

— EA-WIZARD · ARTIFICIAL INTELLIGENCE

Native, optional, granular AI

How EA-Wizard weaves artificial intelligence through every architecture module — while leaving you in full control of provider, scope and data residency.

AI in EA-Wizard is not a bolt-on assistant. It is a thread woven through every module — yet it stays entirely in your hands: choose the provider, activate it service by service, and switch to an air-gapped deployment for the most sensitive contexts.

18

distinct AI services

3

supported providers

9+

covered domains

100%

optional & granular

Native, not bolted on

Most EA platforms added AI recently, as isolated chat assistants sitting beside the real data. EA-Wizard takes the opposite path: its AI services have direct, governed access to the dependency graph, the application portfolio, business capabilities and processes. The consequence is concrete — the AI answers questions about **your** architecture, not generic ones.

This grounding is what turns AI from a novelty into an operational lever: every suggestion is traceable to a node, a control or a metric an architect can verify.

The 18 AI services

Governance & decision

- **AI Demand Qualification** — qualifies new demands automatically: scope, complexity, standards alignment.
- **AI Standards Checker** — evaluates each demand against your standards catalog and golden paths.
- **ADR Draft & Enhance** — drafts and enriches Architecture Decision Records (alternatives, consequences, context) for architect validation.
- **AI Governance Insights** — analyses ARB cadence, SLA gaps and the waiver registry.

Augmented mapping

- **AI Capability Mapping** — proposes an L1→L3 business-capability hierarchy from your application landscape.
- **AI Application Intelligence** — portfolio health: criticality, fragility, recommendations.
- **AI Process Intelligence** — detects gaps between modelled processes and reality.
- **AI Data Intelligence** — surfaces data-catalog gaps and proposes quality rules.
- **AI Integration Intelligence** — detects integration anti-patterns and suggests modernisations.
- **AI Tech Tower Intelligence** — watches the technology radar, alerts on obsolescence.
- **AI Infrastructure Intelligence** — analyses topology, detects single points of failure, proposes hardening.

Impact & risk

- **AI Impact Narrative** — turns a blast-radius simulation into an executive-ready report.

- **AI Risk Analysis** — identifies critical components and fragility points.
- **AI Compliance Assistant** — assists control evaluation: requirement interpretation, evidence generation.

Steering

- **AI Roadmap Assistant** — wave-aware roadmap suggestions that respect dependencies.
- **AI Pulse Radar** — interprets each project’s pulse and warns on anticipated drift.
- **AI KPI Recommendations** — nightly per-metric actions (AVS, ADI, GV, ACR, SAR, EAI).
- **EA Chat** — a transversal conversational assistant in natural language.

Multi-provider architecture

EA-Wizard connects to three families of models through a centralised dispatcher (`AiProviderClient`). Your API keys are encrypted with AES-256-GCM, like every other secret in the platform. You can switch provider without redeployment.

PROVIDER	NOTES
Anthropic Claude	Claude 3.5 Sonnet, Opus. Your API keys, your quota.
OpenAI GPT-4	GPT-4 and variants via the OpenAI API. Your keys, your billing.
Ollama (local)	Open-source models (Llama, Mistral, Mixtral...) running locally. Zero outbound data.

Air-gap mode — sovereign AI

In Ollama air-gap mode, AI services run entirely inside your infrastructure. No outbound calls, no architecture data sent to a third party, no cloud account to approve.

- Open-source models of your choice (Llama, Mistral, ...).
- Runs on your GPUs or CPUs, in your datacentre.
- No external connectivity required for AI.
- Full audit trail — no offshore black box.

WHY IT MATTERS

For supervised banks and the public sector, air-gap mode is often the decisive differentiator: it lets an organisation benefit from architecture AI without a single byte of sensitive topology leaving its walls.

Granular activation

EA-Wizard makes AI explicitly optional. Every AI service exposes an activation flag per tenant. If a client allows AI for the Roadmap module but refuses it for Compliance, the engine honours that. If a client refuses AI globally, no calls are issued — and no critical feature is degraded.

- **Per-module activation** — every AI service can be turned on or off independently.
- **Per-tenant activation** — each organisation picks its own level of AI usage.
- **Provider choice per tenant** — one tenant on Claude, another on air-gapped Ollama, without collision.

AI amplifies the architect; it never replaces the architect's judgement, and it is never a precondition for using the platform.