

# Flynn Data Services

## Fractional Data & AI Workflow Engineering

---

I make your data infrastructure production-grade and your AI workflows safe to deploy.

**Track 1 — Data Infrastructure:** Right-size your data stack, reduce tooling costs, and build pipelines that actually work.

**Track 2 — AI Workflow Architecture:** Structured risk assessment, human checkpoint design, evaluation harnesses, and full cost modeling — so AI prototypes make it to production.

*Every service below is backed by a public open-source artifact you can inspect, run, and evaluate before a single dollar changes hands. The tools are free. The judgment and implementation are what you pay for.*

---

### TRACK 2 — AI WORKFLOW

## AI Memory System Setup

*A platform-agnostic knowledge layer your team and AI agents can read from any client.*

Your decisions, meeting notes, and context persist across sessions and across AI providers. No more re-explaining context every conversation. Your memories follow you from Claude to ChatGPT to whatever comes next.

### What you get:

- Supabase + pgvector database deployed to your infrastructure
- MCP server for agent access (works with Claude, ChatGPT, any MCP-compatible client)
- Multi-user schemas with row-level security isolation
- Semantic search across your team's captured knowledge
- Configuration and onboarding for your team

### Who it's for:

Founders and small teams tired of re-explaining context every AI session. Teams building AI-powered products that need persistent, queryable memory.

[Proof of work: "The Platform Is Catching Up" blog post](#)

---

### TRACK 2 — AI WORKFLOW

# Daily CEO Briefing System

*An automated morning email that compiles decisions, action items, and open questions across all your work streams.*

A GitHub Action collects overnight activity from your data sources, Claude formats it into a decision-ready brief, and it lands in your inbox before you wake up. Reply with your decisions and they're captured and executed. Scannable in 5 minutes.

## What you get:

- GitHub Action collector pulling from your CRM, email, calendar, banking, and code repos
- Claude API formatting into a structured, decision-ready morning brief
- Resend email delivery with reply-to-execute decision loop
- magic\_source tool that discovers what integrations to build next
- Runs on your infrastructure for under \$10/month

## Who it's for:

Solo founders, fractional executives, and anyone managing multiple concurrent work streams who needs a single-pane triage every morning.

[Proof of work: CEO Agent architecture post](#)

---

## TRACK 2 — AI WORKFLOW

# AI Workflow Risk Assessment

*A structured audit of your AI workflow before it hits production — failure modes, human checkpoints, and total cost at scale.*

Describe your AI workflow, and the assessment engine scores it across six failure types and five risk dimensions. The output is a report your engineering lead can take to stakeholders to shift the conversation from "is this safe?" to "here's exactly what could go wrong and what we're doing about it."

## What you get:

- Workflow mapped against six failure types (including ghost data)
- Five risk dimensions scored per step with composite risk score
- Human checkpoint recommendations (pre-flight review, sampling audit, escalation triggers)
- Token + infrastructure cost projection (including Snowflake credits if applicable)
- Evidence BI report for stakeholder presentation

## Who it's for:

Engineering leads shipping AI-powered features who need to answer "is this safe to put in front of customers?" Teams whose AI prototypes stalled between demo and production.

[Proof of work: fds-ai-workflow-audit repo \(Apache 2.0\) | "Your AI Prototype Is the Easy Part" blog post](#)

---

## TRACK 2 — AI WORKFLOW

# AI Workflow Documentation & Diagrams

*Production-quality visual documentation of your AI workflows that non-technical stakeholders can understand.*

Animated architecture diagrams and visual documentation using Remotion and programmatic video. Explain how your AI system works to investors, customers, and new hires without requiring them to read code.

**What you get:**

- Animated workflow diagrams (Remotion / programmatic video)
- Static architecture diagrams for docs and presentations
- Stakeholder-friendly documentation of your AI system's architecture

**Who it's for:**

Product teams, marketing teams, and founders who need to communicate "how our AI works" to people who don't read code.

[Proof of work: Video demos at flynndata.com](https://flynndata.com)

---

## Snowflake Governance & RBAC Setup

*Config-driven Snowflake access control — AI-assisted interview, deterministic Terraform deployment.*

An AI-assisted CLI interview captures your org's access requirements and generates a declarative config. Terraform deploys it. The AI never touches your Snowflake account — deterministic infrastructure-as-code with an AI intake layer.

### What you get:

- Connector roles per integration tool (Fivetran, dbt, Looker, etc.)
- Functional access roles per team function
- ACCOUNTADMIN guardrails and environment isolation
- RSA key-pair auth for service accounts
- Full Terraform config deployed and validated against your environment

### Who it's for:

Teams that have outgrown "everyone uses ACCOUNTADMIN" and need real governance before their next audit, security review, or compliance milestone.

Proof of work: [fds-snowflake-governance-blueprint repo \(Apache 2.0\)](#) | [Blog post](#)

## Data Pipeline Audit & Cost Optimization

*Find out what you're overpaying for and what you're missing — with a clear path to fix it.*

A 10-query Snowflake workload audit profiles your environment across query patterns, warehouse utilization, refresh cadence, and DuckDB compatibility. The Evidence BI report shows where credits are going and recommends whether to optimize, route to DuckDB, or right-size the entire stack.

### What you get:

- Snowflake workload cost audit (10 survey queries, Evidence BI report)
- DuckDB routing assessment (Analyst Accelerator / BI Optimizer / Iceberg-Ready Enterprise)
- Governance gap analysis — what RBAC policies break if you route queries to DuckDB
- Custom ETL pipeline builds with dlt + DuckDB + Evidence for teams that don't need Snowflake
- Credit optimization strategy for teams with excess Snowflake credits

### Who it's for:

Teams spending \$2K–\$10K/month on Snowflake and wondering if it's justified. Teams with excess credits rolling over and no plan for deploying them. Teams building their first analytics stack.

Proof of work: [fds-snowflake-cost-audit repo \(Apache 2.0\)](#) | [BI-in-a-Box repo](#)

### Recent Results

**92%** reduction in search infrastructure costs

**66%** reduction in annual CRM licensing costs

**96%** improvement in CRM data latency

**52,000+** invalid contacts removed from a broken marketing database

**34x** cost underestimation caught in an AI workflow audit before production

---

## Let's Talk

No long-term commitment required to get started. Most engagements begin with a focused assessment that tells you exactly where you stand and what to fix first.

[ken@flynndata.com](mailto:ken@flynndata.com)

[flynndata.com](https://flynndata.com)

[cal.com/flynndata/30min](https://cal.com/flynndata/30min)

*Less Stack. More Signal.*

© 2026 Flynn Data Services LLC • Henrico, Virginia