

The Operating System for Business AI

Platform Overview

Datafi is a vertically integrated AI platform that lets enterprises deploy AI agents and workflows on their own data, with governance, security, and observability built in from the start. The platform connects to your data where it lives, orchestrates across any LLM, and enforces policy at the point of use. No code required. Production-ready in weeks, not years.

Platform Components

Studio **Build**

No-code agent and workflow builder. Visual drag-and-drop composition, natural language interface, and human-in-the-loop approval gates. Grounded in full enterprise data context.

Control Tower **Govern**

Centralized AI governance and observability. Real-time monitoring across all agents and workflows. Immutable audit trails, centralized policy enforcement, and full data lineage.

Orchestrate **Run**

AI runtime and control plane. Multi-step agent execution with model routing, cost control, and performance optimization. Connects models, tools, data, workflows, and guardrails.

Sentinel **Secure**

AI-aware cybersecurity layer. Continuous risk scoring, prompt injection detection, data leakage prevention, and zero-trust enforcement with attribute-based access controls.

How It Works

1

Connect Your Data

Federate across SQL Server, Snowflake, Databricks, Azure Synapse, and more. No ETL.



2

Build Agents & Workflows

No-code Studio to compose AI agents that query, analyze, and act on enterprise data.



3

Route to the Right Model

Direct each task to the right LLM. OpenAI, Azure OpenAI, Bedrock, or Cortex.



4

Govern & Secure

Policy enforcement, lineage tracking, risk scoring, and audit trails. Built in.

Enterprise Ready

- LLM agnostic. Any model, any provider.
- SOC 2 Type II certified data centers
- SSO via SAML 2.0, OIDC, SCIM provisioning
- SaaS, private cloud, on-prem, or hybrid deployment
- TLS 1.3 in transit, AES-256 at rest
- Attribute-based access control (ABAC)
- Docker and Kubernetes ready
- Customer AI data never trains other models

Databases: Snowflake, PostgreSQL, BigQuery, MS SQL Server, MySQL, Redshift, Databricks, Azure Synapse, Oracle, DB2

Identity: Auth0, AWS Cognito, Microsoft Entra ID | Storage: AWS S3, Azure Blob, GCS | AI: MCP for Claude and compatible assistants

Getting Started

Introduction Call

1 hour

Understand your data landscape, AI goals, and requirements.



Readiness Assessment

2 to 3 days

Evaluate data sources, infra, and deployment model fit.



Pre-Production Build

3 to 4 weeks

Connect data, configure agents, validate against workflows.



Production Scale

1 to 3 months

Expand across teams, domains, and use cases.