

SERVICE BRIEF

Microsoft Foundry.



A fixed-scope AI Foundry accelerator: one grounded agent, five knowledge sources, guardrails in place.

THE PROBLEM & THE APPROACH

AI projects start strong but stall before production. Demos look great in dev; no grounding, no governance, no evaluation and no clear path to ship the agent to the business.

CNNECT's Microsoft Foundry Accelerator delivers a production-ready Foundry workspace with one grounded AI agent, five knowledge sources connected through Azure AI Search, responsible-AI guardrails configured, and the agent shipped to your users via Teams or Copilot Chat.

Key challenges addressed

EIGHT COMMON RISKS

- | | |
|---|---|
| 01 AI experiments stuck in notebooks, never reaching production users. | 02 No grounding strategy: models hallucinate against business questions. |
| 03 Sensitive data flowing through unvetted model endpoints. | 04 Cost spirals as token usage scales without budget controls or alerts. |
| 05 No evaluation suite: quality drift and regressions go undetected. | 06 Content safety, PII detection and jailbreak filters absent or disabled. |
| 07 Multiple teams running parallel Foundry experiments with no governance. | 08 Agents demoed in dev with no production deployment path mapped. |

SERVICE OVERVIEW

From prompt to production in weeks.

The Foundry Accelerator establishes a production-ready Microsoft Foundry workspace. One grounded AI agent is built and deployed against five of your knowledge sources, with Azure AI Search providing the retrieval layer and Foundry Agent Service hosting the runtime.

Responsible-AI guardrails: content safety, PII detection and evaluation are configured up-front. The agent ships to your users via Teams or Copilot Chat, with observability and cost controls in place.

■ SCOPE · KNOWLEDGE SOURCES

Up to 5 sources.

We ground the agent against any five of your knowledge sources, indexed through Azure AI Search. Most accelerators draw a mix from across SharePoint, Fabric, files, public web and Dataverse.

- SharePoint & OneDrive
- Microsoft Fabric Lakehouse
- Files (PDF, Word, Markdown)
- Web pages & Dataverse

■ SCOPE · OUTPUTS

1 agent, shipped.

One production-ready agent built on Foundry Agent Service, governed by safety, PII and evaluation controls, with observability and cost guardrails. Shipped to your users via Teams or Copilot Chat.

- 1 grounded AI agent
- Azure AI Search index
- Safety, PII & evaluation
- Teams / Copilot Chat deployment

TYPICAL TIMELINE

4 – 6 weeks

From kick-off to agent in production.

INVESTMENT

From £30,000

Fixed price for the accelerator scope.

WHAT'S INCLUDED

01 Foundry workspace 02 AI Search index 03 Grounded agent 04 Safety & evaluation 05 Teams deployment

FIVE DELIVERABLES

NEXT STEPS

Discuss your use case or request an accelerator scope walkthrough.

hello@cnnect.com

+44 20 7173 5320

CNNECT.COM