Lead AI Engineer – Agentic AI (Azure)

EXPERIENCE **4–6 years** | LOCATION: **India - Remote** | EMPLOYMENT **Full-time**

About AZ4C

AZ4C is a technology company focused on helping organisations apply cloud, data and AI to real business challenges. We design and deliver practical solutions on the Microsoft ecosystem, combining strong engineering with close client collaboration to move ideas from proof of concept to reliable production systems.

Role Summary

Lead the build of production worker agents on Microsoft Agent Framework and Microsoft Foundry: long-running, event-driven agents that act in business systems. Own the build from design to production, and train, guide and mentor junior engineers while collaborating with solutions architects and clients.

Key Responsibilities

- Build worker agents on Microsoft Agent Framework and Foundry Agent Service that plan, act through tools and complete business tasks end to end.
- Design long-running, event-driven execution: durable state, checkpoints, retries and timeouts, triggered by schedules, events and queues (Durable Functions, Container Apps, Service Bus, Event Grid).
- Build tools and connectors (MCP servers, APIs) into Fabric, Azure SQL, Cosmos DB, SharePoint and line-of-business systems, with idempotency and locking so parallel agents and retries never double-act.
- Design human-in-the-loop approvals, escalations and handoffs for decisions agents should not make alone.
- Own production readiness: end-to-end tracing (OpenTelemetry, Application Insights, Foundry tracing), structured logging, an audit trail of every action, cost per task and run budgets.
- Run task- and step-level evaluations with regression gates in CI/CD; automate infrastructure with Bicep or Terraform.
- Apply security by design: agent identity (managed identities, Entra ID), least-privilege tools, Key Vault, private networking, content safety and API Management as the AI gateway.
- Break down work, review code and mentor AI engineers; support presales with feasibility checks, estimates and proof-of-concept builds.

Required Skills & Experience

- 4–6 years in software, data or ML engineering, including 1 year shipping agents that take action in real systems, not only chat or Q&A.
- Hands-on Microsoft Agent Framework (or Semantic Kernel or AutoGen) and Microsoft Foundry, including Foundry Agent Service.
- Azure application services: Durable Functions, Container Apps, Service Bus or Event Grid, API Management.
- Building tools and connectors for agents, including MCP servers.
- Distributed systems fundamentals (concurrency, idempotency, retries, failure recovery) and production habits (tracing, logging, debugging from traces, controlling cost).
- Strong Python (APIs, async, testing, CI/CD) and working SQL.
- Experience leading or mentoring engineers and explaining trade-offs to clients.

Preferred Qualifications

- Bachelor's degree in Engineering, Computer Science, IT or a related field.
- C# and .NET alongside Python; Microsoft Fabric (data agents, OneLake) or Databricks.
- Copilot Studio, Microsoft 365 Copilot extensibility, and A2A or multi-agent patterns in production.
- Microsoft Azure AI or developer certifications.

Success Measures

- Worker agents go live, meet agreed acceptance criteria and run reliably without manual intervention.
- Every release passes evaluation gates, failures are traceable, and cost per task stays within budget.
- Engineers on the team deliver independently, and clients trust the build quality.