

Leo Pitts

VP/Director of Data & Intelligence Engineering

☎ (919) 555-0148 ✉ leo.pitts@protonmail.com 💼 linkedin.com/example/leopitts 📍 Raleigh, NC 27609

SUMMARY

Data and Intelligence Engineering executive with 14 years building modern data platforms and leading multidisciplinary teams across consulting, financial services, and regulated enterprise environments. Brings nine years of people leadership for teams of five to fourteen engineers and senior individual contributors. Translates client problems into governed lakehouse, semantic, retrieval, graph, and agentic architectures that support dependable business outcomes. Production delivery spans Zero-ETL, **real-time streaming**, Python, PySpark, SQL, Databricks, Snowflake, Synapse, Fabric, Azure, vector stores, **knowledge graphs**, MCP connectors, Graph-RAG, RAGAS, **Golden Set testing**, and **policy-as-code**. Partners with architecture, design, security, governance, and executive groups to improve delivery quality, coach technical talent, and advance responsible AI practices.

EXPERIENCE

Vice President, Data & Intelligence Engineering

Blue Ridge Intelligence Partners 📅 March 2021 - Present 📍 Raleigh, NC

Leads a 14-person consulting engineering organization delivering governed **Data Foundation** platforms, retrieval services, semantic systems, and agentic solutions for regulated clients. Owns technical standards, coaching, architecture decisions, governance participation, and delivery quality across multidisciplinary programs.

- Set Layer 1 standards with Zero-ETL and **data contracts** for repeatable client delivery outcomes
- Launched **Pinecone**, Weaviate, and Neo4j foundations for retrieval and graph intelligence programs
- Deployed **MCP connectors** for Jira, CRM, ERP, and lakehouse sources with secure access controls
- Defined Data, Reasoning, and Validation hand-offs with AI architecture and design leaders
- Established RAGAS, Golden Set, and retrieval observability practices for production probabilistic systems
- Coached engineers through architecture reviews and career plans, preparing emerging technical leaders
- Represented engineering in **AI Ethics** forums covering governance, sovereignty, and **responsible AI** guardrails

Director, Applied Data Platforms

Meridian Cloud Advisory 📅 July 2017 - February 2021 📍 Durham, NC

Managed a nine-person consulting delivery group responsible for cloud data platforms, lakehouse programs, semantic roadmaps, and early enterprise retrieval solutions. Converted executive priorities into governed architectures, coached senior contributors, and strengthened production readiness across 11 client programs.

- Directed 11 client programs across Azure, **Databricks**, **Snowflake**, Synapse, and Fabric platforms
- Translated executive objectives into lakehouse roadmaps with ownership, lineage, and service expectations
- Built **Python** and **PySpark** frameworks combining batch, streaming, and Zero-ETL ingestion patterns
- Guided **LangChain**, GPT, and **Claude** retrieval programs with baselines and escalation paths
- Partnered with solutions architecture and interface design leaders on integrated platform delivery
- Introduced Golden Sets and evaluation workflows for low-confidence response handling
- Prepared senior contributors for broader ownership through reviews, mentoring, and architecture workshops

Senior Data Engineering Manager

Crescent Harbor Financial 📅 January 2014 - June 2017 📍 Richmond, VA

Supervised seven engineers delivering governed analytical datasets and streaming services for risk, compliance, and relationship-management stakeholders. Directed warehouse modernization, security controls, semantic modeling, operational reliability, **executive communication**, and structured development for engineering talent.

- Modernized warehouse processing with **Python**, **PySpark**, and expert SQL for regulated reporting
- Implemented Snowflake and **Azure** controls for identity, classification, lineage, and sensitive data access
- Unified CRM and ERP concepts through semantic modeling, reducing reporting reconciliation effort
- Introduced dimensional modeling and **non-relational** access patterns for risk data products
- Established monitoring, SLA reviews, **incident response**, and root-cause analysis across pipelines
- Presented architecture risks to executives, compliance partners, and business sponsors
- Developed engineers through structured feedback, technical pairing, and delivery retrospectives

Data Engineering Lead

Summit Retail Systems 📅 August 2012 - December 2013 📍 Columbus, OH

Led a six-person engineering pod building a cloud analytics foundation for merchandising, inventory, and customer engagement teams. Shaped reusable pipelines, secure data services, graph modeling, migration practices, mentoring, and release consistency across business domains.

- Built Python and SQL components validating schema drift, late data, and duplicate records
- Migrated workloads to Azure and Databricks with documented contracts and lineage
- Created **Neo4j** product and supplier models supporting relationship discovery for analysts
- Exposed curated data services through secure APIs and consistent integration patterns
- Mentored junior engineers through reviews, design walkthroughs, and release retrospectives
- Improved testing and deployment consistency through reusable delivery documentation

Data Engineer

Harborview Analytics Lab 📅 June 2011 - July 2012 📍 Columbus, OH

Built foundational data workflows for customer and operations reporting while developing practical depth in Python, SQL, PySpark, Azure, testing, and production support. Strengthened data quality practices through reconciliation, source mapping, and structured investigation of failed loads.

- Built Python and SQL workflows for customer and operations reporting datasets
- Created PySpark transformations for historical files with recurring reconciliation checks
- Modeled **relational** reporting tables and documented source-to-target business definitions
- Supported Azure storage and notebook analysis through a supervised platform proof of concept
- Added unit tests and deployment notes that clarified analyst hand-offs
- Investigated failed loads through query analysis and structured root-cause documentation

PROJECTS

Enterprise Agent Data Fabric 📅 2025

Aligned engineers, architects, and client stakeholders around semantic layers, vector retrieval, graph context, and secure MCP access. Shared design decisions openly, giving teams a practical foundation for agent workflows and production review.

Probabilistic Retrieval Quality Program 📅 2023

Brought delivery, governance, and client teams together around chunking, RAG optimization, RAGAS, evals, and Golden Set testing. Created a calmer review process for uncertain answers and clearer decisions about production readiness.

Regulated Lakehouse Modernization 📅 2020

Worked with finance and operations partners to connect business definitions with lakehouse architecture, lineage, contracts, and service expectations. The group moved forward with shared ownership rather than isolated platform decisions.

LEADERSHIP & AWARDS

- Responsible AI Practice Recognition, Blue Ridge Intelligence Partners, 2024
- Client Innovation Award, Meridian Cloud Advisory, 2020
- Engineering Mentor of the Year, Crescent Harbor Financial, 2016

EDUCATION

Master of Science in Computer Science

North Carolina State University 🎓 GPA: 4.0 📅 2011 📍 Raleigh, NC

Coursework: Data Engineering, Artificial Intelligence, Database Systems, Distributed Computing

Bachelor of Science in Information Technology

University of Cincinnati 🎓 GPA: 4.0 📅 2009 📍 Cincinnati, OH

Coursework: Database Management, Systems Analysis, Software Engineering, Information Security

CERTIFICATIONS

- Microsoft Certified: Azure Data Engineer Associate 📅 2023
- Databricks Certified Data Engineer Professional 📅 2022
- SnowPro Advanced: Data Engineer 📅 2021
- Microsoft Certified: Azure Solutions Architect Expert 📅 2020

TECHNICAL SKILLS

- **Programming:** Python, PySpark, SQL
- **Data Platforms:** Databricks, Snowflake, Synapse
- **Cloud Services:** Azure, AWS, GCP
- **Microsoft Data:** Fabric, Entra ID, Synapse

- **Vector Stores:** Pinecone, Weaviate, Milvus
- **Graph Systems:** Neo4j, Graph-RAG, knowledge graphs
- **Agent Frameworks:** LangChain, LangGraph, CrewAI
- **Foundation Models:** GPT, Claude, Gemini
- **Retrieval Evaluation:** RAGAS, evals, Golden Set testing
- **Data Architecture:** Lakehouse, Zero-ETL, streaming
- **Security Controls:** OAuth, row-level security, column-level security
- **Governance:** Data contracts, lineage, policy-as-code
- **Integration Sources:** Jira, CRM, ERP
- **Operations:** Observability, SLA design, incident response
- **Responsible AI:** AI Ethics, model governance, data sovereignty

SKILLS

- | | | | |
|---------------------|-------------------|--------------|--------------------|
| • Data architecture | • Data governance | • Expert SQL | • Vector databases |
| • Zero-ETL | • Data lineage | • Azure | • Knowledge graphs |
| • Lakehouse | • Python | • AWS | • MCP connectors |
| • Data contracts | • PySpark | • GCP | |

PROFESSIONAL AFFILIATIONS

- AI Ethics and Governance Working Group, Blue Ridge Intelligence Partners, 2022 - Present
- Triangle Data Engineering Meetup, Volunteer Technical Speaker, 2019 - Present
- Data Architecture Center of Excellence, Chair, Meridian Cloud Advisory, 2018 - 2021

LANGUAGES

- English (Native)
- Spanish (Intermediate)

ADDITIONAL INFORMATION

Work Status : Authorized to work in United States. No sponsorship required.

REFERENCES

AVAILABLE ON REQUEST