



**Mind the gap, connecting high performance systems at a
leading Crypto Exchange**

Bitvavo

2024

Lars Werkman and Marcos Maia



- About us
- Bitvavo
- Stack & High level Arch
- High Performance Systems
- Cold-sync
- Optimizing signing
- Azul Prime
- Q&A

Agenda



Marcos Maia
Senior Staff Engineer @
Bitvavo



From Brazil



Settled in The
Netherlands



Distributed Systems

**Startups/Scaleups, Investments(Trading),
Fintech, Banks and Telecom**



Lars Werkman
Staff Engineer @ Bitvavo



Born and raised in the netherlands

Distributed Systems

Startups/Scaleups, Fintech




bitvavo



Bitvavo

Trading platform



Focus

Trading Platform

Most of the Crypto value comes from investment and belief in the future adoption. Crypto is an asset class similar to commodities or stocks and our role is to give people high quality access and exposure to these assets through trading and investing.

Crypto Bank

As Crypto and blockchain become adopted as an asset class, financial systems will start appearing around them. Stocks, bonds, DeFi, payments - all the typical financial services will be adopted to blockchain and Bitvavo will facilitate access to that financial system based on Crypto and blockchain.

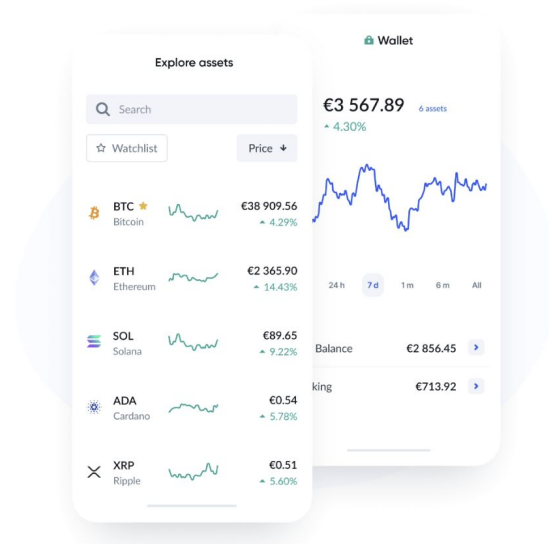
Crypto Gateway



As new use cases are being adopted within Crypto and blockchain rails, Bitvavo will help customers to access non-financial services based on the blockchain and Crypto. For example, in the future you may register the purchase of your real estate in a public registry through Bitvavo.

Trade the future

Buy, sell, and store over 275 digital assets at one of Europe's leading exchanges.

Start Trading →



Asset	Price	Change (24h)	Chart
 Bitcoin BTC	€62,776.5	+0.54%	

BTC-EUR

/

Last price (EUR)
62,819 ▼
24h change
+0.41%
24h volume (EUR)
16,303,920.18

Share feedback
⋮
Marcos ▼

Trade

Buy
Sell

Limit
Market
Stop limit

Available balance 0 EUR

Limit price 62,820 EUR

Amount 0 BTC

Total 0.00 EUR

0% 100%

Preview buy order

Order book

Price (EUR)
Amount (BTC)

62,823	0.03869694
62,822	0.18158529
62,821	0.03870094
62,820	0.31737244
62,819	0.01447000
62,818	0.07939864
62,819 EUR 0.00% spread	
62,815	0.02870754
62,814	0.27126000
62,813	0.14009169
62,812	0.29323000
62,811	0.29009754
62,810	0.03680492

BTC-EUR

D
↑↓
fx Indicators
 ⌵ Show Orders
 ⌵ Show Fills
 ↶
↷ Draft
 ⚙️
📏
📷

Trading view
Depth

BTC-EUR · 1D · Bitvavo ● O 62,563 H 63,286 L 62,066 C 62,819 228.00 (+0.36%)

Volume 252.27212826


9 13 17 21 25 Jun 5
6m 3m 1m 5d 1d 🔄
14:17:32 (UTC+2) | % log auto

Open orders

Asset	Price	24h change	▼ Euro value	Balance	Available balance	Booked for orders
🔗 BTC	62,819 EUR	0.41%	2,694.91 EUR	0.04289957 BTC	0.04289957 BTC	0 BTC

European Market Leader

As the **biggest crypto asset service provider in the Benelux**, we aim to expand into other EU member states to become the biggest European leading crypto asset service provider.



Some Numbers


- Market Leader in Benelux
- Close to 2 million customers
- 24/7 Business
- 230+ Assets
- ~5k orders / second(many more requests)
- ~20k Web Socket connections
- 1 Financial app in Benelux
- Leaders in EUR/Crypto pairs volume





Stack and High Level Architecture

Services



Our Stack






High Performance Systems

Low latency on JVM

- No DB on “hot path”
 - Fully in memory
 - Eventual sync to disk
- Inter machine communication
 - Optimized Inter process communication (IPC)
- Network Communication
 - UDP
 - Dual card machines (2 Nics)
 - Fine tune OS and Network Buffers
 - Tune MTU package size
 - Linux DPDK (Data Plane Development Kit)
 - Fast packet processing



Low latency on JVM

- No containers or Kubernetes
 - EC2 instances
 - Bare Metal Instances
- Distributed State Machine
- Single threaded State Machines
 - Avoid context switch
 - Deterministic
- App thread pinning to dedicated cores
 - Leverages low level CPU caches

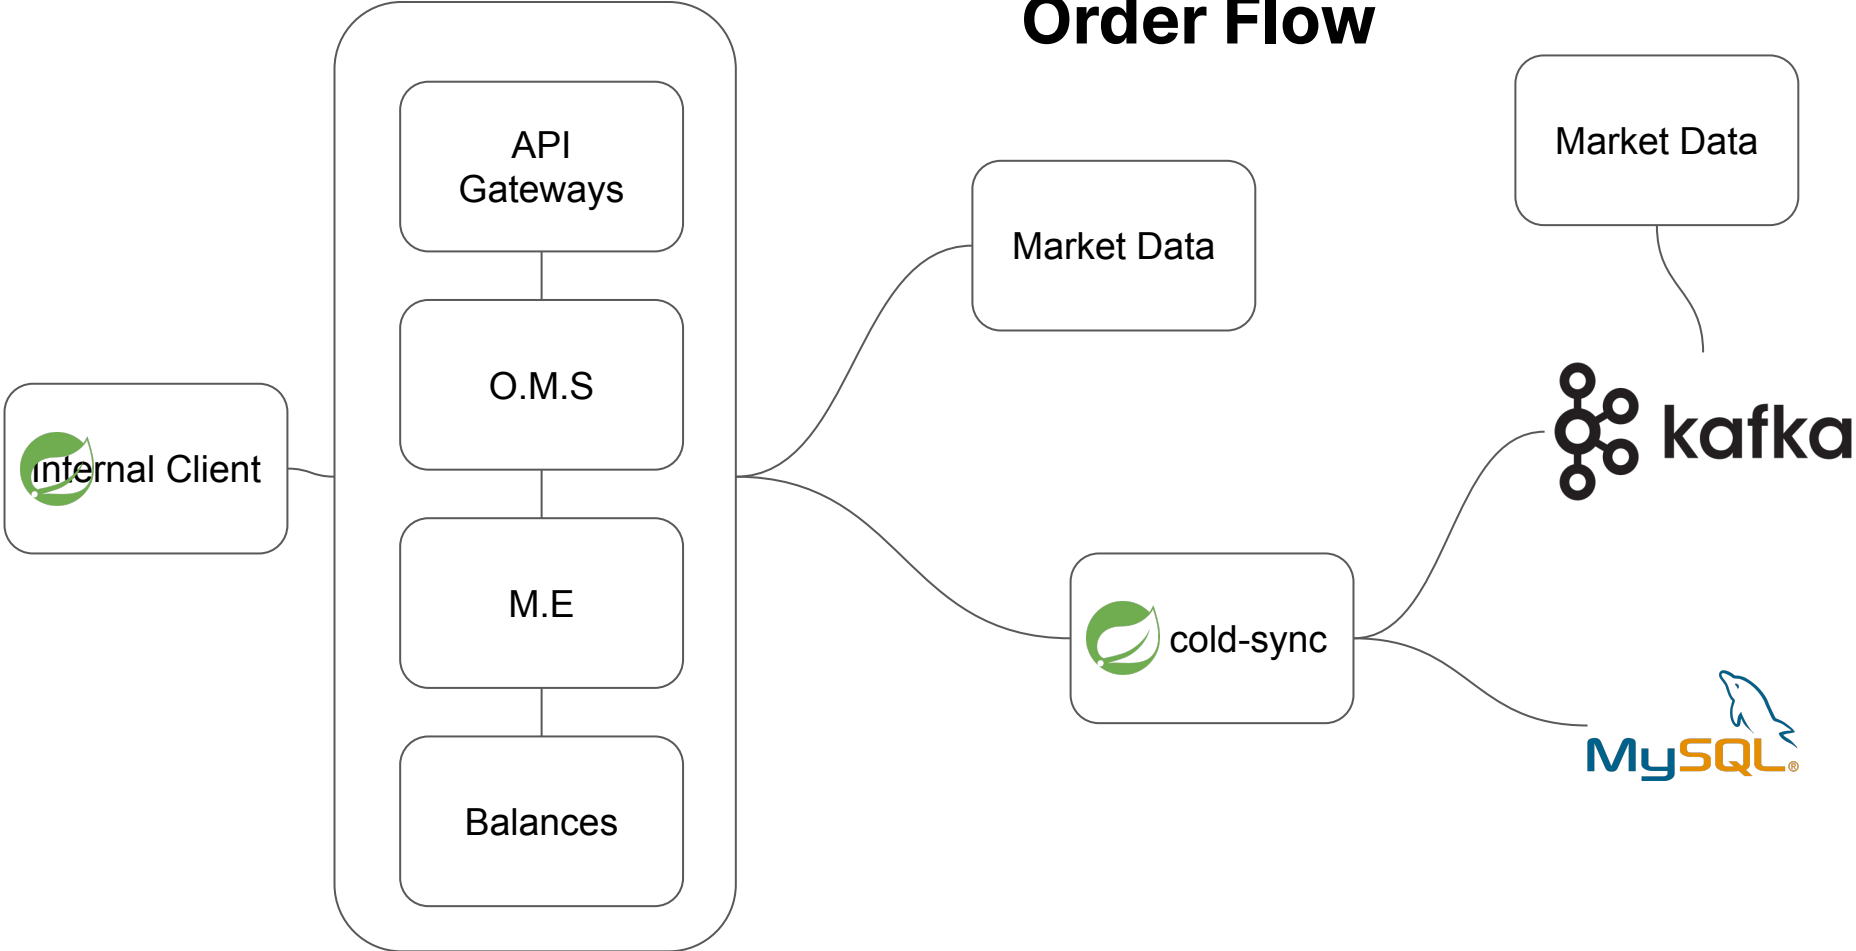


Design Principles

- Garbage Free in steady state running
- Smart batching in the message path
- Lock free algorithms in the message path
- Non blocking IO in the message path
- • No exceptional cases in the message path
- Single Writer principle
- Heavily Prefer unshared state
- Avoid unnecessary data copies
- **No logging**
- **No Strings**



Order Flow

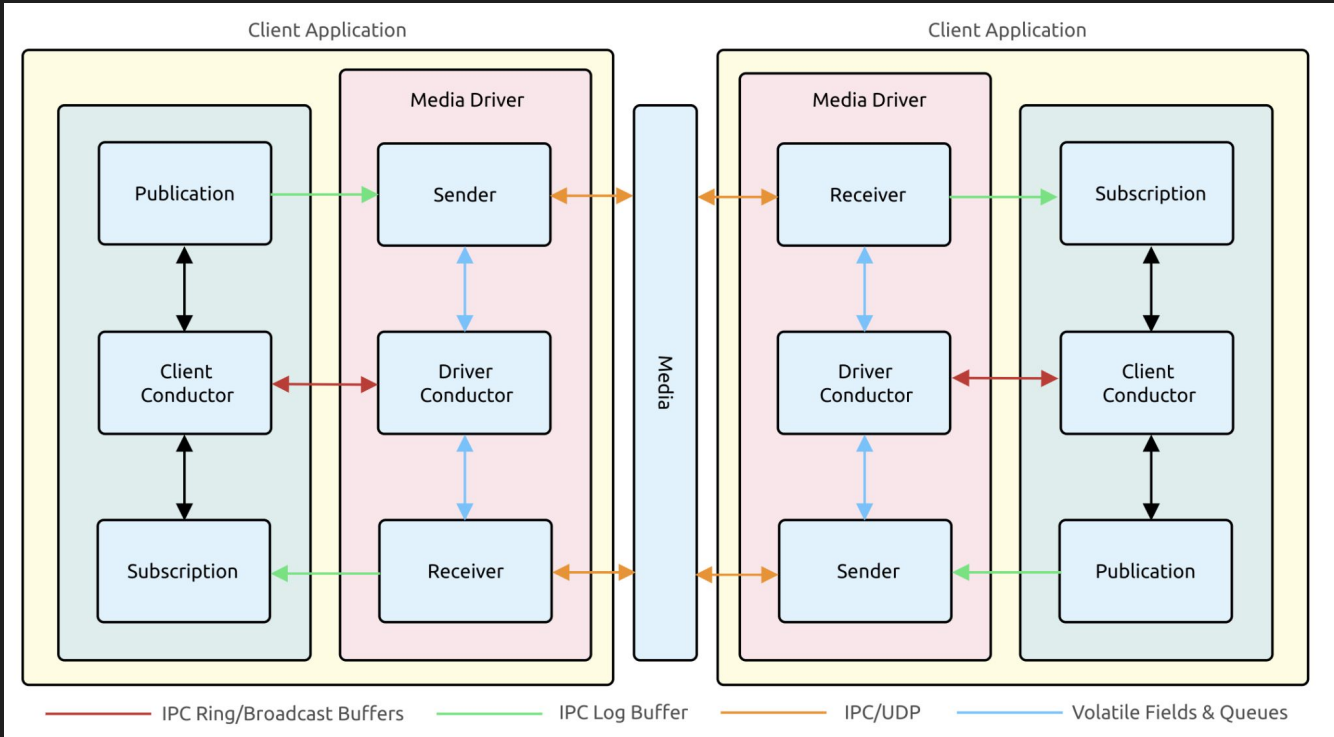




Aeron Framework

Aeron

* replicates ordered log buffers across process or network boundaries with predictable latency.




Aeron Cluster


- Implementation of Raft Consensus
- sequences multiple client connections into a single, replicated log
- provides efficient fault tolerance with 2 or more nodes
- provides Aeron Archive endpoints for writing to snapshots and reading from snapshots on start
- allows 1 or more clustered services to be run, with support for inter-service sequenced messaging
- reliable, sequenced timers
- very high levels of performance




Active Cluster




TUB Node




TUB Node



TUB Node

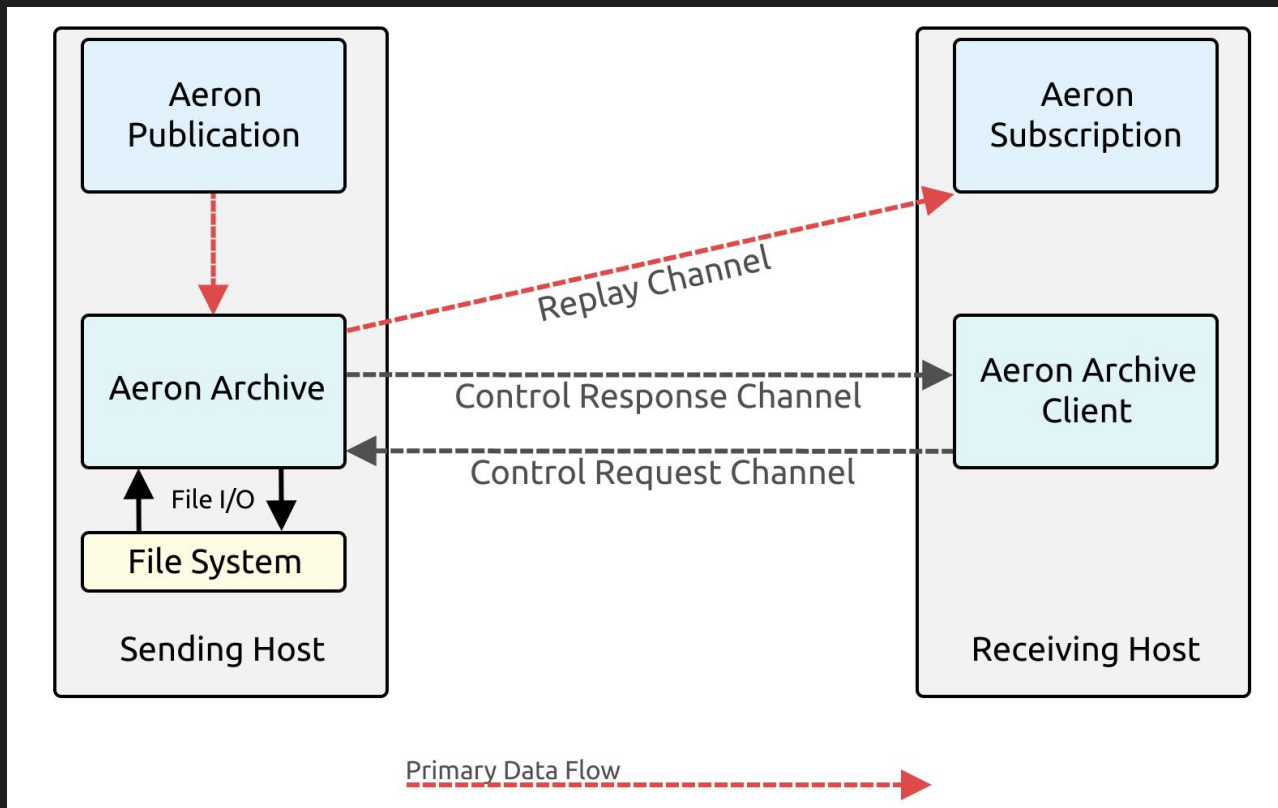


TUB Node



TUB Node


Aeron Archive

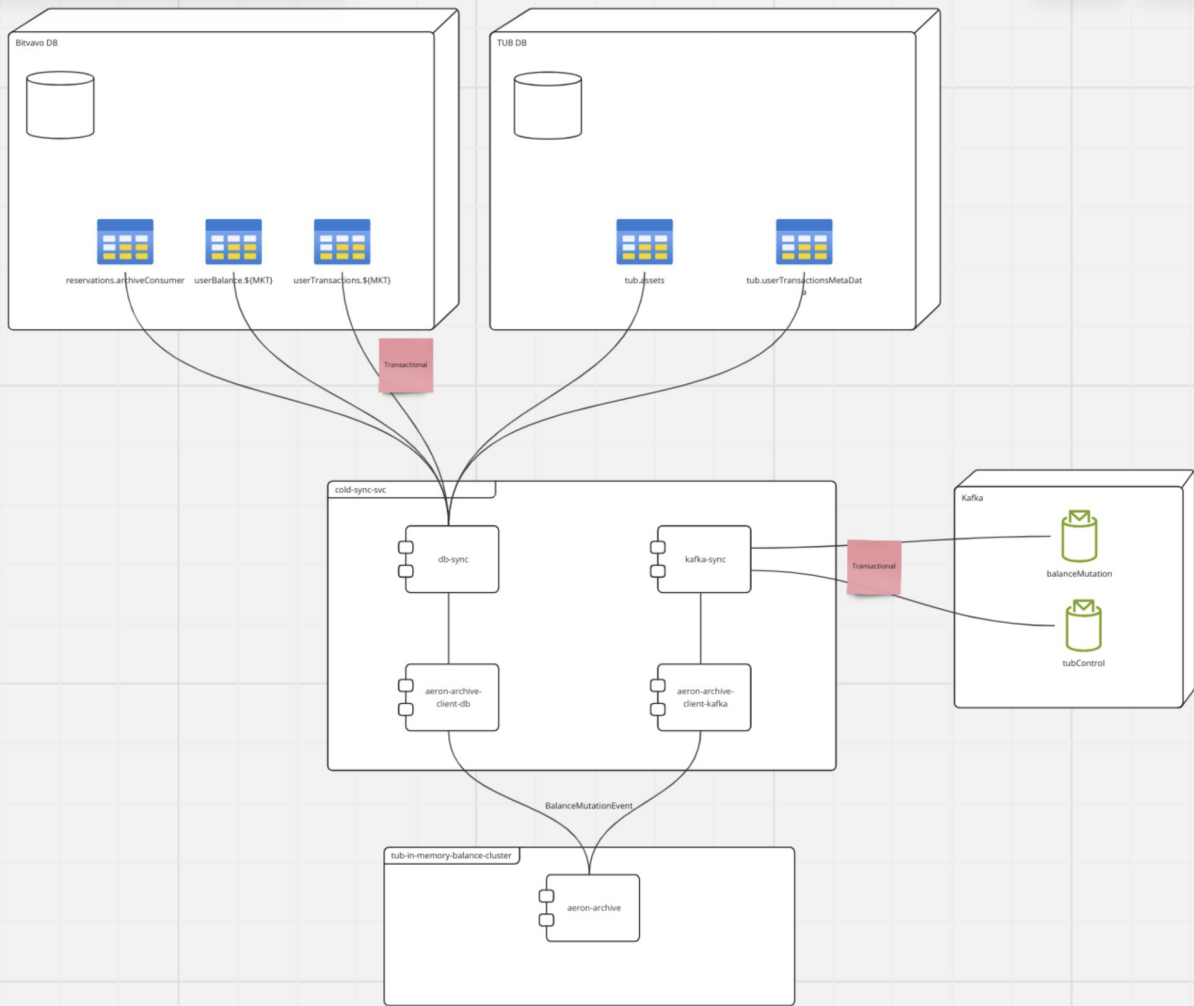




Cold-sync

Order Flow





For performance

- Use Batching
 - Fine tune size
- Use Agents
- Reuse Object Pools
- Apply backpressure
- Smart retries
- Use metrics / counters
- Minimize logging
- Fine Tune your DB
 - For the use case
 - Query
- No JPA
 - Template is fine
- Control offsets





Optimization use case - signing


Optimizing JWT Signing



30x



Goal: P99 > 1 ms



Encryption types

	Symmetric Encryption	Asymmetric Encryption
Speed	Fast	Slow
Security	Single secret	Public and Private key pair
Key length	Shorter key length	Longer key length
Popular Signing Algorithm	HMAC (SHA256 / SHA512)	RSA (SHA256 / SHA512)

Benchmarking using JMH

Benchmark	Mode	Score	Error	Units
HS256	thrpt	334.985 ±	1.313	ops/ms
HS512	thrpt	260.551 ±	2.851	ops/ms
RS256	thrpt	1.462 ±	0.005	ops/ms
HS256	avgt	0.003 ±	0.001	ms/op
HS512	avgt	0.004 ±	0.001	ms/op
RS256	avgt	0.681 ±	0.001	ms/op
HS256:p0.99	sample	0.004		ms/op
HS256:p0.999	sample	0.011		ms/op
HS512:p0.99	sample	0.005		ms/op
HS512:p0.999	sample	0.010		ms/op
RS256:p0.99	sample	0.732		ms/op
RS256:p0.999	sample	0.842		ms/op
HS256	ss	8.275 ±	38.567	ms/op
HS512	ss	8.119 ±	37.721	ms/op
RS256	ss	10.735 ±	40.275	ms/op

```
@Benchmark
fun HS256Signature(blackhole: Blackhole) {
    val token = Jwts.builder()
        .signWith(HS256, Jwts.SIG.HS256)
        .compact()

    blackhole.consume(token)
}
```

*The Blackhole ensure that our token won't get optimized out due to dead code optimization

Optimizing further with Azul Prime

- Enhanced build of OpenJDK (Zing)
 - Low consistent response latency
 - Higher total throughput and carrying capacity
 - Faster Warm-up
- The technologies delivering these results
 - C4 Pauseless garbage collector
 - Falcon JIT Compiler
 - Ready Now / Cloud Native Compiler

Preparing our application for Azul Prime

In 2 easy steps

Building the base image

```
FROM azul/prime:17
```


Configuring the JVM options

```
args: ["-c", "java -jar -Xmx10g application.jar"]
```


bitvavo

Immediate Benefits


Lower request times



Stable Memory Consumption



Predictable CPU Usage



Whats Next

- Integrating ReadyNow in zero downtime deployments in kubernetes
- Fine tuning the JVM options for each individual application



We're hiring!

Start your crypto career with Bitvavo

Innovation seeker? Tech junkie?
Sounds like a great fit.

Join us

<https://jobs.bitvavo.com/>





Thank you!



Pictures used in this presentation

Photos:

by [Sajad Nori](#) on [Unsplash](#)

by Massimo Botturi on Unsplash

by Krzysztof Hepner on Unsplash

by Frankie Lopez on Unsplash