From Blockchain to Dashboard: Realtime on-chain Analytics on <u>Azure</u>

How I built a streaming data pipeline for Ethereum Uniswap transactions



The problem: Blockchain data is huge — but hard to use in real time.

- Ethereum has massive transaction volume (1+ million tx/day)
- Swaps on Uniswap are high-value events (\$500M-\$1B daily)
- How do you monitor this *live*?



I wanted to track Uniswap swaps in real time...

- Stream Ethereum events using Kafka
- Ingest data into ClickHouse, a highperformance OLAP database
- Build a real-time Grafana dashboard
- Deploy to Azure following Cloud
 Design Patterns and best practices



The Architecture on Azure



Real-time on-chain analytics of transactions on the Ethereum Virtual Machine (Architecture on Azure Cloud)



Breaking down the Architecture

- Event Hub: Kafka-compatible and fully managed message queue
- Container Apps: fast deployment for Producer, ClickHouse and Grafana
- User-assigned Identity for increased security between services



What I built

- Dockerized Producer, Consumer (ie. ClickHouse database) & Grafana dashboard
- The Producer app is a python script that reads live blockchain events and sends them in batch to the Event Hub
- Infrastructure-as-Code with Terraform



Grafana Dashboard

Home > Dashboards > Uniswap > Live ETH transactions

🙄 Refresh

5s

Q

② Last 5 minutes

ETH and USD value per block timestamp Swaps per protocol per block 2500 60 20 50 2000 21 40 1500 30 1000 20 10 500 10:20:30 10.18.00 10:18:30 10.19.00 10:19:30 10:20:00 10:21:00 10:21:30 10:22:00 10.22.30 🗕 ETH 🗕 USD Latest transactions on Uniswap V2 and V3 transaction_hash address timestamp block_number protocol from to value 0x76366D95c20164462472 1743841331 0x23D29EecE9c878d35e7b 0x509b52749c2991b7472B3 0x18e96e2d4ae6ff567... 22201490 uniswap v3 0 0x790816379043a68a... 0x7f74c86CD8BA48be647d 1743841331 22201490 uniswap_v3 0x93793Bd1f3e35a0Efd098 0xfBd4cdB413E45a52E2C8 2.22e-11 0x88e6A0c2dDD26FEEb64F 1743841331 0xd17b8cbfee34c6d27... 22201490 uniswap_v3 0 0xE743a49F04F2f77eB2D3I 0xce16F69375520ab01377c 0xC3d1E92A5C1c2f070e6ds 1743841331 0xe3199dbc2cec6d61... 22201490 uniswap_v2 0 0xeD5A36346554F6A7CFD 0x7a250d5630B4cF539739 0xFFf8D5fFF6Ee3226fa2F5 1743841331 0xE1851E1c10147c12a22A7a(0xb300000b72DEAEb607a1 0xec1e37e5d03b7faac... 22201490 uniswap_v2 0.0150 1743841331 0x6b2bb99a1f45c809... 0x6dcba3657EE750A51A13/ 22201490 uniswap_v3 0 0xD6C645Db7Dbce1099A9: 0x66a9893cC07D91D95644

This dashboard allows to monitor in real time:
Number of swaps per block

Total ETH swapped and \$ value per block



Main challenges & What Hearned

How to build a real-time event-driven architecture, how to authenticate and set up the Consumer to read events from the Event Hub

How to deploy a multiple environment cloud architecture using Terraform modules, how to Don't Repeat Yourself (DRY) using Terragrunt, how to automate using the GitOps framework



What's next?

Streaming real-time blockchain data unlocks powerful on-chain insights:

- Understanding market behavior
- Identify emerging trends early
- Detect whales or suspicious activity
- Enable instant alerts and dashboards
- Build better strategies for trading bots



Want to build something similar?

Leave a comment, DM me, or check out my website: <u>www.nicolasguillaume.com</u>

Thanks for reading 😊



Nicolas Guillaume MSc. Data Engineer | Cloud Engineer | DevOps