

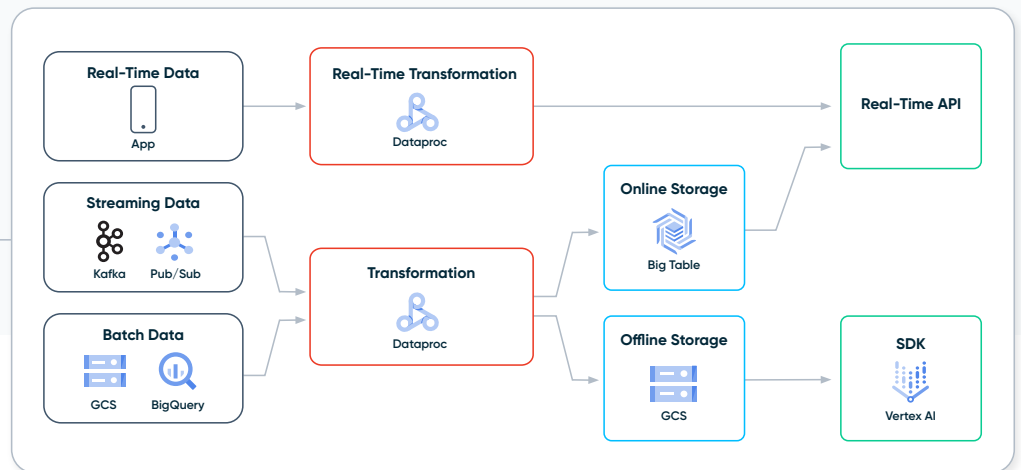
The fastest path to building ML-powered applications on Google Cloud Platform

Tecton and Google Cloud Platform (GCP) partner to accelerate the creation of ML-driven applications. This collaboration offers a cohesive ML solution for data scientists and engineers, combining Tecton's powerful feature engineering capabilities for production ML with Google's robust cloud computing services.

Simple Feature Definition .py file

```
// Declarative Feature Definition
@stream_feature_view(
    inputs=[transactions_batch],
    window='7d',
    feature_start_time='1-1-2020',
    aggregation=(time_window=1d,30d,90d)
    entities=[user],
    online=True,
    offline=True,
    mode="sql"
)
def transaction_counts(transactions_batch):
    return f"""
    SELECT
        user_id,
        1 as transaction,
        timestamp
    FROM
        {transactions}
    """
```

Compiles to Physical Pipeline orchestrated by Tecton



Build & iterate faster

- Deploy production machine learning pipelines in minutes
- Iterate on your ML models faster than ever before



Control & reduce costs

- Avoid overhead of building and maintaining your own infrastructure
- Optimize cloud spend



Maximize performance and reliability

- Increase your model accuracy
- Serve features at low latency and massive scale with enterprise-grade SLAs



Future-proof your ML stack

- Be prepared for the shift toward real-time data and predictions
- Get ready as more AI use cases start needing ML to construct prompts

Key Challenges for Production ML

Great models need great features. High-quality ML features are pivotal for the success of various machine learning applications, from detecting fraudulent transactions to making personalized recommendations. Creating and managing the necessary production-grade data pipelines is complex. As demands for data processing, transformation, training and serving availability, performance monitoring, cost-efficiency, scalability, and high performance grow exponentially, data teams struggle to tie together these services to quickly and

efficiently deploy production ML applications. The introduction of real-time data exacerbates this complexity. Common challenges include:

- deployment delays,
- exponential cost structures,
- inaccurate models, and
- isolated data pipelines that inhibit collaboration, flexible iteration, traceability, and scalability.

Solution Overview

Adopting Tecton alongside Google Cloud Platform's (GCP) powerful suite presents a compelling solution for organizations seeking to build and manage ML applications effectively. While GCP's advanced infrastructures like BigQuery and DataProc form the backbone of many data clouds, Tecton fills the gap, providing an efficient interface to rapidly transform raw data into high-quality ML features for production models.

A Complete Feature Platform: Tecton goes beyond being a simple feature store, offering a comprehensive platform that automates the entire lifecycle of ML features. With a simplified, declarative framework, teams can define features using SQL, Python, or Spark. This accelerates the process of building and managing production-grade features, integrating machine learning decision-making seamlessly into applications. Tecton stores and serves features for real-time inference and offline training, manages features as code,

orchestrates raw data transformation into production-ready features, and monitors feature data quality and operational service levels. This holistic approach streamlines ML feature management, enhancing efficiency and productivity.

A Platform Built for the Enterprise: Tecton provides an enterprise-ready solution for machine learning data management, offering robust performance, strong security measures, and multi-cloud capabilities. It has the scalability to handle over 100,000 requests per second with median latencies of around 5ms while ensuring consistent uptime and reliability. Tecton secures data with comprehensive access controls, and SSO integration, all while being SOC 2 Type 2 compliant and supporting GDPR compliance. With its multi-cloud capabilities, Tecton seamlessly unifies ML data operations across various cloud platforms, thereby improving efficiency and interoperability.

How it Works

Tecton and GCP work together to make machine learning features and model management more efficient. Tecton's feature engineering capabilities help define and manage features in a scalable and automated way, ultimately allowing for faster and more efficient model training and deployment. Together with GCP's leading services for building and running ML

applications, including Vertex AI, Kubernetes, and TensorFlow, and advanced data infrastructure like BigQuery and DataProc, organizations can simplify the process of preparing, managing, and serving data for machine learning models, enabling them to make more accurate predictions and encourage data-driven decision making.

How It Integrates

Underneath the hood, Tecton seamlessly integrates with existing GCP infrastructure, accommodating batch, streaming, and real-time sources such as BigQuery, Google Cloud Storage, Apache Kafka, and Google Cloud PubSub. It leverages fast, scalable, and cost-effective GCP data processing services like DataProc for data transformations. The platform stores feature data offline using Google Cloud Storage, and

offers low-latency online storage via Redis Enterprise or BigTable. Finally, Tecton is not only compatible with preferred notebook environments, including Vertex, or Jupyter, for feature building and training dataset generation, but it also connects seamlessly with preferred model serving layers, including Vertex AI.

Get started today!

Request a demo or a free trial at tecton.ai.