

From AI pilot to production, inside your business

AI Forward Deployed Engineering services for Telco, Banking, Insurance and E-commerce, on Microsoft Azure and Google Cloud

The problem we solve. Most enterprises already have AI licenses, a cloud platform and a few pilots. Few have AI running in production and changing how work gets done. The gap is rarely the model: it is data readiness, integration with legacy systems, security and compliance, evaluation, and adoption by the people who must use it. Our Forward Deployed Engineers (FDEs) embed with your teams and close that gap end to end.

How we are different

- Embedded, not advisory. Our engineers sit with your business and IT teams, inside your cloud tenant and data.
- Production or nothing. Every engagement ends in a running system with evaluation, monitoring and owners, not a slide deck.
- Data-first. We start from your data estate (warehouse, lakehouse, CRM, core systems), because that is where most AI projects stall.
- Dual-cloud. Certified on Microsoft Azure and Google Cloud, so we work with the platform you already run.
- Regulated-industry ready. Built-in patterns for KVKK/GDPR, BDDK/SEDDK and similar regulatory expectations, audit trails and human-in-the-loop controls.
- We leave capability behind. Knowledge transfer and co-building, so your team can run and extend what we ship.

Service catalog

#	Service	What you get	Typical length
1	AI Opportunity Sprint	Workshops with business and IT, use-case backlog scored on value, feasibility and data readiness; 2-3 prioritized use cases with business case.	2-3 weeks
2	AI Data Readiness	Assessment of data sources, quality, governance and access; target architecture; build of the first AI-ready data products (lakehouse, semantic layer, vector index).	4-8 weeks
3	Rapid Prototype (Proof of Value)	Working prototype on your real data, measured against agreed KPIs, demoed weekly to end users.	3-6 weeks
4	Production Deployment	Hardening of the prototype: integration with core systems, security, evaluation pipelines, observability, CI/CD (LLMOps), go-live and hypercare.	6-12 weeks
5	Agentic Workflow Engineering	Multi-agent systems and tool-using agents (MCP, APIs) that automate end-to-end processes with human approval steps.	8-16 weeks
6	Conversational AI & Contact Center	Voice and chat assistants, agent-assist for human agents, call summarization and quality scoring.	6-12 weeks
7	Text-to-Insight / Talk to Your Data	Natural-language questions over your warehouse with governed semantic layer, Text2SQL and answer validation.	4-8 weeks

#	Service	What you get	Typical length
8	AI Governance, Risk & Evaluation	Responsible-AI policy, model and prompt inventory, red-teaming, guardrails, regulatory documentation, eval benchmarks.	3-6 weeks
9	FDE-as-a-Service (Retainer)	Dedicated FDE squad on a monthly retainer: continuous delivery of new use cases and optimization of live ones.	Monthly
10	AI Academy & Adoption	Role-based training for executives, business users and engineers; champions network; adoption metrics.	Ongoing

Our delivery model

Phase	Discover	Prototype	Productionize	Scale
Goal	Find where value is	Prove it on real data	Make it reliable and safe	Repeat and extend
Output	Prioritized use cases, business case	Working demo + KPI baseline	Live system, evals, monitoring	Reusable modules, trained team

Engagement models

- Fixed-scope sprint: Opportunity Sprint, Readiness Assessment, Prototype.
- Outcome-based project: production deployment with agreed KPIs and acceptance criteria.
- FDE squad retainer: 1 lead FDE + 1-2 AI/data engineers + part-time architect, billed monthly.
- Pod embedded in your team: our FDEs work inside your squads and transfer skills as they build.

Squad roles

Role	Responsibility
Lead FDE	Owens business problem, scoping, stakeholders and outcome
AI Engineer	Agents, RAG, prompts, evaluation, guardrails
Data Engineer	Pipelines, data products, semantic layer, vector stores
Cloud / MLOps Engineer	Landing zone, security, CI/CD, monitoring, cost
Solution Architect (part-time)	Target architecture, integration, non-functional requirements

Start with a 2-week AI Opportunity Sprint. We bring a prioritized, costed AI roadmap and one use case ready to prototype. Contact: [email] · [website] · [phone]