

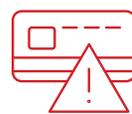
Build and serve production-ready features for real-time machine learning

Power one model or thousands simultaneously with production-ready batch, streaming, and real-time features.

For the few organizations that have been able to power products and services with real-time inference, the value generated goes above and beyond. From customer retention and fighting fraud, to making healthcare more efficient or personalizing pricing, **real-time predictions drive real revenue.**

But a model is only as good as the data that powers it. Building features from slow changing batch data is one thing. Building streaming data pipelines to power models in real time is another. These pipelines add value when they're correct but are very costly if something goes wrong. Simply put, when managed in-house, real-time infrastructure and pipelines are complex to engineer, time consuming to build, costly to maintain, hard to trust (from detecting problems to debugging and validating), and often leave little to no room to focus on what drives value: the data itself.

That's why Tecton does the heavy lifting for you. From a simple feature definition, Tecton compiles and orchestrates production-ready pipelines that transform batch, real-time, and streaming data into predictive features. With Tecton, developers, engineers, and data scientists can spend more time custom-engineering features and delivering models that drive competitive advantage, and less time building bespoke pipelines that rarely ever reach production.



Fraud Prevention



Recommendations



Real-time loan approvals



Ranking & Search



Dynamic Pricing



Personalization

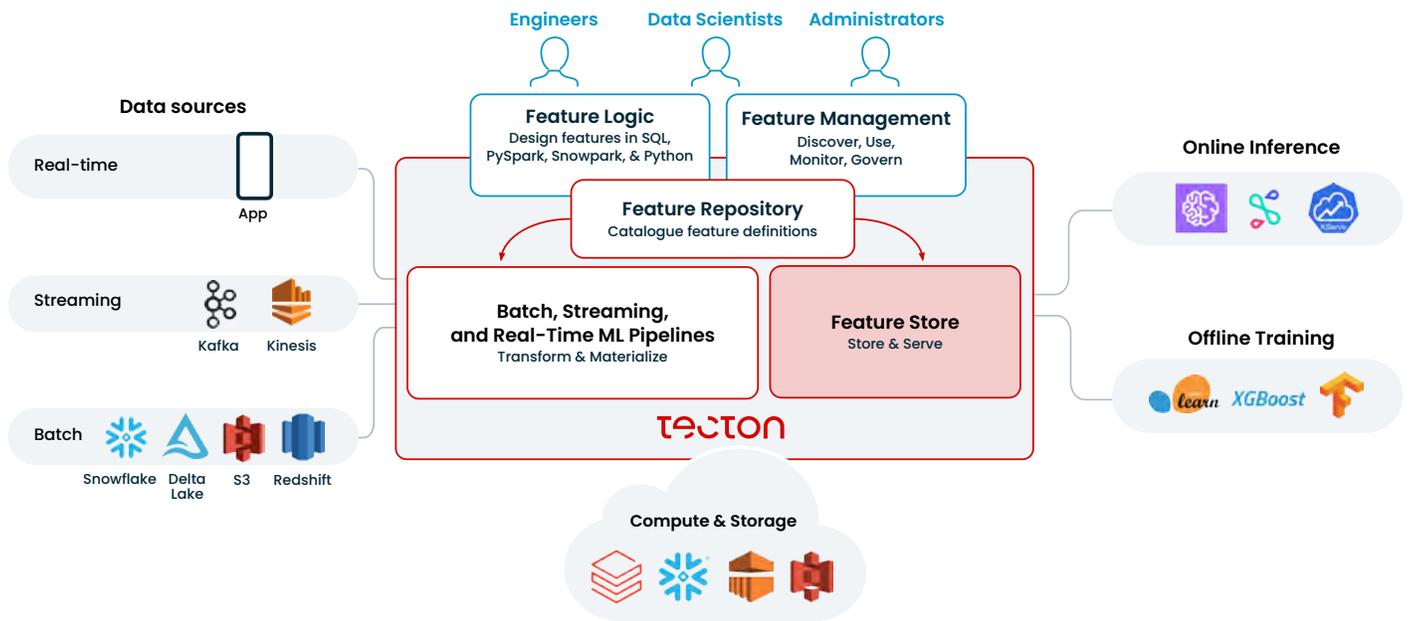
“As a net result of using the Tecton feature platform, we've improved over 200,000 customer interactions every day. This is a monumental improvement for us.”

Geoff Sims, Data Scientist at Atlassian

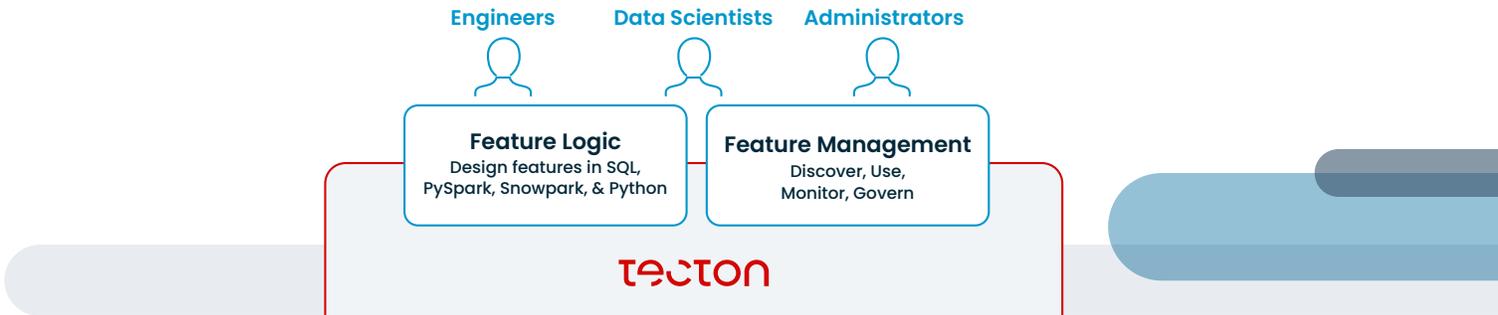
Tecton transforms your tailored feature definitions into real-time, production-ready pipelines

Use Tecton's feature platform for real-time machine learning to transform your feature definitions into production-ready batch, streaming, and real-time data pipelines for machine learning.

Better features means better models.



Custom engineer features using batch, streaming, and real-time data for better predictions



Design Feature Logic

Users can define feature logic using SQL, PySpark, SnowPark, or Python. By using the SDK in a notebook or any other Python environment, they can engineer streaming and real-time transformations in an all-in-one unified feature management framework. With a single line of code that works across batch and real-time data, users can define complex data transformations like time-window aggregations or generate training data with accurate backfills.

“With Tecton, it’s very easy to renew high-quality features, specifically within the same domain. The **overall delivery time has been significantly reduced** when it comes to building additional features and machine learning models.”

Hendrik Brackmann
VP Data at Tide

Manage and Monitor Features

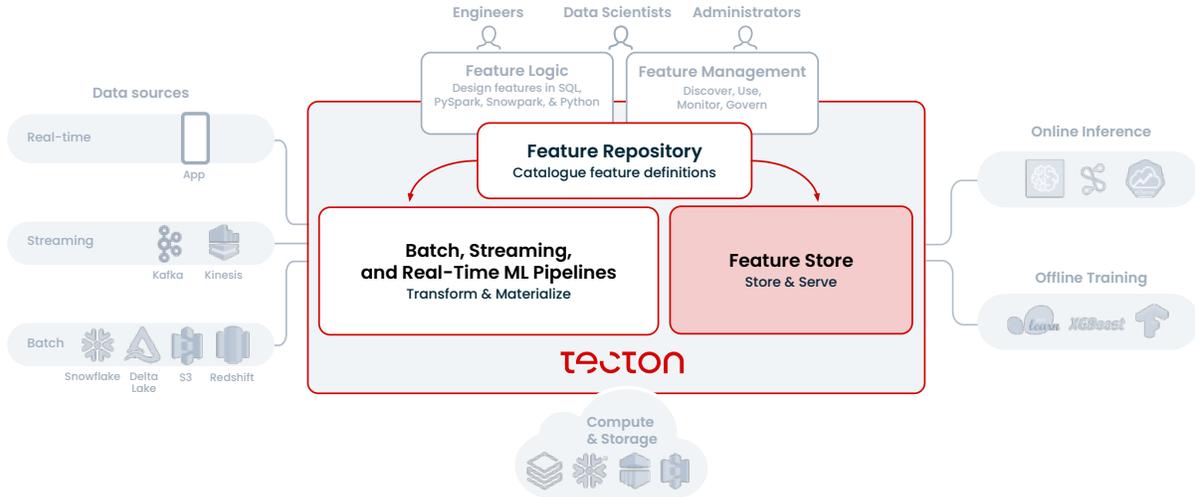
Centralize features: With Tecton, users centralize all the definitions for features in one place, enabling reusability and consistency across environments. Feature definitions and materialized values stored with Tecton are guaranteed consistent across training and serving environments to prevent training-serving skew.

Search for features: Tecton’s intuitive web UI enables users to quickly search and discover existing features across the organization so that they can easily re-use and refine features for new use cases and models.

Serve features: With a REST API, Tecton enables users to serve features at ultra low latency and scale to over 100,000 QPS. Tecton can also automatically backfill features to generate complete training sets through a simple Python SDK all the while ensuring correct time travel across training and serving environments to eliminate the risk of data leakage.

Monitor and govern features and systems: With Tecton, anyone on the team can monitor data pipelines, serving latency, and processing costs of their machine learning applications, and rely on Tecton to automatically resolve or report issues that need to be fixed.

From design to pipeline, in minutes.



Underneath the hood, Tecton stores, compiles, orchestrates, and automates ML feature pipelines for real-time machine learning that scales. Tecton’s platform uniquely manages both the data transformations that define features and the transformed feature data itself. Tecton turns feature logic into high-quality, production-ready data pipelines, exposed as building blocks for ML and analytics workflows.

Tecton registers feature logic

Tecton manages features as files in a git-like repository using a declarative framework. Users integrate CI/CD processes to the feature lifecycle that enables them to unit test features before deploying them to production, isolate workflows in separate production and development workspaces, and manage dependencies and version control of features across models.

Tecton executes transformation pipelines

Because it integrates seamlessly with other parts of the ML stack and connects to existing data processing and storage infrastructure, Tecton leverages your existing data platform and processing engines to automatically compile the underlying data pipelines that compute features, insulating the end user from their complexity.

Tecton orchestrates the underlying architecture

Tecton optimizes resource utilization and scales dynamically based on a user’s pre-defined requirements. Tecton scales compute, storage, and serving independently to adjust to usage patterns, leverages an offline store to optimize for large-scale and low-cost retrieval for training, and an online store for low-latency retrieval for online serving.

Tecton automates fresh feature computations

Customers rely on Tecton for providing uninterrupted access to their machine learning features. Tecton’s core IP implements feature transformations in a cost-effective way, and Tecton is priced on consumption so customers only pay for what they use.

Battle-tested for the enterprise.



Performance and reliability

With Tecton, teams can serve features in production at extreme scale with the confidence that systems will always be up and running. Tecton is built on best-in-class cloud services to maximize resilience and scale, delivering median latencies of ~5ms and supporting over 100,000 requests per second.

Fully managed service

Let Tecton manage your data transformation pipelines and centralized feature store with guaranteed SLAs and enterprise support out of the box. Tecton makes it easy to deploy and operate machine learning with a managed, cloud-native service.

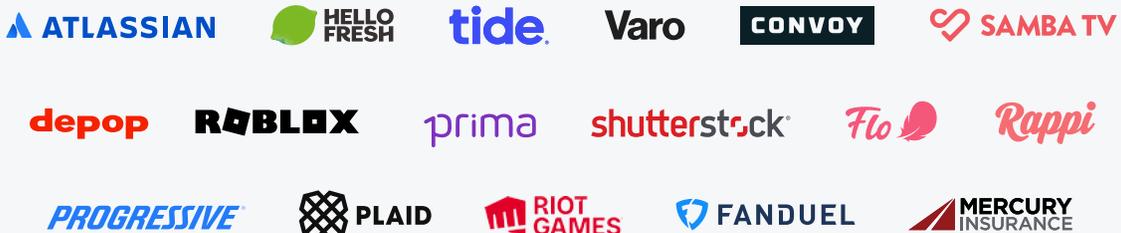
Compliance and security

Tecton meets strict security and compliance standards. Tecton authenticates users via SSO and includes support for access control lists, is SOC2 Type 2 certified, and ensures GDPR compliance in ML applications.

100% code-first solution

Tecton is designed for the engineers and data scientists who enjoy the freedom of a code-first approach. The platform makes any coders life easier by eliminating the many underlying hurdles involved in the path from ML prototype to predictive product.

Join other leading ML teams in their real-time ML journey:



The leading feature platform for real-time ML



Automate production pipelines for real-time ML in minutes

Save months of work by creating robust data pipelines from just a few lines of code—Tecton not only compiles the pipelines, it also orchestrates and maintains them so your teams don't have to. With Tecton, watch melding batch & real-time data, latency constraints, training / serving skew, or complicated backfills become problems of the past.



Increase model accuracy with higher quality features

Free ML engineers from wasting time on building pipelines so that they can focus on designing better features. Automate robust pipelines that materialize fresh features consistently across training and serving environments to drastically improve model performance.



Improve cross team collaboration and scale engineering use case capacity

Help data scientist and engineering teams work together by sharing features across the organization and standardizing all machine learning data workflows on one platform. Support more machine learning use cases without straining engineering investment or resources all the while improving data scientist productivity.



Ensure mission-critical reliability and scalability you can trust

Bring modern ML engineering best practices to your ML powered applications and rely on Tecton to ensure your systems will always be up and running. Serve features in production at extreme scale with median latencies of up to ~5ms and with the ability to support over 100,000 of requests per second.

Learn more at tecton.ai