

Kyndryl Agentic AI Service

Operationalizing Autonomous AI at Enterprise Scale

Kyndryl's Agentic AI Service enables organizations to design, govern, and operate autonomous and semi-autonomous AI agents as trusted components of enterprise workflows. Built on Kyndryl's operational expertise, the service embeds governance, observability, and responsible AI controls to help organizations move confidently from AI pilots to production.

The Challenge

Organizations struggle to scale AI safely due to fragmented pilots, limited visibility into AI behavior, compliance risks, and integration challenges. As AI autonomy increases, so does the need for strong governance, transparency, and control.

The Kyndryl Approach

Kyndryl bridges the gap between experimentation and AI operations through an enterprise-grade Agentic AI Framework. The framework orchestrates AI agents, enforces policy-as-code, and provides real-time observability, ensuring AI operates securely and predictably across complex environments.

Key Capabilities

Secure orchestration and lifecycle management of AI agents. Built-in governance and policy-as-code enforcement. Real-time observability and auditability. Responsible AI guardrails and human-in-the-loop controls. Seamless integration across enterprise platforms.

Business Benefits

The service accelerates time-to-value, improves operational efficiency, strengthens governance and compliance, and enables scalable AI adoption without added complexity. Organizations benefit from better customer and employee experiences while maintaining trust and control.

Why Kyndryl

Kyndryl brings unmatched experience operating mission-critical systems. Unlike AI-only vendors, Kyndryl designs agentic AI for real-world conditions, ensuring security, resilience, and long-term sustainability.