Governed AI agents and models, running in your environment.

Agents do the groundwork; your people make the call, with evidence on every decision. One control plane for models, knowledge, data and governance, and you keep what is built.

Seven layers, one boundary

1. **People and tools**: chat, email, portals or your applications through the API.
2. **AI agents** look it up, check the facts, work it through and stop for a person.
3. **Control plane**: access by role and team, policies, approval routing, evaluation gates.
4. **Inference gateway**: every model call authenticated, logged, metered and routed.
5. **Self-hosted models**: open-weight and custom models, classifiers and embeddings.
6. **Knowledge and Data Hub**: governed documents with lineage; records read where they live.
7. **Your systems of record**, written to only after approval.

Deployment options

- **Cloud**: your tenancy, your region, your keys.
- **Hybrid**: sensitive work on your servers, the rest in your cloud.
- **On-premises**: your data centre, behind your firewall.
- **Air-gapped**: no outside connection, self-hosted models only.
- **Compute**: CPU and GPU serving on your cloud, your servers or air-gapped hardware.

Residency, sovereignty and blast radius usually decide the option before the model is chosen. The same platform, governance and evidence trail follow every option.

Integration

- REST API across the platform, with Python and TypeScript SDKs.
- An OpenAI-compatible inference gateway, so existing applications can call your models.
- MCP, so models are reachable as tools.
- Agents name no vendor: they connect to the ERP, CRM, core banking or property system you already run.
- Records are read at query time, not copied into the platform.

Agents and models

- **1,141 agents** across 11 departments, each with its rules, settings and approvers written down on its own page.
- **77 test runs**, one per department per industry: every run reached its approval step or finished.
- Deterministic and generative models, open-weight or custom, tuned on your data. Route by cost, latency, residency or capability.

Governance across every layer

- **Access**: role- and attribute-based permissions with deny-overrides; isolated projects per team and business line.
- **Approvals**: agents stop for a named person where their rules say so; limits are settings, not code.
- **Evidence**: reason codes, cited sources and a signed trail from each output back to its inputs.
- **Observability**: every inference logged and metered; drift and quality checked continuously, with evaluation gates before promotion.

How an engagement runs

1. **AI assessment**: the decision, the data behind it, and whether AI can honestly move the number.
2. **Pilot**: one use case or department, on your data, with a target agreed at the start.
3. **Production**: deployed into your tenancy, governed and instrumented, with your people approving what matters.

What you keep

- Models tuned on your data, and their weights.
- Your governed document collections.
- The decision history and audit trail.
- Which artefacts transfer, in what format and how often, agreed in writing before the first deployment.