

YOUR YELLOWBRICK DATA WAREHOUSE CAN BE PRODUCTION READY IN A DAY

Today's unprecedented growth in data and user performance needs pushes traditional data warehouses to their limits. Many IT departments face the daunting task of forklift upgrades and over-purchasing storage and compute capacity—all to buy just a few years relief until the next upgrade. Not so with the Yellowbrick Data Warehouse.

The Yellowbrick Data Warehouse gets customers up and running quickly, with as minimal, if any, disruption. While traditional data warehouse migrations require weeks or months of planning and labor, a Yellowbrick Data Warehouse can be installed, migrated, and running meaningful analytics in minutes. In addition, customers can add storage and compute resources incrementally, so they will not have to purchase more than they need or worry about future upgrades.

LIMITATIONS WITH PUBLIC CLOUD DATA WAREHOUSES

While it's acknowledged that the public cloud has many advantages there are concerns and drawbacks that must be taken into account with a uniquely public cloud solution for data warehouses.

- **Data Control and Security:** Enterprises are reluctant to hand over mission critical data to the public cloud due to concerns around data security and control over a shared infrastructure. Although cloud vendors offer assurances, the reality is that data breaches are happening more frequently leading to negative outcomes for the enterprise's end customers.
- **Performance:** Public clouds often don't live up to their performance claims. Not only does it take a long time to fully move all data into a cloud environment, once it is resident in the cloud, the performance to access data can be slow.
- **Support for Large Datasets:** Data sets continue to grow exponentially, and public clouds are unable to scale when datasets exceed a few 10s of terabytes and have limited elasticity with datasets in the petabytes.
- **Data Migration:** With a cloud only solution, enterprises are limited in their ability to move datasets from the cloud back to their on-premises equipment or to another cloud. Cloud vendors charge fees for transferring data out their cloud environment which dramatically increases costs – particularly as datasets continue to grow.

KEY HIGHLIGHTS

- Installation in minutes
- Fast, safe, secure data migration
- Seamless integration with existing tools and apps
- No hidden costs
- Grows to petabytes of capacity
- Eliminate painful upgrades with a highly scalable and incremental expand-as-you grow architecture



- **Hidden Costs:** From a pricing perspective, public cloud data warehouses can be expensive and costs unpredictable. Hidden costs such as long-term contracts, incremental staffing resources and supplementary licensing fees add more operating expenses to the bottom line.
- **Vendor lock-in:** Finally, and most importantly, with a public cloud data warehouse, a customer experiences vendor lock-in since data migration is limited and complex.

ADVANTAGES OF A HYBRID CLOUD DATA WAREHOUSE

Given all the limitations of a public cloud data warehouse, enterprises are realizing they need the best qualities of their private cloud and public cloud combined. The solution is a hybrid cloud data warehouse that offers the control and security of a private cloud combined with the disaster recovery and resource elasticity of a public cloud. According to the 2019 State of the Cloud Report, 84% of enterprise respondents have a multi-cloud strategy and 69% of respondents use at least one public and private cloud.

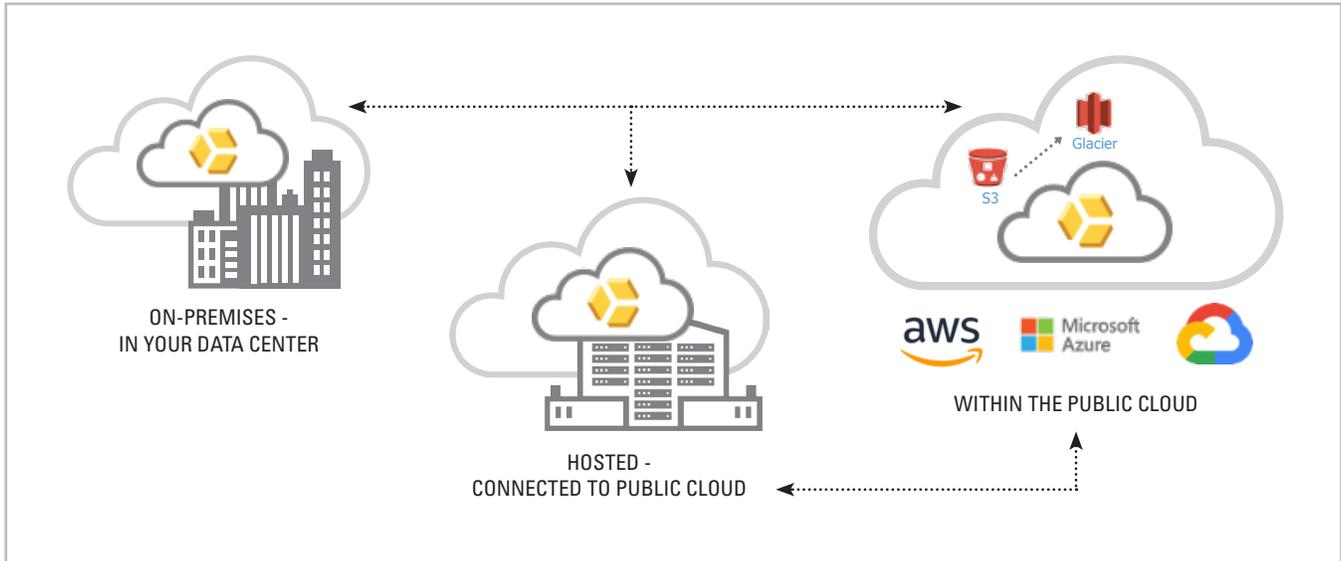
With a hybrid cloud deployment, enterprises can have their data in multiple public and private clouds simultaneously with analytics that operate seamlessly in both environments and that can scale to meet business growth. This flexibility ensures that an enterprise can be agile with their data and meet the needs of all parts of their organization with easy, fast access to data whenever and wherever it's needed.

Enterprises need to use their data to make real-time decisions quickly and efficiently. With a hybrid cloud data warehouse, organizations are no longer bottle-necked with slow performance and can drive to meaningful insights in a much shorter time period.

A hybrid data warehouse also provides data reliability, predictable pricing and compatibility with ETL/BI tools across both the public and private cloud.

YELLOWBRICK – THE NEW STANDARD IN HYBRID CLOUD DATA WAREHOUSE

Yellowbrick Data has architected a solution that not only effectively marries the data found in private and public clouds but also solves the challenges faced by public cloud only data warehouse solutions. Yellowbrick Data provides enterprises with a single tenant appliance whether in a private or public cloud. When customers are operating in multiple environments, Yellowbrick Data can place a single tenant appliance in each cloud (private or public) and establish a direct link between the appliances. Not only does this architecture allow rapid data replication in different cloud environments but also provides an effective disaster recovery and data backup solution.



DIFFERENTIATING FEATURES OF THE YELLOWBRICK DATA WAREHOUSE

Yellowbrick Data has a number of features that differentiates it from the other traditional on-premises and the public cloud data warehouse providers.

Platform Scalability

The Yellowbrick Data Warehouse stack has been engineered to perform at the speed of flash memory. The spinning disk has been replaced with an all memory architecture allowing data to move directly from flash memory to the CPU with virtually no latency. Flash architecture in conjunction with the scalable design means customers are able to expand memory capacity quickly. Also, the modular blade design allows customers to scale from terabytes to petabytes of data by adding compute nodes on the fly.

Performance

Enterprises don't have to sacrifice performance when operating in a hybrid environment, in fact just the opposite is true. The Yellowbrick Data Warehouse is architected to provide superior performance to those systems that only operate in the public or private cloud. The Yellowbrick Data Warehouse uses a flash-based architecture which is built and optimized for tasks such as ETL pushdown, in-place ELT, data wrangling and de-normalization. Overall, Yellowbrick Data sees a 10x price/performance increase and up to 100x better performance of ETL-heavy workloads. In addition, there is increased performance with this solution because the Yellowbrick Data Warehouse is always-on without the setup or take down time that is required with public cloud data warehouses.



Pricing

While the costs of operating a data warehouse in a public cloud appear straight forward, there are often many hidden costs including 1-3-year contract commitments, additional licensing costs for necessary tools and the need for incremental resources to manage this new environment. The Yellowbrick Data Warehouse offers a simplified pricing model with deterministic, predictable pricing. Because of Yellowbrick's architectural advances we also save costs for our private cloud customers by replacing racks of hardware with a 12-inch, fully integrated appliance allowing customers to save on hardware, power, cooling and data center costs.

Compatibility

The Yellowbrick Data Warehouse is ANSI SQL and Postgres compatible with a broad range of Hybrid ETL tools including Informatica PowerCenter, SyncSort, Hadoop, Spark, Kafka, and *DBC-based tools etc. The Yellowbrick Data warehouse also supports the use of standard BI, analytics and mining tools like Visual Studio, MicroStrategy, Tableau, and SAS.

Data Reliability

For most public and private cloud data warehouses, disaster recovery and backup are either an expensive option or don't exist at all. Disaster recovery for on-premises appliances requires buying a second appliance which sits idle most of the time, and disaster recovery for the cloud isn't typically included in enterprise SLAs. With its unique design, the Yellowbrick Data Warehouse provides for efficient disaster recovery and data backup at no incremental cost.

YELLOWBRICK DATA WAREHOUSE- EFFECTIVELY DELIVERING HYBRID DATA

Yellowbrick Data has built the best data warehouse for complex hybrid cloud needs. The Yellowbrick Data Warehouse is the world's fastest and offers the best price performance of any other solution – public cloud or on premise. The Yellowbrick solution is available via AWS, Google and Azure public clouds and we have an extensive partner ecosystem as well as direct sales force to ensure our customers have the right support for their data needs.

About Yellowbrick Data

Yellowbrick powers the data-driven enterprise with the world's only modern data warehouse designed for hybrid cloud and on-premises environments. Yellowbrick sets a new standard for the data warehouse. Built from the ground up to be self-managing and to provide unmatched performance with petabytes of data, Yellowbrick can support thousands of concurrent users performing complex ad hoc queries and running sophisticated reports. Through its unique hybrid cloud architecture, Yellowbrick natively manages large, difficult-to-migrate datasets where they reside—either on premises or in the cloud, all with a price/performance metric that surpasses any data warehouse, traditional or cloud-based.