

Features for Production ML at Scale

Key Benefits

Tecton's Feature Platform for Machine Learning (ML) deployed on top of Snowflake Data Cloud enables organizations to securely and reliably build, store, and serve production-grade features across the enterprise.



Centralize feature logic and simplify feature management



Fast and accurate retrieval of training data sets of any size



Incorporate real-time data into production without training-serving skew

Solution Overview

Together, Tecton and Snowflake provide a simple and fast path to building and serving production-grade features to support a broad range of machine learning applications including fraud detection, recommendation systems, real-time pricing, and much more.

Available in preview to all Snowflake customers on AWS, Tecton acts as the central hub for ML features, allowing data teams to define features as code using Python and SQL and then automating production-grade ML data pipelines to generate accurate training datasets and serve up-to-date features online for real-time inference.

Key Challenges of Production ML

Whether you're building batch pipelines or already including real-time features in your ML initiatives, Tecton solves the many data and engineering hurdles that keep development time painfully high and, in many cases, predictive applications from ever reaching production at all, including:

- Training-serving skew
- Point-in-time correctness
- Productionizing notebooks
- Real-time transformations
- Melding batch + real-time data
- Latency constraints
- Data scientist and data engineering siloed workflows
- Limited discovery and re-use of features across teams

Get started today!

Request a Demo or a Free Trial
www.tecton.ai

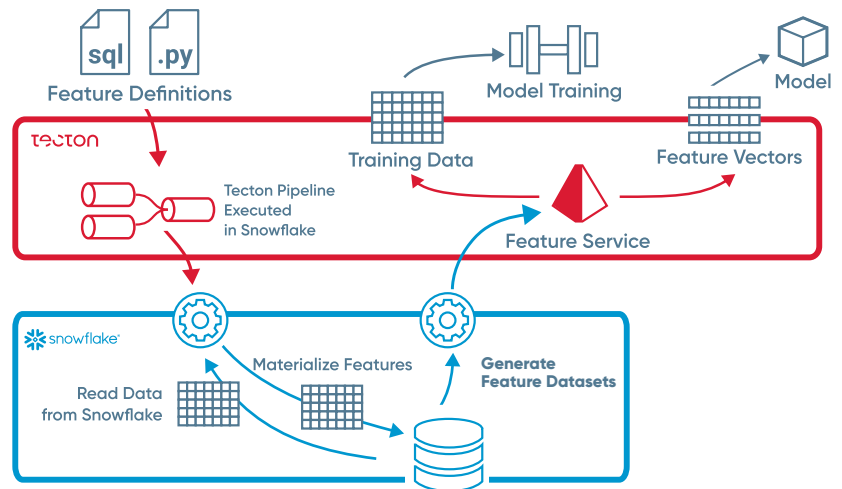


About Tecton

Tecton is the operational ML company. Organizations worldwide—including Tide, Atlassian, and Fortune 500 companies—rely on Tecton’s Feature Platform to bridge the gap between ML development and automated, real-time ML in production. Tecton is headquartered in San Francisco, with offices around the United States and Europe. Founded by the original creators of Uber Michelangelo, the very first enterprise feature store, Tecton is on a mission to help organizations embed automated machine learning into live, customer facing applications and business processes, quickly and at scale.

How it Works

Sitting on top of the Snowflake Data Cloud and its powerful processing engine for Python and SQL, Tecton’s feature platform enables data engineers and data scientists to build production-ready feature pipelines, and serve them at scale across teams, systems, and models, with only a few lines of code.



Under the hood, Tecton abstracts and automates the complex process that transforms raw data from batch, streaming, or real-time sources into features used to train ML models and feed predictive applications in production. Managing the ML feature lifecycle with Tecton not only ensures that feature materializations are always consistent, online for training and offline for inference, but that they are also stored in a searchable repository for easy sharing and re-use across teams and use cases.

“Tecton allowed us to significantly speed up our time to put models in production, enabling us to get faster feedback on our ML applications and in turn build products that deliver substantial value to our customers.”

tide Hendrik Brackmann
Director of Data Science and Analytics