

The 2025 State of Enterprise AI Report

A Professional Toolkit Resource by TheConsultAI

This report provides benchmark insights collected from 150+ enterprise AI deployments across 15 countries between 2023 and 2025. It highlights success patterns, budget distributions, MLOps metrics, and key failure points.

Key Benchmark Findings

- **Deployment Rate:** Only 12% of enterprise AI proofs-of-concept (PoCs) successfully reach production within 12 months. The remaining 88% remain in sandbox phase due to integration complexity or cost.
- **Average Timelines:** Accountable consultancies using pre-built accelerators average 12-16 weeks to production, compared to 18-24 months for internal builds.
- **Failure Triggers:** The #1 cause of deployment failure is data quality (34%), followed by lack of business stakeholder alignment (26%), and unpredicted API token costs (20%).

Budget Allocation Breakdown

Infrastructure & LLM APIs	35%	Vector databases, dedicated hostings, token billing
Data Integration	30%	Cleanings, vectorizing pipelines, API middleware
Advisory & Strategy	15%	KPI mappings, board scoping, process mappings
Training & Alignment	20%	User UAT, operations training, change management

The MLOps Maturity Model

Successful production deployments implement the following three-stage pipeline strategy:

- Continuous Monitoring: Real-time tracking of prompt latency, user feedback, and API costs.
- Guardrails: Intercepting prompts and outputs to prevent data leaks, compliance breach, or hallucinations.
- Iterative Feedback Loops: Logging low-confidence outputs to feed back into model prompt fine-tuning cycles.