

AI Landing Zone Implementation

Motivation

Enterprises adopting AI need more than a model API key. They need a secure, governed, and scalable foundation that integrates AI services into their existing cloud platform. Organizations that already operate a CAF Landing Zone based on the Cloud Adoption Framework have the ideal starting point — but extending it for AI workloads requires specialized expertise in Azure AI Foundry, private networking for AI services, AI-specific governance policies, and secure connectivity to data sources.

This offering delivers a structured, three-phase approach to establish an enterprise-ready AI Landing Zone pattern on top of your existing Cloud Adoption Framework (CAF) foundation. It begins with a focused assessment of your current landing zone, identifying integration constraints and governance gaps for integrating the AI Landing Zone pattern. This is followed by a detailed architecture and design phase that defines the target pattern aligned with Azure AI Landing Zones.

The final phase is a hands-on MVP implementation, including a five-day hackathon to deploy:

- an AI Gateway Landing Zone based on Azure API Management for centralized model access, policy enforcement, traffic control, and content safety, and
- a Foundry Landing Zone for hosting AI applications, services, and data sources.

The outcome is a production-ready AI Landing Zone pattern implementation, with an AI Gateway Landing Zone and a Foundry Landing Zone, active governance, monitoring, and security controls, and a reference team fully onboarded. The setup establishes the operational foundation for GenAIOps, including quota and cost management, model lifecycle governance, content safety configuration, and usage monitoring. Development teams receive a solid, enterprise-grade platform to build and operate AI applications at scale—enabled by clear guardrails, centralized control, and built-in operational readiness.

What We Bring

With more than 20 years of experience delivering enterprise software and cloud architectures for some of the world's largest enterprises, PRODYNA brings a unique combination of cloud platform depth and AI expertise. We have built and operated CAF Landing Zones at scale and know what it takes to extend them securely for AI workloads. Our strengths in this domain include:

- **Proven CAF & IaC Expertise:** We have extensive experience with CAF Landing Zone architectures and Terraform-based infrastructure and have applied them across multiple enterprise engagements.
- **Enterprise Security First:** We design every AI integration with zero-trust principles: private networking, managed identities, Azure Policy enforcement, and no API keys — by default, not as an afterthought.
- **GenAIOps from Day One:** We bring operational readiness into the architecture from the start — quota management, model lifecycle governance, content safety, and cost controls built into the platform, not bolted on later.
- **Simplicity First:** We start with the simplest viable setup and evolve it only as needed, avoiding over-engineering and keeping your platform maintainable.
- **Knowledge Transfer:** We work alongside your platform and AI teams throughout all three phases so that they own and can extend the result independently.



Benefits

- **Reduced risk:** Three phases align governance, architecture, and implementation before rollout.
- **Solid foundation:** Decisions are based on a clear operating, governance, and cost model.
- **Fast delivery:** The MVP includes a hands-on 5-day hackathon and a working setup.
- **Proven approach:** You benefit from PRODYNA's enterprise landing zone experience.
- **Scalable control:** Onboard projects via Foundry Landing Zones while the AI Gateway Landing Zone centralizes model access and chargeback.
- **Strong security:** Private networking, managed identities, and policy-based controls by default.



Quick facts

- **Duration:** ~1 month across three phases
- **Prerequisite:** Existing CAF Landing Zone with shared platform services
- **Model-agnostic:** Any model from the Azure AI Foundry catalog
- **Modular:** Enable only the AI services and data sources you need



What You Need

To make the best use of this offer and permit a fast and efficient start, you will need:

- An existing Azure Landing Zone setup based on the CAF hub-spoke architecture.
- An existing Azure contract e.g. Enterprise Agreement (EA), or any other Cloud Service Provider (CSP).
- Availability of your experts (e.g. Azure platform team, network, DNS, IAM, security, and future AI team).

What You Get

In about 1 month across three phases, we assess your current CAF Landing Zone, define the AI Landing Zone pattern integration approach, design the target architecture, and implement the AI Landing Zone pattern consisting of an AI Gateway Landing Zone and a Foundry Landing Zone. By the end, you will have:

- A clear governance model for AI covering security, compliance, ownership, cost allocation, and model strategy
- A validated AI Gateway and Foundry landing zone architecture, including model catalog, guardrails, policies, and secure connectivity
- A deployed AI Gateway Landing Zone and Foundry Landing Zone
- A reference team onboarded and enabled to operate and extend the platform
- A GenAIOps baseline with runbooks for continued operations
- Full documentation and a scalable pattern for onboarding additional teams and projects

PHASE	ACTIONS
Assessment (3 to 5 Days)	• Define the AI governance target state (security, compliance, ownership) with key stakeholders • Assess the existing CAF Landing Zone setup (firewall, DNS, IAM, policies, shared services, and connectivity) • Decide the operating model for centralized vs. decentralized AI usage • Define cost and billing principles (allocation, chargeback/showback boundaries) • Define model and provider strategy, and align scope, expectations, dependencies, and risks for the next phases
Design (7 to 10 Days)	• Produce target architecture for the AI Gateway Landing Zone, Foundry Landing Zone, and support components • Define how the AI Landing Zone pattern (AI Gateway Landing Zone + Foundry Landing Zone) integrates with existing platform services and management groups • Define AI gateway controls in APIM (token tracking, throttling, routing, and circuit breaker policies) • Specify Azure Policy guardrails, RBAC mappings, managed identity patterns, and approval workflows • Engineer private networking blueprint (private endpoints, private DNS, NSGs, UDR/firewall routing, and restricted outbound) • Prepare implementation backlog, Terraform module plan, and operating model for go-live
MVP (5 Days)	• Build the AI Gateway Landing Zone and Foundry Landing Zone in Terraform in a 5-day hackathon • Deploy core services (Foundry, APIM, private connectivity, identities, policy assignments, and monitoring baseline) • Enable governance and operations controls (quotas, budgets, diagnostics, alerts, content safety, and security posture management) • Execute onboarding of one reference team with end-to-end validation and private-connectivity checks • Hand over runbooks, operations checklist, and repeatable scale-out pattern for additional teams and projects

Note: This is our standard methodology. We are happy to modify and adapt to the specific needs of your organization.

Get Started

To learn about pricing and how to get started, please contact info@prodyna.com

About PRODYNA

PRODYNA designs and operates AI-enabled business systems serving the needs of corporations and enterprises across the European continent. For more information, please visit www.prodyna.com.