

# HARNESS THE POWER OF ALL YOUR DATA

## WHY WE EXIST

Every day, enterprises face difficult business and technical tradeoffs caused by the explosive growth of mission-critical data. They must either pay through the nose for storage, throw away valuable data, or live with poor performance. Hydrolix was founded to change the economics of big data and eliminate these tradeoffs. Our Cloud Analytics Engine can ingest, store and query vast amounts of data at low cost without sacrificing interactive performance.

## FOR DEVOPS, SECOPS, AND DATAOPS

### **Hydrolix is an elastic cloud data platform built for observability.**

- › Designed for high-volume, append-only workloads.
- › Seamlessly combines logs, metrics and traces into multi-schema tables.
- › No-cost ingestion of live stream and out-of-order data.
- › Native support for unlimited dimensions & cardinality.
- › Fast and easy to deploy (and free to try) inside your own VPC.

## UNLIMITED DATA RETENTION


### **Hydrolix's decoupled architecture means you can store all your data at a fraction of the cost of other solutions.**

- › Patented high density compression slashes costs and enables unlimited retention.
- › Eliminates tiered storage by making everything hot.
- › On-prem deployment means zero egress costs and full data control.
- › Scales to zero at any time—no reliance on caching or re-partitioning of data.

›› START YOUR FREE TRIAL TODAY AT [hydrolix.io/try](https://hydrolix.io/try)

## Hydrolix delivers SSD-like query performance from inexpensive cloud storage.

- › Micro-indexing and predicate pushdown techniques replace brute-force column scans.
- › Read throughput of 100M values per second per core.
- › Optimized for ad hoc queries and historical analysis.
- › Instant scale up to take advantage of massively parallel cloud resources.
- › ANSI-compliant SQL API (including ClickHouse ML extensions).
- › Bundled with Grafana for ease of use.



Query latency (seconds)

SOURCE: [tech.marksblogg.com/benchmarks.html](https://tech.marksblogg.com/benchmarks.html)

## Hydrolix + OpenTelemetry + Grafana for Observability

	Hydrolix	Cloud Data Platforms	Time Series Databases	ELK Stack	APM Services
Decoupled storage	✓	✓	✗	✗	✗
Eliminates tiered storage	✓	✓	✗	✗	✗
Unlimited dimensionality, unlimited cardinality	✓	✓	✗	✗	✗
Free stream ingest	✓	✗	✓	✓	✗
Out-of-order data	✓	✓	✗	✓	✗
Self-defined messages	✓	✗	✗	✗	✓
Multi-schema tables	✓	✗	✗	✗	✗
Indexed search	✓	✗	✓	✓	✓
Interactive performance	✓	cache only	✓	✓	✓
Cloud on-prem (VPC)	✓	✗	✓	✓	✗
Zero egress charges	✓	✗	✓	✓	✗
		ADX Athena BigQuery Snowflake	Influx Prometheus M3DB Timescale	Elasticsearch	DataDog New Relic Splunk SumoLogic

