

# PRADIMEAI-CoreQuantum

Enterprise AI Strategy, Advisory & AI Platforms | PradimeAI.com

## AI Architect, Data & Analytics SME — Contract / Fractional

📍 Chicago / Dallas preferred — Remote considered

🕒 Contract based | 1099 only

🏢 Financial Services | Healthcare | Industrial

### Building Our Strategic Architect Network

PradimeAI is building a curated bench of senior AI and data architects ahead of upcoming client engagements. We are proactively identifying the right talent now so we can move quickly when SOWs are signed. If you are a hands-on architect who can design enterprise-grade data and AI platforms and speak strategy to the C-suite in the same breath, we want to meet you.

**Strictly, no agencies please.**

### ABOUT PRADIMEAI

PradimeAI is a boutique enterprise AI strategy and advisory firm founded by a former Digital, Data and AI executive with 25+ years leading AI transformation, data platform engineering, and enterprise strategy across global financial services and regulated industries. We advise Fortune 500 C-suites and Boards on enterprise AI strategy, AI operating model design, data architecture transformation, and responsible AI governance — with the implementation depth to move clients from strategy to execution.

### THE OPPORTUNITY

We are looking for a senior AI Architect specializing in data and analytics who brings a rare combination: the technical depth to design and govern end-to-end enterprise data and AI platform architectures, and the strategic fluency to shape how those platforms serve business goals at the C-suite level. You are equally at home defining a Snowflake and Databricks reference architecture, reviewing a cloud AI stack for production readiness, and presenting platform strategy to a CDO or Board.

*These are project-based, SOW-driven engagements with the expectation of an ongoing relationship as our client portfolio grows. The right candidates will become a core part of PradimeAI's architecture and delivery capability.*

### WHAT YOU WILL DO

- Design and govern end-to-end enterprise data and analytics platform architectures — defining the reference architecture across ingestion, storage, processing, serving, and AI/ML layers for large-scale regulated enterprises
- Lead Snowflake, Databricks, and cloud-native AI stack architecture engagements — assessing current-state platforms, identifying gaps, and defining the target architecture and migration roadmap required to support scalable AI and analytics at enterprise scale
- Define platform services architecture for AI and analytics — including data lake house design, feature store strategy, real-time streaming architecture (Kafka, Spark, Flink), MLOps platform design, and API/data serving layer patterns
- Architect enterprise AI infrastructure for GenAI, LLM, and Agentic AI deployment — including vector database design, RAG pipeline architecture, LLM inference infrastructure, prompt engineering patterns, and responsible AI guardrails at the platform level
- Conduct deep data architecture reviews for enterprise clients — assessing Systems of Record modernization needs, data mesh vs. lake house tradeoffs, data governance and lineage design, and the infrastructure readiness required to sustain AI at scale
- Define cloud AI platform strategy across AWS, Azure, and GCP — including platform service selection, multi-cloud architecture patterns, cost optimization, and security and compliance architecture for regulated industries
- Bridge architecture and strategy — translating complex technical platform decisions into clear strategic choices, investment tradeoffs, and business cases that CDOs, CTOs, and Boards can act on

- Partner with client data engineering, ML, and platform teams — providing architecture governance, design review, and hands-on guidance through the implementation phases of engagements

## WHAT WE ARE LOOKING FOR

### Experience & Background:

- 10–15+ years of hands-on enterprise data and AI architecture experience — spanning senior individual contributor architect roles through principal or staff architect levels at large regulated enterprises; consulting background considered when paired with deep implementation depth at enterprise scale
- Deep domain expertise in one or more of: large-scale banking or financial services, healthcare systems, or industrial / manufacturing enterprises — with firsthand knowledge of the regulatory, data, and operational complexity of those environments
- Proven track record designing and delivering production-grade enterprise data and AI platforms — not just advising on architecture, but governing and validating execution through to live deployment at scale
- Expert-level fluency in Snowflake and Databricks — including architecture design, performance optimization, cost governance, data sharing models, and integration with AI/ML and BI tooling in enterprise environments
- Broad cloud AI stack depth across AWS (SageMaker, Glue, Redshift, Kinesis), Azure (Synapse, Azure ML, Data Factory, Event Hub), and/or GCP (BigQuery, Vertex AI, Dataflow) — with the ability to evaluate and architect multi-cloud data and AI platforms for enterprise use

### Technical & Strategic Fluency:

- Snowflake Architecture & Optimization
- Cloud AI Stack: AWS / Azure / GCP
- Real-Time Streaming: Kafka / Flink / Spark
- GenAI & LLM Infrastructure (RAG, Vector DB)
- Enterprise Data Governance & Lineage
- API & Data Serving Layer Architecture
- Databricks Lakehouse & Delta Lake
- Data Mesh & Data Lakehouse Design
- Feature Store & MLOps Platform Design
- Agentic AI Platform Architecture
- Systems of Record Modernization
- AI Governance & Responsible AI by Design

### The Intangibles:

- You are equally comfortable whiteboarding a Databricks medallion architecture with a data engineering team and presenting a cloud AI platform strategy to a CDO or CTO. Technical credibility and executive communication are both non-negotiable
- You have strong, defensible opinions on architecture tradeoffs; data mesh vs. lake house, build vs. buy on ML platforms, batch vs. real-time pipeline design — and the ability to make those recommendations land with both technical and business audiences
- You are energized by building something from the ground up. PradimeAI is growing and we want architects who see this as a platform, not just a project — people who want their work to define how enterprise AI gets built across industries

## ENGAGEMENT MODEL

**Structure:** Project-based SOW engagements; fractional retainer model available for ongoing architecture advisory and governance

**Location:** Chicago or Dallas metro preferred; remote considered for the right profile

**Travel:** Client-site presence expected for architecture reviews and key delivery milestones; typically 30–40% depending on SOW

**Commitment:** Engagements typically range from 3–9 months; potential for multi-engagement relationship as our client portfolio scales

## APPLY

We review profiles on a rolling basis. To be considered for our founding architect network, send a note directly to **Dimple Thakkar, Founder & Chief AI Officer, PradimeAI** at [Dimple@PradimeAI.com](mailto:Dimple@PradimeAI.com).