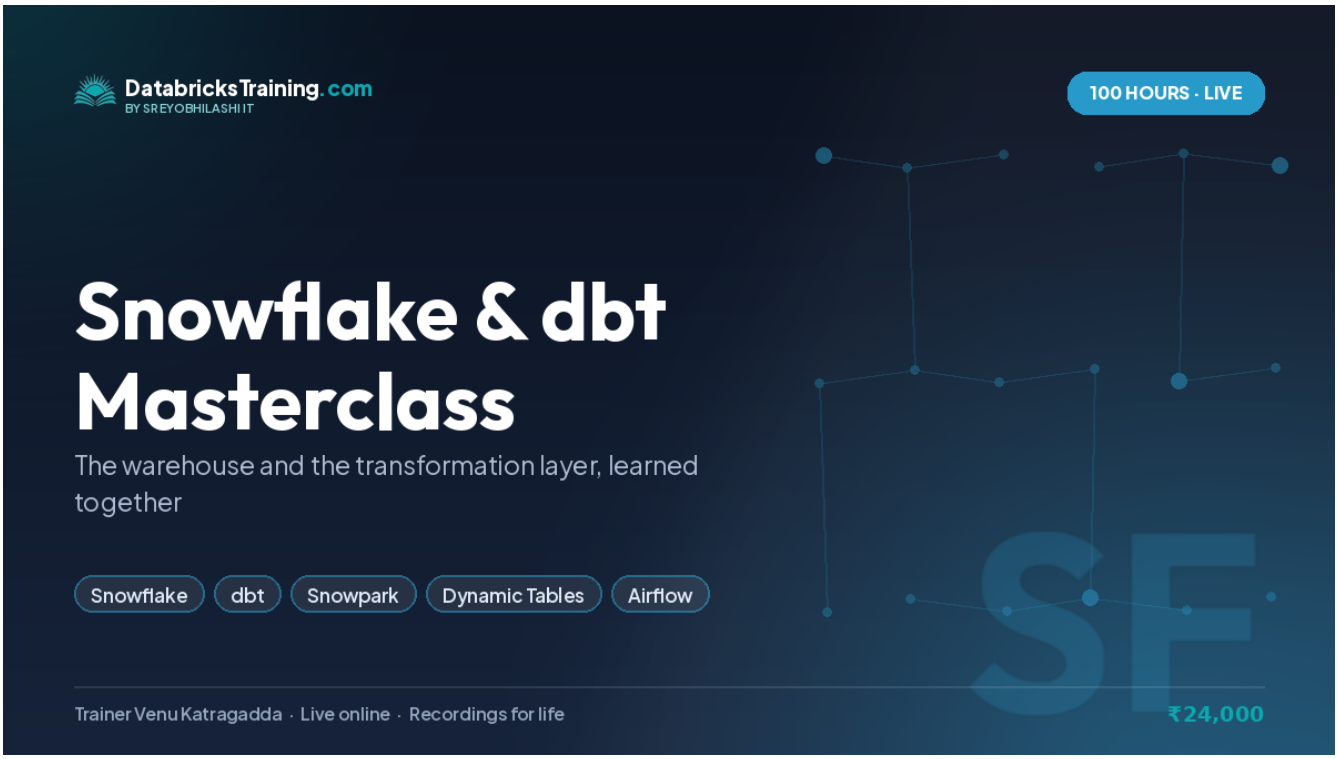

COURSE SYLLABUS

Snowflake & dbt Masterclass

Learn the warehouse and the transformation layer together, the way modern data teams actually build.

100 hours	10 modules	3 projects	**₹24,000**
LIVE INSTRUCTION	STRUCTURED PATH	PORTFOLIO READY	EMI AVAILABLE

Level	Beginner to Advanced
Mode	Live, instructor-led online sessions. Every session recorded; lifetime access.
Trainer	Venu Katragadda — 14+ years in industry (L&T Infotech, Capgemini, Brillio, Manthan Systems). 1,200+ professionals trained.
Batch timings	Morning 6:00–7:30 AM IST · Morning 7:30–9:00 AM IST · Evening 7:30–9:00 PM IST · Weekend (Sat & Sun)
Certification	SnowPro Core · dbt Analytics Engineering concepts
Enquiries	+91-8500002025 / +91-9247159150 (WhatsApp) · info@databrickstraining.com

What you will be able to do

Every outcome below is demonstrated in a lab or a graded project — not just discussed in a slide.

- Explain Snowflake's architecture and make correct warehouse-sizing and clustering decisions
- Write performant SQL and diagnose queries with the profile and query history
- Cut Snowflake credits with warehouse policies, result cache, clustering and resource monitors
- Build a full dbt project: staging → intermediate → marts, with tests, docs, snapshots and macros
- Ship dbt through CI/CD with slim builds, and orchestrate with Airflow or dbt Cloud
- Handle streaming and CDC into Snowflake with Snowpipe, Streams, Tasks and Dynamic Tables

Who this masterclass is for

- Analytics engineers and BI developers
- Data engineers adding Snowflake to their stack
- SQL developers moving to the cloud
- Teams adopting dbt

Prerequisites

- Good SQL — this course assumes JOINs and GROUP BY are comfortable
- Basic Git
- Snowflake 30-day trial ($400 credits) and dbt Core (free)

Tools and technologies

Snowflake · Snowpark · Snowpipe · Streams & Tasks · Dynamic Tables · Iceberg tables · dbt Core · dbt Cloud · Airflow · Kafka · Fivetran/Airbyte · Git · GitHub Actions · Looker/Power BI

Curriculum — 100 hours across 10 modules

Sessions run live in small batches. Roughly 60% of each session is hands-on lab work.

MODULE 1 Snowflake architecture and fundamentals — 10 hrs

- Three-layer architecture: storage, multi-cluster compute, cloud services
- Micro-partitions, metadata, pruning — why Snowflake is fast and when it isn't
- Editions, regions, cloud providers; accounts, roles, RBAC hierarchy
- Virtual warehouses: sizing, scaling policy, multi-cluster, auto-suspend/resume
- Loading basics: stages, COPY INTO, file formats, error handling
- Lab: account setup, role hierarchy, first loaded dataset

MODULE 2 SQL and data modelling in Snowflake — 12 hrs

- Snowflake SQL specifics: QUALIFY, LATERAL FLATTEN, semi-structured VARIANT, ASOF joins
- Window functions, recursive CTEs, pivots, time-series gap filling
- Dimensional modelling: star schema, SCD types, surrogate keys, fact grain, bridge tables
- Data Vault 2.0 overview and when it earns its complexity
- Lab: model a messy source system into a defensible star schema

MODULE 3 Loading, streaming and CDC — 10 hrs

- Snowpipe (auto-ingest) and Snowpipe Streaming; Kafka connector
- Streams (standard, append-only, insert-only) and Tasks (serverless, DAGs)
- Dynamic Tables: declarative pipelines, lag targets, refresh modes
- External tables, Iceberg tables, external volumes
- Fivetran/Airbyte ELT patterns and their cost model
- Lab: Project 2 near-real-time pipeline

MODULE 4 Performance and cost optimisation — 12 hrs

- Query profile reading: spilling, pruning ratio, exploding joins, remote IO
- Clustering keys, automatic clustering, search optimisation service, materialised views
- Caching layers: result cache, warehouse cache, metadata — and how to measure hit rates
- Warehouse strategy: separation by workload, queuing, multi-cluster economics
- Resource monitors, budgets, ACCOUNT_USAGE queries, per-team chargeback
- Lab: Project 3 cost audit with a written savings recommendation

MODULE 5 Security, governance and sharing — 8 hrs

- RBAC design patterns: functional vs access roles; least privilege at scale
- Dynamic data masking, row access policies, tag-based policies, object tagging
- Network policies, SSO/SCIM, key-pair auth, secrets handling
- Time Travel, Fail-safe, cloning, replication and failover
- Secure data sharing, reader accounts, Snowflake Marketplace, clean rooms
- Lab: GDPR-style masking implementation with an access-test matrix

MODULE 6 Snowpark and programmability — 8 hrs

- Snowpark DataFrame API in Python; pushdown execution model
- UDFs, UDTFs, stored procedures; Python packages and Anaconda channel
- Snowpark Container Services overview; Streamlit in Snowflake
- Cortex AI functions and vector search — LLM features inside the warehouse
- Lab: replace a Spark job with Snowpark and compare cost and runtime

MODULE 7 dbt fundamentals — 12 hrs

- Project structure, profiles, targets; dbt Core vs dbt Cloud vs dbt Fusion
- Models, ref(), source(), materialisations (view, table, incremental, ephemeral)
- Staging → intermediate → marts layering; naming and folder conventions
- Tests: generic, singular, dbt_utils, dbt_expectations; severity and thresholds
- Documentation, exposures, the docs site, lineage graph
- Lab: build the staging and marts layers of Project 1

MODULE 8 dbt advanced — 12 hrs

- Incremental models: strategies (merge, insert_overwrite, microbatch), is_incremental, late data
- Snapshots for SCD2; seeds; hooks; custom schemas and aliases
- Jinja and macros; packages; custom generic tests; dispatch and adapters
- Semantic layer and MetricFlow; metrics definitions
- Performance: model materialisation choices, DAG shape, threads, warehouse per model
- Lab: convert a 40-minute full refresh into a 3-minute incremental build

MODULE 9 Orchestration, CI/CD and quality — 8 hrs

- dbt Cloud jobs vs Airflow (Cosmos) vs Dagster — choosing an orchestrator
- Slim CI: state:modified, deferral, PR-scoped builds, blue/green deployments
- GitHub Actions pipeline for dbt: lint (sqlfluff), compile, test, deploy
- Data quality and observability: freshness, anomaly detection, alerting, data contracts
- Lab: full PR-to-production dbt workflow with tests as gates

MODULE 10 Architecture, capstone and certification — 8 hrs

- Snowflake vs Databricks vs BigQuery vs Redshift — an honest comparison you can defend
- Migration playbooks: Teradata/Oracle/SQL Server → Snowflake
- Capstone presentation and architecture review
- SnowPro Core exam guide, 120 practice questions, two timed mocks
- Analytics engineer interview questions and resume framing

Hands-on projects

You leave with work you can demo in an interview. Each project is code-reviewed and returned with written feedback.

PROJECT 1 E-commerce analytics warehouse

Raw → staging → marts in dbt over a multi-source e-commerce dataset, with 60+ tests, exposures, docs site and a Power BI/Looker Studio layer.

PROJECT 2 Near-real-time pipeline

Kafka → Snowpipe Streaming → Streams & Tasks → Dynamic Tables, with latency measurement and a cost comparison against micro-batch.

PROJECT 3 Cost optimisation engagement

Audit a messy account: warehouse right-sizing, query rewrite, clustering decisions, resource monitors — delivered as a written recommendation with projected savings.

Frequently asked questions

Do I need to pay for Snowflake?

No. The 30-day trial includes $400 of credits, which comfortably covers the course if you keep warehouses on XS with 60-second auto-suspend — taught in Module 1. You can start a fresh trial if you need more time.

dbt Core or dbt Cloud?

You will learn dbt Core (free, open source) as the foundation, and use dbt Cloud's free developer tier for the scheduling and CI features. Everything transfers between them.

Is this course useful if my company uses Databricks?

The dbt half is directly transferable — dbt runs on Databricks too. The Snowflake half is platform-specific but the modelling, testing and cost-governance thinking carries over. Many students take this alongside the Databricks course.

What about the SnowPro certifications?

The course maps to SnowPro Core, with 120 practice questions and two mocks. Advanced Data Engineer topics are touched on but that exam needs additional dedicated study.

How to enrol

- Book a free demo. Sit in on a live session with the current batch — no card, no commitment. WhatsApp +91-9247159150 or call +91-8500002025.
- Pick your batch. Weekday morning, weekday evening or weekend. All times IST.
- Pay the fee (₹24,000, instalments and EMI available) and get your workspace, material and calendar invites.
- Start building. Labs from session one; recordings posted the same day.

What is included

- 100 hours of live instruction by Venu Katragadda
- Lifetime access to all session recordings
- 3 code-reviewed portfolio projects
- Preparation for: SnowPro Core · dbt Analytics Engineering concepts
- Resume review, LinkedIn optimisation and one mock interview with written feedback
- Free re-attend of any future batch of the same masterclass
- Doubt-clearing sessions and a private student group

Refund policy in brief

When you cancel	Refund
Before the batch starts	Up to 90%
Within 5 days of start, max 2 classes attended	Up to 70%
After 7 days, or more than 5 classes attended	No refund — free transfer to a later batch

Full policy: databrickstraining.com/refund-policy/

Talk to us

Call	+91-8500002025 · +91-9247159150
WhatsApp	+91-9247159150 — usually answered within the hour
Email	info@databrickstraining.com
Website	databrickstraining.com/courses/snowflake-dbt-masterclass/
Foundation tier	databrickstraining.in — shorter 70–80 hour courses if this one is too advanced
Institute	Sreyobhilashi IT · Udyam UDYAM-TS-02-0334643 · 112, Samshraya Homes, opp. PF Office, Balanagar, Kukatpally, Hyderabad, Telangana 500037