iNikola

ENTERPRISE AI ENGINEERING · SECURE AI ENABLEMENT

Secure enterprise AI, built for *production.*

We take enterprise AI from a blocked use-case to a security-approved system in production — engineered, governed and deployed inside the boundary your data is already allowed to live in.

contact@inikola.com +91 93050 09726 Hyderabad, India inikola.com

AI adoption is not the hard part. *Approval is.*

Enterprise AI projects rarely fail because a model cannot generate text. They stall when the organisation has to answer where the data goes, who is allowed to see what came back, and who owns the system on the day it breaks.

WHERE PROJECTS STOP	WHAT WE PUT IN ITS PLACE
Blocked use case	An approved data path security and legal can sign off
Shadow AI	One gateway under company identity, policy and audit
No audit trail	Per-query evidence: who asked what, of which model, over which documents
Permission leakage	Retrieval inheriting the permissions the source already carries
Model sprawl	Policy-based routing behind a single control plane
Stalled proof-of-concept	A production system that is owned, monitored and supported

The enterprise conversation has moved from *"can AI work?"* to *"can AI be approved?"* That layer — identity, permissions, integration, governance and production ownership — is what we build.

End to end, and *secure at every step.*

01

AI Strategy & Readiness

Use-case discovery, data sensitivity mapping, architecture options, ROI and the approval plan.

02

Enterprise AI Engineering

GenAI applications, ML systems, RAG, multimodal AI, copilots, APIs and integration.

03

Agentic AI & Automation

Tool-using agents, workflow automation, multi-agent patterns and human approval loops.

04

Private & Secure AI

Secure AI gateway, SSO and RBAC, DLP, controlled egress, model routing and audit.

05

AI Security Engineering

Threat modelling, prompt-injection defence, tool and MCP security, red-teaming, release control.

06

Enterprise Knowledge & RAG

Permission-aware retrieval, vector and hybrid search, citations, document intelligence.

07

AI Governance & Quality

Evaluation sets, guardrails, risk testing, human-in-the-loop and auditability.

08

LLMOps & Managed AI

Monitoring, continuous evaluation, model updates, cost tracking and incident support.

The secure AI *control plane.*

We separate enterprise control from the model deliberately. The model will change — probably several times. Your identity, policy, permission and audit requirements will not.

Enterprise users	Employees, teams and applications
Identity & access	SSO · MFA · RBAC / ABAC · source permissions
iNikola secure AI gateway	Policy enforcement · routing · isolation · egress control
Data & safety controls	DLP and PII controls · prompt and response policy · audit logs
Permission-aware RAG & agents	Retrieval honouring source ACLs · tools · orchestration
Enterprise data sources	SharePoint · files · databases · APIs · tickets · code · SOPs
Model abstraction layer	Local SLM · open-weight LLM · VPC model · approved enterprise API
Evaluation & observability	Quality · security tests · cost · latency · risk · usage

The model becomes replaceable. Governance, integration, permissions and security become the part that compounds — which is why the second use case costs a fraction of the first.

Private does not *have to mean on-premise.*

We use the minimum architecture your risk profile actually requires — then design so the boundary can move later without rebuilding the application.

BOUNDARY	FITS WHEN	OUR ROLE
Approved enterprise API	Enterprise terms already cover the data class	Application, access model, data-flow controls, evaluation, governance
Customer VPC / private endpoint	Cloud-first estate needing private connectivity	Secure architecture, routing, retrieval, IAM, observability
Private cloud / on-premise	Regulation or genuinely sensitive IP requires it	Private inference stack, integrations, governance, operations
Disconnected / air-gapped	No external egress permitted at all	Local model, retrieval and tooling with isolated controls
Hybrid / policy router	Different data classes need different answers	Policy routing behind one gateway, with one audit trail

Private Knowledge Copilot

A secure assistant over your own documents, deployed inside your boundary. The fastest route to a governed AI system in production — and a reusable control plane the next five use cases build on.

CONTROLS

- Permission-aware retrieval inheriting source access control
- Citations on every answer, back to the originating document
- SSO, MFA and RBAC through your existing identity provider
- DLP and policy controls on prompts and responses
- Complete audit trail per user, per query, per document
- Model-agnostic: local SLM, open-weight, VPC or approved API

BY VERTICAL

NBFC / lending	Credit & policy copilot
Wealth / AMC / PMS	Compliance & research copilot
Insurance / TPA	Claims & policy copilot
GCC	Engineering knowledge copilot
Pharma	Quality & SOP copilot
BPO / KPO	Client-isolated delivery copilot
Manufacturing	Engineering troubleshooting copilot

PRODUCTS & ACCELERATORS — AVAILABLE TODAY

Posting Expert

AI marketing and social publishing run as an automated agentic workflow.

EvalSea

Evaluation, assessment and verification workflows, automated end to end.

HR Bot

Screening, interview and hiring workflow automation with human checkpoints.

One workflow. One security review. *Four weeks.*

Week 1

Use-case discovery

Workflow, data sensitivity, ROI owner, approval blockers, success criteria.

Week 2

Architecture & risk plan

Boundary, access, egress, model choice, data flow, TCO, control matrix.

Week 3

Private pilot build

Retrieval or agent, authentication, guardrails, logging, integration, evaluation.

Week 4

Security review & handoff

Demo, control evidence, test results, business KPI, rollout plan.

COMMERCIAL LADDER

STAGE	OFFER	WHAT IT IS
Discovery	AI Approval & Architecture Review	30 minutes on one blocked workflow. No charge.
Assessment	Readiness & Security Assessment	Paid engagement producing architecture and approval plan.
Pilot	Fixed-scope Secure AI Pilot	One workflow, real data, controlled users, measurable criteria.
Production	Departmental Deployment	Rollout, integrations, hardening, security sign-off.
Platform	Multi-use-case AI Platform	Shared control plane across business units.
Operate	Managed AI Operations	Monitoring, evaluations, updates, incident support.

Infrastructure and third-party licences are procured in your own accounts on your own commercial terms. We charge for architecture, implementation, controls, integration and operations — never for hardware resale. You own the code, the infrastructure and the documentation.

07 · WHERE WE WORK

- NBFC & digital lending
- Wealth, AMC, PMS & AIF
- Insurance & TPAs
- Global capability centres
- Payments & fintech infrastructure
- Pharma, biotech & CRO
- IT services, BPM & KPO
- Manufacturing & engineering

Which AI workflow does your business want that *security has not approved?*

Bring us that one. In thirty minutes we will map the data path, name the blockers, and tell you the architecture that clears them — or tell you it is already fine as it stands.

contact@inikola.com

+91 93050 09726 · Hyderabad, Telangana, India

iNikola