PILLAR P4 | Weeks 1-10 - first governed agent live by week 4 | EX - Data, Automation & AI

Atlas Intelligence Plane

Lakehouse - Vector + RAG - Agent Runtime - Model Gateway

Lakehouse, vector store, feature ops and an agent runtime - one governed plane.

MODEL TIME-TO-VALUE	TOKEN SPEND PER TASK	EVAL PASS RATE
4	58%	> 94%

MANIFESTO

Why this pillar exists

Most enterprises have data, models and agents - but in three different stacks with three different governance stories. Atlas is one plane: lakehouse, vector store, feature ops, model gateway and agent runtime, all governed end-to-end with the same identity and policy fabric.

KPIS

Outcomes we contract to

- Model time-to-value: 4
- Token spend per task: 58%
- Eval pass rate: > 94%

DELIVERABLES

Everything that ships

> Governed lakehouse
Iceberg / Delta with row-level security, lineage, contracts, freshness SLOs.

> Vector + hybrid retrieval
Per-tenant indexes, freshness gates, citation enforcement, eval harness.

> Feature & eval ops
Online/offline parity, drift detection, golden sets, regression suites.

> Model gateway
Multi-model routing (OpenAI / Anthropic / Gemini / Mistral / OSS) with cost + policy controls.

> Agent runtime
Planner/worker/critic loops, tool registry, deterministic fallbacks, tracing.

POD COMPOSITION

Who shows up

- Data Principal
- ML Platform Engineer
- Agent / RAG Engineer
- Governance + Lineage Lead

TIMELINE

Weeks 1-10 - first governed agent live by week 4

1. **Weeks 1-3 - Plane foundations**
   Lakehouse, identity, lineage, contracts, model gateway, observability.
2. **Weeks 3-6 - First governed agent**
   RAG + tools + evals + tracing live for one P&L line; shadow mode then live.
3. **Weeks 6-10 - Multi-agent + cost optimisation**
   Planner/worker/critic, model routing, cost-per-task scorecard.

FLAGSHIP PODS

Squads we drop in

- Lakehouse Foundations Pod
- RAG + Vector Pod
- Model Gateway Pod
- Multi-Agent Orchestration Pod

PARTNERS

Hyperscalers & ISVs we orchestrate

Databricks - Snowflake - Microsoft Fabric - OpenAI - Anthropic - Google Vertex - Mistral - Qdrant - Pinecone

FAQS

What buyers actually ask

Q. Do you replace Databricks / Snowflake / Fabric?
A. No - we sit on top, orchestrating governance, vector and agent runtime across whichever lakehouse you've chosen.

Q. Can we bring our own models?
A. Yes. The model gateway is multi-vendor and supports OSS (Llama, Mistral, Qwen) on your private GPUs alongside hosted APIs.

Q. How is this different from a chatbot rollout?
A. Every agent is contracted to a P&L line, has an eval harness, a cost ceiling, an escalation path, and a regulator dossier.

Q. Sovereign deployment?
A. Yes - EU, UK, US, GCC and on-prem GPU deployments are first-class. Region pinning is a policy, not a project.

Commission Atlas Intelligence Plane

Outcome-priced. Sovereign by default. Refund-backed if the contracted KPI isn't hit.

Reply to this PDF: alfaxprienz.com/stack/atlas-intelligence-plane#commission