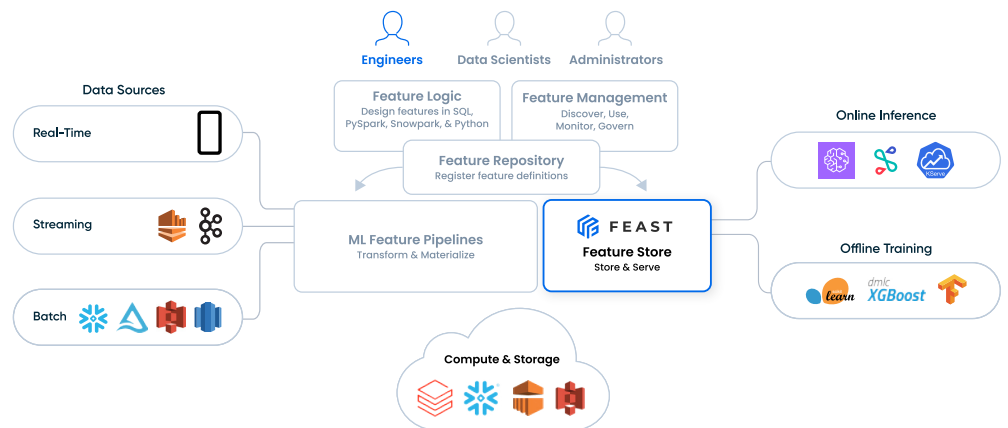


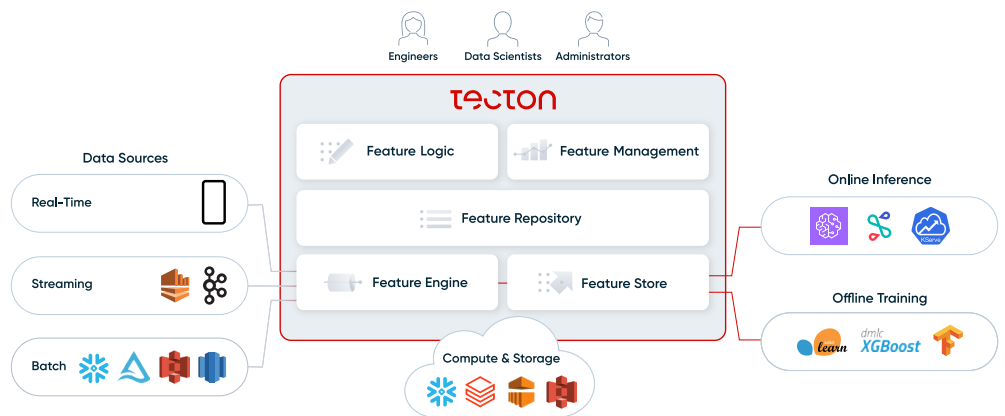
Choosing the Right Feature Store: Feast or Tecton?

OVERVIEW

FEAST
Feast is a self managed Feature Store
 Store and Serve Features for Production ML



TACTON
Tecton is a managed Feature Platform
 Design, Build, Centralize, Serve, and Manage Features for Production ML



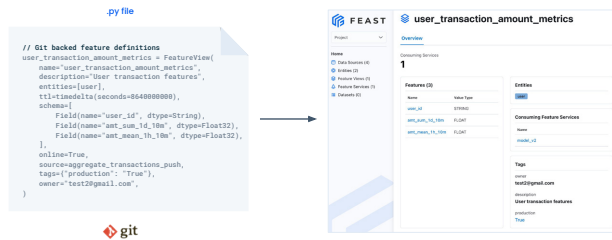
FUNCTIONALITY DIFFERENCES

	 FEAST	
Store and serve features		
Generate training data	✓	✓
Serve features online	✓	✓
Ingest from self-managed pipelines	✓	✓
Ingest from Tecton-managed pipelines	✓	✓
Automate pipelines		
Tecton-managed batch, streaming, and real-time pipelines		✓
Manage features		
Centralize definitions in common		✓
Discover, share and re-use		✓
Monitor data quality and operational service levels		✓

OTHER IMPORTANT CONSIDERATIONS

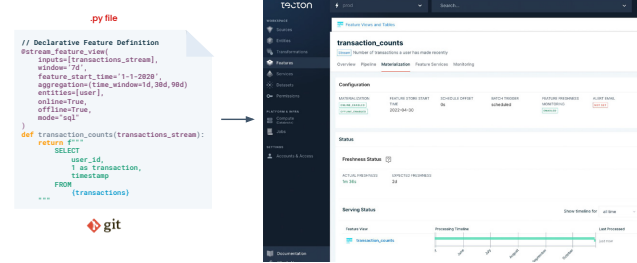
 FEAST	
Open Source Software	Fully-Managed SaaS
Self-configured and deployed	Fastest way to get started
Self-supported	24x7 enterprise support
Highly customizable	Plugs into leading data and ML platforms
Available on-prem or any cloud	Available on AWS today / GCP and Azure soon
Variable scale and reliability	Mission-critical scale and reliability
Serving latency: self-managed	Serving latency: <10ms
Scale: self-managed	Scale: >100,000 requests / second
	Production SLAs
	Discover, share and re-use
	Monitor data quality and operational service levels

CREATING FEATURES



Register a feature in Feast


Feast primarily helps in serving features for online inference and offline training, but **not** in transforming feature values (Feast takes transformed values as input). This is why creating a feature in Feast consists of registering a pre-processed feature, along with its metadata.



Declare a managed feature in Tecton

In contrast, when you create a feature in Tecton, you're **actually defining the feature transformation itself**. Tecton will manage the data pipeline for you and store the materialized data in online and offline stores, as well as run the backfill for you.

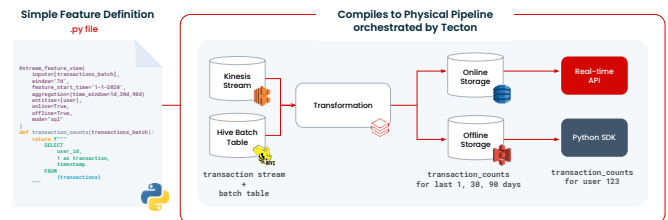
DATA PIPELINES



Feast decouples ML from data infrastructure

As mentioned above, Feast takes transformed values as input. This means you'll have to **manage your own data pipelines**. The typical steps include:

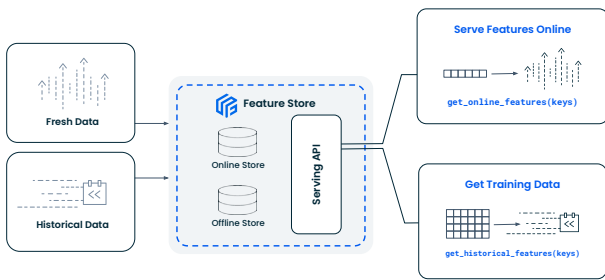
- Creating a streaming pipeline (e.g., with Spark)
- Creating a batch pipeline (e.g., with DBT)
- Setting up monitoring to make sure your features are properly sent to Feast and watch out for train / serve skew
- Setting up a backfill job with Airflow



Tecton creates production-ready data pipelines from simple declarative feature definitions

Because Tecton handles the feature transformations for you, it removes a lot of the complexity. There's no need to create pipelines, set up monitoring, or run Airflow jobs yourself—**Tecton will create these pipelines for you in the background, and handle monitoring and alerting as soon as you create the feature** (as described in the previous step).

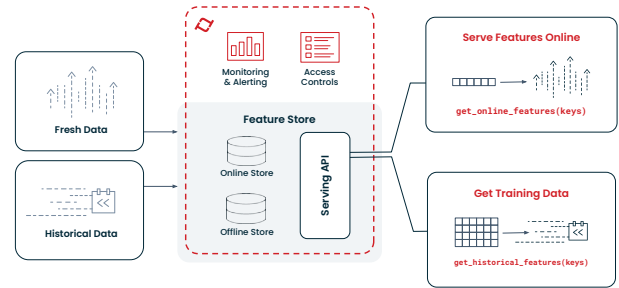
SERVING FEATURES



Serving features with Feast

Both Feast and Tecton allow you to serve features online and create training datasets. They share the same feature retrieval APIs:

- 'get_offline_features' to generate a training dataset with point-in-time correct features with labels
- 'get_online_features' to retrieve the features you need to make a prediction (for instance, to predict whether a transaction is fraudulent).



Serving features with Tecton

Similar to Feast, Tecton allows you to serve features online and create training datasets. Tecton has the same feature retrieval APIs as Feast. **The main difference is that Tecton manages the online store for you, which includes a layer of monitoring and alerting** (for example, if feature values stop coming through and your feature is getting stale, you'll receive an alert that your upstream data source might be down).

SCALING CONSIDERATIONS

Feat requires custom configuration to power production ML at scale

Common considerations include:

- How to reliably and scalably compute batch and / or streaming features
- How to reliably and scalably backfill features to the offline / online store
- How to efficiently retrieve features in a real-time recommender system
- How to manage online store costs given large data volumes

Tecton is a fully hosted solution

Tecton ensures enterprise scalability and reliability, with SLA guarantees.

RECAP

Use FEAST if you...

- Have the resources and skills to manage your own feature store
- Need a highly customizable solution
- Don't need any help with streaming or real-time data pipelines
- Need to deploy on-prem, on GCP, or on Azure

Use TECTON if you:

- Want minimal overhead in managing your own feature store
- Need mission-critical reliability, scalability, and/or support
- Want to automate batch, streaming, and real-time features
- Want to collaborate on, share, and re-use features

LEARN MORE

For more details, check out our webinar comparing Feast and Tecton resources.tecton.ai/choosing-the-right-feature-store-feast-or-tecton

Get a personalized demo of Tecton tecton.ai/product-demo/