

Vishwajeet Bhangare

CONTACT

Email: bhangare.vishwajeet@gmail.com

Phone: +91 9595107293

Location: Pune, Maharashtra

LinkedIn: [linkedin.com/in/vishwajeet-bhangare/](https://www.linkedin.com/in/vishwajeet-bhangare/) | <https://github.com/V4Data>

SKILLS

- Data Engineering & Architecture
- Azure Data Platform
- Snowflake & Databricks
- Microsoft Fabric & Synapse
- ETL / ELT Modernization
- Data Governance & Security
- Real-Time Data Processing
- Cloud Data Warehousing
- Performance Optimization
- Data Modeling & Analytics
- DevOps / CI-CD Automation
- Multi-Cloud Integration

TECHNICAL EXPERTISE

- Cloud & Data Platforms:** Azure | AWS | Snowflake | Microsoft Fabric | Azure Synapse | Databricks | Amazon Redshift | Azure Data Lake | Amazon S3 | dbt
- Data Engineering:** ETL/ELT | Data Pipelines | Data Architecture | Data Warehousing | Data Modeling | Dimensional Modeling | Medallion Architecture | Data Transformation
- Big Data & Analytics:** Apache Spark | Delta Lake | Spark Streaming | Airflow | Real-Time Processing | Batch Processing | POWER BI
- Programming & Databases:** Python | SQL | Pyspark | Microsoft SQL Server | Stored Procedures | REST APIs | DAX
- Governance & Security:** Data Governance | GDPR | Data Quality | Data Security | RBAC | RLS | CLS | Data Masking | Audit Compliance
- DevOps & Engineering:** CI/CD | Azure DevOps | Git | Agile | Jira | Performance Tuning | RCA

Professional Summary

- Strategic Data Engineering professional with 8+ years of experience **architecting enterprise-scale data platforms across Azure, Snowflake, Databricks, Synapse, AWS, Microsoft Fabric, and multi-cloud ecosystems.**
- Proven expertise in designing **high-performance data architectures, modern ETL/ELT frameworks, governance programs, and scalable analytics solutions** supporting thousands of users and petabyte-scale data environments.
- Strong track record of **driving cloud modernization, query optimization, data platform transformation, and real-time analytics** initiatives delivering significant improvements in performance, cost efficiency, compliance, and business agility.
- Specialized in **data governance, security, dimensional modeling, Medallion Architecture, data quality, GDPR compliance, DevOps, CI/CD, and advanced performance engineering** across distributed data platforms.
- Recognized for **ownership-driven delivery, cross-functional stakeholder collaboration,** and enabling self-service analytics, real-time decision intelligence, and enterprise-wide data democratization.

Work Experience

Senior Data Engineer

Aug 2025 – Present

EXL Service India Pvt. Ltd.

- Spearheading Snowflake performance optimization initiatives through advanced query tuning, micro-partition pruning, clustering strategies, and workload efficiency improvements.
- Designed metadata-driven Azure Data Factory solutions to automate schema extraction, column-level transformations, and scalable pipeline orchestration across enterprise datasets.
- Developed innovative automation framework generating 1,500+ empty CSV files with column headers using Snowflake metadata, overcoming native platform export limitations.
- Implemented parameterized ADF pipeline templates enabling reusable, scalable, and standardized data processing across multiple business use cases.
- Reduced pipeline downtime by 30%, improved query execution performance by ~50%, lowered Snowflake credit consumption by 35%, and enhanced overall platform efficiency through proactive performance engineering.

Associate Staff Engineer

Nov 2024 – Aug 2025

Nagarro Software Pvt. Ltd.

- Led Microsoft Fabric Warehouse optimization initiatives utilizing adaptive Liquid Clustering, incremental statistics, predicate pushdown, and advanced PySpark performance tuning.
- Consolidated 20+ enterprise data sources spanning Azure, AWS, and on-premises environments into Microsoft Fabric OneLake, creating a unified multi-cloud data ecosystem.
- Architected scalable ETL frameworks and automated integration pipelines significantly accelerating enterprise data onboarding and accessibility.
- Optimized Databricks Delta Lake workloads through advanced partition pruning, schema enhancements, and ingestion-level transformation strategies.
- Delivered 82% faster query execution, 97% reduction in data scanned, 83% lower compute consumption, reduced integration timelines from days to hours, and improved large-scale analytics performance by 60-70%.

Azure Data Engineer

Sep 2023 – Oct 2024

Linova Info Pvt. Ltd.

- Built governed Snowflake data marts and self-service analytics environments supporting 500+ business users with secure role-based access controls.
- Developed dbt-powered transformation frameworks, automated refresh pipelines, and metadata-driven architectures improving enterprise data accessibility and freshness.
- Optimized Snowflake workloads through clustering strategies, materialized views, query profiling, partition reduction, and compute cost management.
- Implemented Microsoft Fabric Data Factory enhancements and dynamic lookup frameworks capable of reliably managing metadata repositories exceeding 10,000 records.

NOTABLE CONTRIBUTIONS

- Optimized Microsoft Fabric workloads delivering **82% faster query performance**, **97% reduction in scanned data**, and **83% lower compute consumption**.
- Balanced data across executors, eliminating hotspots and **cutting job runtime from 20 minutes to under 4 minutes**. Improved cluster stability production workloads.
- Managed governance and security controls across **10+ PB of enterprise data**, achieving **100% audit compliance** with zero critical findings.
- Consolidated **20+ multi-cloud enterprise data sources** into a unified Microsoft Fabric OneLake ecosystem.
- Reduced Databricks and Synapse workload runtimes by up to **70–80%**, generating substantial compute and operational savings.
- Automated complex metadata-driven data operations including enterprise-scale schema extraction and CSV generation frameworks.
- Delivered real-time and near-real-time analytics capabilities across Azure, Databricks, Snowflake, Synapse, and Microsoft Fabric platforms.

PROFESSIONAL CERTIFICATIONS

- Microsoft Fabric Data Engineer Associate (DP-700)** | Microsoft
- Databricks Certified Data Engineer Associate** | Databricks
- AWS Certified Data Engineer – Associate** | AWS
- HackerRank SQL (Advanced & Intermediate)**
- Snowflake Snowpro Core Certified**

LANGUAGES

English | Hindi

- Enabled enterprise-wide self-service reporting, reduced support requests by 40%, improved dashboard refresh performance by 3x, lowered compute costs substantially, and achieved near 100% pipeline reliability.

Data Engineer

Sep 2021 – Aug 2023

Altech Star Solutions

- Established enterprise-grade governance, security, encryption, masking, row-level security, and column-level security controls across Azure and Snowflake environments managing 10+ PB of sensitive data.
- Designed and optimized Azure Synapse data warehouse architectures utilizing hash distribution, replicated dimensions, and advanced workload management strategies.
- Engineered Databricks Delta Lake optimization frameworks leveraging bucketing, partitioning, compaction strategies, and parallel processing enhancements.
- Resolved large-scale data skew and cluster imbalance challenges through adaptive partitioning, repartitioning, coalescing, and Spark execution optimization techniques.
- Achieved 100% audit compliance, reduced unauthorized access attempts by 99%, improved Synapse query performance by up to 60%, reduced Delta workload runtimes by over 70%, and generated significant daily compute savings.

Application Development Analyst

May 2019 – Oct 2020

Accenture Solutions Pvt. Ltd.

- Designed Azure-based modern data processing ecosystems integrating Azure Data Factory, Azure Data Lake Gen2, Azure Databricks, and streaming architectures.
- Configured secure Azure Databricks storage integrations and enterprise-grade data access frameworks for cloud-native data processing.
- Developed event-driven data pipelines enabling automated orchestration and near real-time data delivery across business domains.
- Standardized PySpark transformation frameworks to ensure consistency, scalability, maintainability, and data quality across enterprise workloads.
- Enabled real-time analytics capabilities, eliminated manual credential management, improved transformation efficiency, reduced latency, and accelerated business decision-making.

Data Consultant

Mar 2017 – Apr 2019

Nano Tech E Services Pvt. Ltd.

- Developed SQL Server database solutions including stored procedures, validation frameworks, transaction management, and audit-ready logging mechanisms.
- Implemented robust data quality, validation, governance, and exception-handling processes ensuring enterprise data integrity.
- Designed and automated ETL/ELT pipelines utilizing Azure Data Factory to streamline data ingestion, transformation, and delivery workflows.
- Built scalable PySpark-based transformation frameworks supporting distributed processing and enterprise reporting requirements.
- Improved operational efficiency through automation, strengthened audit compliance, enhanced data reliability, and accelerated downstream analytics availability.

Education

Master of Science (M.Sc.) – Data Science & Artificial Intelligence

Oct 2025 – Present

O.P. Jindal University

Bachelor of Engineering (B.E.) – Electronics & Telecommunication

Aug 2008 – May 2013

Savitribai Phule Pune University