

Snowpark for Python and ML

From experiment to production – all within Snowflake

Data science teams are stuck in notebooks, and models rarely reach production. Every workflow requires exporting data and juggling tools creating silos, delays, and risk. Snowpark for Python brings the entire ML lifecycle into Snowflake to unlock faster, safer insights.

What's holding you back

- Data movement and fragmented tools create security risks and slow collaboration
- Complex ML infrastructure and inconsistent environments drive high overhead
- Poor cost governance and unoptimized compute limit scalability
- Operational friction keeps models from reaching production

What success looks like

- › 70%+ reduction in data movement across systems
- › Models moving to production in weeks, not months
- › Consistent performance through a unified data and compute layer
- › Broad adoption across data engineers, analysts, and scientists
- › Lower tool sprawl, improved governance, and full observability

How evolv helps



evolv **accelerates your ML journey** by unifying data engineering, model development, and deployment inside Snowflake.



We **eliminate silos, reduce infrastructure complexity, and embed governance** from day one.



Our Snowpark solutions turn **predictive models into real-time business outcomes** securely and scalably.

Why evolv?

- **We build for clarity.** Unified workflows that make sense across teams
- **We embed governance.** Versioning, monitoring, and reproducibility built in
- **We drive adoption.** Enablement across data science, engineering, and analytics
- **We deliver outcomes.** Faster insights, lower costs, and models that actually scale