

# Alex Hernandez

[hernandezalex.dev](http://hernandezalex.dev) · [alexcp.hernandez@gmail.com](mailto:alexcp.hernandez@gmail.com) · (832) 360-5636

## SUMMARY

Systems Engineer with 7+ years of experience designing and implementing high-throughput data systems, scalable databases, and performant data access layers for enterprise applications. Expert in building production-grade ETL pipelines processing TB-scale datasets using Python and Apache Spark, with deep proficiency in data warehousing solutions including Amazon Redshift, Azure Synapse, and Google BigQuery. Proven track record optimizing data architectures for real-time processing, integrating data from enterprise systems (PLM, ERP, MES-adjacent), and collaborating with cross-functional teams to deliver data solutions that drive business value. Strong foundation in cloud platforms (AWS, Azure, GCP), big data technologies, and performance optimization at scale.

## PROFESSIONAL EXPERIENCE

### **AlixPartners, LLP – Senior Systems Engineer**

*Feb 2023 – Present*

- Designed and implemented high-performance data access layers processing 10TB/day across distributed data platforms (Databricks, Azure Synapse, Palantir Foundry), serving 50+ enterprise clients with reliable, scalable performance
- Architected scalable databases and data warehousing solutions on Amazon Redshift and Azure Synapse, optimizing schema design and indexing strategies to support high-speed, high-throughput analytical workloads
- Built production-grade ETL pipelines using Python and PySpark processing 500GB/hour with custom partitioning strategies, achieving 5x throughput improvement through performance optimization and data architecture redesign
- Integrated data from multiple enterprise business systems including ERP platforms, analytics tools, and cloud services, ensuring seamless data flow and maintaining data quality across distributed environments
- Monitored and optimized data pipeline performance continuously, reducing processing time 40% and infrastructure costs \$1.2M annually through query optimization, data partitioning, and resource utilization improvements
- Collaborated with software engineers, data scientists, and business stakeholders to design data architectures supporting diverse analytical needs including real-time dashboards, ML model training, and executive reporting
- Maintained comprehensive documentation for data models, ETL processes, and data integration workflows, establishing best practices for data governance and ensuring system transparency and auditability

### **Visualutions, Inc / Gryphon Healthcare – Data Platform Engineer**

*Sep 2021 – Feb 2023*

- Designed and maintained data pipelines and databases supporting 1,000+ users across Azure and AWS, ensuring 99.9% availability for analytics and machine learning workloads in healthcare and enterprise environments
- Developed Python-based ETL frameworks for data extraction, transformation, and loading from heterogeneous sources, implementing data quality validation and error handling for production reliability
- Implemented database schemas and data models optimized for analytical query performance, utilizing indexing strategies and materialized views to accelerate data retrieval for time-sensitive applications
- Integrated data from cloud-based business systems and APIs, building scalable data ingestion pipelines with automated monitoring and alerting for data quality and pipeline health
- Collaborated with data science teams to support ML model development, creating feature stores and data access patterns that enabled <100ms inference latency at 10K requests/day

## TECHNICAL SKILLS

### Programming & Data Engineering

- Programming Languages: Python (7+ years production experience, expert-level), SQL (expert), Java (working knowledge), Scala (familiar)
- Big Data Technologies: Apache Spark/PySpark (proficient), Hadoop ecosystem (familiar), distributed data processing at scale
- ETL & Data Pipelines: Data pipeline design and optimization, data transformation, data quality validation, workflow orchestration, real-time and batch processing

### Databases & Data Warehousing

- Relational Databases: PostgreSQL, MySQL, SQL Server (expert in query optimization, indexing, schema design)
- NoSQL Databases: MongoDB, Redis, Cassandra concepts, key-value stores
- Data Warehousing: Amazon Redshift (proficient), Azure Synapse Analytics (expert), Google BigQuery (working knowledge), data warehouse design patterns

### Data Platforms & Cloud Infrastructure

- Data Platforms: Databricks (Apache Spark), Palantir Foundry, Azure Synapse, Amazon Redshift, data lakehouse architectures
- Cloud Platforms: AWS (proficient - S3, EC2, Lambda, Redshift, Glue), Azure (expert - Data Factory, Synapse, Storage, Databricks), GCP (working knowledge - BigQuery)
- Stream Processing: Kafka concepts, real-time data processing patterns, event-driven architectures (learning Flink)

### Data Architecture & Best Practices

- Data Architecture: Data modeling (dimensional, normalized), data warehouse design, data lake architecture, scalable database design patterns
- Performance Optimization: Query optimization, partitioning strategies, indexing, caching, profiling and bottleneck identification

- Data Governance: Data quality management, metadata management, documentation practices, audit trails, compliance awareness
- Integration: API integration (REST/GraphQL), enterprise system integration (ERP, analytics platforms), data synchronization patterns

### **DevOps & Tools**

- Containerization: Docker (proficient), Kubernetes concepts (learning), container orchestration patterns
- Infrastructure-as-Code: Terraform (working knowledge), ARM/Bicep, CloudFormation, version-controlled infrastructure
- Version Control & CI/CD: Git/GitHub, automated testing, CI/CD pipelines, code review practices

### **CERTIFICATIONS**

- Microsoft Certified: Azure AI Engineer (AI-102)
- Microsoft Certified: Azure Fundamentals (AZ-900)
- Microsoft Certified: Azure Administrator Associate (AZ-104) - Previously held

### **EDUCATION & CLEARANCE**

#### **Professional Background**

7+ years progressive experience in data engineering, data architecture, and scalable database design with continuous professional development through vendor certifications and hands-on implementation of big data technologies across cloud platforms.

#### **Security Clearance**

U.S. Citizen, eligible to obtain and maintain an active U.S. Secret security clearance.