

KRISHNA NIDADAVOLU

Senior Data Engineer • Cloud Data Modernization • Snowflake | AWS | Azure & Microsoft Fabric

203-804-7966 • nidadavolukrishna@gmail.com • linkedin.com/in/krishnanidadavolu • krishnanidadavolu.com

AWS Certified Solutions Architect – Associate • U.S. work authorized — no visa sponsorship required

PROFESSIONAL SUMMARY

Senior Data Engineer with 15+ years in enterprise data and 7+ years architecting cloud-native data platforms on Snowflake, AWS, Azure Data Factory, and Microsoft Fabric. Proven leader of end-to-end modernization — retiring on-prem SQL Server data marts, replacing legacy ETL with Snowflake/ELT, and standing up secure serverless data products on AWS. Translates complex business requirements into governed, scalable, and cost-efficient pipelines that drive regulatory reporting, analytics, and customer-facing financial products.

CORE COMPETENCIES

Cloud Data Platforms: Snowflake • AWS (Redshift, S3, Glue, Athena, Lambda, API Gateway, Cognito, DynamoDB, Secrets Manager) • Azure Data Factory • Microsoft Fabric (Data Factory, Lakehouse, Warehouse)

Data Engineering: ETL/ELT design • Data lakes & lakehouse • Dimensional modeling (Star/Snowflake) • SCD Type 1/2 • Performance tuning • Data quality

Programming: Python • PySpark • SQL (advanced, performance tuning) • PL/SQL • T-SQL • UNIX shell scripting

Orchestration: Azure Data Factory • Microsoft Fabric Data Factory & Activators • IBM DataStage 11.7 • Apache Airflow • Control-M • UC4 • Zena

Databases: Snowflake • MS SQL Server • Oracle • DB2 • Netezza • DynamoDB

Security & Governance: AWS IAM, Cognito, Secrets Manager • RBAC & Resource Monitors (Snowflake) • Solix data masking • Data quality rules

DevOps & Reporting: Azure DevOps (Pipelines, Releases) • Git-based CI/CD for data • MicroStrategy • Cognos

PROFESSIONAL EXPERIENCE

National Cooperative Bank • Hybrid (Arlington, VA)

Sep 2020 – Present

Senior Data Engineer

Lead modernization of an enterprise data platform supporting analytics, regulatory reporting, and a customer-facing financial portal that aggregates data from 15+ core banking systems and external vendors.

- **Snowflake migration:** Led migration of a legacy on-prem SQL Server Data Mart to Snowflake — designed the data model, virtual warehouses, resource monitors, and RBAC; translated complex T-SQL into optimized Snowflake SQL using Streams and Tasks, improving end-user reporting query performance by ~40%.
- **Serverless data portal:** Architected a secure serverless financial data portal on AWS using Cognito, API Gateway, Lambda, and DynamoDB — implemented multi-tier hierarchical authorization that dynamically restricts data access per user and brokers calls to FIS Global banking APIs.
- **Hybrid cloud pipelines:** Built hybrid cloud orchestration with Azure Data Factory and Microsoft Fabric Data Factory — lift-and-shift of on-prem EDW data into Fabric Warehouse, plus stored-procedure-driven loads from Fabric EDW into a Fabric DataMart.
- **Event-driven orchestration:** Engineered event-driven pipelines using Fabric event streams and activators that trigger Fabric loads on DataStage completion and refresh MicroStrategy reports automatically — eliminating manual hand-offs across the daily batch.
- **Lakehouse engineering:** Developed PySpark notebooks in Databricks to validate, transform, and stage CSV data into the Lakehouse with automated archival, ensuring idempotent reruns and end-to-end traceability.

- **Snowflake integration:** Built end-to-end pipelines extracting Fiserv data from Snowflake into the on-prem EDW; converted DataStage ETLs into Snowflake stored procedures to streamline the Snowflake DataMart refresh.
- **Reliability & monitoring:** Hardened the legacy ETL ecosystem with a real-time DataStage monitoring pipeline and operational dashboard — giving Operations first-time visibility into batch performance across daily SOAP/REST API ingests.
- **Security:** Implemented secure connectivity via private API endpoints and AWS Secrets Manager between internal platforms and FIS Global services, satisfying audit and compliance requirements.

Environment: Snowflake, AWS (Cognito, Lambda, API Gateway, DynamoDB, Secrets Manager, IAM), Azure Data Factory, Microsoft Fabric (Data Factory, Lakehouse, Warehouse), Databricks, PySpark, IBM DataStage 11.7, MS SQL Server, Python, Solix, MicroStrategy, Azure DevOps.

The Children's Place • Secaucus, NJ

Aug 2018 – Aug 2020

Lead DataStage Consultant

Owned ETL design for promotions and POS data feeding the EDW and SAP-integrated point-of-sale platform.

- Partnered with business stakeholders to map operational source systems and design data-warehouse strategies for promotions, promotion categories, and POS data.
- Built complex DataStage parallel jobs leveraging pipeline and partition parallelism to load promotion data marts; modified existing flows to support new business rules.
- Implemented XML integration using Hierarchical and MQ stages — generating XML payloads from SAP for POS upgrades and consuming inbound XML into Netezza tables for downstream analytics.
- Used Basic Routines to pass values from sequencers to jobs at runtime; orchestrated DEV/PROD jobs via UC4 scheduler.
- Drove unit, integration, and UAT testing standards across the team, reducing post-release defects.

Environment: IBM DataStage 11.7, DB2, Netezza, Oracle 12c, UC4, SAP, MQ, SQL Developer, Toad for DB2.

NY State Department of Health • Albany, NY

May 2017 – Aug 2018

Lead DataStage Consultant

Delivered ETL for Medicaid enrollment & eligibility analytics powering Cognos reporting.

- Translated functional design documents into technical mappings; built parallel jobs and sequences to populate Type 1 and Type 2 dimensions.
- Authored PL/SQL procedures, functions, and packages; converted existing PL/SQL packages into modernized DataStage jobs.
- Designed and loaded enrollment / eligibility dimension, fact, and aggregate tables driving Cognos reports.
- Tuned DataStage jobs for performance; orchestrated runs and debugging through the DataStage Director runtime.

Environment: DataStage 11.5, DB2, Oracle 12c, IBM Rational Team Concert.

Health Care Service Corporation (Accenture) • Chicago, IL

Jan 2016 – May 2017

Lead DataStage Consultant

Built parallel ETL frameworks for healthcare data warehousing on IBM Information Server.

- Designed parallel jobs using the full DataStage stage library (Join, Merge, Lookup, Change Capture/Apply, Funnel, Filter, Transformer, Surrogate Key, Aggregator) to cleanse, aggregate, and load source data into the warehouse.
- Implemented Slowly Changing Dimensions for clinical and member data domains using the dedicated SCD stage.
- Tuned partitioning strategies and lookup-file staging, raising throughput on critical nightly loads.
- Built Master Sequencers and parameter-driven jobs with environment variables; orchestrated through Zena and provided 24x7 production support.

Environment: IBM Information Server / DataStage / Quality Stage 8.5, Zena 3.1, IBM CDC 10.2, Toad Data Point, MQ Visual Browse.

Whataburger • San Antonio, TX

Sep 2014 – Oct 2015

Lead DataStage Consultant

- Designed DataStage jobs ingesting MS SQL Server, flat files, and complex feeds into the EDW with business-rule transformations and Quality Stage data matching.
- Built dimension/fact loaders for daily and monthly marts based on source granularity; established naming, parameterization, and environment-promotion standards.
- Coordinated offshore delivery and led DataStage → SSIS migration tasks.

Environment: IBM WebSphere DataStage 8.5, SSIS, MS SQL Server 2012/2008 R2, Active Batch V9.

USAA (Accenture) • San Antonio, TX

Dec 2011 – Sep 2014

DataStage Consultant

- Built DataStage jobs ingesting Oracle, Netezza, DB2, MS SQL Server, and flat files into the EDW with business validations across HRMS / EPM domains.
- Profiled data with Quality Stage Investigate stage and PL/SQL scripts to detect anomalies, patterns, and inconsistencies.
- Migrated DataStage from 8.1 → 9.1 with no production disruption; used Borland Star Team for environment promotion and Control-M for orchestration.

Environment: DataStage / Quality Stage 7.5/8.1/9.1, PeopleSoft HRMS 9.1 / EPM 8.9, Oracle 10g/11g, DB2, MS SQL Server, Control-M, Red Hat Linux.

Payless Shoe Source • Topeka, KS

Aug 2010 – Nov 2011

DataStage Developer

- Loaded staging and warehouse layers using Sequential File, Dataset, Lookup, Modify, and database stages; implemented SCDs against star-schema mappings.
- Built Quality Stage cleansing flows (Investigate, Standardize, Match Frequency, Survive, Reference Match).
- Documented job templates and parameter standards for project reuse; orchestrated loads via Control-M.

Environment: IBM DataStage WebSphere 8.1, IBM AIX, DB2, Oracle 11g, UNIX, Control-M.

EDUCATION & CERTIFICATIONS

Master's Degree • Graduated 2010

AWS Certified Solutions Architect – Associate