CloudBoostUP

# AI Governance Foundation

Tested platform foundations for governing AI at scale in regulated multi-tenant orgs

**WHAT WE HEARD FROM A PLATFORM LEADER AT A FORTUNE-100 PHARMA**

*"bringing governance into the AI topic in general"*

## Why this. Why now.

Large regulated multi-tenant orgs already run AI. Each business unit picks its own Cognitive Services, Azure OpenAI deployments, and ML workspaces. By the time the audit team asks "what AI runs where", nobody can answer in days. The EU AI Act high-risk obligations applied from August 2026. EMA and FDA released joint AI guiding principles in January 2026. Pharma R&D under Article 2(6) is exempt, but anything embedded in MDR or IVDR product is automatic high-risk.

The Azure primitives exist (Microsoft Foundry hub, Azure ML registries, Purview AI Hub, Responsible AI Dashboard, Azure Policy). The gap is the assembly: how those primitives become a tenant pattern that BU teams adopt without slowing research, while satisfying audit-grade lineage and e-records integrity.

## What you get in 6 weeks

| | |
|---|---|
| **Tenant pattern + IaC** | Hub-spoke landing zone extension for AI workloads. Bicep or Terraform module that slots into your existing landing-zone pattern as a new tenancy type. |
| **Central model registry** | Azure ML registry deployed once at the org-tenant scope. Versioned, immutable, accessible across BUs. End-to-end lineage between training job, dataset, code, environment, and deployed endpoint. |
| **Lineage + classification** | Microsoft Purview connected to ML workspaces. Daily auto-publish of model and dataset metadata. AI Hub policies for classification, sensitivity labels, DLP, audit logging. |
| **Stage gate pipelines** | Azure DevOps or GitHub Actions YAML templates. Promotion from dev to test to prod gated by automated validation, RAI scorecard PDF, manual approval. Research workspaces stay unconstrained. |
| **Audit-pack workbook** | Azure Workbook with per-BU drill-down. Live model inventory, deployment state, scorecard status, data classification. Monthly export-to-PDF for audit responses. |
| **Pharma overlay (when applicable)** | GxP variant: high-risk decision tree (EU AI Act + MDR/IVDR triggers), ALCOA+ audit trail spec aligned to 21 CFR Part 11, context-of-use registration tied to model registry entry. |

## Who it is for

- Platform / Cloud Architecture leads at 5,000-plus-employee regulated orgs (pharma, finserv, energy, public sector).
- Orgs where multiple BUs run Azure AI today, with no central catalog and no Policy-level guardrails.
- Teams under audit pressure: EU AI Act, EMA/FDA, NIS2, DORA, ISO 42001, 21 CFR Part 11, GxP.

- Teams already running on Azure landing zones (Microsoft Foundry, Azure Verified Modules, or equivalent IaC pattern).

## How it works

| | |
|---|---|
| **Week 1** | Tenant baseline. Inventory current AI workloads. Map them to subscriptions, BUs, data classes, regulatory class. Gap report against landing-zone target. |
| **Week 2 to 3** | Hub-spoke deployment. Entra group setup. Purview connection. Azure Policy guardrails (approved registry models, mandatory private endpoints, deny public endpoints, tag taxonomy). |
| **Week 4** | Registry contract. Promotion pipeline templates. RAI scorecard automation tied to gated promotions. Required-artifacts checklist enforced. |
| **Week 5** | Audit-pack Workbook with per-BU view. Lineage query runbook. Compliance Manager template mapping (EU AI Act, ISO 42001, 21 CFR Part 11). |
| **Week 6** | Pilot onboard. One BU's AI workload migrated end-to-end through the new pattern. Handoff session (4 hours live + recorded). Wiki documentation published in your ADO. |
| **Post-engagement** | 30 days of async follow-up included. Adrian on Slack/Teams or email for follow-up questions, escalations, scope clarifications. |

## What you leave with

**Code + config**

✓ Bicep / Terraform module for AI tenant pattern
✓ Azure Policy initiative (enterprise-grade) ready to deploy
✓ Promotion pipeline templates (ADO + GitHub Actions)
✓ Workbook JSON + KQL query library

**Documentation + runbooks**

✓ RAI scorecard generation runbook
✓ Lineage taxonomy + Purview classification rules
✓ Audit-response runbook (regulator question to answer in hours)
✓ Wiki documentation published in your ADO

## What we do not do

*We do not build custom models. We do not act as Data Protection Officer or file regulatory submissions on your behalf. We do not replace your existing SIEM, GRC, or cloud security tools (we integrate with them). We do not propose Microsoft AI Landing Zone verbatim where your enterprise agreement steers a different AI runtime; we adapt the pattern to what your platform team already owns.*


### Adrian Komorek

Founder, CloudBoostUP · Platform engineering for regulated multi-tenant Azure estates
adrian.komorek@cloudboostup.com · +48 695 546 826
cloudboostup.com · linkedin.com/in/adriankomorek


**Engagement model**
Fixed scope · 6 weeks · EU rate band
*Reference price on request*