

NTT DATA Managed AI Landing Zone

The Challenge

- AI initiatives remain stuck in isolated PoCs: duplicated setup work, inconsistent security controls, and limited reuse across teams and use cases.
- Missing standardized guardrails and end-to-end transparency (cost, latency, policy hits, audit trail) prevents scaling from “pilot” to “production” and makes it impossible to measure ROI.
- Slow onboarding and unclear operating model: teams spend weeks wiring infrastructure, access, and integrations instead of building business value.
- Compliance: Ad-hoc setups lack the comprehensive, audit-ready evidence required to pass strict regulatory reviews like the EU AI Act.

The Solution

- Transform standard cloud infrastructure into a production-ready “AI Factory” equipped with agentic runtimes, data pipelines, and governance controls. Accelerate the transition from pilot to production by providing a secure, vendor-neutral foundation designed for scaling autonomous agents while considering EU AI Act compliance.
- A ready-made platform with blueprints for RAG pipelines and agentic workflows, moving beyond basic cloud infrastructure setup.
 - Dedicated runtimes and orchestration logic to support autonomous agents that can plan tasks and execute tools, not just chat.
 - Includes an abstraction layer (LLM Gateway) that allows seamless switching between models to prevent lock-in and manage costs.
 - “Compliance-by-Design” technical guardrails that intercept prompts for PII and topic blocking before they reach the model.
 - Provides granular cost visibility per token and request to prevent bill shock and enable chargeback models for high-volume usage.
 - Accelerated delivery using validated MVP and the underlying landing zone infrastructure within weeks.
 - Continuous Governance: “Managed AI Landing Zone” service that handles ongoing model updates, security policy evolution, and operational optimization.
 - Hardened Operating Model (RACI): a clear, federated operational blueprint: the Central Platform Team owns the gateway, guardrails, and model onboarding, while Domain Squads own prompt design and business outcomes.

Key KPIs & Benefits

- Time to onboard a new team / use case: days not weeks
- Time to first prototype / MVP: 2–6 weeks depending on scope
- % use cases using standardized guardrails and logging: target >90%
- Cost per request / use case + budget variance; cost transparency & chargeback readiness
- Operational KPIs: latency, availability, incident MTTR, trust signals (policy violations blocked, PII leak detections)

Team Composition

Joint delivery team (NTT DATA + Customer) aligned to our flexible ecosystem. Client can enter at any point based on their maturity:

Mandatory roles:

- **Phase A “Hackathon” / POC**
 - Agentic AI Architect (platform & reference architecture, patterns, scalability)
- **Phase B “MVP”**
 - Product Owner / Project Lead (client + NTT DATA)
 - Agentic AI Engineers (agents, workflows, RAG, tool integration)
- **Phase C “Full Scale”**
 - AI/ML Ops / DevOps Engineer (CI/CD, deployments, monitoring, run)
 - Security / Compliance lead (guardrails, policy enforcement, auditability)
 - Data Engineer / Integration SME (enterprise data sources, connectors, access patterns)

Optional roles:

- Business / Process Analyst (translate business requirements into technical specification)
- Financial Analyst (cost control and chargeback models)

Timeline

1. Week 0-2
Project Preparation

2. Week 3-5
Platform-Layer-
Integration

3. Week 6-9
Workload-Modul
Development

4. Week 10-11
CI/CD-Pipeline
Implementation

5. Week 12-14
Environment
Deployment

6. Week 15-16
Workload
Deployment & Tests

7. Week 17
Rollout & Go-Live

8. Week 18
Handover & Training