

EXECUTIVE BRIEF

Hybrid Cloud Architecture for Data Analytics

Hybrid cloud architecture is a highly strategic approach to cloud adoption that addresses customer concerns about reliable performance at scale, cloud lock-in, and the ability to support workloads that need to stay on premises due to security or other concerns. According to the Flexera 2020 State of the Cloud Report, more than 80 percent of organizations either have or want to have such a strategy; per Forrester Research, nearly 75 percent of enterprises already describe their strategy as hybrid and/or multi-cloud.

A hybrid cloud data warehouse is a logical and very effective addition to a hybrid cloud environment. As organizations optimize IT architectures comprising best-of-breed clouds and existing on-premises equipment and services, adding in a purpose-built hybrid data warehouse helps organizations meet and exceed their own unique data access and business analytics needs. It provides the best characteristics from cloud solutions (flexibility, data center savings, and scale) with the best characteristics of on-premises data warehouse solutions (security and performance at scale).

The authentic approach to hybrid cloud data warehousing

Not all approaches are created equal, however. For example, legacy on-premises vendors such as Teradata, IBM, and Oracle all have cloud versions of one kind or another in their portfolios, and even Amazon Web Services (AWS) has announced its intention to release on-premises products. However, in every case, customers are either locked into a single cloud platform, and/or the performance or control planes of the on-premises and cloud versions differ substantially. Furthermore, these solutions tend to be very expensive. At Yellowbrick Data, we believe that a true hybrid cloud

analytics architecture will not only unlock multi-cloud deployments but also will offer an identical level of performance across on-premises/private cloud and the public cloud. Furthermore, you should be able to consume and manage your workloads across all deployments in exactly the same way—in fact, the physical location of your data should be transparent to you. And finally, you should be free to immediately access all of that data, all of the time, to support long-running workloads without “surprise” consumption fees or waiting for data caches to rehydrate.

Yellowbrick hybrid cloud data warehouse

The Yellowbrick hybrid cloud data warehouse is the fastest available data warehouse platform for hybrid and multi-cloud environments. Thanks to highly specialized hardware/software that removes the traditional bottlenecks between storage and CPU, it offers breakthrough performance (up to 100x) beyond what either legacy platforms or the cloud providers’ virtualized commodity hardware can achieve, while providing the agility of the cloud—and at fraction of the cost of alternatives.

The Yellowbrick platform provides a reserved, single-tenant, always-on instance type that you consume like any other cloud service, whether from a data center/private cloud and/or from any major public cloud, including AWS, Microsoft Azure, and Google Cloud Platform. At the same time, Yellowbrick lets you avoid hidden fees for things such as data egress, bandwidth, and compute—fees the cloud-only providers love. Just connect to it like any other cloud service; should you ever need to switch clouds, simply reconnect to your instance and go (no egress charges apply). Users residing in different clouds can query a single Yellowbrick instance simultaneously.

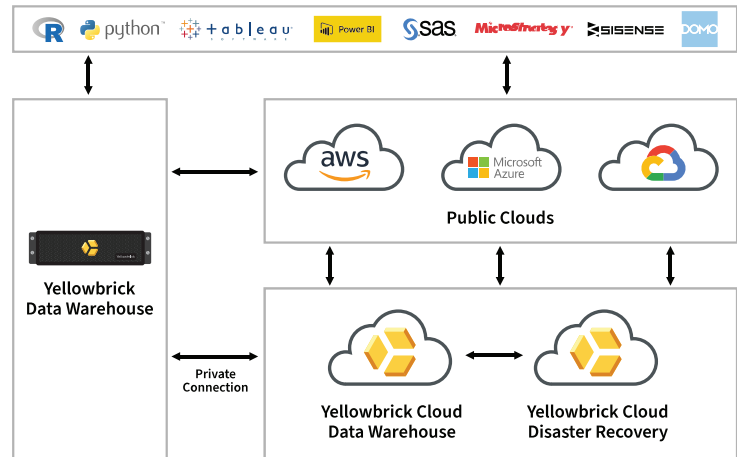
Yellowbrick also provides:

- > **Single-pane management across the entire system.** Yellowbrick provides a single-pane management interface that spans hybrid environments in a consistent, unified way.
- > **Simplified hybrid cloud and workload management.** Yellowbrick automatically takes care of many traditional management and database housekeeping tasks. Its intuitive console provides simplified but powerful workload management in an easy-to-use format.
- > **Cloud disaster recovery services.** Yellowbrick offers low-cost replication-based asynchronous disaster recovery as a service for multi-region fail-over and fail-back support.
- > **Easy integration and migration.** Yellowbrick looks and acts like an industry-standard Postgres RDBMS and integrates with common BI, data science, and ETL tools that speak ANSI SQL. That also makes it easy for users to write queries, as well as for existing queries to be migrated from similar platforms without modification.

Conclusion

Yellowbrick is the fastest and most flexible hybrid cloud data warehouse available, giving you performance and flexibility to meet fast-changing enterprise requirements at a fraction of the cost of alternatives. Benefits include:

- > **100x performance.** The Yellowbrick platform is architected to provide consistent, superior performance whether on premises or via public clouds, giving thousands of users real-time access to data that spans wide historical data, with no compromise in latency.
- > **Clear, predictable pricing.** Yellowbrick has a simplified pricing model with deterministic, predictable pricing that eliminates hidden cloud fees.
- > **Petabyte scale.** Yellowbrick can scale to meet the needs of organizations that have petabytes of data to



Yellowbrick Hybrid Cloud Architecture Overview

manage or that need to allow thousands of concurrent users access to data.

- > **Ease of hybrid cloud management.** Yellowbrick increases performance and flexibility without increasing management complexity—in fact, it simplifies management of hybrid cloud data warehousing, with a single management interface across the entire hybrid environment for managing roles, privileges, and more.

About Yellowbrick Data

Yellowbrick Data provides the world’s fastest data warehouse for hybrid and multi-cloud environments. Enterprises rely on the Yellowbrick hybrid cloud data warehouse to do the impossible in data analytics: get answers to the hardest business questions for improved profitability, better customer loyalty, and faster innovation in near real time, and at a fraction of the cost of alternatives. Yellowbrick offers superior price/performance for thousands of concurrent users on petabytes of data, along with the unique ability to run analytic workloads on premises, in a private cloud, and/or in any public cloud and manage them in a simple, consistent way—all with predictable pricing via annual subscription.

Learn more at yellowbrick.com.