

Yellowbrick for Hedge Funds

The world's fastest data warehouse for hybrid and multi-cloud environments gives Hedge Funds and portfolio managers a sharpened competitive edge with real-time speed, petabyte-scale true elasticity, and industry-leading deployment flexibility.

Your Data Anywhere



Overview

Data has always been core to Hedge Funds' ability to beat the market and stay ahead of the pack. But the volume and breadth of data needed to drive trading analytics have exploded in recent years, driven by the market's embrace of Big Data and the emergence of thousands of new "alternative data" sources.

Hedge Funds, both those pursuing quantitative approaches and fundamentals-based approaches, are adding these datasets alongside traditional trade and quote (TAQ) data. This data can be derived from industry participants' core business from social media, web traffic, consumer reviews, shipping data, weather, and a host of other sources. All this data is being added to trading and investment models as they broaden searches for alpha beyond traditional trade and quote data – in an attempt to gain that unique edge.

Hedge Funds are rethinking their data management strategy. Key drivers include:

- The need to cope with much larger data volumes and real-time data.
- Adding new alternative data sources to the analytics mix.
- Testing strategies, advanced statistical analysis,

and modeling.

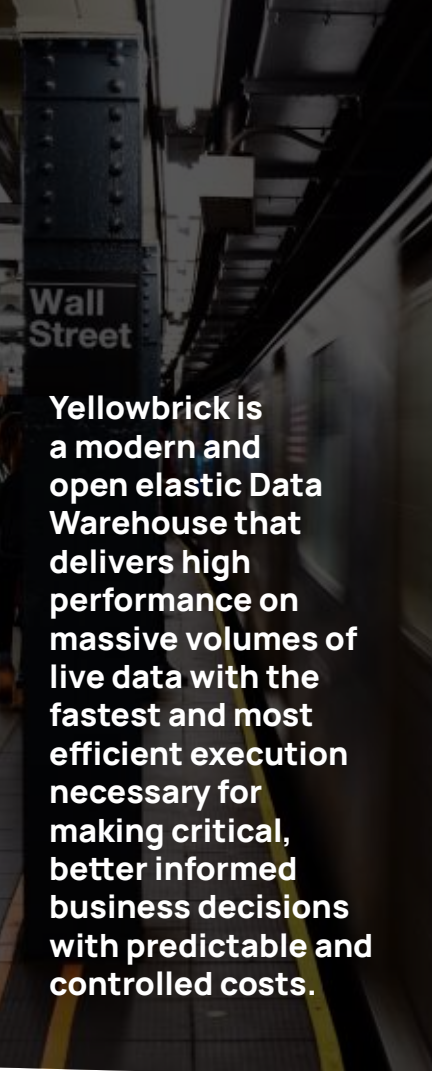
- Applying machine learning and AI techniques.
- Attracting and retaining talent.

Managing Data in a New Way

As these datasets become larger and more complex, it becomes less feasible for individual Portfolio Managers to manage their systems. Solutions built using traditional single server databases cannot cope with the volume or complexity of data or performance expectations. Much of this data arrives in real-time requiring constant attention.

As a result, central IT functions need to provision a central data services hub. They face the daunting challenge of ingesting huge volumes of data but also offering up a service that enables individual Portfolio Managers and their quants to experiment – find, back-test, and refine trading strategies independently.

All of this while executing the mission-critical day-to-day reporting, e.g., operational risk analytics, value-at-risk models, and financial reporting. Meeting the diverse needs of all participants across the hub while ensuring SLAs are met for time-sensitive processes is a non-trivial task. Historically, this results in blackout periods where ad-hoc analytics are restricted.



Yellowbrick is a modern and open elastic Data Warehouse that delivers high performance on massive volumes of live data with the fastest and most efficient execution necessary for making critical, better informed business decisions with predictable and controlled costs.

Data systems designed over a decade ago are unable to meet analytics needs in terms of scale, performance, or concurrency. They were not designed to meet the needs of a central hub. Real-time platforms like kdb+ cannot handle analytics on historical data and requires analysts skilled in kdb's proprietary query language. Attempts to adopt Big Data technologies like Hadoop have failed due to challenges with performance, the complexity of management, and a lack of stability.

The cloud, while it offers the ability to innovate and test new ideas without upfront capital investment in hardware, is not a panacea. Firms are often contractually unable to put client data or data not in the public domain in the public cloud. The skills needed to manage cloud services are also new and hard to come by. Further, cloud OPEX models and unfriendly vendor pricing can result in unpredictable runaway costs, which further constrain agility. In tough times when returns are hard to come by, variable costs are the first to be hit – ironically, this is when analytics are most needed.

Talent retention is also a consideration; recruits want to make money and want to work for firms that deliver the freedom and tools to make use of data. Platforms that restrict when or how much teams can query data to "protect" the system or control costs are no longer acceptable.


How can firms operate their data estate to get the best out of this hybrid-cloud and on-premises world?

Yellowbrick powers the central data hub for Hedge Funds, while addressing the needs of individual Portfolio Managers

That's where Yellowbrick Data Warehouse excels for our Hedge Fund customers, which include top 10 multi-strategy and quantitative strategy funds.

- **Performance, predictability, and familiarity**
Nobody runs queries at scale faster than Yellowbrick – joining multiple datasets, querying billions of rows, combining multiple rolling time windows, and complex calculations are not a challenge for Yellowbrick. The familiar SQL environment ensures there is no productivity gap and no new languages to learn.

Advanced workload management assures predictable response for mission-critical finance, compliance, risk, and regulatory reporting workloads while allowing complex ad hoc queries to run at the same time – addressing conflicts and eliminating restrictive time windows for querying data.



[LEARN MORE >](#)

[CONTACT US >](#)

[TAKE A TEST DRIVE >](#)

Yellowbrick logo are trademarks of Yellowbrick Data. All other trademarks are the property of their respective owners.

- **Addressing Portfolio Managers independence**

To cater to the needs of individual Portfolio Managers, firms provision resource-independent clusters within the same platform with shared data. Alternatively, each team has its own private independent instances, configuring automatic data replication from common pooled data. Non-public data remains private, visible only to individual teams, and remains on-premises if necessary.

- **The best of cloud with predictable pricing**

Yellowbrick instances run anywhere, but always "as-a-service" in both public cloud and on-premises on optimized hardware, solving for sensitive data. With Yellowbrick, an organization's data always runs entirely in its own cloud account or data center environment, eliminating concerns over data leakage (e.g., to cloud providers or service operators). Managed replication brings data from pooled datasets in near real-time and ships reporting data back to the hub.

Yellowbrick's 1- or 3-year subscription pricing model is highly predictable, works across cloud and on-premises, and caters to short-term burst needs via on-demand pricing.

Built to require minimal management, Yellowbrick allows organizations to focus their effort on business challenges and not managing technology. The security model provides both for client data sensitivities and enables appropriate separation between Portfolio Managers. The predictable pricing model enables agility but removes incentives to reduce spending on analytics when times are tough and analytics are most needed.

SUMMARY

Yellowbrick is the modern data warehouse designed to solve the complex challenges of Portfolio Manager analytics. It's easy to migrate to and delivers performance at low cost, giving superpowers to Portfolio and Fund Managers looking to deliver returns for investors, while providing critical internal and regulatory reporting. Yellowbrick is unique in allowing customers to combine private datasets, including those on-premises and data in a central cloud hub seamlessly, in a Distributed Data Cloud. No one runs data warehousing workloads faster and more efficiently than Yellowbrick.