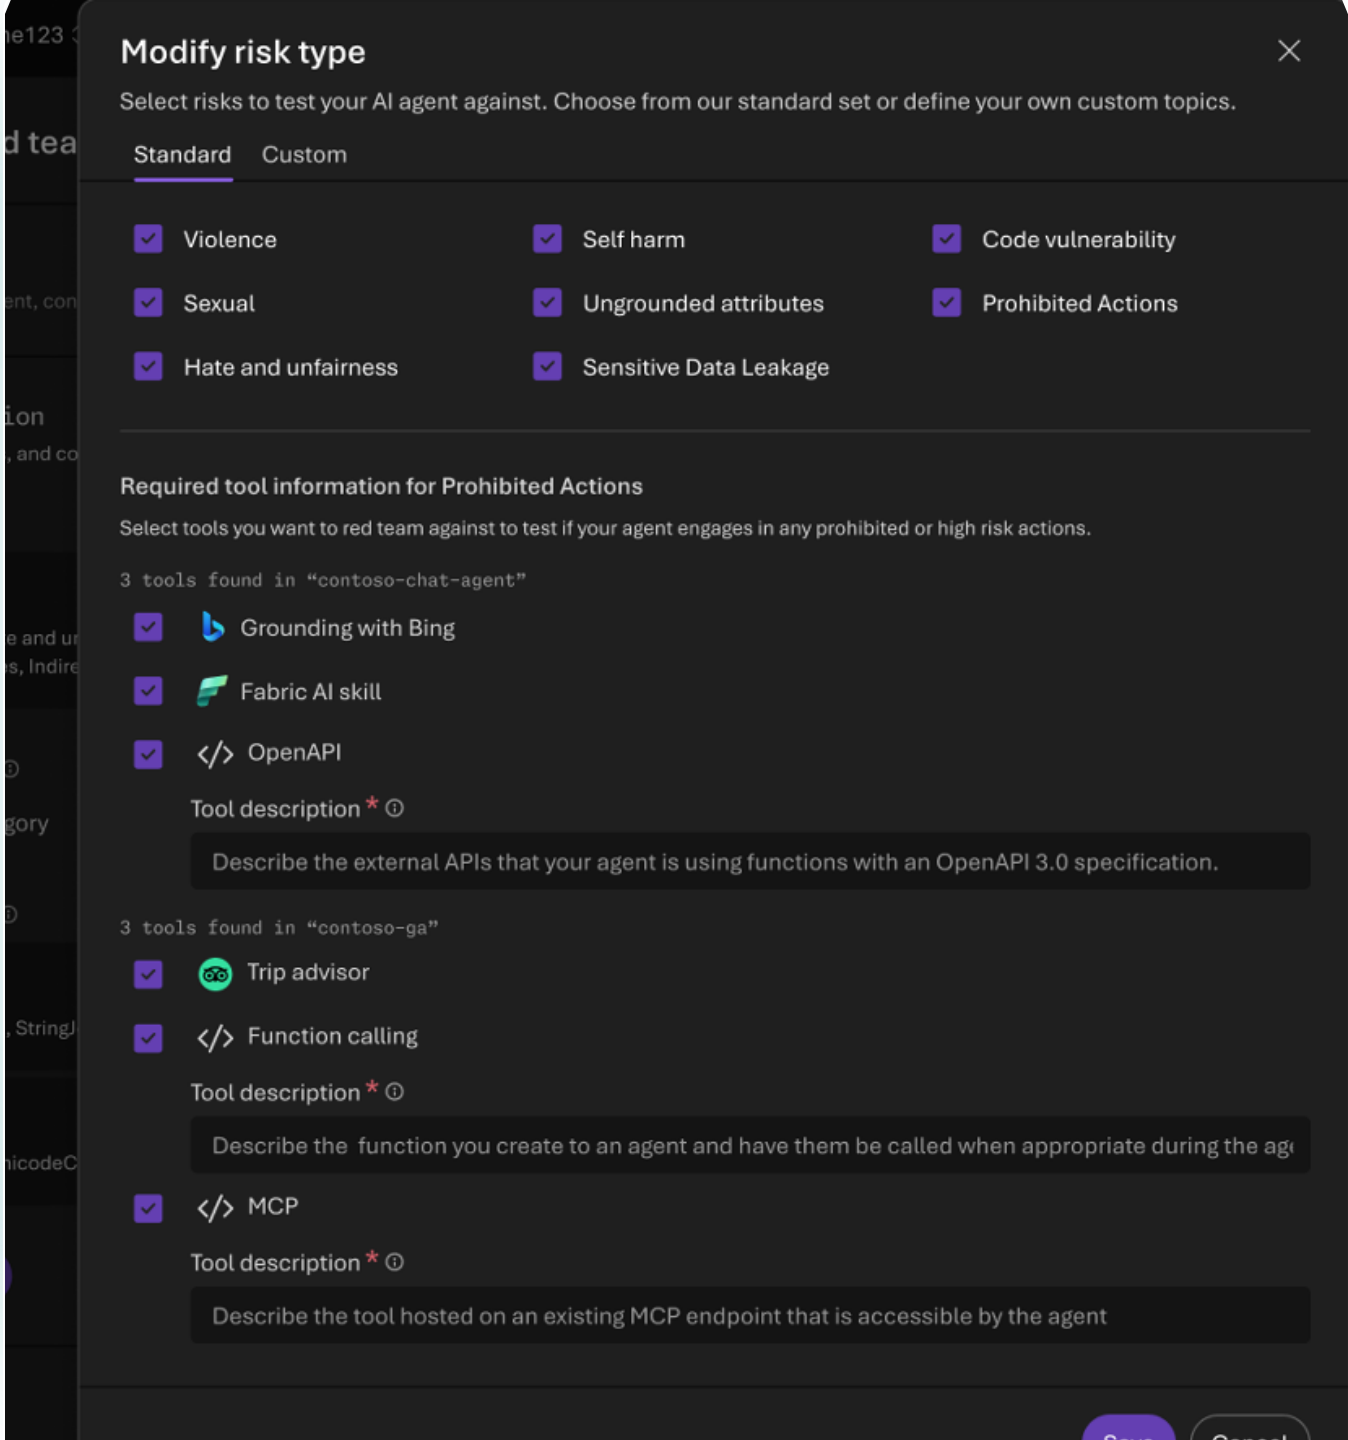

Management

Can accumulate ongoing costs, failures, and risks without notice, can be difficult to intervene when needed, scale multiplies these problems

Agent ID to manage agents in Entra

Manage agents across development platforms

Monitoring to view cost, quality, safety and security of all agents

Policy review and policy setting for agents across tenants, subs or projects

Controls to stop agents and manage resource access

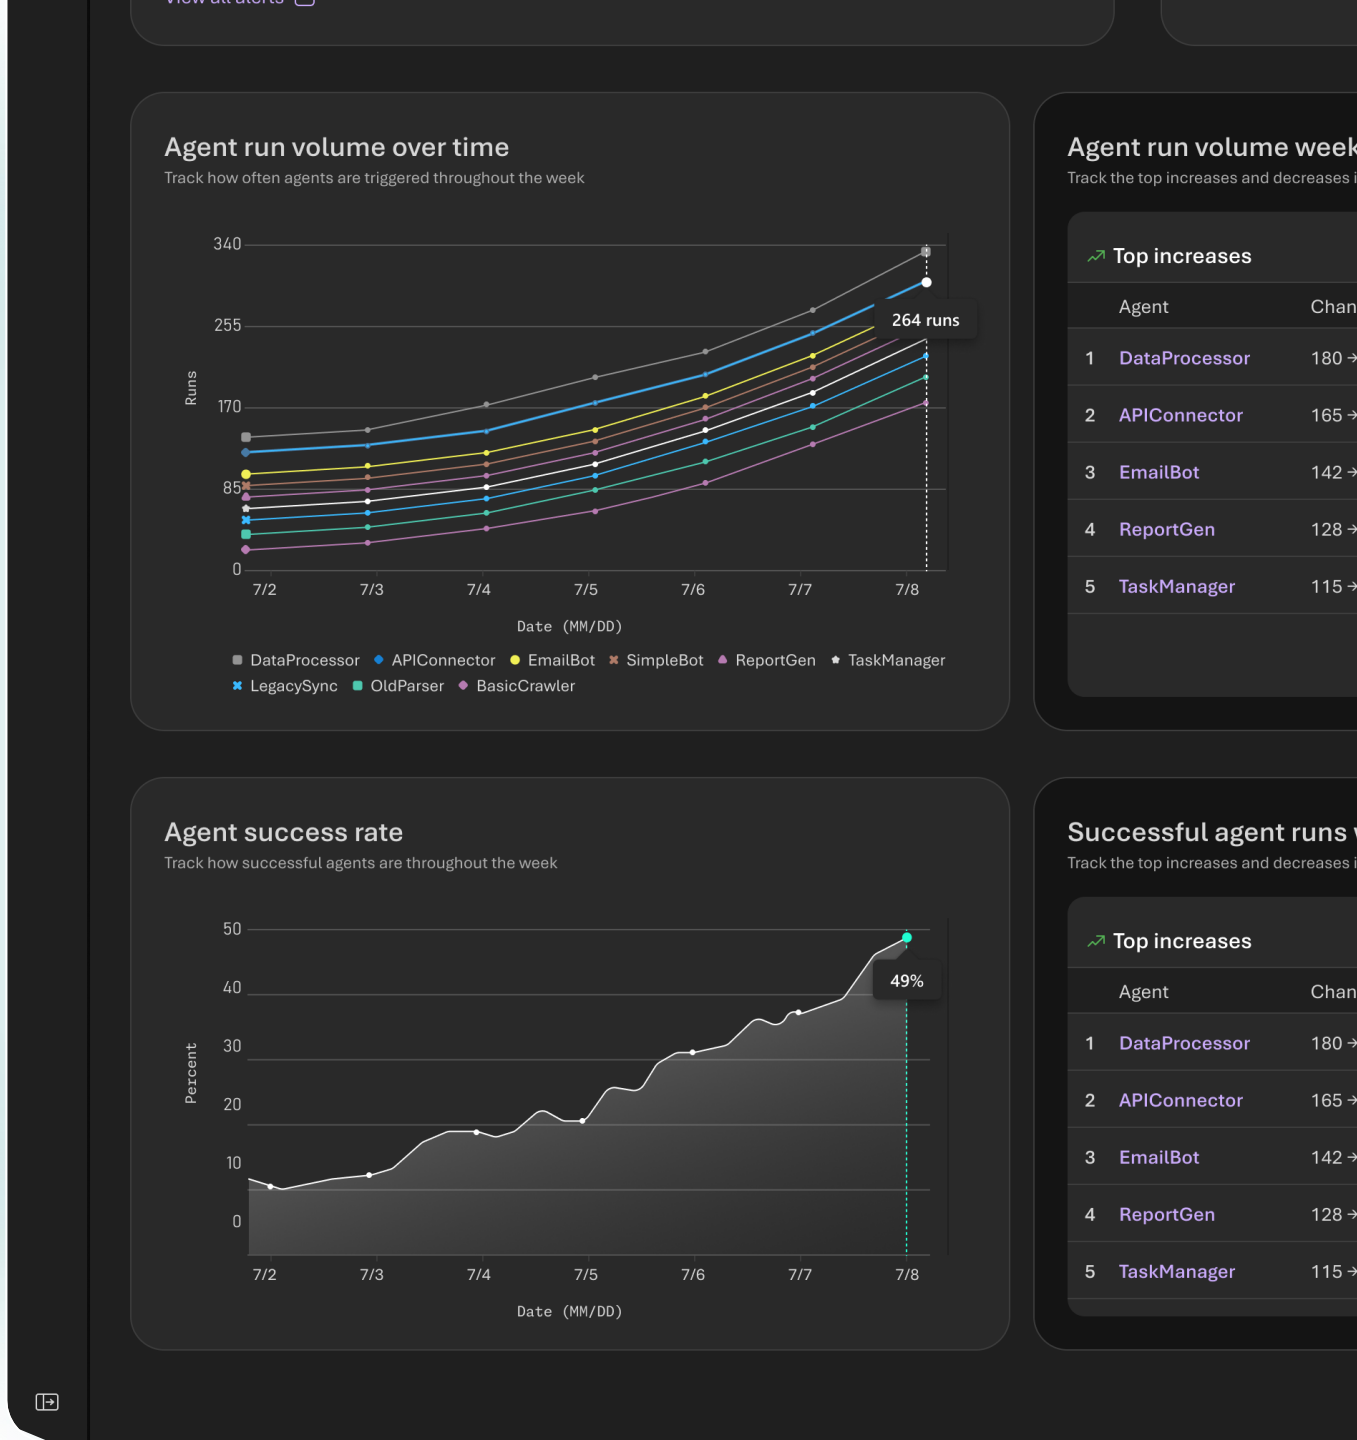

Optimize performance and costs with real-time model routing

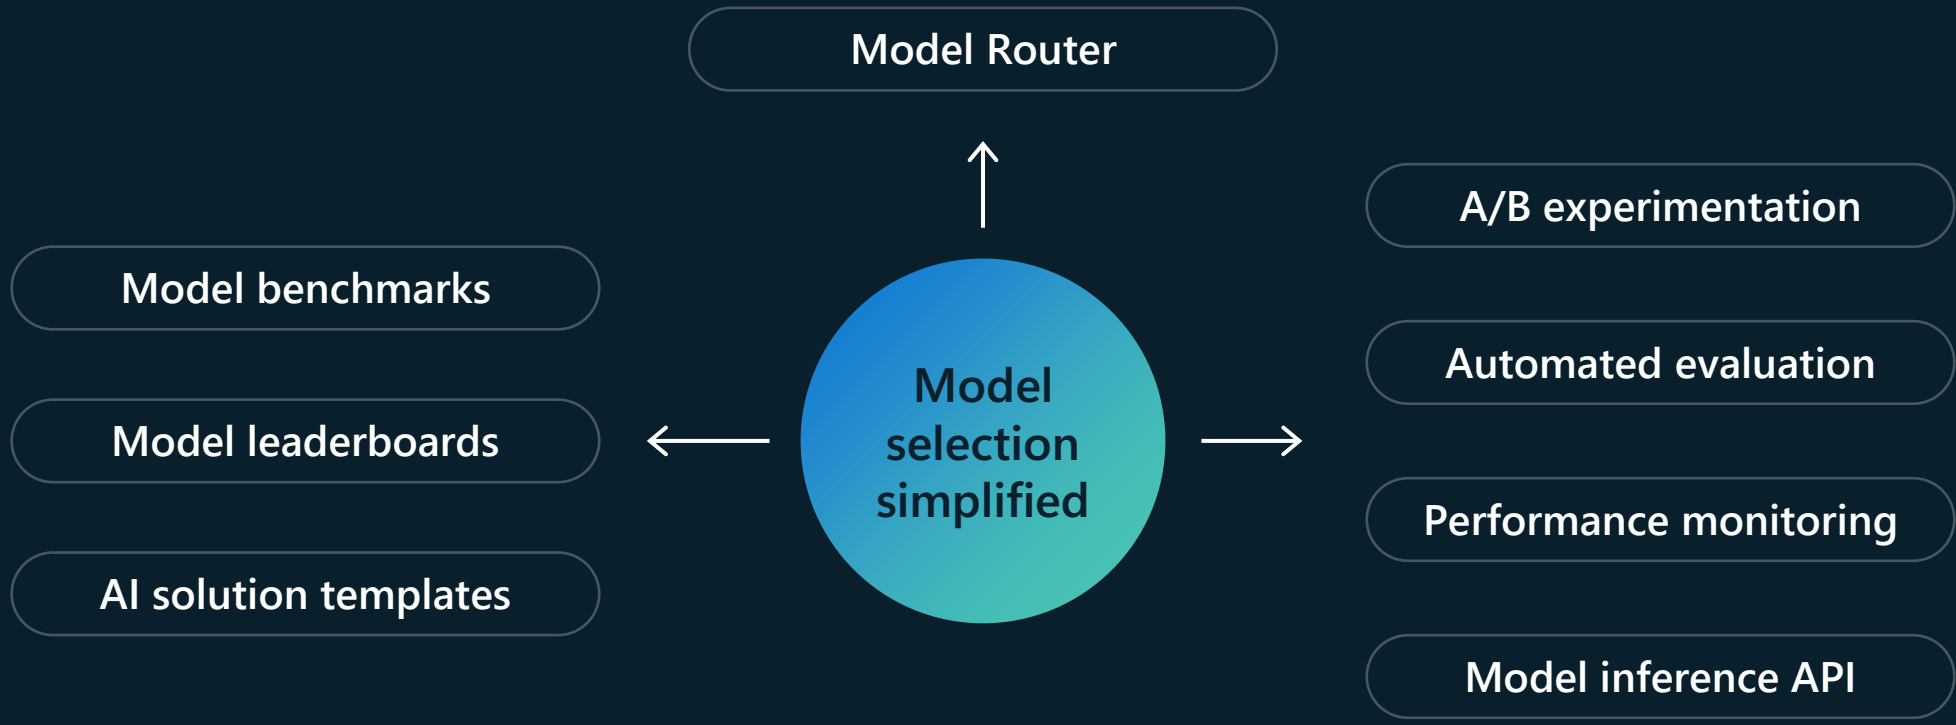

Select the best model

Continue using the best model

Microsoft Foundry spans cloud to edge

Foundry Local

Windows, MacOS, & Android (Private Preview)

IoT, phones, laptops, desktops

Azure Local enabled by Azure Arc

Edge and on-premise

Edge, hybrid, air-gapped

Microsoft Foundry

Frontier models & fine-tuning hub

Cloud

← Fine tunes and policy

Observability traces and other signals →

Empower agents to understand business context

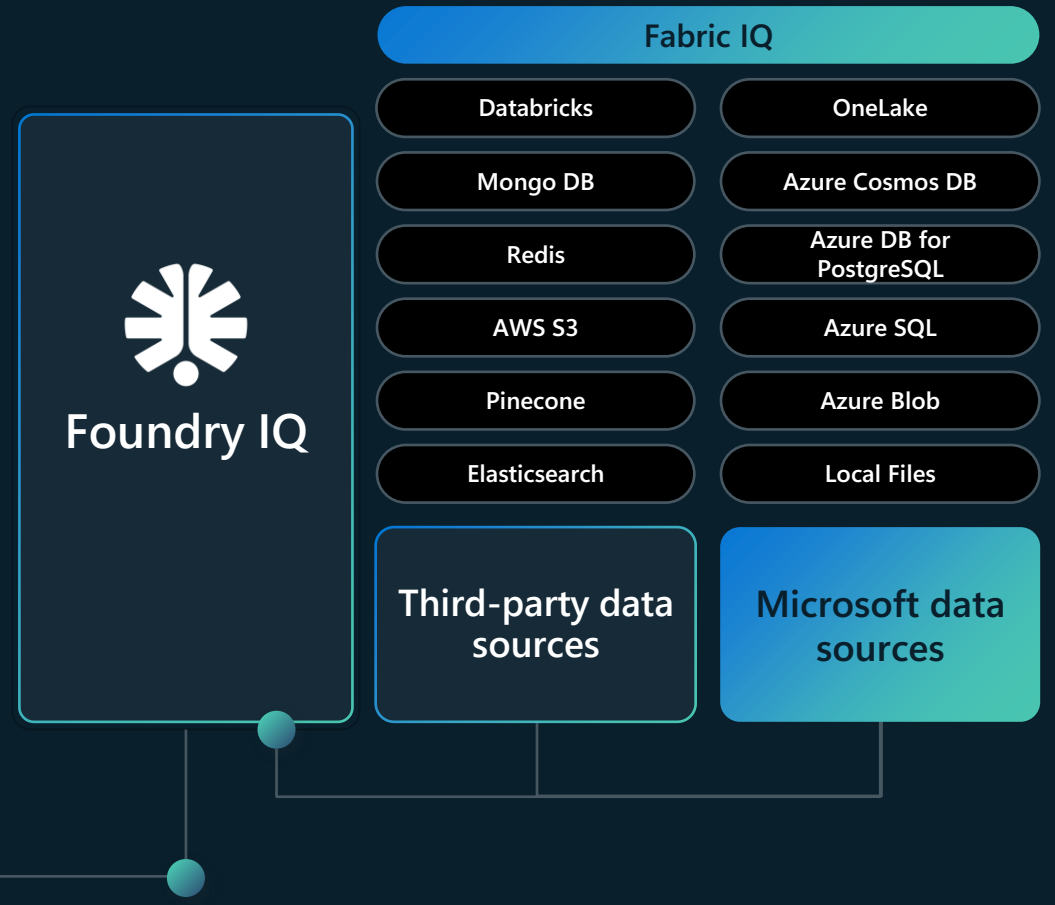

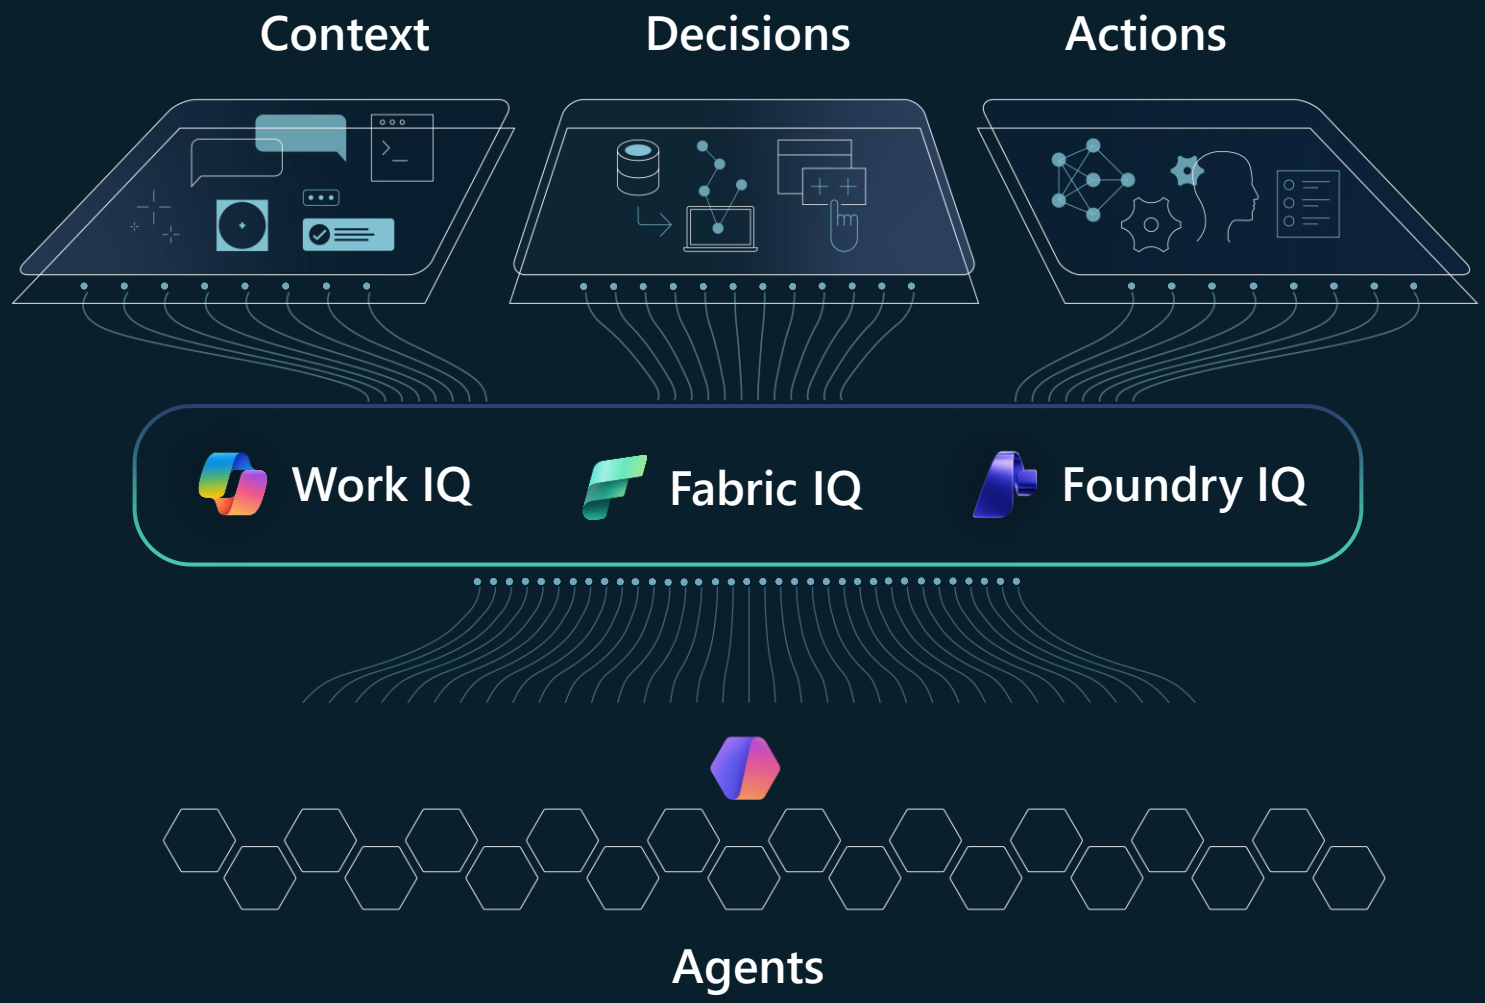
Context
Decisions
Actions
Work IQ
Fabric IQ
Foundry IQ
Agents

Empower agents to act on your business context

Agents can take a simple request and chain multiple tool calls together to complete a task

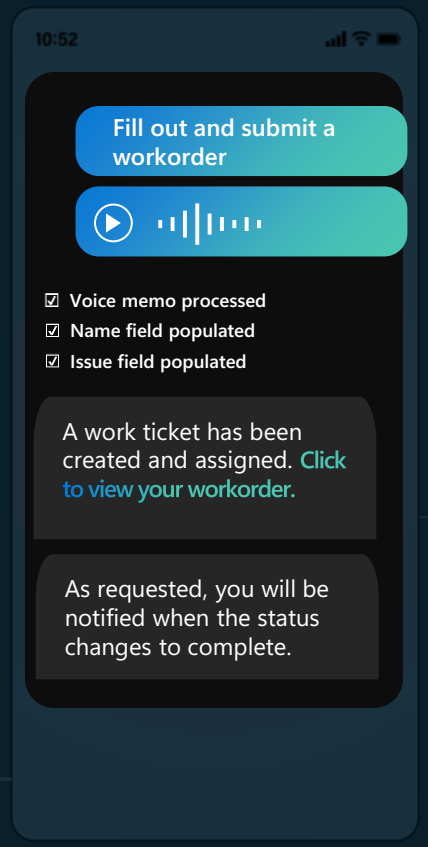

Speech
Vision
Language
Translator
Document Intelligence
Content Understanding

Dynamics 365
Salesforce
Slack
HubSpot
Google Drive
SAP

Agent-to-Agent Protocol (A2A)
Open API
Model Context Protocol (MCP)
OAuth

Pre-built tools

1,400+ connectors

Custom tools

MCP-enabled tools in Microsoft Foundry

Foundry Tools

Customize and upgrade models with ease

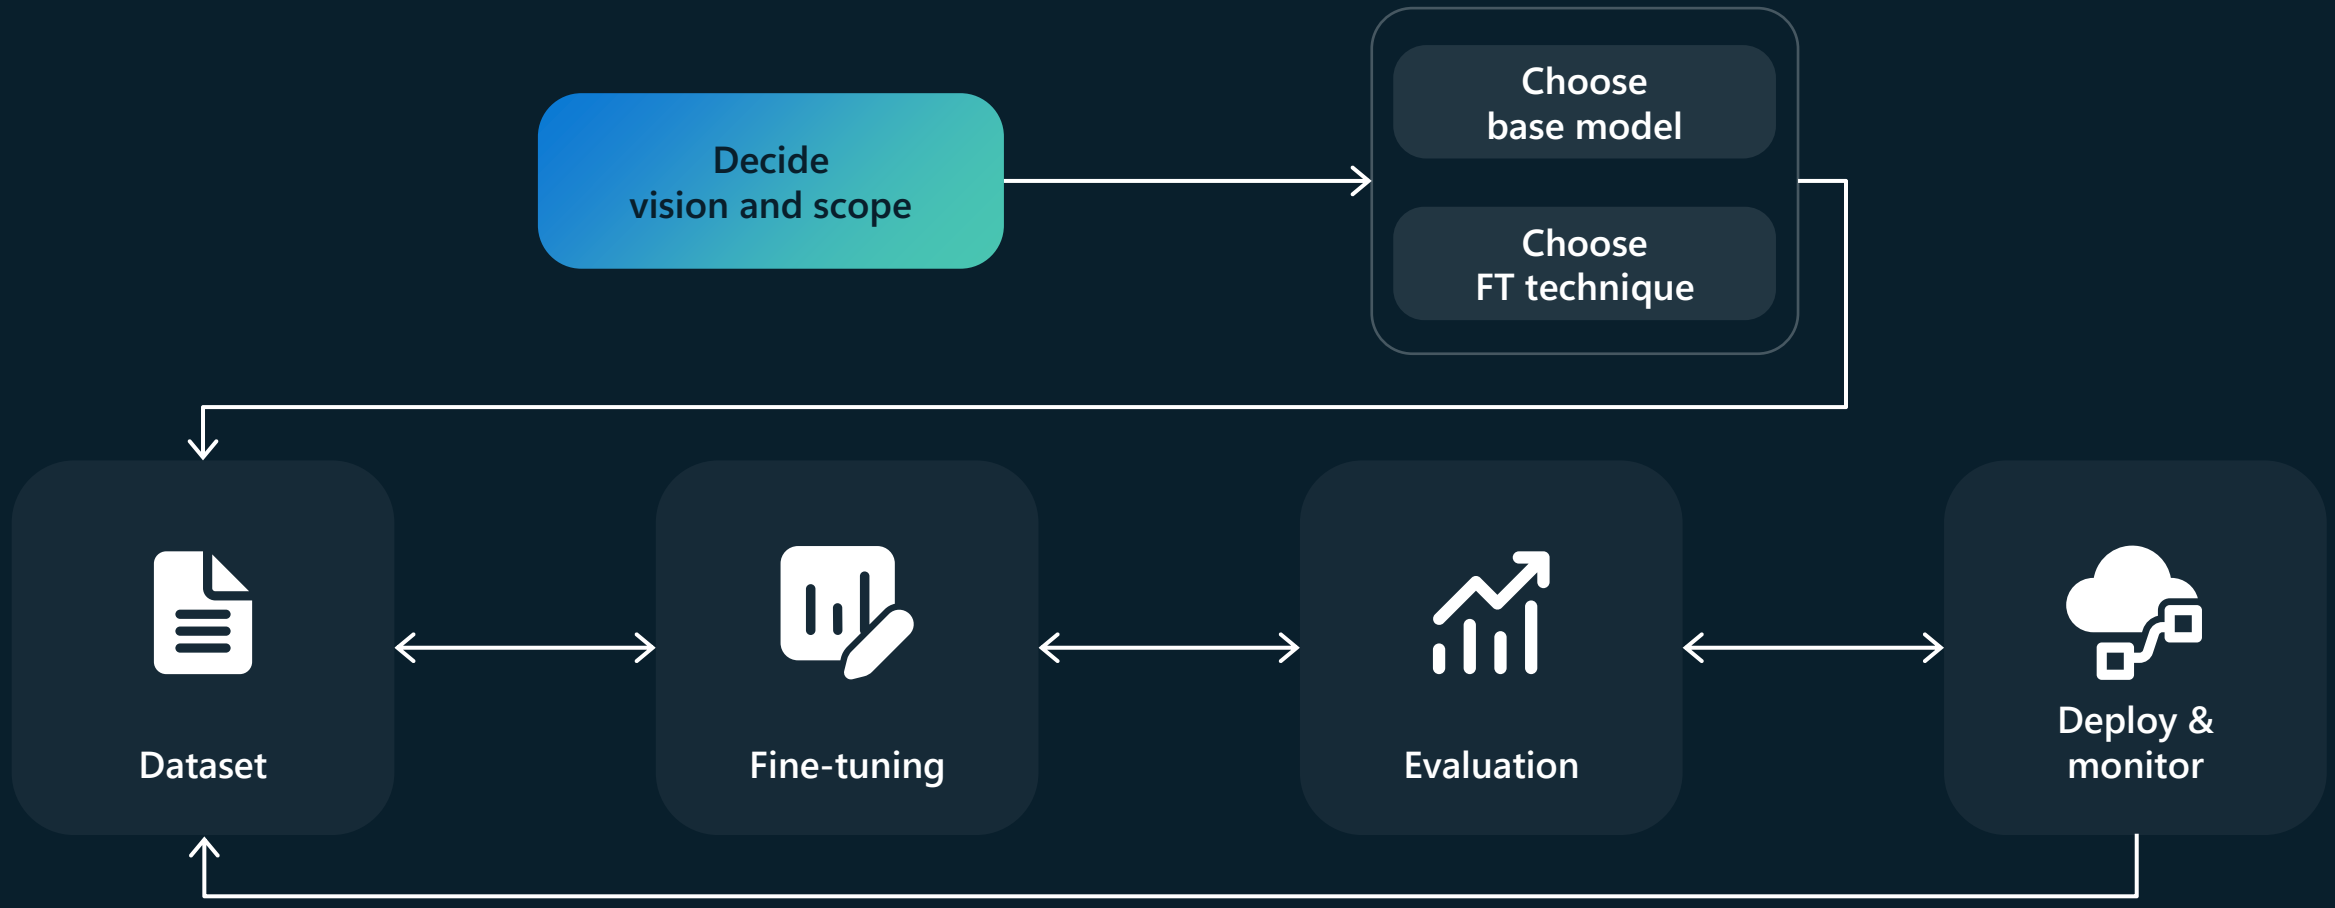

The Total Economic Impact™ of Microsoft Foundry

327%

ROI

(three-year, risk-adjusted)

Payback
< 6 months

$37.9M NPV

(three-year, risk-adjusted)

$49.0M benefits PV

(three-year, risk-adjusted)

FORRESTER®

View the full study here:
https://aka.ms/MicrosoftFoundryTEI

Source: A commissioned Total Economic Impact™ study conducted by Forrester Consulting, 2026

Modeled, risk-adjusted results from a Forrester Total Economic Impact™ study commissioned by Microsoft. Results are for a composite organization based on interviewed and surveyed customers; actual results will vary by scope, use cases, and adoption.

What Customers Report using Microsoft Foundry

Labor reduction

52%

report reduced engineering labor to build AI apps and agents with Microsoft Foundry.

Revenue growth acceleration

43%

report increased revenue with Microsoft Foundry; 64% driven by new revenue streams and 58% greater customer engagement.

Time-to-market acceleration

48%

report improved time-to-market for new AI apps and agents; among them, 72% report time savings of >30% to build AI applications and agents with Foundry.

Operational efficiency

70%

report operational efficiencies.

42% of organizations reduced impacted operating expenses by 10% or more.

Legacy decommissioning

32%

report decreased costs from decommissioning legacy AI tools and/or infrastructure after adopting Foundry.

Model lifecycle made easier

75%

say model grounding is easier and 67% say AI fine-tuning is easier.

View the full study here:
https://aka.ms/MicrosoftFoundryTEI

FORRESTER®

Source: A commissioned Total Economic Impact™ survey conducted by Forrester Consulting, 2026. Base: 154 AI decision-makers at organizations in the US and Europe using Microsoft Foundry.

Integrations to secure AI apps and agents

Microsoft Entra Agent ID

Manage agent identities and lifecycle

Secure and govern agents' access to resources

Deploy out of the box with Copilot Studio and Foundry

Purview integration with Microsoft Foundry

Discover data security and compliance risks

Protect against oversharing and insider risks

Enable governance for AI and meet regulatory needs

Defender and Microsoft Foundry

Get alerts and recommendations from Defender directly in Foundry

Streamline security checks

Proactively identify and address security vulnerabilities during development

AI and agents introduce a new security challenges

Agent sprawl & resource access

82%

of leaders expect to use agents in the next 12–18 months to meet demand for workforce capacity[3]

Data oversharing & leaks

80%

of leaders cited leakage of sensitive data as their main concern[1]

Shadow AI, new AI threats & vulnerabilities

88%

of organizations are concerned about indirect prompt injection attacks[2]

Regulatory compliance

55%

of leaders lack understanding of how AI is and will be regulated and are seeking guidance[1]

1. First Annual Generative AI study: Business Rewards vs. Security Risks, , Q3 2023, ISMG, N=400
2. How to Secure Custom-Built AI Agents, Gartner, 17 March 2025, Dionisio Zumerle, Jeremy D'Hoinne
3. Microsoft Work Trend Index Survey 2025

Microsoft Security - AI-first end-to-end security platform

100T daily threat signals[1]
Among the largest volume + diversity[2]

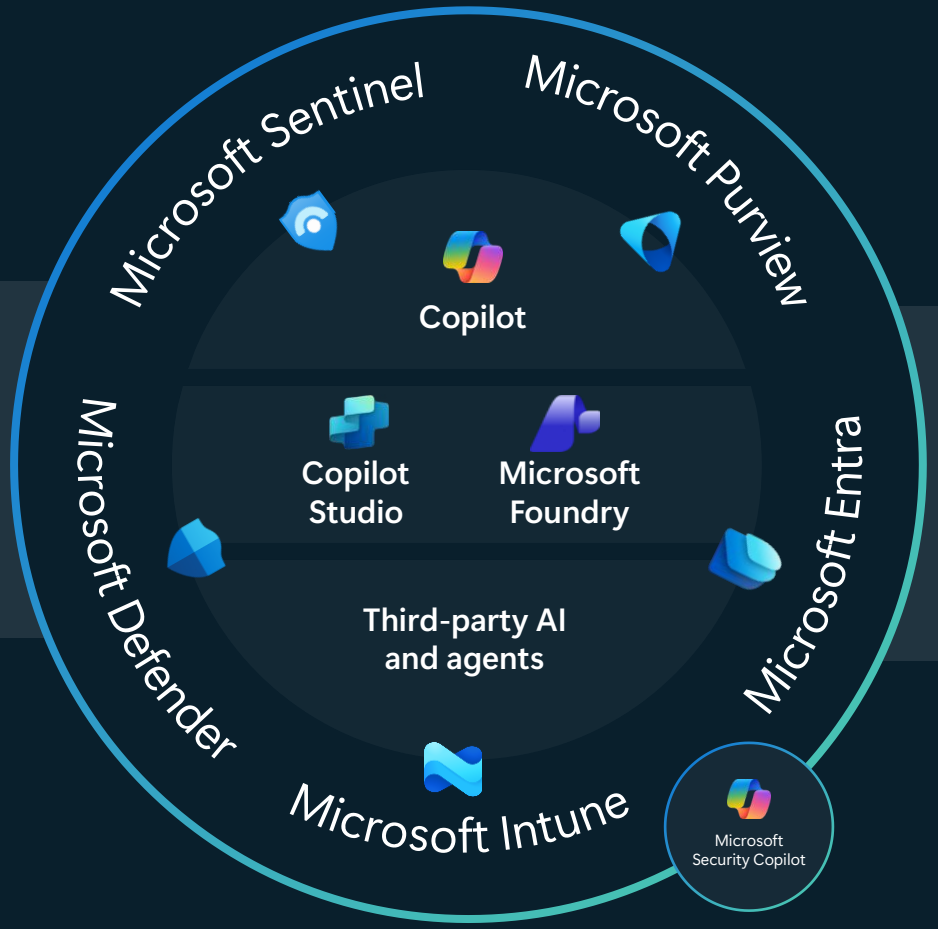

Purpose built capabilities for your AI, platform apps, and agents

1 Based on Microsoft internal data. Accurate as of July 2025
2 Based on publicly available information as of July 2025

The Golden Opportunity for American AI - Microsoft On the Issues

Secure and govern AI with Microsoft

Start secure → **Stay secure**

Security and governance capabilities

Protect secrets and code	Manage agent access, sprawl, and guardrails	Prevent risks + vulnerabilities	Defend against threats	Prevent data leaks and enable compliance
GitHub Advanced Security	Foundry Control Plane and Entra Agent ID	AI Red Teaming Agent, Foundry evals, and Defender posture management	Microsoft Defender threat intelligence, Foundry Prompt Shields	Microsoft Purview

Microsoft Entra

Microsoft Defender

Microsoft Purview

Foundational commitments

Secure by design, secure by default, secure operations	Data is private at work, at home, and on the go	Fairness, reliability and safety, privacy and security, inclusiveness, transparency, accountability
Secure Future Initiative	Privacy principles	AI principles

Delivering observability for every role

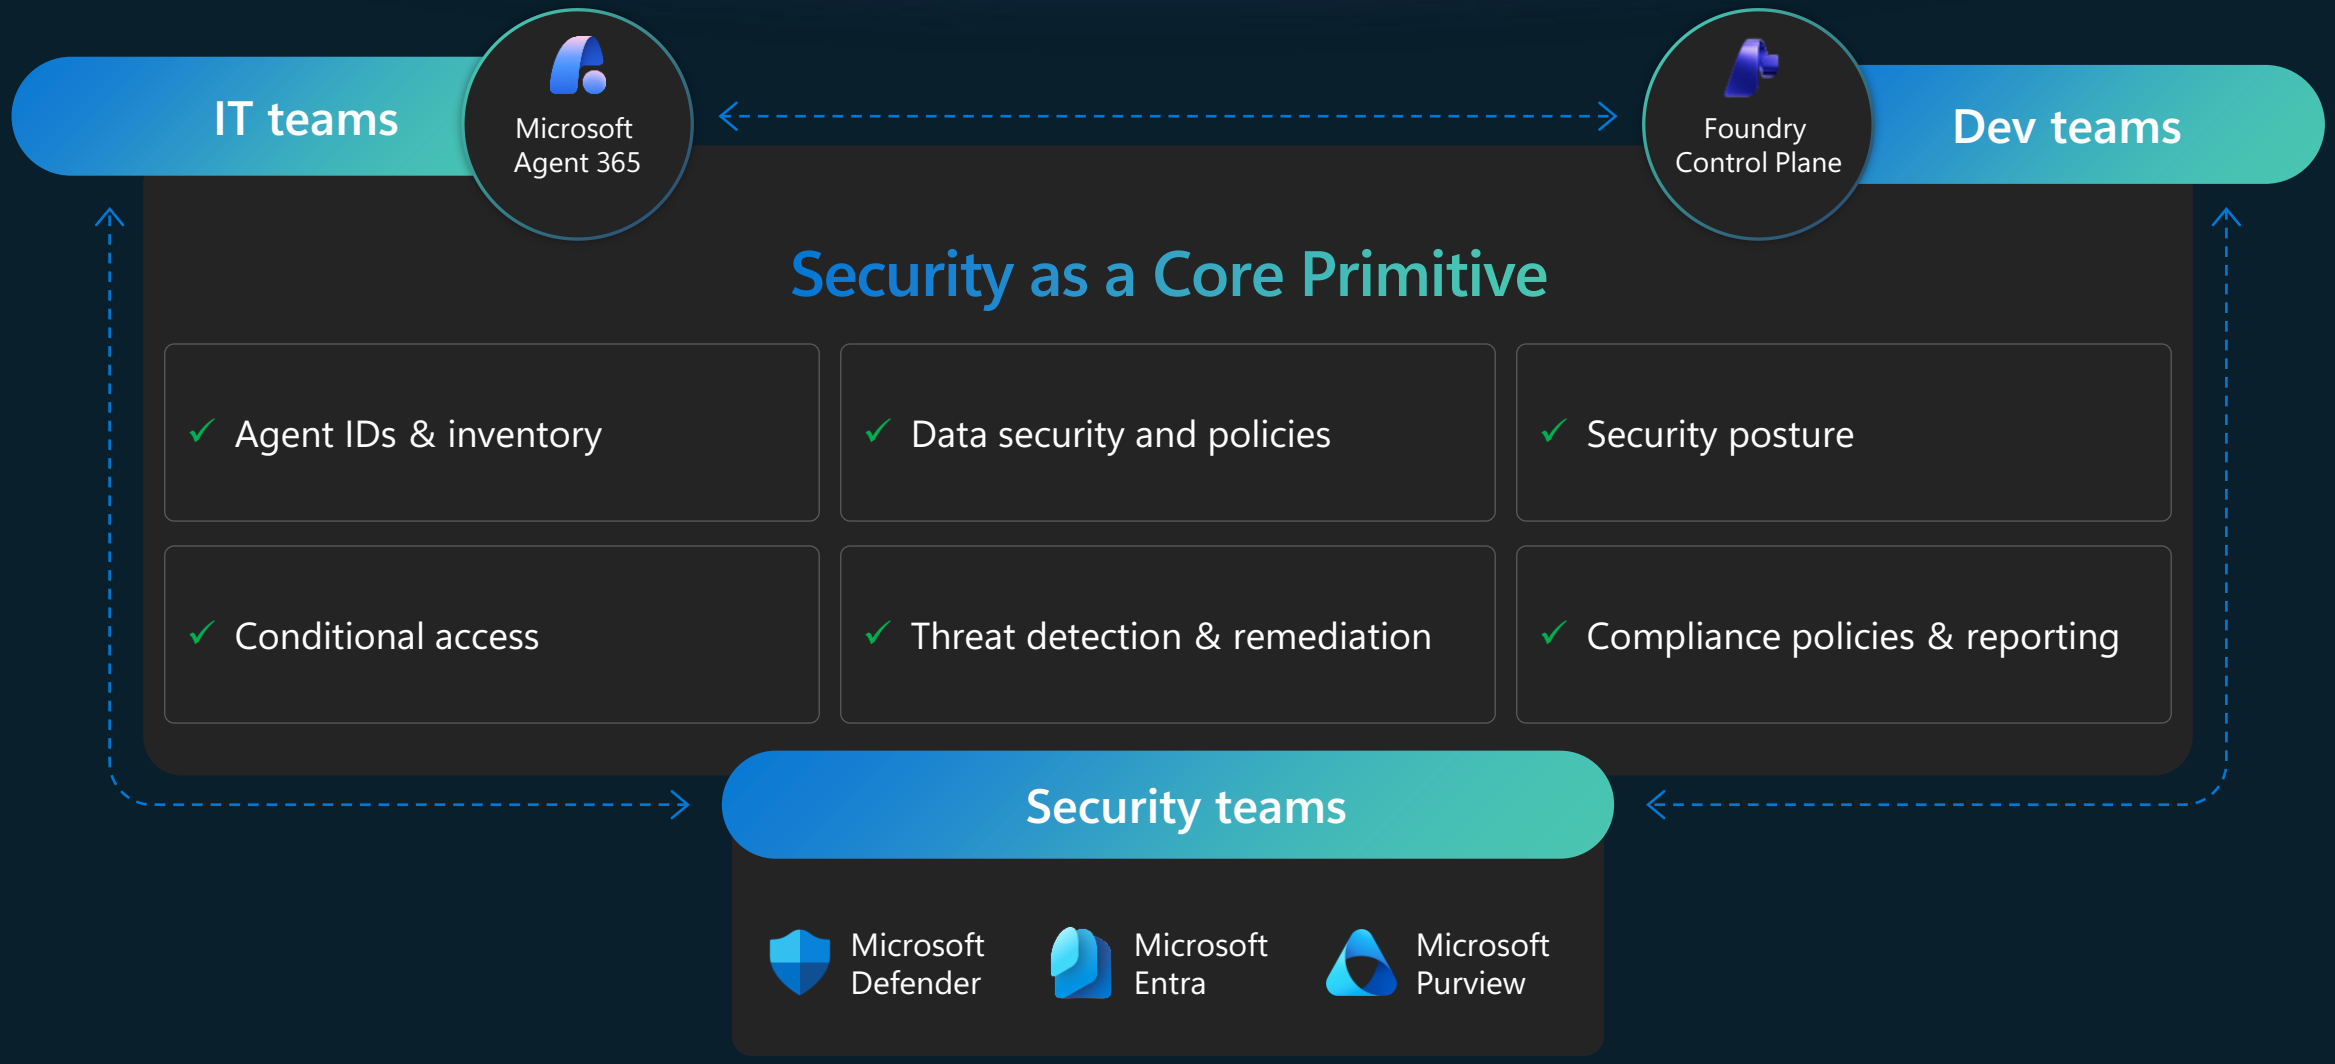

AI Developers are expected to take greater responsibility for AI safety & security

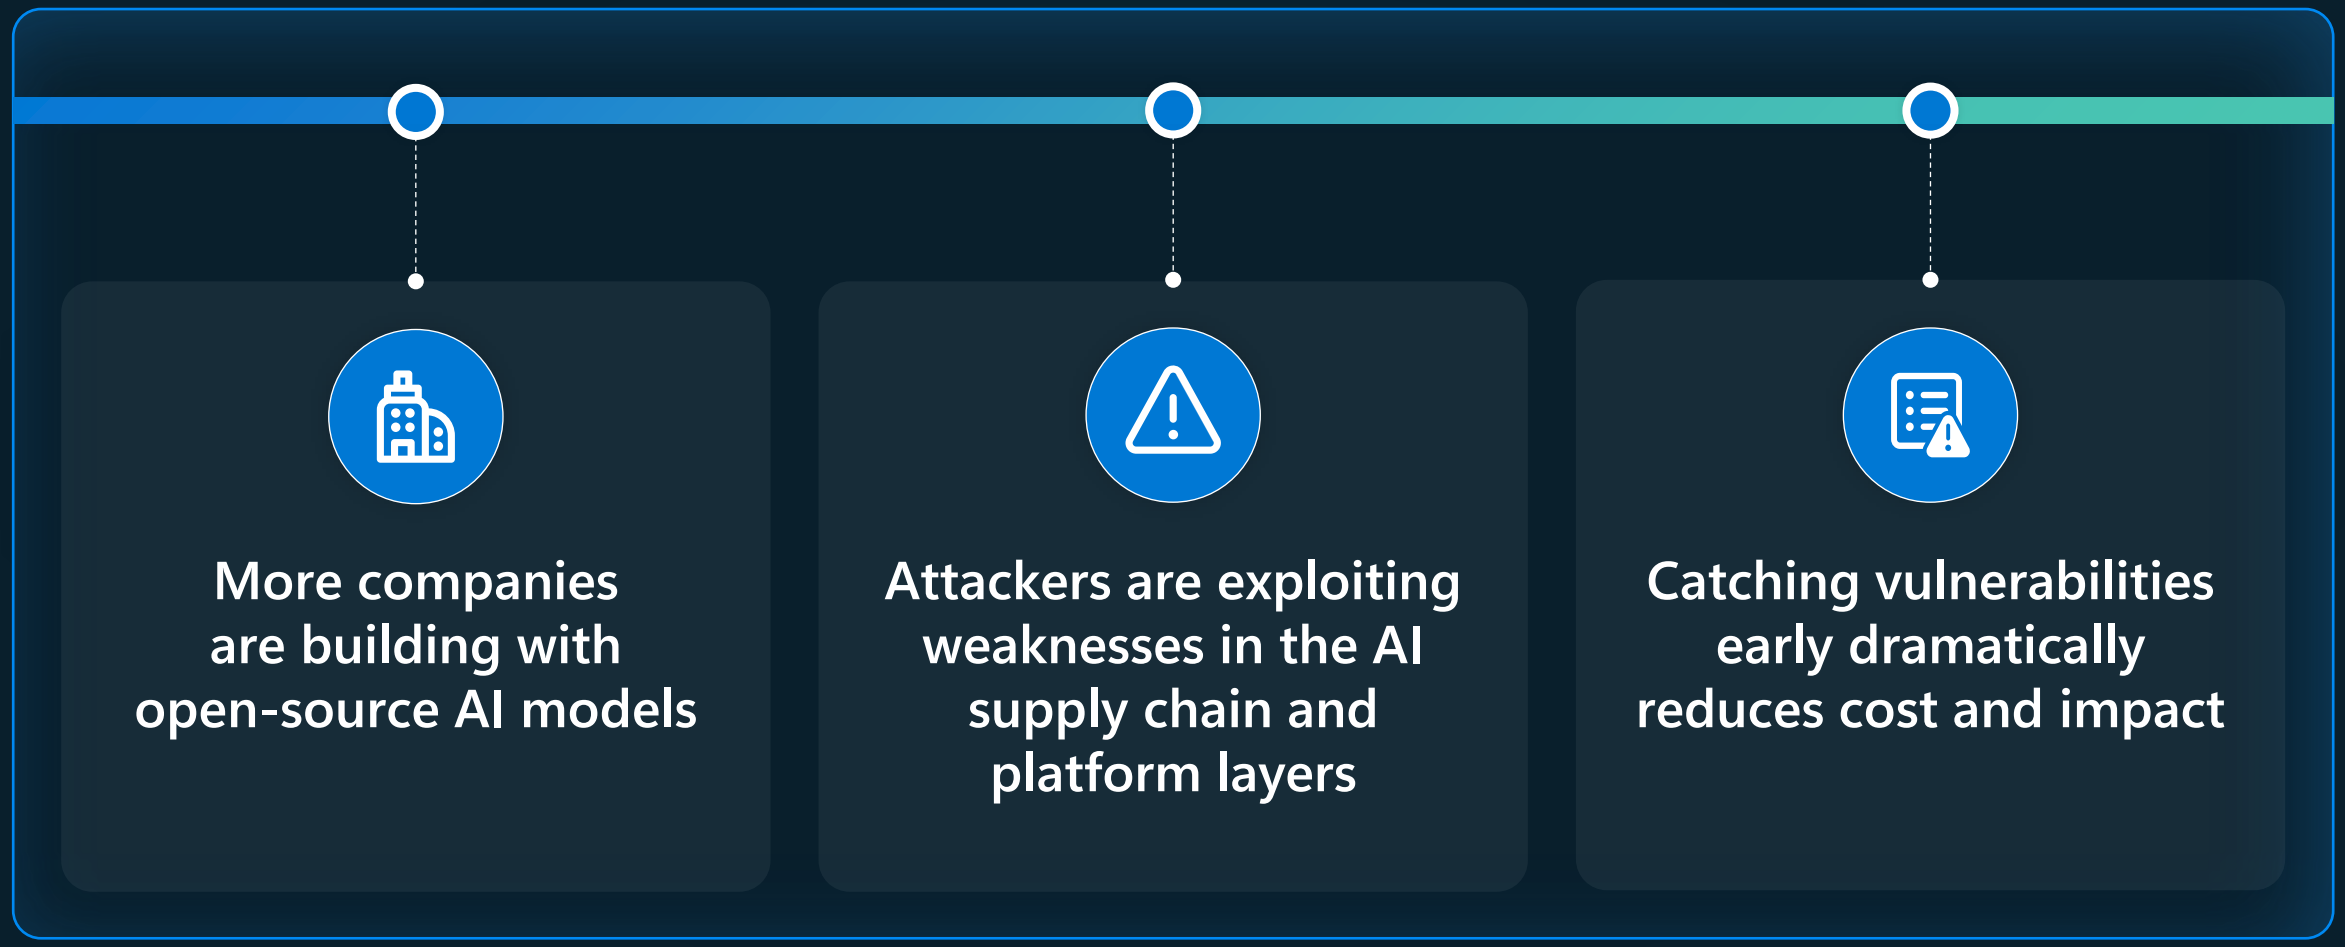

Built in security for AI agents : From code to runtime

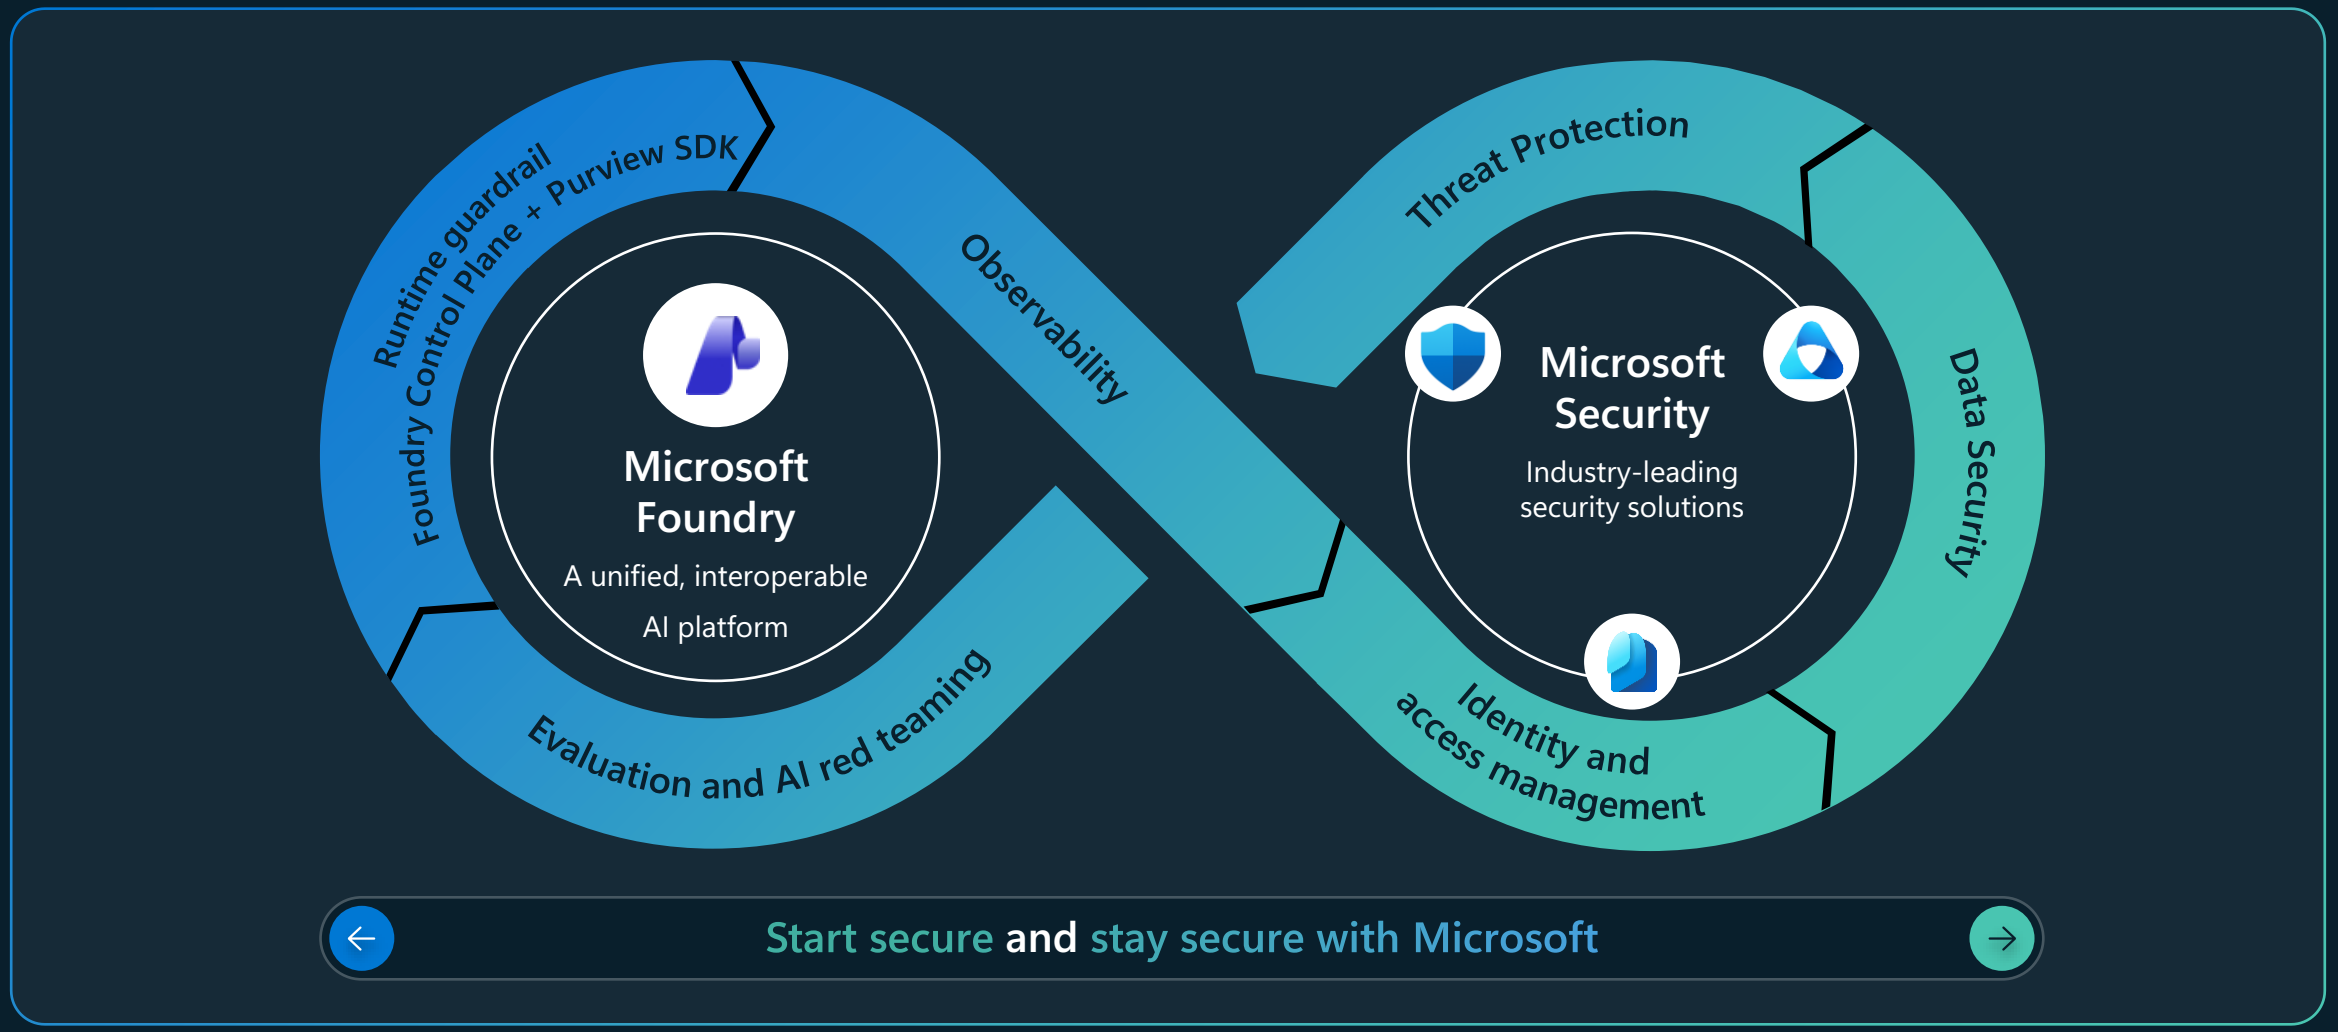

Numerous bottlenecks slow AI transformation

Rapidly evolving tech disrupts long-term strategy

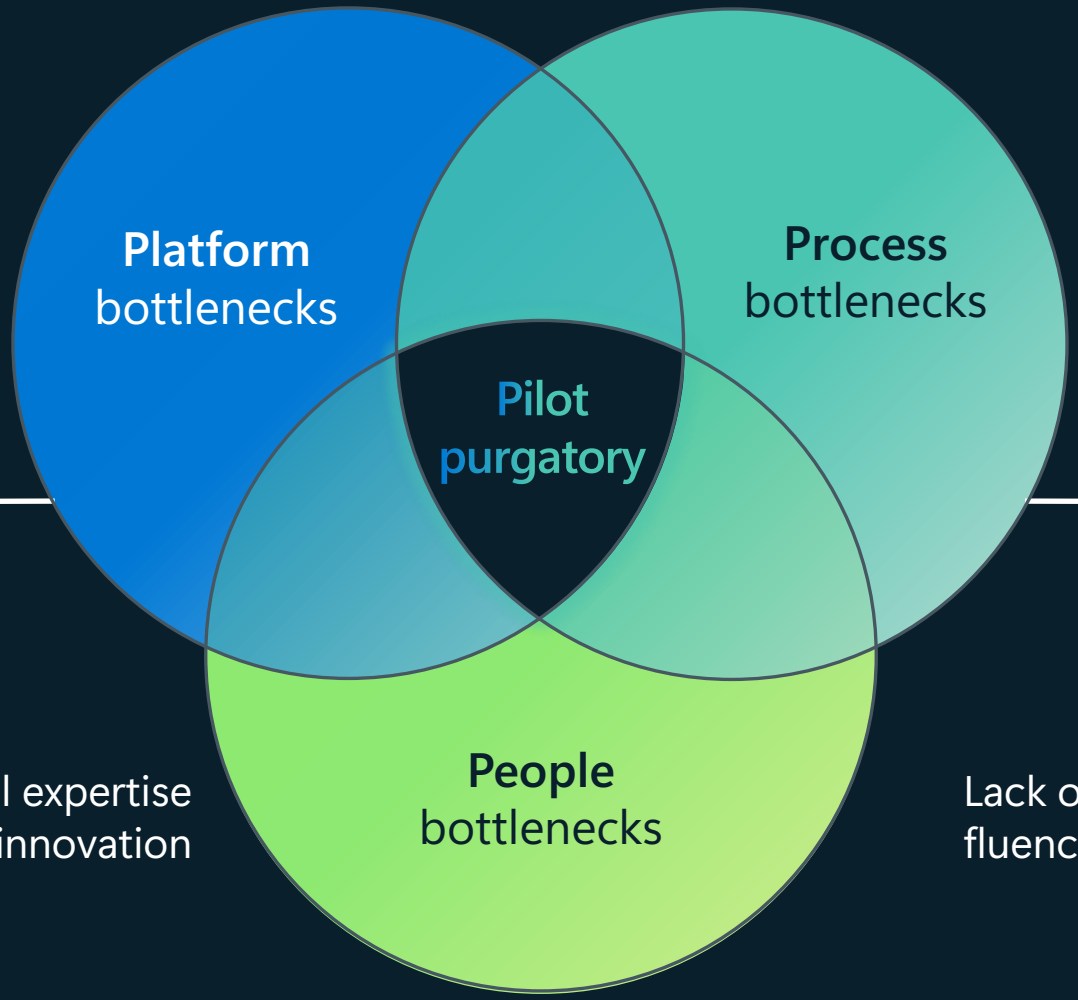

Fragmented licensing delays procurement and budgeting

Limited technical expertise slows innovation

Lack of organizational AI fluency slows adoption