


# Meet the Analytics Lake

Integrate raw data from data lakes and structured data from warehouses for advanced analytics, BI, and accurate predictions.



When it comes to data platforms, data teams are spoiled for choice, with many well-established companies offering powerful, one-size-fits-all solutions. But while these platforms may excel as general-purpose tools, they often lack the integration and automation needed to meet the evolving demands of advanced analytics organizations. The result? Missed opportunities, delayed deadlines, and shockingly high total cost of ownership.

While emerging approaches like **data product thinking** aim to address these challenges, we have to ask ourselves: is our technology ready? In response to this, we've created a solution that shatters the limitations of traditional data platforms.

## What Is the Analytics Lake?

The [GoodData Analytics Lake](#) is created on a foundation of modern, scalable, and highly efficient technology. With built-in automations and end-to-end integrations, it empowers you to rapidly develop data and analytics products, delivering superior outputs and dramatically lowering your TCO.

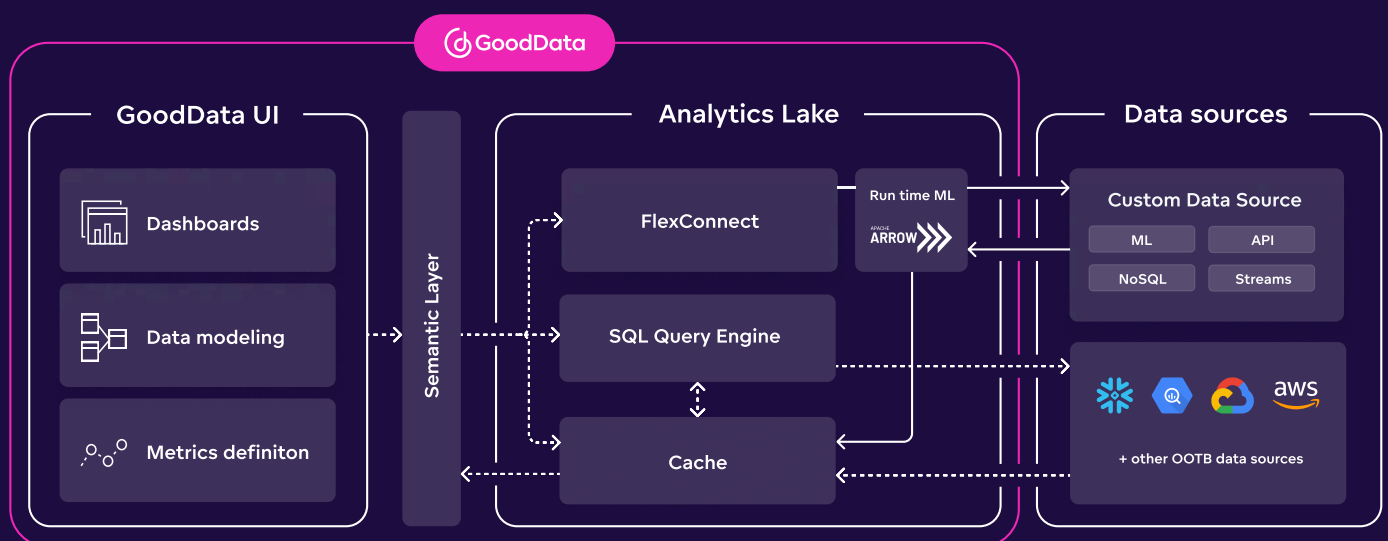
The Analytics Lake combines data storage, compute, and semantics into a single platform designed specifically for analytics purposes and nothing else. It comes pre-configured and integrated with [GoodData Business Intelligence](#) and [GoodData Artificial Intelligence](#) to offer **end-to-end data** and **analytics product capabilities**, but is open to work with any existing upstream and downstream technology.

Most importantly, it is ready-made for teams that embrace product thinking and engineering-style solutions to data problems.

# What Sets Our Approach Apart?

The Analytics Lake moves beyond monolithic data warehouse-only solutions to incorporate many features in a single platform, including in-memory compute, semantics, and CI/CD integrations.

- ✓ The best **open-source foundation** including Apache Arrow, Apache Iceberg, and DuckDB in an integrated system.
- ✓ **Holistic solutions** that combine the data, what it means (semantics), and its user-facing representations (BI & AI) in one platform.
- ✓ Achieving **security and scale** via engineering techniques and analytics-as-code.
- ✓ **Data product & design thinking** as the cornerstone of analytics delivery.
- ✓ One platform **deployed anywhere** — as SaaS, in your cloud, or even on-prem.
- ✓ **Aligning costs**, with value delivered, not compute cycles generated.



**The Analytics Lake is not just a philosophy**, it's a technology platform. It offers engineers and architects a variety of ways to optimize performance and cost with both efficient disk and super high-performing in-memory storage for data:

- 1. Cost-effective storage in open table format** utilizing Apache Iceberg.
- 2. Advanced, high-performance in-memory cache** and data materialization using Apache Arrow.
- 3. Automated, distributed in-memory data marts** driven by DuckDB.
- 4. Semantic, metrics, and metadata modeling layer.**
- 5. Endpoints for developers and users** including APIs, SDKs, BI interfaces, and AI integrations.
- 6. Advanced developer tooling** including CI/CD integration, analytics-as-code, and build your own data source capabilities.

## The Next Step in Your Data Journey

Our goal is to use the best technologies to **improve performance, reduce unnecessary spending**, and **fill in all the gaps** that slow down your data team and frustrate your end users.

By combining our 'built for analytics and nothing else' Analytics Lake with your existing infrastructure you optimize the only thing that truly matters — your ability to access the right data products with the speed and flexibility necessary to make a real impact on your business.

**Want to try it out?** Get in touch for early access now.