

# Reimagining Traditional Data Warehouses with Yellowbrick and Dell Technologies

As data volumes continue to explode, finding that “*needle in the haystack*” becomes increasingly more difficult. Business users continue to demand new ways to drive value from their data – which in turn drives increased demands on IT organizations to meet critical SLA’s.

With Dell EMC PowerScale and Dell EMC ECS, large petabyte and exabyte data consolidation often hold critical data for many lines of business: Market Data for Capital Markets Quant Research, Fraud data for payment fraud risk management. Users are becoming more sophisticated in the types of questions they want/need to ask from this data. Building indexes, aggregations, and views, only scales so far so the ability to do “What-If” analysis becomes next to impossible. So, what’s the answer?

Dell Technologies has partnered with Yellowbrick Data to enable our customers to drive value from structured & semi-structured data. This myriad of data can now be leveraged on premise or in the Cloud without a forklift to another environment.

Based on a deep and persistent innovation in analytic database architecture (including hardware & software), the Yellowbrick data analytics solution enables lightning-fast insights from petabytes of data, even with hundreds of users working in parallel using common BI and data science tools. With Yellowbrick, even the most complex real-time workloads (involving years of historical data run in seconds).



Fig.1 – Capital Markets Example

## Solution Highlights

- **Extreme Performance:** The Yellowbrick Data Warehouse is optimized for challenging analytic environments. It dramatically accelerates the performance of individual workloads, while powering through the mixed workloads that slow traditional data warehouses.
- **Capacity:** Customer can maintain in-memory performance levels for multi petabytes of data. Capacity utilization is efficient as Yellowbrick compresses all data and can also encrypt data at rest.
- **Run Mixed Workloads, all the time:** Run ad-hoc queries with large batch queries, business reports with ETL processes, bulk loads with ODBC inserts, or all of these at the same time
- **Change workload priority and mix on the fly:** Slow jobs no longer result in a fire-drill. Administrators can re-prioritize running workloads at any time and configure alerts to let you know when unexpected things are happening
- **Spend less time tuning:** Data indexing is automatic. The core engine is so fast it plows through data at unprecedented rates. Ad-hoc workloads on the Yellowbrick Database run faster than heavily tuned, indexed queries on other databases.
- **Seamless data migration:** The Yellowbrick Data Warehouse can source data with real-time ingest, including flat ASCII, compressed data, Hadoop, S3, and many more.

## Solutions Overview

Today's data environments are diverse in their processing locations (i.e. cloud, on-prem), as well as their data structures (i.e. relational, JSON, XML, video, image). We have taken the one of the world's fastest Data Warehouse platform and connected it to the one of the world's most flexible and feature-rich storage infrastructure. In the process, we have delivered a Data Consolidation that is next level in its processing capabilities. Using Yellowbrick's modular appliances and its external table functionality, combined with the powerful and scalable Dell EMC storage offerings, this combination is now ready to ingest and process data of any type.

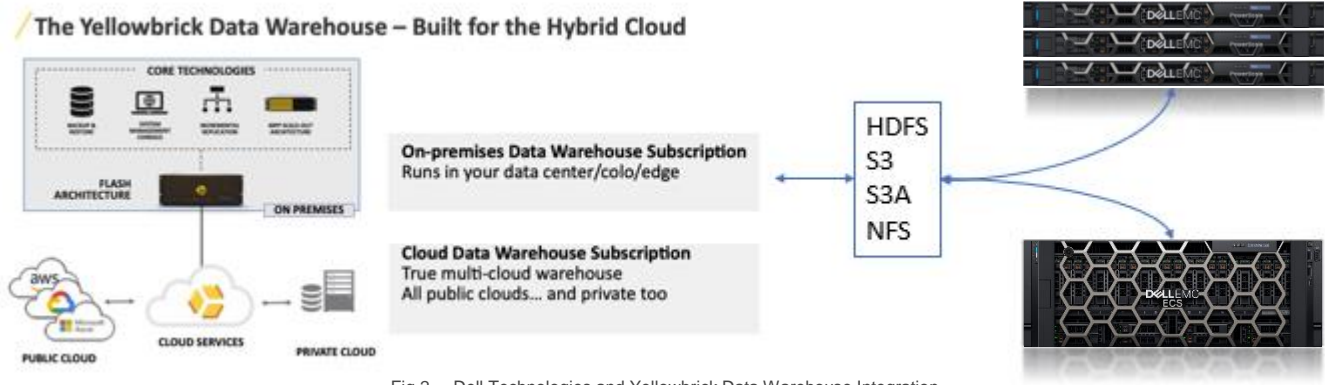


Fig 2. – Dell Technologies and Yellowbrick Data Warehouse Integration

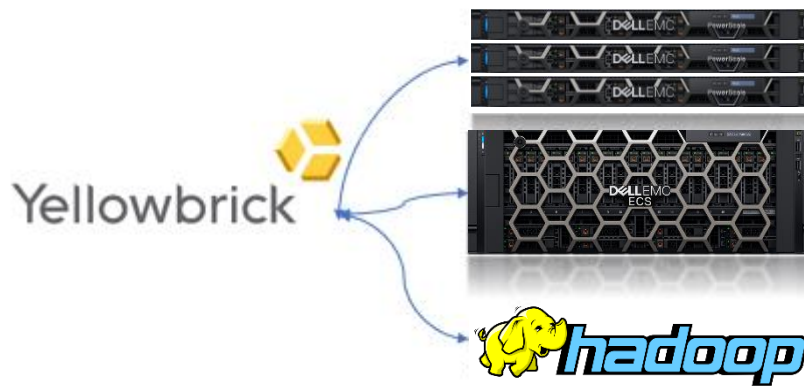


Fig 3. – Multiple sources to Yellowbrick Data Warehouse

## About Dell Technologies

Dell Technologies is a trusted leader in storage for analytics and AI. Dell EMC storage forms the foundation of a simple and future-proof infrastructure, giving you the agility to transform business operations, and the flexibility to adapt to new analytics workflows. Through innovative products and services, Dell EMC storage accelerates analytics and AI, helping businesses realize the true potential of their data.

## About Yellowbrick Data

Yellowbrick Data was founded in 2014 by experts in database and flash memory technologies to simplify data warehousing. We wanted to solve the challenges of high availability, running complex mixed workloads, support for ad-hoc SQL, computing correct answers on any schema, massive scalability and supporting large numbers of concurrent users.

Ultimately, the Yellowbrick Data Warehouse was built to be quick to deploy, easy to expand and simple to manage. And in doing so, it also fundamentally changed the economics of enterprise data warehousing to deliver low acquisition and operating costs, and high performance to customers.



To learn more about our offerings visit our website  
[DellTechnologies.com/StorageforDataAnalytics](https://DellTechnologies.com/StorageforDataAnalytics)