



AI Agent Platform Comparison Matrix

Copilot Studio, Microsoft Foundry Agent Service, OpenClaw and custom-built agents compared on cost, complexity, governance, integration and scale. Use it to shortlist, then test against your own workflow.

Kevin Kaminski, Microsoft MVP

Big Hat Group Inc.

Version 1.0 · October 2026

	Microsoft Copilot Studio Low-code SaaS on Power Platform · built by makers and fusion teams	Microsoft Foundry Agent Service Managed platform for developers · formerly Azure AI Agent Service	OpenClaw Open-source, self-hosted · for teams that run their own infrastructure	Custom-built agents Code-first frameworks · professional developers
BEST FOR	Microsoft 365-centric organizations building help-desk, HR and business-process agents that live in Teams and Copilot.	Azure-centric development teams that need custom orchestration and model choice without running infrastructure.	Multi-channel agents (Teams, Slack, WhatsApp, email), browser automation and persistent memory, under your full control.	Agents that are your product or a core differentiator, and complex multi-agent systems where control beats speed.
COST MODEL	Copilot Credits: prepaid packs (no rollover), pay-as-you-go through Azure, or a one-year prepurchase. Some usage is zero-rated for Microsoft 365 Copilot users.	No charge for the agent service itself. You pay for model tokens and tools (search, code interpreter, storage). Hosted agents add compute per active session.	Free, MIT-licensed (OpenClaw Foundation). You pay for the infrastructure you run plus model and API usage.	Frameworks are free or open source. You pay for compute, models and observability, and engineering time is the biggest line.
COMPLEXITY	Low No infrastructure to run.	Low for prompt agents. Medium-High for hosted agents: containers, SDKs, Azure RBAC and networking.	Medium-High Hosting, sandboxing, secrets and patching are all yours.	Highest You own orchestration, evaluation, security and scaling.
GOVERNANCE & IDENTITY	Power Platform data policies (DLP), environments, Purview audit, data-residency controls. An Entra Agent ID per agent with Agent 365.	An Entra agent identity per agent, Azure RBAC, content filters and prompt-injection guardrails, VNet isolation, bring-your-own storage, tracing.	Data, memory and credentials stay on your infrastructure; sandboxed mode and sender pairing. No native Entra Agent ID, DLP or Purview: governance is built around it.	Whatever you build. Entra identities, Key Vault and Agent 365 registration are available, but nothing is enforced by default.
INTEGRATION & MODELS	Power Platform connectors, MCP servers, native Teams and Microsoft 365 publishing. OpenAI models by default; Anthropic Claude and others selectable.	Foundry model catalog (OpenAI, Anthropic, Llama and more). MCP, OpenAPI and A2A tools. Publishes to Teams and Microsoft 365 Copilot.	Model-agnostic (Claude, GPT, local models). MCP, 20+ chat channels, and a community skills ecosystem.	Unlimited: any model, MCP, A2A, any API. Microsoft Agent Framework, LangGraph, OpenAI Agents SDK, Claude Agent SDK.
SCALE & OPERATIONS	Fully managed by Microsoft. Capacity is enforced monthly, so plan credits ahead.	Microsoft-managed runtime with autoscale. Hosted agents scale to zero between sessions.	Self-hosted. Community and foundation support, no vendor SLA.	Self-managed on AKS, Container Apps or Functions, or on Foundry hosted agents for a managed runtime.
WATCH OUT FOR	Credit consumption is hard to forecast. Deep custom orchestration hits the low-code ceiling.	Costs spread across many meters. "Classic" agents retire on 2027-03-31, so build on the current API.	2026 brought critical vulnerabilities and malicious community skills. Hardening and skill vetting are mandatory, not optional.	Maintenance never ends. Budget for evaluations, model upgrades and security reviews.

Quick decision guide

- Users live in Teams and Microsoft 365, and makers will build → **Copilot Studio**
- Multi-channel or browser agents under your own control → **OpenClaw**, hardened
- Developers need custom logic and model choice on Azure → **Foundry Agent Service**
- The agent is your product or a differentiator → **Custom-built**, hosted on Foundry if you want a managed runtime

Our rule: start with the most managed option that meets the requirement, and move to more control only when a written requirement forces it. Many organizations end up running more than one.

Not sure which fits?

bighatgroup.com

Book a free 30-minute call. We'll recommend the right platform for your workflow, even when it isn't one we deploy.

[Book a discovery call →](#)