

Case Study

Scaling Blockchain Analytics on Avalanche L1 subnets with Codimite & Google Cloud



The Challenge

Feature.io, a dynamic Web3 platform, faced mounting pressure to scale their data infrastructure. With an expanding user base and increasingly complex digital asset campaigns, the company needed a way to:

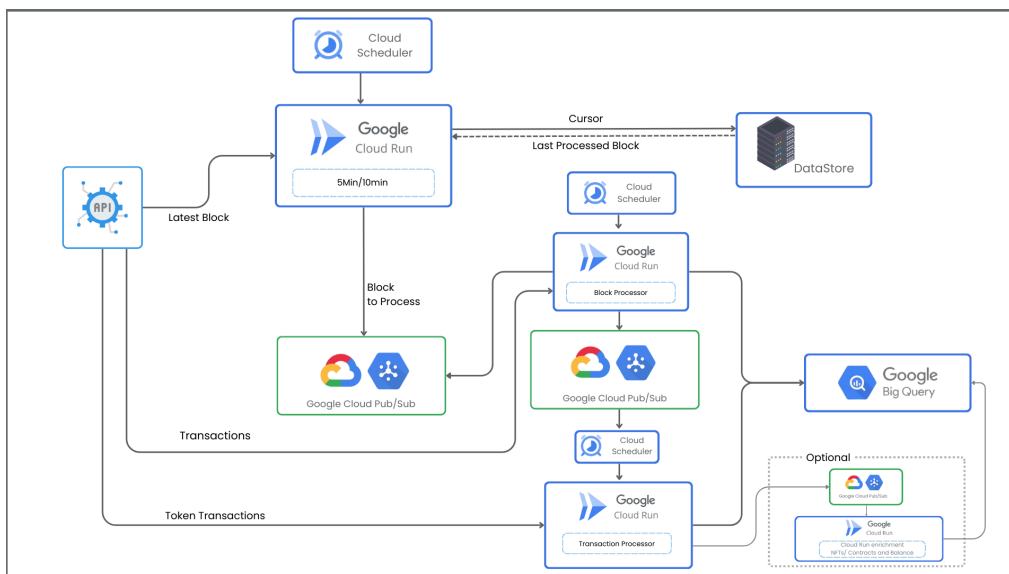
- Reliably **ingest Avalanche L1 / subnet blockchain data** for advanced analytics.
- Ensure **scalability** as blockchain activity grows.
- Data cross-pollination from other information systems to provide **real-time insights** to marketing, operations, and product teams,
- Maintain **tight security, observability, and uptime** in a decentralized environment.

Their existing system lacked the robustness and flexibility to support the demands of data-hungry business users. Key stakeholders couldn't easily access on-chain activity, campaign performance, or user segmentation data—limiting data-driven growth strategies.

The Solution

Codimite engineered a **GCP-native, serverless ETL pipeline**, purpose-built to meet Feature.io’s analytics goals. The solution centered around modern, composable cloud services:

- **Cloud Run** for scalable, containerized services handling block ingestion and enrichment.
- **Cloud Scheduler** to orchestrate timely polling intervals (5–10 minutes).
- **Pub/Sub** as the messaging backbone to decouple ingestion and processing stages.
- **Cloud Datastore** for cursor tracking, ensuring no duplication or missed blocks.
- **BigQuery** as the centralized analytics warehouse, structured with partitioning and clustering for high-performance querying.
- **Secret Manager** for secure key and credential handling.



To complete the ecosystem:

- **Cloud Monitoring & Alerting** was set up for uptime tracking, pipeline health, and error notifications (email).
- **Audit logs** and compliance features were embedded to meet enterprise-grade standards.
- **Dashboards** were prepared in BigQuery and Cloud Monitoring to enable easy tracking of smart contract usage, NFT interactions, and user activity.

This architecture was fully walkthrough’d with the client, receiving **positive feedback** for clarity, robustness, and security. Documentation and handover are set for **September 18th**.

The Results

The implementation delivered a **highly scalable, secure, and reliable blockchain analytics pipeline** for Feature.io:

- **Streamlined ingestion** of blockchain data from the Avalanche network.
- **Stable and reliable operations**, with monitoring and cursor recovery for resiliency.
- **Performance-optimized analytics** using BigQuery for deep event queries.
- **Improved stakeholder visibility**, allowing product and marketing teams to track engagement metrics, campaign completion rates, and token usage patterns.
- Feature.io team praised the architecture for its **clarity, scalability, and operational simplicity**.

Future Opportunities

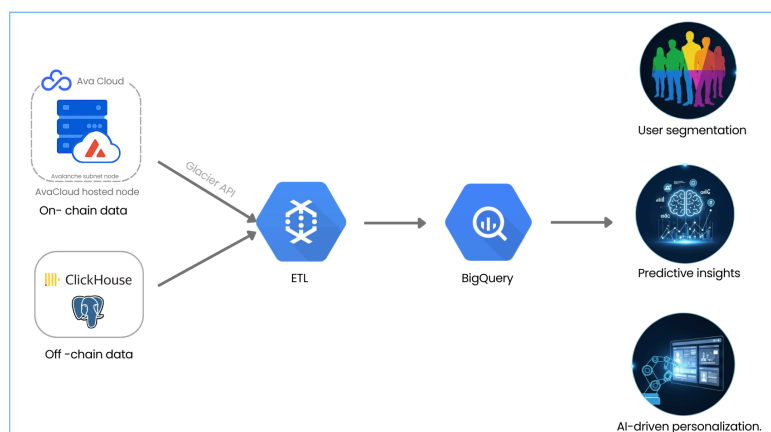
Phase 2 is a possibility on the horizon.

Codimite and Feature.io might plan to **extend the ETL pipeline** to integrate **ClickHouse and PostgreSQL**—enabling a unified view across:

- **On-chain data** (wallets, NFTs, smart contract usage).
- **Off-chain data** (user behavior, campaign metadata, email engagement).

This unlocks the ability to:

- Segment users across **full-funnel journeys** (on- & off-chain).
- Deliver **hyper-personalized marketing** and predictive product insights.
- Power **AI-based recommendation systems** for fan campaigns and brand partnerships.



Conclusion

By rethinking their infrastructure with Codimite and Google Cloud, Feature.io is no longer just storing blockchain data—they're **activating it**. With the handover of Phase 1 imminent and Phase 2 in motion, they're on track to become a **data-first Web3 platform** capable of delivering intelligent, real-time insights to every business team.

Power Your Web3 Analytics with Codimite

Whether you're running a DeFi protocol, launching digital collectibles, or powering fan experiences like Feature.io, your data infrastructure **shouldn't hold you back**.

With Codimite's expertise in **serverless cloud architectures** and **blockchain data engineering**, you can:

- **Ingest and transform data** from any EVM-based chain.
- Build **unified analytics platforms** combining on-chain and off-chain insights.
- Deliver **real-time, actionable intelligence** to your product and marketing teams.

Ready to scale your Web3 analytics?

Let's talk: codimite.com | info@codimite.com